Real-time feedback control of flow-induced cavity tones—Part 2: Adaptive control [☆]

M.A. Kegerise[a,*], R.H. Cabell[b], L.N. Cattafesta III[c]

[a]*NASA-Langley Research Center, Flow Physics and Control Branch, Mail Stop 170, Hampton, VA 23681, USA*
[b]*NASA-Langley Research Center, Structural Acoustics Branch, Mail Stop 463, Hampton, VA 23681, USA*
[c]*University of Florida, Mechanical and Aerospace Engineering, P.O. Box 116250, Gainesville, FL 32611, USA*

Abstract

An adaptive generalized predictive control (GPC) algorithm was formulated and applied to the cavity flow-tone problem. The algorithm employs gradient descent to update the GPC coefficients at each time step. Past input–output data and an estimate of the open-loop pulse response sequence are all that is needed to implement the algorithm for application at fixed Mach numbers. Transient measurements made during controller adaptation at fixed Mach number revealed that the controller coefficients converged to a steady state in the mean, and this implies that adaptation can be turned off at some point with no degradation in control performance. The control algorithm demonstrated multiple Rossiter mode suppression at fixed Mach numbers of 0.275, 0.32, and 0.38, provided the plant model was updated for each Mach number. However, as in the case of fixed-gain GPC, the adaptive GPC was limited by spillover in sidebands around the suppressed Rossiter modes. The algorithm was also able to maintain suppression of multiple cavity tones as the freestream Mach number was varied over a modest range (0.275–0.29). Beyond this range, stable operation of the control algorithm was not possible due to the fixed plant model in the algorithm.
Published by Elsevier Ltd.

1. Introduction

This article is the second of two parts that describes experiments aimed at the development of a real-time adaptive controller for the cavity flow-tone problem. Of particular interest is the self-tuning adaptive controller shown in Fig. 1. The inner loop of this control structure is a dynamic feedback control system comprised of the process and controller, with a reference signal equal to zero for the present application. In the outer loop, system identification is performed to estimate the parameters of a dynamic model for the process and these are subsequently used to design the control law. Running this outer loop at the sample rate of the inner loop feedback control system allows the system to track process changes and maintain the control objective.
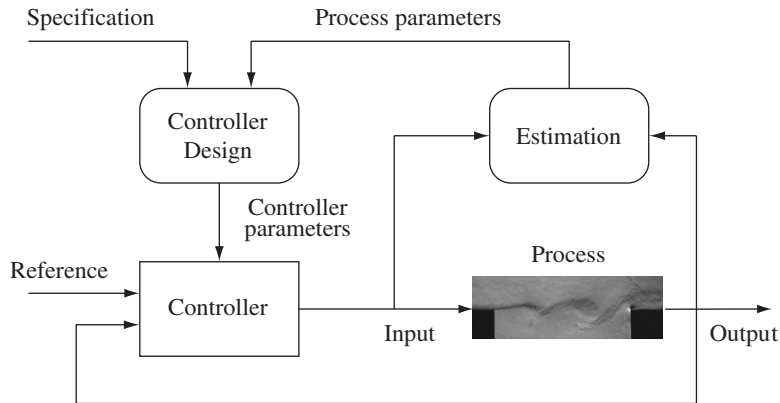
Fig. 1. Block diagram of a self-tuning adaptive controller (after Astrom and Wittenmark [1]).

As a first step towards a self-tuning adaptive controller, Part I considered the application of a dynamic feedback controller to the cavity flow-tone problem, with model identification and control design performed offline in a batch mode. That approach was based on the assumption that the cavity dynamics were linear and time invariant for any given *fixed* Mach number. The result was a series of linear, fixed-gain control laws valid only for the Mach numbers at which they were designed. The particular control design methodology used in Part I was generalized predictive control (GPC). The GPC demonstrated simultaneous suppression of multiple Rossiter modes at fixed Mach numbers ranging from 0.275 to 0.38. The suppression was found to be global since all pressure sensors within the cavity exhibited similar reductions in cavity-tone amplitudes. This success indicates that the approach taken is a sound methodology for cavity flow-tone control and sets the framework for an adaptive controller. The need for an adaptive controller is evident in the results of Part I, where small changes in the Mach number were found to degrade the performance of fixed gain control laws. The Mach number sensitivity of the cavity dynamics therefore necessitates an adaptive control design algorithm that adjusts the controller coefficients to maintain optimal suppression.

Previous attempts at adaptive control of cavity-flow tones have been reported in the literature. Cattafesta et al. [2] applied an adaptive disturbance rejection algorithm to a cavity flow at Mach 0.74 and achieved 10 dB suppression of a single Rossiter mode. Other modes in the cavity spectrum, however, were unaffected by the control system. While the disturbance-rejection algorithm should be able to suppress multiple tonal components, insufficient actuator bandwidth and authority limited the control performance to a single mode. Williams and Morrow [3] applied an adaptive feedforward scheme to the cavity flow-tone problem and demonstrated multiple cavity-tone suppression at Mach numbers up to 0.48. This was accompanied, however, by simultaneous amplification of other cavity tones. In general, a feedforward scheme that does not account for the feedback path from the actuator to the reference sensor inside the cavity will result in a poorly performing or unstable control system. Numerical simulations were used by Kestens and Nicoud [4] to study the adaptive control of cavity-flow tones. In that study, the filtered-$\times$ least-mean squares (LMS) algorithm was used to minimize the output of a single error sensor, i.e., a point measurement of pressure in the computational domain. A reduction in the pressure associated with a single Rossiter mode was achieved, but only within a small region around the error sensor. Several factors contribute to the poor control performance. First, system identification was performed with no flow, and so the process dynamics associated with the aeroacoustic feedback loop are not properly modeled. Second, the finite impulse response filters used to represent the process dynamics are unsuitable for the cavity flow [5], where lightly damped modes dominate the acoustic spectrum. Finally, similar to the feedforward scheme used by Williams and Morrow, the reference signal senses the control input. Although these past efforts have not met with great success, the application of adaptive control to the cavity flow-tone problem should not be ruled out since it is clear that hardware limitations, suboptimal control algorithms, and suboptimal process model selection/identification contributed to poor control performance.

To move further towards a self-tuning adaptive controller for the cavity flow-tone problem, this article focuses on the controller design block shown in Fig. 1. Specifically, the offline GPC design presented in Part I is extended to a recursive algorithm that updates the coefficients of a GPC control law at every time step. The coefficient updates are computed as a convolution of error sensor responses with a fixed plant model that is computed offline in a batch mode. This approach provides some adaptability to changes in the plant. Nonetheless, for a given fixed Mach number, the control performance for a control law designed online is expected to yield similar performance to a control law designed offline.

In the next section, the control methodology and adaptive GPC algorithm are presented. This is followed by a description of the experimental setup. The results section presents the transient and steady-state characteristics of the adaptive controller as applied to the cavity-flow test bed. Comparisons of fixed gain and adaptive GPC results are then made, followed by adaptive control results for a time varying Mach number.

2. Control methodology

In this section, the control methodology and adaptive GPC algorithm are presented. The overall control objective is to minimize the fluctuating pressures on the cavity walls. This is achieved by placing a piezoelectric flap actuator at the cavity leading edge with the aim to cancel shear layer instability waves due to the aeroacoustic feedback. Pressure transducers embedded in the cavity walls provide feedback signals and performance measures for the controller.

In Part I, the cavity dynamics were found to be characterized as the forced response of a lightly damped system. Therefore, an auto-regressive/exogenous-input (ARX) model structure was chosen to represent the open-loop cavity-flow dynamics [6,7]:

$$y(k) = \alpha_1 y(k-1) + \cdots + \alpha_p y(k-p) + \beta_0 u(k) + \cdots + \beta_p u(k-p), \tag{1}$$

where $y(k)$ are the $m \times 1$ wall-pressure outputs, $u(k)$ are the $r \times 1$ actuator inputs, p is the model order, and k is the current time step. The coefficient matrices, α_i $(i = 1, 2, \ldots, p)$ of $m \times m$ and β_i $(i = 0, 1, 2, \ldots, p)$ of $m \times r$, are the ARX model parameters or the observer Markov parameters. To determine the observer Markov parameters of the model, the observer-Kalman filter identification algorithm of Juang [7] was applied to input–output data collected from the baseline cavity flow. The parameters for this model structure will be different for any given fixed Mach number, but are assumed to remain constant at a given condition.

