Optimal Control of Ordinary Differential Equations
SOD 311

Laurent Pfeiffer
Inria and CentraleSupélec, Université Paris-Saclay

Ensta-Paris
Paris-Saclay University

Lecture 3:
Model Predictive Control

- An adaptive numerical method for solving long-horizon optimal control problems.
- Stability and quantitative analysis of the method for stabilization problems.

Bibliography

The following references are related to the lecture:

- Nonlinear Model Predictive Control: Theory and Algorithms, by Lars Grüne and Jürgen Pannek, 2011.
- Slides by Lars Grüne on “Deterministic Stabilizing and Economic Model Predictive Control”, from the 2022 MTNS Conference.

A continuous-time problem

Consider the optimal control problem:

$$\inf_{\substack{y \in W^{1,\infty}(0,\infty;\mathbb{R}^n) \\ u \in L^\infty(0,\infty;U)}} \int_0^\infty \ell(u(t), y(t))\,\mathrm{d}t, \quad \text{s.t.} \begin{cases} \dot{y}(t) = f(u(t), y(t)), \\ y(0) = x, \end{cases} \qquad (\mathcal{P}_\infty(x))$$

in which U, $\ell \colon U \times \mathbb{R}^n \to \mathbb{R}$, and $f \colon U \times \mathbb{R}^n \to \mathbb{R}^n$ are given and fixed. The initial condition x is seen as a parameter.

Given $y_0 \in \mathbb{R}^n$, we want to solve the problem for $x = y_0$. Issues:

- Difficult to solve such a problem on an **infinite** (or extremely large) horizon.
- Sometimes, finding an optimal pair $(\bar{y}, \bar{u})$ is insufficient because of **perturbations** or model uncertainties.

Dynamic programming

The **dynamic programming principle** plays a key role in the treatment of those difficulties.

- Let $(\bar{y}, \bar{u})$ be a solution to $\mathcal{P}_\infty(x)$ for $x = y_0$. Let $\theta > 0$.
- Define $\tilde{y} = \bar{y}_{|[\theta,\infty)}$, $\tilde{u} = \bar{u}_{|[\theta,\infty)}$. Then $(\tilde{y}, \tilde{u})$ is a solution to

$$\inf_{\substack{y \in W^{1,\infty}(\theta,\infty;\mathbb{R}^n) \\ u \in L^\infty(\theta,\infty;U)}} \int_\theta^\infty \ell(u(t), y(t))\,\mathrm{d}t, \quad \text{s.t.} \begin{cases} \dot{y}(t) = f(u(t), y(t)), \\ y(\theta) = \bar{y}(\theta). \end{cases}$$

- Define $\hat{y}$ and $\hat{u}$ by $\hat{y}(t) = \tilde{y}(t-\theta)$ and $\hat{y}(t) = \tilde{y}(t-\theta)$.

Then $(\hat{y}, \hat{u})$ is a **solution** to $\mathcal{P}_\infty(\bar{y}(\theta))$.

Towards MPC

Dynamic programming motivates the following algorithm, for solving $\mathcal{P}_\infty(x)$ with $x = y_0$ and taking into account possible perturbations.

Fix $\tau > 0$. At every time $t = k\tau$, $k \in \mathbb{N}$, a **measurement** of the state is realized.

- Find a solution $(\bar{y}, \bar{u})$ to $\mathcal{P}_\infty(x)$, with $x = y_0$.
- Implement $u(t) = \bar{u}(t)$, for a.e. $t \in (0, \tau)$.
- Measure the value of the state at time τ, call it y_1.
- Find a solution $(\bar{y}, \bar{u})$ to $\mathcal{P}_\infty(x)$, with $x = y_1$.
- Implement $u(t) = \bar{u}(t - \tau)$, for a.e. $t \in (\tau, 2\tau)$.
- Measure the value of the state at time 2τ, call it y_2.
- And so on.

Towards MPC

This approach already raises many questions!

- Benefit with respect to uncertainties in the model or perturbations ?
- Delay in the measurement and in the computation of a solution to $\mathcal{P}_\infty(x)$?

We do **not** address them directly in the lecture.

