Control Using Logic-Based Switching

A. S. Morse[*]

Yale University, New Haven, CT 06520 USA

1 Introduction

Between the well-studied areas of discontinuous control [1], [2] on the one hand and sampled data control [3] on the other lies the largely unexplored area of logic-based switching control systems. By a logic-based switching controller is meant a controller whose subsystems include not only familiar dynamical components {integrators, summers, gains, etc.} but logic-driven elements as well {e.g., [4]}. More often than not the predominately logical component within such a system is called a supervisor [5], a mode changer [6], a gain scheduler, or something similar. Within the last decade a number of analytical studies of such systems have emerged, mainly in the area of self-adjusting control [7, 8, 9, 10, 11, 12, 13, 14, 15, 16]. These studies and others have shown that much can be gained by using logic-based switching together with more familiar techniques in the synthesis of feedback controls. The overall models of systems composed of such logics together with the processes they are intended to control are concrete examples of *hybrid dynamical systems* [17, 18, 19]. The aim of this paper is to give a brief tutorial review of four different classes of hybrid systems of this type - each consists of a continuous-time process to be controlled, a parameterized family of candidate controllers, and an event driven switching logic. Three of the logics, called *prerouted switching*, *hysteresis switching* and *dwell-time switching* respectively, are simple strategies capable of determining in real time which candidate controller should be put in feedback with a process in order to achieve desired closed-loop performance. The fourth, called *cyclic switching*, has been devised to solve the long-standing stabilizability problem which arises in the synthesis of identifier-based adaptive controllers because of the existence of points in parameter space where the estimated model upon which certainty equivalence synthesis is based, loses stabilizability.

In section 2, we discuss several basic issues common to supervised control systems of all types. In most cases of interest, the job of a supervisor is to

[*] The author's research was supported by NSF Grant n. ECS-9206021, AFOSR Grant n. F49620-94-1-0181, and ARO Grant n. DAAH04-95-1-0114

orchestrate the switching of a sequence of candidate controllers into feedback or series with a process, so as to achieve some prescribed goal. No matter what the goal might be, the underlying architecture of the supervised control system is pretty much the same - at least in concept. In section 2 we make the point that such "multi-controller" architectures can usually be implemented most efficiently as "state-shared" parameter-dependent controllers.

In section 3 we briefly discuss two examples of logic-based switching controllers which arise in nonadaptive applications. The first is an 'intelligent' control strategy devised to maximize system performance while at the same guaranteeing that hard-bound saturation constraints are satisfied [20]. The second is a simple, time-invariant, chatter-free, switching logic with one state variable, which is capable of asymptotically stabilizing a particular bilinear system of current interest called the "nonholonomic integrator" [21].

The aim of §4 is to explain the concepts of prerouted, hysteresis, dwell-time, and cyclic switching. Although each of these strategies is applicable to a variety of systems [11, 12, 13, 14, 22, 16, 23], for the sake of uniformity all are reviewed within the context of a single prototype problem - the set-point control of a siso linear system with large-scale parametric uncertainty [24]. The problem is formulated in §4.1.

The concept of "prerouted switching" is closely allied with the idea of a "nonestimator based supervisor"; both topics are discussed in §4.2.

Section 4.3 focuses on the idea of an estimator-based supervisor. It is within this context that the concepts of hysteresis switching and dwell-time switching are explained. The idea of cyclic switching is then reviewed in section 4.4

The logics discussed in §4 are conceptually straight forward. What's interesting about them theoretically is the set of technical questions they generate. Most of the questions have to do with dynamical systems in which switching is non-terminating, non-chattering and asynchronous. Many unanswered questions exist. Some are briefly discussed in §5.

2 Multi-Controllers

Perhaps the simplest architecture one can think of for a feedback system employing a family of controllers is that depicted in Figure 1. That is, the measured output y of a process to be controlled drives a bank of controllers, each controller generating a candidate {possibly vector-valued} feedback signal u_i. The control signal applied to the process at each instant of time is then

$$u \stackrel{\Delta}{=} u_\eta$$

where $\eta : [0, \infty) \to \mathcal{I}$ is a piecewise-constant switching signal taking values in the family's index set $\mathcal{I}$. The generation of such a switching signal is typically carried out by some type of hybrid dynamical system which depending on the situation might be called a tuner, a supervisor, a mode-changer, or something similar. In the sequel we shall refer to such architectures informally as *multi-controllers*.

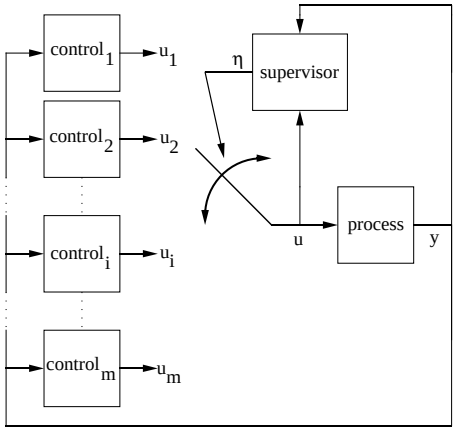

Fig. 1. Multi-Control

Many multi-controller configurations can be implemented using a much simpler architecture than Figure 1 would suggest. The key factor which makes this possible is simply that at any instant of time only *one* of the constituent controller is to be applied to the process. Because of this, at each time t it is only necessary to generate one candidate control signal. Often this means significant simplification can be achieved if all control signals are generated by a single system. In other word, rather than implementing each of the controllers in the family as a separate dynamical system, one can often achieve the same end using a single controller with adjustable parameters. The idea is quite straight forward and is called *state sharing.*

For example suppose that it is desired to implement a finite {or even countable} family of siso linear controllers with reduced transfer functions

$$\kappa_i(s) = \frac{\alpha_i(s)}{\beta_i(s)}, \quad i \in \mathcal{I}$$

where each $\beta_i(s)$ is a monic polynomial. Assuming a fixed upper bound n for the McMillan Degrees of the $\kappa_i(s)$, it is always possible to "cover" this family with a parameter-dependent transfer function $h_q(s)$ whose denominator is of degree n and whose parameter vector q takes values in a linear space of dimension not exceeding $2n+1$. In fact, for any positive integer $\bar{n} \leq 2n+1$, it is always possible to pick a subset $\mathcal{Q} \subset \mathbb{R}^{\bar{n}}$ with the same cardinality as $\mathcal{I}$, and a parameter-dependent transfer function $h_q(s)$ so that for each $i \in \mathcal{I}$ there is a $q_i \in \mathcal{Q}$ such that $\kappa_i(s) = h_{q_i}(s)$ after cancellation of common poles and zeros. Moreover it is always possible to choose $h_q(s)$ in such a way that whenever such pole-zero cancellations occur, they occur at prescribed stable locations.

Having constructed such an $h_q(s)$, the above multi-controller can be implemented as a parameter dependent system $\Sigma_C(\sigma)$ of the form

$$\dot{x_C} = A_\sigma x_C + b_\sigma y \tag{1}$$

$$u_\sigma = f_\sigma x_C + g_\sigma y \tag{2}$$

where $\{A_q, b_q, f_q, g_q\}$ is a n-dimensional realization of $h_q(s)$ and σ is a piecewise constant switching signal taking values in $\mathcal{Q}$. The resulting multi-control system would then appear as in Figure 2.

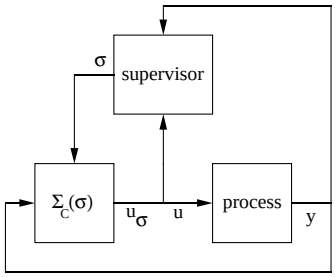

Fig. 2. State-Shared Multi-Controller Implementation

If the supervisor is allowed to re-initialize Σ_C's state at switching times, then this implementation can generate exactly the same feedback control signal as would have been generated had the original architecture been employed.

For multi-controller families consisting of more than just a few controllers, this state-shared implementation is clearly a lot less complicated than a direct implementation of the original multi-controller. Moreover state-sharing frees one from having to be concerned about the boundedness of the out-of-loop control signals which would be present in a direct implementation of the original multi-controller architecture.

There are of course a great many different ways to realize $h_q(s)$. The only essential requirement of any such realization is that it be a "globally detectable, globally stabilizable" system; i.e. for each fixed value of $q \in \mathcal{Q}$, the linear system $\{A_q, b_q, f_q, g_q\}$) should be stabilizable and detectable[2]. For without stabilizability, closed-loop boundedness of u and y cannot be assured and without detectability boundedness of Σ_C's state cannot be assured even if u and y are. One familiar structure which is globally detectable is of the form

$$\left\{ \begin{bmatrix} A & 0 \\ 0 & A \end{bmatrix} + \begin{bmatrix} b \\ 0 \end{bmatrix} f_q,\ \begin{bmatrix} 0 \\ b \end{bmatrix},\ f_q,\ d_q \right\}$$

where (A, b) is a parameter-independent, n-dimensional siso, controllable pair with A stable. Another is $\{A + k_q f, b_q, f, d_q\}$ where (f, A) is an n-dimensional, parameter-independent observable pair. This particular realization is actually observable for all $q \in \mathbb{R}^{\bar{n}}$; moreover in the event that d_q is constant on $\mathcal{Q}$, this

[2] The reader should recognize that any such parameter-dependent system will *always* have points in $\mathcal{Q}$ at which it is not controllable and observable if the transfer functions being realized are not all of the same McMillan Degree.

realization guarantees that there will be a "bumpless" transfer between control signals when σ switches; i.e., $u \stackrel{\Delta}{=} u_\sigma$ is continuous, even at those times at which σ changes values. Of course bumpless transfer can also be achieved with state-reinitialization, whether d_q is constant on $\mathcal{Q}$ or not.

It is fairly clear that the preceding ideas apply to multi-controller families of mimo finite dimensional controllers configured in almost any way imaginable. It is also clear that the number of (fixed-parameter) controllers one might contemplate implementing in a particular multi-controller application need not be finite nor even countable. In other words the complexity of a multi-controller is not so much a function of a number of controllers in a family as it is of the number of algebraically independent gains needed to parameterize the family.

3 Examples

In the sequel are several examples of {nonadaptive} logic-based switching controllers.

3.1 Smart Governors

An important problem of continuing interest is that of developing feedback controllers for linearly modeled processes whose associated inputs and outputs are required to satisfy hard-bound magnitiude constraints. Remarkable advances have recently been made in the development of implementable algorithms for the stabilization of such systems [25, 26]. At the same time there has also been a growing interest in the development of "smart controllers" employing logic aimed not only at maintaining loop stability, but at enhancing system performance as well [27, 28, 29, 20]. One configuration characteristic of this line of research is as follows.

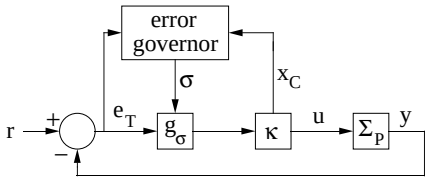

Here Σ_P represents a linear process with an input saturation constraint, g_σ is an adjustable gain, and κ is a linear controller. The idea is to design κ to meet performance specifications in the absence of saturation constraints; this is done for the case $g_\sigma = 1$. The error governor is designed to adjust g's value to give the best performance possible subject to the requirement that the saturation constraints are satisfied. This is accomplished, roughly speaking, by leaving g set at 1 whenever r is 'small' and by reducing g's value when r is 'large' by as much as is required to insure that there is no saturation. The error governor which

accomplishes this is a logical circuit which carries out the required computations in real time. A generalized {discrete-time} version of the preceding with greatly reduced computational requirements has been proposed in [27].

An even more elaborate multi-controller architecture, aimed at a similar problem has been suggested in [20]. The problem addressed is to bring to zero from an admissible start, the state $x \stackrel{\Delta}{=} \{x_P, x_C\}$ of the system $\Sigma(\sigma)$ depicted in the following figure while not violating a set of prespecified state constraints along the way.

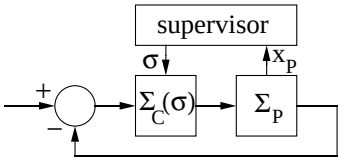

Associated with each fixed control index $q \in \mathcal{Q}$ is a *maximal admissible set* $\mathcal{S}_q$. A state x_0 is in $\mathcal{S}_q$ just in case each point on the closed-loop trajectory of $\Sigma(q)$ emanating from x_0, satisfies the aforementioned state constraints. According to [20], it is possible to use the theory of maximal output admissible sets [30] to design controllers so that $\mathcal{S}_q \subset \mathcal{S}_{q+1}$ and in addition so that controller q achieves better performance than controller $q+1$ when the system is initialized at a state in $\mathcal{S}_q \cap \mathcal{S}_{q+1}$. In [20] it is then explain how to construct a supervisor which successively switches σ to smaller and smaller values to as to achieve better and better performance while satisfying state constraints.

3.2 Nonholonomic Integrators

For more than a decade it has been known that there are nonlinear systems which are locally null controllable but which nevertheless cannot be locally asymptotically stabilized with any smooth, time-invariant controller [31]. A prototypical example of this is the bilinear system

$$\begin{aligned} \dot{x} &= u \\ \dot{y} &= v \\ \dot{z} &= xv - yu \end{aligned}$$

which is sometimes called the "nonholonomic integrator" [32]. Nonholonomic systems such as this have evoked considerable interest in recent years [33]. This has been especially true of the nonholonomic integrator itself. For example, a number of time-varying, periodic controllers have been devised which asymptotically stabilize the above system {cf. [32]}. In addition, by appealing to the theory of sliding modes [1], it has been recently shown that the simple discontinuous control $u = -x + y(\text{sign}(z))$ $v = -y - x(\text{sign}(z))$ will drive x, y, and z

to zero provided one admits generalized solution in the sense of [34]. It turns out to be possible to achieve asymptotic stability without chattering using a time-invariant logic-based switching controller. One strategy which accomplishes this uses a multi-controller of the form

$$\begin{bmatrix} u \\ v \end{bmatrix} = g_\sigma(x, y, z)$$

where

$$g_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \qquad g_2 = \begin{bmatrix} x + yz \\ y - xz \end{bmatrix} \qquad g_3 = \begin{bmatrix} -x + yz \\ -y - xz \end{bmatrix} \qquad g_4 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

and σ is a piece-wise constant switching signal taking values in $\mathcal{I} \stackrel{\Delta}{=} \{1, 2, 3, 4\}$. σ is generated by a supervisor of the form

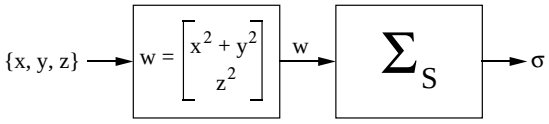

where Σ_S is a switching logic whose input is w and whose state and output are both σ. Σ_S's definition requires one to pick four appropriately structured overlapping regions $\mathcal{R}_q, \quad q \in \mathcal{I}$ which together cover the closed positive quadrant $\Omega \stackrel{\Delta}{=} \{(r_1, r_2) : r_1 \geq 0,\ r_2 \geq 0\}$ in $\mathbb{R}^2$. One possible set of regions is

$$\begin{aligned}
\mathcal{R}_1 &\stackrel{\Delta}{=} \{(r_1, r_2) : r_2 < 2\pi(r_2),\ \ (r_1, r_2) \in \Omega\} \\
\mathcal{R}_2 &\stackrel{\Delta}{=} \{(r_1, r_2) : \pi(r_1) < r_2 < 4\pi(r_1),\ \ (r_1, r_2) \in \Omega\} \\
\mathcal{R}_3 &\stackrel{\Delta}{=} \{(r_1, r_2) : r_2 > 3\pi(r_1),\ \ (r_1, r_2) \in \Omega\} \\
\mathcal{R}_4 &\stackrel{\Delta}{=} \{(0, 0)\}
\end{aligned}$$

where $\pi(r_1) \stackrel{\Delta}{=} (1 - e^{-r_1})$. Σ_S's internal logic is then defined by the computer diagram

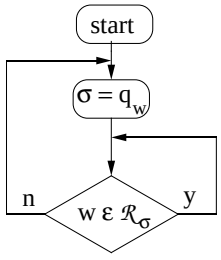

where

$$q_w \stackrel{\Delta}{=} \min_q \{q : w \in \mathcal{R}_q,\ q \in \mathcal{I}\}$$

In interpreting this diagram it is to be understood that σ's value at each of its switching times $\bar{t}$ is its limit from above as $t \downarrow \bar{t}$. Thus if $\bar{t}_i$ and $\bar{t}_{i+1}$ are any two successive switching times, then σ is constant on $[\bar{t}_i, \bar{t}_{i+1})$.

