ETP 4 HPC

EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

European High-Performance Computing Projects

Handbook 2021

Text:
ETP4HPC

Contact:

contact@etp4hpc.eu

www.etp4hpc.eu

@ETP4HPC

@etp4hpc

This handbook was produced with the support of the HPC-GIG project. This project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 824151.

Dear Reader,

Welcome to the seventh edition of the ETP4HPC Handbook of European HPC projects! This 2021 issue is bigger than ever, listing 68 on-going projects (and 34 more finalised projects are referenced). The EuroHPC-funded projects launched earlier this year are now covered. But we are not limiting ourselves to EuroHPC, or to the former HPC cPPP projects: this Handbook aims to provide a comprehensive overview of all projects in our arena. And our arena tends to expand, as HPC no longer lives in isolation but interacts with many related domains such as Big Data, Artificial Intelligence, Internet of Things and other technologies.

The projects featured in this Handbook are funded by 25 different calls. For ease of use, we have grouped them by scope, rather than by call. However, we also provide in *European Exascale Projects* an overview of the different calls, indicating where you will find the corresponding projects in our Handbook. And if you are looking for a project by name, the alphabetical index of projects at the end of the Handbook will point you directly to the right page.

I trust this Handbook will help you navigate the complex European HPC landscape. Our dearest wish is that this document will give you ideas and contacts for many future innovative projects that will nourish and strengthen our European HPC ecosystem and technological sovereignty!

Best regards,

ETP4HPC Chairman
Jean-Pierre Panziera

Table of Contents

European Exascale Projects

Our Handbook features a number of projects, which, though not financed by HPC cPPP or Euro-HPC calls, are closely related to HPC and have their place here. For ease of use, we have grouped projects by scope, instead of by call. However, we also provide in the table below an overview of the different calls, indicating where you will find the corresponding projects in our Handbook.

If you are looking for a project by name, the alphabetical index of projects at the end of the Handbook will point you directly to the right page.

Finally, this Handbook only features on-going projects. Projects that ended before January 2021 are listed in the Annex, and described in detail in previous editions of the Handbook (available on our website at https://www.etp4hpc.eu/european-hpc-handbook.html).

Chapter	Corresponding calls	Comments
Technology: HW and SW building blocks for exascale	EuroHPC-01-2019 Extreme scale computing and data driven technologies	10 projects started in 2021, most of them with a duration of 3 years
	EuroHPC-06-2019 Advanced Experimental Platform Towards Exascale	1 project started in 2020
	FETFLAG-05-2020 Complementary call on Quantum Computing	1 project started in 2020
	FETHPC-01-2016	2 projects started in 2017 and ending at the beginning of 2021
	FETHPC-02-2017 Transition to Exascale	11 projects started in 2018 with a duration of around 3 years
	ICT-05-2017 Customised and low energy computing (including Low power processor technologies)	Prequel to EPI
	ICT-16-2018 Software Technologies	1 project related to HPC
	ICT-42-2017 Framework Partnership Agreement in European low-power microprocessor technologies SGA-LPMT-01-2018 Specific Grant Agreement under the Framework Partnership Agreement “EPI FPA”	The European Processor Initiative (EPI)

Chapter	Corresponding calls	Comments
Centres of Excellence in computing applications	E-INFRA-5-2015 Centres of Excellence for computing applications	1 remaining project ending at the beginning of 2021
	INFRAEDI-02-2018 Centres of Excellence for computing applications	9 projects and 1 Coordination and Support Action started in 2018
	INFRAEDI-05-2020 Centres of Excellence in exascale computing	5 projects started in 2020 and 2021 with a duration of around 3 years
Applications in the HPC Continuum	CEF-TC-2018-5	1 project related to HPC
	EuroHPC-02-2019 HPC and data centric environments and application platforms	5 projects started in 2021
	EuroHPC-03-2019 Industrial software codes for extreme scale computing environments and applications	5 projects started in 2021
	FETHPC-01-2018 International Cooperation on HPC	2 projects started in June 2019
	ICT-11-2018-2019 Subtopic (a) HPC and Big Data enabled Large-scale Test-beds and Applications	4 Innovative Actions started in 2018
	ICT-12-2018-2020 Big Data technologies and extreme-scale analytics	1 project related to HPC
	SC1-PHE-CORONAVIRUS-2020 Advancing knowledge for the clinical and public health response to the 2019-nCoV epidemic	Special call
Ecosystem development	FETHPC-2-2014 HPC Ecosystem Development	2 Coordination and Support Actions
	FETHPC-03-2017 Exascale HPC ecosystem development	2 Coordination and Support Actions
	INFRAIA-01-2016-2017 Integrating Activities for Advanced Communities	HPC-EUROPA3
	EuroHPC-04-2019 Innovating and Widening the HPC use and skills base	2 projects starting in 2020
	EuroHPC-05-2019 Innovating and Widening the HPC use and skills base	1 project starting in 2020
	ICT-01-2019 Computing technologies and engineering methods for cyber-physical systems of systems	1 project related to HPC

Technology: HW and SW building blocks for exascale

ADMIRE

Adaptive multi-tier intelligent data manager for Exascale

malleable data solutions for HPC

The growing need to process extremely large data sets is one of the main drivers for building exascale HPC systems today. However, the flat storage hierarchies found in classic HPC architectures no longer satisfy the performance requirements of data-processing applications. Uncoordinated file access in combination with limited bandwidth make the centralised back-end parallel file system a serious bottleneck. At the same time, emerging multi-tier storage hierarchies come with the potential to remove this barrier. But maximising performance still requires careful control to avoid congestion and balance computational with storage performance. Unfortunately, appropriate interfaces and policies for managing such an enhanced I/O stack are still lacking.

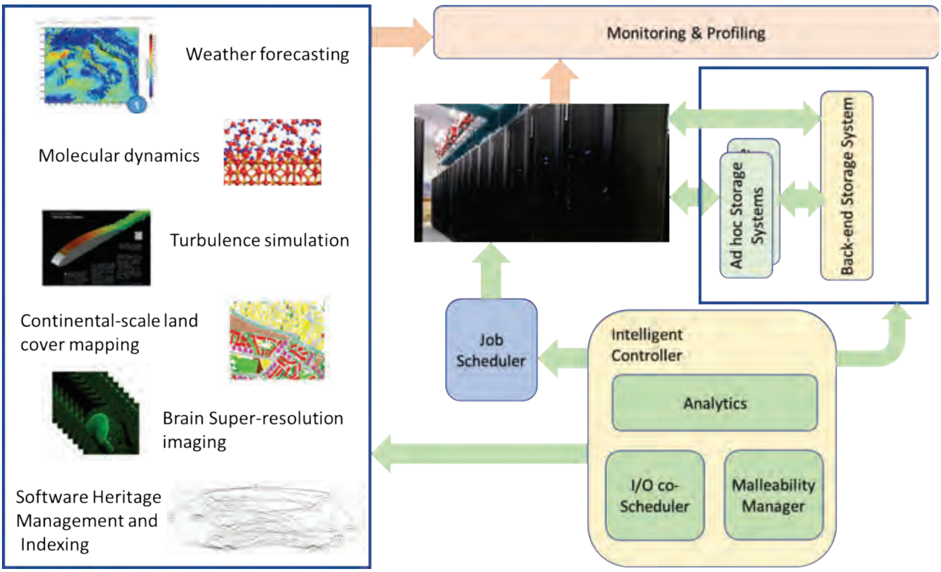

The main objective of the ADMIRE project is to establish this control by creating an active I/O stack that dynamically adjusts computation and storage requirements through intelligent global coordination, malleability of computation and I/O, and the scheduling of storage resources along all levels of the storage hierarchy. To achieve this, we will develop a software-defined framework based on the principles of scalable monitoring and control, separated control and data paths, and the orchestration of key system components and applications through embedded control points.

Our software-only solution will allow the throughput of HPC systems and the performance of individual applications to be substantially increased – and consequently energy consumption to be decreased – by taking advantage of fast and power-efficient node-local storage tiers using novel, European ad-hoc storage systems and in-transit/in-situ processing facilities. Furthermore, our enhanced I/O stack will offer quality-of-service (QoS) and resilience. An integrated and operational prototype will be validated with several use cases from various domains, including climate/weather, life sciences, physics, remote sensing, and deep learning.

Extreme scale computing and data driven technologies

www.admire-eurohpc.eu
Twitter: @admire_eurohpc

COORDINATING ORGANISATION

Universidad Carlos III de Madrid, Spain

OTHER PARTNERS

- Johannes Gutenberg-Universität Mainz, Germany
- Barcelona Supercomputing Centre (BSC), Spain
- Technische Universität Darmstadt, Germany
- DataDirect Networks (DDN), France
- Institut National de Recherche en Informatique et Automatique (INRIA), France
- ParaTools, SAS, France
- Forschungszentrum Jülich, Germany
- Consorzio Interuniversitario Nazionale per l'Informatica (CINI), Italy
- CINECA, Consorzio Interuniversitario, Italy
- E4 Computer Engineering SpA, Italy
- Instytut Chemii Bioorganicznej Polskiej Akademii Nauk, Poland
- Kungliga Tekniska Högskolan (KTH), Sweden

CONTACT

Jesús Carretero, jesus.carretero@uc3m.es

CALL

EuroHPC-01-2019

PROJECT TIMESPAN

01/04/2021- 31/03/2024

ASPIDE

exAScale ProgrammIng models for extreme Data processing

OBJECTIVES

The ASPIDE project will contribute with the definition of a new programming paradigms, APIs, runtime tools and methodologies for expressing data-intensive tasks on Exascale systems, which can pave the way for the exploitation of massive parallelism over a simplified model of the system architecture, promoting high performance and efficiency, and offering powerful operations and mechanisms for processing extreme data sources at high speed and/or real-time.

DCEX PROGRAMMING MODEL

The design of parallel patterns for programming Exascale applications has been one of the main goals of the ASPIDE project. DCEx addresses various features to enhance performance of data intensive computations in Exascale systems. Indeed, the cost of accessing, moving, and processing data across a parallel system can be enormous. The workflow modelling in the DCEx enables set-up a data life cycle management, allowing data locality and data affinity. The elementary workflow units permit to either place the data close to computational node where data is processed (data locality) or to distribute the computation where data was previously generated (data affinity avoids data movements). This way, the proposed solution assists application developers to access and use resources without the need to manage low-level architectural entities. In the same way, we want to provide a way to easily switch among different execution modes or policies without requiring to modify the applications source code.

AD-HOC IN-MEMORY STORAGE SYSTEM (IMSS)

IMSS is a proposal to enhance I/O in both traditional HPC and High-Performance Data Analytics (HPDA) systems. The architectural design follows a client-server design model where the client itself will be responsible of the server entities deployment. We propose an application-attached deployment constrained to application's nodes and an application-detached considering offshore nodes. The client layer is in charge of dealing with data locality exploitation alongside the implementation of multiple I/O patterns providing diverse data distribution policies.

INTELLIGENT ANOMALIES DETECTION SYSTEM

ASPIDE proposes an engine for data analysis for events detection in monitoring data for the purposes of application autotuning. Overall, we developed an anomalies detection approach based on machine learning, capable of detecting anomalies during Exascale application execution, such as hardware failures and communication bottlenecks. We utilise the events and anomalies detection engine to constrain the search space of the optimization problem, thus further improving the execution efficiency of the Exascale applications.

AUTO-TUNING

We introduce the ASPIDE auto-tuning approach based on a multi-objective optimization algorithm that considers multi-dimensional search space with pluggable objectives, including execution time and energy. More- over, to further improve the application execution, the ASPIDE approach utilizes a machine learning (ML) based events detection approach, capable of identifying point and contextual anomalies. In general, the ASPIDE auto-tuner assists developers in understanding the non-functional properties of their applications by making it easy to analyse and experiment with the input parameters. The auto-tuner further supports them in exposing their obtained insights using tunable parameters.

LARGE SCALE MONITORING

One important challenge in ExaScale computing consists of developing scalable components that are able to monitor in a coordinated and efficient manner the use of the hardware resources and the behaviour of the applications. ASPIDE provides a hierarchical monitor system based on well-known components based on data aggregators, machine learning techniques, and time series analysis that aim to reduce the overhead of sensing large scale infrastructures.

Transition to Exascale Computing

www.aspide-project.eu

Twitter: @ASPIDE_PROJECT

COORDINATING ORGANISATION

Universidad Carlos III de Madrid, Spain

OTHER PARTNERS

- Institute e-Austria Timisoara, Romania
- University of Calabria, Italy
- Universität Klagenfurt, Austria
- Institute of Bioorganic Chemistry of Polish Academy of Sciences, Poland
- Servicio Madrileño de Salud, Spain
- INTEGRIS SA Italy
- Atos (Bull SAS), France

CONTACT

Javier GARCIA-BLAS, fjblas@inf.uc3m.es

CALL

FETHPC-02-2017

PROJECT TIMESPAN

15/06/2018 - 14/06/2021

DCoMEX

Data Driven Computational Mechanics at Exascale

DCoMEX aims to provide unprecedented advances to the field of Computational Mechanics by developing novel numerical methods enhanced by Artificial Intelligence, along with a scalable software framework that enables exascale computing.

A key innovation of our project is the development of AISolve, a novel scalable library of AI-enhanced algorithms for the solution of large scale sparse linear system that are the core of computational mechanics.

Our methods fuse physics-constrained machine learning with efficient block-iterative methods and incorporate experimental data at multiple levels of fidelity to quantify model uncertainties. Efficient deployment of these methods in exascale supercomputers will provide scientists and engineers with unprecedented capabilities for predictive simulations of mechanical systems in applications ranging from bioengineering to manufacturing.

DCoMEX exploits the computational power of modern exascale architectures, to provide a robust and user-friendly framework that can be adopted in many applications. This framework is comprised of AI-Solve library integrated in two complementary computational mechanics HPC libraries. The first is a general-purpose multiphysics engine and the second a Bayesian uncertainty quantification and optimisation platform.

We will demonstrate DCoMEX potential by detailed simulations in two case studies:

- patient-specific optimization of cancer immunotherapy treatment, and
- design of advanced composite materials and structures at multiple scales.

We envision that software and methods developed in this project can be further customized and also facilitate developments in critical European industrial sectors like medicine, infrastructure, materials, automotive and aeronautics design.

Extreme scale computing and data driven technologies

mgroup.ntua.gr/dcomex

COORDINATING ORGANISATION

National Technical University of Athens (NTUA), Greece

OTHER PARTNERS

- Eidgenössische Technische Hochschule Zürich, Switzerland
- University of Cyprus, Cyprus
- Technische Universität München, Germany
- National Infrastructures for Research and Technology, Greece

CONTACT

mgroup.ntua.gr/dcomex/

CALL

EuroHPC-01-2019

PROJECT TIMESPAN

01/04/2021- 31/03/2024

DEEP-EST

DEEP PROJECTS
Dynamical Exascale Entry Platform

There is more than one way to build a supercomputer, and meeting the diverse demands of modern applications, which increasingly combine data analytics and artificial intelligence (AI) with simulation, requires a flexible system architecture. Since 2011, the DEEP series of projects (DEEP, DEEP-ER, DEEP-EST) has pioneered an innovative concept known as the Modular Supercomputer Architecture (MSA), whereby multiple modules are coupled like building blocks. Each module is tailored to the needs of a specific class of applications, and all modules together behave as a single machine.

Connected by a high-speed, federated network and programmed in a uniform system software and programming environment, the supercomputer allows an application to be distributed over several hardware modules, running each code component on the one which best suits its particular needs. Specifically, DEEP-EST, finished in March 2021, has built a prototype with three modules: a general-purpose Cluster Module (CM) for low or medium scalable codes, the highly scalable Extreme Booster Module (ESB) comprising a cluster of accelerators, and a Data Analytics Module (DAM), which were tested with six applications combining high-performance computing (HPC) with high-performance data analytics (HPDA) and machine learning (ML).

The DEEP approach is part of the trend towards using accelerators to improve performance and overall energy efficiency – but with a twist. Traditionally, heterogeneity is done within the node, combining a central processing unit (CPU) with one or more accelerators. In DEEP-EST the resources were segregated and pooled into compute modules, as this enables to flexibly adapt the system to very diverse application requirements. In addition to usability and flexibility, the sustained performance made possible by following this approach aims to reach exascale levels.

One important aspect that makes the DEEP architecture stand out is the co-design approach, which is a key component of the project. In DEEP-EST, six ambitious HPC/HPDA applications were used to define and evaluate the hardware and software technologies developed. Careful analysis of the application codes allowed a fuller understanding of their requirements, which informed the prototype's design and configuration.

In addition to traditional compute-intensive HPC applications, the DEEP-EST DAM includes leading-edge memory and storage technology tailored to the needs of data-intensive workloads, which occur in data analytics and ML.

Through the DEEP projects, researchers have shown that combining resources in compute modules efficiently serves applications from multi-physics simulations to simulations integrating HPC with HPDA, to complex heterogeneous workflows such as those in artificial intelligence applications.

The next step in the series of DEEP projects is made with DEEP-SEA – Software for Exascale Architectures.

Co-design of HPC systems and applications

www.deep-projects.eu
Twitter: @DEEPprojects

COORDINATING ORGANISATION

Forschungszentrum Jülich (Jülich Supercomputing Centre), Germany

OTHER PARTNERS

Current partners in DEEP-EST:

- Intel, Germany
- Leibniz Supercomputing Centre, Germany
- Barcelona Supercomputing Centre (BSC), Spain
- Megware Computer Vertrieb und Service GmbH, Germany
- Heidelberg University, Germany
- EXTOLL, Germany
- The University of Edinburgh, United Kingdom
- Fraunhofer ITWM, Germany
- Astron, Netherlands
- KU Leuven, Belgium
- National Center For Supercomputing Applications, Bulgaria
- Norges Miljo-Og Biovitenskaplige, Norway
- Universitet Haskoli Islands, Iceland
- European Organisation for Nuclear Research (CERN), Switzerland
- ParTec, Germany

Partners in DEEP and DEEP-ER:

- CERFACS, France
- CGG, France
- CINECA, Italy
- Eurotech, Italy
- EPFL, Switzerland
- Seagate, UK
- INRIA, Institut National de Recherche en Informatique et Automatique, France
- The Cyprus Institute, Cyprus
- Universität Regensburg, Germany

CONTACT

Dr. Estela Suarez,
pmt@deep-projects.eu

CALL

ETHPC-01-2016

PROJECT TIMESPAN (DEEP-EST)

01/07/2017 - 31/03/2021

DEEP-SEA

DEEP PROJECTS
Dynamical Exascale Entry Platform

DEEP projects go into the next round: After DEEP-EST (Extreme Scale Technologies) was finished successfully in March 2021, DEEP-SEA (Software for Exascale Architectures) started on April 1, 2021. The 4th member of the DEEP project series focusses on the question how future exascale systems can be used in an easy – or partly automated – manner, while at the same time being as energy efficient as possible.

With DEEP-EST, a prototype was established that leverages the benefits of a Modular Supercomputing Architecture (MSA). Here, different components like standard CPUs, accelerators like GPUs, among others, build a complex mesh of heterogeneous technologies. Each technology is best suited for specific workloads, applications can leverage the parts of the MSA that allow for optimal energy/performance ratio. Large-scale simulations, data analytics, machine- and deep learning – different jobs have different requirements to run at the equilibrium of energy usage and compute performance.

The downside: the complexity on system and node level makes it hard to allocate all resources in the best way. Parallelization is not trivial in traditional HPC environment – with the rising complexity of heterogeneous resources is will be hardly manageable by the users and the application developers. They need to understand their codes and the underlying hardware very well in order to decide on what components parts of the code should run. Also, the need grows to port codes to different technologies such as accelerators or even to completely new platforms. Not to mention the burden of optimization for a variety of HPC systems.

DEEP-SEA addresses this challenge. The main goals of this project are:

- Better manage and program compute and memory heterogeneity.
- Targets easier programming for Modular Supercomputers.

To achieve these goals, DEEP-SEA looks at all relevant parts of the system: nodes, clusters, applications, system software and tools.

To provide solutions that are usable on as many HPC systems as possible and that scale to exascale and beyond, DEEP-SEA uses the co-design approach that has already proven effective in earlier DEEP projects. Real-world applications from seven areas form an integral part of DEEP-SEA to prove the practical use of the chosen approach:

- Space Weather (xPic, AIDApy)
- Weather Forecast (IFS)
- Seismic imaging (RTM, BSIT)
- Molecular dynamics (GROMACS)
- Computational fluid dynamics (Nek5000)
- Neutron Monte Carlo transport for nuclear energy (PATMOS)
- Earth System Modelling (TSMP)

To support the co-design approach, an early-access programme will be established as soon as the first relevant milestones are reached.

DEEP-SEA builds upon exascale concepts developed over nearly 10 years. It will build a software stack for heterogeneous compute and memory systems that allows scientists and developers to make best use of all available resources.

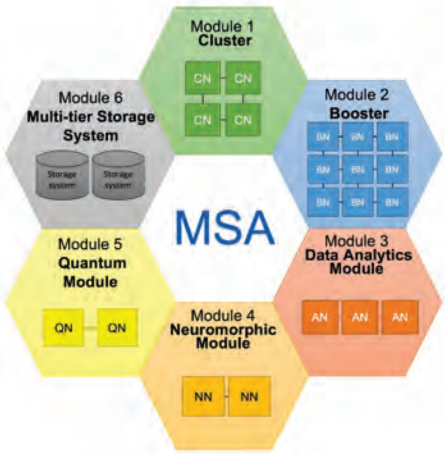

Extreme scale computing and data driven technologies

www.deep-projects.eu
Twitter: @DEEPprojects

COORDINATING ORGANISATION

Forschungszentrum Jülich (Jülich Supercomputing Centre), Germany

OTHER PARTNERS

- Atos (Bull SAS), France
- Leibniz Supercomputing Centre , Germany
- Barcelona Supercomputing Centre (BSC), Spain
- Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V, Germany
- Katholieke Universiteit (KU) Leuven, Belgium
- Eidgenössische Technische Hochschule (ETH) Zürich, Switzerland
- ParTec, Germany
- Idryma Technologias Kai Erevnas (FORTH), Greece
- Commissariat à l'énergie atomique et aux énergies alternatives (CEA), France
- Technische Universität München, Germany
- Technische Universität Darmstadt, Germany
- Kungliga Tekniska Högskolan (KTH), Sweden
- European Centre for Medium-Range Weather Forecasts (ECMWF), United Kingdom

CONTACT

Dr. Estela Suarez,
pmt@deep-projects.eu

CALL

EuroHPC-01-2019

PROJECT TIMESPAN

01/04/2021 - 31/03/2024

EPEEC

European joint Effort toward a Highly Productive Programming Environment for Heterogeneous Exascale Computing

EPEEC's main goal is to develop and deploy a production-ready parallel programming environment that turns upcoming overwhelmingly-heterogeneous exascale supercomputers into manageable platforms for domain application developers. The consortium has significantly advanced and integrated existing state-of-the-art components based on European technology (programming models, runtime systems, and tools) with key features enabling three overarching objectives: high coding productivity, high performance, and energy awareness.

An automatic generator of compiler directives provides outstanding coding productivity from the very beginning of the application developing/porting process. Developers are able to leverage either shared memory or distributed-shared memory programming flavours, and code in their preferred language: C, Fortran, or C++. EPEEC ensures the composability and interoperability of its programming models and runtimes, which incorporate specific features to handle data-intensive and extreme-data applications. Enhanced leading-edge performance tools offer integral profiling, performance prediction, and visualisation of traces.

EPEEC PROGRAMMING ENVIRONMENT

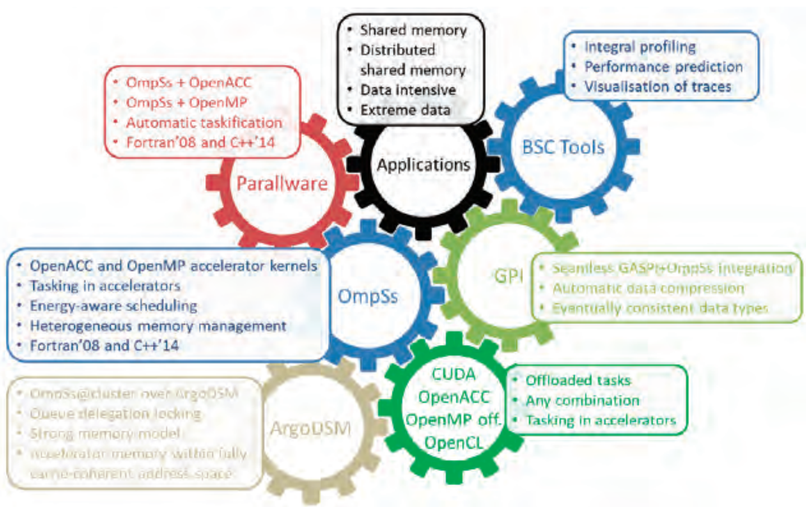

EPEEC exploits results from past FET projects that led to the cutting-edge software components it builds upon, and pursues influencing the most relevant parallel programming standardisation bodies. The consortium is composed of European institutions and individuals with the highest expertise in their field, including not only leading research centres and universities but also SME/start-up companies, all of them recognised as high-tech innovators worldwide. Adhering to the Work Programme's guidelines, EPEEC features the participation of young and high-potential researchers, and includes careful dissemination, exploitation, and public engagement plans.

APPLICATIONS

Five applications representative of different relevant scientific domains serve as part of a strong inter-disciplinary co-design approach and as technology demonstrators:

- AVBP
- DIOGENeS
- OSIRIS
- Quantum ESPRESSO
- SMURFF

PROGRAMMING ENVIRONMENT

The EPEEC programming environment for exascale is developed to achieve high programming productivity, high execution efficiency and scalability, energy awareness, and smooth composability/interoperability. It is formed by five main components:

- Parallelware
- OmpSs
- GASPI
- ArgoDSM
- BSC performance tools (Extrae, Paraver, Dimemas)

Software prototypes and releases are available for download on the EPEEC website.

Transition to Exascale Computing

www.epeec-project.eu
Twitter: #EPEECproject

COORDINATING ORGANISATION
Barcelona Supercomputing Centre (BSC), Spain

OTHER PARTNERS
- Fraunhofer Gesellschaft zur Förderung der angewandten Forschung E.V., Germany
- INESC ID - Intituto de Engenhariade Sistemas e Computadores, Investigacao e Desenvolvimento em Lisboa, Portugal
- INRIA, Institut National de Recherche en Informatique et Automatique, France
- Appentra Solutions SL, Spain
- CINECA Consorzio Interuniversitario, Italy
- ETA SCALE AB, Sweden
- CERFACS, France
- Interuniversitair Micro-Electronica Centrum, Belgium
- Uppsala Universitet, Sweden

CONTACT
Antonio J. Peña, antonio.pena@bsc.es

CALL
FETHPC-02-2017

PROJECT TIMESPAN
01/10/2018 - 31/03/2022

EPI

European Processor Initiative

Europe has an ambitious plan to become a main player in supercomputing. One of the core components for achieving that goal is a processor. The European Processor Initiative (EPI) is a part of a broader strategy to develop and independent European HPC industry based on domestic and innovative technologies as presented in the EuroHPC Joint Undertaking proposed by the European Commission.

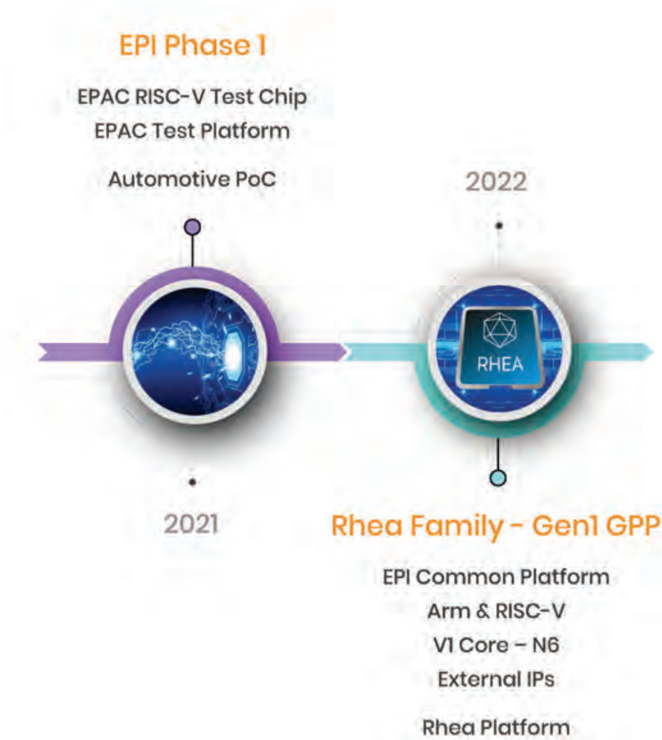

The EPI timeline

The general objective of EPI is to design a roadmap and develop the key Intellectual Property for future European low-power processors addressing extreme scale computing (exascale), high-performance, big-data and emerging verticals (e.g. automotive computing) and other fields that require highly efficient computing infrastructure. More precisely, EPI aims at establishing the key technologies to reach three fundamental goals:

1. Developing low-power processor technology to be included in advanced experimental pilot platforms towards exascale systems for Europe;
2. Ensuring that a significant part of that technology and intellectual property is European;
3. Ensuring that the application areas of the technology are not limited only to HPC, but cover other areas, such as automotive and data centres, thus ensuring the economic viability of the initiative.

EPI gathers 28 partners from 10 European countries, with a wide range of expertise and background: HPC, supercomputing centres, automotive computing, including researchers and key industrial players. The fact that the envisioned European processor is planned to be based on already existing tools either owned by the partners or being offered as open-source with a large community of users,provides three key exploitation advantages: (1) the time-to-market will be reduced as most of these tools are already used in industry and well known. (2) It will enable

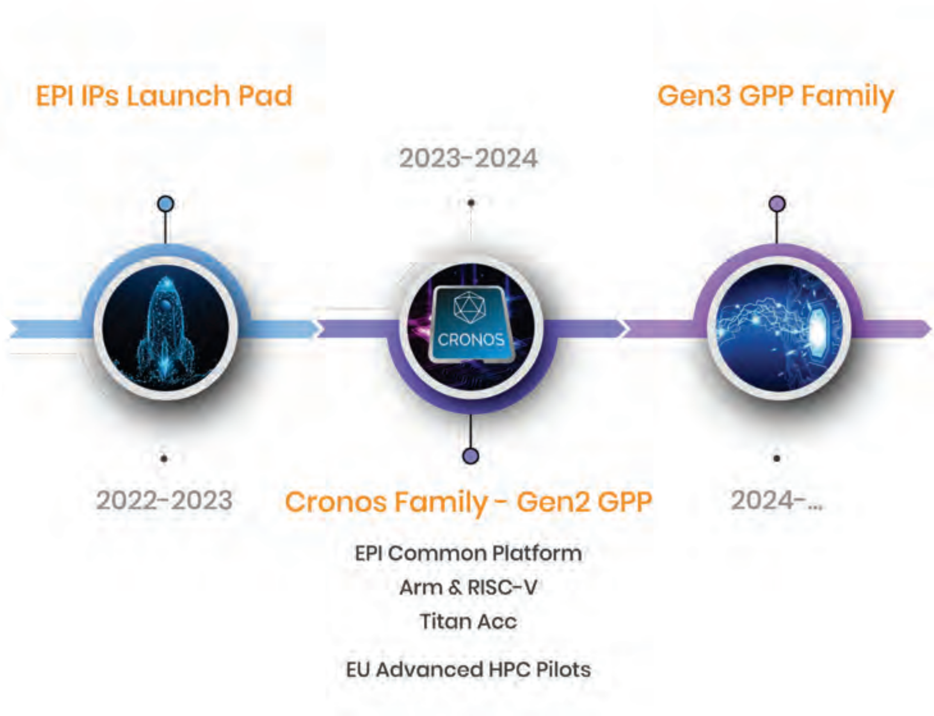

EPI partners to incorporate the results in their commercial portfolio or in their scientific roadmap. (3) it fundamentally reduces the technological risk associated to advanced technologies developed and implemented in EPI.

