MODEL-BASED CONTROL OF CAVITY OSCILLATIONS, PART II: SYSTEM IDENTIFICATION AND ANALYSIS

Clarence W. Rowley[a] David R. Williams[b] Tim Colonius[c] Richard M. Murray[c]

Douglas G. MacMartin[c] Drazen Fabris[d]

[a]*Mechanical & Aerospace Engineering, Princeton University, Princeton, NJ 08544*
[b]*Mechanical, Materials & Aerospace Engineering, Illinois Institute of Technology, Chicago, IL 60616*
[c]*Division of Engineering and Applied Science, California Institute of Technology, Pasadena, CA 91125*
[d]*Mechanical Engineering, Santa Clara University, Santa Clara, CA 95053*

ABSTRACT

Experiments using active control to reduce oscillations in the flow past a rectangular cavity have uncovered surprising phenomena: in the controlled system, often new frequencies of oscillation appear, and often the main frequency of oscillation is split into two sideband frequencies. The goal of this paper is to explain these effects using physics-based models, and to use these ideas to guide control design.

We present a linear model for the cavity flow, based on the physical mechanisms of the familiar Rossiter model. Experimental data indicates that under many operating conditions, the oscillations are not self-sustained, but in fact are caused by amplification of external disturbances. We present some experimental results demonstrating the peak-splitting phenomena mentioned above, use the physics-based model to study the phenomena, and discuss fundamental performance limitations which limit the achievable performance of any control scheme.

1   INTRODUCTION

Recent experiments using feedback to control oscillations in the flow past a cavity have met with limited success. Typical control schemes are able to reduce the steady-state amplitude at one resonant frequency, but increase the amplitude at other frequencies.[10] The goal of this paper is to understand these effects using physics-based models, to use these models to guide future control designs, and to understand any performance limitations of feedback.

The usual description of the mechanism for cavity oscillations involves *self-sustained oscillations*, caused by the familiar Rossiter mechanism:[6–8] small disturbances are amplified by the shear layer, and

produce acoustic waves when they impinge on the downstream corner; these acoustic waves then propagate upstream and excite further instabilities in the shear layer, leading to self-forcing. The usual view is that the cavity is a "globally unstable" flow: in the absence of any external forcing, the cavity would still continue to oscillate. From a dynamical systems point of view, then, the cavity behaves as a dynamical system with an *unstable* equilibrium point (a steady solution of Navier-Stokes), and a stable limit-cycle. The amplitude of the oscillations is thus determined by the nonlinearities.

Here, we present an alternative viewpoint. We demonstrate that for many conditions where oscillations are observed, the cavity behaves as a *stable*, lightly damped system. The flow amplifies noise at certain resonant frequencies, but if the external forcing were removed, the oscillations would disappear. Purely linear models may be used to describe this mechanism, as the final amplitude of oscillations is determined not by nonlinear saturation, but rather by the amplitude of the forcing disturbances (e.g., boundary layer turbulence), and by the linear gain of the system. Feedback may be used to decrease this gain at resonant frequencies, but as we shall see, attenuation at certain frequencies must always be balanced by amplification at other frequencies.

The paper is organized as follows: we give a brief description of the experimental setup in section 2. In section 3, we present the physics-based linear model we use to describe the system, and in section 4 we discuss the stable and unstable regimes observed in the experiment. In section 5 we present the results of a frequency response experiment designed to identify the system experimentally. Finally, in sections 6 and 7 we present some surprising phenomena observed in the closed-loop experiments, and use the physics-based models to explain the observed behavior.

Figure 1: Diagram of experimental apparatus (side view). Location of Kulite pressure transducers is indicated by K1–K8, location of hot films indicated by HW1, HW2.

2 Experimental setup

The experimental apparatus is described in detail in a companion paper,[11] so here we give only a brief overview.

Experiments were performed using the $3\,\text{ft} \times 3\,\text{ft}$ subsonic wind tunnel at the United States Air Force Academy in Colorado Springs. A cavity model was installed in the floor of the test section, and a diagram of the setup is shown in figure 1.

2.1 Sensors

The cavity was instrumented with eight Kulite pressure transducers placed along the cavity walls, one on the upstream wall, one on the downstream wall, and six along the floor, approximately equally spaced. The approximate locations are indicated in figure 1.

