ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

CINECA

EE HPC SOP Kobe meeting at RIKEN R-CCS Activity report from Cineca/UNIBO

***Andrea Bartolini* -** University of Bologna – DEI, Italy
Carlo Cavazzoni - CINECA, Italy

Outline

- Cineca Roadmap & Bologna Science Park
- Datacentre automation
- OoB Fine-Grain Power Sampling (DiG) & Anomaly detection
- Fine-Grain Energy-Management (Countdown)

The University of Bologna

Alma Mater Studiorum
Università di Bologna
is a multi-campus
university based in Bologna,
Cesena, Forlì, Ravenna, and
Rimini.

University Statute, Constituent
Principles, Art. 1 para. 2

Bologna the city of …

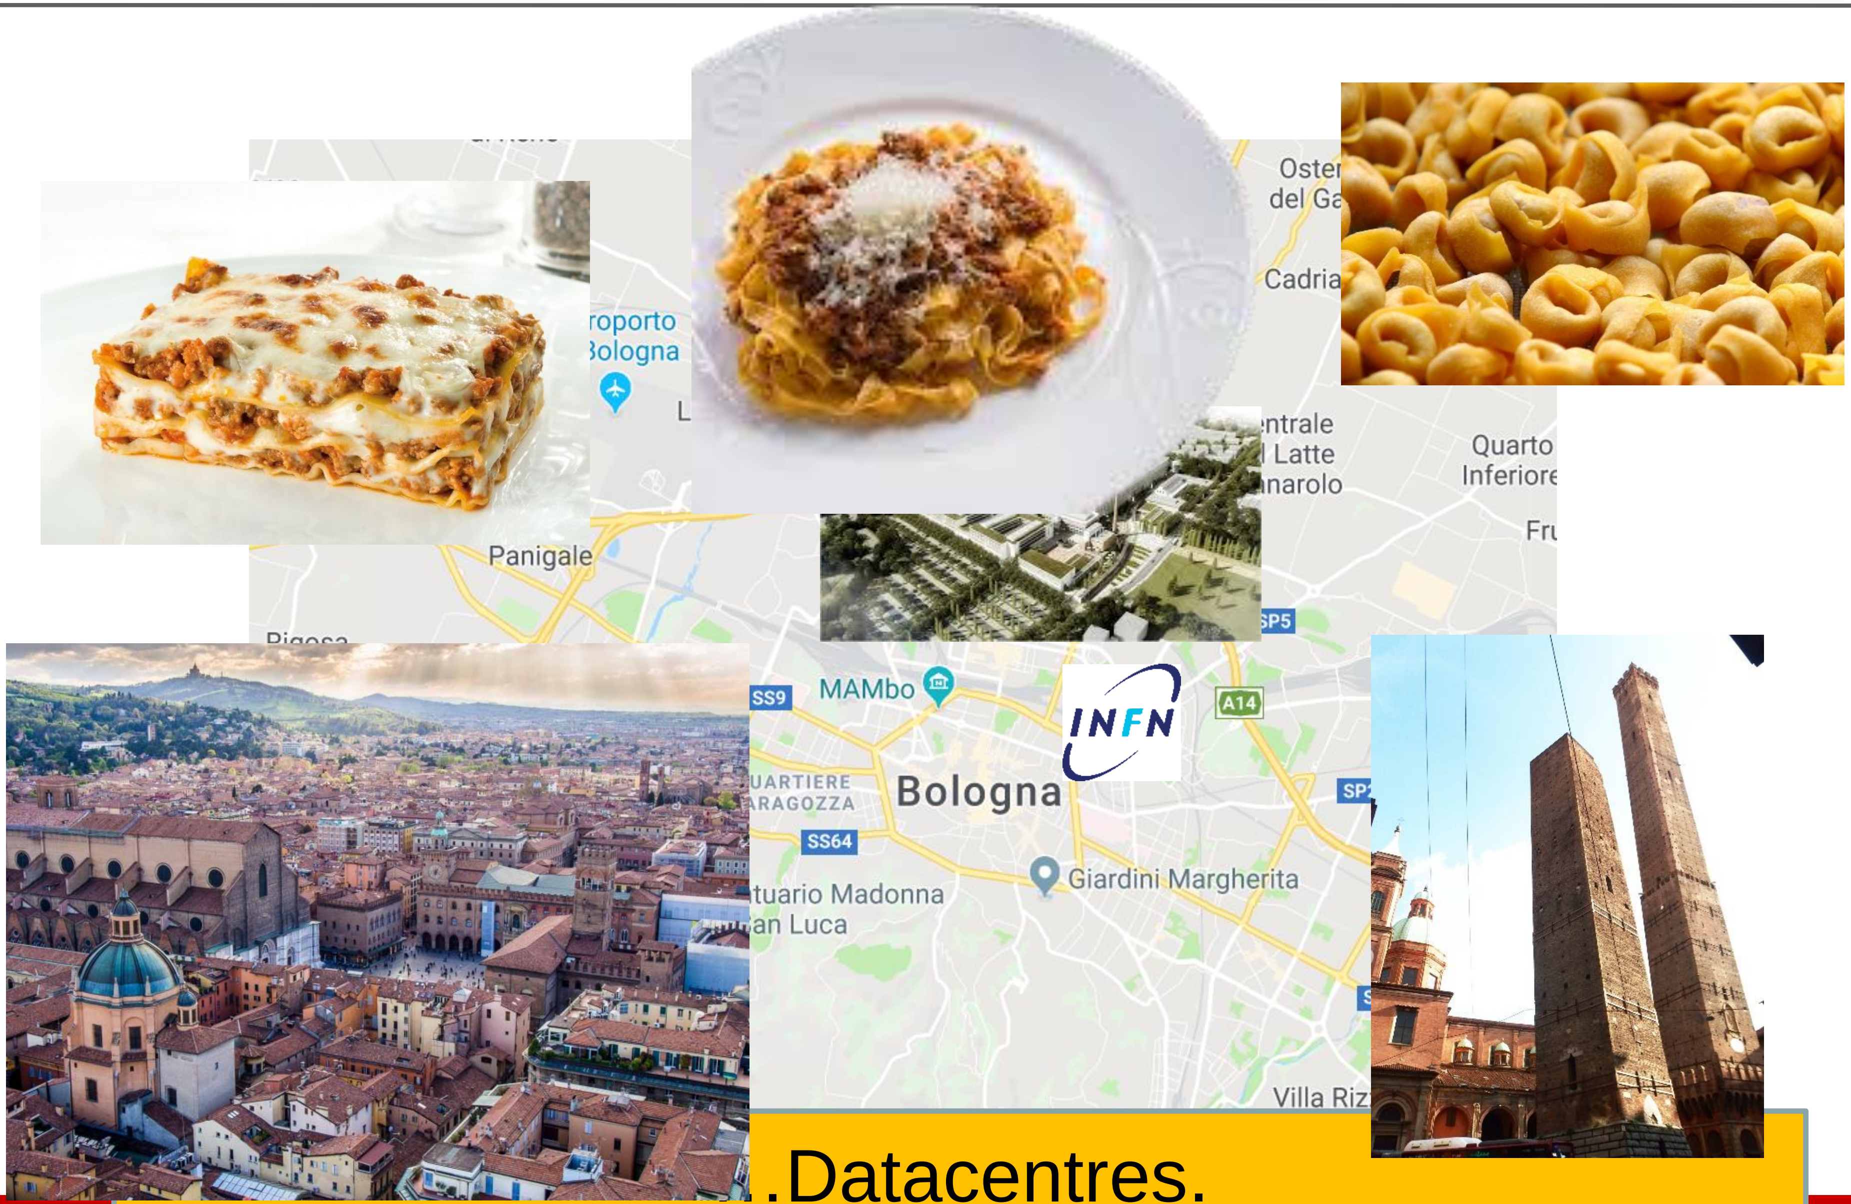

.Datacentres.

Bologna Science Park

EuroHPC Pre-Exascale @2021

Italian Exascale @2025

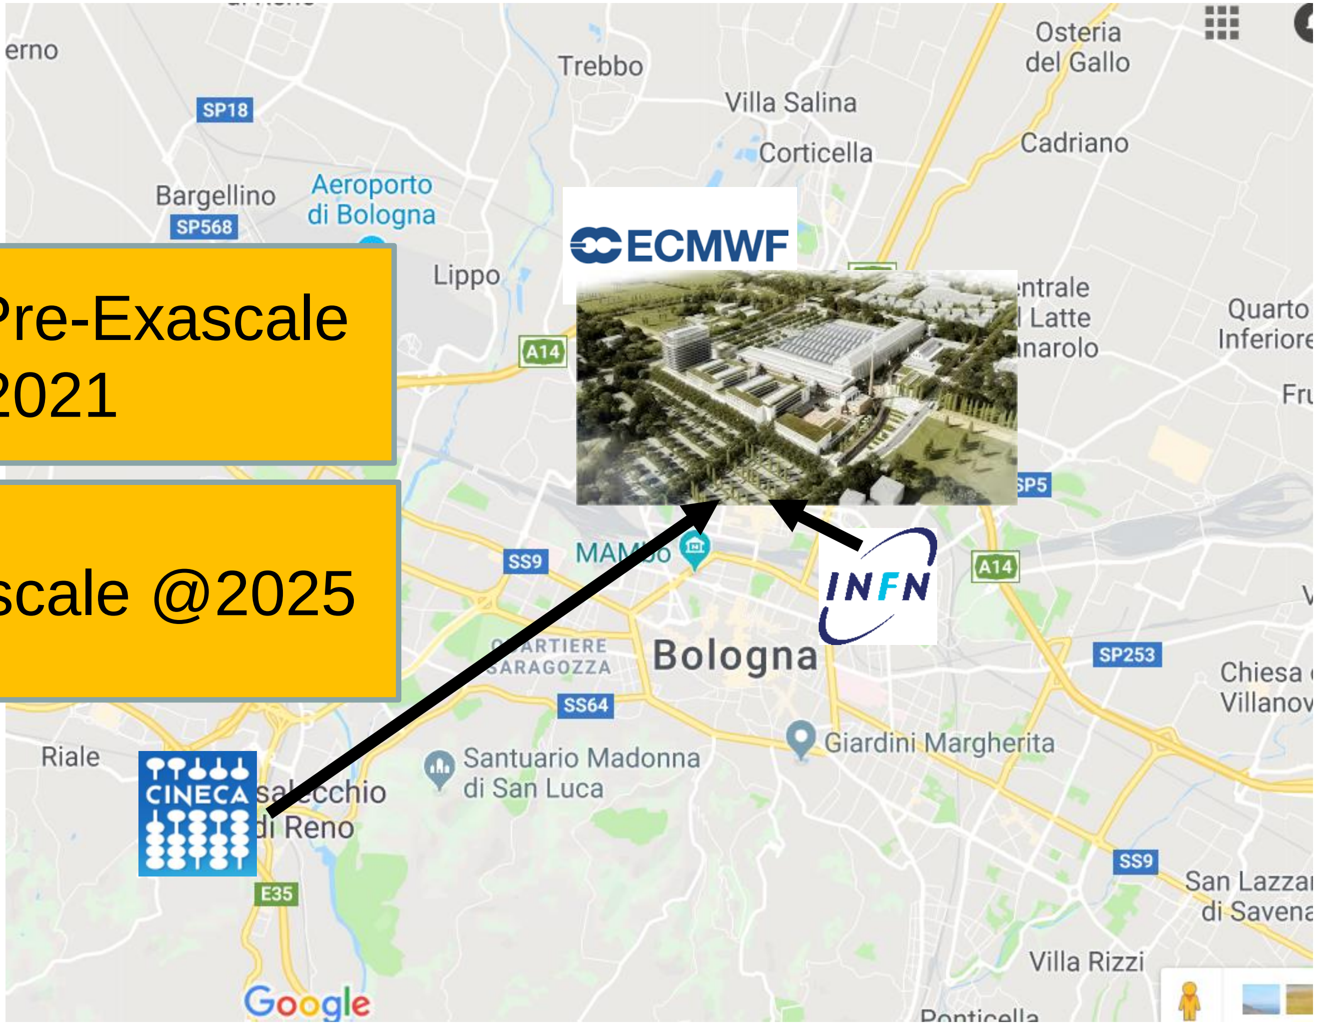

Outline

- **Cineca Roadmap & Bologna Science Park**
- Datacentre automation
- OoB Fine-Grain Power Sampling (DiG) & Anomaly detection
- Fine-Grain Energy-Management (Countdown)