2.1. Adaptive generalized predictive control

The present objective is to develop a recursive update equation for the coefficients of a generalized predictive controller. To that end, a predictive matrix equation based on the model in Eq. (1) is formed, a cost function representing the desired control objective is defined, and a stochastic gradient-descent algorithm is used to iteratively search for the optimal control coefficients that minimize the cost function. As in Part I, the predictive matrix equation formed from the model in Eq. (1) is given by [8–11]

$$\mathbf{y}_s(k) = \mathbf{T}\mathbf{u}_s(k) + \mathbf{\Theta}\mathbf{v}_p(k-p), \tag{2}$$

where $\mathbf{y}_s(k)$ is an $sm \times 1$ vector of current and future wall-pressure outputs:

$$\mathbf{y}_s(k) = \left\{ \begin{array}{c} y(k) \\ y(k+1) \\ \vdots \\ y(k+s-1) \end{array} \right\}, \tag{3}$$

$\mathbf{u}_s(k)$ is an $sr \times 1$ vector of current and future actuator inputs:

$$\mathbf{u}_s(k) = \begin{Bmatrix} u(k) \\ u(k+1) \\ \vdots \\ u(k+s-1) \end{Bmatrix}, \tag{4}$$

and $\mathbf{v}_p(k-p)$ is a $p(m+r) \times 1$ vector of past inputs and outputs running from time step $k-p$ to $k-1$:

$$\mathbf{v}_p(k-p) = \begin{Bmatrix} u(k-p) \\ \vdots \\ u(k-1) \\ y(k-p) \\ \vdots \\ y(k-1) \end{Bmatrix}. \tag{5}$$

The parameter s is referred to as the prediction horizon. The predictive matrix equation states that the future output data depends on the future control inputs and past input–output data. The matrix $\mathbf{T}$ is an $sm \times sr$ Toeplitz matrix:

$$\mathbf{T} = \begin{bmatrix} \beta_0 & & & & \\ \beta_0^{(1)} & \beta_0 & & & \\ \beta_0^{(2)} & \beta_0^{(1)} & \beta_0 & & \\ \vdots & \vdots & \ddots & \ddots & \\ \beta_0^{(s-1)} & \beta_0^{(s-2)} & \cdots & \beta_0^{(1)} & \beta_0 \end{bmatrix}, \tag{6}$$

where β_0, $\beta_0^{(1)}, \ldots, \beta_0^{(s-1)}$ are the pulse response parameters of the open-loop plant. These parameters can be obtained from the observer Markov parameters [8]. The rectangular matrix $\mathbf{\Theta}$ is formed with a set of recursive equations and the observer Markov parameters [8]. As will be shown below, however, this matrix is not needed in the real-time implementation of the adaptive controller.

Previously, the vector for current and future control inputs, $\mathbf{u}_s(k)$, was proportional to the vector of past inputs and outputs, $\mathbf{v}_p(k-p)$. That is, the generalized predictive controller designed in an offline batch mode can be represented by

$$\mathbf{u}_s(k) = \mathbf{H}\mathbf{v}_p(k-p), \tag{7}$$

where $\mathbf{H}$ denotes the $sr \times p(m+r)$ matrix of controller coefficients. While the coefficients in that case were constant values for a given Mach number, the development below will result in a recursive equation that updates the values of $\mathbf{H}$ at each time step, k, always seeking to minimize a desired cost function.

With the control objective being disturbance rejection at the system output, or minimization of wall-pressure output, an appropriate cost function can be defined as

$$J = \mathbf{y}_s^T(k)\mathbf{Q}\mathbf{y}_s(k) + \gamma \operatorname{tr}(\mathbf{H}^T\mathbf{H}), \tag{8}$$

where J is a scalar value, γ is a scalar weighting factor, and $\operatorname{tr}(\cdot)$ denotes the trace of the matrix in parentheses. The first term in Eq. (8) is the sum of the weighted, squared output values over the prediction horizon and $\mathbf{Q}$ is

a $sm \times sm$ block-diagonal matrix of sensor weights:

$$\mathbf{Q} = \begin{bmatrix} Q_1 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & Q_s \end{bmatrix}, \tag{9}$$

where Q_i is an $m \times m$ diagonal matrix of sensor weights:

$$Q_i = \mathrm{diag}(q_1, q_2, \ldots, q_m). \tag{10}$$

The sensor weights $q_1, \ldots, q_m$ take on values between 0 and 1. A value of zero means that the sensor is not included in the cost function, but the sensor information will still be used by the feedback controller.

The second term in the cost function penalizes the controller coefficients and, as will be shown below, gives rise to a leakage factor in the update algorithm. This penalty term differs from the one used in Part I, where an explicit penalty term on the control input was used in the cost function, i.e., $\mathbf{u}_s^{\mathrm{T}} \lambda \mathbf{u}_s$. Both penalty terms have the similar effect of imposing an effort penalty on the control input to avoid large control inputs and actuator saturation. When using the present cost function defined in Eq. (8), however, the resulting update equation is easier to implement and less computationally expensive. With the assumed model given in Eq. (2), and assuming the plant is time-invariant, solving for the optimal controller becomes a linear algebra problem. The result is a fixed-gain feedback controller which is optimal for a fixed Mach number. This is essentially a receding horizon control law [12], and its use was discussed in Part I.

To enable an adaptive controller that maintains good performance while the Mach number is changing, a stochastic gradient descent algorithm is used to minimize the cost function in Eq. (8). This algorithm updates the controller coefficients at each time step with

$$\mathbf{H}(k+1) = \mathbf{H}(k) - \mu \frac{\partial J(k)}{\partial \mathbf{H}(k)}, \tag{11}$$

where the gradient of J is an instantaneous or stochastic estimate and μ sets the adaptive rate of the algorithm. This equation is the basis for the well known LMS algorithm [13,14].

An instantaneous estimate of the gradient of the cost function can be determined by substituting Eqs. (2) and (7) into Eq. (8) and taking the derivative with respect to $\mathbf{H}(k)$. After some algebraic manipulation:

$$\frac{\partial J(k)}{\partial \mathbf{H}(k)} = 2\mathbf{T}^{\mathrm{T}} \mathbf{Q} \mathbf{y}_s(k) \mathbf{v}_p^{\mathrm{T}}(k-p) + 2\gamma \mathbf{H}(k). \tag{12}$$

This gradient estimate assumes $\partial \mathbf{v}_p(k-p)/\partial \mathbf{H}(k)$ is zero, which means the derivatives of past inputs and outputs with respect to the current weights are assumed to be zero. A similar approximation has been used to derive other adaptive filter structures, such as Feintuch's IIR-LMS algorithm [15], Eriksson's filtered-u algorithm [16], and adaptive filters used in active noise control [17]. Like the current algorithm, these other adaptive algorithms are computationally simple and appear to have a self-stabilizing property [17]. Nonetheless, this approximation needs to be examined in the context of the current feedback filter structure as part of a more thorough statistical analysis of the convergence properties of this algorithm.

Substituting this expression into Eq. (11) gives

$$\mathbf{H}(k+1) = \alpha \mathbf{H}(k) - 2\mu \mathbf{T}^{\mathrm{T}} \mathbf{Q} \mathbf{y}_s(k) \mathbf{v}_p^{\mathrm{T}}(k-p), \tag{13}$$

where $\alpha = (1 - 2\mu\gamma)$. Examination of the update equation indicates that it depends on future input values, $\mathbf{u}_s(k)$, and future output values, $\mathbf{y}_s(k)$, and so cannot be implemented in real time. However, assuming the

gradient of the cost function does not change significantly over the prediction horizon, s, the data vectors can be shifted back in time by s time steps and the update equation becomes:

$$\mathbf{H}(k + 1) = \alpha\mathbf{H}(k) - 2\mu\mathbf{T}^{\mathrm{T}}\mathbf{Q}\mathbf{y}_s(k - s)\mathbf{v}_p^{\mathrm{T}}(k - p - s). \tag{14}$$

The variable α in the update equation is referred to as the leakage factor and, as noted earlier, it arises from the $\gamma\,\mathrm{tr}(\mathbf{H}^{\mathrm{T}}\mathbf{H})$ term in the cost function [13,18,19]. The value of α is typically set to slightly less than one (e.g., $\alpha = 0.9999$). Including leakage in the update equation results in a small degradation in the nominal performance that would otherwise be achieved if γ was set to zero in the cost function (Eq. (8)). On the other hand, including leakage in the update equation has several benefits that outweigh this performance degradation. Leakage improves robustness and rate of convergence [13,14,19], counteracts the parameter drift that can occur in the standard LMS algorithm [13], and most importantly, leakage acts like an effort penalty on the control input [18,19].

