




**HPC ALLIANCE FOR APPLICATIONS AND SUPERCOMPUTING
INNOVATION: THE EUROPE – JAPAN COLLABORATION**



**Funded by
the European Union**



EuroHPC
Joint Undertaking

This project received funding from the European High Performance Computing Joint Undertaking (EuroHPC JU) under the European Union's Horizon Europe framework program for research and innovation and Grant Agreement No.101136269. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or EuroHPC Joint Undertaking. Neither the European Union nor the granting authority can be held responsible for them.

Deliverable D6.2

First results and sustainability outlook

Deliverable Information

Deliverable Number:	D6.2
Deliverable Title:	First results and sustainability outlook
Project Title:	Hpc AlliaNce for Applications and supercoMputing Innovation: the Europe - Japan collaboration
Project Ref:	EuroHPC International Cooperation (HORIZON-EUROHPC-JU-2022-INCO-04)
Project Acronym:	HANAMI
Project Number:	101139786
Type of Action:	HORIZON JU Research and Innovation Actions
Topic:	HORIZON JU Research and Innovation Actions
Start Date:	2024-03-01
End Date:	2027-02-28
Duration:	36 months
Website:	http://hanami-project.com
Work Package:	6
Lead Authors:	Daniele Varsano (CNR)
Contributors:	Edoardo Di Napoli (FZJ), Giacomo Giorgi (CNR), Pablo Ordejon (ICN2), Alessandro Pecchia (CNR), Michele Casula (CNRS), Luigi Genovese (CEA)
Peer Reviewers:	Pablo Ordejon (ICN2), Edoardo Di Napoli (FZJ)
Due Date:	31/08/2025 (Month 18)
Submission Date:	01/09/2025

Dissemination Level:

- PU: Public

Disclaimer:

Funded by the European Union. Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or EuroHPC Joint Undertaking. Neither the EU nor the granting authority can be held responsible for them.

Contents

1	Introduction	6
2	Report on Task Development, Partnership Consolidation with Japan and outlook on sustainability of enlargement of collaborative partnership.	6
2.1	Task 6.1. Sustainable materials for photovoltaics. Task Leader: Giacomo Giorgi, CNR, Perugia	6
2.2	Task 6.2. Electrochemistry and battery research from first principles. Task Leader: Pablo Ordejón, ICN2, Barcelona	11
2.3	Task 6.3. Hydrogen storage: adsorption, desorption, reaction barriers from QMC. Task Leader: Michele Casula, CNRS, Paris	16
2.4	Task 6.4. Numerical linear algebra framework for Materials Science. Task Leader: Edoardo Di Napoli, FZJ, Jülich	19
2.5	Task 6.5. Non-equilibrium charge-carrier dynamics in 2d-materials. Task Leader: Alessandro Pecchia, CNR	22
2.6	Task 6.6. Enhancing Materials Science scalability through innovative block sparse libraries. Task Leader: Luigi Genovese, CEA	26
3	Collaborative Actions Across Project Tasks	27
4	Conclusions	30

List of Figures

1	Optimized structure of lead-free perovskite interfaced with (a) crystalline germania and (b) amorphous germania	8
2	Left: Optimized structure of quasi-2D structure of $\text{AVA}_2\text{MA}_{n-1}\text{Sn}_n\text{I}_{3n+1}$ ($n=5$) along with (in the inset) the structure of the spacer. Right: Top and side view of optimized crystal structure of the Dion–Jacobson perovskite (DAN) SnI_4 . Inset: Molecular structure of the DAN cation.	8
3	Strong scaling tests for full GW calculations using <i>yambo</i> . Left: Graphene on cobalt with a coarse k -point grid; Center: Graphene on cobalt with a dense k -point grid; Right: Cs_2PbI_4 perovskite. All simulations were performed with 4 MPI tasks per node and 12 OpenMP threads per task.	9
4	Giacomo Giorgi and Costanza Borghesi (CNR team) in discussion with the research groups of Koichi Yamashita and Azusa Muraoka on novel 2D and quasi-2D lead-free perovskite solar cell candidates and YAMBO code performance on the Fugaku supercomputer. Meeting held at Japan Women’s University, Tokyo on July 4th 2025.	10
5	Scaling tests of SIESTA in Deucalion, in preparation for Fugaku, showing the results obtained with three different Hamiltonian solvers: ScaLAPACK, ELPA and PEXSI. Left: Weak scaling - wall-time vs. system size, increasing the number of nodes proportionally to the system size). Right: Strong scaling - speed-up vs. number of nodes for a system with a fixed number of orbitals (10,014). The speed-up is, for each solver, relative to the compute time for the same solver on one node.	13
6	Cartoon representation of the WISE (water-in-salt electrolyte) containing acetonitrile in contact with a zinc anode and a carbon cathode. Code for colors and chemical species: white - hydrogen; blue - nitrogen; red - oxygen; green - chlorine; yellow - sodium; pink - zinc. Courtesy of Pedro Gomez-Romero (ICN2).	14

7	Radial distribution functions of ClO_4 (left), CH_3CN (center) and H_2O (right) around Na^+ (blue curves) and Zn^{+2} (red curves) ions in the WISE/acetonitrile solution. The shadowed area indicates the range of values obtained in the different MD runs.	15
8	Mean Square Deviation (MSD) as a function of time for the Na^+ (blue curves) and Zn^{+2} (red curves) ions in the WISE/acetonitrile solution. Values for the self-diffusion coefficients, obtained from the linear fits of the MSD as a function of time are also shown. The shadowed area indicates the range of values obtained in the different MD runs.	15
9	Pictures taken during M. Casula, M. Cherubini, and A. Raghav's visit of the NIMS campus (Tsukuba), in early April, together with K. Nakano and other members of his group.	17
10	Temperature versus pressure phase diagrams obtained for the LLT in pristine hydrogen (left panel, LLT line from our work in red color compared with previous works), for the structural transitions in H_3S (central panel) and in LaH_{10} (right panel) under pressure, by a combination of <i>ab initio</i> methods (in this case DFT), machine learning techniques, and molecular dynamics.	19
11	Weak-scaling benchmark measured on Miyabi and Leonardo for diffusion Monte Carlo calculations of the benzene molecule with 30 valence electrons, as a function of the number of GPUs.	20
12	Pictures taken during T. Imamura and K. Nagajima visit at FZJ in early March 2024, together with E. Di Napoli and other members of his group.	21
13	Elapsed time of ChASE in computing 5% of eigenpairs of three matrices - $\text{tol} = 1.0\text{e-}4$ (solid line) / $1.0\text{e-}5$ (dashed line) - $n = 23552$ (red), 64512 (blue), 79488 (green)	22
14	(left) 5x5 supercell of MoS_2/WS_2 with S-vacancy and (Center) DFT bandstructure for the 1x1 bilayer and folded 5x5 supercell. (Left) Comparison between DFT-PBE (orange) and G_0W_0 bandstructures (blue).	24
15	From Left to Right: First dark IL transition; first bright IN transition on MoS_2 ; bright IN transition on WS_2	25
16	A typical atomic structure for a system partitioned for transport calculations with left and right contacts and a central Device region.	25
17	Comparison of different linear algebra strategies integrated in BigDFT, showing the evolution of residual norms over inner iterations for wavefunction (blue), outer SCF cycles (purple), and density or potential (orange). Each panel corresponds to a different solver approach, assessing their convergence behavior, stability, and numerical robustness across representative test cases.	27

Executive Summary

Work Package 6 (WP6) of the HANAMI project focuses on building a sustained collaboration between European and Japanese research teams in the field of computational materials science. The work centers on the co-development of scalable simulation codes, numerical libraries, and scientific workflows tailored to emerging exascale HPC platforms. During this reporting period, WP6 has made significant progress across Tasks, with close coordination between European teams and key Japanese partners from institutions such as RIKEN, NIMS Tsukuba, the University of Tsukuba, and Yokohama University. These collaborations have led to shared advances in both scientific applications and HPC technologies.

In task 6.1 collaborative efforts were initiated on sustainable photovoltaic materials, focusing on lead-free hybrid perovskites and related materials. Simulation workflows based on the YAMBO code were deployed to assess their electronic and properties and thermodynamic stability. In Task 6.2, we have already ported the ESM code to the SIESTA version 5 branch, and have started work on next-generation electrolytes for battery applications, centred on water-in-salt electrolytes (WISE). The modeling activities aim to understand the impact of composition on viscosity, ionic mobility, and electrode–electrolyte interactions, in close coordination with experimental efforts. Task 6.3 involved the application of TurboRVB in Quantum Monte Carlo simulations of hydrogen-bearing materials, with benchmark results shared across teams. In Tasks 6.4 and 6.6, European and Japanese teams collaborated on scalable numerical methods, resulting in progress on the ChASE eigensolver for multi-GPU systems and the implementation and benchmarking of block-sparse matrix solvers within BigDFT. Finally, in Task 6.5, joint workflows based on BigDFT and libNEGF are being used to model carrier dynamics in layered and 2D materials, with emphasis on nonequilibrium conditions and quantum transport phenomena.