EPI covers a complete range of expertise, skills, and competencies needed to design and execute a sustainable roadmap for research and innovation in processor and computing technology, fostering future exascale HPC and emerging applications, including Big Data, and automotive computing for autonomous vehicles. Development of a new processor to be at the core of future computing systems will be divided into several streams:

- Common Platform and global architecture [stream 1]
- HPC general purpose processor [stream 2]
- Accelerator [stream 3]
- Automotive platform [stream 4]

The results from the two streams related to the general-purpose processor and accelerator chips will generate a heterogeneous, energy-efficient CPU for use in both standard and non-traditional and compute-intensive segments, e.g. automotive where SoA in autonomous driving requires significant computational resources.

EPI strives to maximize the synergies between the two streams and will work with existing EU initiatives on technology, infrastructure and applications, to position Europe as a world leader in HPC and emerging markets for exascale era such as automotive computing for autonomous driving.

European Processor

european-processor-initiative.eu

Twitter: @EUProcessor
LinkedIn: @european-processor-initiative
YouTube: www.youtube.com

COORDINATING ORGANISATION

Atos (Bull SAS), France

OTHER PARTNERS

- BSC, Barcelona supercomputing center – centro nacional de supercomputación, Spain
- IFAG, Infineon Technologies AG, Germany
- SemiDynamics, SemiDynamics technology services SL, Spain
- CEA, Commissariat à l'Energie Atomique et aux énergies alternatives, France
- CHALMERS, Chalmers tekniska högskola, Sweden
- ETH Zürich, Eidgenössische Technische Hochschule Zürich, Switzerland
- FORTH, Foundation for Research and Technology Hellas, Greece
- GENCI, Grand Equipement National de Calcul Intensif, France
- IST, Instituto Superior Technico, Portugal
- JÜLICH, Forschungszentrum Jülich GmbH, Germany
- UNIBO, Alma mater studiorum - Universita di Bologna, Italy
- UNIZG-FER, Sveučilište u Zagrebu, Fakultet elektrotehnike i računarstva, Croatia
- Fraunhofer, Fraunhofer Gesellschaaft zur Förderung der angewandten Forschung E.V., Germany
- ST-I, STMICROELECTRONICS SRL, Italy
- E4, E4 Computer Engineering SPA, Italy
- UNIPI, Universita di Pisa, Italy
- SURFsara BV, Netherlands
- Kalray SA, France
- Extoll GmbH, Germany
- CINECA, Cineca Consorzio Interuniversitario, Italy
- BMW Group, Bayerische Motoren Werke Aktiengesellschaft, Germany
- Elektrobit, Automotive GmbH, Germany
- KIT, Karlsruher Institut für Technologie, Germany
- Menta SAS, France
- Prove & Run, France
- SIPEARL SAS, France
- Kernkonzept, Germany

CONTACT

contact@european-processor-initiative.eu

CALL

SGA-LPMT-01-2018

PROJECT TIMESPAN

01/12/2018 - 30/11/2021

EPiGRAM-HS

Exascale Programming Models for Heterogeneous Systems

MOTIVATION

- Increasing presence of heterogeneous technologies on pre-exascale supercomputers
- Need to port key HPC and emerging applications to these systems on time for exascale

OBJECTIVES

- Extend the programmability of large-scale heterogeneous systems with GPUs, FPGAs, HBM and NVM
- Introduce new concepts and functionalities, and implement them in two widely-used HPC programming systems for large-scale supercomputers: MPI and GASPI
- Maximize the productivity of application development on heterogeneous supercomputers by:
  - providing auto-tuned collective communication
  - a framework for automatic code generation for FPGAs
  - a memory abstraction device comprised of APIs
  - a runtime for automatic data placement on diverse memories and a DSL for large-scale deep-learning frameworks

Transition to Exascale Computing

epigram-hs.eu

Twitter: @epigramhs

COORDINATING ORGANISATION
KTH, Sweden

OTHER PARTNERS
- EPCC, United Kingdom
- ETHZ - Eidgenössische Technische Hochschule Zürich , Switzerland
- Fraunhofer, Germany
- Cray, Switzerland
- ECMWF, United Kingdom

CONTACT
Stefano Markidis, markidis@kth.se

CALL
FETHPC-02-2017

PROJECT TIMESPAN
01/09/2018 - 31/08/2021

eProcessor

European, extendable, energy-efficient, energetic, embedded, extensible, Processor Ecosystem

The eProcessor project is an ambitious combination of processor design, based on the RISC-V open source hardware ISA, applications and system software, bringing together multiple partners to leverage and extend pre-existing Intellectual Property (IP), combined with new IP that can be used as building blocks for future HPC systems, both for traditional and emerging application domains. As such, the **eProcessor project's overall goal is to create an open source full stack ecosystem (both software and hardware)** by achieving the following objectives:

1. **Extend open source to include open source hardware for HPC:** The eProcessor technology will be based on the RISC-V open source instruction set architecture (ISA) and will feature high performance computing and data analytics accelerators coupled to a high performance, low energy out-of-order processor (Europe's first high performance out-of-order 64-bit RISC-V platform). This is a major first step in the direction of an open European software/hardware ecosystem, which will guarantee technology independence.
2. **Software/hardware co-design for improved application performance and system energy efficiency:** eProcessor will meet the performance and energy requirements of new and existing HPC applications. eProcessor will co-design solutions to provide high performance, low-power, and fault tolerance. Uniquely, the

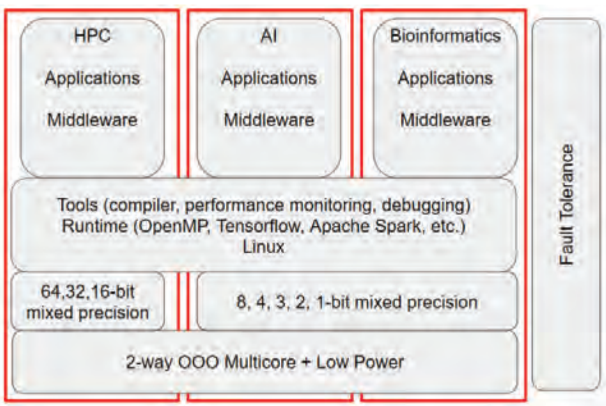

project will specialize all components of the system in the context of a broad application domain: a combination of energy efficient accelerators, adaptive on-chip memory structures, and a flexible and high performance energy-efficient CPU, with the corresponding open source software stack.
3. **HPC and HPDA applications:** eProcessor will use a diverse set of applications in HPC and high performance data analytics (HPDA, which includes Artificial Intelligence (AI), Deep Learning (DL), Machine Learning (ML) and Bioinformatics applications to drive the design of the overall system. eProcessor will extend these applications and their frameworks to support the RISC-V ISA.
4. **Focus on sustained application performance:** Many HPC and HPDA applications use sparse data sets and/or low/mixed-precision. Instead of focusing on the peak performance of dense computations, eProcessor targets a broader collection of applications by developing a system targeting sustained application performance.
5. **Stimulate European collaboration:** The eProcessor partners will leverage their existing IP from multiple European projects such as European Processor Initiative (EPI), Low-Energy Toolset for Heterogeneous Computing (LEGaTO), MareNostrum Experimental Exascale Platform (MEEP), POP2 CoE, Tulipp, EuroEXA, ExaNeSt and DeepHealth and extend their capabilities and improve the Technical Readiness Level (TRL).

By employing software/hardware co-design, as shown in the figure on the previous page, eProcessor covers the whole system stack from HPC and HPDA applications at the top, down to the processor core with accelerators at the bottom while all levels are mutable so as to enable true, unrestricted co-design.

Extreme scale computing and data driven technologies

eprocessor.eu
Twitter: @eprocessor_eu
LinkedIn: @eprocessor

COORDINATING ORGANISATION
Barcelona Supercomputing Center (BSC)

OTHER PARTNERS
- Chalmers tekniska högskola, Sweden
- Idryma Technologias Kai Erevnas, Foundation for Research and Technology, Greece
- Università degli Studi di Roma "La Sapienza", Italy
- Cortus, France
- christmann informationstechnik + medien GmbH & Co. KG, Germany
- Universität Bielefeld, Germany
- EXTOLL GmbH, Germany
- Thales SA, France
- Exascale Performance Systems – EXAPSYS P.C., Greece

CONTACT
Nehir Sonmez, nehir.sonmez@bsc.es

CALL
EuroHPC-01-2019

PROJECT TIMESPAN
01/04/2021 - 31/03/2024

ESCAPE-2

Energy-efficient SCalable Algorithms for weather and climate Prediction at Exascale

OBJECTIVE

ESCAPE-2 will develop world-class, extreme-scale computing capabilities for European operational numerical weather and climate prediction, and provide the key components for weather and climate domain benchmarks to be deployed on extreme-scale demonstrators and beyond. This will be achieved by developing bespoke and novel mathematical and algorithmic concepts, combining them with proven methods, and thereby reassessing the mathematical foundations forming the basis of Earth system models. ESCAPE-2 also invests in significantly more productive programming models for the weather-climate community through which novel algorithm development will be accelerated and future-proofed. Eventually, the project aims at providing exascale-ready production benchmarks to be operated on extreme-scale demonstrators (EsD) and beyond. ESCAPE-2 combines cross-disciplinary uncertainty quantification tools (URANIE) for high-performance computing, originating from the energy sector, with ensemble-based weather and climate models to quantify the effect of model and data related uncertainties on forecasting – a capability, which weather and climate prediction has pioneered since the 1960s.

The mathematics and algorithmic research in ESCAPE-2 will focus on implementing data structures and tools supporting parallel computation of dynamics and physics on multiple scales and multiple levels. Highly-scalable spatial discretization will be combined with proven large time-stepping techniques to optimize both time-to-solution and energy-to-solution. Connecting multi-grid tools, iterative solvers, and overlapping computations with flexible-order spatial discretization will strengthen algorithm resilience against soft or hard failure. In addition, machine learning techniques will be applied for accelerating complex sub-components. The sum of these efforts will aim at achieving at the same time: performance, resilience, accuracy and portability.

Transition to Exascale Computing

www.hpc-escape2.eu/

COORDINATING ORGANISATION

ECMWF - European Centre for Medium-range Weather Forecasts, United Kingdom

OTHER PARTNERS

- DKRZ - Deutsches Klimarechenzentrum GmbH, Germany
- Max-Planck-Gesellschaft zur Förderung der Wissenschaften EV, Germany
- Eidgenössisches Depertment des Innern, Switzerland
- Barcelona Supercomputing Centre (BSC), Spain
- CEA - Commissariat à l'Energie Atomique et aux énergies alternatives, France
- Loughborough University, United Kingdom
- Institut Royal Météorologique de Belgique, Belgium
- Politecnico di Milano, Italy
- Danmarks Meteorologiske Institut, Denmark
- Fondazione Centro Euro-Mediterraneo sui Cambiamenti Climatici, Italy
- Atos (Bull SAS), France

CONTACT

Peter Bauer, peter.bauer@ecmwf.int

CALL

FETHPC-02-2017

PROJECT TIMESPAN

01/10/2018 - 30/09/2021

EuroExa

The Largest Group of ExaScale Projects and Biggest Co-Design Effort

EuroEXA is a co-designed project funded by EU Horizon 2020. We are working to develop technologies that meet the demands of ExaScale computing – delivering a ground-breaking architecture capable of scaling to peak performance of 400 PetaFLOPS, with a peak power system envelope of 30MW that approaches PUE parity through renewables and immersion-based cooling.

The project has a €20m investment over 52 months and is part of a total €50m investment made by the European Commission across the EuroEXA group of projects supporting research, innovation and action across applications, system software, hardware, networking, storage, liquid cooling, and data centre technologies.

Our project group was assembled to drive the development of our members' technologies, working together to co-design the future of EU-based ExaScale supercomputers. To do that, we are building three testbeds, showing new architectures and technologies which can be extrapolated to demonstrate:

- Compute Performance – up to 400 PetaFLOPS, with an estimated system power of around 30MW peak consumption.
- Energy Efficiency – 250 PetaFLOPS per MW for future iterations of semiconductor technologies within the architecture.
- Compactness – showing that an ExaScale machine could be built within 30 shipping containers, with an edge-to-edge distance of less than 40m for the computers.

Partners have ported a rich mix of key applications to the EuroExa FPGA platform – from forecasting the weather and modelling particle physics to simulating the human brain and searching for new drugs to treat disease.

We are working to build very dense PCBs, bringing the electronics as close together as possible to drive performance and reduce energy demands. Similarly, we have developed a three-level interconnect that gives data the shortest possible routes to travel, while allowing for the expansion of the system with commodity OpenFlow switches – offering scalability with fewer bottlenecks.

The project objectives include to develop and deploy Arm Cortex and AMD EPYC technology processing systems with FPGA acceleration at PetaFLOP level an ExaScale procurement for deployment in 2022/23.

The EuroEXA Blade

Co-design of HPC systems and applications

www.euroexa.eu

Twitter: @EuroEXA
LinkedIn: @euroEXA

COORDINATING ORGANISATION

ICCS (Institute of Communication and Computer Systems), Greece

OTHER PARTNERS

- Arm, UK
- The University of Manchester, UK
- Barcelona Supercomputing Centre (BSC), Spain
- FORTH (Foundation for Research and Technology Hellas), Greece
- The Hartree Centre of STFC, UK
- IMEC, Belgium
- ZeroPoint Technologies, Sweden
- Iceotope, UK
- Synelixis Solutions Ltd., Greece
- Maxeler Technologies, UK
- Neurasmus, Netherlands
- INFN (Istituto Nazionale Di Fisica Nucleare), Italy
- INAF (Istituto Nazionale Di Astrofisica, Italy
- ECMWF (European Centre for Medium-Range Weather Forecasts), International
- Fraunhofer, Germany

CONTACT

Peter Hopton (Dissemination & Exploitation Lead)
peter.hopton@iceotope.com

CALL

FETHPC-01-2016

PROJECT TIMESPAN

01/09/2017 - 31/12/2021

EXA2PRO

Enhancing Programmability and boosting Performance Portability for Exascale Computing Systems

Advancing scientific progress in domains such as physics, energy and material science by developing and deploying applications in computing clusters of hundreds of thousand CPU cores is exciting, yet very challenging. The European Commission-funded project EXA2PRO – Enhancing Programmability and boosting Performance Portability for Exascale Computing Systems - designs and develops a framework for enabling the efficient deployment of applications in supercomputing systems.

EXA2PRO project is developing programming models and tools in a co-design activity together with relevant applications from the high energy physics, material science, chemical processes for cleaner environment and energy storage domains. The EXA2PRO goal is to increase the productivity of developing and deploying applications on heterogeneous computing systems and to promote and lower the barrier of access to exascale computing systems to the scientific community and industry.

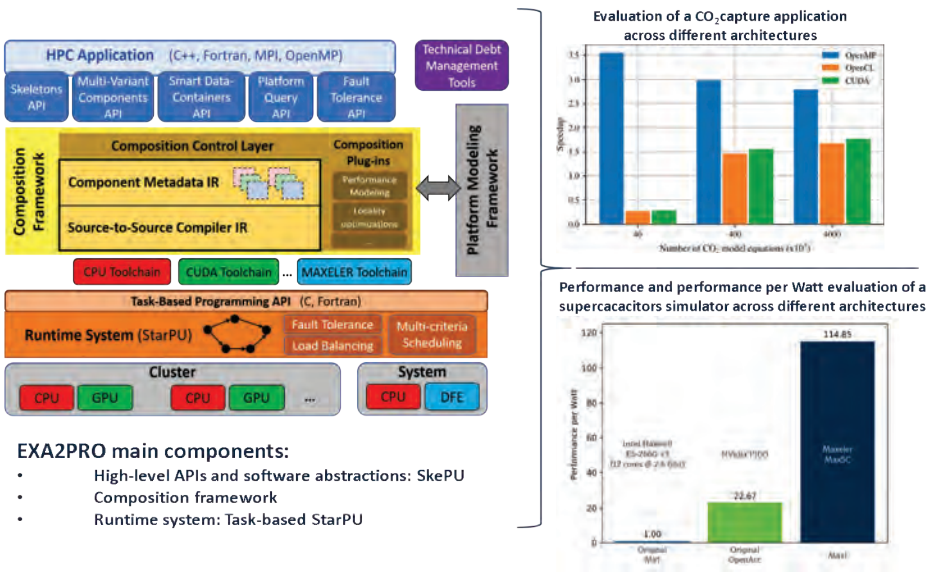

EXA2PRO framework and indicative results

As the project ends within 2021, tools of the EXAPRO framework have been evaluated in a wide variety of applications, generating impressive results. Some recent success stories are the following:

- The EXA2PRO tools have been used to advance CO2 capture technologies, by enabling the generation of CO2 capture solutions 41% faster.
- The performance of a supercapacitor simulation code (MetalWalls) improved by 33% by applying the EXA2PRO runtime system.
- Initial results of porting the supercapacitor simulation code on a dataflow engine accelerator, showed significant performance and energy gains.
- The EXA2PRO programming model has been applied to a multi-node neural simulation code, greatly improving its scalability.

EXA2PRO outcomes have major impact on:

1. the scientific and industrial community that focuses on application deployment in supercomputing centres: The EXA2PRO framework improves productivity. After applying the EXA2PRO API in applications, evaluation across a variety of architectures is performed automatically.
2. application developers that target exascale computing systems: EXA2PRO provides tools for improving source code maintainability/reusability, which will allow application evolution with reduced developers' effort.
3. the scientific community and the industry relevant to the EXA2PRO applications, with significant impact on the materials and processes design for CO2 capture and on the supercapacitors industry.

Transition to Exascale Computing

exa2pro.eu/

Twitter: @exa2pro_h2020

COORDINATING ORGANISATION

ICCS (Institute of Communication and Computer Systems), Greece

OTHER PARTNERS

- Linköpings Universitet, Sweden
- CERTH, Ethniko Kentro Erevnas kai Technologikis Anaptyxis, Greece
- INRIA, Institut National de Recherche en Informatique et Automatique, France
- Forschungszentrum Jülich GmbH, Germany
- Maxeler Technologies Ltd, United Kingdom
- CNRS - Centre National de la Recherche Scientifique, France
- UoM, University of Macedonia, Greece

CONTACT

Prof. Dimitrios Soudris (project coordinator), dsoudris@microlab.ntua.gr
Dr. Lazaros Papadopoulos (technical coordinator), lpapadop@microlab.ntua.gr

CALL

FETHPC-02-2017

PROJECT TIMESPAN

01/05/2018 - 31/07/2021

ExaQUte

ExaQUte - EXAscale Quantification of Uncertainties for Technology and Science Simulation

Exascale Quantification of Uncertainties for Technology and Science Simulation

OBJECTIVE

The ExaQUte project aims at constructing a framework to enable Uncertainty Quantification (UQ) and Optimization Under Uncertainties (OUU) in complex engineering problems using computational simulations on Exascale systems.

The stochastic problem of quantifying uncertainties will be tackled by using a Multi Level MonteCarlo (MLMC) approach that allows a high number of stochastic variables. New theoretical developments will be carried out to enable its combination with adaptive mesh refinement, considering both, octree-based and anisotropic mesh adaptation.
Gradient-based optimization techniques will be extended to consider uncertainties by developing methods to compute stochastic sensitivities. This requires new theoretical and computational developments. With a proper definition of risk measures and constraints, these methods allow high-performance robust designs, also maximizing the solution reliability.

The description of complex geometries will be possible by employing embedded methods, which guarantee a high robustness in the mesh generation and adaptation steps, while allowing preserving the exact geometry representation.

The efficient exploitation of Exascale system will be addressed by combining State-of-the-Art dynamic task-scheduling technologies with space-time accelerated solution methods, where parallelism is harvested both in space and time.

The methods and tools developed in ExaQUte will be applicable to many fields of science and technology. The chosen application focuses on wind engineering, a field of notable industrial interest for which currently no reliable solution exists. This will include the quantification of uncertainties in the response of civil engineering structures to the wind action, and the shape optimization taking into account uncertainties related to wind loading, structural shape and material behaviour.

All developments in ExaQUte will be open-source and will follow a modular approach, thus maximizing future impact.

Transition to Exascale Computing

www.exaqute.eu

Twitter: @ExaQUteEU

COORDINATING ORGANISATION

CIMNE- Centre Internacional de Metodes Numerics en Enginyeria, Spain

OTHER PARTNERS

- Barcelona Supercomputing Centre (BSC), Spain
- Technische Universität München, Germany
- INRIA, Institut National de Recherche en Informatique et Automatique, France
- IT4Innovations, VSB – Technical University of Ostrava, Czech Republic
- Ecole Polytechnique Fédérale de Lausanne, Switzerland
- Universitat Politecnica de Catalunya, Spain
- str.ucture GmbH, Germany

CONTACT

Cecilia Soriano, csoriano@cimne.upc.edu
Riccardo Rossi, rrossi@cimne.upc.edu

CALL

FETHPC-02-2017

PROJECT TIMESPAN

01/06/2018 - 31/05/2021

IO-SEA

IO Software for Exascale Architecture

IO-SEA aims to provide a novel data management and storage platform for exascale computing based on hierarchical storage management (HSM) and on-demand provisioning of storage services. The platform will efficiently make use of storage tiers spanning NVMe and NVRAM at the top all the way down to tapebased technologies. System requirements are driven by data intensive use-cases, in a very strict co-design approach. The concept of ephemeral data nodes and data accessors is introduced that allow users to flexibly operate the system, using various well-known data access paradigms, such as POSIX namespaces, S3/Swift Interfaces, MPI-IO and other data formats and protocols. These ephemeral resources eliminate the problem of treating storage resources as static and unchanging system components – which is not a tenable proposition for data intensive exascale environments. The methods and techniques are applicable to exascale class data intensive applications and workflows that need to be deployed in highly heterogeneous computing environments.

Critical aspects of intelligent data placement are considered for extreme volumes of data. This ensures that the right resources among the storage tiers are used and accessed by data nodes as close as possible to compute nodes – optimising performance, cost, and energy at extreme scale. Advanced IO instrumentation and monitoring features will be developed to that effect leveraging the latest advancements in AI and machine learning to systematically analyse the telemetry records to make smart decisions on data placement. These ideas coupled with in-storage-computation remove unnecessary data movements within the system.

The IO-SEA project (EuroHPC-2019-1 topic b) has connections to the DEEP-SEA (topic d) and RED-SEA (topic c) project. It leverages technologies developed by the SAGE, SAGE2 and NextGEN-IO projects, and strengthens the TLR of the developed products and technologies.

Extreme scale computing and data driv-en technologies

iosea-project.eu

Twitter: @iosea_eu

COORDINATING ORGANISATION

Commissariat à l'énergie atomique et aux énergies alternatives (CEA), France

OTHER PARTNERS

- Atos (Bull SAS), France
- Forschungszentrum Jülich, Germany
- European Centre for Medium-Range Weather Forecasts, United Kingdom
- Seagate Systems (UK) Limited, United Kingdom
- National University of Ireland Galway, Ireland
- IT4Innovations, VSB – Technical University of Ostrava, Czech Republic
- Kungliga Tekniska Högskolan (KTH), Sweden
- Masaryk University, Czech Republic
- Johannes Gutenberg-Universitat Mainz, Germany

CONTACT

Philippe DENIEL, philippe.deniel@cea.fr

CALL

EuroHPC-01-2019

PROJECT TIMESPAN

01/04/2021 - 31/03/2024

MAELSTROM

MAchinE Learning for Scalable meTeoROlogy and cliMate

MAELSTROM

To develop Europe’s computer architecture of the future, MAELSTROM will co-design bespoke compute system designs for optimal application performance and energy efficiency, a software framework to optimise usability and training efficiency for machine learning at scale, and large-scale machine learning applications for the domain of weather and climate science.

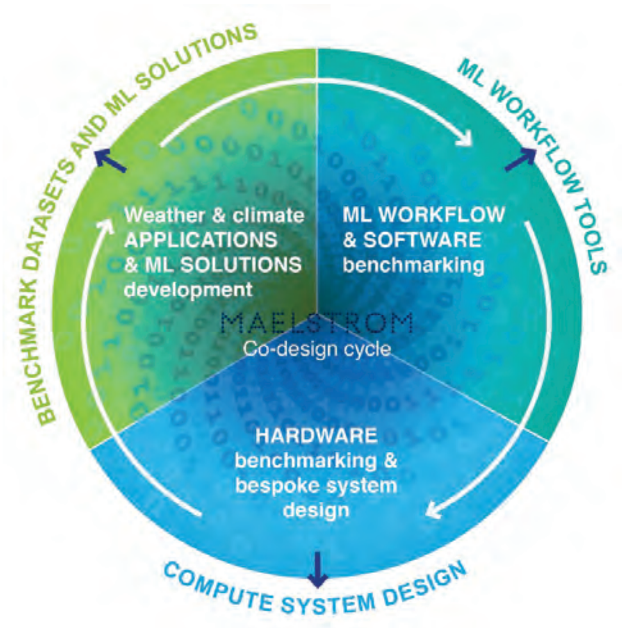

The MAELSTROM compute system designs will benchmark the applications across a range of computing systems regarding energy consumption, time-to-solution, numerical precision and solution accuracy.

Customised compute systems will be designed that are optimised for application needs to strengthen Europe's high-performance computing portfolio and to pull recent hardware developments, driven by general machine learning applications, toward needs of weather and climate applications.

The MAELSTROM software framework will enable scientists to apply and compare machine learning tools and libraries efficiently across a wide range of computer systems. A user interface will link application developers with compute system designers, and automated benchmarking and error detection of machine learning solutions will be performed during the development phase. Tools will be published as open source.

The MAELSTROM machine learning applications will cover all important components of the workflow of weather and climate predictions including the processing of observations, the assimilation of observations to generate initial and reference conditions, model simulations, as well as post-processing of model data and the development of forecast products. For each application, benchmark datasets with up to 10 terabytes of data will be published online for training and machine learning tool-developments at the scale of the fastest supercomputers in the world. MAELSTROM machine learning solutions will serve as blueprint for a wide range of machine learning applications on supercomputers in the future.

Extreme scale computing and data driven technologies

www.maelstrom-eurohpc.eu

Twitter: @MAELSTROM_EU
LinkedIn: @maelstrom-eu

COORDINATING ORGANISATION

European Centre for Medium-Range Weather Forecasts, United Kingdom

OTHER PARTNERS

- 4Cast GmbH & Co. KG, Germany
- E4 Computer Engineering SpA, Italy
- Eidgenössische Technische Hochschule Zürich, Switzerland
- Forschungszentrum Jülich, Germany
- Meteorologisk institutt, Norway
- Université du Luxembourg, Luxembourg

CONTACT

Peter Dueben, peter.dueben@ecmwf.int

CALL

EuroHPC-01-2019

PROJECT TIMESPAN

01/04/2021 - 31/03/2024

MAESTRO

Middleware for memory and data-awareness in workflows

Maestro will build a data-aware and memory-aware middleware framework that addresses ubiquitous problems of data movement in complex memory hierarchies and at many levels of the HPC software stack.

Though HPC and HPDA applications pose a broad variety of efficiency challenges, it would be fair to say that the performance of both has become dominated by data movement through the memory and storage systems, as opposed to floating point computational capability. Despite this shift, current software technologies remain severely limited in their ability to optimise data movement. The Maestro project addresses what it sees as the two major impediments of modern HPC software:

1. Moving data through memory was not always the bottleneck. The software stack that HPC relies upon was built through decades of a different situation – when the cost of performing floating point operations (FLOPS) was paramount. Several decades of technical evolution built a software stack and programming models highly fit for optimising floating-point operations but lacking in basic data handling functionality. We characterise the set of technical issues at missing data-awareness.

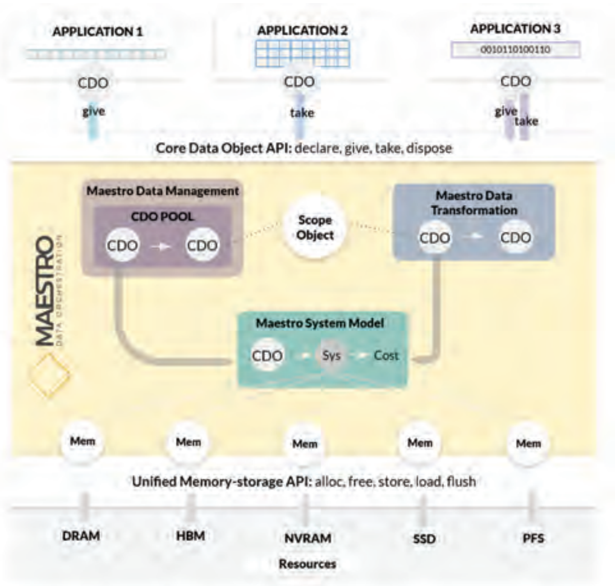

2. Software rightfully insulates users from hardware details, especially as we move higher up the software stack. But HPC applications, programming environments and systems software cannot make key data movement decisions without some understanding of the hardware, especially the increasingly complex memory hierarchy. With the exception of runtimes, which treat memory in a domain-specific manner, software typically must make hardware-neutral decisions which can often leave performance on the table. We characterise this issue as missing memory-awareness.

Maestro proposes a middleware framework that enables memory- and data-awareness. Maestro has developed new data-aware abstractions that can be be used in any level of software, e.g. compiler, runtime or application. Core elements are the Core Data Objects (CDO), which through a give/take semantics provided to the middleware. The middleware is being designed such that it enables modelling of memory and storage hierarchies to allow for reasoning about data movement and placement based on costs of moving data objects. The middleware will support automatic movement and promotion of data in memories and storage and allow for data transformations and optimisation.

Maestro follows a co-design methodology using a set of applications and workflows from diverse areas including numerical weather forecasting, earth-system modelling, materials sciences and in-situ data analysis pipelines for computational fluid dynamics simulations.

Transition to Exascale Computing

www.maestro-data.eu/

Twitter: @maestrodata
LinkedIn: @maestrodata

COORDINATING ORGANISATION
Forschungszentrum Jülich GmbH, Germany

OTHER PARTNERS
- CEA - Commissariat à l’Energie Atomique et aux énergies alternatives, France
- Appentra Solutions SL, Spain
- ETHZ - Eidgenössische Technische Hochschule Zürich, Switzerland
- ECMWF - European Centre for Medium-range Weather Forecasts, United Kingdom
- Seagate Systems UK Ltd, United Kingdom
- Hewlett Packard Enterprise, Switzerland

CONTACT
Dirk Pleiter, d.pleiter@fz-juelich.de

CALL
FETHPC-02-2017

PROJECT TIMESPAN
01/09/2018 - 30/11/2021

MEEP

The MareNostrum Experimental Exascale Platform

MEEP is an exploratory platform, flexible FPGA (field-programmable gate array) based emulation, to develop, integrate, test and co-design hardware and software for exascale supercomputers and other hardware targets, based on European-developed intellectual property (IP).

