

Optimization based stabilization of nonlinear control systems

Lars Grüne

Mathematical Institute, University of Bayreuth, 95440 Bayreuth, Germany

`lars.gruene@uni-bayreuth.de`

WWW home page: <http://www.math.uni-bayreuth.de/~lgruene/>

Abstract. We present a general framework for analysis and design of optimization based numerical feedback stabilization schemes utilizing ideas from relaxed dynamic programming. The application of the framework is illustrated for a set valued and graph theoretic offline optimization algorithm and for receding horizon online optimization.

1 Introduction

The design of feedback controls for nonlinear systems is one of the basic problem classes in mathematical control theory. Among the variety of different design objectives within this class, the design of stabilizing feedbacks is an important subproblem, on the one hand because it captures the essential difficulties and on the other hand because it often occurs in engineering practice. The rapid development of numerical methods for optimization and optimal control which has led to highly efficient algorithms which are applicable even to large scale nonlinear systems, naturally leads to the idea of using such algorithms in feedback stabilization. While the basic principles underlying the relation between optimal control and stabilization have been understood since the 1960s in the context of the linear quadratic regulator design, the application to nonlinear systems poses new problems and challenges. These are caused, for instance, due to the complicated dynamical behavior which even low dimensional nonlinear systems can exhibit, due to hybrid or switching structures incorporated in the dynamics or simply due to the size of the problems to be solved, e.g. when discretized PDEs are to be controlled.

In this paper we investigate the foundation of optimization based feedback stabilization for nonlinear systems and develop a framework which allows to design and analyze different numerical approaches. The main goal of our framework is to give rigorous mathematical stability proofs even in the case when the numerical approximation is inaccurate, in the sense that the optimal control problem which should be solved theoretically is only very coarsely approximated by our numerical algorithm. Our approach was motivated and inspired by two sources: on the one hand by classical Lyapunov function stability theory, and in this context our condition can be understood as a preservation of the Lyapunov function property under (not necessarily small) numerical approximation

errors. On the other hand we make use of relaxed dynamic programming methods, which essentially ensure that even in the presence of errors we can still give precise bounds for the performance of the feedback controller derived from the numerical information.

The organization of this paper is as follows. After describing the setup and the (theoretical) relation between optimization and stabilization in Section 2, we develop our conditions for optimization based stabilization with coarse numerical approximations in Section 3. In Section 4 we apply these conditions to a graph theoretic offline optimization approach while in Section 5 we show how these conditions can contribute to the analysis and design of unconstrained receding horizon control schemes.

2 Setup and preliminaries

We consider a nonlinear discrete time 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 $n \in \mathbb{N}_0$. We denote the space of control sequences $u : \mathbb{N}_0 \rightarrow U$ 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, i.e., it can range from a finite set to an infinite dimensional space.

A typical class of systems we consider are sampled-data systems governed by a controlled differential equation $\dot{x}(t) = g(x(t), \tilde{u}(t))$ with solution $\varphi(t, x_0, \tilde{u})$ for initial value x_0 . These are obtained by fixing a sampling period $T > 0$ and setting

$$f(x, u) := \varphi(T, x, \tilde{u}) \quad \text{with} \quad \tilde{u}(t) \equiv u. \quad (2)$$

Then, for any discrete time control function $u \in \mathcal{U}$ the solutions x_u of (1),(2) satisfy $x_u(n) = \varphi(nT, x_0, \tilde{u})$ for the piecewise constant continuous time control function $\tilde{u} : \mathbb{R} \rightarrow U$ with $\tilde{u}|_{[nT, (n+1)T)} \equiv u(n)$. Note that with this construction the discrete time n corresponds to the continuous time $t = nT$.

2.1 Infinite horizon optimal control

Our goal is to find a feedback control law minimizing the infinite horizon cost

$$J_\infty(x_0, u) = \sum_{n=0}^{\infty} l(x_u(n), u(n)), \quad (3)$$

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

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

Here a (static state) feedback law is a control law $F : X \rightarrow U$ which assigns a control value u to each state x and which is applied to the system according to the rule

$$x_F(n+1) = f(x_F(n), F(x_F(n))), \quad x_F(0) = x_0. \quad (4)$$

From dynamic programming theory (cf. e.g. [1]) it is well known that the optimal value function satisfies Bellman's optimality principle, i.e.,

$$V_\infty(x) = \min_{u \in U} \{l(x, u) + V_\infty(f(x, u))\} \quad (5)$$

and that the optimal feedback law F is given by

$$F(x) := \operatorname{argmin}_{u \in U} \{l(x, u) + V_\infty(f(x, u))\}. \quad (6)$$

Remark 1. In order to simplify and streamline the presentation, throughout this paper it is assumed that in all relevant expressions the minimum with respect to $u \in U^m$ is attained. Alternatively, modified statements using approximate minimizers could be used which would, however, considerably increase the amount of technicalities needed in order to formulate our assumptions and results.

