DIGITAL TWINNING DATACENTERS

OPENDT: EXPLORING DATACENTER PERFORMANCE AND SUSTAINABILITY WITH A **SELF-CALIBRATING DIGITAL TWIN**

@Large Research
Massivizing Computer Systems

These slides:
bit.ly/opendt-hcp


Radu Nicolae


Jules van der Toorn


Stavriana Kraniti


Houcen Liu


Alexandru Iosup

Digital twinning = monitoring + datagen + ODA +modsim platform + goal-oriented infr. steering

VU VRIJE UNIVERSITEIT AMSTERDAM

FUTURE 6G NETWORK SERVICES

SURF

GRAPH MASSIVIZER

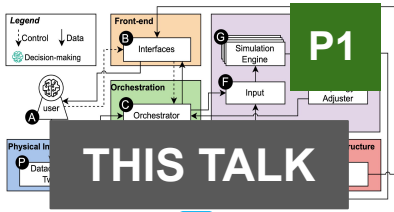

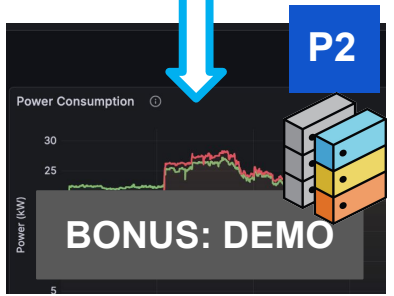

“THIS IS THE GOLDEN AGE OF DATACENTERS”

Datacenters are Crucial to Today's Society

Inside a Google Datacenter

Source: Google (online)

Datacenters are widely used, worldwide, amongst academia, government, and industry.

VU VRIJE UNIVERSITEIT AMSTERDAM
TU Delft
Massachusetts Institute of Technology
amazon web services
NETFLIX
zoom

SUSTAINABILITY	ECONOMICAL	PERFORMANCE
3-5% of global electricity (**2025**) Estimated **12% (2035)**	Impacts >60% of NL GDP Impacts 3.3 million jobs	Modern day Moore's law? (LLMs scale 2x/6 mo)

Sources: Iosup et al., Massivizing Computer Systems, ICDCS 2018 [Online] / Dutch Data Center Association, 2020 [Online] / Growth: NL Gov't, Flexera, Binx 2020. Gartner 2019. IA 2017 / Nicolae et al., Kavier (2025)

Datacenters are Crucial to Today's Society

Gizmodo Australia

People Can't Vacuum Or Use Their Doorbell Because Amazon's Cloud Servers Are Down

The company that produces Roombas, iRobot, confirmed they (along with their robot mops) were no longer working. An Amazon AWS outage is ...

26 nov. 2020

The Register®
Biting the hand that feeds IT

Google goes dark for 2 minutes, kills 40% of world's net traffic www.theregister.co.uk/2013/08/17/google_outage/

Systemwide outage knocks every service offline

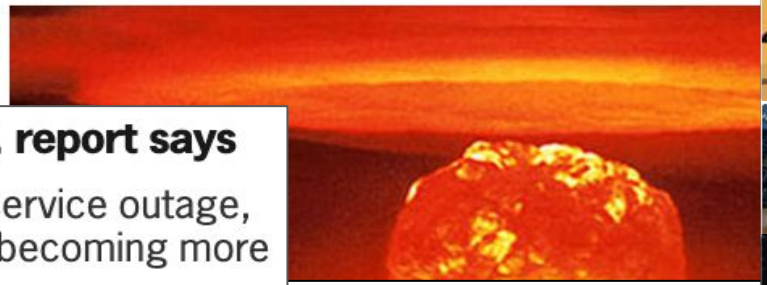

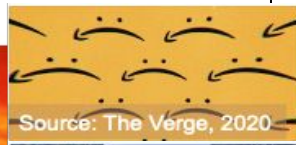

Data-center outages: Causes are changing, report says

Power issues are less likely to cause a major IT service outage, while IT configuration and network problems are becoming more common, according to the Uptime Institute.

Source: Network World, 2021 [Online]

Improving datacenters using simulation

Simulators enable timely and cost-efficient, large-scale experiments and analysis

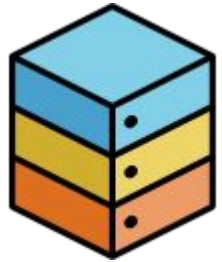

Challenge: simulation is manual, one-time

&

Challenge: need **continuous**, **real-time** operation, monitoring, and adjustment

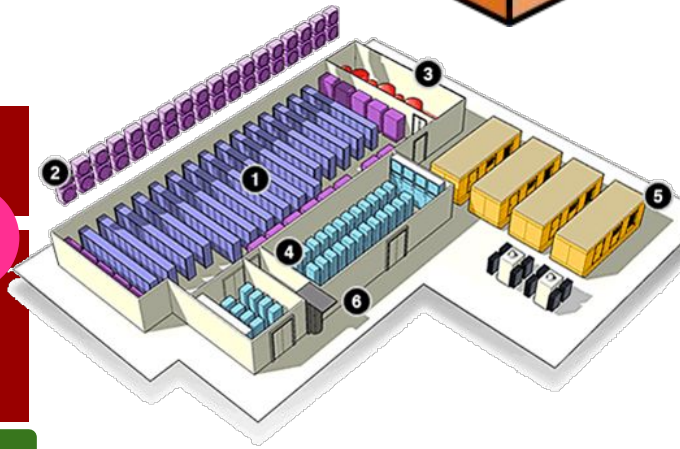

Source: Mastenbroek et al. CCGrid'21 [ONLINE]

Toward: closed-loop digital twinning!

