Hierarchical hysteresis switching*

Daniel Liberzon† João P. Hespanha‡ A. Stephen Morse†

†Dept. of Electrical Engineering, Yale University, New Haven, CT 06520, {liberzon,morse}@cs.yale.edu

‡Dept. of Electr. Eng.–Systems, Univ. of Southern California, Los Angeles, CA 90089, hespanha@usc.edu

Abstract

We describe a new switching logic, called "hierarchical hysteresis switching", and establish a bound on the number of switchings produced by this logic on a given interval. The motivating application is the problem of controlling a linear system with large modeling uncertainty. We consider a control algorithm consisting of a finite family of linear controllers supervised by the hierarchical hysteresis switching logic. In this context, the bound on the number of switchings enables us to prove stability of the closed-loop system in the presence of noise, disturbances, and unmodeled dynamics.

1 Introduction

Suppose that a given process admits a model that contains unknown parameters, and the goal is to design a feedback controller that achieves some desired behavior in the face of noise, disturbances, and unmodeled dynamics. The kind of control algorithm that we have in mind is the one that relies on switching among a family of candidate controllers, and bases controller selection on certainty equivalence. In this framework, one associates to each possible value of the unknown parameters a *monitoring signal*, designed in such a way that a small value of this signal indicates a high likelihood that the corresponding parameters are close to the actual unknown values. The switching algorithm selects, from time to time, a controller that has been designed for the parameter values associated with the smallest monitoring signal.

Questions that suggest themselves are: How to pick the individual controllers? How to design the monitoring signals? Which controller to switch to? When to switch? In addressing the first question, the methods considered here allow one to rely on conventional techniques from linear robust control theory. An answer to the second question emerges from the *supervisory control* architecture, described in [7, 11, 5, 10] and reviewed later in this paper. The certainty equivalence principle provides one way to settle the third question. To deal with the last question, one needs to specify a *switching logic*. It is this last question that is of primary concern here.

A simple switching logic, called *hysteresis switching*, was described in [6, 9]. According to this logic, a switch occurs when the monitoring signal that corresponds to the controller currently in the feedback loop exceeds the smallest monitoring signal by a prespecified positive number, called the *hysteresis constant*. For a finite family of monitoring signals satisfying suitable assumptions, this logic guarantees that the switching stops in finite time. Hysteresis switching allows one to design supervisory control algorithms which are effective in those cases when the unknown parameters take values in a finite set and there are no noise, disturbances, or unmodeled dynamics.

An altogether different way to orchestrate the switching is provided by the *dwell-time switching logic*. In this logic, consecutive switching instants are separated by (at least) a prespecified time interval, called the *dwell time*, which is large enough so that the switching does not destabilize the system; this idea is ubiquitous in the switching control literature. Dwell-time switching was used in [7, 8] to design set-point supervisory control algorithms for linear systems with a continuum of parametric uncertainty, noise, disturbances, and unmodeled dynamics. These results go far beyond what can be established using the hysteresis switching logic.

However, dwell-time switching has its own disadvantages. First, the analysis that is needed to verify the correctness of the algorithms given in [7, 8] is quite tedious. More importantly, if the uncertain process is nonlinear, the existence of a prescribed dwell time may lead to trajectories escaping to infinity in finite time. These considerations motivate further study of hysteresis-based switching algorithms, which are easier to analyze and more suitable for control of nonlinear systems.

The results obtained in [7, 8] relied in part on the fact that the dwell-time switching logic is *scale-independent*, in the sense that its output does not change if all the monitoring signals are multiplied by a positive function of time. The hysteresis switching logic discussed in [6, 9] does not have this desirable property. However, it is not difficult to modify the logic by introducing a multiplicative hysteresis constant instead of an additive one. The resulting *scale-independent hysteresis switching logic* was studied and applied to control of uncertain nonlinear systems in [1] and elsewhere. These results still relied on the termination of switching in finite time, and were thus limited to situations where the parametric uncertainty range is described by a finite set and there are no noise, disturbances, or unmodeled dynamics.

*Research supported by AFOSR, DARPA, NSF, and ONR.

The scale-independent hysteresis switching logic was further studied in the recent paper [2], where a bound on the number of switches on a finite interval was established. Combining this with the results of [4], it is possible to analyze the correctness of the supervisory control algorithm without relying on the termination of switching, which allows a successful treatment of noise, disturbances, and unmodeled dynamics (but the unknown parameters are still restricted to belong to a finite set); see [2, 3]. In [3] a new switching logic was also introduced, called *local priority hysteresis switching*. It was designed primarily for the case when the unknown parameters belong to a continuum. As shown in [3], if this logic is used instead of scale-independent hysteresis, then in the absence of noise, disturbances, and unmodeled dynamics the switching stops in finite time, thus enabling one to generalize some of the previously available results.

