AA203
Optimal and Learning-based Control

Introduction to MPC, persistent feasibility

Stanford University

Autonomous Systems Lab

Roadmap

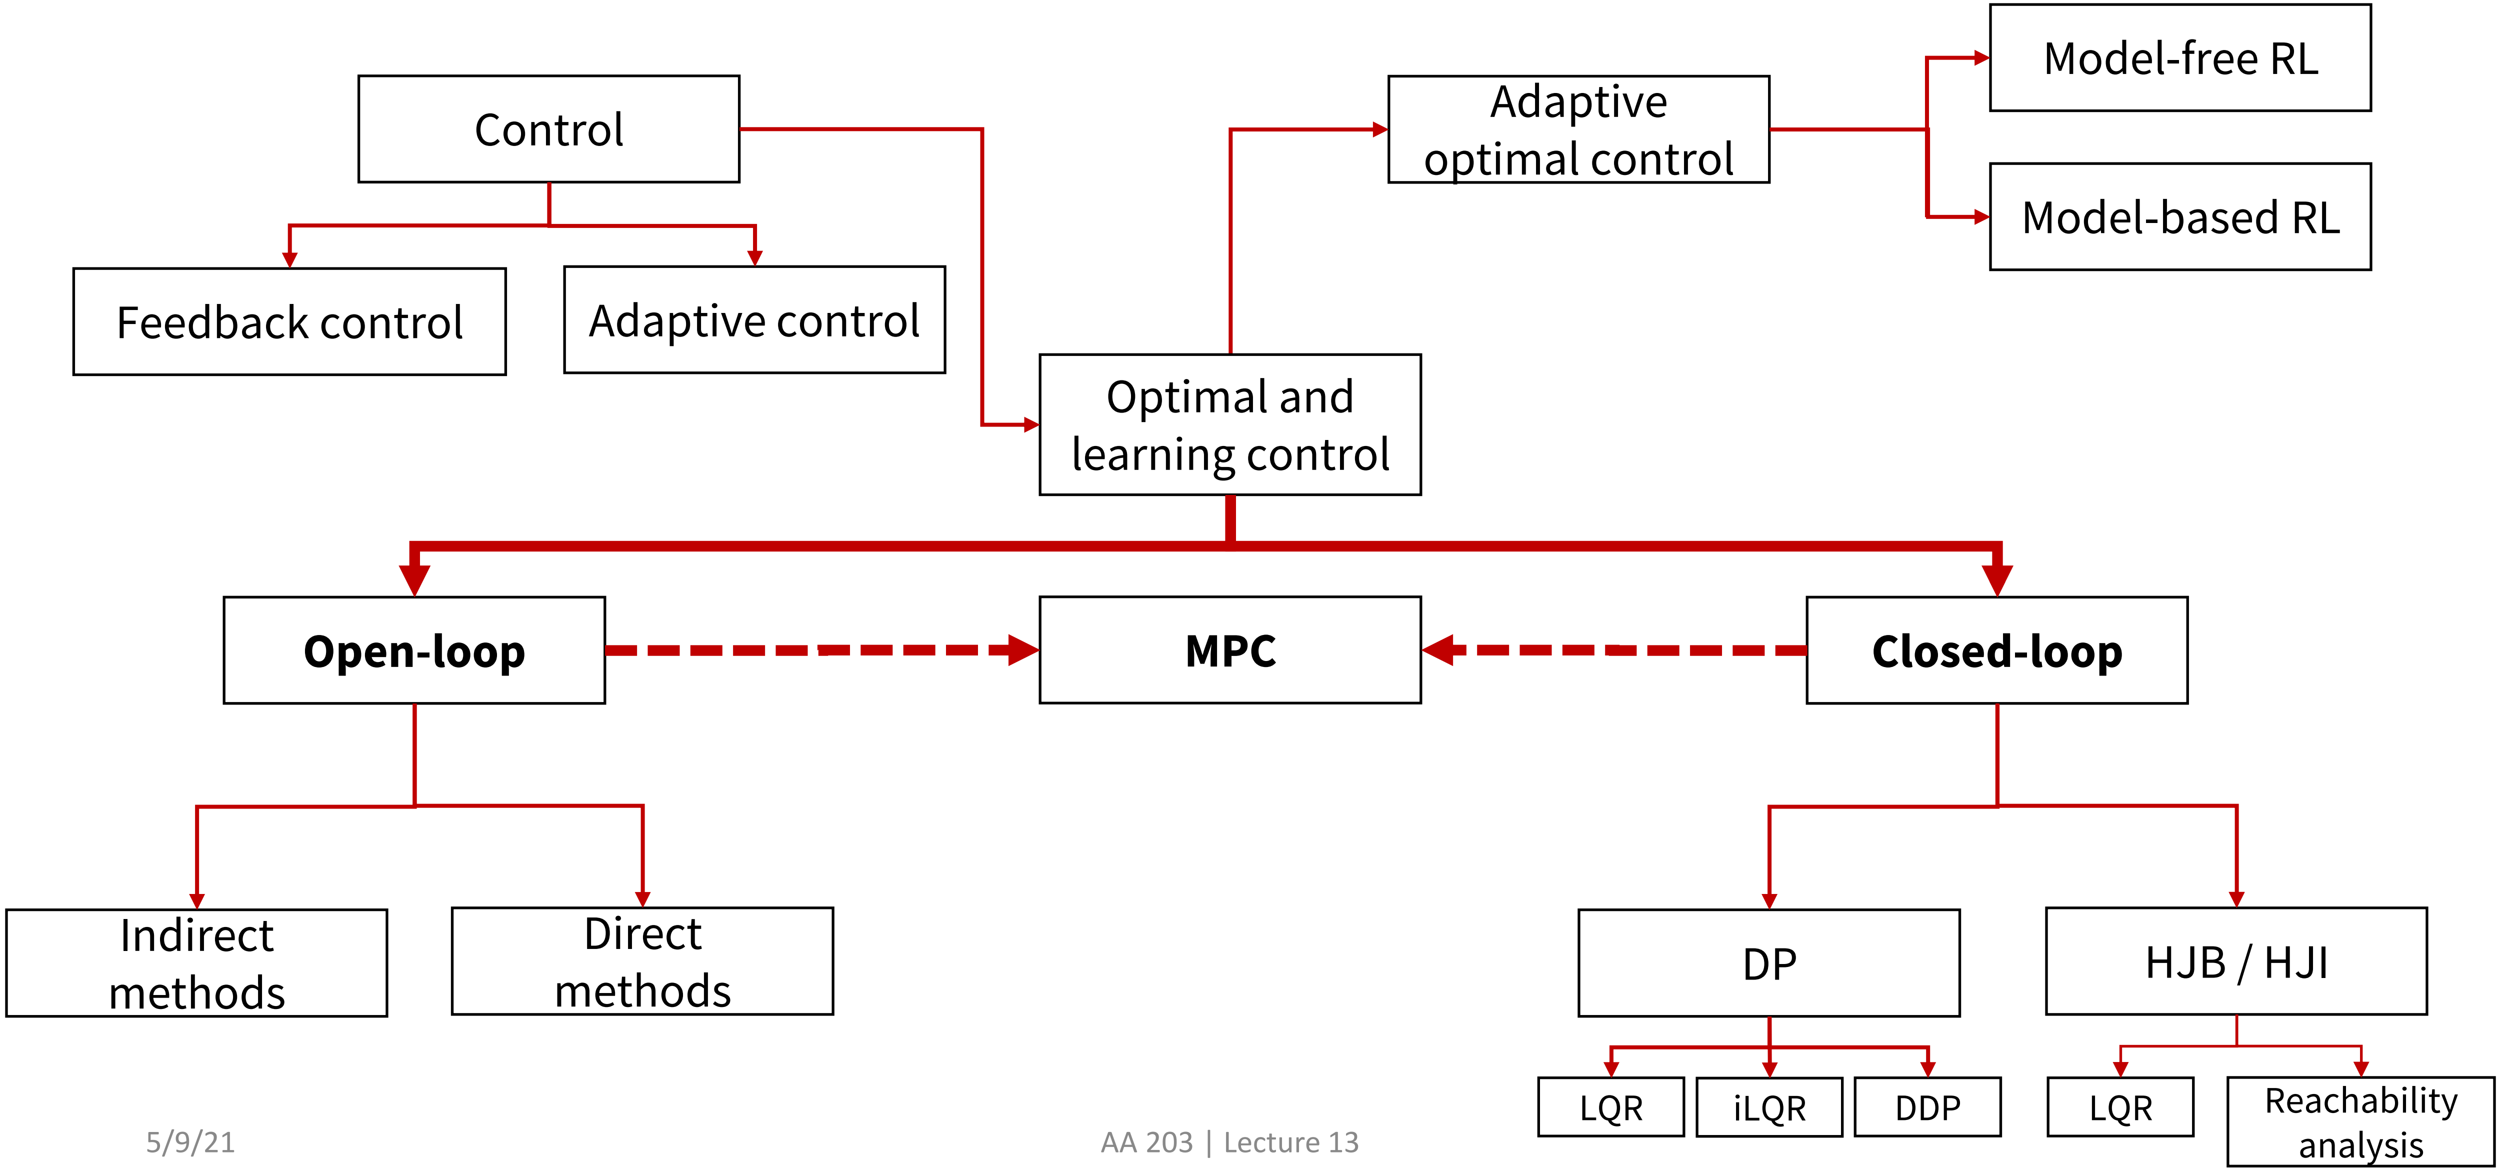

Control
Feedback control
Adaptive control
Optimal and learning control
Adaptive optimal control
Model-free RL
Model-based RL
Open-loop
MPC
Closed-loop
Indirect methods
Direct methods
DP
HJB / HJI
LQR
iLQR
DDP
LQR
Reachability analysis

Model predictive control

- Introduction: basic setting and key ideas
- Persistent feasibility of MPC

- Further reading:
  - F. Borrelli, A. Bemporad, M. Morari. *Predictive Control for Linear and Hybrid Systems*, 2017.
  - J. B. Rawlings, D. Q. Mayne, M. M. Diehl. *Model Predictive Control: Theory, Computation, and Design*, 2017.

Model predictive control

- Model predictive control (or, more broadly, receding horizon control) entails solving finite-time optimal control problems in a receding horizon fashion

