Parametrized Infinite-Horizon Model Predictive Control for Linear Time-Invariant Systems with Input and State Constraints*

Michael Muehlebach and Raffaello D'Andrea[1]

Abstract— **In this work we propose a parametric description of input and state trajectories in order to approximate infinte-horizon optimal control problems encountered in model predictive control. The consequences of applying the parametric description to receding horizon control (model predictive control) are discussed. In particular, recursive feasibility and closed-loop stability are shown. In contrast to the standard, discrete-time finite-horizon model predictive control formulation, the parametric approach provides inherent stability guarantees and has typically fewer optimization variables. A numerical example is used to illustrate the proposed control algorithm.**

I. Introduction

The ability to cope with various operating conditions and to explicitly exploit the full problem structure, including the system dynamics and input/state constraints, makes model predictive control (MPC) a promising and successful control strategy. The underlying principle of MPC is the following: In each sampling interval an optimization problem including a prediction of the state's trajectory is solved, and the first portion of the obtained input is applied to the system. The procedure is then repeated in the next time step. In order to capture the complete evolution of the system, the optimization problem involves ideally an infinite prediction horizon. In practice, this is often intractable and therefore the prediction horizon is typically truncated leading to the standard finite-horizon MPC formulation, [1], where the dynamics are commonly formulated in discrete time. Note that also continuous time finite-horizon MPC formulations have been proposed, see e.g. [2], [3].

Rather than truncating the time horizon, we propose to use a parametric description of input and state trajectories in order to approximate the optimization problem involving an infinite horizon. For polytopic state and input constraints, a quadratic objective function and linear time-invariant dynamics, a convex quadratic program is obtained. We will discuss recursive feasibility and show that the infinite-horizon formulation provides inherent stability guarantees. This contrasts to the finite-horizon formulation, where terminal state constraints and a terminal cost are usually introduced to ensure stability, see e.g. [4]. Moreover, stability is even retained when the optimization routine is terminated prematurely (in fact, after only one iteration provided that the solver is monotonic in the cost).

*This work was supported by ETH-Grant 0-20125-15.

[1]Michael Muehlebach and Raffaello D'Andrea are with the Institute for Dynamic Systems and Control, Swiss Federal Institute of Technology, 8092 Zurich, Switzerland. The contact author is Michael Muehlebach, michaemu@ethz.ch.

In addition, the parametric formulation leads to an optimization problem with typically fewer variables compared to the discrete-time finite-horizon description, making it attractive for online MPC on embedded platforms with limited memory and computational resources. To highlight the numerical efficiency of the proposed parametric formulation simulations of an inverted pendulum system are presented.

A. Related Work

MPC is commonly formulated using a linear time-invariant, discrete-time state-space model and a quadratic cost over a finite horizon, see [1], [5] and references therein. This leads to a quadratic program, which can be solved by standard algorithms and software packages. However, due to the finite horizon formulation, stability and feasibility for all time instants cannot be guaranteed. For the standard MPC formulation, these questions have been extensively studied in the literature, see e.g. [4], [6], [7]. Most stability proofs rely on using the finite-horizon value function as a Lyapunov function. Common techniques to guarantee closed-loop stability are either choosing the prediction horizon large enough, enforcing final state constraints, or introducing a final state cost in the hope to mimic the infinite-horizon problem. In addition, truncating the time horizon leads to a discrepancy between the closed-loop performance objective and the finite-horizon open-loop objective, which is minimized at every time step. An approach to quantify this difference is given in [7], for example.

As opposed to the finite-horizon formulation, stability and feasibility are inherent to the problem formulation, as we will show in the remainder. Moreover, due to the infinite horizon formulation, there is typically a smaller difference between the open-loop objective, which is minimized at every time step, and the resulting closed-loop performance objective.

In [2], [3], the input trajectory is parametrized by a linear combination of basis functions (Laguerre or Kauz). This allows for a more compact problem description, with fewer optimization variables, which potentially provides computational benefits. Still, the finite-horizon formulation is retained.

The use of a parametric description via Laguerre or Kauz basis functions for the identification of linear time-invariant systems has a rich history, see e.g. [8]–[10] and references therein. In system identification, the motivation for using basis functions is to obtain low order models and to incorporate a priori information about the system's time constants.

Likewise, the approach pursued herein includes also a parametric description via linear combinations of basis functions. However, we parametrize input and state trajectories, use a variational formulation of the dynamics (Galerkin approach) to obtain a finite dimensional representation of the state and input trajectories, tackle the infinite horizon problem, and investigate the consequences regarding stability and feasibility of the resulting MPC algorithm. The variational formulation has the advantage that the framework can be extended to nonlinear systems. Moreover, we do not restrict ourselves to a special type of basis functions, e.g. Laguerre or Kauz functions, but provide two conditions which the basis functions have to fulfill in order for the stability and feasibility results to be valid.

