European Exascale Projects and their International Collaboration Potential

Supercomputing Conference 2016 (SC'16), Salt Lake City Utah

Birds-of-a-Feather Session (BOF)

on Wednesday 16th of November 2016 in Room 155-A, from 10:30am to 12pm

Organised by ETP4HPC on behalf of EXDCI

The following project details are included:

● *Project name* ● *Website address* ● *Contact Name* ● *Contact email* ●*Contact phone* ●*Logo* ● *Partners Involved* ● *Which areas does your project target?* ● Please list the main achievements of your project in one sentence ● *Project Profile* ● Please summarise, in one sentence, the potential areas of international collaboration of your project. ● What future partners are you looking to collaborate with?

For additional information and project summaries please visit:

A list of the European Horizon 2020 projects:

http://cordis.europa.eu/projects/home_en.html (search for FETHPC)

EXDCI summaries: https://exdci.eu/collaboration/fethpc https://exdci.eu/collaboration/exascale-projects https://exdci.eu/collaboration/coe

www.etp4hpc.eu - the website of ETP4HPC (the European Technology Platform for HPC)

http://exascale-projects.eu/ - an overview on FP7-funded projects

Table of Contents

Allscale - An Exascale Programming, Multi-objective Optimisation and Resilience Management Environment Based on Nested Recursive Parallelism

Website address http://www.allscale.eu

Contact name Thomas Fahringer, University of Innsbruck

Contact email tf@dps.uibk.ac.at

Contact telephone +43 6768 725 64410

Partners involved: University of Innsbruck, Friedrich-Alexander-University of Erlangen-Nuremberg, Queen's University of Belfast, Royal Institute of Technology, Numeca International S.A. IBM Ireland Limited

Which areas does your project target?

Programming methodologies, environments, languages and tool

Please list the main achievements of your project in one sentence

Develop an API and programming environment for extreme-scale architectures that exploits nested recursive paralelism and which is based on a compiler and runtime system implementing multi-objective optimisations and resilience management.

Project profile

Extreme scale HPC systems impose significant challenges for developers aiming at obtaining applications efficiently utilising all available resources. In particular, the development of such applications is accompanied by the complex and labour-intensive task of managing parallel control flows, data dependencies and underlying hardware resources – each of these obligations constituting challenging problems on its own. The AllScale environment, the focus of this project, will provide a novel, sophisticated approach enabling the decoupling of the specification of parallelism from the associated management activities during program execution. Its foundation is a parallel programming model based on nested recursive parallelism, opening up the potential for a variety of compiler and runtime system based techniques adding to the capabilities of resulting applications. These include the (i) automated porting of application from small- to extreme scale architectures, (ii) the flexible tuning of the program execution to satisfy trade-offs among multiple objectives including execution time, energy and resource usage, (iii) the management of hardware resources and associated parameters (e.g. clock speed), (iv) the integration of resilience management measures to compensate for isolated hardware failures and (v) the possibility of online performance monitoring and analysis. All these services will be provided in an application independent, reusable fashion by a combination of sophisticated, modular, and customizable compiler and runtime system based solutions.

AllScale will boost the development productivity, portability, and runtime, energy, and resource efficiency of parallel applications targeting small to extreme scale parallel systems by leveraging the inherent advantages of nested recursive parallelism, and will be validated with applications from fluid dynamics as well as environmental hazard and space weather simulations provided by SME, industry and scientific partners.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

We see great potential to collaborate with compiler, runtime system, resilience and online performance monitoring developers and with application developers that can benefit by nested recursive parallelism for small to extreme-scale parallel architectures.

What future partners are you looking to collaborate with?

Joint workshops, BOF sessions, joint calls for funding by both the EU and NSF.

ANTAREX AutoTuning and Adaptivity appRoach for Energy efficient eXascale HPC systems

Website address http://www.antarex-project.eu/

Contact name Prof. Cristina Silvano

Contact email cristina.silvano@polimi.it

Contact telephone +39 02 2399 3692

Partners involved: Politecnico di Milano (Italy), ETH Zurich (Switzerland), Universidate do Porto (Portugal), INRIA Rennes (France), Consorzio Interuniversitario CINECA (Italy), IT$Innovations (Czech Republic), Dompé Farmaceutici SPA (Italy), Sygic (Slovakia)

Which areas does your project target?

Programming methodologies, environments, languages and tools APIs and system software for future extreme scale systems

Project profile

To reach Exascale computing (1018 FLOPs), current supercomputers must achieve an energy efficiency "quantum leap" that allows this level of computation to be done at around 20 Megawatts. This will only be possible if we can target all layers of the system, from thesoftware stack to the cooling system.

ANTAREX will solve these challenging problems by proposing a disruptive holistic approach spanning all the decision layers composing the supercomputer software stack and exploiting effectively the full system capabilities (including heterogeneity and energy management). The main goal of the ANTAREX project is to provide a breakthrough approach to express by a Domain Specific Language the application self-adaptivity and to runtime manage and auto tune applications for green and heterogeneous High Performance Computing (HPC) systems upto the Exascale level.

CRESTA - Collaborative Research into Exascale Systemware, Tools and Applications

Website address http://www.cresta-project.eu

Contact name Dr Lorna Smith

Contact email l.smith@epcc.ed.ac.uk

Contact telephone +44 131 651 3408

Partners involved: Åbo Akademi University (ABO,) Allinea, Software, Cray UK Limited (CRAY), DLR, Ecole Centrale Paris (ECP), EPCC, The University of Edinburgh, Kungliga, Tekniska, Högskolan (KTH), The Center for Information Services and High Performance, Computing (ZIH), The European Centre for Medium-Range Weather Forecasts, The University of Jyvaskyla, Department of Physics The University of Stuttgart (USTUTT), Tieteen Tietotekniikan Keskus OY (CSC)'s, IT Center For Science LTD., University College London

Which areas does your project target?

