Overcoming the limitations of adaptive control by means of logic-based switching

João P. Hespanha[*] Daniel Liberzon[†] A. Stephen Morse[‡]

Abstract

In this paper we describe a framework for deterministic adaptive control which involves logic-based switching among a family of candidate controllers. We compare it with more conventional adaptive control techniques that rely on continuous tuning, emphasizing how switching and logic can be used to overcome some of the limitations of traditional adaptive control. The issues are discussed in a tutorial, non-technical manner and illustrated with specific examples.

1 Introduction

Adaptive control is a methodology for controlling systems with large modeling uncertainties which render robust control design tools inapplicable and thus require adaptation. By adaptation we usually mean a combination of on-line estimation and control, whereby a suitable controller is selected on the basis of the current estimate for the uncertain process. More precisely, one chooses a parameterized family of controllers, where the parameter varies over a continuum which corresponds to the process uncertainty range in a suitable way. One then runs an estimation procedure, which at each instant of time provides an estimate of the unknown process model. According to certainty equivalence, one applies a controller that is known to guarantee some desired behavior of the process model corresponding to the current estimate (see, e.g., [20]).

This classical approach to deterministic adaptive control has some inherent limitations which have been well recognized in the literature. Most notably, if unknown parameters enter the process model in complicated ways, it may be very difficult to construct a continuously parameterized family of candidate controllers. Estimation over a continuum may also be a challenging task. These issues become especially severe if robustness and high performance are sought. As a result, design of adaptive control algorithms involves a large number of specialized techniques and often depends on trial and error.

In this paper we discuss an alternative approach to control of uncertain systems, which seeks to overcome some of the above difficulties while retaining the fundamental ideas on which adaptive control is based. The main feature which distinguishes it from conventional adaptive control is that controller selection is carried out by means of logic-based switching rather than continuous tuning. Switching among candidate controllers is orchestrated by a high-level decision maker called a *supervisor*, hence the name *supervisory control*. The supervisor updates controller parameters when a new estimate of the process parameters becomes available, similarly to the adaptive control paradigm, but these events occur at discrete instants of time. This results in a *hybrid* closed-loop system. The idea of using switching in an adaptive context

[*]Dept. of Electr. and Computer Eng., Univ. of California, Santa Barbara, CA 93106. Email: hespanha@usc.edu. Supported by NSF Grant ECS-0093762.

[†]Corresponding author. Coordinated Science Lab., Univ. of Illinois at Urbana-Champaign, Urbana, IL 61801. Email: liberzon@uiuc.edu. Fax: (217) 244-2352. Supported by NSF Grants ECS-0134115 and ECS-0114725.

[‡]Dept. of Electrical Engineering, Yale University, New Haven, CT 06520. Email: morse@cs.yale.edu. Supported by NSF and DARPA.

has been around for some time, and various approaches have been developed; see, e.g., [29, 11, 31, 22]. In the form considered here, supervisory control originates in [36, 38] and is pursued in [25, 43, 4, 19, 17, 14, 27, 1, 16, 40, 6, 44, 15].

In the supervisory control framework, it is not necessary to construct a continuously parameterized family of controllers, since the controller selection is performed in a discrete fashion. This allows one to handle process models that are nonlinearly parameterized over non-convex sets and also to use advanced controllers that are not readily parameterized continuously. If the unknown process model parameters belong to a discrete set, there is no need to embed this set in a continuum (as one would normally do in the context of adaptive control); instead, one can switch among a discrete family of suitably chosen controllers. If the parametric uncertainty is described by a continuum, one has the choice of working with a continuous or a discrete (perhaps even finite) family of controllers. In the latter case, one needs to ensure that every admissible process model is satisfactorily controlled by at least one of these controllers; under appropriate conditions, this can always be achieved with a finite controller family (see [1]).

Another important aspect of supervisory control is modularity. The principles that govern the design of the switching logic, the estimators, and the candidate controllers are mutually independent. In other words, the analysis of the overall system relies on certain basic properties of its individual parts, but not on the particular ways in which these parts are implemented. As a result, one gains the advantage of being able to use "off-the-shelf" control laws, rather than having to design control laws tailored to the specifics of the continuously-tuned adaptive algorithms. This provides greater flexibility in applications (where there is often pressure to utilize existing control structures) and facilitates the use of advanced controllers for difficult problems. Similar remarks apply to the estimation procedure. We will support these claims with examples below.

As is well known, one reason for considering logic-based switching control, and hybrid control in general, is that it enables one to overcome obstructions that are present in systems with continuous controllers. For example, there are quite general classes of systems, including nonholonomic systems, which cannot be stabilized by continuous feedback because they fail to satisfy Brockett's necessary condition [5]. Hybrid control laws provide one way of dealing with this problem. In view of this more general observation, it is perhaps not surprising that supervisory control is capable of overcoming limitations that are characteristic of conventional adaptive control algorithms.

The switching algorithms that seem to be the most promising are those that evaluate on-line the potential performance of each candidate controller and use this to direct their search. These algorithms can roughly be divided into two categories: those based on process estimation, using either certainty equivalence (see, e.g., [36, 41, 25, 43]) or model validation [23, 8, 49]; and those based on direct performance evaluation of each candidate controller [47, 46, 24, 40]. Although these algorithms originate from fundamentally different approaches, they share key structures and exhibit important common properties. In this paper, we mostly address estimator-based supervisory control based on certainty equivalence. However, many of the comments to be made also apply to other forms of supervisory control [13]. We will not address here supervisory control based on a sequential, or "pre-routed", search among a set of controllers. We will also restrict our attention to deterministic algorithms that do not require persistency of excitation. Further discussion and references on the different approaches mentioned above can be found in [34].

2 Supervisory control system

In this section we describe, in general terms, the basic ingredients of a supervisory control system, and sketch the ideas involved in its analysis. Let $\mathbb{P}$ be the uncertain process to be controlled, with input u and output y, possibly perturbed by a bounded disturbance input d and a bounded output noise n. We assume that the model of $\mathbb{P}$ is a member of some family of admissible process models

$$\mathcal{F} := \bigcup_{p \in \mathcal{P}} \mathcal{F}_p \tag{1}$$

where $\mathcal{P}$ is an index set. Here, for each p, $\mathcal{F}_p$ denotes a family of systems "centered" around some known *nominal* process model ν_p. Loosely speaking, 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. Several ways of specifying admissible unmodeled dynamics around the nominal process models are discussed in [36, 16].

A standard problem of interest is output regulation or set-point control of $\mathbb{P}$. Typically, no single controller is capable of solving the problem if the range of modeling uncertainty is large. We therefore consider a parameterized family of *candidate controllers* $\{\mathbb{C}_q : q \in \mathcal{Q}\}$, where $\mathcal{Q}$ is an index set, and switch in real time between the elements of this family, on the basis of observed data. This leads to a switched controller, which we call the *multi-controller* and denote by $\mathbb{C}$. The set $\mathcal{Q}$ may be different from $\mathcal{P}$; for example, $\mathcal{Q}$ might be finite while $\mathcal{P}$ is a continuum (cf. below). The understanding here is that for each $p \in \mathcal{P}$ there is a $q \in \mathcal{Q}$ such that the control input u_q produced by the candidate controller $\mathbb{C}_q$ would yield the desired behavior if $\mathbb{P}$ were known to be a member of $\mathcal{F}_p$. If $\mathcal{Q}$ is a finite set with a relatively small number of elements, the multi-controller can be realized simply as a parallel connection of all the candidate controllers. If the number of elements in $\mathcal{Q}$ is large or infinite, the multi-controller can be implemented using the idea of *state-sharing* [36], leading to a dynamical system whose (finite) dimension is independent of the size of $\mathcal{Q}$.

