Closed-Loop Control of Fluid Flow: a Review of Linear Approaches and Tools for the Stabilization of Transitional Flows

D. Sipp
(Onera)
P. Schmid
(LadHyX Ecole Polytechnique)

E-mail: denis.sipp@onera.fr

Flow control is concerned with the targeted manipulation of intrinsic flow behavior to optimally satisfy prescribed objectives. This article will give an overview of the most common tools for the design of control strategies. We focus on linear control that is aimed at stabilizing fixed points of the Navier-Stokes equations, such as those existing in the case of transitional flows. Key steps to build a Galerkin-based and a data-based model will be presented and illustrated on two generic flow configurations: flow over an open cavity and over a backward-facing step. In the former case, a feedback configuration will result and particular attention will be paid to performance measures and robustness analyses. For the latter case, a feedforward setup has been chosen and a system-identification method will be employed to extract the fluid behavior from data sequences, via statistical learning techniques.

Introduction

The control of fluid flow is a central component of fluid mechanics, since it holds great promise in manipulating the inherent flow behavior of a fluid system. The suppression of instabilities, the exploration of previously inaccessible parameter ranges, efforts to increase stability margins and diminish sensitivity to external noise and attempts to improve efficiency and minimize environmental impact (less pollution, less unburnt gases, less NOx emission, less sound radiation, better and more homogeneous mixing) are just a few technological areas where active flow control could make an appreciable difference. Much effort has thus been concentrated on applying control-theoretic methods to fluid systems and on establishing a theoretical framework and foundation for the design of effective flow control strategies.

The design of flow control strategies relies on a model for the fluid system, but also on a model for the noise environment. For flows that are weakly sensitive to external noise - referred to as oscillator flows - it is difficult to design effective control strategies due to the inherent nonlinearities that drive the system; flow over a cavity at supercritical Reynolds numbers is a typical example of such a flow type. For flows that respond sensitively to environmental noise - known as amplifier flows - an additional difficulty lies in the design of an adequate model of the noise environment; a two-dimensional flow over a backward-facing step would fall into this category.

Past work in flow control has heavily relied on models using a projection of the Navier- Stokes equations onto a reduced basis, for example a POD basis. Both linear (see [1, 2, 3, 4, 5, 6]) and non-linear state-space models (see [7, 8]) can be designed. These models then form the foundation on which controllers can be built. Various control design approaches, from the simple PI (proportional integral) controller [9] to the more complex H_2 or H∞-control (see, for example, [10], [11], [12]) can be followed, each of which can be characterized by a given performance - to what extent has the control objective been reached? - and a given robustness - how does the performance change as the reduced-order models are slightly perturbed?

A different approach to the design of control strategies consists in using data-based models, which are built from observing and evaluating sequences of system input and output signals. A discrete transfer function is extracted from a learning data-set, via system identification techniques [13]. This transfer function then forms the foundation for an optimal control design. First, attempts to design control schemes following this approach have been undertaken [14, 15], which showed that the underlying model - in particular the incorporation of unobservable noise sources - is of critical importance. The designed compensator showed a great amount of robustness to off-design noise environments and even to non-linear effects [15].

In this article, we will present and discuss the critical steps in the design of flow control schemes for both types of flow behavior (oscillators and amplifiers) and compare and contrast models derived from Galerkin projections and data-based models. We stress that this review focuses on the linear stabilization of equilibrium points existing in transitional flows. For more general approaches of closed-loop control - which aim, for example, at manipulating turbulent flows - the reader is referred to [16], [17], [18], [19], [20].

Description of a flow control problem

The motion of the fluid is assumed to be governed by an evolution equation of the general form

$$\dot{x} = f(x,u,w) \tag{1}$$

where x is the state vector containing all of the variables necessary to uniquely describe the state of the fluid system (such as velocities, pressure, density, temperature, mixture fractions, etc.), u is the control vector which contains input variables that are used to influence the flow, and w is the (stochastic) noise vector summarizing the influences of external noise sources or inaccuracies (freestream disturbances, acoustic noise sources, wall roughness distributions, etc.). The time rate of change $\dot{x}$ is related to these vectors by a (generally non-linear) function f. In this article, we assume that an equilibrium point x_0 of these equations exists, such that.

$$f(x_0, u=0, w=0) = 0 \tag{2}$$

The existence of such equilibrium points is assured in the case of transitional flows, as considered in this paper. In light of this, the control is thus aimed at maintaining the system as close as possible to these fixed points, where closeness is measured by a performance sensor $z = c_z\ (x\text{-}x_0)$. If the function c_z represents the deviation of the full state from the equilibrium solution, the perturbation kinetic energy can be targeted as a control objective and transition delay issues may be considered. However, c_z may also represent a local measure of the flow, for example the signal from a wall-shear stress sensor. The controller will minimize the following cost-functional

$$J = J(\mathrm{z},\mathrm{u}) \tag{3}$$

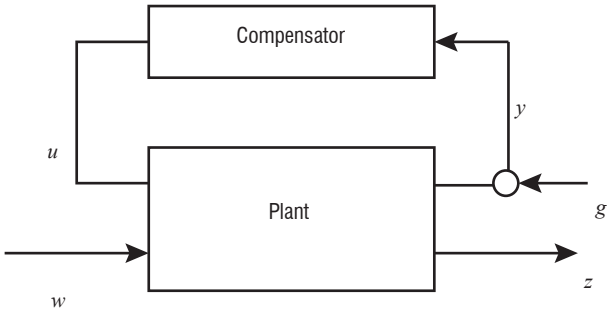

Figure 1 - Block diagram of a typical feedback control setup, including plant, compensator, external noise sources *(w, g)* and objective output z

which expresses the amount of gain from applying control strategies to the fluid system. It also contains the control vector u since, in most cases, we try to avoid excessive control input and thus wish to penalize the expended control energy. For example,

$$J(z,u) = \int_0^\infty \left\| c_z\,(x\text{-}x_0) \right\|^2 + l^2 \left\| u \right\|_u^{\ 2} dt \quad \text{where } \|\cdot\| \text{ and } \|\cdot\|_u$$

are suitable norms and l^2 is a user-specified parameter related to the cost of control. In a closed-loop set-up, the control input u depends via the controller on some measure y of the flow (the estimation measure) given by

$$y = c_y(x,g) \tag{4}$$

where y corresponds to a measurement vector, describing the information that we extract from the state vector x (for example, by shear sensors placed on the wall, by hot-wire probes placed in the wake of an obstacle, by microphones placed in the acoustic far-field, or by TR-PIV measurements in a two-dimensional slice). This measure is commonly different from the performance measure z. Realistically, the estimation measurements must be assumed to be contaminated by measurement noise g. The function c_y is thus a function that extracts exploitable (and measurable) information from the entire flow field.

The objective of our control problem is thus to find a mapping $y \rightarrow u$ (the compensator) that minimizes the cost functional (3) while observing the constraints given by the governing equation (1). Mathematically, we must solve a constrained optimization problem. The complete flow control setup is graphically illustrated in the form of a block diagram in figure 1. The flow system (labeled “plant”) is influenced by external noise w and driven by the control signal u from the compensator unit. The measurement signal y, contaminated by sensor noise g, is the output from the plant and is, in turn, used in the compensator unit to ultimately determine the optimal control signal u. The goal of the compensator is to force the plant in such a way that the performance measurement (z) is minimal, while at the same time excessive control input (u) is avoided.

For effective flow control, the placement of actuators and sensors plays a vital role, and two generic configurations can be distinguished, based on the inherent behavior of the fluid flow. The first category is referred to as oscillator flows, while the second category is known as amplifier flows.

Oscillator flows: the case of an open cavity

In the case of oscillator flows, the flow generally undergoes a supercritical Hopf bifurcation at a critical Reynolds number, becomes globally unstable and establishes limit-cycle behavior. The shedding of vortical structures in the wake of a cylinder above a critical Reynolds number of $R_{ec} = 47$ is a quintessential and much-studied example of an oscillator flow. This type of flow is characterized by a relative insensitivity to external perturbations and the dominant frequency of the flow is little influenced by noise sources. The control objective in this case can be two-fold. We can maintain the flow in the vicinity of the unstable equilibrium points using linear control theory. Alternatively, and more challenging, we can force the fluid system away from its natural limit-cycle behavior towards the equilibrium state; this effort requires a non-linear underlying model and a non-linear approach to flow control, especially if the Reynolds number is chosen well above its critical value. The scope of this article allows us to only concentrate on the first choice: a linear control effort near the critical Reynolds number. As a representative oscillator flow, we choose the flow over an open cavity at $R_e = 7500$ (see figure 2 (a) for a sketch of the geometry), with one actuator placed upstream of the cavity and one sensor located downstream of the cavity edge. In this setup, the upstream control u influences the downstream estimation sensor y, yielding a feedback configuration.

