EPI Forum

Paris, 6-7 October, 2025

European Processor Initiative

WHEN PERFORMANCE MATTERS

ARE SCIENTIFIC APPLICATIONS READY TO HARNESS THE POTENTIAL OF RHEA AND EPAC?

Elisabetta Boella, *HPC Product Specialist*

EPI Forum
October 6th, 2025

E4 COMPUTER ENGINEERING

E4 Computer Engineering designs and manufactures highly technological solutions for **HPC** Clusters, **Cloud**, **Data Analytics**, **Artificial Intelligence**, **Hyper-Converged infrastructure** and **Quantum Computing** for the Academic and Industrial markets. We have been collaborating for years with the main research centers at national and international level (CINECA, CERN, ECMWF, LEONARDO) and we are involved in national and **European projects** in the HPC, Quantum Computing, and AI fields.

Each E4 solution is **UNIQUE**, like every one of our customers; **TESTED** in every single component; **VALIDATED** to verify the actual performance of each system and **SERVICED** by technicians who provide assistance in the most extensive and complex Italian and European computing infrastructures.

E4 ANALYTICS
LET YOUR DATA PAY YOUR GROWTH

Through the sister company E4 Analytics, E4 works to integrate **Artificial Intelligence** and **Data Science** in organizations that undertake the **Digital Transformation** of their business to improve products/processes and optimize resources. We operate at the intersection between business and technology, supporting the customer in the adoption of customized and secure **GenAI solutions**: with E4 Analytics, company data become a **resource for creating value**, enhancing **innovation** and **competitiveness** in the marketplace.

E4 COMPUTER ENGINEERING

NATIONAL AND EU PROJECTS – WHERE WE CONTRIBUTE

2013
PEDRAFORCA CLUSTER

Barcelona Supercomputing Center – Centro Nacional de Supercomputación

2016 - 2017
D.A.V.I.D.E.

PRACE, CINECA

299° in TOP500 (June 2017)
14° in GREEN500 (June 2017)

2018
EPI SGA1

European Processor Initiative

2019
EUROHPC JU CALLS

2020
textarossa, ADMIRE, ATTRACT, exaFOAM, LIGATE, MAELSTROM

2021
EPI SGA2

2022
European Processor Initiative, DYNAMOS, EUPEX

2023
EUMaster4HPC, TRISTAN, ISOLDE, MAX – DRIVING THE EXASCALE TRANSITION, DECICE, SPACE, EDOCC, EXCELLERAT P2

2024
SmartHPC.QC, EoCoE, FUTURE, SEED, MoseGad, RISC-V4ASTRO

NEW SUBMISSIONS UNDERWAY

ARM V8, ARM V9, HBM, AND RISC-V: PREPARING APPLICATIONS FOR RHEA AND EPAC

AMPERE ALTRA MAX
128 cores/processor
Arm Neoverse N1
3.0 GHz
64 MB L3 cache/socket
1 TB DDR4-3200 RAM

TDP 225 W/processor

NVIDIA GRACE
144 cores/Superchip
Arm Neoverse V2
3.1 GHz
114 MB L3 cache/socket
450k MB LPDDR5 RAM

TDP 500 W/Superchip

INTEL SPR
48 cores/processor
x86_64
3.5 GHz
105 MB L3 cache/socket
256 GB DDR5-4800 RAM
128 GB HBM
TDP 350 W/processor

SOPHON SG2042
64 cores/processor
Risc-V
2.0 GHz
64 MB L3 cache
256 GB DDR

TDP 120 W

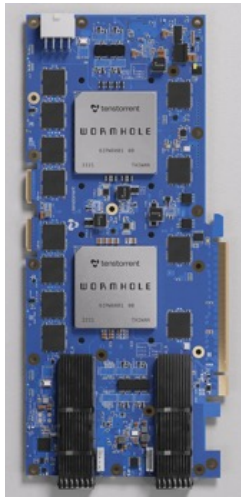

TENSTOREENT n300
accelerator
Risc-V

TDP 300 W/card

SIESTA:
A CODE FOR AB INITIO MOLECULAR DYNAMICS

Based on density-functional theory (DFT)

