On conditions under which receding horizon control delivers approximately optimal feedback strategies

Lars Grüne

Mathematisches Institut, Universität Bayreuth

theory based on joint work with
Marleen Stieler (Bayreuth), Matthias Müller & Frank Allgöwer (Stuttgart)
Anastasia Panin (Bayreuth), Karl Worthmann (Ilmenau)

application based on joint work with
Arthur Fleig (Bayreuth), Roberto Guglielmi (Linz)

supported by DFG, Elitenetzwerk Bayern and Marie-Curie ITN SADCO

International Workshop “From Open to Closed Loop Control”
Mariatrost, June 22–26, 2015

Setup

We consider nonlinear discrete time control systems

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

with $x_{\mathbf{u}}(n) \in X$, $\mathbf{u}(n) \in U$, $\quad X$, U normed spaces

Setup

We consider nonlinear discrete time control systems

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

with $x_{\mathbf{u}}(n) \in X$, $\mathbf{u}(n) \in U$, $\quad X$, U normed spaces

Usual interpretation:

$x_{\mathbf{u}}(n) \quad = \quad$ state of the system at time t_n

Setup

We consider nonlinear discrete time control systems

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

with $x_{\mathbf{u}}(n) \in X$, $\mathbf{u}(n) \in U$, X, U normed spaces

Usual interpretation:

$x_{\mathbf{u}}(n)$ = state of the system at time t_n

$\mathbf{u}(n)$ = control acting from time t_n to t_{n+1}

Setup

We consider nonlinear discrete time control systems

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

with $x_{\mathbf{u}}(n) \in X$, $\mathbf{u}(n) \in U$, $\quad X$, U normed spaces

Usual interpretation:

$x_{\mathbf{u}}(n)$	$=$	state of the system at time t_n
$\mathbf{u}(n)$	$=$	control acting from time t_n to t_{n+1}
f	$=$	solution operator of a controlled ODE/PDE or of a discrete time model

Setup

We consider nonlinear discrete time control systems

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

with $x_{\mathbf{u}}(n) \in X$, $\mathbf{u}(n) \in U$, $\quad X$, U normed spaces

Usual interpretation:

$x_{\mathbf{u}}(n)$	$=$	state of the system at time t_n
$\mathbf{u}(n)$	$=$	control acting from time t_n to t_{n+1}
f	$=$	solution operator of a controlled ODE/PDE or of a discrete time model (or a numerical approximation of one of these)

Setup

We consider nonlinear discrete time control systems

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

with $x_{\mathbf{u}}(n) \in X$, $\mathbf{u}(n) \in U$, $\quad X$, U normed spaces

Setup

We consider nonlinear discrete time control systems

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

with $x_{\mathbf{u}}(n) \in X$, $\mathbf{u}(n) \in U$, $\quad X$, U normed spaces

Problem: infinite horizon optimal control

Setup

We consider nonlinear discrete time control systems

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

with $x_{\mathbf{u}}(n) \in X$, $\mathbf{u}(n) \in U$, $\quad X$, U normed spaces

Problem: infinite horizon optimal control

Prototype problem: For a stage cost $\ell : X \times U \to \mathbb{R}$ solve

$$\text{minimize } J_\infty(x, \mathbf{u}) = \sum_{n=0}^{\infty} \ell(x_{\mathbf{u}}(n), \mathbf{u}(n))$$

Setup

We consider nonlinear discrete time control systems

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

with $x_{\mathbf{u}}(n) \in X$, $\mathbf{u}(n) \in U$, X, U normed spaces

Problem: infinite horizon optimal control

Prototype problem: For a stage cost $\ell : X \times U \to \mathbb{R}$ solve

$$\text{minimize } J_\infty(x, \mathbf{u}) = \sum_{n=0}^{\infty} \ell(x_{\mathbf{u}}(n), \mathbf{u}(n))$$

subject to state/control constraints $x_{\mathbf{u}}(n) \in \mathbb{X}$, $\mathbf{u}(n) \in \mathbb{U}$

Setup

We consider nonlinear discrete time control systems

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

with $x_{\mathbf{u}}(n) \in X$, $\mathbf{u}(n) \in U$, $\quad X$, U normed spaces

Problem: infinite horizon optimal control

Prototype problem: For a stage cost $\ell : X \times U \to \mathbb{R}$ solve

$$\text{minimize } J_\infty(x, \mathbf{u}) = \sum_{n=0}^{\infty} \ell(x_{\mathbf{u}}(n), \mathbf{u}(n))$$

subject to state/control constraints $x_{\mathbf{u}}(n) \in \mathbb{X}$, $\mathbf{u}(n) \in \mathbb{U}$

with optimal control in feedback form $\mathbf{u}(n) = \mu(x_{\mathbf{u}}(n))$

Receding horizon control

Direct solution of the problem is numerically hard

Alternative method: receding horizon or model predictive control (MPC)

Receding horizon control

Direct solution of the problem is numerically hard

Alternative method: receding horizon or model predictive control (MPC)

Idea: replace the infinite horizon problem

$$\text{minimize } J_\infty(x,\mathbf{u}) = \sum_{n=0}^{\infty} \ell(x_\mathbf{u}(n),\mathbf{u}(n))$$

by the iterative solution of finite horizon problems

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

with fixed $N \in \mathbb{N}$ and $x_\mathbf{u}(k) \in \mathbb{X}$, $\mathbf{u}(k) \in \mathbb{U}$

Receding horizon control

Direct solution of the problem is numerically hard

Alternative method: receding horizon or model predictive control (MPC)

Idea: replace the infinite horizon problem

$$\text{minimize } J_\infty(x,\mathbf{u}) = \sum_{n=0}^{\infty} \ell(x_\mathbf{u}(n),\mathbf{u}(n))$$

by the iterative solution of finite horizon problems

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

with fixed $N \in \mathbb{N}$ and $x_\mathbf{u}(k) \in \mathbb{X}$, $\mathbf{u}(k) \in \mathbb{U}$

We obtain a feedback law μ_N by a receding horizon technique

MPC from the trajectory point of view

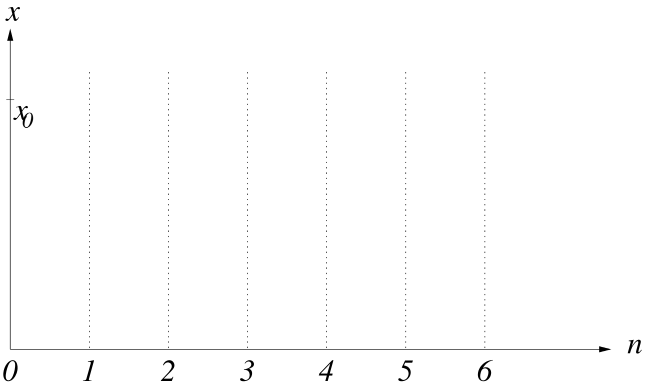

MPC from the trajectory point of view

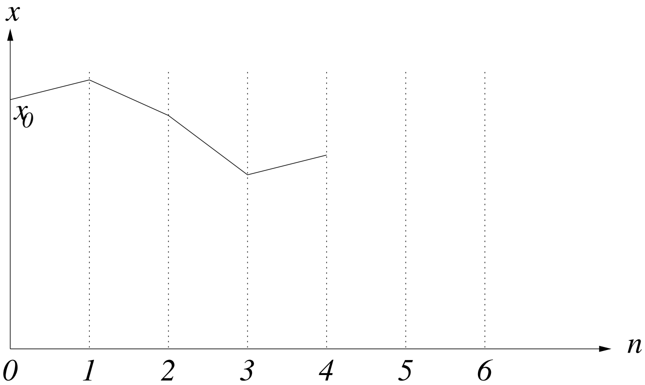

black = predictions (open loop optimization)

MPC from the trajectory point of view

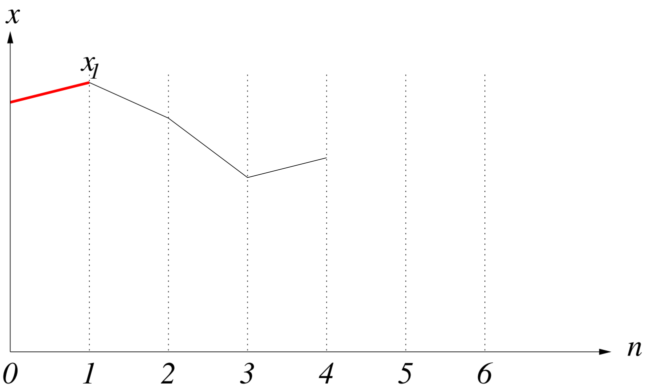

black = predictions (open loop optimization)

red = MPC closed loop

MPC from the trajectory point of view

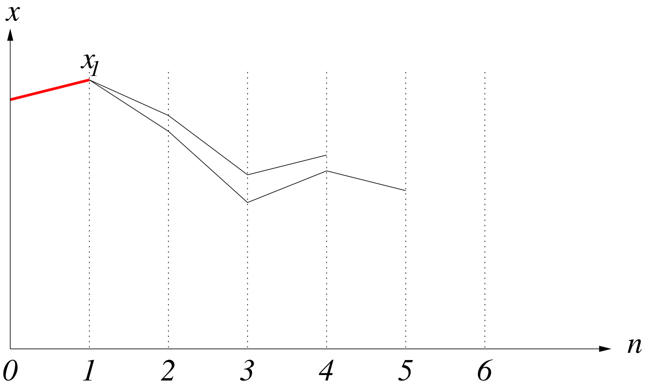

black = predictions (open loop optimization)

red = MPC closed loop

MPC from the trajectory point of view

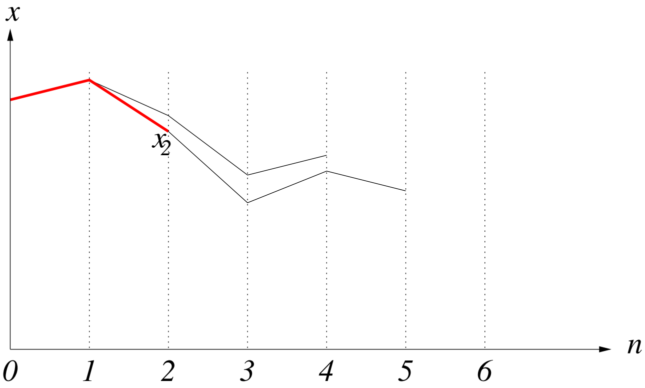

black = predictions (open loop optimization)

red = MPC closed loop

MPC from the trajectory point of view

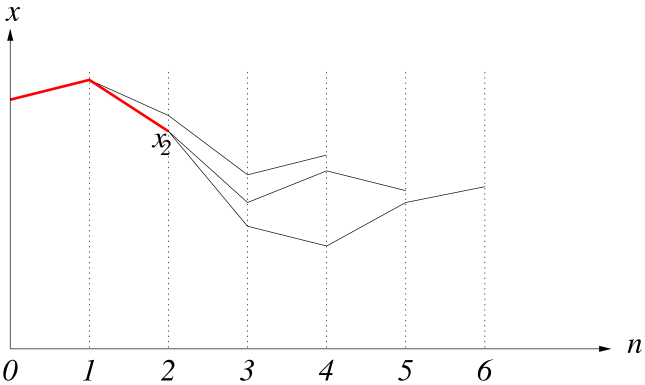

black = predictions (open loop optimization)

red = MPC closed loop

MPC from the trajectory point of view

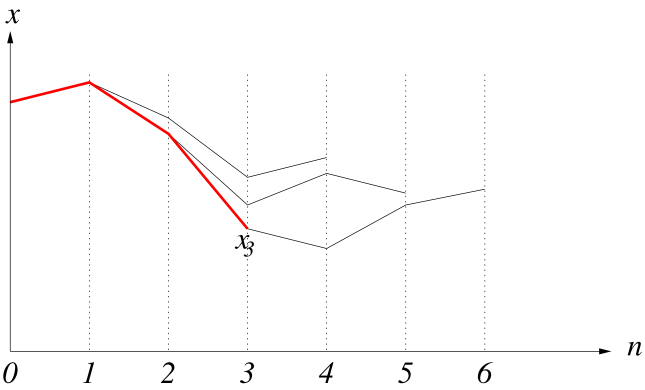

black = predictions (open loop optimization)

red = MPC closed loop

MPC from the trajectory point of view

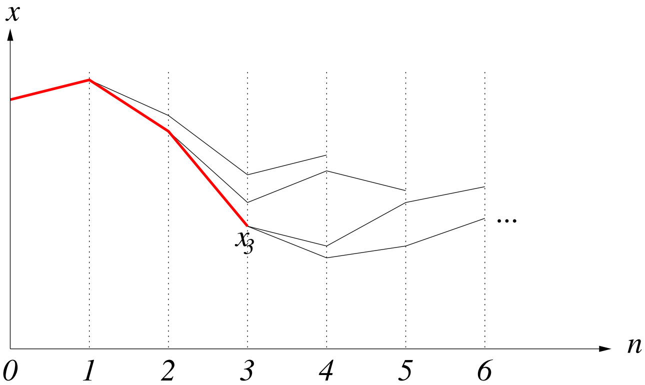

black = predictions (open loop optimization)

red = MPC closed loop

MPC from the trajectory point of view

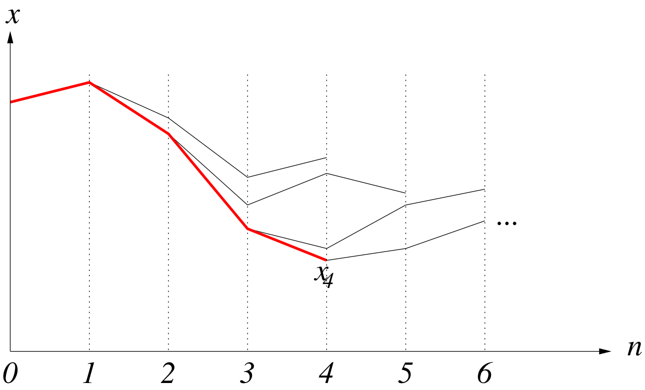

black = predictions (open loop optimization)

red = MPC closed loop

MPC from the trajectory point of view

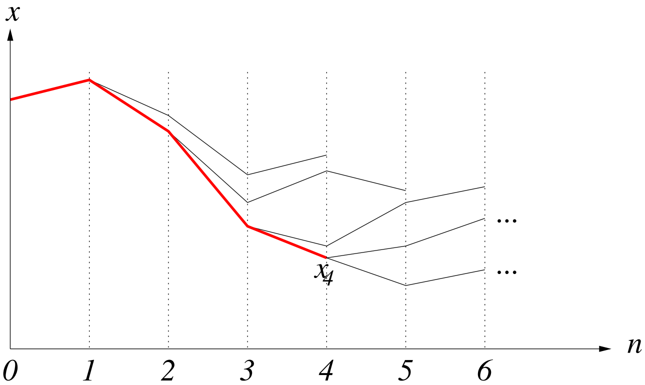

black = predictions (open loop optimization)

red = MPC closed loop

MPC from the trajectory point of view

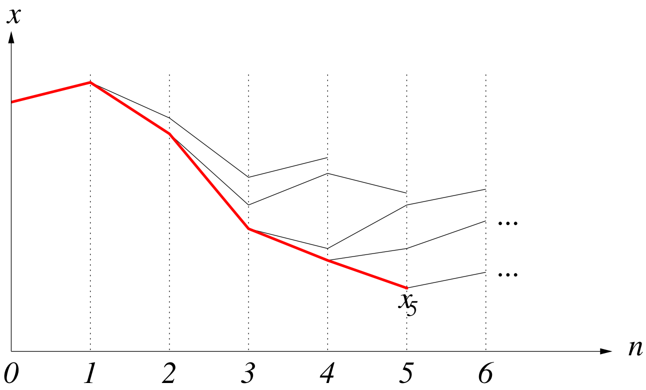

black = predictions (open loop optimization)

red = MPC closed loop

MPC from the trajectory point of view

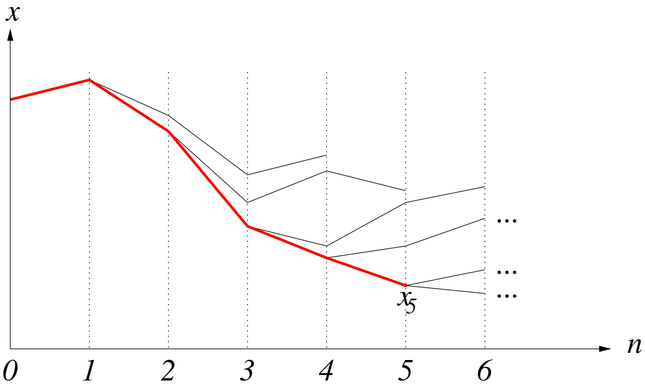

black = predictions (open loop optimization)

red = MPC closed loop

MPC from the trajectory point of view

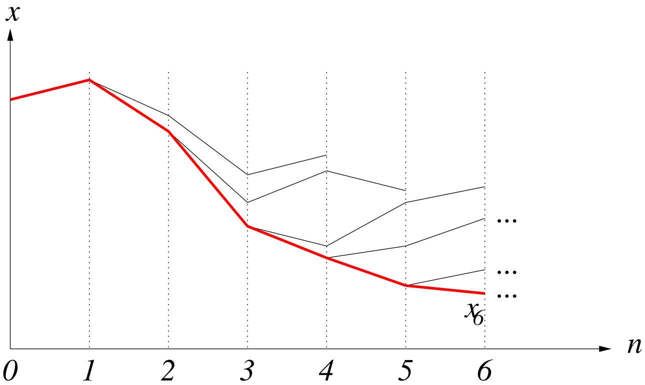

black = predictions (open loop optimization)

red = MPC closed loop

MPC from the trajectory point of view

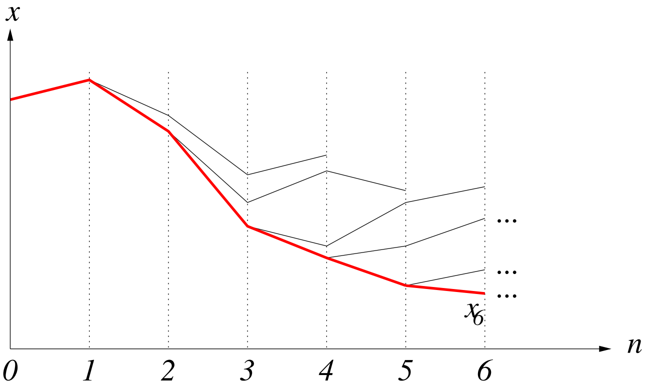

black = predictions (open loop optimization)

red = MPC closed loop $x_{\mu_N}(n, x_0)$

Model predictive control

Basic receding horizon MPC concept:

Model predictive control

Basic receding horizon MPC concept:

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

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

subject to the constraints $x_u(k) \in \mathbb{X}$, $\mathbf{u}(k) \in \mathbb{U}$

Model predictive control

Basic receding horizon MPC concept:

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

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

subject to the constraints $x_u(k) \in \mathbb{X}$, $\mathbf{u}(k) \in \mathbb{U}$

$\rightsquigarrow$ optimal trajectory $\quad x^\star(0), \ldots, x^\star(N-1)$

Model predictive control

Basic receding horizon MPC concept:

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

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

subject to the constraints $x_u(k) \in \mathbb{X}$, $\mathbf{u}(k) \in \mathbb{U}$

⇝ optimal trajectory $x^\star(0), \ldots, x^\star(N-1)$

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

Model predictive control

Basic receding horizon MPC concept:

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

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

subject to the constraints $x_u(k) \in \mathbb{X}$, $\mathbf{u}(k) \in \mathbb{U}$

⇝ optimal trajectory $x^\star(0), \ldots, x^\star(N-1)$

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

⇝ MPC feedback law $\mu_N(x_{\mu_N}(n)) := \mathbf{u}^\star(0)$

Model predictive control

Basic receding horizon MPC concept:

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

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

subject to the constraints $x_u(k) \in \mathbb{X}$, $\mathbf{u}(k) \in \mathbb{U}$

⇝ optimal trajectory $x^\star(0), \ldots, x^\star(N-1)$

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

⇝ MPC feedback law $\mu_N(x_{\mu_N}(n)) := \mathbf{u}^\star(0)$

⇝ closed loop system

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

Why use MPC?

