How Can Datacenters Join the Smart Grid to Address the Climate Crisis?

Using simulation to explore power and cost effects of direct participation in the energy market

Hongyu He
(VU: 2632195, UvA: 12958794)
hongyu.he@vu.nl

Supervisors:
Prof.Dr.ir. Alexandru Iosup (VU Amsterdam)
Fabian Mastenbroek (TU Delft)
Leon Overweel (Dexter Energy Services B.V.)
Dr.ir. Animesh Trivedi (VU Amsterdam)

22 July 2021

Opening Questions

1. Did you know that energy is a **special** commodity?
2. Did you know the power grid is "**smart**" where prosumers can "**communicate**" with the grid?
3. Did you know datacenters can also **directly** participate in the energy market?

- Datacenters (DCs) are large consumers → **Important for the Market!**
- DCs can be more energy-aware → **Important for DCs themselves!**

Route for Today

1 Introduction

2 Problem Statement

3 Design

4 Evaluation

5 Conclusion

Introduction: Energy is a Commodity

Limitations:

1. Balance must be kept at all time.
2. Large-scale storage is uneconomical.
3. Demand cannot be adjusted by setting prices.
4. *And more* …

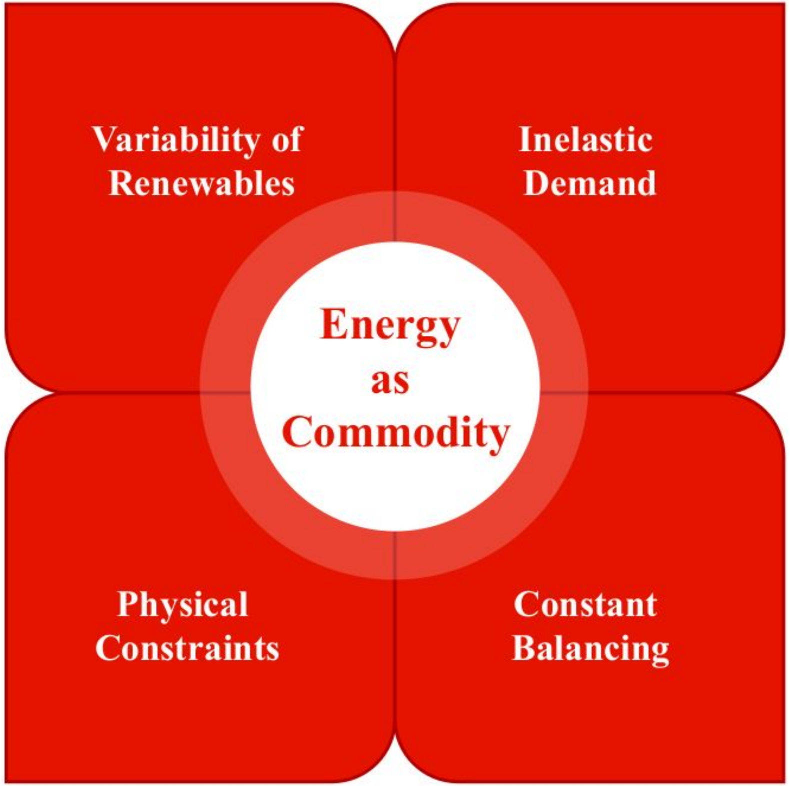

Introduction: Markets around the Power Grid (in the EU)

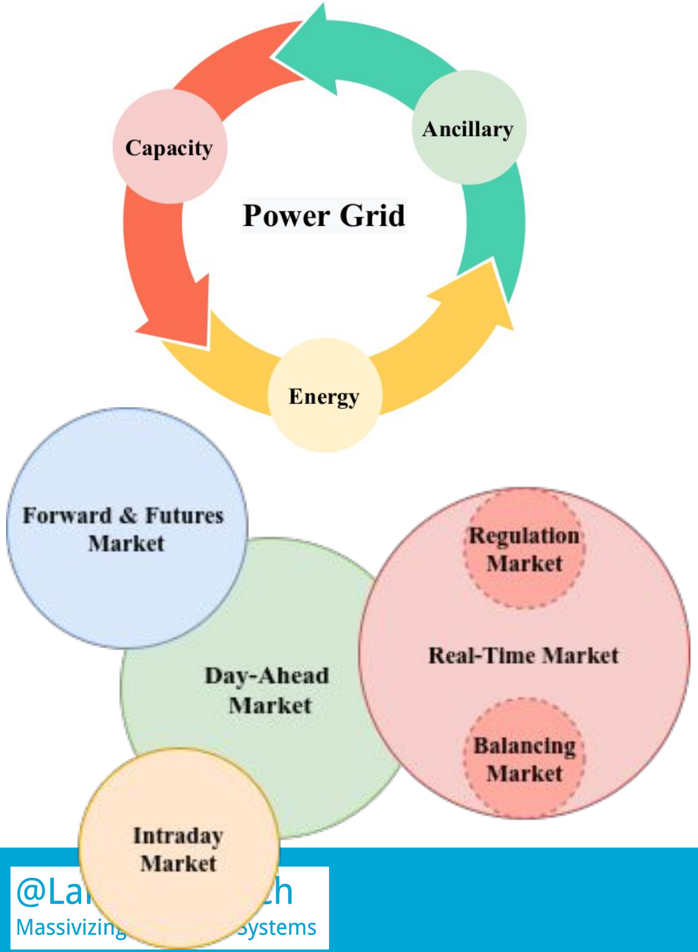

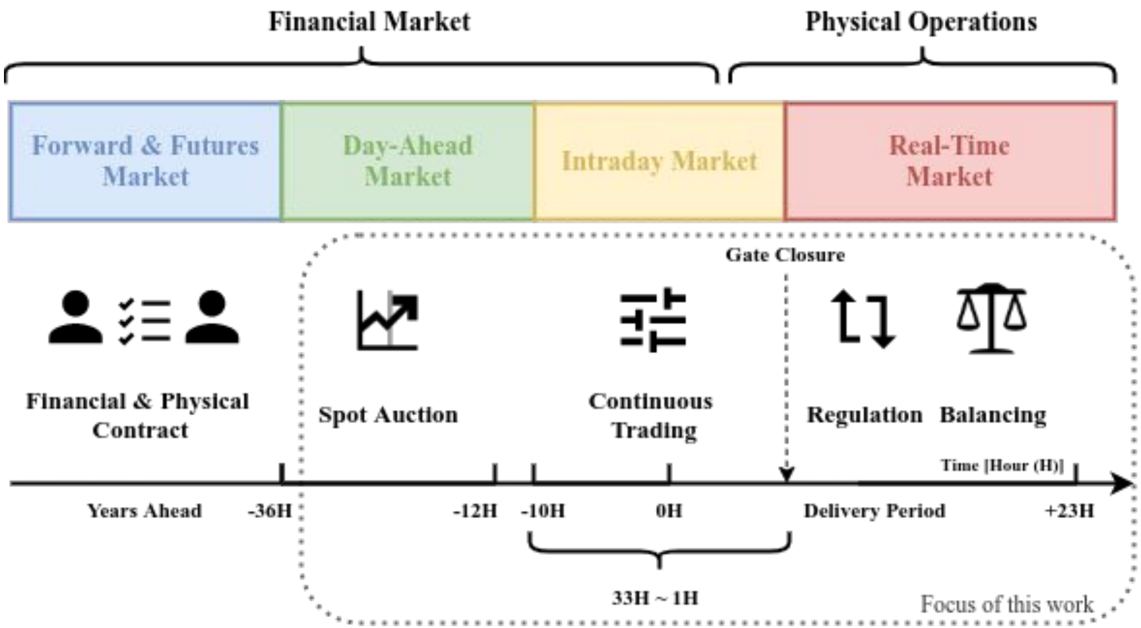

Introduction: Smart Grid

- What if the sun doesn't shine, and the wind doesn't blow?
- The power grid has become "smarter".
- **Demand side management (DSM)**.

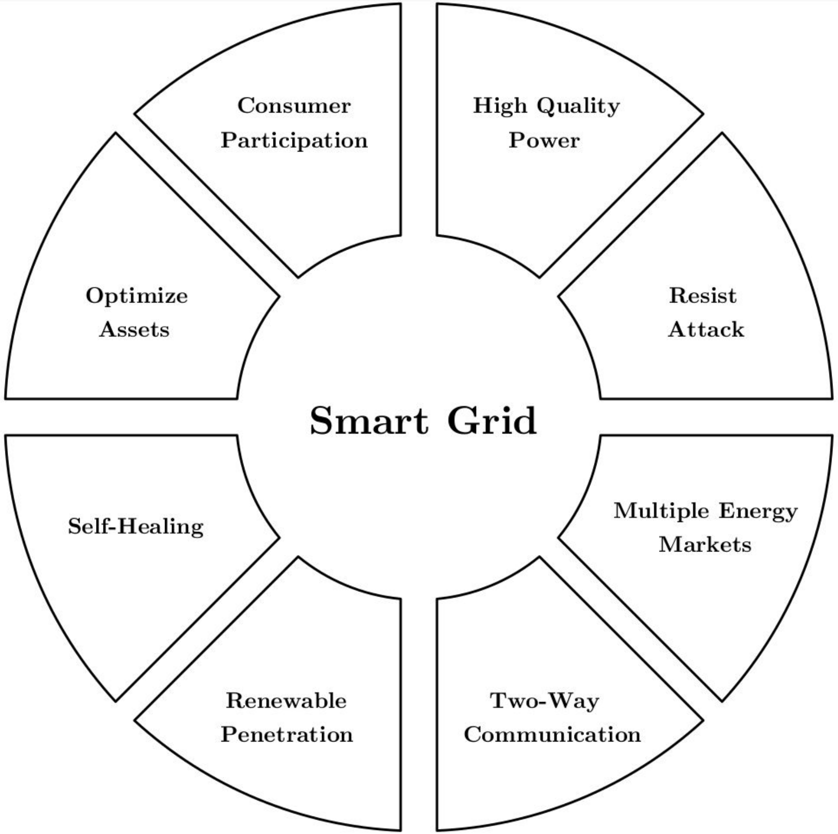

Introduction: Demand Response (DR)

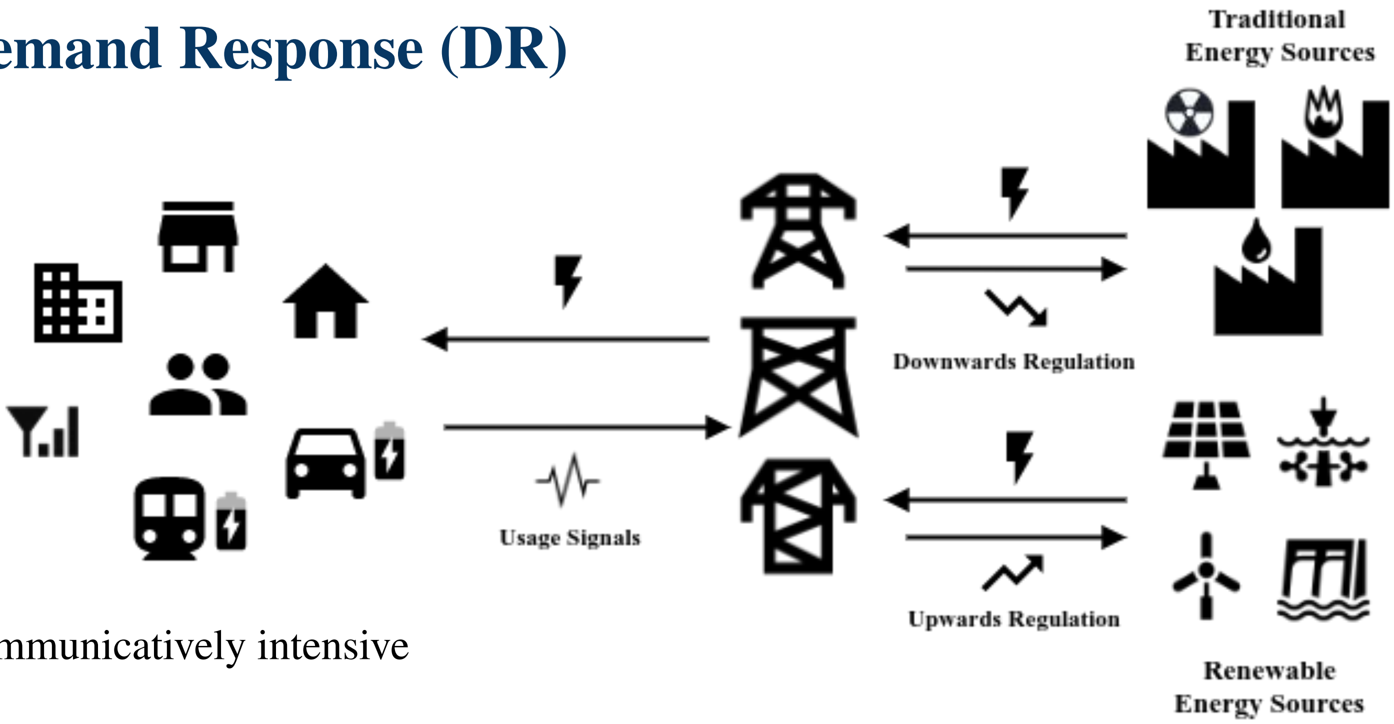

Direct DR

+ Accurate and fast

– Computationally and communicatively intensive

Indirect DR

+ Cheaper and more flexible

– Has uncertainty

Introduction: DVFS

- Active power management.
- DR cannot save energy, but DVFS can!

$$P \propto C \cdot V^2 \cdot F + P^{\mathrm{idle}}$$

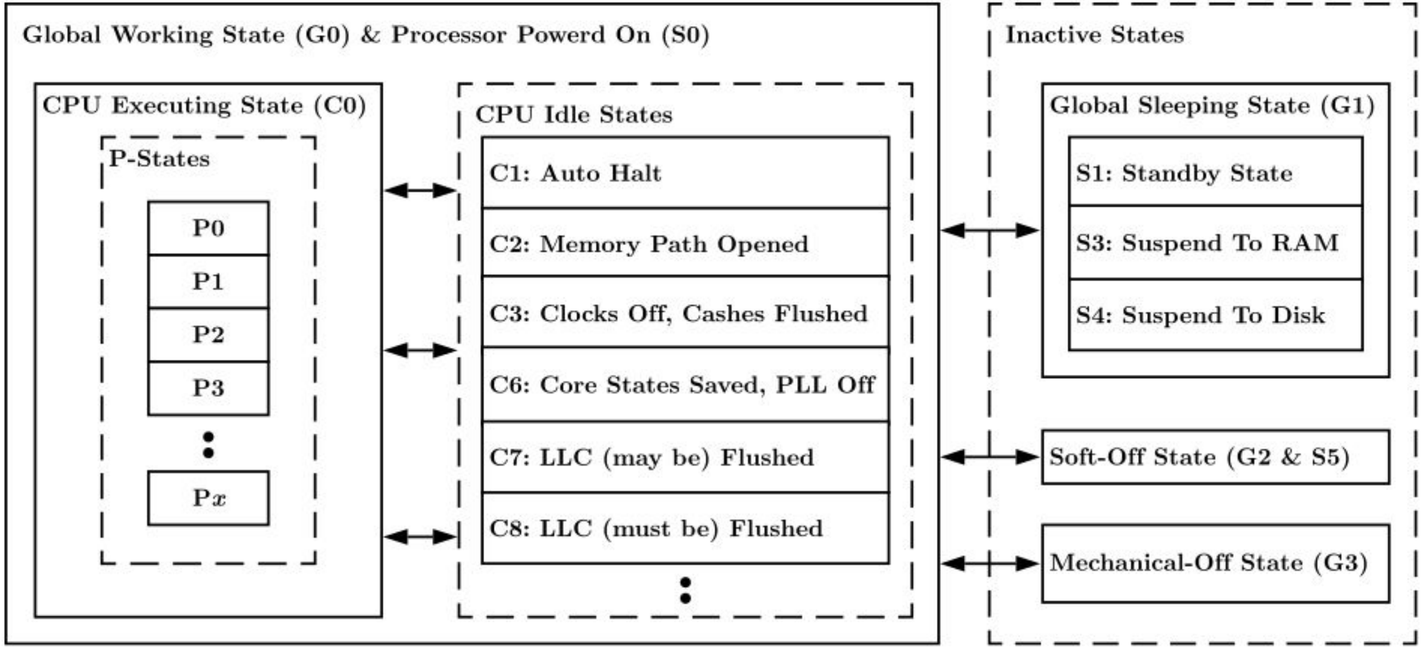

Introduction: Opportunities

DCs are well-suited for providing DR.

But Why?

1. Large capacity ⇒ massive energy storage
2. Elastic load
3. Large redundancy
4. Highly automated
5. *And more …*

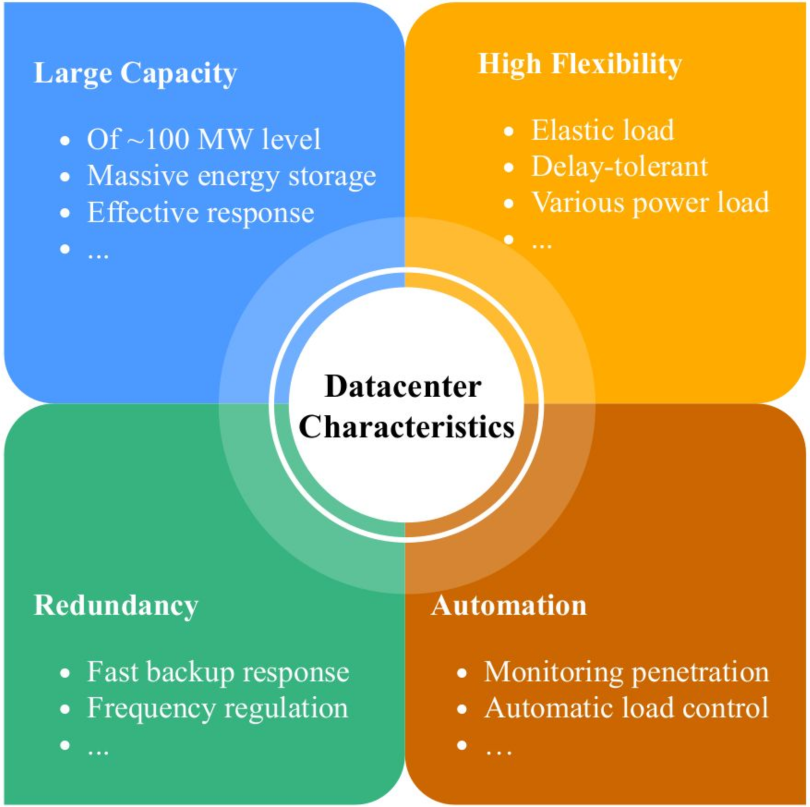

Introduction: Challenges

DCs nowadays provide **little, if any, response** to the power grid [Ghatikar et al. '12] [Glanz et al., '12], [Liu et al., '13].

But why?

1. **Unsuitable** market designs for DCs [Johari et al., '11] [Xu et al., '16];
2. **Limitations** in the current DR programmes [Sle et al., '11] [Liu et al., '14];
3. **High complexity** in proposed methods from existing literature;
4. **Expensive** experimenting, testing, and evaluating energy-saving techniques;

⇒ We need a more instrumental solution to incentivize individual DCs!

Agenda

Problem Statement (1)

MRQ: How feasible and beneficial is it for individual DCs to directly participate in the energy market whilst providing the power grid with indirect DR?

RQ1: How to model the power system of datacenters?

RQ2: Is it beneficial for DCs to participate in the energy market in the first place?
(I.e., why should DCs participate?)

RQ3: How to procure energy in the energy markets according to forecasted power load?
(I.e., how to save energy cost?)

Problem Statement (2)

RQ4: How to optimize energy consumption using DVFS based upon machine learning (ML) methods?

RQ5: How to create an exploratory tool for problems in this domain, to be used by experts in both the IT and the energy industry?

Thesis Statement — **Individual** DCs can and should **directly** participate in the energy market to provide the power grid with **indirect DR**, whilst both **saving their energy costs** and **curbing their energy consumption**.

Agenda

Design: Development Pipeline

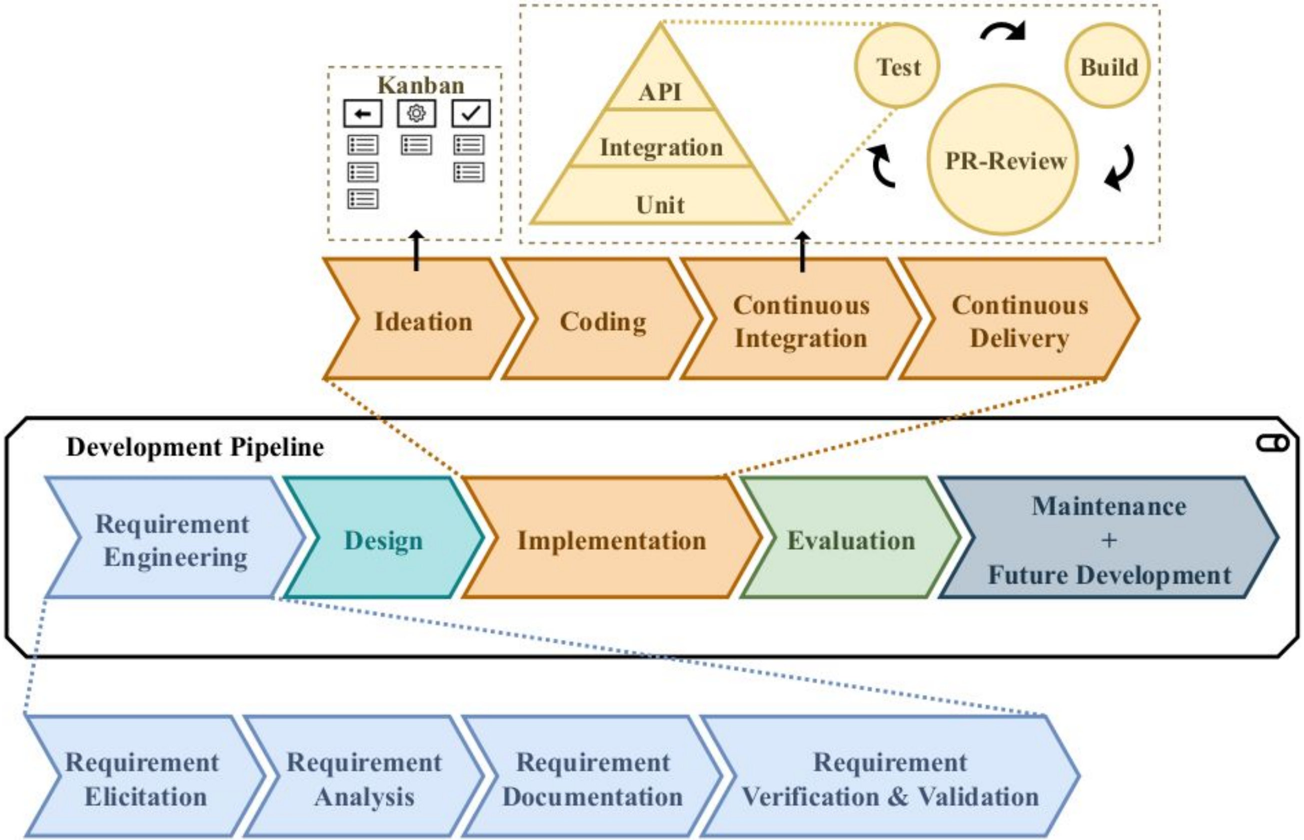

Design: Requirement Engineering (1)

Who cares?

