Model Predictive Control of Data Centers in the Smart Grid Scenario

Luca Parolini, Bruno Sinopoli, Bruce H. Krogh

Electrical and Computer Engineering
Carnegie Mellon University

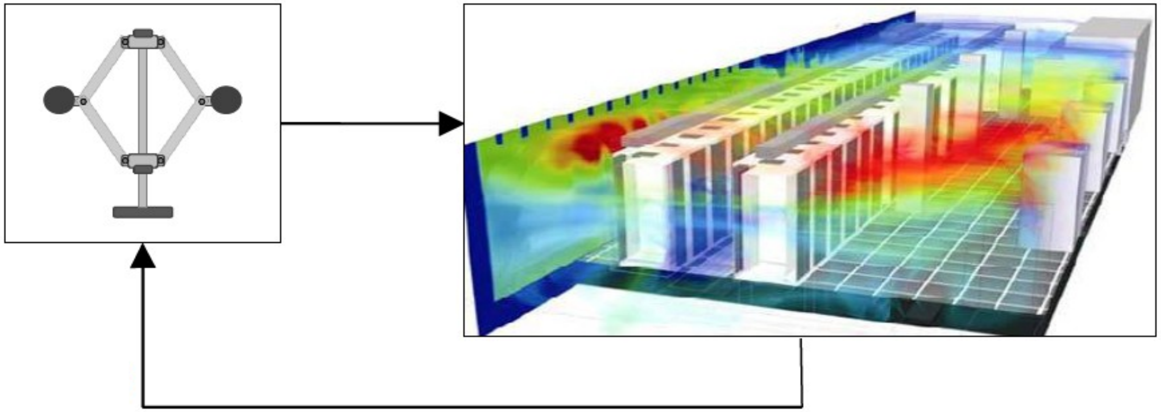

IFAC – 1st Sep. 2011

Electricity consumption

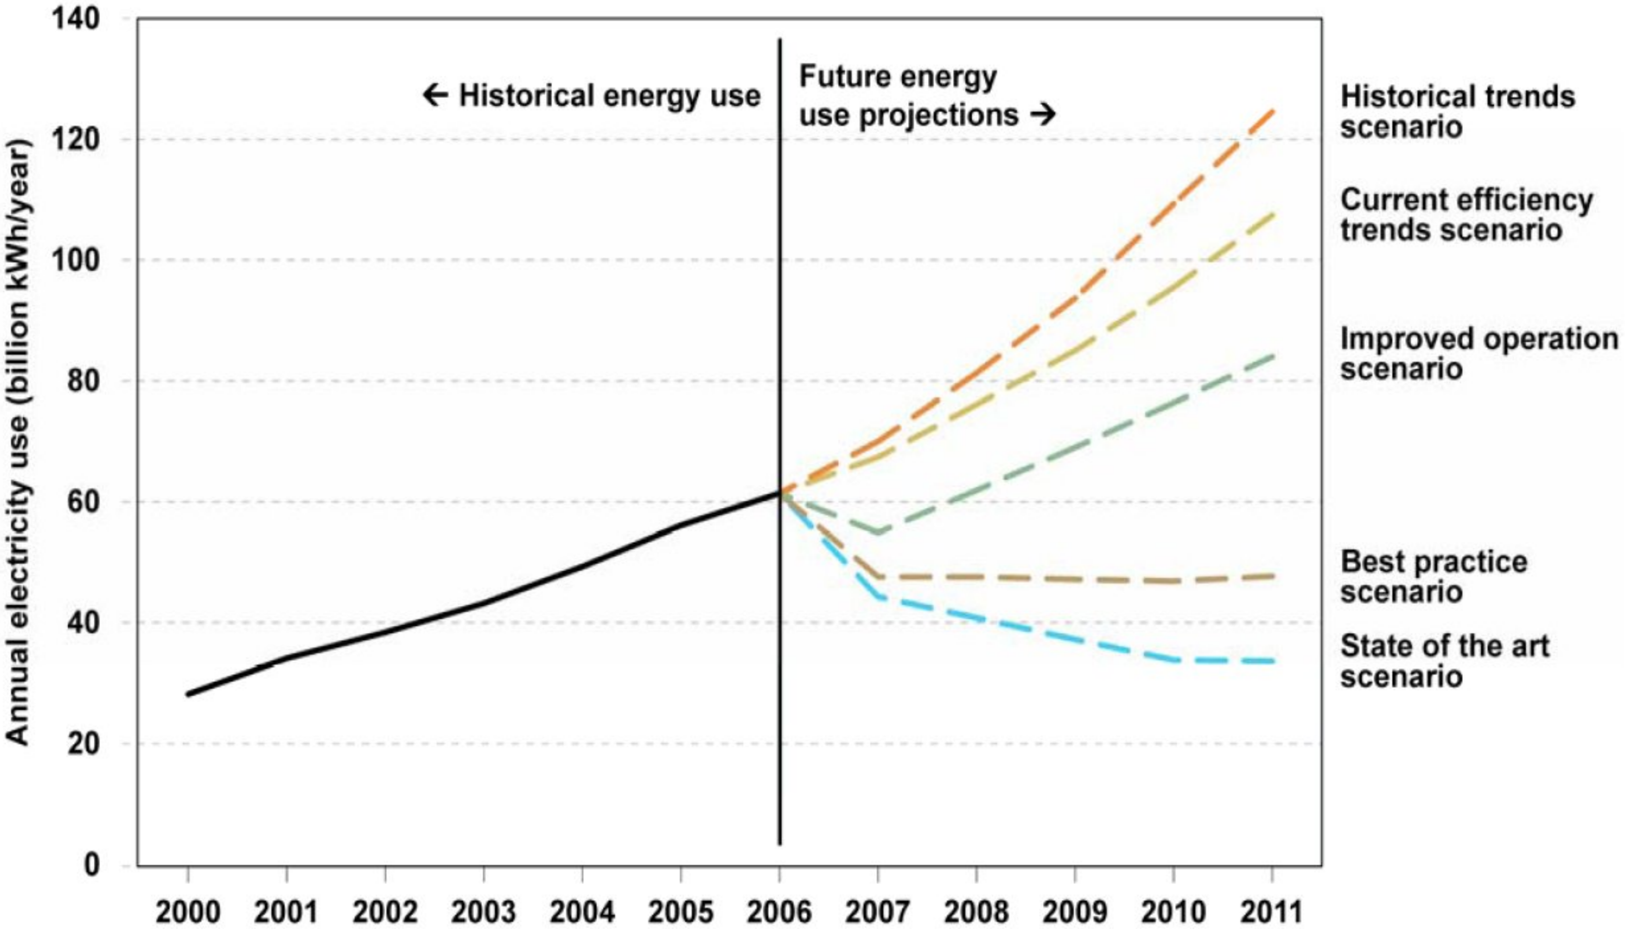

Environmental Protection Agency, Tech. Rep. 2007

Electricity consumption

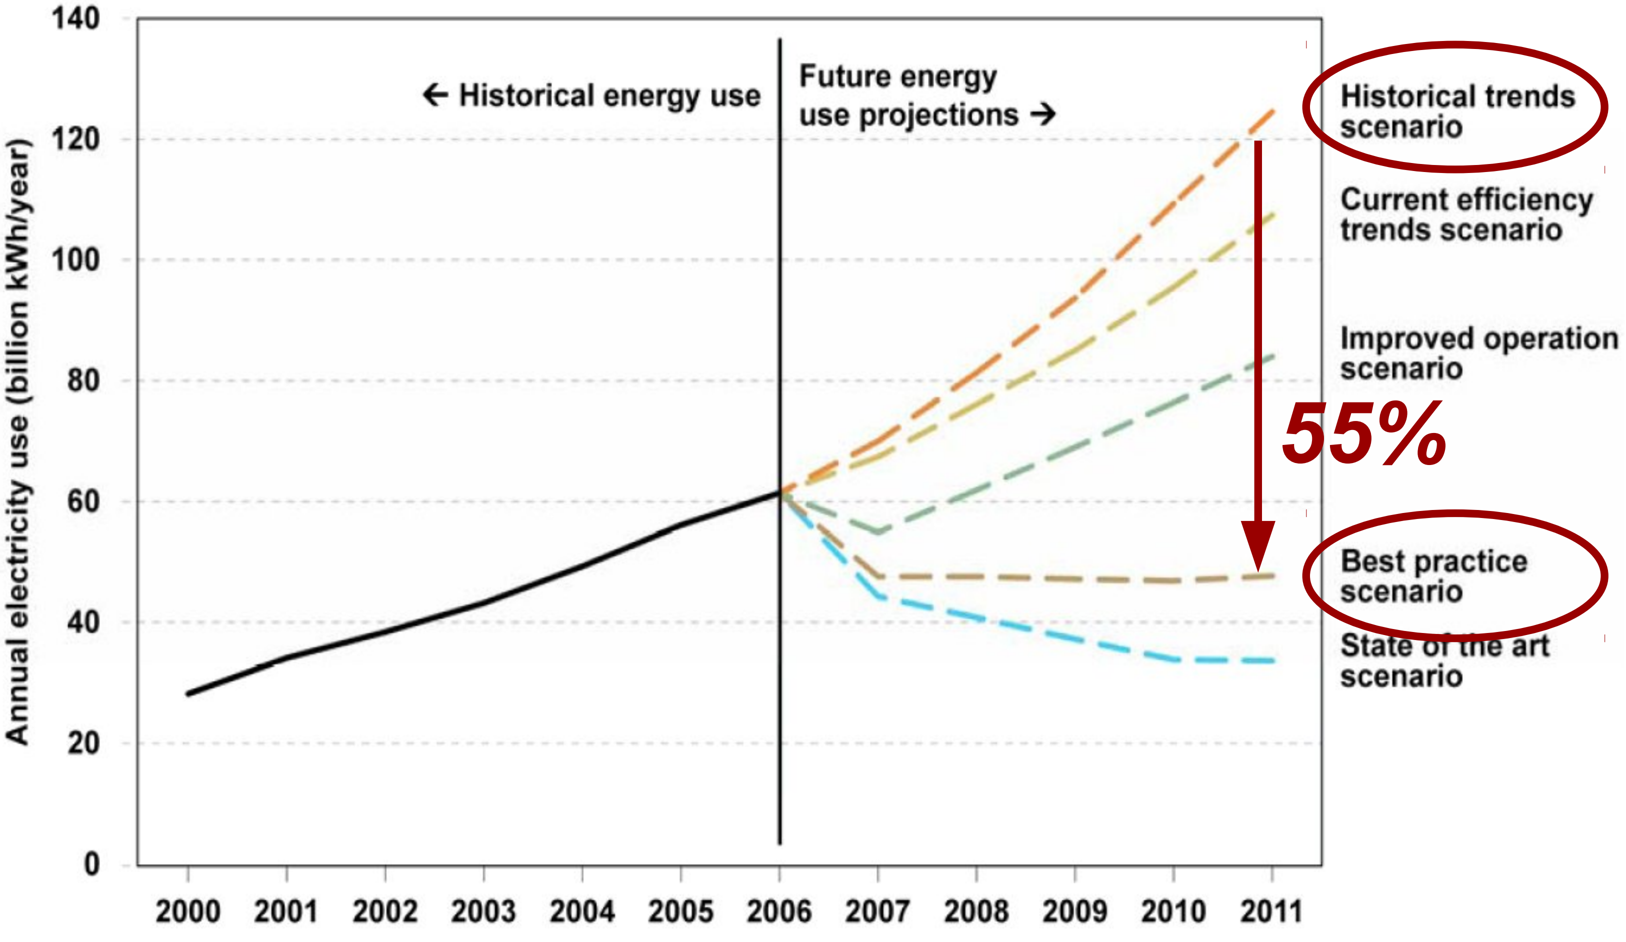

Environmental Protection Agency, Tech. Rep. 2007

Examples of data centers

- **Racks**
  - 42 (1U) servers are stacked in a rack
  - 1U server
    - 480mm x 800mm x 44m
    - Peak power ~380W
  - Rack peak power ~16KW
  - Large density of power consumption
    - 800W/m^2 – 1.6KW/m^2

- **Facebook's data center in North Carolina, US**
  - 450$ million project
  - ~28.000 m^2 (300.000 ft^2)
- **One of the data center in San Jose**
  - 2.300 m^2 (25.000 ft^2), 5 MW load

http://blog.cluttr.be/2011/01/26/the-case-for-data-center-power-efficiency/
http://www.datacenterknowledge.com/

Thermal constraints

- **Without proper cooling, chip temperatures would exceed safe operation limits**
  - *Almost complete transformation of electrical energy into heat*
- **Chip temperature generally unobservable**
- **Industrial approach**
  - Bound server *inlet air temperature*

