Multiple Model Adaptive Control,
Part 1: Finite Controller Coverings[*]

Brian D. O. Anderson[¶], Thomas S. Brinsmead[¶], Franky De Bruyne[¶],
João Hespanha[‡], Daniel Liberzon[§], A Stephen Morse[§]

September 2000

¶ RSISE, Australian National University, Canberra ACT 0200, Australia
‡ EE-Systems, Univ. of Southern California, Los Angeles, CA 90089-2563
§ Dept. Elec. Eng., Yale University, New Haven, CT 06520-8267, USA

Abstract

We consider the problem of determining an appropriate model set on which to design a set of controllers for a multiple model switching adaptive control scheme. We show that, given mild assumptions on the uncertainty set of linear time-invariant plant models, it is possible to determine a finite set of controllers such that for each plant in the uncertainty set, satisfactory performance will be obtained for some controller in the finite set. We also demonstrate how such a controller set may be found. The analysis exploits the Vinnicombe metric and the fact that the set of approximately band- and time-limited transfer functions is approximately finite-dimensional.

Key Phrases- Band-Limited, epsilon-Entropy (ϵ-Entropy), Finite Covering, Multiple Model Control, Time-limited, Vinnicombe Metric

1 Introduction

One of the more recent approaches to adaptive control is termed multiple model adaptive control (Morse, 1996, 1998a,b; Narendra and Balakrishnan, 1997; Pait and Kassab, 1997, 2001). The concept is as follows. There is an unknown plant P which belongs to a set $\mathcal{P}$ which is usually not a finite set. A set of controllers $\{C_1, C_2, \ldots, \}$ is available, and assumed to have the property that each plant in the set $\mathcal{P}$ will be satisfactorily controlled by at least one of the controllers C_i.

The adaptive control algorithm starts with one of the C_i being connected to P, and based on measurements on the closed loop structure, switching among the controllers takes place until a satisfactory one is obtained. Of course it could be that the initially chosen C_i is satisfactory, and no switching actually occurs. Also, after a satisfactory C_i is obtained, it may be (locally) tuned to further improve performance (Narendra and Balakrishnan, 1997). These issues are considered in this work to be a matter of detail. This paper is concerned with the key preliminary question: Which sets that contain the unknown plant are susceptible to treatment by multiple model adaptive control? In particular, how many controllers are needed and how may they be found?

It is not necessary to stipulate at the outset that the controller family is finite (see, for example Morse, 1996). However, there are good reasons to work with a finite family of controllers. Such reasons include ease of implementation; tractability and efficiency of the minimisation procedure; and tractability of stability analysis. In supervisory control, the estimators typically cannot be implemented in cases where there is a continuum of uncertainty unless a certain separability condition is met (Morse, 1996). Also, a minimisation procedure is used to find the controller to use at each instant of time, and minimisation over an infinite set of controllers is more difficult to perform, if indeed possible. Stability analysis too, is simpler in the finite case for obvious reasons. All these issues provide motivation for the desire that a set of controllers $\{C_1, C_2, \ldots, C_N\}$ that is finite, be found. This problem is the focus of this particular paper. We also point to a companion paper (Hespanha et al., 2000), which deals with the design of a controller switching algorithm once a finite set of controllers has been found. Earlier work, Pait and Kassab (2001), also tackles the finite covering problem in a setting very similar to the one discussed in this paper.

The problem of how to determine an initial set of controllers is obviously quite an important one in terms of the practical implementation of Multiple Model Adaptive Control. In addition, in order to keep the computational burden to within tractable limits, consideration needs to be given to the question of what is the minimum number of controllers (or plant models) which is needed in order to give satisfactory performance. Consideration of such problems is necessary since many works on such control algorithms (Morse, 1996, 1998a,b; Narendra and Balakrishnan, 1997; Hespanha et al., 2000) begin with the assumption that such a set has already been found and concentrate attention on the issue of the switching algorithm. Fortunately these two subproblems of Multiple Model Adaptive Control are reasonably independent and so work on each subproblem can proceed in isolation from that on the other.

In addressing the above question, we shall consider three types of set.

1. The set of unknown plants is characterised by a parameter uncertainty: $\mathcal{P} = \{P(\lambda) : \lambda \in \Lambda\}$ Here, λ is a real parameter and Λ is a subset, usually compact, of $\real^m$ for some m, and $P(\lambda)$ is continuous in λ, in a sense to be defined shortly. A useful image is a "line" of plants (Figure 1) for $m = 1$.

2. The set described in the first item is mildly fattened, by allowing some unstructured uncertainty around $P(\lambda)$ for each $\lambda \in \Lambda$. When $m = 1$, we can think of a thin "tube" of plants (Figure 2).

3. The set described in the previous item is substantially fattened, with unstructured uncertainty, to the extent that even if λ is known, there may be no single controller that achieves satisfactory robust performance for all plants described by the uncertainty set. An extreme case could arise when λ is a singleton, and all the unknownness is non-parametric.

These three possibilities are dealt with in separate sections of the paper. We regard the first item as

Figure 1: "Line" of Plants

Figure 2: "Tube" of Plants

3

the core problem, and so develop the ideas in the most detail for that, including an example showing calculations giving rise to a specific finite set of λ_i on which to base controller design. The remaining items are dealt with as extensions to the first.

In the following section we review the Vinnicombe metric as a measure of the distance between models of linear time-invariant plants, and give sufficient conditions for the existence of a controller that gives satisfactory performance for several plants models. In Section 3, we then investigate the first type of uncertainty set mentioned above: where the plants are parametrised by a parameter from within a compact set. Using the concepts developed in the immediately preceding section, we describe explicitly how to construct a finite set of plants so that each plant in the uncertainty set may be satisfactorily controlled by a controller designed for one of the plants in the finite set. This is followed by a simple numerical example where the uncertainty set is described by an unknown scaling of a nominal plant. We then extend the uncertainty set to permit a limited amount of unstructured uncertainty in addition to the parametric uncertainty, and show that if the unstructured uncertainty is sufficiently small, then a relatively minor extension to the theory presented in Section 3 is required.

Finally we look at the more general case where there is no *a priori* bound on unstructured uncertainty. We show that some commonly used uncertainty set models are not susceptible to finite approximation in a way that gives the desired sufficient conditions for the existence of a finite set of controllers that yields satisfactory performance. Using the concept of ϵ-Entropy developed earlier in Zames (1979), we demonstrate that, with modest and physically realistic restrictions on the uncertainty set, the desired finite approximation can be achieved.

2   The basic multiple model problem: Existence

In this section, we consider the following problem. Let $P(\lambda)$ be a set of plants, indexed by a parameter λ in a compact set Λ, with the elements of the set $P(\lambda)$ depending continuously on λ. The sense in which continuity is to be interpreted will be explored further nearer the end of this section. Suppose that for each λ, a stabilising controller $C(\lambda)$ can be identified which, in conjunction with $P(\lambda)$ achieves a desirable performance. For the purposes of this paper, we assume that $C(\lambda)$ is a family of one degree of freedom controllers, see Figure 3. The ideas presented here can readily be extended to two degree of freedom controllers. We can then ask the following question. Under what circumstances does there exist a finite set of controllers $C(\lambda_1), \cdots, C(\lambda_N)$ such that each $P(\lambda)$ is "satisfactorily" controlled by at least one of the $C(\lambda_i)$? An explanation of "satisfactory performance" appears below. Assuming a positive answer to this existence question, the relevant follow-up question is: How may the required number of controllers, N and the set of $C(\lambda_i)$ be determined?

An object playing a key role in our approach to these questions is the closed-loop generalised sensitivity transfer function matrix (see Glover and McFarlane, 1989; Vinnicombe, 1993), denoted by $T(P(\lambda), C(\lambda))$, which links the inputs r_1, r_2 to the loop signals u and y in Figure 3, and is important in defining the generalised stability margin (Vinnicombe, 1993). We have

$$T(P(\lambda), C(\lambda)) = \left[\begin{array}{c} P(\lambda) \\ I \end{array} \right] \left[I - C(\lambda)P(\lambda) \right]^{-1} \left[-C(\lambda)\ I \right]. \tag{2.1}$$

A measure of the closeness of linear plants or controllers is provided by the Vinnicombe distance (Vinnicombe, 1993). This metric induces the same topology as the earlier gap metric (El-Sakkary, 1985; Georgiou et al., 1987; Zames and El-Sakkary, 1980), but has more desirable properties in that it is less conservative than the gap metric in the following sense. If proximity in the Vinnicombe metric is *unable to guarantee* that a perturbed plant will be stabilised by a controller which results in a certain minimum performance level for a given original plant, then there exists some controller which achieves that minimum performance level on the original plant, and some plant which achieves the Vinnicombe proximity condition, which is also *destabilised* by that controller (Vinnicombe, 1999, see Chapter 4). Such lack of conservatism is not a property which is enjoyed by the gap metric, and hence we argue that the Vinnicombe metric is a more *control relevant* metric than the gap metric.
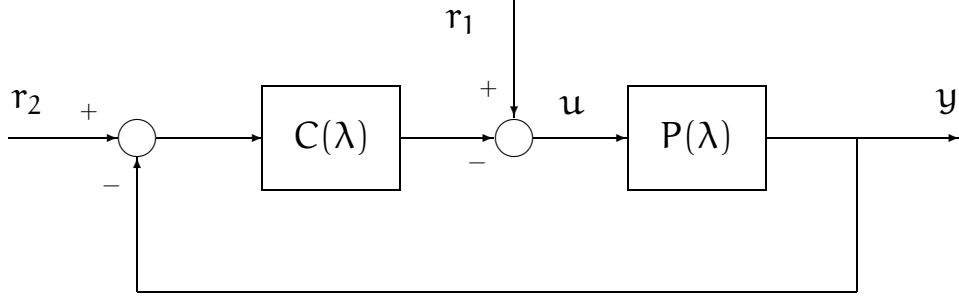
Figure 3: Closed-loop system $(P(\lambda), C(\lambda))$

