Nonlinear Model Predictive Control

Lars Grüne

Mathematical Institute, University of Bayreuth, Germany

AICES Regional School, RWTH Aachen, 12 May 2016

Contents

(1) Introduction
What is Model Predictive Control (MPC)?

Setup

We consider nonlinear discrete time control systems

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

or, briefly

$$x^+ = f(x, u)$$

with $x \in X$, $u \in U$

- we consider discrete time systems for simplicity of exposition
- continuous time systems can be treated by using the discrete time representation of the corresponding sampled data system or a numerical approximation
- X and U depend on the model. These may be Euclidean spaces $\mathbb{R}^n$ and $\mathbb{R}^m$ or more general (e.g., infinite dimensional) spaces. For simplicity of exposition we assume that we have a norm $\|\cdot\|$ on both spaces

Prototype Problem

Assume there exists an equilibrium $x_* \in X$ for $u = 0$, i.e.

$$f(x_*, 0) = x_*$$

Task: stabilize the system

$$x^+ = f(x, u)$$

at x_* via static state feedback, i.e., find $\mu : X \to U$, such that x_* is asymptotically stable for the feedback controlled system

$$x_\mu(n+1) = f(x_\mu(n), \mu(x_\mu(n))), \;\; x_\mu(0) = x_0$$

Additionally, we impose state constraints $x_\mu(n) \in \mathbb{X}$
and control constraints $\mu(x(n)) \in \mathbb{U}$

for all $n \in \mathbb{N}$ and given sets $\mathbb{X} \subseteq X$, $\mathbb{U} \subseteq U$

Prototype Problem

Asymptotic stability means

Attraction: $x_\mu(n) \to x_*$ as $n \to \infty$

plus

Stability: Solutions starting close to x_* remain close to x_*

(we will later formalize this property using $\mathcal{KL}$ functions)

Informal interpretation: control the system to x_* and keep it there while obeying the state and control constraints

Idea of MPC: use an optimal control problem which minimizes the distance to x_* in order to synthesize a feedback law μ

The idea of MPC

For defining the MPC scheme, we choose a stage cost $\ell(x, u)$ penalizing the distance from x_* and the control effort, e.g., $\ell(x, u) = \|x - x_*\|^2 + \lambda\|u\|^2$ for $\lambda \geq 0$

The basic idea of MPC is:

- minimize the summed stage cost along trajectories generated from our model over a prediction horizon N
- use the first element of the resulting optimal control sequence as feedback value
- repeat this procedure iteratively for all sampling instants $n = 0, 1, 2, \ldots$

Notation in what follows:

- general feedback laws will be denoted by μ
- the MPC feedback law will be denoted by μ_N

The basic MPC scheme

Formal description of the basic MPC scheme:

At each time instant n solve for the current state $x_{\mu_N}(n)$

$$\underset{\mathbf{u}\text{ admissible}}{\text{minimize}} \;\; J_N(x_{\mu_N}(n), \mathbf{u}) = \sum_{k=0}^{N-1} \ell(x_\mathbf{u}(k), \mathbf{u}(k)), \;\; x_\mathbf{u}(0) = x_{\mu_N}(n)$$

($\mathbf{u}$ admissible $\;\Leftrightarrow\;$ $\mathbf{u} \in \mathbb{U}^N$ and $x_\mathbf{u}(k) \in \mathbb{X}$)

$\rightsquigarrow$ optimal trajectory $\;\; x^\star(0), \ldots, x^\star(N)$

with optimal control $\mathbf{u}^\star(0), \ldots, \mathbf{u}^\star(N-1)$

Define the MPC feedback law $\;\; \mu(x_{\mu_N}(n)) := \mathbf{u}^\star(0)$

$\rightsquigarrow$ $x_{\mu_N}(n+1) = f(x_{\mu_N}(n), \mu_N(x_{\mu_N}(n))) = f(x_{\mu_N}(n), \mathbf{u}^\star(0)) = x^\star(1)$

MPC from the trajectory point of view

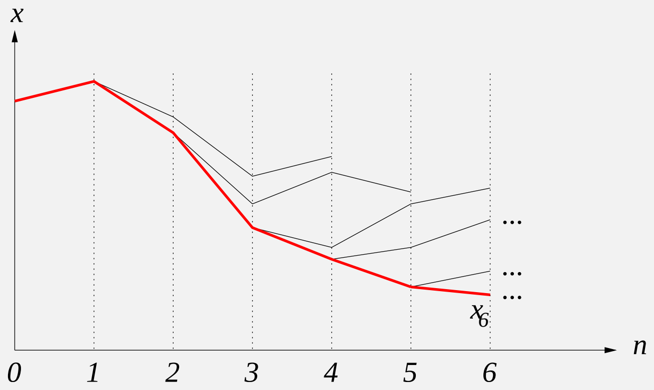

black = predictions (open loop optimization)
red = MPC closed loop, $x_n = x_{\mu_N}(n)$

Model predictive control (aka Receding horizon control)

Idea first formulated by A.I. Propoi in 1963, often rediscovered

used in industrial applications since the mid 1970s, mainly for constrained linear systems [Qin & Badgwell, 1997, 2001]

more than 9000 industrial MPC applications in Germany counted in [Dittmar & Pfeifer, 2005]

development of theory since ∼1980 (linear), ∼1990 (nonlinear)

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

Central questions:

- When does MPC stabilize the system?
- How good is the performance of the MPC feedback law?
- How long does the optimization horizon N need to be?

and, of course, the development of good algorithms (not topic of this course)

An example

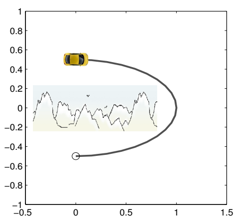

$$x_1^+ = \sin(\varphi + u)$$
$$x_2^+ = \cos(\varphi + u)/2$$

with $\varphi = \begin{cases} \arccos 2x_2, & x_1 \geq 0 \\ 2\pi - \arccos 2x_2, & x_1 < 0, \end{cases}$

$\mathbb{X} = \{x \in \mathbb{R}^2 : \|(x_1, 2x_2)^T\| = 1\}$, $\mathbb{U} = [0, u_{\max}]$

$x_* = (0, -1/2)^T$, $x_0 = (0, 1/2)^T$

MPC with $\ell(x,u) = \|x - x_*\|^2 + |u|^2$ and $u_{\max} = 0.2$ yields asymptotic stability for $N = 11$ but not for $N \leq 10$

Summary of Section (1)

- MPC is an online optimal control based method for computing stabilizing feedback laws
- MPC computes the feedback law by iteratively solving finite horizon optimal control problems using the current state $x_0 = x_{\mu_N}(n)$ as initial value
- the feedback value $\mu_N(x_0)$ is the first element of the resulting optimal control sequence
- MPC can considerably reduce the computation time needed for solving infinite horizon optimal control problems, up to real time capability
  $\rightsquigarrow$ model reduction in time
- the car-and-mountain example shows that MPC does not always yield an asymptotically stabilizing feedback law

(2a) Background material: Lyapunov functions

Purpose of this section

We introduce Lyapunov functions as a tool to rigorously verify asymptotic stability

In the subsequent sections, this will be used in order to establish asymptotic stability of the MPC closed loop

In this section, we consider discrete time systems without input, i.e.,

$$x^+ = g(x)$$

with $x \in X$ or, in long form

$$x(n+1) = g(x(n)), \;\; x(0) = x_0$$

(later we will apply the results to $g(x) = f(x, \mu_N(x))$)

Note: we do not require g to be continuous

UNIVERSITÄT BAYREUTH

Comparison functions

For $\mathbb{R}_0^+ = [0, \infty)$ we use the following classes of comparison functions

$$\mathcal{K} \;:=\; \left\{ \alpha : \mathbb{R}_0^+ \to \mathbb{R}_0^+ \;\middle|\; \begin{array}{l} \alpha \text{ is continuous and strictly} \\ \text{increasing with } \alpha(0) = 0 \end{array} \right\}$$

$$\mathcal{K}_\infty \;:=\; \left\{ \alpha : \mathbb{R}_0^+ \to \mathbb{R}_0^+ \;\middle|\; \alpha \in \mathcal{K} \text{ and } \alpha \text{ is unbounded} \right\}$$

$$\mathcal{KL} \;:=\; \left\{ \beta : \mathbb{R}_0^+ \times \mathbb{R}_0^+ \to \mathbb{R}_0^+ \;\middle|\; \begin{array}{l} \beta \text{ is continuous,} \\ \beta(\cdot, t) \in \mathcal{K} \text{ for all } t \in \mathbb{R}_0^+ \\ \text{and } \beta(r, \cdot) \text{ is strictly de-} \\ \text{creasing to } 0 \text{ for all } r \in \mathbb{R}_0^+ \end{array} \right\}$$

UNIVERSITÄT BAYREUTH

Asymptotic stability revisited

A point x_* is called an equilibrium of $x^+ = g(x)$ if $g(x_*) = x_*$

A set $Y \subseteq X$ is called forward invariant for $x^+ = g(x)$ if $g(x) \in Y$ holds for each $x \in Y$

We say that x_* is asymptotically stable for $x^+ = g(x)$ on a forward invariant set Y if there exists $\beta \in \mathcal{KL}$ such that

$$\|x(n) - x_*\| \leq \beta(\|x(0) - x_*\|, n)$$

holds for all $x \in Y$ and $n \in \mathbb{N}$

How can we check whether this property holds?

