



European High-Performance Computing

Handbook **2018**

Dear Partners,

This is the third edition of our European HPC Handbook. Our HPC landscape is growing (and so is this publication). As new projects and initiatives enter the stage, Europe is establishing itself as home to a world-class HPC ecosystem. We are introducing a new wave of HPC technology projects and application centres of excellence, both due to start in Q4 2018 or Q1 2019. These new projects continue the work initiated in 2015 by the first round of HPC projects. The two co-design projects, DEEP-EST and EuroExa, are gaining speed and have some interesting results to demonstrate. Another flagship project, the European Processor Initiative is being set up. Finally, we have included two projects which, though not part of the FETHPC calls, are focused on HPC and exascale. The future mechanism to finalise the HPC programme will be EuroHPC. In this initiative, the European Member States and the European Commission commit to fund the last, most complicated part of the effort: building European world-class HPC systems. EuroHPC will also continue funding HPC technology and application projects, with an emphasis on synergies with Big Data and AI. We hope this publication will help you understand our complex landscape. Our goal is also to help build international collaborations – the contact details of each project are included herewith.

EXDCI-2 Coordinator and PRACE Managing Director
Serge Bogaerts

ETP4HPC Chairman
Jean-Pierre Panziera

Table of Contents

6 European Exascale Projects

9 Basic Technology Projects started 2015

- 10 AllScale
- 12 ANTAREX
- 14 ComPat
- 16 ECOSCALE
- 18 ESCAPE
- 20 ExaFLOW
- 22 ExaHyPE
- 24 ExaNeSt
- 26 ExaNoDe
- 28 ExCAPE
- 30 EXTRA
- 32 Green Flash
- 34 INTERTWinE
- 36 MANGO
- 38 Mont-Blanc
- 40 NEXTGenIO
- 42 NLAFET
- 44 READEX
- 46 SAGE

49 Basic Technology Projects starting 2018/2019

- 50 ASPIDE
- 52 EPEEC
- 54 EPIGRAM-HS
- 56 ESCAPE-2
- 58 EXA2PRO
- 60 ExaQute
- 62 MAESTRO
- 64 RECIPE
- 66 Sage2
- 68 VECMA
- 70 VESTEC

73 Other HPC projects

- 74 EPI
- 76 Mont-Blanc 2020
- 78 PROCESS

81 Co-design projects

- 82 DEEP projects
- 84 EuroExa

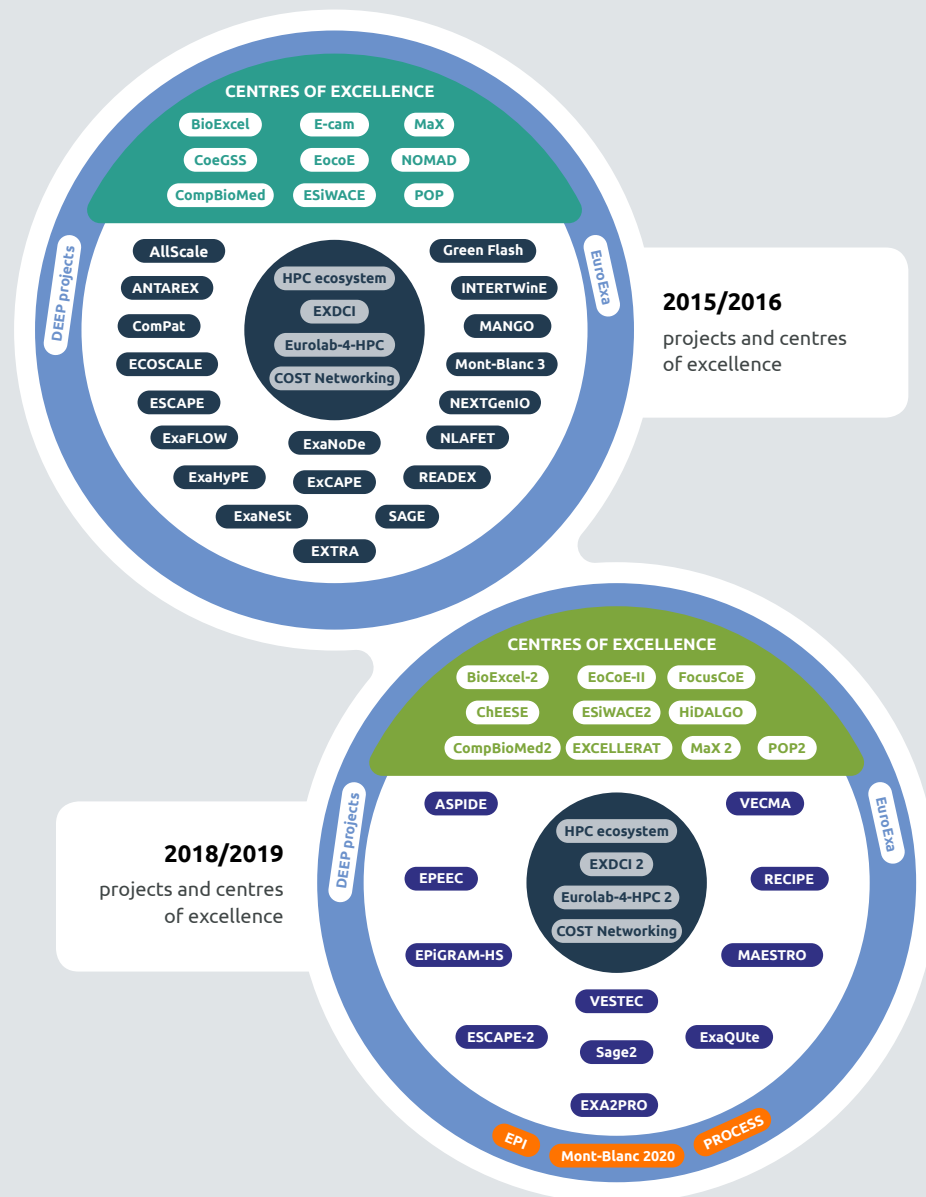
87 Centres of Excellence in computing applications started 2015

- 88 BioExcel
- 90 CoeGSS
- 92 CompBioMed
- 94 E-cam
- 96 EoCoE
- 98 ESIWACE
- 100 MaX
- 102 NOMAD
- 104 POP


107 Centres of Excellence in computing applications starting 2018

- 108 BioExcel-2
- 110 ChEESE
- 112 CompBioMed2
- 114 EoCoE-II
- 116 ESIWACE2
- 118 EXCELLERAT
- 120 FocusCoE
- 122 HiDALGO
- 124 MaX 2
- 126 POP2

EUROPEAN EXASCALE PROJECTS



SUMMARY OF EUROPEAN HPC EFFORT



- Basic Technology 2015**
 19 HPC technology projects starting 2015 with a duration of around 3 years
- Applications Excellence 2016**
 9 Centres of Excellence for Computing Applications - starting in 2016 with a duration of around 3 years
- Co-Design 2017**
 2 Co-Design projects (DEEP-EST and EuroExa) - starting in 2017 with a duration of around 3 years
- Basic Technology 2018**
 11 HPC technology projects starting 2018 with a duration of around 3 years
- Applications Excellence 2018**
 10 Centres of Excellence for Computing Applications to sign their project agreements in Q4 2018
- European Processor 2018**
 The European Processor Initiative - to start operating in Q4 2018 with a duration of around 4 years
- EuroHPC**
 A complex initiative of the EC and Member States with an objective to deliver European Exa-scale machines to start in Q1 2019 with a duration of 7 years.

Basic Technology Projects started 2015



● OBJECTIVES

AllScale is an innovative programming approach for ExaScale, decoupling program specification of parallelism by the programmer from management tasks automatized by the underlying runtime system. The programmer is exposed to a single, generic parallel model that provides a global view of parallelism and locality which is automatically mapped by exploiting recursive parallelism to a variety of parallel resources with the help of the AllScale toolchain that supports:

- Automated applications porting from small- to extreme scale architectures
- Flexible tuning of program execution to fulfil trade-offs among execution time, energy and resource usage
- Efficient management of hardware resources and associated parameters (e.g. clock speed)
- Seamless integration of resilience management measures to compensate for isolated hardware failures,
- Online performance monitoring and analysis.

The key-enablers of this approach are the utilization of nested recursive parallel pri-

AllScale enables developers to be productive and to port applications to any scale of system

mitives, the empowerment of the runtime system to actively manage the distribution of work and data throughout the system, and the utilization of advanced compiler analysis to aid the runtime system in steering program executions.

AllScale is expected to boost the parallel applications development productivity, their performance portability, and runtime efficiency. It will reduce energy consumption, thus improving the resource efficiency utilization of small to extreme scale parallel systems. The achieved outcomes are validated with applications from fluid dynamics, environment hazard and space weather simulations provided from SME, industry and academic institutions.

● ALLSCALE API BASED ON C++ TEMPLATES

User-level API:

- High-level abstractions (e.g. grids, meshes, stencils, channels)
- Familiar interfaces (e.g. parallel for loops, map-reduce)

Core-level API:

- Generic function template for recursive parallelism

www.allscale.eu

📧: @AllScaleEurope

COORDINATING ORGANISATION

Universitaet Innsbruck

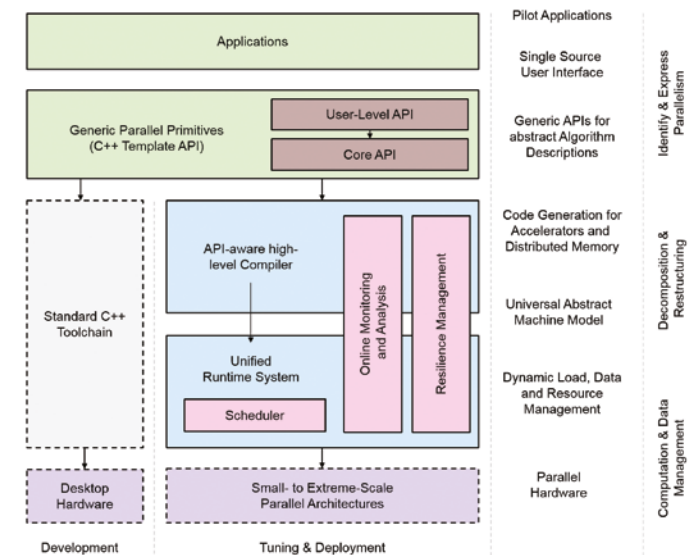
OTHER PARTNERS

Queen's University Belfast (QUB)
IBM Research (IBM)
NUMECA International (NUM)
Royal Institute of Technology (KTH)
Friedrich-Alexander-Universität (FAU)

CONTACT

Thomas Fahringer
tf@dps.uibk.ac.at
+ 43 512 507 53204

- Set of recursive data structure templates
- Synchronization, control- and data-flow primitives



AllScale architecture



● PROJECT HIGHLIGHTS

Energy-efficient heterogeneous supercomputing architectures need to be coupled with a radically new software stack capable of exploiting the benefits offered by the heterogeneity at all the different levels (supercomputer, job, node) to meet the scalability and energy efficiency required by Exascale supercomputers.

ANTAREX will solve these challenging problems by providing a breakthrough approach to express application self-adaptivity at design-time and to runtime manage and autotune applications for green and heterogeneous High Performance Computing (HPC) systems up to the Exascale level.

● WHAT ARE ANTICIPATED TECHNOLOGY (HW/SW/METHODOLOGY) SUGGESTED FOR INCLUSION IN AN ESD PROJECT AND DESCRIBE THE CURRENT MATURITY?

The compiler technology being developed consists of a separation of concerns (where self-adaptivity and energy efficient strategies are specified aside to application functionalities) promoted using LARA, a Domain Specific Language (DSL) inspired

by aspect-oriented programming (AOP) concepts for heterogeneous systems. The project includes the development of standalone libraries, methodologies and tools focused on code analysis and optimization, runtime application autotuning, and resource and power management.

We have different maturity levels among the tools and libraries. The technology behind the DSL has several years of development, the libraries and tools have different levels of maturity (some have started in this project, others started in previous projects and have been extended), and the integration of all the tools and libraries in a single framework is still in development. The framework is based on the following technologies:

- The mARGOt autotuner (autotuning framework which enhances the application with an adaptation layer):
- The Examon framework (A highly scalable framework for performance and energy monitoring of supercomputing machine in production):
- The LibVersioningCompiler (easy dynamic compilation with versioning support)
- The PowerManager (Power capper that

selects the best performance point for each core in order to maintain a power constraint while adapting at the workload parallelization strategy):

- The CLAVA + LARA compiler framework (C/C++ source-to-source tool for code instrumentation and transformations controlled by the LARA language)

● HOW SHOULD THIS TECHNOLOGY BE USED / INTEGRATED

The DSL is being used to control the integration between the several libraries/tools and target applications. Other third-party tools and libraries can be used to output information (e.g., resultant from analysis and profiling) to LARA strategies. These LARA strategies can then consider the input information to decide about some code transformations, instrumentation, and autotuning knobs and calls. In addition, with the ANTAREX approach other instrumentation and monitoring libraries and/or autotuning schemes can be holistically integrated by the use of LARA strategies to transform and inject extra code in the application source code.

