Advances in Turbulence IX
Proceedings of the Ninth European Turbulence Conference
I.P. Castro, P.E. Hancock & T.G.Thomas (Eds.)

[id:,1]

The emerging roles of model-based control theory in fluid mechanics

T. R. Bewley

Dept. of MAE, UC San Diego, La Jolla, CA 92093-0411 USA
email: *bewley@ucsd.edu*

As traditional approaches to the analysis and simulation of fundamental flow systems begin to mature, a suite of powerful new tools leveraging model-based control theory are emerging as viable techniques to provide further understanding and to solve new problems based on these system models. In particular, these tools are contributing to three specific areas in the field of fluid mechanics: 1) open-loop system optimization, 2) closed-loop feedback control, and 3) *a priori* characterization of fundamental limitations on control effectiveness. The present extended abstract reviews some of the essential equations and surveys a few of the recent developments in the first two of these areas.

Preface: framing a suitable question

Before model-based control theory can be applied to a fluid-mechanical system, a canonical problem must first be identified that is amenable to simulation. This is not necessarily easy, as many flows of engineering interest, such as those involving combustion, are still outside the realm of accurate prediction with computational fluid dynamics (CFD). A critical look at the suitability and implications of the assumptions made in order to frame a numerically tractable flow model is essential, as such assumptions sometimes lead to reasonable flow simulations and visualizations but absurd conclusions about the related control problem. Once spatially discretized, the form of a typical such problem is depicted in Figure 1.

For a flow system which may be accurately modeled with available CFD techniques, it is usually possible to apply iterative, gradient-based approaches to optimize the controls on the system for a given purpose, as discussed in §1. However, the computation of model-based feedback rules to coordinate the controls with available measurements in such a system is a significantly more challenging problem. If the appropriate simplifications can be made, effective feedback control rules may sometimes be found, as discussed in §2.

Gut-level intuition plays a significant role in the exploratory process of research, as it guides the framing of suitable questions whose solutions, if found, can lead to fundamental new insight. To supplement this intuition, one can seek quantification of rigorous *performance limitations*, such as the minimum

Full nonlinear state equation. A fairly general ODE form of a nonlinear evolution equation relating the state $\mathbf{q}$ (a spatial discretization of the flow system) to the known external forces $\mathbf{f}$, the controls and uncertain model parameters $\boldsymbol{\phi}$, and the unknown external disturbances $\boldsymbol{\psi}$ may be written:

$$\left\{ \begin{array}{lll} E\dot{\mathbf{q}} = N(\mathbf{q}, \mathbf{f}, \boldsymbol{\phi}, \boldsymbol{\psi}) & \text{on} & 0 < t < T \\ \quad \mathbf{q} = \mathbf{q}_0 & \text{at} & t = 0 \end{array} \right\}. \tag{1}$$

The matrix E may be singular, as is the case with the "descriptor systems" formed by spatial discretization of "differential algebraic equations" such as the incompressible Navier-Stokes equation in the primative $\{u, v, w, p\}$ form.

Linearized perturbation equation. In a well-posed problem, small perturbations $\{\boldsymbol{\phi}', \boldsymbol{\psi}'\}$ to the control and disturbance vectors $\{\boldsymbol{\phi}, \boldsymbol{\psi}\}$ result in small perturbation $\mathbf{q}'$ to the state $\mathbf{q}$. Such perturbations are related by:

$$\left\{ \begin{array}{lll} \mathcal{L}\mathbf{q}' = B_{\boldsymbol{\phi}}\boldsymbol{\phi}' + B_{\boldsymbol{\psi}}\boldsymbol{\psi}' & \text{on} & 0 < t < T \\ \quad \mathbf{q}' = 0 & \text{at} & t = 0 \end{array} \right\}, \tag{2}$$

where $\mathcal{L}\mathbf{q}' = (E\, d/dt - A)\mathbf{q}'$, $B_{\boldsymbol{\phi}}\boldsymbol{\phi}'$, and $B_{\boldsymbol{\psi}}\boldsymbol{\psi}'$ are found by linearization of (1) about a base flow $\mathbf{q}(\boldsymbol{\phi}, \boldsymbol{\psi})$. The base flow may be an unsteady trajectory found by solution of (1), a steady-state solution of this system in the nominal setting $\boldsymbol{\phi} = \boldsymbol{\psi} = 0$, or some other representative time-averaged flow profile.

