Ref. Ares(2020)7820603 - 21/12/2020

Horizon 2020
European Union funding
for Research & Innovation

E-Infrastructures
H2020-INFRAEDI-2018-1

INFRAEDI-02-2018: Centres of Excellence on HPC

EoCoE-II

Energy Oriented Center of Excellence:
toward exascale for energy

Grant Agreement Number: 824158

Deliverable 7.3

Mid-term Impact Assessment report

Project and Deliverable Information Sheet

EoCoE-II	Project Ref:	EINFRA-824158
	Project Title:	Energy Oriented Center of Excellence: toward exascale for energy
	Project Website:	http://www.eocoe2.eu
	Deliverable ID:	D7.3
	Deliverable Nature	Mid-term Impact Assessment Report
	Dissemination Level*:	PU
	Contractual Date of delivery:	M18 – 1 July 2020
	Actual Date of delivery	
	EC Project Officer	Evangelia Markidou

* - The dissemination levels are indicated as follows: PU – Public, CO – Confidential, only for members of the consortium (including the Commission Services) CL – Classified, as referred to in Commission Decision 2991/844/EC.

Document Control Sheet

Document	Title:	Mid-term Impact Assessment Report
	ID:	D7.3
	Available at:	http://www.eocoe2.eu
	Software Tool:	Microsoft Word
Authorship	Written by:	M. Ramalho
	Contributors:	E. Audit, M. Celino, P. Gibbon, G. Good, M. Lobet, E. di Napoli, J. Thelot, M. Wolf, A. B. Walker
	Reviewed by:	PEC, PSB, ECG

Document Keywords:

Table of Contents

1. Introduction

The assessment of the impact of the project is presented in deliverable D7.3 at mid-term (M18) and D7.7 at project end. This mid-term report will be used to re-assess the dissemination and exploitation strategy, as defined under WP6.

The document presents a qualitative assessment and, whenever possible, a quantitative assessment of the target indicators. In the latter case, the success indicator proposed at the creation of the project will be used as a benchmark.

First of all, the potential impact targets, as defined in the project call for proposals, are assessed. After that, the potential impacts specific to the two domain of research and innovation (R&I) in EoCoE-II are presented, namely impacts in computational sciences and in energy related R&I challenges. Finally, the potential impacts related to jobs, training and other socioeconomic parameters are reviewed.

2. Impacts related to the call for projects

2.1 Targets for Impact 1

Impact 1. European leadership in exascale and extreme-scale oriented codes and innovative algorithms/solutions that address societal challenges or are important for key scientific and industrial applications

(i) Integrating new cutting-edge HPC libraries in production applications (Expected: over 10 libraries integrated in production applications)

The list below presents the list of application codes in which cutting-edge HPC numerical libraries were integrated:

- Alya: MUMPS, PaStiX, MaPhyS, AGMG, and PSBLAS (and its preconditioners library MLD2P4). For information on preliminary results see deliverable 3.2
- Parflow: PSBLAS and MLD2P4 was interfaced to KINSOL for seamless use in Parflow. For information on preliminary results see deliverable 3.2
- Shemat-Suite: The integration of AGMG in Shemat-Suite through Petsc interface is work in progress
- SOLEDGE3X (the evolution of the code TOKAM3X): AGMG and PaStiX libraries were integrated.

In context of input/output (I/O) and data handling tools, the following libraries were integrated or in the process to become integrated:

- Alya: FTI (Fault tolerance interface)
- Parflow: PDI (Portable data interface)
- Shemat-Suite: PDI integration work in progress
- Gysela: PDI integration work in progress

(ii) Enabling energy-oriented applications on pre-exascale systems (Expected: five flagship applications able to run on more than 80% of the full system of the future PRACE largest machines)

In the race towards exascale, it is becoming clear that the technology of choice will be CPU-GPU hybridization. This is what makes up the computational power of many of the most powerful computing machines in the world as of today and probably also in the near future (see the most recent update of the 500 most powerful super-computers in the world, the TOP500[1]listing). The use of the GPUs as computing accelerators is a solution that is strongly promoted by the Exascale Computing Project[2] in the United States of America.The rise of GPUs in systems used for scientific computing can be explained by several factors. First, the graphics card industry is strongly driven by the gaming industry, which has contributed to constant innovations now for several years, whereas CPU related innovations have not occurred as often as GPU's. Second, the constant technological improvements in GPUs is the result of a collaboration between NVIDIA™ and AMD™ backed up by the leading industry's foundries. Third, and in terms of "pure" computing, the nature of GPUs computing - large numbers of small, simple computing units and massive parallelism - allows for significantly better theoretical performance than CPUs for the same energy cost envelope. For a decade now, the use of GPUs as a means for boosting the computing performance of scientific calculations has progressed significantly. Last but not least, the market for GPUs for supercomputing has been joined by HPC-AI computing convergence in recent years.

There is another possible way to reach exascale performances today, through CPUs based on the ARM architecture. This is the direction that Japan has taken for its future pre-exascale Fugaku machine. In this exascale computing scenario, a large part of the CPU optimizations carried out in the EoCoE-II project (node-level, vectorization and scaling) will also be beneficial when applied to optimisations in systems based on ARM's CPUs.

Given the rationale above, it is important to note that codes that can run efficiently on both GPUs and CPUs based computing architectures will have better adaptability on future pre-Exascale and Exascale supercomputers.