B. Outline

Section II introduces the notation and formulates the non-parametric infinite-horizon optimal control problem used as a starting point. The problem is approximated by restricting the input and state trajectories to be linear combinations of basis functions. Using a Galerkin approach, a finite dimensional representation of the dynamics is obtained. The section concludes by discussing convergence of the optimizers of the approximated optimal control problem. In Section III the consequences of using the approximated optimal control-problem for MPC are discussed. In particular, recursive feasibility and stability is shown. Section IV presents simulation results of an inverted pendulum system and compares the traditional discrete-time finite-horizon MPC formulation to the parametric formulation proposed herein. The paper concludes with final remarks in Section V.

II. Problem Formulation

In the following we derive an approximation of the following problem:

$$\begin{aligned} J_\infty := \inf \int_0^\infty & l(x(t),u(t),t)\mathrm{d}t \\ \text{s.t. } & \dot{x}(t) = Ax(t) + Bu(t), \quad x(0) = x_0, \\ & x(t) \in \mathcal{X},\ u(t) \in \mathcal{U}, \quad \forall t \in [0,\infty), \\ & x \in L^2([0,\infty),\mathbb{R}^n), \quad u \in L^2([0,\infty),\mathbb{R}^m), \end{aligned} \tag{1}$$

where $\mathcal{X}$, $\mathcal{U}$ are closed, convex subsets of $\mathbb{R}^n$, respectively $\mathbb{R}^m$, and $l(x,u,t)$ is differentiable in each argument, strongly convex and positive definite in x and u (meaning $l(x,u,t) > 0$ for all $x \neq 0$, $u \neq 0$, for all $t \in [0,\infty)$ and $l(0,0,t) = 0$ for all $t \in [0,\infty)$). The space of square integrable functions mapping from $[0,\infty)$ to $\mathbb{R}^p$ is denoted by $L^2([0,\infty),\mathbb{R}^p)$ and $x_0 \in \mathbb{R}^n$ represents the initial condition. We assume that J_∞ is finite.

Next, the original problem is approximated by introducing a parametrization of input and state trajectories using basis functions, and a variational formulation of the dynamics.

A. Parametrization of Input and State Trajectories

We introduce the parametrized state and input trajectories, denoted by $\tilde{x}(t)$ and $\tilde{u}(t)$ as linear combinations of basis functions $\tau \in L^2([0,\infty),\mathbb{R}^s)$, that is

$$\tilde{x}(t) := (I_n \otimes \tau(t))^\mathsf{T}\eta_x,\ \tilde{u}(t) := (I_m \otimes \tau(t))^\mathsf{T}\eta_u, \tag{2}$$

where $\otimes$ denotes the Kronecker product, $I_p \in \mathbb{R}^{p\times p}$ the identity matrix, and $\eta_x \in \mathbb{R}^{ns}$ and $\eta_u \in \mathbb{R}^{ms}$ the parameter vectors. It is assumed that the basis functions fulfill the following two assumptions:

A1) They are linearly independent.
A2) They fulfill $\dot{\tau}(t) = M_\lambda \tau(t)$ for all $t \in [0,\infty)$, for some matrix $M_\lambda \in \mathbb{R}^{s\times s}$.

Assumption A2) will become important in Section II-B where it is shown that the parametrized state and input trajectories fulfill the equations of motion. Moreover, assumption A2) leads to the following time-shift property, which is used to demonstrate recursive feasibility and stability of the MPC algorithm in Section III-B.

Proposition 2.1: From Assumption A2) it follows that for arbitrary parameters $\eta \in \mathbb{R}^s$ and any time-shift t_s there exists a set of parameters $\hat{\eta} \in \mathbb{R}^s$ such that $\tau(t-t_s)^\mathsf{T}\hat{\eta} = \tau(t)\eta$ for all $t \in [t_s,\infty)$.

Proof: From A2) it follows that

$$\tau(t) = \exp(M_\lambda t)\tau(0), \quad \forall t \in [0,\infty).$$

Note that the matrix exponential is well-defined since $\tau \in L^2([0,\infty),\mathbb{R}^s)$. This implies that

$$\begin{aligned} \tau(t-t_s)^\mathsf{T}\hat{\eta} &= \tau(0)^\mathsf{T}\exp(M_\lambda t)^\mathsf{T}\exp(-M_\lambda t_s)^\mathsf{T}\hat{\eta} \\ &= \tau(t)^\mathsf{T}\exp(-M_\lambda t_s)^\mathsf{T}\hat{\eta} = \tau(t)^\mathsf{T}\eta. \end{aligned}$$