UNIVERSITÄT BAYREUTH

Lyapunov function

Let $Y \subseteq X$ be a forward invariant set and $x_* \in X$. A function $V: Y \to \mathbb{R}_0^+$ is called a Lyapunov function for $x^+ = g(x)$ if the following two conditions hold for all $x \in Y$:

(i) There exists $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$ such that

$$\alpha_1(\|x - x_*\|) \leq V(x) \leq \alpha_2(\|x - x_*\|)$$

(ii) There exists $\alpha_V \in \mathcal{K}$ such that

$$V(x^+) \leq V(x) - \alpha_V(\|x - x_*\|)$$

UNIVERSITÄT BAYREUTH

Stability theorem

Theorem: If the system $x^+ = g(x)$ admits a Lyapunov function V on a forward invariant set Y, then x_* is an asymptotically stable equilibrium on Y

Idea of proof: $V(x^+) \leq V(x) - \alpha_V(\|x - x_*\|)$ implies that V is strictly decaying along solutions away from x_*

This allows to construct $\tilde{\beta} \in \mathcal{KL}$ with $V(x(n)) \leq \tilde{\beta}(V(x(0)), n)$

The bounds $\alpha_1(\|x - x_*\|) \leq V(x) \leq \alpha_2(\|x - x_*\|)$ imply that asymptotic stability holds with $\beta(r, t) = \alpha_1^{-1}(\tilde{\beta}(\alpha_2(r), t))$

UNIVERSITÄT BAYREUTH

Lyapunov functions — discussion

While the convergence $x(n) \to x_*$ is typically non-monotone for an asymptotically stable system, the convergence $V(x(n)) \to 0$ is strictly monotone

It is hence sufficient to check the decay of V in one time step

$\rightsquigarrow$ it is typically quite easy to check whether a given function is a Lyapunov function

But it is in general difficult to find a candidate for a Lyapunov function

For MPC, we will use the optimal value functions which we introduce in the next section

UNIVERSITÄT BAYREUTH

(2b) Background material:

Dynamic Programming

Purpose of this section

We define the optimal value functions V_N for the optimal control problem

$$\underset{\mathbf{u}\text{ admissible}}{\text{minimize}} \quad J_N(x_0,\mathbf{u}) = \sum_{k=0}^{N-1} \ell(x_{\mathbf{u}}(k),\mathbf{u}(k)), \quad x_{\mathbf{u}}(0) = x_0$$

used within the MPC scheme (with $x_0 = x_{\mu_N}(n)$)

We present the dynamic programming principle, which establishes a relation for these functions and will eventually enable us to derive conditions under which V_N is a Lyapunov function

Optimal value functions

We define the optimal value function

$$V_N(x_0) := \inf_{\mathbf{u}\text{ admissible}} J_N(x_0,\mathbf{u})$$

setting $V_N(x_0) := \infty$ if x_0 is not feasible, i.e., if there is no admissible $\mathbf{u}$ (recall: $\mathbf{u}$ admissible $\Leftrightarrow$ $x_{\mathbf{u}}(k) \in \mathbb{X}$, $\mathbf{u}(k) \in \mathbb{U}$)

An admissible control sequence $\mathbf{u}^\star$ is called optimal, if

$$J_N(x_0,\mathbf{u}^\star) = V_N(x_0)$$

Note: an optimal $\mathbf{u}^\star$ does not need to exist in general. In the sequel we assume that $\mathbf{u}^\star$ exists if x_0 is feasible

Dynamic Programming Principle

Theorem: (Dynamic Programming Principle) For any feasible $x_0 \in \mathbb{X}$ the optimal value function satisfies

$$V_N(x_0) = \inf_{\substack{u\in\mathbb{U}\\ f(x_0,u)\in\mathbb{X}}} \{\ell(x_0,u) + V_{N-1}(f(x_0,u))\}$$

Moreover, if $\mathbf{u}^\star$ is an optimal control, then

$$V_N(x_0) = \ell(x_0,\mathbf{u}^\star(0)) + V_{N-1}(f(x_0,\mathbf{u}^\star(0)))$$

holds.

Idea of Proof: Follows by taking infima in the identity

$$\begin{aligned} J_N(x_0,\mathbf{u}) &= \ell(x_{\mathbf{u}}(0),\mathbf{u}(0)) + \sum_{k=1}^{N-1} \ell(x_{\mathbf{u}}(k),\mathbf{u}(k)) \\ &= \ell(x_0,\mathbf{u}(0)) + J_{N-1}(f(x_0,\mathbf{u}(0)),\mathbf{u}(\cdot+1)) \end{aligned}$$

Corollaries

Corollary: Let $x^\star$ be an optimal trajectory of length N with optimal control $u^\star$ and $x^\star(0) = x$. Then

(i) The "tail"

$$\big(x^\star(k), x^\star(k+1),\ldots,x^\star(N-1)\big)$$

is an optimal trajectory of length $N-k$.

(ii) The MPC feedback μ_N satisfies

$$\mu_N(x) \in \operatorname*{argmin}_{u\in\mathbb{U}} \{\ell(x,u) + V_{N-1}(f(x,u))\}$$

(i.e., $u = \mu_N(x)$ minimizes this expression),

$$V_N(x) = \ell(x,\mu_N(x)) + V_{N-1}(f(x,\mu_N(x)))$$

and

$$u^\star(k) = \mu_{N-k}(x^\star(k)), \quad k = 0,\ldots,N-1$$

Dynamic Programming Principle — discussion

We will see later, that under suitable conditions the optimal value function will play the role of a Lyapunov function for the MPC closed loop

The dynamic programming principle and its corollaries will prove to be important tools to establish this fact

In order to see why this can work, in the next section we briefly look at infinite horizon optimal control problems

Moreover, for simple systems the principle can be used for computing V_N and μ_N — we will see an example in the excercises

(2c) Background material:

Relaxed Dynamic Programming

Infinite horizon optimal control

Just like the finite horizon problem we can define the infinite horizon optimal control problem

$$\underset{\mathbf{u} \text{ admissible}}{\text{minimize}} \quad J_\infty(x_0, \mathbf{u}) = \sum_{k=0}^{\infty} \ell(x_\mathbf{u}(k), \mathbf{u}(k)), \quad x_\mathbf{u}(0) = x_0$$

and the corresponding optimal value function

$$V_\infty(x_0) := \inf_{\mathbf{u} \text{ admissible}} J_\infty(x_0, \mathbf{u})$$

If we could compute an optimal feedback μ_∞ for this problem (which is — in contrast to computing μ_N — in general a very difficult problem), we would have solved the stabilization problem

Infinite horizon dynamic programming principle

Recall the corollary from the finite horizon dynamic programming principle

$$V_N(x) = \ell(x, \mu_N(x)) + V_{N-1}(f(x, \mu_N(x)))$$

The corresponding result which can be proved for the infinite horizon problem reads

$$V_\infty(x) = \ell(x, \mu_\infty(x)) + V_\infty(f(x, \mu_\infty(x)))$$

$\rightsquigarrow$ if $\ell(x, \mu_\infty(x)) \ge \alpha_V(\|x - x_*\|)$ holds, then we get

$$V_\infty(f(x, \mu_\infty(x))) \le V_\infty(x) - \alpha_V(\|x - x_*\|)$$

and if in addition $\alpha_1(\|x - x_*\|) \le V(x) \le \alpha_2(\|x - x_*\|)$ holds, then V_∞ is a Lyapunov function $\rightsquigarrow$ asymptotic stability

Relaxing dynamic programming

Unfortunately, an equation of the type

$$V_\infty(x) = \ell(x, \mu_\infty(x)) + V_\infty(f(x, \mu_\infty(x)))$$

cannot be expected if we replace "∞" by "N" everywhere (in fact, it would imply $V_N = V_\infty$)

However, we will see that we can establish relaxed versions of this inequality in which we

- relax "$=$" to "$\geq$"
- relax $\ell(x, \mu(x))$ to $\alpha\ell(x, \mu(x))$ for some $\alpha \in (0, 1]$

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

"relaxed dynamic programming inequality" [Rantzer et al. '06ff]

What can we conclude from this inequality?

Relaxed dynamic programming

We define the infinite horizon performance of the MPC closed loop system $x^+ = f(x, \mu_N(x))$ as

$$J^{cl}_\infty(x_0, \mu_N) = \sum_{k=0}^{\infty} \ell(x_{\mu_N}(k), \mu_N(x_{\mu_N}(k))), \ \ x_{\mu_N}(0) = x_0$$

Theorem: [Gr./Rantzer '08, Gr./Pannek '11] Let $Y \subseteq \mathbb{X}$ be a forward invariant set for the MPC closed loop and assume that

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

holds for all $x \in Y$ and some $N \in \mathbb{N}$ and $\alpha \in (0, 1]$

Then for all $x \in Y$ the infinite horizon performance satisfies

$$J^{cl}_\infty(x_0, \mu_N) \leq V_N(x_0)/\alpha$$

Relaxed dynamic programming

Theorem (continued): If, moreover, there exists $\alpha_2, \alpha_3 \in \mathcal{K}_\infty$ such that the inequalities

$$V_N(x) \leq \alpha_2(\|x - x_*\|), \quad \inf_{u \in \mathbb{U}} \ell(x, u) \geq \alpha_3(\|x - x_*\|)$$

hold for all $x \in Y$, then the MPC closed loop is asymptotically stable on Y with Lyapunov function V_N.

Proof: The assumed inequalities immediately imply that $V = V_N$ is a Lyapunov function for $x^+ = g(x) = f(x, \mu_N(x))$ with

$$\alpha_1(r) = \alpha_3(r), \quad \alpha_V(r) = \alpha\, \alpha_3(r)$$

$\Rightarrow$ asymptotic stability

Relaxed dynamic programming

For proving the performance estimate $J^{cl}_\infty(x_0, \mu_N) \leq V_N(x_0)/\alpha$, the relaxed dynamic programming inequality implies

$$\begin{aligned} & \alpha \sum_{n=0}^{K-1} \ell(x_{\mu_N}(k), \mu_N(x_{\mu_N}(k))) \\ \leq & \sum_{n=0}^{K-1} \Big(V_N(x_{\mu_N}(n)) - V_N(x_{\mu_N}(n+1)) \Big) \\ = & \ V_N(x_{\mu_N}(0)) - V_N(x_{\mu_N}(K)) \ \leq \ V_N(x_{\mu_N}(0)) \end{aligned}$$

Since all summands are ≥ 0, this implies that the limit for $K \to \infty$ exists and we get

$$\alpha J^{cl}_\infty(x_0, \mu_N) = \alpha \sum_{n=0}^{\infty} \ell(x_{\mu_N}(k), \mu_N(x_{\mu_N}(k))) \leq V_N(x_{\mu_N}(0))$$

$\Rightarrow$ assertion

Summary of Section (2)

- Lyapunov functions are our central tool for verifying asymptotic stability
- Dynamic programming provides us with equations which will be heavily used in the subsequent analysis
- Infinite horizon optimal control would solve the stabilization problem — if we could compute the feedback law μ_∞
- The performance of the MPC controller can be measured by looking at the infinite horizon value along the MPC closed loop trajectories
- Relaxed dynamic programming gives us conditions under which both asymptotic stability and performance results can be derived

Application of background results

The main task will be to verify the assumptions of the relaxed dynamic programming theorem, i.e.,

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

for some $\alpha \in (0,1]$, and

$$V_N(x) \leq \alpha_2(\|x - x_*\|), \quad \inf_{u \in \mathbb{U}} \ell(x,u) \geq \alpha_3(\|x - x_*\|)$$

for all x in a forward invariant set Y for $x^+ = f(x, \mu_N(x))$

To this end, we present two different approaches:

- modify the optimal control problem in the MPC loop by adding terminal constraints and costs
- derive assumptions on f and ℓ under which MPC works without terminal constraints and costs

(3) Stabilizing Model Predictive Control with stabilizing terminal conditions

V_N as a Lyapunov Function

Problem: Prove that the MPC feedback law μ_N is stabilizing

Approach: Verify the assumptions

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

for some $\alpha \in (0,1]$, and

$$V_N(x) \leq \alpha_2(\|x - x_*\|), \quad \inf_{u \in \mathbb{U}} \ell(x,u) \geq \alpha_3(\|x - x_*\|)$$

of the relaxed dynamic programming theorem for the optimal value function

$$V_N(x_0) := \inf_{\mathbf{u} \text{ admissible}} \sum_{k=0}^{N-1} \ell(x_{\mathbf{u}}(k), \mathbf{u}(k)), \quad x_{\mathbf{u}}(0) = x_0$$

Why is this difficult?

