European High-Performance Computing

Handbook **2019**

Dear Reader,

This Handbook has become a recognised mechanism to provide an overview of the European HPC projects. Its increasing size is a sign of the vitality of our European HPC landscape. We have included a number of projects besides those financed by HPC cPPP-related calls, as they are closely related to HPC. If you know of an EU-funded project related to HPC but not mentioned in this Handbook, do report it to our team, we will be happy to include it in the next edition.

EuroHPC – a joint initiative of the European Commission and 29 European member states – has taken over the task of defining and funding European HPC projects. The objective of EuroHPC is also to fund European HPC systems and that is where we hope the projects presented in this publication will be able to find a venue for the application of their results.

The projects started in 2015 – at the beginning of the H2020 multiannual research programme - have reached their end. They have resulted in an impressive outcome, with over 150 applicable pieces of technology produced by 19 projects: software packages, APIs, demonstrators, libraries, memory types, etc. We hope these results will be reused, leading to synergies and collaborations between European (and overseas) players. (The full list of outcomes and their analysis is available on the website of the EXDCI project: https://exdci.eu/activities/fethpc-results).

On the other hand, there is new work on the horizon. The first EuroHPC call for projects was issued during the summer, taking into account the suggestions presented by our ecosystem (represented in the EuroHPC JU by ETP4HPC and BDVA, the Big Data Value Association). The projects selected in this process will form part of next year's Handbook. See you all then!

EXDCI-2 Coordinator and PRACE Director
Serge Bogaerts

ETP4HPC Chairman
Jean-Pierre Panziera

Table of Contents

EUROPEAN EXASCALE PROJECTS

This edition of our Handbook now features a number of projects, which, though not financed by FET-HPC calls, are closely related to HPC and have their place here. For ease of use, we have grouped projects by scope or objective, instead of by H2020 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.
Finally, this Handbook only features on-going projects – including those that end in 2019. Projects that ended before January 2019 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).

Most of the projects in this Handbook have been funded in the scope of H2020 HPC contractual Public Private Partnership. Notable particularities are:

- "HPC and Big Data testbeds" projects which were funded by a call sitting between HPC and Big Data cPPPs;
- "Other projects" which encompasses a selection of HPC projects funded in calls which are not under the umbrella of HPC cPPP;
- HPC-Europa3 which is not directly related to the cPPP either but was added besides cPPP FETHPC Coordination and Support Actions, because of its ecosystem development dimension.

CHAPTER	CALLS TO WHICH THE PROJECTS BELONG	COMMENTS
Technology Projects **All projects in this chapter are hardware and/or software oriented**	**FETHPC-1-2014** HPC Core Technologies, Programming Environments and Algorithms for Extreme Parallelism and Extreme Data Applications	19 projects starting 2015 with a duration of around 3 years
	FETHPC-01-2016 Co-design of HPC systems and applications	2 projects starting in 2017 with a duration of around 3 years
	FETHPC-02-2017 Transition to Exascale	11 projects starting 2015 with a duration of around 3 years
Centres of Excellence in computing applications	**E-INFRA-5-2015** Centres of Excellence for computing applications	9 projects starting in 2016 with a duration of around 3 years
	INFRAEDI-02-2018 Centres of Excellence for computing applications	10 projects and 1 Coordination and Support Action starting in 2018
European Processor	**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)
	ICT-05-2017 Customised and low energy computing (including Low power processor technologies)	Prequel to EPI
HPC and Big Data testbeds	**ICT-11-2018-2019** Subtopic (a) HPC and Big Data enabled Large-scale Test-beds and Applications	4 Innovative Action starting in 2018
International Cooperation	**FETHPC-01-2018** International Cooperation on HPC	2 projects starting June 2019
Other projects	**ICT-04-2015** Customised and low power computing **ICT-12-2018-2020** Big Data technologies and extreme-scale analytics **FETPROACT-01-2016** Proactive: emerging themes and communities **EINFRA-21-2017** Platform-driven e-infrastructure innovation	A selection of technology projects related to HPC (hardware/software) from various ICT, FET and INFRA calls
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

Technology Projects

ASPIDE

exAScale ProgramIng models for extreme Data processing

www.aspide-project.eu
: @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*
- Bull SAS (Atos Group), *France*

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

CALL
FETHPC-02-2017

PROJECT TIMESPAN
15/06/2018 - 14/12/2020

OBJECTIVE

Extreme Data is an incarnation of Big Data concept distinguished by the massive amounts of data that must be queried, communicated and analysed in (near) real-time by using a very large number of memory/storage elements and Exascale computing systems. Immediate examples are the scientific data produced at a rate of hundreds of gigabits-per-second that must be stored, filtered and analysed, the millions of images per day that must be mined (analysed) in parallel, the one billion of social data posts queried in real-time on an in-memory components database. Traditional disks or commercial storage cannot handle nowadays the extreme scale of such application data.

Following the need of improvement of current concepts and technologies, ASPIDE's activities focus on data-intensive applications running on systems composed of up to millions of computing elements (Exascale systems). Practical results will include the methodology and software prototypes that will be designed and used to implement Exascale applications. 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.

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, 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, the latest project in the series, is building a prototype with three modules: a general-purpose cluster for low or medium scalable codes, a highly scalable booster comprising a cluster of accelerators, and a Data Analytics Module (DAM), which will be tested with six applications combining high-performance computing (HPC) with high-performance data analytics (HPDA) and machine learning (ML). The Cluster Module and the Data Analytics Module are already installed.

Forschungszentrum Jülich / Ralf- Uwe Limbach

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 are 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 will be used to define and evaluate the hardware and software technologies developed. Careful analysis of the application codes allows 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. Based on general-purpose processors with a huge amount of memory per core, this module is boosted by powerful general-purpose graphics processing units (GPGPU) and field-programmable gate array (FPGA) accelerators.

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.

www.deep-projects.eu
: @DEEPprojects

COORDINATING ORGANISATION

Forschungszentrum Jülich
(Jülich Supercomputing Centre)

OTHER PARTNERS

Current partners in DEEP-EST:
- Forschungszentrum Jülich,
- Intel,
- Leibniz Supercomputing Centre,
- Barcelona Supercomputing Centre (BSC),
- Spain,
- Megware Computer Vertrieb und Service GmbH,
- Heidelberg University,
- EXTOLL,
- The University of Edinburgh,
- Fraunhofer ITWM,
- Astron,
- KU Leuven,
- National Center For Supercomputing Applications (*Bulgaria*),
- Norges Miljo-Og Biovitenskaplige Universitet Haskoli Islands,
- European Organisation for Nuclear Research (CERN),
- ParTec.

Partners in DEEP and DEEP-ER:
- CERFACS,
- CGG,
- CINECA,
- Eurotech,
- EPFL,
- Seagate,
- INRIA, Institut National de Recherche en Informatique et Automatique,
- Mellanox,
- The Cyprus Institute,
- Universität Regensburg.

CONTACT

Dr. Estela Suarez
pmt@deep-projects.eu

CALL

ETHPC-01-2016

PROJECT TIMESPAN (DEEP-EST)

01/07/2017 - 31/03/2021

ECOSCALE

Energy-efficient Heterogeneous COmputing at exaSCALE

ECOSCALE tackles the exascale challenge by providing a novel heterogeneous energy-efficient hierarchical architecture, a hybrid distributed OpenCL programming environment and a runtime system. The ECOSCALE architecture, programming model and runtime system follows a hierarchical approach where the system is partitioned into multiple autonomous Workers (i.e. compute nodes). Workers are interconnected in a tree-like structure in order to form larger Partitioned Global Address Space (PGAS) partitions. To further increase the energy efficiency of the system, the Workers employ reconfigurable accelerators implemented on FPGAs that perform coherent memory accesses in the virtual address space while being programmed by OpenCL.

The novel UNILOGIC (Unified Logic) architecture, introduced within ECOSCALE, is an extension of the UNIMEM architecture. UNIMEM provides shared partitioned global address space while UNILOGIC provides shared partitioned reconfigurable resources. The UNIMEM architecture gives the user the option to move tasks and processes close to data instead of moving data around and thus it reduces significantly the data traffic and related energy consumption and delays. The proposed UNILOGIC+UNIMEM architecture and the ECOSCALE design tools partition seamlessly the application into several Worker nodes that communicate through a fat-tree communication infrastructure. The final FPGA-based ECOSCALE prototype can execute real-world applications with 10x higher energy efficiency than conventional CPU- and/or GPU-based parallel systems.

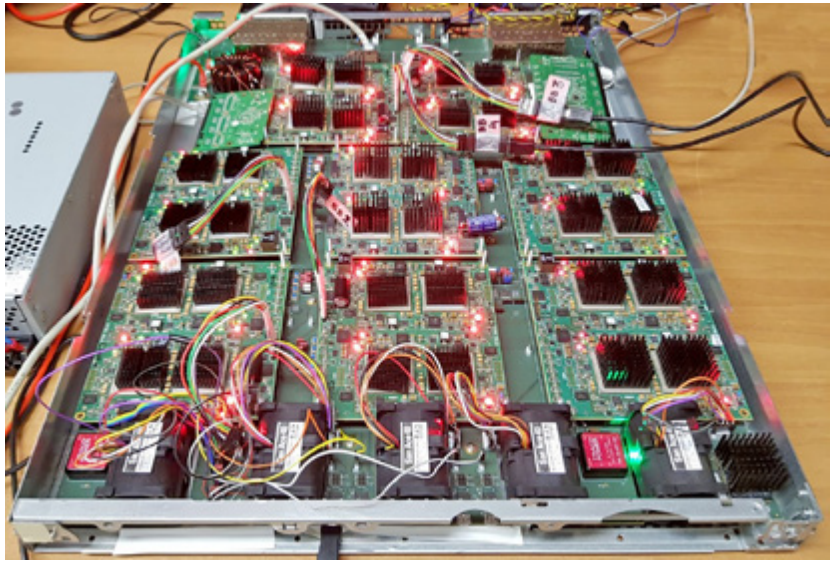

ECOSCALE prototype

www.ecoscale.eu
: @ECOSCALE_H2020

COORDINATING ORGANISATION

Telecommunication Systems Institute (TSI)

OTHER PARTNERS

- Telecommunication Systems Institute
- Queen's University Belfast
- STMicroelectronics
- Acciona
- University of Manchester
- Politecnico di Torino
- Chalmers University of Technology, *Sweden*
- Synelixis Solutions

CONTACT

Yannis Papaefstathiou
ygp@synelixis.com
+30 694427772

CALL

FETHPC-1-2014

PROJECT TIMESPAN

01/10/2015 - 31/07/2019

EPEEC

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

OBJECTIVE

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 will significantly advance and integrate existing state-of-the-art components based on European technology (programming models, runtime systems, and tools) with key features enabling 3 overarching objectives: high coding productivity, high performance, and energy awareness.

An automatic generator of compiler directives will provide outstanding coding productivity from the very beginning of the application developing/porting process. Developers will be able to leverage either shared memory or distributed-shared memory programming flavours, and code in their preferred language: C, Fortran, or C++. EPEEC will ensure the composability and interoperability of its programming models and runtimes, which will incorporate specific features to handle data-intensive and extreme-data applications. Enhanced leading-edge performance tools will offer integral profiling, performance prediction, and visualisation of traces. Five applications representative of different relevant scientific domains will serve as part of a strong inter-disciplinary co-design approach and as technology demonstrators.

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.

www.epeec-project.eu

COORDINATING ORGANISATION

Barcelona Supercomputing Centre (BSC), *Spain*

OTHER PARTNERS

- Frauhofer 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 - 30/09/2021

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

OBJECTIVE

- 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

https://epigram-hs.eu
: @EpigramHs

COORDINATING ORGANISATION

KTH, *Sweden*

OTHER PARTNERS

- EPCC, *UK*
- ETHZ - Eidgenössische Technische Hochschule Zürich, *Switzerland*
- Fraunhofer, *Germany*
- Cray, *Switzerland*
- ECMWF, *UK*

CONTACT

Stefano Markidis
markidis@kth.se

CALL

FETHPC-02-2017

PROJECT TIMESPAN

01/09/2018 - 31/08/2021

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.

www.ecmwf.int/escape

COORDINATING ORGANISATION

ECMWF - European Centre for Medium-range Weather Forecasts, *UK*

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*
- Bull SAS (Atos group), *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

The Horizon 2020 EuroEXA project proposes a ground-breaking design for mind blowing results: Over four times higher performance and four times higher energy efficiency than today's high-performance platforms. Originally the informal names for a group of H2020 research projects, ExaNeSt, EcoScale and ExaNoDe, EuroEXA has its own EU investment as a co-design project to further develop technologies from the project group and support the EU in its bid to deliver EU based ExaScale supercomputers.

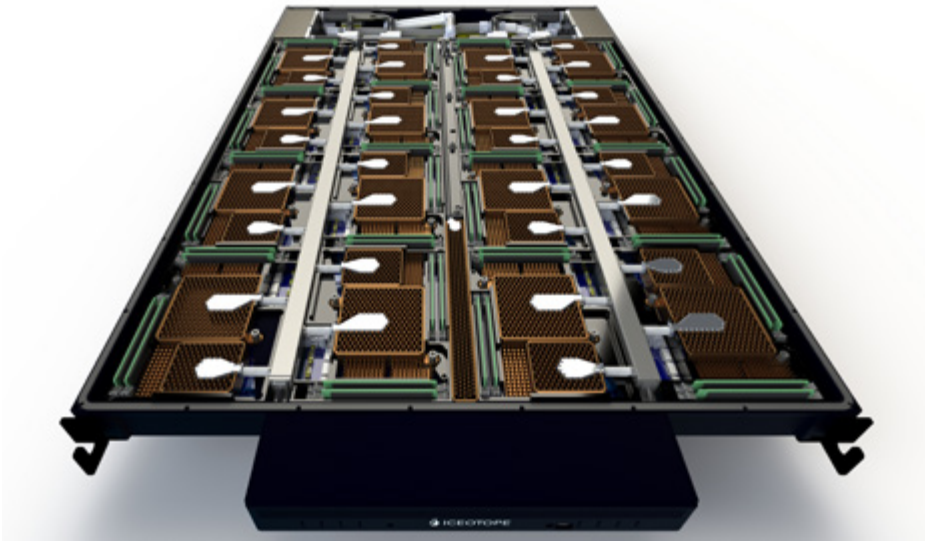

The EuroEXA Blade

VISION:

- First testbed architecture will be shown to be capable of scaling to world-class peak performance in excess of 400 PFLOPS with an estimated system power of around 30 MW peak
- A compute-centric 250 PFLOPS per 15 MW by 2019
- Show that an exascale machine could be built in 2020 within 30 shipping containers with an edge to edge distance of less than 40m

This project has a €20m investment over a 42-month period and is part of a total €50m, investment made be 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.