MEEP provides two important functions:

- An evaluation platform of pre-silicon IP and ideas, at speed and scale;
- Software development and experimentation platform to enable software readiness for new hardware.

MEEP enables software development, accelerating software maturity, compared to the limitations of software simulation. IP can be tested and validated before moving to silicon, saving time and money.

The objectives of MEEP are to leverage and extend projects like the European Processor Initiative (EPI) and the Performance Optimisation and Productivity Centre of Excellence (POP CoE).

The ultimate goal of the project is to create an open full-stack ecosystem that can be used for academic purposes and integrated into a functional accelerator or cores for traditional and emerging HPC applications.

European Processor

meep-project.eu

COORDINATING ORGANISATION
Barcelona Supercomputing Center (BSC), Spain

OTHER PARTNERS
- University of Zagreb, Faculty of Electrical Engineering and Computing, Croatia
- The Scientific and Technological Research Council of Turkey, Turkey

CONTACT
meep-project@bsc.es

CALL
EuroHPC-06-2019

PROJECT TIMESPAN
01/01/2020 - 31/12/2022

Mont-Blanc 2020

European scalable, modular and power-efficient HPC processor

Now that the Mont-Blanc 2020 project has come to an end, it is time to look back on its accomplishments. Mont-Blanc 2020 was the last of a long series of projects. The initial concept when we started in 2011 was very disruptive: leveraging mobile (Arm) chips and their power-efficiency to run high-performance computing (HPC) applications. Successive Mont-Blanc projects have witnessed and accompanied the rise of Arm processors in servers: the current Fugaku system ranked first in the TOP500 has reached the apex of high-end Arm CPUs, vindicating the initial rationale behind the project.

One of the strong points of the Mont-Blanc projects was theirindustry / academia collaboration. Mont-Blanc 2020 was no exception, with a team of three core partners with complementary profiles (Arm, Atos, BSC), three active SMEs (Kalray, Semidynamics, Sipearl) and prominent research partners (BSC, CEA, Jülich Supercomputing Centre).

The focus of Mont-Blanc 2020 was processor design. It essentially addressed system-on-chip (SoC) design and processor intellectual property (IP) to enable the work of the European Processor Initiative (EPI), focusing on the challenge of achieving extremely high performance per Watt. For the compute unit, we selected the Arm instruction set, with its Scalable Vector Extension (SVE) optimized for HPC and artificial intelligence (AI). It has particular technological relevance for high-end cores; more importantly, the availability of a dynamic software ecosystem was necessary to run real applications as required by our co-design methodology.

An important achievement of Mont-Blanc 2020 is the tools and methodology we selected and developed for processor simulation and virtual prototyping, i.e. the tools that allowed our researchers to test applications and evaluate future performance prior to silicon availability. We developed a unique co-design methodology for SoC infrastructure verification and optimization. Co-design is always a challenge, but we faced an additional test, which was to get hardware and software teams from different organizations to work together. We had to build a bridge between the computer-aided design (CAD) tools used by our industrial partners and the open source tools used by our academic partners. Our approach has increased the speed of simulation by a factor of 1,000 and even by 10,000 for some applications.

Many of the features we developed for our Simulation Framework are already used outside of the project. For example, Mont-Blanc

2020 was instrumental in the implementation of SVE instructions in gem5, which is part of the official open source release 20.0 of gem5. Another example is the SVE-related improvements to the MUlti-scale Simulation Approach (MUSA) developed within Mont-Blanc 2020, which are used within EPI.

However, the top Mont-Blanc 2020 achievement is without doubt the IP developed for a low power network on chip (NoC). NoC is critical in a SoC when targeting highly demanding applications. It is also very challenging: in manycore architectures you need to maintain low latency while increasing the number of cores as well as the throughput of each core. The NoC IP developed by Mont-Blanc 2020 will be included in the next-generation EPI processor. Our NoC and related NoC IPs are also integrated in Atos's IP portfolio that will serve future commercial and research projects.

To conclude, European sovereignty in the provision of HPC technology was part of Mont-Blanc's vision from the start. We aimed to contribute to the revival of the European SoC design ecosystem by creating an IP portfolio as well as boosting the skills necessary for chip design. Today we can say: mission accomplished!

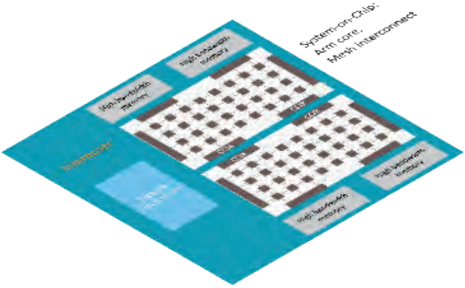

The Mont-Blanc 2020 physical architecture

European Processor

montblanc-project.eu

Twitter: @MontBlanc_Eu
LinkedIn: @mont-blanc-project

COORDINATING ORGANISATION
Atos (Bull SAS), France

OTHER PARTNERS
- Arm, United Kingdom
- Barcelona Supercomputing Centre (BSC), Spain
- CEA - Commissariat à l'Energie Atomique et aux énergies alternatives, France
- Jülich Forschungszentrum, Germany
- Kalray, France
- SemiDynamics, Spain

CONTACT
Etienne Walter,
Etienne.walter@atos.net

CALL
ICT-05-2017

PROJECT TIMESPAN
01/12/2017 - 31/03/2021

NEASQC

NExt ApplicationS of Quantum Computing

The NEASQC project brings together academic experts and industrial end-users to investigate and develop a new breed of Quantum-enabled applications that can take advantage of NISQ (Noise Intermediate-Scale Quantum) systems in the near future. NEASQC is use-case driven, addressing practical problems such as drug discovery, CO2 capture, smart energy management, natural language processing, breast cancer detection, probabilistic risk assessment for energy infrastructures or hydrocarbon well optimisation. NEASQC aims to initiate an active European community around NISQ Quantum Computing by providing a common toolset that will attract new industrial users.

NEASQC aims at demonstrating that, though the millions of qubits that will guarantee fully fault-tolerant quantum computing are still far away, there are practical use cases for the NISQ (Noise Intermediate-Scale Quantum) devices that will be available in the near future. NISQ computing can deliver significant advantages when running certain applications, thus bringing game-changing benefits to users, and particularly industrial users.

The NEASQC consortium has chosen a wide selection of NISQ-compatible industrial and financial use-cases, and will develop new quantum software techniques to solve those use-cases with a practical quantum advantage. To achieve this, the project brings together a multidisciplinary consortium of academic and industry experts in Quantum Computing, High Performance Computing, Artificial Intelligence, chemistry...

The ultimate ambition of NEASQC is to encourage European user communities to investigate NISQ quantum computing. For this purpose, the project consortium will define and make available a complete and common toolset that new industrial actors can use to start their own practical investigation and share their results.

NEASQC also aims to build a much-needed bridge between Quantum Computing hardware activities, particularly those of the Quantum Technologies flagship, and the end-user community. Even more than in classical IT, NISQ computing demands a strong cooperation between hardware teams and software users. We expect our work in use cases will provide strong directions for the development of NISQ machines, what will be very valuable to the nascent quantum hardware industry.

NEASQC OBJECTIVES

1. Develop 9 industrial and financial use cases with a practical quantum advantage for NISQ machines.
2. Develop open source NISQ programming libraries for industrial use cases, with a view to facilitate quantum computing experimentation for new users.
3. Build a strong user community dedicated to industrial NISQ applications.
4. Develop software stacks and benchmarks for the Quantum Technology Flagship hardware platforms.

NEASQC is one of the projects selected within the second wave of Quantum Flagship projects.

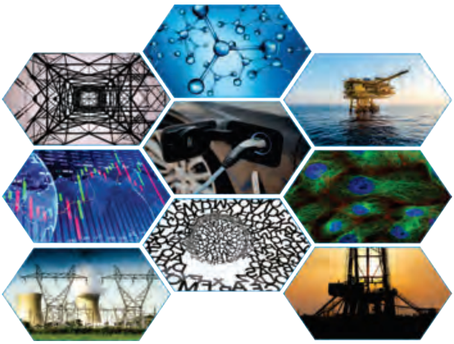

Quantum Computing

neasqc.eu

Twitter: @NEASQC
LinkedIn: @neasqc-project

COORDINATING ORGANISATION
Atos (Bull SAS), France

OTHER PARTNERS
- AstraZeneca AB, Sweden
- CESGA (Fundación Publica Gallega Centro Tecnológico de Supercomputación de Galicia), Spain
- Electricité de France (EDF), France
- HQS Quantum Simulations GmbH, Germany
- HSBC Bank Plc, United Kingdom
- Irish Centre for High-End Computing (ICHEC), Ireland
- Leiden University, Netherlands
- TILDE SIA, Latvia
- TOTAL S.A., France
- Universidade da Coruña – CITIC, Spain
- Université de Lorraine (LORIA), France

CONTACT
Cyril Allouche, cyril.allouche@atos.net

CALL
FETFLAG-05-2020

PROJECT TIMESPAN
01/09/2020 – 31/08/2024

RECIPE

REliable power and time-ConstraInts-aware Predictive management of heterogeneous Exascale systems

OBJECTIVE

The current HPC facilities will need to grow by an order of magnitude in the next few years to reach the Exascale range. The dedicated middleware needed to manage the enormous complexity of future HPC centres, where deep heterogeneity is needed to handle the wide variety of applications within reasonable power budgets, will be one of the most critical aspects in the evolution of HPC infrastructure towards Exascale. This middleware will need to address the critical issue of reliability in face of the increasing number of resources, and therefore decreasing mean time between failures.

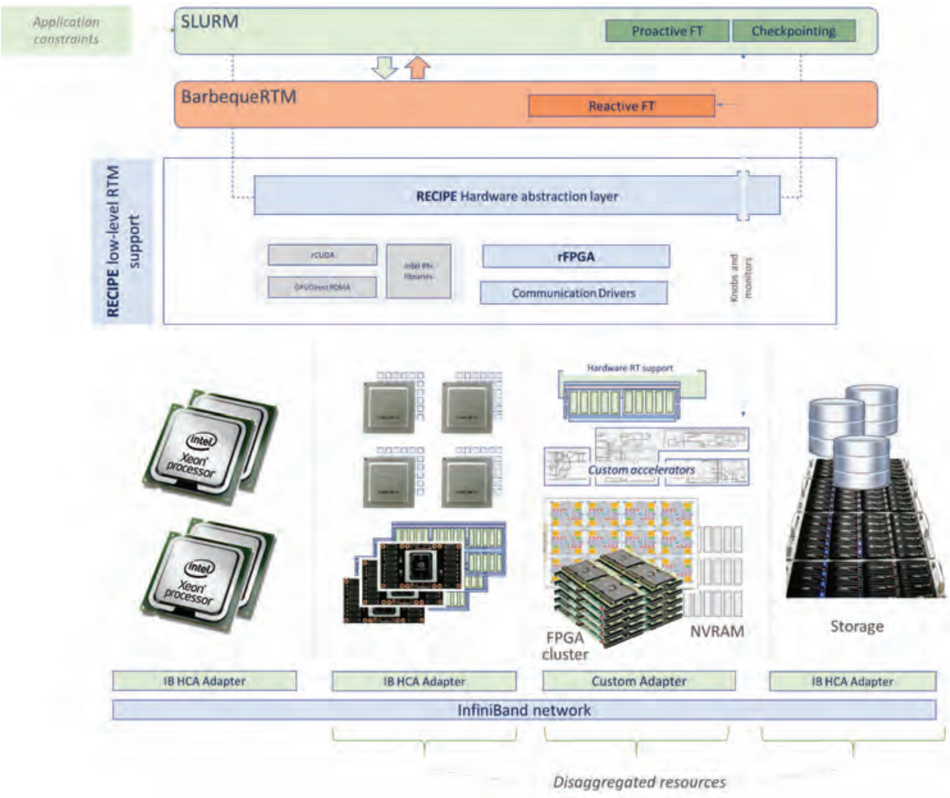

The RECIPE Hardware-Software Stack

To close this gap, RECIPE provides: a hierarchical runtime resource management infrastructure optimizing energy efficiency and ensuring reliability for both time-critical and throughput-oriented computation; a predictive reliability methodology to support the enforcing of QoS guarantees in face of both transient and long-term hardware failures, including thermal, timing and reliability models; and a set of integration layers allowing the resource manager to interact with both the application and the underlying deeply heterogeneous architecture, addressing them in a disaggregate way.

Quantitative goals for RECIPE include: 25% increase in energy efficiency (performance/watt) with an 15% MTTF improvement due to proactive thermal management; energy-delay product improved up to 25%; 20% reduction of faulty executions.

The project will assess its results against the following set of real world use cases, addressing key application domains ranging from well-established HPC applications such as geophysical exploration and meteorology, to emerging application domains such as biomedical machine learning and data analytics.

To this end, RECIPE relies on a consortium composed of four leading academic partners (POLIMI, UPV, EPFL, CeRICT); two supercomputing centres, BSC and PSNC; a research hospital, CHUV, and an SME, IBTS, which provide effective exploitation avenues through industry-based use cases.

Transition to Exascale Computing

www.recipe-project.eu

Twitter: @EUrecipe

COORDINATING ORGANISATION
Politecnico di Milano, Italy

OTHER PARTNERS
- UPV - Universitat Politecnica de Valencia, Spain
- CeRICT - Centro Regionale Information Communication Technology scrl, Italy
- BSC - Barcelona Supercomputing Centre, Spain
- PSNC - Instytut Chemii Bioorganicznej Polskiej Akademii Nauk, Poland
- EPFL - Ecole Polytechnique Fédérale de Lausanne, Switzerland
- IBTS - Intelligence behind Things Solutions SRL, Italy
- CHUV - Centre Hospitalier Universitaire Vaudois, Switzerland

CONTACT
Prof. William Fornaciari, william.fornaciari@polimi.it
Prof. Giovanni Agosta, agosta@acm.org

CALL
FETHPC-02-2017

PROJECT TIMESPAN
01/05/2018 - 31/10/2021

RED-SEA

Network Solution for Exascale Architectures

Network interconnects play an enabling role in HPC systems – and this will be even truer for the coming Exascale systems that will rely on higher node counts and increased use of parallelism and communication. Moreover, next-generation HPC and data-driven systems will be powered by heterogeneous computing devices, including low-power Arm and RISC-V processors, high-end CPUs, vector acceleration units and GPUs suitable for massive single-instruction multiple-data (SIMD) workloads, as well as FPGA and ASIC designs tailored for extremely power-efficient custom codes.

These compute units will be surrounded by distributed, heterogeneous (often deep) memory hierarchies, including high-bandwidth memories and fast devices offering microsecond-level access time. At the same time, modern data-parallel processing units such as GPUs and vector accelerators can crunch data at amazing rates (tens of TFLOPS). In this landscape, the network may well become the next big bottleneck, similar to memory in single node systems.

RED-SEA will build upon the European interconnect BXI (BullSequana eXascale Interconnect), together with standard and mature technology (Ethernet) and previous EU-funded initiatives to provide a competitive and efficient network solution for the exascale era and beyond. This involves developing the key IPs and the software environment that will deliver:

- scalability, while maintaining an acceptable total cost of ownership and power efficiency;
- virtualization and security, to allow various applications to efficiently and safely share an HPC system;
- Quality-of-service and congestion management to make it possible to share the platform among users and applications with different demands;
- reliability at scale, because fault tolerance is a key concern in a system with a very large number of components;
- support of high-bandwidth low-latency HPC Ethernet, as HPC systems increasingly need to interact securely with the outside world, including public clouds, edge servers or third party HPC systems;
- support of heterogeneous programming model and runtimes to facilitate the convergence of HPC and HPDA;
- support for low-power processors and accelerators.

RED-SEA IN THE MODULAR SUPERCOMPUTING ARCHITECTURE

RED-SEA supports the Modular Supercomputing Architecture (MSA) that underpins all of the SEA projects (DEEP-SEA and IO-SEA, two other projects resulting from the EuroHPC-01-2019 call).

In the MSA, BXI is the HPC fabric within each compute module, delivering low-latency, high bandwidth and all required HPC features, whereas Ethernet is the high-performance federative network that offers interface to storage and with other compute modules. RED-SEA will design a seamless interface between BXI and Ethernet via a new Gateway solution.

Extreme scale computing and data driven technologies

redsea-project.eu
Twitter: @redsea_EU
LinkedIn: @redsea-project

COORDINATING ORGANISATION

Atos (Bull SAS), France

OTHER PARTNERS

- Idryma Technologias Kai Erevnas, Greece
- Commissariat à l'énergie atomique et aux énergies alternatives (CEA), France
- Forschungszentrum Jülich, Germany
- EXTOLL GmbH, Germany
- Eidgenössische Technische Hochschule Zürich, Switzerland
- Universitat Politècnica de València, Spain
- Universidad de Castilla - La Manche, Spain
- Istituto Nazionale di Fisica Nucleare – INFN, Italy
- Exapsys - Exascale Performance Systems, Greece
- eXact lab SRL, Italy

CONTACT

redsea-project.eu/contact/

CALL

EuroHPC-01-2019

PROJECT TIMESPAN

01/04/2021 - 31/03/2024

Sage2

Percipient Storage for Exascale Data Centric Computing 2

OBJECTIVE

The landscape for High Performance Computing is changing with the proliferation of enormous volumes of data created by scientific instruments and sensors, in addition to data from simulations. This data needs to be stored, processed and analysed, and existing storage system technologies in the realm of extreme computing need to be adapted to achieve reasonable efficiency in achieving higher scientific throughput.

We started on the journey to address this problem with the SAGE project. The HPC use cases and the technology ecosystem is now further evolving and there are new requirements and innovations that are brought to the forefront. It is extremely critical to address them today without "reinventing the wheel" leveraging existing initiatives and know-how to build the pieces of the Exascale puzzle as quickly and efficiently as we can.

The SAGE paradigm already provides a basic framework to address the extreme scale data aspects of High Performance Computing on the path to Exascale. Sage2 (Percipient StorAGe for Exascale Data Centric Computing 2) intends to validate a next generation storage system building on top of the already existing SAGE platform to address new use case requirements in the areas of extreme scale computing scientific workflows and AI/deep learning leveraging the latest developments in storage infrastructure software and storage technology ecosystem.

Sage2 aims to provide significantly enhanced scientific throughput, improved scalability, and, time & energy to solution for the use cases at scale. Sage2 will also dramatically increase the productivity of developers and users of these systems.

Sage2 will provide a highly performant and resilient, QoS capable multi-tiered storage system, with data layouts across the tiers managed by the Mero Object Store, which is capable of handling in-transit/in-situ processing of data within the storage system, accessible through the Clovis API.

The 3-rack prototype system is now available for use at Jülich Supercomputing Centre.

Transition to Exascale Computing

sagestorage.eu

Twitter: @SageStorage

COORDINATING ORGANISATION

Seagate Systems UK Ltd, United Kingdom

OTHER PARTNERS

- Atos (Bull SAS), France
- CEA - Commissariat à l'Energie Atomique et aux énergies alternatives, France
- United Kingdom Atomic Energy Authority, United Kingdom
- Kungliga Tekniska Hoegskolan (KTH), Sweden
- Forschungszentrum Jülich GmbH, Germany
- The University of Edinburgh, United Kingdom
- Kitware SAS, France
- Arm Ltd, United Kingdom

CONTACT

Sai Narasimhamurthy,
sai.narasimhamurthy@seagate.com

CALL

FETHPC-02-2017

PROJECT TIMESPAN

01/09/2018 - 30/11/2021

SODALITE

SOftware Defined AppLication Infrastructures managemenT and Engineering

Empowering the DevOps World

In recent years the global ICT market has seen a tremendous rise in utility computing, which serves as the backend for practically any new technology, methodology or advancement from healthcare to aerospace. We are entering a new era of heterogeneous, software-defined, high performance compute environments. In this context, SODALITE creates order into chaos, as the way for all these technologies to be harmonized in a single ecosystem extending their focus to other types of hardware and architectures.

SODALITE brings the vast knowledge of performance optimisation accrued by the HPC industry into the cloud computing arena. So, it exploits a model-driven approach, empowered by ontological reasoning, to support experts operating applications or resources over heterogeneous infrastructures. In general lines, the most benefited actors are:

Application Ops Experts. Those in charge of operating (deploy, execute, optimize and monitor) applications, as they will be able to better understand the application requirements

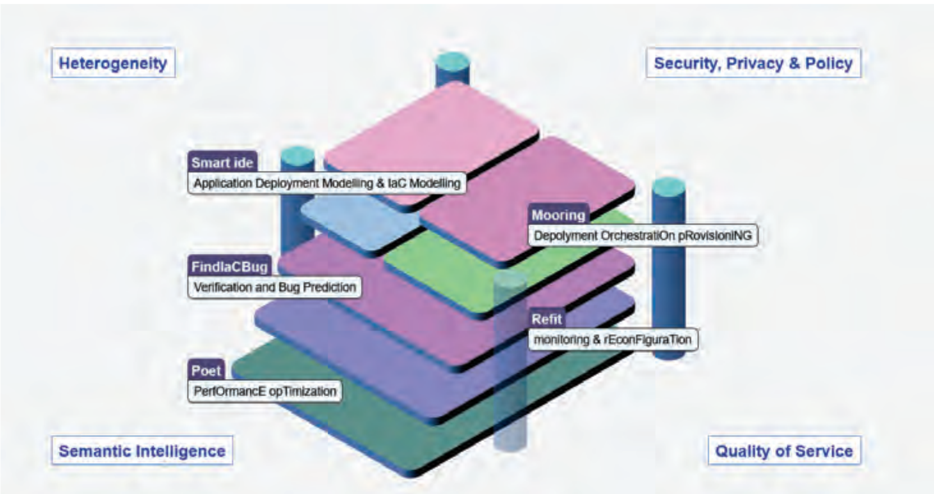

in terms of deployment/execution environment and QoS.
Resource Experts. As from the infrastructure management point of view, they can operate resources on different environments, such as cloud, HPC or edge, in a simplified manner.
Quality Experts. As they can make use of SODALITE libraries of patterns for addressing specific performance and quality problems, both from infrastructure and application perspectives.

SODALITE PILLARS AND LAYERS

SODALITE offers a layered solution (5 layers) based on four pillars: (i) semantic intelligence for describing, deploying and optimising applications; (ii) quality of service based on guaranteed SLAs; (iii) heterogeneity support of underlying infrastructures; and (iv) security, privacy and policy. Each of these layers represents a product itself but can be combined for an enhanced experience.

Layer 1: Smart IDE
It provides the user with an IDE and a semantic reasoner API based on sophisticated knowledge graphs and semantic reasoning.

Layer 2: deployMent OrchestratiOn pRovisioniNG (MOORING)
Consists of a framework for the usage of IaC and concepts of modelling application deployment, ensuring a smooth resource operation.

Layer 3: Verification & Bug Prediction (FindIaCBug)
It is responsible for the verification of the IaC code and the detection of various quality issues such as errors, smells, antipatterns and bugs in IaC artifacts.

Layer 4: monitoRing & rEconFIguraTion (REFIT)
Ensures the monitoring of system and network, collecting application-level and infrastructure-level metrics and events, used for reconfiguration and refactoring.

Layer 5: PerfOrmancE optimization
This layer is composed of static and dynamic application optimisation, during design and runtime.

Software Technologies

www.sodalite.eu

Twitter: @SODALITESW
LinkedIn: @sodalite-eu
GitHub: github.com/SODALITE-EU

COORDINATING ORGANISATION
XLAB, Slovenia

OTHER PARTNERS

- Atos Spain, Spain
- IBM, Israel
- Politecnico di Milano (POLIMI), Italy
- HPE, Switzerland
- ITI-CERTH, Greece
- Adaptant, Germany
- Jheronimus Academy of Data Science (JADS), Netherlands
- High Performance Computing Center - Universität Stuttgart (HLRS), Germany

CONTACT
Project Coordinator:
Daniel Vladušič,
daniel.vladusic@xlab.si +38640352527

CALL
H2020-ICT-2018-2

TIMESPAN
01/02/2019 - 31/01/2022

SPARCITY

An Optimization and Co-design Framework for Sparse Computation

Perfectly aligned with the vision of the EuroHPC Joint Undertaking, the SparCity project aims at creating a supercomputing framework that will provide efficient algorithms and coherent tools specifically designed for maximising the performance and energy efficiency of sparse computations on emerging HPC systems, while also opening up new usage areas for sparse computations in data analytics and deep learning.

The framework enables comprehensive application characterization and modeling, performing synergistic node-level and system-level software optimizations. By creating a digital SuperTwin, the framework is also capable of evaluating existing hardware components and addressing what-if scenarios on emerging architectures and systems in a co-design perspective.

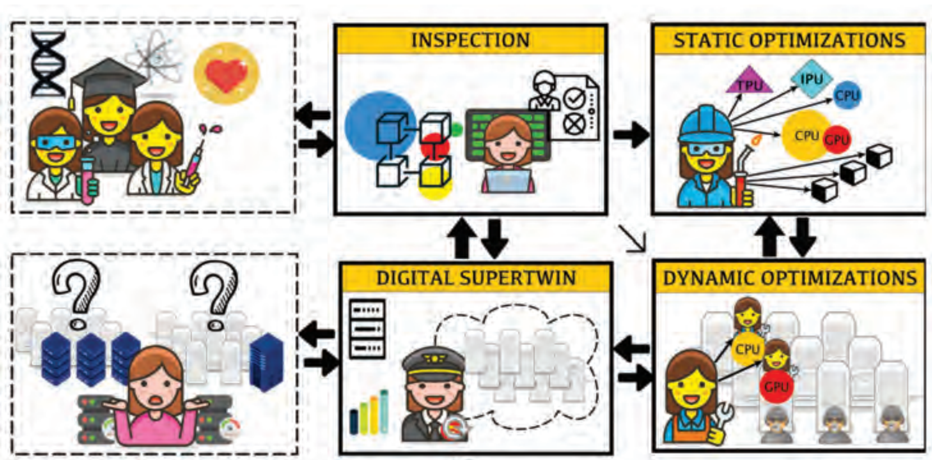

To demonstrate the effectiveness, societal impact, and usability of the framework, the SparCity project will enhance the computing scale and energy efficiency of four challenging real-life applications that come from drastically different domains, namely, computational cardiology, social networks, bioinformatics and autonomous driving. By targeting this collection of challenging applications, SparCity will develop world-class, extreme scale and energy-efficient HPC technologies, and contribute to building a sustainable exascale ecosystem and increasing Europe's competitiveness.

Extreme scale computing and data driven technologies

sparcity.eu

Twitter: @sparcity_eu
LinkedIn: sparcity-project-944b4320a

COORDINATING ORGANISATION
Koç Üniversitesi, Turkey

OTHER PARTNERS
- Sabancı University,Turkey
- Simula Research Laboratory AS, Norway
- INESC ID - Instituto de Engenharia de Sistemas e Computadores, Investigação e Desenvolvimento em Lisboa, Portugal
- Ludwig-Maximilians-Universität München, Germany
- Graphcore AS, Norway

CONTACT
Prof. Didem Unat, dunat@ku.edu.tr

CALL

EuroHPC-01-2019

PROJECT TIMESPAN
01/04/2021 - 31/03/2024

TEXTAROSSA

Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale

To achieve high performance and high energy efficiency on near-future exascale computing systems, a technology gap needs to be bridged: increase efficiency of computation with extreme efficiency in HW and new arithmetics, as well as providing methods and tools for seamless integration of reconfigurable accelerators in heterogeneous HPC multi-node platforms. TEXTAROSSA aims at tackling this gap through applying a co-design approach to heterogeneous HPC solutions, supported by the integration and extension of IPs, programming models and tools derived from European research projects, led by TEXTAROSSA partners.

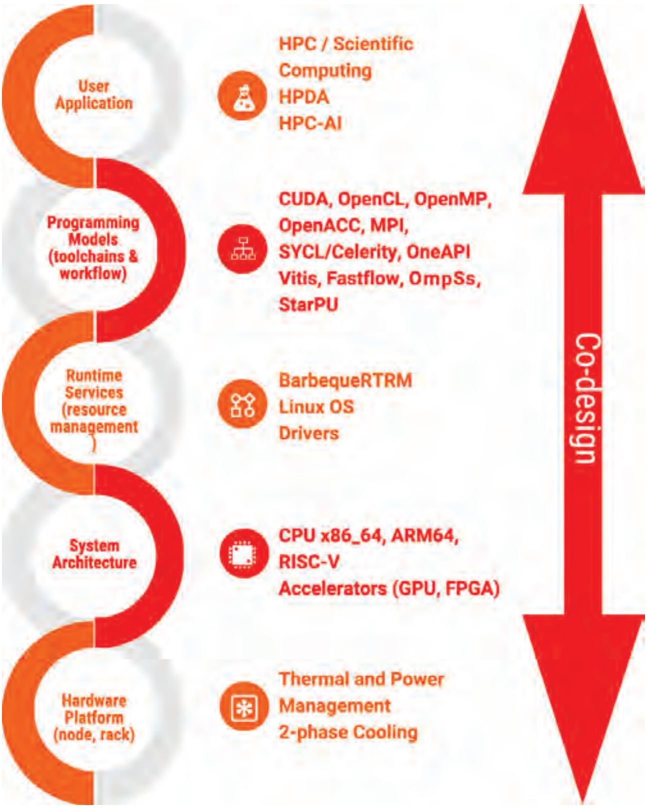

The TEXTAROSSA Co-Design flow

The main directions for innovation are towards:

1. enabling mixed-precision computing, through the definition of IPs, libraries, and compilers supporting novel data types (including Posits), used also to boost the performance of AI accelerators;
2. implementing new multilevel thermal management and two-phase liquid cooling;
3. developing improved data movement and storage tools through compression;
4. ensure secure HPC operation through HW accelerated cryptography;
5. providing RISC-V based IP for fast task scheduling and IPs for low-latency intra/inter-node communication.

These technologies will be tested on the Integrated Development Vehicles mirroring and extending the European Processor Initiative ARM64-based architecture, and on an OpenSequana testbed. To drive the technology development and assess the impact of the proposed innovations TEXTAROSSA will use a selected but representative number of HPC, HPDA and AI demonstrators covering challenging HPC domains such as general-purpose numerical kernels, High Energy Physics (HEP), Oil & Gas, climate modelling, and emerging domains such as High Performance Data Analytics (HPDA) and High Performance Artificial Intelligence (HPC-AI).

The TEXTAROSSA consortium includes: three leading Italian universities (Politecnico di Milano, Università degli studi di Torino and Università di Pisa, Linked Third Parties of CINI), CINECA (Italy, Inkind Third Party of ENEA) and Universitat Politècnica de Catalunya (UPC, Spain, Inkind Third Party of BSC).