It can be shown that with this switching logic, chattering cannot occur and that x, y and z must tend to zero no matter how they and σ are initialized [21]. It can also be shown that the origin $x = y = z = 0$ is 'Lyapunov stable' in an appropriately defined sense. We refer the reader to [35] for a different application of a switching logic similar to the one we've been discussing.

4 Self-Adjusting Control

The aim of this section is to give a brief tutorial overview of four different classes of logic-based switching control systems - each consists of a continuous-time process to be controlled, a parameterized family of linear controllers, and an event driven switching logic. Three of the logics, called *prerouted switching*, *hysteresis switching* and *dwell-time switching* respectively, are simple strategies capable of determining in real time which controller from a family of candidates should be put in feedback with a process in order to achieve desired closed-loop performance. The fourth, called *cyclic switching*, has been devised to solve the long-standing stabilizability problem which arises in the synthesis of identifier-based adaptive controllers because of the existence of points in parameter space where the estimated model upon which certainty equivalence synthesis is based, loses stabilizability. Although each of these strategies is applicable to a variety of problems, the sake of uniformity all are explained within the context of a single prototype problem - the set-point control of a siso linear system with large-scale parametric uncertainty §4.1. The concept of prerouted switching is closely allied with the idea of a "nonestimator based supervisor"; both topics are discussed in §4.2. Hysteresis switching and dwell-time switching are explained in §4.3 in connection with the concept of an estimator-based supervisor. Cyclic switching is discussed in §4.4

4.1 The Problem

The prototype problem we want to consider is basic: to construct a control system capable of driving to and holding at a prescribed set-point, the output of a process modeled by a dynamical system with large scale parametric uncertainty. Assume the process admits the model of a siso controllable, observable linear system Σ_P with control input u and measured output y. Further assume that Σ_P's transfer function from u to y is a member of a known class of admissible strictly proper transfer functions $\mathcal{C}_P$. In view of the requirements of set-point control, assume that the numerator of each transfer in $\mathcal{C}_P$ is nonzero at $s = 0$.

The specific design goal is to construct a positioning or set-point control system capable of causing y to approach and closely track any constant reference input r. Towards this end we introduce a *tracking error*

$$\mathbf{e_T} \stackrel{\Delta}{=} r - y \tag{3}$$

and an integrating subsystem to generate u; i.e.,

$$\dot{u} = v \tag{4}$$

Here v is a control signal which will be defined in the sequel.

As our concern is mainly with supervisory control, we are going to take as given, a parameterized family of proper, reduced controller transfer functions $\mathcal{K} \stackrel{\Delta}{=} \{\kappa_q : q \in \mathcal{Q}\}$ which has the property that for each transfer function $\tau \in \mathcal{C}_P$, there is at least one controller transfer function $\kappa \in \mathcal{K}$ which internally stabilize feedback interconnection shown in Figure 3.

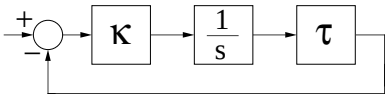

Fig. 3. Feedback Interconnection

The sub-system to be supervised is thus of the form

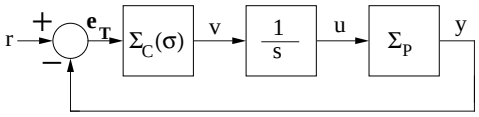

Fig. 4. Supervised Sub-System $\Sigma(\sigma)$

where $\Sigma_C(q)$ is a parameter-dependent, globally detectable/stabilizable realization of κ_q with state x_C. In the sequel we shall describe various types of supervisors capable of generating σ so as to at least achieve set-point regulation {i.e., $\mathbf{e_T} \to 0$} and global boundedness.

4.2 Nonestimator-Based Supervisor

A 'nonestimator-based' supervisor is a hybrid dynamical system whose input is a suitably defined "tuning error" e and whose output is σ.

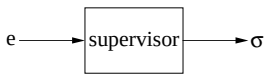

A tuning error is a linear {possibly parameter-dependent} function of the measurable signals in the sub-system $\Sigma(\sigma)$ shown in Figure 4. The key requirements governing the selection of e are as follows.

Tuning Error Requirements:

1. For each fixed $q \in \mathcal{Q}$, $\Sigma(q)$ must be detectable through e.
2. For each constant r and each $q \in \mathcal{Q}$, e must vanish on $\Sigma(q)$'s equilibrium state.

The global detectablilty requirement is fundamental. Its significance has been discussed in a broader context in [36].

One definition of e which satisfies both requirements for the problem under consideration is

$$e \stackrel{\Delta}{=} \begin{bmatrix} \mathbf{e_T} \\ v \end{bmatrix}$$

There are many other acceptable choices as well.

Assume that e has been defined so that the preceding requirements are satisfied. The sub-system depicted in Figure 4 then admits a state space model of the form

$$\begin{aligned} e &= C_\sigma \bar{x} \\ \dot{\bar{x}} &= A_\sigma \bar{x} + b_\sigma r \end{aligned} \tag{5}$$

where $\bar{x}$ is the composite state

$$\bar{x} \stackrel{\Delta}{=} \begin{bmatrix} x_P \\ u \\ x_C \end{bmatrix}$$

and A_q, b_q and C_q are parameter-dependent matrices determined by the definition of e and the coefficient matrices of Σ_P and Σ_C. The position of the integrator in Figure 4 is important [24]: its location guarantees that for each fixed r, the equilibrium state of (5), namely

$$\bar{x}_0 \stackrel{\Delta}{=} -A_q^{-1} b_q r,$$

is independent of $q \in \mathcal{Q}$. Because of this and the assumption that e satisfies Tuning Error Requirement 2, it is possible to write

$$\begin{aligned} e &= C_\sigma x \\ \dot{x} &= A_\sigma x \end{aligned} \tag{6}$$

where $x \stackrel{\Delta}{=} \bar{x} - \bar{x}_0$.

What we want to do is to explain how to construct a supervisor whose output σ causes $x \to 0$ as $t \to \infty$. The *only* properties of (6) which we will exploit are the following:

Properties of $(C_q,\ A_q)$ **:**

1. There exists a parameter value $q^* \in \mathcal{Q}$ for which A_{q^*} is a stability matrix.
2. (C_q, A_q) is detectable for each $q \in \mathcal{Q}$.

The first property is a consequence of the assumption that for each transfer function $\tau \in \mathcal{C}_P$ there is a transfer function $\kappa \in \mathcal{K}$ which stabilizes the system shown in Figure 3. The second property follows from Tuning Error Requirement 1.

Apart from the preceding, nothing is assumed about (6) other than that e can be measured. In particular, neither the parameter-dependent pair (C_q, A_q) nor q^* are presumed to be known. Of course with so little known, one should not expect to come up with a supervisor worthy of actual implementation unless perhaps $\mathcal{Q}$ is a finite set with a small number of elements.

Within the parameter-adaptive framework proposed in [36], a nonestimator-based supervisor would be called a "prerouted" parameter tuner. All parameter tuners, be they prerouted or not, are based on the same underlying strategy which roughly speaking is to keep adjusting σ until e is "small" in some suitably defined sense. Although there are a great many different methods for accomplishing this, in most instances tuning is carried out in one of two fundamentally different ways depending on whether the 'path' σ takes in $\mathcal{Q}$ is 'prerouted' or not. For the prerouted case, tuning is achieved by moving σ through $\mathcal{Q}$ along a prespecified path or route, using on-line {i.e., real-time} data to decide only if and when or how fast to change σ from one value along the path to the next. In contrast, for the non-prerouted case, the path in $\mathcal{Q}$ along which σ is adjusted is not prespecified off-line but instead is determined in real time from the values of various measured signals.

The basic idea of prerouted tuning was devised by Mårtensson with the expressed purpose of delineating the theoretical limits of what might be achieved with any adaptive algorithm [7]. Over the past decade many refinements and modifications of the concept have appeared [8, 9, 10, 37, 38]. Although these modified algorithms differ from each other in many ways, all share certain underlying features in common. In most cases prerouted tuners consist of the cascade connection of two subsystems, one a *scheduling logic* Σ_S and the other a memoryless map $h : \{1, 2, \ldots, \infty\} \to \mathcal{Q}$ called a *routing function*.

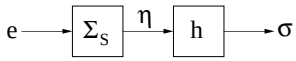

h is invariably required to have the *revisitation property:* That is, for any $q \in \mathcal{Q}$ and any positive integer i there must exist an integer $j \geq i$ at which $h(j) = q$. In other words, h must have the property that the prerouted path $h(1), h(2), \ldots$ *revisits* {i.e., passes through} each point in $\mathcal{Q}$ infinitely often. For this to be possible, $\mathcal{Q}$ must clearly be at least a countable set[3]. Assuming this to be the case, it is always possible to define a routing function with the revisitation property. One way to do this is as follows.

[3] Actually in Martenson's original work $\mathcal{Q}$ is a continuum and the elements of the sequence $h(1), h(2), \ldots$ are only required to get close to {rather than equal} previously visited ones [39].

1. If $\mathcal{Q} = \{q_1, q_2, \ldots, q_m\}$ is a finite set, define h to be the m periodic function whose first m values are $q_1, q_2, \ldots, q_{m-1}$ and q_m respectively.
2. If $\mathcal{Q} = \{q_1, q_2, \ldots, \}$ is not a finite set, define h be the function whose sequence of values $h(1), h(2), \ldots$ are the elements of the sequence $q_1, q_1, q_2, q_1, q_2, q_3, q_1, q_2, q_3, q_4, q_1, \ldots$

There are many possible ways to define Σ_S, depending what one is trying to accomplish. For illustrative purposes, we shall take Σ_S to be a hybrid dynamical system whose input is e and whose output is a piecewise-constant switching signal η taking values in the set of positive integers. Σ_S's state consists of four variables - η, a timing signal τ, a piecewise-continuous 'performance signal' π and a piecewise constant 'sampled performance signal' $\bar{\pi}$. Both π and $\bar{\pi}$ take values in $[0, \infty)$. Timing signal τ takes values in the closed interval $[0, \tau_D]$, where τ_D is a preselected positive number called a *dwell time*. Σ_S's dynamics are defined by the following computer diagram.

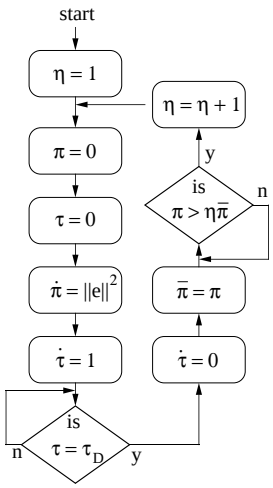

The functioning of Σ_S is as follows. During the first τ_D time units after the algorithm is initiated, τ is increased linearly from 0 to τ_D using a reset integrator and π is increased from 0 according to the rule

$$\dot{\pi} = ||e||^2 \tag{7}$$

Just at the end of this period, τ is reset to zero, the reset integrator is turned off, and $\bar{\pi}$ is set equal to the present value of π. So long as π remains less than or equal to $\sigma\bar{\pi}$, the updating of π continues according to (7). If and when π becomes greater than $\eta\bar{\pi}$, η is incremented by 1, π and $\bar{\pi}$ are reset to zero, and the entire process is repeated. Note that the time between any two successive switchings of η can never be smaller than τ_D. Said differently, η "dwells" at each of its values for at least τ_D time units. Because of this, infinitely fast switching cannot occur

so existence and uniqueness of solutions to the differential equations involved is not an issue.

Analysis: There is a fairly straight forward way to go about analyzing the type of supervisory control system we've just described. The key step is to prove that switching stops in finite time. In particular, for the problem at hand the trick is to show that there must be a finite integer $\bar{\eta}$, depending only on the pair (C_q, A_q) and not on the initial value of x, which η cannot exceed in value. Before we address this claim, let us consider its consequences. What the claim implies is that η can switch at most a finite number of times and therefore that there must be a time $\bar{t}$ beyond which η is constant and $\pi \leq \bar{\eta}\bar{\pi}$. The latter assures that

$$\int_{\bar{t}}^{t} ||e(s)||^2 ds \leq \bar{\eta}\bar{\pi}, \quad t \geq \bar{t}$$

and thus that e has a finite $\mathcal{L}^2[0,\infty)$ norm. Moreover since (C_q, A_q) is detectable on $\mathcal{Q}$ and η is fixed at some value $\bar{q} \in \mathcal{Q}$, $(C_{\bar{q}}, A_{\bar{q}})$ is detectable[4]. As a consequence, for $t \geq \bar{t}$ it is possible to rewrite (6) as

$$\dot{x} = (A_{\bar{q}} + KC_{\bar{q}})x - Ke$$

where K is any matrix which stabilizes $A_{\bar{q}} + KC_{\bar{q}}$. Therefore $x \to 0$ since e has a finite $\mathcal{L}^2[0,\infty)$ norm. In fact, because $\dot{x} = A_{\bar{q}}x$ is a time invariant linear system, x must go to zero as fast at $e^{-\lambda t}$, $-\lambda$ being the largest of the real parts of $A_{\bar{q}}$'s stable eigenvalues. In other words, to prove that $x \to 0$ {and consequently that $\bar{x}$ has a finite limit and that $\mathbf{e_T} \to 0$} it is enough to show that there is an integer $\bar{\eta}$ which η cannot exceed in value.

Here briefly is the idea. Let p be any point in $\mathcal{Q}$ at which A_p is a stability matrix. If $C_p = 0$, let $\bar{\eta}$ be the first positive integer such that $h(\bar{\eta}) = p$. Then either there is an interval $[t_0, t_1)$ of maximal length on which $\eta = \bar{\eta}$ or η never gets as large as $\bar{\eta}$. If the latter is true, then we are done. On the other hand, if the former is true, then $\pi = 0$ on $[t_0, t_1)$ so for such t, $\pi \leq \bar{\eta}\bar{\pi}$. Because of Σ_S's definition, this means that no more switching can occur, that $t_1 = \infty$ and thus that η can grow no larger than $\bar{\eta}$.

Now suppose $C_p \neq 0$. Reduce (C_p, A_p) to an observable pair $(\bar{C}, \bar{A})$ by picking any full rank matrix R whose kernel is the unobservable space of (C_p, A_p) and solving the linear equations $C_p = \bar{C}R$, $RA_p = \bar{A}R$ for $\bar{C}$ and $\bar{A}$ respectively. Note that $\bar{A}$ must be a stability matrix because A_p is.

