



Flow control: new challenges for a new Renaissance

Thomas R. Bewley*

Department of MAE, UC San Diego, La Jolla, CA 92093-0411, USA

Abstract

As traditional scientific disciplines individually grow towards their maturity, many new opportunities for significant advances lie at their intersection. For example, remarkable developments in control theory in the last few decades have considerably expanded the selection of available tools which may be applied to regulate physical and electrical systems. These techniques hold great promise for several applications in fluid mechanics, including the delay of transition and the regulation of turbulence. Such applications of control theory require a very balanced perspective, in which one considers the relevant flow physics when designing the control algorithms and, conversely, takes into account the requirements and limitations of control algorithms when designing both reduced-order flow models and the fluid–mechanical systems to be controlled themselves. Such a balanced perspective is elusive, however, as both the research establishment in general and universities in particular are accustomed only to the dissemination and teaching of component technologies in isolated fields. To advance, we must not toss substantial new interdisciplinary questions over the fence for fear of them being “outside our area”; rather, we must break down these very fences that limit us, and attack these challenging new questions with a Renaissance approach. In this spirit, this paper surveys a few recent attempts at bridging the gaps between the several scientific disciplines comprising the field of flow control, in an attempt to clarify the author’s perspective on how recent advances in these constituent disciplines fit together in a manner that opens up significant new research opportunities. Published by Elsevier Science Ltd.

Contents

Preface	22
1. Linearization: life in a small neighborhood	22
2. Linear stabilization: leveraging modern linear control theory	25
2.1. The $\mathcal{H}_\infty$ approach to control design	25
2.2. Advantages of modern control design for nonnormal systems	28
2.3. Effectiveness of control feedback at particular wavenumber pairs	30
3. Decentralization: designing for massive arrays	30
3.1. Centralized approach	32
3.2. Decentralized approach	32
4. Localization: relaxing nonphysical assumptions	33
4.1. Open questions	34
5. Compensator reduction: eliminating unnecessary complexity	35
5.1. Fourier-space compensator reduction	35
5.2. Physical-space compensator reduction	35
5.3. Non-spatially-invariant systems	36
6. Extrapolation: linear control of nonlinear systems	36
7. Generalization: extending to spatially developing flows	39

* Tel.: + 1-858-534-4287.

E-mail address: bewley@ucsd.edu (T.R. Bewley).

URL: <http://turbulence.ucsd.edu>.

8. Nonlinear optimization: local solutions for full Navier–Stokes.	40
8.1. Adjoint-based optimization approach.	41
8.2. Continuous adjoint vs. discrete adjoint.	45
9. Robustification: appealing to Murphy’s law.	46
9.1. Well-posedness.	47
9.2. Convergence of numerical algorithm.	47
10. Unification: synthesizing a general framework.	47
11. Decomposition: simulation-based system modeling.	48
12. Global stabilization: conservatively enhancing stability.	48
13. Adaptation: accounting for a changing environment.	48
14. Performance limitation: identifying ideal control targets.	49
15. Implementation: evaluating engineering tradeoffs.	50
16. Discussion: a common language for dialog.	51
17. The future: a Renaissance.	51
Acknowledgements.	53
Appendix A. Alternative scaling of $\mathcal{H}_\infty$ problem.	53
Appendix B. A discrete adjoint for the 2D incompressible Navier–Stokes equation.	54
References.	57

Preface

The area of flow control plainly resides at the intersection of disciplines, incorporating essential and nontrivial elements from control theory, fluid mechanics, Navier–Stokes mathematics, and numerical methods. Recent developments in the integration of these disciplines, while grounding us with appropriate techniques to address some fundamental open questions, hint at the solution of several new questions which are yet to be asked. To follow up on these new directions, it is essential to have a clear vision of how recent advances in these fields fit together, and to know where the significant unresolved issues at their intersection lie.

The present paper will attempt to elucidate the utility of an interdisciplinary perspective to this type of problem by focusing on the control of a prototypical and fundamental fluid system: plane channel flow. The control of the flow in this simple geometry embodies a myriad of complex issues and inter-relationships whose understanding requires us to draw from a variety of traditional disciplines. Only when these issues and perspectives are combined is a complete understanding of the state of the art achieved, and a vision of where to proceed next identified.

Though plane channel flow will be the focus problem we will discuss here, the purpose of this work goes well beyond simply controlling this particular flow with a particular actuator/sensor configuration. At its core, the research effort we will describe is devoted to the development of an integrated, interdisciplinary understanding that will allow us to synthesize the necessary tools to attack a variety of flow control problems in the future. The focus problem of control of channel flow is chosen not simply because of its technological relevance or fundamental character, but because it embodies many of the

important unsolved issues to be encountered in the assortment of new flow control problems that will inevitably follow. The primary objective of the present work is to lay a solid, integrated footing upon which these future efforts may be based.

To this end, this paper will describe mostly the efforts with which the author has been directly involved, in an attempt to weave the story which threads these projects together as part of the fabric of a substantial new area of interdisciplinary research. Space does not permit the complete development of these projects in the present paper; rather, the paper will survey a selection of recent results which bring the relevant issues to light. The reader is referred to the appropriate full journal articles for all of the relevant details and careful placement of these projects in context with the works of others. Space limitations also do not allow this brief paper to adequately review the various directions all my friends and colleagues are taking in this field. Rather than attempt such a review and fail, the reader is referred to a host of other recent review papers which, taken together, themselves span only a fraction of the current work being done in this active area of research. From the experimental perspective, the reader is referred specifically to recent reviews of Ho and Tai [1,2], McMichael [3], Gad-El-Hak [4], and Löfdahl and Gad-El-Hak [5]. From the mathematical perspective, the reader is referred to the recent dedicated volumes compiled by Banks [6], Banks et al. [7], Gunzburger [8], Lagnese et al. [9], and Sritharan [10], for a sampling of recent results.

1. Linearization: life in a small neighborhood

As a starting point for the introduction of control theory into the fluid–mechanical setting, we first consider




Fig. 1. Geometry of plane channel flow. The flow is sustained by an externally applied pressure gradient in the x direction. This canonical problem provides an excellent testbed for the study of both transition and turbulence in wall-bounded flows. Note that many of the important flow phenomena in this geometry, in both the linear and nonlinear setting, are fundamentally 3D. A nonphysical assumption of periodicity of the flow perturbations in the x and z directions is often assumed for numerical convenience, with the box size chosen to be large enough that this nonphysical assumption has minimal effect on the observed flow statistics. It is important to evaluate critically the implications of such assumptions during the process of control design, as discussed in detail in Sections 3 and 4.

the linearized system arising from the equation governing small perturbations to a laminar flow. From a physical point of view, such perturbations are quite significant, as they represent the initial stages of the complex process of transition to turbulence, and thus their mitigation or enhancement has a substantial effect on the evolution of the flow.

To be concrete, an enlightening problem which captures the essential physics of many important features of both transition and turbulence in wall-bounded flows is that of plane channel flow, as illustrated in Fig. 1. Without loss of generality, we assume the walls are located at $y = \pm 1$. We begin our study by analyzing small perturbations $\{u, v, w, p\}$ to the (parabolic) laminar flow profile $U(y)$ in this geometry, which are governed by the linearized incompressible Navier–Stokes equation

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z} = 0, \quad (1a)$$

$$\dot{u} + U \frac{\partial}{\partial x} u + U'v = -\frac{\partial p}{\partial x} + \frac{1}{\text{Re}} \Delta u, \quad (1b)$$

$$\dot{v} + U \frac{\partial}{\partial x} v = -\frac{\partial p}{\partial y} + \frac{1}{\text{Re}} \Delta v, \quad (1c)$$

$$\dot{w} + U \frac{\partial}{\partial x} w = -\frac{\partial p}{\partial z} + \frac{1}{\text{Re}} \Delta w. \quad (1d)$$

Eq. (1a), the continuity equation, constrains the solution of (1b)–(1d), the momentum equations, to be divergence

free. This constraint is imposed through the ∇p terms in the momentum equations, which act as Lagrange multipliers to maintain the velocity field on a divergence-free submanifold of the space of square-integrable vector fields. In the discretized setting, such systems are called descriptor systems or differential–algebraic equations and, defining a state vector $\mathbf{x}$ and a control vector $\mathbf{u}$, may be written in the generalized state-space form

$$E\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{u}. \quad (2)$$

Note that if the Navier–Stokes equation (1) is put directly into this form, E is singular. This is an essential feature of the Navier–Stokes equation which necessitates careful treatment in both simulation and control design in order to avoid spurious numerical artifacts. A variety of techniques exist to express system (1) with a reduced set of variables or spatially distributed functions with only two degrees of freedom per spatial location, referred to as a divergence-free basis. In such a basis, the continuity equation is applied implicitly, and the pressure is eliminated from the set of governing equations. All three velocity components and the pressure (up to an arbitrary constant) may be determined from solutions represented in such a basis. When discretized and represented in form (2), the Navier–Stokes equation written in such a basis leads to an expression for E which is nonsingular.

For the geometry indicated in Fig. 1, a suitable choice for this reduced set of variables, which is convenient in terms of the implementation of boundary conditions, is the wall-normal velocity, v , and the wall normal vorticity, $\omega \triangleq \partial u / \partial z - \partial w / \partial x$. Taking the Fourier transform of (1) in the streamwise and spanwise directions and manipulating these equations and their derivatives leads to the classical Orr–Sommerfeld/Squire formulation of the Navier–Stokes equation at each wavenumber pair $\{k_x, k_z\}$:

$$\Delta \hat{v} = \{-ik_x U \Delta + ik_x U'' + \Delta(\Delta/\text{Re})\} \hat{v}, \quad (3a)$$

$$\dot{\hat{\omega}} = \{-ik_z U'\} \hat{v} + \{-ik_x U + \Delta/\text{Re}\} \hat{\omega}, \quad (3b)$$

where the hats ($\hat{\cdot}$) indicate Fourier coefficients and the Laplacian now takes the form $\Delta \triangleq \partial^2 / \partial y^2 - k_x^2 - k_z^2$. Note that particular care is needed when solving this system; in order to invert the Laplacian on the LHS of (3a), the boundary conditions on v must be accounted for properly. By manipulation of the governing equations and casting them in a derivative form, we effectively trade one numerical difficulty (singularity of E) for another (a tricky boundary condition inclusion to make the Laplacian on the LHS of (3a) invertible).

Note the spatially invariant structure of the present geometry: every point on each wall is, statistically speaking, identical to every other point on that wall. Canonical problems with this sort of spatially invariant structure in one or more directions form the backbone of much of the literature on flow transition and turbulence. It is this

structure which facilitates the use of Fourier transforms to *completely decouple* the system state $\{\hat{v}, \hat{\omega}\}$ at each wavenumber pair $\{k_x, k_z\}$ from the system state at every other wavenumber pair, as indicated in (3). Such decoupling of the Fourier modes of the unforced linear system in the directions of spatial invariance is a classical result upon which much of the available linear theory for the stability of Navier–Stokes systems is based. As noted by Bewley and Agarwal [11], taking the Fourier transform of both the control variables and the measurement variables maintains this system decoupling in the control formulation, greatly reducing the complexity of the control design problem to several smaller, completely decoupled control design problems at each wavenumber pair $\{k_x, k_z\}$, each of which requires spatial discretization in the y direction only.

Once a tractable form of the governing equation has been selected, in order to pose the flow control problem completely, several steps remain:

- the state equation must be spatially discretized,
- boundary conditions must be chosen and enforced,
- the variables representing the controls and the available measurements must be identified and extracted,
- the disturbances must be modeled, and
- the “control objective” must be precisely defined.

To identify a fundamental yet physically relevant flow control problem, the decisions made at each of these steps requires engineering judgment. Such judgment is based on physical insight concerning the flow system to be controlled and how the essential features of such a system may be accurately modeled. An example of how to accomplish these steps is described in some detail by Bewley and Liu [12]. In short, we may choose:

- a Chebyshev spatial discretization in y ,
- no-slip boundary conditions ($u = w = 0$ on the walls) with the distribution of v on the walls (the blowing/suction profile) prescribed as the control,
- skin friction measurements distributed on the walls,
- idealized disturbances exciting the system and
- an objective of minimizing flow perturbation energy.

As we learn more about the physics of the system to be controlled, there is significant room for improvement in this problem formulation, particularly in modeling the structure of relevant system disturbances and in the precise statement of the control objective.

Once the above-mentioned steps are complete, the present decoupled system at each wavenumber pair $\{k_x, k_z\}$ may finally be manipulated into the standard state-space form

$$\begin{aligned}\dot{\mathbf{x}} &= \mathbf{A}\mathbf{x} + \mathbf{B}_1\mathbf{w} + \mathbf{B}_2\mathbf{u}, \\ \mathbf{y} &= \mathbf{C}_2\mathbf{x} + \mathbf{D}_{21}\mathbf{w}\end{aligned}\quad (4)$$

with

$$\mathbf{B}_1 \triangleq (\mathbf{G}_1 \ 0), \quad \mathbf{C}_2 \triangleq \mathbf{G}_2^{-1}\mathbf{C}, \quad \mathbf{D}_{21} \triangleq (0 \ \alpha\mathbf{I}), \quad \mathbf{w} \triangleq \begin{pmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \end{pmatrix},$$

where $\mathbf{x}$ denotes the state, $\mathbf{u}$ denotes the control, $\mathbf{y}$ denotes the available measurements (scaled as discussed below), and $\mathbf{w}$ accounts for the external disturbances (including the state disturbances $\mathbf{w}_1$ and the measurement noise $\mathbf{w}_2$, scaled as discussed below). Note that $\mathbf{C}\mathbf{x}$ denotes the raw vector of measured variables, and $\mathbf{G}_1$ and $\alpha\mathbf{G}_2$ represent the square root of any known or expected covariance structure of the state disturbances and measurement noise respectively. The scalar α^2 is identified as an adjustable parameter which defines the ratio of the maximum singular value of the covariance of the measurement noise divided by the maximum singular value of the covariance of the state disturbances; w.l.o.g., we take $\bar{\sigma}(\mathbf{G}_1) = \bar{\sigma}(\mathbf{G}_2) = 1$. Effectively, the matrix $\mathbf{G}_1$ reflects which state disturbances are strongest, and the matrix $\mathbf{G}_2$ reflects which measurements are most corrupted by noise. Small α implies relatively high overall confidence in the measurements, whereas large α implies relatively low overall confidence in the measurements.

Not surprisingly, there is a wide body of theory surrounding how to control a linear system of the standard form (4). The application of one popular technique (to a related 2D problem), called proportional-integral (PI) control and generally referred to as “classical” control design, is presented in Joshi et al. [13]. The application of another technique, called $\mathcal{H}_\infty$ control and generally referred to as “modern” control design, is laid out in Bewley and Liu [12], hereafter referred to as BL98. The application of a related modern control strategy (to the 2D problem), called loop transfer recovery (LTR), is presented in Cortezzi and Speyer [14]. More recent publications by these groups further extend these seminal efforts.

It is useful, to some extent, to understand the various theoretical implications of the control design technique chosen. Ultimately, however, flow control boils down to the *design* of a control that achieves the desired engineering objective (transition delay, drag reduction, mixing enhancement, etc.) to the maximum extent possible. The theoretical implications of the particular control technique chosen are useful only to the degree to which they help attain this objective. Engineering judgment, based both on an understanding of the merits of the various control theories and on the suitability of such theories to the structure of the fluid–mechanical problem of interest, guides the selection of an appropriate control design strategy. In the following section, we summarize the $\mathcal{H}_\infty$ control design approach, illustrate why this approach is appropriate for the structure of the problem at hand, and highlight an important distinguishing

characteristic of the present system when controls computed via this approach are applied.

2. Linear stabilization: leveraging modern linear control theory

As only a limited number of noisy measurements y of the state x are available in any practical control implementation, it is beneficial to develop a filter which extracts as much useful information as possible from the available flow measurements before using this filtered information to compute a suitable control. In modern control theory, a model of the system itself is used as this filter, and the filtered information extracted from the measurements is simply an estimate of the state of the physical system. This intuitive framework is illustrated schematically in Fig. 2. By modeling (or neglecting) the influence of the unknown disturbances in (4), the system model takes the form

$$\dot{\hat{x}} = A\hat{x} + B_1\hat{w} + B_2u - v, \quad (5a)$$

$$\hat{y} = C_2\hat{x} + D_{21}\hat{w}, \quad (5b)$$

where $\hat{x}$ is the state estimate, $\hat{w}$ is a disturbance estimate, and v is a feedback term based on the difference between the measurement of the state y and the corresponding quantity in the model $\hat{y}$ such that

$$v = L(y - \hat{y}). \quad (5c)$$




Fig. 2. Flow of information in a modern control realization. The plant, forced by external disturbances, has an internal state x which cannot be observed. Instead, a noisy measurement y is made, with which a state estimate $\hat{x}$ is determined. This state estimate is then used to determine the control u to be applied to the plant to regulate x to zero. Essentially, the full equation for the plant (or a reduced model thereof) is used in the estimator as a filter to extract useful information about the state from the available measurements.

The control u , in turn, is based on the state estimate $\hat{x}$ such that

$$u = K\hat{x}. \quad (6)$$

Eq. (4) is referred to as the “plant”, (5) is referred to as the “estimator”, and (6) is referred to as the “controller”. The estimator (5) and the controller (6), taken together, will be referred to as the “compensator”. The problem at hand is to compute linear time-invariant (LTI) matrices K and L and some estimate of the disturbance, $\hat{w}$, such that

- (i) the estimator feedback v forces $\hat{x}$ towards x , and
- (ii) the controller feedback u forces x towards zero,

even as unknown disturbances w both disrupt the system evolution and corrupt the available measurements of the system state.

2.1. The $\mathcal{H}_\infty$ approach to control design

Several textbooks describe in detail how the $\mathcal{H}_\infty$ technique determines K , L , and $\hat{w}$ for systems of form (4)–(6) in the presence of structured and unstructured disturbances w . The reader is referred to the seminal paper by Doyle et al. [15], the more accessible textbook by Green and Limebeer [16], and the more advanced texts by Zhou et al. [17] and Zhou and Doyle [18] for derivation and further discussion of these control theories, and to [12] for an extended discussion in the context of the present problem. To summarize this approach briefly, a cost function $\mathcal{J}$ describing the control problem at hand is defined that weighs together the state x , the control u , and the disturbance w such that

$$\begin{aligned} \mathcal{J} &\triangleq E[x^*Qx + \ell^2u^*u - \gamma^2w^*w] \\ &\triangleq E[z^*z - \gamma^2w^*w], \end{aligned} \quad (7a)$$

where

$$z \triangleq C_1x + D_{12}u \quad (7b)$$

and

$$C_1 \triangleq \begin{pmatrix} Q^{1/2} \\ 0 \end{pmatrix}, \quad D_{12} \triangleq \begin{pmatrix} 0 \\ I \end{pmatrix}.$$

The matrix Q shaping the dependence on the state in the cost function, x^*Qx , may be selected to numerically approximate any of a variety of physical properties of the flow, such as the flow perturbation energy, its entropy, the mean square of the drag measurements, etc. The matrix Q may also be biased to place extra penalty on flow perturbations in a specific region in space of particular physical significance. The choice of Q has a profound effect on the final closed-loop behavior, and must be selected with care. Based on our numerical tests to date,

cost functions related to the energy of the flow perturbations have been the most successful for the purpose of transition delay. To simplify the algebra that follows, we have set the matrices R and S shaping the $\mathbf{u}^* \mathbf{R} \mathbf{u}$ and $\mathbf{w}^* \mathbf{S} \mathbf{w}$ terms in the cost function equal to I . It is straightforward to generalize this result to other positive-definite choices for R and S ; this alternative scaling of the $\mathcal{H}_\infty$ control problem is presented in Appendix A. As discussed in Lange and Bewley [19], such a generalization is particularly useful when designing controls for a discretization of a PDE in a consistent manner such that the feedback kernels will converge to continuous functions as the computational grid is refined.

