Nonlinear model predictive control—Inherent robustness and discrete actuators

James B. Rawlings

Department of Chemical Engineering

Systems Control and Optimization Laboratory
Universität Freiburg
Freiburg, Germany
September 15–19, 2025

Outline

1. Is the stability robust?
   - Input-to-state stability (ISS) and ISS Lyapunov functions
   - Inherent robustness of MPC
2. MPC with discrete actuators
3. Conclusions

So far so good; now is the stability robust?

- Consider disturbances to the process (d) and state measurement (e)

$$x^+ = f(x, \kappa_N(x)) \qquad \text{nominal system}$$
$$x^+ = f(x, \kappa_N(x+e)) + d \qquad \text{nominal controller with disturbances}$$

- How does the perturbed system behave?
- Study of *inherent* robustness motivated by Teel (2004) who showed examples for which arbitrarily small perturbations can destabilize the nominally stabilizing controller.
- If we cannot ensure desirable behavior with small disturbances, the control system will not be useful in practice.
- Every control system fails with large disturbances (think Fukushima nuclear reactor and a tsunami). But the inherent robustness of feedback control must ensure tolerance to small disturbances.

Desired behavior with and without disturbance

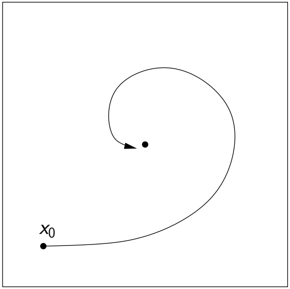

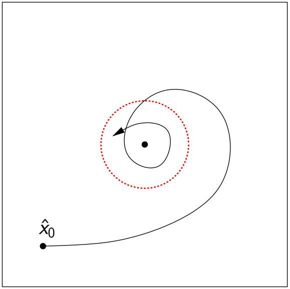

Nominal System

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

System with Disturbance

$$x^+ = f(x, u) + d$$
$$u = \kappa_N(x + e)$$

d is the process disturbance
e is the measurement disturbance

How do we define this desired behavior?

- Nominal controller with disturbances. Note $x_m = x + e$

$$x^+ \in f(x, \kappa_N(x+e)) + d$$
$$x_m^+ \in f(x_m - e, \kappa_N(x_m)) + d + e^+$$
$$x^+ \in F(x, w) \qquad w = (d, e) \text{ or } w = (d, e, e^+)$$

- Inherent robustness: is the origin of the closed-loop system $x^+ \in F(x, w)$ input-to-state stable considering disturbance $w = (d, e)$ as the input?

Input-to-state stability (ISS)

Why ISS?

- Consider a system $x^+ = f(x, w)$ with input w

Definition 1 (Input-to-state stable)

The system $x^+ = f(x, w)$ is (globally) input-to-state stable (ISS) if there exists a $\mathcal{KL}$ function $\beta(\cdot)$ and a $\mathcal{K}$ function $\sigma(\cdot)$ such that, for each $x_0 \in \mathbb{R}^n$, and each bounded disturbance sequence $\boldsymbol{w} = (w(0), w(1), \ldots)$

$$|x(k; x_0, \boldsymbol{w})| \leq \beta(|x_0|, k) + \sigma(\|\boldsymbol{w}\|_{0:k-1})$$

for all $k \in \mathbb{I}_{\geq 0}$, $\|\boldsymbol{w}\|_{a:b} := \max_{j \in \mathbb{I}_{[a:b]}} |w(j)|$

- The main ingredient of robust stability is that the closed-loop system is ISS considering the disturbance as the input

Desired behavior with disturbance

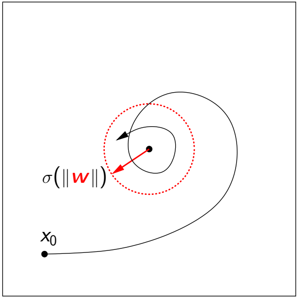

ISS in pictures

$$x^+ \in f(x, w)$$
$$|x(k; x_0, \boldsymbol{w})| \leq \beta(|x_0|, k) + \sigma(\|\boldsymbol{w}\|_{0:k-1})$$

Note also that ISS implies the desirable behavior that if $w(k) \to 0$ as $k \to \infty$, then $x(k; x_0, \boldsymbol{w}) \to 0$ also.