A major step forward of this period resulted from the successful joint EU–Japan proposal that secured about 6 million node-hours on the Fugaku supercomputer. Developed in collaboration with Japanese colleagues, the project provides a unique opportunity to benchmark and exploit HANAMI-relevant codes at scale on the Fugaku architecture. Most of the codes —such as SIESTA, YAMBO, TurboRVB, BigDFT, and libNEGF— have already been ported, and the grant enables their scientific deployment and performance testing on one of the world’s most advanced platforms.

Collaboration has been actively reinforced through joint scientific events. The WP6 session at the HANAMI High-Level Symposium (HLS1) in Barcelona gathered researchers from both regions for a two-day exchange on numerical libraries and energy materials. The upcoming workshop on advanced electronic structure methods, which will see the presence of key Japanese partners and scheduled for November 2025 in Paris, will further strengthen these international collaborations.

1 Introduction

In the pursuit of sustainable energy solutions, computational materials science plays an increasingly central role. With the rapid advancement of High-Performance Computing (HPC), researchers now have the capability to model complex materials systems with unprecedented accuracy and scale. These tools not only support fundamental understanding but also accelerate the discovery and optimization of materials for clean energy technologies such as photovoltaics, batteries, and hydrogen storage systems. Work Package 6 (WP6) of the HANAMI project is designed to reinforce the collaboration between leading European and Japanese teams in computational materials science and software development for HPC. The activities in WP6 are structured into two main pillars:

- **Collaborative research** focused on sustainable materials for energy applications, including photovoltaics, battery electrolytes, and hydrogen storage (Tasks 6.1, 6.2, 6.3);
- **Joint software co-development** targeting scalable algorithms, numerical libraries, and interoperable code frameworks suitable for pre-exascale and exascale supercomputers (Tasks 6.4, 6.5, 6.6).

Since the initial setup described in the previous report, WP6 has moved into the phase of active collaboration, implementation, and first scientific results. Notable progress includes simulation on energy materials, development of interoperable workflows, and shared benchmarking activities on HPC systems such as Fugaku or ARM partition of Deucalion, the Portugal-based EuroHPC-JU supercomputer. The close involvement of Japanese partners in each of these efforts has been essential, with joint proposals, coordinated meetings, and collaborative events further reinforcing the sustainability of these interactions.

This report documents the ongoing activities and progress of WP6. It is structured into task-specific sections, each detailing updates on scientific results, porting strategies, technical developments related to the codes, and collaboration strategies. The final section presents transversal and cross-task actions that have helped strengthen the EU–Japan partnership, including the organization of joint workshops, coordinated access to HPC resources, and emerging opportunities for sustained collaboration.

2 Report on Task Development, Partnership Consolidation with Japan and outlook on sustainability of enlargement of collaborative partnership.

2.1 Task 6.1. Sustainable materials for photovoltaics. Task Leader: Giacomo Giorgi, CNR, Perugia

The primary objective of this task is to investigate the electronic and optical properties of lead-free perovskite materials, a promising class of sustainable candidates for next-generation photovoltaic applications. The aim is to explore several material systems, with particular emphasis on evaluating their structural stability and potential for efficient light-to-electricity conversion. These assessments will be carried out through large-scale Density Functional Theory (DFT) and Many-Body Perturbation Theory (MBPT) simulations.

Given the complexity and size of the target systems, the use of the Fugaku supercomputer is essential to achieve the required computational accuracy and scalability.

This task is carried out by the group of Dr. Daniele Varsano (CNR-NANO) and Prof. Giacomo Giorgi (Univ. Perugia and CNR-NANO) in close collaboration with the Japanese partners, Prof. Koichi Yamashita (University of Tokyo and Yokohama City University) and Dr. Takahito Nakajima (RIKEN Center for Computational Science). Both are internationally recognised experts in first-principles atomistic modelling of materials for technological applications, with a specific focus on lead-free halide perovskites for sustainable energy technologies.

Extensions of collaborations

In line with the objective of strengthening and expanding the EU-Japan research partnership within the HANAMI Project, we have recently taken concrete steps to formalize new institutional collaborations. Notably, we have invited Professor Azusa Muraoka from Japan Women's University (JWU) to join the HANAMI consortium, recognizing the strong alignment of her research with the project's focus on High-Performance Computing (HPC) and its applications.

Professor Muraoka has an established track record in the theoretical investigation of photovoltaic materials, with a particular emphasis on organic photovoltaics, a field highly complementary to ongoing HANAMI research activities. She leads an active research group at JWU comprising one Ph.D. student, four Master's students, and six undergraduate students, fostering both academic excellence and training of young researchers.

To further solidify this collaboration, Japan Women's University and the University of Perugia are currently finalizing a Memorandum of Understanding (MoU). The agreement, initially set for a five-year period, will facilitate the exchange of students, early-career researchers, and faculty members, laying the groundwork for sustained scientific cooperation, joint research initiatives, and knowledge transfer between the two institutions.

Scientific framework

During the initial phase of the HANAMI project, we have actively pursued multiple research directions within our task, focusing on three main areas.

- Investigation of the interface between Sn- and Ge-based perovskites and oxide layers.
- The study of complex mixed 2D/3D halide perovskites incorporating organic spacers, of particular interest.
- The characterization of the optical properties of 2D perovskites, with a specific focus on understanding the relationship between material stability and the chemical nature of the organic spacer.

More specifically, the first research line, examining the interface problem, is based on previous work and aims to evaluate the effects of oxide layers (both crystalline and amorphous) that naturally occur on the surface of these lead-free perovskites (see Figure 1). This study, conducted in collaboration with the Koichi Yamashita group, is progressing well and is expected to culminate in a paper submission in fall 2025. In this context, the effect of perovskite termination and its role in stabilizing properties with respect to oxidation are of particular interest.

Inspired by experimental work that has recently appeared in the literature, the investigation of 2D/3D mixed halide perovskites with complex organic spacers is a more recent direction of research, also conducted in collaboration with Professor Yamashita and Professor Azusa Muraoka (the latter, from Japanese Women University (JWU), recently joined the Hanami project, *vide infra*). This line of work aims to explore novel materials with tailored optoelectronic properties by leveraging the unique structural characteristics of mixed-dimensional perovskites. In particular, we are nowadays focusing on perovskites with 5-Ammonium Valeric Acid, $\text{AVA}_2\text{MA}_{n-1}\text{Sn}_n\text{I}_{3n+1}$ ($n = 1-5$) as spacer (see Figure 2)

Turning to the third topic, the Dion-Jacobson perovskite (DAN) SnI_4 stands out as a particularly intriguing system for theoretical investigation, especially in light of recent experimental advances. Its structure, stabilized by a rigid and bulky diaminonaphthalene (DAN) spacer (shown in Figure 2), has been experimentally demonstrated to offer exceptional air stability among tin-based 2D perovskites¹. Building on this experimental foundation, a theoretical study

¹Triggs et al., ACS Energy Lett. 2024, 9, 1835–1843.




Figure 1: Optimized structure of lead-free perovskite interfaced with (a) crystalline germania and (b) amorphous germania




Figure 2: Left: Optimized structure of quasi-2D structure of $\text{AVA}_2\text{MA}_{n-1}\text{Sn}_n\text{I}_{3n+1}$ ($n=5$) along with (in the inset) the structure of the spacer. Right: Top and side view of optimized crystal structure of the Dion-Jacobson perovskite (DAN) SnI_4 . Inset: Molecular structure of the DAN cation.

of (DAN) SnI_4 provides an opportunity to probe the underlying electronic structure and excitonic properties in greater detail. Calculations are being performed using the YAMBO code on Fugaku supercomputer, with preliminary results already demonstrating the viability of the latest ported version on a realistic and experimentally grounded system like (DAN) SnI_4 .

Porting and First Benchmark campaign of the YAMBO code in the Fugaku supercomputer

The Many Body Perturbation Theory code YAMBO² was ported and benchmarked in the Fugaku machine. This initiative was carried out by the CNR group (D. Varsano, G. Giorgi, N. Spallanzani) as part of a trial access project (project ID:hp240326), in collaboration with Prof. Koichi Yamashita from the Yokohama City University. YAMBO version 5.3.0 was compiled on Fugaku using the YAMBO Spack recipe, provided via an external repository, leveraging the GCC compiler suite (version 12.2.0) and the Fujitsu-MPI library. The benchmarking focused on evaluating scalability while computing quasiparticle energies using the GW approximation,

²D. Sangalli et al., J. Phys.: Cond. Mat. 31, 325902 (2019)

a state-of-the-art method for accurately determining electronic gaps and ionization energies in semiconductors. The benchmarks were conducted on two technologically relevant systems. The first system is a graphene/cobalt interface, consisting of a graphene sheet adsorbed on a cobalt slab with a vacuum layer equal in size to the cobalt slab. This system serves as a tuneable magnetic substrate for spintronic applications³. To evaluate performance under different conditions, two Brillouin zone sampling sizes were tested.