Extreme scale computing and data driven technologies

textarossa.eu

Twitter: @TEXTAROSSA1
Facebook: @textarossa

COORDINATING ORGANISATION

Agenzia nazionale per le nuove tecnologie, l'energia e lo sviluppo economico sostenibile (ENEA), Italy

OTHER PARTNERS

- Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V., Germany
- Consorzio Interuniversitario Nazionale per l’Informatica (CINI), Italy
- Institut National de Recherche en Informatique et en Automatique (INRIA), France
- Atos (Bull SAS), France
- E4 Computer Engineering SpA, Italy
- Barcelona Supercomputing Centre (BSC), Spain
- Instytut Chemii Bioorganicznej Polskiej Akademii Nauk, Poland
- Istituto Nazionale di Fisica Nucleare – INFN, Italy
- Consiglio Nazionale delle Ricerche (CNR), Italy
- In Quattro Srl, Italy

CONTACT

Project Coordinator: Massimo Celino, ENEA, massimo.celino@enea.it
Project Technical Manager: William Fornaciari, Politecnico di Milano – DEIB william.fornaciari@polimi.it

CALL

EuroHPC-01-2019

PROJECT TIMESPAN

01/04/2021 - 31/03/2024

TIME-X

TIME parallelisation: for eXascale computing and beyond

Recent successes have established the potential of parallel-in-time integration as a powerful algorithmic paradigm to unlock the performance of Exascale systems. However, these successes have mainly been achieved in a rather academic setting, without an overarching understanding. TIME-X will take the next leap in the development and deployment of this promising new approach for massively parallel HPC simulation, enabling efficient parallel-in-time integration for real-life applications.

We will:

- Provide software for parallel-in-time integration on current and future Exascale HPC architectures, delivering substantial improvements in parallel scaling;
- Develop novel algorithmic concepts for parallel-in-time integration, deepening our mathematical understanding of their convergence behaviour and including advances in multi-scale methodology;
- Demonstrate the impact of parallel-in-time integration, showcasing the potential on problems that, to date, cannot be tackled with full parallel efficiency in three diverse and challenging application fields with high societal impact: weather and climate, medicine, drug design and electromagnetics.

To realise these ambitious, yet achievable goals, the inherently inter-disciplinary TIME-X Consortium unites top researchers from numerical analysis and applied mathematics, computer science and the selected application domains. Europe is leading research in parallel-in-time integration.

TIME-X unites all relevant actors at the European level for the first time in a joint strategic research effort. The project will enable taking the necessary next step: advancing parallel-in-time integration from an academic/mathematical methodology into a widely available technology with a convincing proof of concept, maintaining European leadership in this rapidly advancing field and paving the way for industrial adoption.

Extreme scale computing and data driven technologies

time-x-eurohpc.eu

COORDINATING ORGANISATION

KU Leuven, Belgium

OTHER PARTNERS

- Sorbonne Université, France
- Forschungszentrum Jülich, Germany
- Technische Universität Hamburg, Germany
- Technische Universität Darmstadt, Germany
- Università della Svizzera italiana, Switzerland
- Bergische Universität Wuppertal, Germany
- Technische Universität München, Germany
- Ecole nationale des ponts et chaussées, France
- Université de Genève, Switzerland

CONTACT

Giovanni Samaey,
giovanni.samaey@kuleuven.be

CALL

EuroHPC-01-2019

PROJECT TIMESPAN

01/04/2021 - 31/03/2024

VECMA

Verified Exascale Computing for Multiscale Applications

The purpose of this project is to enable a diverse set of multiscale, multiphysics applications -- from fusion and advanced materials through climate and migration, to drug discovery and the sharp end of clinical decision making in personalised medicine -- to run on current multi-petascale computers and emerging exascale environments with high fidelity such that their output is «actionable». That is, the calculations and simulations are certifiable as validated (V), verified (V) and equipped with uncertainty quantification (UQ) by tight error bars such that they may be relied upon for making important decisions in all the domains of concern.

The central deliverable is an open source toolkit for multiscale VVUQ based on generic multiscale VV and UQ primitives, to be released in stages over the lifetime of this project, fully tested and evaluated in emerging exascale environments, actively promoted over the lifetime of this project, and made widely available in European HPC centres.

The project includes a fast track that will ensure applications are able to apply available multiscale VVUQ tools as soon as possible, while guiding the deep track development of new capabilities and their integration into a wider set of production applications by the end of the project. The deep track includes the development of more disruptive and automated algorithms, and their exascale-aware implementation in a more intrusive way with respect to the underlying and pre-existing multiscale modelling and simulation schemes.

The potential impact of these certified multiscale simulations is enormous, and we have already been promoting the VVUQ toolkit (VECMAtk) across a wide range of scientific and social scientific domains, as well as within computational science more broadly. We have made ten releases of the toolkit with the dedicated website at https://www.vecma-toolkit.eu/ available for public users to access, download and use.

To further develop and disseminate the VECMAtk, we have been working with the Alan Turing Institute to jointly run the planned event of Reliability and Reproducibility in Computational Science: Implementing Verification, Validation and Uncertainty Quantification in silico. The event comprised as part of it the first VECMA training workshop in January 2020.

Ahead of this event, we had run a VECMA hackathon event in association with VECMAtk in September 2019.

We held an online conference: Multiscale Modelling, Uncertainty Quantification and the Reliability of Computer Simulations in June 2020. The conference was a combination of three events that were due to take place at the SIAM Conference on Uncertainty Quantification (UQ20) and the International Conference on Computational Science (ICCS) 2020. These events associated with VECMAtk have participation from some of our external users from industry and government who use the tools we are developing on HPC at supercomputer centres in Europe.

Transition to Exascale Computing

www.vecma.eu

Twitter: @VECMA4

COORDINATING ORGANISATION

University College London, United Kingdom

OTHER PARTNERS

- Max-Planck-Gesellschaft zur Förderung der Wissenschaften EV, Germany
- Brunel University London, United Kingdom
- Bayerische Akademie der Wissenschaften, Germany
- Atos (Bull SAS), France
- Stichting Nederlandse Wetenschappelijk Onderzoek Instituten, Netherlands
- CBK Sci Con Ltd, United Kingdom
- Universiteit van Amsterdam, Netherlands
- Instytut Chemii Bioorganicznej Polskiej Akademii Nauk, Poland

INTERNATIONAL PARTNER

Rutgers University, USA

CONTACT

Hugh Martin h.martin@cbkscicon.com

CALL

FETHPC-02-2017

PROJECT TIMESPAN

15/06/2018 - 14/12/2021

VESTEC

Visual Exploration and Sampling Toolkit for Extreme Computing

OBJECTIVE

Technological advances in high performance computing are creating exciting new opportunities that move well beyond improving the precision of simulation models. The use of extreme computing in real-time applications with high velocity data and live analytics is within reach. The availability of fast growing social and sensor networks raises new possibilities in monitoring, assessing and predicting environmental, social and economic incidents as they happen. Add in grand challenges in data fusion, analysis and visualization, and extreme computing hardware has an increasingly essential role in enabling efficient processing workflows for huge heterogeneous data streams.

VESTEC will create the software solutions needed to realise this vision for urgent decision making in various fields with high impact for the European community. VESTEC will build a flexible toolchain to combine multiple data sources, efficiently extract essential features, enable flexible scheduling and interactive supercomputing, and realise 3D visualization environments for interactive explorations by stakeholders and decision makers. VESTEC will develop and evaluate methods and interfaces to integrate high-performance data analytics processes into running simulations and real-time data environments. Interactive ensemble management will launch new simulations for new data, building up statistically more and more accurate pictures of emerging, time-critical phenomena. Innovative data compression approaches, based on topological feature extraction and data sampling, will result in considerable reductions in storage and processing demands by discarding domain-irrelevant data.

Three emerging use cases will demonstrate the immense benefit for urgent decision making: wildfire monitoring and forecasting; analysis of risk associated with mosquito-borne diseases; and the effects of space weather on technical supply chains. VESTEC brings together experts in each domain to address the challenges holistically.

Transition to Exascale Computing

www.vestec-project.eu

Twitter: @VESTECproject

COORDINATING ORGANISATION

DLR - Deutsches Zentrum für Luft- und Raumfahrt EV, Germany

OTHER PARTNERS

- The University of Edinburgh, United Kingdom
- Kungliga Tekniska Hoegskolan (KTH), Sweden
- Sorbonne Université, France
- Kitware SAS, France
- Intel Deutschland GmbH, Germany
- Fondazione Bruno Kessler, Italy
- Université Paul Sabatier Toulouse III, France
- Tecnosylva SL, Spain

CONTACT

Andreas Gerndt, andreas.gerndt@dlr.de

CALL

FETHPC-02-2017

PROJECT TIMESPAN

01/09/2018 - 31/08/2021

Centres of Excellence in computing applications

FocusCoE

Concerted action for the European HPC CoEs

FocusCoE contributes to the success of the EU HPC Ecosystem and the EuroHPC Initiative by supporting the EU HPC Centres of Excellence (CoEs) to more effectively fulfil their role within the ecosystem and initiative: ensuring that extreme scale applications result in tangible benefits for addressing scientific, industrial or societal challenges. It achieves this by creating an effective platform for the CoEs to coordinate strategic directions and collaboration (addressing possible fragmentation of activities across the CoEs and coordinating interactions with the overall HPC ecosystem) and provides support services for the CoEs in relation to both industrial outreach and promotion of their services and competences by acting as a focal point for users to discover those services.

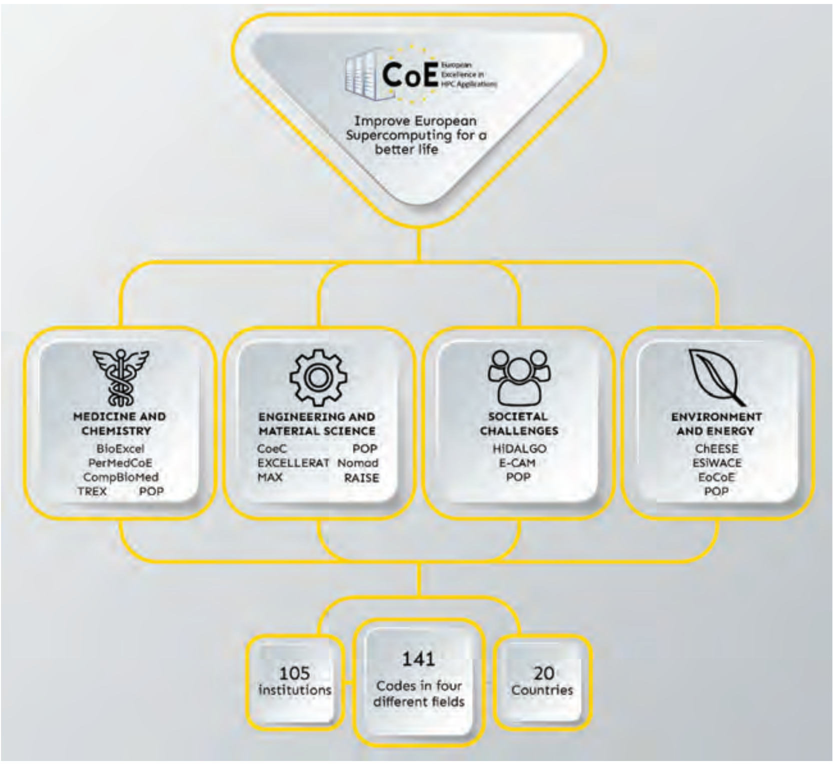

The specific objectives and achievements of FocusCoE are

- Instigating the creation of the *HPC CoE Council* (HPC3) and supporting its operation. HPC3 is a vehicle that allows all European HPC CoEs to collectively define an overriding strategy and realise an effective interaction with the European HPC Ecosystem.
- To support the HPC CoEs to achieve enhanced interaction with industry, and SMEs in particular, through concerted outreach and business development actions. FocusCoE initiated contacts between CoEs and industry by several actions including presenting CoEs at sectorial events and organizing thematic webinars. A first tranche of industrial success stories can be found at www.hpccoe.eu/success-stories/
- To instigate concerted action on training by and for the complete set of HPC CoEs. FocusCoE organised workshops on HPC training development in Europe which also served to promote CoE training, supplemented by CoE training events ranging from pedagogy through benchmarking to innovation development targeted at CoE personnel. The CoE Training registry (www.hpccoe.eu/coe-training-calendar/) provides a unique view of all CoE trainings.
- To promote and concert the capabilities of and services offered by the HPC CoEs and development of the EU HPC CoE "brand" raising awareness with stakeholders and both academic and industrial users. FocusCoE created regular newsletters[1], organised CoE workshops at key HPC events and highlighted CoE achievements in social media. The set of service offerings by CoEs to industrial and academic users are presented at www.hpccoe.eu/technological-offerings-of-the-eu-hpc-coes-3/ .

H2020 Centre of Excellence

www.hpccoe.eu

Twitter: @FocusCoE
LinkedIn: @focus-coe

COORDINATING ORGANISATION
scapos AG, Germany

OTHER PARTNERS
- CEA - Commissariat à l'Energie Atomique et aux Energies Alternatives, France
- Kuningla Tekniska högskolan (KTH), Sweden
- Höchstleistungsrechenzentrum der Universität Stuttgart (HLRS), Germany
- Barcelona Supercomputing Center (BSC), Spain
- University College London, UK
- Agenzia nazionale per le nuove tecnologie, l'energia e lo sviluppo economico sostenibile (ENEA), Italy
- National University of Ireland, Galway, Ireland
- Teratec, France
- Forschungszentrum Jülich GmbH, Germany
- Partnership for advanced computing in Europe (PRACE), Belgium
- Max Planck Gesellschaft zur Forderung der Wissenschaft eV, Germany

CONTACT
Guy Lonsdale, guy.lonsdale@scapos.com

CALL
INFRAEDI-02-2018

PROJECT TIMESPAN
01/12/2018 - 30/11/2021

1. www.focus-coe.eu/index.php/newsletter/

BioExcel-2

Centre of Excellence for Computational Biomolecular Research

BioExcel is the leading European Centre of Excellence for Computational Biomolecular Research. We develop advanced software, tools and full solutions for high-end and extreme scale computing for biomolecular research. The centre supports academic and industrial research by providing expertise and advanced training to end-users and promoting best practices in the field. The centre aspires to operate in the long-term as the core facility for advanced biomolecular modelling and simulations in Europe.

OVERVIEW

Much of the current Life Sciences research relies on intensive biomolecular modelling and simulation. As a result, both academia and industry are facing significant challenges when applying best practices for optimal usage of compute infrastructures. At the same time, increasing productivity of researchers will be of high importance for achieving reliable scientific results at faster pace.

High-performance computing (HPC) and *high-throughput computing* (HTC) techniques have now reached a level of maturity for many of the widely-used codes and platforms, but taking full advantage of them requires that researchers have the necessary training and access to guidance by experts. The necessary ecosystem of services in the field is presently inadequate. A suitable infrastructure is needed to be set up in a sustainable, long-term operational fashion.

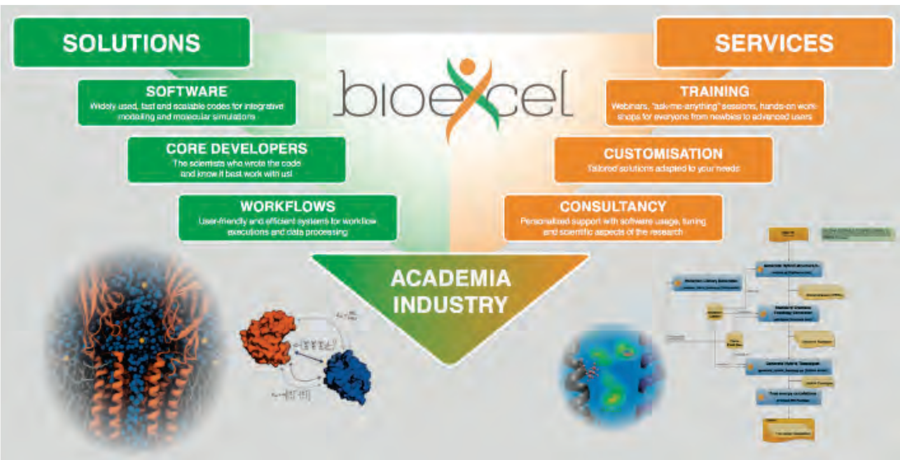

BioExcel CoE was thus established as the go-to provider of a full-range of software applications and training services. These cover fast and scalable software, user-friendly automation workflows and a support base of experts. The main services offered include hands-on training, tailored customization of code and personalized consultancy support. BioExcel actively works with:

1. **academic** and non-profit researchers,
2. **industrial** researchers,
3. **software vendors** and academic **code providers**,
4. non-profit and commercial **resource providers**, and
5. related international **projects and initiatives**.

The centre was established and operates through funding by the EC Horizon 2020 program (Grant agreements 675728 and 823830).

H2020 Centre of Excellence

bioexcel.eu/

Twitter: @BioExcelCoE
LinkedIn: @bioexcel
YouTube: https://goo.gl/5dBzmw

COORDINATING ORGANISATION

KTH Royal Institute of Technology, Sweden

OTHER PARTNERS

- The University of Manchester, UK
- University of Utrecht, the Netherlands
- Institute of Research in Biomedicine (IRB), Spain
- European Molecular Biology Laboratory (EMBL-EBI), UK
- Forschungszentrum Jülich GmbH, Germany
- University of Jyväskylä, Finland
- The University of Edinburgh (EPCC), United Kingdom
- Max Planck Gesellschaft, Germany
- Across Limits, Malta
- Barcelona Supercomputing Centre (BSC), Spain
- Ian Harrow Consulting, UK
- Norman Consulting, Norway
- Nostrum Biodiscovery, Spain

CONTACT

Rossen Apostolov, rossen@kth.se

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/01/2019 - 30/06/2022

ChEESE

Centre of Excellence for Exascale and Solid Earth

ChEESE

ChEESE addresses extreme computing scientific and societal challenges in Solid Earth by harnessing European institutions in charge of operational monitoring networks, tier-0 supercomputing centres, academia, hardware developers and third-parties from SMEs, industry and public-governance. The scientific challenge is to prepare 10 open-source flagship codes to solve Exascale capacity and capability problems on computational seismology, magnetohydrodynamics, physical volcanology, tsunamis, and data analysis and predictive techniques from monitoring earthquake and volcanic activity. The selected codes are periodically audited and optimized at both intranode level (including heterogeneous computing nodes) and internode level on heterogeneous hardware prototypes for the upcoming Exascale architectures, thereby ensuring commitment with a co-design approach. Preparation to Exascale considers also aspects like workflows like data management and sharing, I/O, post-process and visualization. Additionally, ChEESE has devel-

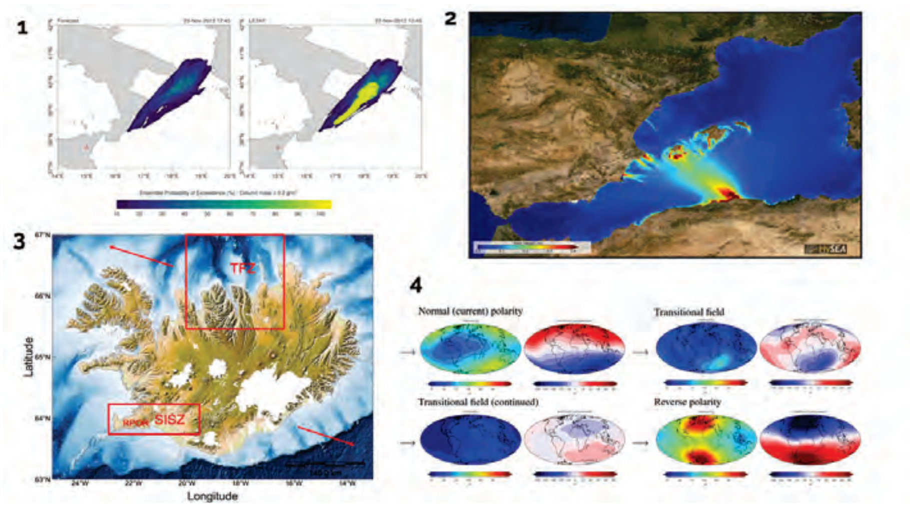

Simulations of (1) volcanic ash dispersal (2) tsunami (3) earthquake and (4) magnetic fields using ChEESE flagship codes.

oped WMS-light, an open source workflow management system for the geoscience community that supports a wide range of typical HPC application scenarios and is easy to use and integrate with existing workflow setups.

In parallel with these transversal activities, ChEESE supports three vertical pillars.
First, it develops pilot demonstrators for challenging scientific problems requiring of Exascale computing in alignment with the vision of European Exascale roadmaps. This includes near real-time seismic simulations and full-wave inversion, ensemble-based volcanic ash dispersal forecasts, faster than real-time tsunami simulations or physics-based hazard assessments for earthquakes, volcanoes and tsunamis.
Second, some pilots are also intended for enabling of operational services requiring of extreme HPC on urgent computing, early warning forecast of geohazards, hazard assessment and data analytics. Pilots with a higher Technology Readiness Level (TRL) will be tested in an operational environment in collaboration with a broader user's community. Additionally, and in collaboration with the European Plate Observing System (EPOS) and Collaborative Data Infrastructure (EUDAT), ChEESE will promote and facilitate the integration of HPC services to widen the access to codes and fostering transfer of know-how to Solid Earth community.
Finally, the third pillar of ChEESE acts as a hub to foster HPC across the Solid Earth Community and related stakeholders and to provide specialized training on services and capacity building measures. Training courses ChEESE has organized include: Modelling tsunamis and volcanic plumes using European flagship codes, Advanced training on HPC for computational seismology and Tools and techniques to quickly improve performances of HPC applications in Solid Earth.

H2020 Centre of Excellence

www.cheese-coe.eu

Twitter: @Cheese_CoE
LinkedIn: @ChEESE-CoE

COORDINATING ORGANISATION

Barcelona Supercomputing Centre (BSC), Spain

OTHER PARTNERS

- Istituto Nazionale di Geofisica e Vulcanologia (INGV), Italy
- Icelandic Meteorological Office (IMO), Iceland
- Swiss Federal Institute of Technology (ETH), Switzerland
- Universität Stuttgart - High Performance Computing Center (HLRS), Germany
- CINECA Consorzio Interuniversitario, Italy
- Technical University of Munich (TUM), Germany
- Ludwig-Maximilians Universität München (LMU), Germany
- University of Málaga (UMA), Spain
- Norwegian Geotechnical Institute (NGI), Norway
- Institut de Physique du Globe de Paris (IPGP), France
- CNRS - Centre National de la Recherche Scientifique, France
- Atos (Bull SAS), France

CONTACT

Arnau Folch, arnau.folch@bsc.es
Rose Gregorio rose.gregorio@bsc.es

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/11/2018 - 31/03/22

CoEC

Centre of Excellence in Combustion

The Centre of Excellence in Combustion (CoEC) is a collaborative effort to exploit Exascale computing technologies to address fundamental challenges encountered in combustion systems. The CoEC vision is aligned with the goals of decarbonization of the European power and transportation sectors and Europe´s vision to achieve net-zero greenhouse gas emissions (GHG) by 2050. The CoEC initiative is a contribution of the European HPC combustion community to a long-term GHG emissions reduction in accordance with the Paris Agreement. The consortium is composed of leading European institutions in computational combustion and HPC and promotes a core of scientific and technological activities aiming to extend the state-of-the-art in combustion simulation capabilities through advanced methodologies enabled by Exascale computing. These advances will help increase the TRL of representative codes from the EU combustion community and increase the EU competitiveness and leadership in power and propulsion technologies.

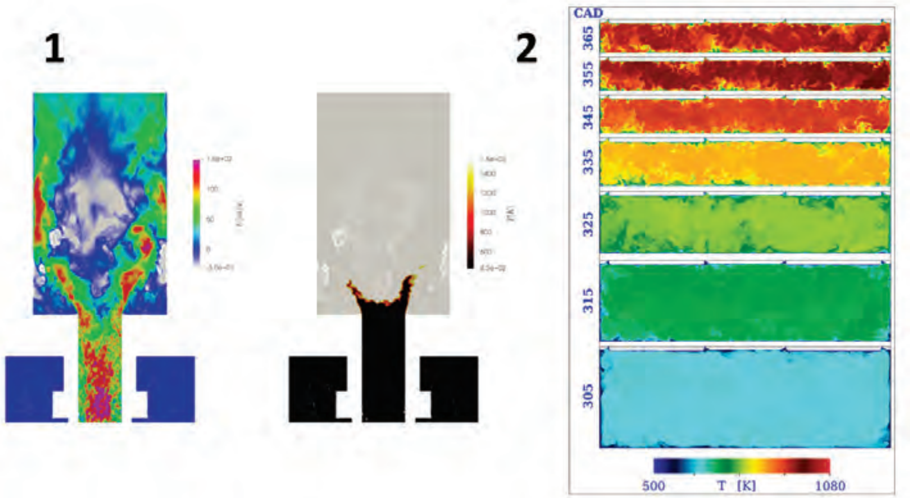

Figure 1. Large eddy simulation of hydrogen. Figure 2. Simulation of turbulent flow using NekRS code.

The outcomes of the project will be the basis of a service portfolio to support the academic and industrial activities in the power and transportation sectors. The Centre will also stimulate the consolidation of the European combustion community, by combining the knowledge of representative actors from the scientific and industrial sectors with the capacity of HPC community including research institutions and supercomputer centres. CoEC will promote the usage of combustion simulations through dissemination, training and user support services.

CONCEPT AND APPROACH

The technical activities integrated into CoEC are defined from the fundamental challenges of relevant application areas where combustion is involved. These areas include transportation (aviation and automotive), energy and power generation, aerospace, fire and safety. All these sectors share common challenges that are associated to scientific and technological problems related with combustion simulation; e.g: instabilities, emissions, noise, plasma, nanoparticles, multiphase flow, autoignition or the use of alternative fuels in practical applications. These research topics have large tradition in combustion science and fundamental progress has been made in the last decade to obtain further understanding on these phenomena using numerical simulations. However, the degree of integration in industry still remains low either because the incomplete fidelity/accuracy of the models, or due to insufficient computing power that hinders the maturity of the technology to be deployed with the level of industrial standards. Exascale computing technologies will enable a new paradigm in physical and numerical modelling, which requires a profound revision of models. For this purpose, two important building blocks regarding Simulation Technologies, and Exascale Methodologies will also be incorporated in the project.

H2020 Centre of Excellence

www.coec-project.eu

Twitter: @CoEC_EU
LinkedIn: @coec-eu

COORDINATING ORGANISATION

Barcelona Supercomputing Centre (BSC), Spain

OTHER PARTNERS

- CERFACS - Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique, France
- Rheinisch-Westfälische Technische Hochschule Aachen (RWTH-AACHEN), Germany
- Technische Universiteit Eindhoven, Netherlands
- University of Cambridge, United Kingdom
- CNRS - Centre National de la Recherche Scientifique, France
- Technische Universität Darmstadt, Germany
- ETHZ - Eidgenössische Technische Hochschule Zürich, Switzerland
- Aristotle University of Thessaloniki , Greece
- Forschungszentrum Jülich GmbH, Germany
- National Centre for Supercomputing Applications, Bulgaria

CONTACT

Daniel Mira, daniel.mira@bsc.es

CALL

INFRAEDI-05-2020

PROJECT TIMESPAN

01/10/2020 – 30/09/2023

CompBioMed2

Centre of Excellence in Computational Biomedicine

CompBioMed is a user-driven Centre of Excellence (CoE) in Computational Biomedicine, designed to nurture and promote the uptake and exploitation of high performance computing within the biomedical modelling community. Our user communities come from academia, industry and clinical practice.

The first phase of the CompBioMed CoE has already achieved notable successes in the development of applications, training and efficient access mechanisms for using HPC machines and clouds in computational biomedicine. We have brought together a growing list of HPC codes relevant for biomedicine which have been enhanced and scaled up to larger machines. Our codes (such as Alya, HemeLB, BAC, Palabos and HemoCell) are now running on several of the world's fastest supercomputers and investigating challenging applications ranging from defining clinical biomarkers of arrhythmic risk to the impact of mutations on cancer treatment.

Our work has provided the ability to integrate clinical datasets with HPC simulations through fully working computational pipelines designed to provide clinically relevant patient-specific models. The reach of the project beyond the funded partners is manifested by our highly effective Associate Partner Programme (all of whom have played an active role in our activities) with cost free, light-weight joining mechanism, and an Innovation Exchange Programme that has brought upward of thirty scientists, industrialists and clinicians into the project from the wider community.

Furthermore, we have developed and implemented a highly successful training programme, targeting the full range of medical students, biomedical engineers, biophysics, and computational scientists. This programme contains a mix of tailored courses for specific groups, webinars and winter schools, which is now being packaged into an easy-to-use training package.

In CompBioMed2 we are extending the CoE to serve the community for a total of 7 years. CompBioMed has established itself as a hub for practitioners in the field, successfully nucleating a substantial body of research, education, training, innovation and outreach within the nascent field of Computational Biomedicine. This emergent technology will enable clinicians to develop and refine personalised medicine strategies ahead of their clinical delivery to the patient. Medical regulatory authorities are currently embracing the prospect of using *in silico* methods in the area of clinical trials and we intend to be in the vanguard of this activity, laying the groundwork for the application of HPC-based Computational Biomedicine approaches to a greater number of therapeutic areas. The HPC requirements of our users are as diverse as the communities we represent. We support both monolithic codes, potentially scaling to the exascale, and complex workflows requiring support for advanced execution patterns. Understanding the complex outputs of such simulations requires both rigorous uncertainty quantification and the embrace of the convergence of HPC and high-performance data analytics (HPDA). CompBioMed2 seeks to combine these approaches with the large, heterogeneous datasets from medical records and from the experimental laboratory to underpin clinical decision support systems. CompBioMed2 will continue to support, nurture and grow our community of practitioners, delivering incubator activities to prepare our most mature applications for wider usage, providing avenues that will sustain CompBioMed2 well-beyond the proposed funding period.

H2020 Centre of Excellence

www.compbiomed.eu

Twitter: @bio_comp
LinkedIn: @compbiomed

COORDINATING ORGANISATION
University College London, UK

OTHER PARTNERS
- University of Amsterdam, Netherlands
- The University of Edinburgh, United Kingdom
- Barcelona Supercomputing Centre (BSC), Spain
- SURFsara, Netherlands
- University of Oxford, UK
- University of Geneva, Switzerland
- The University of Sheffield, UK
- CBK Sci Con Ltd, UK
- Universitat Pompeu Fabra, Spain
- Bayerische Akademie der Wissenschaften, Germany
- Acellera, Spain
- Evotec Ltd, UK
- Atos (Bull SAS), France
- Janssen Pharmaceutica, Belgium
- University of Bologna, Italy

CONTACT
Emily Lumley, e.lumley@ucl.ac.uk

CALL
INFRAEDI-02-2018

PROJECT TIMESPAN
01/10/2019 - 30/09/2023

E-CAM

The E-CAM Centre of Excellence is an e-infrastructure for software development, training and industrial discussion in simulation and modelling. E-CAM is based around the Centre Européen de Calcul Atomique et Moléculaire (CECAM) distributed network of simulation science nodes, and the physical computational e-infrastructure of PRACE. We are a partnership of **13 CECAM Nodes**, **4 PRACE** centres, 12 industrial partners and one Centre for Industrial Computing (Hartree Centre).

