Model Predictive Control

Lecture: Introduction

Colin Jones

Laboratoire d'Automatique, EPFL

Outline

1. Introduction to MPC
    - Concept
    - The math
    - Examples

2. Administration

Optimization in the loop

Classical control loop:

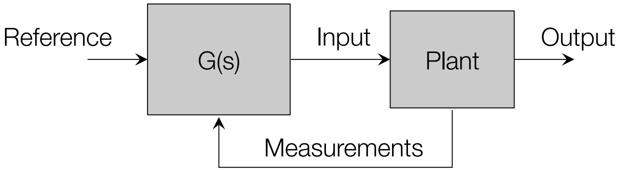

The classical controller is replaced by an optimization algorithm:

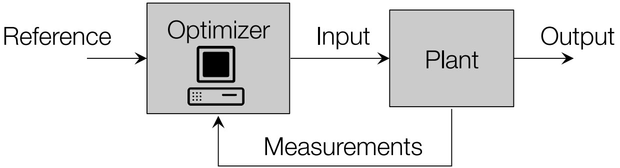

The optimization uses predictions based on a model of the plant.

Optimization-based control: Conceptual Example

Constraints:

- Stay on road
- Don't skid
- Limited acceleration

Intuitive approach:

- Look forward and plan path based on
  - Road conditions
  - Upcoming corners
  - Abilities of car
  - etc...

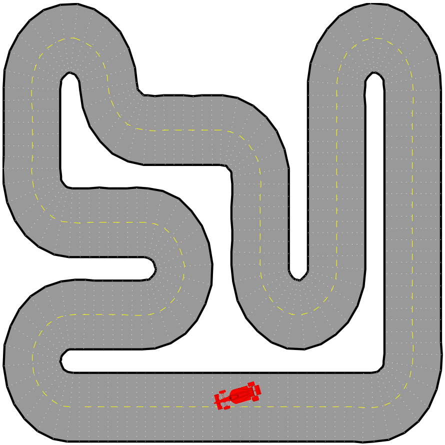

Optimization-based control: Conceptual Example

minimize	(circuit time)
while	avoid other cars
	stay on road
	...

- Solve optimization problem to compute minimum-time path

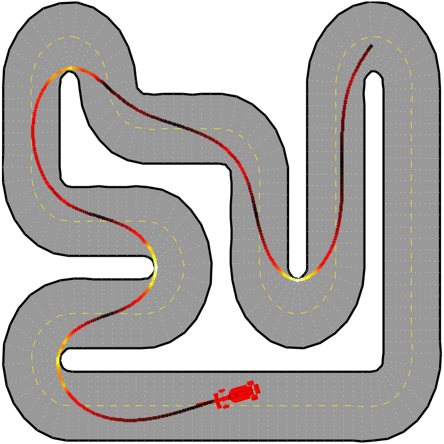

Optimization-based control: Conceptual Example

minimize	(circuit time)
while	avoid other cars
	stay on road
	...

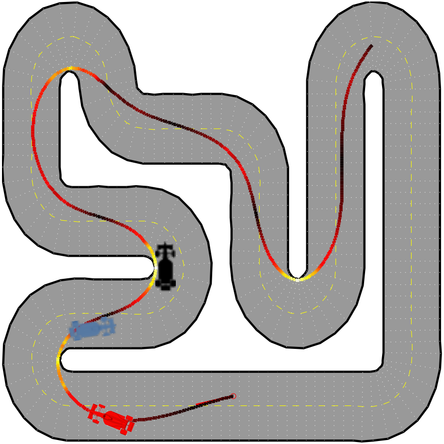

- Solve optimization problem to compute minimum-time path
- What happens if something unexpected happens?
  - We didn't see a car around the corner!
  - Must introduce **feedback**

Optimization-based control: Conceptual Example

minimize	(circuit time)
while	avoid other cars
	stay on road
	...

