On the infinite horizon performance of receding horizon controllers

Lars Grüne*
Mathematisches Institut
Universität Bayreuth
95440 Bayreuth, Germany
lars.gruene@uni-bayreuth.de

Anders Rantzer
Department of Automatic Control
LTH, Lund University
Box 118, 21100 Lund, Sweden
rantzer@control.lth.se

October 5, 2007

Abstract: Receding horizon control is a well established approach for control of systems with constraints and nonlinearities. Optimization over an infinite time-horizon, which is often computationally intractable, is therein replaced by a sequence of finite horizon problems. This paper provides a method to quantify the performance degradation that comes with this approximation. Results are provided for problems both with and without terminal costs and constraints and for both exactly and practically asymptotically stabilizable systems.

Keywords: Receding horizon control, model predictive control, dynamic programming, suboptimality

1 Introduction

Receding horizon control (RHC), often also termed model predictive control (MPC), is by now a well established method for the optimal control of linear and nonlinear systems [1, 3, 16]. One way of interpreting this method in a discrete time setting is the following: In order to approximate the solution to a (computationally intractable) infinite horizon optimal control problem, a sequence of — often suitably constrained — finite horizon optimal control problems is solved. Then in each time step the first element of the resulting optimal control sequence is used as a feedback control value for the current state.

This interpretation immediately leads to the question about the suboptimality of the resulting RHC feedback: how good is the resulting RHC controller with respect to the original infinite horizon cost functional? This question has been addressed in a number of papers and estimates for the infinite horizon performance are given, e.g., for nonlinear systems in [13] with zero endpoint constrained finite horizon problems and in [10] with (local) Lyapunov function terminal cost, respectively, and for linear systems in [19] without terminal costs or constraints. The paper [4], again dealing with linear systems, presents a negative result about the monotonicity of the infinite horizon performance of RHC controllers

*Parts of the research for this paper were carried out while the first author was visiting the Department of Automatic Control at Lund University.

with terminal cost. In this paper we consider discrete time nonlinear systems on arbitrary metric spaces both without and with terminal costs and constraints. For these systems we derive explicit estimates for the degree of suboptimality which in particular induce bounds for the minimal optimization horizon needed for stabilization.

Despite the fact that RHC schemes without stabilizing terminal costs and constraints possess a number of advantages (see the discussion at the beginning of Section 4), in the literature they appear less frequently addressed. Notable exceptions from this rule are, among others, the papers by Shamma and Xiong [19], Primbs and Nevistić [17], Jadbabaie and Hauser [12] and Grimm et al. [5]. In [19], for discrete time linear finite dimensional systems it is shown that the knowledge of the finite horizon optimal value functions can be used in order to compute a bound on the degree of suboptimality of the receding horizon controller and that this result can be applied by using numerical approximations of the optimal value functions. This result was extended in [17] to linear systems subject to linear constraints. For nonlinear continuous time systems with stabilizable linearization, it is shown in [12] that a receding horizon controller stabilizes the system for sufficiently large optimization horizon. The paper [5] addressing discrete time nonlinear systems is in a similar spirit, however, instead of imposing conditions on the linearization it uses rather general and genuinely nonlinear detectability conditions. The papers [5, 12] have in common that suboptimality is not considered and that the stability results are merely asymptotic, i.e., for sufficiently large horizon. Compared to these four papers, our contribution is most similar to [19] and [5]: as in [19] we derive explicit bounds on the degree of suboptimality of the RHC closed loop system and as in [5] we address general discrete time nonlinear systems (here even on arbitrary metric spaces) without imposing conditions on the linearization. The main difference to [19] is that our conditions do not rely on the knowledge of the finite time optimal value functions but rather on suitable bounds on these functions, see Remark 4.10 for details. The main difference to [5] is that we derive explicit bounds for suboptimality and stability instead of asymptotic estimates for sufficiently large horizons; furthermore, some aspects of our conditions are more general, cf. Remark 5.11.

Concerning our results for RHC schemes with terminal costs and constraints, a property which is related to our question is the inverse optimality of RHC controllers: it is well known that under suitable conditions RHC controllers are infinite horizon inversely optimal, i.e., they are optimal for an infinite horizon optimal control problem with a suitably adjusted running cost, see, e.g., [2, 16, 15]. However, this property does not yield estimates for the suboptimality with respect to the original running cost, as the adjusted running cost contains precisely those terms which characterize the mismatch between the infinite horizon optimal and the RHC controller, i.e., those terms that we intend to estimate. A paper which is closer to our approach is [10] in which it is shown that infinite horizon optimality is maintained if the terminal cost equals the infinite horizon optimal value function. In this paper we extend this result by showing how and in which quantitative sense this results remains true if the terminal cost merely approximates the infinite horizon optimal value function, cf. Remark 6.3 for details.

The common technique for all our results in the different settings is a relaxed dynamic programming inequality. Inequalities of such type have been used frequently in the optimal control literature, however, a systematic study seems to have performed only recently in [14, 18]. A specific relaxed dynamic programming inequality, formulated in Proposition

2.2, is the cornerstone of our analysis.

The paper is organized as follows. We start by describing the setup and some preliminary results on relaxed dynamic programming in Section 2. Section 4 contains our main results on RHC without using terminal costs and constraints and presents an assumptions under which we can prove suboptimality for RHC schemes for asymptotically stabilizable systems. Section 5 shows how to relax this assumption in order to deal with only practically stabilizable systems and Section 6 presents our results for RHC with terminal cost and terminal constraints. The final Section 7 concludes our paper.

2 Setup and preliminary results

We consider a nonlinear discrete time system given by

$$x(n+1) = f(x(n), u(n)), \quad x(0) = x_0 \tag{2.1}$$

with $x(n) \in X$ and $u(n) \in U$ for $n \in \mathbb{N}_0$. Here we denote the space of control sequences $u : \mathbb{N}_0 \to U$ by $\mathcal{U}$ and the solution trajectory for some $u \in \mathcal{U}$ by $x_u(n)$. Here the state space X is an arbitrary metric space, which in particular means that our results also apply to the discrete time dynamics induced by a sampled infinite dimensional system, cf. [11] for a continuous time analysis of this setting.

Our goal is to find a feedback control law $u = \mu(x)$ minimizing the infinite horizon cost

$$J_\infty(x_0, u) = \sum_{n=0}^{\infty} l(x_u(n), u(n)), \tag{2.2}$$

with running cost $l : X \times U \to \mathbb{R}_0^+$. We denote the optimal value function for this problem by

$$V_\infty(x_0) = \inf_{u \in \mathcal{U}} J_\infty(x_0, u).$$

If this optimal value function is known, it is easy to prove using Bellman's optimality principle that the optimal feedback law μ is given by

$$\mu(x) := \operatorname*{argmin}_{u \in U} \{V_\infty(f(x, u)) + l(x, u)\}.$$

Remark 2.1 We assume throughout this paper that in all relevant expressions the minimum with respect to $u \in U$ is attained. Although it is possible to give modified statements using approximate minimizers, we decided to make this assumption in order to simplify and streamline the presentation. □

Since infinite horizon optimal control problems are often computationally infeasible, we use a receding horizon approach in order to compute a controller by considering the finite horizon problem given by

$$J_N(x_0, u) = \sum_{n=0}^{N-1} l(x_u(n), u(n)) \tag{2.3}$$

for $N \in \mathbb{N}_0$ (using $\sum_{n=0}^{-1} = 0$) with optimal value function

$$V_N(x_0) = \inf_{u \in \mathcal{U}} J_N(x_0, u). \tag{2.4}$$

A variant of this approach often considered in the literature is obtained when we add a terminal cost $F : X \to \mathbb{R}_0^+$ to the problem. In this case, (2.3) is changed to

$$J_N(x_0, u) = \sum_{n=0}^{N-1} l(x_u(n), u(n)) + F(x_u(N)). \tag{2.5}$$

Another common extension is the introduction of a terminal constraint set $X_f \subset X$ for the finite horizon optimization, which amounts to replacing (2.4) by

$$V_N(x_0) = \inf_{u \in \mathcal{U} : x_u(N) \in X_f} J_N(x_0, u). \tag{2.6}$$

Here we assume that the set X_f is forward invariant, i.e., for each $x \in X_f$ there exists $u \in U$ with $f(x, u) \in X_f$. When imposing such a terminal constraint, the domain of points on which V_N is defined is restricted to the feasible set X_N, which is the set of initial values which can be controlled to X_f in at most N steps, i.e.,

$$X_N := \{x_0 \in X \,|\, \text{there exists } u \in \mathcal{U} \text{ with } x_u(N) \in X_f\}.$$

Note that (2.3) is a special case of (2.5), with $F \equiv 0$, and that (2.4) is a special case of (2.6) with $X_f = X$. Here we have stated (2.3) and (2.4) explicitly because it is the simplest version of receding horizon control and a major part of our results apply particularly to this case.