The project is also supported by a high value donation of IP from ARM and Synopsys. Funded under H2020-EU.1.2.2. FET Proactive (FETHPC-2016-01) as a result of a competitive selection process, the consortium, partners bring a rich mix of key applications from across climate/weather, physics/energy and life-science/bioinformatics. The project objectives include to develop and deploy an ARM Cortex technology processing system with FPGA acceleration at peta-flop level an Exa-scale procurement for deployment in 2022/23.

www.euroexa.eu
: @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, *SE*
- 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
+44 1142 245500

CALL

FETHPC-01-2016

PROJECT TIMESPAN

01/09/2017 - 28/02/2021

EXA2PRO

Enhancing Programmability and boosting Performance Portability for Exascale Computing Systems

OBJECTIVE

The vision of EXA2PRO is to develop a programming environment that will enable the productive deployment of highly parallel applications in exascale computing systems. EXA2PRO programming environment will integrate tools that will address significant exascale challenges. It will support a wide range of scientific applications, provide tools for improving source code quality, enable efficient exploitation of exascale systems' heterogeneity and integrate tools for data and memory management optimization. Additionally, it will provide fault-tolerance mechanisms, both user-exposed and at runtime system level and performance monitoring features. EXA2PRO will be evaluated using 4 applications from 3 different domains, which will be deployed in JUELICH supercomputing centre: High energy physics, materials and supercapacitors. The applications will leverage the EXA2PRO toolchain and we expect:

1. Increased programmability that enables the efficient exploitation of heterogeneity of modern supercomputing systems, which allows the evaluation of more complex problems.
2. Effective deployment in an environment that provides features critical for exascale computing systems such as fault tolerance, flexibility of execution and performance monitoring based on EXA2PRO optimization tools.
3. Identification of trade-offs between design qualities (source code maintainability/reusability) and run-time constraints (performance/energy consumption).

EXA2PRO outcome is expected to have major impact on:

1. the scientific and industrial community that focuses on application deployment in supercomputing centres: EXA2PRO environment will allow efficient application deployment with reduced effort.
2. application developers that target exascale systems: EXA2PRO will provide 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: Significant impact is expected on the materials and processes design for CO2 capture and on the supercapacitors industry.

https://exa2pro.eu

COORDINATING ORGANISATION

ICCS (Institute of Communication and Computer Systems), *Greece*

OTHER PARTNERS

- Linkopings 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*

CONTACTS

Prof. Dimitrios Soudris
dsoudris@microlab.ntua.gr

CALL

FETHPC-02-2017

PROJECT TIMESPAN

01/05/2018 - 30/04/2021

ExaHyPE

An Exascale Hyperbolic PDE Engine

EXAHYPE VISION

Hyperbolic systems of PDE resulting from conservation laws are the basis to simulate a wide range of problems, ranging from complex earthquake physics, via atmospheric processes and all kinds of flow problems up to the most catastrophic events in the universe. The ExaHyPe engine shall enable teams of computational scientists to more quickly realize grand challenge simulations on exascale hardware based on hyperbolic PDE models.

EXAHYPE ALGORITHMS

The engine implements a high-order discontinuous Galerkin approach with ADER time stepping and a-posteriori finite-volume limiting, which combines high-order accuracy with the robustness of finite volumes. Problems are discretised on tree-structured fully adaptive Cartesian meshes, for which advanced parallelisation approaches and load-balancing algorithms are available.

EXAHYPE ENGINE

The hyperbolic PDE engine is available as open source software, hosted at www.exahype.org.

The consortium provides a guidebook coming along with the released code which contains, besides documentation, rationale and further documentation links.

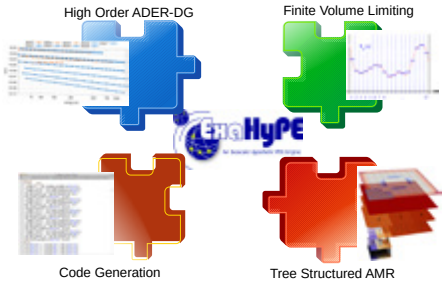

EXAHYPE DEMONSTRATOR APPLICATIONS

The ExaHyPE project evaluates the engine's capabilities on two exascale candidate scenarios stemming from computational seismology and astrophysics, respectively: regional earthquake simulation, particularly in alpine areas, and the fully relativistic simulation of (systems of) neutron stars and black holes. Demonstrator scenarios from these application areas and further use-cases are available on exahype.org.

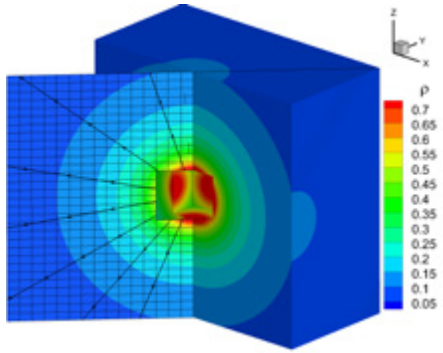

ADER-DG solution for the 3D Michel accretion test.

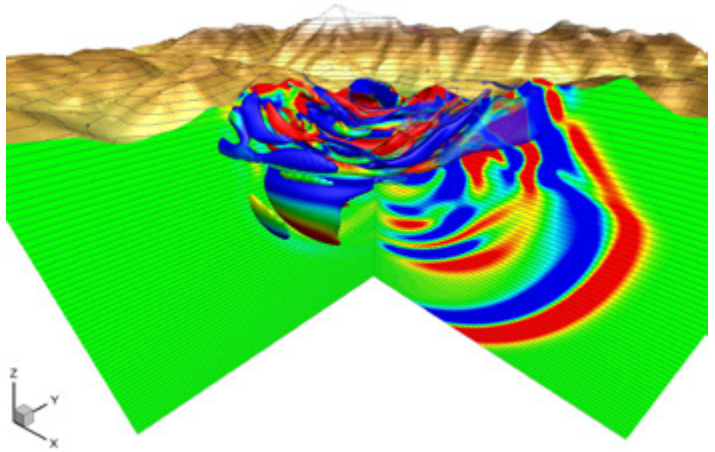

Seismic wave propagation in a complex alpine topography.

www.exahype.eu

COORDINATING ORGANISATION

Technical University of Munich

OTHER PARTNERS

- Universita degli Studi di Trento
- Durham University
- Frankfurt Institute for Advanced Studies
- Ludwig-Maximilians-Universität München
- RSC Technologies
- Bavarian Research Alliance

CONTACTS

Michael Bader
bader@in.tum.de

CALL

FETHPC-1-2014

PROJECT TIMESPAN

01/10/2015 - 30/09/2019

ExaNeSt

European Exascale System Interconnect and Storage (ExaNeSt) project has successfully built a prototype for Exascale. ExaNeSt has taken a sensible integrated approach to co-design the hardware and software that is needed to enable the prototype to run in real-life evaluations. The energy-efficient solutions are expected to spread to standard off the shelf computing systems, assisting the efforts of the EU to reduce carbon dioxide emissions and reach environmental targets.

ACHIEVEMENTS:

- Unified, Optimized, Low-power, Hierarchical Interconnect
  - Unimem (zero-copy, user-level), low-cost; Top-of-Rack; Optical
- Distributed, in-node NVM's, scalable Storage
  - Modified BeeGFS parallel FS, ported data base; Resilience; Virtualization
- Scalable dense packaging based on Liquid Cooling
  - Ultra-dense daughter boards; next-generation cooling technology
- Not only simulations for scalability, but also built a rack prototype
- Design tailored to entire Applications
  - Ported, re-engineered; WP6: integrated, evaluated
- Exploitation Objective: New products; follow on projects

INTERCONNECT:

- Multi-tiered 4D topology
- Custom, low-cost, resilient, virtualised, RDMA-capable interconnect
- Sub-microsecond process-to-process communication
- Low-latency packet-based prototype (ExaNet)
- Running real HPC app's (ExaNet MPI port) 10x or more improvement over Ethernet

STORAGE:

- UNIMEM – PGAS
- BeeGFS Filesystem
- HPC Virtualization
- Profiling Tools
- Checkpointing
- Status Monitoring
- FPGA Acceleration

ExaNeSt has contributed to a new generation of European Supercomputers. The project's advances in performance and efficiency will enable SME's in several sectors of the economy to utilize HPC and data analytics, with an attractive trade-off between usability and affordability.

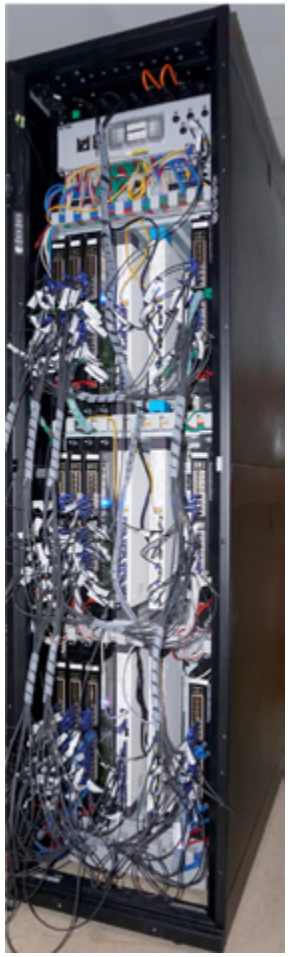

Exanest Prototype

www.exanest.eu
: @ExaNeSt_H2020

COORDINATING ORGANISATION

FORTH

OTHER PARTNERS

- Iceotope
- Allinea
- EnginSoft
- eXact lab
- MonetDB
- Virtual Open Systems
- Istituto Nazionale di Astrofisica (INAF)
- National Institutefor Nuclear Physics(INFN)
- The University of Manchester
- Universitat Politècnica de València
- Fraunhofer

CONTACT

Manolis Katevenis
(Project Coordinator)
kateveni@ics.forth.gr
+30 2811 39 1078

Peter Hopton
(Dissemination & Exploitation Lead)
peter.hopton@iceotope.com
+44 7437 01 2186

CALL

FETHPC-1-2014

PROJECT TIMESPAN

01/12/2015 - 31/05/2019

ExaNoDe

European Exascale Processor & Memory Node Design

ExaNoDe designed core technologies for a highly energy efficient and highly integrated heterogeneous compute node directed towards Exascale computing. ExaNoDe has built a compute node prototype available for collaborative working comprising:

- Technology and design solutions for an interposer-based computing device targeting HPC applications,
- Integration of devices in a Multi- Chip-Module (MCM) to increase compute density and take advantage of heterogeneity,
- System and middleware SW stack.

CONTRIBUTION TO EXASCALE COMPUTING

The ExaNoDe final prototype consists of an integrated HW/SW daughter board with two MCMs, each including two chiplets stacked on one silicon interposer and two Xilinx Zynq Ultrascale+ FPGAs. It targets prototype-level use by system integrators, software teams and its subsequent evaluation through industrial deployment.

ExaNoDe's main achievements cover all aspects of heterogeneous integration from silicon technology to system software including:

- Design of innovative, high-speed and low-power interconnect for the heterogeneous integration of chiplets via a silicon interposer.
- Design of a Convolutional Neural Network (CNN) accelerator hardware IP for a use case demonstrating heterogeneous integration.
- Design and manufacturing of a chiplet System-on-Chip (SoC) in 28FDSOI technology node
- 3D integration of chiplets on an active silicon interposer with about 50,000 high density (20 µm pitch) connections.

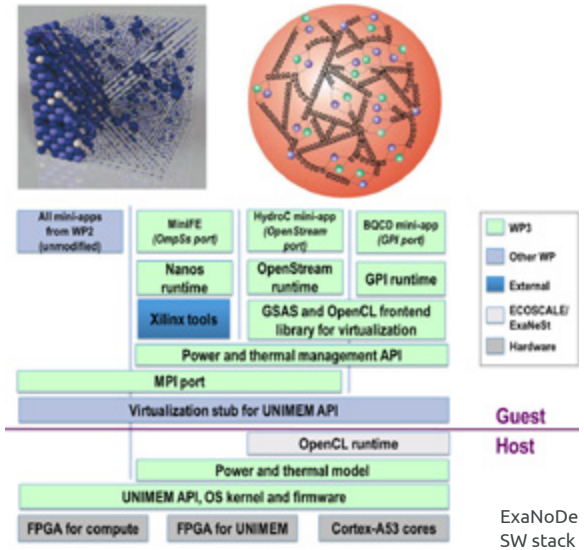

ExaNoDe SW stack

- Advanced package integration with two FPGA bare dies including ARMv8 cores, one interposer and 43 decoupling capacitors in a 68.5 mm × 55 mm Multi-chip-Module (MCM).
- Integration of two MCMs on a 260 mm x 120 mm daughter board. Development of a complete SW stack including UNIMEM-based system software and middleware; Runtimes libraries optimized for the UNIMEM architecture (OmpSs, MPI, OpenStream, GPI); Checkpointing technology for virtualisation; A set of mini-applications for benchmarking purposes.
- A performance projection of ExaNoDe technologies into a strawman architecture representative of upcoming HPC processors.