- Solve optimization problem to compute minimum-time path
- Obtain planned control actions
- Apply first control move
- Repeat the planning procedure

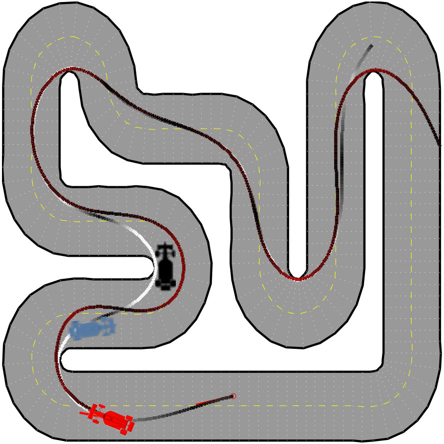

Receding horizon control

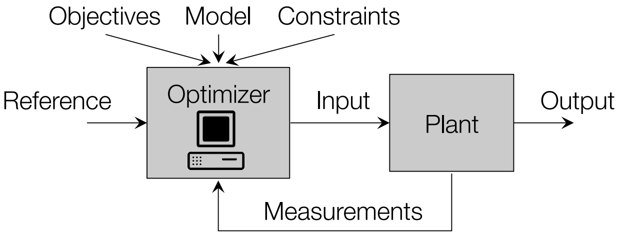

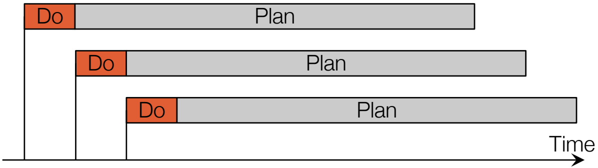

Receding horizon strategy introduces feedback.

Constraints in Control

All physical systems have **constraints**.

- Physical constraints, e.g. actuator limits
- Performance constraints, e.g. overshoot
- Safety constraints, e.g. temperature/pressure limits

Optimal operating points are often near constraints.

Classical control methods:

- No knowledge of constraints
- Set point sufficiently far from constraints
- Suboptimal plant operation

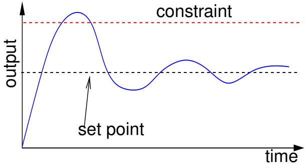

Predictive control:

- Constraints included in the design
- Set point optimal
- Efficient plant operation

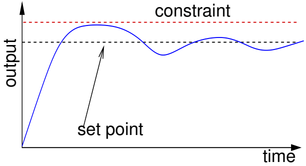

Reasons to Use Predictive Control

When to use predictive control?

1. Constraints drive performance
2. Strongly nonlinear system dynamics
3. Complex objectives
4. Some future knowledge

System matches *any* of these conditions: MPC could be the right solution

When not to use predictive control?

1. Whenever something simpler would work!

MPC: Mathematical formulation

$$
\begin{aligned}
u^\star(x) := \text{argmin} \quad & x_N^T Q_f x_N + \sum_{i=0}^{N-1} x_i^T Q x_i + u_i^T R u_i \\
\text{s.t.} \quad & x_0 = x && \text{measurement} \\
& x_{i+1} = A x_i + B u_i && \text{system model} \\
& C x_i + D u_i \leq b && \text{constraints} \\
& R \succ 0, Q \succ 0 && \text{performance weights}
\end{aligned}
$$

Problem is defined by

- **Objective** that is minimized,
  e.g., distance from origin, sum of squared/absolute errors, economic,...
- Internal **system model** to predict system behavior
  e.g., linear, nonlinear, single-/multi-variable, ...
- **Constraints** that have to be satisfied
  e.g., on inputs, outputs, states, linear, quadratic,...

