



A new approach to adaptive control: no nonlinearities[☆]

Daniel E. Miller*

Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1

Abstract

In adaptive control the goal is to design a controller to control an uncertain system whose parameters may be changing with time. Typically the controller consists of an identifier (or tuner) which is used to adjust the parameters of a linear time-invariant (LTI) compensator, and under suitable assumptions on the plant model uncertainty it is proven that good asymptotic behaviour is achieved, such as model matching (for minimum phase systems) or stability. However, a typical adaptive controller does not track time-varying parameters very well, and it is often highly nonlinear, which can result in undesirable behaviour, such as large transients or a large control signal. Furthermore, most adaptive controllers provide only asymptotic tracking, with no ability to design for a pre-specified settling time.

Here we propose an alternative approach, which yields a linear periodic controller. Rather than estimating the plant or compensator parameters, instead we estimate what the control signal would be if the plant parameters were known. In this paper we argue the utility of this approach and then examine the first order case in detail, including a simulation. We also explore the benefits and limitations of the approach.

© 2003 Elsevier Science B.V. All rights reserved.

Keywords: Model reference control; Linear time-varying control; Adaptive control

1. Introduction

Adaptive control is an approach used to deal with plant uncertainty and/or time-varying parameters. The basic idea is to have a controller which tunes itself to the plant being controlled; such controllers can usually be described by a nonlinear time-varying differential or difference equation. An adaptive controller typically consists of an LTI compensator together with an identifier (or tuner) which is used to adjust the compensator parameters; a common approach to tuning is to invoke the *Certainty Equivalence Principle*, whereby it is assumed at each instance of time that the plant parameter estimate is correct and the controller gains are updated accordingly.

The problem of adaptive control has been considered for some time, certainly as far back as the 1950s, e.g. see [4] and [26]. While the initial objective of adaptive control was to track time varying parameters, it soon became apparent that was an overly demanding task. Hence, the problem was simplified to that of controlling the system with *fixed but unknown* parameters, and even that turned out to be hard; it took until the period 1978–1980 for the famous Model Reference Adaptive Control Problem (MRACP) to be solved,

[☆] This work was supported by the Natural Sciences and Engineering Research Council of Canada via a research grant.

* Tel.: +519-885-1211; fax: +519-746-3077.

E-mail address: miller@hobbes.uwaterloo.ca (D.E. Miller).

e.g. see Goodwin et al. [1], Morse [16], and Narendra et al. [19], as well as the nice overview paper by Morse on the history of the problem [17]. At that point a number of researchers turned to the original (and more natural) problem of handling time-varying parameters, but with mixed results. While it has been proven that adaptive control can be carried out with time-varying parameters, typically either a persistently exciting disturbance must be added to the control signal (which degrades the tracking quality) or the time-variations have to be slow or of a known form, e.g. see [8,25,3] and [5], and the references contained therein. As far as the author is aware there are only a few methods which allow for quite general time-variations, each with its own drawbacks:

- The approach of [21] can handle rapidly varying parameters but there is a tradeoff between tracking error accuracy and control signal chattering; furthermore, the method requires knowledge of the sign of the high-frequency gain.
- The approach of [6] requires the reference signal to be the output of an unforced stable exosystem, and good tracking is proven only if the plant parameters are of bounded variation (on $[0, \infty)$).
- The high-gain approach of [13] can potentially be used to handle rapidly varying parameters for minimum phase systems; the significant drawback is that the control signal could get extremely large.

Another property that most adaptive controllers possess is that they are nonlinear, some more so than other, which makes the behaviour hard to predict; the closed loop system can exhibit undesirable behaviour, such as large transients, especially if the initial parameter estimates are poor.¹ In the identifier-based approach to adaptive control, either the plant parameters (indirect) or controller parameters (direct) are first estimated, and then the controller signal is constructed accordingly; the nonlinearity enters two ways—via the estimation and by the multiplication needed to map the parameter estimates to the control signal. Of course, there are other approaches to adaptive control as well, perhaps the most notable being logic based switching. This can be traced back to the mid-1980s, e.g. see [27] and [7], in which the idea is to traverse a set of candidate controllers, one of which will work. This early work was pre-routed, which typically meant poor transients.² In the 1990s researchers started looking at more cleverly designed logic, e.g. see [18,20] and [2]. Even there the controllers are nonlinear, and it is hard to separate the effect of the initial conditions from that of the input.

To sum up, most adaptive controllers to date have several key features:

- (a) they do not track time-varying parameters very well,
- (b) typically only asymptotic results are proved—the transients may be poor,
- (c) they are nonlinear, so the effect of the ICs and the exogenous input are coupled, and
- (d) the control signal can get quite large (in comparison to what the control signal would be if the plant parameters were known and the “correct” LTI controller were applied).

Of course, every rule has its exceptions: some of the controllers in the literature have alleviated one or more of the above, but, as far as the author is aware, none are immune to them all. This calls for a new approach.

