Automatica 50 (2014) 2561–2569

Contents lists available at ScienceDirect

Automatica

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

Brief paper

On finite-time and infinite-time cost improvement of economic model predictive control for nonlinear systems[☆]

Matthew Ellis[a], Panagiotis D. Christofides[a,b,1]

[a] Department of Chemical and Biomolecular Engineering, University of California, Los Angeles, CA 90095-1592, USA

[b] Department of Electrical Engineering, University of California, Los Angeles, CA 90095-1592, USA

ARTICLE INFO

Article history:
Received 8 October 2013
Received in revised form
9 May 2014
Accepted 29 May 2014
Available online 31 August 2014

Keywords:
Economic model predictive control
Process economics
Nonlinear systems
Dynamic process optimization

ABSTRACT

A novel two-layer economic model predictive control (EMPC) structure that addresses provable finite-time and infinite-time closed-loop economic performance of nonlinear systems in closed-loop with EMPC is presented. In the upper layer, a Lyapunov-based EMPC (LEMPC) scheme is formulated with performance constraints by taking advantage of an auxiliary Lyapunov-based model predictive control (LMPC) problem solution formulated with a quadratic cost function. The lower layer LEMPC uses a shorter prediction horizon and smaller sampling period than the upper layer LEMPC and involves explicit performance-based constraints computed by the upper layer LEMPC. Thus, the two-layer architecture allows for dividing dynamic optimization and control tasks into two layers for a computationally manageable control scheme at the feedback control (lower) layer. A chemical process example is used to demonstrate the performance and stability properties of the two-layer LEMPC structure.

1. Introduction

Within process control, economic model predictive control (EMPC) has ignited wide-spread interest because of its unique ability to dynamically regulate processes to achieve closed-loop economic performance not attainable through traditional tracking control techniques (Adetola & Guay, 2010; Amrit, Rawlings, & Angeli, 2011; Angeli, Amrit, & Rawlings, 2012; Baldea & Touretzky, 2013; Diehl, Amrit, & Rawlings, 2011; Fagiano & Teel, 2013; Ferramosca, Rawlings, Limon, & Camacho, 2010; Grüne, 2013; Guay & Adetola, 2013; Heidarinejad, Liu, & Christofides, 2012, 2013; Huang, Biegler, & Harinath, 2012; Idris & Engell, 2012; Ma, Qin, Salsbury, & Xu, 2012; Müller, Angeli, & Allgöwer, 2013; Omell & Chmielewski, 2013). The fundamental difference between EMPC and conventional model predictive control (MPC) is the cost function used in the formulations of these two control schemes. Typically, in conventional MPC schemes, a quadratic cost function that penalizes a weighted error of states and inputs from their economically optimal steady-state values is typically used, while, EMPC schemes use a general cost function that is derived from the process economics (e.g., operating cost or profit). As a result of the type of cost function used, EMPC can handle both dynamic process economic optimization and process control. To utilize EMPC for the computation of optimal inputs in real-time, EMPC is formulated with a finite prediction horizon.

An important, albeit not well understood property, is the closed-loop performance of systems under EMPC since EMPC is formulated with a finite prediction horizon. The main results on closed-loop performance with EMPC include: (1) EMPC formulated with a terminal constraint has asymptotic (infinite-time) average performance at least as good as the economically optimal steady-state (Angeli et al., 2012) (others have extended asymptotic average performance to various EMPC formulations Amrit et al., 2011, Fagiano & Teel, 2013, Müller et al., 2013), (2) EMPC formulated without any (terminal) constraints was shown to be approximately optimal for both finite-time (i.e., transient) and infinite-time when a sufficiently long prediction horizon is used and certain controllability assumptions are satisfied (Grüne, 2013), and (3) a Lyapunov-based EMPC (LEMPC) which uses performance constraints derived from an auxiliary conventional (tracking) MPC and a shrinking horizon to guarantee that over a finite operating window the closed-loop performance under LEMPC is at least as good as the auxiliary conventional MPC (Heidarinejad et al., 2013). In Angeli

☆ Financial support from the National Science Foundation and the U.S. Department of Energy is gratefully acknowledged. The material in this paper was partially presented at the American Control Conference (ACC2014), June 4–6, 2014, Portland, Oregon, USA. This paper was recommended for publication in revised form by Associate Editor David Angeli under the direction of Editor Andrew R. Teel.