In practice, it is not reasonable to assume that $\mathcal{P}_\infty(x)$ is solvable instantaneously at each time $k\tau$. Instead, a **truncated** version of it is solved $\rightarrow$ source of error.

Objective: We want to analyze the impact of truncation in the ideal case without model uncertainty or perturbations.

Finite-horizon approximation

Model Predictive Control (MPC) is a numerical method which provides an approximate solution by solving **a sequence of short horizon problems**.

Parameters of the method:

- A sampling time $\tau > 0$.
- A prediction horizon $T = N\tau$, where $N \in \mathbb{N}^*$.

Consider the finite-horizon problem:

$$\inf_{\substack{y \in W^{1,\infty}(0,T;\mathbb{R}^n) \\ u \in L^\infty(0,T;U)}} \int_0^T \ell(u(t), y(t))\, dt, \quad \text{s.t.} \begin{cases} \dot{y}(t) = f(u(t), y(t)), \\ y(0) = x. \end{cases} \tag{$\mathcal{P}_T(x)$}$$

MPC method

The MPC method generates a pair $(y_{\text{MPC}}, u_{\text{MPC}})$ iteratively.
At iteration k, the pair is defined for $t \in [k\tau, (k+1)\tau]$.

MPC method

Set $y_{\text{MPC}}(0) = y_0$.
For $k = 0, ...,$ do

- Find a solution $(\bar{y}, \bar{u})$ of $\mathcal{P}_T(x)$, with $x = y_{\text{MPC}}(k\tau)$.
- Set $y_{\text{MPC}}(k\tau + t) = \bar{y}(t)$, $u_{\text{MPC}}(k\tau + t) = \bar{u}(t)$, $\forall t \in [0, \tau]$.

Remarks:

- Can be implemented in real time, assuming $\mathcal{P}_T(x)$ can be solved instantaneously.
- Intuition: replacing ∞ by T, we neglect the far future. Acceptable since we only utilize the solutions on $[0, \tau]$.

Discrete-time reformulation

Problem $\mathcal{P}$ can be put in the form:

$$\inf_{\substack{\tilde{y}\in\tilde{Y}^{\mathbb{N}} \\ \tilde{u}\in\tilde{U}^{\mathbb{N}}}} \sum_{k=0}^{N} \tilde{\ell}(\tilde{u}(k),\tilde{y}(k)), \quad \text{s.t.} \begin{cases} \forall k = 0, \dots \\ \tilde{y}(k+1) = \tilde{f}(\tilde{u}(k),\tilde{y}(k)) \\ \tilde{y}(0) = y_0. \end{cases}$$

Similarly, problem $\mathcal{P}_T(x)$ can be put in the form:

$$\inf_{\substack{\tilde{y}\in\tilde{Y}^{N+1} \\ \tilde{u}\in\tilde{U}^{N}}} \sum_{k=0}^{N-1} \ell(\tilde{u}(k),\tilde{y}(k)), \quad \text{s.t.} \begin{cases} \forall k = 0, \dots, N-1, \\ \tilde{y}(k+1) = \tilde{f}(\tilde{u}(k),\tilde{y}(k)), \\ \tilde{y}(0) = x. \end{cases}$$

Discrete-time reformulation

In this context, the MPC method reads:

Set $\tilde{y}_{MPC}(0) = y_0$.
For $k = 0, \ldots,$ do

- Find a solution $(\tilde{y}, \tilde{u})$ of $\mathcal{P}_T(x)$, with $x = \tilde{y}_{MPC}(k)$.
- Define $\tilde{y}_{\text{MPC}}(k+1) = \tilde{y}(1)$ and $\tilde{u}_{\text{MPC}}(k) = \tilde{u}(0)$.

We stick to this **discrete-time framework** and forget about the original continuous-time motivation.

Data problem

We introduce now the general setting utilized in the rest of the lecture.

Data of the problem:

- a set Y: the state space
- a set U: the control space
- a subset $\mathbb{Y} \subseteq Y$: the set of feasible states
- a multivalued map $\mathbb{U} \colon y \in \mathbb{Y} \mapsto \mathbb{U}(y) \subseteq U$
- a dynamics $f \colon U \times Y \to Y$
- a running cost $\ell \colon U \times Y \to \mathbb{R}$.