RoadMap

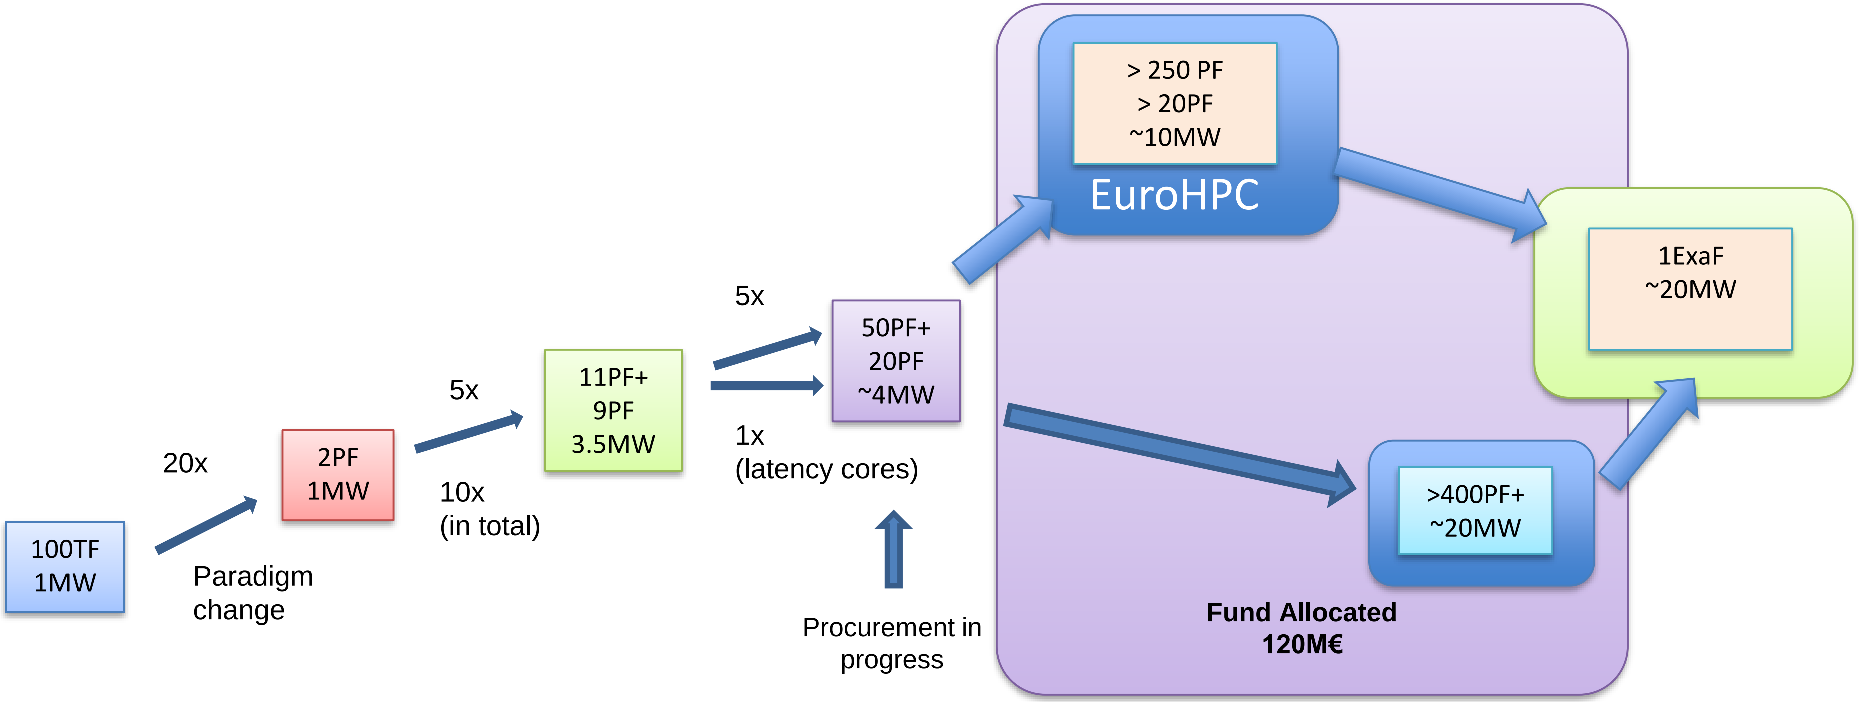

2009	2012	2016	2019/2020	2021	2023-2025	2025-2026
IBM SP6 Power6	Fermi IBM BGQ PowerA2	Marconi Lenovo Xeon+KNL	Marconi PPI4HPC ICEI - PPIHBP	With EuroHPC contribution	Post-Marconi exascale technology	Post-Exa with EuroHPC

Partnership

EuroHPC hosting @ Bologna Science Park

Cooling equipment
3MW (2020) -> 5MW(2023)

Computer Rooms
10MW (2020) -> 20MW (2023)

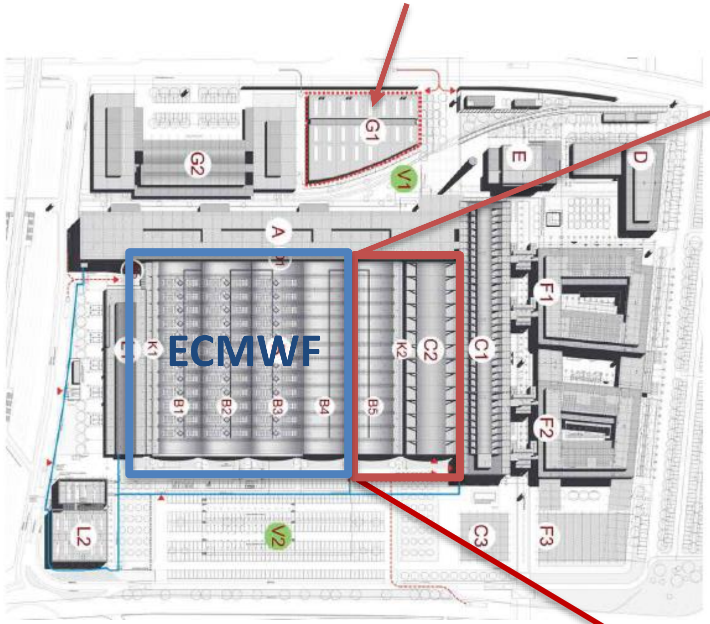

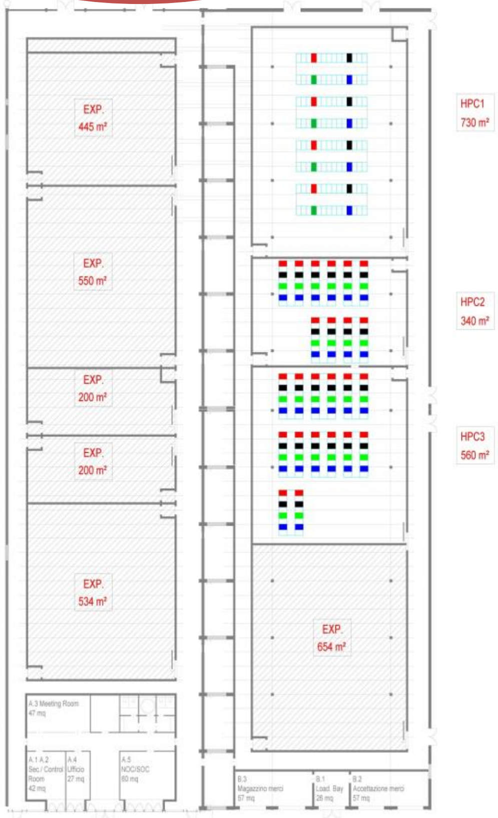

8MW hot water DLC
Compute nodes

2MW AIR Cooled
Storage + Ancillary

PUE < 1.1

DATA ROOM STAGE 1: 1600 sqm
DATA ROOM STAGE 2: 2600 sqm
ANCILLARY SPACES: 900 sqm

Partnership

Energy Efficiency and Management Objectives:

Objectives of interest:

(i) Enable correlation between power consumption and system workload;
(ii) Enable dynamic power capping with graceful performance degradation of the system;
(iii) Provide capability to optimize the job execution environment for better energy efficiency;
(iv) Provide energy accounting mechanism;
(v) Allow energy profiling of applications to enable EtS optimization without TtS degradation

The HPC solution:

- *Reliable power and energy measurement at different level (CPU, node, rack) & at high frequency*
- *Interfaces for integration with:*
    - *resource scheduler for energy accounting mechanism & power capping*
    - *holistic monitoring frameworks for datacentre automation*

The Datacentre solution:

- *Features an energy management system (EMS) to monitor, measure and control the loads.*
- *Centrally control cooling devices (HVAC type, etc.) and lighting systems*
- *Enables site's energy reporting and optimization w. measurement, submetering and monitoring functions*

Outline

- Motivation & Datacentre Automation
- Cineca Roadmap & Bologna Science Park
- **Datacentre automation**
- OoB Fine-Grain Power Sampling (DiG) & Anomaly detection
- Fine-Grain Energy-Management (Countdown)

Datacentre Automation

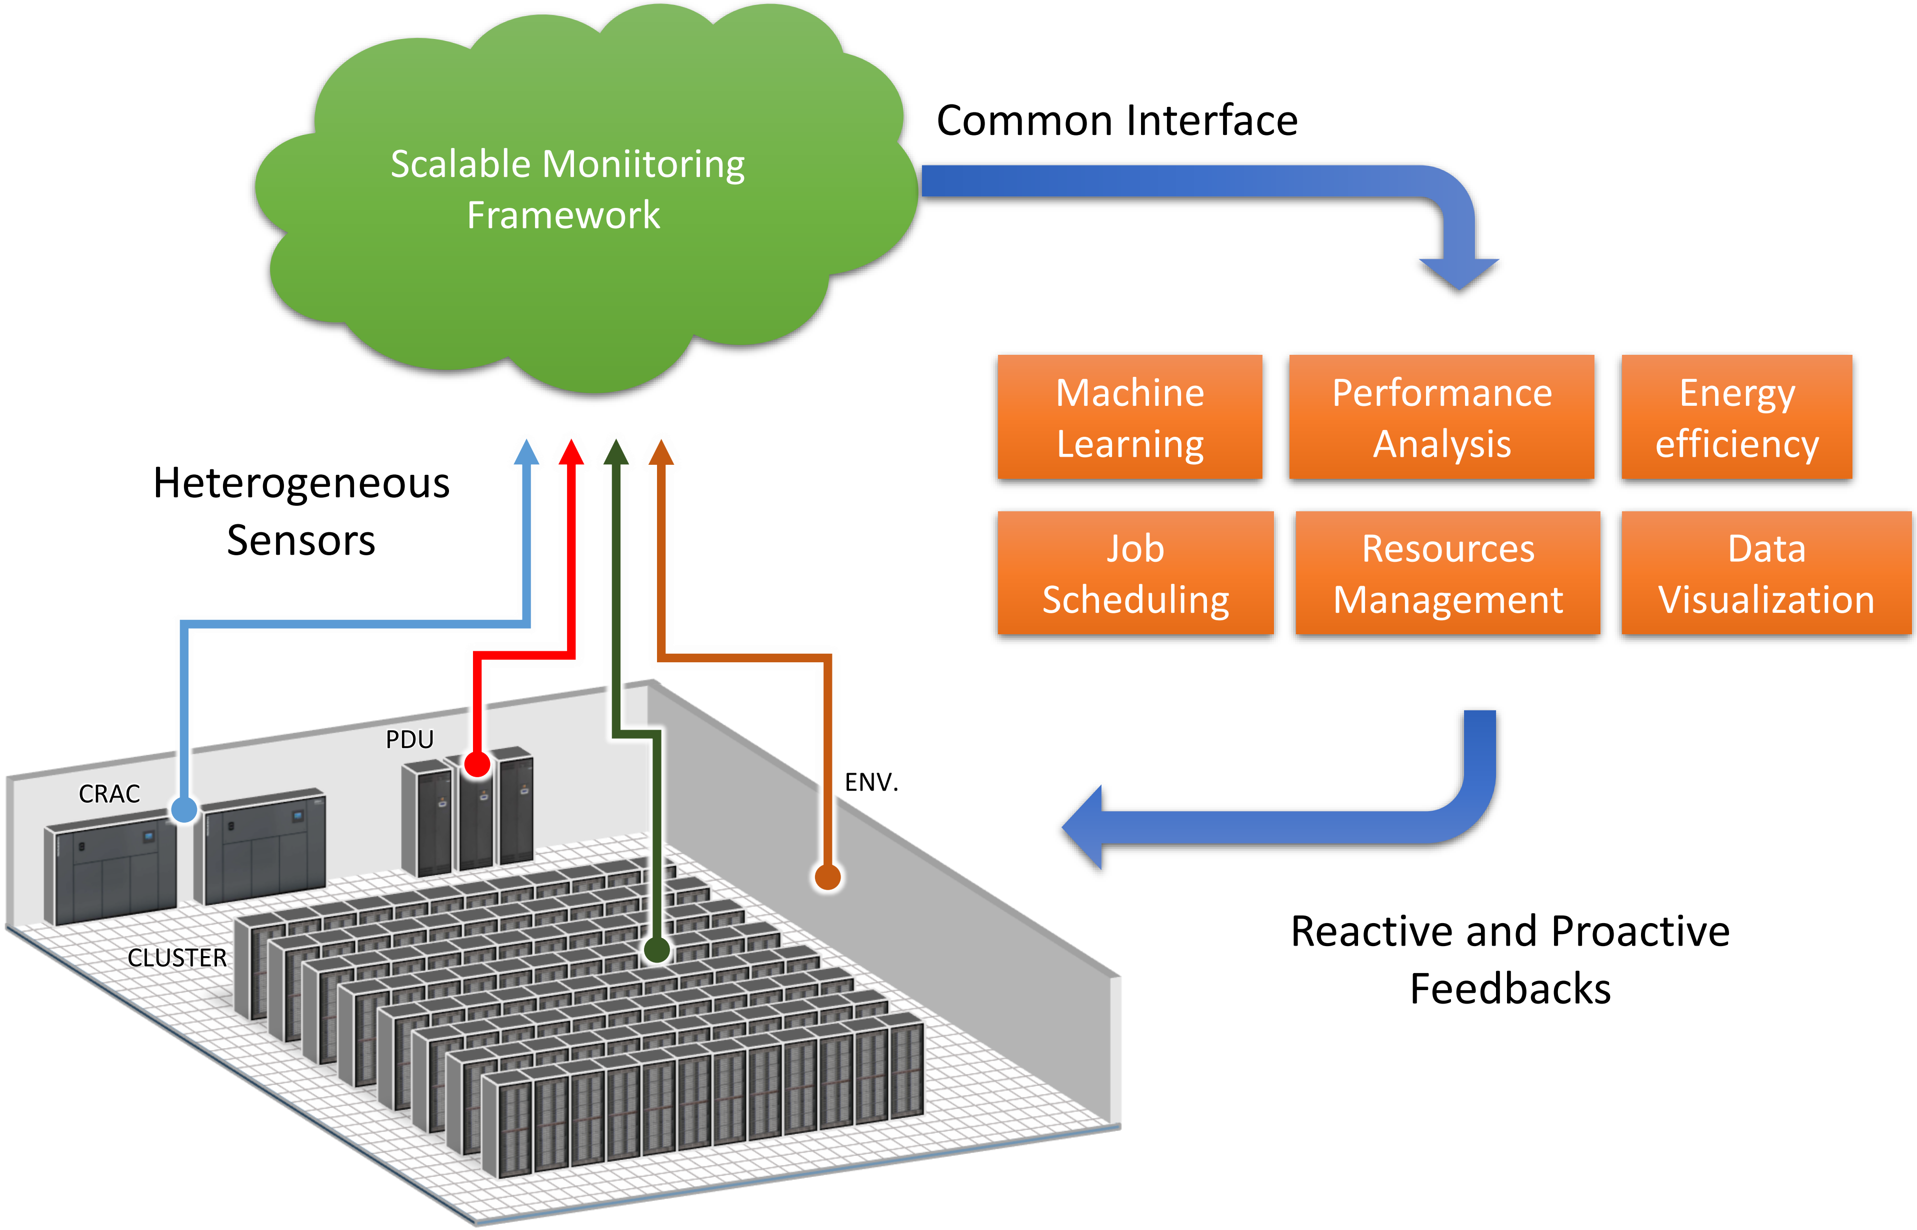

Usage Scenario #1 – Anomaly detection

Fine Grain Power and Performance Measurements:

