Automatica 82 (2017) 269–277

Contents lists available at ScienceDirect

Automatica

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

Quadratic costs do not always work in MPC[☆]

Matthias A. Müller[a], Karl Worthmann[b,1]

[a] Institute for Systems Theory and Automatic Control, University of Stuttgart, 70550 Stuttgart, Germany
[b] Institute for Mathematics, Technische Universität Ilmenau, 98693 Ilmenau, Germany

ARTICLE INFO

Article history:
Received 10 July 2016
Received in revised form 22 December 2016
Accepted 10 April 2017

Keywords:
Model predictive control
Nonlinear systems
Mobile robots
Asymptotic stabilization
Quadratic costs

ABSTRACT

We consider model predictive control (MPC) without terminal costs and constraints. Firstly, we rigorously show that MPC based on quadratic stage costs may fail, i.e., there does not exist a prediction horizon length such that a (controlled) equilibrium is asymptotically stable for the MPC closed loop although the system is, e.g., finite time controllable. Hence, stability properties of the infinite horizon optimal control problem are, in general, *not* preserved in MPC as long as purely quadratic costs are employed. This shows the necessity of using the stage cost as a design parameter to achieve asymptotic stability. Furthermore, we relax the standard controllability assumption employed in MPC without terminal costs and constraints to alleviate its verification.

1. Introduction

Model predictive control (MPC) is nowadays a well-established control methodology — both from a more theoretical point of view (Grüne & Pannek, 2017; Lee, 2011; Rawlings & Mayne, 2009) and in many different fields of application (Camacho & Bordons, 2012; Qin & Badgwell, 2003; Rodriguez, Kazmierkowski, Espinoza, Zanchetta, Abu-Rub, Young, & et al., 2013). One of the key drivers for its success story is the simplicity of the basic idea: measure the current state, solve a finite-horizon optimal control problem online, and implement the first portion of the computed control strategy. Then, this loop is iteratively repeated ad infinitum to generate an input signal on the infinite time horizon. MPC is particularly attractive due to its capability to deal with constrained multi-input multi-output systems. However, its stability analysis is far from being trivial and closed-loop stability is not necessarily guaranteed if the MPC controller is not designed appropriately, see, e.g., Raff, Huber, Nagy, and Allgöwer (2006). In order to establish asymptotic stability of the MPC closed loop, there are two main approaches available in the literature. The first is to impose suitable additional terminal constraints and terminal costs in the repeatedly solved optimal control problem, see, e.g., Chen and Allgöwer (1998) and Mayne, Rawlings, Rao, and Scokaert (2000) or the textbook (Rawlings & Mayne, 2009). Alternatively, a certain *controllability assumption* is required (Grimm, Messina, Tuna, & Teel, 2005; Grüne, 2009; Primbs & Nevistić, 2000) in order to avoid the necessity to use such (artificial) terminal ingredients.

The goal of this paper is to shed light on various aspects and assumptions of MPC without terminal cost and terminal constraints. Within this setting, typically a certain controllability assumption is used, which is formulated in terms of an upper bound on the optimal value function, see, e.g., Grüne, Pannek, Seehafer, and Worthmann (2010), Grüne and Rantzer (2008), Reble and Allgöwer (2012) and Tuna, Messina, and Teel (2006). While such a controllability condition can be used to establish asymptotic stability of the MPC closed loop, its verification is in general a difficult task, see, e.g., Worthmann, Mehrez, Zanon, Gosine, Mann, and Diehl (2016) for a non-trivial example. As a first main contribution, we weaken this controllability condition, which might help to alleviate this difficulty. In doing so, we also show that this new relaxed condition is *sharp*.