Given the structure of the system defined in (4)–(6) and the control objective defined in (7), the $\mathcal{H}_\infty$ compensator is determined by simultaneously minimizing the cost function $\mathcal{J}$ with respect to the control $\mathbf{u}$ and maximizing $\mathcal{J}$ with respect to the disturbance $\mathbf{w}$. In such a way, a control $\mathbf{u}$ is found which maximally attains the control objective even in the presence of a disturbance $\mathbf{w}$ which maximally disrupts this objective. For *sufficiently large* γ and a system which is both *stabilizable* and *detectable* via the controls and measurements chosen,¹ this results in finite values for $\mathbf{u}$, $\mathbf{v}$ and $\mathbf{w}$, the magnitudes of which may be adjusted by variation of the three scalar parameters ℓ , α and γ respectively. Reducing ℓ , modeling the “price of the control” in the engineering design, generally results in increased levels of control feedback $\mathbf{u}$. Reducing α , modeling the “relative level of corruption” of the measurements by noise, generally results in increased levels of estimator feedback $\mathbf{v}$. Reducing γ , modeling the “price” of the disturbance to Nature (in the spirit of a noncooperative game), generally results in increased levels of disturbances $\mathbf{w}$ of maximally disruptive structure to be accounted for during the design of the compensator.

The $\mathcal{H}_\infty$ control solution [15] may be described as follows: a compensator which minimizes $\mathcal{J}$ in the presence of that disturbance which simultaneously maximizes $\mathcal{J}$ is given by

$$K = -\frac{1}{\ell^2} B_2^* X, \quad L = -\frac{1}{\alpha^2} Z Y C_2^*,$$

$$\hat{\mathbf{w}} = \frac{1}{\gamma^2} B_1^* X \hat{\mathbf{x}}, \quad (8)$$

where

$$X = \text{Ric} \begin{pmatrix} A & \frac{1}{\gamma^2} B_1 B_1^* - \frac{1}{\ell^2} B_2 B_2^* \\ -C_1^* C_1 & -A^* \end{pmatrix},$$

¹ Further description of these important technical requirements for solvability of the control problem is deferred to the above mentioned texts. These requirements are easily met in many practical settings.

$$Y = \text{Ric} \begin{pmatrix} A^* & \frac{1}{\gamma^2} C_1^* C_1 - \frac{1}{\alpha^2} C_2^* C_2 \\ -B_1 B_1^* & -A \end{pmatrix},$$

$$Z = \left(I - \frac{Y X}{\gamma^2} \right)^{-1},$$

where $\text{Ric}(\cdot)$ denotes the positive-definite solution of the associated Riccati equation [20].² The simple structure of the above solution, and its profound implications in terms of the performance and robustness of the resulting closed-loop system, is one of the most elegant results of linear control theory. We comment below on a few of the more salient features of this result.

Algebraic manipulation of (4)–(8) leads to the closed-loop form

$$\dot{\tilde{\mathbf{x}}} = \tilde{A} \tilde{\mathbf{x}} + \tilde{B} \mathbf{w},$$

$$\mathbf{z} = \tilde{C} \tilde{\mathbf{x}}, \quad (9)$$

where

$$\tilde{A} = \begin{pmatrix} A + B_2 K & -B_2 K \\ -\gamma^{-2} B_1 B_1^* & A + L C_2 + \gamma^{-2} B_1 B_1^* \end{pmatrix},$$

$$\tilde{\mathbf{x}} = \begin{pmatrix} \mathbf{x} \\ \mathbf{x} - \hat{\mathbf{x}} \end{pmatrix}, \quad \tilde{B} = \begin{pmatrix} B_1 \\ B_1 + L D_{21} \end{pmatrix},$$

$$\tilde{C} = (C_1 + D_{12} K - D_{12} K).$$

Taking the Laplace transform of (9), it is easy to define the transfer function $T_{zw}(s)$ from $\mathbf{w}(s)$ to $\mathbf{z}(s)$ (the Laplace transforms of $\mathbf{w}$ and $\mathbf{z}$) such that

$$\mathbf{z}(s) = \tilde{C}(sI - \tilde{A})^{-1} \tilde{B} \mathbf{w}(s) \triangleq T_{zw}(s) \mathbf{w}(s).$$

Norms of the system transfer function $T_{zw}(s)$ quantify how the system output of interest, $\mathbf{z}$, responds to disturbances $\mathbf{w}$ exciting the closed-loop system.

The expected value of the rms of the output $\mathbf{z}$ over the rms of the input $\mathbf{w}$ for disturbances $\mathbf{w}$ of maximally disruptive structure is denoted by the ∞ -norm of the system transfer function,

$$\|T_{zw}\|_\infty \triangleq \sup_{\omega} \bar{\sigma}[T_{zw}(j\omega)].$$

$\mathcal{H}_\infty$ control is often referred to as “robust” control, as $\|T_{zw}\|_\infty$, reflecting the worst-case amplification of disturbances by the system from the input $\mathbf{w}$ to the output $\mathbf{z}$, is in fact bounded from above by the value of γ used in the problem formulation. Subject to this ∞ -norm bound,

² Note that, for the control problem to be soluble, γ must be sufficiently large so that: (a) X and Y may be found that are positive definite, and (b) $\rho(XY) < \gamma^2$, where $\rho(\cdot)$ denotes the spectral radius. An approximate lower bound on γ which meets these conditions, denoted γ_0 , may be determined by trial and error.

$\mathcal{H}_\infty$ control minimizes the expected value of the rms of the output $\mathbf{z}$ over the rms of the input $\mathbf{w}$ for white Gaussian disturbances $\mathbf{w}$ with identity covariance, denoted by the 2-norm of the system transfer function

$$\|T_{zw}\|_2 \triangleq \left(\frac{1}{2\pi} \int_{-\infty}^{\infty} \text{trace}[T_{zw}(j\omega)^* T_{zw}(j\omega)] d\omega \right)^{1/2}.$$

Note that $\|T_{zw}\|_2$ is often cited as a measure of performance of the closed-loop system, whereas $\|T_{zw}\|_\infty$ is often cited as a measure of its robustness. Further motivation for consideration of control theories related to these particular norms is elucidated by Skogestad and Postlethwaite [21]. Efficient numerical algorithms to solve the Riccati equations for X and Y in the compensator design and to compute the transfer function norms $\|T_{zw}\|_2$ and $\|T_{zw}\|_\infty$ quantifying the closed-loop system behavior are well developed, and are discussed further in the standard texts.

Note that, for high-dimensional discretizations of infinite dimensional systems, it is not feasible to perform a parametric variation on the individual elements of the matrices defining the control problem. The control design approach taken here represents a balance of *engineering judgment* in the construction of the matrices defining in the structure of the control problem, $\{B_1, B_2, C_1, C_2\}$, and *parametric variation* of the three scalar parameters involved, $\{\ell, \alpha, \gamma\}$, in order to achieve the desired tradeoffs between performance, robustness, and the control effort required. This approach retains a sufficient but not excessive degree of flexibility in the control design process. In general, intermediate values of the three parameters $\{\ell, \alpha, \gamma\}$ are found to lead to the most suitable control designs.

$\mathcal{H}_2$ control (also known as linear quadratic Gaussian control, or LQG) is an important limiting case of $\mathcal{H}_\infty$ control. It is obtained in the present formulation by relaxing the bound γ on the infinity norm of the closed-loop system, taking the limit as $\gamma \rightarrow \infty$ in the controller formulation. Such a control formulation focuses solely on performance, i.e., minimizing $\|T_{zw}\|_2$. As LQG does not provide any guarantees about system behavior for disturbances of particularly disruptive structure ($\|T_{zw}\|_\infty$), it is often referred to as “optimal” control. Though one might confirm a posteriori that a particular LQG design has favorable robustness properties, such properties are not guaranteed by the LQG control design process. When designing a large number of compensators for an entire array of wavenumber pairs $\{k_x, k_z\}$ via an automated algorithm, as is necessary in the present problem, it is useful to have a control design tool which inherently builds in system robustness, such as $\mathcal{H}_\infty$. For isolated low-dimensional systems, as often encountered in many industrial processes, a posteriori robustness checks on hand-tuned LQG designs are often sufficient.

It is also interesting to note that certain favorable robustness properties may be assured by the LQG approach by strategies involving either:

- (a) setting $B_1 = (B_2 \ 0)$ and taking $\alpha \rightarrow 0$, or
- (b) setting $C_1 = \begin{pmatrix} C_2 \\ 0 \end{pmatrix}$ and taking $\ell \rightarrow 0$.

These two approaches are referred to as loop transfer recovery (LQG/LTR), and are further explained in Stein and Athans [22]. Such a strategy is explored by Cortezzi and Speyer [14] in the 2D setting of the present problem. In the present system, both B_2 and C_2 are very low rank, as there is only a single control variable and a single measurement variable at each wall in the Fourier-space representation of the physical system at each wavenumber pair $\{k_x, k_z\}$. However, the state itself is a high-dimensional approximation of an infinite-dimensional system. It is beneficial in such a problem to allow the modeled state disturbances $\mathbf{w}_1$ to input the system, via the matrix B_1 , at more than just the actuator inputs, and to allow the response of the system $\mathbf{x}$ to be weighted in the cost function, via the matrix C_1 , at more than just the sensor outputs. The LQG/LTR approach of assuring closed-loop system robustness, however, requires us to sacrifice one of these features in the control formulation, in addition to taking $\alpha \rightarrow 0$ or $\ell \rightarrow 0$, in order to apply one of the two strategies listed above. It is noted here that the $\mathcal{H}_\infty$ approach, when soluble, allows for the design of compensators with inherent robustness guarantees without such sacrifices of flexibility in the definition of the control problem of interest, thereby giving significantly more latitude in the design of a “robust” compensator.

The names $\mathcal{H}_2$ and $\mathcal{H}_\infty$ are derived from the system norms $\|T_{zw}\|_2$ and $\|T_{zw}\|_\infty$ which these control theories address, with the symbol $\mathcal{H}$ denoting the particular “Hardy space” in which these transfer functions norms are well defined. It deserves mention that the difference between $\|T_{zw}\|_2$ and $\|T_{zw}\|_\infty$ might be expected to be increasingly significant as the dimension of the system is increased. Neglecting, for the moment, the dependence on ω in the definition of the system norms, the matrix Frobenius norm, $(\text{trace}[T^*T])^{1/2}$, and the matrix 2-norm, $\bar{\sigma}[T]$, are “equivalent” up to a constant. Indeed, for scalar systems, these two matrix norms are identical, and for low-dimensional systems, their ratio is bounded by a constant related to the dimension of the system. For high-dimensional discretizations of infinite-dimensional systems, however, this norm equivalence is relaxed, and the differences between these two matrix norms may be substantial. The temporal dependence of the two system norms $\|T_{zw}\|_2$ and $\|T_{zw}\|_\infty$ distinguishes them even for low-dimensional systems; the point here is only that, for high-dimensional systems, the important differences between these two system norms is even more pronounced, and control techniques, such as $\mathcal{H}_\infty$, which account for both such norms might prove to be beneficial.

Techniques (like $\mathcal{H}_\infty$) which bound $\|T_{zw}\|_\infty$ are especially appropriate for the present problem, as transition is often associated with the triggering of a “worst-case” phenomenon, which is well characterized by this measure.

2.2. Advantages of modern control design for nonnormal systems

Matrices A arising from the discretization of systems in fluid mechanics are often highly “nonnormal”, which means that the eigenvectors of A are highly nonorthogonal. This is especially true for transition in a plane channel. Important characteristics of this system, such as $O(1000)$ transient energy growth and large amplification of external disturbance energy in stable flows at subcritical Reynolds numbers, cannot be explained by examination of its eigenvalues alone. Discretizations of (3), when put into the state-space form (4), lead to system matrices of the form

$$A = \begin{pmatrix} L & 0 \\ C & S \end{pmatrix}. \quad (10)$$

For certain wavenumber pairs (specifically, those with $k_x \approx 0$ and $k_z = O(1)$), the eigenvalues of A are real and stable, the matrices L and S are quite similar in structure, and $\bar{\sigma}(C)$ is disproportionately large.

In order to illustrate the behavior of a system matrix with such structure, consider a reduced system matrix of the above form but where L , C , and S are scalars. Specifically, compare the two stable closed-loop system matrices

$$A_1 = \begin{pmatrix} -0.01 & 0 \\ 0 & -0.011 \end{pmatrix}, \quad A_2 = \begin{pmatrix} -0.01 & 0 \\ 1 & -0.011 \end{pmatrix}.$$

Both matrices have the same eigenvalues. However, the eigenvectors of A_1 are orthogonal, whereas the eigenvectors of A_2 are

$$\xi_1 = \begin{pmatrix} 0.001 \\ 1.000 \end{pmatrix} \quad \text{and} \quad \xi_2 = \begin{pmatrix} 0 \\ 1.000 \end{pmatrix}.$$

Even though its eigenvalues differ by 10%, the eigenvectors of A_2 are less than 0.06° from being exactly parallel. It is in this sense that we define this system as being “nonnormal” or “nearly defective”.³ This severe non-


orthogonality of the system eigenvectors is a direct result of the disproportionately large coupling term C . Compensators which reduce C will make the eigenvectors of A_2 closer to orthogonal without necessarily changing the system eigenvalues.

The consequences of nonorthogonality of the system eigenvectors are significant. Though the “energy” (the L^2 -norm) of the state of the system $\dot{\mathbf{x}} = A_1 \mathbf{x}$ uniformly decreases in time from all initial conditions, the “energy” of the state of the system $\dot{\mathbf{x}} = A_2 \mathbf{x}$ from the initial condition $\mathbf{x}(0) = \xi_1 - \xi_2$ grows by a factor of over a thousand before eventually decaying due to the stability of the system. This is referred to as the transient energy growth of the stable nonnormal system, and is a result of the reduced destructive interference exhibited by the two modes of the solution as they decay at different rates. In fluid mechanics, transient energy growth is thought to be an important linear mechanism leading to transition in subcritical flows, which are linearly stable but nonlinearly unstable [23].


The excitation of such systems by external disturbances is well described in terms of the system norms $\|T_{zw}\|_2$ and $\|T_{zw}\|_\infty$, which (as described previously) quantify the rms amplification of Gaussian and worst-case disturbances by the system. For example, consider a closed-loop system of form (9) with $\tilde{B} = \tilde{C} = I$. Taking the system matrix $\tilde{A} = A_1$, the norms of the system transfer function are $\|T_{zw}\|_2 = 9.8$ and $\|T_{zw}\|_\infty = 100$. On the other hand, taking the system matrix $\tilde{A} = A_2$, the 2-norm of the system transfer function is 48 times larger and the ∞ -norm is 91 times larger, though the two systems have identical closed-loop eigenvalues. Large system transfer function norms and large values of maximum transient energy growth are often highly correlated, as they both come about due to nonnormality in a stable system.

Graphical interpretations of $\|T_{zw}\|_2$ and $\|T_{zw}\|_\infty$ for the present channel flow system are given in Figs. 3 and 4 by examining contour plots of the appropriate matrix norms of $T_{zw}(s)$ in the complex plane s . Recall that $T_{zw}(s) \triangleq \tilde{C}(sI - \tilde{A})^{-1}\tilde{B}$, so these contours approach infinity in the neighborhood of each eigenvalue of $\tilde{A}$. Contour plots of this type have recently become known as the pseudospectra of an input/output system, and have become a popular generalization of plots of the eigenvalues of $\tilde{A}$ in recent efforts to study nonnormality in uncontrolled fluid systems [24]. For the open-loop systems depicted in these figures, we define $\tilde{A} = A$, $\tilde{B} = B_1$, and $\tilde{C} = C_1$. The severe nonnormality of the present fluid system for Fourier modes with $k_x \approx 0$ is reflected by the elliptical isolines surrounding each pair of eigenvalues with nearly parallel eigenvectors in these pseudospectra, a feature which is much more pronounced in the system depicted in Fig. 3 than in that depicted in Fig. 4. The severe nonnormality of the system depicted in Fig. 3 is also reflected by its much larger value of $\|T_{zw}\|_\infty$. As

³Though this definition is dependent on the coordinate system and norm chosen to define the orthogonality of the eigenvectors, the physical systems we will consider lend themselves naturally to preferred norm definitions motivated by the energetics of the system. In particular, the nonlinear terms of the Navier–Stokes equation (which are neglected in the present linear analysis) are orthogonal only under certain norms related to the kinetic energy of the system at hand, suggesting a natural, physically motivated choice of norm for the systems we will consider.



(a) Isocontours of $\bar{\sigma}[T_{\mathbf{zw}}(s)]$ in the complex plane s . The peak value of this matrix norm on the $j\omega$ axis is defined as the system norm $\|T_{\mathbf{zw}}\|_{\infty}$, and corresponds to the solid isoline with the smallest value.



(b) Isocontours of $[\text{trace}(T_{\mathbf{zw}}^*(s)T_{\mathbf{zw}}(s))]^{1/2}$ in the complex plane s . The system norm $\|T_{\mathbf{zw}}\|_2$ is related to the integral of the square of this matrix norm over the $j\omega$ axis.

Fig. 3. Graphical interpretations (a.k.a. “pseudospectra”) of the transfer function norms $\|T_{\mathbf{zw}}\|_{\infty}$ (top) and $\|T_{\mathbf{zw}}\|_2$ (bottom) for the present system in open loop, obtained at $k_x = 0$, $k_z = 2$ and $\text{Re} = 5000$. The eigenvalues of the system matrix A are marked with an $\times$. All isoline values are separated by a factor of 2, and the isolines with the largest value are those nearest to the eigenvalues. For this system, $\|T_{\mathbf{zw}}\|_{\infty} = 2.6 \times 10^5$.

$\{\tilde{A}, \tilde{B}, \tilde{C}\}$ may be defined for either the open-loop or the closed-loop case, this technique for analysis of nonnormality extends directly to the characterization of controlled fluid systems.

The $\mathcal{H}_{\infty}$ control technique is in fact based on minimizing the 2-norm of the system transfer function while simultaneously bounding the ∞ -norm of the system transfer function. In the present transition problem, our

control objective is to inhibit the (linear) formation of energetic flow perturbations that can lead to nonlinear instability and transition to turbulence. It is natural that control techniques such as $\mathcal{H}_{\infty}$, which are designed upon the very transfer function norms which quantify the excitation of such flow perturbations by external disturbances, will have a distinct advantage for achieving this objective over control techniques which account for the




Fig. 4. Pseudospectra interpretations of $\|T_{zw}\|_\infty$ (left) and $\|T_{zw}\|_2$ (right) for the open loop system at $k_x = -1, k_z = 0$ and $\text{Re} = 5000$. For plotting details, see Fig. 3. For this system, $\|T_{zw}\|_\infty = 1.9 \times 10^4$.

eigenvalues only, such as those based on the analysis of root-locus plots.

