AIAA 2000-1969

POD BASED MODELS OF SELF-SUSTAINED OSCILLATIONS IN THE FLOW PAST AN OPEN CAVITY

Clarence W. Rowley Tim Colonius Richard M. Murray

Division of Engineering and Applied Science
California Institute of Technology
Pasadena, CA 91125

Abstract

The goal of this work is to provide accurate dynamical models of oscillations in the flow past a rectangular cavity, for the purpose of bifurcation analysis and control. We have performed an extensive set of direct numerical simulations which provide the data used to derive and evaluate the models. Based on the method of Proper Orthogonal Decomposition (POD) and Galerkin projection, we obtain low-order models (from 6 to 60 states) which capture the dynamics very accurately over a few periods of oscillation, but deviate for long time.

1 Introduction

Open cavities on aircraft are subject to intense resonant pressure fluctuations. Resulting internal acoustic loads with sound pressure levels (SPL) in excess of 160 dB have been measured and these can damage stores, fatigue nearby surfaces and components, and lead to intense noise radiation.[30,31] The reduction of these fluctuations via active flow control is the motivation for this work.

We present the results of a modeling effort for describing the flow over a rectangular cavity, aimed at providing a low-order dynamic model suitable for bifurcation study and model-based active control.

Our modeling uses as input an extensive set of two-dimensional direct numerical simulations[5] of the flow for a range of Mach numbers, cavity aspect ratios, upstream boundary layer thicknesses, and Reynolds numbers. At present, the flow is restricted to sufficiently low Reynolds number such that the upstream boundary layer is laminar.

1.1 Previous control strategies

It has long been known that passive devices, such as spoilers and ramps, can attenuate cavity oscillations under certain conditions.[12] Recently, active flow control has been used to reduce resonant oscillations in subsonic and transonic flows over open cavities. The (open-loop) introduction of flow disturbances has been studied both experimentally and numerically.[3,10,19,30,34] Significant reductions in sound pressure levels (around 20 dB) have been achieved through upstream mass injection[30] as well as piezoelectric flaps.[3] Closed-loop feedback control has also been studied,[3,15,21] and promises similar reductions in noise, but with much lower power input.[3]

So far, control strategies used on the cavity have been relatively simple, typically gain/phase/delay controllers[31,35] with notch filters, or adaptive controllers[4] which are not model-based. These have been reasonably successful at attenuating oscillations at certain frequencies, but if one wants to improve these results by making use of more sophisticated control strategies, it is necessary to have an accurate model of the flow physics.

1.2 Conceptual modeling

In compressible flows, cavity resonance is thought to arise from a feedback loop involving: (i) shear layer instability and the growth of disturbances (vortices) in the shear layer, (ii) the impingement of the vortices at the downstream edge, and subsequent scattering of acoustic waves, (iii) the transmission of acoustic waves upstream, and (iv) their conversion to vortical fluctuations at the cavity leading edge (receptivity). The first description of this feedback process is credited to Rossiter.[27] His semi-empirical formula to predict the measured resonant frequencies remains widely used:

$$\mathrm{St}_n = \frac{n-\alpha}{M+1/\kappa}, \qquad n = 1, 2, 3, \ldots, \tag{1}$$

where $\mathrm{St}_n := f_n L/U$ is the Strouhal number corresponding to the n-th mode frequency f_n, L is the cavity length, U is the freestream velocity, M the Mach number, κ is the average phase speed of the vortical disturbances, and α is an empirical constant, typically taken as 0.25. The formula and, indeed,

the underlying instability process are similar to that which occurs in edgetones.[24]

More detailed models of the processes (i) to (iv) have also been used successfully to predict the frequency.[33] Predicting the overall amplitude of the resonant oscillations is more challenging. The amplitude depends critically on the amplification of disturbances in the shear layer region and the way in which these disturbances plus any small scale turbulence cause spreading of the mean flow. We refer the reader to several reviews,[1,25,26] and more recent analyses,[2,33] for a detailed account of the different models that have been proposed. Prediction of the mean flow spreading rate based on the amplification of the disturbances, for example using a growth model based on an integral mechanical energy balance of the type used by Morris et al,[22] would appear to provide a reasonable approach for predicting the amplitudes, and these ideas have recently been incorporated into the Cavity Acoustics Modeling Simulation (CAMS) design tool at Boeing.[2] The amplitude of the resonant modes is, of course, not only a function of the growth of instabilities in the shear layer, but depends also on the conversion of vortical disturbances into acoustic waves at the trailing edge, the propagation of the waves, and their re-conversion to vortical disturbances at the leading edge (or along the shear layer in a distributed way). The modeling of similar processes using Wiener-Hopf techniques was recently attempted for the edgetone by Crighton,[8] and for supersonic cavities by Kerschen.[14]

1.3 Low-order phenomena

Cavity oscillations exhibit several phenomena characteristic of a low-dimensional dynamical system. For instance, for a fixed flow velocity U and momentum thickness θ upstream of the cavity, there is a minimum cavity length L below which oscillations do not occur, as seen in low Mach-number experiments by Sarohia.[29] Simiarly, if U is varied, keeping other parameters fixed, there is a minimum U below which oscillations do not occur, and if the momentum thickness θ is varied, there is a maximum thickness above which oscillations do not occur. In cavity flows, no hysteresis has been observed, but in other edge tones hysteresis has been observed.