Robust positive invariance

We also require that the system not leave an invariant set due to the disturbance.

Definition 2 (Robust Positive Invariance)

A set $\mathcal{X} \subseteq \mathbb{R}^n$ is robustly positive invariant with respect to a difference inclusion $x^+ \in f(x, w)$ if there exists some $\delta > 0$ such that $f(x, w) \subseteq \mathcal{X}$ for all $x \in \mathcal{X}$ and all disturbance sequences $\boldsymbol{w}$ satisfying $\|\boldsymbol{w}\| \leq \delta$.

Robust asymptotic stability

So, we define robust asymptotic stability as input-to-state stability on a robust positive invariant set.

Definition 3 (Robust Asymptotic Stability)

The origin of a perturbed difference inclusion $x^+ \in f(x, w)$ is robustly asymptotically stable in $\mathcal{X}$ if there exists functions $\beta(\cdot) \in \mathcal{KL}$ and $\gamma(\cdot) \in \mathcal{K}$ and $\delta > 0$ such that for all $x \in \mathcal{X}$ and $\|\boldsymbol{w}\| \leq \delta$, $\mathcal{X}$ is robustly positive invariant and all solutions $\phi(k; x, \boldsymbol{w})$ satisfy

$$|\phi(k; x, \boldsymbol{w})| \leq \beta(|x|, k) + \gamma(\|\boldsymbol{w}\|) \tag{1}$$

for all $k \in \mathbb{I}_{\geq 0}$.

Input-to-state stability Lyapunov function

In order to establish ISS, we define an ISS Lyapunov function for a difference inclusion, similar to ISS Lyapunov function defined in Jiang and Wang (2001) and Lazar, Heemels, and Teel (2013).

Definition 4 (ISS Lyapunov Function)

$V(\cdot)$ is an ISS Lyapunov function in the robust positive invariant set $\mathcal{X}$ for the difference inclusion $x^+ \in f(x, w)$ if there exists some $\delta > 0$, functions $\alpha_1(\cdot), \alpha_2(\cdot), \alpha_3(\cdot) \in \mathcal{K}_\infty$, and function $\sigma(\cdot) \in \mathcal{K}$ such that for all $x \in \mathcal{X}$ and $\|\boldsymbol{w}\| \leq \delta$

$$\alpha_1(|x|) \leq V(x) \leq \alpha_2(|x|) \tag{2}$$

$$\sup_{x^+ \in f(x,w)} V(x^+) \leq V(x) - \alpha_3(|x|) + \sigma(\|\boldsymbol{w}\|) \tag{3}$$

ISS Lyapunov function implies ISS

Proposition 5 (ISS Lyapunov stability theorem)

If a difference inclusion $x^+ \in f(x, w)$ admits an ISS Lyapunov function in a robust positive invariant set $\mathcal{X}$ for all $\|\boldsymbol{w}\| \leq \delta$ for some $\delta > 0$, then the origin is robustly asymptotically stable in $\mathcal{X}$ for all $\|\boldsymbol{w}\| \leq \delta$.

- This is a valuable result to know when trying to establish robustness of stability.
- Let's skip this proof (hooray!), but it's not difficult (Jiang and Wang, 2001; Allan, Bates, Risbeck, and Rawlings, 2017).

Inherent robustness of nominal MPC

- Our strategy now is to establish that $V_N^0(x)$ is an ISS Lyapunov function for the perturbed closed-loop system.
- We have already established the upper and lower bounding inequalities

$$\alpha_1(|x|) \leq V_N^0(x) \leq \alpha_2(|x|)$$

- So we require only

$$\sup_{x^+ \in f(x,w)} V_N^0(x^+) \leq V_N^0(x) - \alpha_3(|x|) + \sigma(\|\boldsymbol{w}\|)$$

- That plus robust positive invariance, and we've established RAS of the controlled system.