■

B. Variational Formulation of the Dynamics

Due to the fundamental Lemma of the calculus of variations, [11], the equivalence between

$$\dot{x}(t) = Ax(t) + Bu(t),\ \forall t \in [0,\infty) \tag{3}$$

and

$$\int_0^\infty \delta p(t)^\mathsf{T}(Ax(t) + Bu(t) - \dot{x}(t))\mathrm{d}t = 0, \tag{4}$$

for all variations $\delta p \in L^2([0,\infty),\mathbb{R}^n)$, holds (in the almost everywhere sense). By restricting the variations δp to be linear combinations of basis functions[1], i.e. $\delta\tilde{p}(t) = (I_n \otimes \tau(t))^\mathsf{T}\delta\eta_p$, with $\delta\eta_p \in \mathbb{R}^{ns}$, and inserting the parametrized input and state trajectories, (4) is simplified to

$$\delta\eta_p^\mathsf{T}\int_0^\infty (I_n \otimes \tau(t))(A\tilde{x}(t) + B\tilde{u}(t) - \dot{\tilde{x}}(t))\mathrm{d}t = 0, \tag{5}$$

for all $\delta\eta_p \in \mathbb{R}^{ns}$. Combined with (2), a linear relationship between the parameters of the state and input trajectories is found,

$$(A \otimes U_\tau - I_n \otimes U_\tau M_\lambda^\mathsf{T})\eta_x + (B \otimes U_\tau)\eta_u = 0, \tag{6}$$

where $U_\tau := \int_0^\infty \tau(t)\tau(t)^\mathsf{T}\mathrm{d}t$.

[1]We use the same basis functions to parametrize state and input trajectories, as well as variations $\delta\tilde{p}$.

We show next that the parameters η_x and η_u satisfying (6) fulfill the equations of motion.

Proposition 2.2: The following are equivalent:

1) $\dot{\tilde{x}}(t) = A\tilde{x}(t) + B\tilde{u}(t)$, for all $t \in [0, \infty)$.
2) $(A \otimes U_\tau - I_n \otimes U_\tau M_\lambda^\mathsf{T})\eta_x + (B \otimes U_\tau)\eta_u = 0$.

Proof: Without loss of generality we assume that the basis functions τ are orthonormal. Note that orthonormal basis functions can be constructed using the Gram-Schmidt process, [12, p.50]. This implies that U_τ reduces to the identity and therefore the condition 2) reads as

$$(A \otimes I_s - I_n \otimes M_\lambda^\mathsf{T})\eta_x + (B \otimes I_s)\eta_u = 0. \tag{7}$$

From $\dot{\tilde{x}}(t) = A\tilde{x}(t) + B\tilde{u}(t)$ for all $t \in [0, \infty)$ it can be inferred that $\tilde{x}^{(n+1)}(t) = A\tilde{x}^{(n)}(t) + B\tilde{u}^{(n)}(t)$ is fulfilled, where $\tilde{x}^{(n)}$ denotes the n'th derivative with respect to time. The dynamics $\dot{\tilde{x}}(t) = A\tilde{x}(t) + B\tilde{u}(t)$ are therefore equivalent to

$$\begin{aligned} I_n \otimes &(\tau(t), \dot{\tau}(t), \dots, \tau^{(s-1)}(t))^\mathsf{T} \\ &\left[(A \otimes I_s - I_n \otimes M_\lambda^\mathsf{T})\eta_x + (B \otimes I_s)\eta_u\right] = 0, \end{aligned} \tag{8}$$

which is obtained by inserting $\tilde{x}(t) = (I_n \otimes \tau(t))^\mathsf{T}\eta_x$ and $\tilde{u}(t) = (I_m \otimes \tau(t))^\mathsf{T}\eta_u$ and separating time-dependent and time-independent terms. Due to the linear independence of the basis functions τ, the matrix $I_n \otimes (\tau(t), \dot{\tau}(t), \dots, \tau^{(s-1)}(t))^\mathsf{T}$ is full rank for all $t \in [0, \infty)$, [13, p.1079]. Thus, (8) and (7) are equivalent, which concludes the proof. ■

C. Approximation of (1)

In the previous section it was shown that, given the parametrization of input and state trajectories as linear combination of basis functions, the equations of motion can be simplified to the linear relationship (6). This motivates the approximation of the original optimal control problem (1) by

$$\begin{aligned} J_s := \min \int_0^\infty & l((I_n \otimes \tau(t))^\mathsf{T}\eta_x, (I_m \otimes \tau(t))^\mathsf{T}\eta_u, t)\mathrm{d}t \\ \text{s.t. } & (A \otimes U_\tau - I_n \otimes U_\tau M_\lambda^\mathsf{T})\eta_x + (B \otimes U_\tau)\eta_u = 0, \\ & (I_n \otimes \tau(0))^\mathsf{T}\eta_x = x_0, \quad (I_n \otimes \tau(t))^\mathsf{T}\eta_x \in \mathcal{X}, \\ & (I_m \otimes \tau(t))^\mathsf{T}\eta_u \in \mathcal{U}, \ \forall t \in [0, \infty). \end{aligned} \tag{9}$$

Note that the subscript s refers to the number of basis functions used, which is linked to the dimensions of the parameter vectors η_x and η_u. In general the integral of the running cost cannot be solved analytically unless the running cost $l(x, u, t)$ has a special form (e.g. quadratic in x and u).