E-mail addresses: matthewellis@ucla.edu (M. Ellis), pdc@seas.ucla.edu (P.D. Christofides).

[1] Tel.: +1 310 794 1015; fax: +1 310 206 4107.

http://dx.doi.org/10.1016/j.automatica.2014.08.011
0005-1098/

et al. (2012), the effect of the initial condition is essentially neglected since it is insignificant when considering operation for an infinite-time period. Given the power of EMPC to yield dynamically optimal regulation of systems operating away from steady-state, the importance of considering the effect of the initial condition on closed-loop performance should be considered as this is an important property of EMPC. In Heidarinejad et al. (2013), on the other hand, guarantees on closed-loop performance can only be made over finite operating windows. Therefore, introducing an EMPC structure that provides provable finite-time (i.e., accounts for the effect of the initial condition) and infinite-time performance guarantees on closed-loop economic performance is an important issue.

Another challenge of EMPC is that the achievable closed-loop economic performance benefit of EMPC over conventional tracking MPC may strongly depend on the prediction horizon length (e.g., Grüne, 2013). A long prediction horizon (i.e., many decision variables), however, may make it difficult to solve the EMPC optimization problem for real-time applications. To address guaranteed closed-loop economic performance while formulating a computationally efficient control structure, a novel two-layer LEMPC structure is proposed in this work. The core idea of the two-layer EMPC is to solve the upper layer LEMPC infrequently (i.e., not every sampling period) over a long horizon. Then, take advantage of the solution generated by the upper layer LEMPC in the formulation of a lower layer LEMPC used for feedback control. Specifically, in the upper layer, an LEMPC, formulated with a sufficiently large prediction horizon, is used to compute economically optimal trajectories which are sent down to the lower layer LEMPC. The lower layer LEMPC uses a shorter prediction horizon and smaller sampling time than the upper layer LEMPC to compute control actions for the process in real-time while maintaining operation around the economically optimal trajectories computed in the upper layer. For guaranteed performance improvement with the proposed LEMPC scheme, both layers are formulated with explicit performance-based constraints computed from an auxiliary Lyapunov-based model predictive control (LMPC) problem solution formulated with a quadratic cost which allows for provable finite-time and infinite-time closed-loop economic performance and effectively merges the provable performance guarantees on finite-time and infinite-time performance compared to a conventional (tracking) MPC. The two-layer LEMPC structure is applied to a chemical process example to demonstrate the closed-loop performance, stability, and robustness properties of the two-layer LEMPC structure.

2. Preliminaries

2.1. Class of systems

The class of continuous-time nonlinear systems considered is described by the following state-space form:

$$\dot{x}(t) = f(x(t), u(t)) \tag{1}$$

where the state is $x(t) \in \mathbf{R}^n$ and the input is $u(t) \in \mathbf{R}^m$. The vector function $f : \mathbf{R}^n \times \mathbf{R}^m \to \mathbf{R}^n$ is a locally Lipschitz vector function on $\mathbf{R}^n \times \mathbf{R}^m$. The available control effort is defined by the convex set $U = \{u_{\min} \le u \le u_{\max}\} \subset \mathbf{R}^m$. The state x of the system is synchronously sampled at time instances $t_0 + k\Delta$ with $k = 0, 1, 2, \ldots$ where t_0 is the initial time and Δ is the sampling period. Without loss of generality, the initial time is taken to be zero ($t_0 = 0$). To distinguish between the continuous time and the discrete sampling instances, the notation t will be used for the continuous time and the time sequence $\{\tau_k\}_{k=0}^{\infty}$ is the partitioning of t with $\tau_k = k\Delta$.