For two plants with the same input and output dimension, the Vinnicombe distance (Vinnicombe metric, or ν-gap metric) is defined by

$$\delta_\nu(P_1, P_2) = \|(I + P_2 P_2^*)^{-\frac{1}{2}} (P_2 - P_1) (I + P_1^* P_1)^{-\frac{1}{2}}\|_\infty \tag{2.2}$$

provided the following winding number condition is satisfied

$$\det(I + P_1 P_2^*)(j\omega) \neq 0, \quad \forall \omega \quad \text{and} \tag{2.3}$$
$$\text{wno} \det(I + P_1 P_2^*) + \eta(P_1) - \tilde{\eta}(P_2) = 0.$$

If equation (2.3) is not satisfied then $\delta_\nu(P_1, P_2) = 1$. Here $\eta(P_i)$ denote the number of poles of P_i in the *open* right half complex plane $\text{Re}[s] > 0$ and $\tilde{\eta}(P_j)$ is the number of poles of P_j in the *closed* right half plane $\text{Re}[s] \geq 0$, counted according to multiplicity, and wno denotes the winding number evaluated on the standard Nyquist contour indented into the right half plane around any imaginary axis poles of P_1 and P_2. The notation $X^*(s)$ denotes the conjugate system $\bar{X}^T(-\bar{s})$ which, for real rational systems equals $X^T(-s)$ and $\|X\|_\infty$ denotes the $\mathcal{L}_\infty$ norm of the transfer matrix X, given by $\|X\|_\infty = \sup_{\omega \in} \sigma_{\max}[X(j\omega)]$, where $\sigma_{\max}(X)$ denotes the maximum singular value of a matrix.

Now let $\bar{P}$ be an arbitrary plant with

$$\delta_\nu(\bar{P}, P(\lambda)) < \frac{1}{\|T(P(\lambda), C(\lambda))\|_\infty}. \tag{2.4}$$

It is known that if $C(\lambda)$ stabilises $P(\lambda)$, then $C(\lambda)$ also stabilises $\bar{P}$ (see Vinnicombe, 1993). Moreover, a simple calculation set out in the appendix (see Appendix A.1) yields

$$\|T(\bar{P}, C(\lambda)) - T(P(\lambda), C(\lambda))\|_\infty \leq \frac{\|T(P(\lambda), C(\lambda))\|_\infty^2 \, \delta_\nu(\bar{P}, P(\lambda))}{1 - \|T(P(\lambda), C(\lambda))\|_\infty \, \delta_\nu(\bar{P}, P(\lambda))} \tag{2.5}$$

With equation (2.4), we can identify a range of plants $P(\mu)$ for μ near λ which are stabilised by $C(\lambda)$, and with (2.5) we can identify a range of plants which give similar performance as the one achieved on $P(\lambda)$ when connected in closed-loop with $C(\lambda)$. Here, we are implicitly identifying performance with the achievement of a particular target shape for $T(P, C(\lambda))$. We might agree for example that an acceptable performance with $\bar{P}$ occurs provided

$$\|T(\bar{P}, C(\lambda)) - T(P(\lambda), C(\lambda))\|_\infty \leq r \|T(P(\lambda), C(\lambda))\|_\infty, \tag{2.6}$$

for some r with $0 < r < 1$. Noting that the following two conditions are equivalent:

$$\|T(P(\lambda), C(\lambda))\|_\infty \, \delta_\nu(\bar{P}, P(\lambda)) < r(1 + r)^{-1} = \hat{r}, \tag{2.7}$$
$$\frac{\|T(P(\lambda), C(\lambda))\|_\infty \, \delta_\nu(\bar{P}, P(\lambda))}{1 - \|T(P(\lambda), C(\lambda))\|_\infty \, \delta_\nu(\bar{P}, P(\lambda))} < r,$$

it is easy to see, using the fact that equation (2.5) holds, that equations (2.6) and (2.7) are also equivalent conditions. (Note that if $r \ll 1$ then $\hat{r} \approx r$). Either condition (2.7) or (2.6) also automatically assures (2.4), and therefore any $\bar{P}$ satisfying equation (2.7) is certainly stabilised by $C(\lambda)$. Thus when we tighten up condition (2.4)— which only guarantees that when $C(\lambda)$ stabilises $P(\lambda)$ it will then stabilise $\bar{P}$— then we actually obtain the property that the $[P(\lambda), C(\lambda)]$ loop and the $[\bar{P}, C(\lambda)]$ loop will exhibit similar performance. This allows us to state the following desirable result.

Theorem 2.1 *For any given controller performance threshold r in equation (2.6), there exists a finite set of controllers $C(\lambda_i)$, such that each plant in the set, parametrised by λ, given by $\mathcal{P} = \{P(\lambda) : \lambda \in \Lambda\}$ is controlled satisfactorily by some $C(\lambda_i)$ in the sense that equation (2.6) is satisfied. In the definition of $\mathcal{P}$, the set Λ is understood to be a compact set, and the plants $P(\lambda)$ depend continuously on the parameter λ,*

Proof : Given any r, by the previous arguments, we see that there exists an open ball $\mathcal{B}(\lambda)$ around each λ such that for any $\mu \in \mathcal{B}(\lambda)$ the equation (2.7) with $P(\mu)$ substituted for $\bar{P}$ is satisfied and so an infinite cover of Λ exists. Because the set Λ is compact, we may appeal to the Heine-Borel property (see Loomis and Sternberg, 1968, pg 214.) to conclude that the set Λ may be covered by a finite set of balls, $\mathcal{B}(\lambda_1), \cdots, \mathcal{B}(\lambda_N)$ say, with the property that if $\mu \in \mathcal{B}(\lambda_i)$, then $P(\mu)$ is satisfactorily controlled by $C(\lambda_i)$. The result follows immediately. $\qquad\square$

The above arguments provide sufficient conditions only, that is, there may well be plants $\bar{P}$ not so close in the Vinnicombe metric to $P(\lambda)$ as indicated by (2.7) for which $C(\lambda)$ provides satisfactory performance. There is, in fact a comparatively easy way to obtain a less conservative result. Let

$$\widehat{\delta_\nu}(P_1, P_2, j\omega) = \sigma_{\max}\left\{[I+P_2(j\omega)P_2^*(j\omega)]^{-\frac{1}{2}}[P_2(j\omega)-P_1(j\omega)][I+P_1(j\omega)P_1^*(j\omega)]^{-\frac{1}{2}}\right\} \qquad (2.8)$$

Then provided that the winding number condition (2.3) is satisfied $C(\lambda)$ will also stabilise $\bar{P}$ if instead of (2.4), one has

$$\widehat{\delta_\nu}(P, P(\lambda), j\omega) \quad < \quad \frac{1}{\sigma_{\max}\{T(P(\lambda), C(\lambda), j\omega)\}}, \text{ for all } \omega$$

and acceptable performance in the sense that

$$\sigma_{\max}\{T(P, C(\lambda), j\omega) - T(P(\lambda), C(\lambda), j\omega)\} \quad < \quad r \cdot \sigma_{\max}\{T(P, C(\lambda), j\omega)\}, \text{ for all } \omega$$

is assured if

$$\sigma_{\max}\{T(P(\lambda), C(\lambda), j\omega)\}\widehat{\delta_\nu}(P, P(\lambda), j\omega) \quad < \quad \hat{r}, \text{ for all } \omega.$$

In short, one may replace δ_ν and $\|T\|_\infty$ by frequency dependent quantities and invoke (2.4) and (2.7) on a point-wise in frequency basis. Further reduction of conservatism can be obtained by introducing frequency weighting terms in the above calculations (see Vinnicombe, 1999, 1992). Particularly if the parametric uncertainty is simply a scaling factor, then the appropriate frequency weighting is approximately proportional to the frequency-dependent Vinnicombe metric $\widehat{\delta_\nu}$ between the central plant and those plants corresponding to the neighbourhood of the uncertain parameter. Another, quite different way of obtaining a less conservative result is to restrict the complexity (Vinnicombe, 1996) of the controllers $C(\lambda_i)$.

Let us now begin to explore what we mean by continuity. It is clear from the above analysis that we actually want continuity in the Vinnicombe metric. Equivalently (Georgiou and Smith, 1990) one would need to think about continuity in the graph topology.

There are a number of sufficient conditions for this which reflect common situations. First, if $P(\lambda)$ is stable for all λ, and $P(\lambda, j\omega)$ is continuous in λ in the $\mathcal{H}_\infty$ norm, then we have continuity in the Vinnicombe metric because $\delta_\nu(P(\lambda), P(\mu)) \leq \|P(\lambda) - P(\mu)\|_\infty$. This can be seen from (2.2).

Next, suppose that $P(\lambda) = N(\lambda)\, D^{-1}(\lambda)$ is a fractional description of $P(\lambda)$ using a normalised coprime realization (Vardulakis, 1991). Suppose also that $[N(\lambda)^\top \ D(\lambda)^\top]^\top$ is continuous in λ in the $\mathcal{H}_\infty$ norm, in the sense that given arbitrary $\epsilon > 0$, there exists $\delta(\epsilon)$ such that

$$\left\| \begin{bmatrix} N(\lambda) \\ D(\lambda) \end{bmatrix} - \begin{bmatrix} N(\mu) \\ D(\mu) \end{bmatrix} \right\|_\infty \quad < \quad \epsilon$$
$$\text{for all } \mu \text{ with } |\mu - \lambda| \quad < \quad \delta. \tag{2.9}$$

Note that if λ lies in a compact set, then continuity is uniform. If a particular factorisation is given which is not normalised, but the transformation to a normalised fraction is achieved in a manner depending continuously on λ, again, the requisite continuity in equation (2.9) is guaranteed. To understand why this gives continuity in the Vinnicombe metric, note that one characterisation of $\delta_v(P(\lambda), P(\mu))$ is

$$\delta_v(P(\lambda), P(\mu)) = \inf_{\substack{Q,\ Q^{-1} \in \mathcal{RL}_\infty \\ \text{wno}\det(Q) = 0}} \left\| \begin{bmatrix} N(\lambda) \\ D(\lambda) \end{bmatrix} - \begin{bmatrix} N(\mu) \\ D(\mu) \end{bmatrix} Q \right\|_\infty.$$

If it holds that

$$\left\| \begin{matrix} N(\lambda) - N(\mu) \\ D(\lambda) - D(\mu) \end{matrix} \right\|_\infty < \epsilon, \tag{2.10}$$

then (by choosing $Q = I$, rather than the infimum)

$$\delta_v(P(\lambda), P(\mu)) \leq \left\| \begin{matrix} N(\lambda) - N(\mu) \\ D(\lambda) - D(\mu) \end{matrix} \right\|_\infty < \epsilon.$$

In summary, continuity of the normalised coprime fraction description in the $\mathcal{H}_\infty$ norm is sufficient for continuity in the gap metric.