Let us first consider the inequality

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

The dynamic programming principle for V_N yields

$$V_N(x) \geq \ell(x, \mu_N(x)) + V_{N-1}(f(x, \mu_N(x)))$$

$\rightsquigarrow$ we have V_{N-1} where we would like to have V_N

$\rightsquigarrow$ we would get the desired inequality if we could ensure

$$V_{N-1}(f(x, \mu_N(x))) \geq V_N(f(x, \mu_N(x))) + \text{"small error"}$$

(where "small" means that the error can be compensated replacing $\ell(x, \mu_N(x))$ by $\alpha \ell(x, \mu_N(x))$ with $\alpha \in (0, 1)$)

Why is this difficult?

Task: Find conditions under which

$$V_{N-1}(x) \geq V_N(x) + \text{"small error"}$$

holds

For

$$V_N(x_0) := \inf_{\mathbf{u} \text{ admissible}} \sum_{k=0}^{N-1} \ell(x_{\mathbf{u}}(k), \mathbf{u}(k)), \quad x_{\mathbf{u}}(0) = x_0$$

this appeared to be out of reach until the mid 1990s

Note: $V_{N-1} \leq V_N$ by non-negativity of ℓ; typically with strict "$<$"

$\rightsquigarrow$ additional stabilizing constraints were proposed

Terminal constraints and cost

Optimal control problem

$$\underset{\mathbf{u} \text{ admissible}}{\text{minimize}} \quad J_N(x_0, \mathbf{u}) = \sum_{k=0}^{N-1} \ell(x_{\mathbf{u}}(k), \mathbf{u}(k)), \quad x_{\mathbf{u}}(0) = x_0$$

We want V_N to become a Lyapunov function

Idea: add local Lyapunov function $F : \mathbb{X}_0 \to \mathbb{R}_0^+$ as terminal cost

$$J_N(x_0, u) = \sum_{k=0}^{N-1} \ell(x_{\mathbf{u}}(k), \mathbf{u}(k)) + F(x_{\mathbf{u}}(N))$$

F is defined on a region $\mathbb{X}_0$ around x_* which is imposed as terminal constraint $\quad x(N) \in \mathbb{X}_0$

[Chen & Allgöwer '98, Jadbabaie et al. '98 ...]

Terminal constraints and cost

We thus change the optimal control problem to

$$\underset{\mathbf{u} \in \mathbb{U}^N_{\mathbb{X}_0}(x_0)}{\text{minimize}} \quad J_N(x_0, \mathbf{u}) = \sum_{k=0}^{N-1} \ell(x_{\mathbf{u}}(k), \mathbf{u}(k)) \; + \; F(x_{\mathbf{u}}(N))$$

with

$$\mathbb{U}^N_{\mathbb{X}_0}(x_0) := \{\mathbf{u} \in \mathbb{U}^N \text{ admissible and } x_{\mathbf{u}}(N) \in \mathbb{X}_0\}$$

Which properties do we need for F and $\mathbb{X}_0$ in order to ensure

$$V_{N-1}(x) \geq V_N(x) + \text{"small error" ?}$$

With suitable assumptions we can even avoid the "small error"

Terminal constraints and cost

Assumptions on $F : \mathbb{X}_0 \to \mathbb{R}_0^+$ and $\mathbb{X}_0$

There exists a controller $\kappa : \mathbb{X}_0 \to \mathbb{U}$ with the following properties:

(i) $\mathbb{X}_0$ is forward invariant for $x^+ = f(x, \kappa(x))$:
for each $x \in \mathbb{X}_0$ we have $f(x, \kappa(x)) \in \mathbb{X}_0$

(ii) F is a Lyapunov function for $x^+ = f(x, \kappa(x))$ on $\mathbb{X}_0$ which is compatible with the stage cost ℓ in the following sense:
for each $x \in \mathbb{X}_0$ the inequality

$$F(f(x, \kappa(x))) \leq F(x) - \ell(x, \kappa(x))$$

holds

Simplest choice: $\mathbb{X}_0 = \{x_*\}$, $F \equiv 0$, $\kappa \equiv 0$ (if $f(x_*, 0) = x_*$)

Prolongation of control sequences

Let $\tilde{\mathbf{u}} \in \mathbb{U}^{N-1}_{\mathbb{X}_0}(x_0) \quad \Rightarrow \quad \tilde{x} := x_{\tilde{\mathbf{u}}}(N-1) \in \mathbb{X}_0$

Define $\mathbf{u} \in \mathbb{U}^N$ as $\mathbf{u}(k) := \begin{cases} \tilde{\mathbf{u}}(k), & k = 0, \ldots, N-2 \\ \kappa(\tilde{x}), & k = N-1 \end{cases}$

with κ from (i)

$\Rightarrow \quad x_{\mathbf{u}}(N) = f(x_{\tilde{\mathbf{u}}}(N-1), \mathbf{u}(N-1)) = f(\tilde{x}, \kappa(\tilde{x})) \in \mathbb{X}_0$

$\Rightarrow \quad \mathbf{u} \in \mathbb{U}^N_{\mathbb{X}_0}(x_0)$

$\rightsquigarrow$ every $\tilde{\mathbf{u}} \in \mathbb{U}^{N-1}_{\mathbb{X}_0}(x_0)$ can be prolonged to an $\mathbf{u} \in \mathbb{U}^N_{\mathbb{X}_0}(x_0)$

By (ii) the stage cost of the prolongation is bounded by

$$\ell(x_{\mathbf{u}}(N-1), \mathbf{u}(N-1)) \leq F(x_{\mathbf{u}}(N-1)) - F(x_{\mathbf{u}}(N))$$

Reversal of $V_{N-1} \leq V_N$

Let $\tilde{\mathbf{u}}^\star \in \mathbb{U}^{N-1}_{\mathbb{X}_0}(x_0)$ be the optimal control for J_{N-1}, i.e.,

$$V_{N-1}(x_0) = J_{N-1}(x_0, \tilde{\mathbf{u}}^\star)$$

Denote by $\mathbf{u} \in \mathbb{U}^N_{\mathbb{X}_0}(x_0)$ its prolongation

$$\begin{aligned} \Rightarrow \quad V_{N-1}(x_0) &= J_{N-1}(x_0, \tilde{\mathbf{u}}^\star) \\ &= \sum_{k=0}^{N-2} \ell(x_{\tilde{\mathbf{u}}^\star}(k), \tilde{\mathbf{u}}^\star(k)) + \underbrace{F(x_{\tilde{\mathbf{u}}^\star}(N-1))}_{\geq \ell(x_{\mathbf{u}}(N-1), \mathbf{u}(N-1)) + F(x_{\mathbf{u}}(N))} \\ &\geq \sum_{n=0}^{N-1} \ell(x_{\mathbf{u}}(k), \mathbf{u}(k)) + F(x_{\mathbf{u}}(N)) \\ &= J_N(x_0, \mathbf{u}) \ \geq \ V_N(x_0) \end{aligned}$$

$\rightsquigarrow$ as desired, we obtain $V_{N-1} \geq V_N$

Feasible sets

Define the feasible set

$$\mathbb{X}_N := \{x \in \mathbb{X} \,|\, \mathbb{U}^N_{\mathbb{X}_0}(x) \neq \emptyset\}$$

On $\mathbb{X}_N$ one can ensure the inequality

$$V_N(x) \leq \alpha_2(\|x - x_*\|)$$

for some $\alpha_2 \in \mathcal{K}_\infty$ under mild conditions, while outside $\mathbb{X}_N$ we get $V_N(x) = \infty$

$\rightsquigarrow$ the MPC control is only defined on $\mathbb{X}_N$!

Stability theorem

Theorem: Consider the MPC scheme with regional terminal constraint $x_{\mathbf{u}}(N) \in \mathbb{X}_0$ and Lyapunov function terminal cost F compatible with ℓ. Assume that

$$V_N(x) \le \alpha_2(\|x - x_*\|), \quad \inf_{u \in \mathbb{U}} \ell(x,u) \ge \alpha_3(\|x - x_*\|)$$

holds for all $x \in \mathbb{X}_N$.

Then $\mathbb{X}_N$ is forward invariant, the MPC closed loop is asymptotically stable on $\mathbb{X}_N$ and the performance estimate

$$J_\infty^{cl}(x, \mu_N) \le V_N(x)$$

holds.

Proof: Combine dynamic programming with $V_{N-1} \ge V_N$

UNIVERSITÄT BAYREUTH

Stabilizing terminal conditions — Discussion

Stabilizing terminal conditions

- yield an elegant stability theory
- can indeed improve the stability properties
- can be derived in a systematic way, e.g., by linearization

But:

- large feasible set usually needs a large optimization horizon N
  (see again the car-and-mountains example)
- additional analytical effort for computing F
- hardly ever used in industrial practice

In Section (4) we will see how stability can be proved without stabilizing terminal constraints

UNIVERSITÄT BAYREUTH

Suboptimality

Recall: For both stabilizing terminal constraints the relaxed dynamic programming theorem yields the estimate

$$J_\infty^{cl}(x_0, \mu_N) \le V_N(x_0)$$

But: How large is V_N ?