Other interesting phenomena arise involving the frequency of oscillation. If, for instance, the length L is gradually increased, the frequency of oscillation decreases linearly, but once a critical value of L is reached, the frequency jumps up, as the cavity switches to a higher Rossiter mode. Mode switching has also been observed, where the cavity periodically switches between two different Rossiter frequencies.

In addition, experiments by Gharib and Roshko[11] and recent simulations we have performed[5] have demonstrated transition to a "wake mode," where the cavity flow resembles the wake behind a bluff body, rather than a mixing layer. In our simulations, any of the parameters considered (L/D, L/θ, M, and Re) can cause transition to wake mode. All of the experiments and simulations which demonstrate wake mode have laminar boundary layers upstream of the cavity. Colonius et al[6] have developed a criterion for predicting the onset of wake mode in flows with laminar boundary layers.

Many of these features of the cavity flow resemble qualitative features characteristic of low-dimensional nonlinear dynamical systems: limit cycles, bifurcations, and heteroclinic orbits. These qualitative similarities suggest the possibility of describing the cavity flow with a low-order model. The goal is to understand the various transitions more rigorously, as bifurcations in a dynamical system, and ultimately to control the cavity oscillations with a model-based feedback law.

2 Direct Numerical Simulations

The fully compressible, unsteady Navier-Stokes equations are solved using 6th-order-accurate compact finite-difference schemes[20] and 4th-order Runge-Kutta time advancement. Nonreflecting boundary conditions are applied at all open boundaries, and are presently based on buffer zone techniques[7,9] together with the one-dimensional boundary conditions of Poinsot and Lele.[23] Further details may be found in Colonius et al.[5]

The code has been validated by a careful convergence study, and the results are independent of grid spacing and boundary placement.[5] Figure 1 shows a comparison of the DNS with experiments of Krishnamurty[18] for two different Mach numbers.

As mentioned in the previous section, the simulations show an interesting transition to a wake mode. In the present study, we restrict our attention to the shear layer mode oscillations, and use the data from the numerical experiments both to develop low-order models of cavity resonance, and as a reference with which to compare the resulting models.

3 Model reduction

In this section, we discuss the POD/Galerkin method, for obtaining reduced-order models using

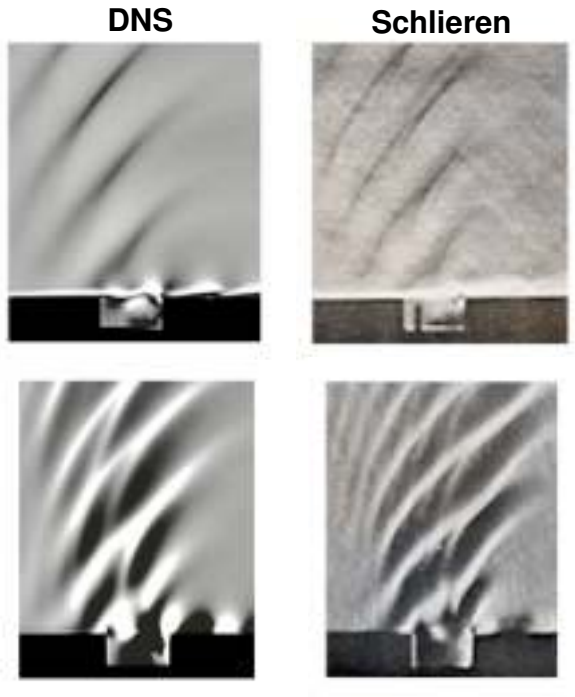

Figure 1: Density gradients of DNS (left) compared with Schlieren photographs[18] (right), at $M = 0.7$ (top) and $M = 0.8$ (bottom).

data from computations or experiments.

3.1 Proper orthogonal decomposition

The Proper Orthogonal Decomposition (POD) is a commonly used tool for extracting coherent structures from data, either experimental or computational.[13,32] Given a set of data, represented as a function of space and time, the POD determines a basis set of orthogonal functions of space which span the data optimally in the L^2 sense. More precisely, if $u(x,t)$ is a function of space and time, POD determines orthogonal functions $\varphi_j(x)$, $j = 1, 2, \ldots$, such that the projection onto the first n functions

$$\hat{u}(x,t) = \sum_{j=1}^{n} a_j(t)\varphi_j(x) \tag{2}$$

has the smallest error, defined by $E(\|u-\hat{u}\|^2)$. Here, $E(\cdot)$ denotes time average and $\|\cdot\|$ denotes the L^2 norm on functions of space. The functions φ_j are computed by solving the Fredholm integral equation

$$\int K(x,y)\varphi(y)\,dy = \lambda\varphi(x), \tag{3}$$

where the kernel $K(x,y) = E\big(u(x,t)u(y,t)\big)$. The functions φ_j are called the *POD modes* (also called *Karhunen-Loève eigenfunctions*, or *empirical eigenfunctions*).