The main contribution of the present paper is a new switching logic, which we call *hierarchical hysteresis switching*. It relies on a partition of the parametric uncertainty set (typically a continuum) into a finite number of subsets. The name of the logic comes from the fact that the minimization of the monitoring signals is carried out on two levels: first, the smallest one is taken in each of the subsets that form the partition, and then these signals are compared with each other. In the supervisory control context, the subsets in the partition are chosen to be sufficiently small in a suitable sense. We show that this switching logic leads to a supervisory control algorithm whose stability can be analyzed in the presence of noise, disturbances, and unmodeled dynamics. Thus we are able to handle the same class of systems as that treated in [7, 8], without sacrificing the valuable advantages of hysteresis-based switching algorithms.

The hierarchical hysteresis switching logic is presented in the next section. The supervisory control system is described in Section 3. Its analysis is given in Section 4.

2 Hierarchical hysteresis switching

We now describe the hierarchical hysteresis switching logic. Its inputs are some given continuous signals μ_p, $p \in \mathcal{P}$, where $\mathcal{P}$ is a compact index set. For m a positive integer, we will let $\mathbf{m}$ denote the set $\{1, 2, \ldots, m\}$. We assume that we are given a positive integer m and a family of closed subsets $D_i, i \in \mathbf{m}$ of $\mathcal{P}$, whose union is the entire $\mathcal{P}$. The output of the switching logic will be a *switching signal* σ taking values in $\mathbf{m}$. Pick a number $h > 0$, called the *hysteresis constant*. First, we select some $j_0 \in \mathbf{m}$ such that D_{j_0} contains $\arg\min_{p\in\mathcal{P}}\{\mu_p(0)\}$, and set $\sigma(0) = j_0$. Suppose that at a certain time t_i the value of σ has just switched to some $j_i \in \mathbf{m}$. We then keep σ fixed until a time $t_{i+1} > t_i$ such that the following inequality is satisfied:

$$(1+h)\min_{p\in\mathcal{P}}\{\mu_p(t_{i+1})\} \leq \min_{p\in D_{j_i}}\{\mu_p(t_{i+1})\}.$$

At this point, we select some $j_{i+1} \in \mathbf{m}$ such that D_{j+1} contains $\arg\min_{p\in\mathcal{P}}\{\mu_p(t_{i+1})\}$, and set $\sigma(0) = j_{i+1}$. When the indicated $\arg\min$ is not unique, a particular index among those that achieve the minimum can be chosen arbitrarily. We refer the reader to [7] for a discussion of tractability issues regarding minimization over $\mathcal{P}$. The understanding here is that minimization over D_i's is computationally tractable if these sets are sufficiently small.

The above procedure yields a piecewise constant signal σ which is continuous from the right everywhere. By the same argument as in [1], one can show that chattering is avoided if all μ_p, $p \in \mathcal{P}$ are bounded below by some positive number. In fact, there exists a maximal interval $[0, T_{\max})$ on which σ is defined, and there can only be a finite number of switches on each proper subinterval of $[0, T_{\max})$. In the supervisory control application treated below, we will always have $T_{\max} = \infty$.

Remark 1. The signal σ produced by this logic coincides with the signal that would be produced by the scale-independent hysteresis switching logic of [1] with inputs $\min_{p\in D_i}\{\mu_p(t)\}$, $i \in \mathbf{m}$. □

The above switching logic is scale-independent, i.e., its output would not be affected if we replaced the signals μ_p, $p \in \mathcal{P}$ by their scaled versions

$$\bar{\mu}_p(t) := \Theta(t)\mu_p(t), \qquad p \in \mathcal{P} \tag{1}$$

where Θ is some positive function of time. In the sequel, we assume that it is possible to choose Θ in such a way that the scaled signals $\bar{\mu}_p$, $p \in \mathcal{P}$ are strictly positive and monotone increasing. Scaled signals $\bar{\mu}_p$, $p \in \mathcal{P}$ with these properties will be useful for analysis purposes. For $0 \leq t_0 < t < T_{\max}$, we denote by $N_\sigma(t, t_0)$ the number of discontinuities of σ on the interval (t_0, t).