Temperature distribution

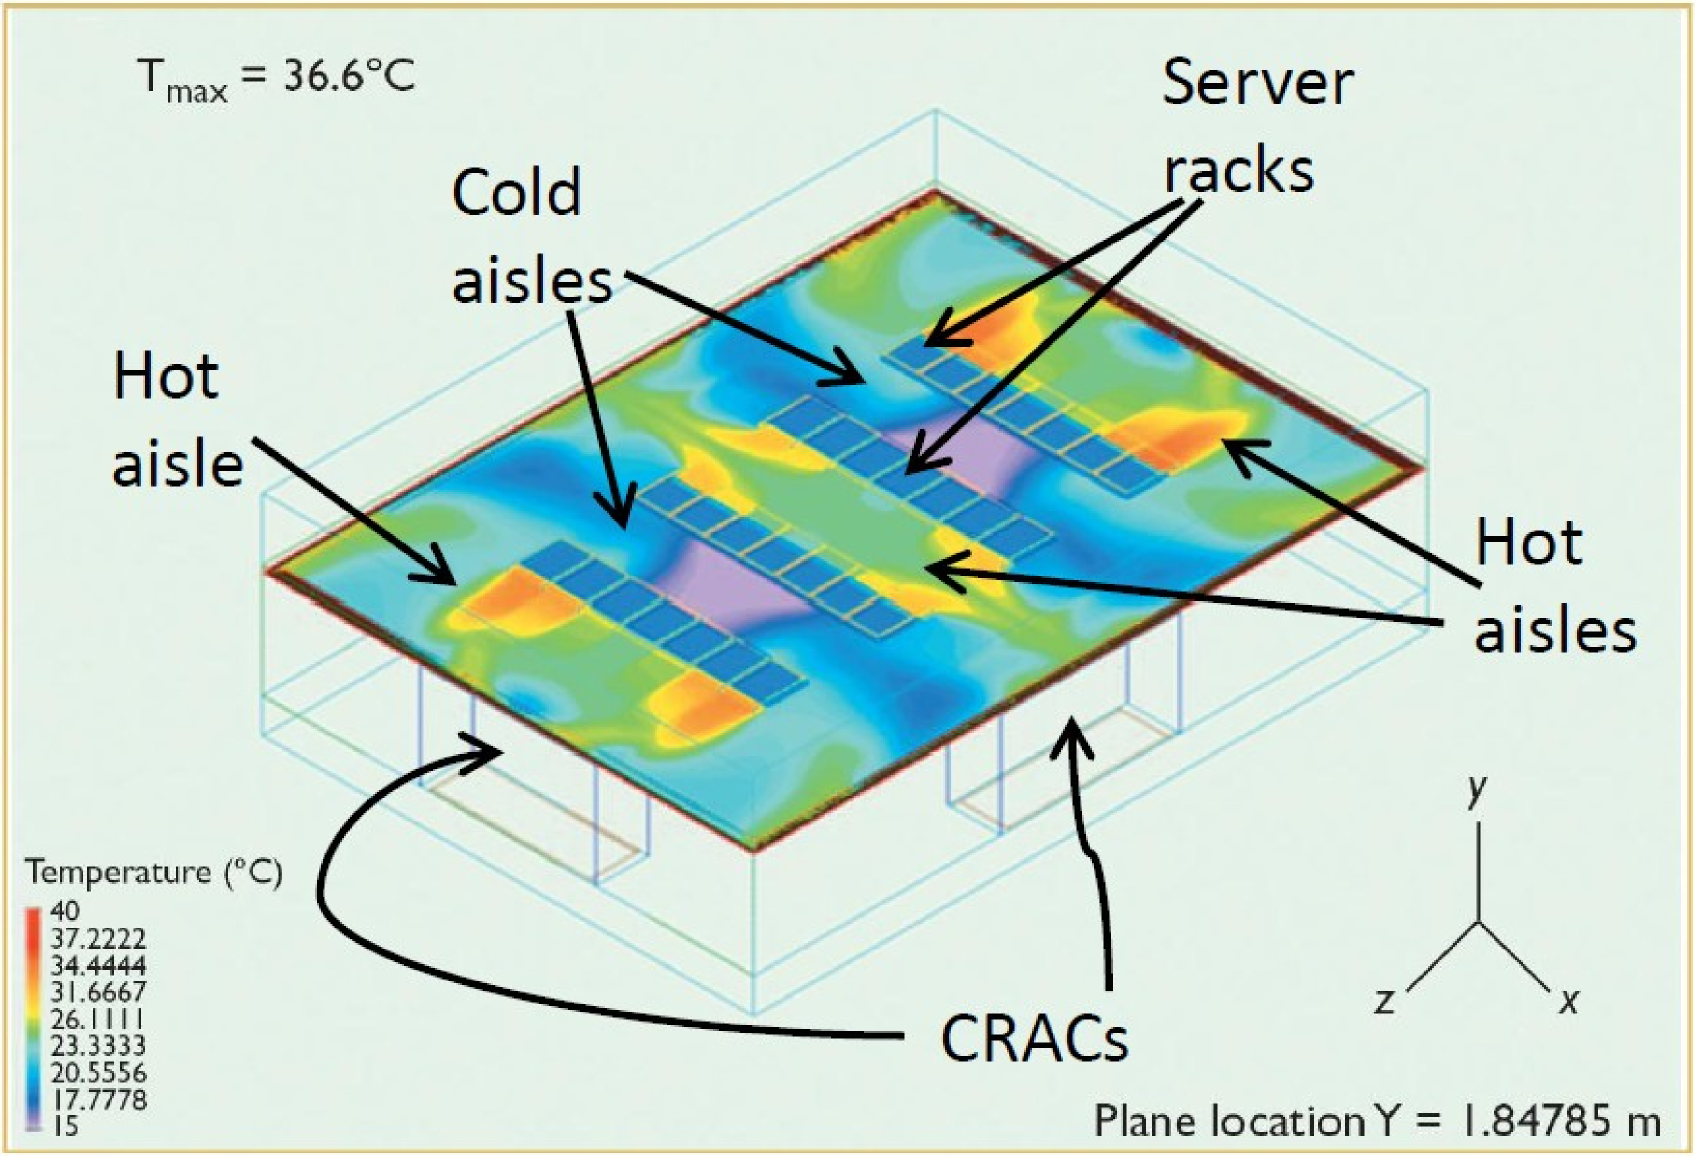

R. K. Sharma et al. IEEE Internet Computing 2005

Run-time cost

- **Difference between income due to the workload processing and the cost of powering the data center**
- **Depends on two service level agreements (SLAs)**
  - SLA_U: sets the income based on the quality of service (QoS)
    - Approximated by the ratio between required and assigned hardware resources

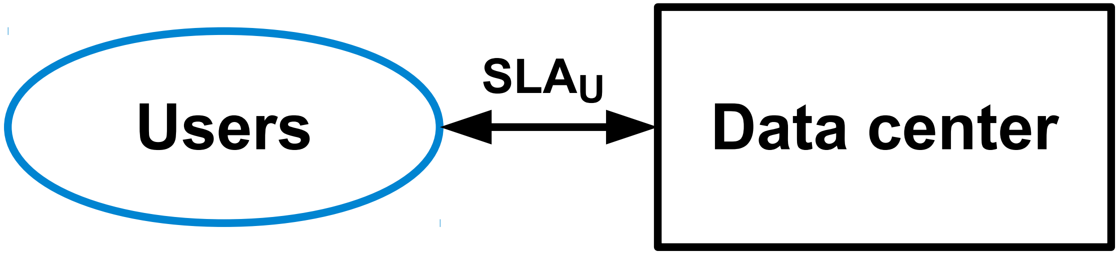

Run-time cost

- **Difference between income due to the workload processing and the cost of powering the data center**
- **Depends on two service level agreements (SLAs)**
  - SLA_U: sets the income based on the quality of service (QoS)
    - Approximated by the ratio between required and assigned hardware resources
  - SLA_G: sets the data center's powering cost
    - The energy cost is time-varying and power consumption dependent

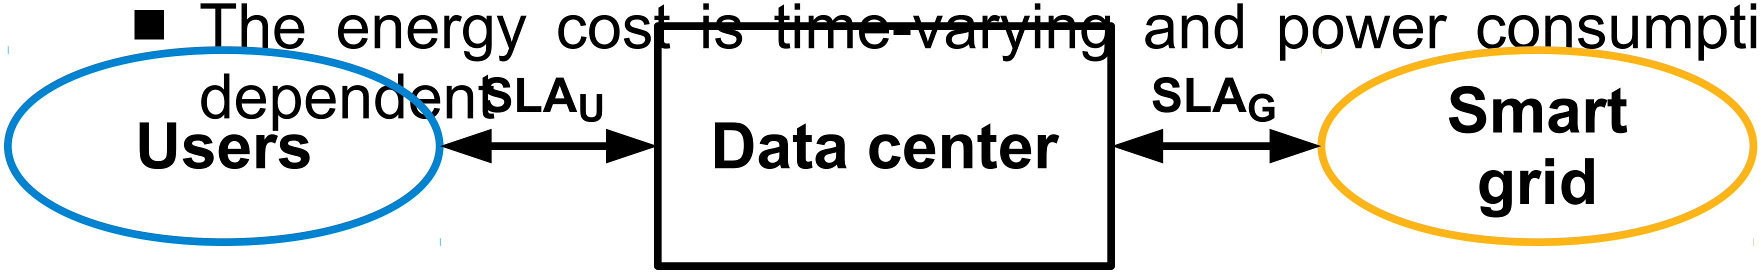

Users

SLA_U

Data center

SLA_G

Smart grid

Interaction with the smart grid

- **Time-varying electricity price**
  - Used by the smart-grid to cap the average power consumption of the data center

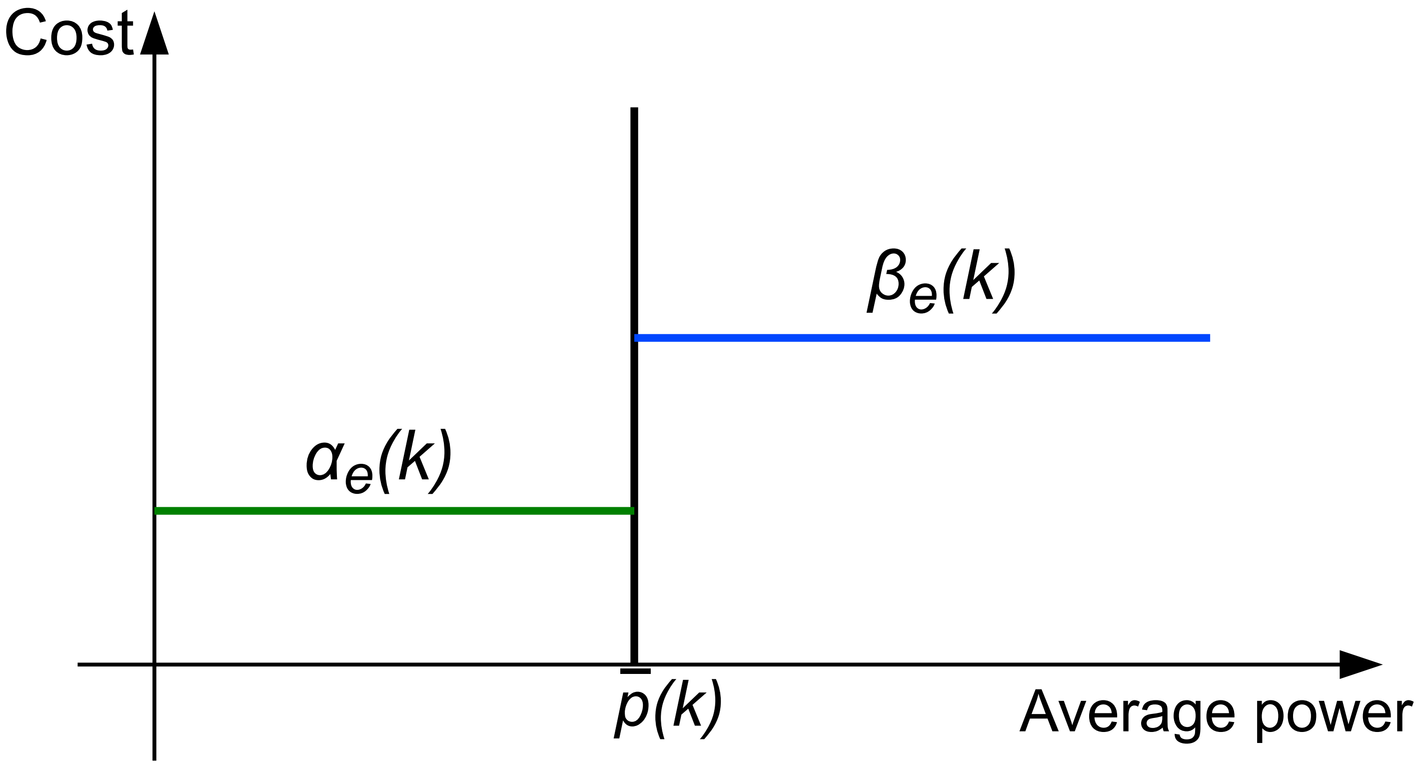

Outline

- Introduction
- **Control-oriented model**
- Simulation results
- Conclusion and future work