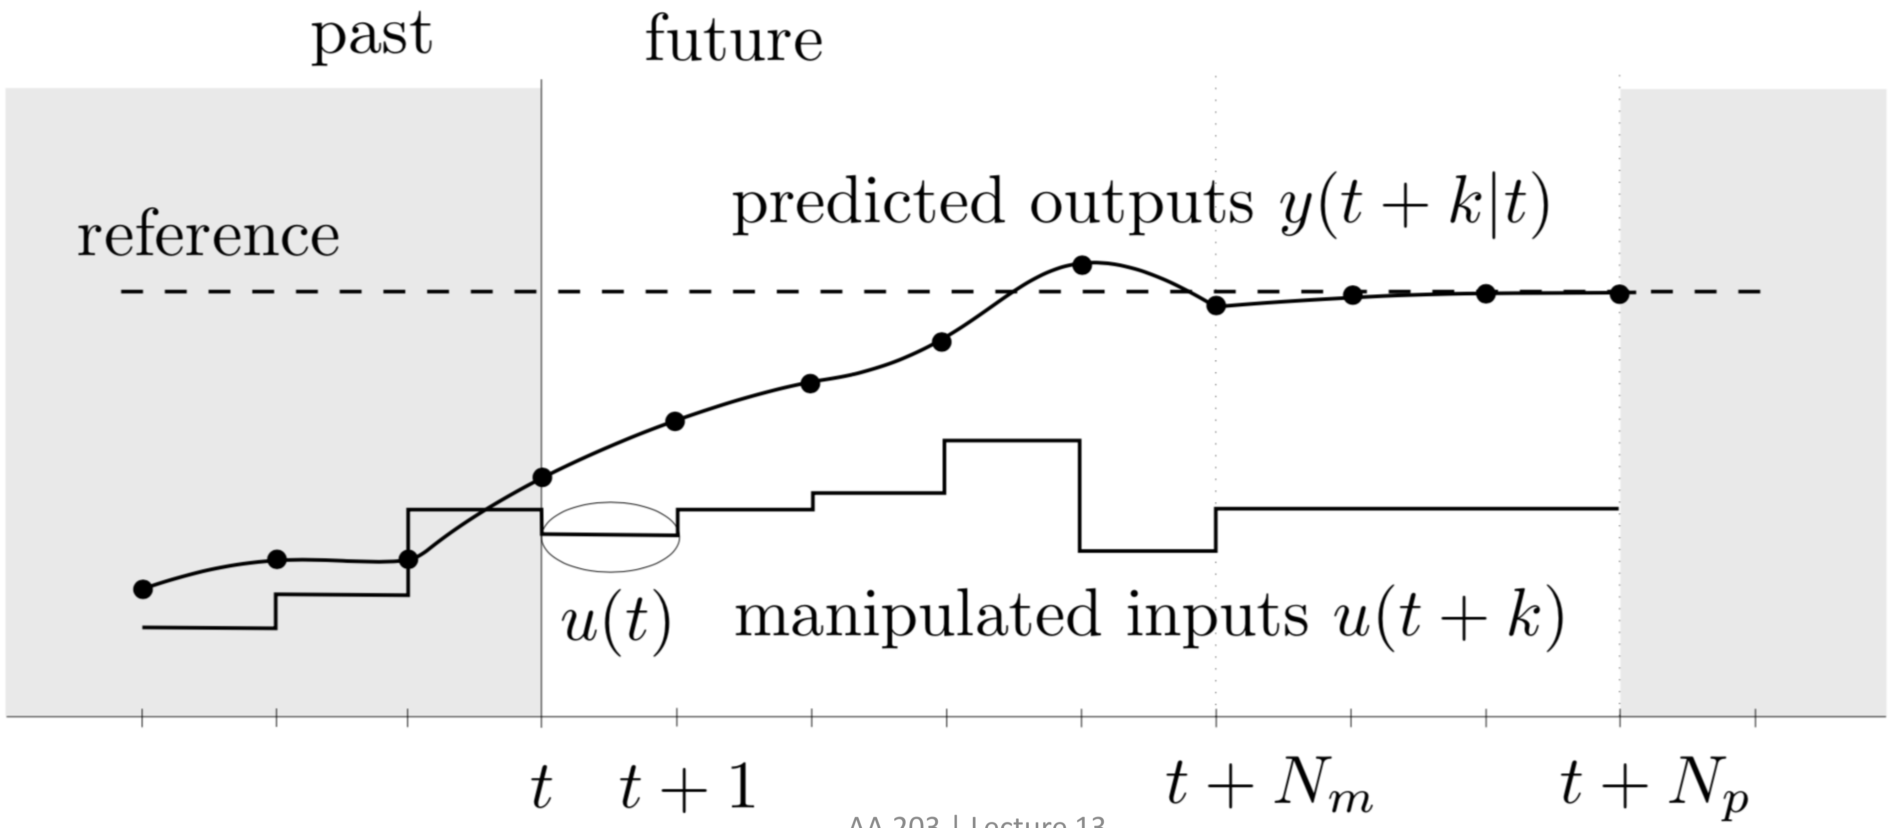

Model predictive control

- Model predictive control (or, more broadly, receding horizon control) entails solving finite-time optimal control problems in a receding horizon fashion

