Available online at www.sciencedirect.com
SCIENCE DIRECT®
Systems & Control Letters 49 (2003) 21–36
SYSTEMS & CONTROL LETTERS
www.elsevier.com/locate/sysconle

Randomized algorithms for the synthesis of cautious adaptive controllers☆

M.C. Campi*, Maria Prandini

Department of Electronics for Automation, University of Brescia, Via Branze, 38 25123 Brescia, Italy

Abstract

We introduce a new methodology for the design of cautious adaptive controllers based on the following two-step procedure: (i) a probability measure describing the likelihood of different models is updated on-line based on observations, and (ii) a controller with certain robust control specifications is tuned to the updated probability by means of randomized algorithms. The robust control specifications are assigned as average specifications with respect to the estimated probability measure, and randomized algorithms are used to make the controller tuning computationally tractable.

This paper provides a general overview of the proposed new methodology. Still, many issues remain open and represent interesting topics for future research.

Keywords: Adaptive control; Cautious control; Randomized algorithms; Statistical learning theory; Stochastic systems

1. Introduction

1.1. The optimal control problem

Consider a standard optimal control problem setting, where a plant described by a transfer function $G^{\circ}(z^{-1})$ has to be regulated by a controller to be chosen in a *feasible controller set* $\{K(\gamma, z^{-1}), \gamma \in \Gamma\}$. For a generic system $G(z^{-1})$, the control performance achieved by applying controller $K(\gamma, z^{-1})$ to $G(z^{-1})$ is measured by a (positive) cost criterion $J(G(z^{-1}), K(\gamma, z^{-1}))$ (this can be e.g. an H_2 or an H_∞ cost). The control objective is to determine the control parameter γ that minimizes $J(G^{\circ}(z^{-1}), K(\gamma, z^{-1}))$.

1.2. The cautious adaptive control problem

Throughout the paper we assume that $G^{\circ}(z^{-1})$ is not known. Then, according to the indirect approach to adaptive control, the controller is selected based on some model of $G^{\circ}(z^{-1})$. Typically, a *parameterized family*

☆ Research supported by MIUR under the project "New techniques for the identification and adaptive control of industrial systems".

* Corresponding author. Department of Electrical Engineering, University of Brescia, Via Branze, 38 25123 Brescia, Italy.
E-mail addresses: campi@ing.unibs.it (M.C. Campi), prandini@ing.unibs.it (M. Prandini).

0167-6911/03/$ - see front matter © 2003 Elsevier Science B.V. All rights reserved.
doi:10.1016/S0167-6911(02)00340-7

of models $\{G(\vartheta, z^{-1}), \vartheta \in \Theta\}$ is considered, and the parameter ϑ is estimated on-line based on observations.

In this paper, we do not only update the parameter estimate, but also compute on-line a probability distribution $\mathscr{P}_t$ describing the likelihood of the different ϑ's at time t. The probability $\mathscr{P}_t$ is used in the definition of the robust control cost

$$c_t(\gamma) := E_{\mathscr{P}_t}[J(\cdot, \gamma)] \tag{1}$$

(here, $J(\cdot, \gamma) := J(G(\cdot, z^{-1}), K(\gamma, z^{-1}))$ for ease of notation). Function $c_t(\gamma)$ is the average performance of controller $K(\gamma, z^{-1})$ for a set of plants distributed according to $\mathscr{P}_t$. Minimizing $c_t(\gamma)$ corresponds to optimizing the average control system behavior where different models are given different weights according to their likelihood at time t. When distribution $\mathscr{P}_t$ is spread over the set Θ, the corresponding optimal controller is expected to exhibit robust characteristics and, hence, it will be conservative. As time goes by, more information is accumulated. Then, distribution $\mathscr{P}_t$ becomes more sharply peaked and the minimization of function $c_t(\gamma)$ leads to controllers better tailored to the true plant. This will ultimately result in an improvement of the control performance.

The idea of minimizing average costs such as (1) is not new and, in fact, it has been previously introduced under the name of *cautious control*, see e.g. [2,12]. On the other hand, average costs have been used only in very specific and simple cases. The reason for this has to be found in mere computational issues. Specifically, a main obstacle in the use of (1) is that integrating the cost criterion $J(\cdot, \gamma)$ over Θ with respect to measure $\mathscr{P}_t$ is a very difficult computational task.

1.3. Randomized algorithms

Our main thrust in this paper is to propose the use of randomized algorithms to overcome the computational difficulties involved in the implementation of the average cost approach. The use of randomized algorithms for the controller selection in a non-adaptive context has been proposed in [31]. Roughly, the idea is to substitute the expectation in the average cost with a finite sum, that is $c_t(\gamma)$ in (1) is substituted by a sampling version (see Section 3 for details). On the one hand, this makes the problem computationally tractable. On the other hand, one can resort to very powerful results from the statistical learning literature (and, more precisely, from the theory of uniform convergence of empirical means, [11,26–30]) to quantify the level of approximation introduced by substituting the expectation with a finite sum.

In the context of adaptive control, we design algorithms with two main features: (i) the controller at time t is selected via randomized methods and, therefore, exhibits robustness characteristics, and (ii) the probability distribution $\mathscr{P}_t$ is updated through time starting from an initial probability distribution $\mathscr{P}$, thus enabling the controller to tune to the real plant.

Point (ii) is a key distinguishing feature with respect to [31] and has important implications: (a) it allows the controller to tune to the plant characteristics; and (b) it makes the choice of the initial probability distribution $\mathscr{P}$ much less critical than in the robust control setting of [31], where $\mathscr{P}$ is fixed and a shrewd selection of it is fundamental.

The paper is organized as follows. In Section 2, some notions of robust and tuning controllers are introduced. A general algorithm for the controller selection via randomized methods is discussed in Section 3. Some implementation aspects are left open in the presentation of the algorithm, namely: (a) the updating of $\mathscr{P}_t$; (b) the computation of the so-called Pollard-dimension. These aspects are the subjects of Sections 4 and 5, respectively. Section 6 presents the tuning properties of the proposed algorithm. In Section 7, we consider LQG control as an application example. Finally, the issue of minimizing the sampling version of $c_t(\cdot)$ is discussed in Section 8.

2. Robustness and tuning notions

Let $\mathscr{P}_t$ be the probability distribution on Θ describing the parameter likelihood at time t. According to the cautious control philosophy, at time t the control objective is to minimize the cost function $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$. Given that an exact minimization of this cost function is computationally excessive, we introduce general notions of an approximate minimizer of $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$.

Definition 1. Given $\varepsilon > 0$, γ_t is an approximate minimizer of $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$ to accuracy ε if $E_{\mathscr{P}_t}[J(\cdot,\gamma_t)] \leqslant \inf_{\gamma\in\Gamma} E_{\mathscr{P}_t}[J(\cdot,\gamma)] + \varepsilon$.

Suppose now that γ_t is selected according to a randomized procedure so that it is stochastic. Then, we have the following definition:

Definition 2. Given $\varepsilon > 0$ and $\delta > 0$, a random γ_t is an approximate minimizer of $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$ to accuracy ε with confidence $1-\delta$ if $E_{\mathscr{P}_t}[J(\cdot,\gamma_t)] \leqslant \inf_{\gamma\in\Gamma} E_{\mathscr{P}_t}[J(\cdot,\gamma)] + \varepsilon$, with probability not less than $1-\delta$.

Here the specification "with probability not less than $1-\delta$" makes reference to the probability involved in the stochastic procedure for selecting γ_t.

Definitions 1 and 2 are of robustness type. Basically, they require that the controller associated with γ_t performs well in average over the uncertainty model set at a given time t.

In adaptive control applications, $\mathscr{P}_t$ is updated on-line. Therefore, in the instances in which the model class is rich enough to contain the true transfer function $G^\circ(z^{-1})$, one can expect that, under suitable excitation conditions, $\mathscr{P}_t$ converges (in some sense) to the true parameter $\vartheta^\circ \in \Theta$ and the adaptive controller self-tunes to the optimal controller. This idea is formalized in the following definition.

Definition 3. $K(\gamma_t, z^{-1})$ self-tunes if $\lim_{t\to\infty} J(\vartheta^\circ,\gamma_t) = \inf_{\gamma\in\Gamma} J(\vartheta^\circ,\gamma)$, with probability 1.