Let us return to the essential insight of Certainty Equivalence: the idea is to estimate the parameters and then update the control signal accordingly. However, given that the end goal is to generate the correct control signal, perhaps we are estimating the wrong entity: why not simply estimate the control signal directly? The essential ideas proposed here have been utilized by the author in solving the simultaneous stabilization

¹ Philosophically one might expect that while the controller is learning about the system we will have poor performance. However, this does not make it any more palatable!

² An exception is the high-gain approach of [13], which can provide an arbitrarily good transient and steady-state response for *minimum phase systems*. The downside is that extremely large control signals are often required.

problem with near optimal performance [12,15,14], as well as to solve the MRACP in the 1-D case with fixed parameters [9]. Here we will focus on the MRACP in the 1-D case, but with time-varying parameters, providing the intuition for the approach and only a sketch of the proof.³ The important features that this approach provides are:

- (a)' we can tolerate time-varying parameters as long as we have an upper bound on the derivatives,
- (b)' if the plant/controller ICs are zero, then the tracking is immediate (we no longer get exact tracking, but we do get tracking as good as we would like),
- (c)' the controller is linear, so the effect of the ICs and the exogenous input are decoupled,
- (d)' the control signal is modest in size (it is very much like the control signal one would use if the plant parameters were known).

Of course, a price must be paid for these extremely nice features:

- (i) While the control signal tends to be modest in size, the controller gains tend to be large, which may yield poor noise tolerance, and
- (ii) fast actuators may be required.

However, we emphasize that, as far as we aware, this is the first controller with all of the features listed above. This hopefully will provide some impetus for researchers to further investigate this approach and perhaps enhanced techniques can be developed which provide the same features with a smaller downside.

Since this paper was written, the details for the high order, high relative degree, case have been developed [10].⁴ We emphasize that this approach is not limited to the MRACP. Indeed, we have been recently investigating its use for a first order optimal pole placement problem (with fixed but unknown parameters), for which the goal is to design a controller so that the closed loop behaviour is near optimal in the LQR sense. We are presently working on extending the ideas to the problem of pole placement of time-varying high order systems.

Before proceeding with details, let us first provide an overview of the controller. The controller that we propose is sampled-data and periodic. Each period is split into two parts: an estimation phase and a control phase. During the estimation phase we linearly estimate what the control signal would be if the plant parameters were known. In the control phase we apply a suitably scaled version of the estimate. We prove that as the sampling period goes to zero, we will tend toward exact model matching.

2. Preliminary mathematics

With $x \in \mathbf{R}^n$, the norm of x is defined by

$$\|x\| := \max\{|x_i|: i = 1, \dots, n\}.$$

The norm of $A \in \mathbf{R}^{n \times m}$, denoted $\|A\|$, is the corresponding induced norm.

We let $PC(\mathbf{R}^{n \times m})$ denote the set of piecewise continuous functions from \mathbf{R}^+ to $\mathbf{R}^{n \times m}$, and $PC^1(\mathbf{R}^{n \times m})$ to denote the subset of absolutely continuous signals. For $f \in PC(\mathbf{R}^{n \times m})$, define

$$\|f\|_\infty := \text{esssup}_{t \in \mathbf{R}^+} \|f(t)\|.$$

Let $PC_\infty(\mathbf{R}^{n \times m})$ denote the set of $f \in PC(\mathbf{R}^{n \times m})$ for which $\|f\|_\infty < \infty$, and let $PC_\infty^1(\mathbf{R}^{n \times m})$ denote the set of $f \in PC^1(\mathbf{R}^{n \times m})$ for which $\|f\|_\infty < \infty$ and $\|\dot{f}\|_\infty < \infty$. Henceforth we drop the $\mathbf{R}^{n \times m}$ and simply write PC , PC_∞ , PC^1 and PC_∞^1 .

³ A proof for this case as well as for a general high order, high relative degree, plant is provided in [10].

⁴ A condensed version is available in the conference paper [11].

In this paper we will be dealing with linear time-varying systems, and it will be convenient to discuss the gain of such a system when starting with zero initial conditions at time zero. To this end, the gain of $G: PC_\infty \rightarrow PC$ is defined by

$$\|G\| := \sup \left\{ \frac{\|Gu_m\|_\infty}{\|u_m\|_\infty} : u_m \in PC_\infty, \|u_m\|_\infty \neq 0 \right\}.$$

We say that the system is *stable* or *bounded* if $\|G\| < \infty$.