E-CAM aims at

- **Developing software** to solve important simulation and modelling problems in industry and academia, with applications from drug development, to the design of new materials
- **Tuning those codes to run on HPC**, through application co-design and the provision of HPC oriented libraries and services;
- **Training scientists from industry and academia**;
- **Supporting industrial end-users** in their use of simulation and modelling, via workshops and direct discussions with experts in the CECAM community.

Our approach is focused on four scientific areas, critical for HPC simulations relevant to key societal and industrial challenges. These areas are classical molecular dynamics, electronic structure, quantum dynamics and meso- and multi-scale modelling. E-CAM develops new scientific ideas and transfers them to algorithm development, optimisation, and parallelization in these four respective areas, and delivers the related training. Postdoctoral researchers are employed under each scientific area, working closely with the scientific programmers to create, oversee and implement the different software codes, in collaboration with our industrial partners.

So far E-CAM has

- Certified more than 250 software modules that are open access and easily available for the industrial and academic communities through our software repository;
- Trained about 500 people at our Extended Software Development Workshops in advanced computational methods, good practices for code development, documentation and maintenance;
- Deployed an online training infrastructure to support the development of software for extreme-scale hardware, contributed to the LearnHPC platform for scalable HPC training;

- Worked on 10 pilot projects in direct collaboration with our industrial partners;
- Organized 10 State of the Art workshops and 6 Scoping workshops, that brought together industrialists, software developers and academic researchers, to discuss the challenges they face;
- Worked on software development projects that enable an HPC practice with potential for transfer into industry. Examples include: GPU re-write of the DL_MESO code for mesoscale simulations; the development of a load-balancing library specifically targeting MD applications, and of an HTC library capable of handling thousands of tasks each requiring peta-scale computing; among other efforts;
- Developed original dissemination material for the general public, including a Comics issue presented at Festivals and Schools;
- Ensured sustainability of its main outputs via community effort, memberships, and the CECAM network.

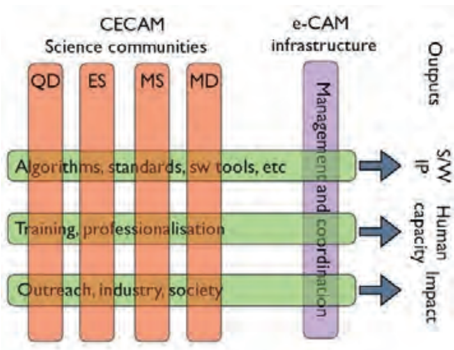

Positioning of the E-CAM infrastructure

H2020 Centre of Excellence

www.e-cam2020.eu

Twitter: @ECAM2020
LinkedIn: @e-cam

COORDINATING ORGANISATION

Ecole Polytechnique Fédérale de Lausanne, Switzerland

OTHER PARTNERS

- University College Dublin, Ireland
- Freie Universität Berlin, Germany
- Universita degli Studi di Roma La Sapienza, Italy
- CNRS - Centre National de la Recherche Scientifique, France
- Technische Universität Wien, AT
- University of Cambridge, UK
- Max-Planck-Institut für Polymerforschung, Germany
- ENS - École Normale Supérieure de Lyon, France
- Forschungszentrum Jülich, Germany
- Universitat de Barcelona, Spain
- Daresbury Laboratory, Scientific and Technology Facilities Council, UK
- Scuola Internazionale Superiore Di Trieste, Italy
- Universiteit van Amsterdam, Netherlands
- Scuola Normale Superiore Pisa, Italy
- Aalto University, FI
- CSC-IT Centre for Science Ltd, Finland
- Irish Centre for High-End Computing, Ireland

CONTACT

Luke Drury (Project Chair)
ld@cp.dias.ie
Ignacio Pagonabarraga (Project Technical Manager) ignacio.pagonabarraga@epfl.ch

CALL

EINFRA-5-2015

PROJECT TIMESPAN

01/10/2015 - 31/03/2021

EoCoE-II

Energy Oriented Centre Of Excellence, Phase II

The Energy-oriented Centre of Excellence (EoCoE) applies cutting-edge computational methods in its mission to accelerate the transition to the production, storage and management of clean, decarbonized energy. EoCoE is anchored in the High Performance Computing (HPC) community and targets research institutes, key commercial players and SMEs who develop and enable energy-relevant numerical models to be run on exascale supercomputers, demonstrating their benefits for low-carbon energy technology. The present project draws on a successful proof-of-principle phase of EoCoE-I, where a large set of diverse computer applications from four such energy domains achieved significant efficiency gains thanks to its multidisciplinary expertise in applied mathematics and supercomputing. During this 2nd round, EoCoE-II channels its efforts into 5 scientific Exascale challenges in the low-carbon sectors of Energy Meteorology, Materials, Water, Wind and Fusion. This multidisciplinary effort harnesses innovations in computer science and mathematical algorithms within a tightly integrated co-design approach to overcome performance bottlenecks and to anticipate future HPC hardware developments. A world-class consortium of 18 complementary partners from 7 countries forms a unique network of expertise in energy science, scientific computing and HPC, including 3 leading European supercomputing centres. New modelling capabilities in selected energy sectors will be created at unprecedented scale, demonstrating the potential benefits to the energy industry, such as accelerated design of storage devices, high-resolution probabilistic wind and solar forecasting for the power grid and quantitative understanding of plasma core-edge interactions in ITER-scale tokamaks. These flagship applications will provide a high-visibility platform for high-performance computational energy science, cross-fertilized through close working connections to the EERA and EUROfusion consortia.

EoCoE is structured around a central Franco-German hub coordinating a pan-European network, gathering a total of 7 countries and 21 teams. Its partners are strongly engaged in both the HPC and energy fields. The primary goal of EoCoE is to create a new, long lasting and sustainable community around computational energy science. EoCoE resolves current bottlenecks in application codes; it develops cutting-edge mathematical and numerical methods, and tools to foster the usage of Exascale computing. Dedicated services for laboratories and industries are established to leverage this expertise and to develop an ecosystem around HPC for energy.

We are interested in collaborations in the area of HPC (e.g., programming models, exascale architectures, linear solvers, I/O) and also with people working in the energy domain and needing expertise for carrying out ambitious simulation.
See our service page for more details: http://www.eocoe.eu/services

H2020 Centre of Excellence

www.eocoe.eu

LinkedIn: @hpc-energy

COORDINATING ORGANISATION

Maison de la Simulation (MdlS) and Institut de Recherche sur la Fusion par confinement Magnétique (IRFM) at CEA, France

OTHER PARTNERS

- Forschungszentrum Jülich (FZJ), with RWTH Aachen University, Germany
- Agenzia nazionale per le nuove tecnologie, l'energia e lo sviluppo economico sostenibile (ENEA), Italy
- Barcelona Supercomputer Centre (BSC), Spain
- Centre National de la Recherche Scientifique (CNRS) with Inst. Nat. Polytechnique Toulouse (INPT), France
- Institut National de Recherche en Informatique et Automatique (INRIA), France
- Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS), France
- Max-Planck Gesellschaft (MPG), Germany
- Fraunhofer Gesellschaft, Germany
- Friedrich-Alexander Univ. Erlangen-Nürnberg (FAU), Germany
- Consiglio Nazionale delle Ricerche (CNR), with Univ. Rome, Tor Vergata (UNITOV), Italy
- Università degli Studi di Trento (UNITN), Italy
- Instytut Chemii Bioorganicznej Polskiej Akademii Nauk (PSNC), Poland
- Université Libre de Bruxelles (ULB), Belgium
- University of Bath (UBAH), United Kingdom
- Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Spain
- IFP Energies Nouvelles (IFPEN), France
- DDN Storage, France

CONTACT

Edouard Audit, contact@eocoe.eu

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/01/2019 - 31/12/2021

ESiWACE2

Centre of Excellence in Simulation of Weather and Climate in Europe, Phase 2

esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER AND CLIMATE IN EUROPE

ESiWACE2 will leverage the potential of the envisaged EuroHPC pre-exascale systems to push European climate and weather codes to world-leading spatial resolution for production-ready simulations, including the associated data management and data analytics workflows. For this goal, the portfolio of climate models supported by the project has been extended with respect to the prototype-like demonstrators of ESiWACE1. Besides, while the central focus of ESiWACE2 lies on achieving scientific performance goals with HPC systems that will become available within the next four years, research and development to prepare the community for the systems of the exascale era constitutes another project goal.

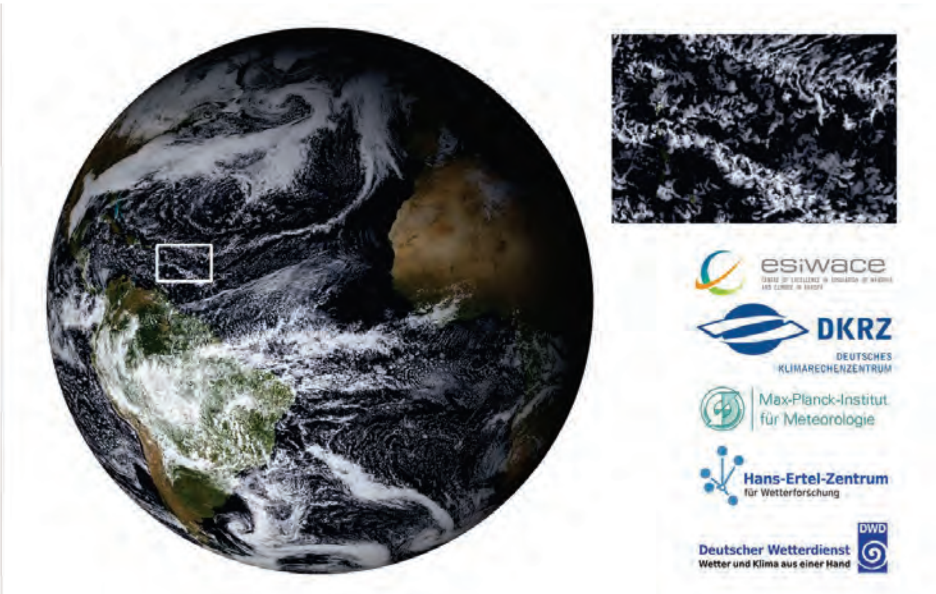

The illustration above shows clouds simulated by the ICON model at 2.5 km resolution in a DYAMOND Winter simulation. The inset shows the Barbados region, where DYAMOND Winter connects to the EUREC4A measurement campaign. The background image is created by NASA.

These developments will be complemented by the establishment of new technologies such as domain-specific languages and tools to minimise data output for ensemble runs. Co-design between model developers, HPC manufacturers and HPC centres is to be fostered, in particular through the design and employment of High-Performance Climate and Weather benchmarks, with the first version of these benchmarks feeding into ESiWACE2 through the FET-HPC research project ESCAPE-2. Additionally, training and dedicated services will be set up to enable the wider community to efficiently use upcoming pre-exascale and exascale supercomputers.

H2020 Centre of Excellence

www.esiwace.eu

Twitter: @ESIWACE

COORDINATING ORGANISATION

Deutsches Klimarechenzentrum (DKRZ), Germany

OTHER PARTNERS

- CNRS - Centre National de la Recherche Scientifique, France
- ECMWF - European Centre for Medium-Range Weather Forecasts, United Kingdom
- Barcelona Supercomputing Center (BSC), Spain
- Max-Planck-Institut für Meteorologie, Germany
- Sveriges meteorologiska och hydrologiska institut, Sweden
- CERFACS - Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique, France
- National University of Ireland Galway (Irish Centre for High End Computing), Ireland
- Met Office, United Kingdom
- Fondazione Centro Euro-Mediterraneo sui Cambiamenti Climatici, Italy
- The University of Reading, United Kingdom
- STFC - Science and Technology Facilities Council, United Kingdom
- Atos (Bull SAS), France
- Seagate Systems UK Limited, United Kingdom
- ETHZ - Eidgenössische Technische Hochschule Zürich, Switzerland
- The University of Manchester, United Kingdom
- Netherlands eScience Centre, Netherlands
- Federal Office of Meteorology and Climatology, Switzerland
- DataDirect Networks, France
- Mercator Océan, France

CONTACT

esiwace@dkrz.de

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/01/2019 - 31/12/2022

EXCELLERAT

Engineering applications will be among the first exploiting Exascale. Not only in academia but also in industry. In fact, the industrial engineering field is *the* field with the highest Exascale potential. Thus the goal of EXCELLERAT is to enable the European engineering industry to advance towards Exascale technologies and to create a single entry point to services and knowledge for all stakeholders (industrial end users, ISVs, technology providers, HPC providers, academics, code developers, engineering experts) of HPC for engineering. In order to achieve this goal, EXCELLERAT brings together key players from industry, research and HPC to provide all necessary services. This is in line with the European HPC Strategy as implemented through the EuroHPC Joint Undertaking.

Core competences of the EXCELLERAT Centre of Excellence

To fulfil its mission, EXCELLERAT focuses its developments on six carefully chosen reference applications (Nek5000, Alya, AVBP, TPLS, FEniCS, Coda), which were analysed on their potential to support the aim to achieve Exascale performance in HPC for Engineering. Thus, they are promising candidates to act as showcases for the evolution of applications towards execution on Exascale Demonstrators, Pre-Exascale Systems and Exascale Machines.

EXCELLERAT addresses the setup of a Centre as an entity, acting as a single hub, covering a wide range of issues, from «non-pure-technical» services such as access to knowledge or networking up to technical services as e.g. co-design, scalability enhancement or code porting to new (Exa) Hardware.

As the consortium contains key players in HPC, HPDA, AI and experts for the reference applications, impact (e.g. code improvements and awareness raising) is guaranteed. The scientific excellence of the EXCELLERAT consortium enables evolution, optimisation, scaling and porting of applications towards disruptive technologies and increases Europe's competitiveness in engineering. Within the frame of the project, EXCELLERAT will prove the applicability of the results not only for the six chosen reference applications, but even going beyond. EXCELLERAT extends the recipients of its developments beyond the consortium and interacts via diverse mechanisms to integrate external stakeholders of its value network into its evolution.

The EXCELLERAT team is developing solutions that support the advances towards Exascale technologies. These achievements are reflected in a collection of success stories that are available online.

H2020 Centre of Excellence

www.excellerat.eu
EXCELLERAT services: services.excellerat.eu

Twitter: @EXCELLERAT_CoE
LinkedIn: @excellerat
SoundCloud: soundcloud.com/user-61965686

COORDINATING ORGANISATION

Universität Stuttgart (HLRS), Germany

OTHER PARTNERS

- The University of Edinburgh (EPCC), United Kingdom
- CINECA Consorzio Interuniversitario, Italy
- SICOS BW GmbH, Germany
- Kungliga Tekniska Hoegskolan (KTH), Sweden
- ARCTUR DOO, Slovenia
- DLR - Deutsches Zentrum für Luft- und Raumfahrt EV , Germany
- Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS), France
- Barcelona Supercomputing Centre (BSC), Spain
- SSC-Services GmbH, Germany
- Fraunhofer Gesellschaft zur Förderung der Angewandten Forschung E.V., Germany
- TERATEC, France
- Rheinisch-Westfälische Technische Hochschule Aachen, Germany

CONTACT

Dr.-Ing. Bastian Koller, koller@hlrs.de

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/12/2018 - 28/02/2022

HiDALGO

HPC and Big Data Technologies for Global Challenges

OBJECTIVE

Developing evidence and understanding concerning Global Challenges and their underlying parameters is rapidly becoming a vital challenge for modern societies. Various examples, such as health care, the transition of green technologies or the evolution of the global climate up to hazards and stress tests for the financial sector demonstrate the complexity of the involved systems and underpin their interdisciplinary as well as their globality. This becomes even more obvious if coupled systems are considered: problem statements and their corresponding parameters are dependent on each other, which results in interconnected simulations with a tremendous overall complexity. Although the process for bringing together the different communities has already started within the Centre of Excellence for Global Systems Science (CoeGSS), the importance of assisted decision making by addressing global, multidimensional problems is more important than ever.

Global decisions with their dependencies cannot be based on incomplete problem assessments or gut feelings anymore, since impacts cannot be foreseen without an accurate problem representation and its systemic evolution. Therefore, HiDALGO bridges that shortcoming by enabling highly accurate simulations, data analytics and data visualisation, but also by providing technology as well as knowledge on how to integrate the various workflows and the corresponding data.

H2020 Centre of Excellence

hidalgo-project.eu/

Twitter: @EU_HiDALGO

COORDINATING ORGANISATION

Atos Spain, Spain

OTHER PARTNERS

- Universität Stuttgart – High Performance Computing Center Stuttgart, Germany
- Instytut Chemii Bioorganicznej Polskiej Akademii Nauk – Poznan Supercomputing and Networking Center, Poland
- ICCS (Institute of Communication and Computer Systems), Greece
- Brunel University London, United Kingdom
- Know Center GmbH, Austria
- Szechenyi Istvan University, Hungary
- Paris-Lodron-Universität Salzburg, Austria
- European Centre for Medium-Range Weather Forecasts, United Kingdom
- MoonStar Communications GmbH, Germany
- DIALOGIK Gemeinnützige Gesellschaft für Kommunikations- und Kooperationsforschung MBH, Germany
- Magyar Kozut Nonprofit Zrt., Hungary
- ARH Informatikai Zrt., Hungary

CONTACT

Francisco Javier Nieto (francisco.nieto@atos.net)

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/12/2018 - 30/11/2021

MaX

MAterials design at the eXascale

MaX - Materials design at the eXascale is a Centre of Excellence with focus on driving the evolution and exascale transition of materials science codes and libraries, and creating an integrated ecosystem of codes, data, workflows, and analysis tools for high-performance (HPC) and high-throughput computing (HTC). Particular emphasis is on co-design activities to ensure that future HPC architectures are well suited for the materials science applications and their users.

The focus of MaX is on first principles materials science applications, i.e., codes that allow predictive simulations of materials and their properties from the laws of quantum physics and chemistry, without resorting to empirical parameters.

The exascale perspective is expected to boost the massive use of these codes in designing materials structures and functionalities for research and manufacturing.

MaX works with code developers and experts from HPC centres to support such transition.
It focuses on selected complementary open-source codes: QUANTUM ESPRESSO, SIESTA, Yambo, FLEUR, CP2K, BigDFT. In addition, it contributes to the development of the AiiDA materials informatics infrastructure and the Materials Cloud environment.

The main actions of the project include:

1. code and library restructuring: including modularizing and adapting for heterogeneous architectures, as well as adoption of new algorithms;
2. co-design: working on codes and providing feedback to architects and integrators; developing workflows and turn-key solutions for properties calculations and curated data sharing;
3. enabling convergence of HPC, HTC, and high-performance data analytics;
4. addressing the skills gap: widening access to codes, training, and transferring know-how to user communities;
5. serving end-users in industry and research: providing dedicated services, support, and high-level consulting.

Results and lessons learned on MaX flagship codes and projects are made available to the developers and user communities at large.

H2020 Centre of Excellence

www.max-centre.eu

Twitter: @max_center2
LinkedIn: @max-centre

COORDINATING ORGANISATION

CNR Nano, Modena, Italy (Elisa Molinari)

OTHER PARTNERS

- SISSA Trieste, Italy (Stefano Baroni)
- ICN2 Barcelona, Spain (Pablo Ordejón)
- FZJ Jülich, Germany (Stefan Blügel, Dirk Pleiter)
- EPFL Lausanne, Switzerland (Nicola Marzari)
- CINECA Bologna, Italy (Fabio Affinito)
- Barcelona Supercomputing Centre (BSC), Spain (Stephan Mohr)
- CSCS ETH Zürich, Switzerland (Joost VandeVondele)
- CEA Grenoble, France (Thierry Deutsch)
- UGent, Belgium (Stefaan Cottenier)
- E4 Computer Engineering, Italy (Fabrizio Magugliani)
- Arm, UK (Conrad Hillairet)
- ICTP UNESCO, Trieste, Italy (Ivan Girotto)
- Trust-IT Pisa, Italy (Silvana Muscella)

CONTACT

Elisa Molinari, Director
elisa.molinari@max-centre.eu

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/12/2018 - 31/05/2022

NOMAD

Novel Materials Discovery

Predicting novel materials with specific desirable properties is a major aim of ab initio computational materials science (aiCMS) and an urgent requirement of basic and applied materials science, engineering and industry. Such materials can have immense impact on the environment and on society, e.g., on energy, transport, IT, medical-device sectors and much more. Currently, however, precisely predicting complex materials is computationally infeasible.

NOMAD CoE will develop a new level of materials modelling, enabled by upcoming HPC exascale computing and extreme-scale data hardware.

In close contact with the R&D community, including industry, we will

- develop exascale algorithms to create accurate predictive models of real-world, industrially-relevant, complex materials;
- provide exascale software libraries for all code families (not just selected codes); enhancing today's aiCMS to take advantage of tomorrow's HPC computing platforms;
- develop energy-to-solution as a fundamental part of new models. This will be achieved by developing novel artificial-intelligence (AI) guided work-flow engines that optimise the modelling calculations and provide significantly more information per calculation performed;
- offer extreme-scale data services (data infrastructure, storage, retrieval and analytics/AI);
- test and demonstrate our results in two exciting use cases, solving urgent challenges for the energy and environment that cannot be computed properly with today's hard- and software (water splitting and novel thermoelectric materials);
- train the next generation of students, also in countries where HPC studies are not yet well developed.

NOMAD CoE is working closely together with POP, and it is synergistically complementary to and closely coordinated with the EoCoE, ECAM, BioExcel and MaX CoEs. NOMAD CoE will push the limits of aiCMS to unprecedented capabilities, speed and accuracy, serving basic science, industry and thus society.

H2020 Centre of Excellence

www.nomad-coe.eu

Twitter: @NoMaDCoE

COORDINATING ORGANISATION

Max-Planck-Gesellschaft zur Förderung der Wissenschaften EV, Germany

OTHER PARTNERS

- Humboldt-Universität zu Berlin, Germany
- Danmarks Tekniske Universitet, Denmark
- Technische Universität Wien, Austria
- The University of Warwick, United Kingdom
- Université Catholique de Louvain, Belgium
- Barcelona Supercomputing Centre (BSC), Spain
- CSC – IT Center for Science Ltd., Finland
- University of Cambridge, United Kingdom
- CEA - Commissariat à l'Energie Atomique et aux Energies Alternatives, France
- Aalto University Helsinki, Finland
- University of Latvia, Latvia

CONTACT

contact@nomad-coe.eu

CALL

INFRAEDI-05-2020

PROJECT TIMESPAN

01/10/2020 – 30/09/2023

PerMedCoE

HPC / Exascale Centre of Excellence in Personalised Medicine

Personalised Medicine (PerMed) opens unexplored frontiers to treat diseases at the individual level combining clinical and omics information. However, the performances of the current simulation software are still insufficient to tackle medical problems such as tumour evolution or patient-specific treatments. The challenge is to develop a sustainable roadmap to scale-up the essential software for the cell-level simulation to the new European HPC/Exascale systems. Simulation of cellular mechanistic models are essential for the translation of omic data to medical relevant actions and these should be accessible to the end-users in the appropriate environment of the PerMed-specific big confidential data.

The goal of the HPC/Exascale Centre of Excellence in Personalised Medicine (PerMedCoE) is to provide an efficient and sustainable entry point to the HPC/Exascale-upgraded methodology to translate omics analyses into actionable models of cellular functions of medical relevance. It will accomplish so by

1. optimising four core applications for cell-level simulations to the new pre-exascale platforms;
2. integrating PerMed into the new European HPC/Exascale ecosystem, by offering access to HPC/Exascale-adapted and optimised software;
3. running a comprehensive set of PerMed use cases;
4. building the basis for the sustainability of the PerMedCoE by coordinating PerMed and HPC communities, and reaching out to industrial and academic end-users, with use cases, training, expertise, and best practices.

The PerMedCoE cell-level simulations will fill the gap between the molecular- and organ-level simulations from the CompBioMed and BioExcel CoEs with which this proposal is aligned at different levels. It will connect methods' developers with HPC, HTC and HPDA experts (at POP and HiDALGO CoEs). Finally, the PerMedCoE will work with biomedical consortia (i.e. ELIXIR, LifeTime initiative) and pre-exascale infrastructures (BSC and CSC), including a substantial co-design effort.

H2020 Centre of Excellence

permedcoe.eu

Twitter: @PerMedCoE
LinkedIn: @permedcoe

COORDINATING ORGANISATION

Barcelona Supercomputing Centre (BSC), Spain

OTHER PARTNERS

- CSC-IT Centre for Science Ltd, Finland
- Université du Luxembourg, Luxembourg
- Institut Curie, France
- Universitätsklinikum Heidelberg, Germany
- IBM Research GmbH, Switzerland
- Atos Spain, Spain
- Kungliga Tekniska Hoegskolan (KTH), Sweden
- European Molecular Biology Laboratory, Germany
- Fundacio Centre de Regulacio Genomica, Spain
- Max Delbrück Centre for Molecular Medicine, Germany
- University of Ljubljana, Slovenia
- ELEM BIOTECH SL, Spain

CONTACT

Alfonso Valencia alfonso.valencia@bsc.es

CALL

INFRAEDI-05-2020

PROJECT TIMESPAN

01/10/2020 – 30/09/2023

POP2

Performance Optimisation and Productivity 2

The growing complexity of parallel computers is leading to a situation where code owners and users are not aware of the detailed issues affecting the performance of their applications. The result is often an inefficient use of the infrastructures. Even in the cases where the need to get further performance and efficiency is perceived, the code developers may not have insight enough on its detailed causes so as to properly address the problem. This may lead to blind attempts to restructure codes in a way that might not be the most productive ones.

POP2 extends and expands the activities successfully carried out by the POP Centre of Excellence since October 2015. The effort in the POP and POP2 projects resulted up to now in close to 400 assessment services provided to customers in academia, research and industry helping them to better understand the behaviour and improve the performance of their applications.

The external view, advice, and help provided by POP have been extremely useful for many of these customers to significantly improve the performance of their codes by factors of 20% in some cases but up to 2x or even 10x in others. The POP experience was also extremely valuable to identify issues in methodologies and tools that if improved will reduce the assessment cycle time.

POP2 continues to focus on a transversal (horizontal) service to offer highly specialized expertise and know-how to all other CoEs facilitating cross-fertilisation between them and other sectors. This could lead to wider access to codes, including specific and targeted measures for industry & SMEs. Collaborations with other CoEs and Projects help to reach potential users for services and promote support to the governance of HPC Infrastructures. Centres could adopt the POP methodology, helping it embed in the wider HPC community, leveraging the use of existing infrastructures.

H2020 Centre of Excellence

pop-coe.eu

Twitter: @POP_HPC

COORDINATING ORGANISATION

Barcelona Supercomputing Centre (BSC), Spain

OTHER PARTNERS

- Universität Stuttgart HLRS, Germany
- Forschungszentrum Jülich GmbH, Germany
- Numerical Algorithms Group Ltd, UK
- Rheinisch-Westfälische Technische Hochschule Aachen (RWTH-AACHEN), Germany
- Teratec, France
- Université de Versailles Saint Quentin, France
- IT4Innovations, VSB – Technical University of Ostrava, Czech Republic

CONTACT

Jesus Labarta, jesus.labarta@bsc.es
Judit Gimenez, judit@bsc.es

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/12/2018 - 31/05/2022

RAISE

Research on AI- and Simulation-Based Engineering at Exascale

Compute- and data-driven research encompasses a broad spectrum of disciplines and is the key to Europe's global success in various scientific and economic fields. The massive amount of data produced by such technologies demands novel methods to post-process, analyse, and to reveal valuable mechanisms. The development of artificial intelligence (AI) methods is rapidly proceeding and they are progressively applied to many stages of workflows to solve complex problems. Analysing and processing big data require high computational power and scalable AI solutions. Therefore, it becomes mandatory to develop entirely new workflows from current applications that efficiently run on future high-performance computing architectures at Exascale.

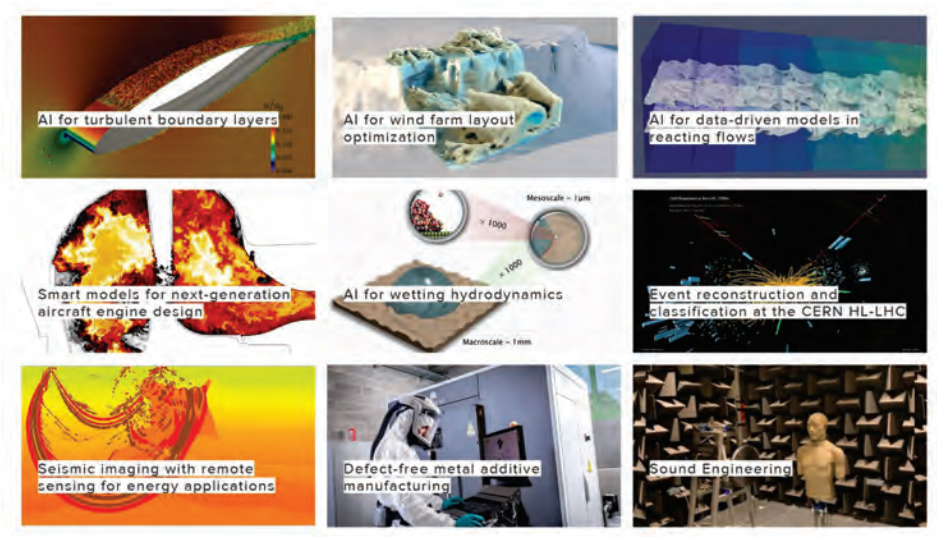

Compute- and data-driven use cases of CoE RAISE

The Centre of Excellence for AI- and Simulation-based Engineering at Exascale (RAISE) will be the excellent enabler for the advancement of such technologies in Europe on industrial and academic levels, and a driver for novel intertwined AI and HPC methods. These technologies will be advanced along representative use-cases, covering a wide spectrum of academic and industrial applications, e.g. coming from wind energy harvesting, wetting hydrodynamics, manufacturing, physics, turbomachinery, and aerospace. It aims at closing the gap in full loops using forward simulation models and AI-based inverse inference models, in conjunction with statistical methods to learn from current and historical data. In this context, novel hardware technologies, i.e. Modular Supercomputing Architectures, Quantum Annealing, and prototypes from the DEEP project series will be used for exploring unseen performance in data processing. Best practices, support, and education for industry, SMEs, academia, and HPC centres on Tier-2 level and below will be developed and provided in RAISE's European network attracting new user communities. This goes along with the development of a business providing new services to various user communities.