What is the advantage of MPC over other methods of solving optimal control problems?

Why use MPC?

What is the advantage of MPC over other methods of solving optimal control problems?

- significantly reduced computational complexity

Why use MPC?

What is the advantage of MPC over other methods of solving optimal control problems?

- significantly reduced computational complexity
  $\rightsquigarrow$ real time capability

Why use MPC?

What is the advantage of MPC over other methods of solving optimal control problems?

- significantly reduced computational complexity
  $\rightsquigarrow$ real time capability
- ability to react to perturbations

Why use MPC?

What is the advantage of MPC over other methods of solving optimal control problems?

- significantly reduced computational complexity
  $\rightsquigarrow$ real time capability
- ability to react to perturbations
- applicability to problems in which data becomes available online

Why use MPC?

What is the advantage of MPC over other methods of solving optimal control problems?

- significantly reduced computational complexity
  $\rightsquigarrow$ real time capability
- ability to react to perturbations
- applicability to problems in which data becomes available online

But: The trajectory delivered by MPC can be far from optimal!

Why use MPC?

What is the advantage of MPC over other methods of solving optimal control problems?

- significantly reduced computational complexity
  $\rightsquigarrow$ real time capability
- ability to react to perturbations
- applicability to problems in which data becomes available online

But: The trajectory delivered by MPC can be far from optimal!

$\rightsquigarrow$ Key question in this talk: When does MPC yield closed loop trajectories with approximately optimal performance?

Why use MPC?

What is the advantage of MPC over other methods of solving optimal control problems?

- significantly reduced computational complexity
  $\rightsquigarrow$ real time capability
- ability to react to perturbations
- applicability to problems in which data becomes available online

But: The trajectory delivered by MPC can be far from optimal!

$\rightsquigarrow$ Key question in this talk: When does MPC yield closed loop trajectories with approximately optimal performance?

First question: How to define performance?

Performance

In this talk we do not want to limit ourselves to tracking type functionals, i.e., $\ell(x,u) = \|x - x_*\|^2 + \lambda\|u - u_*\|^2$

Performance

In this talk we do not want to limit ourselves to tracking type functionals, i.e., $\ell(x,u) = \|x - x_*\|^2 + \lambda\|u - u_*\|^2$

MPC with more general ℓ is often termed economic MPC. In this setting, performance of μ_N can be measured in two ways

Performance

In this talk we do not want to limit ourselves to tracking type functionals, i.e., $\ell(x,u) = \|x - x_*\|^2 + \lambda\|u - u_*\|^2$

MPC with more general ℓ is often termed economic MPC. In this setting, performance of μ_N can be measured in two ways

Infinite horizon averaged performance:

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

Performance

In this talk we do not want to limit ourselves to tracking type functionals, i.e., $\ell(x,u) = \|x - x_*\|^2 + \lambda\|u - u_*\|^2$

MPC with more general ℓ is often termed economic MPC. In this setting, performance of μ_N can be measured in two ways

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:

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

Performance

In this talk we do not want to limit ourselves to tracking type functionals, i.e., $\ell(x,u) = \|x - x_*\|^2 + \lambda\|u - u_*\|^2$

MPC with more general ℓ is often termed economic MPC. In this setting, performance of μ_N can be measured in two ways