Amplifier flows: the case of a backward-facing step

Amplifier flows are globally stable flows that are characterized by a strong response to external forces. Their flow behavior is typically influenced and driven by upstream (unknown) noise sources and the

observed flow structures are the result of a frequency-selective amplification/damping of this noise environment. Even though globally stable, the flow is convectively unstable, and perturbations are amplified locally and convected downstream by the base flow. For sufficiently weak noise sources, the dynamics of amplifier flows can be suitably described by a linear model. In such a configuration, we can estimate the state downstream of the estimation sensor and manipulate the flow downstream of the actuator. A sensor-actuator setup (with a sensor upstream of the actuator) is thus appropriate for amplifier flows, whereas an actuator-sensor setup (with an actuator upstream of the sensor) is more suited for oscillator flows. A prototypical example of an amplifier flow is the flow over a backward-facing step at $R_e = 500$. Figure 2(b) illustrates the geometry of the flow and the appropriate control setup. The controller will be designed to cancel the perturbation field estimated by the sensor measurements. Due to the convective nature of the flow, the actuator signal u has no impact on the measurement y. This layout is known as a feedforward configuration.

Linear control

Even though techniques exist for designing effective control strategies for the above system of non-linear equations, it is often convenient and sufficient to apply control to a linear system describing the evolution of small perturbations around the equilibrium point x_0. In this case, the above equations simplify to

$$\dot{x} = Ax + B_u u + B_w w, \tag{5}$$

$$y = C_y x + g \tag{6}$$

$$z = C_z x \tag{7}$$

Here $A = \partial f / \partial x \,|\, x_0$ and analogously for the other matrices $B_{u,w}$ and $C_{y,z}$. In what follows, we will work with the linearized set of equations (5)-(7).

A control strategy that minimizes our cost objective J (involving the quantity z) simply consists in finding a linear mapping between the measurement y and the control signal u. The design of this strategy depends on a representation of the governing equations (5)-(7) but cannot rely on explicit knowledge of the state vector x. The controller must operate in real-time mode and thus a reduced-order model (ROM) of the governing equations (5)-(7) with significantly lower degrees of freedom $\bar{x}$ is needed that still reproduces the input-output relation of the full system with sufficient fidelity. More precisely, for

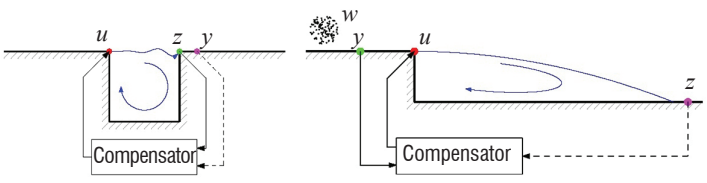

a) feedback configuration b) feedforward configuration

Figure 2 - Sketch of flow over an open cavity and a backward-facing step - two generic flow configurations representing an oscillator and noise-amplifier flow, respectively. The actuator (u), flow sensor (y) and performance sensor (z) are marked by colored symbols for each configuration. External upstream noise sources are indicated by w in the case of an amplifier flow. The cavity and backward-facing step configurations are defined in more detail in Barbagallo et al. [5] and Hervé et al. [15]

an effective controller design it is mandatory for the reduced-order model to recover the dynamic links between all inputs (w and u) and all outputs (y and z). When building this model, the state vector x can be thought of as an internal (hidden) variable describing the full state of the fluid system; typically, we do not have access to the state vector x.

Two approaches for the construction of a reduced-order representation emerge. First, a Galerkin-based approach can be taken that uses appropriately chosen flow structures to reduce the governing equations, via a Galerkin projection, followed by the design of an estimator that computes an approximate (reduced) state $\hat{x}$ by optimally processing any mismatch in the measurement vectors from the reduced-order system ($\bar{y}$) and a duplicate system ($\hat{y}$). Alternatively, we can eliminate the state vector x and postulate a direct relation between (u,w) and (y, z). The coefficients of this relation can be determined by observing and processing temporal sequences of control and measurement signals. The first approach yields the classical Kalman filter and LQG-compensator, while the second approach utilizes solely input-output information and produces disturbance-rejection control strategies. Both approaches - their underlying assumptions, strengths and shortcomings - will be discussed in more detail below.

Approach I : state-space model, observer and feedback control

The governing equations (5)-(7) constitute a mapping from a low-dimensional signal u via a high-dimensional state x to low-dimensional signals y and z. In the first approach, we follow a similar mapping and try to find a reduced-order model of the form

$$\dot{\bar{x}} = \bar{A}\bar{x} + \bar{B}_u u + \bar{B}_w w \tag{8}$$

$$y = \bar{C}_y \bar{x} + g \tag{9}$$

$$z = \bar{C}_z \bar{x} \tag{10}$$

The new state vector $\bar{x}$ represents a vector of coefficients in a linear combination of flow structure, i.e., $x = V\bar{x}$. The flow structures form the columns of the rectangular matrix V and are typically chosen from proper orthogonal decomposition (POD) modes or balanced POD (bPOD) modes. By selecting only a small number of modes, we arrive at the reduced-order system (8)-(10).

As mentioned previously, the reduced-order model must accurately match the unreduced dynamics from all input to all output variables. A quality measure for this match is based on transfer functions. The transfer functions can be derived by applying a Laplace transform to (8)-(10), which yields an algebraic expression between input and output variables. For example, the transfer function from u to y reads $\bar{T}_{yu}(s) = \bar{C}_y (s\bar{I} - \bar{A})^{-1} \bar{B}_u$, with s as the Laplace variable. Evaluating the transfer function $\bar{T}_{yu}(s)$ along the imaginary s-axis describes the frequency response of the input-output system. For control design purposes, a close match (in a norm to be specified) between the transfer function of the reduced system $\bar{T}_{yu}(s)$ and the transfer function of the original system $T_{yu}(s) = C_y (sI - A)^{-1} B_u$ is required. This quality measure, $\left\| T_{yu}(s) - \bar{T}_{yu}(s) \right\|$, can be applied to stable and unstable configurations; in the case of unstable systems, it requires that the unstable eigenvalues of the reduced and unreduced systems coincide and that a close match for the stable subspace-dynamics be attained. The quality of the match between the two transfer functions is influenced by the number and type of flow structures used in the reduction basis V.

In figure 3 (a), we have represented with a black dashed line the magnitude of the full-plant transfer function $|T_{yu}|$. This result can be obtained by directly evaluating the quantity $C_y(sI - A)^{-1}B_u$ for various frequencies $s = i\omega$, where A is the large-scale Jacobian matrix and appropriate numerical methods are required to perform the inversion. The peaks in the transfer function are linked to a series of global modes displaying Kelvin-Helmholtz type instabilities on the shear-layer : the first peak corresponds to a stable response, while the four subsequent peaks are linked to an unstable response.

How to build a reduced state-space model

Two approaches can be distinguished. The first approach is based on the reduction of the original system via a Petrov-Galerkin projection, as mentioned earlier. A bi-orthogonal basis for the state vector x consists of two rectangular matrices V, W satisfying $W^H V = \bar{I}$, with $\bar{I}$ as the identity matrix and ^{H} denoting the conjugate transpose. Expressing the state vector in this basis yields a reduced-order model for the coefficient vector $\bar{x} = W^H x$. The governing equations (5)-(7) can then be replaced by the reduced system (8)-(10) with $\bar{A} = W^H A V$. The expressions for the remaining matrices in (8)-(10) follow analogously. The critical question remains on how to choose the basis vectors in V to arrive at a meaningful reduced-order system that serves as a proper substitute for the original system, for the subsequent compensator design. A common approach of computing the bases V and W proceeds as follows: once a base-flow is computed, the direct and adjoint unstable global modes are extracted; this is followed by computing impulse responses for the various inputs (using the direct Navier-Stokes operator) and outputs (using the adjoint Navier-Stokes operator), each projected onto the stable subspaces. Processing of the snapshots with the balanced POD (bPOD) technique [2] straightforwardly yields the various matrices for the reduced-order system [5]. Alternatively, balanced modes can also be computed in the frequency domain [21], which directly produces the stable and unstable dynamics (without) having to compute the unstable global modes beforehand). The reduced state $\bar{x}$ then consists of the amplitudes of the unstable global modes and the amplitudes of the balanced modes.

