

Adaptive stabilization and resonance suppression

ANDREW J. TELFORD† and JOHN B. MOORE‡

Control systems can drift into instability or, less catastrophically, exhibit resonance behaviour. One role for adaptive controllers is to learn sufficient information concerning the dominant closed-loop system mode so as to apply effective feedback to dampen these modes. In such situations the adaptive loop augments the fixed controller feedback loop. This paper presents an algorithm for adaptive resonance suppression and provides simulation results to study its behaviour in the presence of high-order unmodelled dynamics. The algorithm appears particularly useful for enhancing existing fixed controller designs.

1. Introduction

With precise knowledge of a linear multivariable plant, controller design techniques are effective at achieving a high performance in terms of disturbance response and control energy trade-offs. However, since plants are invariably uncertain objects, possibly drifting in their characteristics with time, the performance of a nominal model is usually compromised in a design procedure to achieve robustness to plant uncertainty. Even with a fairly robust design, there is a possibility that the control system can catastrophically drift into instability or approach such a condition by exhibiting resonance behaviour. In such cases, one lightly damped mode often dominates, and if this dominant mode can be dampened by adaptive techniques, and the resonance suppressed without otherwise inordinately influencing the control system, robustness/performance enhancement will result.

One area of application of resonance suppression is in aircraft control. For example, the ride quality can deteriorate if a structural resonance is excited by turbulence. Indeed, catastrophic failure can arise in the presence of wing flutter, which can occur in an emergency situation when the flutter speed is exceeded. The work of this paper is motivated by the need for adaptive resonance suppression, rather than the possibility of devising an elegant solution to such a problem. Clearly, if low-order adaptive schemes are applied to high-order uncertain plants such as an aircraft body or wing, there are inevitably unmodelled dynamics. How then can the problem of unmodelled dynamics be overcome? The methods proposed in this paper are presented more as a challenge for following researchers, rather than to give an optimal approach. We do not seek global convergence results for our methods because of unmodelled dynamics, and since there are high-order unmodelled dynamics we do not seek to calculate regions of local convergence. Rather we assess our methods by simulation studies with random model selection.

The adaptive scheme proposed here differs from those treated previously in the literature more in terms of its objectives and orchestration than in terms of its building blocks. It has evolved from case studies such as by Chakravarty and Moore (1986).

Received 17 November 1989.

† BHP Melbourne Research Laboratories, P.O. Box 264, Clayton 3168, Australia.

‡ Department of Systems Engineering, Research School of Physical Sciences, Australian National University, GPO Box 4, Canberra, ACT 2601, Australia.

Most adaptive control designs in the literature tend to replace an off-line designed controller with an adaptive version of the off-line design. Thus, instead of a pole assignment or linear quadratic controller, there is an adaptive pole assignment scheme or an adaptive linear quadratic controller. These utilize linear parameter estimation schemes, based often on input-output models with little or no incorporation of *a priori* model information.

The scheme proposed here is to provide an additional loop on an existing off-line designed control system, associated with a nominal model. The *a priori* information is that the control system is stabilizing for the nominal model, but may drift into a resonant condition or instability. Thus, in the simplest case we may reasonably assume that all modes are stable, except for one dominant mode which is near instability or is just unstable. The desired control objective is merely to dampen the dominant mode somewhat. This suggests estimating its frequency and damping, and applying an adaptive pole assignment scheme to force this mode to a location at the same frequency, but with greater damping. The intention is to achieve this without driving the other modes into instability or exciting other lightly damped modes.

The algorithms are presented in § 2, based on low-order plant idealizations and known least-squares identification and pole assignment techniques. In § 3, simulation studies are presented for both idealized low-order plants and random high-order ones, and stabilization and resonance suppression properties are observed. Section 4 presents an application of the new algorithm, which allows it to enhance adaptively fixed controller designs. Conclusions are drawn in § 5.