Picture of the argument we are going to make

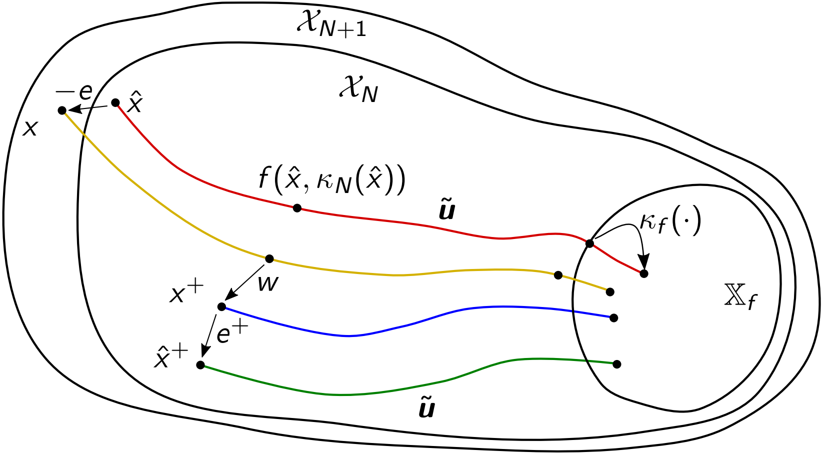

We have that $\hat{x}^+ = f(\hat{x} - e, \kappa_N(\hat{x})) + w + e^+$

We next compute difference in cost of red and green using $\tilde{\boldsymbol{u}}$

Note that $\tilde{\boldsymbol{u}}$ is feasible also for green, i.e., terminates in $\mathbb{X}_f := \text{lev}\, V_f$.

A useful tool for invoking continuity

Continuity in the language of K-functions

The usual ϵ-δ definition of continuity is equivalent to the following K-function definition (Rawlings and Risbeck, 2015).

Definition 6 (Continuity: K-function)

A function $f : \mathbb{R}^n \to \mathbb{R}^m$ is continuous at x if there exists a K-function $\gamma(\cdot)$ (note that the function $\gamma(\cdot)$ may depend on x) such that

$$|f(x+p) - f(x)| \leq \gamma(|p|) \quad \text{for all } |p| \in \text{Dom}(\gamma) \tag{4}$$

OK, let's jump in (Allan et al., 2017)

Since $V_N(x, \boldsymbol{u})$ is a continuous function

$$\left|V_N(\hat{x}^+, \tilde{\boldsymbol{u}}) - V_N(f(\hat{x}, \kappa_N(\hat{x})), \tilde{\boldsymbol{u}})\right| \leq \sigma_V(\left|\hat{x}^+ - f(\hat{x}, \kappa_N(\hat{x}))\right|)$$

with $\sigma_V(\cdot) \in \mathcal{K}$ (note we are *not* using the possibly discontinuous $V_N^0(x)$ here). Since $f(x, u)$ is also continuous

$$\begin{aligned}
\left|\hat{x}^+ - f(\hat{x}, \kappa_N(\hat{x}))\right| &= \left|f(\hat{x} + e, \kappa_N(\hat{x})) + w + e^+ - f(\hat{x}, \kappa_N(\hat{x}))\right| \\
&\leq \left|f(\hat{x} + e, \kappa_N(\hat{x})) - f(\hat{x}, \kappa_N(\hat{x}))\right| + |w| + \left|e^+\right| \\
&\leq \sigma_f(|e|) + |w| + \left|e^+\right| \\
&\leq \sigma_f(|d|) + 2\,|d| \leq \tilde{\sigma}_f(|d|)
\end{aligned}$$

with $d := (e, w, e^+)$ and $\tilde{\sigma}_f(\cdot) := \sigma_f(\cdot) + 2(\cdot) \in \mathcal{K}$. Therefore

$$\begin{aligned}
\left|V_N(\hat{x}^+, \tilde{\boldsymbol{u}}) - V_N(f(\hat{x}, \kappa_N(\hat{x})), \tilde{\boldsymbol{u}})\right| &\leq \sigma_V(\tilde{\sigma}_f(|d|)) := \sigma(|d|) \\
V_N(\hat{x}^+, \tilde{\boldsymbol{u}}) &\leq V_N(f(\hat{x}, \kappa_N(\hat{x})), \tilde{\boldsymbol{u}}) + \sigma(|d|)
\end{aligned}$$

with $\sigma(\cdot) \in \mathcal{K}$.