Written in Fortran

Parallelised with MPI and OpenMP

Uses BLAS, LaPACK, ScaLAPACK, and ELSI

Simulation: quantum dot

Compiler: GNU suite various versions

Optimisation: -O3

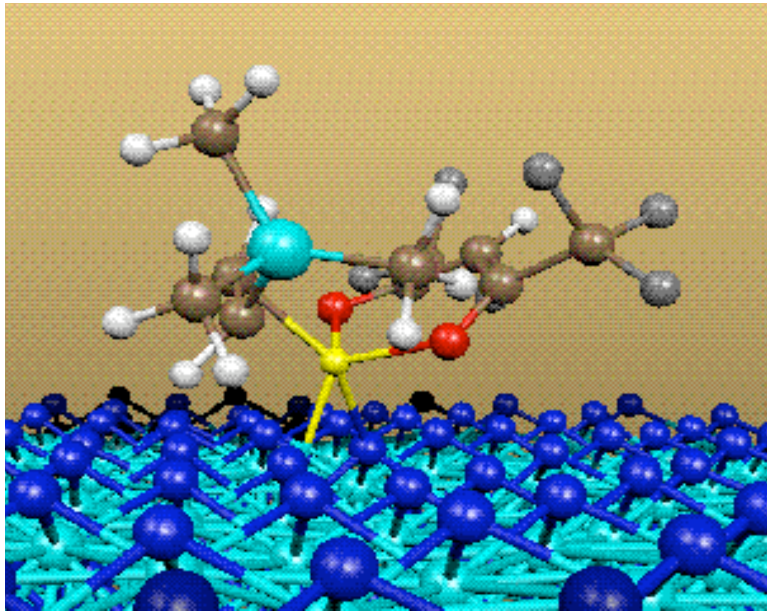

P. Ordejon, SIESTA (nano) TUTORIAL (2010).

IMPROVED PARALLEL EFFICIENCY ON GRACE WITH REDUCED ENERGY CONSUMPTION

L.Riha et al. (2025), In Proceedings of the 22nd ACM International Conference on Computing Frontiers. ACM, New York, NY, USA, 150–156.