- Verify and classify node performance
  - In spec / out of spec behaviour
  - Miss configuration
  - Aging and wear out
- Detect security hazards
- Predictive maintenance

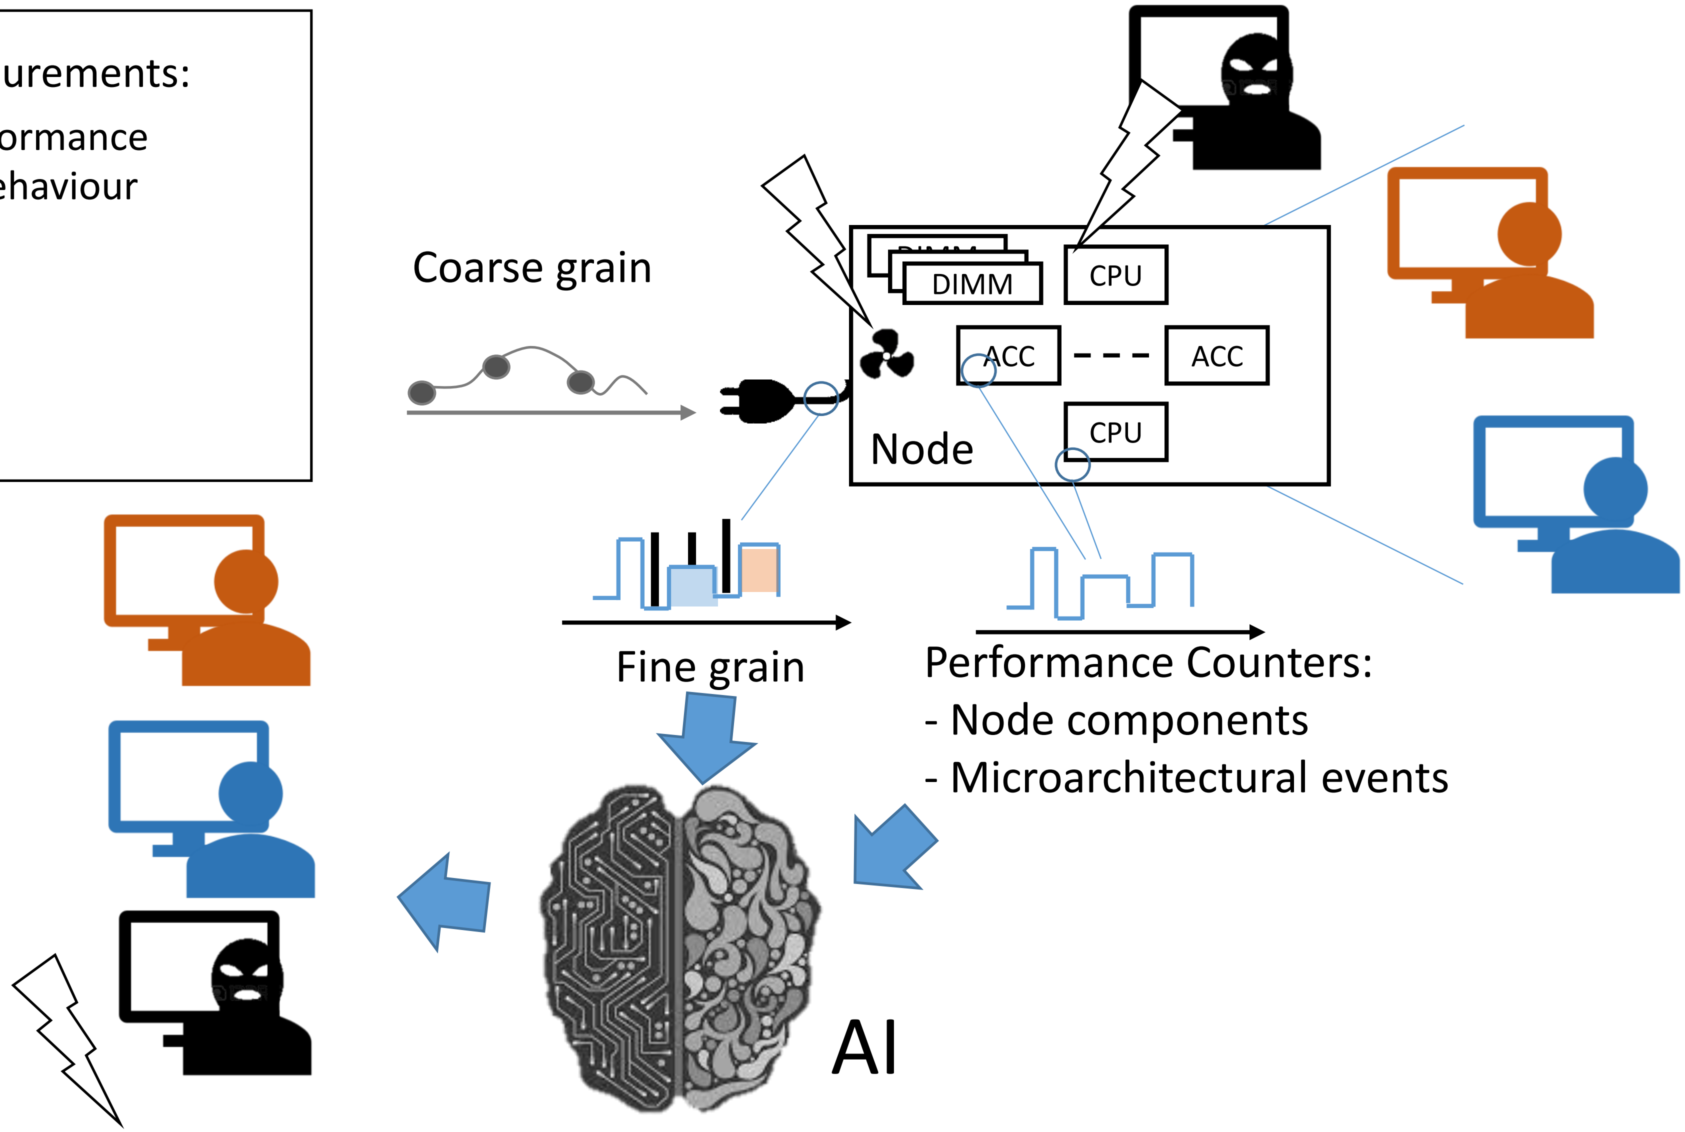

Usage Scenario #2 – Energy Efficiency

Energy
Saving !!!