Last steps

Note that for the candidate sequence
$V_N(f(\hat{x}, \kappa_N(\hat{x})), \tilde{\boldsymbol{u}}) \leq V_N^0(\hat{x}) - \ell(\hat{x}, \kappa_N(\hat{x}))$ so we have that

$$V_N(f(\hat{x}, \kappa_N(\hat{x})), \tilde{\boldsymbol{u}}) \leq V_N^0(\hat{x}) - \alpha_1(|\hat{x}|)$$

since $\alpha_1(|x|) \leq \ell(x, \kappa_N(x))$ for all x. Therefore, we finally have

$$\begin{aligned} V_N(\hat{x}^+, \tilde{\boldsymbol{u}}) &\leq V_N^0(\hat{x}) - \alpha_1(|\hat{x}|) + \sigma(|d|) \\ V_N^0(\hat{x}^+) &\leq V_N^0(\hat{x}) - \alpha_1(|\hat{x}|) + \sigma(\|\boldsymbol{d}\|) \end{aligned}$$

and we have established that $V_N^0(\cdot)$ is an ISS-Lyapunov function!
That plus robust invariance gives robust asymptotic stability of $\hat{x}$. Since $x = \hat{x} + e$, that gives also RAS of x.
Notice that neither $V_N^0(\cdot)$ nor $\kappa_N(\cdot)$ need be continuous for MPC to be inherently robust.

Discrete actuators

In addition to continuous actuators, many process systems also have discrete actuators that are constrained to be *integers*.

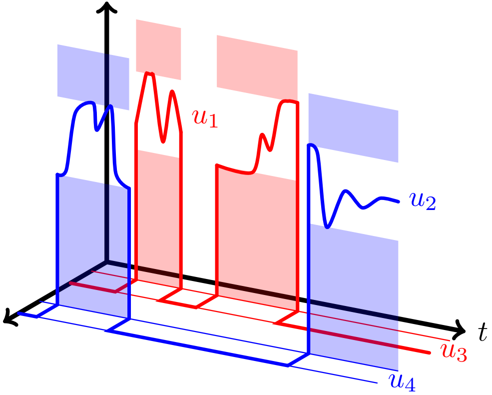

- Processes with banks of furnaces, heaters, chillers, etc.
- Scheduling models with discrete decisions.
- Switched systems with input-dependent dynamics.
- Semi-continuous variables (e.g. $u \in \{0\} \cup [1, 2]$).

Continuous and mixed continuous-discrete actuators

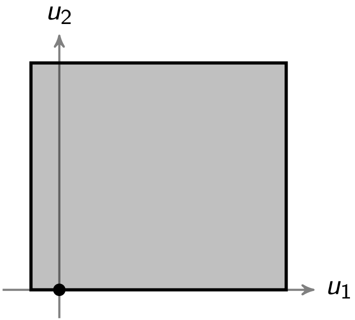

(a) (b)

Typical input constraint sets $\mathbb{U}$ for (a) continuous actuators and (b) mixed continuous-discrete actuators; the origin (•) is the equilibrium of interest.