Let $G(t)$ denote the observability Gramian

$$G(t) \stackrel{\Delta}{=} \int_0^t e^{\bar{A}s}\bar{C}\bar{C}e^{\bar{A}s}ds$$

Note that $G(\infty)$ must exist because of $\bar{A}$'s stability. Moreover, $G(\tau_D)$ must be positive definite because of the observability of $(\bar{C}, \bar{A})$. This implies that

$$\mu \stackrel{\Delta}{=} \sup_x \frac{x'G(\infty)x}{x'G(\tau_D)x} < \infty$$

[4] There is of course no reason to assume that $A_{\bar{q}}$ is a stability matrix.

and that

$$G(\infty) \leq \mu G(\tau_D) \tag{8}$$

Now let η_p be the least integer no smaller than μ for which $h(\eta_p) = p$. Because of the revisitation property, η_p must necessarily exist. In view of (8)

$$G(\infty) \leq \eta_p G(\tau_D) \tag{9}$$

We claim that $\bar{\eta} \stackrel{\Delta}{=} \eta_p$ has the desired property. To prove that this is so, we may as well assume that there is an interval on which $\eta = \bar{\eta}$. For if this were not the case then σ could not exceed $\bar{\eta}$ and we would be done.

Let $[t_0,\ t_1)$ denote the largest interval on which $\eta = \bar{\eta}$. For $t \in [t_0, t_1)$, $\sigma = h(\eta) = p$ and

$$\pi(t) = \int_{t_0}^{t} ||C_p e^{A_p(w-t_0)} x(t_0)||^2 dw \leq \int_0^{\infty} ||\bar{C} e^{\bar{A}s} R x(t_0)||^2 ds = ||\sqrt{G(\infty)} R x(t_0)||^2$$

From this, (9) and the definitions of $\bar{\eta}$ and $\bar{\pi}$ it follows that for $t \in [t_0, t_1)$,

$$\pi(t) \leq \eta_q ||\sqrt{G(\tau_D)} R x(t_0)||^2 = \bar{\eta} \int_0^{\tau_D} ||\bar{C} e^{\bar{A}s} R x(t_0)||^2 ds = \bar{\eta}\bar{\pi}$$

Thus because of Σ_S's definition, no more switching can occur, $t_1 = \infty$ and thus η can grow no larger than $\bar{\eta}$. ■

There are many provably correct versions of the algorithm we've just analyzed [7, 8, 9, 10, 37, 38]. All employ a tuning error satisfying the aforementioned requirements, a performance signal, a routing function and a switching logic similar to the one we've described. Usually τ_D is an increasing function of η rather than a constant. In most cases, the proof technique employed relies on the cessation of switching in finite time. The selection of π and the definition of Σ_S are made to insure that this is so.

Although nonestimator based supervisors are prerouted tuners, the converse is not necessarily true. For example, it is quite possible for a supervisor employing "estimators" to use prerouted tuning to generate σ. Supervisors admitting this structure have in fact been studied in [40]. This reference actually examines the convergence properties of a variety of estimator-based switching logics.

The findings of [40] and earlier work clearly suggest that some of the concepts we've covered here have a universal character and may well be extendable to significantly broader classes of problems than have been considered so far. In the sequel we briefly summarize some preliminary thoughts along these lines.

Generalization: Let Q be a countable set. Suppose that for each $q \in \mathcal{Q}$, $A_q : \mathbb{R}^n \to \mathbb{R}^n$ is a smooth, possibly nonlinear function and that for some $q^* \in \mathcal{Q}$, the zero state of

$$\dot{x} = A_{q^*}(x)$$

is a globally asymptotically stable equilibrium. Assume that for each piecewise constant switching signal $\sigma : [0, \infty) \to \mathcal{Q}$, all solutions to the differential equation

$$\dot{x} = A_\sigma(x) \tag{10}$$

exist on $[0,\infty)$. Our aim is to briefly outline how one might go about constructing a nonestimator based supervisor, not depending on q^* or precise knowledge of the A_q, which cause all "supervised " (i.e., closed-loop) solutions to (10) to tend to zero as $t \to \infty$.

Suppose it is possible to construct a smooth function $b : \mathbb{R}^n \mapsto \mathbb{R}$ such that

$$||A_q(x)|| \leq ||b(x)||, \quad \forall x \in \mathbb{R}^n, \quad q \in \mathcal{Q} \tag{11}$$

and for some $q \in \mathcal{Q}$

$$\sup_{z \in \mathbb{R}^n} \frac{\int_0^\infty (||\phi(t,z)||^2 + ||b(\phi(t,z))||^2)dt}{\int_0^{\tau_D} (||\phi(t,z)||^2 + ||b(\phi(t,z))||^2)dt} = \mu < \infty \tag{12}$$

where $\phi : [0,\infty) \times \mathbb{R}^n \to \mathbb{R}^n$ is the flow of

$$\dot{x} = A_q(x) \tag{13}$$

initialized at z. Requirement (11) is relatively mild and can typically be satisfied without precise knowledge of the A_q. Implicit in (12) is the requirement that the zero state of (13) is {at least} an asymptotically stable equilibrium; in fact, for the requirement to make sense as it stands, all solutions to (13) would have to have finite $\mathcal{L}^2[0,\infty)$ norms.

We claim that the supervisor we've already described will accomplish the prescribed task provided

$$e \stackrel{\Delta}{=} \begin{bmatrix} x \\ b(x) \end{bmatrix} \tag{14}$$

The reasoning upon which this claim is based is as follows.

First of all note that satisfaction of (12) guarantees that η cannot exceed the least integer $\bar{\eta}$ no smaller than μ for which $h(\bar{\eta}) = q$. The argument which justifies this assertion exploits the inequality

$$\int_{t_0}^{\infty} \left\| \begin{bmatrix} \phi(t-t_0,z) \\ b(\phi(t-t_0,z)) \end{bmatrix} \right\|^2 dt \leq \bar{\eta} \int_{t_0}^{t_0+\tau_D} \left\| \begin{bmatrix} \phi(t-t_0,z) \\ b(\phi(t-t_0,z)) \end{bmatrix} \right\|^2 dt, \quad t_0 \geq 0,\ z \in \mathbb{R}^n$$

and is essentially the same as before. The inequality is a consequence of (12).

At this point we need a good working definition of detectability for nonlinear systems. Suppose we agree to call a smooth dynamical system of the form

$$\begin{aligned} \dot{x} &= A(x) \\ e &= C(x) \end{aligned} \tag{15}$$

detectable if there exists a positive definite, radially unbounded, continuously differentiable function $V : \mathbb{R}^n \to \mathbb{R}$ such that

$$\frac{\partial V}{\partial x} A(x) - ||C(x)||^2 < 0, \quad x \in \mathbb{R}^n,\ x \neq 0 \tag{16}$$

- The definition characterizes detectability more as a generalization of stability than of observability; note for example, that if $C(x) = x$, (15) may not satisfy the definition, even though for this example (15) would certainly have to be considered an observable system.
- In the linear case when $C(x) = Cx$ and $A(x) = Ax$, the standard definition of detectability is known to be equivalent to the existence of a positive definite matrix P which satisfies the matrix inequality $PA + A'P - C'C < 0$ [41]; since $V \stackrel{\Delta}{=} x'Px$ satisfies (16), the definition of detectability proposed here thus has the virtue of reducing to the standard one in the linear case.
- It can be easily shown that if (15) is a detectable nonlinear system and e has a finite $\mathcal{L}^2[0,\infty)$ norm along some solution x, then x must tend to zero. Thus the proposed definition fulfills the intuitively appealing requirement that smallness of the output of a detectable system ought to imply smallness of the system's state.

Returning to our problem we point out that (11) implies that for each fixed $q \in \mathcal{Q}$, the dynamical system

$$\begin{aligned} \dot{x} &= A_q(x) \\ e &= \begin{bmatrix} x \\ b(x) \end{bmatrix} \end{aligned} \tag{17}$$

is detectable through e. This can be verified using the function $V \stackrel{\Delta}{=} \frac{1}{2}||x||^2$.

The steps involved in showing that $x \to 0$ are clear. Since switching stops, π is bounded which means that e must have a finite $\mathcal{L}^2[0,\infty)$ norm. Suppose q is σ's final value. Then (17) governs the evolutions of x and e after switching stops. Because (17) is a detectable system and e has a finite $\mathcal{L}^2[0,\infty)$ norm, x must tend to zero as claimed.□

There are of course plenty of practical reasons why one would not want to seriously consider implementing the system just described. On the other hand, there are components of the preceding {e.g., the notion of detectability and how to use it} which will no doubt prove useful in the analysis of more meaningful algorithms.

One drawback of many "switched" control systems including the ones we've discussed so far, is that they make use of signal which grows monotonically with time. For the supervisor we've described this would be η. Since bounded monotone signals converge, switched systems which employ them tend to be fairly easy to analyze. The problem is that when $\mathcal{L}^\infty$ bounded noise and or exogenous disturbances signals are present, monotone signals tend to blow up. To get around this, it is generally necessary to eliminate monotone signals altogether, usually by introducing "forgetting factors" or "exponential weighting" of some form [8, 38, 24]. What this means is that with such modifications in place, switching can no longer be expected to terminate in finite time. As a result one is usually confronted with an analysis problem which is very much more challenging than that encountered in the noise-free case when monotone convergence

could be counted on. Because of this *there is a specific need for technical results appropriate to the analysis of systems within which switching never terminates.*

Perhaps the most serious criticism of the nonestimator approach is its reliance on prerouted tuning. Clearly if $\mathcal{Q}$ is a large set, one should not expect a prerouted supervisory control system to perform very well.

4.3 Estimator-Based Supervisors

The overall responsibility of any multi-controller supervisor can be divided into a scheduling task - deciding *when* to switch controllers - and a routing task - deciding *which* controller to switch to next. Nonestimator-based supervisors have the routing question decided for them and are thus designed to deal only with scheduling. It is natural to expect that improved overall performance can be achieved by employing a supervisor endowed with the capability of making *both* scheduling and routing decisions in real time. An important class of supervisors possessing this capability are those which are estimator-based. Estimator-based supervisors utilize a form of certainty equivalence and as such are in some ways quite similar to conventional estimator-based tuners encountered in parameter adaptive control.

Since an estimator-based supervisor is responsible for both scheduling and routing, it is not surprising that defining one should require a more detailed description of $\mathcal{C}_P$ then we've assumed so far. For illustrative purposes suppose $\mathcal{C}_P$ to be of the form

$$\mathcal{C}_P = \bigcup_{p \in \mathcal{P}} \mathcal{C}(p)$$

where $\mathcal{P}$ is a closed, bounded {possibly finite} subset of a real, finite-dimensional linear space. Here $\mathcal{C}(p)$ denotes the subclass

$$\mathcal{C}(p) = \{\nu_p + \delta : ||\delta||_\infty \leq \epsilon_p\}$$

where ν_p is a preselected, reduced, strictly proper *nominal transfer function*, ϵ_p is a real non-negative number and δ is a stable, strictly proper norm-bounded perturbation representing unmodelled dynamics of the additive type; $||\cdot||_\infty$ denotes the shifted infinity norm

$$||\delta||_\infty \stackrel{\Delta}{=} \sup_{s \in C(\lambda_u)} |\delta(s)|,$$

where λ_u is a prespecified positive number called the *unmodelled dynamics stability margin*, and $C(\lambda_u)$ is the subset of the complex plane consisting of all points on and to the right of the vertical line $s = -\lambda_u$. Assume for each $p \in \mathcal{P}$, that the allowable values of δ exclude transfer functions for which $\nu_p + \delta$ has unstable poles and zeros in common. All transfer functions in $\mathcal{C}_P$ are thus strictly proper, but not necessarily stable rational functions.

As before, we take as given a parameterized a family of admissible controller transfer functions $\mathcal{K}$ which has the property that for each transfer function τ in $\mathcal{C}_P$ there is at least one controller transfer function $\kappa \in \mathcal{K}$ which internally

stabilizes the interconnection shown in Figure 3. Because estimator-based supervisors base decision-making on the idea of certainty equivalence, to configure such a supervisor it is necessary to first specify a well-defined function F from the nominal process model transfer function class $\mathcal{N} \stackrel{\Delta}{=} \{\nu_p : p \in \mathcal{P}\}$ to $\mathcal{K}$ in such a way that the assignment $\nu_p \longmapsto F(\nu_p)$ meets prescribed specifications. Given F, a natural way to make this assignment explicit is to stipulate that $\mathcal{P}$ be a subset of $\mathcal{K}$'s parameter space $\mathcal{Q}$ and then to define $\kappa_p \stackrel{\Delta}{=} F(\nu_p)$ for each $p \in \mathcal{P}$. For the present we shall actually take $\mathcal{Q} = \mathcal{P}$. The reader should realize however that there are situations in which it is advantageous to choose $\mathcal{Q}$ larger than $\mathcal{P}$. For example, picking $\mathcal{Q}$ larger than $\mathcal{P}$ makes it possible to define generalized supervisors whose controller selection strategies are not based just on certainty equivalence alone {c.f. §4.4}.

Assume that the transfer functions in $\mathcal{K}$ satisfy the
Stability Margin Requirement: For each $p \in \mathcal{P}$ the real parts of the closed-loop poles of the feedback interconnection shown in Figure 5 are less than $-\lambda_S$ where λ_S some prespecified positive number called a *stability margin.*

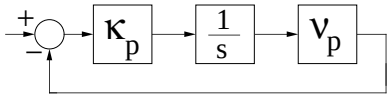

Fig. 5. Feedback Interconnection

In concept, an estimator-based supervisor can be explained in terms of the "multi-estimator" architecture shown in Figure 6.

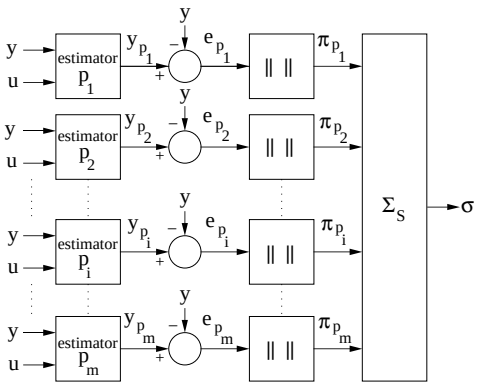

Fig. 6. Multi-Estimator Configured Supervisor

where each y_p is a suitably defined estimate of y which would be asymptotically

correct if ν_p were the process model's transfer function. For each $p \in \mathcal{P}$,

$$e_p \stackrel{\Delta}{=} y_p - y \tag{18}$$

denotes the pth *output estimation error*; π_p is a "normed" value of e_p or a "performance signal" which is used by the supervisor assess the potential performance of controller p. Σ_S is a switching logic whose function is to determine σ on the basis of the current values of the π_p.

The underlying decision making strategy used by an estimator-based supervisor of the 'non-prerouted' type is basically this: From time to time select for v, that candidate control signal v_p whose corresponding performance signal π_p is the smallest among the $\pi_p,\ p \in \mathcal{P}$. What makes a non-prerouted supervisor such as this distinctly different from a prerouted one is thus the philosophy underlying the method it uses to carry out its task. In particular, a non-prerouted supervisor decides which controller to put in the feedback loop, not by search along a predetermined route in $\mathcal{K}$, but rather by continuously comparing in real time suitably defined normed output estimation errors or performance signals associated with the admissible nominal process models. Motivation for this idea is obvious: the process model whose associated performance signal is the smallest, "best" approximates what the process is and thus the candidate controller designed on the basis of that model ought to be able to do the best job of controlling the process. The origin of this idea is of course the concept of certainty equivalence from parameter adaptive control.

