Monitoring, Modelling and Identification of Data Center Servers

Martin Eriksson

Sep. 14th, 2017

Acknowledgement

Damiano Varagnolo

Arash Mousavi

Riccardo Lucchese

I would also like to thank Filip Blylod and all the staff at Swedish Institute of Computer Science ICE (SICS)

Contents

Monitoring, Modelling and Identification
of Data Center Servers

Data centers

Server blades

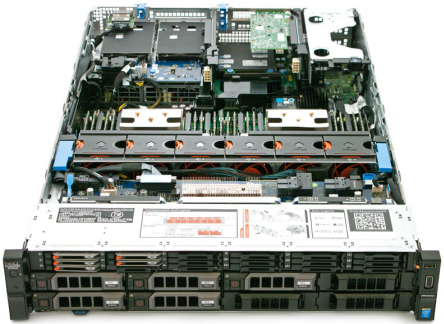

Introduction

What is a data center?

- specific for hosting computer systems
- runs IT system applications
- consumes a huge amount of electricity

Categorising data centers

Categorising data centers

courtesy of Pär Bäckström

Footprint of data centers

- $\sim 2\%$ of the global electrical energy production
- comparable to Italy's electrical capacity
- yearly CO_2 emissions in 2013: 38.6 million metric tonnes

ABB Review - Datacenters (2013)
The corporate technical journal

How are data centers organized?

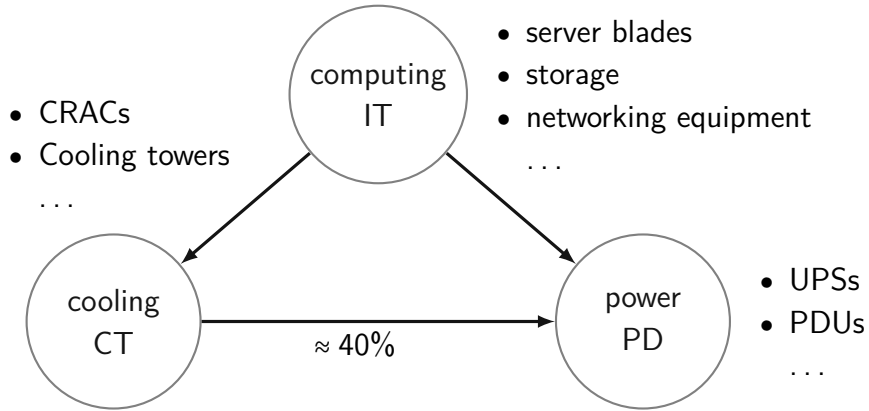

- cooling of the server halls responsible for around 40%

Scope

Scope

Problem: cooling accounts for a large part of the energy budget ($\approx 40\%$)

Objective: make it more efficient

How: implement predictive control strategies

Why predictive control can be beneficial?

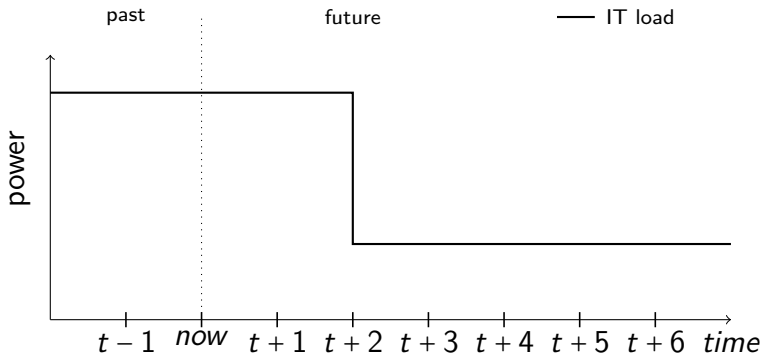

Why predictive control can be beneficial?