The second system is a 2D halide perovskite, Cs_2PbI_4 , a prototype of an entirely inorganic perovskite with a Ruddlesden-Popper structure. This material has gained significant attention for its enhanced stability and performance in photovoltaic and photoluminescent applications. The test consisted of a full run to calculate a few quasiparticle energies, which correspond to the energy levels required to add or remove an electron from the system. These quantities are experimentally accessible through photoelectron spectroscopy. The complete workflow included the computation of dipoles, the independent-particle polarizability (X_0), the interacting polarizability (screened potential) obtained by solving a Dyson equation, the evaluation of the interaction self-energy, and finally, the solution of a linearized quasiparticle equation.

The results demonstrated excellent scalability, even at many nodes. For the graphene/cobalt system with 7 k points, the efficiency reached 75% at 256 nodes. The same system with 61k points maintained 75% efficiency at 1,000 nodes. For the Cs_2PbI_4 benchmark, the efficiency was 87% at 128 nodes.

While the benchmark results indicate strong scalability, further improvements in time-to-solution could be achieved by compiling with the Fujitsu compiler and implementing OpenMP dynamic scheduling. These optimizations will be the focus of the next benchmark campaign, which will also test on Fugaku the recent OpenMP parallelization developments carried out within the MaX project. These enhancements are particularly well suited for the Fugaku's ARM-based architecture, making it an ideal platform to exploit the optimized parallelization strategy

This benchmark campaign represented a significant step in extending YAMBO's reach to Fugaku, thereby enlarging its user base in the Japanese materials science community and enabling the exploitation of the resources on one of the world's most powerful supercomputers.




Figure 3: Strong scaling tests for full GW calculations using yambo. Left: Graphene on cobalt with a coarse k -point grid; Center: Graphene on cobalt with a dense k -point grid; Right: Cs_2PbI_4 perovskite. All simulations were performed with 4 MPI tasks per node and 12 OpenMP threads per task.

Mini-app extraction and reduced precision exploration.

Since January 2025, on the occasion of the First High-Level Symposium in Barcelona, the CNR team (D. Varsano) initiated a collaboration with William Dawson from RIKEN, aimed at exploring the use of reduced precision techniques within selected modules of the YAMBO code. This line of research is expected to be instrumental in optimizing performance for modern accelerated hardware architectures.

³G. Avvisati et al. Nano Lett., 18, 2268 (2018)

As an initial step, a dedicated mini-application (mini-app) has been developed to isolate the computational kernel associated with the response function, one of the most computationally intensive components of YAMBO. This kernel is central to several critical calculations, including quasiparticle corrections within the GW approximation and the computation of optical spectra via the Bethe-Salpeter Equation (BSE) framework.

The main objective of the mini-app is to restructure the computational workflow to explicitly expose GEMM-like (General Matrix-Matrix Multiplication) operations, thereby facilitating hardware-specific tuning and optimization. In particular, the mini-app will serve as a key testbed to assess the potential of reduced precision strategies, such as double-precision emulation techniques within the Ozaki scheme, for improving performance without compromising numerical stability.

The evaluation of these reduced precision approaches will be carried out by the Japanese partners at RIKEN, who bring specific expertise in this area. This task is currently ongoing.

Collaboration strategy



Figure 4: Giacomo Giorgi and Costanza Borghesi (CNR team) in discussion with the research groups of Koichi Yamashita and Azusa Muraoka on novel 2D and quasi-2D lead-free perovskite solar cell candidates and YAMBO code performance on the Fugaku supercomputer. Meeting held at Japan Women’s University, Tokyo on July 4th 2025.

Since the beginning of the project, we have systematically maintained a regular schedule of meetings (both virtual and in person) between project participants to report on the progress of research activities. These meetings complement the visits that the researchers involved in the project make to the projects’ units in Japan. In addition to the progress made in the study of photoconversion materials (and the possible submission/publication of articles in internationally renowned scientific journals), the collaboration also includes the discussion on the research activities of the Japanese partners’ students. In this framework, in June-July 2025, Giacomo Giorgio and Costanza Borghesi (a postdoc hired to work in this task) visited the Japanese Women’s University to meet with the research group led by Professor Azusa Muraoka and to discuss with her about the eventual development of “green” materials to focus on for the photoconversion mechanism.

To reinforce this joint effort and promote long-term sustainability, the partners jointly prepared a bilateral proposal under the JSPS-CNR funding scheme, aimed at supporting dissemination and training activities on the use of the YAMBO code and its adaptation to Japanese HPC environments. The proposal was positively evaluated, but was not funded in the current round. However, both sides have expressed strong interest in resubmitting the project in a future call, as it would directly support collaboration and long-term software co-development.

2.2 Task 6.2. Electrochemistry and battery research from first principles. Task Leader: Pablo Ordejón, ICN2, Barcelona

The project focuses on developing tools to deal with problems related to the interfaces between electrified electrodes (at a given external voltage) and electrolytes. These problems are relevant for a number of scientific and technological applications, from batteries to electrocatalysis/photocatalysis, corrosion, fuel cells, and many others. The aim is to leverage the first-principles codes developed by the European HANAMI members with the methods and software libraries specialized in describing electrochemical systems developed by the Japan partners.

In particular, the main part of the work is a collaboration between the group of Prof. Pablo Ordejón at ICN2 in Barcelona, and the groups of Prof. Minoru Otani at the University of Tsukuba and Prof. Yoshitaka Tateyama at NIMS. The group of Prof. Ordejón is one of the main developers of SIESTA, a DFT code capable to solve very large systems (with many thousands of atoms), which are needed to deal with these complex electrochemical interfaces.

In the first months of the project, SIESTA has been ported to Fugaku, the main HPC infrastructures of Japan, and combined with the ESM (Effective Screening Medium) method developed by Prof. Otani and collaborators, which is able to describe electrolytes in contact with electrified electrodes. We have also started applying these computational tools for the research on materials for batteries, in collaboration with the group of Prof. Tateyama (NIMS Tsukuba).

Collaboration strategy and extensions on Code co-development

Code co-development in Task 6.2 revolves around the implementation of the Effective Screening Medium (ESM) method in SIESTA. ESM was developed by Profs. O. Sugino and M. Otani and co-workers⁴, and is implemented in the ESM software package⁵, maintained mainly by Prof. M. Otani at Tsukuba University.

The ESM package was formerly interfaced with SIESTA, and was only compatible with old versions of the SIESTA code (versions 2.0.2 and 3.1). Modern versions of SIESTA (5.0 and beyond) contain a myriad of new developments, new property calculations, performance optimizations, modularization, and interfacing with other codes. Many of those improvements were achieved within the context of the EuroHPC-JU MaX Centre of Excellence project⁶ (in its past and current editions, 2015-to date). In particular, MaX produced important optimizations in performance in modern HPC infrastructures (including heterogeneous architectures of processors and GPUs). All these advances, however, were not compatible with the ESM package until HANAMI.

In October 2024, ICN2 hired a postdoctoral researcher, Dr. Nils Wittemeier, to work on merging of ESM to the most recent branch of SIESTA. He has worked in close contact with Prof. Otani, who has provided the old interface with the deprecated SIESTA versions, and the advice on how to build the new interface. The collaboration has taken place mainly remotely, by videoconference and by email. The implementation also benefited from the presential meetings of the ICN2 team with Prof. Otani when he visited Barcelona in January 2025 for his participation in the HANAMI High-Level Symposium (HLS1). During that visit, we also discussed the extension of the collaboration to implement interfaces between SIESTA and other packages developed by Prof. Otani and his collaborators, like the “Reference Interaction Site Method” - RISM⁷, which enables electrolyte simulations using continuum models. Work along these lines is currently underway, and we have programmed a research visit of Prof. Otani and one of his

⁴M. Otani and O. Sugino, Phys. Rev. B **73**, 115407 (2006)

⁵<https://sugino.issp.u-tokyo.ac.jp/esm>

⁶<https://www.max-centre.eu>

⁷S. Nishihara and M. Otani, Phys. Rev. B **96**, 115429 (2017).

PhD students to Barcelona on December 15-19, 2025, to wrap up this implementation and make a first release of the merged SIESTA + RISM code.

The porting of ESM to the SIESTA version 5 branch has been achieved, and is now in the batch of merging requests of SIESTA. The merge request and the code can be found in the SIESTA Gitlab site⁸. The link to the merge request (including its status and all the related codes) can be found here. The interface has been built in such a way that it is almost completely agnostic to the specific SIESTA versions, so that it can be easily adapted to any future version of SIESTA that we or other developers distribute. The merging and its incorporation to the tip version of SIESTA are now pending a final validation.