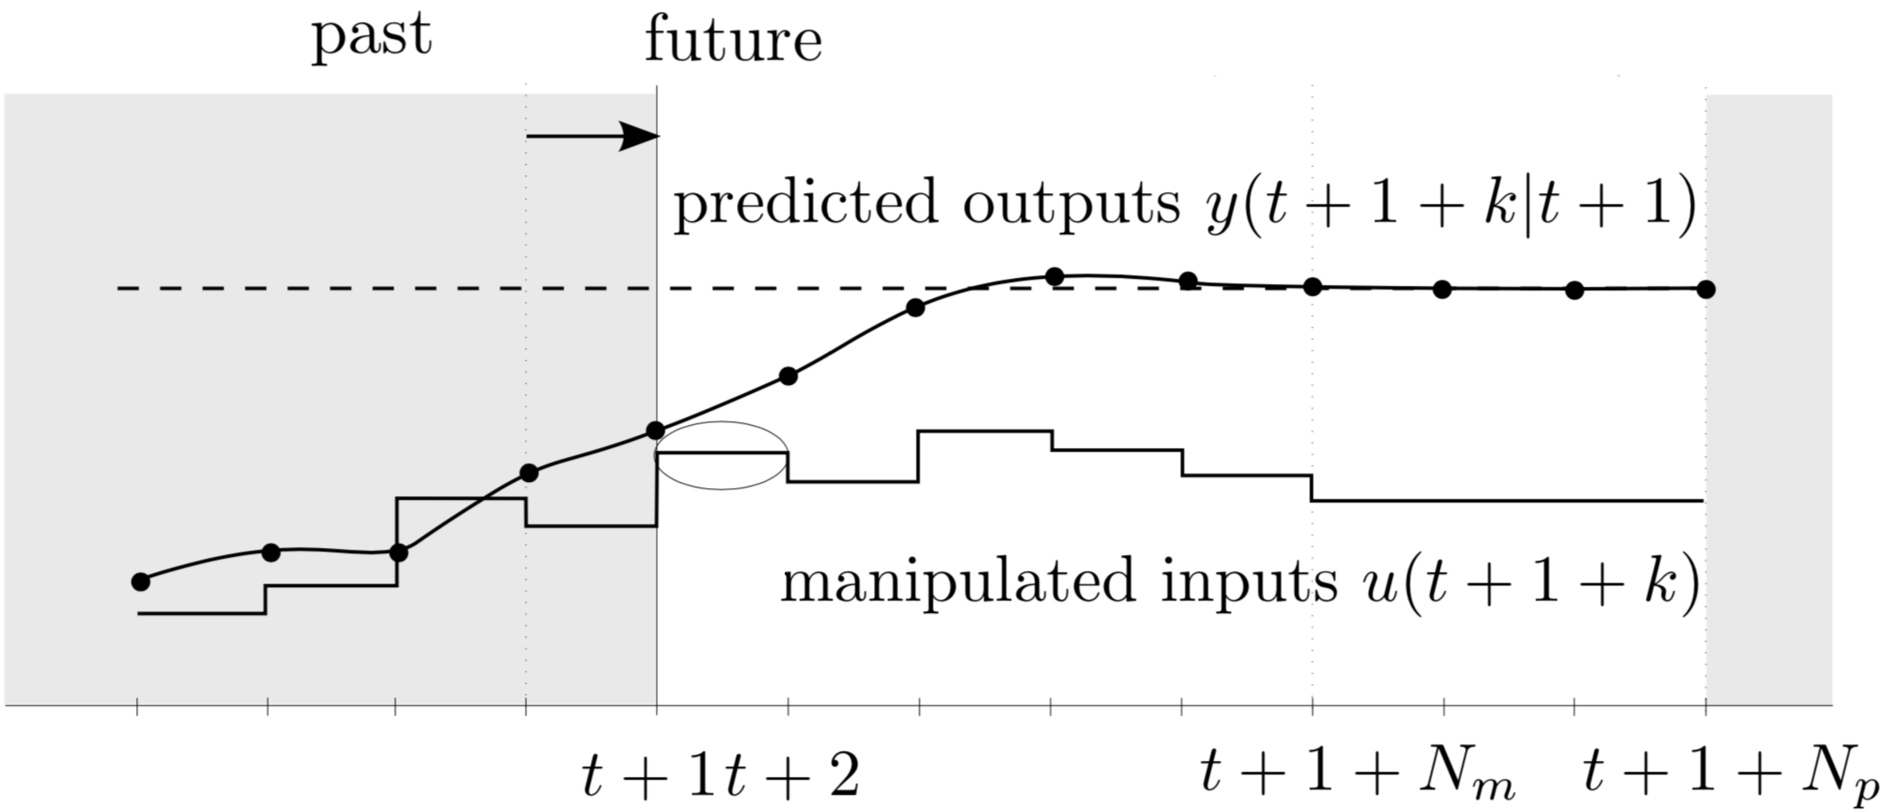

Model predictive control

Key steps:

1. At each sampling time t, solve an *open-loop* optimal control problem over a finite horizon
2. Apply optimal input signal during the following sampling interval $[t, t+1)$
3. At the next time step $t+1$, solve new optimal control problem based on new measurements of the state over a shifted horizon

Basic formulation

- Consider the problem of regulating to the origin the discrete-time linear time-invariant system

$$\mathbf{x}(t+1) = A\mathbf{x}(t) + B\mathbf{u}(t), \qquad \mathbf{x}(t) \in \mathbb{R}^n, \quad \mathbf{u}(t) \in \mathbb{R}^m$$

subject to the constraints

$$\mathbf{x}(t) \in X, \qquad \mathbf{u}(t) \in U, \qquad t \geq 0$$

where the sets X and U are *polyhedra*

Basic formulation

- Assume that a full measurement of the state $\mathbf{x}(t)$ is available at the current time t
- The finite-time optimal control problem solved at each stage is

$$J_t^*\big(\mathbf{x}(t)\big) = \min_{\mathbf{u}_{t|t},\ldots,\mathbf{u}_{t+N-1|t}} p\big(\mathbf{x}_{t+N|t}\big) + \sum_{k=0}^{N-1} c(\mathbf{x}_{t+k|t}, \mathbf{u}_{t+k|t})$$

$$\begin{aligned}
\text{subject to} \quad & \mathbf{x}_{t+k+1|t} = A\mathbf{x}_{t+k|t} + B\mathbf{u}_{t+k|t}, \quad && k = 0, \ldots, N-1 \\
& \mathbf{x}_{t+k|t} \in X, \quad \mathbf{u}_{t+k|t} \in U, && k = 0, \ldots, N-1 \\
& \mathbf{x}_{t+N|t} \in X_f \\
& \mathbf{x}_{t|t} = \mathbf{x}(t)
\end{aligned}$$