Industry	Stakeholders
IT	datacenter managers, datacenter operators, datacenter technicians, cloud architects, cloud tenants
Energy	consulting firms, energy market operators, power grid system operators, renewable energy suppliers
Others	legislators, end-users of cloud services

Category	Stakeholders
Active stakeholder	datacenter managers, datacenter operators, consulting firms, energy market operators, renewable energy suppliers
Passive stakeholder	datacenter technicians, cloud architects, power grid system operators, legislators, end-users of cloud services

Design: Requirement Engineering (2)

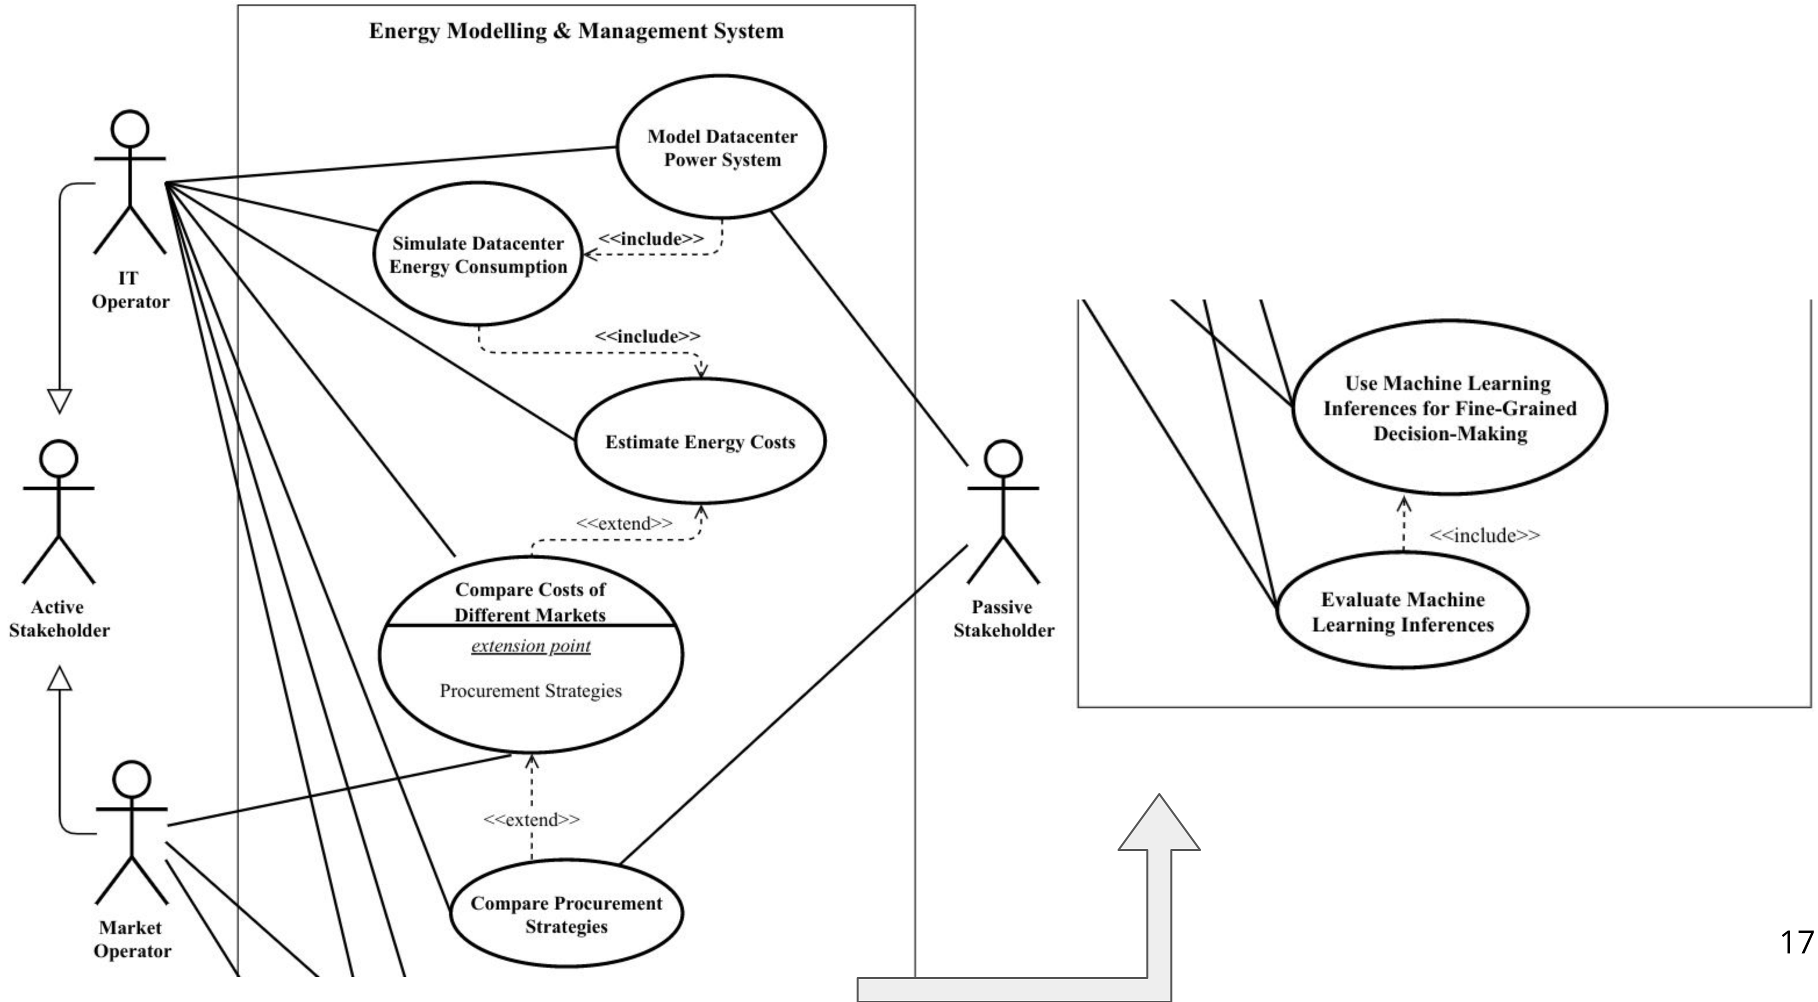

Design: Power Modelling & Management System*

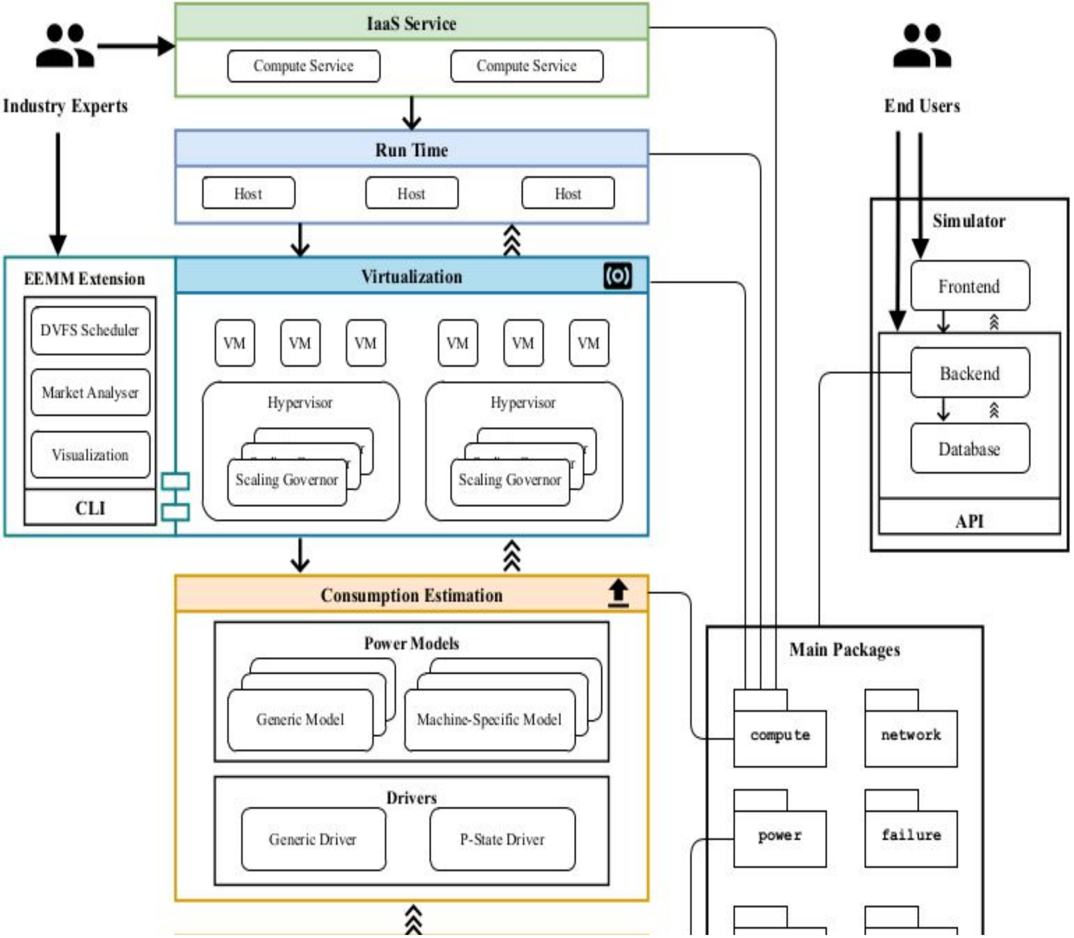

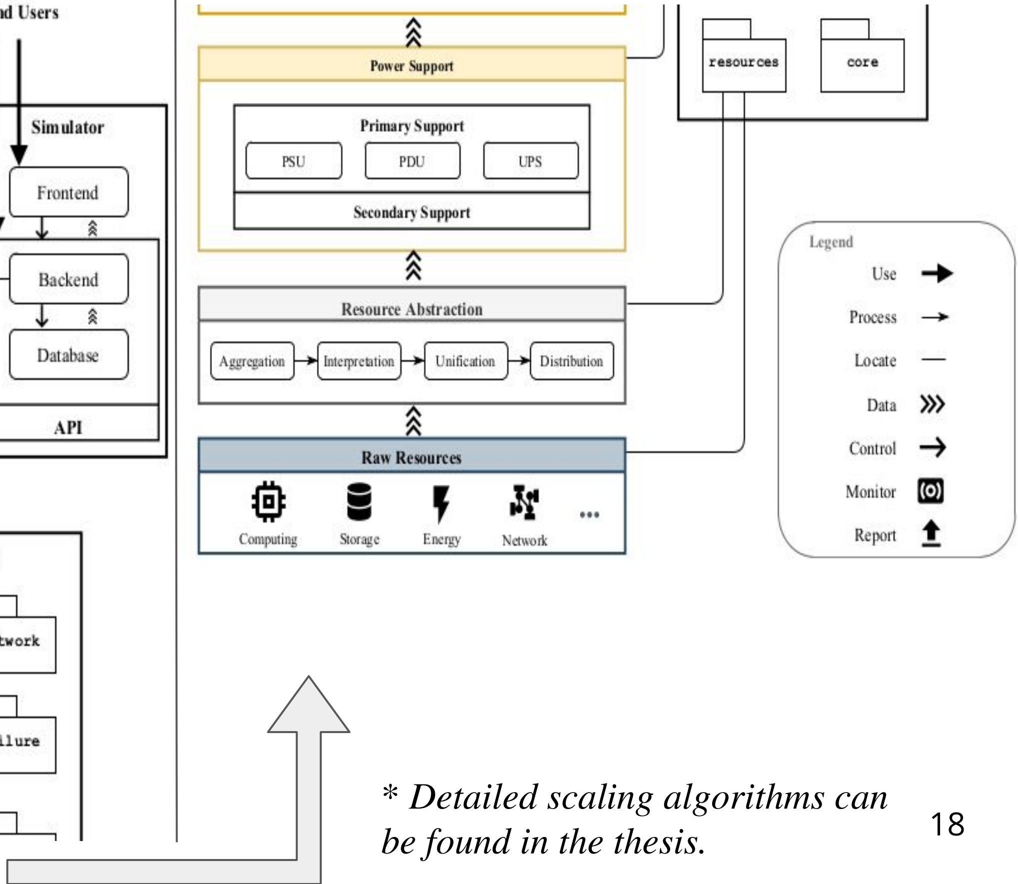

** Detailed scaling algorithms can be found in the thesis.*

Design: DC Power System Model

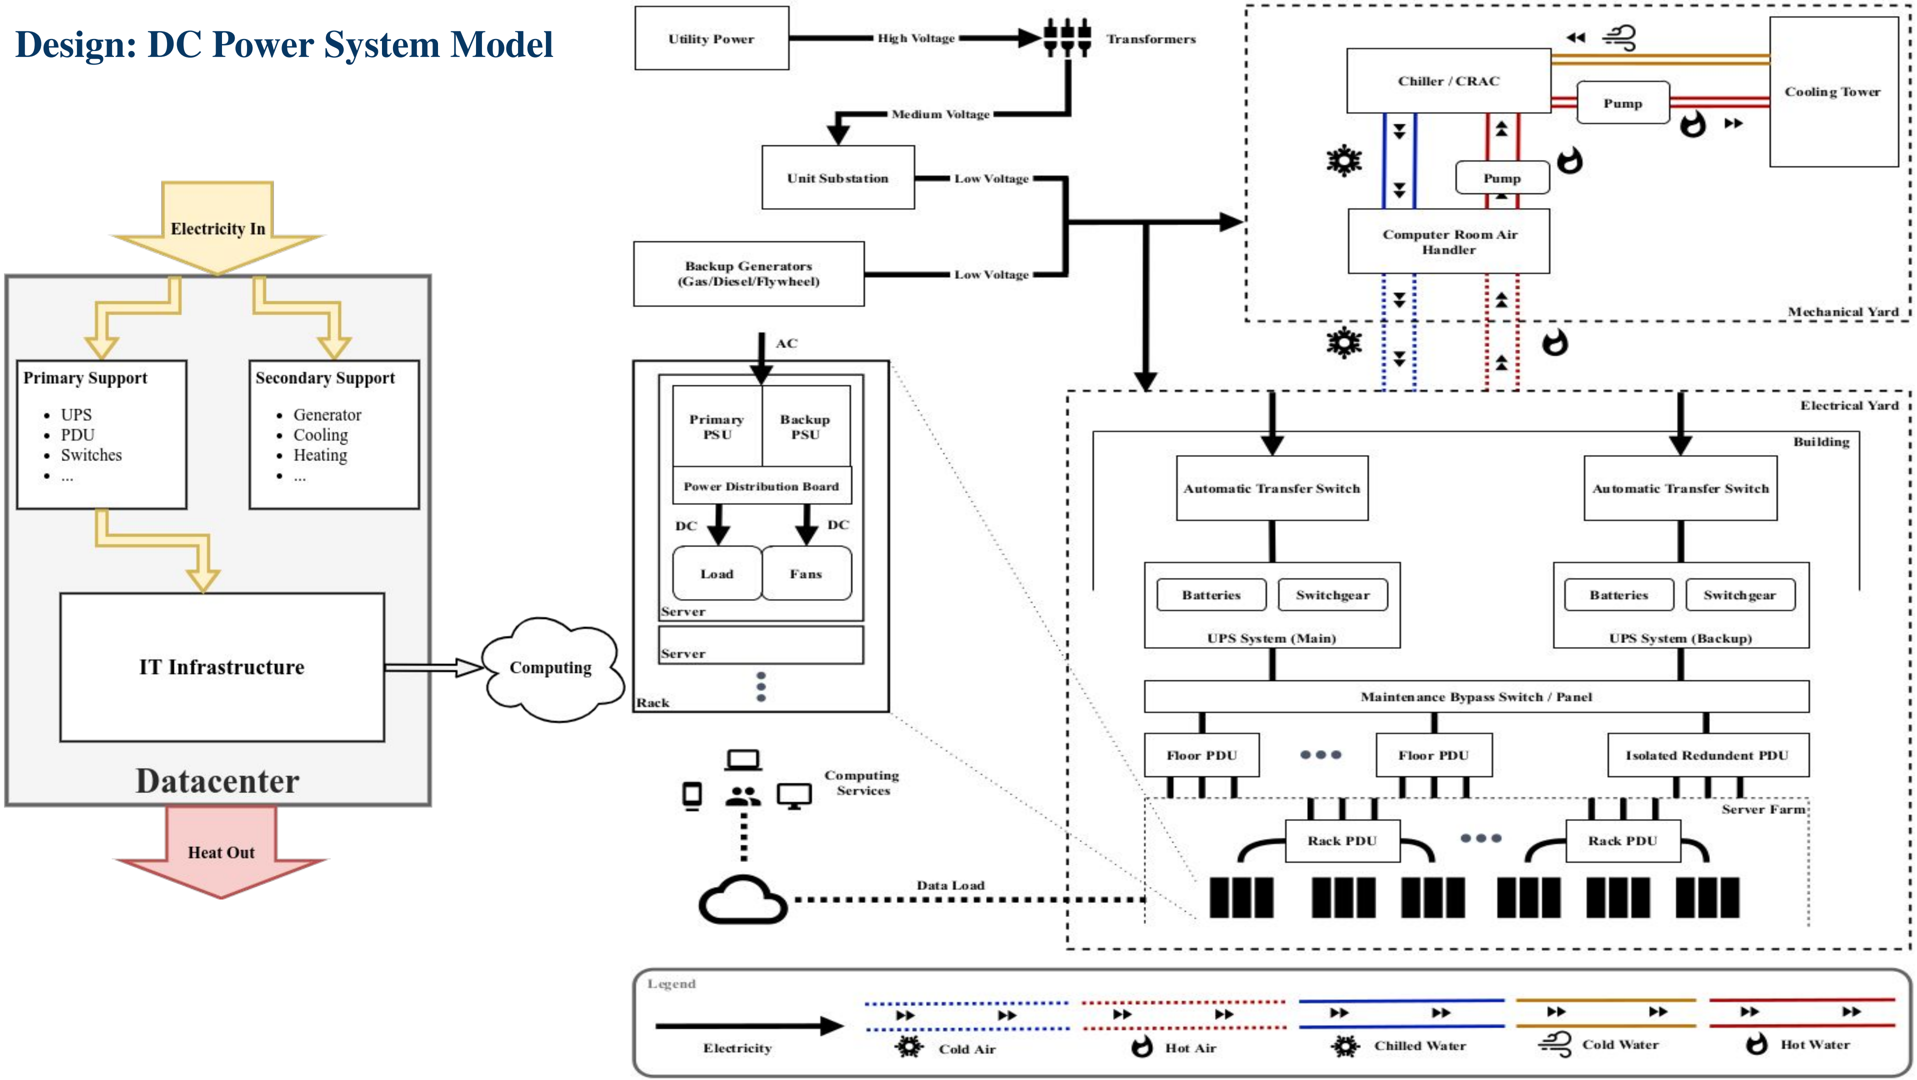

Design: Power Support Subsystem*

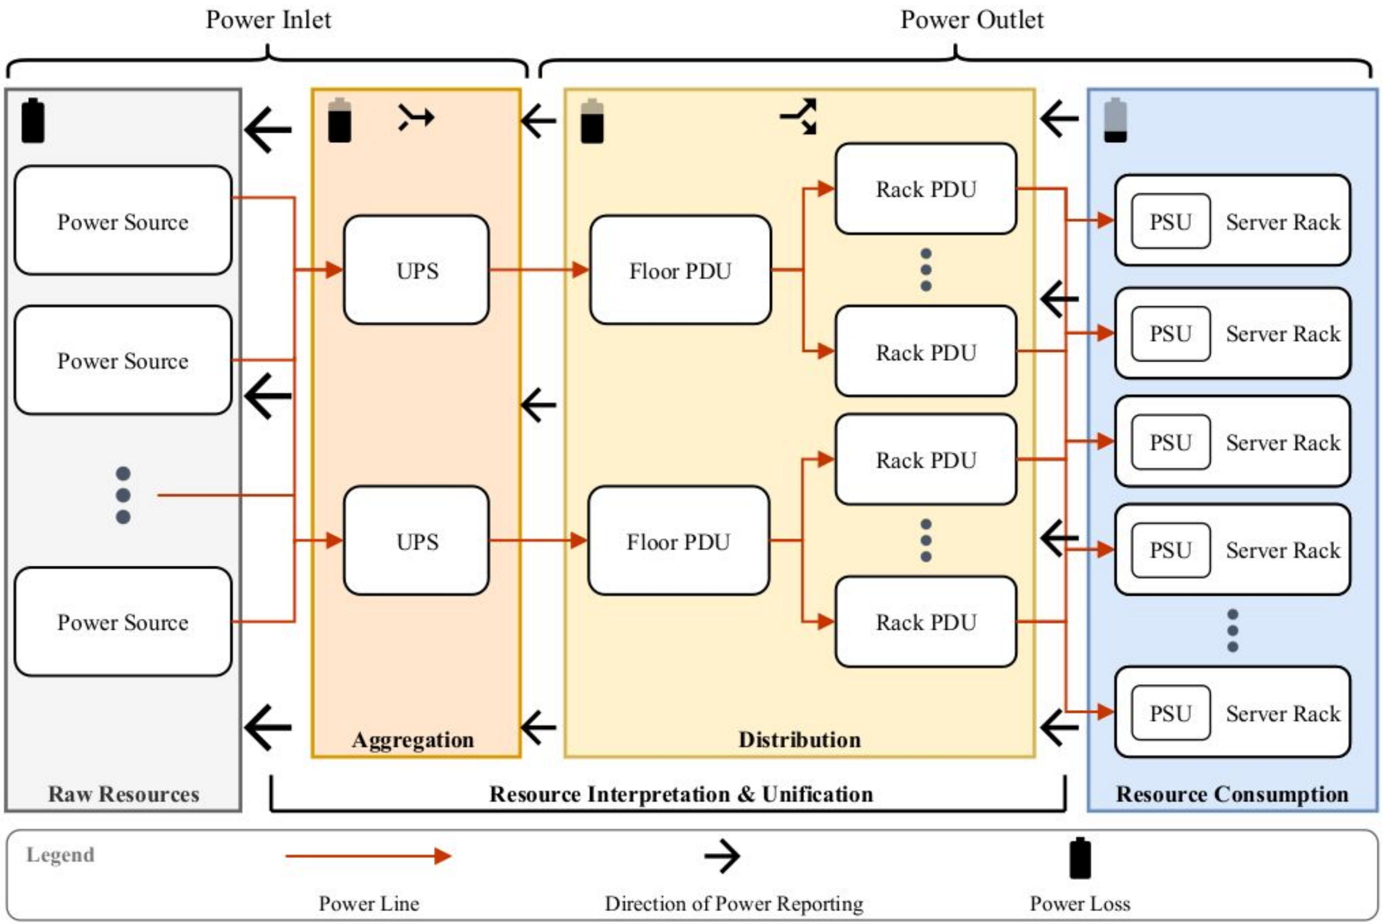

** Detailed power distribution algorithms can be found in the thesis.*

Design: Market Extension*

Installation:

```
$ pip install opendc-eemm
```

Code:

https://github.com/hongyuhe/opendc-eemm

Doc:

https://opendc-eemm.rtfd.io

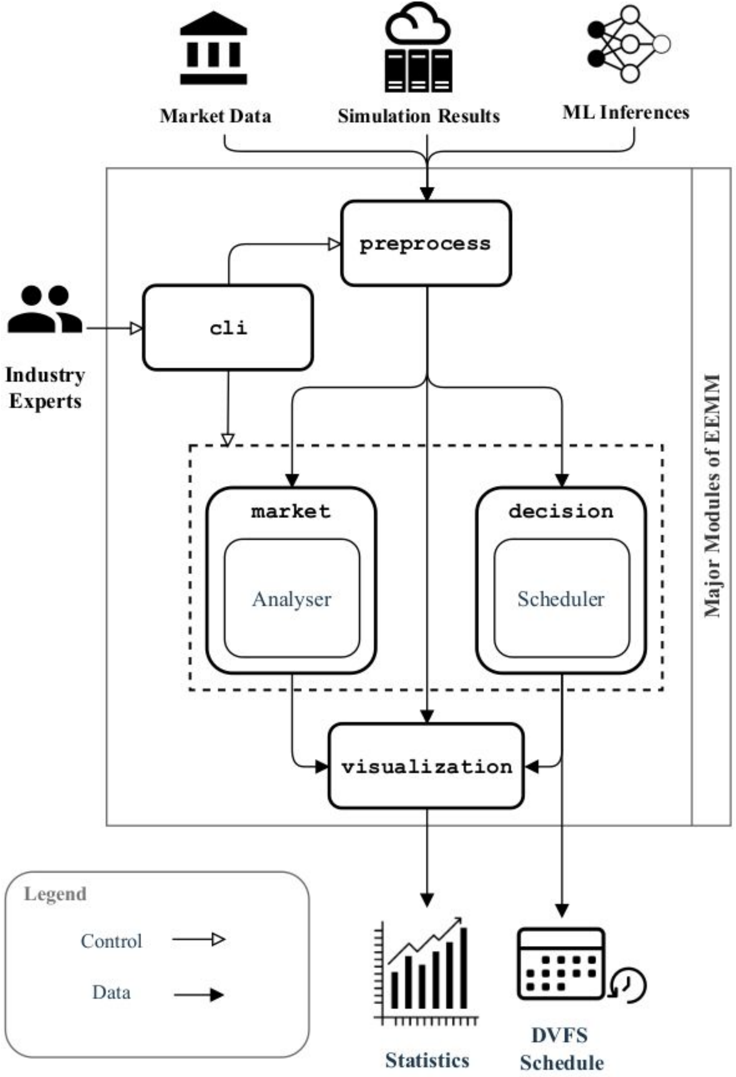

** Detailed scheduling algorithm can be found in the thesis.*

Design: Comparison with the State of the Art

Simulator	IT Infrasturcture		Primary Support		Secondary Support	Energy Market Integration
	Critical Load	DVFS	UPS	PDU		
DCSim [67]	✓	✗	✗	✗	✗	✗
CloudSim [30]	✓	✓	✗	✗	✗	✗
GDCSim [67]	✓	✗	✗	✗	✓+	✗
CloudSched [153]	✓	✓	✗	✗	✗	✗
DISSECT-CF [115, 90, 89]	✓	✓	✗	✗	✓+	✗
GreenCloud [23, 155, 94]	✓+	✓+	✗	✗	✗	✗
iCanCloud/E-mc² [32, 123]	✓+	✓+	✗	✗	✗	✗
SimGrid [31, 72, 47]	✓+	✓+	✗	✗	✗	✗
OpenDC [82, 119]	✓	✓+	✓+	✓+	✓	✓+

Table 2.2: Overview of the eight surveyed datacenter simulators, where the ✓ symbol means that the corresponding energy model is available, the ✗ symbol means that it is unavailable, and + represents advanced support.

Agenda

1 Introduction

2 Problem Statement

3 Design

4 Evaluation

5 Conclusion

Evaluation: Setup (1)

Machine Model

Machine Model	Year of Release	CPU	Base Frequency	Cache	#Cores	#Threads
Old	2007	Intel® Core™2 Quad Q6700	2.66 GHz	8 MB	4	4
New	2021	Intel® Xeon® Platinum 8380	2.30 GHz	60 MB	40	80

$$N_{\text{Hosts}} = \left\lceil \left\lceil \frac{\Psi_d}{\Psi_f} \right\rceil \Big/ c \right\rceil$$

$$N_{\text{Units}} = \left\lceil \left\lceil \frac{\Psi_m}{N_{\text{Hosts}}} \right\rceil \Big/ m \right\rceil$$

Machine Model	#Cores/Host	#Host	Size of Memory Unit [MB]	#Memory Units/Host
Old	4	284	4,000	4
New	160	9	3,200	48

Market Model

- Markets of interest: on-demand, day-ahead, and balancing
- Pricing systems
  - Balancing: two-price system (introduce later).
  - On-demand:

Price Level	Price [€/MWh]	Source
Low	38.0	NieuweStroom B.V. (2021, average) [29]
Medium	56.5	PricewaterhouseCoopers (2017, average) [133]
High	80.4	Essent N.V. (2021, fixed) [124]

Evaluation: Setup (2)

Energy Model

- Critical load:
  - Square-root model (SQRT) → old machine model (upper bound)
  - Linear model (LINEAR) → old machine model (lower bound)
  - Interpolation (INTERPOLATION) → new machine model

- PSU: 870 W (AC) of 80 Plus Titanium standard
- PUE: 1.58 (global average*)

Traces: Bitbrains http://gwa.ewi.tudelft.nl/datasets/gwa-t-12-bitbrains

* https://journal.uptimeinstitute.com/data-center-pues-flat-since-2013/

Evaluation: Setup (3)

Energy Model

1. Critical load:

`LinearPowerModel`	$P(u) = P_{\text{idle}} + (P_{\max} - P_{\text{idle}})\,u$
`SqrtPowerModel`	$P(u) = P_{\text{idle}} + (P_{\max} - P_{\text{idle}})\sqrt{u}$
`AsymptoticPowerModel`	$P(u) = P_{\text{idle}} + \frac{(P_{\max} - P_{\text{idle}})}{2}\left(1 + u - e^{-\frac{u}{\alpha}}\right)$

$$\eta_l = \frac{100}{\tau} \cdot P_{\text{server}} \tag{5.3}$$

$$\eta_e = \begin{cases} 90\% & \text{if } 0 \leqslant \eta_l \leqslant 10\% \\ 94\% & \text{if } 10 < \eta_l \leqslant 20\% \\ 96\% & \text{if } 20\% < \eta_l \leqslant 50\% \\ 91\% & \text{if } 50\% \leqslant \eta_l < 100\% \end{cases} \tag{5.4}$$

$$E_{\text{PSU}} = \int_{t_0}^{t_{\text{PSU}}} \frac{(P_{\text{server}} \cdot 100)}{\eta_e} - P_{\text{server}} dt \tag{5.5}$$

2. Primary support:

$$\begin{cases} \alpha = \pi_{\text{UPS}} - \lambda_{\text{UPS}} \\ P_{\text{UPS}}^{\text{tare}} = \lambda_{\text{UPS}} \cdot P_{\text{UPS}}^{\text{rated}} \\ P_{\text{UPS}}^{\text{loss}} = P_{\text{UPS}}^{\text{tare}} + \alpha \cdot (\sum_i^{N_{\text{PDU}}} P_{\text{PDU}_i}^{\text{in}}) \end{cases}$$

$$\begin{cases} \beta = \pi_{\text{PDU}} - \lambda_{\text{PDU}} \\ P_{\text{PDU}}^{\text{tare}} = \lambda_{\text{PDU}} \cdot P_{\text{PDU}}^{\text{rated}} \\ P_{\text{PDU}}^{\text{loss}} = P_{\text{PDU}}^{\text{tare}} + \beta \cdot (\sum_i^{N_{\text{server}}} P_{\text{server}_i}^{\text{in}})^2 \end{cases}$$

3. Secondary support:

$$P^{2\text{nd}} = \text{PUE} \cdot \sum_i^{N_{\text{server}}} P_{\text{server}_i} - \left(\sum_i^{N_{\text{PDU}}} P_{\text{PDU}_i} + \sum_i^{N_{\text{UPS}}} P_{\text{UPS}_i} \right)$$

Evaluation: Power loads

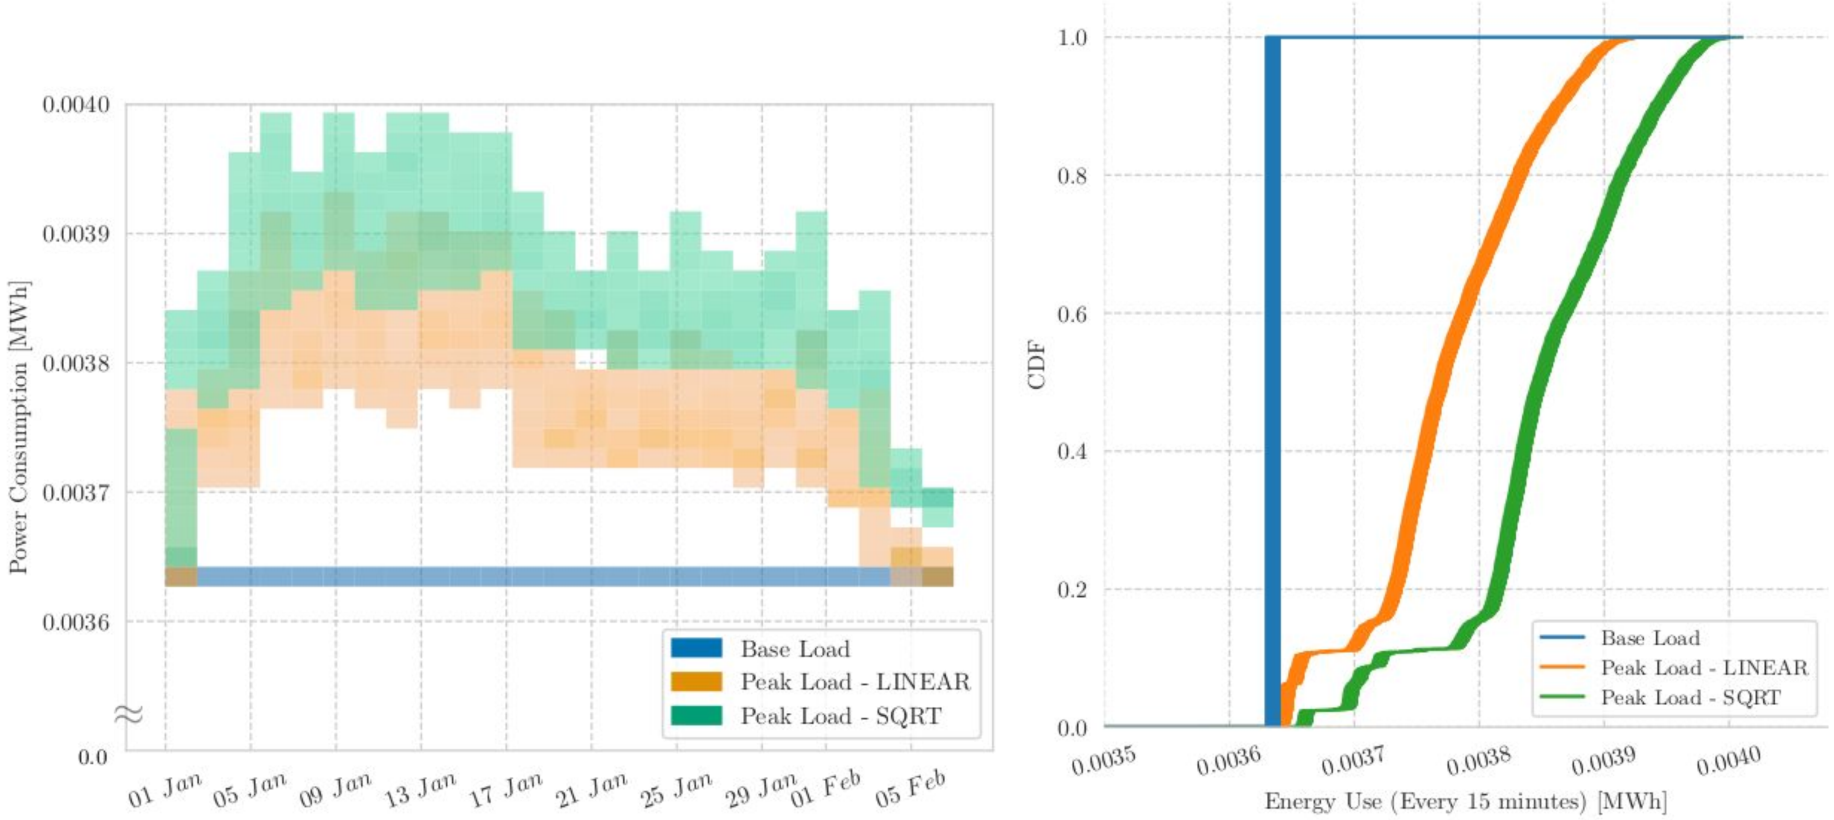

(a) Instant power loads.

(b) CDF of the power loads.

Evaluation: Energy Costs

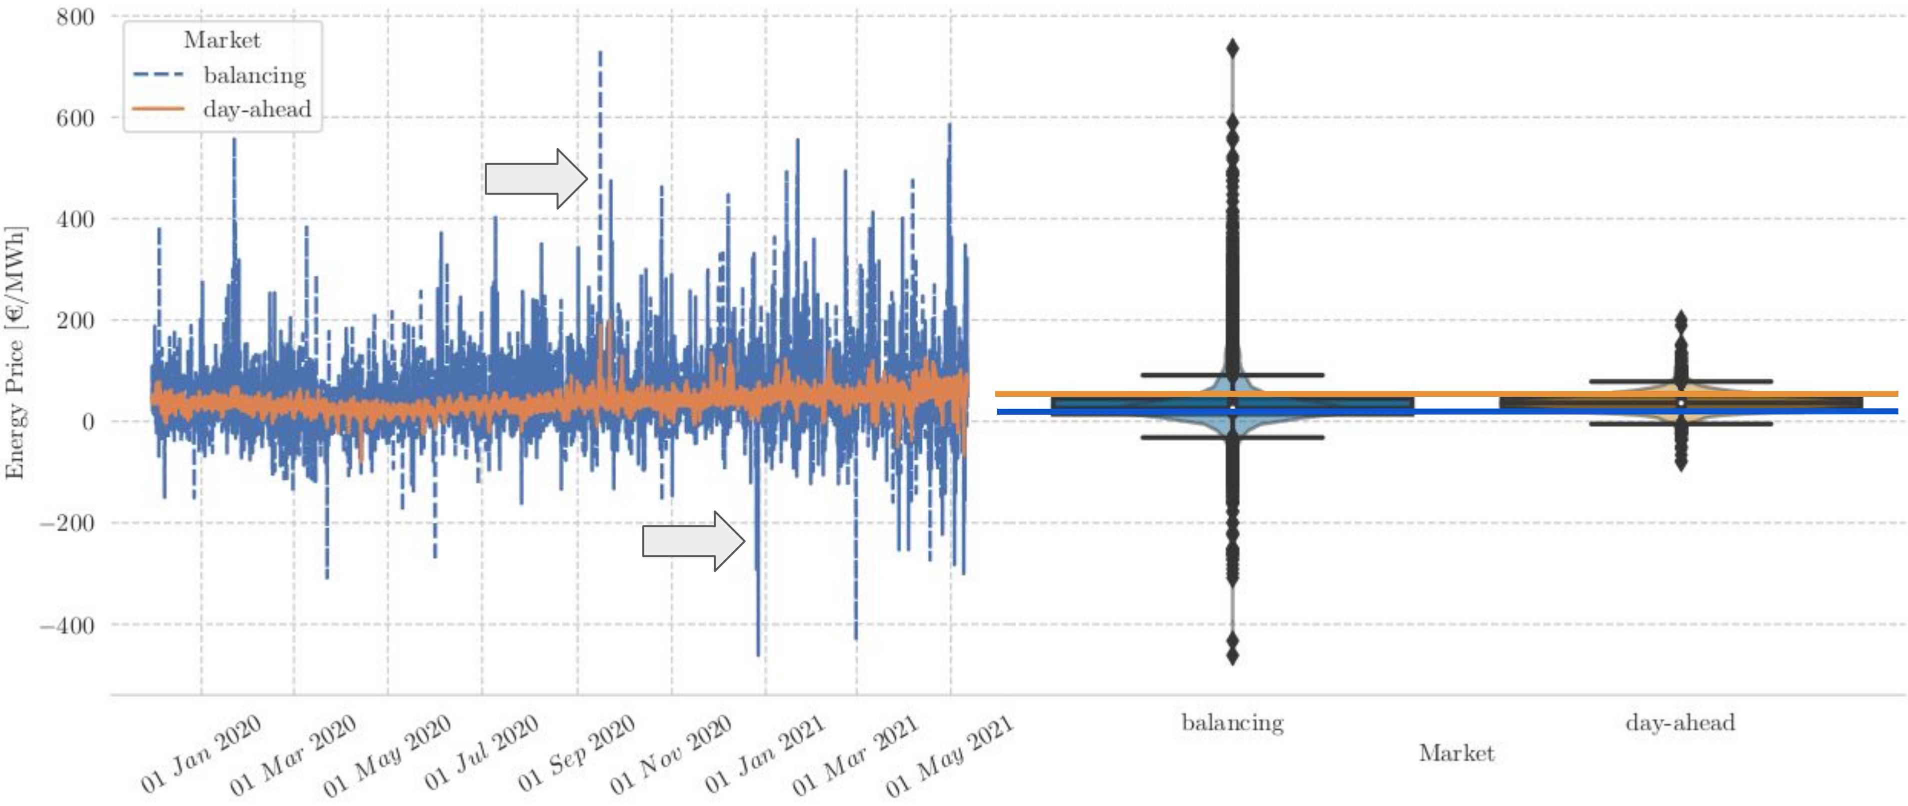

Evaluation: Summary 1

It is **financially beneficial** for DCs to participate in both the day-ahead and the balancing market.