Example: Driving a manual transmission

- State: vehicle velocity v
- Inputs: engine RPM $\omega \in [0, \omega_{\max}]$
  gear $\gamma \in \{1, 2, 3, 4, 5\}$

$$\frac{dv}{dt} = a_{\max}(\gamma)\, \sigma\left(R(\gamma)\omega - v\right)$$

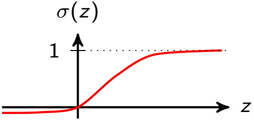

- Maximum acceleration $a_{\max}(\gamma)$ decreases for higher gears
- Final velocity $v = R(\gamma)\omega$ increases for higher gears

Choose setpoint v_{sp} and use tracking stage cost

$$\ell(v, \omega, \gamma) = \underbrace{20\left(\frac{v}{v_{\mathrm{sp}}} - 1\right)^2}_{\text{Track } v_{\mathrm{sp}}} + \underbrace{8 \max\left(0, \frac{\omega - \omega_{\mathrm{ss}}}{\omega_{\max}}\right)}_{\text{Minimize excessive } \omega} + \underbrace{(\Delta\gamma)^2}_{\text{Restrict switching}}$$

Example Simulation

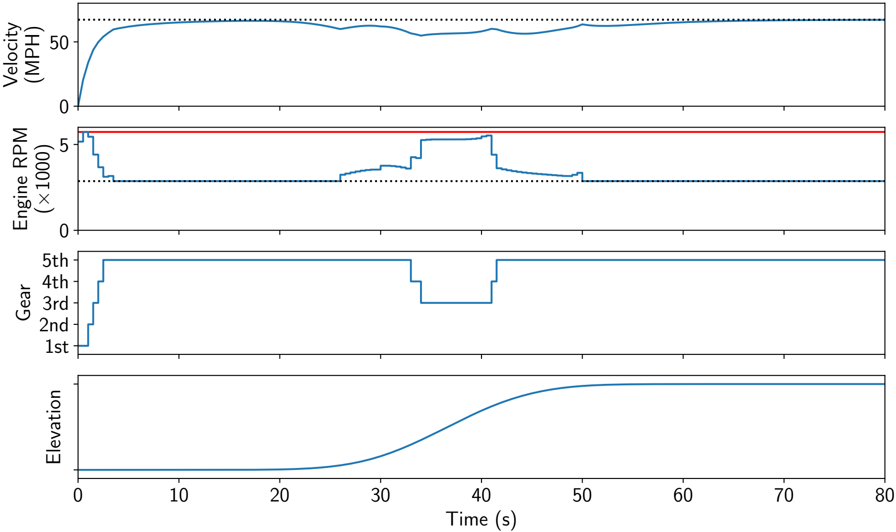

Figure 1: Closed-loop evolution of car system. Optimization performed using Bonmin.

Inherent Robustness—Extension to discrete actuators

- The extension to discrete actuators is immediate
- The set $\mathbb{U}$ need not be convex, connected, etc.—it need only contain the origin

However, design choices become more striking with discrete actuators:

- Theory forbids “large” control action near the setpoint
  - System must be locally stabilizable using only unsaturated actuators
  - Discrete actuators are always saturated
- Single setpoint stabilization may no longer be an appropriate goal

Feasible Sets

- MPC is stabilizing on $\mathcal{X}_N$ but $\mathcal{X}_N$ may not be what you expect