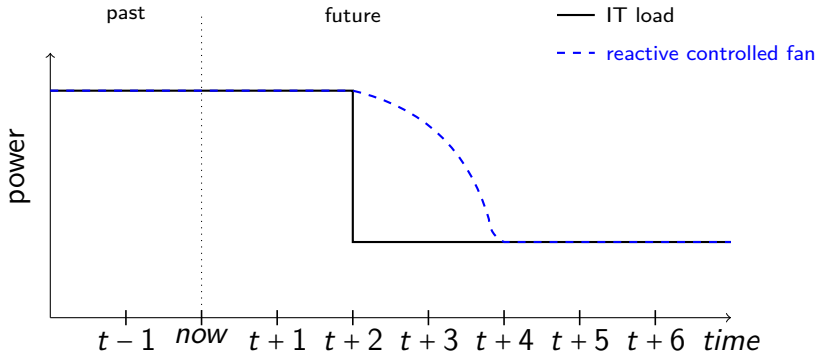

Why predictive control can be beneficial?

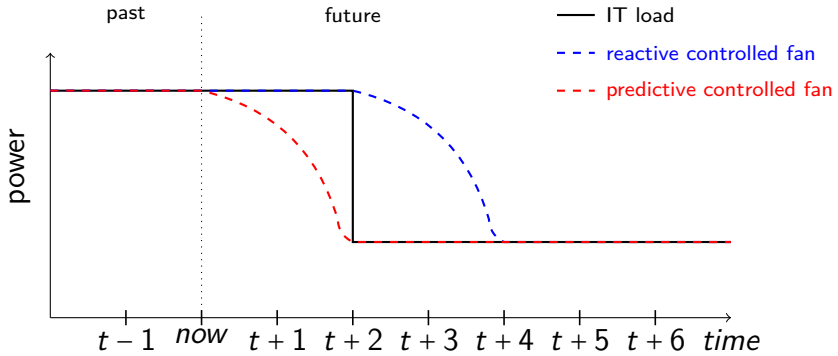

Why predictive control can be beneficial?

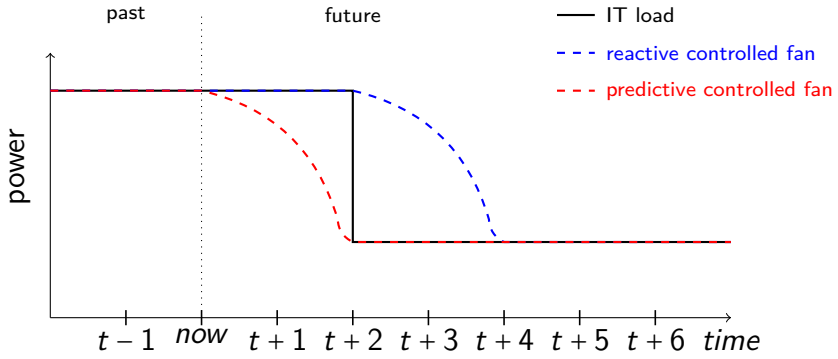

Main ingredients:

1. be able to predict the future behaviour of the system
2. have an accurate model of the system

Aim

Aims

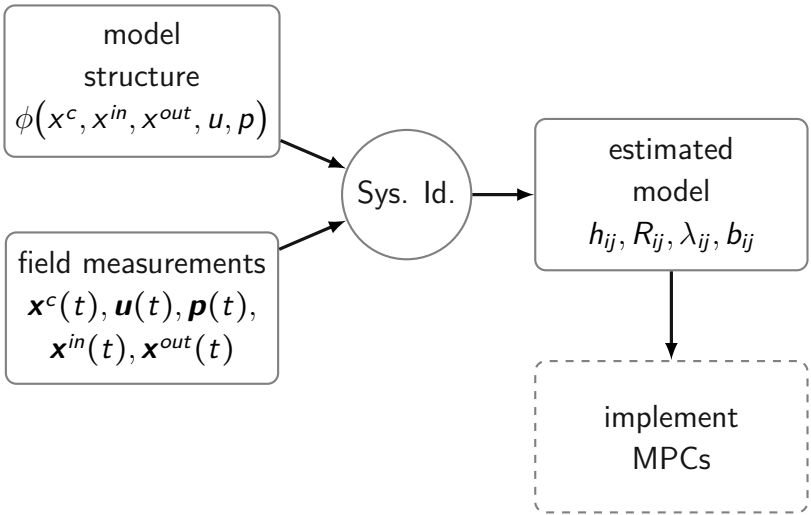
model
structure
$\phi\left(x^c, x^{in}, x^{out}, u, p\right)$
field measurements
$\boldsymbol{x}^c(t), \boldsymbol{u}(t), \boldsymbol{p}(t),$
$\boldsymbol{x}^{in}(t), \boldsymbol{x}^{out}(t)$
Sys. Id.
estimated
model
$h_{ij}, R_{ij}, \lambda_{ij}, b_{ij}$
implement
MPCs