Control approach

- **Hierarchical-distributed control approach**
  - Takes advantage of the modularity found in data centers
- **We focus on the highest level of the hierarchy**
  - Based on a model predictive control (MPC) approach
  - Simplified modeling approach
    - Single servers are grouped into zones
    - Power consumption of a zone is proportional to the amount of workload executed in the zone
  - Considers computational and thermal dynamics
  - Considers the nonlinear efficiency of the CRAC units

Computational model

- **Computational resources modeled via a network of queues**

jobs in every zone

arrivals in the interval

scheduling action

departures in the interval

$$\hat{\boldsymbol{l}}(\nu+1|k) = \hat{\tilde{\boldsymbol{l}}}(\nu|k) + \hat{\mathrm{A}}(\nu|k)\hat{\boldsymbol{s}}(\nu|k) - \hat{\boldsymbol{d}}(\nu|k)$$

$$\hat{\boldsymbol{d}}(\nu|k) = \min\left\{\hat{\tilde{\boldsymbol{l}}}(\nu|k) + \hat{\mathrm{A}}(\nu|k)\hat{\boldsymbol{s}}(\nu|k),\ M\hat{\boldsymbol{\rho}}(\nu|k)\right\}$$

allocated resources

- **Power consumption proportional to the amount of jobs processed**

$$\hat{\boldsymbol{p}}_{\mathcal{N}}(\nu|k) = \mathsf{B}_d\hat{\boldsymbol{d}}(\nu|k)$$

Thermal network

- **Linear model**

Output temperature of zones and CRAC units

$$\hat{\boldsymbol{T}}_{\text{out}}(\nu+1|k) = A_{T,D}\hat{\boldsymbol{T}}_{\text{out}}(\nu|k) + B_{T,D}\left[\hat{\mathbf{p}}_{\mathcal{N}}(\nu|k)^T\ \hat{\boldsymbol{T}}_{\text{ref}}(\nu|k)^T\right]^T$$

Power consumption of zones

Reference temperature of CRAC units

$$\hat{\boldsymbol{T}}_{\text{in}}(\nu|k) = \Psi\hat{\boldsymbol{T}}_{\text{out}}(\nu|k) \leq \overline{\boldsymbol{T}_{\text{in}}}$$

- **CRAC power consumption depends on the coefficient of performance (COP)**

$$\mathsf{p}_i(t) = \frac{\dot{Q}_i(t)}{COP_i(T_{\text{out},i}(t))}$$

Heat removed rate (W)

Computational variables at time *k*

			Variables
Input	Controllable	Job scheduling	$\boldsymbol{s}(k)$
		Resource allocation	$\boldsymbol{\rho}(k)$
	Uncontrollable	Job arrival	$\boldsymbol{a}(k)$
Ouput	Zone power consumption		$\boldsymbol{p}_{\mathbf{N}}(k)$
State	Number of jobs in zones		$\boldsymbol{l}(k)$

Thermal network variables at time *k*

			Variables
Input	Controllable	CRAC unit reference temperature	$\boldsymbol{T}_{ref}(k)$
	Uncontrollable	Zone power consumption	$\boldsymbol{p}_{N}(k)$
Ouput	Power consumption of CRAC nodes		$\boldsymbol{p}_{C}(k)$
	Input temperatures of zones		$\boldsymbol{T}_{in}(k)$
State	Output temperatures of CRACs and zones		$\boldsymbol{T}_{out}(k)$

Outline

- Introduction
- Control-oriented model
- **Simulation results**
- Conclusion and future work

Simulation parameters

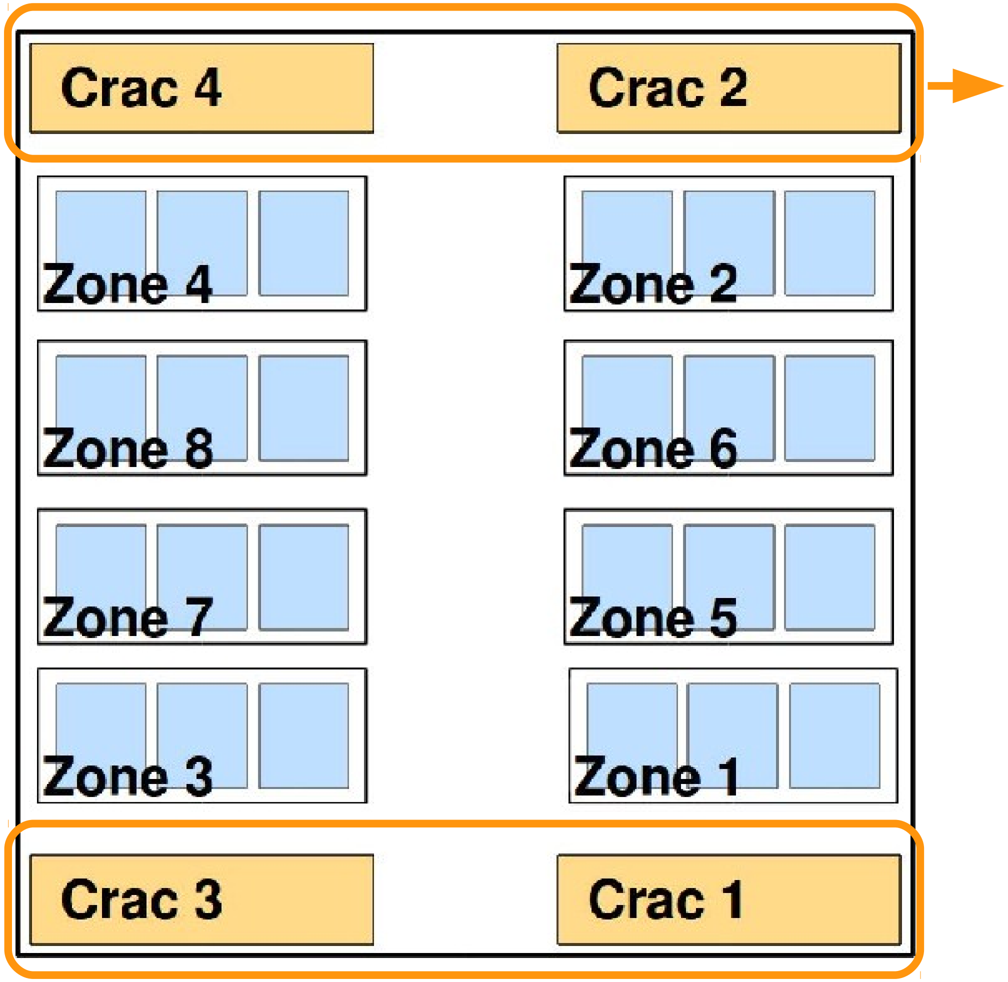

- **4 CRAC units**
  - Identical each other

Simulation parameters

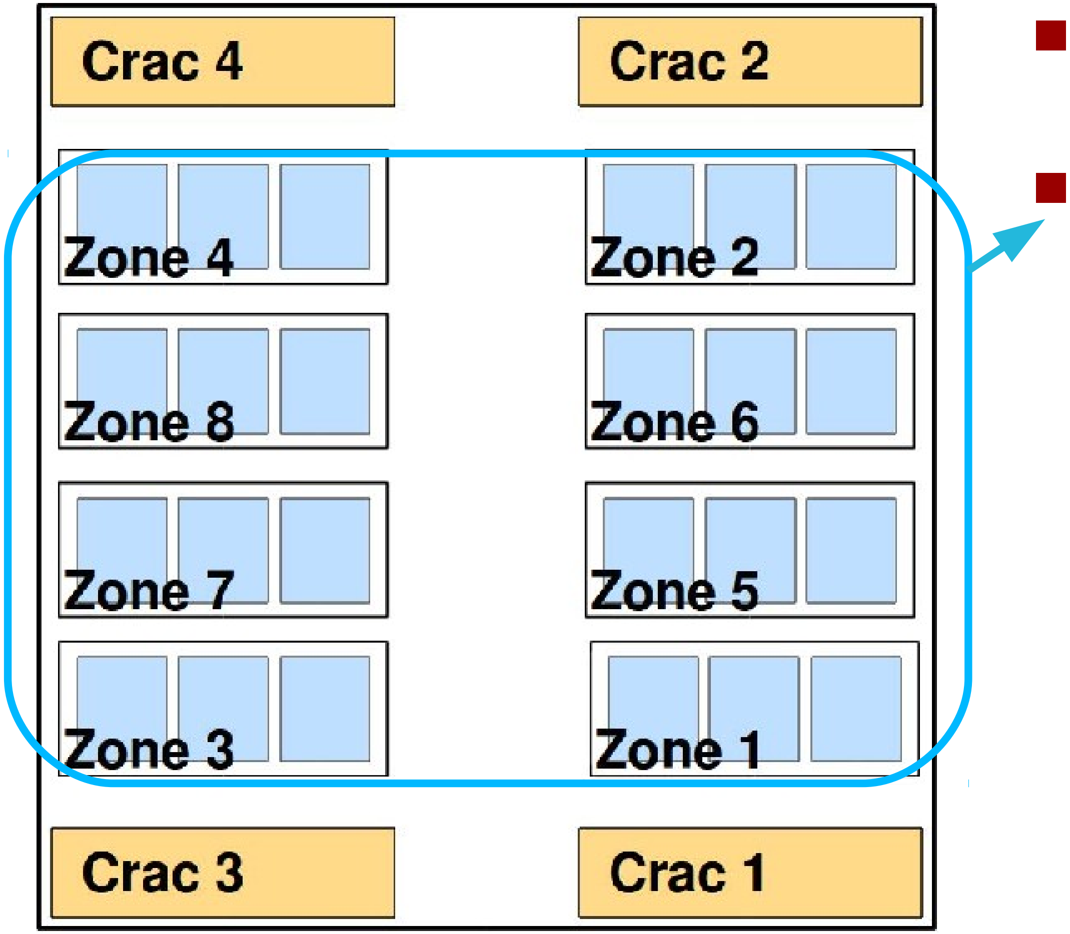

- **4 CRAC units**
  - Identical each other
- **8 Zones**
  - 3 Racks each (126 servers per zone)

Simulation parameters

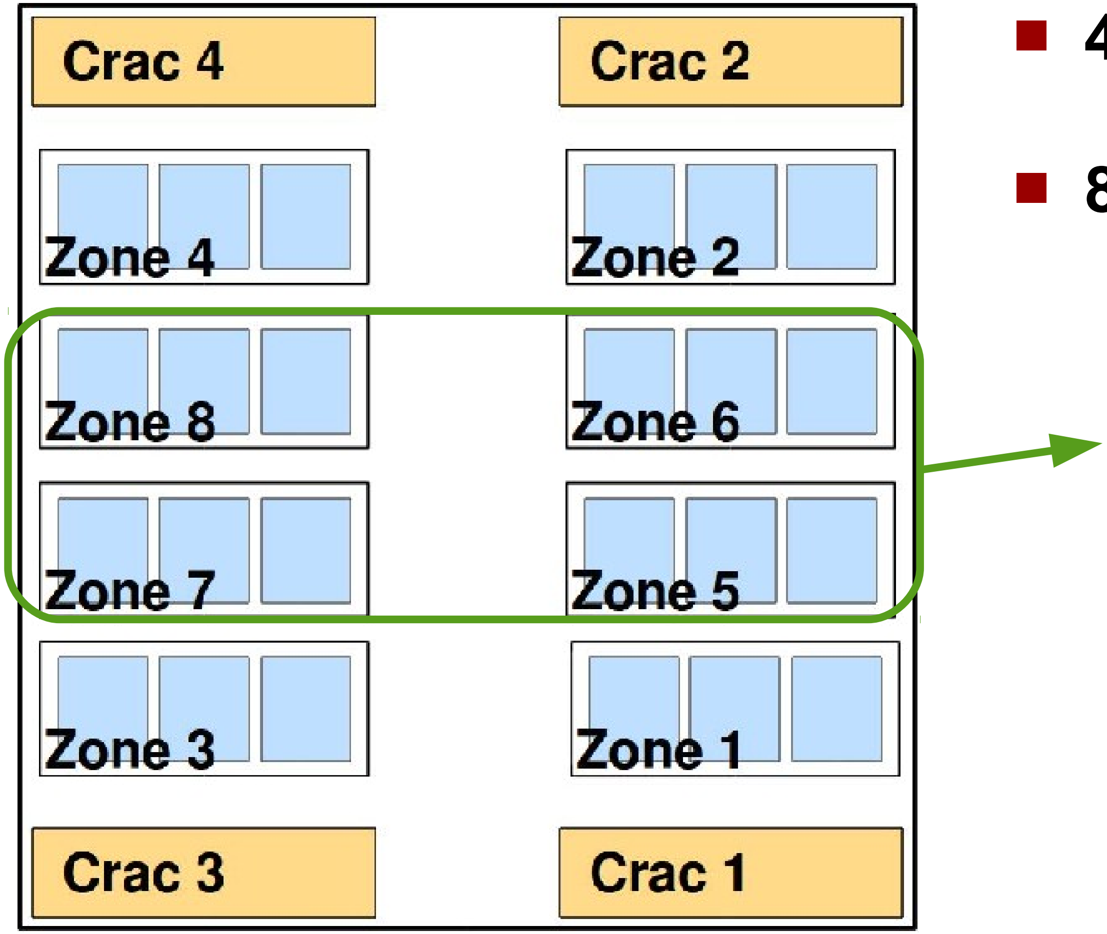

- **4 CRAC units**
  - Identical each other
- **8 Zones**
  - 3 Racks each (126 servers per zone)
  - Energy efficient servers