● ARE THERE ANY PRE- OR CO-REQUIREMENT ITEMS

The approach aims at avoiding manual modifications on source code, so requisites would mostly be writing strategies in the LARA DSL or adapting existing ones. Besides that, the DSL framework (CLAVA) needs a Java runtime, and each tool/library has its own specific dependences.

www.antarex-project.eu

📧: @antarex_project

COORDINATING ORGANISATION

Politecnico di Milano, *Italy*

OTHER PARTNERS

- CINECA, *Italy*
- Dompé, *Italy*
- ETH Zürich, *Switzerland*
- INRIA, Institut National de Recherche en Informatique et Automatique, *France*
- IT4Innovations, *Czech Republic*
- Sygic, *Slovaquie*
- University of Porto, *Portugal*

CONTACT

Cristina Silvano
cristina.silvano@polimi.it

● ANY EXTRA WORK/INTERACTION NEEDED TO MAKE THEM READY?

Some of the technologies in development are already deployed in HPC systems in production, however to have a production-ready framework that integrates all the tools, we consider that it will need extra work beyond the project.

● WHAT INFORMATION / ACTIONS ARE NEEDED TO BEST PREPARE FOR ESD PROJECTS?

Ensure that each individual component has a clear interface and is well documented, start the integration to have a proof-of-concept and assess the current state, to prepare a plan of what would be necessary to have a fully-fledged product ready for HPC systems.



Multiscale phenomena are ubiquitous, and they are the key to understanding the complexity of our world. **Our Main objective** is to develop generic and reusable High Performance Multiscale Computing algorithms to address the exascale challenges of heterogeneous architectures and be able to run multiscale applications with extreme data requirements while achieving scalability, robustness, resiliency, and energy efficiency.

Nine grand challenge applications from four different science domains demonstrate the viability of our approach. **Nuclear fusion:** the interactions between turbulence at very small scales and the large-scale plasma behaviour holds the key to control its magnetic confinement in order to produce clean and carbon free energy for the indefinite future. **Astrophysics:** the formation processes of stars in their clustered environment, as well as the origin and propagation of structure in the stellar disk of the Milky Way Galaxy. **Materials science:** prediction of the materials properties of macroscopic samples of matter based on the specification of the atoms and molecules comprising it. **Biomedicine:** pathophysiology of vascular disease, as well as

to provide personalised models of the vasculature in near to or real time to support clinical decision-making.

The approach is to identify three multiscale computing patterns, realise generic algorithms for all three, implement the selected grand challenge applications as tailor-made instantiations of the computing patterns, which will serve to demonstrate the efficacy of the new algorithms. **The insight** we gain en-route will be used to develop performance prediction models to anticipate the efficiency of the applications on future exascale machines.

The Multiscale Modelling and Simulation Framework (MMSF) is a theoretical and practical way to model, characterise and simulate multiscale phenomena. MMSF currently comprises a 4-stage pipeline, going from developing a multiscale model to executing a multiscale simulation.

Our approach is based on **Generic Multiscale Computing Patterns** that allow us to implement customized algorithms to optimise under generic exascale application scenarios. The extreme Scaling computing pattern where one (or a few) of the single

www.compat-project.eu

Twitter: [compatproject](https://twitter.com/compatproject)

COORDINATING ORGANISATION

University of Amsterdam

OTHER PARTNERS

- University of Amsterdam
- University College London
- Institute of Bioorganic Chemistry – Poznan Supercomputing and Networking Centre
- Bayerische Akademie der Wissenschaften – Leibniz-Rechenzentrum
- Max-Planck-Gesellschaft zur Förderung der Wissenschaften e.V
- University Leiden
- The Hartree Centre/STFC
- Allinea Software
- CBK Sci Con
- National Research University ITMO
- Brunel University

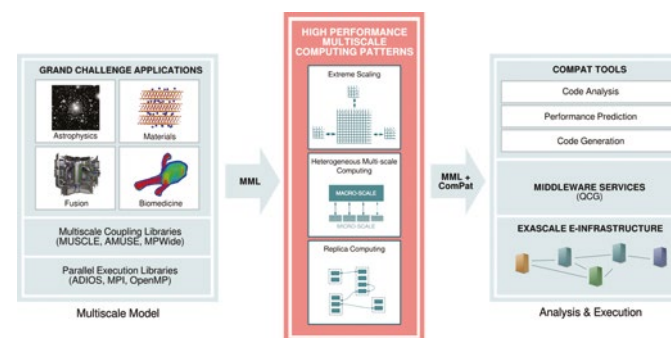
CONTACT

Hugh Martin
h.martin@cbkscicon.com
02076795300

scale models in the overall multiscale model dominates all others by far, in terms of required computing power. In the **Heterogeneous Multiscale Computing pattern**, we couple a macroscopic model to a large and dynamic number of microscopic models. **Replica Computing** is a multiscale computing pattern that combines a potentially very large number of terascale and petascale simulations (also known as ‘replicas’) to produce robust outcomes

The multiscale computing patterns determine the ordering and composition of models that are coupled within a multiscale application. To orchestrate the execution

of the application on HPC resources and specifically exascale systems, we require a **comprehensive technology stack**.





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 that perform coherent memory accesses in the virtual address space while being programmed by OpenCL.

The novel UNIOLOGIC (Unified Logic) architecture, introduced within ECOSCALE, is an extension of the UNIMEM architecture. UNIMEM provides shared partitioned global address space while UNIOLOGIC 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 UNIOLOGIC+UNIMEM architecture partitions the design into several Worker nodes that communicate through a fat-tree communication infrastructure.

www.ecoscale.eu

📧: @ECOSCALE_H2020

COORDINATING ORGANISATION

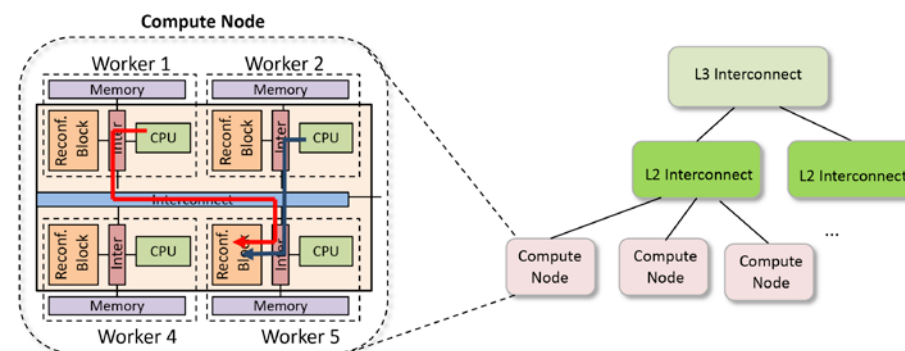
Telecommunication Systems Institute (TSI)

OTHER PARTNERS

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

CONTACT

Yannis Papaefstathiou
ygp@synelixis.com
+30 694427772



ECOSCALE architecture

ESCAPE



ESCAPE stands for Energy-efficient Scalable Algorithms for Weather Prediction at Exascale. The project will develop world-class, extreme-scale computing capabilities for European operational numerical weather prediction (NWP) and future climate models. The biggest challenge for state-of-the-art NWP arises from the need to simulate complex physical phenomena within tight production schedules. Existing extreme-scale application software of weather and climate services is ill-equipped to adapt to the rapidly evolving hardware. This is exacerbated by other drivers for hardware development, with processor arrangements not necessarily optimal for weather and climate simulations. ESCAPE will redress this imbalance through innovation actions that fundamentally reform Earth system modelling.

ESCAPE addresses the ETP4HPC Strategic Research Agenda 'Energy and resiliency' priority topic, developing a holistic understanding of energy-efficiency for extreme-scale applications using heterogeneous architectures, accelerators and special compute units. The three key reasons why this

Energy efficient Scalable Algorithms for weather Prediction at Exascale (ESCAPE)

project will provide the necessary means to take a huge step forward in weather and climate modelling as well as interdisciplinary research on energy-efficient high-performance computing are:

- Defining and encapsulating the fundamental algorithmic building blocks ('Weather & Climate Dwarfs') underlying weather and climate services. This is the prerequisite for any subsequent co-design, optimization, and adaptation efforts.
- Combining ground-breaking frontier research on algorithm development for use in extreme-scale, high-performance computing applications, minimizing time- and cost-to-solution.
- Synthesizing the complementary skills of all project partners. ECMWF and leading European regional forecasting consortia are teaming up with excellent university research and experienced high-performance computing centres, two world-leading hardware companies, and one European start-up SME, providing entirely new knowledge and technology to the field.

www.hpc-escape.eu

COORDINATING ORGANISATION

European Centre for Medium-Range Weather Forecasts (ECMWF), UK

OTHER PARTNERS

- Météo-France, *France*
- Royal Météorologique de Belgique, *Belgium*
- Danmarks Meteorologiske Institut, *Denmark*
- Federal Office of Meteorology and Climatology, *Switzerland*
- Deutscher Wetterdienst, *Germany*
- Loughborough University, *UK*
- Irish Centre for High-End Computing, *Ireland*
- Poznan Supercomputing and Networking Centre, *Poland*
- Bull SAS (Atos Group), *France*
- NVIDIA, *Switzerland*
- Optalysys, *UK*

CONTACT

Dr. Peter Bauer
peter.bauer@ecmwf.int
+44 118 949 9080



Computational Fluid Dynamics (CFD) is a prime contender for reaching exascale performance: In fluid dynamics there is virtually no limit to the size of the systems to be studied via numerical simulations, which can be exploited for extreme parallel scaling. Moreover, fluid flows are an essential element of many industrial and academic problems: A crude estimate shows that 10% of the energy use in the world is spent overcoming turbulent friction. As such, collaboration between software creators and European industries within automotive, manufacturing, aerospace, energy and

health care is crucial. The goals of ExaFLOW comprise four key innovation areas, including aspects of:

- Accurate error control, mesh adaptivity
- Strategies to ensure fault tolerance and resilience
- Strong scaling at exascale through novel CG-HDG discretisations
- I/O techniques for extreme scale data

In all these areas, significant progress has been made since the project start. For instance, we have developed fault tolerance mechanisms, such that the new algorithm



Simulation using Nektar++ using ARCHER on the RP1 Elemental Track Car

www.exaflow-project.eu

🐦: @ExaFLOWproject

COORDINATING ORGANISATION

KTH Royal Institute of Technology, Sweden

OTHER PARTNERS

- University of Stuttgart, USTUTT, Germany (*Institute of Aerodynamics and Gas Dynamics, IAG, High Performance Computing Centre Stuttgart, HLRS*)
- Ecole Polytechnique Fédérale de Lausanne, EPFL, Switzerland
- Imperial College London, IC, UK
- University of Southampton, SOTON, UK
- The University of Edinburgh, UEDIN, UK
- McLaren Racing Ltd, McLaren, UK
- Automotive Simulation Centre Stuttgart, ASCS, Germany

CONTACT

Erwin Laure
exaflow-contact@pdc.kth.se
+46 8-790 65 14

“survives” >90% of the errors that would otherwise have resulted in an execution failure; this all with very little overhead both in fault-free execution and in recovery. New I/O techniques and compression algorithms allowing very high compression ratios (>1:100) with very small error rates due to lossy compression. Scalable numerical methods, including efficient AMG based pre-conditioners, CG-HDG discretisations to reduce communication requirements and h- p-type mesh refinement. Finally, development of new algorithmic techniques for error control based on spectral error indicators or adjoint error estimators to steer mesh refinement in an adaptive simulation. In the final year of the project, ExaFLOW conducted three flagship runs with high academic and industrial relevance to assess the outcome of the project.



● EXAHYPE VISION

Hyperbolic systems of PDE resulting from conservation laws are the basis to simulate a wide range of problems, ranging from weather forecast, complex earthquake physics, hematic flows in patients up to the most catastrophic events in the universe. In this project, we develop an exascale-ready “engine” that shall enable teams of computational scientists to more quickly realize grand challenge simulations based on hyperbolic PDE models.

● EXAHYPE ALGORITHMS


The developed engine implements a high-order discontinuous Galerkin approach with ADER time stepping and a-posteriori finite-volume limiting, which promises to combine high-order accuracy and excellent robustness of solutions. Problems are discretised on tree-structured fully adaptive Cartesian meshes, for which advanced parallelisation approaches and load-balancing algorithms are developed in the project.