The following are the codes that, amongst the flagship codes selected at the start of the project, will have the capabilities to run on the biggest machines to come :

- Alya : Alya is a high-performance computational mechanics code that solves complex coupled multi-physics problems, mostly coming from the engineering realm. It is the flagship code of the Wind Scientific challenge. The main goal for Alya is to bring the code to Exascale, to tackle the simulation of full wind farm over complex terrain with up to 100 wind turbines. To achieve this goal, Alya has one of the most ambitious optimization projects. The code will be able to run very efficiently on both CPU and GPU machines.
- EURAD-IM*:* EURAD-IM simulates the formation and transportation of atmospheric chemical species and particles (aerosols) on the regional to continental scale. At the end of the project EoCoE-II, this application code will be able to use both CPU- and GPU-based computing architectures.
- libNEGF: Although optimization work on libNEGF is still at an early stage, previous work on similar prototypes (1D-NEGF during EoCoE-1) has already shown a very good scale-up on the Blue Gene Q machines. Once optimized and extended, libNEGF has the potential to do the same on the current and future PRACE machines. Early results show scalability of a medium-sized system to be able to execute with very good parallel efficiency up to 375 compute nodes of the JUWELS cluster. Optimization efforts will be initially focused on the CPU version and, depending on securing a new GPU resource specialist, will be ported on a multi-GPU cluster. The final target is to execute the code on one of the largest PRACE pre-exascale cluster.
- ParFlow: ParFlow is a parallel, integrated hydrologic model, which simulates surface and subsurface flow. At the end of the project, ParFlow will be very advanced in terms of CPU and GPU optimization.
- Gysela-X: GYSELA is a 5D full-F (regarding Vlasov equations) and flux-driven gyro-kinetic Fortran parallel code that solves Vlasov and Poisson equations to simulate electrostatic plasma turbulence and transport in the core of Tokamak devices. Gysela-X will be highly optimized for ARM CPU architecture, thanks to an extensive collaboration with Japan which has already started to bear fruits.

Some so-called "satellite" codes in EoCoE-II also have a sufficient degree of optimization to run efficiently on future pre-exascale machines:

- WaLBerla: WaLBerla is a fluid simulation code that uses the lattice Boltzmann method. The code has a high-level of parallelism and has been optimized for CPU-based architectures.

(iii) Porting applications to new European exascale technologies (expected: over two applications supported by EoCoE-II participating to the Extreme Scale demonstrator projects

Since mid 2019, a new set of EU Research and Innovation Actions has been launched under the umbrella of the EuroHPC Joint Undertaking (JU), aimed at supporting the procurement of European Exascale systems (e.g. EuroHPC JU Pre-Exascale and pilot systems, EuroHPC JU HPC system procurements). This initiative specifically aims to maximise application performance and efficiency in next generation supercomputers, with the help of software engineering techniques, programming tools and libraries, and explicitly encourages interaction with the Centres of Excellence. From the call H2020-JTI-EuroHPC-2019-1, two projects - DEEP-IO and DEEP-SEA - are connected to EoCoE-II through work on the TSMP (alias TerrSysMP) application suite for which ParFlow is the principle software component.
Another call was issued in April 2020 on *Advanced Pilots Towards the European Supercomputers* with a closing date of September 15th (H2020-JTI-EuroHPC-2020-1), where EoCoE flagship applications are again expected to feature. Further upcoming measures can be found in the EuroHPC workplan.

2.2 Targets for Impact 2

Impact 2. Improved access to computing applications and expertise that enables researchers and industry to be more productive, leading to scientific excellence and economic and social benefit

(i) Spread the use of exascale applications outside EoCoE-II consortium (Expected: over ten users of flagship applications outside the consortium (at least three from industry)

In the following lines, we describe the current status on the spread of use of EoCoE's flagship application codes outside the consortium (external users), both for academic and industrial users.

Code Alya
Most users of Alya are protected by privacy agreement and their details cannot be disclosed. We provide their numbers divided by categories

- Academic users
  - Universities (not affiliated to EoCoE-II) users: 25
  - Research centers: 9
- Industry users
  - 4 Private companies

Code EURAD

- Academic users
  - Rheinish Institute for Environmental Research at the University of Köln, Germany
  - Federal University of Rio Grande do Norte, Natal, Brazil.
- Industry users
  - None currently

Code libNEGF

- Academic users
  - libNEGF is included in the DFTB+ software and used by all users of this popular DFT code based on Tight Binding. Their number is once again not quantifiable but surely > 10.
  - libNEGF is also interfaced with the simulator TiberCAD and used extensively by the members of the group of Prof. Di Carlo at the University of Rome (Tor Vergata)

    - The phonon transport within DFTB+ is also used at the University of Dresden (prof. Cuniberti group), University of Lyon (prof. Niehaus group) and at the University of Shen-Zhen (prof. Frauenheim)
- Industry users
    - LibNEGF is included inside Materials Studio (owned initially by Accelrys and now Biovia). It is also the driver for the quantum transport of DMol3 and it is used by all their customers. Their number cannot be quantified but it is well over 10

Code KMC / DMC
- Academic users
    - Chemistry department at the University of Bath (group of Dr. Benjamin Morgan) which is outside the EoCoE-II project
- Industry users
    - the KMC/DMC codes have been developed almost from scratch as part of EoCoE-II, the scope for them to have external users has been limited, but we are in discussions with a few groups external to the University of Bath who are interested in using them.

Code ParFlow
- Academic users
    - 9 International partners outside EoCoE-II: Princeton University, Lawrence Livermore National Laboratory, Colorado School of Mines, Universität Bonn, Washington State University, Syracuse University, Institut des Geociences de l'environment, HPCS TerrSyS, Integrated Groundwater Modeling Centre.
    - over 100 users in Universities and Research labs worldwide.

Code SHEMat
- Academic users
    - SHEMAT-Suite is used by six research projects/researchers at the E.ON Energy Research Center, RWTH that are not related to the EoCoE-II project
- Industry users
    - SHEMAT-Suite is used by a SME: Geophysica Beratungsgesellschaft mbH, Aachen, Germany.