Based on this finite horizon optimal value function we define a feedback law μ_N by picking the first element of the optimal control sequence for this problem. Since Bellman's optimality principle for the functions V_N reads

$$V_N(x) = \min_{u \in U_N(x)} \{V_{N-1}(f(x, u)) + l(x, u)\}, \tag{2.7}$$

with $U_N(x) := \{u \in U \,|\, f(x, u) \in X_{N-1}\}$ this amounts to defining

$$\mu_N(x) := \operatorname*{argmin}_{u \in U_N(x)} \{V_{N-1}(f(x, u)) + l(x, u)\}. \tag{2.8}$$

Note that the feedback law μ_N is not the optimal controller for the problem (2.4). However, the optimal trajectory for this problem can be expressed via the controllers $\mu_1, \ldots, \mu_N$ in the following inductive fashion

$$x(0) = x_0, \quad x(n + 1) = f(x(n), \mu_{N-n}(x(n))), \;\; n = 0, \ldots, N - 1. \tag{2.9}$$

The goal of the present paper is to give estimates about the suboptimality of the feedback μ_N for the infinite horizon problem. More precisely, if x_{μ_N} denotes the solution of the closed loop system

$$x_{\mu_N}(n + 1) = f(x_{\mu_N}(n), \mu_N(x_{\mu_N}(n))), \quad x_{\mu_N}(0) = x_0$$

and we define the infinite horizon cost corresponding to μ_N by

$$V_\infty^{\mu_N}(x_0) := \sum_{n=0}^{\infty} l(x_{\mu_N}(n), \mu_N(x_{\mu_N}(n))$$

then we are interested in upper bounds for this infinite horizon value, either in terms of the finite horizon optimal value function V_N or in terms of the infinite horizon optimal value function V_∞. In particular, the latter will give us estimates about the "degree of suboptimality" of the controller μ_N.

A result closely related to our problem is the inverse optimality of receding horizon schemes, see [16, Section 3.5], [15] or [2]. This result states that the controller μ_N is an infinite horizon optimal controller for the cost

$$\bar{l}(x,u) := l(x,u) + V_{N-1}(f(x,\mu_N(x))) - V_N(f(x,\mu_N(x)))$$

and that V_N is the corresponding infinite horizon optimal value function. The importance of this result lies in the fact that it establishes infinite horizon optimality for the resulting controller. However, its disadvantage is that — unless one has informations about the relation between l and $\bar{l}$ — it only does so for the new running cost $\bar{l}$. Thus, in general this result does not tell us much about the performance of μ_N with respect to the *original* cost l, which is what we are interested in.

Note that in (undiscounted) infinite horizon optimal control one is in general interested in nonnegative running cost functions, in order to be able to conclude, e.g., stability of the closed loop system. Thus, in this context the inverse optimality result is only useful if $\bar{l}(x,u) \geq 0$, implying the condition $V_{N-1}(f(x,\mu_N(x))) - V_N(f(x,\mu_N(x))) \geq -l(x,u)$. Essentially, our approach is based on deriving estimates of this type (for precise formulations see the Lemmas 4.1 and 5.3) from appropriate conditions on either the functions V_N or on the running cost l.

The approach we take in this paper relies on recently developed results on relaxed dynamic programming [14, 18]. In the remainder of this section we present a variant of the basic relaxed dynamic programming inequality in Proposition 2.2 and give a sufficient condition for it which is adapted to our receding horizon setting in Lemma 4.1.

Proposition 2.2 Consider a feedback law $\tilde{\mu} : X \to U$ and a function $\widetilde{V} : X \to \mathbb{R}_0^+$ satisfying the inequality

$$\widetilde{V}(x) \geq \widetilde{V}(f(x,\tilde{\mu}(x))) + \alpha l(x,\tilde{\mu}(x)) \tag{2.10}$$

for some $\alpha \in [0,1]$ and all $x \in X$. Then for all $x \in X$ the estimate

$$\alpha V_\infty(x) \leq \alpha V_\infty^{\tilde{\mu}}(x) \leq \widetilde{V}(x)$$

holds.

Proof: The proof is similar to that of [18, Proposition 3]: Consider $x_0 \in X$, the trajectory $x(n) = x_{\tilde{\mu}}(n)$ generated by the closed loop system using $\tilde{\mu}$, and the control sequence generated by $u(n) = \tilde{\mu}(x(n))$. Then from (2.10) we obtain

$$\alpha l(x(n),u(n)) \leq \widetilde{V}(x(n)) - \widetilde{V}(x(n+1)).$$

Thus, summing over n yields

$$\alpha \sum_{n=0}^{K-1} l(x(n), u(n)) \leq \widetilde{V}(x(0)) - \widetilde{V}(x(K)) \leq \widetilde{V}(x(0)).$$

Thus, $\widetilde{V}$ is an upper bound on $\alpha V_\infty^{\tilde{\mu}}$ and we immediately obtain

$$\alpha V_\infty(x) \leq \alpha V_\infty^{\tilde{\mu}}(x) \leq \widetilde{V}(x).$$

□

Remark 2.3 All of our results remain true if the set U of admissible control values is subject to — possibly state dependent — constraints.

Similarly, the set of states X does not necessarily need to be the original state space of the system. Indeed, all of our results immediately carry over if X is an arbitrary subset of the state space which is forward invariant under the receding horizon feedback. Another generalization would be to choose X as the feasible set of a state constrained problem, similar to the construction used in [17] in the linear setting in order to carry over the stability results from [19] to constrained systems. □

3 Summary of the main results

Before we turn to the technical presentation, in this section we give a brief non–technical summary of our results.

The main theme of this paper is to apply relaxed dynamic programming techniques to receding horizon control schemes for discrete time nonlinear systems on arbitrary metric spaces. The basis for our results is Proposition 2.2 which we are going to apply to $\widetilde{V} = V_N$, i.e., to the finite time optimal value function. Hence, we need to establish checkable conditions under which V_N satisfies (2.10). As we will see, these techniques are applicable to schemes with and without terminal costs and constraints as well as to "mixed" forms.

Our results for RHC schemes without terminal cost in Section 4 are in the spirit of [19, 17, 5, 12] and give conditions on the running cost under which Proposition 2.2 can be applied and thus precise suboptimality estimates (and thus stability, cf. Remark 4.6(i)) can be derived. Intuitively, our results state that if the instantaneous running cost contains sufficient information about the optimal value function, then the resulting controller will also be sub–optimal. Here, the term "sufficient information" is formalized by an inequality, requiring that the scaled running cost $(\gamma + 1)l$ for some $\gamma \geq 0$ be larger than the optimal value functions V_N, cf. Assumption 4.2. If this is the case, then the α in Proposition 2.2 and consequently the degree of suboptimality of the controller μ_N can be readily computed from the factor γ and the horizon length N. These conditions can be either checked analytically using a priori controllability information about the system, cf. Proposition 4.7 or numerically at the time the scheme (or a numerical simulation) is running, cf. Remark 4.6(ii).

A variant of these results is presented in Section 5 in which we do no longer assume that the system is controllable to the 0–level set of the running cost l. In this situation we can show the existence of a set around the 0–level which is forward invariant for the RHC controller and that the resulting trajectories are sub–optimal until they enter this set. This is similar to practical asymptotic stability results, in which a system behaves like an asymptotically stable system until it reaches a forward invariant set, which is why we call this property "practical suboptimality"

Our final set of results is given in Section 6 in which we deal with RHC schemes with Lyapunov function terminal costs, the situation most widely discussed in the literature. Again, we use a condition on the running cost l, however, in contrast to the results without terminal costs now the condition is on the relation between l and the terminal cost F. The results in Theorem 6.2, which extend those in [10], show that if the terminal cost F approximates the infinite horizon optimal value function V_∞ and is in a suitable sense compatible with the running cost l (cf. Assumption 6.1), then the degree of suboptimality precisely equals the difference $F - V_\infty$. Again, the proof is based on Proposition 2.2, now with $\alpha = 1$. As we only assume F to be a local Lyapunov function, these results are only valid on those regions of the state space from which the optimal trajectories enter the region on which the local Lyapunov function is valid, which may be small if the optimization horizon is small. However, this local terminal cost can be coupled with the conditions on the running cost outside the domain of the terminal cost. Thus, combining the two conditions in Theorem 6.4 we obtain a scheme which is locally stable and suboptimal by virtue of the local terminal cost and globally suboptimal by virtue of properties of the running cost away from the origin.

