
Performance of NMPC schemes without stabilizing terminal constraints

Nils Altmüller, Lars Grüne, and Karl Worthmann

Mathematical Institute, University of Bayreuth, 95440 Bayreuth, Germany
nils.altmueller, lars.gruene, karl.worthmann@uni-bayreuth.de

Summary. In this paper we investigate the performance of unconstrained nonlinear model predictive control (NMPC) schemes, i.e., schemes in which no additional terminal constraints or terminal costs are added to the finite horizon problem in order to enforce stability properties. The contribution of this paper is twofold: on the one hand in Section 3 we give a concise summary of recent results from [7, 3, 4] in a simplified setting. On the other hand, in Section 4 we present a numerical case study for a control system governed by a semilinear parabolic PDE which illustrates how our theoretical results can be used in order to explain the differences in the performance of NMPC schemes for distributed and boundary control.

1 Introduction

Model predictive control (MPC) is a well established method for approximating the optimal control of linear and nonlinear systems [1, 8, 9]. MPC approximates the optimal solutions of in general computationally intractable infinite horizon optimal control problems by the iterative solution of finite horizon problems, the so called receding horizon strategy. This interpretation of MPC immediately leads to the question of how good the performance of the MPC scheme is compared to the original infinite horizon optimization criterion. Since infinite horizon problems are often formulated in order to obtain stabilizing feedback laws, another important question is whether the resulting MPC feedback law will still stabilize the system.

In this paper we investigate these issues for so called unconstrained nonlinear MPC (NMPC) schemes. Here *unconstrained* refers to those terminal constraints or terminal costs which are added to the finite horizon problem in order to enforce stability properties; other constraints like, e.g., state and control constraints motivated by physical considerations can easily be included in our analysis although for simplicity of exposition we do not elaborate on this aspect in this paper and refer to, e.g., [9] for an extensive treatment of feasibility issues. Such unconstrained schemes are appealing in many ways, cf. the discussion at the end of the introductory Section 2.

The contribution of this paper is twofold: on the one hand in Section 3 we give a concise summary of recent results from [3, 4, 7] in a simplified setting, restricting the reasoning to the special case of exponential controllability and classical NMPC feedback laws. For an extended setting including networked control systems, finite time controllability and additional weights in the cost functional we refer to [3, 4] and [5]. On the other hand, in Section 4 we present a numerical case study for a control system governed by a semilinear parabolic PDE. This case study illustrates how our theoretical results can be used in order to explain the differences in the performance of NMPC schemes for distributed and boundary control.

2 Setup and Preliminaries

We consider a nonlinear discrete time control system given by

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

with $x(n) \in X$ and $u(n) \in U$ for $n \in \mathbb{N}_0$. Here the state space X and the control value space U are arbitrary metric spaces with metrics denoted by $d(\cdot, \cdot)$. We denote the space of control sequences $u : \mathbb{N}_0 \rightarrow U$ by \mathcal{U} and the solution trajectory for given $u \in \mathcal{U}$ by $x_u(\cdot)$. State and control constraints can be incorporated by replacing X and U by appropriate subsets of the respective spaces, however, for brevity of exposition we will not address feasibility issues in this paper.

A typical class of such discrete time systems are sampled–data systems induced by a controlled — finite or infinite dimensional — differential equation with sampling period $T > 0$ where the discrete time control value $u(n)$ corresponds to the constant control value $u_c(t)$ applied in the sampling interval $[nT, (n+1)T)$.

Our goal is to minimize the infinite horizon cost functional $J_\infty(x_0, u) = \sum_{n=0}^{\infty} \ell(x_u(n), u(n))$ with running cost $\ell : X \times U \rightarrow \mathbb{R}_0^+$ by a static state feedback control law $\mu : X \rightarrow U$ which is applied according to the rule

$$x_\mu(0) = x_0, \quad x_\mu(n+1) = f(x_\mu(n), \mu(x_\mu(n))). \quad (2)$$

We denote the optimal value function for this problem by $V_\infty(x_0) := \inf_{u \in \mathcal{U}} J_\infty(x_0, u)$. The motivation for this problem stems from stabilizing the system (1) at a fixed point, i.e., at a point $x^* \in X$ for which there exists a control value $u^* \in U$ with $f(x^*, u^*) = x^*$ and $\ell(x^*, u^*) = 0$. Under mild conditions on ℓ it is known that the optimal feedback for J_∞ indeed asymptotically stabilizes the system with V_∞ as a Lyapunov function, see, e.g., [6].