Simulation parameters

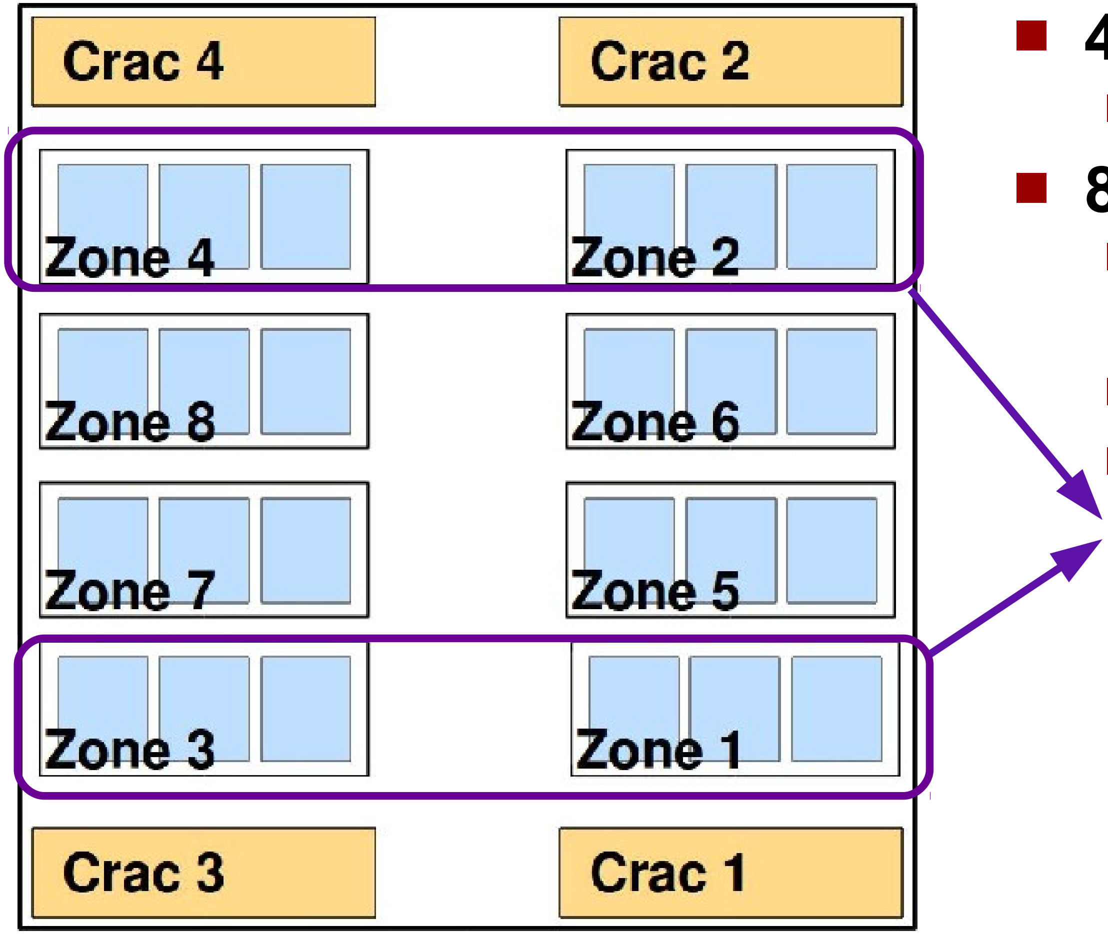

- **4 CRAC units**
  - Identical each other
- **8 Zones**
  - 3 Racks each (126 servers per zone)
  - Energy efficient servers
  - Efficiently cooled servers

Simulation parameters

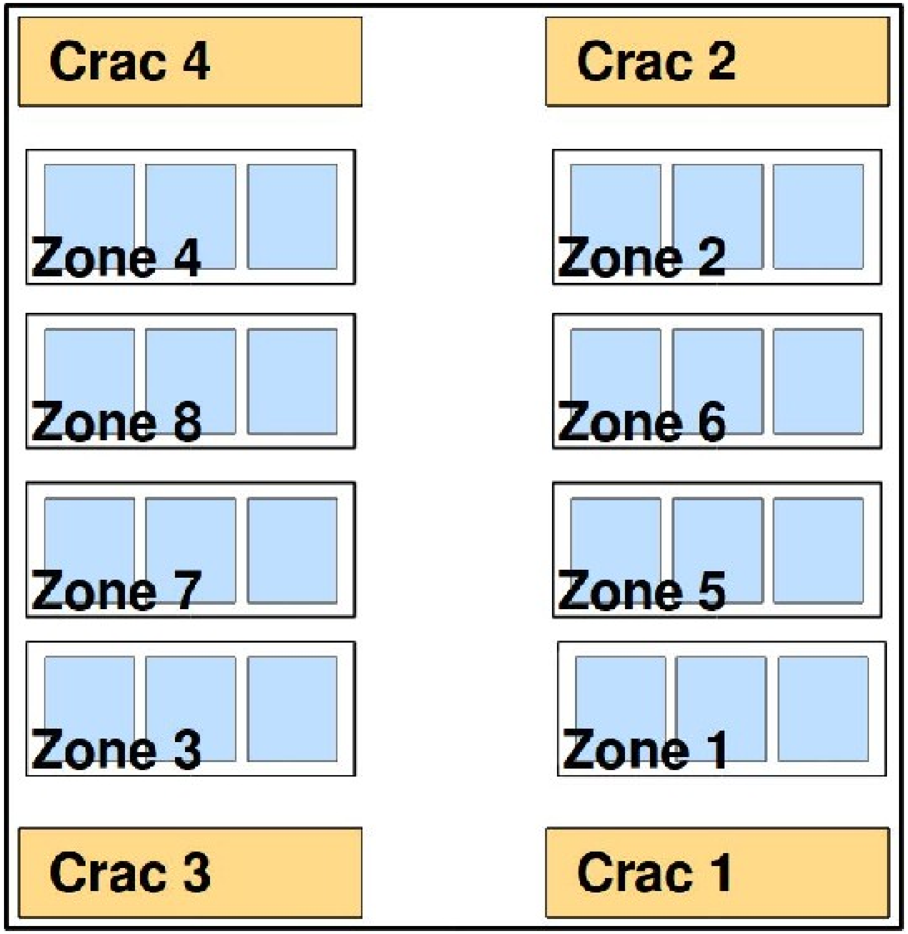

- **4 CRAC units**
  - Identical each other
- **8 Zones**
  - 3 Racks each (126 servers per zone)
  - Energy efficient servers
  - Efficiently cooled servers
- **Simulate data center over 24hr time**
  - Two control approaches
  - Controllers' time step: 10 min
  - Horizon: 6 steps (1hr)

Uncoordinated & Coordinated strategies

- **Uncoordinated strategy**
  - Two-step control approach

1. $$\min_{\mathcal{R},\mathcal{S}} F_1(\mathcal{R},\mathcal{S};\, \boldsymbol{l}(k))$$
$$\text{subject to}$$
$$\qquad \text{computational dynamics}$$

→

2. $$\min_{\mathcal{T}_{\text{ref}}} F_2(\mathcal{T}_{\text{ref}};\, \boldsymbol{l}(k), \boldsymbol{T}_{\text{out}}(k))$$
$$\text{subject to}$$
$$\qquad \text{thermal dynamics}$$
$$\qquad \hat{\boldsymbol{T}}_{\text{in}}(\nu|k) \le \overline{\boldsymbol{T}_{\text{in}}}$$

- **Coordinated strategy**

$$\min_{\mathcal{R},\mathcal{S},\mathcal{T}_{\text{ref}}} F(\mathcal{R},\mathcal{S},\mathcal{T}_{\text{ref}};\, \boldsymbol{l}(k), \boldsymbol{T}_{\text{out}}(k))$$
$$\text{subject to}$$
$$\qquad \text{computational dynamics}$$
$$\qquad \text{thermal dynamics}$$
$$\qquad \hat{\boldsymbol{p}}_{\mathcal{N}}(\nu|k) = \mathrm{B}_d \hat{\boldsymbol{d}}(\nu|k)$$
$$\qquad \hat{\boldsymbol{T}}_{\text{in}}(\nu|k) \le \overline{\boldsymbol{T}_{\text{in}}}$$

- $\mathcal{R} = \{\hat{\boldsymbol{\rho}}(k|k), \ldots, \hat{\boldsymbol{\rho}}(k+T-1|k)\}$
- $\mathcal{S} = \{\hat{\boldsymbol{s}}(k|k), \ldots, \hat{\boldsymbol{s}}(k+T-1|k)\}$
- $\mathcal{T}_{\text{ref}} = \{\hat{\boldsymbol{T}}_{ref}(k|k), \ldots, \hat{\boldsymbol{T}}_{ref}(k+T-1|k)\}$