2.3. Effectiveness of control feedback at particular wavenumber pairs

The application of the modern control design approach described in Section 2.1 to the Orr-Sommerfeld/Squire problem laid out in Section 1 was explored extensively in [12] for two particular wavenumber pairs and Reynolds numbers. The control effectiveness was quantified using several different techniques, including eigenmode analysis, transient energy growth, and transfer function norms. The control was remarkably effective and the trends with $\{\ell, \alpha, \gamma\}$ were all as expected; the reader is referred to the journal article for complete tabulation of the results. One of the most notable features of this paper is that the application of the control resulted in the closed-loop eigenvectors becoming significantly closer to orthogonal, as illustrated in Fig. 5; note especially the high degree of correlation between the second and third eigenvectors of Fig. 5a, and how this correlation is disrupted in Fig. 5b. This was accompanied by concomitant reductions in both transient energy growth

and the system transfer function norms in the controlled system. Note that the nearly parallel nature of the pairs of eigenvectors $\{\xi_2, \xi_3\}, \{\xi_4, \xi_5\}, \{\xi_6, \xi_7\}$ and $\{\xi_8, \xi_9\}$ in the uncontrolled case (Fig. 5a) is also reflected by the elliptical isolines surrounding the corresponding eigenvalues illustrated by the pseudospectra of Fig. 3.

Note the nonzero value of $\hat{v}$ at the walls in Fig. 5b; this reflects the wall blowing/suction applied as the control. Note also that half of the eigenvectors in Fig. 5a have zero $\hat{v}$ component. These are commonly referred to as the Squire modes of the system, and are decoupled from the perturbations in $\hat{v}$ because of the block of zeros in the upper-right corner of A . Such decoupling is not seen in Fig. 5b, because the closed-loop system matrix $A + B_2K$ is full.

3. Decentralization: designing for massive arrays

As illustrated in Figs. 6 and 7, there are two possible approaches for experimental implementation of linear compensators for this problem:

- (1) a centralized approach, applied in Fourier space, or
- (2) a decentralized approach, applied in physical space.




Fig. 5. The nine least stable eigenmodes of the closed-loop system matrix $A + B_2K$ for $k_x = 0$, $k_z = 2$ and $\text{Re} = 5000$. Plotted are the nonzero part of the $\hat{w}$ component of the eigenvectors (solid) and the nonzero part of the $\hat{v}$ component of the eigenvectors (dashed) as a function of y from the lower wall (bottom) to the upper wall (top). In (a), the dashed line is magnified by a factor of 1000 with respect to the solid line; in (b), the dashed line is magnified by a factor of 300. Note that the eigenvectors become significantly closer to orthogonal by the application of the control [12].




Fig. 6. Centralized approach to the control of plane channel flow in Fourier space.




Fig. 7. Decentralized approach to the control of plane channel flow in physical space.

Both of these approaches may be used to apply boundary control (such as distributions of blowing/suction) based on wall information (such as distributions of skin friction measurements). Both approaches may be used to implement the $\mathcal{H}_\infty$ compensators developed in Section 2, LQG/LTR compensators, PID feedback, or a host of other types of control designs. However, there are important differences in terms of the applicability of these two approaches to physical systems. The pros and cons of these approaches are now presented.

3.1. Centralized approach

The centralized approach is simplest in terms of its derivation, as most linear compensators in this geometry are designed in Fourier space, leveraging the spatially invariant structure of this system mentioned previously and the complete decoupling into Fourier modes which this structure provides [11]. As indicated in Fig. 6, implementation of this approach is straightforward. This type of experimental realization was recommended by Cortelezzi and Speyer [14] in related work. There are two major shortcomings of this approach:

- (A) the approach requires an on-line 2D FFT of the entire measurement vector and an on-line 2D iFFT of the entire control vector, and
- (B) the approach assumes spatial periodicity of the flow perturbations.

With regard to point A, it is important to note that the expense of centralized computations of 2D FFTs and iFFTs will grow rapidly with the size of the array of sensors and actuators; to be specific, the computational expense is proportional to $N_x N_z \log(N_x N_z)$. This will rapidly decrease the bandwidth possible as the array size (and the number of Fourier modes) is increased for a fixed speed of the central processing unit (CPU). Communication of signals to and from the CPU is also an important limiting factor as the array size grows. Thus, this approach does not extend well to massive arrays of sensors and actuators.

With regard to point B, it is important to note that transition phenomena in physical systems, such as boundary layers and plane channels, are not spatially periodic, though it is often useful to characterize the solutions of such systems with Fourier modes. The application of Fourier-space controllers which assume spatial periodicity in their formulation to physical systems which are not spatially periodic will be corrupted by Gibbs phenomenon, the well-known effect in which a Fourier transform is spoiled across all frequencies when the data one is transforming is not itself spatially periodic. In order to correct for this phenomenon in formulations which are based on Fourier-space computations of the control, windowing functions such as the Hanning window are appropriate. Windowing functions filter the

signals coming into the compensator such that they are driven to zero near the edges of the physical domain under consideration, thus artificially imposing spatial periodicity on the non-spatially-periodic measurement vector.

3.2. Decentralized approach

The decentralized approach, applied in physical space, is not as convenient to derive. Riccati equations of the size of the entire discretized 3D system pictured in Fig. 1 and governed by (1), represented in physical space, appear to be numerically intractable.

However, if such a problem could be solved, one would expect that the controller feedback kernels relating the state estimate $\hat{\mathbf{x}}$ inside the domain to the control forcing $\mathbf{u}$ at some point on the wall should decay quickly as a function of distance from the control point, as the control authority of any blowing/suction hole drilled into the wall on the surrounding flow decays rapidly with distance in a distributed viscous system.

Similarly, the estimator feedback kernels relating measurement errors $(\mathbf{y} - \hat{\mathbf{y}})$ at some point on the wall to the estimator forcing terms $\mathbf{v}$ on the system model inside the domain should decay as a function of distance from the measurement point, as the correlation of any two flow perturbation variables are known to decay with distance in a distributed viscous system.

Finally, due to the spatially invariant structure of the problem at hand, the control and estimation kernels for each sensor and actuator on the wall should be identical, though spatially shifted.

In other words, the physical-space kernels sought to determine the control and estimator feedback are spatially localized convolution kernels. If their spatial decay rate is rapid enough (e.g., exponential), then we will be able to truncate them at a finite distance from each actuator and sensor while maintaining a prescribed degree of accuracy in the feedback computation, resulting in spatially compact convolution kernels with finite support.

With such spatially compact convolution kernels, decentralized control of the present system becomes possible, as illustrated in Fig. 7. In such an approach, several tiles are fabricated, each with sensors, actuators, and an identical logic circuit. The computations on each tile are limited in spatial extent, with the individual logic circuit on each tile responsible for the (physical-space) computation of the state estimate only in the volume immediately above that tile. Each tile communicates its local measurements and state estimates with its immediate neighbors, with the number of tiles over which such information propagates in each direction depending on the tile size and spatial extent of the truncated convolution kernels. By replication, we can extend such an approach to arbitrarily large arrays of sensors and actuators. Though additional truncation of the kernels will disrupt the effectiveness of this control strategy near




Fig. 8. Localized controller gains relating the state estimate $\hat{\mathbf{x}}$ inside the domain to the control forcing $\mathbf{u}$ at the point $\{x = 0, y = -1, z = 0\}$ on the wall: visualized are a positive and negative isosurface of the convolution kernels for (left) the wall-normal component of velocity and (right) the wall-normal component of vorticity [27].




Fig. 9. Localized estimator gains relating the measurement error $(\mathbf{y} - \hat{\mathbf{y}})$ at the point $\{x = 0, y = -1, z = 0\}$ on the wall to the estimator forcing terms $\mathbf{v}$ inside the domain: visualized are a positive and negative isosurface of the convolution kernels for (left) the wall-normal component of velocity and (right) the wall-normal component of vorticity [27].

the edges of the array, such edge effects are limited to the edges in this case (unlike Gibbs phenomenon), and should become insignificant as the array size is increased.

4. Localization: relaxing nonphysical assumptions

As discussed previously, though the physical-space representation of the 3D linear system is intractable in the controls setting, the (completely decoupled) 1D systems at each wavenumber pair $\{k_x, k_z\}$ in the Fourier-space representation of this problem are easily managed. Remarkably, these two representations are *completely equivalent*. Performing a Fourier transform (which is simply a linear change of variables) of the *entire* 3D system (including the state, the controls, the measurements, and the disturbances) block diagonalizes *all* of the matrices involved in the 3D physical-space control problem. With such block diagonal structure, the constituent $\mathcal{H}_\infty$ control problems at each wavenumber pair $\{k_x, k_z\}$

may be solved independently and, once solved, reassembled in physical space with an inverse Fourier transform. If the numerics are handled properly, this approach is equivalent to solving the 3D physical-space control problem directly.

Recent theoretical work on this problem by Bamieh et al. [25], and related work by D'Andrea and Dullerud [26], further support the notion that an array of $\mathcal{H}_\infty$ compensators developed at each wavenumber pair, when inverse-transformed back to the physical domain, should in fact result in spatially localized convolution kernels with exponential decay. This exponential decay, in turn, allows truncation of the kernels to any prescribed degree of accuracy. Thus, if the truncated kernels are allowed to be sufficiently large in streamwise and spanwise extent, favorable closed-loop system properties, such as robust stability and reduced system transfer function norms, may be retained. Until very recently, however, it has not been possible to obtain such kernels for Navier–Stokes systems, due to an assortment of numerical challenges.

In Högberg and Bewley [27], spatially localized convolution kernels for both the control and estimation of plane channel flow have finally been obtained. The technique used was based on that described previously, deriving (in our initial efforts) $\mathcal{H}_2$ compensation at an array of wavenumber pairs $\{k_x, k_z\}$ and then inverse transforming the lot, with special attention paid to the details of the control formulation and the numerical method. In particular, a numerical discretization technique which was not plagued by spurious eigenvalues was chosen, and the control formulation was slightly modified such that the time derivative of the blowing/suction velocities is penalized in the cost function. The resulting localized kernels are illustrated in Figs. 8 and 9. Such kernels facilitate the decentralized control implementation discussed in Section 3.2 and depicted in Fig. 7, paving the way for experimental implementation with massive arrays of tiles integrating sensing, actuating, and the control logic.

Note that the control convolution kernels shown in Fig. 8 angle away from the wall in the *upstream* direction. Coupled with the mean flow profile indicated in Fig. 1, this accounts for the convective delay which requires us to anticipate flow perturbations on the interior of the domain with actuation on the wall somewhere downstream. The estimation convolution kernels shown in Fig. 9, on the other hand, extend well *downstream* of the measurement point. This accounts for the delay between the motions of the convecting flow structures on the interior of the domain and the eventual influence of these motions on the local drag profile on the wall; during this time delay, the flow structures responsible for these motions convect downstream. Note that the upstream bias of the control kernels and the downstream bias of the estimation kernels, though physically tenable, were not prescribed in the problem formulation. A posteriori study of the streamwise, spanwise, and wall-normal extent, the symmetry, and the shape of such control and estimation kernels provides us with a powerful new tool with which the fundamental physics of this distributed fluid-mechanical system may be characterized.

The localized convolution kernels illustrated in Figs. 8 and 9 are approximately independent of computational box size in which they were computed, so long as this box is sufficiently large. Thus, when implementing these kernels, we may effectively assume that they were derived in an *infinite*-sized box, relaxing the nonphysical assumption of spatial periodicity used in the problem formulation and modeling the physical situation of spatially evolving flow perturbations in a spatially invariant geometry and mean flow.

The localized convolution kernels illustrated in Figs. 8 and 9 are also approximately independent of the computational mesh resolution with which they were computed, so long as this computational mesh is sufficiently fine. Indeed, a computational mesh which is

sufficient to resolve the flow under consideration also adequately resolves these convolution kernels.

4.1. Open questions

As we have shown, the framework for decentralized $\mathcal{H}_\infty$ control of the fully resolved transition problem in the geometry depicted in Fig. 1 is now established. Obtaining spatial localization of the convolution kernels in physical space was the final remaining conceptual/numerical hurdle to be overcome. This work paves the way for decentralized application of such compensation with massive arrays of identical control tiles integrating sensing, actuation, and the control logic (Fig. 7). Though in some sense “complete”, this effort has also exposed several fundamental open questions, which will now be briefly discussed.

For a given choice of the matrices $\{B_1, B_2, C_1, C_2\}$ and design parameters $\{\ell, \alpha, \gamma > \gamma_0\}$ selected, decentralized $\mathcal{H}_\infty$ compensators may be determined using the procedure described above, and performance and robustness benchmarks may be obtained via simulation. As a final step in the control design process, it is of interest to explore how much the computational effort required by the logic on each tile may be reduced without significant degradation in the closed-loop system behavior. This can lead to a significant reduction in the number of floating point operations per second required by the logic circuit on each tile. However, as will be discussed in Section 5, compensator reduction in the decentralized setting remains a significant unsolved problem; standard reduction strategies which have been developed for finite, closed systems are not applicable and new research is motivated.

With the decentralized linear control framework established and prototypical numerical examples solved, we are now in a position to explore the effectiveness of compensators computed via this framework to the finite-amplitude perturbations that actually lead to transition, and to the “large”-amplitude perturbations of fully developed turbulence, in the nonlinear equations of fluid motion. An extensive analytical and numerical study within this framework is underway. Issues regarding our preliminary efforts in this direction are briefly reviewed in Section 6. As emphasized in the preface, such a study should be guided by an interdisciplinary perspective in order to be maximally successful. Specifically, such a study should fully incorporate the known or postulated linear mechanisms leading to transition or, in the case of turbulence, the linear mechanisms thought to be at least partially responsible for sustaining the turbulent cascade of energy. In addition, this effort motivates the development of new analytical tools which might help clarify the types of state disturbances and flow perturbations which are particularly important in such phenomena. Armed with such an understanding, large benefits might be realized in the compensator design, as

the modeling of the structure of the state disturbances exciting the system, G_1 , and the weighting on the flow perturbations of interest in the cost function, Q , are important design criteria. In fact, we fully expect that the transfer of information between our physical understanding of fundamental flow phenomena and our knowledge of how to control such phenomena will be a two-way transfer. Such a strategy promises to provide powerful new tools for obtaining fundamental physical understanding of classical problems in fluid mechanics as we gain new insight in how to modify these phenomena by the action of control feedback.

A host of other canonical flow control problems, including the control of spatially developing boundary layers, bluff body flows, and free shear layers, should also be amenable to linear control application using the framework outlined here. A few such extensions are discussed briefly in Section 7.

5. Compensator reduction: eliminating unnecessary complexity

Strategies for the development of reduced-order decentralized compensators of the present form remain a key unsolved issue. With the present approach, as described previously, a physical-space state estimate in the volume immediately above each tile must be updated online by the logic circuit on each tile as the flow evolves. However, it is not at all necessary for the compensator to compute an accurate state estimate as an intermediate variable; indeed, our only requirement is that, based on whatever filtered information the dynamic compensator does extract from the noisy system measurements, suitable controls may be determined to achieve the desired closed-loop system behavior. It should be possible to reduce substantially the complexity of the dynamic compensator and still achieve this more modest objective.

There are two possible representations in which the complexity of the compensator can be reduced: in Fourier space (where the compensator is designed), or in physical space (where the decentralized compensation is applied).

5.1. Fourier-space compensator reduction

At any particular wavenumber pair $\{k_x, k_z\}$, there is one actuator variable at each wall, one sensor variable at each wall, and a spatial discretization in y of the state variables across the domain stretching between these walls. Due to the complete decoupling of the control problem into separate Fourier modes, the system model used in the estimator at each particular wavenumber pair is not referenced by the compensator at any other wavenumber pair. Thus, the compensators at each wavenumber pair are completely decoupled and may be

reduced independently. At certain wavenumber pairs, it might be important to retain several degrees of freedom in the dynamic compensator, while at other wavenumber pairs, it might be possible to retain significantly fewer degrees of freedom without significant degradation in the closed-loop system behavior. Several existing compensator reduction strategies are well suited to this problem, and their application in this setting is straightforward. Cortezzi and Speyer [14] successfully applied the balanced truncation technique of open-loop model reduction in this Fourier-space framework in order to facilitate the design of a reduced-complexity dynamic compensator.

As mentioned earlier, it is the nonorthogonality of the entire set of system eigenvectors which leads to the peculiar (and important) possibilities for energy amplification in these systems, so compensator reduction techniques which are mindful of the relevant transfer function norms are necessary. In addition, as eloquently described by Obinata and Anderson [28], it is most appropriate when designing low-order compensators for high-order plants to reduce the compensator while accounting for how it performs in the closed loop. An assortment of closed-loop compensator reduction techniques are now available, and should be tested in future work.

In the setting of designing a decentralized compensator, there is an important shortcoming to performing standard compensator reductions in Fourier space. As the compensator reduction problem is independent at each wavenumber pair, we might be left with a different number of degrees of freedom in the reduced-order compensator at each wavenumber pair, leaving us with a dynamical system model which is impossible to inverse transform back into the physical domain. Even if we restrict the compensator reduction algorithm to reduce to the same number of degrees of freedom at each wavenumber pair (a restrictive assumption which should be unnecessary), there appears to be no appropriate strategy currently available to coordinate this reduction process across all wavenumbers in a consistent manner such that the inverse transform of the reduced dynamic model is spatially localized. Without such coordination, it seems inevitable that the ordering and representation of the various modes of this dynamic model will be scrambled during the process of compensator reduction at each wavenumber pair, resulting in an inverse transform back in physical space that does not exhibit the spatial localization which is essential to facilitate decentralized control.

5.2. Physical-space compensator reduction

As an alternative to Fourier-space compensator reduction, one might consider instead the reduction of the physical-space model and its associated localized convolution kernels. This has several advantages linked

to the fact that this is the actual compensation to be computed on each tile. The first advantage is that spatial localization will be retained, as compensator reduction is applied after the localized kernels are obtained. Another important advantage is that this setting allows us to keep more degrees of freedom in the dynamical system model to represent streamwise and spanwise fluctuations of the state near the wall than we retain to represent the behavior of the state on the interior of the domain. This effectively relaxes the restrictive assumption referred to in the previous paragraph. Such an emphasis on resolving the state near the wall is motivated by inspection of the convolution kernels plotted in Figs. 8 and 9, in which it is clear that the details of the flow near the wall are of increased importance when computing the feedback.

However, note that the system model simulated on each individual tile is not self-contained, due to the interconnections with neighboring tiles indicated in Fig. 7. Thus, if one reduces the system model above a single tile, all neighboring tiles which reference this state estimate will be affected. As the system model is not self-contained, as it was in the Fourier-space case, existing compensator reduction approaches are not applicable.

An important observation, however, is that the structure of the system model carried by each tile is identical. Due to the repeated structure of the model represented on the array, it is sufficient to optimize the system model carried by a single tile. The repeated structure of the distributed physical-space model should make the compensator reduction problem tractable. This fundamental problem of reducing distributed, interconnected dynamic compensators in the decentralized closed-loop setting remains, as yet, unsolved.