2. Adaptive algorithm for resonance suppression

Consider now a scalar stochastic plant or closed-loop system $G(q^{-1})$, modelled in terms of polynomials in the backwards shift operator q^{-1} , with one dominant resonant mode associated with a pair of poles in the vicinity of the unit circle. The adaptive resonance suppression algorithm presented here seeks to determine the frequency band of the resonant mode via an inner identification loop, and exploit this information in an outer adaptive control loop. The algorithm is shown in Fig. 1, and is now developed in some detail.

2.1. Plant model

It is assumed that the plant $G(q^{-1})$ can be accurately modelled by a high-order moving-average autoregressive model (ARMAX):

$$y_k + a_1 y_{k-1} + \dots + a_{n_0} y_{k-n_0} = b_1 u_{k-1} + \dots + b_{m_0} u_{k-m_0} + c_1 w_{k-1} + \dots + c_{p_0} w_{k-p_0} \quad (2.1)$$

where w is a zero mean gaussian noise disturbance process. The parameters of the ARMAX model will be represented by the vector

$$\Theta = [a_1 \quad \dots \quad a_{n_0} \quad b_1 \quad \dots \quad b_{m_0} \quad c_1 \quad \dots \quad c_{p_1}]' \quad (2.2)$$

2.2. Inner loop identification

The recursive least-squares identifier RLS₁ is based on a low-order ARMA model:

$$y_k + a_1 y_{k-1} + \dots + a_{n_1} y_{k-n_1} = b_1 u_{k-1} + \dots + b_{m_1} u_{k-m_1} \quad (2.3)$$

where typically $n_1 = 4$. This identification block has two outputs, $\Theta_{1,k}$ and $\omega_{1,k}$. The

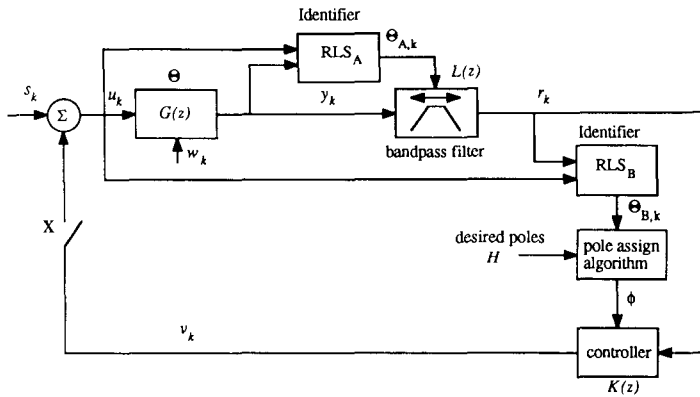


Figure 1. Main control system.

value $\Theta_{1,k}$ is a recursive estimate at time k of the true model Θ . The estimate will not in general asymptotically approach any of the true model parameters, because the model is underparametrized. However, in the presence of a dominant resonant mode (or unstable mode) $\Theta_{1,k}$ can allow a crude estimate of $\omega_1 \in [0, \pi]$, the radian frequency normalized by the sampling frequency, of the least stable pole pair. Here the least stable pole pair is the pole pair belonging to Θ , which is furthest from the origin. The estimate $\omega_{1,k}$ of ω_1 is the frequency of the least stable poles of the ARMA model given by $\Theta_{1,k}$. The idea is that RLS_1 will give a rapidly converging estimate of the frequency ω_1 of the dominant resonance, and this information can be used to adjust the bandpass filter at the plant output. Here the dominant resonant mode is always considered to be due to a complex pole pair; the possibility of a nearly stable or unstable pole on the real axis is excluded.

2.3. Frequency shaping filter

Using the frequency information from the inner loop, our approach is for the bandpass filter at the plant output to be adjusted to accentuate signals in the frequency band associated with the resonant mode, and to attenuate signals at other frequencies. With the series bandpass filter in place, the outer identification/control loop can concentrate on stabilizing the resonant mode, ignoring the now well damped stable modes. The underlying assumption of this approach is that the augmented plant, formed by the combination of the plant and the frequency shaping filter, can be closely approximated by a lower-order plant. Whether this is achieved in practice depends on the nature of the plant and filter. Certainly the frequency shaping attenuates plant modes at a frequencies far from the frequency of the dominant resonant mode.