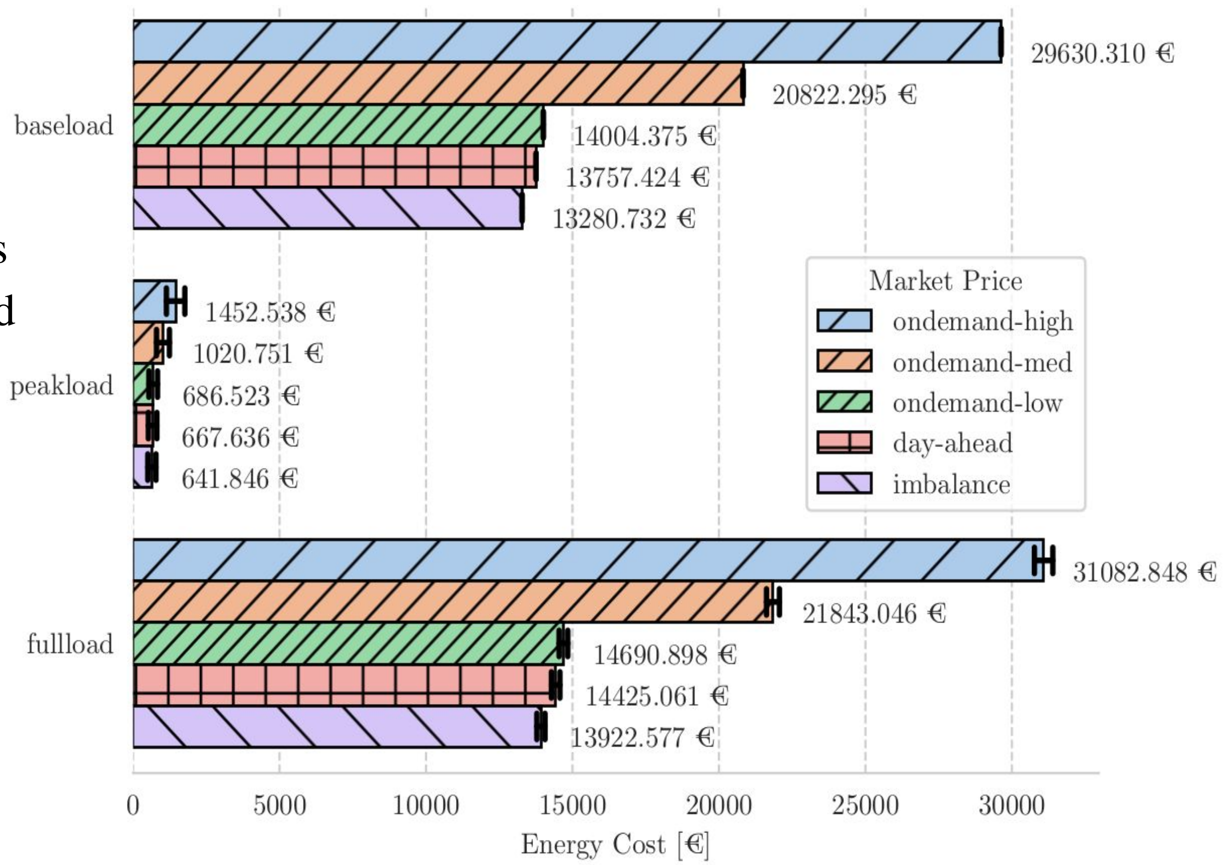

Evaluation: The benefit of having more recent machines (1)

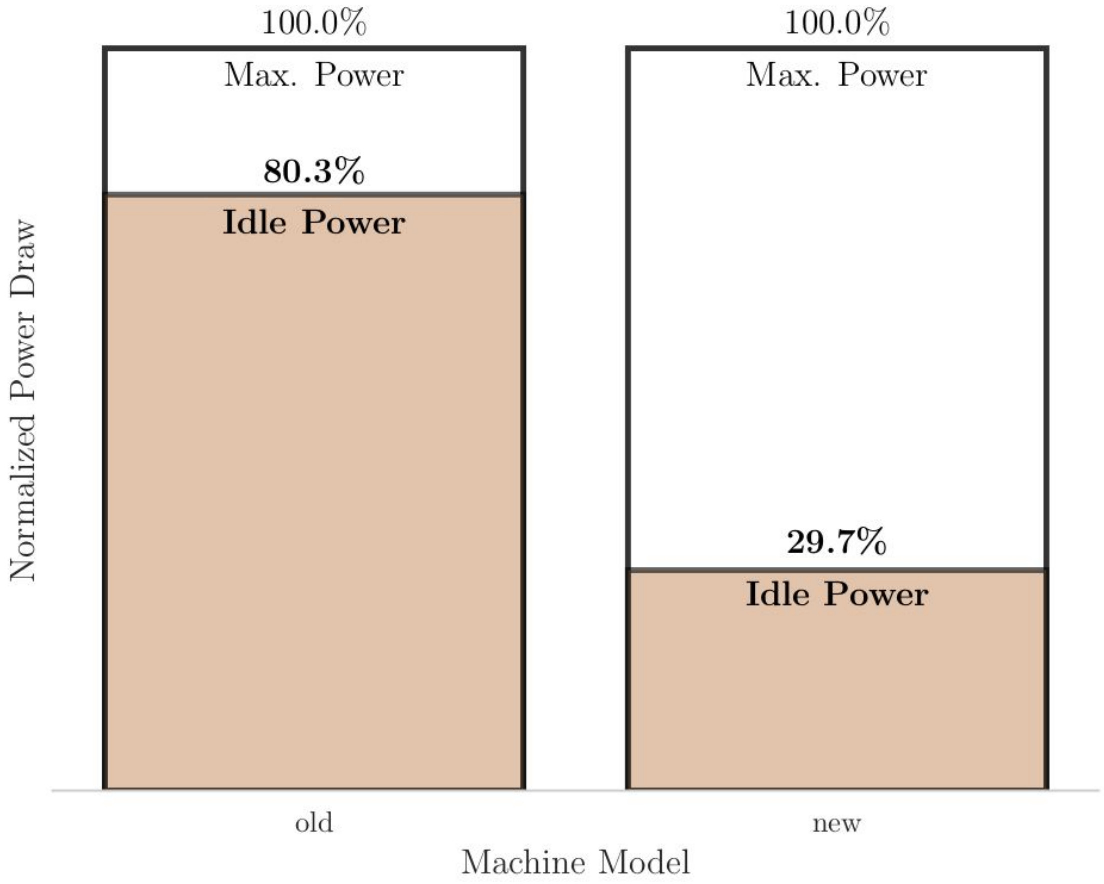

Evaluation: The benefit of having more recent machines (2)

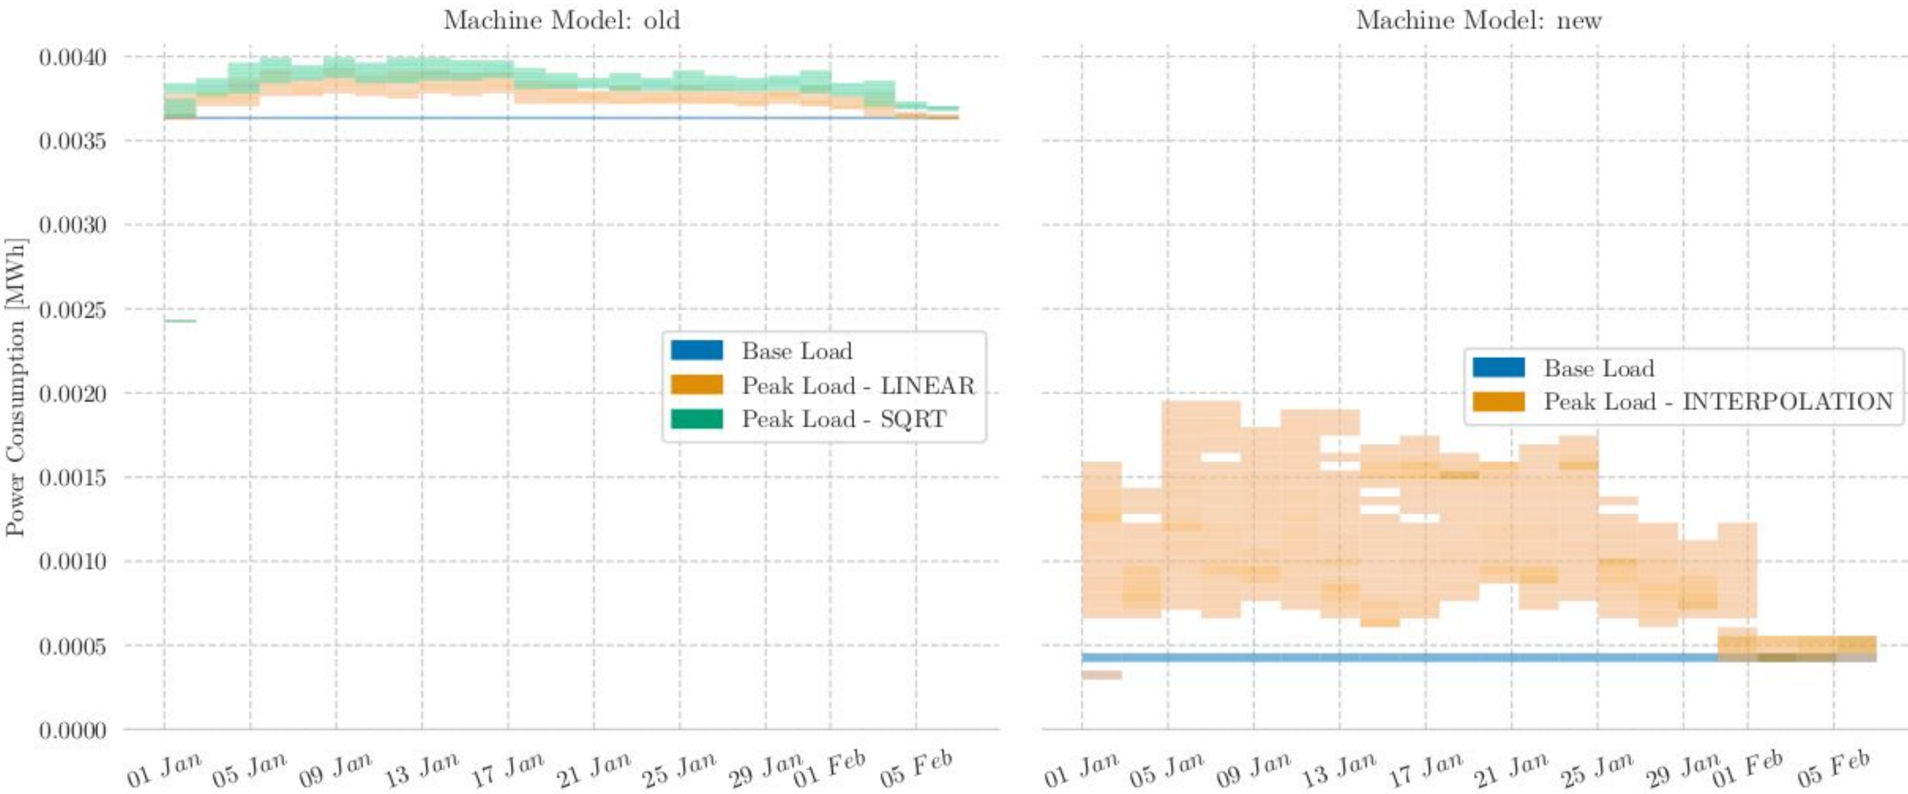

Evaluation: The benefit of having more recent machines (3)

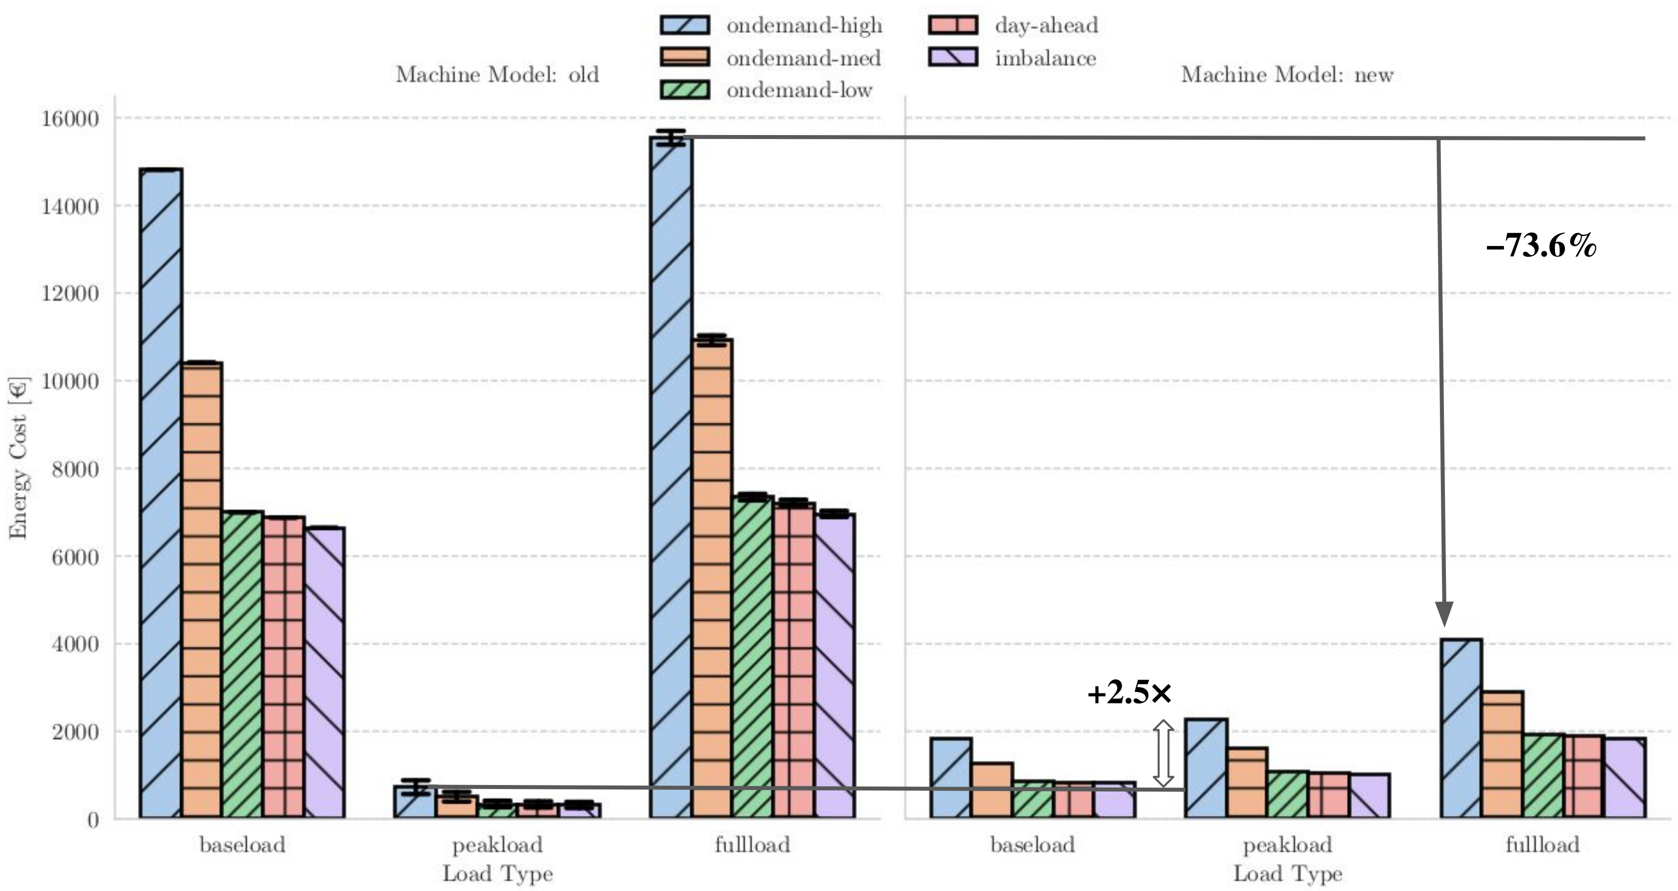

Evaluation: Energy Procurement Strategy (1)

1. Assumptions (**A**s):

    A1 Datacenter operators purchase energy in the day-ahead market based on the load forecast of the corresponding *day*.

    A2 The load forecast is *perfect*. In other words, the load predictions always precisely match the actual loads of the corresponding day.

    A3 Datacenter operators do not deliberately schedule even less energy than the bare-minimum quantity — the base load.

    A4 Datacenters' participation is abided by the two-price balancing system.

2. Procurement model:

$$Q^S = \mathbb{Q}_q(l_f) \cdot s, \tag{5.9}$$

where l_f denotes the load forecast of the next day, q is the quantile of the quantile function $\mathbb{Q}$, and s is a scalar to apply.

Evaluation: Energy Procurement Strategy (2)

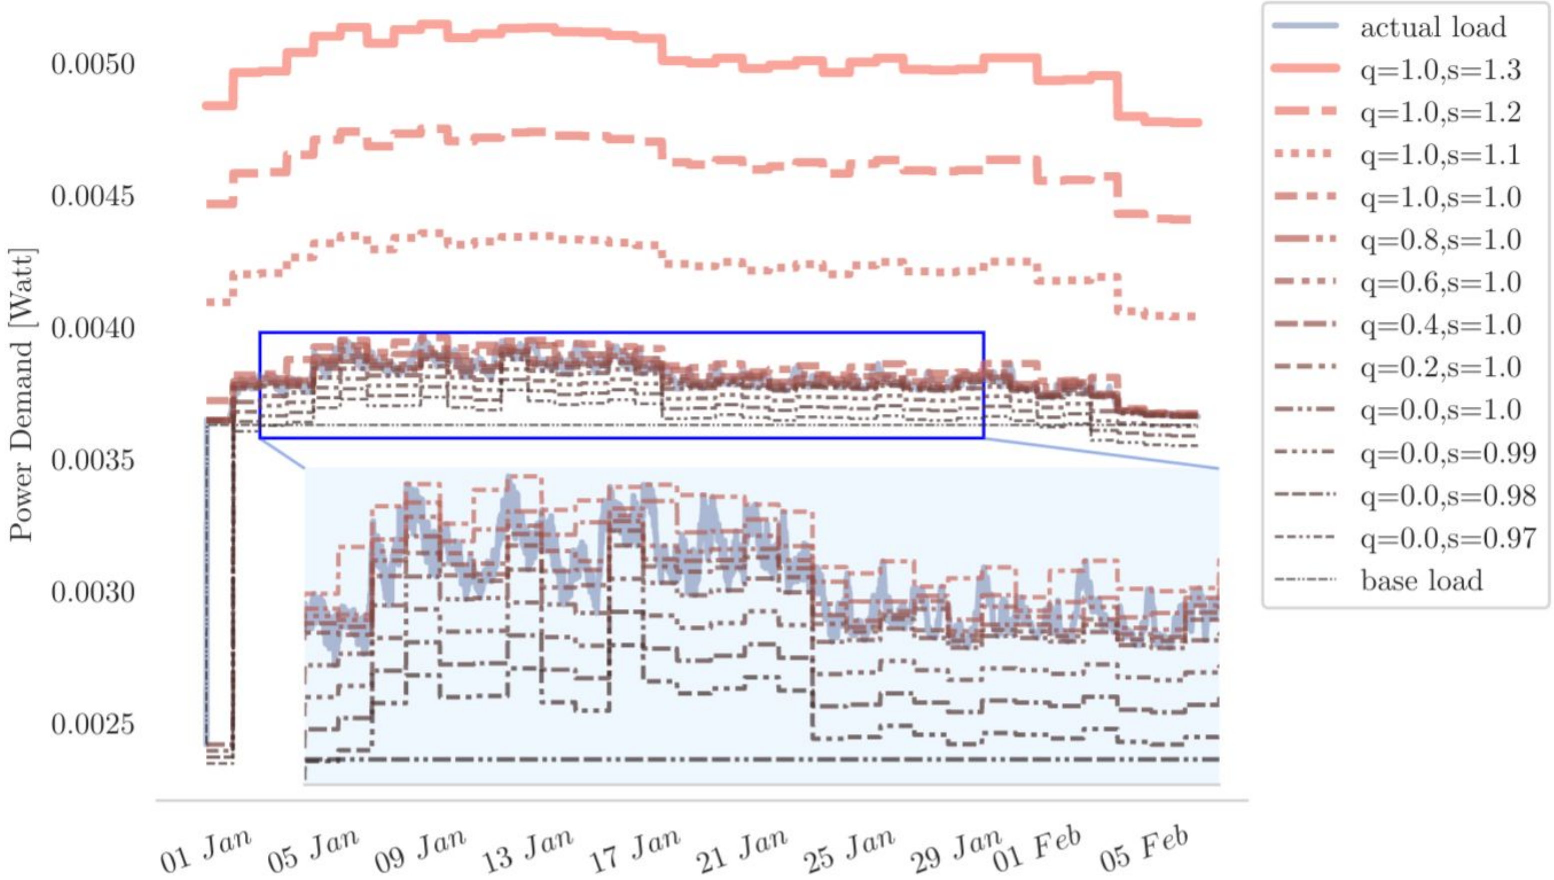

Evaluation: Energy Procurement Strategy (3)