2.2 Asymptotic feedback stabilization

Our main motivation for considering infinite horizon optimal control problems is the fact that these problems yield asymptotically stabilizing feedback laws. In order to make this statement precise, we first define what we mean by an asymptotically stabilizing feedback law.

Let us assume that the control system under consideration has an equilibrium $x^* \in X$ for some control $u^* \in U$, i.e.,

$$f(x^*, u^*) = x^*.$$

Asymptotic stability can be elegantly formulated using the concept of comparison functions. To this end, as usual in nonlinear stability theory, we define the class \mathcal{K} of continuous functions $\delta : \mathbb{R}_0^+ \rightarrow \mathbb{R}_0^+$ which are strictly increasing and satisfy $\delta(0) = 0$ and the class \mathcal{K}_∞ of functions $\delta \in \mathcal{K}$ which are unbounded and hence invertible with $\delta^{-1} \in \mathcal{K}_\infty$. We also define the (discrete time) class \mathcal{KL} of continuous functions $\beta : \mathbb{R}_0^+ \times \mathbb{N}_0 \rightarrow \mathbb{R}_0^+$ which are of class \mathcal{K} in the first argument and strictly decreasing to 0 in the second argument. Examples for $\beta \in \mathcal{KL}$ are, for instance,

$$\beta(r, n) = Ce^{-\sigma n r} \quad \text{or} \quad \beta(r, n) = \frac{C\sqrt{r}}{1+n}.$$

Then we say that a feedback law $F : X \rightarrow U$ asymptotically stabilizes the equilibrium x^* , if there exists $\beta \in \mathcal{KL}$ such that for all initial values $x_0 \in X$ the solution of (4) satisfies

$$\|x_F(n)\|_{x^*} \leq \beta(\|x_0\|_{x^*}, n) \quad (7)$$

using the brief notation $\|x\|_{x^*} = d(x, x^*)$ for the distance of a point $x \in X$ to the equilibrium x^* , where $d(\cdot, \cdot)$ is a metric on X . Note that $\|\cdot\|_{x^*}$ does not need to be a norm. In less formal words, Condition (7) demands that, by virtue of β

being of class \mathcal{K} in its first argument, any solution starting close to x^* remains close to x^* for all future times and that, since β is decreasing to 0 in its second argument, any solution converges to x^* as $n \rightarrow \infty$. This \mathcal{KL} characterization of asymptotic stability is actually equivalent to the ε - δ formulation often found in the literature.

In order to obtain an asymptotically stabilizing optimal feedback F from (6) we proceed as follows: For the running cost l we define

$$l^*(x) := \inf_{u \in U} l(x, u)$$

and choose l in such a way that there exist $\gamma_1 \in \mathcal{K}_\infty$ satisfying

$$\gamma_1(\|x\|_{x^*}) \leq l^*(x). \quad (8)$$

Then, if an asymptotically stabilizing feedback law exists, under suitable boundedness conditions on l (for details see [2, Theorem 5.4]; see also [3] for a treatment in continuous time) there exist $\delta_1, \delta_2 \in \mathcal{K}_\infty$ such that the inequality

$$\delta_1(\|x\|_{x^*}) \leq V_\infty(x) \leq \delta_2(\|x\|_{x^*}) \quad (9)$$

holds. Furthermore, from the optimality principle we can deduce the inequality

$$V_\infty(f(x, F(x))) \leq V_\infty(x) - l(x, F(x)) \leq V_\infty(x) - l^*(x) \leq V_\infty(x) - \gamma_1(\|x\|_{x^*}).$$

Since $V_\infty(x) \leq \delta_2(\|x\|_{x^*})$ implies $\|x\|_{x^*} \geq \delta_2^{-1}(V_\infty(x))$ we obtain

$$\gamma_1(\|x\|_{x^*}) \geq \gamma_1(\delta_2^{-1}(V_\infty(x)))$$

and thus

$$V_\infty(f(x, F(x))) \leq V_\infty(x) - \gamma_1(\delta_2^{-1}(V_\infty(x))).$$

For the solution $x_F(n)$ from (4) this implies

$$V_\infty(x_F(n)) \leq \sigma(V_\infty(x_0), n)$$

for some suitable $\sigma \in \mathcal{KL}$. Thus, using (9), we eventually obtain

$$\|x_F(n)\|_{x^*} \leq \delta_1^{-1}(\sigma(\delta_2(\|x_0\|_{x^*}), n)) =: \beta(\|x_0\|_{x^*}, n).$$

Using the monotonicity of the involved \mathcal{K}_∞ functions it is an easy exercise to show that $\beta \in \mathcal{KL}$. This proves that the infinite horizon optimal feedback indeed asymptotically stabilizes the equilibrium x^* .

