J. Fluid Mech. (2012), *vol.* 709, *pp.* 223–248.
doi:10.1017/jfm.2012.330

Feedback control of cavity flow oscillations using simple linear models

Simon J. Illingworth[1]†, Aimee S. Morgans[2] and Clarence W. Rowley[3]

[1] Department of Engineering, University of Cambridge, Cambridge CB2 1PZ, UK

[2] Department of Aeronautics, Imperial College London, South Kensington SW7 2AZ, UK

[3] Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

(Received 9 March 2011; revised 23 May 2012; accepted 28 June 2012;
first published online 29 August 2012)

Using data from direct numerical simulations, linear models of the compressible flow past a rectangular cavity are found. The emphasis is on forming simple models which capture the input–output behaviour of the system, and which are useful for feedback controller design. Two different approaches for finding a linear model are investigated. The first involves using input–output data of the linearized cavity flow to form a balanced, reduced-order model directly. The second approach is conceptual, and involves modelling each element of the flow physics separately using simple analytical expressions, the parameters of which are chosen based on simulation data at salient points in the cavity's computational domain. Both models are validated: first in the time domain by comparing their impulse responses to that of the full system in direct numerical simulations; and second in the frequency domain by comparing their frequency responses. Finally, the validity of both linear models is shown most clearly by using them for feedback controller design, and then applying each controller in direct numerical simulations. Both controllers completely eliminate oscillations, and demonstrate the advantages of model-based feedback controllers, even when the models upon which they are based are very simple.

Key words: aeroacoustics, control theory, instability control

1. Introduction

The oscillations induced by the compressible flow past a cavity has become a canonical problem in flow control. Some of the earliest control studies used either passive means (Heller & Bliss 1975) or open-loop actuation (Sarohia & Massier 1977). In more recent years the possibility of using feedback (or closed-loop) control for their suppression has been investigated, the idea dating back to Gharib (1987). Early studies of feedback control used simple (trial-and-error) phase-locked controllers (Shaw & Northcraft 1999; Williams, Fabris & Morrow 2000), which do not require a model of the cavity flow dynamics. These studies suppressed individual tones by up to 20 dB, but other tones were unaffected or even made worse by the feedback.

Feedback control is most effective when an accurate model of the system to be controlled is available, and this has been acknowledged for cavity oscillations in recent years, with model-based feedback control of cavity resonances being promoted

† Email address for correspondence: si250@cam.ac.uk

(Rowley *et al.* 2006). Such model-based approaches can take advantage of methods from linear control theory to design robust controllers, which can give rise to improved closed-loop performance. This could include the suppression of multiple tones; the reduction of the peak-splitting phenomenon seen in previous studies (e.g. Rowley *et al.* 2002*b*; Cabell *et al.* 2006); or the continued suppression over a wide range of operating conditions.

Clearly, to use such model-based methods, a model of the system must be available. Dynamical models of the cavity flow have typically been found using either system identification methods (Kook, Mongeau & Franchek 2002; Cabell *et al.* 2006; Kegerise, Cabell & Cattafesta 2007) or the proper orthogonal decomposition (Samimy *et al.* 2007), and standard techniques such as linear quadratic control (Cabell *et al.* 2006) and $\mathscr{H}_\infty$ control (Yan *et al.* 2006) have been used for controller design. For a recent review of the control of cavity oscillations, see Cattafesta *et al.* (2008).

The focus of this paper is on finding simple, linear models of cavity oscillations that are useful for feedback control purposes, and builds on the work of Rowley *et al.* (2006). The focus is not on capturing every aspect of the cavity flow's dynamics, but rather on capturing those dynamics that have the greatest influence on the cavity's input–output behaviour. One of the main conclusions will be that when one prioritizes capturing this input–output behaviour, very simple models suffice for feedback control design.

The paper is organized as follows. In § 2 the conceptual modelling approach proposed by Rowley *et al.* (2006) is introduced. In § 3, details of the cavity's geometry, the numerical scheme used for its solution, and a method for finding its linear dynamics are described. In § 4, the two methodologies for forming linear models are presented. In § 5, the two linear models are validated, first in the time domain using their impulse responses, and second in the frequency domain using their frequency responses. The validation of the models continues in § 6, where two model-based feedback controllers (one for each linear model) are designed and tested in direct numerical simulations.

2. Linear modelling approach

Rowley *et al.* (2006) propose a linear model for cavity oscillations. The model, composed of four constituent transfer functions, is simple and is intended for feedback control design purposes. Each of the constituent transfer functions corresponds to one of the four elements in the mechanism commonly attributed to Rossiter (1964) and also described by Powell (1953): the amplification of disturbances by the free shear layer; the scattering of the disturbances into pressure fluctuations at the downstream corner; the propagation of these pressure fluctuations inside the cavity; and the receptivity to pressure perturbations at the upstream corner. This mechanism is depicted in figure 1(*a*), and a block diagram of the linear model proposed is shown in figure 1(*b*). Here s is the Laplace variable, and the 'o' and 'L' subscripts are used for the upstream and downstream corners respectively.

In figure 1(*b*) a forcing term, f_c, has been introduced, which will be provided by our feedback controller. This forcing term enters the momentum equation (in the direction perpendicular to the free stream) like $\partial \rho v/\partial t = (\text{Navier–Stokes} + \text{forcing})$, where ρ is the fluid density and v is a velocity component which is perpendicular to the free stream. We therefore model the relationship between the forcing term and the change in ρv that it produces as an integrator. In the frequency domain, this means that $\rho v_c = f_c/s$. Summing the contributions from ρv_c and from the natural receptivity

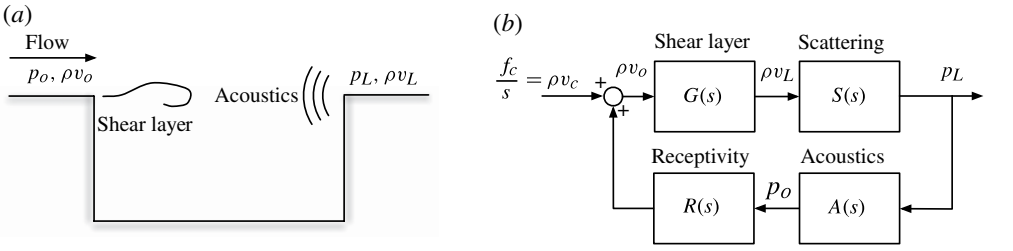

FIGURE 1. Linear model of the cavity flow: (*a*) schematic of the mechanism; and (*b*) block diagram of the linear model proposed by Rowley *et al.* (2006). '*o*' and '*L*' subscripts are used for the upstream and downstream corners, and *s* is the Laplace variable.

process then gives ρv_o (see figure 1*b*). For control purposes we are interested in the transfer function between the forcing, f_c, and the pressure at some point in the cavity. For a pressure measurement at the upstream corner (p_o in figure 1), this is given by

$$\frac{p_o(s)}{f_c(s)} = \frac{1}{s}\frac{G(s)S(s)A(s)}{1 - G(s)S(s)A(s)R(s)} = P_1(s). \tag{2.1a}$$

If we are interested in a pressure measurement at the downstream corner instead (p_L in figure 1), then we have

$$\frac{p_L(s)}{f_c(s)} = \frac{1}{s}\frac{G(s)S(s)}{1 - G(s)S(s)A(s)R(s)} = P_2(s). \tag{2.1b}$$

(It will become clear in the next section why $P_{1,2}(s)$ on the right-hand side of the two equations have been introduced.)

3. Details of the cavity flow

The two-dimensional compressible flow past a rectangular cavity is considered, shown in figure 2(*a*). The flow conditions are for a Mach number $M = 0.6$, length to depth ratio $L/D = 2.0$, a Reynolds number based on momentum thickness θ at the cavity leading edge of $Re_\theta = 56.8$, and $L/\theta = 52.8$.

Two pressure sensors are positioned in the cavity: one half-way up the upstream wall, denoted p_1, and a second located half-way up the downstream wall, denoted p_2. Actuation is provided by a forcing term, f_c, as introduced in the previous section. The actuator is positioned in the shear layer a distance $0.05L$ from the cavity's upstream corner. The distribution of the actuation is Gaussian in space, specified such that the forcing has fallen to 5 % of its maximum at a radial distance of $0.045L$ from its centre.

With the forcing, f_c, as the system input, the two transfer functions between that forcing and the two pressure measurements p_1 and p_2 are given by (2.1*a*) and (2.1*b*) respectively. (Hence the introduction of $P_{1,2}(s)$ in those equations.) This can be seen most clearly by comparing the pressures (p_o, p_L) in figure 1 to (p_1, p_2) in figure 2. (The pressures p_1 and p_2 are measured half-way along the upstream and downstream walls, whilst the linear models (2.1*a*) and (2.1*b*) are for the pressure at the upstream and downstream corners, but this should make little difference, since the Rossiter mechanism is governed by the longitudinal acoustic resonances of the cavity.)

In § 5, pressure measurements at both sensor locations will be used to validate the linear models found, both in the time domain and in the frequency domain. In § 6,

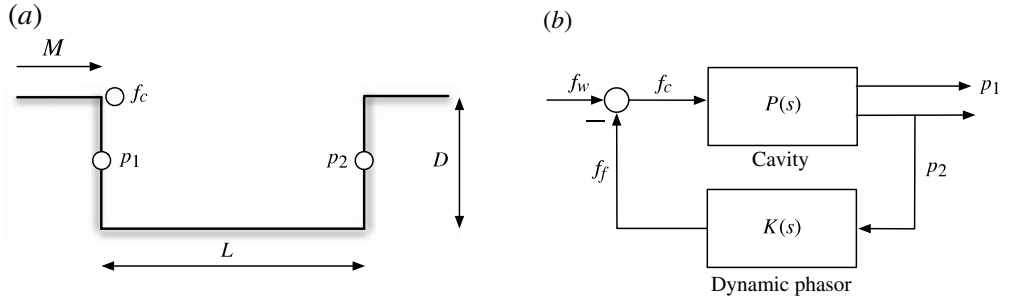

FIGURE 2. Details of the cavity flow: (*a*) cavity geometry and actuator and sensor locations; and (*b*) feedback arrangement used for acquiring input–output data of the linearized system.

the pressure measurement at probe 2 will be used as a sensor signal for model-based feedback control of the cavity.