By an estimator of y, based on transfer function ν_p, is meant a linear system of the form

$$\dot{x}_p = A_p x_p + d_p y + b_p u \tag{19}$$
$$y_p = c_p x_p \tag{20}$$

where $\{A_p + d_p c_p, b_p, c_p\}$ is a realization of ν_p and A_p is a stability matrix. It is easy to verify that any such realization necessarily fulfills the requirement that y_p be an asymptotically correct estimate of y if the process model transfer function were ν_p. Notice that such realizations are invariably detectable because of A_p's stability. For the present we are only going to consider realizations which are stabilizable as well, even though by doing so we are sidestepping some subtle but important issues{cf., §4.4}. There are many ways to construct estimators which meet these requirements. For example, if n is an upper bound on ν_p's McMillan Degree, y_p can always be be generated by an observer-based estimator of the form

$$\begin{aligned} \dot{x}_p &= A_O x_p + d_p y + b_p u \\ y_p &= c_O x_p \end{aligned} \tag{21}$$

where (c_O, A_O) is an n-dimensional, parameter-independent observable, stable pair and $\{A_O + d_p c_O, b_p, c_O\}$ is a stabilizable realization of ν_p. It is also possible

to generate y_p using an identifier-based estimator of the form

$$\begin{aligned}\dot{x}_I &= \begin{bmatrix} A_I & 0 \\ 0 & A_I \end{bmatrix} x_I + \begin{bmatrix} b_I \\ 0 \end{bmatrix} y + \begin{bmatrix} 0 \\ b_I \end{bmatrix} u \\ y_p &= c_p x_I \end{aligned}$$

where (A_I, b_I) is a parameter-independent, n-dimensional siso, controllable pair with A_I stable and

$$\left\{ \begin{bmatrix} A_I & 0 \\ 0 & A_I \end{bmatrix} + \begin{bmatrix} b_I \\ 0 \end{bmatrix} c_p \,,\ \begin{bmatrix} 0 \\ b_I \end{bmatrix},\ c_p \right\}$$

is a stabilizable realization of ν_p. Note that the state of this estimator is independent of p, whereas the state of the observer-based estimator in (21) is not. What this means is that if n is an upper bound on the McMillan Degrees of all of the nominal transfer functions in $\mathcal{N}$, then all of the y_p can be generated using a single estimator with shared state x_I and parameter-dependent readout map c_p.

There is a third way to generate y_p which is very similar to the second but which is especially well-suited to the set-point control problem under consideration. In this case one uses an identifier-based estimator Σ_E of the form

$$\dot{x}_E = \begin{bmatrix} A_E & 0 \\ 0 & A_E \end{bmatrix} x_E + \begin{bmatrix} b_E \\ 0 \end{bmatrix} y + \begin{bmatrix} 0 \\ b_E \end{bmatrix} v \tag{22}$$

$$y_p = c_p x_E \tag{23}$$

where (A_E, b_E) is a parameter-independent, $(n+1)$-dimensional siso, controllable pair with A_E stable and

$$\left\{ \begin{bmatrix} A_E & 0 \\ 0 & A_E \end{bmatrix} + \begin{bmatrix} b_E \\ 0 \end{bmatrix} c_p \,,\ \begin{bmatrix} 0 \\ b_E \end{bmatrix},\ c_p \right\}$$

is a stabilizable realization of $\frac{1}{s}\nu_p$. A state-shared implementation based on this estimator would then appear as in Figure 7. Naturally this architecture can only be implemented as it stands if the number of output estimation errors is finite; i.e., if $\mathcal{P}$ is a finite set. It turns out however that such a supervisor can often be implemented using a simpler architecture - one which permits $\mathcal{P}$ to contain a continuum of points. To explain why this is so, it is useful to formalize the idea of a supervisor.

By an *estimator-based supervisor* {cf, Figure 8} is meant a specially structured hybrid dynamical system whose output σ is a switching signal taking values in $\mathcal{Q}$ and whose inputs are v and y. Internally such a supervisor consists of three subsystems: a state-shared estimator Σ_E, a *performance weight generator* Σ_W and a *switching logic* Σ_S. Σ_W is a causal dynamical system whose inputs are x_E and y and whose state and output W is a "weighting matrix" which takes values in a linear space $\mathcal{W}$. W together with a suitably defined *performance function* $\Pi : \mathcal{W} \times \mathcal{P} \to \mathbb{R}$ determine, for each $p \in \mathcal{P}$, scalar-valued *performance signals* of the form

$$\pi_p = \Pi(W, p),\ p \in \mathcal{P} \tag{24}$$

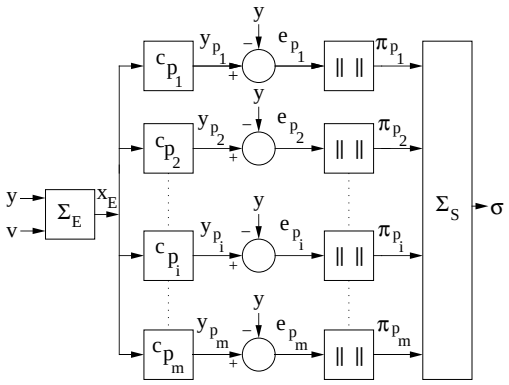

Fig. 7. State-Shared Estimator-Based Supervisor

These performance signals play the same role as before; i.e., π_p is considered to be a measure of the expected performance of control signal p. One possible pair of definitions for Σ_W and Π is

$$\dot{W} = -2\lambda W + \begin{bmatrix} x_E \\ y \end{bmatrix} \begin{bmatrix} x_E \\ y \end{bmatrix}' \tag{25}$$

and

$$\Pi(W,p) = [\,c_p \quad -1\,]\, W\, [\,c_p \quad -1\,]' \tag{26}$$

respectively where λ is a prespecified nonnegative number. In the light of (18) and (23) it is easy to see that these definitions imply that

$$\dot{\pi}_p = -2\lambda\pi_p + e_p^2, \quad p \in \mathcal{P} \tag{27}$$

Although we will deal here exclusively with such "exponentially weighted $\mathcal{L}^2$" performance signal, it should be noted that it is possible to realize other types of performance signals by defining W and Π in other ways. For example, if $\mathcal{P}$ is a finite set {say $\mathcal{P} = \{1, 2, \ldots, m\}$} and if Σ_W is the dynamical system $\dot{w}_p = |e_p|$, $p \in \mathcal{P}$ with state $w \stackrel{\Delta}{=} [\,w_1 \quad w_2 \quad \cdots \quad w_m\,]'$, defining $\Pi(w,p) \stackrel{\Delta}{=} w_p$ would realize the $\mathcal{L}^1$ performance signal $\dot{\pi}_p = |e_p|$. Note however that if $\mathcal{P}$ were not finite, this particular performance signal could not be realized with $\mathcal{W}$ finite dimensional.

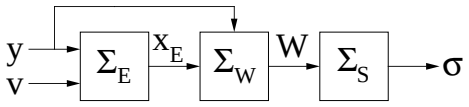

Fig. 8. Estimator-Based Supervisor

Hysteresis Switching: There are a number of different ways to define switching logic Σ_S. In the sequel we shall consider two. The first, called "Hysteresis Switching," was originally devised for switching between the members of a finite family of parameter adaptive controllers [11, 12, 13]. We shall explain this logic's basic attributes in the following manner.

Suppose $\{f_q : q \in \mathcal{Q}\}$ is a family of functions $f_q : \mathbb{R}^n \times [0, \infty) \to \mathbb{R}^n$. Our aim is to study the behavior of the dynamical system

$$\dot{x} = f_\sigma(x, t), \qquad x(0) = x_0 \tag{28}$$

where σ is a switching signal taking values in $\mathcal{Q}$[5]. Suppose that $\mathcal{P} = \mathcal{Q}$, and that W is a function of x and t which takes values in $\mathcal{W}$; i.e.,

$$W = g(x, t) \tag{29}$$

As before, suppose that $\Pi : \mathcal{W} \times \mathcal{P} \to \mathbb{R}$ is a performance function and that for $p \in \mathcal{P}$, $\pi_p \stackrel{\Delta}{=} \Pi(w, p)$ is a performance signal. What we want to do is to explain how to generate a switching signal σ which under certain conditions, converges to a value $\bar{q} \in \mathcal{Q}$ at which $\pi_{\bar{q}}$ is a bounded signal. The algorithm which generates σ is called a "hysteresis switching logic."

By a *hysteresis switching logic* is meant a hybrid dynamical system Σ_H whose input is W and whose state and output are both σ.

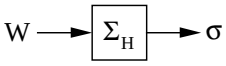

To specify Σ_H it is necessary to first pick a positive number $h > 0$ called a *hysteresis constant*. Σ_H's internal logic is then defined by the computer diagram shown in Figure 9 where for $X \in \mathcal{W}$, q_X denotes a value of $q \in \mathcal{Q}$ which minimizes $\Pi(X, q)$.

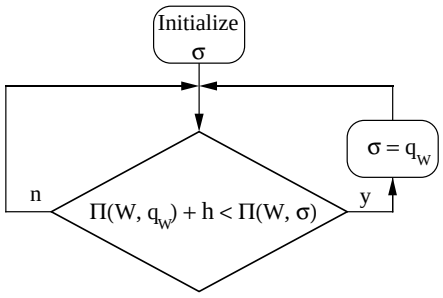

Fig. 9. Computer Diagram of Σ_H

[5] For the set-point control problem under consideration, x would represent the composite state $\{x_P, u, x_E, x_C, W\}$.

In interpreting this diagram it is to be understood that σ's value at each of its switching times $\bar{t}$ is its limit from above as $t \downarrow \bar{t}$. Thus if $\bar{t}_i$ and $\bar{t}_{i+1}$ are any two successive switching times, then σ is constant on $[\bar{t}_i, \bar{t}_{i+1})$. Note that the definition of Σ_H implies that $\pi_{\sigma(t)}(t) \leq \pi_q(t) + h, \quad t \geq 0, \quad q \in \mathcal{Q}$ and that $\pi_{\sigma(\bar{t})}(\bar{t}) \leq \pi_q(\bar{t}), \quad q \in \mathcal{Q}$ if $\bar{t}$ is a switching time.

The functioning of Σ_H is roughly as follows. Suppose that at some time t_0, Σ_H has just changed the value of σ to q. σ is then held fixed at this value unless and until there is a time $t_1 > t_0$ at which $\pi_p + h < \pi_q$ for some $p \in \mathcal{Q}$. If this occurs, σ is set equal to p and so on.

Note that since all the supervisor has to do is to compute values of $p \in \mathcal{P}$ which minimize $\Pi(W, p)$ at various times, there is in principle nothing to prevent $\mathcal{P}$ from containing a continuum of points. Of course the minimization problems to be solved must be tractable and the time it takes to compute these minima needs to be taken into account. We will discuss both of these points further in the sequel.

For the present our objective is to describe some of the properties of the closed-loop system determined by (28), (29) and Σ_H assuming that g and each f_q is at least locally Lipschitz in x and piecewise-continuous in t. Observe that because of the hysteresis constant h and the assumed smoothness of g and the f_q, there must exist an interval $(0, t_1)$ of maximal length on which σ is constant. Either this interval is the maximal interval of existence for x or it is not in which case x is bounded on $[0, t_1)$. If the latter is true, a switch must occur at t_1 and again because of the hysteresis constant h, the continuity of x and the smoothness of g and the f_q, there must be an interval $[t_1, t_2)$ of maximal length on which σ is constant. Continuing this reasoning we conclude that there must be an interval $[0, T)$ of maximal length on which there is a unique pair $\{x, \sigma\}$ with x continuous and σ piecewise constant, which satisfies (28) and (29). Moreover, on each proper subinterval $[0, \tau) \subset [0, T)$, σ can switch at most a finite number of times.

Our aim now is to characterize the limiting behavior of σ as $t \to T$. For this we need to make certain "open-loop" assumptions. Let $\mathcal{S}$ denote the class of all piecewise-constant functions $s : [0, \infty) \to \mathcal{Q}$. In what follows, for each $s \in \mathcal{S}$, T_s is the length of the maximal interval of existence for the equations

$$\dot{x} = f_{s(t)}(x, t), \qquad x(0) = x_0$$

and x_s is the corresponding solution. We make the following

Assumption 1 (Open-Loop)

1. For each $s \in \mathcal{S}$ and each $q \in \mathcal{Q}$, performance signal $\pi_q(t) = \Pi(g(x_s(t), t), q)$ has a limit (which may be infinite) as $t \to T_s$.
2. There exists at least one point $q^* \in \mathcal{Q}$ such that for each $s \in \mathcal{S}$, performance signal $\pi_{q^*}(t) = \Pi(g(x_s(t), t), q^*)$ is bounded on $[0, T_s)$.

These assumptions enable one to prove the following [12].

Lemma 1 Hysteresis Switching . *For fixed initial state* $(x_0, \sigma_0) \in \mathbb{R}^n \times \mathcal{Q}$, *let* (x, σ) *denote the unique solution to (28) and (29) with* σ *the output of* Σ_H - *and suppose* $[0, T)$ *is the largest interval on which this solution is defined. If the open-loop assumptions hold, there is a time* $T^* < T$ *beyond which* σ *is constant and no more switching occurs. Moreover,* $\pi_{\sigma(T^*)}$ *is bounded on* $[0, T)$.

Analysis: What we want to do next is to very briefly sketch how one might use the Hysteresis Switching Lemma to to analyze the closed-loop behavior of the supervisory control system shown in Figure 10.

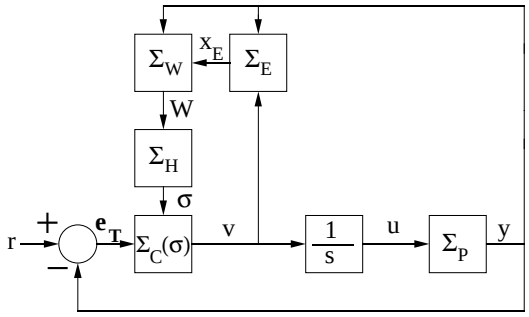

Fig. 10. Supervisory Control System Using Hysteresis Switching

Here $\Sigma_C(q)$ is a globally detectable/stabilizable realization of κ_q with state x_C, Σ_E is the globally detectable/stabilizable estimator defined by (22) and (23), and Σ_H is a hysteresis switching logic. Assume that $\lambda = 0$ and that Σ_W and Π are defined by (25) and (26) respectively. Therefore in this case, π_p is the $\mathcal{L}^2$ performance signal

$$\dot{\pi}_p = e_p^2, \quad p \in \mathcal{P} \tag{30}$$

Note that the Open-Loop Assumption 1 automatically holds because all of the π_p are monotone functions.

It can be shown [24] that there are constant vectors b and h, singly indexed matrices A_p, d_p, g_p, and $\bar{c}_p$ and doubly indexed matrices f_{qp} and c_{qp} such that for all constant r

$$\begin{bmatrix} x_E \\ x_C \end{bmatrix} = x + hr \tag{31}$$

where

$$\dot{x} = (A_l + b f_{\sigma l})x + d_l e_l \tag{32}$$

$$e_p = c_{pl} x + e_l \quad p \in \mathcal{P} \tag{33}$$

$$v = f_{\sigma l} x + g_\sigma e_l \tag{34}$$

$$\mathbf{e_T} = e_l + \bar{c}_l x \tag{35}$$

Because of the Stability Margin Requirement, it also turns out to be true that the matrix pairs $(c_{pl}, A_l + bf_{pl}),\ \ p, l \in \mathcal{P}$ are each detectable. These claims can be verified in a straight forward manner by direct analysis of the equations under consideration.

What we want to do next is to very briefly outline how one might use the Hysteresis Switching Lemma to analyze the closed-loop behavior of the supervisory control system shown in Figure 10 under the assumption that for some $p^* \in \mathcal{P}$, nominal transfer function ν_{p^*} matches or equals that of Σ_P. The exact matching assumption provides exactly one new piece of information, namely that e_{p^*} must go to zero as fast as $e^{-\lambda_E t}$ where $-\lambda_E$ is the largest of the real parts of the eigenvalues of A_E. Because of (30) this means that

$$\lim_{t\to\infty} \pi_{p^*}(t) \stackrel{\Delta}{=} C^* < \infty$$

Thus Open-Loop Assumption 2 is satisfied.