We say that $f: \mathbf{R}^+ \rightarrow \mathbf{R}^{n \times m}$ is of order T^j , and write $f = \mathcal{O}(T^j)$, if there exist constants $c_1 > 0$ and $T_1 > 0$ so that

$$\|f(T)\| \leq c_1 T^j, \quad T \in (0, T_1).$$

On occasion we have a function f which depends not only on $T \geq 0$ but also on a pair (a, b) restricted to a set $\Gamma \subset PC_\infty^1$; we say that $f = \mathcal{O}(T^j)$ if there exist constants $c_1 > 0$ and $T_1 > 0$ so that

$$\|f(T)\| \leq c_1 T^j, \quad T \in (0, T_1), \quad (a, b) \in \Gamma.$$

3. The problem

Our first order plant P is described by

$$\dot{y}(t) = a(t)y(t) + b(t)u(t), \quad y(t_0) = y_0, \quad (1)$$

with $y(t) \in \mathbf{R}$ the plant state (and measured output) and $u(t) \in \mathbf{R}$ the plant input; we associate this plant with the pair (a, b) .⁵ If a and b happen to be constant then the associated transfer function from u to y is given by

$$P(s) := \frac{b}{s - a}.$$

We assume that (a, b) are absolutely continuous, bounded with a bounded derivative, and with b bounded away from zero, i.e. there exist positive constants c_0, c_1, c_2, c_3 and c_4 so that the set of admissible parameters is given by

$$\Gamma \subset \{(a, b) \in PC_\infty^1 : \|a\|_\infty \leq c_0, \|b\|_\infty \leq c_1, \|\dot{a}\|_\infty \leq c_2, \|\dot{b}\|_\infty \leq c_3, |b(\cdot)| \geq c_4\}.$$

While our controller can tolerate the occasional parameter jump, we shall omit that from the paper. If both the LTI plants $c_1/(s - c_0)$ and $-c_1/(s - c_0)$ are admissible, since they are not simultaneously stabilizable by an LTI controller it follows that we would need either a nonlinear or time-varying controller to stabilize all admissible systems.

Our stable SISO reference model P_m is described by

$$\dot{y}_m(t) = a_m y_m(t) + b_m u_m(t), \quad y_m(t_0) = y_{m_0}, \quad (2)$$

with $y_m(t) \in \mathbf{R}$ the reference model state and $u_m(t) \in \mathbf{R}$ the reference model input.⁶ The model is chosen to embody the desired behaviour of the closed loop system; clearly we require it to be stable. Informally, our goal is to design a controller to make the plant output track the reference model output in the face of plant uncertainty. To this end, we define the (*tracking*) error by

$$e(t) := y_m(t) - y(t).$$

⁵ We let P denote the input-output map of (1) when $y_0 = t_0 = 0$.

⁶ We let P_m denote the input-output map of (2) when $y_{m_0} = t_0 = 0$.

Our goal is to construct a single linear time-varying controller which not only provides closed loop stability but also provides near optimal performance for each possible model. It is our intention to use a sampled data controller, so it is natural to use an anti-aliasing filter at the input. Hence, our control law has two parts: with $\alpha > 0$ we have an anti-aliasing filter at the input of the form

$$\dot{\bar{u}}_m = -\alpha \bar{u}_m + \alpha u_m, \quad \bar{u}_m(t_0) = \bar{u}_{m_0} \quad (3)$$

whose input-output map is labeled F_α ; this is followed by a sampled-data controller of the form

$$\begin{aligned} z[k+1] &= F(k)z[k] + G(k)y(kh) + H(k)\bar{u}_m(kh), \quad z[k_0] = z_0 \in \mathbf{R}^l, \\ u(kh + \tau) &= J(k)z[k] + L(k)y(kh), \quad \tau \in [0, h) \end{aligned} \quad (4)$$

whose input-output map is labeled K and whose gains F , G , H , J and L are periodic of period $p \in \mathbf{N}$; the period of the controller is $T := ph$, and we associate this system with the 7-tuple (F, G, H, J, L, h, p) . Observe that (4) can be implemented with a sampler, a zero-order-hold, and an l th order periodically time-varying discrete-time system of period p .

At this point we would like to define a notion of stability. Since we are mixing continuous and discrete-time signals, we have to be careful: we define our combined state by

$$x_{sd}(t) = \begin{bmatrix} y(t) \\ y_m(t) \\ \bar{u}_m(t) \\ z[k] \end{bmatrix}, \quad t \in [kh, (k+1)h).$$

Definition 1. The controller (3) and (4) exponentially stabilizes Γ if there exist constants $\gamma > 0$ and $\lambda < 0$ so that, for every $(a, b) \in \Gamma$, set of initial conditions y_0 , y_{m_0} , \bar{u}_{m_0} , and z_0 , and set of initial times $k_0 \in \mathbf{Z}^+$ and $t_0 = k_0$, with $u_m(t) = 0$ we have

$$\|x_{sd}(t)\| \leq \gamma e^{\lambda(t-t_0)} (\|y_0\| + \|y_{m_0}\| + \|\bar{u}_{m_0}\| + \|z_0\|), \quad t \geq t_0.$$