3.1. *Numerical method*

The flow is solved by direct numerical simulation. A sixth-order-accurate compact finite-difference scheme is used for the spatial discretization (Lele 1992), and a fourth-order Runge–Kutta method performs the time integration. Measured from the cavity's upstream corner, the computational domain has its upstream boundary at $-3.9D$, its downstream boundary at $+7.6D$, and extends to $+9.2D$ in the normal direction. The simulation uses 240×96 grid points inside the cavity and 1008×384 grid points above the cavity, which is sufficient to resolve all scales at this Reynolds number. The initial condition for simulations is a (laminar) flat-plate boundary layer along the wall, which spans the cavity, and for which the (incompressible) Blasius profile and spreading rate are used. The wall is assumed to be isothermal at the same temperature as the free stream. For the inflow, outflow and normal boundaries, the one-dimensional boundary conditions of Poinsot & Lele (1992) are used, together with artificial damping terms in a buffer region (Freund 1997). The simulations have been carefully validated using grid resolution and boundary placement studies, together with comparison with experimental data (Rowley, Colonius & Basu 2002*a*).

The compressible Navier–Stokes equations are non-dimensionalized using the speed of sound, a_∞, and the cavity depth, D. In the results presented later, then, the non-dimensional pressure, p, is related to the dimensional pressure, p^d, by $p = p^d/\rho_\infty a_\infty^2$. (Times and frequencies, however, will be non-dimensionalized by the free-stream velocity, U_∞, and the cavity length, L, to make them consistent with a Strouhal number, $St = f^d L/U_\infty$.)

3.2. *Unforced behaviour*

At these conditions, the flow develops into a limit cycle, with the two most energetic peaks in the spectra occurring at Strouhal numbers of 0.40 and 0.70. (We will see in § 5.3 that each of these peaks corresponds to an unstable mode of the cavity.) These peaks are in reasonable agreement with the first two modes predicted by the semi-empirical formula of Rossiter (1964):

$$St_n = \frac{f_n^d L}{U_\infty} = \frac{n-\gamma}{M+1/\kappa}, \quad n = 1, 2, \ldots, \tag{3.1}$$

which predicts the first two modes at Strouhal numbers of 0.32 and 0.74. Here n is the mode number, and κ and γ are empirical constants corresponding to the average convection speed of disturbances in the shear layer and a phase delay. Typical values are $\gamma = 0.25$ and $1/\kappa = 1.75$, which are the values used here.

3.3. *Linearization of the cavity dynamics*

We are interested in finding linear models of cavity resonances directly from simulation data. The cavity flow under consideration is unstable, however, and this makes finding a linear model challenging for two reasons. First, the impulse response of an unstable system is unbounded, meaning that some of the standard techniques for forming a reduced-order model of a high-dimensional system (such as balanced truncation) cannot be used. Second, the growing amplitudes of the cavity's unstable Rossiter modes will ultimately give rise to nonlinear, limit-cycling behaviour, which is certainly not linear.

One solution to this problem, which can be applied in numerical studies, is to linearize the flow about an unstable steady state, and to develop a reduced-order model of the unstable flow by partitioning it into two subspaces. The first subspace describes the unstable dynamics using the unstable global modes, while the second subspace describes the flow's stable dynamics. This method has been applied to the incompressible flow past an open cavity (Barbagallo, Sipp & Schmid 2009) and to the flow past a flat plate (Ahuja & Rowley 2010).

A second solution is to acquire the model at a stable operating point which is close to the unstable case of interest. This stable system will have a linear response to (sufficiently small) forcing, and one hopes that a controller designed for this stable system will also be successful for the unstable system of interest. This approach was successfully used for the model-based feedback control of combustion oscillations by Langhorne, Dowling & Hooper (1990) and by Tierno & Doyle (1992).

A third solution, and the solution used in this paper, is to first stabilize the flow using feedback, and to measure the response to forcing of the (linear) system operating in closed-loop. This approach was used successfully by Morgans & Dowling (2007) for combustion oscillations, and is shown schematically in figure 2(*b*), where $P(s)$ is the open-loop system of interest and $K(s)$ is a feedback controller. The controller used for the linearization is the dynamic phasor model developed for the cavity flow by Rowley & Juttijudata (2005). This dynamic phasor postulates a model that captures the relevant dynamical features of the nonlinear, limit-cycling flow. It was found that this controller eliminated oscillations in direct numerical simulations, and so it is used for the present linearization of the system.

The total input, f_c in figure 2(*b*) is given by $f_c = f_w - f_f$ (negative feedback convention), where f_w is a forcing signal and f_f is the control signal from the feedback controller. The system's *open-loop* transfer function at either probe, $P_{1,2}(s)$, is then given by

$$P_{1,2}(s) = \frac{p_{1,2}(s)}{f_c(s)}. \tag{3.2}$$

Therefore despite the system's operating in closed-loop, its open-loop dynamics are recovered. (If instead the forcing signal f_w were used as the system input, the *closed-loop*, stabilized transfer function would be found.)

One must ensure that the system is responding linearly to the forcing signal, f_w, for the forcing amplitudes used. This was verified by checking that the controlled cavity

satisfied both the additive property, $y(u_1 + u_2) = y(u_1) + y(u_2)$, and the homogeneous property, $y(\alpha u) = \alpha y(u)$.

4. Identification of linear models

We now look at two ways of determining a linear model of the cavity flow from direct numerical simulations. In both cases, the cavity flow is first stabilized using the dynamic phasor described in § 3.3. The linearized system is forced with a broadband signal, f_w, at the input (see figure 2*b*), and this allows input–output data to be obtained. The input–output data can then be used to form linear models of the cavity flow. The emphasis here is on measuring linear models which are useful for feedback control design purposes.

For convenience, the linear models found from the two methods will be referred to as 'LM1' and 'LM2'. The first method, which gives LM1, involves measuring the cavity's dynamics directly from input–output data, while the second method, which gives LM2, involves finding approximations to the constituent transfer functions of the conceptual model of § 2.

4.1. *Identifying a linear model directly:* LM1

The first method involves directly measuring the cavity's linear dynamics (its transfer function) from input–output data between the actuation and the pressure measured at either of the pressure sensor locations to give directly $P_{1,2}(s)$. Since this first linear model is measured directly, it should be able to capture the input–output behaviour of the cavity very accurately, and this will be seen to be the case in § 5.1, where its impulse response matches very closely the full system. Therefore this model is intended to show what is possible when an accurate model of the system to be controlled is available, and in some ways can be used as a benchmark for the second linear model, LM2.

To form LM1 we use the eigensystem realization algorithm (ERA) (Juang & Pappa 1985), which gives a balanced, reduced-order state-space model of the system (in discrete-time) of the form:

$$x(k+1) = Ax(k) + Bu(k), \tag{4.1a}$$

$$y(k) = Cx(k). \tag{4.1b}$$

Here $u \in \mathbb{R}^p$ is the system input; $y \in \mathbb{R}^q$ is the system output; $x \in \mathbb{R}^n$ is the system state; k is the time step index; and A, B, C are suitably dimensioned matrices.

Cattafesta *et al.* (1997) and Cabell *et al.* (2006) use the ERA in their studies on feedback control of cavity oscillations. In those studies, however, the ERA is used for system identification purposes, whereas in this study we use it specifically to find a low-order model of the cavity flow.

The state-space model used is eighth-order ($n = 8$), which is adequate to capture the most pertinent input–output dynamics of the cavity, as we will see in § 5.1. The formation of reduced-order models of the cavity flow is not the focus of this paper, but for further details, see Illingworth, Morgans & Rowley (2011). For an overview of the ERA for fluid systems, see Ma, Ahuja & Rowley (2011).

4.2. *Identifying constituent transfer functions:* LM2

The second method involves using flow data at salient points in the cavity's computational domain to measure approximations to each element of the conceptual model introduced in § 2. (The excitation of the cavity is still provided by the

broadband forcing term, f_w, as it was in § 4.1.) The simple constituent models proposed by Rowley *et al.* (2006), with suitable parameters, are then compared to the measurements. An overall transfer function for the cavity is then formed by combining the constituent transfer functions according to (2.1*a*) and (2.1*b*). This second approach allows one to quantify to what extent a simple linear model can reproduce the cavity's linear dynamics (to be answered in § 5), and in particular to what extent such a model is useful for feedback control purposes (to be answered in § 6).

We now look at each of the four constituent transfer functions in turn. In each case, a simple model based on the constituent transfer function forms suggested by Rowley *et al.* (2006) is compared to a direct measurement of the transfer function using data from direct numerical simulations. Each measured constituent transfer function comes from simulation data of the forced flow using spectral analysis. This is achieved using MATLAB's `spafdr` command. This spectral approach may be likened to taking the ratio of the discrete Fourier transforms of the input and output signals, but with an important refinement: to form the frequency response at some frequency ω, it uses not just the ratio of the discrete Fourier transforms at that frequency, but also a weighted contribution from neighbouring frequencies. This weighting, often called the frequency window, provides some smoothing of the frequency response, and is founded on the principle that the true frequency response is a smooth function of ω. See Ljung (1999, chap. 6) for more details on this approach.

For each of the four measured transfer functions, a simple conceptual model, as suggested by Rowley *et al.* (2006), is fitted to the measurements by trial-and-error. (For the scattering and receptivity, we also consider the Rossiter mode eigenvalues that the transfer functions give rise to.) The parameters stated for each constituent model are non-dimensionalized by Strouhal number. For a time delay, for example, the non-dimensional value that is stated, τ, is related to the dimensional value, τ^d, by $\tau = \tau^d U_\infty / L$, where U_∞ is the free-stream velocity. The measurement locations given are all measured from the cavity's upstream corner.

4.2.1. *Shear layer*

The shear layer transfer function, $G(s)$, describes the amplification of velocity perturbations in the shear layer as they travel downstream. Since the input used for control includes the density (ρv_c in figure 1*b*), we include the density in the input and output. Referring to figure 1(*b*), the input to the shear layer is the quantity ρv_o, and so this is what we must measure. It is given by the summation of the contributions from the natural receptivity process *and* from the forcing (ρv_c). We must therefore measure at a point which is just downstream of the actuation region to ensure that its contribution is included (the actuation region ends at $x = 0.095L$). For the shear layer's output, meanwhile, we want to measure close to the downstream corner. We cannot measure exactly there, though, since then ρv would be zero. Subject to these constraints, and to the measurement data available, the input and output are measured at the locations $(x, y) = (0.1210L, -0.0136L)$ and $(x, y) = (0.9903L, -0.0136L)$, respectively.

