Introduction to codesign

Matteo Turisini (CINECA)
m.turisini@cineca.it

Barcelona, 2024 Oct 10th,

EUPEX

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101033975. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Germany, Italy, Greece, United Kingdom, Czech Republic, Croatia.

Introduction

- Trends in HPC
  - Hardware specialization using chiplet technology (shifting from x86 to ARM)
  - Applications are becoming workflows (heterogeneity at system level)
  - European R&D project (gather and integrate results of a decade of projects)

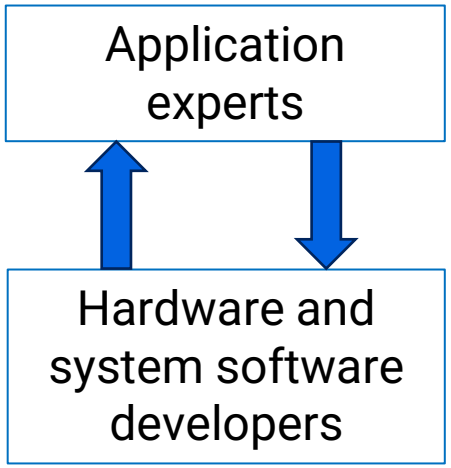

⑩Codesign: bidirectional and iterative interaction process

- **Applications** employed to identify impact of **design decisions** onto **performances**
- Idea born in the 80s
- Direct: apps developed in tandem (rare, ultra-specialized)
- Indirect: software tested on early-release hardware (pragmatic and wider)

Success = how well a developed technology fullfill the needs of end users

Methodology and exploitable products

- Application selection
  - Selection criteria: exascale readiness, workflow complexity, ability to stress subsystems, use of accelerator, presence in CoE and size of user community, relevance for societal impact (industial and academic)
  - Initial profiling on reference platform
- Code porting/optimization
  - Adaptation or improvements investigated
  - Focus on hardware features SVE and HBM
  - Focus on system level (MSA)

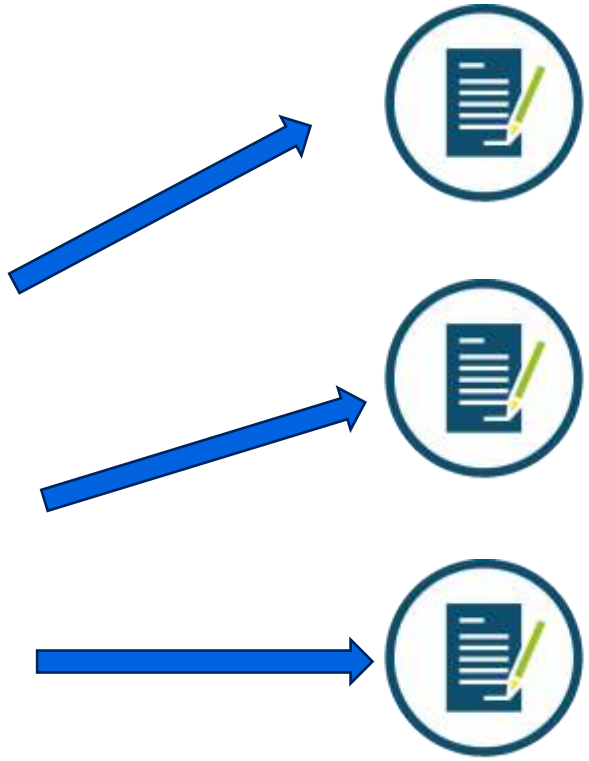

D3.1 (2023 January)

D3.2 (2024 January)

D3.3 (2025 July)

Towards european exascale performance Pilot system!

Codesign output for the EUPEX platform

- Outputs Based on a survey from application leaders
- Other parameters fixed by Rhea design specs e.g. number of cores, cache/HBM size
- Fat tree topology (cost effective at Pilot scale, <128 nodes)
- GPU Technology

CPU Node	Codesign outputs
DRAM	512GB (4x128GB DDR5 DIMM)
Number of sockets	2
Network bandwidth per node	200 Gbit/s
Number of NICs per node	2 (1 NIC per socket)
Local disk	Not required for the applications

GPU Node	Codesign outputs
Number of GPUs	4
DDR Memory	512GB (4x128GB DDR5 DIMM)
Number of sockets	2
Network bandwidth per node	400 Gbit/s
Number of NICs per node	4
Local disk	No required for the applications

Software Development Vehicles (SDVs)

Used for porting (on ARM), optimization activities (SVE, HBM) and to study software environment

SDV1

A64FX + Nvidia H100
IRENE cluster at CEA

A64FX first CPU with SVE support

IRENE will be part of the EUPEX Early Access Program (EAP)
See talk by Mario Kovac at 17:00

SDV2

Nvidia Grace and GH200 nodes @E4

Grace CPU: more recent than A64FX
Based on Neoverse V2 (with SVE and HBM)
Closer to the target, halfway between Rhea1 (Neoverse V1) and Rhea2 (Neoverse V3)