In practice, the data is usually discretized in both space and time. In this case, the integral equation (3) reduces to a standard eigenvalue problem, and in the common case where the number of snapshots in time is smaller than the number of gridpoints, the POD modes are most easily computed using the *method of snapshots*, in which one constructs the *data matrix*

$$A = \begin{pmatrix} u(x_1,t_1) & u(x_2,t_1) & \cdots & u(x_n,t_1) \\ u(x_1,t_2) & u(x_2,t_2) & \cdots & u(x_n,t_2) \\ \vdots & \vdots & \ddots & \vdots \\ u(x_1,t_m) & u(x_2,t_m) & \cdots & u(x_n,t_m) \end{pmatrix}, \tag{4}$$

where here n is the number of gridpoints in x, m the number of snapshots, and x_j and t_j denote the discretized spatial points and times, respectively. The POD modes are then the right singular vectors of A, easily computed by standard, efficient algorithms for singular value decomposition.

If a nonuniform grid is used for the discretization in x then, of course, the solutions to the integral equation (3) for the continuous system will not correspond to the right eigenvectors of the data matrix of the discretized system, even as the grid spacing goes to zero, unless the columns of the data matrix are weighted appropriately, according to the volume of the computational element corresponding to x_j. In this paper, data is always interpolated onto a uniform grid before the POD modes are computed, so this subtlety is avoided.

Note that we may easily take $u(x,t)$ to be vector valued, by introducing the appropriate inner product on vector-valued functions. In this case, the POD modes φ_j are also vector valued, but the time coefficients a_j are scalars.

Also note that for systems with symmetry (for instance, axisymmetric or 3D cylindrical jets), it is important to take symmetries into consideration. Proper treatment of continuous symmetry in the POD method has been shown to lead to dramatic computational savings.[28] See Sirovich[32] or Holmes et al[13] for more details on the POD method.

3.2 Galerkin projection

Assume we have a system governed by a partial differential equation (PDE) of the form

$$\partial_t u = D_\lambda(u) \tag{5}$$

where $u(x,t)$ is a function of space and time, and D_λ is a nonlinear spatial differential operator which depends on some parameters λ. In our case, equation (5) will be the Navier-Stokes equations or the Euler equations, and λ will contain the parameters we are interested in studying (for instance, $\lambda = (M, L/\theta, \mathrm{Re})$).

We wish to compute approximate solutions of (5) by solving a simpler set of ordinary differential equations (ODEs). To do this, we start with a truncated series expansion (2), where the POD modes φ_j are known. Substituting this expansion into (5), and taking an inner product with φ_k gives

$$\dot{a}_k = \langle D_\lambda(u), \varphi_k \rangle, \qquad k = 1, \ldots, n, \tag{6}$$

where $\langle \cdot, \cdot \rangle$ denotes the L^2 inner product. Since u is written in terms of the time coefficients by (2), this gives a set of ODEs

$$\dot{a} = f(a, \lambda), \tag{7}$$

where $a = (a_1, \ldots, a_n)$.

3.3 Equations of motion

The system of PDEs we start with for the purposes of model reduction is the isentropic Euler equations, given in dimensional form by

$$u_t + uu_x + vu_y + h_x = 0 \tag{8}$$
$$v_t + uv_x + vv_y + h_y = 0 \tag{9}$$
$$h_t + uh_x + vh_y + (\gamma - 1)h(u_x + v_y) = 0, \tag{10}$$

where u and v are velocities, h is the enthalpy, and γ is the ratio of specific heats. We will also consider these equations with approximate viscous terms added:

$$u_t + uu_x + vu_y + h_x = \nu(u_{xx} + u_{yy}) \tag{11}$$
$$v_t + uv_x + vv_y + h_y = \nu(v_{xx} + v_{yy}) \tag{12}$$
$$h_t + uh_x + vh_y + (\gamma - 1)h(u_x + v_y) = 0, \tag{13}$$

where in making this approximation we neglect heat dissipation and density variations with temperature. Considering the full viscous equations would add cubic terms (e.g., $\rho u u_x$) and an additional (energy) equation, so this approximation represents a significant reduction in complexity, and is reasonable since temperature fluctuations are small in this flow.

Note that several of the parameters we are interested in (here, M and L/θ, where θ is the momentum thickness at the upstream edge of the cavity) do not appear in the equations. They will appear either in the initial conditions (M, θ) or in the boundary conditions (L). This is undesirable because there is no clear way to include parameters in the Galerkin equations (7) unless they appear directly in the governing PDEs, like λ did in the previous section.

We can resolve this difficulty by appropriately nondimensionalizing the equations, such that the relevant parameters do appear. If we nondimensionalize velocities by the free-stream velocity U, the enthalpy by the square of the sound speed a, x by L, y by θ, and t by L/U, the equations become

$$u_t + uu_x + \frac{L}{\theta}vu_y + \frac{1}{M^2}h_x = \frac{1}{\mathrm{Re}_\theta}\left(\frac{\theta}{L}u_{xx} + \frac{L}{\theta}u_{yy}\right)$$
$$v_t + uv_x + \frac{L}{\theta}vv_y + \frac{1}{M^2}\frac{L}{\theta}h_y = \frac{1}{\mathrm{Re}_\theta}\left(\frac{\theta}{L}v_{xx} + \frac{L}{\theta}v_{yy}\right)$$
$$h_t + uh_x + \frac{L}{\theta}vh_y + (\gamma - 1)h(u_x + \frac{L}{\theta}v_y) = 0,$$

where $\mathrm{Re}_\theta = U\theta/\nu$ is the Reynolds number based on momentum thickness at the cavity leading edge. Now all of the relevant parameters appear in the equations, instead of lying hidden in the initial conditions or boundary conditions. This idea of rescaling the equations is a special case of a more general idea of mapping different domains (geometries) into a canonical geometry, originally suggested by I. Kevrekidis and P. Holmes for unsteady flow through diffusers with different geometries.[16,17]