For the moment, we only require that $\ell(u, y) \geq 0$, for any $(u, y) \in U \times Y$.

Feasible controls

- Given $n \in \mathbb{N} \cup \{\infty\}$, we call **control sequence** (of length n) any $u = (u(k))_{k=0,\dots,n-1} \in U^n$.
- Given an initial condition $x \in \mathbb{Y}$, we call **associated trajectory** $y[u,x] = (y[u,x](k))_{k=0,\dots,n} \in Y^{n+1}$ the solution to

$$\begin{aligned} y[u,x](k+1) &= f(y[u,x](k), u(k)), \quad \forall k = 0, \dots n-1 \\ y[u,x](0) &= x. \end{aligned}$$

- Given $x \in \mathbb{Y}$, we define the set of **feasible infinite control sequence**

$$\mathbb{U}^\infty(x) = \left\{ u \in U^\infty \;\middle|\; \begin{array}{ll} y[u,x](k) \in \mathbb{Y}, & \forall k \in \mathbb{N}, \\ u(k) \in \mathbb{U}(y[u,x](k)), & \forall k \in \mathbb{N} \end{array} \right\}.$$

Optimal control problem

- Given an initial condition x and a control sequence $(u(k))_{k=0,1,\ldots} \in U^{\infty}$, we consider the **cost**

$$J_{\infty}(u,x) = \sum_{k=0}^{\infty} \ell(u(k), y[u,x](k)) \in \mathbb{R} \cup \{\infty\}.$$

  Note that J_{∞} can be infinite.

- The **optimal control problem** of interest is:

$$\inf_{u \in \mathbb{U}^{\infty}(x)} J_{\infty}(u,x). \qquad (\mathcal{P}_{\infty}(x))$$

Stabilization problems

We restrict ourselves to **stabilization problems**: we assume the existence of a pair (y^*, u^*) such that: $y^* \in \mathbb{Y}$, $u^* \in \mathbb{U}(y^*)$, and

$$y^* = f(u^*, y^*)$$
$$\ell(u, y) = 0 \Longleftrightarrow (u, y) = (u^*, y^*), \quad \forall (u, y) \in U \times Y.$$

In this framework, the cost function J_∞ steers the state to y^*.
If $y_0 = y^*$, then the solution to $\mathcal{P}$ is the constant sequence $\bar{u}$ equal to u^*, since $\bar{u}$ is feasible and $J_\infty(\bar{u}, y^*) = 0$.

MPC method

The **MPC method** requires:

- a time horizon $N \in \mathbb{N}$
- a terminal set $\mathbb{Y}_0 \subseteq \mathbb{Y}$
- a terminal cost $F \colon \mathbb{Y}_0 \to \mathbb{R}_+$.

Given $x \in \mathbb{Y}$, we denote

$$\mathbb{U}^N(x) = \left\{ u \in U^N \;\middle|\; \begin{array}{ll} y[u,x](k) \in \mathbb{Y}, & \forall k = 0, \dots, N-1 \\ u(k) \in \mathbb{U}(y[u,x](k)), & \forall k = 0, \dots, N-1 \\ y[u,x](N) \in \mathbb{Y}_0 & \end{array} \right\}.$$

For any $N \in \mathbb{N} \cup \{\infty\}$, we consider the set of **feasible initial conditions** (for $\mathcal{P}_N(x)$): $\mathbb{Y}_N = \left\{ x \in \mathbb{Y} \,\middle|\, \mathbb{U}^N(x) \neq \emptyset \right\}$.

MPC method

Moreover, given $u \in \mathbb{U}^N(x)$, we define

$$J_N(u,x) = \left(\sum_{k=0}^{N-1} \ell(u(k), y[u,x](k)) \right) + F(y[u,x](N)).$$

The finite-horizon problem utilized in the MPC method is:

$$\inf_{u \in \mathbb{U}^N(x)} J_N(u,x). \tag{$\mathcal{P}_N(x)$}$$

Introduction

The MPC method essentially consists in computing a **feedback function** $\mu_N \colon \mathbb{Y} \to U$ in the following fashion: for any $x \in \mathbb{Y}_N$,

- Find a solution $\bar{u}$ to $\mathcal{P}_N(x)$.
- Set $\mu_N(x) = \bar{u}(0)$.