Contributions

Contributions

- control-oriented thermal model tuned using field data

Contributions

- control-oriented thermal model tuned using field data
- *CISSI*: A software suite that allows real-time monitoring of servers

Contributions

- control-oriented thermal model tuned using field data
- *CISSI*: A software suite that allows real-time monitoring of servers
- tailored identification algorithms for estimating thermal model from field data

Modelling

Facebook's Windmill V2 server blade

Schematic of the thermal model

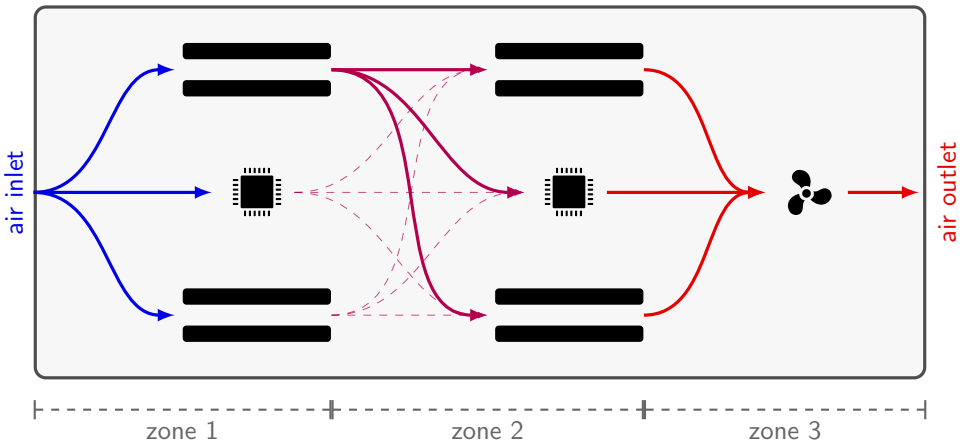

Mathematical notation

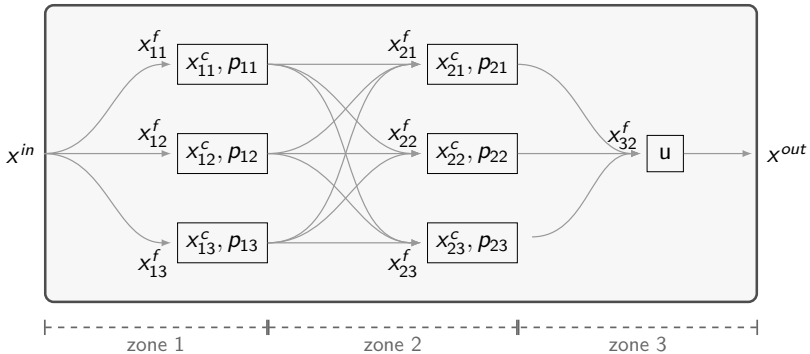

- *states*: x^c_{ij}, x^f_{ij}, x^{out}
- *exogenous inputs*: p_{ij}, x^{in}
- *auxiliary variables*: $f_{i,j\to k}$
- *manipulable input*: u (influences $f_{i,j\to k}$)

Ingredients to be modelled

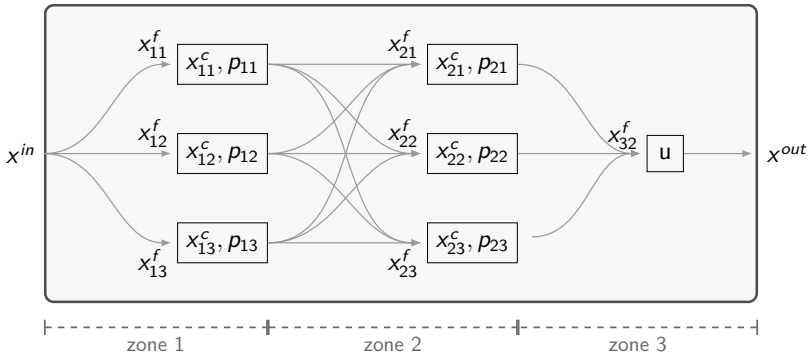

- mass of the air flows f
- temperature of the air flows x^f
- temperature of the IT components x^c