Essentially, this proof uses that V_∞ is a Lyapunov function for the closed loop system (4) controlled by the infinite horizon optimal feedback law F .

3 The relaxed optimality principle

The relation between asymptotic feedback stabilization and infinite horizon optimal control paves the way for applying powerful numerical algorithms from

the area of optimal control to the feedback stabilization problem. Thus, it is no surprise that in the literature one can find numerous approaches which attempt to proceed this way, i.e., find a numerical approximation $\tilde{V} \approx V_\infty$ and compute a numerical approximation \tilde{F} to the optimal feedback law using

$$\tilde{F}(x) := \operatorname{argmin}_{u \in U} \left\{ l(x, u) + \tilde{V}(f(x, u)) \right\}. \quad (10)$$

Examples for such approaches can be found, e.g., in [4–6] for general nonlinear systems and in [7, 8] for homogeneous systems; the approach is also closely related to semi-Lagrangian finite element discretization schemes for Hamilton-Jacobi PDEs, see, e.g., [9, Appendix 1], [10, 11]. All these schemes rely on the fact that the numerically computed function \tilde{V} closely approximates the true optimal value function V_∞ and typically fail in case of larger numerical errors.

Unfortunately, however, except for certain special situations like, e.g., unconstrained linear quadratic problems, even sophisticated numerical techniques can only yield good approximations $\tilde{V} \approx V_\infty$ in low dimensional state spaces X . Hence, in general it seems too demanding to expect a highly accurate numerical approximation to the optimal value function and thus we have to develop concepts which allow to prove stability of numerically generated feedback laws even for rather coarse numerical approximations \tilde{V} .

The main tool we are going to use for this purpose is a rather straightforward and easily proved “relaxed” version of the dynamic programming principle. This fact, which we are going to formalize in the following theorem, has been used implicitly in many papers on dynamic programming techniques during the last decades. Recently, it has been extensively studied and used by Lincoln and Rantzer in [12, 13].

In order to formulate the theorem we need to define the *infinite horizon value function* $V_\infty^{\tilde{F}}$ of a feedback law $\tilde{F} : X \rightarrow U$, which is given by

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

where $x_{\tilde{F}}$ is the solution from (4) with \tilde{F} instead of F .

Theorem 1. (i) Consider a feedback law $\tilde{F} : X \rightarrow U$ and a function $\tilde{V} : X \rightarrow \mathbb{R}_0^+$ satisfying the inequality

$$\tilde{V}(x) \geq \alpha l(x, \tilde{F}(x)) + \tilde{V}(f(x, \tilde{F}(x))) \quad (11)$$

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

$$\alpha V_\infty(x) \leq \alpha V_\infty^{\tilde{F}}(x) \leq \tilde{V}(x) \quad (12)$$

holds.

(ii) If, in addition, the inequalities (8) and (9) with \tilde{V} instead of V_∞ hold, then the feedback law \tilde{F} asymptotically stabilizes the system for all $x_0 \in X$.

Proof. (i) Using (11) for all $x = x_{\tilde{F}}(n)$, $n \in \mathbb{N}_0$ we obtain

$$\alpha l(x_{\tilde{F}}(n), \tilde{F}(x(n))) \leq \tilde{V}(x_{\tilde{F}}(n)) - \tilde{V}(x_{\tilde{F}}(n+1)).$$

Summing over n yields

$$\alpha \sum_{n=0}^m l(x_{\tilde{F}}(n), \tilde{F}(x_{\tilde{F}}(n))) \leq \tilde{V}(x(0)) - \tilde{V}(x(m)) \leq \tilde{V}(x(0)).$$

For $m \rightarrow \infty$ this yields that \tilde{V} is an upper bound for $\alpha V_{\infty}^{\tilde{F}}$ and hence (12), since the first inequality in (12) is obvious.

(ii) From (11) we immediately obtain

$$\tilde{V}(f(x, \tilde{F}(x))) \leq \tilde{V}(x) - \alpha l(x, \tilde{F}(x)).$$

Now we can proceed exactly as in Section 2.2 using $\alpha\gamma_1$ instead of γ_1 in order to conclude asymptotic stability.