As a third example, related to the previous one, suppose that $P(\lambda) = n(\lambda, s)\, d^{-1}(\lambda, s)$ where $n(\lambda, s)$ and $d(\lambda, s)$ are polynomials in s with coefficients depending continuously on λ with $d(\lambda, s)$ monic. Suppose that there does not exist $\lambda \in \Lambda$ with the property that $n(\lambda, s)$ and $d(\lambda, s)$ have a common zero in $\text{Re}[s] \geq 0$. Then the continuity condition is able to be fulfilled.

A special case is $P(\lambda) = n(\lambda, s)\, d^{-1}(\lambda, s)$ where $n(\lambda, s)$ and $d(\lambda, s)$ are polynomials in s with coefficients depending continuously on λ, and $P(\lambda)$ is stable for all $\lambda \in \Lambda$.

3 Basic Multiple Model Problem: Constructive Aspects

In the previous section, we observed that if a set of plants $P(\lambda)$ indexed by a parameter λ in a compact set Λ obeys a continuity condition, then one can determine a finite set of controllers $C(\lambda_i), i = 1 \ldots N$ such that each plant is stabilised, and even satisfactorily controlled, by one of these controllers. The argument depends on a covering of the set $\mathcal{P}$ by a finite set of balls.

The basic multiple model construction problem is then to determine what that covering should be, that is, to determine both what is an acceptable value of N and what values the λ_i should assume?

The argument already given in the last section indicates what size a ball around λ_i can be, once a controller has been determined. The determination of a single controller for $P(\lambda_i)$ of course depends on the purpose of the controller design. It is the controller which will determine $T(P(\lambda_i, C(\lambda_i))$ and its norm. There is, however, one important point to make, which is that, for a fixed $P(\lambda_i)$, there is an under-bound for $\|T(P(\lambda_i, C(\lambda_i))\|_\infty$, when $C(\lambda_i)$ ranges over all stabilising controllers. This is obtained as follows (Glover and McFarlane (1989), see also Zhou, Doyle, and Glover (1996)). Let $P(\lambda) = N(\lambda)D(\lambda)^{-1}$ where $[N, D]$ constitute a normalised right coprime factorisation of the transfer

matrix $P(\lambda)$. Then

$$\inf_{\substack{C \text{ is stabilising}}} \|T(P_\lambda, C)\|_\infty \;=\; \left[1 - \left\| \begin{array}{c} N(\lambda) \\ D(\lambda) \end{array} \right\|_H^2 \right]^{-\frac{1}{2}} \tag{3.1}$$

where $\|\cdot\|_H$ denotes the Hankel norm and evaluates to be the largest magnitude Hankel singular value. The infimum is attained by a finite dimensional controller.

In the case that Λ is a finite interval in , say $\Lambda = [\lambda_{min}, \lambda_{max}]$ it is comparatively easy to outline how N and the λ_i should be chosen.

- For each λ determine an acceptable $C(\lambda)$ and the associated $\|T(P(\lambda), C(\lambda))\|_\infty$. Note that the value of equation (3.1) could be used.

  Note that in practice, because $P(\lambda)$ is continuous in λ, it may be sufficient to investigate only a finite set of $P(\lambda_k)$. It is possible to choose the density of the λ_k according to the δ of equation (2.9) to guarantee that $P(\lambda_k)$ approximates $P(\lambda)$ for $|\lambda_k - \lambda| < \delta$ within an arbitrarily chosen ϵ degree of accuracy.

- Determine the value of $\delta_\nu(P(\lambda), P(\mu))$ for values of λ, μ near λ_{min}. Choose λ_1, so that for all $\lambda \in [\lambda_{min}, \lambda_1]$, the controller $C(\lambda_1)$ satisfactorily controls $P(\lambda)$:

$$\lambda_1 \;=\; \max\{\bar{\lambda} : \delta_\nu(P(\lambda), P(\bar{\lambda}))\|T(P(\bar{\lambda}), C(\bar{\lambda}))\|_\infty \leq \hat{r}, \ \forall \lambda \in [\lambda_{min}, \bar{\lambda}]\}$$

  Remark: If stability rather than "satisfactory control" is needed, the $\hat{r}$ bound in the inequality above may be replaced by $1 - \epsilon$ for some arbitrarily small ϵ

- Notice that $C(\lambda_1)$ will also satisfactorily control each plant $P(\lambda)$ with $\lambda > \lambda_1$ for which

$$\delta_\nu(P(\lambda), P(\lambda_1))\|T(P(\lambda_1), C(\lambda_1))\|_\infty \;\leq\; \hat{r}.$$

  Let

$$\hat{\lambda}_1 \;=\; \max\left\{\bar{\lambda} : \delta_\nu(P(\lambda), P(\lambda_1))\|T(P(\lambda_1), C(\lambda_1))\|_\infty \leq \hat{r} \ \forall \lambda \in [\lambda_1, \bar{\lambda}]\right\}$$

  Then we choose λ_2 to ensure that all plants with $\lambda \in [\hat{\lambda}_1, \lambda_2]$ are satisfactorily controlled:

$$\lambda_2 \;=\; \max\{\bar{\lambda} : \delta_\nu(P(\lambda), P(\bar{\lambda}))\|T(P(\bar{\lambda}), C(\bar{\lambda}))\|_\infty \leq \hat{r}, \ \forall \lambda \in [\hat{\lambda}_1, \bar{\lambda}]\},$$

- More generally, having chosen λ_i, we have

$$\hat{\lambda}_i \;=\; \max\{\bar{\lambda} : \delta_\nu(P(\lambda), P(\lambda_i))\|T(P(\lambda_i), C(\lambda_i))\|_\infty \leq \hat{r}, \ \forall \lambda \in [\lambda_i, \bar{\lambda}]\}.$$

$$\lambda_{i+1} \;=\; \max\{\bar{\lambda} : \delta_\nu(P(\lambda), P(\bar{\lambda}))\|T(P(\bar{\lambda}), C(\bar{\lambda}))\|_\infty \leq \hat{r}, \ \forall \lambda \in [\hat{\lambda}_i, \bar{\lambda}]\}.$$

- We stop when $\lambda_{max} \in [\hat{\lambda}_i, \hat{\lambda}_{i+1}]$ and set $N = i + 1$. The above procedure assures that $C(\lambda_i)$ satisfactorily controls plants $P(\lambda)$ for $\lambda \in [\hat{\lambda}_{i-1}, \hat{\lambda}_i]$

Obvious variations to the above can take place when Λ is not single-dimensional, but still finite dimensional.

4 Numerical example

We examine a set of plants which has already been investigated in earlier works (see, for example Doyle et al., 1992; Morse, 1996). The set of plants is

$$\mathcal{P} = \left\{ P(\lambda) = \lambda \frac{s-1}{(s+1)(s-2)}, \lambda \in [1, 40] \right\} \tag{4.1}$$

With the aim of developing a multiple model control algorithm, a set of values λ_i was sought, so that the controllers $C(\lambda_i)$ could be found such that each $P(\lambda)$ would be controlled "satisfactorily" by one of the $C(\lambda_i)$.

Because the parametric uncertainty in this system is a very simple scalar multiplicative gain, it is possible, and indeed straightforward, to work with an infinite continuum of controllers. In this example, it is easy to design a parameterized family of controllers, with one controller for each possible value of the parameter, even though the parameter takes infinitely many values. In fact we just need to design one controller for $\lambda = 1$ (by whatever method we like best) and then, controllers for different values of λ are automatically obtained by normalising the gain of the original controller by λ.

In this sense, then the use of a finite number of fixed controllers for this particular system is actually hard to justify. However, in the general case, when the unknown parameter enters the model in a more complicated fashion, it may be very difficult to design a parameterized family of controllers as the parameter takes values in an infinite set. Thus, the choice of a finite set of controllers for this simple model of scalar multiplicative uncertainty is merely to illustrate that such a scheme is possible, despite the fact that, for this example, it is not absolutely necessary.

A trial and error approach, which is necessarily tedious and non-systematic, led to the conclusion that 21 values of λ_i would suffice, with $\lambda_i = (1.2)^{i-1}$ for $i = 1, 2, \ldots, 21$. A family of controllers C_i was also proposed (Morse, 1996), with

$$C_i = \lambda_i^{-1} \frac{448s^2 + 450s - 18}{31s(s-9)} \tag{4.2}$$

Using Vinnicombe metric techniques, we verify that this was a good solution. Table 1 shows the results of these calculations. The second column shows the λ_i corresponding to each plant in the model set, and the next column shows a range of λ which corresponds to plants deemed to be "close to" λ_i (these were calculated as $\overrightarrow{\lambda} = (1.2)^{i-0.5}$ and $\overleftarrow{\lambda} = (1.2)^{i+0.5}$). The Vinnicombe distance between $P(\lambda_i)$ and the plants corresponding to the lower and upper bounds of the corresponding λ range appears next as $\delta_\nu(P(\lambda_i), P(\overrightarrow{\lambda}))$ and $\delta_\nu(P(\lambda_i), P(\overleftarrow{\lambda}))$. It can be seen these distances are quite consistent across the data set. The sixth column gives a value of the optimal stability margin (calculated from equation (3.1)) for each plant, and the next columns shows the quantity in equation (2.7). We note that for the last few examples, this quantity exceeds unity, and so equation (2.4) does not guarantee that there exists a controller which yields stability for all $\lambda \in [\overrightarrow{\lambda}, \overleftarrow{\lambda}]$. The frequency dependent quantities appear in the next two columns, corresponding to the optimal controller of equation (3.1). Note, however, that as may be anticipated (Green and Limebeer, 1995), the closed loop generalised sensitivity function corresponding to the optimal $\mathcal{H}_\infty$ problem (3.1), is flat. And so the optimal frequency dependent quantities differ little from the more conservative, frequency independent calculations.