The ExaNoDe prototype has been delivered at the end of June 2019. It has been designed to demonstrate and to validate all technologies developed within the project. The methodology, the know-how and the building blocks related to advanced packaging (MCM and interposer) and system software for heterogeneous integration provide an IP baseline for the next generation of European processors to be used in future Exascale systems.

ExaNoDe prototype with Multi-Chip-Modules, 3D Integrated Circuit and FPGA bare dies

www.exanode.eu
: @ExanodeProject

COORDINATING ORGANISATION

CEA - Commissariat à l'Energie Atomique et aux énergies alternatives, France

OTHER PARTNERS

- Arm Limited, *UK*
- ETHZ - Eidgenössische Technische Hochschule Zürich, *Switzerland*
- FORTH, *Greece*
- Fraunhofer ITWM, *Germany*
- Scapos AG, *Germany*
- University of Manchester, *UK*
- Bull SAS (Atos Group), *France*
- Virtual Open Systems, *France*
- Barcelona Supercomputing Centre (BSC), *Spain*
- Forschungszentrum Jülich, *Germany*
- Kalray SA, *France*
- CNRS - Centre National de la Recherche Scientifique, *France*

CONTACTS

Denis Dutoit
denis.dutoit@cea.fr
+33677126676

Guy Lonsdale
guy.lonsdale@scapos.com
+492241142818

CALL

FETHPC-1-2014

PROJECT TIMESPAN

01/10/2015 - 30/06/2019

ExaNoDe collaborated closely with the ExaNeSt and EcoScale projects on the basis of a joint MoU.

ExaQUte

ExaQUte
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.

www.exaqute.eu
: @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*
- Vysoka Skola Banska - Technicka Univerzita Ostrava, *Czech Republic*
- Ecole Polytechnique Fédérale de Lausanne, *Switzerland*
- Universitat Politecnica de Catalunya, *Spain*
- str.ucture GmbH, *Germany*

CONTACTS

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

CALL

FETHPC-02-2017

PROJECT TIMESPAN

01/06/2018 - 31/05/2021

MAESTRO

Middleware for memory and data-awareness in workflows

MAESTRO
DATA ORCHESTRATION

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:

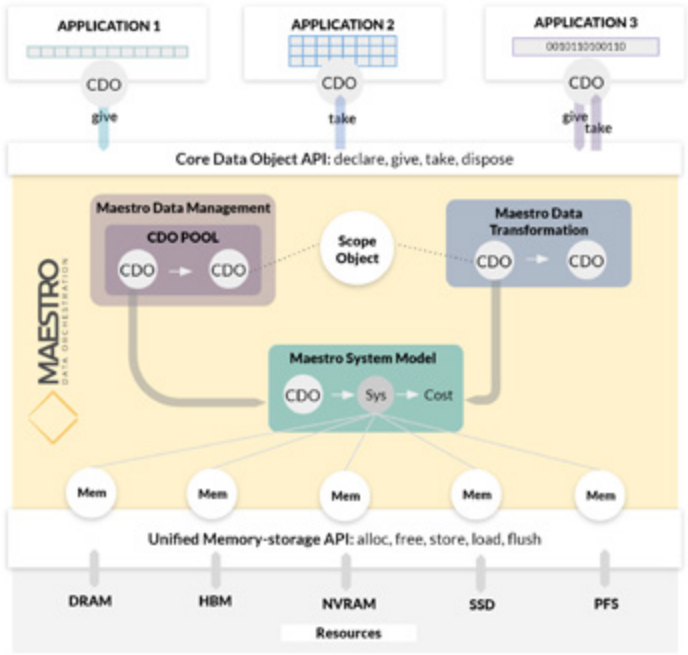

1. Moving data through memory was not always the bottleneck. The software stack that HPC relies upon was built through decades of a different situation – when the cost of performing floating point operations (FLOPS) was paramount. Several decades of technical evolution built a software stack and programming 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.

www.maestro-data.eu
: @maestrodata

COORDINATING ORGANISATION

Forschungszentrum Jülich GmbH, *Germany*

OTHER PARTNERS

- CEA - Commissariat à l'Energie Atomique et aux énergies alternatives, *France*
- Appentra Solutions SL, *Spain*
- CSCS - Swiss National Supercomputing Centre, *Switzerland*
- ETHZ - Eidgenössische Technische Hochschule Zürich, *Switzerland*
- ECMWF - European Centre for Medium-range Weather Forecasts, *United Kingdom*
- Seagate Systems UK Ltd, *United Kingdom*
- Cray Computer GmbH, *Switzerland*

CONTACT

Dirk Pleiter
d.pleiter@fz-juelich.de

CALL

FETHPC-02-2017

PROJECT TIMESPAN

01/09/2018 - 31/08/2021

MANGO

Exploring Manycore Architectures for Next- GeneratiOn HPC systems

MANGO project has been focused on designing and building a large-system prototype for efficient exploration of future HPC manycore architectures including CPU cores, GPU cores and FPGA/HW Accelerator cores.

The prototype embeds highly heterogeneous accelerators in a common infrastructure (a network) which provides 3P design approach (Performance/Power/Predictability) to the applications.

MANGO is deploying all the software stack (at both server side and accelerator side) to let users easily adapt and port their applications to new emerging scenarios with multitude of highly heterogeneous accelerators. For all components an API and the resource manager are provided. MANGO system also includes innovative cooling techniques using termoshypons.

Final MANGO prototype offers versatile exploration platform for future HPC system architectures as well as HPC HW/SW building components to be possibly included in future commercial designs.

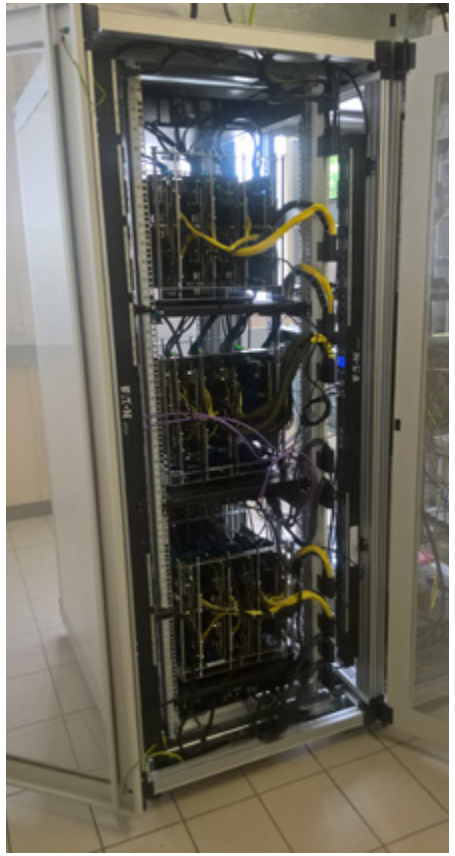

The Mango prototype

www.mango-project.eu
: @mangoEU

COORDINATING ORGANISATION

Universitat Politecnica de Valencia, Spain

OTHER PARTNERS

- Universitat Politècnica de València, *Spain*
- Centro Regionale Information Communication Technology, *Italy*
- Faculty of Electrical Engineering and Computing, University of Zagreb, *HR*
- Politecnico di Milano / Dipartimento di Elettronica, Informazione e Bioingegneria, *Italy*
- École polytechnique fédérale de Lausanne, *Switzerland*
- Pro Design Electronic GmbH, *Germany*
- Eaton Corporation, *France*
- Thales Group, *France*
- Philips, *Netherlands*

CONTACT

José Flich
jflich@disca.upv.es
+34 96 387 70 07 (ext. 75753)

CALL

FETHPC-1-2014

PROJECT TIMESPAN

01/10/2015 - 31/03/2019

NEXTGenIO

Next Generation I/O for the Exascale

The NEXTGenIO project addresses the **I/O performanc**e challenge not only for Exascale, but also for HPC and data intensive computing in general.

NEXTGenIO is developing a **prototype computing platform** that uses on-node non-volatile memory, bridging the latency gap between DRAM and disk. In addition to the hardware that will be built as part of the project, NEXTGenIO is developing the software stack that goes hand-in-hand with this new hardware architecture, and is testing the developments using a set of applications that include both traditional HPC (e.g. CFD and weather) and novel workloads (e.g. machine learning and genomics).

In addition to a detailed understanding of the wide range of use cases for non-volatile memory, key outcomes of NEXTGenIO are:

- Prototype compute nodes with non-volatile memory
- Profiling and debugging tools
- Data and power/energy aware job scheduling system
- Filesystem and object store
- Library for persistent data structures
- Workload benchmark generator & I/O workflow simulator

www.nextgenio.eu
: @nextgenio

COORDINATING ORGANISATION

EPCC – The University of Edinburgh, *UK*

OTHER PARTNERS

- Intel Deutschland GmbH, Germany
- Fujitsu
- Barcelona Supercomputing Center (BSC), *Spain*
- Technische Universität Dresden
- Allinea, now part of Arm
- European Centre for Medium Range Weather Forecasts
- ARCTUR

CONTACTS

Prof Mark Parsons
(Project Coordinator)
m.parsons@epcc.ed.ac.uk

Dr Michèle Weiland
(Project Manager)
m.weiland@epcc.ed.ac.uk
+44 (0)131 651 3580

CALL

FETHPC-1-2014

PROJECT TIMESPAN

01/10/2015 - 30/09/2019

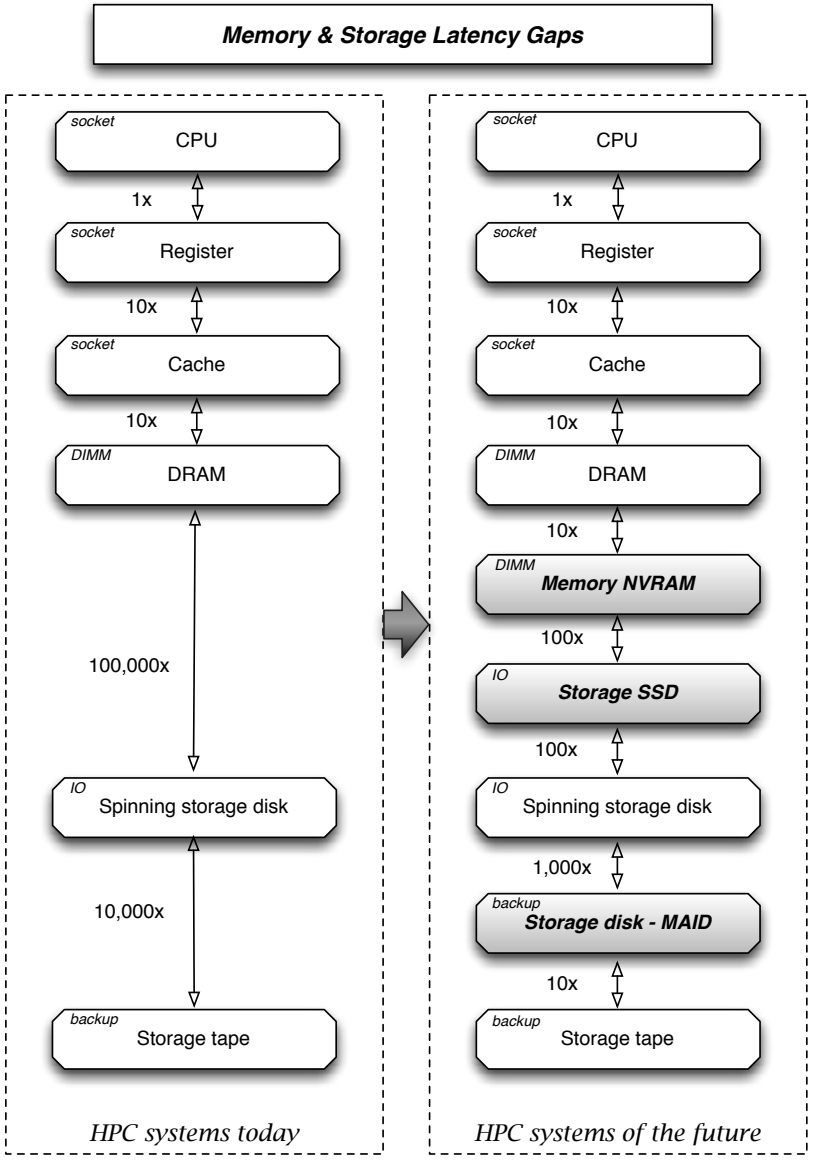

Memory and Storage Latency Gaps

NLAFET

Parallel Numerical Linear Algebra for Future Extreme Scale Systems

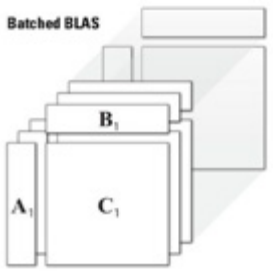

NLAFET MISSION

Today's largest HPC systems have a serious gap between the peak capabilities of the hardware and the performance realized by high-performance computing applications. NLAFET is a direct response to this challenge. NLAFET will enable a radical improvement in the performance and scalability of a wide range of real-world applications, by developing novel architecture-aware algorithms, and the supporting runtime capabilities to achieve scalable performance and resilience on heterogeneous architectures.

The validation and dissemination of results will be done by integrating new software solutions into challenging scientific applications in materials science, power systems, study of energy solutions, and data analysis in astrophysics. The software will be packaged into open-source library modules.

NLAFET SAMPLE RESULTS

The main scientific and technological achievements include advances in the development of Parallel Numerical Linear Algebra algorithms for dense and sparse linear systems and eigenvalue problems, the development of communication avoiding algorithms, an initial assessment of application use cases, and an evaluation of different runtime systems and auto-tuning infrastructures.