TODAY'S AGENDA

1. Simulators for datacenters
2. **Digital twins** for datacenters
3. A self-calibrating digital twin: **OpenDT**
4. Exploring digital twin vs. simulation
5. Exploring calibration in the wild
6. Live demo

An analogy: Digital twins in aviation

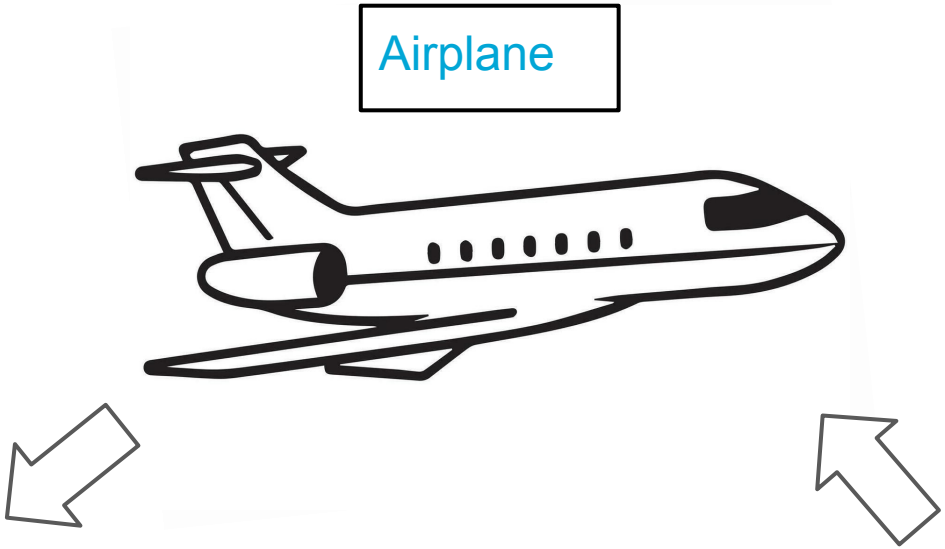

Digital twins in ICT

Datacenter

Measured data

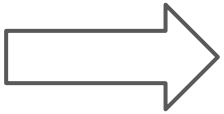

