Systems & Control Letters 97 (2016) 91–97

Contents lists available at ScienceDirect

Systems & Control Letters

journal homepage: www.elsevier.com/locate/sysconle

The role of convexity in the adaptive control of rapidly time-varying systems

Daniel E. Miller

Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, Ontario, Canada

ARTICLE INFO

Article history:
Received 1 May 2016
Received in revised form
8 August 2016
Accepted 11 September 2016

Keywords:
Adaptive control
Convexity
Time-varying systems

ABSTRACT

In classical adaptive control the parameters are assumed to be fixed or slowly time-varying. In order to facilitate parameter estimation/tuning it is desirable to have the set of admissible parameters lie in a convex set; if this set is not convex, a common trick is to replace it with its convex hull, but the adaptive control problem is challenging if stabilizability of the set of admissible parameters is lost. However, such a convexity assumption is an artifact of the approach to the problem, rather than an inherent constraint, since most logic-based and supervisory approaches to the problem make no such requirement. On the other hand, here we show that losing stabilizability on the convex hull of the set of admissible parameters plays an important role in the adaptive control of rapidly time-varying systems.

1. Introduction

In classical parameter adaptive control, restricting the set of admissible parameters to a convex set is useful in carrying out parameter estimation/tuning, e.g. see [1]. Of course, if the set of admissible parameters is not convex, it is natural to replace it with its convex hull; however, this can create the problem of introducing uncontrollable or unobservable modes, which can create difficulty in proving that the associated adaptive controller is stabilizing. This has spurred a fair bit of effort to get around this problem, e.g. see [2,3], and [4]. These methods have been successful, and are effective in controlling plant models whose parameters are either fixed or slowly time-varying; this is also true of most logic-based and supervisory approaches to adaptive control, e.g. see [5,6].

Now let us turn to the adaptive control of rapidly time-varying systems. This problem is very difficult, and only limited results have been obtained, each of which requires fairly rich structure on the plant model:

(i) the form of the time-variations (or at least of the fast terms) is assumed to be known (e.g. see [7,8]);
(ii) the only uncertainty is a gain at the input, e.g. see [9];
(iii) the plant has stable zero dynamics (roughly speaking, this is the time-varying counterpart of minimum phase), e.g. see [10–15];
(iv) the plant has unstable zero dynamics but several stringent matching requirements must hold—see [16,17].

In this paper our goal is to ascertain performance limitations in the adaptive control of rapidly time-varying systems. To avoid imposing unnecessary structure on the set of admissible plant parameters (such as connectedness), we restrict our attention to that of jumps in the plant parameters. We demonstrate that, in two important cases, if the convex hull of the set of admissible parameters does not possess a weak notion of stabilizability, then regardless of the controller used, the performance must necessarily degrade rapidly as the time between parameter jumps decreases. This provides an inviolable bound on the achievable performance of any adaptive controller for such a rapidly time-varying uncertain system.

2. Mathematical preliminaries

Let $\mathbf{Z}$ denote the set of integers, $\mathbf{Z}^+$ represent the set of non-negative integers, $\mathbf{N}$ denote the set of natural numbers, $\mathbf{R}$ denote the set of real numbers, $\mathbf{R}^+$ represent the set of non-negative real numbers, and $\mathbf{C}$ represent the set of complex numbers. We will use the Euclidean norm to measure the size of a vector: for $x \in \mathbf{C}^n$, we define $\|x\| := (\sum_{i=1}^n |x_i|)^{1/2}$. The corresponding induced norm of a matrix $A \in \mathbf{C}^{m\times n}$ is defined in a usual manner: $\|A\| = \sup_{\|x\|\neq 0} \frac{\|Ax\|}{\|x\|}$. If $x \in \mathbf{C}^n$ we use x^T to denote the transpose and x^* to denote the complex conjugate transpose.