We now investigate the plant-controller controller combination suggested by equation (4.2). The results appear in Table 2. The stability margin corresponding to the plant and controller combination is necessarily inferior to the optimal stability margin b_{opt} (compare the fourth column of the Table 2 to the sixth column of Table 1) and so the frequency independent quantities in the fifth and sixth columns of Table 2, which display $\delta_\nu(P(\lambda_i), P(\overrightarrow{\lambda}))b_{[p,C]}^{-1}$ and $\delta_\nu(P(\lambda_i), P(\overleftarrow{\lambda}))b_{[p,C]}^{-1}$ are necessarily worse than the corresponding data in Table 1. However, the frequency dependent quantities with the controller (4.2) are better, and in fact are less than unity for the entire range of λ. Hence by equation

i	λ_i	$[\overrightarrow{\lambda}, \overleftarrow{\lambda}]$	$\delta_\nu(P_\lambda, P_{\overrightarrow{\lambda}})$ $\times 10^{-2}$	$\delta_\nu(P_\lambda, P_{\overleftarrow{\lambda}})$ $\times 10^{-2}$	b_{opt} $\times 10^{-2}$	$\dfrac{\overrightarrow{\delta_\nu}}{b_{opt}}$	$\dfrac{\overleftarrow{\delta_\nu}}{b_{opt}}$	$\sup\limits_\omega \dfrac{\widehat{\overrightarrow{\delta_\nu}}}{\widehat{b}_{opt}}$	$\sup\limits_\omega \dfrac{\widehat{\overleftarrow{\delta_\nu}}}{\widehat{b}_{opt}}$
1	1.00	[0.91, 1.10]	3.54	3.74	7.45	0.48	0.50	0.48	0.50
2	1.20	[1.10, 1.31]	3.94	4.10	8.55	0.46	0.48	0.46	0.48
3	1.44	[1.31, 1.58]	4.26	4.38	9.66	0.44	0.45	0.44	0.45
4	1.73	[1.58, 1.89]	4.47	4.53	10.70	0.42	0.42	0.42	0.42
5	2.07	[1.89, 2.27]	4.56	4.56	11.60	0.39	0.39	0.39	0.39
6	2.49	[2.27, 2.73]	4.56	4.56	12.24	0.37	0.37	0.37	0.37
7	2.99	[2.73, 3.27]	4.55	4.55	12.56	0.36	0.36	0.36	0.36
8	3.58	[3.27, 3.93]	4.55	4.55	12.52	0.36	0.36	0.36	0.36
9	4.30	[3.93, 4.71]	4.55	4.56	12.14	0.38	0.38	0.38	0.38
10	5.16	[4.71, 5.65]	4.55	4.56	11.46	0.40	0.40	0.40	0.40
11	6.19	[5.65, 6.78]	4.56	4.56	10.57	0.43	0.43	0.43	0.43
12	7.43	[6.78, 8.14]	4.56	4.56	9.56	0.48	0.48	0.48	0.48
13	8.92	[8.14, 9.77]	4.56	4.55	8.51	0.54	0.54	0.54	0.54
14	10.70	[9.77, 11.7]	4.56	4.56	7.47	0.61	0.61	0.61	0.61
15	12.84	[11.7, 14.1]	4.55	4.55	6.49	0.70	0.70	0.70	0.70
16	15.40	[14.1, 16.9]	4.56	4.56	5.60	0.81	0.81	0.81	0.81
17	18.49	[16.9, 20.3]	4.55	4.55	4.79	0.95	0.95	0.95	0.95
18	22.19	[20.3, 24.3]	4.55	4.55	4.08	1.12	1.12	1.11	1.11
19	26.62	[24.3, 29.2]	4.56	4.56	3.46	1.32	1.32	1.31	1.31
20	31.94	[29.2, 35.0]	4.56	4.56	2.93	1.56	1.56	1.55	1.55
21	38.34	[35.0, 42.0]	4.56	4.56	2.47	1.85	1.85	1.84	1.84

Table 1: Data for $\mathcal{P}$ using the Optimal Stability Margin Controller

(2.4), stability is guaranteed for all plants with $\lambda \in [\overrightarrow{\lambda}, \overleftarrow{\lambda}]$. The fact that the frequency dependent ratios vary little with i also suggests that the choice of λ_i is a reasonable one.

We note that the above is still a conservative result, since it does not take into account restrictions on controller complexity (Vinnicombe, 1999, see Chapter 6).

5 Multiple Model Problem in the Presence of Parametric and Small Unstructured Uncertainty

In Section 2, we introduced a set $\mathcal{P}$ of plants $P(\lambda)$, which depend continuously on $\lambda \in \Lambda$ for a compact set Λ. It was shown that various forms of continuity are sufficient to be able to derive the results on robust stability and performance. What is necessary is continuity in the Vinnicombe metric. Given any $\epsilon > 0$, there must be a δ such that $\delta_\nu(P(\lambda), P(\mu)) < \epsilon$ whenever $|\lambda - \mu| < \delta$. In this section, we postulate that the set of plants of interest is wider again, allowing unstructured uncertainty around $P(\lambda)$. There are various ways of doing this. For example, we could allow $\bar{P}(j\omega)$ to vary such that

$$\left| \bar{P}(j\omega) - P(\lambda, j\omega) \right| < \epsilon(\lambda, j\omega) \tag{5.1}$$

where some appropriate point-wise norm is used in the multi-variable case. Alternatively, we could use

$$\| \bar{P} - P(\lambda) \|_\infty \leq \epsilon(\lambda). \tag{5.2}$$

and require $\bar{P} - P(\lambda)$ to be stable. This would permit unstable $P(\lambda)$, but the unstable (anti-causal) parts of $P(\lambda)$ and $\bar{P}$ would have to be identical. One way that is particularly tailored for proving

i	λ_i	$[\overrightarrow{\lambda}\,,\overleftarrow{\lambda}\,]$	$b_{[P,C]}$ $\times 10^{-2}$	$\dfrac{\overrightarrow{\delta_\nu}}{b_{[P,C]}}$	$\dfrac{\overleftarrow{\delta_\nu}}{b_{[P,C]}}$	$\sup\limits_\omega \dfrac{\widehat{\overrightarrow{\delta_\nu}}(j\omega)}{\hat{b}_{[P,C]}(j\omega)}$	$\sup\limits_\omega \dfrac{\widehat{\overleftarrow{\delta_\nu}}(j\omega)}{\hat{b}_{[P,C]}(j\omega)}$
1	1.00	[0.91, 1.10]	6.14	0.58	0.61	0.51	0.51
2	1.20	[1.10, 1.31]	7.14	0.55	0.57	0.52	0.52
3	1.44	[1.31, 1.58]	8.02	0.53	0.55	0.53	0.53
4	1.73	[1.58, 1.89]	8.28	0.54	0.55	0.54	0.54
5	2.07	[1.89, 2.27]	8.27	0.55	0.55	0.55	0.55
6	2.49	[2.27, 2.73]	7.99	0.57	0.57	0.56	0.56
7	2.99	[2.73, 3.27]	7.51	0.61	0.61	0.57	0.57
8	3.58	[3.27, 3.93]	6.85	0.66	0.66	0.58	0.58
9	4.30	[3.93, 4.71]	6.12	0.74	0.75	0.59	0.59
10	5.16	[4.71, 5.65]	5.37	0.85	0.85	0.60	0.60
11	6.19	[5.65, 6.78]	4.64	0.98	0.98	0.60	0.60
12	7.43	[6.78, 8.14]	3.97	1.15	1.15	0.61	0.61
13	8.92	[8.14, 9.77]	3.36	1.36	1.35	0.61	0.61
14	10.70	[9.77, 11.7]	2.84	1.60	1.60	0.61	0.61
15	12.84	[11.7, 14.1]	2.39	1.90	1.90	0.62	0.62
16	15.40	[14.1, 16.9]	2.01	2.27	2.27	0.62	0.62
17	18.49	[16.9, 20.3]	1.68	2.71	2.71	0.62	0.62
18	22.19	[20.3, 24.3]	1.40	3.25	3.25	0.62	0.62
19	26.62	[24.3, 29.2]	1.17	3.89	3.89	0.62	0.62
20	31.94	[29.2, 35.0]	0.98	4.66	4.66	0.62	0.62
21	38.34	[35.0, 42.0]	0.82	5.58	5.59	0.62	0.62

Table 2: Data for $\mathcal{P}$ using the Suggested Controller (equation 4.2)

stability results is to require

$$\delta_\nu(\bar{P}, P(\lambda)) < \epsilon(\lambda). \tag{5.3}$$

We assume that $\epsilon(\lambda)$ is continuous with λ. Note that the metrics in equations (5.2) and (5.3) can be related. For example, if $\bar{P}$ and $P(\lambda)$ are stable, then

$$\delta_\nu(\bar{P}, P(\lambda)) \le \|\bar{P} - P(\lambda)\|_\infty \le \left\| \begin{array}{c} \bar{P} \\ I \end{array} \right\|_\infty \left\| \begin{array}{c} P(\lambda) \\ I \end{array} \right\|_\infty \delta_\nu(\bar{P}, P(\lambda)).$$

Also, as noted in Section 2, virtually the whole of the Vinnicombe theory can be carried through on a frequency-by-frequency basis. This involves using a quantity $\widehat{\delta_\nu}(\bar{P}, P(\lambda), j\omega)$ with the property $\delta_\nu(\bar{P}, P(\lambda)) = \sup_\omega \widehat{\delta_\nu}(\bar{P}, P(\lambda), j\omega)$ and allows us to connect equations (5.2) and (5.3) to equation (5.1).