HPC core technologies and architectures, Programming methodologies, environments, languages and tools APIs and system software for future extreme scale systems ,New mathematical and algorithmic approaches Enabling Applications for future extreme scalesystems

Please list the main achievements of your project in one sentence

CRESTA has created a tailored software collection for future Exascale platforms. Designed to allow applications to exploit future Exascale platforms effectively, this collection also shows significant benefit on current petascale resources, producing a set of applications able to exploit today's largest systems and tomorrow's Exascale platforms, through the use of a co-design process. These applications have demonstrated clear socio-economic benefits from the work on CRESTA and has also produced a significant research portfolio of novel algorithms and techniques for Exascale systems, which have been utilized within the project's software and applications.

Project profile

The Collaborative Research into Exascale Systemware, Tools and Applications (CRESTA) project develops techniques and solutions for some of the most difficult technical challenges that computing at the exascale can present. The project has two integrated strands: one focused on enabling a key set of applications for exascale, the other focused on building and exploring appropriate systemware for exascale platforms. All of the applications had demonstrated a need for exascale performance with associated scientific challenges of global significancethat cannot be solved on current petascale systems, but require exascale performance.

At the heart of the project is the co-design process, with the co-design applications providing guidance and feedback to the exascale software development process, and integrating and benefitting from this development. CRESTA employs incremental and disruptive approaches throughout – sometimes following both paths for a particular problem to compare and contrast the challenges associated with the approach.

ComPat - Computing Patterns for High Performance Multiscale Computing

Website address	http://www.compat-project.eu
Contact name	Prof. Dr. Alfons G. Hoekstra
Contact email	a.g.hoekstra@uva.nl
Contact telephone	N/A

Partners involved: University of Amsterdam, University Leiden, University College London, The Hartree Centre /STFC, Poznan Supercomputing and Networking Centre, Allinea Software, Leibniz Supercomputing Centre, CBK Sci Con Limited, Max-Planck-Institute für Plasmaphysik, ITMO University.

Which areas does your project target?

New mathematical and algorithmic approaches

Project profile

Multiscale phenomena are ubiquitous and they are the key to understanding the complexity of our world. Despite the significant progress achieved through computer simulations over the last decades, we are still limited in our capability to accurately and reliably simulate hierarchies of interacting multiscale physical processes that span a wide range of time and length scales, thus quickly reaching the limits of contemporary high performance computing at the tera- and petascale. Exascale supercomputers promise to lift this limitation, and in this project we will develop multiscale computing algorithms capable of producing high-fidelity scientific results and scalable to exascale computing systems. Our main objective is to develop generic and reusable High Performance Multiscale Computing algorithms that will address the exascale challenges posed by heterogeneous architectures and will enable us to run multiscale applications with extreme data requirements while achieving scalability, robustness, resiliency, and energy efficiency.

Our approach is based on generic multiscale computing patterns that allow us to implement customized algorithms to optimise load balancing, data handling, fault tolerance and energy consumption under generic exascale application scenarios. We will realise an experimental execution environment on our pan-European facility, which will be used to measure performance characteristics and develop models that can provide reliable performance predictions for emerging and future exascale architectures. The viability of our approach will be demonstrated by implementing nine grand challenge applications which are exascale-ready and pave the road to unprecedented scientific discoveries. Our ambition is to establish new standards for multiscale computing at exascale, and provision a robust and reliable software technology stack that empowers multiscale modellers to transform computer simulations into predictive science.

DEEP-(Dynamical Exascale Entry Platform) and Deeper (Dynamical Exascale Entry Platform - Extended Reach)

Website address	http://www.deep- project.eu http://www.deep-er.eu
Contact name	Dr. Estela Suarez
Contact email	pmt@deep- project.eu pmt@deep-er.eu
Contact telephone	N/A

Partners involved: Jülich Supercomputing Centre, Intel, Eurotech, Barcelona Supercomputing Center, Leibniz Supercomputing Centre, University Heidelberg (via spin-off EXTOLL), ParTec, Seagate, Fraunhofer, ITWM, University of Regensburg KU Leuven, École polytechnique fédérale de Lausanne, The Cyprus Institute, German Research School for Simulation Sciences, CERFACS, CINECA, CGG, Astron, Inria, Mellanox

Which areas does your project target?

HPC core technologies and architecture
Programming methodologies, environments, languages and tools
APIs and system software for future extreme scale systems

Please list the main achievements of your project in one sentence

The DEEP and DEEP-ER consortium has developed a novel, Exascale-enabling supercomputing architecture with a matching software stack and a set of optimized grand-challenge simulation applications.

Project profile

The DEEP project and its follow-up project DEEP-ER present an innovative solution for next generation supercomputer addressing various Exascale challenges by following a stringent Co-Design approach. The consortium has developed a novel, Exascale-enabling supercomputing architecture with a matching software stack and a set of optimized grand-challenge simulation applications.

DEEP takes the concept of compute acceleration to a new level: instead of attaching accelerator cards to cluster nodes, DEEP has built a cluster of accelerators, called Booster, to complement a conventional HPC system and increase its compute performance. Accompanied by a software stack focused on meeting Exascale requirements - comprising adapted programming models, libraries and performance tools - the architecture enables unprecedented scalability. The cluster-level heterogeneity of the system attenuates the consequences of Amdahl's law allowing users to run applications with kernels of high scalability alongside kernels of low scalability concurrently on different sides of the system.