2.4. Outer-loop control

The outer identification loop consists of a recursive least-squares identifier RLS_O and a pole assignment controller. There is no need to be restricted to RLS identification or pole positioning algorithms: these algorithms were chosen in the simulations for simplicity. Notice that the control loop can be broken at the point X. It may initially be necessary to run the identification algorithms without feedback applied, to allow the estimates $\Theta_{1,k}$, $\Theta_{O,k}$ to get close to their final values.

Since the frequency of the resonant mode is assumed to be initially unknown, a standard pole assignment controller may attempt to move the resonant mode to an assigned pole location far from its initial position. This will result in a large control energy, and possibly destabilize the closed-loop system.

As an alternative to using standard pole assignment, we propose an algorithm in which the location to which the closed-loop poles are assigned is a function of the estimate $\omega_{1,k}$ of the frequency of the resonant mode. An attempt is made to assign the resonant mode radially inwards, so that it has a higher damping than originally, but is still not heavily damped. This idea is depicted in Fig. 2. If $\omega_{1,k}$ is a good estimate of ω_1 , then this pole assignment strategy minimizes the distance that the pole is moved, and hence the control energy to achieve more damping.




Figure 2. Pole assignment with the dominant resonant mode assigned radially inwards.

In the following section, simulation results are discussed and improvements to the algorithm are suggested.

3. Simulation study

3.1. Preliminary results with standard pole assignment

In all of the following simulations, a time-invariant plant $G(q^{-1})$ is used, even though the algorithm is ultimately intended for use when plants have slow time variations. The parameter estimates $\Theta_{1,k}$, $\Theta_{0,k}$ are given arbitrary initial values, and the covariance matrices associated with the RLS estimation are initialized to matrices of the form αI , with α a large real number, and I the identity matrix.

It is also necessary to choose the relative magnitudes of the external excitation s_k , and the noise w_k inherent in the ARMAX model. In most of the following, the magnitude of s_k is ten times the magnitude of $c_1 w_{k-1}$ ($p_0 = 1$).

With $n_1 = 4$, and with various different plants and initial conditions, it is observed that $\omega_{1,k}$ is a good estimate of the frequency of the least stable pole ω_1 . As one would intuitively expect, increasing n_1 results in $\omega_{1,k}$ being a more accurate estimate; the disadvantage is that the computational effort becomes much greater. As a result of these initial tests, it seemed reasonable to set $n_1, m_1 = 1$ for the remaining simulations. This corresponds to RLS₁ identifying a model of the form

$$\frac{b_1 q^{-1} + b_2 q^{-2} + b_3 q^{-3} + b_4 q^{-4}}{1 + a_1 q^{-1} + a_2 q^{-2} + a_3 q^{-3} + a_4 q^{-4}} \quad (3.1)$$

The loop filter $L(q^{-1})$ is designed by bilinear transformation of a low-pass continuous-time prototype of the form

$$L(s) = \frac{\omega^2}{s^2 + 2\zeta\omega s + \omega^2} \quad (3.2)$$

Figure 3 shows the frequency response of such a system. The magnitude of the frequency peak can be increased by reducing the damping coefficient ζ . In the outer loop, a model of the form (3.1) is also identified. In the simulation, the loop is closed at X after a fixed number of iterations and trials run with different values of ζ . For simplicity, and to enable comparison with the adaptive pole positioning in the following section, the closed-loop pole assignment is initially made non-adaptive. Two of the closed-loop poles are assigned to $0.5 \pm 0.5j$ and the rest are assigned to the origin. Since maintaining stability is the most important criterion for a successful controller, many simulations must be run, and a check made of how many runs result in closed-loop stability.




Figure 3. Frequency response of second-order lightly damped transfer function.