Now suppose that (3)–(4) exponentially stabilizes (1) and that $t_0 = k_0 = y_0 = 0$ and $z_0 = 0$; we let $\mathcal{F}(P, K)$ denote the closed loop map from $\bar{u}_m \rightarrow y$. Our goal is to design K so that $\|\mathcal{F}(P, K)F_\alpha - P_m\|$ is small for every possible P ; however, notice that

$$\begin{aligned} \|\mathcal{F}(P, K)F_\alpha - P_m\| &\leq \|\mathcal{F}(P, K)F_\alpha - P_m F_\alpha\| + \|P_m(F_\alpha - 1)\| \\ &\leq \|\mathcal{F}(P, K)F_\alpha - P_m F_\alpha\| + \frac{2|b_m|}{|a_m + \alpha|}. \end{aligned}$$

Hence, we can first choose a sufficiently large α to make the second term small, and then we can proceed to design K to make the first term small for all admissible P . To this end, we define

$$\dot{\bar{y}}_m = a_m \bar{y}_m + b_m \bar{u}_m, \quad \bar{y}_m(t_0) = y_{m_0}. \quad (5)$$

Notice that for $\alpha > |a_m|$, if $t_0 = 0$ we have

$$|\bar{y}_m(t) - y_m(t)| \leq \frac{|b_m|}{|a_m + \alpha|} |\bar{u}_{m_0}| e^{a_m t} + \frac{2|b_m|}{|a_m + \alpha|} \|u_m\|_\infty, \quad t \geq 0. \quad (6)$$

In the next section we will provide a high level description of the design approach.

4. The approach

Here we explain the motivation of the approach to the problem. Our goal is to make the difference between y and \bar{y}_m small, so let's form a differential equation describing this quantity:

$$(\dot{\bar{y}}_m - \dot{y}) = a_m(\bar{y}_m - y) + (b_m\bar{u}_m - b(t)u + (a_m - a(t))y).$$

Since we would like the plant to act like the reference model, we may as well require that the error caused by a mis-match in initial condition go to zero like $e^{a_m t}$, i.e. we may as well set⁷

$$b_m\bar{u}_m - b(t)u + (a_m - a(t))y = 0 \leftrightarrow u = \frac{1}{b(t)} (b_m\bar{u}_m + (a_m - a(t))y), \quad (7)$$

which yields

$$y(t) = e^{a_m t} y(0) + \int_0^t e^{a_m(t-\tau)} b_m\bar{u}_m(\tau) d\tau. \quad (8)$$

Of course, a and b are unknown and time-varying, so some form of estimation is required. Since we would like to end up with a linear controller, we would like to get rid of the $1/b$ term. It will turn out that we will be able to deal with polynomials. From the form of Γ , it is clear that there exist positive \underline{b} and \bar{b} so that

$$(a, b) \in \Gamma \Rightarrow 0 < \underline{b} \leq |b(\cdot)| \leq \bar{b}, \quad (a, b) \in \Gamma.$$

From the Stone–Weierstrass Approximation Theorem [23] we know that we can approximate $1/b$ arbitrarily well over the compact set $[-\bar{b}, -\underline{b}] \cup [\underline{b}, \bar{b}]$ via a polynomial. So for every $\varepsilon > 0$ we can choose a polynomial $\hat{f}_\varepsilon(b)$ so that

$$|1 - b\hat{f}_\varepsilon(b)| < \varepsilon, \quad b \in [-\bar{b}, -\underline{b}] \cup [\underline{b}, \bar{b}]. \quad (9)$$

Now we consider our second approximation (the first one being our use of an anti-aliasing filter). Rather than applying the feedback control (7), instead we apply

$$u = \hat{f}_\varepsilon(b(t))(b_m\bar{u}_m + (a_m - a(t))y). \quad (10)$$

It is intuitively reasonable that for small $\varepsilon > 0$, this is a good approximation to the optimal controller given by (7). *At this point we freeze $\varepsilon > 0$ small and choose a polynomial*

$$\hat{f}_\varepsilon(b) = \sum_{i=1}^q c_i b^i$$

so that (9) is satisfied. (It is easy to see that we do not need a constant term.)

Now the goal is to design a sampled-data control law of the form (4) so that we can approximately implement (10) regardless of which admissible system that we are controlling. We use h small and with q the order of our polynomial approximation to $1/b$, we let $p > 2q + 1$; recall that the controller period is $T = ph$. Here we provide a conceptual description of the controller and a high-level description of why it should work. We first do so in open loop and then explain how to convert it to a controller of the form (4). To motivate our approach, first observe that the sampled-data control law

$$u(t) = \begin{cases} 0, & t \in [kT, kT + (2q + 1)h), \\ \frac{p}{p - (2q + 1)} \hat{f}_\varepsilon(b(kT)) [b_m\bar{u}_m(kT) + (a_m - a(kT))y(kT)], & t \in [kT + (2q + 1)h, (k + 1)T) \end{cases}$$

⁷ This is not essential—we can require the dynamics of the mis-match in initial conditions to be as fast as we like; here we have made the most natural choice.

should be a good approximation to (10) if h and $T = ph$ are both small. Here we will implement something similar to this: every period $[kT, (k+1)T)$ is divided into two phases:

- identification/estimation phase: on the interval $[kT, kT + (2q+1)h)$ we estimate

$$\hat{f}_\varepsilon(b(kT))[b_m \bar{u}_m(kT) + (a_m - a(kT))y(kT)].$$

While we do not set $u(t)$ equal to zero here, we ensure that the effect of the probing used in the estimation yields only a second order effect.