In view of the Hysteresis Switching Lemma there there must be a time T^* beyond which σ is constant and no more switching occurs. Moreover, $\pi_{\sigma(T^*)}$ must be bounded on the maximal interval of existence $[0, T)$ for solution to the overall system of equations involved. Because switching has stopped, it can be shown that the solution in question in fact exists globally {i.e., $T = \infty$}.

Suppose that $\bar{q}$ is the final value of σ. Since $\pi_{\bar{q}}$ is bounded on $[0, \infty)$, $e_{\bar{q}}$ must have a finite $\mathcal{L}^2[0, \infty)$ norm because of (30). Next observe that for t sufficiently large and $l \stackrel{\Delta}{=} p^*$, (32) can be written as

$$\dot{x} = (A_{p^*} + bf_{\bar{q}p^*})x + d_{p^*}e_{p^*} \tag{36}$$

In view of the detectability of $(c_{\bar{q}p^*}, A_{p^*} + bf_{\bar{q}p^*})$, there must exist a matrix k which stabilizes $A_{p^*} + bf_{\bar{q}p^*} + kc_{\bar{q}p^*}$. Thus because of (33), (36) can be rewritten as

$$\dot{x} = (A_{p^*} + bf_{\bar{q}p^*} + kc_{\bar{q}p^*})x - ke_{\bar{p}} + (k + d_{p^*})e_{p^*}$$

Since $A_{p^*} + bf_{\bar{q}p^*} + kc_{\bar{q}p^*}$ is a stability matrix and both $e_{\bar{q}}$ and e_{p^*} have finite $\mathcal{L}^2[0, \infty)$ norms, x must have a limit of zero as $t \to \infty$. Therefore x_E and x_C must have a finite limits because of (31). So also must v because of (34). Moreover, since x and e_{p^*} both tend to zero, so must $\mathbf{e_T}$ because of (35). Therefore $y \to r$. Since y and v have finite limits, and Σ_P's transfer function is nonzero at $s = 0$, u must have a finite limit as well. In other words, y, u, v, x_E, and x_C all tend to finite limits and $\mathbf{e_T} \to 0$.

Note how detectability has once again played a central role in the analysis. Together with the Hysteresis Switching Lemma it has enabled us to establish the limiting behavior of y, u, v, x_E, and x_C in a very elementary way.

The preceding is less than satisfactory for at least four important reasons:

1. If $r \neq 0$, W will grow without bound.
2. If noise and disturbances are present W will almost certainly grow without bound.
3. The analysis fails to account for unmodelled process dynamics

4. The analysis fails to account for computation time; i.e., the time it takes the supervisor to carry out the calculations necessary to select a new control.

A possible remedy for the first two problems would be to introduce a forgetting factor or exponential weighting in the definition of Σ_W in (25). For example, one might pick $\lambda > 0$. Of course any such change would make the resulting system substantially more difficult to analyze than the one we've been considering since π_{p^*} would no longer be monotone and switching would not necessarily terminate in finite time. Add in a small amount of unmodelled dynamics, and the analysis problem would become even more difficult because it would no longer possible to presume at the outset that e_{p^*} tends to zero or even that it is bounded. Some progress in dealing with these difficulties has recently been announced in [42].

Taking into account computation time makes things even more difficult. On the other hand, the reality of a positive computation time - however small - to some extent mitigates the need for hysteresis, since the only reason for introducing hysteresis in the first place was to prevent unbounded chatter [11]. Rather than further pursue this topic, we turn instead to an alternative switching logic which takes computation time directly into account and which results in a supervisory control system which can be shown to perform its function in the face of unmodelled dynamics and exogenous disturbances [43].

Dwell-Time Switching: By a *dwell-time switching logic* [15] Σ_D, is meant a hybrid dynamical system whose input and output are W and σ respectively, and whose state is the ordered triple $\{X, \tau, \sigma\}$.

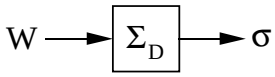

Here X is a discrete-time matrix which takes on sampled values of W, and τ is a continuous-time variable called a *timing signal*. τ takes values in the closed interval $[0, \tau_D]$, where τ_D is a prespecified positive number called a *dwell time*. Also assumed prespecified is a *computation time* $\tau_C \leq \tau_D$ which bounds from above for any $X \in \mathcal{W}$, the time it would take a supervisor to compute a value $p = p_X \in \mathcal{P}$ which minimizes $\Pi(X, p)$. Between "event times" τ is generated by a reset integrator according to the rule $\dot{\tau} = 1$. Event times occur when the value of τ reaches either $\tau_D - \tau_C$ or τ_D; at such times τ is reset to either 0 or $\tau_D - \tau_C$ depending on the value of Σ_D's state. Σ_D's internal logic is defined by the computer diagram shown in Figure 11 where p_X denotes a value of $p \in \mathcal{P}$ which minimizes $\Pi(X, p)$.

The functioning of Σ_D can be explained as follows. Suppose that at some time t_0, Σ_D has just changed the value of σ to p. At this instant τ is reset to 0. After $\tau_D - \tau_C$ time units have elapsed, W is sampled and X is set equal to this value. During the next τ_C time units, a value $p = p_X$ is computed which minimizes $\Pi(X, p)$. At the end of this period, when $\tau = \tau_D$, if $\Pi(X, p_X)$ is smaller than

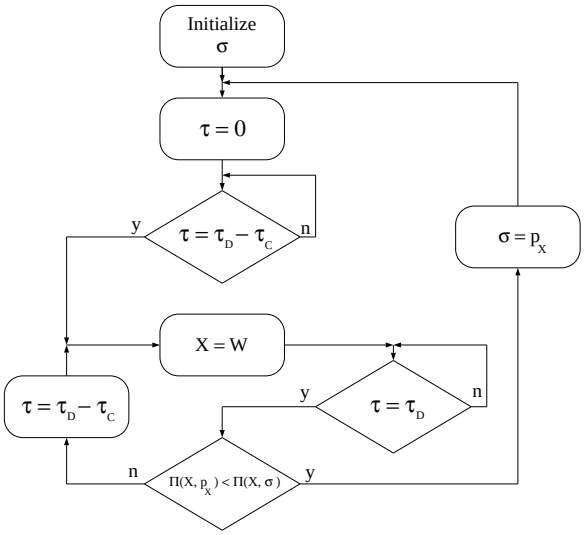

Fig. 11. Computer Diagram of Σ_D

$\Pi(X,\sigma)$, then σ is set equal to p_X, τ is reset to zero and the entire process is repeated. If on the other hand, $\Pi(X,\sigma)$ is less than or equal to $\Pi(X,p_X)$, τ is reset to $\tau_D - \tau_C$, W is again sampled, X takes on this new sampled value, minimization is again carried out over the next τ_C time units..... and so on.

Note that Σ_D is *scale independent* in that its output σ remains unchanged if its performance function-weighting matrix pair (Π, W) is replaced by another performance function-weighting matrix pair $(\bar{\Pi}, \bar{W})$ satisfying $\bar{\Pi}(\bar{W}, p) = \theta\Pi(W,p)$, $p \in \mathcal{P}$, where $\theta : [0,\infty) \to \mathbb{R}$ is a positive time function. This is because for any fixed t, the values of p which minimize $\Pi(W(t),p)$ are exactly the same as the values of p which minimize $\theta(t)\Pi(W(t),p)$.

Let us agree to call a piecewise-constant function $\sigma : [0,\infty) \to \mathcal{P}$ *admissible* if it either switches values at most once, or if it switches more than once and the set of time differences between each two successive switching times is bounded below by a positive number μ. The supremum of such values of μ is σ's *dwell time*. Because of the definition of Σ_D, it is clear its output σ will be admissible with dwell time no smaller than that of Σ_D. This means that switching cannot occur infinitely fast and thus that existence and uniqueness of solutions to the differential equations involved is not an issue.

Analysis: What we want to do next is to very briefly outline how one might analyze the closed-loop behavior of the supervisory control system shown in Figure 12 under the assumption that for some $p^* \in \mathcal{P}$, nominal transfer function ν_{p^*} matches or equals that of Σ_P. Unlike the supervisory control system considered in the last section, we will not {and probably cannot} prove that switching terminates in finite time.

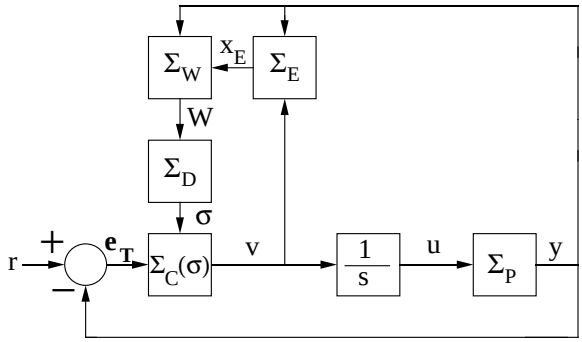

Fig. 12. Supervisory Control System Using Dwell-Time Switching

For the present we continue to assume $\lambda = 0$ and that Σ_W and Π are defined by (25) and (26) respectively. Thus

$$\dot{\pi}_p = e_p^2, \quad p \in \mathcal{P} \tag{37}$$

just as before. For simplicity we focus only on the case in which $\mathcal{P}$ is a finite set. Since (31)-(35) still hold, we can write

$$\begin{bmatrix} x_E \\ x_C \end{bmatrix} = x + hr \tag{38}$$

and

$$\dot{x} = (A_{p^*} + bf_{\sigma p^*})x + d_{p^*}e_{p^*} \tag{39}$$

$$e_p = c_{pp^*}x + e_{p^*} \quad p \in \mathcal{P} \tag{40}$$

$$v = f_{\sigma p^*}x + g_\sigma e_{p^*} \tag{41}$$

$$\mathbf{e_T} = e_{p^*} + \bar{c}_{p^*}x \tag{42}$$

Since the exact matching hypothesis implies that e_{p^*} goes to zero as fast as $e^{-\lambda_E t}$, it must be true that the set

$$\mathcal{P}^* \triangleq \left\{ p : \int_0^\infty ||e_p||^2 dt < \infty, \quad p \in \mathcal{P} \right\} \tag{43}$$

is nonempty. The assumption that $\mathcal{P}$ is a finite set can be used to prove that there must be a finite time t^* beyond which σ takes values only in $\mathcal{P}^*$ [24].

Let $\{c_{p_1p^*},\ c_{p_2p^*}, \ldots, c_{p_mp^*}\}$ be a basis for the span of $\{c_{pp^*} : p \in \mathcal{P}^*\}$. Define $C = [\,c'_{p_1p^*} \quad c'_{p_2p^*} \quad \cdots \quad c'_{p_mp^*}\,]'$ and

$$\bar{e} = Cx \tag{44}$$

These definitions together with (40) imply that $e_{p_i} - e_{p^*}$ is the ith entry of $\bar{e}$. Since each such entry has a finite $\mathcal{L}^2[0,\ \infty)$ norm, $\bar{e}$ must have a finite $\mathcal{L}^2[0,\ \infty)$ norm as well. Note also that the definition of C implies that there must be a bounded

function $s : \mathcal{P}^* \to \mathbb{R}^{m\times 1}$ for which $s(p)C = c_{pp^*},\ p \in \mathcal{P}^*$. In view of this and the previously noted detectability of the matrix pairs $(c_{pl}, A_l + bf_{pl}),\ \ p, l \in \mathcal{P}$, it must be that the matrix pair $(C, A_{p^*} + bf_{pp^*})$ is detectable for each $p \in \mathcal{P}^*$.

Note that for any appropriately sized, matrix $p \longmapsto K_p$ which is bounded on $\mathcal{P}^*$, (39) can be rewritten as

$$\dot{x} = (A_{p^*} + bf_{\sigma p^*} + K_\sigma C)x - K_\sigma \bar{e} + d_{p^*} e_{p^*}$$

for $t \geq t^*$. Suppose that such a function K_p can be shown to exist for which the time-varying matrix $A_{p^*} + bf_{\sigma p^*} + K_\sigma C$ is exponentially stable. Then because $\bar{e}$ and e_{p^*} have finite $\mathcal{L}^2[0,\infty)$ norms, x would tend to zero. Hence x_C and x_E would tend to finite limits because of (38). Moreover since e_{p^*} tends to zero, (41) and (42) would imply that v and $\mathbf{e_T}$ tend to zero as well. As a consequence, y would tend to r and u would tend to finite limit; the latter would be true because of the converging of y and v to constant values and because Σ_P's transfer function is nonzero at $s = 0$. In other words, to show that $y \to r$ and that x_C, x_E and u tend to finite limits its enough to show that $A_{p^*} + bf_{\sigma p^*} + K_\sigma C$ is exponentially stable for some suitably defined function K_p.

We claim that a function K_p exists provided τ_D is sufficiently large. To understand why this is so, first recall that $(C, A_{p^*} + bf_{pp^*})$ is detectable for each $p \in \mathcal{P}^*$. Thus for each such p there must be a constant matrix K_p which stabilizes $A_{p^*} + bf_{pp^*} + K_pC$. Therefore for each $p \in \mathcal{P}^*$ it is possible to find numbers $a_p \geq 0$ and $\lambda_p > 0$ for which

$$|e^{(A_{p^*} + bf_{pp^*} + K_pC)t}| \leq e^{(a_p - \lambda_p t)} \quad t \geq 0$$

Since $\frac{a_p}{\lambda_p}$ is an upper bound on the time it takes for $|e^{(A_{p^*} + bf_{pp^*} + K_pC)t}|$ to drop below one in value, it is not surprising that the state transition matrix of $A_{p^*} + bf_{\sigma p^*} + K_\sigma C$ will be exponentially stable provided

$$\tau_D > \sup_{p\in\mathcal{P}^*} \left\{ \frac{a_p}{\lambda_p} \right\}$$

This in fact can be shown to be true [24]. Thus we may conclude that if τ_D is chosen large enough, then u, x_C and x_E must converge to finite limits and and y must tend to r.

Performance Signals: One of the problems with the preceding is that W will not remain bounded if $r \neq 0$. One easy way to remedy this problem is as follows.

Under the exact matching hypothesis, $e_{p^*} \to 0$ as fast as $e^{-\lambda_E t}$. Thus there is a non-negative constant C_0 such that $e^2_{p^*}(t) \leq C_0 e^{-2\lambda_E t}$. Pick $\lambda \in (0, \lambda_E)$. Let Π and π_p be defined as in (24) and (26) respectively, but rather than using (25) to generate W, use the equation

$$\dot{W} = e^{2\lambda t} \begin{bmatrix} x_E \\ y \end{bmatrix} \begin{bmatrix} x_E \\ y \end{bmatrix}' \tag{45}$$

instead. Clearly

$$\dot{\pi}_p = e^{2\lambda t} e_p^2$$

As defined, π_p has three crucial properties:

1. For each $p \in \mathcal{P}$, π_p is monotone nondecreasing.
2. $\lim_{t\to\infty} \pi_{p^*} \stackrel{\Delta}{=} C^* \leq \pi_{p^*}(0) + \int_0^\infty C_0 e^{-2(\lambda_E - \lambda)t} dt < \infty$
3. If $\mathcal{P}^*$ is defined as before, then σ must take values only within $\mathcal{P}^*$ beyond some finite time.

These are precisely the properties needed to define C and $\bar{e}$ as in (44) so that $\bar{e}$ has a finite $\mathcal{L}^2[0,\infty)$ norm and that $(C, A_{p^*} + bf_{pp^*})$ is detectable for each $p \in \mathcal{P}^*$. In other words, if one were to use (45) to generate W, then the convergence properties of y, x_E, x_C and u would still hold.