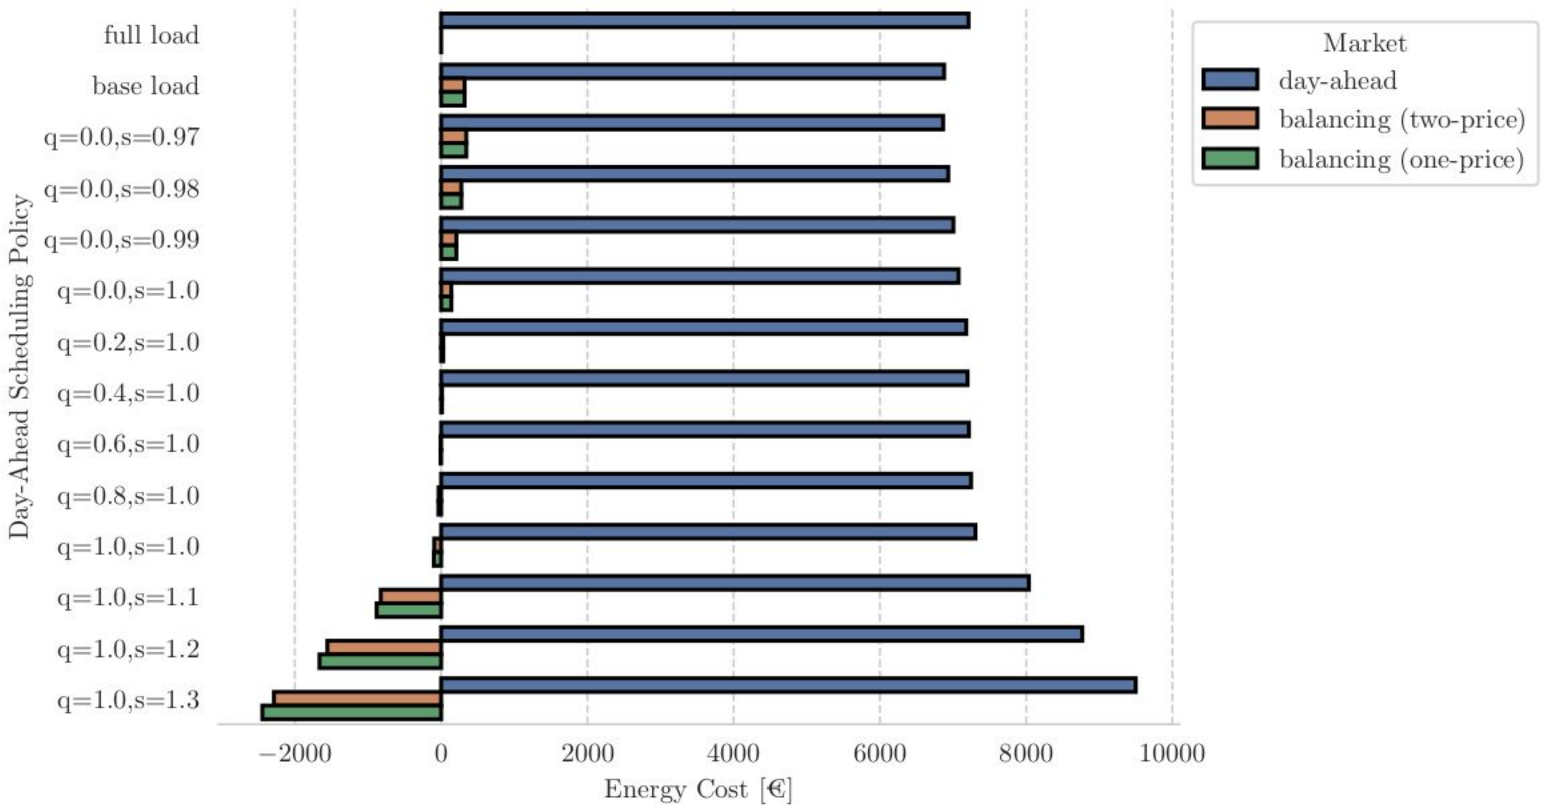

Evaluation: Imbalance Pricing System

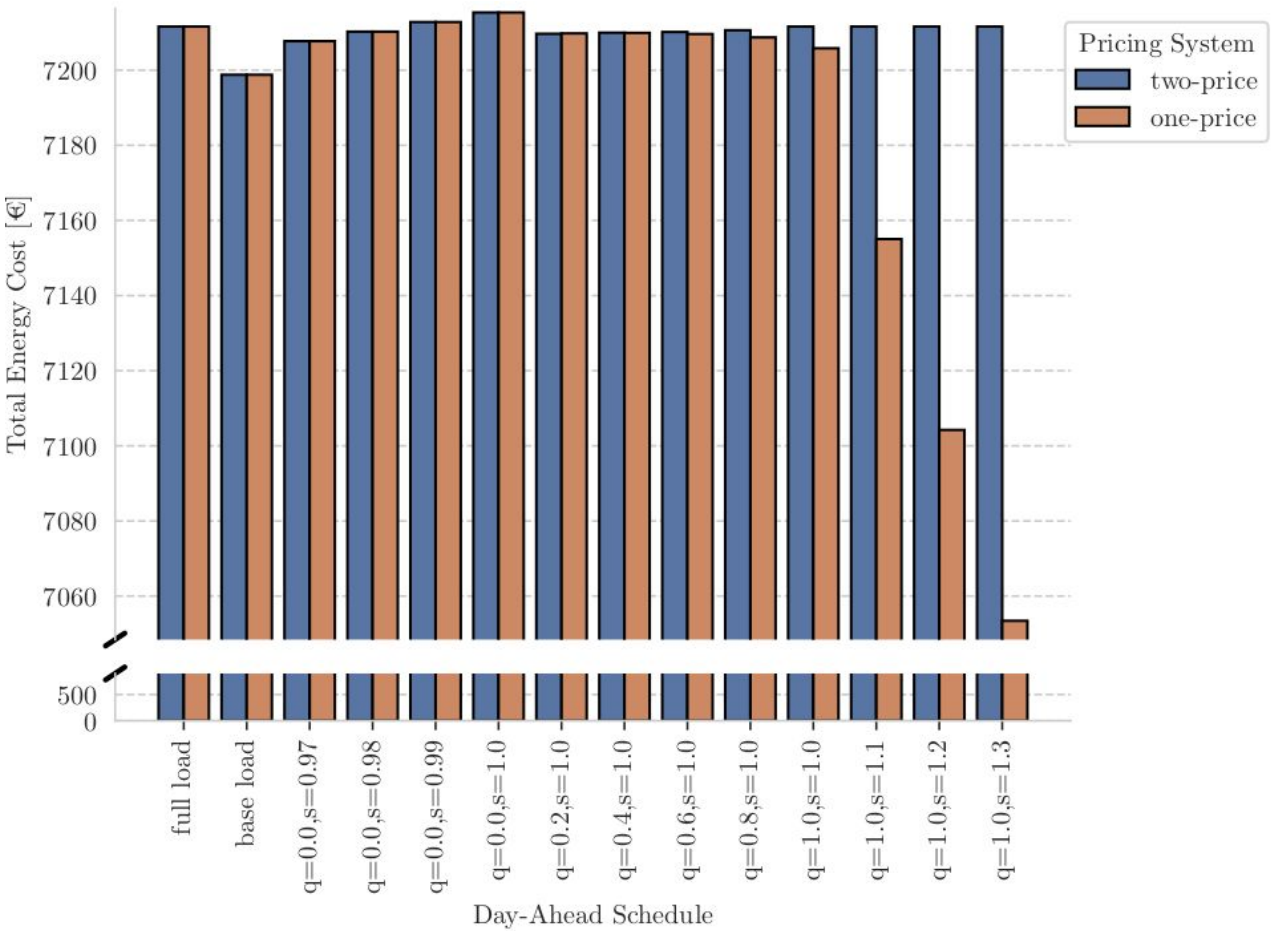

Evaluation: Summary 2

1. The difference is small.
2. The total cost steadily increases when $s > 1.0$.
3. The base-load strategy is preferred.
4. The above conclusions apply also to the new machine model due to **A2**.

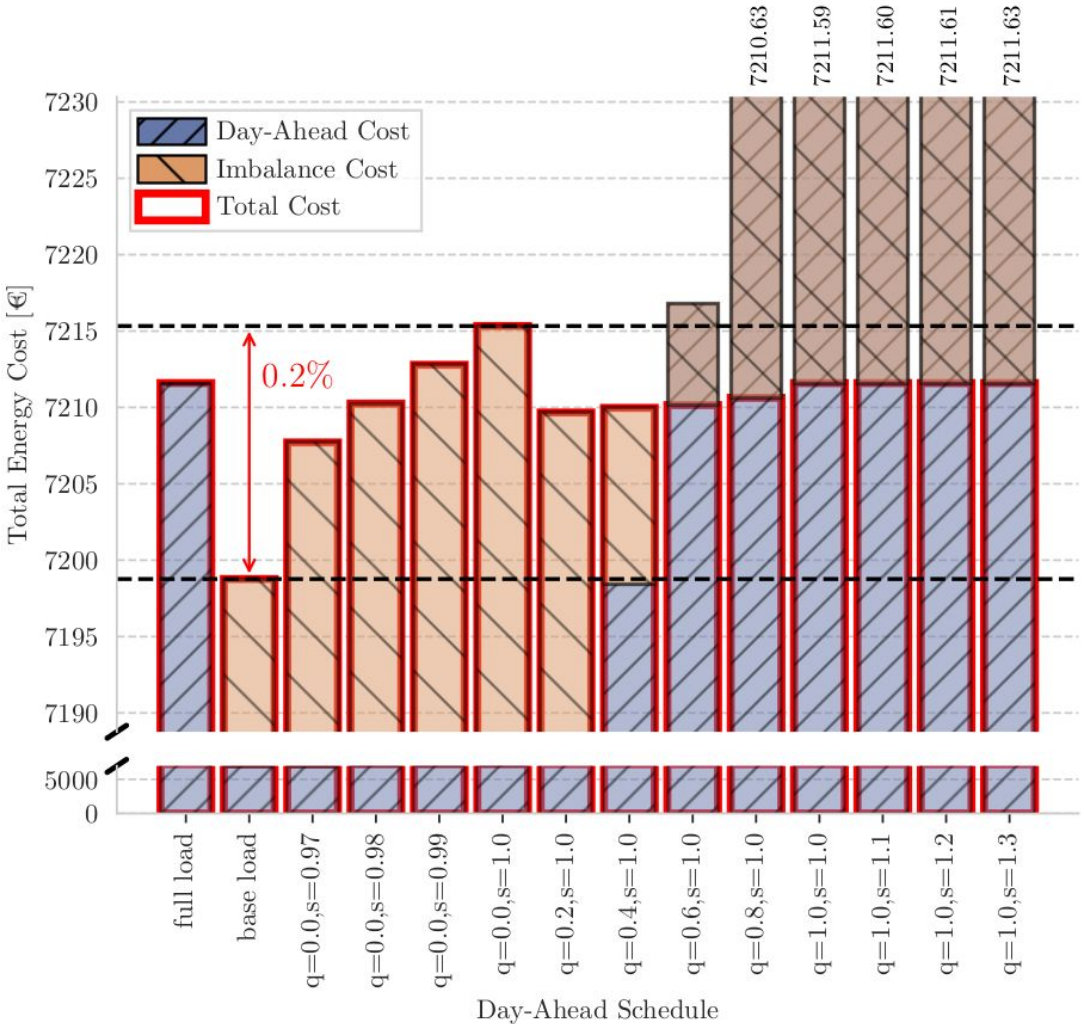

Evaluation: Why Use ML Methods?

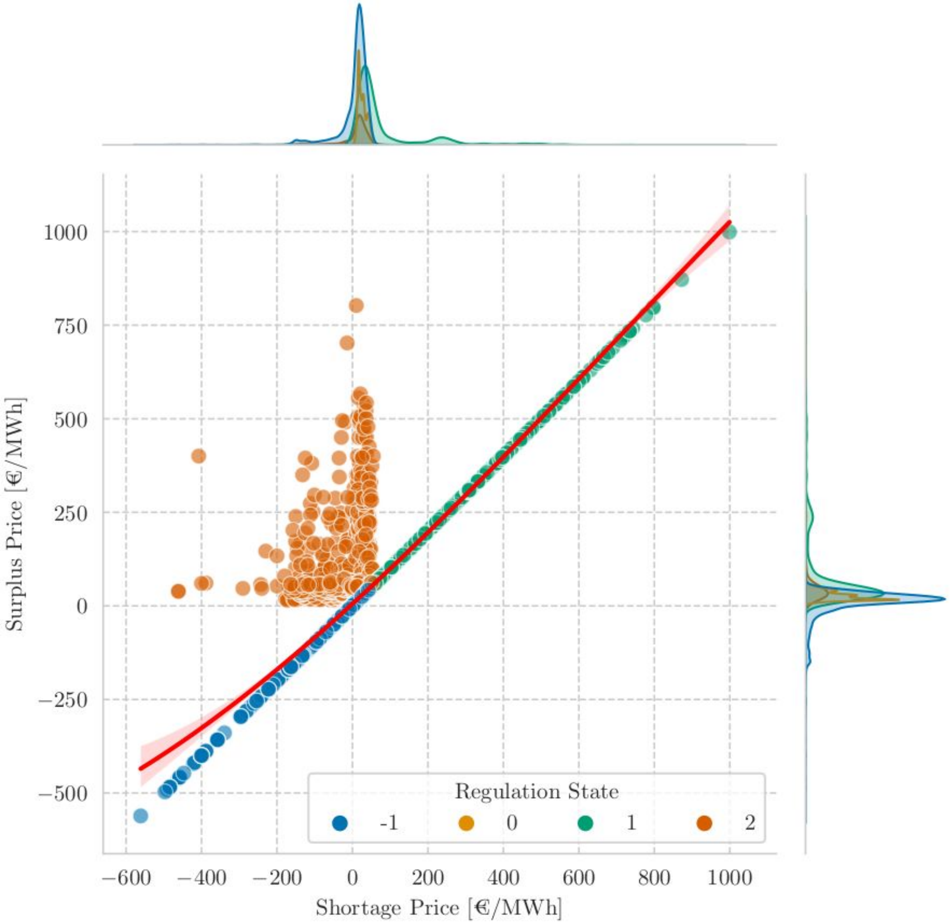

Evaluation: Summary 3

1. There is barely any correlation between prices of the two markets.

   ⇒ Making decisions by heuristics is not feasible.

2. Large profit can be obtained by using (early) ML inferences.

   ⇒ ML methods can be of help in leveraging profits.

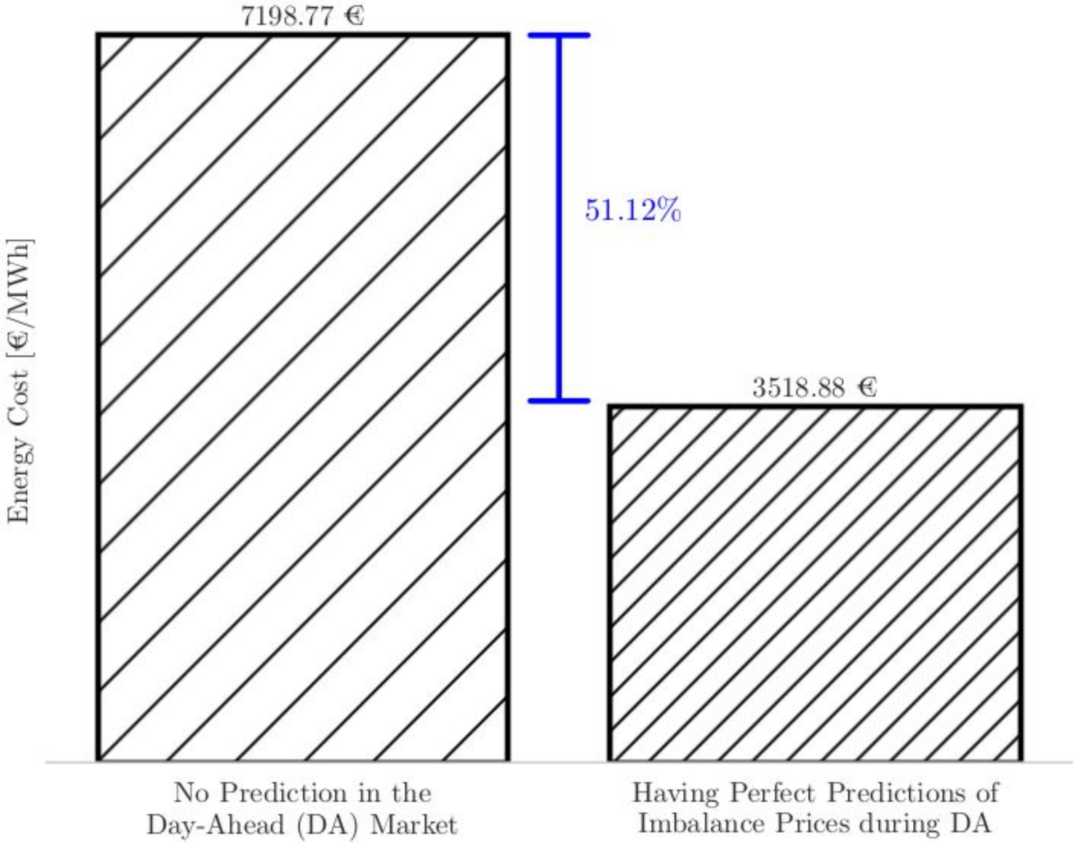

Evaluation: DVFS CPU Usage (1)

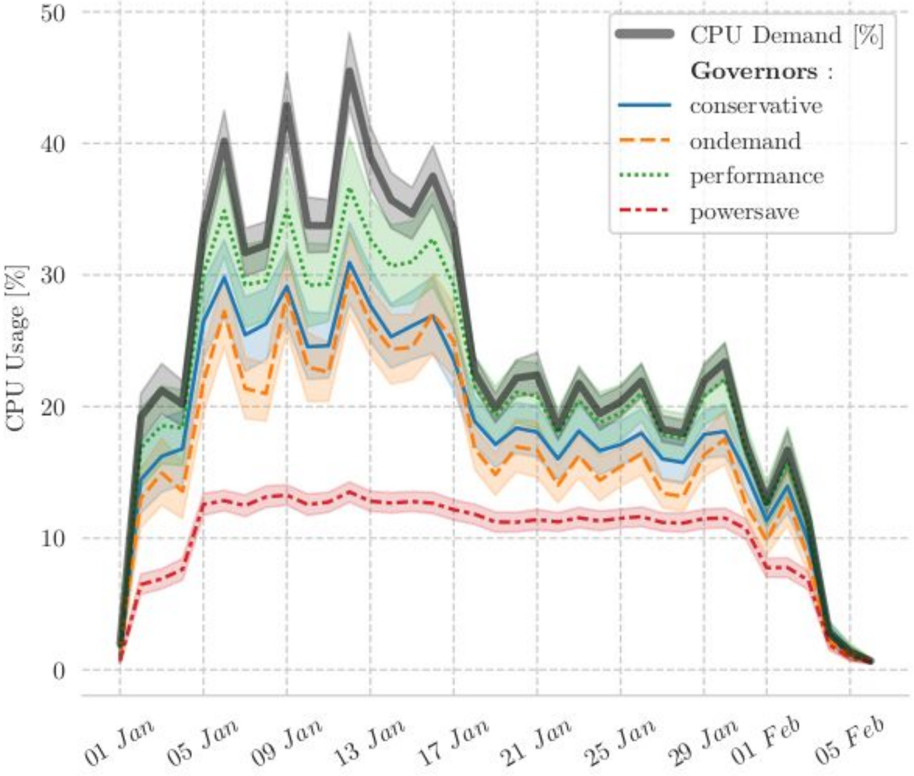

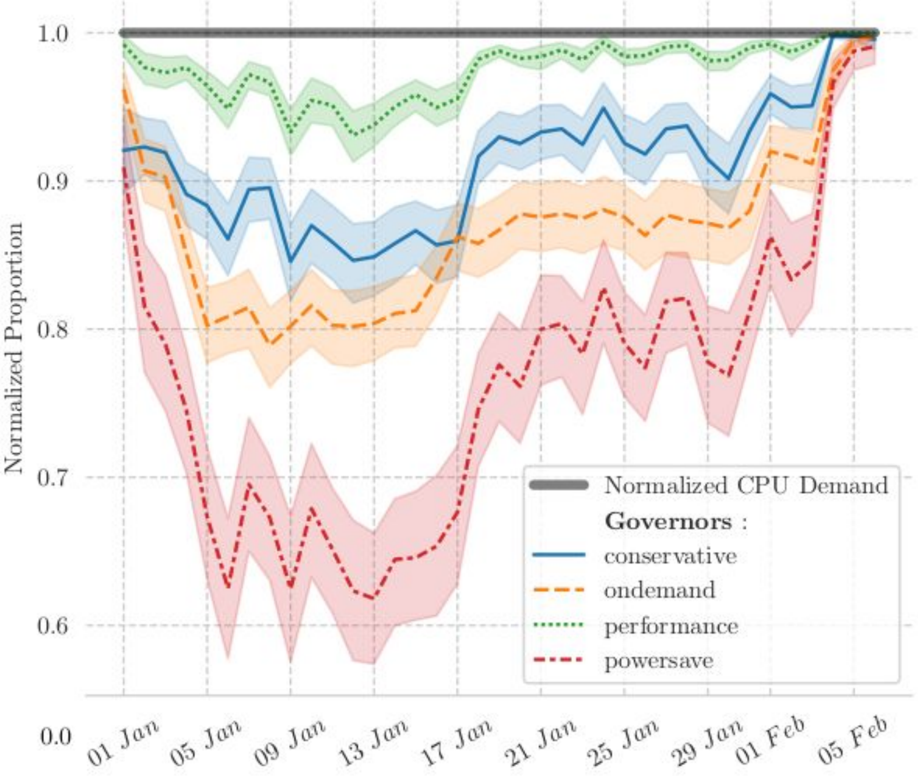

Evaluation: DVFS CPU Usage (2)

