MaX - Materials design at the exascale - a EuroHPC Centre of Excellence: Recent selected results

Invited Paper

Lubomir Riha
IT4Innovations, VSB-TU Ostrava
Ostrava, Czech Republic
lubomir.riha@vsb.cz

Ada Böhm
IT4Innovations, VSB-TU Ostrava
Ostrava, Czech Republic
ada.bohm@vsb.cz

Ondrej Vysocky
IT4Innovations, VSB-TU Ostrava
Ostrava, Czech Republic
ondrej.vysocky@vsb.cz

Filip Vaverka
IT4Innovations, VSB-TU Ostrava
Ostrava, Czech Republic
filip.vaverka@vsb.cz

Elisabetta Boella
E4 Computer Engineering SpA
Scandiano, Italy
elisabetta.boella@e4company.com

Daniele Gregori
E4 Computer Engineering SpA
Scandiano, Italy
daniele.gregori@e4company.com

Mattia Paladino
E4 Computer Engineering SpA
Scandiano, Italy
mattia.paladino@e4company.com

Pietro Delugas
SISSA
Trieste, Italy
pdelugas@sissa.it

Oscar Baseggio
SISSA
Trieste, Italy
obaseggi@sissa.it

Stefano Baroni
SISSA
Trieste, Italy
baroni@sissa.it

Francesco Andreucci
SISSA
Trieste, Italy
fandreuc@sissa.it

Kaveh Haghighi Mood
Forschungszentrum Jülich
Jülich, Germany
k.haghighi.mood@fz-juelich.de

Andreas Herten
Forschungszentrum Jülich
Jülich, Germany
a.herten@fz-juelich.de

Gregor Michalicek
Forschungszentrum Jülich
Jülich, Germany
g.michalicek@fz-juelich.de

Daniel Wortmann
Forschungszentrum Jülich
Jülich, Germany
d.wortmann@fz-juelich.de

Stefan Blügel
Forschungszentrum Jülich
Jülich, Germany
s.bluegel@fz-juelich.de

Pablo Ordejon
Catalan Institute of Nanoscience and Nanotechnology - ICN2 (CSIC and BIST)
Barcelona, Spain
pablo.ordejon@icn2.cat

Federico N. Pedron
Catalan Institute of Nanoscience and Nanotechnology - ICN2 (CSIC and BIST)
Barcelona, Spain
federiconicolas.pedron@icn2.cat

Jose Julio Gutierrez Moreno
Barcelona Supercomputing Center
Barcelona, Spain
julio.gutierrez@bsc.es

Rogeli Grima Torres
Barcelona Supercomputing Center
Barcelona, Spain
rogeli.grima@bsc.es

Fabio Affinito
CINECA Consorzio Interuniversitario
Casalecchio di Reno, Italy
f.affinito@cineca.it

Mariella Ippolito
CINECA Consorzio Interuniversitario
Rome, Italy
m.ippolito@cineca.it

Paolo Giannozzi
Università di Udine
Undine, Italy
paolo.giannozzi@uniud.it

Luigi Genovese
CEA / IRIG / MEM / L_Sim
Grenoble, France
luigi.genovese@cea.fr

Fabrizio Ferrari Ruffino
CNR-IOM
Trieste, Italy
ferrariruffino@iom.cnr.it

Ivan Carnimeo
CNR-IOM
Trieste, Italy
carnimeo@iom.cnr.it

Marc Sergent
CEPP, Eviden
Echirolles, France
marc.sergent@eviden.com

Erwan Raffin
CEPP, Eviden
Rennes, France
erwan.raffin@eviden.com

Ivan Girotto
ICTP
Trieste, Italy
igirotto@ictp.it

Alberto Garcia
Institut de Ciencia de Materials de Barcelona (ICMAB-CSIC)
Bellaterra, Spain
albertog@icmab.es

Jan Jona Javoršek
Jozef Stefan Institute
Ljubljana, Slovenia
jan.javorsek@ijs.si

Barbara Krašovec
Jozef Stefan Institute
Ljubljana, Slovenia
barbara.krasovec@ijs.si

Alja Prah
Jozef Stefan Institute
Ljubljana, Slovenia
alja.prah@ijs.si

Sebastien Strban
Jozef Stefan Institute
Ljubljana, Slovenia
sebastien.strban@ijs.si