Now consider replacing W with the "scaled" weighting matrix

$$\bar{W} \stackrel{\Delta}{=} e^{-2\lambda t} W \tag{46}$$

Note that $\Pi(\bar{W}, p) = e^{-2\lambda t}\Pi(W, p), \quad p \in \mathcal{P}$. In the light of the scale independence property of Σ_D noted previously, it must be that replacing W with $\bar{W}$ has no effect on σ and consequently on y, x_E, x_C and u. The key point here is that the weighting matrix $\bar{W}$ defined by (46) can also be generated directly by the stable dynamical system

$$\dot{\bar{W}} = -2\lambda\bar{W} + \begin{bmatrix} x_E \\ y \end{bmatrix} \begin{bmatrix} x_E \\ y \end{bmatrix}' \tag{47}$$

Moreover, since y and x_E, tend to finite limits, it must be that $\bar{W}$ {and therefore its sampled state $\bar{X}$} tend to finite limits as well. Thus at this point we may conclude that if τ_D is chosen large enough, if λ is picked in $(0, \lambda_E)$, and if W is generated by (25), then u, x_C, x_E, *and* W must converge to to finite limits and and y must tend to r.

Fast Switching: A key step in the analysis just given was to show that for the family of detectable pairs $\{(C, A_{p^*} + bf_{pp^*}) : p \in \mathcal{P}^*\}$, there exists a a bounded, output injection function K_p and a dwell time τ_D for which $A_{p^*} + bf_{\sigma p^*} + K_\sigma C$ is exponentially stable for any admissible switching function σ with dwell time no smaller than τ_D. It turns out that for *any* given positive dwell time τ_D, it is possible to find a function K_p which exponentially stabilizes $A_{p^*} + bf_{\sigma p^*} + K_\sigma C$ for any admissible switching function σ with dwell time no smaller than τ_D [24].

To reader should realize that detectability of such matrix pairs is by itself *not* sufficient for the existence of a function K_p with the aforementioned property. To understand why, just consider the situation in which a family of detectable pairs of the form $\{(C, A_p) : p \in \mathcal{P}\}$ has a zero readout matrix C; in this case each A_p must be a stability matrix and $A_p + K_p C = A_p$ for all K_p. It is well known that if the A_p do not commute with each other, exponential stability of A_σ cannot in general be assured unless τ_D is large enough; for an example see [44]. In other words, there are families of detectable pairs of the form $\{(C, A_p) : p \in \mathcal{P}\}$ for which no stabilizing function K_p exists if τ_D is too small. What's especially interesting is that if $\{(C, A_p) : p \in \mathcal{P}\}$ is a family of *observable* matrix pairs, then no matter how small τ_D is, there does in fact exist a matrix function K_p

with the required stabilizing property. This is an immediate consequence of the following result [14].

Squashing Lemma: *Let (C, A) be a fixed, constant, observable matrix pair, and let τ_0 be a positive number. For each positive number δ there exists a positive number λ and a constant output-injection matrix K for which*

$$|e^{(A+KC)t}| \leq \delta e^{-\lambda(t-\tau_0)}, \quad t \geq 0 \tag{48}$$

The way to construct K_p for a family of observable pairs such as $\{(C, A_p) : p \in \mathcal{P}\}$, is as follows. Pick $\delta \in (0, 1)$, set $\tau_0 = \tau_D$ and for each $p \in \mathcal{P}$ use the Squashing Lemma to find a value of K_p for which

$$|e^{(A_p+K_pC)t}| \leq \delta e^{-\lambda(t-\tau)}, \quad t \geq 0$$

It can be shown that with K_p so chosen, $A_\sigma + K_\sigma C$ will be exponentially stable if σ is any admissible switching signal with dwell time no smaller than τ_0 [24].

Unfortunately, for the problems of interest in this paper, the matrix pairs in $\{(C, A_{p^*} + bf_{pp^*}) : p \in \mathcal{P}^*\}$ cannot be assumed to be observable without a definite loss of generality. On the other hand, observability is in general sufficient for stabilizability whereas detectability is not. The way out of this dilemma has been to make use of additional properties of the matrices under consideration. A typical result along these lines is the following.

Switching Theorem: *Let $\lambda_0 > 0$ and $\tau_0 > 0$ be fixed. Let $(C_{q_0\times n}, A_{n\times n}, B_{n\times m})$ be a left invertible system. Suppose that $\{(C_p, F_p) : p \in \mathcal{P}\}$ is a closed, bounded subset of matrix pairs in $\mathbb{R}^{q\times n} \oplus \mathbb{R}^{m\times n}$ with the property that for each $p \in \mathcal{P}$, $(C_p, \lambda_0 I + A + BF_p)$ is detectable. There exist a constant $a \geq 0$ and bounded, matrix-valued output injection functions $p \longmapsto H_p$ and $p \longmapsto K_p$ on $\mathcal{P}$ which, for any admissible switching signal $\sigma : [0, \infty) \to \mathcal{P}$ with dwell time no smaller than τ_0, causes the state transition matrix of*

$$A + K_\sigma C_\sigma + H_\sigma C + BF_\sigma$$

to satisfies

$$|\Phi(t, \mu)| \leq e^{(a-\lambda_0(t-\mu))}, \quad t \geq \mu \geq 0$$

Using this theorem it has been possible prove that for any dwell time greater than zero and any value of $\lambda \in (0, \lambda_E)$, the supervisory control system we've been discussing achieves set-point regulation and global boundedness [24]. It has also been possible to show that these results continue to hold in the face of norm bounded unmodelled dynamics provided λ is further constrained to be smaller than both the stability margin λ_S and the unmodelled dynamics stability margin λ_u [43]. Moreover the introduction of $\mathcal{L}^\infty$ bounded noise and disturbance inputs cannot destabilize the system.

4.4 Cyclic Switching

As we have just explained, estimator-based supervisors generate control signals in accordance with the idea of certainty equivalence; i.e., at each instant of time the controller in feedback with the process is based on a current estimate of what the nominal process model transfer function is; such estimates are selected from a suitably defined admissible nominal process model transfer function set $\mathcal{N}$. Because $\mathcal{N}$ must be finitely parameterized, it can always be regarded as a subset of a finite dimensional linear space. In practice, $\mathcal{N}$ is typically chosen to best satisfy a number of conflicting requirements. For example, $\mathcal{N}$ should be "big" enough to ensure that $\mathcal{C}_P$ includes a transfer function model of the process. If $\mathcal{N}$ contains a continuum of transfer functions, then for on-line model estimation {i.e., minimization of $\Pi(W,p)$ } to be tractable, $\mathcal{N}$ should be convex or at least the union of a finite number of convex sets. Since each transfer function in $\mathcal{N}$ is a candidate process model transfer function, for the formulated problem to make sense, each such transfer function should be at least stabilizable {i.e., without any unstable poles and zeros in common}.

It is not very difficult to see that these are conflicting requirements. In particular, stabilizability, convexity and largeness of $\mathcal{N}$ are at odds. If stabilizability and largeness are required, then convexity and consequently tractability must be sacrificed. If convexity and stabilizability are required, then $\mathcal{N}$ must be "small."

A way out of this dilemma, which enables one to achieve tractability while retaining stabilizability and largeness, is to embed $\mathcal{N}$ in a larger set of 'admissible' transfer functions $\bar{\mathcal{N}}$ which is convex, but which is not restricted to have only stabilizable transfer functions. Naturally those transfer functions in $\bar{\mathcal{N}}$ which are not stabilizable cannot be candidate process model transfer functions. Nevertheless, because of the tractability issue it is useful to consider such transfer functions to be *admissible for estimation purposes.* Therefore an alternative to certainty equivalence is needed for selecting controllers when such transfer functions are encountered during the on-line estimation process. Such an alternative, based on the concept of "cyclic switching," has recently been proposed for applications in parameter-adaptive control where the same problem also arises [45, 14]. The aim of this section is to explain what cyclic switching is within the context of the set-point problem we've been considering.

We will be concerned exclusively with the case when $\mathcal{N}$ contains a continuum of reduced transfer functions. For simplicity assume that each such transfer function has the same McMillan Degree n. This means that $\mathcal{N}$ can be viewed as a subset of the $2n$-dimensional linear space of strictly proper {unreduced} rational functions whose denominators are monic and of degree n.

As before we assume that $\mathcal{P}$ is a closed, bounded subset of a finite dimensional linear space. Assume in addition that the coefficients of ν_p are defined on this space as affine linear functions. Assuming

$$\left\{ \begin{bmatrix} A_E & 0 \\ 0 & A_E \end{bmatrix} + \begin{bmatrix} b_E \\ 0 \end{bmatrix} c_p \,,\ \begin{bmatrix} 0 \\ b_E \end{bmatrix},\ c_p \right\}$$

again realizes ν_p, this means that c_p will also be an affine linear function. As a

consequence, the parameterized performance signal $\Pi(W,p)$ defined by

$$\Pi(W,p) = \begin{bmatrix} c_p & -1 \end{bmatrix} W \begin{bmatrix} c_p & -1 \end{bmatrix}' \tag{49}$$

$$\dot{W} = -2\lambda W + \begin{bmatrix} x_E \\ y \end{bmatrix} \begin{bmatrix} x_E \\ y \end{bmatrix}' \tag{50}$$

will be a quadratic function of p.

We are interested in the case when $\mathcal{P}$ is not necessarily convex since convexity of $\mathcal{P}$ would imply convexity of $\mathcal{N}$. To ensure a tractable minimization problem, we presume that $\mathcal{P}$ has been embedded in a conveniently chosen, closed, bounded convex subset $\bar{\mathcal{P}}$ {e.g., the convex hull of $\mathcal{P}$} and that the set of admissible nominal transfer functions has been enlarged to $\bar{\mathcal{N}} \stackrel{\Delta}{=} \{\nu_p : p \in \bar{\mathcal{P}}\}$. This reduces the problem of minimizing $\Pi(W,p)$ over $\bar{\mathcal{P}}$ to a finite dimensional convex, quadratic programming problem. Such problems are highly tractable and many fast algorithms for solving them are known.

We shall assume that all of the points $p \in \bar{\mathcal{P}}$ {if any} at which $\frac{1}{s}\nu_p$ has a pole-zero cancellation are in the interior of a specified closed set $\mathcal{S} \subset \bar{\mathcal{P}}$, called a *singular region*. {Therefore $\frac{1}{s}\nu_p$ can't have any pole-zero cancellations on the closure of $\bar{\mathcal{P}} - \mathcal{S}$.} It is reasonable to require $\mathcal{C}_P$ and $\{\nu_p : p \in \mathcal{S}\}$ to be disjoint.

In the sequel we will define a generalized supervisor whose decision making strategy takes into account the possibility that there may be times at which the best possible admissible transfer function estimate, determined by minimizing $\Pi(W,p)$ over $\bar{\mathcal{P}}$, falls within the singular set $\{\nu_p : p \in \mathcal{S}\}$. To define such a supervisor two things are needed:

Controller Requirements:

1. A bounded set of controller transfer functions $\{\kappa_q : q \in (\bar{\mathcal{P}} - \mathcal{S})\}$ which satisfies the *Stability Margin Requirement* on $\bar{\mathcal{P}} - \mathcal{S}$; i.e., for each $p \in (\bar{\mathcal{P}} - \mathcal{S})$ the real parts of the closed-loop poles of the feedback interconnection shown in Figure 13 are less than $-\lambda_S$.

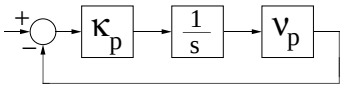

Fig. 13. Feedback Interconnection

   Since $\frac{1}{s}\nu_p$ has no pole-zero cancellations on the closure of $\bar{\mathcal{P}} - \mathcal{S}$, such a family clearly exists.
2. A set of real gains $\{g_1, g_2, \ldots, g_{n_S}\}$ which fulfills the *Observation Requirement*; i.e., for each $l \in \mathcal{P}$ and each $p \in \mathcal{S}$, there is a value of $q \in \{1, 2, \ldots, n_S\}$ for which the feedforward interconnection of controllable, observable realizations of ν_p and ν_l shown in Figure 14 is observable through e_{ff}. It can be shown that such a family exists because of the assumed disjointness of $\mathcal{N}$ and $\mathcal{S}$ [14].

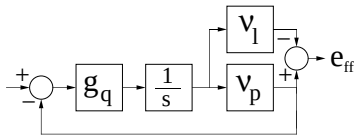

Fig. 14. Feedforward Interconnection

Assume that $\{\kappa_q : q \in (\bar{\mathcal{P}} - \mathcal{S})\}$ is a bounded set of controller transfer functions which satisfies the Stability Margin Requirement and that $\{\kappa_1, \kappa_2, \ldots, \kappa_{n_S}\}$ is a finite family of gains which satisfies the Observation Requirement. In addition, adopt the notation $\mathcal{I} \stackrel{\Delta}{=} \{1, 2, \ldots, n_S\}$ and write $\mathcal{Q}$ for the disjoint union $\mathcal{Q} \stackrel{\Delta}{=} (\bar{\mathcal{P}} - \mathcal{S}) \cup \mathcal{I}$. Suppose $\Sigma_C(q)$ is a globally detectable/stabilizable realization of κ_q on $\mathcal{Q}$.

The overall structure of the supervisory control system we want to consider is the same as before, except that now instead of Σ_D, the supervisor uses a yet-to-be-defined "dwell-time/cyclic logic" Σ_{DC}.

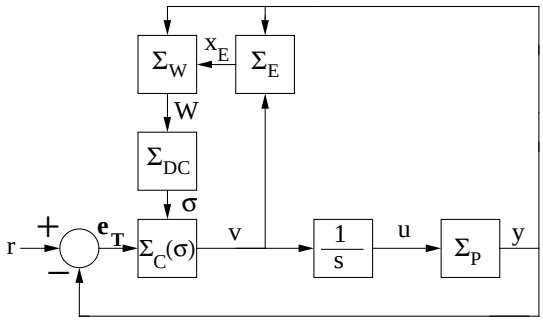

Fig. 15. Supervisory Control System Using Dwell-Time/Cyclic Switching

Σ_{DC} is essentially a combined version of Σ_D and the cyclic switching logic of [14]. The underlying strategy upon which Σ_{DC}'s logic is predicated can be explained roughly as follows. Consider again the by now familiar equations

$$\begin{bmatrix} x_E \\ x_C \end{bmatrix} = x + hr \tag{51}$$

$$\dot{x} = (A_l + bf_{\sigma l})x + d_l e_l \tag{52}$$

$$e_p = c_{pl}x + e_l \quad p \in \bar{\mathcal{P}} \tag{53}$$

$$v = f_{\sigma l}x + g_\sigma e_l \tag{54}$$

$$\mathbf{e_T} = e_l + \bar{c}_l x \tag{55}$$

which hold for all constant r and all $l, p \in \mathcal{P}$. Σ_{DC}'s strategy stems from two

facts, each a direct consequence of one of the two corresponding **Controller Requirement** stipulated above:

1. For each $l \in \mathcal{P}$ and each $p \in \bar{\mathcal{P}} - \mathcal{S}$ the matrix pair $(c_{pl}, A_l + bf_{ql})|_{q=p}$ is detectable.
2. For each $l \in \mathcal{P}$ and each $p \in \mathcal{S}$ there exists a $q \in \mathcal{I}$ such that $(c_{pl}, A_l + bf_{ql})$ is detectable.

