Automatica 78 (2017) 266–273

Contents lists available at ScienceDirect

Automatica

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

Brief paper

On the convergence of iterative learning control☆

M. Mahdi Ghazaei Ardakani[a], Sei Zhen Khong[b,a], Bo Bernhardsson[a]

[a] Department of Automatic Control, Faculty of Engineering (LTH), Lund University, SE 221 00 Lund, Sweden

[b] Institute for Mathematics and its Applications, University of Minnesota, Minneapolis, MN 55455, USA

ARTICLE INFO

Article history:
Received 11 February 2016
Received in revised form
23 August 2016
Accepted 16 November 2016

Keywords:
Learning control
Iterative improvement
Convergence analysis
Transient stability analysis
Stability criteria

ABSTRACT

We derive frequency-domain criteria for the convergence of linear iterative learning control (ILC) on finite-time intervals that are less restrictive than existing ones in the literature. In particular, the former can be used to establish the convergence of ILC in certain cases where the latter are violated. The results cover ILC with non-causal filters and provide insights into the transient behaviors of the algorithm before convergence. We also stipulate some practical rules under which ILC can be applied to a wider range of applications.

1. Introduction

The main application of iterative learning control (ILC) is to improve the reference tracking performance of a system. In order to reduce the tracking error, the control signal to the system is adjusted in each iteration by using feedback information from previous iterations. In effect, ILC finds an approximate system inverse for a specific reference (Moore, Dahleh, & Bhattacharyya, 1989). An advantage of ILC is that it does not require an explicit model of the transfer function or even linearity of the system for finding the inverse. Instead, it often uses the actual system as a part of the algorithm. ILC has found successful applications in many different fields (Ahn, Chen, & Moore, 2007; Freeman, Rogers, Hughes, Burridge, & Meadmore, 2012; Sörnmo, Bernhardsson, Kröling, Gunnarsson, & Tenghamn, 2016), where accurate models of the system and disturbances are difficult to obtain.

While the frequency domain is the preferred approach for filter design and analysis of linear ILC (Wang, Ye, & Zhang, 2014), the widely used convergence criterion applies, only to strictly monotone convergence of the algorithm (the 2-norm of the error between the current control signal and its final value strictly decreases in each iteration). Moreover, it is not theoretically clear to what extent the frequency criterion is applicable to a practical ILC system where each iteration runs only over a finite-time interval and to ILC systems with non-causal filters. To motivate this study, we demonstrate examples for which the ILC converges but the classical frequency condition cannot provide any indication of the convergence property. Our analysis gives an explanation for this mode of convergence.

We extend the work of Norrlöf and Gunnarsson (2002) by introducing a less conservative criterion, hence reducing the gap between the existing time-domain and frequency-domain criteria. We also provide an analysis of the transient behavior of the algorithm, which proves useful when the convergence is not monotone. The contributions of this article can be summarized as follows:

- Analysis of "convergence on finite-time interval" motivated by practical ILC where the trial length is finite.
- A less conservative frequency domain convergence criterion than the one by Norrlöf and Gunnarsson (2002) is derived (see Theorem 8)

  $$\inf_{\rho>0} \sup_{\omega} \left|G(\rho e^{i\omega})\right| < 1.$$

  The criterion is applicable to ILC systems with *causal* as well as *non-causal* filters and for strictly monotone convergence coincides with the classical result.

☆ This work was supported by the Swedish Research Council through LCCC Linnaeus Center and the eLLIIT Excellence Center at Lund University and Institute for Mathematics and its Applications with funds provided by the National Science Foundation. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Changyun Wen under the direction of Editor Miroslav Krstic.

E-mail addresses: mahdi@control.lth.se (M.M. Ghazaei Ardakani), szkhong@umn.edu (S.Z. Khong), bob@control.lth.se (B. Bernhardsson).

http://dx.doi.org/10.1016/j.automatica.2016.12.030
0005-1098/

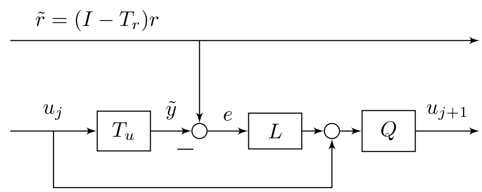

Fig. 1. Block diagram of an iterative learning controller. Here, $\tilde{y} = y - T_r r$.

- The connection between time-domain and frequency-domain criteria is established in a rigorous manner using Toeplitz operators.
- A frequency domain tool for understanding the transient behavior of ILC – i.e., the wave of convergence/divergence – is introduced.
- A strategy to limit the growth of the transient errors when the convergence is not monotone is proposed.

1.1. Previous work

ILC is a two-dimensional process, in the sense that the dynamics are indexed by both time and iteration variables (Kurek & Zaremba, 1993). A standard approach to analysis of linear and a certain class of nonlinear ILC algorithms relies on the lifted-system framework, i.e., considering a time series as a vector (Bristow, Tharayil, & Alleyne, 2006). Norrlöf (2000) has extensively studied the theory and applications of linear ILC. Time-domain criteria as well as a classical frequency-domain criterion for the convergence of the linear ILC algorithm have been derived by Norrlöf and Gunnarsson (2002).