5.3. *Non-spatially-invariant systems*

Finally, it should be stated that the Fourier-space decoupling leveraged at the outset of this problem formulation has been one of the key ingredients which has permitted accurate solution of well-resolved canonical flow control problems to date. The linear control technique we have used to solve these control problems involves the solution of matrix Riccati equations, which are accurately soluble for state dimensions only up to $O(10^3)$. As we move to more applied flow control problems in which such Fourier-space decoupling is either more restrictive or not available, if we continue to use Riccati-based control approaches, creative new compensator reduction strategies will be required. We might need to apply “open-loop” model reduction strategies (in advance of computing the control feedback and closing the loop) in order to make manageable the dimension of the Riccati equations to be solved in the compensator design. As mentioned earlier, it is most appropriate when designing low-order compensators for high-order plants

to reduce the compensator while accounting for how it performs in the closed loop. Unfortunately, extremely high-order discretizations of non-spatially-invariant PDE systems will not likely afford us this luxury, as such systems do not decouple (via Fourier transforms) into constituent lower-order control problems amenable to matrix-based compensator design strategies.

6. *Extrapolation: linear control of nonlinear systems*

Once a decentralized linear compensator of the present form is developed, a verification of its utility for the transition problem may be obtained by applying it to the laminar flow depicted in Fig. 1 with either finite-amplitude (but sufficiently small) initial flow perturbations and/or finite-amplitude (but sufficiently small) applied external disturbances. The resulting finite-amplitude flow perturbations are governed by the fully nonlinear Navier–Stokes equation, and have been simulated in well resolved direct numerical simulations (DNS) with the code benchmarked in Bewley et al. [30]. Representative simulations are indicated in Fig. 10, indicating that linear compensators can indeed relaminarize perturbed flows that would otherwise proceed rapidly towards transition to turbulence. With the framework presented here, extensive numerical studies promise to significantly extend our fundamental understanding of the process of transition and how this process may be inhibited by control feedback.

It is also of interest to consider the application of decentralized linear compensation to the fully nonlinear problem of a turbulent flow, such as that shown in Fig. 11. The first reason to try such an approach is simply because we can: linear control theory leads to implementable control algorithms and grants a lot of flexibility in the compensator design. Nonlinear turbulence control strategies, though currently under active development (see Sections 8–11), are much more difficult to design and implement, and require substantial further research before they will provide implementable control strategies as flexible and powerful as those which we currently have at our disposal in the linear setting.

There is at least some evidence in the fluids literature that applying linear control feedback to turbulence might be at least partially effective. Though the significance of this result has been debated in the fluid mechanics community, Farrell and Ioannou [31] have clearly shown that linearized Navier–Stokes systems in plane channel flows, when excited with the appropriate stochastic forcing, exhibit behavior which is reminiscent of the streamwise vortices and streamwise streaks which characterize actual near-wall turbulence. Whatever information the linearized Navier–Stokes equation actually contains about the mechanisms sustaining these turbulence structures, the present linear control framework (perhaps




Fig. 10. Evolution of oblique waves (left) and an initially random flow perturbation (right) added to a laminar flow at $Re = 2000$, with and without decentralized linear control feedback. The magnitude of the initial flow perturbations in these simulations greatly exceed the thresholds reported by Reddy et al. [29] that lead to transition in an uncontrolled flow (by a factor of 225 for the oblique waves and by a factor of 15 for the random initial perturbation). Solid lines indicate the energy evolution in the controlled case, dashed lines indicate the energy evolution in the uncontrolled case. Both of the uncontrolled systems lead quickly to transition to turbulence whereas, both of the controlled systems relaminarize. For the controlled cases, initial perturbations with greater energy fail to relaminarize, whereas initial perturbations with less energy relaminarize earlier [27].




Fig. 11. Visualization of the coherent structures of uncontrolled near-wall turbulence at $Re_\tau = 180$. Despite the geometric simplicity of this flow (see Fig. 1), it is phenomenologically rich, and is characterized by a large range of length scales and time scales over which energy transport and scalar mixing occur. The relevant spectra characterizing these complex nonlinear phenomena are continuous over this large range of scales, and thus such flows have largely eluded accurate description via dynamic models of low state dimension. The nonlinearity, the distributed nature, and the inherent complexity of its dynamics make turbulent flow systems particularly challenging for successful application of control theory. (Simulation by Bewley et al. [30].)




Fig. 12. Example of the spectacular failure of linear control theory to stabilize a simple nonlinear chaotic convection system governed by the Lorenz equation. Plotted are the regions of attraction to the desired stationary point (blue) and to an undesired stationary point (red) in the linearly controlled nonlinear system, and typical trajectories in each region (black and green, respectively). The cubical domain illustrated is $\Omega = (-25, 25)^3$ in all subfigures; for clarity, different viewpoints are used in each subfigure [34].

restricted to a finite horizon) should be able to exploit. Though the life cycle of the near-wall coherent structures of turbulence appears to involve important nonlinear phenomena (see, e.g. [32]), that in itself does not disqualify the utility of linear control strategies to effectively disrupt critical linear terms of this nonlinear process. Indeed, recent numerical experiments by Kim and Lim [33] support this idea by conclusively demonstrating the importance of the coupling term C in the linearized system matrix A (see Eq. (10)) for maintaining near-wall turbulence in nonlinear simulations.

In order to understand the possible pitfalls of applying linear feedback to nonlinear systems, a low-order nonlinear convection problem governed by the Lorenz equation was studied by Bewley [34]. As with the problem of turbulent channel flow, but in a low-order system easily amenable to analysis, control feedback was determined with linear control theory by linearizing the governing equation about a desired fixed point. Once a linear controller was determined by such an approach, it was then applied directly to the fully nonlinear system. The result is depicted in Fig. 12.

For control feedback determined by linear control theory with a large weighting ℓ on the control effort, direct application of linear feedback to the full nonlinear system stabilizes both the desired state and an undesired state, indicated by the two trajectories marked in Fig. 12(a). An unstable manifold exists between these two states, indicated by the contorted surface shown. Any initial state on one side of this manifold will converge to the desired state, and any initial state on the other of this manifold will converge to the undesired state.

As seen in Figs. 12(b) and (c), as the weighting on the control effort, ℓ , is turned down and the desired stationary state is stabilized more aggressively, the domain of convergence to the undesired state remains large. This undesired state is “aggravated” by the enhanced control feedback, moving farther from the origin. The undesired state eventually escapes to infinity for sufficiently small ℓ , indicating instability of the nonlinear system from a wide range of initial conditions even though the desired stationary point is endowed with a high degree of linear stability. Implication: strong linear stabilization of a desired system state (such as laminar flow) will not necessarily eliminate undesired nonlinear system behavior (such as turbulence) in a chaotic system.

Some form of nonlinearity in the feedback rule was required to eliminate this undesired behavior. One effective technique is to apply a switch such that the linear control feedback is turned on only when the state $\mathbf{x}(t)$ is within some sufficiently small neighborhood of the desired stabilized state $\bar{\mathbf{x}}$ in the linearly-controlled system. The chaotic dynamics of the uncontrolled Lorenz system will bring the system into this neighborhood in finite time, after which control may be applied to “catch” the system at the desired equilibrium state.

Thus, even in this simple model problem, linear feedback can have a destabilizing influence if applied outside the neighborhood for which it was designed. For the full Navier–Stokes problem, though a certain set of linear feedback gains might stabilize the laminar state, on the “other side of the manifold” might lie a turbulent state which is aggravated by the same linear controls. Application of linear control to nonlinear chaotic systems must therefore be done with vigilance, lest nonlinearities




Fig. 13. Evolution of fully developed turbulence at $Re_\tau = 100$ with and without decentralized linear control feedback. Note that this flow has approximately the same mass flux as the laminar flow at $Re = 2000$. Top: energy of flow perturbation. Middle: drag (note approximately 25% reduction in the controlled cases). Bottom: control effort used. The uncontrolled energy and drag are the (upper) solid lines in the top and middle figures. A gain scheduling approach is used to tune the control feedback gains to the instantaneous mean flow profile [27].

destabilize the closed-loop system, as shown here. The easy fix found for this low-order model problem (that is, simply turn off the control until the chaotic dynamics bring the state into a neighborhood of the desired state) might not be available for the (high-dimensional) problem of turbulence, as fully turbulent flows appear to remain at all times far from the laminar state.

In our preliminary attempts at applying the decentralized compensators developed above to turbulence, we have succeeded in reducing the drag of a fully developed turbulent flow by 25% with state-feedback controllers, as shown in Fig. 13. Interestingly, for the choice of control parameters selected here, there is no evidence of an aggravated turbulent state. A 25% drag reduction, though significant, is comparable to the drag reductions obtained with a variety of other ad hoc control approaches in this flow. We are actively pursuing modification of this linear control feedback to improve upon this result. Interdisciplinary considerations, such as those involved in the design of linear compensation for the problem of transition, are essential in this effort. Specifically, the (unmodeled) nonlinear terms in the Navier–Stokes equation provide insight as to structure of the disturbances,

G_1 , to be accounted for in the linear control formulation in order to best compensate for their unmodeled effects. Additionally, the coherent structures of fully developed near-wall turbulence, believed to be a major player in the self-sustaining nonlinear process of turbulence generation near the wall, provide a phenomenological target which may be exploited in the selection of the weighting on the flow perturbations, Q , in the cost function.

7. Generalization: extending to spatially developing flows

Extension of the decentralized linear control framework developed here to a large class of slightly nonparallel flows is heuristic but straightforward. To accomplish this, the parabolic mean flow profile $U(y)$ indicated in Fig. 1 is replaced with an appropriate “quasi-1D” profile, such as the Blasius boundary layer profile. As long as the mean flow profile evolves slowly enough in space (as compared to the wavelengths of the significant instabilities in the problem), it may be assumed to be constant in space for the purpose of developing

the linear control feedback. Such an assumption of slow spatial divergence forms the foundation of the study of local and global modes used in the characterization of absolute and convective instabilities [35], and has proven to be a quite powerful concept. For the appropriate flows, we believe this concept is also appropriate in the context of the development of control feedback.

Implementation of the decentralized control concept in this setting is a heuristic extension of the approach presented in Fig. 7: gradual variations in the mean flow are accounted for by local extension of the mean flow profile in the compensator derivation for each tile, gradually scaling the compensation rules from one tile to the next as the flow develops downstream. For example, we may consider developing this strategy for the laminar boundary layer (LBL) solutions of the Falkner–Skan–Cooke family, found by solving the ODE:

$$f''' + ff'' + \beta(1 - f'^2) = 0$$

with $f(0) = f'(0) = 0$ and $f'(\infty) \rightarrow 1$ and defining

$$U = U_0 f'(\eta) \quad \text{and} \quad V = \sqrt{\frac{\nu U_0}{2x}} \left[\eta f'(\eta) - f(\eta) \right].$$

Cases of interest include the Blasius profile, modeling a zero-pressure-gradient flat-plate LBL with

$$U_0 = U_\infty, \quad \beta = 0, \quad \eta = y \sqrt{\frac{U_0}{2\nu x}},$$

the Falkner–Skan profile, modeling a nonzero-pressure-gradient LBL or wedge flow by taking

$$U_0 = Kx^m, \quad \beta = \frac{2m}{1+m}, \quad \eta = y \sqrt{\frac{(m+1)U_0}{2\nu x}},$$

and Falkner–Skan–Cooke profile, which models the addition of sweep to the leading edge by solving the supplemental ODE

$$g'' + fg' = 0$$

with $g(0) = 0$ and $g(\infty) \rightarrow 1$ and defining $W = W_\infty g(\eta)$. Note that the self-similarity of the LBL profiles might lead to simplified parameterizations of the convolution kernels for the control and estimation problems. Extension of this approach to a variety of other spatially developing flows (self-similar or otherwise) should also be straightforward.

8. Nonlinear optimization: local solutions for full Navier–Stokes

Given an idealized setting of full state information, no disturbances, and extensive computational resources, significant finite-horizon optimization problems may be formulated and (locally) solved for complex nonlinear systems using iterative, adjoint-based, gradient optimiza-

tion strategies. Such optimization problems can now be solved for high-dimensional discretizations of turbulent flow systems, incorporating the full nonlinear Navier–Stokes equation, locally minimizing cost functionals representing a variety of control problems of physical interest within a given space of feasible control variables. The mathematical framework for such optimizations will be reviewed briefly in Section 8.1, and is described in greater detail by Bewley et al. [30].

The optimizations obtained via this approach are, strictly speaking, only “local” over the domain of feasible controls (that is, unless restrictive assumptions are made in the formulation of the control problem). Thus, the performance obtained via this approach can usually not be guaranteed to be “globally optimal”. However, the performance obtained with such nonlinear optimizations often far exceeds that possible with other control design approaches (see, e.g. Fig. 14). In addition, this approach is quite flexible, as it can be used to iteratively improve high-dimensional control distributions directly, as will be illustrated below, or, alternatively, to optimize open-loop forcing schedules, shape functions, or the coefficients of practical, implementable, and possibly nonlinear feedback control rules. Thus, interest in adjoint-based optimization strategies for turbulent flow systems goes far beyond that of establishing performance benchmarks via predictive optimizations of the control distribution itself. Establishing such benchmarks is only a first step towards a much wider range of applications for adjoint-based tools in turbulent flow systems.

The general idea of this approach, often referred to as model predictive control, is well motivated by comparing and contrasting it to massively parallel brute-force algorithms recently developed to play the game of chess. The goal when playing chess is to capture the other player’s king through an alternating series of discrete moves with the opponent: at any particular turn, a player has to select one move out of at most 20 or 30 legal alternatives.

To accomplish its optimization, a computer program designed to play the comparatively “simple” game of chess, such as *Deep Blue* [36], must, in the worst case, plan ahead by iteratively examining a tree of possible evolutions of the game several moves into the future [37], a strategy based on “function evaluations” alone. At each step, the program selects that move which leads to its best expected outcome, given that the opponent is doing the same in a truly noncooperative competition. The version of *Deep Blue* that defeated Garry Kasparov in 1997 was able to calculate up to 200 billion moves in the 3 min it was allowed to conduct each turn. Even with this extreme number of function evaluations at its disposal on this relatively simple problem, the algorithm was only about an even match with Kasparov’s human intuition.

An improved algorithm to those based on function evaluations alone, suitable for optimizing the present problem in a reasonable amount of time, is available




Fig. 14. Performance of optimized blowing/suction controls for formulations based on minimizing $\mathcal{J}_o(\phi)$, case *c* (see Section 8.1.2), as a function of the optimization horizon T^+ . The direct numerical simulations of turbulent channel flow reported here were conducted at $Re_\tau = 100$. For small optimization horizons ($T^+ = O(1)$, sometimes called the “suboptimal approximation”), approximately 20% drag reduction is obtained, a result which can be obtained with a variety of other approaches. For sufficiently large optimization horizons ($T^+ \gtrsim 25$), the flow is returned to the region of stability of the laminar flow, and the flow relaminarizes with no further control effort required. No other control algorithm tested in this flow to date has achieved this result with this type of flow actuation [30].

because (i) we know the equation governing the evolution of the present system, and (ii) we can state the problem of interest as a functional to be minimized. Taking these two facts together, we may devise an iterative procedure based on gradient information, derived from an *adjoint field*, to optimize the controls for the desired purpose on the prediction horizon of interest in an efficient manner. Only by exploiting such gradient information can the high-dimensional optimization problem at hand (up to $O(10^7)$ control variables per optimization horizon in some of our simulations) be made tractable.

8.1. Adjoint-based optimization approach

8.1.1. Governing equation

The problem we consider here is the control of a fully developed turbulent channel flow with full flow-field information and copious computational resources available to the control algorithm. The flow is governed by the incompressible Navier–Stokes equation inside a three-dimensional rectangular domain (Fig. 15) with unsteady wall-normal velocity boundary conditions ϕ applied on the walls as the control. Three vector fields are first defined: the flow state $\mathbf{q}$, the flow perturbation state $\mathbf{q}'$, and the adjoint state $\mathbf{q}^*$:

$$\mathbf{q}(\mathbf{x}, t) = \begin{pmatrix} p(\mathbf{x}, t) \\ \mathbf{u}(\mathbf{x}, t) \end{pmatrix}, \quad \mathbf{q}'(\mathbf{x}, t) = \begin{pmatrix} p'(\mathbf{x}, t) \\ \mathbf{u}'(\mathbf{x}, t) \end{pmatrix},$$

$$\mathbf{q}^*(\mathbf{x}, t) = \begin{pmatrix} p^*(\mathbf{x}, t) \\ \mathbf{u}^*(\mathbf{x}, t) \end{pmatrix}.$$




Fig. 15. Channel flow geometry. The interior of the domain is denoted Ω and the boundaries of the domain in the x_i direction are denoted $\Gamma_i^\pm$. Unsteady wall-normal velocity boundary conditions are applied on the walls $\Gamma_2^\pm$ as the control, with periodic boundary conditions applied in the streamwise direction x_1 and spanwise direction x_3 . An external pressure gradient is applied to induce a mean flow in the x_1 direction.

Each of these vector fields is composed of a pressure component and a velocity component, all of which are continuous functions of space, $\mathbf{x}$, and time, t . The velocity components themselves are also vectors, with components in the streamwise direction x_1 , the wall-normal direction x_2 , and the spanwise direction x_3 . Partial differential equations governing all three of these fields will be derived in due course, and the motivation for introducing $\mathbf{q}'$ and $\mathbf{q}^*$ will be given as the need for these fields arises in the control derivation. Only after the optimization approach has been derived completely in differential form is it discretized in space and time; an alternative strategy, discretizing the state equation in space

before determining the adjoint operator, is discussed in Section 8.2.

The governing equation is written as

$$\mathcal{N}(\mathbf{q}) = \mathbf{F} \quad \text{in } \Omega, \quad (11a)$$

$$\mathbf{u} = -\phi \mathbf{n} \quad \text{on } \Gamma_2^\pm, \quad (11b)$$

$$\mathbf{u} = \mathbf{u}_0 \quad \text{at } t = 0, \quad (11c)$$

where $\mathcal{N}(\mathbf{q})$ is the (nonlinear) Navier–Stokes operator

$$\mathcal{N}(\mathbf{q}) = \left(\frac{\partial u_i}{\partial t} + \frac{\partial u_j u_i}{\partial x_j} - \nu \frac{\partial^2 u_i}{\partial x_j^2} + \frac{\partial p}{\partial x_i} \right),$$

$\mathbf{F}$ is a forcing vector accounting for an externally applied mean pressure gradient driving the flow in the streamwise direction, and $\mathbf{n}$ is the unit outward normal to the boundary $\partial\Omega$. The boundary conditions on the state $\mathbf{q}$ are periodic in the streamwise and spanwise directions. A wall-normal control velocity ϕ is distributed over the walls as indicated, and is constrained to inject zero net mass such that, $\forall t$, $\int_{\Gamma_2^+} \phi \, d\mathbf{x} = \int_{\Gamma_2^-} \phi \, d\mathbf{x} = 0$. Initial conditions on the velocity, $\mathbf{u}_0$, of fully developed turbulent channel flow are prescribed.

