HPC Research Project
CINECA - Unibo

Dr. Daniele Cesarini, HPC Software Engineer and Analyst, CINECA

School of Engineering, University of Bologna, 19 February 2020

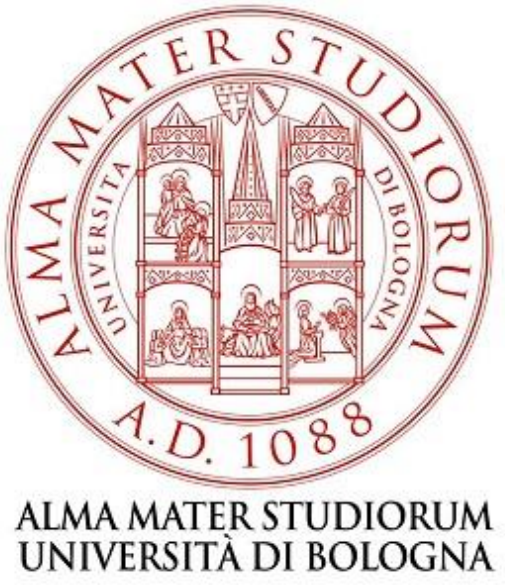

CINECA
Italian National Supercomputing Center

Cineca is a non-profit consortium, made up of 70 Italian universities, 4 national research centres, and the Ministry of Universities and Research (MIUR)
The consortium's institutional mission is to support the Italian scientific community through supercomputing and scientific visualisation tools.