There have been many attempts to understand and improve the convergence properties of the linear ILC. Longman and Huang (2002) have noted that the algorithm might practically diverge after an initial substantial decay of the tracking error. Elci, Longman, Phan, Juang, and Ugoletti (2002) have introduced a non-causal filter, namely a zero-phase filter, in the algorithm to improve the transient behavior. The transient properties of the convergence have been studied in more detail by Longman and Huang (2002) and Wang et al. (2014). Longman (2000) and Norrlöf and Gunnarsson (2002) have commented on the potential convergence of the algorithm despite a transient growth of the norm of the error, i.e., when the classical frequency condition is not fulfilled.

1.2. Problem description

A general form of the discrete linear first-order ILC algorithm is

$$y_j = T_r r + T_u u_j \tag{1}$$

$$e_j = r - y_j \tag{2}$$

$$u_j = Q\left(u_{j-1} + L e_{j-1}\right); \tag{3}$$

see Norrlöf (2000). Here $j \in \mathbb{Z}_{\geq 0}$ is the iteration index, r, y_j, and $e_j \in \ell_2$ are the reference, output, and tracking error signals, respectively, $u_j \in \ell_2$ is the control signal. The stable systems from reference to output and control signal to output are denoted by T_r and T_u, respectively, and Q and L are filters to be designed. The choice of u_0 is free. Fig. 1 depicts the ILC algorithm. Note that in practice the trial length is finite, i.e., the system is stopped after N samples and signal values at time $n \in \{0, \dots, N-1\}$ are stored. The filters Q and L do not need to be causal since they operate on the signals of the previous iteration.

Let us define $G(e^{i\omega}) := Q(e^{i\omega})(1 - L(e^{i\omega})T_u(e^{i\omega}))$. The classical sufficient condition for strictly monotone convergence of ILC requires that (see for example Norrlöf & Gunnarsson, 2002)

$$|G(e^{i\omega})| < 1, \quad \forall \omega \in [0, 2\pi), \tag{4}$$

where $L(e^{i\omega})$, $T_u(e^{i\omega})$, and $Q(e^{i\omega})$ are the frequency representations of the respective filters.

Given the definition of the ILC algorithm in (1)–(3) and the fact that each iteration runs only over a finite-time interval, $n \in \{0, \dots, N-1\}$, our purpose is to find less restrictive conditions for G that guarantee the convergence of the algorithm, i.e., that the limits $u_j \to u_\infty$ and $e_j \to e_\infty$ exist for the finite trial length.

The rest of the article is organized as follows: In Section 2, we present a motivating example for which the ILC converges but the classical condition cannot provide any indication of the convergence property. The iteration-domain dynamics for ILC are derived in Section 3 before we delve into the issue of convergence. Section 4 starts with a formal definition of convergence for iterative procedures and states our convergence results. In Section 5, practical aspects concerning the transient behavior of ILC when the convergence is non-monotone are discussed. We propose qualitative measures that characterize the convergence, and discuss the gap between the time- and frequency-domain criteria in Section 6. We draw conclusions in Section 7. Additionally, a list of useful results and definitions as the background is collected in the Appendix.

2. Motivating example

Let us consider the following transfer functions

$$T_u(s) = \frac{1}{(s+1)(s^2+0.8s+16)}, \qquad T_r(s) = 0, \tag{5}$$

$$Q(s) = \frac{10}{s+10}, \qquad L_d(z) = 10k(1 - 0.9z^{-1})z^a. \tag{6}$$

We discretize $T_u(s)$ and filter $Q(s)$ by the zero-order-hold (ZOH) method (see Åström & Wittenmark, 1997) with sampling time $h = 0.1$ s. Fig. 2 compares the time responses of the systems corresponding to two ILC scenarios where in (1) $k = 0.8, a = 5$ (System I) and in scenario (2) $k = 0.5, a = 8$ (System II). After discretization, Q is implemented as a zero-phase filter and hence we get

$$G(e^{i\omega}) = Q_d(e^{i\omega})Q_d(e^{-i\omega})\left(1 - L_d(e^{i\omega})T_{ud}(e^{i\omega})\right). \tag{7}$$

In Fig. 3, the Bode plots for $G(e^{i\omega})$ are illustrated. We see that in both scenarios the condition $|G(e^{i\omega})| < 1$ is violated. Nevertheless, System I appears to converge, at least for the time region of interest, while System II does not. The Bode diagrams corresponding to convergent and non-convergent scenarios may seem counterintuitive at first glance since the one with the highest peak in the gain $|G(e^{i\omega})|$ corresponds to the convergent case.

Our result in Theorem 8 explains the situation and says that if there exists a $\rho > 0$ such that $\sup_\omega \left|G(\rho e^{i\omega})\right| < 1$, then we have convergence in the sense that u_∞ and e_∞ exist on the finite interval $[0, \dots, N)$. In Fig. 4, where $\sup_\omega |G(\rho e^{i\omega})|$ is plotted against ρ, it can be seen that the curve for System I goes below 1 for some ρ and thus the ILC algorithm converges.

3. Iteration-domain dynamics

In order to analyze the convergence of the ILC system (1)–(3), we derive the dynamics of the system in the iteration domain. Furthermore, to take into account the assumption of the finite-time intervals, we define the truncated counterparts of the original operators.

Define the truncation operator as

$$(\Pi_k x)[n] = \begin{cases} x[n], & n < k \\ 0, & \text{otherwise.} \end{cases} \tag{8}$$

For an operator $G : \ell_2 \to \ell_2$, we define the truncated operator

$$\bar{G} := \Pi_N(I - \Pi_0)G\Pi_N(I - \Pi_0), \tag{9}$$