Basic formulation

- Assume that a full measurement of the state $\mathbf{x}(t)$ is available at the current time t
- The finite-time optimal control problem solved at each stage is

$$J_t^*(\mathbf{x}(t)) = \min_{\mathbf{u}_{t|t},\ldots,\mathbf{u}_{t+N-1|t}} \boxed{p(\mathbf{x}_{t+N|t})} + \sum_{k=0}^{N-1} c(\mathbf{x}_{t+k|t}, \mathbf{u}_{t+k|t})$$

$$\begin{aligned} \text{subject to} \quad & \mathbf{x}_{t+k+1|t} = A\mathbf{x}_{t+k|t} + B\mathbf{u}_{t+k|t}, && k = 0, \ldots, N-1 \\ & \mathbf{x}_{t+k|t} \in X, \quad \mathbf{u}_{t+k|t} \in U, && k = 0, \ldots, N-1 \\ & \boxed{\mathbf{x}_{t+N|t} \in X_f} \\ & \mathbf{x}_{t|t} = \mathbf{x}(t) \end{aligned}$$

Key MPC design choices!

Basic formulation

Notation:

- $\mathbf{x}_{t+k|t}$ is the state vector at time $t + k$ predicted at time t (via the system's dynamics)
- $\mathbf{u}_{t+k|t}$ is the input $\mathbf{u}$ at time $t + k$ computed at time t

Note: $\mathbf{x}_{3|1} \neq \mathbf{x}_{3|2}$

Basic formulation

- Let $U^*_{t\rightarrow t+N|t} \coloneqq \{\mathbf{u}^*_{t|t}, \mathbf{u}^*_{t+1|t}, \dots, \mathbf{u}^*_{t+N-1|t}\}$ be the optimal solution, then
$$\mathbf{u}(t) = \mathbf{u}^*_{t|t}(\mathbf{x}(t))$$
- The optimization problem is then repeated at time $t+1$, based on the new state $\mathbf{x}_{t+1|t+1} = \mathbf{x}(t+1)$
- Define $\pi_t\big(\mathbf{x}(t)\big) \coloneqq \mathbf{u}^*_{t|t}(\mathbf{x}(t))$
- Then the closed-loop system evolves as
$$\mathbf{x}(t+1) = A\mathbf{x}(t) + B\pi_t\big(\mathbf{x}(t)\big) \coloneqq \mathbf{f}_{\mathrm{cl}}(\mathbf{x}(t), t)$$
- Central question: characterize the behavior of closed-loop system

Simplifying the notation

- Note that the setup is time-invariant, hence, to simplify the notation, we can let $t = 0$ in the finite-time optimal control problem, namely

$$J_0^*(\mathbf{x}(t)) = \min_{\mathbf{u}_0,\ldots,\mathbf{u}_{N-1}} p(\mathbf{x}_N) + \sum_{k=0}^{N-1} c(\mathbf{x}_k, \mathbf{u}_k)$$

$$\begin{aligned} \text{subject to} \quad & \mathbf{x}_{k+1} = A\mathbf{x}_k + B\mathbf{u}_k, && k = 0, \ldots, N-1 \\ & \mathbf{x}_k \in X, \quad \mathbf{u}_k \in U, && k = 0, \ldots, N-1 \\ & \mathbf{x}_N \in X_f \\ & \mathbf{x}_0 = \mathbf{x}(t) \end{aligned}$$

- Denote $U_0^*(\mathbf{x}(t)) = \{\mathbf{u}_0^*, \ldots, \mathbf{u}_{N-1}^*\}$

Simplifying the notation

- With new notation,

$$\mathbf{u}(t) = \mathbf{u}_0^*\big(\mathbf{x}(t)\big) = \pi(\mathbf{x}(t))$$

and closed-loop system becomes

$$\mathbf{x}(t+1) = A\mathbf{x}(t) + B\pi\big(\mathbf{x}(t)\big) \coloneqq \mathbf{f}_{\mathrm{cl}}(\mathbf{x}(t))$$

Typical cost functions

- 2-norm:

$$p(\mathbf{x}_N) = \mathbf{x}_N^{\mathrm{T}} P \mathbf{x}_N, \quad c(\mathbf{x}_k, \mathbf{u}_k) = \mathbf{x}_k^{\mathrm{T}} Q \mathbf{x}_k + \mathbf{u}_k^{\mathrm{T}} R \mathbf{u}_k, \quad P \succcurlyeq 0, Q \succcurlyeq 0, R \succ 0$$