The shear layer transfer function can be found using linear stability theory. As a simpler alternative for a linear model, Rowley *et al.* (2006) propose using a second-order system with a time delay,

$$G(s) = \frac{\rho v_L(s)}{\rho v_o(s)} = k_g \frac{\omega_o^2}{s^2 + 2\zeta\omega_o s + \omega_o^2} \mathrm{e}^{-s\tau_g}. \tag{4.2}$$

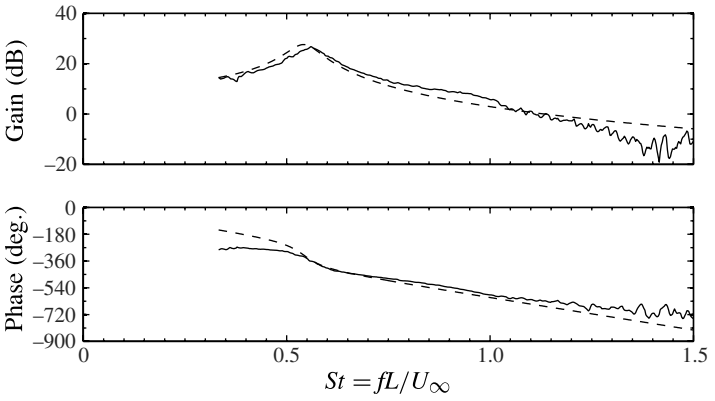

FIGURE 3. Shear-layer transfer function $G(s)$: as measured in direct numerical simulations (——); and approximated as a second-order transfer function with time delay (– –).

Figure 3 compares the measured shear-layer transfer function to that given by the linear model (4.2) for $k_g = 3.35$, $\omega_o = 3.44$ (corresponding to a Strouhal number of $3.44/2\pi = 0.547$), $\zeta = 0.07$ and $\tau_g = 1.43$ (which corresponds to $\kappa = 0.70$ in Rossiter's formula (3.1)). The simple model captures the important dynamical features of the shear layer: near-unity gain at low frequencies; amplification of disturbances at mid-frequencies; attenuation of disturbances at high frequencies; and an appropriate phase. (The phase is given by the time delay τ_g, the time taken for a disturbance generated at the upstream corner to propagate to the downstream corner.) One observes that the maximum gain, cross-over frequency and phase characteristics are all well-captured by the model. (The cross-over frequency is the frequency at which the gain passes through 0 dB.)

4.2.2. *Acoustics*

The acoustic transfer function, $A(s)$, relates pressure perturbations at the downstream corner, p_L, to pressure perturbations at the upstream corner, p_o. Since the two pressure probes are located half-way down the cavity walls, we measure the pressure at this y-location for both the input and the output. (Since the Rossiter mechanism is caused by the longitudinal acoustic cavity resonances, the measured transfer function should not be overly sensitive to the y-location chosen, and this has been verified.) The input and output for the acoustic transfer function are therefore measured at the locations $(x, y) = (L, -0.25L)$ and $(x, y) = (0, -0.25L)$, respectively.

The linear model proposed by Rowley *et al.* (2006) accounts for longitudinal acoustic cavity resonances by modelling acoustic propagation and reflection in the cavity:

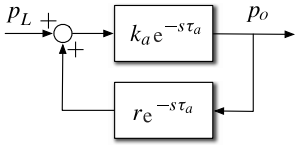

$$A(s) = \frac{p_o(s)}{p_L(s)} = \frac{k_a e^{-s\tau_a}}{1 - k_a r e^{-2s\tau_a}}. \tag{4.3}$$

An acoustic wave generated at the downstream corner propagates upstream. Some of it reflects off the upstream wall, propagates downstream, and arrives back at the downstream wall, where it is again reflected. The coefficients k_a and r measure the total efficiency of the reflection process, including losses to the far field. The time

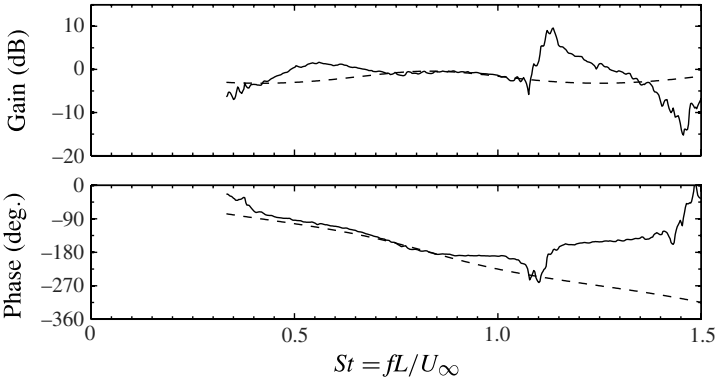

FIGURE 4. Acoustic transfer function $A(s)$: as measured in direct numerical simulations (——); and approximated as a time delay with reflection (− −).

delay τ_a represents the time taken for an acoustic wave to propagate the length of the cavity.

Figure 4 compares the measured acoustic transfer function with that given by the linear model (4.3) for $k_a = 0.8$, $r = 0.16$ and $\tau_a = 0.60$. (This acoustic time delay is the value expected and, in its non-dimensional form, is given simply by the Mach number M.) One observes reasonable agreement, with the phase characteristics better captured than the gain. It is likely that the discrepancy in the gain at lower frequencies is caused by the acoustic field generated by the shear layer's evolving vorticity. The frequency at which the greatest discrepancy is seen coincides with the frequency at which the shear layer's gain is largest, which seems to support this observation.

We briefly note the behaviour of $A(s)$ at high frequencies here. Above Strouhal numbers of approximately 1.0, the gain of $A(s)$ is poorly captured by the linear model, and the phase of the measured $A(s)$ has a positive gradient, which suggests a negative time delay (since a pure time delay τ is represented in the frequency domain by $\mathrm{e}^{-s\tau}$). We will see similar behaviour for the scattering and receptivity transfer functions: that the agreement between measurements and the conceptual model is poor above Strouhal numbers of approximately 1.0. An explanation for all of these observations will be given in § 5.5.

4.2.3. *Scattering and receptivity*

After Rowley *et al.* (2006), the scattering and receptivity transfer functions are both modelled as a constant gain and constant phase lag,

$$S(s) = k_s \mathrm{e}^{-\mathrm{i}2\pi\theta_s}, \tag{4.4}$$

$$R(s) = k_r \mathrm{e}^{-\mathrm{i}2\pi\theta_r}. \tag{4.5}$$

We will see shortly that these models provide only a crude approximation of the scattering and receptivity processes. In particular, the measured transfer functions show sizeable variations in gain and phase with frequency, meaning that the most appropriate choices of k_s, k_r, θ_s and θ_r are not obvious. A slightly different approach is therefore taken to approximating the scattering and receptivity transfer functions. To determine them we note two things. First, the eigenvalues of the two open-loop transfer functions (2.1*a*) and (2.1*b*) are given by the roots of their denominator, i.e. the roots of $1 - G(s)S(s)A(s)R(s) = 0$. Second, that with the shear-layer and acoustic

transfer functions already determined, it remains only to determine the *product* of the scattering and receptivity transfer functions, $S(s)R(s) = k_s k_r \mathrm{e}^{-\mathrm{i}2\pi(\theta_r+\theta_s)}$ to completely determine the eigenvalues of (2.1*a*) and (2.1*b*).

We therefore proceed in three stages. First, find the values of $(k_s k_r)$ and $(\theta_s + \theta_r)$ which give the best match between the eigenvalues of LM1 and LM2. Since it is the first two Rossiter modes that are unstable and that we must control, we consider the eigenvalues corresponding to these two modes when choosing $(k_s k_r)$ and $(\theta_s + \theta_r)$, and we minimize the 2-norm of the distance between the eigenvalues of LM1 and LM2 in the complex plane. Second, with $(k_s k_r)$ and $(\theta_s + \theta_r)$ determined, we must choose the individual values of k_s, k_r, θ_s and θ_r. We note that the scattering transfer function, $S(s)$ appears in the numerator of (2.1*a*) and (2.1*b*), and therefore that the individual values of k_s and θ_s also affect the two open-loop transfer functions. With the eigenvalues already fixed, we find the values of k_s and θ_s which give the best match with the directly measured transfer function, LM1. This is achieved by minimizing the $\mathscr{H}_\infty$ norm between LM1 and LM2. Third and finally, with the values of k_s, k_r, θ_s and θ_r all chosen, we compare the resulting model transfer functions with those measured in direct numerical simulations. (We could have chosen to simultaneously optimize over all four parameters k_s, k_r, θ_s and θ_r. But this requires optimization over a four-dimensional parameter space, which is much more expensive than the current approach, which optimizes over two two-dimensional parameter spaces.)

The parameter values given by this procedure are $k_s = 0.140$, $k_r = 1.821$, $\theta_s = 0.002$ and $\theta_r = 0.140$, which we will now compare to measurements from direct numerical simulations.

4.2.4. *Scattering*

The scattering transfer function, $S(s)$, describes the generation of acoustic waves when velocity perturbations impinge on the downstream corner. In keeping with the acoustic transfer function, we measure the pressure, p_L, half-way down the cavity's downstream wall at $(x, y) = (L, -0.25L)$, which serves as the output. For the input, ρv_L, we must be careful. If one measures exactly on the downstream corner, then clearly ρv will be zero. We must therefore choose a point slightly away from the surface. To ensure that the measured transfer function is not overly sensitive to the exact distance chosen, the measured scattering transfer function is plotted for two locations, which correspond to measurement locations of $(x, y) = (0.9773L, -0.0136L)$ and $(x, y) = (0.9903L, -0.0136L)$. The comparison is shown in figure 5. Whilst the model is crude, its constant gain and phase lag give reasonable agreement up to Strouhal numbers of $\sim$1.0.