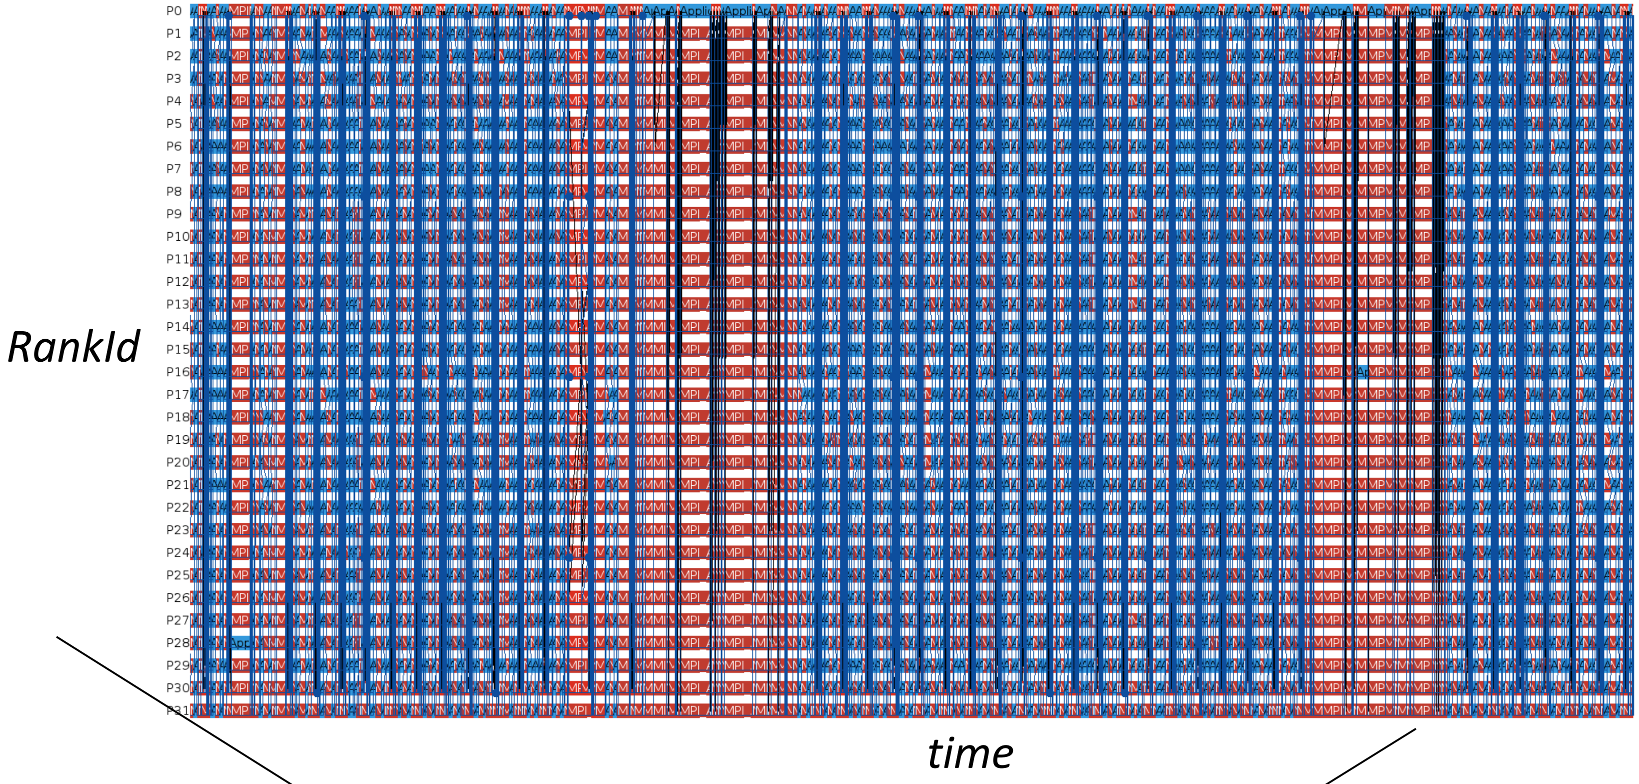

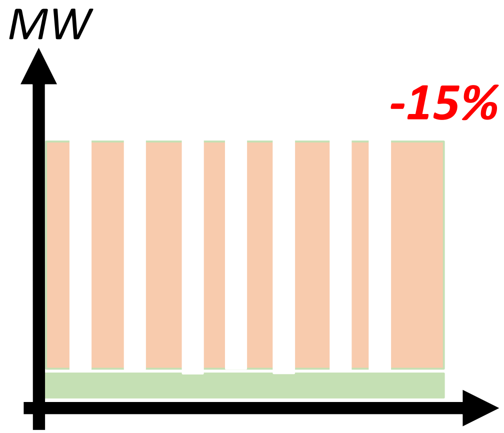

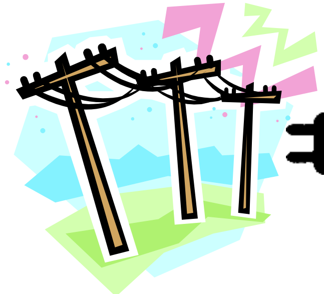

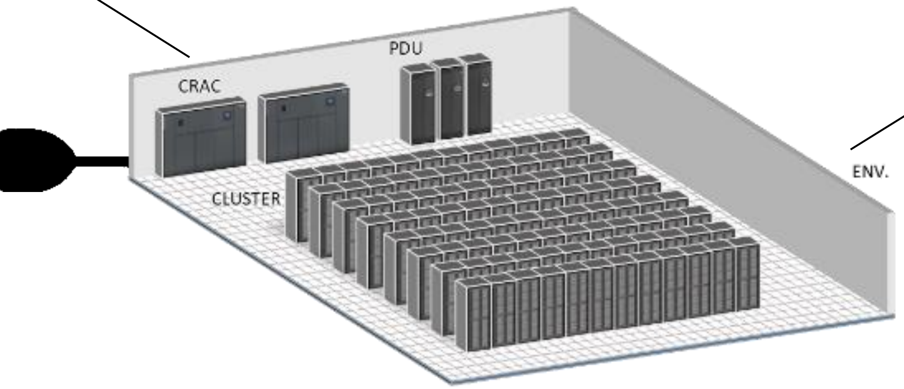

Scalable Data Collection, Analytics

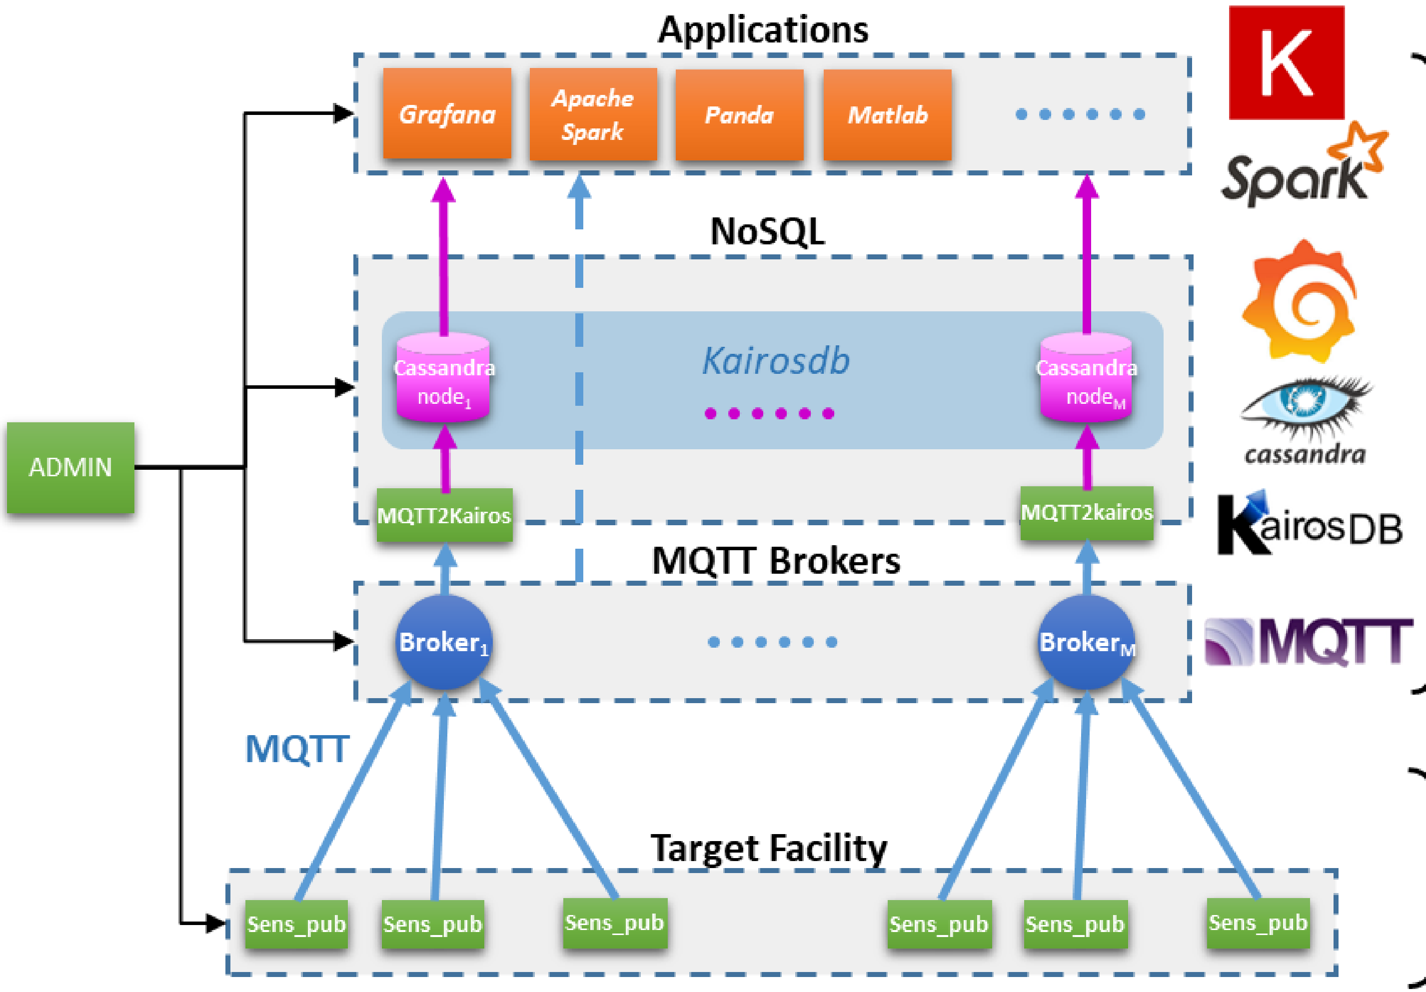

Front-end

- **MQTT** Brokers
- Data **Visualization**
- **NoSQL** Storage
- **Big Data** Analytics

Back-end

- MQTT–enabled sensor collectors

https://github.com/EEESlab/examon

F. Beneventi et al., "Continuous learning of HPC infrastructure models using big data analytics and in-memory processing tools"
A. Bartolini et al, "The DAVIDE Big-Data-Powered Fine-Grain Power and Performance Monitoring Support"

Scalable Data Collection and Analytics

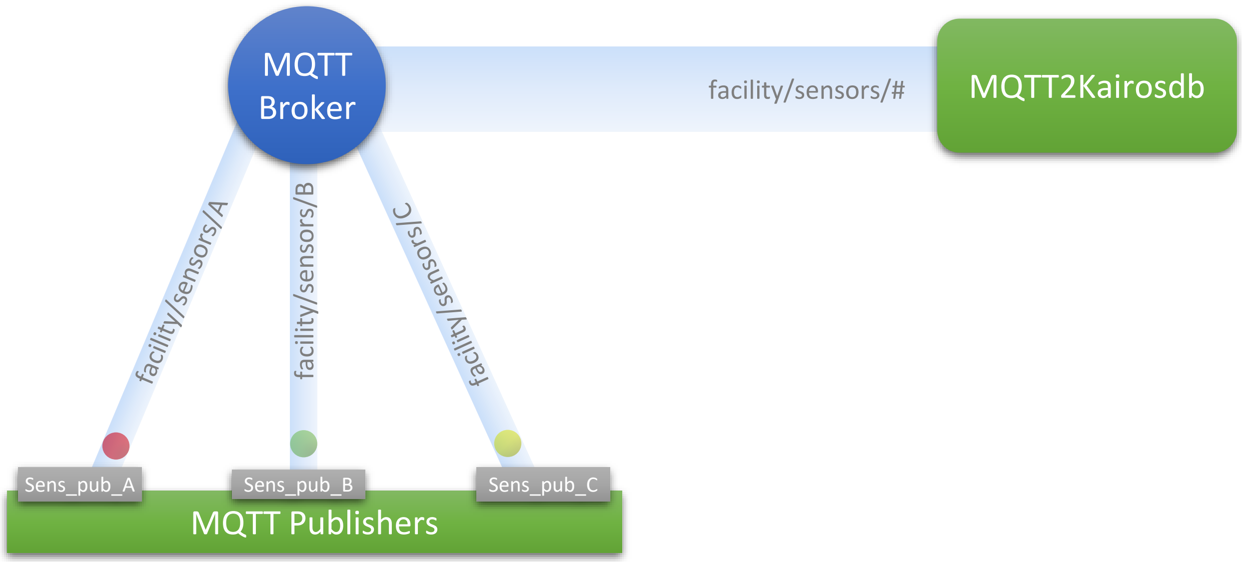

{Key,Value} = TS, Measurement
Topic = /davide/node1/Metric

EXAMON

ExaMon: Batch & Streaming & Edge

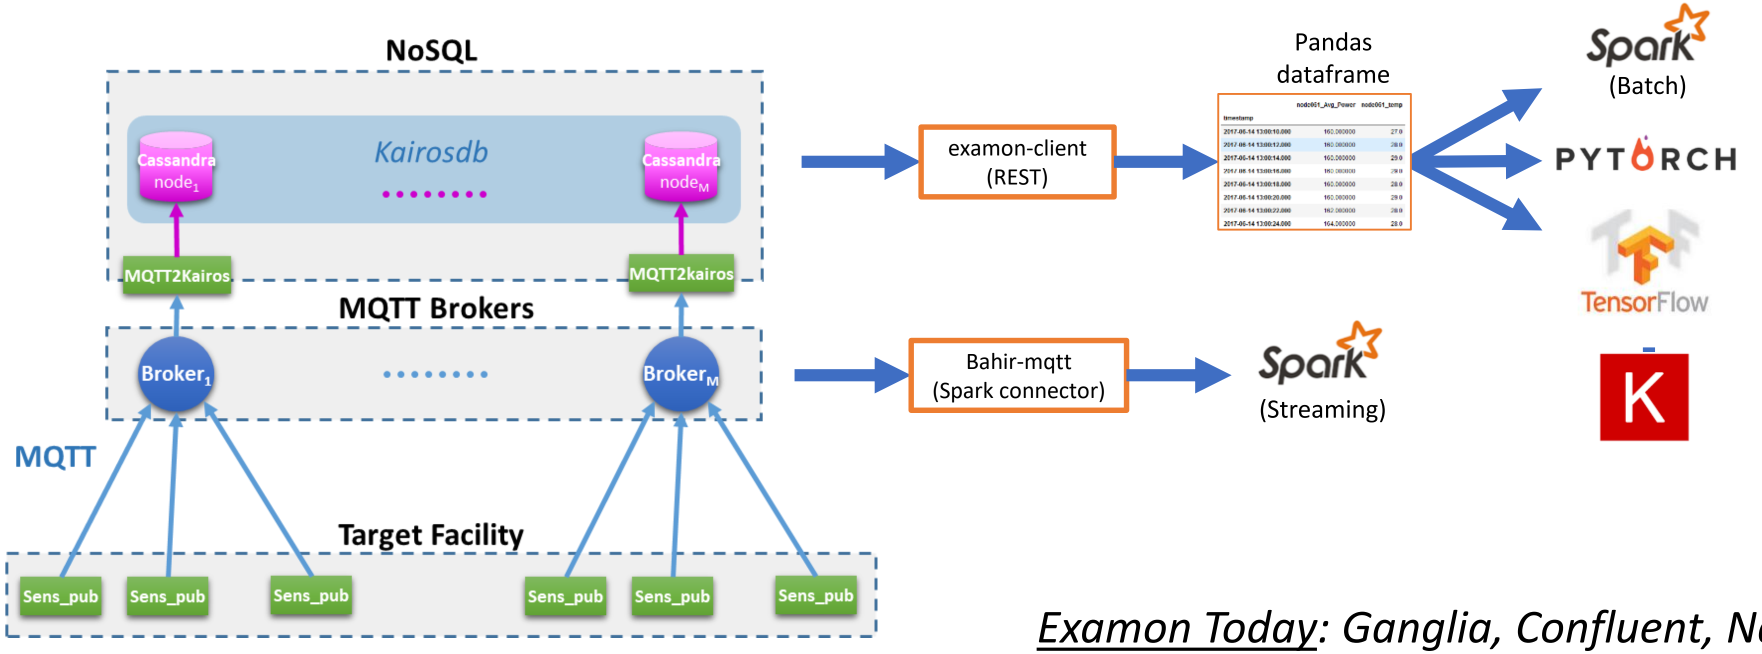

<u>Examon Today</u>: Ganglia, Confluent, Nagios, IPMI, DiG
1M metrics monitored
~8K computing nodes
70GB/day of Data

Datacenter Automation Design and Bottlenecks

[AICAS18] Borghesi et al. **Online Anomaly Detection in HPC Systems**

[DAAC18] Libri et al. **DiG: Enabling Out-of-Band Scalable High-Resolution Monitoring nalytics, Automation, and Control**