4 Results without terminal cost

The presumably simplest version of RHC schemes are those in which the infinite horizon functional is truncated to a finite horizon functional and no constraints are imposed. In the literature, this version appears less frequently addressed than versions with terminal costs and constraints, although it has a number of advantages compared to more complicated schemes: for nonlinear systems the Lyapunov function property of the terminal cost — which is typically needed in order to ensure stability — is difficult to establish globally. Local Lyapunov functions are easily computed at equilibria if the linearization is controllable but they require terminal constraints making the optimization problem computationally harder and restricting the operating region to the feasible set, often leading to large optimization horizons for large operating regions. Local Lyapunov functions are also difficult to obtain if more general sets than equilibria are to be stabilized or if the problem is time varying, like in tracking problems, where in our own numerical experiments we observed RHC without terminal cost to perform very well, see [7, 8].

It seems that one reason for not exploiting these advantages in practice — at least in the practice reported in the literature — is the lack of theoretical foundation, in particular a lack of rigorous stability proofs, given that within the vast amount of papers on RHC/MPC there are only few papers addressing this issue. Notable papers among these few are, e.g., the papers [19, 17] for linear and [5, 12] for nonlinear systems, whose contributions were already discussed in the introduction.

When using RHC schemes without terminal cost, the first straightforward observation is the inequality

$$V_M(x) \le V_N(x) \le V_\infty(x) \tag{4.1}$$

for all $M, N \in \mathbb{N}_0$ with $M \le N$.

Our approach now is to apply Proposition 2.2 to $\widetilde{V} = V_N$. More precisely, our goal is to provide checkable conditions under which V_N satisfies (2.10). For this purpose, the following straightforward observation is useful.

Lemma 4.1 Consider $N \in \mathbb{N}$ and the receding horizon feedback law μ_N. Assume that

$$V_N(f(x,\mu_N(x))) - V_{N-1}(f(x,\mu_N(x))) \le (1-\alpha)l(x,\mu_N(x)) \tag{4.2}$$

holds for some $\alpha \in [0,1]$ and all $x \in X$. Then $\widetilde{V} = V_N$ satisfies (2.10) and, in particular,

$$\alpha V_\infty^{\mu_N}(x) \le V_N(x)$$

holds for all $x \in X$.

Proof: Combining (2.7) and (2.8) and inserting (4.2) yields

$$\begin{aligned} V_N(x) &= \min_u \{V_{N-1}(f(x,u)) + l(x,u)\} = V_{N-1}(f(x,\mu_N(x))) + l(x,\mu_N(x)) \\ &\ge V_N(f(x,\mu_N(x))) - (1-\alpha)l(x,\mu_N(x)) + l(x,\mu_N(x)) \\ &= V_N(f(x,\mu_N(x))) + \alpha l(x,\mu_N(x)), \end{aligned}$$

which shows (2.10). Now Proposition 2.2 yields the assertion. □

The following assumption contains our central condition.

Assumption 4.2 For a given $N \in \mathbb{N}$ there exists $\gamma > 0$ such that the inequalities

$$V_2(x) \le (\gamma+1)V_1(x) \quad \text{and} \quad V_k(x) \le (\gamma+1)l(x,\mu_k(x)), \ \ k = 2,\dots,N$$

hold for all $x \in X$. □

Remark 4.3 If the inequality $V_N(x) \le (\gamma+1)l(x,u)$ holds for all $x \in X$ and all $u \in U$, then (4.1) immediately implies Assumption 4.2. □

Proposition 4.4 Let $N \ge 2$ and assume that Assumption 4.2 holds for this N. Then the inequality

$$\frac{(\gamma+1)^{N-2}}{(\gamma+1)^{N-2} + \gamma^{N-1}} V_N(x) \le V_{N-1}(x)$$

holds for all $x \in X$.

Proof: We first show that Assumption 4.2 implies the estimate

$$V_{k-1}(f(x,\mu_k(x)) \le \gamma l(x,\mu_k(x)) \tag{4.3}$$

for all $k = 2, \dots, N$ and all $x \in X$. In order to prove (4.3), we use the optimality principle

$$V_k(x) = V_{k-1}(f(x, \mu_k(x)) + l(x, \mu_k(x)).$$

Now Assumption 4.2 implies

$$\begin{aligned} V_{k-1}(f(x,\mu_k(x)) &= V_k(x) - l(x,\mu_k(x)) \\ &\le (\gamma+1)l(x,\mu_k(x)) - l(x,\mu_k(x)) \;=\; \gamma l(x,\mu_k(x)), \end{aligned}$$

which shows (4.3).

By induction over $k = 2, \dots, N$ we prove

$$\eta_k V_k(x) \le V_{k-1}(x) \qquad (4.4)$$

for

$$\eta_k = \frac{(\gamma+1)^{k-2}}{(\gamma+1)^{k-2} + \gamma^{k-1}}.$$

For $k = 2$ (4.4) follows directly from Assumption 4.2 because

$$V_2(x) \le (\gamma+1)V_1(x) = \frac{1}{\eta_2} V_1(x).$$

For the induction step $k \to k+1$, for $x \in X$ we obtain

$$\begin{aligned} V_k(x) &= V_{k-1}(f(x,\mu_k(x))) + l(x,\mu_k(x)) \\ &\ge \left(1 + \frac{1-\eta_k}{\gamma+\eta_k}\right) V_{k-1}(f(x,\mu_k(x))) + \left(1 - \gamma\frac{1-\eta_k}{\gamma+\eta_k}\right) l(x,\mu_k(x)) \\ &\ge \eta_k \left(1 + \frac{1-\eta_k}{\gamma+\eta_k}\right) V_k(f(x,\mu_k(x))) + \left(1 - \gamma\frac{1-\eta_k}{\gamma+\eta_k}\right) l(x,\mu_k(x)) \\ &= \eta_k \frac{\gamma+1}{\gamma+\eta_k}\{V_k(f(x,\mu_k(x))) + l(x,\mu_k(x))\} \\ &\ge \eta_k \frac{\gamma+1}{\gamma+\eta_k} \min_{u\in U}\{V_k(f(x,u)) + l(x,u)\} \;=\; \eta_k \frac{\gamma+1}{\gamma+\eta_k} V_{k+1}(x), \end{aligned}$$

where we have used (4.3) in the first inequality and the induction assumption in the second inequality. This implies (4.4) because

$$\eta_k \frac{\gamma+1}{\gamma+\eta_k} = \frac{(\gamma+1)^{k-2}}{(\gamma+1)^{k-2} + \gamma^{k-1}} \frac{\gamma+1}{\gamma + \frac{(\gamma+1)^{k-2}}{(\gamma+1)^{k-2}+\gamma^{k-1}}} = \frac{(\gamma+1)^{k-1}}{(\gamma+1)^{k-1} + \gamma^k} = \eta_{k+1}.$$

□

Combining Proposition 4.4 with Lemma 4.1 we can now derive the main result of this section.

Theorem 4.5 Consider $\gamma > 0$ and $N \in \mathbb{N}$ such that $(\gamma+1)^{N-2} > \gamma^N$ holds. Assume that Assumption 4.2 holds for these γ and N. Then the inequality

$$V_\infty^{\mu_N}(x) \le \frac{(\gamma+1)^{N-2}}{(\gamma+1)^{N-2} - \gamma^N} V_\infty(x)$$

holds for all $x \in X$. In particular, the inequality

$$\frac{V_\infty^{\mu_N}(x) - V_\infty(x)}{V_\infty(x)} \le \frac{\gamma^N}{(\gamma+1)^{N-2} - \gamma^N}$$

holds for the relative difference between $V_\infty^{\mu_N}$ and V_∞.

Proof: From Proposition 4.4 we obtain

$$V_N(x) - V_{N-1}(x) \le \left(\frac{(\gamma+1)^{N-2} + \gamma^{N-1}}{(\gamma+1)^{N-2}} - 1\right) V_{N-1}(x) = \frac{\gamma^{N-1}}{(\gamma+1)^{N-2}} V_{N-1}(x).$$

Using this inequality for $x = f(x, F_N(x))$ and combining it with inequality (4.3) for $k = N$ we can conclude

$$V_N(f(x, F_N(x)) - V_{N-1}(f(x, F_N(x))) \le \frac{\gamma^N}{(\gamma+1)^{N-2}} l(x, F_N(x)).$$

Hence we can apply Lemma 4.1 with

$$\alpha = 1 - \frac{\gamma^N}{(\gamma+1)^{N-2}} = \frac{(\gamma+1)^{N-2} - \gamma^N}{(\gamma+1)^{N-2}}$$

in order to obtain the assertion. □

The following remark summarizes a number of observations for our result.