Velocity measurements were obtained using two hot-film sensors placed in the shear layer spanning the cavity, one near the upstream corner, and one near the downstream corner, both in line with the lip of the cavity, and close to the center in the spanwise direction. These are indicated in figure 1.

All signals were passed through anti-aliasing filters prior to sampling by a digital data acquisition system. Data were sampled at $6\,\text{kHz}$, typically for 65,536 samples (10.9 sec). The anti-aliasing filters were 4th-order Butterworth bandpass filters, with a pass band of $0.6\,\text{Hz}$–$2.2\,\text{kHz}$. (The high-pass filter was necessary to remove the DC offset, and when needed, the DC component was measured using a digital multimeter.)

2.2 Actuator

The flow was forced using zero-net-mass blowing through a slot in the upstream wall of the cavity, shown in figure 1. The actuator was a pair of 500-Watt $8\,\text{in}$ diameter loudspeakers in an enclosed chamber. Though the actuator injects zero net mass through the slot, a nonzero net momentum is induced by spanwise vortices generated by periodically blowing through the slot (the "synthetic jet" effect).

2.3 Control implementation

Both analog and digital controllers were implemented. The analog controller consisted of a band-pass filter and a phase shifter. The gain and phase could be continuously adjusted, and the frequencies of the passbands could be adjusted in discrete increments.

Digital controllers were implemented using a dSPACE interface board, running on a separate computer from the data acquisition system. For typical controllers we were running, the maximum sample rate of the dSPACE system was about $20\,\text{kHz}$.

3 Physics-based model

Our model for the cavity dynamics is based on the familiar Rossiter mechanism described in the introduction. A block diagram of the model is shown in figure 2, where we represent each component of the physical mechanism as a separate transfer function.

Here, $G(s)$ represents the shear-layer transfer function, i.e., the transfer function from velocity disturbances v_0 at the leading edge to velocity disturbances v_L at the trailing edge. Transfer functions for acoustic scattering, propagation, and receptivity are given by S, A, and R, and in the diagram p_0 and p_L denote pressure disturbances at the leading and trailing edges, respectively. These quantities may be measured from the experiment: v_0 is measured by hot film 1, v_L by hot film 2, and p_L and p_0 by Kulites 2 and 8, respectively (see figure 1). Here, we do not use Kulite 1, as this sensor measures substantial pressure fluctuations from the impinging shear layer.

The other transfer functions depicted in figure 2 represent the influence of the actuator and controller. The controller transfer function (which we choose) is given by C, and the actuator dynamics are described by a transfer function V. The control signal u is the voltage to the amplifier, and we use the pressure signal from Kulite 8 as the plant output y. The plant is excited by external noise (e.g., turbulent boundary layer fluctuations), modeled by a stochastic input n.

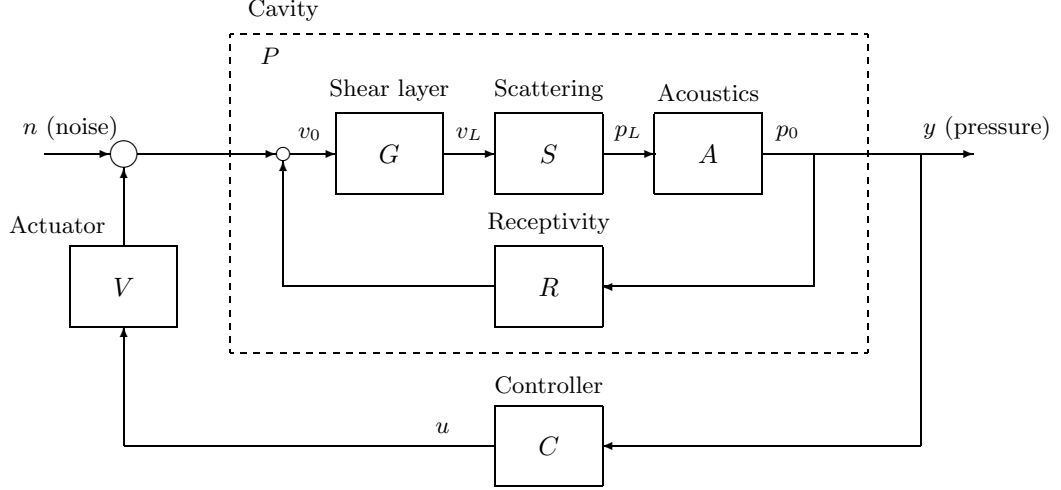The overall transfer function P for the cavity is

