Worst case vs. average performance estimates for unconstrained NMPC schemes

Lars Grüne*

Lehrstuhl für Angewandte Mathematik, Mathematisches Institut, Universität Bayreuth, 95440 Bayreuth, Germany

We consider performance estimates for unconstrained NMPC schemes under a controllability assumption. By means of worst case and average estimates we investigate the effect of varying control horizons on the performance and compare the results with numerical simulations.

Copyright line will be provided by the publisher

1 Introduction

In this paper we analyze nonlinear model predictive control (NMPC) schemes without stabilizing terminal constraints. More precisely, we investigate how the performance of the scheme depends on the length of the control horizon: in NMPC, at each time instant an optimal control problem on a finite horizon of length N is solved. Then the first piece of the resulting optimal control sequence is applied to the system and the procedure is repeated iteratively. The *control horizon* $m \in \mathbb{N}$ now denotes the number of elements of the sequence which is applied to the system before a re-optimization is performed.

Formally, we consider nonlinear discrete time control systems

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

with $x(n) \in X$ and $u(n) \in U$ for $n \in \mathbb{N}_0$, where X and U are arbitrary metric spaces. We denote the space of control sequences $u : \mathbb{N}_0 \to U$ by $\mathcal{U}$ and the solution trajectory for given $u \in \mathcal{U}$ by $x_u(n)$.

At each optimization time we measure the current state x_0 and solve the finite horizon optimal control problem with horizon length $N \in \mathbb{N}$

$$\text{minimize} \quad J_N(x_0, u) = \sum_{n=0}^{N-1} l(x(n), u(n)), \qquad \text{where } x(\cdot) \text{ solves (1).} \tag{2}$$

From the resulting optimal control sequence $u^*(0), \ldots, u^*(N-1)$ we apply the first $m \in \{1, \ldots, N-1\}$ elements and then continue iteratively solving (2) again. Formally, this amounts to defining $\mu(x_0, n) := u^*(n)$, $n = 0, \ldots, m-1$ and applying μ to (1) according to the rule $x_\mu(0) = x_0$, $x_\mu(n+1) = f(x_\mu(n), \mu(x_\mu(km), n-km))$, where $k \in \mathbb{N}$ is maximal with $km \le n$.

Note that the concept can be generalized to time varying m which is particularly useful in the case of networked control systems, cf. [1]. In order to assess the performance of the resulting trajectory x_μ, we introduce the infinite horizon performance criterion $J_\infty^\mu(x) = \sum_{n=0}^\infty l(x_\mu(n), \mu(x_\mu(km), n-km))$, where again $k \in \mathbb{N}$ is maximal with $km \le n$.

Furthermore, we consider the optimal value functions $V_N(x_0) = \inf_{u \in \mathcal{U}} J_N(x_0, u)$ and $V_\infty(x_0) = \inf_{u \in \mathcal{U}} J_\infty(x_0, u)$, where $J_\infty(x_0, u)$ is defined analogously to $J_N(x_0, u)$ with the infinite optimization horizon $N = \infty$. From these definitions it is immediately clear that $J_\infty^\mu(x) \ge V_\infty(x)$ holds. The question we want to investigate now is: how much larger is $J_\infty^\mu(x)$?

2 Worst case performance

In [2] an analytic answer to this question is given under the following controllability assumption.

Assumption 2.1 Given constants $C > 0$, $\sigma \in (0, 1)$, for each $x_0 \in X$ there exists a control function $u_{x_0} \in \mathcal{U}$ satisfying $l(x_{u_{x_0}}(n), u_{x_0}(n)) \le C\sigma^n \inf_{u \in U} l(x_0, u)$ for all $n \in \mathbb{N}_0$.

Under this assumption, Theorem 5.4 in [2] ensures the inequality $J_\infty^\mu(x) \le V_\infty(x)/\alpha$ for

$$\alpha = 1 - \frac{\prod_{i=m+1}^{N}(\gamma_i - 1) \prod_{i=N-m+1}^{N}(\gamma_i - 1)}{\left(\prod_{i=m+1}^{N}\gamma_i - \prod_{i=m+1}^{N}(\gamma_i - 1)\right)\left(\prod_{i=N-m+1}^{N}\gamma_i - \prod_{i=N-m+1}^{N}(\gamma_i - 1)\right)}, \text{ where } \gamma_i = C\frac{1-\sigma^i}{1-\sigma}. \tag{3}$$

When plotting α depending on the control horizon m, one observes an interesting symmetry shown in Figure 1 (left and middle) which shows that the performance estimate for optimization horizon m is identical to $N - m$. In particular, for large control horizons we obtain the same performance estimate as for short horizons.