Code GyselaX
- Academic users
    - David Zarzoso: Aix Marseille Université, CNRS, PIIM, UMR 7345, Marseille, France
    - Etienne Gravier and Maxime Lesur: Université de Lorraine, CNRS, IJL, F-54000 Nancy, France
    - Laure Vermare: LPP, CNRS, Ecole polytechnique, UPMC Univ Paris 06, Univ. Paris-Sud, Univ. Paris-Saclay, Sorbonne Univ., 91128 Palaiseau France
    - Yuuichi Asahi: Center for Computational Science and e-Systems (CCSE), Japan Atomic Energy Agency (JAEA), 178-4-4, Wakashiba, Kashiwa, Chiba, Japan, 277-0871

2.3 Targets for Impact 3

Impact 3: Improved competitiveness for European companies and SMEs through access to CoE expertise and services

(i) Industry and SMEs dedicated to low-carbon energy and its variability, accessed through EoCoE-II pan-European network (Expected: over 10 consortia / industries / SMEs supported via consultancy, application optimisation and/or usage of EoCoE supported libraries)

Please read above under Impact 2, Target (i) "Spread the use of exascale applications outside EoCoE-II consortium" (Pages 8 and 9) the description of collaboration with academic consortia and industry (including SMEs) outside the consortia of partners in project EoCoE-II

(ii) Provide a valuable SaaS service

At this point in time, it is not yet possible to make an impact assessment on the value of the services provided via the SaaS portal developed in the EoCoE-II project. This is because the portal is not yet available online, although its launch is expected soon after the summer in 2020. The portal would serve as a tool to provide working examples of EoCoE applications to different types of stakeholders, including early adopters and researchers in academia and industrial environments. For more details on the status of deployment of the portal see Deliverable 6.3 pages 31-34).

After several months of uneasy exchanges with ATOS / Bull, a company that was, on paper, perfectly suited to provide a ready-made portal to showcase the EoCoE-II codes, the EoCoE team came to the conclusion that we needed to go in another direction. The offer from ATOS/Bull proved to be extremely expensive, and the portal would not have received any technical support after the project's end. This lack of technical support would have included critical updates, leaving the portal vulnerable - an unacceptable outcome given the need to protect critical resources. Given the role the portal will be playing in EoCoE's long-term future, as a mean to showcase the project's applications to stakeholders and to convince industrial partners of HPC's potential, we determined we needed to change our approach and find a secure, sustainable, long-lasting solution.

This alternative solution is a in-house PSNC portal, which is already operational. The PSNC team is confident it will be up and running early 2021, after they start integrating the EoCoE applications on it, and adapt the portal's web interface to the need of both applications developers and end-users.
This change of strategy will imply a change of resources allocation, as PSNC will need to move the € 140 000 originally allocated to the "other costs" category under "Software as a Service platform" to personal resources, to cover the personal costs needed to configure and update the portal.

The EoCoE team intended to precisely determine how this original tool could achieve maximum efficiency and be exploited to its fullest. This process was made possible through EoCoE's collaboration with META Group, a company that specializes in maximizing the exploitation of projects' results. Several workshops were organized to work on the SaaS portal, and ensure its creation would follow a market-conscious process, dressing questions such as potential competitors, target market, end goal, lifespan, etc.
As it stands, and as a result of this process, the general frame of the SaaS portal is as follows:
- The SaaS portal's goal is to allow users to easily test EoCoE HPC applications, and therefore see concrete examples of how to run applications on HPC. Users will be able to run predefined example problems and modify some parameters to see how they influence the computation, performances, and results.

- The SaaS portal will allow users to access application codes, as well as resources, without requiring that these users have the technical background normally required to run HPC code demonstrations. It will also allow users to circumvent the complexity of applying individually for HPC accounts on local HPC sites.
- The SaaS portal is a web-based service that dedicates HPC resources to demonstrations, where all users can directly register and test EoCoE applications for themselves. Its target market is wide, since it is comprised of all users interested by EoCoE results.
- There is no current competition for such a portal, and it could ultimately be extended to host HPC applications from other CoEs looking for a ready-made demonstration platform.
- The SaaS portal's time to market is short, given that it will be up and running early 2021, and fully operational during the EoCoE-II project's lifetime.
- The SaaS portal will be EoCoE's main demonstrator when it comes to industrial outreach, and will figure prominently in the project's dissemination materials once it is online.

2.4 Targets for Impact 4

Impact 4: Federating capabilities and integrating communities around computational science in Europe

(i)Integrating HPC and domain scientist teams in order to adopt a co-design approach (Expected: over 10 new collaborations triggered by EoCOE-II)

The structural separation of the EoCoE-II project into scientific and technical challenges means that new collaborations can be forged both within natural energy domain-centered communities (wind, meteorology, water, materials etc), but also between HPC-oriented teams, and via outreach to new renewable energy partners. Some examples of these follow:

- (FZJ) Cooperation between JSC and IEK-14 (Institute for Conversion Technologies; Prof. Steven Beale) on performance optimisation of fuel cell modelling since the beginning of 2019.
- (FZJ) New cooperation with IEK-13 (Prof. Michael Eikerling) on modelling and simulation of materials for emerging energy technologies (fuel cells, batteries and electrolysers)
- (IPP, FZJ) Strengthened bilateral cooperation between FZJ and IPP within the Helmholtz Programme Supercomputing. Here, the established contact through the EoCoE project served as a strong precedent to encourage the structural relocation of the IPP Garching theory groups from the energy programme to supercomputing.
- (CEA, FZJ) The strong established cooperation through EoCoE-I and -II naturally helped to establish several new bilateral collaborations within the freshly inaugurated Franco-German virtual lab AIDAS (AI, Data Analytics and Scalable Simulation) in areas including both HPC and energy modelling.
- (FZJ, ITER) Contacts between JSC and the modelling group of Peter de Vries at ITER established at the start of EoCoE-II have led to a joint PhD project on first-principles tokamak breakdown modelling
- (FZJ, DDN) The cooperation between JSC and DDN within the EoCoE-II project, to leverage the I/O capabilities of the IME system, helps to utilize and distribute this knowledge, as part of the general JSC support, towards various additional HPC communities.
- (CEA, AMU) A new collaboration between CEA and Aix-Marseille University aims at implementing shaped plasma geometries in the GyselaX code with state-of-the-art and HPC relevant numerical techniques. A co-supervised PhD thesis is ongoing, making use of non-equidistant splines.
- The electrostatic solver"coulomb_kmc" developed at University of Bath was used as an internal benchmark to assess the Azure cloud computing as a HPC platform.