MPC: Mathematical formulation

$$
\begin{aligned}
u^\star(x) := \text{argmin} \quad & x_N^T Q_f x_N + \sum_{i=0}^{N-1} x_i^T Q x_i + u_i^T R u_i \\
\text{s.t.} \quad & x_0 = x && \text{measurement} \\
& x_{i+1} = Ax_i + Bu_i && \text{system model} \\
& Cx_i + Du_i \leq b && \text{constraints} \\
& R \succ 0, Q \succ 0 && \text{performance weights}
\end{aligned}
$$

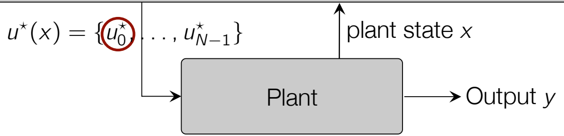

At each sample time:

- Measure /estimate current state
- Find the optimal input sequence for the entire planning window N
- Implement only the **first** control action

MPC: Applications

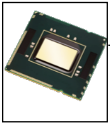

Computer control — ns

μs — Power systems

Traction control — ms

Seconds — Buildings

Refineries — Minutes

Hours — Nurse rostering

Train scheduling — Days

Weeks — Production planning

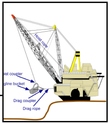

Example: Ball on Plate

- Movable plate (0.66m x 0.66m)
- Can be revolved around two axis $[+17°; -17°]$ by two DC motors
- Angle is measured by potentiometers
- Position of the ball is measured by a camera
- Model: Linearized dynamics, 4 states, 1 input per axis
- Input constraints: Voltage of motors
- State constraints: Boundary of the plate, angle of the plate

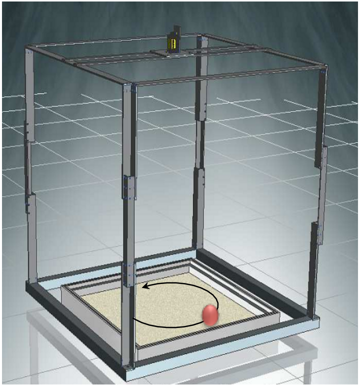

[Master thesis R. Waldvogel, ETH, 2011]

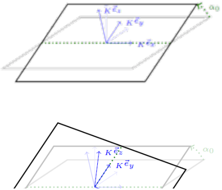

Example: Ball on Plate

Controller comparison:
LQR vs. MPC in the presence of input constraints

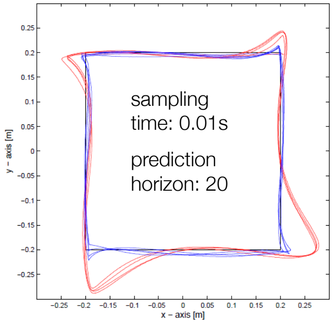

(a) LQR (red) vs MPC Controller (blue)

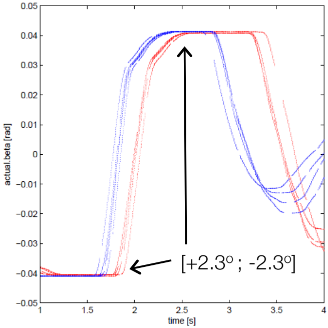

(b) Input β for the upper left corner

MPC introduces preview by predicting the state over a finite horizon

[Master thesis by R. Waldvogel, ETH, 2011]

Path Following

MPC Control of a crane along a **known** path:

crane

[Jan Swevers, KU Leuven]

Example: Autonomous Quadrocopter flight

Quadrocopters:

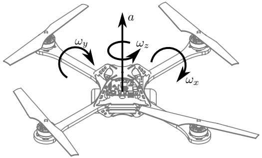

- Highly agile due to fast rotational dynamics
- High thrust-to-weight ratio allows for large translational accelerations
- Motion control by altering rotation rate and/or pitch of the rotors
- High thrust motors enable high performance control

Control Problem:

- Nonlinear system in 6D (position, attitude)
- Constraints: limited thrust, rates,...
- Task: Hovering, trajectory tracking
- Challenges: Fast unstable dynamics

Example: Autonomous Quadrocopter flight

[IDSC, ETH Zurich]

dandrea mpc

Example: Energy Efficient Building Control

- Buildings account for $\approx 40\%$ of global energy use
- Most of the energy is consumed during the use of the buildings
- Building sector has large potential for cost-effective reduction of CO_2 emissions
- Most investments in buidlings are expected to pay back through reduced energy bills

