

First part: what is a datacenter and what is its importance in our society

Data centers control:
challenges and opportunities

Damiano Varagnolo
Luleå University of Technology

Aug. 24th, 2015



Data centers control: challenges and opportunities

Damiano Varagnolo

Luleå University of Technology

Aug. 24th, 2015



└─What is a datacenter?

Datacenters are the pillar of our IT society: banks, hospitals, governments, everything that is IT-based relies on datacenters.

What is a datacenter?

"a facility that centralizes an organization's IT operations and equipment, and where it stores, manages, and disseminates its data"

 Google (Aug. 2015)

What is a datacenter?

“a facility that centralizes an organization’s IT operations and equipment, and where it stores, manages, and disseminates its data”



Google (Aug. 2015)

└─What is a datacenter?

“a data center is a complex production facility, more similar to a chemical plant, power plant or other mission critical facility, than a simple building with a room full of servers” [?].

A datacenter is a composition of IT systems and support infrastructure supplying power and cooling:

- IT systems: servers, storage devices, networking devices, middleware, software
- support infrastructure: backup power generators, UPSs, power distribution units, batteries and CT systems (server fans, CRACs, chillers, and cooling towers)

What is a datacenter?
TierPoint datacenters, Dallas



What is a datacenter?

TierPoint datacenters, Dallas



2015-10-22

Datacenters Control

└ What is a datacenter?

FB in Lulea = 60MW

What is a datacenter?
Facebook, Luleå, Sweden



What is a datacenter?

Facebook, Luleå, Sweden



└─What is a datacenter?

What is a datacenter?
Sun Microsystems Modular Datacenter



What is a datacenter?

Sun Microsystems Modular Datacenter



└ Why do they exist?

Why do they exist?

- lower delays and high bandwidth between the servers
- easier to maintain
- higher level of data security and privacy

Why do they exist?

- lower delays and high bandwidth between the servers
- easier to maintain
- higher level of data security and privacy

└─What is their footprint?

What is their footprint?

In 2013 for EU-28:

- electrical energy consumption: 103.4 GWh (~ 3% of the total generation)
- industry market: 18.85 billion €
- CO₂ emissions: 38.6 million tonnes (347g/kWh)

 Pan European Datacenter Academy Project (2014)
Final Report Summary

What is their footprint?

In 2013 for EU-28:

- electrical energy consumption: 103.4 GWh (~ 3% of the total generation)
- industry market: 18.85 billion €
- CO₂ emissions: 38.6 million tonnes (347g/kWh)



Pan European Datacenter Academy Project (2014)

Final Report Summary


└─What will their footprint be?

What will their footprint be?



Cisco Global Cloud Index 2013 - 2018

What will their footprint be?



Cisco Global Cloud Index 2013 - 2018

└─What will their footprint be?

What will their footprint be?

"forecasted energy savings from implementing *best practices* in datacenters in EU is 15,500GWh per year, approximatively equivalent to the energy consumed by 1M EU households yearly, 1.1 billion euro in electricity costs, 5.4 Mtonnes of CO₂"

 [Pan European Datacenter Academy Project \(2014\)](#)
Final Report Summary

What will their footprint be?

“forecasted energy savings from implementing *best practices* in datacenters in EU is 15,500GWh per year, approximatively equivalent to the energy consumed by 1M EU households yearly, 1.1 billion euro in electricity costs, 5.4 Mtonnes of CO₂”



[Pan European Datacenter Academy Project \(2014\)](#)

Final Report Summary

└ The best practices & Automatic Control

Automatic control is important for the datacenters industry because modern and extremely efficient data centers are a minority in the global panorama of data centers.