● EXAHYPE ENGINE

The hyperbolic PDE engine developed in the ExaHyPE project is available as open

source 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 simulation of neutron star mergers and the resulting gravitational waves.

www.exahype.eu

COORDINATING ORGANISATION


Technical University of Munich

OTHER PARTNERS


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

CONTACT


Michael Bader
bader@in.tum.de



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



Seismic wave propagation in a complex alpine topography.




● PROJECT HIGHLIGHTS

- ARMv8 with *UNIMEM* (global sharable coherent memory scheme)
 - Supporting MPI and Partitioned Global Address Space (PGAS)
 - low energy compute, low overhead communications
 - UNIOLOGIC FPGA acceleration
- Networking – unified compute & storage, low latency
- Storage – converged offering locality while distributed, non-volatile memories
- Real Applications – Scientific, Engineering, Data Analytics
- Data Centre Infrastructure – Power, Cooling, Mechanicals, Total Liquid Cooling

● PROTOTYPE OUTCOMES

- Networking, Storage and Application Developments – Ready to support the Nest, Nodes and the Ecosystem,
 - The Nest – Ready to power, connect and liquid cool different distributed compute nodes (@3.2kW/u, 762mmx600mm cabinets)
 - FPGA/ARM Nodes – Initial Nodes for the Nest with 4 FPGAs
- Commercialization outlook, in cooperation with external partners, including KALEAO



120mm x 130mm Quad FPGA/ARM Node with NVRAM and DRAM

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 Institute for Nuclear Physics (INFN)
- The University of Manchester
- Universitat Politècnica de València
- Fraunhofer

CONTACTS

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

Peter Hopton
(Dissemination & Exploitation Lead)
peter.hopton@iceotope.com
+447437012186



PROJECT HIGHLIGHTS

ExaNoDe designs core technologies for a highly energy efficient and highly integrated heterogeneous compute node directed towards Exascale computing. ExaNoDe's goal is to deliver an HW/SW integrated 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 μ m pitch) connections.
- Advanced package integration with two FPGA bare dies including ARMv8 cores, one interposer and 43 decoupling capacitors in a 68.5 mm \times 55 mm Multi-chip-Module (MCM).
- Integration of two MCMs on a 260 mm \times 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.

At the beginning of September 2018, the ExaNoDe prototype, scheduled for end of March 2019, entered its final integration phase. 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 closely collaborates with the ExaNeSt and EcoScale projects on the basis of a joint MoU.



Multi-Chip-Module

Exanode prototype

www.exanode.eu

📧: @ExanodeProject

COORDINATING ORGANISATION

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


OTHER PARTNERS

- Arm Limited, UK
- ETH 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, 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



Die stacking with active interposer

ExCAPE



The ExCAPE project is scaling up machine learning techniques to solve challenging problems on HPC machines. Our driving application is compound-activity prediction for drug discovery. We are preparing data, developing the state of the art for machine learning algorithms, and researching programming models to implement them.

Our main achievements so far include:

- Public release of a data set that resembles industry data in terms of size and distribution of hits and misses, and experiments showing the potential of novel compound descriptors for doing multi-target predictions at scale
- Open source release of HyperLoom, a programming model and task execution system that runs on HPC systems and is designed to cope with machine learning
- Open source release of SMURFF, a matrix factorization package allowing sophisticated combinations of techniques such as GFA and Macau

- Exploration of the use of and hardware implications of sparse matrix techniques to deal with large sparse feature vectors, and the impact on deep learning
- Benchmarking of ML techniques on large scale data sets using HyperLoom
- Novel algorithms to improve scalability of matrix factorisation on large machines
- Demonstrating the computational requirements of conformal prediction
- Novel clustering implementations for pre-processing compound data

● PROJECT AREAS OF INTERNATIONAL COLLABORATION

- Experts in scheduling machine learning tasks on HPC systems
- Pharma companies to compare scalability of learning approaches
- Machine learning experts to explore scalability of different classes of algorithms
- Other users and developers of large scale multi-target learning

www.excape-h2020.eu

COORDINATING ORGANISATION


Imec

OTHER PARTNERS

- IT4I Czech National Supercomputing Centre
- AstraZeneca
- Janssen Pharmaceutica
- IDEAConsult
- Intel
- JKU Linz
- University of Aalto
- Royal Holloway UL

CONTACT

Tom Ashby, ashby@imec.be
Tom Vander Aa, vanderaa@imec.be
+32-16-288500





Our objective is to create a flexible exploration platform for the development and optimization of architectures, tools, and applications targeting Exascale systems using reconfigurable technology. We aim to provide an open-source research platform to enable sustainable research on reconfigurable architectures and tools, and prepare all stakeholders (application, systems, and tool developers) for the emergence of efficient reconfigurable accelerators for Exascale systems. Our work revolves around 5 activities:

● UNDERSTAND THE REQUIREMENTS

We use three different case-studies - Asian option pricing, retinal image segmentation, and Diffusion Monte-Carlo - to investigate the performance, scalability, and energy efficiency challenges that emerge in HPC-like applications at Exascale level. We aim to provide mechanisms to cope with these challenges in both hardware and software.

● DESIGN, DEVELOP, AND DEPLOY THE OPEN RESEARCH PLATFORM

We focus on the design and implementation

of the Open Research Platform, covering its reconfigurability-specific features and design tools. This work includes, for example, research on Virtual Coarse Grained Reconfigurable Arrays (VCGRAs), static and dynamic workload partitioning, efficient data offloading, and parallel memory systems. Furthermore, the CAOS toolchain is developed to support the design and deployment of reconfigurable accelerators, and will integrate the main analysis and optimization tools developed in EXTRA, with support for community extensions in the future.

● RESEARCH NOVEL TOOLS.

We strive to advance the state-of-art compilation tools for reconfigurable hardware, covering just-in-time synthesis, hardware monitoring, emergency management, and hardware debugging.

● INVESTIGATE RECONFIGURABLE APPLICATIONS

We search for generalizable optimization opportunities in HPC applications, to leverage reconfiguration, as well as implement, integrate and evaluate their impact on HPC systems.

www.extrahpc.eu

📧: @extrahpc

COORDINATING ORGANISATION

Ghent University, Belgium

OTHER PARTNERS

- Imperial College London, *UK*
- University of Cambridge, *UK*
- University of Amsterdam, *Netherlands*
- Politecnico di Milano, *Italy*
- Ruhr University Bochum, *Germany*
- Telecommunication Systems Institute, *Greece*
- Synelxis, *Greece*
- Maxeler, *UK*

CONTACTS

Prof. Dirk Stroobandt, U. Ghent
(project coordinator)
Dirk.Stroobandt@UGent.be
+32 (0)9 264 34 01

Dr. ir. Ana Lucia Varbanescu, UvA
A.L.Varbanescu@uva.nl
+31 (0)20 525 8683

● PROPOSE NEW IDEAS FOR RECONFIGURABLE TECHNOLOGY

Based on the combined analysis of applications, programming technologies and tools, and hardware platforms, we aim to propose technological improvements that will increase the efficiency of future reconfigurable systems. The main aspects of this activity include the improvement of the reconfiguration process and finding efficient interface solutions for coupling the FPGA fabric and the processing units.

● PLATFORM OVERVIEW

The Open Research Platform has six main components to support the optimization of applications for the next-generation reconfigurable HPC systems. For this purpose, the platform provides an integrated system that combines models for various reconfigurable architectures, tools and applications, allowing researchers to focus on any of these aspects separately or together. The *platform input* includes information about the application (e.g., source-code, performance requirements, profiling information) and the specification of the target platform (e.g., size, resources, simulated/

hardware). The *Hardware Platform Model* guides the design and optimization process for a specific target platform. The Tools Platform (e.g., CAOS) generates a reconfigurable design through a set of analysis and transformations steps on the application code. The *Performance Modeling* component checks the fit of the expected performance with the user requirements at every stage of the development process. One or *more candidate reconfigurable* design versions is generated to support design-space exploration and select the most suitable implementation to execute on the *Reconfigurable Platform*, which in turn combines the hardware architecture and runtime of the reconfigurable HPC system.

Green Flash



The main goal of Green Flash is to design and build a prototype for a Real-Time Controller targeting the European Extremely Large Telescope (E-ELT) adaptive optics instrumentation.

On one hand, we propose to demonstrate the scalability of accelerator-based solutions for real-time HPC data-intensive applications. On the other hand, an alternative to the commodity accelerator solution for real-time control is to build tailored processor boards, based on high cell density FPGAs. The emergence of new FPGA products, integrating an ARM-based Hard Processor

System (HPS) with the FPGA fabric, could enable a new approach, merging the concept of compute node in a HPC cluster and specialized processor board in an application specific facility. Our strategy is based on a strong interaction between academic and industrial partners. A prototype harbouring all the features is used to assess the performance. It also provides the proof of concept for a resilient modular solution to equip a large scale European scientific facility, while containing the development cost by providing opportunities for return on investment.

www.greenflash-h2020.eu

COORDINATING ORGANISATION

Observatoire de Paris

OTHER PARTNERS

- University of Durham
- Microgate
- PLDA

CONTACTS

Damien Gratadour (OdP),
damien.gratadour@obspm.fr

James Osborn (UoD),
james.osborn@durham.ac.uk


Roberto Biasi (Mic),
roberto.biasi@microgate.it

Hugues Deneux (PLDA),
hdeneux@plda.com

Current trend in heterogeneous HPC



Accelerators, ARM64 & smart interconnect



FPGA-based microserver





The first Exascale computers will be very highly parallel systems, consisting of a hierarchy of architectural levels. To program such systems effectively and portably, application programming interface (APIs) with efficient and robust implementations must be ready in the appropriate timescale. A single, “silver bullet” API which addresses all the architectural levels does not exist and seems very unlikely to emerge soon enough. We must therefore expect that using combinations of different APIs at different system levels will be the only practical solution in the short to medium term. The INTERTWinE project focused on the challenges that lie in interoperability between APIs, both at the specification level and at the implementation level.

● PROJECT DESCRIPTION

INTERTWinE brought together the principal European organisations driving the evolution of programming models and their implementations. We focused on six key programming APIs: MPI, GASPI, OpenMP, OmpSs, StarPU, and PaRSEC, each of which had a project partner with extensive experience in API design and implementation.

INTERTWinE partners have been very active in pursuing interoperability issues and new API developments in three international standards bodies: the MPI Forum, the GASPI Forum, and the OpenMP ARB. Project partners have participated in the relevant working groups to propose and test modifications and extensions to these important APIs.

Interoperability requirements, and evaluation of implementations, were driven by porting a set of kernels and applications to various API combinations. These codes have been released as a series of Resource Packs on the Developer Hub section of the project webpage.

Driven by these application requirements, INTERTWinE has designed and implemented two new runtime APIs: a Resource Manager to allow multiple runtimes to negotiate resource sharing within a node, and a Directory/Cache to support execution of task-based programming models on distributed systems by abstracting the communication layer.

www.intertwine-project.eu

📧: @intertwine_eu

COORDINATING ORGANISATION

EPCC, University of Edinburgh,
United Kingdom

OTHER PARTNERS

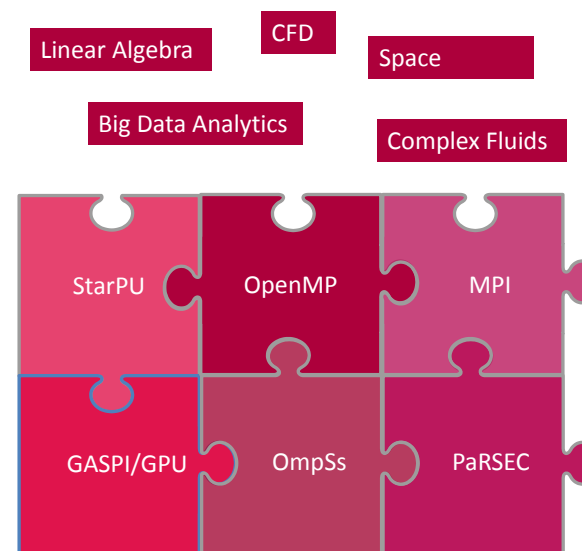
- BSC – Barcelona Supercomputing Centre, Spain
- KTH, Sweden
- Fraunhofer ITWM, Germany
- DLR, Germany
- T-Systems SFR, Germany
- Universitat Jaume I, Spain
- Inria, France
- University of Manchester, UK

CONTACT

Dr George Beckett, EPCC
info@intertwine-project.eu
+44 (0)131 651 3577

The project has also provided a number of resources to help developers combine multiple APIs in their applications, by running a programme of training courses, producing a series of Best Practice Guides, and creating a dedicated Developer Hub section of the webpage.