Porting of the SIESTA code to the Fugaku supercomputer

A significant part of the effort done in this Task 6.2 during the reporting period was the porting and installation of SIESTA to the Fugaku supercomputer, and preparing for the future optimization, benchmarking and full deployment there. Although a development version of SIESTA built by the Fugaku support team was available on Fugaku, this was an old and outdated version (pre-SIESTA 5). Its compilation parameters and libraries had not been tuned for the specific technical specifications of the machine. SIESTA 5 incorporates a larger number of libraries and dependencies and comes with a re-engineered build system based on cmake. Therefore, we have focused on installing and deploying this version in Fugaku.

Before gaining access to Fugaku, and in preparation for the porting and full deployment of SIESTA, in July 2024 the ICN2 team applied to access the ARM partition of Deucalion, the Portugal-based EuroHPC-JU supercomputer. As explained in HANAMI Deliverable D1.2 "Recommendations for a sustainable cross access for the supercomputing facilities", Deucalion is the only EuroHPC supercomputer with an architecture similar to Fugaku (both using Fujitsu A64FX processors). Our Deucalion-ARM development access project, EHPC-DEV-2024D08-062 (PI: J.M. Escartín), is running from 22nd August 2024 to 22nd August 2025. On Deucalion, SIESTA 5.2.2 was compiled using open source tools (GCC 11.3.0 + OpenMPI 4.1.4). While Deucalion also offers Fujitsu Software Compiler Package (which includes compilers and an MPI library), we decided to start on the first instance with the GCC compiler, since we knew that it can compile SIESTA for other (x86) architectures, while with an untested compiler (Fujitsu) we were more likely to find issues in Fortran language support. This proved to be the case, and, when we actually tried to compile SIESTA with the Fujitsu compiler on Deucalion-ARM, we hit a Fujitsu compiler bug. This compiler bug is already fixed on the latest versions of the Fugaku compiler, that are not installed on Deucalion yet. Therefore, we expect, once moving to Fugaku, to be able to use the full Fujitsu optimized software stack.

Some preliminary scaling tests of our non-optimized, GCC+OpenMPI compilation in Deucalion were performed. We used systems relevant for the planned work in HANAMI, involving interfaces of water with Na⁺ and Cl⁻ ions in solution, in contact with gold electrodes. The system sizes were moderate in size, reaching only up to some 4,000 atoms, and a small (SZP) basis set was used. Realistic calculations planned for the coming period will be significantly larger both in terms of number of atoms and, most importantly, of basis set size. Figure 5 shows the weak and strong scaling results obtained with three different Hamiltonian solvers: two dense diagonalization libraries: ScaLAPACK and ELPA⁹), and a sparse solver based on a pole expansion of the density matrix (PEXSI¹⁰).

To study the weak scaling, calculations were done for different systems sizes (from 5,070 to 40,560 orbitals), while the number of nodes was increased in the same proportion. The diagonalization data show a weak scaling that is slightly more favorable than what should be

⁸<https://gitlab.com/siesta-project/siesta>

⁹<https://elpa.mpcdf.mpg.de>

¹⁰L. Lin *et al.*, J. Phys.: Condens. Matter **26**, 305503 (2014)

expected from the theoretical $O(N^3)$ scaling of the total CPU versus system size (which would translate into a $O(N^2)$ scaling in the walltime in the weak scaling test). This is due to the improvement of the parallel performance with system size. In turn, the theoretical scaling of the total CPU by the PEXSI algorithm is $O(N^2)$, which should lead to a linear scaling of the wall time in the weak scaling test. However, the actual tests also show a scaling below the theoretical value (that is, sublinear), especially for larger systems, which again indicates that the algorithm performs better in parallel for larger systems. We can also observe that the diagonalization schemes are less time consuming than the PEXSI approach for smaller systems, but PEXSI becomes favorable when the size of the system reaches some 40,000 orbitals.

The strong scaling tests were performed in an admittedly small system of only 10,014 orbitals, which does not show the parallel performance potential for the larger systems to be computed in the work planned for the future. The results show an acceptable strong scaling for a modest number of nodes. We must notice that the system is so small that, for 8 nodes, the number of orbitals assigned to each core is just 26. However, the weak scaling tests discussed above make us confident that we will be able to use efficiently a significantly larger number of nodes when our system sizes reach more realistic values (both in terms of the number of atoms and the number of orbitals per atom in the basis set). Finally, we note that ELPA performs better than ScaLAPACK for large number of nodes, so it would be the preferred solver for diagonalization.




Figure 5: Scaling tests of SIESTA in Deucalion, in preparation for Fugaku, showing the results obtained with three different Hamiltonian solvers: ScaLAPACK, ELPA and PEXSI. Left: Weak scaling - wall-time vs. system size, increasing the number of nodes proportionally to the system size). Right: Strong scaling - speed-up vs. number of nodes for a system with a fixed number of orbitals (10,014). The speed-up is, for each solver, relative to the compute time for the same solver on one node.

The team in Task 6.2 participated in the application for access to Fugaku through a proposal presented as part of the Work Package 6 (see the detailed explanation in Section 3: Collaborative Actions Across Project Tasks, Subsection “Joint Access to Fugaku – Strengthening EU–Japan Collaboration through WP6” in this report). We obtained an allocation of 5,960,088 node hours, about 30% of which (around 3 M node-hours) are reserved for this Task 6.2. The period of execution runs from April 1, 2025 to March 31, 2026. At the time of writing this report, we have already compiled SIESTA 5.4.0 in Fugaku, where the open source software infrastructure is based on spack. Given our Deucalion experience, on the first instance we compiled SIESTA 5.4.0 on Fugaku using GCC 12.2.0 and the Fujitsu MPI library. This also proved tricky, since the setup of the fujitsu-mpi spack package provided by Fugaku seemed to have issues with the paths to the wrappers for the gcc compilers, and this affected the compilation of one of our dependencies (ELPA). While a support request is still open for this matter, we eventually managed to compile a version of SIESTA that supports most functionalities, the main missing ones being PEXSI (which we are trying to fix) and NetCDF support. Work is in progress to solve these issues, porting the compilation to the native Fujitsu Software Compiler Package, making realistic benchmarks and, eventually, optimizing the performance of the code in Fugaku.

Scientific results

As a first topic of research in applying our computational approach to problems related to batteries and energy storage systems, we have chosen the modeling of new electrolytes, and their interaction with the electrode materials. In particular, we are focusing on "water-in-salt electrolytes" (WISE), which have demonstrated good potential to reduce H_2 evolution reactions (thus suppressing self-discharge and corrosion processes), and at the same time extending the electrochemical stability window (thus allowing for larger operation voltages). In particular, we are considering the mixture of WISE electrolytes containing $NaClO_4$ and $Zn(ClO_4)_2$ and water at different concentrations, which are being used by experimental colleagues at ICN2 (Prof. Pedro Gomez-Romero) for Zn-based storage devices. The problems associated with high viscosity of these electrolytes are addressed experimentally by adding organic solvents such as acetonitrile (CH_3CN). We are therefore making simulations of how the presence of acetonitrile modifies the viscosity and the ionic mobility of the electrolyte, and what is the optimal mixture composition (salt concentration and water/acetonitrile ratio) to improve these properties. Figure 6 shows a cartoon of the electrolyte in contact with zinc and carbon electrodes (anode and cathode, respectively).




Figure 6: Cartoon representation of the WISE (water-in-salt electrolyte) containing acetonitrile in contact with a zinc anode and a carbon cathode. Code for colors and chemical species: white - hydrogen; blue - nitrogen; red - oxygen; green - chlorine; yellow - sodium; pink - zinc. Courtesy of Pedro Gomez-Romero (ICN2).

We plan to run full-scale simulations of the interface of the electrolyte and the Zn and C electrodes using the Fugaku allocation obtained as described in Section 3. To prepare those simulations, we have started with a study of the electrolytes in the bulk, without the presence of the electrodes. This will allow us to have a reference of the structure of the electrolyte far from the electrode surfaces, and quantify the effect of its composition (salt concentration and water/acetonitrile ratio) on the diffusion constants of the cations (Na^+ and Zn^{2+}). Our first calculations have considered a composition used by the experimental group: 1 $NaClO_4$; 0.1 $Zn(ClO_4)_2$; 3 H_2O ; 2.4 CH_3CN . We have performed classical MM simulations of 100 ns, and repeated them for 20 different initial configurations in order to obtain converged averaged properties. Figure 7 shows the coordination shells of ClO_4^- , CH_3CN , and H_2O species around the Na^+ and Zn^{2+} ions. Significant differences in the Na and Zn solvation shells are observed. Figure 8 displays the mean square deviation for the Na^+ and Zn^{2+} ions as a function of time. The self-diffusion coefficient, which measures the mobility of the ions in the solution, can be obtained from the slope of these curves. We find that the mobility of the Na^+ ions is about four times larger than that for Zn^{2+} . This result can be correlated with the differences in the

solvation shells of both ions, as Zn^{+2} is more tightly bound to the solvent molecules (both ionic and neutral), hindering their motion. The diffusion coefficients are also considerably smaller than those for the ions in pure water ($1.3 \times 10^{-5} \text{ cm}^2/\text{s}$ and $7.0 \times 10^{-6} \text{ cm}^2/\text{s}$ for Na^+ and Zn^{+2} , respectively). These results are in full agreement with the experimental observations, and are an excellent starting point to study the effect of composition in the bulk, on one side, and of the effect of the interaction with the electrodes on the other. We are currently setting up these calculations in Fugaku. While the simulations in the bulk are done using classical interatomic potentials, the interaction with the electrodes will be modeled using the QM/MM capability of SIESTA, using DFT to describe the electrodes, and eventually the solvent or electrolyte species that participate in chemical reactions with the electrode (such as the solution of Zn atoms from the anode in the form of Zn^{+2} ions).




