Dynamic Programming and Model Predictive Control

Anders Rantzer

Automatic Control LTH
Lund University
Sweden

Course Outline

Tuesday 08:45
- Optimal control
- Dynamic programming
- Bellman's equation
- Value iteration
- Approximate DynP

Wednesday 08:45
- Tuesday summary
- MPC with terminal cost
- MPC without terminal cost
- Adaptive horizon
- Receding horizon estimation

Tuesday 09:45
- Model Predicitve Control
- Stability/feasibility
- Terminal constraints
- Introduction to exercise

Wednesday 09:45
- Dual Decomposition
- Distributed MPC
- Gradient methods
- Large-scale systems

Wednesday afternoon: Computer exercise

The MPC Control Law

Define the MPC control law μ_N through the minimization

$$V_N(x_0) = \inf_{u,x} \sum_{t=0}^{N-1} \ell(x(t), u(t))$$

where infimum is taken over $x(t) \in X$, $u(t) \in U$ satisfying $x(t+1) = f(x(t), u(t))$ and $x(0) = x_0$.

Then

$$V_1 \leq V_2 \leq \ldots \leq V_N \leq \ldots \leq V_\infty \leq V_\infty^{\mu_N}$$

Example 2 — Things can go bad

$$\inf_{u,x} \sum_{t=0}^{N-1} (|x(t)|^2 + u(t)^2)$$

$$x(t+1) = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x(t) + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u(t)$$

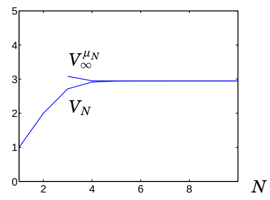

Marginally unstable for $N \leq 2$.

$$\inf_{u,x} \sum_{t=0}^{N-1} (|x(t)|^2 + u(t)^2)$$

$$x(t+1) = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} x(t) + \begin{bmatrix} -2 \\ 1 \end{bmatrix} u(t)$$

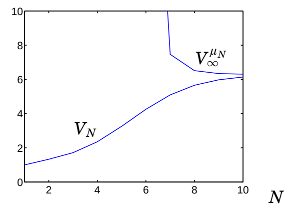

Exponentially unstable for $3 \leq N \leq 5$!

MPC with Equilibrium Terminal Constraint

Define the MPC control law μ_N through the minimization

$$\overline{V}_N(x_0) = \inf_{u,x} \sum_{t=0}^{N-1} \ell(x(t), u(t))$$

subject to $x(t) \in X$, $u(t) \in U$, $x(t+1) = f(x(t), u(t))$, $x(0) = x_0$ and the *terminal constraint* $x(N) = 0$.

Then $\overline{V}_N$ is a Lyapunov function that proves stability! Moreover $V_\infty \leq V_\infty^{\mu_N} \leq \overline{V}_N \leq \ldots \leq \overline{V}_2 \leq \overline{V}_1$.

Can this idea be generalized?

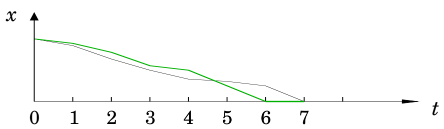

Lecture 3

- Tuesday summary
- **MPC with terminal cost**
- MPC without terminal cost
- Adaptive prediction horizon
- Reference tracking and estimation

Upper Bound on the Optimal Cost

Recall from Lecture 1 that if $W(x) \geq 0$ and a control law $\mu : X \to U$ is given such that

$$W(f(x, \mu(x))) + \ell(x, \mu(x)) \leq W(x)$$

then W is a Lyapunov function for the closed loop system and the infinite horizon cost for the control law is bounded by W:

$$V_\infty^\mu(x) \leq W(x)$$

Such a "control Lyapunov function" can be used to get performance guarantees in MPC.

MPC with Terminal cost

Assume that

$$W(f(x, \mu(x))) + \ell(x, \mu(x)) \leq W(x) \qquad \text{for all } x$$

Define the MPC control law μ_N using the minimization

$$\overline{V}_N(x_0) = \inf_{u,x} \left[\sum_{t=0}^{N-1} \ell(x(t), u(t)) \underbrace{+W(x(N))}_{\textit{terminal cost}} \right]$$

with $x(t) \in X$, $u(t) \in U$, $x(t+1) = f(x(t), u(t))$, $x(0) = x_0$.
Then μ_N is stabilizing and $V_\infty \leq V_\infty^{\mu_N} \leq \overline{V}_N \leq \ldots \leq \overline{V}_2 \leq \overline{V}_1$.