INTERTWinE also had collaborations with runtime development teams in the US, with close connections to the PaRSEC team at the University of Tennessee, and also to the Centre for Computing Research at Sandia National Laboratories, working on the OpenMPI library



INTERTWinE's approach




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 thermoshypons.

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 Politècnica 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
jflinch@disca.upv.es
+34 96 387 70 07 (ext. 75753)

Mont-Blanc



Project Mont-Blanc is now in its third phase. All phases of the Mont-Blanc project share the vision of developing a European Exascale approach leveraging commodity power-and cost-efficient embedded technologies.

The key outcome of the project is the deployment of Arm-based computing platforms enabling Arm architecture in HPC, which boost system software development and allow to test real scientific applications at scale.

Based on the experience gained from the development of various platforms since 2011 and implementing a co-design approach, the Mont-Blanc project now aims to define the architecture of an Exascale-class compute node based on the Arm architecture, and capable of being manufactured at industrial scale.

● OUR TOP ACHIEVEMENTS

1. Demonstrating that it is possible to run HPC workloads with European embedded technology
2. Contributing to the design of a next-generation exascale-class machine with a co-design approach
3. Testing and scaling REAL scientific applications on a non-conventional HPC architecture



Mont-Blanc 3 prototype blade, equipped with Cavium ThunderX2 ArmV8 processors

The Mont-Blanc team at SC17 with our HPCwire Awards:

1. Best Collaborative Project for Mont-Blanc and
2. TOP 5 new products for the blade derived from our prototype and commercialized by Atos in their BullSequana X1000 range).

www.montblanc-project.eu

📧: @MontBlanc_EU

COORDINATING ORGANISATION

Bull SAS (Atos Group),s

OTHER PARTNERS

- Arm, *UK*
- AVL, *Austria*
- BSC, *Spain*
- CRNS (Centre National de la Recherche Scientifique) / LIRMM, *France*
- ETH Zurich, *Switzerland*
- HLRS, *Germany*
- Universidad de Cantabria (Santander), *Spain*
- University of Graz, *Austria*
- Université de Versailles Saint Quentin, *France*

CONTACT

Etienne Walter
Etienne.walter@atos.net
+33 1 30 80 74 96





The NEXTGenIO project addresses the **I/O performance** 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

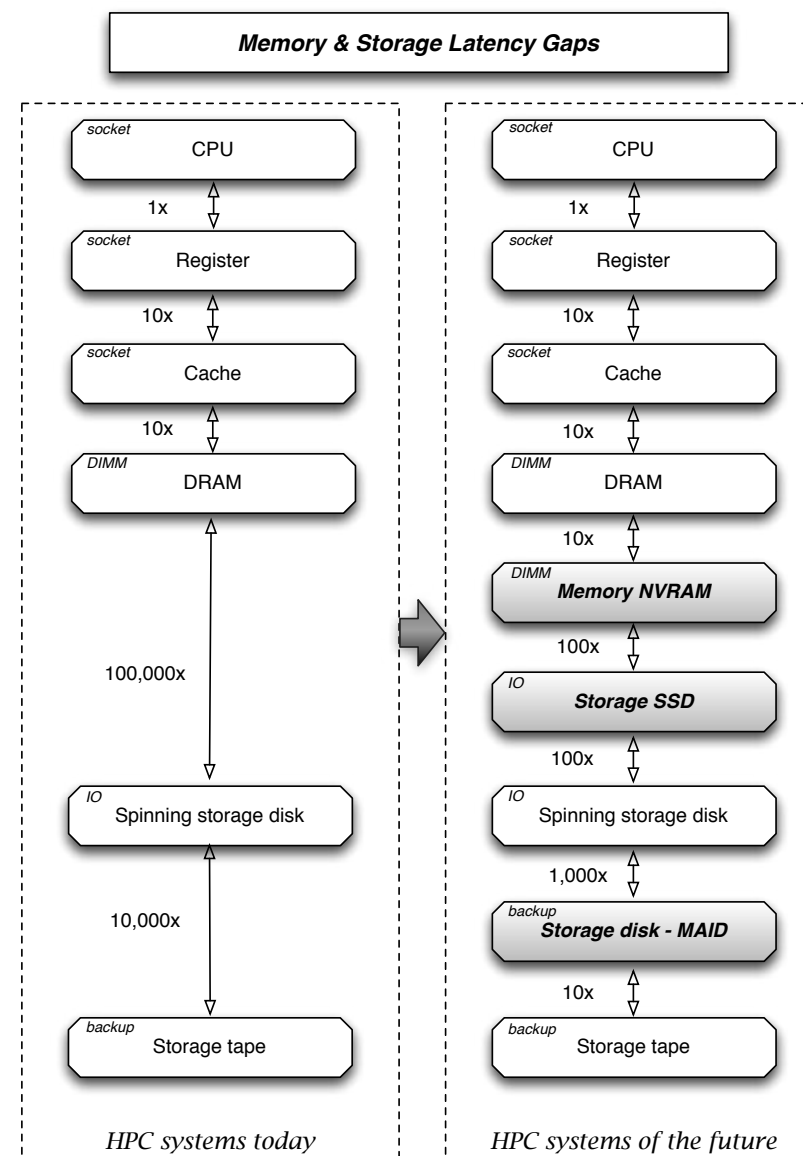
OTHER PARTNERS

- Intel Deutschland GmbH, Germany
- Fujitsu
- Barcelona Supercomputing Center
- 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



Memory and Storage Latency Gaps



● 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 parallel programming models used include MPI, OpenMP, StarPU, and PaRSEC. 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.

● NLAFET IMPACT

The main impact is to develop, deploy and make 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

www.nlafet.eu

COORDINATING ORGANISATION

Umeå University, Sweden

OTHER PARTNERS

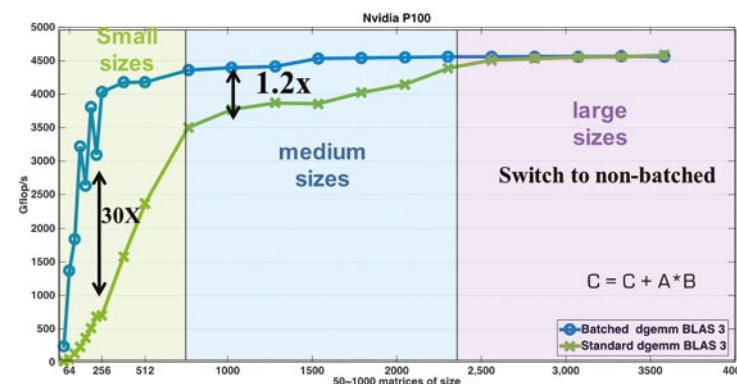

- University of Manchester, UK
- Institut National de Recherche en Informatique et en Automatique, Inria, France
- Science and Technology Facilities Council, UK

CONTACT

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

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



The importance of energy efficiency is constantly increasing in High Performance Computing (HPC). While systems can be adapted to individual applications in order to reduce energy consumption, manual tuning of platform parameters is a tedious and often neglected task.

The READEX project automates this by developing a tools-aided methodology for dynamic auto-tuning that combines technologies from two ends of the computing spectrum: system scenario methodology from the embedded world and auto-tuning from the field of HPC.

● READEX TOOLS-AIDED METHODOLOGY

The READEX methodology has been designed for exploiting the dynamic behaviour of software. At design time different runtime situations (RTS) are detected and optimized system configurations are determined. RTSs with the same configuration are grouped into scenarios, forming the tuning model. At runtime, the tuning model is used to switch system configurations dynamically.

Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing

● IMPACT AND VALIDATION

To validate the impact of the READEX project, several real-world applications are employed. In a co-design approach, selected applications are being hand-tuned and both the improvements in energy efficiency and the effort spent will be compared with the automatic tuning approach.

● DESIGN TIME ANALYSIS

Design time analysis (DTA) is carried out with the Periscope Tuning Framework (PTF). It uses a multi-agent-based approach to identify RTSs and to determine optimized system configurations. These are settings for tuning parameters, e.g., core and uncore frequencies. It also provides means for the specification of domain knowledge (DK) to improve the automatic tuning results. Part of the DK is the specification of application tuning parameters, which allows users to offload computation to accelerated devices. The result of DTA is a tuning model that guides runtime tuning.

www.readex.eu

📧: @readex_eu

COORDINATING ORGANISATION

Technische Universität Dresden
Center for Information Services
and High Performance Computing

OTHER PARTNERS

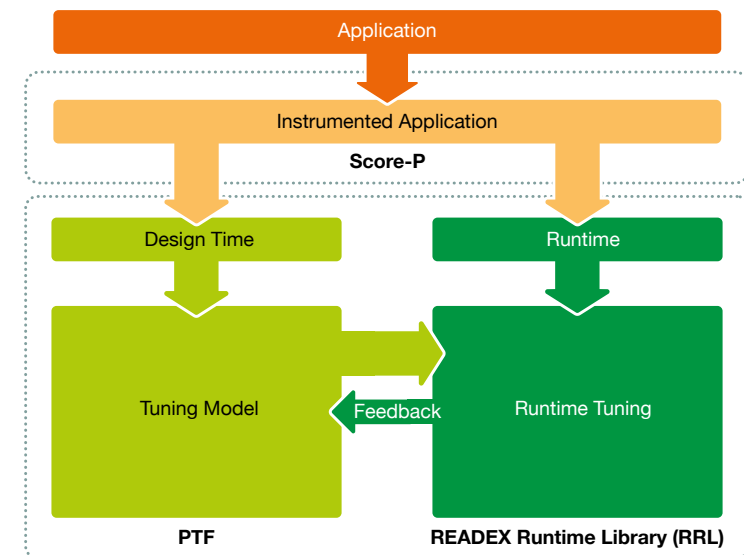
- Gesellschaft für Numerische Simulation mbH
- Intel Exascale Labs Paris
- IT4Innovations, VSB - Technical University of Ostrava
- National University of Ireland Galway
- Norwegian University of Science and Technology
- Technische Universität München

CONTACT

Prof. Dr. Wolfgang E. Nagel
wolfgang.nagel@tu-dresden.de
+49 (0) 351 463-35450

● RUNTIME-TUNING

During production runs of the user's application, the READEX Runtime Library takes control. It is designed to apply the different configuration in a lightweight manor. Moreover, the READEX Runtime Library will be able to adapt to changing application behaviour. The latter is to be implemented by state-of-the-art machine learning mechanisms, which are currently under development.



An overview of the READEX architecture.



The StorAGe for Exascale Data Centric Computing (SAGE) system, researched and built as part of the SAGE project, aims to implement a Big Data/ Extreme Computing

(BDEC) and High Performance Data Analytics (HPDA) capable infrastructure suitable for Extreme scales - Exascale and beyond. Increasingly, overlaps occur between Big

Data Analysis and High Performance Computing (HPC), caused by the proliferation of massive data sources, such as large, dispersed scientific instruments, sensors, and social media data, whose data needs to be processed, analysed and integrated into computational simulations to derive scientific and innovative insights. The SAGE storage system, will be capable of efficiently storing and retrieving immense volumes of data at Extreme scales, with the added functionality of "Percipience" or the ability to accept and perform user defined computations integral to the storage system. The SAGE system is built around the Mero object storage software platform, and its API Clovis, and its supporting ecosystem of tools and techniques, that will work together to provide the required functionalities and scaling desired by Extreme scale workflows. The SAGE system will seamlessly integrate a new generation of storage device technologies, including non-volatile memories as they become available. The SAGE system will also offer a very flexible API and a powerful software framework suitable for easy extensibility by third parties.