Typically, MPC is applied to solve set point stabilization (tracking) problems. To this end, stage (running) costs are constructed such that the deviation from the desired state and the control effort are penalized. A typical choice in industrial practice is to use quadratic cost functions, i.e., the distance from the set point is weighted quadratically. As a second main contribution, we illustrate via a simple example (the nonholonomic integrator/robot) that when using such a quadratic stage cost, MPC might in general *not* be stabilizing — independent of the length of the prediction horizon. This means that stability properties of the infinite horizon optimal control problem are not necessarily preserved — even for

☆ M.A. Müller and K. Worthmann are supported by the Deutsche Forschungsgemeinschaft, Grants WO 2056/1-1 and WO 2056/4-1. The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Giancarlo Ferrari-Trecate under the direction of Editor Ian R. Petersen.

E-mail addresses: matthias.mueller@ist.uni-stuttgart.de (M.A. Müller), karl.worthmann@tu-ilmenau.de (K. Worthmann).
[1] Fax: +49 3677 693270.

http://dx.doi.org/10.1016/j.automatica.2017.04.058

a very large prediction horizon; rather the stage cost has to be suitably chosen to ensure asymptotic stability.

Finally, we discuss sufficient conditions under which quadratic stage cost functions can be used, i.e., under which the above described situation that quadratic stage cost functions might fail cannot occur. Besides the well known case where the linearization is stabilizable (Chen & Allgöwer, 1998), we present a (nonlinear) local controllability condition which is suitable to this end.

The remainder of this paper is structured as follows. In Section 2, we introduce the considered problem setup and briefly recall stability results in MPC without terminal constraints and costs. In Section 3, we show how the standard controllability condition used in this setting can be relaxed and discuss implications of this relaxation for closed-loop performance statements. Section 4 shows that MPC with a quadratic stage cost might *not* be stabilizing independent of the length of the prediction horizon, before a sufficient condition is derived under which quadratic cost functions *work*. Finally, Section 5 concludes the paper.

Notation. $\mathbb{N}$ and $\mathbb{R}_{\geq 0}$ denote the natural and the non-negative real numbers. $\mathcal{B}_r(x)$ denotes the ball $\{y \in \mathbb{R}^n \,|\, \|y - x\| \leq r\}$ of radius $r \in \mathbb{R}_{>0}$ centered around $x \in \mathbb{R}^n$, where $\|\cdot\|$ is the Euclidean norm of the vector x.

2. Problem formulation

In this work, nonlinear systems governed by ordinary differential equations of the form

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

with continuous vector field $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ are considered. The state and the control input at time $t \in \mathbb{R}_{\geq 0}$ are denoted by $x(t) \in \mathbb{R}^n$ and $u(t) \in \mathbb{R}^m$, respectively. In addition, the inputs of system (1) are subject to pointwise in time input constraints, i.e. $u(t) \in U \subseteq \mathbb{R}^m$, where U is supposed to be closed. For the sake of completeness, the control functions $u : \mathbb{R}_{\geq 0} \to \mathbb{R}^m$ are assumed to be measurable and locally (Lebesgue-)integrable, i.e. $u \in \mathcal{L}^1_{\mathrm{loc}}([0, \infty), \mathbb{R}^m)$. Moreover, the vector field f is supposed to be continuous and locally Lipschitz with respect to its first argument such that existence and uniqueness of the solution $x(\cdot; x^0, u)$ of (1) for given control function u and initial state x^0 is at least locally ensured. To simplify the notation, the solution is denoted by $x(\cdot)$ if there is no ambiguity.

The control objective is to asymptotically stabilize a (controlled) equilibrium, which without loss of generality is assumed to be the origin, i.e. $f(0, 0) = 0$ and $0 \in U$. We want to fulfill this control task with model predictive control. To this end, the cost functional $J_T : \mathcal{L}^1([0, T), U) \times \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is defined as

$$J_T(u, \hat{x}) := \int_0^T \ell(x(t; \hat{x}, u), u(t))\, \mathrm{d}t \tag{2}$$

based on the positive definite stage cost function $\ell : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}_{\geq 0}$. The corresponding (optimal) value function $V_T : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is given by

$$V_T(\hat{x}) := \inf_{u \in \mathcal{L}^1([0,T),U)} J_T(u, \hat{x}). \tag{3}$$

Note that (2) is only well-defined if the solution trajectory $x(\cdot)$ exists on $[0, T]$. In the following, it is tacitly assumed that, for each state $\hat{x} \in \mathbb{R}^n$, there is at least one admissible control, i.e.