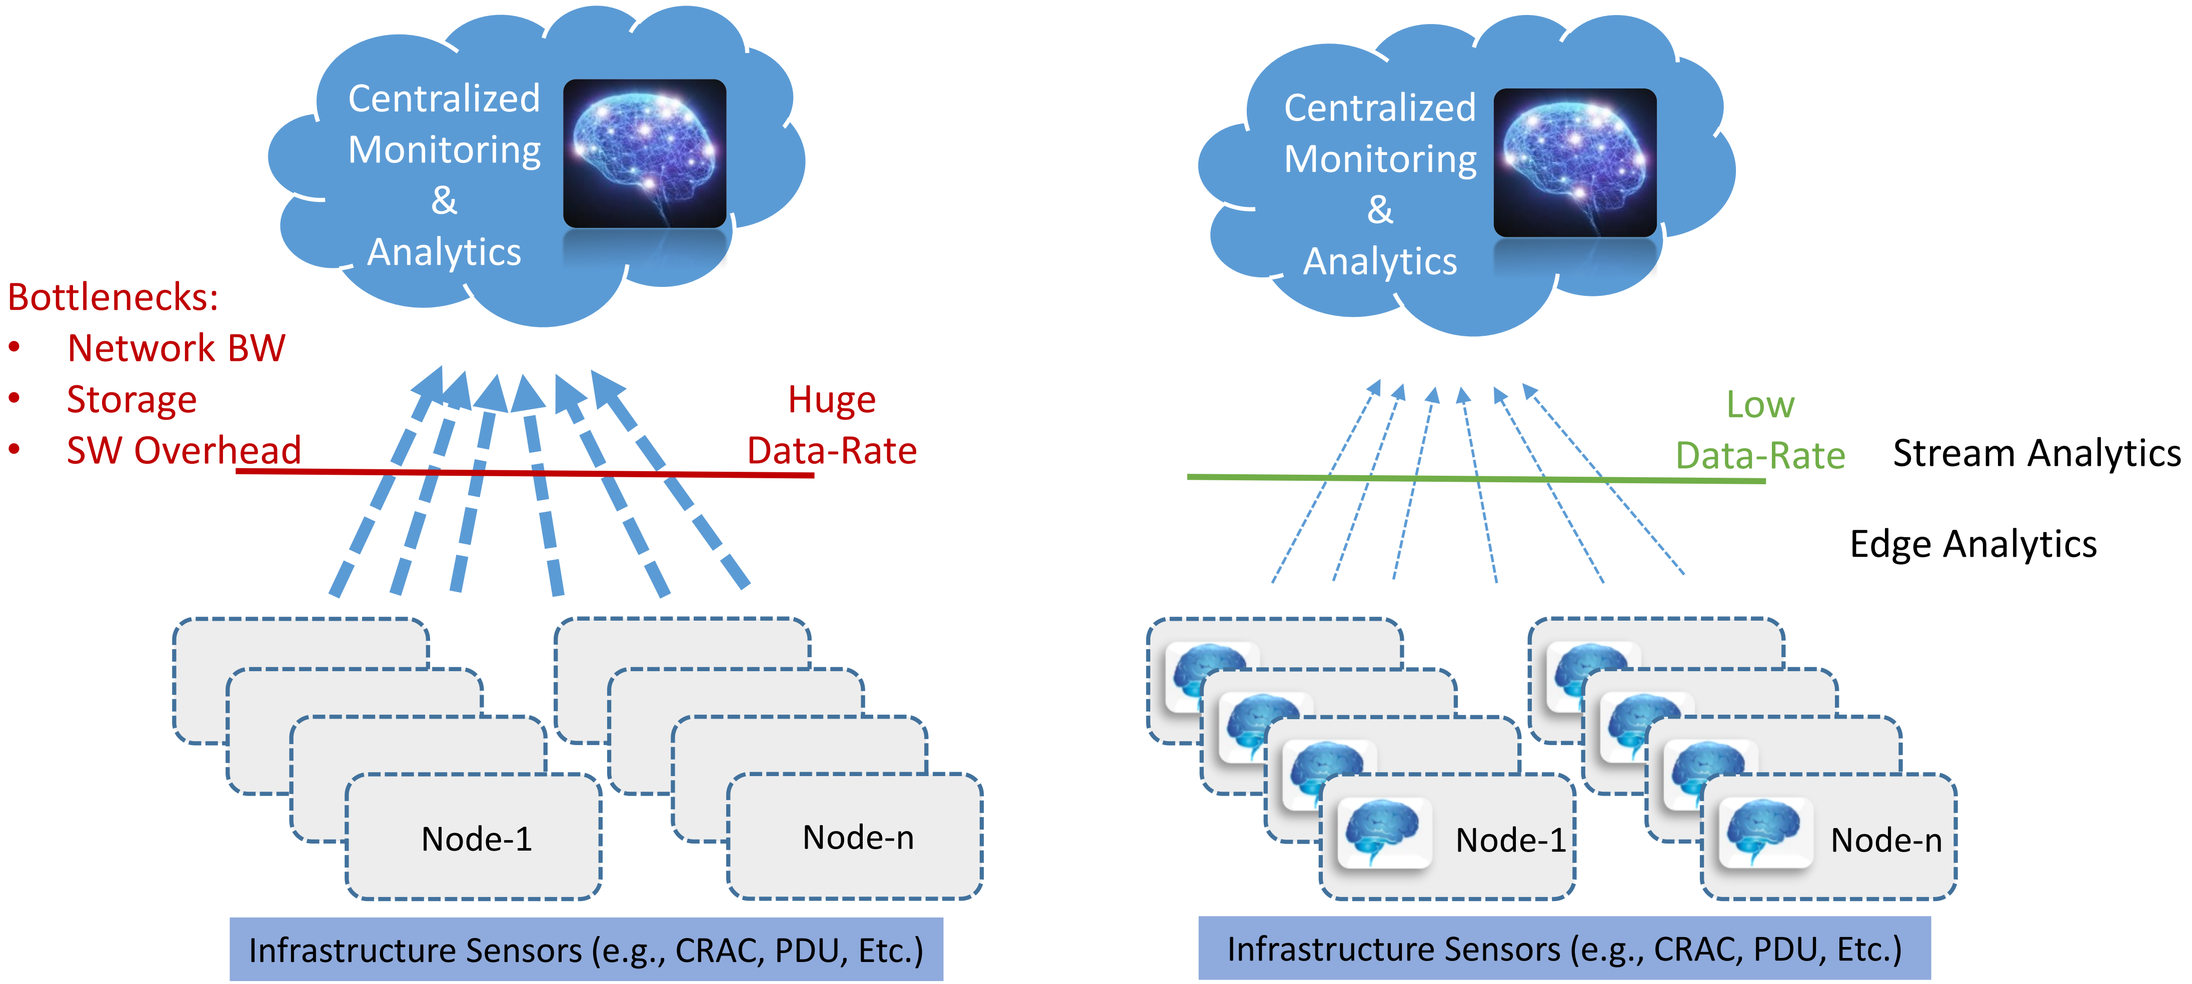

Outline

- Motivation & Datacentre Automation
- Cineca Roadmap & Bologna Science Park
- Datacentre automation
- **OoB Fine-Grain Power Sampling (DiG) & Anomaly detection**
- Fine-Grain Energy-Management (Countdown)

DiG = High Frequency Monitoring on D.A.V.I.D.E.

E4 COMPUTER ENGINEERING
PRACE
wistron
CINECA

OCP form factor compute node
based on IBM Minsky

4x NVIDIA Tesla
P100 HSMX2

2 x IBM POWER8
with NVLink

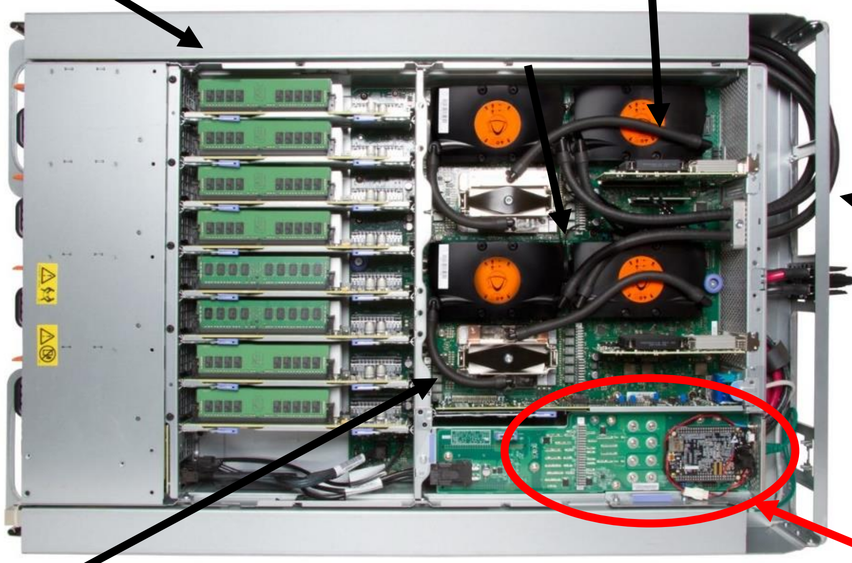

2xIB EDR

BusBar

D.A.V.I.D.E - PRACE PCP III
x45 ~1PFlop

LIQUID COOLING

DiG

ETH Zurich / Univ. of Bologna
SoA out-of-band
High Resolution Power
and Performance Monitoring

[DAAC18] https://arxiv.org/abs/1806.02698

DiG = High Frequency Monitoring on D.A.V.I.D.E.

[DAAC18] Libri et al. 2018. **DiG: Enabling Out-of-Band Scalable High-Resolution Monitoring for Data-Center Analytics, Automation and Control**
[ANDARE18] Libri et al. 2018. **Evaluation of NTP / PTP Fine-Grain Synchronization Performance in HPC Clusters**
[AICAS18] Borghesi et al. **Online Anomaly Detection in HPC Systems**

Broker

MQTT

MQTT

DiG SW Daemons

Pow_pub | IPMI_pub | OCC_pub

D.A.V.I.D.E. Front-End

PSU_pub | Cooling_pub | Slurm_pub

IPMI

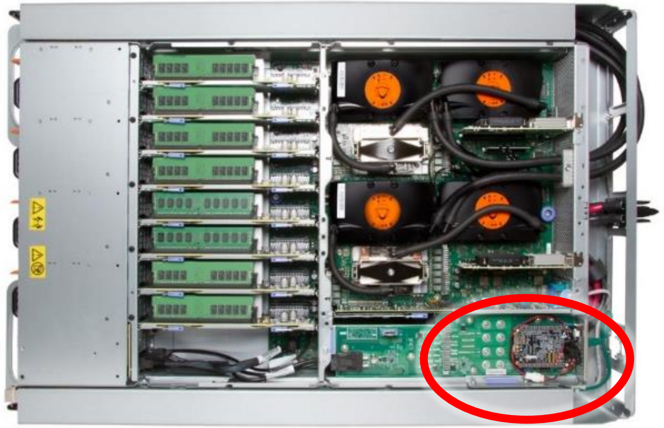

DiG

EXAMON

Liteon

Asetek

slurm workload manager

ETH zürich

Low overhead, accurate monitoring

ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

Real-time Frequency analysis on power supply and more…a live oscilloscope

- For instance, using the FFT we plot the power spectral density of the power benchmark of two applications, and we can distinguish them by the harmonics present in each of the signals

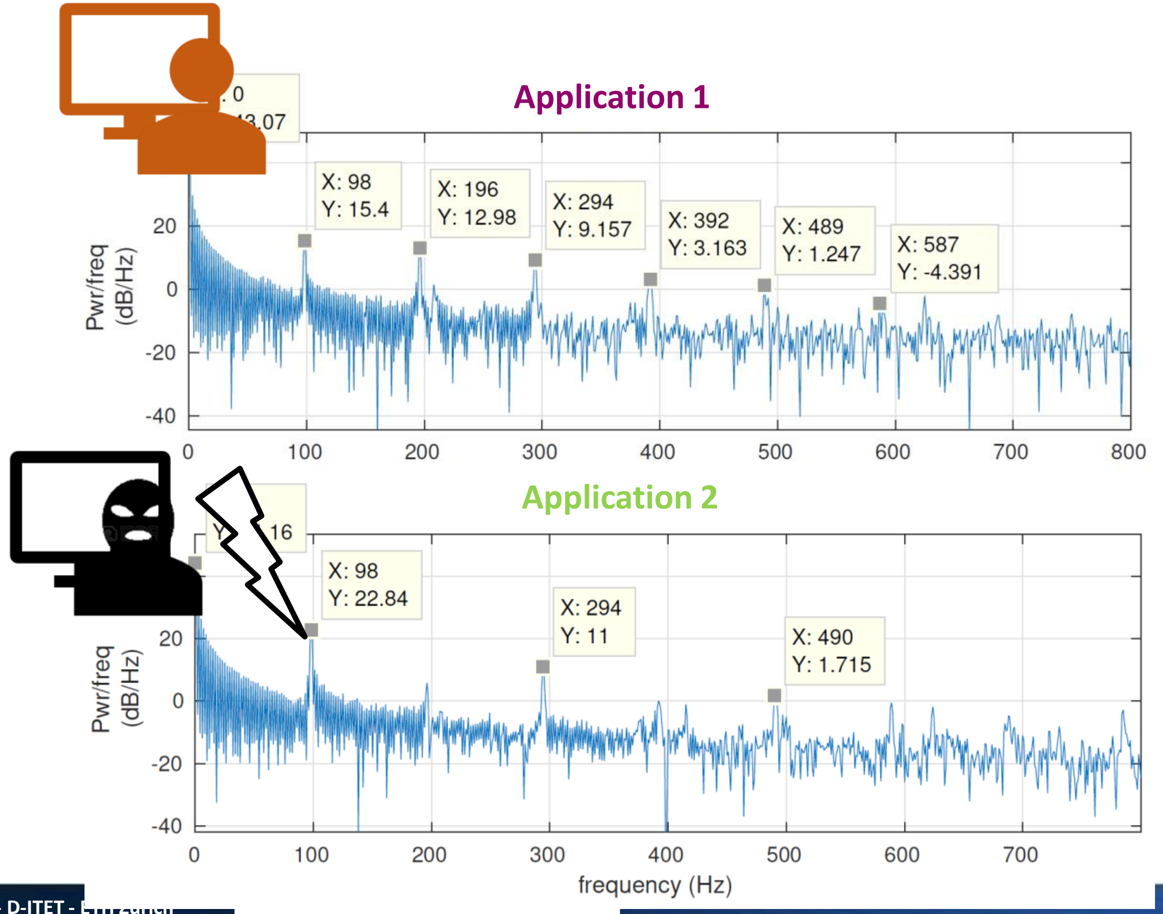

Spectral signature of an application!

Interesting feature for node level and system level Intrusion Detection System [IDS] !

PSD can be computed on the edge
On-going research: IDS based on in real-time PSD

Anomaly Detection

Not only power spectral densities, but also heterogenous sensors.

How to leverage them in real-time for anomaly detection ?

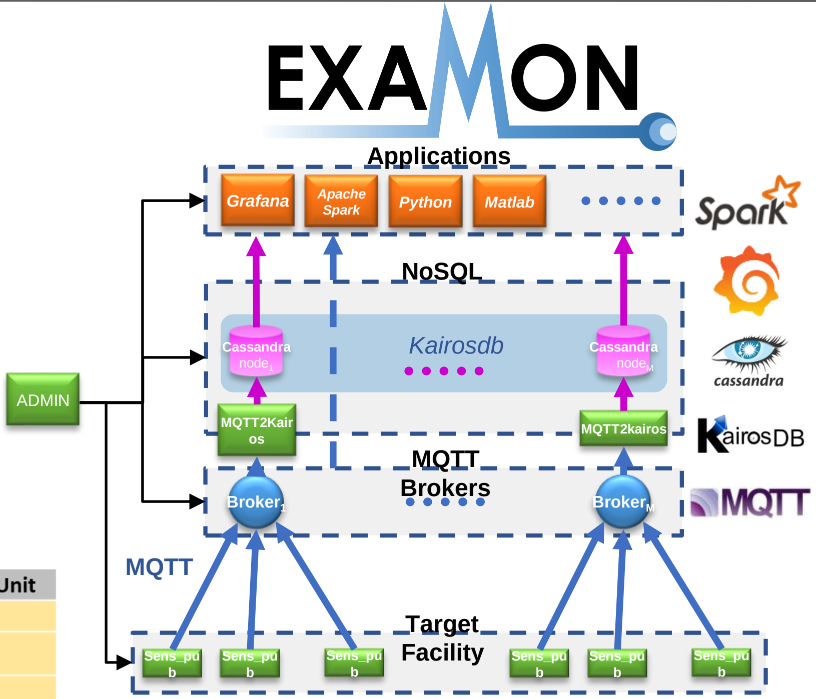

Metric Name	Description	Unit
Ambient_Temp	Node ambient temperature	°C
CPU_Core_Temp_1,...,CPU_Core_Temp_24	Core temperature	°C
CPU_Diode_1, CPU_Diode_2	Package temperature (Diode)	°C
CPU1_Temp, CPU2_Temp	Package temperature	°C
DIMM1_Temp,...,DIMM32_Temp	DIMMs temperature	°C
GPU_Temp_1,...,GPU_Temp_4	GPU temperature	°C
Mem_Buf_Temp_1,...,Mem_Buf_Temp_8	Memory temperature (Centaur)	°C
CPU_VDD_Curr	CPU current	A
Fan_1,...,Fan_4	Fan speed	RPM
CPU_VDD_Volt	CPU Voltage	V
Fan_Power	Fan power	W
GPU_Power	GPU power	W
Mem_Cache_Power	Memory power (Centaur)	W
Mem_Proc0_Pwr, Mem_Proc1_Pwr	DIMMs power	W
PCIE_Proc0_Pwr, PCIE_Proc1_Power	PCIExpress power	W
Proc0_Power, Proc1_Power	CPU Power	W
System_Power	Node total power	W