Remark 4.6 (i) If the running cost l is positive definite and proper with respect to some compact set A and X is a finite dimensional space, then Theorem 4.5 implies asymptotic stability of A if $(\gamma+1)^{N-2} > \gamma^N$ or, equivalently, $N > 2 + 2\log(\gamma)/(\log(\gamma+1) - \log(\gamma))$. This follows from inequality (4.2) which holds by virtue of Theorem 4.5 and immediately yields that V_N is a Lyapunov function for the closed loop system. More generally, one can obtain stability replacing the positive definiteness of l by a detectability condition, for details we refer, e.g., to [5].

(ii) Our condition depends on the knowledge of γ which can be computed in two different ways: on the one hand it is possible to give sufficient analytical conditions on the dynamics of the system seen through the "output" $l(x, u)$. Such conditions are developed in the remainder of this section. On the other hand, our conditions can be checked during the runtime of the RHC algorithm or respective numerical simulations: although we have supposed the inequalities in Assumption 4.2 to hold for all $x \in X$, in the proofs of our results we only need these inequalities along the optimal trajectories which are known once the optimization problem has been solved. This way we can estimate the degree of suboptimality of the RHC feedback scheme a posteriori. Details of such run time estimates and algorithms which use this a posteriori information for an online adaptation of the optimization horizon N are currently under investigation.

(iii) Note that our results give precise suboptimality bounds on the performance once the quantity γ has been computed according to (ii). In particular, the results are not merely asymptotic ("for sufficiently large N") but provide precise quantitative information for the size of N needed in order to ensure a desired performance.

(iv) As an alternative to the inequalities in Assumption 4.2 one could use

$$V_k(f(x, \mu_k(x))) \leq \gamma l(x, \mu_k(x))$$

in order to prove a result similar to Theorem 4.5, see [9] for details. □

In the remainder of this section we investigate analytical conditions on the running cost l which ensure Assumption 4.2.

Proposition 4.7 Assume that there exists a function $W : X \to \mathbb{R}_0^+$ and constants $\alpha, \beta > 0$, and $0 \leq \lambda < 1$ such that for all $x \in X$ the following two conditions hold:

(i) $l(x, u) \geq \alpha W(x)$ for all $u \in U$

(ii) there exists a control sequence $u^* \in \mathcal{U}$ such that the corresponding solution $x(n)$ with $x(0) = x$ satisfies

$$l(x(n), u^*(n)) \leq \beta \lambda^n W(x).$$

Then Assumption 4.2 holds with $\gamma = \frac{\beta}{\alpha(1-\lambda)} - 1$.

Proof: Condition (ii) implies

$$\begin{aligned} V_k(x) &\leq V_\infty(x) \\ &\leq \sum_{n=0}^{\infty} l(x(n), u^*(n)) \leq \sum_{n=0}^{\infty} \beta\lambda^n W(x) = \frac{\beta}{1-\lambda} W(x). \end{aligned}$$

Combining this with condition (i) yields

$$V_k(x) \leq \frac{\beta}{1-\lambda} W(x) \leq \frac{\beta}{\alpha(1-\lambda)} l(x, u)$$

for all $u \in U$ and thus in particular for $u = \mu_1(x)$, implying Assumption 4.2 for $k = 2$ and for $u = \mu_k(x)$, implying Assumption 4.2 for $k \geq 3$. □

Remark 4.8 (i) The condition is a combination of an observability condition (i) and a controllability condition (ii). More precisely, condition (ii) is an exponential controllability condition for the running cost l. Note that exponentially converging cost functions can always be constructed from control Lyapunov functions, however, since such control Lyapunov functions are hard to find, this approach may not be feasible. In an RHC context, exponential controllability conditions for the running cost are discussed in [5, Section III], in particular for homogeneous systems. Note that if the system is exponentially controllable to the origin, then the condition is always satisfied for $l(x, u) = \|x\|^2 + \|u\|^2$ and $W(x) = \|x\|^2$, although other choices of l and W may yield better constants α, β and γ.

(ii) In order to check the condition we need to know a null controlling control sequence u^*. Note, however, that this sequence does not need to be optimal and that it does not need to be in feedback form in order to apply the proposition. Example 4.9, below, illustrates this procedure. □

For Theorem 4.5 we do now have three types of assumptions and conditions, which differ in the type of information used:

- an assumption involving the optimal value functions and the RHC controllers (Assumption 4.2)
- a sufficient condition involving the optimal value functions (Remark 4.3)
- a sufficient condition involving an auxiliary function W (Proposition 4.7)

The following simple example highlights the difference between these assumptions.

Example 4.9 Consider the linear one–dimensional control system

$$x(n+1) = 2x(n) + u(n) =: f(x(n), u(n))$$

with $x(n) \in X = \mathbb{R}$ and $u(n) \in U = \mathbb{R}$. We first consider the running cost

$$l(x) = x^2.$$

Here it is easy to solve the infinite horizon optimal control problem, because for $\mu(x) = -2x$ the related optimal value function

$$V_\infty^\mu(x) = x^2$$

satisfies the optimality principle, because

$$x^2 = V_\infty^\mu(x) = \inf_{u\in\mathbb{R}} \{l(x) + V_\infty^\mu(f(x,u))\} = \inf_{u\in\mathbb{R}} \{x^2 + (2x+u)^2\} = x^2.$$

Using the same argument one also sees that the finite time optimal value functions are given by

$$V_N(x) = x^2, \quad N \geq 1$$

with corresponding RHC feedback laws

$$\mu_N(x) = -2x, \quad N \geq 2.$$

Thus, for $N \geq 2$, the RHC controller is indeed optimal for the infinite horizon problem.

This optimality property can be obtained from Theorem 4.5 using Assumption 4.2: this Assumption is satisfied for each $N \in \mathbb{N}$ with $\gamma = 0$ for $X = \mathbb{R}$. Thus, for each $N \geq 2$ we obtain the estimate

$$\frac{V_\infty^{\mu_N}(x) - V_\infty(x)}{V_\infty(x)} \leq \frac{\gamma^N}{(\gamma+1)^{N-2} - \gamma^N} = 0,$$

i.e., a sharp estimate.

Note that for checking Assumption 4.2 directly we have used information about the RHC controller, which we cannot expect to know in general. If this information is not available, Remark 4.3 can be used instead: its assumptions are easily verified by estimating V_k using the control sequence $u^* = (-2x(0), 0, 0, \ldots)$, yielding $V_k(x) \leq x^2$ and thus again $\gamma = 0$.

In order to illustrate the use of Proposition 4.7 we alter the running cost to

$$l(x,u) = x^2 + u^2.$$

Then, using $W(x) = x^2$ one obtains (i) with $\alpha = 1$. Applying the control sequence $u^* = (-x(0), -x(0)/2, -x(0)/4, \ldots)$ yields (ii) with $\beta = 2$ and $\lambda = 1/4$, resulting in $\gamma = 8/3$. Table 4.1 shows the minimal horizon length N needed according to Theorem 4.5 in order to ensure the given values for the relative accuracy.

relative accuracy	needed horizon length N
0.50	12
0.10	16
0.01	23

Table 4.1: Performance for running cost $l(x,u) = x^2 + u^2$

Note that we do not claim that these estimates in Table 4.1 are tight or even optimal. In particular, the use of other sequences u^* might lead to smaller values of γ and hence tighter estimates. We have chosen the given sequence u^* because it allows for easy computations. In fact, the possibility to use arbitrary null controlling control sequences allowing for easy computations can be considered as a particular advantage of our approach. □

Remark 4.10 Our results bear some similarities with those obtained for linear finite dimensional systems in [19]. Apart from the fact that here we treat nonlinear systems on arbitrary metric spaces, the main difference of our results to [19] is that we provide sufficient conditions on the running cost l in order to ensure that the difference between V_{N-1} and V_N is small, while in [19] this is a condition which is to be verified by computing numerical approximations to the optimal value functions V_{N-1} and V_N. While for linear systems — at least in low dimensions — the numerical computation of $\|V_{N-1} - V_N\|$ is a feasible task, in our nonlinear setting on arbitrary metric spaces this is typically impossible, which is why our conditions, which can be either verified analytically or checked numerically at run time, are preferable. It may, however, be possible to modify the approach in [19] in order to yield conditions which are also checkable at run time. □

5 Practical optimality

There are cases where it is too optimistic to expect that the conditions presented in the last section hold. For instance, it may happen that the discrete time system considered is obtained from sampling a continuous time system with zero order hold and even if this continuous time system is controllable to some fixed point x^* it is likely that the corresponding sampled–data system is only controllable to a neighborhood of this fixed point (see the example in [6, Section 9.4]), i.e., it is only *practically* stabilizable at x^*. In this case, for a positive definite running cost with respect to x^*, i.e., $l(x,u) = 0 \Leftrightarrow x = x^*$, it is not possible to find a control sequence yielding $l(x(n), u(n)) \to 0$ and it is easily seen that in this case Assumption 4.2 is not satisfied because the functions V_k grow unboundedly for $k \to \infty$.