The novel software developed and deployed are available via the NLAFET website www.nlafet.eu/software and associated 13 public GitHub repositories structured in five groups: Dense matrix factorizations and solvers; Solvers and tool for standard and generalized eigenvalue problems; Sparse direct factorizations and solvers; Communication optimal algorithms for iterative methods; Cross-cutting tools. The parallel programming models used include MPI, OpenMP, PaRSEC and StarPU. NLAFET software implementations are tested and evaluated on small- to large-scale problems executing on homogeneous systems and some on heterogeneous systems with accelerator hardware. Scientific progress and results of NLAFET have been presented at several international conferences and in scientific journals (see www.nlafet.eu/publications).

NLAFET IMPACT

The main impact is the software available, via GitHub NLAFET repositories, to the scientific community and industry, and thereby providing novel tools for their computational challenges. The work on the batched BLAS specification has already achieved considerable impact with industry in reaching a community standard. The idea is to group multiple independent BLAS operations on small matrices as a single routine (see the figure). The graph shows performance on an Nvidia P100 GPU for 50 to 1000 GEMM operations on matrices of size 16-by-16 up to 4000-by-4000. For example, for matrices of size 256-by-256 around 30 times speedup is obtained compared to CuBLAS GEMM. For large matrix sizes we switch to non-batched BLAS.

Sample applications for batched BLAS can be found in Structural mechanics, Astrophysics, Direct sparse solvers, High-order FEM simulations, and Machine learning.

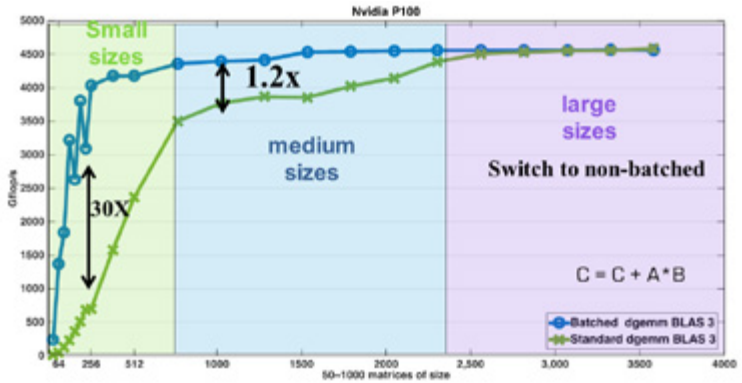

Batched Computations: How do we design and optimise

www.nlafet.eu

COORDINATING ORGANISATION

Umeå University, *Sweden*

OTHER PARTNERS

- University of Manchester, *UK*
- Institut National de Recherche en Informatique et an Automatique, Inria, *France*
- Science and Technology Facilities Council, *UK*

CONTACT

Prof. Bo Kågström
info@nlafet.eu - bokg@cs.umu.se
+46-90-786 5419

CALL

FETHPC-1-2014

PROJECT TIMESPAN

01/11/2015 - 30/04/2019

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.

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.

www.recipe-project.eu
: @EUrecipe

COORDINATING ORGANISATION

Politecnico di Milano, *Italy*

OTHER PARTNERS

- Universitat Politecnica de Valencia, *Spain*
- Centro Regionale Information Communication Technology scrl, *Italy*
- Barcelona Supercomputing Centre (BSC), *Spain*
- Instytut Chemii Bioorganicznej Polskiej Akademii Nauk, *Poland*
- EPFL - Ecole Polytechnique Fédérale de Lausanne, *Switzerland*
- Intelligence behind Things Solutions SRL, *Italy*
- Centre Hospitalier Universitaire Vaudois, *Switzerland*

CONTACT

Prof. William Fornaciari
william.fornaciari@polimi.it

CALL

FETHPC-02-2017

PROJECT TIMESPAN

01/05/2018 - 30/04/2021

SAGE2

Percipient Storage for Exascale Data Centric Computing 2

www.sagestorage.eu
: @SageStorage

COORDINATING ORGANISATION

Seagate Systems UK Ltd, *United Kingdom*

OTHER PARTNERS

- Bull SAS (Atos Group), *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 - 31/08/2021

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.

VECMA

Visual Exploration and Sampling Toolkit for Extreme Computing

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 under¬lying 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 six 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 will comprise as part of it the first VECMA training workshop in January 2020. Ahead of this event, we have run a VECMA hackathon event in association with VECMAtk in September 2019. These events associated with VECMAtk have participation from some of our users from industry and government who use the tools we are developing on HPC at supercomputer centres in Europe at this time.

www.vecma.eu
: @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*
- Bull SAS (Atos Group), *France*
- Stichting Nederlandse Wetenschappelijk Onderzoek Instituten, *Netherlands*
- CBK Sci Con Ltd, United Kingdom
- Universiteit van Amsterdam, *Netherlands*
- Instytut Chemii Bioorganicznej Polskiej Akademii Nauk, Poland

CONTACT

Xuanye Gu
xuanye.gu@ucl.ac.uk

CALL

FETHPC-02-2017

PROJECT TIMESPAN

15/06/2018 - 14/06/2021

VESTEC

Visual Exploration and Sampling Toolkit for Extreme Computing

www.vestec-project.eu
: @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

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.

Centres of Excellence in computing applications

FocusCoE

Concerted action for the European HPC CoEs

European Excellence in HPC Applications

OBJECTIVE

FocusCoE contributes to the success of the EU HPC Ecosystem and the EuroHPC Initiative by supporting the EU HPC 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

- The creation of the HPC CoE Council (HPC3) in May 2019 in Poznań that allows all European HPC CoEs to collectively define an overriding strategy and collaborative implementation for interactions with, and contributions to, the EU 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 has already initiated contacts between CoEs and industry. The focus in the coming period will be on sectorial industrial events and industry associations. A first tranche of industrial success stories from the CoEs is in preparation.
- To instigate concerted action on training by and for the complete set of HPC CoEs: providing consolidating vehicle for user training offered by the CoEs and by PRACE (PATCs) and providing cross-area training to the CoEs (e.g. on sustainable business development). An HPC training stakeholder was organised in October 2019 in Brussels, defining training and education needs and the assessing requirements for HPC training programmes in the EU.
- 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. First steps have been regular newsletters[1] and highlighting CoE achievements in social media. Next steps will be a handbook and a dedicated web site presenting the set of service offerings by CoEs to industrial and academic users.

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

www.focus-coe.eu

COORDINATING ORGANISATION

Scapos AG, *Germany*

OTHER PARTNERS

- 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*
- CEA - Commissariat à l'Energie Atomique et aux Energies Alternatives, *France*

CONTACT

Guy Lonsdale
guy.lonsdale@scapos.com

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/12/2018s - 30/11/2021

BioExcel

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 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 asprires 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.

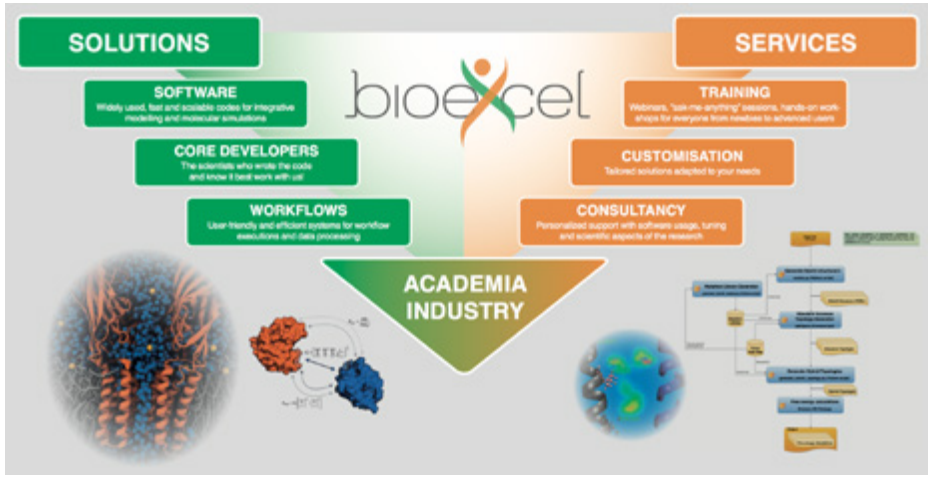

High-performance computing (HPC) and high-throughput computing (HTC) techniques have now reached a level of maturity for many of the widely-used codes and platforms, but taking full advantage of them requires that researchers have the necessary training and access to guidance by experts. The necessary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

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

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

https://bioexcel.eu
: @BioExcelCoE
: 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, *Germany*
- University of Jyväskylä, *Finland*
- The University of Edinburgh, *UK*
- Max Planck Gesellschaft, *Germany*
- Across Limits, *Malta*
- Barcelona Supercomputing Centre (BSC), *Spain*
- Ian Harrow Consulting, *UK*
- Norman Consulting, *Norway*
- Nostrum Biodiscovery, *Spain*

CONTACT

Erwin Laure
erwinl@pdc.kth.se

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/01/2019 - 31/12/2021

ChEESE

Centre of Excellence for Exascale and Solid Earth

ChEESE

OBJECTIVE

ChEESE will address extreme computing scientific and societal challenges 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 challenging ambition is to prepare 10 open-source flagship codes to solve Exascale problems on computational seismology, magnetohydrodynamics, physical volcanology, tsunamis, and data analysis and predictive techniques, including machine learning and predictive techniques from monitoring earthquake and volcanic activity. The selected codes will be audit 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 will consider also code inter-kernel aspects of simulation workflows like data management and sharing, I/O, post-process and visualization.

In parallel with these transversal activities, ChEESE will sustent on three vertical pillars. First, it will develop pilot demonstrators for scientific challenging 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 and physics-based hazard assessments for seismics, volcanoes and tsunamis.

Second, 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. Selected pilots will be tested in an operational environment to make them available to a broader user community. Additionally, and in collaboration with the European Plate Observing System (EPOS), 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 user communities.
Finally, the third pillar of ChEESE aims at acting 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.

www.cheese-coe.eu

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*
- Bull SAS (Atos Group), *France*

CONTACT

Arnau Folch
arnau.folch@bsc.es

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/11/2018 - 31/10/2021

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, lightweight 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.

www.compbiomed.eu
: @bio_comp

COORDINATING ORGANISATION

University College London, *UK*

OTHER PARTNERS

- University of Amsterdam, *Netherlands*
- The University of Edinburgh, *UK*
- 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*
- Bull SAS (Atos Group), *France*
- Janssen Pharmaceutica, *Belgium*

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 CECAM's distributed European Network of simulation science nodes and the physical computational e-infrastructure of PRACE. We are a partnership of **16 CECAM Nodes, 4 PRACE** centres, 12 industrial partners and one Centre for Industrial Computing (Hartree Centre).

Our goals are to:

- **Develop software modules** to be used in academia and industry to solve simulation and modelling problems. Interface those software modules with standard codes and tune them to run on the next generation of exascale computers;
- **Train scientists from industry and academia** on the development of methods and software scaling towards the high end of HPC systems;
- **Provide multidisciplinary, coordinated, top level applied discussions to support industrial end-users** (both large multinationals and SMEs) in their use of simulation and modelling. This includes workshops with industry to identify areas of mutual interest, training courses for industrialists, on specific families of methods, software and HPC tools, as well as the possibility to engage in direct discussions.

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.

The impact of E-CAM so far:

- More than one hundred certified software modules that are open access and easily available for the industrial and academic communities through our software repository.
- Our Extended Software Development Workshops (ESDWs) trained 230 people so far, in advanced computational methods, good practices for code development, documentation and maintenance
- An online training infrastructure is in place to support the develop software for extreme-scale hardware
- Several industrially relevant pilot projects have started in E-CAM. Developments are reported on our website
- Our State of the Art and Scoping Workshops are important access points for companies and SMEs to expertise and discussions on leading-edge simulation and modelling techniques
- E-CAM works on software development projects that enable an HPC practice with potential for transfer into industry. Examples are: GPU re-write of DL_MESO code for mesoscale simulations; the development of a load-balancing library specifically targeting MD applications, and of an HTC library; among other efforts.

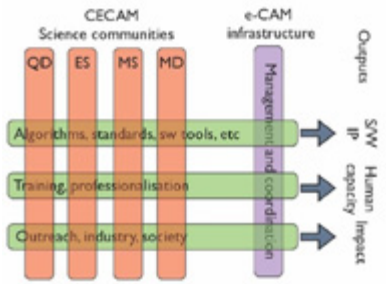

Positioning of the e-CAM infrastructure

www.e-cam2020.eu
: @ECAM2020

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, *s*
- CSC-IT Centre for Science Ltd, *Finland*
- Irish Centre for High-End Computing, *Ireland*

CONTACT

Luke Drury
(Project Chair)
ld@cp.dias.ie
+353 1 6621333 ext 321

Ana Catarina Mendonça
(Project Coordinator)
ana.mendonca@epfl.ch
+41 21 693 0395

CALL

EINFRA-5-2015

PROJECT TIMESPAN

01/10/2015 - 30/09/2020

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:
www.eocoe.eu/services

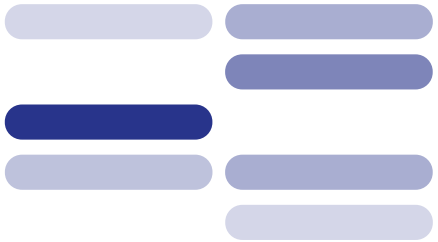

www.eocoe.eu
: @EoCoE

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

CONTACT

Edouard Audit
contact@eocoe.eu
+33 1 69 08 96 35

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/01/2019 - 31/12/2021

ESiWACE

Centre of Excellence in Simulation of Weather and Climate in Europe

www.esiwace.eu
: @ESIWACE

COORDINATING ORGANISATION

Deutsches Klimarechenzentrum (DKRZ), *Germany*

OTHER PARTNERS

- CNRS- Centre National de la Recherche Scientifique, *France*
- Barcelona Supercomputing Centre (BSC), *Spain*
- European Centre for Medium-Range Weather Forecasts, *United Kingdom*
- 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, *UK*
- Fondazione Centro Euro-Mediterraneo sui Cambiamenti Climatici, *Italy*
- The University of Reading, *UK*
- STFC - Science and Technology Facilities Council, *UK*
- Bull SAS (Atos Group), *France*
- Seagate Systems UK Limited, *UK*
- DWD - Deutscher Wetterdienst, *Germany*
- Arm Limited, *UK*

CONTACT

Joachim Biercamp
esiwace@lists.dkrz.de

CALL

EINFRA-5-2015

PROJECT TIMESPAN

01/09/2015 - 31/08/2019

OBJECTIVE

Science driver for ESiWACE is the establishment of global weather and climate models that allow simulating convective clouds and small-scale ocean eddies to enable reliable simulation of high-impact regional events. This is not affordable today, considering the required compute power, data loads and tight production schedules (for weather forecasts), and will require exascale computing. We address and quantify the computational challenges in achieving this science case. The centre of excellence ESiWACE further leverages two European networks for this purpose: ENES (European Network for Earth System Modelling) and ECMWF (European Centre for Medium-Range Weather Forecasts).