3. A randomized algorithm for cautious adaptive control

In this section, we introduce a general algorithm for the implementation of the cautious adaptive control approach described in Section 1.

To start with, we focus attention on the problem of computing an estimate of $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$. Later below, we shall integrate such an estimate in the general algorithm.

3.1. A sampling estimate of $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$

Our discussion here is strictly related to the approach of [31]. Suppose that at time t distribution $\mathscr{P}_t$ has been computed (the actual computation of $\mathscr{P}_t$ is discussed in Section 4). By definition (1), $c_t(\gamma)$ is the integral of function $J(\cdot,\gamma)$ over the parameter space Θ with respect to probability measure $\mathscr{P}_t$. Since an exact evaluation of this integral is difficult, we introduce a sampling version of it based on random extractions.

Denote by $\theta := (\vartheta_1, \vartheta_2, \ldots, \vartheta_M)$ a sequence of M system parameters independently extracted from Θ according to $\mathscr{P}_t$. Then, a sampling version of $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$ is given by

$$\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)] := \frac{1}{M}\sum_{i=1}^{M} J(\vartheta_i,\gamma). \tag{2}$$

In principle, the original problem of minimizing the expected value $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$ can be replaced by the problem of minimizing the easy-to-compute approximant $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$. However, for guaranteeing that this minimization leads to an approximate minimizer of $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$, $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$ must be an *uniformly* good approximation to $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$ over the set Γ, namely

$$\sup_{\gamma\in\Gamma}|\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]-E_{\mathscr{P}_t}[J(\cdot,\gamma)]|\leqslant\frac{\varepsilon}{2} \tag{3}$$

for some $\varepsilon>0$ (ε is the so-called accuracy parameter). Indeed, letting

$$\gamma_t:=\arg\min_{\gamma\in\Gamma}\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)],\qquad \gamma_t^\circ:=\arg\min_{\gamma\in\Gamma}E_{\mathscr{P}_t}[J(\cdot,\gamma)] \tag{4}$$

we have that

$$\begin{aligned}E_{\mathscr{P}_t}[J(\cdot,\gamma_t)]&\leqslant\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma_t)]+\frac{\varepsilon}{2}\leqslant\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma_t^\circ)]+\frac{\varepsilon}{2}\\&\leqslant E_{\mathscr{P}_t}[J(\cdot,\gamma_t^\circ)]+\frac{\varepsilon}{2}+\frac{\varepsilon}{2}=\inf_{\gamma\in\Gamma}E_{\mathscr{P}_t}[J(\cdot,\gamma)]+\varepsilon,\end{aligned}$$

which proves that γ_t is an approximate minimizer of $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$ according to Definition 1. Observe that Eq. (4) implicitly assumes that $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$ and $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$ have a minimum over Γ. Should this not be the case, nearly minimal points can be used instead.

Note that the sampling estimate (2) is based on a random selection of parameters $\vartheta_i\in\Theta$ and, as such, it is a random variable over the space $\Theta^M:=\Theta\times\Theta\times\cdots\times\Theta$, M times. Even for one fixed parameter $\bar{\gamma}$, it may happen that the randomly extracted multisample $\theta\in\Theta^M$ provides a bad approximation of $E_{\mathscr{P}_t}[J(\cdot,\bar{\gamma})]$, i.e., $|\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\bar{\gamma})]-E_{\mathscr{P}_t}[J(\cdot,\bar{\gamma})]|$ is large. This may happen a fortiori if we evaluate the supremum over Γ of this quantity. Thus, in general, the best that one can hope for is that Eq. (3) holds true with high probability in the space Θ^M and not for all multisamples $\theta=(\vartheta_1,\vartheta_2,\ldots,\vartheta_M)$. This is precisely formalized as follows. Set

$$q(M,\varepsilon):=\mathscr{P}_t^M\left\{\theta:\ \sup_{\gamma\in\Gamma}|\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]-E_{\mathscr{P}_t}[J(\cdot,\gamma)]|>\frac{\varepsilon}{2}\right\}. \tag{5}$$

We say that $J(\cdot,\cdot)$ has the uniform convergence of empirical means (UCEM) property if $q(M,\varepsilon)\to 0$, as $M\to\infty$, $\forall\varepsilon>0$. This is a fundamental property since it implies that we can select M so that Eq. (3) holds with probability arbitrarily close to 1.

In the last two decades, the UCEM property for general classes of functions has been largely studied in the statistical learning literature and general conditions for this property to hold are now available (see e.g. [15,16,26,27,30]). For our purposes the main result is that the UCEM property of $J(\cdot,\cdot)$ holds if the function $J(\cdot,\cdot)$ takes values in $[0,1]$ (which can be guaranteed by rescaling any J' to $J=\beta J'/(1+\beta J')$ with $\beta>0$) and the class $J(\cdot,\gamma)$ of functions from Θ to $[0,1]$ parameterized by $\gamma\in\Gamma$ has finite Pollard(P)-dimension (cf. [30, p. 74]). Moreover, letting d be the P-dimension, $q(M,\varepsilon)$ defined in (5) is upper bounded as follows [30, Theorem 7.1]:

$$q(M,\varepsilon)\leqslant 8\left(\frac{32\mathrm{e}}{\varepsilon}\ln\frac{32\mathrm{e}}{\varepsilon}\right)^d\mathrm{e}^{-M\varepsilon^2/128}. \tag{6}$$

By choosing M such that the right-hand side in (6) is smaller than or equal to δ, we are sure that $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$ is an uniformly good approximant of $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$ to accuracy $\varepsilon/2$ with confidence $1-\delta$. Thus, in view of the discussion after Eq. (3), minimizing $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$ leads to an approximate minimizer of $E_{\mathscr{P}_t}[J(\cdot,\gamma)]$ according to Definition 2 in Section 2.

3.2. The general algorithm for cautious control

Based on the previous discussion, we are now in a position to state the following adaptive control algorithm.

Algorithm 1. Let d be the P-dimension of $\{J(\cdot,\gamma), \gamma \in \Gamma\}$. Given $\varepsilon > 0$ and $\delta > 0$, at any instant t do the following:

1. compute the probability distribution $\mathscr{P}_t$;
2. extract at random $M(\varepsilon,\delta) \geqslant 128/\varepsilon^2[\ln 8/\delta + d \ln 32\mathrm{e}/\varepsilon + d \ln\ln 32\mathrm{e}/\varepsilon]$ independent model parameters $\vartheta_1, \vartheta_2, \ldots, \vartheta_{M(\varepsilon,\delta)}$ according to probability $\mathscr{P}_t$ (the bound on $M(\varepsilon,\delta)$ follows straightforwardly from (6) with $q(M,\varepsilon) = \delta$);
3. compute $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)] = (1/M(\varepsilon,\delta)) \sum_{i=1}^{M(\varepsilon,\delta)} J(\vartheta_i,\gamma)$;
4. choose $\gamma_t = \arg\min_{\gamma\in\Gamma} \hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$;
5. apply the controller $K(\gamma_t, z^{-1})$.

As a consequence of how γ_t is chosen in point 2, Algorithm 1 has robustness properties. Moreover, as time goes by, uncertainty on the true system description is reduced by the observations and the adaptive controller better tunes to the true plant characteristics. Under suitable conditions, this eventually leads to a self-tuning controller (see Section 6).

Two issues need to be addressed for the actual implementation of Algorithm 1, namely:

(a) the updating of $\mathscr{P}_t$;
(b) the computation of the Pollard-dimension of $\{J(\cdot,\gamma), \gamma \in \Gamma\}$.

Such issues are dealt with in the next two sections.

4. Updating $\mathscr{P}_t$

As it is expected, no general recipe for computing the probability distribution $\mathscr{P}_t$ can be given. In this section, we discuss two different situations.

Situation 1. When the system has a very specific structure, $\mathscr{P}_t$ can be computed at a reasonable computational cost as the a posteriori probability of the system parameter given observations.

