Finanziato dall'Unione europea NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

ICSC
Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

ICSC
Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

Spoke 1 Future HPC & Big Data
Stato delle attivita' e prospettive future

Luca Benini (UNIBO), Marco Aldinucci (UNITO)

Kick off meeting 25/26 novembre 2022, Bologna

Struttura e Organizzazione

- Two "**Spoke labs**" → research + infrastructure investment and living-labs
  - Bologna (UNIBO/CINECA →TECNOPOLO) HWS: future HPC HW and system prototyping [co-located with HPC facilities]
  - Torino SWI SW integration lab
- 4/5 Multi-partner "**Flagship Initiatives"**
  - Similar to EU projects →objectives, KPIs
- Cascade funding projects

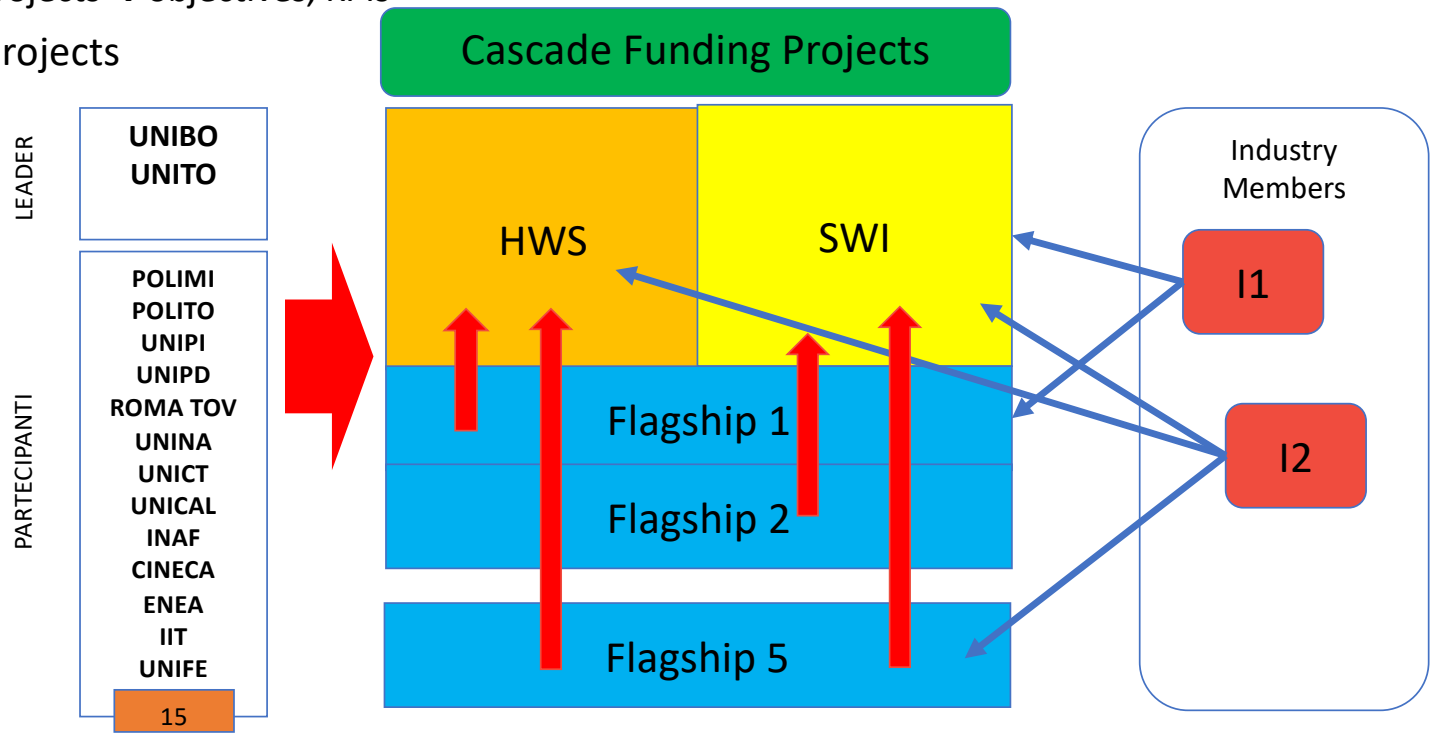

Flagships e Laboratori

Kick off meeting

- **HWS** Hardware and System Lab
- **SWI** Software and Integration Lab
- **FL1** Non-functional properties - Design exploration: energy, power reliability (performance portability)
- **FL2** Heterogeneous acceleration - architecture, tools, software
- **FL3** Workflow Management System, integration cloud-HPC convergence, High-Performance Storage &IO (filesystems, data movement), Tools for digital twins
- **FL4** Trustworthiness, security, privacy (TEEs vs. homomorphic encryption…) open reference architecture for a Trusted Execution Environment
- **FL5** Codesign - application + SW + HW targeting, benchmarking, patterns, microkernels