#532 #68 #167 #23

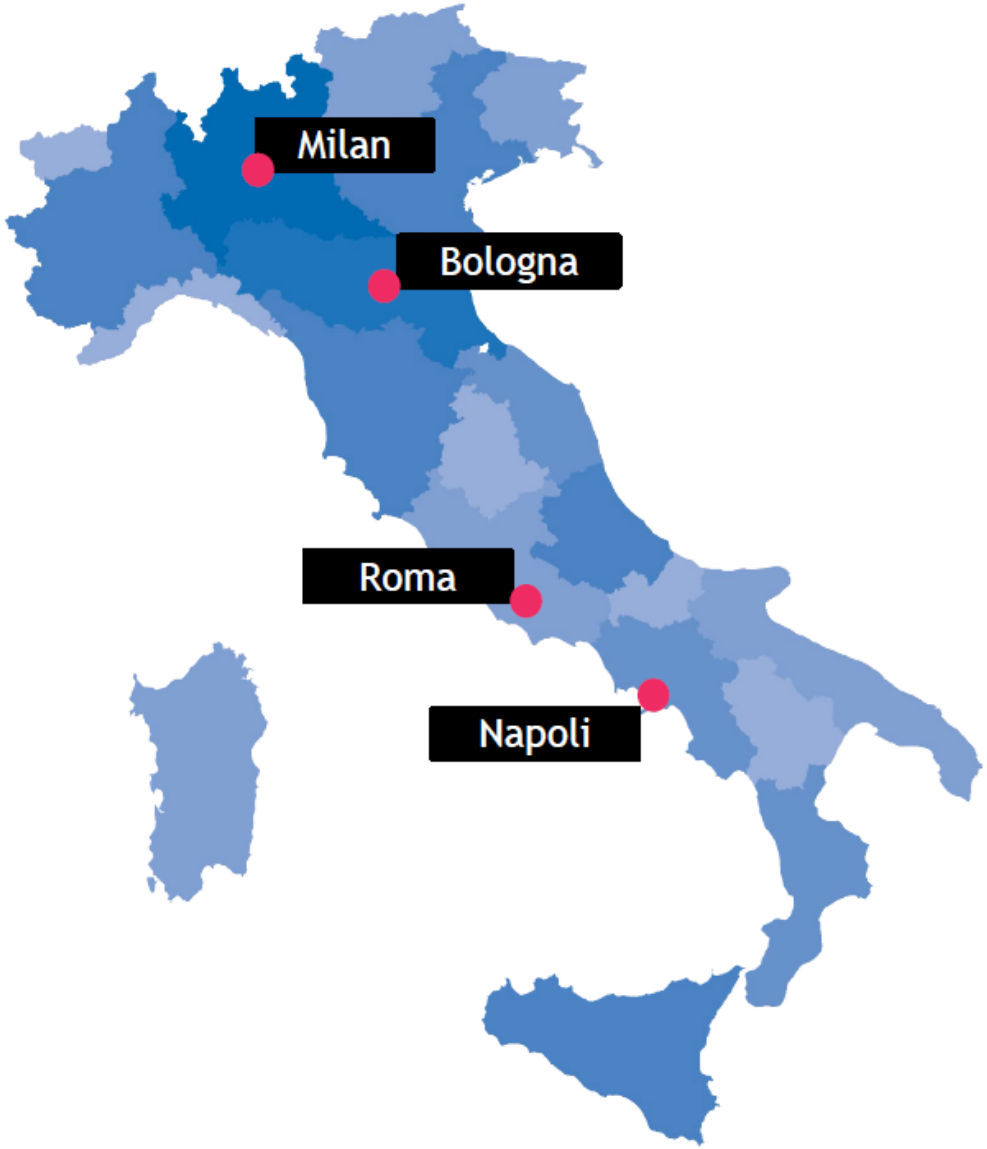

High-Performance Computing (Supercomputing)

> *High Performance Computing most generally refers to the practice of aggregating computing power in a way that delivers much higher performance than one could get out of a typical desktop computer or workstation in order to solve large problems in science, engineering, or business.*

Science

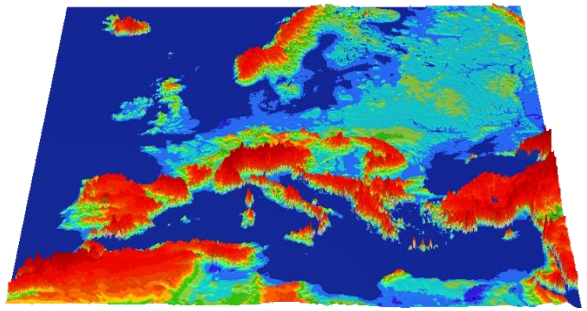

Engineering

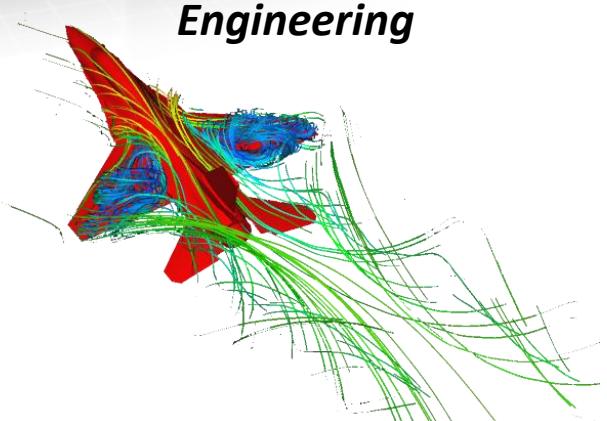

Business

CINECA HPC Systems

Marconi

- **Partition CPU:**
  - Nodes: 2,188
  - CPU: Intel Xeon Platinum 8160 CPU @ 2.10GHz
  - Cores: 105,024
  - DRAM: 410TB
  - Network: Omni-path 100Gbit/s
- **Partition GPU:**
  - Nodes: 1,000
  - GPU: 4,000 Nvidia Volta Tesla
  - Cores: 44,000
  - DRAM: 250TB
  - Network: Infiniband 200Gbit/s

Galileo

- Nodes: 1,022
- CPU: Intel Xeon E5-2697 v4 @ 2.30GHz
- Cores: 36,792
- DRAM: 128TB
- Network: Omni-path 100Gbit/s

CINECA – Data Center Power Consumption

Data Center at Casalecchio di Reno (BO)

Data Hall: 1.000 m²

Racks: ≈100

HPC nodes: ≈4000

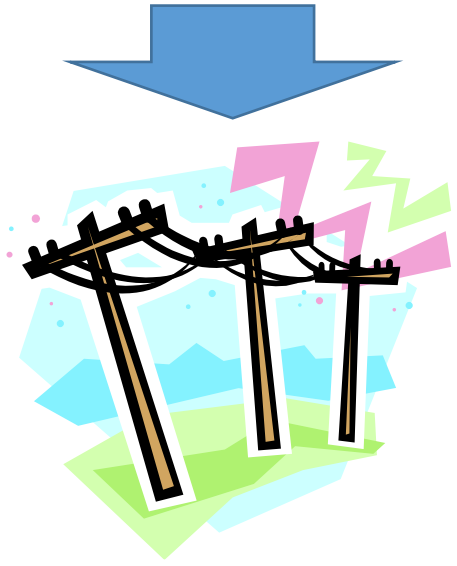

Electrical power committed:

5.0 MW - 30.000 MWh / year

EuroHPC hosting @ Bologna Science Park

CINECA + INFN Istituto Nazionale di Fisica Nucleare

ALMA MATER STUDIORUM UNIVERSITÀ DI BOLOGNA

Cooling equipment:

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

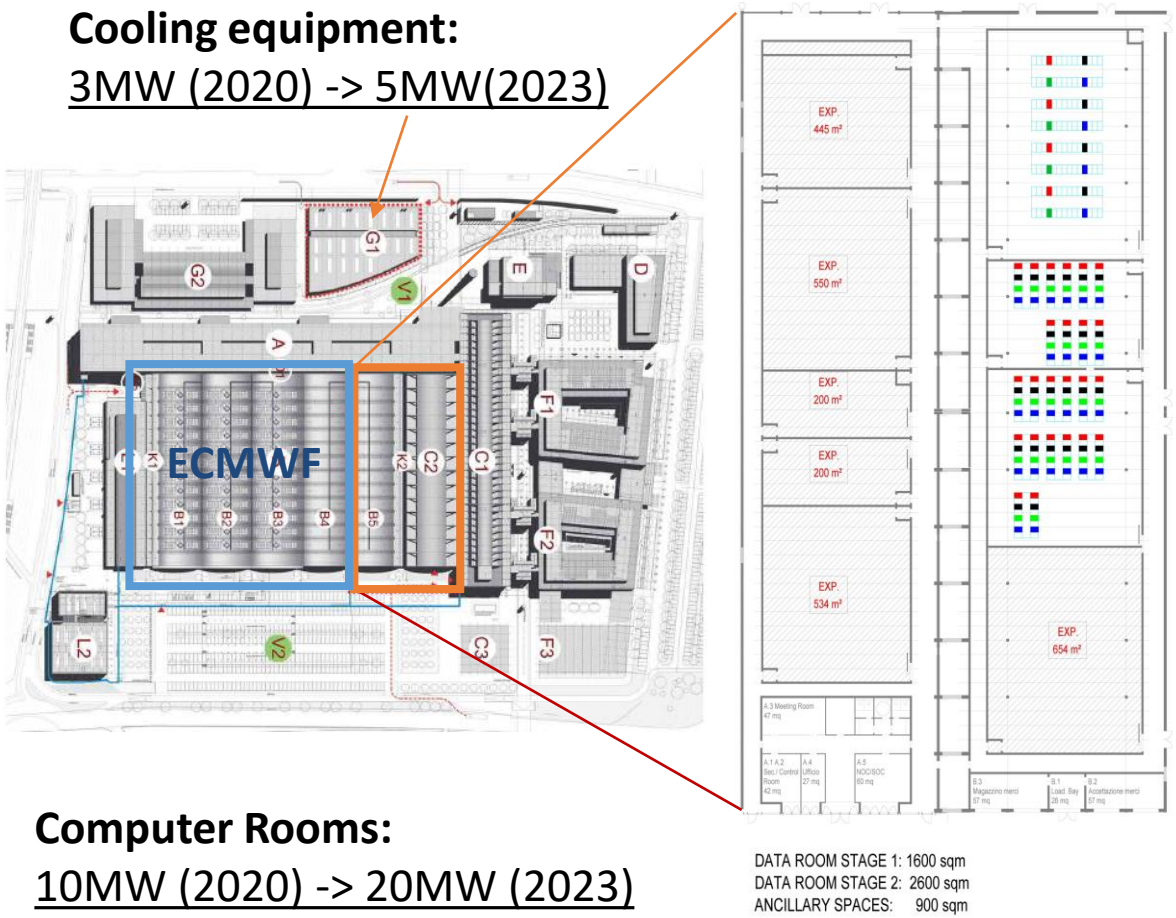

Computer Rooms:

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

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

Energy Saving

Energy Saving!!!