Figure 2: Block diagram of cavity model.

then

$$P = \frac{ASG}{1 - RASG}. \qquad (1)$$

For the purposes of studying the dynamical features of this model, we ignore the actuator dynamics, setting $V = 1$. (These actuator dynamics may in principle be measured from the experiment, and once measured, their effects may be inverted out of the control laws we obtain.) Theoretical models for the remaining transfer functions are discussed below.

3.1   Shear layer

The shear layer transfer function $G(s)$ may be determined from linear stability theory. We begin with velocity profiles measured in experiments by Williams and Fabris,[9] shown in figure 3. These profiles are from a run with Mach number $M = 0.35$, in a cavity with aspect ratio $L/D = 5$. Figure 3 shows the experimental data along with hyperbolic tangent profiles with the same vorticity thickness. The spreading rate of the shear layer is determined from a linear fit to the data, and used as an input to a linear stability calculation to determine the amplification and phase of shear layer disturbances. We then fit a rational function to the resulting transfer function, and the result is shown in figure 4.

As an alternative model, we also consider the shear layer modeled as a second-order system with a time delay

$$G(s) = G_0(s)e^{-s\tau_s} = \frac{\omega_0^2}{s^2 + 2\zeta\omega_0 s + \omega_0^2}\, e^{-s\tau_s}, \quad (2)$$


Figure 3: Velocity profiles for the cavity shear layer. Hot wire measurements ($\square$) and tanh profiles with same vorticity thickness and deflection (——).

Figure 4: Bode plot of shear layer transfer function $G_s(s)$, determined from linear stability theory.

Figure 5: Block diagram of transfer function $A(s)$ for cavity acoustics.

where ω_0 is (approximately) the frequency of the most unstable Kelvin-Helmholtz mode, and ζ is the damping, related to the maximum shear-layer amplification. The time delay τ_s is the convection time for a disturbance to travel the length of the cavity, and is given by $\tau_s = L/c_p$, where c_p is the mean phase speed.

Though the model (2) has less physical justification than the model obtained from the linear stability calculation, it provides the same general features, is easier to analyze, and has adjustable parameters, which allow it to be tuned to better match experimentally observed results.

3.2 Acoustics

The model we use for acoustic propagation in the cavity is shown in figure 5. Here, $\tau_a = L/a$ is a time delay which represents the acoustic lag between the trailing edge and leading edge (here, L is the cavity length and a is the sound speed inside the cavity). An acoustic wave emanating from the downstream corner $x = L$ propagates upstream, and some of


Figure 6: Frequency response of acoustic transfer function, for $M = 0.6$. Reflection coefficient $r = 0$ (———); $r = 0.25$ (-----); $r = 0.5$ (—·—); $r = 0.75$ (········). Note the linear frequency scale.

it reflects off the upstream wall, propagates downstream, and again reflects off the downstream wall. The reflection coefficient r measures the attenuation in these reflections (e.g., if both reflections are perfect, then $r = 1$; if each reflection reduces the amplitude by 0.5, then $r = 0.25$). This model therefore captures longitudinal modes of acoustic resonance, but ignores depth modes. It is probably reasonable to ignore the depth modes for such a shallow cavity ($L/D = 5$).

The overall transfer function is given by

$$A(s) = \frac{e^{-s\tau_a}}{1 - re^{-2s\tau_a}} \qquad (3)$$

and the Bode plot is shown in figure 6, for $M = U/a = 0.6$, and various values of r, ranging from 0 to 0.75. The resonant peaks are clearly apparent, and for this simple model, all of the harmonics are equally strong. (Harmonics of different strengths could easily be introduced by making r frequency-dependent.) Note that these resonant peaks represent the longitudinal acoustic modes of the cavity, and not the Rossiter frequencies. However, when these resonant acoustic frequencies approach the Rossiter frequencies, they may influence the mode selection, determining which Rossiter mode is dominant.[10]

3.3 Scattering and Receptivity

The scattering and receptivity effects, which couple vortical and acoustic disturbances at the trailing and

leading edge, are the least well understood aspects of the cavity model. Details of these effects have been studied by Crighton[2] for edge tones, and Kerschen[4] for cavity flows. However, both of these models describe the scattering by a sharp edge, rather than a corner; scattering at a sharp edge produces a dipole source, as is well known for edge tones, while the acoustic source in the cavity is more closely represented by a monopole.[7,8] Furthermore, these previous scattering models are quite detailed, employing a Wiener-Hopf factorization (which does not extend easily to corners).