Luca Benini, Marco Aldinucci

Collaborazione con le Aziende

Kick off meeting

LEADER

UNIBO
UNITO

PARTECIPANTI

POLIMI
POLITO
UNIPI
UNIPD
ROMA TOV
UNINA
UNICT
UNICAL
INAF
CINECA
ENEA
IIT
UNIFE

15

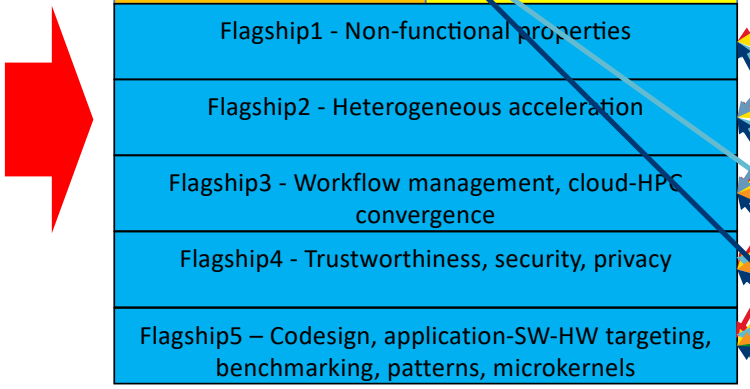

Cascade funding +
Innovation projects

Aziende ed enti di Ricerca Esterni

autostrade per l'italia
1.Evaluate future HPC HW architectures and edge computing
2.Video surveillance system using ML/AI

LEONARDO
1.Green Computing
2.Accelerated Computing, OpenSource HW and SW
3.Containers for Cloud continuum, Edge and IoT
4.HPC -Cloud Convergence for Data driven Digital Twin, BigData and AI

Eni
1.Scouting future HPC HW architectures and system prototyping
2.Parallel programming models and HPC system software

ENGINEERING
1.AI andBigData frameworks, integration in HPC ecosystem, federated, distributed learning
2. Mini-applications and benchmarking from multiple domain;

INTESA SANPAOLO
1.Enterprise Multi-tenancy Distributed learning
2.Enterprise AI and BigData frameworks for HPC ecosystem

ThalesAlenia Space
1.High-performance RISC-V processors for satellite
2.HPC technologies for SAR ground processing of Earth observation data
3.Cyber-security on-board spacecraft, including embedded AI framework and radiation tolerant AI processing in space

UnipolSai ASSICURAZIONI
1.Cloud-HPC convergence: federated and distributed learning
2.Hybrid MLOps Workflow management & cross-application streaming tools

FINCANTIERI
1.Hybrid HPC Solutions
2.HPC smart sensors and big data
3.Mini applications in naval domain
4.Solutions for onboard power monitoring and control