If the optimization problem (9) is feasible for at least s_0 basis functions, then the sequence of optimal values J_s, J_{s+1}, ..., with $s \geq s_0$, is decreasing. This is because the optimizer with cost J_s is a solution candidate for the optimization problem over $s+1$ basis functions (the optimization over $s+1$ basis functions has the optimal cost J_{s+1}). Note that feasible candidates of (9) need to satisfy the equations of motion, as well as input and state constraints; therefore extending the vectors η_x and η_u obtained from the optimization over s basis functions with zeros allows to construct a feasible solution for the problem (9) with $s + 1$ basis functions. Likewise, the minimizer of (9) is also a candidate for the optimization problem (1), which implies that $J_\infty \leq J_s$ for all integers $s > 0$.

The fact that the sequence of optimal values is decreasing implies the existence of a weakly convergent subsequence of the state and input trajectories, $\tilde{x}(t)$ and $\tilde{u}(t)$. This is because from the positive definiteness and strong convexity of $l(x, u, t)$ it follows that

$$l(x, u, t) \geq \alpha_x x^\mathsf{T} x + \alpha_u u^\mathsf{T} u, \quad \forall t \in [0, \infty), \tag{10}$$

with $\alpha_x > 0$, $\alpha_u > 0$, which implies that

$$J_{s_0} \geq \alpha_x ||\tilde{x}||_2^2 + \alpha_u ||\tilde{u}||_2^2 \tag{11}$$

for any $s \geq s_0$, where $||\cdot||_2$ denotes the L^2-norm. This implies that $||\tilde{x}||_2$ and $||\tilde{u}||_2$ are bounded sequences (in s). Combined with the fact that $L^2([0, \infty), \mathbb{R}^n) \times L^2([0, \infty), \mathbb{R}^m)$ is a Hilbert space it implies the existence of a weakly convergent subsequence of $(\tilde{x}, \tilde{u})$, [14, p.163].

The results are summarized by the following proposition.

Proposition 2.3: Let (9) be feasible for s_0 basis functions, let $\tilde{x}^s$ and $\tilde{u}^s$ denote the optimizers of (9), and let J_s denote the corresponding optimal cost. If the basis functions fulfill assumptions A1) and A2) it holds that

1) $J_s \geq J_{s+1} \geq J_{s+2} \geq \cdots \geq J_\infty$, for all $s \geq s_0$.
2) There exists a subsequence $s(k)$, $k = 1, 2, \dots$, such that $\tilde{x}^{s(k)}$ and $\tilde{u}^{s(k)}$ converge weakly.

III. Application to MPC

In the following section we apply the finite dimensional approximation (9) to MPC. In particular, recursive feasibility, closed-loop stability and the implementation of the constraints $x(t) \in \mathcal{X}$ and $u(t) \in \mathcal{U}$ will be discussed.

A. Notation and Definition of the MPC Algorithm

The following notation is used: The closed-loop state and input trajectories are denoted by $x(t)$ and $u(t)$. We refer to the predicted trajectories as $\tilde{x}(t_p|t)$, $\tilde{u}(t_p|t)$, where $t_p \geq 0$ denotes the prediction horizon. For $t_p = 0$ the prediction matches the exact trajectory, that is $\tilde{x}(0|t) = x(t)$, $\tilde{u}(0|t) = u(t)$ for all $t \in [0, \infty)$. The predictions $\tilde{x}(t_p|t)$ and $\tilde{u}(t_p|t)$ are obtained by solving (9) subject to the initial condition $x_0 = x(t)$, which yields the parameters η_x and η_u defining $\tilde{x}(t_p|t)$, $\tilde{u}(t_p|t)$ by $\tilde{x}(t_p|t) = (I_n \otimes \tau(t_p))^\mathsf{T}\eta_x$ and $\tilde{u}(t_p|t) = (I_m \otimes \tau(t_p))^\mathsf{T}\eta_u$. In order to highlight the dependence of the optimal cost in (9) on the initial condition, the optimal cost is denoted by $J_s = J_s(x(t))$.

The MPC algorithm consists of the following steps: The optimization problem (9) is solved at the sampling intervals $t = kT_s$ (T_s is the sampling period), with respect to the current state as initial condition, i.e. $x_0 = x(t)$. This yields the optimal cost $J_s(x(t))$, together with the state and input predictions $\tilde{x}(t_p|t)$ and $\tilde{u}(t_p|t)$. In between the time intervals kT_s, that is $t \in [kT_s, (k+1)T_s)$ the predicted input $u(t) = \tilde{u}(t - kT_s|t)$ is applied to the system, until the optimization is repeated at time $t = (k+1)T_s$.