Without terminal constraints, the inequality $V_N \le V_\infty$ is immediate

However, the terminal constraints also reverse this inequality, i.e., we have $V_N \ge V_\infty$ and the gap is very difficult to estimate

UNIVERSITÄT BAYREUTH

Suboptimality — example

We consider two examples with $\mathbb{X} = \mathbb{R}$, $\mathbb{U} = \mathbb{R}$ for $N = 2$

Example 1: $x^+ = x + u, \quad \ell(x,u) = x^2 + u^2$

Terminal constraints $\quad x_{\mathbf{u}}(N) = x_* = 0$

$V_\infty(x) \approx 1.618x^2, \quad J_\infty^{cl}(x, \mu_2) = 1.625x^2$

Example 2: as Example 1, but with $\quad \ell(x,u) = x^2 + u^4$

$V_\infty(20) \le 1726, \quad J_\infty^{cl}(x, \mu_2) \approx 11240$

General estimates for fixed N appear difficult to obtain. But we can give an asymptotic result for $N \to \infty$

UNIVERSITÄT BAYREUTH

Asymptotic Suboptimality

Theorem: For both types of terminal constraints the assumptions of the stability theorems ensure

$$V_N(x) \to V_\infty(x)$$

and thus

$$J_\infty^{cl}(x, \mu_N) \to V_\infty(x)$$

as $N \to \infty$ uniformly on compact subsets of the feasible sets, i.e., the MPC performance converges to the optimal one

Idea of proof: uses that any approximately optimal trajectory for J_∞ converges to x_* and can thus be modified to meet the constraints with only moderately changing its value

UNIVERSITÄT BAYREUTH

Summary of Section (3)

- terminal conditions reverse the usual inequality $V_{N-1} \leq V_N$ to $V_{N-1} \geq V_N$
- this enables us to derive the relaxed dynamic programming inequality (with $\alpha = 1$) from the dynamic programming principle
- the operating region is restricted to the feasible set $\mathbb{X}_N$
- $J_\infty^{cl}(x, \mu_N) \leq V_N(x)$ holds and $V_N \to V_\infty$ for $N \to \infty$, but $V_N \gg V_\infty$ is possible

UNIVERSITÄT BAYREUTH

(4) Stabilizing Model Predictive Control without stabilizing terminal conditions

MPC without stabilizing terminal constraints

We return to the basic MPC formulation

$$\underset{\mathbf{u}\text{ admissible}}{\text{minimize}} \quad J_N(x_0, u) = \sum_{k=0}^{N-1} \ell(x_{\mathbf{u}}(k), \mathbf{u}(k)), \quad x_{\mathbf{u}}(0) = x_0 = x_{\mu_N}(n)$$

without any stabilizing terminal constraints and costs

In order to motivate why we want to avoid terminal constraints and costs, we consider an example of P double integrators in the plane

UNIVERSITÄT BAYREUTH

Motivation for avoiding terminal conditions

Example: [Annunziato/Borzì '10ff., Fleig '14ff.] The Fokker-Planck Equation

$$\partial_t y(x,t) - \sum_{i,j=1}^{d} \partial^2_{x_i x_j}\Big(a_{ij}(x,t)y(x,t)\Big) + \sum_{i=1}^{d} \partial_{x_i}\Big(b_i(x,t;u))y(x,t)\Big) = 0$$

$$y(\cdot,0) = y_0$$

is a parabolic PDE describing the evolution of a probability density function for stochastic control systems

By solving a Fokker-Planck control problem controls for large ensembles of stochastic systems can be computed

Due to the complexity of the problem the derivation of terminal conditions meeting the theoretical assumptions is hardly feasible — but even without MPC works

Reference density function

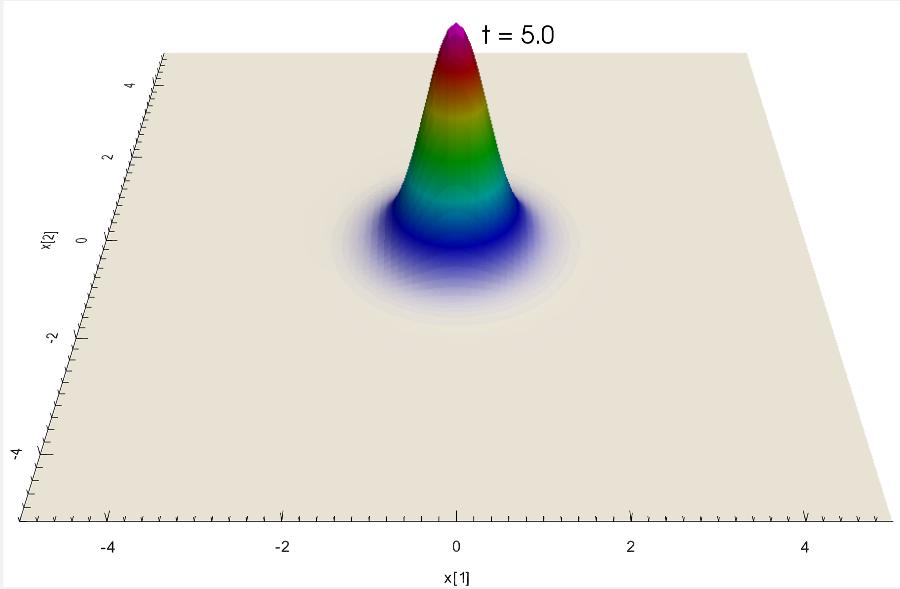

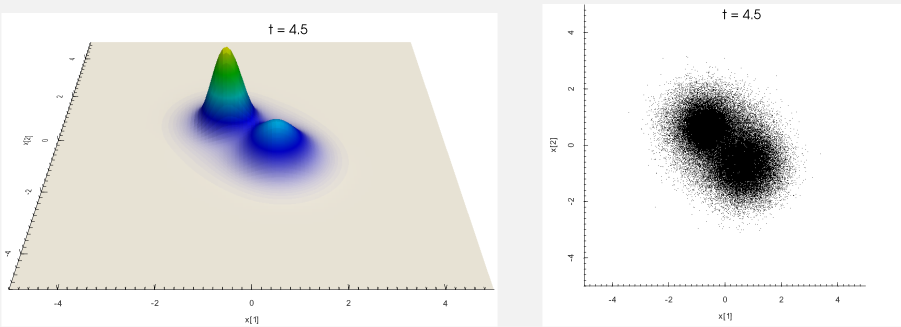

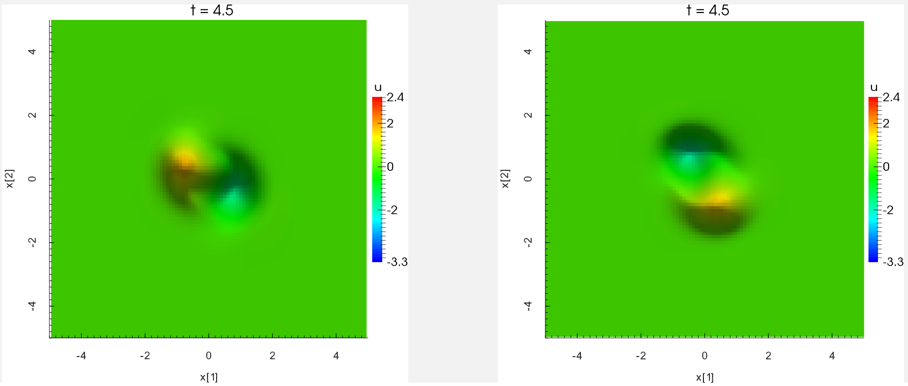

Stabilizing NMPC without terminal constraint

(Some) stability and performance results known in the literature:

[Alamir/Bornard '95]
use a controllability condition for all $x \in \mathbb{X}$

[Shamma/Xiong '97, Primbs/Nevistić '00]
use knowledge of optimal value functions

[Jadbabaie/Hauser '05]
use controllability of linearization in x_*

[Grimm/Messina/Tuna/Teel '05, Tuna/Messina/Teel '06, Gr./Rantzer '08, Gr. '09, Gr./Pannek/Seehafer/Worthmann '10]
use bounds on optimal value functions

Here we explain the last approach

Bounds on the optimal value function

Recall the definition of the optimal value function

$$V_N(x) := \inf_{\mathbf{u}\text{ admissible}} \sum_{k=0}^{N-1} \ell(x_{\mathbf{u}}(k,x),\mathbf{u}(k))$$

Boundedness assumption: there exists $\gamma > 0$ with

$$V_N(x) \le \gamma \ell^\star(x) \quad \text{for all } x \in \mathbb{X}, N \in \mathbb{N}$$

where $\ell^\star(x) := \min_{u\in\mathbb{U}} \ell(x,u)$

(sufficient conditions for and relaxations of this bound will be discussed later)

Stability and performance index

We choose ℓ, such that

$$\alpha_3(\|x - x_*\|) \le \ell^\star(x) \le \alpha_4(\|x - x_*\|)$$

holds for $\alpha_3, \alpha_4 \in \mathcal{K}_\infty$ (again, $\ell(x,u) = \|x - x_*\|^2 + \lambda\|u\|^2$ works)

Then, the only inequality left to prove in order to apply the relaxed dynamic programming theorem is

$$V_N(f(x,\mu_N(x))) \le V_N(x) - \alpha_N \ell(x,\mu_N(x))$$

for some $\alpha_N \in (0,1)$ and all $x \in \mathbb{X}$

We can compute α_N from the bound $V_N(x) \le \gamma \ell^\star(x)$

Computing α_N

We assume $\quad V_N(x) \le \gamma \ell^\star(x) \quad$ for all $\quad x \in \mathbb{X},\ N \in \mathbb{N} \quad (*)$

We want $\quad V_N(x^\star(1)) \le V_N(x^\star(0)) - \alpha_N \ell(x^\star(0),\mathbf{u}^\star(0))$