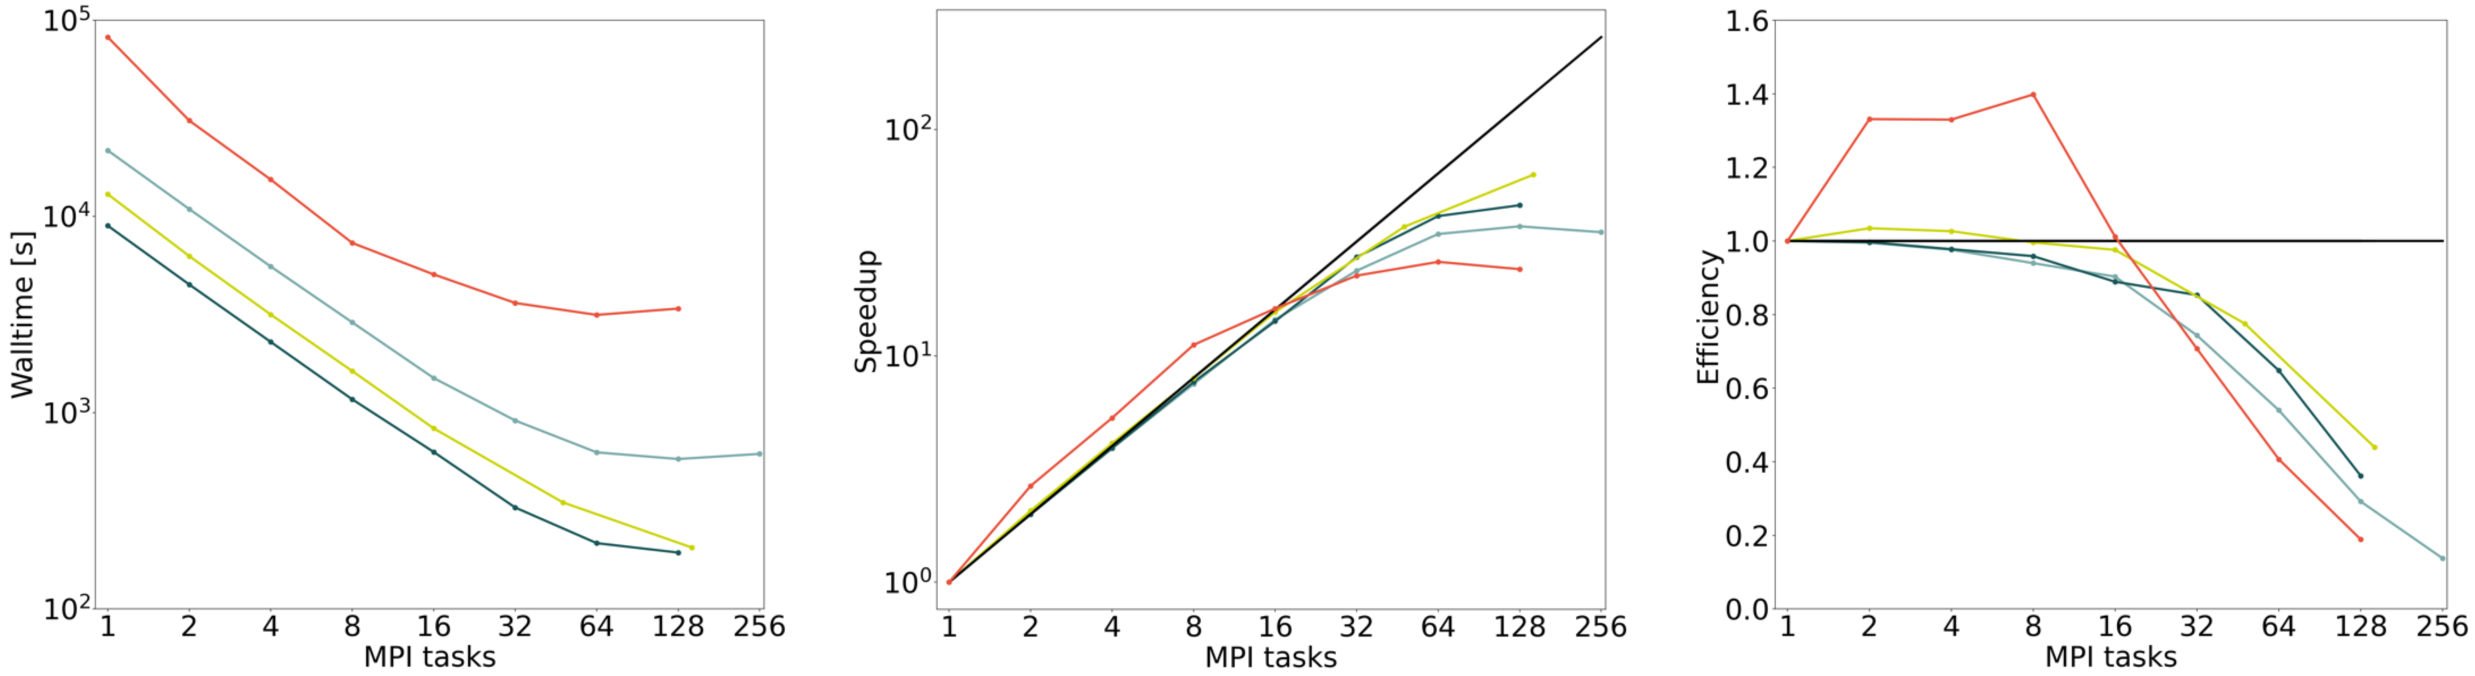

-- SOPHON SG2042, -- AMD GENOA, -- AMPERE ALTRA MAX, -- NVIDIA GRACE

Energy-to-Solution full node estimate: AMD GENOA: 139.5 kJ, NVIDIA GRACE: 102.7 kJ