Proof

Each feasible trajectory $u(0), u(1), \dots, u(N-2)$ can be prolonged by setting $u(N-1) = \mu(x(N-1))$ to get

$$\begin{aligned} J_{N-1}(x_0, u) &= \sum_{t=0}^{N-2} \ell(x(t), u(t)) + W(x(N-1)) \\ &\geq \sum_{t=0}^{N-1} \ell(x(t), u(t)) + W(x(N)) = J_N(x_0, u) \end{aligned}$$

Minimization gives

$$\overline{V}_{N-1}(x) \geq \overline{V}_N(x)$$

and stability follows as for the equilibrium constraint.

Next: If control Lyapunov function W is only valid near $x = 0$...

Terminal cost and terminal constraint

Assume existence of a function $W(x) \geq 0$, a control law $u = \mu(x)$ and a number $\epsilon > 0$ such that $W(x) \leq \epsilon \Rightarrow W(f(x, \mu(x)) + \ell(x, \mu(x)) \leq W(x)$.

Define the MPC control law μ_N using the minimization

$$\overline{V}_N(x_0) = \inf_{u,x} \left[\sum_{t=0}^{N-1} \ell(x(t), u(t)) \underbrace{+W(x(N))}_{\textit{terminal cost}} \right]$$

subject to $x(t) \in X$, $u(t) \in U$, $x(t+1) = f(x(t), u(t))$, $x(0) = x_0$ and the *terminal constraint* $W(x) \leq \epsilon$.

Then μ_N is stabilizing and $V_\infty \leq V_\infty^{\mu_N} \leq \overline{V}_N \leq \dots \leq \overline{V}_2 \leq \overline{V}_1$.

Dynamic Programming versus MPC

- Dynamic Programming (Explicit MPC)
  - Corresponds to MPC with $N = 2$ and accurate terminal cost
  - Heavy off-line computations and memory requirements
  - Extremely fast on-line
- Model Predictive Control
  - No off-line computations
  - Heavy on-line computations
  - Wide range of industrial applications exist

Lecture 3

- Tuesday summary
- MPC with terminal cost
- **MPC without terminal cost**
- Adaptive prediction horizon
- Reference tracking and estimation

When is MPC Stabilizing Without Terminal Cost?

Consider

$$V_N(x_0) = \inf_{u,x} \sum_{t=0}^{N-1} \ell(x(t), u(t))$$

where infimum is taken over $x(t) \in X$, $u(t) \in U$ satisfying $x(t+1) = f(x(t), u(t))$ and $x(0) = x_0$. The MPC control law

$$\mu_N(x) := \arg\min_u \{V_{N-1}(f(x, u)) + \ell(x, u)\}$$

gives

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

so V_N is a Lyapunov function provided that the right hand side is bigger than $V_N(f(x, \mu_N(x)))$.

Such comparisons (value iteration convergence) were done in the previous lecture.

Exponential stabilizability

Suppose there exist numbers $C > 0$ and $\sigma \in (0, 1)$ such that for every $x_0 \in X$ there exists a sequence $u(0), u(1), \dots \in U$ with

$$\ell(x(t), u(t)) \leq C\sigma^t \ell^*(x_0) \qquad \text{for all } t \geq 0$$

where $\ell^*(x_0) = \min_v \ell(x_0, v)$. This can be viewed as a condition of exponential stabilizability.

Then the MPC control law $\mu_N(x)$ is stabilizing provided that

$$N \geq 2\gamma \ln \gamma$$

where $\gamma = \frac{C}{1-\sigma}$.

[Grüne and Rantzer, TAC 53:9, 2009, Proposition 4.7]

Lecture 3

- Tuesday summary
- MPC with terminal cost
- MPC without terminal cost
- **Adaptive prediction horizon**
- Receding horizon estimation

Choice of Prediciton Horizon

- Should correspond to time constant of closed loop
- Fundamental bounds on achievable time constants
  - Unstable zeros
  - Time-delays
  - Input saturations
- Fast sampling, but long horizon:
  Limit optimization to inputs that change more seldom.

Bounding Performance Versus Horizon

Computing V_N without terminal costs and $\overline{V}_N$ with terminal costs/constraints as stated before gives

$$V_1 \leq V_2 \leq \ldots \leq V_N \leq \ldots \leq V_\infty \leq V_\infty^{\mu_N} \leq \overline{V}_N \leq \ldots \leq \overline{V}_2 \leq \overline{V}_1$$

In particular, the deviation between the MPC performance $V_\infty^{\mu_N}$ and the optimal cost V_∞ is bounded above as

$$V_\infty^{\mu_N}(x) - V_\infty(x) \leq \overline{V}_N(x) - V_N(x)$$

where the right hand side is computed by solving two optimization problems with horizon N.