Luca Benini, Marco Aldinucci

Stato Milestones Organizzative

Luca Benini, Marco Aldinucci

Bologna, 25/26 novembre 2022

Spoke organizzazione amministrativa

Kick off meeting

Spoke Leader UNIBO: struttura di supporto amm/gest già operativa

ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

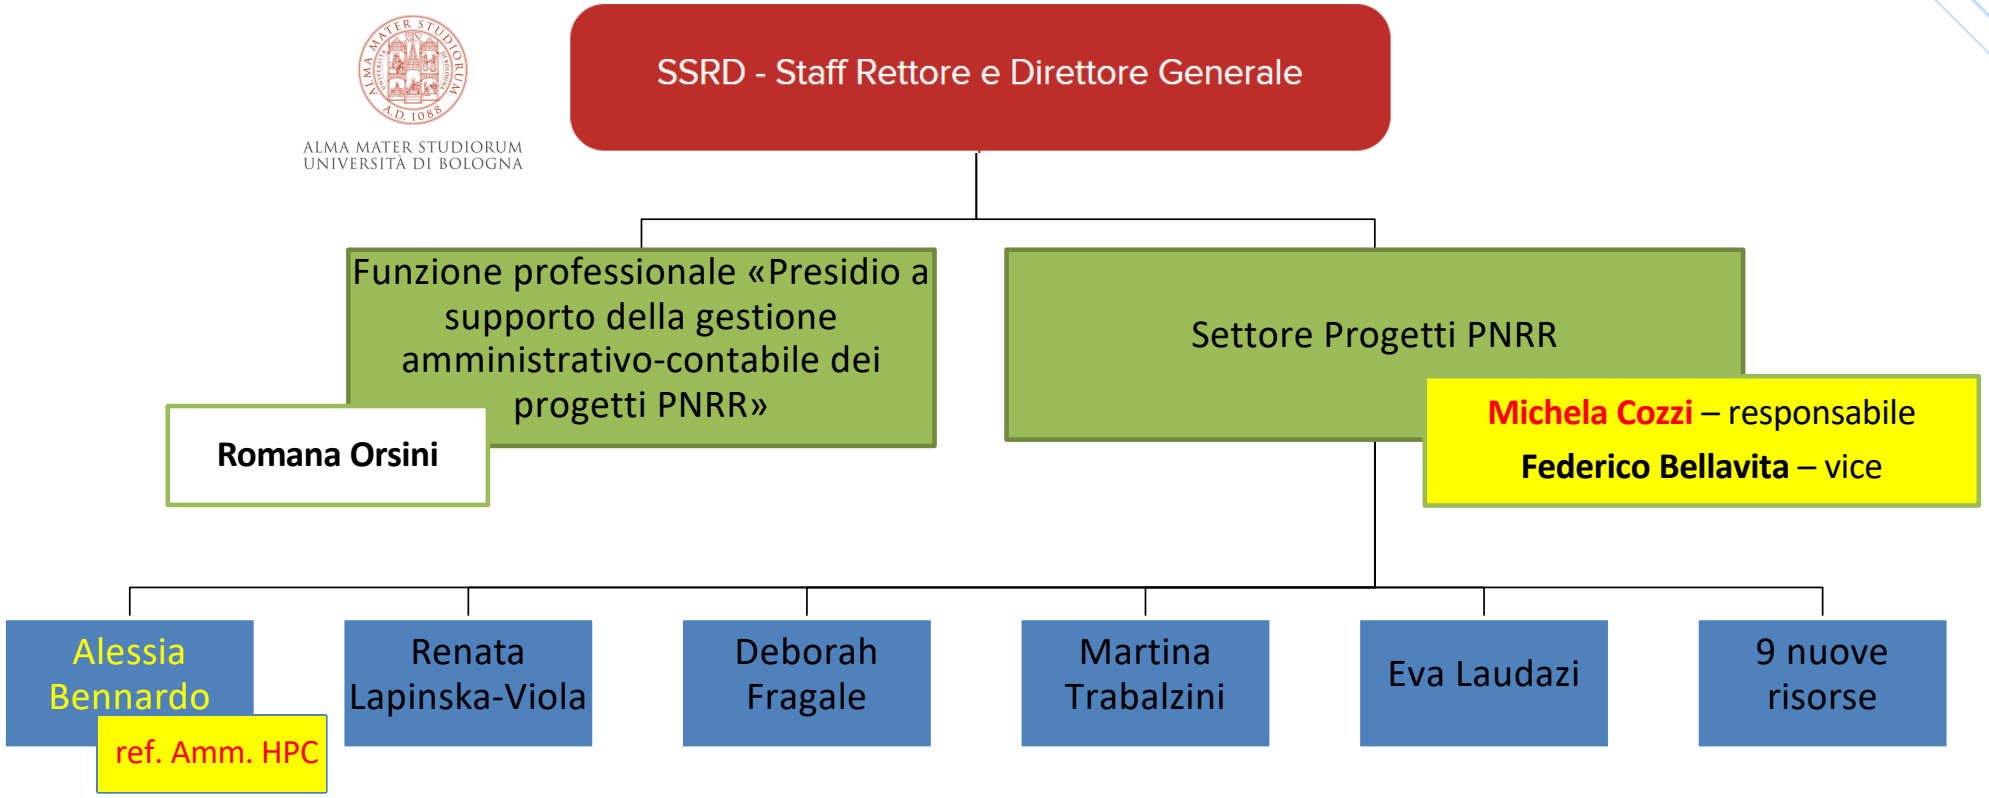

Luca Benini, Marco Aldinucci

Bologna, 25/26 novembre 2022

Milestone: Personale e Assunzioni

Kick off meeting

https://docs.google.com/spreadsheets/d/1ctex6T8ypsx4sViD6AE0kkB7weUFEkXwWeFCWybQhEk/edit#gid=0

LEADER

UNIBO
UNITO

PARTECIPANTI

POLIMI
POLITO
UNIPI
UNIPD
ROMA TOV
UNINA
UNICT
UNICAL
INAF
CINECA
ENEA
IIT
UNIFE

15

1. **UNIBO** 2 RTDA (2 banditi, procedure in corso), 3 Dott. (1 Bandito 38°, 2 39°) 1 AdR (bando 12/22)
2. **UNITO** 3 RTDA (3 banditi 11/22(=, 1 Dott (bandito 38°), 2 Adr (banditi 11/22)
3. **POLIMI** 2 RDTA (1 bandita, 1 da approvare in CdD), 2 Dott. (da bandire Q1-23), Attivata LM (HPC eng)
4. **POLITO** 2 RDTA (Attesa presa servizio), 2 Dott. (una presa servizio, un bando aperto)
5. **UNIPI** 3 RDTA (Banditi e in corso), 2 Dott. (Bandito 38° una assegnata)
6. **UNIPD** 1 RDTA (Approvato CdD), 1 Dott. (Bandito 38°)
7. **ROMA TOV** 1 RDTA (Approvato CdD), 1 Dott. (Bandito 38°), 1 Adr (attesa presa servizio)
8. **UNINA** 1 Dott. (Bandito 38°), 7 annualità di Contratti di Ricerca (3.5 posizioni per 2 anni – previste 23)
9. **UNICT** 2 RTDA (2 banditi, procedure in corso), 3 Dott. (da bandire 39o ciclo)
10. **UNICAL** 1 RTDA e 1 PhD (bandi in Cda 29/11/22, in uscita 12/22).
11. **INAF** 4 posizioni TD (in preparazione), 2 Dott (assegnati)
12. **CINECA** 2 Dott (da bandire nel 39°)
13. **ENEA** 2 RTDA (Banditi)
14. **IIT** 1 posizione bandita (9/22)
15. **UNIFE** 2 RTDA (Attesa prese servizio), 1.5 Dott (selezione in Corso)

Luca Benini, Marco Aldinucci

Stato delle Attività Scientifiche

Luca Benini, Marco Aldinucci

Bologna, 25/26 novembre 2022

Milestones - HWS Lab (UNIBO)

