Systems & Control Letters 103 (2017) 23–31

Contents lists available at ScienceDirect

Systems & Control Letters

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

A parameter adaptive controller which provides exponential stability: The first order case

Daniel E. Miller

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

ARTICLE INFO

Article history:
Received 8 December 2016
Received in revised form 28 January 2017
Accepted 23 February 2017

Keywords:
Adaptive control
Projection algorithm
Projection
Exponential stability
Bounded gain

ABSTRACT

Classical discrete-time adaptive controllers provide asymptotic stabilization. While the original adaptive controllers did not handle noise or unmodelled dynamics well, redesigned versions were proven to have some tolerance; however, neither exponential stabilization nor a bounded noise gain is typically proven. Here we consider the first order case and prove that if the original, ideal, projection algorithm is used in the estimation process (subject to the common assumption that the plant parameters lie in a convex, compact set and that the parameter estimates are restricted to that set), then it guarantees linear-like convolution bounds on the closed loop behaviour, which implies exponential stability and a bounded noise gain, as well an easily proven tolerance to unmodelled dynamics and plant parameter variation.

1. Introduction

The idea of adaptive control is for the controller to adapt to the plant being controlled. The first general proofs that parameter adaptive controllers could work came around 1980, e.g. see [1–5]. However, such controllers were typically not robust to unmodelled dynamics, did not tolerate time-variations well, had poor transient behaviour, and did not handle noise (or disturbances) well, e.g. see [6].

During the 1980s and 1990s a great deal of effort was made to address these shortcomings. The most common approach was to make modifications to the estimator used in the adaptive control laws so that the resulting controllers tolerated a small amount of tolerance to unmodelled dynamics, slow time-variations, and/or noise or disturbances, e.g. see [7–12]. Indeed, simply using projection (onto a convex set of admissible parameters) has proved quite powerful, and the resulting controllers typically provide a bounded-noise bounded-state property, and often tolerance of some degree of unmodelled dynamics and time-variations — see [11,13–17]. As far as the author is aware, a bounded gain on the noise is proven only in one special case of [14]; however, a convolution-like bound on the behaviour is not proven (this is critical to proving that the approach can tolerate slow time-variations — see Section 6), and a crisp bound on the effect of the initial condition is not provided.

In this paper we revisit the discrete-time approach, such as that of [3,18]. It is common to carry out parameter estimation with either a least-squares algorithm or a projection algorithm. However, it is the norm in the literature to use a modified version of the projection algorithm to avoid the issue of dividing by zero; it turns out that an unexpected consequence of this minor adjustment is that some inherent properties of the scheme are destroyed. Here we consider the first order case and demonstrate that if the original, ideal, projection algorithm is used (subject to the common assumption that the plant parameters lie in a convex, compact set and that the parameter estimates are restricted to that set) in conjunction with the common one-step-ahead adaptive control law, then we can obtain linear-like convolution bounds on the closed-loop behaviour. This immediately confers exponential stability as well as a bounded gain on the noise; this can be leveraged to prove a degree of tolerance to parameter variation as well as unmodelled dynamics. Here we have focused on the first order case since its rich structure allows for a direct and self-contained analysis.

We use standard notation throughout the paper. We use the Euclidean 2-norm for vectors and the corresponding induced norm for matrices, and denote the norm of a vector or matrix by $\|\cdot\|$. We let l_∞ denote the set of real-valued bounded sequences. If $\mathcal{S} \subset \mathbf{R}^p$ is a convex and compact set, we define $\|\mathcal{S}\| := \max_{x\in\mathcal{S}} \|x\|$ and the function $\pi_{\mathcal{S}} : \mathbf{R}^p \to \mathcal{S}$ denotes the projection onto $\mathcal{S}$; it is well-known that $\pi_{\mathcal{S}}$ is well-defined. With $\varepsilon > 0$, we let $s(\mathcal{S}, \varepsilon)$ denote the set of sequences in $x \in l_\infty$ taking values in $\mathcal{S}$ and satisfying $|x(i+1) - x(i)| \le \varepsilon$ for $i \in \mathbf{Z}$.