www.sagestorage.eu

📧: @SageStorage

COORDINATING ORGANISATION

Seagate

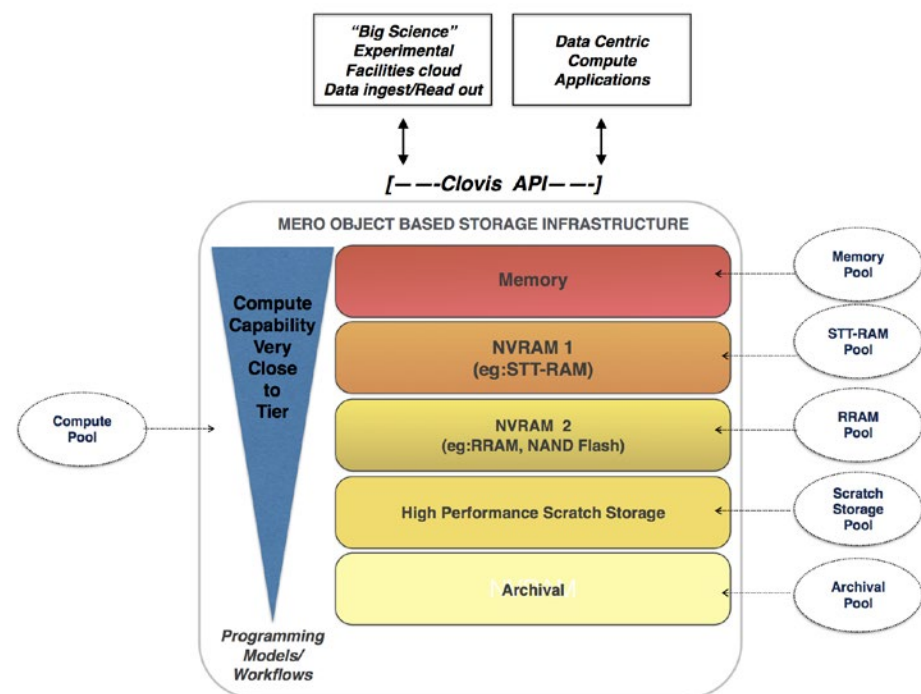
OTHER PARTNERS

- Seagate Systems, *UK*
- Allinea Software [ARM], *UK*
- Bull SAS (Atos Group), *France*
- United Kingdom Atomic Energy Authority, *UK*
- Commissariat à l'énergie atomique et aux énergies alternatives (CEA), *France*
- Deutsche Forschungszentrum für Künstliche Intelligenz GmbH (DFKI), *Germany*
- Diamond Light Source, *UK*
- Forschungszentrum Jülich, *Germany*
- The Royal Institute of Technology, *KTH, S*
- Science and Technology Facilities Council, *UK*

CONTACT

Sai Narasimhamurthy
sai.narasimhamurthy@seagate.com
+44 (0)2392496648

The two racks of the SAGE system are now integrated into the Jülich Supercomputing Centre and undergoing software integration and application testing. Any communities interested in testing the SAGE prototype are advised to write to: info@sagestorage.eu.



The SAGE system

Basic Technology Projects starting 2018/2019



● 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.

exAScale Programing models for extreme Data processing

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.

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



● 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-in-

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

tensive 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, *Spain*

OTHER PARTNERS

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

CONTACT

Antonio J. Peña,
antonio.pena@bsc.es



● MOTIVATION

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

● OBJECTIVES

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

<https://epigram-hs.eu/>
🐦: @EpigramHs

COORDINATING ORGANISATION

KTH, Sweden

OTHER PARTNERS

- EPCC, UK
- ETH Zürich, Switzerland
- Fraunhofer, Germany
- Cray, Switzerland
- ECMWF, UK

CONTACT

Apostolos Vasileiadis,
apovas@kth.se

ESCAPE-2



● 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

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

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 Departement des Innern, *Switzerland*
- Barcelona Supercomputing Centre, *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

EXA2PRO



● 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 various fault-tolerance mechanisms, both user-exposed and at runtime system level and performance monitoring features. EXA2PRO will be evaluated using 4 use cases from 4 different domains, which will be deployed in JUELICH supercomputing centre. The use cases will leverage the EXA2PRO tool-chain and we expect:

1. Increased applications performance based on EXA2PRO optimization tools (data and memory management)
2. Efficient exploitation of heterogeneity by the applications that will allow the evaluation of more complex problems.

Enhancing Programmability and boosting Performance Portability for Exascale Computing Systems

3. Identification of trade-offs between design qualities (source code maintainability/reusability) and run-time constraints (performance/energy consumption).
4. Evaluation of various fault-tolerance mechanisms for applications with different characteristics.

EXA2PRO outcome is expected to have major impact in

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. On application developers of exascale application: EXA2PRO will provide tools for improving source code maintainability/reusability, which will allow application evaluation with reduced developers' effort.
3. On the scientific community and the industry relevant to the EXA2PRO use cases. At least two of the EXA2PRO use cases will have significant impact to the CO2 capture and to the Supercapacitors industry.

<https://exa2pro.eu/>

COORDINATING ORGANISATION

Institute of Communication and Computer Systems, *Greece*

OTHER PARTNERS

- Linköpings Universitet, *Sweden*
- 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*

CONTACT

Prof. Dimitrios Soudris,
dsoudris@microlab.ntua.gr

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

ExaQute - EXAscale Quantification of Uncertainties for Technology and Science Simulation

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, *Spain*
- Technische Universität München, *Germany*
- INRIA, Institut National de Recherche en Informatique et Automatique, *France*
- Vysoka Skola Banská - Technická Univerzita Ostrava, *Czech Republic*
- Ecole Polytechnique Fédérale de Lausanne, *Switzerland*
- Universitat Politècnica de Catalunya, *Spain*
- str.ucture GmbH, *Germany*

CONTACTS

Cecilia Soriano,
csoriano@cimne.upc.edu

Riccardo Rossi,
rrossi@cimne.upc.edu



● OBJECTIVE

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.

<https://www.maestro-data.eu/>

COORDINATING ORGANISATION

Forschungszentrum Jülich GmbH, *Germany*

OTHER PARTNERS

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

CONTACT

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

RECIPE



● 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

RELIABLE power and time-Constraints-aware Predictive management of heterogeneous Exascale systems

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.

<http://www.recipe-project.eu/>
📧: @EUrecipe

COORDINATING ORGANISATION

Politecnico di Milano, *Italy*

OTHER PARTNERS

- Universitat Politècnica de Valencia, *Spain*
- Centro Regionale Information Communication Technology srl, *Italy*
- Barcelona Supercomputing Centre, *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

Sage2



● 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 (Perceptient StorAGe for Exascale Data Centric

Perceptient Storage for Exascale Data Centric Computing 2

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.

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 Högskolan, *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



● OBJECTIVE

The purpose of this project is to enable a diverse set of multiscale, multiphysics applications -- from fusion and advanced materials through climate and migration, to drug discovery and the sharp end of clinical decision making in personalised medicine -- to run on current multi-petascale computers and emerging exascale environments with high fidelity such that their output is «actionable». That is, the calculations and simulations are certifiable as validated (V), verified (V) and equipped with uncertainty quantification (UQ) by tight error bars such that they may be relied upon for making important decisions in all the domains of concern. The central deliverable is an open source toolkit for multiscale VVUQ based on generic multiscale VV and UQ primitives, to be released in stages over the lifetime of this project, fully tested and evaluated in emerging exascale environments, actively promoted over the lifetime of this project, and made widely available in European HPC centres.

The project includes a fast track that will ensure applications are able to apply available multiscale VVUQ tools as soon as possible,

while guiding the deep track development of new capabilities and their integration into a wider set of production applications by the end of the project. The deep track includes the development of more disruptive and automated algorithms, and their exascale-aware implementation in a more intrusive way with respect to the underlying and pre-existing multiscale modelling and simulation schemes.

The potential impact of these certified multiscale simulations is enormous, and so we aim to promote the VVUQ toolkit across a wide range of scientific and social scientific domains, as well as within computational science more broadly.

www.vecma.eu/

🐦: @VECMA4

COORDINATING ORGANISATION

University College London, UK

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,
uccaxg0@ucl.ac.uk

● 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.

Coordinating Organisation

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

OTHER PARTNERS

- The University of Edinburgh, *United Kingdom*
- Kungliga Tekniska Högskolan, *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

Other HPC projects



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 EPI initiative is part of a broader strategy to develop an 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 the European Processor Initiative (EPI) is to design a roadmap for future European low power

processors for extreme scale computing (exascale), high-performance, big-data and emerging applications (e.g. automotive computing) and other fields that require highly efficient computing infrastructure. More precisely, EPI aims at establishing a roadmap to reach three fundamental goals:

1. Developing low-power processor technology to be included in a pre-exascale system for Europe in 2020-2021 and exascale in 2022-2023;
2. Ensuring that a significant part of that technology is European;
3. Ensuring that the application areas of the technology are not limited only to HPC, but cover other areas, such as automotive and data centres, thus ensuring the economic viability of the initiative.

EPI gathers together 23 partners from 10 European countries with a wide range of expertise and background: HPC, supercomputing centres, automotive computing,

including researchers and key industrial players. The fact that the envisioned European processor is planned to be based on already existing tools either owned by the partners or being offered as open-source with a large community of users, provides two 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 into their commercial portfolio or in their scientific roadmap.

EPI covers the 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. First stream is focused on a general, HPC processor chip. This generation is aiming to go into pre-exascale machines with industry proven technologies. Second stream focuses on accelerator chips and systems with the RISC-V architecture as the enabling component for exascale, low-power computing systems. The results from two streams 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

<http://european-processor-initiative.eu/>

COORDINATING ORGANISATION

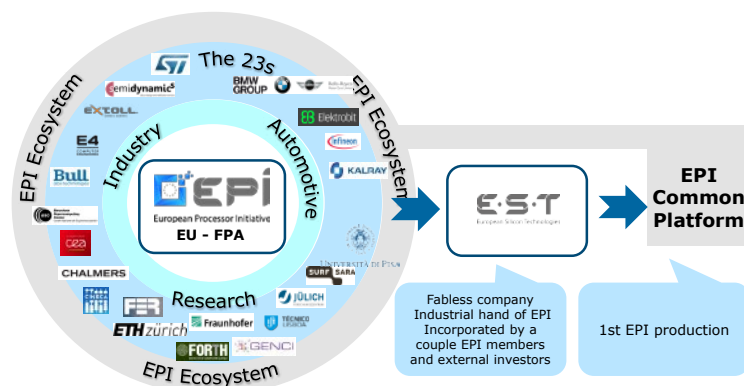
Bull SAS (Atos group), France

OTHER PARTNERS

- BSC, Barcelona Supercomputing Centre, Spain
- Infineon, Germany
- Semidynamics, Spain
- CEA, Commissariat à l'Energie Atomique et aux énergies alternatives, France
- Chalmers, Sweden
- ETHZ, Eidgenössische Technische Hochschule Zürich, Switzerland
- Forth, Foundation for Research and Technology Hellas, Greece
- GENCI, Grand Equipement National de Calcul Intensif, France
- IST, Instituto Superior Técnico, Portugal
- FZJ, Forschungszentrum Jülich GmbH, Germany
- Unibo, Università di Bologna, Italy
- UNIZG, Sveučilište u Zagrebu Fakultet Elektronike i Računarstva, Croatia
- Fraunhofer Gesellschaft zur Förderung der angewandten Forschung E.V., Germany
- STM, STMICROELECTRONICS SRL, Italy
- E4 Computer Engineering SPA, Italy
- Unipi, Università di Pisa, Italy
- SURFsara, Netherlands
- Kalray SA, France
- Extoll GmbH, Germany
- Cineca, Italy
- BMW, Bayerische Motoren Werke AG, Germany
- Elektrobit, Automotive GmbH, Germany

CONTACT

contact@european-processor-initiative.eu



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.

Mont-Blanc 2020

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 has a budget of 10.1 million Euros, funded by the European Commission under the Horizon2020 program. It was launched in December 2017.

It 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:

European scalable, modular and power-efficient HPC processor

1. define a low-power System-on-Chip architecture targeting Exascale;
2. implement new critical building blocks (IPs) and provide a blueprint for its first-generation implementation;
3. deliver initial proof-of-concept demonstration of its critical components on real life applications;
4. 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.

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 produced in the European Processor Initiative (EPI).

EUROPEAN HIGH-PERFORMANCE COMPUTING

HANDBOOK 2018



The Mont-Blanc 2020 physical architecture

montblanc-project.eu/
📧: @MontBlanc_Eu

COORDINATING ORGANISATION

Bull (Atos group), France

OTHER PARTNERS

- Arm (United Kingdom)
- BSC (Spain)
- CEA (France)
- Jülich Forschungszentrum (Germany)
- Kalray (France)
- SemiDynamics (Spain)

CONTACT

Etienne Walter
Etienne.walter@atos.net



A Key Output: the MB2020 demonstrator

PROCESS



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 five 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 industry pilot 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 Square Kilometre Array (SKA) project – 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.

PROviding Computing solutions for ExaScale Challenges

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 medical informatics, 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.