The control law in Eq. (7) yields a vector of current and future inputs. In the real-time implementation of this control law, however, the current control effort, $u(k)$, is applied to the process and the future values are discarded at each time step. Therefore, it is sufficient to compute the current control effort as

$$u(k) = \mathbf{h}\mathbf{v}_p(k - p), \tag{15}$$

where $\mathbf{h}$ denotes the first r rows of the matrix $\mathbf{H}$. In turn, the update equation can be rewritten to update the $r \times p(m + r)$ matrix $\mathbf{h}$ as

$$\mathbf{h}(k + 1) = \alpha\mathbf{h}(k) - 2\mu\{\mathbf{T}^{\mathrm{T}}\mathbf{Q}\}_r\mathbf{y}_s(k - s)\mathbf{v}_p^{\mathrm{T}}(k - p - s), \tag{16}$$

where the subscript r denotes the first r rows of the matrix $\mathbf{T}^{\mathrm{T}}\mathbf{Q}$.

Eqs. (15) and (16) are the key equations for implementation of the adaptive GPC algorithm. No assumptions have been made with regard to the number of inputs and outputs in the development of these equations. Therefore, they can be applied to general multiple input–output systems. A listing of the control algorithm for real-time implementation is shown below. As this listing indicates, all that is needed to implement this algorithm is an estimate of the observer Markov parameters and measurements of input–output data. The primary advantage of this algorithm is that it is computationally simple and therefore consumes a relatively small portion of the digital signal processor turnaround time during each sample update. Since the eventual goal for cavity-tone control is to include online system identification, the low computational cost for updating the controller gains is a definite advantage.

There are several parameters in the algorithm that must be tuned to achieve a balance between optimal performance and stability of the controller: the model order p, the prediction horizon s, the adaptive rate μ, the leakage factor α, the sensor weights Q_i, and the sampling rate. Although the parameter values are problem dependent, experience provides some guidelines for their selection. The model order is selected through system identification. The key issue in this case is to choose p large enough such that all of the pertinent open-loop dynamics are captured by the model. The prediction horizon should be at least equal to the model order, but in practice, is typically taken as 2–3 times the model order [10]. For the sampling rate, experience indicates that a value 2–3 times the highest frequency results in the best performance [10]. The highest frequency in the present context is the highest Rossiter frequency of interest.

The parameter μ controls the convergence rate of the algorithm. As the value of μ is increased, the convergence time of the algorithm decreases, but for too large a value, the algorithm will become unstable. Ideally, μ should be chosen such that the convergence time of the algorithm is smaller than the time scale over which the process dynamics change.

As mentioned above, the numerical value of α is typically slightly less than 1. As the leakage factor approaches 1, the control effort penalty decreases, resulting in a more aggressive controller. It is important to recall that α also depends on μ; i.e., $\alpha = 1 - 2\mu\gamma$. For example, if μ is decreased but α remains the same, the implication is that γ has increased. A larger value of γ results in a larger control-effort penalty and therefore, a more sluggish, stable controller. The upshot of this interdependence is μ and α must both be tuned to achieve a balance between the optimal performance, convergence rate, and algorithm stability.

Algorithm listing

System identification
Determine the model parameters and form the matrix, $\{\mathbf{T}^T\mathbf{Q}\}_r$
Initialization
Initialize the algorithm by setting

$$\mathbf{h}(0) = \mathbf{0}$$

For each time step
Update the data vectors $\mathbf{y}_s(k-s)$ and $\mathbf{v}_p(k-p-s)$
Update the controller coefficients

$$\mathbf{h}(k+1) = \alpha\mathbf{h}(k) - 2\mu\{\mathbf{T}^T\mathbf{Q}\}_r\mathbf{y}_s(k-s)\mathbf{v}_p^T(k-p-s)$$

Compute the current control effort
$$u(k) = \mathbf{h}(k)\mathbf{v}_p(k-p)$$

---

3. Experimental details

The experimental hardware used for the adaptive GPC experiments was the same as that described in Part I. All experiments were performed in the NASA-Langley Research Center Probe Calibration Tunnel. The rectangular cavity model was fixed to an aspect ratio of $\ell/d = 5$. As before, the cavity model was instrumented with a pair of piezoresistive pressure transducers; one sensor located in the front cavity wall centerline 12 mm down from the leading edge and the second located on the floor centerline 15 mm upstream from the rear wall. The actuator was a piezoelectric bimorph cantilever beam with its tip situated at the cavity leading edge. The actuator response was found to be characteristic of a 2nd-order under-damped system with a natural frequency of ≈ 1200 Hz and a DC gain of $\approx 0.25\,\mu\text{m/V}$. The control hardware consisted of anti-alias filters, a 16 bit A/D, a floating-point digital signal processor, a 14-bit D/A, and a reconstruction filter. The adaptive GPC algorithm was coded and compiled to run on the digital signal processor at a sample-time interval of 250 μs. Further details on the experimental hardware and setup can be found in Part I.

4. Results and discussion

The results of the control experiments are presented in this section. Three freestream Mach numbers were considered: $M_\infty = 0.275$, 0.32, and 0.38. Throughout the tests, the cavity geometry was fixed to an aspect ratio of $\ell/d = 5$, the total pressure was 138 kPa, and the total temperature was 297 K. The Reynolds number based on cavity length ($Re_\ell = \ell U_\infty/v$)) for the three conditions ranged from 1.2×10^6 to 1.6×10^6. Baseline pressure spectra for the three test conditions are presented in Fig. 2. A single-input/two-output, ARX model of the cavity dynamics was formed for each Mach number using experimental input–output data and the observer-Kalman identification algorithm of Juang [7]. A model order of $p = 80$ was used for all test conditions. Data were collected to describe both the transient behavior of the controller, while the control coefficients were converging, and the steady-state behavior, after the coefficients had converged.

4.1. Transient behavior of the control algorithm

The time evolution of the closed-loop performance of the algorithm was examined through measurement of unsteady pressures and controller coefficients while the control algorithm was adapting. Unsteady pressure time series were collected in a pre-/post-trigger mode, where the trigger condition was control on. The adaptive controller coefficients were also sampled at every 100th time step. Ten triggered data sets were acquired for
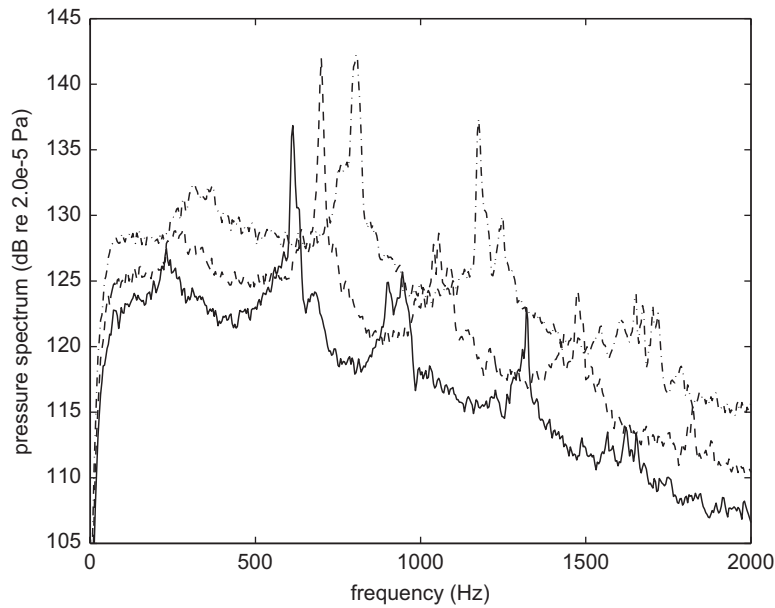
Fig. 2. Baseline pressure spectra of rear-floor sensor for $\ell/d = 5$ cavity. — $M_\infty = 0.275$, – – – $M_\infty = 0.32$, – · – $M_\infty = 0.38$.

each test condition so stochastic variations in convergence behavior could be averaged out. In each case, the initial controller coefficients were set to zero.