4.2.5. *Receptivity*

The receptivity transfer function, $R(s)$, describes the perturbations in ρv generated by incident pressure perturbations at the upstream corner. In keeping with the acoustic transfer function we measure the input, p_o, half-way down the cavity's upstream wall at $(x, y) = (0, -0.25L)$. Referring to figure 1(*b*), notice that the output of $R(s)$ is not the same as ρv_o, since the quantity ρv_o also includes the contribution from the forcing, ρv_c. We must therefore ensure that the output of the receptivity transfer function does not include this forced contribution, and we do this by measuring at a point which is upstream of the actuation region. (The actuation region is centred at $0.05L$ from the cavity's upstream corner, and has a radius of $0.045L$.) Like the scattering transfer function, we also cannot measure exactly on the upstream corner, since then ρv will be zero. Subject to these constraints, and subject to the locations where data are available, the measured transfer functions are shown for the two locations

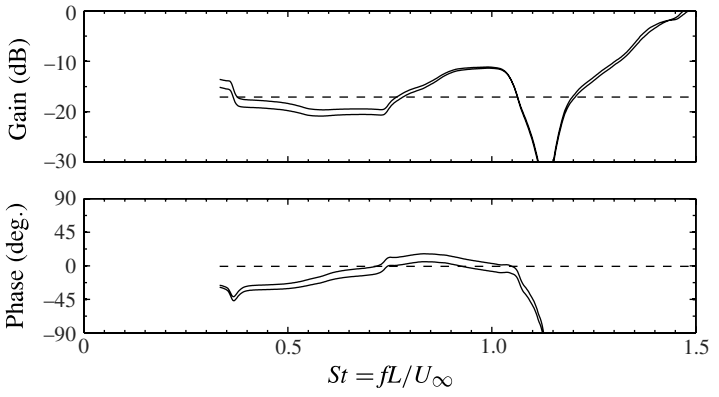

FIGURE 5. Scattering transfer function $S(s)$: as measured in direct numerical simulations at two locations (——); and approximated as a complex-valued constant (− −).

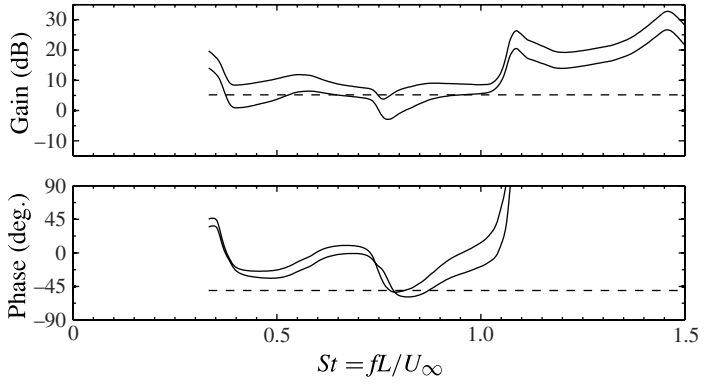

FIGURE 6. Receptivity transfer function $R(s)$: as measured in direct numerical simulations at two locations (——); and approximated as a complex-valued constant (− −).

$(x, y) = (0.0129L, -0.0495L)$ and $(x, y) = (0.0129L, -0.0135L)$. The comparison is shown in figure 6. Again, whilst crude, the model's constant gain and phase lag give reasonable agreement up to Strouhal numbers of $\sim$1.0.

5. Validation of the linear models

This section is concerned with validating the linear models found in § 4, and this is done in two ways. First, both models are validated in the time domain by comparing their impulse responses to that of the full system in direct numerical simulations (DNS). Second, the models are validated in the frequency domain by comparing their frequency responses. Staying in the frequency domain, we then look at the importance of the conceptual model's 'loop gain', before looking at how its inadequacies at higher frequencies are explained in terms of the Rossiter mechanism being 'cut-off'.

5.1. *Impulse response of* LM1

We now compare the impulse response given by LM1 to that found directly in DNS. To prevent the large gradients that would be introduced by a pure impulse, we instead

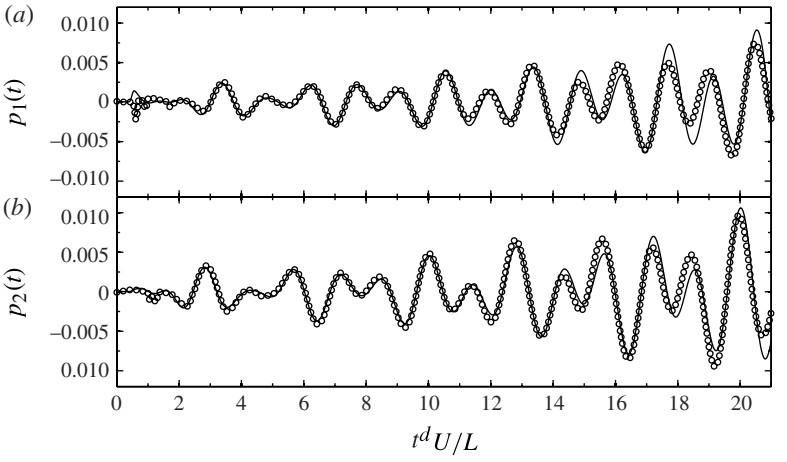

FIGURE 7. Impulse response for the pressure measurement: (*a*) at probe 1; and (*b*) at probe 2. In each plot, the impulse response of the directly measured linear model, LM1 (——) is compared to that of the full system in direct numerical simulations (○).

inject a Gaussian pulse of the form $f_c = A\exp\{-(t-t_0)^2/\sigma^2\}$, with $A = 0.5$, $t_0 = 0.484$ and $\sigma = 0.0638$. This pulse is injected at the actuator location in the unforced DNS. Similarly, that same signal acts as the input to the eighth-order ERA model, LM1, which is simulated in MATLAB.

These impulse responses for both sensor locations are compared in figure 7. The agreement is very good, and indicates that the eighth-order model captures the input–output dynamics of the cavity flow successfully. (As expected, for longer time, the two solutions diverge as the growing amplitudes give rise to nonlinear dynamics which are not captured by the linear model.) This state-space model will be used in § 6 to design the first of two feedback controllers for the cavity flow.

5.2. *Impulse response of* LM2

We now study the impulse response of LM2 using the same Gaussian pulse as in § 5.1 as its input, comparing it to that found in DNS. Now, although the linear model (2.1) given by the interconnection of the four constituent transfer functions is simple, the time delays used in its development mean that it is infinite-dimensional. To remedy this, a state-space model is formed by first finding a high-order rational fit using the `fitfrd` command in MATLAB, and then using balanced truncation to reduce the system order. The resulting state-space model is also eighth-order: this order was chosen for consistency with the state-space model of LM1, and because eight states were sufficient to capture the dominant dynamics of the model.

The impulse responses for LM2 and the DNS, again for both sensor locations, are compared in figure 8. The agreement is reasonable. As one expects, it is not as good as that given by LM1, which was found directly. This state-space model will be used in § 6 to design a second feedback controller for the cavity flow.

5.3. *Comparison in the frequency domain*

We now compare the two linear models found in §§ 4.1 and 4.2 in the frequency domain by comparing their frequency responses (transfer functions). Again, the models are compared at both sensor locations, p_1 and p_2. Strictly speaking, both LM1 and

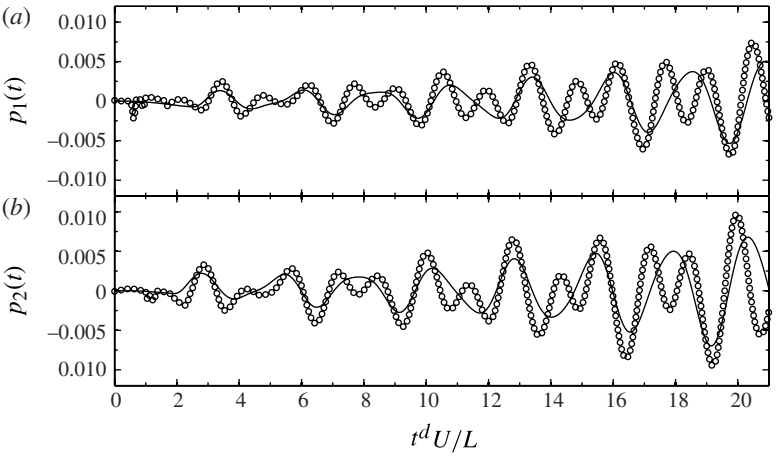

FIGURE 8. Impulse response for the pressure measurement: (*a*) at probe 1; and (*b*) at probe 2. In each plot, the impulse response of the conceptual linear model, LM2 (——) is compared to that of the full system in direct numerical simulations (○).

LM2 are approximations of the full system, but LM1 is used here to validate LM2, since LM1 was measured directly.

Being now in the frequency domain, it makes less sense to plot the (low-order) state-space models that were used for the impulse responses in figures 7 and 8. This is because although both state-space models capture the most dominant dynamics (as we saw in their impulse responses), they do so only over the frequency range of those dominant dynamics, and not at higher frequencies where their responses are small. One cannot really label this a disadvantage of the low-order state-space models, since this is exactly what they set out to do. But here we are interested in comparing the linear models over a wider frequency range. For the directly measured linear model (LM1), this means using spectral methods (previously described in § 4.2) to find its frequency response directly, which gives the complex-valued response at a discrete set of frequencies (Ljung 1999, chap 6). For the conceptual model (LM2), it means plotting the (simple but infinite-dimensional) linear model directly in the frequency domain.

The frequency domain comparison is shown in figure 9, where the two transfer functions $P_1(s)$ and $P_2(s)$ are plotted in a Bode diagram in figure 9(*a,c*), and in a Nyquist diagram in figure 9(*b,d*). (These two diagrams show the same information but in different ways: the Bode diagram plotting gain and phase as a function of frequency, and the Nyquist diagram showing the response in the complex plane.) In all four plots one observes three Rossiter modes captured by LM1 at Strouhal numbers of 0.40, 0.70 and 0.98. The increase in phase seen across the first two Rossiter modes indicates that the corresponding eigenvalues (or poles) of those modes are in the right half-plane, and therefore unstable. The phase decrease across the third Rossiter mode indicates that its corresponding poles are in the left half-plane, and therefore stable. This is shown more clearly in appendix A, where the locations of the eigenvalues of the two (state-space) linear models are compared in the complex plane. Each Rossiter mode corresponds to a complex-conjugate pair of eigenvalues, and therefore the three Rossiter modes seen in figure 9 can be represented using six states. With this in mind, appendix A also explains why eight states were needed for both state-space models.