If we are able to compute a forward invariant stabilizable neighborhood $\mathcal{N}$ of x^* then it is straightforward to apply our results from the previous section to running costs l which are positive definite with respect to $\mathcal{N}$. In practice, however, it may be impossible to compute such a set $\mathcal{N}$ which makes the design of an appropriate cost function a difficult task. In this case, a much simpler approach is to choose l positive definite with respect to x^* ignoring the lack of exact stabilizability. Since such a function is smaller near x^* than far away from x^* one may expect that the RHC controller will still drive the system to a neighborhood of x^*, thus yielding the closed loop system practically stable and — more important in our context — yielding suboptimal trajectories before reaching this neighborhood.

In this section we relax our Assumption 4.2 in order to cope with this setting. Before we investigate practical versions of our relaxed dynamic programming assumptions and statements and illustrate them in Example 5.10, below, we prove appropriate practical versions of the preliminary results from Section 2. We start with a practical version of Proposition 2.2.

Proposition 5.1 Consider a feedback law $\tilde{\mu} : X \to U$ and a nonnegative function $\widetilde{V} : X \to \mathbb{R}_0^+$ satisfying the inequality

$$\widetilde{V}(x) \geq \widetilde{V}(f(x, \tilde{\mu}(x))) + \min\{\alpha(l(x, \tilde{\mu}(x)) - \varepsilon),\, l(x, \tilde{\mu}(x)) - \varepsilon\} \tag{5.1}$$

for some $\alpha \in [0, 1]$, some $\varepsilon > 0$ and all $x \in X$.

Let $\mathcal{L} \subset X$ be the minimal set which is invariant for $\tilde{\mu}$ such that for all $x \in X \setminus \mathcal{L}$ the inequality

$$\widetilde{V}(f(x, \tilde{\mu}(x))) \leq \widetilde{V}(x)$$

holds and let

$$\sigma := \inf\{\widetilde{V}(f(x, \tilde{\mu}(x))) \,|\, x \in X \setminus \mathcal{L}\}.$$

Consider the modified running cost

$$\bar{l}(x, u) = \begin{cases} \max\{l(x, u) - \varepsilon, 0\}, & x \notin \mathcal{L} \\ 0, & x \in \mathcal{L} \end{cases},$$

the corresponding infinite horizon optimal value function $\overline{V}_\infty$ and the corresponding functional $\overline{V}_\infty^{\tilde{\mu}}$ using the controller $\tilde{\mu}$.

Then for all $x \in X$ the estimate

$$\alpha \overline{V}_\infty(x) \leq \alpha \overline{V}_\infty^{\tilde{\mu}}(x) \leq \widetilde{V}(x) - \sigma$$

holds.

Proof: Consider $x_0 \in X$, the trajectory $x(n) = x_{\tilde{\mu}}(n)$ generated by the closed loop system using $\tilde{\mu}$, and the control sequence generated by $u(n) = \tilde{\mu}(x(n))$. Let $n_0 \in \mathbb{N}_0$ be minimal with $x(n_0+1) \in \mathcal{L}$ setting $n_0 = \infty$ if this never happens. Then, from the definition of $\bar{l}$ and (5.1) for $n = 1, \ldots, n_0$ we obtain

$$\alpha \bar{l}(x(n), u(n)) = \max\{\alpha l(x(n), u(n)) - \alpha\varepsilon, 0\} \le \widetilde{V}(x(n)) - \widetilde{V}(x(n+1)),$$

where we have used in the inequality that the right hand side is nonnegative because of $x(n) \notin \mathcal{L}$ and the definition of $\mathcal{L}$. For $n \ge n_0 + 1$ the invariance of $\mathcal{L}$ implies $x(n) \in \mathcal{L}$, thus $\bar{l}(x(n), u(n)) = 0$. Hence, for any $K \in \mathbb{N}$ and $K_0 = \min\{K, n_0\}$ summing over n yields

$$\alpha \sum_{n=0}^{K} \bar{l}(x(n), u(n)) = \alpha \sum_{n=0}^{K_0} \bar{l}(x(n), u(n)) \le \widetilde{V}(x(0)) - \widetilde{V}(x(K_0+1)) \le \widetilde{V}(x(0)) - \sigma.$$

Since $K \in \mathbb{N}$ was arbitrary, $\widetilde{V} - \sigma$ is an upper bound for $\overline{V}_\infty^{\tilde{\mu}}$ and we obtain the assertion. ∎

Remark 5.2 The definition of $\mathcal{L}$ in this proposition is implicit and the precise shape of $\mathcal{L}$ can not be deduced from (5.1). However, we can obtain an estimate for $\mathcal{L}$ from (5.1) by defining

$$\eta := \sup\{\widetilde{V}(x) \,|\, x \in X, \min_{u \in U} l(x,u) \le \varepsilon\} + \varepsilon.$$

Then the inclusion

$$\mathcal{L} \subseteq \mathcal{V} := \{x \in X \,|\, \widetilde{V}(x) \le \eta\}.$$

holds.

In order to prove this claim it is sufficient to show that $\mathcal{V}$ is invariant under $\tilde{\mu}$ and

$$\widetilde{V}(f(x, \tilde{\mu}(x))) \le \widetilde{V}(x)$$

for all $x \in X \setminus \mathcal{V}$. The stated inequality follows immediately from (5.1) because outside $\mathcal{V}$ the inequality $l(x, \tilde{\mu}(x)) - \varepsilon \ge 0$ holds. In order to show invariance of $\mathcal{V}$ under $\tilde{\mu}$, i.e., $f(x, \tilde{\mu}(x)) \in \mathcal{V}$ for $x \in \mathcal{V}$, we distinguish two cases:

Case 1: $\widetilde{V}(x) \le \eta - \varepsilon$. In this case (5.1) and $l(x, \tilde{\mu}(x)) \ge 0$ yields

$$\widetilde{V}(f(x, \tilde{\mu}(x))) \le \widetilde{V}(x) - \max\{\alpha(l(x, \tilde{\mu}(x)) - \varepsilon),\, l(x, \tilde{\mu}(x)) - \varepsilon\} \le \eta - \varepsilon + \varepsilon = \eta,$$

implying $f(x, \tilde{\mu}(x)) \in \mathcal{V}$

Case 2: $\widetilde{V}(x) \in (\eta - \varepsilon, \eta]$. In this case the definition of η yields $l(x, \tilde{\mu}(x)) \ge \varepsilon$, hence the maximum in (5.1) is attained in the first term and thus we obtain

$$\widetilde{V}(f(x, \tilde{\mu}(x))) \le \widetilde{V}(x) - \alpha(l(x, \tilde{\mu}(x)) + \varepsilon) \le \widetilde{V}(x) \le \eta,$$

again implying $f(x, \tilde{\mu}(x)) \in \mathcal{V}$.

The reason for formulating Proposition 5.1 using the implicitly defined set $\mathcal{L}$ lies in the fact that the estimate via $\mathcal{V}$ may be conservative and using $\mathcal{L}$ in general yields a sharper estimate.

We would also like to point out that there are other methods for estimating $\mathcal{L}$, like, e.g., the techniques developed in [5]. □

The next lemma is a practical version of Lemma 4.1.

Lemma 5.3 Consider $N \in \mathbb{N}$ and the receding horizon feedback law μ_N. Assume that

$$V_N(f(x,\mu_N(x))) - V_{N-1}(f(x,\mu_N(x))) \le \max\{(1-\alpha)l(x,\mu_N(x)) + \alpha\varepsilon,\, \varepsilon\} \qquad (5.2)$$

holds for some $\alpha \in [0,1]$ and all $x \in X$. Then $\widetilde{V} = V_N$ satisfies (5.1) and, in particular,

$$\alpha \overline{V}_\infty^{\mu_N}(x) \le V_N(x) - \sigma$$

holds for all $x \in X$ using the notation from Proposition 5.1.