The supervisor consists of three subsystems (see Figure 1):

multi-estimator $\mathbb{E}$ – a dynamical system whose inputs are the input u and the output y of the process $\mathbb{P}$ and whose outputs are denoted by y_p, $p \in \mathcal{P}$.

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[1] of the estimation errors, called *monitoring signals*.

switching logic $\mathbb{S}$ – a dynamical system whose inputs are the monitoring signals μ_p, $p \in \mathcal{P}$ and whose output is a piecewise constant *switching signal* σ, taking values in $\mathcal{Q}$, which is used to define the control law $u = u_\sigma$.

We now explain the basic requirements that need to be placed on the different parts of the supervisory control system. For simplicity, we first consider the case when $\mathcal{P} = \mathcal{Q}$, and later explain what modifications are needed to handle the general situation. Consider the switched system that describes the combined dynamics of the process, the multi-controller, and the multi-estimator. Denoting its state by x and ignoring the noise and disturbances, we can write this system as

$$\dot{x} = f_\sigma(x). \tag{2}$$

From the perspective of the remaining components of the system, the outputs of (2) are the estimation errors, which can be generated by equations of the form

$$e_p = h_p(x), \quad p \in \mathcal{P}. \tag{3}$$

[1]One can also process the estimation errors in a more sophisticated way (see, e.g., [25, 40, 2]).

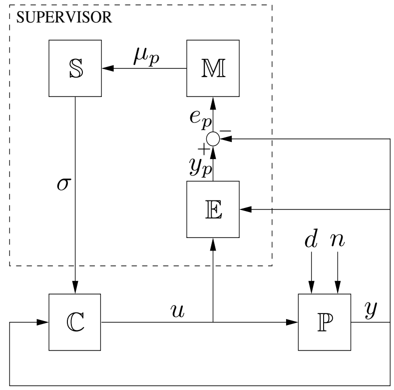

Figure 1: Supervisory control system.

The above switched system is required to have two basic properties, which are crucial for the analysis of the overall system: the Detectability Property and the Matching Property. The former is essentially a property of the multi-controller, whereas the latter is a property of the multi-estimator.

We begin with the multi-controller. The *Detectability Property* that we impose on the candidate controllers is that for every fixed $q \in \mathcal{Q}$, the switched system (2)–(3) must be detectable with respect to the corresponding estimation error e_q when the value of the switching signal is frozen at q. Adopted to the present context, the results proved in [33] imply that in the linear case the Detectability Property holds if the controller asymptotically stabilizes the multi-estimator and the process is detectable ("certainty equivalence stabilization theorem"), or if the controller asymptotically output-stabilizes the multi-estimator and the process is minimum-phase ("certainty equivalence *output* stabilization theorem"). These conditions are useful because they decouple the properties that need to be satisfied by the parts of the system constructed by the designer from the properties of the unknown process. Extensions of these results to nonlinear systems are discussed in [17, 15, 28]. In particular, it is shown in [17] that detectability, defined in a suitable way for nonlinear systems, is guaranteed if the process is detectable and the controller *input-to-state stabilizes* (in the sense of Sontag [48]) the multi-estimator with respect to the estimation error. The design of candidate controllers is thereby reduced to a disturbance attenuation problem well studied in the nonlinear control literature. The paper [15] develops an integral variant of this result, and the recent work [28] contains a nonlinear version of the certainty equivalence output stabilization theorem.

The *Matching Property* refers to the fact that the multi-estimator should be designed so that each particular y_p provides a "good" approximation to the process output y–and therefore e_p is "small"–whenever the actual process model is inside the corresponding $\mathcal{F}_p$. Since the process is assumed to match one of the models in the set (1), we should then expect at least one of the estimation errors, say e_{p^*}, to be small in some sense. For example, we may require that in the absence of unmodeled dynamics, noise, and disturbances e_{p^*} converge to zero exponentially fast for every control input u. It is also desirable to have an explicit characterization of e_{p^*} in the presence of unmodeled dynamics, noise, and disturbances. For linear systems, a multi-estimator satisfying such requirements can be obtained as explained in [36]. In that paper it is also shown how the multi-estimator can be constructed in a *state-shared* fashion (so that it is finite-dimensional even if $\mathcal{P}$ is infinite), using standard results from realization theory. Multi-estimators with similar properties can also be designed for some useful classes of nonlinear systems, as discussed in [17]. State-sharing is always possible if the parameters enter the process model "separably" (but not necessarily linearly).

The index σ of the controller in the feedback loop is determined by the switching logic, whose inputs are produced by the monitoring signal generator. In accordance with certainty equivalence, the value of σ at each instant of time should coincide with the index of the smallest monitoring signal. In this way, supervisory control selects the candidate controller that corresponds to an estimator error that has been small for some time (in an integral sense). To prevent chattering, one approximates this mechanism by introducing a *dwell-time* [36] or *hysteresis* [12, 17, 16, 27]. As discussed in these references, different versions of the latter approach can be applied to nonlinear systems with $\mathcal{P}$ finite or infinite and $\mathcal{Q}$ being equal to $\mathcal{P}$ or a finite subset of $\mathcal{P}$; an overview of various switching logics is given in [27]. Two properties need to be satisfied by the switching logic and the monitoring signal generator: the Non-Destabilization Property and the Small Error Gain Property.

Recall that, in view of the Detectability Property, for every *fixed* value of σ the system (2)–(3) is detectable with respect to the corresponding estimation error e_σ. The switching signal σ is said to have the *Non-Destabilization Property* if it preserves the detectability, i.e., if the switched system (2)–(3) is detectable with respect to the output e_σ. The Non-Destabilization Property trivially holds if the switching stops in finite time (which is the case if the scale-independent hysteresis switching logic of [12] or its variants proposed in [16, 27] are applied in the absence of noise, disturbances, and unmodeled dynamics). In the linear case, a standard output injection argument shows that detectability is not destroyed by switching if the switching is sufficiently slow (so as not to destabilize the injected switched system). According to the results of [18], it actually suffices to require that the switching be slow *on the average*. However, it should be noted that the Non-Destabilization Property does not necessarily amount to a slow switching condition; for example, the switching can be fast if the systems being switched are in some sense "close" to each other. For another fast switching result that exploits the structure of linear multi-controllers and multi-estimators, see [36, Section VIII].

The *Small Error Gain Property* calls for a bound on e_σ in terms of the smallest of the signals e_p, $p \in \mathcal{P}$. For example, if $\mathcal{P}$ is a finite set and the monitoring signals are defined as $\mu_p(t) = \int_0^t e_p^2(s)ds$, then the scale-independent hysteresis switching logic of [12] guarantees that for every $p \in \mathcal{P}$,

$$\int_0^t e_\sigma^2(s)ds \leq C \int_0^t e_p^2(s)ds \tag{4}$$

where C is a constant (which depends on the number of controllers and the hysteresis parameter) and the integral on the left is to be interpreted as the sum of integrals over intervals on which σ is constant. If e_{p^*} decays exponentially as discussed earlier, then (4) guarantees that the signal e_σ is in $\mathcal{L}_2$. At the heart of the switching logic, there is a conflict between the desire to switch to the smallest estimation error to satisfy the Small Error Gain Property and the concern that too much switching may violate the Non-Destabilization Property.