A time-invariant economic cost measure $l_e(x, u)$ is assumed to describe the real-time economics of the system of equation (1) and is assumed to be continuous on $X \times U$ where $X \subset \mathbf{R}^n$ is the set of admissible operating states. The optimal steady-state x_s^* and steady-state input u_s^* with respect to the economic cost function is $(x_s^*, u_s^*) = \arg\max_{u_s \in U,\, x_s \in X} \{l_e(x_s, u_s) : f(x_s, u_s) = 0\}$. For the sake of simplicity, the optimal steady-state is assumed to be unique and to be $(x_s^*, u_s^*) = (0, 0)$. Furthermore, the notation $|\cdot|$ denotes the Euclidean norm of a vector, the notation $|\cdot|_Q$ denotes the square of a weighted Euclidean norm of a vector (i.e., $|x|_Q = x^T Q x$ where Q is a positive definite matrix), and the symbol Ω_ρ denotes a level set of a Lyapunov function (i.e., $\Omega_\rho = \{y \in \mathbf{R}^n : V(y) \le \rho\}$).

2.2. Existence of a stabilizing controller

Assumption 1. There exists a locally Lipschitz feedback controller $u = h(x)$ with $h(0) = 0$ for the system of equation (1) that renders the origin of the closed-loop system under continuous implementation of the controller $h(x)$ locally exponentially stable. More specifically, there exist constants $\rho > 0$, $c_i > 0$, $i = 1, 2, 3, 4$ and a continuously differentiable Lyapunov function $V : \mathbf{R}^n \to \mathbf{R}_+$ such that the following inequalities hold for all $x \in \Omega_\rho$:

$$c_1 |x|^2 \le V(x) \le c_2 |x|^2, \tag{2a}$$

$$\frac{\partial V(x)}{\partial x} f(x, h(x)) \le -c_3 |x|^2, \tag{2b}$$

$$\left| \frac{\partial V(x)}{\partial x} \right| \le c_4 |x|, \tag{2c}$$

for all $x \in \Omega_\rho$.

Explicit feedback controllers that may be designed to satisfy Assumption 1 are, for example, feedback linearizing controller and some Lyapunov-based controllers (e.g., Khalil, 2002, Kokotović & Arcak, 2001). With the controller $h(x)$, the following results hold for the closed-loop system of equation (1) under the controller $h(x)$ implemented in a zero-order sample-and-hold fashion with sampling period Δ (i.e., $h(x)$ is applied as an emulation controller).

Proposition 2. *Suppose Assumption 1 holds. Then, there exists $\Delta^* > 0$ and $M, \sigma > 0$ such that for the partition $\{\tau_i\}_{i=0}^{\infty}$ of $\mathbf{R}_+$ with $\tau_{i+1} - \tau_i = \Delta \le \Delta^*$ the closed-loop system of equation (1) with the input trajectory*

$$u(t) = h(x(\tau_i)) \quad \text{for } t \in [\tau_i, \tau_{i+1}) \text{ and integers } i \ge 0 \tag{3}$$

and arbitrary initial condition $x(0) = x_0 \in \Omega_\rho$ satisfies the estimate:

$$|x(t)| \le M \exp(-\sigma t) |x_0| \tag{4}$$

for all $t \ge 0$.

The proof of Proposition 2 may be found in Ellis et al. (2014, Corollary 1) and shows that V is a Lyapunov function for the closed-loop sampled-data system in the sense that there exists a constant $\hat{c}_3 > 0$ such that

$$\frac{\partial V(x(t))}{\partial x} f(x(t), h(x(\tau_i))) \le -\hat{c}_3 |x(t)|^2 \tag{5}$$

for all $t \in [\tau_i, \tau_{i+1})$ and integers $i \ge 0$ where $x(t)$ is the solution of Eq. (1) starting from $x(\tau_i) \in \Omega_\rho$ and with the input $u(t) = h(x(\tau_i))$ for $t \in [\tau_i, \tau_{i+1})$. The stability region of the closed-loop system under the controller $h(x)$ is defined as $\Omega_\rho \subseteq X$.

Remark 3. Assumption 1 is stronger than the one imposed in our previous works (e.g., Christofides, Liu, & Muñoz de la Peña, 2011, Heidarinejad et al., 2012). In the present works, the existence of a controller $h(x)$ that renders the origin of the closed-loop system locally exponentially stable under continuous implementation is assumed whereas, in Christofides et al. (2011) and Heidarinejad et al.