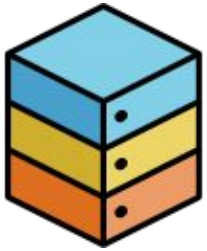

Simulator

Digital twins in ICT

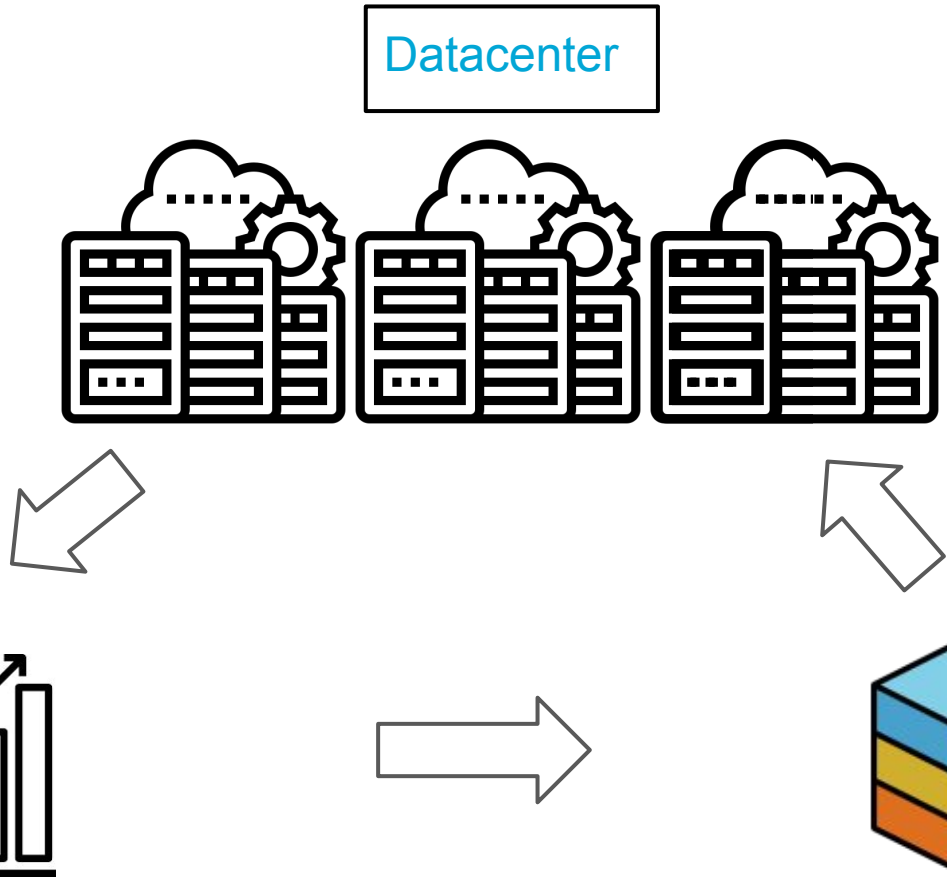

Improving datacenters using digital twinning

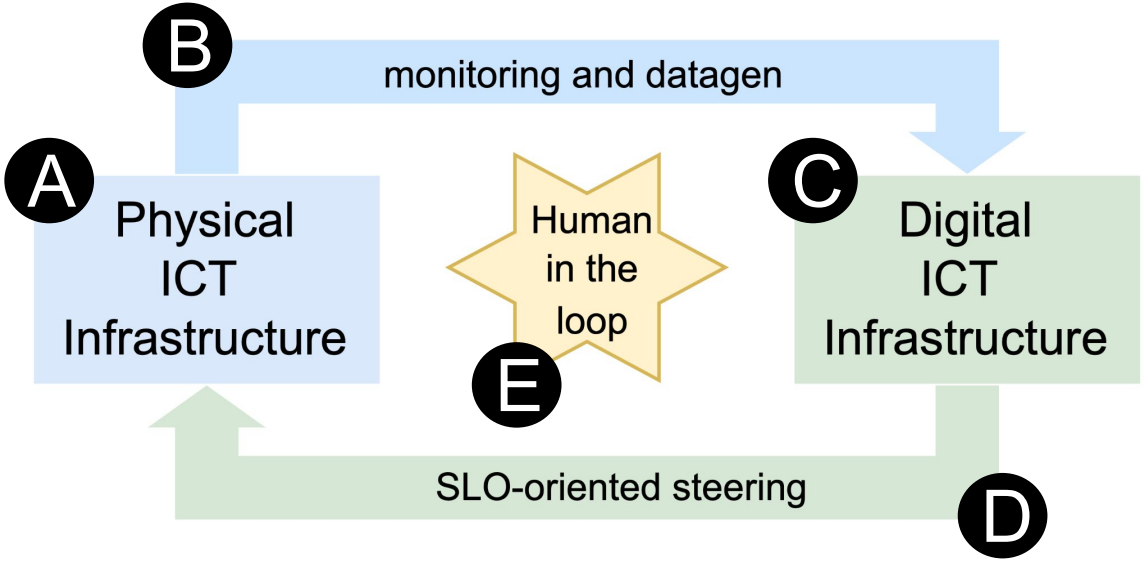

Use cases (UC)

UC1 Monitoring & operating datacenters

UC2 Research ICT

UC3 CompSys Education

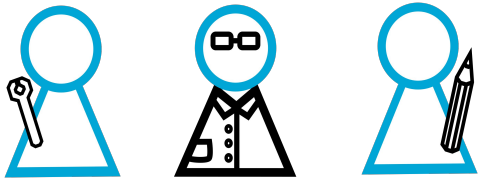

Main functional requirement (MFR)