- use $(*)$ to find $\eta_N > 0$, $k^\star \ge 1$ with $\ell^\star(x^\star(k^\star)) \le \eta_N \ell^\star(x^\star(0))$
- concatenate $x^\star(1),\dots,x^\star(k^\star)$ and the optimal trajectory starting in $x^\star(k^\star)$ $\rightsquigarrow$ $\tilde{x}(\cdot)$, $\tilde{\mathbf{u}}(\cdot)$

$$\Rightarrow\ V_N(x^\star(1)) \le J_N(x^\star(1),\tilde{\mathbf{u}}) \le V_N(x^\star(0)) - \underbrace{(1-\gamma\eta_N)}_{=\alpha_N}\ell(x^\star(0),\mathbf{u}^\star(0))$$

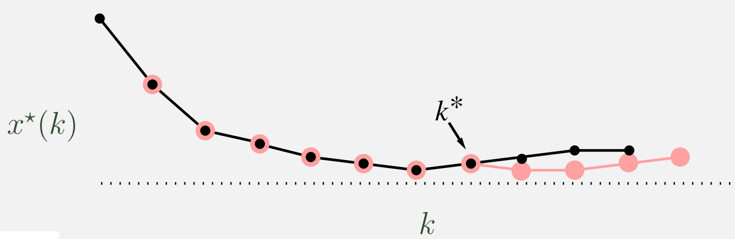

Decay of the optimal trajectory

We assume $\quad V_N(x) \le \gamma \ell^\star(x) \quad$ for all $\quad x \in \mathbb{X},\ N \in \mathbb{N}$

We want $\quad \eta_N > 0$, $k^\star \ge 1$ with $\ell^\star(x^\star(k^\star)) \le \eta_N \ell^\star(x^\star(0))$

Variant 1 [Grimm/Messina/Tuna/Teel '05]

$V_N(x) \le \gamma \ell^\star(x) \quad \Rightarrow \quad \ell(x^\star(k),u^\star(k)) \le \gamma\ell^\star(x)/N$ for at least one $k^\star \quad \Rightarrow \quad \alpha_N = 1 - \gamma(\gamma-1)/N$

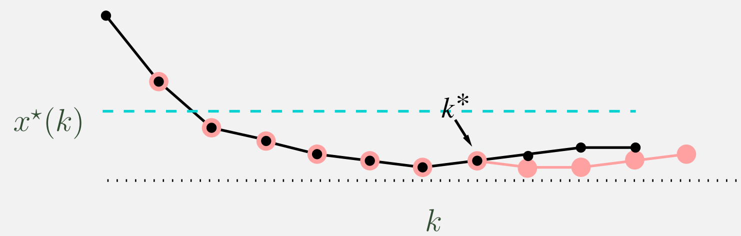

Decay of the optimal trajectory

We assume $V_N(x) \le \gamma \ell^\star(x)$ for all $x \in \mathbb{X}$, $N \in \mathbb{N}$

We want $\eta_N > 0$, $k^\star \ge 1$ with $\ell^\star(x^\star(k^\star)) \le \eta_N \ell^\star(x^\star(0))$

Variant 2 [Tuna/Messina/Teel '06, Gr./Rantzer '08]

$$V_N(x) \le \gamma \ell^\star(x) \quad \Rightarrow \quad \ell(x^\star(k), u^\star(k)) \le \gamma \left(\frac{\gamma-1}{\gamma}\right)^k \ell^\star(x)$$

$$\Rightarrow \quad k^\star = N-1 \quad \Rightarrow \quad \alpha_N = 1 - (\gamma-1)^N / \gamma^{N-2}$$

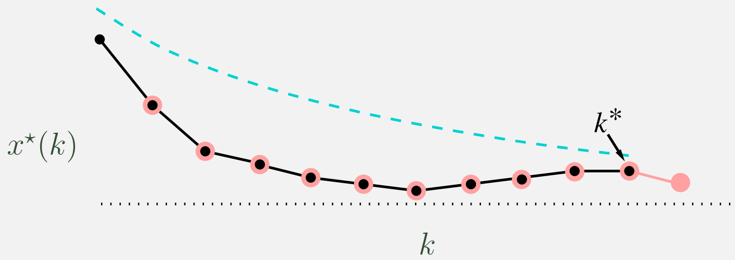

Decay of the optimal trajectory

We assume $V_N(x) \le \gamma \ell^\star(x)$ for all $x \in \mathbb{X}$, $N \in \mathbb{N}$

We want $\eta_N > 0$, $k^\star \ge 1$ with $\ell^\star(x^\star(k^\star)) \le \eta_N \ell^\star(x^\star(0))$

Variant 3 [Gr. '09, Gr./Pannek/Seehafer/Worthmann '10]

$V_N(x) \le \gamma \ell^\star(x) \quad \Rightarrow$ formulate all constraints and trajectories

$\Rightarrow$ optimize for $\alpha_N \quad \Rightarrow \quad \alpha_N = 1 - \frac{(\gamma-1)^N}{\gamma^{N-1} - (\gamma-1)^{N-2}}$

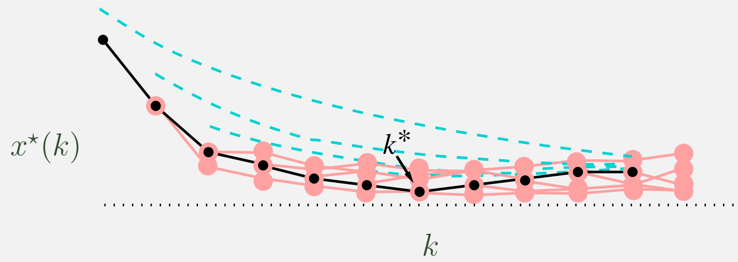

Optimization approach to compute α_N

We explain the optimization approach (Variant 3) in more detail. We want α_N such that

$$V_N(x^\star(1)) \le V_N(x^\star(0)) - \alpha_N \ell(x^\star(0), \mathbf{u}^\star(0))$$

holds for all optimal trajectories $x^\star(n), \mathbf{u}^\star(n)$ for V_N

The bound and the dynamic programming principle imply:

$$\begin{aligned}
V_N(x^\star(1)) &\le \gamma \ell^\star(x^\star(1)) \\
V_N(x^\star(1)) &\le \ell(x^\star(1), \mathbf{u}^\star(1)) + \gamma \ell^\star(x^\star(2)) \\
V_N(x^\star(1)) &\le \ell(x^\star(1), \mathbf{u}^\star(1)) + \ell(x^\star(2), \mathbf{u}^\star(2)) + \gamma \ell^\star(x^\star(3)) \\
\vdots \quad & \quad \vdots \quad \quad \vdots
\end{aligned}$$

Optimization approach to compute α_N

$\leadsto V_N(x^\star(1))$ is bounded by sums over $\ell(x^\star(n), \mathbf{u}^\star(n))$

For sums of these values, in turn, we get bounds from the dynamic programming principle and the bound:

$$\begin{aligned}
\sum_{n=0}^{N-1} \ell(x^\star(n), \mathbf{u}^\star(n)) &= V_N(x^\star(0)) &&\le \gamma \ell^\star(x^\star(0)) \\
\sum_{n=1}^{N-1} \ell(x^\star(n), \mathbf{u}^\star(n)) &= V_{N-1}(x^\star(1)) &&\le \gamma \ell^\star(x^\star(1)) \\
\sum_{n=2}^{N-1} \ell(x^\star(n), \mathbf{u}^\star(n)) &= V_{N-2}(x^\star(2)) &&\le \gamma \ell^\star(x^\star(2)) \\
&\vdots && \vdots
\end{aligned}$$

Verifying the relaxed Lyapunov inequality

Find α_N, such that for all optimal trajectories $x^\star$, $\mathbf{u}^\star$:

$$V_N(x^\star(1)) \le V_N(x^\star(0)) - \alpha_N \ell(x^\star(0), \mathbf{u}^\star(0)) \qquad (*)$$

Define $\quad \lambda_n := \ell(x^\star(n), \mathbf{u}^\star(n)), \quad \nu := V_N(x^\star(1))$

Then: $\quad (*) \;\Leftrightarrow\; \nu \le \sum_{n=0}^{N-1} \lambda_n - \alpha_N \lambda_0$

The inequalities from the last slides translate to

$$\sum_{n=k}^{N-1} \lambda_n \le \gamma \lambda_k, \quad k = 0, \ldots, N-2 \qquad (1)$$

$$\nu \le \sum_{n=1}^{j} \lambda_n + \gamma \lambda_{j+1}, \quad j = 0, \ldots, N-2 \qquad (2)$$

We call $\lambda_0, \ldots, \lambda_{N-1}, \nu \ge 0$ with (1), (2) admissible

Optimization problem

$\Rightarrow$ if α_N is such that the inequality

$$\nu \le \sum_{n=0}^{N-1} \lambda_n - \alpha_N \lambda_0 \;\Leftrightarrow\; \alpha_N \le \frac{\sum_{n=0}^{N-1} \lambda_n - \nu}{\lambda_0}$$

holds for all admissible λ_n and ν, then the desired inequality will hold for all optimal trajectories

The largest α_N satisfying this condition is

$$\alpha_N := \min_{\lambda_n, \nu \text{ admissible}} \frac{\sum_{n=0}^{N-1} \lambda_n - \nu}{\lambda_0}$$

This is a linear optimization problem whose solution can be computed explicitly (which is nontrivial) and reads

$$\alpha_N = 1 - \frac{(\gamma - 1)^N}{\gamma^{N-1} - (\gamma - 1)^{N-1}}$$

Stability and performance theorem

Theorem: [Gr./Pannek/Seehafer/Worthmann '10]: Assume $V_N(x) \le \gamma \ell^\star(x)$ for all $x \in \mathbb{X}$, $N \in \mathbb{N}$. If

$$\alpha_N > 0 \;\Leftrightarrow\; N > 2 + \frac{\ln(\gamma - 1)}{\ln \gamma - \ln(\gamma - 1)} \;\sim\; \gamma \ln \gamma$$

then the NMPC closed loop is asymptotically stable with Lyapunov function V_N and we get the performance estimate $J_\infty^{cl}(x, \mu_N) \le V_\infty(x)/\alpha_N$ with

$$\alpha_N = 1 - \frac{(\gamma - 1)^N}{\gamma^{N-1} - (\gamma - 1)^{N-1}} \to 1 \quad \text{as} \quad N \to \infty$$

Conversely, if $N < 2 + \frac{\ln(\gamma-1)}{\ln \gamma - \ln(\gamma-1)}$, then there exists a system for which $V_N(x) \le \gamma \ell^\star(x)$ holds but the NMPC closed loop is not asymptotically stable.