Figure 7: Radial distribution functions of ClO_4 (left), CH_3CN (center) and H_2O (right) around Na^+ (blue curves) and Zn^{+2} (red curves) ions in the WISE/acetonitrile solution. The shadowed area indicates the range of values obtained in the different MD runs.




Figure 8: Mean Square Deviation (MSD) as a function of time for the Na^+ (blue curves) and Zn^{+2} (red curves) ions in the WISE/acetonitrile solution. Values for the self-diffusion coefficients, obtained from the linear fits of the MSD as a function of time are also shown. The shadowed area indicates the range of values obtained in the different MD runs.

2.3 Task 6.3. Hydrogen storage: adsorption, desorption, reaction barriers from QMC. Task Leader: Michele Casula, CNRS, Paris

This task is focused on the quantitative understanding of hydrogen and its chemistry in hydrogen-bearing systems, such as superconducting hydrides (the so-called “super-hydrides”) and hydrogen-hosting clathrate hydrates. Modeling hydrogen-mediated or hydrogen-involving interactions in a predictive way is a challenging problem because it requires an accurate evaluation of binding energies and a proper account of nuclear quantum effects (NQE). The difficulty of finding a consensus in the determination of the phase diagram of a system as simple as pristine hydrogen under pressure witnesses the difficulties present in the field. However, a recent work (Nature Physics **19**, 845 (2023)) has shown how the use of many-body quantum Monte Carlo (QMC) methods for the electronic structure coupled with an advanced treatment of NQE can overcome the main bottlenecks. The combination of these state-of-the-art methods delivered the so-far most accurate high-pressure phase diagram of hydrogen, finally reconciling previous experimental findings. Within this framework, the high-performance computing (HPC) resources are badly needed, due to the large computational cost of both QMC and the proposed methods dealing with NQE.

In this task, we address both points, namely the accurate evaluation of the electronic structure and the precise account of NQE and related lattice vibrations, by using QMC as method of choice to evaluate the chemical binding energy, and the path integral molecular dynamics (PIMD) to take into account NQE. We employ the TurboRVB code, which has both capabilities, as it implements efficient QMC algorithms and it also includes a built-in PIMD code. M. Casula (CNRS, France) and K. Nakano (NIMS, Tsukuba, Japan), both involved in the WP6.3 task, have been the main co-developers of the code. More specifically, the plan is to combine QMC with Machine Learning (ML) techniques to generate fast and accurate potential energy surfaces, and enhance the portability of TurboRVB towards HPC and its interoperability with existing ML software. This is done also thanks to the recent deployment of TurboGenius, a workflow manager seamlessly interconnected with TurboRVB, and developed by K. Nakano at NIMS, Tsukuba. The enhanced workflow streamlines the complex sequence of steps required to generate the first set of nuclear configurations, perform QMC calculations of electronic energies and forces, train a machine learning interatomic potential (MLIP), and ultimately carry out PIMD simulations. The scheme is general and portable, and can be run on various computing systems across Europe and Japan.

Collaboration strategy and its extensions

Thanks to the fertile environment provided by HANAMI, this research line has seen a significant boost. Indeed, the underlying collaboration, built around the previously developed scientific interaction between M. Casula (CNRS, IMPMC Sorbonne, France) and K. Nakano (NIMS, Tsukuba, Japan), has been developed with the arrival of two HANAMI postdocs, M. Cherubini and A. Raghav, hired in the IMPMC institute. They worked in the realm of generating new MLIP from *ab initio* simulations for two super-hydrides, namely H_3S and LaH_{10} , respectively. Their results will be published soon and will be briefly presented in the following subsection. In April 2025, we organized a visit to the NIMS institute in Tsukuba. Some pictures taken during our visit are shown in Fig. 9. That was a nice occasion to present our work in the institute and to broaden the collaboration by involving also R. Arita from the First-principles Materials Science Research Team, RIKEN Center for Emergent Matter Science and from the Department of Physics, University of Tokyo. He got interested in the project and he started to contribute during our discussions. Both Cherubini and Raghav left the IMPMC institute after the end of their HANAMI contract; Raghav applied to a position in Arita’s and Nakano’s groups and he got accepted. Not only from the scientific viewpoint but also from the management of

human resources, HANAMI has thus been a fundamental catalyzer for the development of this collaboration.

Another chapter of this collaboration has been the co-supervision of an external PhD student, G. Tenti, affiliated to SISSA, co-advised by both M. Casula and K. Nakano. The topic and the frame of this thesis co-direction have been profoundly inspired by HANAMI. G. Tenti successfully defended his PhD thesis on January 21st, 2025. His scientific results will be shortly summarized in the next subsection.

Besides these latest developments, we are thinking on ways to make this collaboration sustainable in the long term. We do not have a well defined plan yet, however the remaining HANAMI period will be useful to establish new collaborative platforms. One of them is already there: the co-development of the QMC TurboRVB code. Indeed, its exploitation, its usage dissemination, and the communications of its scientific outcome are going to be fundamental to define one of the possible pillars of a sustainable scientific relation between the CNRS and NIMS on these specific topics. See the thematic workshop description to find a working-case example on how the dissemination aspect of this pillar can be constructed and managed, under the HANAMI project. A suggestive perspective is represented by the Laboratory for Innovative Key Materials and Structures (LINK) in Tsukuba, an international laboratory born from the institutional collaboration between the CNRS and NIMS. Extending the LINK activities to those of interest for the HANAMI WP6.3 task could be an appealing possibility, to explore in the near future.



Figure 9: Pictures taken during M. Casula, M. Cherubini, and A. Raghav's visit of the NIMS campus (Tsukuba), in early April, together with K. Nakano and other members of his group.

Scientific framework and its outcome

Within the project, we obtained so far highly accurate potential energy surfaces (PES) for pristine hydrogen under high pressure conditions. This is thanks to atomistic potentials derived from machine learning schemes applied to training sets generated by the QMC TurboRVB code, by exploiting the accrued interoperability reached by the workflow development done in HANAMI.

To learn a good MLIP, accurate QMC energies and forces are needed for a extended set of ionic configurations. This work has been carried out during G. Tenti's PhD project, under the supervision of M. Casula and K. Nakano. In particular, we have been able to:

- obtain a very accurate MLIP from the QMC optimization of the many-body electronic wave function, and apply it to molecular dynamics simulations of the Hugoniot curve (Phys. Rev. B **110**, L041107 (2024));
- show that there is a cheaper way to get accurate QMC forces without fully optimizing the wave function, by introducing an *ad hoc* correction (Physical Review B **109**, 205151 (2024));
- demonstrate the accuracy of the MLIP obtained through the force correction above in the Hugoniot working case scenario, where fully optimized wave function data were already available for direct comparison (arXiv 2504.20481);
- improve the MLIP scheme based on the MACE neural network in difficult cases such as the liquid-liquid transition (LLT) in high pressure hydrogen (arXiv 2502.02447; see left panel of Fig. 10).

We also paved the way to generate accurate PES from QMC for the super-hydrides, such as H_3 and LaH_{10} , compounds that bear the highest superconducting critical temperature measured so far. For these compounds, we have been able to generate MLIP from density functional theory (DFT) calculations, and run accurate PIMD on top of these DFT-MLIP. The work is underway to upgrade the DFT-MLIP to QMC-generated MLIP is under way. For the time being, we have been able to characterize the structural phase transition that affects both compounds as a function of temperature and pressure, as reported in Fig. 10 (central and left panel, respectively). The interesting physical outcome has been the localization and full characterization of the quantum critical point, separating a symmetric high-pressure structure (paraelectric) from a low-pressure broken-symmetry structure (bearing some multi-polar order). The superconducting behavior is supposed to develop on top of this quantum criticality affecting the structure. The ongoing work has three main goals:


- to characterize the multi-polar order developing at lower pressures and temperatures;
- to understand the link between quantum criticality and superconductivity;
- to improve the DFT-derived MLIP by using QMC.

We highlight the fact that the PIMD calculations run with our ML interatomic potentials are challenging for the large size of the system together with the full quantum treatment of nuclear motion needed to describe the structural transition. In the central and right panels of Fig. 10, one can appreciate the remarkable impact of NQEs by comparing the huge displacement between the classical (red line) and quantum (black line) transition boundaries found in both compounds.