Since infinite horizon optimal control problems are in general computationally infeasible, we use a receding horizon NMPC method in order to compute an approximately optimal feedback law. To this end, we consider the finite horizon functional

$$J_N(x_0, u) = \sum_{n=0}^{N-1} \ell(x_u(n), u(n)) \quad (3)$$

with *optimization horizon* $N \in \mathbb{N}_{\geq 2}$ and optimal value function $V_N(x_0) := \inf_{u \in \mathcal{U}} J_N(x_0, u)$. By minimizing (3) over $u \in \mathcal{U}$ we obtain an optimal control sequence¹ $u^*(0), u^*(1), \dots, u^*(N-1)$ depending on the initial value x_0 . Implementing the first element of this sequence, i.e., $u^*(0)$, yields a new state $x_{u^*}(1, x_0)$ for which we redo the procedure, i.e., at the next time instant we minimize (3) for $x_0 := x_{u^*}(1, x_0)$. Iterative application of this procedure provides a control sequence on the infinite time interval. A corresponding closed loop representation of the type (2) is obtained as follows.

Definition 1. For $N \geq 2$ we define the MPC feedback law $\mu_N(x_0) := u^*(0)$, where u^* is a minimizing control for (3) with initial value x_0 .

In many papers in the (N)MPC literature additional stabilizing terminal constraints or terminal costs are added to the optimization objective (3) in order to ensure asymptotic stability of the NMPC closed loop despite the truncation of the horizon (see, e.g., the monograph [9] for a recent account of this theory). In contrast to this approach, here we investigate (3) without any changes. This is motivated by the fact that this “plain” NMPC scheme is the most easy one to implement and appears to be predominant in practical applications, cf. [8]. Another reason appears when looking at the infinite horizon performance of the NMPC feedback law μ_N given by $J_\infty(x_0, \mu_N) := \sum_{n=0}^{\infty} \ell(x_{\mu_N}(n), \mu_N(x_{\mu_N}(n)))$. As we will see, under a suitable controllability condition for NMPC without stabilizing constraints we can establish an upper bound for this value in terms of the optimal value function $V_\infty(x_0)$, which is in general not possible for schemes with stabilizing constraints.

3 Performance and stability analysis

In this section we summarize the main steps of the stability and suboptimality analysis of unconstrained NMPC schemes from [3, 4, 7] in a simplified setting. The cornerstone of our analysis is the following proposition which uses ideas from relaxed dynamic programming.

Proposition 1. Assume there exists $\alpha \in (0, 1]$ such that for all $x \in X$ the inequality

$$V_N(x) \geq V_N(f(x, \mu_N(x))) + \alpha \ell(x, \mu_N(x)) \quad (4)$$

holds. Then for all $x \in X$ the estimate

¹ For simplicity of exposition we assume that a minimizing control sequence u^* exists for (3). However, given that in this abstract formulation \mathcal{U} may be infinite dimensional we do not assume uniqueness of u^* .

$$\alpha V_\infty(x) \leq \alpha J_\infty(x, \mu_N) \leq V_N(x) \leq V_\infty(x) \quad (5)$$

holds. If, in addition, there exist $x^* \in X$ and \mathcal{K}_∞ -functions² α_1, α_2 such that the inequalities

$$\ell^*(x) := \min_{u \in \mathcal{U}} \ell(x, u) \geq \alpha_1(d(x, x^*)) \quad \text{and} \quad V_N(x) \leq \alpha_2(d(x, x^*)) \quad (6)$$

hold for all $x \in X$, then x^* is a globally asymptotically stable equilibrium for (2) with $\mu = \mu_N$ with Lyapunov function V_N .