Definition. We will say that a piecewise constant signal ζ taking values in $\mathcal{P}$ is $\{D_i\}$-*consistent* with σ on an interval $[t_0, t]$ if:

1. For all $s \in [t_0, t]$ we have $\zeta(s) \in D_{\sigma(s)}$.

2. The set of discontinuities of ζ on $[t_0, t]$ is a subset of the set of discontinuities of σ.

Crucial properties of the switching signal produced by the hierarchical hysteresis switching logic are expressed by the following result.

Lemma 1 (Hierarchical Hysteresis Switching Lemma) *Take an arbitrary index $l \in \mathcal{P}$ and arbitrary numbers t_0 and t satisfying $0 \leq t_0 < t < T_{\max}$. We have*

$$N_\sigma(t, t_0) \leq 1 + m + \frac{m}{\log(1+h)} \log\left(\frac{\bar{\mu}_l(t)}{\min_{p\in\mathcal{P}} \bar{\mu}_p(t_0)}\right). \tag{2}$$

In addition, there exists a signal ζ, which is $\{D_i\}$-

consistent with σ on $[0,t]$, such that

$$\sum_{k=0}^{N_\zeta(t,t_0)} \Big(\bar\mu_{\zeta(t_k)}(t_{k+1}) - \bar\mu_{\zeta(t_k)}(t_k)\Big) \le m\Big((1+h)\bar\mu_l(t) - \min_{p\in\mathcal{P}} \bar\mu_p(t_0)\Big) \quad (3)$$

where $t_1 < t_2 < \cdots < t_{N_\zeta(t,t_0)}$ are the discontinuities of ζ on (t_0,t) and $t_{N_\zeta(t,t_0)+1} := t$.

PROOF. The inequality (2) follows at once from the Scale-Independent Hysteresis Switching Theorem (Theorem 1) of [2] and Remark 1. A signal ζ that satisfies the second statement of the lemma can be defined as follows: for each $s \in [t_0,t]$, let $\zeta(s) := \arg\min_{p\in D_{\sigma(s)}}\{\bar\mu_p(\tau)\}$, where τ is the largest number in the interval $[t_0,t]$ for which $\sigma(\tau) = \sigma(s)$. Then ζ is $\{D_i\}$-consistent with σ on $[t_0,t]$ by construction. Grouping all the terms in the summation on the left-hand side of (3) for which σ is the same, and reasoning exactly as in the proof of Theorem 1 in [2], we arrive at (3). □

Remark 2. The signal ζ depends on the choice of the time t. If the signals $\bar\mu_p$, $p\in\mathcal{P}$ are differentiable, then the left-hand side of the inequality (3) equals the integral $\int_{t_0}^t \dot{\bar\mu}_{\zeta(\tau)}(\tau)d\tau$, which is to be interpreted as the sum of integrals over intervals on which ζ is constant. □

3 Supervisory control system

We assume that the uncertain process $\mathbb{P}$ to be controlled admits the model of a SISO finite-dimensional stabilizable and observable linear system with control input u and measured output y, perturbed by a bounded disturbance input d and a bounded output noise signal n. It is assumed known that the transfer function of $\mathbb{P}$ from u to y belongs to a family of admissible process model transfer functions $\bigcup_{p\in\mathcal{P}}\mathcal{F}(p)$, where p is a parameter taking values in some index set $\mathcal{P}$. Here each $\mathcal{F}(p)$ denotes a family of transfer functions "centered" around some known *nominal* process model transfer function ν_p. Throughout the paper, we will take $\mathcal{P}$ to be a compact subset of a finite-dimensional normed linear vector space.

The problem of interest is to design a feedback controller that achieves output regulation, i.e., drives the output y of $\mathbb{P}$ to zero, whenever the noise and disturbance signals are zero. Moreover, all system signals must remain bounded in response to arbitrary bounded noise and disturbance inputs. Everything that follows can be readily extended to the more general problem of set-point control with the help of adding an integrator in the feedback loop, as in [7, 8].