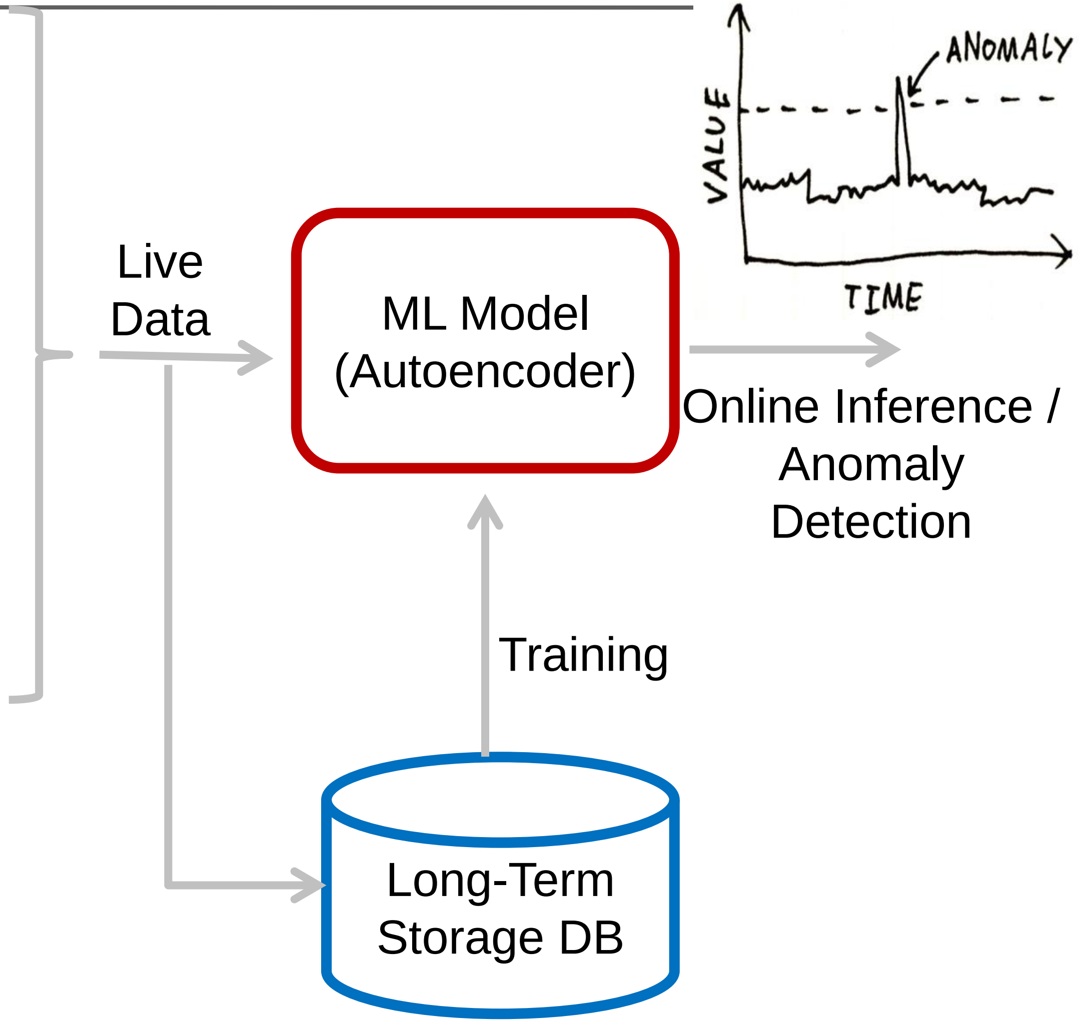

[EAAI19] Borghesi et. al «***A semisupervised autoencoder-based approach for anomaly detection in high performance computing systems***»
[AICAS18] Borghesi et al. «***Online Anomaly Detection in HPC Systems***»

Auto-encoders

An **auto-encoder** is a neural network that learns a representation of its input and is capable to reconstruct it

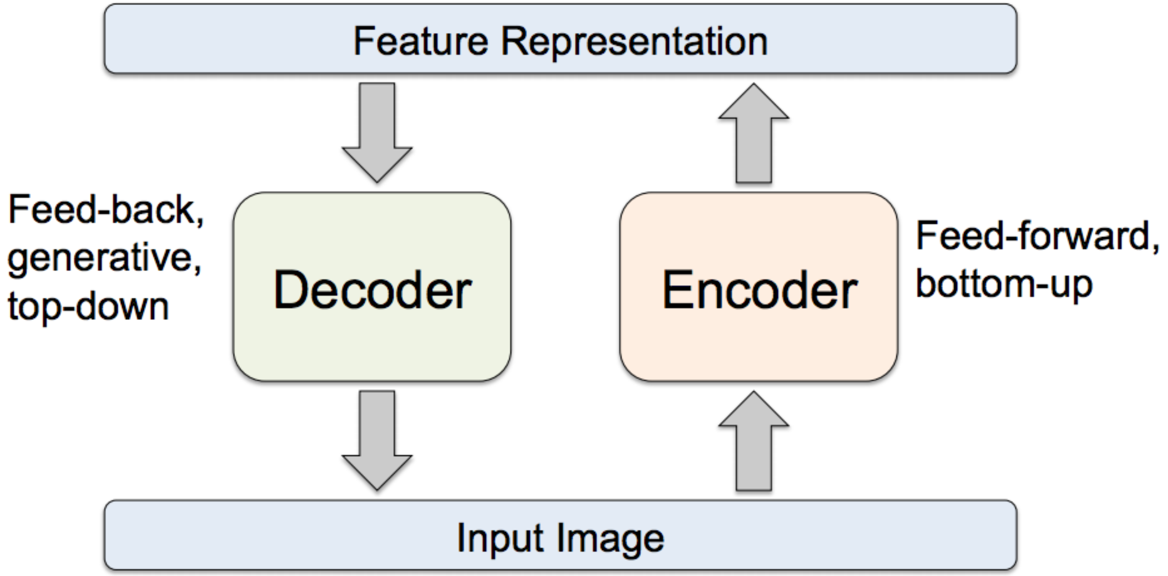

[Figure by R. Salakhutdinov]

Unsupervised learning

Minimize the reconstruction error

IDEA: train an autoencoder with the normal behavior of a HPC system and use its reconstruction error to detect anomalies

Threshold-Based Detection (1)

- How to discriminate anomalous data points from normal ones using the reconstruction error?
- Key observation: the errors distributions of normal examples and anomalous examples are very different

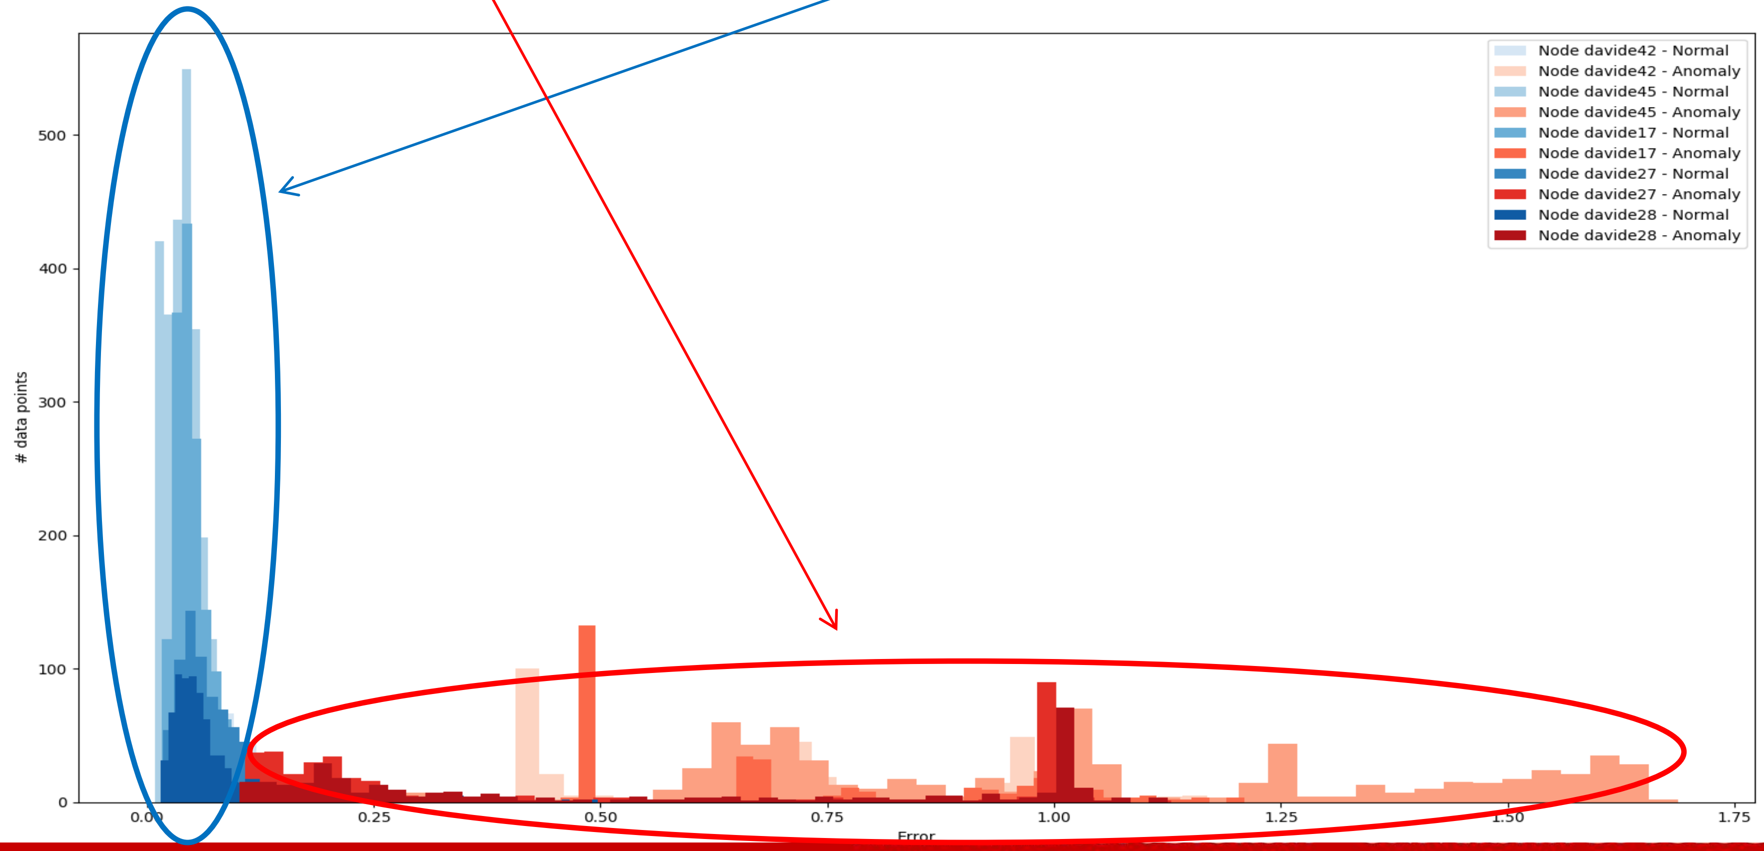

Comparison with semi-supervised techniques

[EAAI19] Borghesi et. al «A semisupervised autoencoder-based approach for anomaly detection in high performance computing systems»

Node	*GMM*				*EE*	*IF*	*SVM*		*AE*
	Diag	Spher	Tied	Full			Poly	RBF	Dedicated
davide10	0.916	0.927	0.924	0.922	0.911	0.941	0.262	0.811	**0.946**
davide11	0.91	0.909	0.921	0.923	0.801	0.946	0.237	0.65	**0.963**
davide12	0.864	0.865	**0.887**	0.191	0.154	0.278	0.216	0.608	0.747
davide13	0.897	0.876	0.892	0.904	0.434	0.953	0.16	0.66	**0.959**
davide16	0.854	0.515	0.855	0.884	0.915	0.923	0.606	0.926	**0.99**
davide17	0.272	0.267	0.269	0.509	0.914	0.929	0.613	0.931	**0.991**
davide18	0.882	0.858	0.888	0.875	0.715	0.923	0.614	0.933	**0.99**
davide19	0.887	0.523	0.524	0.909	0.762	0.919	0.624	0.941	0.99
davide26	0.893	**0.895**	0.894	**0.895**	0.376	0.701	0.218	0.609	0.846
davide27	0.165	0.162	0.161	0.922	0.825	0.652	0.389	0.656	0.9
davide28	0.926	0.756	0.788	**0.939**	0.773	0.912	0.406	0.635	0.892
davide29	0.843	0.841	0.799	0.842	0.882	0.89	0.455	0.92	**0.981**
davide42	0.722	0.356	0.718	0.295	0.793	0.853	0.727	0.935	**0.99**
davide45	0.394	0.561	0.518	0.752	0.627	0.67	0.661	0.933	**0.99**
Average	0.745	0.665	0.717	0.769	0.706	0.821	0.442	0.796	**0.935**

12% improvement compared to 2nd best technique

Outline

- Motivation & Datacentre Automation
- Cineca Roadmap & Bologna Science Park
- Datacentre automation
- OoB Fine-Grain Power Sampling (DiG) & Anomaly detection
- **Fine-Grain Energy-Management (Countdown)**

Usage Scenario #2 – Energy Efficiency

Energy Saving !!!