A standard case is when the system is described by equation

$$\mathscr{A}(\vartheta^\circ, z^{-1}) y_{t+1} = \mathscr{B}(\vartheta^\circ, z^{-1}) u_t + n_{t+1}, \tag{7}$$

where $\mathscr{A}(\vartheta^\circ,z^{-1}) = 1 - \sum_{i=1}^{n} a_i^\circ z^{-i}$ and $\mathscr{B}(\vartheta^\circ,z^{-1}) = \sum_{i=1}^{m} b_i^\circ z^{-(i-1)}$, $\vartheta^\circ = [a_1^\circ,\ldots,a_n^\circ,\ b_1^\circ,\ldots,b_m^\circ]^{\mathrm{T}} \in \Theta = \Re^{n+m}$, and $\{n_t\}$ is a white process with Gaussian $\mathscr{G}(0,\sigma^2)$ distribution. As is well known, in this case, if ϑ° is assumed to be Gaussian, computing $\mathscr{P}_t$ is a finite dimensional estimation problem that is solved by the Kalman filter [9,17].

Situation 2. Unfortunately, in many cases assuming Gaussian distribution can be too restrictive. Then, the problem of estimating the distribution $\mathscr{P}_t$ becomes infinite dimensional and one has to compromise precision in estimation towards computational effort.

Suppose again that the model class is rich enough so as to include an accurate description of the true system, but drop the Gaussianity assumption. Then, the asymptotic theory of prediction error identification methods can be used for computing $\mathscr{P}_t$. In fact, under mild regularity conditions, $\mathscr{P}_t$ tends asymptotically to be Gaussian with the nominal model as mean and a variance which can be estimated from data (see [18,19]).

5. Computing the Pollard-dimension

It is a known fact that the computation of a P-dimension on the basis of its definition is a hard task. On the other hand, an upper bound for the P-dimension can be determined by appealing to some recent and very powerful results developed in [15,16]. The results of [15,16] have been applied to cost criteria arising in connection with stabilization problems and H_∞ control problems in [31].

In Section 7, as an application example, we present a randomized algorithm for cautious LQG adaptive control. There, among other things, we shall compute the P-dimension for this specific control problem.

6. Self-tuning properties

In this section, we study the tuning properties of Algorithm 1 in the ideal case when the model class is sufficiently rich to include the system description.

At each time t, Algorithm 1 selects controller $K(\gamma_t, z^{-1})$, where γ_t minimizes the average of $J(\vartheta,\gamma)$ over the set of model parameters $\vartheta_1,\dots,\vartheta_{M(\varepsilon,\delta)}$ extracted from Θ according to $\mathscr{P}_t$. Suppose that $\mathscr{P}_t$ becomes sharply peaked around ϑ° at a sufficiently fast rate so that $\vartheta_1,\dots,\vartheta_{M(\varepsilon,\delta)}$ get arbitrarily close to the true parameter ϑ° as $t\to\infty$. Then, under appropriate regularity conditions on the control cost $J(\cdot,\cdot)$, $K(\gamma_t, z^{-1})$ is expected to self-tune according to Definition 3 in Section 2.

Here, we determine sufficient conditions for the self-tuning result to be valid in Situation 1 of Section 4, that is in the case when the true system is described by (7) where $\{n_t\}$ is a white Gaussian noise with zero mean and variance σ^2. From Section 4 we know that the a posteriori distribution $\mathscr{P}_t$ is Gaussian, say $\mathscr{G}(\hat{\vartheta}_t, V_t)$. Also, in [5] it is shown that, for all ϑ° apart from a set $\mathscr{N}$ with zero Lebesgue measure, $\hat{\vartheta}_t \to \vartheta^\circ$ with probability 1, provided that $V_t \to 0$ with probability 1. Yet, in order to prove tuning properties, some extra care to convergence rate issues must be paid in the present context that do not enter the game in adaptive control schemes based on the certainty equivalence principle.

We start by proving a proposition which states that the random extractions ϑ_i stay close to ϑ° for t large when $V_t\to 0$ at a certain rate. Throughout we assume that $\vartheta^\circ \notin \mathscr{N}$.

Proposition 1. *Suppose that* $\|V_t\| = \mathrm{O}(1/t^\mu)$ *with probability* 1 *for some* $\mu>0$. *Fix* $\rho>0$. *Then, with probability* 1, *there exists* $\bar{t}$ *such that, for any* $t\geqslant\bar{t}$, *the* $M(\varepsilon,\delta)$ *model parameters* $\vartheta_1,\dots,\vartheta_{M(\varepsilon,\delta)}$ *extracted at step* 2 *of Algorithm* 1 *according to distribution* $\mathscr{P}_t \sim \mathscr{G}(\hat{\vartheta}_t, V_t)$ *belong to the* $(n+m)$*-ball* $B(\vartheta^\circ,\rho)$ *of center* ϑ° *and radius* ρ.

Proof. We start by observing that assumption $\|V_t\| = \mathrm{O}(1/t^\mu)$ with probability 1 entails that function $t^\mu\|V_t\|$ is pathwise bounded with probability 1. Thus, for any fixed $v>0$, one can determine a positive constant k such that

$$\Pr\left\{\sup_t t^\mu\|V_t\| > k\right\} \leqslant v. \tag{8}$$

On the other hand, condition $\hat{\vartheta}_t\to\vartheta^\circ$ with probability 1 implies the existence of an instant point t_1 such that

$$\Pr\left\{\sup_{t\geqslant t_1}\|\hat{\vartheta}_t - \vartheta^\circ\| > \rho/2\right\} \leqslant v. \tag{9}$$

We shall now compute the probability of the event in which at least one of the parameters $\vartheta_1,\vartheta_2,\dots,\vartheta_{M(\varepsilon,\delta)}$ extracted at point 2 of Algorithm 1 at the generic instant $t\geqslant t_1$ is outside the ball $B(\vartheta^\circ,\rho)$, conditioned to event $A := \{\sup_t t^\mu\|V_t\|\leqslant k\}\cap\{\sup_{t\geqslant t_1}\|\hat{\vartheta}_t-\vartheta^\circ\|\leqslant\rho/2\}$. Then, the thesis will be drawn by using this result in conjunction with estimates (8) and (9).

Let $B_t := \{\vartheta_1 \notin B(\vartheta^\circ,\rho)$ or $\dots$ or $\vartheta_{M(\varepsilon,\delta)} \notin B(\vartheta^\circ,\rho)$ at time $t\}$.
Fix a matrix $V \in \Re^{(n+m)\times(n+m)}$ and a vector $\vartheta \in \Re^{n+m}$ such that $\|\vartheta-\vartheta^\circ\| \leqslant \rho/2$. We have

$$\begin{aligned}
&\Pr\{B_t|\hat\vartheta_t=\vartheta, V_t=V\}\\
&\quad\leqslant M(\varepsilon,\delta)\Pr\{\vartheta_1 \notin B(\vartheta^\circ,\rho) \text{ at time t}|\hat\vartheta_t=\vartheta, V_t=V\}\\
&\quad= M(\varepsilon,\delta)\Pr\{\|\vartheta_1-\vartheta^\circ\| > \rho \text{ at time t}|\hat\vartheta_t=\vartheta, V_t=V\}\\
&\quad\leqslant M(\varepsilon,\delta)\Pr\{\|\vartheta_1-\hat\vartheta_t\| > \rho/2 \text{ at time t}|\hat\vartheta_t=\vartheta, V_t=V\}.
\end{aligned}\tag{10}$$

Probability $\Pr\{\|\vartheta_1-\hat\vartheta_t\| > \rho/2$ at time t$|\hat\vartheta_t=\vartheta, V_t=V\}$ can be easily estimated by observing that parameter ϑ_1 is extracted at random from a Gaussian probability distribution with mean $\hat\vartheta_t$ and variance V_t. Letting σ_i^2, $i=1,2,\dots,n+m$ be the eigenvalues of matrix V, we have:

$$\begin{aligned}
&\Pr\{\|\vartheta_1-\hat\vartheta_t\| > \rho/2 \text{ at time t}|\hat\vartheta_t=\vartheta, V_t=V\}\\
&\quad= \int_{\|z\|>\rho/2}\prod_{i=1}^{n+m}\frac{1}{(2\pi)^{1/2}\sigma_i}\exp\left(-\frac{z_i^2}{2\sigma_i^2}\right)\mathrm{d}z\\
&\quad\leqslant \int_{z:\ \exists i\in\{i=1,\dots,n+m\}\ s.t.\ |z_i|>\rho/\{2(n+m)^{1/2}\}}\prod_{i=1}^{n+m}\frac{1}{(2\pi)^{1/2}\sigma_i}\exp\left(-\frac{z_i^2}{2\sigma_i^2}\right)\mathrm{d}z\\
&\quad\leqslant \sum_{i=1}^{n+m}\int_{|w_i|>\rho/\{[8(n+m)]^{1/2}\sigma_i\}}\frac{1}{(2\pi)^{1/2}}\exp(-w_i^2)2^{1/2}\,\mathrm{d}w_i \quad (\text{letting } w_i=z_i/(2^{1/2}\sigma_i))\\
&\quad\leqslant \sum_{i=1}^{n+m}2\int_{w_i^2>\rho^2/\{8(n+m)\sigma_i^2\}}\frac{2(n+m)^{1/2}\sigma_i}{(2\pi)^{1/2}\rho}\exp(-w_i^2)\,\mathrm{d}w_i^2\\
&\quad= \sum_{i=1}^{n+m}\frac{4(n+m)^{1/2}\sigma_i}{(2\pi)^{1/2}\rho}\exp\left(-\frac{\rho^2}{8(n+m)\sigma_i^2}\right).
\end{aligned}\tag{11}$$

