Control Engineering Practice 26 (2014) 125–135

Contents lists available at ScienceDirect

Control Engineering Practice

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

Pareto iterative learning control: Optimized control for multiple performance objectives ☆

Ingyu Lim, Kira L. Barton *

Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI 48109, USA

ARTICLE INFO

Article history:
Received 30 April 2013
Accepted 10 January 2014
Available online 11 February 2014

Keywords:
Learning control
Pareto optimization
Multi-objective control

ABSTRACT

Iterative learning control (ILC) is a 2-degree-of-freedom technique that seeks to improve system performance along the time and iteration domains. Traditionally, ILC has been implemented to minimize trajectory-tracking errors across an entire cycle period. However, there are applications in which the necessity for improved tracking performance can be limited to a few specific locations. For such systems, a modified learning controller focused on improved tracking at the selected points can be leveraged to address multiple performance metrics, resulting in systems that exhibit significantly improved behaviors across a wide variety of performance metrics. This paper presents a pareto learning control framework that incorporates multiple objectives into a single design architecture.

1. Introduction

Iterative learning control (ILC) is an adaptive control approach in which the adaptation occurs at the input signal rather than as a system or control parameter update (Bristow, Tharayil, & Alleyne, 2006; Moore, Dahley, & Bhattacharyya, 1992). ILC has been successfully applied to repetitive applications in robotics (Arimoto, Kawamura, & Miyazaki, 1984; Tayebi & Islam, 2006), manufacturing (Barton & Alleyne, 2011; Kim & Kim, 1993; Rotariu, Steinbuch, & Ellenbroek, 2008), and chemical processing (Lee & Lee, 2007; Mezghani et al., 2002). In these applications, ILC was implemented to improve the trajectory tracking performance of the system through iterative updates to the control signal. Conventional ILC approaches use the complete error signal from previous iterations to generate an updated control signal for improved system performance (Bristow et al., 2006).

As an alternative to using the complete error signal, a point-based controller focuses on improving the error at discrete locations or times for performance enhancements in applications such as robotic pick n' place tasks (Dijkstra et al., 2001), patient stroke rehabilitation (Freeman et al., 2009), and reconnaissance missions with UAVs (Lim & Bang, 2010). In these application examples, specific locations (e.g. the start and end positions for pick n' place robots) are critical to the success of the task, while the motion profile between the locations is irrelevant. Recent work by Freeman, Cai, Rogers, and Lewin (2011) has resulted in an ILC algorithm termed point-to-point learning control that focuses on specific times or locations of a predetermined motion profile. In point-to-point ILC, the selected points define a subset of the motion profile, $\chi(n_i) \subseteq y_d(k)$, where n_i are the selected points for all $i=1,\ldots,M$, $y_d(k)$ defines the motion profile, and k is the time index. The learning controller only applies a feedforward update to these specified points, $\chi(n_i)$. By removing the unnecessary constraint of a predefined path between the points, additional control freedom can be obtained and redirected towards achieving multiple performance objectives (Fig. 1).

The introduction of multiple performance objectives into a learning framework provides an opportunity to leverage underutilized control actuation to improve system performance, in addition to using learning as a means of improving the system performance. Examples of applications that perform repetitive tasks with multiple performance metrics can be found in manufacturing (metrics: throughput, part quality, material waste); robotics (metrics: speed, precision motion control, power utilization, vibration isolation); and unmanned air/ground vehicles (metrics: path following, patrol efficiency, energy consumption, sensor transmission strength).

Pareto optimization is a commonly employed multi-objective approach in which two or more conflicting objectives are weighted (Yang & Catthoor, 2003) within a single framework. Solutions to this class of problems require a tradeoff in the performance objectives based on the desired design criteria. Tradeoff within a control design is frequently made as a tradeoff between performance and robustness (Boulet & Duan, 2007; Jin & Sendhoff, 2003), or as a single performance objective optimization within a constrained system (Mishra, Topcu, & Tomizuka, 2011). Recent work by the authors presented a pareto learning controller for addressing multiple objectives with systems that perform repetitive tasks. This initial work presented the basic framework, but did not provide a tradeoff analysis or experimental validation (Lim & Barton, 2013).