- (FAU, BSC) Cooperation between FAU/RRZE and the Alya team on an in-depth performance analysis of the Alya code, including the extraction of hotspots into proxy apps. We are working towards an improved understanding of the requirements of the application to optimize performance and, eventually, the selection of an optimal hardware architecture. The progress is discussed in regular meetings.

(ii) Develop an integrated long-term vision for HPC in the energy sector(Expected: Produce a position paper with EERA)

EoCoE has been an active player in the creation of the Joint Programme on "Digitalization for Energy" within the European Energy Research Alliance (EERA). This programme will be a transversal programme including research roadmapping for HPC and Data Science and Artificial Intelligence as well as applicative programmes on modelling and simulation for Wind research, Materials research, and systemic studies on production and transport of energy. This joint programme is likely to be kicked-off in the Fall of 2020, once the EERA Executive Committee will officially approve it (based on a document describing its management and coordination/networking actions edited by collaborators of EoCoE-II and other EU projects). EoCoE-II will participate in the coordination of this EERA Joint Programme during its second half term. The Joint Programme "Digitalization for Energy" will involve all the key actors required to draw the "*Position paper on HPC for Energy research*" that includes the main research targets to be pursued from 2021 to 2027.

(iii) Implement joint actions with other European funded HPC projects (Expected: over six actions)

In the context of co-operation events with related projects on HPC technology at EU level, and wherever possible, meetings were organised with the aim to promote joint actions and create critical mass. Some of these meetings occurred in the framework of the project FocusCoE (Coordination and Support Action of the Exascale Centers of Excellence).

Out of such meetings, the following actions are those that are highlighted as actions with a potential impact on research and industrial endeavors in the medium to long-term (for a complete list of the over 25 joint actions with EU projects see Deliverable 6.3 from page 27):

- EERAdata project: Co-editing the terms of reference for the creation of the transversal Joint Programme "Digital for Energy" of the EERA association. This is expected to be the basis for the writing of the Position Paper on HPC for Energy involving key stakeholders from the EERA association.
- HPC CoE Council (HPC3): FocusCoE created the HPC3 platform to better align the European CoEs to fulfil their role within the HPC ecosystem. HPC3 allows CoEs to collectively define an overriding strategy, as well as a collaborative implementation of their interactions with the European HPC ecosystem. HPC3 also allows CoEs to promote and better concert their HPC capabilities and services towards both academic and industrial users. EoCoE-II played a central role in the creation of HPC3 and the EoCoE coordinator has been elected as the General Assembly Chair, where two other CoEs (MaX and BioExcel) are also members.
- EOSC Enhance project: contribution to the EOSC Catalogue with a description of EoCoE's services and featured in the EOSC portal. This participation aims to enhance the visibility of EoCoE's services towards new user communities and the project EOSC Enhance will provide valuable feedbacks on how to make EoCoE-II's services easily discoverable and enlarge its user community.
- HiDALGO: EoCoE and HiDALGO have initiated the process of becoming associate partners. This collaboration will include the organization of joint events, including training, workshops, and public presentations. EoCoE and HiDALGO will co-host a workshop in July 2021 and EoCoE-II will give a presentation on its application areas as well as a general presentation of the project. Given the

connection between the two CoEs, we are confident that this collaboration will yield interesting results and will be of great benefit for the promotion of EoCoE's achievements to a larger community.

- MAX: Participation in EUSEW2020 with a common booth at the Networking Village. MAX would focus on its work on materials for Energy in this event and EoCoE-II on the presentation of the SaaS portal. Due to the pandemic episode, the conference was organised as a virtual conference. Deeming that a face-to-face event would be beneficial, participation to EUSEW2020 with MAX and other CoEs will be organised for EUSEW2021 instead. EUSEW 2020 is the networking event related to Energy policy at EU level and has attendees from all types of stakeholders related to energy and energy transition to clean energy sources.
- ENERXICO is a joint project between Europe and Mexico that will develop high performance simulation tools beyond state-of-the-art in order to modernize the energy sector and meet future energy demands. The BSC- CSE department that develops Alya is involved in this project.
- There are also collaborations with other CoE projects COMPBIOMED and EXCELLERAT via the BSC-CSE department in the context of research on code Alya.

2.5 Targets for Impact 5

Impact 5: A large number of scientists and engineers, in particular female and young ones, trained in the use of computational methods and optimisation of applications

(i) Train researchers/engineers to new HPC technologies (Expected: over 200 trainees) (ii) Exascale performance optimisation high-level training (Expected: over 50 trainees)

EoCoE-II has offered a programme of webinars which have attracted a significant number of attendees. By 'significant' we mean that we compare better to the number of attendees expected to participate in training actions at the start of the project. to date the attendance has been over 200 participants of which 117 were participants external to the project.

An impact assessment analysis based on quantitative data defined at the start of the project cannot be as convincing as a benchmark of values for success indicators on training impact agreed amongst all CoEs. Still, it is positive that the number of participants has been higher than what was expected (467 attendees) for this first reporting period alone, and included a reasonable number of external attendees.