Horizon dependent γ-values

The theorem remains valid if we replace the bound condition

$$V_N(x) \le \gamma \ell^\star(x)$$

by

$$V_N(x) \le \gamma_N \ell^\star(x)$$

for horizon-dependent bounded values $\gamma_N \in \mathbb{R}$, $N \in \mathbb{N}$

$$\rightsquigarrow \quad \alpha_N = 1 - \frac{(\gamma_N - 1) \prod_{i=2}^{N} (\gamma_i - 1)}{\prod_{i=2}^{N} \gamma_i - \prod_{i=2}^{N} (\gamma_i - 1)}$$

This allows for tighter bounds and a refined analysis

Controllability condition

A refined analysis can be performed if we compute γ_N from a controllability condition, e.g., exponential controllability:

Assume that for each $x_0 \in \mathbb{X}$ there exists an admissible control $\mathbf{u}$ such that

$$\ell(x_{\mathbf{u}}(k), \mathbf{u}(k)) \leq C\sigma^k \ell^\star(x_0), \quad k = 0, 1, 2, \ldots$$

for given overshoot constant $C > 0$ and decay rate $\sigma \in (0, 1)$

$$\rightsquigarrow \quad V_N(x) \leq \gamma_N \ell^\star(x) \quad \text{for} \quad \gamma_N = \sum_{k=0}^{N-1} C\sigma^k$$

This allows to compute the minimal stabilizing horizon

$$\min\{N \in \mathbb{N} \,|\, \alpha_N > 0\}$$

depending on C and σ

Stability chart for C and σ

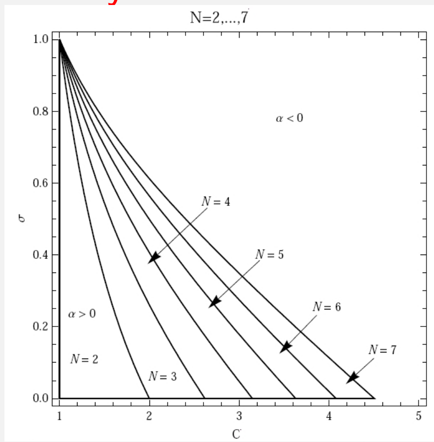

(Figure: Harald Voit)

Conclusion: for short optimization horizon N it is
more important: small C ("small overshoot")
less important: small σ ("fast decay")

(we will see at the end of the section how to use this information)

Comments and extensions

- for unconstrained linear quadratic problems:
  existence of γ $\Leftrightarrow$ (A, B) stabilizable
- additional weights on the last term can be incorporated into the analysis [Gr./Pannek/Seehafer/Worthmann '10]
- instead of using γ, α can be estimated numerically online along the closed loop [Pannek et al. '10ff]
- positive definiteness of ℓ can be replaced by a detectability condition [Grimm/Messina/Tuna/Teel '05]
- under appropriate uniformity assumptions, the results are easily carried over to tracking time variant references $x_{\text{ref}}(n)$ instead of an equilibrium x_* [Gr./Pannek '11]

Comments and extensions

The "linear" inequality $V_N(x) \leq \gamma\ell^\star(x)$ may be too demanding for nonlinear systems under constraints

Generalization: $\quad V_N(x) \leq \rho(\ell^\star(x)), \quad \rho \in \mathcal{K}_\infty$

- there is $\gamma > 0$ with $\rho(r) \leq \gamma r$ for all $r \in [0, \infty]$
  $\Rightarrow$ global asymptotic stability

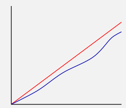

- for each $R > 0$
  there is $\gamma_R > 0$ with $\rho(r) \leq \gamma_R r$ for all $r \in [0, R]$
  $\Rightarrow$ semiglobal asymptotic stability

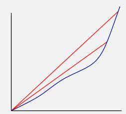

- $\rho \in \mathcal{K}_\infty$ arbitrary
  $\Rightarrow$ semiglobal practical asymptotic stability

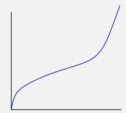

[Grimm/Messina/Tuna/Teel '05, Gr./Pannek '11]

The car-and-mountains example reloaded

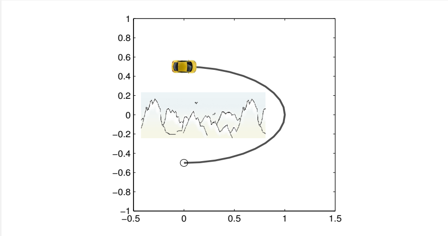

MPC with $\ell(x,u) = \|x - x_*\|^2 + |u|^2$ and $u_{\max} = 0.2$
$\rightsquigarrow$ asymptotic stability for $N = 11$ but not for $N \leq 10$

Reason: detour around mountains causes large overshoot C

Remedy: put larger weight on x_2:

$\ell(x,u) = (x_1 - x_{*,1})^2 + 5(x_2 - x_{*,2})^2 + |u|^2 \rightsquigarrow$ as. stab. for $N = 2$

A PDE example

We illustrate this with the 1d controlled PDE

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

with

domain $\Omega = [0,1]$
solution $y = y(t,x)$
boundary conditions $y(t,0) = y(t,1) = 0$
parameters $\nu = 0.1$ and $\mu = 10$
and distributed control $u : \mathbb{R} \times \Omega \to \mathbb{R}$

Discrete time system: $y(n) = y(nT, \cdot)$, sampling time $T = 0.025$

The uncontrolled PDE

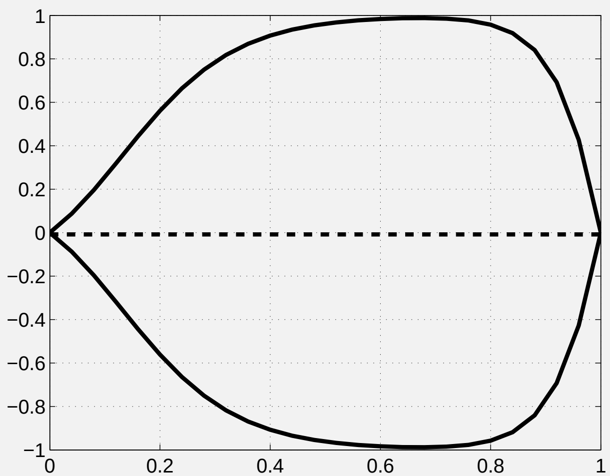

all equilibrium solutions

MPC for the PDE example

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

Goal: stabilize the sampled data system $y(n)$ at $y \equiv 0$

Usual approach: quadratic L^2 cost

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

For $y \approx 0$ the control u must compensate for $y_x \rightsquigarrow u \approx -y_x$

$\rightsquigarrow$ controllability condition

$$\begin{aligned} \ell(y(n), u(n)) &\leq C\sigma^n \ell^*(y(0)) \\ \Leftrightarrow \quad \|y(n)\|^2_{L^2} + \lambda\|u(n)\|^2_{L^2} &\leq C\sigma^n \|y(0)\|^2_{L^2} \\ \approx \quad \|y(n)\|^2_{L^2} + \lambda\|y_x(n)\|^2_{L^2} &\leq C\sigma^n \|y(0)\|^2_{L^2} \end{aligned}$$

for $\|y_x\|_{L^2} \gg \|y\|_{L^2}$ this can only hold if $C \gg 0$

MPC for the PDE example

Conclusion: because of

$$\|y(n)\|_{L^2}^2 + \lambda \|y_x(n)\|_{L^2}^2 \le C\sigma^n \|y(0)\|_{L^2}^2$$

the controllability condition may only hold for very large C

Remedy: use H^1 cost

$$\ell(y(n), u(n)) = \underbrace{\|y(n)\|_{L^2}^2 + \|y_x(n)\|_{L^2}^2}_{=\|y(n)\|_{H^1}^2} + \lambda \|u(n)\|_{L^2}^2.$$

Then an analogous computation yields

$$\|y(n)\|_{L^2}^2 + (1+\lambda)\|y_x(n)\|_{L^2}^2 \le C\sigma^n \Big(\|y(0)\|_{L^2}^2 + \|y_x(0)\|_{L^2}^2 \Big)$$

MPC with L_2 vs. H_1 cost

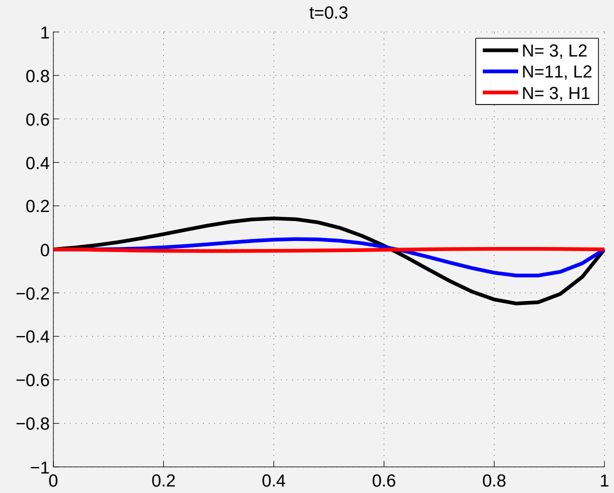

MPC with L_2 and H_1 cost, $\lambda = 0.1$, sampling time $T = 0.025$

Boundary Control

Now we change our PDE from distributed to (Dirichlet-) boundary control, i.e.