In a second model-reduction approach, the same model can be obtained using system identification methods. Given the structure of the model, its unknown coefficients are identified from time-series of the input and output signals. Theoretically, this yields the same model as determined by the first (Petrov-Galerkin) approach. If the model is stable, the eigenvalue realization algorithm (ERA) technique directly yields the reduced system [22] without the need for a Jacobian matrix A and adjoint simulations; if the system is unstable, prior knowledge of a controller that maintains the dynamics of the system in the vicinity of the equilibrium point is required [23]. Then, the ERA technique is applied to identify the closed-loop dynamics, from which the open-loop dynamics can be deduced in a straightforward manner.

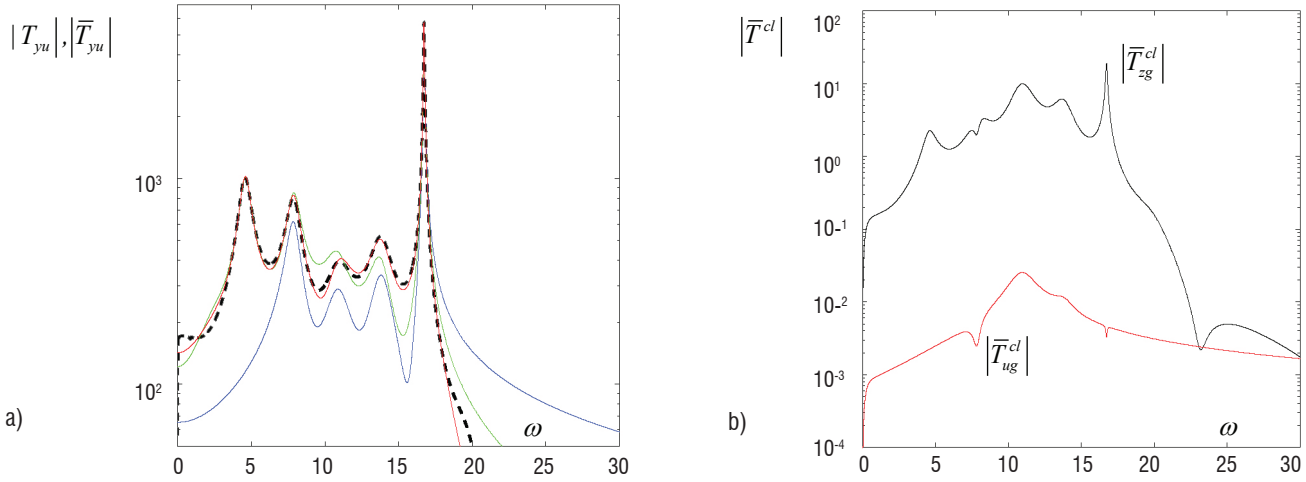

Figure 3 - a) Open-loop transfer function $|T_{yu}(\omega)|$ of the full plant (black dashed line) and of reduced- order models based on the four unstable global modes and no (blue), four (green) and eight (red) balanced POD modes to represent the stable subspace. b) Closed-loop transfer functions from the measurement noise g to the sensor output z (black) and from the measurement noise g to the control input u (red).

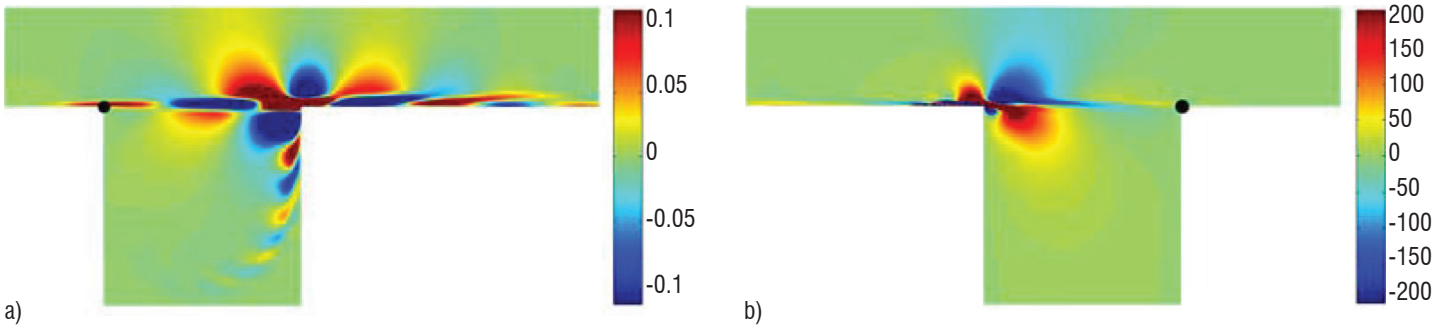

Figure 4 - Dominant balanced modes computed by the bPOD method for flow over an open cavity (visualized by streamwise velocity contours). a) The dominant direct mode, determined from the impulse responses of the actuator location (black symbol); b) the adjoint mode, determined from the adjoint impulse response of the sensor location (black symbol).

In figure 4 (a, b), we show the dominant direct and adjoint bPOD modes. These structures are respectively located near the downstream and upstream edges of the cavity, a consequence of the non-normality of the linearized Navier-Stokes operator. In figure 3 a, the magnitude of the transfer functions of models comprising the four unstable global modes plus zero (blue line), four (green line) and eight (red line) bPOD modes are displayed. A large error is observed if only the unstable global modes are considered in the model. This justifies the fact that additional modes must be considered in the model, to take into account features of the stable subspace. If more than eight bPOD modes are used, the reduced transfer function coincides with the full-plant transfer function.

Design of a compensator

In general, the reduced-order state $\bar{x}$ is not accessible to the controller; for this reason, an approximation to $\bar{x}$ must be determined directly from the available measurements y. The estimated state vector will be denoted by $\hat{x}$, which is governed by a Kalman filter written as

$$\dot{\hat{x}} = \bar{A}\hat{x}+\bar{B}_u u+L\left(y-\hat{y}\right) \tag{11}$$

$$\hat{y} = \bar{C}_y\hat{x} \tag{12}$$

This system of equations for the estimated state $\hat{x}$ resembles (8)-(9), except for the absence of noise terms and for the additional term $L\left(y-\hat{y}\right)$. This latter term represents a driving term based on the discrepancy between the true measurement y and the measurement from the proxy system (11)-(12). The manner in which this discrepancy is processed is given by the quantity L, which is determined by posing an optimization problem : we wish to minimize the covariance of the estimation error $\bar{x}-\hat{x}$, while observing the governing equations (11)- (12). Covariances appear due to the fact that y and thus $\hat{x}$ are stochastic variables and the statistical moments of $\hat{x}$ are the appropriate quantities to focus on. The optimization problem can be formulated variationally, yielding the following algebraic matrix Riccati equation for an auxiliary variable P, according to

$$\bar{A}P + P\bar{A}^H - P\bar{C}_y^H G^{-1}\bar{C}_y P + \bar{B}_w W\bar{B}_w^H = 0 \tag{13}$$

from which the Kalman gain L follows as

$$L = P\bar{C}_y^H G^{-1} \tag{14}$$

Note that equation (13) yields the steady-state Riccati-solution, which results when the objective functional to be minimized is given by the infinite-time integral of the covariance of the estimation error. In equation (13) information about the two noise sources (system noise w and measurement noise g) enters via their respective covariances $W = E(ww^H)$ and $G = E(gg^H)$, with $E(\bullet)$ denoting the expected value. The statement that the Kalman gain L, and consequently the quality of the estimated state $\hat{x}$, crucially depends on the fidelity of the noise covariances seems self-evident. In the limit of large ratios, $G/W \gg 1$, we have either highly contaminated measurement sensors or very weak external noise levels. In this limit, we can be particularly confident in our model which is reflected in very low Kalman gains L. For small ratios, $G/W \ll 1$, we deal with very accurate measurement