Proof. See [7, Prop. 2.2] or [3, Prop. 2.4] and [3, Theorem 5.2]. \square

In order to compute α in (4) we use the following controllability property: we call the system (1) *exponentially controllable* with respect to the running cost ℓ if there exist constants $C \geq 1$ (overshoot bound) and $\sigma \in [0, 1)$ (decay rate) such that

$$\begin{aligned} &\text{for each } x \in X \text{ there exists } u_x \in \mathcal{U} \text{ with} \\ &\ell(x_{u_x}(n, x), u_x(n)) \leq C\sigma^n \ell^*(x) \quad \text{for all } n \in \mathbb{N}_0. \end{aligned} \quad (7)$$

This condition implies

$$V_N(x) \leq J_N(x, u_x) \leq \sum_{n=0}^{N-1} C\sigma^n \ell^*(x) = C \frac{1 - \sigma^N}{1 - \sigma} \ell^*(x) =: B_N(\ell^*(x)). \quad (8)$$

Hence, in particular (6) follows for $\alpha_2 = B_N \circ \alpha_3$ if the inequality

$$\alpha_1(d(x, x^*)) \leq \ell^*(x) \leq \alpha_3(d(x, x^*)) \quad (9)$$

holds for some $\alpha_1, \alpha_3 \in \mathcal{K}_\infty$ and all $x \in X$. Now consider an arbitrary $x \in X$ and let $u^* \in \mathcal{U}$ be an optimal control for $J_N(x, u)$, i.e., $J_N(x, u^*) = V_N(x)$. Note that by definition of μ_N the identity $x_{u^*}(1, x) = f(x, \mu_N(x))$ follows.

For the following lemma we abbreviate

$$\lambda_n = \ell(x_{u^*}(n, x), u^*(n)), \quad n = 0, \dots, N-1 \quad \text{and} \quad \nu = V_N(x_{u^*}(1, x)). \quad (10)$$

Lemma 1. *Assume (7) holds. Then the inequalities*

$$\sum_{n=k}^{N-1} \lambda_n \leq B_{N-k}(\lambda_k) \quad \text{and} \quad \nu \leq \sum_{n=0}^{j-1} \lambda_{n+1} + B_{N-j}(\lambda_{j+1}) \quad (11)$$

hold for $k = 0, \dots, N-2$ and $j = 0, \dots, N-2$.

Proof. See [3, Section 3 and Proposition 4.1]. \square

² A continuous function $\alpha : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ is said to be of class \mathcal{K}_∞ if it is strictly increasing and unbounded with $\alpha(0) = 0$.

The inequalities from Lemma 1 now lead to the following theorem.

Theorem 1. *Assume that the system (1) and ℓ satisfy the controllability condition (7). Then inequality (4) holds for all $x \in X$ with*

$$\alpha = \min_{\lambda_0, \dots, \lambda_{N-1}, \nu} \sum_{n=0}^{N-1} \lambda_n - \nu \quad (12)$$

subject to the constraints (11), $\lambda_0 = 1$ and $\lambda_1, \dots, \lambda_{N-1}, \nu \geq 0$.

Proof. See [3, Section 4]. \square

The consequence of this theorem for the performance of the NMPC closed loop, i.e., (2) with $\mu = \mu_N$, is as follows: if (1) and ℓ satisfy (7) and (9), then global asymptotic stability and the suboptimality estimate (5) are guaranteed whenever α from (12) is positive. In fact, regarding stability we can show more: the construction of an explicit example yields that whenever α from (12) is negative, then there is a system (1) and an ℓ which satisfy (7) and (9) but for which (2) with $\mu = \mu_N$ is not asymptotically stable, cf. [3, Theorem 5.3].

The key observation for computing an explicit expression for α in (4) is that the linear program in Theorem 1 can be solved explicitly.