Hopper GPU: same GPU of the initial SDV (H100), smallest porting effort

Academic, industrial, societal compute domains

Explored by EUPEX for co-design and benchmarking

Climatology, meteorology

- ECMWF, CybeleTech, Atos

Engineering

- IT4I, CINECA

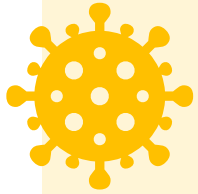

Biology and health

- CINECA, CINI, CEA

Astrophysics

- FORTH, INAF, CEA

Seismology

- INAF, CINECA, CINI, GENCI

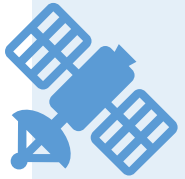

Remote sensing analysis

- FZJ

The benchmark suite

A set of key applications in the European landscape relevant in different scientific fields

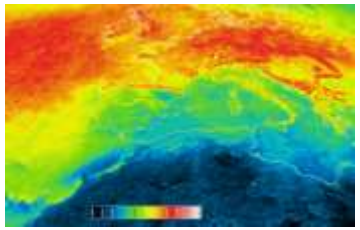

Weather Forecasting

IFS

ECMWF

AI4EO

PyTorch DDP, HPDBSCAN

Forschungszentrum Jülich

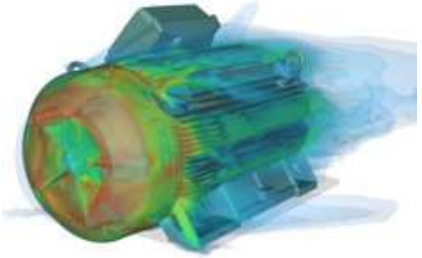

Engineering

ESPRESO FEM

IT4I

Biology and Health

LIGEN

CINECA

Mini Applications

Dyablo, HEVC

CEA, Atos, UNIZG

Precision Agricolture Forecast

Cybeletech

Astrophysics

Gadget

INAF, FORTH

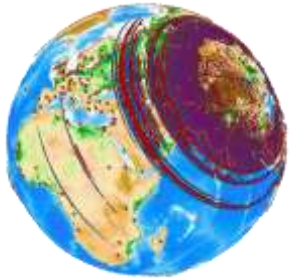

Seismology

SPECFEM3D

HPC4NDR

Material Science for Biology

BigDFT

CEA

Synthetic Benchmarks for EUPEX platform

Atos

EUPEX

Today's presentations

A set of key applications in the European landscape relevant in different scientific fields

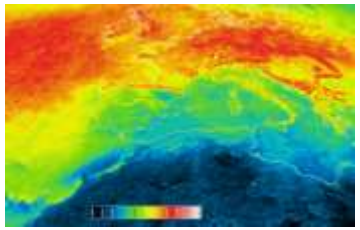

Weather Forecasting

IFS

ECMWF

See talk by Andrew Beggs at 15:00

"Manual and Automated vectorization Techniques for Integrated Forecasting system"

See talk by Filippo Barbari at 16:00

"High throughput drug discovery on the Fujitsu A64FX"

Biology and Health

LIGEN

CINECA

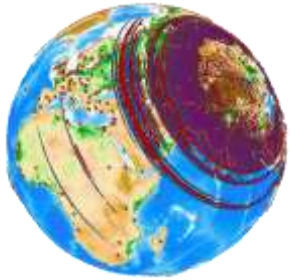

Seismology

SPECFEM3D

HPC4NDR

See talk by Piero Lanucara at 16:30

"EuPEX-ChEESE-2P cooperation: the SPECFEM3D example; context, preliminary results and next steps"

Incoming webinars

Dedicated dissemination activities to share lesson learned and good practices

AI4EO

PyTorch DDP, HPDBSCAN

Forschungszentrum Jülich

2024 November 6th

"Advances in Geospatial Foundation Models for Earth Observation: Scaling Machine Learning with Supercomputing on Large Remote Sensing Datasets"

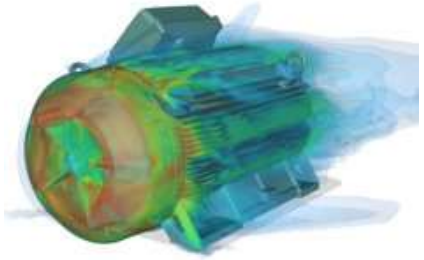

Engineering

ESPRESO FEM

IT4I

2025 March (to be determined)

"Vectorization of FEM kernels"

Mini Applications

HEVC

UNIZG

2024 December 10th

"Bolt65-AI fusion: A case study in object detection"

System software webinars also available from 2025 January to May
e.g. StreamFlow (CINI), Ocean (CEA), Parastation Modulo (Partec)

A cloud of CoEs and beyond

2023-2027 Centers of Excellence (CoEs) and new R&I activities

ChEESE https://cheese-coe.eu/

esiwace https://www.esiwace.eu/

PLASMA PEPSC https://plasma-pepsc.eu/

bioexcel https://bioexcel.eu/

EXCELLERAT P2 https://www.excellerat.eu/

HiDALGO https://hidalgo-project.eu/

MAX DRIVING THE EXASCALE TRANSITION https://www.max-centre.eu/

MultiXscale https://www.multixscale.eu/

SPACE https://www.space-coe.eu/

EUPEX

Thank you and good continuation!

Matteo Turisini (CINECA)
m.turisini@cineca.it

Barcelona, 2024 Oct 10th,

EUPEX

EuroHPC Joint Undertaking

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101033975. The JU receives support from the European Union's Horizon 2020 research and innovation programme and France, Germany, Italy, Greece, United Kingdom, Czech Republic, Croatia.