MFR Active replication of real-world system state

OPENDT: THE FIRST MODEL+SIM BASED, OPEN-SOURCE DIGITAL TWIN FOR OPERATING AND MONITORING DATACENTERS

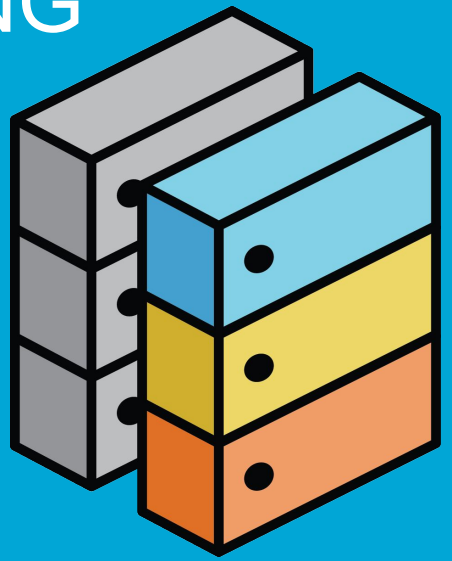

- **C1** DESIGN
- **C2** IMPLEMENTATION
- **C3** REAL-WORLD EXPERIMENTATION

High-level design of OpenDT

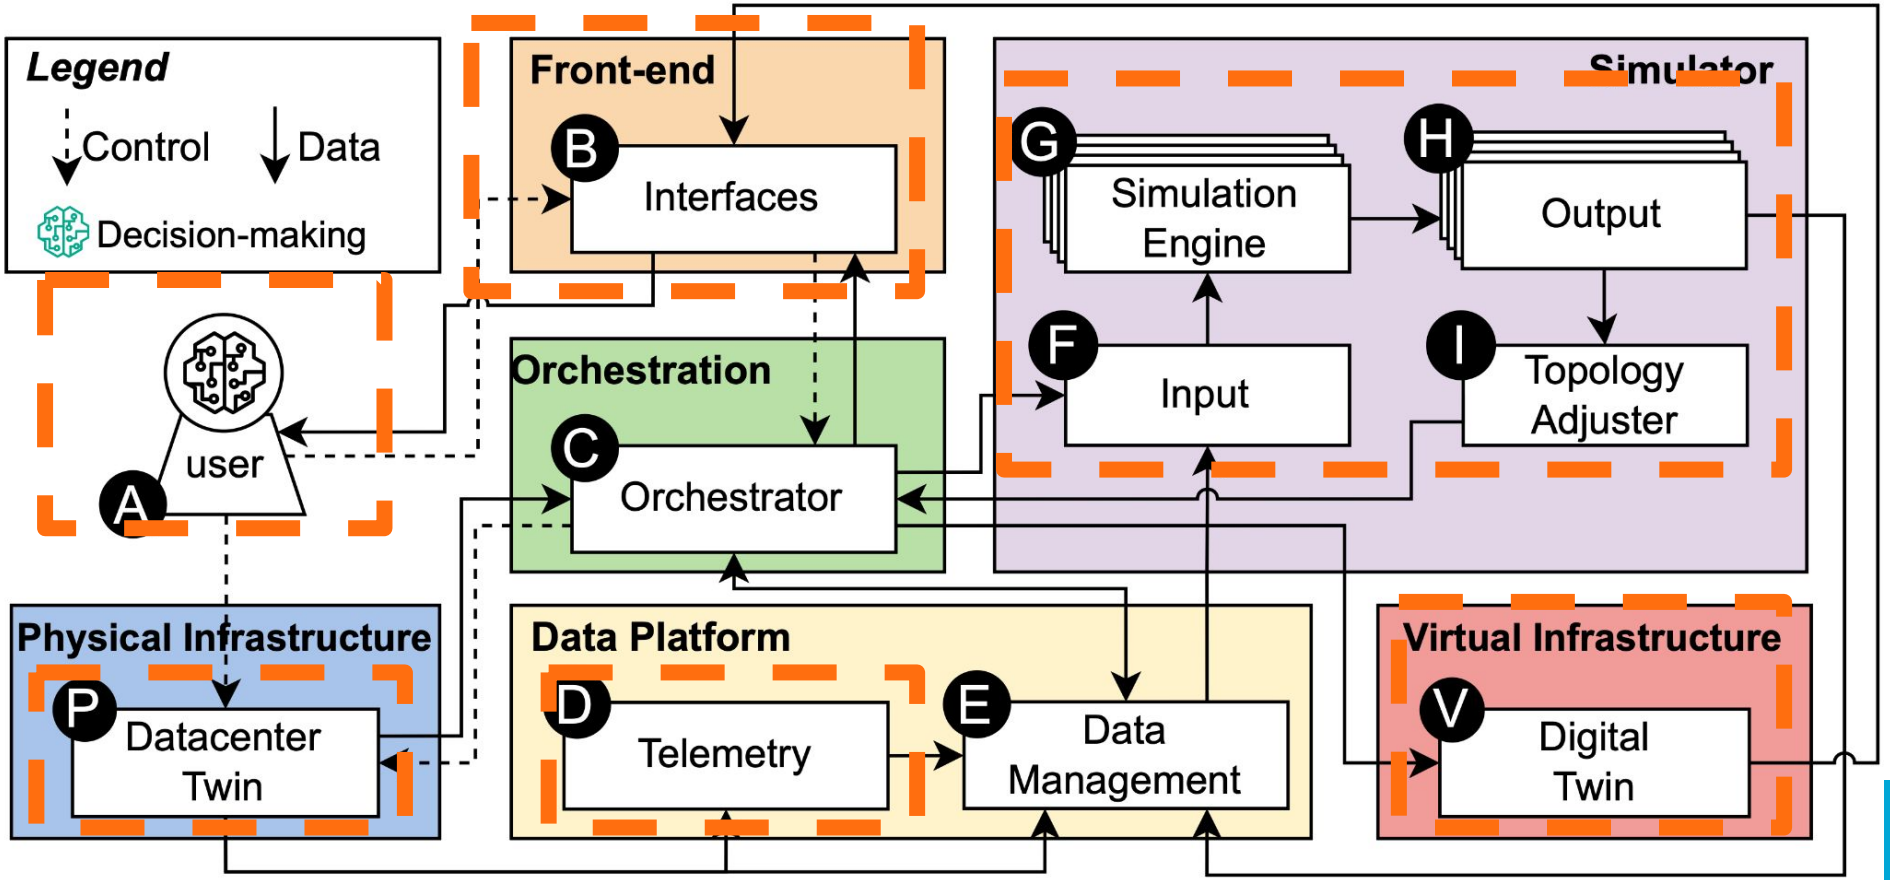

HL implementation of OpenDT

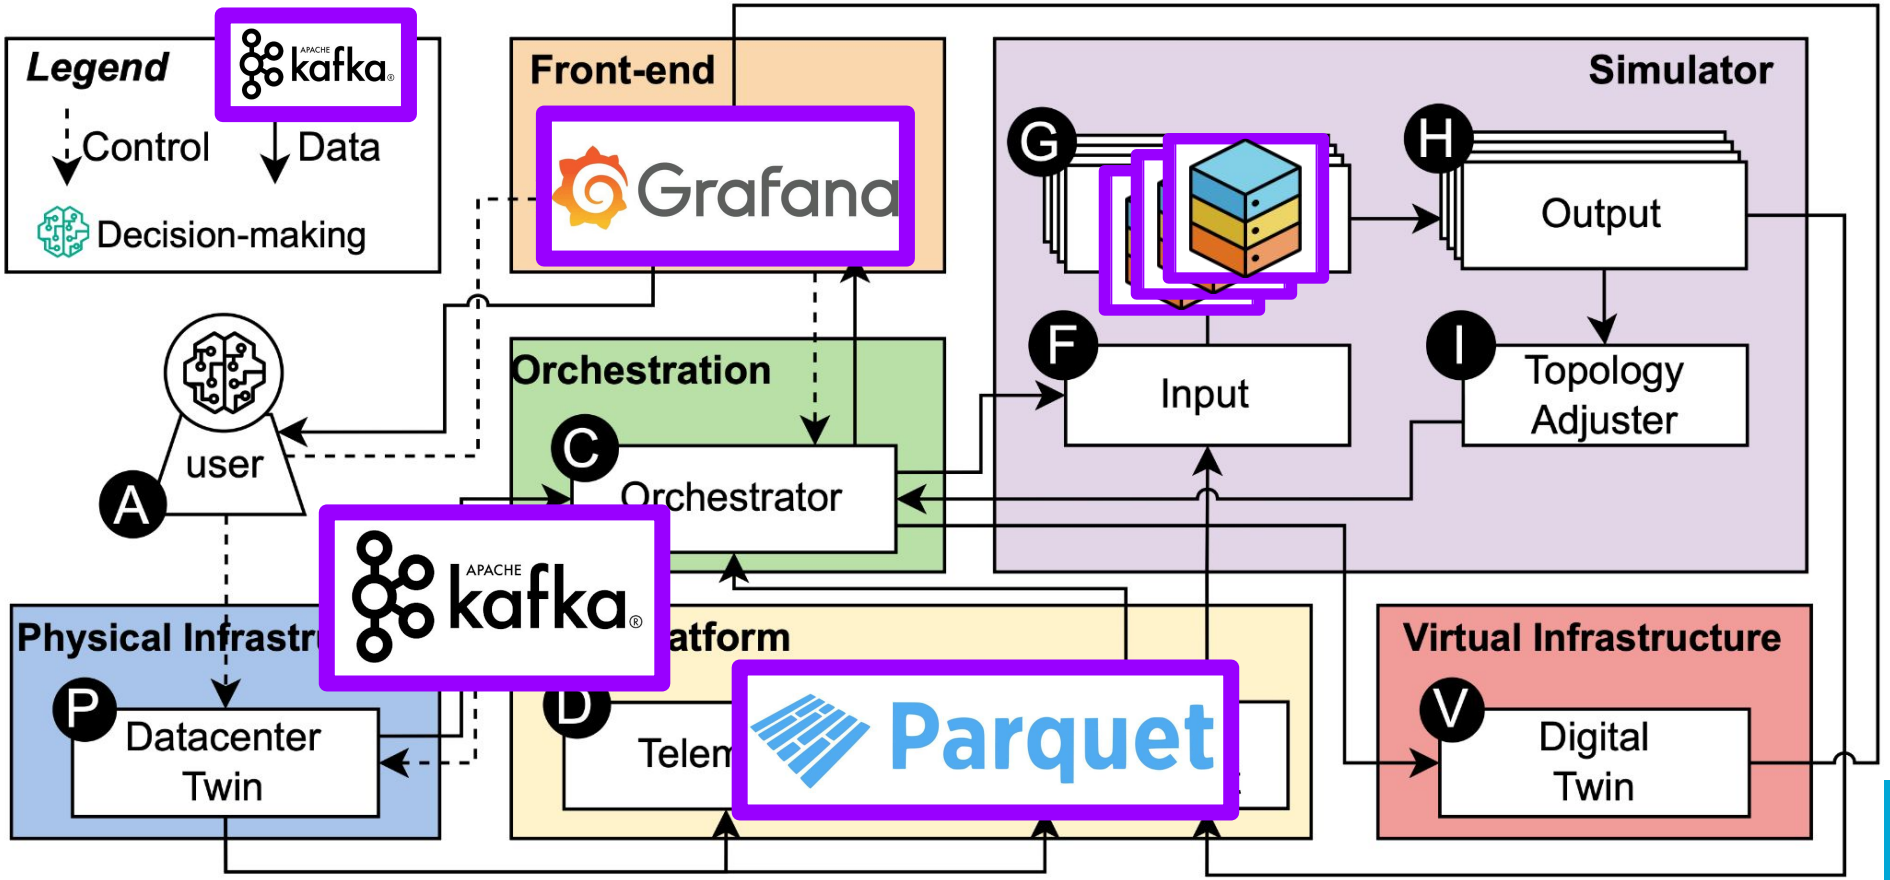

