Model and Controller Selection Policies Based on Output Prediction Errors

Sanjeev R. Kulkarni, *Senior Member, IEEE*, and Peter J. Ramadge, *Fellow, IEEE*

Abstract—**Based on observations of the past inputs and outputs of an unknown system Σ, a countable set of predictors, $O_p, p \in \mathcal{P}$, is used to predict the system output sequence. Using performance measures derived from the resultant prediction errors, a decision rule is to be designed to select a $p \in \mathcal{P}$ at each time k. We study the structure and memory requirements of decision rules that converge to some $q \in \mathcal{P}$ such that the qth prediction error sequence has desirable properties, e.g., is suitably bounded or converges to zero. In a very general setting we give a positive result that there exist stationary decision rules with countable memory that converge (in finite time) to a "good" predictor. These decision rules are robust in a sense made precise in the paper. In addition, we demonstrate that there does not exist a decision rule with finite memory that has this property. This type of problem arises in a variety of contexts, but one of particular interest is the following. Based on the decision rule's selection at time k, a controller for the system Σ is chosen from a family $\Gamma_p, p \in \mathcal{P}$ of predesigned control systems. We show that for certain multi-input/multi-output linear systems the resultant closed-loop controlled system is stable and can asymptotically track an exogenous reference input.**

I. Introduction

Suppose we can observe the inputs and outputs of a fixed but unknown discrete-time system Σ, and we are presented with two problems: the selection of a good predictor for Σ from a given family of predictors $\{O_p, p \in \mathcal{P}\}$ and the selection of a good controller for Σ from a set of controllers derived from these predictors.

We assume that the predictors are driven by the past inputs and outputs of Σ and produce sequences of output predictions $\{\hat{y}_p(k)\}$. For each predictor we compute a real-valued performance measure $J_p(k)$ based on the prediction errors $e_p(k) = \hat{y}_p(k) - y(k)$, $p \in \mathcal{P}$. For example, for fixed $\lambda \in (0,1)$ we might set $J_p(k) = \sum_{j=0}^{k} \lambda^{k-j} \|e_p(j)\|$. Then we seek a decision rule that at each time k uses the available predictor performance measures to select $d(k) \in \mathcal{P}$ such that $q = \lim_{k\to\infty} d(k)$ exists and $J_q(k)$ is suitably bounded or converges to zero. In general, the performance measures $J_p(k)$ will not converge and may not be consistently ordered. Hence our selection criterion will need to be based on features such as limit suprema or local variations. We consider some specific performance measures and selection criteria and analyze the properties of the final predictor selected.

Our first main goal is to study the basic capabilities and limitations of decision rules in a general framework which can then be applied to the predictor selection problem as a special case. For this we define three relevant decision problems and study the existence and complexity of successful decision rules for these problems. For example, we demonstrate that there is a stationary rule with countable memory that will converge (in finite time since $\mathcal{P}$ is countable) to a parameter $q \in \mathcal{P}$ such that the qth data sequence has finite limit supremum. Furthermore, we demonstrate that there is no decision rule with finite memory that has this property. Similarly, for our other decision problems we show that there exist successful decision rules with countable memory, and we give tight lower bounds on the memory requirements of all successful rules.

An application of particular interest is the following. For each $p \in \mathcal{P}$ we are given a controller Γ_p related to the predictor O_p. We use these controllers in a control policy of the following form. If $d(k)$ is the decision rule's predictor selection at time k, then controller $\Gamma_{d(k)}$ is connected in feedback with the unknown plant Σ at time k. This "supervised" control system must ensure that the plant state, input, and output trajectories are well behaved, e.g., remain bounded, and that the output of the system asymptotically satisfies a performance criterion with respect to some class of admissible exogenous signals.

For multi-input/multi-output (MIMO) linear time-invariant (LTI) systems, we show that there exist convergent controller selection rules based on the predictor performance measures such that if the individual controllers have been designed to adequately track an admissible exogenous signal, then within the limits imposed by disturbances so will the supervised system.

Predictor-based controller switching policies have recently been examined in a variety of contexts, e.g., [9], [7], [10], [11], [1], and [13]. Of particular relevance here is the work of [7]. In [7] a family of concurrently operating LTI predictors are used to predict the output of the plant and the resulting prediction errors are used to form a performance measure for each predictor. Then, at a sequence of sample times, one compares the performance of the predictors and selects the controller corresponding to the best predictor at that time. The sequence of selected predictors is not required to converge and in general will not do so. Nevertheless, the system variables remain bounded, and the output of the single-input/single-output system asymptotically tracks a constant set-point. In [9] and [13] switching is used to select a controller structure

Manuscript received January 5, 1995; revised May 10, 1996. Recommended by Associate Editor, W.-B. Gong. This work was supported in part by the National Science Foundation under Grants IRI-9457645 and ECS-9216450, by EPRI under Grant RP8030-18, and by the Army Research Office under Grant DAAL03-92-G-0320.

The authors are with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: kulkarni@princeton.edu).

Publisher Item Identifier S 0018-9286(96)08253-0.

matched to the similarity invariants of the plant, and in [1], [10], and [11] it is used to improve the transient performance of stable adaptive control schemes.

Alternative controller search methods, e.g., [12], [5], [3], [2], and [6], have used a "prerouted" search through the controllers. The "prerouted" search procedures fall into the following basic pattern. A scheduling or routing sequence $\pi: \mathbb{N} \rightarrow \mathcal{P}$ is designed so that it has the "revisitation property" that every $p \in \mathcal{P}$ appears infinitely often in π. The sequence π is the predetermined order in which the candidate controllers will be tried. Then a switching logic is designed that uses the plant inputs and outputs to determine if and when the next element of π should be tried. So the actual time spent "dwelling" on controller $\pi(i)$ will depend on the observed plant variables, but the order in which the candidate controllers are examined does not. For obvious reasons this form of search is usually impractical unless $\mathcal{P}$ is a small finite set. However, these methods have been useful for showing the existence of convergent adaptive control algorithms under weak *a priori* plant assumptions. Good reviews of this work are given in [6] and [4].

In the work reported here we combine both attributes of the predictor performance-based search with the finite time-convergence property of a prescheduled search. Our results are complementary to those of [7] where convergence is not required. The advantage of a convergent policy is that it is possible to examine the limiting controlled system. However, convergence does not imply that the decision rule terminates; if the plant is changed, the rule can again begin switching. Although our results are developed in the context of discrete-time systems, they apply to continuous-time systems as well.

In summary, the main contributions of the paper are as follows.

- We demonstrate successful decision rules for three general decision problems relevant for on-line model and controller selection (Theorem 3.1).
- We provide tight lower bounds for the memory requirements of successful decision rules for the above problems (Theorem 3.3).
- We illustrate the application of these results to the predictor selection problem.
- For LTI systems we give a precise relationship between the performance of a given predictor on a plant and the concurrent performance of a corresponding controller (Proposition 4.1).
- We prove the existence of convergent controller selection policies based on prediction errors that stabilize a plant and ensure that asymptotically the plant adequately tracks a reference signal (Theorem 4.2).

The remainder of the paper is organized as follows. In Section II we set out the framework for the predictor selection problem. We consider specific choices for performance measures and selection requirements and show the implications for the selected prediction error sequence. The design of selection rules for these problems follows from results in Section III, where we formalize and study the existence of decision rules for three basic decision problems. The main results are Theorem 3.1, proving existence, and Theorem 3.3, demonstrating the minimum memory requirements. In Section IV we consider controller selection policies. The main results are Proposition 4.1 relating prediction to control and Theorem 4.2 demonstrating the existence of successful controller switching policies for several MIMO LTI control problems.

II. Predictor Selection

We are given a fixed but unknown discrete-time system Σ described by

$$\begin{aligned} x(k+1) &= A(x(k), u(k), k), \quad x(k_0) = x_0 \\ y(k) &= C(x(k), k) \end{aligned} \tag{1}$$

where $x(k) \in \mathbb{R}^n$, $u(k) \in \mathbb{R}^m$, $y(k) \in \mathbb{R}^l$ and $A: \mathbb{R}^{n \times m} \times \mathbb{N} \rightarrow \mathbb{R}^n$, $C: \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}^l$. The initial condition x_0 is an unknown point in $\mathbb{R}^n$. We make no assumptions on the input sequence to Σ. For simplicity we often assume that the initial time k_0 is zero.

A predictor O_p for Σ is a sequence of functions $O_k^p: \mathbb{R}^{kl} \times \mathbb{R}^{km} \rightarrow \mathbb{R}^l$ with the prediction at time k given by $\hat{y}_p(k) = O_k^p(y(k-1), \cdots, y(0), u(k-1), \cdots, u(0))$. The associated prediction error is $e_p(k) \triangleq \hat{y}_p(k) - y(k)$, $k \geq 0$. In practice it may be convenient to let $\hat{y}_p(k)$ be a function of a state variable $w_p(k)$ taking values in a finite-dimensional vector space, i.e., a system O_p with

$$\begin{aligned} w_p(k+1) &= M_p(w_p(k), y(k), u(k), k), \quad w(0) = w_0 \in \mathbb{R}^n(p) \\ \hat{y}_p(k) &= C_p(w_p(k), k) \end{aligned} \tag{2}$$

where $w_p(k) \in \mathbb{R}^{n(p)}$, $u(k) \in \mathbb{R}^m$, $y(k) \in \mathbb{R}^l$, and $M_p: \mathbb{R}^{n(p) \times m \times l} \times \mathbb{N} \rightarrow \mathbb{R}^{n(p)}$, $C_p: \mathbb{R}^{n(p)} \times \mathbb{N} \rightarrow \mathbb{R}^l$. A system of the form (2) is easy to design, for example, if the systems of interest are linear and time invariant.