Theorem 2. *Under the assumptions of Theorem 1 the value α from (12) is given by*

$$\alpha = 1 - \frac{(\gamma_N - 1) \prod_{i=2}^N (\gamma_i - 1)}{\prod_{i=2}^N \gamma_i - \prod_{i=2}^N (\gamma_i - 1)} \quad \text{with} \quad \gamma_i = C \frac{1 - \sigma^i}{1 - \sigma}. \quad (13)$$

Proof. See [4, Theorem 5.3]. \square

The explicit formula thus derived for α allows us to visualize the impact of the parameters C, σ in (7) on the value of α in (4). As an example, Figure 1 shows the regions in the C, σ -plane for which $\alpha > 0$ and thus asymptotic stability holds for optimization horizons $N = 2, 4, 8$, and 16. Note that since α is increasing in N the stability region for N is always contained in the stability region for all $\tilde{N} > N$.

Figure 1 clearly shows the different roles of the parameters C and σ in (7): While for fixed C the minimal stabilizing N for varying σ is usually larger than 2, for fixed σ it is always possible to achieve stability with $N = 2$ by reducing C . Thus, the overshoot bound C plays a decisive role for the stability and performance of NMPC schemes.

An important observation in this context is that C and σ do not only depend on the control system but also on the running cost ℓ . Hence, ℓ can be used as a design parameter in order to “tune” C and σ with the goal to obtain good closed loop performance with small control horizons N by reducing C as much as possible. For examples see, e.g., [3] and [2] and the following section in which we will illustrate and explain this procedure for a semilinear parabolic PDE control system.

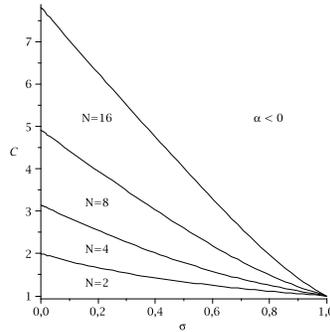


Fig. 1. Stability regions for various optimization horizons N depending on C and σ from (7)

4 A numerical case study

In practice, for many complex control systems and associated running cost functions ℓ it is difficult if not impossible to exactly determine the constants C and σ . However, by means of a controlled semilinear parabolic PDE, in this section we demonstrate that an exact computation of these constants is not necessarily needed in order to understand differences in the NMPC closed loop behavior for different running costs ℓ .

The first model we are considering is the semilinear parabolic PDE

$$y_t(t, x) = \nu y_{xx}(t, x) - y_x(t, x) + \mu (y(t, x) - y(t, x)^3) + u(t, x) \quad (14)$$

with distributed control $u \in L^\infty(\mathbb{R} \times \Omega, \mathbb{R})$ and $\Omega = (0, 1)$ and real parameters $\nu = 0.1$, $\mu = 10$. Here y_t and y_x denote the partial derivatives with respect to t and x , respectively and y_{xx} denotes the second partial derivative with respect to x .

The solution y of (14) is supposed to be continuous in $\overline{\Omega}$ and to satisfy the boundary and initial conditions

$$y(t, 0) = 0, y(t, 1) = 0 \text{ for all } t \geq 0 \text{ and } y(0, x) = y_0(x) \text{ for all } x \in \Omega \quad (15)$$

for some given continuous function $y_0 : \overline{\Omega} \rightarrow \mathbb{R}$ with $y_0(0) = y_0(1) = 0$.

Observe that we have changed notation here in order to be consistent with the usual PDE notation: $x \in \Omega$ is the independent space variable while the unknown function $y(t, \cdot) : \Omega \rightarrow \mathbb{R}$ in (14) is the state now. Hence, the state is now denoted by y (instead of x) and the state space of this PDE control system is a function space, more precisely the Sobolev space $H_0^1(\Omega)$, although the specific form of this space is not crucial for the subsequent reasoning.