In Rossiter's empirical formula for predicting cavity frequencies, the scattering and receptivity effects are treated together as a simple phase lag, independent of frequency. Here, we follow Rossiter's approach and model scattering and receptivity each as constant gains, essentially neglecting them altogether. This model is admittedly crude, but for the purposes of control we are not concerned with detailed flow features, but merely the overall phase and amplitude effects, and neglecting scattering and receptivity provides a reasonable starting point.

3.4 Overall cavity model

The overall cavity transfer function P is formed from equation (1). To gain some insight into the model, first we consider some special cases. In particular, for certain choices of parameters, we recover the Rossiter formula for the frequencies of oscillation. For the shear layer model (2), suppose $G_0(s) = e^{i2\pi\gamma}$, a constant phase, and take $\tau_s = L/c_p$, with $c_p/U = \kappa$. Assuming no reflections ($r = 0$) in the acoustic model (3), the overall transfer function becomes

$$P(s) = \frac{e^{i2\pi\gamma}e^{-s\tau_s}}{1 - e^{i2\pi\gamma}e^{-s(\tau_s+\tau_a)}}$$

which has poles at $s = i\omega$, with

$$\frac{\omega L}{2\pi U} = \frac{\gamma + n}{M + 1/\kappa}, \qquad n = 1, 2, \ldots, \qquad (4)$$

which is the familiar Rossiter formula for the frequencies of oscillation. The other features of the model include the effects of longitudinal acoustic modes in the cavity (with $r > 0$), as well as amplification effects by the shear layer (with $G_0(i\omega) \neq$ const).

Parameters are chosen to make the model approximately agree with the experimental conditions at $M = 0.35$, and the resulting frequency response is shown in Figure 7. For the shear layer, equation (2)


Figure 7: Magnitude and phase of $P(i\omega)$ for cavity model, with $M = 0.35$. Also shown is the phase with a time delay of $\tau = 8$ ms removed.

is used, with $\omega_0 = 350$ Hz, $\zeta = .2$, and $\tau_s = L/c_p$, with $\kappa = c_p/U = 0.625$. (Here, $U \approx 117.5$ m/s is the freestream velocity.) For the time delay, we use an 8th-order Padé approximation to obtain a rational transfer function. The acoustics are modeled by equation (3) with $r = 0.1$ and $\tau_a = L/a$, where a is the sound speed in the freestream, and a 6th-order Padé approximation is used for the time delay. Variations in the sound speed are assumed small for this relatively low Mach number. The scattering gain is taken to be 0.2, and changing this parameter adjusts the stability of the system: for larger values of this gain, the system becomes unstable, and for smaller values, the system is more heavily damped.

The model shown in figure 7 is *stable* (all poles in the left half plane), so the magnitude of the frequency response may be viewed precisely as the amount the flow amplifies disturbances at each frequency. The peaks predicted by the model (imaginary parts of the lightly-damped poles) are at 114 Hz, 234 Hz, and 336 Hz, which correspond to the first three Rossiter frequencies. The third peak is the strongest because the shear layer amplification is the greatest for this frequency, and because the cavity acoustics reinforce oscillations at this frequency.

4 STABLE AND UNSTABLE REGIMES

As mentioned in the introduction, two possible mechanisms may lead to finite-amplitude oscillations. The conventional view[7,8] is that the system (i.e., the plant $P(s)$ from the previous section) is linearly *unstable*, so tiny perturbations will grow in

time, and eventually saturate once nonlinearities become important. An alternative view, considered recently for combustion instabilities,[1] is that the system is linearly *stable*, but lightly damped, and constantly excited by external disturbances. These disturbances are then amplified, causing oscillations at the resonant frequency of the plant, but if the disturbances were removed, the oscillations would also disappear. In this mechanism, nonlinearities may not be important: the amplitude of the oscillations is determined by the amplitude of the excitation noise, not by any saturation effects of nonlinearities. In this section, we demonstrate that the cavity may operate in either regime, depending on the Mach number.