It is now easy to see how the above properties of the various blocks of the supervisory control system can be put together to analyze its behavior. Because of the Matching Property, there exists some $p^* \in \mathcal{P}$ for which e_{p^*} is small (e.g., converges to zero exponentially fast as above). The Small Error Gain Property implies that e_σ is small. The Detectability Property and the Non-Destabilization Property then guarantee that the state x is small as well. Proceeding in this fashion, it is possible to analyze stability and robustness of supervisory control algorithms for quite general classes of uncertain systems [36, 38, 18, 16].

Often, it is not convenient[2] to take the process index set $\mathcal{P}$ to be equal to the controller index set $\mathcal{Q}$. When these sets are different, we need to have a *controller assignment map* $\chi : \mathcal{P} \to \mathcal{Q}$. Let us say that a piecewise constant signal ζ taking values in $\mathcal{P}$ is *σ-consistent* if $\chi(\zeta) \equiv \sigma$ and the set of discontinuities of ζ is a subset of the set of discontinuities of σ. The Detectability Property and the Non-Destabilization Property need to be strengthened to guarantee that the switched system (2)–(3) is detectable with respect to the output e_ζ, for every σ-consistent signal ζ. The Small Error Gain Property can then be relaxed as follows: there must exist a σ-consistent signal ζ such that e_ζ is bounded in terms of the smallest of the signals e_p, $p \in \mathcal{P}$ (for example, in the sense of the $\mathcal{L}_2$ norm as before). The above analysis goes through with suitable minor modifications; see [27] for details.

[2] The reasons for this will be explained shortly and have precisely to do with overcoming limitations of adaptive control.

Not surprisingly, the four properties that were just introduced for supervisory control have direct counterparts in classical adaptive control. The Detectability Property was first recognized in the context of adaptive control in [32], where it was called *tunability*. The Matching Property is usually implicit in the derivation of the error model equations, where one assumes that, for a specific value of the parameter, the output estimate matches the true output. Both the Small Error Gain Property and the Non-Destabilization Property are pertinent to the tuning algorithms, being typically stated in terms of the smallness (most often in the $\mathcal{L}_2$ sense) of the estimation error and the derivative of the parameters estimate, respectively.

3 Modular design

One of the key features of supervisory control is that no specific structure is imposed on the multi-controller or on the multi-estimator, as long as they possess the properties described in the previous section. We proceed by presenting a few examples that utilize candidate controllers and estimators which—for different reasons—cannot be used by traditional forms of adaptive control. Specifically, Example 1 illustrates the use of advanced nonlinear controllers; Example 2 illustrates the use of hybrid controllers; Example 3 studies a process given by a PDE whose associated approximate nominal models are not parameterizable; and Example 4 illustrates the use of advanced estimators. It is interesting to note that the constructions proposed for all the examples are based on non-adaptive designs, and yet can be used in the context of supervisory control without any significant changes.

Example 1 Consider the two-dimensional system

$$\begin{aligned} \dot{x}_1 &= p_1^* x_1^3 + p_2^* x_2 \\ \dot{x}_2 &= u \end{aligned} \tag{5}$$

where u denotes the control input and $p^* := \{p_1^*,\ p_2^*\}$ is an uncertain parameter taking values in a subset $\mathcal{P}$ of $[-1,1] \times ([-1,1] \setminus \{0\})$. The state of (5) is assumed to be available for measurement, and the control objective is to drive x_1 to a known constant set point r. The system (5) is feedback linearizable for every possible value of $p \in \mathcal{P}$. However, for $p_1^* < 0$ the nonlinear term $p_1^* x_1^3$ actually provides desirable damping and should not be canceled. Feedback linearization should then be avoided for negative values of p_1^*. Another option is to use *pointwise min-norm* candidate control laws (see [10]). These are feedback laws of the smallest pointwise magnitude that yield the Detectability Property (as verified by a suitable Lyapunov function). Such control laws possess desirable robustness and performance properties associated with inverse optimality. Both the feedback linearization and the pointwise min-norm control designs can be applied in the supervisory control context.

Since in this example the entire state is available for measurement, it is not difficult to construct a multi-estimator with the property that the estimation error e_{p^*} converges to zero exponentially fast for every control input, so that the Matching Property is satisfied. Moreover, this can be done using state-sharing, so that the multi-estimator and the monitoring signal generator are finite-dimensional dynamical systems even if the parametric uncertainty set $\mathcal{P}$ is infinite. This issue will be addressed with more detail in the examples in the next section. The details of the design and analysis for this example are provided in [12]. Figures 2(a) and (b) show a simulation of the supervised system with the feedback linearizing and pointwise min-norm control laws, respectively, and with the scale-independent hysteresis switching logic. For the design of the controllers and estimators, the parameter set $\mathcal{P}$ was taken to be

$$\mathcal{P} := \{-1, -.9, -.8, \ldots, -.1, 0, .1, .2, \ldots, .9, 1\} \times \{-1, 1\}$$

To demonstrate the robustness of the closed loop system, in the simulations, the values of the actual parameters p_1^* and p_2^* are not exactly in $\mathcal{P}$. In the simulations shown in Figure 2, deviations of up to 1%

were allowed. Moreover, the values of p_1^* and p_2^* were changing with time. Comparing Figures 2(a) and (b), one observes that when $p_1^* < 0$, the pointwise min-norm control laws result in control signals about 10 times smaller than those produced by feedback linearizing control laws, without sacrificing the performance. □

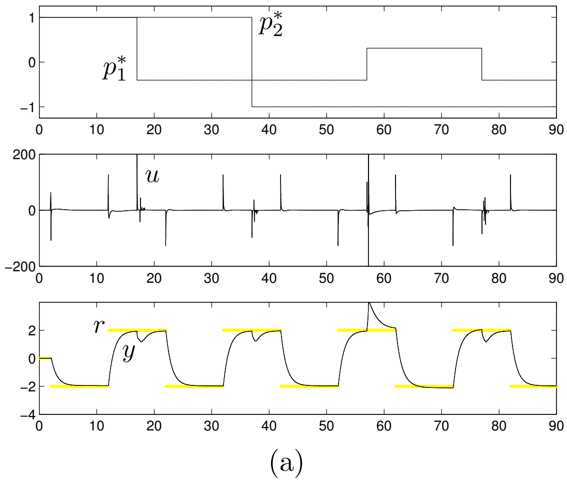

Figure 2: Supervisory control simulation for Example 1 with (a) feedback linearization candidate control laws and (b) pointwise min-norm candidate control laws.

Example 2 The next example addresses the problem of parking a wheeled mobile robot of the unicycle type, shown in Figure 3. Here x_1, x_2 are the coordinates of the point in the middle of the rear axle, and

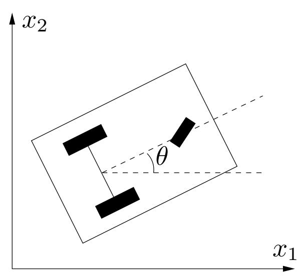

Figure 3: A wheeled mobile robot for Example 2.