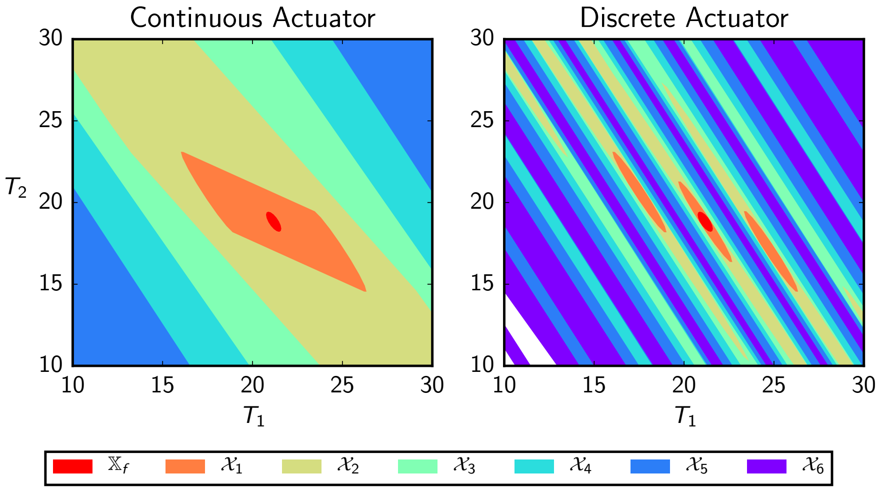

Conclusion

- We have extended standard MPC theory to handle discrete actuators for robust stabilization of an equilibrium point
- This theory extends to periodic trajectories and economic MPC
- Based on these results we offer the following conjecture:

Theorem 7 (Folk theorem)

Any result that holds for standard MPC holds also for MPC with discrete actuators. (Rawlings and Risbeck, 2017)

- Applications include a rich class of commercial building energy optimization problems
- A current challenge is to develop better software tools for efficient, reliable *online* solution of the mixed-integer optimal control problems. See `casadi.org`

Further reading I

D. A. Allan, C. N. Bates, M. J. Risbeck, and J. B. Rawlings. On the inherent robustness of optimal and suboptimal nonlinear MPC. *Sys. Cont. Let.*, 106:68 – 78, 2017. ISSN 0167-6911. doi: 10.1016/j.sysconle.2017.03.005.

Z.-P. Jiang and Y. Wang. Input-to-state stability for discrete-time nonlinear systems. *Automatica*, 37:857–869, 2001.

M. Lazar, W. Heemels, and A. Teel. Further input-to-state stability subtleties for discrete-time systems. *IEEE Trans. Auto. Cont.*, 58(6): 1609–1613, Jun 2013. ISSN 0018-9286. doi: 10.1109/TAC.2012.2231611.

J. B. Rawlings and M. J. Risbeck. On the equivalence between statements with epsilon-delta and K-functions. Technical Report 2015–01, TWCCC Technical Report, December 2015. URL https://engineering.ucsb.edu/~jbraw/jbrweb-archives/tech-reports/twccc-2015-01.pdf.

Further reading II

J. B. Rawlings and M. J. Risbeck. Model predictive control with discrete actuators: Theory and application. *Automatica*, 78:258–265, 2017.

J. B. Rawlings, D. Q. Mayne, and M. M. Diehl. *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, Santa Barbara, CA, 2nd, paperback edition, 2020. 770 pages, ISBN 978-0-9759377-5-4.

A. R. Teel. Discrete time receding horizon control: is the stability robust. In Marcia S. de Queiroz, Michael Malisoff, and Peter Wolenski, editors, *Optimal control, stabilization and nonsmooth analysis*, volume 301 of *Lecture notes in control and information sciences*, pages 3–28. Springer, 2004.

Review

Recommended exercises

- Stability definitions. Exercise B.8.[1]
- Lyapunov functions. Exercise B.2–B.3.[1]
- Dynamic programming. Exercise C.1–C.2.[1]
- MPC stability results. Exercises 2.12, 2.13[1]

[1] Rawlings, Mayne, and Diehl (2020, Chapter 2, Appendices B and C). Downloadable from `engineering.ucsb.edu/~jbraw/mpc`.

Computational Exercise

Consider the following system:

$$\frac{d}{dt}x = f(x) + g(x)u$$
$$\frac{d}{dt}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} -1 & -1 \\ 1 & 1 \end{pmatrix}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + \begin{pmatrix} -x_2 & 0 \\ x_1 & 1 \end{pmatrix}\begin{pmatrix} u_1 \\ u_2 \end{pmatrix}$$
$$\begin{pmatrix} -1 \\ -1 \end{pmatrix} \leq \begin{pmatrix} u_1 \\ u_2 \end{pmatrix} \leq \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

- For fixed u_1, system is linear.
- Far from the origin, system is difficult to stabilize along the x_2-axis.

Exercise

Design a nonlinear MPC controller to regulate the system to the origin.

- Cost functions: $\ell(x,u) = 100x'x + u'u$, $P_f(x) = 1000x'x$
- State is measured.
- No disturbances.