In this presentation we do not touch the concepts of:

- "green cloud computing", i.e., how to improve algorithms so that they require less CPU efforts for solving the same task
- how to remove useless AC/DC conversions
- where to locate the datacenter
- how to design the datacenter

- Optimize Air Management
- Right-Size the Design
- Optimize the Central Plant
- Design Efficient Air Handling
- Improve Humidification Systems and Controls
- Specify Efficient Power Supplies
- Consider On-Site Generation
- Employ Liquid Cooling
- Reduce Standby Losses
- Improve Design, Operations, and Maintenance Processes

 Best Practices for Data Centers: Lessons Learned from Benchmarking 22 Data Centers
Greenberg et al. (2006)
ACEEE Summer Study on Energy Efficiency in Buildings

The best practices & Automatic Control

- *Optimize Air Management*
- Right-Size the Design
- *Optimize the Central Plant*
- Design Efficient Air Handling
- *Improve Humidification Systems and Controls*
- Specify Efficient Power Supplies
- Consider On-Site Generation
- Employ Liquid Cooling
- *Reduce Standby Losses*
- *Improve Design, Operations, and Maintenance Processes*



Best Practices for Data Centers: Lessons Learned from Benchmarking 22 Data Centers

Greenberg et al. (2006)

ACEEE Summer Study on Energy Efficiency in Buildings


Next part: what is a datacenter from automatic control perspectives

└ Control perspectives on a single server


"After the power draw from IT equipment, cooling the data center frequently consumes the largest amount of energy." [?]

Liquid cooling has high installation costs and safety concerns, so it is not so widely used.

Control perspectives on a single server



Control perspectives on a single server




└ Control perspectives on a single server

Problem: workloads are stochastic (and bursty and non-Gaussian at all). There is a need to predict them.

Most important thing for a datacenter is uptime and QoS.

Example of the importance of QoS: adding a 100 ms on the time to complete a user search requests induces a 1% of sales loss to Amazon.



Control perspectives on a single server

Control problem

Control aim

- maximize IT QoS + minimize electric consumptions

Inputs

Controllable:

- CPUs frequencies
- number of VMs
- fan speed

Non-controllable:


- IT load

Disturbances

- temperature of the air inlet

Control perspectives on a single server

Control perspectives on a single server
Hints on the dynamics



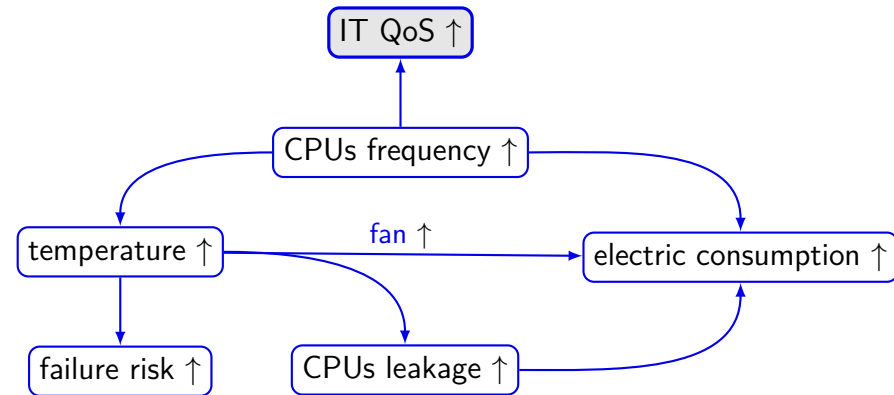
There are strong IT-CT couplings; thus operating IT and CT separately is suboptimal. Timescales: temperatures evolve in the order of minutes, while CPU power states can be changed as frequent as milliseconds.

Consequence of raising ambient temperatures too high:

1. as the ambient air temperature increases, the fan motor consumes significantly more energy to ensure the critical CPUs and memory modules stay within the acceptable operating temperature range
2. as soon as the ambient temperature rises to the point where the CPU chips enter the “leakage zone” leakage power grows exponentially with further increases in ambient temperature
3. acoustic noise from servers increases with the 5th power of the fan RPMs, and this implies indirect costs to face the additional noise
4. the time between a failure of the data center cooling system and the IT equipment reaching critical temperatures diminishes with increased starting ambient temperature


Control perspectives on a single server

Hints on the dynamics




Control perspectives on a servers rack

Control perspectives on a servers rack



Control perspectives on a servers rack



Control perspectives on a servers rack

Control aim

- maximize IT QoS + minimize electric consumptions

Inputs

Controllable:

- CRAC fans speed
- servers & VMs ON / OFF
- IT loads assignment

Non-controllable:

- IT load

Disturbances

- temperature of the air inlet
- thermal couplings

Again strong IT-CT couplings, so that operating IT and CT separately is suboptimal. Racks are modular, and this suggests the implementation of hierarchical / distributed estimation and control strategies.