Job arrival rate and electricity costs

- **Real data**
  - Job arrival rate

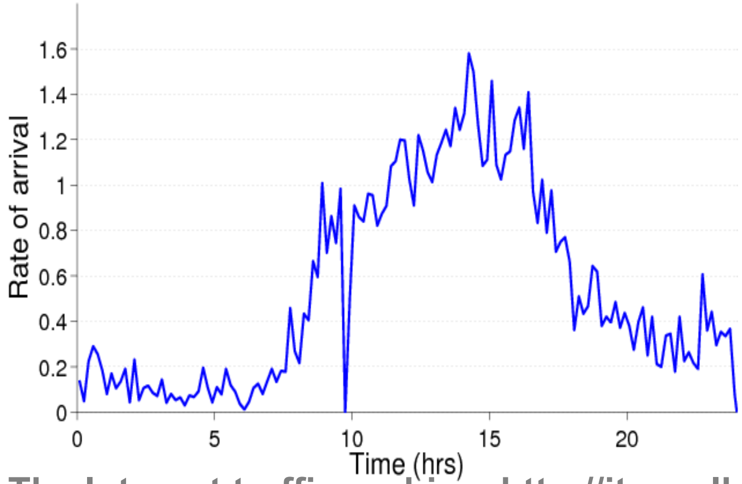

The Internet traffic archive, http://ita.ee.lbl.gov/

  - Day-ahead electricity cost

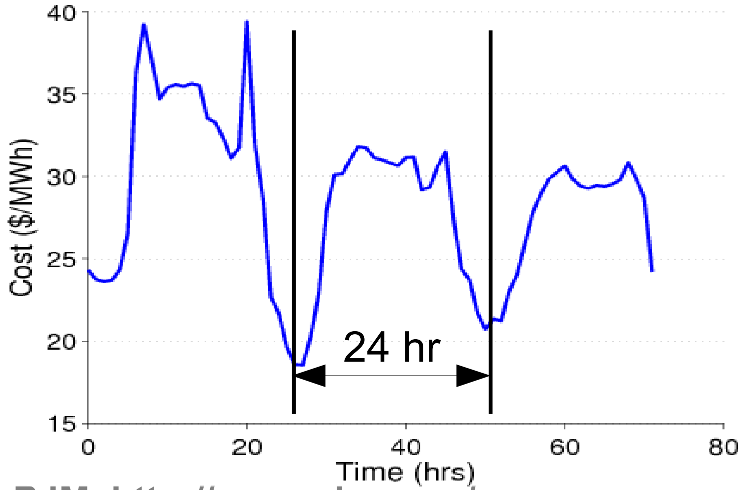

PJM, http://www.pjm.com/

- **Parameters in the simulation**
  - Job arrival rate, 2 classes of jobs

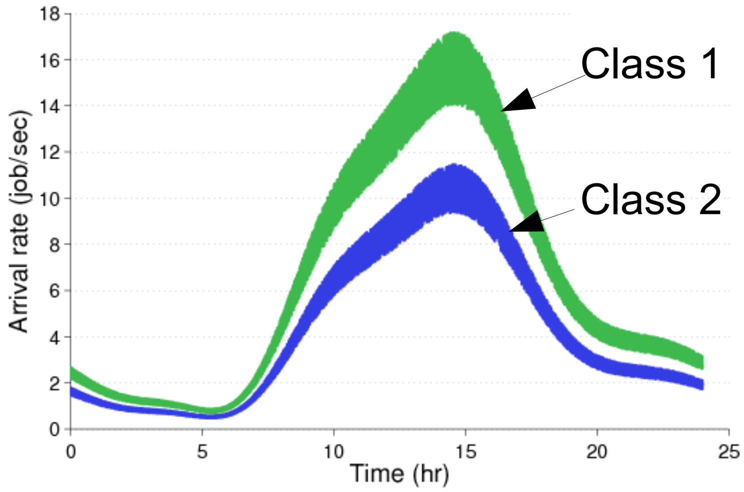

  - Electricity costs

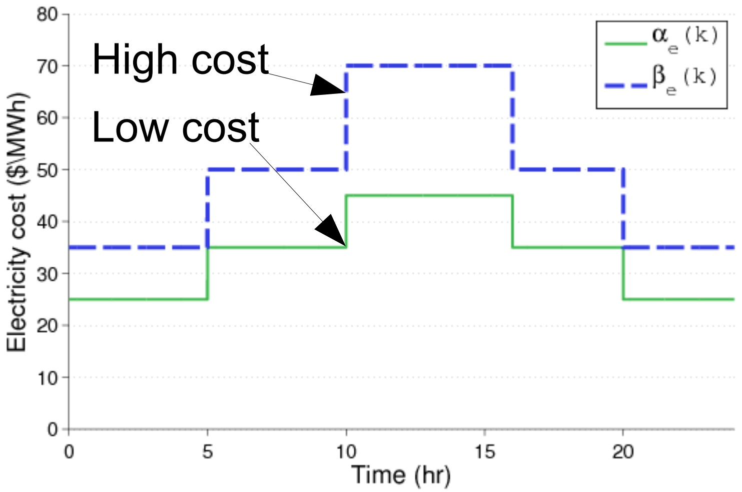

Total data center power consumption