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

with

domain $\Omega = [0,1]$

solution $y = y(t,x)$

boundary conditions $y(t,0) = u_0(t)$, $y(t,1) = u_1(t)$

parameters $\nu = 0.1$ and $\mu = 10$

with boundary control, stability can only be achieved via large gradients in the transient phase
$\rightsquigarrow$ L^2 should perform better that H^1

Boundary control, L_2 vs. H_1, $N = 20$

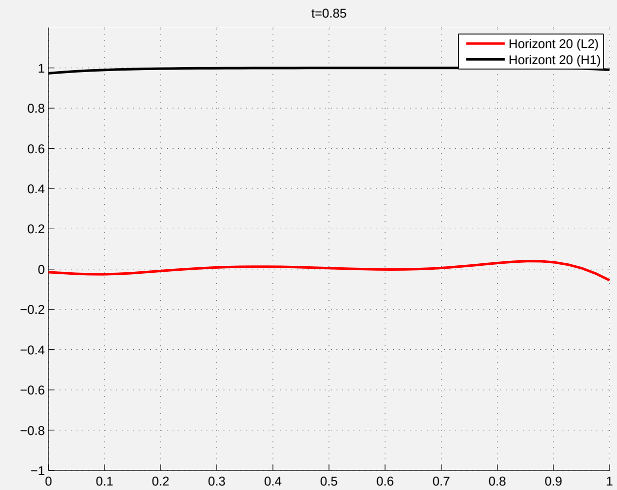

Boundary control, $\lambda = 0.001$, sampling time $T = 0.025$
Can be made rigorous for many PDEs [Altmüller et al. '10ff]

Summary of Section (4)

- Stability and performance of MPC without terminal constraints can be ensured by suitable bounds on V_N
- An optimization approach allows to compute the best possible α_N in the relaxed dynamic programming theorem
- The γ or γ_N can be computed from controllability properties, e.g., exponential controllability
- The overshoot bound $C > 0$ plays a crucial role or obtaining small stabilizing horizons
- Computing tight estimates for C is in general a difficult if not impossible task
- But structural knowledge of the system behavior can be sufficient for choosing a "good" ℓ

(5) Economic Model Predictive Control

Motivation for economic MPC

Typical approach in practice (e.g., in chemical process control):

(1) compute an economically good equilibrium (x_*, u_*)
("good" $=$ high yield, small energy consumption, etc.)

(2) design a controller stabilizing (x_*, u_*), e.g., by MPC

This works fine as long as the system state is close to x_* but on the way towards x_* performance in the sense of the chosen criterion may be bad

Idea: Use a stage cost ℓ which does not penalize the distance to some x_* but directly encodes the desired economic criterion

Mathematical difference of stabilizing and economic MPC

In stabilizing MPC, the stage cost $\ell(x, u)$ penalizes the distance to some equilibrium $(x_*, u_*) \in \mathbb{X} \times \mathbb{U}$. In particular, we required

$$\ell(x, u) > \ell(x_*, u_*) \quad \text{for all } (x, u) \in \mathbb{X} \times \mathbb{U}$$

In economic MPC, we remove this requirement. We use the same algorithm as in stabilizing MPC, but allow for more general ℓ to have more freedom to model economic objectives

We still consider equilibria, but they are now implicitly defined via the optimization criterion. In order to distinguish them from (x_*, u_*) in stabilizing MPC, they are denoted by (x^e, u^e)

Performance

As ℓ is no longer sign definite, we need to redefine how we measure the performance of μ_N

Infinite horizon averaged performance:

$$\overline{J}^{cl}_\infty(x,\mu_N) = \limsup_{K\to\infty} \frac{1}{K}\sum_{n=0}^{K-1} \ell(x_{\mu_N}(n,x),\mu_N(x_{\mu_N}(n,x)))$$

Finite horizon (or transient) performance:

$$J^{cl}_K(x,\mu_N) = \sum_{n=0}^{K-1} \ell(x_{\mu_N}(n,x),\mu_N(x_{\mu_N}(n,x)))$$

Only in special cases $K\to\infty$ makes sense

Example: minimum energy control

Example: Keep the state of the system inside the admissible set $\mathbb{X}$ minimizing the quadratic control effort

$$\ell(x,u) = u^2$$

with dynamics

$$x(n+1) = 2x(n) + \mathbf{u}(n)$$

and constraints $\mathbb{X} = [-2,2]$, $\mathbb{U} = [-3,3]$

For this example, a good strategy is to control the system to $x^e = 0$ and keep it there with $u^e = 0$

$\rightsquigarrow$ (x^e,u^e) is an optimal equilibrium with $\ell(x^e,u^e) = 0$

(recall: (x^e,u^e) equilibrium $\Leftrightarrow$ $f(x^e,u^e) = x^e$)

Example: trajectories

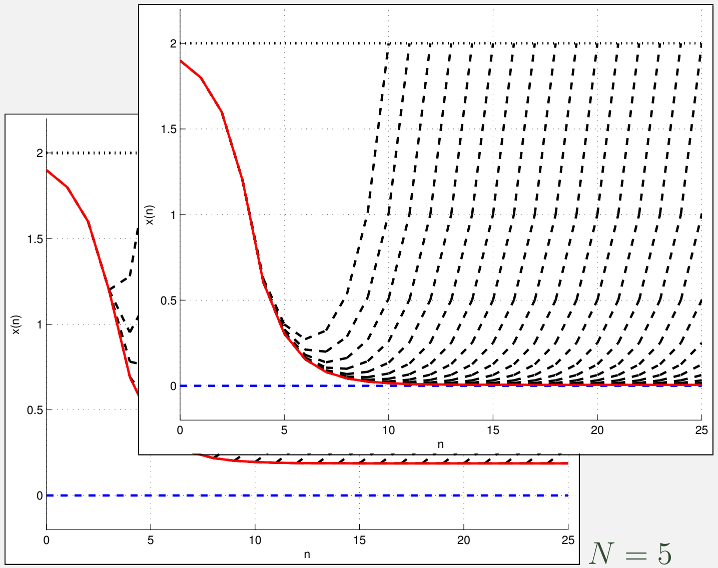

$N = 10$

$N = 5$

Example: averaged closed loop performance

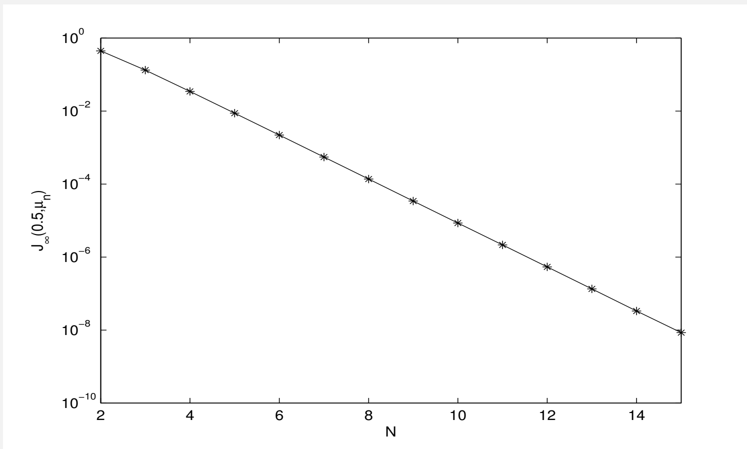

$\overline{J}^{cl}_\infty(0.5,\mu_N) - \ell(x^e,u^e)$ depending on N, logarithmic scale

Observations

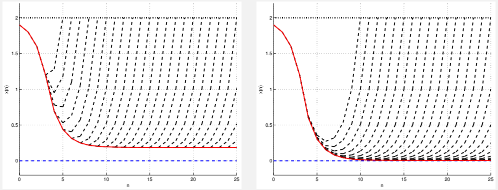

- optimal open loop trajectories approach the optimal equilibrium, stay near it for a while, and turn away – "turnpike property"
- closed loop trajectories converge to a neighborhood of the optimal equilibrium whose size tends to 0 as $N \to \infty$
- the averaged closed loop performance satisfies $\overline{J}^{cl}_{\infty}(x, \mu_N) \to \ell(x^e, u^e)$ as $N \to \infty$ (exponentially fast)

Can we prove this behavior?

The first property will turn out to be the crucial one

Towards a performance estimate

Defining the optimal value function $V_N(x) := \inf_{\mathbf{u}} J_N(x, \mathbf{u})$, the "trick" in all MPC proofs lies in relating V_N and V_{N-1}

In economic MPC, the desired inequality is

$$\underbrace{V_N(x)} \leq V_{N-1}(x) + \ell(x^e, u^e) + \varepsilon$$

$$\Rightarrow \quad \ell(x, \mu_N(x)) + V_{N-1}(f(x, \mu_N(x))) \leq V_{N-1}(x) + \ell(x^e, u^e) + \varepsilon$$

Using this inequality for $x = x_{\mu_N}(0), \ldots, x_{\mu_N}(K-1)$ yields

$$\begin{aligned} \overline{J}^{cl}_K(x, \mu_N) &= \frac{1}{K} \sum_{n=0}^{K-1} \ell(x_{\mu_N}(n), \mu_N(x_{\mu_N}(n))) \\ &\leq \frac{1}{K}(V_{N-1}(x_{\mu_N}(0)) - V_{N-1}(x_{\mu_N}(K))) + \ell(x^e, u^e) + \varepsilon \end{aligned}$$

$$\Rightarrow \quad \overline{J}^{cl}_{\infty}(x, \mu_N) = \limsup_{K \to \infty} \overline{J}^{cl}_K(x, \mu_N) \leq \ell(x^e, u^e) + \varepsilon$$

Towards a performance estimate

Similarly, estimates for the non averaged J^{cl}_K can be obtained

Hence, the desired inequality is

$$V_N(x) \leq V_{N-1}(x) + \ell(x^e, u^e) + \varepsilon$$

for a small $\varepsilon > 0$

In order to obtain this inequality, one

- takes an optimal trajectory corresponding to $V_{N-1}(x)$
- prolongs this trajectory such that its value increases by no more than $\ell(x^e, u^e) + \varepsilon$
- uses the resulting $J_N(x, u)$ as an upper bound for $V_N(x)$

This can be achieved by prolonging the trajectory close to x^e

Prolonging near x^e

Sketch of the idea:

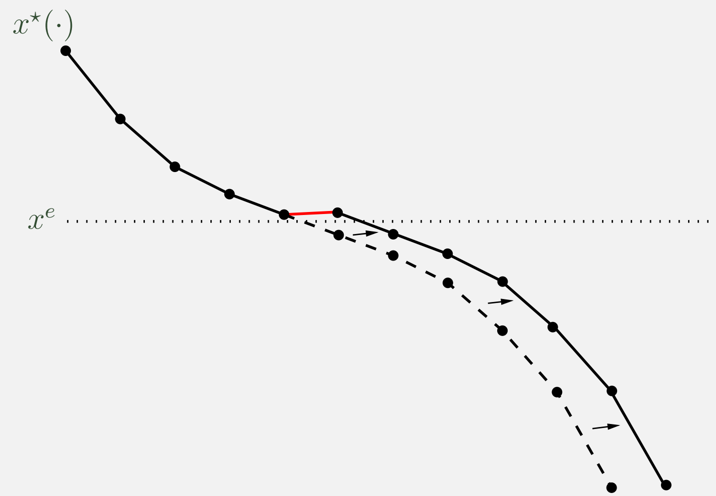

Assumptions needed for this construction

What do we need to make this construction work? [Gr. '13]