8.1.2. Cost functional

As in the linear setting, an essential step in the framing of the nonlinear optimization problem is the representation of the control objective as a cost functional to be minimized. Several cases of physical interest may be represented by a cost functional of the generic form

$$\begin{aligned} \mathcal{J}_0(\phi) = & \frac{1}{2} \int_0^T \int_\Omega |\mathcal{C}_1 \mathbf{u}|^2 \, d\mathbf{x} \, dt + \frac{1}{2} \int_\Omega |\mathcal{C}_2 \mathbf{u}(x, T)|^2 \, d\mathbf{x} \\ & - \int_0^T \int_{\Gamma_2^\pm} \mathcal{C}_3 \nu \frac{\partial \mathbf{u}}{\partial \mathbf{n}} \cdot \mathbf{r} \, d\mathbf{x} \, dt + \frac{\ell^2}{2} \int_0^T \int_{\Gamma_2^\pm} |\phi|^2 \, d\mathbf{x} \, dt. \end{aligned}$$

Four cases of particular interest are:

- (a) $\mathcal{C}_1 = d_1 I$ and $\mathcal{C}_2 = \mathcal{C}_3 = 0 \Rightarrow$ regulation of turbulent kinetic energy;
- (b) $\mathcal{C}_1 = d_2 \nabla \times$ and $\mathcal{C}_2 = \mathcal{C}_3 = 0 \Rightarrow$ regulation of the square of the vorticity;
- (c) $\mathcal{C}_2 = d_3 I$ and $\mathcal{C}_1 = \mathcal{C}_3 = 0 \Rightarrow$ terminal control of turbulent kinetic energy;
- (d) $\mathcal{C}_3 = d_4 I$ and $\mathcal{C}_1 = \mathcal{C}_2 = 0 \Rightarrow$ minimization of the time-average skin friction in the direction $\mathbf{r}$ integrated over the boundary of the domain, where $\mathbf{r}$ is a unit vector in the streamwise direction.

All four of these cases, and many others, may be considered in the present framework, and the extension to other cost functionals is straightforward. The dimensional constants d_i (which are the appropriate functions of the kinematic viscosity, the channel width, and the bulk velocity), as well as ℓ , are included to make the cost

functional dimensionally consistent and to account for the relative weight of each individual term.

Clearly, in both the chess problem and the turbulence problem, the further into the future one can optimize the problem the better (Fig. 14); however, both problems get exponentially harder to optimize as the prediction horizon is increased. Since only intermediate-term optimization is tractable, it is not always the best approach to represent the final objective in the cost functional. In the chess problem, though the final aim is to capture the other player's king, it is most effective to adopt a mid-game strategy of establishing good board position and achieving material advantage. Similarly, if the turbulence control objective is reducing drag, it was found in [30] that it is most effective along the way to minimize a finite-horizon cost functional related to the turbulent kinetic energy of the flow, as the turbulent transport of momentum is responsible for inducing a substantial portion of the drag in a turbulent flow. In a sense, turbulence is the “cause” and high drag is the “effect”, and it is most effective to target the “cause” in the cost functional when optimizations on only intermediate prediction horizons are possible.

In addition, a smart optimization algorithm allows for excursions in the short term if it leads to a long-term advantage. For example, in chess, a good player is willing to sacrifice a lesser piece if, by so doing, a commanding board position is attained and/or a restoring exchange is forced a few moves later. Similarly, by allowing a turbulence control scheme to increase (temporarily) the turbulent kinetic energy of a flow, a transient may ensue which, eventually, effectively diminishes the strength of the near-wall coherent structures. It was found in [30] that terminal control strategies, aimed at minimizing the turbulence only at the end of each optimization period, have a decided advantage over regulation strategies, which penalize excursions of the turbulent kinetic energy over the entire prediction horizon.

8.1.3. Gradient of cost functional

As suggested by Abergel and Temam [38], a rigorous procedure may be developed to determine the sensitivity of a cost functional $\mathcal{J}$ to small modifications of the control ϕ for nonlinear problems of this sort. To do this, consider the perturbation to the cost functional resulting from a small perturbation to the control ϕ in the direction ϕ' . (Note that this control perturbation direction ϕ' is arbitrary and scaled to have unit norm.) Define $\mathcal{J}'$ as the Fréchet differential [39] of a cost functional $\mathcal{J}$ such that

$$\begin{aligned} \mathcal{J}' &\triangleq \lim_{\varepsilon \rightarrow 0} \frac{\mathcal{J}(\phi + \varepsilon \phi') - \mathcal{J}(\phi)}{\varepsilon} \\ &\triangleq \int_0^T \int_{\Gamma_2^\pm} \frac{\mathcal{D} \mathcal{J}(\phi)}{\mathcal{D} \phi} \phi' \, dt \, d\mathbf{x}. \end{aligned}$$

The quantity $\mathcal{J}'$ is the cost functional perturbation due to a control perturbation $\varepsilon\phi'$ scaled by the inverse of the control perturbation magnitude ε in the limit that $\varepsilon \rightarrow 0$. The above relation, considered for arbitrary ϕ' , also defines the gradient of the cost functional $\mathcal{J}$ with respect to the control ϕ , which is written $\mathcal{D}\mathcal{J}(\phi)/\mathcal{D}\phi$.

In the present approach, the cost functional perturbation $\mathcal{J}'$ defined above will be expressed as a simple linear function of the direction of the control perturbation ϕ' through the solution of an adjoint problem. By the above formula, such a representation then reveals the gradient direction $\mathcal{D}\mathcal{J}(\phi)/\mathcal{D}\phi$ directly. With this gradient information, the control ϕ is updated on $(0, T]$ in the direction that, at least locally (i.e. for infinitesimal control updates), most effectively reduces the cost functional. The finite distance the control is updated in this direction is then found by a line search routine, which makes this iteration procedure stable even when controlling nonlinear phenomena. The flow resulting from this modified control is then computed according to the (nonlinear) Navier–Stokes equation (11), the sensitivity of this new flow to further control modification is computed, and the process repeated. Upon convergence of this iteration, the flow is advanced over the interval $(0, T_1]$, where $T_1 \leq T$, and an iteration for the optimal control over a new time interval $(T_1, T_1 + T]$ begins anew.

The cost functional perturbation $\mathcal{J}'_0$ resulting from a control perturbation in the direction ϕ' is given by

$$\begin{aligned} \mathcal{J}'_0(\phi) &= \int_0^T \int_{\Omega} \mathcal{C}_1^* \mathcal{C}_1 \mathbf{u} \cdot \mathbf{u}' \, d\mathbf{x} \, dt + \int_{\Omega} (\mathcal{C}_2^* \mathcal{C}_2 \mathbf{u} \cdot \mathbf{u}')_{t=T} \, d\mathbf{x} \\ &\quad - \int_0^T \int_{\Gamma_{\pm}^{\pm}} \nu \mathcal{C}_3^* \mathbf{r} \cdot \frac{\partial \mathbf{u}'}{\partial \mathbf{n}} \, d\mathbf{x} \, dt + \ell^2 \int_0^T \int_{\Gamma_{\pm}^{\pm}} \phi \phi' \, d\mathbf{x} \, dt \\ &\triangleq \int_0^T \int_{\Gamma_{\pm}^{\pm}} \frac{\mathcal{D}\mathcal{J}_0(\phi)}{\mathcal{D}\phi} \phi' \, d\mathbf{x} \, dt, \end{aligned}$$

where $\mathbf{u}'$ is the Fréchet differential of $\mathbf{u}$, as defined in the following subsection. Adjoint calculus is used simply to re-express the integrals involving $\mathbf{u}'$ as a linear function of ϕ' . Once this is accomplished, ϕ' is factored out of the integrands and, as the equation holds for arbitrary ϕ' , an expression for the gradient $\mathcal{D}\mathcal{J}_0(\phi)/\mathcal{D}\phi$ is identified.

8.1.4. Linearized perturbation field

Now consider the linearized perturbation $\mathbf{q}'$ to the flow $\mathbf{q}$ resulting from a perturbation ϕ' to the control ϕ . Again, the quantity $\mathbf{q}'$ may be defined by the limiting process of a Fréchet differential such that

$$\mathbf{q}' \triangleq \lim_{\varepsilon \rightarrow 0} \frac{\mathbf{q}(\phi + \varepsilon\phi') - \mathbf{q}(\phi)}{\varepsilon}.$$

For the purpose of gaining physical intuition, it is useful to note that the quantity $\mathbf{q}'$, described above as a differential quantity, may instead be defined as the small perturbation to the state $\mathbf{q}$ arising from a small control

perturbation ϕ' to the control ϕ . In such derivations, the notations $\delta\phi$ and $\delta\mathbf{q}$, denoting small perturbations to ϕ and $\mathbf{q}$, are used instead of the differential quantities ϕ' and $\mathbf{q}'$. The two derivations are roughly equivalent, though the present derivation does not assume that primed quantities are small.

The equation governing the dependence of the linearized flow perturbation $\mathbf{q}'$ on the control perturbation ϕ' may be found by taking the Fréchet differential of the state equation (11). The result is

$$\mathcal{N}'(\mathbf{q})\mathbf{q}' = 0 \quad \text{in } \Omega, \quad (12a)$$

$$\mathbf{u}' = -\phi' \mathbf{n} \quad \text{on } \Gamma_{\pm}^{\pm}, \quad (12b)$$

$$\mathbf{u}' = 0 \quad \text{at } t = 0, \quad (12c)$$

where the linearized Navier–Stokes operation $\mathcal{N}'(\mathbf{q})\mathbf{q}'$ is given by

$$\mathcal{N}'(\mathbf{q})\mathbf{q}' = \begin{pmatrix} \frac{\partial u'_j}{\partial x_j} \\ \frac{\partial u'_i}{\partial t} + \frac{\partial}{\partial x_j} (u_j u'_i + u'_j u_i) - \nu \frac{\partial^2 u'_i}{\partial x_j^2} + \frac{\partial p'}{\partial x_i} \end{pmatrix}.$$

The operation $\mathcal{N}'(\mathbf{q})\mathbf{q}'$ is a linear operation on the perturbation field $\mathbf{q}'$, though the operator $\mathcal{N}'(\mathbf{q})$ is itself a function of the solution $\mathbf{q}$ of the Navier–Stokes problem. Eq. (12) thus reflects the linear dependence of the perturbation field $\mathbf{q}'$ in the interior of the domain on the control perturbation ϕ' at the boundary. However, the implicit linear relationship $\mathbf{q}' = \mathbf{q}'(\phi')$ given by this equation is not yet tractable for expressing $\mathcal{J}'_0$ in a simple form from which $\mathcal{D}\mathcal{J}_0(\phi)/\mathcal{D}\phi$ may be deduced. For the purpose of determining a more useful relationship with which we may determine $\mathcal{D}\mathcal{J}_0(\phi)/\mathcal{D}\phi$, we now appeal to an adjoint identity.

8.1.5. Statement of adjoint identity

This subsection derives the adjoint of the linear partial differential operator $\mathcal{N}'(\mathbf{q})$. For readers not familiar with this approach, a review of the derivation of an adjoint operator for a very simple case in the present notation is given in Appendix A of [30]. The adjoint derivation presented below extends in a straightforward manner to more complex equations, such as the compressible Euler equation, as shown in Appendix B of [30] (again, using the same notation). Such generality highlights the versatility of the present approach.

Define an inner product over the domain in space-time under consideration such that

$$\langle \mathbf{q}^*, \mathbf{q}' \rangle = \int_0^T \int_{\Omega} \mathbf{q}^* \cdot \mathbf{q}' \, d\mathbf{x} \, dt$$

and consider the identity

$$\langle \mathbf{q}^*, \mathcal{N}'(\mathbf{q})\mathbf{q}' \rangle = \langle \mathcal{N}'(\mathbf{q})^* \mathbf{q}^*, \mathbf{q}' \rangle + b. \quad (13)$$

Integration by parts may be used to move all differential operations from $\mathbf{q}'$ on the left-hand side of (13) to $\mathbf{q}^*$ on the right-hand side, resulting in the derivation of the adjoint operator

$$\mathcal{N}'(\mathbf{q})^* \mathbf{q}^* = \left(\begin{array}{c} -\frac{\partial u_j^*}{\partial x_j} \\ -\frac{\partial u_i^*}{\partial t} - u_j \left(\frac{\partial u_i^*}{\partial x_j} + \frac{\partial u_j^*}{\partial x_i} \right) - v \frac{\partial^2 u_i^*}{\partial x_j^2} - \frac{\partial p^*}{\partial x_i} \end{array} \right),$$

where, again, the operation $\mathcal{N}'(\mathbf{q})^* \mathbf{q}^*$ is a linear operation on the adjoint field $\mathbf{q}^*$, and the operator $\mathcal{N}'(\mathbf{q})^*$ is itself a function of the solution $\mathbf{q}$ of the Navier–Stokes problem. From the integrations by parts, we also get several boundary terms

$$b = \int_{\Omega} \left(u_j^* u_j' \right) \Big|_{t=0}^{t=T} \mathbf{d}\mathbf{x} + \int_0^T \int_{\Gamma_{\pm}^{\pm}} n_j \left[u_i^* (u_j u_i' + u_j' u_i) + p^* u_j' - v \left(u_i^* \frac{\partial u_i'}{\partial x_j} - u_i' \frac{\partial u_i^*}{\partial x_j} \right) + u_j^* p' \right] \mathbf{d}\mathbf{x} \, dt.$$

The identity (13) is the key to expressing $\mathcal{J}'$ in the desired form. An adjoint field $\mathbf{q}^*$ is first defined using the operator $\mathcal{N}'(\mathbf{q})^*$ together with appropriate forcing on an interior equation with appropriate boundary conditions and initial conditions. There is here some flexibility which we exploit to obtain a simple expression of $\mathcal{J}'$. Indeed, combining this definition of $\mathbf{q}^*$ with the definitions of $\mathbf{q}$ in (11) and $\mathbf{q}'$ in (12), the identity (13) reveals the desired expression, as will now be shown.

8.1.6. Definition of adjoint field

Consider an adjoint state defined (as yet, arbitrarily) by

$$\mathcal{N}'(\mathbf{q})^* \mathbf{q}^* = \begin{pmatrix} 0 \\ \mathcal{C}_1^* \mathcal{C}_1 \mathbf{u} \end{pmatrix} \quad \text{in } \Omega, \quad (14a)$$

$$\mathbf{u}^* = \mathcal{C}_3^* \mathbf{r} \quad \text{on } \Gamma_2^{\pm}, \quad (14b)$$

$$\mathbf{u}^* = \mathcal{C}_2^* \mathcal{C}_2 \mathbf{u} \quad \text{at } t = T, \quad (14c)$$

where the adjoint operation $\mathcal{N}'(\mathbf{q})^* \mathbf{q}^*$ is derived in the previous section. Note that, depending on where the cost functional weighs the flow perturbations (see Section 8.1.2), the adjoint problem may be driven by the initial conditions (14c), by the boundary conditions (14b), or by the RHS of the adjoint PDE (14a) itself. Note also that the “initial” conditions in (14c) are defined at $t = T$, and are thus best referred to as “terminal” conditions. With this definition, the adjoint field must be marched *backward* in time over the optimization horizon — due to the sign of the time derivative and viscous terms in the adjoint operator $\mathcal{N}'(\mathbf{q})^*$, this is the natural direction for this time march. However, as both the operator $\mathcal{N}'(\mathbf{q})^*$ and the RHS forcing on (14a) are functions of $\mathbf{q}$, computation of the adjoint field $\mathbf{q}^*$ requires storage of the flow field $\mathbf{q}$ on $t \in [0, T]$, which itself must be computed with

a *forward* march. This storage issue presents one of the numerical complications which preclude solution of the present optimization problem for large optimization intervals T . However, this storage issue is not insurmountable for intermediate values of $T^+ \lesssim O(100)$. The adjoint problem (14), though linear, has complexity similar to that of the Navier–Stokes problem (11), and may be solved with similar numerical methods.

8.1.7. Identification of gradient

The identity (13) is now simplified using the equations defining the state field (11), the perturbation field (12), and the adjoint field (14). Due to the judicious choice of the forcing terms driving the adjoint problem, the identity (13) reduces (after some manipulation) to

$$\begin{aligned} & \int_0^T \int_{\Omega} \mathcal{C}_1^* \mathcal{C}_1 \mathbf{u} \cdot \mathbf{u}' \, \mathbf{d}\mathbf{x} \, dt + \int_{\Omega} (\mathcal{C}_2^* \mathcal{C}_2 \mathbf{u} \cdot \mathbf{u}')_{t=T} \, \mathbf{d}\mathbf{x} \\ & - \int_0^T \int_{\Gamma_{\pm}^{\pm}} v \mathcal{C}_3^* \mathbf{r} \cdot \frac{\partial \mathbf{u}'}{\partial \mathbf{n}} \, \mathbf{d}\mathbf{x} \, dt = \int_0^T \int_{\Gamma_{\pm}^{\pm}} p^* \phi' \, \mathbf{d}\mathbf{x} \, dt. \end{aligned}$$

Using this equation, the cost functional perturbation $\mathcal{J}'$ may be rewritten as

$$\begin{aligned} \mathcal{J}'_o(\phi; \phi') &= \int_0^T \int_{\Gamma_{\pm}^{\pm}} (p^* + \ell^2 \phi) \phi' \, \mathbf{d}\mathbf{x} \, dt \\ &\triangleq \int_0^T \int_{\Gamma_{\pm}^{\pm}} \frac{\mathcal{D} \mathcal{J}_o(\phi)}{\mathcal{D} \phi} \phi' \, \mathbf{d}\mathbf{x} \, dt. \end{aligned}$$

As ϕ' is arbitrary, we may identify (weakly) the desired gradient as

$$\frac{\mathcal{D} \mathcal{J}_o(\phi)}{\mathcal{D} \phi} = p^* + \ell^2 \phi.$$

The desired gradient $\mathcal{D} \mathcal{J}_o(\phi) / \mathcal{D} \phi$ is thus found to be a simple function of the solution of the adjoint problem proposed in (14); specifically, in the present case of boundary forcing by wall-normal blowing and suction, it is found to be a simple function of the adjoint pressure on the walls.

In fact, this simple result hints at the more fundamental physical interpretation of what the adjoint field actually represents:

The adjoint field $\mathbf{q}^*$, when properly defined, is a measure of the sensitivity of the terms of the cost functional which appraise the state $\mathbf{q}$ to additional forcing of the state equation.

Note that there are exactly as many components of the adjoint field $\mathbf{q}^*$ as there are components of the state PDE on the interior of the domain, and that the adjoint field may take nontrivial values at the initial time $t = 0$ and on the boundaries $\Gamma_{\pm}^{\pm}$. Depending upon where the control is applied to the state equation (11), (i.e., on the RHS of the mass or momentum equations on the interior of the

domain, on the boundary conditions, or on the initial conditions), the adjoint field will appear in the resulting expression for the gradient accordingly.

To summarize: the forcing on the adjoint problem is a function of where the flow perturbations are weighed in the cost functional. The dependence of the gradient $\mathcal{D}\mathcal{J}(\phi)/\mathcal{D}\phi$ on the resulting adjoint field, on the other hand, is a function of where the control enters the state equation.

8.1.8. Gradient update to control

A control optimization strategy using a steepest descent algorithm may now be proposed such that

$$\phi^k = \phi^{k-1} - \alpha^k \frac{\mathcal{D}\mathcal{J}_o(\phi^{k-1})}{\mathcal{D}\phi}$$

over the entire time interval $t \in (0, T]$, where k indicates the iteration number and α^k is a parameter of descent which governs how large an update is made, which is adjusted at each iteration step to be that value which minimizes $\mathcal{J}$. This algorithm updates ϕ at each iteration in the direction of maximum decrease of $\mathcal{J}$. As $k \rightarrow \infty$, the algorithm should converge to some local minimum of $\mathcal{J}$ over the domain of the control ϕ on the time interval $t \in (0, T]$. Note that convergence to a global minimum will not in general be attained by such a scheme, and that, as time proceeds, $\mathcal{J}$ will not necessarily decrease.