Intuition: Here roughly is the idea upon which cyclic switching is based. Think of the supervisor as performing two separate tasks - one estimation and the other controller selection. The estimation task amounts to minimizing $\Pi(W, p)$ over $\bar{\mathcal{P}}$ and goes on over and over without interruption; this generates a sequence of values $\widehat{p} \in \mathcal{P}$. Meanwhile the supervisor tries to select controller's in such as way as to maintain "detectability" through $e_{\widehat{p}}$, at least on the average. Why? Because detectability through $e_{\widehat{p}}$ implies smallness of x whenever $e_{\widehat{p}}$ is small - and smallness of $e_{\widehat{p}}$ ought to be a consequence of the estimation process. So here's how the supervisor achieves "detectability" through $e_{\widehat{p}}$: If $\widehat{p} \in \bar{\mathcal{P}} - \mathcal{S}$, the supervisor relies on property 1 above and certainty equivalence: detectability is achieved by setting $\sigma = \widehat{p}$. On the other hand, if $\widehat{p}$ enters $\mathcal{S}$, the supervisor relies on property 2: in this case "detectability" is achieved on the average by stepping σ through each of the values $\mathcal{I}$, holding fixed on each such value for a prespecified amount of time.

Formally a *Dwell-Time/Cyclic Switching Logic* Σ_{DC} is a hybrid dynamical system whose input and output are W and σ respectively, and whose state is the ordered quintuple $\{X, \widehat{p}, \tau, \beta, \sigma\}$. X is a discrete-time matrix which takes on sampled values of W, $\widehat{p}$ is a discrete-time variable taking values in $\bar{\mathcal{P}}$, τ is a continuous-time timing signal as before, and β is a logic variable taking values in $\{0,\ 1\}$. τ takes values in the closed interval $[0,\ \max\{n_S\tau_S, \tau_D\}]$, where τ_D and τ_S are a prespecified positive numbers called a *dwell time* and a *cycle dwell time* respectively. As before $\tau_C \leq \max\{n_S\tau_S, \tau_D\}$ is a prespecified computation time which bounds from above for any $X \in \mathcal{W}$, the time it would take the supervisor to compute a value $p = p_X \in \mathcal{P}$ which minimizes $\Pi(X, p)$. Between "event times" τ is generated by a reset integrator according to the rule $\dot{\tau} = 1$. Such event times occur for $\beta \in \{0, 1\}$, when the value of τ reaches either $T(\beta) - \tau_C$ or $T(\beta)$ where $T(0) \stackrel{\Delta}{=} \tau_D$ and $T(1) \stackrel{\Delta}{=} n_S\tau_S$; at such times τ is reset to either 0 or $T(\beta) - \tau_C$ depending on the value of Σ_{DC}'s state. Σ_{DC}'s internal logic is defined by the computer diagram shown in Figure 16 where p_X denotes a value of $p \in \mathcal{P}$ which minimizes $\Pi(X, p)$.

The functioning of Σ_{DC} can be explained as follows. Suppose that at some time t_0, Σ_S has just changed the value of $\widehat{p}$. Depending on whether $\widehat{p} \in \mathcal{S}$ or not, one of two different epochs can occur:

- Suppose $\widehat{p} \notin \mathcal{S}$. In this case β is set equal to 0, σ is set equal to $\widehat{p}$ and τ is reset to 0. After $\tau_D - \tau_C$ time units have elapsed, W is sampled and X is set equal to this value. During the next τ_C time units, a value $p = p_X$ is

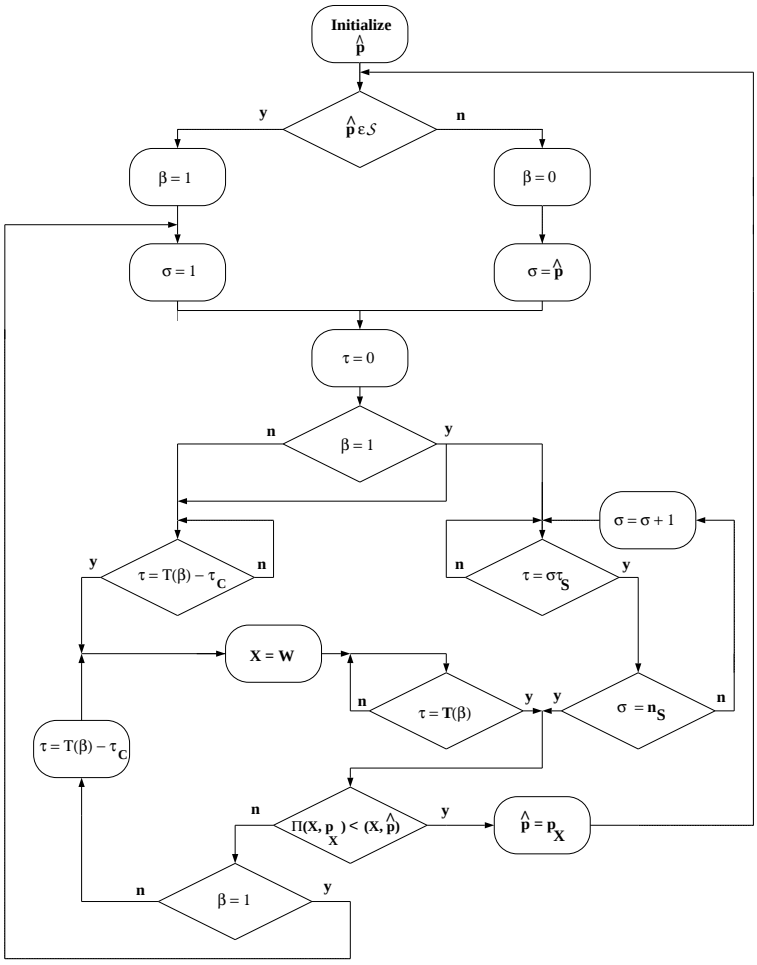

Fig. 16. Dwell-Time/Cyclic Switching Logic Σ_{DC}

computed which minimizes $\Pi(X,p)$. At the end of this period, when $\tau = \tau_D$, if $\Pi(X,p_X)$ is smaller than $\Pi(X,\widehat{p})$, then $\widehat{p}$ is set equal to p_X and the logic goes back to again test whether or not $\widehat{p} \in \mathcal{S}$. If, on the other hand, $\Pi(X,\widehat{p})$ is less than or equal to $\Pi(X,p_X)$, τ is reset to $\tau_D - \tau_C$, W is again sampled, X takes on this new sampled value, minimization is again carried out over the next τ_C time units..... and so on.

– Suppose $\widehat{p} \in \mathcal{S}$. In this case β is set equal to 1, τ is reset to 0, and two distinct sequences of events occur simultaneously, each lasting $n_S \tau_S$ time units:

1. At $\tau = 0$, a switching cycle is executed[6].
2. At $\tau = n_S\tau_S - \tau_C$, W is sampled and X is set equal to this value. During the next τ_C time units, a value $p = p_X$ is computed which minimizes $\Pi(X,p)$. At the end of this period, when $\tau = n_S\tau_S$, if $\Pi(X,p_X)$ is smaller than $\Pi(X,\widehat{p})$, then $\widehat{p}$ is set equal to p_X and the logic goes back to again test whether or not $\widehat{p} \in \mathcal{S}$. If, on the other hand, $\Pi(X,\widehat{p})$ is less than or equal to $\Pi(X,p_X)$, τ is reset to 0, another switching cycle is executed ....and so on.

An analysis of the supervisory control system just described can be found in [46]. For analyses in more traditional adaptive control contexts see [14, 22, 47]. The techniques used [14, 22] are similar to those outlined in the last section. The discrete time case is completely analyzed in [47].

5 Switched Linear Systems

Existing results concerned with the types of switched systems we've been discussing deal mainly with questions of stability, global boundedness and convergence. Interesting as they may be, these results are in many ways less than one might hope for. Especially lacking we think, are results of a more *quantitative* nature. What are needed are good norm bound estimates for allowable unmodelled dynamics. Also needed is a clearer understanding of the relationships between these estimates and design parameters.

Resolution of issues such as these calls for a better understanding of the basic properties of switched systems than we have at present. Needed is a catalog of basic results analogous to those for non-switched linear systems. In the sequel we briefly discuss some of the technical questions suggested by types of problems we've been discussing.

5.1 Stability of Switched Linear Systems

Let $\mathcal{P}$ be either a finite set or a closed, bounded subset of a finite dimensional linear space and let $\mathcal{A} = \{A_p : p \in \mathcal{P}\}$ be a closed, bounded set of real $n \times n$ matrices. Within this context one can formulate a number of different stability problems. For example, one can seek to find a switching logic Σ_S, with input x and piecewise-constant output σ which uniformly exponentially stabilizes

$$\dot{x} = A_\sigma x \tag{56}$$

in the sense that there are positive constants a and λ such that all solutions x to (56) {in closed-loop with Σ_S} exist and are norm bounded in time by $ae^{-\lambda t}$

[6] The supervisor *executes* a switching cycle at clock time $\tau = 0$ by setting $\sigma(t_0 + \tau) = s(\tau)$, $\tau \in [0, n_S\tau_S)$ where t_0 is the actual time τ was reset to 0 and $s : [0, n_S\tau_S) \to \mathcal{I}$ is the piecewise-constant function whose value is i on the subinterval $[(i-1)\tau_S, i\tau_S)$, $i \in \mathcal{I}$.

times the initial normed value of x. The synthesis of such a state driven switching logic seems to be very challenging. See [48] for a discussion of some recent results along these lines.

A somewhat less ambitious goal would be as follows: For a given set $\mathcal{A} = \{A_p : p \in \mathcal{P}\}$ and a given class of switching signals $\mathcal{S}$, find conditions under which there exist positive numbers λ and a such that for each $\sigma \in \mathcal{S}$ the state transition matrix of A_σ satisfies

$$|\Phi(t,\tau)| \leq ae^{-\lambda(t-\tau}, \quad t \geq \tau \geq 0 \tag{57}$$

Interesting choices for $\mathcal{S}$ would include

1. $\mathcal{S}_1 \stackrel{\Delta}{=}$ all piecewise constant switching signals
2. $\mathcal{S}_2 \stackrel{\Delta}{=}$ all piecewise constant switching signals with dwell times no less than some positive number τ_D.
3. $\mathcal{S}_3 \stackrel{\Delta}{=}$ the type of switching signals generated by Σ_{DC} {cf., [46]}.

Here are some easy to derive sufficient conditions under which there exist a and λ such that (57) holds for every piecewise constant $\sigma \in \mathcal{S}_1$

1. There exists a norm $||\cdot||$ on $\mathbb{R}^{n\times n}$ and a positive number λ such that
$$||e^{A_p t}|| \leq e^{-\lambda t}, \quad t \geq 0,\ p \in \mathcal{P}$$
2. The elements of $\mathcal{A}$ share a common quadratic Lyapunov function; i.e., there exists a positive definite matrix Q such that
$$QA_p + A_p'Q < 0, \quad \forall p \in \mathcal{P}$$
3. $\mathcal{P}$ {and therefore $\mathcal{A}$} are finite sets and $A_pA_q = A_qA_p$ for all $p, q \in \mathcal{P}$.
4. The matrices in $\mathcal{A}$ are row diagonally dominant with negative diagonal elements ; i.e., for $A = [a_{ij}] \in \mathcal{A}$,
$$2a_{ii} + \sum_{j=1}^{n} |a_{ij}| < 0,\ i \in \{1, 2, \ldots n\}$$
5. There exists an integer $\bar{n} \geq n$, a full rank $\bar{n} \times n$ matrix V and a family of $\bar{n} \times \bar{n}$ matrices $\bar{\mathcal{A}} = \{\bar{A}_p : p \in \mathcal{P}\}$ such that
   (a) $\bar{A}_pV = VA_p,\ p \in \mathcal{P}$
   (b) the matrix $\bar{A}_\sigma$ is exponentially stable for each $\sigma \in \mathcal{S}_1$

The sufficiency of conditions 1 and 2 are more or less obvious; the same is true of condition 3 since it implies that the associated matrix exponentials of the A_p commute. An explicit construction is given in [49] of a matrix Q satisfying condition 2 for a set of matrices $\mathcal{A}$ satisfying condition 3. Recently condition 1 was shown to be necessary for (57) to hold for all $\sigma \in \mathcal{S}_1$ [50]; of course one is still faced with the problem of deciding when such a norm exists. Condition 4 can be easily established using an simple estimate based on the idea of a matrix

measure {c.f., [51], p. 47}. Condition 5 is a consequence of the fact that for any solution x to $\dot{x} = A_\sigma x$, $\bar{x} \stackrel{\Delta}{=} Vx$ is a solution to $\dot{\bar{x}} = \bar{A}_\sigma \bar{x}$

Note that conditions 4 and 5 imply that the following is sufficient for (57) to hold for all $\sigma \in \mathcal{S}_1$

Property: *There exists an integer $\bar{n} \geq n$, a full rank $\bar{n} \times n$ matrix V and a family of $\bar{n} \times \bar{n}$ matrices $\bar{\mathcal{A}} = \{\bar{A}_p : p \in \mathcal{P}\}$ such that*

1. $\bar{A}_p V = V A_p, \quad p \in \mathcal{P}$
2. *each matrix $\bar{A}_\sigma$ is row diagonally dominant with negative diagonal elements.*

Under certain conditions this property proves to be necessary as well [52]; of course one still needs to figure out when V and the $\bar{A}_p$ exist. For some results concerning the stability of switched *nonlinear* systems, see [44].

The kind of stability questions we've been discussing are very closely related to the problem of deciding when for a given class of matrices $\{M_p : p \in \mathcal{P}\}$ and a given class $\bar{\mathcal{S}}$ of infinite sequences mapping the nonnegative integers into $\mathcal{P}$,

$$\lim_{i \to \infty} \prod_{j=1}^{i} M_{\bar{\sigma}(j)} = 0, \quad \forall \bar{\sigma} \in \bar{\mathcal{S}}$$

This and related questions have been addressed in [53] and [54].

5.2 Other Questions

Here are several other questions involving switched linear systems:

- Given a family of switching signals $\mathcal{S}$ and a family of detectable pairs $\{(C_p, A_p) : p \in \mathcal{P}\}$ when does there exist a matrix function $p \longmapsto K_p$ on $\mathcal{P}$ for which $A_\sigma + K_\sigma C_\sigma$ is 'exponentially stable' for each $\sigma \in \mathcal{S}$?
- With reference to the preceding suppose $\mathcal{K}$ is a nonempty class of matrix functions $p \longmapsto K_p$ on $\mathcal{P}$ with the property that for each function $K_p \in \mathcal{K}$, $A_\sigma + K_\sigma C_\sigma$ is exponentially stable for each $\sigma \in \mathcal{S}$. For a given class of appropriately sized matrices $\{F_p : p \in \mathcal{P}\}$ compute {or at least tightly estimate} the sup over $\mathcal{S}$ of the inf over $\mathcal{K}$ of the induced $\mathcal{L}^2[0, \infty)$ norm of the linear operator

  $$y \longmapsto \int_0^t F_{\sigma(t)} \phi(t, s) K_{\sigma(s)} y(s) ds$$

  where ϕ is the state transition matrix of $A_\sigma + K_\sigma C_\sigma$.
- Given a class of linear systems $\{(A_p, B_p, C_p, D_p) : p \in \mathcal{P}\}$, each with property $\mathbf{P}$, and a class of switching signals $\mathcal{S}$, when is it true that for each $\sigma \in \mathcal{S}$ the switched system $(A_\sigma, B_\sigma, C_\sigma, D_\sigma)$ also has property $\mathbf{P}$? Interesting choices for $\mathbf{P}$ include *stability*, *stabilizability*, *passivity*. Some findings related to the last of these appear in [55].