Cost function. In the "robust" setting, the control objective is modeled mathematically as a noncooperative game in which a "cost function" $\mathcal{J}$ is minimized with respect to the control $\boldsymbol{\phi}$ and maximized with respect to the disturbance $\boldsymbol{\psi}$, where:

$$\mathcal{J} = \frac{1}{2}\int_0^T (\mathbf{q}^* Q \mathbf{q} + \ell^2 \boldsymbol{\phi}^* \boldsymbol{\phi} - \gamma^2 \boldsymbol{\psi}^* \boldsymbol{\psi})\, dt.$$

The cost function balances a term penalizing the flow quantity of interest (turbulent kinetic energy, drag, etc.) with weighted measures of the control and the disturbance. Note that the special case obtained by taking $\gamma \to \infty$ results in $\boldsymbol{\psi} \to 0$, referred to as the "optimal" setting.

Figure 1: Typical form of a flow control problem. For simplicity, the derivation presented here is written in the (spatially discretized) ODE form. Including explicit statements of the appropriate boundary conditions in (1) and (2) and defining the relevant inner products over space-time, the PDE expressions corresponding to the derivations in Figures 1 and 2 follow in a natural fashion.

sustainable drag, and *stabilization limitations*, such as the maximum Reynolds number at which vortex shedding can be eliminated. For space considerations, discussion of our recent results on this topic must be deferred to separate papers.

1 Open-loop system optimization

1.1 Central issue: the value of gradient information

In optimization problems in which evaluations of the function $\mathcal{J}(\boldsymbol{\phi})$ are relatively "cheap" and $\boldsymbol{\phi}$ is relatively low dimensional, iterative strategies to determine a suitable value of $\boldsymbol{\phi}$ based on function evaluations alone are tractable. Thus, thousands of function evaluations $\mathcal{J}(\boldsymbol{\phi})$ may be computed, coordinated by an adaptive or "genetic" strategy, to optimize a low-dimensional $\boldsymbol{\phi}$ in turbulent flow systems when one of the following two options is available:

a) an experiment which can run unattended over a long period of time, or

b) an efficient simulation code together with a huge supercomputer allocation.

However, it is often the case in flow control problems that function evaluations $\mathcal{J}(\boldsymbol{\phi})$ are relatively expensive and $\boldsymbol{\phi}$ is relatively high dimensional. In such cases, simulation-based optimization strategies which leverage knowledge of the gradient $\mathcal{DJ}/\mathcal{D}\boldsymbol{\phi}$ are the only tractable option.

1.2 Essential theory: adjoint analysis

The principal mathematical tool for determining the gradient information central to efficient high-dimensional optimization strategies is *adjoint analysis*, as reviewed in Figure 2. Visualizations of a representative flow/adjoint simulation of a 3D turbulent flow system are depicted in Figure 3. Once the gradient $\mathcal{DJ}/\mathcal{D}\boldsymbol{\phi}$ is determined via adjoint analysis, it is straightforward to use it to optimize $\boldsymbol{\phi}$ to achieve a desired effect using existing gradient-based optimization algorithms, such as conjugate gradient, SQP, BFGS, etc.

The principal application for adjoint analysis in turbulent flow systems is *open-loop optimization*. That is, iterative computer simulations may be used to optimize, for example, the shape of a bluff body, the compliance properties of a flexible wall, or the schedule of time-periodic control inputs near the nozzle of a round jet. Once optimized via computer simulation, the result may be applied to a physical turbulent flow in hopes of achieving the desired effect with no measurements of the flow system and no further optimizations required. The success or failure of such a strategy hinges on the *generalization* of the optimized result to different turbulent flow realizations [besides the particular flow realization(s) for which the optimization was performed] at the same bulk flow conditions. Though not mathematically guaranteed, such generalization is likely if the initial optimization was performed on a sufficiently large ensemble of different turbulent flow realizations at the same operating conditions, or, equivalently (due to the ergodic nature of turbulence), on a single turbulent flow realization over a sufficiently long time period.