Augustin Degomme
SiPearl
Grenoble, France
augustin.degomme@sipearl.com

Cristiano Malica
Bremen University
Bremen, Germany
cmalica@uni-bremen.de

Lorenzo Bastonero
Bremen University
Bremen, Germany
lbastone@uni-bremen.de

Nicola Marzari
Bremen University
Bremen, Germany
marzari@uni-bremen.de

Nicola Spallanzani
S3 Center, Istituto Nanoscienze, CNR
Modena, Italy
nicola.spallanzani@nano.cnr.it

Daniele Varsano
S3 Centre, Istituto Nanoscience
Modena, Italy
daniele.varsano@nano.cnr.it

Andrea Ferretti
Centro S3, CNR-Istituto Nanoscienze
Modena, Italy
andrea.ferretti@cnr.it

Elisa Molinari
FIM Dept
University of Modena and Reggio Emilia
Modena, Italy
CNR NANO
Modena, Italy
elisa.molinari@unimore.it

Abstract

The third phase of the MaX Centre of Excellence is designed to provide the European materials simulation community with the necessary tools to adapt to the technological advancements in massively parallel heterogeneous computing systems, enabling them to tackle previously unattainable scientific challenges.

This paper outlines the primary goals of MaX-3: (*i*) Enhancing and developing lighthouse applications; (*ii*) Implementing exascale workflows, including the creation of workflow tools and foundational technologies; (*iii*) Overcoming technical challenges arising during benchmarking, deployment on EuroHPC systems, and CI/CD; (*iv*) Exploiting co-design technology and improving energy efficiency, including evaluating codes on new processing architectures related to European hardware development projects (EPI, EUPEX, etc) and assessing potential energy savings on EuroHPC production systems and new architectures.

CF Companion '25, Cagliari, Italy

ACM ISBN 979-8-4007-1393-4/25/05
https://doi.org/10.1145/3706594.3727577

CCS Concepts

• **Computing methodologies** → **Parallel computing methodologies**; **Modeling and simulation**; • **Computer systems organization** → **Parallel architectures**.

Keywords

MaX, Material science, HPC, GPU acceleration, exascale

ACM Reference Format:
Lubomir Riha, Ada Böhm, Ondrej Vysocky, Filip Vaverka, Elisabetta Boella, Daniele Gregori, Mattia Paladino, Pietro Delugas, Oscar Baseggio, Stefano Baroni, Francesco Andreucci, Kaveh Haghighi Mood, Andreas Herten, Gregor Michalicek, Daniel Wortmann, Stefan Blügel, Pablo Ordejon, Federico N. Pedron, Jose Julio Gutierrez Moreno, Rogeli Grima Torres, Fabio Affinito, Mariella Ippolito, Paolo Giannozzi, Luigi Genovese, Fabrizio Ferrari Ruffino,

Ivan Carnimeo, Marc Sergent, Erwan Raffin, Ivan Girotto, Alberto Garcia, Jan Jona Javoršek, Barbara Krašovec, Alja Prah, Sebastien Strban, Augustin Degomme, Cristiano Malica, Lorenzo Bastonero, Nicola Marzari, Nicola Spallanzani, Daniele Varsano, Andrea Ferretti, and Elisa Molinari. 2025. MaX - Materials design at the exascale - a EuroHPC Centre of Excellence: Recent selected results: Invited Paper. In *22nd ACM International Conference on Computing Frontiers (CF Companion '25), May 28–30, 2025, Cagliari, Italy.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3706594.3727577

1 Introduction

The overarching goal of the third phase of the MaX Centre of Excellence [3] is to provide the European materials simulation community – both developers and end users – with the capabilities necessary, not only to deal with the technological disruption resulting from the advent of massively parallel heterogeneous computing systems, but above all to leverage this disruption to tackle scientific challenges deemed thus far to be forbidding. In order to achieve this goal, we will follow in the footsteps of the first two phases — which have paved the way to porting, maintaining, and scaling up complex state-of-the-art computer codes to multiple heterogeneous architectures — and take previous achievements to the exascale level, not only by turning flagship codes into lighthouse applications able to exploit up to thousands of accelerated nodes, but also enabling them to work cooperatively within tightly bound exascale workflows, which have so far been considered simply unfeasible. In this process, we have also put in place co-design actions aimed, e.g., at testing and profiling electronic structure codes and mini-apps on advanced HW or at improving their energy efficiency, in order to contribute to the evolution of a stronger European technology and HPC ecosystem.