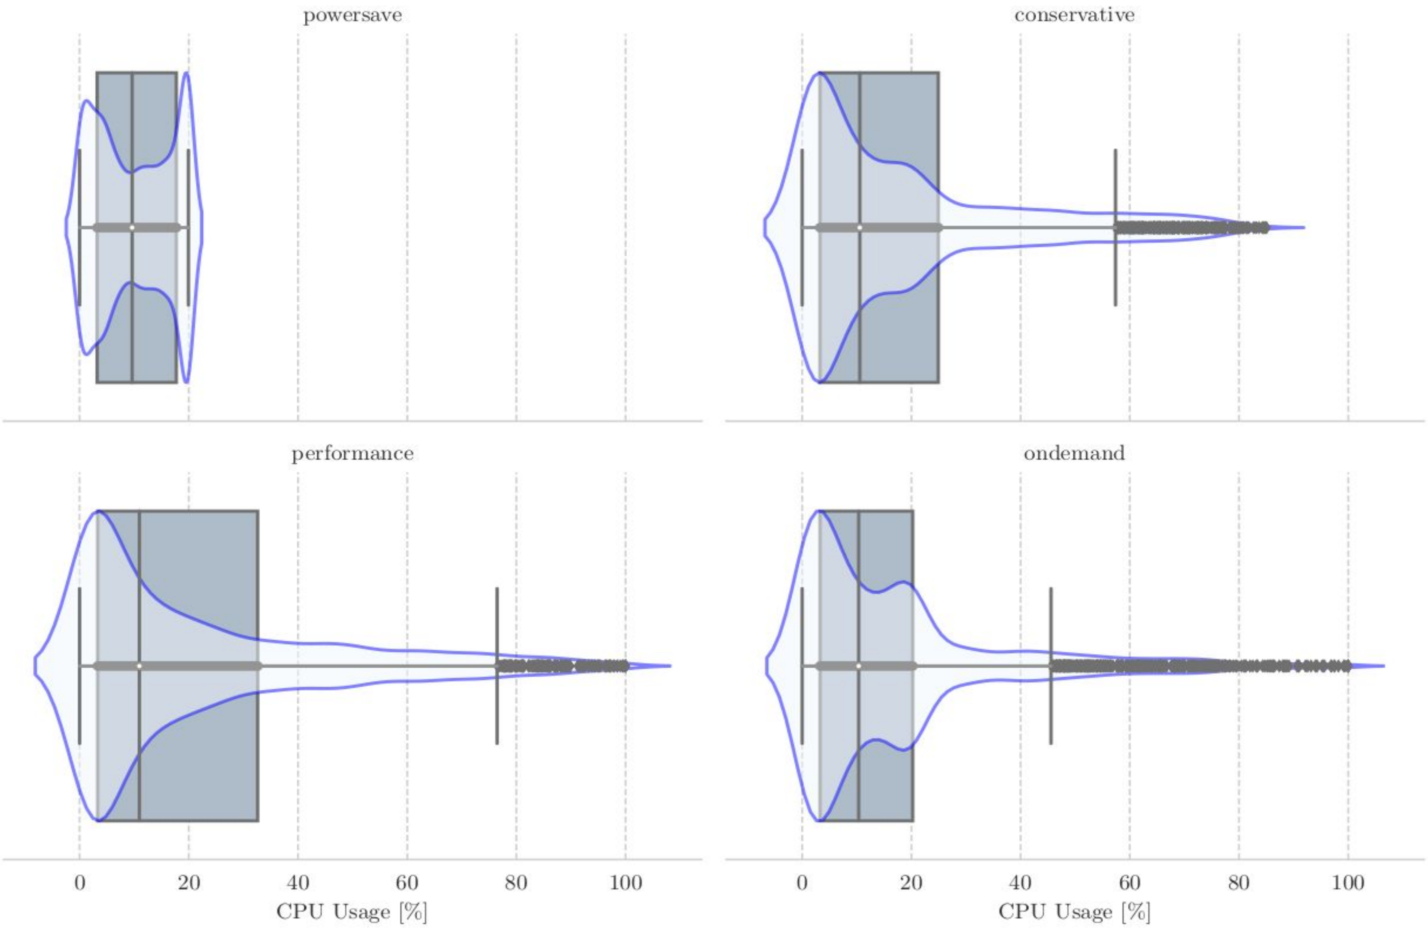

Evaluation: DVFS Power Draw (1)

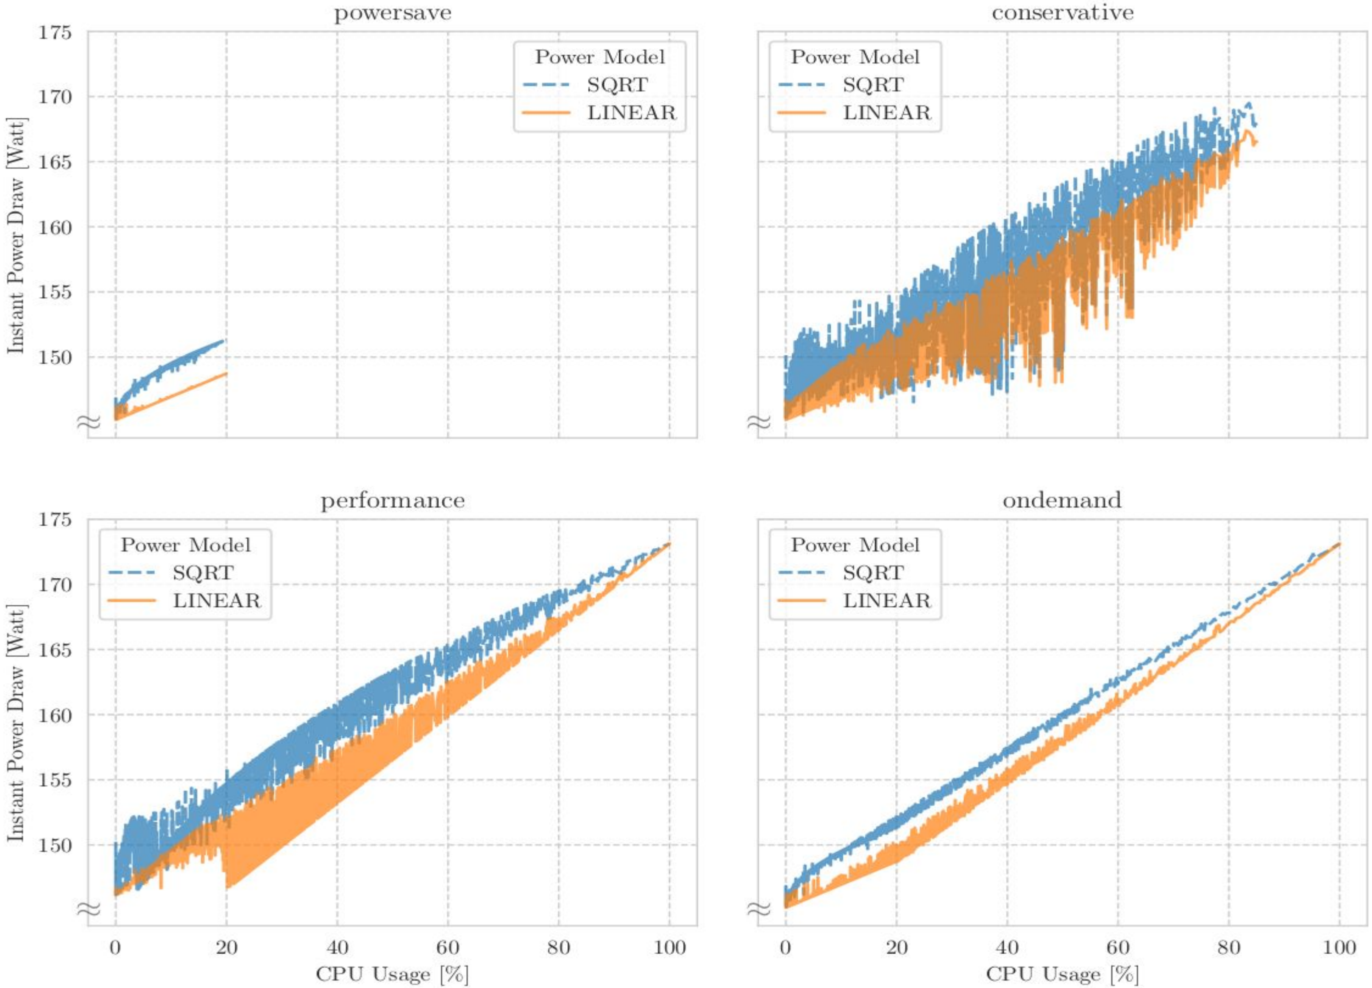

Evaluation: DVFS Power Draw (2)

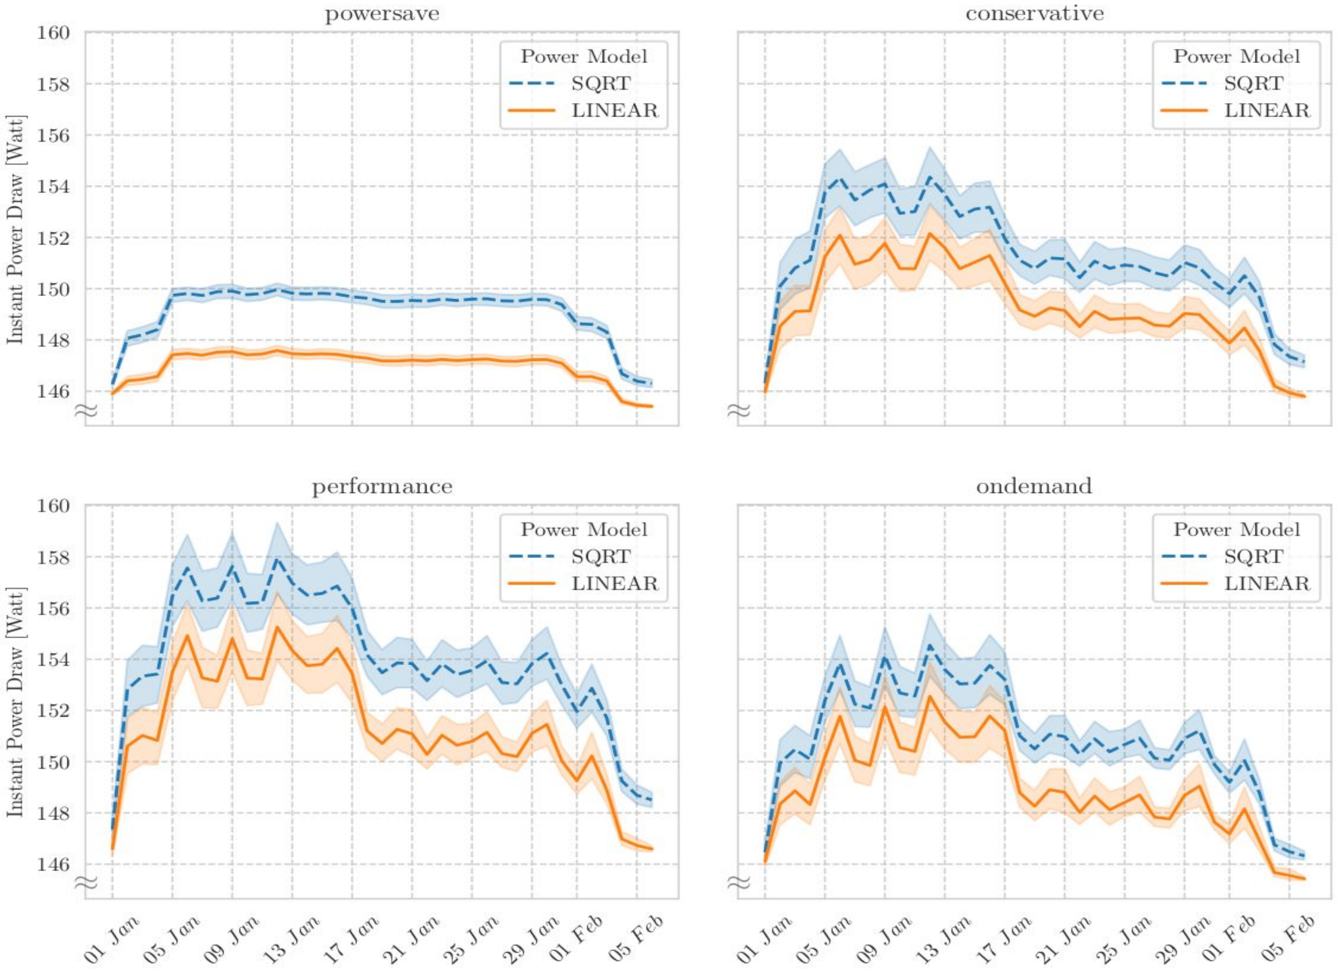

Evaluation: DVFS Over-Commission (1)

1. When hosting VM traces, we do not scale the capacity of each VM according to the frequency changes of the hosts.
2. Instead, we record the over-commissioned CPU cycles.

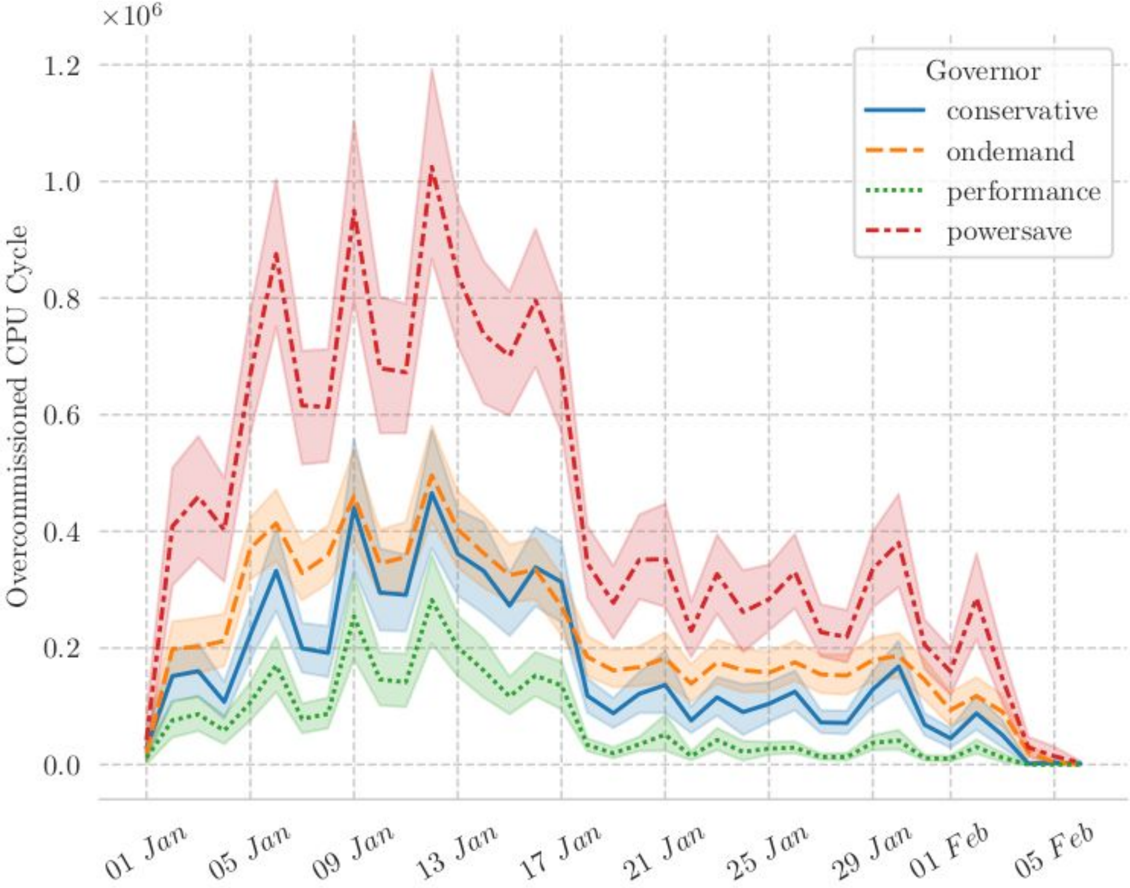

Evaluation: DVFS Over-Commission (2)

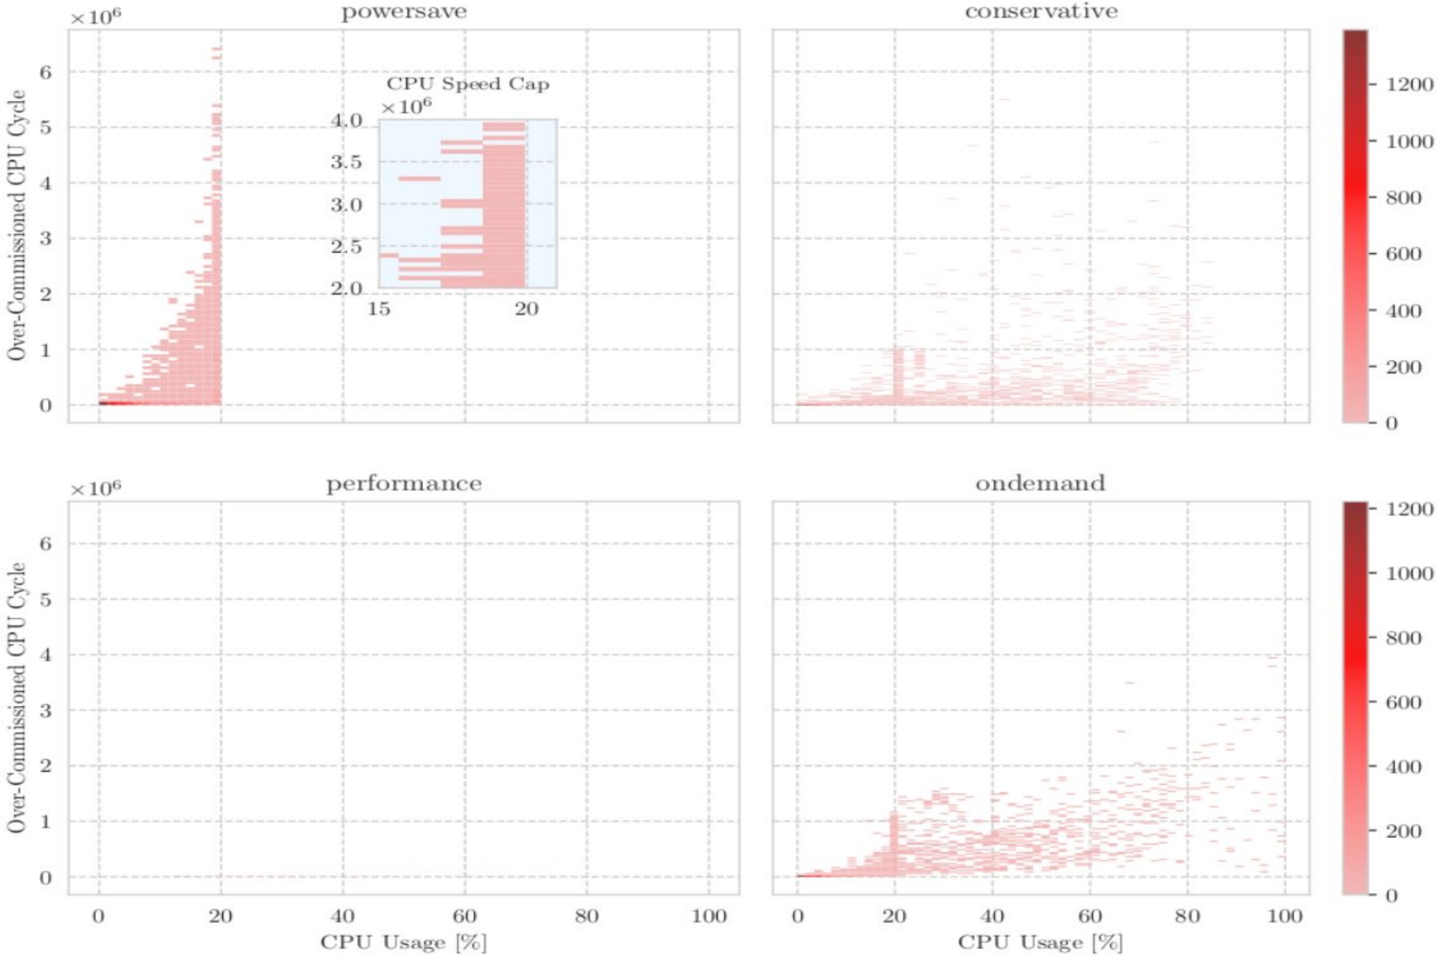

Evaluation: DVFS Over-Commission (3)