Compare results to linear MPC.

- Why might linear MPC be a bad idea for this system?
- Can linear MPC stabilize the system? Where?

Hints

Start with the linearized problem.

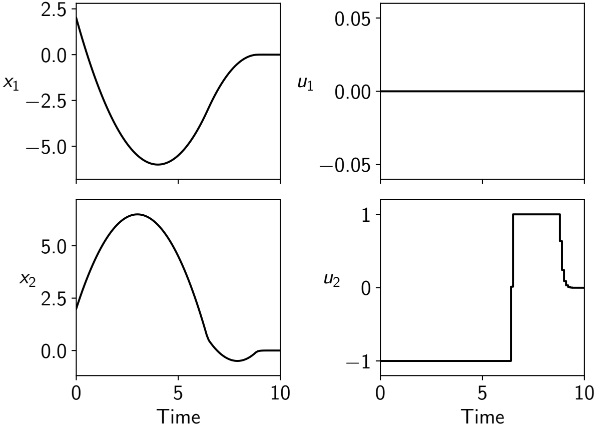

Figure 2: Trajectory using linearized system and linear MPC.

Hints

Adding nonlinearities, you should get something like this:

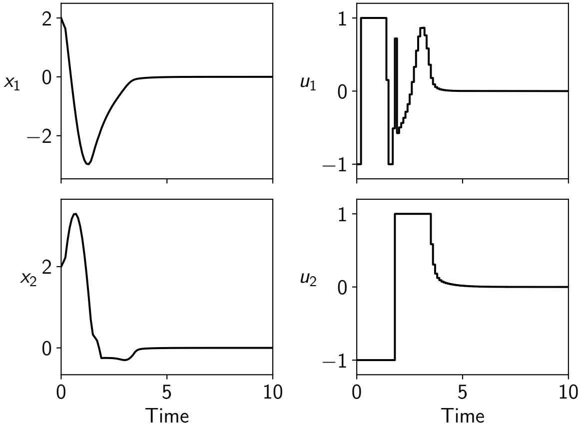

Figure 3: Trajectory using nonlinear MPC.

Finally, you can compare both on the same axes:

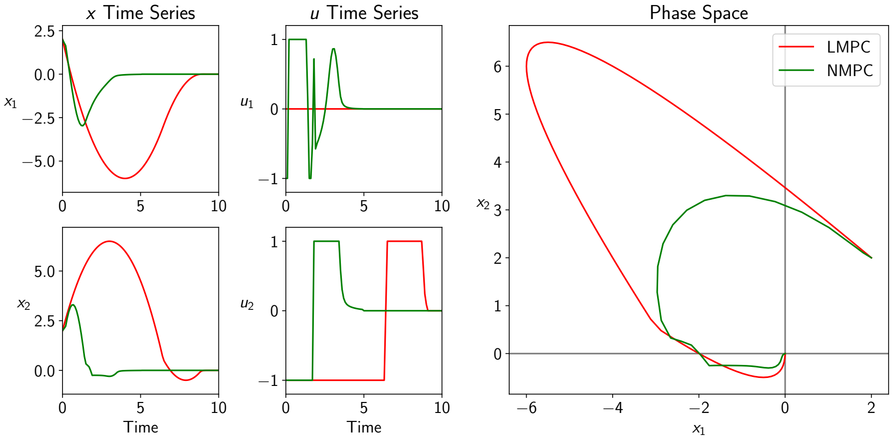

Figure 4: Comparison of linear and nonlinear MPC trajectories.

Computational Exercise 2

Consider the CSTR Example from earlier

Nonlinear CSTR

An irreversible, first-order reaction A→B occurs in the liquid phase and the reactor temperature is regulated with external cooling.