The set $\mathcal{P}$ represents the range of parametric uncertainty, while for each fixed $p\in\mathcal{P}$ the subfamily $\mathcal{F}(p)$ accounts for unmodeled dynamics. There are several ways to specify allowable unmodeled dynamics around the nominal process model transfer functions ν_p (see [3]). For example, take two arbitrary numbers $\delta > 0$ and $\lambda_u \ge 0$. Then for each $p\in\mathcal{P}$ we can define

$$\mathcal{F}(p) := \{\nu_p(1+\delta_p^m)+\delta_p^a : \|\delta_p^m\|_{\infty,\lambda_u} \le \delta, \|\delta_p^a\|_{\infty,\lambda_u} \le \delta\}$$

where $\|\cdot\|_{\infty,\lambda_u}$ denotes the $e^{\lambda_u t}$-weighted $\mathcal{H}_\infty$ norm of a transfer function: $\|\nu\|_{\infty,\lambda_u} = \sup_{\omega\in\mathbb{R}}|\nu(j\omega-\lambda_u)|$. This yields the class of admissible process models treated in [7, 8]. In the sequel, we assume for concreteness that unmodeled dynamics are specified in this way; we will refer to the positive parameter δ as the *unmodeled dynamics bound*.

Modeling uncertainty of the kind described above may be associated with unpredictable changes in operating environment, component failure, or various external influences. Typically, no single controller is capable of solving the regulation problem for the entire family of admissible process models. Therefore, one needs to develop a controller whose dynamics can change on the basis of available real-time data. Within the framework of supervisory control discussed here, this task is carried out by a "high-level" controller, called a *supervisor*, whose purpose is to orchestrate the switching among a parameterized family of *candidate controllers* $\{\mathbb{C}_q : q\in\mathcal{Q}\}$, where $\mathcal{Q}$ is an index set. We require this controller family to be sufficiently rich so that every admissible process model can be stabilized by placing in the feedback loop the controller $\mathbb{C}_q$ for some index $q\in\mathcal{Q}$. In this paper, we focus on the case when $\mathcal{Q} = \mathbf{m}$ for some positive integer m.

The supervisor consists of three subsystems (Fig. 1):

multi-estimator $\mathbb{E}$ – a dynamical system whose inputs are the output y and the input u of the process $\mathbb{P}$ and whose outputs are the signals y_p, $p\in\mathcal{P}$. Each y_p would converge to y asymptotically if the transfer function of $\mathbb{P}$ were equal to the nominal process model transfer function ν_p and there were no noise or disturbances.

monitoring signal generator $\mathbb{M}$ – a dynamical system whose inputs are the *estimation errors*

$$e_p := y_p - y, \qquad p\in\mathcal{P}$$

and whose outputs μ_p, $p\in\mathcal{P}$ are suitably defined integral norms of the estimation errors, called *monitoring signals*.

switching logic $\mathbb{S}$ – a switched system whose inputs are the monitoring signals μ_p, $p\in\mathcal{P}$ and whose output is a *switching signal* σ, taking values in $\mathbf{m}$, which is used to define the control law u.

We now briefly recall from [7] the key state-space equations for the different subsystems appearing in Fig. 1. As i ranges over $\mathbf{m}$, let

$$\dot x_{\mathbb{C}} = A_i x_{\mathbb{C}} + b_i y$$
$$u = \bar c_i x_{\mathbb{C}} + \bar d_i y$$

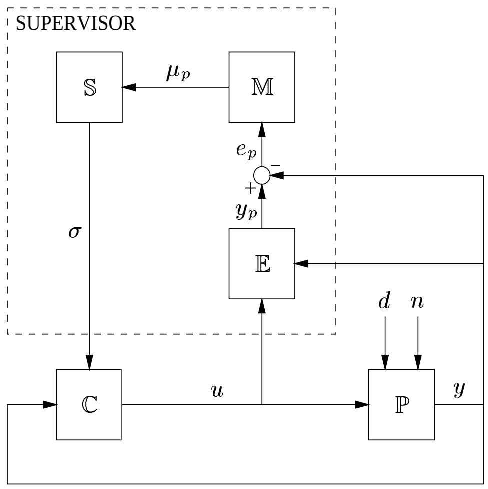

Figure 1: Supervisory control architecture

be realizations of the transfer functions of the candidate controllers, all sharing the same state $x_{\mathbb{C}}$. See [7] for more details on constructing such realizations. We then define the *multi-controller* $\mathbb{C}$ to be the system

$$\begin{aligned} \dot{x}_{\mathbb{C}} &= A_\sigma x_{\mathbb{C}} + b_\sigma y \\ u &= \bar{c}_\sigma x_{\mathbb{C}} + \bar{d}_\sigma y \end{aligned}$$