HBM MEMORY DELIVERS UP TO 1.5× PERFORMANCE GAINS

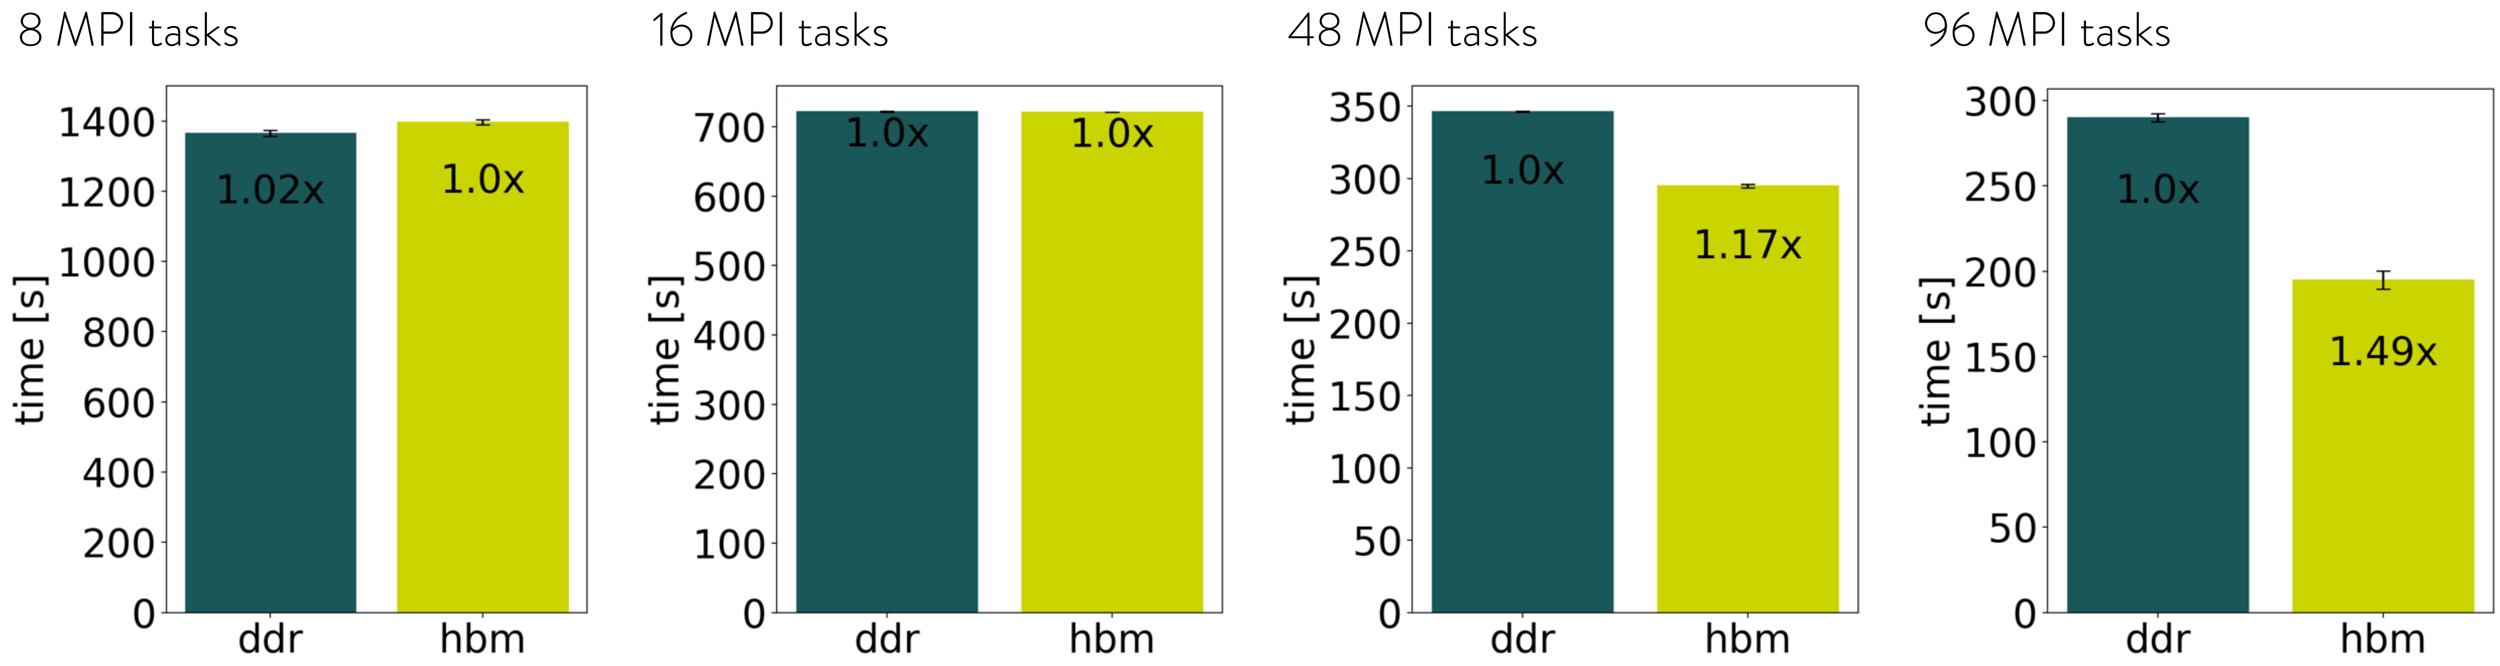

ACCELERATING N-BODY SIMULATIONS ON RISC-V: TWICE THE SPEED, HALF THE ENERGY VS X86

J.L. Almerol et al. (2025), arXiv pre-print, arXiv:2509.19294

Time-to-solution (TTS) over ~50 simulations

Energy-to-solution (ETS) over ~50 simulations

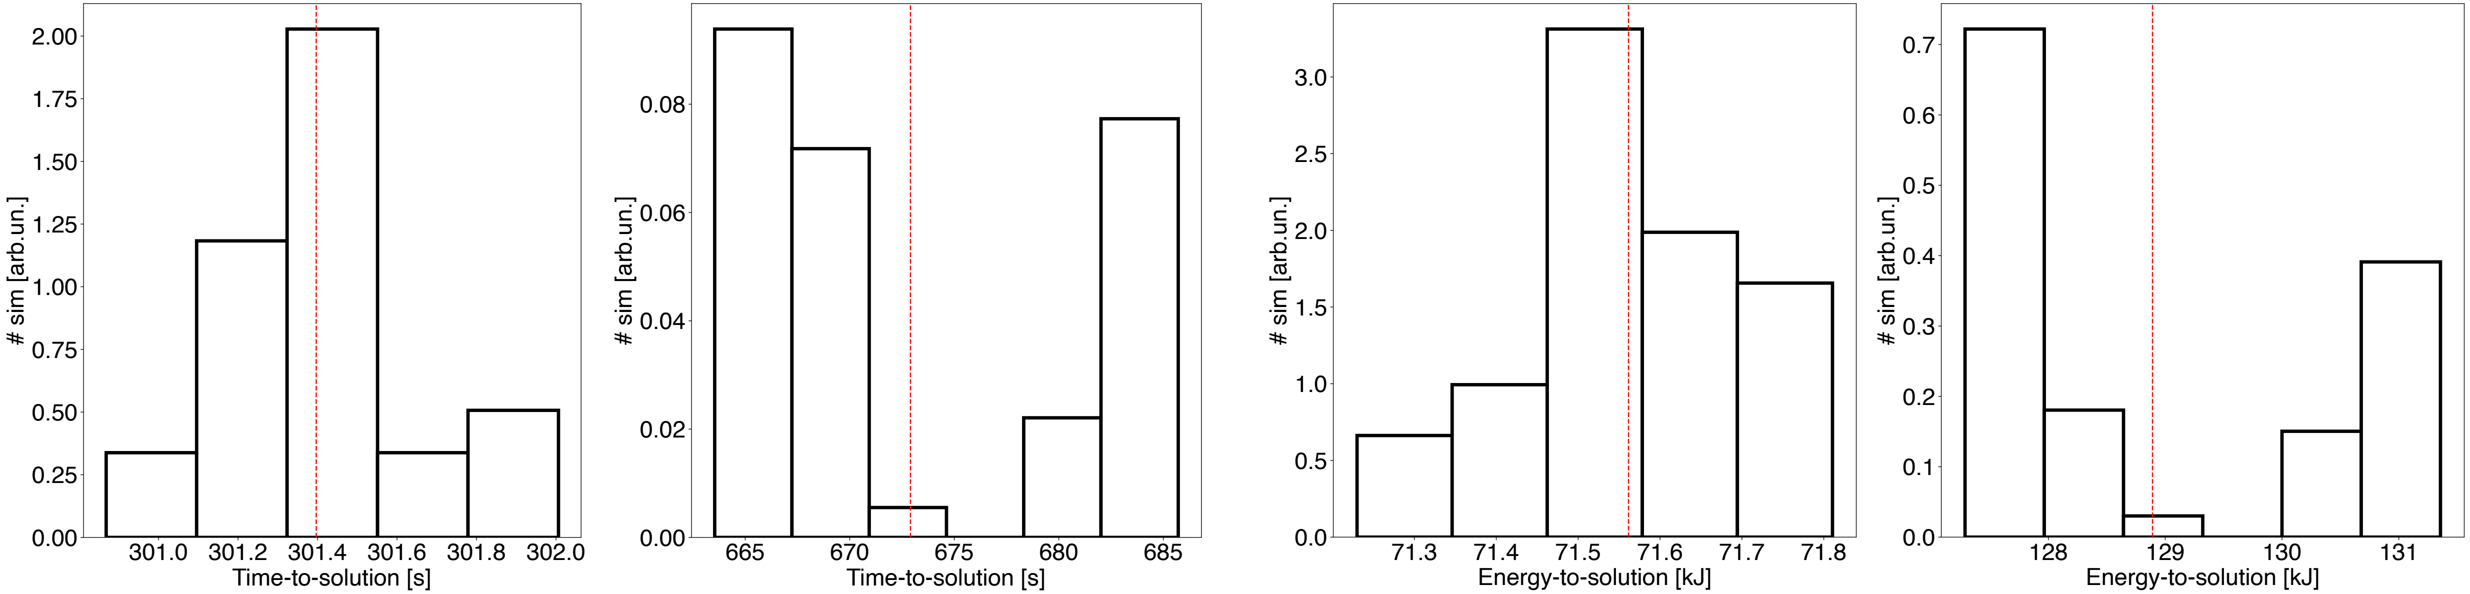

1 core +
1 Tenstorrent n300

32 cores +
AVX-512 intrinsics

1 core +
1 Tenstorrent n300

32 cores +
AVX-512 intrinsics

TTS_{avg} = 301.40 ± 0.24 s

TTS_{avg} = 672.90 ± 7.83 s

ETS_{avg} = 71.56 ± 0.13 kJ

ETS_{avg} = 128.89 ± 1.52 kJ

SUMMARY & PERSPECTIVES

- Scientific applications have been evaluated on multiple architectures with features comparable to Rhea.
- HBM yields ~1.5× speedup with proper compilation flags
- SVE boosts ARM performance, even simply via autovectorisation
- RISC-V CPUs still lag behind in raw performance, but reliable compilation and correct results provide a strong basis for rapid progress.
- RISC-V accelerators achieve impressive results, with 2× performance and energy efficiency compared to x86.