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

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 $u^e = 0$

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 $\ell(x^e, u^e) = 0$

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

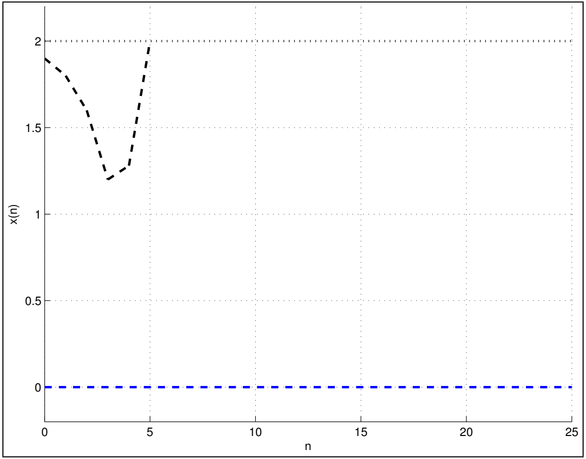

$N = 5$

Example: trajectories

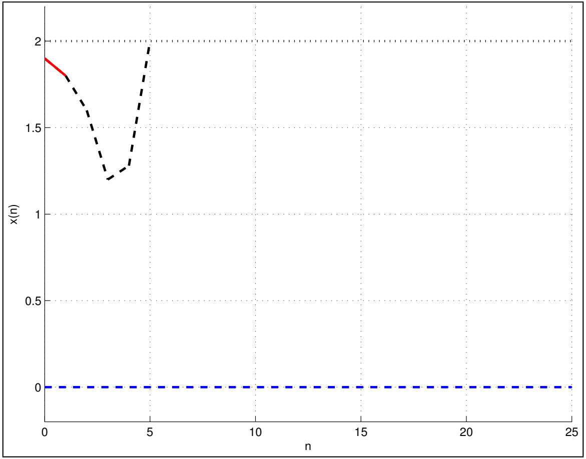

$N = 5$

Example: trajectories

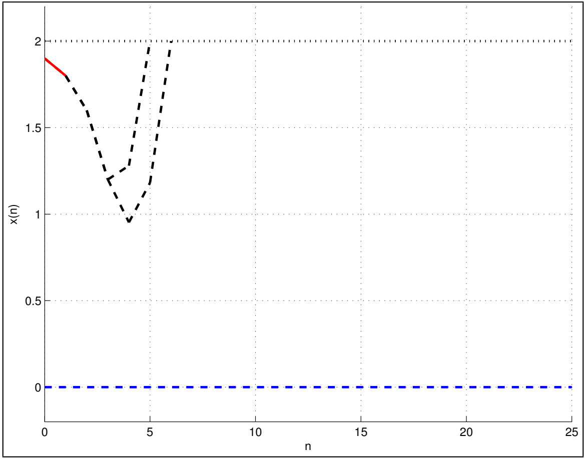

$N = 5$

Example: trajectories

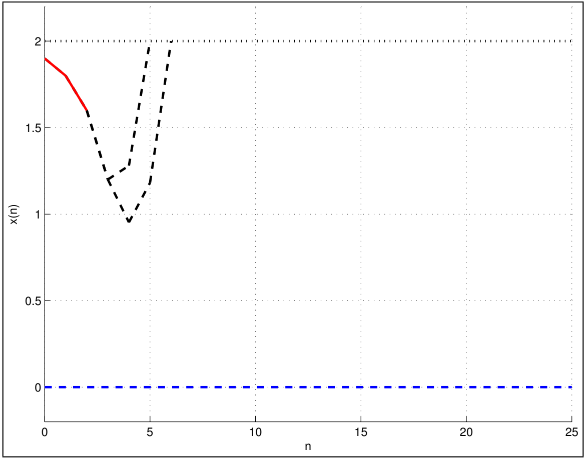

$N = 5$

Example: trajectories

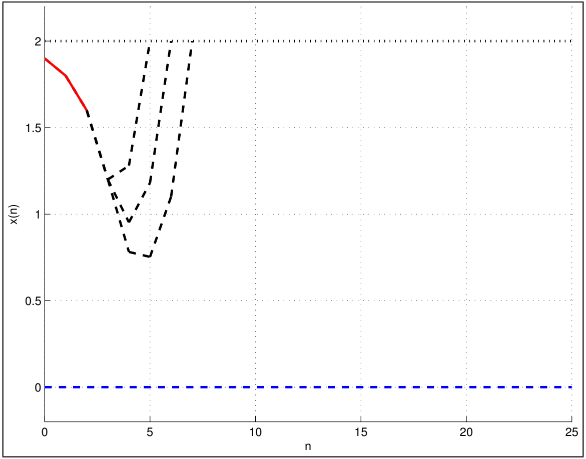

$N = 5$

Example: trajectories

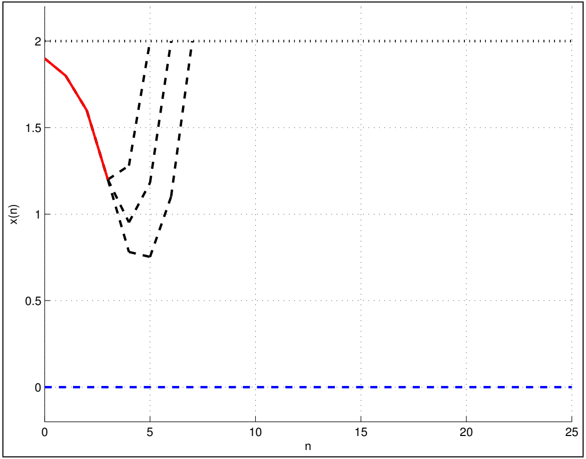

$N = 5$

Example: trajectories

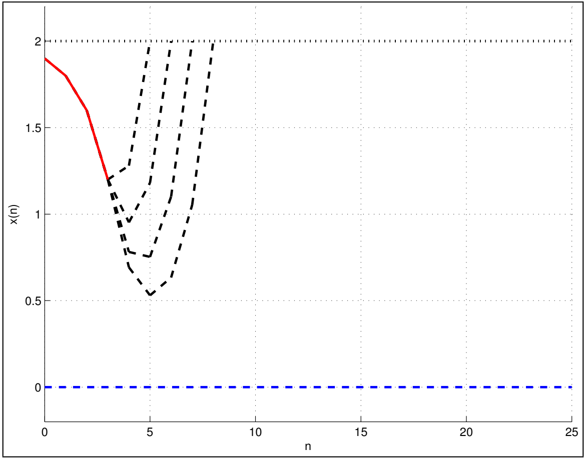

$N = 5$

Example: trajectories

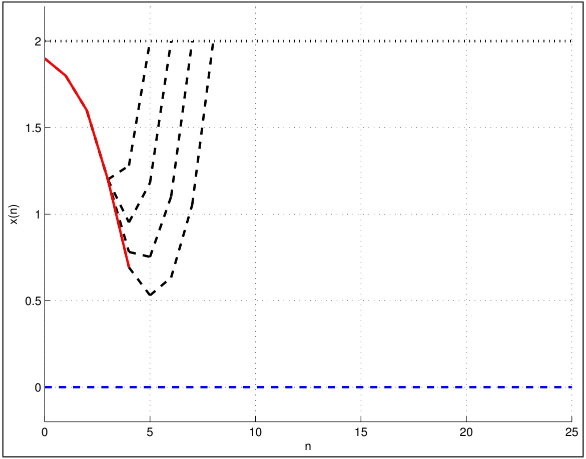

$N = 5$

Example: trajectories

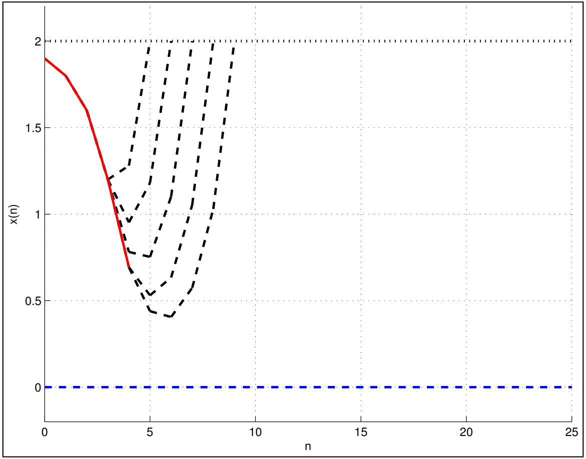

$N = 5$

Example: trajectories

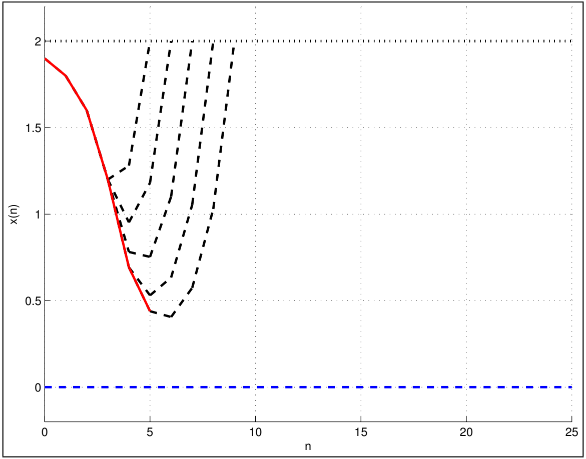

$N = 5$

Example: trajectories

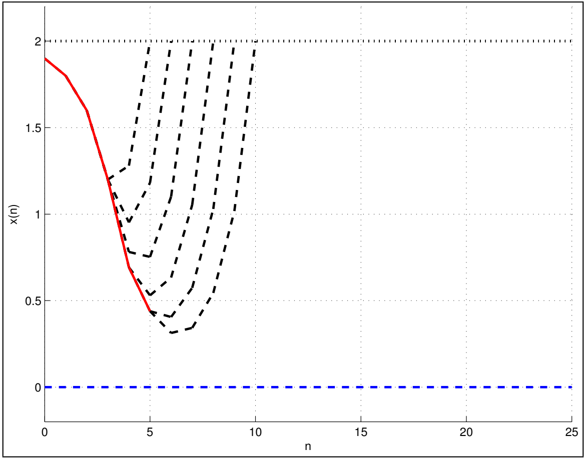

$N = 5$

Example: trajectories

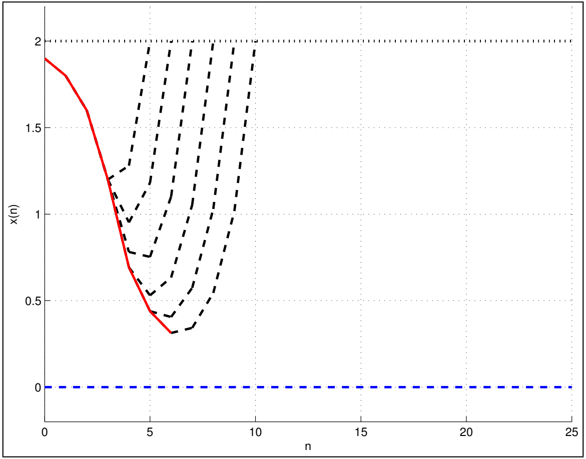

$N = 5$

Example: trajectories

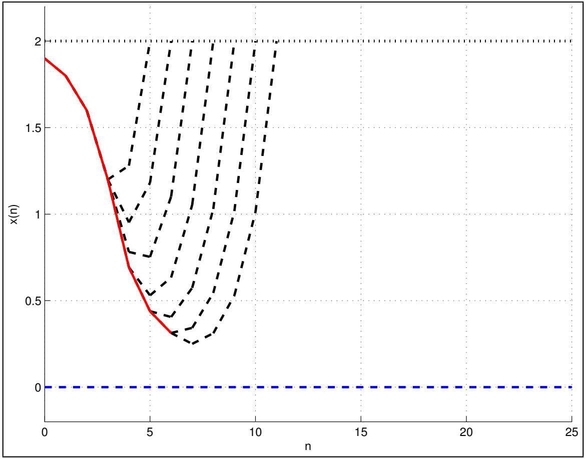

$N = 5$

Example: trajectories

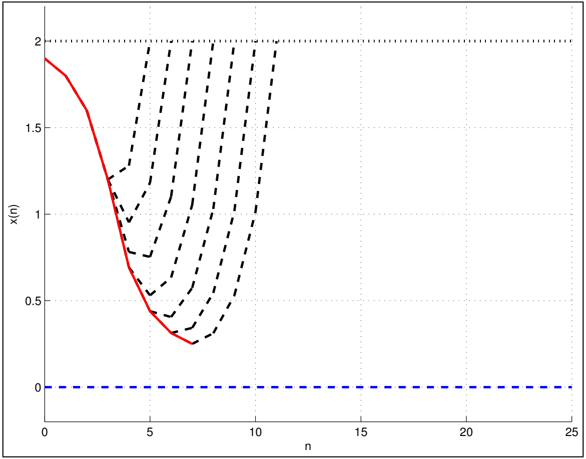

$N = 5$

Example: trajectories

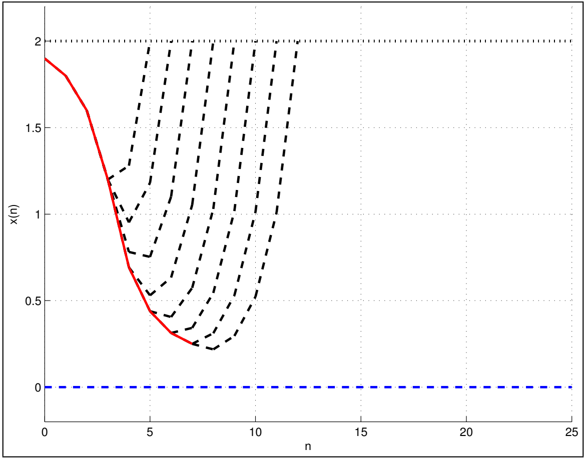

$N = 5$

Example: trajectories

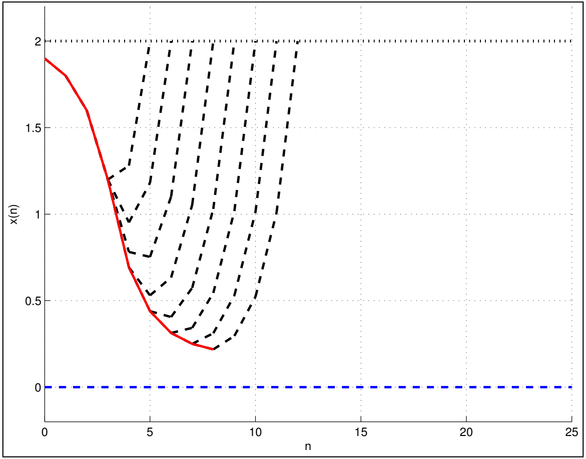

$N = 5$

Example: trajectories

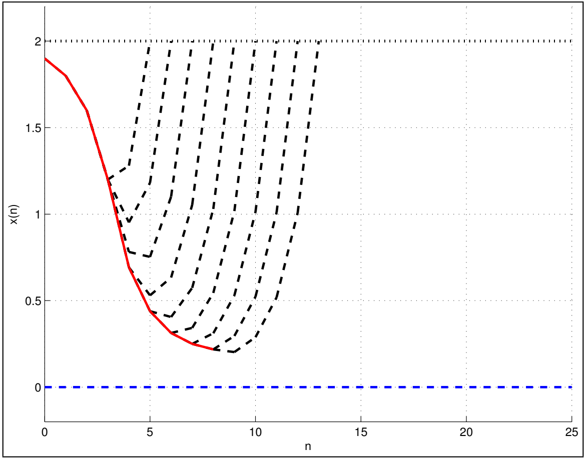

$N = 5$

Example: trajectories

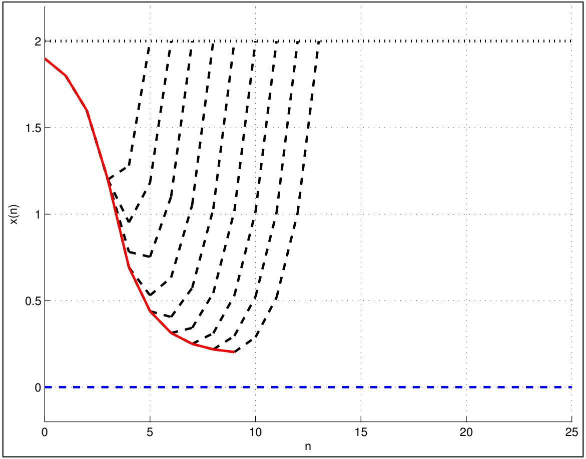

$N = 5$

Example: trajectories

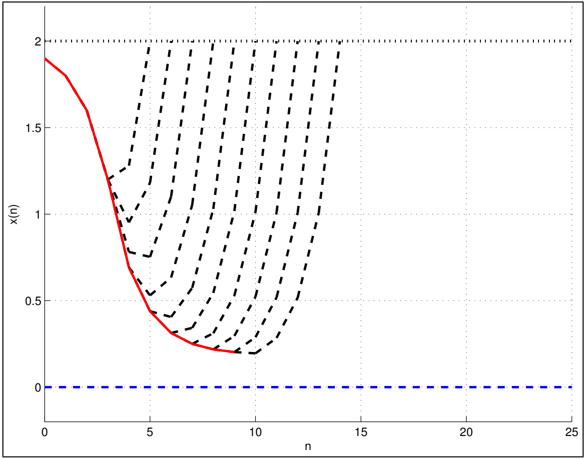

$N = 5$

Example: trajectories

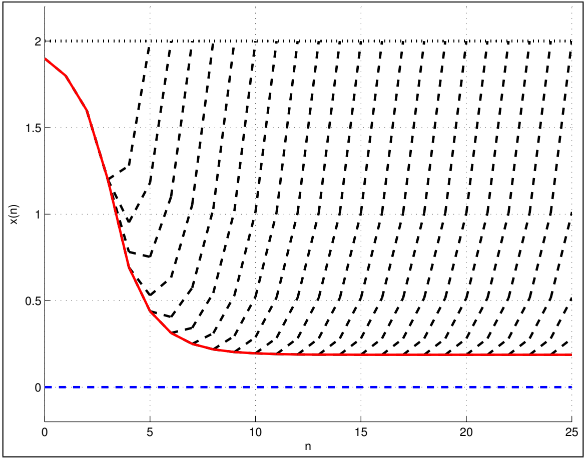

$N = 5$

Example: trajectories

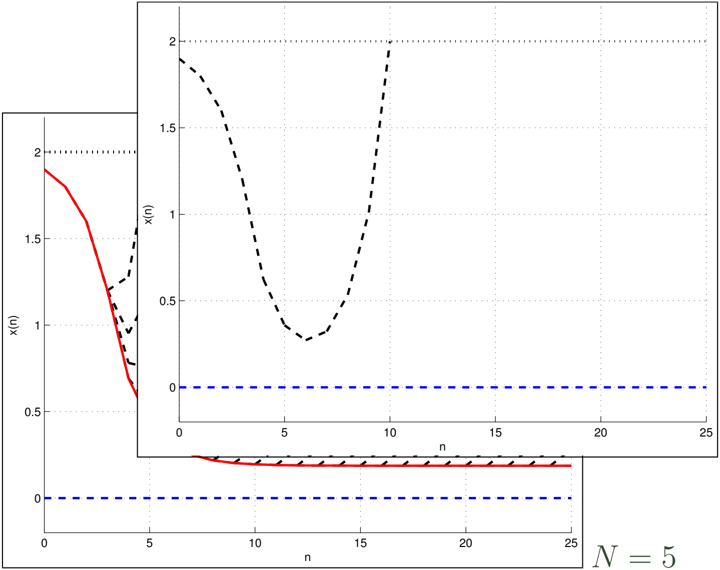

$N = 10$

$N = 5$

Example: trajectories

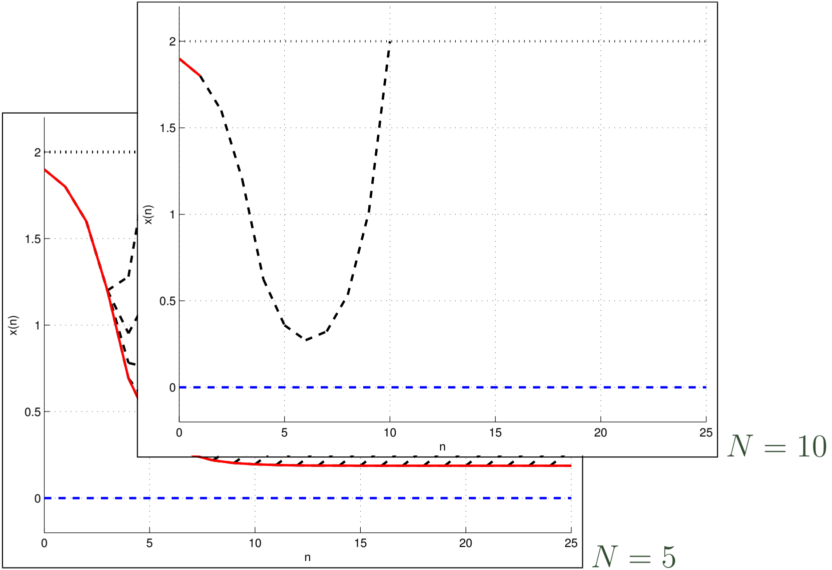

Example: trajectories

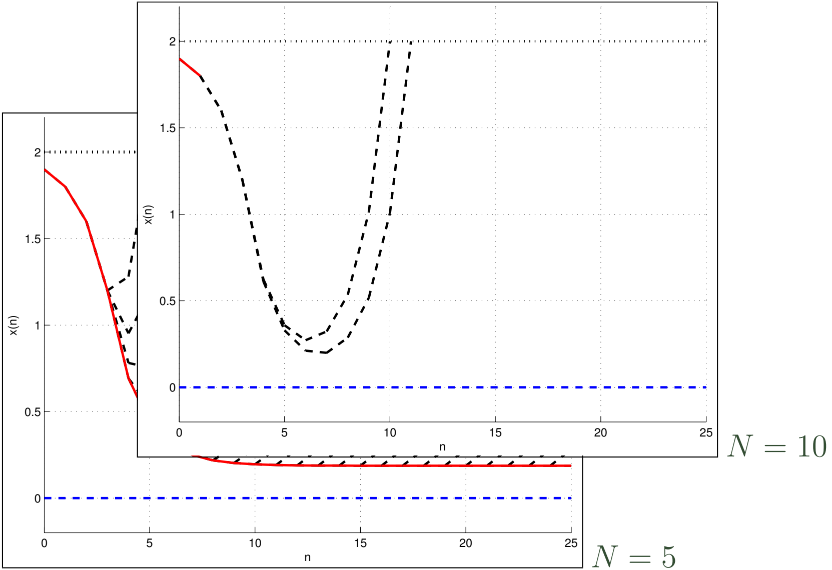

Example: trajectories

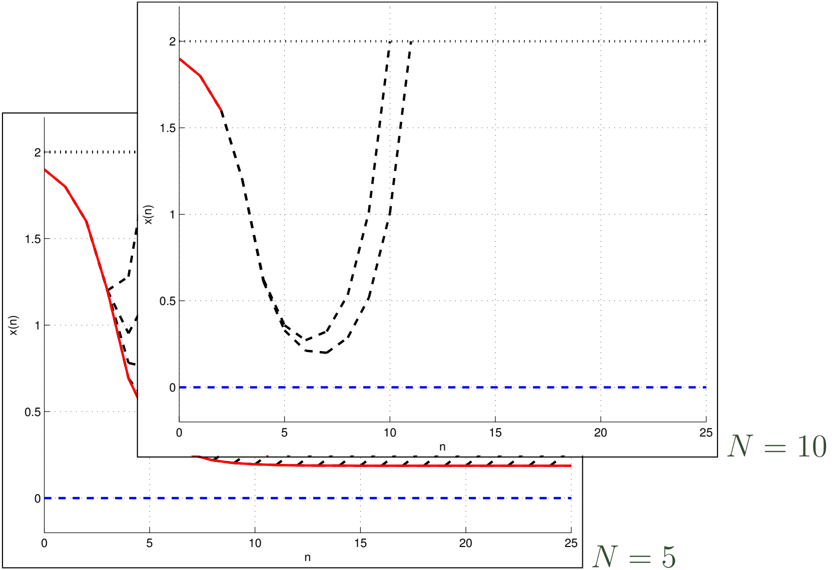

UNIVERSITÄT BAYREUTH

Example: trajectories

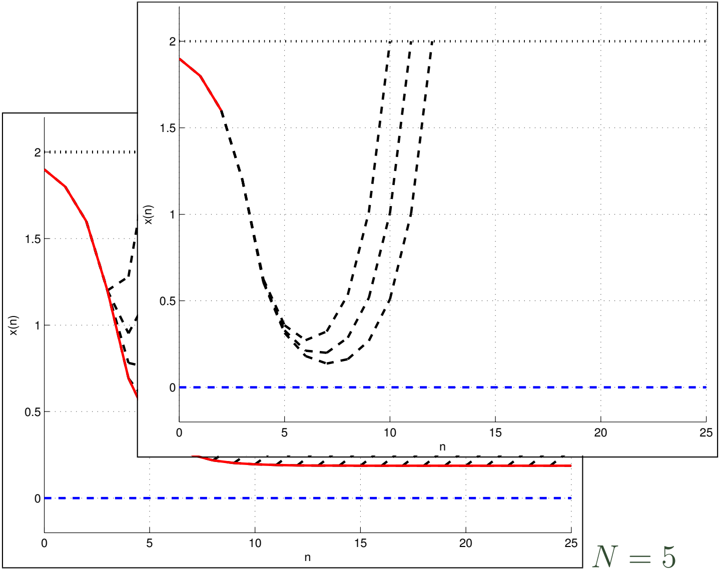

$N = 10$

$N = 5$

Example: trajectories

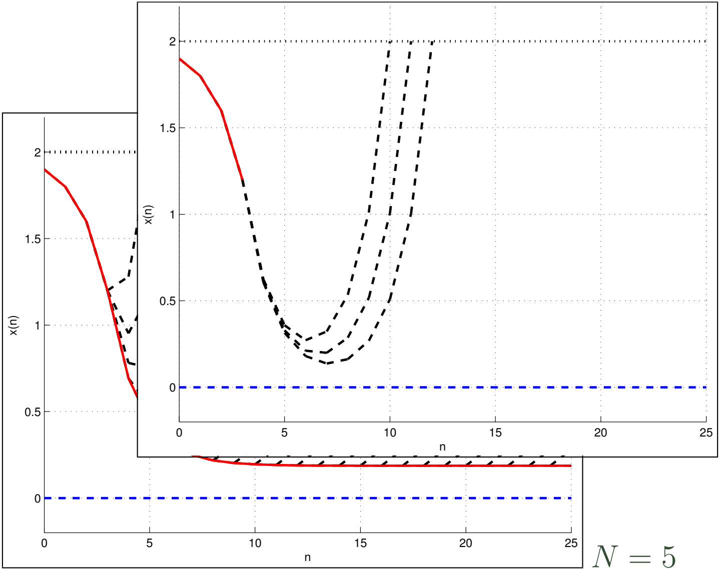

$N = 10$

$N = 5$

Example: trajectories

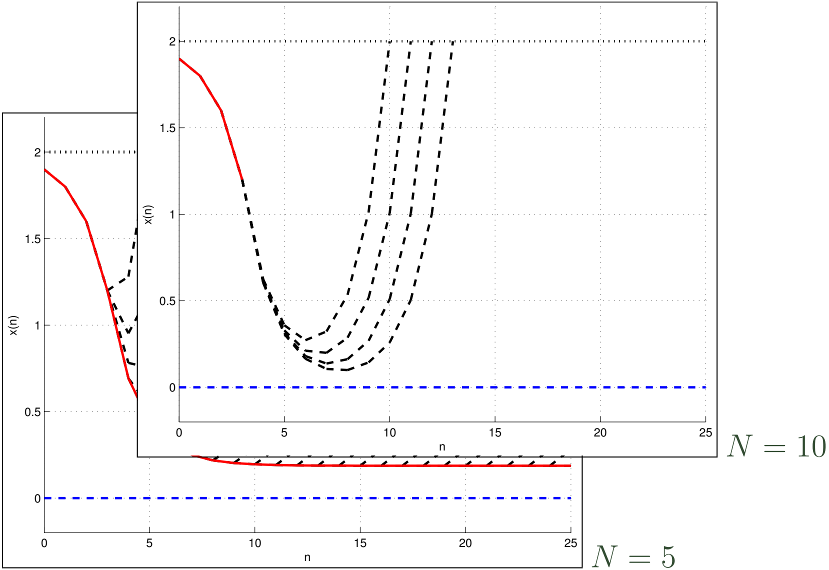

Example: trajectories

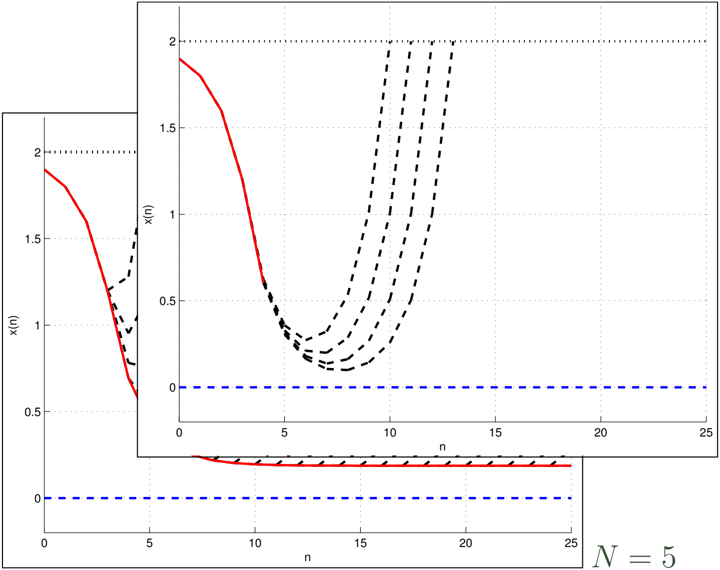

$N = 10$

$N = 5$

Example: trajectories

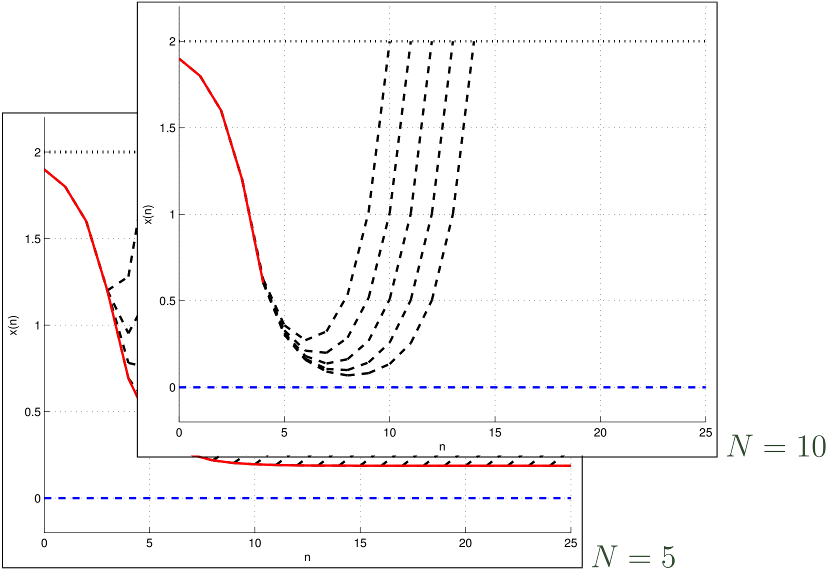

Example: trajectories

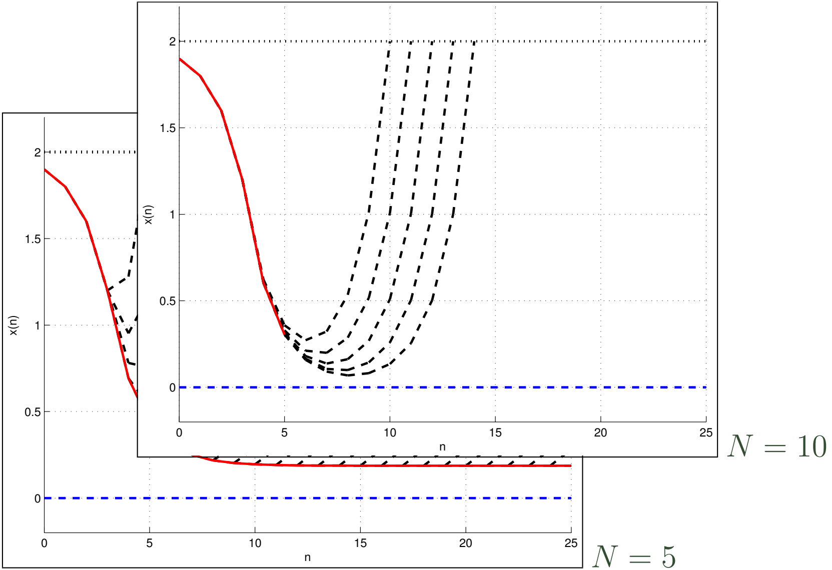

Example: trajectories

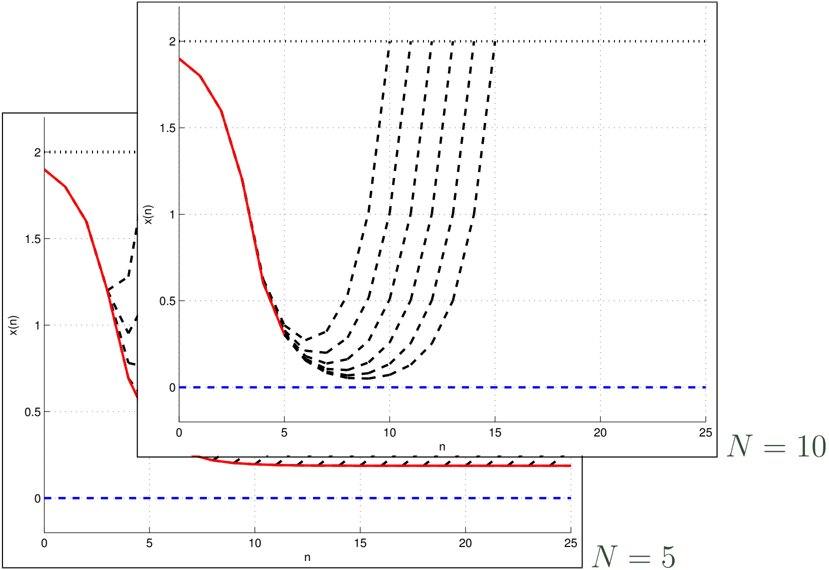

Example: trajectories

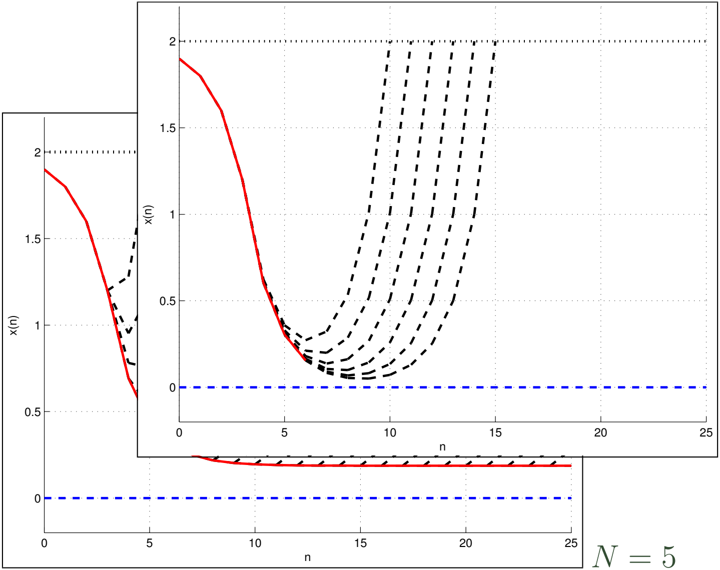

$N = 10$

$N = 5$

Example: trajectories

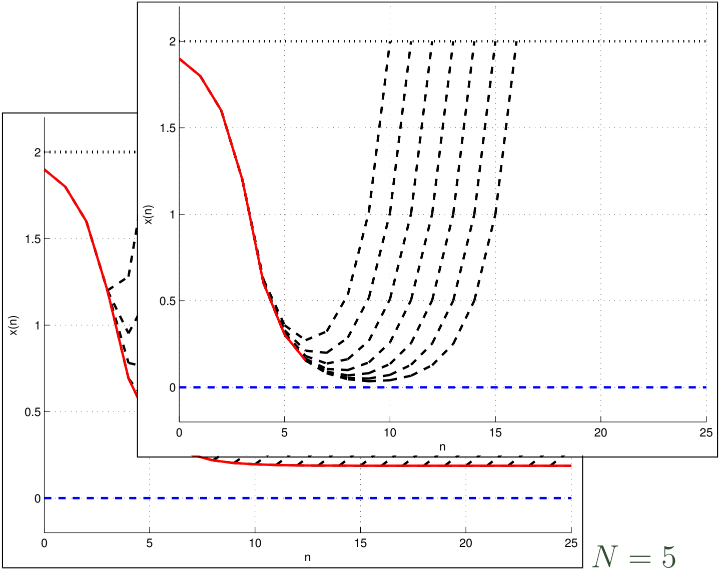

$N = 10$

$N = 5$

Example: trajectories

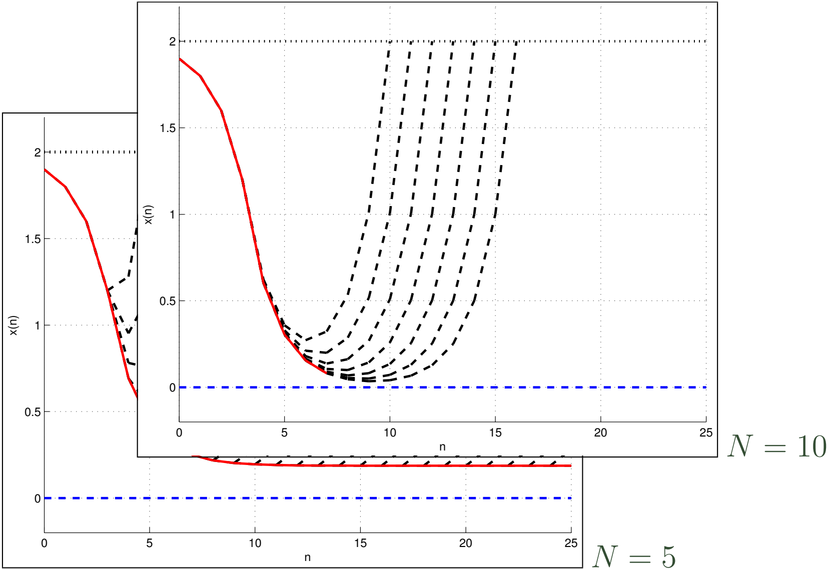

Example: trajectories

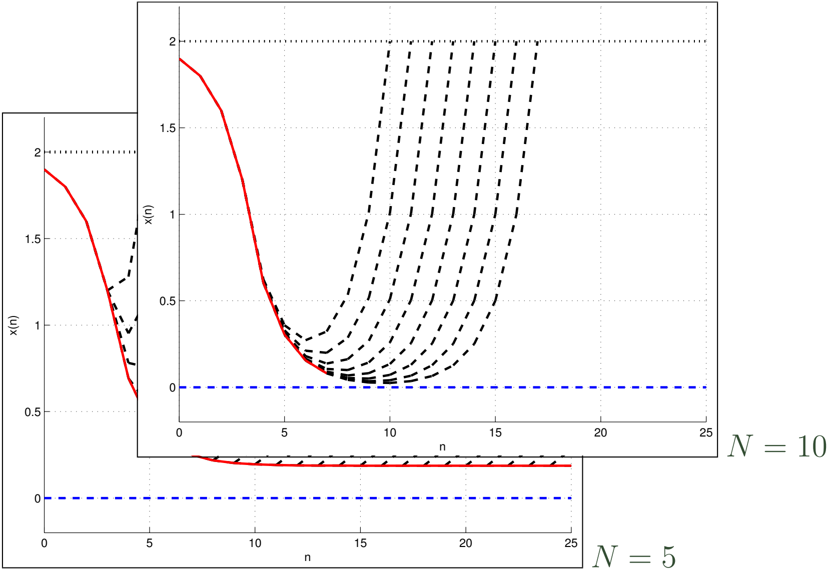

Example: trajectories

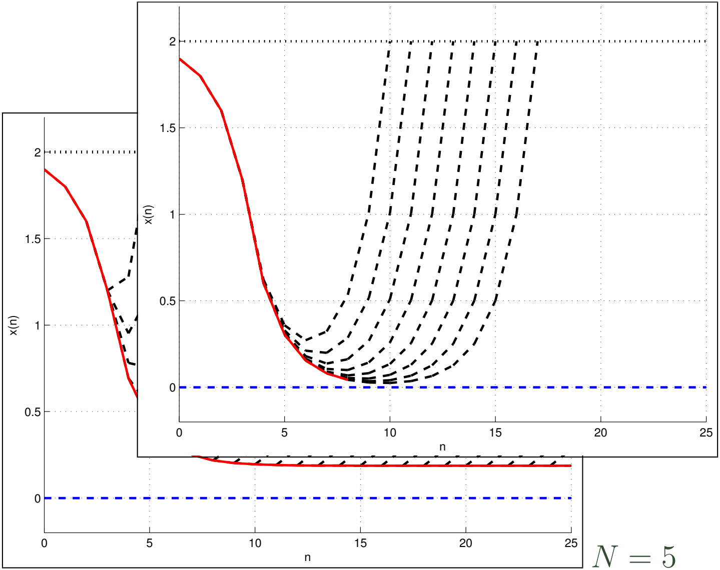

$N = 10$

$N = 5$

Example: trajectories

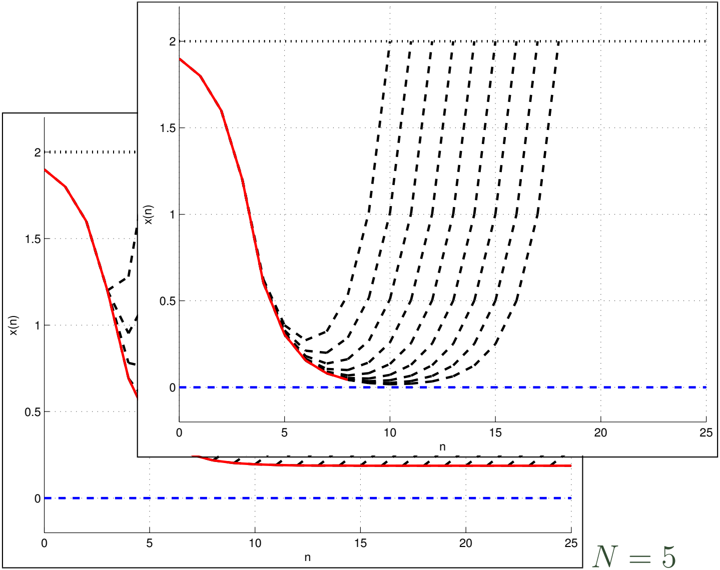

$N = 10$

$N = 5$

Example: trajectories

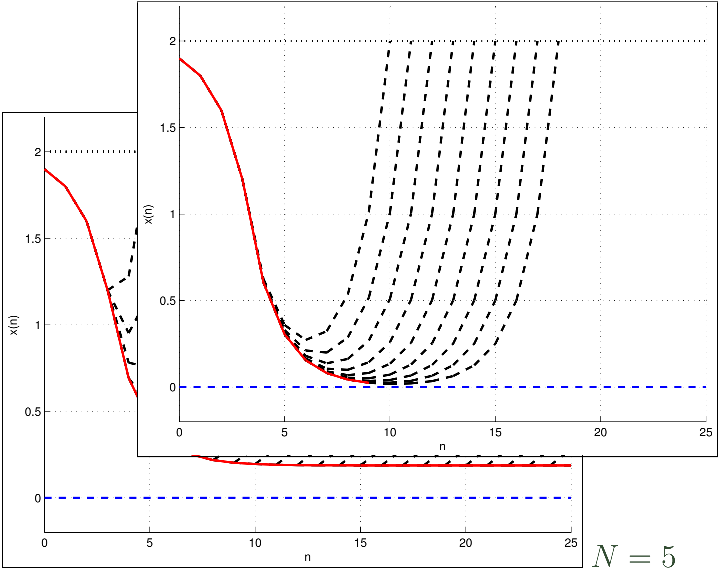

$N = 10$

$N = 5$

Example: trajectories

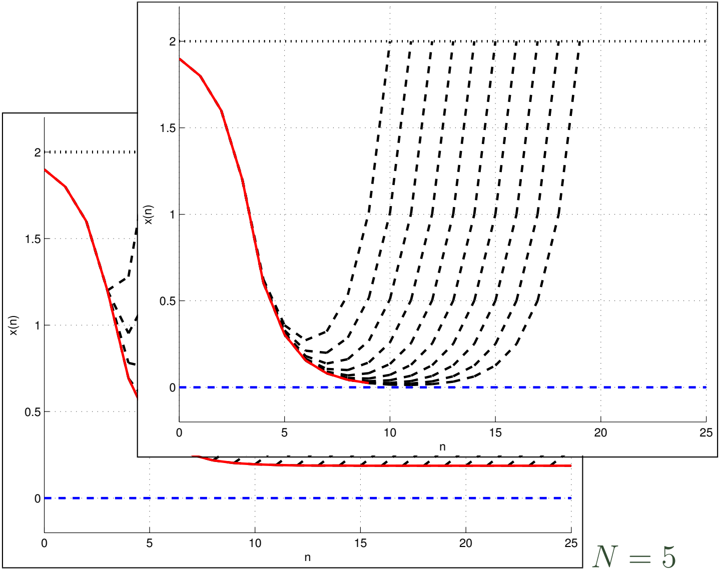

$N = 10$

$N = 5$

Example: trajectories

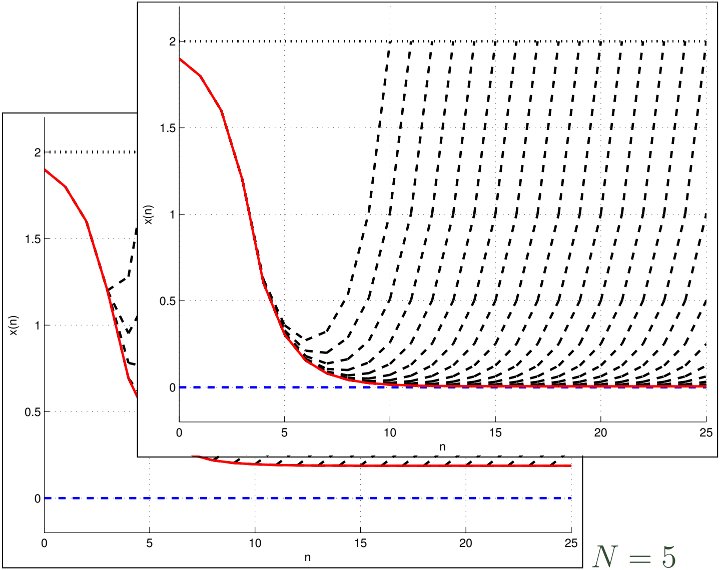

$N = 10$

$N = 5$