Weather and climate models are being pushed towards global 1-2.5 km resolution, cf. Figure 3 2.5-km global simulation using the ICON model. Performance predictions concerning scalability of the models at exascale are made, and efficient tools for data and workflow management are developed. Other contributions of ESiWACE are handbooks on system and software stacks that are required for installation and operation of the various complex models on HPC platforms. This will substantially improve usability and portability of the weather and climate models.

Future work in the scope of ESiWACE comprises very high-resolution runs of coupled atmosphere-ocean models based on the models ICON and EC-Earth.

2.5-km global simulation using the ICON model

ESiWACE2

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

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.

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.

www.esiwace.eu
: @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*
- Bull SAS (Atos Group), *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

Joachim Biercamp
esiwace@lists.dkrz.de

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/01/2019 - 31/12/2022

EXCELLERAT

OBJECTIVE

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 EXCELLERAT activity brings together the premium European expertise to establish a Centre of Excellence in Engineering Applications on HPC with a broad service portfolio, paving the way for the evolution towards Exascale. All perfectly in line with the European HPC Strategy as implemented through the EuroHPC Joint Undertaking. To fulfil its mission, EXCELLERAT will focus its developments on six carefully chosen reference applications (Nek5000, Alya, AVBP, Fluidity, FEniCS, Flucs), 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 evolution of applications towards execution on high scale 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, optimization, 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 for not only the six chosen reference applications, but even going beyond. Thus (but not only for that purpose), EXCELLERAT will extend the recipients of its developments beyond the consortium and interact via diverse mechanisms to integrate external stakeholders of its value network into its evolution.

Base pillars of the EXCELLERAT Centre of Excellence

www.excellerat.eu
: @EXCELLERAT_CoE

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 - 30/11/2021

HiDALGO

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, multi-dimensional 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.

Project website pending

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 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.

PROJECT DESCRIPTION

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: BigDFT, CP2k, FLEUR, Quantum ESPRESSO, SIESTA, YAMBO. In addition, it contributes to the development of the AIIDA materials informatics infrastructure and the Materials Cloud environment.

The main actions include:

1. code and library restructuring: including modularizing and adapting for heterogeneous architectures, as well as adoption of new algorithms;
2. codesign: 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. widening access to codes, training, and transferring know-how to user communities.

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

http://www.max-centre.eu
: @max_center2

COORDINATING ORGANISATION

CNR Nano, Modena, *Italy*

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* (Carlo Cavazzoni)
- Barcelona Supercomputing Centre (BSC), *Spain* (Stephan Mohr)
- CSCS ETH Zürich, *Switzerland* (Joost VandeVondele)
- CEA, *France* (Thierry Deutsch)
- UGent, *Belgium* (Stefaan Cottenier)
- E4 Computer Engineering, *Italy* (Fabrizio Magugliani)
- Arm, *UK* (Filippo Spiga)
- ICTP UNESCO, Trieste, *Italy* (Ivan Girotto)
- Trust-IT, *Italy* (Silvana Muscella)

CONTACT

Elisa Molinari
Director
elisa.molinari@max-centre.eu
+39 059 2055 629

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/12/2018 - 30/11/2021

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 project resulted in more than 120 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.

The objective of POP2 is to continue and improve the POP project operating a Centre of Excellence in Computing Applications in the area of Performance Optimisation and Productivity with a special focus on very large scale towards exascale. POP2 will continue the service oriented activity, giving code developers, users and infrastructure operators an impartial external view that will help them improve their codes. We will still target the same categories of users and focus on identifying performance issues and proposing techniques that can help applications in the direction of exascale. Our objective is to perform 180 services over a three-year period.

https://pop-coe.eu
: @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*
- Vysoka Skola Banska - Technicka Univerzita Ostrava, *Czech Republic*

CONTACT

Jesus Labarta
jesus.labarta@bsc.es

Judit Gimenez
judit@bsc.es

CALL

INFRAEDI-02-2018

PROJECT TIMESPAN

01/12/2018 - 30/11/2021

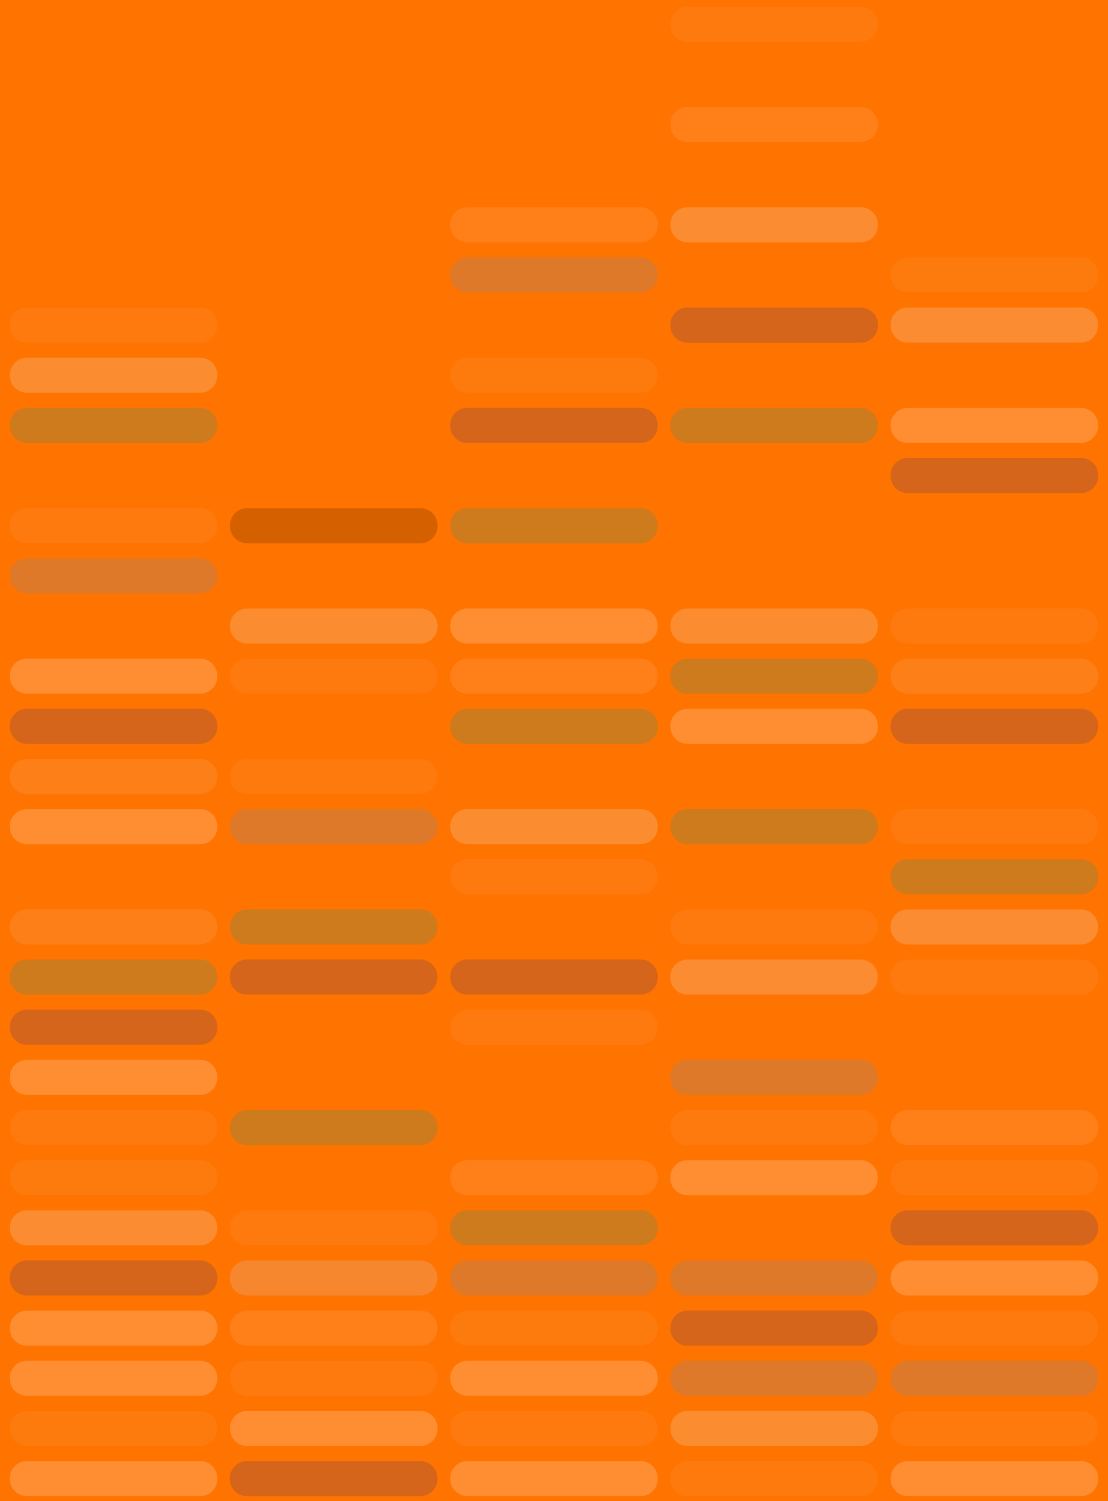

European Processor

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 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 an advanced experimental platform towards exascale system for Europe in 2021 and exascale in 2022-2023;
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 centers, thus ensuring the economic viability of the initiative.

EPI gathers 27 partners from 10 European countries, with a wide range of expertise and background: HPC, supercomputing centers, 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-sources 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.

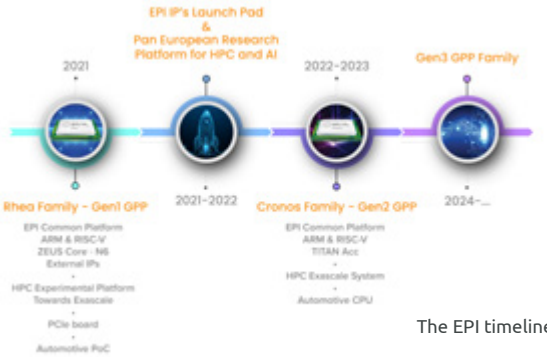

The EPI timeline

CONTACT
contact@european-processor-initiative.eu

CALL
SGA-LPMT-01-2018

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

http://european-processor-initiative.eu
: @EUProcessor
: european-processor-initiative
: www.youtube.com/channel/UCGvQcTosJdWhd013SHnIbpA/featured

COORDINATING ORGANISATION
Bull SAS (Atos group), *France*

OTHER PARTNERS

- BSC, Barcelona supercomputing center - centro nacional de supercomputacion, *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*

Mont-Blanc 2020

European scalable, modular and power-efficient HPC processor

MONT-BLANC 2020

Following on from the three successive Mont-Blanc projects since 2011, the three core partners Arm, Barcelona Supercomputing Center and Bull (Atos Group) have united again to trigger the development of the next generation of industrial processor for Big Data and High Performance Computing. The Mont-Blanc 2020 consortium also includes CEA, Forschungszentrum Jülich, Kalray, and SemiDynamics.

The Mont-Blanc 2020 project intends to pave the way to the future low-power European processor for Exascale. To improve the economic sustainability of the processor generations that will result from the Mont-Blanc 2020 effort, the project includes the analysis of the requirements of other markets. The project's strategy based on modular packaging would make it possible to create a family of SoCs targeting different markets, such as "embedded HPC" for autonomous driving. The project's actual objectives are to:

- define a low-power System-on-Chip architecture targeting Exascale;
- implement new critical building blocks (IPs) and provide a blueprint for its first-generation implementation;
- deliver initial proof-of-concept demonstration of its critical components on real life applications;
- explore the reuse of the building blocks to serve other markets than HPC, with methodologies enabling a better time-predictability, especially for mixed-critical applications where guaranteed execution & response times are crucial.

Mont-Blanc 2020 will provide new IPs, such as a new low-power mesh interconnect based on the Coherent Hub Interface (CHI) architecture. Mont-Blanc 2020 will demonstrate the performance with a prototype implementation in RTL for some of its key components, and demonstration on an emulation platform.

The Mont-Blanc 2020 project is at the heart of the European exascale supercomputer effort, since most of the IP developed within the project will be reused and productized in the European Processor Initiative (EPI).