E-mail address: miller@uwaterloo.ca.

http://dx.doi.org/10.1016/j.sysconle.2016.09.008

For a given set $\mathcal{S} \subseteq \mathbf{R}^{m\times n}$, we let $PC(\mathcal{S})$ denote the set of all piecewise continuous functions $f : \mathbf{R}^+ \to \mathcal{S}$. To measure the size of $f \in PC(\mathcal{S})$, we define

$$\|f\|_\infty := \sup_{t\in\mathbf{R}^+} \|f(t)\|;$$ [1]

we let $PC_\infty(S)$ denote the set of all $f \in PC(S)$ for which $\|f\|_\infty < \infty$. With $T_s > 0$, we let $PC_{con}(S, T_s)$ denote the set of all $f \in PC(S)$ which are piecewise constant with a minimum time of T_s between discontinuities. Last of all, we let $conv(S)$ denote the convex hull of S.

3. The setup

Here we will model the plant uncertainty as follows. For a suitable $l \in \mathbf{N}$ we start with a compact set $\Theta \subset \mathbf{R}^l$. With $A : \Theta \to \mathbf{R}^{n\times n}$, $B : \Theta \to \mathbf{R}^{n\times m}$ and $C : \Theta \to \mathbf{R}^{r\times n}$ continuous functions, and $\theta \in PC(\Theta)$, we consider the time-varying plant

$$\dot{x}(t) = A(\theta(t))x(t) + B(\theta(t))u(t), \qquad x(0) = x_0 \tag{1}$$

$$y(t) = C(\theta(t))x(t); \tag{2}$$

here $x(t) \in \mathbf{R}^n$ is the state, $y(t) \in \mathbf{R}^r$ is the measured output, and $u(t) \in \mathbf{R}^m$ is the control input, and we associate the plant with the triple $(A(\theta(\cdot)), B(\theta(\cdot)), C(\theta(\cdot)))$, or simply P. While the set Θ as well as the functions A, B, and C are known, the variable $\theta \in PC(\Theta)$ *is neither known nor measurable*. The control objective here is a form of stability, so it is reasonable to restrict θ to a subset of $PC(\Theta)$; we consider the case of time-variations which are simply jumps, with a minimum distance separating them: with $T_s > 0$, we consider the subset $PC_{con}(\Theta, T_s)$. Associated with this set of admissible θ's is the set of possible plant models:

$$\mathcal{P}(T_s) := \{(A(\theta(\cdot)), B(\theta(\cdot)), C(\theta(\cdot))) : \theta \in PC_{con}(\Theta, T_s)\}.$$

Remark 1. It turns out that the choice of C plays no role in the results which we prove here. However, we will allow a general form for $C(\theta)$ to emphasize the general applicability of the result.

Remark 2. In classical adaptive control many results are proven for the case of fixed parameters. The setup that we adopt here allows this—it corresponds to the case of $T_s = \infty$.

The control objective is to stabilize the system even though there are rapid variations in $\theta(t)$, which is an adaptive control problem. It is traditional in adaptive control to prove very weak notions of stability, often proving only that the system is well behaved asymptotically, with no uniformity over the admissible models in $\mathcal{P}(T_s)$. However, more recently, techniques such as the supervisory control method of Morse, e.g. see [5,6] (see the Concluding Remarks Section of the latter), and the periodic probing, estimation and control technique of the author, e.g. see [11,18,15,9,19,17], have been used to prove stronger uniform notions of stability, even when the parameters are varying, as illustrated in the following definition.

Definition 3. With $K : PC(\mathbf{R}^r) \to PC(\mathbf{R}^m)$ and $T_s > 0$, we say that the controller

$$u = K(y) \tag{3}$$

is **admissible** for $\mathcal{P}(T_s)$ if, for every $P \in \mathcal{P}(T_s)$, the closed-loop system is well-posed: for every $x_0 \in \mathbf{R}^m$ there are unique $u \in PC(\mathbf{R}^m)$ and $y \in PC(\mathbf{R}^r)$ which satisfy the plant model Eqs. (1)–(2) and the controller Eq. (3), in which case we let $\Phi(x_0, P)$ denote the map $x_0 \to \begin{bmatrix} x \\ u \end{bmatrix}$ from $\mathbf{R}^n \to PC(\mathbf{R}^n) \times PC(\mathbf{R}^m)$. If K is admissible for $\mathcal{P}(T_s)$ then we say that K **stabilizes** $\mathcal{P}(T_s)$ if

(i) $\Phi(0, P) = 0$ for every $P \in \mathcal{P}(T_s)$ and
(ii) the following quantity

$$\gamma(K, \mathcal{P}(T_s)) := \sup\left\{ \frac{\|\Phi(x_0, P)\|_\infty}{\|x_0\|} : x_0 \in \mathbf{R}^n \text{ is nonzero and } P \in \mathcal{P}(T_s)\right\}$$

is finite.