It is not our task here to deal with the existence or design of the predictors O_p. Rather, we focus on decision rules that at each time utilize the output errors of the predictors to select a model from $\mathcal{P}$.

Specifically, let $J^k: \mathbb{R}^{l \times (k+1)} \rightarrow \mathbb{R}^+$ be a sequence of real-valued performance functions, and for each $p \in \mathcal{P}$ let $J_p(k) \triangleq J^k(e_p(0), \cdots, e_p(k))$. The real number $J_p(k)$ measures the performance of predictor p at time k, with a smaller value indicating better performance. At each time k we observe the performance indexes of all the predictors $p_1, p_2, \cdots$, i.e., we observe $\bar{J}(k) \triangleq [J_{p_1}(k), J_{p_2}(k), \cdots]$ and then select a model from $\mathcal{P}$. In general, the performance measures $J_p(k)$ will not converge and may not be consistently ordered. Nevertheless, the sequence of selections is required to converge to a fixed element of $\mathcal{P}$ (in finite time since $\mathcal{P}$ is discrete) with the corresponding performance measure being acceptable.

For example, assume that for some $p^* \in \mathcal{P}$, $\|e_{p^*}(k)\| \leq Lg_1(k) + g_2(k)$ for known nonnegative sequences g_1, g_2, and some (possibly unknown) $L \in \mathbb{R}$. The term $Lg_1(k)$ models the effect of unknown initial conditions and $g_2(k)$ models asymptotic behavior. We assume throughout that $g_1 \in l^1$.

To form a predictor performance measure let $\{a(k), k \geq 0\}$ be a nonnegative sequence with not all terms zero and set $J^k(e(0),\cdots,e(k)) = \sum_{j=0}^k a(k-j)\|e(j)\|$. If $h_i(k) = \sum_{j=0}^k a(k-j)g_i(j)$, $i=1,2$, then

$$J_{p^*}(k) = \sum_{j=0}^{k} a(k-j)\|e_{p^*}(j)\| \le Lh_1(k) + h_2(k).$$

There are several possibilities depending on the assumptions we place on g_2 and a.

- If $g_2 \equiv 0$ and $a \in l^\infty$, then $h_2 \equiv 0$ and h_1 is bounded. Thus $J_{p^*}(k)$ has a finite limit supremum. In this case, we might seek a decision rule that will converge to some $q \in \mathcal{P}$ with $\limsup_{k\to\infty} J_q(k) < \infty$. Now if a is persistent in the sense that there exists $\nu > 0$ and an integer N such that for every $k \ge 0$, $\sum_{j=k}^{k+N-1} a(j) > \nu$, it can then be shown that $e_q \in l^1$ and hence $\lim_{k\to\infty} e_q(k) = 0$.
- If $a \in l^1$ and again $g_2 \equiv 0$, then $h_1 \in l^1$ and $J_{p^*}(k)$ converges to zero at rate $h = h_1$. In this case, we might seek a decision rule that converges to $q \in \mathcal{P}$ with $J_q(k) \le Lh(k)$ for some $L > 0$. Since $h \in l^1$, it would then follow that $J_q \in l^1$ and hence that $e_q \in l^1$.
- If g_2 is a bounded sequence and $a \in l^1$, then h_2 is a bounded sequence and $J_{p^*}(k)$ has a finite limit supremum. If a bound M on $\limsup_{k\to\infty} g_2(k)$ is known, then for a given $\varepsilon > 0$, we might seek a decision rule that will converge to $q \in \mathcal{P}$ with $\limsup_{k\to\infty} J_q(k) \le M + \varepsilon$. Since a is nonnegative with not all terms zero, it would then follow that e_q is a bounded sequence with $\limsup_{k\to\infty} \|e_q(k)\| \le K(M+\varepsilon)$, where $K = 1/\|a\|_\infty$.

Reasoning similar to that above applies to any performance function of the form $J^k(e(0),\cdots,e(k)) = \sum_{j=0}^k a(k-j)f(\|e(j)\|)$, where $f\colon [0,\infty) \to [0,\infty)$ with $f(0) = 0$, f nondecreasing, and $f(x) > 0$ for $x > 0$. In this case, we assume that $\sum_{k=0}^\infty f(g_1(k)) < \infty$. For example, if $f(x) = x^2$, $g_2 \equiv 0$ and $\sum_{k=0}^\infty g_1^2(k) < \infty$, then we seek a decision rule that will converge to some $q \in \mathcal{P}$ with $e_q \in l^2$.

The existence and complexity of decision rules to solve the above problems is examined in the next section. It is shown in Theorem 3.1 below, for example, that there exist decision rules based on the performance functions discussed above that converge to a predictor with good performance.

III. Decision Rules

In this section, we pose the design of decision rules in a general framework which will then allow us to treat the predictor selection problem outlined in the previous section as a special case. Our goal is to obtain results on the basic limitations and complexity requirements of convergent decision rules.

A data sequence $\bar{J} \triangleq \{\bar{J}(k), k \ge 0\}$, with $\bar{J}(k) \in (\mathbb{R}^+)^{|\mathcal{P}|}$, is given, and at each time $k \ge 0$ we are told the value of $\bar{J}(k)$. Based on the presented data we must select a parameter $d(k) \in \mathcal{P}$ such that the sequence $\{d(k), k \ge 0\}$ converges, and the limit satisfies a given criterion of success. To complete the problem specification we need to define how a decision rule is to work, what constraints are placed on the data sequence $\bar{J}$, and what is the criterion of success.

We require each admissible data sequence $\bar{J}$ to be an element of some given collection $\mathcal{F}$ of data sequences. Associated with $\mathcal{F}$ is a criterion of correct selection. This is a map S that maps a pair $\bar{J} \in \mathcal{F}$ and $p \in \mathcal{P}$ to the value one if p is a suitable limit for our decision rule under the input sequence $\bar{J}$ and the value zero otherwise.

Let $\mathcal{P}^{\mathbb{N}}$ denote the set of all sequences of elements of $\mathcal{P}$. Then a decision rule is defined to be a causal map from $\mathcal{F}$ into $\mathcal{P}^{\mathbb{N}}$. We regard $\bar{J} \in \mathcal{F}$ as the input to the rule and the sequence of decisions $d(k), k \ge 0$, as the output.

A decision rule *with memory* may depend on the entire past history of observations and explicitly on the time k. Such rules can be written as a sequence of mappings $\phi_k\colon \mathbb{R}^{|\mathcal{P}|\times(k+1)} \to \mathcal{P}$, where $\phi_k(\bar{J}(0),\cdots,\bar{J}(k))$ represents the decision at time k. A natural restriction is to insist that past observations be summarized by an element of a set V. In this case, a decision rule can be realized as a sequence of state transition mappings $s_k\colon \mathbb{R}^{|\mathcal{P}|} \times V \to V$ and an initial state $v_0 \in V$, together with a sequence of mappings $\phi_k\colon \mathbb{R}^{|\mathcal{P}|} \times V \to \mathcal{P}$. The decisions are then formed by setting

$$v(k+1) = s_k(\bar{J}(k), v(k)), \quad v(0) = v_0 \tag{3}$$

$$d(k) = \phi_k(\bar{J}(k), v(k)). \tag{4}$$

Here $v(k+1)$ denotes the state of the rule after the kth observation and $d(k)$ denotes the decision at time k using the new observation $\bar{J}(k)$. If for some $s\colon \mathbb{R}^{|\mathcal{P}|} \times V \to V$ and $\phi\colon \mathbb{R}^{|\mathcal{P}|} \times V \to \mathcal{P}$ we have $s_k = s$ and $\phi_k = \phi$, then the decision rule is said to be stationary.

To understand the structure and complexity of supervisory control systems necessary to perform given tasks, we impose restrictions on V. We are particularly interested in the case when V is a finite or countable set. When V is finite, the decision rule can be implemented by a finite automaton. In this case, the entire past data must be summarized by a finite state and the cardinality of V represents the size of the memory.

We now formalize the concept of a successful decision rule as follows.

Definition 1: We call a decision rule *successful* for $(\mathcal{F}, \mathcal{S})$ if for every observation sequence $\bar{J} \in \mathcal{F}$ the corresponding sequence of decisions $d(k)$ satisfies:

i) $q \triangleq \lim_{k\to\infty} d(k)$ exists;
ii) $\mathcal{S}(\bar{J}, q) = 1$.

Clearly, as the size of the class $\mathcal{F}$ of admissible sequences is increased, the more demanding becomes the task of the decision rule. Similarly, for each $\bar{J}$ as the set of $p \in \mathcal{P}$ for which $\mathcal{S}(\bar{J}, p) = 1$ decreases, the requirements for a successful decision rule become more stringent. Thus, the complexity of successful decision rules may increase as the size of $\mathcal{F}$ increases or, for a fixed $\mathcal{F}$, as the set on which $\mathcal{S}(\bar{J}, p) = 1$ decreases. On the other hand, the larger the size of $\mathcal{F}$, the more robust, in an intuitive sense, is a successful rule, and the smaller the set on which $\mathcal{S}(\bar{J}, p) = 1$, the stronger are the properties of the limit.

We now give three specific examples of admissible data sequences and their corresponding success criteria. These classes are sufficiently general to cover a variety of interesting situations. They will be the focus of attention in the sequel.

Finite Limit Supremum Observations: This set, denoted $\mathcal{F}_{\text{FLS}}$, consists of all data sequences $\bar{J}$ with the property that there exists at least one $p \in \mathcal{P}$ such that $\limsup_{k\to\infty} J_p(k) < \infty$. The success criterion $\mathcal{S}_{\text{FLS}}$ for $\mathcal{F}_{\text{FLS}}$ is simply $\mathcal{S}_{\text{FLS}}(\bar{J},p) = 1 \Leftrightarrow \limsup_{k\to\infty} J_p(k) < \infty$. So, a decision rule is successful for $(\mathcal{F}_{\text{FLS}}, \mathcal{S}_{\text{FLS}})$ if the decisions converge to some $q \in \mathcal{P}$ with finite limsup (but not necessarily the minimum limsup).