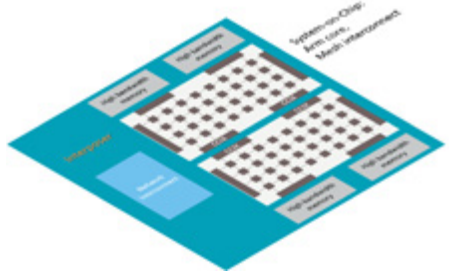

The Mont-Blanc 2020 physical architecture

www.montblanc-project.eu
: @MontBlanc_Eu
: @mont-blanc-project

COORDINATING ORGANISATION

Bull (Atos group), *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

PROJECT TIMESPAN

01/12/2017 - 30/11/2020

CALL

ICT-05-2017

A Key Output: the MB2020 demonstrator

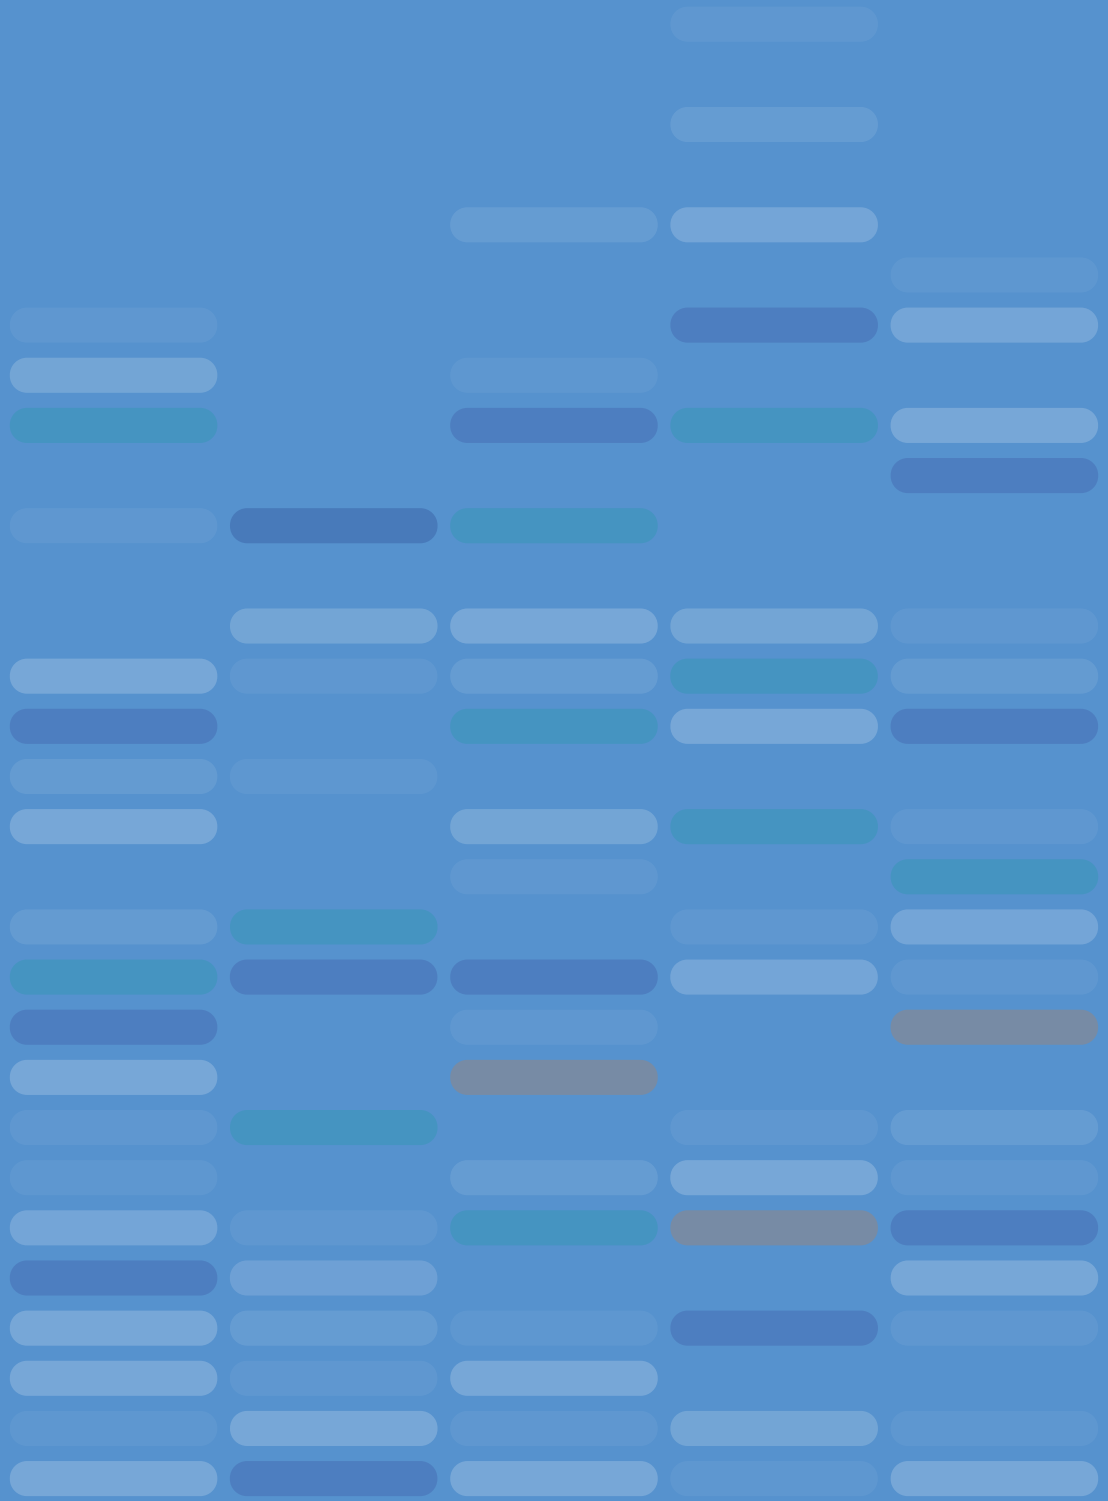

HPC and Big Data testbeds

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

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.

www.cybele-project.eu
: @CYBELE_H2020
: @Cybele-H2020
: @Cybele-project

COORDINATING ORGANISATION

Waterford Institute of Technology, *Ireland*

OTHER PARTNERS

- Barcelona Supercomputing Centre (BSC), *Spain*
- Bull SAS (Atos group), *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*
- Oloklиromena Pliroforiaka Sistimataae, *Greece*
- Kobenhavns Universitet, *Denmark*
- EVENFLOW, *Belgium*
- Open Geospatial Consortium (Europe) Ltd, *United Kingdom*

CALL

ICT-11-2018-2019

PROJECT TIMESPAN

01/01/2019 - 31/12/2021

CONTACT

cybele.h2020@gmail.com
sdavy@tssg.org

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 will be ready to be integrated in current and new biomedical platforms or applications.

The toolkit is composed of two core libraries, the European Distributed Deep Learning Library (EDDLL) and the European Computer Vision Library (ECVL) and a dedicated front-end for facilitating the use of the libraries to computer and data scientists.

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.

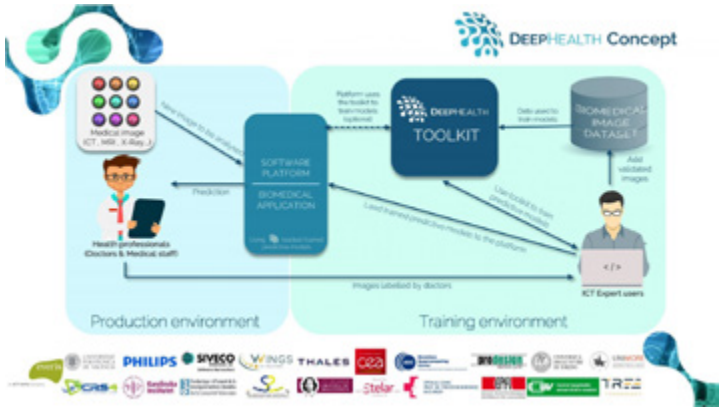

14 PILOTS AND SEVEN PLATFORMS TO VALIDATE THE DEEPHEALTH PROPOSED INNOVATIONS

The DeepHealth innovations will be validated in 14 pilot test-beds through the use of seven different biomedical and AI software platforms provided by partners. The libraries will be 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.

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 825111, DeepHealth Project.

CALL
ICT-11-2018-2019

PROJECT TIMESPAN
01/01/2019 - 31/12/2021

https://deephealth-project.eu
: @DeepHealthEU
: @DeepHealthEU
: @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 Politecnica de Valencia, *Spain*
- WINGS ICT Solutions Information & Communication Technologies IKE, *Greece*

CONTACT
Monica Caballero
(Project manager)
monica.caballero.galeote@everis.com

Jon Ander Gómez
(Technical manager)
jon@dsic.upv.es

EVOLVE

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

evolve

Leading the Big Data Revolution

EVOLVE is a pan European Innovation Action with 19 key partners from 11 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.

Therefore, one key differentiator of EVOLVE is the ambition to tackle heterogeneity: in the infrastructure itself, from data processing to data movement aspects, for workflow optimization. Heterogeneity is also addressed at the level of user needs with support of cloud and big data tools on HPC platform.

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 intends to build and demonstrate the proposed testbed with real-life, massive datasets from demanding applications areas. To realize this vision, EVOLVE brings together technologies and pilot partners from EU industry with demonstrated experience, established markets, and vested interest. Furthermore, EVOLVE will conduct a set of 10-15 Proof-of-Concepts with stakeholders from the Big Data value chain to build up digital ecosystems to achieve a broader market penetration.

www.evolve-h2020.eu
: @Evolve_H2020
: @evolve-h2020

COORDINATING ORGANISATION

Datadirect Networks France, *France*

OTHER PARTNERS

- AVL LIST GmbH, *Austria*
- BMW - Bayerische Motoren Werke AG, *Germany*
- Bull SAS (Atos group), *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
Pierre-Aimé Agnel
pagnel@ddn.com

CALL

ICT-11-2018-2019

PROJECT TIMESPAN

01/12/2018 - 30/11/2021

LEXIS

Large-scale EXecution for Industry & Society

The increasing quantities of data generated by modern industrial and business processes pose enormous challenges for organisations seeking to glean knowledge and understanding from the data. Combinations of HPC, Cloud and Big Data technologies are key to meet the increasingly diverse needs of large and small organisations alike. Critically, access to powerful compute platforms for SMEs - which has been difficult due to both technical and financial reasons - may now be possible.

LEXIS (Large-scale EXecution for Industry & Society) project will build an advanced engineering platform at the confluence of HPC, Cloud and Big Data which will leverage large-scale geographically-distributed resources from existing HPC infrastructure, employ Big Data analytic solutions and augment them with Cloud services. Driven by the requirements of the pilots, the LEXIS platform will build on best of breed data management solutions (EUDAT) and advanced, distributed orchestration solutions (TOSCA), augmenting them with new efficient hardware capabilities in the form of Data Nodes and federation, usage monitoring and accounting/billing supports to realise an innovative solution oriented to stimulate the interest of European industries and creating an ecosystem of organisations that could benefit from the implemented platform of LEXIS that sets up a connection between HPC, HPDA and Data Management.

The consortium will develop a demonstrator with a significant Open Source dimension including validation, test and documentation. It will be validated in the large-scale pilots – in the 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. LEXIS will organise a specific call stimulating the project framework adoption and stakeholders' engagement on the targeted pilots.

LEXIS will promote the solution to the HPC, Cloud and Big Data sectors maximising impact through targeted and qualified communication. LEXIS brings together a consortium 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.

http://lexis-project.eu
: @LexisProject
: @LEXIS2020
: @lexis-project

COORDINATING ORGANISATION

IT4Innovations, VSB – Technical University of Ostrava, *Czechia*

OTHER PARTNERS

- Alfred Wegener Institute Helmholtz Centre for Polar and Marine Research, *Germany*
- Atos BDS (Bull), *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*
- GE AVIO SRL, *Italy*
- Helmholtz Zentrum Potsdam Deutsches Geoforschungs Zentrum GFZ, *Germany*
- 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 - 30/06/2021

International Cooperation

ENERXICO

This project is going to apply exascale HPC techniques to different energy industry simulations of critical interest for Mexico. ENERXICO will give solutions for oil & gas industry in upstream, midstream and downstream problems, wind energy industry and combustion efficiency for transportation. This project brings together the main stakeholders of the energy industry in Mexico and European energy companies working at the Mexican market, jointly with the main European HPC company. The main objectives of the project are: - to develop beyond the state of the art high performance simulation tools that can help the modernization of the Mexican energy industry and are also of interest for European companies. These simulation tools should be ready to be used in exascale computers that are in the roadmap of European IT companies. - to improve the cooperation between industries from EU and Mexico. - to improve the cooperation between the leading research groups in EU and Mexico

www.enerxico-project.eu
: @enerxico-project

COORDINATING ORGANISATION

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

OTHER PARTNERS

- Bull SAS (Atos Group), *France*
- Centro de Investigacion y de Estudios Avanzados del Instituto Politecnico Nacional (CINVESTAV), *Mexico*
- Centro de Investigaciones Energeticas, Medioambientales y Tecnologicas (CIEMAT), *Spain*
- Instituto Mexicano del Petroleo (IMP), *Mexico*
- Instituto Nacional de Investigaciones Nucleares (ININ), *Mexico*
- Instituto Politecnico Nacional (IPN-ESUA), *Mexico*
- Technische Universität München (TUM), *Germany*
- Universidad Autonoma Metropolitana (UAM-A), *Mexico*
- Instituto de Ingenieria de la Universidad Nacional Autonoma de Mexico (UNAM-IINGEN), *Mexico*
- Universitat Politecnica de Valencia (UPV), *Spain*
- Université Grenoble Alpes (UGA), *France*
- IBERDOLA , *Spain*
- REPSOL SA, *Spain*

PROJECT OBSERVER

Petroleos Mexicanos Exploracion y Produccion (PEMEX), *Mexico*

CONTACT

Rose Gregorio
rose.gregorio@bsc.es

Madeleine Gray
madeleine.gray@bsc.es

CALL

FETHPC-01-2018

PROJECT TIMESPAN

01/06/2019 - 31/05/2021

HPCWE

High performance computing for wind energy

EU and Brazil -
Together Bringing Wind Energy Simulation towards Exascale

HPC WE

Wind energy has become an 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. 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; and (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 13 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.