B. Recursive Feasibility and Stability

The time-shift property given by Proposition 2.1 is of paramount importance for both, stability and recursive feasibility. In fact, it implies the existence of a feasible candidate of (9) at time $(k+1)T_s$ given a feasible solution of (9) at time $t = kT_s$. This is because the predictions $\tilde{x}(t_p|t)$ and $\tilde{u}(t_p|t)$, that is the optimizer of (9) at time $t = kT_s$, fulfill the equations of motion and the initial condition $\tilde{x}(0|t) = x(t)$. They are therefore exact, i.e. $x(t+t_p) = \tilde{x}(t_p|t)$, $u(t+t_p) = \tilde{u}(t_p|t)$ for all $t_p \in [0, T_s)$. Hence, by the time-shift property, the feasible candidates

$$\tilde{x}(t_p|(k+1)T_s) = (I_n \otimes \tau(t_p)^{\mathsf{T}})(I_n \otimes \exp(M_\lambda T_s)^{\mathsf{T}})\eta_x$$
$$\tilde{u}(t_p|(k+1)T_s) = (I_m \otimes \tau(t_p)^{\mathsf{T}})(I_m \otimes \exp(M_\lambda T_s)^{\mathsf{T}})\eta_u$$

for the optimization at time $t = (k+1)T_s$ are obtained from $\tilde{x}(t_p|t) = (I_n \otimes \tau(t_p))^{\mathsf{T}}\eta_x$ and $\tilde{u}(t_p|t) = (I_m \otimes \tau(t_p))^{\mathsf{T}}\eta_u$. This proves readily:

Proposition 3.1: Provided that the optimization problem (9) is feasible at time $t = 0$, it remains feasible for all times $t > 0$.

Moreover, choosing the optimal cost of (9) as Lyapunov function (as frequently done in the discrete-time finite-horizon setting, see e.g. [4]) can be used to demonstrate closed-loop stability.

Proposition 3.2: The resulting closed-loop system is asymptotically stable, provided that the optimization problem (9) is feasible at time $t = 0$.

Proof: By Proposition 3.1 the problem (9) remains feasible for all times $t > 0$. We will show that the cost $J_s(x)$ (given by (9)) acts as a Lyapunov function. Note that $J_s(x) > 0$, for all $x \neq 0$ and $J_s(x) = 0$ if and only if $x = 0$, since the evolution of the predicted state $\tilde{x}$ is continuous and $l(x, u, t)$ is positive definite. Without loss of generality we fix $k > 0$ and let the solution of (9), with $x(kT_s)$ as initial condition, be denoted by $\tilde{x}(t_p|t)$, $\tilde{u}(t_p|t)$. The corresponding minimum cost is given by $J_s(x(kT_s))$. Since the prediction $\tilde{x}$ is exact, the state evolves according to $x(t) = \tilde{x}(t - kT_s|kT_s)$, $\forall t \in [kT_s, (k+1)T_s)$. At time $(k+1)T_s$ the problem (9) is solved again, this time with initial condition $x((k+1)T_s)$. The time-shift property implies that the predicted trajectories $\tilde{x}(t_p|t)$, $\tilde{u}(t_p|t)$ are feasible candidates, which yields the upper bound

$$J_s(x((k+1)T_s)) \leq \int_{T_s}^{\infty} l(\tilde{x}(t_p|kT_s), \tilde{u}(t_p|kT_s), t_p)\mathrm{d}t_p.$$

Note that the right-hand side can be rewritten as

$$J_s(x(kT_s)) - \int_0^{T_s} l(\tilde{x}(t_p|kT_s), \tilde{u}(t_p|kT_s), t_p)\mathrm{d}t_p,$$

resulting in

$$J_s(x((k+1)T_s)) - J_s(x(kT_s)) \leq - \int_0^{T_s} l(\tilde{x}(t_p|kT_s), \tilde{u}(t_p|kT_s), t_p)\mathrm{d}t_p.$$

Positive definiteness of the running cost $l(x, u, t)$ implies that the right-hand side is always less than zero (even strictly less than zero except for $x(kT_s) = 0$), which concludes the proof. ■

Note that the cost $J_s(x)$ even acts as a Lyapunov function if the optimization (9) is not solved completely. It is merely required that the cost $J_s(x((k+1)T_s))$ is smaller than $J_s(x(kT_s))$. If the optimization at time $t = (k+1)T_s$ is initialized with the shifted version of the previous solution (obtained at time $t = kT_s$), such a decrease is obtained after only one iteration of the optimization routine, provided that the solver is monotone in the costs. Thus, even if the optimization is stopped prematurely, stability can still be guaranteed.