Finite Limsup Observations with Known Bound: For a known $M < \infty$, let $\mathcal{F}_{\text{FLS}}^{M}$ be the subset of $\mathcal{F}_{\text{FLS}}$ for which there exists at least one $p \in \mathcal{P}$ such that $\limsup_{k\to\infty} J_p(k) \le M$. For a fixed $\epsilon > 0$, let $\mathcal{S}_{\text{FLS}}^{M+\epsilon}$ be defined as $\mathcal{S}_{\text{FLS}}^{M+\epsilon}(\bar{J},p) = 1 \Leftrightarrow \limsup_{k\to\infty} J_p(k) \le M+\epsilon$. Hence, for this problem, we know that at least one sequence has limsup less than or equal to M, but to succeed we need only find one that has limsup less than or equal to $M+\epsilon$.

Finite Limsup Observations with Known Bound and Known Rate: Given a known $M < \infty$ and a known positive sequence $\bar{g} = \{g(k)\}$ with $\lim_{k\to\infty} g(k) = 0$, let $\mathcal{F}_{\text{FLS}}^{M,\bar{g}}$ consist of all data sequences $\bar{J}$ for which there exists at least one $p \in \mathcal{P}$ and some constant $L < \infty$ (that is unknown and can depend on $\bar{J}$) such that $J_p(k) \le M + Lg(k)$. The success criterion corresponding to this class will be denoted $\mathcal{S}_{\text{FLS}}^{M,\bar{g}}$ and is defined by $\mathcal{S}_{\text{FLS}}^{M,\bar{g}}(\bar{J},p) = 1 \Leftrightarrow \exists L < \infty$ such that $J_p(k) \le M + Lg(k)$. For this class we know that for some $p \in \mathcal{P}$, the tails of $J_p(k)$ decay to less than or equal to M at some known rate, but we may not know the constant. This will be useful for problems in which we have guarantees on the rate at which data sequences converge, but the constant may depend on some unknown parameters. The success criterion requires that we converge to some $q \in \mathcal{P}$ for which $J_q(k)$ also decays to less than or equal to M at the same rate as for p (although the constants may differ).

A. Existence Results

We now show some positive results for the three decision problems discussed above. Theorem 3.1 asserts the existence of successful decision rules for these problems and the proof is constructive. The rules constructed in the proof are kept simple to aid clarity; they are neither unique nor optimal. They contain, as a component, a prerouted path. However, this can be used in conjunction with the performance data to determine the actual search path through the candidate predictors. Thus the actual trajectory $d(k)$ through $\mathcal{P}$ is, in general, "sample path" dependent.

Theorem 3.1: Let $\mathcal{P}$ be a finite or countable index set.

1) There exists a stationary decision rule with countable memory that is successful for $(\mathcal{F}_{\text{FLS}}, \mathcal{S}_{\text{FLS}})$.
2) For every $M < \infty$ and every $\epsilon > 0$, there exists a stationary decision rule with memory of size $|\mathcal{P}|$ that is successful for $(\mathcal{F}_{\text{FLS}}^{M}, \mathcal{S}_{\text{FLS}}^{M+\epsilon})$.
3) For every $M < \infty$ and every positive sequence $\bar{g} = \{g(k)\}$ with $g(k) \to 0$, there exists a decision rule with countable memory that is successful for $(\mathcal{F}_{\text{FLS}}^{M,\bar{g}}, \mathcal{S}_{\text{FLS}}^{M,\bar{g}})$.

Proof (Part 1): Let $\pi_1, \pi_2, \cdots \in \mathcal{P}$ be such that for all $p \in \mathcal{P}$, $\pi_i = p$ for infinitely many i, and let $h_1, h_2, \cdots$ be any nondecreasing positive sequence with $\lim_{m\to\infty} h_m = \infty$. Then define a stationary decision rule as follows. Let the initial memory state be $v(0) = 1$, and for $k \ge 1$ let $s(\bar{J}(k), v(k)) = \min\{i \ge v(k) : J_{\pi_i}(k) \le h_i\}$ and let $\phi(\bar{J}(k), v(k)) = \pi_{v(k)}$. Note that the properties of the π and h sequences ensure that the minimum in the definition of $s(\bar{J}(k), v(k))$ exists. The update rule for $v(k)$ is simply to take $v(k+1) = v(k)$ if the observation corresponding to the current selection $\pi_{v(k)}$ does not exceed the current bound $h_{v(k)}$ and to increment $v(k)$ as little as possible otherwise.

We now show that this decision rule is successful for $(\mathcal{F}_{\text{FLS}}, \mathcal{S}_{\text{FLS}})$. Let $\bar{J}$ be any data sequence such that $J^* \triangleq \limsup_{k\to\infty} J_{p^*}(k) < \infty$ for some $p^* \in \mathcal{P}$. Then there exists $N_1 < \infty$ such that $J_{p^*}(k) < J^* + 1$ for all $k \ge N_1$. Note that on any input sequence we have $v(k) < \infty$ for all $k < \infty$, where the exact value of $v(k)$ depends on $\bar{J}$. In particular, $v(N_1) = v(N_1; \bar{J}) < \infty$. Let $N_2 = \min\{n \ge v(N_1): h_n > J^* + 1 \textit{ and } \pi_n = p^*\}$. That $N_2 < \infty$ is well defined follows from the properties of the sequences h and π. Now we have $v(N_1) \le N_2$, and for all $k \ge N_1$ we have

$$J_{\pi_{N_2}}(k) = J_{p^*}(k) \le J^* + 1 \le h_{N_2}.$$

Therefore, by the definition of the state transition rule, $v(k) \le N_2$ for all $k \ge N_1$. Moreover, since $v(k)$ is nondecreasing, it follows that $\lim_{k\to\infty} v(k) = v \le N_2$. Thus, $\lim_{k\to\infty} \phi(\bar{J}(k), v(k)) = \pi_v \triangleq q$, which shows that the corresponding sequence of decisions converges.

Finally, we need to show that $\mathcal{S}_{\text{FLS}}(\bar{J}, q) = 1$. The fact that $v(k)$ converges to v implies that $J_q(k) = J_{\pi_v}(k) \le h_v$ for all sufficiently large k. Hence, $\limsup_{k\to\infty} J_q(k) \le h_v < \infty$, which is the desired result.

(Part 2): Let $M < \infty$ and $\epsilon > 0$ be fixed. Our decision rule will be analogous to that above, except that the h sequence is replaced by the fixed and known constant $M + \epsilon$. When $\mathcal{P}$ is infinite, let $\pi_1, \pi_2, \cdots$ be any recurrent sequence of elements of $\mathcal{P}$. The countable memory in this case will keep track of the index into π. When $\mathcal{P}$ is finite assume without loss of generality that $\mathcal{P} = \{0, 1, \cdots, |\mathcal{P}| - 1\}$. The sequence π is then the finite sequence $0, 1, \cdots, |\mathcal{P}| - 1$. In this case, the $|\mathcal{P}|$ states will correspond in a one-to-one fashion to the elements of $\mathcal{P}$, and the state will be incremented modulo $|\mathcal{P}|$.

Our decision rule is defined as follows. Let $v(0) = 1$ be the initial state. For $k \ge 1$, if $J_p(k) \le M + \epsilon$ for some $p \in \mathcal{P}$, then let $s(\bar{J}(k), v(k)) = (v(k) + j) \bmod |\mathcal{P}|$, where $j = \min\{i \ge 0 : J_{\pi_{(v(k)+i) \bmod |\mathcal{P}|}}(k) \le M + \epsilon\}$. If $J_p(k) > M + \epsilon$ for all $p \in \mathcal{P}$, then for simplicity we let the state remain unchanged, i.e., $v(k+1) = v(k)$. Finally, we take $\phi(\bar{J}(k), v(k)) = \pi_{v(k)}$. The proof that this decision rule is successful for $(\mathcal{F}_{\text{FLS}}^{M}, \mathcal{S}_{\text{FLS}}^{M+\epsilon})$ is similar to the proof in Part 1.

(Part 3): Let $M < \infty$ and a positive sequence $\bar{g} = \{g(k)\}$ with $g(k) \to 0$ be fixed. As in Part 1, let π be a recurrent sequence of elements of $\mathcal{P}$. Let $L_1, L_2, \cdots$ be a positive sequence with $L_m \to \infty$. Again, our decision rule will be very similar to Part 1 with $M + L_{v(k)} g(k)$ replacing the $h_{v(k)}$ from Part 1. Let the initial state be $v(0) = 1$, and for $k \ge 1$ let $s(\bar{J}(k), v(k)) = \min\{i \ge v(k) : J_{\pi_i}(k) \le M + L_i g(k)\}$ and let $\phi(\bar{J}(k), v(k)) = \pi_{v(k)}$. The proof that this decision rule is successful for $(\mathcal{F}_{\text{FLS}}^{M,\bar{g}}, \mathcal{S}_{\text{FLS}}^{M,\bar{g}})$ again follows the corresponding proof from Part 1. □

Note that Part 2 has a smaller class of admissible sequences but a stronger performance criterion than Part 1. The resulting decision rule is less complex in that a memory of only size $|\mathcal{P}|$ is used as opposed to a countable memory. On the other hand, Part 3 again has a smaller class of admissible sequences and a stronger performance criterion than Part 2, but the decision rule for Part 3 is more complex in that it uses a countable memory and is nonstationary.