OpenDT's window-based self-calibration

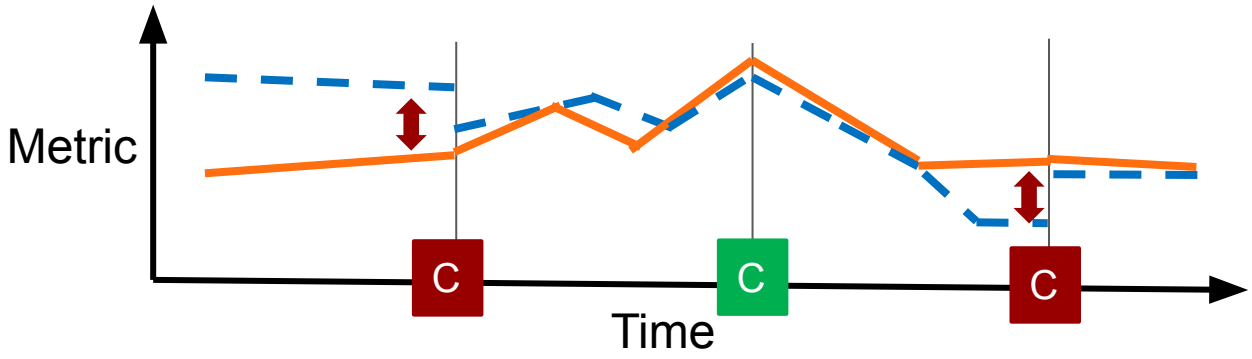

C1

C2

Simulation

Telemetry

Power draw, using a **calibration factor r**

$$P(u) = P_{\text{idle}} + (P_{\text{max}} - P_{\text{idle}})(2u - u^r)$$

P = power draw
u = utilization
r = calibration factor
Pidle = power draw while idle
Pmax = power draw at maximum utilization

N parallel runs

N=1; 3; **5**; 7; 9

N values for r

R=0.5; 0.8; 1.1; **1.4**, 1.7
Error%=8.5; 8.8; 5.1; **2.0**, 3.5

Sweep: best performing factor

R=1.4 with **Error=2.0%**

Sources: Mastenbroek et al., OpenDC 2.0, CCGrid '21 [Online] / Nicolae et al., M3SA, CF'26 [Online]

REAL-WORLD TRACE-BASED EXPERIMENTS WITH OPENDT

4

C3

Experiment 1

Reproducing peer-reviewed experiments with OpenDT capabilities

Setup: SURF-22, FootPrinter, & OpenDT

MAPE error ratio (lower=better)