Our task is to choose a finite set of controllers $C(\lambda_1), \cdots, C(\lambda_N)$ such that each plant $\bar{P}$ satisfying (5.3) for some $\lambda \in \Lambda$ is satisfactorily controlled by one of the $C(\lambda_i)$. By satisfactory control or performance, we follow equation (2.7) in Section 2, that is, we mean that for some $\hat{r} \in (0, \frac{1}{2})$

$$\delta_\nu(\bar{P}, P(\lambda_i)) \, \|T(P(\lambda_i), C(\lambda_i))\|_\infty < \hat{r}.$$

Now as before, suppose that $C(\lambda)$ satisfactorily stabilises $P(\lambda)$. In order that $C(\lambda)$ stabilises all $\bar{P}$ for which $\delta_\nu(\bar{P}, P(\lambda)) < \epsilon(\lambda)$, by the main robust stability theorem of Vinnicombe (1999), it is necessary and sufficient that

$$\epsilon(\lambda) \, \|T(P(\lambda), C(\lambda))\|_\infty < 1, \; \forall \lambda. \tag{5.4}$$

In order that $C(\lambda)$ satisfactorily controls all $\bar{P}$ for which $\delta_\nu(\bar{P}, P(\lambda)) < \epsilon(\lambda)$ it is sufficient that

$$\epsilon(\lambda) \, \|T(P(\lambda), C(\lambda))\|_\infty < \hat{r}, \; \forall \lambda. \tag{5.5}$$

If equation (5.5) is not fulfilled, it means that the combination of the performance objective that gives rise to $T(P(\lambda), C(\lambda))$, the performance robustness objective that gives rise to r, and the size of the unmodeled dynamics under consideration that gives rise to $\epsilon(\lambda)$, are, as sufficiency conditions, mutually incompatible, and something must be relaxed.

In this section, however, we assume that equation (5.5) holds. We investigate the case of large unstructured uncertainty, where this does not hold, in the following section. Since equation (5.5) holds, it follows that $C(\lambda)$ will also satisfactorily control a $\bar{P}$ satisfying

$$\delta_\nu(\bar{P}, P(\mu)) < \epsilon(\mu) \tag{5.6}$$

for some μ (almost certainly close to λ) if we have that equation (5.6) implies

$$\delta_\nu(\bar{P}, P(\lambda)) \, \|T(P(\lambda), C(\lambda))\|_\infty < \hat{r}. \tag{5.7}$$

We now find some sufficient conditions on μ so that equation (5.6) implies (5.7). Suppose that μ is restricted to be such that

$$\delta_\nu(P(\mu), P(\lambda)) \, \|T(P(\lambda), C(\lambda))\|_\infty < \hat{r} - \epsilon(\mu) \, \|T(P(\lambda), C(\lambda))\|_\infty. \tag{5.8}$$

Note that for $\mu = \lambda$, the right hand side of (5.8) is positive by (5.5) above. The continuity of $P(\mu)$ with μ means that $\delta_\nu(P(\mu), P(\lambda))$ is continuous with μ, and the right hand side of (5.8) is continuous with μ. Hence around every λ there is guaranteed to be an open nontrivial ball $\mathcal{B}(\lambda)$ (with nonzero radius) in Λ such that $\mu \in \mathcal{B}(\lambda)$ implied that (5.8) holds. Suppose that a particular μ is in this ball. It follows from (5.8) that

$$\|T(P(\lambda), C(\lambda))\|_\infty \, (\delta_\nu(P(\mu), P(\lambda)) + \epsilon(\mu)) < \hat{r}.$$

Now by the triangle inequality we have that

$$\delta_\nu(P(\mu), P(\lambda)) + \delta_\nu(P(\mu), \bar{P}) \ge \delta_\nu(\bar{P}, P(\lambda)).$$

Hence by the definition of $\epsilon(\mu)$ in (5.8), it can be seen that, for every $\bar{P}$ such that $\delta_\nu(\bar{P}, P(\mu)) < \epsilon(\mu)$ there holds

$$\|T(P(\lambda), C(\lambda))\|_\infty \, \delta_\nu(\bar{P}, P(\lambda)) < \hat{r}, \tag{5.9}$$

that is, every $\bar{P}$ such that $\delta_\nu(\bar{P}, P(\mu)) < \epsilon(\mu)$, with μ obeying equation (5.8), is satisfactorily controlled not only by $C(\mu)$ (by definition) but also by $C(\lambda)$.

Theorem 5.1 *Let $\mathcal{P}$ be a set of plants with uncertainty parametrised by λ in a compact set Λ such that $\mathcal{P} = \{\bar{P} : \delta_\nu(\bar{P}, P(\lambda)) < \epsilon(\lambda), \lambda \in \Lambda\}$. The plants $P(\lambda)$ are understood to depend continuously on λ. For any given controller performance threshold $\hat{r}$, if there exists a nominal controller $C(\lambda)$ for each $P(\lambda), \lambda \in \Lambda$ such that equation (5.5) is satisfied, then there exists a finite set of controllers $C(\lambda_i)$ such that each plant $\bar{P} \in \mathcal{P}$ is satisfactorily controlled by some $C(\lambda_i)$ in the sense that equation (2.6) is satisfied.*

Proof : By the above arguments, for an arbitrary λ there is an open ball $\mathcal{B}(\lambda)$ around it such that $C(\lambda)$ will satisfactorily control all $\bar{P}$ such that $\delta_\nu(\bar{P}, P(\bar{\lambda}) < \epsilon(\bar{\lambda}), \bar{\lambda} \in \mathcal{B}(\lambda)$, and hence an infinite cover of Λ, corresponding to satisfactorily performing controllers for each plant in $\mathcal{P}$, exists. Once again, arguments involving the Heine-Borel property (Loomis and Sternberg, 1968) imply that there is a finite set of λ_i such that the associated balls cover Λ. Under these circumstances, the associated controllers for $i = 1, \ldots, N$ have the property that each plant $\bar{P} \in \mathcal{P}$ is satisfactorily controlled by one of the $C(\lambda_i)$. $\qquad\qquad\square$

Again, one could also just focus on stability, and maximise the ball sizes, by replacing equation (5.4) by equation (5.5) and by replacing equation (5.7) by

$$\delta_\nu(\bar{P}, P(\lambda)) \, \|T(P(\lambda), C^\star(\lambda))\|_\infty < 1,$$

with $C^\star(\lambda)$ defined as the stabilising controller achieving the infimum in equation (5.10) below

$$C^\star(\lambda) = \arg\min_C \|T(P(\lambda), C)\|_\infty. \tag{5.10}$$

6 Multiple model problem with large unstructured uncertainty

In the last section, we considered a situation in which there was parametric plant variation together with unstructured uncertainty. The unstructured uncertainty was limited in the sense that for any fixed parameter value, λ say, we assumed that a satisfactory controller could be found that not only gave satisfactory performance for $P(\lambda)$, but also for a ball of plants around $P(\lambda)$ with unstructured uncertainty, such as $\{\bar{P} : \delta_\nu(\bar{P}, P(\lambda)) < \epsilon(\lambda)\}$ around $P(\lambda)$. We pointed out that $\epsilon(\lambda)$ had to satisfy an upper bound (equation (5.5)) for this problem to be solvable.

What if this bound is not met? Then we might ask whether we could find a finite set of plants, in the ball $\{\bar{P} : \delta_\nu(\bar{P}, P(\lambda)) < \epsilon(\lambda)\}$ around $P(\lambda)$, such that satisfactory performance for each $\bar{P}$ in the ball could be obtained by connecting a controller giving satisfactory performance for one (or more) of the plants in this finite set.

Note that we are not talking here about finding a specific subset $\lambda_1, \lambda_2, \ldots \lambda_N$ of the uncertain parameter λ such that for each $\bar{P}$ with $\delta_\nu(\bar{P}, P(\lambda)) < \epsilon(\lambda)$ there exists a λ_i such that $\delta_\nu(\bar{P}, P(\lambda_i))$ is sufficiently small. Instead, we are looking at a finite set of plants with unstructured uncertainty balls of radius less than $\epsilon(\lambda)$ around them such that these smaller balls cover $\{\bar{P} : \delta_\nu(\bar{P}, P(\lambda)) < \epsilon(\lambda)\}$.

A slightly more abstract statement of this problem is as follows. Given a plant P_0 and a ball $\{\bar{P} : \delta_\nu(\bar{P}, P_0) < \epsilon\}$ of plants $\bar{P}$ around P_0, is there a finite M and a set $\{P_1, P_2, \cdots, P_M\}$ of P_j with $\delta_\nu(P_j, P_0) < \epsilon$ such that each $\bar{P}$ in the ball $\{\bar{P} : \delta_\nu(\bar{P}, P_0) < \epsilon\}$ satisfies $\delta_\nu(\bar{P}, P_j) < \frac{\epsilon}{2}$ for one or more j, that is

$$\{\bar{P} : \delta_\nu(\bar{P}, P_0) < \epsilon\} \subset \bigcup_{j=1}^{M} \left\{\bar{P} : \delta_\nu(\bar{P}, P_j) < \frac{\epsilon}{2}\right\}. \tag{6.1}$$

We will show that the answer to this question is negative; and then we argue that with additional modest restrictions, the answer is positive.

We will actually first study a similar, but not identical, problem, almost as relevant as the above, but with more intuitive content: given a stable scalar plant P_0 and an $\mathcal{H}_\infty$ ball $\{\bar{P} : \|\bar{P} - P_0\|_\infty < \epsilon\}$, is there a finite M and a set $\{P_1, P_2, \cdots, P_M\}$ of P_j with $\|P_j - P_o\|_\infty < \epsilon$ such that each $\bar{P}$ in the ball $\{\bar{P} : \|\bar{P} - P_0\|_\infty < \epsilon\}$ satisfies $\|P_j - \bar{P}\|_\infty < \frac{\epsilon}{2}$ for some j? The answer to this question is in the negative.