- control phase: on the interval $[kT + (2q+1)h, (k+1)T)$ we apply $p/(p-(2q+1))$ times the above estimate.

The tricky part is figuring out a way to carry out the identification/estimation phase in a linear fashion.

Let us look at the first period $[0, T)$. We would like to first form an approximation of

$$\hat{f}_\varepsilon(b(0))[b_m \bar{u}_m(0) + (a_m - a(0))y(0)] = \sum_{i=1}^q c_i b(0)^i \underbrace{[b_m \bar{u}_m(0) + (a_m - a(0))y(0)]}_{=: \phi_i(0)}.$$

Suppose that we initially set

$$u(t) = 0, \quad t \in [0, h).$$

Since a is constrained to a compact set and its derivative is bounded, it follows that

$$y(h) = e^{\int_0^h a(\theta) d\theta} y(0) = [1 + a(0)h + \mathcal{O}(h^2)]y(0).$$

Hence,

$$\frac{1}{h} [y(h) - y(0)] = a(0)y(0) + \mathcal{O}(h)y(0).$$

So at this point we have a good estimate of $a(0)y(0)$, with the quality of the estimate improving as $h \rightarrow 0$.⁸ Since we can measure \bar{u}_m , we can form a good estimate of

$$\phi_0(0) = b_m \bar{u}_m(0) + (a_m - a(0))y(0),$$

namely

$$\begin{aligned} \hat{\phi}_0(0) &:= b_m \bar{u}_m(0) + a_m y(0) - \frac{1}{h} [y(h) - y(0)] \\ &= b_m \bar{u}_m(0) + (a_m - a(0))y(0) + \mathcal{O}(h)y(0) \\ &= \phi_0(0) + \mathcal{O}(h)y(0). \end{aligned}$$

To form estimates of $\phi_i(0) = b(0)^i \phi_0(0)$, we will carry out some experiments. With $\rho > 0$ a scaling factor (we make this small so that it does not disturb the system very much), set

$$u(t) = \rho \hat{\phi}_0(0), \quad t \in [h, 2h).$$

⁸ The skeptics will be concerned that we are differentiating the plant output. We are indeed carrying out a discrete-time approximation to differentiation, with the estimate improving but the noise tolerance degrading as $h \rightarrow 0$. This is akin to the problem arising in PID design where one has to roll off the differentiator term at the appropriate frequency: the higher the rolloff frequency the better the approximation to a pure differentiator and the worse the noise behaviour.

Then

$$\begin{aligned}
y(2h) &= e^{\int_0^{2h} a(\theta) d\theta} y(0) + \left[\int_h^{2h} e^{\int_\tau^{2h} a(\theta) d\theta} b(\tau) d\tau \right] u(h) \\
&= [1 + 2a(0)h + \mathcal{O}(h^2)]y(0) + \int_h^{2h} [1 + \mathcal{O}(h)][b(0) + \mathcal{O}(h)] d\tau \times u(h) \\
&= [1 + 2a(0)h + \mathcal{O}(h^2)]y(0) + [b(0)h + \mathcal{O}(h^2)]\rho\hat{\phi}_0(0) \\
&= (1 + 2a(0)h)y(0) + \rho b(0)h\phi_0(0) + \mathcal{O}(h^2)y(0) + \mathcal{O}(h^2)\bar{u}_m(0).
\end{aligned} \tag{11}$$

Hence, we can define

$$\begin{aligned}
\hat{\phi}_1(0) &:= \frac{1}{\rho h} [y(2h) - 2y(h) + y(0)] \\
&= b(0)\phi_0(0) + \mathcal{O}(h)y(0) + \mathcal{O}(h)\bar{u}_m(0) \\
&= \phi_1(0) + \mathcal{O}(h)y(0) + \mathcal{O}(h)\bar{u}_m(0).
\end{aligned}$$

Of course, in carrying out this experiment we have given the state a boost—see (11). This can be largely undone by applying

$$u(t) = -\rho\hat{\phi}_0(0), \quad t \in [2h, 3h),$$

for then

$$y(3h) = (1 + 3a(0)h)y(0) + \mathcal{O}(h^2)y(0) + \mathcal{O}(h^2)\bar{u}_m(0).$$

This can be repeated $q - 1$ more times: for $i = 1, \dots, q - 1$ we set

$$u(t) = \begin{cases} \rho\hat{\phi}_i(0), & t \in [(2i + 1)h, (2i + 2)h), \\ -\rho\hat{\phi}_i(0), & t \in [(2i + 2)h, (2i + 3)h), \end{cases}$$

and estimate $\phi_i(0)$ via

$$\hat{\phi}_{i+1}(0) := \frac{1}{\rho h} [y((2i + 2)h) - y((2i + 1)h) - y(h) + y(0)].$$