Figure 2 shows the solution of the uncontrolled system (14), (15), i.e., with $u \equiv 0$. For growing t the solution approaches an asymptotically stable

steady state $y^{**} \neq 0$. The figure (as well as all other figures in this section) was computed numerically using a finite difference scheme with 50 equidistant nodes on $(0, 1)$ (finer resolutions did not yield significantly different results) and initial value y_0 with $y_0(0) = y_0(1) = 0$, $y_0|_{[0.02, 0.3]} \equiv -0.1$, $y_0|_{[0.32, 0.98]} \equiv 0.1$ and linear interpolation in between.

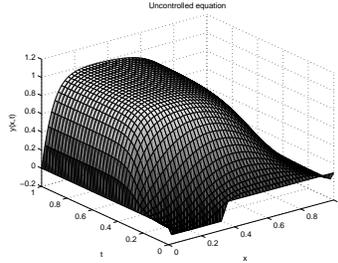


Fig. 2. Solution $y(t, x)$ of (14), (15) with $u \equiv 0$.

By symmetry of (14) the function $-y^{**}$ must be an asymptotically stable steady state, too. Furthermore, from (14) it is obvious that $y^* \equiv 0$ is another steady state, which is, however, unstable. Our goal is now to use NMPC in order to stabilize the unstable equilibrium $y^* \equiv 0$.

To this end we consider the sampled-data system corresponding to (14) with sampling period $T = 0.025$ and denote the state of the sampled-data system at the n -th sampling instant, i.e., at time nT by $y(n, \cdot)$. For penalizing the distance of the state $y(n, \cdot)$ to $y^* \equiv 0$ a popular choice in the literature is the L^2 -functional

$$\ell(y(n, \cdot), u(n, \cdot)) = \|y(n, \cdot)\|_{L^2(\Omega)}^2 + \lambda \|u(n, \cdot)\|_{L^2(\Omega)}^2 \quad (16)$$

with $\lambda = 0.1$ which penalizes the mean squared distance from $y(n, \cdot)$ to $y^* \equiv 0$.

Another possible choice of measuring the distance to $y^* \equiv 0$ is obtained by using the H^1 norm for $y(n, \cdot)$ in ℓ , i.e,

$$\ell(y(n, \cdot), u(n, \cdot)) = \|y(n, \cdot)\|_{L^2(\Omega)}^2 + \|y_x(n, \cdot)\|_{L^2(\Omega)}^2 + \lambda \|u(n, \cdot)\|_{L^2(\Omega)}^2, \quad (17)$$

which in addition to the L^2 distance (16) also penalizes the mean squared distance from $y_x(n, \cdot)$ to $y_x^* \equiv 0$. Figures 3 and 4 show the respective NMPC closed loop solutions with optimization horizons $N = 3$ and $N = 11$.³

Figure 3 indicates that for $N = 3$ the NMPC scheme with ℓ from (16) does not stabilize the system at $y^* \equiv 0$ while for ℓ from (17) it does. For (16) we need an optimization horizon of at least $N = 11$ in order to obtain a stable

³ The computations were performed with PCC, <http://www.nonlinearmpc.com/>

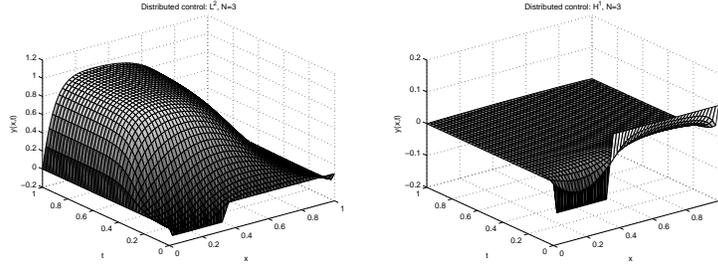


Fig. 3. NMPC closed loop for (14) with $N = 3$ and ℓ from (16)(left) and (17)(right)