Kick off meeting

Mission: prototyping HW and system integration (software stack + power management and monitoring) **Technical focus**: Open (RISC-V) &Heterogeneous (accelerated) systems

luca.Benini@unibo.it

- Locali 300mq -500k€ per ristrutturazione, adeguamento, arredi, attrezzature (2023)
- Accesso aperto al primo prototipo cluster HPC RISC-V a livello mondiale (Monte Cimone)
- Primo Corso Big Data on RISC-V su MC (87195 - LAB OF BIG DATA ARCHITECTURES M - 3 cfu Docente Andrea Bartolini)

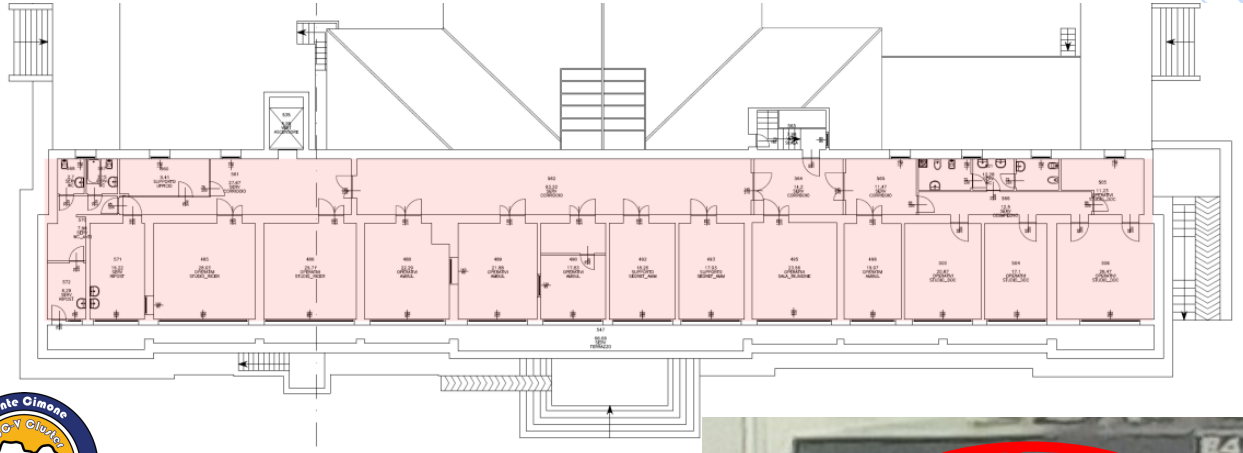

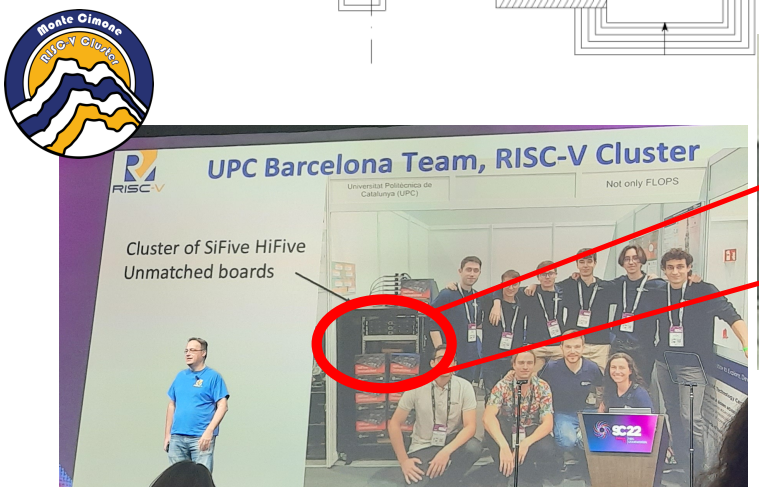