The short-time Fourier transform, which provides a measure of time evolution in the frequency domain, was estimated for each of the 10 triggered data sets and block averaged to reduce random uncertainty in the estimate. The short-time Fourier transform of a pressure signal is defined as [20]

$$\text{STFT}(t, \omega) = \int p(\tau) w^*(\tau - t) \mathrm{e}^{-\mathrm{j}\omega\tau} \, \mathrm{d}\tau, \tag{17}$$

where $*$ denotes the complex conjugate and $w(t)$ is a window function that has a short duration. The short-time Fourier transform is typically visualized by plotting the square of the magnitude as a contour plot. This "spectrogram" represents the time-dependent power spectrum of the time series. The spectrogram was estimated by dividing the time series into several overlapping blocks, each one advancing in time. The blocks were then multiplied by a hanning window function and Fourier transformed. The time and frequency resolution of the short-time Fourier transform was 13 ms and 7.8 Hz, respectively. The short-time Fourier transform was estimated for each of the 10 triggered data sets and then block averaged to reduce the random uncertainty in the estimate.

Spectrograms of the rear-floor pressure at $M_\infty = 0.275$ for two different values of the adaptation parameter are shown Figs. 3 and 4. All algorithm parameters were the same between these two runs except the adaptation parameter, μ, was ten times higher in Fig. 3 than in Fig. 4, and the leakage factor, α, was adjusted accordingly for the two cases to maintain the same value for the control effort penalty factor, γ. In each plot the vertical line at 0 s denotes control on. The spectral lines associated with the first four Rossiter modes are clearly evident in the figures. After control is on, as the control coefficients adapt, the amplitude of each Rossiter mode is reduced, eventually reaching an asymptotic level. An examination of the sampled controller coefficients reveals that they too reach a steady-state level in the mean. However, the coefficients do exhibit stochastic fluctuations about these mean levels.

The primary difference between the spectrograms is the rate at which the Rossiter modes reach their final suppressed levels. In fact, for the smaller μ value data presented in Fig. 4, it takes an order of magnitude longer for all Rossiter modes to reach their final state. To see this more clearly, the values of the two spectrograms at a particular frequency versus time were plotted. The spectrogram data at $f = 613$ Hz, plotted in Fig. 5, represents the time evolution of Rossiter mode 2 for the two adaptive rate values. The baseline level of the
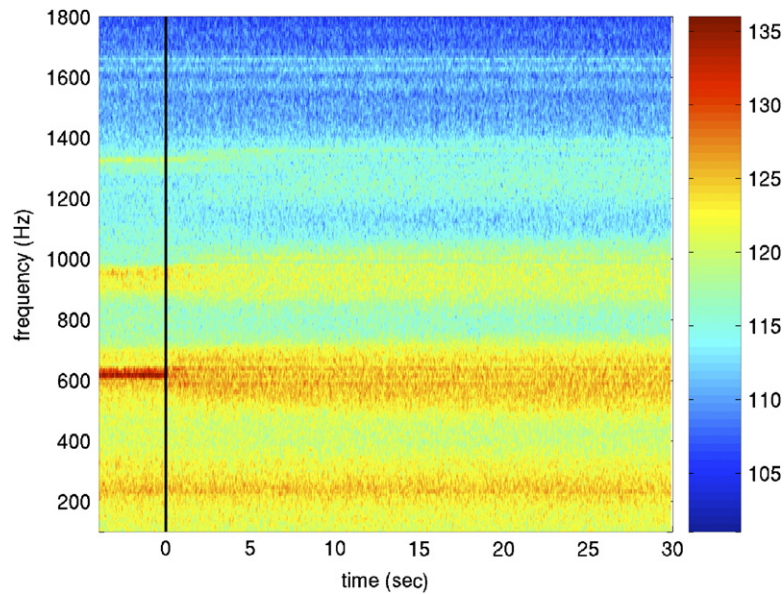
Fig. 3. Spectrogram of the rear-floor pressure sensor at $M_\infty = 0.275$ and $\mu = 1 \times 10^{-4}$. Adaptive controller was turned on at $t = 0$. Colormap scale is in dB re 2×10^{-5} Pa.
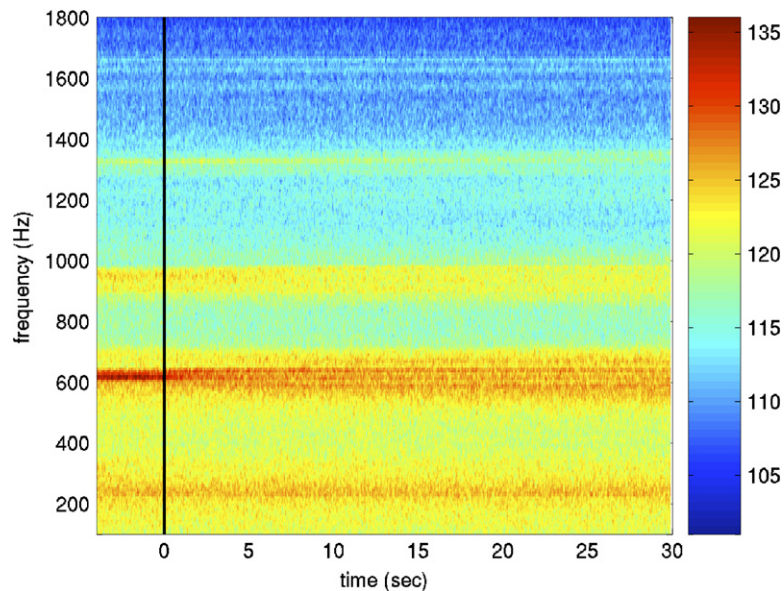

Fig. 4. Spectrogram of the rear-floor pressure sensor at $M_\infty = 0.275$ and $\mu = 1 \times 10^{-5}$. Adaptive controller was turned on at $t = 0$. Colormap scale is in dB re 2×10^{-5} Pa.

spectrogram has been shifted so that it reads 0 dB prior to control on. For $\mu = 1 \times 10^{-4}$, Rossiter mode 2 is suppressed by approximately 10 dB within a few seconds of control on. Decreasing the value of μ by an order of magnitude increases the convergence time by the same factor, but the final level of suppression is essentially the same.

In Part I, the input sensitivity function was shown to provide a good measure of control performance. To see if this is the case for the adaptive controller, the sampled controller coefficients and the open-loop plant
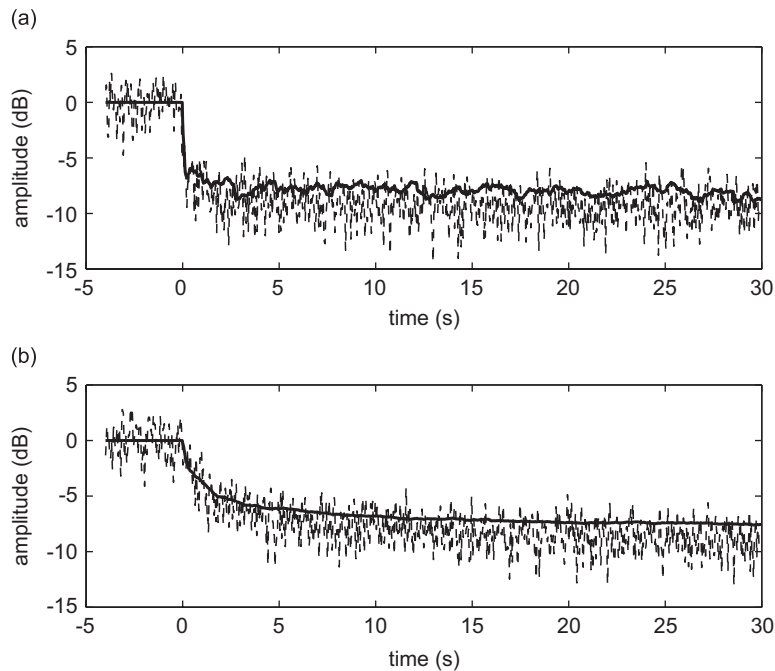
Fig. 5. Time evolution of Rossiter mode 2 amplitude (– – –) for two adaptive rates: (a) $\mu = 1 \times 10^{-4}$, (b) $\mu = 1 \times 10^{-5}$. Controller was turned on at $t = 0$. The log magnitude of the input sensitivity at $f = 613\,\text{Hz}$ (—) is plotted for comparison. $M_\infty = 0.275$.

model were used to compute the input sensitivity at $f = 613\,\text{Hz}$ for every 100th time step. In doing so, a key assumption is that the controller coefficients are "frozen" at their current values for any given time step. This was performed for each of the 10 data sets and then block averaged. The results are shown in Fig. 5 for comparison to the spectrogram data. Like the measured mode amplitude, the input sensitivity also reaches a steady-state level. This occurs because the controller transfer function converges to a steady state. The convergence of the controller transfer function to a steady-state suggests that adaptation can be stopped after some time to freeze the controller coefficients with no degradation in performance. Pressure spectra measured with a fixed set of converged controller coefficients confirmed this.