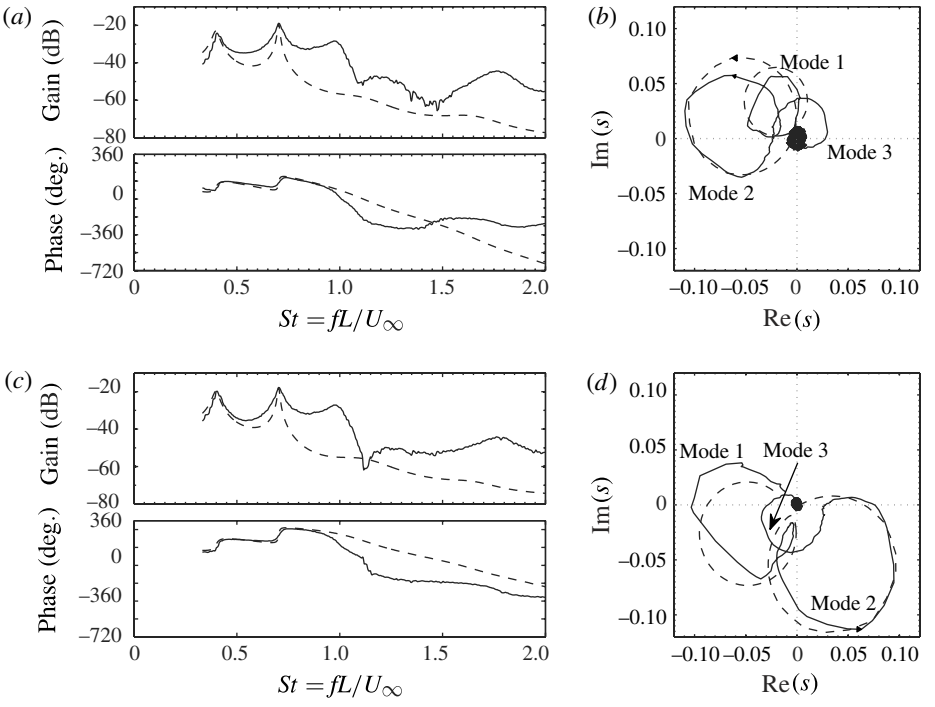

FIGURE 9. Comparison of the open-loop transfer functions: (*a,b*) for probe 1, $P_1(s)$; and (*c,d*) for probe 2, $P_2(s)$. The comparison is shown in Bode diagrams in (*a,c*) and in Nyquist diagrams in (*b,d*). In all four plots, the transfer function given by LM1 (——) is compared with that given by LM2 (– –). Each of the three Rossiter modes is indicated on each Nyquist diagram.

It is clear from all four plots that, over the frequency range of the first two Rossiter modes, the agreement between LM1 and LM2 is very good. This is encouraging, since it is the dynamics over this frequency range, where the unstable Rossiter modes are found, that drives the design of a model-based controller. It is also encouraging that the agreement is good at both probe locations. However, the third Rossiter mode is not well-captured by the conceptual model LM2. We now look at understanding these findings in terms of LM2's loop gain.

5.4. *Importance of the loop gain*

Much of the behaviour of the conceptual model LM2 – and in particular its prediction of the Rossiter modes – can be explained in terms of the *loop gain*, $L(s)$, where

$$L(s) = G(s)S(s)A(s)R(s). \tag{5.1}$$

This appears in the denominator of both transfer functions (2.1*a*) and (2.1*b*), from which it is clear that the model's poles occur at those values of s in the complex plane satisfying $L(s) = 1 + \mathrm{i}0$, since then the denominator of (2.1*a,b*) vanishes.

Let us now restrict our attention to $s = \mathrm{i}\omega$, i.e. the frequency response. This will allow us to relate the frequency response of the loop gain, $L(\mathrm{i}\omega)$, to the frequency response of the linear model, $P_{1,2}(\mathrm{i}\omega)$, which was plotted in figure 9. We can say that we will observe a 'mode' ('mode' being loosely defined here as a peak in $P_{1,2}(\mathrm{i}\omega)$) at some frequency ω when $L(\mathrm{i}\omega)$ passes near the point $1 + \mathrm{i}0$, since then the

denominator of (2.1*a*,*b*) becomes large. This in turn implies that $|L(\mathrm{i}\omega)| \approx 1$ (and that $\angle L(\mathrm{i}\omega) \approx 0°, -360°, \ldots$).

This argument can be used to explain the poor prediction of the third Rossiter mode. LM2 predicts a third Rossiter mode at a Strouhal number of 1.115 (see appendix A), which is higher than that seen in simulations. Since the loop gain decreases with increasing frequency, the over-prediction of the third Rossiter mode's frequency has a negative impact on the prediction of its loop gain as well, and so the prediction of both its frequency and its loop gain is poor: the loop gain at this Strouhal number is much smaller than 1, and therefore the third Rossiter mode is barely visible in LM2's transfer function.

5.5. *High-frequency dynamics*

The conceptual model LM2 shows reasonable success in predicting the gain and phase characteristics of the unforced cavity flow over the frequency range of the Rossiter tones, and it is the dynamics over this frequency range which dominate the cavity's input–output behaviour (as made clear by the impulse responses). Therefore LM2 is clearly useful for feedback control design purposes, and this will be seen concretely in § 6. However, LM2 performs more poorly at higher frequencies, and so first we seek an explanation for this.

The conceptual model LM2 neglects the acoustic field generated by the actuation f_c directly. At low frequencies, this is perfectly valid: the acoustics generated by the Rossiter mechanism dominate, and the effect of the actuation's direct acoustic field is negligible. At higher frequencies, though, the shear layer attenuates disturbances rather than amplifying them, i.e. $|G(s)| < 1$. This means that the velocity fluctuations arriving at the downstream corner are small, and the Rossiter mechanism is 'cut-off'. Indeed, these considerations lead to a convenient definition of low frequencies and high frequencies: low frequencies are those for which the shear layer amplifies disturbances, whilst high frequencies are those for which disturbances are attenuated. (From figure 3, the cross-over frequency of the shear layer – where the magnitude of $G(s)$ passes through 0 dB – is at a Strouhal number of $\sim$1.1.)

This explains the positive gradient of the acoustic transfer function's phase seen in figure 4, which suggested a negative time delay. At frequencies above the shear layer's cross-over frequency, the actuation dominates the acoustic field. Since the actuator is closer to the upstream corner than to the downstream corner, its effect is seen at p_o before it is seen at p_L, and this corresponds to $A(s) = p_o(s)/p_L(s)$ having a negative time delay.

The acoustic transfer function makes clear what is true for all four constituent transfer functions: it is only valid to seek them over the frequency range at which the Rossiter modes can occur, i.e. at frequencies where $|G(s)| \gtrsim 1$. To do so at higher frequencies is to seek a mechanism that no longer exists, the open-loop dynamics being instead dominated by the acoustics directly generated by the actuation f_c. This helps to explain the poor agreement seen at higher frequencies between measurements and LM2 for the transfer functions representing the acoustics $A(s)$, the scattering $S(s)$ and the receptivity process $R(s)$.

6. Model-based feedback control

The focus of this paper is on finding simple, linear models of cavity oscillations which are useful for feedback control design. With this in mind, the utility of the two

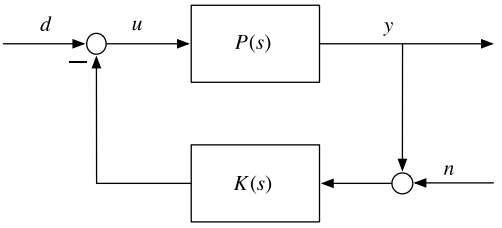

FIGURE 10. General closed-loop control arrangement. The effect of noise, n, and disturbances, d, is included.

linear models is now tested by using them to design model-based feedback controllers. First, the specific type of feedback controller used, the linear quadratic Gaussian (LQG) regulator, is described. Then closed-loop control results using two different LQG regulators – one for each linear model – are presented.

6.1. *The LQG regulator*

We now look at the specific choice of controller that will be used for feedback control of the cavity – the linear quadratic Gaussian (LQG) regulator. A closed-loop control arrangement for some general plant $P(s)$ is shown in figure 10. Here $K(s)$ is the closed-loop regulator to be designed. The effect of noise (n) and disturbances (d) is included. For the discrete-time state-space model defined by (4.1), a linear quadratic (LQ) regulator uses the state feedback law

$$u = -\Gamma x, \tag{6.1}$$

where Γ is a matrix of control gains, to minimize the quadratic cost function

$$J = \sum_{i=0}^{\infty} (x^{\mathrm{T}} Q_x x + u^{\mathrm{T}} Q_u u). \tag{6.2}$$

Here $Q_x \in \mathbb{R}^{n\times n}$ and $Q_u \in \mathbb{R}^{p\times p}$ are weighting matrices used to penalize large system states and large control inputs respectively. Substituting (6.1) into (4.1), we obtain the closed-loop system

$$x(k+1) = (A - B\Gamma)x(k), \tag{6.3a}$$

$$y(k) = Cx(k), \tag{6.3b}$$

and so by suitable choice of the feedback gain matrix Γ, the eigenvalues of $(A - B\Gamma)$ – and therefore the dynamics of the closed-loop system – can be made stable.

Linear Quadratic (LQ) control assumes that the full state x is available to the regulator – this state is used in the feedback law (6.1). In many cases, though, the full state is not available for measurement, the regulator having direct access only to the inputs u and the outputs y. Linear quadratic *Gaussian* (LQG) control remedies this by using an observer to form an estimate $\hat{x}$ of the real state x. This estimate of the state is used in the LQG control feedback law

$$u = -\Gamma \hat{x}. \tag{6.4}$$

It remains to see how the estimate of the state $\hat{x}$ is formed. The Gaussian term in linear quadratic Gaussian control refers to the type of observer used, the Kalman filter. A Kalman filter amounts to a specific choice of the observer which is optimal, in the

sense that the error converges in the presence of stochastic disturbances d and noise n (see figure 10), assumed to be zero-mean, Gaussian, white-noise processes.

The LQG and optimal estimation problems can be solved using standard routines in MATLAB. The resulting LQG controller and observer can then be combined to give an overall regulator $K(s)$, which is of the same order as the plant model used $P(s)$.

It should be noted that although the LQ control problem has guaranteed stability margins, there are no such guaranteed margins for the LQG control problem (Doyle 1978), and therefore the stability margins must be checked for a given control design. In practice, though, LQG control often gives very satisfactory results, and this will be seen to be the case for the cavity flow.