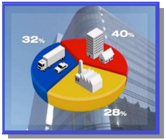

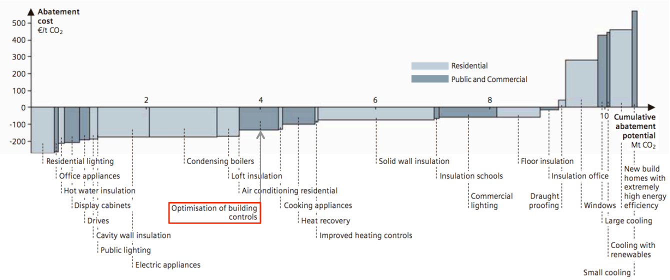

Greenhouse gas abatement cost curve for London buildings (2025, decision maker perspective)

Source: Watson, J. (ed.) (2008): Sustainable Urban Infrastructure, London Edition – a view to 2025. Siemens AG, Corporate Communications (CC) Munich, 71pp.

Example: Energy Efficient Building Control

Application "Integrated Room Automation":
Integrated control of heating, cooling, ventilation, electrical lighting, blinds,... of a single room/zone

Control Task: Use minimum amount of energy (or money) to keep room temperature, illuminance level and CO_2 concentration in prescribed comfort ranges

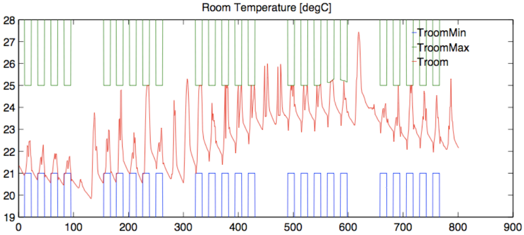

Example: One-Zone Building

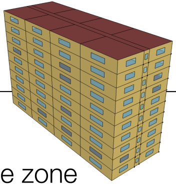

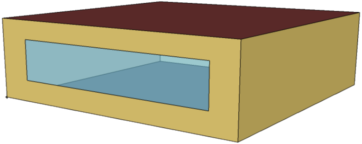

- EnergyPlus model of a single zone
- Automatic extraction and linearization: openBuild tool
- Electric heating
- Weather: Jan, 2007 in San Francisco

$$x^+ = Ax + Bu + E_{\text{rad}} v_{\text{rad}} + E_{\text{amb}} T_{\text{amb}}$$

Four states:
x_1 Zone temp
$x_2..x_4$ Wall temps

Input:
Heat flux to zone

Disturbance:
Solar radiation

Disturbance:
External temperature

Problem Formulation: Simplest Configuration

Minimize total thermal energy use

Forecast temperature / solar radiation i steps in the future

$$\min_{u} \sum_{i=0}^{N} |u_i|$$

$$\text{s.t. } x_{i+1} = Ax_i + Bu_i + E_{rad} v_{i,\text{rad}} + E_{amb} T_{i,\text{amb}}$$