DEEP-ER advances the Cluster-Booster architecture developed in DEEP from a hardware point of view in terms of processor technology, network interconnect, and storage. On the software side DEEP-ER focuses on two central research topics: highly scalable parallel I/O and resiliency.

Both DEEP and DEEP-ER validate their concepts on the prototype systems built within the projects. The DEEP prototype system with a peak performance of 500 TFlop/s is already up and running at Jülich Supercomputing Centre.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

For DEEP and DEEP-ER it would be very valuable to get feedback from a diverse group of application scientists on how to use and to benefit from the Cluster-Booster architecture.

What future partners are you looking to collaborate with?

Application developers / domain scientists

ECOSCALE - Energy-efficient Heterogeneous Computing at exascale

Website address http://www.ecoscale.eu/

Contact name Iakovos Mavroidis

Contact email iakovosmavro@gmail.com

Contact telephone +30 2810254899

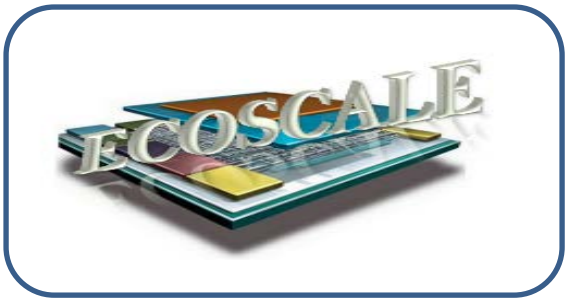

Partners involved: Telecommunication Systems Institute, Greece, Queen's University Belfast, United Kingdom, STMicroelectronics, France, Acciona Infraestructuras S.A., Spain , University of Manchester, United Kingdom, Politecnico di Torino, Italy, Chalmers University of Technology, Sweden, Synelixis, Greece

Which areas does your project target?

HPC core technologies and architectures, Programming methodologies, environments, languages and tools

Project profile

In order to reach exascale performance, current HPC systems need to be improved. Simple hardware scaling is not a feasible solution due to the increasing utility costs and power consumption limitations. Apart from improvements in implementation technology, what is needed is to refine the HPC application development flow as well as the system architecture of future HPC systems.

ECOSCALE tackles these challenges by proposing a scalable programming environment and architecture, aiming to substantially reduce energy consumption as well as data traffic and latency. ECOSCALE introduces a novel heterogeneous energy-efficient hierarchical architecture, as well as a hybrid many-core+OpenCL programming environment and runtime system. The ECOSCALE approach is hierarchical and is expected to scale well by partitioning the physical system into multiple independent Workers (i.e. compute nodes). Workers are interconnected in a tree-like fashion and define a contiguous global address space that can be viewed either as a set of Partitioned Global Address Space (PGAS) partitions, or as a set of nodes hierarchically interconnected via an MPI protocol.

To further increase energy efficiency, as well as to provide resilience, the Workers employ reconfigurable accelerators mapped into the virtual address space utilizing a dual stage System Memory Management Unit with coherent memory access. The architecture supports shared partitioned reconfigurable resources accessed by any Worker in a PGAS partition, as well as automated hardware synthesis of these resources from an OpenCL-based programming model.

EPiGRAM - Exascale Programming Models

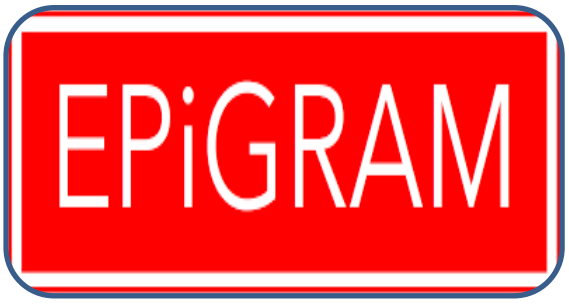

Website address http://www.epigram-project.eu/

Contact name Stefano Markidis

Contact email markidis@kth.se

Contact telephone +46 87907896

Partners involved: KTH Royal Institute of Technology, Technische Universität Wien, EPCC - University of Edinburgh, Cray UK, Fraunhofer

Which areas does your project target?

Programming methodologies, environments, languages and tools

Please list the main achievements of your project in one sentence

EPiGRAM as introduced and implemented new concepts in MPI, such as MPI streams, persistent collectives, endpoints and new derived data-type to improve MPI scalability in execution time and memory consumption. The effectiveness of these new concepts has been proven in the EPiGRAM applications, iPIC3D and Nek5000. It has improved the interoperability between MPI and GPI-2 by new developments in GPI, and designed and implemented a research PGAS-based MPI, called EMPI4Re.

Project profile

We are preparing Message Passing and PGAS programming models for exascale systems by addressing their main current limitations. By providing prototype implementations for both MP and PGAS concepts we will contribute to advancement in programming models forextreme-scale computing.

ESCAPE - Energy-efficient Scalable Algorithms for weather Prediction at Exascale

Website address http://www.hpc-escape.eu

Contact name Dr Peter Bauer

Contact email peter.bauer@ecmwf.int

Contact telephone +44 1189499080

Partners involved: European Centre for Medium-Range Weather Forecasts, Meteo-France, Institut Royal Meteorologique de Belgique, Danmarks Meteorologiske Institut, Eidgenoessisches Departement des Innern Deutscher Wetterdienst, Loughborough University National University of Ireland, Galway, Instytut Chemii, Bioorganicznej Polskiej Akademii Nauk, Bull SAS, AGEIA Technologies Switzerland AG, Optalysys

Which areas does your project target?