Proof: Combining (2.7) and (2.8) and inserting (5.2) yields

$$\begin{aligned} V_N(x) &= \min_u \{V_{N-1}(f(x,u)) + l(x,u)\} = V_{N-1}(f(x,\mu_N(x))) + l(x,\mu_N(x)) \\ &\ge V_N(f(x,\mu_N(x))) - \max\{(1-\alpha)l(x,\mu_N(x)) + \alpha\varepsilon,\, \varepsilon\} + l(x,\mu_N(x)) \\ &= V_N(f(x,\mu_N(x))) + \min\{\alpha l(x,\mu_N(x)) - \alpha\varepsilon,\, l(x,\mu_N(x)) - \varepsilon\}, \end{aligned}$$

which shows (5.1). Now Proposition 5.1 yields the assertion. □

Having derived the appropriate practical versions of our preliminary results we can now turn to the formulation of a practical version of Assumption 4.2.

Assumption 5.4 For given $N \in \mathbb{N}$ there exists $\gamma > 0$ and $\varepsilon > 0$ such that the inequalities

$$V_2(x) \le \max\{V_1(x) + \varepsilon,\, (\gamma+1)V_1 + (1-\gamma)\varepsilon\}$$

and

$$V_k(x) \le \max\{l(x,\mu_k(x)) + (k-1)\varepsilon,\, (\gamma+1)l(x,\mu_k(x)) + (k-1-\gamma)\varepsilon\}, \quad k = 2,\dots,N$$

hold for all $x \in X$. □

Remark 5.5 Essentially, this assumption relaxes Assumption 4.2 by requiring (suitably shifted versions of) the inequalities from Assumption 4.2 only in those parts of the state space where $V_1(x)$ or $l(x,\mu_k(x))$ are relatively large compared to ε, while on the rest of the state space we only require a weaker bound which does not scale linearly with $V_1(x)$ or $l(x,\mu_k(x))$. An important feature of this assumption is that it also applies if the V_k grow unboundedly in k on the whole state space X, cf. also Remark 5.11. Furthermore, in practical examples there may be a tradeoff between ε and γ. Example 5.10, below, illustrates this situation. □

The next proposition is the modification of Proposition 4.4 for Assumption 5.4.

Proposition 5.6 Let $N \in \mathbb{N}$ and assume that Assumption 5.4 holds for this N. Then the inequality

$$\min\left\{ \frac{(\gamma+1)^{N-2}}{(\gamma+1)^{N-2} + \gamma^{N-1}}(V_N(x) - N\varepsilon),\;\; V_N(x) - N\varepsilon \right\} \le V_{N-1}(x) - (N-1)\varepsilon$$

holds for all $x \in X$.

Proof: Consider the optimal control problem with running cost $\tilde{l}$ defined by

$$\tilde{l}(x,u) = l(x,u) - \varepsilon.$$

The corresponding optimal value functions $\widetilde{V}_k$ satisfy

$$\widetilde{V}_k(x) = V_k(x) - k\varepsilon,$$

and the optimal controls for $\widetilde{V}_k$ and V_k coincide. This implies

$$\begin{aligned}\widetilde{V}_k(x) &\leq \max\{l(x,\mu_k(x)) + (k-1)\varepsilon,\, (\gamma+1)l(x,\mu_k(x)) + (k-1-\gamma)\varepsilon\} - k\varepsilon \\ &= \max\{l(x,\mu_k(x)) - \varepsilon,\, (\gamma+1)l(x,\mu_k(x)) - (\gamma+1)\varepsilon\} \\ &= \max\{\tilde{l}(x,\mu_k(x)),\, (\gamma+1)\tilde{l}(x,\mu_k(x))\}\end{aligned}$$

and similarly

$$\widetilde{V}_2(x) \leq \max\{\widetilde{V}_1(x),\, (\gamma+1)\widetilde{V}_1(x)\}.$$

Now we can proceed inductively as in the proof of Proposition 4.4 (note that nonnegativity of V_k and l is not needed in this induction) using either γ from Assumption 5.4 or $\gamma = 1$ in each induction step. This proves the estimate

$$\min\{\eta_k \widetilde{V}_k(x),\, \widetilde{V}_k\} \leq \widetilde{V}_{k-1}(x) \tag{5.3}$$

for η_k as in the proof of Proposition 4.4. Translated back to V_k, this yields the asserted inequality. □

Remark 5.7 Note that the inequality from Proposition 5.6 implies the more explicit estimate

$$\frac{(\gamma+1)^{N-2}}{(\gamma+1)^{N-2} + \gamma^{N-1}} V_N(x) \leq V_{N-1}(x) + \varepsilon,$$

however, for our subsequent calculations the sharper inequality stated in Proposition 5.6 is more suitable. □

Finally, we can state the practical version of Theorem 4.5.

Theorem 5.8 Consider $\gamma > 0$ and $N \in \mathbb{N}$ such that $(\gamma+1)^{N-2} > \gamma^N$ holds. Assume that Assumption 5.4 holds for these γ and N and some $\varepsilon > 0$. Then

$$\overline{V}_\infty^{\mu_N}(x) \leq \frac{(\gamma+1)^{N-2}}{(\gamma+1)^{N-2} - \gamma^N}(V_N(x) - \sigma) \leq \frac{(\gamma+1)^{N-2}}{(\gamma+1)^{N-2} - \gamma^N}(V_\infty(x) - \sigma)$$

using the notation from Proposition 5.1.

Proof: From Proposition 5.6 we obtain the inequality

$$\min\left\{\frac{(\gamma+1)^{N-2}}{(\gamma+1)^{N-2} + \gamma^{N-1}}(V_N(x) - N\varepsilon),\, V_N(x) - N\varepsilon\right\} \leq V_{N-1}(x) - (N-1)\varepsilon$$

which implies

$$V_N(x) - V_{N-1}(x) - \varepsilon \le \max\left\{\frac{\gamma^{N-1}}{(\gamma+1)^{N-2}}(V_{N-1}(x) - (N-1)\varepsilon),\, 0\right\}. \tag{5.4}$$

Analogous to (4.3), from Assumption 5.4 for $k = N$ we obtain

$$V_{N-1}(f(x,\mu_N(x)) - (N-1)\varepsilon \le \max\{\gamma(l(x,\mu_N(x)) - \varepsilon),\, 0\}.$$

Combining this with (5.4) for $x = f(x,\mu_N(x))$ yields

$$V_N(f(x,\mu_N(x)) - V_{N-1}(f(x,\mu_N(x)) - \varepsilon \le \max\left\{\frac{\gamma^N}{(\gamma+1)^{N-2}}(l(x,\mu_N(x)) - \varepsilon),\, 0\right\}.$$

Hence, we obtain (5.2) with

$$\alpha = 1 - \frac{\gamma^N}{(\gamma+1)^{N-2}} = \frac{(\gamma+1)^{N-2} - \gamma^N}{(\gamma+1)^{N-2}}$$

which using Lemma 5.3 implies

$$\overline{V}_\infty^{\mu_N}(x) \le \frac{1}{\alpha}(V_N(x) - \sigma) \le \frac{1}{\alpha}(V_\infty(x) - \sigma).$$

□

Remark 5.9 It should be noted that in the motivating example at the beginning of this section V_∞ is unbounded, in which case the final inequality in Theorem 5.8 is useless. □

We illustrate Theorem 5.8 by a simple example.

Example 5.10 Consider the discrete time system

$$x(n+1) = x(n) + u(n)h_\beta(x(n)) =: f(x(n), u(n))$$

with

$$h_\beta(x) = \max\{x - \beta,\, -\beta - x,\, 0\}$$

for some $\beta > 0$, $X = \mathbb{R}$ and $U = [-1, 1]$ (this discrete time system mimics the behavior of the radial component of the 2d sampled–data system discussed in [6, Section 9.4]).

Using the feedback law $u(x) = \mathrm{sgn}(x)$ one easily sees that the set $A_\beta = [-\beta, \beta]$ is asymptotically (even finite time) stable, however, it is not possible to steer the system into the interior of this set.

We choose the running cost $l(x,u) = |x|(1 + |u|)$ and fix $\varepsilon \ge \beta$. Using the control sequence $u = (0, 0, 0, \ldots)$, for $|x| \le \varepsilon$ one obtains

$$V_k(x) \le k|x| \le V_1(x) + (k-1)\varepsilon \le l(x,u) + (k-1)\varepsilon$$

which yields Assumption 5.4 for arbitrary $\gamma \ge 1$.