The contribution of Theorem 1 is twofold: On the one hand, in (i) it gives an estimate for the infinite horizon value $V_{\infty}^{\tilde{F}}$ based on \tilde{V} and α , on the other hand, in (ii) it ensures that the corresponding feedback \tilde{F} is indeed asymptotically stabilizing.

We emphasize the fact that no relation between \tilde{V} and V_{∞} is needed in order to obtain these results. Hence, we can use this theorem even if \tilde{V} is only a very rough approximation to V_{∞} , provided, of course, that our numerical scheme is such that $\alpha \in (0, 1]$ satisfying (11) can be found.

In the following two sections we present two numerical approaches for which this is indeed possible. In Section 4 we discuss an offline optimization method particularly suitable for low dimensional systems, whose main advantages are the cheap online evaluation of the feedback and its capability to be easily extended to hybrid systems, i.e., systems with additional discrete states and switching rules. In this context we will see that a suitable extension of the basic algorithm results in a method for which the assumptions of Theorem 1 can be verified.

In Section 5 we investigate receding horizon control, an online optimization technique particularly suitable for smooth dynamics for which fast online optimization is possible even for large scale systems. For these schemes we will see that Theorem 1 induces conditions on the infinite horizon running cost l which can be used in order to considerably reduce the complexity of the online optimization problem.

4 A set oriented and graph theoretic approach

In this section we describe an offline optimization method which is based on a set oriented discretization method followed by graph theoretic optimization methods. Since the method is described in detail in [14], here we only sketch the main ideas and in particular the relevance of Theorem 1 in this context.

In order to apply the method, we assume that the state space X is a compact subset of a finite dimensional Euclidean space and consider a partition \mathcal{P} of X consisting of finitely many disjoint sets P , called cells or boxes. For each cell $P \in \mathcal{P}$ we define a map $F : \mathcal{P} \times U \rightarrow 2^{\mathcal{P}}$ by setting

$$F(P, u) := \{Q \in \mathcal{P} \mid f(x, u) \in Q \text{ for some } x \in P\}.$$

Furthermore, we fix a target region T consisting of cells in \mathcal{P} and containing the desired equilibrium x^* .

The basic idea of the approach as presented in [15] is to define a weighted directed graph $G = (\mathcal{P}, E, w)$ consisting of *nodes* \mathcal{P} , *edges* $E \subset \mathcal{P} \times \mathcal{P}$ and *weights* $w : E \rightarrow \mathbb{R}_0^+$ capturing the dynamics of F and the running cost l . This is accomplished by setting

$$E := \{(P, Q) \subset \mathcal{P} \times \mathcal{P} \mid Q \in F(P, U)\}$$

and

$$w((P, Q)) := \inf\{l(x, u) \mid x \in P, u \in U : f(x, u) \in Q\}.$$

For $P, Q \in \mathcal{P}$ a *path* $p(P, Q)$ joining P and Q is a sequence of edges $e_1 = (P_0, P_1), e_2 = (P_1, P_2), \dots, e_l = (P_{l-1}, P_l)$ with $P_0 = P$ and $P_l = Q$. Its *length* is defined as

$$L(p(P, Q)) := \sum_{k=1}^l w(e_k).$$

The *shortest path problem* then consists of computing the value

$$V_{\mathcal{P}}(P) := \inf\{L(p(P, Q)) \mid Q \subset T\}$$

with the convention $\inf \emptyset = \infty$, i.e., it computes the length of the shortest path in the graph joining P and some cell Q in the target T . Such a shortest path problem can be efficiently solved using, e.g., Dijkstra's algorithm, see [15].

This shortest path problem assigns a value to each node of the graph G and thus to each cell P of the partition \mathcal{P} . If for each $x \in X$ we denote by $\rho(x) \in \mathcal{P}$ the cell containing x , then we can define the numerical value function on X as $\tilde{V}_{\mathcal{P}}(x) = V_{\mathcal{P}}(\rho(x))$. For this function it turns out that

$$\tilde{V}_{\mathcal{P}}(x) \leq V_{\infty}(x) \quad \text{and} \quad \tilde{V}_{\mathcal{P}_i}(x) \rightarrow V_{\infty}(x),$$

where the convergence holds for subsequently finer partitions \mathcal{P}_i with target sets $T_i \rightarrow \{x^*\}$, see [15]. Furthermore, as shown in [6], under suitable conditions this convergence is even uniform and the feedback defined using (10) asymptotically stabilizes the system — due to the use of the target set T not necessarily exactly at x^* but at least at a neighborhood of x^* , a property called *practical stabilization*.