C. Implementation of Constraints

A certain difficulty associated with solving the optimization (9) is the implementation of the input and state constraints, i.e. ensuring that input and state trajectories fulfill

$$\tilde{x}(t) \in \mathcal{X}, \quad \tilde{u}(t) \in \mathcal{U}, \quad t \in [0, \infty). \tag{12}$$

It turns out that the requirement (12) induces closed and convex finite dimensional sets in the parameter space, which can be approximated using polyhedra. These approximations, denoted by $\tilde{\mathcal{X}}$ and $\tilde{\mathcal{U}}$, could be chosen such that $\eta_x \in \tilde{\mathcal{X}}$ and $\eta_u \in \tilde{\mathcal{U}}$ implies (12), and the previous stability and feasibility results would remain valid. However, in order to simplify the presentation, we use constraint sampling in the following, as proposed in [2]. We therefore enforce the constraints only at certain time instants t_{ci} (finitely many) in the hope that the constraints are also satisfied for all other time instants. Alternatively one could check constraint violations a posteriori, refine the constraint sampling if necessary and solve (9) again to ensure that the constraints are not violated. If the constraints are violated in between these time instants, the theoretic guarantees are no longer valid in general. Numerical experiments indicate that the method performs well, provided that the time intervals $[t_{ci}, t_{c(i+1)})$ are chosen to be sufficiently small, see Section IV.

IV. Simulation Example

In the following, the effectiveness of the proposed parametric MPC approach is illustrated with simulations of an inverted pendulum system as depicted in Fig. 1.

We compare the standard (discrete-time, finite-horizon) MPC strategy to the proposed parametric MPC approach. The dynamics are linearized about the upright equilibrium given by $\varphi = \dot{\varphi} = r = \dot{r} = u = 0$, and discretized with a sampling time of $T_s = 20\,\mathrm{ms}$ in the standard MPC formulation. Note that the unstable pole lies at $4.43\,\mathrm{rad/s}$ for the given parameters (as listed in the caption of Fig. 1), which makes the sampling time of $T_s = 20\,\mathrm{ms}$ a sensible choice. The goal is to drive the cart form $r_0 = 0.5\,\mathrm{m}$ back to $r = 0$, while balancing the pendulum in the upright position. The running cost $l(x, u, t)$ is chosen to be quadratic, more precisely

$$l(x, u, t) := \frac{1}{2}x^{\mathsf{T}}x + \frac{0.05}{2}u^2, \tag{13}$$

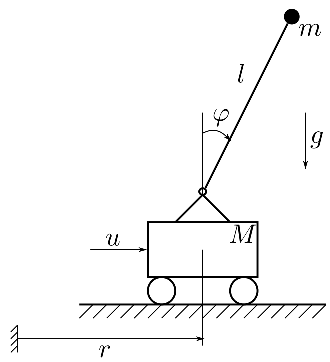

Fig. 1. Inverted pendulum on a cart. The cart mass M was chosen to be $1\,\mathrm{kg}$, the pendulum length l, $1\,\mathrm{m}$, the pendulum mass m, $1\,\mathrm{kg}$. The cart position is parametrized by r, the inclination angle by φ, whereas g denotes the gravitational acceleration. The control input is given by the force u applied to the cart. It is limited to $\pm 0.5\,\mathrm{N}$.

where the state vector x is defined as $x := (r, \dot{r}, \varphi, \dot{\varphi})^{\mathsf{T}}$. In the standard discrete-time MPC formulation the running cost is weighted by the sampling time and a terminal cost matching the infinite-horizon LQR cost-to-go is added. For simplicity, state constraints are not considered, while the control input u is restricted to lie within $[-0.5\,\mathrm{N}, 0.5\,\mathrm{N}]$. The basis functions τ are chosen to be orthonormal and spanned by

$$\tau(t) \in \exp(-\lambda t)\ \mathrm{span}\left(1, t, t^2, \ldots, t^{s-1}\right), \tag{14}$$

where the time constant of the exponential decay is chosen to be $\lambda = 3$ 1/s (corresponding approximately to the closed-loop poles of an LQR design). This leads naturally to so called Laguerre functions, [2], that is

$$\tau(t) = (l_1(t), l_2(t), \ldots, l_s(t))^{\mathsf{T}}, \tag{15}$$

where

$$l_i(t) := \sqrt{2\lambda} \exp(-\lambda t) \sum_{k=0}^{i-1} \binom{i-1}{k} \frac{(-1)^k}{k!} (2\lambda t)^k. \tag{16}$$

The constraints are sampled at the time instants t_{ci} which are determined by solving

$$t_{c1} = 0, \quad \tau(t_{c1})^{\mathsf{T}} \tau(t_{ci}) = 0, \quad i = 2, \ldots, s. \tag{17}$$

It turns out that these sampling instants fulfill

$$\tau(t_{ci})^{\mathsf{T}} \tau(t_{cj}) = 0, \quad \forall i \neq j, \tag{18}$$

which can be used to simplify the numerical optimization routines.[2]

For $s = 10$, the closed loop trajectories are depicted in Fig. 2 and indicate a reasonable control performance. For the given initial condition, starting at rest at $r_0 = 0.5\,\mathrm{m}$, the problem becomes feasible for $s = 7$.