Plugging (11) into (10), we finally obtain

$$\Pr\{B_t|\hat\vartheta_t=\vartheta, V_t=V\} \leqslant M(\varepsilon,\delta)\sum_{i=1}^{n+m}\frac{4(n+m)^{1/2}\sigma_i}{(2\pi)^{1/2}\rho}\exp\left(-\frac{\rho^2}{8(n+m)\sigma_i^2}\right).\tag{12}$$

We are now in a position to bound the probability of event B_t given A. Since in set A we have $\|V_t\| \leqslant k/t^\mu$—and, thereby, all eigenvalues of matrix V_t are bounded by k/t^μ—from (12) we can conclude that $\sum_{t\geqslant t_1}\Pr\{B_t|A\} < \infty$. The thesis can now be easily proven by using this result along with (8) and (9).

Since $\sum_{t=0}^{\infty}\Pr(B_t\cap A) \leqslant \sum_{t=0}^{\infty}\Pr(B_t|A) \leqslant t_1+\sum_{t\geqslant t_1}\Pr(B_t|A) < \infty$, by Borel-Cantelly lemma, [10], we have (i.o.=infinitely often) $\Pr(B_t\cap A \text{ i.o.})=0$. On the other hand, by (8) and (9), $\Pr(A) > 1-2v$ and so, owing to the arbitrariness of v, $\Pr(B_t \text{ i.o.})=0$. This proves that with probability 1 one can determine a $\bar t$ such that all model parameters $\vartheta_1,\dots,\vartheta_{M(\varepsilon,\delta)}$ selected at point 2 of Algorithm 1 are in the ball $B(\vartheta^\circ,\rho)$, $\forall t \geqslant \bar t$. □

In the following theorem we show that if $J(\cdot,\gamma)$ is continuous in ϑ°, uniformly in γ, then the controller selected by Algorithm 1 self-tunes when $\mathscr{P}_t$ shrinks around ϑ° at the rate described in Proposition 1. The uniform continuity property is somehow restrictive and is assumed here for convenience. Other less restrictive—though more complicated to formulate—regularity conditions could also be given.

Theorem 1. *Suppose that $J(\cdot,\gamma)$ is continuous in ϑ°, uniformly in $\gamma \in \Gamma$. If $\|V_t\| = \mathrm{O}(1/t^\mu)$ with probability* 1 *for some $\mu > 0$, then the controller computed in Algorithm* 1 *self-tunes according to Definition* 3 *in Section* 2.

Proof. Set $\gamma^\circ := \arg\min_{\gamma\in\Gamma} J(\vartheta^\circ,\gamma)$ (we assume for simplicity it exists).

Due to the uniform continuity assumption of $J(\cdot,\gamma)$ in ϑ°, for any fixed $\alpha > 0$, there exists $\rho > 0$ such that

$$|J(\vartheta,\gamma) - J(\vartheta^\circ,\gamma)| \leqslant \alpha, \quad \forall \vartheta \in B(\vartheta^\circ,\rho), \ \forall \gamma \in \Gamma. \tag{13}$$

By Proposition 1, we know that with probability 1 we can find a time $\bar{t}$ such that the model parameters $\vartheta_1,\ldots,\vartheta_{M(\varepsilon,\delta)}$ extracted at point 2 of Algorithm 1 satisfy $\vartheta_j \in B(\vartheta^\circ,\rho)$, $j = 1,\ldots,M(\varepsilon,\delta)$, for all $t \geqslant \bar{t}$. Then, the following holds for $t \geqslant \bar{t}$. First, Eq. (13) computed in $\gamma = \gamma^\circ$ implies that

$$J(\vartheta_j,\gamma^\circ) \leqslant J(\vartheta^\circ,\gamma^\circ) + \alpha, \quad j = 1,\ldots,M(\varepsilon,\delta). \tag{14}$$

From the definition of γ_t and Eq. (14), it then follows that there exists at least a j such that $J(\vartheta_j,\gamma_t) \leqslant J(\vartheta^\circ,\gamma^\circ) + \alpha$. Since (13) with $\gamma = \gamma_t$ entails that $J(\vartheta^\circ,\gamma_t) \leqslant J(\vartheta_j,\gamma_t) + \alpha$, we can conclude that $J(\vartheta^\circ,\gamma_t) \leqslant J(\vartheta^\circ,\gamma^\circ) + 2\alpha$. Thus $\limsup_{t\to\infty} J(\vartheta^\circ,\gamma_t) \leqslant J(\vartheta^\circ,\gamma^\circ) + 2\alpha$. Due to the arbitrariness of α, it follows that $\lim_{t\to\infty} J(\vartheta^\circ,\gamma_t) = J(\vartheta^\circ,\gamma^\circ)$. □

It is important to note that, in adaptive control applications, there is in general no guarantee that the input signal is sufficiently exciting in such a way that it probes the system and reduces uncertainty. In particular, in our context, the cautious control law of Algorithm 1 can fail to ensure the desired convergence rate of $\|V_t\|$ to zero. In [8, Theorem 6.2] this issue is addressed in the context of certainty equivalence control, and it is shown that, under a certain growth assumption on the input sequence $\{u_t\}$, the required rate can be obtained by adding to the control variable u_t an asymptotically vanishing dither noise. This approach can be used in our context as well. Details are omitted and the interested reader is referred to [7,23] for more discussion.

7. Application example: LQG control

Consider the LQG control problem where a system of the form (7) has to be controlled so as to minimize the cost function

$$J' = \limsup_{N\to\infty} \frac{1}{N} \sum_{t=0}^{N-1} [y_t^2 + r u_t^2], \quad r > 0. \tag{15}$$

In the sequel, we shall refer to some basic facts on LQG control that are reviewed in the appendix.

Suppose that the a priori distribution $\mathscr{P}$ is Gaussian, i.e., $\mathscr{P} \sim \mathscr{G}(\bar{\vartheta},\bar{V})$. Then, from the discussion in Section 4 we know that the Gaussian distribution $\mathscr{P}_t$ can be computed by a Kalman filter initialized with $\bar{\vartheta}$ and $\bar{V}$.

The feasible controller set is described by equation $\mathscr{C}(\gamma,z^{-1})u_t = \mathscr{D}(\gamma,z^{-1})y_t$, with $\mathscr{C}(\gamma,z^{-1}) = 1 - \sum_{i=1}^{m-1} c_i z^{-i}$ and $\mathscr{D}(\gamma,z^{-1}) = \sum_{i=0}^{n-1} d_i z^{-i}$, where the parameter γ is defined as $\gamma := [d_0,\ldots,d_{n-1},c_1,\ldots,c_{m-1}]$ and takes values in $\Gamma = \Re^q$, with $q := n + m - 1$. Note that such a set includes the optimal LQG controller for the true system.

The LQG performance $J'(\vartheta,\gamma)$ achieved for model (7) with parameter ϑ controlled by the controller with parameter γ is obtained by solving a Lyapunov equation in the case when the closed-loop system is stable.