The main limitation of this approach is that typically rather fine partitions \mathcal{P} are needed in order to make the resulting feedback work, thus the method often exceeds the computer's memory capability. In other words, we are exactly in the situation described at the beginning of Section 3.

A remedy for this problem can be obtained by analyzing the construction of $\tilde{V}_{\mathcal{P}}$ in more detail. In fact, the shortest path problem leads to the optimality principle

$$V_{\mathcal{P}}(P) = \inf_{(P,Q) \in E} \{w((P,Q)) + V_{\mathcal{P}}(Q)\}$$

which by construction of the graph is equivalent to

$$\tilde{V}_{\mathcal{P}}(x) = \inf_{u \in U, x' \in \rho(x)} \{l(x', u) + \tilde{V}_{\mathcal{P}}(f(x', u))\}.$$

If we could change this last equation to

$$\tilde{V}_{\mathcal{P}}(x) = \inf_{u \in U} \sup_{x' \in \rho(x)} \{l(x', u) + \tilde{V}_{\mathcal{P}}(f(x', u))\}, \quad (13)$$

and if $\tilde{F}(x) \in U$ denotes the minimizer of (13), then we immediately obtain

$$\tilde{V}_{\mathcal{P}}(x) \geq l(x, \tilde{F}(x)) + \tilde{V}_{\mathcal{P}}(f(x, \tilde{F}(x))),$$

i.e., (11) with $\alpha = 1$ — independently of how fine the partition \mathcal{P} is chosen.

This is indeed possible by modifying the shortest path problem defining $\tilde{V}_{\mathcal{P}}$:

Instead of a set E of edges $e = (P, Q)$ we now define a set H of *hyperedges*, i.e., pairs $h = (P, \mathcal{N})$ with $\mathcal{N} \in 2^{\mathcal{P}}$, by

$$H := \{(P, \mathcal{N}) \subset \mathcal{P} \times 2^{\mathcal{P}} \mid \mathcal{N} = F(P, u) \text{ for some } u \in U\}$$

with weights

$$w((P, \mathcal{N})) := \inf_{x \in P} \{\sup_{u \in U} l(x, u) \mid \mathcal{N} = F(P, u)\}$$

leading to a *directed hypergraph*. It turns out (see [16, 17]) that Dijkstra's algorithm can be efficiently extended to hypergraphs, leading to values $V_{\mathcal{P}}(P)$ of the nodes satisfying the optimality principle

$$V_{\mathcal{P}}(P) = \inf_{(P, \mathcal{N}) \in H} \sup_{Q \in \mathcal{N}} \{w((P, Q)) + V_{\mathcal{P}}(Q)\}$$

which is equivalent to the desired equation (13).

In order to illustrate the benefit of this hypergraph approach whose solution “automatically” satisfies (11) with $\alpha = 1$, we consider the classical inverted pendulum on a cart given by

$$\left(\frac{4}{3} - m_r \cos^2 \varphi\right) \ddot{\varphi} + \frac{1}{2} m_r \dot{\varphi}^2 \sin 2\varphi - \frac{g}{\ell} \sin \varphi = -u \frac{m_r}{m\ell} \cos \varphi,$$

where we have used the parameters $m = 2$ for the pendulum mass, $m_r = m/(m+M)$ for the mass ratio with cart mass $M = 8$, $\ell = 0.5$ as the length of the pendulum and $g = 9.8$ for the gravitational constant. We use the corresponding sampled-data system (2) with $T = 0.1$ and the running cost

$$l(\varphi_0, \dot{\varphi}_0, u) = \frac{1}{2} \int_0^T 0.1\varphi(t, \varphi_0, \dot{\varphi}_0, u)^2 + 0.05\dot{\varphi}(t, \varphi_0, \dot{\varphi}_0, u)^2 + 0.01u^2 dt. \quad (14)$$

and choose $X = [-8, 8] \times [-10, 10]$ as the region of interest.

The following figure compares the two approaches on the respective coarsest possible partitions on which stabilization was achieved. It is clearly visible that the hypergraph approach (right) leads to both considerably fewer partition cells and to a much faster convergence of the controlled trajectory.




Fig. 1. Approximate optimal value function and resulting feedback trajectory for the inverted pendulum on a 2^{18} box partition using the graph approach (left) and on a 2^{14} box partition using the hypergraph approach (right)

5 A receding horizon approach

Receding horizon control — also known as model predictive control — is probably the most successful class of optimization based control methods and is widely used in industrial applications.