The steepest descent algorithm described above illustrates the essence of the approach, but is usually not very efficient. Even in linear low-dimensional problems, for cases in which the cost functional has a long, narrow “valley”, the lack of a momentum term from one iteration to the next tends to cause the steepest descent algorithm to bounce from one side of the valley to the other without turning to proceed along the valley floor. Standard nonlinear conjugate gradient algorithms (see, e.g. [40]) improve this behavior considerably with relatively little added computational cost or algorithmic complexity, as discussed further in [30].

As mentioned previously, the dimension of the control in the present problem (once discretized) is quite large, which precludes the use of second-order techniques which are based on the computation or approximation of the Hessian matrix $\partial^2 \mathcal{J} / \partial \phi_i \partial \phi_j$ or its inverse during the control optimization. The number of elements in such a matrix scales with the square of the number of control variables, and is unmanageable in the present case. However, reduced-storage variants of variable metric methods [41], such as the Davidon–Fletcher–Powell (DFP) method, the Broydon–Fletcher–Goldfarb–Shanno (BFGS) method, and the sequential quadratic programming (SQP) method, approximate the inverse Hessian information by outer products of stored gradient vectors, and thus achieve nearly second-order convergence without storage of the Hessian matrix itself. Such

techniques should be explored further for very large scale optimization problems such as the present in future work.

8.2. Continuous adjoint vs. discrete adjoint

Direct numerical simulations (DNS) of the present three-dimensional nonlinear system necessitate carefully chosen numerical techniques involving a stretched, staggered grid, an energy-conserving spatial discretization, and a mixture of implicit and multi-step explicit schemes for accurate time advancement, with incompressibility enforced by an involved fractional step algorithm. The optimization approach described above, which will be referred to as “optimize then discretize” (OTD), avoids all of these cumbersome numerical details by deriving the gradient of the cost functional in the continuous setting, discretizing in time and space only as the final step before implementation in numerical code. The remarkable similarity of the flow and adjoint systems allow both to be coded with similar numerical techniques. For systems which are well resolved in the numerical discretization, this approach is entirely justifiable, and is found to yield adjoint systems which are easy to derive and implement in numerical code.

Unfortunately, many PDE systems, such as high Reynolds-number turbulent flows, are difficult or impossible to simulate with sufficient resolution to capture accurately all of the important dynamic phenomena of the continuous system. Such systems are often simulated on coarse grids, usually with some “subgrid-scale model” to account for the unresolved dynamics. This setting is referred to as large eddy simulation (LES), and a variety of techniques are currently under development to model the significant subgrid-scale effects.

There are important unresolved issues concerning how to approach large eddy simulations in the optimization framework. If we continue with the OTD approach, in which the optimization equations are determined before the numerical discretization is applied, it is not yet clear at what point the LES model should be introduced. Prof. Scott Collis’s group (Rice U.) have modified the numerical code of [30] in order to study this issue; Chang and Collis [42] report on their preliminary findings.

An alternative approach to the OTD setting, in which one spatially discretizes the governing equation before determining the optimization equations, may also be considered.⁴ After spatially discretizing the governing equation, this approach, which will be referred to as

⁴Note that we may also consider temporally discretizing the governing equation before determining the optimization equations. Generally, the spatial discretization of a turbulent system is the most restrictive issue, however, so we focus on that problem here.

“discretize then optimize” (DTO), follows an analogous sequence of steps as the OTD approach presented previously, with these steps now applied in the discrete setting. As shown in Appendix B, derivation of the adjoint operator is significantly more cumbersome in this discrete setting. In general, the processes of optimization and discretization do not commute, and thus the OTD and DTO approaches are not necessarily equivalent even upon refinement of the space/time grid [43]. However, by carefully framing the discrete identity defining the DTO adjoint operator as a discrete approximation of the identity given in (13), these two approaches can be posed in an equivalent fashion for Navier–Stokes systems, as shown in Appendix B.

It remains the topic of some debate whether or not the DTO approach is better than the OTD approach for marginally resolved PDE systems. The argument for DTO is that it clearly is the most direct way to optimize the discrete problem actually being solved by the computer. The argument against DTO is that one really wants to optimize the continuous problem, so gradient information which identifies and exploits deficiencies in the numerical discretization which can lead to performance improvements in the discrete problem might be misleading when interpreting the numerical results in terms of the physical system.

9. Robustification: appealing to Murphy’s law

Though optimal control approaches possess an attractive mathematical elegance and are now proven to provide excellent results in terms of drag and turbulent kinetic energy reduction in fully developed turbulent flows, they are often impractical. One of the most significant drawbacks of this nonlinear optimization approach is that it tends to “over-optimize” the system, leaving a high degree of design-point sensitivity. This phenomena has been encountered frequently in, for example, the adjoint-based optimization of the shape of aircraft wings. Overly optimized wing shapes might work quite well at exactly the flow conditions for which they were designed, but their performance is often abysmal at off-design conditions. In order to abate such system sensitivity, the noncooperative framework of robust control provides a natural means to “detune” the optimized results. This concept can be applied easily to a broad range of related applications.

The noncooperative approach to robust control, one might say, amounts to Murphy’s law taken seriously:

If a worst-case disturbance can disrupt a controlled closed-loop system, it will.

When designing a robust controller, therefore, one might plan on a finite component of the worst-case disturbance




Fig. 16. Schematic of a saddle point representing the neighborhood of a solution to a robust control problem with one scalar disturbance variable ψ and one scalar control variable ϕ . When the robust control problem is solved, the cost function $\mathcal{J}_r$ is simultaneously maximized with respect to ψ and minimized with respect to ϕ , and a saddle point such as $(\bar{\psi}, \bar{\phi})$ is reached. An essentially infinite-dimensional extension of this concept may be formulated to achieve robustness to disturbances and insensitivity to design point in fluid-mechanical systems. In such approaches, the cost $\mathcal{J}_r$ is related to a distributed disturbance ψ and a distributed control ϕ through the solution of the Navier–Stokes equation.

aggravating the system, and design a controller which is suited to handle even this extreme situation. A controller which is designed to work even in the presence of a finite component of the worst-case disturbance will also be robust to a wide class of other possible disturbances which, by definition, are not as detrimental to the control objective as the worst-case disturbance. This concept is exactly that which leads to the $\mathcal{H}_\infty$ control formulation discussed previously in the linear setting, and can easily be extended to the optimization of nonlinear systems.

Based on the ideas of $\mathcal{H}_\infty$ control theory presented in Section 2, the extension of the nonlinear optimization approach presented in Section 8 to the noncooperative setting is straightforward. A disturbance is first introduced to the governing equation (11); as an example, we may consider disturbances which perturb the state PDE itself such that

$$\mathcal{N}(\mathbf{q}) = \mathbf{F} + \mathbf{B}_1(\psi) \quad \text{in } \Omega.$$

Accounting for disturbances to the boundary conditions and initial conditions of the governing equation is also straightforward. The cost functional is then extended to penalize these disturbances in the noncooperative framework, as was also done in the linear setting

$$\mathcal{J}_r(\psi, \phi) = \mathcal{J}_o - \frac{\gamma^2}{2} \int_0^T \int_\Omega |\psi|^2 \, \mathbf{d}\mathbf{x} \, \mathbf{d}t.$$

This cost functional is simultaneously minimized with respect to the controls ϕ and maximized with respect to

the disturbances ψ (Fig. 16). The parameter γ is used to scale the magnitude of the disturbances accounted for in this noncooperative competition, with the limit of large γ recovering the optimal approach discussed in Section 8 (i.e., $\psi \rightarrow 0$). A gradient-based algorithm may then be devised to march to the saddle point, such as the simple algorithm given by

$$\begin{aligned}\phi^k &= \phi^{k-1} - \alpha^k \frac{\mathcal{D} \mathcal{J}_r(\psi^{k-1}; \phi^{k-1})}{\mathcal{D} \phi}, \\ \psi^k &= \psi^{k-1} + \beta^k \frac{\mathcal{D} \mathcal{J}_r(\psi^{k-1}; \phi^{k-1})}{\mathcal{D} \psi}.\end{aligned}$$

The robust control problem is considered to be solved when a saddle point $(\bar{\psi}, \bar{\phi})$ is reached; note that such a solution, if it exists, is not necessarily unique.

The gradients $\mathcal{D} \mathcal{J}_r(\psi; \phi)/\mathcal{D} \phi$ and $\mathcal{D} \mathcal{J}_r(\psi; \phi)/\mathcal{D} \psi$ may be found in a manner analogous to that leading to $\mathcal{D} \mathcal{J}_0(\phi)/\mathcal{D} \phi$ discussed in Section 8. In fact, both gradients may be extracted from the single adjoint field defined by (14). Thus, the additional computational complexity introduced by the noncooperative component of the robust control problem is simply a matter of updating and storing the appropriate disturbance variables.

9.1. Well-posedness

Based on the extensive mathematical literature on the Navier–Stokes equation, Abergel and Temam [38] established the well-posedness of the mathematical framework for the optimization problem presented in Section 8. This characterization was generalized and extended to the noncooperative framework of Section 9 in Bewley et al. [44].

Due to the fact that the inequalities currently available for estimating the magnitude of the various terms of the Navier–Stokes equation are limited, the mathematical characterizations in both of these articles are quite conservative. In our numerical simulations, we regularly apply numerical optimization techniques to control problems which are well outside the range over which we can mathematically establish well-posedness. However, such mathematical characterizations are still quite important, as they give us confidence that, for example, if ℓ and γ are at least taken to be large enough, a saddle point of the noncooperative optimization problem will exist. Once such mathematical characterizations are derived, numerically determining the values of ℓ and γ for which solutions of the control problem may still be obtained is reduced to a simple matter of implementation.

9.2. Convergence of numerical algorithm

Saddle points are typically more difficult to find than minimum points, and particular care needs to be taken to

craft efficient but stable numerical algorithms for finding them. In the approach described above, sufficiently small values of α^k and β^k must be selected in order to insure convergence. Fortunately, the same mathematical inequalities used to characterize well-posedness of the control problem can also be used to characterize convergence of proposed numerical algorithms. Such characterizations lend valuable insight when designing practical numerical algorithms. Preliminary work in the development of such saddle point algorithms is reported by Tachim Medjo [45].

10. Unification: synthesizing a general framework

The various cost functionals considered previously led to three possible sources of forcing for the adjoint




Fig. 17. Schematic of the space-time domain over which the flow field $\mathbf{q}$ is defined. The possible regions of forcing in the system defining $\mathbf{q}$ are: (1) the right-hand side of the PDE, indicated with shading, representing flow control by interior volume forcing (e.g., externally applied electromagnetic forcing by wall-mounted magnets and electrodes); (2) the boundary conditions, indicated with diagonal stripes, representing flow control by boundary forcing (e.g., wall transpiration); (3) the initial conditions, indicated with checkerboard, representing optimization of the initial state in a data assimilation framework (e.g., the weather forecasting problem).




Fig. 18. Schematic of the space-time domain over which the adjoint field $\mathbf{q}^*$ is defined. The possible regions of forcing in the system defining $\mathbf{q}^*$, corresponding exactly to the possible domains in which the cost functional $\mathcal{J}$ can depend on $\mathbf{q}$, are: (1) the right-hand side of the PDE, indicated with shading, representing regulation of an interior quantity (e.g., turbulent kinetic energy); (2) the boundary conditions, indicated with diagonal stripes, representing regulation of a boundary quantity (e.g., wall skin friction); (3) the terminal conditions, indicated with checkerboard, representing terminal control of an interior quantity (e.g., turbulent kinetic energy).

problem: the right-hand side of the PDE, the boundary conditions, and the initial conditions. Similarly, three different locations of forcing may be identified for the flow problem. As illustrated in Figs. 17 and 18 and discussed further in [44], the various regions of forcing of the flow and adjoint problems together form a general framework which can be applied to a wide variety of problems in fluid mechanics including both flow control (e.g., drag reduction, mixing enhancement, and noise control) and flow forecasting (e.g., weather prediction and storm forecasting). Related techniques, but applied to the time-averaged Navier–Stokes equation, have also been used extensively to optimize the shapes of airfoils (see, for example, [46]).

By identifying a range of problems which all fit into the same general framework, we can better understand how to extend, for example, the idea of noncooperative optimizations to a full suite of related problems in fluid mechanics. Though advanced research projects must often be highly focused and specialized in order to obtain solid results, the importance of making connections of such research to a large scope of related problems must be recognized in order to realize fully the potential impact of the techniques developed.

11. Decomposition: simulation-based system modeling

For the purpose of developing model-based feedback control strategies for turbulent flows, reduced-order nonlinear models of turbulence which are effective in the closed-loop setting are highly desired. Recent work in this direction, using proper orthogonal decompositions (POD) to obtain these reduced-order representations, is reviewed by Lumley and Blossey [47].

The POD technique uses analysis of a simulation database to develop an efficient reduced-order basis for the system dynamics represented within the database [48]. One of the challenges of this approach is that the dynamics of the system in closed loop (after the control is turned on) is often quite different than the dynamics of the open-loop (uncontrolled) system. Thus, development of simulation-based reduced-order models for turbulent flows should probably be coordinated with the design of the control algorithm itself in order to determine system models which are maximally effective in the closed-loop setting. Such coordination of simulation-based modeling and control design is largely an unsolved problem. A particularly sticky issue is the fact that, as the controls are turned on, the dynamics of the turbulent flow system are nonstationary (they evolve in time). The system eventually relaminarizes if the control is sufficiently effective. In such nonstationary problems, it is not clear which dynamics the POD should represent (that of the flow shortly after the control is turned on, that of the nearly relaminarized flow, or something in between), or if in fact,

several PODs should be used in a scheduled approach in an attempt to capture several different stages of the nonstationary relaminarization process.

Reduced-order models which are effective in the closed-loop setting need not capture the majority of the energetics of the unsteady flow. Rather, the essential feature of a system model for the purpose of control design is that the model capture the important effects of the control on the system dynamics. Future control-oriented modeling efforts might benefit by deviating from the standard POD mindset of attempting to capture the energetics of the system dynamics, instead focusing on capturing the significant effects of the control on the system in a reduced-order fashion.

12. Global stabilization: conservatively enhancing stability

Global stabilization approaches based on Lyapunov analysis of the system energetics have recently been explored for 2D channel-flow systems (in the continuous setting) by Balogh et al. [49]. In the setting considered there, localized tangential wall motions are coordinated with local measurements of skin friction via simple proportional feedback strategies. Analysis of the flow at $Re \leq 0.125$ motivate such feedback rules, indicating appropriate values of proportional feedback coefficients which enhance the L^2 stability of the flow. Though such an approach is very conservative, rigorously guaranteeing enhanced stability of the channel-flow system only at extremely low Reynolds numbers, extrapolation of the feedback strategies so determined to much higher Reynolds numbers also indicates effective enhancements of system stability, even for 3D systems up to $Re = 2000$ (A. Balogh, private communication).

An alternative approach for achieving global stabilization of a nonlinear PDE is the application of nonlinear backstepping to the discretized system equation. Bošković and Krstić [50] reports on recent efforts in this direction (applied to a thermal convection loop). Backstepping is typically an aggressive approach to stabilization. One of the primary difficulties with this approach is that proofs of convergence to a continuous, bounded function upon refinement of the grid is difficult to attain due to increasing controller complexity as the grid is refined. Significant advancements will be necessary before this approach will be practical for turbulent flow systems.

13. Adaptation: accounting for a changing environment

Adaptive control algorithms, such as least mean squares (LMS), neural networks (NN), genetic algorithms (GA), simulated annealing, extremum seeking, and the

like, play an important role in the control of fluid–mechanical systems when the number of undetermined parameters in the control problem is fairly small ($\lesssim O(10)$) and individual “function evaluations” (i.e., quantitative characterizations of the effectiveness of the control) can be performed relatively quickly. Many control problems in fluid mechanics are of this type, and are readily approachable by a wide variety of well-established adaptive control strategies. A significant advantage of such approaches over those discussed previously is that they do not require extensive analysis or coding of localized convolution kernels, adjoint fields, etc., but may instead be applied directly “out of the box” to optimize the parameters of interest in a given fluid–mechanical problem. This also poses a bit of a disadvantage, however, because the analysis required during the development of model-based control strategies can sometimes yield significant physical insight which black-box optimizations fail to provide.

To apply the adaptive approach, one needs a fast simulation code or an experimental apparatus in which the control parameters can be altered by an automated algorithm. Any of a number of established methodological strategies can then be used to search the parameter space for favorable closed-loop system behavior. Given enough function evaluations and a small enough number of control parameters, such strategies usually converge to effective control solutions. Koumoutsakos et al. [51] demonstrate this approach (computationally) to determine effective control parameters for exciting instabilities in a round jet. Rathnasingham and Breuer [52] demonstrate this approach (experimentally) for the feedforward reduction of turbulence intensities in a boundary layer.

Unfortunately, due to an effect known as “the curse of dimensionality”, as the number of control parameters to be optimized is increased, the ability of adaptive strategies to converge to effective control solutions based on function evaluations alone is diminished. For example, in a system with 1000 control parameters, it takes 1000 function evaluations to determine the gradient information available in a single adjoint computation. Thus, for problems in which the number of control variables to be optimized is large, the convergence of adaptive strategies based on function evaluations alone is generally quite poor. In such high-dimensional problems, for cases in which the control problem of interest is plagued by multiple minima, a blend of an efficient adjoint-based gradient optimization approach with GA-type management of parameter “mutations” or the simulated annealing approach of varying levels of “noise” added to the optimization process might prove to be beneficial.

Adaptive strategies are also quite valuable for recognizing and responding to changing conditions in the flow system. In the low-dimensional setting, they can be used online to update controller gains directly as the system evolves in time (for instance, as the mean speed or direc-




Fig. 19. An enticing picture: fundamental restructuring of the near-wall unsteadiness to insulate the wall from the viscous effects of the bulk flow. It has been argued [53,54] that it might be possible to maintain a series of so-called “fluid rollers” to effectively reduce the drag of a near-wall flow. Such rollers are depicted in the cartoon above by indicating total velocity vectors in a reference frame convecting with the vortices themselves; in this frame, the generic picture of fluid rollers is similar to a series of stationary Kelvin–Stuart cat’s eye vortices. A possible mechanism for drag reduction might be akin to a series of solid cylinders serving as an effective conveyor belt, with the bulk flow moving to the right above the vortices and the wall moving to the left below the vortices. It is still the topic of some debate whether or not a continuous flow can be maintained in such a configuration by an unsteady control in such a way as to sustain the mean skin friction below laminar levels. Such a control might be implemented either by interior electromagnetic forcing (applied with wall-mounted magnets and electrodes) or by boundary controls such as zero-net mass flux blowing/suction.

tion of the flow changes or as the sensitivity of a sensor degrades). In the high-dimensional setting, adaptive strategies can be used to identify certain critical aspects of the flow (such as the flow speed), and, based on this identification, an appropriate control strategy may be selected from a look-up table of previously computed controller gains.

The selection of what level of adaptation is appropriate for a particular flow control problem of interest is, again, a consideration that must be guided by physical insight of the particular problem at hand.