2. The setup

Here we start with the first order linear time-invariant discrete-time plant

$$y(t+1) = ay(t) + bu(t) + n(t), \quad y(t_0) = y_0 \tag{1}$$

E-mail address: miller@hobbes.uwaterloo.ca.

http://dx.doi.org/10.1016/j.sysconle.2017.02.009
0167-6911/

where $y(t) \in \mathbf{R}$ is the state, $u(t) \in \mathbf{R}$ is the control input, and $n(t) \in \mathbf{R}$ is the noise (or disturbance). Here the pair $\begin{bmatrix} a \\ b \end{bmatrix}$ lies in a **closed, convex, bounded** set $\mathcal{S} \subset \mathbf{R}^2$; the closed and convex properties are used to facilitate projection, while the bounded property is used to prove uniform gains, decay rates, etc. To ensure controllability we require that $\begin{bmatrix} a \\ 0 \end{bmatrix} \notin \mathcal{S}$ for any $a \in \mathbf{R}$.

Here we have an exogenous reference signal $r(t)$ and the objective is to track it asymptotically; we assume that we know it one step ahead, i.e. that we know $r(t)$ at time $t-1$. We will be interested in analysing the corresponding one-step-ahead control law when the plant parameters are unknown. We use an estimator together with the Certainty Equivalence Principle[1] to form our control law.

With $\phi(t) := \begin{bmatrix} y(t) \\ u(t) \end{bmatrix}$ and $\theta^* := \begin{bmatrix} a \\ b \end{bmatrix}$, we can rewrite the plant as

$$y(t+1) = \phi(t)^T \theta^* + n(t).$$

Given an estimate $\hat{\theta}(t)$ of θ^* at time t, we define the **prediction error** by

$$e(t+1) := y(t+1) - \phi(t)^T \hat{\theta}(t); \tag{2}$$

this is a measure of the error in $\hat{\theta}(t)$. The common way to obtain a new estimate is from solving the optimization problem

$$argmin_\theta \{\|\theta - \hat{\theta}(t)\| : y(t+1) = \phi(t)^T \theta\},$$

yielding the **ideal (projection) algorithm**

$$\hat{\theta}(t+1) = \begin{cases} \hat{\theta}(t) & \text{if } \phi(t) = 0 \\ \hat{\theta}(t) + \dfrac{\phi(t)}{\phi(t)^T \phi(t)} e(t+1) & \text{otherwise.} \end{cases} \tag{3}$$

Of course, if $\phi(t)$ is close to zero, numerical problems can occur, so it is the norm in the literature (e.g. [3,18]) to replace this by the following **classical algorithm**: with $0 < \alpha < 2$ and $\beta > 0$, define

$$\hat{\theta}(t+1) = \hat{\theta}(t) + \frac{\alpha \phi(t)}{\beta + \phi(t)^T \phi(t)} e(t+1). \tag{4}$$

This latter algorithm is widely used and plays a role in many discrete-time adaptive control algorithms; however, when this algorithm is used all of the results are asymptotic, and exponential stability and a bounded gain on the noise are never proven. It is not hard to guess why — a careful look at the estimator shows that the gain on the update law is small if $\phi(t)$ is small; we will discuss this further once the control law is defined.

Here we will be using the ideal projection algorithm (3) together with projection onto $\mathcal{S}$. With initial conditions of $\hat{\theta}(t_0) = \theta_0 \in \mathcal{S}$ and $y(t_0) = y_0 \in \mathbf{R}$, for $t \geq t_0 + 1$ we set

$$\check{\theta}(t+1) = \begin{cases} \hat{\theta}(t) & \text{if } \phi(t) = 0 \\ \hat{\theta}(t) + \dfrac{\phi(t)}{\phi(t)^T \phi(t)} e(t+1) & \text{otherwise,} \end{cases} \tag{5}$$

which we then project onto $\mathcal{S}$:

$$\hat{\theta}(t+1) := \pi_{\mathcal{S}}(\check{\theta}(t+1)). \tag{6}$$

We partition $\hat{\theta}(t+1)$ in a natural way as

$$\hat{\theta}(t+1) =: \begin{bmatrix} \hat{a}(t+1) \\ \hat{b}(t+1) \end{bmatrix}.$$