In percentage terms, the majority of attendees to EoCoE-II training events are from academia (51%) or from research institutes (38%). There are also 8% of attendees from industry. The percentage of women that took part in trainings accounts for approximately 17%. More details on training actions can be found in Deliverable 6.3 (pages 34-38).

[1] see https://www.top500.org

[2] see https://www.exascaleproject.org/

3. Impacts on computational challenges

In addition to the impacts specific to this call for proposals, EoCoE-II develops new computational technologies and addresses an important societal challenge (the transition to low-carbon and clean energy sources) that will have broader impacts than what we would be able to infer at this stage. The EoCoE-II consortium focus its efforts on enabling scientific breakthroughs in the Energy domain. It does this by designing and developing cutting-edge computational methods and high-end production-ready HPC software to bring the numerical tools supported in EoCoE-II to exploit the "soon-to-come" exascale computing systems. These new innovative algorithms/solutions will have an impact on key scientific and industrial applications, namely by producing better modeling tools.

By adopting cutting-edge software technologies for Exascale, the project explores multiple solutions that it will be able to promote to the community. Work Packages 2, 3, 4 and 5 focus respectively on different technical expertise and promote the best tools in each of these areas: computational performance, numerical solvers, Input/Output, ensemble runs. The project also contributes to the development of several tools and libraries (see the deliverables of WPs 2, 3, 4 and 5 in Month 18 for more details).

Through workshops, WP2 participates in the dissemination of good practices in the field of HPC. The methods adopted are presented to the community through EoCoE webinars, participation in workshops and publications. WP2 is also involved in the exploration of future Exascale technologies by working on GPU or on ARM architectures. The optimization successes encountered by EoCoE applications will serve as an example for other European codes, in particular for engaging in the path for Exascale computing.

A first example is the development done in the ParFlow to port the code to GPU-based architectures. The developers adopted a strategy of abstracting memory allocation and loop management by designing an in-house Domain Specific Language[1] (DSL). This was done for several reasons. Firstly, it allows to hide from scientific developers (physicists and numeric experts) the complexity of the code related to parallelism and HPC. Physicists can implement algorithms without worrying about which platform they will run on. Conversely, this allows HPC developers to optimize parallelism and support for new architectures seamlessly. This goes in the direction of a software engineering method oriented towards making the best use of cutting edge expertise and promoting the portability of code between computing architectures. In addition, developers can now test the Kokkos programming model to do this work with a C++ formalism. This will have an impact on computational sciences since the method used for this use case will serve as an example for porting many other applications to exascale systems.

Alya is another example of innovative optimization that can meet Exascale requirements. When talking about hybrid parallelism on both CPU and GPU, it is in fact rare that the two architectures work together asynchronously. In most codes, the CPU acts as a handler: it handles initialization, Input/Output tasks, communications and some other specific tasks that are better suited to this computing architecture. The GPU handles computationally intensive tasks. However, when one is running, the other is often on standby. What could formerly be done by the CPU is simply shifted to the GPU. This means that, in practice, part of the allocated computing power is systematically lost.

[1] A Domain Specific Language is (DSL) is a computer language specialized to a particular application domain. This is in contrast to a general-purpose language (GPL), which is broadly applicable across domains.

Today, task-based methods make it possible to use heterogeneous architectures asynchronously. Nevertheless these methods are still under research and are not yet integrated in the most common software frameworks. Alya uses a load-balancing method to distribute some of its calculations equitably between the CPU and the GPU. Calculations in both type of units are then performed at the same time. This allows to better manage the use of the resources and, therefore, to obtain the maximum of computing power out of the computation nodes. These developments have an impact since they can be seen as successful examples and may encourage other codes to adopt similar techniques.

Gysela is another example that can lead to potential impacts out of the actions in WP2. The Gysela development team has started a double partnership with CINES in France and RIKEN in Japan to optimize the application code on ARM-based computing architectures. RIKEN has chosen to develop a first pre-exascale demonstrator by developing its own HPC processors based on the ARM architecture. Europe is currently testing this possibility also. By working with the Japanese teams, Gysela is gaining valuable experience on this cutting-edge technology. The community will have learned a lot from this work to adapt to future ARM machines.

The impact is also magnified in terms of the development of new tools. WP2 and WP4 contribute to the development of the Parallel Data Interface (PDI) library. The Parallel Data Interface (PDI) is not a library itself but an interface that enables users to decouple all Input/Output processes from a code through a single API. PDI offers a declarative API for simulation codes to expose information required by the implementation of I/O processes. The latter are encapsulated inside plugins that access the exposed information. In an efficient software engineering approach, the codes should be split in independent parts in order to separate the different expertise (HPC, IO, physics). This approach is essential to the development of large simulation codes, especially industrial ones. For smaller development teams composed mainly of science experts, this facilitates access to effective parallel Input/Output without having any expertise in the field.

WP4 worked on the integration of the Infinite Memory Engine (IME) API, as designed by partner DDN, into higher-level APIs such as SIONlib and FTI, to allow easier utilization of the IME ecosystem. IME itself surfs as an intermediate caching layer between the application and the underlying file system and helps to deal with an increasing amount of input and output data while scaling an application to the Exascale level.