The results indicate that there is, in general, no advantage in using the loop filter $L(q^{-1})$. There seemed to be little correlation between the damping of $L(q^{-1})$ (and hence the corresponding frequency selectivity) and the effectiveness of the controller. It is important to note that for a continuous-time plant with a broad spectral response sampled prior to applying adaptive techniques, there is already an in-built pre-filtering which focuses on any frequency band of interest. Our simulations suggest that, in general, there is no merit in a further pre-filtering, although there are certainly situations where pre-filtering does help.

3.2. Closed-loop poles adaptively assigned

We now allow the location to which the closed-loop poles are assigned to be a function of $\omega_{1,k}$. In the simulations, one pair of closed-loop poles is assigned to frequencies of $\pm \omega_{1,k}$ and at a radius of 0.7. The other closed-loop poles are assigned to the origin. The plant models in the simulation have a lightly damped dominant

pole with a randomly chosen frequency, and the other poles with random locations close to the origin. The design parameter that is changed is the damping coefficient ζ associated with the loop filter design. Some simulations are also run with no loop filter; here the only purpose of the inner identifier is to allow the location to which the closed-loop poles are assigned to be adaptive. In all cases, the control algorithm is run in open loop for the first fifty iterations, then closed. The results, based on a small number of simulations, are given in the Table.

	No. runs	No. successful runs	Percentage successful
No filtering	24	18	75
$\zeta = 0.6$	11	2	18
$\zeta = 0.2$	5	11	45

Simulation results: adaptive pole positioning.

3.2.1. *Loop filter.* The results of the Table indicate that the use of the loop filter can actually reduce the chance of obtaining a stabilizing controller. One reason for this is that the loop filter introduces additional poles into the control loop, and these must be taken into account in the pole assignment. A refinement of the algorithm is to assign two closed-loop pole pairs to a radius of approximately 0.7 and the rest to the origin; one to assign the dominant resonant pole radially inwards, and the second to take into account the filter poles. Simulations show that this assignment of a double pole pair results in a more reliable controller.

In the following, we describe some enhancements to the controller design.

3.3. Cautious control

In the simulations above, the control loop is not closed until after the first fifty iterations. The resulting transient sometimes causes the system to become unstable. To lessen the effect of this start-up transient, concepts from Åström's cautious control (Åström 1983) can be implemented. Indirect adaptive control algorithms employing caution use not only the parameter estimate, but also information about the covariance of the estimate, when designing the control law. In our case, the algorithm is modified by replacing the control signal v_k by a scaled control signal $Q_k v_k$, where Q_k is a non-decreasing positive sequence bounded above by one. Such a sequence can be generated as follows. Given outcomes r_k and a regression vector ϕ_k of past measurements, the RLS algorithm identifies the parameters θ of the best model of the form

$$r_k = \phi_k \theta \quad (3.3)$$

The RLS algorithm also calculates a matrix P_k , which, under certain noise assumptions, can be interpreted as an estimate of the covariance matrix associated with the estimate $\hat{\theta}_k$ of the true system model. We therefore propose Q_k given by (3.4), where k_s is some empirically determined positive constant,

$$Q_k = \frac{1}{1 + k_s / (\theta_k' P_k^{-1} \theta_k)} \quad (3.4)$$

Simulations show that there is little variation in Q_k for different simulation runs. As a result of this a Q_k is made independent of θ_k , P_k , and increasing in a linear fashion

where

$$|\det J| = \left| \frac{\det S_{EF}}{\det S_{AB}} \right| \quad (3.11)$$

The minimization above is not possible in practice, because it is impossible to evaluate $f[\psi(\theta)|\mathcal{F}_{k-1}]$ at all values of θ . As a compromise, we instead evaluate this expression only over the set of θ for which $\psi(\theta)$ is of necessity evaluated, that is $\hat{\theta}_k, \hat{\theta}_{k-1}, \dots, \hat{\theta}_{k-M}$ for some M .