θ denotes the angle that the vehicle makes with the x_1-axis. The kinematics of the robot can be modeled by the equations

$$\begin{aligned}\dot{x}_1 &= p_1^* u_1 \cos\theta \\ \dot{x}_2 &= p_1^* u_1 \sin\theta \\ \dot{\theta} &= p_2^* u_2\end{aligned}$$

where p_1^* and p_2^* are positive parameters determined by the radius of the rear wheels and the distance between them, and u_1 and u_2 are the control inputs (the forward and the angular velocity, respectively). The case we are interested in is when the actual values of p_1^* and p_2^* are unknown, so that the parameter vector $p^* = \{p_1^*, p_2^*\}$ belongs to some subset $\mathcal{P} = \mathcal{P}_1 \times \mathcal{P}_2$ of $(0, \infty) \times (0, \infty)$. Without loss of generality, by parking the vehicle we mean making x_1, x_2, and θ tend to zero by applying state feedback.

What makes this problem especially interesting is that for every set of values for the parameters, the corresponding system is nonholonomic and so cannot be stabilized by any continuous state feedback law.

Thus even in the non-adaptive situation the problem of parking the vehicle is a challenging one. This brings us outside the scope of most of the available adaptive control algorithms, despite the fact that the unknown parameters enter the model linearly[3]. When one considers hybrid control laws, this obstruction disappears. For the case of a known p^*, the system can be stabilized by a variety of hybrid control laws, such as the simple switching control law given in [14] which makes every trajectory approach the origin after at most one switch. In the presence of parametric modeling uncertainty, it is thus natural to proceed with the supervisory control design, using switching candidate controllers and a supervisor that orchestrates the switching among them. In this example, switching occurs at two levels with distinct purposes: at the lower level, each candidate controller utilizes switching to overcome the smooth nonstabilizability of the nonholonomic process, and at the higher level, the supervisor switches to deal with the process' parametric uncertainty. As shown in [14], it is indeed possible to solve the problem in this way, which illustrates once again the flexibility of the supervisory control in incorporating advanced control techniques developed in the non-adaptive control literature.

The full details of design and analysis are given in [14]. The switching logic that we used is a suitable version of the scale-independent hysteresis switching logic from [12]. If the set $\mathcal{P}$ is finite, all signals in the resulting supervisory control system remain bounded, the switching stops in finite time, and the state converges to zero. If $\mathcal{P}$ is infinite, the analysis breaks down because the switching can no longer be guaranteed to stop and stability results for the case of persistent switching in nonlinear systems are lacking. However, simulations indicate that the system still displays the desired behavior. A typical trajectory for the case $\mathcal{P} = [0.1, 2] \times [0.1, 3]$ is depicted in Figure 4. A parking movie generated with MATLAB/Simulink which illustrates the corresponding motion of the robot is available from the web [26]. □

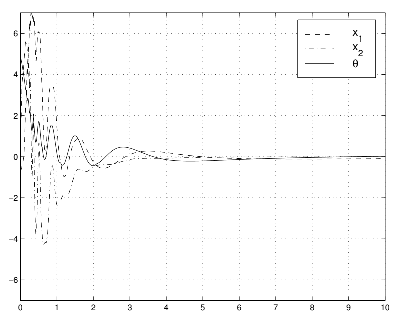

Figure 4: Parking the robot in Example 2.

Example 3 Consider the flexible manipulator shown in Figure 5. For small bending, this system can be modeled by the following PDE

$$\ddot{y}(x,t) + \frac{EI}{\rho} y''''(x,t) = -x\ddot{\theta}(t), \tag{6}$$

with boundary conditions

$$y(0,t) = y'(0,t) = 0, \qquad y''(L,t) = y'''(L,t) + \frac{m_t}{\rho} y''''(L,t) = 0, \qquad T(t) = I_H \ddot{\theta}(t) - EIy''(0,t),$$

where E denotes elasticity of the beam, I the inertia of a transversal slice, I_H the axis' inertia, L the beam's length, ρ the beam's mass density, and m_t the load mass at the tip of the manipulator. Four measurements are used for control: the base angle $\theta(t)$, the base angular velocity $\dot{\theta}(t)$, the tip position $y_{tip}(t) := L\theta(t) + y(L,t)$, and the bending $y''(x_{sg}, t)$ measured by a strain gauge attached to the flexible

[3]Notable exceptions include the work reported in [7, 21].

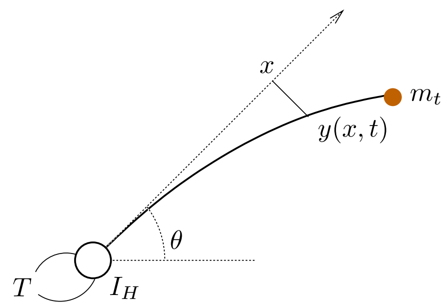

Figure 5: Flexible manipulator for Example 3.

bar at position x_{sg}. The actuator is a direct drive motor that applies the torque $T(t)$ at the base. The control objective is to drive the tip position to a pre-specified set-point r. Two parameters are assumed unknown: the value of the load mass m_t, assumed to be in the range $[0, .1\text{Kg}]$, and the exact position x_{sg} of the strain gauge, assumed to be in the range $[40\text{cm}, 60\text{cm}]$.

The PDE (6) can be solved by expanding the solution into the series $y(x,t) = \sum_{k=1}^{\infty} \phi_k(x) q_k(t)$, where the $\phi_k(x)$ are the eigenfunctions of the beam and the time-varying coefficients $q_k(t)$ are the solutions of an infinite-dimensional system of ODEs. This series is truncated for the design of the candidate controllers, resulting in a nominal finite-dimensional model $y(x,t) \approx \sum_{k=1}^{N} \phi_k(x) q_k(t)$. The ignored terms are treated as unmodeled dynamics. The main difficulty in designing an adaptive controller for (6) is that it is not possible to explicitly write a finite-dimensional nominal model in terms of the unknown parameters. With supervisory control this is not a problem, since we can work with a finite grid $\mathcal{P} \subset [0, .1\text{Kg}] \times [40\text{cm}, 60\text{cm}]$ and still cover the entire parameter space by considering appropriately large families $\mathcal{F}_p$ around each nominal process model. For each nominal model, one controller can then be designed using standard methods. Figure 6 shows a simulation of the closed loop system that utilized 18 candidate controllers designed using the LQR/LQE techniques. As in Example 1, the unknown parameters (m_t and x_{sg}) were taken to be time-varying. $\square$

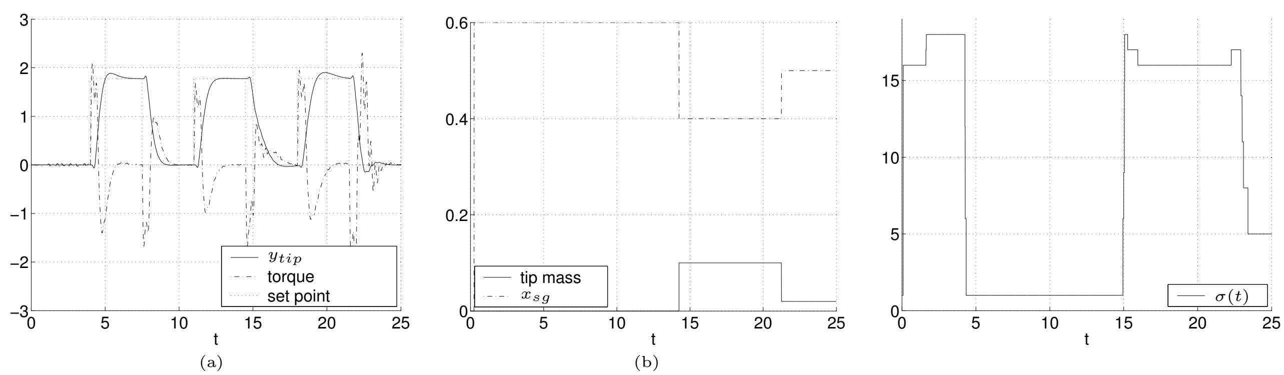

Figure 6: Supervisory control of a flexible manipulator in Example 3: (a) set point, tip position, and control signal; (b) tip mass and position of the strain gauge; (c) index of the candidate controller in use.