FootPrinter	OpenDT
7.86%	5.13%

OpenDT Twinning Power-Draw

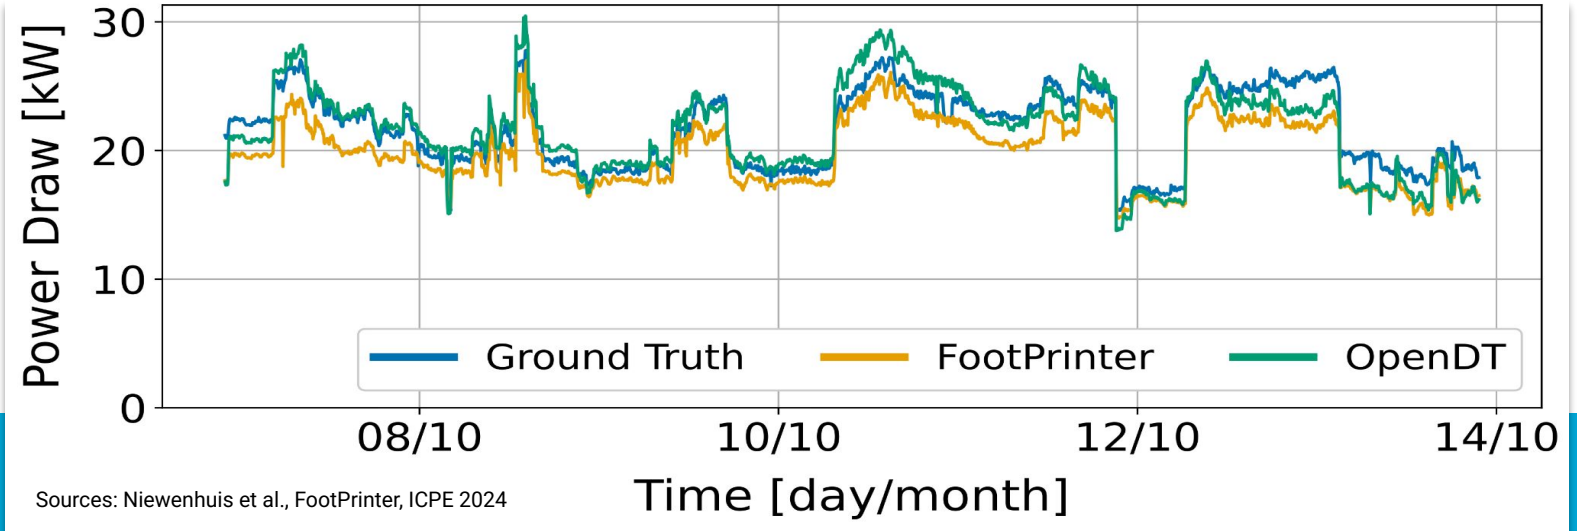

Sources: Niewenhuis et al., FootPrinter, ICPE 2024

Experiment 1

Setup: SURF-22, OpenDT

OpenDT twinning **performance (CPU + Latency)**

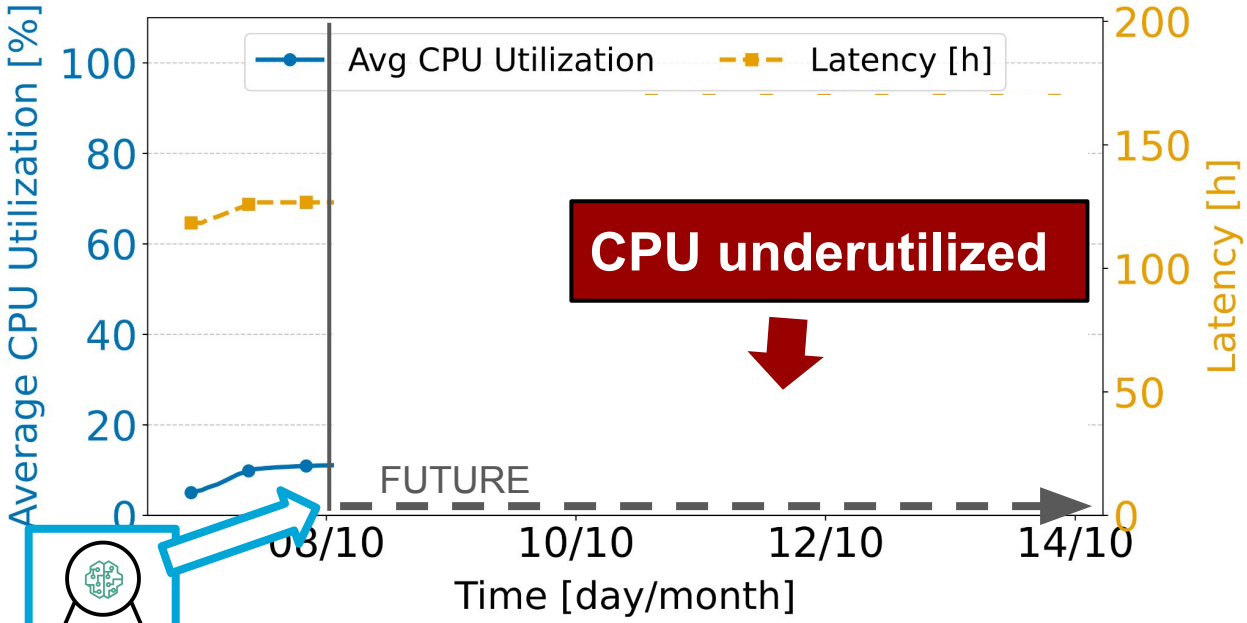

CPU underutilized

Consuming power, while idle

Pidle = 30-50% x Pmax

Radu Nicolae, Jules van der Toorn, Stavriana Kraniti, Houcen Liu, Alexandru Iosup (2026) OpenDT: Exploring Datacenter Performance and Sustainability with a Self-Calibrating Digital Twin. ICPE Companion 2026: 142-147

C3

Experiment 2

Ask about these![2]