We assume that the multi-estimator is also realized in a state-shared fashion, as given by

$$\begin{aligned} \dot{x}_{\mathbb{E}} &= A_{\mathbb{E}} x_{\mathbb{E}} + b_{\mathbb{E}} y + d_{\mathbb{E}} u \\ y_p &= c_p x_{\mathbb{E}}\,, \qquad p \in \mathcal{P} \end{aligned}$$

with $A_{\mathbb{E}}$ a stable matrix. This type of structure is quite common in adaptive control. Denote by x the composite state $(x_{\mathbb{E}}', x_{\mathbb{C}}')'$ of the multi-estimator and the multi-controller, and let p^* be an (unknown) element of $\mathcal{P}$ such that the transfer function of $\mathbb{P}$ belongs to $\mathcal{F}(p^*)$, i.e., a "true" parameter value. For every $l \in \mathcal{P}$, the evolution of x can be described by a system of the form

$$\dot{x} = A_{\sigma l} x + d_\sigma e_l \tag{4}$$

$$y = (c_{p^*}\ \ 0)\, x - e_{p^*} \tag{5}$$

$$u = f_\sigma x + g_\sigma e_{p^*} \tag{6}$$

We assume that a partition $\mathcal{P} = \bigcup_{i \in \mathbf{m}} D_i$ is given, such that the matrices A_{ip}, $i \in \mathbf{m}$, $p \in \mathcal{P}$ have the following property: for every $i \in \mathbf{m}$ and every $p \in D_i$ the matrix $A_{ip} + \lambda_0 I$ is stable, where λ_0 is a fixed positive number. This property is a direct consequence of the construction described in [7], provided that the sets D_i, $i \in \mathbf{m}$ are chosen to be sufficiently small.

Fix a number $\lambda \in (0, \min\{\lambda_u, \lambda_0\})$. As shown in [7, 8], there exist positive constants δ_1, δ_2 that only depend on the unmodeled dynamics bound δ and go to zero as δ goes to zero, positive constants B_1, B_2 that only depend on the noise and disturbance bounds and go to zero as these bounds go to zero, and positive constants C_1, C_2 that only depend on the system's parameters and on initial conditions, such that along all solutions of the closed-loop system we have

$$\int_0^t e^{2\lambda\tau} e_{p^*}^2(\tau)d\tau \le B_1 e^{2\lambda t} + C_1 + \delta_1 \int_0^t e^{2\lambda\tau} u^2(\tau)d\tau \tag{7}$$

and

$$|e_{p^*}(t)| \le B_2 + C_2 e^{-\lambda t} + \delta_2 e^{-\lambda t} \sqrt{\int_0^t e^{2\lambda\tau} u^2(\tau)d\tau}. \tag{8}$$

The above inequalities represent the basic requirements being placed on the multi-controller and the multi-estimator, upon which the subsequent analysis depends.

The constant λ will play the role of a "weighting" design parameter in the definition of the monitoring signals. Fix an arbitrary constant $\epsilon_\mu \ge 0$ (its role will become clear later). We generate the monitoring signals μ_p, $p \in \mathcal{P}$ by the equations

$$\begin{aligned} \dot{W} &= -2\lambda W + \begin{pmatrix} x_{\mathbb{E}} \\ y \end{pmatrix} \begin{pmatrix} x_{\mathbb{E}} \\ y \end{pmatrix}', \qquad W(0) \ge 0 \\ \mu_p &:= \begin{pmatrix} c_p & -1 \end{pmatrix} W \begin{pmatrix} c_p & -1 \end{pmatrix}' + \epsilon_\mu, \qquad p \in \mathcal{P} \end{aligned} \tag{9}$$

where $W(t)$ is a symmetric nonnegative-definite $k \times k$ matrix, $k := \dim(x_{\mathbb{E}}) + 1$. Since $c_p x_{\mathbb{E}} - y = e_p\ \forall p \in \mathcal{P}$, this yields

$$\mu_p(t) = e^{-2\lambda t}\tilde{\mu}_p(0) + \int_0^t e^{-2\lambda(t-\tau)} e_p^2(\tau)d\tau + \epsilon_\mu, \quad p \in \mathcal{P}$$

where $\tilde{\mu}_p(0) := \begin{pmatrix} c_p & -1 \end{pmatrix} W(0) \begin{pmatrix} c_p & -1 \end{pmatrix}'$.