Adjoint analysis may also be used to coordinate the controls with knowledge of the state of the flow system, a problem generally referred to as "feedback control". In the "optimal" setting, the adjoint-based approach was in fact the first control strategy to achieve relaminarization in the benchmark problem of

Cost function perturbation. Perturbations $\{\boldsymbol{\phi}', \boldsymbol{\psi}'\}$ to the control and disturbance $\{\boldsymbol{\phi}, \boldsymbol{\psi}\}$ result in a perturbation $\mathcal{J}'$ to the cost function $\mathcal{J}$, where:

$$\mathcal{J}' = \int_0^T (\mathbf{q}^* Q \mathbf{q}' + \ell^2 \boldsymbol{\phi}^* \boldsymbol{\phi}' - \gamma^2 \boldsymbol{\psi}^* \boldsymbol{\psi}')\, dt = \int_0^T \left[\left(\frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\boldsymbol{\phi}} \right)^* \boldsymbol{\phi}' + \left(\frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\boldsymbol{\psi}} \right)^* \boldsymbol{\psi}' \right] dt.$$

Adjoint calculus is used to re-express $\mathcal{J}'$ in the integral form depicted on the right-hand side, from which the desired gradient vectors may be identified.

Statement of adjoint identity. Define an inner product $\langle \mathbf{r}, \mathbf{q}' \rangle = \int_0^T \mathbf{r}^* \mathbf{q}'\, dt$. By integration by parts, it follows that:

$$\langle \mathbf{r}, \mathcal{L}\mathbf{q}' \rangle = \langle \mathcal{L}^* \mathbf{r}, \mathbf{q}' \rangle + \mathbf{b},$$

where $\mathcal{L}^* \mathbf{r} = \Big(- E^* \frac{d}{dt} - A^* \Big) \mathbf{r}$ and $\mathbf{b} = \mathbf{r}^* E \mathbf{q}' \Big|_{t=T} - \mathbf{r}^* E \mathbf{q}' \Big|_{t=0}$.

Definition of adjoint equation. Consider the following definition:

$$\left\{ \begin{array}{llll} \mathcal{L}^* \mathbf{r} = Q\mathbf{q} & \Rightarrow \quad -E^* \dot{\mathbf{r}} = A^* \mathbf{r} + Q \mathbf{q} & \text{on} & 0 < t < T \\ \quad \mathbf{r} = 0 & & \text{at} & t = T \end{array} \right\}. \tag{3}$$

The "adjoint field" $\mathbf{r}$ so defined is fairly easy to compute, requiring a simulation code of complexity and dimension similar to the original DNS code. Note that the adjoint system marches *backward* in time, from $T \to 0$, though $A^* = A^*(\mathbf{q})$, where $\mathbf{q}$ is computed by marching forward in time, from $0 \to T$. This presents a storage problem, which is easily averted with a "checkpointing" algorithm.

Extraction of gradients. Combining equations, it follows that

$$\langle \mathbf{r}, B_{\boldsymbol{\phi}} \boldsymbol{\phi}' + B_{\boldsymbol{\psi}} \boldsymbol{\psi}' \rangle = \langle Q\mathbf{q}, \mathbf{q}' \rangle \quad \Rightarrow \quad \int_0^T \mathbf{q}^* Q \mathbf{q}'\, dt = \int_0^T \mathbf{r}^* (B_{\boldsymbol{\phi}} \boldsymbol{\phi}' + B_{\boldsymbol{\psi}} \boldsymbol{\psi}')\, dt,$$

and thus $\mathcal{J}'$ may be written in the form:

$$\mathcal{J}' = \int_0^T \left[\left(B_{\boldsymbol{\phi}}^* \mathbf{r} + \ell^2 \boldsymbol{\phi} \right)^* \boldsymbol{\phi}' + \left(B_{\boldsymbol{\psi}}^* \mathbf{r} - \gamma^2 \boldsymbol{\psi} \right)^* \boldsymbol{\psi}' \right] dt.$$

As $\boldsymbol{\phi}'$ and $\boldsymbol{\psi}'$ are arbitrary, it follows that the desired gradient information may be identified as a simple function of the solution of the adjoint problem (3):

$$\frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\boldsymbol{\phi}} = B_{\boldsymbol{\phi}}^* \mathbf{r} + \ell^2 \boldsymbol{\phi} \qquad \text{and} \qquad \frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\boldsymbol{\psi}} = B_{\boldsymbol{\psi}}^* \mathbf{r} - \gamma^2 \boldsymbol{\psi}. \tag{4}$$

Figure 2: The essential steps in the gradient calculation for the iterative optimization of the problem framed in Figure 1 for the full nonlinear system (1).