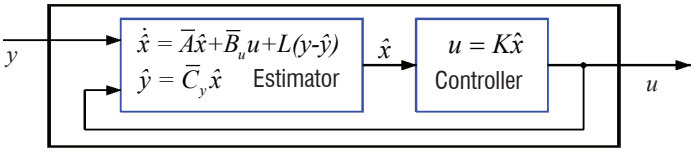

Figure 5 - Block diagram of a compensator consisting of an estimator and a controller module

sensors or a strong response of our system to external noise. In this case, frequent sensing is necessary to accurately estimate the state of the system and high Kalman gains L are common. In a final design step, the controller will be based on the state vector, given by a proportional law according to $u = K\bar{x}$. In this manner, it represents a mapping from the state vector $\bar{x}$ to the low-dimensional control signal u. The control gain K is again obtained from an optimization problem : we wish to minimize our cost objective, while observing the governing equations. When considering a cost objective involving an infinite-time integral, the solution of this optimization problem, as before, leads to the following algebraic matrix Riccati equation for an auxiliary variable Q, according to

$$\bar{A}^H Q + Q\bar{A} - Q\bar{B}_u R^{-1}\bar{B}_u^H Q + \bar{C}_z^H\bar{C}_z = 0 \tag{15}$$

from which the control gain K follows as

$$K = -R^{-1}\bar{B}_u^H Q \tag{16}$$

where R is a weight matrix appearing in the chosen quadratic form of the cost functional $J = \int_0^\infty \bar{x}^H\bar{C}_z^H\bar{C}_z\bar{x} + u^H R u\, dt$; it accounts for the cost associated with the expended control. If the cost of control is high and any control effort is heavily penalized ($R \gg 1$), the associated control gains K are very low and the control signal u is expended rather parsimoniously. In the Riccati equation (15) for the control problem, we notice the absence of noise-related terms. It thus appears that the controller is not concerned with noise sources; they are strictly dealt with by the estimator. Any performance degradation due to noise (or any failure to model noise sources accurately) must be ascribed to the estimator, not the controller. In the above analysis, we assumed a proportional control involving the true state vector $\bar{x}$; of course, under realistic conditions, this must be substituted by the estimated state $\hat{x}$. The resulting control law thus reads $u = K\hat{x}$. The separation principle assures that the control gain K is still optimal, after this substitution has been made. The combination of estimator and controller, referred to as the compensator, is sketched as a block diagram in figure 5 : it processes the measurement vector y and produces the optimal control signal u. The noise sources enter into the Kalman gain L, the cost objective is accounted for in the control gain K and the estimated (reduced) state vector $\hat{x}$ appears as an internal (hidden) variable.

The full closed-loop control problem, including true and estimated states, is finally written (in matrix form) as:

$$\begin{pmatrix}\dot{\bar{x}}\\ \dot{\hat{x}}\end{pmatrix} = \underbrace{\begin{pmatrix}\bar{A} & \bar{B}_u K\\ L\bar{C}_y & \bar{A}-L\bar{C}_y+\bar{B}_u K\end{pmatrix}}_{\bar{A}^{cl}}\begin{pmatrix}\bar{x}\\ \hat{x}\end{pmatrix} + \begin{pmatrix}\bar{B}_w w\\ Lg\end{pmatrix} \tag{17}$$

The composite block-matrix $\overline{A}^{cl}$, describing the dynamics of the controlled and estimated system, is stable by design. The performance of the compensated system (17) can be measured by the norm of the transfer functions from the noise sources (w, g) to the performance measurement signal z.

Returning to the configuration of the open cavity flow, if the noise sources are negligibly small, only very few measurements y are required to fully capture the system dynamics. In this case, the Kalman gain L can be determined in the small-gain limit. If furthermore the cost of control is considerable, the control gain K can also be computed in the small-gain limit. The performance of the resulting controller is then given by the magnitude of the transfer function of the compensated $\overline{T}_{zg}^{cl}$ system, from the measurement noise g to the performance measure z (see the black line in figure 3 b). The magnitude of the resulting control signal is obtained by inspection of the transfer function $\overline{T}_{ug}^{cl}$ of the compensated system, from the measurement noise g to the control input u (see the red line in figure 3 b).

Robustness

It is important to realize that the design of the compensator, i.e., the estimator and the controller, is based on the reduced-order model (see figure 6 a). Ultimately, however, the compensator is applied to the full system on which it must prove its effectiveness (see figure 6 b). Owing to residual degrees of freedom that have not been included in the reduced-order model, it is conceivable that problems may arise when applying a compensator to a system that it was not exactly designed for. The issue of robustness thus naturally arises in this case.

If the compensator is applied to the full system, the closed-loop transfer function from the sensor noise g to the performance measure z is written as $T_{zg}^{cl} = T_{zu}K_{uy}/\left(1-T_{yu}K_{uy}\right)$, where T_{zu} and T_{yu} are the full-plant transfer functions (see figure 3 a) and K_{uy} is the transfer function of the compensator, which is given by $K_{uy}(s) = K\left(sI-\overline{A}-\overline{B}_uK+L\overline{C}_y\right)^{-1}L$. By inspection of the transfer function T_{zg}^{cl}, the stability margins can be obtained by scrutinizing the appearance of unstable zeros in the denominator of T_{zg}^{cl}, i.e, $1-T_{yu}(s)K_{uy}(s)$. If the full-plant transfer function T_{yu} coincides with the reduced transfer function $\overline{T}_{yu}$, then by the design of K_{uy}, all zeros of $1-T_{yu}(s)K_{uy}(s)$ are stable (these zeros correspond to the eigenvalues of $\overline{A}^{cl}$). However, since $\overline{T}_{yu}$ was based on an approximation procedure, deviations from the transfer function of the original system, T_{yu}, arise. This slight mismatch may cause instabilities in the feedback loop, in other words, unstable zeros of $1-T_{yu}(s)K_{uy}(s)$. A robustness analysis encompasses a modification of the reduced-plant transfer function $\overline{T}_{yu}$, followed by an analysis of its influence on the stability of the compensated system. More specifically, we consider modified transfer functions $T_{yu} = g\overline{T}_{yu}$, with g as a complex constant and track the appearance of unstable zeros of $1-g\overline{T}_{yu}(s)K_{uy}(s)$ as a function of g. In particular, we seek critical gain margins $g = a \in \mathbb{R}^+$ and phase margins $g = e^{i\varphi}$ that render the closed-loop system unstable. Physically, the gain margin a refers to an error in the estimation of the amplification rate of an instability in the system, while the phase margin φ describes an error in the estimation of the convection speed of the perturbations.

It is noteworthy that this issue is only relevant when $T_{yu} \neq 0$, that is, in the case of a feedback configuration where the control u affects the estimation measurement y. This is the case for oscillator flows, where a global structure is destabilized and, in turn, synchronizes the entire dynamics. In the case of convective (amplifier) flows, however, with the measurement y located upstream of the actuator u, the issue of robustness is inconsequential.

For orientation we first display the eigenvalues of the original reduced-order plant, i.e., the eigenvalues of $\overline{A}$, in figure 7a by blue symbols; only positive frequencies are shown. Four eigenvalues can be identified in the unstable (gray) half-plane. Applying the compensator, the originally unstable flow will be stabilized, which manifests itself in the reflection of the unstable eigenvalues around the neutral line ($\lambda_r = 0$) into the stable half-plane. The compensated eigenvalues are depicted by red symbols. In the small-gain limit (expensive control and $G/W \rangle\rangle 1$), control efforts focus only on the stabilization of the system, while leaving the stable sub-dynamics unaffected : the stable compensated eigenvalues therefore coincide with the uncontrolled eigenvalues.