In the following, we list the key objectives of the third phase of the MaX Centre of Excellence. For each objective we highlight selected key results obtained within the first two years of activity. Extended descriptions and data are available on the MaX website, project and benchmark repositories [1–3].

2 Lighthouse applications

As a core mission to MaX, selected community codes for quantum materials modelling, largely adopted by a broad body of end users and already working on pre-exascale systems, are progressively turned into lighthouse applications, ready to run on new exascale HPC platforms, further developing their scientific capabilities, while keeping engaged their communities of developers and users.

2.1 Quantum ESPRESSO

Description and Scientific Objectives. Quantum ESPRESSO[1] [4] is a powerful tool for ab-initio calculations on materials reliant on DFT, based on plane waves and pseudoptentials. It consists of a suite of codes allowing one to simulate and predict the behaviour of a wide range of materials, from crystals to molecules. It is equipped with quantum engines, property calculators, and post-processing tools that give access to total energies, atomistic structures, charge distributions, bonding, vibrational and dielectric properties, electronic excitation spectra and many others.

[1]https://www.quantum-espresso.org

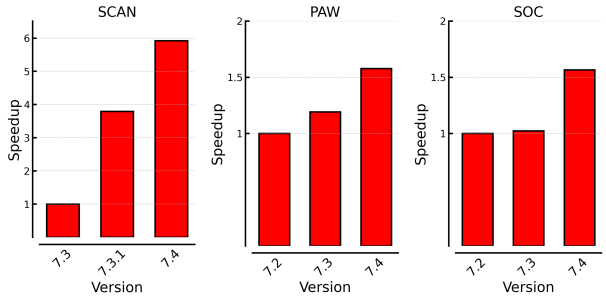

Figure 1: Speedups of recent QE versions with MaX optimizations of the local GPU parallelism.

Community and Dissemination. Quantum ESPRESSO is freely distributed[1] under GPL license and the development is open to external contributions.[2] The suite has a worldwide community of users, fostered over the years by recurrent dissemination and training initiatives, and has become one of the most used tools in the field. An equally diffused community of developers from many institutions and countries contributes to the scientific packages of the suite while the maintenance, HPC enhancement, and porting to new architecture have been, in the last decade, strongly supported by the MaX CoE. Thanks to the layered modular structure of the suite, the many community-contributed scientific developments can straightforwardly, or with minimal effort, benefit from the parallelization, optimization, and porting work done on the general kernels.

Parallel Computing Performance. The QE suite is designed to efficiently utilize high-performance computing (HPC) machines through multiple levels of parallelization. Depending on the node architecture, workload distribution can be managed using MPI combined with OpenMP multithreading or offloading tasks to GPGPUs. This strategy allows for computations involving systems with up to approximately 10,000 atoms and minimizes the computational cost associated with 3D Fast Fourier Transforms (FFTs) and operations on 3D data grids. Additional MPI levels provide further scalability, while band parallelization allocates operations across different wave functions. Optimized local GPU parallelism of the various kernels is then crucial to improve performance and throughput. As an example we report in Fig. 1 the speedup of time to solution for selected scientific use-cases, involving the use of PAW, SCAN functional, or spin-orbit coupling (SOC). Data are available in the MaX benchmark repository [1].

2.2 Siesta

Description and Scientific Objectives. SIESTA[3] [7] is a DFT code designed from the outset to perform efficient simulations for large systems. SIESTA's efficiency stems from the use of a basis set of strictly-localised atomic orbitals, which leads to moderate matrix sizes in diagonalization and enables the use of low-complexity methods (e.g., linear scaling and the PEXSI scheme) that exploit sparsity. A very important feature of the code is that its accuracy