Air flow model

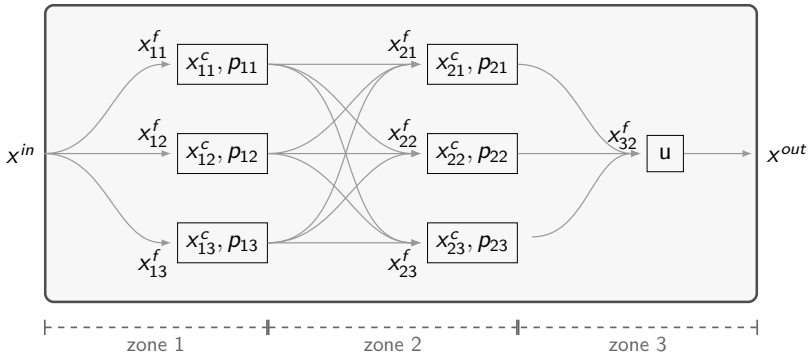

Flow interactions = linear mapping of the total air flow mass:

$$\boldsymbol{f} = \boldsymbol{\Lambda} u$$

Air flow temperature model

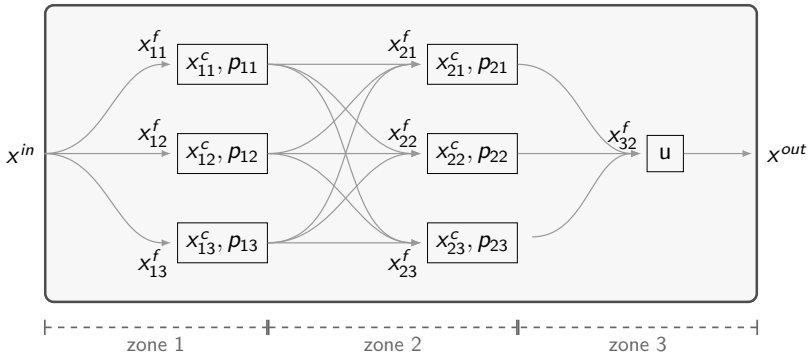

- perfect flow mixing
- heat energy conservation

$$x^f_{2j} = \frac{\sum_{k=1}^{3} \lambda_{2,k\to j} f_{1,in\to k}\left(x^{in} + \frac{h_{1k}}{c_p}\left(x^c_{1k} - x^{in}\right)\right)}{f_{2,j}}$$

Model of the temperatures of the IT components

$$\dot{x}_{ij}^{c} = \underbrace{-h_{ij} f_{i,j} \left(x_{ij}^{c} - x_{ij}^{f} \right)}_{\text{convection}} + \underbrace{\begin{bmatrix} R_{(ij)} & \rho_{ij} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}^{c} \\ x^{in} \end{bmatrix}}_{\text{conduction}} + \underbrace{b_{ij} p_{ij}}_{\text{electrical power}}$$