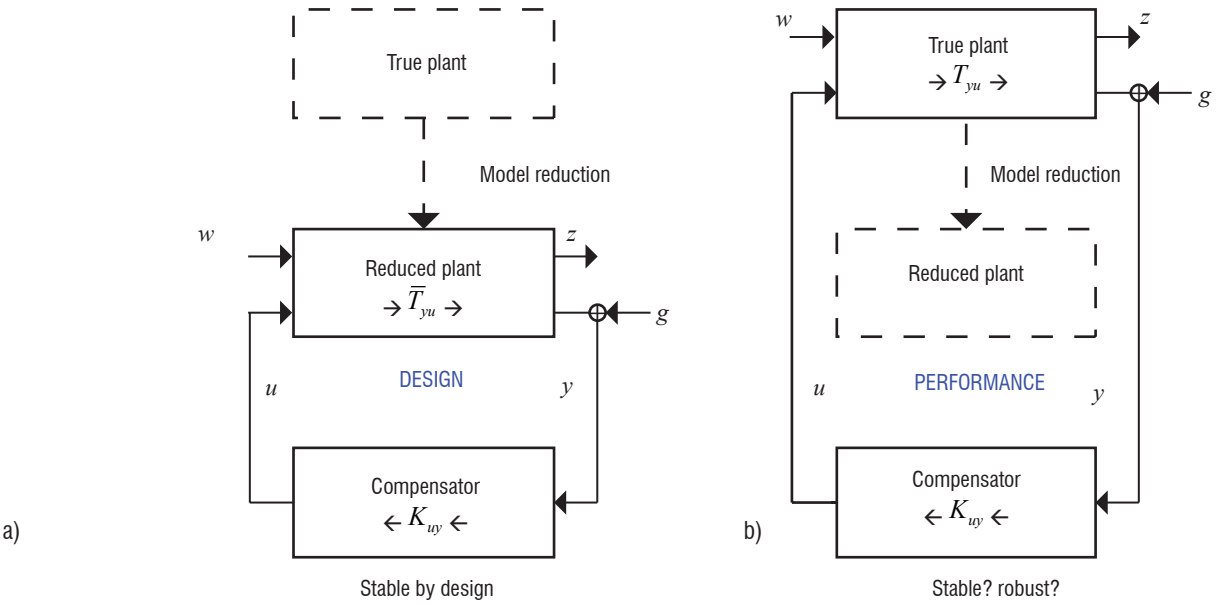

Figure 6 - a) Design of a compensator based on a reduced-order model of the true plant. b) Application of the designed compensator to the full model of the true plant

Following our robustness analysis, we will explore the movement of these eigenvalues as the parameter g is varied. We will concentrate on the gain margins of the compensated system and choose $g = a \in \mathbb{R}^+$, ranging from 0.75 to 1.5. For the case a = 1, we recover the true design configuration which, by construction, yields a stable compensated system (the case shown in figure 7a). The eigenvalues that yield the critical gain margins are indicated by the green box; we will focus on their trace through the complex plane for $0.75 \leq a \leq 1.5$. Figure 7b displays the locus of the two eigenvalues of the compensated system (i.e., the eigenvalues of $\overline{A}^{cl}$) for the specified parameter range. We observe that for a = 0.75 the two eigenvalues start out in the unstable half-plane and gradually move into the stable one. At a critical value of $a^- = 0.931$ both eigenvalues are contained in the stable half-plane, indicating stability of the compensated system. The design-state is recovered for a = 1 (green box). By further increasing a, the two eigenvalues move in opposite direction: one eigenvalue tends to higher damping rates, while the other re-crosses into the unstable half-plane at a critical value of $a^+ = 1.05$, rendering the compensated system unstable again. We can thus conclude that the stability of the compensated system can be guaranteed only for $0.931 \leq a \leq 1.05$. Obviously, this interval contains the design point a = 1; however, it shows rather narrow margin values and thus rather poor robustness of the compensator.

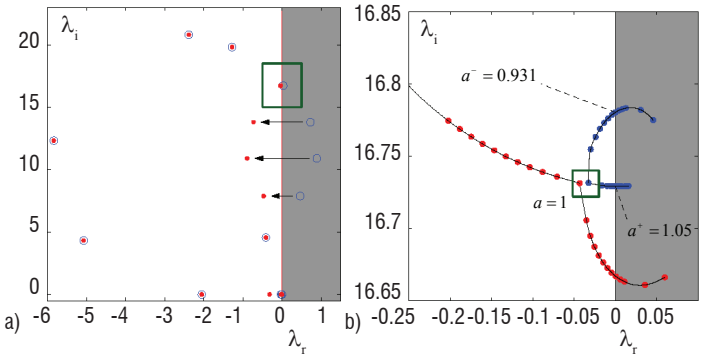

Figure 7 - a) Eigenvalues of the uncontrolled reduced-order model (blue symbols) and of the compensated system (red symbols), showing the reflection of the (four) unstable eigenvalues into the stable half-plane, as well as the invariance of the stable eigenvalues. (b) Robustness analysis: tracing the critical eigenvalues (indicated by the green box in (a)) as a function of $g = a \in \mathbb{R}^+$. The gain margins are given by $a^+ = 1.05$ and $a^- = 0.931$. For $0.931 \leq a \leq 1.05$ the compensated system is stable. The case a = 1 corresponds to the red eigenvalues inside the green box in Subfigure (a).

An analogous analysis can be conducted by choosing $g = \exp(i\varphi)$, thus exploring phase margins and robustness with respect to a mismatch in the estimated convection speed. The results are very similar and thus omitted here: only a small phase margin of $\varphi = \pm 3°$ can be tolerated before instabilities in $\overline{A}^{cl}$ are encountered.

Performance of the compensator in full-plant mode

The compensator is now evaluated in the direct numerical simulation code. Despite the narrow stability margins, we expect the compensator to perform well, since the error between the full-plant transfer function T_{yu} and the reduced-plant transfer function $\overline{T}_{yu}$, on which the compensator has been designed, is very small. The initial condition of the simulation consists of the base-flow and of the most unstable global mode with a low amplitude. The estimation sensor is contaminated by Gaussian white noise of variance $G = 1$. The simulation has been performed from time $t = 0$ to time $t = 16$. Results are shown in figures 8 a,b, where the performance measure z and the control input u are shown as a function of time t. After a short transient, the performance measure z decreases and oscillates erratically with a variance and spectrum in accordance with the performance predicted in figure 3 b. Similar comments can be made for the control input u and for the variance and spectrum of the resulting control input. At time $t = 16$, control is switched off : the signal u is set to and maintained at zero, and the performance measure z is seen to increase dramatically for some time before saturating, in agreement with the expected limit-cycle behavior.

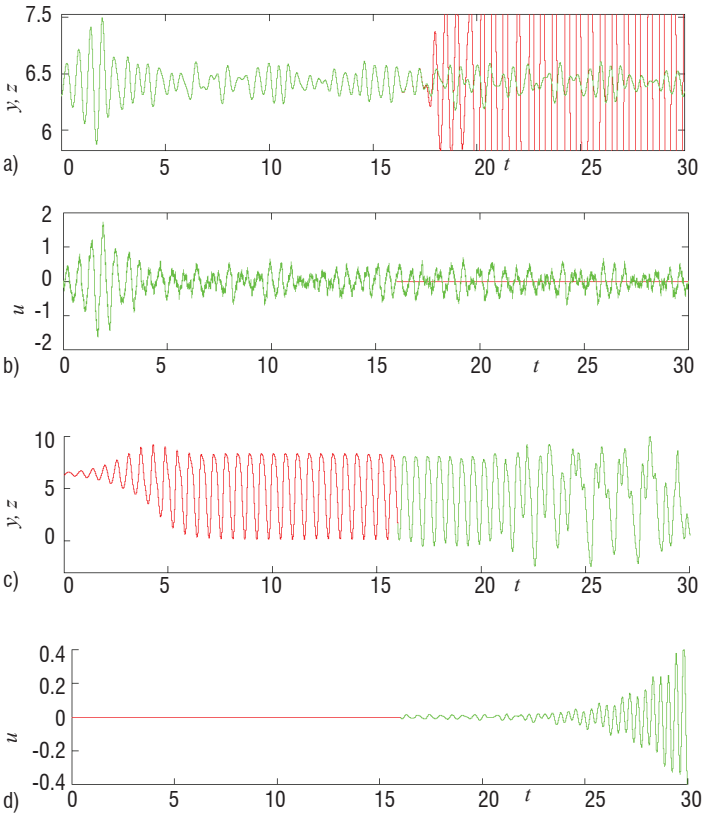

Figure 8 - Controller performance for flow over an open cavity. a) Applying control to the system dynamics near the base flow, with corresponding control signal b), c) Applying linear control to the saturated limit-cycle behavior, with corresponding control signal d).

We have shown that the designed compensator manages to maintain the flowfield in the vicinity of the equilibrium point. We will now test whether the same compensator is able to drive the system from the limit cycle towards the equilibrium point. For this, we have first performed (from $t = 0$ to $t = 16$) an uncontrolled simulation leading to the limit-cycling behavior; then, at $t = 16$, control has been switched on. Results are shown in figures 8 c,d, where the performance measure z and the control input u have been represented as a function of time t. We observe that both the performance measure and the control input display oscillations that continuously grow in amplitude, indicating that this compensator is ineffective once limit-cycle behavior is established. We note that a more robust compensator - one that may have been obtained, for example, using an $H\infty$-approach [10] - may have given better results.