☆ This work was supported by start-up funds from the University of Michigan.
* Corresponding author.
E-mail addresses: ingyulim@umich.edu (I. Lim), bartonkl@umich.edu (K.L. Barton).

0967-0661/$ - see front matter © 2014 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.conengprac.2014.01.011

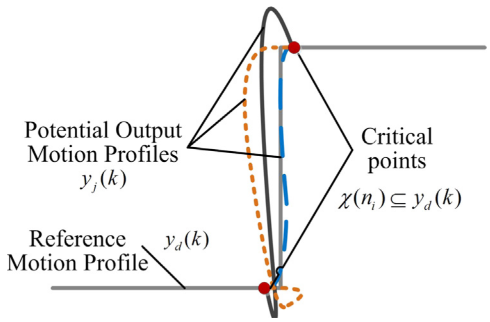

Fig. 1. Illustration of the point-to-point tracking problem. The control objective requires precise tracking only at the critical points. This additional flexibility results in multiple potential solutions that can be optimized for speed, energy usage, attack angle, and robustness.

Recently published papers presented a dual optimization approach to multi-objective learning (Freeman & Tan, 2013; Freeman, 2012; Owens, Freeman, & Chu, 2013). These papers utilize a two-step approach to optimizing the overall performance as close to zero steady-state tracking as possible. In step 1, the framework obtains an optimal control solution for zero steady-state trajectory tracking. In step 2, the framework seeks to optimize the performance of an additional objective through the use of a cost function that considers the additional objective while simultaneously minimizing the difference between a new control input and the optimal control signal determined in step 1. This iterative learning sequence involves multiple steps, while bounding the range of the new solution to be arbitrarily close to the initial optimal input.

In this paper, we present a generalized multi-objective learning control framework for systems that require the optimization of multiple performance objectives simultaneously. To address the performance requirements for these types of systems, the control objectives are posed as a pareto optimization-based learning problem where the controller seeks to optimize a cost function containing multiple performance objectives. As a result of our one-step optimization approach, tradeoffs between trajectory tracking and additional objectives can be clearly observed. Additionally, the optimization search is implemented over a broad set of potential solutions, thus enabling a greater variety of possible outcomes. This research extends the work provided in Lim and Barton (2013) through several notable modifications: (1) the modified learning controller with cost function convergence and bounded system outputs is provided, (2) a performance tradeoff analysis is included to evaluate potential design choices for energy reduction, (3) a detailed design methodology has been provided, and (4) simulation and experimental results validate the controller performance and provide a means for verifying trends in the system behavior.

2. Class of systems

For clarity of exposition, the class of systems considered for this work includes linear, causal, discrete-time single-input, single-output (SISO) systems, given as

$$\mathbf{H} \triangleq \begin{cases} \mathbf{x}_j(k+1) = \mathbf{A}\mathbf{x}_j(k) + \mathbf{B}u_j(k) \\ y_j(k) = \mathbf{C}\mathbf{x}_j(k) \end{cases} \tag{1}$$

where $\mathbf{x}(k) \in \mathbb{R}^p$ are the system states, $u(k) \in \mathbb{R}$ is the control input, $y(k) \in \mathbb{R}$ is the output, $k \in \mathbb{Z}^{N+1}$ is the time index, and $j = 1, 2, \ldots$ is the iteration index. ($\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$) are appropriately sized iteration-invariant real-valued matrices. It is assumed that $\mathbf{x}_j(0) = \mathbf{x}_0$ for all j. As defined by the matrices ($\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$), $\mathbf{H}$ is time-invariant over a single profile and iteration-invariant from trial-to-trial. In the lifted-domain, the discrete-time behavior of the system is represented by its convolution matrix using impulse response data $\mathbf{H}_{m,n}$. The lifted-system representation transforms a two-dimensional (time and iteration) system into a one-dimensional (iteration) system. The lifted system representation is given by

$$\mathbf{H} = \begin{bmatrix} H_{0,0} & & 0 \\ \vdots & \ddots & \\ H_{N-1,0} & \cdots & H_{N-1,N-1} \end{bmatrix} \tag{2}$$