4.1 Notion of stability

First, we clarify what we mean by stability of a system. In dynamical systems, one typically discusses stability of an equilibrium point, or a limit cycle. In fluid mechanics, one must further specify the reference frame (e.g., fixed or traveling with the fluid), and this leads to different notions of stability (e.g., convective vs. absolute).

In this paper, we view the cavity as an input-output system, and we say that the system is *stable* if the input-output map $u \mapsto y$ is bounded in L_2. For instance, suppose that the system is excited by noise for a finite amount of time, so that the input u is bounded. If the output y decays after the noise is removed, then the output will be bounded, so we regard the system as stable. If finite-amplitude oscillations persist after the noise is removed, then the output will not be bounded in L_2, so we regard the system as unstable.

4.2 Cavity flow regimes

It is difficult (and perhaps impossible) to distinguish between stable and unstable regimes using only frequency spectra. Both regimes are characterized by peaks at the resonant frequencies, and one cannot tell whether the system is in a limit cycle (with noise on top of it), or whether it is stable, merely amplifying disturbances at certain frequencies. However, it is possible to distinguish between the two regimes using the probability density function (PDF) of the output signal.[5]

If the input disturbances have a Gaussian distribution, the PDF of the stable system excited by these disturbances will also be Gaussian. By contrast, the PDF of a limit cycling system (say $y(t) = \sin(t)$) will have two peaks, because the system spends more


Figure 8: Two different Mach numbers, without control: comparison of spectra, phase portraits, and probability density functions at $M = 0.34$ (– – – –) and $M = 0.45$ (———). The $M = 0.34$ case is unstable and limit cycling; the $M = 0.45$ case is stable, driven by noise.

time near the extrema of the limit cycle. In addition, the phase portrait of a limit cycle will look like an ellipse, so the phase portrait of a noisy limit cycle will look like a "fuzzy ellipse," while the phase portrait of a stable system forced by noise will be concentrated about a point.

Measurements from the cavity experiment at two different Mach numbers are shown in Figure 8. At $M = 0.34$, the system appears to be unstable, in a limit cycle. The phase portrait indeed looks like a fuzzy ellipse, and the PDF has two distinct peaks. However, at $M = 0.45$, the system appears to be stable, driven by noise. The phase portrait is concentrated about a point, and the PDF has a single peak which closely resembles a Gaussian.

A sweep of Mach numbers from 0.1 to 0.45 revealed that $M = 0.34$ is the only Mach number where the nonlinear regime is observed. Furthermore, at this Mach number, only a single frequency

is observed, while at most other Mach numbers, multiple modes exist simultaneously. This is probably because at $M = 0.34$, the longitudinal acoustic modes in the cavity reinforce the Rossiter modes: the frequency of the first longitudinal mode coincides with the frequency of the third Rossiter mode.[10] Presumably, this reinforcement increases the amplitude of oscillations enough to cause the system to become unstable.

5 MODEL IDENTIFICATION FROM EXPERIMENT

In this section, we describe a frequency-response experiment designed to identify the cavity model directly. The experiment was performed at $M = 0.34$, where only a single frequency of oscillation was observed. Because the system appears to be unstable and limit cycling (see the previous section), nonlinearities are present and active, so a frequency response experiment on the limit cycling system would not make sense. To remove (or reduce) the effect of nonlinearities, we first stabilize the system using a feedback controller, reducing the amplitude of oscillations enough, we hope, that nonlinearities are no longer active; and then add sinusoidal disturbances to this stable, closed-loop system.

An analog controller was used to stabilize the system, and pressure measurements for the baseline and controlled cases are compared in Figure 9. From the phase portrait and the PDF, it appears that the unforced case is limit cycling, but with control the system is stable. The frequency response shows that the closed-loop system does excite oscillations at a new frequency (about 420 Hz), and we discuss these adverse effects of control later, but from the PDF it appears that these oscillations are the result of disturbance amplification, not instability.