An example is the electrostatic solver developed as part of the research programme at UBAH, in which kinetic Monte Carlo (KMC) techniques are used to study charge carrier dynamics in solid state systems, such as electrons and holes in organic semi-conductors, or ions in solid electrolytes. In those simulations, large numbers of calculations of the electrostatic energy must be performed to evaluate the probabilities of the many possible particle “hopping” events proposed at each step of the simulation, of which only one is accepted. Due to the long ranged nature of the electrostatic interaction, the application of standard algorithms to calculate the electrostatic energy lead to KMC codes which scale sub-optimally with system size; for example, using a typical FFT accelerated Ewald summation leads to $O(N^2\log(N))$ scaling, while the standard Fast Multipole Method (FMM) leads to $O(N^2)$ scaling, where N is the number of particles in the system. However, by exploiting the fact that it is only necessary to calculate energy *differences* due to small changes in the configuration of the system, we have developed a modified version of the FMM algorithm which allows for $O(N)$ scaling in KMC simulations. Furthermore, the algorithm is highly suited to parallel architectures as 1) all energy differences for proposed moves can be computed independently and 2) when an event is accepted all data structures, that store the potential everywhere in the simulation domain, can be updated independently. The only communication that must occur is the specific details of the charge that hops between sites, and hence there is a very high ratio of computation to communication. The algorithm is implemented in a performing software library "coulomb_kmc", written in Python, with which we have demonstrated efficient scaling to systems containing 10^6–10^8 particles. Hence this library is suitable for simulating physically relevant systems. We have also demonstrated efficient parallel scaling on 4096 Intel Skylake cores and 8192 Thunder X2 ARM cores. The details of the algorithm, its implementation and scaling have been reported in a recent publication [W. R. Saunders *et al*., J. Comp. Phys. 410, 109379 (2020)].

In the field of Materials for energy, the new adopted flagship code, libNEGF, has the potential to become one among three codes worldwide (the other tow being OMEN and NEMO5, which are proprietary codes) to be capable of running full non-balistic quantum transport simulations at scale. In the current status libNEGF can include elastic electron-phonon scattering and simple corrections to the transmission probability due to absorptions/emissions of photons. The planned implementation of inelastic electron/phonon scattering (e.g., optical phonons), on the other hand, will introd uce important code improvements in the parallel communications in order to tackle problems of realistic size. This work will be carried out in the next half of the project within the WP1 and the WP2.

The resulting new flagship code (which will be re-named neXGf) will have several other applications behond the scope of the project. The applications of the NEGF techniques have been extensively applied in recent years in quantum optics, quantum corrections to the Boltzmann transport equation, high field transport in bulk systems and electron transport through nano-scaled devices. For instance, the possibility to simulate structures with lateral dimensions of 5-10 nm will open the possibility to study non trivial interfaces, alloy disorder, stacking faults and in general defects that typically require large super-cells in order to properly include the effects of charge localizations and charged defects. The problem is relevant both in modelling exact broadening of the PL spectrum as well as considering correctly injection and well to well charge transfer, including defect-mediated tunneling. In the emerging field of 2D materials, neXGf will be able to describe out of plane transport properties of stacked multilayers, where coherent transport might still play an important role.

In the Scientific Challenge “Water for Energy”, great emphasis is put on data assimilation with impact on climate change analysis and numerical weather predictions. In this context, the flagship code ParFlow together with CLM has been setup at a spatial resolution of 3 km over Europe. The meteorological forcing variables (variables such as temperature, precipitation, wind, vapor pressure, and downward longwave and shortwave radiation) were used as input to the model for the time period of the spinup and subsequent years simulated. Overall, multiple validations over the high-resolution European surface soil moisture reanalysis dataset (ESSMRA) revealed that data were consistent with the in-situ soil moisture data and other existing global reanalysis products. The relatively long time series and fine spatial resolution of this new European gridded ESSMRA dataset can provide a valuable data source for many hydrological applications. For example, it can be used as an initial input data for climate change analysis and for numerical weather prediction models to improve the model forecast in terms of location and amount of extreme precipitation events. Because of the scarcity of the in-situ soil moisture observations over large areas, this dataset can also be used for validation of SM outputs in modeling studies. This dataset will be also useful to understand the development and persistence of extreme weather events such as droughts, floods and heatwaves.

The results in the Scientific Challenge "Meteo for Energy" are already expected to have a key impact on meteorological research and by extension, associated industrial applications. The scale of the ESIAS simulations have allowed us to test 672 combinations of micro-, cumulative, and boundary layer physics schemes in the open-source Weather Research Model (WRF), fulfilling the need to investigate the stochastic weather forecasting skills of multiple schemes. WRF is widely used in the scientific community, including in Europe, where it was not developed and is not well calibrated. By identifying the optimal combination of schemes in WRF, we expect a broad impact on the accuracy of future meteorological research in Europe.

Similarly, research we have been conducting on next-generation forecasting technologies could have substantial impacts. ESIAS facilitates ensemble forecasts of ~1000 ensembles, whereas contemporary forecasts have less than ~100, and typically underestimate even ensemble spread. Capturing extreme cases and events can help to signal expensive forecasting errors for e.g. renewable energy. The confluence of cloud-tracking and ESIAS is hoped to greatly improve short-term wind power forecasting, which currently lack any sort of spatially-resolved information in real time. Finally, the integration of the MELISSA middleware is meant to improve simulation accessibility and speed up the impact on science and society.

4. Impacts on Energy Societal Challenges

EoCoE-II accesses and develops, at the fundamental research stage, new innovative algorithms and solutions that will strongly support the energy transition toward a strategic “energy mix”. As a mid-term impact, the knowledge generated by EoCoE-II will be valuable to further develop, master and commercialize both batteries and photovoltaic (PV) modules. In the mid- to long-term, the research developed by EoCoE-II’s partners has a high potential impact in the transition to a reliable, low-carbon and clean energy supply (see also work referred to regarding the creation of a new Joint Programme in the EERA association in chapters 2 under Impact 4).