6.2. *LQG regulator design*

Having introduced the LQG regulator in the previous section, this section outlines some specific details concerning the regulator design for the cavity flow. Two LQG regulators are formed: one for each of the two linear models, LM1 and LM2. For convenience, the regulator based on LM1 will be referred to as 'LQG1', while the regulator based on LM2 will be referred to as 'LQG2'. (See appendix B for further details on these regulators, and their relation to the model orders chosen for LM1 and LM2.)

The pressure at probe 2, p_2 is used for feedback control. Therefore the cavity has a single input, f_c, and a single output, p_2. The two models of the cavity flow found in § 4 do not include an actuator transfer function. The actuator dynamics can be important for cavity oscillations, especially when the actuator has limited bandwidth, which makes closed-loop amplification of disturbances over some parts of its bandwidth inevitable (Rowley *et al.* 2006). Therefore an actuator is included here in the regulator design, and will be included in direct numerical simulations in the next section. The actuator transfer function used is a second-order, under-damped system:

$$k_L \frac{\omega_L^2}{s^2 + 2\zeta_L \omega_L s + \omega_L^2}, \tag{6.5}$$

with $k_L = 1.0$, $\omega_L = 6.28$ (which corresponds to a non-dimensional frequency of $f_L = 1.0$) and $\zeta_L = 0.3$. See Kegerise *et al.* (2007) for example, where the frequency response of the piezoelectric actuator used is of this type. The LQG regulator is then designed for the system including this actuator. (The actuator's frequency response is shown in appendix B.)

The cavity flow has a single input (f_c), and so Q_u in (6.2) is simply a scalar. Choosing the output measurement $y = Cx$ $(= p_2)$ as the output cost (giving $Q_x = C^{\mathrm{T}}C$), the quadratic cost function has the simple form

$$J = \sum_{i=0}^{\infty} (p_2^2 + q_u f_c^2), \tag{6.6}$$

where the scalar q_u is used to specify the relative importance of maintaining a small input signal f_c and maintaining a small output measurement p_2. (The same value of q_u is used for both regulators, LQG1 and LQG2.)

Very little noise is encountered in the simulations, but one expects more in a real experiment, so Gaussian white noise is artificially added to the pressure measurement p_2. The covariance of this noise is used in the design of the observer.

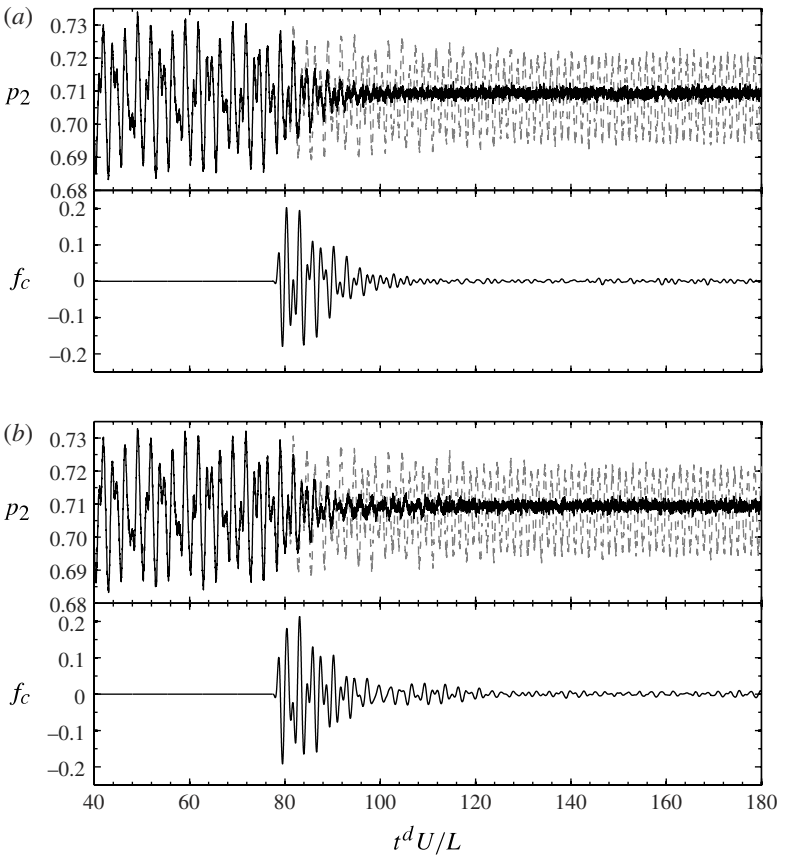

FIGURE 11. Closed-loop control of the cavity flow: (*a*) using LQG1; and (*b*) using LQG2. In both figures, the pressure measurement (top) and the control signal (bottom) are shown. For the pressure measurements, the closed-loop case (——) is compared to the unforced case (– –).

6.3. *Results*

In this section the two regulators LQG1 and LQG2 are applied to the full system in direct numerical simulations.

Closed-loop control results for the two regulators are shown in figure 11. The regulators have been designed using the directly measured linear model (LQG1) in figure 11(*a*), and using the conceptual linear model (LQG2) in figure 11(*b*). Encouraging results are seen: in both cases the regulator completely eliminates oscillations. It is not surprising that the LQG1 regulator eliminates oscillations slightly more quickly than the LQG2 regulator, or that it does so with a slightly smaller control input, when one considers the more accurate model (LM1) of the cavity dynamics used in its design. Nevertheless, the LQG2 regulator performs well, and compares well with the LQG1 regulator.

The closed-loop behaviour can also be studied by looking at the time-delay phase portrait of the pressure signal, which can help to distinguish between a stable, noise-driven regime and a limit-cycling regime (Mezić & Banaszuk 2004). The phase portrait of a limit-cycling system will be a closed curve (and, with noise, a fuzzy

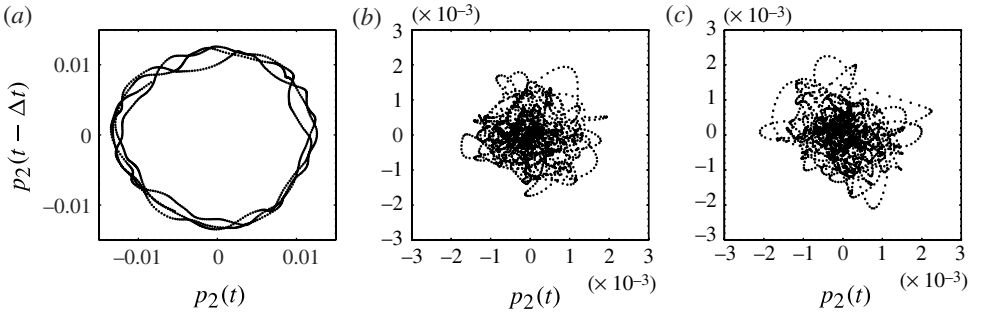

FIGURE 12. Phase portraits of the pressure measurement, p_2: (*a*) without control; (*b*) for LQG1; and (*c*) for LQG2. Note the larger scale used for the unforced data in (*a*).

curve). The phase portrait of a stable system, forced by noise, will be concentrated about a point. Figure 12 shows the time-delay phase portraits for the unforced system, for LQG1 and for LQG2. The pressure data have first been passed through a narrowband filter at the frequency where the unforced flow has its greatest power spectral density, which is at a Strouhal number of $St = 0.7$ (the second Rossiter mode). The time delay Δt used in the phase portraits is chosen to be one quarter of the period corresponding to this frequency (i.e. $\Delta t = 1/(4 \times 0.7)$). For the closed-loop cases in figure 12 (*b*,c), the pressure data are plotted for a short window of time after oscillations have been completely eliminated. The plots help confirm the closed-loop stability brought by both regulators, figure 12(*a*) being characteristic of a limit-cycling regime, and figures 12(*b*) and 12(*c*) both characteristic of stable, noise-driven systems. Comparing figures 12(*b*) and 12(*c*), what is perhaps surprising is that the responses of the two closed-loop systems to noise appear very similar in amplitude, even though LQG2 has been designed using a less accurate model of the cavity.

We will now see that regulator LQG1 does in fact provide better closed-loop attenuation than LQG2 by comparing the two regulators' closed-loop spectra. We will also compare the two LQG regulators to the original dynamic phasor compensator, again using their spectra to perform the comparison.

6.4. *Closed-loop spectra and comparison with the dynamic phasor compensator*

Figure 13 shows the power spectral density (PSD) of the non-dimensional pressure (in dB, relative to unity) for the unforced and closed-loop cavities. It is clear that both regulators achieve significant reductions in the spectra when compared to the unforced cavity. As expected, LQG1 achieves greater attenuation than LQG2. Figure 13 also shows the spectrum for the original dynamic phasor controller. We see that, as we might expect, LQG1 achieves greater attenuation than the dynamic phasor. The attenuation achieved by LQG2, meanwhile, lies somewhere between the two.

These observations are made quantitative in table 1, which supplements figure 13 by showing the peak in the power spectral densities at each of the three Rossiter modes, and does so for each of the three compensators. From table 1 we see that LQG1 provides the best attenuation across all three Rossiter modes, and we can attribute this to the quality of the model used to synthesize LQG1. Consistent with the less accurate model used to synthesize it, LQG2 performs less well than LQG1 for all three Rossiter

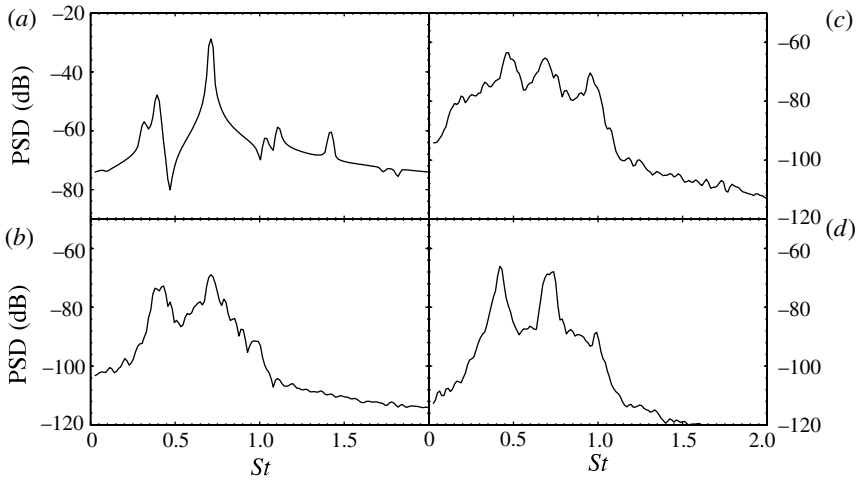

FIGURE 13. Power spectra of the non-dimensional pressure signal (in dB, relative to unity): (*a*) without control; and closed-loop spectra for (*b*) the dynamic phasor; (*c*) LQG1; and (*d*) LQG2. Notice the larger scale used for the unforced spectrum in (*a*).