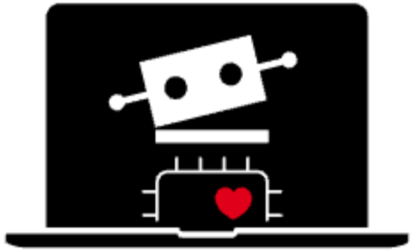

WHEN PERFORMANCE MATTERS

EXTRA SLIDES

Elisabetta Boella, *HPC Product Specialist*

EPI Forum
October 6th, 2025

SUCCESS STORIES

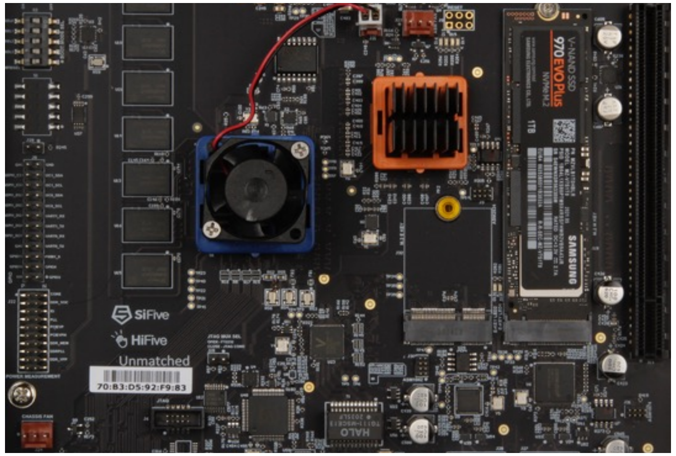

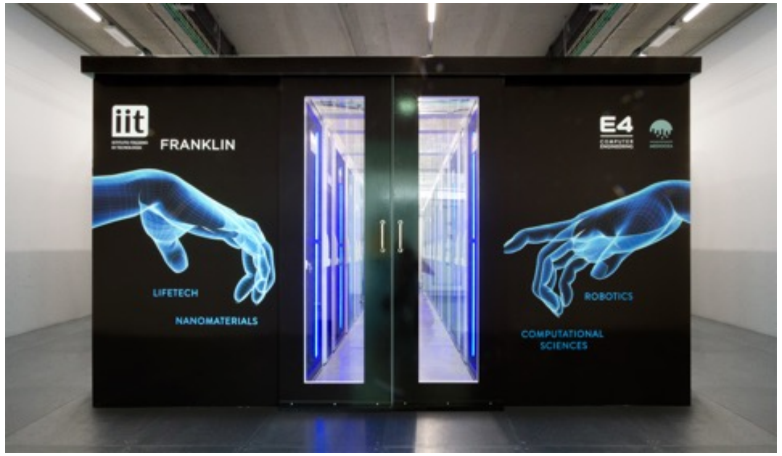

MONTE CIMONE
First HPC cluster
based on RISC-V

IIT FRANKLIN
HW and SW

MEDOOZA ICE4HPC

INAF ASTRI-MA
Non-supervised real time
data analysis

HETEROGENEITY IS KEY FOR MODERN HPC

AMPERE ALTRA MAX

128 cores/processor
Arm Neoverse N1
3.0 GHz
64 KB L1 cache
1 MB L2 cache
4 MB L3 cache
225 W

NVIDIA GRACE

144 cores/processor
Arm Neoverse V2
3.1 GHz
64 KB L1 cache
1 MB L2 cache
228 MB L3 cache
500 W

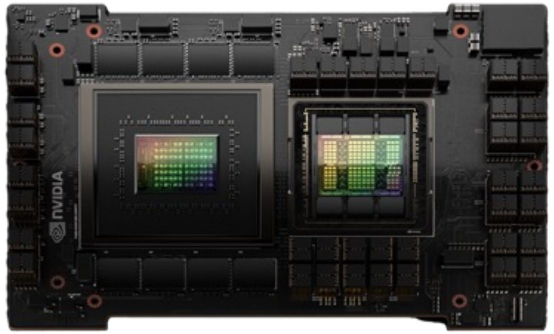

NVIDIA GRACE HOPPER

72 cores/processor
117 MB L3 cache
1 GPU H100
450 - 1000 W

SOPHON SG 2042

64 cores/processor
Risc-V
2.0 GHz
64 KB L1 cache
1 MB L2 cache
64 MB L3 cache
120 W

TESTED PLATFORMS @E4 DATA CENTRE

Architecture	CPU Model	Frequency	Cores/node	Memory/node	L3 cache
x86_64	Intel(R) Xeon(R) Gold 6226R (Cascade Lake)	2.9 GHz	32	192 GB	22 MB
x86_64	AMD EPYC 7313 16-Core Processor (Milan)	3.3 GHz	32	256 GB	**128 MB**
aarch64	Ampere Altra Max ARM Neoverse N1	3.0 GHz	256	512 GB	4 MB
aarch64	NVIDIA Grace Hopper Arm Neoverse-V2	3.1 GHz	72	480 GB	117 MB

gcc 8.5.0 + OpenMPI 4.1.4 - OpenFOAM v2212 or OpenFOAM v2306

MICROBENCHMARKS SELECTED

Microbenchmarks	Top-Level Solver	Mesh generation - Cell count - Cell type
MB1 Cavity 3D	icoFoam	blockMesh - 8M - Hexahedra
MB2 Compressible starting square jet	rhoPimpleFoam	blockMesh - 2M - Hexahedra
MB4 DLR-JHC burner	reactingFoam	blockMesh - 400k - Hexahedra
MB5 ERCOFTAC Conical diffuser	simpleFoam	blockMesh - 3M - Hexahedra
MB6 Two cylinders in line	adjointOptimisationFoam	blockMesh - 24500 - Hexahedra
MB8 Rotating Wheel	pimpleFoam	snappyHexMesh - 20M - Polyhedra
MB9 High-lift airfoil	rhoPimpleFoam	snappyHexMesh - 19796480 - Polyhedra
MB11 Pitz&Daily Combustor	XiFoam	blockMesh - 200k - Hexahedra
MB12 Model Wind Farm	pimpleFoam	blockMesh - 8M - Hexahedra
MB17 1D Aeroacoustic Wave Train	rhoPimpleFoam	1D blockMesh - 0.05M - Hexahedra
MB19 Viscoelastic polymer melt flow	viscoelasticFluidFoam	cfMesh - 1M -Polyhedra