One can decide which servers should be turned on / off and where the loads should be served; this is not negligible because different servers are more or less close to the internet fibers \Rightarrow different latencies and energy efficiencies.

Idle servers consume about 60% of their peak power consumption but operate between 10% and 50% of their maximum utilization. Aim is thus make active servers run at 100%, turn off the other ones.

Problem: setup time (time required to turn a server back on) is on the order of minutes, while migrating VMs is on the order of seconds.

Turning on / off things make the control problem a mixed-integer control problem.

Moreover thermal dynamics are nonlinear and costs are non convex, so the problem is the worst possible \rightarrow need for simplifications and intelligent approximations.

Control perspectives on a servers rack

Control problem

Control aim

- maximize IT QoS + minimize electric consumptions

Inputs

Controllable:

- CRAC fans speed
- servers & VMs ON / OFF
- IT loads assignment


Non-controllable:

- IT load

Disturbances

- temperature of the air inlet
- thermal couplings

Control perspectives on a whole datacenter




Again strong IT-CT couplings, so that operating IT and CT separately is suboptimal. Both how and where the workload is executed affect the total energy consumption: the way workload is executed impacts the efficiency at which IT operates, while the location where the workload is executed impacts the efficiency at which CT operates. Datacenters are modular, and this suggests the implementation of hierarchical / distributed estimation and control strategies. Nonetheless this is problematic because cooling resources are usually shared, so distributed controllers need data from a whole-datacenter-level point of view.

Datacenters are also large-scale systems: the number of state variables ranges in the order of ten of thousands for a medium-size data center, since it is useful to let each CPU's temperature, workload, and amount of resources used be part of the state of the system.

Every different datacenter has very specific thermal dynamics, and there is no general “whole datacenter model”. This complicates the development of general-purpose MPCs – one should learn each model of each single datacenter independently.


An other problem: there is no algorithm that accurately models the air flows and that is fast enough to be used for real-time control.

Control perspectives on a whole datacenter



Control perspectives on a smart grid

Control perspectives on a smart grid



Exchanging IT loads enables exploiting geographically-varying electricity prices.

Interacting with the power-grid enables taking advantage of a time-varying electricity prices and performing ancillary services like demand response for electrical networks (but this requires predictive control approaches).


Interacting with the district heating enables selling the exhaust thermal power and performing ancillary services like demand response for thermal networks.

Control perspectives on a smart grid

computational

electrical

thermal



Next part: current states of practice and art, and research directions

└ State of the practice

Big need for monitoring: ("you cannot improve what you are not measuring").


Importance of model predictive approaches is that preemptive over-cooling may be used to take advantage of time-varying electricity prices, external weather and variability of the load.

Problem is that to develop a model predictive controller one must first have the models, and the models should be learned automatically.

State of the practice



State of the practice

control level


└ State of the practice

Big need for monitoring: ("you cannot improve what you are not measuring").

Importance of model predictive approaches is that preemptive over-cooling may be used to take advantage of time-varying electricity prices, external weather and variability of the load.

Problem is that to develop a model predictive controller one must first have the models, and the models should be learned automatically.

State of the practice



State of the practice

*control level**control technology*

single server

servers rack

whole datacenter

smart grid

no IT-CT couplings

IT, CT control: mostly PIDs
no learning, no robustness
big focus on DCIMs

└ State of the practice

Big need for monitoring: ("you cannot improve what you are not measuring").