The result of Part 3 can actually be extended to the case in which each $p \in \mathcal{P}$ may have a different rate of convergence to its limsup (assuming the limsup of the sequence $J_p(k)$ is finite). In this case, the class of admissible sequences consists of those for which for some $p \in \mathcal{P}$, $J_p(k) \leq M + Lg_p(k)$, where the convergence rate $g_p(k)$ can depend on $p \in \mathcal{P}$. In this case, the success criterion $\mathcal{S}$ can be modified so that if the decision rule converges to $q \in \mathcal{P}$, then we will be guaranteed that the sequence $J_q(k)$ has limsup bounded by M with corresponding convergence rate $g_q(k)$. The decision rule for this case simply uses $M + L_{v(k)} g_{\pi_{v(k)}}(k)$ in place of $M + L_{v(k)} g(k)$ in the memory state transition rule for $s_k(\bar{J}(k), v(k))$.

B. Performance-Based Search

It is possible to construct, in an ad hoc fashion, many successful rules for the problems of Theorem 3.1. Some of these will be more efficient than others. For example, common sense suggests that the indexes to the sequences π and h in the proof of Theorem 3.1 should be updated separately. In this case the state $v(k)$ consists of a pair of integers (n, m) where n is the index into h and m is the index into π. However, good characterizations of efficiency and the degree to which the selections of successful rules can be sample-path dependent are not available at this point.

In the predictor selection problem it is natural to assume that the smaller the value of $J_p(k)$ the better the performance of predictor p as measured over the interval $[0, k]$. Hence at time k it would be natural to regard the minimum value of $J_p(k)$ over $\mathcal{P}$ (if it exists) as the "best" current predictor performance and the minimizing value of p as the "best" predictor. One way to formalize this is to let $\epsilon_k \geq 0$ and $\sigma_k \geq 1$ and define

$$J_*(k) \triangleq \inf_{p \in \mathcal{P}} J_p(k)$$

and

$$G_k(\bar{J}) \triangleq \{p\colon J_p(k) \leq \sigma_k(J_*(k) + \epsilon_k)\}. \tag{5}$$

It is clear that for each k, $J_*(k) \geq 0$, and that under the assumption that $\bar{J} \in \mathcal{F}_{\rm FLS}$, we have $\limsup J_*(k) = J_+ < \infty$. It is also clear that for each k and each pair (σ_k, ϵ_k) with $\epsilon_k > 0$ or $\sigma_k > 1$, the set $G_k(\bar{J})$ is nonempty.

This leads us to consider the following decision problem. Let $\bar{\epsilon} \triangleq \{\epsilon_k, k \geq 0\}$ and $\bar{\sigma} \triangleq \{\sigma_k, k \geq 0\}$ be sequences of positive real numbers with $\sigma_k \geq 1$. Then let $G_k(\bar{J})$ be given by (5) and define

$$G_\infty(\bar{J}) \triangleq \cup_{N \geq 0} \cap_{k \geq N} G_k(\bar{J}).$$

Now consider the family $\mathcal{F}_{\rm FLS}^{\bar{\sigma},\bar{\epsilon}}$ of all $\bar{J}$ for which there is at least one $p \in \mathcal{P}$ with $\limsup_{k \to \infty} \bar{J}_p(k) < \infty$ and with the property that $G_\infty(\bar{J}) \neq \emptyset$. The corresponding success criterion is $\mathcal{S}^{\bar{\sigma},\bar{\epsilon}}(\bar{J}, q) = 1$ iff $q \in G_\infty(\bar{J})$.

It is easy to see how to modify decision rule 2) of Theorem 3.1 to obtain the following corollary.

Corollary 3.2: Let $\mathcal{P}$ be a finite or countable set. Then there exists a decision rule with memory of size $|\mathcal{P}|$ that is successful for $(\mathcal{F}_{\rm FLS}^{\bar{\sigma},\bar{\epsilon}}, \mathcal{S}^{\bar{\sigma},\epsilon})$.

The "hysteresis switching rule" of [9] and [13] can be formulated as a special case of this corollary.

C. Negative Results

In this subsection we show that the memory requirements of the decision rules presented in Theorem 3.1 cannot be simplified. If we regard the size of the state space (i.e., the memory) as a rough measure of complexity, then these results give tight lower bounds for the complexity of successful rules.

The main result is the following.

Theorem 3.3: Let $\mathcal{P}$ be a finite or countable index set with $|\mathcal{P}| \geq 2$.

1) There is no finite memory decision rule that is successful for $(\mathcal{F}_{\rm FLS}, \mathcal{S}_{\rm FLS})$.
2) For any $M < \infty$ and $\epsilon > 0$, there is no decision rule with memory of size less than $|\mathcal{P}|$ that is successful for $(\mathcal{F}_{\rm FLS}^{M}, \mathcal{S}_{\rm FLS}^{M+\epsilon})$.
3) For any $M < \infty$ and positive sequence $\bar{g} = \{g(k)\}$ with $g(k) \to 0$, there is no finite memory decision rule that is successful for $(\mathcal{F}_{\rm FLS}^{M,\bar{g}}, \mathcal{S}_{\rm FLS}^{M,\bar{g}})$.

Before proving the above theorem, we provide a lemma which simplifies the analysis. The lemma shows that by augmenting the memory of a decision rule one can take the output map to be stationary and a function of only the state.

Lemma 3.4: If (s_k, ϕ_k) is any decision rule with state set V and output set $\mathcal{P}$, then there exists a decision rule $(\hat{s}_k, \hat{\phi})$ with state set $\hat{V} = V \times \mathcal{P}$ and

$$\begin{aligned} &\hat{s}_k : \mathbb{R}^{|\mathcal{P}|} \times \hat{V} \to \hat{V} \\ &\hat{\phi} : \hat{V} \to \mathcal{P} \end{aligned} \tag{6}$$

such that the two rules realize the same input–output map.

Proof: The proof is elementary and hence omitted. The modified rule requires the following minor change in notation: $d(k) = \hat{\phi}(\hat{v}(k+1))$. □

We are now ready for the proof of Theorem 3.3.

Proof (Part 1): For simplicity we assume $\mathcal{P} = \{1, 2\}$. The case where $|\mathcal{P}| > 2$ can be handled by simply considering the subclass of sequences in $\mathcal{F}_{\rm FLS}$ for which $J_p(k) \to \infty$ for all but two $p \in \mathcal{P}$.

Suppose there is a decision rule with finite memory that is successful for $(\mathcal{F}_{\rm FLS}, \mathcal{S}_{\rm FLS})$. Then since $\mathcal{P}$ is finite, by Lemma 3.4 there exists a successful decision rule (s_k, ϕ) with finite memory V of the form (6). Now, define a new decision rule as follows. Let the state set be $W = V \times \{0, 1, 2, \cdots\}$, the transition function $S\colon \mathbb{R}^2 \times W \to W$ be $S((a, b), (v, k)) = (s_k((a, b), v), k + 1)$, and the output map $\Phi\colon W \to \{1, 2\}$ be $\Phi((v, k)) = \phi(v)$. Finally, remove from W any state that is not reachable from the initial state $W_0 = (v_0, 0)$. Clearly

(s_k, ϕ) is a successful decision rule iff (S, Φ) is a successful decision rule.

We now analyze the properties of the decision rule (S, Φ). For $i = 1, 2$, define $W_i \triangleq \{w: \Phi(w) = i\} \subseteq W$, and let W_i^* denote the largest invariant subset of W_i under the input $(0, 0)$. So W_1^* is the largest subset of W_1 such that $w \in W_1^*$ implies $S((0,0), w) \in W_1^*$. Similarly for W_2^*. If the constant input $(0,0), (0,0), (0,0), \cdots$ is applied to (S, Φ) then at some finite time the state must enter and thereafter remain in one of W_1 or W_2. So at least one of W_1^* and W_2^* is nonempty. Without loss of generality assume W_1^* is nonempty.

Consider the input sequence $(1,0), (2,0), (3,0), \cdots$ applied at $w \in W_1^*$. Since (S, Φ) is successful, the resultant trajectory must eventually enter and thereafter remain in W_2. Hence for each state $w \in W_1^*$, there is a finite length input sequence of the form $\bar{U} = ((a_1, 0), \cdots, (a_m, 0))$ which causes the state trajectory starting from w to leave W_1^*. For each such input sequence let $C(w, \bar{U})$ denote the maximum value of a_i applied along the initial segment of the trajectory until the first state outside W_1^* is reached. $C(w, \bar{U})$ represents the cost to leave W_1^* using the input sequence $\bar{U}$. Then for each $w \in W_1^*$ define

$$C(w) \triangleq \inf\left\{C(w, \bar{U}): \text{ state leaves } W_1^* \text{ from } w \text{ under input } \bar{U}\right\}.$$

So $C(w)$ is the greatest lower bound on the cost to leave W_1^* starting from state w. Clearly, $0 \leq C(w) < \infty$. For each $w \in W_1^*$, if an input of the form $(u, 0)$ with $u = 0$ or $u < C(w)$ is applied, then the next state $w' = S((u,0), w)$ remains in W_1^* and $C(w) \leq C(w')$.

Fix $0 < \gamma < 1$ and at each $w \in W_1^*$ apply the input $(\gamma C(w), 0)$. If $C(w) > 0$, then $\gamma C(w) < C(w)$ so the state remains in W_1^*. On the other hand, if $C(w) = 0$, then $\gamma C(w) = 0$ and by definition the input $(0,0)$ keeps the state in W_1^*. So under this feedback law the trajectory that starts at $w \in W_1^*$ stays in W_1^* passing through states $w, w^1, w^2, w^3, \cdots$ with $C(w) \leq C(w^1) \leq C(w^2) \leq \cdots$ using the input sequence $(\gamma C(w), 0), (\gamma C(w^1), 0), (\gamma C(w^2), 0), \cdots$. Since the decision along this trajectory is always one and (S, Φ) is successful, we must have $\limsup_{j \to \infty} \gamma C(w^j) < \infty$, and hence, since the $C(w^j)$ sequence is nondecreasing, $K(w) \triangleq \lim_{j \to \infty} C(w^j) < \infty$.