$$y_i = c^T \cdot x_i$$

$$|y_i - 21| \leq 2$$

$$0 \leq u_i \leq 1$$

Output : Room temperature

Comfort constraints

Input: Heating 0-100%

Simple MPC

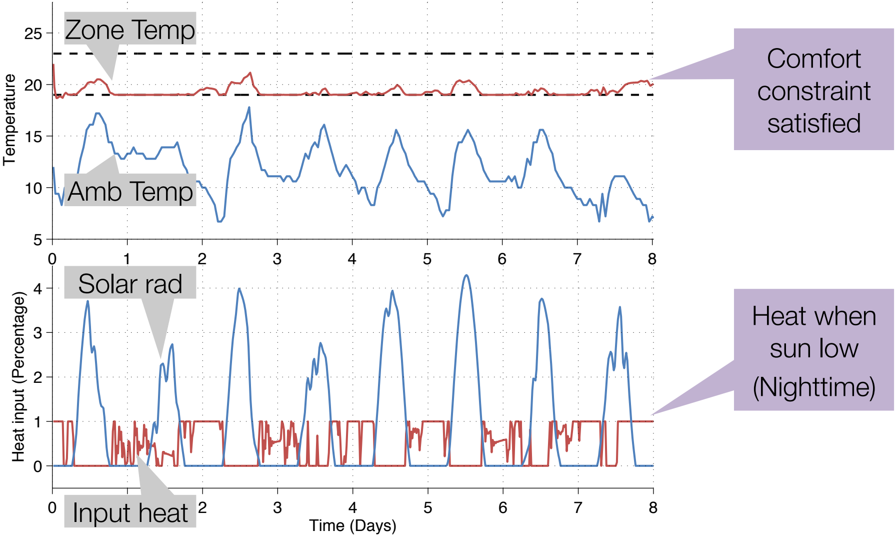

Annualized energy used: 81.6 kWh / m^2

Problem Formulation: Nighttime Setbacks & Pre-cooling

$$\min_u \sum_{i=0}^{N} |u_i|$$

$$\text{s.t. } x_{i+1} = Ax_i + Bu_i + E_{rad} v_{i,\text{rad}} + E_{amb} T_{i,\text{amb}}$$

$$y_i = c^T \cdot x_i$$

$$|y_i - 21| \leq 2 + \sigma_i$$

$$0 \leq u_i \leq 1$$

Nighttime setbacks

$$\sigma_i = \begin{cases} 0 & i \in \text{'daytime'} \\ 6 & i \in \text{'nighttime'} \end{cases}$$

Night Setback

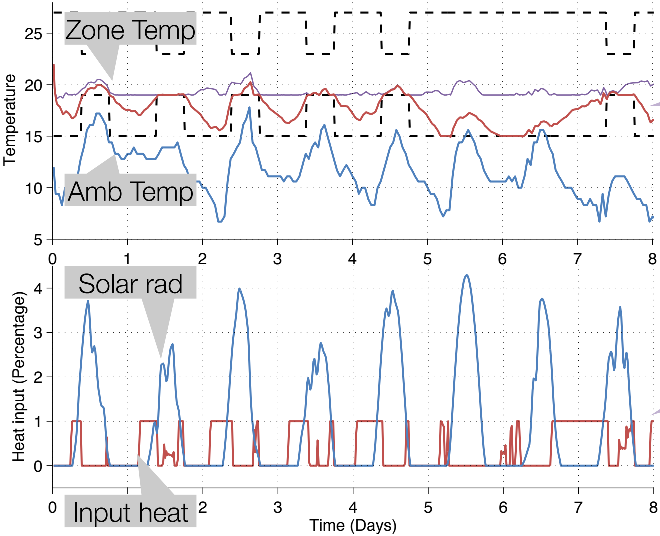

Timing of setback is automatic

Early-morning pre-heating

Annualized energy used: 55.7 kWh / m^2

Problem Formulation: Time-of-Use Pricing

$$\min_{u} \sum_{i=0}^{N} c_i \cdot |u_i|$$

Time-of-Use Tariff

$$c_i := \begin{cases} c_{\text{day}} & \text{Daytime } 9h - 18h \\ c_{\text{night}} & \text{Nighttime } 18h - 9h \end{cases}$$

$$\begin{aligned}
\text{s.t. } x_{i+1} &= Ax_i + Bu_i + E_{rad} v_{i,\text{rad}} + E_{amb} T_{i,\text{amb}} \\
y_i &= c^T \cdot x_i \\
|y_i - 21| &\leq 2 + \sigma_i \\
0 \leq u_i &\leq 1
\end{aligned}$$