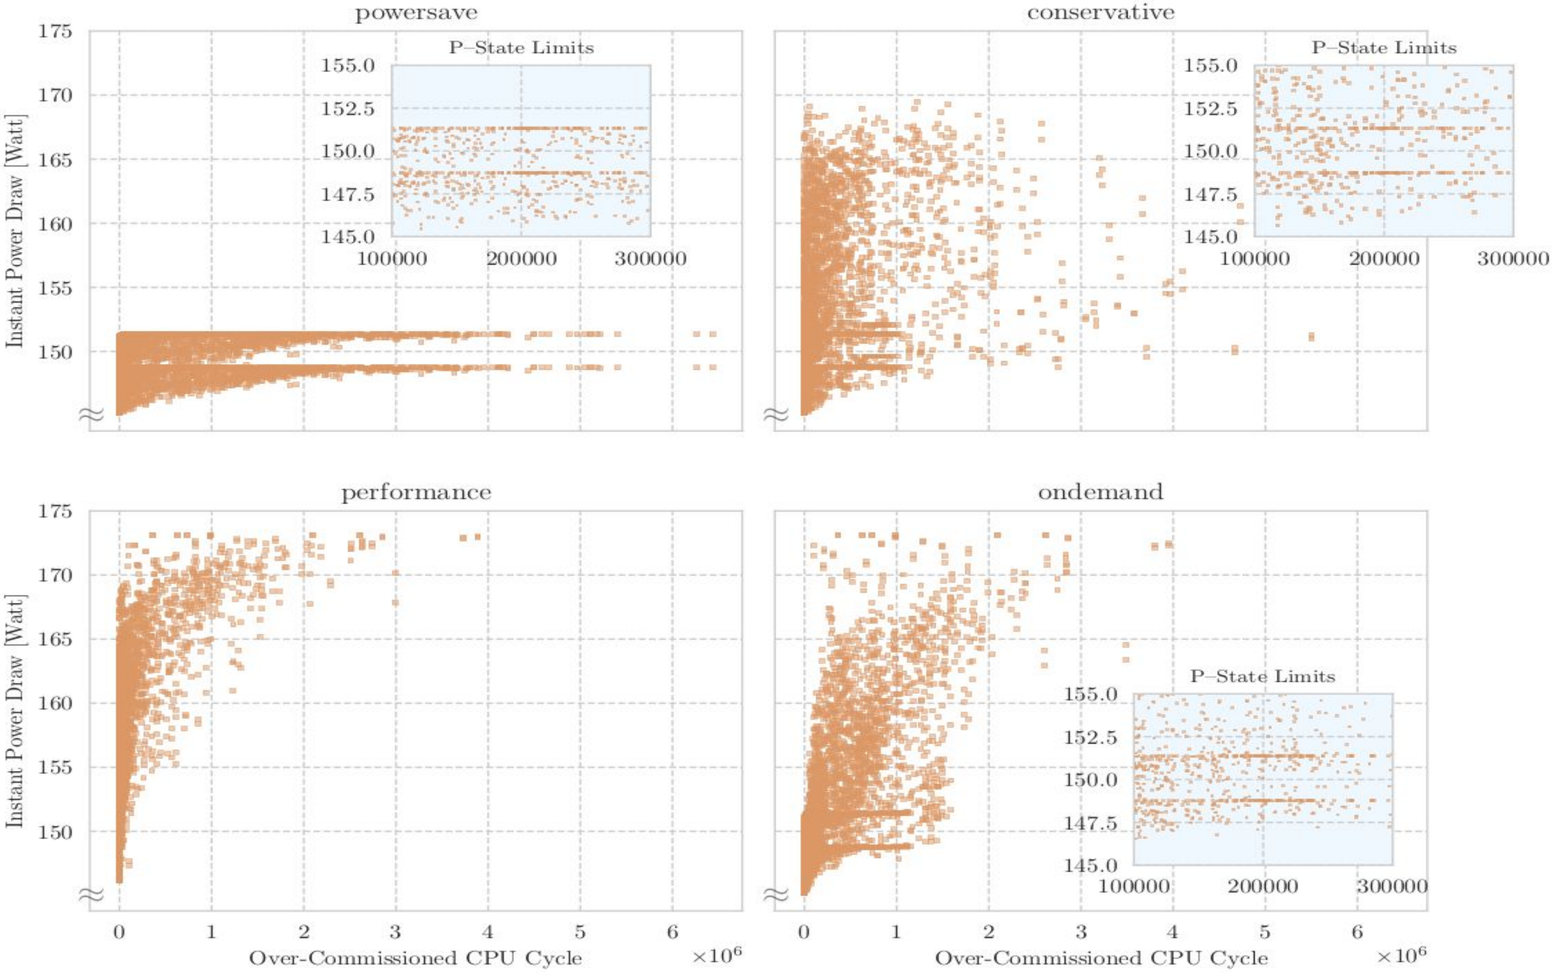

Evaluation: DVFS Consumption of Committed Work

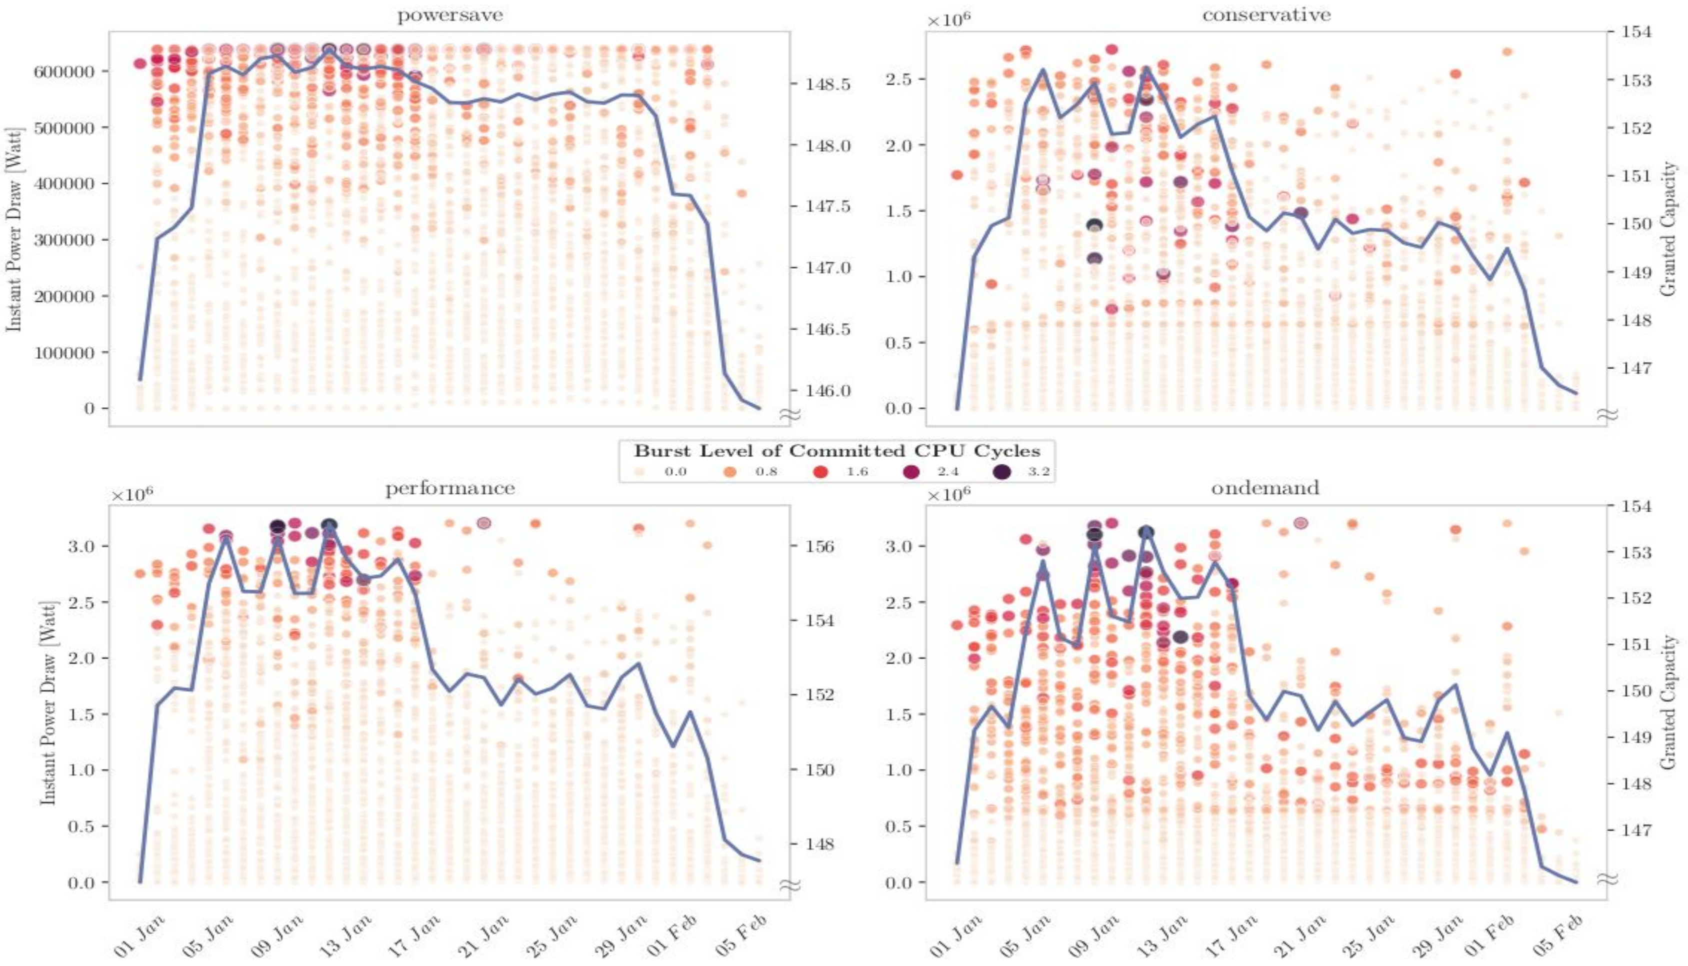

Evaluation: Summary 4

1. The **overhead** of using DVFS is the prolonged execution time, which is captured by the over-commission in our experiments.

2. The **benefit** of saving energy will eventually overrun the overhead, as the execution time is **linearly** scaled to the frequency, whilst the power draw is **quadratically** scaled.

3. The DVFS scheduler should aim to **proactively strike a balance** between the benefit and the overhead.

Evaluation: Scheduler Tuning

1. Finding the *"Sweet Spot"* is <u>not</u> the only way to tune the scheduler.

   ⇒ Users should adjust the factor according to their needs.

2. We set the factor to 12 for the following experiments.

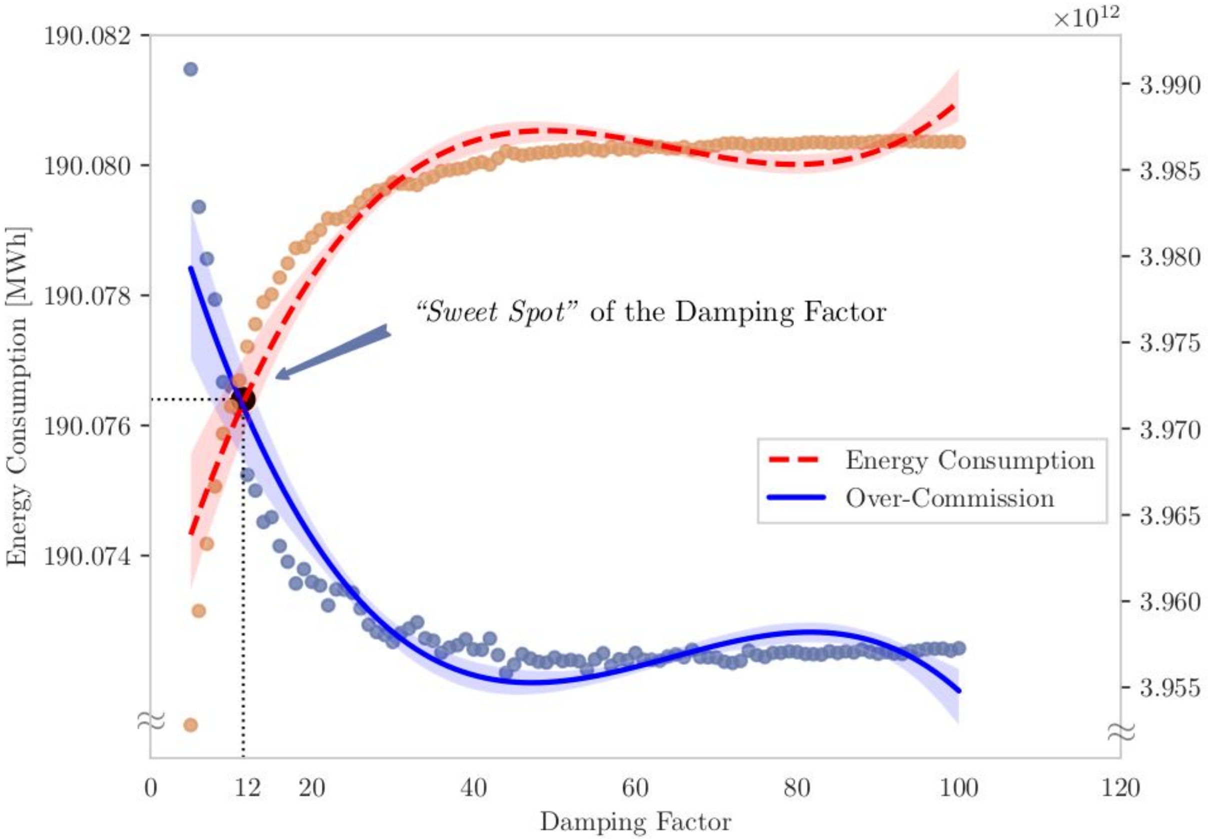

Evaluation: Metric for ML Methods

Timeline of ML prediction:

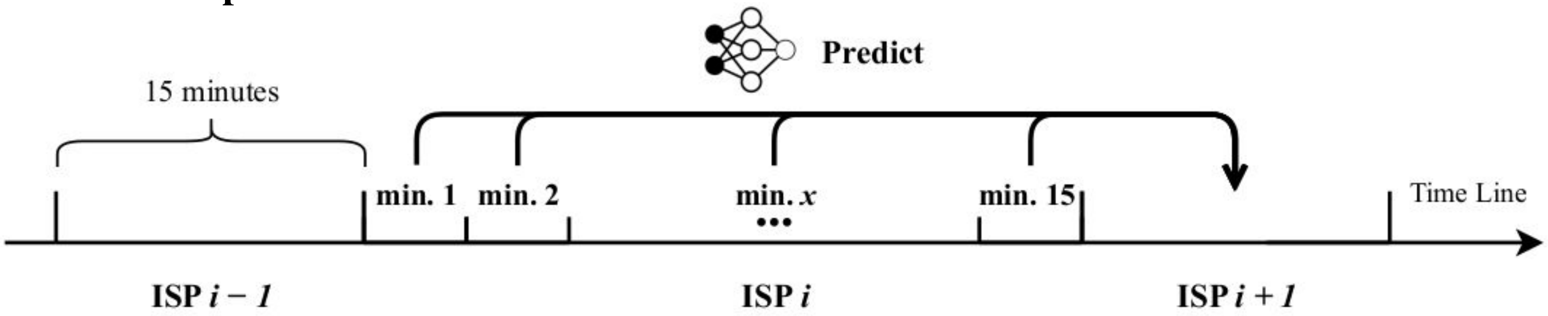

Agreement Accuracy (AA):

N_{ISP}	Number of ISPs
p^S	Spot price in the day-ahead market
p^B	Shortage price in the balancing market
p^F	Forecasted imbalance price

$$\mathbb{S}(x) = \begin{cases} +1 & x > 0 \\ 0 & x = 0 \\ -1 & x < 0 \end{cases} \tag{5.12}$$

$$AA = \frac{\sum_i^{N_{\text{ISP}}} \mathbb{1}\left\{\mathbb{1}\left[\mathbb{S}(p_i^B) = \mathbb{S}(p_i^F)\right] = \mathbb{1}\left[\mathbb{S}(p_i^B - p_i^S) = \mathbb{S}(p_i^F - p_i^S)\right]\right\}}{N_{\text{ISP}}} \tag{5.13}$$

Evaluation: Synthetic Predictors

Adding Gaussian errors:

$$p^f = p^B + E,$$

$$E \sim \mathcal{N}(0, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \cdot e^{-\frac{1}{2}\left(\frac{x}{\sigma}\right)^2}$$

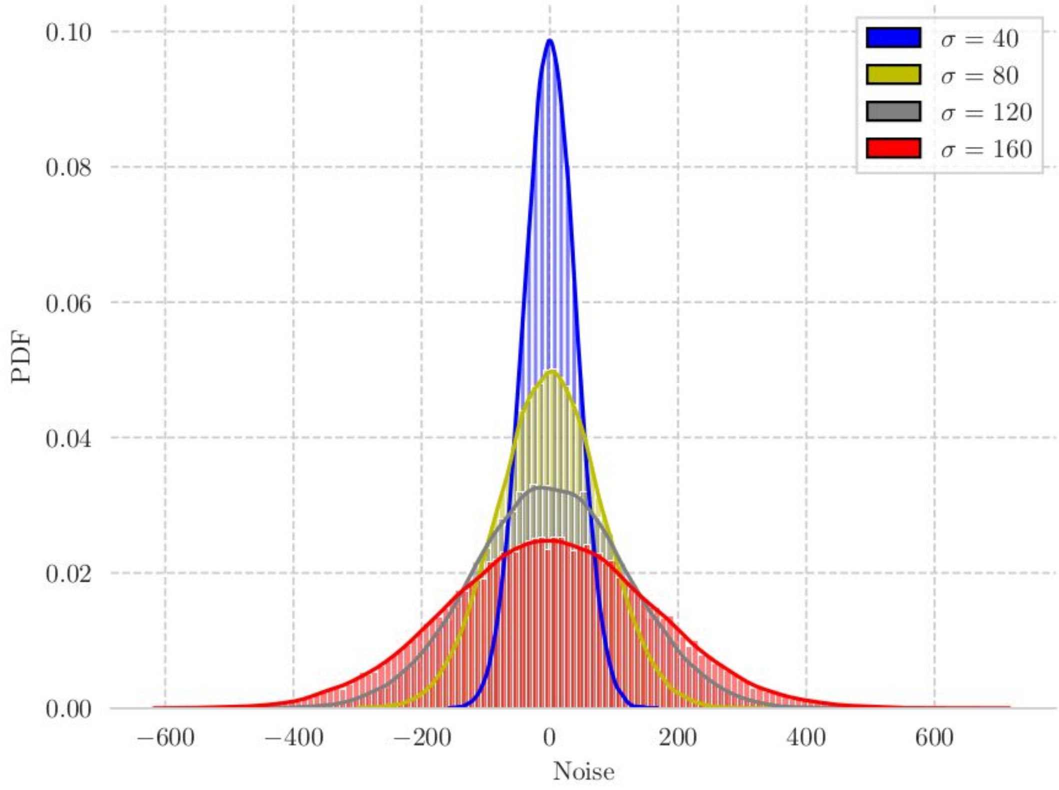

Evaluation: Bounded Comparison (1)

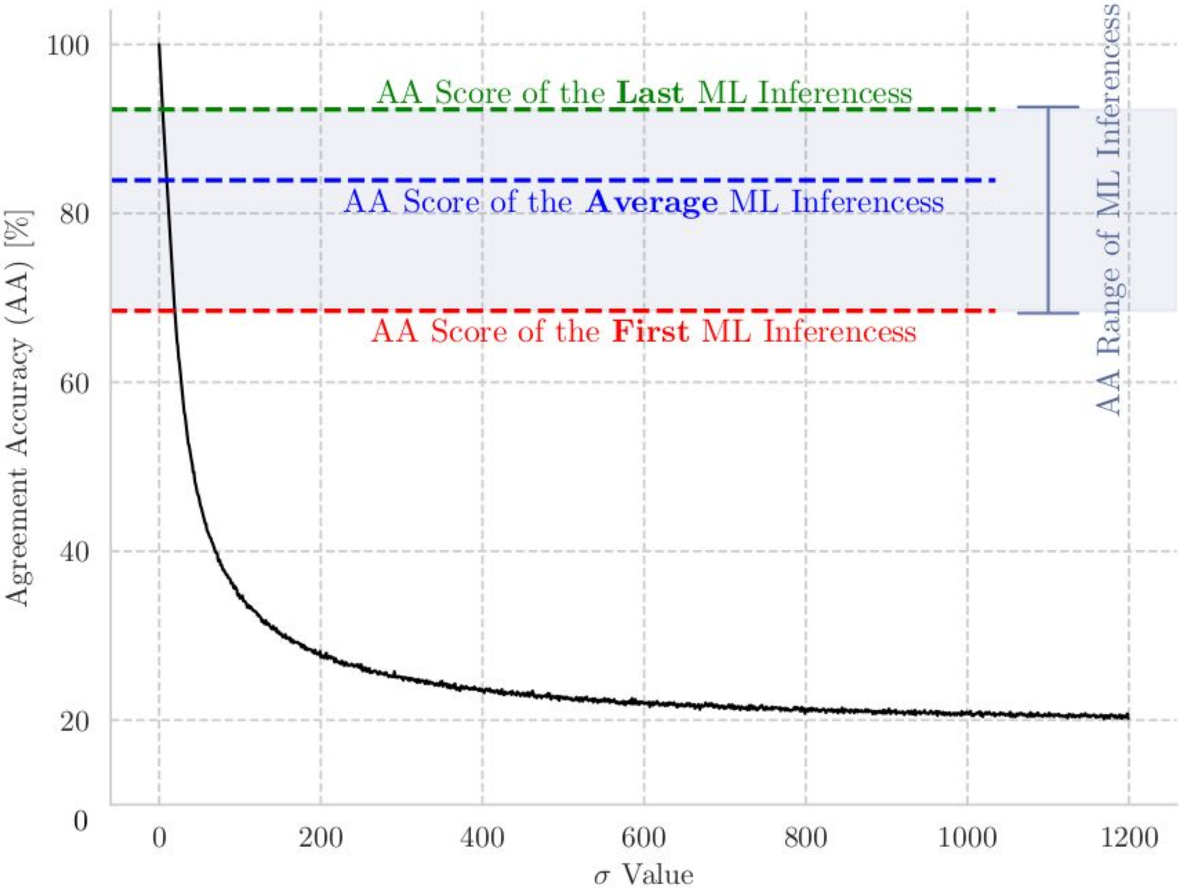

Evaluation: Bounded Comparison (2)

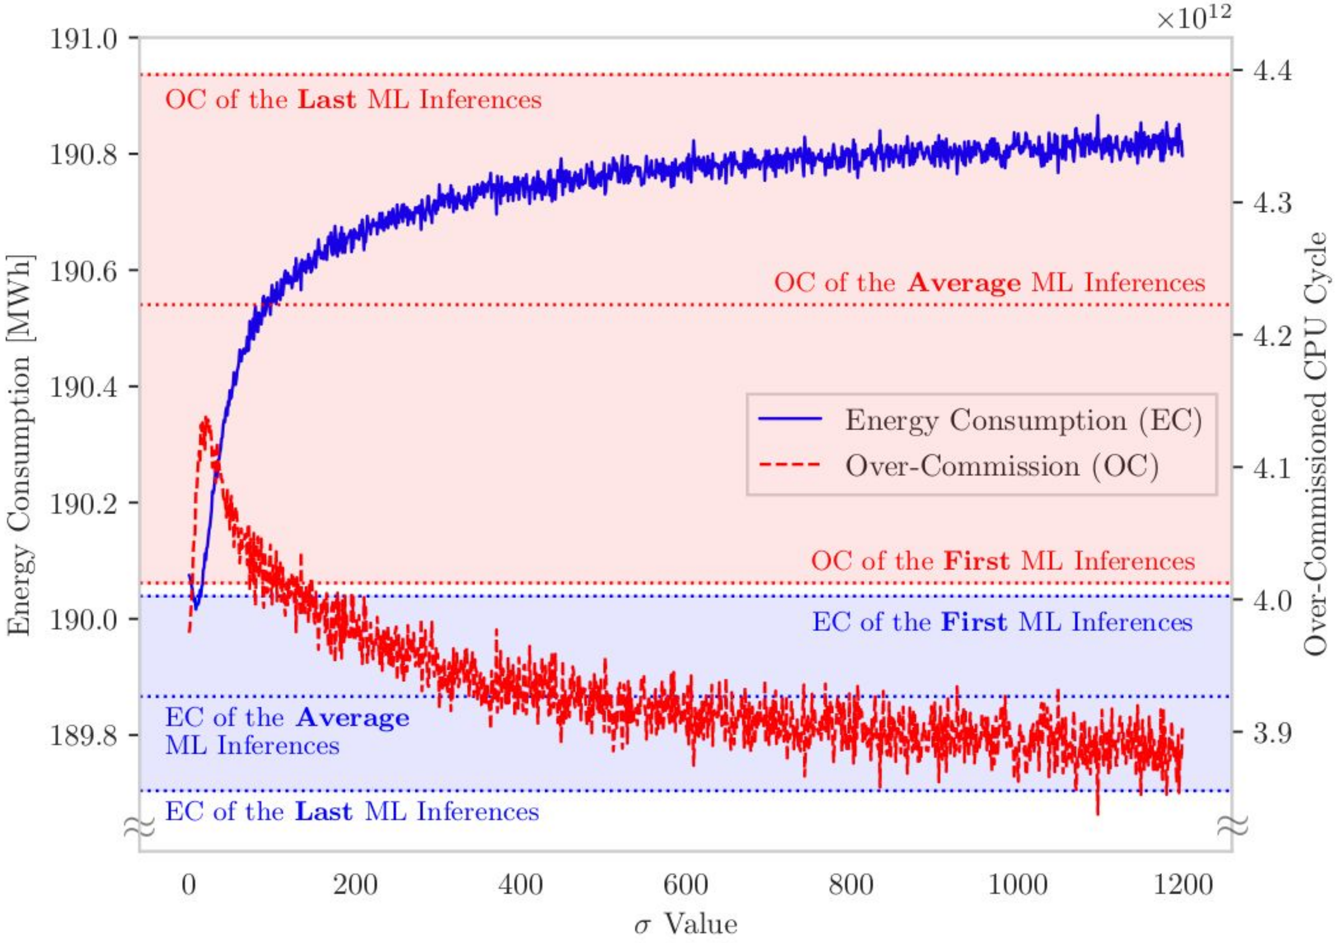

Evaluation: Bounded Comparison (3)