Importance of model predictive approaches is that preemptive over-cooling may be used to take advantage of time-varying electricity prices, external weather and variability of the load.

Problem is that to develop a model predictive controller one must first have the models, and the models should be learned automatically.

State of the practice

control level	control technology
single server servers rack whole datacenter	<i>no IT-CT couplings</i> IT, CT control: mostly PIDs no learning, no robustness big focus on DCIMs
smart grid	some experiments on demand response and district heating integration

State of the practice

*control level**control technology*

single server

servers rack

whole datacenter

no IT-CT couplings

IT, CT control: mostly PIDs
no learning, no robustness
big focus on DCIMs

smart grid

some experiments on
demand response and
district heating integration

└ State of the art


Load shedding can be done via temporarily changing the settings of the CT and exploit the buildings thermal dynamics.

Geographical load balancing routes user requests to locations with cheaper and cleaner electricity; should also account for the status (i.e., current load and temperature) of the various datacenters.

State of the art



State of the art

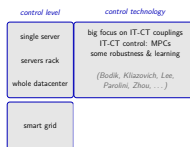
control level

└ State of the art

Load shedding can be done via temporarily changing the settings of the CT and exploit the buildings thermal dynamics.

Geographical load balancing routes user requests to locations with cheaper and cleaner electricity; should also account for the status (i.e., current load and temperature) of the various datacenters.

State of the art



State of the art

*control level**control technology*

single server

servers rack

whole datacenter

big focus on IT-CT couplings
IT-CT control: MPCs
some robustness & learning

*(Bodik, Kliazovich, Lee,
Parolini, Zhou, . . .)*

smart grid

└ State of the art

Load shedding can be done via temporarily changing the settings of the CT and exploit the buildings thermal dynamics.

Geographical load balancing routes user requests to locations with cheaper and cleaner electricity; should also account for the status (i.e., current load and temperature) of the various datacenters.

State of the art

control level	control technology
single server servers rack whole datacenter	big focus on IT-CT couplings IT-CT control: MPCs some robustness & learning (Bodik, Kliazovich, Lee, Parolini, Zhou, ...)
smart grid	load shedding geographical load balancing integration with renewables

State of the art

*control level**control technology*

single server

servers rack

whole datacenter

big focus on IT-CT couplings
IT-CT control: MPCs
some robustness & learning

(Bodik, Kliazovich, Lee,
Parolini, Zhou, ...)

smart grid

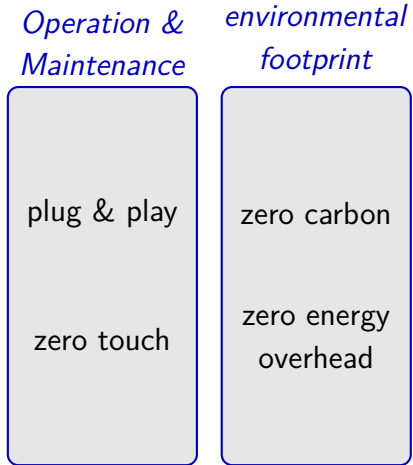
load shedding
geographical load balancing
integration with renewables