The general reason for the negative answer is that the set $\{\bar{P} : \|\bar{P} - P_0\|_\infty < \epsilon\}$ is not compact. Let us understand this in more detail using a contradiction argument. Pick any frequency, say ω_1. Then there are stable plants $P_{\omega_1}^+$ and $P_{\omega_1}^-$ such that

$$P_{\omega_1}^+(j\omega_1) - P_0(j\omega_1) = -\left[P_{\omega_1}^-(j\omega_1) - P_0(j\omega_1)\right] = d$$
$$\text{where } |d| = \frac{3\epsilon}{4}$$

and $\|P_{\omega_1}^\pm - P_0\| < \epsilon$. Observe that $|P_{\omega_1}^+(j\omega_1) - P_{\omega_1}^-(j\omega_1)| = \frac{3\epsilon}{2}$ and evidently, there cannot exist a single member of the set $P_1(j\omega_1), P_2(j\omega_1), \cdots, P_M(j\omega_1)$ such that $\left|P_{\omega_1}^+(j\omega_1) - P_j\right| < \frac{\epsilon}{2}$ and $\left|P_{\omega_1}^-(j\omega_1) - P_j\right| < \frac{\epsilon}{2}$ simultaneously. Hence there does not exist a single $P_j \in \{P_1, P_2, \ldots, P_M\}$ such that $\|P_{\omega_1}^+ - P_j\|_\infty < \frac{\epsilon}{2}$ and $\|P_{\omega_1}^- - P_j\|_\infty < \frac{\epsilon}{2}$. More generally, pick any m distinct frequencies $\omega_1, \cdots, \omega_m$ and recognise that one can find 2^m different plants, call them $\bar{P}_1, \cdots, \bar{P}_{2^m}$, that take the values of $P_0(j\omega_i) \pm \frac{3\epsilon}{4}$ at the frequencies $\omega_1, \omega_2, \ldots, \omega_m$ of $P_0(j\omega_i) \pm \frac{3\epsilon}{4}$ (in all possible combinations), with all plants lying in $\|P - P_0\|_\infty < \epsilon$. Moreover, these require at least 2^m different P_j if for each $i = 1, \cdots, 2^m$ there exists $j(i) \in \{1, 2, \cdots, M\}$ such that $\|P_j - \bar{P}_i\|_\infty < \frac{\epsilon}{2}$. Hence $M > 2^m$. But since m is arbitrary, this shows that a finite set of covering P_j can in general not exist. This is the lack of compactness mentioned above.

Let us now return to the first problem we posed, embodied in equation (6.1) for the Vinnicombe metric. We indicate first the negative result.

Theorem 6.1 *Let P_0 be a plant, with an associated ball $\{\bar{P} : \delta_\nu(\bar{P}, P_0) < \epsilon\}$. It is **not** possible to determine a finite M and plants P_j, $j = 1, \cdots, M$ in the ball such that for each $\bar{P}$ in the ball, $\{\bar{P} : \delta_\nu(\bar{P}, P_j) < \frac{\epsilon}{2}\}$ for at least one j.*

Proof : We assume that $\bar{P}$ is scalar. Recall that, subject to a winding number condition,

$$\delta_\nu(\bar{P}, P_0) = \sup_\omega \left[\frac{|\bar{P} - P_0|}{[1 + |\bar{P}|^2]^{\frac{1}{2}} [1 + |P_0|^2]^{\frac{1}{2}}}\right].$$

(with indentations into the right half plane at frequencies where $\bar{P}$ or P_0 has an imaginary axis pole.) Suppose that, at some frequency ω_1, $P_0(j\omega_1) = x_0 + jy_0$. Consider the set of plants for which $P(j\omega_1) = \lambda(x_0 + jy_0)$ where λ is a real constant. Define

$$g(\lambda_1, \lambda_2) = \frac{|(\lambda_2 - \lambda_1)(x_0 + jy_0)|}{[1 + \lambda_1^2 |x_0 + jy_0|^2]^{\frac{1}{2}} [1 + \lambda_2^2 |x_0 + jy_0|^2]^{\frac{1}{2}}}$$

As shown in Appendix A, Lemma A.1 and Corollary A.4, $\frac{\partial}{\partial \lambda_1} g(\lambda_1, 1) > 0$ for $\lambda_1 > 1$, and $\lim_{\lambda_1 \to \infty} g(\lambda_1, 1) = [1 + |x_0 + jy_0|^2]^{-\frac{1}{2}}$. When λ_1 moves in a negative direction below 1, $g(\lambda_1, 1)$ increases until some negative value of λ_m such that $\lambda_m = -|x_0 + jy_0|^{-2}$, where there holds $g(\lambda_m, 1) = 1$. It follows from the continuity of $g(\lambda_1, 1)$ and Corollary A.4 that one can choose a value of $\lambda_1 = \lambda_- < 1$ with such that

$$\frac{|(1 - \lambda_-)(x_0 + jy_0)|}{[1 + \lambda_-^2 |x_0 + jy_0|^2]^{\frac{1}{2}} [1 + |x_0 + jy_0|^2]^{\frac{1}{2}}} = \frac{3\epsilon}{4}.$$

14

Also, by taking $\lambda_1 = \lambda_-$ and increasing λ_2 through 1 to λ_+ for some $\lambda_+ > 1$, the monotonic increase of $g(\lambda_1, \lambda_2)$ with respect to λ_2 (see again Corollary A.4) gives

$$\frac{|(\lambda_+ - \lambda_-)(x_0 + jy_0)|}{[1 + \lambda_-^2 |x_0 + jy_0|^2]^{\frac{1}{2}} [1 + \lambda_+^2 |x_0 + jy_0|^2]^{\frac{1}{2}}} > \frac{3\,\epsilon}{4}.$$

Now imagine a smooth perturbation of $P_0(j\omega)$ in the vicinity of $\omega = \omega_1$, to see that there exist two plants $P^\pm$ with $P^\pm(j\omega_1) = \lambda_\pm (x_0 + jy_0)$, $\delta_\nu(P^\pm, P_0) < \epsilon$ and $\delta_\nu(P^+, P^-) > \frac{3\epsilon}{4}$. From this point, the lack of compactness argument follows the one used for $\mathcal{H}_\infty$ balls in the early part of this section. $\square$

Compactness may be secured by adding additional restrictions. The heuristic idea is that the set of transfer functions which are simultaneously (and necessarily approximately) band-limited and time-limited is approximately finite dimensional. We combine this idea with the fact that any bounded set in a finite dimensional space is compact, and thus any covering with open sets of such a bounded set has a finite sub-cover. Zames (1979) obtained relevant results. Let $\mathbf{B}$ denote the Banach space over the real field of complex valued, essentially bounded functions $F(j\omega)$ for $\omega \in (-\infty, \infty)$ with $F(j\omega) = \overline{F(-j\omega)}$ and $\|F\| = \sup_\omega |F(j\omega)|$. Let $\mathcal{H}(C, K, a)$ denote the subset of $\mathbf{B}$ satisfying the following conditions

- $F(j\omega) = \displaystyle\int_0^\infty f(t) e^{-j\omega t} dt$ for some exponentially bounded impulse response $f(t)$ obtained from the inverse Laplace Transform $\mathcal{L}^{-1}\{F(s)\}$, with $|f(t)| \le C e^{-at}, t \ge 0$.

- $|\omega F(j\omega)| \le K, \forall \omega \in \ \ .$

It is shown in Zames (1979) that given $\epsilon > 0$, there exists a finite integer $N(\epsilon, C, K, a)$, and N functions $G_1, G_2, \ldots G_N \in \mathbf{B}$ such that each $F \in \mathcal{H}(C, K, a)$ lies within an ϵ-ball of a least one of the G_j, that is, given arbitrary $F \in \mathcal{H}(C, K, a)$ there is at least one j for which $\sup_\omega |F - G_j| < \epsilon$.

Now the unhelpful feature of this result is that the G_j are not themselves necessarily in $\mathcal{H}(C, K, a)$, and we would like them to have this property. This defect may be remedied quite easily.

In order to obtain such a result, let us assume that the set $\{G_1, G_2, \ldots, G_M\}$ is minimal in the sense that for each j, the ball around G_j contains at least one plant $F_j \in \mathcal{H}(C, K, a)$ such that $\|F_j - G_j\|_\infty < \epsilon$. If the set is non-minimal, it can obviously be replaced by a smaller minimal set.

This is weaker than the condition that set $\{G_1, G_2, \ldots, G_M\}$ is not redundant, in the sense that for each j, there is some $F_j \in \mathcal{H}(C, K, a)$ such that $\|F_j - G_j\|_\infty < \epsilon$, and in addition that all the F_j are chosen so that $|F_i - G_k| \ge \epsilon$ for all $i \ne k$. If the set is redundant, then it can obviously also be replaced by a smaller set which is not redundant.

Corollary 6.2 *Suppose the set $\{G_1, G_2, \ldots, G_N\}$ is minimal. Then, for arbitrary ϵ there exists a finite set $\{F_1, F_2, \ldots, F_N\} \subset \mathcal{H}(C, K, a)$ such that each $F \in \mathcal{H}(C, K, a)$ satisfies $\|F - F_j\| < 2\epsilon$ for at least one j.*

Proof : There exists G_j with $\|F - G_j\|_\infty < \epsilon$. By minimality, there exists $F_j \in \mathcal{H}(C, K, a)$ with $\|F_j - G_j\|_\infty < \epsilon$. By the triangle inequality this implies the desired result. $\square$

Let us now identify some classes of unstructured uncertainty and restrictions on that uncertainty following from Zames (1979), which will recover the compactness property for the *Vinnicombe metric*. First, we list the classes. Then we shall establish the claims.