* e-mail: lars.gruene@uni-bayreuth.de, Phone: +49 921 55 3270, Fax: +49 921 55 5361

Copyright line will be provided by the publisher

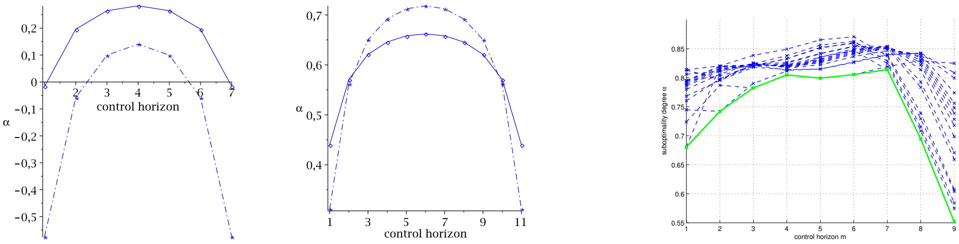

Fig. 1 Left and middle: α from (3) depending on m for $C = \frac{3}{2}$, $\sigma = \frac{4}{5}$ (solid line, o) and $C = \frac{7}{2}$ and $\sigma = \frac{3}{5}$ (dashed line, *) and optimization horizons $N = 8$ (left) and $N = 12$ (middle). Right: α measured along simulated trajectories, for details see [1, Section IV]

In contrast to this, Figure 1(right) shows α-values computed along simulated trajectories of a stabilized inverted pendulum using the algorithm from [3]. Clearly, the performance for large m is significantly worse than the performance for small m. How can this difference between theory and numerical experiment be explained?

3 Average performance

Formula (3) is the outcome of an optimization problem which minimizes α over all NMPC problems satisfying Assumption 2.1. More precisely, the optimization problem behind (3) minimizes $\alpha = \sum_{n=0}^{N-1} \lambda_n - \nu$ over λ_i, $\nu \geq 0$ subject to the constraints

$$(a)\ \sum_{n=k}^{N-1} \lambda_n \leq \gamma_{N-k}\lambda_k, \qquad (b)\ \sum_{n=0}^{m-1} \lambda_n = 1, \qquad (c)\ \nu \leq \sum_{n=0}^{j-1} \lambda_{n+m} + \gamma_{N-j}\lambda_{j+m}, \tag{4}$$

cf. [2]. Hence, a possible explanation for the difference is that a "randomly" chosen system is more likely to be close to the worst case system for large m than for small.

In order to validate this conjecture, instead of computing α by minimization, we compute α by generating random admissible sequences $\lambda_0, \ldots, \lambda_{N-1}$ satisfying (4)(a,b), minimize α over ν satisfying (4)(c) and average over the resulting α in a Monte-Carlo simulation. We generate these random sequences by backward induction: we set

$$\lambda_{N-1} := 1, \qquad \lambda_k := \frac{\sum_{n=k}^{N-1} \lambda_n}{(\gamma_{N-k} - 1)z}, \quad k = N-2, \ldots, 0$$

and normalize the resulting sequence such that (4)(b) holds. Here $z \in (0, 1)$ are uniformly distributed random numbers generated in each iteration step; experiments with other distributions yielded qualitatively similar results.

Figure 2 shows the result of this Monte-Carlo simulation with 1000 samples for $C = 3$, $\sigma = 0.6$ and $N = 10$.

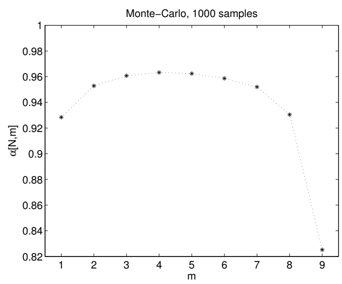

Fig. 2 Monte-Carlo simulation for α with 1000 randomly generated λ_i sequences.

We observe a similar qualitative behavior as for the simulations in Figure 1(right). In particular, this supports our conjecture that the average performance for large m is closer to the worst case than for small m. Still, the result also shows that for control horizons m up to about $0.8N$ we do not observe a significant loss of performance compared to the classical MPC case $m = 1$.

References

[1] L. Grüne, J. Pannek, and K. Worthmann, A networked unconstrained nonlinear MPC scheme, in: Proceedings of the European Control Conference 2009, (Budapest, Hungary, 2009), pp. 371–376.
[2] L. Grüne, J. Pannek, M. Seehafer, and K. Worthmann, Analysis of unconstrained nonlinear MPC schemes with varying control horizon, Preprint, Universität Bayreuth, 2009, Submitted.
[3] L. Grüne and J. Pannek, Syst. Contr. Lett. **58**, 161–168 (2009).