Programming methodologies, environments, languages and tools, New mathematical

and algorithmic approaches

Project profile

ESCAPE is funded under FET-HPC and is a Research and Innovation project. ESCAPE will develop world-class, extreme-scale computing capabilities for European operational numerical weather prediction (NWP) and future climate models. ESCAPE aims at developing next generation IFS numerical building blocks and compute intensive algorithms, implement and compute/energy efficiency diagnostics; identify new approaches and implementation on novel architectures and perform testing in operational configurations.

EXA2CT - Exascale Algorithms and Advanced Computational Techniques

Website address http://www.exa2ct.eu

Contact name Tom Vander Aa

Contact email tom.vanderaa@imec.be

Contact telephone +32 16 28 80 53

Partners involved: Interuniversitair Micro-Electronica Centrum Vzw Imec Belgium, Universiteit Antwerpen Ua Belgium, Universita Della Svizzera Italiana Usi Switzerland, Vysoka Skola Banska - Technicka Univerzita Ostrava Vsb-Tuo Czech Republic, Institut National De Recherche En Informatique Et En Automatique Inria France, Universite De Versailles Saint-Quentin-En-Yvelines. Uvsq France, T-Systems Solutions For Research Gmbh Ts-Sfr Germany, Fraunhofer-Gesellschaft Zur Foerderung Der Angewandten Forschung E.V, Intel Corporation Sas France, Numerical Algorithms Group Ltd Nag United Kingdom

Which areas does your project target?

Programming methodologies, environments, languages and tools APIs and system software for future extreme scale systems, new mathematical and algorithmic approaches

Please list the main achievements of your project in one sentence

Our scalable pipelined solvers are now available in the PETSc library and we have shown the GASPI programming model can genuinely outperform MPI, whilst many of our achievements have been ported to the industrial applications from our scientific and industrial board (SIB).

Project profile

The EXA2CT project brings together experts at the cutting edge of the development of solvers, related algorithmic techniques, and HPC software architects for programming models and communication. We will produce modular open source proto-applications that demonstrate the algorithms and programming techniques developed in the project, to help boot-strap the creation of genuine exascale codes.

ExaFLOW Enabling Exascale Fluid Dynamics Simulations

Website address http://exaflow-project.eu

Contact name Erwin Laure,
KTH Royal Institute of Technology

Contact email erwinl@pdc.kth.se

Contact telephone N/A

Partners involved: The Royal Institute of Technology,KTH, Sweden; -PDC Center for High Performance Computing, Department of Mechanics Imperial College, IC, United Kingdom; University of Southampton, SOTON, United Kingdom; The University of Edinburgh, EPCC, United Kingdom; Universitaet Stuttgart, USTUTT, Germany; Institut für Aero und Gasdynamik, IAG; High Performance Computing Center Stuttgart; HLRS; Ecole Polytechnique Federale de Lausanne; EPFL, Switzerland; McLaren Racing Ltd, McLaren,United Kingdom; Automotive Simulation Center Stuttgart ,ASCS ,Germany

Which areas does your project target?

New mathematical and algorithmic approaches

Please list the main achievements of your project in one sentence

Algorithmic improvements for increased performance, scalability and exascale readiness of major high order, open source computational fluid dynamics codes.

Project profile

We are surrounded by moving fluids (gases and liquids), be it during breathing or the blood flowing in arteries; the flow around cars, ships, and airplanes; the changes in cloud formations or the plankton transport in oceans; even the formation of stars and galaxies are closely modelled as phenomena in fluid dynamics. Fluid Dynamics (FD) simulations provide a powerful tool for the analysis of such fluid flows and are an essential element of many industrial and academic problems.

The complexities and nature of fluid flows, often combined with problems set in open domains, implies that the resources needed to computationally model problems of industrial and academic relevance is virtually unbounded. FD simulations therefore are a natural driver for exascale computing and have the potential for substantial societal impact, like reduced energy consumption, alternative sources of energy, improved health care, and improved climate models.

The main goal of this project is to address algorithmic challenges to enable the use of accurate simulation models in exascale environments. Driven by problems of practical engineering interest we focus on important simulation aspects including:

- error control and adaptive mesh refinement in complex computational domains,
- resilience and fault tolerance in complex simulations
- heterogeneous modeling

• evaluation of energy efficiency in solver design

• parallel input/output and in-situ compression for extreme data.

The algorithms developed by the project will be prototyped in major open-source simulation packages in a co-design fashion, exploiting software engineering techniques for exascale. We are building directly on the results of previous exascale projects (CRESTA, EPiGRAM, etc.) and will exploit advanced and novel parallelism features required for emerging exascale architectures. The results will be validated in a number of pilot applications of concrete practical importance in close collaboration with industrial partners.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

Extreme scale computational fluid dynamics simulations; Algorithmics improvements towards exascale

What future partners are you looking to collaborate with?

Fluid dynamics communities of the codes used by the project; HPC experts for development of exascale software capabilities; Industrial users (energy, aerospace, automotive)

ExaHyPE - An Exascale Hyperbolic PDE Engine

Website address http://www.exahype.eu

Contact name Michael Bader, Technical University of Munich

Contact email bader@in.tum.de

Contact telephone +49 89 35831 7810 (contact via email strongly preferred)

Partners involved: Technische Universität München (Coordinator) - Prof. Dr. Michael Bader (High Performance Computing); Università degli Studi di Trento - Prof. Dr.-Ing. Michael Dumbser (Numerical Analysis) Durham University Dr. Tobias Weinzierl (High Performance Computing); Frankfurt Institute for Advanced Studies - Prof. Dr. Luciano Rezzolla (Theoretical Astrophysics); Ludwig-Maximilians-Universität München - Dr. Alice-Agnes Gabriel, Prof. Dr. Heiner Igel (Computational Seismology); RSC Technologies, Moscow - Alexander Moskovsky; Bavarian Research Alliance GmbH, - Dipl.-Ing. Robert Iberl, Teresa Kindermann

Which areas does your project target?