In its simplest form, receding horizon control consists in truncating the infinite horizon functional, i.e., for $N \in \mathbb{N}$ we consider the functional

$$J_N(x_0, u) = \sum_{n=0}^{N-1} l(x_u(n), u(n)), \quad (15)$$

with optimal value function $V_N(x_0) := \inf_{u \in \mathcal{U}} J_N(x_0, u)$.

This problem can be solved by various numerical techniques, e.g. by converting the problem into a static optimization problem with the dynamics as constraints, which can be solved by the SQP method, cf., e.g., [18].

In order to get a feedback law F_N from this finite horizon problem, at each time instant we measure the current state x_n and (online) minimize (15) with $x_0 = x_n$. This yields an optimal control sequence $u^*(0), \dots, u^*(N-1)$ from which we obtain the feedback by setting

$$F_N(x_n) := u^*(0),$$

i.e., by taking the first element of the optimal control sequence.

The questions we want to investigate now is whether this scheme does yield a stabilizing feedback control F_N and, if yes, what is the performance of this controller, i.e, what is $V_\infty^{F_N}$. Our main motivation for this analysis is to derive conditions on the running cost l under which we can ensure stability and good performance even for short optimization horizons N , leading to low complexity and thus short computational time for solving the online optimization problem.

In the literature, the majority of papers dealing with stability issues of receding horizon control use additional terminal constraints and costs, typically requiring $x(N)$ to lie in a neighborhood of the equilibrium to be stabilized. This modification is known to enhance stability both in theory and in practice, however, its main disadvantage is that the operating region of the resulting controller is restricted to the feasible set, i.e., to the set of initial conditions for which the terminal constraints are feasible. This set, in turn, depends on the optimization horizon N and, thus, in order to obtain large operating regions, typically large optimization horizons N are needed leading to complex optimization problems and high computational effort.

Stability results for receding horizon problems without terminal costs have been presented in [19, 20] using convergence $V_N \rightarrow V_\infty$. Another way to obtain such results has been pursued in [21] based on the convergence $|V_N - V_{N+1}| \rightarrow 0$ as $N \rightarrow \infty$. The next proposition is one of the main results from [21], which will allow us to apply Theorem 1 to $\tilde{V} = V_N$ and $\tilde{F} = F_N$.

Proposition 1. *Consider $\gamma > 0$ and $N \geq 2$ and assume that the inequalities*

$$V_2(x) \leq (\gamma + 1)l(x, F_1(x)) \quad \text{and} \quad V_k(x) \leq (\gamma + 1)l(x, F_k(x)), \quad k = 3, \dots, N$$

hold for all $x \in X$. Then the inequality

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

holds for $x \in X$.

The proof can be found in [21] and relies on a inductive application of the optimality principle for V_k , $k = 1, \dots, N$.

Combining Proposition 1 and Theorem 1 we immediately arrive at the following theorem, which was first proved in [21].

Theorem 2. *Consider $\gamma > 0$, let $N \in \mathbb{N}$ be so large that $(\gamma + 1)^{N-2} > \gamma^N$ holds and let the assumption of Proposition 1 holds. Then the inequality*

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

holds. If, in addition, the inequalities (8) and (9) hold, then the feedback law F_N asymptotically stabilizes the system.

Proof. First note that the first and the last inequality in the assertion are obvious. In order to derive the middle inequality, from Proposition 1 we obtain

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

Using the optimality principle for V_N

$$V_N(x) = l(x, F_N(x)) + V_{N-1}(f(x, F_N(x)))$$

and the assumption of Proposition 1 for $k = N$, we obtain the inequality $V_{N-1}(f(x, F_N(x))) \leq \gamma l(x, F_N(x))$ and can conclude

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

Thus, using the optimality principle for V_N once again yields

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

implying the assumption of Theorem 1(i) for $\tilde{V} = V_N$ with

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

and thus the asserted inequalities.

Asymptotic stability now follows immediately from Theorem 1(ii) since the proved inequality together with (8) and (9) implies (9) for $\tilde{V} = V_N$.

We emphasize that the decisive condition for stability is $(\gamma + 1)^{N-2} > \gamma^N$. In particular, the larger γ is, the larger the optimization horizon N must be in order to meet this condition. Hence, in order to ensure stability for small N , we need to ensure that γ is small.