Real-time simulation re-calibration

Setup: SURF-22, OpenDT

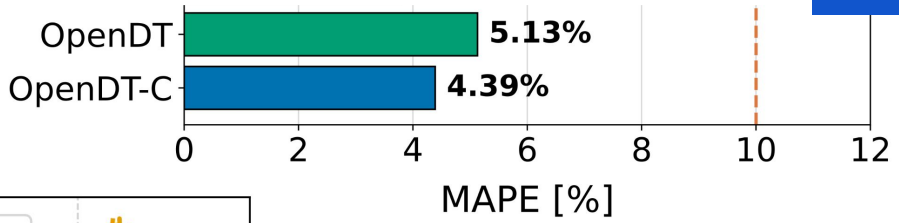

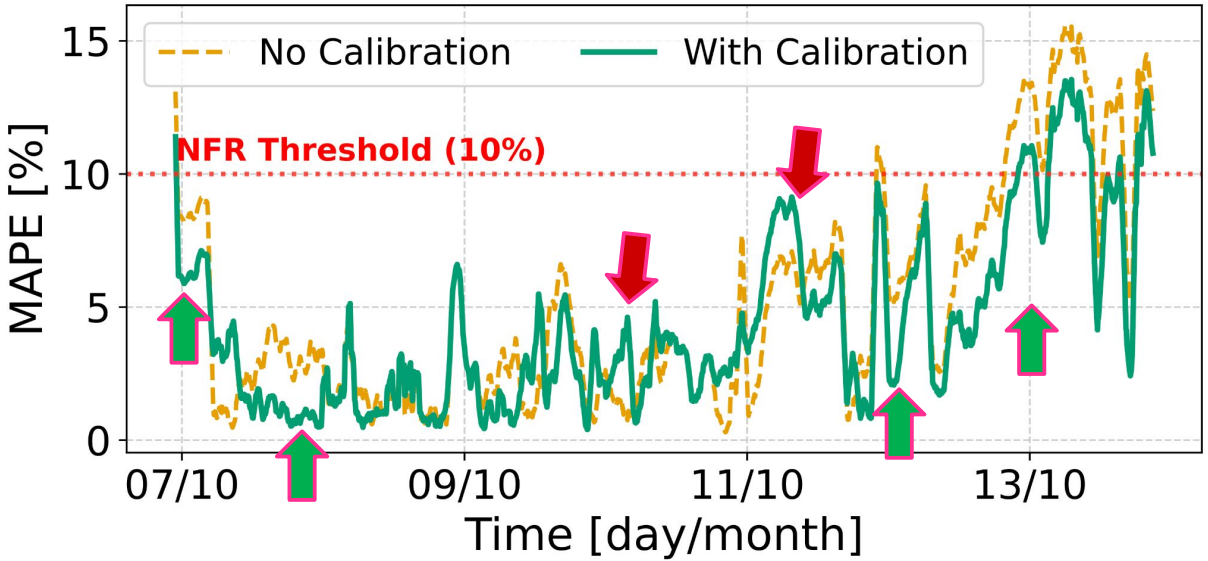

Calibration more accurate

Not always more accurate

[2] Nicolae et al., M3SA, CF'26 [Online]
[3] Sources: Niewenhuis et al., ICPE '24 [Online];

[1] Radu Nicolae, Jules van der Toorn, Stavriana Kraniti, Houcen Liu, Alexandru Iosup (2026) OpenDT: Exploring Datacenter Performance and Sustainability with a Self-Calibrating Digital Twin. ICPE Companion 2026: 142-147

TODAY'S AGENDA

1. Simulators for datacenters
2. **Digital twins** for datacenters
3. A self-calibrating digital twin: **OpenDT**
4. Exploring digital twin vs. simulation
5. Exploring calibration in the wild
6. Up next: (Live) Demo

OPENDT: DIGITAL TWINNING for DATACENTERS

Talk: bit.ly/opendt-hcp

VU VRIJE UNIVERSITEIT AMSTERDAM

@Large Research
Massivizing Computer Systems

FUTURE NETWORK SERVICES 6G

SURF

GRAPH MASSIVIZER

C1 Design of a self-calibrating DT

C2 Prototyping OpenDT

C3 Real-world experimentation

PAPER

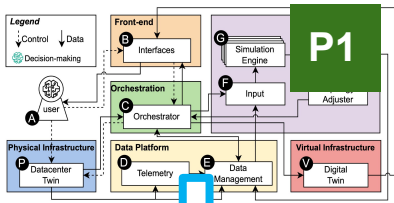

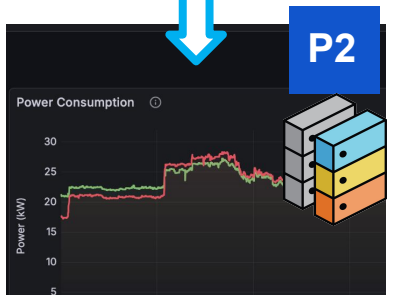

Future work

1. Closed-loop digital twin, coupled with real-world DC
2. Expanded simulation of the compute continuum
3. Digital twin for DCs under large-scale AI workloads
4. & more
5. Starting: SPEC RG activity on digital twin