[2]https://gitlab.com/QEF/q-e
[3]https://siesta-project.org

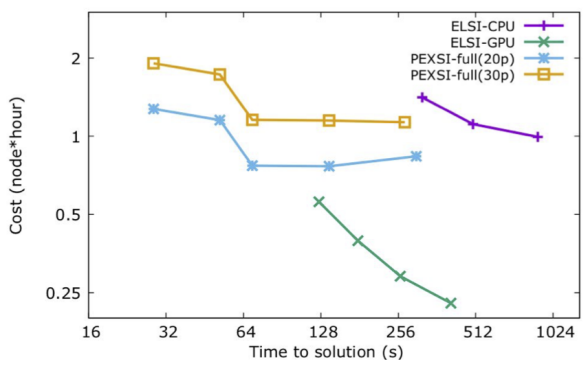

Figure 2: Relative performance of SIESTA with diagonalization (with and without GPU acceleration) and the massively parallel PEXSI solver for a problem with around 60000 orbitals. In this figure, horizontal lines denote perfect scaling.

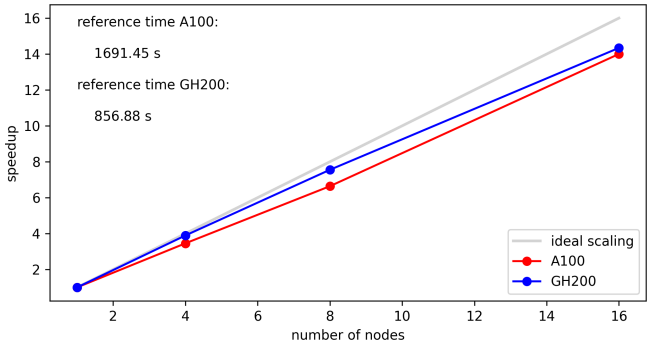

Figure 3: Scaling of FLEUR on compute clusters with GPU-accelerated nodes. Each node features either 2 AMD EPYC 7742 CPUs together with 4 NVIDIA A100 GPUs or a 4 NVIDIA GH200 Grace-Hopper superchip, respectively.

and cost can be tuned in a wide range, from very fast exploratory calculations to highly accurate simulations matching the quality of other approaches, such as plane-wave methods. SIESTA includes the TranSIESTA module for simulations under bias using the non-equilibrium Green's functions formalism. In addition to increased performance, the use of a basis set of atomic orbitals affords a very natural analysis of the electronic structure and the bonding effects.

Community and Dissemination. The possibility of treating large systems with first-principles electronic-structure methods has enabled new opportunities in many disciplines. The SIESTA program is open source (development and distribution hosted in GitLab[4]) and has become quite popular, being increasingly used by researchers in geosciences, biology, and engineering (apart from those in its natural habitat of materials science, condensed matter physics, and chemistry). Currently there are several thousand users all over the world, and the papers describing the method have received more than 12,000 citations so far.

Parallel Computing Performance. SIESTA supports very diverse hardware architectures, including GPU acceleration of diagonalization via the ELPA library and massive parallelization with the PEXSI solver (Fig. 2), which has also a small memory footprint. SIESTA's cost-effectiveness enables DFT calculations of hundreds of atoms even in personal computers. In HPC infrastructures, SIESTA can deal easily with systems involving thousands of atoms, giving it a distinct competitive advantage.

2.3 FLEUR

Description and Scientific Objectives. FLEUR[5] is an all-electron DFT code based on the full-potential linearized augmented plane wave (FLAPW) method, treating all electrons equally, unlike most pseudopotential-based DFT codes. It targets ground-state properties of 2D and 3D multicomponent solids, with emphasis on noncollinear magnetism, spin-orbit phenomena (e.g., Rashba effect, magnetic anisotropies), and magnon dispersion. Its interface with WANNIER90 enables computation of spin and charge transport properties (e.g., Hall effects, spin-orbit torque). FLEUR supports LDA+U, hybrid functionals, and advanced many-body methods (GW, GW+T) through SPEX. FLEUR's strength lies in modeling electronically and magnetically complex systems, such as those involving transition metals or rare-earths, making it ideal for studying spin-related properties and for generating parameters for multiscale simulations.
Community and Dissemination. FLEUR is distributed under the MIT license and supported by a growing user community.

Parallel Computing Performance. FLEUR's main computational challenges are related to the setup and diagonalization of dense matrices, which can reach sizes up to several $100k \times 100k$ elements. MPI parallelization allows to distribute the matrices over many compute nodes, enabling scaling from workstations up to the HPC infrastructures (Fig. 3). OpenMP and OpenACC parallelizations are implemented for individual compute kernels. The code is able to treat unit cells of several thousand atoms with local and semilocal exchange-correlation functionals and several hundred atoms with hybrid functionals.