Example: averaged closed loop performance

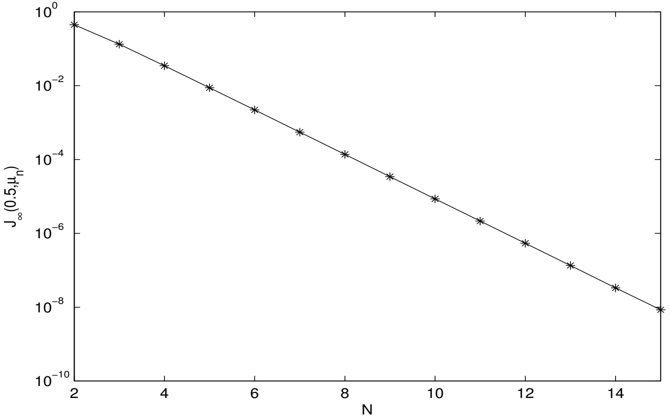

$\overline{J}^{cl}_\infty(0.5, \mu_N) - \ell(x^e, u^e)$ depending on N, logarithmic scale

Observations

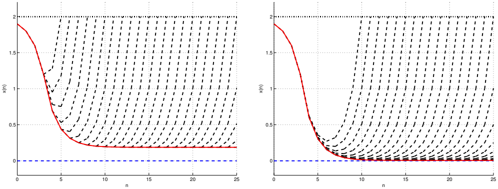

Observations

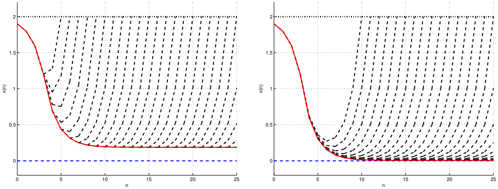

- optimal open loop trajectories approach the optimal equilibrium, stay there for a while, and turn away

  – "turnpike property"

Observations

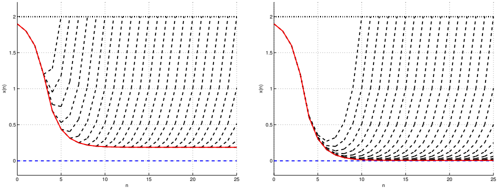

- optimal open loop trajectories approach the optimal equilibrium, stay there for a while, and turn away – "turnpike property"
- closed loop trajectories converge to a neighborhood of the optimal equilibrium whose size tends to 0 as $N \to \infty$

Observations

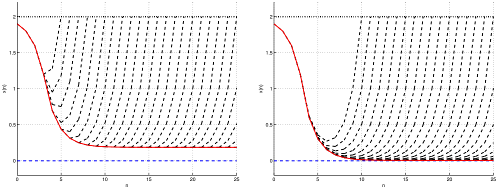

- optimal open loop trajectories approach the optimal equilibrium, stay there for a while, and turn away

  – "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)

Observations

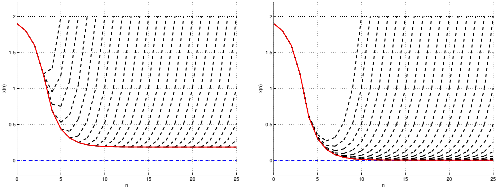

- optimal open loop trajectories approach the optimal equilibrium, stay there for a while, and turn away – "turnpike property"
- closed loop trajectories converge to a neighborhood of the optimal equilibrium whose size tends to 0 as $N \to \infty$
- the averaged closed loop performance satisfies $\overline{J}^{cl}_{\infty}(x,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

Can we prove this behavior?

Observations

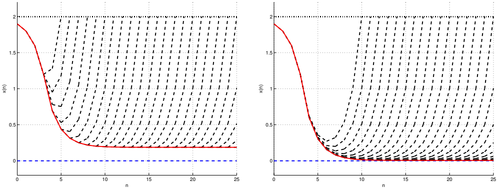

- optimal open loop trajectories approach the optimal equilibrium, stay there for a while, and turn away – "turnpike property"
- closed loop trajectories converge to a neighborhood of the optimal equilibrium whose size tends to 0 as $N \to \infty$
- the averaged closed loop performance satisfies $\overline{J}^{cl}_{\infty}(x,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

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

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

for a small error term $\varepsilon > 0$

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Towards a performance estimate

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

Towards a performance estimate

Similarly, estimates for the non averaged J_K^{cl} 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 $\varepsilon > 0$

Towards a performance estimate

Similarly, estimates for the non averaged J_K^{cl} 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 $V_{N-1}(x)$

Towards a performance estimate

Similarly, estimates for the non averaged J_K^{cl} 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 $\ell(x^e, u^e) + \varepsilon$

Towards a performance estimate

Similarly, estimates for the non averaged J_K^{cl} 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 $V_N(x)$

Towards a performance estimate

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

Hence, the desired inequality is

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

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:

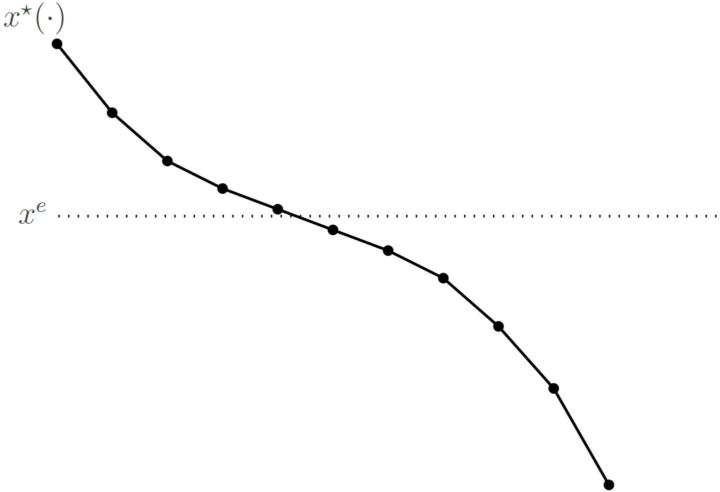

Prolonging near x^e

Sketch of the idea:

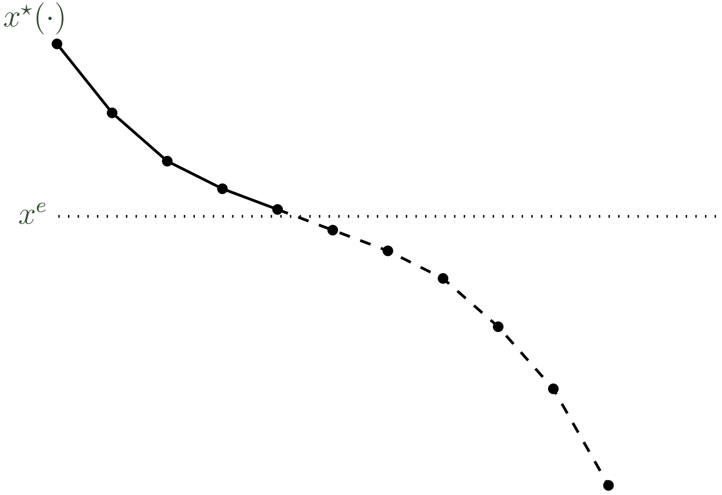

Prolonging near x^e

Sketch of the idea:

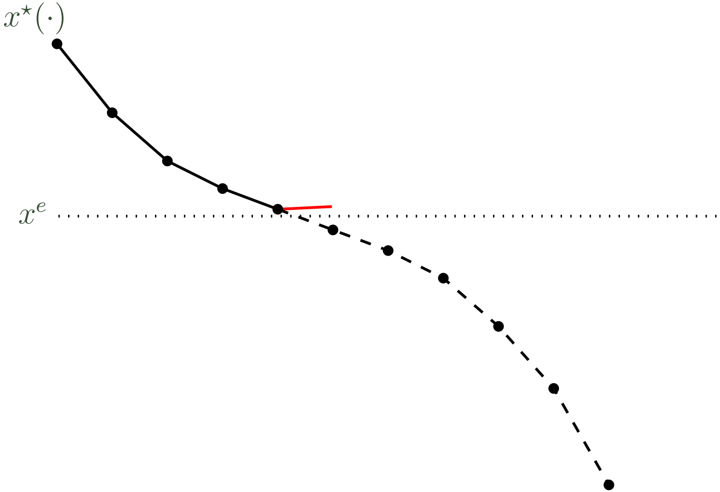

Prolonging near x^e

Sketch of the idea:

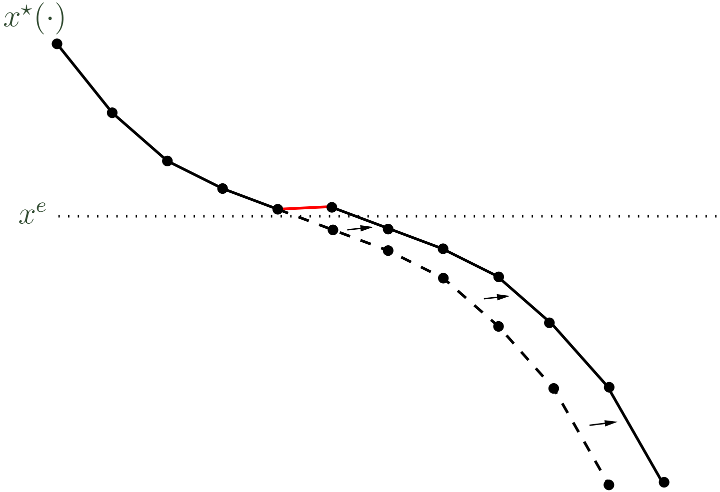

Assumptions needed for this construction

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

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

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

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

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

$$\ell(x,u) - \ell(x^e,u^e) + \lambda(x) - \lambda(f(x,u)) \geq \alpha(\|x - x^e\|)$$

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

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

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 '15]: Under suitable controllability conditions, strict dissipativity is equivalent to the turnpike property plus optimality of the equilibrium (x^e, u^e)

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 '15]: Under suitable controllability conditions, strict dissipativity is equivalent to the turnpike property plus optimality of the equilibrium (x^e, u^e)

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

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 '15]: Under suitable controllability conditions, strict dissipativity is equivalent to the turnpike property plus optimality of the equilibrium (x^e, u^e)

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

Tracking type 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 $x \in \mathbb{X}$

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 $x \in \mathbb{X}$

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

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

} $\Rightarrow$ uniform continuity of V_N

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

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

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

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

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

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

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

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

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

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

(i)–(iii) $\Rightarrow$ turnpike property

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

(i)–(iv) $\Rightarrow$ exponential turnpike
[Damm/Gr./Stieler/Worthmann '14]

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 } $\Rightarrow$ uniform continuity of V_N

(ii) strict dissipativity } $\Rightarrow$ turnpike property

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

(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

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:

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

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:

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

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

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

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)

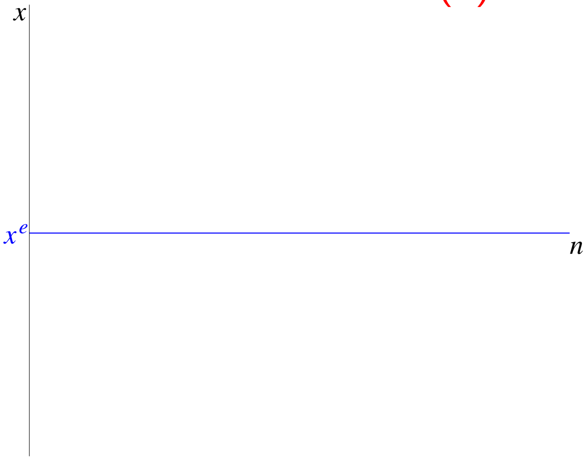

Illustration of (2) and (3)

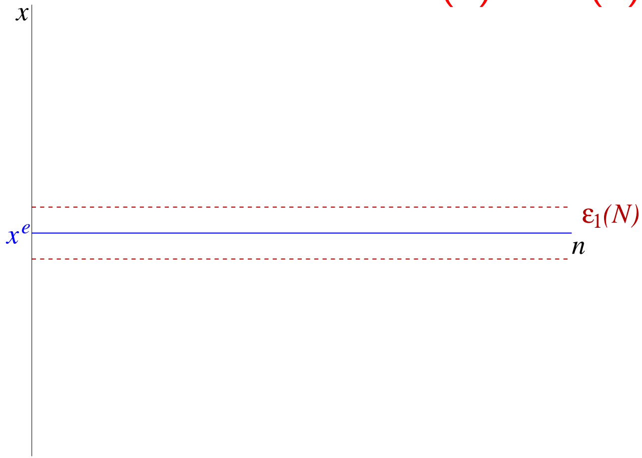

Illustration of (2) and (3)

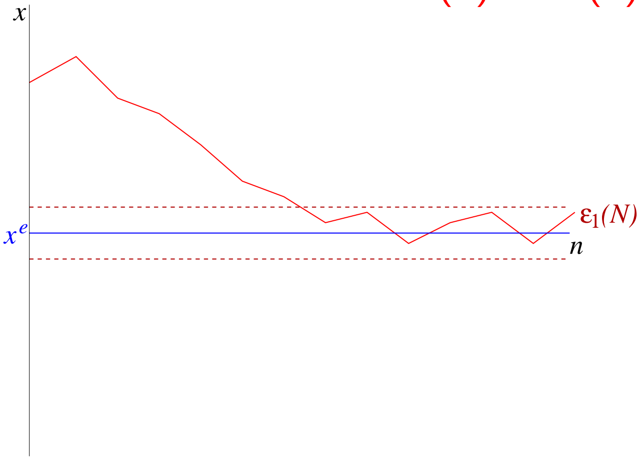

(2): $x_{\mu_N}(n)$ converges to the $\varepsilon_1(N)$-ball around x^e

Illustration of (2) and (3)

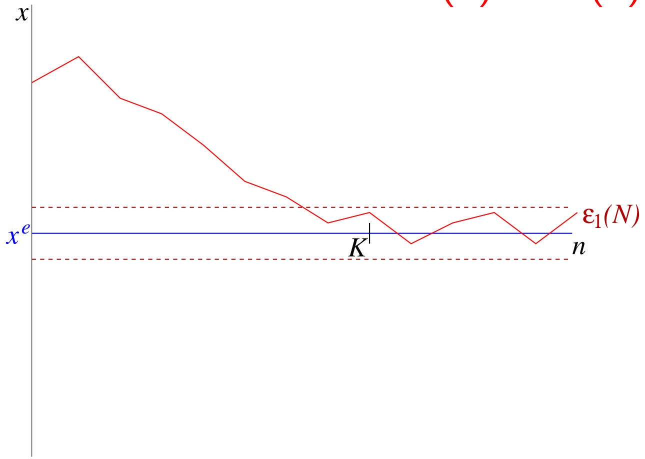

(2): $x_{\mu_N}(n)$ converges to the $\varepsilon_1(N)$-ball around x^e

Illustration of (2) and (3)

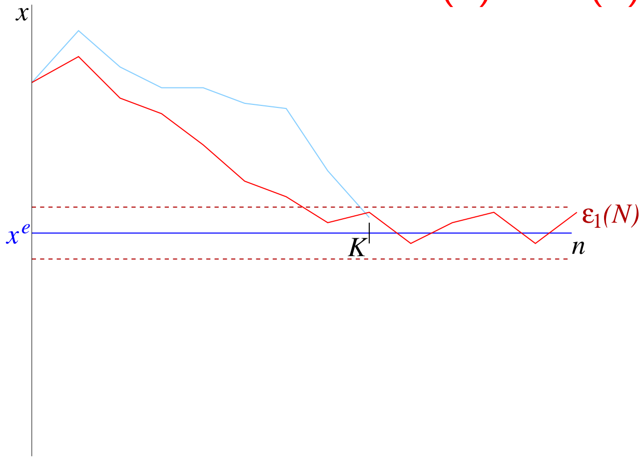

(2): $x_{\mu_N}(n)$ converges to the $\varepsilon_1(N)$-ball around x^e

Illustration of (2) and (3)

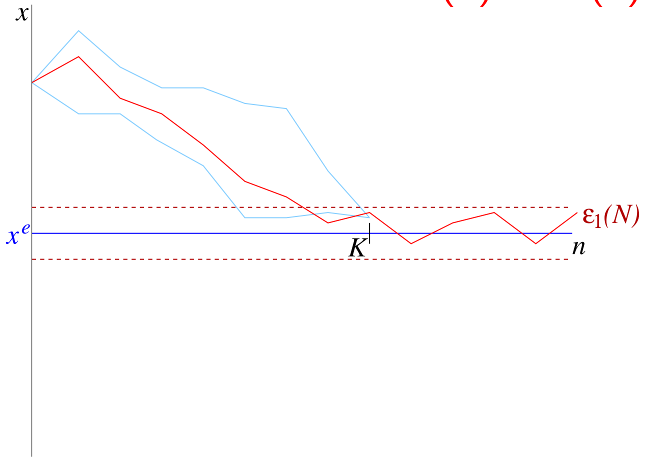

(2): $x_{\mu_N}(n)$ converges to the $\varepsilon_1(N)$-ball around x^e

Illustration of (2) and (3)

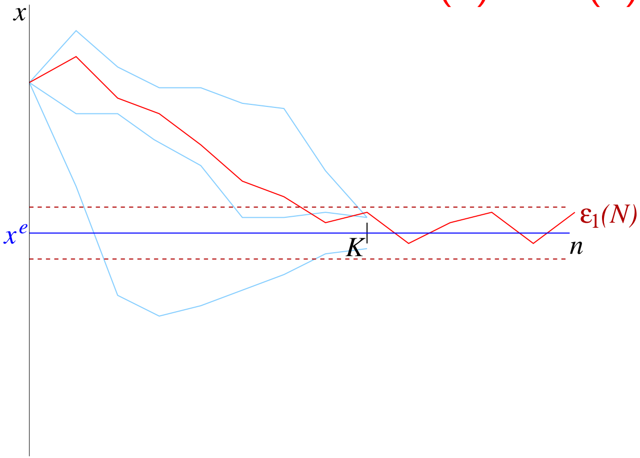

(2): $x_{\mu_N}(n)$ converges to the $\varepsilon_1(N)$-ball around x^e

Illustration of (2) and (3)

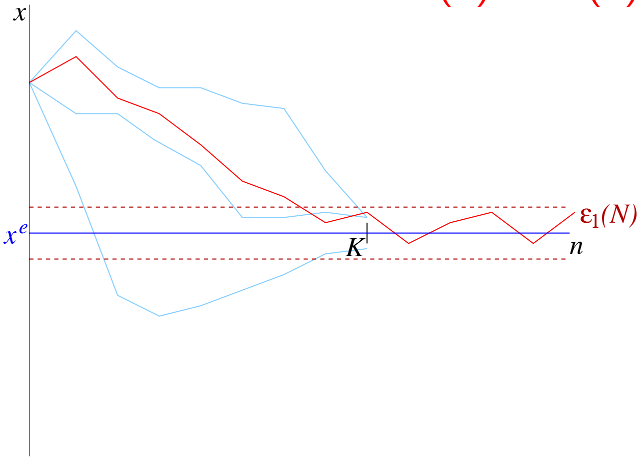

(2): $x_{\mu_N}(n)$ converges to the $\varepsilon_1(N)$-ball around x^e

(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 $K\varepsilon_1(N) + \varepsilon_2(K)$

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $\quad x_{\mathbf{u}}(N) = x^e$

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

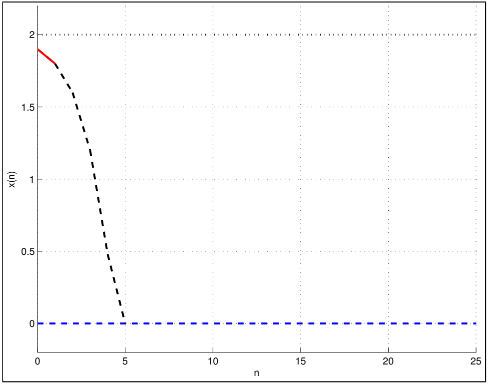

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

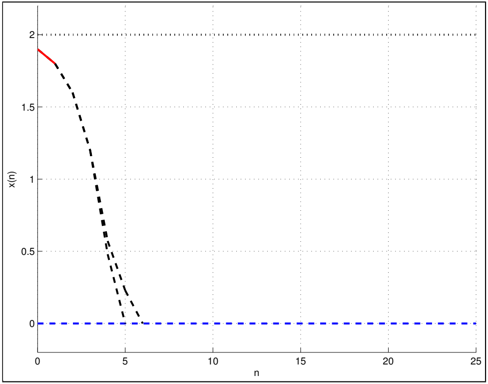

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $\quad x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

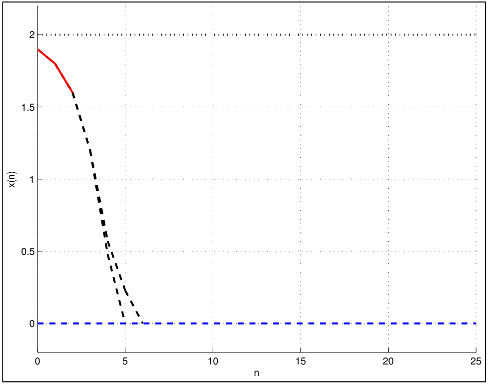

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

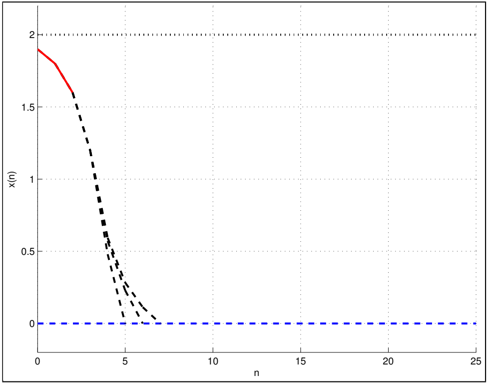

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

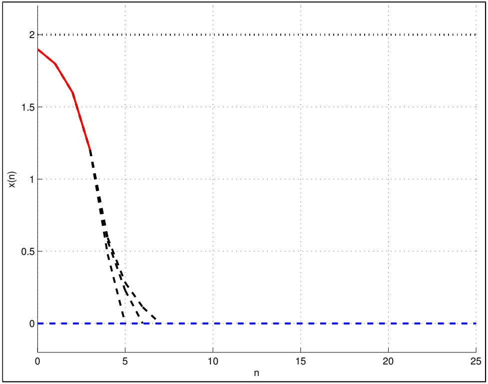

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $\quad x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

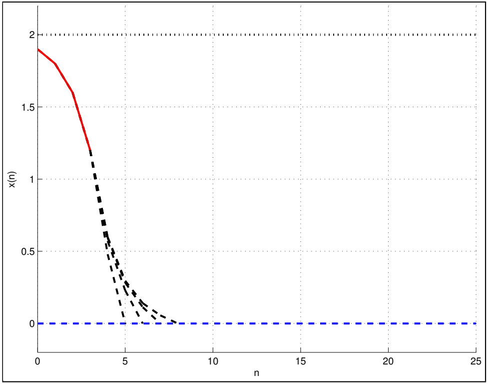

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

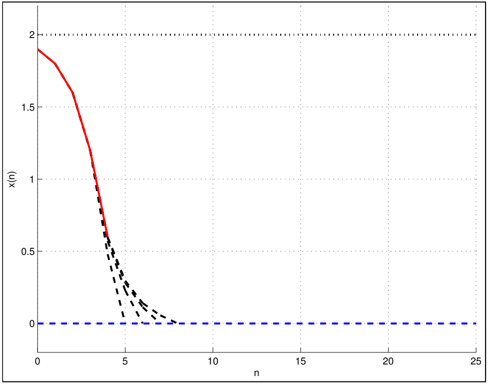

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

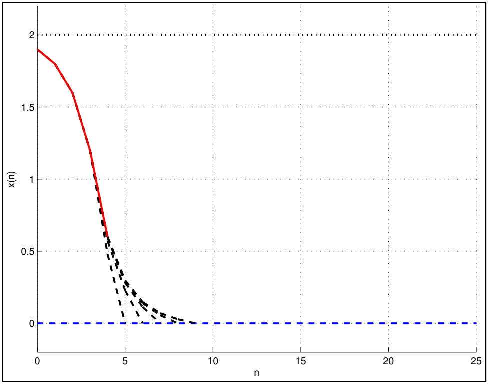

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

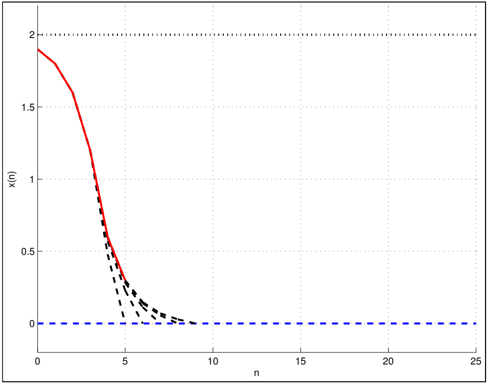

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

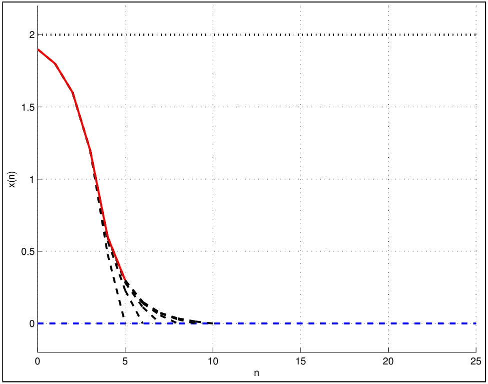

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

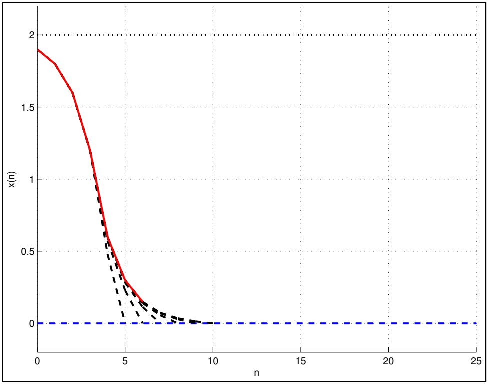

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

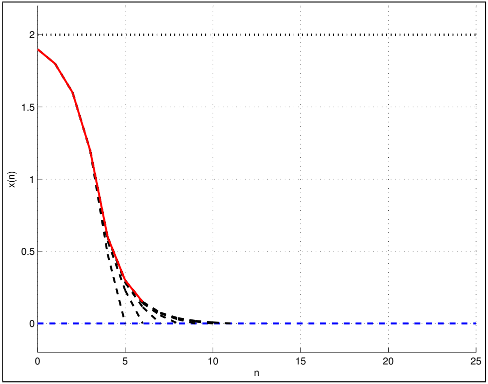

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

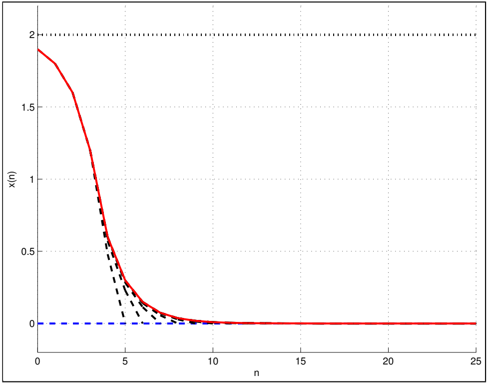

Schemes with terminal constraints

If we know the equilibrium x^e, we may use it as a terminal constraint, i.e., in each step of the MPC scheme we optimize only over those trajectories satisfying $\quad x_{\mathbf{u}}(N) = x^e$

Example:
$N = 5$

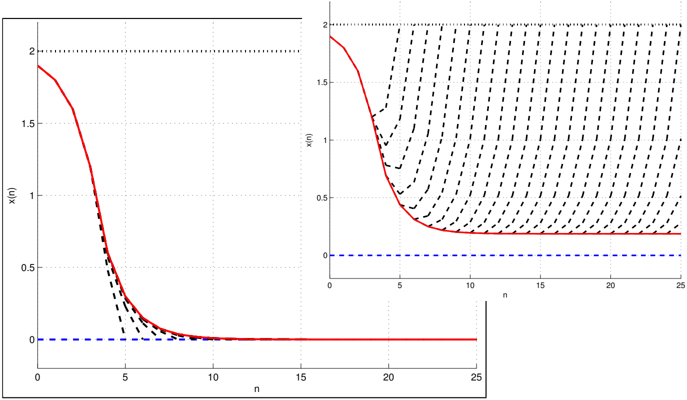

Schemes with terminal constraints

Imposing $x_{\mathbf{u}}(N) = x^e$ improves the previous results

Schemes with terminal constraints

Imposing $x_{\mathbf{u}}(N) = x^e$ 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

Schemes with terminal constraints

Imposing $x_{\mathbf{u}}(N) = x^e$ 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,

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

Schemes with terminal constraints

Imposing $x_{\mathbf{u}}(N) = x^e$ 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,

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

Schemes with terminal constraints

Imposing $x_{\mathbf{u}}(N) = x^e$ 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,

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

for which x^e is asymptotically stable

Schemes with terminal constraints

Imposing $x_{\mathbf{u}}(N) = x^e$ 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,

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

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

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

Schemes with terminal constraints

Imposing $x_{\mathbf{u}}(N) = x^e$ 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.,

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

Schemes with terminal constraints

Imposing $x_{\mathbf{u}}(N) = x^e$ 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,

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

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

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

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

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

Schemes with terminal constraints

Imposing $x_{\mathbf{u}}(N) = x^e$ 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,

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

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

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

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

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

Example: closed loop cost

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

vs.

