In Proc. of the 30th Annual Conf. on Information Sciences and Syst., vol. 2, pages 765-768, Mar. 1996.

Towards the Supervisory Control of Families of Nonlinear Regulators*

João Hespanha & A. S. Morse
Department of Electrical Engineering
Yale University, New Haven, CT 06520

Abstract

A simply-structured high-level controller called a "supervisor," has recently been proposed in [1] for the purpose of orchestrating the switching of a sequence of candidate controllers into feedback with an imprecisely modeled process so as to stabilize the process and make its output track a constant reference input. In [1], the process is assumed to be modeled by a SISO linear system which is known only to the extent that it is an element within the union of a number of subclasses, each subclass being small enough so that one of the candidate controllers would solve the problem, were the process's model to be one of the subclass's members. Each subclass contains a "nominal process model" about which the subclass is centered. The supervisor's unique feature, distinguishing it from other logics which might be used for the same purpose, is that controller selection is made by (i) continuously comparing in real time suitably defined "normed" output estimation errors or "performance signals" determined by the nominal models and (ii) by placing in the feedback-loop, from time to time, that candidate controller whose corresponding performance signal is the smallest. The paradigm is a manifestation of the idea of certainty equivalence from parameter adaptive control. The aim of this paper is to discuss some preliminary results dealing with the extension of these ideas to the supervisory control of families of nonlinear regulators.

1 Introduction

The process Σ_P to be controlled is assumed to be modeled by a one-dimensional differential equation of the form

$$\dot{y} = a(y,\omega) + b(\omega)u \tag{1}$$

where ω is an uncertain parameter taking values in a finite set $\mathcal{P} = \{p_1, \ldots, p_m\} \subset \mathbb{R}$. We assume that $a(\cdot)$ and $b(\cdot)$ are given functions such that, for each fixed $p \in \mathcal{P}$, $b(p)$ is nonzero, $a(0,p) = 0$, and $a(y,p)$ is a locally Lipschitz continuous function in y which satisfies the *growth condition*

$$|a(y,p)| \leq g\left(|y| + y^2\right), \qquad \forall y \in \mathbb{R}$$

for some non-negative constant g not depending on p. An example of such a pair of functions is $a(y,p) = py^2 \sin(py)$ and $b(p) = \cos(p)$ where $\mathcal{P}$ is any finite subset of $[0, 2\pi]$ not containing any zeros of $\cos(p)$. Despite its apparent simplicity, for some choices of $a(\cdot)$ and $b(\cdot)$, the problem of developing stabilizing controllers for such one-dimensional systems, seems to be beyond the scope of existing adaptive and/or robust control techniques.

In the sequel we take as given, a family of candidate control laws of the form

$$u_p = f(y,p) \triangleq -\frac{1}{b(p)}\Big(\mu\, y + a(y,p)\Big), \quad p \in \mathcal{P}$$

where μ is a positive number. Note that the control

$$u \triangleq u_\omega \tag{2}$$

is a conventional {i.e., "off-the-shelf"} feedback law which both stabilizes and linearizes the model of Σ_P defined by (1). In fact, with u so defined, y satisfies $\dot{y} = -\mu\, y$. Since the value of ω is not presumed to be know, implementation of (2) is not possible. What's to be used instead is the "multi-controller" $\Sigma_C(\sigma)$ defined by the feedback law

$$u = u_\sigma$$

where σ is a suitably defined piecewise-constant *switching signal* taking values in $\mathcal{P}$. The system to be supervised thus consists of the feedback interconnection of $\Sigma_C(\sigma)$ and the process Σ_P modeled by (1).

The problem of interest is to develop a provably correct, "high level" controller called a *supervisor* {c.f. Figure 1} which, without a-priori knowledge of ω, is capable of choosing σ so as to cause y to tend to zero provided $y(0)$ lies within a prescribed bounded set containing the origin.

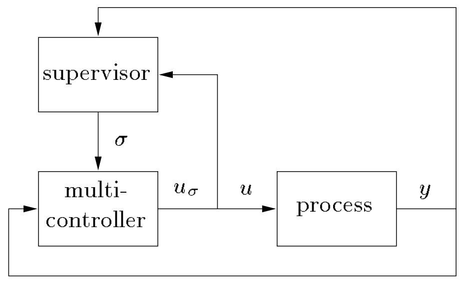

Figure 1: Supervisory Control.