An estimate for γ can, e.g., be obtained if a null-controlling control sequence is known, i.e., if for each x_0 we can find a sequence $u \in \mathcal{U}$ such that $l(x_u(n), u(n))$ converges to 0 sufficiently fast. In this case, for each $k \in \mathbb{N}$ we can estimate

$$V_k(x_0) \leq V_\infty(x_0) \leq J_\infty(x_0, u) \quad \text{and} \quad l^*(x_0) \leq l(x_0, F_k(x_0))$$

and an estimate for γ can then be computed comparing $J_\infty(x_0, u)$ and $l^*(x_0)$. In particular, such an analysis can be used for the design of running costs l which lead to small values of γ and thus to stability for small optimization horizons N .

We illustrate this procedure for a control system governed by a reaction-advection-diffusion PDE with distributed control given by

$$y_t = y_x + \nu y_{xx} + \mu y(y + 1)(1 - y) + u \tag{16}$$

with solutions $y = y(t, x)$ ¹ for $x \in \Omega = (0, 1)$, boundary conditions $y(t, 0) = y(t, 1) = 0$, initial condition $y(0, x) = y_0(x)$ and distributed control $u(t, \cdot) \in L^2(\Omega)$. The corresponding discrete time system (1), whose solutions and control functions we denote by $y(n, x)$ and $u(n, x)$, respectively, is the sampled-data system obtained according to (2) with sampling period $T = 0.025$.

For our numerical computations we discretized the equation in space by finite differences on a grid with nodes $x_i = i/M$, $i = 0, \dots, M$, using backward (i.e., upwind) differences for the advection part y_x . Figure 2 shows the equilibria of the discretized system for $u \equiv 0$, $\nu = 0.1$, $\mu = 10$ and $M = 25$.




Fig. 2. Equilibria for $u \equiv 0$; solid=asymptotically stable, dashed=unstable

Our goal is to stabilize the unstable equilibrium $y^* \equiv 0$, which is possible because with the additive distributed control we can compensate the whole dynamics of the system. In order to achieve this task, a natural choice for a running cost l is the tracking type functional

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

which we implemented with $\lambda = 10^{-3}$ for the discretized model in MATLAB using the `lsqnonlin` solver for the resulting optimization problem.

The simulations shown in Figure 3 reveal that the performance of this controller is not completely satisfactory: for $N = 11$ the solution remains close to $y^* = 0$ but does not converge while for $N = 3$ the solution even grows.

The reason for this behavior lies in the fact that in order to control the system to $y^* = 0$, in (16) the control needs to compensate for y_x , i.e., any stabilizing control must satisfy $\|u(n, \cdot)\|_{L^2(\Omega)}^2 \gtrsim \|y_x(n, \cdot)\|_{L^2(\Omega)}^2$. Thus, for any stabilizing control sequence u we obtain $J_\infty(y_0, u) \gtrsim \lambda \|y_x(n, \cdot)\|_{L^2(\Omega)}^2$ which — even for small values of λ — may be considerably larger than $l^*(y) = \|y\|_{L^2(\Omega)}^2$, resulting in a large γ and thus the need for a large optimization horizon N in order to achieve stability.

¹ Note the change in the notation: x is the independent state variable while $y(t, \cdot)$ is the new state, i.e., X is now an infinite dimensional space.




Fig. 3. Receding horizon with l from (17), $N = 3$ (left) and $N = 11$ (right)

This effect can be avoided by changing l in such a way that $l^*(y)$ includes $\|y_x\|_{L^2(\Omega)}^2$, e.g., by setting

$$l(y(n, \cdot), u(n, \cdot)) = \|y(n, \cdot)\|_{L^2(\Omega)}^2 + \|y_x(n, \cdot)\|_{L^2(\Omega)}^2 + \lambda \|u(n, \cdot)\|_{L^2(\Omega)}^2. \quad (18)$$

For this l the control effort needed in order to control (16) to $y^* = 0$ is proportional to $l^*(y)$. Thus, γ is essentially proportional to λ and thus, in particular, small for our choice of $\lambda = 10^{-3}$ which implies stability even for small optimization horizon N . The simulations using the corresponding discretized running cost illustrated in Figure 4 show that this is indeed the case: we obtain asymptotic stability even for the very small optimization horizons $N = 2$ and $N = 3$.