In order to apply the cautious control approach, we rescale $J'(\vartheta,\gamma)$ as follows

$$J(\vartheta,\gamma)=\frac{\beta J'(\vartheta,\gamma)}{1+\beta J'(\vartheta,\gamma)},$$

where β is a real positive number. When the controller destabilizes the model, we set J equal to its maximal value ($J(\vartheta,\gamma)=1$). Thus,

$$J(\vartheta,\gamma)=\begin{cases}\dfrac{\beta J'(\vartheta,\gamma)}{1+\beta J'(\vartheta,\gamma)} & \text{if controller } \gamma \text{ stabilizes model } \vartheta,\\ 1 & \text{otherwise.}\end{cases} \tag{16}$$

As for the Pollard-dimension of the family of functions $J(\cdot,\gamma)$ parameterized by $\gamma\in\Re^q$, the following result is proven in the appendix.

Theorem 2.

$$\text{P-dimension}(\{J(\cdot,\gamma),\gamma\in\Re^q\})\leqslant\begin{cases}2\log_2(16\mathrm{e}) & \textit{if } q=1,\\ 2q\log_2(16\mathrm{e}q(2q+1)) & \textit{otherwise.}\end{cases}$$

We next compare the performance achieved when applying the standard certainty equivalence and the cautious adaptive controllers in a simple numerical example. This is in order to point out the effectiveness of the proposed cautious approach and to get insight into such a control strategy.

7.1. A simulation example

Consider the model class

$$y_{t+1}=a_1y_t+b_1u_t+b_2u_{t-1}+n_{t+1},\quad \vartheta:=[a_1\ b_1\ b_2]^{\mathrm{T}}\in\Re^3, \tag{17}$$

where $\{n_t\}$ is $\mathscr{G}(0,1)$. The true system has parameter $\vartheta^\circ=[0.8\ 1-0.9]^{\mathrm{T}}$. For the estimation of $\mathscr{P}_t$, we assume the a priori distribution $\mathscr{P}\sim\mathscr{G}(\bar\vartheta,\bar V)$ with $\bar\vartheta=[-0.4\ 1.4-0.1]^{\mathrm{T}}$ and $\bar V=I_3$.

In (16), we take $\beta=0.25$ and the control cost coefficient r is set to 1. Finally, the controller set is $\{u_t=d_0y_t+c_1u_{t-1},\ \gamma=[d_0\ c_1]\in\Gamma=\Re^2\}$.

In Fig. 1, the control system behavior is expressed in terms of the "sample LQG performance index" $(1/t)\sum_{i=0}^{t-1}[y_i^2+u_i^2]$. Graph (a) shows that in the certainty equivalence case the input and output variables assume high values in the transient phase. Such an undesirable phenomenon is due to the fact that the certainty equivalent controller puts an infinite trust in the most probable model. In the transient phase, such a model may be a poor description of the system, which may ultimately result in an excessive control action applied to the system or even a transiently unstable closed-loop system.

As shown by graph (b), the cautious adaptive controller is able to overcome the above difficulty and it significantly enhances the transient behavior with respect to the certainty equivalence controller. However, the control system performance obtained by applying the cautious adaptive controller are strictly suboptimal even in the long run. The reason why this happens is that the cautious control has no probing features and, in fact, it tends not to excite the unknown dynamics of the true system. Thus, the controller parameter are not consistently estimated.

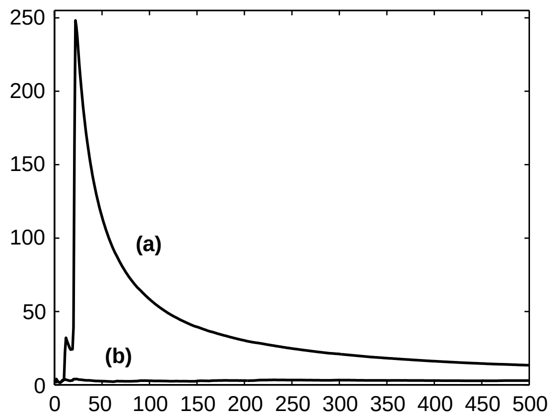

Fig. 1. Sample LQG performance index: (a) certainty equivalence control; (b) cautious adaptive control ($\varepsilon = \delta = 0.1$).

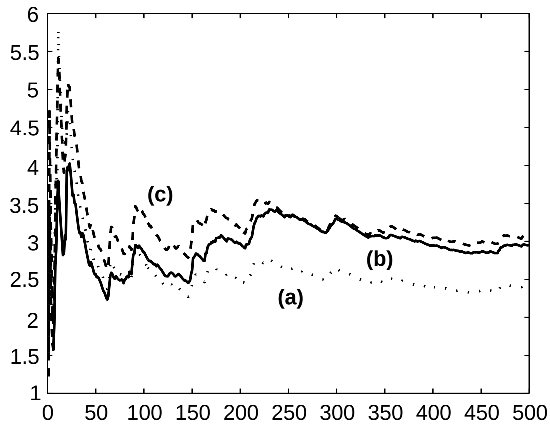

Fig. 2. Sample LQG performance index: (a) cautious adaptive control with dither; (b) cautious adaptive control; (c) cautious robust control. For all control laws, $\varepsilon = \delta = 0.1$.

In order to overcome this difficulty a dither noise can be added to the cautious control input. If the dither noise is suitably selected, then one can enforce optimality. The adopted control law with dither takes the form

$$u_t = d_0 y_t + c_1 u_{t-1} + \text{dither}_t,$$

where $\{\text{dither}_t\}$ is a white process, independent of $\{n_t\}$, such that its distribution at time t is uniform in $[-\sqrt{0.3}/(t+1)^{1/15}, \sqrt{0.3}/(t+1)^{1/15}]$.

The control system behavior obtained by applying the cautious LQG adaptive control algorithm with dither noise is displayed in Fig. 2. Fig. 2 also shows the cautious LQG adaptive control without dither (same as (b) in Fig. 1) for easy comparison. The controller with dither outperforms that without dither and is in fact asymptotically optimal (the sample cost tends to 2.13—the optimal cost—when $t \to \infty$). As an additional element of comparison, in Fig. 2 the performance achieved by the cautious robust controller—i.e., the controller designed on the basis of the initial distribution $\mathscr{P}$ and maintained fixed through time—is also represented.

8. Discussion and open problems: minimizing $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$

In this section we discuss the issues involved in the minimization of $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$ over the set Γ.

Typically, Γ is a continuum set of parameterizations and $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$ is a non-convex function of the controller parameter γ. Consequently, its minimization is a non-trivial task whenever the space in which Γ is embedded is high dimensional.

In this section, we discuss two methods for addressing this issue. The key idea in both cases is to focus attention on a finite subset of the controller parameter set and perform an exhaustive search on it.

8.1. Randomized minimization

Introduce a probability distribution $\mathscr{Q}_t$ on the controller parameter set Γ. Then, in order to minimize $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$, we can resort to the following stochastic algorithm, previously introduced in [20,25]:

1. extract at random N independent controller parameters $\gamma_1,\gamma_2,\ldots,\gamma_N$ according to $\mathscr{Q}_t$;
2. choose $\gamma_t = \arg\min_{\gamma\in\{\gamma_j\}_{j=1}^N} \hat{E}_{\mathscr{P}_t}[J(\cdot,\gamma)]$.

Clearly, since we are testing just N values of γ, we cannot expect that γ_t is the global minimizer of $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$ over Γ. In addition, the quality of the result is random due to the stochastic selection of parameters γ_j's. Nevertheless, the following quantitative statement can be proven showing that γ_t can be considered optimal in some probabilistic sense (see [30, Lemma 11.1] for a proof of this result): fix an arbitrary real number $\alpha > 0$. Then,

$$\mathscr{Q}_t\{\gamma\colon \hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)] < \hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma_t)]\} \leqslant \alpha \tag{18}$$

with probability at least equal to $1-(1-\alpha)^N$.

Here, the specification "with probability at least equal to $1-(1-\alpha)^N$" makes reference to the probability involved in the random extractions of the γ_j's. Note that Eq. (18) does not prevent the existence of γ's such that $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)] < \hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma_t)]$. It only requires that this happens in an exceptional set in Γ whose probability $\mathscr{Q}_t$ is upper bounded by α. From this, we see that parameter α quantifies the difficulty of improving $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$ over $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma_t)]$ by resorting to a stochastic selection (based on probability $\mathscr{Q}_t$) of the controller parameter γ. On the other hand, a decision maker having at his disposal a perfect deterministic minimization procedure might be able to gain a large improvement over $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma_t)]$. This is more likely to happen if the controller parameter space is high dimensional, since in this case the $\mathscr{Q}_t$ measure of the nearly optimal controller parameter set is generally small (curse of dimensionality).