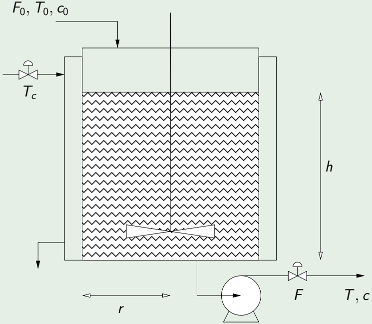

Computational Exercise 2

Simulation Parameters

1. Initial Condition and Sample Time

$$x_0 = \begin{bmatrix} 0.05c^s \\ 0.75\,T^s \\ 0.5h^s \end{bmatrix} \qquad \Delta = 0.25\,\mathrm{min}$$

2. Input Constraints

$$\begin{bmatrix} 0.975\,T_c^s \\ 0.75F^s \end{bmatrix} \leq u \leq \begin{bmatrix} 1.025\,T_c^s \\ 1.25F^s \end{bmatrix}$$

Computational Exercise 2

Reactor Startup

Using the model and parameters provided previously,

1. Simulate the performance of an uncontrolled startup by injecting the steady-state input into the system. Does the system reach the desired operating point?
2. Use linear MPC to simulate the same startup. Does the system reach the desired operating point with a linear controller?
3. Repeat the startup, but with nonlinear MPC. Does the system reach the desired operating point with a nonlinear controller? Comment on the performance differences between the various approaches.

Reactor Startup

The uncontrolled startup does not drive the reactor to the desired steady state, however both the linear and nonlinear MPC controllers do.

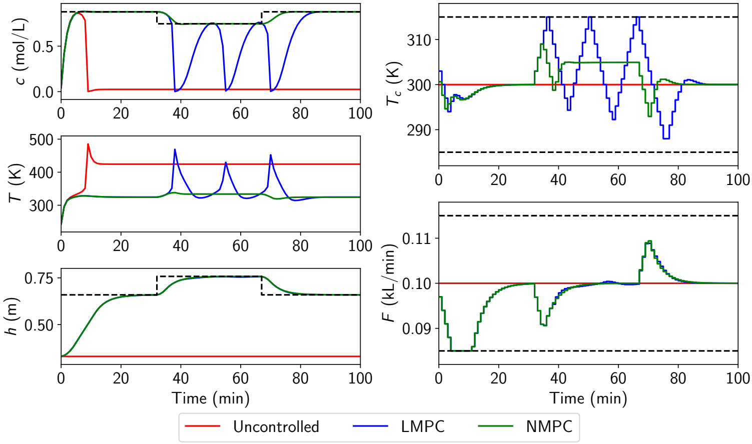

Figure 5: Solution for Reactor Startup Exercise.

Ball Maze

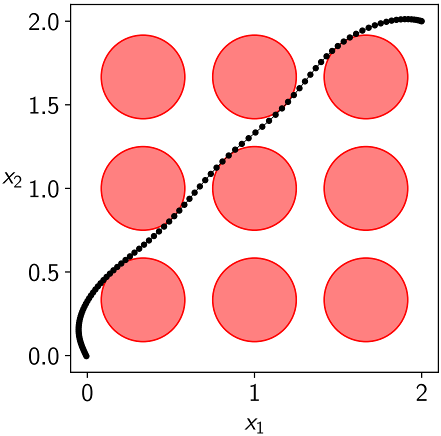

Figure 6: MPC navigating a ball maze. Although the constraints are nonconvex, we can still find a local solution.

Airplane Descent

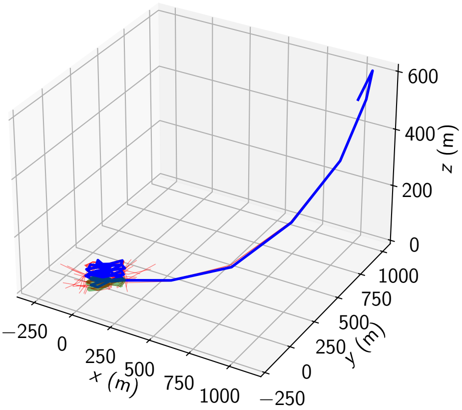

Figure 7: MPC for guiding a descending plane. While the goal is to reach a periodic holding pattern, the optimizer does not find that solution due to nonconvexity.