OPENFOAM: AN OPEN SOURCE CFD APPLICATION

Based on the cell centered finite volume method

Developed by OpenCFD Ltd @ESI Group

Written in c++

Parallelised with MPI

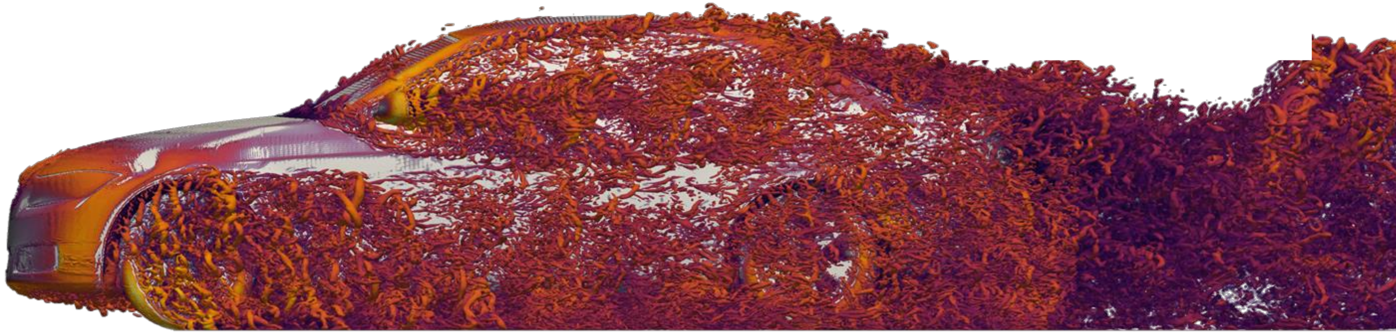

upstreamCFD©

Simulation: open-closed cooling drivaer

Compiler: GNU suite various versions

Optimisation: -O3

HBM ACCELERATES COMPUTATION BUT FACES MEMORY BOTTLENECKS ON LARGE GRIDS

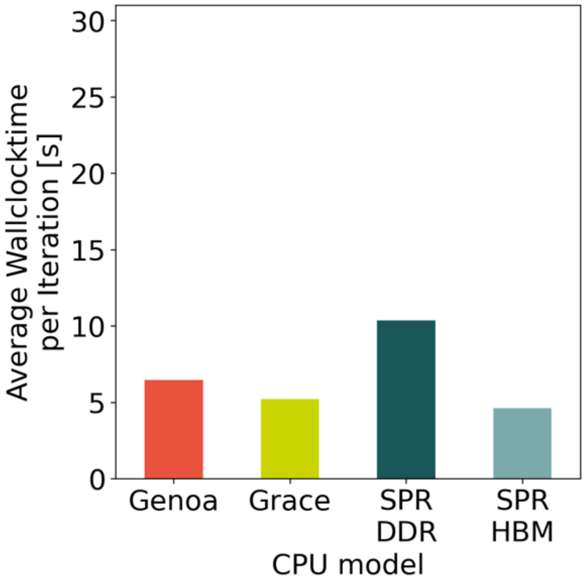

Architecture	Cores	TTS [s]	Energy [MJ]
AMD Genoa	64	26435.4	19
NVIDIA Grace	144	21187.7	15.3
INTEL SPR	96	20957.1	13.7