$$\{u \in \mathcal{L}^1([0, T), U) : x(\cdot; \hat{x}, u) \text{ exists on } [0, T]\} \neq \emptyset$$

holds for all $\hat{x} \in \mathbb{R}^n$. Moreover, let us suppose that the infimum of the right hand side in (3) is attained, i.e. existence of an admissible control $u^\star$ such that $V_T(\hat{x}) = J_T(u^\star, \hat{x})$ holds. Note that both the cost functional J_T and the value function V_T depend on the horizon length $T > 0$.

Using the introduced notation, the MPC scheme is as follows.

Algorithm 1 MPC Algorithm

Given: Prediction horizon length $T > 0$ and sampling period $\delta \in (0, T)$.
Set $t = 0$.

(1) Measure the current state $\hat{x} := x(t)$.

(2) Compute a minimizer $u^\star : [0, T) \to \mathbb{R}^m$ of (2).

(3) Implement $u^{MPC}(t + \tau) = u^\star(\tau)$ for $\tau \in [0, \delta)$, set $t = t + \delta$, and goto Step 1.

Algorithm 1 is an MPC scheme without terminal constraints and costs.

Remark 1. The above problem formulation can be extended to also include state constraints, i.e., $x(t) \in X \subseteq \mathbb{R}^n$ is required to hold for all $t \geq 0$. When considering MPC schemes without additional terminal constraints, the presence of such state constraints necessitates some additional assumptions or techniques in order to ensure recursive feasibility of the MPC algorithm, see, e.g., Boccia, Grüne, and Worthmann (2014) or (Grüne & Pannek, 2017, Chapter 7). The following results then also hold in such a setting including state constraints.

Without terminal constraints and costs, the prediction horizon T has to be chosen large enough such that in combination with a certain controllability assumption, asymptotic stability can be concluded. This controllability assumption is typically stated as the existence of a suitable upper bound on the optimal value function, see, e.g. Tuna et al. (2006) and Reble and Allgöwer (2012) for its counterpart in continuous time.

Assumption 2 (*Growth Bound*)**.** Let a continuous, monotonically increasing, and bounded function $B : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ be given such that, for each $x \in \mathbb{R}^n$ and each $t \in \mathbb{R}_{\geq 0}$, the following inequality holds for the optimal value function V_t as defined in (3):

$$V_t(x) \leq B(t) \cdot \inf_{u \in U} \ell(x, u). \tag{4}$$

Typically, $\inf_{u \in U} \ell(x, u) = \ell(x, 0)$ holds, e.g. if the control effort is penalized by an additive term in the stage cost ℓ. Using Assumption 2, one can obtain the following result, see Reble and Allgöwer (2012).

Theorem 3. *Suppose that Assumption 2 is satisfied and that there exists a $\mathcal{K}_\infty$-function*[2] *η such that $\eta(\|x\|) \leq \inf_{u \in U} \ell(x, u)$ holds for all $x \in \mathbb{R}^n$. Then there exists $T > 0$ such that the MPC closed loop resulting from Algorithm 1 is (globally) asymptotically stable.*

In fact, the results in Grüne et al. (2010) (discrete time) and Reble and Allgöwer (2012) (continuous time setting) also provide a technique to estimate the prediction horizon length T such that asymptotic stability holds (see Worthmann, Reble, Grüne, and Allgöwer, 2014 for the connection of both approaches). The basic idea is to interpret the value function V_T as a Lyapunov function. In addition, a performance estimate of the MPC closed loop compared to the infinite horizon optimal solution can be concluded (degree of suboptimality α).

[2] A function $\eta : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is said to be of class $\mathcal{K}_\infty$ if it is continuous, zero at zero, strictly monotonically increasing, and unbounded.