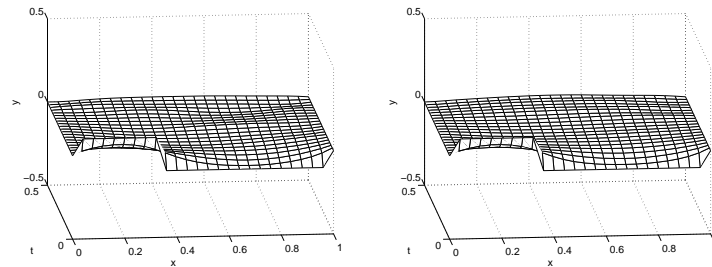


Fig. 4. Receding horizon with l from (18), $N = 2$ (left) and $N = 3$ (right)

Acknowledgment: I would like to thank Oliver Junge and Karl Worthmann for computing the examples in Sections 4 and 5, respectively.

References

1. Bertsekas, D.P.: Dynamic Programming and Optimal Control. Vol. 1 and 2. Athena Scientific, Belmont, MA (1995)

2. Grüne, L., Nešić, D.: Optimization based stabilization of sampled–data nonlinear systems via their approximate discrete–time models. *SIAM J. Control Optim.* **42** (2003) 98–122
3. Camilli, F., Grüne, L., Wirth, F.: Control Lyapunov functions and Zubov’s method. *SIAM J. Control Optim.* (2008) To appear.
4. Kreisselmeier, G., Birkhölzer, T.: Numerical nonlinear regulator design. *IEEE Trans. Autom. Control* **39**(1) (1994) 33–46
5. Johansen, T.A.: Approximate explicit receding horizon control of constrained nonlinear systems. *Automatica* **40**(2) (2004) 293–300
6. Grüne, L., Junge, O.: A set oriented approach to optimal feedback stabilization. *Syst. Control Lett.* **54**(2) (2005) 169–180
7. Grüne, L.: Homogeneous state feedback stabilization of homogeneous systems. *SIAM J. Control Optim.* **38** (2000) 1288–1314
8. Tuna, E.S.: Optimal regulation of homogeneous systems. *Automatica* **41**(11) (2005) 1879–1890
9. Bardi, M., Capuzzo Dolcetta, I.: *Optimal Control and Viscosity Solutions of Hamilton–Jacobi–Bellman equations.* Birkhäuser, Boston (1997)
10. Grüne, L.: An adaptive grid scheme for the discrete Hamilton–Jacobi–Bellman equation. *Numer. Math.* **75**(3) (1997) 319–337
11. Camilli, F., Grüne, L., Wirth, F.: A regularization of Zubov’s equation for robust domains of attraction. In Isidori, A., Lamnabhi-Lagarigue, F., Respondek, W., eds.: *Nonlinear Control in the Year 2000, Volume 1. Lecture Notes in Control and Information Sciences 258*, NCN, Springer Verlag, London (2000) 277–290
12. Lincoln, B., Rantzer, A.: Relaxing dynamic programming. *IEEE Trans. Autom. Control* **51** (2006) 1249–1260
13. Rantzer, A.: Relaxed dynamic programming in switching systems. *IEE Proceedings — Control Theory and Applications* **153** (2006) 567–574
14. Grüne, L., Junge, O.: Approximately optimal nonlinear stabilization with preservation of the Lyapunov function property. In: *Proceedings of the 46th IEEE Conference on Decision and Control*, New Orleans, Louisiana. (2007)
15. Junge, O., Osinga, H.M.: A set oriented approach to global optimal control. *ESAIM Control Optim. Calc. Var.* **10**(2) (2004) 259–270
16. Grüne, L., Junge, O.: Global optimal control of perturbed systems. *J. Optim. Theory Appl.* **136** (2008) To appear.
17. von Lossow, M.: A min-max version of Dijkstra’s algorithm with application to perturbed optimal control problems. In: *Proceedings of the GAMM Annual meeting*, Zürich, Switzerland. (2007) To appear.
18. Grüne, L., Nešić, D., Pannek, J.: Model predictive control for nonlinear sampled–data systems. In Allgöwer, F., Biegler, L., Findeisen, R., eds.: *Assessment and Future Directions of Nonlinear Model Predictive Control. Volume 358 of Springer LNCIS.* (2007)
19. Jadbabaie, A., Hauser, J.: On the stability of receding horizon control with a general terminal cost. *IEEE Trans. Automat. Control* **50**(5) (2005) 674–678
20. Grimm, G., Messina, M.J., Tuna, S.E., Teel, A.R.: Model predictive control: for want of a local control Lyapunov function, all is not lost. *IEEE Trans. Automat. Control* **50**(5) (2005) 546–558
21. Grüne, L., Rantzer, A.: On the infinite horizon performance of receding horizon controllers. Preprint, Universität Bayreuth (2006) www.math.uni-bayreuth.de/~lgruene/publ/infhorrhc.html.