The magnitude and phase of the measured frequency response are shown in Figure 10, along with the coherence between the input and output signals. Here, the input is the voltage to the actuator and the ouput is the pressure measured by Kulite 8, on the upstream wall of the cavity. We found that when the system was forced at a frequency near the resonant frequency (337 Hz), the data revealed a PDF with two peaks, indicating that the system again became unstable, despite the presence of the stabilizing controller. Therefore, measurements near 340 Hz are not accurate, despite the good coherence, because nonlinearities have become important. This frequency range is, of course, the most important for control analysis, as it is the frequency range where we desire good performance (i.e., reduction in dis-
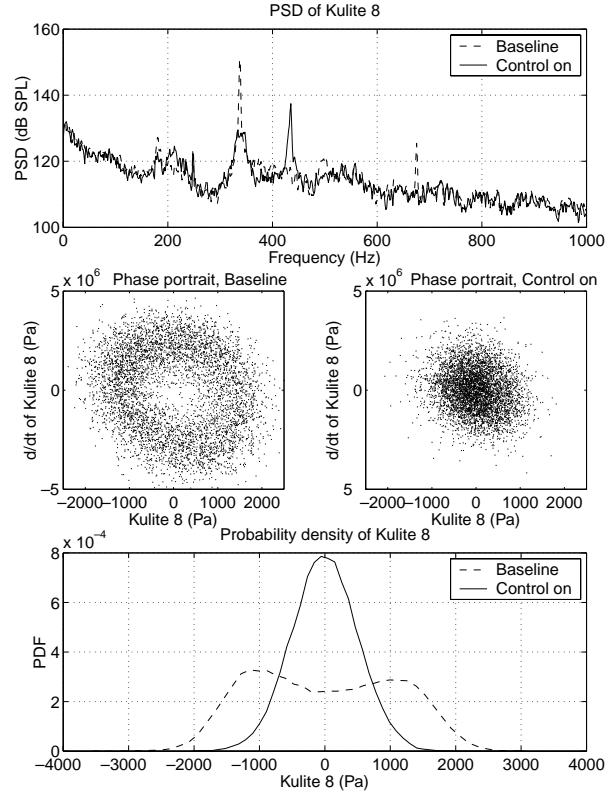

Figure 9: With and without control, $M = 0.34$: comparison of spectra, phase portraits, and probability density functions for unforced system (- - - -), and closed-loop system with gain-delay feedback (———).

Name	Order	Passband
Filter 1	2	320–360 Hz
Filter 2	2	290–390 Hz
Filter 3	2	215–465 Hz

Table 1: Parameters of digital Butterworth filters used.

Figure 10: Transfer function from input voltage to pressure measured by Kulite 8, measured from frequency response experiment at $M = 0.34$. Also shown is the phase with a time delay of $\tau = 8\,\text{ms}$ removed.

turbance amplification).

For a reliable control design, we must not have large uncertainties over the same frequency range where we desire good performance, so in this sense the frequency response experiment was not fruitful. Presumably, it would be much easier to perform such a system identification for other Mach numbers (e.g., $M = 0.45$), where nonlinearities are not important (see Figure 8). The $M = 0.34$ case at first may seem the easiest to control, because only one oscillation frequency is present, but it may actually be the hardest to control, because it is the one Mach number we have observed which is actually unstable and limit cycling.

6 CLOSED-LOOP RESULTS

Here, we present some results of different control schemes applied to the cavity at $M = 0.34$. We point out some surprising phenomena we have observed, and in the next section we present an analysis which explains these phenomena, at least qualitatively, using the physics-based model.

The control schemes all involved bandpass filters, with an adjustable gain and a time delay. We tried several digital Butterworth filters of different orders and passbands, and the orders and passbands for the filters used here are shown in Table 1.

6.1 New oscillation frequencies

Figure 11 shows the results of the closed-loop experiments with the different filters. In this figure, the gain and delay were tuned for the best suppression. The narrow band filter showed very little attenuation, and the broadband filters better attenuated the main cavity frequency, but a higher frequency peak appears. This same effect was observed with the analog controller, as can be seen from Figure 9. The frequency of this peak shifts with different controllers: with filter 2, the peak is at 430 Hz; with filter 3, 449 Hz.