H2020 Centre of Excellence

www.coe-raise.eu

Twitter: @CoeRaise
LinkedIn: @coe-raise
Facebook: @CoERAISE2021
YouTube: www.youtube.com

COORDINATING ORGANISATION

Forschungszentrum Jülich GmbH, Germany

OTHER PARTNERS

- Universitet Haskoli Islands, Iceland
- The Cyprus Institute, Cyprus
- RWTH Aachen University, Germany
- Barcelona Supercomputing Centre (BSC), Spain
- European Organisation for Nuclear Research (CERN), Switzerland
- Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS), France
- Atos (BULL SAS), France
- Riga Technical University, Latvia
- Flanders Make, Belgium
- SAFRAN Helicopter Engines, France
- ParTec AG, Germany
- Delft University of Technology (TU Delft), Netherlands

CONTACT

Dr.-Ing. Andreas Lintermann
A.Lintermann@fz-juelich.de

CALL

INFRAEDI-05-2020

PROJECT TIMESPAN

01/01/2021 - 31/12/2023

TREX

Targeting Real chemical accuracy at the Exascale

Targeting Real chemical accuracy at the EXascale

TREX federates European scientists, HPC stakeholders, and SMEs to develop and apply quantum mechanical simulations in the framework of stochastic quantum Monte Carlo methods. This methodology encompasses various techniques at the high-end in the accuracy ladder of electronic structure approaches and is uniquely positioned to fully exploit the massive parallelism of the upcoming exascale architectures. The marriage of these advanced methods with exascale will enable simulations at the nanoscale of unprecedented accuracy, targeting a fully consistent description of the quantum mechanical electron problem.

TREX's main focus will be the development of a user-friendly and open-source software suite in the domain of stochastic quantum chemistry simulations, which integrates TREX community codes within an interoperable, high-performance platform. In parallel, TREX will work on showcases to leverage this methodology for commercial applications as well as develop and implement software components and services that make it easier for commercial operators and user communities to use HPC resources for these applications.

H2020 Centre of Excellence

trex-coe.eu/

Twitter: @trex_eu
LinkedIn: @trex-eu

COORDINATING ORGANISATION

University of Twente, Netherlands

OTHER PARTNERS

- CNRS - Centre National de la Recherche Scientifique, France
- Scuola Internazionale Superiore Di Studi Avanzati Di Trieste, Italy
- CINECA Consorzio Interuniversitario, Italy
- Forschungszentrum Jülich GmbH, Germany
- Max-Planck-Gesellschaft zur Förderung der Wissenschaften EV, Germany
- Université de Versailles Saint Quentin-en-Yvelines, France
- Megware Computer Vertrieb und Service GmbH, Germany
- Slovak Academy of Sciences, Bratislava, Slovakia
- Universität Wien, Austria
- Politechnika Lodzka, Poland
- TRUST-IT SRL, Italy

CONTACT

info@trex-coe.eu
Claudia Filippi, Project Coordinator, c.filippi@utwente.nl

CALL

INFRAEDI-05-2020

PROJECT TIMESPAN

01/10/2020 – 30/09/2023

Applications in the HPC Continuum

ACROSS

HPC Big DAta ArtifiCial Intelligence cross-stack platfoRm tOwardS ExaScale

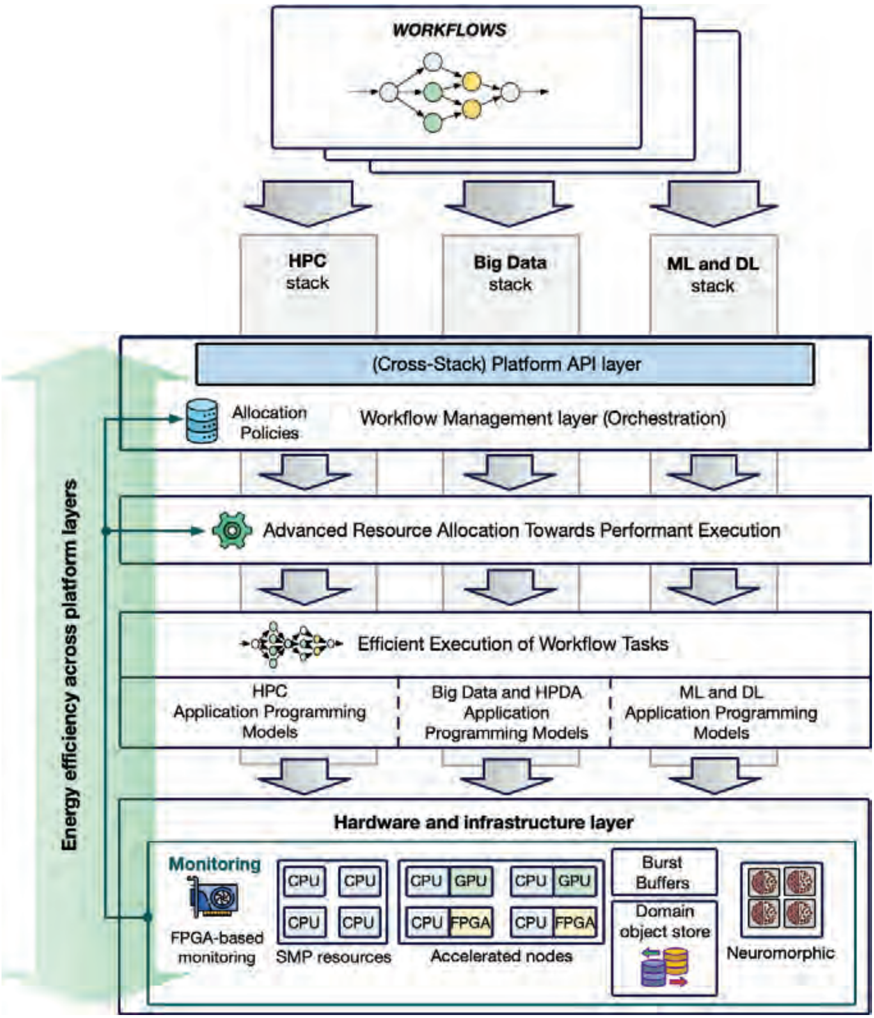

Artificial intelligence and Big Data are going to revolutionize the HPC domain, by offering new tools for analyzing the immense amount of data generated by scientific experiments and modern simulation tools.
To this end, ACROSS aims at developing an innovative execution platform for serving emerging workflows that mix large and complex numerical simulations along with Big Data analytics, machine learning and deep learning techniques. ACROSS will co-design this innovative platform revolving around a large set of advanced hardware and software technologies towards the Exascale level of performance. As such, hardware acceleration technologies will be of primary interest for ACROSS to get the maximum performance while consuming the least amount of energy. In detail, ACROSS will move beyond traditional GPUs, integrating and exploiting FPGAs and Neuromorphic hardware.
Looking forward, ACROSS will envision the adoption of the future European Processor Initiative (EPI) outcomes, also to promote European first-class technology. Energy-efficiency, efficient usage of computing resources and application performance will be achieved through the design and implementation of an innovative workflow and resource management system. As such, a multi-level approach will be used to define the workflows, manage the resource allocation beyond the traditional queuing systems by using a machine learning based approach, efficiently execute workload tasks on the allocated heterogeneous resources. Three pilots have been selected to demonstrate the capabilities of the ACROSS platform in three relevant industrial and scientific domains: (1) aeronautics pilot; (2) weather, climate, hydrological and farming pilot; and (3) energy and carbon sequestration pilot.
The ACROSS platform will be backed by supercomputing and Cloud resources provided by European HPC centres (IT4Innovations and CINECA), including a new generation of supercomputers.

HPC and data centric environments and application platforms

www.acrossproject.eu

Twitter: @across_project
LinkedIn: @acrossproject
Facebook: @acrossprojecteu

COORDINATING ORGANISATION

Fondazione LINKS, Italy

OTHER PARTNERS

- Atos (Bull SAS), France
- IT4Innovations, VSB – Technical University of Ostrava, Czech Republic
- CINECA, Consorzio Interuniversitario, Italy
- GE Avio Aero, Italy
- European Centre For Medium-Range Weather Foreccasts (ECMWF), United Kingdom
- Consorzio Interuniversitario Nazionale per l'Informatica (CINI), Italy
- Institut National de Recherche en Informatique et en Automatique (INRIA), France
- SINTEF AS, Norway
- Neuropublic AE Pliroforikis & Epikoinonion, Greece
- Stichting DELTARES, Netherlands
- Max-Planck-Gesellschaft Zur Forderung Der Wissenschaften EV - MPG, Germany
- Morfo Design SRL , Italy

CONTACT

Olivier Terzo (project coordinator),
olivier.terzo@linksfoundation.com

CALL

EuroHPC-02-2019

PROJECT TIMESPAN

01/03/2021 - 29/02/2024

CYBELE

Fostering precision agriculture and livestock farming through secure access to large-scale HPC-enabled virtual industrial experimentation environment empowering scalable big data analytics

CYBELE generates innovation and creates value in the domain of agri-food, and its verticals in the sub-domains of Precision Agriculture (PA) and Precision Livestock Farming (PLF) in specific. This will be demonstrated by the real-life industrial cases, empowering capacity building within the industrial and research community. CYBELE aims at demonstrating how the convergence of HPC, Big Data, Cloud Computing and the Internet of Things can revolutionise farming, reduce scarcity and increase food supply, bringing social, economic, and environmental benefits.

CYBELE intends to safeguard that stakeholders have integrated, unmediated access to a vast amount of large-scale datasets of diverse types from a variety of sources. By providing secure and unmediated access to large-scale HPC infrastructures supporting data discovery, processing, combination and visualisation services, Stakeholders shall be enabled to generate more value and deeper insights in operations.

CYBELE develops large scale HPC-enabled test beds and delivers a distributed big data management architecture and a data management strategy providing:

1) integrated, unmediated access to large scale datasets of diverse types from a multitude of distributed data sources,
2) a data and service driven virtual HPC-enabled environment supporting the execution of multi-parametric agri-food related impact model experiments, optimising the features of processing large scale datasets and
3) a bouquet of domain specific and generic services on top of the virtual research environment facilitating the elicitation of knowledge from big agri-food related data, addressing the issue of increasing responsiveness and empowering automation-assisted decision making, empowering the stakeholders to use resources in a more environmentally responsible manner, improve sourcing decisions, and implement circular-economy solutions in the food chain.

HPC and Big Data enabled Large-scale Test-beds and Applications

www.cybele-project.eu

Twitter: @CYBELE_H2020
Facebook: @Cybele-H2020
LinkedIn: @Cybele-project

COORDINATING ORGANISATION

Waterford Institute of Technology, Ireland

OTHER PARTNERS

- Barcelona Supercomputing Centre (BSC), Spain
- Atos (Bull SAS), France
- CINECA Consorzio Interuniversitario, Italy
- Instytut Chemii Bioorganicznej Polskiej Akademii Nauk, Poland
- RYAX Technologies, France
- Universität Stuttgart, Germany
- EXODUS Anonymos Etaireia Pliroforikis, Greece
- LEANXCALE SL, Spain
- ICCS (Institute of Communication and Computer Systems), Greece
- Centre for Research and Technology Hellas - Information Technologies Institute (ITI), Greece
- Tampereen Korkeakoulusaatio SR, Finland
- UBITECH, Greece
- University of Piraeus Research Center, Greece
- SUITE5 Data Intelligence Solutions Ltd, Cyprus
- INTRASOFT International SA, Luxembourg
- ENGINEERING - Ingegneria Informatica SPA, Italy
- Wageningen University, Netherlands
- Stichting Wageningen Research, Netherlands
- BIOSENSE Institute, Serbia
- Donau Soja Gemeinnutzige GmbH, Austria
- Agroknow IKE, Greece
- GMV Aerospace and Defence SA, Spain
- Federacion de Cooperativas Agroalimentares de la Comunidad Valenciana, Spain
- University of Strathclyde, United Kingdom
- ILVO - Instituut voor Landbouw- en Visserrijonderzoek, Belgium
- VION Food Nederland BV, Netherlands
- Olokliromena Pliroforiaka Sistimataae, Greece
- Kobenhavns Universitet, Denmark
- EVENFLOW, Belgium
- Open Geospatial Consortium (Europe) Ltd, United Kingdom

CONTACT

cybele.h2020@gmail.com

CALL

ICT-11-2018-2019

PROJECT TIMESPAN

01/01/2019 - 31/12/2021

DeepHealth

Deep-Learning and HPC to Boost Biomedical Applications for Health

The main goal of the DeepHealth project is to put HPC computing power at the service of biomedical applications; and apply Deep Learning (DL) and Computer Vision (CV) techniques on large and complex biomedical datasets to support new and more efficient ways of diagnosis, monitoring and treatment of diseases.

THE DEEPHEALTH TOOLKIT: A KEY OPEN-SOURCE ASSET FOR EHEALTH AI-BASED SOLUTIONS

At the centre of the proposed innovations is the DeepHealth toolkit, an open-source free software that will provide a unified framework to exploit heterogeneous HPC and big data architectures assembled with DL and CV capabilities to optimise the training of predictive models. The toolkit is composed of two core libraries plus a back end offering a RESTful API to other applications to facilitate the use of the libraries, and a dedicated front end that interacts with the back end for facilitating the use of the libraries to data scientists. The libraries are the European Distributed Deep Learning Library (EDDL) and the European Computer Vision Library (ECVL). Version 1.0 of both libraries has been released during second semester of 2021, and both are ready to be integrated in current and new biomedical platforms or applications.

THE DEEPHEALTH CONCEPT – APPLICATION SCENARIOS

The DeepHealth concept focuses on scenarios where image processing is needed for diagnosis. In the training environment IT experts work with datasets of images for training predictive models. In the production environment the medical personnel ingests an image coming from a scan into a platform or biomedical application that uses predictive models to get clues to support them during diagnosis. The DeepHealth toolkit will allow the IT staff to train models and run the training algorithms over hybrid HPC + big data architectures without a profound knowledge of DL, CV, HPC or big data and increase their productivity reducing the required time to do it.

14 PILOTS AND SEVEN PLATFORMS TO VALIDATE THE DEEPHEALTH PROPOSED INNOVATIONS

The DeepHealth innovations are being validated in 14 pilots through the use of seven different biomedical and AI software platforms provided by partners. The libraries are being integrated and validated in seven AI and biomedical software platforms: commercial platforms: (everis Lumen, PHILIPS Open Innovation Platform, THALES PIAF) and research-oriented platforms (CEA's ExpressIFTM, CRS4's Digital Pathology, WINGS MigraineNet).

The use cases cover three main areas: (i) Neurological diseases, (ii) Tumour detection and early cancer prediction and (iii) Digital pathology and automated image annotation. The pilots will allow evaluating the performance of the proposed solutions in terms of the time needed for pre-processing images, the time needed to train models and the time to put models in production. In some cases, it is expected to reduce these times from days or weeks to just hours. This is one of the major expected impacts of the DeepHealth project.

deephealth-project.eu

Twitter: @DeepHealthEU
Facebook: @DeepHealthEU
LinkedIn: @deephealtheu

COORDINATING ORGANISATION

Everis Spain SL, Spain

OTHER PARTNERS

- Azienda Ospedaliera Citta della Salute e della Scienza di Torino, Italy
- Barcelona Supercomputing Centre (BSC), Spain
- Centre Hospitalier Universitaire Vaudois, Switzerland
- Centro di Ricerca, Sviluppo e Studi Superiori in Sardegna SRL, Italy
- CEA - Commissariat à l'Energie Atomique et aux énergies alternatives, France
- EPFL - Ecole Polytechnique Fédérale de Lausanne, Switzerland
- Fundacion para el Fomento de la Investigacion Sanitaria y Biomedica de la Comunitat Valenciana, Spain
- Karolinska Institutet, Sweden
- Otto-von-Guericke-Universität Magdeburg, Germany
- Philips Medical Systems Nederland BV, Netherlands
- Pro Design Electronic GmbH, Germany

- SIVECO Romania SA, Romania
- Spitalul Clinic Prof Dr Theodor Burghele, Romania
- STELAR Security Technology Law Research UG, Germany
- Thales SIX GTS France SAS, France
- TREE Technology SA, Spain
- Università degli Studi di Modena e Reggio Emilia, Italy
- Università degli Studi di Torino, Italy
- Universitat Politècnica de València, Spain
- WINGS ICT Solutions Information & Communication Technologies IKE, Greece

CONTACT

PROJECT MANAGER: Monica Caballero monica.caballero.galeote@nttdata.com
TECHNICAL MANAGER:
Jon Ander Gómez jon@dsic.upv.es

CALL

ICT-11-2018-2019

PROJECT TIMESPAN

01/01/2019 - 30/06/2022

eFlows4HPC

Enabling dynamic and Intelligent workflows in the future EuroHPC ecosystem

Today, developers lack tools that enable the development of complex workflows involving HPC simulation and modelling with data analytics (DA) and machine learning (ML). The eFlows4HPC project aims to deliver a workflow software stack and an additional set of services to enable the integration of HPC simulation and modelling with big data analytics and machine learning in scientific and industrial applications.

The software stack will allow to develop innovative adaptive workflows that efficiently use the computing resources and also innovative data management solutions. In order to attract first-time users of HPC systems, the project will provide HPC Workflows as a Service (HPCWaaS), an environment for sharing, reusing, deploying and executing existing workflows on HPC systems. The project will leverage workflow technologies, machine learning and big data libraries from previous open-source European initiatives. Specific optimizations of application kernels tackling the use of accelerators (FPGAs, GPUs) and the European Processor Initiative (EPI) will be performed. To demonstrate the workflow software stack, use cases from three thematic pillars have been selected.

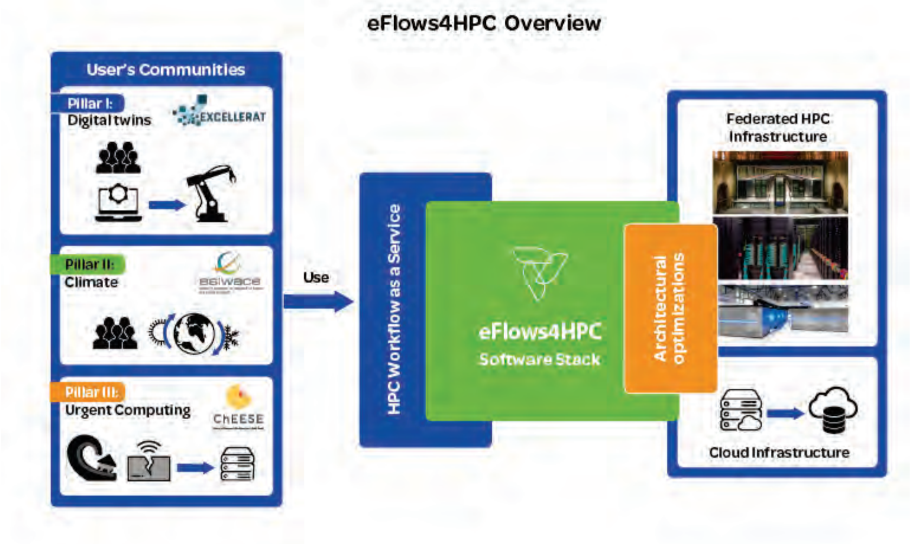

Pillar I focuses on the construction of DigitalTwins for the design of complex manufactured objects integrating state-of-the-art adaptive solvers with machine learning and data-mining, contributing to the Industry 4.0 vision. Pillar II develops innovative adaptive workflows for climate and the study of Tropical Cyclones (TC) in the context of the CMIP6 experiment. Pillar III explores the modelling of natural catastrophes - in particular, earthquakes and their associated tsunamis shortly after such an event is recorded. Leveraging two existing workflows, Pillar III will work on integrating them with the eFlows4HPC software stack and on producing policies for urgent access to supercomputers. The results from the three Pillars will be demonstrated to the target community CoEs to foster adoption and receive feedback.

HPC and data centric environments and application platforms

eflows4hpc.eu

Twitter: @eFlows4HPC
LinkedIn: @eflows4hpc

COORDINATING ORGANISATION

Barcelona Supercomputing Center (BSC), Spain

OTHER PARTNERS

- Centre Internacional de Mètodes Numèrics a l'Enginyeria (CIMNE), Spain
- Forschungszentrum Jülich (FZJ), Germany
- Universidad Politécnica de Valencia (UPV), Spain
- Atos (Bull SAS), France
- DtoK Lab S.r.l., Italy
- Fondazione Centro Euro-Mediterraneo sui Cambiamenti Climatici (CMCC), Italy
- Institut national de recherche en informatique et en automatique (INRIA), France
- Scuola Internazionale Superiore di Studi Avanzati di Trieste (SISSA), Italy
- Instytut Chemii Bioorganicznej Polskiej Akademii Nauk (PSNC), Poland
- Universidad de Málaga (UMA), Spain
- Istituto nazionale di geofisica e vulcanología (INGV), Italy
- Alfred-Wegener-Institut, Helmholtz-Zentrum für Polar- und Meeresforschung (AWI), Germany
- Eidgenössische Technische Hochschule Zürich (ETHZ), Switzerland
- SIEMENS AG, Germany
- Norges Geotekniske Institutt (NGI), Norway

CONTACT

Rosa M. Badia, rosa.m.badia@bsc.es

CALL

EuroHPC-02-2019

PROJECT TIMESPAN

01/01/2021 - 31/12/2023

ENERXICO

Supercomputing and Energy for Mexico

The ENERXICO project is applying exascale HPC techniques to different energy industry simulations of critical interest for Mexico. ENERXICO works to provide solutions for the oil and gas industry, improve wind energy performance and increase the efficiency of biofuels for transportation. Composed of some of the best academic and industrial institutions in the EU and Mexico, ENERXICO demonstrates the power of cross-border collaboration to put supercomputing technologies at the service of the energy sector. The specific objectives of the project are:

- To develop beyond state-of-the-art high performance simulation tools for the energy industry
- To increase oil & gas reserves using geophysical exploration for subsalt reservoirs
- To improve refining and transport efficiency of heavy oil
- To develop a strong wind energy sector to mitigate oil dependency
- To improve fuel generation using biofuels
- To improve the cooperation between industries from the EU and Mexico
- To improve the cooperation between the leading research groups in the EU and Mexico

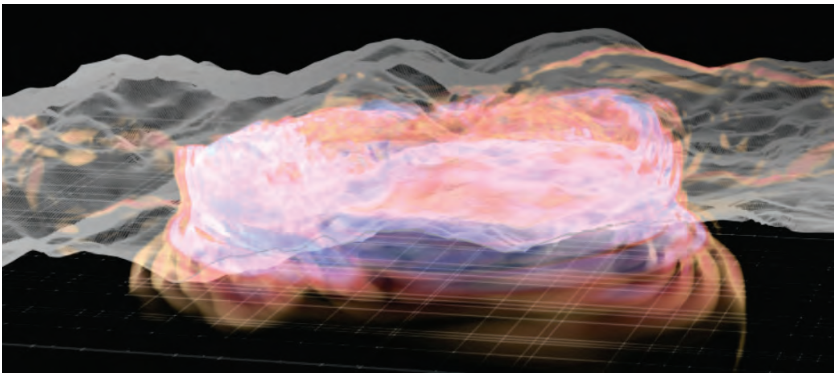

ENERXICO works to optimise numerical simulation codes used in the energy industry for the exascale, such as the Barcelona Subsurface Imaging Tools (BSIT), which simulates seismic wavefields in a complex topography region.

ENERXICO is composed of partners from Mexico, Spain, France and Germany.

International cooperation

www.enerxico-project.eu
Twitter: @ENERXICOproject
LinkedIn: @enerxico-project

COORDINATING ORGANISATIONS

- Barcelona Supercomputing Center (BSC), Spain
- Instituto Nacional de Investigaciones Nucleares (ININ), Mexico

OTHER PARTNERS

- Atos (Bull SAS), France
- Centro de Investigación y de Estudios Avanzados del Instituto Politécnico Nacional (CINVESTAV) , Mexico
- Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas (CIEMAT), Spain
- Instituto Mexicano del Petróleo (IMP) , Mexico
- Instituto Politécnico Nacional (IPN-ESIA), Mexico
- Technische Universität München (TUM), Germany
- Universidad Autónoma Metropolitana (UAM-A), Mexico
- Instituto de Ingeniería de la Universidad Nacional Autónoma de México (UNAM-IINGEN), Mexico
- Universitat Politècnica de València (UPV), Spain
- Université Grenoble Alpes (UGA), France
- IBERDROLA , Spain
- REPSOL SA, Spain

PROJECT OBSERVER

Petróleos Mexicanos (PEMEX), Mexico

CONTACT

Rose Gregorio rose.gregorio@bsc.es
Mireia Cos mireia.cospique@bsc.es

CALL

FETHPC-01-2018

PROJECT TIMESPAN

01/06/2019 - 31/08/2021

EVOLVE

HPC and Cloud-enhanced Testbed for Extracting Value from Diverse Data at Large Scale

Leading the Big Data Revolution

EVOLVE is a pan European Innovation Action with 19 key partners from 10 European countries introducing important elements of High-Performance Computing (HPC) and Cloud in Big Data platforms taking advantage of recent technological advancements to enable cost-effective applications in 7 different pilots to keep up with the unprecedented data growth we are experiencing. EVOLVE aims to build a large-scale testbed by integrating technology from:

The HPC world: An advanced computing platform with HPC features and systems software.

The Big Data world: A versatile big-data processing stack for end-to-end workflows.

The Cloud world: Ease of deployment, access, and use in a shared manner, while addressing data protection.

EVOLVE aims to take concrete and decisive steps in bringing together the Big Data, HPC, and Cloud worlds, and to increase the ability to extract value from massive and demanding datasets. EVOLVE aims to bring the following benefits for processing large and demanding datasets:

- Performance: Reduced turn-around time for domain-experts, industry (large and SMEs), and end-users.
- Experts: Increased productivity when designing new products and services, by processing large datasets.
- Businesses: Reduced capital and operational costs for acquiring and maintaining computing infrastructure.
- Society: Accelerated innovation via faster design and deployment of innovative services that unleash creativity.

EVOLVE is now reaching its final stretch, as the projects ends in November 2021.

- At this stage the hardware platform is fully integrated with heterogeneity at the processing and storage level. Even if this part of the work plan is achieved, some infrastructures enhancements are still scheduled, such as the installation of nodes with the latest generation of GPU and the addition of local NVMe to some computing nodes.
- Thanks to our implementation of a cooperation between Kubernetes with Slurm, the software stack is also completed, with the ability to run complex containerized workflows even if they include MPI stages.
- Most of the focus is now on the pilot applications integration and the ecosystem development. The results collected on the fully ported pilot applications are more than encouraging: with speed-up in the range of 1 to 2 order of magnitudes. More importantly similar gains are observed on the manageability of the system, (even though they are harder to

quantify). HPC level of performance is reached while keeping DevOps Cloud oriented approached.

Therefore, the central promise of EVOLVE delivering performance while keeping the user experience at the centre will be fulfilled.

HPC and Big Data enabled Large-scale Test-beds and Applications

www.evolve-h2020.eu

Twitter: @Evolve_H2020
LinkedIn: @evolve-h2020

COORDINATING ORGANISATION

Datadirect Networks France, France

OTHER PARTNERS

- AVL LIST GmbH, Austria
- BMW - Bayerische Motoren Werke AG, Germany
- Atos (Bull SAS), France
- Cybeletech, France
- Globaz S.A., Portugal
- IBM Ireland Ltd, Ireland
- Idryma Technologias Kai Erevnas, Greece
- ICCS (Institute of Communication and Computer Systems), Greece
- KOOLA d.o.o., Bosnia and Herzegovina
- Kompetenzzentrum - Das Virtuelle Fahrzeug, Forschungsgesellschaft mbH, Austria
- MEMEX SRL, Italy
- MEMOSCALE AS, Norway
- NEUROCOM Luxembourg SA, Luxembourg
- ONAPP Ltd, Gibraltar
- Space Hellas SA, Greece
- Thales Alenia Space France SAS, France
- TIEMME SPA, Italy
- WEBLYZARD Technology GmbH, Austria

CONTACT

Catarina Pereira catarina@loba.pt
Jean-Thomas Acquaviva
jtacquaviva@ddn.com

CALL

ICT-11-2018-2019

PROJECT TIMESPAN

01/12/2018 - 30/11/2021

EXA MODE

EXtreme-scale Analytics via Multimodal Ontology Discovery & Enhancement

Exascale volumes of diverse data from distributed sources are continuously produced. Healthcare data stand out in the size produced (production 2020 >2000 exabytes), heterogeneity (many media, acquisition methods), included knowledge (e.g. diagnostic reports) and commercial value. The supervised nature of deep learning models requires large labelled, annotated data, which precludes models to extract knowledge and value.

ExaMode solves this by allowing easy & fast, weakly supervised knowledge discovery of heterogeneous exascale data provided by the partners, limiting human interaction. Its objectives include the development and release of extreme analytics methods and tools that are adopted in decision making by industry and hospitals. Deep learning naturally allows to build semantic representations of entities and relations in multimodal data. Knowledge discovery is performed via semantic document-level networks in text and the extraction of homogeneous features in heterogeneous images. The results are fused, aligned to medical ontologies, visualized and refined. Knowledge is then applied using a semantic middleware to compress, segment and classify images and it is exploited in decision support and semantic knowledge management prototypes.

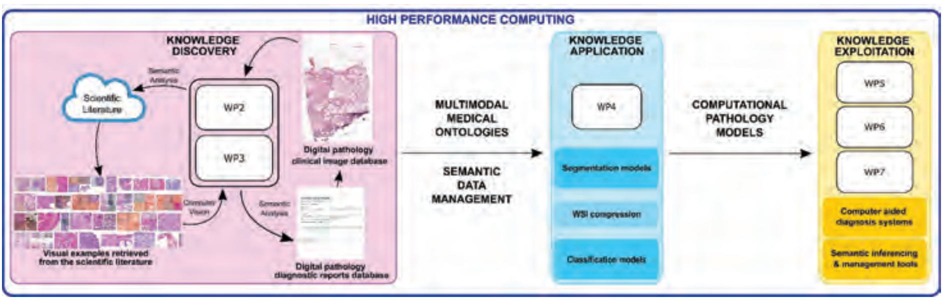

ExaMode is relevant to ICT12 in several aspects:

1. Challenge: it extracts knowledge and value from heterogeneous quickly increasing data volumes.
2. Scope: the consortium develops and releases new methods and concepts for extreme scale analytics to accelerate deep analysis also via data compression, for precise predictions, support decision making and visualize multi-modal knowledge.
3. Impact: the multi-modal/media semantic middleware makes heterogeneous data management & analysis easier & faster, it improves architectures for complex distributed systems with better tools increasing speed of data throughput and access, as resulting from tests in extreme analysis by industry and in hospitals.

Big Data technologies and extreme-scale analytics

www.examode.eu

Twitter: @examode
LinkedIn: **@examode**
Facebook: @examode.eu

COORDINATING ORGANISATION
University of Applied Sciences Western Switzerland, HES-SO, Sierre, Switzerland

OTHER PARTNERS
- Azienda Ospedaliera per l'Emergenza Cannizzaro, Italy
- MicroscopeIT Sp. z o.o., Poland
- SIRMA AI, Bulgaria
- Radboud University Medical Center, Nijmegen, Netherlands
- SURFsara BV, Netherlands
- Università degli Studi di Padova, Italy

CONTACT
Dr. Manfredo Atzori
manfredo.atzori@hevs.ch
Prof. Henning Müller
henning.mueller@hevs.ch

CALL
ICT-12-2018-2020

PROJECT TIMESPAN
01/01/2019 - 31/12/2022

exaFOAM

Exploitation of Exascale Systems for Open-Source Computational Fluid Dynamics by Mainstream Industry

Computational Fluid Dynamics (CFD) has become a mature technology in engineering design, contributing strongly to industrial competitiveness and sustainability across a wide range of sectors (e.g. transportation, power generation, disaster prevention). Future growth depends upon the exploitation of massively parallel HPC architectures, however this is currently hampered by performance scaling bottlenecks.