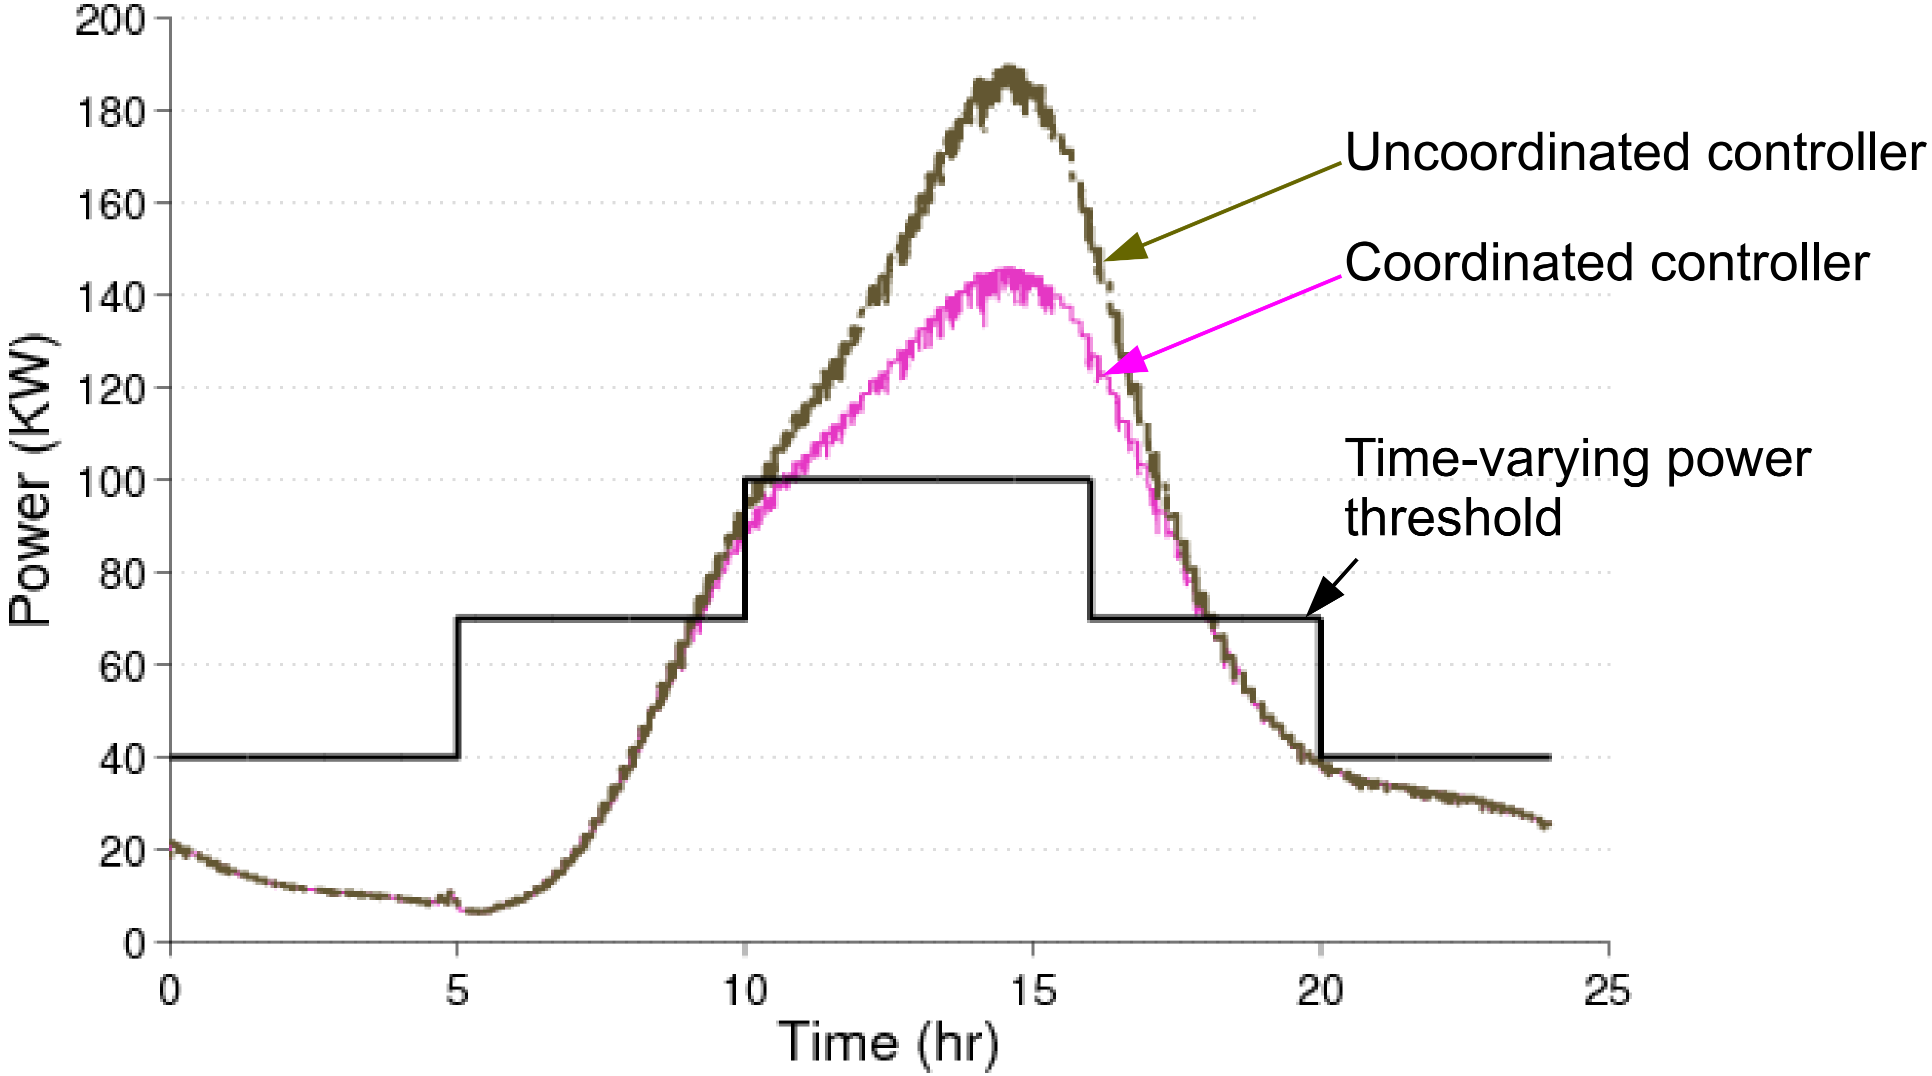

Total data center power consumption

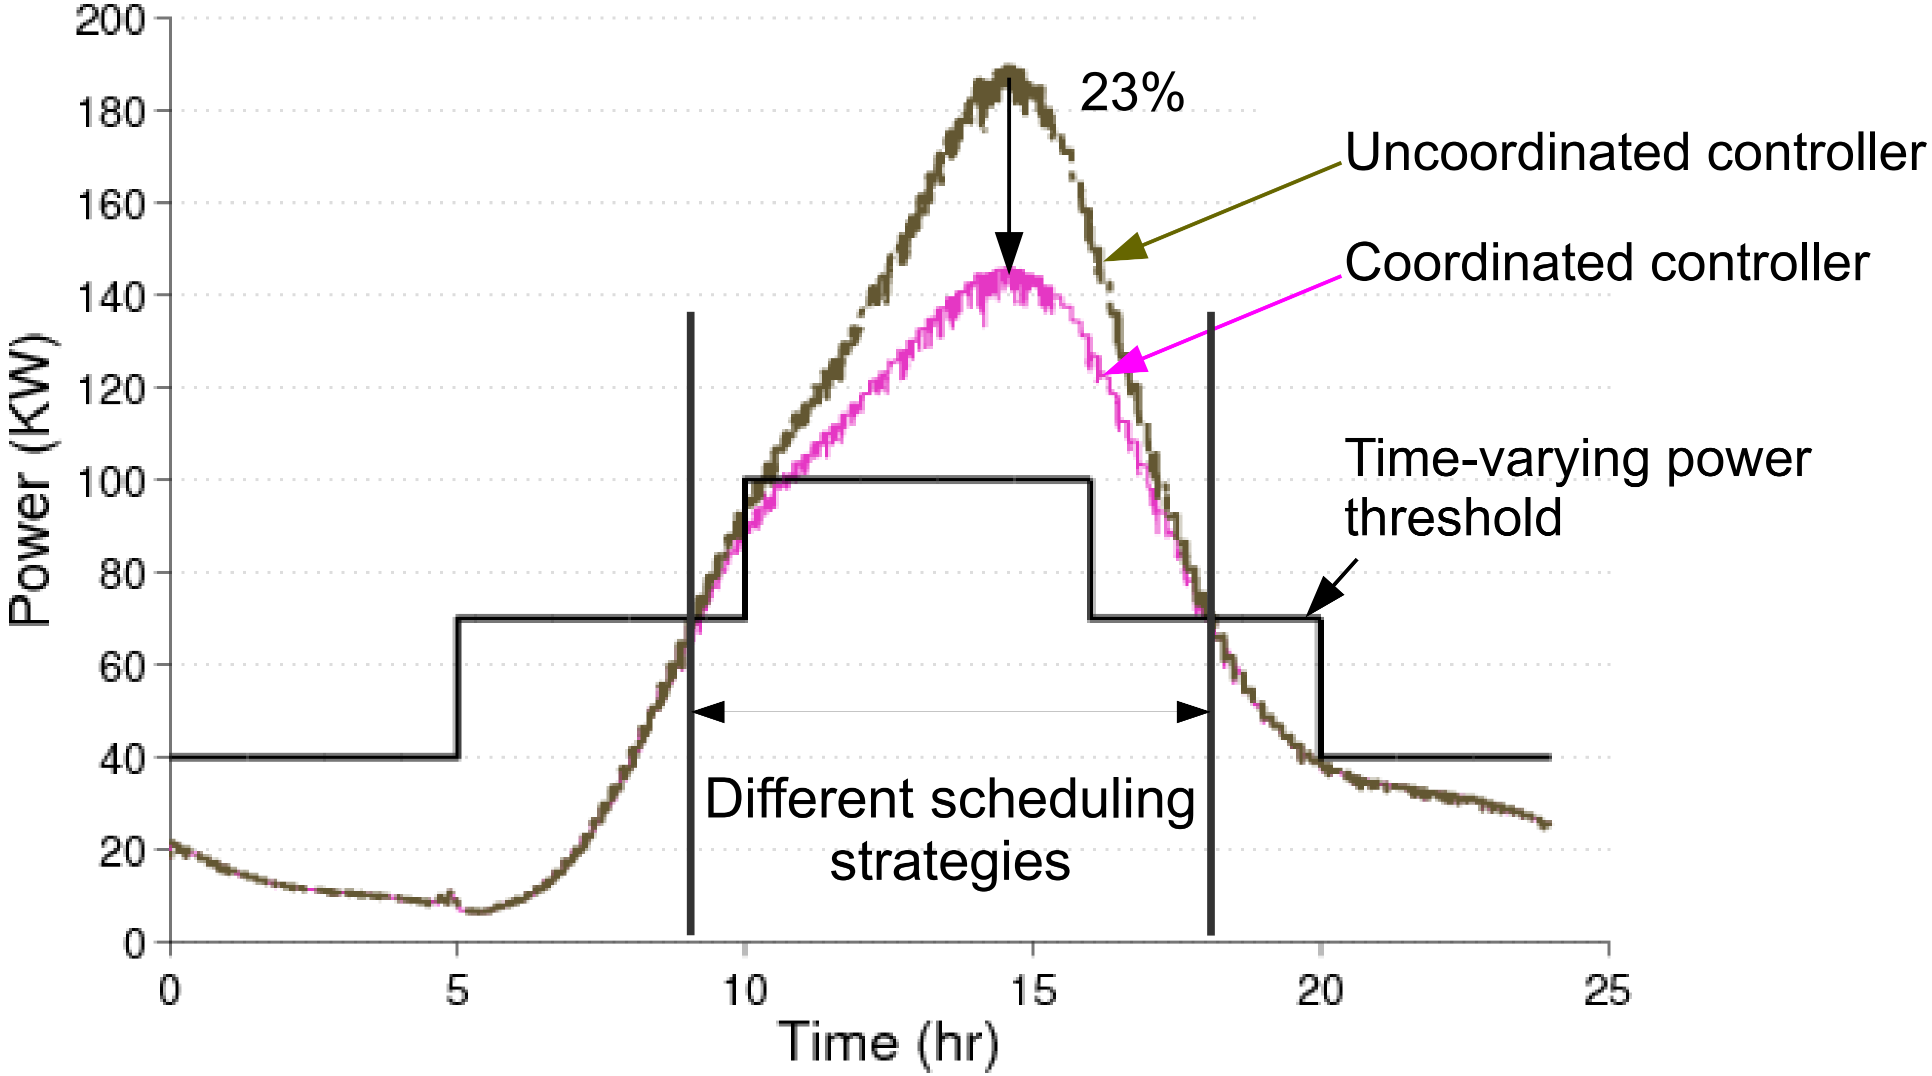

23%

Uncoordinated controller

Coordinated controller

Time-varying power threshold

Different scheduling strategies

Average reference temperatures

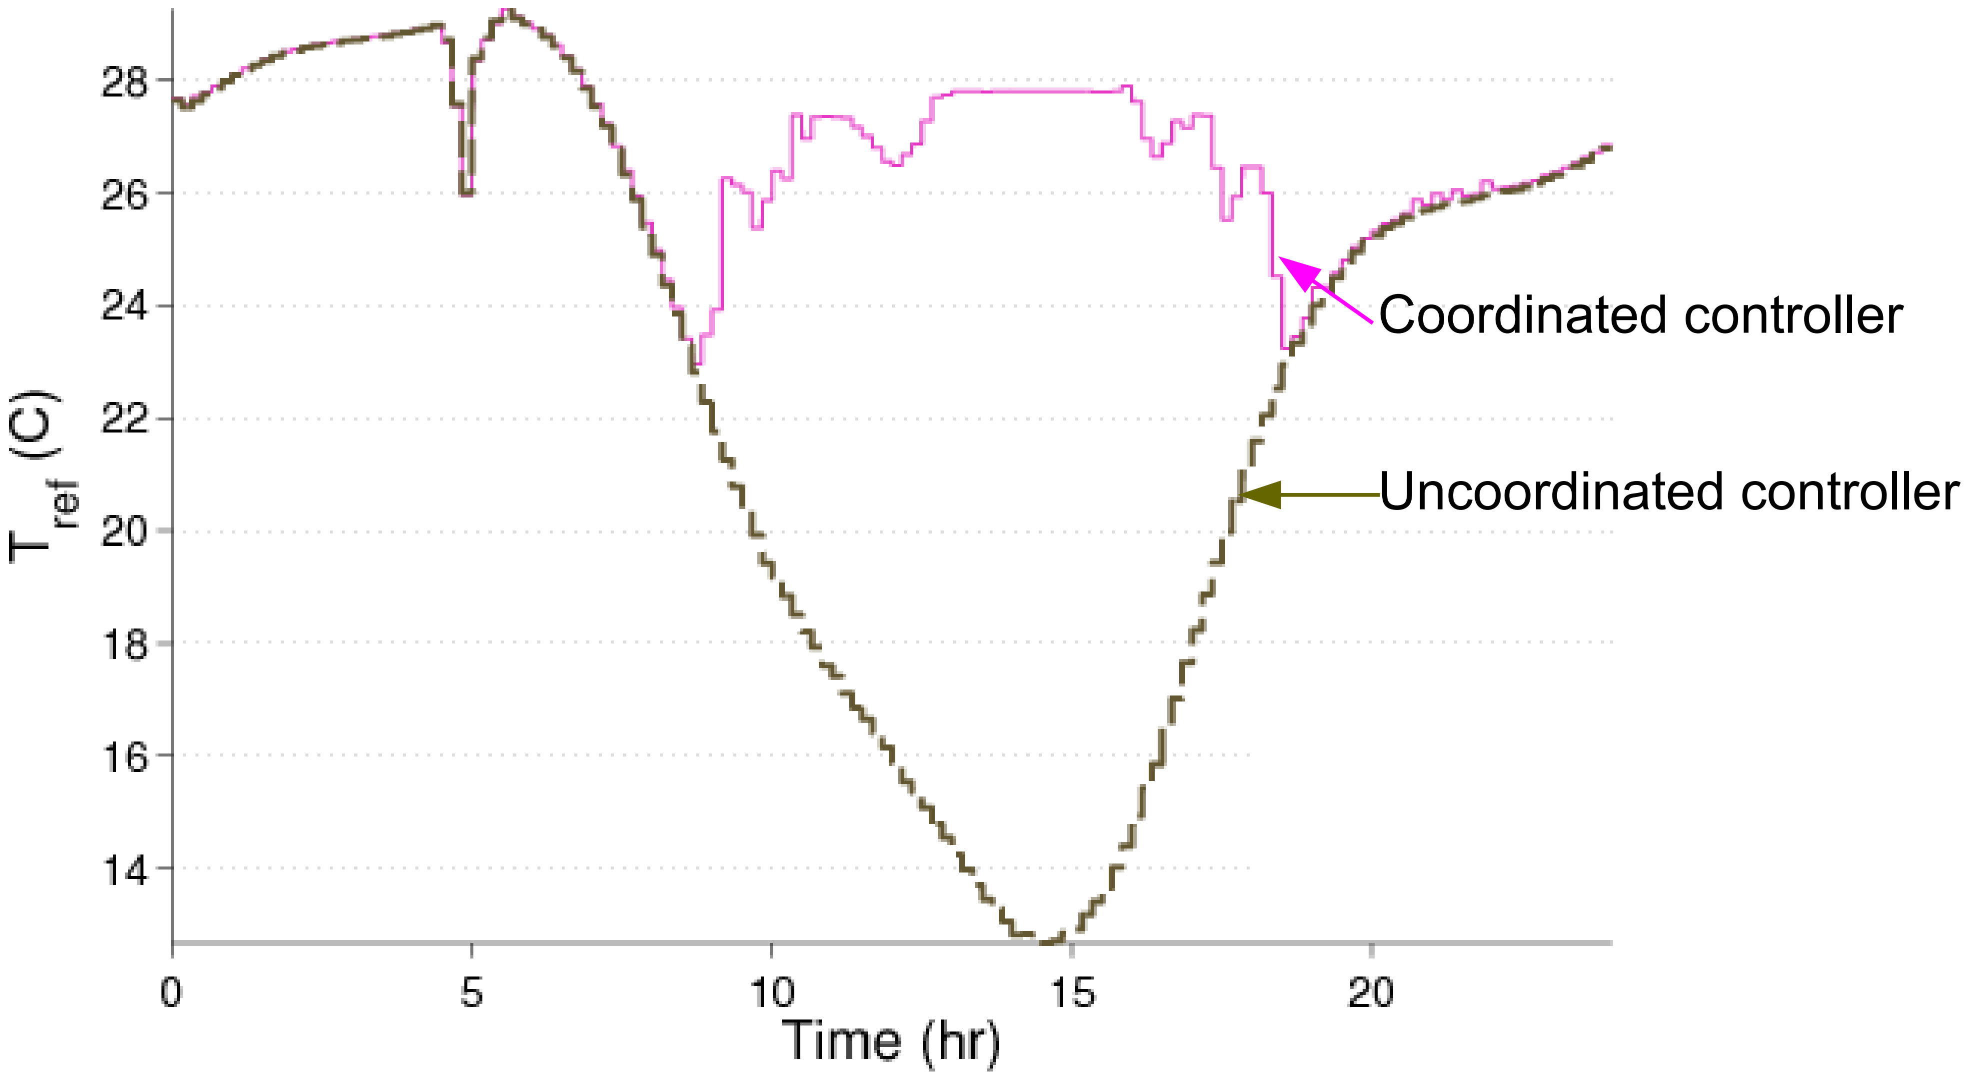

Average reference temperatures

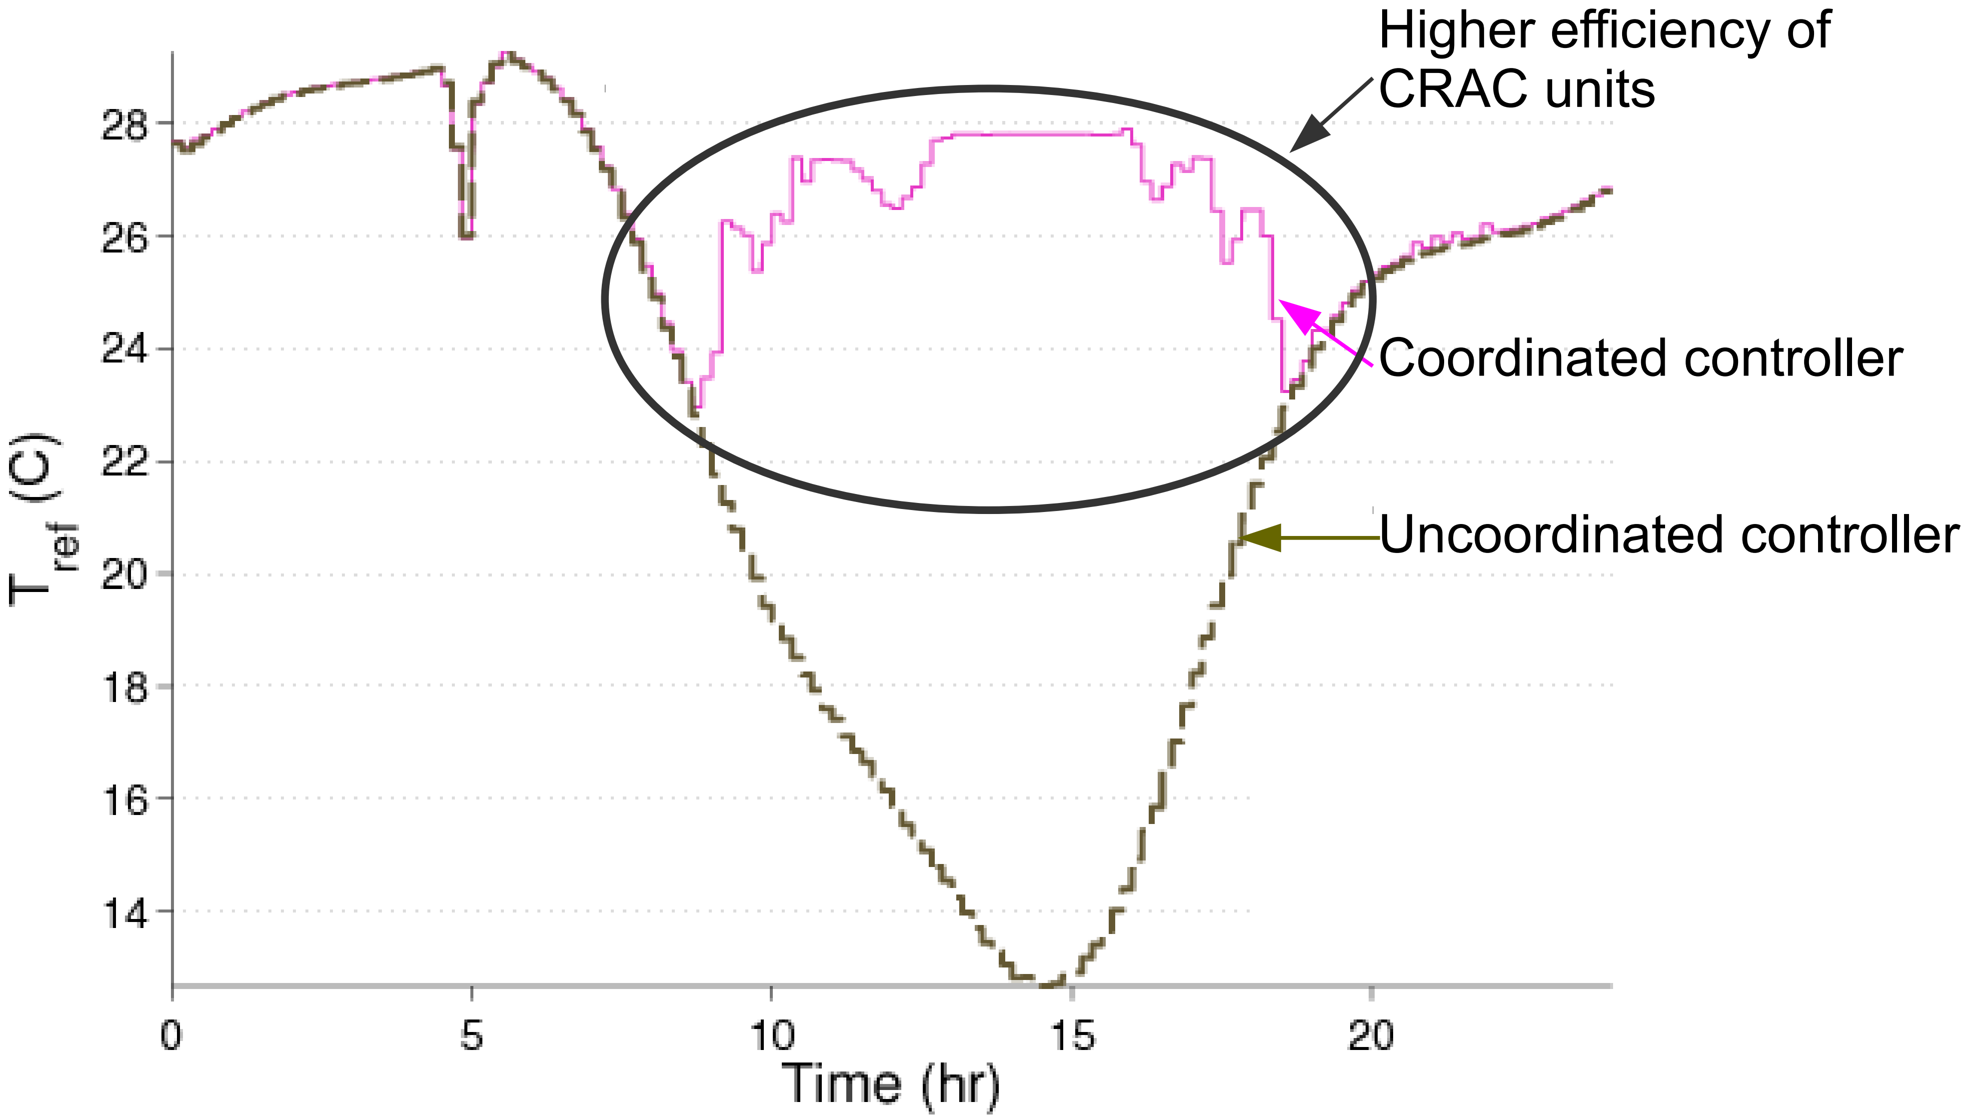

Higher efficiency of CRAC units

Coordinated controller

Uncoordinated controller

Average reference temperatures