2.4 BigDFT

Description and Scientific Objectives. BigDFT[6] [5] is a DFT code based on a systematic wavelet basis, enabling accurate simulations from molecules to extended systems. A key feature is its linear-scaling mode, which allows the treatment of large, complex systems—such as biomolecules, solid-state defects, or electrochemical interfaces with systematic precision. Within MAX, BigDFT evolves as a lighthouse code targeting exascale simulations, particularly of macromolecular systems. Its capabilities make it an ideal tool for interdisciplinary research, as exemplified by a recent quantum-mechanical study on SARS-CoV-2 mutations involving full electronic simulations of viral spike–receptor binding.

Community and Dissemination. Initially part of ABINIT, BigDFT became a stand-alone code in 2008 and has since developed into a modular suite of libraries. Its Poisson solver, based on interpolating scaling functions, has been adopted as a plugin by other DFT codes,

[4]https://gitlab.com/siesta-project
[5]https://www.flapw.de
[6]https://bigdft.org

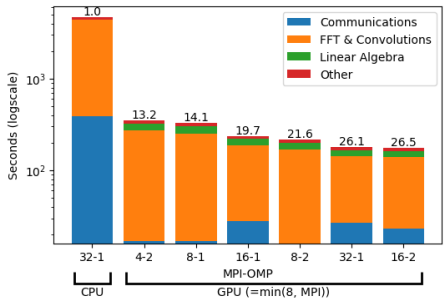

Figure 4: Performance breakdown and speedup of BigDFT for a dense system with highly accurate DFT calculation on an Intel Ice Lake + NVIDIA A100 architecture. The numbers above each bar indicate the speedup relative to the baseline CPU run (normalized to 1.0).

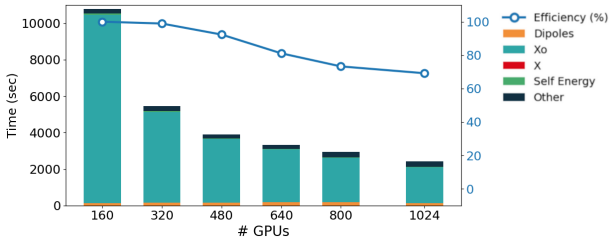

Figure 5: Quasi-particle corrections for a GrCo interface computed using YAMBO v5.2.3 on Leonardo-Booster. Numerical details can be found in the MAX benchmark repository [1].

including ABINIT and Octopus, due to its real-space accuracy. The user and developer community spans HPC, materials science, and quantum chemistry, with the wavelet formalism enabling both periodic and molecular simulations. BigDFT is distributed under GNU (L)GPL, with some libraries under more permissive licenses.

Parallel Computing Performance. BigDFT's use of wavelet convolutions as a core kernel led to the early development of custom GPU-accelerated routines. The code supports a hybrid MPI/OpenMP/GPU model, enabling efficient scaling across heterogeneous and multi-node systems, see Fig. 4 for a recent single-node example. Its flexible infrastructure adapts to different simulation types—linear-scaling, fragment-based, or cubic-scaling—making it well-suited to current and emerging exascale platforms.

2.5 Yambo

Description and Scientific Objectives. YAMBO[7] [6] is a community a open-source software designed for ab-initio calculations of both ground- and excited-state properties. It implements equilibrium and non-equilibrium Green's function methods to compute key physical properties such as band gaps, band alignments, defect quasi-particle energies, optical and non-equilibrium phenomena. YAMBO relies on precomputed electronic structures and interfaces with DFT codes like QUANTUM ESPRESSO. It enables calculations of quasi-particle energies, excitonic effects, temperature-dependent properties, non-equilibrium dynamics, and offers advanced post-processing tools. The code is actively developed and well-documented.