The ambitious exaFOAM project aims to overcome these limitations through the development and validation of a range of algorithmic improvements. Improvements across the entire CFD process chain (preprocessing, simulation, I/O, post-processing) will be developed. Effectiveness will be demonstrated via a suite of HPC Grand Challenge and Industrial Application Challenge cases. All developments will be implemented in the open-source CFD software OpenFOAM, one of the most successful open-source projects in the area of computational modelling, with a large industrial and academic user base.

To ensure success, the project mobilises a highly-capable consortium of 12 beneficiaries consisting of experts in HPC CFD algorithms and industrial applications and includes universities, HPC centres, SMEs and code release authority OpenCFD Ltd (openfoam.com) as a linked third party to the PI. Project management will be facilitated by a clear project structure and quantified objectives enable tracking of the project progress.

Special emphasis will be placed on ensuring a strong impact of the exaFOAM project. The project has been conceived to address all expected impacts set out in the Work Programme. All developed code and validation cases will be released as open-source to the community in coordination with the OpenFOAM Governance structure. The involvement of 17 industrial supporters and stakeholders from outside the consortium underscores the industrial relevance of the project outcomes. A well-structured and multichannelled plan for dissemination and exploitation of the project outcomes further reinforces the expected impact.

Industrial software codes for extreme scale computing environments and applications

exafoam.eu

COORDINATING ORGANISATION

ESI Group, France

OTHER PARTNERS

- CINECA, Consorzio Interuniversitario, Italy
- E4 Computer Engineering SpA, Italy
- Politecnico di Milano, Italy
- Fakultet strojarstva i brodogradnje, Sveučilište u Zagrebu, Croatia
- Upstream CFD GmbH, Germany
- Technische Universität Darmstadt, Germany
- Universität Stuttgart, Germany
- Barcelona Supercomputing Centre (BSC), Spain
- Wikki Gesellschaft für numerische Kontinuumsmechanik mbH, Germany
- National Technical University of Athens - NTUA, Greece
- Universidade do Minho, Portugal

CONTACT

Fred Mendonca,
fred.mendonca@esi-group.com

CALL

EuroHPC-03-2019

PROJECT TIMESPAN

01/04/2021 - 31/03/2024

EXSCALATE4COV

EXaSCale smArt pLatform Against paThogEns for Corona Virus

The EXSCALATE4CoV (E4C) project aims to exploit the most powerful computing resources currently based in Europe to empower smart in-silico drug design. Advanced Computer-Aided Drug Design (CADD) in combination with the high throughput bio- chemical and phenotypic screening allows the rapid evaluation of the simulations results and the reduction of time for the discovery of new drugs.

Against a pandemic crisis, the immediate identification of effective treatments has a paramount importance. First, E4C selected through the Dompé EXSCALATE platform, the most promising commercialized and developing drugs safe in man. Second, selected from >500 billion molecules new pan coronavirus inhibitors.

Huge computational resources are needed, therefore the activities are supported and empowered by four of the most powerful computer centres in Europe: CINECA, BSC, JÜLICH and Eni HPC5. The Swiss Institute of Bioinformatics (SIB) provides the homology 3D models for the viral proteins. The Fraunhofer IME provides the BROAD Repurposing Library and biochemical assays for the most relevant viral proteins. Phenotypic screenings have been run by KU LUEVEN to identify molecules capable of blocking virus replication in in-vitro models. IIMCB and ELECTRA determines the crystal structure of at least one coronavirus functional proteins to evaluate the structural similarities with other viral proteins.

All the other partners support the project by studying the mechanism of action, synthetizing the most relevant candidate drugs and enhancing the HPC and AI infrastructure of the consortium. The worldwide biggest virtual screening simulation was performed by evaluating more than 1 Trillion molecular docking. Chelonia leads the communication and dissemination activities.

EXSCALATE4CoV consortium identified safe in man drugs repurposed as 2019-nCoV antiviral and proposed to the EMA innovation task force (ITF) to define a preliminary development strategy and a proposal for a registration path. The E4C project shared promptly its scientific outcomes with the research community by using established channels: ChEMBL portal for the biochemical data, the SWISS-MODEL portal for the homology models of viral proteins WT and mutants, the Protein Data Bank, the EUDAT for the data generated by in-silico simulations and the E4C project website.

Combining HPC simulations and AI driven system pharmacology E4C has identified and in vitro experimentally validated a first promising drug for the treatment of mildly symptomatic Covid19 patients. Raloxifene Phase II clinical trial was completed in July 2021. The Spallanzani hospital is the Coordinating Investigator in the clinical trial. Several initiatives have been launched in the framework of E4C as well as MEDIATE also supported by SAS, SPIKE MUTANTS, VIRALSEQ and DRUGBOX. All of them are accessible through the main web portal.

Fight against COVID

www.exscalate4cov.eu

LinkedIn: @exscalate4cov
Twitter: @exscalate4cov
Facebook: @exscalate4cov

COORDINATING ORGANISATION

Dompé farmaceutici S.p.A., Italy

OTHER PARTNERS

- CINECA Consorzio Interuniversitario, Italy
- Politecnico di Milano, Italy
- Università degli Studi di Milano, Italy
- KU Leuven, Belgium
- International Institute of Molecular and Cell Biology, Poland
- Elettra Sincrotrone Trieste, Italy
- Fraunhofer-Institut für Molekularbiologie und Angewandte Oekologie IME
- Barcelona Supercomputing Center, Spain
- Forschungszentrum Jülich GmbH, Germany
- Universita degli Studi di Napoli Federico II, Italy
- Università degli Studi di Cagliari, Italy
- SIB Institute Suisse de Bioinformatique, Switzerland
- Kungliga Tekniska Hoegskolan (KTH), Sweden
- Istituto Nazionale per le Malattie Infettive Lazzaro Spallanzani, Italy
- Associazione Big Data, Italy
- Istituto Nazionale di Fisica Nucleare (INFN), Italy
- Chelonia SA, Switzerland

CONTACT

info@exscalate4cov.eu

CALL

SC1-PHE-CORONAVIRUS-2020

PROJECT TIMESPAN

06/04/2020 – 05/10/2021

HEROES

Hybrid Eco Responsible Optimized European Solution

Bridging the gap between HPC & AI/ML user communities and HPC centres is key to unleash Europe's innovation potential. A lot of effort is done to build the European technologies able to deliver centralised, petascale/exascale HPC & ML. Fostered by increasing access from the edge of the network, it is equally important to make such resources easily and responsibly consumable.

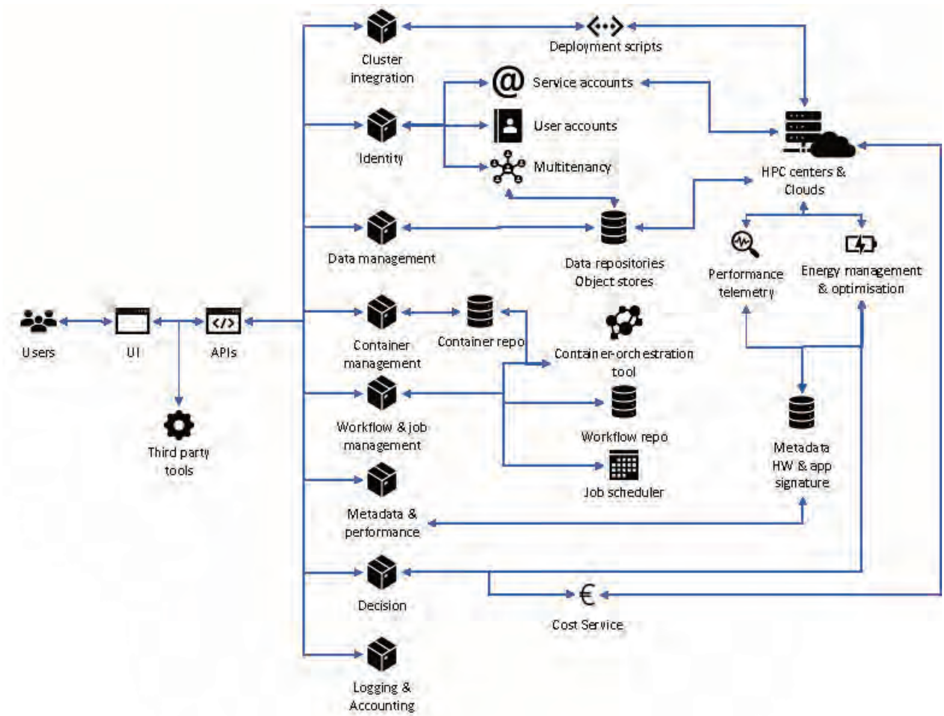

Modules of the HEROES' architecture. End-users will be able to submit workflows to multiple HPC Centres and Cloud Providers, while taking informed decisions on the placement of the tasks based on their performance, energy and costs constraints.

The HEROES project aims at developing an innovative software solution addressing both the industrial and scientific user communities. The platform will allow end users to submit their complex simulation and ML workflows to both HPC data centres and Cloud infrastructures and provide them with the possibility to choose the best option to achieve their goals in time, within budget and with the best energy efficiency.

Although the project will be able, in the future, to provide value creation across multiple industrial sectors, the initial focus for the demonstrator will address workflows of strategic importance in the field of "Renewable Energy" and in "Manufacturing" applications where HPC is involved for the design of more energy efficient products (for example in the design of energy efficient vehicles).

HEROES major innovations reside in its platform selection decision module and its application of marketplace concepts to HPC. The consortium involves 4 European SMEs which bring HPC to their clients and are facing everyday this market demand. A major Supercomputing Research Centre complements the project with its specific expertise in energy management and resource optimisation. HEROES is supported by Meteosim, UL Renewables, EDF and Dallara which will advise on workflows relevant to such use cases. At the end of the project, its outcomes will be commercialised and further reinforce Europe's capacity to lead and nurture innovation in HPC.

HPC and data centric environments and application platforms

heroes-project.eu

Twitter: @heroes_hpc

COORDINATING ORGANISATION

UCit, France

OTHER PARTNERS

- Neovia Innovation, France
- HPCNow! CONSULTING SL, Spain
- Do IT Systems S.r.l., Italy
- Barcelona Supercomputing Centre (BSC), Spain

CONTACT

Philippe Bricard,
Philippe.bricard@ucit.fr

CALL

EuroHPC-02-2019

PROJECT TIMESPAN

01/03/2021 - 28/02/2023

HPCWE

High performance computing for wind energy

Wind energy has become increasingly important as a clean and renewable alternative to fossil fuels in the energy portfolios of both Europe and Brazil. At almost every stage in wind energy exploitation ranging from wind turbine design, wind resource assessment to wind farm layout and operations, the application of HPC is a must.

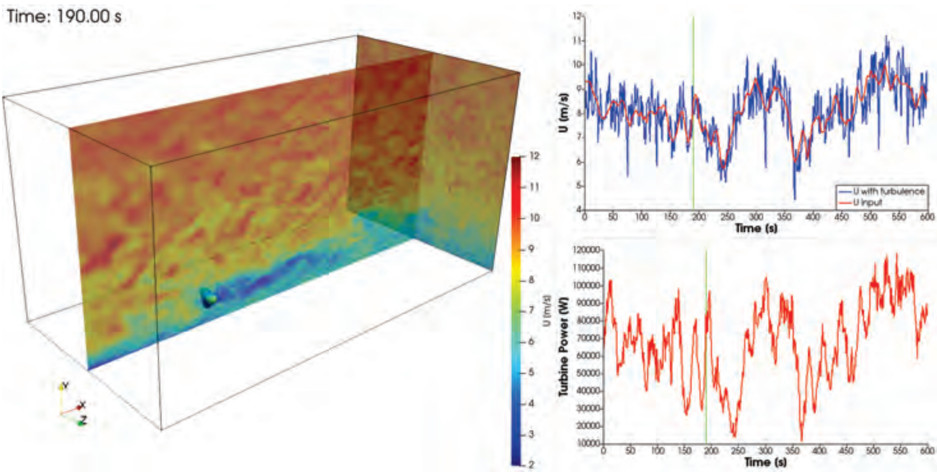

Simulation of the power generated by a wind turbine, which responds non-linearly to the incoming turbulent wind speed

The goal of HPCWE is to address the following key open challenges in applying HPC on wind energy:

i. Efficient use of HPC resources in wind turbine simulations, via the development and implementation of novel algorithms. This leads to the development of methods for verification, validation, uncertainty quantification (VVUQ) and in-situ scientific data interpretation.
ii. Accurate integration of meso-scale atmosphere dynamics and micro-scale wind turbine flow simulations, as this interface is the key for accurate wind energy simulations. In HPCWE a novel scale integration approach will be applied and tested through test cases in a Brazil wind farm.
iii. Adjoint-based optimization, which implies large I/O consumption as well as storing data on large-scale file systems. HPCWE research aims at alleviating the bottlenecks caused by data transfer from memory to disk.

The HPCWE consortium consists of 11 partners representing the top academic institutes, HPC centres and industries in Europe and Brazil. By exploring this collaboration, this consortium will develop novel algorithms, implement them in state-of-the-art codes and test the codes in academic and industrial cases to benefit the wind energy industry and research in both Europe and Brazil.

International Cooperation

hpcwe-project.eu

Twitter: @HPCWE_project
LinkedIn: @hpcwe

COORDINATING ORGANISATION

The University of Nottingham, United Kingdom

OTHER PARTNERS

- Danmarks Tekniske Universitet, Denmark
- Electricité de France (EDF), France
- Imperial College of Science Technology and Medicine, United Kingdom
- The University of Edinburgh (EPCC), United Kingdom
- Universidade de Sao Paulo, Brazil
- Universidade Estadual de Campinas, Brazil
- Universidade Federal de Santa Catarina, Brazil
- Universität Stuttgart, Germany
- Universiteit Twente, Netherlands
- VORTEX Factoria de Calculs SL, Spain

CONTACT

Prof. Atanas Popov (Project Coordinator)
Atanas.Popov@nottingham.ac.uk

CALL

FETHPC-01-2018

PROJECT TIMESPAN

01/06/2019 - 30/11/2021

LEXIS

Large-scale EXecution for Industry & Society

The ever-increasing quantity of data generated by modern industrial and business processes poses an enormous challenge for organisations seeking to glean knowledge and understanding from the data. Combinations of HPC, Cloud and Big Data technologies are key to meeting the increasingly diverse needs of large and small organisations alike, however, access to powerful computing platforms for SMEs which has been difficult due to both technical and financial considerations may now be possible.

The LEXIS (Large-scale EXecution for Industry & Society) project is building an advanced engineering platform at the confluence of HPC, Cloud and Big Data. To run application workflows, including Big Data analytics and simulations, it leverages large-scale geographically-distributed resources from the existing

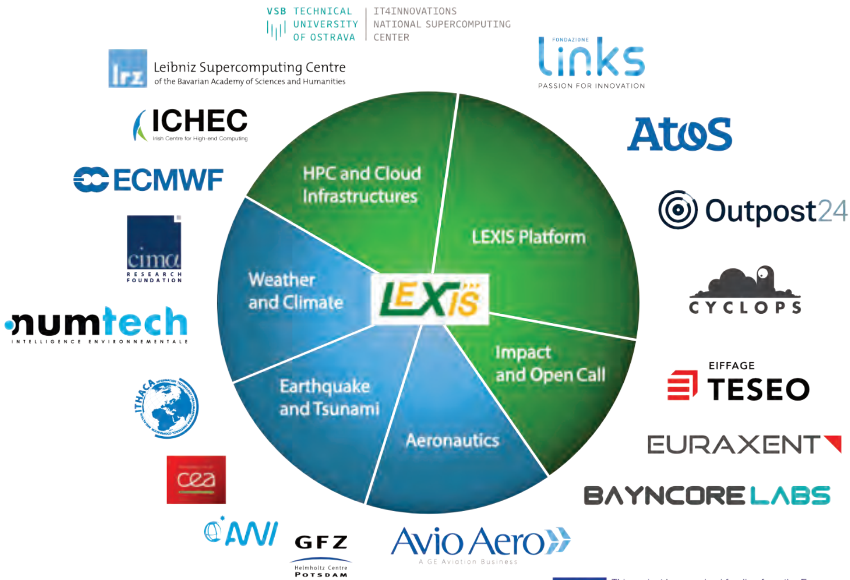

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 825532.

HPC infrastructure and augments them with Cloud services. Driven by the requirements of several pilot test cases, the LEXIS platform relies on best-in-class data management solutions (EUDAT) and advanced, distributed orchestration solutions (TOSCA), augmenting them with new, efficient hardware and platform capabilities (e.g., in the form of Data Nodes and federation, usage monitoring and accounting/billing support). Thus, LEXIS realises an innovative solution aimed at stimulating the interest of the European industry and at creating an ecosystem of organisations that benefit from the LEXIS platform and its well-integrated HPC, HPDA and Data Management solutions.

The consortium is thus developing a demonstrator with a significant Open Source dimension, including validation, tests and documentation. Platform validation is in progress with large-scale pilots – applications from industrial and scientific sectors (Aeronautics, Earthquake and Tsunami, Weather and Climate), where significant improvements in KPIs including job execution time and solution accuracy are anticipated. To further stimulate adoption of the LEXIS framework and increase stakeholders' engagement with the targeted pilots, the project consortium has organised an Open Call for platform validation, attracting more use cases.

LEXIS is promoting its platform to the HPC, Cloud and Big Data sectors, maximising impact through targeted and qualified communication. The project consortium brings together partners with the skills and experience to deliver a complex multi-faceted project, spanning a range of complex technologies across seven European countries, including large industry, flagship HPC centres, industrial and scientific compute pilot users, technology providers and SMEs.

HPC and Big Data enabled Large-scale Test-beds and Applications

lexis-project.eu

Twitter: @LexisProject
Facebook: @LEXIS2020
LinkedIn: @lexis-project
YouTube: bit.ly/youtube_LEXIS

COORDINATING ORGANISATION
IT4Innovations, VSB – Technical University of Ostrava, Czech Republic

OTHER PARTNERS
- Alfred Wegener Institute Helmholtz Centre for Polar and Marine Research, Germany
- Atos (Bull SAS), France
- Bayerische Akademie der Wissenschaften / Leibniz Rechenzentrum der BAdW, Germany
- Bayncore Labs Ltd, Ireland
- CEA, Commissariat à l'Energie Atomique et aux énergies alternatives, France
- CYCLOPS Labs GmbH, Switzerland
- Centro Internazionale in Monitoraggio Ambientale - Fondazione CIMA, Italy
- ECMWF, European Centre for Medium Range Weather Forecasts, United Kingdom
- EURAXENT, United Kingdom
- GE AVIO SRL, Italy
- Helmholtz Zentrum Potsdam Deutsches Geoforschungs Zentrum GFZ, Germany
- Irish Centre for High-End Computing ICHEC, National University of Ireland Galway, Ireland
- Information Technology for Humanitarian Assistance Cooperation and Action, Italy
- LINKS Foundation, Italy
- NUMTECH SARL, France
- OUTPOST24, France
- Teseo Eiffage Energie Systèmes, Italy

CONTACT
Jan Martinovic jan.martinovic@vsb.cz

CALL
ICT-11-2018-2019

PROJECT TIMESPAN
01/01/2019 - 30/12/2021

LIGATE

LIgand Generator and portable drug discovery platform AT Exascale

Today digital revolution is having a dramatic impact on the pharmaceutical industry and the entire healthcare system. The implementation of machine learning, extreme scale computer simulations, and big data analytics in the drug design and development process offer an excellent opportunity to lower the risk of investment and reduce the time to patent and time to patient.

LIGATE aims to integrate and co-design best in class European open-source components together with European Intellectual Properties (whose development has already been co-funded by previous Horizon 2020 projects). It will support Europe to keep worldwide leadership on Computer-Aided Drug Design (CADD) solutions, exploiting today's high-end supercomputers and tomorrow's exascale resources, while fostering the European competitiveness in this field. The project will enhance the CADD technology of the drug discovery platform EXSCALATE.

The proposed LIGATE solution enables to deliver the result of a drug design campaign with the highest speed along with the highest accuracy. This predictability, together with the fully automation of the solution and the availability of the exascale system, will let run the full in silico drug discovery campaign in less than one day to respond promptly for example to worldwide pandemic crisis. The platform will also support European projects in repurposing drugs, natural products and nutraceuticals with therapeutic indications to answer high unmet medical needs like rare, metabolic, neurological and cancer diseases, and emerging ones as new non infective pandemics.

Since the evolution of HPC architectures is heading toward specialization and extreme heterogeneity, including future exascale architectures, the LIGATE solution focuses also on code portability with the possibility to deploy the CADD platform on any available type of architecture in order not to have a legacy in the hardware.

The project plans to make the platform available and open to support the discovery a novel treatment to fight virus infections and multidrug-resistant bacteria. The project will also make available to the research community the outcome of a final simulation.

Industrial software codes for extreme scale computing environments and applications

www.ligateproject.eu

Twitter: @ligateproject
LinkedIn: @ligateproject
Facebook: @ligateproject

COORDINATING ORGANISATION

Dompé farmaceutici S.p.A., Italy

OTHER PARTNERS

- Politecnico di Milano, Italy
- CINECA, Consorzio Interuniversitario Italy
- Kungliga Tekniska Högskolan, Sweden
- Università degli Studi di Salerno (UNISA), Italy
- Universität Innsbruck, Austria
- E4 Computer Engineering SpA, Italy
- Chelonia SA, Switzerland
- tofmotion GmbH, Austria
- IT4Innovations, VSB – Technical University of Ostrava, Czech Republic
- Universität Basel, Switzerland

CONTACT

info@ligateproject.eu

CALL

EuroHPC-03-2019

PROJECT TIMESPAN

01/01/2021 - 31/12/2023

MICROCARD

Numerical modelling of cardiac electrophysiology at the cellular scale

The MICROCARD project will develop an exascale simulation platform to study the mechanisms of cardiac arrhythmia with models that represent the heart cell by cell.

Cardiac arrhythmia, which are among the most frequent causes of death, are disorders of the electrical synchronization system of the heart. Numerical models of this complex system are already highly sophisticated and widely used, but to match observations in aging and diseased hearts they need to move from a continuum approach to a representation of individual cells and their interconnections. This implies a different, harder numerical problem and a 10,000-fold increase in problem size, and consequently a transition from tier 1 to exascale systems. Moreover, the project aims to limit its climate impact by developing energy-efficient code, and operates in a domain where numerical simulation is all but commonplace.

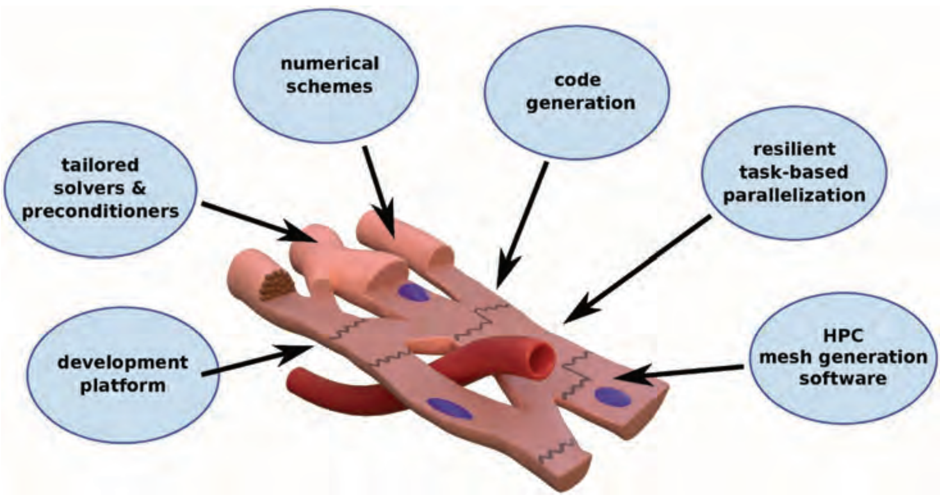

To meet these challenges the platform is co-designed by HPC experts, numerical scientists, biomedical engineers, and biomedical scientists, from academia and industry. We develop numerical schemes suitable for exascale parallelism together with matching linear-system solvers and preconditioners, and a dedicated compiler to translate high-level model descriptions into optimized, energy-efficient system code for heterogeneous computing systems. The code will be resilient to hardware failures and will use an energy-aware task placement strategy.

At the same time, we develop highly parallel mesh generation software to build the extremely large and complex models that we will need for our simulations. These meshes will be based on large volumes of confocal microscopy data that were recently acquired with the latest clearing and labelling techniques.

Our simulation platform will be applied to real-life use cases and will be made accessible for a wide range of users. It will be adaptable to similar biological systems such as nerves, and some of the components and techniques will be reusable in a wide range of applications.

MICROCARD is coordinated by the IHU Liryc, an institute dedicated to research on cardiac arrhythmia that is part of the University of Bordeaux (France) and its university hospital. The project also has strong ties with the Institute for Experimental Cardiovascular Medicine in Freiburg (Germany), and several other clinical and experimental research groups. Together with an end-user advisory board composed of prominent experimental and clinical cardiologists these links will help us to have a real impact on cardiology research and ultimately on patient care.

HPC and data centric environments and application platforms

microcard.eu

COORDINATING ORGANISATION

Université de Bordeaux, France

OTHER PARTNERS

- Université de Strasbourg, France
- Simula Research Laboratory AS, Norway
- Università degli Studi di Pavia, Italy
- USI: Università della Svizzera italiana, Switzerland
- KIT - Karlsruhe Institute for Technology , Germany
- Konrad-Zuse-Zentrum für Informationstechnik Berlin, Germany
- MEGWARE Computer Vertrieb und Service GmbH, Germany
- NumeriCor GmbH, Austria
- Orobix srl, Italy

CONTACT

Mark Potse, mark.potse@inria.fr (scientific coordinator)

CALL

EuroHPC-02-2019

PROJECT TIMESPAN

01/04/2021 - 30/09/2024

NextSim

CODA: Next generation of industrial aerodynamic simulation code

NextSim partners, as fundamental European players in Aeronautics and Simulation, recognise that there is a need to increase the capabilities of current Computational Fluid Dynamics tools for aeronautical design by re-engineering them for extreme-scale parallel computing platforms.

The backbone of NextSim is centred on the fact that, today, the capabilities of leading-edge emerging HPC architectures are not fully exploited by industrial simulation tools. Current state-of-the-art industrial solvers do not take sufficient advantage of the immense capabilities of new hardware architectures, such as streaming processors or many-core platforms. A combined research effort focusing on algorithms and HPC is the only way to make possible to develop and advance simulation tools to meet the needs of the European aeronautical industry.

NextSim will focus on the development of the numerical flow solver CODA (Finite Volume and high-order discontinuous Galerkin schemes), that will be the new reference solver for aerodynamic applications inside AIRBUS group, having a significant impact in the aeronautical market. To demonstrate NextSim market impact, AIRBUS has defined a series of market relevant problems. The numerical simulation of those problems is still a challenge for the aeronautical industry and their solution, at a required accuracy and an affordable computational cost, is still not possible with the current industrial solvers.

Following this idea, three additional working areas are proposed in NextSim: algorithms for numerical efficiency, algorithms for data management and the efficiency implementation of those algorithms in the most advanced HPC platforms.

Finally, NextSim will provide access to project results through the “mini-apps” concept, small pieces of software, seeking synergies with open-source components, which demonstrate the use of the novel mathematical methods and algorithms developed in CODA but that will be freely distributed to the scientific community.

Industrial software codes for extreme scale computing environments and applications

nextsimproject.eu

Twitter: @NextSim_EU
LinkedIn: @nextsim-eu

COORDINATING ORGANISATION

Barcelona Supercomputing Center (BSC), Spain

OTHER PARTNERS

- Universidad Politécnica de Madrid, Spain
- CIMNE - Centre Internacional de Mètodes Numèrics a l'Enginyeria, Spain
- Office National d'Etudes et de Recherches Aérospatiales (ONERA), France
- Deutsches Zentrum für Luft- und Raumfahrt (DLR), Germany
- Centre européen de recherche et de formation avancée en calcul scientifique (CERFACS), France
- Airbus Operations SAS, France

CONTACT

Oriol Lehmkuhl, oriol.lehmkuhl@bsc.es
Jordi Pons, jpons@cimne.upc.edu

CALL

EuroHPC-03-2019

PROJECT TIMESPAN

01/03/2021 - 29/02/2024

OPTIMA

Optimizing Industrial Applications for Heterogeneous HPC systems

In order to support the expanding demands for processing power from emerging HPC applications, within a pragmatic energy envelope, the future HPC systems will incorporate accelerators. One promising approach, towards this end, is the utilization of FPGAs; the main advantage of those devices is that, since they can be reconfigured at any time so as to implement tailor-made application accelerators, their energy efficiency and/or performance, in most of the cases, is much higher than that of CPUs and GPUs.

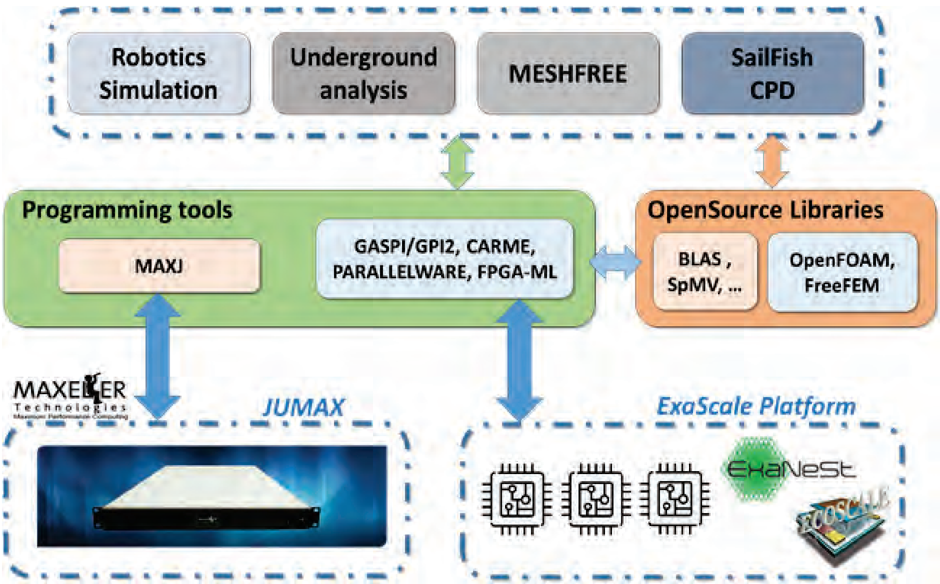

OPTIMA is an SME-driven project aiming to port and optimize a number of industrial applications as well as a set of open-source libraries, utilized in at least 3 different application domains, to two novel FPGA-populated HPC systems: a Maxeler MPC-X node located at the Jülich Supercomputer Centre and the ExaNest prototype located at the Institute of Computer Science - FORTH. Several innovative programming environments, toolchains and runtimes will be utilized including Maxeler’s MaxCompiler, Xilinx Runtime (XRT) and Fraunhofer's GASPI/GPI-2. It is expected that the applications and the libraries will be executed, in those heterogeneous HPC systems at significantly higher energy-efficiency as described by the Energy Delay Product metric (EDP); in particular, the EDPof the OPTIMA applications and libraries when executed on the targeted FPGA-based HPC systems, is expected to be more than 10x higher than that triggered by CPU-based systems and more than 3x higher than the GPU-based ones.