Fig. 3 compares the achieved closed-loop cost and the required average execution time for solving the optimization problem (9) of the parametric MPC approach with the discrete-time finite-horizon formulation. The closed-loop trajectory is initialized with $r_0 = 0.5\,\mathrm{m}$, $\varphi_0 = \dot{\varphi}_0 = \dot{r}_0 = 0$. The optimization routines are run on a Laptop with an IntelCore i7-4710MQ (2.50Ghz) CPU and 8.0GB RAM using Matlab, and to obtain a better picture, 1st and 2nd order solution methods are evaluated. The resulting quadratic program is solved with the generalized fast dual gradient method, [15], which represents a state of the art 1st order optimization routine, and Mosek[3] (2nd order). We use Mosek despite the fact that 2nd order optimization routines tailored to the standard MPC problem formulation, exploiting sparsity and achieving significant speedup are available, [16], [17]. This is because via coordinate transformation resulting in a set of slightly modified basis functions (still spanned by (14)), a sparse problem description can be obtained with the parametric MPC formulation too. This allows for similar strategies to exploit sparsity as in the standard MPC formulation, which could potentially lead to similar speedups. Adopting a 2nd order optimization routine to the specific problem structure given by the parametric MPC approach is beyond the scope of this paper. Therefore the 2nd order timing results should be seen as indication of meaningful trends, rather than absolute timings.

The resulting optimal cost is depicted in Fig. 4. Note that in the standard finite-horizon formulation, the optimal cost J_∞ is significantly underestimated for short prediction horizons (i.e. prediction horizons of less than $4\,\mathrm{s}$ resulting in less than 1004 optimization variables), see Fig. 5.[4] Instead, the parametric MPC formulation overestimates the cost J_∞ when using only a reduced number of basis functions ($s = 8, 9$). However, to obtain the same accuracy in terms of optimum cost, significant higher computational costs are necessary in case of the standard finite-horizon formulation.

V. Conclusion

In this article the state and input trajectories of a linear time-invariant system were represented as linear combinations of basis functions. This allowed to simplify the infinite-horizon optimal control problems encountered in MPC. The consequences of using such an approximation regarding recursive feasibility and closed-loop stability were discussed. A numerical simulation example shows the efficiency of the proposed approach by providing a sensible approximation of the underlying infinite-horizon optimal control problem at relatively low computational expenses.

References

[1] M. Morari and J. H. Lee, "Model predictive control: past, present and future," *Computers & Chemical Engineering*, vol. 23, no. 4, pp. 667–682, 1999.

[2] L. Wang, "Continuous time model predictive control design using orthonormal functions," *International Journal of Control*, vol. 74, no. 16, pp. 1588–1600, 2001.

[3] ——, *Model predictive control system design and implementation using MATLAB®*. Springer Science & Business Media, 2009.

[4] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[2] A detailed discussion is beyond the scope of this paper.

[3] `https://www.mosek.com`.

[4] Even though a terminal cost matching the infinite horizon LQR cost was included.

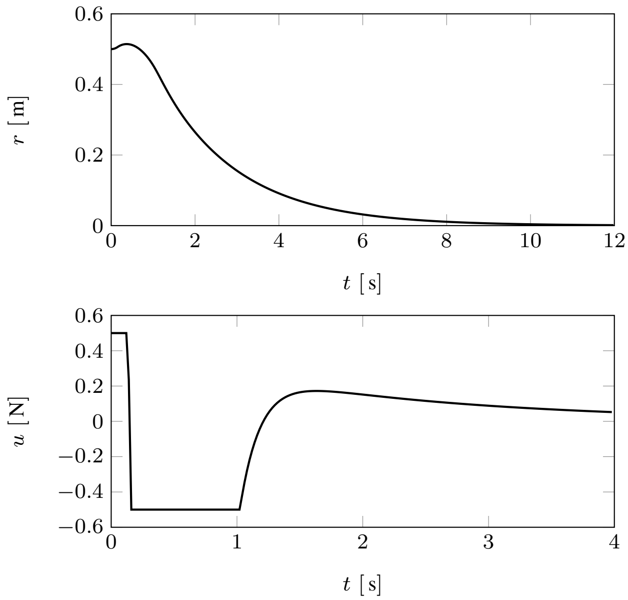

Fig. 2. Resulting closed-loop trajectories, position $r(t)$ (top) and input $u(t)$ (bottom). Note that the control input saturates for a significant amount of time.