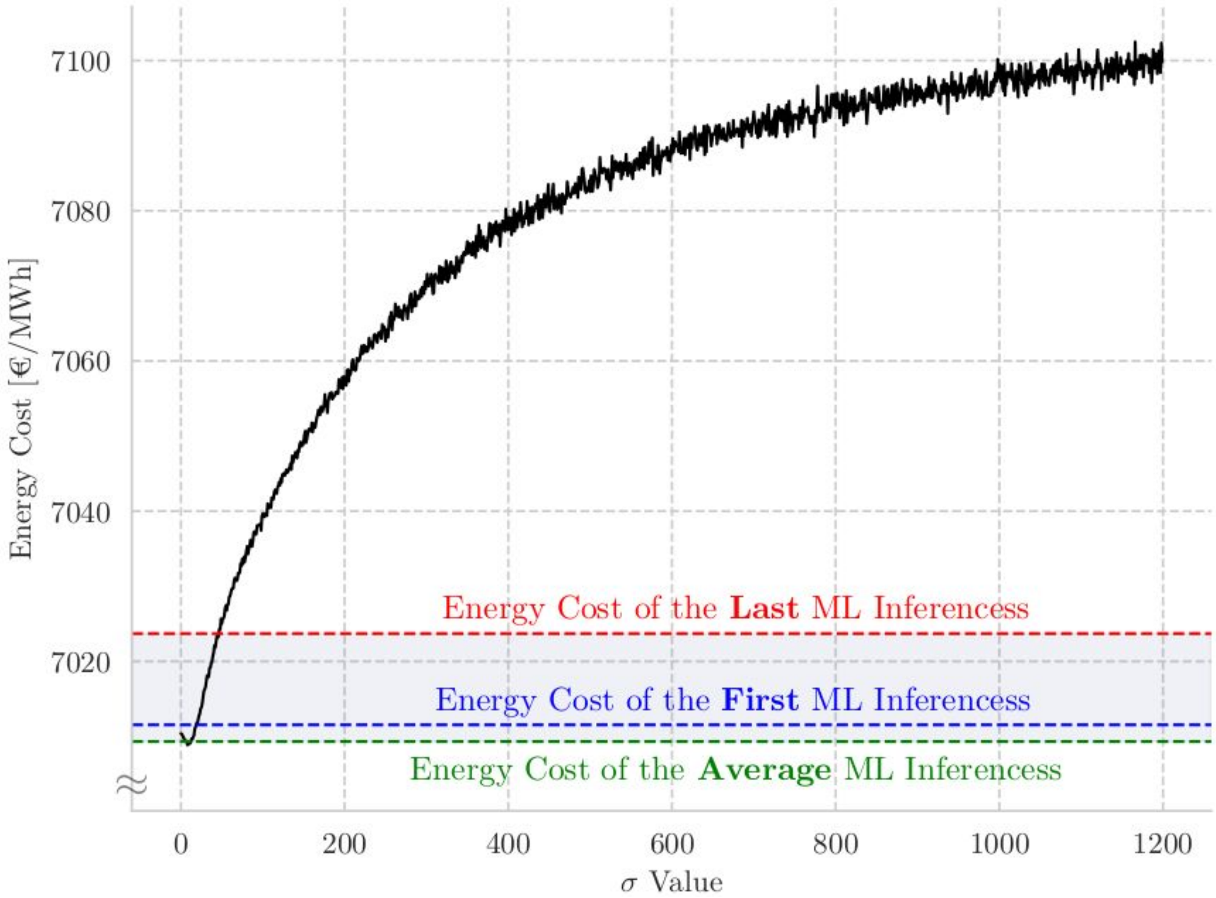

Evaluation: Indirect DR

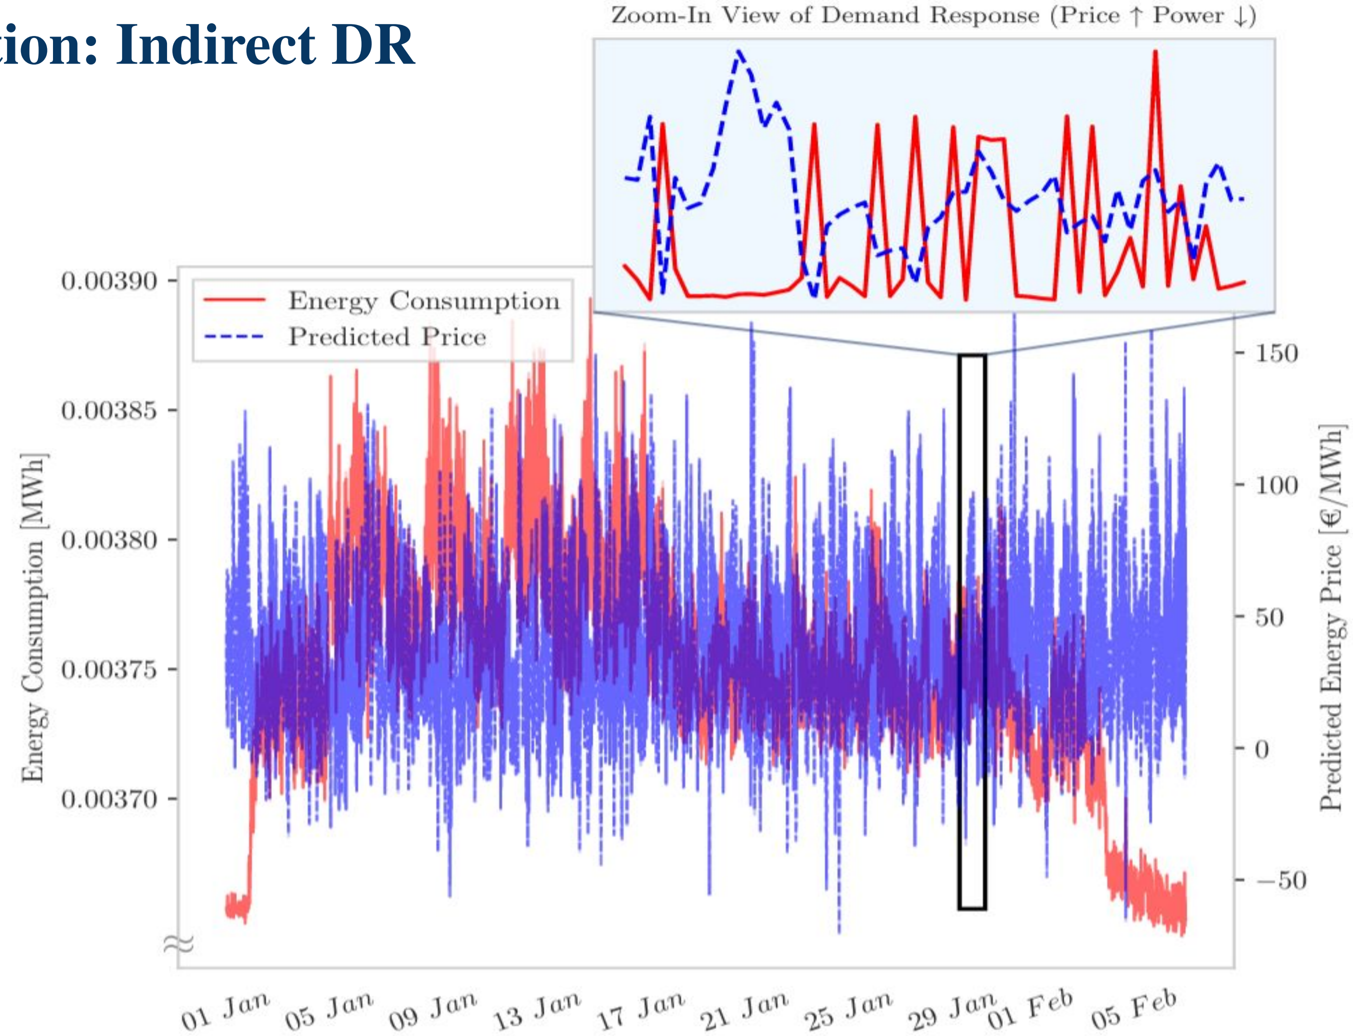

Evaluation: Total Energy Saving

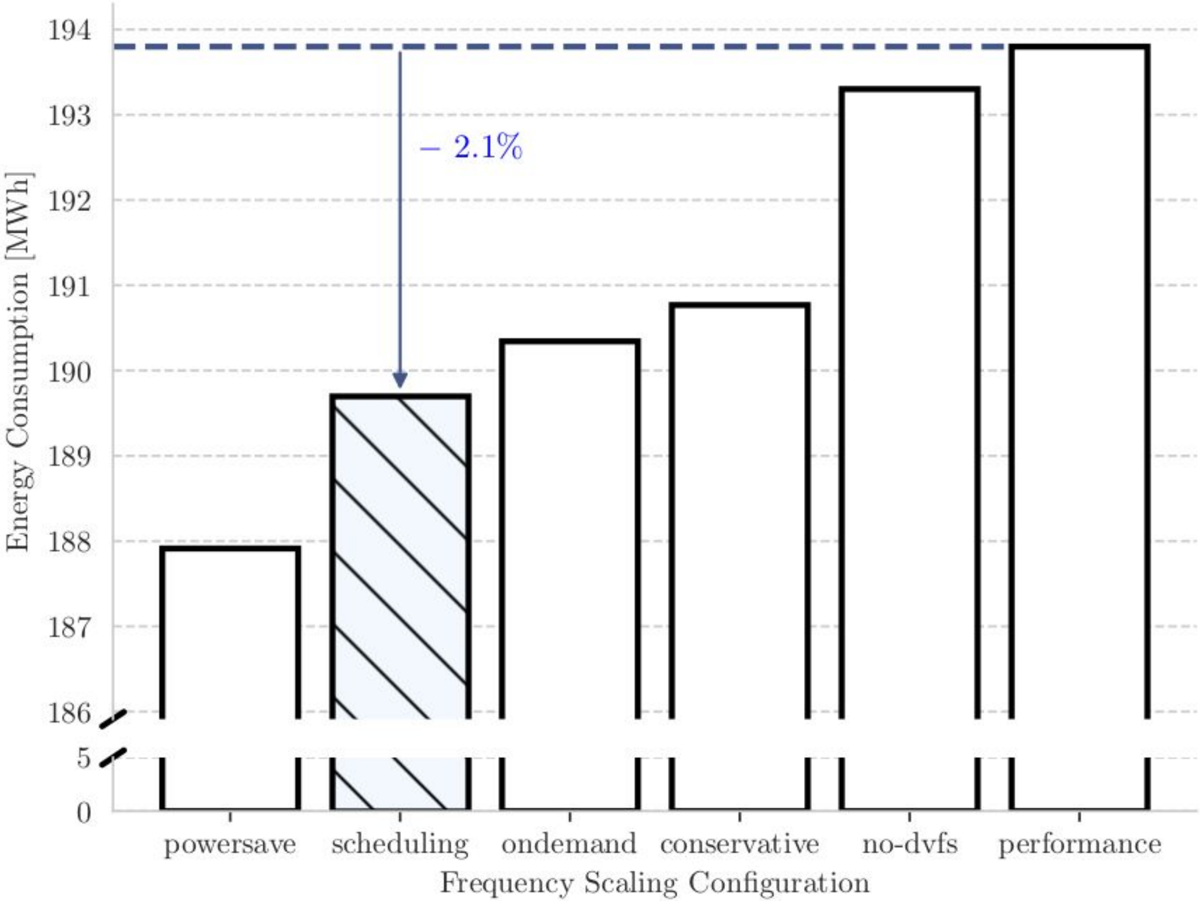

Evaluation: Total Over-Commission Improvement

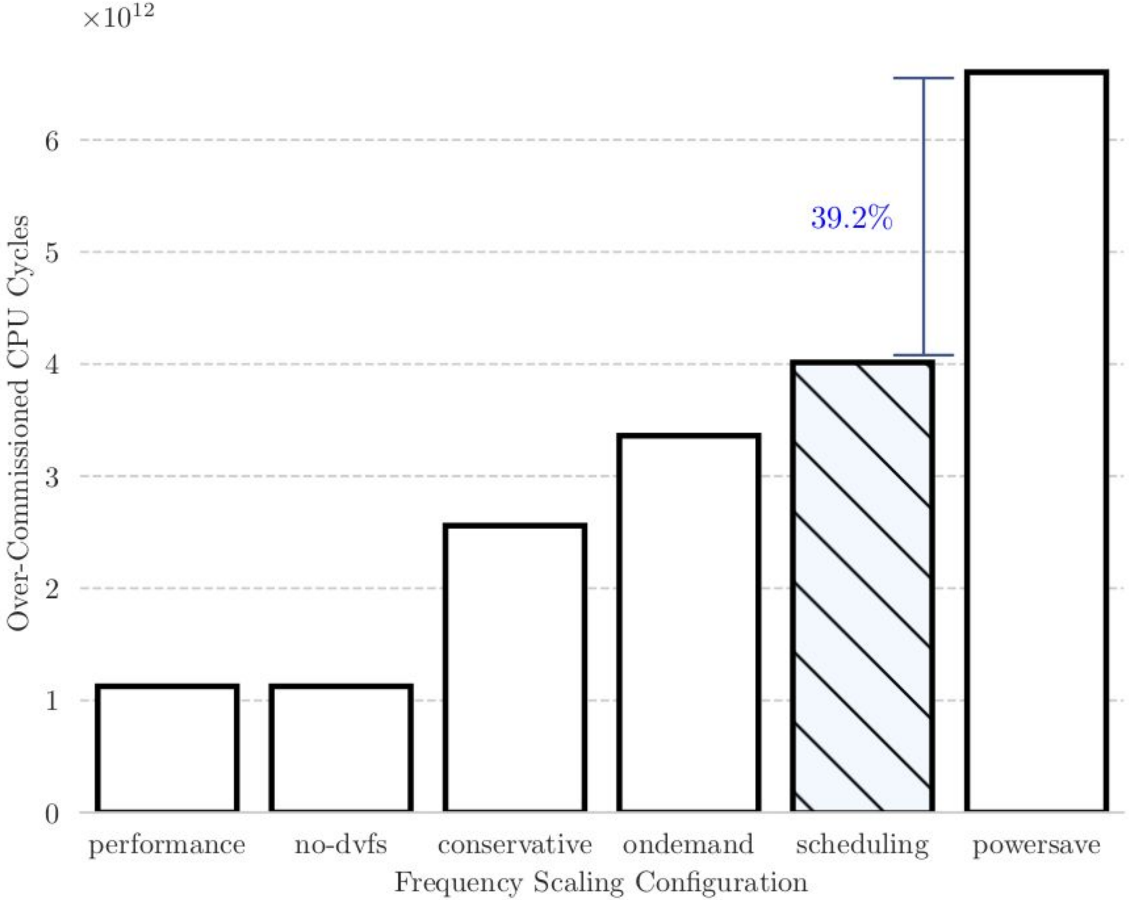

Evaluation: Total Financial Benefit

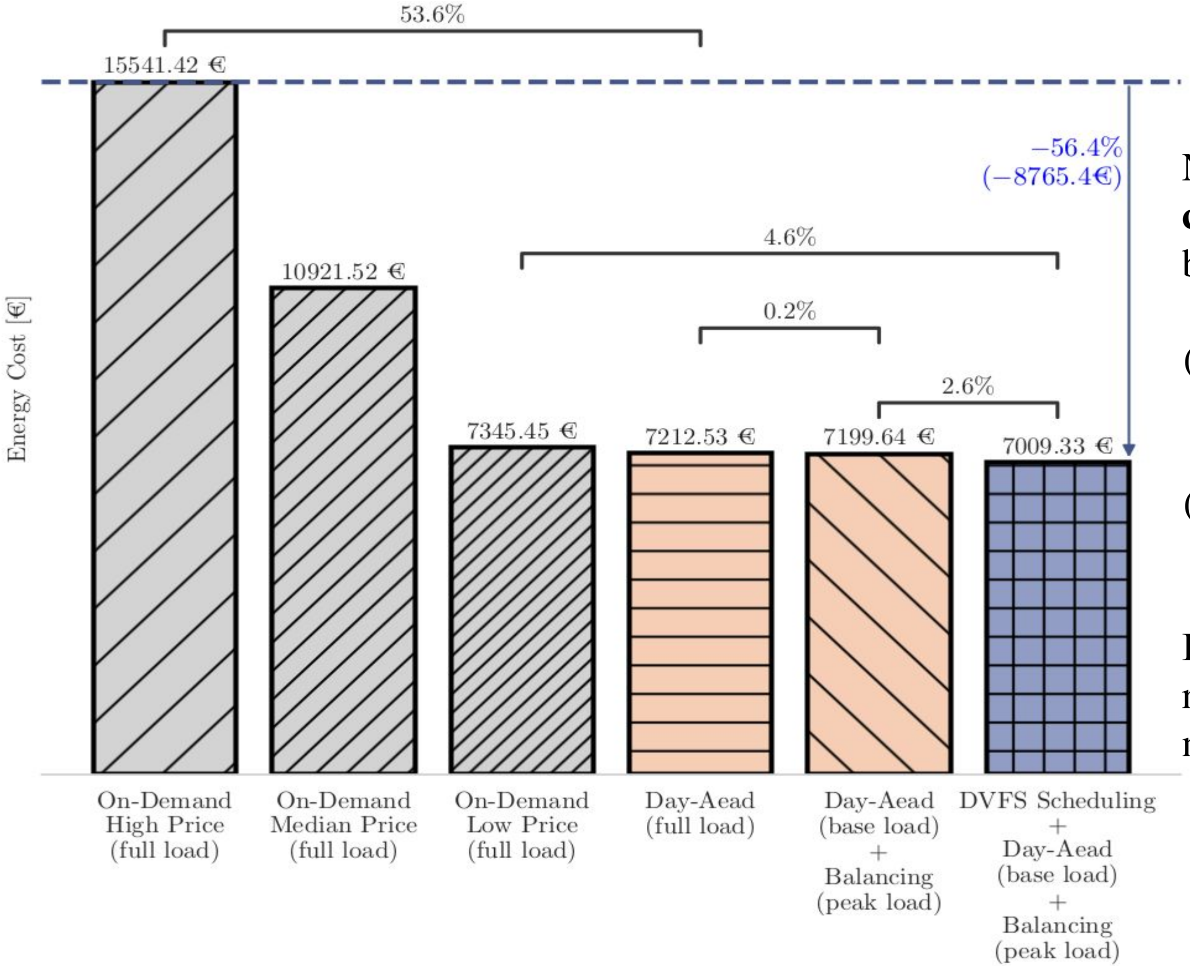

Note that we believe these results are **conservative**, since they are produced by the old machine model where

(1) The peak load only takes up a tiny portion of the overall power load.
(2) The variations in power load between energy states is small.

In other word, we would expect the results to be **more significant** if a newer machine model were available!

Agenda

Conclusion: Answering RQs

RQ1: OpenDC now is able to model the DC power system in a ***flexible*** and ***highly customizable*** way

RQ2: There is a ***strong financial incentive*** for DCs to participate in the energy market.

RQ3: The ***base-load procurement strategy*** is preferred.

RQ4: The ***proactive DVFS scheduler*** powered by the ML methods that can reduce energy cost and consumption whilst constraining the over-head of DVFS.

RQ5: Both the simulator and its extension are ***ready-to-use*** and ***user-friendly.***

Conclusion: Future Work — A New Research Line!

1. How feasible and beneficial is it for individual DCs to serve as BSPs instead of BRPs?
2. How can we develop a convenient and liable tool to measure P-state consumption levels, which will enable us to use more machine models?
3. What is the impact of employing not only the energy price but also the power grid frequency in DCs' decision-making?
4. How to improve the algorithms of resource allocation (e.g., power distribution, VM scaling/placement) in response to market signals?
5. What is the effect of core-level P-state frequency scaling on DC energy consumption?
6. How can we improve the market design to incentivize DCs to participate?
7. How can we improve the design of the energy market to incentivize datacenters' active participation?
8. How to orchestrate redundancies (.e.g., PSU, UPS, etc.) to provide DR?

And more …

Acknowledgement

I am grateful for all the help from my supervisors who actually made all these happen!

As a '18 HP student, I appreciate all the training and support from our AtLarge Research group over the past three years!

Individual DCs can and should **directly** participate in the energy market to provide the power grid with **indirect DR**, whilst both **saving their energy costs** and **curbing their energy consumption**.

Thank you!

OpenDC.org

hongyu.nl
hongyu.he@vu.nl
github.com/hongyuhe
www.linkedin.com/in/hongyuhe