Community and Dissemination. Over time, YAMBO has fostered a growing user and developer community. It has contributed to numerous scientific publications, with more than a thousand citations, with a user forum counting over 900 members. Currently, the project is maintained and publicly available on GitHub.[8]

Parallel Computing Performance. YAMBO features a user-friendly interface, flexible I/O, and a hybrid MPI/OpenMP parallelization strategy, efficiently supporting GPGPU-based architectures. Its parallel workload distribution spans multiple levels (k/q grids, bands, quasi-particles), enabling large-scale simulations with up to tens of thousands of MPI tasks. Initially ported to GPUs using CUDA-Fortran, in the latest version (5.3.0) full and extended support for the deviceXlib library has been integrated, enabling beta-phase support for OpenACC and OpenMP offloading paradigms. Preprocessor macros minimize code duplication while optimizing integration of MPI-OpenMP with GPU programming models. Scaling tests on a graphene/Co (GrCo) interface confirm excellent scalability up to 1024 GPUs (Fig. 5). This is a prototype calculation that performs a complete GW workflow to compute quasi-particle corrections.

[7] https://www.yambo-code.eu
[8] https://github.com/yambo-code/yambo

3 Exascale workflows and data

The whole complexity needed to estimate materials' properties will be encoded in exascale-oriented workflows, allowing lighthouse applications to be driven automatically and properly orchestrated to exploit HPC capabilities at the exascale, to deliver resilience and fault tolerance and to ensure the full and meaningful dissemination of the entire simulation protocol, results, and data. These exascale workflows will tackle the complex logic, data movement, and possibly concurrent execution of multiple flagship codes, as needed by the targeted simulations.

Scientific applications that will be enabled by these multi-code exascale workflows include, e.g., the simulation of devices where the optical read-out of qubits is driven by excitons coupled with phonons; the non-adiabatic dynamics of photocatalytic electron-transfer reactions; the real-time multi-scale dynamics of complex topological structures in magnetic devices; the real-time spectroscopic characterization of amorphous materials or liquids; the mechanical evolution of a battery cathode during charge and discharge; the training on the fly of deep neural networks to drive the simulations of materials processing and manufacturing.

Key Results. (i) A set of target scientific grand challenges to be addressed by means of MAX exascale workflows has been identified, refined, and selected. The involved workflows have been rationalised and mapped to a few archetype workflow structures; (ii) Code interoperability requirements as posed by the selected exascale workflows have been identified; (iii) Early execution of selected steps of the workflows have started in order to provide a demonstration of the calculation feasibility and technical data to identify technical issues; (iv) Concerning FAIR data and storage, we have implemented a mirroring mechanism that enables the Materials Cloud data stored at CSCS to be also available from CINECA (and soon from the Jülich supercomputing centre).

Figure 6: MAX codes deployment status on EuroHPC architectures in January 2025.

Technical tools aimed at workflow implementation Tools such as the AiiDA orchestrator and the HyperQueue (HQ) meta-scheduler have been further developed and adapted. In particular: new AiiDA basic workflows have been coded (GW convergence, determination of Hubbard U+V parameters, core level spectroscopy), while new public releases of HQ offer new exascale-oriented features (improved scheduler, server checkpointing, a data transfer layer, and abstracted semantics for task completion). At the same time, more software tools aimed at encoding complex workflows have been identified and field tested (including, e.g., ZeroMQ for socket management or BigDFT's RemoteManager for task automation).

4 Benchmarking and deployment

One of the objectives of MAX is the continuous assessment and analysis of the parallel performance of the flagship codes, pointing up the direction for the development aimed at the effective exploitation of the existing technologies. To this aim, we make use of tools for code profiling and optimisation, and of the most recent programming models. An important output of the benchmarking and profiling activity is to discover and monitor code bottlenecks, to identify the code or architecture feature responsible for them (memory bandwidth, communication imbalance, latency, bandwidth to GPU, etc.), and to propose dedicated solutions also through the engineering of ad-hoc proof-of-concepts.

We adopted a common and homogeneous strategy for benchmarking and profiling activities of all MAX codes. In this regard, we have decided to use the same benchmarking and profiling tools for all MAX codes, and, to coordinate these activities, we organise periodic meetings between HPC experts and each MAX code community. These side-by-side meetings should have the purpose of continuously evaluating the code performances, optimising them, identifying possible bottlenecks and finding solutions, but also to train MAX code developers on benchmarking and profiling tools by gaining experience on these.