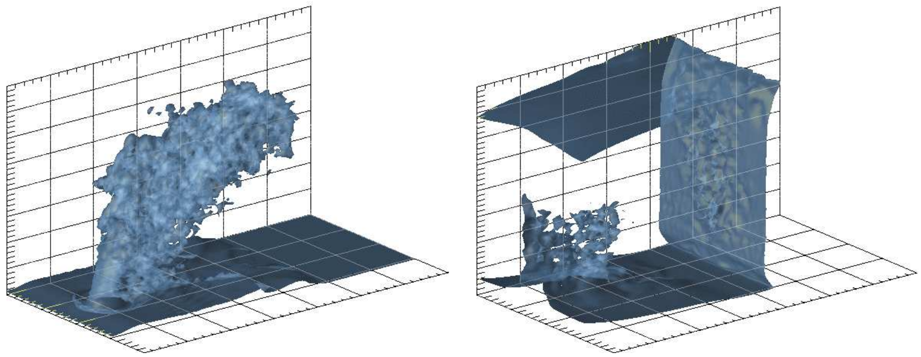

Figure 3: State simulation (left) for the problem of mixing enhancement in a cross-flow jet. The bulk flow is from left to right, and a cross-flow jet flow emanates from a hole in the bottom wall. In this system, a cost function based on the mixing of these two flows is defined downstream. The relevant adjoint simulation (right) is then marched backward in time and upstream in space in order to determine where additional forcing on this system may be applied to enhance mixing. Simulation by G. Hagen; animations are available at: `http://turbulence.ucsd.edu/gallery.html`.

stabilization of low Reynolds number channel-flow turbulence using distributed blowing/suction on the walls as the control coordinated with full information about the flow state. In this work, effective unsteady controls were optimized iteratively a short time segment into the future, these controls were applied to the flow, then the optimization process begun anew on the following time segment, etc. This framework is referred to in the controls community as receding-horizon Model Predictive Control (MPC).

Though successfully demonstrating that small amounts of blowing/suction controls on the wall are in fact able to relaminarize near-wall turbulence if coordinated appropriately with the flow state, there are fundamental shortcomings with the MPC approach in the optimal setting when one considers the practical issues of controlling turbulence in physical systems. The primary shortcoming of this algorithm is simply its computational expense. The algorithm requires online iterative direct numerical simulations (DNS) and convergence of the iterative control optimization process on a time scale which is small with respect to the evolution of the physical turbulent flow system. For most systems of interest (global weather systems being the notable exception) this requirement is not reasonable, even if one extends Moore's law well into the next century. Another shortcoming of this approach is that it assumes a perfect computational model and complete information about the flow state, neither of which is available in the practical setting. Due to these shortcomings, alternative techniques to solve the feedback control problem are required, as discussed in §2.

1.3 Key elements: regularization and noncooperation

A subtle aspect of the adjoint framework is that inner products are used (or implied, if not explicitly stated) in three distinct steps of the adjoint-based optimization process. Flexibility in the definition of these three inner products provides the capability of "regularizing", or making smooth, the different fields involved in the optimization procedure, allowing one to:

1) focus the cost function on the dynamics of interest,
2) define a "smooth" adjoint field that may be computed accurately, and
3) appropriately "precondition" the gradient for rapid convergence.

When performing adjoint simulations such as that depicted in Figure 3, we have found that adjoint fields tend to "blow up" in thin shear layers unless the issue of regularization is attended to carefully. Thus, the issue of regularization is not just a mathematical curiosity; rather, proper attention to this issue is a numerical necessity. Recent attempts to develop a comprehensive framework which clarifies the relationship between the various inner products used in the optimization problem and the three regularization objectives listed above are discussed further in the abstract by B. Protas included elsewhere in these proceedings.

The adjoint-based optimization procedure in the optimal setting assumes a perfect computational model of the flow system. Unfortunately, a common outcome of optimizations performed under this assumption is excellent system performance at exactly the design conditions but a terrible degradation of performance at "off-design" conditions, a situation known as "over-optimization" of the system. The idea of noncooperative optimization is a natural cure for this problem, as illustrated by the following example.