Example 4 The behavior of an induction motor in the so-called current-fed operation mode is described by a third-order model which expresses the rotor flux $\lambda \in \mathbb{R}^2$ and the stator currents $u \in \mathbb{R}^2$ in a reference frame rotating at the rotor angular speed ω as

$$\begin{aligned}\dot{\lambda} &= -R\lambda + Ru \\ \dot{\omega} &= \tau - \tau_L \\ \tau &= u' J\lambda\end{aligned} \tag{7}$$

where $\tau \in \mathbb{R}$ is the generated torque, R is the rotor resistance, taking values in an interval $[R_m, R_M]$, with R_m a positive number, τ_L is the load torque taking values in another interval $[\tau_{Lm}, \tau_{LM}]$, and $J := \left[\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right]$. The only measured output is ω.

Let $p := \{R, \tau_L\}$ be the vector of unknown parameters, taking values in $\mathcal{P} := [R_m, R_M] \times [\tau_{Lm}, \tau_{LM}]$. A multi-estimator for (7) can be constructed using the equations

$$\dot{\lambda}_R = -R\lambda_R + Ru, \qquad R \in [R_m, R_M] \tag{8}$$

$$\dot{\omega}_p = u^T J\lambda_R - \tau_L - (1 + |u|^2)(\omega_p - \omega), \qquad p = \{R, \tau_L\} \in \mathcal{P}. \tag{9}$$

With λ_R and ω_p so defined, one can show that if the actual value of the unknown parameter vector is $p^* := \{R^*, \tau_L^*\}$, then λ_{R^*} and ω_{p^*} converge exponentially fast to λ and ω, respectively, for every input signal u and thus the Matching Property holds. Note that (9) can be regarded as a linear time-varying system with inputs $u^T J\lambda_R + \kappa(1 + |u|^2)\omega$ and $-\tau_L$. Thus the same estimate ω_p could also be generated by the 2-dimensional system

$$\begin{aligned}\dot{\mu} &= -\kappa(1 + |u|^2)\mu + u^T J\lambda_R + \kappa(1 + |u|^2)\omega \\ \dot{\nu} &= -\kappa(1 + |u|^2)\nu - 1 \\ \omega_p &= \mu + \tau_L \nu, \qquad \tau_L \in [\tau_{Lm}, \tau_{LM}].\end{aligned}$$

We are therefore able to state-share the multi-estimator with respect to the unknown parameter τ_L, i.e., the above two-dimensional system handles the entire range $[\tau_{Lm}, \tau_{LM}]$. However, state-sharing with respect to R does not seem possible, and so instead of (8) we can only implement a finite number of equations corresponding to a suitable finite subset of the interval $[R_m, R_M]$.

Figure 7 shows a simulation where the above multi-estimator is used. This simulation uses off-the-shelf *field-oriented* candidate controllers. Field-oriented control is the *de-facto* industry standard for high performance application induction motors. In its indirect formulation, this is a nonlinear dynamic output feedback controller with a cascaded structure, where the inner loop control consists of a rotation and the outer loop is typically defined via a PI regulator around the velocity error. The reader is referred to [6] for the numerical values used for the parameters. In the simulation shown in Figure 7, we considered a situation where both the rotor resistance and the load torque change. The rotor resistance changes from its initial value $R = 6$ to $R = 8$ at $t = 60$ and the load torque changes from its initial value $\tau_L = 2$ to $\tau_L = 3$ at $t = 20$ and changes again to $\tau_L = 4$ at $t = 40$. One can observe that tracking as well as correct estimation of the rotor resistance and the load torque are achieved with only brief bursting at the instants when the parameters change. □

4 Admissible classes of process models

Another important feature of supervisory control is its flexibility in dealing with fairly general process model classes of the form (1). In this section we explain some of the limitations of traditional adaptive control in this area and illustrate how supervisory control can be used to overcome them.

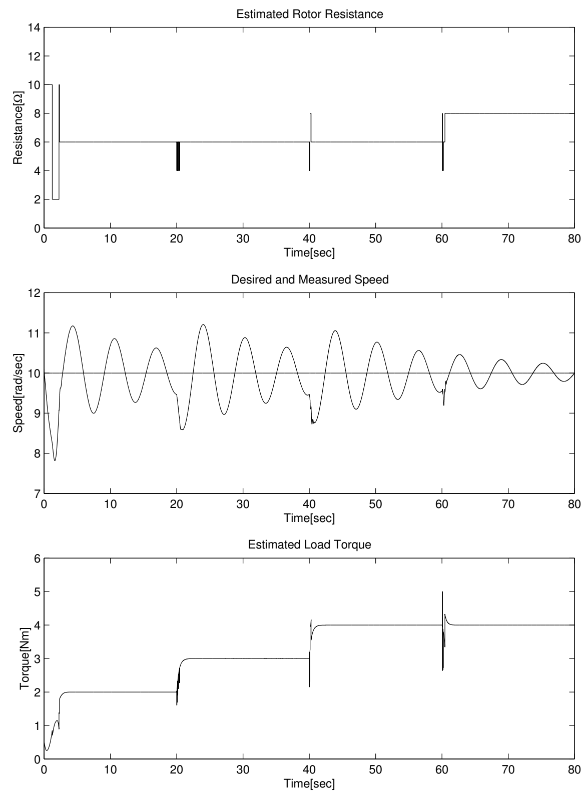

Figure 7: Simultaneous change of rotor resistance and load torque in Example 4.

4.1 Process model parameterizations

For simplicity, let us assume that for each $p \in \mathcal{P}$, the family $\mathcal{F}_p$ consists of a single nominal process model ν_p. This case is often called "exact matching" and corresponds to the absence of unmodeled dynamics.

In deterministic adaptive control, severe constraints are imposed on the parameterization map $p \mapsto \nu_p$. The reason for this is that continuous tuning based on gradient descent-like methods is very fragile with respect to this map. To gain insight into this issue, consider the following standard parameter estimation problem. Two signals $y : [0, \infty) \to \mathbb{R}^m$ and $z : [0, \infty) \to \mathbb{R}^k$ are related by the equation

$$y(t) = F(p^*, z(t)), \tag{10}$$

where p^* is an unknown parameter taking values in $\mathcal{P}$ and F is a function from $\mathcal{P} \times \mathbb{R}^k$ to $\mathbb{R}^m$. A standard way to estimate p^* is to solve the optimization problem

$$\min_{p \in \mathcal{P}} \int_0^\infty \|y(t) - F(p, z(t))\|^2 dt.$$

Adaptive control attempts to perform this minimization on-line by moving the current estimate $\hat{p}$ of p^* in the direction of steepest descent of $\|y(t) - F(\hat{p}, z(t))\|^2$ (for frozen time), namely,

$$\dot{\hat{p}} = -\frac{\partial}{\partial \hat{p}} \frac{\|y - F(\hat{p}, z)\|^2}{2} = (y - F(\hat{p}, z))' \frac{\partial F}{\partial \hat{p}}(\hat{p}, z). \tag{11}$$