ECSIM: A MASSIVELY PARALLEL PLASMA PHYSICS CODE

Kinetic plasma physics code

Based on the Particle-In-Cell method

Written in c/c++

Parallelised with MPI

Includes OpenACC directives

Uses PETSc to solve fields

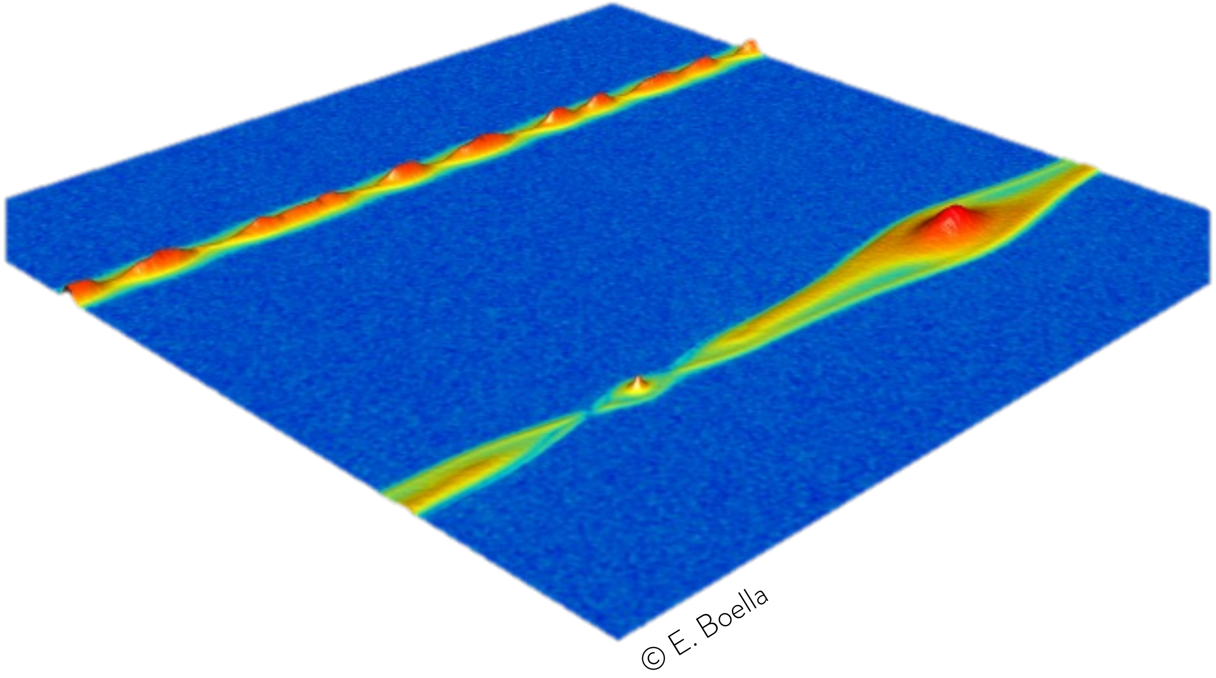

KERNELS WITH DATA MOVEMENTS BENEFIT FROM GH

Current Filamentation Instability

NVHPC 23.5 or NVHPC 23.11, OpenMPI 3.1.5

Moment Gathering

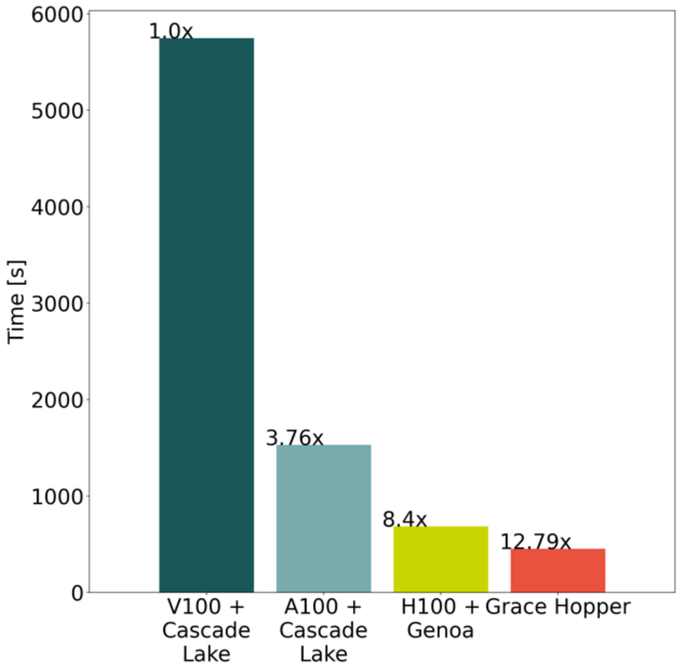

Particle Mover

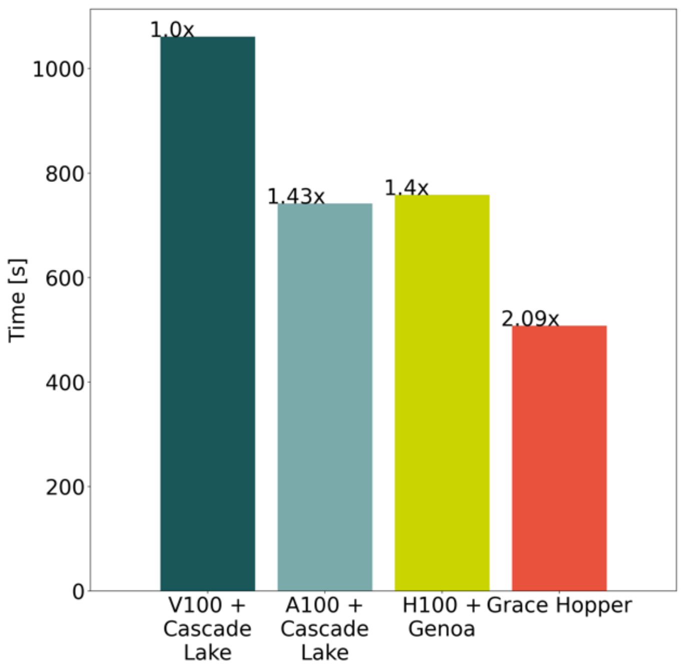

Cascade Lake + V100 @CINECA - Cascade Lake + A100, AMD Genoa + H100 and GH200 @E4

Tests performed using 1 MPI task + 1 GPU