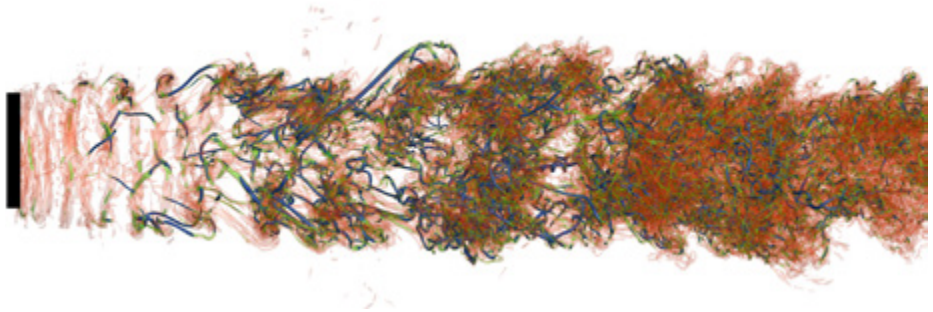

Simulation of the wake flow downstream of a wind turbine

www.hpcwe-project.eu

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

Xuerui Mao
Project Coordinator
Xuerui.Mao@nottingham.ac.uk
+441157486020

CALL

FETHPC-01-2018

PROJECT TIMESPAN

01/06/2019 - 31/05/2021

Other projects

ARGO

WCET-Aware Parallelization of Model-Based Applications for Heterogeneous Parallel Systems

Increasing performance and reducing cost, while maintaining safety levels and programmability are the key demands for embedded and cyber-physical systems in European domains, e.g. aerospace, automation, and automotive. For many applications, the necessary performance with low energy consumption can only be provided by customized computing platforms based on heterogeneous many-core architectures. However, their parallel programming with time-critical embedded applications suffers from a complex toolchain and programming process.

The ARGO (WCET-Aware Parallelization of Model-Based Applications for Heterogeneous Parallel Systems) research project addressed this challenge with a holistic tool chain for programming heterogeneous multi- and many-core architectures using automatic parallelization of model-based real-time applications. ARGO enhances WCET-aware automatic parallelization by a cross-layer programming approach combining automatic tool-based and user-guided parallelization to reduce the need for expertise in programming parallel heterogeneous architectures. An integrated multi-core WCET analysis computes safe bounds for the WCET of parallelized programs, which qualifies the approach for the use in the hard real-time domain. The ARGO approach was assessed and demonstrated by prototyping comprehensive time-critical applications from both aerospace and industrial automation domains on customized heterogeneous many-core platforms.

The challenging research and innovation action was implemented by the unique ARGO consortium that brings together industry, leading research institutes and universities. High class SMEs such as Scilab Enterprises, AbsInt GmbH and emmtrix Technologies GmbH contributed their diverse know-how in model-based design environments, WCET calculation and automated software parallelization. The academic partners contributed their outstanding expertise in code transformations, automatic parallelization and multi-core WCET analysis.

http://argo-project.eu

COORDINATING ORGANISATION

Karlsruher Institut für Technologie (KIT), *Germany*

OTHER PARTNERS

- ABSINT Angewandte Informatik GmbH, *Germany*
- DLR - Deutsches Zentrum für Luft - und Raumfahrt EV, *Germany*
- EMMTRIX Technologies GmbH, Germany
- Frauhofer Gesellschaft zur Förderung der angewandten Forschung E.V., *Germany*
- SCILAB Entreprises, *France*
- Technologiko Ekpaideftiko Idryma (TEI) Dytikis Elladas, *Greece*
- Université de Rennes I, *France*

CONTACT

Prof. Dr.-Ing. Dr. h. c. Jürgen Becker
Coordinator
mail@argo-project.eu

CALL

ICT-04-2015

PROJECT TIMESPAN

01/01/2016 - 31/03/2019

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.

www.examode.eu
: @examode
: @examode.eu

COORDINATING ORGANISATION

University of Applied Sciences Western Switzerland, HES-SO, *Switzerland*

OTHER PARTNERS

- Azienda Ospedaliera per l'Emergenza Cannizzaro, *Italy*
- MicroscopeIT Sp. z o.o., *Poland*
- Ontotext AD, *Bulgaria*
- Radboud University Medical Center, *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

M2DC

The Modular Microserver DataCentre (M2DC) project developed a new class of energy-efficient and Total Cost of Ownership (TCO)-optimized appliances with built-in efficiency and dependability enhancements. The appliances are easy to integrate with a broad ecosystem of management software and fully software-defined to enable optimization for a variety of future demanding applications in a cost-effective way. The highly flexible M2DC server platform enables customization and smooth adaptation to various types of applications, while advanced management strategies and system efficiency enhancements (SEE) can be used to improve energy efficiency, performance, security and reliability. Data centre capable abstraction of the underlying heterogeneity of the server is provided by an OpenStack-based middleware.

In more detail, M2DC developed a resource-efficient, highly scalable heterogeneous microserver platform able to integrate various microservers and hardware accelerators. Within the project, the consortium worked on microservers that use high-performance ARMv8 server processors; low-power ARMv8 embedded/mobile processors; x86 processors; and accelerators (MPSoCs, GPUs, and FPGAs). This includes a new microserver designs based on ARMv8 and Intel Stratix 10. The platform includes an optimized middleware for deployment of the optimized appliances. It may contain up to 144 low power and 27 high performance microservers.

The main features of the platform include energy efficiency, built-in hardware accelerated functions and lower costs of deployment and upgrades. Baseline benchmarks show the high potential of M2DC appliances for the targeted applications including photo finishing systems, IoT data processing, cloud computing, artificial neural networks and HPC. They reached in many cases significant efficiency and Total Cost of Ownership (TCO) gains, higher than 50%.

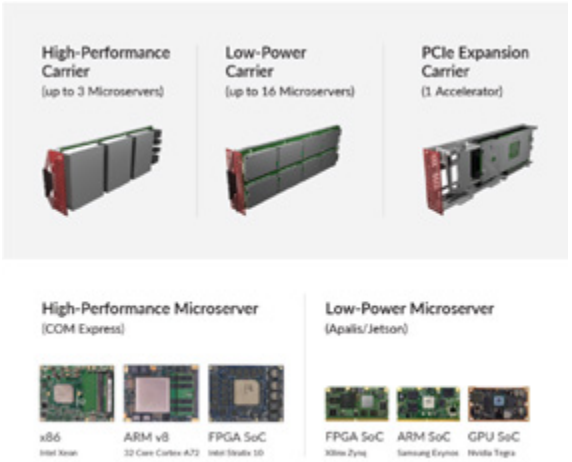

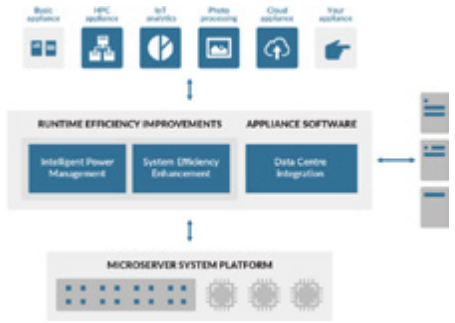

http://m2dc.eu
: @M2DC_Project
: @M2DCproject
: @m2dc

COORDINATING ORGANISATION

Instytut Chemii Bioorganicznej Polskiej Akademii Nauk – Poznań Supercomputing and Networking Center (PSNC), *Poland*

OTHER PARTNERS

- Alliance Services Plus, *France*
- Arm Limited, *United Kingdom*
- Beyond.PL SP ZOO, *Poland*
- CEWE Stiftung & CO KGaA, *Germany*
- Christmann Informationstechnik + Medien GmbH & CO KG, *Germany*
- CEA - Commissariat à l'Energie Atomique et aux énergies alternatives, *France*
- Huawei Technologies Düsseldorf GmbH, *Germany*
- Offis EV, *Germany*
- Politecnico di Milano, *Italy*
- REFLEX CES, *France*
- Universität Bielefeld, *Germany*
- Vodafone Automotive Italia S.P.A., *Italy*
- XLAB d.o.o., *Slovenia*

CONTACT

Ariel Oleksiak
Project Coordinator
ariel@man.poznan.pl

CALL

ICT-04-2015

PROJECT TIMESPAN

01/01/2016 - 30/06/2019

OPRECOMP

Open transPREcision COMPuting

Guaranteed numerical precision of each elementary step in a complex computation has been the mainstay of traditional computing systems for many years. This era, fueled by Moore's law and the constant exponential improvement in computing efficiency, is at its twilight: from tiny nodes of the Internet-of-Things, to large HPC computing centers, sub-picoJoule/operation energy efficiency is essential for practical realizations. To overcome the "power wall", a shift from traditional computing paradigms is now mandatory.

OPRECOMP aims at demolishing the ultra-conservative "precise" computing abstraction and replacing it with a more flexible and efficient one, namely transprecision computing. OPRECOMP will investigate the theoretical and practical understanding of the energy efficiency boost obtainable when accuracy requirements on data being processed, stored and communicated can be lifted for intermediate calculations. While approximate computing approaches have been used before, in OPRECOMP for the first time ever, a complete framework for transprecision computing, covering devices, circuits, software tools, and algorithms, along with the mathematical theory and physical foundations of the ideas will be developed. It not only will provide error bounds with respect to full precision results, but it will also enable major energy efficiency improvements even when there is no freedom to relax end-to-end application quality-of-results.

The mission of OPRECOMP is to demonstrate, by using physical demonstrators, that this idea holds in a huge range of application scenarios in the domains of IoT, Big Data Analytics, Deep Learning, and HPC simulations: from the sub-milliWatt to the MegaWatt range, spanning nine orders of magnitude. In view of industrial exploitation, we will prove quality and reliability and demonstrate that transprecision computing is the way to think about future systems.

http://oprecomp.eu
: @oprecompProject

COORDINATING ORGANISATION

IBM Research GmbH, *Switzerland*

OTHER PARTNERS

- Alma Mater Studiorum - Università di Bologna, *Italy*
- CINECA Consorzio Interuniversitario, *Italy*
- CEA - Commissariat à l'Energie Atomique et aux énergies alternatives, *France*
- ETHZ - Eidgenössische Technische Hochschule Zürich, *Switzerland*
- Greenwavs Technologies, *France*
- Technische Universität Kaiserslautern, *Germany*
- The Queen's University of Belfast, *United Kingdom*
- Università degli Studi di Perugia, *Italy*
- Universitat Jaume I de Castellon, *Spain*

CONTACT

Cristiano MALOSSI, PhD
acm@zurich.ibm.com

CALL

FETPROACT-01-2016

PROJECT TIMESPAN

01/01/2017 - 31/12/2020

PROCESS

PROviding Computing solutions for ExaScale ChallengeS

The PROCESS demonstrators will pave the way towards exascale data services that will accelerate innovation and maximise the benefits of these emerging data solutions. The main tangible outputs of PROCESS are very large data service prototypes, implemented using a mature, modular, generalizable open source solution for user friendly exascale data. The services will be thoroughly validated in real-world settings, both in scientific research and in advanced industry deployments.

To achieve these ambitious objectives, the project consortium brings together the key players in the new data-driven ecosystem: top-level HPC and big data centres, communities – such as LOFAR – with unique data challenges that the current solutions are unable to meet and experienced e-Infrastructure solution providers with an extensive track record of rapid application development. In addition to providing the service prototypes that can cope with very large data, PROCESS addresses the work programme goals by using the tools and services with heterogeneous use cases, including exascale learning on medical images, airline revenue management and open data for global disaster risk reduction. This diversity of user communities ensures that in addition to supporting communities that push the envelope, the solutions will also ease the learning curve for broadest possible range of user communities. In addition, the chosen open source strategy maximises the potential for uptake and reuse, together with mature software engineering practices that minimise the efforts needed to set up and maintain services based on the PROCESS software releases.

The PROCESS project has achieved its mid-term objectives and accelerated the development in the use case's communities. Implementations for LOFAR will also be valuable for the upcoming Square Kilometre Array (SKA) project, which increases the underlying challenge in the next years to Exabytes. The PROCESS solutions are deployed in three supercomputing centres across Europe, in Munich, Amsterdam and Krakow, enabling already today a testbed for exascale applications of tomorrow.

The figure shows the detailed PROCESS architecture, with the main components end-user portal, orchestration, data and compute service. The heart of the data service is LOBCDER and its Micro-Infrastructure-Containers, enabling a low-overhead solution for the data processing challenge. The already connected centres allow PROCESS to include applications based on HPC, Cloud and Accelerator-systems. The future main tasks include the support of EOSC-based solutions and the optimisation of the already deployed services.