Finally, we define the switching signal using the hierarchical hysteresis switching logic described in Section 2, where the sets D_i, $i \in \mathbf{m}$ are chosen as explained earlier. Setting $\Theta(t) := e^{2\lambda t}$ in (1), we see that the signals $\bar{\mu}_p$, $p \in \mathcal{P}$ are monotone increasing, because they satisfy

$$\bar{\mu}_p(t) = \tilde{\mu}_p(0) + \int_0^t e^{2\lambda\tau} e_p^2(\tau)d\tau + \epsilon_\mu e^{2\lambda t}, \quad p \in \mathcal{P}. \tag{10}$$

Moreover, it is easy to ensure that $\mu_p(0) > 0\ \forall p \in \mathcal{P}$, either by setting $\epsilon_\mu > 0$ or by requiring $W(0)$ to be positive definite. Therefore, we can apply Lemma 1 and conclude that the inequalities (2) and (3) are valid. Since in this case the signals $\bar{\mu}_p$, $p \in \mathcal{P}$ are differentiable, the left-hand side of the inequality (3) equals $\int_{t_0}^t \dot{\bar{\mu}}_{\zeta(\tau)}(\tau)d\tau$ (see Remark 2). From (10) we have the following formula:

$$\dot{\bar{\mu}}_p(t) = e^{2\lambda t} e_p^2(t) + 2\lambda\epsilon_\mu e^{2\lambda t}, \qquad p \in \mathcal{P}. \tag{11}$$

4 Analysis

We now proceed to the analysis of the supervisory control system defined by (4), (5), (6), (9), and the hierarchical switching logic. We will sometimes appeal to the state of the uncertain process $\mathbb{P}$, which we denote by $x_{\mathbb{P}}$. Following [4], we will say that a switching signal σ has an *average dwell time* $\tau_{AD} > 0$ if there exists a nonnegative number N_0 such that the number of discontinuities of σ on an arbitrary interval (t_0, t) satisfies

$$N_\sigma(t, t_0) \leq N_0 + \frac{t - t_0}{\tau_{AD}}. \tag{12}$$

We will need the following result, which in view of the present assumptions is a straightforward corollary of the main result of [4]. It states that if σ has a sufficiently large average dwell time, then the switched system

$$\dot{x} = A_{\sigma\zeta} x \tag{13}$$

is exponentially stable with stability margin λ, uniformly over all signals ζ that are $\{D_i\}$-consistent with σ.

Lemma 2 *There exist positive constants τ^* and c such that for every switching signal σ with an average dwell time $\tau_{AD} \geq \tau^*$ and every signal ζ which is $\{D_i\}$-consistent with σ on a given interval $[t_0, t]$, solutions of* (13) *satisfy $|x(t)| \leq ce^{-\lambda(t-t_0)}|x(t_0)|$.*

Let τ^* be the number specified by this lemma; it can be calculated explicitly from the proof of the main result in [4]. Consider the system obtained from (4) by substituting ζ for l:

$$\dot{x} = A_{\sigma\zeta} x + d_\sigma e_\zeta. \tag{14}$$

An immediate corollary of Lemma 2 is that this system has a finite $e^{\lambda t}$-weighted $\mathcal{L}_2$-to-$\mathcal{L}_\infty$ induced norm, uniform over σ and ζ.

Corollary 3 *There exist positive constants g and g_0 such that for every $t > 0$, every switching signal σ with an average dwell time $\tau_{AD} \geq \tau^*$, and every signal ζ which is $\{D_i\}$-consistent with σ on $[0, t]$, solutions of* (14) *satisfy*

$$e^{2\lambda t}|x(t)|^2 \leq g \int_0^t e^{2\lambda\tau} e^2_{\zeta(\tau)}(\tau) d\tau + g_0|x(0)|^2. \tag{15}$$

With these results in place, the analysis is similar to that given in [3]; some details will be omitted.

No noise, disturbances, or unmodeled dynamics

We first consider the simple situation where there are no unmodeled dynamics ($\delta = 0$), i.e., the process $\mathbb{P}$ exactly matches one of the nominal process models, and where the noise and disturbance signals are zero ($n = d \equiv 0$). In this case, the constants $B_1, B_2, \delta_1, \delta_2$ in (7) and (8) are all zero. Let us take ϵ_μ in the definition of the monitoring signals to be zero as well ($W(0)$ must then be taken positive definite). The inequality (7) gives $\int_0^t e^{2\lambda\tau} e^2_{p^*}(\tau) d\tau \leq C_1$, which together with (10) implies $\bar{\mu}_{p^*} \leq \tilde{\mu}_{p^*}(0) + C_1$. It follows from (2), applied with $l = p^*$, that $N_\sigma(t, t_0)$ is bounded by a fixed constant for arbitrary $t > t_0 \geq 0$. This means that the switching stops in finite time, i.e., there exist a time T^* and an index $i^* \in \mathbf{m}$ such that $\sigma(t) = i^*$ for $t \geq T^*$. In this case (12) holds for every τ_{AD} if N_0 is large enough. Fix an arbitrary $t > 0$. In view of Lemma 1, Remark 2, and the formula (11), there exists a signal ζ which is $\{D_i\}$-consistent with σ on $[0, t]$ and satisfies $\int_0^t e^{2\lambda\tau} e^2_{\zeta(\tau)}(\tau) d\tau \leq m(1 + h)(\tilde{\mu}_{p^*}(0) + C_1)$. Using (15), we have $e^{2\lambda t}|x(t)|^2 \leq gm(1+h)(\tilde{\mu}_{p^*}(0)+C_1)+g_0|x(0)|^2$, thus $x \to 0$. Since $e_{p^*} \to 0$ by virtue of (8), we conclude from (5) that $y \to 0$. Therefore, the output regulation problem is solved. In light of (6), (9), and detectability of $\mathbb{P}$, all the other signals remain bounded for all $t \geq 0$. We summarize this as follows.