The backward-facing step

As mentioned above, the configuration of the backward- facing step is a feedforward configuration, for which the control input u has only a weak influence on the estimation sensor y. Therefore, the feedback framework as described in the previous sections seems inappropriate here. It is thus futile to perform control based on a state- space model and to conduct a robustness analysis. More generally, this observation is typical of advection-dominated noise amplifiers with a sensor location upstream of the actuator. In addition, the prevalence of the noise input w and its inaccessibility in a real control experiment further rule out the use of a state-space model or Galerkin-framework, which would require precise knowledge of B_w and of the covariance W (see Bagheri et al. [4] for an example of a feedforward configuration that has been treated within a feedback framework). In essence, for any flow situations, in which noise amplification through the fluid system is the primary source of perturbation dynamics, the necessity of an accurate noise model constitutes a major design difficulty and often the weakest point of the first (state-space) approach. The next section will introduce techniques and tools to address this problem.

Approach II: input-output map and feedforward control

The first approach replicated the mapping of the system between low-dimensional signals (u and z), via a reduced state vector, by estimating the state. The alternative approach is to directly extract this mapping between the low-dimensional signals by eliminating the state. This approach is easily motivated by the formal solution to (4-5) for vanishing initial condition ($x(0) = 0$) and in the absence of noise ($w = 0$, $g = 0$). We have

$$z = \int_0^t \underbrace{C_z \exp\left(\left(t-\tau\right)A\right)B_u}_{h_{zu}(t-\tau)} u(\tau)\, d\tau \qquad (18)$$

where we recognize the low-high-low-dimensional expression $C_z \exp\left(\left(t-\tau\right)A\right)B_u$ but replace it with the low-dimensional temporal convolution kernel $h_{zu}\left(t-\tau\right)$.This expression then suggests that the measurement z can be modeled by a time-history of control signals u, appropriately weighted by the kernel function h_{zu}, which implicitly contains information about the system dynamics (given by A), the control (given by B_u) and the measurement (given by C_z). In the presence of noise sources ($\mathrm{w} \neq 0$, $g \neq 0$), modifications to the previous expression must be made; the general approach, however, persists. For convenience (but without loss of generality), we convert to a time-discrete version of (5-6) and further assume a scalar control signal u, a scalar measurement signal z and scalar noise signals w and g. We obtain

$$x_{n+1} = Ax_n + B_u u_n + B_w w_n \qquad (19)$$

$$z_n = C_z x_n \qquad (20)$$

where $x_n = x(n\Delta t)$ and similar for the other time-dependent variables. Note that the system, control and measurement matrices A, B, C are not identical (but related) to their continuous equivalents in (5-6). Pursuing the approach outlined above, we continue by expressing the discrete measurement vector z_n in general as a function of the control and noise sources, according to

$$z_n + \sum_{i=1} a_i z_{n-i} = \sum_{j=0} b_j u_{n-j} + \sum_{k=0} c_k w_{n-k} + R \qquad (21)$$

with (yet) unknown coefficients {b}, {c}. An additional dependence on the history of z has been included, in the form of an auto-regressive term with unknown coefficients {a}. By observing temporal sequences of measurement, control and noise signals, the above coefficients can be determined using linear regression techniques. Two problems arise: (i) measurements of the noise signal w are difficult to obtain and (ii) linear regression techniques assume a uniform (white) distribution for the residual error R. The first difficulty will be remedied by introducing additional sensor input upstream of the actuator that provides information about the incoming disturbance environment; this additional (scalar) sensor will be denoted by y and will be treated as a substitute for w. We thus have

$$z_n + \sum_{i=1} a_i z_{n-i} = \sum_{j=0} b_j u_{n-j} + \sum_{k=0} c_k y_{n-k} + R \qquad (22)$$

A sketch of the control configuration is shown in figure 9, together with the paths of information transfer in the flow. The influence of the control u on z is labeled by (1), while the influence of the upstream measurement y on z is indicated by (2). Even though the y-signal will capture most of the incoming noise (given by path (3) in figure 9), it is conceivable that part of the environmental noise w will not be detected by the y-sensor upstream, will pass through the fluid system and will affect the measurement z further downstream (see path (4) in figure 9). During this passage through the system, the (perhaps) white noise upstream will be modified by the system - some frequency components will be amplified, others will be damped - and will have a distinct color when it impacts the measurement z downstream. This component, at the moment represented by the residual error R, however, will be treated as white by any linear regression technique; instead, a colored noise model is needed to express R. This is accomplished by allowing a relation between time-instants of R. We thus postulate the revised model

$$z_n + \sum_{i=1} a_i z_{n-i} = \sum_{j=0} b_j u_{n-j} + \sum_{k=0} c_k y_{n-k} + \sum_{p=1} d_p R_{n-p} + R_n \qquad (23)$$

The newly introduced coefficients d_p capture a temporal correlation of the error term and thus model the color of the noise affecting the downstream sensor. The model above is known as an ARMAX model : consisting of an auto-regressive (AR) part for the measurement signal z, a moving-average (MA) part for R and two exogenous (X) inputs for u and y.

Model obtained by statistical learning

Once the structure of the model has been selected on physical grounds, time sequences of the various signals will be used as a learning data-set to determine the set of coefficients {a}, {b}, {c} and {d}. To this end, in addition to the excitations w arising from external noise sources, the system is forced by an input signal u, which is chosen as sufficiently rich in frequencies to excite the relevant time scales of the fluid system. The response to this forcing in the sensors y and z will be recorded, together with the input signal u, and will form the learning data-set. The coefficients will be determined using a least-squares technique, by matching the predicted response in z to the true measured response. Once the coefficients are known, the generality of the ARMAX- model must be verified by processing a testing data-set. This validation will ensure that input signals and/or noise environments that have not been part of the learning data-set will nevertheless yield accurate predictions of the output signal z. It is no trivial task to determine the proper complexity of the model, or the

length of the learning samples. Too short data-sequences or models displaying too high a complexity will result in rather small errors for the learning data-set, but often in large deviations for the testing data-set. This phenomenon is known as overlearning and must be avoided by striking a compromise between the error norm from the learning and testing data-set [24].

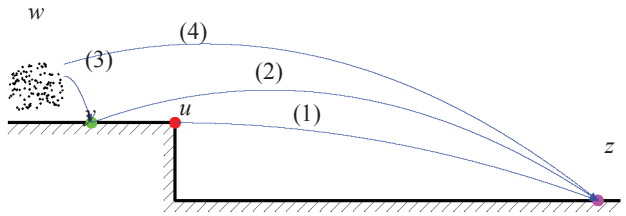

Figure 9 - Sketch indicating the transfer of information to be modeled by the ARMAX structure. Transfer of information (1) from the control u to the performance sensor z, (2) from the upstream sensor y to the performance sensor z, (3) the observable part of the disturbance environment w measured by y and (4) the part of w, unobservable by y but affecting the performance sensor z

For our case of the backward-facing step, the relevant signals y, u and z from the learning data-set are shown as a function of time in figures 10 a,b,c. We force the system in u by finite-width pulses from the actuator and record the upstream measurement y (note the rather stochastic nature of the signal), as well as the signal from the downstream sensor z. The identified model is then validated by subjecting it to a different forcing u and a different noise environment w and by comparing the model output in z to the true output from a direct numerical simulation. Comparison between true z_t and predicted z_p performance measures are shown in figure 10 d: after a transient, which corresponds to the time taken for a perturbation to travel from sensor y and actuator u to the performance sensor z, the two curves collapse indicating that the model is able to perfectly predict the performance measure z from knowledge of the estimation measure y and control input u.

Design of a compensator

Once the ARMAX model has been identified by determining the coefficients such that the predicted sensor output matches the actual one, a compensator can be designed. The underlying principle of this design is disturbance rejection, which is based on the relation $z = T_{zy}y + T_{zu}u$ linking the measurement from the upstream sensor y and the control input u to the downstream sensor z via two transfer functions. Our goal is to minimize the signal z, and, by postulating $z = 0$, we can derive the control law $u = -T_{zu}^{-1}T_{zy}y$ based on the upstream sensor measurements y. It is conceivable, and often the case, that the transfer function T_{zu} is small or zero for certain frequencies; in this case, the inversion in the expression for the control law would fail, or lead to undesirable large control gains. To avoid this issue, a pseudo-inversion (rather than an exact inversion) is taken, where a lower threshold can be specified, below which no control action is taken. This regularization step is necessary to obtain effective control action.