The main outcomes of OPTIMA will be that:

- the participating SMEs will gain a significant advantage since they will be able to execute their applications much more efficiently than the competition,
- it will be further proved that Europe is at the forefront of developing efficient FPGA-populated HPC systems and application/libraries taking advantage of them,
- the open-source libraries as well as the open-source applications developed within OPTIMA will allow third parties to easily target FPGA-based HPC systems for their application developments,
- there will be an open-to-use HPC infrastructure supported by a specially formed sustainability body.

Industrial software codes for extreme scale computing environments and applications

optima-hpc.eu

Twitter: @optima_hpc
LinkedIn: @optima-hpc

COORDINATING ORGANISATION

Telecommunication Systems Research Institute, Greece

OTHER PARTNERS

- Cyberbotics SARL, Switzerland
- Fraunhofer-Gesellschaft zur Förderung der angewandten Forschung e.V., Germany
- Exapsys | Exascale Performance Systems, Greece
- Institute of Communications and Computer Systems, Greece
- M3E S.r.l. Italy
- Maxeler IoT-Labs B.V., Netherlands
- Forschungszentrum Jülich, Germany
- EnginSoft SpA, Italy
- Appentra Solutions SL, Spain

CONTACT

info@optima-hpc.eu

CALL

EuroHPC-03-2019

PROJECT TIMESPAN

01/03/2021 - 30/11/2023

Phidias

Prototype of HPC/Data Infrastructure for On-demand Services

In today's world, more and more data are constantly being generated, resulting in changing the nature of computing, with an increasing number of data-intensive critical applications.

As the Council adopted a regulation on establishing the European High Performance Computing Joint Undertaking (EuroHPC) in July 2021, the regulation paves the way for the development within Europe of the next generation of supercomputers, strengthens research and innovation capabilities, the development of a supercomputing infrastructure ecosystem and the acquisition of world-class supercomputers.

The PHIDIAS project addresses a big challenge for Earth sciences. Today, many datasets are disseminated, and researchers are not always aware of their existence, their level of quality and more generally how to access and reuse them for analysis, in combination with their datasets to provide new information and knowledge.

The PHIDIAS project, funded by the European Union's Connecting Europe Facility (CEF), is built within this framework, aiming to become a reference point for the Earth science community enabling them to discover, manage and process spatial and environmental data, through the development of a set of high-performance computing (HPC) based services and tools exploiting large satellite datasets.

The project foresees the development of three use cases:

- Intelligent screening of a large amount of satellite data for detection and identification of anomalous atmospheric composition events;
- Processing on-demand services for environmental monitoring; and.
- Improving the use of cloud services for marine data management.

"With the efficient support of the consortium, I am working to ensure that the PHIDIAS objectives are being consistently pursued, to deliver a catalogue that will implement interoperable services for the discovery, access and processing of data, guaranteeing the largest degree of reusability of data as possible, and the improvement of the FAIRisation of satellite and environmental datasets. Essentially, paving the way and making life easier for the next generation of HPC and the Computational Scientific community."

Boris Dintrans, Director of CINES and PHIDIAS project coordinator

By federating different data sources, PHIDIAS will provide, on one hand, an interoperable and easy to use catalogue of environmental resources accessible and browsable by anybody, and, on the other hand, will provide researchers and practitioners with a platform on top of which new knowledge and new business models can be developed.

Synergies between earth science communities, high performance computing, and data-driven operations have the potential of enabling novel innovation opportunities and pave the way for the realisation of brand-new applications and services. PHIDIAS foresees a large social impact in terms of overall environmental monitoring capabilities enabled by the pooling of different stakeholders, namely, researchers, public authorities, private players, and citizen scientists.

Finally, it is also worth stressing that the social dimension cannot be underestimated. It is, however, our standpoint that the transparent access to environmental data enabled by PHIDIAS can allow us to provide due care to the players and countries that could be negatively affected by the low-carbon economy.

www.phidias-hpc.eu

Twitter: @PhidiasHpc
LinkedIn: @phidias-hpc

COORDINATING ORGANISATION

Centre Informatique National de l'Enseignement Supérieur (CINES), France

OTHER PARTNERS

- Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS), France
- Centre National de la Recherche Scientifique (CNRS), France
- IT Centre for Science Ltd. (CSC), Finland
- Geomatys, France
- Institut de Recherche pour le Développement (IRD), France
- Institut Français de Recherche pour l'exploitation de la Mer (IFREMER), France
- Mariene Informatie Service MARIS BV (MARIS), Netherlands
- Neovia Innovation, France
- SPASCIA, France
- Finnish Environment Institute (SYKE), Finland
- Trust-IT, Italy
- Université de Liège, Belgium

CONTACT

Boris Dintrans (Project Coordinator)
boris.dintrans@cines.fr
info@phidias-hpc.eu

CALL

CEF-TC-2018-5

PROJECT TIMESPAN

01/09/2019 - 30/09/2022

REGALE

An open architecture to equip next generation HPC applications with exascale capabilities

Supercomputers in the near future will be much more than just larger and more powerful machines than the ones used today. We are approaching exascale system – High Performance Computers that will be able to do 1018 floating point operations per second. With that, some paradigms of HPC will become obsolete. And energy consumption will become one of the key factors for the systems to come.

These systems will be very heterogeneous and complex on node level. They will leverage very different technologies, such as general-purpose CPUs as well as GPU and FPGA accelerators. For energy efficient computing without significant trade-off on the performance side, application will have to leverage these technologies dynamically. To achieve this, a new software stack is needed.

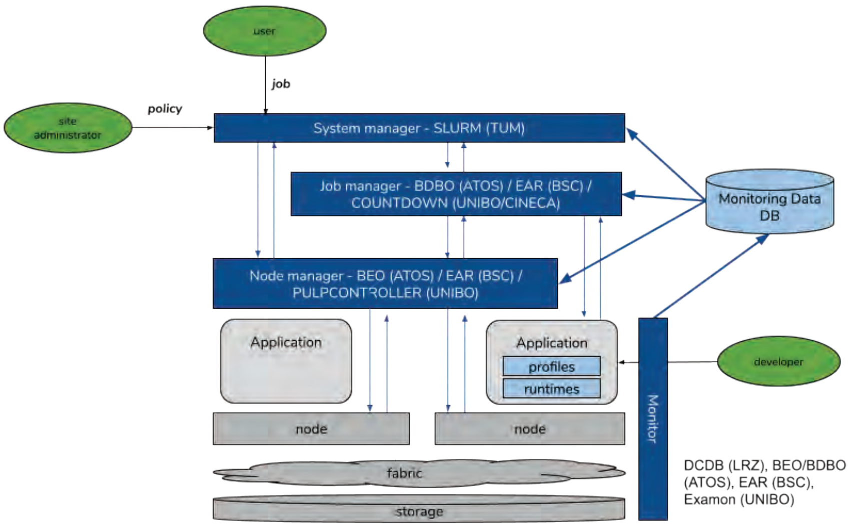

One of the new challenges for exascale is the total cost of ownership (TCO), and more specifically, the operating costs of exascale systems, the environmental footprint associated with the energy and power consumed, not only by the machine itself but also by the infrastructure supporting it, and power and performance inhomogeneity, stemming from manufacturing variability. The current state of practice in system wide resource management is rather passive in the sense that it merely serves user resource requests and applies brute-force enforcement of administrator-driven policies. What is needed is a holistic coordinator that will be able to support power, energy, performance, and throughput optimizations across the whole system in a scalable manner.

In the race to greater performance using less energy, the focus has often been on hardware. However, system software can now play a crucial role towards a controlled balance between application performance, power constraints and energy consumption, taking advantage of software interfaces for hardware configuration, without neglecting application needs. REGALE intends to realize these benefits under real operational environments by equipping the HPC system software stack with the mechanisms and policies for effective performance, power and energy control and optimization, bridging the gap between research and practice.

The REGALE project started in April 2021. It aims to build a software stack that is not only capable of bringing efficiency to tomorrow's HPC systems. It also will be open and scalable for massive supercomputers. REGALE brings together leading supercomputing stakeholders, prestigious academics, top European supercomputing centers and end users from critical target sectors, covering the entire value chain in system software and applications for extreme scale technologies.

HPC and data centric environments and application platforms

regale-project.eu

Twitter: @EuRegale

COORDINATING ORGANISATION

Institute of Communication and Computer Systems (ICCS), Greece

OTHER PARTNERS

- Technische Universität München (TUM), Germany
- Leibniz Supercomputing Centre, Germany
- ANDRITZ HYDRO GMBH, Austria
- Université Grenoble Alpes, France
- Atos (Bull SAS), France
- Barcelona Supercomputing Centre (BSC), Spain
- CINECA, Consorzio Interuniversitario, Italy
- E4 Computer Engineering SpA, Italy
- National Technical University of Athens - NTUA, Greece
- Ryax Technologies, France
- Electricité de France (EDF), France
- Ubitech, Greece
- SCiO Private Company, Greece
- TWT GmbH Science & Innovation, Germany
- Alma Mater Studiorum - Università di Bologna, Italy

CONTACT

Georgios Goumas, regale@cslab.ece.ntua.gr

CALL

EuroHPC-02-2019

PROJECT TIMESPAN

01/04/2021 - 31/03/2024

SCALABLE

SCAlable LAttice Boltzmann Leaps to Exascale

In SCALABLE, eminent industrials and academic partners will team up to improve the performance, scalability, and energy efficiency of an industrial LBM-based computational fluid dynamics (CFD) software.

Lattice Boltzmann methods (LBM) have already evolved to become trustworthy alternatives to conventional CFD. In several engineering applications, they are shown to be roughly an order of magnitude faster than Navier-Stokes approaches in a fair comparison and in comparable scenarios.

In the context of EuroHPC, LBM is especially well suited to exploit advanced supercomputer architectures through vectorization, accelerators, and massive parallelization.

In the public domain research code waLBerla, superb performance and outstanding scalability has been demonstrated, reaching more than a trillion (1012) lattice cells already on Petascale systems. WaLBerla performance excels because of its uncompromising unique, architecture-specific automatic generation of optimized compute kernels, together with carefully designed parallel data structures. However, is not compliant with industrial applications due to lack of a geometry engine and user friendliness for non-HPC experts.

On the other hand, the industrial CFD software LaBS already has such industrial capabilities at a proven high level of maturity, but it still has performance worthy of improvement.

The Karolina supercomputer (IT4Innovations, VSB-TUO) – 3 CFD numerical simulations on aircraft from various perspectives

SCALABLE will make the most of these two existing CFD tools (waLBerla and LaBS, the ProLB software) notably by transferring cutting-edge technology and aiming to break down the silos between the scientific computing world and that of physical flow modelling in order to deliver improved efficiency and scalability for the upcoming European Exascale systems.

The project outcomes will directly benefit to the European industry as confirmed by the active involvement of Renault and Airbus in the SCALABLE project, and will additionally contribute to fundamental research.

Industrial software codes for extreme scale computing environments and applications

scalable-hpc.eu

Twitter: @scalable_hpc
LinkedIn: @scalable-hpc

COORDINATING ORGANISATION

CS Group, France

OTHER PARTNERS

- Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany
- Centre européen de recherche et de formation avancée en calcul scientifique (CERFACS), France
- Forschungszentrum Jülich, Germany
- Neovia Innovation, France
- IT4Innovations, VSB – Technical University of Ostrava, Czech Republic
- AIRBUS Operations GmbH, Germany
- RENAULT SAS, France

CONTACT

Jérôme Texier, Project Coordinator,
jerome.texier@csgroup.eu

CALL

EuroHPC-03-2019

PROJECT TIMESPAN

01/01/2021 - 31/12/2023

Ecosystem development

CASTIEL

Connecting the countries NCCs to each other and the European Ecosystem

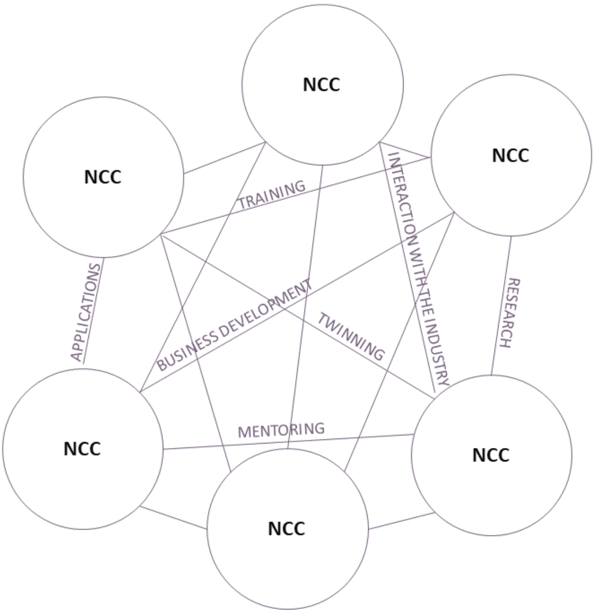

WHAT IS CASTIEL?

The Coordination and Support Action (CSA) CASTIEL leads to cross-European networking activities between National Competence Centres (NCCs) in HPC-related topics addressed through the EuroCC project. CASTIEL emphasises training, industrial interaction and cooperation, business development, raising awareness of high-performance computing (HPC)-related technologies and expertise. As a hub for information exchange and training, CASTIEL promotes networking among NCCs and strengthens idea exchange by developing best practices. The identification of synergies, challenges, and possible solutions is implemented through the close cooperation of the NCCs at a European level.

WHAT IS THE RELATIONSHIP BETWEEN CASTIEL AND EUROCC?

CASTIEL, the Coordination and Support Action (CSA) closely associated with EuroCC, combines the National Competence Centres (NCC) formed in EuroCC into a pan-European network. The aggregation of HPC, high-performance data analytics (HPDA), and artificial intelligence (AI) competencies demonstrates the global competitiveness of the EU partners. The two activities are the beginning of a strategic positioning of the European HPC competence and will contribute to the comprehensive independence of the mentioned technologies in Europe.

EuroHPC Innovating and Widening the HPC use and skills base

eurocc-access.eu

Twitter: @CASTIEL_project
LinkedIn: @CASTIEL-project

COORDINATING ORGANISATION

High-Performance Computing Center Stuttgart (USTUTT/HLRS), Germany

OTHER PARTNERS

- Gauss Centre for Supercomputing e.V. (GCS), Germany
- Consorzio Interuniversitario (CINECA), Italy
- TERATEC, France
- Barcelona Supercomputing Center – Centro Nacional De Supercomputación (BSC), Spain
- Partnership for Advanced Computing in Europe AISBL (PRACE), Belgium

CONTACT

Dr.-Ing. Bastian Koller, castiel@hlrs.de

CALL

EuroHPC-04-2019

PROJECT TIMESPAN

01.09.2020 – 31.08.2022

EuroCC

European Competence Centres

Within the EuroCC project under the European Union's Horizon 2020 (H2020), participating countries are tasked with establishing a single National Competence Centre (NCC) in the area of high-performance computing (HPC) in their respective countries. These NCCs coordinate activities in all HPC-related fields at the national level and serve as a contact point for customers from industry, science, (future) HPC experts, and the general public alike. The EuroCC project is funded 50 percent through H2020 (EuroHPC Joint Undertaking [JU]) and 50 percent through national funding programmes within the partner countries.

WHAT ARE THE NCCS' TASKS?

The tasks of the National Competence Centres include mapping existing HPC, high-performance data analytics (HPDA) and artificial intelligence (AI) competences, HPC training and the identification of HPC, HPDA and AI experts, as well as the coordination of services such as business development, application support, technology transfer, training and education, and access to expertise. Researchers from academia and industry both benefit from this competence concentration, and more efficient research ultimately benefits state and national governments and society as a whole.

The National Competence Centres shall act as the first points of contact for HPC, HPDA and AI in their respective countries. The NCCs will bundle information and provide experts, offer access to expertise and HPC training and education.

The overall objective of EuroCC — to create a European base of excellence in HPC by filling existing gaps — comes with a clear vision: offering a broad portfolio of services in all HPC, HPDA, and AI-related areas, tailored to the needs of industry, science, and public administrations.

COLLABORATION BETWEEN EUROCC AND CASTIEL

EuroCC works closely together with the Coordination and Support Action (CSA) CASTIEL, which links the national centres throughout Europe and ensures successful transfer and training possibilities between the participating countries.

EuroHPC Innovating and Widening the HPC use and skills base

eurocc-access.eu

Twitter: @EuroCC_project
LinkedIn: @EuroCC
Newsletter: subscribe on eurocc-access.eu

COORDINATING ORGANISATION

High-Performance Computing Center Stuttgart (USTUTT/HLRS), Germany

OTHER PARTNERS

- Gauss Centre for Supercomputing (GCS) Germany
- Institute of Information and Communication Technologies at Bulgarian Academy of Sciences (IICT-BAS), Bulgaria
- Universität Wien (UNIVIE), Austria
- University of Zagreb University Computing Centre (SRCE), Croatia
- Computation-based Science and Technology Research Centre, The Cyprus Institute (CaSToRC-CyI), Cyprus
- IT4Innovations, VSB – Technical University of Ostrava, Czech Republic
- Technical University of Denmark (DTU), Denmark
- University of Tartu HPC Center (UTHPC), Estonia
- CSC – IT Center for Science Ltd (CSC), Finland
- National Infrastructures for Research and Technology S.A. (GRNET S.A.), Greece
- Kormányzati Informatikai Fejlesztési Ügynökség (KIFÜ), Hungary
- National University of Ireland, Galway (NUI Galway), Ireland
- CINECA – Consorzio Interuniversitario (CINECA), Italy
- Vilnius University (LitGrid-HPC), Lithuania
- Riga Technical University (RTU), Latvia
- UNINETT Sigma2 AS (Sigma2), Norway
- Norwegian Research Centre AS (NORCE), Norway
- SINTEF AS (SINTEF), Norway
- Academic Computer Centre Cyfronet AGH (CYFRONET), Poland
- Fundação para a Ciência e a Tecnologia (FCT), Portugal
- National Institute for Research-Development in Informatics – ICI Bucharest (ICIB), Romania
- Academic and Research Network of Slovenia (ARNES), Slovenia
- Barcelona Supercomputing Center – Centro Nacional de Supercomputación (BSC), Spain
- Uppsala University (UU), Sweden
- Eidgenössische Technische Hochschule Zürich (ETH Zurich), Switzerland
- The Scientific and Technological Research Council of Turkey (TUBITAK), Turkey
- The University of Edinburgh (EPCC), United Kingdom
- TERATEC (TERATEC), France
- SURFSARA BV (SURFSARA), Netherlands
- Centre de recherche en aéronautique asbl (Cenaero), Belgium
- Luxinnovation GIE (LXI), Luxembourg
- Center of Operations of the Slovak Academy of Sciences (CC SAS), Slovak Republic
- University of Ss. Cyril and Methodius, Faculty of computer science and engineering (UKIM), Republic of North Macedonia
- Háskóli Íslands – University of Iceland (UICE), Iceland
- University of Donja Gorica (UDG), Montenegro

CONTACT

Dr.-Ing. Bastian Koller, eurocc@hlrs.de

CALL

EuroHPC-04-2019

PROJECT TIMESPAN

01.09.2020 – 31.08.2022

FF4EUROHPC

Connecting SMEs with HPC

BASE PILLARS OF THE FF4EUROHPC PROJECT

FF4EuroHPC has the general objective of supporting the EuroHPC initiative to promote the industrial uptake of high-performance computing (HPC) technology, in particular by small and medium-sized enterprises (SMEs), and thus increase the innovation potential of European industry. In order to achieve that general objective, FF4EuroHPC has defined a set of specific objectives.

THE FF4EUROHPC GENERAL OBJECTIVES ARE:

- Realise a portfolio of business-oriented application "experiments" that are driven by the SME end-users needs and executed by teams covering all required actors in the value-chain, with the innovation potential for HPC use and service provision of utmost priority.

- Support the future national HPC Competence Centres (resulting from the call EuroHPC-04-2019) to more effectively collaborate with SMEs through involvement in the FF4EuroHPC experiment portfolio and integration in related outreach activities.
- Support the participating SMEs in the establishment of HPC-related innovation either by using HPC infrastructures or services for their business needs or by providing new HPC-based services.
- Facilitate the widening of industrial HPC user communities and service providers in Europe by delivering compelling success stories for the use of HPC by SMEs; ensuring maximal awareness via communication and dissemination in collaboration with relevant Digital Innovation Hubs (DIHs) and industry associations.

The open calls target highest quality experiments involving innovative, agile SMEs and putting forward work plans built around innovation targets arising from the use of advanced HPC services. Proposals are sought that address business challenges from European SMEs from varied application domains, with strong preference being given to engineering and manufacturing, or sectors able to demonstrate fast economic growth or particular economic impact for Europe. The highest priority is given to proposals directly addressing the business challenges of manufacturing SMEs.

EuroHPC Innovating and Widening the HPC use and skills base

ff4eurohpc.eu

Twitter: @FF4EuroHPC
LinkedIn: @FF4EuroHPC
Newsletter: subscribe on ff4euroHPC.eu

COORDINATING ORGANISATION

High Performance Computing Centrer Stuttgart (USTUTT/HLRS), Germany

OTHER PARTNERS

- Scapos AG,Germany
- Consorzio Interuniversitario (CINECA), Italy
- Teratec (TERATEC), France
- Fundacion publica gallega centro tecnologico de supercomputacion de galicia (CESGA), Spain
- ARCTUR racunalniski inzeniring doo (ARCTUR), Slovenia

CONTACT

Dr.-Ing. Bastian Koller, ff4eurohpc@hlrs.de

CALL

EuroHPC-04-2019

PROJECT TIMESPAN

01.09.2020 – 31.08.2023

HiPEAC

High Performance Embedded Architecture and Compilation

HiPEAC is a coordination and support action (CSA) that aims to structure, connect and cross-fertilise the European academic and industrial research and innovation communities in Embedded Computing and Cyber-Physical Systems (i) by attracting members from the Cyber-Physical Systems community, industry and innovation community, (ii) by organising quarterly networking activities to connect the different communities, (iii) by

Extract from the HiPEAC comic
Past, present and future of the internet and digitally-augmented humanity - A HiPEAC vision

attracting computing talent and by retaining it in Europe (iv) by producing a high-quality vision on the future of computing systems in Europe, (v) by producing an impact analysis, and (vi) by professionally disseminating the research outcomes in and beyond the European computing systems community.

This HiPEAC CSA project is the continuation of five successful projects with the same name (HiPEAC1-5). HiPEAC will leverage the existing community, the expertise and the set of instruments that were developed since 2004 and work on the objectives of call ICT-01-2019, mainly in the field of Cyber-physical Systems of Systems (CPSoS).

The overall approach of the HiPEAC CSA is that it brings together all actors and stakeholders in the CPSoS and computing systems community in Europe in one well-managed structure where they can interact, disseminate/share information, transfer knowledge/technology, exchange human resources, think about future challenges, experiment with ideas to strengthen the community, etc.

The HiPEAC CSA will support its members and projects with tasks that are too difficult/complex to carry out individually like vision building, professional communications, recruitment, event management at the European level, and cross-fertilise the European academic and industrial research. By offering such services, a burden is taken away from the projects and members. They can then focus on the content, and the impact of their efforts is amplified.

European network in computing systems

hipeac.net
Twitter: @hipeac
LinkedIn: hipeac.net/linkedin
YouTube: hipeac.net/youtube

COORDINATING ORGANISATION

Universiteit Gent, Belgium

OTHER PARTNERS

- Barcelona Supercomputing Centre (BSC), Spain
- Chalmers Tekniska Högskola AB, Sweden
- Idryma Technologias Kai Erevnas, Greece
- Institut National de Recherche en Informatique et en Automatique (INRIA), France
- The University of Edinburgh (EPCC), United Kingdom
- Universita di Pisa, Italy
- Rheinisch-Westfälische Technische Hochschule Aachen (RWTH-AACHEN), Germany
- Commissariat à l'énergie atomique et aux énergies alternatives (CEA), France
- Arm Ltd, United Kingdom
- THALES, France
- Artemisia Vereniging, Netherlands

CONTACT

Koen De Bosschere, info@hipeac.net

CALL

ICT-01-2019

PROJECT TIMESPAN

01/12/2019 - 28/02/2023

HPC-EUROPA3

Pan-European Research Infrastructure on High Performance Computing

TRANSNATIONAL ACCESS

The core activity of HPC-Europa3 is the Transnational Access programme, which funds short collaborative research visits in any scientific domain using High Performance Computing (HPC).

Visitors gain access to some of Europe's most powerful HPC systems, as well as technical support from the relevant HPC centre to enable them to make the best use of these facilities.

Applicants identify a "host" researcher working in their own field of research, and are integrated into the host department during their visit. Visits can be made to any research institute in Finland, Germany, Greece, Ireland, Italy, the Netherlands, Spain, Sweden or the UK. (Note that project partner CNRS does not participate in the Transnational Access activity and therefore it is not possible to visit research groups in France).

HPC-Europa3 visits can last between 3 and 13 weeks.

The programme is open to researchers of any level, from academia or industry, working in any area of computational science. Priority is given to researchers working in the EU and Associated States (see http://bit.ly/AssociatedStates), but limited places are available for researchers from third countries.

HPC-Europa3 has introduced a "Regional Access Programme" to encourage applications from researchers working in the Baltic and South-East Europe regions. The Regional Access Programme specifically targets researchers with little or no HPC experience who need more powerful computing facilities than they have, but not the most powerful systems offered by HPC-Europa3. Such researchers can apply respectively to KTH-PDC in Sweden or GRNET in Athens, and will be given priority over more experienced applicants.

NETWORKING AND JOINT RESEARCH ACTIVITIES

The main HPC-Europa3 Networking Activity, External co-operation for enhancing the best use of HPC, aims to build stronger relationships with other European HPC initiatives, e.g., PRACE, ETP4HPC, and the HPC Centres of Excellence. The goal is to provide HPC services in a more integrated way across Europe. This activity also aims to increase awareness of HPC within SMEs. A series of workshops targeted at SMEs has been organised to encourage their uptake of HPC.

The Joint Research Activity, Container-as-a-service for HPC, aims to enable easy portability of end-user applications to different HPC centres, providing clear benefits for HPC-Europa3 visitors. This approach enables different applications and runtime environments to be supported simultaneously on the same hardware resources with no significant decrease in performance. It also provides the capability to fully capture and package all dependencies of an application so that the application can be migrated to other machines or preserved to enable a reproducible environment in the future.

The results from this JRA can be found in various deliverables which are available at http://www.hpc-europa.eu/public_documents.

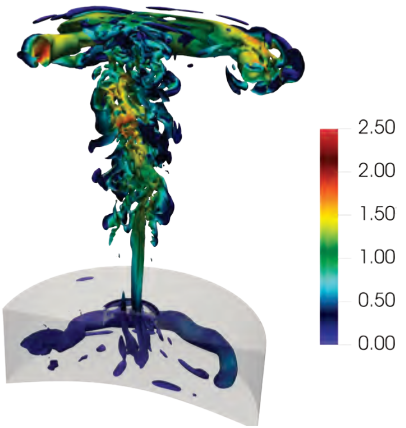

Flow of a circular synthetic jet represented using the Q-criterion and coloured by the magnitude of the velocity (from research carried out by HPC-Europa3 visitor Arnau Miró, Universitat Politècnica de Catalunya)

www.hpc-europa.eu

Twitter: @hpceuropa3
Facebook: @hpceuropa
LinkedIn: @hpc-europa3-eu-project
YouTube: www.youtube.com

COORDINATING ORGANISATION

CINECA Consorzio Interuniversitario, Italy

OTHER PARTNERS

- BSC - Barcelona Supercomputing Centre, Spain
- CNRS - Centre National de la Recherche Scientifique, France
- CSC – IT Center for Science Ltd., Finland
- EPCC - The University of Edinburgh, United Kingdom
- GRNET - Ethniko Diktyo Erevnas Technologias AE, Greece
- HLRS - Universität Stuttgart, Germany
- KTH - Kungliga Tekniska Hoegskolan, Sweden
- ICHEC - National University of Ireland Galway, Ireland
- SURF BV, The Netherlands

CONTACT

staff@hpc-europa.org

CALL

INFRAIA-01-2016-2017

PROJECT TIMESPAN

01/05/2017 - 30/04/2022

Index of on-going projects

Index of completed projects

Projects completed before 1st January 2021 are not featured in this edition of the Handbook, but you can find the descriptions of completed projects in the 2018, 2019 or 2020 Handbooks, available on our website:

Project name	Project description	Handbook Edition
ALLScale	An Exascale Programming, Multi-objective Optimisation and Resilience Management Environment Based on Nested Recursive Parallelism	2018
ANTAREX	AutoTuning and Adaptivity appRoach for Energy efficient eXascale HPC systems	2018
AQMO	Air Quality and MObility	2020
ARGO	WCET-Aware Parallelization of Model-Based Applications for Heterogeneous Parallel Systems	2019
BioExcel 1	Centre of Excellence for Biomolecular Research	2019
COEGSS	Centre of Excellence for Global Systems Science	2018
ComPat	Computing Patterns for High Performance Multiscale Computing	2018
CompBioMed 1	A Centre of Excellence in Computational Biomedicine	2018
ECOSCALE	Energy-efficient heterogeneous COmputing at exaSCALE	2019
EoCoE 1	Energy oriented Centre of Excellence for computer applications	2018
ESCAPE 1	Energy-efficient SCalable Algorithms for weather Prediction at Exascale	2018
ESIWACE 1	Centre of Excellence in Simulation of Weather and Climate in Europe	2019
EUROLAB-4-HPC 2	Consolidation of European Research Excellence in Exascale HPC Systems	2020
ExaFLOW	Enabling Exascale Fluid Dynamics Simulations	2018
ExaHyPE	An Exascale Hyperbolic PDE Engine	2019
ExaNeST	European Exascale System Interconnect and Storage	2019

Project name	Project description	Handbook Edition
ExaNoDe	European Exascale Processor and Memory Node Design	2019
ExCAPE	Exascale Compound Activity Prediction Engine	2018
EXDCI-2	European eXtreme Data and Computing Initiative	2020
EXTRA	Exploiting eXascale Technology with Reconfigurable Architectures	2018
greenFLASH	Green Flash, energy efficient high performance computing for real-time science	2018
INTERTWINE	Programming Model INTERoperability ToWards Exascale (INTERTWinE)	2018
MANGO	Exploring Manycore Architectures for Next-GeneratiOn HPC systems	2018
MaX 1	Materials design at the eXascale	2019
M2DC	Modular Microserver DataCentre	2019
Mont-Blanc 3	Mont-Blanc 3, European scalable and power efficient HPC platform based on low-power embedded technology	2019
NEXTGenIO	Next Generation I/O for the Exascale	2018
NLAFET	Parallel Numerical Linear Algebra for Future Extreme Scale Systems	2018
NoMaD 1	The Novel Materials Discovery Laboratory	2018
POP 1	Performance Optimisation and Productivity	2018
PROCESS	PROviding Computing solutions for ExaScale ChallengeS	2020
OPRECOMP	Open Transprecision COmputing	2020
READEX	Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing	2018
SAGE 1	Percipient StorAGe for Exascale Data Centric Computing	2018