└ Research & Development directions


Research & Development directions



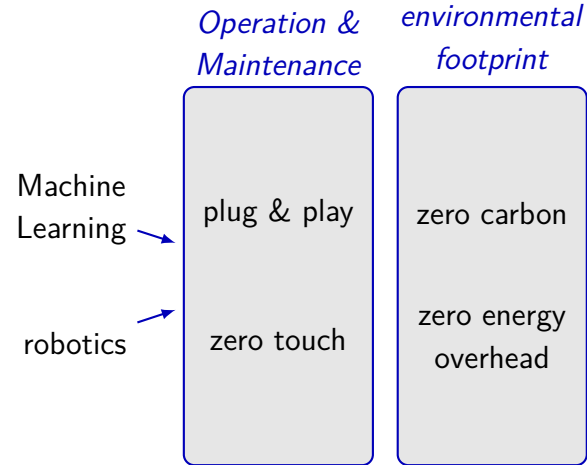
Research & Development directions



└ Research & Development directions



Research & Development directions

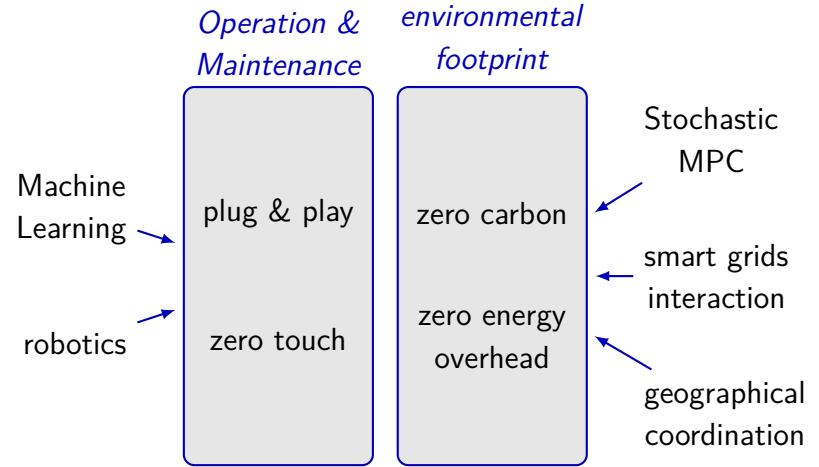


Research & Development directions

Research & Development directions



Research & Development directions






Next part: the SICS ICE testbed


└ The SICS ICE datacenter

The SICS ICE datacenter




The SICS ICE datacenter

Infrastructure to enable both small and large scale experiments.
Supports growing datacenter industry and research efforts.




2015-10-22

Datacenters Control



2015-10-22

Datcenters Control



└ The SICS ICE datacenter

“Experiment-as-a-Service” is a way to describe the value delivered by the facility: different kinds of experiments have different costs. E.g., experiments on a single server for one hour will cost less than experiments on the whole datacenter for a month.

The SICS ICE datacenter

Characteristics

- 3000 - 4000 servers (2MW)
- 160 m² lab
- biogas back up generators
- connections with the urban district heating network
- Generality, Flexibility and Expandability

The SICS ICE datacenter

Characteristics

- 3000 - 4000 servers (2MW)
- 160 m² lab
- biogas back up generators
- connections with the urban district heating network
- Generality, Flexibility and Expandability

└ The SICS ICE datacenter

“Experiment-as-a-Service” is a way to describe the value delivered by the facility: different kinds of experiments have different costs. E.g., experiments on a single server for one hour will cost less than experiments on the whole datacenter for a month.

The SICS ICE datacenter

Characteristics

- 3000 - 4000 servers (2MW)
- 160 m² lab
- biogas back up generators
- connections with the urban district heating network
- Generality, Flexibility and Expandability

Experiment-as-a-Service

The SICS ICE datacenter

Characteristics

- 3000 - 4000 servers (2MW)
- 160 m² lab
- biogas back up generators
- connections with the urban district heating network
- Generality, Flexibility and Expandability

Experiment-as-a-Service

└ In conclusion

In conclusion

- big "market", even bigger in the future

In conclusion

- big “market”, even bigger in the future

└ In conclusion

In conclusion

- big "market", even bigger in the future
- distributed optimization, system identification and stochastic control are *the* tools

In conclusion

- big "market", even bigger in the future
- distributed optimization, system identification and stochastic control are *the* tools

└ In conclusion

In conclusion

- big "market", even bigger in the future
- distributed optimization, system identification and stochastic control are *the* tools
- everybody can do experiments

In conclusion

- big "market", even bigger in the future
- distributed optimization, system identification and stochastic control are *the* tools
- everybody can do experiments

Data centers control: challenges and opportunities

Damiano Varagnolo

Luleå University of Technology

Aug. 24th, 2015

`damiano.varagnolo@ltu.se`

References