New mathematical and algorithmic approaches

Please list the main achievements of your project in one sentence

We have implemented the ADER-DG numerical approach on adaptive spacetree meshes within a first prototype of our engine for solving hyperbolic systems of PDE with high performance on future exascale supercomputers.

Project profile

ExaHyPE develops a novel Exascale-ready engine to simulate large-scale problems that maybe expressed via hyperbolic systems of conservation laws. It relies on latest and further developments of the ADER-DG (Arbitrary high-order derviative discontinuous Galerkin) numerical scheme, dynamically adaptive Cartesian meshes, agile load balancing and hardware- oriented optimzation of the respective algorithms and implementations. While a generic PDE engine is targeted, ExaHyPE will in its project period focus on two well-defined grand challenge scenarios from computational geo- and astrophysics: the first scenario is the simulation of rotating (and collapsing) binary neutron stars, which are primary suspects for causing phenomena such as gamma ray bursts or gravitational waves; the second scenario considers regional earthquake simulation with a special focus on dynamic rupture processes, which may lead to better understanding of earthquake processes and thus to improved seismic hazard assessment for critical infrastructure, e.g. On the methodological side, ExaHyPE seeks to demonstrate the necessity for high-order approximation in space and time to optimize time- and energy-to-solution, and the excellent suitability of these methods for future supercomputing platforms.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

ExaHyPE welcomes potential users of the hyperbolic PDE engine to discuss and prototypically implement possible applications, but also seeks collaboration with providers of emerging supercomputing technology, including actual supercomputers in the same way as programming environments.

What future partners are you looking to collaborate with?

ExaHyPE seeks collaboration with:

- providers of supercomputing hardware (existing petascale machines, prototypes for emerging/future exascale technology, etc.)
- potentials users of the hyperbolic PDE engine (esp. computational astrophysics and seismology, but also other fields)
- developers of exascale-aware programming environments (offering resiliency, energy efficiency, etc.)
- providers of high-performance visualisation solutions, esp. for high-order methods

ExaNeSt - European Exascale System Interconnect and Storage

Website address http://www.exanest.eu

Contact name Manolis G.H. Katevenis (FORTH), Coordinator

Contact email nickpap@ics.forth.gr

Contact telephone +30 2811.39.1078

Partners involved: Foundation for Research and Technology – Hellas (FORTH), Iceotope Technologies Ltd, Allinea Software Ltd, EnginSoft S.p.A., eXact lab srl, MonetDB Solutions (MDBS), Virtual Open Systems (VOSYS), Istituto Nazionale di Astrofisica (INAF), National Institute for Nuclear Physics (INFN), The University of Manchester (UoM), Universitat Politècnica de València (UPV), Fraunhofer-Gesellschaft Zur Foerderung Der Angewandten Forschung E.V (Fraunhofer)

Which areas does your project target?

HPC core technologies and architectures, Programming methodologies, environments, languages and tools, APIs and system software for future extreme scale systems

Project profile

ExaNeSt develops and prototypes solutions for Interconnection Networks, Storage, and Cooling, as these have to evolve in order for the production of exascale-level supercomputers to become feasible. We tune real HPC Applications, and we use them to evaluate our solutions.

Interconnection Network: exascale performance can only be reached by interconnecting millions of computing cores, their (volatile) memories and (non-volatile) storage, special- purpose accelerator hardware, and their input/output (I/O) devices, in a way such that all of them can cooperate tightly and effectively in solving one huge problem in a reasonable amount of time. This amounts to huge challenge for the network that implements this interconnection and its interface to the hardware and software components of the entire system: it has to be fast, resilient, and low-cost, both in term of cost-to-build and energy-to- operate.

We develop and prototype innovative hardware and software for such networks to become tightly integrated with the system components, to become faster, to offer better quality-of- service (QoS) – especially congestion mitigation, to be resilient to failures, and to consume less energy.

Storage: traditional supercomputers used a large number of magnetic disks for storing non- volatile and permanent checkpoints and data, where these disks appeared as I/O devices to the computing cores. Storage technologies now change to flash and non-volatile memories (NVM), featuring dramatically lower latencies; interconnection and software architecture have to also change, in order to take advantage of such much faster access times.

We develop and prototype a distributed storage system where NVM's are local to the compute cores hence fast to access at low energy cost, yet the aggregate NVM's in the entire system form a unified storage.

Cooling: communicating at low delay and energy cost requires physical proximity, i.e. packing thousands of

cores and their components into a blade board and packing about a hundred blades into a rack (which also economizes on installation floor area). The by-product, unfortunately, is a large heat density to be removed.

We develop and prototype innovative Packaging and Cooling technology, based on total immersion in a sophisticated, non-conductive, engineered coolant fluid that allows the highest possible packing density while maintaining reliability.

Applications: we evaluate all these technologies using real High-Performance Computing (HPC) and Big Data Applications –from HPC simulations to Business Intelligence support– running on a real prototype at the scale of many hundred nodes containing thousands of compute cores.

Furthermore, we tune our firmware, the systems software, libraries, and such applications so that they take the best possible advantage of our novel communication and storage architecture: we support task-to-data software locality models, to ensure minimum data communication energy overheads and property maintenance in databases; and we provide a platform management scheme for big-data I/O to our resilient, unified distributed storage compute architecture.