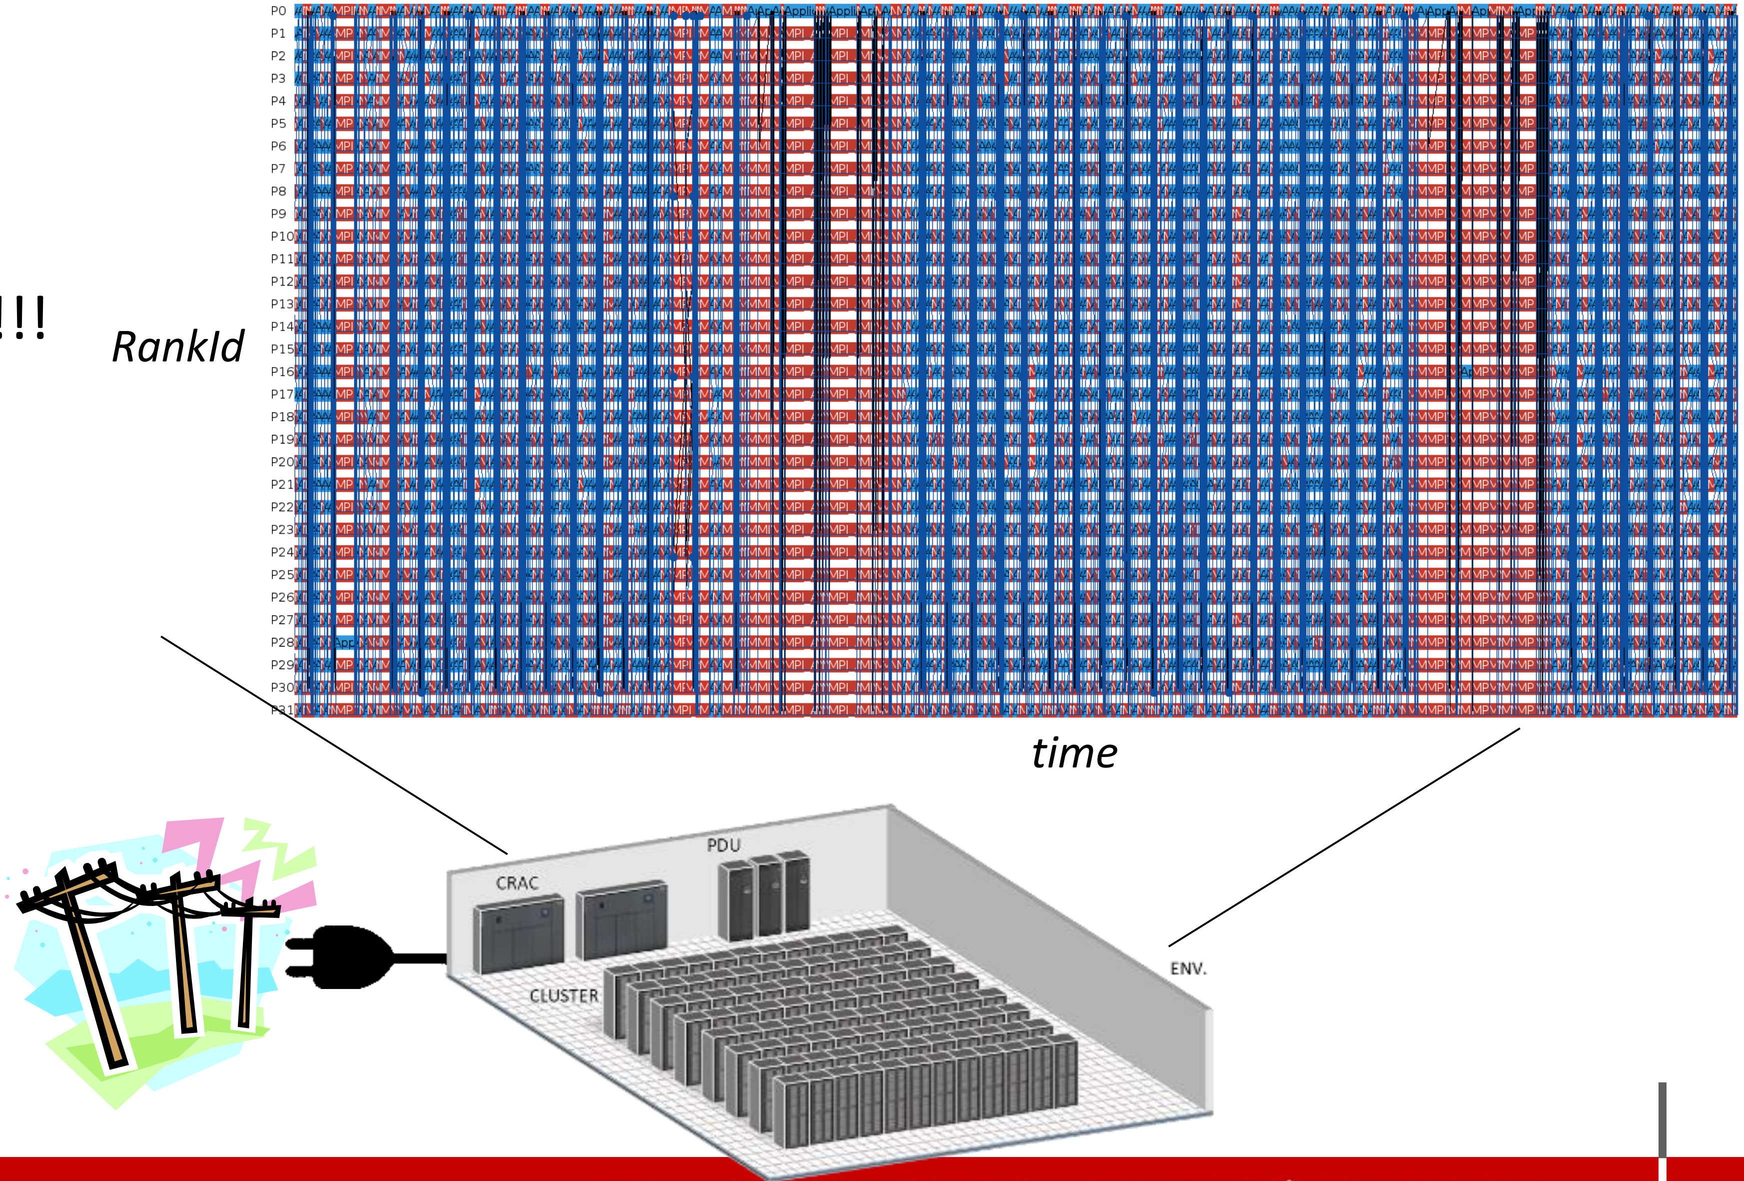

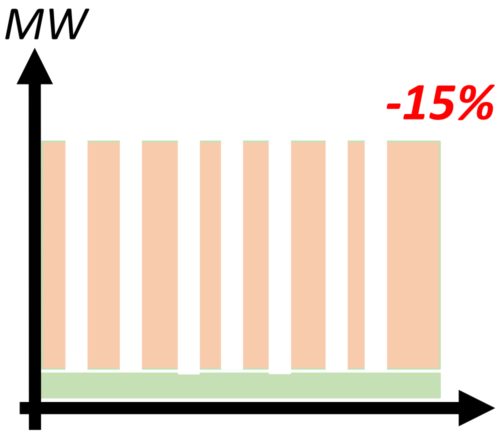

Application-aware energy management

COUNTDOWN is a tool to identify and automatically reduce the power consumption of the computing elements during communication and MPI primitives. It is based on ultra fine grain capabilities for profiling.

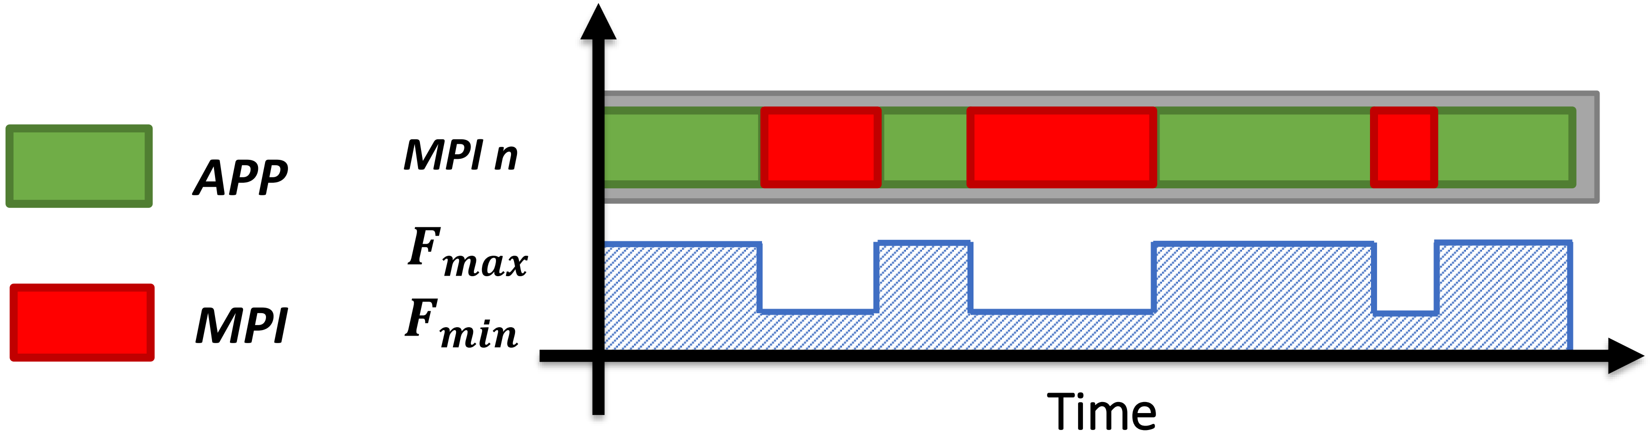

COUNTDOWN does not impact on the application tasks but only on the communication phases!

Reactive policy without MPI phase awareness!

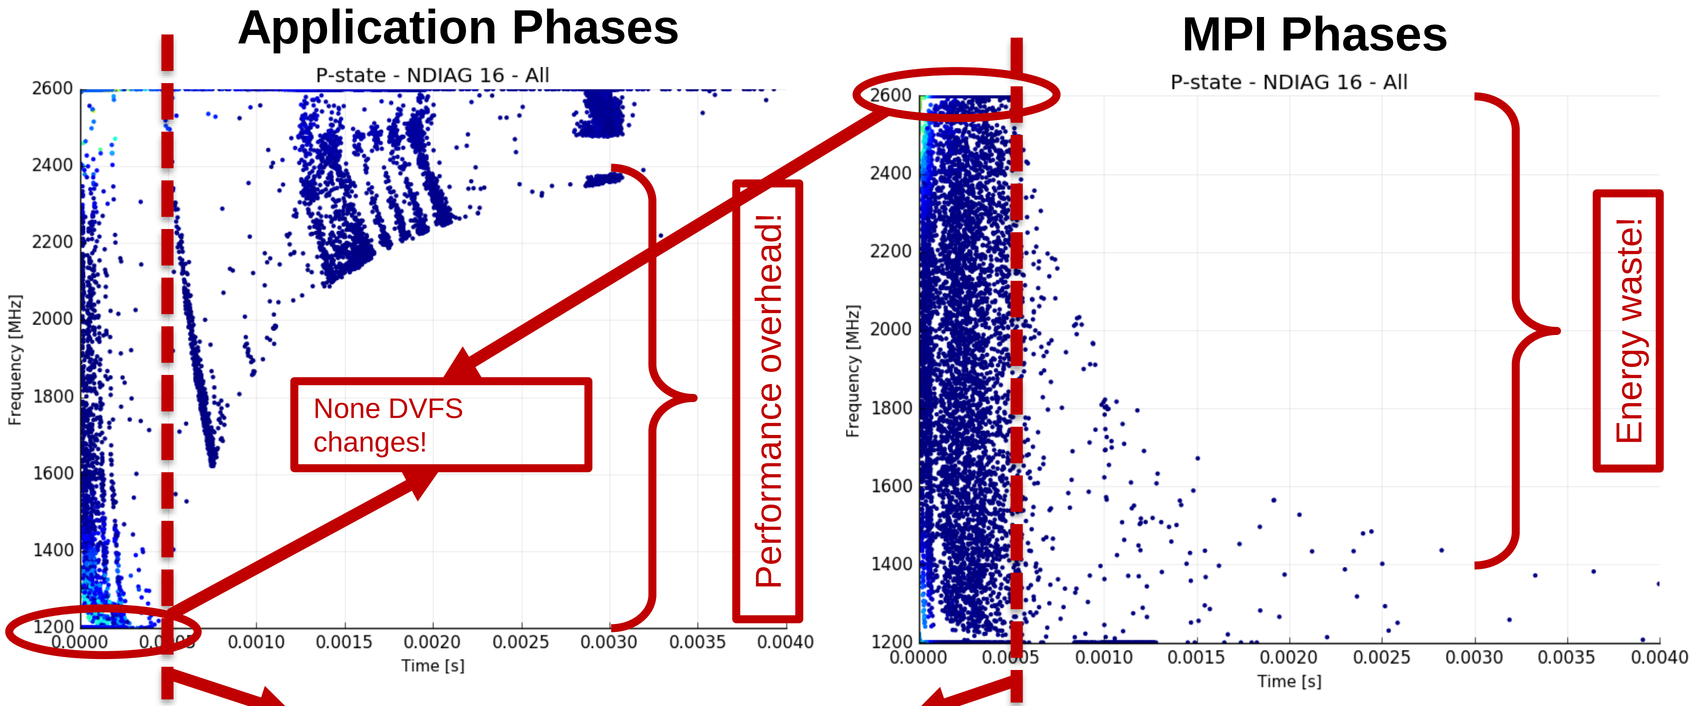

> 500us APP/MPI phases manifest sensitivity to DVFS changes

QUANTUMESPRESSO + DVFS for MPI call +

Dual socket Intel Haswell E5-2630 v3, 8 cores at 2.4 GHz (16 cores),
85W TDP, 128 GB DDR3 RAM

Intel DVFS Power Manager

Today’s HW power manager of Intel Architectures is quite slow in frequency variation! Literatures studied this mechanism and, for reverse engineering, discovered a 500us latency!