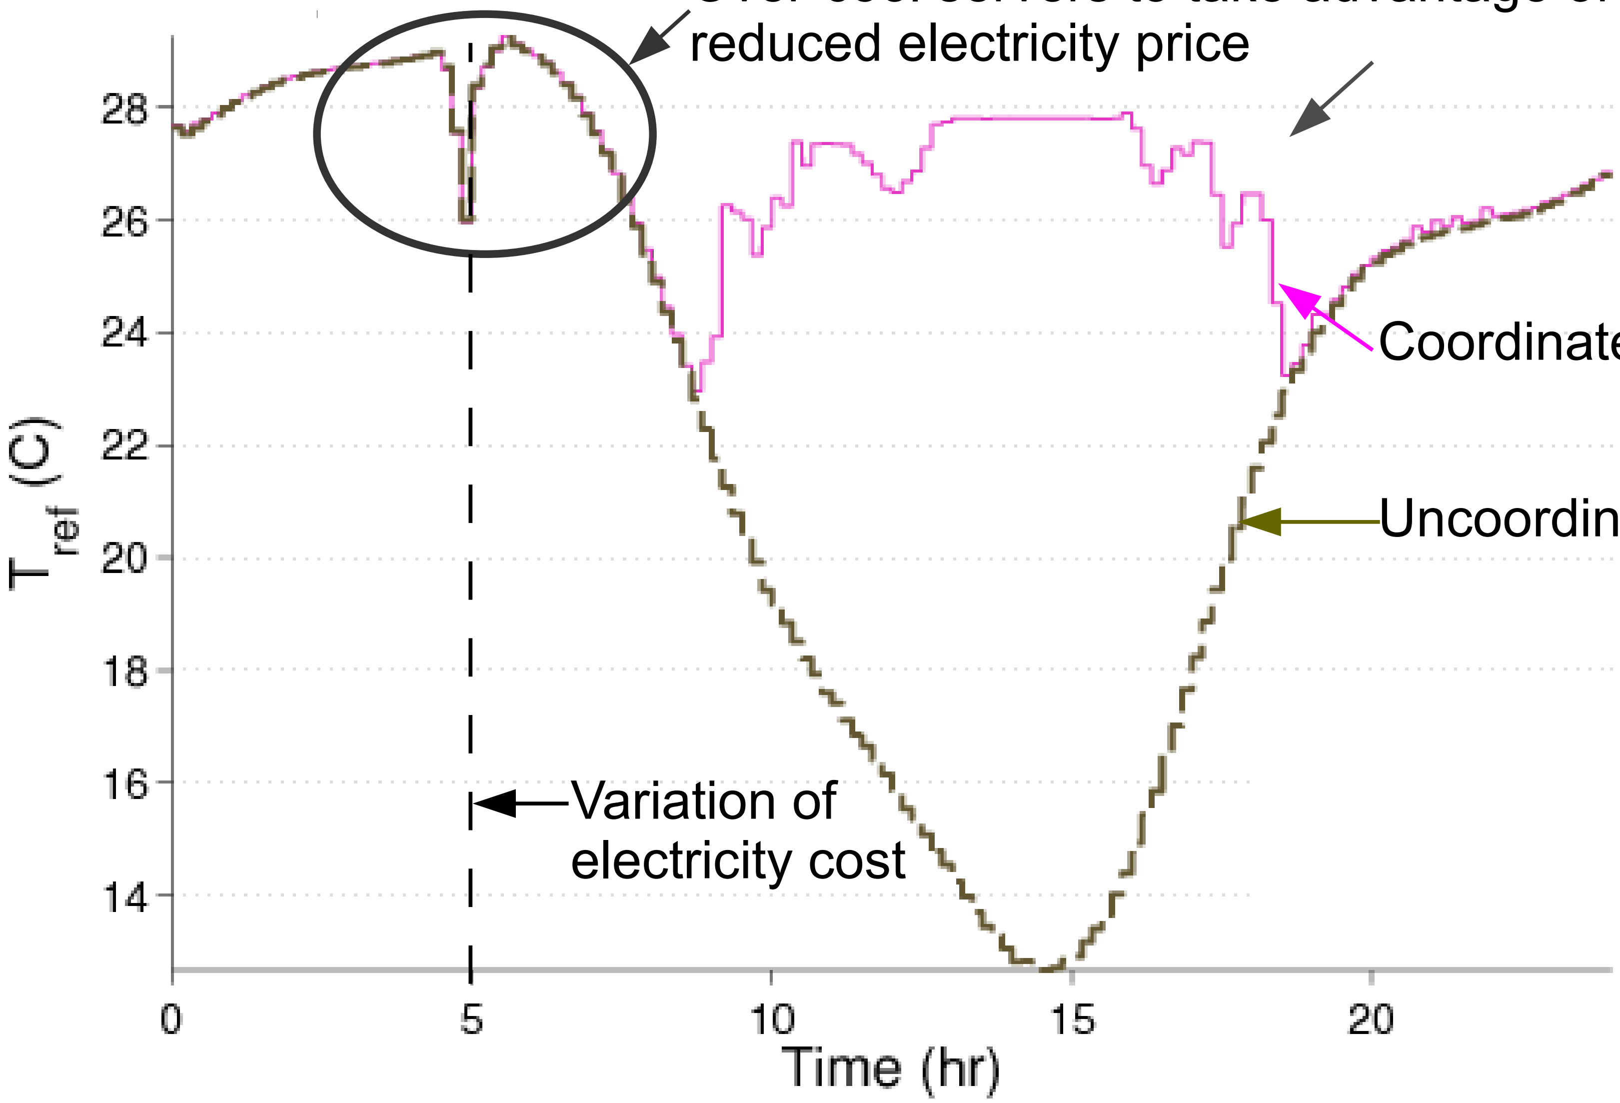

Over-cool servers to take advantage of reduced electricity price

Coordinated controller

Uncoordinated controller

Variation of electricity cost

Simulation parameters

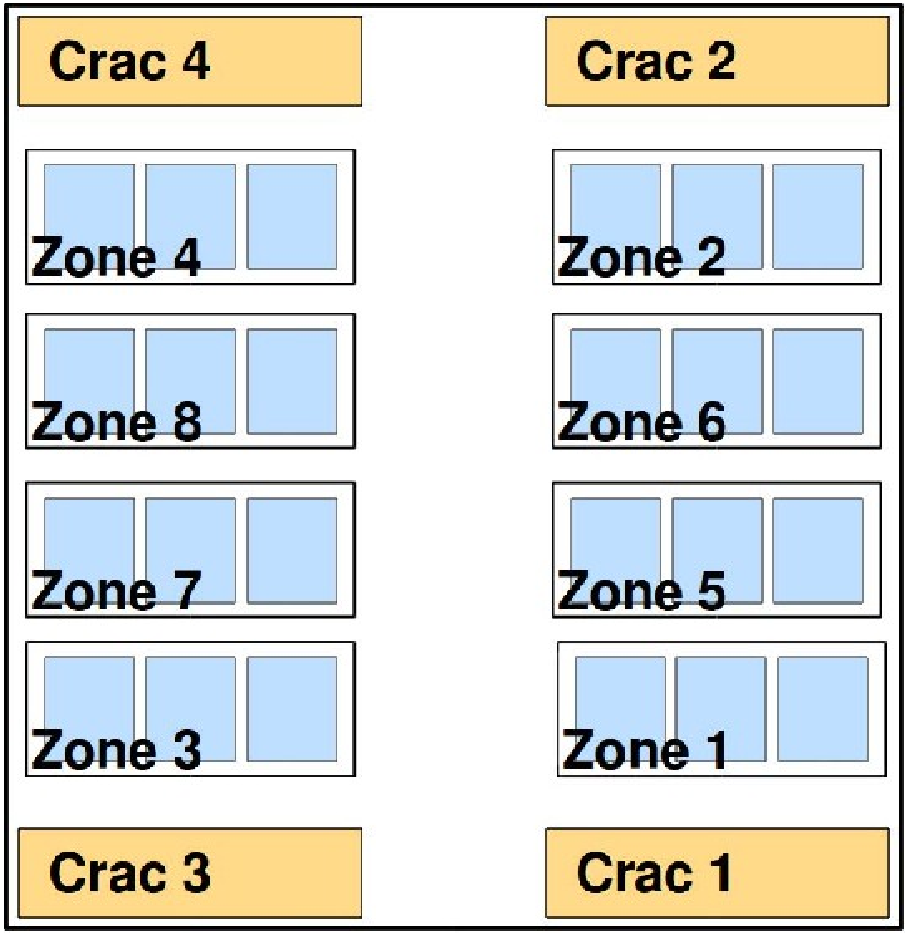

- **4 CRAC units**
  - Identical each other
- **8 Zones**
  - 3 Racks each (126 servers per zone)
  - Energy efficient servers
  - Efficiently cooled servers
- **Simulate data center over 24hr time**
  - *Compare two coordinated control strategies*