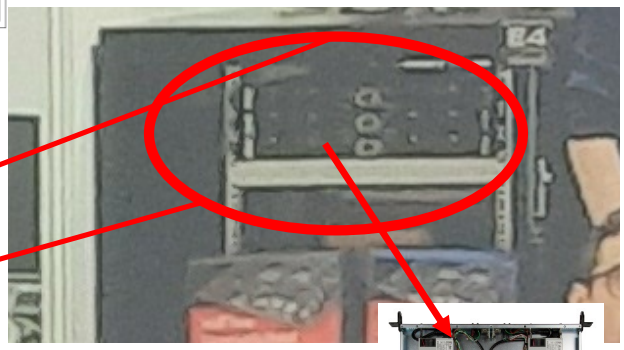

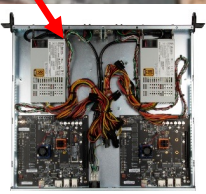

Luca Benini, Marco Aldinucci

Bologna, 25/26 novembre 2022

HWS Lab: UNIBO-CINECA → Examon

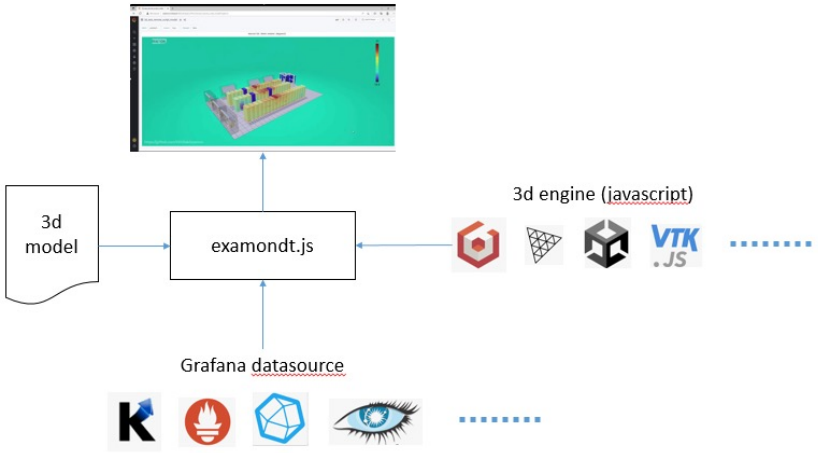

Data collection + 3D model, towards a digital twin of the HPC facility

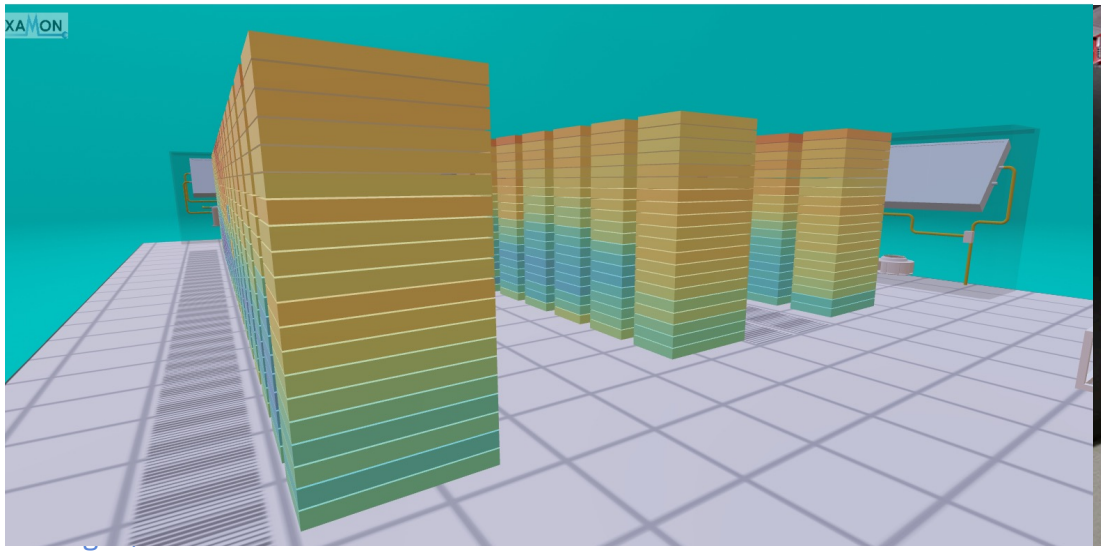

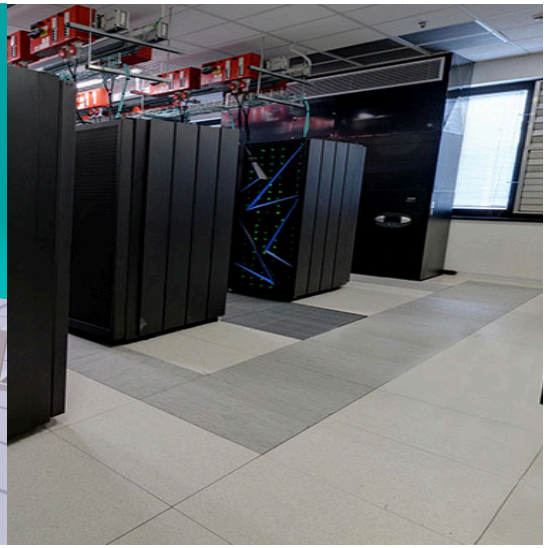

Examon in Production

- Examon is operational on Marconi100
- Deployment on Leonardo is under discussion
- Open-data release of multi TB dataset planned in the next weeks