(1) Continuity of V_N near x^e (uniform in x and N)

- ensures that we can prolong the trajectory in the middle without changing the value of the tail too much

(2) Turnpike property

- ensures that the finite horizon optimal trajectories stay for a certain time near the optimal equilibrium x^e
- note: in numerical examples we often observe exponential turnpike, i.e., the minimum distance to x^e shrinks exponentially fast as N increases

Instead of the turnpike property, in the MPC literature another property is usually imposed: strict dissipativity

Strict dissipativity [Willems '72]

The optimal control problem is called strictly dissipative if there exists $\lambda : \mathbb{X} \to \mathbb{R}$ bounded from below and $\alpha \in \mathcal{K}_\infty$ with

$$\tilde{\ell}(x,u) := \ell(x,u) - \ell(x^e,u^e) + \lambda(x) - \lambda(f(x,u)) \geq \alpha(\|x - x^e\|)$$

for all $x \in \mathbb{X}$, $u \in \mathbb{U}$

While originally introduced as a sufficient condition guaranteeing the turnpike property, a recent result shows:

Theorem [Gr./Müller '16]: Under suitable controllability conditions, strict dissipativity is equivalent to a robust turnpike property plus optimality of the equilibrium (x^e, u^e)

The previous example is strictly dissipative with $\lambda(x) = -x^2/2$

Stabilizing functionals are strictly dissipative with $\lambda \equiv 0$

Economic MPC theorem

Theorem: [Gr./Stieler '14]

Let f and ℓ be Lipschitz, $\mathbb{X}$ and $\mathbb{U}$ be compact and assume

(i) local controllability near x^e

(ii) strict dissipativity

(iii) reachability of x^e from all $x \in \mathbb{X}$

(i), (ii) $\Rightarrow$ uniform continuity of V_N

(ii), (iii) $\Rightarrow$ turnpike property

(iv) polynomial growth conditions for $\tilde{\ell}$

(i)–(iv) $\Rightarrow$ exponential turnpike
[Damm/Gr./Stieler/Worthmann '14]
(for alternative conditions see also [Porretta/Zuazua '13]
[Trelat/Zuazua '14])

Economic MPC theorem

Under assumptions (i)–(iii), there exist $\varepsilon_1(N), \varepsilon_2(K) \to 0$ as $N \to \infty$ and $K \to \infty$, exponentially fast if additionally (iv) holds, such that the following properties hold

(1) Approximate average optimality:

$$\overline{J}^{cl}_\infty(x, \mu_N) \leq \ell(x^e, u^e) + \varepsilon_1(N)$$

(2) Practical asymptotic stability: there is $\beta \in \mathcal{KL}$:

$$\|x_{\mu_N}(k,x) - x^e\| \leq \beta(\|x - x^e\|, k) + \varepsilon_1(N) \text{ for all } k \in \mathbb{N}$$

(3) Approximate transient optimality: for all $K \in \mathbb{N}$:

$$J^{cl}_K(x, \mu_N(x)) \leq J_K(x, \mathbf{u}) + K\varepsilon_1(N) + \varepsilon_2(K)$$

for all admissible $\mathbf{u}$ with $\|x_{\mathbf{u}}(K,x) - x^e\| \leq \beta(\|x - x^e\|, K) + \varepsilon_1(N)$

Illustration of (2) and (3)

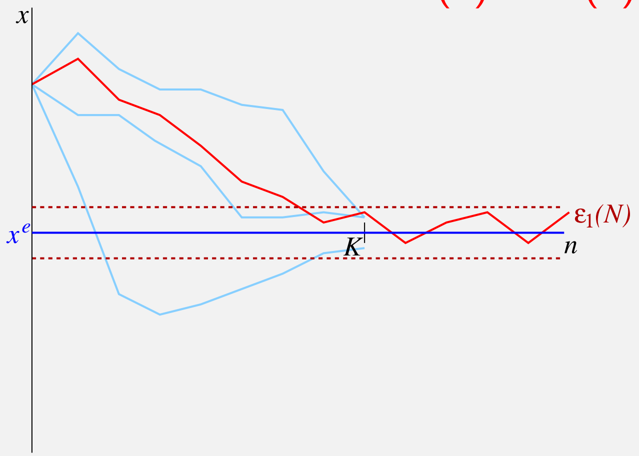

(2): $x_{\mu_N}(n)$ converges to the $\varepsilon_1(N)$-ball around x^e

(3): cost of all other trajectories reaching the ball at time K is higher than that of $x_{\mu_N}(n)$ up to the error $K\varepsilon_1(N) + \varepsilon_2(K)$

Economic MPC with terminal conditions

Like in stabilizing MPC, terminal conditions can be defined

The assumptions on the terminal constraint and cost are — somewhat surprisingly — identical to the stabilizing case [Amrit/Angeli/Rawlings '12]

We illustrate the effect of terminal condition for the previous example with the simplest choice of the terminal constraints $\mathbb{X}_0 = \{x^e\}$

Example with terminal conditions

Example:
$N = 5$

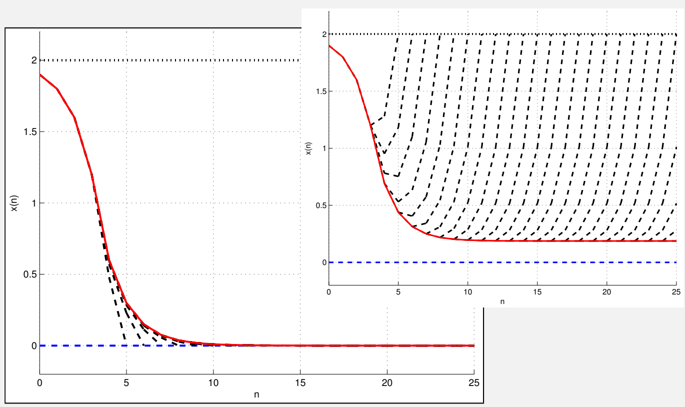

Schemes with terminal constraints

Imposing terminal conditions improves the previous results

Theorem: [Angeli/Amrit/Rawlings '12; Diehl/Rawlings '11] Under strict dissipativity and controllability, the resulting MPC scheme yields averaged optimal trajectories, i.e,

$$\overline{J}^{cl}_\infty(x, \mu_N) \le \ell(x^e, u^e) + \cancel{\varepsilon_1(N)}$$

for which x^e is asymptotically stable, i.e.,

$$\|x_{\mu_N}(k, x) - x^e\| \le \beta(\|x - x^e\|, k) + \cancel{\varepsilon_1(N)}.$$

In addition [Gr./Panin '15] we get approx. transient optimality

$$J^{cl}_K(x, \mu_N(x)) \le J_K(x, \mathbf{u}) + \cancel{K}\tilde{\varepsilon}_1(N) + \varepsilon_2(K)$$

Example: closed loop cost

$$J_K^{cl}(x,\mu_N(x)) \le J_K(x,\mathbf{u}) + K\varepsilon_1(N) + \varepsilon_2(K)$$

vs.

$$J_K^{cl}(x,\mu_N(x)) \le J_K(x,\mathbf{u}) + \tilde{\varepsilon}_1(N) + \varepsilon_2(K)$$

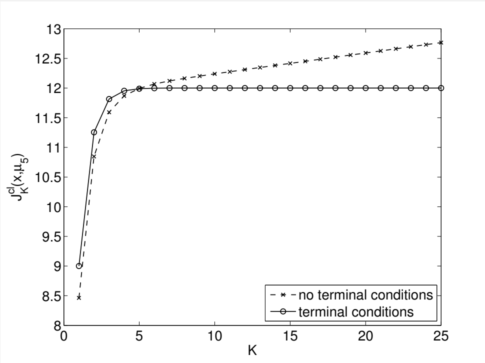

But: terminal constraints can cause infeasibility and severe numerical problems

Extensions, further results

- In the affine linear quadratic case our conditions are equivalent to the system being stabilizable [Gr./Stieler '14]
- The optimal equilibrium can be replaced by an optimal periodic orbit [Zanon/Gr. '16, Müller/Gr. '16]
- The results can be formulated directly in continuous time [Faulwasser/Bonvin '15, Alessandretti/Aguiar/Jones '15]
- First results for time varying systems are available [Zanon/Gros/Diehl '13, Alessandretti/Aguiar/Jones '15]
- First results for discounted optimal control problems [Gr./Semmler/Stieler '15, Gr./Kellett/Weller '16]

Summary of Section (5)

- Model predictive control works not only for stabilizing problems
- The turnpike property can be used as a replacement for positive definiteness of ℓ
- Strict dissipativity is essentially equivalent to this property and may be used as a checkable condition
- The required uniform continuity may pose a problem for control systems governed by PDEs $\rightsquigarrow$ work in progress

Selected literature

D.Q. Mayne, J.B. Rawlings, C.V. Rao, P.O.M. Scokaert, Constrained model predictive control: stability and optimality, Automatica, 36(2000), 789–814 ("The" classical reference for the results from (3))

L. Grüne and J. Pannek, Nonlinear Model Predictive Control, Springer, 2011 (contains most of the material from (2)–(4))

L. Grüne, NMPC without terminal constraints, Proceedings of the IFAC Conference on Nonlinear Model Predictive Control, 2012, 1–13 (survey with some results from (4) and (5))

D. Angeli, R. Amrit, J.B. Rawlings, On average performance and stability of economic model predictive control, IEEE Trans. Autom. Control, 57 (2012), 1615–1626 (some results from (5))

L. Grüne and M. Stieler, Asymptotic stability and transient optimality of economic MPC without terminal conditions, Journal of Process Control, 24 (2014), 1187–1196 (some results from (5))

L. Grüne, Approximation properties of receding horizon optimal control, DMV Nachrichten, 118 (2016), 3–38 (survey with results from Section (5) and variants)