From Definition 3 we see that if K stabilizes $\mathcal{P}(T_s)$, then

$$\|\Phi(x_0, P)\|_\infty \le \gamma(K, \mathcal{P}(T_s))\|x_0\|$$

for every $x_0 \in \mathbf{R}^n$ and $P \in \mathcal{P}(T_s)$. Here the goal is to bound $\gamma(K, \mathcal{P}(T_s))$ in certain circumstances. As observed in Remark 1, $C(\theta(\cdot))$ plays no role in our result. To this end, we now define the convex hull of the admissible $(A(\cdot), B(\cdot))$ pairs: with μ playing the role of a dummy variable, we define

$$\mathcal{H} := \text{conv}\{(A(\mu), B(\mu)) : \mu \in \Theta\}. \tag{4}$$

The question at hand is: if there is a pair in $\mathcal{H}$ which loses stabilizability, what is the consequence on stabilizing the corresponding set $\mathcal{P}(T_s)$? Of course, if $T_s = \infty$, then we have the classical adaptive control setup of no time-variations, and there are general techniques such as supervisory control [5,6] as well as the periodic probing, estimation and control technique of [18] which yield stability. So the real concern is that this loss of stabilizability may impact the situation when $T_s < \infty$, measured in terms of a lower bound on $\gamma(K, \mathcal{P}(T_s))$. Here we will show, under suitable assumptions, that $\gamma(K, \mathcal{P}(T_s))$ must necessarily be large if T_s is small. We consider three situations:

- In Section 4, we consider the case of B being fixed, and we prove that if "weak stabilizability"[2] is lost then $\gamma(K, \mathcal{P}(T_s)) \to \infty$ as $T_s \to 0$.
- In Section 5, we assume that A is fixed and B is variable, and provide an example from the literature which demonstrates that no general result is provable.
- In Section 6, we consider the general case of allowable variations in both A and B, but consider a special controller structure associated with step tracking; in this situation a result similar to that of Section 4 can be proven.

4. The Case of time-variations in $A(\theta(\cdot))$ Only

In this case we assume that the only variation is in A, i.e. we assume that $B(\theta(\cdot))$ is constant, so we simply represent it by B. To proceed, we first define

$$\mathcal{A} := \{A(\mu) : \mu \in \Theta\} \subset \mathbf{R}^{n\times n}.$$

We now introduce a weak notion of stabilizability, which differs from the classical notion of stabilizability by not deeming eigenvalues on the imaginary to be in the "bad region ".

Definition 4. (A, B) is **weakly stabilizable** if

$$\text{rank}\begin{bmatrix} A - \lambda I & B \end{bmatrix} = n \tag{5}$$

for all $\lambda \in \mathbf{C}$ satisfying Re $\lambda > 0$; $\mathcal{H} \subset \mathbf{R}^{n\times n} \times \mathbf{R}^{n\times m}$ is **weakly stabilizable** if every pair $(A, B) \in \mathcal{H}$ is **weakly stabilizable**.

We now prove that if $(\text{conv}(\mathcal{A}), B)$ is **not** weakly stabilizable, then the closed-loop performance provided by a controller for $\mathcal{P}(T_s)$ is bounded below by a function of T_s. Before proceeding, define

$$\bar{a} := \sup_{\theta\in\Theta} \|A(\theta)\|, \qquad \bar{b} := \|B\|.$$

[1] Here we will be allowing sampled-data controllers, so we cannot use "ess sup" here.

[2] This is a slightly weaker version of the usual notion of stabilizability and will be defined shortly.