In the following points (i) and (ii), we indicate two possible ways to overcome the difficulty posed by the presence of the exceptional set of measure α.

(i) *A wise selection of the probability* $\mathscr{Q}_t$: A possibility is to select the probability distribution $\mathscr{Q}_t$ to be the probability induced on Γ by $\mathscr{P}_t$ through the function $\vartheta \to \gamma^\circ(\vartheta)$ that maps each ϑ in the parameterization γ that is optimal for the model with parameter ϑ. In mathematical terms, this corresponds to set $\mathscr{Q}_t(A) = \mathscr{P}_t\{\vartheta\colon \gamma^\circ(\vartheta)\in A\}$, $\forall A\subseteq\Gamma$ (measurability issues are neglected here). This choice of $\mathscr{Q}_t$ leads to a selection of parameters γ_j's better suited to the control situation at hand at least in the long run, when $\mathscr{P}_t$ becomes more sharply peaked around the true plant parameter ([7,23]).

(ii) *A structured controller*: Consider a controller composed by two parts: a nominal controller plus a detuning filter (e.g. a low pass filter) with very few (say 1 or 2) parameters. In this setting one can conceive to determine the controller according to a two-step procedure. The nominal controller is first tuned according to the certainty equivalence principle; then, the detuning filter is selected in a cautious fashion, according to the average cost approach so as to decrease the crossover frequency of the control system. This way, due to the low dimensionality of γ, the problem of tuning the cautious controller parameter is easy to solve. One can e.g. use the stochastic minimization method where $\mathscr{Q}_t$ can be simply chosen to be the uniform distribution, or sweep the entire Γ set in some deterministic fashion. This approach has been adopted in [4].

8.2. Predefined finite controller covering

The idea here is to reduce the feasible controller set from the very beginning to an appropriately designed finite set so as to overcome the difficulties involved in the minimization of $\hat{E}_{\mathscr{P}_t,M}[J(\cdot,\gamma)]$, while not degradating too much the best achievable level of performance. Such an approach is inspired by the literature on switching adaptive control (see e.g. [14,21,24]), where the need for a finite set of candidate controllers has been motivated from both a practical and a theoretical point of view.

In a switching control scheme, at event-driven time instants, a supervisory control system decides which one of the candidate controllers in a feasible finite set should be placed in feedback with the system so as to achieve a certain performance level. In particular, in the estimator-based approach to switching control, the supervisor relies on a model set for the controller selection. Specifically, it switches in the loop the candidate controller that is optimal for the model better resembling the system behavior according to some identification criterion. The feasible set then has to be designed so that for each model there is a candidate controller which ensures a certain known performance level when placed in closed-loop with it. Hence, the name "finite controller cover" for the feasible controller set.

In our cautious control context, we suggest to consider many controller covers at the same time. The first controller cover corresponds to controllers with a high robustness level, while penalizing performance. Subsequent covers correspond to controllers with progressively increasing performance at the price of a decreasing level of robustness. All the controllers in these covers are put together in a single large set and the controller selection within this set is performed by means of an exhaustive search based on our randomized method. When $\mathscr{P}_t$ is spread over the model parameter set Θ, the chosen controller will belong to a cover with high robustness and low performance, while, as $\mathscr{P}_t$ becomes more concentrated around the true system parameter, controllers belonging to covers with lower robustness and higher performance will be automatically preferred.

The main difficulty in adopting this approach is determining the finite controller cover. This issue has in fact been addressed only recently in the literature. In [1], the existence of a finite controller cover has been proven for H_∞ control and a compact Θ, and a procedure for determining it has been presented. Similar results have been shown also in [22]. Yet, in many cases a computationally efficient method to determine a finite cover may be difficult to work out.

Preliminary results along this multi-cover approach are reported in [6].

9. Conclusions

In this paper, we have introduced a general methodology for adaptive control that combines randomized methods for the minimization of average cost criteria with the updating of the probability distribution representing the uncertainty on the system description. After introducing the main ideas in a general control setup, we have concentrated on LQG control. A simulation example has illustrated the efficacy of the method.

Still, as discussed in the paper, important open issues need to be further investigated. This includes:

- the study of how $\mathscr{P}_t$ can be updated and, in particular, how the asymptotic theory of system identification can be used in this context;
- unmodeled dynamics have not been considered in the updating of $\mathscr{P}_t$. For real applicability of the method, the presence of unmodeled dynamics has somehow to be taken care of;

- the combined use of cautious randomized control and switching control represents a way to overcome existing difficulties of both methods, while preserving their positive features. This study is still at the outset, though.

These points represent a stimulus for future research.

Appendix A. Optimal LQG control

We next review basic facts on the solution to the optimal LQG control introduced in Section 7. The interested reader is referred to, e.g., [3,8] for a more comprehensive treatment.

Suppose that $n > 0$ (for $n=0$, the optimal control law is trivially $u_t = 0$, $t \geqslant 0$). Set $x_t := [y_t, \ldots, y_{t-(n-1)}, u_{t-1}, \ldots, u_{t-(m-1)}]^{\mathrm{T}}$. Then, model (7) with parameter ϑ can be given the state space representation:

$$x_{t+1} = A(\vartheta)x_t + B(\vartheta)u_t + Cn_{t+1}, \qquad y_t = Hx_t, \tag{A.1}$$

where

$$A(\vartheta) = \begin{bmatrix} a_1 & \ldots & a_{n-1} & a_n & b_2 & \ldots & b_{m-1} & b_m \\ 1 & 0 & \ldots & & & & & \\ & \ddots & \ddots & & & & & \\ & & 1 & 0 & & & & \\ 0 & \ldots & \ldots & \ldots & 0 & \ldots & \ldots & 0 \\ & & & & 1 & 0 & & \\ & & & & & \ddots & \ddots & \\ & & & & & & 1 & 0 \end{bmatrix}, \quad B(\vartheta) = \begin{bmatrix} b_1 \\ 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \quad C = H^{\mathrm{T}} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \\ 0 \\ \vdots \\ \vdots \\ 0 \end{bmatrix}.$$

If system (A.1) is stabilizable and detectable, then there exists a unique positive semidefinite solution $P(\vartheta)$ to the discrete time algebraic Riccati equation

$$P = A(\vartheta)^{\mathrm{T}}[P - PB(\vartheta)(B(\vartheta)^{\mathrm{T}}PB(\vartheta) + r)^{-1}B(\vartheta)^{\mathrm{T}}P]A(\vartheta) + H^{\mathrm{T}}H$$

and the optimal LQG control law is given by $u_t = \gamma^\circ(\vartheta)x_t$, where $\gamma^\circ(\vartheta) = -(B(\vartheta)^{\mathrm{T}}P(\vartheta)B(\vartheta) + r)^{-1} B(\vartheta)^{\mathrm{T}}P(\vartheta)A(\vartheta)$. Letting $\gamma^\circ(\vartheta) = [d_0^\circ(\vartheta), \ldots, d_{n-1}^\circ(\vartheta),\ c_1^\circ(\vartheta), \ldots, c_{m-1}^\circ(\vartheta)]$, the optimal LQG control can be rewritten in the input–output form $\mathscr{C}(\gamma^\circ(\vartheta), z^{-1})\,u_t = \mathscr{D}(\gamma^\circ(\vartheta), z^{-1})\,y_t$, where $\mathscr{C}(\gamma^\circ(\vartheta), z^{-1}) = 1 - \sum_{i=1}^{m-1} c_i^\circ(\vartheta)z^{-i}$ and $\mathscr{D}(\gamma^\circ(\vartheta), z^{-1}) = \sum_{i=0}^{n-1} d_i^\circ(\vartheta)z^{-i}$.