It is of interest to note that the averaged sensitivity data presented in Fig. 5 is much smoother than the averaged short-time Fourier transform data. This is because the adaptive GPC, due to its recursive nature, inherently averages noise in the pressure responses. The controller coefficients, although they do vary stochastically, will vary much less than the actual unsteady pressures, for most reasonable values of adaptation rate, μ. The relationship between S_i variance and μ can be observed in Fig. 5, where lower values of μ exhibit less variance in the input sensitivity.

The input sensitivity is observed to closely match the trend of the measured mode amplitude with time and therefore, it can be used to study the transient control performance at other Rossiter frequencies. In Fig. 6, the input sensitivity at Rossiter modes 1 through 4 for the $M_\infty = 0.275$ case is plotted for two adaptive rates. The plot demonstrates that the different Rossiter modes are suppressed at different rates. The dominant mode is suppressed first, followed by the second largest mode, and so on. Nevertheless, given sufficient time, all Rossiter modes reach a steady-state level of suppression for a fixed Mach number.

4.2. Steady-state control performance

As discussed in the previous section, the controller coefficients were found to converge to steady-state values in the mean. Repeated runs of the adaptive control algorithm at a fixed Mach number further revealed that the controller coefficients converge to the same values. Once the controller converged, time series of the unsteady
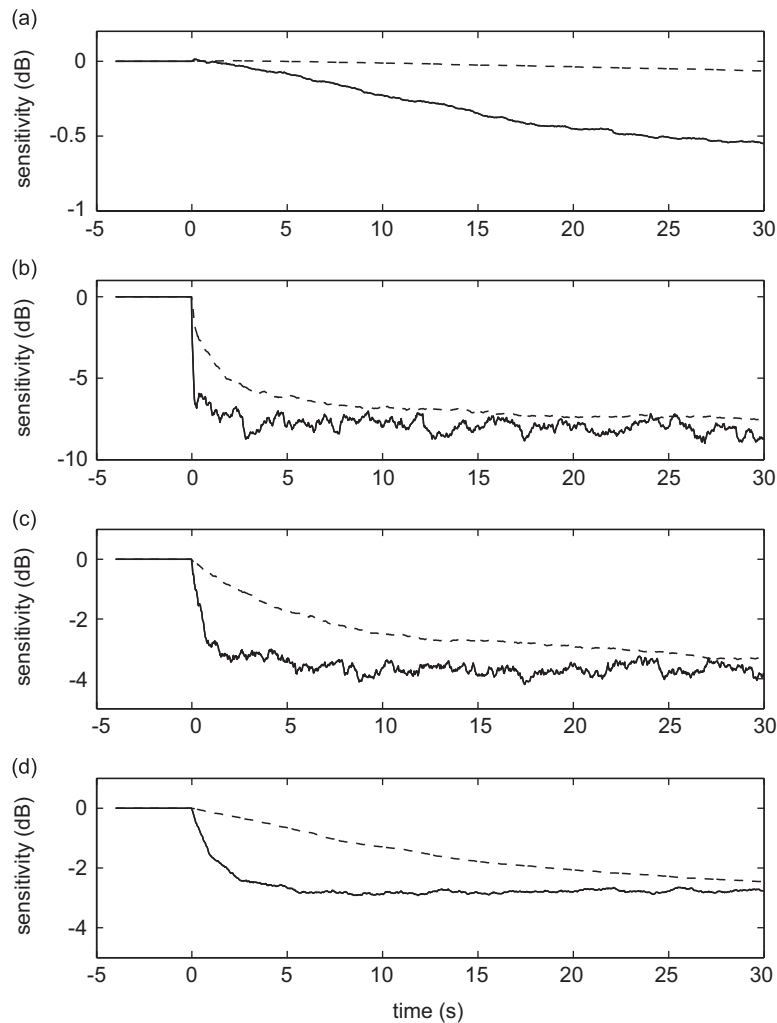
Fig. 6. Time evolution of input sensitivity for two adaptive rates (— $\mu = 1.0 \times 10^{-4}$, $- - -$ $\mu = 1.0 \times 10^{-5}$): (a) Rossiter mode 1 at 240 Hz, (b) Rossiter mode 2 at 613 Hz, (c) Rossiter mode 3 at 945 Hz, (d) Rossiter mode 4 at 1320 Hz. $M_\infty = 0.275$.

wall pressures were collected and pressure spectra were computed. The spectra were computed using 1024 point FFTs, a hanning window, 50% overlap, and 160 block averages. The frequency resolution for the power spectrum estimates was 3.9 Hz.

Control results for the three Mach number conditions are presented in Figs. 7a, 8a, and 9a. The figures show baseline and controlled pressure spectra measured at the rear-floor sensor. These results are representative of the best performance achieved for each flow condition, given several runs of the controller for a range of algorithm parameters. Table 1 lists the algorithm parameters for the results shown in Figs. 7–9. For all cases, the model order used in system identification was 80 and the prediction horizon was set to 240. The sensor weights for the two outputs were zero for the front sensor (q_1) and one for the rear sensor (q_2). The control performance was found to be weakly dependent on the sensor weights, and only subtle differences were observed with different values of sensor weightings. At each flow condition the convergence rate, μ, and leakage parameter, α were adjusted to achieve the best performance while maintaining stability and avoiding actuator saturation. Multiple Rossiter modes were suppressed by the control algorithm at all three conditions. For the $M_\infty = 0.275$ case in particular, the first four Rossiter modes exhibit some suppression. The broadband levels of the pressure fluctuations are not significantly altered in any of the control runs.
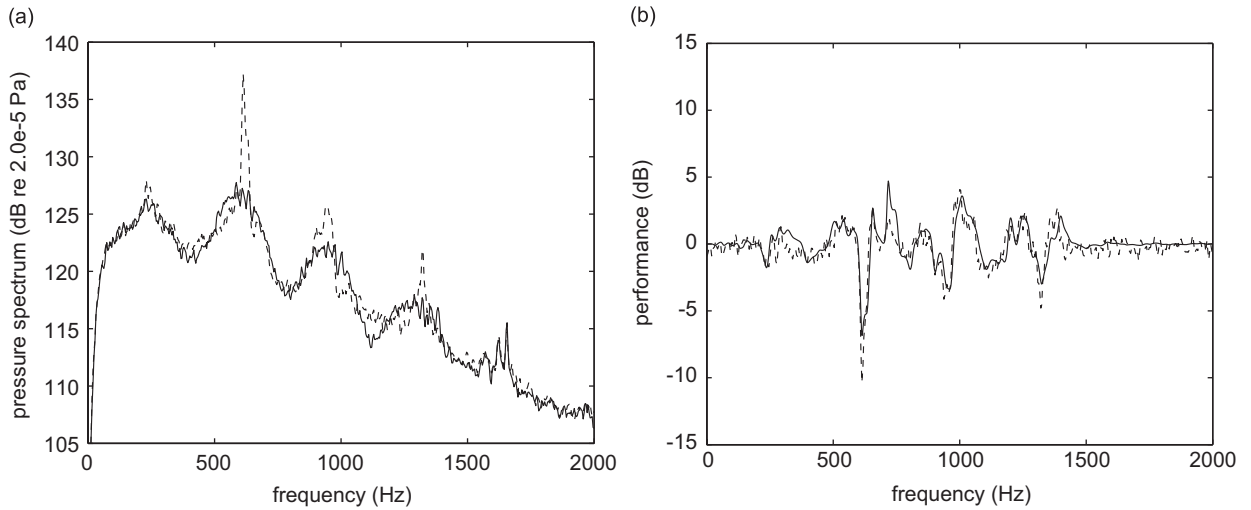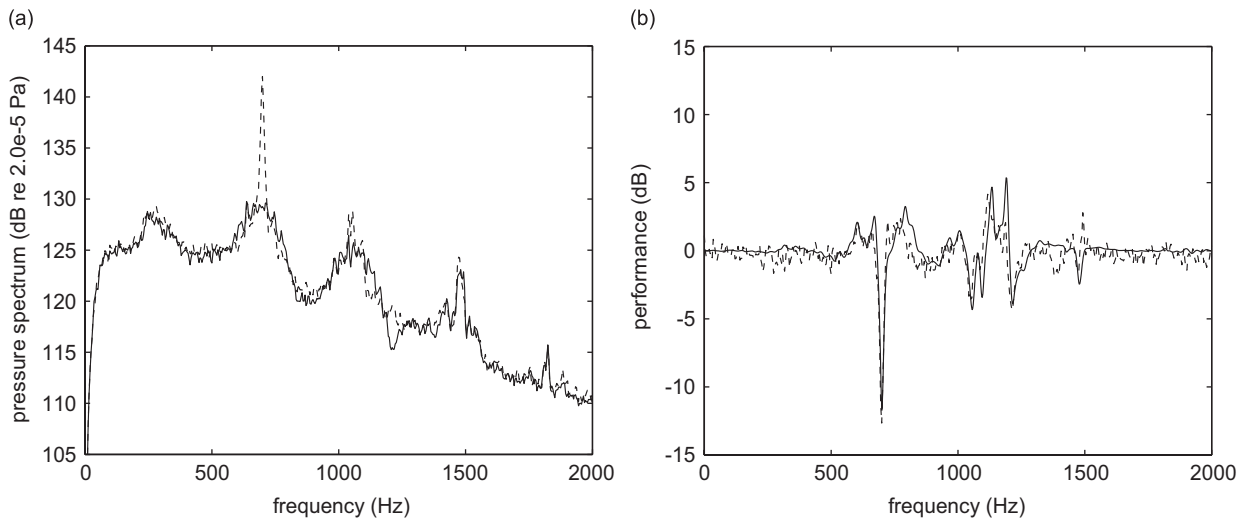
Fig. 7. Adaptive GPC results at $M_\infty = 0.275$: (a) Baseline (– – –) and controlled (—) pressure spectra measured at rear-floor sensor, (b) measured performance (– – –) and input sensitivity (—) for GPC.