The control design can also be performed in the time-domain, rather than frequency domain. In this case, a time-horizon must be specified, over which the measurement signal z is minimized. Also within this setting, a regularized inversion (via a pseudo-inverse) is required to produce an operative and practical control strategy.

Performance of the compensator in full-plant mode

Following its design, the compensator is now applied to the direct numerical simulations: measurements from the upstream sensor y are taken and processed by the compensator, which provides the input signal u. This signal has been designed to optimally reduce the downstream measurement signal z. Even though not explicitly stated in the control objective, the total perturbation energy of the fluid system is expected to be reduced as well by this action. This expectation is not assured, but rather depends on the specifics of the flow configuration and must be verified for each case. Figure 11 presents results from uncontrolled and controlled simulations. A significant reduction of the perturbation energy can be observed, where control efforts lead to a lowering by about two orders of magnitude (see figure 11(a)). The spatial extent of the perturbation energy is obtained by locally time-averaging the fluctuation energy. Figure 11(b,c) juxtaposes the controlled and uncontrolled case, using identical color maps for the displayed contours. In the uncontrolled case, a substantial concentration of high perturbation energy (about 25 step heights from the step) is observed downstream of the base-flow reattachment point, which eventually decays due to the global stability of the flow. This concentration could be drastically decreased by applying flow control, as illustrated in figure 11(c). Even though the location of the perturbation-energy peak has not changed significantly, the amplitude has been reduced by two orders of magnitude, as evidenced in figure 11(a). It is noteworthy that fluctuation energy is reduced throughout the

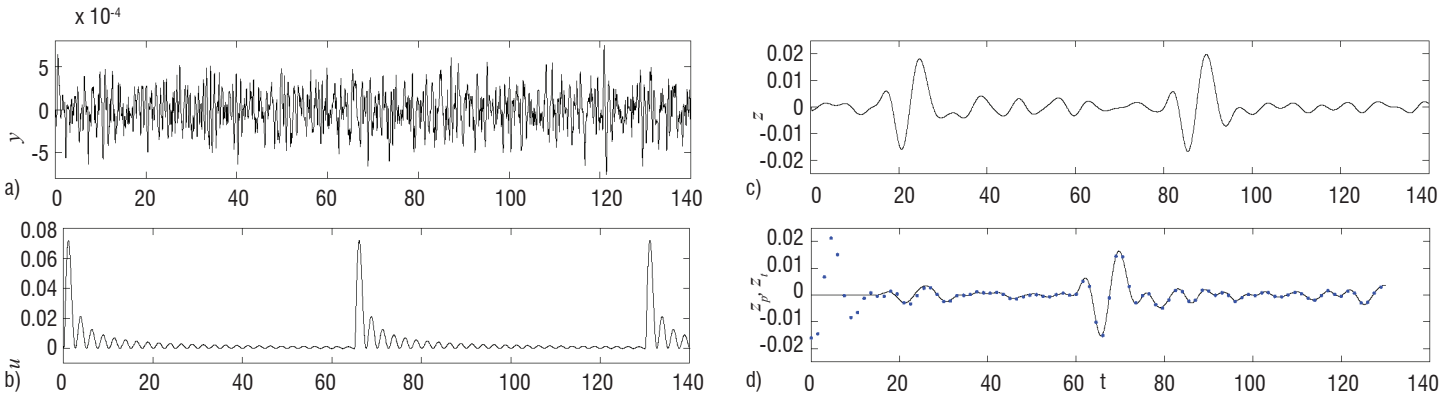

Figure 10 - Learning data-set consisting of the recorded measurements from the upstream estimation sensor y (a), input signal u (b) and downstream performance sensor z (c). Note that a pulse in u (b) yields a pulse in z (c) after a time-delay corresponding to the travelling time of a perturbation between the actuator and the performance sensor location. The validation of the model is shown in (d), where the predicted output (solid black line), for a forcing different from the learning set and for a different noise environment, is compared to the true signal (blue symbols) from the full system.

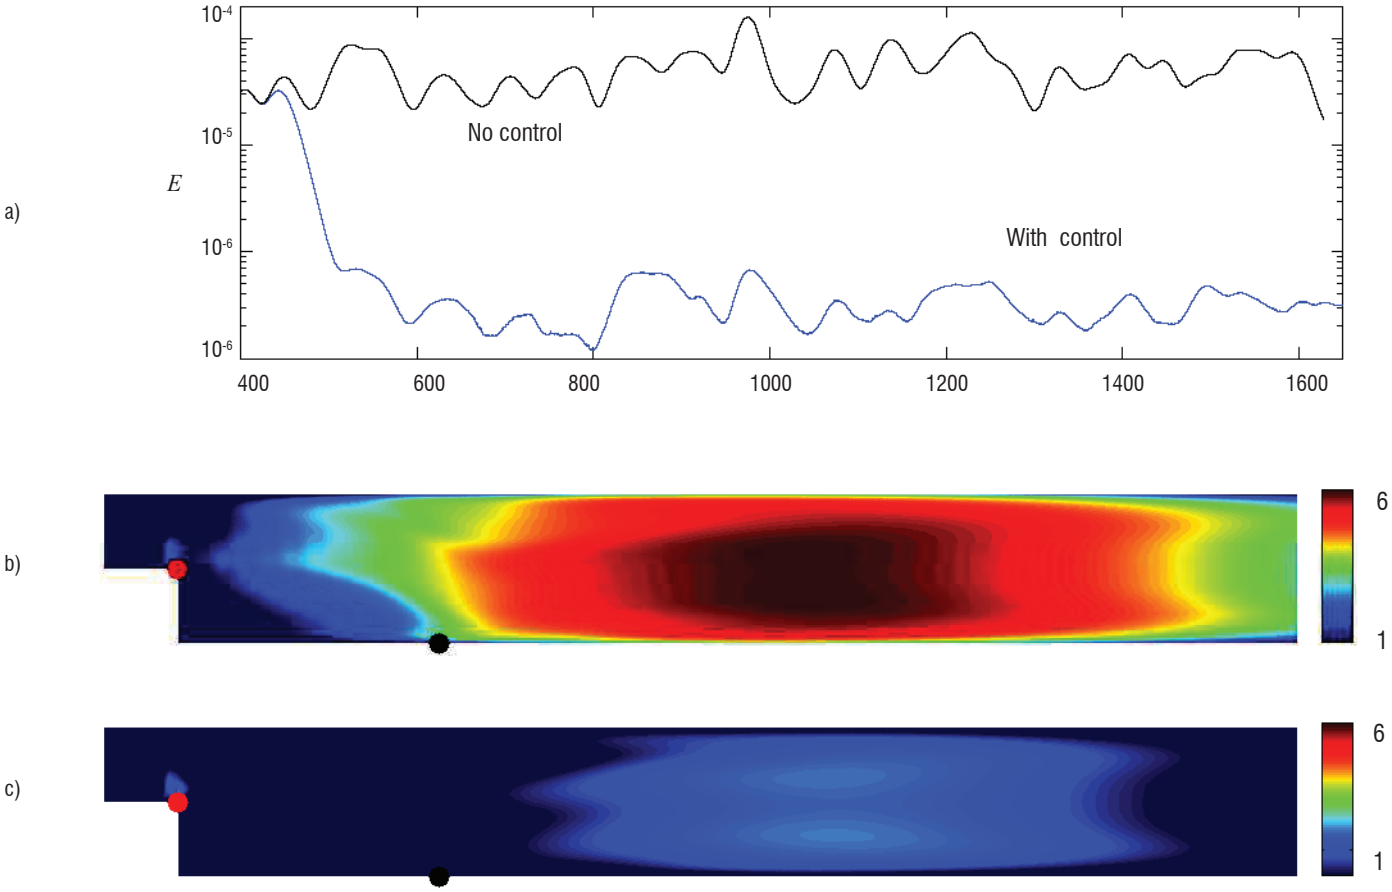

Figure 11 -(a) Global perturbation energy versus time for the uncontrolled (black) and controlled (blue) case; a reduction of two orders of magnitude is accomplished, by applying feedforward flow control. (b,c) Time-averaged local perturbation energy for the uncontrolled (b) and controlled (c) flow over a backward-facing step. The vertical coordinate has been scaled by a factor of three; the controller location (red symbol) and objective sensor location (black symbol) are also indicated. Logarithmic contour levels have been used, with identical color maps.

computational domain downstream of the reattachment point, even though only a reduction of the measurement signal from the downstream sensor (black symbol in figures 11(b,c)) has been included in the design of the compensator.