It is not possible, without a more sophisticated stretching, to match both L/θ and L/D. For the cavity flow, it has been demonstrated by Sarohia[29] that for sufficiently deep cavities, the cavity depth has negligible effect, compared to the effect of the boundary layer momentum thickness (at least for low Mach number). For the purpose of modeling, therefore, it seems reasonable to ignore the effect of L/D, and focus on L/θ.

4 Results

We consider the results of three of our DNS runs, with parameters given in Table 4. Runs A, B, and C here correspond to runs L2, H2, and TK4b, respectively, in Colonius et al,[6] where more details about the flow field are given.

First, we show the POD modes obtained for several different run conditions, and then we show how various low-order models behave, for a particular run condition.

Run	L/D	L/θ	Re_θ	M
A	2	52.8	56.8	0.6
B	2	58.4	68.5	0.6
C	4	60.2	58.8	0.6

Table 1: Parameters for different runs considered.

4.1 POD modes

We computed POD modes for the runs listed in Table 4, using the method of snapshots, taking between 51 and 101 snapshots over 5–10 periods of the cavity oscillation. In all cases discussed here, we took the snapshots after the initial transients had settled out. The mean of all the snapshots is subtracted before the POD modes are computed, so the actual expansion considered is

$$q(x,y,t) = \sum_{j=1}^{n} a_j(t)\varphi_j(x,y) + \bar{q}(x,y) \qquad (14)$$

where q is any one of $\{u, v, h\}$, and $\bar{q}$ denotes the time average of the snapshots.

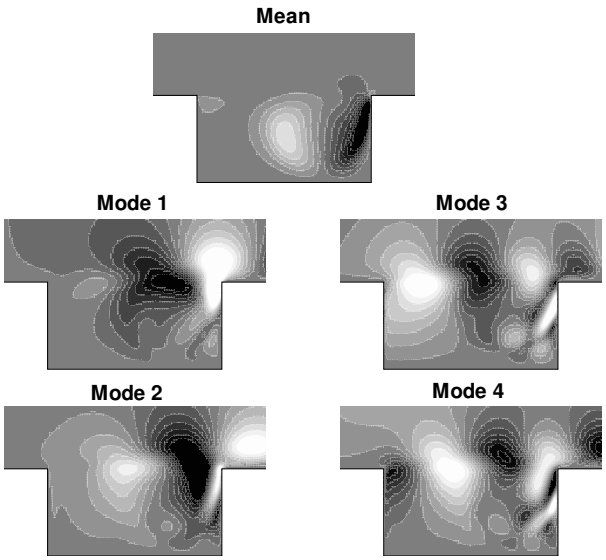

Figure 2: Mean and first four POD modes for Run A. (Contours of v-velocity.)

Figure 2 shows v-velocity contours of the mean and POD modes for Run A. The modes shown in Figure 2 are typical of the runs considered here, and show one or two wavelengths in the shear layer, indicating the presence of the first two Rossiter modes. The POD modes for Run B (not shown) look similar, although the first two modes shown in Figure 2 are not present, because in Run B the cavity is oscillating primarily in Rossiter mode 2, while in Run A the cavity oscillates in Rossiter modes 1 and 2. In general, the higher POD modes show finer and finer spatial scales.

Figure 3 shows the percent of the fluctuating energy ($\|q - \bar{q}\|^2$) captured in each mode, for Runs A and B. The first 4 modes capture a total of 50% of the energy in Run A, and 76% in Run B. The difference between the two runs is to be expected, since Run A has more complicated behavior. The higher modes still capture a significant amount of energy, but since this energy is at the smaller spatial scales, it is likely that they will be unimportant for modeling the overall global behavior of the cavity. This hypothesis is tested and verified in section 4.3.2.