The phenomenon of over-optimization is well known in the field of airfoil shape optimization, where adjoint-based optimizations in the optimal setting often lead to airfoils that work quite well at exactly the design point but perform quite badly when there is a small perturbation to the Mach number, Reynolds number, or angle of attack. The current operational "patch" for this problem is to increase the number of "design points" from one to about six, though such "multi-point" optimization strategies can still lead to degraded performance between the several design points. An alternative strategy to "detune" the optimization procedure is to determine the "best" airfoil shape in the presence of a disturbance to the uncertain parameters with the "worst" possible structure. This is the essence of *noncooperative optimization*, and may be achieved using the framework laid out in Figures 1 and 2 simply by choosing a finite (but sufficiently large) value for γ, then optimizing the system with a saddle-point algorithm, simultaneously minimizing $\mathcal{J}(\boldsymbol{\phi},\boldsymbol{\psi})$ with respect to the shape parameters $\boldsymbol{\phi}$ and maximizing $\mathcal{J}(\boldsymbol{\phi},\boldsymbol{\psi})$ with respect to the uncertain model parameters $\boldsymbol{\psi}$. By so doing, an airfoil shape is found which is maximally effective even in the presence of variations of the uncertain parameters which are maximally disruptive.

Adjoint-based optimization of turbulent flow systems requires iterative direct numerical simulations (or, perhaps, large eddy simulations). Efficient numerical techniques and ample computational resources are thus essential.

2 Closed-loop feedback control

2.1 Central issue: simplification of the control problem

In mechanical systems comprised of a few dozen masses, springs, and dashpots, or electrical systems comprised of as many inductors, capacitors, and resistors, equations for matrices with N^2 unknowns (where N is the dimension of the state) are easily managed. Linear control theory was born out of the study of such systems. There is now a wide body of literature and several textbooks which discuss how to estimate the state of linear systems based on limited noisy measurements and how to control linear systems based on these state estimates, a framework known as $\mathcal{H}_2/\mathcal{H}_\infty$ control theory. Further, there is a growing body of literature on how to extend such control strategies to the nonlinear regime.

ODE discretizations of turbulent flow systems typically have $N \geq O(10^6)$. Such systems are often called *multiscale*, as they are characterized by a complex cascade of energy over a broad range of length scales and time scales. In such systems, direct application of matrix-based control strategies is completely intractable. Further, due to the multiscale nature of turbulent flows, the most appropriate technique to reduce such system descriptions to a manageable size while still retaining their essential features is not obvious. This is one of the central difficulties with model-based control of fluid-mechanical systems. We now survey some of the relevant issues when simplifying such system models.

Modal decoupling with Fourier transforms. In the linear setting, the 3D Navier-Stokes equation completely decouples on a mode-by-mode basis when *all* variables with spatial variation are Fourier transformed in the homogeneous directions of the flow. The allows the complete decoupling of an extremely large linear control problem into several small linear control problems. This is one of the best techniques available to simplify the control problem at hand, and has been in use in our group since 1996.

Boundary-layer approximation. A exciting new strategy to simplify certain (otherwise intractable) $\mathcal{H}_2/\mathcal{H}_\infty$ control problems is to apply the boundary-layer approximation to the system as it develops parabolicly in the streamwise coordinate. This technique, currently under development by P. Cathalifaud, is quite interesting from the control theoretic perspective, as *noncausal feedback* (that is, using downstream measurements to update the controls upstream) is necessary for maximum closed-loop system performance.

Model reduction: open-loop versus closed-loop. The qualities required of a simplified system model for accurate dynamic simulation are quite different than the qualities required for effective closed-loop feedback design. What is necessary for the latter problem is that the model capture the relevant "input-output" relationship between the control inputs and the quantities of interest in the system model. Unfortunately, control theory is found to be rather inferior in the area of open-loop model reduction; that is, it is difficult to know how to reduce a system description while retaining its "essential features" until after the feedback control problem is solved. However, the strategies available to reduce the system model in closed loop (after the feedback control is applied) in order to simplify the implementation are quite promising, and should be quite useful in the field of flow control when such strategies can be made tractable.