Parameters to be estimated: h_{ij}, R_{ij}, λ_{ij} and b_{ij}

Monitoring

Overview of the logical structure of *CISSI*

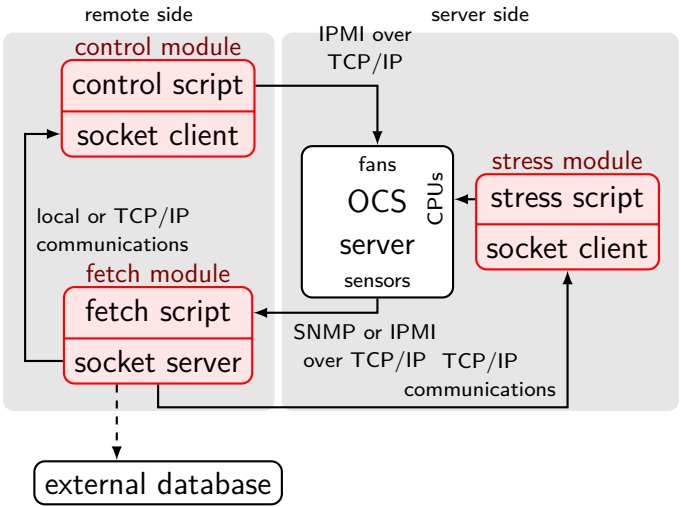

Data acquisition module - fetch

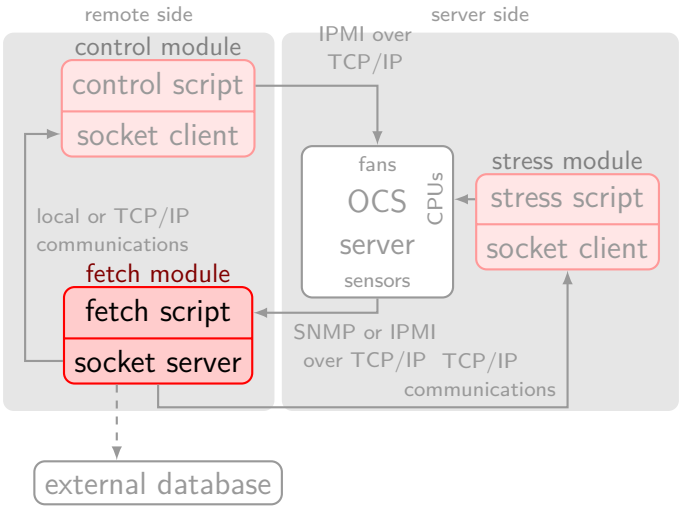

- collects temperatures, fan speeds, and CPU loads

CPU stressing module - stress

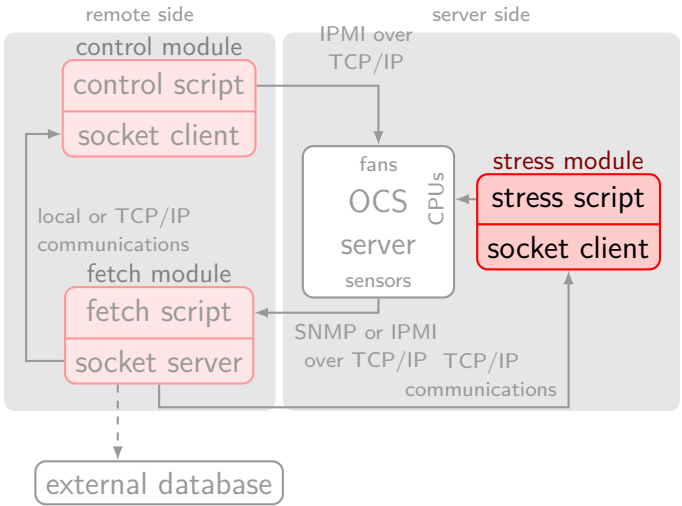

- stresses the CPUs, collects and save the data

Server control module - control

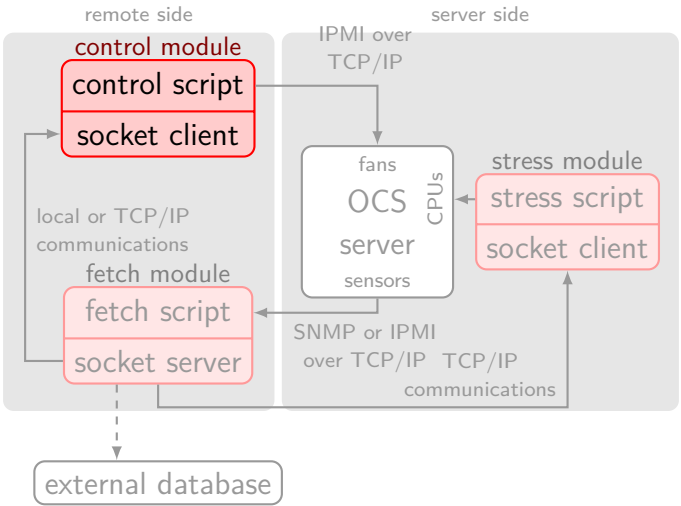