- All plants $\bar{P} \in \mathcal{P}$ defined as $\bar{P}$ such that

$$\delta_\nu(\bar{P}, P(\lambda)) \quad < \quad \eta \tag{6.2}$$

for a given $P(\lambda), \eta$, where both $\bar{P}$ and $P(\lambda)$ are restricted to be in $\mathcal{H}(C, K, a)$ for some given C, K, a (and are therefore stable).

- All plants $\bar{P} \in \mathcal{P}$ defined as $\bar{P}$ of the form

$$\bar{P} \quad = \quad (N(\lambda) + \Delta_N)(D(\lambda) + \Delta_D)^{-1} \tag{6.3}$$

$$\text{where } \begin{bmatrix} \Delta_N \\ \Delta_D \end{bmatrix} \quad \in \quad \Delta \subset \mathcal{H}(C, K, a) \text{ for some } C, K, a$$

$$\text{and } \begin{bmatrix} N(\lambda) \\ D(\lambda) \end{bmatrix} \text{ is normalised.}$$

Additionally, we require that Δ is a connected set (and is obviously bounded), and for all ω and all $\Delta_N, \Delta_D \in \Delta$ there holds

$$[N(\lambda) + \Delta_N]^*[N(\lambda) + \Delta_N] + [D(\lambda) + \Delta_D]^*[N(\lambda) + \Delta_D] \geq \alpha^2 I \tag{6.4}$$

for some $\alpha > 0$. (This condition ensures that all fractional representations of the form (6.3) are coprime, although not, of course, normalised.)

Theorem 6.3 *Let $P(\lambda)$ be a given plant which defines a set of plants $\mathcal{P}$ via equation (6.2) or (6.3) and the associated conditions. Given any $\epsilon > 0$, there exists a finite set of plants $P_1, P_2, \ldots, P_M$ with the P_j themselves in $\mathcal{P}$ such that each plant $\bar{P} \in \mathcal{P}$ above satisfies $\delta_\nu(\bar{P}, P_j) < \epsilon$ for some j.*

Proof : The proof for $\mathcal{P}$ defined by equation (6.2) is an immediate consequence of results in Zames (1979) applied to $\mathcal{H}(C, K, a)$ and the fact that

$$\delta_\nu(\bar{P}, P(\lambda)) \quad \leq \quad \|\bar{P} - P(\lambda)\|_\infty$$

when $\bar{P}, P(\bar{\lambda})$ are stable. We may ignore the condition (6.2) and just use the fact that $\mathcal{P} \subset \mathcal{H}(C, K, a)$. By Corollary 6.2, there exists a finite set of plants $\hat{\mathcal{F}} = \{\hat{F}_1, \hat{F}_2, \ldots, \hat{F}_M\} \subset \mathcal{H}(C, K, a)$ such that for all $\hat{P} \in \mathcal{H}(C, K, a)$ there is a $\hat{F}_j$ with $\|\hat{P} - \hat{F}_j\|_\infty < \epsilon$. By the above equation such a $\hat{F}_j$ also gives $\delta_\nu(\hat{P}, \hat{F}_j) < \epsilon$, and so $\mathcal{H}(C, K, a)$ is clearly shown to be able to be covered by a finite set of Vinnicombe metric balls of radius ϵ. Since the uncertainty set $\mathcal{P}$ defined by equation (6.2) and the associated conditions is a *subset* of $\mathcal{H}(C, K, a)$, the same assertion holds for $\mathcal{P}$.

For the proof for $\mathcal{P}$ defined by equation (6.3) and associated conditions, let P_1, P_2 be two plants in the set with

$$P_1 \quad = \quad [N(\bar{\lambda}) + \Delta_{N1}][D(\bar{\lambda}) + \Delta_{D1}]^{-1},$$
$$P_2 \quad = \quad [N(\bar{\lambda}) + \Delta_{N2}][D(\bar{\lambda}) + \Delta_{D2}]^{-1},$$

and let U be such that

$$\hat{P}_1 \quad = \quad \begin{bmatrix} (N(\bar{\lambda}) + \Delta_{N1})U \\ (D(\bar{\lambda}) + \Delta_{D1})U \end{bmatrix} \tag{6.5}$$

is normalised. Now

$$\delta_\nu(P_1, P_2) \quad = \quad \inf_{\substack{Q \in \mathcal{L}_\infty \\ wno \det Q \neq 0}} \left\| \begin{bmatrix} (N(\bar{\lambda}) + \Delta_{N1}) \\ (D(\bar{\lambda}) + \Delta_{D1}) \end{bmatrix} U - \begin{bmatrix} (N(\bar{\lambda}) + \Delta_{N2}) \\ (D(\bar{\lambda}) + \Delta_{D2}) \end{bmatrix} Q \right\|_\infty$$

$$\leq \quad \left\| \begin{bmatrix} (N(\bar{\lambda}) + \Delta_{N1}) \\ (D(\bar{\lambda}) + \Delta_{D1}) \end{bmatrix} U - \begin{bmatrix} N(\bar{\lambda}) + \Delta_{N2} \\ D(\bar{\lambda}) + \Delta_{D2} \end{bmatrix} U \right\|_\infty .$$

The condition (6.4) ensures that $\|U\|_\infty \leq \alpha^{-1}$ and hence

$$\delta_\nu(P_1, P_2) \quad \leq \quad \alpha^{-1} \left\| \begin{array}{c} \Delta_{N1} - \Delta_{N2} \\ \Delta_{D1} - \Delta_{D2} \end{array} \right\|_\infty . \tag{6.6}$$

Now consider the set $\hat{\mathcal{F}}$ such that $\hat{F} \in \hat{\mathcal{F}}$ if

$$\hat{F} = \begin{bmatrix} N(\bar{\lambda}) + \Delta_N \\ D(\bar{\lambda}) + \Delta_D \end{bmatrix}$$

for Δ_N, Δ_D satisfying both equation (6.4) and the conditions associated with equation (6.3). Thus by Corollary 6.2 there is a finite set of plants $\hat{\mathcal{F}} = \{\hat{F}_1, \hat{F}_2, \ldots, \hat{F}_M\} \subset \hat{\mathcal{F}}$ such that for all $\hat{F} \in \hat{\mathcal{F}}$ there is a $\hat{F}_j$ with

$$\|\hat{F} - \hat{F}_j\|_\infty = \left\| \begin{array}{c} \Delta_{NF} - \Delta_{NFj} \\ \Delta_{DF} - \Delta_{DFj} \end{array} \right\|_\infty < \alpha\epsilon, \tag{6.7}$$

for the $\Delta_{NF}, \Delta_{DF}, \Delta_{NFj}, \Delta_{DFj}$ corresponding to $\hat{F}, \hat{F}_j \in \mathcal{F}$. Now let F_j be defined by $F_j = F_{Nj}F_{Dj}^{-1}$ where

$$\hat{F}_j = \begin{bmatrix} F_{Nj} \\ F_{Dj} \end{bmatrix}$$

so that F_{Nj}, F_{Dj} are right coprime by condition (6.4). For any $P \in \mathcal{P}$, it follows by equation (6.6) that $\delta_v(P, F_j) < \epsilon$. Thus the set of F_j is a finite set with the desired property. $\qquad\square$

Remark 6.4 *The above restriction of allowable uncertainty using the concept of band- and time-limited transfer functions allows us to overcome the difficulties suggested in Georgiou, Pascoal, and Khargonekar (1987) where Theorems 4.1 and 5.1 show that there are some balls around a nominal plant (additive or multiplicative uncertainty) which cannot be stabilized by a single linear time-invariant controller. This occurs when then there is a sequence of plants P_n within the ball such that, as n goes to infinity, the Hankel norm of P_n converges to zero. It seems, however, that the above constraints on the uncertainty asserted by (i) the uniform exponential decay of the impulse response and (ii) by the 20dB decay of the Bode plot, are permitting us to avoid this possibility.*

Remark 6.5 *There are earlier results which are related. For example, it was shown in Vinnicombe (1996) that ball of plants in the Vinnicombe metric which also admit a given complexity bound ζ, that is $\left\{ \bar{P} : \delta_v(\bar{P}, P_0) < \eta, \sup_{\omega_1} \lim_{\omega_2 \to \omega_1} \frac{\kappa(\bar{P}, j\omega_1), \bar{P}, j\omega_2)}{\omega_1 - \omega_2} < \zeta \right\}$ is also relatively compact in the gap topology. In the above expression for complexity $\kappa(X, Y) = \sigma_{\max} \left((I + XX^*)^{-\frac{1}{2}} (X - Y)(I + Y^*Y)^{-\frac{1}{2}} \right).$*

7   Conclusions

We have shown that given an (possibly infinite) uncertainty set of plants where the uncertainty satisfies certain minor conditions, it is possible to find a finite set of plants such that at least one element of a finite set of corresponding controllers, is able to satisfactorily control each plant in the uncertainty set. We considered uncertainty sets where parametrised uncertainty is continuous in the Vinnicombe metric for the parameters, which are constrained to a compact set, and where unstructured uncertainty is constrained to be essentially bounded and approximately time- and band- limited. A method for finding one such finite set of plants with the desired property was presented. This places an upper bound on the number of plant models which it is sufficient to use for a multiple model switching adaptive control scheme. It further remains to find a computationally efficient method to determine such a finite set, so that the number of models needed is not overly conservative. The significance of these results in the switching adaptive control context will be further explored in the sequel to this paper (Hespanha et al., 2000).