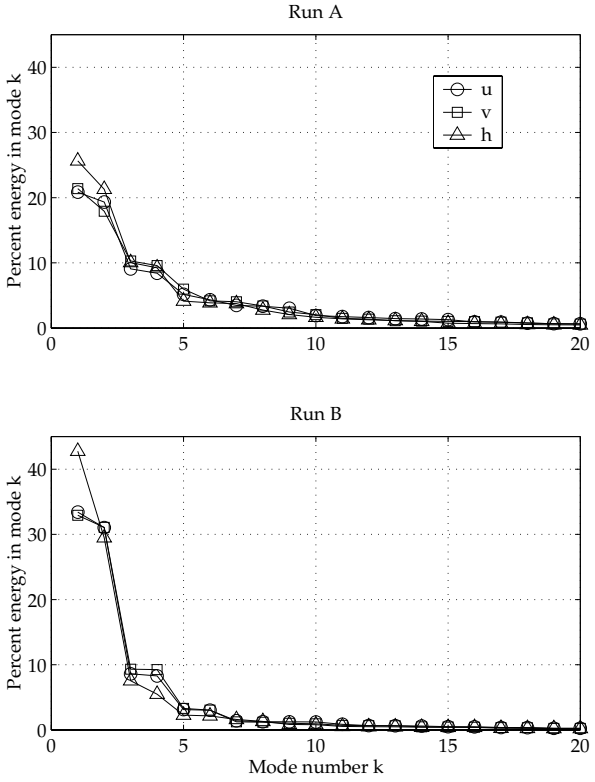

Figure 3: Percent energy in the fluctuations captured by POD modes for u, v, and h, for two different run conditions. In Run A, the cavity is oscillating in Rossiter modes 1 and 2, and in Run B, the cavity is in Rossiter mode 2.

4.1.1 Scaling

It is interesting to compare (qualitatively) POD modes for different operating conditions, to evaluate

the scaling ideas described in the previous section. Figure 4 shows the mean and the first POD mode from Run B, compared to those from Run C. These runs have nearly the same value of L/θ, but different values of L/D.

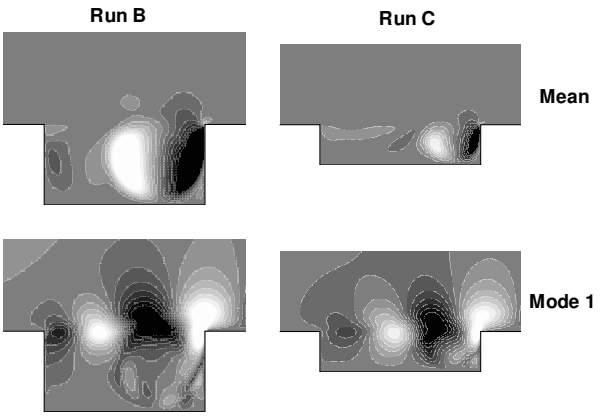

Figure 4: Mean and first POD mode for Runs B and C. (Contours of v-velocity.)

The pictures are indeed qualitatively similar, especially in the shear-layer region, where the dynamics are important. This provides further evidence that L/θ is a reasonable choice of scaling parameter (or bifurcation parameter), and is perhaps more relevant than L/D.

4.2 Linear stability modes

The POD modes also look qualitatively similar to eigenfunctions determined from linear stability theory (although these eigenfunctions are obviously not orthogonal). Figure 5 shows a comparison three different types of modes, all of which look qualitatively similar: POD modes 1 and 3 from Run A; modes computed from a quasi-parallel linear stability calculation, for the two primary frequencies observed, those of Rossiter modes 1 and 2 (St = 0.4 and 0.7, respectively); and a discrete Fourier transform of the DNS data, for the same two frequencies.

The quasi-parallel inviscid linear stability calculation uses a tanh profile for the shear layer, and the (slowly varying) shear layer thickness at each x is measured from the DNS. The bottom wall is imposed as a boundary condition, but the side walls of the cavity are not considered in the quasi-parallel analysis. For a given frequency, spatial growth rates and eigenfunctions are computed at each x-location, and then the resulting eigenfunctions are put together, and the growth rates integrated, to give the modes in Figure 5.

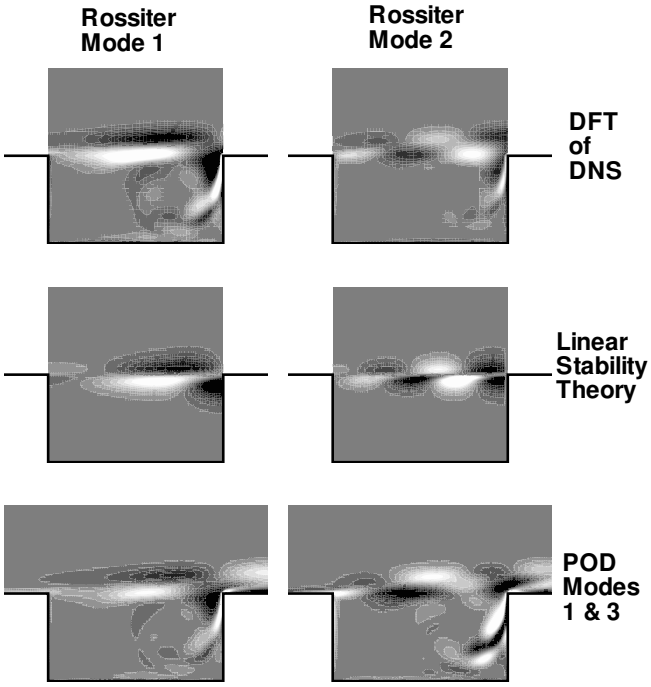

Figure 5: Vorticity eigenfunctions computed from linear stability theory, for Rossiter mode 1 (St = 0.40) and Rossiter mode 2 (St = 0.70); discrete Fourier transform of DNS for the same frequencies; and comparison with vorticity POD modes 1 and 3 for Run A.