Proposition 4 *Suppose that the noise and disturbance signals are zero and there are no unmodeled dynamics, and set $\epsilon_\mu = 0$. Then all the signals in the supervisory control system remain bounded for every set of initial conditions such that $W(0) > 0$. Moreover, the switching stops in finite time, and we have $y(t) \to 0$ as $t \to \infty$.*

Remark 3. Since the evolution of $x_{\mathbb{P}}$ and x for $t \geq T^*$ is described by a linear time-invariant system, the rate of convergence in the above statement is actually exponential. □

Noise and disturbances, no unmodeled dynamics

We now assume that bounded noise n and disturbance d are present, but there are no unmodeled dynamics. In this case the switching typically will not stop in finite time. The inequalities (7) and (8) hold with some unknown but finite constants B_1, B_2. The parameters δ_1 and δ_2 are still zero, and C_1 and C_2 are positive constants as before. We take ϵ_μ to be a positive number. From (7) and (10) we have

$$\bar{\mu}_{p^*}(t) \leq \tilde{\mu}_{p^*}(0) + B_1 e^{2\lambda t} + C_1 + \epsilon_\mu e^{2\lambda t} \tag{16}$$

The formula (2), applied with $l = p^*$, yields

$$N_\sigma(t, t_0) \leq N_0 + \frac{t - t_0}{\tau_{AD}}$$

where

$$\tau_{AD} = \frac{\log(1 + h)}{2\lambda m}$$

and

$$N_0 = 1 + m + \frac{m}{\log(1 + h)} \log\left(\frac{\tilde{\mu}_{p^*}(0) + B_1 + C_1 + \epsilon_\mu}{\epsilon_\mu}\right).$$

We can guarantee that $\tau_{AD} \geq \tau^*$ by increasing the hysteresis constant h and/or decreasing the weighting constant λ if necessary. In the sequel, we assume that h and λ have been chosen in this way.

Using (3), (11), and (16), we obtain

$$\int_0^t e^{2\lambda\tau} e^2_{\zeta(\tau)}(\tau)d\tau \le m\big((1+h)(\tilde{\mu}_{p^*}(0) + B_1 e^{2\lambda t} + C_1 + \epsilon_\mu e^{2\lambda t}) - \epsilon_\mu e^{2\lambda t_0}\big)$$

where $t > 0$ is arbitrary and ζ is the signal provided by Lemma 1. Together with (15) this implies that

$$|x(t)|^2 \le \big(gm(1+h)(\tilde{\mu}_{p^*}(0) + C_1) + g_0|x(0)|^2\big)e^{-2\lambda t} + gm(1+h)(B_1 + \epsilon_\mu)$$

Two conclusions can be drawn from the last formula. First, x is bounded, and as in the previous subsection we can easily deduce from (6), (8), (9), and detectability of $\mathbb{P}$ that all system signals remain bounded. Note that the choice of the design parameters λ, h and ϵ_μ did not depend on the noise or disturbance bounds, in other words, explicit knowledge of these bounds is not necessary (we are merely requiring that such bounds exist). Secondly, if n and d equal or converge to zero, then x will approach a neighborhood of the origin whose size is proportional to $g\epsilon_\mu$. A close examination of the last quantity reveals that it decreases to 0 as ϵ_μ goes to 0, which means that we can make this neighborhood as small as desired by choosing ϵ_μ sufficiently small. Moreover, e_{p^*} will converge to zero because of (8), hence y will also become arbitrarily small in view of (5). We arrive at the following result.