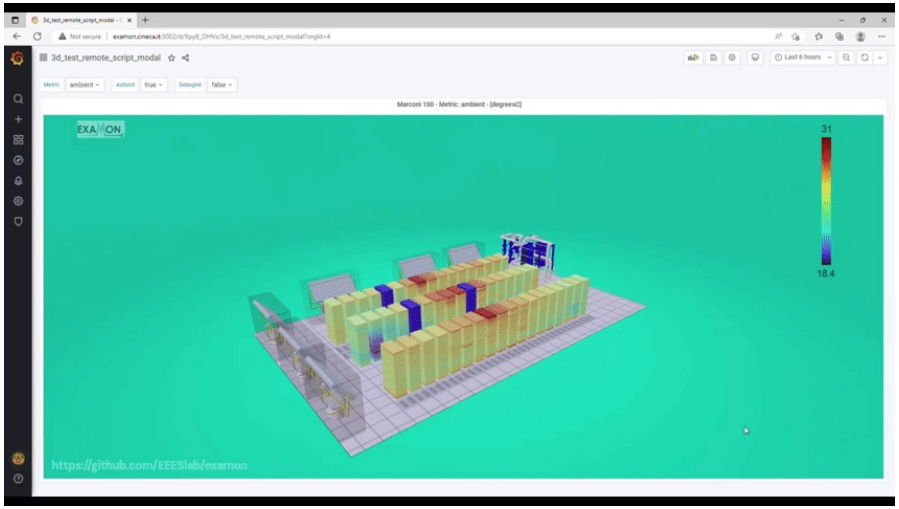

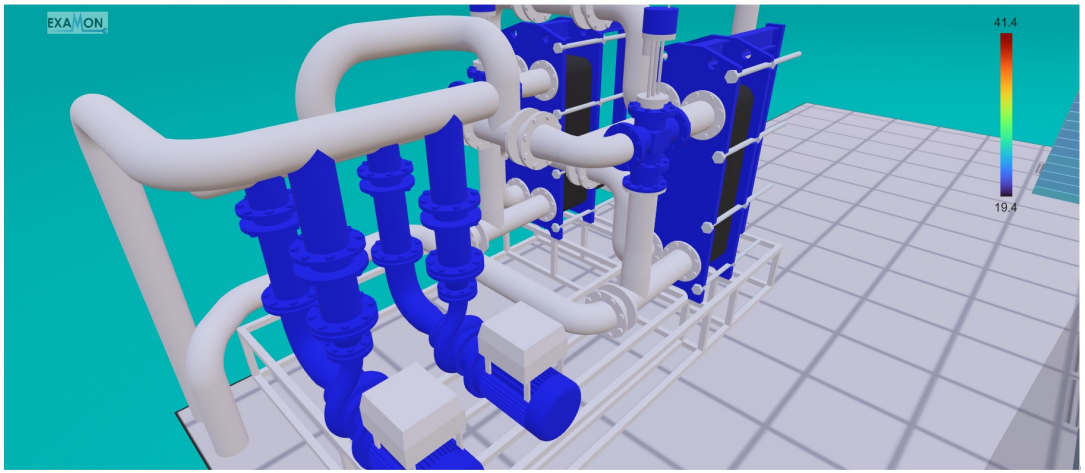

Luca Benini, Marco Aldinucci

Bologna, 25/26 novembre 2022

Milestones- SWI Lab (UNITO)

Kick off meeting

Mission: designing and implementing programming models and system software - Workflows, **Technical Focus**: run-time systems, I/O and data management, HPC-cloud/AI/Big Data convergence -

aldinuc@di.unito.it

- Spoke lab software & Integration (SWI) – 570k€ (Apr 2022)
  - Affitto locali 400mq 3 anni - 120k€ (Ott 2022)
  - Progettazione esecutiva + ristruttazione + arredamento - 350k€ (Apr 2022)
  - Attrezzatura ICT ufficio - 100k€ (Apr 2022)
- HPC4AI@UNITO
  - Attualmente 10k core (Intel, Arm, RISC-V), 120GPU (Nvidia), 4PB
  - Estenzione programmata 500k€ (Q3 2023)

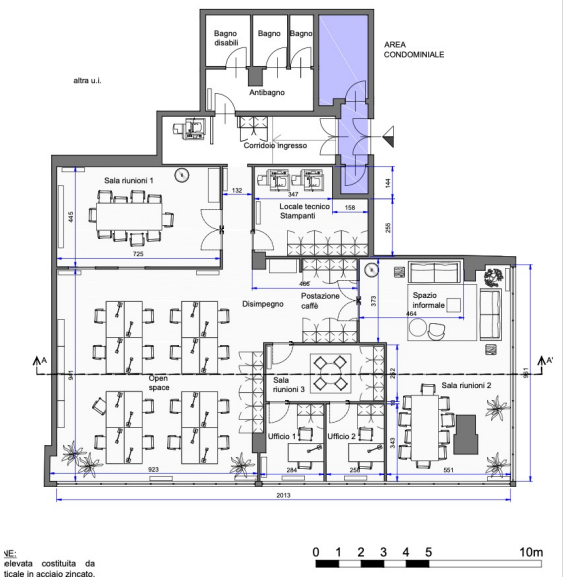

Luca Benini, Marco Aldinucci

Bologna, 25/26 novembre 2022

Federated Learning on RISC-V (UNITO+UNIBO)

	Energy/FLOP (CPU only)	Avg CPU power (idle)	TDP (per socket)	Avg system power (idle)
Intel	5 nJ	44 W	125 W	190 W
ARM	1 nJ	15 W	250 W	290 W
RISC-V	12 nJ	3.4 W	5 W	5 W

TABLE III: **Comparison of the different systems for power employed by CPU and overall systems.** The Intel system has two sockets, whereas the others have one socket.