Then the pair $(u_{\text{MPC}}, y_{\text{MPC}})$ is recursively defined by:

$$\begin{cases} y_{\text{MPC}}(0) = y_0, & \\ u_{\text{MPC}}(k) = \mu_N(y_{MPC}(k)), & \forall k = 0, 1, \dots \\ y_{\text{MPC}}(k+1) = f(u_{\text{MPC}}(k), y_{\text{MPC}}(k)), & \forall k = 0, 1, \dots \end{cases}$$

Equivalently, $y_{\text{MPC}}(k+1) = f_N(y_{\text{MPC}}(k))$, where f_N is defined by

$$f_N(x) = f(\mu_N(x), x).$$

Viability

Lemma 1

Assume that for any $x \in \mathbb{Y}_0$*, there exists* $u \in \mathbb{U}(x)$ *such that* $f(u, x) \in \mathbb{Y}_0$*. Then*

$$\mathbb{Y}_N \subseteq \mathbb{Y}_{N+1} \subseteq \dots \subseteq \mathbb{Y}_\infty.$$

Moreover, for any $x \in \mathbb{Y}_N$*, it holds that*

$$f_N(x) = f(\mu_N(x), x) \in \mathbb{Y}_N.$$

Remark: under the assumption of the lemma, if $\mathcal{P}_N(x)$ is feasible in the first step of the method, then it is for all other steps.

Objectives

The issues related to the existence of a solution to all optimization problems is eluded here. It can be addressed with standard arguments.

We will investigate:

- the **qualitative** behavior of MPC
  $\rightarrow$ convergence of y_{MPC} to y^*?
- the **quantitative** behavior of MPC
  $\rightarrow$ bound of $J_\infty(u_{\text{MPC}}, x)$?

Stability analysis

- We forget about the MPC method introduced before and focus on **discrete-time dynamical systems** of the form

$$y(0) = x, \quad y(k+1) = g(y(k)), \quad \forall k = 0, 1, \dots$$

for a given initial condition $x \in Y$ and $g\colon Y \to Y$. We denote the solution by $(y[x](k))_{k=0,1,\dots}$.

- We fix an **equilibrium point** y^* of the dynamics g, that is to say, we assume that $y^* = g(y^*)$.
- We investigate here the **convergence** of $y[x](k)$ to y^*. To this purpose we assume that $\mathbb{Y}$ is metric space and we denote the **distance** of an arbitrary point $y \in Y$ to y^* by $|y|$.

Comparison functions

Definition 2

We define here several classes of functions:

- $\mathcal{K} = \left\{ \alpha \colon [0,\infty) \to [0,\infty) \;\middle|\; \begin{array}{l} \alpha \text{ is continuous} \\ \alpha \text{ is strictly increasing} \\ \alpha(0) = 0 \end{array} \right\}$
- $\mathcal{K}_\infty = \left\{ \alpha \in \mathcal{K} \;\middle|\; \alpha(r) \underset{r\to\infty}{\longrightarrow} \infty \right\}$
- $\mathcal{L} = \left\{ \delta \colon \mathbb{N} \to [0,\infty) \;\middle|\; \begin{array}{l} \delta \text{ is strictly decreasing} \\ \delta(t) \underset{t\to\infty}{\longrightarrow} 0 \end{array} \right\}$
- $\mathcal{KL} = \left\{ \beta \colon [0,\infty) \times \mathbb{N} \to \mathbb{R} \;\middle|\; \begin{array}{l} \beta(\cdot,t) \in \mathcal{K},\ \forall t \in \mathbb{N} \\ \beta(r,\cdot) \in \mathcal{L},\ \forall r \in [0,\infty) \end{array} \right\}.$

Remark: the t variable in the definition of $\mathcal{L}$ and $\mathcal{KL}$ is usually supposed to lie in $\mathbb{R}$ in the literature.