ExaNoDe - European Exascale Processor Memory Node Design

Website address	http://exanode.eu/
Contact name	Christian Pinto, Virtual Open Systems
Contact email	c.pinto@virtualopensystems.com
Contact telephone	N/A

Partners involved: Cea, Arm, Eth Zürich, Forth, Fraunhofer, , , Scapos, University Of Manchester, Bull, Vosys, Bsc, Juelich, Kalray, Cnrs

Which areas does your project target?

HPC core technologies and architectures

Please list the main achievements of your project in one sentence

ExaNoDe designs a high energy efficient and highly integrated heterogeneous compute node enabling Pre-Exascale level computing, mixing low-power processors, co-processors and advanced hardware integration technologies with the novel UNIMEM Global Address Space memory system.

Project profile

ExaNoDe is a collaborative European project within the "Horizon 2020 Framework Programme", that investigates and develop a highly integrated, high-performance, heterogeneous System-on-a-Chip (SoC) aimed towards exascale computing. It is addressing these important challenges through the coordinated application of several innovative solutions recently deployed in HPC: ARM-v8 low-power processors for energy efficiency, 3D interposer integration for compute density and an advanced memory scheme for exabyte level capacities. The ExaNoDe SoC will embed multiple silicon "chiplets", stacked on an active silicon interposer in order to build an innovative 3D-Integrated-Circuit (3D-IC). A full software stack allowing for multi-node capability will be developed within the project. The project will deliver a reference hardware that will enable the deployment of multiple 3D-IC System-on-Chips and the evaluation, tuning and analysis of HPC mini-apps along with the associated software stack.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

By aiming at creating an extremely dense and high-energy efficient compute node, ExaNoDe could create links with either other research projects or companies active in hardware design and manufacturing in both USA, Japan and other countries.

ExCAPE - Exascale Compound Activity Prediction Engine

Website address http://www.excape-h2020.eu

Contact name Tom Ashby, Imec vzw

Contact email ashby@imec.be

Contact telephone +32-16-28-1251

Partners involved: Imec vzw, (Belgium); JANSSEN CILAG SA, (Spain); IT4I, (Czech Republic); ASRRAZENECA AB (Sweden);UNIVERSITAT LINZ (Austria); Aalto University (Finland) Intel (Belgium); IDEAconsult (Bulgaria) Royal Holloway (UK)

Which areas does your project target?

Programming methodologies, environments, languages and tools, New mathematical and algorithmic approaches

Please list the main achievements of your project in one sentence

The ExCAPE project has prepared machine learning data sets for the pharmaceutical industry, made them accessible on the Czech supercomputing cluster, and applied machine learning techniques using the cluster.

Project profile

Our project focuses on developing better machine learning algorithms for predicting biologicalactivity of drugs, by taking advantage of large amounts of computation. A secondary concern is making easier the deployment of new machine learning algorithms on supercomputers, whilst keeping the implementations efficient so that the models benefit from the computation available. The project brings together machine learning experts to develop the new algorithms, scale-up experts to make prototype implementations and pharmaceutical industry partners to test the new algorithms and implementations.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

ExCAPE would like international collaboration in the areas of large (10M to 100M training point) multi-task, sparse-evidence and sparse feature vector machine learning algorithms that can be run on supercomputer hardware, especially in the area of compound activity prediction for the pharmaceutical industry.

What future partners are you looking to collaborate with?

1) Companies from the world-wide pharmaceutical sector. We aim to give targeted information to companies that join an industrial board that we're currently putting together.
2) Other groups doing machine learning at scale, both industrial and academic. We are especially interested in multi-task solutions and learning techniques that can make good use of large amounts of computing power to improve predictions, including deep learning, Bayesian techniques and confidence estimation.
3) Companies from different industrial sectors that have other large scale machine learning problems to solve.
4) Groups that are providing tools to make it easier to scale machine learning on HPC hardware.

EXTRA - Exploiting exascale Technology with Reconfigurable Architectures

Website address http://www.extrahpc.eu

Contact name Dirk Stroobandt

Contact email Dirk.Stroobandt@UGent.be

Contact telephone +32 9 264 34 01

Partners involved: Ghent University Belgium, Telecommunications Systems Institute, Greece Imperial College London UK, Politecnico di Milano, Italy, University of Amsterdam, Netherlands, Rurh-University Bochum, Germany, Maxeler, UK, Synelixis, Greece, University of Cambridge, UK

Which areas does your project target?

HPC core technologies and architectures, Programming methodologies, environments, languages and tools, Reconfigurable Computing for HPC

Project profile

EXTRA focuses on the fundamental building blocks for run-time reconfigurable exascale HPC systems: new reconfigurable architectures with very low reconfiguration overhead, new tools that truly take reconfiguration as a design concept, and applications that are tuned to maximally exploit run-time reconfiguration techniques. Our goal is to provide the European platform for run-time reconfiguration to maintain Europe's competitive edge and leadership in run-time reconfigurable computing.

Green Flash - energy efficient high performance computing for real-time science

Website address	http://green-flash.lesia.obspm.fr/
Contact name	Damien Gratadour
Contact email	damien.gratadour@obspm.fr
Contact telephone	+33 145077757

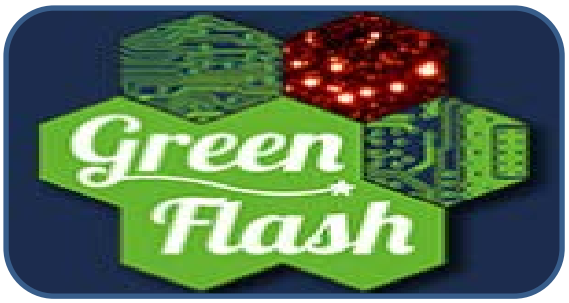

Partners involved: Observatoire de Paris, University of Durham, Microgate, PLDA

Which areas does your project target?