Summary and outlook

An overview of approaches and tools for the design of effective control strategies has been given. Two generic flow configurations have been used throughout this article: (i) the two-dimensional flow over an open cavity, which above a critical Reynolds number is globally unstable and thus acts as a flow-oscillator and (ii) the two-dimensional flow over a backward-facing step, which is globally stable but acts as an amplifier of ambient noise sources. Due to the prevailing global instability, the first case is rather insensitive to noise sources and can thus be accurately described by a Galerkin-based approach within a state- space formulation; in fact, the noise terms can safely be discarded in the analysis. An LQG (optimal control) approach for designing a feedback controller yields optimal performance for the compensated system, but does not guarantee robustness. Instead, a separate robustness analysis of the closed-loop transfer function must be performed, to establish stability margins. The second case, flow over a backward-facing step, responds sensitively to a noisy environment. It is thus less amenable to a Galerkin-based model, since a sufficiently accurate model of the environmental noise is difficult to come by. A data-based model is more suitable in this case and for amplifier flows in general. It involves the processing of noisy input and output signals, and the recovery of a discrete transfer function, which is subsequently used to design a feedforward controller.

In summary, a Galerkin-based approach is efficient for globally unstable flows, where an internal instability synchronizes the flow and makes it insensitive to external noise sources. Any control effort focuses on the suppression of this instability; performance is optimal, but robustness needs to be assessed a posteriori. A data-based model is suited for globally stable noise-amplifier flows. No user-supplied environmental noise model is required, since the upstream sensor acts as a substitute for the incoming noise background and a moving-average residual captures the remaining unobservable part of the upstream noise sources. This approach appears robust and efficient across a considerable range of off-design parameter values ■

Acknowledgments

The authors would like to thank Alexandre Barbagallo and Aurelien Hervé for their support and constructive criticism.

References

[1] A.C. ANTOULAS - *Approximation of Large-Scale Dynamical Systems*. Advances in Design and Control, SIAM (2005).
[2] C.W. ROWLEY - *Model Reduction for Fluids, Using Balanced Proper Orthogonal Decomposition.* Int. J. Bifur- cation Chaos, 15 (2005), 997.
[3] E. AKERVIK, J. HOEPPFNER, U. EHRENSTEIN and D.S. HENNINGSON - *Optimal Growth, Model Reduction and Control in a Separated Boundary-Layer Flow Using Global Eigenmodes*. J. Fluid Mech., 579 (2007), pp. 305–314.
[4] S. BAGHERI, L. BRANDT and D.S. HENNINGSON - *Input-Output Analysis, Model Reduction and Control of the Flat-Plate Boundary Layer.* J. Fluid Mech., 620 (2009), pp. 263–298.
[5] A. BARBAGALLO, D. SIPP, and P.J. SCHMID - *Closed-Loop Control of an Open Cavity Flow Using Reduced - Order Models*. J. Fluid Mech., 641 (2009), pp. 1-50.
[6] S. AHUJA and C.W. ROWLEY - *Feedback Control of Unstable Steady States of Flow Past a Flat Plate Using Reduced-Order Estimators*. J. Fluid Mech., 645 (2010), pp. 447-478.
[7] B.R. NOACK, K. AFANASIEV, M. MORZYN´SKI, G. TADMOR, and F. THIELE - *A Hierarchy of Low-Dimensional Models for the Transient and Post-Transient Cylinder Wake*. J. Fluid Mech., 497 (2003), pp. 335–363.
[8] E. CARABALLO, J. LITTLE, M. DEBIASI and M. SAMIMY - *Development and Implementation of an Experimental-Based Reduced-Order Model for Feedback Control of Subsonic Cavity Flows*. J. Fluids En- gin., 129 (2007), pp. 813–824.
[9] S.S. JOSHI, J.L. SPEYER and J. KIM - *A Systems Theory Approach to the Feedback Stabilization of in Finitesimal and Finite-Amplitude Disturbances in Plane Poiseuille Flow*. J. Fluid Mech., 332 (1997), pp. 157–184.
[10] J.C. DOYLE, K. GLOVER, P.P. KHARGONEKAR and B.A. FRANCIS - *State-Space Solutions to Standard H2 and H∞ Control Problems*. IEEE Trans. Automat. Control, 34(8) (1989), pp. 831–847.
[11] T.R. BEWLEY and S. LIU - *Optimal and Robust Control and Estimation of Linear Paths to Transition.* J. Fluid Mech., 365 (1998), pp. 305–349.
[12] J. KIM and T.R. BEWLEY - *A Linear Systems Approach to Flow Control*. Annu. Rev. Fluid Mech., 39 (2007), pp. 383–417.
[13] L. LJUNG - *System Identification: Theory for the User.* 2nd edition, Prentice Hall (1998).
[14] S.-C. HUANG and J. KIM - *Control and System Identification of a Separated Flow*. Phys. Fluids, 20(10) (2008), pp. 101509.
[15] A. HERVÉ, D. SIPP, and P.J. SCHMID - *A Physics-Based Approach to Flow Control Using System Identification.* J. Fluid Mech., 702 (2012), pp. 26–58.
[16] S.S. COLLIS, R.D. JOSLIN, A. SEIFERT and V. THEOFILIS - *Active Flow Control: Theory, Control, Simulation, and Experiment*. Prog. Aerosp. Sci., 40 (2004), pp. 237–289.
[17] M. PASTOOR, B.R. NOACK, R. KING and G. TADMOR - *Spatiotemporal Waveform Observers and Feedback in Shear Layer Control*. AIAA paper, 2006-1402 (2006).
[18] M. SAMIMY, J.H. KIM, J. KASTNER, I. ADAMOVICH and Y. UTKIN - *Active Control of High-Speed and High-Reynolds-Number Jets Using Plasma actuators*. J. Fluid Mech., 578 (2007), pp. 305–330.
[19] L.N. CATTAFESTA III, Q. SONG, D.R. WILLIAMS, C.W. ROWLEY, F.S. ALVI - *Active Control of Flow- Induced Cavity Oscillations*. Prog. Aerosp. Sci., 44 (2008), pp. 479-502.
[20] R. KING (Editor) - *Active Flow Control II*. Springer Berlin Heidelberg, (2010).
[21] G. DERGHAM, D. SIPP, J.-C. ROBINET, and A. BARBAGALLO - *Model Reduction for Fluids Using Frequential Snapshots*. Phys. Fluids, 23 (2011), 064101.
[22] Z. MA, S. AHUJA, and C.W. ROWLEY - *Reduced-Order Models for Control of Fluids Using the Eigensystem Realization Algorithm*. Theot. Comp. Fluid Dyn., 25, 1-4 (2011), pp. 233–247.
[23] S.J. ILLINGWORTH, A.S. MORGANS, and C.W. ROWLEY - *Feedback Control off Low Resonances Using Balanced Reduced-Order Models*. J. Sound Vib., 330 (2011), pp. 1567–1581.
[24] G. DREYFUS, J.-M. MARTINEZ, et M. SAMUELIDES - *Apprentissage statistique*. Editions Eyrolles, Paris (2011)

Acronyms

ARMAX	(Auto Regressive Moving Average with eXternal inputs)
bPOD	(Balanced Proper Orthogonal Decompostion)
ERA	(Eigenvalue Realization Algorithm)
LQG	(Linear Quadradic Gaussian)
PI	(Proportional Integral)
POD	(Proper Orthogonal Decomposition)
ROM	(Reduced Order Model)

AUTHORS

Denis Sipp has been a researcher at Onera since 2002. He is the head of the Fluid Mechanics unit in the Department of Fundamental and Experimental Aerodynamics. He obtained a PhD degree from Ecole Polytechnique in 1999, on the stability of vortex pairs. He obtained his Habilitation Degree in 2009 at the Pierre et Marie Curie University in Paris. He has been a Professor (PCC) in the Department of Mechanics at Ecole Polytechnique since 2003.

Peter Schmid is currently a research director with the French National Research Agency (Cnrs) and Professor (PCC) of Mechanics at the Ecole Polytechnique in Paris. Previously, he held a faculty position in Applied Mathematics at the University of Washington in Seattle. He received his Ph.D. in Mathematics from the Massachusetts Institute of Technology (MIT) and his Engineer's degree in Aeronautics and Astronautics from the Technical University Munich.