(a) Filter 1 (b) Filter 2 (c) Filter 3

Figure 11: Power spectrum with 3 different digital filters, tuning gain and phase manually for best suppression: control off ($----$); control on ($-\!\!-\!\!-$).

6.2 Peak splitting

Note that, especially evident in Figure 11b, the main cavity frequency is sometimes split into two sidebands when the control is turned on. This effect is explored further in figure 12, where filter 2 was used, and the time delay was adjusted. The main resonant frequency at 337 Hz is almost completely attenuated, but sidebands appear very close in frequency, at about 320 Hz and 341 Hz. As the time delay is changed, the relative strength of the sidebands changes, and the frequency changes slightly—the lower frequency shifts from 320 to 325 Hz in figures (a)–(c). This peak-splitting phenomenon has also been observed in combustion experiments at UTRC,[1] and have been explained using the ideas presented in the next section.

7 Sensitivity function analysis

To understand the effects described in the previous section, we need to understand how feedback affects the amplification of disturbances. Without control, the transfer function from disturbances to measured pressure is simply $P(s)$ (see Figure 2). With (negative) feedback, the transfer function is $P(s)/(1+P(s)C(s))$, so the open-loop transfer function is modified by the amount

$$S(s) = \frac{1}{1+P(s)C(s)} \qquad (5)$$

called the *sensitivity function*. If $|S(i\omega)| < 1$, then disturbances are attenuated, so feedback is beneficial, but if $|S(i\omega)| > 1$, then disturbances are amplified further, and the controlled system is worse than open-loop.

The sensitivity function may determined easily from a Nyquist plot of the system, which is just


Figure 13: Nyquist plot for the feedback system, along with corresponding frequencies (in Hz). Points outside the dashed circle correspond to a performance benefit ($|S(i\omega)| < 1$), and points inside the circle correspond to a penalty ($|S(i\omega)| > 1$).

(a) Delay = 4×10^{-4} sec   (b) Delay = 5×10^{-4} sec   (c) Delay = 6×10^{-4} sec

Figure 12: Power spectra with digital controllers, showing sidebands: control off (----); control on, using filter 2 (——).

a plot of $P(i\omega)C(i\omega)$ in the complex plane, as ω varies from $0 \to \infty$. Figure 13 shows the Nyquist plot for the plant $P(s)$ given by the physics-based model from section 3, and for $C(s)$ given by filter 2, with a gain and time delay.

Graphically, the magnitude of the sensitivity function $S(i\omega)$ is just 1 over the distance from $P(i\omega)C(i\omega)$ to the -1 point. Figure 13 then demonstrates why peak splitting occurs: because $C(i\omega)$ has a narrow passband, the magnitude of $P(i\omega)C(i\omega)$ is large only over a narrow frequency range. Over this frequency range, the phase of $P(i\omega)C(i\omega)$ changes rapidly, because of the large time delay in $P(i\omega)$, and because of the steep phase change of $C(i\omega)$ in the passband. Thus, if the phase of $C(i\omega)$ is chosen to rotate the resonant frequency ($\approx 337\,\mathrm{Hz}$) far away from the -1 point, nearby frequencies will move closer to the -1 point, and feedback will amplify disturbances at these frequencies.

From Figure 13, we would therefore expect the main frequency at $337\,\mathrm{Hz}$ to be attenuated by feedback, but sidebands at about $315\,\mathrm{Hz}$ and $360\,\mathrm{Hz}$ to be amplified. These effects are confirmed by Figure 14, which shows the magnitude of the disturbance-output transfer function, with and without control. As expected, feedback attenuates the main frequency, but produces sidebands.

For a slightly smaller time delay than that used in Figure 13, the Nyquist plot will be rotated slightly counter-clockwise, so the lower frequency peak ($\approx 315\,\mathrm{Hz}$) will be amplified even more, and the higher frequency ($\approx 360\,\mathrm{Hz}$) will not be amplified as much. conversely, for a larger time delay, the Nyquist plot will be rotated clockwise, so the lower frequency peak will decrease, and the higher frequency peak will increase. These effects also agree with those measured in experiment (cf. Figure 12).


Figure 14: Magnitude of disturbance-output transfer function, without control (----), and with control (——), demonstrating peak splitting.