For the benchmarking activity, we adopted JUBE (Jülich Benchmarking Environment)[9], a workflow automation tool designed to standardise and automate performance benchmarking across varying software configurations, hardware architectures, and compiler setups. JUBE provides a script-based framework to easily create benchmark sets, enabling reproducible and systematic testing of code variants for different computer systems, allowing a straightforward evaluation of the results. All the necessary ingredients to perform the benchmarks of the MAX flagship codes (on the available HPC systems) have been collected in the JUBE4MAX repository[10]. The benchmark results, obtained in this way, are organized in the MAX benchmark repository[11] so that the data can be easily reachable from the users and developers. Thanks to the standardisation of the benchmarking data, we have set up a visualization tool that allows users to quickly see plots of the performance figures of different MAX codes and use cases.

We are committed to deploying the MAX flagship codes in all the EuroHPC systems, monitoring and continuously reporting their efficiency and performance. The status of deployment is summarised in Fig.6. In order to streamline the installation of the MAX flagship codes, we worked to create and maintain Spack recipes for each of the codes, which are adapted to the different computing infrastructures. These permit the codes to be deployed in all the major European HPC infrastructures. We are looking forward to extending this approach by including the EasyBuild tool.

5 Co-design, technology exploitation, and energy efficiency

MAX continuously monitors new HW and SW platforms and evaluates their suitability for the flagship codes. Also, the composition of the consortium allows MAX to co-design at different levels (chip, node, system) on HW that is relevant to the development of EuroHPC exascale machines.

Co-design for European technologies and Advanced HW exploitation Relevant parts applications are optimized for the Rhea processor, and the EUPEX prototype. These parts were identified and extracted as mini-apps from the MAX codes. Mini-apps are available and used for co-design, focusing on the efficient use of HBM and vectorization on ARM processors supporting SVE.

For the advanced HW exploitation task, interesting technologies that come up in the MAXCoE time frame to potentially improve the application's performance are identified. The technology provider partners supplied the roadmaps of major HW vendors. Supercomputing centers and technology providers gathered the details of available prototypes, and application owners provided the performance characteristics of the most significant kernels of the MAX codes. These efforts are summarized in the D4.1 deliverable [2]. MAX flagship codes have also been deployed, benchmarked, and tested on a selected advanced HW platform, like RISC-V or ARM-based processors, and compared to the latest x86 CPUs. Examples of such tests are presented in Table 1 and 2.

[9] https://www.fz-juelich.de/en/ias/jsc/services/user-support/software-tools/jube?expand=translations,fzjsettings,nearest-institut
[10] https://gitlab.com/max-centre/JUBE4MaX
[11] https://gitlab.com/max-centre/benchmarks-max3

Table 1: FLEUR benchmarks on advanced HW.

Platform	Peak performance [TFLOP/s]	TDP [W]	Time [sec]
AMD EPYC 7442	2 × 3.48	2 × 225	145.19
Sapphire Rapids	2 × 1.4	2 × 350	258.1
GH200(ARM, CPU)	1 × 3.57	~300	172.61
Grace (ARM)	2 × 3.57	2 × 250	79.79

Table 2: SIESTA benchmarks on advanced HW.

Platform	Runtime 1 core [sec]	Runtime 1 node [sec]	Cores per node / Speedup
Genoa (x86_64)	8971	194	128 / 46x
Grace (ARM)	12988	205	144 / 63x
SG2042 (RISCV)	81975	3392	128 / 24x

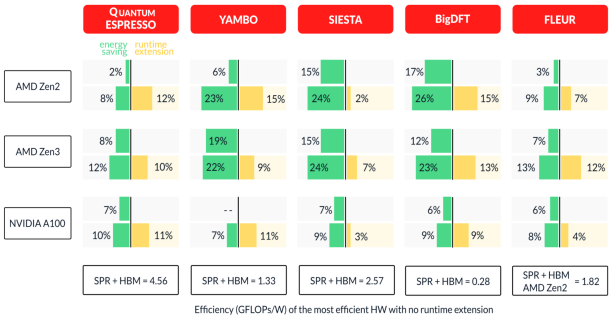

Figure 7: Energy efficiency analysis for various HW platforms.

For each code, we report the platform employed and its underlining architecture, execution time on one core and results on one node. In the latter case, besides the execution time, we indicate also the number of cores/node and the strong scaling speedup (in parenthesis).