tree(k,d) - DL edge inference schema

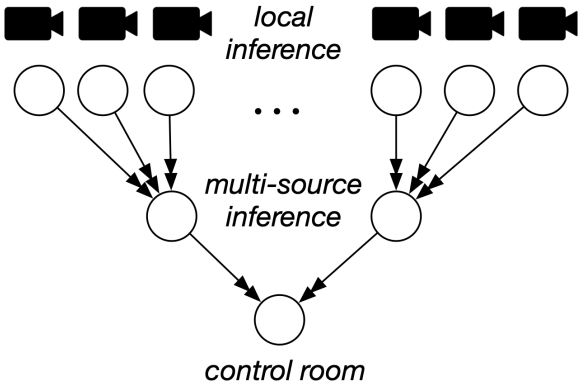

master-worker(n) - FL client-server

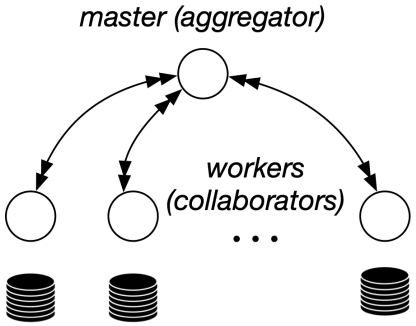

p2p(n) - FL distributed

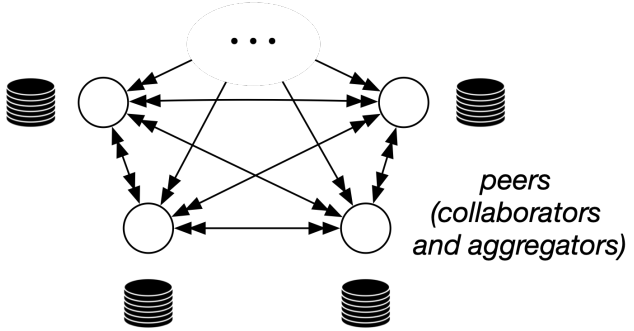

- First port of Pytorch on RISC-V
- First deployment of FL on heterogenous RISC-V+Arm+Intel
- Performance/Energy study

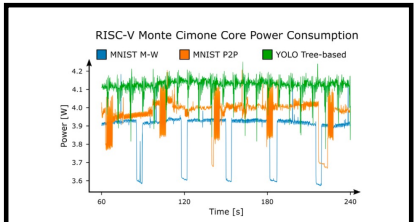

Fig. 3: **RISC-V power consumption traces** Three minutes power consumption on the Monte Cimone RISC-V system.

Experimenting with Systems for Decentralized Machine Learning. Submitted for peer-review sep 2022

Luca Benini, Marco Aldinucci

Bologna, 25/26 novembre 2022

Milestones- Flagship 1 (POLITO)

Kick off meeting

Non-functional properties - Design exploration: energy, power, reliability – **Mission**: Evaluating and exploring design solutions in terms of non-functional properties (energy, power, reliability). matteo.sonzareorda@polito.it

Expected TRL6 outcomes

- Holistic power and reliability-aware management for a pre-exascale Tier-zero HPC machine.
- Dynamic Autotuner and Compiler toolchain for precision tuning.
- Design exploration environment supporting the analysis of different architectures and solutions wrt the target power, reliability and performance metrics.

Start-up activities (on-going)

- Initial actions: Stimulate awareness and dissemination
- Proposal for a focus group within the CINI HPC-KTT lab
- dissemination to general audiences (a first article appeared on Agenda Digitale magazine)
- initial liaisons with academic and industrial partners
- follow RISC-V standardization activities (e.g. RISC-V Application Platform-TEE Task Group, AP-TEE TG)

Milestones- Flagship 2 (POLIMI)

Kick off meeting

Heterogeneous acceleration, architecture, tools, software **Mission**: Design, modeling, simulation of heterogeneous accelerators and runtime resource management, libraries, and tools for heterogeneous HPC-cloud-edge systems cristina.silvano@polimi.it

Expected TRL6 outcomes

- Energy-efficient RISC-V based accelerators based on open European technology derived from the EuroHPC EPI (European Processor Initiative).
- Energy-efficient neural network accelerator chip set based on in-memory computing.
- High-level synthesis -based workflow for reconfigurable HW accelerators including automatic generation of power monitors

Start-up activities (on-going)

- Design space exploration and mapping of low-power DNN accelerators
- Analog matrix computing with crosspoint resistive memory arrays for in-memory DNN acceleration
- High-Level Synthesis of multi-threaded accelerators for DNNN and graph analytics;
- Runtime resource management and thermal management for heterogeneous HPC-Cloud-edge systems;
- Task scheduling and time predictability on heterogeneous HPC systems with probabilistic real-time

Milestones - Flagship 3 (UNIPI)

Kick off meeting

Workflows, I/O, and HPC-cloud convergence **Misssion & Outcomes**

- Portability across systems and new processors, scalability matching end-user expectations, modularity, and composability to accommodate different programming needs and styles while enabling the reuse and integration of legacy code.
- Umbrella community contributing, developing, and *maintaining an integrated set of SW tools* implementing a unifying software stack for HPC-oriented workflows, integrating tools and methodologies for cloud-HPC and high-performance storage & IO, and providing suitable design and operation tools for digital twins.

marco.danelutto@unipi.it

Attività preliminari

- Set up comunità (inclusa propaganda per le posizioni che si stavano per aprire)
- Primi scambi di interessi e inizio raccolta contributi per il deliverable della prima milestone
- Revisione interna dei contributi dei singoli gruppi nella prospettiva Spoke1/FL3
- Contatti con industrie per definizione di possibili dimostratori

Milestones- Flagship 4 (UNINA)

Kick off meeting

Trustworthy High-Performance Computing - **Mission:** put open-source HPC technologies on par with commercial offerings (Intel SGX, TDX, AMD SME, SEV, …) acilardo@unina.it

Expected TRL6 outcomes

- RISC-V TEE targeted at HPC demonstrated in a relevant environment
- Software support and toolflows for application development jointly demonstrated with the TEE
- At least two privacy-sensitive compute-intensive applications in multi-tenancy HPC/cloud

Start-up activities (on-going)

- Identification of local reference persons
- Hiring of young researchers
- Training of young researchers
- Identification of suitable high-level fault models for reliability-oriented design exploration

Milestones - Flagship 4 (UNINA)