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

Example: closed loop cost

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

vs.

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

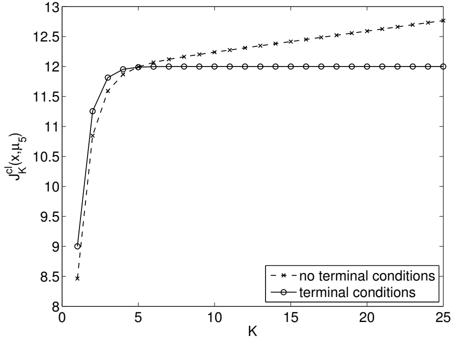

Example: closed loop cost

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

vs.

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

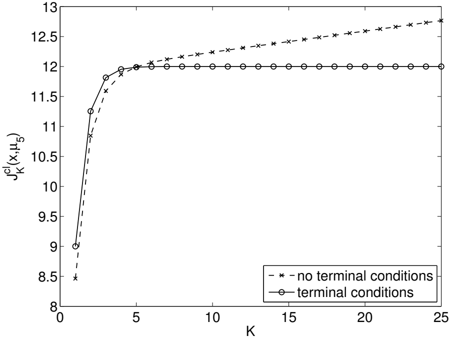

But: terminal constraints can cause infeasibility and severe numerical problems

Extensions, further results

- In the affine linear quadratic case our conditions are equivalent to the system being stabilizable [Gr./Stieler '14]

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. '15, Müller/Gr. '15]

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. '15, Müller/Gr. '15]
- The terminal constraint $x_{\mathbf{u}}(N) = x^e$ can be relaxed to $x_{\mathbf{u}}(N) \in \mathbb{X}_0$ for a neighborhood $\mathbb{X}_0$ of x^e if the functional J_N is appropriately modified [Amrit/Rawlings/Angeli '12, Gr./Panin '15]

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. '15, Müller/Gr. '15]
- The terminal constraint $x_{\mathbf{u}}(N) = x^e$ can be relaxed to $x_{\mathbf{u}}(N) \in \mathbb{X}_0$ for a neighborhood $\mathbb{X}_0$ of x^e if the functional J_N is appropriately modified [Amrit/Rawlings/Angeli '12, Gr./Panin '15]
- The results can be formulated directly in continuous time [Faulwasser/Bonvin '15, Alessandretti/Aguiar/Jones '15]

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. '15, Müller/Gr. '15]
- The terminal constraint $x_{\mathbf{u}}(N) = x^e$ can be relaxed to $x_{\mathbf{u}}(N) \in \mathbb{X}_0$ for a neighborhood $\mathbb{X}_0$ of x^e if the functional J_N is appropriately modified [Amrit/Rawlings/Angeli '12, Gr./Panin '15]
- 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]

Example: Fokker-Planck Equation

Consider a stochastic process governed by a controlled Itô stochastic differential equation (SDE)

$$dX_t = b(X_t, t; u)dt + \sigma(X_t, t)dW_t, \qquad X_{t_0} = x_0$$

Example: Fokker-Planck Equation

Consider a stochastic process governed by a controlled Itô stochastic differential equation (SDE)

$$dX_t = b(X_t, t; u)dt + \sigma(X_t, t)dW_t, \qquad X_{t_0} = x_0$$

where the random variable $X_t \in \mathbb{R}^d$ represents the state

Example: Fokker-Planck Equation

Consider a stochastic process governed by a controlled Itô stochastic differential equation (SDE)

$$dX_t = b(X_t, t; u)dt + \sigma(X_t, t)dW_t, \qquad X_{t_0} = x_0$$

where the random variable $X_t \in \mathbb{R}^d$ represents the state

Idea: control the statistical properties of X_t by controlling its probability density function $y(x, t)$

The Fokker-Planck Equation

The probability density function (PDF) $y(x,t)$ of X_t solves the Fokker-Planck Equation

$$\begin{aligned} \partial_t y(x,t) - \sum_{i,j=1}^{d} \partial^2_{x_i x_j}\Big(a_{ij}(x,t)y(x,t)\Big) + \sum_{i=1}^{d} \partial_{x_i}\Big(b_i(x,t;u))y(x,t)\Big) &= 0 \\ y(\cdot,0) &= y_0 \end{aligned}$$

The Fokker-Planck Equation

The probability density function (PDF) $y(x,t)$ of X_t solves the Fokker-Planck Equation

$$\begin{aligned} \partial_t y(x,t) - \sum_{i,j=1}^{d} \partial^2_{x_i x_j}\Big(a_{ij}(x,t)y(x,t)\Big) + \sum_{i=1}^{d} \partial_{x_i}\Big(b_i(x,t;u))y(x,t)\Big) &= 0 \\ y(\cdot,0) &= y_0 \end{aligned}$$

where $y\colon \mathbb{R}^d \times [0,\infty[\to \mathbb{R}_{\geq 0}$ is the PDF

$y_0\colon \mathbb{R}^d \to \mathbb{R}_{\geq 0}$ is the initial PDF

$a = \sigma\sigma^T/2$ is a positive definite symmetric matrix

$b_i\colon \mathbb{R}^d \times [0,\infty[\times U \to \mathbb{R}$, $i = 1,\dots,d$.

MPC for the Fokker-Planck equation

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

Idea: [Annunziato/Borzì '10ff.] Prescribe a desired PDF $y_d(x,t)$ and use MPC for the FP equation in order to track this PDF

$$\rightsquigarrow \quad J_N(y,u) = \frac{1}{2}\sum_{n=0}^{N-1}\Big(\|y(t_{n+1}) - y_d(t_{n+1})\|^2_{L^2(\Omega)} + \lambda\|u(t_n)\|^2\Big)$$

$$t_n = nT$$

MPC for the Fokker-Planck equation

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

Idea: [Annunziato/Borzì '10ff.] Prescribe a desired PDF $y_d(x,t)$ and use MPC for the FP equation in order to track this PDF

$$\rightsquigarrow \;\; J_N(y,u) = \frac{1}{2}\sum_{n=0}^{N-1}\Big(\|y(t_{n+1}) - y_d(t_{n+1})\|^2_{L^2(\Omega)} + \lambda\|u(t_n)\|^2\Big)$$

$$t_n = nT$$

[Annunziato/Borzì '10ff.] used this idea with $N = 2$ and u independent of the space variable x

MPC for the Fokker-Planck equation

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

Idea: [Annunziato/Borzì '10ff.] Prescribe a desired PDF $y_d(x,t)$ and use MPC for the FP equation in order to track this PDF

$$\rightsquigarrow \quad J_N(y,u) = \frac{1}{2}\sum_{n=0}^{N-1}\Big(\|y(t_{n+1}) - y_d(t_{n+1})\|^2_{L^2(\Omega)} + \lambda\|u(t_n)\|^2\Big)$$

$$t_n = nT$$

[Annunziato/Borzì '10ff.] used this idea with $N = 2$ and u independent of the space variable x

We extended this to arbitrary N and u depending on t and x

Numerical Example in 2D

2d Ornstein-Uhlenbeck type process on $\Omega = (-5,5)^2$

$$dX_t = b(X_t, t; u)dt + \sigma(X_t, t)dW_t, \qquad X_{t_0} = x_0$$

with

$$\sigma(x,t) = \begin{pmatrix} 0.8 & 0 \\ 0 & 0.8 \end{pmatrix}, \quad b(x,t;u) = \begin{pmatrix} -\mu_1 x_1 + u_1 \\ -\mu_2 x_2 + u_2 \end{pmatrix}$$

$\rightsquigarrow$ Fokker-Planck equation

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

with

$$a(x,t) = \begin{pmatrix} 0.32 & 0 \\ 0 & 0.32 \end{pmatrix}, \quad b(x,t;u) = \begin{pmatrix} -\mu_1 x_1 + u_1 \\ -\mu_2 x_2 + u_2 \end{pmatrix}$$

Numerical Example in 2D

Reference PDF is a bi-modal Gaussian given by

$$y_d(x,t) = \frac{1}{2}\frac{\exp\left(-\frac{(x_1+\mu(t))^2}{2\sigma_{11}^2} - \frac{(x_2-\mu(t))^2}{2\sigma_{21}^2}\right)}{2\pi\sigma_{11}\sigma_{21}} + \frac{1}{2}\frac{\exp\left(-\frac{(x_1-\mu(t))^2}{2\sigma_{12}^2} - \frac{(x_2+\mu(t))^2}{2\sigma_{22}^2}\right)}{2\pi\sigma_{12}\sigma_{22}}$$

with $\mu(t) = 2\sin(\frac{\pi t}{5}), \sigma_{11} = \sigma_{21} = 0.4, \sigma_{12} = \sigma_{22} = 0.6.$

Reference PDF y_d

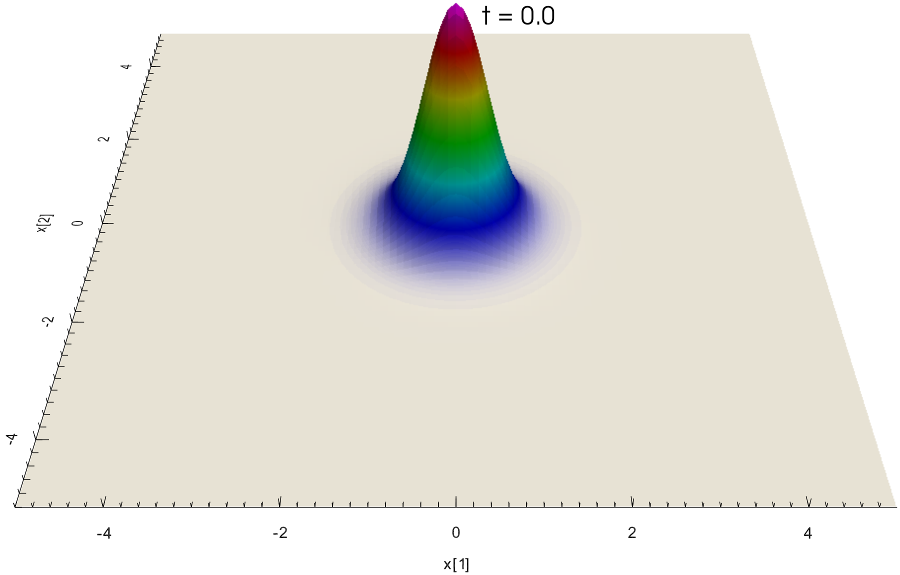

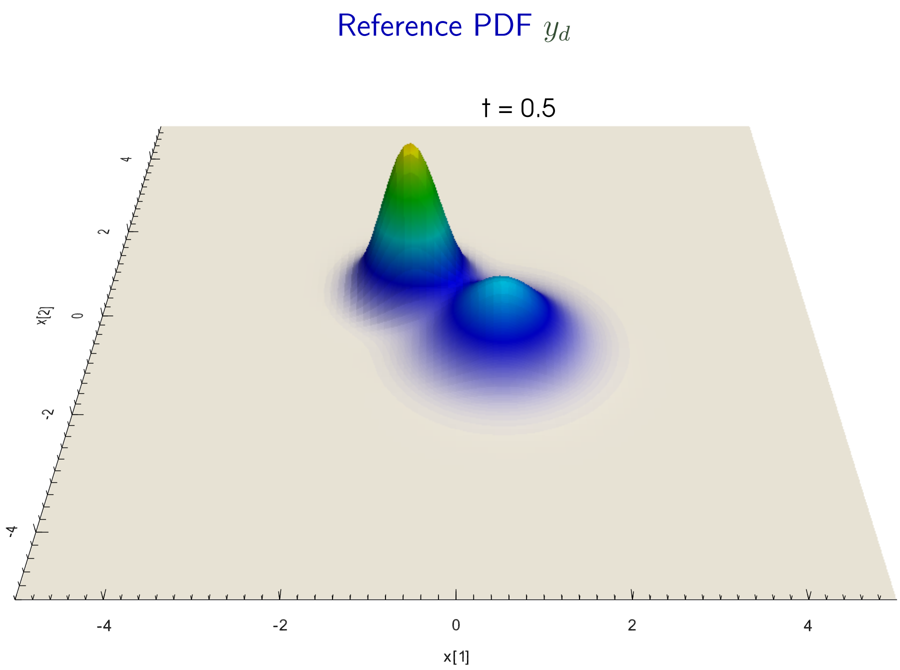

Reference PDF y_d
t = 0.5
x[1]
x[2]
-4
-2
0
2
4

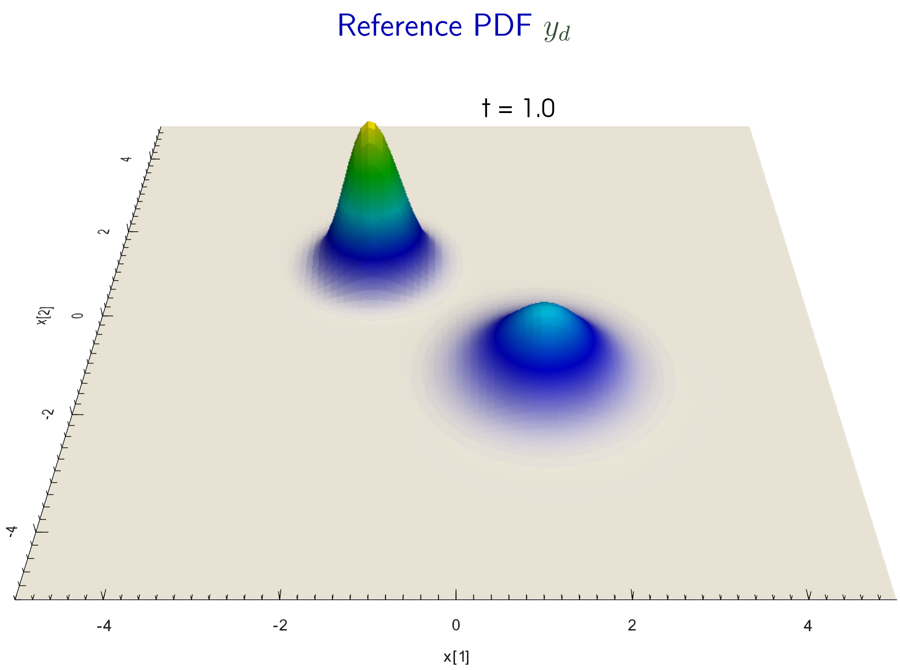
Reference PDF y_d
t = 1.0
4
2
0
-2
-4
x[2]
-4
-2
0
2
4
x[1]

Reference PDF y_d

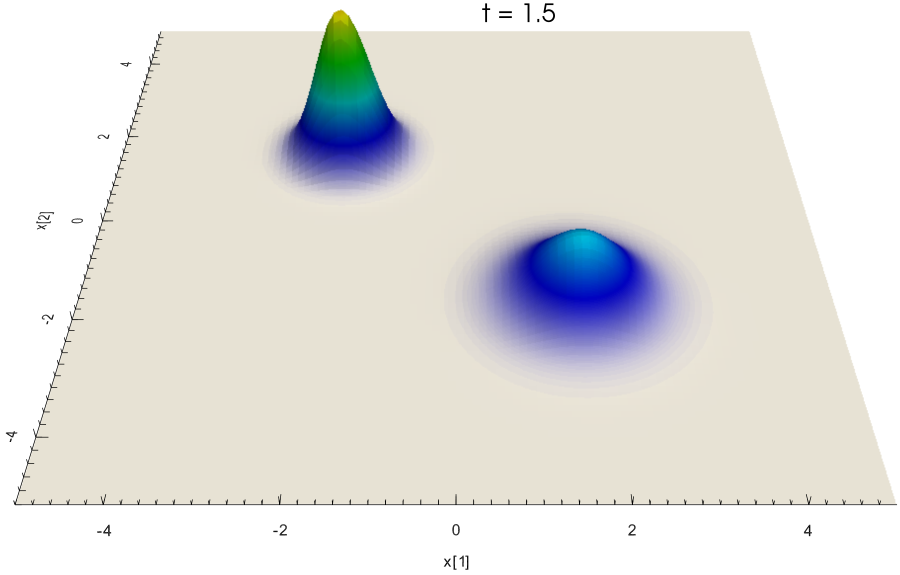

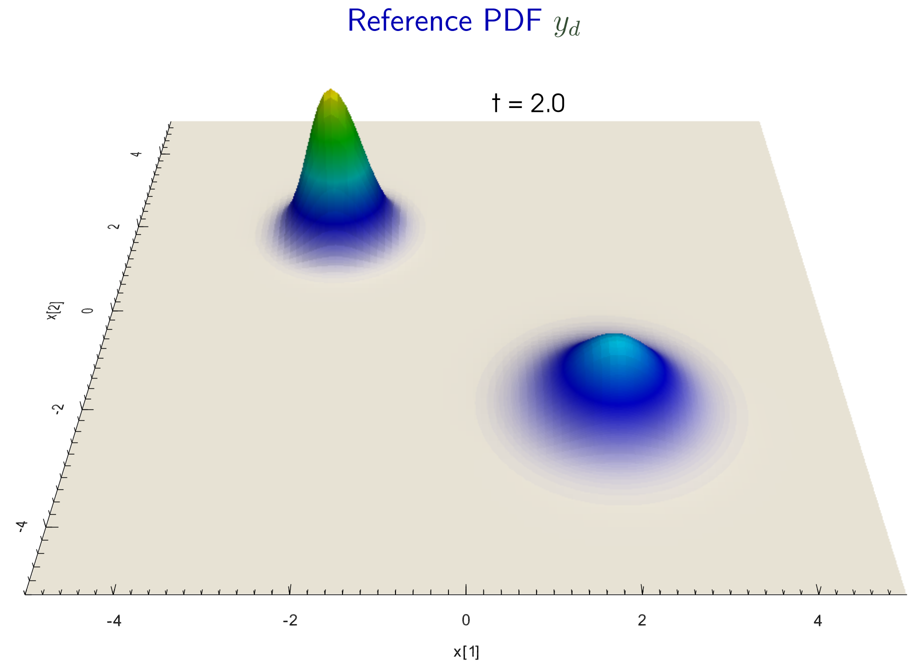
Reference PDF y_d
t = 2.0
4
2
0
-2
-4
x[2]
-4
-2
0
2
4
x[1]