*This research was supported by NSF Grant n. ECS-9206021, AFOSR Grant n. F49620-94-1-0181, and ARO Grant n. DAAH04-95-1-0114

2 Estimator-based Supervisor

The supervisor proposed consists of three subsystems {c.f. Figure 2}:

multi-estimator Σ_E – a dynamical system whose inputs are u and y and whose outputs are *estimation errors* of the form $e_p = y_p - y$, $p \in \mathcal{P}$; here each y_p is a suitably defined estimate of y which would be asymptotically correct if p were equal to ω.

performance signal generator Σ_Π – a dynamical system whose inputs are the estimation errors and whose outputs $\{\pi_p : p \in \mathcal{P}\}$ are suitably "normed" values of the estimation errors called *performance signals*.

switching logic Σ_S – a hybrid dynamical system whose inputs are the performance signals and whose output is σ.

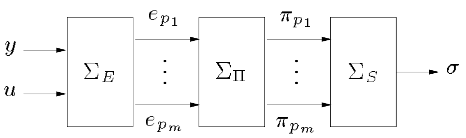

Figure 2: Estimator-based Supervisor

The underlying decision making strategy used by this supervisor is basically this: From time to time select for σ, that candidate control index q whose corresponding performance signal π_q is the smallest among the π_p, $p \in \mathcal{P}$. Motivation for this idea is as follows: the nominal process model whose associated performance signal is the smallest, "best" approximates what the process is and thus the candidate controller designed on the basis of that model ought to be able to do the best job of controlling the process. The origin of this idea is the concept of certainty equivalence from parameter adaptive control.

2.1 Multi-Estimator

For the problem at hand, the multi-estimator Σ_E consists of a family of one-dimensional estimators of the form

$$\dot{y}_p = -\lambda(y_p - y) + a(y,p) + b(p)u, \quad p \in \mathcal{P} \tag{3}$$

$$e_p = y_p - y \tag{4}$$

where λ is a positive number. Thus Σ_E is a dynamical system with state $x_E = [y_{p_1} \ \cdots \ y_{p_m}]'$, and inputs u and y. Note that because of (1) and the defining equations (3) and (4),

$$\dot{e}_\omega = -\lambda\, e_\omega \tag{5}$$

In other words, for any input u, the estimation error associated with ω converges to zero exponentially fast. For simplicity we shall henceforth assume that $\mu = \lambda$.

2.2 Performance Signals Generator

The performance signal generator Σ_Π to be considered consists of the following family of one-dimensional differential equations

$$\dot{\pi}_p = -\gamma\pi_p + e_p^2, \qquad \pi_p(0) = 0, \qquad p \in \mathcal{P} \tag{6}$$

where γ is a positive number smaller than λ $(= \mu)$. The definition implies that for each $p \in \mathcal{P}$

$$\pi_p(t) = \int_0^t e^{-\gamma(t-\tau)} e_p^2(\tau)\, d\tau, \qquad t \geq 0$$

which shows that the performance signal π_p is just an exponentially weighted $\mathcal{L}_2$-norm of the estimation error e_p.

State Sharing

In the event that $a(y,p)$ is a *separable function* of the form

$$a(y,p) = a_1(y)a_2(p),$$

the estimation errors can be generated more efficiently using "state sharing." In this case one could replace the individual estimator equations (3), (4) by the three-dimensional system

$$\dot{x}_E = -\lambda x_E + \begin{bmatrix} \lambda y \\ a_1(y) \\ u \end{bmatrix}$$

together with the estimation errors

$$e_p = c_p \begin{bmatrix} x_E \\ y \end{bmatrix}, \quad p \in \mathcal{P}$$

where c_p is the 1×3 matrix $c_p = \begin{bmatrix} 1 & a_2(p) & b(p) & -1 \end{bmatrix}$. Note that this change would not affect the validity of (5). Moreover, since

$$\left(c_p \begin{bmatrix} x_E \\ y \end{bmatrix} \right)^2$$

is a separable function of the form

$$\rho_1(c_p)\rho_2\left(\begin{bmatrix} x_E \\ y \end{bmatrix} \right),$$

it is also possible to generate the performance signals governed by (6) using state-shared equations of the form

$$\dot{w} = -\gamma w + \rho_2\left(\begin{bmatrix} x_E \\ y \end{bmatrix} \right)$$

together with the readout equations

$$\pi_p = \rho_1(c_p)w$$

State sharing is useful if $\mathcal{P}$ contains a large number of points. It is even possible to exploit these ideas in the case when $\mathcal{P}$ contains a continuum of points [2]. In this case however, tractability considerations would typically dictate that $\mathcal{P}$ be convex and that for each fixed value of w, $\rho_1(c_p)w$ be a convex function on $\mathcal{P}$.