Kick off meeting

Codesign - application + SW + HW targeting, benchmarking, patterns, microkernels
Mission: Design, modeling and simulation of heterogeneous accelerators for HPC-cloud systems and edge servers. Sebastiano.battiato@unict.it

Planned activities

- Benchmarking of accelerators, software optimizazion, performance analysis and portability, profiling of HPC codes
- Mini-applications and benchmarking from multiple domains. Algorithmic prototyping, and algorithmic co-design: simulation, modeling, optimization;
- Codesign HW Development and exploitation of next generation of HPC systems

Start-up activities (on-going)

- Identification of local reference persons
- Hiring of young researchers
- Training of young researchers
- Identification of suitable case-studies, tools and benchmarking

Stato delle Attività (Industry PoCs)

Luca Benini, Marco Aldinucci

Bologna, 25/26 novembre 2022

Rationale

Existing relations of collaboration and trust

- Design Methodologies
- Technologies
- Ecosystem of
  - HW IPs
  - SW tools

- Tech Needs
- Use cases
- Areas of interest
- PoCs

Technological transfers can only happen via trust relations among people and bodies.

The often happens locally: in a city, in a region, in a scientific body, …

Luca Benini, Marco Aldinucci

Bologna, 25/26 novembre 2022

Innovation grant & OpenCalls (regolamento attualmente in definizione)

- Innovation 1,8M€ (0,726M€ sud)
  - PoC di interesse industriale, possibili temi cross-spoke
  - Ogni PoC si implementa in alcune sedi dello spoke, non tutte
    - Personale TD assunto da università e/o ente
    - Possibile coinvolgere PMI
  - Gennaio 2023; Tornata 2: novembre 2023
- Opencalls 3,2M€ (1,6M€ sud)
  - Categorie: Experiments, Use cases, Deployment of demonstrators, Proof of Concepts, Pilot applications
  - Su temi di interesse industriale e accademico
  - Taglio: 50-100(-200) k€ - cofin imprese e accademia TBD
  - Finanziate (B/C) Mag 2023, Dic 2023, Mag 2024
  - Finanziate con compute-hours da spoke0 (A): asap e con tornate periodiche
  - 1 grant per tornata, 2 grant nelle 3 tornate previste
- Creare una filiera HPC: industrie, PMI, accademia/enti

Luca Benini, Marco Aldinucci

Bologna, 25/26 novembre 2022

PoC pianificati

Kick off meeting

- Fincantieri
  - High-performance Digital Twin based on Cloud-Edge Continuum and access to Cloud-based HPC running large-scale models and/or process-based models
- ENI
  - HPC Next-Gen Full Waveform Inversion
  - Job duration prediction
  - Supervised learning beyond minibatches
- Intesa
  - Enterprise Multi-tenancy Distributed learning (FL3 + FL4 + FL5)
  - Enterprise AI and BigData frameworks for HPC ecosystem (FL1 + FL3 + FL4 + FL5)
- Leonardo
  - RISC-V co-design, benchmarking, programming modes, targeting HPC applications at the EDGE
  - Convergence of HPC and Cloud using containers. Cloud Continuum with workflows enabling applications HPC in the loop
  - Green Computing (both HW and SW), develop containers/applications with embedded energy monitoring capabilities + Cloud security (zero trust for container technologies)
- Unipol
  - Cloud-HPC convergence: federated and distributed learning
- Engineering
  - Automating HPC benchmarking
- ThalesAlenia Spazio
  - High-performance RISC-V processor for satellite
  - HPC technologies for SAR ground processing of Earth Observation data
  - Cyber-securityon-board spacecraft,including embedded AI framework and radiation tolerant AI processing in space
- Autostrade & Sogei (TBD)

Luca Benini, Marco Aldinucci

Bologna, 25/26 novembre 2022

Esempio-1: ENI

Kick off meeting

Next-Gen Full Waveform Inversion

Full-waveform inversion (FWI) is a challenging data-fitting procedure based on full-wavefield modeling to extract quantitative information from seismograms. Although some implementations have been developed already, new frameworks and software technologies now available make it possible to build a POC providing new features.

Cross-platform solver

Frameworks such as Devito Codes (https://www.devitocodes.com/) let the developer describe a finite-differences based algorithm using a high-level language. This description can be compiled just-in-time on different platforms, providing unprecedented portability. The POC is expected to prove this portability effective.

Workflow manager

A workflow manager like StreamFlow (https://streamflow.di.unito.it/) allows a distributed, platform-independent implementation of the FWI algorithm. The workflow manager shall provide system resilience to faults (i.e. job failures) and multi-platform support.

Expected POC duration: >1 year

Human resources: 1 FTE PHD student + 0.2 FTE supervisor

Hardware requirements: Development environment: 5+ computing nodes with Nvidia GPU (required by devito), Stress-test environment: 100+ computing nodes

Luca Benini, Marco Aldinucci

Esempio-2: TASI

High-performance RISC-V processor for satellite.

Study and development of high-performance computing architecture based on RISC-V processor suitable for on-board computing. The processor will be designed to operate at high clock speeds and to support high-performance computing applications. We will also investigate the feasibility of using this processor in a harsh space environment.

Expected POC duration: 3 years

Human resources: 2 FTE RTDA/equivalent PhD

Hardware requirements: Cluster RISC-V