Now consider two distinct states $w_1, w_2 \in W_1^*$ with $K(w_1) \neq K(w_2)$. The two trajectories that start at $w_1 = (v_1, k_1)$ and $w_2 = (v_2, k_2)$ must thus remain distinct. For $i = 1, 2$, let $(v_i, k_i), (v_i^1, k_i + 1), (v_i^2, k_i + 2), \cdots$ denote the state trajectory that starts from w_i. Then for all $k \geq \max\{k_1, k_2\}$, we have $(v_1^{k-k_1}, k) \neq (v_2^{k-k_2}, k)$. Since there are only $|V|$ choices for v_1^j, this means that there can be at most $|V|$ distinct values for $K(w)$. It follows that there exists a finite B such that $C(w) \leq B$ for all $w \in W_1^*$.

A symmetric argument can be applied to W_2^*. If W_2^* is nonempty, then there exists a finite bound B such that for all $w \in W_2^*$, $C(w) \leq B$.

We complete the proof of Part 1 by arguing as follows. Starting from the initial condition of (S, Φ), select the constant input $(0,0)$ until the state enters W_1^*. Then apply an input segment that drives the state out of W_1^*. This requires a first input size of at most B. When the state first leaves W_1^*, switch again to the constant input $(0,0)$. From this point the state must eventually leave W_1 and enter W_2; either it enters only $W_2 - W_2^*$ or it enters W_2^*. In either case, the state has visited W_2 and eventually settles in W_1^* or W_2^*. At this point, we repeat the above. Thus, using an input sequence in which both components have a finite limsup we have produced a trajectory that does not converge on a decision.

(Part 2): It is enough to consider the case when $\mathcal{P}$ is finite. Without loss of generality, let $M = 0$ and $\epsilon = 1/2$.

Suppose there exists a successful decision rule with memory V such that $|V| < |\mathcal{P}|$. For each $k \geq 1$, if we apply the input $\bar{J}(k) = (0, \cdots, 0)$ to all $v \in V$, then the set of decisions obtained is $D(k) = \{\phi_k((0, \cdots, 0), v): v \in V\}$. Since $|V| < |\mathcal{P}|$, for each $k \geq 1$ there exists $p_k \in \mathcal{P}$ with $p_k \notin D(k)$. Then since $\mathcal{P}$ is finite there exists some fixed $p^* \in \mathcal{P}$ such that $p^* \notin D(k)$ for an infinite set of times $K = \{k_1, k_2, \cdots\}$. Without loss of generality suppose $p^* = 1$. By taking a subsequence of the k_i (if necessary), we can select an infinite set of times $K_0 \subset K$ such that $\mathbb{N} - K_0$ is also infinite.

Now, consider the input sequence $\bar{J}$ defined as follows:

$$\bar{J}(k) = \begin{cases} (0, 0, \cdots, 0) & \text{for } k \in K_0 \\ (0, 1, \cdots, 1) & \text{for } k \notin K_0. \end{cases} \tag{7}$$

Since $\mathbb{N} - K_0$ is infinite, we have $\limsup_{k \to \infty} J_q(k) < 1/2 = M + \epsilon$ iff $q = 1 = p^*$. However, by construction of K_0, $d(k) \neq 1$ infinitely often (namely, for $k \in K_0$). This contradicts the assumption that the decision rule is successful.

(Part 3): The proof closely follows the proof of Part 1. The main difference is that instead of keeping the magnitudes of the input sequence bounded, we need to keep the constant multiplying the rate sequence $g(k)$ bounded. This difference manifests itself primarily in the cost of an input sequence, with the rest of the proof being essentially identical to that of Part 1.

Let $\bar{g} = \{g(k)\}$ be a fixed positive rate sequence with $g(k) \to 0$, and without loss of generality, take $M = 0$. Suppose there exists a decision rule with finite memory that is successful for $(\mathcal{F}_{\text{FLS}}^{0,\bar{g}}, S_{\text{FLS}}^{0,\bar{g}})$. Proceed as in the proof of Part 1, but define the cost of an input sequence that drives a state in W_1^* to a state outside of W_1^* as follows. Given $w = (v, k) \in W_1^*$ and an appropriate input sequence $\bar{U} = ((a_1, 0), \cdots, (a_m, 0))$, define $C(w, \bar{U}) = C((v, k), \bar{U})$ as

$$C(w, \bar{U}) = \max_{1 \leq i \leq m} \frac{a_i}{g(k + i - 1)}.$$

Then for each $w \in W_1^*$, define

$$C(w) = \inf\left\{C(w, \bar{U}): \text{ state leaves } W_1^* \text{ from } w \text{ under input } \bar{U}\right\}.$$

As before, for each $w = (v, k) \in W_1^*$ we have $0 \leq C(w) < \infty$, and if an input $(u, 0)$ is applied with $u = 0$ or $u < g(k)C(w)$, then the next state w' is in W_1^* and $C(w) \leq C(w')$.

For any fixed $0 < \gamma < 1$, we select the input at $w = (v, k) \in W_1^*$ to be $(\gamma g(k) C(w), 0)$. Then by following the proof of Part 1 exactly, we construct an input sequence $(u_1(k), u_2(k))$ such that for some finite B we have $u_i(k) \leq B g(k)$ for

$i = 1, 2$. Thus, this input sequence belongs to $\mathcal{F}^{0,\bar{g}}_{\mathrm{FLS}}$, and yet by construction the decision rule fails to converge. □

IV. Controller Selection Policies

In this section we consider the second problem discussed in the introduction: the selection of a suitable controller for the plant Σ from a family of predesigned controllers $\Gamma_p, p \in \mathcal{P}$. As our strongest results are for linear time-invariant MIMO systems, we focus attention to this case. We begin by defining the families of predictors, models, and controllers.

Assume that each predictor O_p, $p \in \mathcal{P}$, is a linear time-invariant system with state-space representation

$$\begin{aligned} O_p\colon w_p(k+1) &= M_p w_p(k) + B_p u(k) + K_p y(k), \\ & \qquad w_p(0) = w_{p0} \\ \hat{y}_p(k) &= C_p w_p(k) \\ e_p(k) &= C_p w_p(k) - y(k). \end{aligned} \tag{8}$$

Similarly, each controller Γ_p is described by

$$\begin{aligned} \Gamma_p\colon z_p(k+1) &= F_p z_p(k) + G_p y(k) + R_p r(k), \quad z_p(0) = z_{p0} \\ u(k) &= H_p z_p(k) \end{aligned} \tag{9}$$

where r is an exogenous reference input in an admissible class.

The predictor O_p and controller Γ_p will be constrained through the following model:

$$\begin{aligned} \Sigma_p\colon x_p(k+1) &= (M_p + K_p C_p) x_p(k) + B_p u_p(k), \\ & \qquad x_p(0) = x_{p0} \\ y_p(k) &= C_p x_p(k). \end{aligned} \tag{10}$$

This is obtained by setting $y(k) = \hat{y}_p(k)$ in (8). If we let $A_p \triangleq M_p + K_p C_p$, then (C_p, A_p) is a detectable pair and O_p is an observer for Σ_p. Equivalently, one could regard the detectable model (10) as given and then derive $M_p = A_p - K_p C_p$ by stabilizing A_p using output injection. Either point of view is valid; the first is more convenient for our purposes here.

The equations for the closed-loop system that result when Γ_q is connected in feedback with Σ_q are

$$\begin{aligned} \begin{pmatrix} z_q \\ x_q \end{pmatrix}(k+1) &= \begin{pmatrix} F_q & G_q C_q \\ B_q H_q & A_q \end{pmatrix} \begin{pmatrix} z_q \\ x_q \end{pmatrix}(k) + \begin{pmatrix} R_q \\ 0 \end{pmatrix} r(k) \\ y^r_{qq}(k) &= (0 \; C_q) \begin{pmatrix} z_q \\ x_q \end{pmatrix}(k). \end{aligned} \tag{11}$$

We let

$$A_{qq} \triangleq \begin{pmatrix} F_q & G_q C_q \\ B_q H_q & A_q \end{pmatrix}. \tag{12}$$

The predictors and controllers are required to satisfy the following basic constraints. For each $p \in \mathcal{P}$:

L1) the matrices M_p, $p \in \mathcal{P}$, have their eigenvalues inside the circle of radius $\sigma < 1$;

L2) A_{pp} is stable, i.e., Γ_p stabilizes Σ_p.

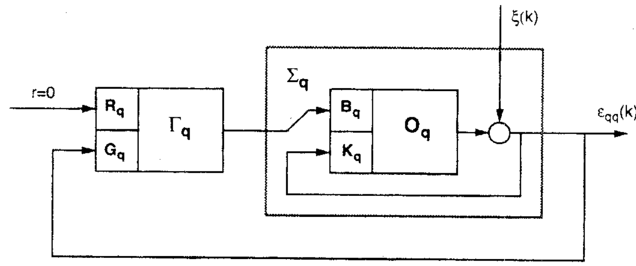

Fig. 1. Closed-loop system (Γ_q, Σ_q) with a disturbance.

We will assume that in addition to satisfying L2), Γ_q has been designed so that the controlled system (Γ_q, Σ_q) gives good tracking performance over the admissible class of reference signals. For example, Γ_q might be designed to stabilize Σ_q and to keep an appropriate weighted induced norm of the mapping from r to the tracking error $y^r_{qq} - r$ of (11) small.