Compensator	Rossiter 1	Rossiter 2	Rossiter 3
Phasor	−63.5	−65.3	−70.4
LQG1	−72.8	−68.9	−91.4
LQG2	−66.1	−67.9	−88.5

TABLE 1. Comparison of the power spectral densities (in dB, relative to unity) of each Rossiter mode for each of the three compensators.

modes but, encouragingly, still provides better attenuation than the dynamic phasor, again for all Rossiter modes.

In this study we have focused attention on control at the design Mach number of $M = 0.6$, but an additional benefit of LQG1 is its greater robustness to changes in the free-stream Mach number. Using the dynamic phasor at off-design Mach numbers, oscillations were reduced but not stabilized at $M = 0.65$, and at $M = 0.55$ oscillations were made worse by control (Rowley & Juttijudata 2005). For LQG1, on the other hand, closed-loop stability is still provided at Mach numbers in the range $0.4 \leqslant M \leqslant 0.7$ (Illingworth *et al.* 2011). The off-design performance of LQG2, meanwhile, sits somewhere between the two. Applying LQG2 at a Mach number of $M = 0.50$, oscillations are completely suppressed, while at $M = 0.70$ the limit cycle persists, but is reduced in amplitude by control by approximately 30 % (3.1 dB). Notice that this off-design performance of LQG2 is for the regulator designed for $M = 0.60$. An alternative approach would be to anticipate the cavity's dynamics at a different Mach number by using the constituent transfer functions of LM2 to form a new linear model, using which a new regulator could be formed, and this would be an interesting area for further work.

6.5. *Discussion*

The system under consideration is not an easy one to control: the cavity flow is governed by high-order, nonlinear dynamics; there are two unstable modes that must be stabilized simultaneously; there is an appreciable time delay; and the controller must contend with a limited-bandwidth actuator, which makes closed-loop amplification of disturbances over some parts of its bandwidth inevitable (Rowley *et al.* 2006). Considering these challenges it is quite remarkable that the regulators – designed using simple, low-order, linear models of the cavity's dynamics – perform so well.

At this point we may ask what effect the eighth-order models used for LM1 and LM2 have on the closed-loop control results. In particular, how sensitive are the results to the model order, and would we see any improvement if higher-order models were used? These questions are answered in appendix B, where the LQG regulators produced by the low-order models are compared in the frequency domain to those given by higher-order representations of LM1 and LM2. It is shown that the regulators produced by the eighth-order and higher-order models are almost identical, and this tells us that the cavity's dynamics that are important for control purposes have been captured by the eighth-order models. (An alternative approach would have been to form LQG1 and LQG2 based on the high-order models of appendix B, and then perform model reduction on the regulators themselves. The results of appendix B, however, tell us that these regulators would be almost identical to the regulators we have used.) This means that, should we choose to use these higher-order models as our starting point for the regulator design, the closed-loop performance that we would achieve would be the same as that already seen, and this is borne out in practice: closed-loop simulations have been performed using the higher-order regulators shown in appendix B, and the closed-loop spectra that they produce are indistinguishable from those shown in figure 13.

We have seen that the feedback controller LQG2 performs well despite the inaccuracies in the model (LM2) used to design it, and therefore that the controller exhibits a certain level of robustness to model uncertainty. We have not tried to quantify this robustness, but this could be achieved using the ν-gap metric (Vinnicombe 2000), which is the most useful measure of distance between dynamical systems when one is concerned specifically with applying feedback to those systems. The ν-gap metric fits very naturally into an $\mathscr{H}_\infty$ loop-shaping control design, and this would be an interesting area for future work.

7. Conclusions

This paper has focused on forming linear models of cavity resonances that are useful for feedback control purposes. Two linear models have been found using two quite different methodologies: the first using input–output data of the cavity flow directly, and the second using a simple conceptual model to represent each element of the flow physics separately, as suggested by Rowley *et al.* (2006). Both models are simple and can be represented using few degrees of freedom (or equivalently, states).

The linear models have been validated: first in the time domain by comparing their impulse responses to that of the full system in direct numerical simulations; and second in the frequency domain by comparing their frequency responses. The conceptual linear model does not capture the high-frequency dynamics of the cavity, owing to the Rossiter mechanism upon which it is based being 'cut-off'. Nevertheless,

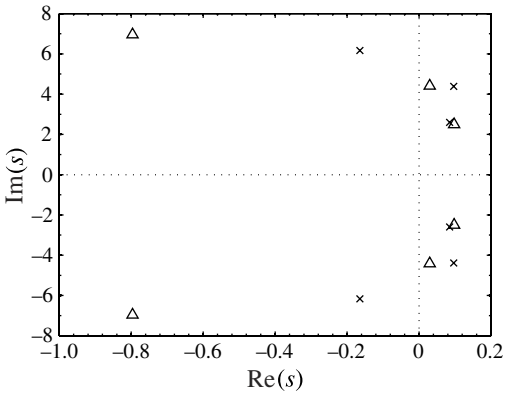

FIGURE 14. Comparison of the system eigenvalues (or poles) corresponding to the first three Rossiter modes for LM1 (×) and LM2 (△).

by modelling only the dominant dynamics, it captures the important input–output characteristics of the cavity reasonably well.

Finally, the utility of these linear models has been demonstrated most clearly by using them to design model-based feedback controllers, and applying these controllers in direct numerical simulations. Both controllers completely eliminate oscillations, and their success suggests two things. First, that linear-model-based control can be very effective in eliminating cavity oscillations; and second, that the linear models used can be very simple when one concentrates on modelling those dynamics which are most important for feedback control.

Acknowledgements

S.J.I. gratefully acknowledges financial support from the Engineering and Physical Sciences Research Council (EPSRC) and Rolls-Royce plc. A.S.M. gratefully acknowledges the Royal Academy of Engineering and the EPSRC, who supported her as a Research Fellow throughout this work.

Appendix A. Rossiter mode eigenvalues

Figure 14 shows the eigenvalues (or poles) corresponding to the three Rossiter modes for the state-space models of LM1 and LM2. Each Rossiter mode gives a complex-conjugate pair of poles, so there are six plotted altogether. The poles of the first two Rossiter modes (for both models) are in the right half-plane, $\mathrm{Re}(s) > 0$, and are therefore unstable. The poles of the third Rossiter mode are in the left half-plane, $\mathrm{Re}(s) < 0$, and are therefore stable.

The agreement between the two linear models for the first two Rossiter modes is good, and this goes some way to explaining the success of LM2 for feedback controller design. The agreement for the third Rossiter mode is poor though, with the poles of LM2 for this mode lying much deeper in the left half-plane, which corresponds to that mode being more heavily damped. This explains why the third Rossiter mode is barely visible in LM2's frequency responses in figure 9.

As noted in § 5.3, each Rossiter mode corresponds to a complex-conjugate pair of eigenvalues, and therefore the three Rossiter modes seen in figure 9 can be represented using six states. One may now ask why eight states were needed for the two linear

models, then? This is most easily answered by looking at LM1, whose state-space model came from the eigensystem realization algorithm (ERA), although similar arguments apply to LM2. The ERA is a form of approximate balanced truncation, and balanced truncation has a very useful property: for a given model order n, it guarantees a minimum level of 'goodness' of the reduced-order model. That is, an upper bound for the error in the reduced-order model can be derived. Let P_n denote the transfer function of the reduced-order model of order n, and let P_r denote the transfer function of the full system, of order r. Then the $\mathscr{H}_\infty$-norm of the error satisfies

$$\|P_r - P_n\|_\infty \leqslant 2 \sum_{j=n+1}^{r} \sigma_j \tag{A 1}$$

('twice the sum of the tails'), where σ_j are the Hankel singular values of P_r.

Balanced truncation is not optimal though: for a given order n, there is no guarantee that the model is as good as is possible. It is difficult to know whether the two extra states in the state-space models represent some 'lumped' approximation to some other important dynamics of the cavity, or if they are simply a by-product of the non-optimality of the method, but in any case it is not surprising (or troublesome) that eight states were needed. In fact, it is not at all surprising that more than six states were required when one considers that, as well as the Rossiter modes that one expects to recover, there are also time delays introduced by the convection of disturbances in the shear layer and by acoustic propagation times (see §§ 4.2.1 and 4.2.2). It is crucial that these time delays are captured by the state-space models, and therefore one can say that suitable low-dimensional approximations of these delay dynamics are 'embedded' in the state-space models for LM1 and LM2.

Appendix B. Model order and its relation to LQG regulator design

In this appendix we provide further details relating to our choice of eighth-order models for both LM1 and LM2, together with details of the LQG regulators that those models produce. We do so for two reasons. First, some consideration of the regulators produced by the LQG design procedure is useful in its own right. In particular, the LQG regulator problem is posed in the time domain, but an understanding of the regulator is best gained by considering its behaviour in the frequency domain using its transfer function. Second, by comparing the regulators produced by the low-order representations of LM1 and LM2 to those produced by higher-order representations, we provide justification for our choice of those low-order models.

B.1. *The actuator transfer function*

The actuator transfer function (gain only) given by (6.5) is shown in figure 15. Notice that the LQG regulators presented next are designed for the cavity including this actuator.

B.2. *The overall open-loop plants and their LQG regulators*

Recall that, for control, we have a single input and a single output, the output given by the pressure at probe 2. We therefore only consider this second probe location in this section. We do this first for LM1, and then for LM2.

B.2.1. *Linear model LM1*

Figure 16(*a*) compares the transfer function for the eighth-order model of LM1, which was generated by the ERA, to a hundredth-order ($n = 100$) model found using

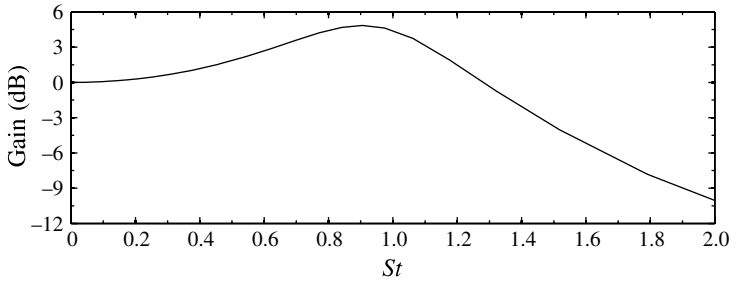

FIGURE 15. Actuator transfer function (only gain shown) given by (6.5) with $k_L = 1.0$, $\omega_L = 6.28$ and $\zeta_L = 0.3$.