For LTI systems, $\mathbf{H}_{m,n}$ contains the impulse response data and can be derived using the matrices in (1)

$$\mathbf{H}_{m,n} : \{\mathbf{C}\mathbf{A}^{m-n}\mathbf{B}, \quad m \geq n \tag{3}$$

While the results presented in this paper are for an LTI system, the same design process can be applied to LTV systems. In the case of LTV systems, $\mathbf{H}_{m,n}$ is of the form

$$\mathbf{H}_{m,n} : \begin{cases} \mathbf{C}(n)\mathbf{B}(n), & m = n \\ \mathbf{C}(m)\mathbf{A}(m-1)\mathbf{A}(m-2)\ldots\mathbf{A}(n)\mathbf{B}(n), & m > n \end{cases} \tag{4}$$

3. Norm optimal ILC

This work adopts the widely used norm optimal iterative learning control (NOILC) approach (Amann, Owens, & Rogers, 1996; Bristow & Hencey, 2008; van de Wijdeven & Bosgra, 2008). The norm optimal approach was chosen for its monotonic convergence guarantees and design tradeoff abilities, such as the intuitive weighting structure and modal architecture that enable weighting of multiple objectives. The general and point-to-point based norm optimal frameworks are briefly described here, and will be extended to enable design modifications for systems with multiple performance objectives.

3.1. Conventional ILC

In this paper, a well-known norm optimal ILC update law is adopted (Bristow, Barton, & Alleyne, 2010)

$$\mathbf{u}_{j+1} = \mathbf{L}_u\mathbf{u}_j + \mathbf{L}_e\mathbf{e}_j \tag{5}$$

where

$$\mathbf{e}_j = \mathbf{y}_d - \mathbf{y}_j = \mathbf{y}_d - \mathbf{H}\mathbf{u}_j \tag{6}$$

with

$$\mathbf{e}_j = [e_j^T(1) \;\; e_j^T(2) \;\; \cdots \;\; e_j^T(N)]^T \tag{7}$$

$$\mathbf{u}_j = [u_j^T(0) \;\; u_j^T(1) \;\; \cdots \;\; u_j^T(N-1)]^T. \tag{8}$$

The norm optimal ILC algorithm is designed to minimize a quadratic optimization problem, in which the objective is to minimize a cost function (Phan & Longman, 1988)

$$J = \mathbf{e}_{j+1}^T\mathbf{Q}\mathbf{e}_{j+1} + \mathbf{u}_{j+1}^T\mathbf{S}\mathbf{u}_{j+1} + (\mathbf{u}_{j+1} - \mathbf{u}_j)^T\mathbf{R}(\mathbf{u}_{j+1} - \mathbf{u}_j). \tag{9}$$

($\mathbf{Q}$, $\mathbf{S}$, $\mathbf{R}$) are symmetric positive definite matrices with a common form given as $(\mathbf{Q}, \mathbf{S}, \mathbf{R}) \triangleq (q\mathbf{I}, s\mathbf{I}, r\mathbf{I})$. Minimizing the cost function J with respect to $\mathbf{u}_{j+1}$ yields the norm optimal ILC update algorithm filters with respect to the weighting matrices ($\mathbf{Q}$, $\mathbf{S}$, $\mathbf{R}$) and the plant $\mathbf{H}$ (van de Wijdeven & Bosgra, 2008).

$$\mathbf{L}_u = (\mathbf{H}^T\mathbf{Q}\mathbf{H} + \mathbf{S} + \mathbf{R})^{-1}(\mathbf{H}^T\mathbf{Q}\mathbf{H} + \mathbf{R}) \tag{10}$$

$$\mathbf{L}_e = (\mathbf{H}^T\mathbf{Q}\mathbf{H} + \mathbf{S} + \mathbf{R})^{-1}\mathbf{H}^T\mathbf{Q} \tag{11}$$

Note that for $(\mathbf{L}_u, \mathbf{L}_e)$ to ensure convergence, $\mathbf{H}^T\mathbf{Q}\mathbf{H} + \mathbf{S} + \mathbf{R}$ must be positive definite (Barton & Alleyne, 2011).