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. Stanisława Staszica W Krakowie, *Poland*

CONTACT

Nils gentschen Felde, Coordinator
Rubén Riestra, Communications Manager
contact@process-project.eu

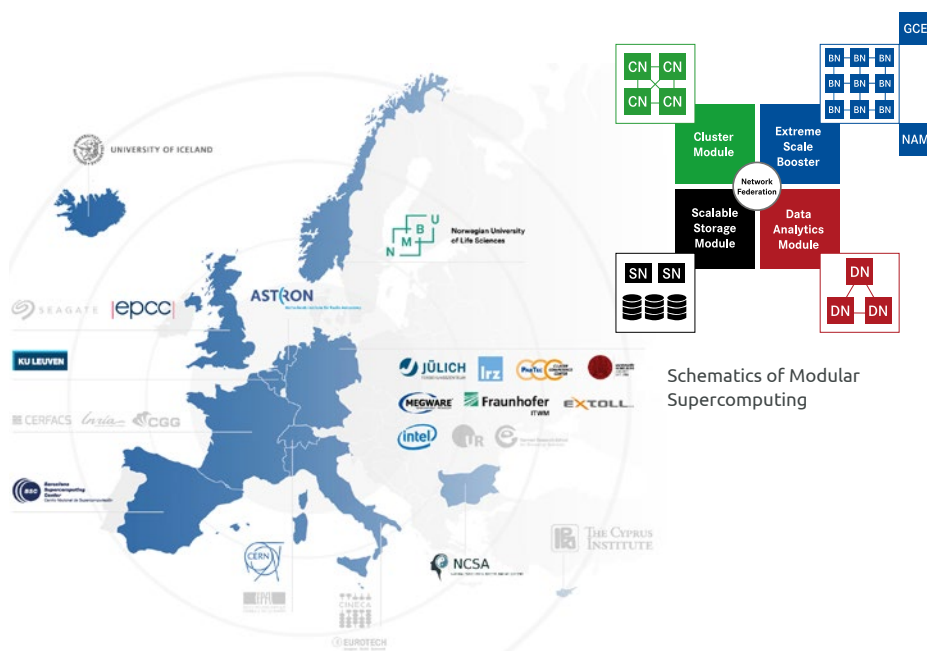
Co-design projects



The DEEP projects (DEEP, DEEP-ER and DEEP-EST) present an innovative solution for next generation supercomputers, aiming at most efficiently organising heterogeneous resources. This is achieved by addressing the main Exascale challenges – including scalability, programmability, end-

to-end performance, resiliency, and energy efficiency – through a stringent co-design approach.

The DEEP projects developed the Cluster-Booster architecture – which combines a standard HPC Cluster with the Booster,



a unique cluster of high-throughput many-core processors – and extended it by including a multi-level hierarchy based on innovative memory technologies. Additionally, a full software stack has been created by extending MPI – the de-facto standard programming model in HPC – and complementing it with task-based I/O, and resiliency functionalities. The next step in the DEEP project’s roadmap is the generalisation of the Cluster-Booster concept towards the so-called “Modular Supercomputing Architecture”, in which the Cluster and the Booster are complemented by further computing modules with characteristics tailored to the needs of new workloads, such as the ones present in high-performance data analytics (HPDA).

The developments cut across the complete HPC stack and amount to a fully integrated system prototype combining hardware with system software, programming environments and highly tuned applications. The latter are a total of 15 ambitious and highly relevant applications from HPC and HPDA domains, which drive co-design and serve to evaluate the projects' ideas and demonstrate their benefits.

● PROJECT AREAS OF INTERNATIONAL COLLABORATION

1. Prototype development: exchange of lessons learned
2. Application porting: Exchange of experience in code modernisation, especially KNL optimisation
3. Domain scientists / application developers: invited to test the project developments with their codes

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 Center
- 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 Miljø-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

EuroExa



EuroEXA is a program that represents a significant EU investment to innovate across a new ground-breaking platform for computing in its support to deliver Exa-Scale computing. Originally the informal name for a group of H2020 research projects,




A view of the EuroEXA Liquid Cooled Blade

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

ExaNeSt, EcoScale and ExaNoDe, EuroEXA now 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.

This project has a €20m investment over a 42-month period and is part of a total €50m investment made by the EC across the EuroEXA group of projects supporting research, innovation and action across applications, system software, hardware, networking, storage, liquid cooling and data centre technologies. Together bringing the technologies required to enable the digital economy, the future of computers, and the drive towards Exa-Scale ca-

EUROPEAN HIGH-PERFORMANCE COMPUTING



The EuroEXA Liquid Cooled Blade, 16 x Quad Processor Nodes in 1u Short Depth

pability. 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 by 2020, it is hoped that this will enable an Exa-Scale procurement for deployment in 2022/23.

"To deliver the demands of next generation computing and Exa-Scale HPC, it is not possible to simply optimize the components of the existing platform. In EuroEXA, we have taken a holistic approach to break-down the inefficiencies of the historic abstractions and bring significant innovation and co-design across the entire computing stack." John Goodacre, Professor of Computer Architectures at the University of Manchester

"This is a world class program that aims to increase EU computing capabilities by 100

HANDBOOK 2018

www.euroexa.eu

📧: @EuroEXA

COORDINATING ORGANISATION

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

OTHER PARTNERS

- Arm, *UK*
- The University of Manchester, *UK*
- BSC (Barcelona Supercomputing Centre), *Spain*
- FORTH (Foundation for Research and Technology Hellas), *Greece*
- The Hartree Centre of STFC, *UK*
- IMEC, *Belgium*
- ZeroPoint Technologies, *SE*
- Iceotope, *UK*
- Synelxis Solutions Ltd., *Greece*
- Maxeler Technologies, *UK*
- Neurasimus, *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

times, the EuroEXA project is truly an exceptional collection of EU engineering excellence in this field. We have all set our ultimate goal – to enable the power-efficient delivery of the world's biggest supercomputer" Peter Hopton, Founder of Iceotope and Dissemination Lead for EuroEXA

As part of the H2020 competitive process, the 16 organisations of EuroEXA (above) have been selected for their technologies and capabilities from across 8 Countries.

Centres of Excellence in computing applications started 2015



BioExcel is the European Centre of Excellence (CoE) for provision of support to academic and industrial researchers in the use of high-end computing in biomolecular research. The centre works on enabling better science by improving the most popular biomolecular software and spreading best practices and expertise among the communities through consultancy and training. The long-term vision of the centre is to be a central hub for biomolecular modelling and simulations.

PROJECT DESCRIPTION

Much of the current Life Sciences research relies on intensive biomolecular modelling and simulation. As a result of this, both academic and industrial researchers are facing significant challenges when it comes to applying best practices for optimal resource usage and workflow productivity, and to finding a faster path to achieve results. *High-performance computing* (HPC) and *high-throughput computing* (HTC) techniques have now reached a level of maturity in widely used codes and platforms, but taking full advantage of these requires training and guidance by experts. The services

ecosystem required for that is presently inadequate, so a suitable infrastructure needs to be set up in a sustainable way. BioExcel CoE was thus established to provide the necessary solutions for long-term support of the biomolecular research communities: fast and scalable software, user-friendly automation workflows and a support base of expert core developers. The main services offered by the centre include hands-on training, tailored customization of code and personalized consultancy support. The BioExcel CoE actively engages with a number of interest groups, formed by members of academic and industrial communities, which lay the foundation of the long-term basis for user-driven governance of the centre:

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

www.bioexcel.eu

Twitter: @BioExcelCoE

Youtube: <https://goo.gl/5dBzmm>

AREAS OF INTERNATIONAL COLLABORATION

The centre engages with a number of international organizations and initiatives which have complementary efforts in the area of computational biomolecular research. Some of them include:

- ELIXIR (www.elixir-europe.org)
- Molecular Science Software Institute (www.molssi.org)
- Software Sustainability Institute (SSI) (www.software.ac.uk)
- Open PHACTS foundation (www.openphacts.org)
- PRACE (www.prace-ri.eu)
- Common Workflow Language (www.commonwl.org)
- HPC-Europa3 (www.hpc-europa.eu)

COORDINATING ORGANISATION

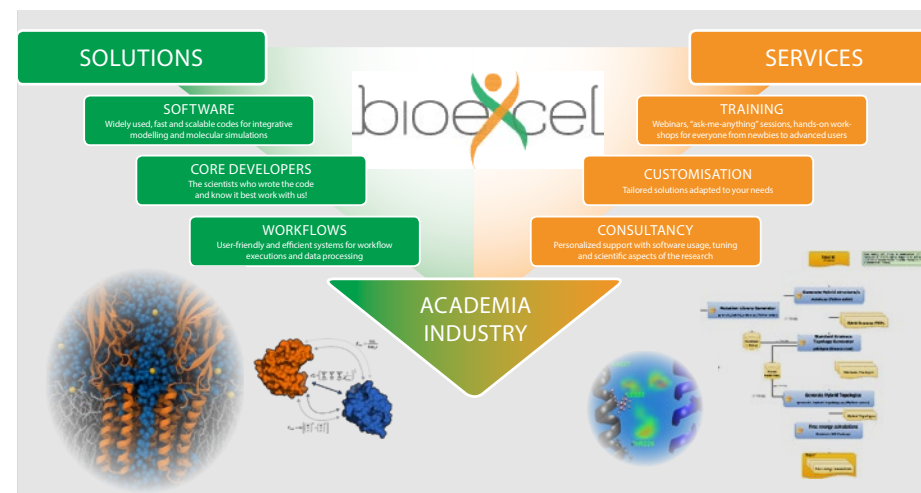
KTH Royal Institute of Technology, SE

OTHER PARTNERS

- The University of Manchester, UK
- University of Utrecht, Netherlands
- Institute of Research in Biomedicine (IRB), Spain
- European Molecular Biology Laboratory (EMBL-EBI), UK
- Forschungszentrum Jülich, Germany
- The University of Edinburgh, UK
- Max Planck Gesellschaft, Germany
- Forward Technologies, SE
- Barcelona Supercomputing Centre, Spain
- Ian Harrow Consulting, UK

CONTACT

Erwin Laure
erwinl@pdc.kth.se
+46 70-404 09 27





The globalisation of humanity's social and industrial activities, as observed over the past decades, has caused a growing need to address the global risks and opportunities involved. Some of these prominent challenges include:

- The global health risks – from diabetes to pandemics – involved in the spread of unhealthy social habits as well as the opportunity to achieve major global health improvements through healthy behaviour.
- The global diffusion of green growth initiatives, including policy initiatives, business strategies and lifestyle changes for successful as well as inefficient pathways.
- The challenges of global urbanisation, with special focus on the impact of infrastructure decisions regarding indicators like congestion, real estate prices and greenhouse gas emissions.

Approaches that address the above-mentioned challenges are investigated by a newly-emerging research area: Global Systems Science (GSS). However, with these transdisciplinary problems the demand for compute performance due to data and time constraints increases drastically so that the assistance of High Performance Computing

(HPC) is mandatory. Therefore, the CoeGSS Project brought together both communities and developed an environment for successful collaboration between the stakeholders dealing with Global Challenges on the one hand and the High Performance Computing institutions that provide the mandatory capabilities to address those complex challenges on the other.

So far, the use of HPC in GSS studies for processing, simulating, analysing, and visualizing large and complex data has been very limited due to a lack of tailored HPC-enabled tools and technologies. Within CoeGSS, those issues have been addressed by developing models, simulations, but also the technology and tools to interpret the output data efficiently. Typical GSS applications are data-bound, whereas traditional HPC tools and libraries are optimized to solve computationally intensive problems so that their applicability is limited for this emerging domain. Consequently, one major goal of the CoeGSS project was to establish a robust technology base in order to execute network, and in particular memory intensive applications in a scalable manner. Nevertheless, the main difference between

www.coegss-project.eu

📧: @CoeGSS

COORDINATING ORGANISATION

- Coordinator: University of Potsdam, *Germany*
- Technical Coordinator: High Performance Computing Centre Stuttgart, University of Stuttgart, *Germany*

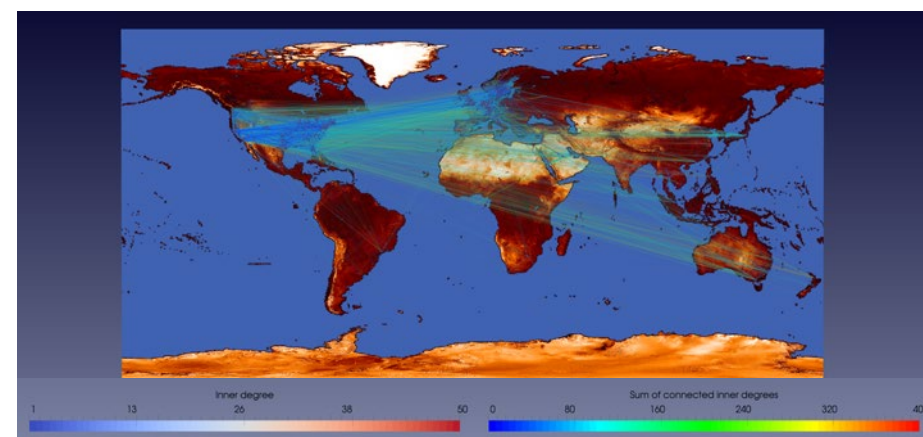
OTHER PARTNERS

- Atos Spain SA, *Spain*
- PSNC - Poznań Supercomputing and Networking Centre, *Poland*
- The CoSMo Company, *France*
- Chalmers Tekniska Högskola AB, *Sweden*
- CSP s.c.a r.l., *Italy*
- DIALOGIK, Non-profit Institute for Communication and Cooperation Research, *Germany*
- Global Climate Forum EV, *Germany*
- IMT School for Advanced Studies Lucca, *Italy*
- ISI, Institute for Scientific Interchange, *Italy*
- TOP-IX Consortium UP, *Italy*

CONTACT

Prof. Christoph Kreitz,
kreitz@uni-potsdam.de
+49 331 977 3060

typical HPC and GSS applications lies in the data sources and outputs as well as the used algorithms. Whereas lots of traditional high-performance application codes, like those of computational fluid dynamics, require massive parallelism and high computational power, GSS applications demand additional capabilities, for instance huge and varying data or in a generic manner, data-centric computation. Thus, looking for a trade-off between the data-centric programming models of the Cloud infrastructures and the highly efficient and scalable HPC technologies has been one of the key challenges for the CoeGSS project.