7.1 Performance limitations

Ideally, one would like to design a compensator $C(s)$ such that the sensitivity function $|S(i\omega)| \ll 1$ for all frequencies. Unfortunately, this is not possible, because of the *area rule*, which states that any decrease in sensitivity over one frequency range must be balanced by an increase for some other frequencies.[3] More precisely, for a system with relative degree at least 2, the area rule states that

$$\int_0^\infty \log |S(i\omega)| \, d\omega = \pi \sum_k \text{Re}(p_k), \qquad (6)$$

where p_k are the unstable poles of PC. So for a stable plant, any negative area ($|S(i\omega)| < 1$) in the log-linear plot of S versus ω must be balanced by an equal positive area ($|S(i\omega)| > 1$), as shown in Figure 15. For unstable plants, the situation is worse, and the net area must be positive.

The area rule in itself does not imply any peaking of $|S(i\omega)|$, as the positive area may be spread out over a large frequency band, as $\omega \to \infty$. However, for narrow bandwidth controllers, and plants with significant time delays, the area rule does imply a peaking of $|S(i\omega)|$, since all of the amplification must occur within the narrow bandwidth of the controller.[1] The more narrow the bandwidth, the greater the amount of peaking. This implies a strong argument in favor of large bandwidth actuators, and suggests that narrow-bandwidth actuators (such as piezoelectrics) might not be suitable for feedback control.

8 Conclusions

We have presented a linear model for cavity oscillations, incorporating the effect of external disturbances. Under some conditions, the system is unstable, so perturbations will grow until nonlinearities become important and the linear model is no longer valid. However, for other conditions, the system is stable, but lightly damped, acting as a noise amplifier. Phase portraits and probability density functions of experimental data indicate that for most flow regimes observed in our experiment, the cavity is a stable noise amplifier, oscillating at several different Rossiter modes. For the $M = 0.35$ case, however, the flow is in a limit cycle, oscillating at a single Rossiter mode.

For this Mach number, the flow may be stabilized using a controller consisting of a bandpass filter and time delay. When control is introduced, however, surprising effects are observed, including new peaks


Figure 15: Magnitude of sensitivity function $S(i\omega)$ for the model plant, with gain-delay controller. The area rule says that the area corresponding to attenuation (dark shaded region) must equal the area corresponding to amplification (light shaded region).

at different frequencies, and a peak splitting phenomenon, where the main peak splits into two side-bands. These same effects are found in the linear model. The peak splitting effect has been observed in experiments in combustion instabilities, and is a common feature of systems with limited bandwidth and large time delay.

If the noise-amplification model of cavity oscillations is correct, one cannot expect to be able to reduce the amplitude of oscillations at all frequencies using feedback, because of fundamental limitations imposed by the area rule. However, given an accurate model of the system (e.g., from a frequency response experiment as described in section 5), it is straightforward to design a compensator to minimize these adverse effects, and reduce oscillations over important frequency ranges, while paying the penalty over less important frequency ranges, or ranges where the plant itself is not so sensitive to disturbances.

References

[1] A. Banaszuk, P. G. Mehta, C. A. Jacobson, and A. I. Khibnik. Limits of achievable performance of controlled combustion processes. *IEEE Trans. Automat. Contr.*, (submitted), 2001.

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

[3] J. C. Doyle, B. A. Francis, and A. R. Tannenbaum. *Feedback Control Theory*. Macmillan, 1992.

[4] 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.

[5] I. Mezić and A. Banaszuk. Comparison of systems with complex behavior: Spectral methods. Conference on Decision and Control, Dec. 2000.

[6] 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.

[7] C. W. Rowley, T. Colonius, and A. J. Basu. On self-sustained oscillations in two-dimensional compressible flow over rectangular cavities. *J. Fluid Mech.*, to appear, 2001.

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

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

[10] D. R. Williams, D. Fabris, and J. Morrow. Experiments on controlling multiple acoustic modes in cavities. AIAA Paper 2000-1903, June 2000.

[11] D. R. Williams, C. W. Rowley, D. Fabris, T. Colonius, and R. M. Murray. Model-based control of cavity oscillations, part I: Experiments. AIAA Paper 2002-0971, Jan. 2002.