4.1 Organic photovoltaic (PV) and perovskite solar cells (PSC)

Organic PV and PSC cells have the potential to provide abundant and low-energy-production PV modules at a lower cost than first- and second-generation solar technologies. By the use of flexible electrodes and substrates, fully flexible cells have been manufactured. There are markets for these cells in transparent solar windows, where there are versions optimised for underwater or indoor use and self-adhesive tape, on weak roofing and on building facades. Organic PV cells have achieved power efficiencies near 17%, but efficiency limitations as well as long-term reliability remain significant barriers. PSCs have power efficiencies near 25% but exploitation is hindered by short lifetimes and toxicity (due to soluble lead) of the most widely used perovskite materials. It is unlikely that perovskite cells will reach the prototyping stage unless lead-free perovskites are found and that have efficiencies exceeding 20%. Currently, organic solar cells and perovskite cells are each 1% of the global PV technology share. The impact in EoCoE-II comes from fundamental research that will allow EU companies exploiting these technologies to bring them closer to market. The EoCoE-II research on PV is undertaken at University of Bath who have developed mesoscale models of organic PV and PSC that are HPC friendly and are being used to model charge transport in the light harvesting and charge transport layers. EoCoE-II will aid the search for new materials, dopants and device architectures that have higher efficiencies and longer lifetimes than state-of-the-art cells.

5. Impacts on Training

The EoCoE-II project is a renowned platform able to train young scientists with both an interdisciplinary profile and a theoretical & applied sciences profile. This is because of the many different types of expertise present among the various highly qualified members of the EoCoE-II consortium. This training is very attractive for young scientists on the employment market, as many domains can be addressed by such multidisciplinary, top-level competencies.

Table 1 – Training impact

Partner	Impact
CEA	José Fonseca (postdoc) is gaining high-level skills in HPC programming methods, performance portability, numerical schemes dedicated to new architectures Kyle Reeves (Postdoc) has completed online training in machine learning with a specialization in deep learning (deeplearning.ai) and is currently continuing his training in reinforcement learning. Kevin Obrejan (postdoc) has formerly developed – from scratch – the 5-dimensional full-f gyrokinetic code GKNET to address plasma turbulence and transport in arbitrary tokamak geometries. His expertise further broadens in EoCoE-II by exploring the impact of non-circular geometry on collisional and turbulent transport in GYSELA-X simulations, with a focus on non-equidistant spline interpolation as a route towards exascale scalability. Julien Thélot (project manager), while not a researcher, has been trained on the European HPC ecosystem, including major actors, programs and challenges.
CNRS	Two young scientists (post-docs) will be trained, respectively, on the subjects of meteorology for energy (in the context of WP1) and scalable linear solvers and their use within high performance energy applications (in the context of WP3).
BSC	At the BSC-CASE department there are now 12 Ph.D students (one is female). Most of them will be trained in using the tools developed within the project.
CNR	1 Ph.Doc has been hired for 2 years and he is working on WP3
FZJ	As noted in Table 3 below, the staff hired at FZ-Juelich are situated in four separate institutes (Geosphere, Troposphere, Photovoltaics and the Supercomputing Centre). Each of these staff are embedded with teams working on related topics with the same or similar energy applications, so that a high synergistic exchange takes place both within and between all four participating FZJ institutes. PhD and MSc students also benefit from this arrangement, receiving direct on-the-job training in performance tools, parallel IO, programming models and reengineering of energy models from the postdocs and senior scientists associated with EoCoE-II.
DDN	The whole R&D team of the Paris office (15 persons) on the developments conducted within EoCoE-II will be trained. Some of the DDN engineers are also involved in teaching at the University and we expect some of the results of the EoCoE project to find their way into their teaching material.
INRIA	On WP3: 1 permanent senior researcher, 1 engineer (junior) has been hired, one 6 month internship has been working within the EoCoE-2 context sharing one month of his time co-working with BSC On WP5: 1 PhD
UBAH	In UBAH, there are 6-7 postdoctoral scientists and postgraduate students being trained in the methods being developed in EoCoE-II. Ages are 24-35 (postdocs), 21-24 (postgrads). UBAH also trains undergraduate project students who have made valuable contributions to EoCoE-II related research.

PSNC	1 senior system architect (engineer) and 2 junior software developers have been trained in usage of scalable, paralel computing and I/O libraries as well as in code optimisation for exascale computing and I/O processing. The knowlege aquired will be applied while consulting HPC end-users and user application developers as well as system administrators at PSNC and elsewhere in HPC computing environment in Poland.

6. Impact on social, industrial and academic domains

In general terms, methods and software developed in EoCoE-II will bring impacts at all levels of the social, industrial and academic spectra, some examples are given in the following table:

Table 2 - innovative algorithms/solutions created and impacts

Project result	Impact	Scope
Deterministic and probabilistic day-ahead and short-term regional and grid node forecasts for wind and solar power	Increase the quality of day-ahead and short-term solar and wind power forecasts. More stable power grid despite larger fractions of weather- dependent power plants, with benefits in power and weather forecasting	Industries and the general public will benefit from this impact. More specifically, better probabilistic forecasts are important in the power industry, as they facilitate more economical exchanges on the energy markets. There will be a long-term impact in pricing that will involve the general public.
Alya Multiphysics code – CFD + Fluid Structure Interaction	Increase efficiency of wind farms by: making energy cheaper and thus the utilities more competitive, allowing smaller turbines to produce same amount of energy as large ones and saving materials (Wind energy is 5cts/ kWh, which is cheaper than fossil sources which cost an average of 5.4 cts/ kWh and could be lowered to 2 cts/kWh); understand noise production and help decrease it (3dB reduction without loss of energy production) The algorithms developed will also have application in the automotive, aerospace and health industries.	Industry (energy utilities, manufacturers, etc.) and general public (i.e. wind farms' neighbour communities that will benefit from noise reduction). The algorithms will have application in any simulating rotating system (such as wheels), in simulating combustion turbines (such as airplanes), general CFD codes, and simulation of the human heart and respiratory systems. Therefore, this brings a wide range of societal, scientific and industrial applications.
Climate and hydrology scenarios/ simulation	These simulations will provide among other projections of state variables like soil water content, groundwater storage, etc and stream flows at selected large locations at a better resolution (normally locations taken are small and cover 12 km, while the project can take it down to 3 km). In the long term, these scenarios will allow to simulate water savings in agriculture and allow predicting extreme events e.g. floods.	Industry (hydropower utilities, water utilities, etc.) and general public (agricultural communities and communities affected by floods)
Algorithms for precise physical description of plasma	Uncertainty quantification techniques, will enable to have knowledge on the precision of the models	Main application is magnetic fusion energy. Unlimited source of clean energy.
Testing of innovative numerical algorithms and software for Linear Algebra kernels	Integration of the linear algebra software packages in some of the EoCoE-II flagship applications and analysis of robustness and efficiency within realistic simulations	These new algorithms and software will be re-used in other areas spreading best practices.
Improved weather models and power models reliant on meteorological data	To improve weather and renewable power forecasts by O(1) percentage points in the short-term (<6 hr) and O(0.1) percents for day-ahead.	Better probabilistic forecasts are important in the power industry, as they facilitate more economical exchanges on the energy markets. German control zones spent

		around half a billion on control power in 2017. Control power is also only applied after redispatch is used to balance the grid, costing another 1 B€. A perfect renewables forecast would thus eliminate balancing costs.
Algorithms for fast accurate electrostatic solvers and kinetic Monte Carlo simulations, adaptation of device Monte Carlo codes to hybrid organic-inorganic materials	Kinetic Monte Carlo simulations of 106–108 particles with accurate electrostatics, facilitating studies of particle dynamics in systems exibiting structural and stoichiometric heterogeneity on multiple length scales. Current applications are to organic and perovskite solar cells, in future also to batteries.	Software sellers of simulation codes for organic and perovskite devices. Firms selling materials and devices for these technologies. Firms exploiting these technologies in e.g. built in photovoltaics and to power sensors in the Internet of Things.
Optimal micro-physics, cumulative physics, and boundary layer schemes for Weather Research and Forecasting model (WRF)	Many European research groups rely on the public WRF model, which was developed and calibrated in North America and performs badly in Europe compared to operational models. EoCoE-II has enabled to test an unprecedented 672 combinations of WRF physical schemes over many months to determine the optimal combination in Europe.	Stakeholders of public research in meteorology, including the public and industry partners like in the energy sector will benefit from better and more applicable results.

• The ability to simulate devices at exascale will have a major impact on the device modelling capability for organic and perovskite solar cells, by allowing us to study more complex systems e.g. interfaces, and at longer length scales (100 nm), so we can explain a much greater variety of currently poorly understood behaviour in these devices.
• The development of the scenarios concerning expected changes in hydropower potential over the Italian Alpine region will help understanding the risk associated to future alterations of hydrologic regime, support companies with sustainable energy production planning and provide information to develop adaptation strategies for climate-change.
• Storage is about 10-15% of the energy footprint of a computing system. With EoCoE-II partner DDN, we expect a significant energy saving of the storage part following the optimization brought by the EoCoE-II project. Speeding up storage part by a factor of 3 should reduce the total energy envelop of HPC system by 5%.

7. Impact on job creation

The impact of the project EoCoE-II in terms of new jobs created is depicted in table 3 below.

Table 3 - Impact on job creation per partner in the EoCoE-II consortium

Partner	Impact
Faunhofer-IEE	One researcher full-time and one post-doc.
CNRS	One Post-doc position (CNRM/ Meteo France from September 2020 - September 2021) One Post-doc position (IRIT-INPT from July 2020 - December 2021)
CEA	One Post-doc position (CEA-IRFM, Kevin Obrejan, from mid-April 2020 to December 2021) One Post-doc position (CEA-MDLS, José Fonseca, from May 2019 to December 2021) One Post-doc position (CEA-CNRS, Kyle Reeves, from January 2020 to December 2020) One Project manager position (CEA, Julen Thélot, from June 2019 to December 2021)
CNR	We plan to transform the Post-Doc position funded by EoCoE II in a CNR permanent position
BSC	There will be at least 1 PhD student working closely with PostDocs working on EoCoE-II at BSC
FZJ	FZ-Juelich has hired the equivalent of 5-6 full-time staff working on the scientific challenges of WP1. In particular: 1 senior scientist, female, in the Water4Energy challenge (IBG-3); 1 postdoc, male, on the Meteo challenge (IEK-8); 1 scientist, female, on the solar power application (IEK-5). At JSC, funding has enabled the hiring or retention of 4-5 early-career HPC experts working on performance optimisation and parallel IO in the technical challenges of WP2 and WP4 respectively.
DDN	During the life span of the project, EoCoE will lead to the creation of 3+ full time R&D jobs
INRIA	WP3 : one engineer and one internship made their first professional steps and discover the R&D activity to build on their future career. One engineer will continue for an additional year working on the software package and its applications WP5: one PhD has been hired and there is an on-going process to hire an engineer.
ENEA	The recruitment procedures have been delayed. The next one is expected in the end of 2020. One researcher will be hired in a permanent position and 1or 2 will be taken as eligible to be hired permanently. Next year 2 postdoc will be hired.
UBAH	1.5 post-doctoral researchers full time.

Table of the Abbreviations

ARM	Trademark of Arm Holdings that develops the architecture and the licenses of a family of RISC-based architectures for computer processors.
CPU	Central Processing Unit
DSL	Domain Specific Language
HPC	High-Performance Computing
GPU	Graphics Processing Unit
RISC	Reduced Instruction Set Computing