Time-of-Use Tariff

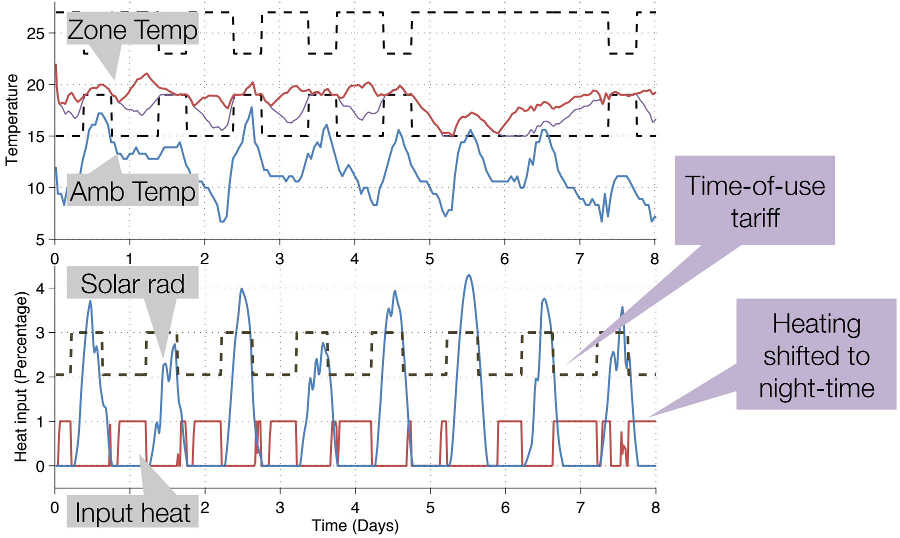

Annualized energy used: 71.3 kWh / m^2

Annualized Comparison

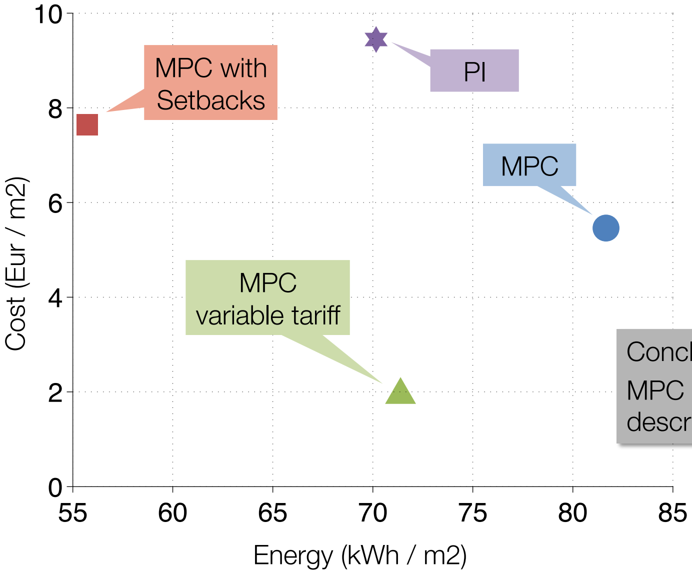

Conclusion:

MPC is easy to tune / describe complex goals

Playing the Market: New York Demand Response

Can bid 'negawatts' on open market - Paid to reduce consumption

Question: Reduce from what?!

Complex regulations define 'baseline': function of usage over x previous days

Can we 'control' our benchmark to gain income?

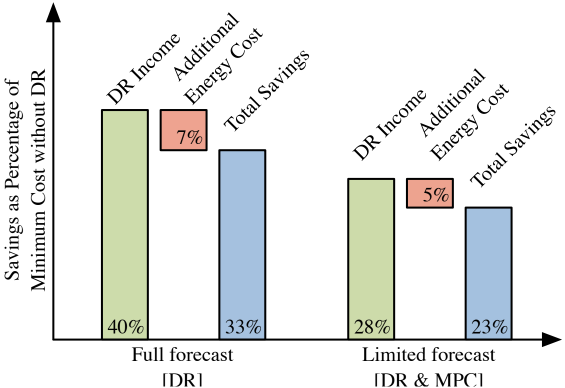

Mathematical formulation

$$
\begin{aligned}
u^\star(x) := \text{argmin} \quad & x_N^T Q_f x_N + \sum_{i=0}^{N-1} x_i^T Q x_i + u_i^T R u_i \\
\text{s.t.} \quad & x_0 = x && \text{measurement} \\
& x_{i+1} = A x_i + B u_i && \text{system model} \\
& C x_i + D u_i \le b && \text{constraints} \\
& R \succ 0, Q \succ 0 && \text{performance weights}
\end{aligned}
$$