- uses *ipmi-raw*

Acquisition example results

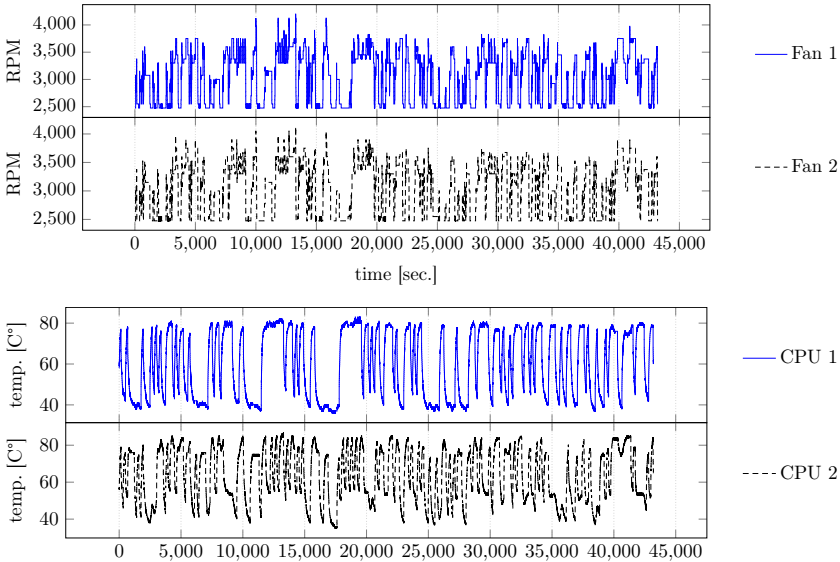

Acquisition example results

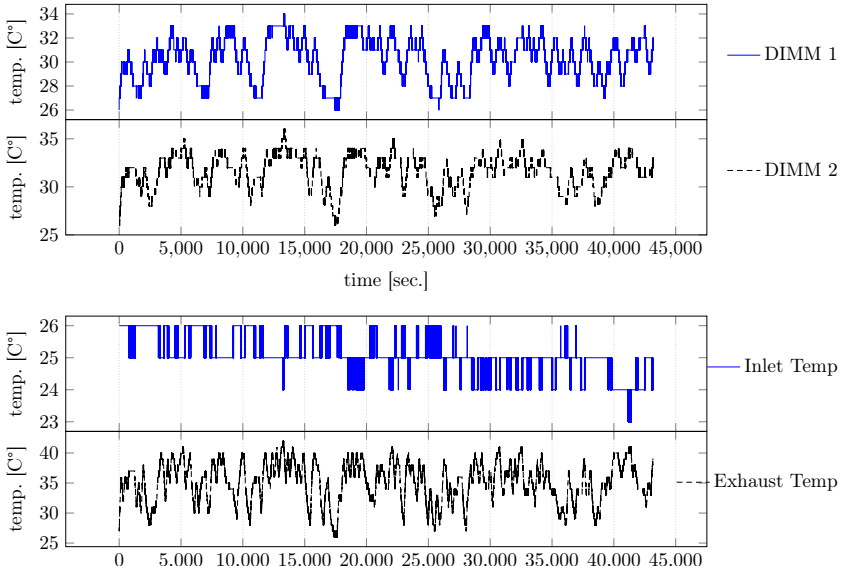

System identification

Background on system identification

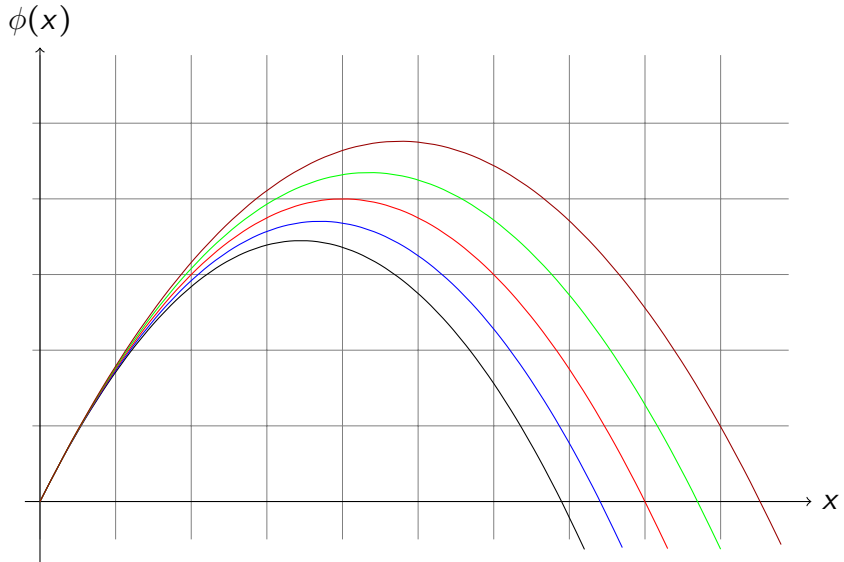

Background on system identification

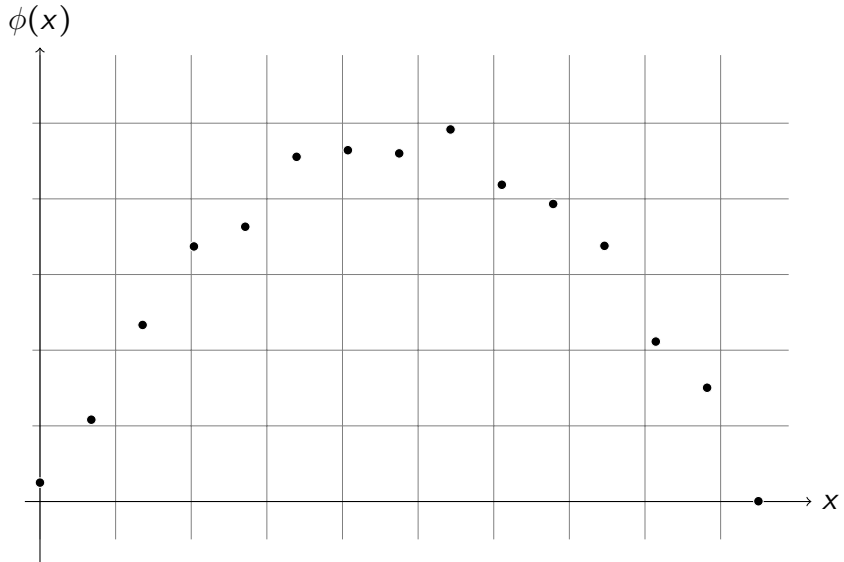

Background on system identification

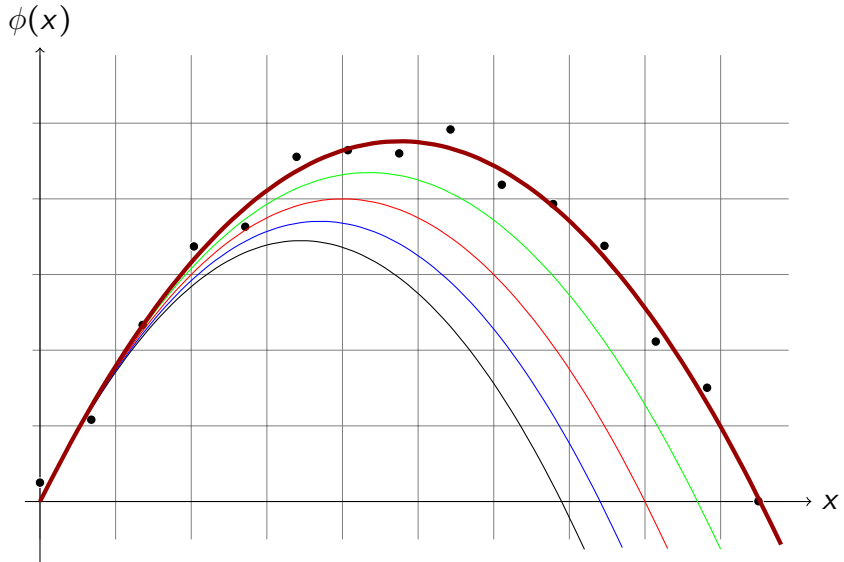

System identification

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}\in\Theta} \sum_{i,j} J_{ij}(\boldsymbol{\theta})$$

$$\text{s.t.} \quad \sum_{j=1}^{3} \lambda_{1,in\to j} = 1$$

$$\sum_{k=1}^{3} \lambda_{2,k\to j} = 1 \quad j = 1,2,3$$

$$\sum_{j=1}^{3} \lambda_{3,j\to out} = 1$$

$$J_{ij}(\boldsymbol{\theta}) := \sum_{t=1}^{N-1} \left(x_{ij}^c(t+1) - x_{ij}^c(t) - \Delta\Psi_{ij}(t\,;\,\boldsymbol{\theta}) \right)^2$$

$$\boldsymbol{x}^c(t), u(t), \boldsymbol{p}(t), x^{in}(t), x^{out}(t) \quad \text{for} \quad t = 1,\ldots,N$$