- 1-norm or ∞-norm:

$$p(\mathbf{x}_N) = \|P\mathbf{x}_N\|_p \quad c(\mathbf{x}_k, \mathbf{u}_k) = \|Q\mathbf{x}_k\|_p + \|R\mathbf{u}_k\|_p, \qquad p = 1 \text{ or } \infty$$

where P, Q, R are full column ranks

Online model predictive control

repeat

measure the state $\mathbf{x}(t)$ at time instant t

obtain $U_0^*(\mathbf{x}(t))$ by solving finite-time optimal control problem

if $U_0^*(\mathbf{x}(t)) = \emptyset$ **then** ‘problem infeasible’ **stop**

apply the first element $\mathbf{u}_0^*$ of $U_0^*(\mathbf{x}(t))$ to the system

wait for the new sampling time $t + 1$

Main implementation issues

1. The controller may lead us into a situation where after a few steps the finite-time optimal control problem is infeasible → *persistent feasibility issue*

2. Even if the feasibility problem does not occur, the generated control inputs may not lead to trajectories that converge to the origin (i.e., closed-loop system is unstable) → *stability issue*

Key question: how do we guarantee that such a "short- sighted" strategy leads to effective long-term behavior?

Analysis approaches

1. Analyze closed-loop behavior directly → generally very difficult

2. Derive conditions on terminal function p, and terminal constraint set X_f so that persistent feasibility and closed-loop stability are guaranteed

Addressing persistent feasibility

Goal: design MPC controller so that feasibility for all future times is guaranteed

Approach: leverage tools from *invariant set theory*

Set of feasible initial states

- Set of feasible initial states

$$X_0 \coloneqq \{\mathbf{x}_0 \in X \mid \exists\ (\mathbf{u}_0, \dots, \mathbf{u}_{N-1}) \text{ such that } \mathbf{x}_k \in X, \mathbf{u}_k \in U, k = 0, \dots, N-1, \\ \mathbf{x}_N \in X_f \text{ where } \mathbf{x}_{k+1} = A\mathbf{x}_k + B\mathbf{u}_k, k = 0, \dots, N-1\}$$

- A control input can be found only if $\mathbf{x}(0) \in X_0$!

Controllable sets

- For the autonomous system $\mathbf{x}(t+1) = \phi(\mathbf{x}(t))$ with constraints $\mathbf{x}(t) \in X, \mathbf{u}(t) \in U$, the one-step controllable set to set S is defined as

$$\mathrm{Pre}(S) \coloneqq \{\mathbf{x} \in \mathbb{R}^n : \ \phi(\mathbf{x}) \in S\}$$

- For the system $\mathbf{x}(t+1) = \phi\big(\mathbf{x}(t), \mathbf{u}(t)\big)$ with constraints $\mathbf{x}(t) \in X$, $\mathbf{u}(t) \in U$, the one-step controllable set to set S is defined as

$$\mathrm{Pre}(S) \coloneqq \{\mathbf{x} \in \mathbb{R}^n : \exists \mathbf{u} \in U \text{ such that } \phi(\mathbf{x}, \mathbf{u}) \in S\}$$

Control invariant sets

- A set $C \subseteq X$ is said to be a control invariant set for the system $\mathbf{x}(t+1) = \phi(\mathbf{x}(t), \mathbf{u}(t))$ with constraints $\mathbf{x}(t) \in X,\ \mathbf{u}(t) \in U$, if:

$$\mathbf{x}(t) \in C \Rightarrow \exists \mathbf{u} \in U \text{ such that } \phi(\mathbf{x}(t), \mathbf{u}(t)) \in C, \text{ for all } t$$

- The set $C_\infty \subseteq X$ is said to be the maximal control invariant set for the system $\mathbf{x}(t+1) = \phi(\mathbf{x}(t), \mathbf{u}(t))$ with constraints $\mathbf{x}(t) \in X,\ \mathbf{u}(t) \in U$, if it is control invariant and contains all control invariant sets contained in X

- For autonomous systems: a set $A \subseteq X$ is said to be a positive invariant set for the system $\mathbf{x}(t+1) = \phi(\mathbf{x}(t))$ if $\mathbf{x}(t) \in A \Rightarrow \phi(\mathbf{x}(t)) \in A$; the maximal positive invariant set contains all other positive invariant sets.