Data-based model reduction. Closing a control loop significantly alters the global dynamics of a given system. The introduction of a single sensor/actuator pair into a spatially-distributed system at least imposes a new length scale (the distance between the sensor and the actuator). The new dynamics created by closing the control loop often dominate the behavior of the entire system. Thus, data-based models, such as those based on Proper Orthogonal Decompositions (POD), are quite tricky to use, as the decomposition itself must be evolve as the control loop is closed in order to be of practical use.

2.2 Essential theory: Riccati analysis

The principal mathematical tool for model-based closed-loop feedback control is the numerical solution of the *Riccati equation* characterizing the relationship between the relevant perturbation and adjoint systems, as shown in Figure 4.

In recent years, the important issue of non-normality of the eigenvectors of the system matrix A [see (2)] in linearized descriptions of shear flows has received a growing amount of attention. It is important to note that the control theory laid out in Figure 4 effectively addresses this issue from an appropriate input-output perspective. Controls determined by the $\mathcal{H}_2/\mathcal{H}_\infty$ approach are found to make the eigenvectors of the closed-loop system matrix significantly closer to orthogonal, thereby reducing the sensitivity of the system to excitation by external disturbances.

Leveraging the modal decoupling of the Fourier representation (as mentioned previously) to solve the control problem, and then inverse transforming the resulting feedback gains to the physical domain, spatially-localized convolution kernels relating, e.g., the state of the system inside a flow domain to control forcing on the wall may be obtained, as shown in Figure 5. Theoretical motivation for this result is provided by the analysis of B. Bamieh and coworkers.

2.3 Key element: decentralized control logic

With spatially-localized feedback structure, such as that implied by Figure 5, decentralized control of wall-bounded flows becomes possible, as illustrated in Figure 6. 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.

Characterization of saddle point. By (4), the control $\boldsymbol{\phi}$ which minimizes $\mathcal{J}$ and the disturbance $\boldsymbol{\psi}$ which simultaneously maximizes $\mathcal{J}$ are given by

$$\frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\boldsymbol{\phi}} = 0, \quad \frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\boldsymbol{\psi}} = 0 \qquad \Rightarrow \qquad \boldsymbol{\phi} = -\frac{1}{\ell^2} B_{\boldsymbol{\phi}}^* \mathbf{r}, \quad \boldsymbol{\psi} = \frac{1}{\gamma^2} B_{\boldsymbol{\psi}}^* \mathbf{r}. \tag{5}$$

Combined matrix form. Combining the perturbation and adjoint equations (2) and (3) at the saddle point (5), assuming $E = I$, gives:

$$\begin{array}{r} \text{Perturbation equation} \rightarrow \\ \text{Adjoint equation} \rightarrow \end{array} \begin{bmatrix} \dot{\mathbf{q}}' \\ \mathbf{r} \end{bmatrix} = \begin{bmatrix} A & \overbrace{-\frac{1}{\ell^2} B_{\boldsymbol{\phi}} B_{\boldsymbol{\phi}}^* + \frac{1}{\gamma^2} B_{\boldsymbol{\psi}} B_{\boldsymbol{\psi}}^*}^{\text{control and disturbance}} \\ -Q & -A^* \end{bmatrix} \begin{bmatrix} \mathbf{q}' \\ \mathbf{r} \end{bmatrix}.$$

Solution Ansatz. To make the feedback control computation possible, we now assume a general form of the relationship between the perturbation $\mathbf{q}' = \mathbf{q}'(t)$ and the adjoint $\mathbf{r} = \mathbf{r}(t)$: $\quad \mathbf{r} = X\mathbf{q}' \quad$ where $\quad X = X(t)$.

Riccati equation. Inserting solution ansatz into the combined matrix form to eliminate $\mathbf{r}$ and combining rows to eliminate $\dot{\mathbf{q}}'$ gives:

$$\left[-\dot{X} = A^* X + XA - X \left(\tfrac{1}{\ell^2} B_{\boldsymbol{\phi}} B_{\boldsymbol{\phi}}^* - \tfrac{1}{\gamma^2} B_{\boldsymbol{\psi}} B_{\boldsymbol{\psi}}^* \right) X + Q \right] \mathbf{q}'.$$