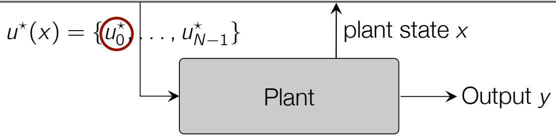

Each sample time:

- Measure /estimate current state
- Find the optimal input sequence for the entire planning window
- Implement only the **first** control action

Summarizing

Need:

- A model of the system
- A state observer
- Define the optimal control problem
- Set up the optimization problem
- Get the optimal control sequence (solve the optimization problem)
- Verify that the closed-loop system performs as desired,
  e.g., check performance criteria, robustness, real-time aspects,...

Important aspects of Model Predictive Control

Main advantages:

- Systematic approach for handling **constraints**
- High **performance** controller

Main challenges:

- **Feasibility**:
  Optimization problem may become infeasible at some future time step, i.e. there does not exist a plan satisfying all constraints
- **Stability**:
  Closed-loop stability, i.e. convergence, is not guaranteed
- **Robustness**:
  The closed-loop system is not robust against uncertainties or disturbances
- **Implementation**:
  MPC problem has to be solved in real-time, i.e. within the sampling time of the system, and with available hardware (storage, processor,...).

Outline

1. Introduction to MPC
    - Concept
    - The math
    - Examples

2. Administration

Course information

Professor: Colin Jones, Room ME C2 408
colin.jones@epfl.ch

Lectures: Pre-recorded
- Videos and weekly schedule on moodle

Supervision: Fridays 15h15 - 17h00
- In person in Room CE 4
- Asynchronously via Ed Discussion (link on moodle)

Lecture Notes: On moodle

All details will be updated on Moodle

Exam & Grades

Written Exam	60%
Mini-Project	40%
Exercises	Not graded

Class Schedule

Week 1	Introduction
Week 2	Unconstrained control
Week 3	Optimization
Week 4	Constrained systems
Week 5	MPC
Week 6	Practical MPC
Week 7	Robust MPC
Week 8	Robust MPC
Week 9	Advanced topics in MPC
Week 10	Advanced topics in MPC
Week 11	Advanced topics in MPC
Week 12	Mini-project and guest speakers
Week 13	Mini-project and guest speakers
Week 14	Mini-project and guest speakers

Mini-project

- Groups of three
- **Report worth 40% of the final grade**
- Three-week project

Literature

No required textbook

Model Predictive Control:

- Model Predictive Control: Theory and Design, James B. Rawlings and David Q. Mayne, 2009 Nob Hill Publishing
- Predictive Control with Constraints, Jan Maciejowski, 2000 Prentice Hall
- Predictive Control for linear and hybrid systems, 2014 F. Borrelli, A. Bemporad and M. Morari
  Available for free at
  `http://www.mpc.berkeley.edu/mpc-course-material`

Optimization:

- Convex Optimization, Stephen Boyd and Lieven Vandenberghe, 2004 Cambridge University Press
- Numerical Optimization, Jorge Nocedal and Stephen Wright, 2006 Springer

[0] Parts of the notes in this lecture are based on or have been extracted from: Linear Dynamical Systems, Stephen Boyd, Stanford; Convex Optimization, Stephen Boyd, Stanford; Model Predictive Control, Manfred Morari, ETH Zurich; Model Predictive Control, Francesco Borrelli, Berkeley

Summary

- MPC uses a model of the system to predict the future trajectory
- We minimize a value function to choose the ‘best’ of these future trajectories
- Benefit: Nonlinear, constrained systems with complex objectives