Proposition 5 *Suppose that the noise and disturbance signals are bounded and there are no unmodeled dynamics. Then for an arbitrary $\epsilon_\mu > 0$ all the signals in the supervisory control system remain bounded for every set of initial conditions. Moreover, for every number $\epsilon_y > 0$ there is a value of ϵ_μ leading to the property that if the noise and disturbance signals converge to zero, then for each solution there is a time $\bar{T}$ such that $|y(t)| \le \epsilon_y$ for all $t \ge \bar{T}$.*

Remark 4. We cannot simply let $\epsilon_\mu = 0$, as this would invalidate the above analysis even if $W(0) > 0$. However, by decreasing ϵ_μ on-line (e.g., in a piecewise constant fashion), it is possible to recover asymptotic convergence of y to zero when the noise and disturbance signals converge to zero. □

Noise, disturbances, and unmodeled dynamics

If unmodeled dynamics are present, i.e., if the parameter δ is positive, then δ_1 and δ_2 in (7) and (8) are also positive. In this case, the analysis becomes more complicated, because we can no longer deduce from (7) that the switched system must possess an average dwell time. However, it is possible to prove that the above control algorithm, without any modification, is robust with respect to unmodeled dynamics in the following, "semi-global", sense. The proof uses a small-gain argument, and is almost identical to the proof of Theorem 4 in [3].

Theorem 6 *For arbitrary bounds on the noise and disturbance signals, the supervisory control system has the following properties:*

1. *For every positive value of ϵ_μ and every number $E > 0$ there exists a number $\bar{\delta} > 0$ such that if the unmodeled dynamics bound δ is smaller than $\bar{\delta}$, then all signals remain bounded for every set of initial conditions such that $|x_{\mathbb{P}}(0)|, |x(0)| \le E$.*
2. *For arbitrary positive numbers E and ϵ_y there exist a value of ϵ_μ and a number $\bar{\delta} > 0$ such that if the noise and disturbance signals converge to zero and the unmodeled dynamics bound δ is smaller than $\bar{\delta}$, then for each solution with $|x_{\mathbb{P}}(0)|, |x(0)| \le E$ there is a time $\bar{T}$ such that $|y(t)| \le \epsilon_y$ for all $t \ge \bar{T}$.*

References

[1] J. P. Hespanha, Logic-based switching algorithms in control, Ph.D. Thesis, Dept. of Electrical Engineering, Yale University, 1998.

[2] J. P. Hespanha, D. Liberzon, A. S. Morse, Bounds on the number of switchings with scale-independent hysteresis: applications to supervisory control, in *Proc. 39th Conf. on Decision and Control*, 2000, to appear.

[3] J. P. Hespanha, D. Liberzon, A. S. Morse, B. D. O. Anderson, T. S. Brinsmead, F. De Bruyne, Multiple model adaptive control, part 2: Switching, submitted to *Int. J. Robust Nonlinear Control*, 2000.

[4] J. P. Hespanha, A. S. Morse, Stability of switched systems with average dwell-time, in *Proc. 38th Conf. on Decision and Control*, 1999, pp. 2655–2660.

[5] J. Hocherman-Frommer, S. R. Kulkarni, P. J. Ramadge, Controller switching based on output prediction errors, *IEEE Trans. Automat. Control*, vol. 43, 1998, pp. 596–607.

[6] R. H. Middleton, G. C. Goodwin, D. J. Hill, D. Q. Mayne, Design issues in adaptive control, *IEEE Trans. Automat. Control*, vol. 33, 1988, pp. 50–58.

[7] A. S. Morse, Supervisory control of families of linear set-point controllers, part 1: Exact matching, *IEEE Trans. Automat. Control*, vol. 41, 1996, pp. 1413–1431.

[8] A. S. Morse, Supervisory control of families of linear set-point controllers, part 2: Robustness, *IEEE Trans. Automat. Control*, vol. 42, 1997, pp. 1500–1515.

[9] A. S. Morse, D. Q. Mayne, G. C. Goodwin, Applications of hysteresis switching in parameter adaptive control, *IEEE Trans. Automat. Control*, vol. 37, 1992, pp. 1343–1354.

[10] E. Mosca, F. Capecchi, Designing predictors for MIMO switching supervisory control, *Int. J. Robust and Nonlinear Control*, to appear.

[11] F. M. Pait, F. Kassab Jr., Parallel algorithms for adaptive control: Robust stability, in *Control Using Logic-Based Switching*, (A. S. Morse, ed.), Springer-Verlag, London, 1997.