Reference PDF y_d

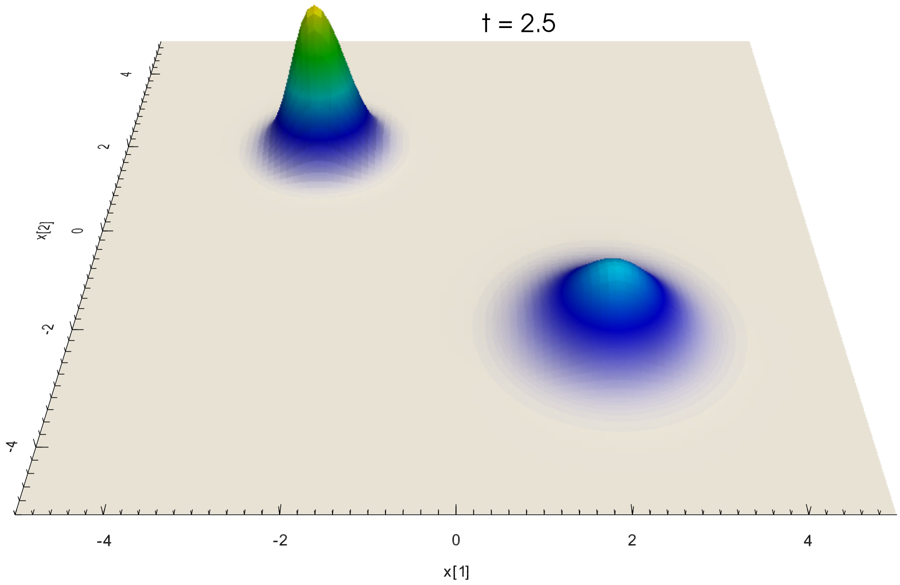

Reference PDF y_d

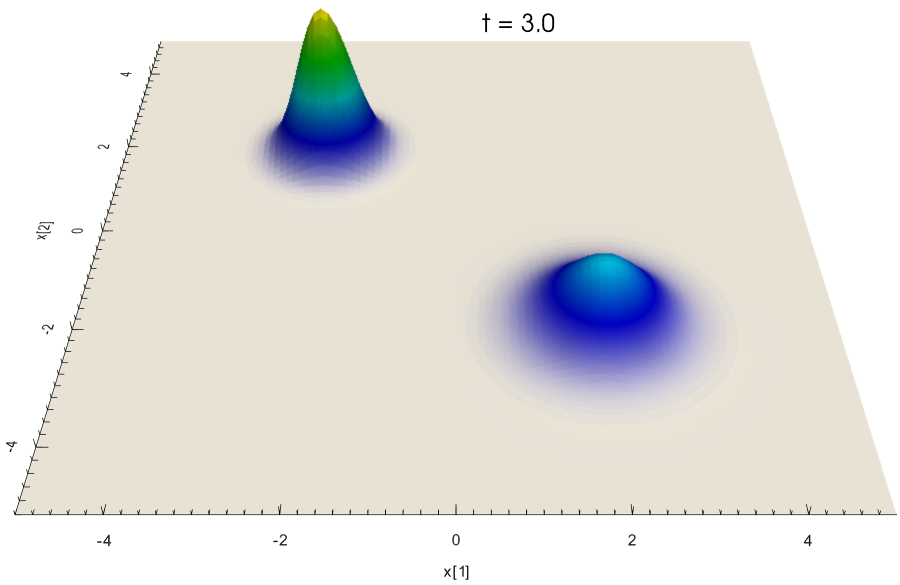

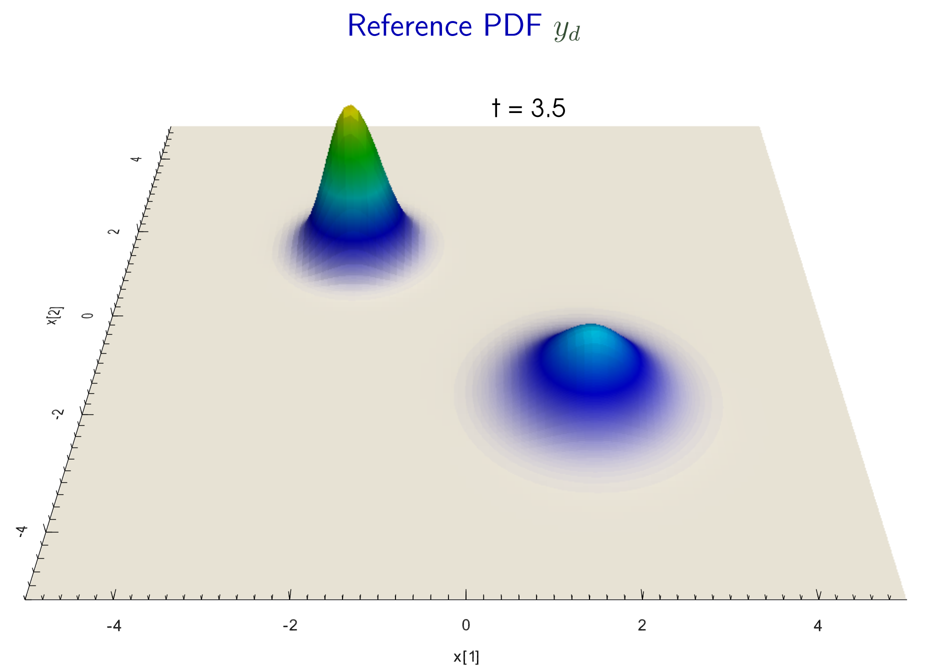

Reference PDF y_d
t = 3.5
4
2
0
-2
-4
x[2]
-4
-2
0
2
4
x[1]

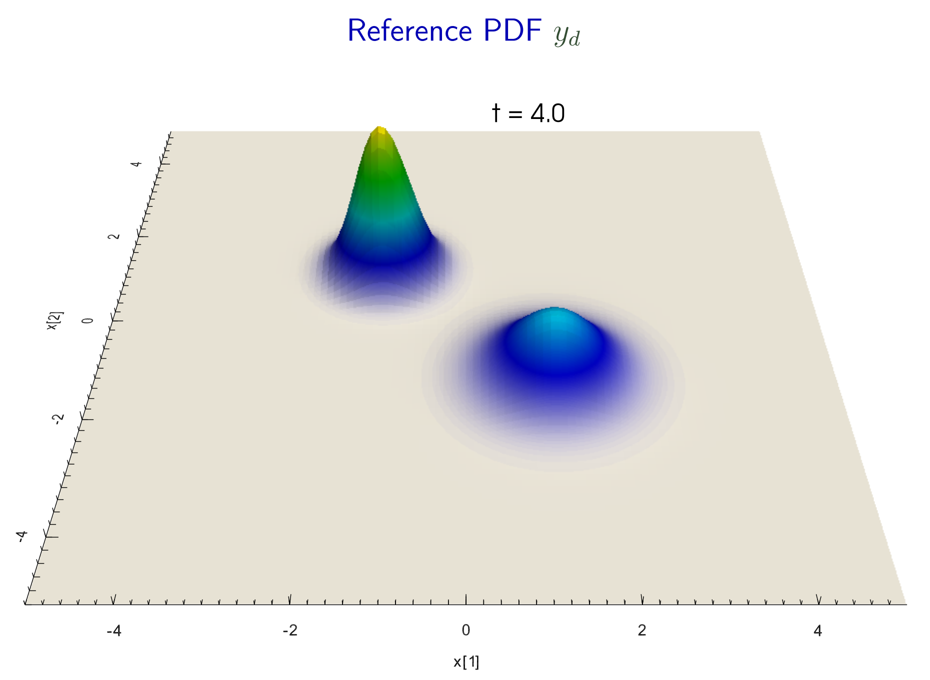

Reference PDF y_d
t = 4.0
x[2]
4
2
0
-2
-4
x[1]
-4
-2
0
2
4

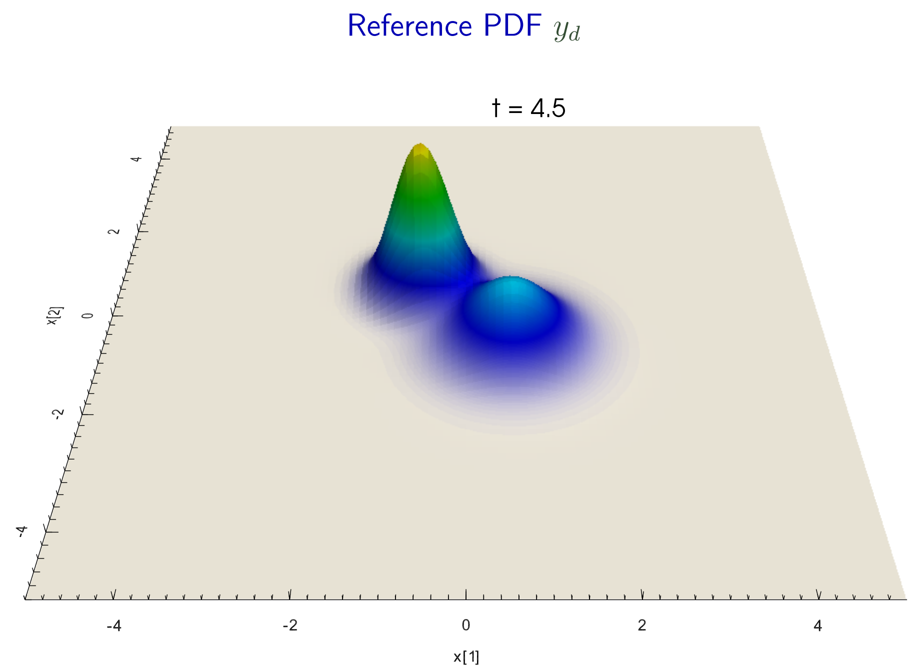
Reference PDF y_d
t = 4.5
x[2]
4
2
0
-2
-4
x[1]
-4
-2
0
2
4

Reference PDF y_d

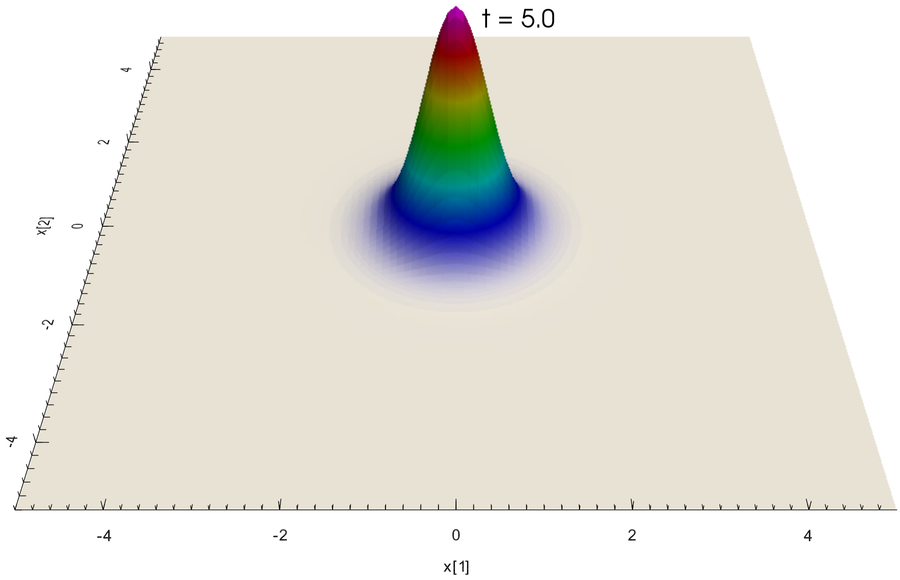

Numerical Example in 2D

Cost functional

$$J(y,u) := \frac{1}{2}\left\|y(t+T) - y_d(t+T)\right\|^2_{L^2(\Omega)} + \frac{\lambda}{2}\left\|u(t)\right\|^2_{L^2(\Omega)}$$

Numerical Example in 2D

Cost functional

$$J(y,u) := \frac{1}{2}\,\|y(t+T) - y_d(t+T)\|^2_{L^2(\Omega)} + \frac{\lambda}{2}\,\|u(t)\|^2_{L^2(\Omega)}$$

Simulation parameters

- initial distribution $y_0(x) = y_d(x,0)$
- optimization horizon $N = 2$
- sampling time $T = 0.5$
- control penalization $\lambda = 0.001$
- control range $u_{1/2} \in [-10, 10]$

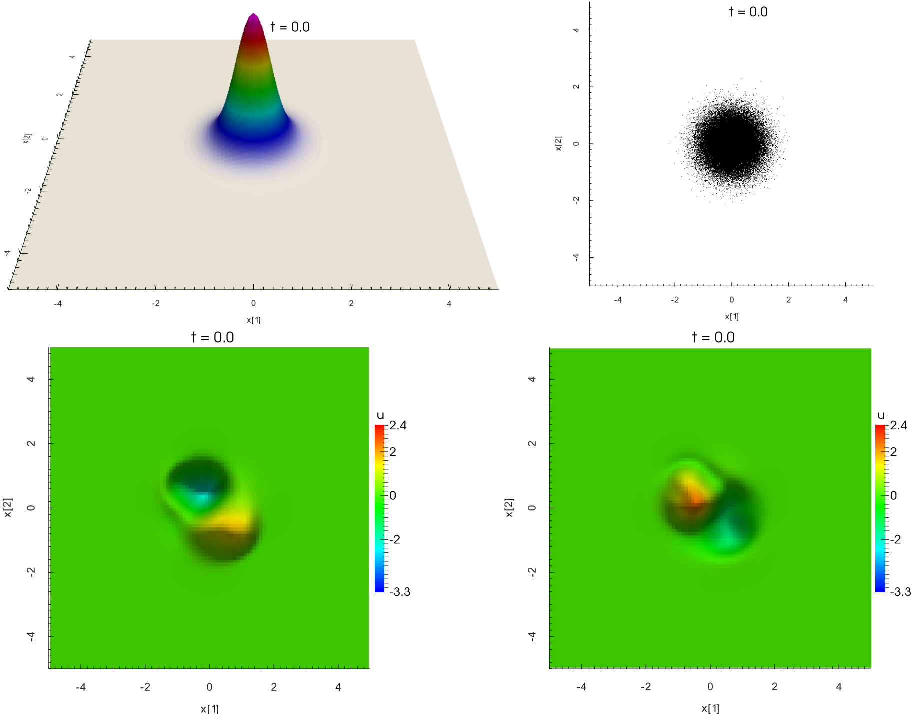
t = 0.0
x[1]
x[2]
t = 0.0
x[1]
x[2]
t = 0.0
u
2.4
2
0
-2
-3.3
x[1]
x[2]
t = 0.0
u
2.4
2
0
-2
-3.3
x[1]
x[2]

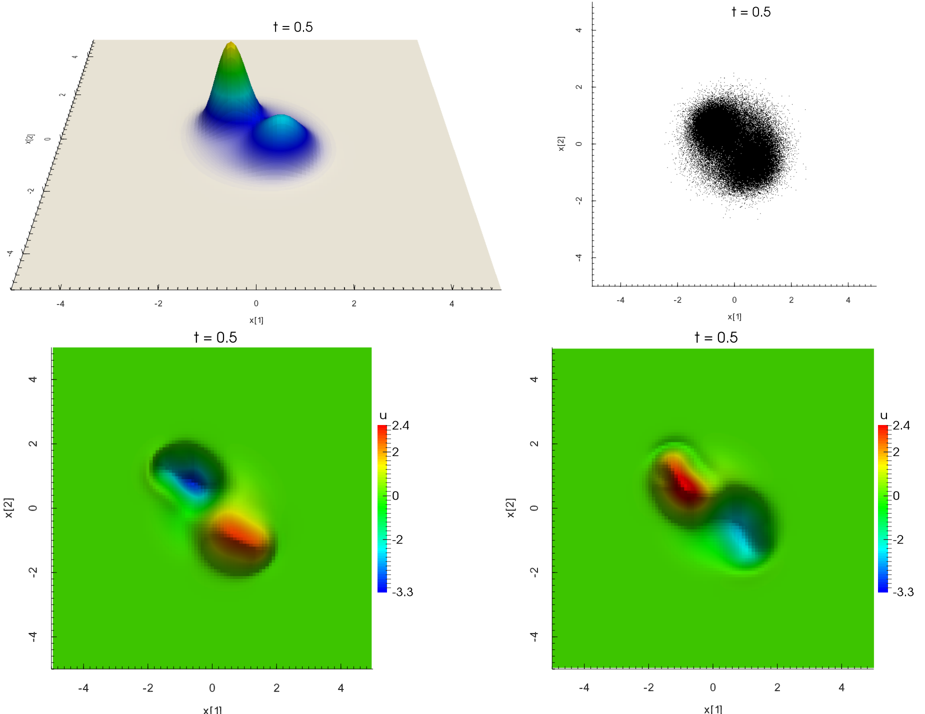

t = 0.5
x[1]
x[2]
u
2.4
2
0
-2
-3.3

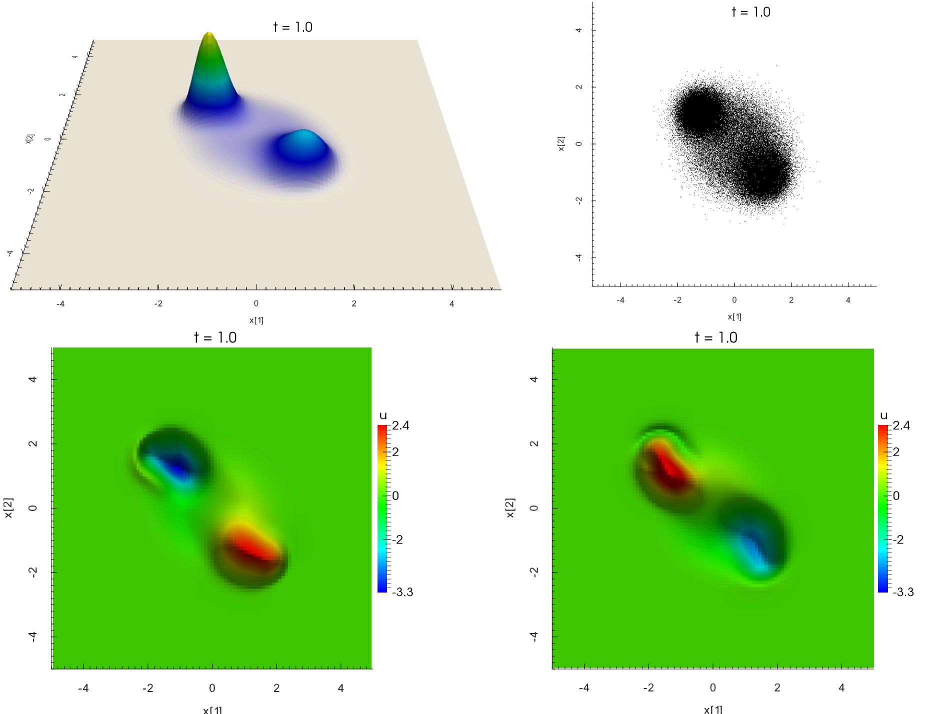

t = 1.0
x[1]
x[2]
u
2.4
2
0
-2
-3.3

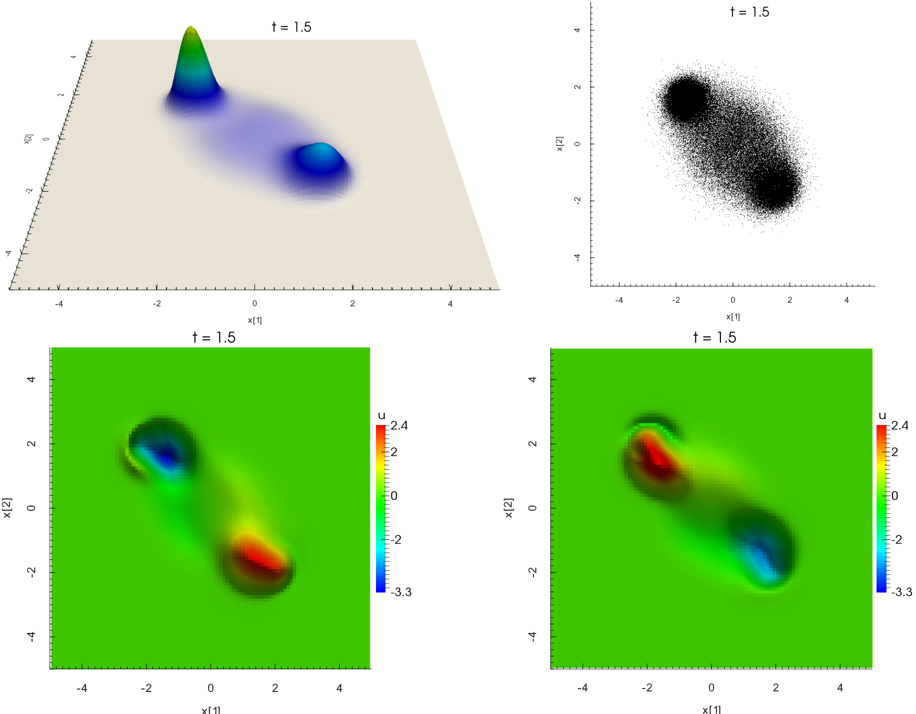
t = 1.5
x[1]
x[2]
-4
-2
0
2
4
t = 1.5
x[1]
x[2]
-4
-2
0
2
4
t = 1.5
u
2.4
2
0
-2
-3.3
x[1]
x[2]
t = 1.5
u
2.4
2
0
-2
-3.3
x[1]
x[2]