Improving Energy Efficiency. Selected HW platforms that best represent the EuroHPC systems have been tested and evaluated in terms of energy efficiency using the MERIC tool developed at IT4Innovations. The optimization process focuses on finding the optimal HW configuration (CPU core and uncore frequencies and GPU streaming multiprocessor frequency) for the given application and the use case. The goal is to reduce energy consumption while maintaining the performance of the code, see Fig. 7 for results.

6 Training and Education in HPC for Materials Science

A key focus of the MAX CoE is the training and education of the next generation of users and developers of materials science codes in high-performance computing (HPC) environments. To address the diverse needs of its audience, MAX has designed a range of training formats. All events followed specific guidelines focused on equity, inclusion, and the effective dissemination of training materials. In recent years, it has organized specialized schools and hands-on tutorials aimed at equipping young scientists from academia and industry with best practices for the efficient use of MAX flagship codes in HPC. To complement this, and to provide comprehensive training for future developers, MAX has adopted the hackathon format, offering participants the opportunity to engage directly with leading code developers, HPC experts, and industry specialists in an interactive, hands-on setting.

Many events have been organized in close collaboration with National Competence Centres (NCCs) and leading institutions such as Psi-k and CECAM. In addition, MAX has actively contributed to Master, PhD and undergraduated programs offering specialized courses, seminars, and flexible teaching modules centred on its flagship codes. A key achievement of these efforts is the establishment of a sustainable repository of training materials, hosted on the Lhumos platform (https://www.lhumos.org), which serves as an enduring educational resource for the wider community.

7 Conclusions

The MAX CoE revolves around a set of selected open source, community codes from the domain of electronic structure simulation – namely QUANTUM ESPRESSO, SIESTA, FLEUR, BigDFT, and YAMBO. Among MAX core missions is the optimization, porting, further development, and extension of these applications in order to transform them into exascale-ready lighthouse codes. These are then the building blocks of exascale workflows aimed at addressing selected scientific challenges in the materials science domain. Enabling MAX flagship codes on large-scale heterogeneous HPC machines (featuring diverse GPUs and software stacks) requires facing a number of technical challenges (from benchmarking and profiling at scale, to the choice of programming models, to the deployment of the codes). Interestingly, these activities allow us also to directly interface with HW experts, using MAX flagship codes and the related mini-apps as co-design vehicles to test advanced HW solutions, as well as to track and optimize their energy efficiency.

In this paper, in addition to a presentation of the overall structure of the MAX CoE, we report some key achievements resulting from the first two years of CoE activity. These include porting and optimizing the flagship codes on EuroHPC machines, developing and deploying selected exascale workflows, setting up and populating a benchmark repository, testing novel hardware solutions and optimizing the energy efficiency of the flagship codes, as well as running an extended set of hands-on training events.

Acknowledgments

This work has been supported by the European Union through the MAX "MAterials design at the eXascale" Centre of Excellence (Grant agreement No. 101093374, co-funded by the EuroHPC JU and its participating countries).

References

[1] 2025. *MAX benchmarks.* https://gitlab.com/max-centre/benchmarks-max3
[2] 2025. *MAX public deliverables.* https://www.max-centre.eu/project-repository
[3] 2025. *MAX website.* https://www.max-centre.eu/
[4] P Giannozzi et al. 2017. Advanced capabilities for materials modelling with Quantum ESPRESSO. *J. Phys.: Condens. Matter* 29, 46 (Oct 2017), 465901. doi:10.1088/1361-648x/aa8f79
[5] L. E. Ratcliff et al. 2020. Flexibilities of wavelets as a computational basis set for large-scale electronic structure calculations. *J. Chem. Phys.* 152, 19 (05 2020), 194110. doi:10.1063/5.0004792
[6] D Sangalli et al. 2019. Many-body perturbation theory calculations using the yambo code. *J. Phys.: Condens. Matter* 31, 32 (may 2019), 325902. doi:10.1088/1361-648x/ab15d0
[7] J. M. Soler et al. 2002. The SIESTA method for ab-initio order-N materials simulation. *J. Phys.: Condens. Matter* 14, 11 (mar 2002), 2745–2779. doi:10.1088/0953-8984/14/11/302