For the simulations above based on random plant parameter selection, there is a low probability of introducing near pole-zero cancellations, so that we do not expect any improvement in an average sense as a result of introducing central tendency modifications. However, as shown by Moore *et al.* (1989) non-generic cases can arise where dramatic improvements to performance can be expected.

3.5. Transient performance simulation results

Some simulation results are now presented to show the typical transient behaviour of the controller algorithm. The tenth-order plant is randomly chosen with one unstable pole pair at a radius of 1.1, and four other pole pairs randomly distributed inside a circle of radius 0.7 centred at the origin. The frequency of all the plant poles is uniformly distributed on $[0, \pi)$, and the radius of the stable plant poles is uniformly distributed on $[0, 0.7)$. It is suggested here that the above class of randomly selected plants be used as a benchmark, enabling comparison of our resonance suppression algorithm with those of other authors. The identifiers RLS_1 and RLS_0 identify fourth-order ARMA models. Cautious control and central tendency concepts (with $M = 5$) are used as described in § 3.

One pole pair is assigned to a radius of 0.7 and a frequency given by $\omega_{1,k}$. The control loop does not include a loop filter $L(z)$. Figure 4 shows the estimation of the frequency of the least stable pole pair; the actual frequency based on the true plant parameters is also marked on the graph. The parameter estimates given by RLS_0 are shown in Fig. 5. The plant output y_k increases initially in an unstable manner until the control system learns the plant parameters, after which time y_k settles down again.

4. Preconditioning methods

The methods proposed here have largely been motivated for use in conjunction with other control methods. As an example, consider the indirect adaptive techniques of Tay *et al.* (1989), based on the theory on the class of all stabilizing controllers (Vidyasagar 1985, Doyle 1984). In this work, the real plant is embedded in a control loop, as in Fig. 7. The design of J_K is based on a control system with a nominal plant G_0 and a stabilizing controller K_0 . With stable proper coprime factorizations

$$G_0 = N_0 M_0^{-1} = \tilde{M}_0^{-1} \tilde{N}_0 \quad (4.1 a)$$

$$K_0 = U_0 V_0^{-1} = \tilde{V}_0^{-1} \tilde{U}_0 \quad (4.1 b)$$

then

$$J_K = \begin{bmatrix} K_0 & \tilde{V}_0^{-1} \\ V_0^{-1} & -V_0^{-1} N_0 \end{bmatrix} \quad (4.2)$$

With $Q = 0$, J_K forms a stabilizing controller for the nominal plant; such a controller




Figure 4. Estimate of the frequency ω_1 of the least stable pole pair.




Figure 5. Estimate $\Theta_{0,k}$ generated by outer identifier (RLS₀).

can be designed to achieve specific performance or robustness objectives for the nominal plant. The block Q represents an additional adaptive feedback path over that of the nominal controller. In fact, the nominal plant G_0 is stabilized if and only if Q itself is stable. Furthermore, as Q spans the class of all stable transfer functions, the controller class

$$K(Q) = K_0 + \tilde{V}_0^{-1} Q (I + V_0^{-1} N_0 Q)^{-1} V_0^{-1} \tag{4.3}$$

is the class of all stabilizing controllers for G_0 .

One result of Tay *et al.* (1989) is that G will be stabilized if and only if Q stabilizes




Figure 6. Plant output.

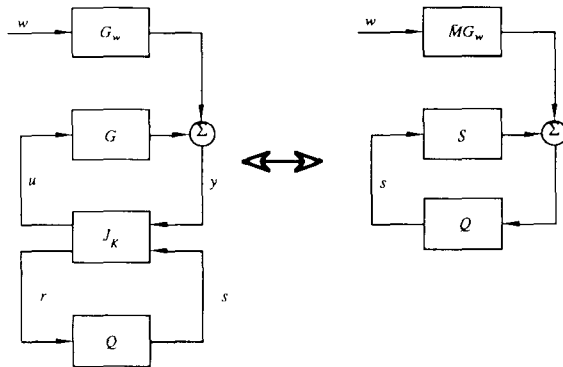


Figure 7. Plant/noise model.