Let $\gamma = [d_0, \ldots, d_{n-1}, c_1, \ldots, c_{m-1}]$ and suppose that controller $\mathscr{C}(\gamma, z^{-1})\,u_t = \mathscr{D}(\gamma, z^{-1})\,y_t$ is applied to the system with parameter ϑ. Under the assumption that $A(\vartheta) + B(\vartheta)\gamma$ is stable, the LQG control performance of the closed-loop system where the model with parameter ϑ is controlled by the controller with parameter γ is given by

$$J'(\vartheta, \gamma) = \mathrm{trace}(P(\vartheta, \gamma)CC^{\mathrm{T}}\sigma^2), \tag{A.2}$$

where $P(\vartheta,\gamma)$ is the solution to the Lyapunov equation

$$P=(A(\vartheta)+B(\vartheta)\gamma)^{\mathrm{T}}P(A(\vartheta)+B(\vartheta)\gamma)+r\gamma^{\mathrm{T}}\gamma+H^{\mathrm{T}}H. \tag{A.3}$$

Appendix B. Proof of Theorem 2

Consider the family of functions $\mathscr{F}=\{J(\cdot,\gamma):\Theta=\Re^{q+1}\to[0,1],\gamma\in\Gamma=\Re^{q}\}$. Given a function $J(\cdot,\gamma)$ in $\mathscr{F}$, let $g((\vartheta,c),\gamma):=H(J(\vartheta,\gamma)-c)$, where $c\in[0,1]$ is an additional variable and $H(\cdot)$ is the Heaviside function ($H(x)=1$, if $x\geqslant 0$, $H(x)=0$, if $x<0$). Also, let $\mathscr{G}:=\{g((\cdot,\cdot),\gamma):\Theta\times[0,1]\to\{0,1\},\gamma\in\Gamma\}$. Then, by [30, Lemma 10.1] P-dimension($\mathscr{F}$)= VC-dimension($\mathscr{G}$) (see e.g. [30, p. 69] for the definition of the VC-dimension). Thus, the original problem of computing P-dimension($\mathscr{F}$) is reduced to the one of computing VC-dimension($\mathscr{G}$). This computation can be carried out by resorting to [30, Corollary 10.2] as explained next.

Suppose that $q>1$ (the derivation for the case $q=1$ is simple, and is not developed in detail). Given any set S, let I_S be the indicator function of S. We prove below that $g((\vartheta,c),\gamma)$ can be written as

$$g((\vartheta,c),\gamma)=I_S((\vartheta,c),\gamma), \tag{B.1}$$

where $S\subset\Re^{q+1}\times[0,1]\times\Re^{q}$ is a set with a particular structure. Precisely,

$$S=\text{Boolean formula applied to }\{S_i\}_{i=1}^{2q+1},$$

where S_i, $i=1,\ldots,2q+1$, are sets given by $S_i=\{\tau_i((\vartheta,c),\gamma)>0\}$ with τ_i polynomials in γ whose largest degree is $v=4q$ (a Boolean formula is any set expression containing union, intersection, and complementation). Before proving (B.1)), we note that the statement of Theorem 2 can be obtained from (B.1) by applying the bound on the VC-dimension($\mathscr{G}$) in Corollary 10.2 in [30], which, in our notations, says that VC-dimension($\mathscr{G}$) $\leqslant 2q\log_2(4\mathrm{e}v(2q+1))=2q\log_2(16\mathrm{e}q(2q+1))$.

The proof is now completed by showing (B.1).

Observe first that by definition (16) of $J(\vartheta,\gamma)$ the following equality holds ($^{\mathrm{c}}$ stands for complement):

$$\{J(\vartheta,\gamma)-c\geqslant 0\}=\{A(\vartheta)+B(\vartheta)\gamma\ \text{stable}\}^{\mathrm{c}}\cup\left(\{A(\vartheta)+B(\vartheta)\gamma\ \text{stable}\}\cap\left\{J'(\vartheta,\gamma)\geqslant\frac{c}{(1-c)\beta}\right\}\right),$$

where $J'(\vartheta,\gamma)=\mathrm{trace}(P(\vartheta,\gamma)CC^{\mathrm{T}}\sigma^2)$ (see (A.2)). We then need to consider the sets $\{A(\vartheta)+B(\vartheta)\gamma\ \text{stable}\}$ and $\{\mathrm{trace}(P(\vartheta,\gamma)CC^{\mathrm{T}}\sigma^2)\geqslant c/(1-c)\beta\}$.

Let us start with $\{A(\vartheta)+B(\vartheta)\gamma\ \text{stable}\}$.

According to Schur–Cohn criterion, the roots of

$$\begin{aligned}f_{(\vartheta,\gamma)}(\lambda)&=\det(\lambda I_q-(A(\vartheta)+B(\vartheta)\gamma))\\&=\lambda^q+f_1(\vartheta,\gamma)\lambda^{q-1}+\cdots+f_q(\vartheta,\gamma)\end{aligned}$$

will lie in the open unit circle if and only if the following $2q$ conditions are satisfied:

$$\begin{aligned}&f_{(\vartheta,\gamma)}(1)>0,\\&(-1)^q f_{(\vartheta,\gamma)}(-1)>0,\\&\det(X_i(\vartheta,\gamma)+Y_i(\vartheta,\gamma))>0,\quad i=1,2,\ldots,q-1,\\&\det(X_i(\vartheta,\gamma)-Y_i(\vartheta,\gamma))>0,\quad i=1,2,\ldots,q-1\end{aligned}$$

with

$$X_i := \begin{bmatrix} 1 & f_1 & f_2 & \dots & f_{i-1} \\ & 1 & f_1 & \dots & f_{i-2} \\ & & & & \vdots \\ & & & 1 & f_1 \\ & & & & 1 \end{bmatrix}, \quad Y_i := \begin{bmatrix} & & & & f_q \\ & & & f_q & f_{q-1} \\ & & & & \vdots \\ & f_q & f_{q-1} & \dots & f_{q-i+2} \\ f_q & f_{q-1} & f_{q-2} & \dots & f_{q-i+1} \end{bmatrix},$$

where X_i, Y_i, f_i are short for $X_i(\vartheta,\gamma)$, $Y_i(\vartheta,\gamma)$, $f_i(\vartheta,\gamma)$. Since the degree in γ of $f_i(\vartheta,\gamma)$, $i=1,\dots,q$, is at most 2, such $2q$ conditions can be rewritten in the form $\tau_i(\vartheta,\gamma)>0$, $i=1,2,\dots,2q$, where the largest degree of each polynomial $\tau_i(\vartheta,\gamma)$ as a function of γ is $2(q-1)$.

Therefore $\{A(\vartheta)+B(\vartheta)\gamma \text{ stable}\} = \bigcap_{i=1}^{2q} S_i$, with $S_i=\{\tau_i(\vartheta,\gamma)>0\}$ such that $\text{degree}(\tau_i(\vartheta,\gamma)) \leqslant 2(q-1)$, $i=1,2,\dots,2q$.

Consider now the set $\{\text{trace}(P(\vartheta,\gamma)CC^{\mathrm{T}}\sigma^2) \geqslant c/(1-c)\beta\}$.

Matrix $P(\vartheta,\gamma)$ is the unique solution to the Lyapunov equation (A.3) which can be reformulated as a system of linear equations in the components of P by means of the Kronecker product $\otimes$, [13]. This is explained next.

Given a $h\times k$ matrix V, denote by $\text{vec}(V)$ the $hk\times 1$ column vector $\text{vec}(V)=[v_1^{\mathrm{T}}\ v_2^{\mathrm{T}}\ \dots\ v_k^{\mathrm{T}}]^{\mathrm{T}}$, where v_i is the ith column of V. Eq. (A.3) can be rewritten as follows:

$$\begin{aligned}\text{vec}(P) &= \text{vec}((A(\vartheta)+B(\vartheta)\gamma)^{\mathrm{T}}P(A(\vartheta)+B(\vartheta)\gamma)) + \text{vec}(r\gamma^{\mathrm{T}}\gamma + H^{\mathrm{T}}H)\\ &= ((A(\vartheta)+B(\vartheta)\gamma)^{\mathrm{T}} \otimes (A(\vartheta)+B(\vartheta)\gamma)^{\mathrm{T}})\,\text{vec}(P) + \text{vec}(r\gamma^{\mathrm{T}}\gamma + H^{\mathrm{T}}H),\end{aligned}$$

where we have used the property that, given matrices V, W and Z of appropriate dimensions, $\text{vec}(VWZ) = (Z^{\mathrm{T}}\otimes V)\,\text{vec}(W)$. From this it follows that

$$\text{vec}(P(\vartheta,\gamma))^{\mathrm{T}} = \text{vec}(r\gamma^{\mathrm{T}}\gamma + H^{\mathrm{T}}H)^{\mathrm{T}}(I_{q^2} - (A(\vartheta)+B(\vartheta)\gamma)\otimes(A(\vartheta)+B(\vartheta)\gamma))^{-1}.$$

Notice now that $P(\vartheta,\gamma)CC^{\mathrm{T}}\sigma^2$ has all entries equal to zero except for the first column due to the fact that $C=[1\ 0\ \dots 0]^{\mathrm{T}}$. Thus, $\text{trace}(P(\vartheta,\gamma)CC^{\mathrm{T}}\sigma^2)$ only extracts the 1,1 entry of $P(\vartheta,\gamma)CC^{\mathrm{T}}\sigma^2$.