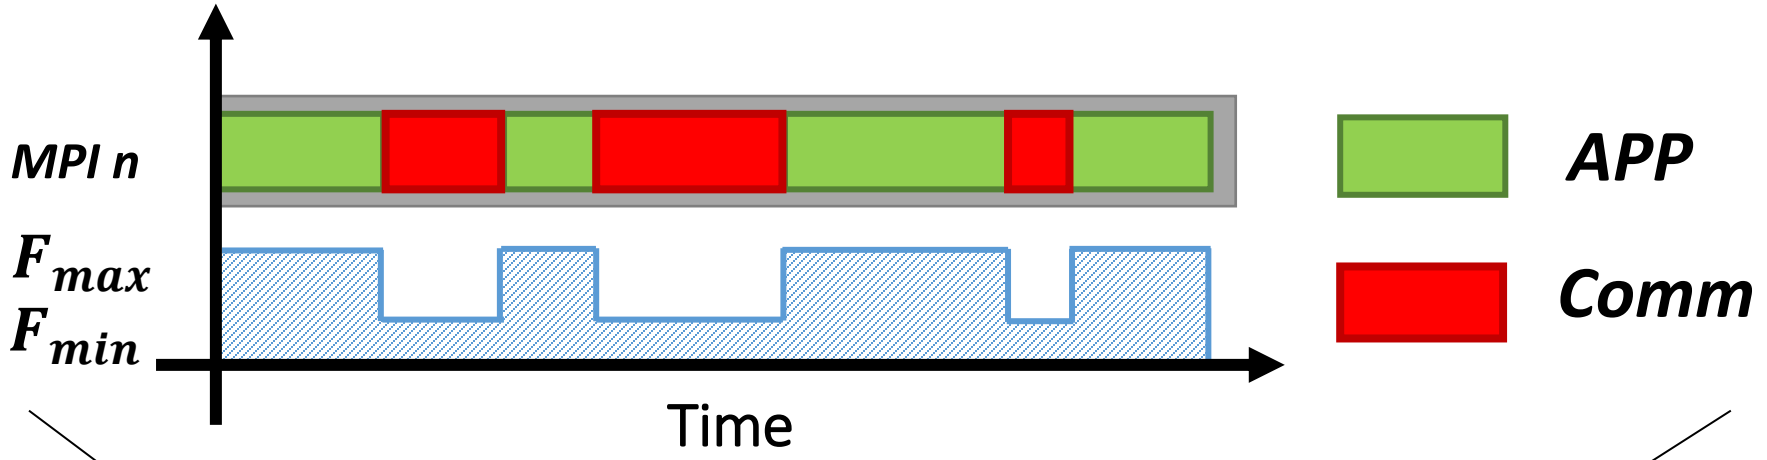

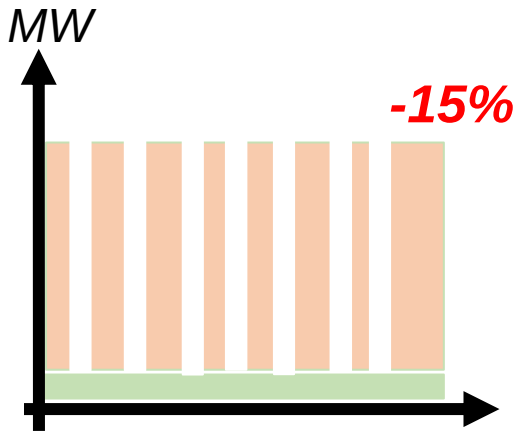

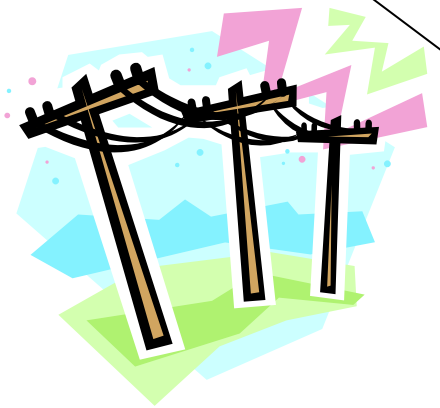

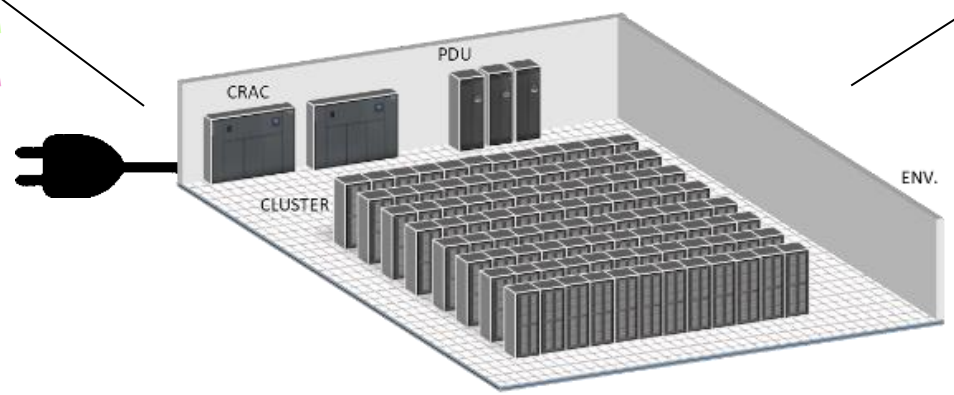

A new trend - Datacentre Automation

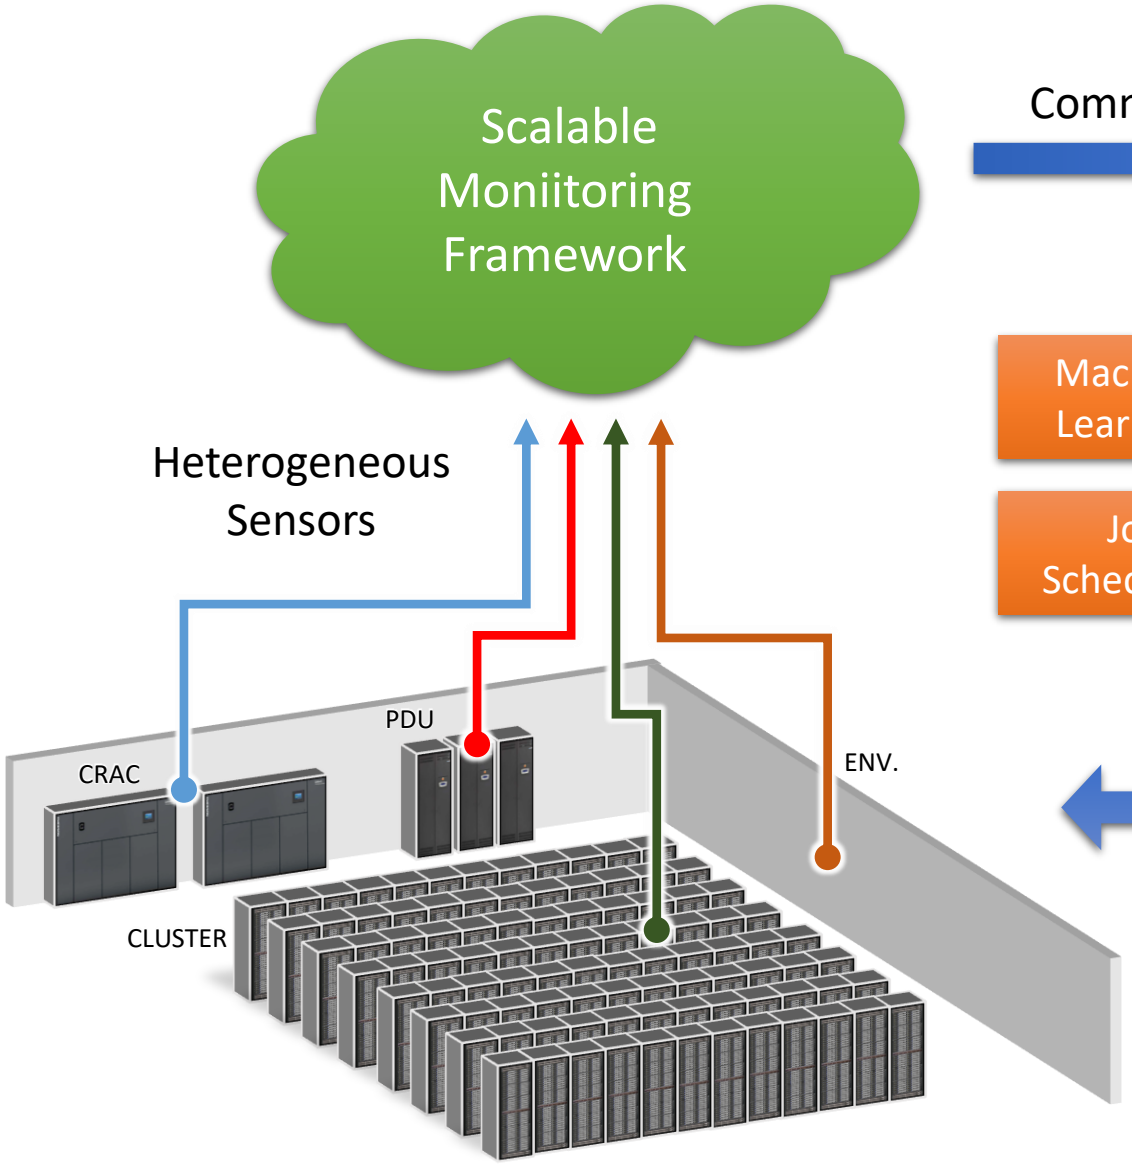

Common Interface

Machine Learning	Performance Analysis	Energy efficiency
Job Scheduling	Resources Management	Data Visualization

Reactive and Proactive
Feedbacks

Anomaly Detection

(1) how to overpass the lack of faulty data in production systems?