This approach is validated by a Lyapunov-like argument based on the partial Lyapunov function

$$V(\hat{p}) := \frac{1}{2}\|p^* - \hat{p}\|^2.$$

Since p^* is assumed constant, along trajectories of (11) we have

$$\dot{V} = -(y - F(\hat{p}, z))' \frac{\partial F}{\partial \hat{p}}(\hat{p}, z)(p^* - \hat{p}).$$

If $F(\hat{p}, z)$ is linear in $\hat{p}$, we conclude from (10) that $y - F(\hat{p}, z) = \frac{\partial F}{\partial \hat{p}}(\hat{p}, z)(p^* - \hat{p})$ and therefore

$$\dot{V} = -\|y - F(\hat{p}, z)\|^2.$$

This means that $\hat{p}$ approaches p^* as long as y does not match $F(\hat{p}, z)$. However, the decrease of V along trajectories is very much dependent on the linearity of F with respect to p. Indeed, even if F depends on p in a very regular way, e.g., if it is quadratic (and hence convex) in p, it may happen that V no longer decreases. For example, if y, z, and p are scalars and

$$F(p, z) = (p + z)^2, \tag{12}$$

we obtain

$$\dot{V} = -2(\hat{p} + z)(\hat{p} + p^* + 2z)(p^* - \hat{p})^2$$

which is positive for $\hat{p}$ between $-z$ and $-p^* - 2z$.

In adaptive control, limitations are imposed not only on the form of the parameterization map $p \mapsto \nu_p$, but also on the shape of its domain $\mathcal{P}$. The set $\mathcal{P}$ is invariably required to be convex (and is usually taken to be a ball or a hypercube). The reason for this can also be seen from the above discussion. If the set $\mathcal{P}$ is non-convex, the parameter estimate can get trapped in a local minimum, at which the steepest descent vector points outside $\mathcal{P}$. The problem of how to address non-convex parameter sets has been a major stumbling block in adaptive control. Although it is uncommon to find physical parameters (such as masses, damping coefficients, etc.) that take values in a non-convex set, the need to make the parameterization $p \mapsto \nu_p$ linear often leads to reparameterizations that artificially introduce non-convexity. We will say more about reparameterizations later.

In supervisory control, we have a much greater flexibility in selecting a parameter estimate $\hat{p}$, because we are no longer forced to continuously tune this variable and can instead directly choose for $\hat{p}$, e.g., the value that minimizes some integral norm of $e_{\hat{p}} := y - F(\hat{p}, z)$. In the context of the above parameter estimation problem, this corresponds to setting, from time to time,

$$\hat{p}(t) = \arg\min_{p \in \mathcal{P}} \int_0^t \|y(t) - F(p, z(t))\|^2 dt. \tag{13}$$

When the minimum is achieved at multiple points, any one of them can be used for $\hat{p}$. Note that (13) can produce very rapidly varying and possibly discontinuous estimates. In the context of certainty equivalence, this could lead to very rapid changes in the controller and even chattering. To prevent this, one introduces dwell-time or hysteresis (cf. Section 2).

The integrals on the right-hand side of (13) correspond to the monitoring signals μ_p, so the above expression can be rewritten as

$$\hat{p}(t) = \arg\min_{p \in \mathcal{P}} \mu_p(t). \tag{14}$$

Thus it is possible to relax the assumption that the family of process models be linearly parameterized over a convex set, which is required by standard adaptive control techniques[4]. Computational issues associated with the on-line optimization (13)–(14) still need attention; they can be handled in one of the following ways.

[4]We note that in the context of continuously-tuned adaptive control, there have also been research efforts directed at removing the linearity assumption; see, e.g., [3].

1. In some cases, a closed-form solution can be found. This typically occurs when the unknown parameters enter the nominal process model in a polynomial fashion. For example, in the case of (12) it is straightforward to show that the optimization simply amounts to computing the roots of a cubic polynomial. Similarly, in the context of Example 2 one needs to find the roots of a polynomial of degree 5. See also Example 5 below where we provide some supporting calculations.

2. Numerical optimization can be performed when a closed-form solution is not available. In this case, convexity of the parameter set is usually needed. In general, the computational time required by the optimization will determine the maximum rate at which switching can occur. This can be taken directly into account with the dwell-time switching logic (see [36]).

3. The set $\mathcal{P}$ can be partitioned into a finite number of subsets, so that on each subset the optimization problem is convex and therefore computationally tractable. (This was in fact one of the motivations for the introduction of the hysteresis switching logic in [30], which was later used in supervisory control.) The *hierarchical* hysteresis switching logic developed in [27] relies on such a partition. The minimization is carried out on two levels: first, the smallest monitoring signal is taken in each of the subsets that form the partition, and then these signals are compared with each other. The work on combining switching and tuning, i.e., switching among adaptive controllers, is also relevant in this regard (see, e.g., [41]).

4. In many cases, one can take the parameter set $\mathcal{P}$ to be finite and absorb the remaining parameter values into unmodeled dynamics, so that the optimization problem is reduced to comparing a finite number of monitoring signals (see [1] for a detailed discussion of this issue). This additional flexibility is an advantage that logic-based switching algorithms have over continuously-tuned adaptive control techniques.

It should be emphasized that in supervisory control, computational issues of the kind mentioned above only arise with respect to the parameterization of the nominal process models $p \mapsto \nu_p$. The parameterization of the candidate controllers $q \mapsto \mathbb{C}_q$, on the other hand, does not require any special care.

Example 5 Following [36], we consider the family of SISO linear time-invariant processes with transfer functions

$$\nu_p := \frac{s - \frac{1}{6}(p+2)}{s^3 + ps^2 - \frac{2}{9}p(p+2)s}, \qquad p \in \mathcal{P} := [-1, 1]. \tag{15}$$

It is possible to construct a state-shared multi-estimator of the form

$$\begin{aligned} \dot{z} &= \begin{bmatrix} A & 0 \\ 0 & A \end{bmatrix} z + \begin{bmatrix} b \\ 0 \end{bmatrix} y + \begin{bmatrix} 0 \\ b \end{bmatrix} u \\ y_p &= \begin{bmatrix} a_1 & \left(a_2 + \frac{2}{9}p(p+2)\right) & (a_3 - p) & \left(\frac{1}{6}(p+2)\right) & 1 & 0 \end{bmatrix}, \qquad p \in \mathcal{P} \end{aligned} \tag{16}$$

where $A \in \mathbb{R}^{3\times 3}$, $b \in \mathbb{R}^{3\times 1}$ are in the controllable canonical form and A is stable with characteristic polynomial $s^3 + a_3 s^2 + a_2 s + a_1$. The nonlinear term $p(p+2)$, which propagated from the original parameterization of ν_p to (16), prevents the use of conventional adaptive control techniques. However, supervisory control has no difficulties in dealing with this example. As in most linear problems, we can use quadratic monitoring signals μ_p, $p \in \mathcal{P}$ with an exponential forgetting factor, namely,

$$\mu_p(t) = \int_o^t e^{-2\lambda(t-\tau)} \|e_p(\tau)\|^2 d\tau, \qquad \lambda > 0. \tag{17}$$

Straightforward calculations show that these signals can be realized as

$$\mu_p(t) = w_5(t)p^4 + w_4(t)p^3 + w_3(t)p^2 + w_2(t)p + w_1(t),$$

where the w_i are generated by appropriate differential equations. Since for each fixed time t, μ_p is a polynomial in p, the value of p that minimizes $\mu_p(t)$ is either a real root of $\frac{\partial \mu_p}{\partial p}(t)$ or one of the two boundary points of $\mathcal{P}$, so the minimization is straightforward. The use of the negative exponential in (17) provides a forgetting factor, which is desirable for processes whose unknown parameters drift slowly. □

A reader familiar with adaptive control would probably argue that reparameterization can be used to circumvent the difficulty in this example. However, there is a strong reason against doing this, as we explain below.