HPC core technologies and architectures, Programming methodologies, environments, languages and tools

Project profile

The main goal of Green Flash is to design and build a prototype for a Real-Time Controller (RTC) targeting the European Extremely Large Telescope (E-ELT) Adaptive Optics (AO) instrumentation. The E-ELT is a 39m diameter telescope to see first light in the early 2020s.To build this critical component of the telescope operations, the astronomical community is facing technical challenges, emerging from the combination of high data transfer bandwidth, low latency and high throughput requirements, similar to the identified critical barriers on the road to Exascale. With Green Flash, we will propose technical solutions, assess these enabling technologies through prototyping and assemble a full scale demonstrator to be validated with a simulator and tested on sky. With this R&D program we aim at feeding the E-ELT AO systems preliminary design studies, led by the selected first-light instruments consortia, with technological validations supporting the designs of their RTC modules. Our strategy is based on a strong interaction between academic and industrial partners. Components specifications and system requirements are derived from the AO application. Industrial partners lead the development of enabling technologies aiming at innovative tailored solutions with potential wide application range. The academic partners provide the missing links in the ecosystem, targeting their application with mainstream solutions. This increases both the value and market opportunities of the developed products. A prototype harboring 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.

INTERWinE - Programming Model Interoperability towards Exascale

Website address	http://www.intertwine-project.eu/
Contact name	Dr George Beckett, EPCC (University of Edinburgh)
Contact email	info@intertwine-project.eu
Contact telephone	+44 (0) 131 651 3577
Partners involved:	EPCC (University of Edinburgh), UK; BSC, Spain; KTH, Sweden; Fraunhofer ITWM, Germany; DLR, Germany; T-Systems SfR, Germany; Universitat Jaume I, Spain; Inria, France; University of Manchester, UK

Which areas does your project target?

Programming methodologies, environments, languages and tools

Please list the main achievements of your project in one sentence

INTERTWinE has established a Europe-wide programme of advanced training on parallel and interoperable programming for extreme scale, supported by Best Practice Guides that distill the expertise of leading European computational scientists.

Project profile

The INTERTWinE project addresses the problem of programming model design and implementation for the Exascale. 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 interfaces (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 main challenges lie in interoperability between APIs. It is this interoperability, both at the specification level and at the implementation level, which we are addressing and making substantial progress towards. We are focusing on six key programming APIs: MPI, GASPI, OpenMP, OmpSs, StarPU and PaRSEC, each of which is represented by a project partner with extensive experience in API design and implementation. Our interoperability requirements, and evaluation of implementations, are driven by a set of kernels and applications, each of which is represented by a project partner with a major role in their development. The project employs a co-design cycle, in which advances in API design and implementation are fed into the applications and kernels, which in turn drives new requirements and hence further advances.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

INTERTWinE has established, and is addressing, the need for a model-agnostic interoperability forum at which representatives for key parallel APIs and scientific libraries can work together to solve the exascale programming problem.

What future partners are you looking to collaborate with?

Stakeholders in the development of strategically important parallel programming models and environments in the international community

MANGO - Exploring Many-core Architectures for Next-Generation HPC systems

Website address	www.mango-project.eu/
Contact name	José Flich Cardo
Contact email	jflich@disca.upv.es
Contact telephone	N/A

Partners involved: The Swiss Federal Institute of Technology Lausanne (EPFL), PRO DESIGN, Germany, Politecnico di Milano (Milan, Italy), The University of Zagreb, Croatia, The Centro Regionale Information Communication Technology scrl (CeRICT), Italy, EATON, France, Universitat Politecnica de Valencia, Spain, Thales Communications & Security S.A.S., Royal Philips Electronics

Project profile

MANGO targets to achieve extreme resource efficiency in future QoS-sensitive HPC through ambitious cross-boundary architecture exploration for performance/power/predictability (PPP) based on the definition of new-generation high-performance, power-efficient, heterogeneous architectures with native mechanisms for isolation and quality-of-service, and an innovative two-phase passive cooling system. Its disruptive approach will involve many interrelated mechanisms at various architectural levels, including heterogeneous computing cores, memory architectures, interconnects, run-time resource management, power monitoring and cooling, to the programming models. The system architecture will be inherently heterogeneous as an enabler for efficiency and application-based customization, where general-purpose compute nodes (GN) are intertwined with heterogeneous acceleration nodes (HN), linked by an across-boundary homogeneous interconnect. It will provide guarantees for predictability, bandwidth and latency for the whole HN node infrastructure, allowing dynamic adaptation to applications. MANGO will develop a toolset for PPP and explore holistic pro-active thermal and power management for energy optimization including chip, board and rack cooling levels, creating a hitherto inexistent link between HW and SW effects at all layers. Project will build an effective large-scale emulation platform. The architecture will be validated through noticeable examples of application with QoS and high-performance requirements.

Ultimately, the combined interplay of the multi-level innovative solutions brought by MANGO will result in a new positioning in the PPP space, ensuring sustainable performance as high as 100 PFLOPS for the realistic levels of power consumption (<15MWatt) delivered to QoS-sensitive applications in large-scale capacity computing scenarios providing essential building blocks at the architectural level enabling the full realization of the ETP4HPC strategic research agenda.