Asymptotic stability

Definition 3

We say that y^* is **asymptotically stable** if there exists $\beta \in \mathcal{KL}$ such that for any $x \in \mathbb{Y}$,

$$|y[x](k)| \leq \beta(|x|, k), \quad \forall k \in \mathbb{N}.$$

Remark: this implies that $|y[x](k)| \to 0$ as $k \to \infty$ (not an equivalence).

Lyapunov functions

Definition 4

We call $V\colon Y \to [0,\infty)$ a **Lyapunov function** (associated with g) if

- there exists α_1 and $\alpha_2 \in \mathcal{K}_\infty$ such that

$$\alpha_1(|y|) \leq V(y) \leq \alpha_2(|y|), \quad \forall y \in Y$$

- there exists $\alpha_3 \in \mathcal{K}$ such that

$$V(g(y)) \leq V(y) - \alpha_3(|y|), \quad \forall y \in Y.$$

Remark: if V is a Lyapunov function, then $V(y) \geq 0$, for any $y \in Y$. Moreover, $V(y) = 0 \iff y = y^*$ and y^* is the unique equilibrium.

Stability under Lyapunov

Theorem 5

Assume the existence of a Lyapunov function. Then y^* is asymptotically stable.

Our **objective** for the rest of the lecture: constructing a Lyapunov function for the dynamic system

$$y(k+1) = f_N(y(k)) = f(\mu_N(y(k)), y(k))$$

corresponding to the MPC method.

Value function

For $N \in \mathbb{N} \cup \{\infty\}$, the **value function** $V_N \colon Y \to \mathbb{R} \cup \{\infty\}$ is defined by:

$$V_N(x) = \inf_{u \in \mathbb{U}^N(x)} J_N(u, x).$$

We set $V_N(x) = +\infty$ if $x \notin \mathbb{Y}_N$ (i.e. $\mathbb{U}^N(x)$ is empty).

Theorem 6

For $N \geq 1$ (possibly $N = \infty$), we have

$$V_N(x) = \inf_{u \in \mathbb{U}(x)} \ell(u, x) + V_{N-1}(f(u, x)), \quad \forall x \in \mathbb{Y}. \qquad \text{(DP)}$$

For $N = 0$, it holds that

$$V_0(x) = \begin{cases} F(x) & \text{if } x \in \mathbb{Y}_0 \\ \infty & \text{otherwise.} \end{cases}$$

Dynamic programming

Corollary 7

Let $N \geq 1$ (possibly $N = \infty$). Let $x \in \mathbb{Y}_N$.

- *Let $\bar{u}$ be a solution to $\mathcal{P}_N(x)$. Then $\bar{u}(0)$ is a solution to (DP). Define $u' \in U^\infty$ by $u'(k) = \bar{u}(k+1)$. Then u' is a solution to $\mathcal{P}_{N-1}(f(u,x))$.*

  In particular, $\mu_N(x)$ is a solution to (DP).

- *Let u be a solution to (DP). Let u' be a solution to $\mathcal{P}_{N-1}(f(u,x))$. Define $\bar{u} \in U^\infty$ by*

  $$\bar{u}(0) = u, \quad \bar{u}(k+1) = u'(k), \quad \forall k \in \mathbb{N}.$$

  Then $\bar{u}$ is a solution to $\mathcal{P}_N(x)$.

From DP to Lyapunov

A key consequence of the dynamic programming principle is the following:

$$V_{N-1}(f_N(x)) = V_N(x) - \ell(\mu_N(x), x), \quad \forall x \in \mathbb{Y}_N.$$

Observation: this relation is close to the decay condition satisfied by Lyapunov functions.

Can we find **structural assumptions** allowing to use V_N as a Lyapunov function?