It is also reasonable to expect that Γ_q has been designed so that the closed-loop system (Γ_q, Σ_q) has good disturbance and noise attenuation properties. Of particular interest in this regard is the system

$$\begin{aligned} \begin{pmatrix} z_q \\ x_q \end{pmatrix}(k+1) &= \begin{pmatrix} F_q & G_q C_q \\ B_q H_q & A_q \end{pmatrix} \begin{pmatrix} z_q \\ x_q \end{pmatrix}(k) + \begin{pmatrix} G_q \\ K_q \end{pmatrix} \xi(k) \\ \varepsilon_{qq}(k) &= (0 \; C_q) \begin{pmatrix} z_q \\ x_q \end{pmatrix}(k) + \xi(k). \end{aligned} \tag{13}$$

This is obtained from (11) by setting r to zero and adding a disturbance ξ as an additive term in the output. See Fig. 1. Let β_q denote the induced ∞-norm of the mapping from ξ to ε_{qq}. Since A_{qq} is stable, β_q is finite. For reasons to be seen shortly, a desirable design requirement of Γ_q is that β_q is small.

The plant Σ is assumed to be a linear time-invariant system with state-space representation

$$\begin{aligned} \Sigma\colon x(k+1) &= Ax(k) + Bu(k) + Dv(k), \quad x(0) = x_0 \\ y(k) &= Cx(k) + n(k). \end{aligned} \tag{14}$$

Here v and n are exogenous disturbance and sensor noise signals, respectively.

We assume that the plant satisfies:

P1) for some $p^* \in \mathcal{P}$, (Γ_{p^*}, Σ) is stable.

A consequence of P1) is that (C, A) is detectable, i.e., $A - KC$ is stable for suitable K, and (A, B) is stabilizable.

A. Good Prediction Implies Good Control

We now relate the performance of Γ_q when connected in feedback with the plant Σ to the concurrent performance of predictor O_q. The framework of this result is quite general and includes unmodeled dynamics, disturbances, and sensor noise. The result applies without change to continuous-time systems and easily extends to time-varying linear systems.

Proposition 4.1: Let (14), driven by admissible disturbances and noise signals v and n, be connected in feedback with the control system (9) driven by an admissible reference signal r (see Fig. 2). Assume that P1), L1), and L2) are

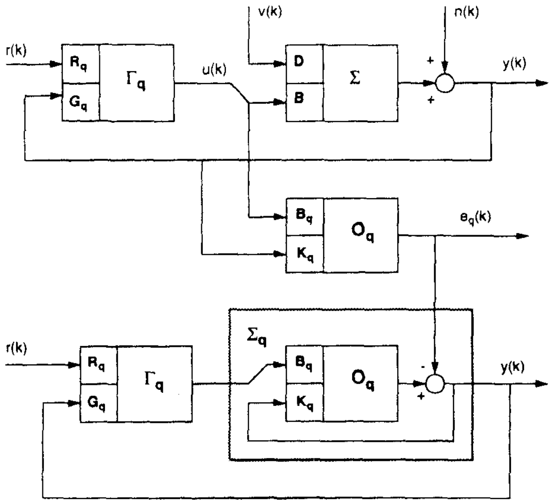

Fig. 2. The systems considered in Proposition 4.1.

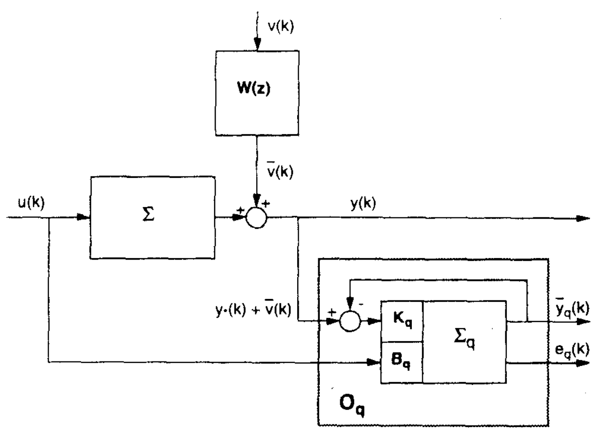

Fig. 3. The setup for additive output disturbances with the details of the predictor O_q.

satisfied. Then

$$y(k) = y_{qq}^r(k) + \varepsilon_{qq}(k), \quad k \geq 0$$

where y_{qq}^r is an output sequence of the disturbance and noise-free stable closed-loop system (11), ε_{qq} is the output of the stable system (13) with $\xi(k) = -e_q(k)$, y is the output of Σ, and e_q is the prediction error sequence of O_q along the trajectory.

In particular, we have the following.

1) If $e_q \to 0$ exponentially, then $\varepsilon_{qq} \to 0$ exponentially. If, in addition, we have $v = n = r \equiv 0$, then $x(k), z_q(k), w_q(k) \to 0$ exponentially, and if r, v, n are bounded, then $x(k), z_q(k), w_q(k)$ are bounded.
2) If $\|e_q\|_{\bar{p}} < \infty$ for some $1 \leq \bar{p} \leq \infty$, then $\|\varepsilon_{qq}\|_{\bar{p}} < \infty$ and

$$\limsup_{k\to\infty} \|\varepsilon_{qq}(k)\| \leq \beta_q \limsup_{k\to\infty} \|e_q(k)\|.$$

In addition, if r, v, n are bounded, then x, z_q, and w_q are bounded.

Proof: The system of interest is

$$\begin{pmatrix} x \\ z_q \\ w_q \end{pmatrix}(k+1) = \begin{pmatrix} A & BH_q & 0 \\ G_qC & F_q & 0 \\ K_qC & B_qH_q & M_q \end{pmatrix} \begin{pmatrix} x \\ z_q \\ w_q \end{pmatrix}(k) + \begin{pmatrix} 0 \\ R_q \\ 0 \end{pmatrix} r(k) + \begin{pmatrix} D \\ 0 \\ 0 \end{pmatrix} v(k) + \begin{pmatrix} 0 \\ G_q \\ K_q \end{pmatrix} n(k)$$

$$y(k) = (C \quad 0 \quad 0) \begin{pmatrix} x \\ z_q \\ w_q \end{pmatrix}(k) + n(k)$$

$$e_q(k) = (-C \quad 0 \quad C_q) \begin{pmatrix} x \\ z_q \\ w_q \end{pmatrix}(k) - n(k) \tag{15}$$

which we will write more compactly as

$$\tilde{x}(k+1) = \tilde{A}\tilde{x}(k) + \tilde{R}r(k) + \tilde{D}v(k) + \tilde{N}n(k) \tag{16}$$

$$y(k) = \tilde{C}\tilde{x}(k) + n(k) \tag{17}$$

$$e_q(k) = \tilde{E}\tilde{x}(k) - n(k) \tag{18}$$

with the obvious definitions of $\tilde{A}, \tilde{R}, \tilde{D}, \tilde{N}, \tilde{C}$, and $\tilde{E}$.

Let

$$\tilde{A}_2 \triangleq \begin{pmatrix} A - KC & BH_q & KC_q \\ 0 & F_q & G_qC_q \\ 0 & B_qH_q & A_q \end{pmatrix} \text{ and } \tilde{K} \triangleq \begin{pmatrix} K \\ G_q \\ K_q \end{pmatrix}. \tag{19}$$

From P1) and L2) it follows that by suitable choice of K we can ensure that $\tilde{A}_2$ is stable, and it is easy to check that $\tilde{A} + \tilde{K}\tilde{E} = \tilde{A}_2$, so $(\tilde{E}, \tilde{A})$ is detectable. Now from (16) and (18)

$$\begin{aligned} \tilde{x}(k+1) &= (\tilde{A} + \tilde{K}\tilde{E})\tilde{x}(k) - \tilde{K}\tilde{E}\tilde{x}(k) + \tilde{R}r(k) \\ &\quad + \tilde{D}v(k) + \tilde{N}n(k) \\ &= \tilde{A}_2\tilde{x}(k) - \tilde{K}e_q(k) + \tilde{R}r(k) + \tilde{D}v(k) \\ &\quad + (\tilde{N} - \tilde{K})n(k). \end{aligned} \tag{20}$$

Let $\tilde{C}_2 \triangleq (\tilde{C}+\tilde{E})$ and $\tilde{x}_1(k)$ and $\varepsilon_{qq}(k)$ denote the solutions to the equations

$$\begin{aligned} \tilde{x}_1(k+1) &= \tilde{A}_2\tilde{x}_1(k) - \tilde{K}e_q(k), \quad \tilde{x}_1(0) = 0 \\ \varepsilon_{qq}(k) &= \tilde{C}_2\tilde{x}_1(k) - e_q(k). \end{aligned} \tag{21}$$

By writing these equations out in detail, one sees that ε_{qq} is the output of (13) with $\xi = -e_q$.

Similarly, let $\tilde{x}_2(k)$ and $y_{qq}^r(k)$ be the solutions of

$$\begin{aligned} \tilde{x}_2(k+1) &= \tilde{A}_2\tilde{x}_2(k) + \tilde{R}r(k) + \tilde{D}v(k) + (\tilde{N} - \tilde{K})n(k), \\ &\qquad \tilde{x}_2(0) = \tilde{x}(0) \end{aligned} \tag{22}$$

$$y_{qq}^r(k) = \tilde{C}_2\tilde{x}_2(k).$$

It is easily checked by writing these equations out in detail that y_{qq}^r is a solution of (11).

Now from (20) it is clear that $\tilde{x}(k) = \tilde{x}_1(k) + \tilde{x}_2(k)$. Then simple algebra together with (17), (18), (22), and (21) yields $y(k) = y_{qq}^r(k) + \varepsilon_{qq}(k)$.