Generally speaking and in a longer-term perspective, through the framework we developed and used, one can carry out predictive calculations of hydrogen-rich compounds from first principles to engineer high-temperature superconductors and hydrogen storage materials.

Code co-development

We would like to report here some code co-development activity done under the stimulus and synergy ignited by the HANAMI collaboration. With the idea of building a more flexible QMC computational platform able to natively run on GPUs HPC architectures, we started developing a python version of the FORTRAN TurboRVB code. The python implementation called jQMC



2.4 Task 6.4. Numerical linear algebra framework for Materials Science. Task Leader: Edoardo Di Napoli, FZJ, Jülich

To achieve these goals, the project will implement a series of coordinated steps that combine technical development with strategic collaboration. Central to this effort will be the close partnership with Japanese colleagues at RIKEN (T. Imamura and W. Dawson) and the University of Tokyo (K. Nagajima), whose expertise in high-performance computing and algorithm design will play a pivotal role in shaping the framework. In parallel, the consortium will rely on the active participation of all EU partners participating in the first three tasks of Work Package 6, ensuring that knowledge, tools, and best practices are seamlessly shared between institutions. Together, these collaborations will guide the incremental design of the framework, its integration with existing libraries, and the validation of its performance through benchmarking activities.

As an initial step in the technical development, we have successfully ported our in-house developed numerical linear algebra library, ChASE, to ARM-based supercomputers in Japan, including Fugaku and Wisteria. In addition to this porting effort, we have implemented new functionalities that extend ChASE’s applicability to Materials Science codes such as YAMBO, thereby broadening its scientific impact and demonstrating the framework’s potential to accel-




Figure 11: Weak-scaling benchmark measured on Miyabi and Leonardo for diffusion Monte Carlo calculations of the benzene molecule with 30 valence electrons, as a function of the number of GPUs.

erate domain-specific applications.

Existing collaborations with Japan and their extension

The collaboration between the European partners and Japanese institutions has already been firmly established through ongoing joint activities with the University of Tokyo and RIKEN. These collaborations are evidenced by several successful Computational Time projects awarded within the HPCI framework, starting in 2023 and continuing at present. The awarded projects have provided access to flagship Japanese supercomputers such as Fugaku and Wisteria, enabling the execution of large-scale numerical experiments and the porting and optimization of the ChASE library to ARM-based architectures. These activities have not only advanced the technical understanding of how to develop the framework but also demonstrated the high degree of complementarity and synergy between the EU and Japanese research teams.

Building on this foundation, the collaboration is now set to be further consolidated and extended through long-term research visits. In particular, the FZJ team will undertake an extended stay at RIKEN, reinforcing ties with Japanese colleagues and advancing the co-design of numerical methods for next-generation HPC systems. This team will consist not only of the group leader, Dr. E. Di Napoli, but also of Dr. C. Richefort, a postdoctoral researcher hired with the support of the HANAMI project funding. This arrangement will ensure continuity and capacity building, while also facilitating deeper engagement in joint research activities, training, and the exchange of expertise. Such measures will strengthen the EU–Japan partnership in the area of numerical linear algebra for exascale systems and will serve as a cornerstone for future collaborative efforts beyond the scope of the current project.



Figure 12: Pictures taken during T. Imamura and K. Nagajima visit at FZJ in early March 2024, together with E. Di Napoli and other members of his group.

Porting ChASE on Japanese clusters

The porting of the ChASE library to ARM-based supercomputers Fugaku and Wisteria represented a key milestone in the technical development of the project. The process was carefully staged to minimize risks and maximize efficiency. Wisteria was first used as a staging ground, which significantly reduced the complexity of the subsequent deployment on Fugaku. As part of this effort, the build environment was adapted to the Fujitsu compiler toolchain. In particular, the CMake system was extended to accommodate the `-Nclang` flag and to switch from `stdc++` to `libc++`. Further modifications included patching `FindScaLAPACK.cmake` to ensure correct linking against the Fujitsu SSL2BLAMP libraries. In addition, the Boost library was replaced with Popl for command-line parsing due to incompatibilities with `libc++`.

Following the successful porting, a focused phase of performance engineering was undertaken. During this phase, new initialization bottlenecks were identified, particularly in the random number generation routines. Benchmarking studies were carried out to compare the performance of GEMV and GEMM operations between Fujitsu BLAS and ARMPL, providing a clearer view of the trade-offs between the available implementations. Furthermore, optimization of MPI process mapping based on the topological connections of Fugaku nodes resulted in a reduction of communication overhead by approximately 20%. These optimizations enabled the execution of strong-scaling benchmarks on real-world matrices from FHI-aims code, achieving problem sizes of up to 640,000 dimensions and thousands of computed eigenpairs.

Finally, a number of algorithmic innovations were introduced into ChASE to expand its applicability. The library was extended to support quasi-Hermitian eigenvalue problems (see next subsection), particularly those arising in the full Bethe–Salpeter Equation (BSE). In addition, the Ozaki scheme was incorporated into the GEMM operations, allowing efficient and accurate solutions on modern GPUs through the exploitation of Integer cores. These advances not only broaden the range of scientific applications of ChASE, but also demonstrate its capability to serve as a cutting-edge building block for exascale-ready numerical software in Materials Science and related domains.

Upgrading ChASE for pseudo-Hermitian Hamiltonians

The study of the optical properties and electronic structures of materials is of strategic importance. In order to improve energy production and storage, the energies of quasiparticles called excitons are extracted by solving the Bethe-Salpeter Equation (BSE). The computation of these

energies can be carried out by calculating a portion of the eigenspectrum of the Hamiltonian matrix. When the coupling term of the Hamiltonian can be neglected, the Tamm-Dancoff approximation reduces the BSE to a Hermitian eigenvalue problem. In that case, the literature already contains an extensive list of Hermitian eigensolvers that leverage the computational power of modern parallel machines. In particular, the *Chebyshev Accelerated Subspace Eigensolver* (ChASE) has good properties in terms of scalability and convergence to compute the most extreme eigenvalues of Hermitian matrices.

In certain cases, the coupling term of the Hamiltonian cannot be neglected, and therefore solving the BSE equation leads to computing the eigenpairs of a pseudo-Hermitian matrix. These matrices are twice larger and the spectral properties are less straightforward than in the Hermitian case. This task aims at upgrading ChASE for the computation of several thousands of eigenpairs of pseudo-hermitian Hamiltonians and with similar performance as in the Hermitian case. Each of the main steps of ChASE algorithm, the matrix-multiplication- enabled Chebyshev filter, the orthonormalization of the search space, and the Rayleigh-Ritz procedure, have been upgraded accordingly and for four different implementations: single CPU, single GPU, distributed CPUs, distributed GPUs. Figure 13 represents the performance of the multi-GPUs implementation in




Figure 13: Elapsed time of ChASE in computing 5% of eigenpairs of three matrices - $\text{tol} = 1.0e-4$ (solid line) / $1.0e-5$ (dashed line) - $n = 23552$ (red), 64512 (blue), 79488 (green)

computing 5% of eigenpairs of three different matrices. The method demonstrates relatively good scalability, except for the Rayleigh-Ritz part which relies on a non-Hermitian direct eigensolver GEEV for the diagonalization of the Rayleigh-Quotient. A recent improvement incorporates the formulation of an Hermitian Rayleigh-quotient to exploit the more efficient HEEVD kernel, which improves the overall scalability. In particular, we demonstrate that the convergence of the Ritz-values is quadratic as is in the Hermitian case. New benchmarks will soon be available in an upcoming publication.

2.5 Task 6.5. Non-equilibrium charge-carrier dynamics in 2d-materials. Task Leader: Alessandro Pecchia, CNR

Two dimensional materials such as transition metal dichalcogenides (TMDC) are promising candidates as energy harvesting materials, in particular heterostructures of TMDC have potential applications as thermoelectrics or as efficient photoconversion materials. For this reason, the primary objective of this task is to analyze the electronic and optical properties of hetero-bilayers

of MoS₂/WS₂ and perform non-equilibrium quantum transport simulations to investigate the charge carrier dynamics. This will better assess their potentiality as PV materials, especially in the presence of defects. The calculations rely on large-scale atomistic models and employ advanced electronic structure methods, including Density Functional Theory (DFT), Many-Body Perturbation Theory (MBPT)—such as the GW approximation and the Bethe-Salpeter Equation (BSE)—as well as Non-Equilibrium Green’s Functions (NEGF) to compute charge transport properties under illumination.

Given the complexity and scale of the calculations, we are leveraging the Fugaku supercomputer to achieve the required level of scalability and performance. This effort is conducted in close collaboration with Japanese partners, in particular the research groups led by Dr. T. Nakajima and Dr. W. Dawson at RIKEN.