Fig. 8. Adaptive GPC results at $M_\infty = 0.32$: (a) Baseline (– – –) and controlled (—) pressure spectra measured at rear-floor sensor, (b) measured performance (– – –) and input sensitivity (—) for GPC.

In Part I, a performance measure was defined as

$$\frac{\sqrt{\mathbf{Y}_{cl}^H \mathbf{Y}_{cl}}}{\sqrt{\mathbf{Y}_{ol}^H \mathbf{Y}_{ol}}}, \tag{18}$$

where $\mathbf{Y}_{cl}$ and $\mathbf{Y}_{ol}$ are vectors of the Fourier transform for the controlled and baseline cases, respectively, and H denotes the complex-conjugate transpose. Eq. (18) essentially provides a scalar measure of disturbance rejection for the multiple output sensors, where a value less than one indicates disturbance attenuation and a value greater then one indicates disturbance amplification. The performance measure was calculated for each of the control cases and the results are represented by the dashed lines in Figs. 7b, 8b, and 9b.
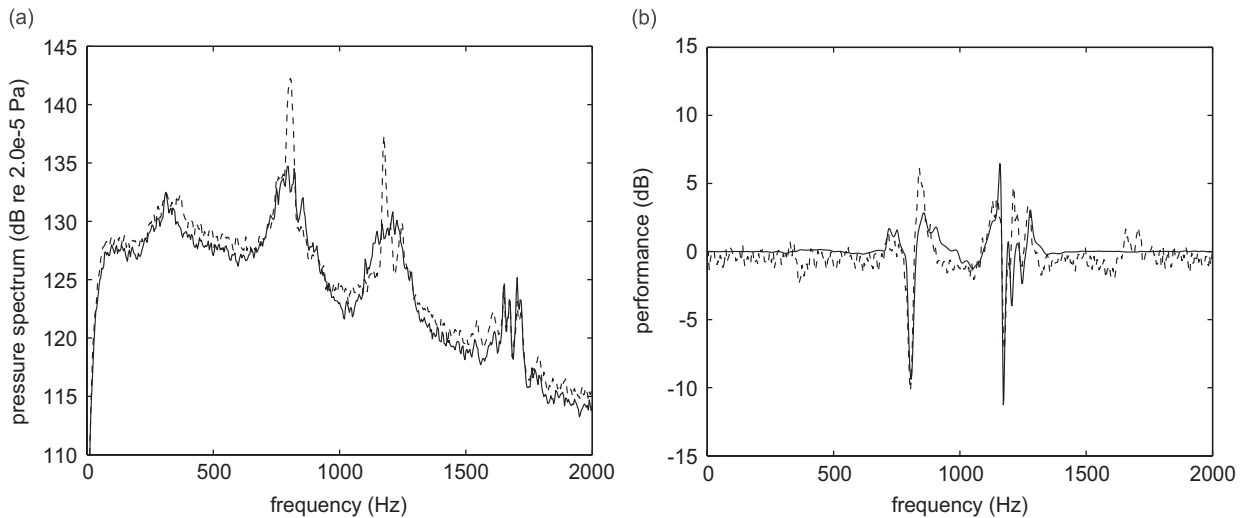
(a)

(b)

Fig. 9. Adaptive GPC results at $M_\infty = 0.38$: (a) Baseline $(- - -)$ and controlled $(—)$ pressure spectra measured at rear-floor sensor, (b) measured performance $(- - - -)$ and input sensitivity $(—)$ for GPC.

Table 1
Adaptive GPC algorithm parameters

M_∞	Model order p	Pred. horizon s	Sensor weights		Leakage factor α	Adaptive rate μ
			q_1	q_2		
0.275	80	240	0	1	0.999995	1×10^{-4}
0.32	80	240	0	1	0.99999	1×10^{-5}
0.38	80	240	0	1	0.999999	1×10^{-6}

Like the fixed-gain control results, the performance measure for adaptive GPC is less than one at the Rossiter modes, but is always accompanied by amplification in sideband frequencies, i.e., spillover is present.

Using the converged controller coefficients and the open-loop plant model, the input sensitivity function was computed for each of the test cases. The results are presented in Figs. 7b, 8b, and 9b. Again, the input sensitivity is found to be in close agreement with the performance measurement. The character of these control-performance measurements is as expected in view of the results presented in Part I. The physical collocation of the actuator with the disturbance, in part, gives rise to the spillover that is observed in the control results.