4.2 Loss of stabilizability

In the above discussion, we were concerned with the restrictions imposed on the classes of admissible process models by the estimation procedure. Of course, even if the estimation issues are resolved, additional conditions are needed to ensure that control objectives can be met. To apply the certainty equivalence stabilization theorem from [33], one must have a stabilizing controller for each estimated nominal model. Although it may seem quite reasonable to require that every candidate nominal model be stabilizable, this requirement is often at odds with the linearity and convexity properties mentioned earlier. In fact, to ensure the properties needed for continuous tuning, one often needs to "overparameterize" the candidate nominal models, thus introducing loss of stabilizability. This can be avoided with supervisory control.

To illustrate this point, let us revisit Example 5. One could obtain a linear parameterization of (15) by introducing a fictitious parameter, say $q := p^2$, and then considering a new two-dimensional convex parameter set $\bar{\mathcal{P}}$. However, the smallest such set turns out to be $\bar{\mathcal{P}} = \{(p, q) : p^2 \leq q \leq 1, p \in \mathcal{P}\}$, and for $(p, q) = (0, 1/2)$ the corresponding transfer function

$$\frac{s - \frac{1}{6}(p + 2)}{s^3 + ps^2 - \frac{2}{9}(q + 2p)s}$$

has an unstable pole-zero cancellation. It turns out that any other linear reparameterization that utilizes a convex parameter set has similar problems. This is because, as p ranges over $\mathcal{P}$, the three parameter-dependent coefficients of the transfer function (15) form a curve in $\mathbb{R}^3$ whose convex hull contains points corresponding to unstable pole-zero cancellations [36]. It follows that for this example, *every* linear reparameterization that utilizes a convex parameter set $\bar{\mathcal{P}}$ will necessarily introduce new transfer functions with unstable pole-zero cancellations. These correspond to processes not stabilizable by linear output feedback, which prevents the use of the certainty equivalence stabilization theorem. This is known as the *loss of stabilizability* problem.

In the context of adaptive control, the above difficulty is usually overcome by introducing additional assumptions on the family of process models or by using techniques such as parameter projection and persistent excitation (see [20, Section 7.6]). However, with supervisory control one can avoid these measures altogether by working with families of process models that are nonlinearly parameterized over non-convex sets, instead of using reparameterizations that may lead to unstable pole-zero cancellations and consequent loss of stabilizability. As we explained earlier, logic-based switching in the supervisory control framework provides significant flexibility in addressing these issues. It is also important to emphasize that even when a reparameterization does not introduce unstable pole-zero cancellations, it enlarges the parameter set, thereby introducing "artificial uncertainty" which typically leads to performance degradation.

Another technique that uses switching and logic was proposed in [42, 45] for dealing with the loss of stabilizability problem directly. The key observation is that one can find candidate controllers for which the Detectability Property holds, even when not all nominal models are stabilizable. This is because the certainty equivalence stabilization theorem provides a condition for detectability that is sufficient but not necessary[5]. In [42, 45] it is shown how to find candidate controllers that guarantee detectability when the value of the parameter estimate $\hat{p}$ corresponds to a nominal model that is not stabilizable. A single controller is in general not sufficient, and therefore *cyclic switching* among several controllers is needed. In the above example featuring a reparameterization, a switching cycle would be executed every time the parameter estimate hits the singular value $(0, 1/2)$. As explained in [35], the same technique can also be applied in the supervisory control context.

5 Conclusions

In this paper we described the basic features of the supervisory control framework for uncertain systems, which relies on logic-based switching among a family of candidate controllers. We demonstrated that it allows one to overcome fundamental difficulties associated with more standard continuously-tuned adaptive control algorithms. Specifically, these difficulties include a limited ability to use off-the-shelf, non-adaptive designs and severe constraints imposed on the classes of admissible process models.

Adding to all the benefits of supervisory control discussed above, it is actually our belief that one of the most important ones is the high performance of this type of adaptation algorithms in terms of both transient and steady-state behavior. A precise comparison of the performance of supervisory and traditional adaptive control methods is beyond the scope of this paper. However, in view of the quantitative results obtained so far (see especially [39]), the remarkably good performance of supervisory control in simulations and applications (see the examples given in this paper), and recent advances in the theory of hybrid systems, a quantitative performance-based theory of adaptive control based on switching and logic may be within our reach.

It should be mentioned that, when state-sharing is not possible, supervisory control can lead to high-dimensional controllers (as in Example 4). In these cases, one should certainly think twice before using supervisory control and only take this route if a simpler robust controller with acceptable performance cannot be found. However, it is important to note that when state-sharing fails, adaptation based on continuous tuning is generally not applicable because the unknown parameters enter the nominal model in a nonlinear fashion.

Finally, it should be added that many problems remain open in this area. These include: establishing tight bounds on the closed-loop performance of supervised systems; proving stability for the supervision of nonlinear systems with infinite parametric uncertainty and/or unmodeled dynamics; developing methodologies for "gridding" the process and/or controller parameter spaces when state-sharing is not possible; avoiding switching to candidate controllers that may not stabilize the process (even if these controllers are only kept in the loop for a short period of time); etc. Promising results on the last two topics can be found in [1] and [2], respectively.

References

[1] B. D. O. Anderson, T. S. Brinsmead, F. D. Bruyne, J. P. Hespanha, D. Liberzon, and A. S. Morse. Multiple model adaptive control, part 1: finite controller coverings. *Int. J. Robust Nonlinear Control*, 10:909–929, 2000.

[2] B. D. O. Anderson, T. S. Brinsmead, D. Liberzon, and A. S. Morse. Multiple model adaptive control with safe switching. *Int. J. Adaptive Control Signal Proc.*, 15:445–470, 2001.

[5] Another manifestation of this fact is the certainty equivalence output stabilization theorem, useful in model reference adaptive control, which states that the Detectability Property still holds for a minimum-phase process if the controller only output-stabilizes the estimated model.

[3] A. M. Annaswamy, F. P. Skantze, and A.-P. Loh. Adaptive control of continuous time systems with convex/concave parametrization. *Automatica*, 34:33–49, 1998.

[4] D. Borrelli, A. S. Morse, and E. Mosca. Discrete-time supervisory control of families of two-degrees-of-freedom linear set-point controllers. *IEEE Trans. on Automat. Contr.*, 44(1):178–181, Jan. 1998.

[5] R. W. Brockett. Asymptotic stability and feedback stabilization. In R. W. Brockett, R. S. Millman, and H. J. Sussmann, editors, *Differential Geometric Control Theory*, pages 181–191. Birkhäuser, Boston, 1983.

[6] G. Chang, J. P. Hespanha, A. S. Morse, M. Netto, and R. Ortega. Supervisory field-oriented control of induction motors with uncertain rotor resistance. *Int. J. of Adaptive Control and Signal Processing* Special Issue on Switching and Logic, 15(3):353–375, May 2001.

[7] R. Colbaugh and K. Glass. Stabilization of nonholonomic robotic systems using adaptation and homogeneous feedback. *Journal of Intelligent and Robotic Systems*, 26:1–27, 1999.