2.3 Switching Logic

The switching logic Σ_S we shall consider is a simplified version of the dwell time switching logic introduced in [1]. Σ_S is a hybrid system whose inputs are the performance signals and whose state and output is the piecewise-constant switching signal σ whose value is held constant between successive "event times." For the logic considered here event times are equally spaced by a pre-specified positive number τ_D called *dwell time*. To be more precise, event times are simply the set of times $\{k\,\tau_D : k \in \mathbb{N}\}$. The initial value of σ on the interval $[0, \tau_D)$ can be any point in $\mathcal{P}$ and for each $k \geq 1$, σ's value on $[k\,\tau_D, (k+1)\tau_D)$ is defined by the rule

$$\sigma = \begin{cases} \min\{p \in \mathcal{P} : \pi_p(k\,\tau_D) < \pi_{\sigma'}(k\,\tau_D)\} & \text{if nonempty} \\ \sigma' & \text{otherwise} \end{cases}$$

where σ' is the value of σ on the preceding interval.

3 Motivation

At this point we digress briefly to discuss the significance of the implication

$$\textit{certainty equivalence} \Rightarrow \textit{detectability} \tag{7}$$

since this is what's motivated both the architecture and analysis of the nonlinear supervisory control system just described. We begin by remarking that (7) has already played a key role in the development and analysis of supervisory control systems for processes whose models are linear [1]. It turns out however that linearity is not crucial and that (7) is also valid for a broader class of {nonlinear} problems which encompasses the one we've been considering. In order to explain what (7) means we need first of all, a good working definition of detectability appropriate to nonlinear systems. One possibility has been suggested in [2] and is as follows[1]. Suppose we agree to call a smooth n-dimensional dynamical system of the form

$$e = C(x) \qquad\qquad \dot{x} = A(x) \tag{8}$$

detectable if there exist positive definite, radially unbounded, continuously differentiable functions $V : \mathbb{R}^n \to \mathbb{R}$ and $W : \mathbb{R} \to \mathbb{R}$ such that

$$\frac{\partial V}{\partial x}A(x) - W(||C(x)||) < 0, \quad x \in \mathbb{R}^n,\ x \neq 0 \tag{9}$$

- The definition characterizes detectability more as a generalization of stability than of observability. In other words, rather than considering a detectable system to be a system whose unobservable modes {if any} are stable, it is just as valid to consider a detectable system to be one whose unstable modes {if any} are observable. This prompts one to ask how Lyapunov's definition of a stable system should be modified to account for unstable but observable modes. The aforementioned definition of detectability provides a natural answer.

- In the linear case when $C(x) = Cx$ and $A(x) = Ax$, the standard definition of detectability is known to be equivalent to the existence of a positive definite matrix P which satisfies the matrix inequality $PA + A'P - C'C < 0$ [4]; since $V \triangleq x'Px$ and $W(\mu) \triangleq \mu^2$ satisfy (9), this definition of detectability has the virtue of reducing to the standard one in the linear case.

- It can be easily shown that if (8) is a detectable nonlinear system and $W(e)$ has a finite $\mathcal{L}^1[0,\infty)$ norm along some solution x, then x must tend to zero. Thus this definition fulfills the intuitively appealing requirement that convergence to zero of the output of a detectable system ought to imply convergence to zero of the system's state.

To proceed, let us recall that in an adaptive context, the heuristic idea of *certainty equivalence* advocates that the feedback control to a poorly modeled process should, at each instant of time, be designed on the basis of a current process model estimate with the understanding that this estimate is to be viewed as correct, even though it may not be. Suppose for the problem at hand, that at some specific time, the qth "estimated model"

$$\dot{\bar{y}} = a(\bar{y}, q) + b(q)\bar{u} \tag{10}$$

is held to be the best current candidate for the actual process model (1). Suppose in addition that the stabilizing feedback law $\bar{u} = f(\bar{y}, q)$ is what one would use, were (10) in fact the correct process model.

Let us examine the consequence of using the certainty equivalence control $u = f(y, q)$. Application of this feedback law to (1) and (3) yields the closed-loop differential equations

$$\dot{y} = a(y, \omega) + b(\omega)f(y, q) \tag{11}$$

$$\dot{y}_p = -\lambda(y_p - y) + a(y, p) + b(p)f(y, q), \quad p \in \mathcal{P} \tag{12}$$

Meanwhile from (4) one can get the equation for the relevant output estimation error, namely

$$e_q = y_q - y \tag{13}$$

Equations (11)-(13) are of the form (8) where $e = e_q$ and $x = [\,y \quad y_{p_1} \quad \cdots \quad y_{p_m}\,]'$. It can be shown that there is a positive definite matrix Q such that (9) holds with $V(x) \triangleq x'Qx + y_q^4$ and $W(\mu) \triangleq \mu^2 + \mu^4$. In other words (11)-(13) is detectable through e_q. This consequence of the fact that $\bar{u} = f(\bar{y}, q)$ stabilizes (10), explains why certainty equivalence implies detectability for the problem

[1] See also [3] in this proceedings.

of interest. It is worth noting that detectability of (11)-(13) does *not* depend on ω being an element of $\mathcal{P}$ or, for that matter, on the relation between u and y being given by the one-dimensional system (1). In other words, the validity of (7) has nothing to do with "exact matching."