S , where

$$S = \tilde{M}(G - G_0)M_0 \tag{4.4}$$

Here \tilde{M} and M_0 provide a natural frequency weighting for $(G - G_0)$ in the frequency bands of interest.

We now consider the problem of finding a suitable adaptive Q to stabilize the augmented plant S . Since S emphasizes frequencies in the passband of the closed loop system (G, K_0) , it follows that S may often be a transfer function with a dominant resonant mode. This is an ideal opportunity to utilize the resonance suppression techniques of § 2. The proposed scheme is shown in Fig. 8.

5. Conclusions

We present here a preliminary investigation into the problem of designing an adaptive controller when there is *a priori* knowledge that the plant has one dominant

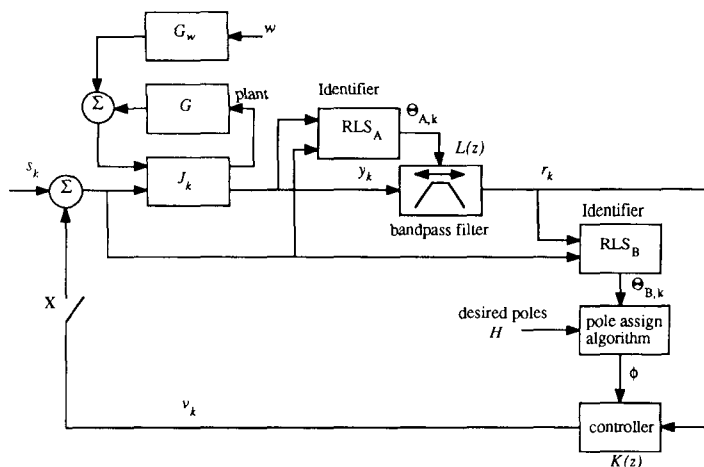


Figure 8. Adaptive scheme to enhance fixed controller design.

lightly damped, or possibly unstable, mode. Such a problem could be tackled without making any use of the knowledge that a dominant mode exists, as is the case when standard adaptive control schemes are applied.

The algorithm in § 2 makes use of an inner underparametrized identification of the plant, which enables fast estimation of the frequency of the resonant mode. This estimate can be used to adjust filters in the control loop, or even to adjust the position to which the closed-loop poles are assigned. The latter possibility seems to be particularly attractive, as it prevents a situation where the controller tries to change the frequency of the dominant resonant mode by arbitrary pole assignment. The controller can instead simply apply feedback to increase the damping of the resonant mode. In practice, techniques such as central tendency adaptive control (Moore *et al.* 1989) and cautious control (Åström 1983) can be used to improve the robustness of the algorithm.

The problem which has been studied is one that we believe is important and that arises in many engineering situations. It is not shown that our algorithm is a universal or optimal resonance suppression algorithm, but instead that *ad hoc* modifications to existing adaptive control algorithms can improve their reliability.

ACKNOWLEDGMENT

The work was partially supported by Boeing (BCAC) and by the DSTO, Australia.

REFERENCES

- ÅSTRÖM, K. J., 1983, Theory and applications of adaptive control—a survey. *Automatica*, **5**, 471–484.
- CHAKRAVARTY, A., and MOORE, J. B., 1986, Flutter suppression via central tendency pole assignment. *Proceedings of the Control Engineering Conference*, Sydney.
- DOYLE, J., 1984, ONR/Honeywell Workshop on Advances in Multivariable Control. Lecture Notes, Minneapolis.
- MOORE, J. B., RYALL, T., and XIA, L., 1989, Central tendency adaptive pole assignment. *I.E.E.E. Transactions on Automatic Control*, **34**, 363–367.

- TAY, T., MOORE, J., and HOROWITZ, R., 1989, Indirect adaptive techniques for fixed controller performance enhancement, *International Journal of Control*, **50**, 1941–1959.
- VIDYASAGAR, M., 1985, *Control Systems Synthesis: a Factorization Approach* (Cambridge, Mass: MIT).