Previous control results obtained with fixed-gain GPC were based on the assumption of linear and time invariant cavity dynamics for a given fixed Mach number. That assumption was used in the present development and application of an adaptive GPC algorithm. The agreement between the measured control performance and input sensitivity function, both in the transient and steady-state cases, tends to further support these assumptions. It is also of interest to compare the performance of adaptive GPC to that of the fixed-gain GPC designed offline. Fixed-gain GPC and adaptive GPC results for $M_\infty = 0.275$ are compared in Figs. 10a–c. The pressure spectra for the two cases are nearly identical, both yielding suppression of the first four Rossiter modes. The measured performance and input sensitivity are also observed to be in close agreement, suggesting that the controller transfer functions for the two cases are essentially the same. Examination of these performance measures at other test Mach numbers indicates similar agreement.
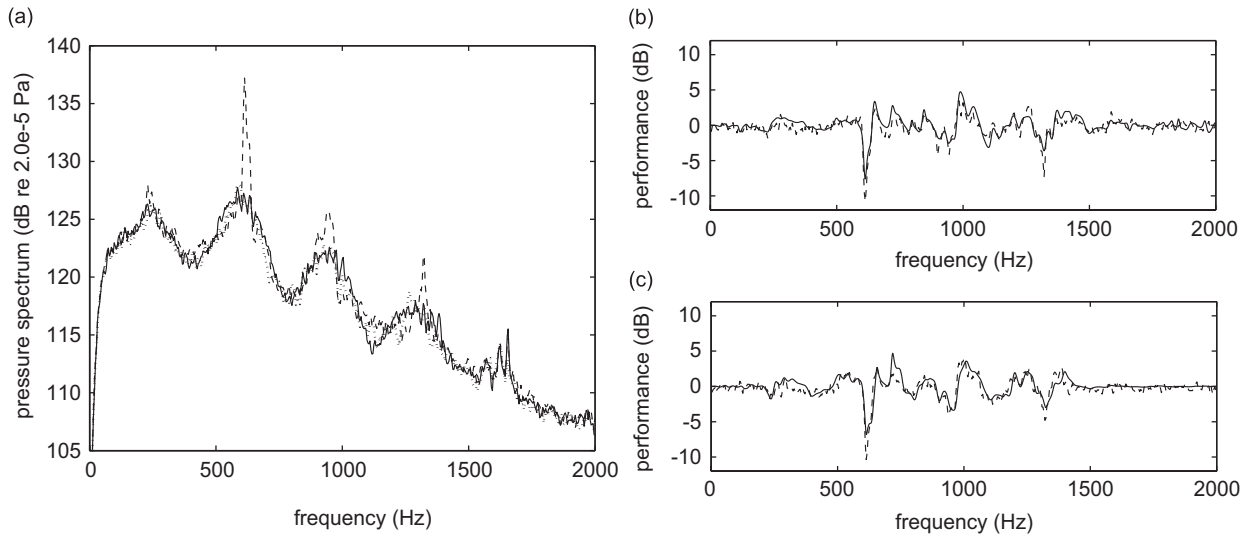
Fig. 10. Comparison of adaptive GPC and fixed-gain GPC control results at $M_\infty = 0.275$: (a) baseline $(- - -)$, adaptive GPC $(—)$, and fixed-gain GPC $(\cdots)$ pressure spectra at rear-floor sensor, (b) measured performance $(- - -)$ and input sensitivity $(—)$ for fixed-gain GPC, (c) measured performance $(- - -)$ and input sensitivity $(—)$ for adaptive GPC.

4.3. Adaptation to time varying Mach number

The adaptive GPC algorithm presented in this paper requires an estimate of the pulse response sequence of the open-loop plant. For a fixed Mach number, the results of the previous section indicate that the pulse response sequence can be determined offline from open-loop, input–output data and then subsequently used in the real-time control algorithm. If the Mach number changes, however, the open-loop dynamics change. Then, the open-loop dynamics will have to be identified with a recursive algorithm. This has not yet been implemented in our experimental program. It is still of interest to consider the robustness of the adaptive algorithm to changes in the freestream Mach number without re-identification of the system. To that end, system identification was performed at $M_\infty = 0.275$. The control algorithm was then run at that Mach number and the controller was allowed to converge. The Mach number was then slowly increased to 0.28, while the controller continued to run, and time-series data were collected. Then, the Mach number was increased to 0.29 and data were again collected. Beyond this Mach number, a stable controller could not be maintained.

Baseline and controlled pressure spectra measured at the rear-floor sensor for $M_\infty = 0.28$ and 0.29 are shown in Fig. 11. For this small change in Mach number ($\sim 9\%$) the adaptive GPC is observed to maintain suppression in the first four Rossiter modes. For comparison, results for a fixed-gain GPC controller designed at $M_\infty = 0.275$ are included in Fig. 11. That controller was designed with the methodology outlined in Part I, with $p = 80$, $s = 240$, $q_1 = 0$, $q_2 = 1$, and $\lambda = 0.05$. For $M_\infty = 0.28$, the performance of the fixed-gain algorithm is similar to that for the adaptive algorithm. At $M_\infty = 0.29$, however, the performance of the fixed-gain algorithm begins to degrade, and beyond that Mach number, the controller became unstable.

Previous authors [21] have shown that for gradient descent algorithms where computation of the gradient requires a model of the plant from actuator to error sensor, a sufficient condition for convergence is that the ratio of the plant model to the true plant be strictly positive real at all frequencies. For the current application, this condition can be written as

$$\phi(\omega) = \text{Re}[\hat{G}_{0.275}(j\omega)G_M^\star(j\omega)] > 0 \quad \forall \omega, \tag{19}$$

where $G_M^\star(j\omega)$ denotes the complex conjugate of the true plant from actuator to error sensors at the Mach number M, and $\hat{G}_{0.275}(j\omega)$ denotes the transfer function model for $M_\infty = 0.275$, which was the model used in the gradient update equation.
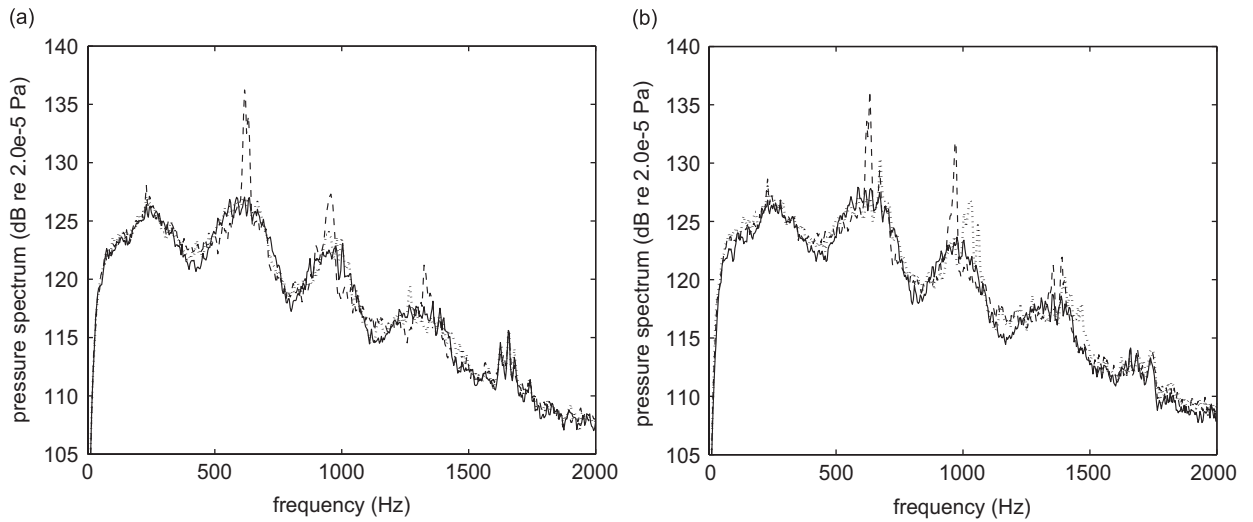
Fig. 11. Adaptive GPC control performance as freestream Mach number is increased: (a) $M_\infty = 0.28$, (b) $M_\infty = 0.29$. System identification was performed at $M_\infty = 0.275$. — — — Baseline, — adaptive GPC, $\cdots$ fixed-gain GPC.
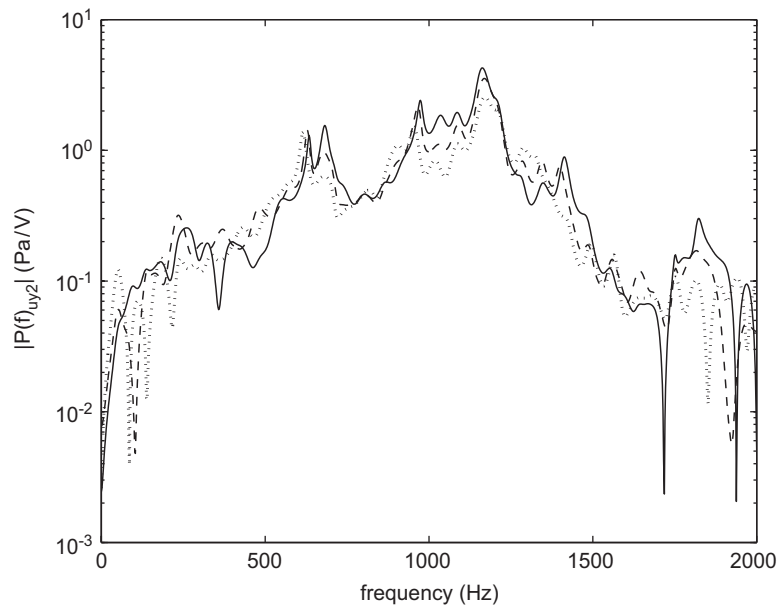

Fig. 12. Identified frequency response functions from actuator input voltage to output pressure at the rear-floor sensor for three free stream Mach numbers. $\cdots M_\infty = 0.275$, — — — $M_\infty = 0.29$, — $M_\infty = 0.30$.