The mode shapes from linear theory show good agreement with the DNS results. The growth rate of each mode, when scaled with the DNS data near $x = 0$, is also well predicted. This result indicates that the saturation of the linear stability wave is determined primarily by the spreading of the mean flow (which in turn modifies the local growth rate of the mode), rather than any direct nonlinear interaction with other frequencies. This supports the modeling approach used by Cain et al[2] for predicting cavity oscillation amplitudes, but it remains to be determined how well the spreading of the mean flow could be predicted *a priori*.

The linear stability modes agree well even near the downstream corner, indicating that the presence of the corner does not significantly affect the behavior of the free shear layer (although it clearly impacts the overall flow, by scattering acoustic waves).

4.3 Low-order models

In this section, we discuss low-order models for the cavity, with POD modes taken from Run B. The ini-

tial conditions for all runs are obtained by projecting a snapshot of the DNS run onto the POD modes. In sections 4.3.1 and 4.3.2, the snapshot for the initial condition is taken after transients have settled out, and the behavior in the transient region is addressed in section 4.3.3.

4.3.1 Effects of viscosity

Since many of the important phenomena of cavity oscillations are essentially inviscid phenomena (e.g., shear layer instability, acoustic wave propagation), one might expect that the inviscid equations would model the flow nearly as well as the viscous equations.

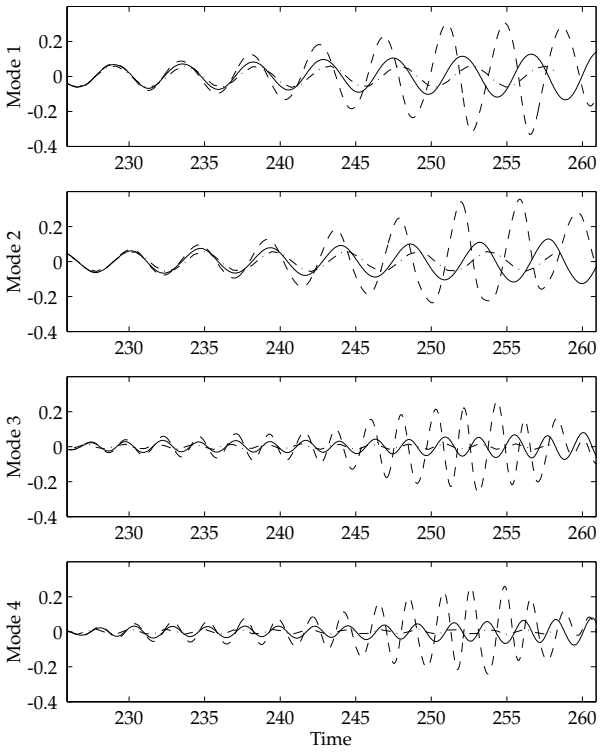

Figure 6: Time traces of first 4 POD modes for u-velocity: 4-mode Galerkin simulations with viscosity (———), without viscosity (– – – –), and projection of DNS (—·—).

Figure 6 shows time traces from two different 4-mode Galerkin models, with and without the approximate viscous terms discussed in section 3.3. The viscous model uses the same value of Re_θ used in Run B.

Both models match the DNS very well for the first two periods of oscillation, but then oscillations start to grow, and the inviscid model starts to perform worse. For longer times, the inviscid calculation eventually blows up, while the viscous model remains stable, at least until $t = 200$, with no sign of blow-up. For the remainder of this paper, we consider the viscous model exclusively.

4.3.2 Effects of higher modes

Figure 7 shows how the models behave as we increase the number of modes retained from 2 to 4 to 20. (These correspond to ODEs with 6, 12, and 60 states, respectively, since modes for u, v, and h are taken separately.) Initially, all of the models track well for the first three or four periods. Though it is not apparent in the figure, the 20-mode case follows the DNS the closest for the first period, but later the 2-mode case actually performs best of all.

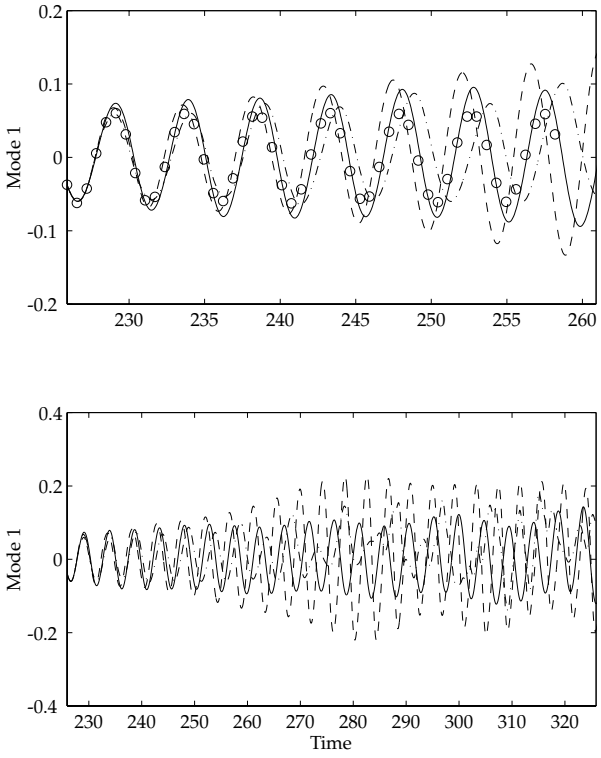

Figure 7: The top plot shows time traces of the first POD mode for u-velocity: Galerkin simulations 2 modes (———), 4 modes (– – – –), and 20 modes (—·—), and projection of DNS (○). The bottom plot is identical, but for longer time.