We can show inductively that

$$\hat{\phi}_i(0) = \phi_i(0) + \mathcal{O}(h)y(0) + \mathcal{O}(h)\bar{u}_m(0), \quad i = 1, \dots, q.$$

At the end of the estimation phase, we are at $t = (2q + 1)h$, and we have estimates of $\phi_1(0), \dots, \phi_q(0)$ to form our control signal to be applied during the control phase. There is a lot of flexibility in how long one can make the control phase: the higher the percentage that we carry out control the closer that the magnitude will be to the ideal one; however, if we make the percentage too large then we will need a very small h to ensure that T is small enough, which will exacerbate noise tolerance. In any event, at this point fix

$$p > 2q + 1$$

with a corresponding controller period of $T = ph$. Then we set

$$u(t) = \frac{p}{p - 2q - 1} \sum_{i=1}^q c_i \hat{\phi}_i(0), \quad t \in [(2q + 1)h, ph).$$

It follows that

$$\begin{aligned}
y(T) &= e^{\int_0^T a(\theta) d\theta} y(0) + \int_0^T e^{\int_\tau^T a(\theta) d\theta} b(\tau) u(\tau) d\tau \\
&= [1 + a(0)T + \mathcal{O}(T^2)]y(0) + Tb(0) \sum_{i=1}^q c_i \hat{\phi}_i(0) + \mathcal{O}(T^2)y(0) + \mathcal{O}(T^2)\bar{u}_m(0) \\
&= [1 + a(0)T]y(0) + Tb(0) \hat{f}_\varepsilon(b(0)) [b_m \bar{u}_m(0) + (a_m - a(0))y(0)] + \mathcal{O}(T^2)y(0) + \mathcal{O}(T^2)\bar{u}_m(0) \\
&\approx [1 + a(0)T]y(0) + T[b_m \bar{u}_m(0) + (a_m - a(0))y(0)] \\
&= (1 + a_m T)y(0) + Tb_m \bar{u}_m(0) \\
&\approx e^{a_m T} y(0) + \int_0^T e^{a_m(T-\tau)} b_m \bar{u}_m(\tau) d\tau.
\end{aligned}$$

Of course, at time $T = ph$ the procedure is repeated, but now using $y(T)$ and $\bar{u}_m(T)$ instead of $y(0)$ and $\bar{u}_m(0)$. We end up with

$$y[(k+1)T] \approx e^{a_m T} y[kT] + \int_0^T e^{a_m(T-\tau)} b_m \bar{u}_m(kT + \tau) d\tau,$$

which means that

$$y(t) \approx e^{a_m t} y(0) + \int_0^t e^{a_m(t-\tau)} b_m \bar{u}_m(\tau) d\tau,$$

which is the desired goal given in (8).

The above controller is provided in open loop. However, with a little thought we can easily see how to implement this controller using a sampled-data system of the form (4). Specifically, we need a state z of dimension 4:

- z_1 keeps track of $y(kT + h) - y(kT)$ for later use,
- z_2 keeps track of the current $\hat{\phi}_i$ being constructed,
- z_3 is used to construct the control signal which is applied during $[kT, kT + (2q+1)h)$, which is the estimation phase, and
- z_4 is used to construct the control signal which is applied during $[kT + (2q+1)h, (k+1)T)$, which is the control phase.

We omit the details of the construction, and refer the reader to [10].

Remark 1. Observe that during the identification phase we can make the control signal small by letting ρ be small. Also, during the control phase we have the control signal approximately equal to $p/(p-2q-1)$ times what it would be if the plant parameters were known. Since we can choose p as large as we'd like, we see that we can ensure that the control signal does not get too large.

Theorem 1. For every $\delta > 0$ there exists a controller of the form (3) and (4) with the following properties:

- (i) it exponentially stabilizes Γ , and
- (ii) if it is applied to $(a, b) \in \Gamma$ with $k_0 = t_0 = 0$, then the closed loop system satisfies

$$|[y(t) - y_m(t)] - e^{a_m t} [y_0 - y_{m_0}]| \leq \delta e^{(a_m + \delta)t} (|y_0| + |y_{m_0}| + |\bar{u}_{m_0}|) + \delta \|u_m\|_\infty, \quad t \geq 0.$$

Proof. Here we provide a sketch of the proof; the complete version can be found in [10].⁹

First, we let the aliasing filter parameter α be large so that the RHS of (6) is small. Next, we let the parameter $\varepsilon > 0$ be small so that the closed loop behaviour provided by the approximate ideal LTI controller (10) is close to that provided by the ideal LTI controller (7). Last of all, we let the parameter T in the proposed sampled-data controller be small so that the closed loop behaviour provided by it is close to that provided by the approximate ideal LTI controller (10). \square

Remark 2. Hence, we see that the effect of the mismatch in initial conditions between the plant and the controller decays to zero at a pre-specified rate (since the value of a_m is determined by the control system designer) while the tracking of the reference signal is immediate.