A Calculations

A.1 Bounds on Performance Change due to Plant Change

Proof :[Proof of equation (2.5)] It is known (see Vinnicombe, 1993), that

$$
\begin{aligned}
\delta_\nu(\bar{P}, P(\lambda)) &\leq \|T(\bar{P}, C(\lambda)) - T(P(\lambda), C(\lambda))\|_\infty \\
&\leq \|T(\bar{P}, C(\lambda))\|_\infty \|T(P(\lambda), C(\lambda))\|_\infty \delta_\nu(\bar{P}, P(\lambda))
\end{aligned}
$$

The second inequality implies

$$
\|T(\bar{P}, C(\lambda)) - T(P(\lambda), C(\lambda))\|_\infty \leq
$$
$$
\left\{ \|T(\bar{P}, C(\lambda)) - T(P(\lambda), C(\lambda))\|_\infty + \|T(P(\lambda), C(\lambda))\|_\infty \right\} \|T(P(\lambda), C(\lambda))\|_\infty \delta_\nu(\bar{P}, P(\lambda))
$$
$$
\text{hence,} \ \|T(\bar{P}, C(\lambda)) - T(P(\lambda), C(\lambda))\|_\infty \leq \frac{\|T(P(\lambda), C(\lambda))\|_\infty^2 \delta_\nu(\bar{P}, P(\lambda)}{1 - \|T(P(\lambda), C(\lambda))\|_\infty \delta_\nu(\bar{P}, P(\lambda))}
$$

as required. □

A.2 Vinnicombe Metric Between Scaled Plants

Lemma A.1 *Let* $M(\zeta)$ *be a positive semidefinite matrix whose entries are a continuous function of* ζ. *If at a particular* $\zeta = \zeta_0$, $\frac{\partial}{\partial \zeta} M$ *is positive semi-definite then it follows that* $\frac{\partial}{\partial \zeta} \sigma_{\max}(M(\zeta)) \geq 0$ *at* ζ_0, *with strict inequality if* $\frac{\partial}{\partial \zeta} M$ *is positive definite.*

Proof : This can be derived from the definition of the maximum singular value and the properties of positive definite and positive semi-definite matrices (Horn and Johnson, 1985). □

Proposition A.2 *Let* M *be a complex matrix with* $M = P_0(j\omega)$ *and* $P_1(j\omega) = (\zeta + 1) M$ *with* ζ *real. Then* $\widehat{\delta_\nu}(P_0, P_1, j\omega)$ *is monotone strictly increasing with* $|\zeta|$ *on both* $\zeta > 0$ *and the interval* $-1 - [\sigma_{\max}(M)]^{-2} < \zeta < 0$.

Proof : Let $X = M^* M$, $\hat{X} = MM^*$ for the complex matrix M and let

$$
\begin{aligned}
Y &= [I + P_0(j\omega)P_0(j\omega)^*]^{-\frac{1}{2}} [P_1(j\omega) - P_0(j\omega)] [I + P_1(j\omega)^* P_1(j\omega)]^{-\frac{1}{2}} \ \text{and} \\
\hat{Y} &= YY^* \\
&= \zeta^2 (I + \hat{X})^{-\frac{1}{2}} M (I + (\zeta + 1)^2 X)^{-1} M^* (I + \hat{X})^{-\frac{1}{2}} \ \text{after some algebra.}
\end{aligned}
$$

$$
\begin{aligned}
\text{Define} \ \bar{Y} &= \zeta^2 (I + (\zeta + 1)^2 X)^{-1} \\
\text{so that} \ \hat{Y} &= (I + \hat{X})^{-\frac{1}{2}} M \bar{Y} M^* (I + \hat{X})^{-\frac{1}{2}}. \\
\text{Then} \ \frac{\partial}{\partial \zeta} \bar{Y} &= 2\zeta \left[I - \zeta(\zeta + 1)(I + (\zeta + 1)^2 X)^{-1} \right] (I + (\zeta + 1)^2 X)^{-1} \\
&= 2\zeta (I + (\zeta + 1)^2 X)^{-1} (I + (\zeta + 1) X)(I + (\zeta + 1)^2 X)^{-1} \\
\frac{\partial}{\partial \zeta} \bar{Y} & \begin{cases} \geq 0 & \text{for } \zeta \geq 0, \\ \leq 0 & \text{for } -1 - [\sigma_{\max}(X)]^{-1} \leq \zeta \leq 0 \end{cases} \\
\text{Hence} \ \frac{\partial}{\partial \zeta} \hat{Y} & \begin{cases} \geq 0 & \text{for } \zeta \geq 0, \\ \leq 0 & \text{for } -1 - [\sigma_{\max}(P_0(j\omega))]^{-2} \leq \zeta \leq 0 \end{cases}
\end{aligned}
$$

By the previous lemma, Lemma A.1 this implies that the same holds for $\frac{\partial}{\partial \zeta} \sigma_{\max}(\hat{Y})$ and hence for $\frac{\partial}{\partial \zeta} \sigma_{\max}(Y)$ for the range of ζ in the proposition statement. Since $\widehat{\delta_\nu}(P_0, P_1, j\omega) = \sigma_{\max}(Y)$ the result follows directly. □

Corollary A.3 *For* $P_1(s) = (1 + \zeta)P_0(s)$, $\zeta \in$, *the Vinnicombe distance* $\delta_\nu(P_0, P_1(s)) \geq 0$, *is also monotone strictly increasing with* $|\zeta|$ *for all* ζ *in the range* $-1 - \|P_0(s)\|_\infty^{-2} < \zeta$.

Proof : We have $\delta_\nu(P_1, P_0) = \sup_{\omega \in} \widehat{\delta_\nu}(P_0, P_1, j\omega)$. At each frequency $\widehat{\delta_\nu}(P_0, P_1, j\omega)$ is monotone strictly increasing with $|\zeta|$ for $-1 - \sigma_{\max}(P_0(j\omega))^{-2} < \zeta$. The Corollary statement follows immediately.
$\square$

Corollary A.4 *For a scalar plant* $P_0(s)$, *let* $P_0(j\omega) = p_0$ *be an arbitrary point in the complex plane. Let* $P_1(j\omega) = (1 + \zeta)p_0$ *with* ζ *real. As* ζ *increases from 0 to* $+\infty$, *the distance* $\delta_\nu(P_0, P_1, j\omega)$ *increases strictly monotonically to* $[1 + |p_0|^2]^{-\frac{1}{2}}$. *When* ζ *decreases from unity and then through* -1, $\delta_\nu(P_0, P_1, j\omega)$ *increases strictly monotonically to* 1 *when* $\zeta = -1 - |p_0|^{-2}$.

Proof : By substitution of $\zeta \to \infty$ and $1 + \zeta = -|p_0|^{-2}$ into the expression for $\hat{Y}$ in Proposition A.2.
$\square$

References

J. Doyle, B. Francis, and A. Tannenbaum. *Feedback Control Theory*. Macmillan Publishing Company, New York, 1992.

A. K. El-Sakkary. The gap metric: Robustness of stabilization of feedback systems. *IEEE Trans. Auto. Control*, AC-30(3):240–247, March 1985.

T. T. Georgiou, A. Pascoal, and P. P. Khargonekar. On the robust stabilizability of uncertain linear time-invariant plants using nonlinear time-varying controllers. *Automatica*, 35(5):617–624, September 1987.

T. T. Georgiou and M. C. Smith. Optimal robustness in the gap metric. *IEEE Trans. Auto. Control*, 35(6):673–687, June 1990.

K. Glover and D. McFarlane. Robust stabilization of normalized coprime factor plant descriptions with $\mathcal{H}_\infty$-bounded uncertainty. *IEEE Trans. Auto. Control*, 34(8):821–830, August 1989.

M. Green and D. J. N. Limebeer. *Linear Robust Control*. Prentice Hall, NJ, 1995.

J. P. Hespanha, D. Liberzon, A. S. Morse, B. D. O. Anderson, T. S. Brinsmead, and F. De Bruyne. Multiple model adaptive control, part 2: Switching. *International Journal on Robust and Nonlinear Control*, 2000. submitted.

R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, 1985.

L. H. Loomis and S. Sternberg. *Advanced Calculus*. Addison-Wesley, 1968.

A. S. Morse. Supervisory control of families of linear set-point controllers— part 1: Exact matching. *IEEE Trans. Auto. Control*, 41(10):1413–1431, October 1996.

A. S. Morse. A bound for the disturbance-to-tracking-error gain of a supervised set-point control system. In D. Normand-Cyrot, editor, *Perspectives in Control, Theory and Applications*, pages 23–41. Springer-Verlag, 1998a.

A. S. Morse. Control using logic-based switching. In A. Isidori, editor, *Trends in Control: A European Perspective*, pages 69–113. Springer-Verlag, 1998b.

K. S. Narendra and J. Balakrishnan. Adaptive control using multiple models. *IEEE Trans. Auto. Control*, 42(2):171–187, February 1997.

F. M. Pait and F. Kassab. Parallel algorithms for adaptive control: robust stability. In A. S. Morse, editor, *Control Using Logic-Based Switching*. Springer-Verlag, London, 1997.

F. M. Pait and F. Kassab. On a class of switched, robustly stable, adaptive systems. *Int. J. on Adaptive Control and Signal Processing*, 2001. to appear.

A. I. G. Vardulakis. *Linear Multivariable Control: Algebraic Analysis and Synthesis Methods*. Wiley, Chichester, 1991.

G. Vinnicombe. On the frequency response interpretation on an indexed $\mathcal{L}_2$ gap metric. *Proc. Amer. Control Conf.*, 1992.

G. Vinnicombe. Frequency domain uncertainty and the graph topology. *IEEE Trans. Auto. Control*, 38(9):1371–1383, September 1993.

G. Vinnicombe. The robustness of feedback systems with bounded complexity controllers. *IEEE Trans. Auto. Control*, 41(6):795–803, June 1996.

G. Vinnicombe. *Uncertainty and Feedback: $\mathcal{H}_\infty$ Uncertainty and the ν-gap metric*. to appear, 1999.

G. Zames. On the metric complexity of causal linear systems: ϵ-entropy and ϵ-dimension for continuous time. *IEEE Trans. Auto. Control*, AC-24(2):222–230, April 1979.

G. Zames and A. K. El-Sakkary. Unstable systems and feedback: The gap metric. *Proc. Allerton Conf.*, pages 380–385, 1980.

K. Zhou, J. C. Doyle, and K. Glover. *Robust and Optimal Control*. Prentice Hall, NJ, 1996.