14. Performance limitation: identifying ideal control targets

Another important, but as yet largely unrealized, role for mathematical analysis in the field of flow control is in the identification of fundamental limitations on the performance that can be achieved in certain flow control problems. For example, motivated by the active debate surrounding the proposed physical mechanism for channel-flow drag reduction illustrated in Fig. 19, we formally state the following, as yet unproven, conjecture:

Conjecture. *The lowest sustainable drag of an incompressible constant mass-flux channel flow, in either 2D or 3D, when controlled via a distribution of zero-net mass-flux blowing/suction over the channel walls, is exactly that of the laminar flow.*

Note that, by “sustainable drag”, we mean the long-time average of the instantaneous drag, given by

$$D_{\infty} = \lim_{T \rightarrow \infty} \frac{-1}{T} \int_0^T \int_{\Gamma_{\pm}^{\pm}} v \frac{\partial u_1}{\partial n} dx dt.$$

Proof (by mathematical analysis) or disproof (by counterexample) of this conjecture would be quite significant and lead to greatly improved physical understanding of the channel flow problem. If proven to be correct, it would provide rigorous motivation for targeting flow re-laminarization when the problem one actually seeks to solve is minimization of drag. If shown to be incorrect, our target trajectories for future flow control strategies might be substantially altered.

Similar fundamental performance limitations may also be sought for exterior flow problems, such as the minimum drag of a circular cylinder subject to a class of zero-net control actions, such as rotation or transverse oscillation (B. Protas, private communication).

15. Implementation: evaluating engineering tradeoffs

We are still some years away from applying the distributed control techniques discussed herein to micro-electro-mechanical-systems (MEMS) arrays of sensors and actuators, such as that depicted in Fig. 20. One of the primary hurdles left to be tackled in order to bring us closer to actual implementation is that of accounting for practical designs of sensors and actuators in the control formulations, rather than the idealized distributions of blowing/suction and skin-friction measurements which we have assumed here. Detailed simulations, such as that shown in Fig. 21, of proposed actuator designs will be essential for developing reduced-order models of the effects of the actuators on the system of interest in order to




Fig. 20. A MEMS tile integrating sensors, actuators and control logic for distributed flow control applications, developed by Profs. Chih-Ming Ho (UCLA) and Yu-Chong Tai (Caltech).

make control design for realistic arrays of sensors and actuators tractable.

By performing analysis and control design in a high-dimensional, unconstrained setting, as discussed in the present article, it is believed that we can obtain substantial insight into the physical characteristics of highly effective control strategies. Such insight naturally guides the engineering tradeoffs that follow in order to make the design of the turbulence control system practical. Particular traits of the present control solutions in which we are especially interested include the times scales and the streamwise and spanwise length scales which are dominant in the optimized control computations (which shed insight on suitable actuator bandwidth, dimensions, and spacing) and the extent and structure of the convolution




Fig. 21. Simulation of a proposed driven-cavity actuator design by Prof. Rajat Mittal (U. of Florida). The fluid-filled cavity is driven by vertical motions of the membrane along its lower wall. Numerical simulation and reduced-order modeling of the influence of such flow-control actuators on the system of interest will be essential for the development of feedback control algorithms to coordinate arrays of realistic sensor/actuator configurations.

kernels (which indicate the distance and direction over which sensor measurements and state estimates should propagate when designing the communication architecture of the tiled array).

It is recognized that the control algorithm finally to be implemented must be kept fairly simple for its realization in the on-board electronics to be feasible. We believe that an appropriate strategy for determining implementable feedback algorithms which are both effective and simple is to learn how to solve the high-dimensional, fully resolved control problem first, as discussed herein. This results in high-dimensional compensator designs which are highly effective in the closed-loop setting. Compensator reduction strategies combined with engineering judgment may then be used to distill the essential features of such well-resolved control solutions to implementable feedback designs with minimal degradation of the closed-loop system behavior.

16. Discussion: a common language for dialog

It is imperative that an accessible language be developed which provides a common ground upon which people from the fields of fluid mechanics, mathematics, and controls can meet, communicate, and develop new theories and techniques for flow control. Pierre-Simon de Laplace (quoted by Rose [55]) once said

Such is the advantage of a well constructed language that its simplified notation often becomes the source of profound theories.

Similarly, it was recognized by Gottfried Wilhelm Leibniz (quoted by Simmons [56]) that

In symbols one observes an advantage in discovery which is greatest when they express the exact nature of a thing briefly ... then indeed the labor of thought is wonderfully diminished.

Profound new theories are still possible in this young field. To a large extent, however, we have not yet homed in on a common language in which such profound theories can be framed. Such a language needs to be actively pursued; time spent on identifying, implementing, and explaining a clear “compromise” language which is approachable by those from the related “traditional” disciplines is time well spent.

In particular, care should be taken to respect the meaning of certain “loaded” words which imply specific techniques, qualities, or phenomena in some disciplines, but only general notions in others. When both writing and reading papers on flow control, one must be especially alert, as these words are sometimes used outside of their more narrow, specialized definitions, creating undue confusion. With time, a common language will

develop. In the meantime, avoiding the use of such words outside of their specialized definitions, precisely defining such words when they are used, and identifying and using the existing names for specialized techniques already well established in some disciplines when introducing such techniques into other disciplines, will go a long way towards keeping us focused and in sync as an extended research community.

There are, of course, some significant obstacles to the implementation of a common language. For example, fluid mechanics have historically used $\mathbf{u}$ to denote flow velocities and $\mathbf{x}$ to denote spatial coordinates, whereas the controls community overwhelmingly adopts $\mathbf{x}$ as the state vector and $\mathbf{u}$ as the control. The simplified 2D systems fluid mechanics often study examines the flow in a vertical plane, whereas the simplified 2D systems meteorologists often study examines the flow in a horizontal plane; thus, when studying 3D problems such as turbulence, those with a background in fluid mechanics usually introduce their third coordinate, z , in a horizontal direction, whereas those with a background in meteorology normally have “their zed in the clouds”. Writing papers in a manner which is conscious to such different backgrounds and notations, elucidating, motivating, and distilling the suitable control strategies, the relevant flow physics, the useful mathematical inequalities, and the appropriate numerical methods to a general audience of specialists from other fields, is certainly extra work. However, such efforts are necessary to make flow control research accessible to the broad audience of scientists, mathematicians, and engineers whose talents will be instrumental in advancing this field in the years to come.

17. The future: a Renaissance

The field of flow control is now poised for explosive growth and exciting new discoveries. The relative maturity of the traditional scientific disciplines contributing to this field provides us with key elements which future efforts in this field may leverage. The work described herein represents only our first, preliminary steps towards laying an integrated, interdisciplinary footing upon which future efforts in this field may be based. Many technologically significant and fundamentally important problems lie before us, awaiting analysis and new understanding in this setting. With each of these new applications come significant new questions about how best to integrate the constituent disciplines. The answers to these difficult questions will only come about through a broad knowledge of what these disciplines have to offer and how they can best be used in concert. A few problems which might be studied in the near future in the present interdisciplinary framework are highlighted in Fig. 22.




Fig. 22. Future interdisciplinary problems in flow control amenable to adjoint-based analysis: (a) minimization of sound radiating from a turbulent jet (simulation by Prof. Jon Freund, UCLA), (b) maximization of mixing in interacting cross-flow jets (simulation by Dr. Peter Blossey, UCSD), (c) optimization of surface compliance properties to minimize turbulent skin friction and (d) accurate forecasting of inclement weather systems.

Unfortunately, there are particular difficulties in pursuing truly interdisciplinary investigations of fundamental problems in flow control in our current society, as it is impossible to conduct such investigations from the perspective of any particular traditional discipline alone. Though the language of interdisciplinary research is in vogue, many university departments, funding agencies, technical journals, and, therefore, college professors fall back on the pervasive tendency of the 20-century scientist to categorize and isolate difficult scientific questions, often to the exclusion of addressing the fundamentally interdisciplinary issues. The proliferation and advancement of science in the twentieth century was, in fact, largely due to such an approach; by isolating specific and difficult problems with single-minded focus into narrowly defined scientific disciplines, great advances could once be achieved. To a large extent, however, the opportunities which were once possible with such a narrow focus have stagnated in many fields, though we are left with the scientific infrastructure in which that approach once flourished. To advance, we must courageously lead our research groups outside of the various neatly defined scientific domains into which this infrastructure injects us, and pursue the significant new opportunities appearing at their intersection. University departments and technical journals can and will follow suit as increasingly successful interdisciplinary efforts, such as those in the field of flow control, gain momentum. The endorsement which professional societies, technical journals, and funding agencies might bring to such interdisciplinary efforts holds the potential to significantly accelerate this reformation of the scientific infrastructure.

In order to promote interdisciplinary work in the scientific community at large, describing oneself as working at the intersection of disciplines X and Y (or, where they are still disjoint, the bridge between such disciplines) needs to become more commonplace. People often resort to the philosophy “I do X ... oh, and I also sometimes dabble a bit with Y ”, as the philosophy “I do $X*Y$ ”, where $*$ denotes something of the nature of an integral convolution, has not been in favor since the Renaissance. Perhaps the primary reason for this is that X and Y (and Z , W , etc.) have gotten progressively more and more difficult. By specialization (though often to the point of isolation), we are able to “master” our more and more narrowly defined disciplines. In the experience of the author, not only is it often the case that X and Y are not immiscible, but the solution sought may often not be formulated with the ingredients of X or Y alone. In order to advance, the essential ingredients of X and Y must be crystallized and communicated across the artificial disciplinary boundaries. New research must then be conducted at the intersection of X and Y . To be successful in the years to come, we must prepare ourselves and our students with the training, perspective, and resolve to seize the new opportunities appearing at such intersections with a Renaissance approach.

Acknowledgements

The author is thankful to Prof. Robert Skelton, for his vision to promote integrated controls research at the intersection of disciplines, to Profs. Roger Temam and Mohammed Ziane, for their patient attention to the new mathematical challenges laying the foundation for this development, and to Profs. Dan Henningson, Patrick Huerre, John Kim, and Parviz Moin, for their physical insight and continued support of this sometimes unconventional effort within the fluid mechanics community. The author has also benefited from numerous technical discussions with Profs. Jeff Baggett, Bassam Bamieh, Bob Bitmead, Scott Collis, Brian Farrell, Jon Freund, Petros Ioannou, and Miroslav Krstić, and remains indebted to the several graduate students and post-docs who are making it all happen, including Dr. Peter Blossey, Markus Högberg, Eric Lauga, and Scott Miller. We also thank the AFOSR, NSF, and DARPA for the foresight to sponsor several of these investigations.

This article is dedicated to the memory of Prof. Mohammed Dahleh, whose charisma and intellect endure as a continual inspiration to all who knew him.

Appendix A. Alternative scaling of $\mathcal{H}_\infty$ problem

As mentioned previously, it is useful to pose spatially discretized $\mathcal{H}_\infty$ control problems based on PDE systems in such a way that the feedback gains converge to continuous functions upon refinement of the spatial discretization. In order to accomplish this, consider the discretized plant

$$\dot{\mathbf{x}} = A\mathbf{x} + G_1\mathbf{w}_1 + B\mathbf{u},$$

$$\mathbf{y} = C\mathbf{x} + \alpha G_2\mathbf{w}_2.$$

As before, $C\mathbf{x}$ denotes the raw vector of measured variables, and G_1 and αG_2 represent the square root of any known or expected covariance structure of the state disturbances and measurement noise, respectively. The cost function considered now is

$$\mathcal{J} \triangleq E[\mathbf{x}^*Q\mathbf{x} + \ell^2\mathbf{u}^*R\mathbf{u} - \gamma^2\mathbf{w}^*S\mathbf{w}],$$

$$\mathbf{w} = \begin{pmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \end{pmatrix}, \quad S = \begin{pmatrix} S_1 & 0 \\ 0 & S_2 \end{pmatrix},$$

for R , S_1 , and S_2 which are positive definite and nonsingular. Appropriate definition of the matrices R and S allow for a problem definition that will converge upon refinement of the space-time grid [19]. Assume

a compensator of the form

$$\dot{\hat{\mathbf{x}}} = A\hat{\mathbf{x}} + G_1\hat{\mathbf{w}}_1 + B\mathbf{u} - \mathbf{v},$$

$$\hat{\mathbf{y}} = C\hat{\mathbf{x}} + \alpha G_2\hat{\mathbf{w}}_2,$$

$$\mathbf{v} = L(\mathbf{y} - \hat{\mathbf{y}}),$$

$$\mathbf{u} = K\hat{\mathbf{x}}.$$

The $\mathcal{H}_\infty$ control solution minimizing $\mathcal{J}$ with respect to $\mathbf{u}$ while maximizing $\mathcal{J}$ with respect to $\mathbf{w}$, which exists only for stabilizable and detectable systems for sufficiently large γ , is given by

$$K = -\frac{1}{\ell^2}R^{-1}B^*X, \quad \hat{\mathbf{w}}_1 = \frac{1}{\gamma^2}S_1^{-1}G_1^*X\hat{\mathbf{x}},$$

$$L = -\frac{1}{\alpha^2}ZYC^*G_2^{*-1}S_2G_2^{-1}, \quad \hat{\mathbf{w}}_2 = 0,$$

where

$$X = \text{Ric} \begin{pmatrix} A & \frac{1}{\gamma^2}G_1S_1^{-1}G_1^* - \frac{1}{\ell^2}BR^{-1}B^* \\ -Q & -A^* \end{pmatrix},$$

$$Y = \text{Ric} \begin{pmatrix} A^* & \frac{1}{\gamma^2}Q - \frac{1}{\alpha^2}C^*G_2^{*-1}S_2G_2^{-1}C \\ -G_1S_1^{-1}G_1^* & -A \end{pmatrix},$$

$$Z = \left(I - \frac{YX}{\gamma^2} \right)^{-1}.$$

This solution may be obtained by the change of variables $\bar{\mathbf{u}} = R^{1/2}\mathbf{u}$, $\bar{\mathbf{w}} = S^{1/2}\mathbf{w}$, and $\bar{\mathbf{y}} = S_2^{1/2}G_2^{-1}\mathbf{y}$, which reduces the above control problem to the standard form of the $\mathcal{H}_\infty$ control problem presented in Section 2, then changing the variables back to the present form once the solution is obtained.

Appendix B. A discrete adjoint for the 2D incompressible Navier–Stokes equation

In order to solve for the adjoint of a poorly resolved flow field, there is a significant numerical advantage to computing the adjoint of the spatially-discretized state equation (an approach we will refer to as discretize then optimize, or DTO) rather than attempting to spatially discretize the continuous expression for the adjoint field (an approach we will refer to as optimize then discretize, or OTD). In low-dimensional discretizations of PDE systems, the DTO discrete adjoint equation may be determined by simply taking the conjugate transpose of the (sparse) matrix representing the spatially discretized, linear equation governing small perturbations to the flow. However, for high-dimensional discretizations of PDE systems, it is not efficient (or indeed even tractable) to express this sparse matrix directly. Instead, it is




Fig. 23. A uniform 2D staggered grid. The pressure p is stored at the locations indicated by the dots, and the horizontal and vertical velocities, u and v , are stored at the locations indicated by the horizontal and vertical arrows, respectively. The continuity equation is evaluated at the p locations, and the horizontal and vertical momentum equations are evaluated at the u and v locations, respectively.

desirable to derive the discrete adjoint equation by hand. Recent evidence indicates that such an exercise, though tedious, may be beneficial for the optimization of poorly resolved PDE systems. In order to illustrate the approach, we derive here the discrete adjoint of the incompressible Navier–Stokes equation in conservation form for the case in which this equation is discretized with a second-order method on the uniform staggered 2D grid shown in Fig. 23. Extension of this approach to noncartesian geometries might prove to be difficult, though extensions of the approach to 3D systems, large eddy simulations with subgrid-scale models, stretched grids, and mixed finite-difference/Fourier discretizations should be straightforward.

For clarity, we now adopt the notation that $\{u, v\}$ denote the velocity components, $\{x, y\}$ denote the spatial directions, and the subscripts $\{i, j\}$ denote the gridpoint location. Defining

$$\eta_j \triangleq \begin{cases} \frac{1}{2}, & j = 0, NY \\ 1 & \text{otherwise,} \end{cases}$$

and denoting by N_e , N_h and N_v the components of the spatially-discretized nonlinear 2D Navier–Stokes equation corresponding to the continuity equation, the horizontal momentum equation, and the vertical momentum equation, respectively, it follows that:

$$(N_e)_{i,j} = \left(\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} \right)_{i,j},$$

$$(N_h)_{i+1/2,j} = \left(\frac{\partial u}{\partial t} + \frac{\partial u^2}{\partial x} + \frac{\partial vu}{\partial y} - v\Delta u + \frac{\partial p}{\partial x} \right)_{i+1/2,j},$$

$$(N_v)_{i,j+1/2} = \left(\frac{\partial v}{\partial t} + \frac{\partial uv}{\partial x} + \frac{\partial v^2}{\partial y} - v\Delta v + \frac{\partial p}{\partial y} \right)_{i,j+1/2},$$

where the terms involving discretized spatial derivatives are given by

$$\begin{aligned}
 \left(\frac{\delta u}{\delta x} + \frac{\delta v}{\delta y} \right)_{i,j} &= \frac{1}{\Delta x} [u_{i+1/2,j} - u_{i-1/2,j}] \\
 &\quad + \frac{1}{\Delta y} [v_{i,j+1/2} - v_{i,j-1/2}], \\
 \left(\frac{\delta u^2}{\delta x} \right)_{i+1/2,j} &= \frac{1}{\Delta x} \left[\left(\frac{u_{i+3/2,j} + u_{i+1/2,j}}{2} \right)^2 \right. \\
 &\quad \left. - \left(\frac{u_{i+1/2,j} + u_{i-1/2,j}}{2} \right)^2 \right], \\
 \left(\frac{\delta vu}{\delta y} \right)_{i+1/2,j} &= \frac{1}{\Delta y} \left[\frac{u_{i+1/2,j+1} + u_{i+1/2,j}}{2} \frac{v_{i+1,j+1/2} + v_{i,j+1/2}}{2} \right. \\
 &\quad \left. - \frac{u_{i+1/2,j} + u_{i+1/2,j-1}}{2} \frac{v_{i+1,j-1/2} + v_{i,j-1/2}}{2} \right], \\
 (\Delta u)_{i+1/2,j} &= \frac{1}{(\Delta x)^2} [u_{i+3/2,j} - 2u_{i+1/2,j} + u_{i-1/2,j}] \\
 &\quad + \frac{1}{(\Delta y)^2} [u_{i+1/2,j+1} - 2u_{i+1/2,j} + u_{i+1/2,j-1}], \\
 \left(\frac{\delta p}{\delta x} \right)_{i+1/2,j} &= \frac{1}{\Delta x} [p_{i+1,j} - p_{i,j}], \\
 \left(\frac{\delta uv}{\delta x} \right)_{i,j+1/2} &= \frac{1}{\Delta x} \left[\frac{u_{i+1/2,j+1} + u_{i+1/2,j}}{2} \frac{v_{i+1,j+1/2} + v_{i,j+1/2}}{2} \right. \\
 &\quad \left. - \frac{u_{i-1/2,j+1} + u_{i-1/2,j}}{2} \frac{v_{i+1,j+1/2} + v_{i-1,j+1/2}}{2} \right], \\
 \left(\frac{\delta v^2}{\delta y} \right)_{i,j+1/2} &= \frac{1}{\eta_j \Delta y} \left[\left(\frac{v_{i,j+3/2} + v_{i,j+1/2}}{2} \right)^2 \right. \\
 &\quad \left. - \left(\frac{v_{i,j+1/2} + v_{i,j-1/2}}{2} \right)^2 \right], \\
 (\Delta v)_{i,j+1/2} &= \frac{1}{(\Delta x)^2} [v_{i+1,j+1/2} - 2v_{i,j+1/2} + v_{i-1,j+1/2}] \\
 &\quad + \frac{1}{\eta_j (\Delta y)^2} [v_{i,j+3/2} - 2v_{i,j+1/2} + v_{i,j-1/2}], \\
 \left(\frac{\delta p}{\delta y} \right)_{i,j+1/2} &= \frac{1}{\Delta x} [p_{i,j+1} - p_{i,j}].
 \end{aligned}$$

Near the walls, located at $j = \frac{1}{2}$ and $j = NY + \frac{1}{2}$, the expressions for u and v must be treated carefully to

account for the boundary conditions properly. In the present work, we assume that the boundary values $v_{*,1/2}$ and $v_{*,NY+1/2}$ are prescribed as the control, that the boundary values of u are equal to zero, and that the boundary values of p are equal to their value at the first interior gridpoint. For convenience, define the following (nonphysical) variables outside of the physical domain

$$\begin{aligned}
 p_{*,0} &\triangleq p_{*,1}, \quad p_{*,NY+1} \triangleq p_{*,NY}, \quad u_{*,0} \triangleq -u_{*,1}, \quad u_{*,NY+1} \\
 &\triangleq -u_{*,NY}, \quad v_{*, -1/2} \triangleq v_{*,1/2}, \quad v_{*,NY+3/2} \triangleq v_{*,NY+1/2}.
 \end{aligned}$$

Periodic boundary conditions in x are assumed, so that $u_{NX+1/2,*} \triangleq u_{1/2,*}$, etc. With these definitions, the above expressions for the discretized spatial derivatives are applied at all grid points for which they are defined on the staggered grid, including the p and u gridpoints adjacent to the walls (at $j = 1$ and $j = NY$) and the v gridpoints on the walls (at $j = \frac{1}{2}$ and $j = NY + \frac{1}{2}$).