CompBioMed is a user driven CoE designed to nurture and promote the up- take and exploitation of HPC within the biomedical modelling community. Three distinct exemplar research areas will be pursued: cardiovascular, molecularly-based and neuro-musculoskeletal medicine. This will feed directly back to the clinic enabling healthcare providers to understand the mass of available data and provide clinical decision support

● PROJECT DESCRIPTION

CompBioMed is comprised of 15 Core Partners from across Europe and within academic, industrial and clinical research institutions. We have supplemented this with a growing number of Associate Partners, with whom we can work to provide HPC-solutions and draw on their knowledge to grow our own resources. The research in the project investigates three main pipelines: cardiovascular,



www.compbioimed.eu

🐦: @bio_comp

COORDINATING ORGANISATION

University College London (UCL)

OTHER PARTNERS

- University of Amsterdam, *Netherlands*
- The University of Edinburgh, *UK*
- SURFsara, *Netherlands*
- Barcelona Supercomputing Centre (BSC), *Spain*
- University of Oxford, *UK*
- University of Geneva, *Switzerland*
- The University of Sheffield, *UK*
- CBK Sci Con Ltd, *UK*
- Universitat Pompeu Fabra, *Spain*
- LifeTec Group, *Netherlands*
- Acellera, *Spain*
- Evotec, *UK*
- Bull SAS (Atos Group), *France*
- Janssen Pharmaceutica, *Belgium*

CONTACT

Emily Lumley
e.lumley@ucl.ac.uk
+4420 7679 5300

molecular-based medicine and neuro-musculoskeletal medicine. Functional codes are currently being produced from these pipelines that are being shared and tested. The Centre of Excellence has so far employed 50 codes, of which 1/3 are open source and a similar number can be classified as HPC.

Through a symbiotic relationship with clinical partners, we will use the imaging and medical data provided by them to feed in to our codes and establish programs that will enable healthcare providers to find the most applicable solution for time-limited medical decisions.

● AREAS OF INTERNATIONAL COLLABORATION

The number of academic and industrial Associate Partners is constantly growing. However, we would be keen to engage further with healthcare providers and clinical practitioners. We are also looking to involve further countries classified as having fewer HPC resources through the Innovation Exchange Program.

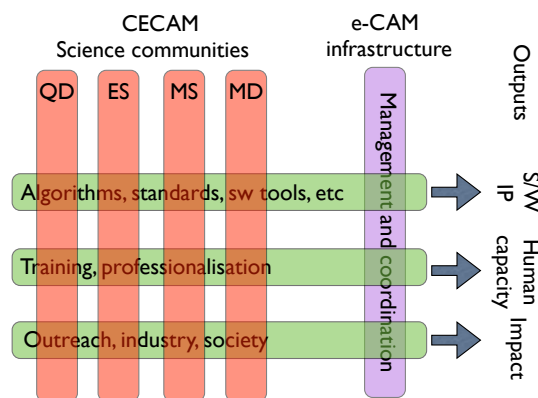


The E-CAM Centre of Excellence is an e-infrastructure for software development, training and industrial discussion in simulation and modelling. E-CAM is based around the experience and scientific community of the extended CECAM family, established over more than four decades and central in promoting fundamental research on advanced computational methods, as well as the computational and hardware expertise of the European partnership PRACE. We are a partnership of 16 CECAM nodes, 3 PRACE centres, 13 industrial partners and

one Centre for Industrial Computing (the Hartree Centre at Daresbury). Our training and software development activities are spread across Europe at the different Node locations.

The goals of E-CAM are pursued via a set of coordinated actions and networking. Its main tasks are as follows:

- software development targeted at specific end-users needs, and including testing of scaling and portability;
- development of the E-CAM repository, an open source repository of software-modules and packages for simulations in academy, material and life science, engineering. Modules include up to date documentation and benchmarks;
- training and dissemination in the field of computational science via a series of workshops and online training modules;
- extended software development workshops for production of modules for the repository based on input from the



Positioning of the e-CAM infrastructure

community and the industrial partners;

- scoping workshops for discussion and problem definition with industrial end-users;
- state-of-the-art workshops for monitoring developments, identifying new directions and transferring knowledge to industry;
- support for academic and industrial research via a set of pilot projects supervised by scientists in the team and sustained by E-CAM funded post-doctoral fellows;
- keeping application developers up-to-date in software developments for current and future HPC systems;
- service-on-demand model to support industry community towards extreme scale HPC applications.

Our approach is focused on four scientific areas, critical for high-performance simulations relevant to key societal and industrial challenges. These areas are classical molecular dynamics (MD), electronic structure calculations (ES), quantum dynamics (QD) and meso- and multi-scale modelling (MS). 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.

www.e-cam2020.eu

📧: @ECAM2020

COORDINATING ORGANISATION

Ecole Polytechnique Fédérale de Lausanne, Switzerland

OTHER PARTNERS

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

CONTACTS

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




The Energy oriented Centre of Excellence in computing applications (EoCoE: read as “Echo”) is focusing on simulation for renewable energy production, storage and delivery and on related transversal high-end expertise in applied mathematics and HPC. EoCoE leverages HPC to foster and accelerate the European transition to a reliable and low carbon energy supply, harnessing the coming revolution in hardware and HPC architectures, calling for a similar paradigm change in the way application codes are designed. EoCoE assists the energy transition via targeted support to four numerical-modelling-hungry areas related to renewable energy:

- **Meteorology:** making real time weather “nowcasting” possible in order to efficiently couple solar and wind farm energy production to the grid,
- **Materials:** enabling high-throughput computing to discover novel materials for batteries, photo-voltaic cells or supercapacitors,
- **Water:** geothermal and hydro-power – using HPC for the management of resources, strategy of usage and estimating the influence of climate change,

- **Fusion:** code coupling and mesh generation at the Exascale for fusion for energy simulation.

These four pillars are anchored within a strong transversal multidisciplinary basis providing high-end expertise in applied mathematics and HPC. This basis is using and developing new programming tools, models and algorithms in order to prepare the applications to the next generation of hardware. EoCoE is structured around a central Franco-German hub coordinating a pan-European network, gathering a total of 8 countries and 23 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.



● SUCCESS STORIES

FUSION— Tokamak Physics at IRFM, Cadarache: “Without the help of the EoCoE network, engaging in these new paths toward exascale would be much more difficult.”

GYSELA code describes a set of major phenomena that take place inside Tokamaks: instabilities, turbulent transport, plasma-wall interaction, and heating. EoCoE project and the systematic code auditing procedure it enabled has largely contributed to the improvement of GYSELA performance.

“EoCoE: Making realtime weather nowcasting possible in post-Moore era”

SolarNowcast aims to forecast the solar irradiation from fisheye lens from webcam images.

www.eocoe.eu

🐦: @EoCoE_project

COORDINATING ORGANISATION

CEA — Maison de la Simulation, France

OTHER PARTNERS

- BATH University
- BSC
- CEA
- CERFACS
- CNR
- CNRS
- CYI
- EDF R&D
- ENEA
- Fraunhofer / IWES
- FZJ
- INRIA, Institut National de Recherche en Informatique et Automatique
- MPG
- PSNC
- RWTH
- AACHEN
- ULB
- UNITN
- UNITOV
- UPSud
- UVSQ
- SUN

CONTACT

Edouard Audit,
edouard.audit@cea.fr

A full performance evaluation on the Forecast code allowed to identify a big optimization potential, both on serial and parallel levels. Optimization efforts improved the execution time by 2 to 5, getting closer to real time effective use.



● 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 compu-

tational 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 below 2.5-km global simulation using the ICON model. Performance predic-

tions 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

www.esiwace.eu

📍: @ESIWACE

COORDINATING ORGANISATION

Deutsches Klimarechenzentrum (DKRZ),
Germany

OTHER PARTNERS

- CNRS- Centre National de la Recherche Scientifique, *France*
- Barcelona Supercomputing Centre, *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, *UK*

CONTACT

Joachim Biercamp,
esiwace@lists.dkrz.de



MaX - Materials design at the eXascale is a Centre of Excellence with focus on: driving the evolution and exascale transition of materials science codes; creating an integrated ecosystem of codes, data, workflows and analysis tools for HPC and high-throughput computing; supporting and training developers and end-users in academia and industry.

● PROJECT DESCRIPTION

The starting point of MaX is the recognized strength of Europe in 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. The effort to prepare such codes for the exascale transition is beyond the possibilities of individual research teams.

MaX addresses this challenge by focusing on five complementary open-source codes (Siesta, Quantum Espresso, Yambo, Fleur, and the Aiida materials informatics infrastructure). Five developing teams of such codes work with experts from five HPC centres, in addition to three partners focused on business/dissemination. With them, MaX has undertaken major code refactoring, released new versions with enhanced performance, modularity and inter-operability as well as new capabilities, developed workflows and turn-key solutions for properties calculations and curated data sharing, organized major training events. It has produced first kernel and domain specific libraries, progressed towards green computing ('energy to solution' measures), and in hardware-software co-design. MaX Users Portal offers basic and advanced services supporting science and industry. The MaX Centre of Excellence, started in 2015, was recently approved for a second phase running until 2021.

www.max-centre.eu

📧: @max_center2

COORDINATING ORGANISATION

CNR Nano, Modena, *Italy* (Elisa Molinari)


OTHER PARTNERS

- SISSA Trieste, *Italy* (Stefano Baroni)
- ICN2 Barcelona, *Spain* (Pablo Ordejón)
- FZJ Jülich, *Germany* (Stefan Blügel, Dirk Pleiter)
- EPFL Lausanne, *Switzerland* (Nicola Marzari)
- CINECA Bologna, *Italy* (Carlo Cavazzoni)
- BSC Barcelona, *Spain* (José Cela)
- CSCS ETH Zürich, *Switzerland* (Thomas Schülthess)
- KTH Stockholm, *Sweden* (Erwin Laure)
- E4 Computer Engineering, *Italy* (Fabrizio Magugliani)
- Cloudweavers Ltd, *London UK* (Carlo Daffara)
- ICTP UNESCO, Trieste, *Italy* (Ivan Girotto)

CONTACT

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

NOMAD



NOVEL MATERIALS DISCOVERY

The *NOMAD Laboratory CoE* maintains the largest *repository* of open-access data derived from the most important materials-science codes available today, with more than 50 million open-access total-energy calculations. It provides these data also in a code-independent normalized form. Artificial Intelligence (AI) tools and data-mining services were developed for the discovery of novel information to advance materials science and engineering.

PROJECT DESCRIPTION

As data in the *NOMAD Repository* is generated from different codes, it is heterogeneous and difficult to use for data analytics and AI. To overcome this obstacle, the NOMAD team developed methods to convert heterogeneous data from over 40 major computational materials-science codes to a homogeneous, code-independent format. In parallel, tools were developed to intuitively access, query, and visualize this data. The *NOMAD Encyclopedia* is a user-friendly, public access point to NOMAD's code-independent data that allows users to explore the data in order to understand the structural, mechanical, and thermal behaviours

The Novel Materials Discovery Laboratory - a European Centre of Excellence

of a large variety of materials, as well as their opto-electronic properties, and more. The team also developed a variety of tools for the *NOMAD Big-Data Analytics Toolkit*. These tools will help scientists and engineers to select materials that will be most useful for specific applications, and to identify or predict promising new materials worth further exploration. In order to allow for interactive data exploration, as well as for enhanced training and dissemination, prototypes for *advanced visualization tools* (remote visualization and virtual reality) were also developed. Delivery of these tools and services is made possible by our high-performance computing expertise and hardware. The NOMAD concept is also highlighted in a recent invited review paper: C. Draxl, M. Scheffler, «NOMAD: The FAIR Concept for Big-Data-Driven Materials Science» (MRS Bulletin 43, 676 (2018)).