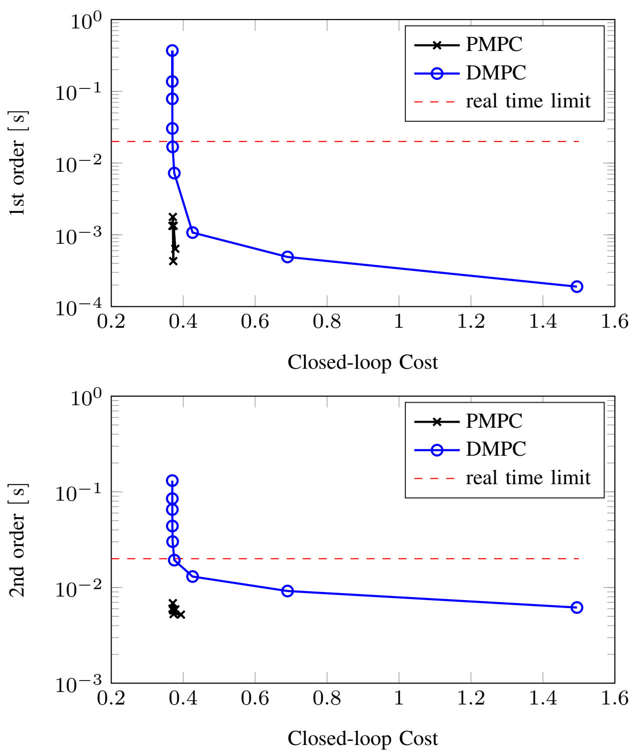

Fig. 3. Execution times versus achieved closed-loop cost using a first order method (top) and a second order method (bottom). The proposed parametric MPC approach (black, crosses) is compared to the standard discrete-time finite-horizon MPC formulation (blue, circles). The real-time boundary corresponding to an execution time of 20 ms is indicated by a dashed line.

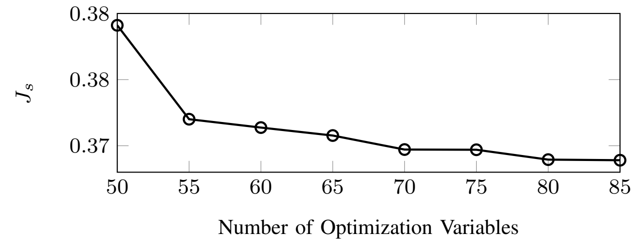

Fig. 4. Computed cost J_s with the parametric MPC approach over the number of optimization variables (which corresponds to the number of basis functions). The cost is monotonically decreasing with an increasing number of basis functions.

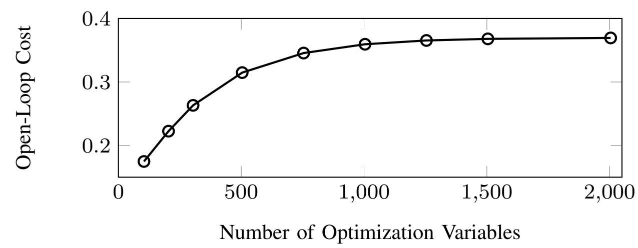

Fig. 5. Computed cost of the standard discrete-time finite-horizon MPC approach over the number of optimization variables (which is proportional to the prediction horizon). Note that a terminal cost matching the LQR infinite-horizon cost was added.

[5] M. Morari, C. Garcia, J. Lee, and D. Prett, *Model predictive control*. Prentice Hall, 1993.

[6] C. Chen and L. Shaw, "On receding horizon feedback control," *Automatica*, vol. 18, no. 3, pp. 349–352, 1982.

[7] J. Primbs and V. Nevistic, "Constrained finite receding horizon linear quadratic control," in *Proceedings of the IEEE Conference on Decision and Control*, vol. 4, 1997, pp. 3196–3201.

[8] B. Wahlberg, "System identification using Laguerre models," *IEEE Transactions on Automatic Control*, vol. 36, no. 5, pp. 551–562, 1991.

[9] ——, "System identification using Kautz models," *IEEE Transactions on Automatic Control*, vol. 39, no. 6, pp. 1276–1282, 1994.

[10] B. Wahlberg and P. Mäkilä, "On approximation of stable linear dynamical systems using Laguerre and Kautz functions," *Automatica*, vol. 32, no. 5, pp. 693–708, 1996.

[11] I. M. Gelfand and S. V. Fomin, *Calculus of variations*. Courier Corporation, 2000.

[12] R. Courant and D. Hilbert, *Methods of Mathematical Physics*. Interscience Publishers, 1953, vol. 1.

[13] A. Jeffrey and D. Zwillinger, *Table of integrals, series, and products*. Academic Press, 2007.

[14] J. B. Conway, *A Course in Functional Analysis*, 2nd ed. Springer, 1990.

[15] P. Giselsson, "Improved fast dual gradient method for embedded model predictive control," *Proceedings of the IFAC World Congress*, pp. 2303–2309, 2014.

[16] A. Domahidi, A. U. Zgraggen, M. N. Zeilinger, M. Morari, and C. Jones, "Efficient interior point methods for multistage problems arising in receding horizon control," in *Proceedings of the IEEE Conference on Decision and Control*, 2012.

[17] Y. Wang and S. Boyd, "Fast model predictive control using online optimization," *IEEE Transactions on Control Systems Technology*, vol. 18, no. 2, pp. 267–278, 2010.