The perturbation of the convective terms (indicated with a prime) that arise when the control is perturbed a small amount (ϕ') follow immediately from the formulae given previously. For example,

$$(N_h)'_{i+1/2,j} = \left(\frac{\partial u}{\partial t} + \frac{\delta u^2}{\delta x} + \frac{\delta vu}{\delta y} - v \Delta u + \frac{\delta p}{\delta x} \right)'_{i+1/2,j}$$

where

$$\begin{aligned}
 \left(\frac{\delta vu}{\delta y} \right)'_{i+1/2,j} &= \frac{1}{\Delta y} \left[\frac{u'_{i+1/2,j+1} + u'_{i+1/2,j}}{2} \frac{v_{i+1,j+1/2} + v_{i,j+1/2}}{2} \right. \\
 &\quad + \frac{u_{i+1/2,j+1} + u_{i+1/2,j}}{2} \frac{v'_{i+1,j+1/2} + v'_{i,j+1/2}}{2} \\
 &\quad - \frac{u'_{i+1/2,j} + u'_{i+1/2,j-1}}{2} \frac{v_{i+1,j-1/2} + v_{i,j-1/2}}{2} \\
 &\quad \left. - \frac{u_{i+1/2,j} + u_{i+1/2,j-1}}{2} \frac{v'_{i+1,j-1/2} + v'_{i,j-1/2}}{2} \right],
 \end{aligned}$$

etc. The several other terms of the discrete perturbation equation follow by inspection. In order to derive the adjoint, we take the inner product of all of the terms of the perturbation equation with the associated adjoint variables $p_{i,j}^*$, $u_{i+1/2,j}^*$, and $v_{i,j+1/2}^*$, perform a discrete integration over the domain of interest, and rearrange, grouping separately all terms which multiply $p_{i,j}'$, $u'_{i+1/2,j}$ and $v'_{i,j+1/2}$. This leads to an identity that may be written in the symbolic form

$$\begin{aligned}
 \int_0^T \sum_{i=1}^{NX} \left\{ \sum_{j=1}^{NY} [p_{i,j}^* (N_c)'_{i,j} + u_{i+1/2,j}^* (N_h)'_{i+1/2,j}] \Delta y \right. \\
 \left. + \sum_{j=0}^{NY} [v_{i,j+1/2}^* (N_v)'_{i,j+1/2}] \eta_j \Delta y \right\} \Delta x dt
 \end{aligned}$$

$$= \int_0^T \sum_{i=1}^{NX} \left\{ \sum_{j=1}^{NY} [(N_c)^*_{i,j} p'_{i,j} + (N_h)^*_{i+1/2,j} u'_{i+1/2,j}] \Delta y \right. \\ \left. + \sum_{j=0}^{NY} [(N_v)^*_{i,j+1/2} v'_{i,j+1/2}] \eta_j \Delta y \right\} \Delta x \, dt + b.$$

The tedious rearrangement of the summations necessary to re-express the LHS of this equation in the form of the RHS comprise what is effectively “discrete integration by parts”. Carrying through all of the terms, it is eventually found that the three parts of the adjoint operator may be written in the suggestive form

$$(N_c)^*_{i,j} = \left(-\frac{\delta u^*}{\delta x} - \frac{\delta v^*}{\delta y} \right)_{i,j},$$

$$(N_h)^*_{i+1/2,j} = \left(-\frac{\partial u^*}{\partial t} - 2u \frac{\delta u^*}{\delta x} - v \frac{\delta u^*}{\delta y} \right. \\ \left. - v \frac{\delta v^*}{\delta x} - v \Delta u^* - \frac{\delta p^*}{\delta x} \right)_{i+1/2,j},$$

$$(N_v)^*_{i,j+1/2} = \left(-\frac{\partial v^*}{\partial t} - u \frac{\delta v^*}{\delta x} - u \frac{\delta u^*}{\delta y} \right. \\ \left. - 2v \frac{\delta v^*}{\delta y} - v \Delta v^* - \frac{\delta p^*}{\delta y} \right)_{i,j+1/2},$$

where the terms of these expressions which are found to be discrete approximations of spatial derivatives are given by

$$\left(\frac{\delta u^*}{\delta x} + \frac{\delta v^*}{\delta y} \right)_{i,j} = \frac{1}{\Delta x} [u^*_{i+1/2,j} - u^*_{i-1/2,j}] \\ + \frac{1}{\Delta y} [v^*_{i,j+1/2} - v^*_{i,j-1/2}],$$

$$\left(2u \frac{\delta u^*}{\delta x} \right)_{i+1/2,j} = \left(\frac{u_{i+3/2,j} + u_{i+1/2,j}}{2} \frac{u^*_{i+3/2,j} - u^*_{i+1/2,j}}{\Delta x} \right. \\ \left. + \frac{u_{i+1/2,j} + u_{i-1/2,j}}{2} \frac{u^*_{i+1/2,j} - u^*_{i-1/2,j}}{\Delta x} \right),$$

$$\left(v \frac{\delta u^*}{\delta y} \right)_{i+1/2,j} = \frac{1}{2} \left(\frac{v_{i+1,j+1/2} + v_{i,j+1/2}}{2} \frac{u^*_{i+1/2,j+1} - u^*_{i+1/2,j}}{\Delta y} \right. \\ \left. + \frac{v_{i+1,j-1/2} + v_{i,j-1/2}}{2} \frac{u^*_{i+1/2,j} - u^*_{i+1/2,j-1}}{\Delta y} \right),$$

$$\left(v \frac{\delta v^*}{\delta x} \right)_{i+1/2,j} = \frac{1}{2} \left(\frac{v_{i+1,j+1/2} + v_{i,j+1/2}}{2} \frac{v^*_{i+1,j+1/2} - v^*_{i,j+1/2}}{\Delta x} \right. \\ \left. + \frac{v_{i+1,j-1/2} + v_{i,j-1/2}}{2} \frac{v^*_{i+1,j-1/2} - v^*_{i,j-1/2}}{\Delta x} \right),$$

$$(\Delta u^*)_{i+1/2,j} = \frac{1}{(\Delta x)^2} [u^*_{i+3/2,j} - 2u^*_{i+1/2,j} + u^*_{i-1/2,j}] \\ + \frac{1}{(\Delta y)^2} [u^*_{i+1/2,j+1} - 2u^*_{i+1/2,j} + u^*_{i+1/2,j-1}],$$

$$\left(\frac{\delta p^*}{\delta x} \right)_{i+1/2,j} = \frac{1}{\Delta x} [p^*_{i+1,j} - p^*_{i,j}],$$

$$\left(u \frac{\delta v^*}{\delta x} \right)_{i,j+1/2} = \frac{1}{2} \left(\frac{u_{i+1/2,j+1} + u_{i+1/2,j}}{2} \frac{v^*_{i+1,j+1/2} - v^*_{i,j+1/2}}{\Delta x} \right. \\ \left. + \frac{u_{i-1/2,j+1} + u_{i-1/2,j}}{2} \frac{v^*_{i,j+1/2} - v^*_{i-1,j+1/2}}{\Delta x} \right),$$

$$\left(u \frac{\delta u^*}{\delta y} \right)_{i,j+1/2} = \frac{1}{2} \left(\frac{u_{i+1/2,j+1} + u_{i+1/2,j}}{2} \frac{u^*_{i+1/2,j+1} - u^*_{i+1/2,j}}{\Delta y} \right. \\ \left. + \frac{u_{i-1/2,j+1} + u_{i-1/2,j}}{2} \frac{u^*_{i-1/2,j+1} - u^*_{i-1/2,j}}{\Delta y} \right),$$

$$\left(2v \frac{\delta v^*}{\delta y} \right)_{i,j+1/2} = \frac{1}{\eta_j} \left(\frac{v_{i,j+3/2} + v_{i,j+1/2}}{2} \frac{v^*_{i,j+3/2} - v^*_{i,j+1/2}}{\Delta y} \right. \\ \left. + \frac{v_{i,j+1/2} + v_{i,j-1/2}}{2} \frac{v^*_{i,j+1/2} - v^*_{i,j-1/2}}{\Delta y} \right),$$

$$(\Delta v^*)_{i,j+1/2} = \frac{1}{(\Delta x)^2} [v^*_{i+1,j+1/2} - 2v^*_{i,j+1/2} + v^*_{i-1,j+1/2}] \\ + \frac{1}{\eta_j (\Delta y)^2} [v^*_{i,j+3/2} - 2v^*_{i,j+1/2} + v^*_{i,j-1/2}],$$

$$\left(\frac{\delta p^*}{\delta y} \right)_{i,j+1/2} = \frac{1}{\Delta x} [p^*_{i,j+1} - p^*_{i,j}].$$

Note that, when properly reorganized, the various terms of the discrete adjoint equation are particular numerical approximations of the continuous adjoint terms derived in Section 8, though not necessarily the numerical approximation of these terms which one might first think to implement in code if one derived only the continuous adjoint form. Similarly, noting the (nonphysical) variable definitions outside of the physical domain which were enumerated previously, the boundary terms associated with these “discrete integrations by parts” are simply a discrete approximation of the corresponding expression for b given in Section 8; the derivation of this discrete approximation of b is left as an exercise for the student.

When the flow system one is trying to optimize is well resolved on the numerical grid chosen, both the OTD and DTO approaches give approximately the same result for the estimation of the gradient. However, deriving the adjoint operator in this discrete fashion results in adjoint field computations which still reveal reliable gradient information for the discrete optimization problem even when the grid is too coarse for the continuous PDEs to be accurately represented by the discretized equations.

References

- [1] Ho C-M, Tai Y-C. Review: MEMS and its applications for flow control. *ASME J Fluid Engng* 1996;118:437–47.
- [2] Ho C-M, Tai Y-C. Micro-electro-mechanical systems (MEMS) and fluid flows. *Annu Rev Fluid Mech* 1998;30:579–612.
- [3] McMichael JM. Progress and prospects for active flow control using microfabricated electro-mechanical systems. *AIAA* 96-0306, 1996.
- [4] Gad-El-Hak M. Modern developments in flow control. *Appl Mech Rev* 1996;49(7):365–79.
- [5] Löfdahl L, Gad-El-Hak M. MEMS applications in turbulence and flow control. *Prog Aerosp Sci* 1999;35:101–203.
- [6] Banks HT, editor. Control and estimation in distributed parameter systems. *Frontiers in applied mathematics*, vol. 11. Philadelphia, PA: SIAM, 1992.
- [7] Banks HT, Fabiano RH, Ito K, editors. Identification and control in systems governed by partial differential equations. *Proceedings in Applied Mathematics*, vol. 68. Philadelphia, PA: SIAM, 1993.
- [8] Gunzburger MD, editor. Flow control. *The Institute for Mathematics and its Applications*, vol. 68. Berlin: Springer, 1995.
- [9] Lagnese JE, Russell DL, White LW, editors. Control and optimal design of distributed parameter systems. *The Institute for Mathematics and its Applications*, vol. 70. Berlin: Springer, 1995.
- [10] Sritharan SS. Optimal control of viscous flows. Philadelphia, PA: SIAM, 1998.
- [11] Bewley TR, Agarwal R. Optimal and robust control of transition. Center for Turbulence Research, *Proceedings of the Summer Program*, 1996. p. 405–32.
- [12] Bewley TR, Liu S. Optimal and robust control and estimation of linear paths to transition. *J Fluid Mech* 1998;365:305–49.
- [13] Joshi SS, Speyer JL, Kim J. A systems theory approach to the feedback stabilization of infinitesimal and finite-amplitude disturbances in plane Poiseuille flow. *J Fluid Mech* 1997;332:157–84.
- [14] Cortezzi L, Speyer JL. Robust reduced-order controller of laminar boundary layer transitions. *Phys Rev E* 1998;58:1906–10.
- [15] Doyle JC, Glover K, Khargonekar PP, Francis BA. State-space solutions to standard $\mathcal{H}_2$ and $\mathcal{H}_\infty$ control problems. *IEEE Trans Automat Control* 1989;34(8):831–47.
- [16] Green M, Limebeer DJN. Linear robust control. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [17] Zhou K, Doyle JC, Glover K. Robust and optimal control. Englewood Cliffs, NJ: Prentice-Hall, 1996.
- [18] Zhou K, Doyle JC. Essentials of robust control. Englewood Cliffs, NJ: Prentice-Hall, 1998.
- [19] Lauga E, Bewley TR. Robust control of linear global instability in models of non-parallel shear flows. 2000, in preparation.
- [20] Laub AJ. Invariant subspace methods for the numerical solution of Riccati equations. In: Bittaini, Laub, Willems, editors. *The Riccati equation*. Berlin: Springer, 1991, p. 163–96.
- [21] Skogestad S, Postlethwaite I. Multivariable feedback control. New York: Wiley, 1996.
- [22] Stein G, Athans M. The LQG/LTR procedure for multivariable feedback control design. *IEEE Trans Automat Control* 1987;32(2):105.
- [23] Butler KM, Farrell BF. Three-dimensional optimal perturbations in viscous shear flows. *Phys Fluids A* 1992;4(8):1637–50.
- [24] Trefethen LN, Trefethen AE, Reddy SC, Driscoll TA. Hydrodynamic stability without eigenvalues. *Science* 1993;261(30):578.
- [25] Bamieh B, Paganini F, Dahleh M. Distributed control of spatially-invariant systems. *IEEE Trans Automat Control* 2000, submitted for publication.
- [26] D’Andrea R, Dullerud GE. Distributed control of Spatially Interconnected Systems. *IEEE Trans Automat Control* 2000, submitted for publication.
- [27] Höglberg M, Bewley TR. Spatially-compact convolution kernels for decentralized control and estimation of transition in plane channel flow. *Automatica* 2000, submitted for publication.
- [28] Obinata G, Anderson BDO. Model reduction for control system design. 2000, in press.
- [29] Reddy SC, Schmid PJ, Baggett JS, Henningson DS. On stability of streamwise streaks and transition thresholds in plane channel flows. *J Fluid Mech* 1998;365:269–303.
- [30] Bewley TR, Moin P, Temam R. DNS-based predictive control of turbulence: an optimal target for feedback algorithms. *J Fluid Mech* 2000, submitted for publication.
- [31] Farrell BF, Ioannou PJ. Stochastic forcing of the linearized Navier–Stokes equation. *Phys Fluids A* 1993;5(11):2600–9.
- [32] Hamilton JM, Kim J, Waleffe F. Regeneration mechanisms of near-wall turbulence structures. *J Fluid Mech* 1995;287:317–48.
- [33] Kim J, Lim J. A linear process in wall-bounded turbulent shear flows. *Phys Fluids* 2000;12(8):1885–8.
- [34] Bewley TR. Linear control and estimation of nonlinear chaotic convection: harnessing the butterfly effect. *Phys Fluids* 1999;11:1169–86.
- [35] Huerre P, Monkewitz PA. Local and global instabilities in spatially developing flows. *Annu Rev Fluid Mech* 1990;22:473–537.
- [36] Newborn M. Kasparov versus Deep Blue: computer chess comes of age. Berlin: Springer, 1997.
- [37] Atkinson GW. Chess and machine intuition. Norwood, NJ: Ablex, 1993.
- [38] Abergel F, Temam R. On some control problems in fluid mechanics. *Theor Comput Fluid Dyn* 1990;1:303–25.
- [39] Vainberg M. Variational methods for the study of nonlinear operators. San Francisco, CA: Holden-Day, 1964.

- [40] Press WH, Flannery BP, Teukolsky SA, Vetterling WT. Numerical Recipes. Cambridge, 1986.
- [41] Vanderplaats G. Numerical optimization techniques for engineering design. New York: McGraw-Hill, 1984.
- [42] Chang Y, Collis SS. Active control of turbulent channel flows based on large-eddy simulation. Proceedings of the Third ASME/JSME Joint Fluids Engineering Conference, FEDSM99-6929, July 18–23, San Francisco, 1999.
- [43] Vogel, Wade. Analysis of costate discretizations in parameter estimation for linear evolution equations. SIAM J Control Optim 1995;33:227–54.
- [44] Bewley TR, Temam R, Ziane M. A general framework for robust control in fluid mechanics. Physica D 2000;138: 360–92.
- [45] Tachim Medjo T. Iterative methods for robust control problems. 2000, in preparation.
- [46] Reuther J, Jameson A, Farmer J, Martinelli L, Saunders D. Aerodynamic shape optimization of complex aircraft configurations via an adjoint formulation. AIAA 96-0094, 1996.
- [47] Lumley J, Blossey P. Control of turbulence. Annu Rev Fluid Mech 1998;30:311–27.
- [48] Holmes P, Lumley JL, Berkooz G. Turbulence, coherent structures, dynamical systems and symmetry. Cambridge, 1996.
- [49] Balogh A, Liu W-J, Krstić M. Stability enhancement by boundary control in 2D channel flow. IEEE Trans Automat Control 2000, submitted for publication.
- [50] Bošković DM, Krstić M. Global stabilization of a thermal convection loop. Automatica 2000, submitted for publication.
- [51] Koumoutsakos P, Freund J, Parekh. Evolution strategies for parameter optimization in controlled jet flows. Proceedings of the 1998 CTR Summer Program, Center for Turbulence Research, Stanford University/NASA Ames, 1998.
- [52] Rathnasingham R, Breuer KS. System identification and control of a turbulent boundary layer. Phys Fluids 1997;9(7):1867–9.
- [53] Nosenchuck DM. Electromagnetic turbulent boundary-layer control. Bull Am Phys Soc 1994;39(9):1938.
- [54] Koumoutsakos P. Vorticity flux control for a turbulent channel flow. Phys Fluids 1999;11(2):248–50.
- [55] Rose N, editor. Mathematical maxims and minims. Raleigh, NC: Rome Press Inc., 1998.
- [56] Simmons G. Calculus gems. New York: McGraw-Hill, 1992.