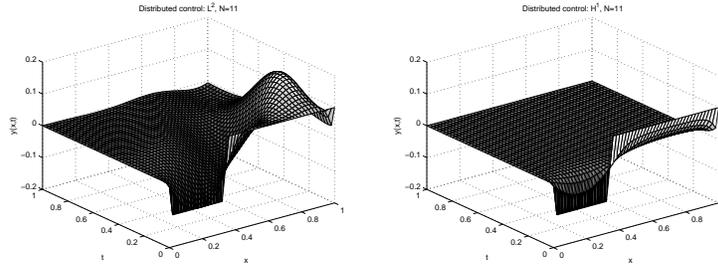


Fig. 4. NMPC closed loop for (14) with $N = 11$ and ℓ from (16)(left) and (17)(right)

closed loop solution, cf. Figure 4. For ℓ from (17) the right images in Figure 3 and 4 show that enlarging the horizon does not improve the solution.

Using our theoretical results we can explain why ℓ from (17) performs much better for small horizons N . For this example our controllability condition (7) reads

$$\ell(y(n, \cdot), u(n, \cdot)) \leq C\sigma^n \ell^*(y(0, \cdot)). \tag{18}$$

For ℓ from (16) this becomes

$$\|y(n, \cdot)\|_{L^2(\Omega)}^2 + \lambda \|u(n, \cdot)\|_{L^2(\Omega)}^2 \leq C\sigma^n \|y(0, \cdot)\|_{L^2(\Omega)}^2. \tag{19}$$

Now in order to control the system to $y^* \equiv 0$, in (14) the control needs to compensate for y_x and $\mu(y(t, x) - y(t, x)^3)$, i.e., any control steering $y(n, \cdot)$ to 0 must satisfy

$$\|u(n, \cdot)\|_{L^2(\Omega)}^2 \approx \|y_x(n, \cdot)\|_{L^2(\Omega)}^2 + \|\mu(y(n, \cdot) - y(n, \cdot)^3)\|_{L^2(\Omega)}^2. \tag{20}$$

This implies — regardless of the value of σ — that the overshoot bound C in (19) is large if $\|y_x(n, \cdot)\|_{L^2(\Omega)}^2 \gg \|y(0, \cdot)\|_{L^2(\Omega)}^2$ holds, which is the case in our example.

For ℓ from (17) inequality (18) becomes

$$\begin{aligned} & \|y(n, \cdot)\|_{L^2(\Omega)}^2 + \|y_x(n, \cdot)\|_{L^2(\Omega)}^2 + \lambda \|u(n, \cdot)\|_{L^2(\Omega)}^2 \\ & \leq C \sigma^n \left(\|y(0, \cdot)\|_{L^2(\Omega)}^2 + \|y_x(0, \cdot)\|_{L^2(\Omega)}^2 \right). \end{aligned} \quad (21)$$

Due to the fact that $\|y_x(0, \cdot)\|_{L^2(\Omega)}^2 \gg \|y(0, \cdot)\|_{L^2(\Omega)}^2$ holds in our example, the approximate equation (20) does not imply large C in (21), which explains the considerable better performance for ℓ from (17).

The fact that the H^1 -norm penalizes the distance to $y^* \equiv 0$ in a “stronger” way might lead to the conjecture that the better performance for this norm is intuitive. Our second example shows that this is not necessarily the case. This example is similar to the equation (14), (15), except that the distributed control is changed to Dirichlet boundary control. Thus, (14) becomes

$$y_t(t, x) = \nu y_{xx}(t, x) - y_x(t, x) + \mu (y(t, x) - y(t, x)^3), \quad (22)$$

again with $\nu = 0.1$ and $\mu = 10$, and (15) changes to

$$y(t, 0) = u_0(t), \quad y(t, 1) = u_1(t) \text{ for all } t \geq 0, \quad y(0, x) = y_0(x) \text{ for all } x \in \Omega$$

with $u_0, u_1 \in L^\infty(\mathbb{R}, \mathbb{R})$. The cost functions (16) and (17) change to

$$\ell(y(n, \cdot), u(n, \cdot)) = \|y(n, \cdot)\|_{L^2(\Omega)}^2 + \lambda(u_0(n)^2 + u_1(n)^2) \quad (23)$$

and

$$\ell(y(n, \cdot), u(n, \cdot)) = \|y(n, \cdot)\|_{L^2(\Omega)}^2 + \|y_x(n, \cdot)\|_{L^2(\Omega)}^2 + \lambda(u_0(n)^2 + u_1(n)^2), \quad (24)$$

respectively, again with $\lambda = 0.1$.