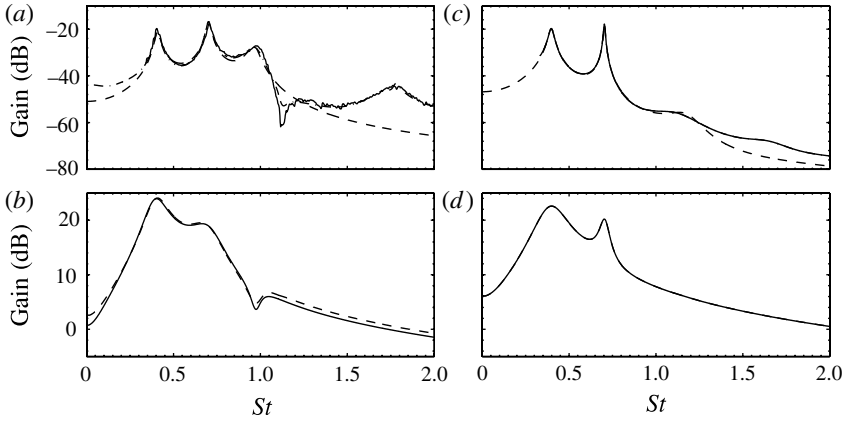

FIGURE 16. Different representations of LM1 and LM2, and their corresponding LQG regulators. In (*a*), the spectral estimate of LM1 (——) is compared to eighth-order (− −) and hundredth-order (− · −) models given by the ERA. In (*b*), the LQG regulators given by the eighth-order (——) and hundredth-order (− −) ERA models are compared. A similar analysis is performed in (*c,d*) for LM2: the original conceptual model (——) is compared to its eighth-order (− −) and thirtieth-order (− · −) state-space approximations in (*c*); and the LQG regulators given by the eighth-order (——) and thirtieth-order (− −) models are compared in (*d*).

the same ERA procedure. Both of these models are compared to the spectral model that was previously plotted in figure 9(*a*). We see that the agreement between the hundredth-order model and the spectral model is excellent, while the eighth-order model captures only the dominant dynamics – that is, the three Rossiter modes.

Figure 16(*b*) shows the LQG regulators that are produced using the eighth-order and hundredth-order representations of LM1. An important point here is that, despite the two models being of vastly different orders, the two regulators which they give rise to are almost identical. This tells us that the control objectives, which are encapsulated in the regulator, remain the same across the two models, and therefore that the eighth-order model, although unable to capture all of the cavity's dynamics, still captures the dynamics which are important for control.

For both regulators, the response is largest at the frequencies of the two unstable Rossiter modes, where the benefits of feedback (i.e. for closed-loop stability and disturbance rejection) are sought.

B.2.2. *Linear model LM2*

Figure 16(*c*) shows the transfer function for the conceptual model (2.1*b*) which, we recall, is given by the interconnection of the four constituent transfer functions (4.2)–(4.5), and which we saw previously in figure 9(*c*). This is compared to its thirtieth-order state-space approximation, found using the `fitfrd` command in MATLAB. By performing balanced truncation on this thirtieth-order model, the eighth-order representation of LM2 – which is then used for the regulator design – is found, which is also plotted. We see that the original conceptual model and its thirtieth-order approximation are indistinguishable, while there are some small differences introduced by the eighth-order approximation.

Despite this, the control objectives are still captured by the eighth-order approximation of LM2. This is demonstrated in figure 16(*d*), where the regulators given by the thirtieth- and eighth-order models are compared: the two are indistinguishable. Like LQG1, we see that LQG2's response is largest at the frequencies of the two unstable Rossiter modes, where the benefits of feedback are sought.

REFERENCES

AHUJA, S. & ROWLEY, C. W. 2010 Feedback control of unstable steady states of flow past a flat plate using reduced-order estimators. *J. Fluid Mech.* **645**, 447–478.

BARBAGALLO, A., SIPP, D. & SCHMID, P. J. 2009 Closed-loop control of an open cavity flow using reduced-order models. *J. Fluid Mech.* **641**, 1–50.

CABELL, R. H., KEGERISE, M. A., COX, D. E. & GIBBS, G. P. 2006 Experimental feedback control of flow-induced cavity tones. *AIAA J.* **44** (8), 1807–1815.

CATTAFESTA, L. N. III, GARG, S., CHOUDHARI, M. & LI, F. 1997 Active control of flow-induced cavity resonance. *AIAA Paper* 97–1804.

CATTAFESTA, L. N. III, SONG, Q., WILLIAMS, D. R., ROWLEY, C. W. & ALVI, F. S. 2008 Active control of flow-induced cavity oscillations. *Prog. Aerosp. Sci.* **44**, 479–502.

DOYLE, J. 1978 Guaranteed margins for LQG regulators. *IEEE Trans. Autom. Control* **23** (4), 756–757.

FREUND, J. B. 1997 Proposed inflow/outflow boundary condition for direct computation of aerodynamic sound. *AIAA J.* **35** (4), 740–742.

GHARIB, M. 1987 Response of the cavity shear layer oscillations to external forcing. *AIAA J.* **25** (1), 43–47.

HELLER, H. H. & BLISS, D. B. 1975 The physical mechanism of flow-induced pressure fluctuations in cavities and concepts for their suppression. *AIAA Paper* 75–491, pp. 281–296.

ILLINGWORTH, S. J., MORGANS, A. S. & ROWLEY, C. W. 2011 Feedback control of flow resonances using balanced reduced-order models. *J. Sound Vib.* **330**, 1567–1581.

JUANG, J. N. & PAPPA, R. S. 1985 Eigensystem realization algorithm for modal parameter identification and model reduction. *J. Guid. Control Dyn.* **8** (5), 620–627.

KEGERISE, M. A., CABELL, R. H. & CATTAFESTA, L. N. III 2007 Real-time feedback control of flow-induced cavity tones—Part 1: Fixed-gain control. *J. Sound Vib.* **307**, 906–923.

KOOK, H., MONGEAU, L. & FRANCHEK, M. A. 2002 Active control of pressure fluctuations due to flow over Helmholtz resonators. *J. Sound Vib.* **255** (1), 61–76.

LANGHORNE, P. J., DOWLING, A. P. & HOOPER, N. 1990 A practical active control system for combustion oscillations. *J. Propul. Power* **6** (3), 324–333.

LELE, S. K. 1992 Compact finite difference schemes with spectral-like resolution. *J. Comput. Phys.* **103** (1), 16–42.

LJUNG, L. 1999 *System Identification: Theory for the User*, 2nd edn. Prentice-Hall.

MA, Z., AHUJA, S. & ROWLEY, C. W. 2011 Reduced order models for control of fluids using the Eigensystem Realization Algorithm. *Theor. Comput. Fluid Mech.* **25**, 233–247.

MEZIĆ, I. & BANASZUK, A. 2004 Comparison of systems with complex behaviour. *Physica* D: *Nonlinear Phenom.* **197**, 101–133.

MORGANS, A. S. & DOWLING, A. P. 2007 Model-based control of combustion instabilities. *J. Sound Vib.* **299**, 261–282.

POINSOT, T. J. & LELE, S. K. 1992 Boundary conditions for direct simulations of compressible viscous flows. *J. Comput. Phys.* **101** (1), 104–129.

POWELL, A. 1953 On edge tones and associated phenomena. *Acustica* **3**, 233–243.

ROSSITER, J. E. 1964 Wind–tunnel experiments on the flow over rectangular cavities at subsonic and transonic speeds. In *Aeronaut. Res. Counc. R& M* 3438.

ROWLEY, C. W., COLONIUS, T. & BASU, A. J. 2002*a* On self-sustained oscillations in two-dimensional compressible flow over rectangular cavities. *J. Fluid Mech.* **455**, 315–346.

ROWLEY, C. W. & JUTTIJUDATA, V. 2005 Model-based control and estimation of cavity flow oscillations. *44th IEEE Conference on Decision and Control*, vol. 15, pp. 512–517.

ROWLEY, C. W., WILLIAMS, D. R., COLONIUS, T., MURRAY, R. M., MACMARTIN, D. G. & FABRIS, D. 2002*b* Model-based control of cavity oscillations part II: system identification and analysis. *AIAA Paper* 2002–0972.

ROWLEY, C. W., WILLIAMS, D. R., COLONIUS, T., MURRAY, R. M. & MACMYNOWSKI, D. G. 2006 Linear models for control of cavity flow oscillations. *J. Fluid Mech.* **547**, 317–330.

SAMIMY, M., DEBIASI, M., CARABALLO, E., SERRANI, A., YUAN, X., LITTLE, J. & MYATT, J. H. 2007 Feedback control of subsonic cavity flows using reduced-order models. *J. Fluid Mech.* **579**, 315–346.

SAROHIA, V. & MASSIER, P. F. 1977 Control of cavity noise. *J. Aircraft* **14** (9), 833–837.

SHAW, L. & NORTHCRAFT, S. 1999 Closed loop active control for cavity acoustics. *AIAA Paper* 99–1902. pp. 683–689.

TIERNO, J. E. & DOYLE, J. C. 1992 Multi mode active stabilization of a Rijke tube. In *Active Control of Noise and Vibration, 1992: presented at the ASME Winter Annual Meeting*, vol. DSC-Vol. 38, pp. 65–68.

VINNICOMBE, G. 2000 *Uncertainty and Feedback*: $\mathscr{H}_\infty$ *Loop-Shaping and the* ν*-gap Metric*. Imperial College Press.

WILLIAMS, D. R., FABRIS, D. & MORROW, J. 2000 Experiments on controlling multiple acoustic modes in cavities. *AIAA Paper* 2000–1903.

YAN, P., DEBIASI, M., YUAN, X., LITTLE, J., ÖZBAY, H. & SAMIMY, M. 2006 Experimental study of linear closed-loop control of subsonic cavity flow. *AIAA J.* **44** (5), 929–938.