Establishing the collaboration

Unlike other partners within the HANAMI consortium, the CNR-ISMN group did not have a pre-existing collaboration with Japanese research institutions. The HANAMI project has thus provided an ideal framework to initiate and strengthen such collaboration. During the first High-level Symposium held in Barcelona (January 2025), initial contacts were made with the group of Dr. Nakajima leading to the identification of a mutually beneficial collaboration with Dr. W. Dawson, centred around the BigDFT code—a tool co-developed with the group of Luigi Genovese at CEA Grenoble. To further solidify this collaboration, Dr. Alessandro Pecchia visited RIKEN for two weeks (April 21 – May 3, 2025). During this visit, the teams began interfacing BigDFT with the libNEGF code, developed jointly by CNR and JSC, and conducted initial testing on the Fugaku supercomputer. More details about this work are given below.

Scientific framework

During the first phase of the HANAMI project, we set the foundation with some preliminary tasks. In particular, we have performed

- DFT calculations of isolated MoS₂ and WS₂ materials and MoS₂/WS₂ bilayers without defects and in the presence of 1 sulfur vacancy located in different positions.
- GW calculations of the isolated slabs and the bilayer in order to extract the screened potential. This is important in order to later include the correct e-h excitonic interaction in the bilayer.
- BSE calculations of the bilayer with and without vacancy defects in order to understand the optical absorption.
- NEGF transport calculations of MoS₂, also contacted with Au, to understand charge injection in the material and to better understand the role of defects.

DFT calculations

Concerning preliminary DFT calculations of the bilayers, we have carefully reproduced energies and bandstructures of different stacking geometries for the 1x1 unit cell. We relaxed the cell with PBE functional and compared the effect of various van der Waals corrections, i.e., Grimme D2, D3 and DF3-B86R. The latter was chosen as it provided a good physical consistency. We confirm that the AA’ stacking is the most energy-favourable configuration, as previously reported in the literature. The lattice constant is found 0.31719 nm and interlayer distance is 0.3643 nm.

The S-vacancy defects have been studied considering a 5x5 supercell, which is at the limit of the GW and BSE calculations. Figure 14 shows the defected 5x5 supercell and the band

structure of the bilayer for the 1x1 and 5x5 cells. The orange/blue colors reflect the projection of the bands on the MoS₂ and WS₂. The defect states appear within the energy gap and are mostly localized on the 3 Mo atoms around the missing S. The Spin-Orbit coupling induces a splitting of 60 meV between spin states, potentially interesting for spintronics.




Figure 14: (left) 5x5 supercell of MoS₂/WS₂ with S-vacancy and (Center) DFT bandstructure for the 1x1 bilayer and folded 5x5 supercell. (Left) Comparison between DFT-PBE (orange) and G₀W₀ bandstructures (blue).

GW and optical properties

We have performed GW calculations using the YAMBO code. The non self-consistent G₀W₀ approximation gives the bandstructure shown in the left panel of Figure 14. As expected, the GW bands open the bandgap, shifting the virtual states upward. We observe that the gap of WS₂ is opened slightly more than that of MoS₂ by about 150 meV, changing the position of the relative absorption peaks. GW also introduces a relative shift downward of the conduction Q valleys located along the Γ -K path. This is certainly relevant when non vertical optical transitions are included.

The optical excitation spectrum has been obtained solving the Bethe-Salpeter equation implemented in YAMBO. The lowest transition states are obtained around the K-point. Only vertical transitions have been considered, i.e., no electron-phonon couplings are included to explore indirect transitions Γ -K or K-Q. The lowest transitions involve the 4 band minima in valence (VB) and conduction bands (CB), around the K-point. The CB bands are essentially localized on the d_{xz} , d_{yz} of Mo or W, whereas the VB involve the in-plane d orbitals of the transition metal and p_x , p_y orbitals on S. A degree of delocalization to the other layer is observed for the VB states. The lowest bright transition is an intra-layer (IN) localized on the MoS₂, followed by an intra-layer transition on WS₂. An almost dark inter-layer (IL) transition is also observed. This involves the transition from WS₂ to MoS₂ with an amplitude about 2 orders of magnitude lower than the bright transitions. Exciton state projected on the non-interacting bands are shown in Figure 15.

Interfacing libNEGF to BigDFT

We started the interfacing of the libNEGF code developed at CNR-ISMN by Dr. Pecchia with the BigDFT. This work was carried during a first visit of Dr. Pecchia to RIKEN and closely




Figure 15: From Left to Right: First dark IL transition; first bright IN transition on MoS₂; bright IN transition on WS₂.

involved a collaboration with Dr. W. Dawson. The first step is to create a prototype of the interfacing, in order to explore feasibility and limitations.

BigDFT can attain linear scaling by expressing the Density Matrix (DM) on localized support functions expanded in an underlying basis set of Daubechies wavelets. An in-situ optimization of the support function is performed also enforcing quasi-orthogonality. This scheme provides an optimized and quasi systematic basis set of minimal size, ideal for interfacing to libNEGF. The libNEGF code can read an externally provided Hamiltonian and Overlap matrices and computes the system Density Matrix using the NEGF formalism, assuming open boundary conditions imposed by an infinite repetition of the contact Hamiltonians. An example of a structure used in transport calculation is shown in Figure 16.




Figure 16: A typical atomic structure for a system partitioned for transport calculations with left and right contacts and a central Device region.

The externally provided Hamiltonian contains information of the Device and the 2 Principal Layers (PL) of each contact. A PL is a sufficiently large supercell with interactions restricted to first-neighbor PLs. This must be accompanied with metadata providing information of the partitioning. Internally libNEGF assumes that the contact PLs are repeated to form semi-infinite leads, effectively providing an open boundary condition to the central (Device) region of atoms. The output of the code is the DM on the Device block and the off-diagonal terms Device/Contact and Contact/Device blocks.

In the current implementation, the DM of the contact blocks are not provided. This is an extension that during the development of the interface turned out useful and is under development.

libNEGF also expects a specific ordering of the blocks, where the Device comes first and then the Hamiltonian of the 2 contacts, with the first PL near the Device region, as shown in Fig. 16.

A Fortran code has been developed to create simple structures with the correct ordering and input metadata. The test interface is currently based on simple example driver codes provided with libNEGF.

Another output of the library is the lead-to-lead electronic current. At the status of development of the interface, we are able to perform an equilibrium DFT calculation within BigDFT and feed the resulting Hamiltonian and Overlap to libNEGF to compute the coherent lead-to-lead transmission and the resulting current. This type of calculation can be understood as a quasi-equilibrium ballistic current calculation that is already of interest in systems with small applied bias.

In the following steps of developments we plan to

- Obtain the equilibrium DM from libNEGF and compare it to the one computed by BigDFT. In principles, these should differ only slightly.
- Drive the self-consistent (equilibrium) DFT loop calling sequentially BigDFT and libNEGF.
- Extend the libNEGF calculator to provide the DM in the contact regions. This should make the calculation of the contact charge density easier.
- Include explicitly an external bias ramp across the contacts and drive the self-consistent charge density loop.

2.6 Task 6.6. Enhancing Materials Science scalability through innovative block sparse libraries. Task Leader: Luigi Genovese, CEA

In the first year of the project, significant progress was made toward the integration and benchmarking of advanced block-sparse linear algebra libraries within the BigDFT codebase. This work leverages the CheSS and NTPoly libraries, respectively developed at CEA and RIKEN, in the context of a long-standing collaboration aimed at scaling quantum mechanical simulations on exascale architectures such as Fugaku.

A series of representative benchmark tests were conducted to compare the performance and accuracy of different solver strategies integrated into BigDFT. These include purification-based methods (with standard and tight convergence), Fermi Operator Expansion (FOE), dense diagonalization via the EigenExa library with finite-temperature smearing, and wave operator approaches at various temperatures.

The results, summarized in the accompanying table and figure, highlight important trade-offs between computational cost (memory, wall time, CPU minutes per atom) and physical accuracy (total energy). While dense solvers and FOE show excellent performance in terms of walltime and efficiency, purification and wave operator methods are more flexible and tunable, offering enhanced stability for large-scale or low-temperature simulations. The purification method with tight convergence, in particular, achieves the highest precision at the cost of increased resource usage.

The table below reports key computational metrics for each approach: peak memory usage (in MB), total wall time (in seconds), CPU time per atom (in CPU-minutes), and the final total energy (in Hartree). These results illustrate the trade-offs between computational cost and numerical precision across methods based on purification, wave operators, Fermi Operator Expansion, and dense diagonalization using EigenExa. All simulations were performed on the 1UAO protein, a biologically relevant system from the RCSB PDB database, used here as a representative benchmark for large-scale quantum mechanical calculations.

	Mem	Walltime	CPUmin/at	energy
Purification Method, tight convergence	5227	3833	16.66	-762.786
Purification Method	5211	2485	10.81	-762.781
Wave operator, Finite Temperature	5153	7063	30.71	-762.772
Wave operator, Finite Temperature 1k K	5045	1248	5.43	-762.769
Dense solver (Eigenexa) + smearing	4896	725	3.15	-762.764
Fermi Operator Expansion	4756	953	4.14	-762.764




Figure 17: Comparison of different linear algebra strategies integrated in BigDFT, showing the evolution of residual norms over inner iterations for wavefunction (blue), outer SCF cycles (purple), and density or potential (orange). Each panel corresponds to a different solver approach, assessing their convergence behavior, stability, and numerical robustness across representative test cases.