(2) Is it possible to train anomaly detection models without domain knowledge?

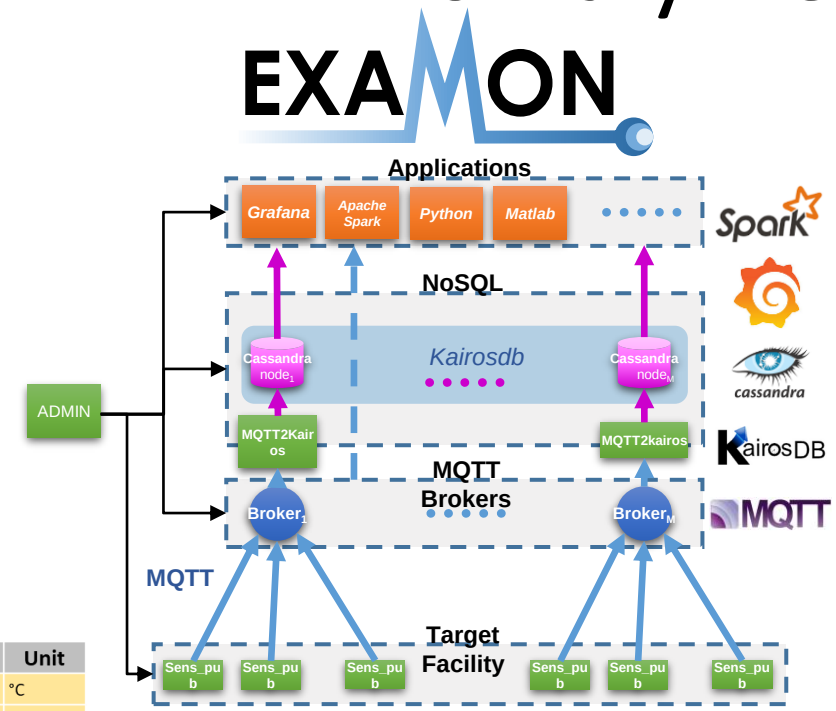

Metric Name	Description	Unit
Ambient_Temp	Node ambient temperature	°C
CPU_Core_Temp_1,...,CPU_Core_Temp_24	Core temperature	°C
CPU_Diode_1, CPU_Diode_2	Package temperature (Diode)	°C
CPU1_Temp, CPU2_Temp	Package temperature	°C
DIMM1_Temp,...,DIMM32_Temp	DIMMs temperature	°C
GPU_Temp_1,...,GPU_Temp_4	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