For $x > \varepsilon$, using the control sequence $u = (-(x-\varepsilon)/(x-\beta), 0, 0, \dots)$ we obtain

$$x(n) = \varepsilon$$

for all $n \geq 1$ and thus

$$V_k(x) \leq x\left(1 + \frac{x-\varepsilon}{x-\beta}\right) + (k-1)\varepsilon = \left(1 + \frac{x}{x-\beta}\right)x + \left(k - 1 - \frac{x}{x-\beta}\right)\varepsilon. \qquad (5.5)$$

Since $x \leq l(x,u)$ for all $u \in U$, for $\varepsilon > \beta$ this implies Assumption 5.4 for

$$\gamma = \sup_{x \geq \varepsilon} \frac{x}{x-\beta} = \frac{\varepsilon}{\varepsilon - \beta}.$$

For symmetry reasons, the same estimate holds for $x < -\varepsilon$. Thus, Assumption 5.4 holds for all $x \in X$ with $\gamma = \varepsilon/(\varepsilon-\beta)$. Note that $\varepsilon \searrow \beta$ results in $\gamma \to \infty$, i.e., there is a tradeoff between ε and γ.

Using the fact that no trajectory from outside A_β can be controlled into the interior of A_β we immediately obtain $A_\beta \subseteq \mathcal{L}$ and $\sigma \geq N\beta^2$. Thus, Theorem 5.8 yields the estimate

$$\overline{V}_\infty^{\mu_N}(x) \leq \frac{\left(1+\frac{\varepsilon}{\varepsilon-\beta}\right)^{N-2}}{\left(1+\frac{\varepsilon}{\varepsilon-\beta}\right)^{N-2} - \left(\frac{\varepsilon}{\varepsilon-\beta}\right)^{N}}(V_N(x) - \sigma) \leq \frac{\left(1+\frac{\varepsilon}{\varepsilon-\beta}\right)^{N-2}}{\left(1+\frac{\varepsilon}{\varepsilon-\beta}\right)^{N-2} - \left(\frac{\varepsilon}{\varepsilon-\beta}\right)^{N}}2x$$

for the modified cost function $\bar{l}$ from Proposition 5.1, provided N is so large that the denominator of the fraction is positive. Here we have used the first inequality in (5.5) for $k = N$ and $\varepsilon = \beta$ for estimating V_N in the last step.

It remains to estimate the stabilized set $\mathcal{L}$. Here we proceed similar to Remark 5.2 exploiting, however, the symmetry of the problem: since the problem is symmetric with respect to the origin and $l(x,u)$ is monotonically increasing with respect to $|x|$ we obtain that V_N is symmetric and monotonically increasing in $|x|$, too. From this it is easy to conclude that the set

$$\{x \in X, \min_{u \in U} l(x,u) \leq \varepsilon\} = [-\varepsilon, \varepsilon]$$

is a sublevel set for each V_N which implies that it is forward invariant. Thus, we do not have to construct the forward invariant set $\mathcal{V}$ from Remark 5.2 but can conclude directly that the receding horizon controller will stabilize the set $[-\varepsilon, \varepsilon]$. Since for $N \to \infty$ we can choose ε arbitrarily close to β if N, this stabilized set converges to A_β as $N \to \infty$. □

Remark 5.11 In some references, an inequality of the form

$$V_k(x) \leq \Phi(x)$$

for some function $\Phi : X \to \mathbb{R}_0^+$, all $k \in \mathbb{N}$ and all $x \in X$ is imposed in order to conclude stability or practical stability of the RHC closed loop, cf. e.g. [5, Standing Assumption 4]. While our Assumption 4.2 fits into this framework, Assumption 5.4 is more general, since in the example above we obtain global practical asymptotic stability of the set $[-\beta, \beta]$ even though $V_k(x) \geq k\beta^2$ holds for all $k \in \mathbb{N}$ and all $x \in \mathbb{R}$ with $|x| \geq \beta$. □

6 Results for terminal costs being Lyapunov functions

Many RHC schemes make use of a suitable terminal cost in order to ensure closed loop asymptotic stability of the RHC controller. Often, in these settings the terminal costs are chosen as Lyapunov functions with respect to the running cost l, see [16] and the references therein. In this section we discuss the consequences on suboptimality of these choices. Here we make the following assumption on the terminal cost F.

Assumption 6.1 For a neighborhood Y_0 of 0 and each $x \in Y_0$ there exists $u \in U$ such that

$$f(x,u) \in Y_0 \quad \text{and} \quad F(f(x,u)) \leq F(x) - l(x,u).$$

□

This condition is often imposed in receding horizon schemes in order to ensure asymptotic stability of the closed loop, see [10, Section II] or [16, Section 3.3 and the references therein]. Note that Assumption 6.1 implies (2.10) for $\widetilde{V} = F$ with $\alpha = 1$. Hence, Proposition 2.2 implies $F(x) \geq V_\infty(x)$ on Y_0 and we can define the positive difference $\eta := \sup_{x\in X} F(x) - V_\infty(x) \geq 0$. Observe that in order to simplify the presentation we assume F to be defined on the whole state space X. If F is only defined locally then, for instance, one could choose Y_0 as a sublevel set of F and extend F continuously outside Y_0 by a constant function, cf. [10, Formula (8)].

A typical situation in which F meeting Assumption 6.1 can be found is if the linearization of f is controllable to 0 and l is close to a quadratic function around the origin. In this case, F can be chosen as the optimal value function of the linear quadratic problem for a quadratic cost function $\tilde{l}$ which is strictly smaller than l. Then, the closer l and $\tilde{l}$ are and the smaller the neighborhood Y_0 is chosen, the smaller η becomes, see also the discussion after Lemma 3 in [12].

In the following theorem we distinguish the case with and without terminal constraint set.

Theorem 6.2 Assume that the terminal cost in (2.5) satisfies Assumption 6.1 on some neighborhood Y_0 of the origin. Define $\eta := \sup_{x\in X} F(x) - V_\infty(x)$ and let $N \in \mathbb{N}$.

(i) Consider the optimal receding horizon controller μ_N from (2.8) based on V_N from (2.4), i.e., without terminal constraint. Let $Y_N \subset X$ be the set of initial values for which the optimal solution $x(n)$ for the finite horizon functional (2.5) satisfies $x(N) \in Y_0$. Then the inequality

$$V_\infty^{\mu_N}(x) \leq V_N(x) \leq V_\infty(x) + \eta$$

holds for each $x \in Y_N$.

(ii) Consider the optimal receding horizon controller μ_N from (2.8) based on V_N from (2.6) with terminal constraint set $X_f = Y_0$. Then the inequality

$$V_\infty^{\mu_N}(x) \leq V_N(x)$$

holds on the feasible set X_N. Let, furthermore, $Y_\infty^N \subset X_N$ be the set of initial values for which the optimal solution $x(n)$ for the infinite horizon functional (2.2) satisfies $x(N) \in Y_0$. Then the inequality

$$V_N(x) \leq V_\infty(x) + \eta$$

holds for each $x \in Y_\infty^N$.

Proof: (i) For $x \in Y_N$ abbreviate $x^+ = f(x, \mu_N(x))$. Then, from the optimality principle we obtain $x^+ \in Y_{N-1}$. Now consider an optimal control sequence $u_{N-1} \in \mathcal{U}$ for the problem (2.4) with horizon length $N-1$ and the corresponding trajectory $x_{u_{N-1}}$ with initial value $x_{u_{N-1}}(0) = x^+$. Since $x^+ \in Y_{N-1}$ we obtain $\bar{x} := x_{u_{N-1}}(N-1) \in Y_0$. Let $\bar{u}$ denote the control value from Assumption 6.1 for $\bar{x}$ and define a control sequence $\tilde{u} = (u_{N-1}(0), \ldots, u_{N-1}(N-1), \bar{u}, \ldots)$. This sequence yields

$$V_N(x^+) \le J_N(x^+, \tilde{u}) = V_{N-1}(x^+) - F(\bar{x}) + l(\bar{x}, \bar{u}) + F(f(\bar{x}, \bar{u})) \le V_{N-1}(x^+).$$

Thus, (4.2) follows with $\alpha = 1$ which implies

$$V_\infty^{\mu_N}(x) \le V_N(x).$$

The inequality $V_N(x) \le V_\infty(x) + \eta$ follows immediately from the definition of J_N and J_∞ and $F \ge V_\infty$, which was observed in the discussion after Assumption 6.1.