These benchmarks provide a strong foundation for selecting optimal solver strategies for future production runs on Japanese and European HPC systems, while informing ongoing efforts to merge and optimize the CheSS and NTPoly components for long-term sustainability within the BigDFT ecosystem.

In the following we plan to introduce input parameter profiles that the user can select to choose the optimal algorithm depending on systems' electronic structure. The collaboration between RIKEN-CCS and the Atomistic Simulation Laboratory is ongoing, fostered also by the CEA-RIKEN collaboration, and regular interactions are taking place along the CheSS-NTPoly interoperability. Also, the weak-scaling behaviour of the different algorithms will be the object of a forthcoming benchmark.

3 Collaborative Actions Across Project Tasks

WP6 session at HLS1 in Barcelona

At the High-Level Symposium (HLS1) held in Barcelona in January 2025, a two-day session dedicated to Materials Science and High-Performance Computing (HPC) was organized under WP6, bringing together key European and Japanese contributors to the HANAMI project. The first day was focused on HPC codes and foundational numerical libraries in computational materials science. This included presentations on the performance and capabilities of the codes underpinning the project's various research tasks (SIESTA, Yambo, TurboRVB, BigDFT, lib-NEGF).

Xinzhe Wu and Clement Richefort (FZJ) presented the ChASE eigensolver on multi-GPU systems, illustrating its scalability and potential for handling large-scale dense eigenvalue problems, and practical integration in Many Body Perturbation Theory code (Yambo). A talk by Toshiyuki Imamura (RIKEN) on the future of Eigenvalue solver for Materials Science followed, which addressed the challenges and perspectives in designing eigensolvers adapted to emerging GPU architectures and their integration into materials science workflows. William Dawson (RIKEN) concluded the session by showcasing innovative strategies to reduce numerical precision requirements in electronic structure calculations, an increasingly-relevant action to leverage next-generation HPC platforms.

The second day focused on application-driven research in the field of energy materials. Topics discussed included new generations of photovoltaic materials, hydrogen storage technologies, and electrochemical interfaces relevant to battery applications. These sessions put the computational developments discussed on the first day into a broader scientific context.

The WP6 session greatly benefited from the presence of several Japanese partners, including K. Yamashita, K. Nakano, M. Otani, T. Imamura, and W. Dawson. Their participation fostered in-depth technical discussions across all project tasks and helped to start new collaborative opportunities. Notably, the reduced-precision benchmarking strategies demonstrated for the BigDFT code were identified as a promising approach and are now being considered for extension to other simulation codes within the consortium.



Figure 18: Prof. Minoru Otani (University of Tsukuba) presenting the Effective screening medium (ESM) technique for treating electrochemical interfaces and possible strategies of its integration in the SIESTA code.

Joint Access to Fugaku – Strengthening EU–Japan Collaboration through WP6

As part of HANAMI Work Package 6 (WP6), a joint EU–Japan proposal was successfully awarded 5,960,088 node hours on the Fugaku supercomputer under the Period-A call. This award marks a significant achievement in the context of WP6's objectives to advance EU–Japan cooperation in computational materials science and software co-development for pre-exascale and exascale systems.

The proposal was developed and submitted through a coordinated effort among all WP6 partners in Europe and key Japanese collaborators, including Koichi Yamashita (Yokohama University), Yoshitaka Tateyama and Kosuke Nakano (NIMS Tsukuba), Minoru Otani (Tsukuba University), Toshiyuki Imamura and William Dawson (RIKEN). This strong collaborative consortium brings together domain scientists and HPC experts from both sides, enabling a synergistic exchange of knowledge and technological capabilities. The project focuses on pushing the frontiers of large-scale simulations for energy-relevant materials, specifically targeting topics

such as: Sustainable materials for photovoltaics, Electrochemistry and battery materials from first principles, Quantum Monte Carlo for hydrogen storage compounds, Scalable numerical libraries and abstraction interfaces, and Non-equilibrium charge-carrier dynamics in 2D systems. The central aim is to fully leverage the recent porting efforts that have successfully brought nearly all major European and Japanese materials science codes—including SIESTA, YAMBO, TurboRVB, BigDFT, and libNEGF—as well as advanced numerical libraries such as ChASE, EigenExa, NTPoly, CheSS, to Fugaku’s ARM-based architecture. This allocation provides a unique opportunity to benchmark, optimize, and scientifically exploit these tools at scale on one of the world’s most advanced supercomputers. Such effort clearly shows how HANAMI fosters practical and sustained collaboration between European and Japanese teams, through shared scientific objectives, as well as through joint technical work on code interoperability, scalability, and performance optimization. Thanks to the awarded computing time, partners can test and co-develop applications and libraries under realistic leadership-class HPC conditions. Ultimately, this initiative reveals the importance of HANAMI in feeding concrete collaboration between European and Japanese teams, enabling the coordinated use of HPC resources and domain expertise to pursue shared scientific objectives and address key challenges in materials science.

Thematic Workshop on Advanced Electronic Structure Methods: Strengthening EU–Japan Collaboration in Materials Science

As part of WP6 activities aimed at fostering EU–Japan collaboration in computational materials science, a thematic workshop will take place on November 6–7, 2025, at Sorbonne University in Paris. Organized with contributions from both European and Japanese researchers, the event marks a concrete step in reinforcing collaboration within the HANAMI framework, with a focus on electronic structure methods and their implementation on high-performance computing (HPC) platforms. The workshop is co-organized by researchers from CNRS/Sorbonne Université, CNR-NANO, ICN2, CEA, and FZJ. Several Japanese colleagues have already confirmed their participation, including Sergei Manzhos (Tokyo Institute of Technology), Takao Kotani (Tottori University), Toshiyuki Imamura (RIKEN), Ayako Nakata (NIMS/MANA), Azusa Muraoka (Japan Women’s University), and Terumasa Tadano (NIMS). Their expertise spans scalable GW and Bethe–Salpeter methods, quantum Monte Carlo (QMC) techniques, and high-performance linear algebra solvers—topics directly relevant to the workshop. The event will feature contributions from leading Japanese and European researchers on a range of topics, including large-scale simulations, recent advances in linear algebra libraries for materials science, Green’s function methods for spectroscopy, modeling and simulation of energy materials, machine learning in materials research, and the electronic and optical properties of emerging materials. It will also provide an opportunity to share recent developments in key HANAMI-relevant codes—including SIESTA, YAMBO, and TurboRVB—and to support ongoing collaborations between European and Japanese teams.

The workshop is preceded by a three-day training school (Nov 3–5), supported by the CECAM-FR-MOSER node and other funding sources. While the school falls outside HANAMI’s direct scope, it serves as useful preparation for attendees and has helped secure additional support, including computing time from HP-CAVe at Sorbonne Université.

Overall, this initiative illustrates how HANAMI is helping to build practical and sustained scientific collaboration, supporting both joint research and software development for advanced HPC simulations.

4 Conclusions

The progress achieved in WP6 over this period reflects a strong commitment to both scientific excellence and Europe–Japan collaboration. Technical and scientific milestones include high-accuracy simulations of photovoltaic and hydrogen storage materials, simulations of water-in-salt electrolytes for battery applications, the implementation of scalable solvers, and the development of workflows for non-equilibrium processes in low-dimensional materials. The allocation of nearly 6 million node-hours on Fugaku obtained via a competitive application highlights the maturity of the European–Japanese collaboration and offers a valuable opportunity to test and exploit HANAMI-relevant codes on the Fugaku platforms. Most of the targeted codes have already been ported to Fugaku’s ARM architecture, and the current focus is on benchmarking and scientific use. In parallel, thematic events such as the HLS1 symposium and the upcoming Paris workshop continue to foster sustained interactions across WP6 tasks and with Japanese partners. These events help consolidate technical discussions and provide the opportunity to set up new collaboration on both scientific and technical aspects.

Looking ahead, the sustainability of these efforts is supported at multiple levels. On the institutional side, national funding schemes and bilateral initiatives between European and Japanese agencies are being pursued to reinforce collaborations, even when individual calls are not immediately successful. Strong synergies are also developing with European HPC Centres of Excellence, which provide a natural framework for sharing software development practices and benchmarking strategies. Finally, long-term support from established scientific organizations such as Psi-k and CECAM ensures continuity in community building, and dissemination activities.

Overall, WP6 has effectively supported collaboration between European and Japanese materials science teams by facilitating the joint development and application of interoperable simulation tools. These efforts have led to concrete progress in modeling challenges related to photovoltaics, batteries, and hydrogen materials, making practical use of advanced HPC infrastructure to address scientifically and technologically relevant problems.