Annual Reviews in Control 41 (2016) 218–224

Contents lists available at ScienceDirect

Annual Reviews in Control

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

Full Length Article

Theoretical advances on Economic Model Predictive Control with time-varying costs

David Angeli[a,b,*], Alessandro Casavola[c], Francesco Tedesco[c]

[a] Department of Electrical and Electronic Engineering, Imperial College London, SW7 2AZ, UK
[b] Department of Information Engineering, University of Florence, Italy
[c] Department of Information Engineering, Modeling, Electronic and Systems, University of Calabria, Rende(CS), Italy

ARTICLE INFO

Article history:
Received 22 December 2015
Revised 25 February 2016
Accepted 5 March 2016

Keywords:
Economic model predictive control
Time-varying costs
Constrained control
Optimization based control

ABSTRACT

Economic Model Predictive Control is a technique for optimization of economic revenues arising from controlled dynamical processes that has established itself as a variant of standard Tracking Model Predictive Control. It departs from the latter in that arbitrary cost functions are allowed in the formulation of the stage cost. This paper takes a further step in expanding the applicability of Economic Model Predictive Control by illustrating how the paradigm can be adapted in order to accommodate time-varying or parameter-varying costs.

1. Introduction and motivations

Model Predictive Control is a model-based control design technique for MIMO systems subject to input and state constraints. In its classical formulation it allows to formulate general tracking problems for nonlinear and/or linear systems by taking into account model-based predictions throughout a finite control horizon and setting up the control selection algorithm as an on-line optimization problem where the adopted cost function is a measure of the discrepancy of the predicted trajectory with respect to the desired set-point signal.

In recent years, Economic MPC has emerged as a variant of Model Predictive Control where the primary control task is profitability enhancement, rather than minimization of a tracking error. From a mathematical perspective, this amounts to considering cost functionals, defined over a typically finite control horizon, which can take arbitrary shape, rather than being limited to some (positive definite) distance function from a set-point of interest. In particular, for nonlinear systems and/or non convex cost functionals, profitability may be maximal away from equilibrium states and this may in turn lead to complex regimes of operation or transient behaviors which may exhibit highly nonlinear features, such as asymmetry with respect to initial conditions or slow and highly oscillatory decays.

Motivated by applications in areas where price variations are comparable in speed with process dynamics, the case of time-varying costs or parameter-varying costs were recently explored in Ellis and Christofides (2014). The method developed is a Lyapunov-based Economic MPC scheme which allows to guarantee boundedness of solutions as well as constraints satisfaction while attempting to optimize a time-varying cost functional.

As a matter of fact, in recent years, many efforts have been devoted to investigate Economic MPC variants allowing time-varying costs in several domains of application: management of energy in buildings (Ma, Qin, & Salsbury, 2014; Touretzky & Baldea 2014), control of chemical plants (Ellis & Christofides, 2014) and supervision of distribution networks, such as water networks (Grosso, Ocampo-Martinez, Puig, Limon, & Pereira, 2014), power grids (Adeodu & Chmielewski, 2013; Cole, Morton, & Edgar, 2014; Hovgaard, Edlund, & Bagterp Jorgensen 2010) gas networks, etc (Gopalakrishnan & Biegler, 2013). Other Economic MPC approaches dealing either time-varying cost or cyclic plant operations from a theoretical perspective can be found in Ferramosca, Limon, and Camacho (2014); Huang, Biegler, and Harinath (2012); Limon, Pereira, Muñoz de la Peña, Alamo, and Grosso (2014). Indeed, the extension of Economic Model Predictive Control to encompass cost variability appears to be a natural question, both from a practical and theoretical perspective, and conceptually similar to the traditional departure from set-point tracking towards tracking of more general reference trajectories.

* Corresponding author at: Department of Electrical and Electronic Engineering, Imperial College London, SW7 2AZ, UK. Tel.: +44 (0)20 7594 6283.
E-mail addresses: d.angeli@imperial.ac.uk (D. Angeli), alessandro.casavola@unical.it (A. Casavola), francesco.tedesco@unical.it (F. Tedesco).

http://dx.doi.org/10.1016/j.arcontrol.2016.04.003
1367-5788/

Two recent papers, Angeli, Casavola, and Tedesco (2015a, 2015b), have attempted to extend the Economic Model Predictive Control analysis framework of Angeli, Amrit, and Rawlings (2012) to this set-up. Hereby we recall the control algorithms and the main results of Angeli et al. (2015a, 2015b) while providing a comparison of their main merits and limitations, together with a simulated case study where both approaches are tested and evaluated against each other. This manuscript is an extended version of Angeli, Casavola, and Tedesco (2015c) suitably edited for this special issue.

2. Preliminaries and problem set-up

The basic formulation of Economic Model Predictive Control deals with discrete-time systems of the following form:

$$x(t+1) = f(x(t), u(t)), \qquad x(0) = x_0 \tag{1}$$

with $t \in \mathbb{N}$, state variable $x \in \mathcal{X} \subset \mathbb{R}^n$, control input $u \in \mathcal{U} \subset \mathbb{R}^m$ and continuous state-transition map $f : \mathcal{X} \times \mathcal{U} \to \mathcal{X}$. Additionally, system's evolutions are subject to pointwise-in-time constraints involving both states and input variables,

$$(x(t), u(t)) \subset \mathcal{Z} \qquad \forall t \geq 0 \tag{2}$$

for some compact set $\mathcal{Z} \subset \mathcal{X} \times \mathcal{U}$. The control task is to fulfill constraints (2) while, at the same time, minimizing a cost functional defined integrating over time an instantaneous (continuous) stage cost ℓ defined as:

$$\ell(x, u) : \mathcal{X} \times \mathcal{U} \to \mathbb{R}. \tag{3}$$

For the case of Tracking Model Predictive Control, the stage cost ℓ typically takes the form of a quadratic function $x'Qx + u'Ru$, or a shifted version of this, if an equilibrium either than 0 is the desired target state.