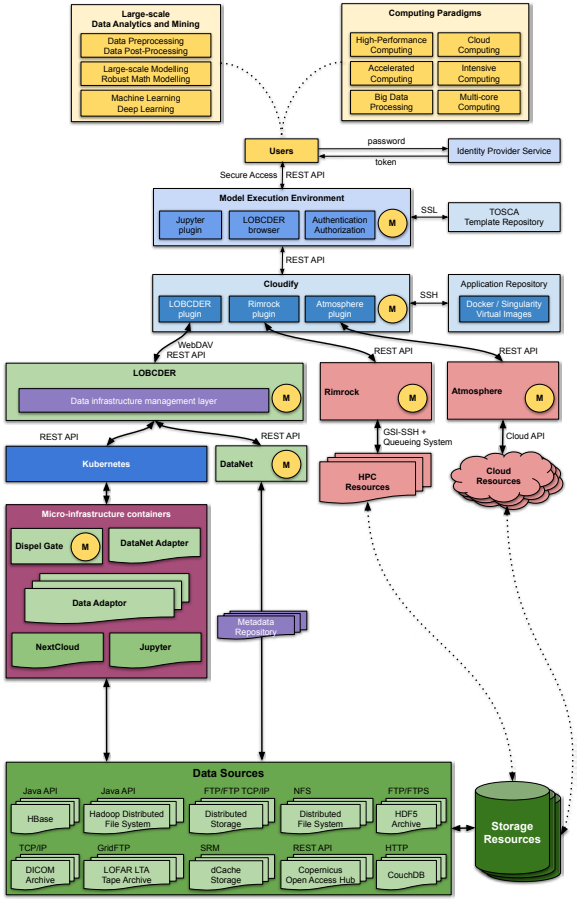

www.www.process-project.eu
: @PROCESS_H2020

COORDINATING ORGANISATION

Ludwig-Maximilians-Universität München, *Germany*

OTHER PARTNERS

- Universiteit van Amsterdam, *Netherlands*
- Stichting Netherlands EScience Centre, *Netherlands*
- Haute Ecole Spécialisée de Suisse Occidentale, *Switzerland*
- Lufthansa Systems GmbH & CO KG, *Germany*
- Inmark Europa SA, *Spain*
- Ustav Informatiky, Slovenska Akademia Vied, *Slovakia*
- Akademia Gorniczo-Hutnicza IM. Stanislawa Staszica W Krakowie, *Poland*

CONTACT

Dieter Kranzlmüller
Coordinator

Maximilian Höb
Technical-Coordinator and Contact
contact@process-project.eu

CALL

EINFRA-21-2017

PROJECT TIMESPAN

01/11/2017 - 31/10/2020

Ecosystem development

EUROLAB-4-HPC-2

Consolidation of European Research Excellence in Exascale HPC Systems

Eurolab4HPC

Europe has made significant progress in becoming a leader in large parts of the HPC ecosystem: from industrial and scientific application providers via system software to Exascale systems, bringing together technical and business stakeholders. Despite such gains, excellence in research in high performance computing systems is fragmented across Europe and opportunities for synergy are missed.

At this moment, there is fierce international competition to sustain long-term leadership in HPC technology and there remains much to do. Since high-performance computing (HPC) is a vital technology for European economic competitiveness and scientific excellence, it is indispensable for innovating enterprises, scientific researchers, and government agencies while it generates new discoveries and breakthrough products and services. By uniting the HPC community, Eurolab4HPC aims to build a robust ecosystem, hence building a sustainable foundation for HPC innovation in Europe.

This will be realised by means of four short-term objectives:

STRUCTURE THE HPC COMMUNITY Uniting the HPC Community, Eurolab4HPC aims to build a robust ecosystem, hence building a sustainable foundation for HPC innovation in Europe.	PROMOTE ENTREPRENEURSHIP Entrepreneurial training, business prototyping, business plan development and helping with funding.
DISSEMINATE COMMUNITY NEWS Investing in dissemination activities, creating a stronger Eurolab4HPC brand and HPC ecosystem.	STIMULATE TECHNOLOGY TRANSFER Connecting with other technology transfer activities and providing seed funding for HPC technology transfer projects.

OPEN SOURCE

Another mission of Eurolab4HPC is to bring together researchers, industry and users for the development and the use of open source hardware and software creating a community to work on long-term projects. We believe that open source projects provide a natural forum for testing ideas and discovering innovation opportunities thereby promoting promote entrepreneurship, fostering industry take-up, and providing an avenue to train more experts in Exascale hardware and software.

LONG-TERM VISION ON HPC

Radical changes in computing are foreseen for the near future. The next Eurolab4HPC Long-Term Vision on High-Performance Computing, planned release in January 2020, will present an assessment of potential changes for High-Performance Computing in the next decade. You will be able to download it here: https://www.eurolab-4hpc.eu/vision/.

https://eurolab4hpc.eu
: @eurolab4hpc

COORDINATING ORGANISATION

Chalmers University of Technology, *Sweden*

OTHER PARTNERS

- Barcelona Supercomputing Centre (BSC), *Spain*
- Ecole Polytechnique Fédérale de Lausanne, *Switzerland*
- ETHZ - Eidgenössische Technische Hochschule Zürich, *Switzerland*
- FORTH Institute for Computer Science, *Greece*
- INRIA, Institut National de Recherche en Informatique et Automatique, *France*
- RWTH-AACHEN Rheinisch-Westfälische Technische Hochschule Aachen (RWTH-AACHEN), *Germany*
- TECHNION - Israel Institute of Technlogy, *Israel*
- The University of Edinburgh, *United Kingdom*
- The University of Manchester, *United Kingdom*
- Universität Augsburg, *Germany*
- Universität Stuttgart, *Germany*
- Universiteit Gent, *Belgium*

CONTACT

info@eurolab4hpc.eu

Per Stenström
per.stenstrom@chalmers.se

CALL

FETHPC-03-2017

PROJECT TIMESPAN

01/05/2018 - 30/04/2020

EXDCI-2

European eXtreme Data and Computing Initiative - 2

The project EXDCI-2 builds upon the success of EXDCI and will continue the coordination of the HPC ecosystem with important enhancements to better address the convergence of big data, cloud and HPC.

The strategic objectives of EXDCI-2 are:
a) Development and advocacy of a competitive European HPC Exascale Strategy
b) Coordination of the stakeholder community for European HPC at the Exascale.

EXDCI-2 mobilizes the European HPC stakeholders through the joint action of PRACE and ETP4HPC. It will promote global community structuring and synchronization in HPC, Big Data, Cloud and embedded computing, for a more competitive related value chain in Europe. It will develop a HPC technology roadmap addressing the convergence with HPDA and the emergence of new HPC uses. It will deliver application and applied mathematics roadmaps that will pave the road towards exascale simulation in academic and industrial domains. It will develop a shared vision for the future of HPC that increases the synergies and prepares for targeted research collaborations. EXDCI-2 will work to increase the impact of the H2020 HPC research projects, by identifying synergies and supporting market acceptance of the results.

At the international level, EXDCI-2 will contribute to the international visibility of Europe, develop contacts with the world leading HPC ecosystems and increase European impact on HPC standards.

EXDCI-2 will improve the HPC awareness by developing international event such as the EuroHPC Summit Week and by targeting specific audience through dedicated media by disseminating the achievements of the European HPC ecosystem.

http://exdci.eu
: @exdci_eu

COORDINATING ORGANISATION

Partnership for Advanced Computing in Europe (PRACE) AISBL, *Belgium*

OTHER PARTNERS

European Technology Platform for High Performance Computing (ETP4HPC), *Netherlands*

CONTACT

Marjolein Oorsprong
M.Oorsprong@staff.prace-ri.eu
Renata Gimenez
renata.gimenez@bsc.es

CALL

FETHPC-03-2017

PROJECT TIMESPAN

01/03/2018 - 31/08/2020

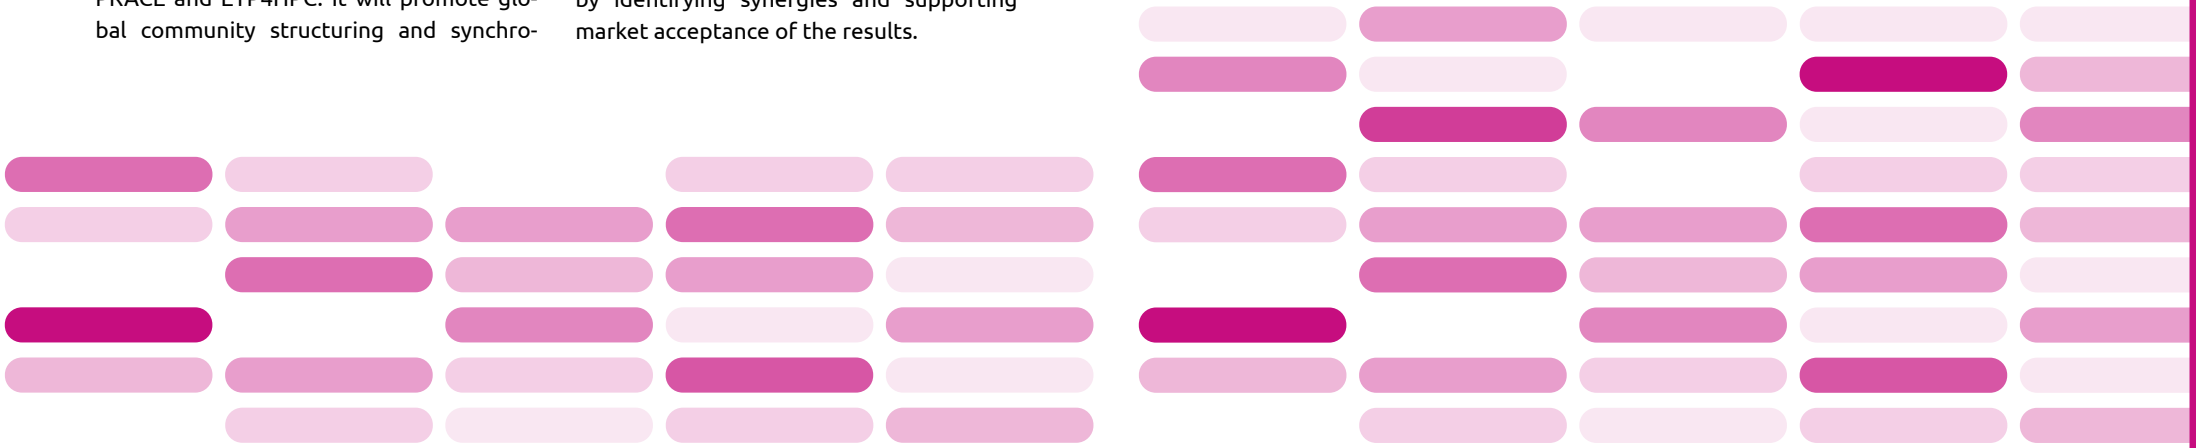

HPC-EUROPA3

Pan-European Research Infrastructure on High Performance Computing

Infrastructure on High Performance Computing
Pan-European Research
HPC Europa3

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 facilities, 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 www.hpc-europa.eu/public_documents.

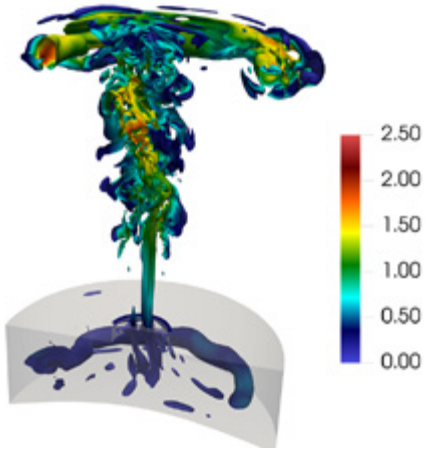

Flow of a circular synthetic jet represented using the Q-criterion and coloured by the magnitude of the velocity (from research carried out by HPC-Europa3 visitor Arnau Miró, Universitat Politècnica de Catalunya)

www.hpc-europa.eu
: @hpceuropa3
: @hpceuropa
: @hpc-europa3-eu-project
: www.youtube.com/channel/UC9uOpFQGP9V0TQPXFUOgs_A

COORDINATING ORGANISATION

CINECA Consorzio Interuniversitario, *Italy*

OTHER PARTNERS

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

CONTACT

staff@hpc-europa.org

CALL

INFRAIA-01-2016-2017

PROJECT TIMESPAN

01/05/2017 - 30/04/2021

Appendix - Completed projects

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

ALLScale	An Exascale Programming, Multi-objective Optimisation and Resilience Management Environment Based on Nested Recursive Parallelism
ANTAREX	AutoTuning and Adaptivity appRoach for Energy efficient eXascale HPC systems
BioExcel	Centre of Excellence for Biomolecular Research
COEGSS	Center of Excellence for Global Systems Science
ComPat	Computing Patterns for High Performance Multiscale Computing
CompBioMed	A Centre of Excellence in Computational Biomedicine
EoCoE	Energy oriented Centre of Excellence for computer applications
ESCAPE	Energy-efficient SCalable Algorithms for weather Prediction at Exascale
ExaFLOW	Enabling Exascale Fluid Dynamics Simulations
ExCAPE	Exascale Compound Activity Prediction Engine
EXTRA	Exploiting eXascale Technology with Reconfigurable Architectures
greenFLASH	Green Flash, energy efficient high performance computing for real-time science
INTERTWINE	Programming Model INTERoperability ToWards Exascale (INTERTWinE)
MaX	Materials design at the eXascale
Mont-Blanc 3	Mont-Blanc 3, European scalable and power efficient HPC platform based on low-power embedded technology
NoMaD	The Novel Materials Discovery Laboratory
POP	Performance Optimisation and Productivity
READEX	Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing
SAGE	Percipient StorAGe for Exascale Data Centric Computing

Imprint:
© Etp4Hpc

Text:
Etp4Hpc

Graphic design and layout:
www.maiffret.net

Paper:
Genyous digital

Printed and bound in november 2019:
Burlet Graphics, France

Contact
contact@etp4hpc.eu

The EXDCI-2 project has received funding from the European Union's Horizon 2020 research and innovation programme under the grant agreement No 800957.