An abstract result

Theorem 8

1. *Assume that there exists* $\alpha \in (0,1]$ *such that*

$$V_N(f_N(x)) \le V_N(x) - \alpha \ell(\mu_N(x), x), \quad \forall x \in \mathbb{Y}_N. \qquad \text{(ADP)}$$

Then the control generated by the MPC method satisfies

$$J_\infty(u^N_{MPC}(x), x) \le V_N(x)/\alpha.$$

2. *Assume moreover that there exist* α_2 *and* α_3 *in* $\mathcal{K}_\infty$ *such that*

$$V_N(x) \le \alpha_2(|x|) \quad \textit{and} \quad \inf_{u \in \mathbb{U}(x)} \ell(u, x) \ge \alpha_3(|x|), \quad \forall x \in \mathbb{Y}_N.$$

Then y^* *is* asymptotically stable, *for the dynamical system* $y(k+1) = f_N(y(k))$.

Methodology

The previous result relies on **non-explicit assumptions**, which remain to be established (under more explicit assumptions...).

We focus on the verification of the **ADP inequality** (for "approximate dynamic programming").

We distinguish two cases:

- MPC **without** terminal cost and terminal constraints: $\mathbb{Y}_0 = \mathbb{Y}$ and $F = 0$.
- MPC **with** terminal cost and constraints.

Each of these two cases requires specific assumptions.

Assumption

Assumption 1

There exists a map $\kappa\colon \mathbb{Y}_0 \to U$ satisfying, for any $x \in \mathbb{Y}_0$,

- $\kappa(x) \in \mathbb{U}(x)$ and $f(\kappa(x), x) \in \mathbb{Y}_0$
- $F(f(\kappa(x), x)) \leq F(x) - \ell(\kappa(x), x)$.

Remarks.

- This assumption is in particular satisfied for $\mathbb{Y}_0 = \{y^*\}$, with $\kappa(y^*) = u^*$ and $F = 0$. But the resolution of $\mathcal{P}_N(x)$ is difficult.
- In general, good candidates for F are approximations of V_∞ obtained through linearization techniques.

Result

Lemma 9

Under Assumption 1, we have

$$V_0(x) \geq \ldots \geq V_{N-1}(x) \geq V_N(x) \geq \ldots V_\infty(x).$$

In particular, $\mathbb{Y}_0 \subseteq \ldots \subseteq \mathbb{Y}_{N-1} \subseteq \mathbb{Y}_N \subseteq \ldots \mathbb{Y}_\infty$.

Corollary 10

Under Assumption 1, V_N *satisfies (ADP) with* $\alpha = 1$. *Therefore*

$$J_\infty(u^N_{MPC}(x), x) \leq V_N(x), \quad \forall x \in \mathbb{Y}_N.$$

Discussion

- The analysis is relatively easy and natural.
- Computation of suitable $\mathbb{Y}_0$ and F is more complex.
- Resolution of $\mathcal{P}_N(x)$ possibly difficult.
- Given a feasible initial condition $x \in \mathbb{Y}_\infty$, one possibly needs to have N quite large so that $x \in \mathbb{Y}_N$.
- No general bound on $V_N(x)$ (in comparison with $V_\infty(x)$, yet convergence (w.r.t. N can be achieved).

Result

For any $x \in \mathbb{Y}$, we denote $\ell^*(x) = \inf_{u \in \mathbb{U}(x)} \ \ell(u, x)$.

Assumption 2

- There is no terminal condition: $\mathbb{Y}_0 = \mathbb{Y}$ and $F = 0$.
- There exists $\gamma > 0$ such that

$$V_N(x) \le \gamma \ell^*(x), \quad \forall x \in \mathbb{Y}, \ \forall N \in \mathbb{N}.$$

- There exist α_3 and $\alpha_4 \in \mathcal{K}_\infty$ such that

$$\alpha_3(|x|) \le \ell^*(x) \le \alpha_4(|x|).$$

Lemma 11

Under Assumption 2, the ADP ineq. is satisfied for any $N \geq 1$ with

$$\alpha_N = 1 - \gamma(\gamma - 1)/(N - 1).$$

Discussion

- Resolution of $\mathcal{P}_N(x)$ easier without terminal condition (used in practice).
- The method also requires N to be sufficiently large (so that $\alpha_N > 0$).
- Refinement are possible (i.e. sharper estimates of α_N on coefficients γ which depend on N).
- Proof of existence of γ doable on a case-by-case basis.