The bottom plot of Figure 7 shows the same data, for a longer time interval, and here it is clear that none of the models correctly predict the amplitude of the oscillation, although the 2-mode model comes closest. This result is disappointing, but is not sur-

prising, since the theory guarantees only that the projected Galerkin system will track the full system for arbitrarily small times, and says nothing about the global behavior of the low-order model. Also, it is worth pointing out that for many control purposes it is not necessary to have a model which correctly predicts long-time behavior.

4.3.3 Transient behavior

In this section, we investigate how well the model performs at early times when the cavity oscillations are still developing, and have not yet saturated. Since the POD modes were taken after the transients had died out, we would not expect this region of phase space to be as well captured by the projection.

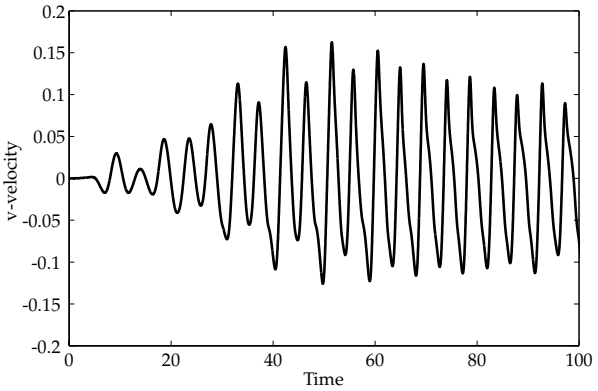

Figure 8: Time trace of a v-velocity probe (in the DNS), placed near the downstream corner of the cavity.

Figure 8 shows data taken from a v-velocity probe placed in the simulation (Run B), near the downstream corner of the cavity. The initial condition for the previous runs were taken much later ($t \approx 225$), long after the transients had died out. Here, we consider an initial condition at time $t = 16$, when the cavity oscillations are just developing.

Figure 9 shows the results of the 4-mode Galerkin model, compared to the projection of the DNS. The model tracks reasonably well for short time, demonstrating that even though the Galerkin model is accurate only locally in time, it is valid over a relatively large portion of phase space. One could, of course, improve this transient response by including snapshots from this region when computing the POD modes.

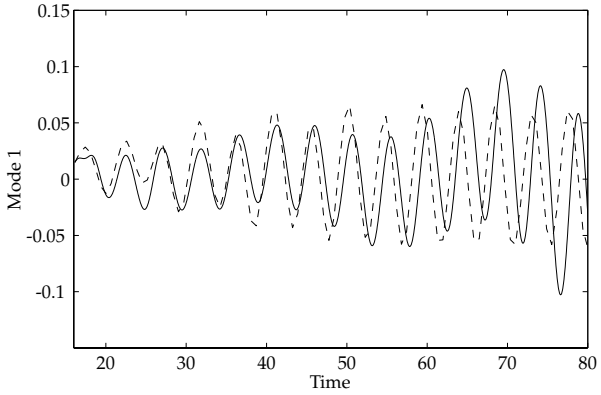

Figure 9: Time trace of first POD mode for u-velocity in transient region: 4-mode Galerkin simulation (——), and projection of DNS (- - - -).

5 Conclusions

We have obtained and evaluated several dynamical models for oscillations in the flow past an open cavity. The models are low-order (6 to 60 states), and are obtained by Galerkin projection of the Euler equations onto spatial modes determined by Proper Orthogonal Decomposition of data from direct numerical simulations. The models work much better when viscous terms are included, but including higher order POD modes does not improve the performance. The models perform very well for short time (two or three periods of oscillation) but then begin to deviate, and do not accurately predict the amplitude of the resulting limit cycles.

A possible reason why the long-time behavior is not better predicted can be understood on physical reasoning. As oscillations grow in amplitude, the mean flow will change, spreading the shear layer, and therefore reducing the growth rate of disturbances. Such an effect is not present in our model, since the spatial shape of the mean and the POD modes is fixed by the snapshots used for the POD. Thus this saturation mechanism (spreading the shear layer) may not be present in our model, and may explain why the model overpredicts the final amplitude of the limit cycle.

This work is part of an ongoing effort for modeling cavity oscillations. Other important issues to be addressed include the effects of the parameters in the equations (particularly L/θ and M), and including the effects of actuators in the model.

Acknowledgements

Supported by AFOSR under grant F49620-98-1-0095 with technical monitor Dr. Thomas Beutner.

References

[1] W. K. Blake and A. Powell. The development of contemporary views of flow-tone generation. In *Recent Advances in Aeroacoustics*, pages 247–345. Springer-Verlag, 1986.

[2] A. B. Cain, W. W. Bower, F. McCotter, and W. W. Romer. Modeling and prediction of weapons bay acoustic amplitude and frequency. Technical report, VEDA Inc., 1996.

[3] L. N. Cattafesta III, S. Garg, M. Choudhari, and F. Li. Active control of flow-induced cavity resonance. AIAA Paper 97-1804, June 1997.