Due to the more limited possibilities to control the equation the problem obviously becomes more difficult, hence we expect to need larger optimization horizons for stability of the NMPC closed loop. However, what is surprising at the first glance is that ℓ from (23) stabilizes the system for smaller horizons than ℓ from (24), as the numerical results in Figure 5 confirm.

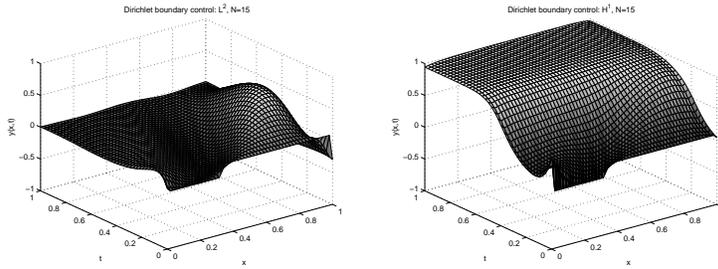


Fig. 5. NMPC closed loop for (22) with $N = 15$ and ℓ from (16)(left) and (17)(right)

A closer look at the dynamics reveals that we can again explain this behaviour with our theoretical results. In fact, steering the chosen initial solution to $y^* = 0$ requires u_1 to be such that a rather large gradient appears close to 1. Thus, during the transient phase $\|y_x(n, \cdot)\|_{L^2(\Omega)}^2$ becomes large which in turn causes ℓ from (24) to become large and thus causes a large overshoot bound C in (18). In ℓ from (23), on the other hand, these large gradients are not “visible” which is why the overshoot in (18) is smaller and thus allows for stabilization with smaller N .

5 Conclusions

In this paper we have shown how performance of NMPC schemes can be analyzed on basis of a controllability condition involving both the system dynamics and the cost function used in the optimization. The example of a semilinear parabolic PDE with distributed and boundary control illustrates how our theoretical results can be used for analyzing concrete systems.

References

1. Allgöwer F, Zheng A, eds. (2000), Nonlinear model predictive control, Birkhäuser, Basel
2. Altmüller N, Grüne L, Worthmann K (2010), Instantaneous control of the linear wave equation, Proceedings of MTNS 2010, Budapest, Hungary, to appear
3. Grüne L (2009) Analysis and design of unconstrained nonlinear MPC schemes for finite and infinite dimensional systems, SIAM J. Control Optim., 48, pp. 1206–1228
4. Grüne L, Pannek J, Seehafer M, Worthmann K (2009), Analysis of unconstrained nonlinear MPC schemes with time varying control horizon, Preprint, Universität Bayreuth; submitted
5. Grüne L, Pannek J, Worthmann K (2009), A networked unconstrained nonlinear MPC scheme, Proceedings of ECC 2009, Budapest, Hungary, pp. 371–376
6. Grüne L, Nešić D (2003), Optimization based stabilization of sampled-data nonlinear systems via their approximate discrete-time models, SIAM J. Control Optim., 42, pp. 98–122
7. Grüne L, Rantzer A (2008), On the infinite horizon performance of receding horizon controllers, IEEE Trans. Automat. Control, 53, pp. 2100–2111
8. Qin S, Badgwell T (2003), A survey of industrial model predictive control technology, Control Engineering Practice, 11, pp. 733–764
9. Rawlings JB, Mayne DQ (2009), Model Predictive Control: Theory and Design, Nob Hill Publishing, Madison

Index

asymptotic stability, 4

boundary control, 9

controllability, 4

distributed control, 6

H^1 -norm, 7

L^2 -norm, 7

model predictive control, 1

nonlinear model predictive control, 1

optimization horizon, 3

overshoot, 4

semilinear parabolic PDE, 6

suboptimality, 3