If $e_q \to 0$ exponentially, then by (21) and the stability of $\tilde{A}_2$, $\tilde{x}_1 \to 0$ and $\varepsilon_{qq} \to 0$ exponentially. If $v = n = r \equiv 0$, then the stability of $\tilde{A}_2$ implies that $\tilde{x}_2 \to 0$ exponentially. Hence $x \to 0$, $w_q \to 0$, and $z_q \to 0$ exponentially. The case when r, v, n are bounded is similar.

If $\|e_q\|_{\bar{p}} < \infty$, then (21) and the stability of $\tilde{A}_2$ imply that $\|\tilde{x}_1\|_{\bar{p}} < \infty$ and $\|\varepsilon_{qq}\|_{\bar{p}} < \infty$. Thus $\tilde{x}_1$ and ε_{qq} are bounded sequences. The stated inequality between the limit suprema of these sequences follows from the definition of β_q. If r, v, n are bounded, then (22) and the stability of $\tilde{A}_2$ imply that $\tilde{x}_2$ is bounded. Hence x, w_q, z_q are all bounded. □

B. Switched Controllers

We now apply the results of Section III and Proposition 4.1 to prove a result for prediction error-based controller switching policies. We first define and comment on several special situations of interest.

Exact Matching: When $\Sigma = \Sigma_{p^*}$ for some unknown $p^* \in \mathcal{P}$, we follow [7] and say we have exact matching. Ideally we would like to have exact matching with $v = n \equiv 0$. However, this is clearly a mathematical idealization.

Output Disturbances: If the mapping from the disturbance v to the output y of Σ is stable, then the effect of the disturbance can be modeled as an additive term $\hat{v}$ at the output of the plant. A commonly used model in this situation is that $\hat{v}$ is the output of a known stable LTI filter $W(z)$ with the input v satisfying $\|v\|_\infty \leq 1$. For simplicity we lump n with $\hat{v}$ (or equivalently assume $n \equiv 0$). To assess the effect of this disturbance on the prediction errors let $G_p(z) = C_p(zI - M_p)^{-1}K_p$ and α_p denote the induced ∞-norm of the stable mapping defined by the transfer matrix $G_p(z)W(z)$ from v to e_p. Since the predictors and $W(z)$ are given, we assume that the α_p are known. The setup for additive output disturbances is shown in Fig 3.

SISO Set Point Regulation: In this case assume that all systems are single-input/single-output (SISO), the disturbances can be modeled as an additive term at the output, the reference signals and disturbance signals are constants, and $n \equiv 0$. To ensure asymptotic set-point regulation we assume that the plant has a pole at one and no zeros at one. If necessary, this can be achieved by the inclusion of a summer in the forward path. Since this structure is known, we also include a summer in the observer dynamics by ensuring that each model A_p has an eigenvalue of one and that the transfer function of the triple (C_p, A_p, K_p) does not have a zero at one. If in addition, $\Sigma = \Sigma_{p^*}$ for some $p^* \in \mathcal{P}$, then we have an exact matching SISO set point problem. This is the equivalent tracking problem to that considered by Morse in [7].

We are now ready to state the main result of this section.

Theorem 4.2: Let Σ_p, Γ_p and O_p, $p \in \mathcal{P}$ be linear systems specified by (8)–(10), where r, v, n are elements of the admissible classes of exogenous reference, disturbance, and noise signals, respectively. Assume that for each $p \in \mathcal{P}$ conditions L1) and L2) are satisfied. Let Σ be an unknown plant of the form (14) satisfying P1) and with unknown initial state $x(0)$. Then we have the following.

1) *Ideal case:* There exists a convergent controller selection rule such that if $\Sigma = \Sigma_{p^*}$ for some $p^* \in \mathcal{P}$, and $v = n \equiv 0$, then for some $q \in \mathcal{P}$ (depending on the initial condition and r), $y(k) - y^r_{qq}(k) \to 0$ exponentially. In addition, if $r \equiv 0$, then $x, w_q, z_q \to 0$ exponentially, and if r is bounded, then x, w_q, and z_q are all bounded.
2) *Exact matching with output disturbances:* For each $\epsilon > 0$ and $\rho \geq 1$ there exists a convergent controller selection rule such that if $\Sigma = \Sigma_{p^*}$ for some $p^* \in \mathcal{P}$, r is bounded, and the assumptions of output disturbances are satisfied, then for some $q \in \mathcal{P}$ (depending on the initial condition, r, and v) $\limsup_{k\to\infty} \|y(k) - y^r_{qq}(k)\| \leq \beta_q(\rho\alpha_q + \epsilon) < \infty$, and x, w_q, and z_q are bounded.
3) *Exact matching SISO set-point problem:* There exists a convergent controller selection rule such that if $\Sigma = \Sigma_{p^*}$ for some $p^* \in \mathcal{P}$, and the assumptions of the SISO set point problem are satisfied, then for some $q \in \mathcal{P}$, $\lim_{k\to\infty} d(k) = q$, $y(k) - r \to 0$ exponentially, and the sequences x, w_q, and z_q are all bounded. In addition, if $v = r \equiv 0$, then $x, w_q, z_q \to 0$ exponentially.

Proof (Part 1): Fix $0 < \lambda < 1$, and let the performance measure be $J^k(e(0), \cdots, e(k)) = \sum_{j=0}^k \lambda^{k-j}\|e(j)\|^2$. By assumption L1), $\|e_{p^*}(k)\| \leq L_{p^*}\sigma^k$ for some positive constant L_{p^*}. Let $h(k) \triangleq \sum_{j=0}^k \lambda^{k-j}\sigma^{2j} = (\lambda^{k+1} - \sigma^{2(k+1)})/(\lambda - \sigma^2)$, where for convenience we assume $\lambda \neq \sigma^2$. So $h(k) \geq 0$ and $h(k) \to 0$ exponentially.

Now

$$J_{p^*}(k) = \sum_{j=0}^k \lambda^{k-j}\|e_{p^*}(j)\|^2 \leq L^2_{p^*} \sum_{j=0}^k \lambda^{k-j}\sigma^{2j} \leq L^2_{p^*} h(k).$$

Thus $J_{p^*}(k)$ converges to zero at rate $h(k)$. Hence we can apply the decision rule of Theorem 3.1 Part 3 with $M = 0$ and rate function h. This rule will converge to $q \in \mathcal{P}$ with $J_q(k) \leq Lh(k)$ for some $L > 0$. It follows that $J_q \to 0$ exponentially, and hence by the definition of $J_q(k)$, that $e_q \to 0$ exponentially. Thus for each joint initial state $X_0 = \{x(0), w_{p0}, z_{p0}, p \in \mathcal{P}\}$ and input r, there is a finite time $N = N(X_0, r)$ and a parameter $q = q(X_0, r) \in \mathcal{P}$ such that for $k \geq N$, $d(k) = q$ and $e_q \to 0$ exponentially. Then by Proposition 4.1, for $k \geq N$, $y(k) - y^{qq}(k) = \varepsilon_{qq}(k)$, where y^{qq} is an output trajectory of the system (Γ_q, Σ_q) driven by r and $\varepsilon_{qq} \to 0$ exponentially. In addition, if $r \equiv 0$, then $x, w_q, z_q \to 0$ exponentially; and if r is bounded, x, w_q, z_q are bounded.

(Part 2): The assumption of exact matching and output disturbances implies $\limsup_{k\to\infty} \|e_{p^*}(k)\| \leq \alpha_{p^*} < \infty$. Fix $0 < \lambda < 1$ and let $J_p(k) = \sum_{j=0}^k \lambda^{k-j}\|e_p(j)\|$. Then $\limsup_{k\to\infty} J_{p^*}(k) \leq \alpha_{p^*}/(1 - \lambda)$. Set $\rho = 1/(1 - \lambda)$. Then replacing M by $\rho\alpha_{d(k)}$ in the proof of Part 2 of Theorem 3.1 yields a decision rule that will converge to $q \in \mathcal{P}$ with $\limsup_{k\to\infty} J_q(k) \leq \rho\alpha_q + \epsilon$. It follows that $\limsup_{k\to\infty} \|e_q(k)\| \leq \rho\alpha_q + \epsilon$. After convergence, we are using controller q. Hence by Proposition 4.1, $\limsup_{k\to\infty} \|y(k) - y^r_{qq}(k)\| \leq \beta_q(\rho\alpha_q + \epsilon) < \infty$. Since the closed-loop system is stable and driven by bounded inputs, x, w_q, and z_q are bounded.

(Part 3): For $k \geq 0$ the inputs to the predictors are $u(k)$ and $y(k) = y_{p^*}(k) + \hat{v}(k)$ where u is the plant input, y_{p^*} is the output of Σ_{p^*}, and $\hat{v}$ is the constant additive output disturbance. Hence by linearity $e_{p^*}(k) = e^*(k) + \bar{e}(k)$, where

e^* is the response of the predictor to the input pair (u, y_{p^*}) and the predictor initial state, and $\bar{e}$ is the zero state response to the input pair $(0, \hat{v})$. Clearly, $e^*(k) \to 0$ exponentially.

Since $\hat{v}$ is a constant, $\bar{e}$ is just the response of predictor p^* to a step of size $\hat{v}$ applied to the y input with the u input held at zero. By the Final Value theorem

$$\lim_{k\to\infty} \bar{e}(k) = \lim_{z\to 1} \left(C_{p^*}(zI - M_{p^*})^{-1}K_{p^*} - 1\right)\hat{v}.$$

Now the transfer function $C_{p^*}(zI - M_{p^*})^{-1}K_{p^*}$ is formed by a unity gain negative feedback loop about the model Σ_{p^*}. The relevant model transfer function is $G(z) = C_{p^*}(zI - A_{p^*})^{-1}K_{p^*}$. Hence

$$C_{p^*}(zI - M_{p^*})^{-1}K_{p^*} - 1 = \frac{-1}{1+G(z)}.$$

By assumption, $G(z)$ has a pole at one and no zero at one. Thus

$$\lim_{k\to\infty} \bar{e}(k) = \lim_{z\to 1} \frac{-1}{1+G(z)}\hat{v} = 0.$$

Thus $\bar{e}(k)$ converges to zero exponentially. The rate of convergence is determined by the poles of the predictor and hence is no slower than the rate of predictor convergence.