[4] L. N. Cattafesta III, D. Shukla, S. Garg, and J. A. Ross. Development of an adaptive weapons-bay suppression system. AIAA Paper 99-1901, May 1999.

[5] T. Colonius, A. J. Basu, and C. W. Rowley. Numerical investigation of the flow past a cavity. AIAA Paper 99-1912, May 1999.

[6] T. Colonius, A. J. Basu, and C. W. Rowley. On self-sustained oscillations in two-dimensional compressible flow over rectangular cavities. *Journal of Fluid Mechanics*, submitted, 1999.

[7] T. Colonius, S. K. Lele, and P. Moin. Boundary conditions for direct computation of aerodynamic sound. *AIAA Journal*, 31(9):1574–1582, 1993.

[8] D. G. Crighton. The jet edge-tone feedback cycle; linear theory for the operating stages. *Journal of Fluid Mechanics*, 234:361–391, 1992.

[9] J. B. Freund. Proposed inflow/outflow boundary condition for direct computation of aerodynamic sound. *AIAA Journal*, 35(4):740–742, Apr. 1997.

[10] M. Gharib. Response of the cavity shear layer oscillations to external forcing. *AIAA Journal*, 25(1), Jan. 1987.

[11] M. Gharib and A. Roshko. The effect of flow oscillations on cavity drag. *Journal of Fluid Mechanics*, 177:501–530, 1987.

[12] H. H. Heller and D. B. Bliss. The physical mechanism of flow induced pressure fluctuations in cavities and concepts for their suppression. AIAA Paper 75-491, 1975.

[13] P. Holmes, J. L. Lumley, and G. Berkooz. *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press, 1996.

[14] E. J. Kerschen. The acoustic field generated by interaction of a shear-layer instability with the downstream lip of a cavity. Technical report, Flow Dynamics Research, Tucson, AZ, 1996.

[15] T. Kestens and F. Nicoud. Active control of an unsteady flow over a rectangular cavity. AIAA Paper 98-2348, 1998.

[16] I. Kevrekidis and P. Holmes. Personal communication.

[17] A. I. Khibnik, S. Narayanan, C. A. Jacobson, and K. Lust. Analysis of low dimensional dynamics of flow separation. *Notes in Computational Fluid Dynamics*, submitted, 1998.

[18] K. Krishnamurty. *Sound Radiation from Surface Cutouts in High Speed Flow*. PhD thesis, California Institute of Technology, 1956.

[19] A. Lamp and N. Chokani. Active control of compressible cavity flows by using a small jet. AIAA Paper 96-0446, Jan. 1996.

[20] S. K. Lele. Compact finite difference schemes with spectral-like resolution. *Journal of Computational Physics*, 103(1):16–42, 1992.

[21] L. Mongeau, H. Kook, and M. A. Franchek. Active control of flow-induced cavity resonance. AIAA/CEAS Paper 98-2349, 1998.

[22] P. J. Morris, M. G. Giridharan, and G. M. Lilley. On the turbulent mixing of compressible free shear layers. *Proceedings of the Royal Society of London A*, 431:219–243, 1990.

[23] T. Poinsot and S. K. Lele. Boundary conditions for direct simulation of compressible viscous flows. *Journal of Computational Physics*, 101:104–129, 1992.

[24] A. Powell. On the edgetone. *Journal of the Acoustical Society of America*, 33:395, 1961.

[25] D. Rockwell and E. Naudascher. Review—self-sustaining oscillations of flow past cavities. *Transactions of the ASME, Journal of Fluids Engineering*, 100, June 1978.

[26] D. Rockwell and E. Naudascher. Self-sustaining oscillations of impinging free shear layers. *Annual Review of Fluid Mechanics*, 11:67–94, 1979.

[27] J. E. Rossiter. Wind-tunnel experiments on the flow over rectangular cavities at subsonic and transonic speeds. Aeronautical Research Council Reports and Memoranda, No. 3438, Oct. 1964.

[28] C. W. Rowley and J. E. Marsden. Reconstruction equations and the Karhunen-Loève expansion for systems with symmetry. *Physica D*, 142:1–19, 2000.

[29] V. Sarohia. *Experimental and Analytical Investigation of Oscillations in Flows Over Cavities*. PhD thesis, California Institute of Technology, 1975.

[30] L. Shaw. Active control for cavity acoustics. AIAA Paper 98-2347, 1998.

[31] L. Shaw and S. Northcraft. Closed loop active control for cavity acoustics. AIAA Paper 99-1902, May 1999.

[32] L. Sirovich. Turbulence and the dynamics of coherent structures, parts I–III. *Quarterly of Applied Mathematics*, XLV(3):561–590, Oct. 1987.

[33] C. K. W. Tam and P. J. W. Block. On the tones and pressure oscillations induced by flow over rectangular cavities. *Journal of Fluid Mechanics*, 89(2):373–399, 1978.

[34] A. D. Vakili and C. Gauther. Control of cavity flow by upstream mass injection. AIAA Paper 91-1645, 1991.

[35] D. R. Williams and D. Fabris. Closed-loop control in cavities with unsteady bleed forcing. AIAA Paper 2000-0470, Jan. 2000.