- These sets can be computed by using the MPT toolbox https://www.mpt3.org/

Persistent feasibility lemma

- Define "truncated" feasibility set:

$$X_1 := \{\mathbf{x}_1 \in X \mid \exists\ (\mathbf{u}_1, \ldots, \mathbf{u}_{N-1}) \text{ such that } \mathbf{x}_k \in X, \mathbf{u}_k \in U, k = 1, \ldots, N-1, \\ \mathbf{x}_N \in X_f \text{ where } \mathbf{x}_{k+1} = A\mathbf{x}_k + B\mathbf{u}_k, k = 1, \ldots, N-1\}$$

- Feasibility lemma: if set X_1 is a *control invariant set* for system:

$$\mathbf{x}(t+1) = A\mathbf{x}(t) + B\mathbf{u}(t), \qquad \mathbf{x}(t) \in X, \ \ \mathbf{u}(t) \in U, \ \ t \geq 0$$

then the MPC law is persistently feasible

Persistent feasibility lemma

- Proof:

1. $\text{Pre}(X_1) = \{\mathbf{x} \in \mathbb{R}^n : \exists \mathbf{u} \in U \text{ such that } A\mathbf{x} + B\mathbf{u} \in X_1\}$
2. Since X_1 is control invariant
$$\forall \mathbf{x} \in X_1 \;\; \exists \mathbf{u} \in U \;\; \text{such that } A\mathbf{x} + B\mathbf{u} \in X_1$$
3. Thus $X_1 \subseteq \text{Pre}(X_1) \cap X$
4. One can write
$$X_0 = \{\mathbf{x}_0 \in X \mid \exists \mathbf{u}_0 \in U \text{ such that } A\mathbf{x}_0 + B\mathbf{u} \in X_1\} = \text{Pre}(X_1) \cap X$$
5. Thus, $X_1 \subseteq X_0$

Persistent feasibility lemma

- Proof:

6. Pick some $\mathbf{x}_0 \in X_0$. Let U_0^* be the solution to the finite-time optimization problem, and $\mathbf{u}_0^*$ be the first control. Let

$$\mathbf{x}_1 = A\mathbf{x}_0 + B\mathbf{u}_0^*$$

7. Since U_0^* is clearly feasible, one has $\mathbf{x}_1 \in X_1$. Since $X_1 \subseteq X_0$, one has

$$\mathbf{x}_1 \in X_0$$

hence the next optimization problem is feasible!

Practical significance

- For $N = 1$, we can set $X_f = X_1$. If we choose the terminal set to be control invariant, then MPC will be persistently feasible *independent* of chosen control objectives and parameters
- Designer can choose the parameters to affect performance (e.g., stability)
- How to extend this result to $N > 1$?

Persistent feasibility theorem

- Feasibility theorem: if set X_f is a *control invariant set* for system:

$$\mathbf{x}(t+1) = A\mathbf{x}(t) + B\mathbf{u}(t), \qquad \mathbf{x}(t) \in X, \;\; \mathbf{u}(t) \in U, \;\; t \geq 0$$

  then the MPC law is persistently feasible

Persistent feasibility theorem

- Proof

1. Define "truncated" feasibility set at step $N-1$:

$$X_{N-1} := \{\mathbf{x}_{N-1} \in X \mid \exists\, \mathbf{u}_{N-1} \text{ such that } \mathbf{x}_{N-1} \in X, \mathbf{u}_{N-1} \in U, \mathbf{x}_N \in X_f \text{ where } \mathbf{x}_N = A\mathbf{x}_{N-1} + B\mathbf{u}_{N-1}\}$$

2. Due to the terminal constraint

$$A\mathbf{x}_{N-1} + B\mathbf{u}_{N-1} = \mathbf{x}_N \in X_f$$

Persistent feasibility theorem

- Proof

3. Since X_f is a control invariant set, there exists a $\mathbf{u} \in U$ such that
$$\mathbf{x}^{+} = A\mathbf{x}_N + B\mathbf{u} \in X_f$$
4. The above is indeed the requirement to belong to set X_{N-1}
5. Thus, $A\mathbf{x}_{N-1} + B\mathbf{u}_{N-1} = \mathbf{x}_N \in X_{N-1}$
6. We have just proved that X_{N-1} is control invariant
7. Repeating this argument, one can recursively show that $X_{N-2}, X_{N-3}, \cdots, X_1$ are control invariant, and the persistent feasibility lemma then applies

Practical aspects of persistent feasibility

- The terminal set X_f is introduced *artificially* for the sole purpose of leading to a *sufficient condition* for persistent feasibility
- We want it to be large so that it does not compromise closed-loop performance
- Though it is simplest to choose $X_f = \{0\}$, this is generally undesirable
- We'll discuss better choices in the next lecture

Next time

- Stability of MPC
- Explicit MPC
- Practical considerations