Numerical results

Numerical results

h_{11}	0.00000		$R_{(12),11}$	0.00013
h_{12}	0.00068		$R_{(21),11}$	0.00000
h_{21}	0.00000		$R_{(22),11}$	0.00000
h_{22}	0.00009		$R_{(11),12}$	0.00493
$\widetilde{h}_{11}$	359.83539		$R_{(21),12}$	0.00000
$\widetilde{h}_{12}$	906.39287		$R_{(22),12}$	0.00150
$\widetilde{h}_{21}$	1459.44161		$R_{(11),21}$	0.00000
$\widetilde{h}_{22}$	1459.44161		$R_{(12),21}$	0.00019
λ_{11}	0.99291		$R_{(22),21}$	0.00000
λ_{12}	0.00709		$R_{(11),22}$	0.00001
λ_{211}	0.47267		$R_{(12),22}$	0.00001
λ_{212}	0.06944		$R_{(21),22}$	0.00234
λ_{221}	0.52733		ρ_{11}	0.00000
λ_{222}	0.93056		ρ_{12}	0.00000
λ_{31}	0.50000		ρ_{21}	0.00000
λ_{32}	0.50000		ρ_{22}	0.00001
b_{11}	0.01167		b_{21}	0.01246
b_{12}	0.68022		b_{22}	0.83668

Numerical results

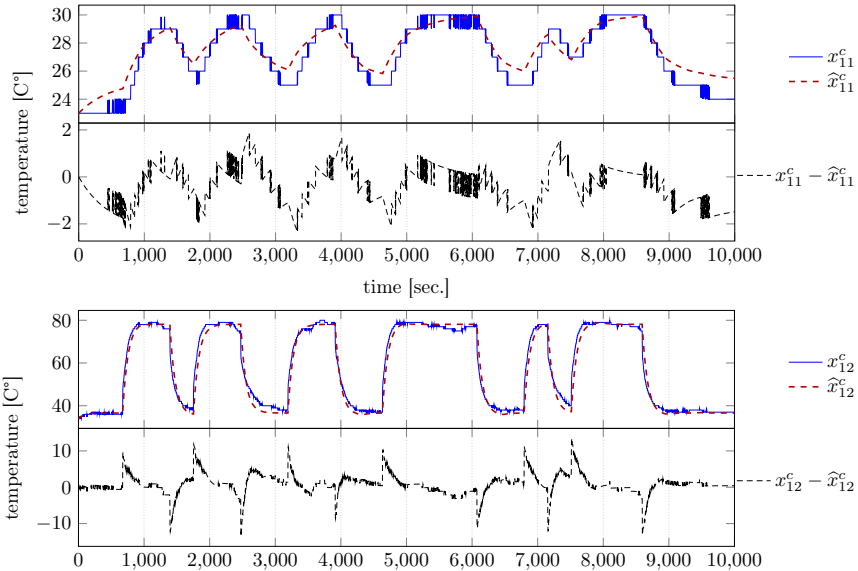

Conclusions and future work

Conclusions

- To implement predictive control solutions, one needs accurate models

Conclusions

- To implement predictive control solutions, one needs accurate models
- A key part to obtain accurate models is estimate the unknown parameters of the system

Conclusions

- To implement predictive control solutions, one needs accurate models
- A key part to obtain accurate models is estimate the unknown parameters of the system
- This implies a need for collecting real-data from the system

Future work

Modeling: improving or comparing the obtained models

- non-linear terms and less assumptions
- Computational Fluid Dynamics (CFD) tools

Future work

Monitoring: continuously comparing the model with the measurements

- sanity checks and signal faults

Future work

Monitoring: continuously comparing the model with the measurements

- sanity checks and signal faults

Controlling: building model-based control strategies with our results as a basis

- LQRs and MPCs

Future work

Monitoring: continuously comparing the model with the measurements

- sanity checks and signal faults

Controlling: building model-based control strategies with our results as a basis

- LQRs and MPCs

CISSI: improve the software suite *CISSI*

- the stress module
- the control module

Monitoring, Modelling and Identification of Data Center Servers

Martin Eriksson

Sep. 14th, 2017