Remark 3. Since the controller is linear and exponentially stabilizing and the plant parameters are uniformly bounded, we automatically have tolerance to noise in the following sense. Suppose that n_1 and n_2 are two noise signals and we replace (4) with

$$\begin{aligned} y_n(t) &= y(t) + n_1(t), \\ z[k+1] &= F(k)z[k] + G(k)y_n(kh) + H(k)\bar{u}_m(kh), \quad z[k_0] = z_0 \in \mathbf{R}^l, \\ u_n(kh + \tau) &= J(k)z[k] + L(k)y_n(kh), \quad \tau \in [0, h), \\ u(t) &= u_n(t) + n_2(t). \end{aligned} \tag{12}$$

When the controller given by (3) and (12) is applied to the plant, then the map from $\begin{bmatrix} n_1 \\ n_2 \end{bmatrix} \mapsto \begin{bmatrix} u \\ y \end{bmatrix}$ is uniformly (w.r.t. Γ) bounded.

Remark 4. Many classical adaptive controllers perform poorly when there are unmodeled dynamics, e.g. see [22]. Fortunately, since our controller is linear (time-varying) and exponentially stabilizes the plant, we automatically get some tolerance to unmodeled dynamics.

5. An example

Here we consider an example to illustrate the proposed design methodology. The set of plant uncertainty is given by

$$\Gamma = \{(a, b) \in PC_\infty^1 : a(t) \in [-1, 1], \quad b^2(t) \in [1, 2], \quad |\dot{a}(t)| \leq 1, \quad |\dot{b}(t)| \leq 1\}$$

and the reference model is

$$\dot{x}_m = -x_m + u_m.$$

Recall that there is no LTI controller to stabilize $\{1/(s-1), -1/(s-1)\}$, so it follows that there does not exist a single LTI controller to stabilize every $(a, b) \in \Gamma$, let alone provide good performance.

We choose an anti-aliasing filter of the form

$$\dot{\bar{u}}_m = -50\bar{u}_m + 50u_m,$$

which means that this degrades the performance by at most 4%.

⁹ A similar controller has been designed and analysed in detail in [9]. It deals solely with the first order (fixed parameter) case, so it is easier to follow than the general proof provided in [10].




Fig. 1. The closed loop response for the example.

Here we base our polynomial approximation of $1/b$ on the optimal approximation technique discussed in [24]: we have

$$\hat{f}_{0.01}(b) = 2.1647b - 1.5153b^3 + 0.3433b^5.$$

So $q = 5$ and we choose $p = 22 = 2(2q + 1)$, i.e. the control phase is 50% of the time. We set $T = 0.048$ (so that $h = 0.002$) and $\rho = 1$. A simulation was carried out (see Fig. 1) with $t_0 = k_0 = 0$, $y_0 = 2$, $y_{m_0} = \bar{u}_{m_0} = 0$, u_m a square wave of period 10,

$$a(t) = \cos(t/2),$$

and

$$b(t) = [1.2 + 0.2 \cos(t/2)] * \text{sign}[\cos(t/4)];$$

for the last 10 s of the simulation we add a noise signal of the form

$$0.5h * \text{random sequence uniformly distributed between } \pm 1$$

to the output measurement. We see that the effect of the initial condition decays exponentially to zero, the tracking is quite good, and the control signal is modest in size, even in the face of significant time variations and noisy measurements; indeed, it even works when there are jumps in the parameters.

6. Summary and conclusions

In classical adaptive control, the controller consists of an identifier (or tuner) which is used to adjust the parameters of an LTI compensator, and under suitable assumptions on the plant model uncertainty it is proven that good asymptotic behaviour is achieved. However, the controller is often poor at tracking rapidly moving parameters, the tracking is often asymptotic only, and due to the nonlinear nature of the controller there may be large transients.

Here we propose an alternative approach. Rather than estimating the plant or compensator parameters, instead *we estimate what the control signal would be if the plant parameters were known*; we are able to do so in a linear fashion, and we end up with a linear periodic controller. We consider the first order case and construct a controller for the MRACP with the following features:

- it handles rapidly moving parameters (we assume an upper bound on rate of change of the parameters),
- the tracking of the reference input is immediate rather than asymptotic,
- the effect of the initial conditions decays exponentially to zero, and
- the control signal is modest in size.

An undesirable feature is that in order to obtain tight tracking performance, we may require a small sampling period and large gains, which may yield poor noise tolerance; furthermore, we may need fast actuators. While the approach is illustrated here only for the first order case, recently it has been extended to the high order, high relative degree, case [10]. We are investigating the use of the technique for adaptive pole placement.

As far as the author is aware, this is the first adaptive controller with all of the features listed above; it is surprising this can be achieved using a linear time-varying controller. This hopefully will provide some impetus for researchers to investigate this approach further and perhaps an enhanced version can be developed which provides the same features with a smaller downside.