Although Eq. (19) requires knowledge of the true plant transfer function at each Mach number, an approximation to the condition in that equation can be obtained by using identified models of the plant at each Mach number. The magnitude of the identified frequency response functions from actuator input voltage to output pressure at the rear-floor sensor for Mach numbers of 0.275, 0.29, and 0.30 are presented in Fig. 12. Even for this small change in freestream Mach number, the cavity-flow dynamics associated with the Rossiter modes are observed to change. Fig. 13 shows the real part of the product in Eq. (19) for the path from actuator to rear error sensor at these Mach numbers. At $M_\infty = 0.275$ and 0.29 the real part is greater than zero across the entire frequency range, indicating convergence should be stable. However, at $M_\infty = 0.30$, the real part is
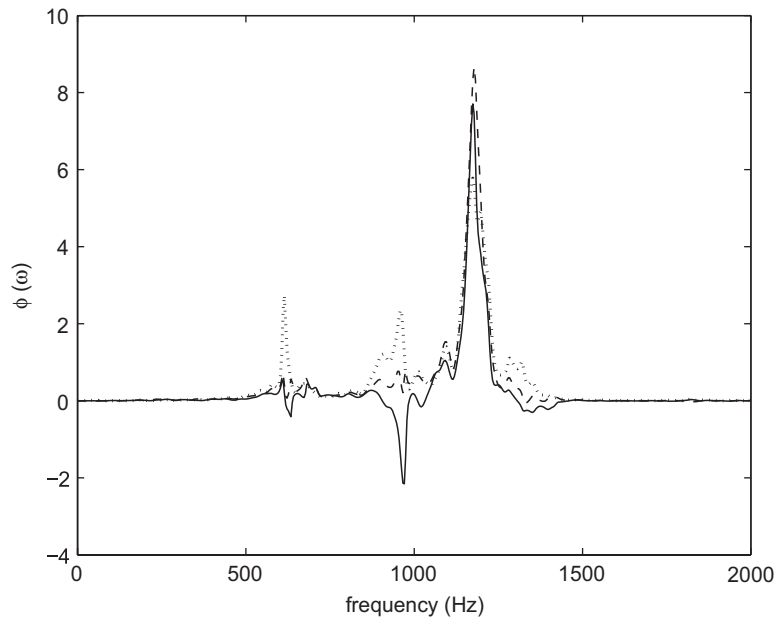
Fig. 13. Real part of the product in Eq. (19) for the path from rear error sensor at three Mach numbers. $\cdots M_\infty = 0.275$, $- - -$ $M_\infty = 0.29$, — $M_\infty = 0.30$.

less than zero just above 600 Hz and just below 1000 Hz. These results thus confirm the experimental observation that a stable controller cannot be maintained beyond $M_\infty = 0.29$ when using a fixed plant model obtained at $M_\infty = 0.275$.

5. Conclusions

In this article, self-tuning adaptive control of cavity flow-tones was considered, with particular emphasis on the development of an online controller design algorithm. The open-loop cavity dynamics were assumed to be linear and time invariant for any given *fixed* Mach number. A stochastic gradient decent algorithm was employed to update the coefficients of a generalized predictive controller at each time step. Past input–output data and an estimate of the open-loop pulse response sequence are needed to implement the algorithm for application at fixed Mach numbers.

Measurements of unsteady cavity pressures and controller coefficients during adaptation were made to examine the transient characteristics of the control algorithm. Joint time–frequency methods showed that multiple Rossiter modes were suppressed and eventually reach a reduced steady-state level. The controller coefficients were also found to converge to a steady-state level, and this implies that controller adaptation can be turned off at some point with no degradation in control performance. Spectra of unsteady pressure time-series acquired after controller convergence exhibited multiple Rossiter mode suppression for each of the Mach numbers tested. However, as in the case of fixed-gain GPC, the adaptive GPC control performance was limited by spillover in sidebands around the suppressed Rossiter modes. No significant reduction in the broadband pressure levels was achieved in any of the control runs. The control performance for fixed-gain GPC and adaptive GPC was compared and found to be nearly identical. This agreement further supports the linear, time-invariant treatment of the open-loop cavity flow dynamics at fixed Mach numbers.

Finally, adaptive GPC was applied to the cavity flow-tone problem as the freestream Mach number was varied. The algorithm was able to maintain suppression of multiple cavity tones over a modest change in Mach number (0.275–0.29). Beyond this range, stable operation of the control algorithm was not possible. This limitation was the result of using a fixed plant model in the algorithm. To achieve the ultimate goal of self-tuning adaptive control for cavity flow tones, the system dynamics must be recursively updated. The

present paper series has demonstrated that real-time feedback control of multiple cavity tones with online control design is feasible. Previous work by Kegerise et al. [5] considered various recursive system identification algorithms that can be used for ARX models of the cavity flow dynamics. With online system identification, a self-tuning adaptive controller that can maintain suppression over a range of Mach numbers should be realized.

References

[1] K.J. Astrom, B. Wittenmark, *Adaptive Control*, Addison-Wesley, Reading, MA, 1995.

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

[3] D.R. Williams, J. Morrow, Adaptive control of multiple acoustic modes in cavities, AIAA Paper 2001-2769, 2001.

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

[5] M.A. Kegerise, L.N. Cattafesta, C. Ha, Adaptive identification and control of flow-induced cavity oscillations, AIAA Paper 2002-3158, 2002.

[6] L. Ljung, *System Identification: Theory for the User*, Prentice-Hall, Upper Saddle River, NJ, 1999.

[7] J. Juang, *Applied System Identification*, Prentice-Hall, Englewood Cliffs, NJ, 1994.

[8] J. Juang, M. Phan, Deadbeat predictive controllers, NASA TM-112862, 1997.

[9] J. Juang, K.W. Eure, Predictive feedback and feedforward control for systems with unknown disturbances, NASA TM-1998-208744, 1998.

[10] K.W. Eure, J. Juang, Broadband noise control using predictive techniques, NASA TM-110320, 1997.

[11] M.Q. Phan, R.K. Lim, R.W. Longman, Unifying input–output and state-space perspectives of predictive control, MAE Technical Report Number 3044, Princeton University, Princeton, NJ, 1998.

[12] R.R. Bitmead, M. Gevers, V. Wertz, *Adaptive Optimal Control: The Thinking Man's GPC*, Prentice-Hall, Upper Saddle River, NJ, 1990.

[13] S. Haykin, *Adaptive Filter Theory*, Prentice-Hall, Upper Saddle River, NJ, 1996.

[14] B. Widrow, S.D. Stearns, *Adaptive Signal Processing*, Prentice-Hall, Upper Saddle River, NJ, 1985.

[15] P.L. Feintuch, An adaptive recursive LMS filter, *Proceedings of the IEEE* 64 (11) (1976) 1622–1624.

[16] L.J. Eriksson, M.C. Allie, R.A. Greiner, The selection and application of an IIR adaptive filter for use in active sound attenuation, *IEEE Transactions on Acoustics, Speech, and Signal Processing* ASSP-35 (1987) 433–437.

[17] D.H. Crawford, R.W. Stewart, Adaptive IIR filtered-v algorithms for active noise control, *Journal of the Acoustical Society of America* 101 (4) (1997) 2097–2103.

[18] P. Darlington, Performance surfaces of minimum effort estimators and controllers, *IEEE Transactions on Signal Processing* 43 (2) (1995) 536–539.

[19] S.J. Elliott, *Signal Processing for Active Control*, Academic Press, San Diego, CA, 2001.

[20] S. Qian, D. Chen, *Joint Time–Frequency Analysis: Methods and Applications*, Prentice-Hall, Upper Saddle River, NJ, 1996.

[21] W. Ren, P.R. Kumar, Adaptive active noise control: structures, algorithms, and convergence analysis, *Proceedings of Inter-Noise* 89 (1989) 435–440.