If a and b were known, then the one-step-ahead control law is

$$u(t) = \frac{-a}{b} y(t) + \frac{1}{b} r(t+1),$$

which ensures that $y(t+1) = r(t+1)$. Since a and b are not known here, we adopt the Certainty Equivalence counterpart:

$$u(t) = \underbrace{-\frac{\hat{a}(t)}{\hat{b}(t)}}_{=:\hat{f}(t)} y(t) + \underbrace{\frac{1}{\hat{b}(t)}}_{=:\hat{g}(t)} r(t+1). \tag{7}$$

This yields a closed-loop system equation of

$$y(t+1) = [a + b\hat{f}(t)] y(t) + n(t) + b\hat{g}(t) r(t+1). \tag{8}$$

The tracking error is defined by

$$\varepsilon(t) := y(t) - r(t);$$

from (7) we see that $\phi(t)^T \hat{\theta}(t) = r(t+1)$, which, when combined with (2), yields

$$e(t+1) = \varepsilon(t+1), \quad t \geq t_0.$$

Before proceeding, we define some constants:

$$\bar{a} := \max\left\{|a| : \begin{bmatrix} a \\ b \end{bmatrix} \in \mathcal{S}\right\}, \quad \bar{b} := \max\left\{|b| : \begin{bmatrix} a \\ b \end{bmatrix} \in \mathcal{S}\right\},$$

$$\bar{f} := \max\left\{\left|\frac{a}{b}\right| : \begin{bmatrix} a \\ b \end{bmatrix} \in \mathcal{S}\right\}, \quad \bar{g} := \max\left\{\left|\frac{1}{b}\right| : \begin{bmatrix} a \\ b \end{bmatrix} \in \mathcal{S}\right\},$$

$$\underline{g} := \min\left\{\left|\frac{1}{b}\right| : \begin{bmatrix} a \\ b \end{bmatrix} \in \mathcal{S}\right\}.$$

3. Preliminary analysis

Analysing the closed-loop system will require a careful analysis of the estimation algorithm. We define the parameter estimation error by

$$\tilde{\theta}(t) := \hat{\theta}(t) - \theta^*,$$

and the corresponding Lyapunov function associated with $\tilde{\theta}(t)$:

$$V(t) := \tilde{\theta}(t)^T \tilde{\theta}(t).$$

In the following result we list a property of $V(t)$; it is a generalization of what is well-known for the classical algorithm (4).

Proposition 1. *For every $t_0 \in \mathbf{Z}$, $y_0 \in \mathbf{R}$, $\theta_0 \in \mathcal{S}$, $\theta^* \in \mathcal{S}$, $n \in l_\infty$ and $r \in l_\infty$, when the adaptive controller (5)–(7) is applied to the plant (1), the following holds:*

$$V(t) \leq V(t_0) - \frac{1}{2} \sum_{j=t_0, \phi(j) \neq 0}^{t-1} \frac{[\varepsilon(j+1)]^2}{\|\phi(j)\|^2} + 2 \sum_{j=t_0, \phi(j) \neq 0}^{t-1} \frac{[n(j)]^2}{\|\phi(j)\|^2}, \quad t \geq t_0 + 1.$$

Proof. See the Appendix.

In order to use this to prove the main result, we need to analyse the time-varying first-order closed-loop system (8). The following technical result proves useful.

Lemma 2. (i) With $m \in \mathbf{N} \cup \{\infty\}$, *suppose that $a_i \in \mathbf{R}$ and $c > 0$ satisfy*

$$\sum_{i=0}^{m} a_i^2 \leq c. \tag{9}$$

Then for every $\lambda \in (0,1)$, if we define $\gamma := c^{\frac{c+1}{2}} (\frac{1}{\lambda})^{\frac{c}{\lambda^2}+1}$, then the following holds:

$$|\pi_{i=0}^{j-1} a_i| \leq \gamma \lambda^j, \quad j = 0, 1, \ldots, m.$$

[1] The Certainty Equivalence Principle mandates that you use the parameter estimates in the control law as if they are the true parameters — see [18] for a more detailed discussion.