References

- [1] G.C. Goodwin, P.J. Ramadge, P.E. Caines, Discrete time multivariable control, IEEE Trans. Automatic Control AC-25 (1980) 449–456.
- [2] J. Hocherman-Frommer, S.R. Kulkarni, P.J. Ramadge, Controller switching based on output prediction errors, IEEE Trans. Automatic Control AC-43 (1998) 596–607.
- [3] P.A. Ioannou, J. Sun, Robust Adaptive Control, Prentice-Hall, Englewood Cliffs, NJ, 1996.
- [4] R.E. Kalman, Design of self-optimizing systems, Trans. ASME 80 (1958) 468–478.
- [5] S. Limanond, K.S. Tsakalis, Model reference adaptive and nonadaptive control of linear time-varying systems, IEEE Trans. Automatic Control AC-45 (2000) 1290–1300.
- [6] R. Marino, P. Tomei, Robust adaptive regulation of linear time-varying systems, IEEE Trans. Automatic Control AC-45 (2000) 1301–1311.
- [7] B. Martensson, The order of any stabilizing regulator is sufficient a priori information for adaptive stabilization, Systems Control Lett. 6 (1985) 87–91.
- [8] R.H. Middleton, G.C. Goodwin, Adaptive control of time-varying systems, IEEE Trans. Automatic Control AC-33 (1988) 150–155.
- [9] D.E. Miller, A time-varying approach to model reference adaptive control, in: S.O.R. Moheimani (Ed.), Lecture Notes in Control and Information Sciences, Vol. 268, Springer, London, 2001, pp. 199–219.
- [10] D.E. Miller, A new approach to model reference adaptive control, Technical Report UW-E& CE#2001-15, Dept. of Electrical and Computer Engineering, University of Waterloo, Ontario, Canada, December 20, 2001.
- [11] D.E. Miller, A new approach to model reference adaptive control, Proceedings of the 2002 American Control Conference, Anchorage, Alaska, May 2002, pp. 2929–2934.
- [12] D.E. Miller, T. Chen, Simultaneous stabilization with near optimal H-inf performance, Proceedings of the 39th IEEE Conference on Decision and Control, Sydney, Australia, December 2000, pp. 2722–2727.
- [13] D.E. Miller, E.J. Davison, An adaptive controller which provides an arbitrarily good transient and steady-state response, IEEE Trans. Automatic Control AC-36 (1991) 68–81.

- [14] D.E. Miller, D.C. Kennedy, Simultaneous stabilization with near optimal model matching, Proceedings of the 2001 American Control Conference, Washington, DC, June 2001, pp. 3159–3164.
- [15] D.E. Miller, M. Rossi, Simultaneous stabilization with near optimal LQR performance, IEEE Trans. Automatic Control AC-46 (2001) 1543–1555.
- [16] A.S. Morse, Global stability of parameter-adaptive control systems, IEEE Trans. Automatic Control AC-25 (1980) 433–439.
- [17] A.S. Morse, Overcoming the obstacle of high relative degree, Eur. J. Control 2 (1996) 29–35.
- [18] A.S. Morse, Supervisory control of families of linear set-point controllers—Part 1: exact matching, IEEE Trans. Automatic Control AC-41 (1996) 1413–1431.
- [19] K.S. Narendra, Y.H. Lin, L.S. Valavani, Stable adaptive controller design, Part II: proof of stability, IEEE Trans. Automatic Control AC-25 (1980) 440–448.
- [20] K.S. Narendra, J. Balakrishnan, Adaptive control using multiple models, IEEE Trans. Automatic Control AC-42 (1997) 171–187.
- [21] Z. Qu, E. Kamen, J. Dorsey, Continuous I/O robust control of SISO time-varying systems, Automatica 33 (1997) 533–550.
- [22] C. Rohrs, L. Valavani, M. Athans, G. Stein, Robustness of adaptive control algorithms in the presence of unmodeled dynamics, IEEE Trans. Automatic Control AC-33 (1985) 881–889.
- [23] W. Rudin, Principles of Mathematical Analysis, 3rd Edition, McGraw-Hill, New York, 1976.
- [24] J. Todd (Ed.), Survey of Numerical Methods, McGraw-Hill, New York, 1962.
- [25] K.S. Tsakalis, P.A. Ioannou, Linear Time-Varying Plants: Control and Adaptation, Prentice-Hall, Englewood Cliffs, NJ, 1993.
- [26] H.P. Whitaker, J. Yamron, A. Kezer, Design of model reference adaptive control systems for aircraft, Report R-164, Instrumentation Laboratory, MIT Press, Cambridge, MA, 1958.
- [27] C.I. Byrnes, J.C. Willems, Adaptive stabilization of multivariable linear systems, Proceedings of the IEEE 23rd Conference on Decision and Control, 1984, pp. 1574–1577.