Total data center power consumption

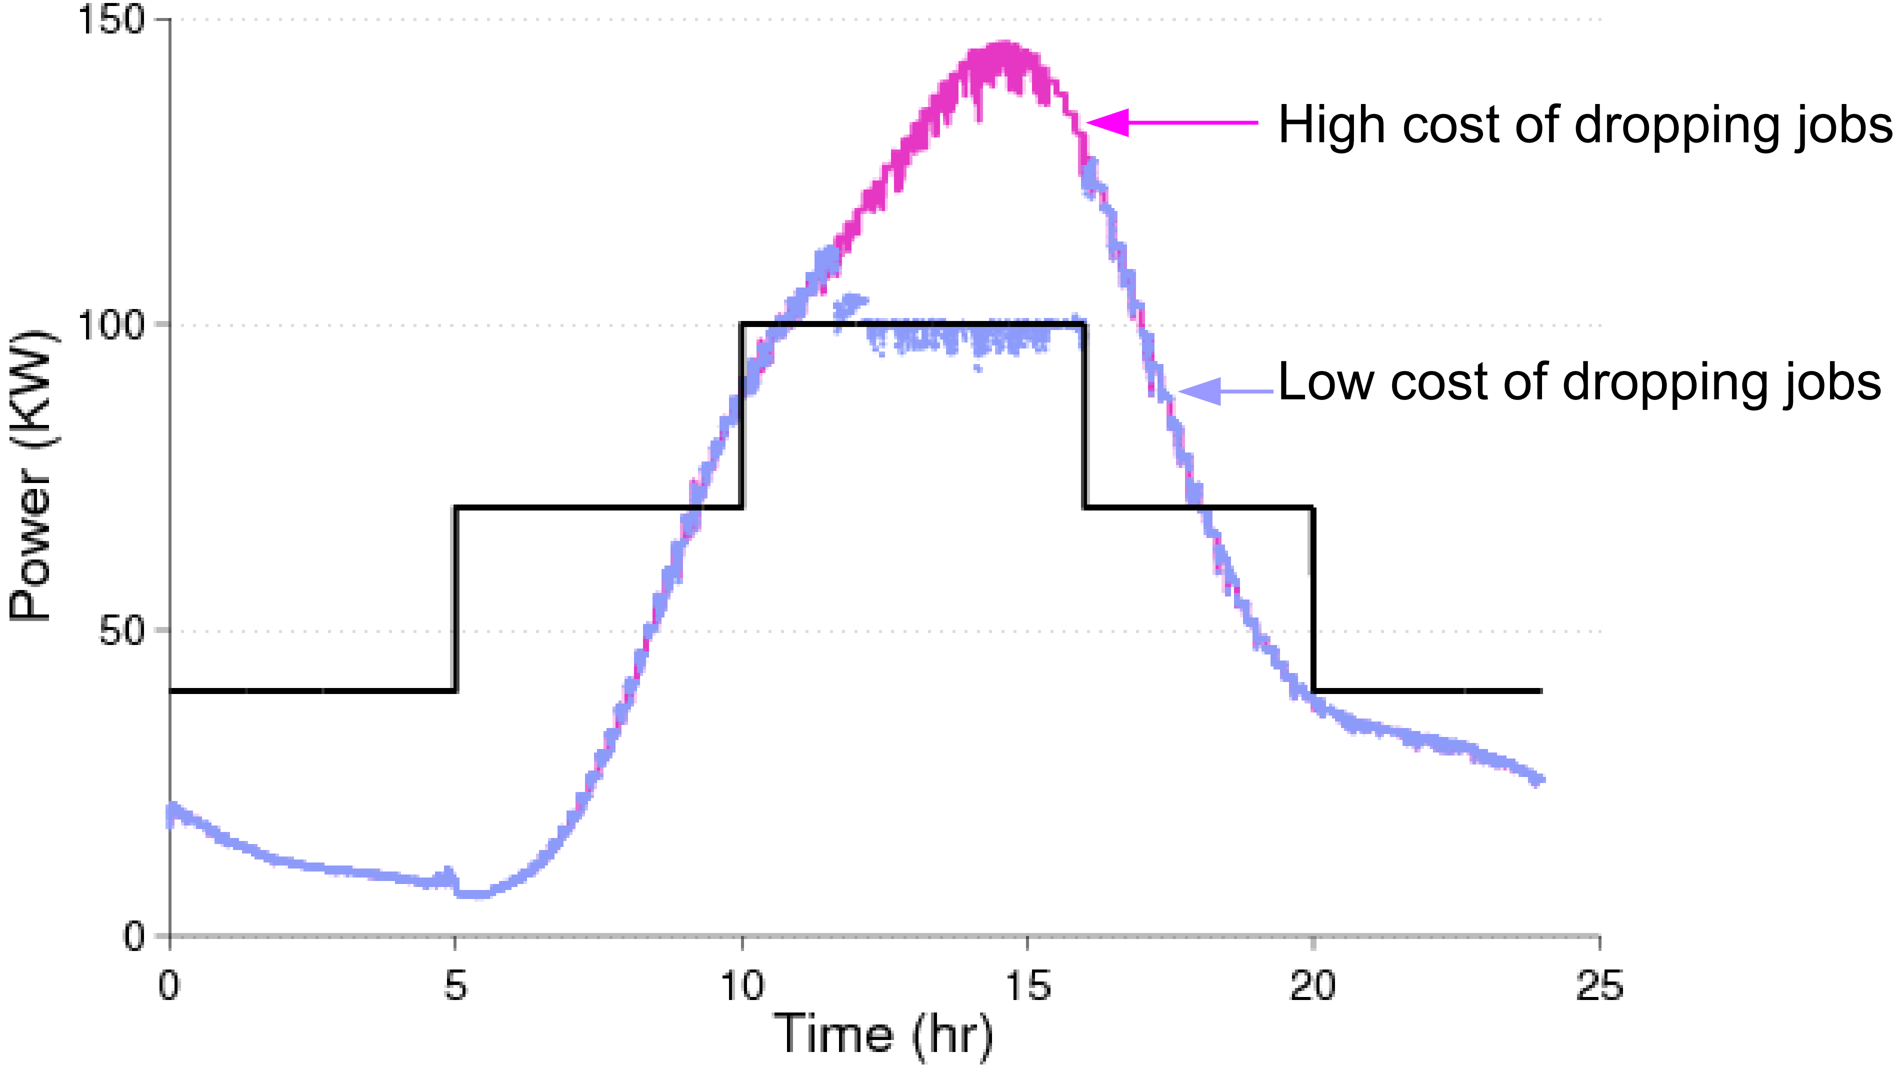

High cost of dropping jobs

Low cost of dropping jobs

Total data center power consumption

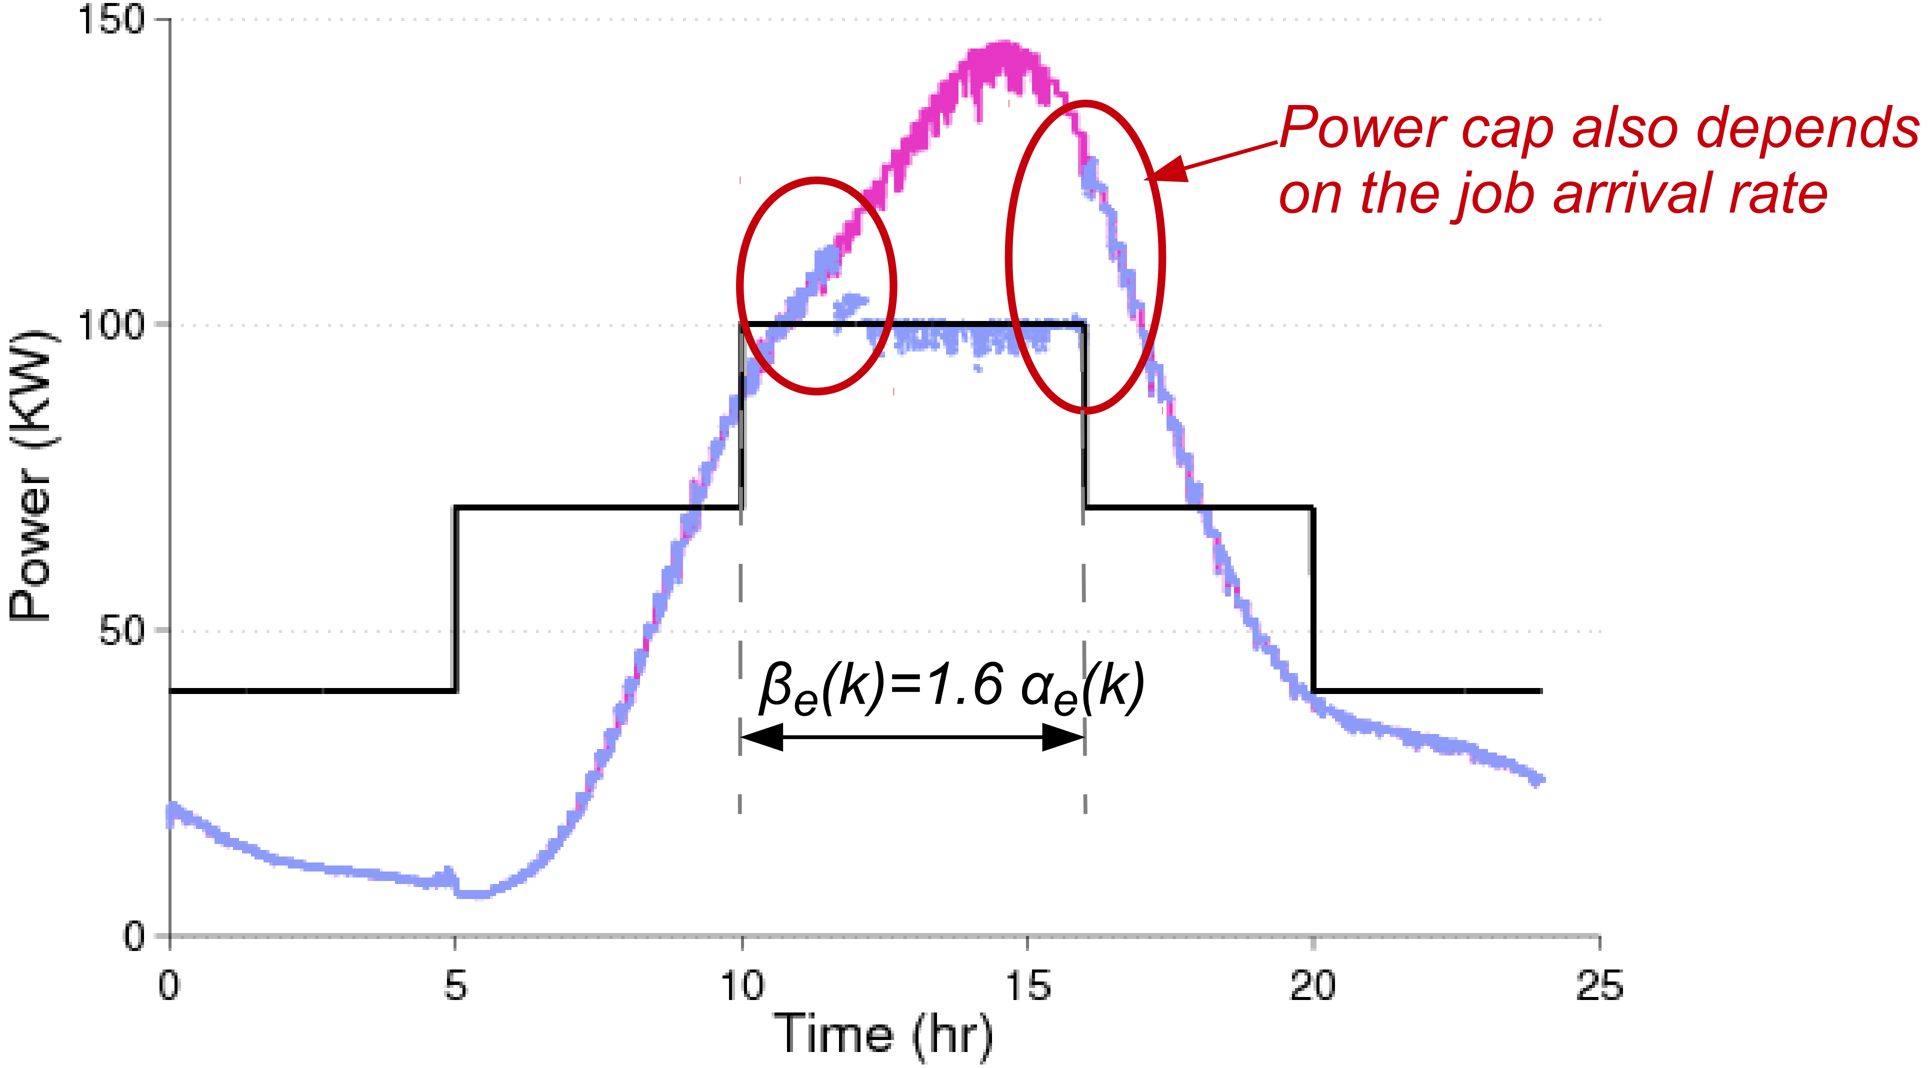

Power cap also depends on the job arrival rate

Conclusion and future work

- **Discussed a control-oriented model which considers both the computational and the thermal characteristics of a data center**
- **Proposed two control strategies which take advantage of SLA with both the users and the power grid**
  - Uncoordinated control approach can be as optimal as the coordinated approach
  - Depends on the thermal-computational characteristics of the data center

Thank you!