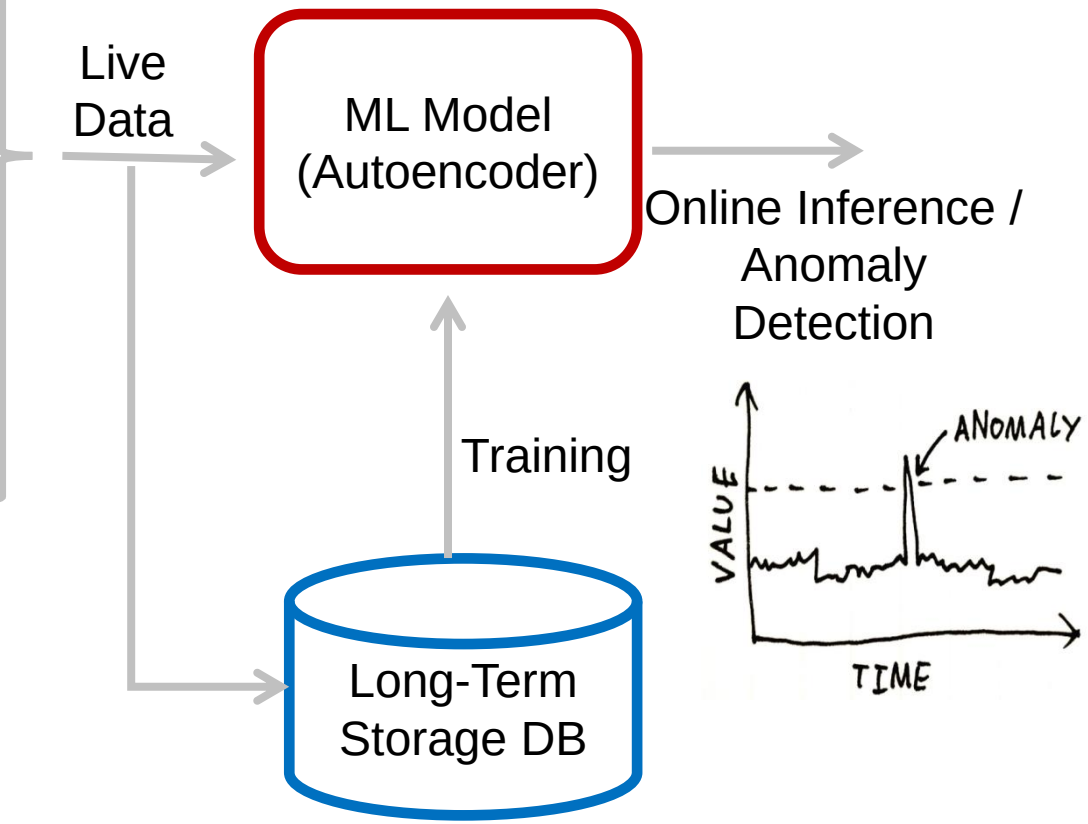

*[EAAI19] Borghesi et. al «**A semisupervised autoencoder-based approach for anomaly detection in high performance computing systems**»*

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

European Processor Initiative (EPI)

- On-board modularity:
  - CCIX links to adjacent sockets
- High-speed standard interconnect:
  - PCIe gen5
- Interposer/chiplet approach:
  - D2D links to adjacent chiplets:
    - Scale-out
    - Accelerators
- “near” memory partition:
  - High Bandwidth Memory on interposer