In Economic Model Predictive Control, on the contrary, ℓ may take an arbitrary shape, and this, in turn, can affect considerably the optimal regimes of operations for the system. Notice that the basic formulation of Economic Model Predictive Control only entails time-invariant "ingredients", viz. dynamics, (1), operational constraints, (2), and operational costs, (3). While it is conceivable to allow all of them to be time-varying, we argue that, in many applications of interest, dynamics are in fact time-invariant, while the only significant source of variability happens at the level of both cost and constraints.

This is because a plant often operates in a manner that does not change in time, apart from deteriorating phenomena that are normally much slower than the time-scales of interest. In this respect, only the environment the plant is interacting with may experience faster and significant variations. Moreover, if we are talking about an 'economic environment', rather than a physical one, time-varying constraints are typically not safety critical, and can often be modeled as *soft constraints*, that is, as suitable cost penalties incurred only in case of constraints violation.

These considerations allow to remarkably simplify the set-up of a time-varying Economic Model Predictive Control scheme and the associated analysis. They allow, in fact, to avoid feasibility issues and associated technical complications that are known to occur whenever time-variability affects constraints or dynamics. On the grounds of such considerations, we consider next two possible modifications of (3), in order to accommodate time-varying or parameter-varying stage costs. Namely, we allow ℓ to directly depend on time:

$$\ell(t, x, u) : \mathbb{N} \times \mathcal{X} \times \mathcal{U} \to \mathbb{R} \tag{4}$$

or indirectly through a time-varying parameter θ taking up finitely many values in $\Theta := \{\theta_1, \theta_2, \ldots, \theta_N\}$:

$$\tilde{\ell}(\theta, x, u) : \Theta \times \mathcal{X} \times \mathcal{U} \to \mathbb{R}. \tag{5}$$

As expected, we are going to formulate the control selection policy as the solution of an associated optimization problem to be performed on-line at each sampling time on the basis of the current knowledge of state and future predictions of both systems trajectories and stage costs variations.

Notice that, also in the case of parameter-varying stage costs, once the time evolution of the parameter θ is assigned, one may define a corresponding time-varying cost simply by composition of the functions $\tilde{\ell}$ and $\theta(t)$, by letting:

$$\ell(t, x, u) := \tilde{\ell}(\theta(t), x, u). \tag{6}$$

This notation allows to formulate a unified cost functional for both cases, by considering:

$$J_H(t, \mathbf{x}, \mathbf{u}) = \sum_{k=0}^{H-1} \ell(t+k, x(k), u(k)), \tag{7}$$

where $\mathbf{x} := [x(0), x(1), \ldots, x(H-1), x(H)]$ denotes the sequence of predicted states, $\mathbf{u} = [u(0), \ldots, u(H-1)]$ that of predicted controls and H denotes the prediction horizon. In both scenarios, the input u is selected at each sampling time t by finding a solution to the following optimization problem:

$$\begin{aligned} J_H^\star(t, x(t), x_F(t)) := &\min_{\mathbf{z},\mathbf{v}} J_H(t, \mathbf{z}, \mathbf{v}) \\ &\text{subject to :} \\ &z(k+1) = f(z(k), v(k)), \\ &(z(k), v(k)) \in \mathcal{Z}, k = 0, \ldots, H-1 \\ &z(H) = x_F(t) \\ &z(0) = x(t) \end{aligned} \tag{8}$$

and letting $u(t)$ correspond to the first input value of any optimal sequence $\mathbf{v}^\star$ solution of (8). Notice that, while continuity and compactness considerations guarantee existence of an optimal solution, this might in general be non unique.

The terminal equality constraint, often used in MPC as a means to achieve recursive feasibility (and possibly stability), is specified in (8) as a function of a suitably defined feasible trajectory $x_F(t)$. Indeed, $x_F(t)$ can be selected according to different criteria. However, existence of an input $u_F(t)$ fulfilling:

(1) $x_F(t+1) = f(x_F(t), u_F(t)) \qquad \forall t \geq 0,$
(2) $(x_F(t), u_F(t)) \in \mathcal{Z} \qquad \forall t \geq 0,$

is crucial to the following developments. Defining the sequence of terminal constraints according to a feasible solution, in fact, allows to easily achieve two important goals: recursive feasibility and guaranteed optimal performance. We recall next one of the technical Lemmas in Angeli et al. (2015a).

Lemma 1. *Consider system (1), controlled by the following time-varying state-feedback,*

$$u(t) = \mathbf{v}_0^\star(t, x(t))$$

where $\mathbf{v}^\star(t, x(t))$ denotes the solution, at each time t, of the optimization problem (8). Then, if a feasible solution exists at time 0, problem (8) is feasible at each subsequent time t. Moreover, the asymptotic average cost of closed-loop trajectories is bounded from above as follows:

$$\limsup_{\tau \to +\infty} \frac{\sum_{t=0}^{\tau-1} \ell(t, x(t), u(t))}{\tau} \leq \limsup_{\tau \to +\infty} \frac{\sum_{t=0}^{\tau-1} \ell(t+H, x_F(t), u_F(t))}{\tau}.$$

Proof. Let $\mathbf{z}^\star(t, x(t))$ and $\mathbf{v}^\star(t, x(t))$ denote the optimal state and control sequences for problem (8) at time t. Then, at time $t+1$, the *shifted* state and control sequences $\tilde{\mathbf{z}} = [\mathbf{z}^\star_{1:H}(t, x(t)), x_F(t+1)]$,