(ii) The inequality $V_\infty^{\mu_N}(x) \le V_N(x)$ is concluded as in (i). The second inequality again follows from the definition of J_N and J_∞ and $F \ge V_\infty$, observing that for $x \in Y_\infty^N$ the optimal control sequence u for (2.2) satisfies the constraint in (2.6). □

Remark 6.3 (i) Note for terminal cost functions which are continuous and constant outside Y_0, Theorem 6.2(i) remains true under the terminal state constraint $X_f = Y_0$. This follows from [10, Theorem 1] where it is shown that Y_N is forward invariant under the RHC feedback μ_N. For this specific structure of the terminal cost the proof of this theorem can also be used as an alternative proof of the first inequality in Theorem 6.2(i), because there $V_N \ge V_N(x^+) + l(x, \mu_N(x))$ is shown for $x \in Y_N$. Summing up this inequality again yields the desired inequality.

(ii) For the particular case $F = V_\infty$ we get $\eta = 0$ and thus Theorem 6.2 yields infinite horizon optimality of the receding horizon controller on Y_N or Y_∞^N, respectively. Thus we recover the result from [10, Theorem 2]. Our more general result proves the conjecture posed at the end of the introduction of [10], i.e., that one obtains suboptimality with guaranteed error bounds when F approximates V_∞. □

The last result we are going to present shows how Theorem 6.2(i) may be combined with the results from Section 4. For the purpose of brevity we restrict ourselves to Assumption 4.2.

The motivation for such a combination stems from the fact that Assumption 4.2 may hold away from the origin but may fail to hold in a neighborhood of the origin. Then, on this neighborhood, a local Lyapunov function terminal cost may be used in order to stabilize the RHC scheme. Since this construction works without imposing terminal constraints, we can combine local stability with global operating region because we do not need to consider any feasible sets.

Theorem 6.4 Assume that the terminal cost in (2.5) satisfies Assumption 6.1 on some neighborhood Y_0 of the origin. Define $\eta := \sup_{x \in X} F(x) - V_\infty(x)$ and let $N \in \mathbb{N}$.

Consider the optimal receding horizon controller μ_N from (2.8) based on V_N from (2.4), i.e., without terminal constraint. Let $Y_k \subset X$ be the set of initial values for which the optimal solution $x(n)$ for the finite horizon functional (2.5) with $N = k$ satisfies $x(k) \in Y_0$ and assume that Assumption 4.2 holds for some $\gamma > 0$ and each $k = 1, \dots, N$ on $X \setminus Y_k$.

Then the inequality

$$V_\infty^{\mu_N}(x) \le \frac{(\gamma+1)^{N-2}}{(\gamma+1)^{N-2} - \gamma^N} V_N(x) \le \frac{(\gamma+1)^{N-2}}{(\gamma+1)^{N-2} - \gamma^N} (V_\infty(x) + \eta)$$

holds for each $x \in X$.

Proof: For $x \in Y_k$ the proof of Theorem 6.2(i) yields

$$V_k(x^+) \le V_{k-1}(x^+)$$

for $x^+ = f(x, \mu_k(x))$. This implies

$$V_k(x) = V_{k-1}(x^+) + l(x, \mu_k(x)) \ge V_k(x^+) + l(x, \mu_k(x)) \ge V_{k+1}(x), \quad x \in Y_k. \tag{6.1}$$

For $x \notin Y_k$ we can proceed as in the proof of Proposition 4.4 in order to conclude

$$V_k(x) \ge \frac{(\gamma+1)^{k-1}}{(\gamma+1)^{k-1} + \gamma^k} V_{k+1}(x), \quad x \in X \setminus Y_k. \tag{6.2}$$

Combining (6.1) and (6.2) yields the inequality in (6.2) for all $x \in X$. From this the first inequality of the assertion follows similarly to the proof of Theorem 4.5 while the second follows as in the proof of Theorem 6.2(i). □

7 Conclusions

We have derived rigorous suboptimality estimates for the infinite horizon performance of RHC controllers. In particular, we have shown that suitable exponential controllability assumptions for the running cost allow for obtaining suboptimality estimates for RHC schemes without terminal cost and constraints, a setting which to the best of our knowledge is not covered by the existing inverse optimality results. These results are complemented by novel estimates for the case where the RHC terminal cost is a Lyapunov functions, which is the classical setting for inverse optimality results. In both cases, techniques from relaxed dynamic programming are the main tool for establishing our results.

Acknowledgement: We thank the anonymous referees for their constructive remarks and for bringing to our attention relevant references to the literature. Furthermore, we are grateful to Jürgen Pannek for reporting a subtle error in an earlier version of the paper.

References

[1] F. Allgöwer and A. Zheng, eds., *Nonlinear model predictive control*, vol. 26 of Progress in Systems and Control Theory, Birkhäuser Verlag, Basel, 2000.

[2] R. R. Bitmead, M. Gevers, and V. Wertz, *Adaptive optimal control. The thinking man's GPC*, Prentice-Hall, New York etc., 1990.

[3] E. F. Camacho and C. Bordons, *Model predictive control*, Springer–Verlag, London, 2nd ed., 2004.

[4] F. Di Palma and L. Magni, *On optimality of nonlinear model predictive control*, in Proceedings of the 16th International Symposium on Mathematical Theory of Networks and Systems, Leuven, Belgium, 2004. CD–ROM, Session MP8, Paper No. 238.

[5] G. Grimm, M. J. Messina, S. E. Tuna, and A. R. Teel, *Model predictive control: for want of a local control Lyapunov function, all is not lost*, IEEE Trans. Automat. Control, 50 (2005), pp. 546–558.

[6] L. Grüne, *Stabilization by sampled and discrete feedback with positive sampling rate*, in Stability and Stabilization of Nonlinear Systems, Proceedings of the 1st NCN Workshop, D. Ayels, F. Lamnabhi-Lagarrigue, and A. van der Schaft, eds., Lecture Notes in Control and Information Sciences 246, Springer-Verlag, London, 1999, pp. 165–182.

[7] L. Grüne, D. Nešić, and J. Pannek, *Model predictive sampled–data redesign for nonlinear systems*, in Proceedings of the 44th IEEE Conference on Decision and Control and European Control Conference, Sevilla, Spain, 2005, pp. 36–41.

[8] L. Grüne, D. Nešić, and J. Pannek, *Model predictive control for nonlinear sampled–data systems*, in Assessment and Future Directions of Nonlinear Model Predictive Control, F. Allgöwer, L. Biegler, and R. Findeisen, eds., Springer Lecture Notes in Control and Information Sciences, 2006. To appear.

[9] L. Grüne and A. Rantzer, *Suboptimality estimates for receding horizon controllers*, in Proceedings of the 17th International Symposium on Mathematical Theory of Networks and Systems MTNS2006, Kyoto, Japan, 2006, pp. 120–127.

[10] B. Hu and A. Linnemann, *Toward infinite-horizon optimality in nonlinear model predictive control*, IEEE Trans. Automat. Control, 47 (2002), pp. 679–682.

[11] K. Ito and K. Kunisch, *Receding horizon optimal control for infinite dimensional systems*, ESAIM Control Optim. Calc. Var., 8 (2002), pp. 741–760. A tribute to J. L. Lions.

[12] A. Jadbabaie and J. Hauser, *On the stability of receding horizon control with a general terminal cost*, IEEE Trans. Automat. Control, 50 (2005), pp. 674–678.

[13] S. S. Keerthy and E. G. Gilbert, *Optimal infinite horizon feedback laws for a general class of constrained discrete-time systems: stability and moving horizon approximations*, J. Optimiz. Theory Appl., 57 (1988), pp. 265–293.

[14] B. Lincoln and A. Rantzer, *Relaxing dynamic programming*, IEEE Trans. Autom. Control, 51 (2006), pp. 1249–1260.

[15] L. Magni and R. Sepulchre, *Stability margins of nonlinear receding-horizon control via inverse optimality*, Systems Control Lett., 32 (1997), pp. 241–245.

[16] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, *Constrained model predictive control: stability and optimality*, Automatica, 36 (2000), pp. 789–814.

[17] J. A. Primbs and V. Nevistić, *Feasibility and stability of constrained finite receding horizon control*, Automatica, 36 (2000), pp. 965–971.

[18] A. Rantzer, *Relaxed dynamic programming in switching systems*, IEE Proceedings — Control Theory and Applications, 153 (2006), pp. 567–574.

[19] J. S. Shamma and D. Xiong, *Linear nonquadratic optimal control*, IEEE Trans. Autom. Control, 42 (1997), pp. 875–879.