Letting $W(\vartheta,\gamma) := I_{q^2} - (A(\vartheta)+B(\vartheta)\gamma)\otimes(A(\vartheta)+B(\vartheta)\gamma)$, the condition $\text{trace}(P(\vartheta,\gamma)CC^{\mathrm{T}}\sigma^2) \geqslant c/(1-c)\beta$ can therefore be rewritten as follows:

$$c\det(W(\vartheta,\gamma)) - (1-c)\beta\sigma^2\,\text{vec}(r\gamma^{\mathrm{T}}\gamma + H^{\mathrm{T}}H)^{\mathrm{T}}[\text{ac}(W(\vartheta,\gamma))]_1 \leqslant 0,$$

where $[\text{ac}(W(\vartheta,\gamma))]_1$ is the first column of the algebraic complement transpose of matrix $W(\vartheta,\gamma)$. By using the definition of Kronecker product, one can see that the degree of polynomial $\det(W(\vartheta,\gamma))$ as a function of γ is less than or equal to $4q$. As for polynomial $\text{vec}(r\gamma^{\mathrm{T}}\gamma + H^{\mathrm{T}}H)^{\mathrm{T}}[\text{ac}(W(\vartheta,\gamma))]_1$, it has the same degree $4q$ since all the terms of vector $[\text{ac}(W(\vartheta,\gamma))]_1$ have at most degree $4q-2$.

Therefore, by setting $\tau_{2q+1}((\vartheta,c),\gamma) = c\det(W(\vartheta,\gamma)) - (1-c)\beta\sigma^2\,\text{vec}(r\gamma^{\mathrm{T}}\gamma + H^{\mathrm{T}}H)^{\mathrm{T}}[\text{ac}(W(\vartheta,\gamma))]_1$, we have that $\{\text{trace}(P(\vartheta,\gamma)CC^{\mathrm{T}}\sigma^2) \geqslant c/(1-c)\beta\} = S_{2q+1}^{\mathrm{c}}$ with $S_{2q+1} = \{\tau_{2q+1}((\vartheta,c),\gamma)>0\}$ and $\text{degree}(\tau_{2q+1}((\vartheta,c),\gamma)) = 4q$.

Then, the sought set S is given by $(\bigcap_{i=1}^{2q} S_i)^{\mathrm{c}} \cup ((\bigcap_{i=1}^{2q} S_i)\cap S_{2q+1}^{\mathrm{c}})$, which concludes the proof.

References

[1] B.D.O. Anderson, T.S. Brinsmead, F. de Bruyne, J. Hespanha, D. Liberzon, A.S. Morse, Multiple model adaptive control. I: Finite controller coverings, Internat. J. Robust Nonlinear Control 10 (2000) 909–929 (George Zames Special Issue).

[2] Y. Bar-Shalom, Stochastic dynamic programming: caution and probing, IEEE Trans. Automat. Control AC-26 (1981) 1184–1195.

[3] D.P. Bertsekas, Dynamic Programming and Optimal Control, Vols. I and II, Athena Scientific, Belmont, MA, 1995.

[4] S. Bittanti, M. C. Campi, S. Garatti, Iterative feedback controller design based on average robust control, in: Proceedings of the 15th IFAC World Congress on Automatic Control, Barcelona, Spain, 2002.

[5] M.C. Campi, Adaptive control of non-minimum phase systems, Internat. J. Adapt. Control Signal Process. 9 (1995) 137–149.

[6] M.C. Campi, J. Hespanha, M. Prandini, Cautious hierarchical switching control of stochastic linear systems, in: Proceedings of the Mathematical Theory of Networks and Systems Conference, Notre Dame, IN, 2002.

[7] M.C. Campi, M. Prandini, Randomized algorithms for the synthesis of adaptive controllers, in: Proceedings of the Mathematical Theory of Networks and Systems Conference, Padova, Italy, 1998, pp. 723–726.

[8] H.F. Chen, L. Guo, Identification and Stochastic Adaptive Control, Birkhäuser, Basel, 1991.

[9] H.F. Chen, P. Kumar, J. van Schuppen, On Kalman filtering for conditionally Gaussian systems with random matrices, Systems Control Lett. 13 (1989) 397–404.

[10] Y.S. Chow, H. Teicher, Probability Theory, Springer, Berlin, 1978.

[11] L. Devroye, L. Györfi, G. Lugosi, A Probabilistic Theory of Pattern Recognition, Springer, Berlin, 1996.

[12] G.C. Goodwin, K.S. Sin, Adaptive Filtering, Prediction and Control, Prentice-Hall, Englewood Cliffs, NJ, 1984.

[13] A. Graham, Kronecker Products on Matrix Calculus, with Applications, Halsted Press, Wiley, New York, 1981.

[14] J. Hespanha, D. Liberzon, A.S. Morse, B.D.O. Anderson, T.S. Brinsmead, F. de Bruyne, Multiple model adaptive control. II: Switching, Internat. J. Robust Nonlinear Control 11 (5) (2001) 479–496.

[15] M. Karpinski, A.J. Macintyre, Polynomial bounds for VC dimension of sigmoidal neural networks, in: Proceedings of the 27th ACM Symposium on Theory of Computing, ACM Press, New York, 1995, pp. 200–208.

[16] M. Karpinski, A.J. Macintyre, Polynomial bounds for VC dimension of sigmoidal and general Pfaffian neural networks, J. Comput. System Sci. 54 (1997) 169–176.

[17] R.S. Liptser, A.N. Shiryayev, Statistics of Random Processes, II: Applications, Springer, Berlin, 1977.

[18] L. Ljung, Identification, model validation and control, in: Proceedings of the 36th Conference on Decision and Control, San Diego, CA, 1997, Plenary Lecture.

[19] L. Ljung, System Identification: Theory for the User, Prentice-Hall, Englewood Cliffs, NJ, 1999.

[20] C. Marrison, R. Stengel, The use of random search and genetic algorithms to optimize stochastic robustness functions, in: Proceedings of the American Control Conference, Baltimore, MD, 1994, pp. 1484–1489.

[21] A.S. Morse, Control using logic-based switching, in: A. Isidori (Ed.), Trends in Control, Springer, Berlin, 1995, pp. 69–113.

[22] F.M. Pait, F. Kassab, On a class of switched, robustly stable, adaptive systems, Internat. J. Adapt. Control Signal Process. 15 (3) (2001) 213–238.

[23] M. Prandini, Adaptive LQG control: optimality analysis and robust controller design, Ph.D. Thesis, University of Brescia, February 1998.

[24] M. Prandini, M.C. Campi, Logic-based switching for the stabilization of stochastic systems in presence of unmodeled dynamics, in: Proceedings of the 40th Conference on Decision and Control, Orlando, FL, 2001, pp. 393–398.

[25] L.R. Ray, R.F. Stengel, Stochastic robustness of linear time-invariant control systems, IEEE Trans. Automat. Control AC-36 (1991) 82–87.

[26] V.N. Vapnik, The Nature of Statistical Learning Theory, Springer, Berlin, 1995.

[27] V.N. Vapnik, Statistical Learning Theory, Wiley, New York, 1998.

[28] V.N. Vapnik, A.Y. Chervonenkis, On the uniform convergence of relative frequencies to their probabilities, Theory Probab. Appl. 16 (1971) 264–280.

[29] V.N. Vapnik, A.Y. Chervonenkis, Necessary and sufficient conditions for the uniform convergence of means to their expectations, Theory Probab. Appl. 26 (1981) 532–553.

[30] M. Vidyasagar, A Theory of Learning and Generalization: With Applications to Neural Networks and Control Systems, Springer, Berlin, 1997.

[31] M. Vidyasagar, Randomized algorithms for robust controller synthesis using statistical learning theory, Automatica 37 (10) (2001) 1515–1528.