6 Concluding Remarks

Switching logics such as those discussed in this paper are typically derived without appealing to any formal notion of state or state transition. Explaining such logics informally can be beneficial and we have sought to do this throughout the paper. On the other hand we've found that it is also worthwhile to go through the exercise of formally modeling such logics as hybrid systems with inputs, states, state transitions and outputs. Formal modeling can clarify an algorithm's functioning by reducing ambiguity and in so doing can help to preclude erroneous conclusions. For example, when cyclic switching was first devised without the benefit of a formal model, it was thought to be a time-varying system which it clearly is not. The formal modeling of the hysteresis switching logic in [12] made it easier to explain. There are no doubt many existing switching logics whose behaviors could more easily grasped if they were both explained informally and modelled formally as hybrid dynamical systems.

In §2 it was emphasized that multi-controller architectures can usually be efficiently implemented as state-shared parameter dependent controllers. There are of course situations when it is advantageous *not* to explicitly compute off-line the parameter-dependent coefficient matrices of such a controllers. This is especially true if the controller in question is of the certainty equivalence type and if the associated family of nominal process models is a continuum. For example, suppose that the set $\mathcal{N}$ of nominal transfer functions considered in §4.3 is a continuum; suppose in addition that for each $p \in \mathcal{P}$, controller transfer function κ_p is to be designed via LQ-theory applied to ν_p. Because solutions to matrix Riccati equations depend on the equation's coefficient matrices in a complicated {nonrational} manner, the dependence of κ_p on p will be at least as complex, even if ν_p depends linearly on p. What this means is that the problem of explicitly parameterizing $\mathcal{K}$ assuming an LQ-based controller design is hopelessly intractable. For MIMO process models the problem is far worse even for simple pole-placement designs. There are at least two ways to avoid this problem. One is to compute controller coefficient matrices in real time; this is feasible with a dwell-time switched supervisory control system provided the computations can be carried out quickly enough.

Another way to avoid the parameterization problem is to settle for a smaller nominal process model transfer function class $\bar{\mathcal{N}}$ containing only finitely many elements. In fact a strong case can be made for doing this not just to avoid the parameterization problem, but for other reasons as well. For example, finiteness of $\bar{\mathcal{N}}$'s parameter space $\bar{\mathcal{P}}$ can greatly simplify the required minimization $\Pi(W, p)$ over $\bar{\mathcal{P}}$ even if $\bar{\mathcal{P}}$ is very large. We refer the reader to [56] for an interesting discussion of how one would go about covering a process model transfer function class of the form

$$\mathcal{C}_P = \bigcup_{p \in \mathcal{P}} \{\nu_p + \delta : ||\delta||_\infty \leq \epsilon_p\}$$

assuming $\mathcal{N} \stackrel{\Delta}{=} \{\nu_p : \mathcal{P}\}$ is a compact continuum with a transfer function class

$\bar{\mathcal{C}}_P \supset \mathcal{C}_P$ of the form

$$\bar{\mathcal{C}}_P = \bigcup_{p \in \bar{\mathcal{P}}} \{\nu_p + \delta : ||\delta||_\infty \leq \bar{\epsilon}_p\}$$

where $\bar{\mathcal{P}}$ is a finite subset of $\mathcal{P}$ and $\epsilon_p \leq \bar{\epsilon}_p,\ p \in \bar{\mathcal{P}}$.

One of the underlying ideas exploited in section 4 is that modeling uncertainty can be dealt with by switching between the members of a family of *fixed gain* controllers. The idea has been around for a long time. For example there is an extensive literature on the "multiple-model" {i.e., multi-estimator} approach to uncertainty which goes back almost thirty years; see for example [57] and the many references therein. The key feature of the estimator-based approach discussed in §4.3 which distinguishes it from the classical multiple-model approach is that switching is orchestrated by a supervisor using a logic which selects controllers on the basis of normed output estimation errors. Surprisingly, a provably correct version of this simple idea does not seem to have found its way into the literature until quite recently [15] - this despite the fact that the idea is a natural extension of the original concept of hysteresis switching [11] which appeared in 1988.

For the case when $\mathcal{P}$ contains a continuum of points, the dwell-time switched, estimator-based supervisory control discussed in §4.3 can be thought of as a form of estimator-based parameter adaptive control in which the supervisor plays the role of parameter tuner [58]. In this context, the concept of a supervisor represents a significant departure from more traditional estimator-based tuning algorithms which typically employ recursive or dynamical parameter tuning. Most closely related to what we've been discussing seems to be the type of adaptive algorithm studied by Naik, Kumar and Ydstie in [59]. Both the NKY algorithm and the dwell-time switched supervisor discussed in §4.3 search on compact parameter spaces; both are inherently robust to unmodelled dynamics in that dynamic normalization [60] is not employed. Perhaps the most significant differences between the two are 1. that the NKY algorithm employs recursive parameter tuning {i.e., pseudo-gradient/projection search} whereas the dwell-time switched supervisor does not and 2. the dwell-time switched supervisor allows time for computation whereas the NKY algorithm does not.

It is reasonable to suspect that many of the ideas covered in §4.3 can be successfully applied to specific classes of nonlinear systems. The well-known obstacles to generalization imposed by a limited nonlinear observer theory can almost certainly be side-stepped by focusing attention on systems whose states can be measured. It is quite likely that a supervised family consisting of a finite number of fixed nonlinear controllers will prove easier to analyze than a continuously parameterized family of controllers under the control of a parameter tuner. Understanding such switched systems calls for new methods of analysis which go beyond the partial Lyapunov function approach commonly used in the study of conventional parameter adaptive systems.

Simulations

The general concept of a dwell-time switched, estimator-based supervisory control system described in this paper has been tested in simulation under a variety of conditions. The reader wishing to experiment with these simulations can obtain Matlab Simulink files via the internet at the address

http://www.cis.yale.edu/~wchang/workshop.html

The simulations were designed and implemented by Wen-Chung Chang and João Hespanha.

References

1. V. I. Utkin. *Sliding Modes in Control Optimization*. Springer-Verlag, 1992.
2. M. I. Zelikin and V. F. Borisov. *Theory of Chattering Control*. Birkhauser, 1994.
3. G. F. Franklin, J. D. Powell, and M. L. Workman. *Digital control of Dynamic Systems*. Addison-Wesley, 1990.
4. C. R. Walli and I. S. Reed. Asynchronous finite-state machines - a novel control system class. *AIAA Journal*, 7(3):385–393, March 1969.
5. C. L. Smith. *Digital Computer Process Control*. Intext Educational Publishers, 1972.
6. R. A. Hilhorst. *Supervisory Control of Mode-Switch Processes*. PhD thesis, University of Twente, 1992.
7. B. Mårtensson. The order of any stabilizing regulator is sufficient a priori information for adaptive stabilization. *Systems and Control Letters*, 6(2):87–91, 1985.
8. M. Fu and B. R. Barmish. Adaptive stabilization of linear systems via switching controls. *IEEE Transactions on Automatic Control*, pages 1079–1103, December 1986.
9. S. J. Cusumano and Poola. Adaptive control of uncertain systems: A new approach. In *Proceedings of the American Automatic Control Conference*, pages 355–359, June 1988.
10. D. E. Miller. *Adaptive Control of Uncertain Systems*. PhD thesis, University of Toronto, 1989.
11. R. H. Middleton, G. C. Goodwin, D. J. Hill, and D. Q. Mayne. Design issues in adaptive control. *IEEE Transactions on Automatic Control*, 33(1):50–58, January 1988.
12. A. S. Morse, D. Q. Mayne and G. C. Goodwin. Applications of hysteresis switching in parameter adaptive control. *IEEE Transactions on Automatic Control*, 37(9):1343–1354, September 1992.
13. S. R. Weller and G. C. Goodwin. Hysteresis switching adaptive control of linear multivariable systems. *IEEE Transactions on Automatic Control*, 39(7):1369–1375, July 1994.
14. F. M. Pait and A. S. Morse. A cyclic switching strategy for parameter-adaptive control. *IEEE Transactions on Automatic Control*, 39(6):1172–1183, June 1994.
15. A. S. Morse. Dwell-time switching. In *Proceedings of the 2nd European Control Conference*, pages 176–181, 1993.
16. K. S. Narendra and J. Balakrishnan. Improving transient response of adaptive control systems using multiple models and switching. Technical report, Yale University, October 1992.

17. R. W. Brockett. Hybrid models for motion control systems. In H. L. Ttentleman and J. C. Willems, editors, *Essays on Control: Perspectives in the Theory and its Applications*, pages 29–54. Birkhäuser, July 1993.
18. M. S. Branicky, V. S. Borkar, and S. K. Mitter. A unified framework for hybrid control: Background, model, and theory. In *Proceedings of the 33rd Conference on Decision and Control*, pages 4228–4234. Control Systems Society, IEEE, December 1994.
19. R. L. Grossman, A. Nerode, A. P. Ravn, and H. Rischel, editors. *Hybrid Systems*, volume 736 of *Lecture Notes in Computer Science*. Springer-Verlag, 1993.
20. K. T. Tan. *Maximal Output Admissible Sets and the Nonlinear Control of Linear Discrete-Time Systems with State and Control Constraints*. PhD thesis, University of Michigan, 1991.
21. J. P. Hespanha. Stabilization of the nonholonomic integrator via logic-based switching. Technical report, Yale, 1995.
22. A. S. Morse and F. M. Pait. Mimo design models and internal regulators for cyclicly-switched parameter-adaptive control systems. *IEEE Transactions on Automatic Control*, 39(9):1809–1818, September 1994.
23. K. S. Narendra and J. Balakrishnan. Adaptive control using multiple models and switching. Technical report, Yale University, July 1994.
24. A. S. Morse. Supervisory control of families of linear set-point controllers - part 1: Exact matching. Technical report, Yale University, Mar 1993. submitted for publication.
25. A. R. Teel. Global stabilization and tracking for multiple integrators with bounded controls. *System and Control Letters*, pages 165–171, January 1992.
26. H. J. Sussmann, E. D. Sontag, and Y. Yang. A general result on the stabilization of linear systems using bounded controls. *IEEE Transactions on Automatic Control*, pages 2411–2425, December 1994.
27. E. G. Gilbert, I. Kolmanovsky, and K. T. Tan. Discrete-time reference governors and the nonlinear control of systems with state and control constraints. *Journal of Robust control and Nonlinear Systems*, August 1995. to appear.
28. P. Kapasouris, M. Athans, and G. Stein. Design of feedback control systems for stable plants with saturating actuators. In *Proc. 27th IEEE Conference on Decision and Control*, pages 469–479, December 1988.
29. G. F. Wredenhagen and P. R. Bélanger. Piecewise-linear lq control for systems with input constraints. *Automatica*, pages 403–416, March 1994.
30. E. G. Gilbert and K. T. Tan. Linear systems with state and control constraints: The theory and application of maximal output admissible sets. *IEEE Transactions on Automatic control*, pages 1008–1020, September 1991.
31. R. W. Brockett. Asymptotic stability and feedback stabilization. In R. Brockett, R. Millman, and H. Sussmann, editors, *Differential Geometric Control theory*, pages 181–191, 1983.
32. R. W. Brockett. Pattern generation and feedback control of nonholonomic systems. In *Proceedings of the IEEE Workshop on Mechanics, Holonomy and Control*, 1993.
33. O.J. Sødalen and O. Egeland. Exponential stabilization of nonholonomic chained systems. *IEEE Transactions on Automatic Control*, 40(1):35–49, 1995.
34. A. F. Filippov. Differential equations with discontinuous right-hand side. *Amer. Math. Soc. Translations*, pages 199–231, 1964.
35. J. Guckenheimer. A robust hybrid stabilization strategy for equilibria. *IEEE Transactions on Automatic Control*, 40(2):321–325, February 1995.

36. A. S. Morse. Towards a unified theory of parameter adaptive control - tunability. *IEEE Transactions on Automatic Control*, 35(9):1002–1012, September 1990.
37. D. E. Miller and E. J. Davison. An adaptive controller which provides an arbitrary good transient and steady-state response. *IEEE Transactions on Automatic Control*, 36(1):68–81, January 1991.
38. D. E. Miller. Adaptive stabilization using a nonlinear time-varying controller. *IEEE Transactions on Automatic Control*, 39(7):1347–1359, July 1994.
39. B. Mårtensson and J. W. Polderman. Correction and simplification to "the order of any stabilizing regulator is sufficient a priori information for adaptive stabilization". *Systems and Control Letters*, 20:465–470, 1993.
40. S. R. Kulkarni and P. J. Ramadge. Model and controller selection policies based on output prediction errors. Technical report, Princeton University, 1995.
41. S. Boyd, L. E. Ghaoui, E. Feron, and V. Balakrishnan. *Linear matrix Inequalities in Systems and Control Theory*. SIAM, 1994.
42. K. S. Narendra and J. Balakrishnan. Intelligent control using fixed and adaptive models. Technical report, Yale University, October 1994.
43. A. S. Morse. Supervisory control of families of linear set-point controllers - part 2: Robustness. Technical report, Yale University, Nov 1994. submitted for publication.
44. M. S. Branicky. Stability of switched and hybrid systems. In *Proceedings of the 33rd Conference on Decision and Control*, pages 3498–3503. Control Systems Society, IEEE, December 1994.
45. F. M. Pait. *Achieving Tunability in Parameter Adaptive Control*. PhD thesis, Yale University, 1993.
46. A. S. Morse. Logic-based switching and control. In B. A. Francis and A. R. Tannenbaun, editors, *Feedback Control, Nonlinear Systems, and Complexity*, pages 173–195. Springer-Verlag, 1995.
47. L. Praly and B. E. Ydstie. *Adaptive Linear Control: Bounded Solutions and Their Properties*. to be published, $\approx$ 1995.
48. M. A. Wicks, P. Peleties, and R. A. Decarlo. Construction of piecewise lyapunov functions for stabilizing switched systems. In *Proceedings of the 1994 IEEE Conference on Decision and Control*, pages 3492–3497, December 1994.
49. K. S. Narendra and J. Balakrishnan. A common lyapunov function for stable lti systems with commuting a-matrices. *IEEE Trans. Auto. Control*, pages 2469–2470, December 1994.
50. L. Gurvits. Untitled manuscript. in preparation, January 1995.
51. M. Vidyasagar. *Nonlinear Systems Analysis*. Prentice-Hall, 1993.
52. A. P. Molchanov and Ye.S. Pyatnitskiy. Criteria of asymptotic stability of differential and difference inclusions encountered in control theory. *Systems & Control Letters*, pages 59–64, 1989.
53. Ingrid Daubechies and J. C. Lagarias. Sets of matrices all infinite products of which converge. *Linear Algebra and Its Applications*, pages 227–263, 1992.
54. L. Gurvits. Stability of discrete linear inclusion. *Linear Algebra and Its Applications*, 1995. to appear.
55. W. A. Sethares, B. D. O. Anderson, and Jr. C. R. Johnson. Adaptive algorithms with filtered regressor and filtered error. *Mathematics of control, Signals, and Systems*, pages 381–403, February 1989.
56. F. M. Pait. On the topologies of spaces on linear dynamical systems commonly employed as models for adaptive and robust control design. In *Proceedings of the Third SIAM Conference on Control and Its Applications*, 1995.

57. P. S. Maybeck and D. L. Pogoda. Multiple model adaptive controller for the stol f-15 with sensor/actuator failures. In *Proceedings of the 28th IEEE Conference on Decision and Control*, pages 1566–1572. IEEE, 1989.
58. A. S. Morse. Towards a unified theory of parameter adaptive control - part 2: Certainty equivalence and implicit tuning. *IEEE Transactions on Automatic Control*, 37(1):15–29, January 1992.
59. S. M. Naik, P. R. Kumar, and B. E. Ydstie. Robust continuous-time adaptive control by parameter projection. *IEEE Transactions on Automatic Control*, 37(2):182–197, February 1992.
60. L. Praly. Global stability of direct adaptive control schemes with respect to a group topology. In *Adaptive And Learning Systems*, pages 5–16, February 1986.

This article was processed using the LaTeX macro package with LMAMULT style