With the preceding observations in mind, recall the underlying decision making strategy of an estimator-based supervisor: From time to time select for σ, that value $q \in \mathcal{P}$ such that the performance signal π_q is the smallest among the $\pi_p,\ p \in \mathcal{P}$. Justification for this strategy is now clear: By choosing σ to maintain smallness of π_σ and consequently e_σ, the supervisor is also maintaining smallness of the composite state of the process Σ_P and the multi-estimator Σ_E because of detectability through e_σ.

4 Main Results

The results which follow are for the special case when the set $\mathcal{P}$ has only two elements, i.e. $\mathcal{P} = \{p_1, p_2\}$. In the sequel h is any given constant such that

$$\frac{|b(p)|}{|b(q)|} \leq h, \qquad \forall p, q \in \mathcal{P}$$

The existence of such a constant is assured by the aforementioned requirement that $b(\cdot)$ be nonzero on $\mathcal{P}$.

Theorem 4.1 *Let τ_D, λ, and γ be positive constants such that $\gamma < \lambda$ and $4(1+h)^2\left(1+2g/\lambda^2\right) < e^{(\lambda-\gamma)\tau_D}$. There is a positive constant δ, depending only on τ_D, λ, and γ, as well as on g and h, such that*

$$y \to 0, \qquad \textit{and} \qquad y_p \to 0,\ p \in \mathcal{P} \tag{14}$$

as $t \to \infty$, provided

$$|y(0)| + \sum_{p \in \mathcal{P}} |y_p(0)| \leq \delta \tag{15}$$

Moreover, given any pre-specified positive value for δ it is possible to find values for τ_D, λ, and γ for which (14) is true provided (15) holds.

The proof of this theorem is given in [5].

5 Concluding Remarks

The main shortcoming of the approach we've outlined, when compared with existing adaptive techniques, is that the derived results do not hold globally in the space in which y and the $\{y_p : p \in \mathcal{P}\}$ take values. Although local, the results are of a "semi-global" nature in the sense that the neighborhood of the origin for which asymptotic stability is guaranteed can be made arbitrarily large by suitable choice of the design parameters.

On the basis of what is known so far about supervisory control of *linear* regulators {e.g., see [2]}, it is reasonable to expect that the approach will prove to have several advantages over conventional adaptive techniques. Among these are the following.

1. Unlike conventional adaptive control, supervisory control can deal with nonlinearly parameterized models.
2. Since supervisory control does not necessarily involve nonlinear "zero" cancelling, the approach is not restricted to "minimum phase" systems. For example, the approach applies {without requiring the measurement of state} to the nonminimum phase, output feedback linearizable system
$$\ddot{y} = \left(1 - \frac{d}{dt}\right)\left(\omega y^2 - u\right)$$
where ω is an uncertain parameter in $\mathcal{P} \triangleq [-1, 1]$.
3. Supervisory control makes use of off-the-shelf {non-adaptive} controllers rather than specialized controllers tailored to the specific needs of continuously-tuned conventional adaptive algorithms.
4. Supervisory control is almost certainly inherently robust to small modeling errors[2] whereas conventional adaptive techniques typically are not unless tuner modifications are introduced.
5. It is almost certainly true that supervisory control can be shown to perform satisfactory in the face of small disturbance and measurements signals whereas, without substantial modification, more familiar adaptive techniques cannot.

References

[1] A. S. Morse. Supervisory control of linear set-point controllers - part 1: exact matching. Technical report, Lab. for Control Science & Eng., Yale University, Mar. 1993.

[2] A. S. Morse. Control using logic-based switching. In A. Isidori, editor, *Trends in Control: A European Perspective*, pp. 69–113. Springer-Verlag, 1995.

[3] E. D. Sontag and Y. Wang. Detectability of nonlinear systems. In *Proc. of the 30th CISS*, vol. 2, pp. 1031–1036, Mar. 1996.

[4] S. Boyd, L. E. Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*, vol. 15 of *SIAM Studies in Applied Mathematics*. SIAM, 1994.

[5] J. P. Hespanha and A. S. Morse. Supervisory control of families of nonlinear regulators. Technical Report 9601, Lab. for Control Science & Eng., Yale University, Feb. 1996.

[2] This is, in part, because the validity of (7) has nothing to do with exact matching; i.e., with (1) being a precise model of Σ_P.