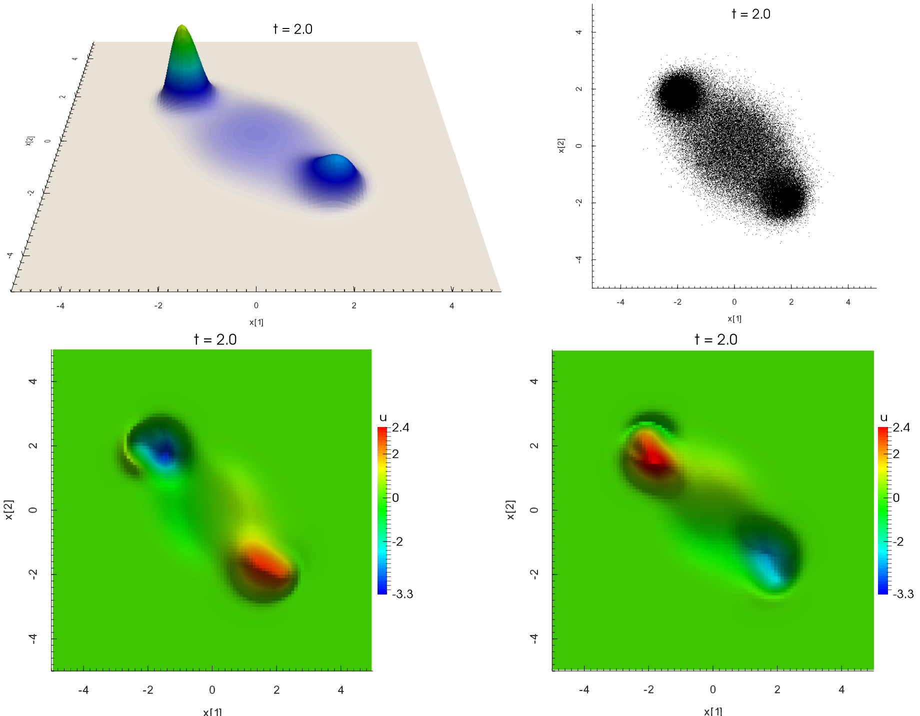
t = 2.0
x[1]
x[2]
t = 2.0
x[1]
x[2]
t = 2.0
u
2.4
2
0
-2
-3.3
x[1]
x[2]
t = 2.0
u
2.4
2
0
-2
-3.3
x[1]
x[2]

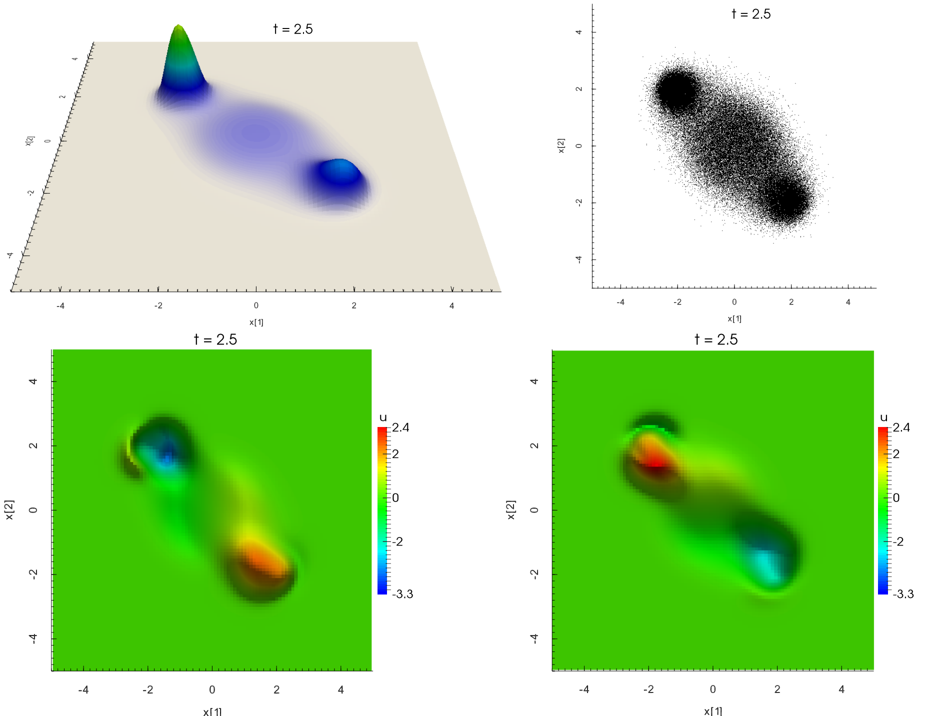

t = 2.5
x[1]
x[2]
-4
-2
0
2
4
u
2.4
2
0
-2
-3.3

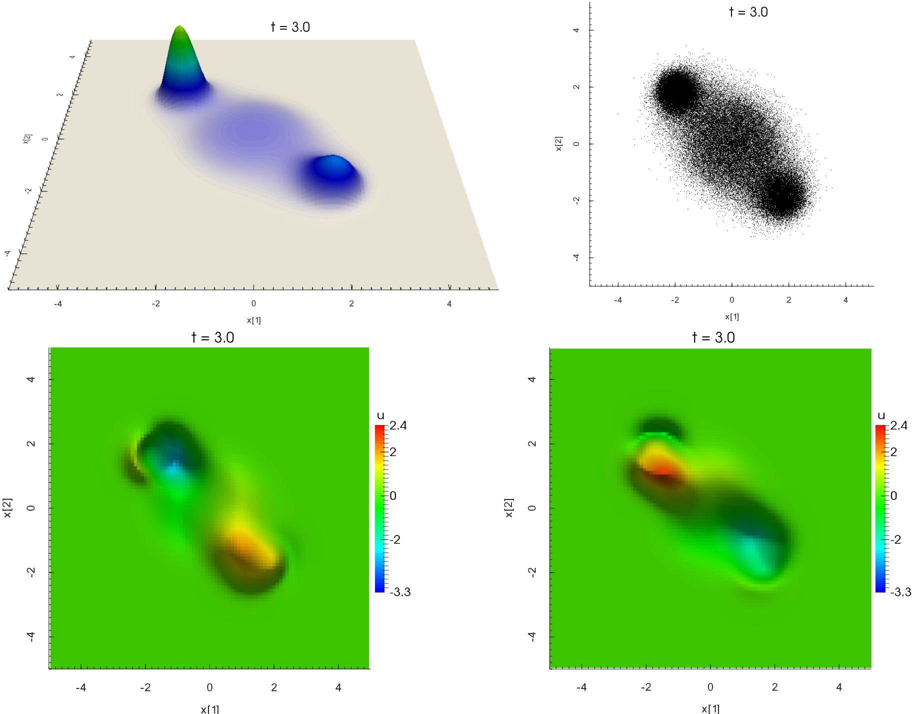
t = 3.0
x[1]
x[2]
-4
-2
0
2
4
t = 3.0
x[1]
x[2]
t = 3.0
x[1]
x[2]
u
2.4
2
0
-2
-3.3
t = 3.0
x[1]
x[2]
u
2.4
2
0
-2
-3.3

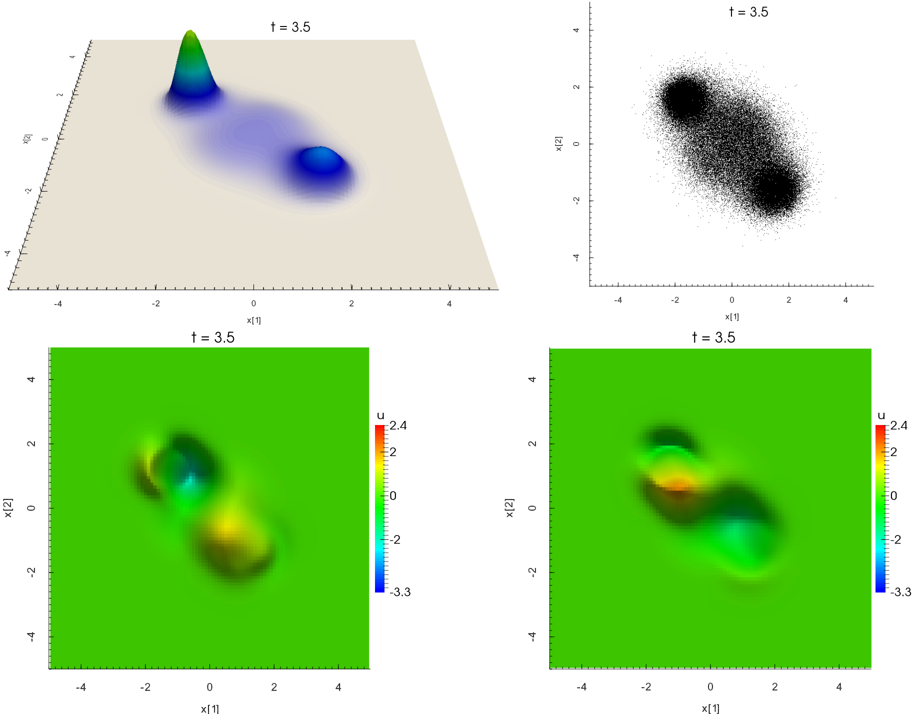

t = 3.5
x[1]
x[2]
-4
-2
0
2
4
t = 3.5
x[1]
x[2]
-4
-2
0
2
4
t = 3.5
x[1]
x[2]
-4
-2
0
2
4
u
2.4
2
0
-2
-3.3
t = 3.5
x[1]
x[2]
-4
-2
0
2
4
u
2.4
2
0
-2
-3.3

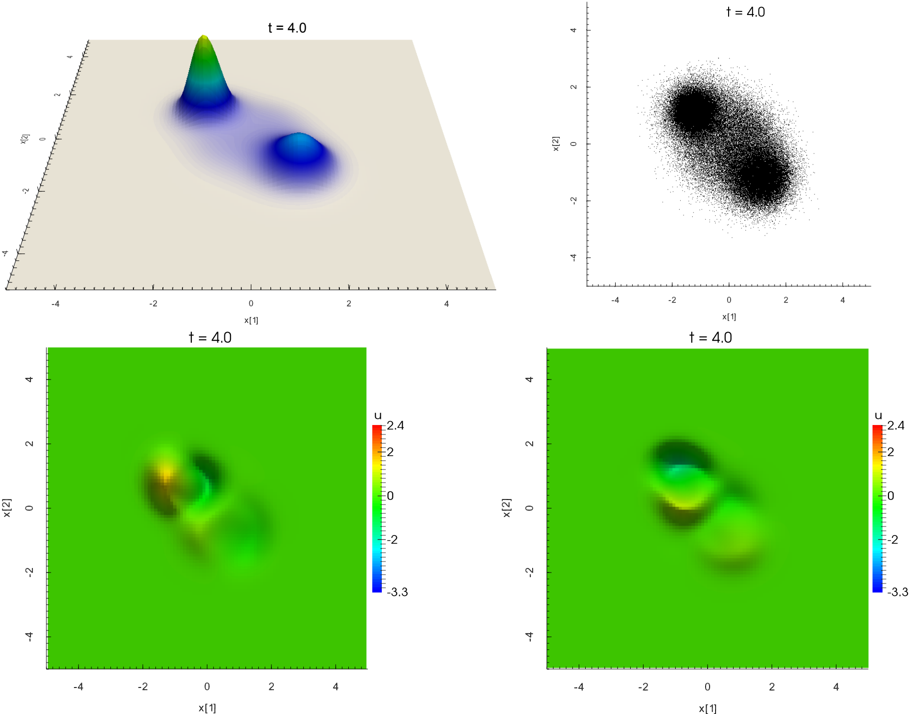

t = 4.0
x[1]
x[2]
t = 4.0
x[1]
x[2]
t = 4.0
u
2.4
2
0
-2
-3.3
x[1]
x[2]
t = 4.0
u
2.4
2
0
-2
-3.3
x[1]
x[2]

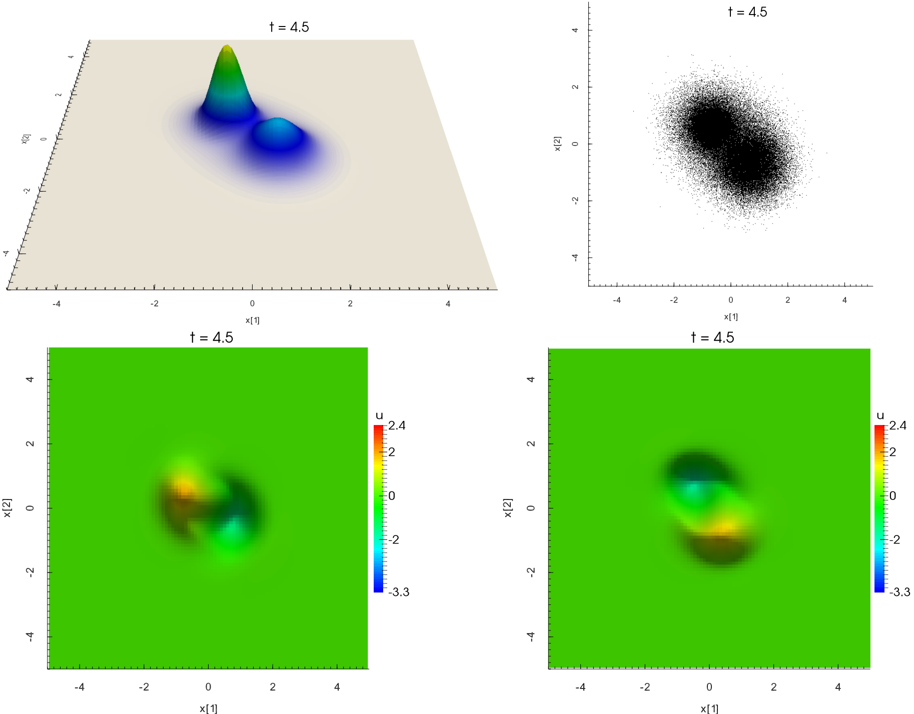
t = 4.5
x[1]
x[2]
t = 4.5
x[1]
x[2]
t = 4.5
u
2.4
2
0
-2
-3.3
x[1]
x[2]
t = 4.5
u
2.4
2
0
-2
-3.3
x[1]
x[2]

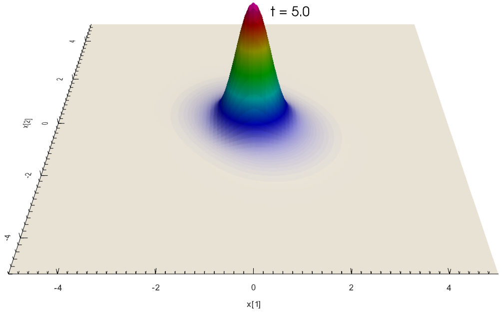
t = 5.0
x[2]
4
2
0
-2
-4
x[1]

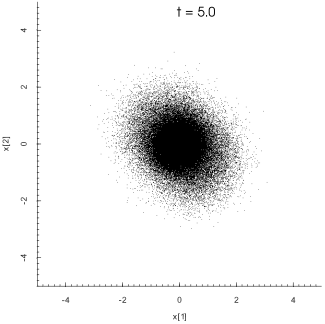
t = 5.0
x[2]
4
2
0
-2
-4
x[1]

Summary and conclusion

- Model Predictive Control can be seen as a method for splitting up an infinite horizon optimal control problem into the iterative solution of finite horizon problems

Summary and conclusion

- Model Predictive Control can be seen as a method for splitting up an infinite horizon optimal control problem into the iterative solution of finite horizon problems

- The existence of the turnpike property at an optimal equilibrium is the key ingredient to make this approach work

Summary and conclusion

- Model Predictive Control can be seen as a method for splitting up an infinite horizon optimal control problem into the iterative solution of finite horizon problems
- The existence of the turnpike property at an optimal equilibrium is the key ingredient to make this approach work
- Strict dissipativity is essentially equivalent to this property and may be used as a checkable condition

Summary and conclusion

- Model Predictive Control can be seen as a method for splitting up an infinite horizon optimal control problem into the iterative solution of finite horizon problems
- The existence of the turnpike property at an optimal equilibrium is the key ingredient to make this approach work
- Strict dissipativity is essentially equivalent to this property and may be used as a checkable condition
- Good news: if MPC works, then it works regardless of whether we checked the conditions

Summary and conclusion

- Model Predictive Control can be seen as a method for splitting up an infinite horizon optimal control problem into the iterative solution of finite horizon problems
- The existence of the turnpike property at an optimal equilibrium is the key ingredient to make this approach work
- Strict dissipativity is essentially equivalent to this property and may be used as a checkable condition
- Good news: if MPC works, then it works regardless of whether we checked the conditions — but if we want to be sure we need to check

Summary and conclusion

- Model Predictive Control can be seen as a method for splitting up an infinite horizon optimal control problem into the iterative solution of finite horizon problems
- The existence of the turnpike property at an optimal equilibrium is the key ingredient to make this approach work
- Strict dissipativity is essentially equivalent to this property and may be used as a checkable condition
- Good news: if MPC works, then it works regardless of whether we checked the conditions — but if we want to be sure we need to check
- If we want to do this for more complex examples, the theory still needs appropriate extension:

Open questions

For instance, for the Fokker-Planck equation the following questions are still open:

Open questions

For instance, for the Fokker-Planck equation the following questions are still open:

- the problem is time varying

Open questions

For instance, for the Fokker-Planck equation the following questions are still open:

- the problem is time varying — but even for a constant reference PDF open questions remain

Open questions

For instance, for the Fokker-Planck equation the following questions are still open:

- the problem is time varying — but even for a constant reference PDF open questions remain
- is the problem strictly dissipative?

Open questions

For instance, for the Fokker-Planck equation the following questions are still open:

- the problem is time varying — but even for a constant reference PDF open questions remain
- is the problem strictly dissipative?
- it is not clear whether the problem satisfies the controllability properties needed to apply our general results

Open questions

For instance, for the Fokker-Planck equation the following questions are still open:

- the problem is time varying — but even for a constant reference PDF open questions remain
- is the problem strictly dissipative?
- it is not clear whether the problem satisfies the controllability properties needed to apply our general results $\rightsquigarrow$ alternative approaches may be needed

Open questions

For instance, for the Fokker-Planck equation the following questions are still open:

- the problem is time varying — but even for a constant reference PDF open questions remain
- is the problem strictly dissipative?
- it is not clear whether the problem satisfies the controllability properties needed to apply our general results $\rightsquigarrow$ alternative approaches may be needed

Overall conclusion: conceptually, the turnpike property appears to be the right tool to understand when receding horizon control works, but: technical issues remain in order to establish the regularity which is needed on top, particularly for PDEs

References

L. Grüne, *Economic receding horizon control without terminal constraints*, Automatica, 49, 725–734, 2013

T. Damm, L. Grüne, M. Stieler, K. Worthmann, *An exponential turnpike theorem for dissipative discrete time optimal control problems*, SIAM J. Control Optim., 52, 1935–1957, 2014

L. Grüne, M. Stieler, *Asymptotic stability and transient optimality of economic MPC without terminal conditions*, Journal of Process Control, 24 (Special Issue on Economic MPC), 1187–1196, 2014

A. Fleig, L. Grüne, R. Guglielmi, *Some results on Model Predictive Control for the Fokker-Planck equation*, Extended Abstract, Proceedings of the 21st MTNS, Groningen, 1203–1206, 2014