Increase N until the accuracy is sufficient!

MPC with Adaptive Horizon

Recall Bellman's equation

$$\overline{V}_N(x) = \ell(x, \mu_N(x)) + \overline{V}_{N-1}(f(x, \mu_N(x)))$$

With time-varying horizon

$$\overline{V}_{N(t)}(x(t)) \geq \ell(x(t), u(t)) + \overline{V}_{N(t+1)}(x(t+1))$$

as long as $N(t+1)$ is at least as big as $N(t)$.

Summing both over t gives

$$\overline{V}_{N(0)}(x_0) \geq \sum_{t=0}^{\infty} \ell(x(t), u(t))$$

Lecture 3

- Tuesday summary
- MPC with terminal cost
- MPC without terminal cost
- Adaptive prediction horizon
- **Reference tracking and estimation**

Reference Tracking and Anti-windup

Zero tracking error for a constant reference requires *integral action*. This can be achieved in many ways, for example

1. Penalize *input changes* $u(t) - u(t-1)$ rather than $u(t)$.
2. Integrate the output error and penalize the integral.
3. Constant load disturbances on control inputs are assumed and estimated. The estimates are used for feedback.

Integral action *must be* combined with anti-windup:

1. Penalize $u(t) - \bar{u}(t-1)$, where $\bar{u}$ is "true" input, not "intended".
2. Stop integration of output error when the input saturates.

Do you see advantages/disadvantages with the two alternatives for multi-input-multi-output systems?

Receding Horizon Estimation

For dynamics with process noise w and measurement error v

$$x(t+1) = f(x(t)) + w(t)$$
$$y(t) = h(x(t)) + v(t)$$

an "optimal" state estimate at time t can be defined as the solution to

$$\hat{x} = \arg\min_x \sum_{\tau=-\infty}^{t-1} \ell(x(\tau+1) - f(x(\tau)), y(\tau) - h(x(\tau)))$$

(For linear dynamics and quadratic ℓ, a Kalman filter is optimal.)

A receding horizon estimate $\hat{x}^N$ with horizon N is defined by

$$\hat{x}^N = \arg\min_x \sum_{\tau=t-N}^{t-1} \ell(x(\tau+1) - f(x(\tau)), y(\tau) - h(x(\tau)))$$

Compare a time-varying $y(t)$ to a time-varying reference signal!

Lecture 3

- Tuesday summary
- MPC with terminal cost
- MPC without terminal cost
- Adaptive prediction horizon
- Reference tracking and estimation
- **Introduction to exercise**

Control of a Quadruple Tank

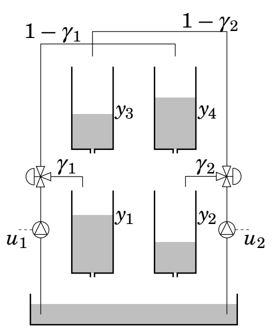

The transfer matrix from (u_1, u_2) to (y_1, y_2)

$$\begin{pmatrix} \dfrac{\gamma_1 c_1}{1+sT_1} & \dfrac{(1-\gamma_2)c_1}{(1+sT_1)(1+sT_3)} \\ \dfrac{(1-\gamma_1)c_2}{(1+sT_2)(1+sT_4)} & \dfrac{\gamma_2 c_2}{1+sT_2} \end{pmatrix}$$

has unstable zero if and only if

$$0 \leq \gamma_1 + \gamma_2 < 1$$

Is there a step response with wrong direction?

Lab material

can be found on `www.control.lth.se/disc-summer-school/`

Lab Exercise

```
% limit on delta u
du_max = [inf inf];
du_min = [-inf -inf];

% limit absolute value of u
u_max = [10-v1 10-v2];
u_min = [-v1 -v2];

% limit controlled outputs
z_max = kc*[15-h1 15-h2 15-h3 15-h4];
z_min = -Inf*[1 1 1 1];
```

Course Outline

Tuesday 08:45

- Optimal control
- Dynamic programming
- Bellman's equation
- Value iteration
- Approximate DynP

Tuesday 09:45

- Model Predicitve Control
- Stability/feasibility
- Terminal constraints
- Introduction to exercise

Wednesday 08:45

- Tuesday summary
- MPC with terminal cost
- MPC without terminal cost
- Adaptive horizon
- Receding horizon estimation

Wednesday 09:45

- Dual Decomposition
- Distributed MPC
- Gradient methods
- Large-scale systems

Wednesday afternoon: Computer exercise