As this equation is valid for all $\mathbf{q}'$, it follows that:

$$-\dot{X} = A^* X + XA - X \left(\tfrac{1}{\ell^2} B_{\boldsymbol{\phi}} B_{\boldsymbol{\phi}}^* - \tfrac{1}{\gamma^2} B_{\boldsymbol{\psi}} B_{\boldsymbol{\psi}}^* \right) X + Q.$$

This is a matrix differential equation, and may be solved with standard techniques. Due to the terminal conditions on $\mathbf{r}$, $X = 0$ at $t = T$. By the solution ansatz together with (5), we have

$$\boldsymbol{\phi} = K\mathbf{q}' \qquad \text{where} \qquad K = -\frac{1}{\ell^2} B_{\boldsymbol{\phi}}^* X.$$

The matrix K is referred to as the finite-horizon $\mathcal{H}_\infty$ control feedback gain.

Linear Time Invariant (LTI) systems. If the system is LTI and the closed-loop system is stable, it follows after some analysis that X approaches a constant symmetric positive definite value in the backwards march for large T. Taking the limit as $T \to \infty$, X solves the quadratic matrix equation

$$0 = A^* X + XA - X \left(\tfrac{1}{\ell^2} B_{\boldsymbol{\phi}} B_{\boldsymbol{\phi}}^* - \tfrac{1}{\gamma^2} B_{\boldsymbol{\psi}} B_{\boldsymbol{\psi}}^* \right) X + Q.$$

Efficient techniques to solve this quadratic equation for X are well developed.

Figure 4: The essential steps in the Riccati-based feedback calculation for the control problem framed in Figure 1 for the linearized system (2). Similar ("dual") procedures can be framed to compute the appropriate estimator feedback.

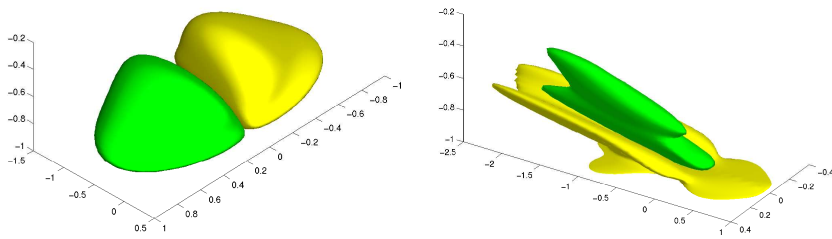

Figure 5: Spatially-localized convolution kernels for $\mathcal{H}_\infty$ control of Poiseuille flow. Calculation by M. Högberg.

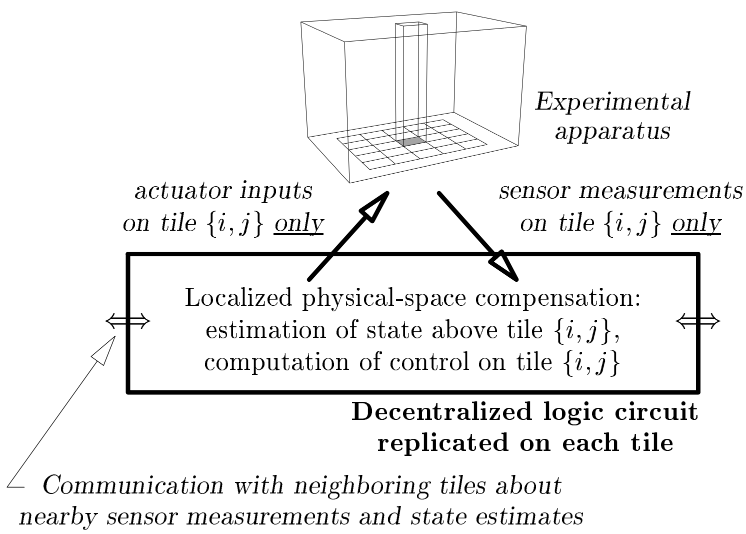

Figure 6: Decentralized communication structure facilitated by spatially-localized feedback rules.

Acknowledgments

Generous funding from Belinda King's and Steve Walker's programs at AFOSR is gratefully acknowledged.

References

Space limitations do not permit a survey of the broad literature on flow control in the present abstract. A recent survey on this topic may be found here:

T. R. Bewley, Flow control: new challenges for a new Renaissance, *Progress in Aerospace Sciences* 37:21-58, 2001.