Fix $0 < \lambda < 1$ and let the performance measure be $J^k(e(0), \cdots, e(k)) = \sum_{j=0}^{k} \lambda^{k-j}|e(j)|^2$. We can now apply the same rule and reasoning as in Part 1. □

Part 1 of Theorem 4.2 is included to indicate what can be achieved in an idealized case. The output of the supervised control system is eventually the sum of an output trajectory of the closed-loop system (Γ_q, Σ_q), for some $q \in \mathcal{P}$, and a term that converges exponentially to zero. Hence, if the controlled systems (Γ_p, Σ_p) all do an adequate job of asymptotically tracking r, then under ideal assumptions so will the supervised control system for Σ. A similar statement can be made when the signal ε_{qq} is bounded except that to be useful the bound must be small compared to the magnitude of the signal r. This can be assured even if e_q does not converge to zero, but an asymptotic bound on e_q is possible.

Part 2 indicates that under the common additive output disturbances model it is still possible to achieve convergence. If the disturbances are small, a reasonable assumption in many cases, the quality of the resultant tracking will still be good. Instead of using a worse-case disturbance model, we could also have used a finite power model. Part 3 is a discrete-time version of the problem considered in [7]. The new contribution is that we show that set-point regulation is possible even when controller selection is required to converge. It is possible to relax the rule used in the proof of Part 3 so that in the limit $J_q(k) < \epsilon$. This simplifies the rule but sacrifices tracking performance.

C. Continuous-Time Systems

We end this section with a brief discussion showing how our results can also be applied to continuous-time LTI systems.

Suppose we have a set of continuous-time LTI predictors O_p, $p \in \mathcal{P}$ and corresponding controllers Γ_p with state-space representations analogous to those in (8) and (9), respectively. We also assume that we have a continuous-time performance measure that is keeping track of how well predictor p is performing. That is, for each p we have a real-valued performance measure $J_p(t)$ that depends on $e_p(\tau)|_{0\le\tau\le t}$. For example, in the case of SISO systems we might fix $\lambda \in \mathbb{R}$ and take $J_p(t) = \int_0^t e^{-\lambda(t-\tau)} e_p^2(\tau)\, d\tau$.

Now, let $t_1 < t_2 < \cdots$ be a sequence of times with $t_n \to \infty$. The t_n denote the potential switching times. That is, at time t_n, we will apply an appropriate decision rule based on the performance indexes sampled at time t_n: $\bar{J}(t_n) = \{J_p(t_n)\}_{p\in\mathcal{P}}$. Based on the output of the decision rule, the corresponding controller will be applied during the interval $t_n < t \le t_{n+1}$. Since the controller can be switched only at a discrete set of times, convergence of a switching policy, etc., can be reduced to a discrete-time analysis. Then, by suitable choices for the performance functions we can apply our results from Section III. Of course, to conclude properties of the resulting continuous-time closed-loop system like those in Theorem 4.2, assumptions analogous to those in Theorem 4.2 must be made.

V. Conclusion

We have analyzed in a general framework the existence and memory requirements of discrete decision rules based on real valued data sequences. We have shown that it is possible to combine the two extremes of performance-based search and prerouted search to obtain some of the desirable characteristics of both methods. The rules contain a mechanism to generate a potential search through the candidate predictors with the revisitation property. However, this is combined with the available performance measure so that only those candidates which are classified as exhibiting "acceptable" performance at time k are considered.

Our two main results on decision rules proved the existence of successful decision rules for three general problems and gave tight lower bounds for the memory requirements of successful rules. In view of the very mild assumptions placed on the signals of interest, these rules should be robust to data perturbation. The rules were applied to the problems of on-line model selection and on-line controller selection. The results are complementary to those of Morse [7]. The added requirement of convergence brings some benefits in terms of controller design but can result in more complex rules and more conservative performance. The example applications did not incorporate all features desirable for a real implementation but nevertheless may provide a starting point for applications. The need to run simultaneous predictors and controllers, for example, is clearly an important implementation issue. Under suitable assumptions, e.g., SISO systems of bounded degree, it is possible to use standard methods to construct a common state realization for the family predictors and controllers; see, e.g., [7].

Open problems that remain to be addressed include:

- a good characterization of the efficiency of convergent selection rules;

- a characterization of the sample path dependence of the selection sequence, e.g., when is a built-in prerouted search path necessary?
- a treatment of more general classes of plant uncertainty.

ACKNOWLEDGMENT

The second author would like to acknowledge discussions with S. Morse which yielded useful insight on prior work. The authors' interest in this area was fostered by a Fall 1994 reading group at Princeton. The authors would like to acknowledge the presentations and comments of reading group members J. Hochermann, C. Horn, and S. E. Posner.

REFERENCES

[1] K. Ciliz and K. Narendra, "Multipole model based adaptive control of robotic manipulators," in *Proc. 33rd IEEE Conf. Decision Contr.*, Lake Buena Vista, FL, Dec. 1994, pp. 1305–1310.
[2] S. J. Cusumano and K. Poolla, "Adaptive control of uncertain systems: A new approach," in *Proc. 1988 Amer. Contr. Conf.*, Atlanta, GA, June 1988, pp. 355–359.
[3] M. Fu and B. R. Barmish, "Adaptive stabilization of linear systems via switching controls," *IEEE Trans. Automat. Contr.*, vol. 31, no. 12, pp. 1079–1103, Dec. 1986.
[4] A. Ilchmann, "Non-identifier-based adaptive control of dynamical systems," *IMA J. Math. Contr. Information*, vol. 39, pp. 321–366, 1991.
[5] B. Mårtensson, "The order of any stabilizing regulator is sufficient *a priori* information for adaptive stabilization," *Syst. Contr. Lett.*, vol. 6, no. 2, pp. 87–91, July 1985.
[6] D. E. Miller, "Adaptive stabilization using a nonlinear time-varying controller," *IEEE Trans. Automat. Contr.*, vol. 39, no. 7, pp. 1347–1359, July 1994.
[7] A. S. Morse, "Supervisory control of families of linear set-point controllers—Part 1: Exact matching," preprint, Mar. 1993.
[8] ——, "Supervisory control of families of linear set-point controllers—Part 2: Robustness," preprint, Nov. 1994.
[9] A. S. Morse, D. Q. Mayne, and G. C. Goodwin, "Applications of hysteresis switching in parameter adaptive control," *IEEE Trans. Automat. Contr.*, vol. 37, no. 9, pp. 1343–1354, Sept. 1992.
[10] K. Narendra and J. Balakrishnan, "Improving transient response of adaptive control systems using multiple models and switching," in *Proc. 32rd IEEE Conf. Decision Contr.*, San Antonio, TX, Dec. 1993.
[11] ——, "Intelligent control using fixed and adaptive models," in *Proc. 33rd IEEE Conf. Decision Contr.*, Lake Buena Vista, FL, Dec. 1994, pp. 1680–1685.
[12] R. D. Nussbaum, "Some remarks on a conjecture in parameter adaptive control," *Syst. Contr. Lett.*, vol. 3, no. 5, pp. 243–246, Nov. 1983.
[13] S. R. Weller and G. C. Goodwin, "Hysteresis switching adaptive control of linear multivariable systems," *IEEE Trans. Automat. Contr.*, vol. 39, no. 7, pp. 1360–1375, July 1994.

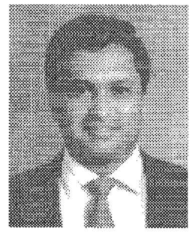

Sanjeev R. Kulkarni (M'91–SM'96) was born in Bombay, India, in 1963. He received the B.S. degree in mathematics, the B.S. degree in electrical engineering, and the M.S. degree in mathematics from Clarkson University, Potsdam, NY, in 1983, 1984, and 1985, respectively. He also received the M.S. degree from Stanford University, Stanford, CA, in 1985 and the Ph.D. from M.I.T., Cambridge, MA, in 1991, both in electrical engineering.

From 1985 to 1991 he was a Member of the Technical Staff at the M.I.T. Lincoln Laboratory, and in the spring of 1986 he was a Part-Time Faculty Member at the University of Massachusetts, Boston. He joined the faculty of Princeton University, Princeton, NJ, in 1991, where he is currently an Assistant Professor of Electrical Engineering. During 1996 he was a Visiting Research Fellow at the Australian National University. His research interests include statistical pattern recognition, machine learning, systems theory, information theory, and image processing.

Dr. Kulkarni received an ARO Young Investigator Award in 1992 and an NSF Young Investigator Award in 1994.

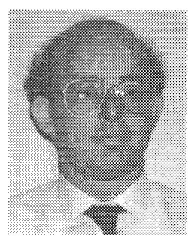

Peter J. Ramadge (S'79–M'82–SM'92–F'96) received the Bachelor of Science and Bachelor of Engineering (electrical) Hons. Class I degrees from the University of Newcastle, Australia, in 1976 and 1978, respectively, the Master of Engineering degree from the University of Newcastle in 1980, and the Ph.D. degree in electrical engineering from the University of Toronto, Canada, in 1983.

He joined Princeton University, Princeton, NJ, in September 1984, where he is currently a Professor of Electrical Engineering. His current research interests are in the theoretical aspects of computer science and control theory with emphasis on applications of computers in signal processing, learning, and control.

Dr. Ramadge is a recipient of the University Medal and the Convocation Medal from the University of Newcastle, Australia. He is a member of SIAM.