AREAS OF INTERNATIONAL COLLABORATION

The *NOMAD Laboratory CoE* is well-placed to collaborate with other international initiatives in computational materials science, within the EU but also from other conti-

nents. Furthermore, NOMAD initiated and participates in the founding of the non-profit association "FAIR Data Infrastructure for Physics, Chemistry, Materials Science, and Astronomy e.V." (<https://fairdi.eu>). FAIR's mission is to build a worldwide data infrastructure for big data from materials science, engineering, and astrophysics that follows the FAIR principles. FAIR stands for Findable, Accessible, Interoperable, Re-purposable/Reusable as suggested by Wilkinson *et al.* (Sci. Data 3, 160018, 2016). The association will start with 5 pillars: a) Computational materials science (NOMAD); b) Experimental materials science; c) Biophysical and soft-matter simulations; d) Astrophysics and space-situational awareness; e) User management, intellectual property rights, cyber security, etc. (horizontal).

www.nomad-coe.eu

Twitter: @NoMaDCoE

COORDINATING ORGANISATION

Fritz Haber Institute
of the Max Planck Society

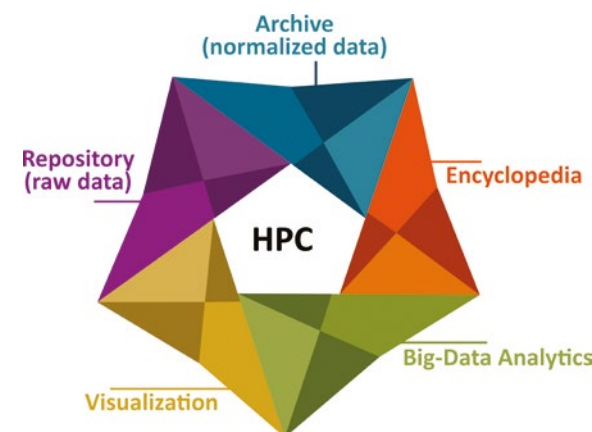
OTHER PARTNERS

- Leibniz Supercomputing Centre, Garching
- Barcelona Supercomputing Centre
- Max Planck Computing and Data Facility
- CSC — IT Center for Science
- Aalto University
- University of Barcelona
- University of Cambridge
- Humboldt-Universität zu Berlin
- King's College London
- Max Planck Institute for the Structure of Dynamics of Matter
- Technical University of Denmark
- Pintail Ltd

CONTACTS

Matthias Scheffler
scheffler@fhi-berlin.mpg.de
+49 30 8413 4711

Kylie O'Brien
kylie.obrien@pintailservices.com
+353 1 841 8046





The POP objective is to help improve the performance obtained by applications on current systems, help identify the really important issues in a quantitative way and help the community maximize the productivity of their development efforts when addressing very large scales.

POP is providing a service to customers from all application domains (engineering, sciences, data analytics, climate, medical, media, deep learning...). Customers apply by filling a small request service form in the project web page (www.pop-coe.eu). The main service is an “assessment” report where an external view is given to the customer on the performance issues experience by the application and recommendations on what would be the most productive code refactoring directions. A unified set of metrics and methodology have been set up to homogenize the reports across partners. In some cases, a “proof of concept (POC)” follow up service is performed where POP experts help customers to apply the proposed techniques and demonstrate the actual performance gains.

The customer gets an assessment report between 5 and 15 pages with a description of the identified issues and recommendations on how to address them. Performance gains from just 10% up to 10x have already been observed on different POC services. Close to 100 assessments have already been performed. UK, Germany and Spain are the countries with more services received, but almost 25% of the services are to countries outside the consortium. The POP CoE is actually a hub establishing links and sharing experiences among a very large community.

One of the activities of the project is to gather feedback from customers on the quality of the service provided so that we can improve our operation. Overall, the customers are very happy and value very positively the assessments and recommendations received. Under their permission, reports are made available in the project web page. We also include there summarized success stories and customer testimonials.

www.pop-coe.eu

📧: @POP_HPC

COORDINATING ORGANISATION

Barcelona Supercomputing Centre, *Spain*

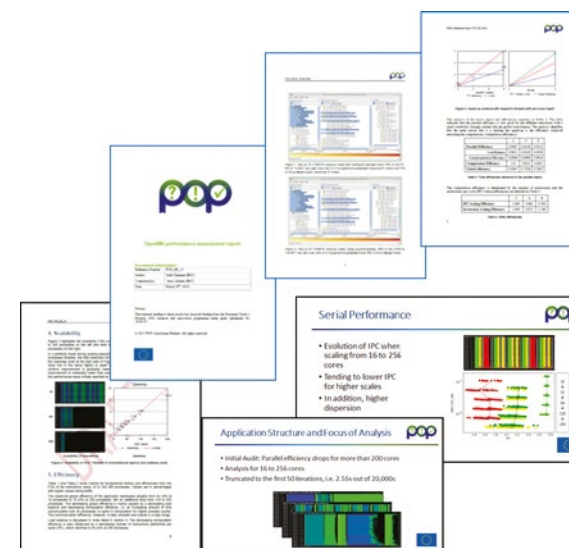
PARTNERS

- Universitaet Stuttgart HLRS, *Germany*
- Forschungszentrum Jülich GmbH, *Germany*
- Numerical Algorithms Group Ltd, *UK*
- Rheinisch-Westfaelische Technische Hochschule Aachen (RWTH-AACHEN), *Germany*
- Teratec, *France*

CONTACTS

Jesus Labarta
jesus.labarta@bsc.es
+34 934 127 642

Judit Gimenez
judit.gimenez@bsc.es
+34 934 017 178



Centres of Excellence in computing applications starting 2018



BioExcel is the European Centre of Excellence (CoE) for providing support to academic and industrial researchers in the use of high-end computing for biomolecular research. The centre facilitates better research by improving the most popular biomolecular software and spreading best practices and expertise among the biomolecular research communities through consultancy and training. The long-term vision of the centre is to be a central hub for biomolecular modelling and simulations.

PROJECT DESCRIPTION

Much of the current Life Sciences research relies on intensive biomolecular modelling and simulation. As a result of this, both academic and industrial researchers are facing significant challenges when it comes to applying best practices for optimal resource usage and workflow productivity, and in terms of finding faster paths to achieving results.

High-performance computing (HPC) and *high-throughput computing (HTC)* techniques have now reached a level of maturity for the widely-used codes and platforms that are available, but taking full advantage

of these requires that researchers have access to training and guidance by experts. The services ecosystem necessary for that is presently inadequate, so a suitable infrastructure needs to be set up in a sustainable way.

BioExcel CoE is providing the necessary **solutions** for the long-term support of the biomolecular research communities: fast and scalable software, user-friendly automation workflows and a support base of expert core developers. The main **services** offered by the centre include hands-on training, tailored customization of code and personalized consultancy support.

The BioExcel CoE actively engages with a number of interest groups, consisting of members of academic and industrial research communities, including:

- **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**.

<https://bioexcel.eu/>
 @BioExcelCoE
 Youtube: <https://goo.gl/5dBzmmw>

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*
- The University of Edinburgh, *UK*
- Max Planck Gesellschaft, *Germany*
- Across Limits, *Malta*
- Barcelona Supercomputing Centre, *Spain*
- Ian Harrow Consulting, *UK*
- Norman Consulting, *Norway*
- Nostrum Biodiscovery, *Spain*

CONTACT

Erwin Laure,
erwinl@pdc.kth.se



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 sustain 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.

www.cheese-coe.eu

COORDINATING ORGANISATION

Barcelona Supercomputing Centre, *Spain*

OTHER PARTNERS

- Istituto Nazionale di Geofisica e Vulcanologia (INGV), *Italy*
- Icelandic Meteorological Office (IMO), *Iceland*
- Swiss Federal Institute of Technology (ETH), *Switzerland*
- University of Stuttgart - High Performance Computing Center (HLRS), *Germany*
- CINECA, *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*
- National Center for Scientific Research Marseille (CNRS), *France*
- Bull SAS (Atos Group), *France*

CONTACT

Arnau Folch,
arnau.folch@bsc.es



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.

In CompBioMed2 we are extending the CoE to serve the community for a total of more than 6 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. Computatio-

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

nal Biomedicine is an emergent technology that 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.



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

www.eocoe.eu/
 @EoCoE_project

COORDINATING ORGANISATION

CEA – Maison de la Simulation, France

OTHER PARTNERS

- BATH University
- BSC – Barcelona Supercomputing Centre
- CERFACS
- CIEMAT
- CNR
- CNRS
- DDN
- ENEA
- FAU
- Fraunhofer/IEEE
- FZJ
- IFPEN
- INRIA, Institut National de Recherche en Informatique et Automatique
- MPG
- PSNC
- ULB
- UNITN

CONTACT

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

new, long lasting and sustainable community around computational energy science. EoCoE resolves current bottlenecks in application codes; it develops cutting-edge mathematical and numerical methods, and tools to foster the usage of Exascale computing. Dedicated services for laboratories and industries are established to leverage this expertise and to develop an ecosystem around HPC for energy.

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



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 – *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*
- ETH 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



● 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* industrial field with the highest Exascale potential, thus EXCELLERAT brings together the necessary European expertise to establish a Centre of Excellence in Engineering with a broad service portfolio, paving the way for the evolution towards EXASCALE. All with-

in the frame of the European HPC Strategy realization just pushed forward with the activities on the EuroHPC Joint Undertaking.

To fulfil its mission, EXCELLERAT will base on six carefully chosen reference applications (Nek5000, Alya, AVBP, Fluidity, FENICS, Fluc), which were analysed on their potential to support the aim to achieve EXASCALE performance in HPC for Engineering. Thus, they are promising can-



Areas of expertise forming the EXCELLERAT baseline

www.excellerat.eu/

📧: @EXCELLERAT_CoE

COORDINATING ORGANISATION

Universität Stuttgart (HLRS), *Germany*

OTHER PARTNERS

- The University of Edinburgh (EPCC), *United Kingdom*
- CINECA, *Italy*
- SICOS BW GmbH, *Germany*
- Kungliga Tekniska Högskolan (KTH), *Sweden*
- ARCTUR DOO, *Slovenia*
- Deutsches ZENTRUM FUER LUFT - UND RAUMFAHRT EV, *Germany*
- Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS), *France*
- Barcelona Supercomputing Centre, *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

didates to be executed on the Exascale Demonstrators, Pre-Exascale Systems and Exascale Machines.

EXCELLERAT addresses the setup of a centre, 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 or Knowledge Transfer and having for all reference applications the developers on board, 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 so-called interest groups to integrate external stakeholders of its value network into its evolution.

● OBJECTIVE

FocusCoE will contribute 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 will do 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 will provide 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 OF FOCUS-COE ARE

- To create a platform, the EU HPC CoE General Assembly, that allows all 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.
- 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)
- 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.

www.focus-coe.eu

COORDINATING ORGANISATION

Scapos AG, *Germany*

OTHER PARTNERS

- Kungliga 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*

CONTACT

Guy Lonsdale,
guy.lonsdale@scapos.com



● 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.

COORDINATING ORGANISATION

Atos Spain, *Spain*

OTHER PARTNERS

- Universitaet Stuttgart – High Performance Computing Center Stuttgart, *Germany*
- Instytut Chemii Bioorganicznej Polskiej Akademii Nauk – Poznan Supercomputing and Networking Center, *Poland*
- 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



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 will focus on selected comple-

mentary 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 planned 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.

Work and lessons learned on MaX flagship codes and projects will be made available to the developers and users communities at large.

www.max-centre.eu/
: @max_center2

COORDINATING ORGANISATION

CNR Nano, Modena, *Italy* (Elisa Molinari)

OTHER PARTNERS

- SISSA Trieste, *Italy* (Stefano Baroni)
- ICN2 Barcelona, *Spain* (Pablo Ordejón)
- FZJ Jülich, *Germany* (Stefan Blügel, Dirk Pleiter)
- EPFL Lausanne, *Switzerland* (Nicola Marzari)
- CINECA Bologna, *Italy* (Carlo Cavazzoni)
- BSC Barcelona, *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



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, 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 Banská - Technická Univerzita Ostrava, Czech Republic

CONTACTS

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

Imprint:

© Etp4Hpc

Text:

Etp4Hpc

Graphic design and layout:

www.maiffret.net

Paper:

Arcoprint

Printed and bound in november 2018:

ADM, Paris

Contact

contact@etp4epc.eu

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