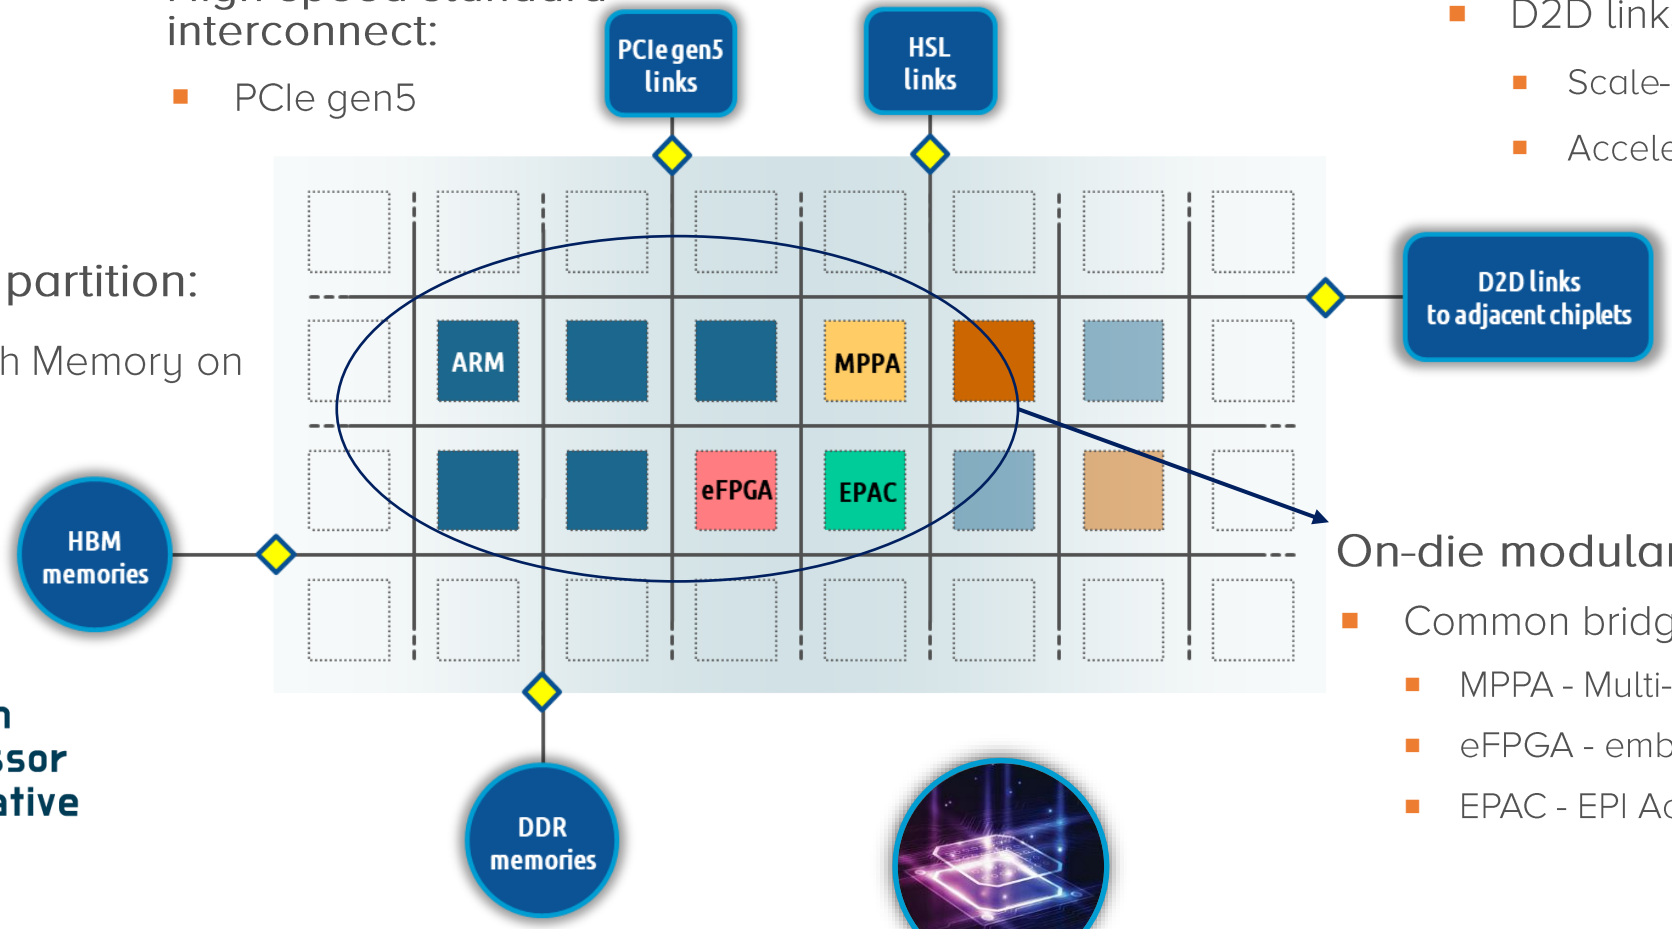

On-die modularity:

- Common bridge for:
  - MPPA - Multi-Purpose Processing Array
  - eFPGA - embedded FPGA
  - EPAC - EPI Accelerator

European Processor Initiative

Research Projects and Thesis

1. **Energy Efficient Runtime (COUNTDOWN):**
   a) Implementing power management strategy in OpenMP runtime and ARM+GPU systems.
2. **Porting and performance evaluation of HPC applications on high-performance ARM and Risc-V systems:**
   a) Numerical weather prediction (NWP)

Strong requirements:

- Good knowledge of C language

Topics, tools, and languages used in the project/thesis:

- CMAKE, GIT and Python
- C++/Fortran language and compilators (GNU, Intel, LLVM, etc.)
- HPC environment, parallel and distribute programming model (OpenMP, MPI, etc.)
- Performance evaluation and modeling
- Power management system for high-performance processors (DVFS, RAPL, P/C states, etc.)

Research Projects and Thesis

1. **Big Data & Deep Learning for HPC:**
   a) Explore AI solution for application acceleration on GPU.
   b) AI/DL on big data for datacenter automation.

2. **Deep Learning and edge AI:**
   a) Deep Reinforcement learning for power management
   b) HW acceleration for on-chip anomaly detection
   c) Edge AI for enhanced security

Strong requirements:

- ML and/or DL and/or AI and/or control theory background

Topics, tools, and languages used in the project/thesis:

- Python, C, C++, Bash, O.S., Cassandra DB, KairosDB.
- Power management system for high-performance processors (DVFS, RAPL, P/C states, etc.)

Work Location & Contact

Work Location

Energy-Efficient Embedded Systems Laboratory (EEES Lab)

Viale Carlo Pepoli 3/1, Bologna, Italy

Unibo Supervisor

Prof. Andrea Bartolini

a.Bartolini@unibo.it

CINECA Co-Supervisor

Dr. Daniele Cesarini

d.cesarini@unibo.it