2015 IEEE International Parallel and Distributed Processing Symposium Workshop

An Energy Efficiency Feature Survey of the Intel Haswell *Processor

Daniel Hackenberg, Robert Schöne, Thomas Ilsche, Daniel Molka, Joseph Schuchart, Robin Geyer
Center for Information Services and High Performance Computing (ZIH)
Technische Universität Dresden – 01062 Dresden, Germany
Email: {daniel.hackenberg, robert.schoene, thomas.ilsche, daniel.molka, joseph.schuchart, robin.geyer}@tu-dresden.de

VI. P-STATE AND C-STATE TRANSITION LATENCIES

A. P-State Transition Latencies

The introduction of integrated voltage regulators, per core frequency domains, and improvements in the power control unit (PCU) have a direct influence on the latency and duration of ACPI processor state [25] transitions. To examine the new architecture, we use FTaLaT [26] for p-states and the tools developed by Schöne et al. [27] for c-states. We modified FTaLaT in the following ways:

- The original FTaLaT reads `scaling_cur_freq` from the Linux *cpufreq* subsystem to verify frequency settings. However, these readings are not reliable indicator for an actual frequency switch in hardware. We therefore add a verification by reading the `PERF_COUNT_HW_CPU_CYCLES` perform...

500us

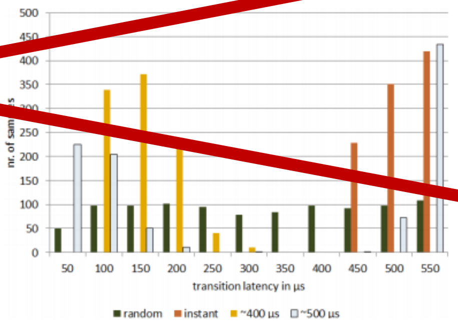

Fig. 3. Histogram of frequency transition latencies for switching between 1.2 and 1.3 GHz, depending on the time since the last frequency change.

therefore take 1,000 measurements for a single pair of start and target frequencies. We chose 1.2 and 1.3 GHz, but other frequency pairs yield similar results.

Figure 3 depicts the results of four experiments with 1,000 results each as a histogram. With frequency change requested at random times, the resulting latency is evenly distributed between a minimum of 21 µs and a maximum of 524 µs. Requesting a frequency transition instantly after a frequency change has been detected leads to around 500 µs in the majority of the results. If we introduce a 400 µs delay after the last frequency change, the transition time is typically about 100 µs. If the delay is in the order of 500 µs, the transition latencies can be split into two different classes–some yield an immediate frequency change while others require over 500 µs.

These results indicate that frequency changes only occur in regular intervals of about 500 µs. The distance between the start

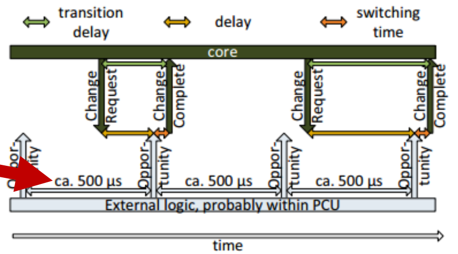

and the target frequency has negligible influence compared to the 500 µs delay. The assumed frequency changing mechanism is depicted in Figure 4.

* Intel Broadwell architectures as well!

COUNTDOWN Approach

COUNTDOWN implement an asynchronous mechanism based on a callback/timer to reduce the core's frequency after 500us in MPI primitives.

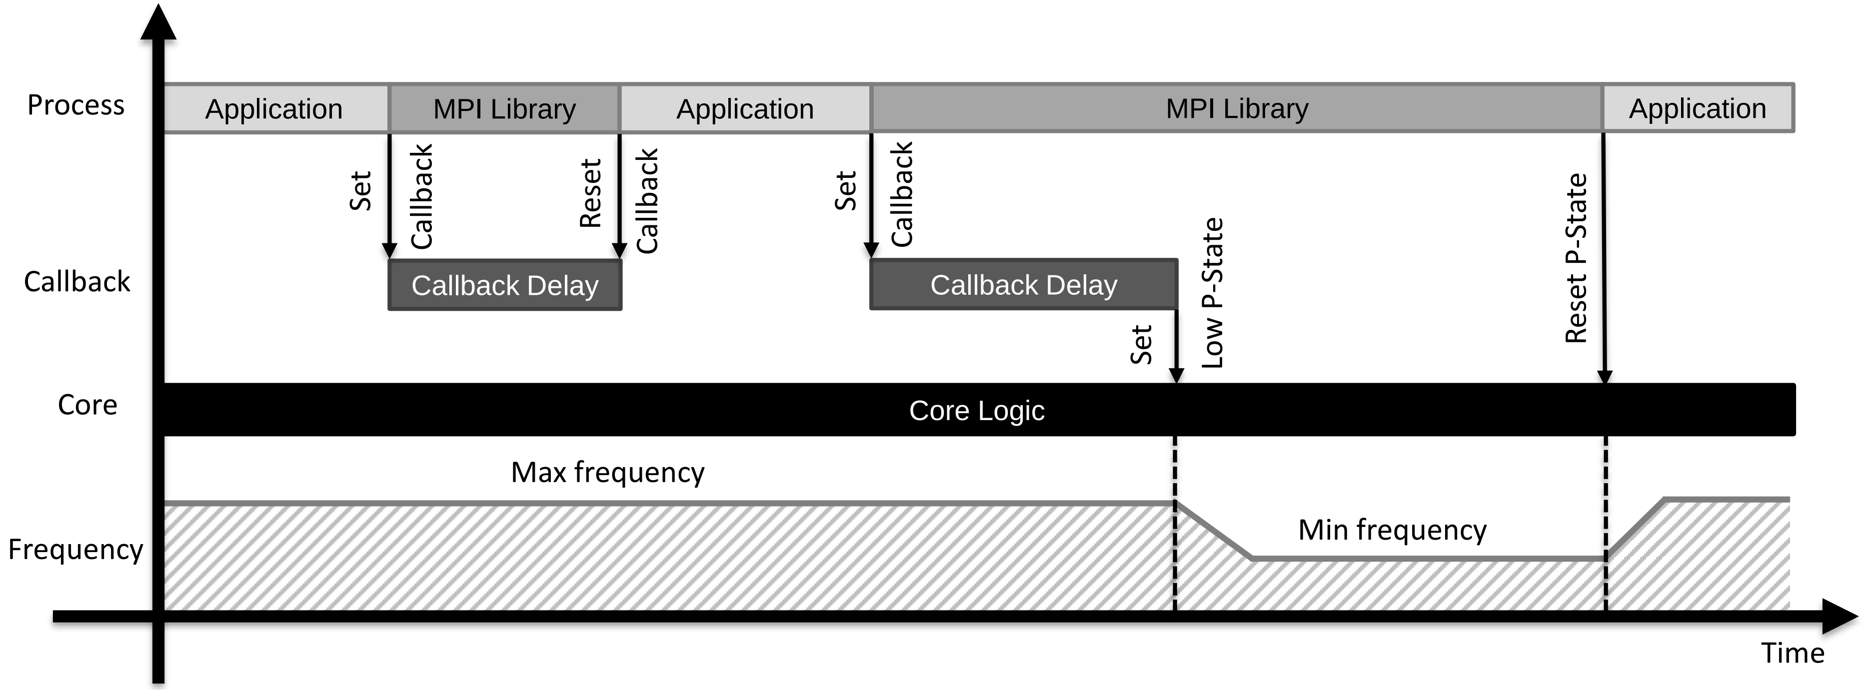

https://github.com/eeeslab/countdown

Multi Node - Target System

Galileo v1: Tier-1 HPC system based on an Lenovo NeXtScale cluster

516 Nodes: Dual socket Intel Haswell E5-2630 v3 CPUs with 8 cores at 2.4 GHz (85W TDP), DDR3 RAM 128 GB

Marconi A1 (Galileo v2): Tier-0 HPC system based on an Lenovo NeXtScale cluster

CINECA

400 Nodes: Dual socket Intel Broadwell E5-2697 v4 CPUs with 18 cores at 2.3 GHz (130W TDP), DDR4 RAM 128 GB

QUANTUMESPRESSO

NAS Benchmarks

NASA

OMEN

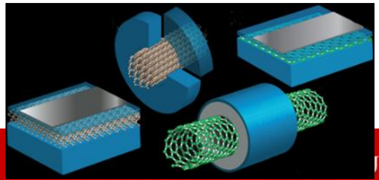

FORECASTING SYSTEM

COSMOS

Experimental results - QE PWscf

http://arxiv.org/abs/1806.07258
https://github.com/eeslab/countdown

QuantumEspresso PWscf

- **QE-PWscf-EU:** Expert User
- **QE-PWscf-NEU:** Not Expert User

Overhead: <6%

Energy/Power Saving: 22% - 43%

Computing Resources: 3456 Cores (96 nodes)

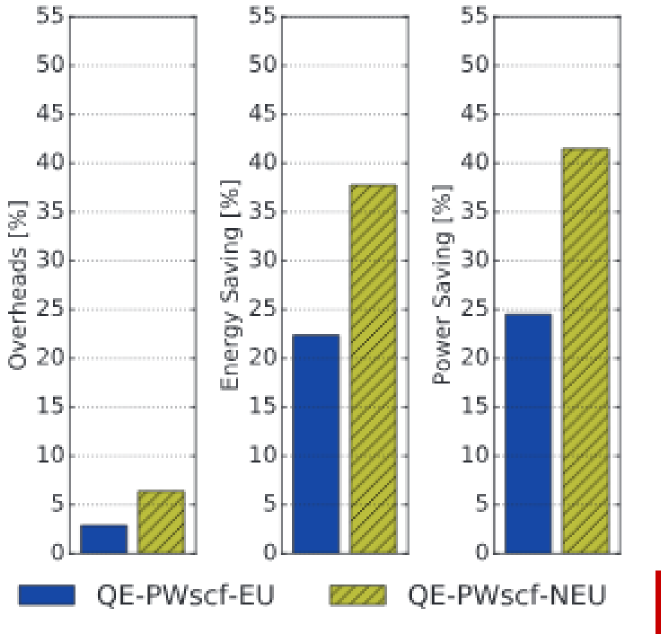

(a) QE-PWscf-EU

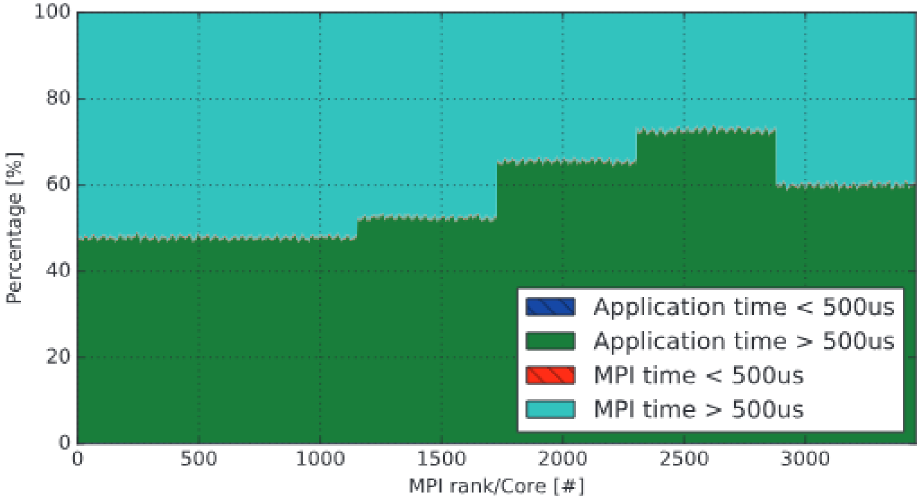

(b) QE-PWscf-NEU

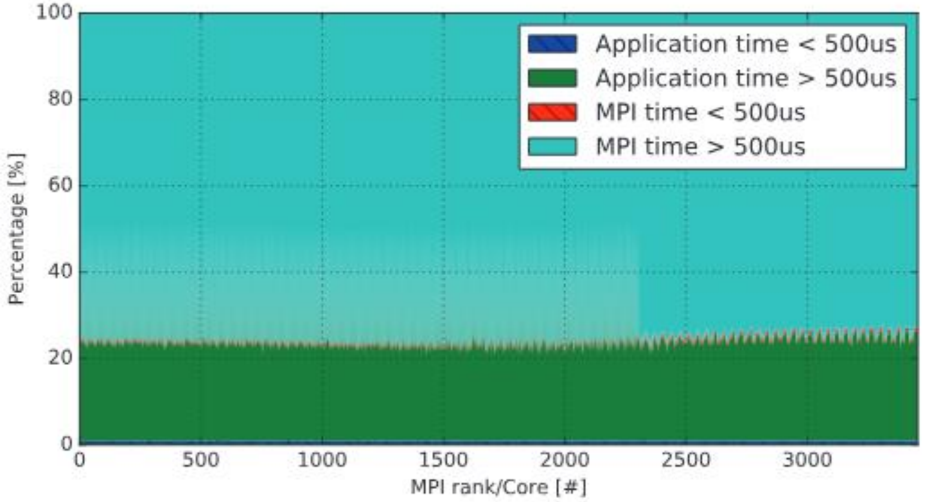

Promising results.

In production, only for a restricted set of users and low SLA.

Working to create safe job scheduler integration.

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Conclusion & Future Works

- Holistic and Fine Grain Monitoring feasible w. Open Source and Scalable of the shelf tools
- Challenge in the effective usage and knowledge extraction from the monitored data
- Fine Grain monitoring - AI and ML can leverage the infrastructure toward datacentre automation.
- Fine Grain power management can lead to important energy saving without sacrificing performance.
- Future works:
  - Scale up anomaly detection to datacentre level, and to security hazards
  - Fine-grain energy saving for other architectures and heterogenous nodes

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

Multitherman

ACKNOWLEDGE

EU FP7 ERC Project
MULTITHERMAN
(g.a.291125)

EU H2020 FETHPC
project ANTAREX
(g.a. 671623)

The UNIBO/ETHZ Datacenter Automation TEAM

- Luca Benini, Michela Milano, Andrea Tilli, Roberto Diversi, Cristian Conficoni, Andrea Borghesi, Daniele Cesarini, Antonio Libri, Francesco Beneventi, Moshen S. Ardebili

CINECA:

- Sanzio Bassini, Carlo Cavazzoni, Elda Rossi, Daniela Galletti, Alessio Mauri, Alessandro De Federico, Isabella Baccarelli, Marco Sbrighi, Giuseppa Muscianisi

ETH zürich

PRACE

CINECA

E4 COMPUTER ENGINEERING