[8] M. A. Dahleh and M. M. Livstone. A unified framework for identification and control. In Francis and Tannenbaum [9], pages 196–214.

[9] B. A. Francis and A. R. Tannenbaum, editors. *Feedback Control, Nonlinear Systems and Complexity.* Number 202 in Lecture Notes in Control and Information Sciences. Springer-Verlag, London, 1995.

[10] R. A. Freeman and P. V. Kokotović. Inverse optimality in robust stabilization. *SIAM J. Control Optim.*, 34:1365–1391, 1996.

[11] M. Fu and B. R. Barmish. Adaptive stabilization of linear systems via switching control. *IEEE Trans. Automat. Control*, 31:1097–1103, 1986.

[12] J. P. Hespanha. *Logic-Based Switching Algorithms in Control.* PhD thesis, Dept. of Electrical Engineering, Yale University, New Haven, CT, 1998.

[13] J. P. Hespanha. Tutorial on supervisory control. Lecture Notes for the workshop *Control using Logic and Switching* for the 40th Conf. on Decision and Contr., Orlando, Florida, Dec. 2001.

[14] J. P. Hespanha, D. Liberzon, and A. S. Morse. Logic-based switching control of a nonholonomic system with parametric modeling uncertainty. *Syst. & Contr. Lett.* Special Issue on Hybrid Systems, 38(3):167–177, Nov. 1999.

[15] J. P. Hespanha, D. Liberzon, and A. S. Morse. Supervision of integral-input-to-state stabilizing controllers. *Automatica*, 38(8), Aug. 2002. To appear.

[16] J. P. Hespanha, D. Liberzon, A. S. Morse, B. D. O. Anderson, T. S. Brinsmead, and F. de Bruyne. Multiple model adaptive control, part 2: switching. *Int. J. of Robust and Nonlinear Control* Special Issue on Hybrid Systems in Control, 11(5):479–496, Apr. 2001.

[17] J. P. Hespanha and A. S. Morse. Certainty equivalence implies detectability. *Systems Control Lett.*, 36(1):1–13, Jan. 1999.

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

[19] J. Hockerman-Frommer, S. R. Kulkarni, and P. J. Ramadge. Controller switching based on output prediction errors. *IEEE Trans. on Automat. Contr.*, 43(5):596–607, May 1998.

[20] P. A. Ioannou and J. Sun. *Robust Adaptive Control.* Prentice-Hall, New Jersey, 1996.

[21] Z. P. Jiang and J.-B. Pomet. Global stabilization of parametric chained-form systems by time-varying dynamic feedback. *Int. J. Adaptive Control Signal Proc.*, 10:47–59, 1996.

[22] E. B. Kosmatopoulos and P. Ioannou. A switching adaptive controller for feedback linearizable systems. *IEEE Trans. on Automat. Contr.*, 44(4):742–750, Apr. 1999.

[23] R. Kosut, M. Lau, and S. Boyd. Set-membership identification of systems with parametric and nonparametric uncertainty. *IEEE Trans. on Automat. Contr.*, 37(7):929–941, July 1992.

[24] R. L. Kosut. Iterative unfalsified adaptive control: Analysis of the disturbance-free case. In *Proc. 1999 American Control Conf.*, pages 566–570, June 1999.

[25] S. R. Kulkarni and P. J. Ramadge. Model and controller selection policies based on output prediction errors. *IEEE Trans. on Automat. Contr.*, 41(11):1594–1604, Nov. 1996.

[26] D. Liberzon's homepage, http://black.csl.uiuc.edu/˜liberzon/parkingmovie.mpg.

[27] D. Liberzon, J. P. Hespanha, and A. S. Morse. Hierarchical hysteresis switching. In *Proc. 39th IEEE Conf. on Decision and Control*, pages 484–489, 2000.

[28] D. Liberzon, A. S. Morse, and E. D. Sontag. Output-input stability and minimum-phase nonlinear systems. *IEEE Trans. Automat. Control*, 47:422–436, 2002.

[29] B. Mårtensson. The order of any stabilizing regulator is sufficient a priori information for adaptive stabilization. *Systems Control Lett.*, 6:87–91, 1985.

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

[31] D. Miller and E. J. Davison. An adaptive controller which provides an arbitrary good transient and steady-state response. *IEEE Trans. on Automat. Contr.*, 36(1):68–81, Jan. 1991.

[32] A. S. Morse. Towards a unified theory of parameter adaptive control: tunability. *IEEE Trans. Automat. Control*, 35:1002–1012, 1990.

[33] A. S. Morse. Towards a unified theory of parameter adaptive control, part II: certainty equivalence and implicit tuning. *IEEE Trans. Automat. Control*, 37:15–29, 1992.

[34] A. S. Morse. Control using logic-based switching. In A. Isidori, editor, *Trends in Control*, pages 69–113. Springer, New York, 1995.

[35] A. S. Morse. Logic-based switching and control. In B. A. Francis and A. R. Tannenbaum, editors, *Feedback Control, Nonlinear Systems, and Complexity*, pages 173–195. Springer, New York, 1995.

[36] A. S. Morse. Supervisory control of families of linear set-point controllers—part 1: exact matching. *IEEE Trans. on Automat. Contr.*, 41(10):1413–1431, Oct. 1996.

[37] A. S. Morse, editor. *Control Using Logic-Based Switching.* Number 222 in Lecture Notes in Control and Information Sciences. Springer-Verlag, London, 1997.

[38] A. S. Morse. Supervisory control of families of linear set-point controllers—part 2: robustness. *IEEE Trans. on Automat. Contr.*, 42(11):1500–1515, Nov. 1997.

[39] A. S. Morse. A bound for the disturbance-to-tracking-error gain of a supervised set-point control system. In D. Normand-Cyrot and Y. D. Landau, editors, *Perspectives in Control: Theory and Applications*, pages 23–41. Springer, London, 1998.

[40] E. Mosca, F. Capecchi, and A. Casavola. Designing predictors for MIMO switching supervisory control. *Int. J. of Adaptive Control and Signal Processing* Special Issue on Switching and Logic, 15(3):265–286, May 2001.

[41] K. S. Narendra and J. Balakrishnan. Adaptive control using multiple models. *IEEE Trans. on Automat. Contr.*, 42(2):171–187, Feb. 1997.

[42] F. M. Pait. *Achieving Tunability in Parameter-Adaptive Control.* PhD thesis, Dept. of Electrical Engineering, Yale University, New Haven, CT, 1993.

[43] F. M. Pait and F. Kassab, Jr. Parallel algorithms for adaptive control: Robust stability. In Morse [37], pages 262–276.

[44] F. M. Pait and F. Kassab, Jr. On a class of switched, robustly stable, adaptive systems. *Int. J. of Adaptive Control and Signal Processing* Special Issue on Switching and Logic, 15(3):213–238, May 2001.

[45] F. M. Pait and A. S. Morse. A cyclic switching strategy for parameter-adaptive control. *IEEE Trans. Automat. Control*, 39:1172–1183, 1994.

[46] M. G. Safonov. Focusing on the knowable: Controller invalidation and learning. In Morse [37], pages 224–233.

[47] M. G. Safonov and T.-C. Tsao. The unfalsified control concept: A direct path from experiment to controller. In Francis and Tannenbaum [9], pages 196–214.

[48] E. D. Sontag. Smooth stabilization implies coprime factorization. *IEEE Trans. Automat. Control*, 34:435–443, 1989.

[49] G. Zames. Towards a general complexity-based theory of identification and adaptive control. In Morse [37], pages 208–223.