Mont-Blanc - European Approach Towards Energy Efficient HPC

Website address http://www.montblanc-project.eu/

Contact name Filippo Mantovani (BSC)
Etienne Walter (Atos)

Contact email Ufilippo.mantovani@bsc.es
Uetienne.walter@atos.net

Contact telephone +34 934137716,
+33 130807496

Partners involved: ARM, Allinea, AVL. BSC, Bull/Atos, CEA, CRNS (Centre national de la Recherche Scientifique) / LIRMM, ETH Zurich, GENCI, HLRS, Inria, Juelich, LRZ, Universidad de Cantabria (Santander,) University of Bristol, University of Graz, Université de Versailles Saint Quentin

Which areas does your project target?

HPC core technologies and architectures, Programming methodologies, environments, languages and tools

Please list the main achievements of your project in one sentence

Among the most relevant results of the Mont-Blanc project are:
a) the deployment of the world's first ARM-based HPC cluster, a prototype that largely contributed to validate the concept of using ARM technology for HPC, and which, together with various ARM-based prototypes, allowed the project to substantially contribute to the development of the HPC system software stack for ARM-based platforms
b) a strong involvement in the realization of a complete software ecosystem allowing the execution of production HPC applications on ARM based clusters

Project profile

The Mont-Blanc and Mont-Blanc2 projects are FP7 EU projects sharing the vision of developing a European Exascale approach leveraging commodity power- and cost-efficientembedded technologies.

In the frame of these two projects the first large ARM-based prototype dedicated to HPC has been deployed. The 1000+ computational nodes are operational since May 2015 at Barcelona Supercomputing Center and open to the partner access. The projects have also developed a complete HPC system software stack for ARM-based compute nodes tested on prototype and commercial platforms.

The rapid progress of Mont-Blanc towards defining a scalable and efficient pre-Exascale platform has revealed a number of challenges and opportunities to broaden the scope of investigations and developments. Particularly, the growing interest of the HPC community in accessing the Mont-Blanc platform calls for increased efforts to setup a production-ready environment. Within the Mont-Blanc 2 project, therefore effort is focused on:

- improvement of the system software stack, with emphasis on programmer tools (debugger, performance analysis)
- research in system resiliency (from applications to architecture support)
- test and support of ARM 64 bitarchitecture
- porting and testing of new industrial and academic applications

Within the H2020 Programme, the brand new Mont-Blanc 3 project is continuing on the way paved by the previous two projects, targeting the development of a scalable and balanced high-end compute node for future HPC systems, based on low-power embedded technology. It will adopt a co-design approach to make sure that the hardware and system innovations are readily translated into benefits for HPC applications. This approach integrates architectural aspects together with simulation efforts to feed and validate the architectural studies, in parallel with work on system software ecosystem and applications. Effort is focused on the following targets:

- Defining the architecture of an Exascale-class compute node based on the ARM architecture, and capable of being manufactured at industrial scale;
- Assessing the available options for maximum compute efficiency;
- Developing the matching software ecosystem to pave the way for market acceptance of ARM solutions.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

As the project offers a fairly complete ARM-based ecosystem of hardware/software infrastructure for HPC, international collaborations can involve tests of production codes on the Mont-Blanc platforms, development and test of tools for power monitoring, joint resiliency studies. Also, collaborations can be at dissemination level, with the involvement of the project in events, workshops, training or conferences.

What future partners are you looking to collaborate with?

Anybody with the interest in promoting and improving the ARM ecosystem with special focus in HPC and scientific computing.

NESUS - Network for Sustainable Ultrascale Computing

Website address www.nesus.eu

Contact name Prof. Jesus Carretero,
University Carlos III of Madrid

Contact email jesus.carretero@uc3m.es

Contact telephone +34 916249458

Partners involved: http://www.nesus.eu/institutions

Please list the main achievements of your project in one sentence

Development of a major research network with more than 80 institutions belonging to 45 countries.

Project profile

NESUS is the acronym of the COST Action IC1305. The Action is a research network composed by almost 80 institutions from 45 countries (most EU, Russia, USA, India, Australia and Colombia) created to research on the evolution of today´s large-scale system, not only in HPC world, but also in HTC, and the challenges arising that could stop adoption of those systems in the future.

The project has created a strong network of researchers that are cooperating in shared publications, workshops and applications related to the topics of the action. A first research roadmap will be issued by the end of 2016. It also promotes a winter school every year to share experiences and to train early researchers.

NESUS strongly promotes international cooperation of the Action´s member, as cooperation is one of its goals. The Action itself is a cooperation tool, but it could be enhanced by means of researchers exchanges, development of joint applications, participation in the winter school, and organization of scientific events related to the Action.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

Programming environment, Data-intensive computing, HPC and Energy HPC and Big Data.

What future partners are you looking to collaborate with?

NEXTGenIO - Next Generation I/O for Exascale

Website address www.nextgenio.eu

Contact name Dr. Michele Weiland, EPCC

Contact email m.weiland@epcc.ed.ac.uk

Contact telephone +44 (0)131 651 3580

Partners involved: EPCC-UK, Allinea-UK, ECMWF-UK, Intel-Belgium, BSC-Spain, Fujitsu-Germany, Technische Universität Dresden, Arctur Universe-Slovenia.

Which areas does your project target?

HPC core technologies and architectures, Programming methodologies, environments, languages and tools APIs and system software for future extreme scale systems

Please list the main achievements of your project in one sentence

To date, our main achievements are the definition of a new HPC architecture (hardware and software) that can make use of NVDIMMs; the project has now entered its implementation phase.

Project profile

The NEXTGenIO project addresses a key challenge not only for Exascale, but also for HPC and data intensive computing in general: the challenge of I/O performance. As core-counts have massively increased over the past few years, the performance of I/O subsystems have struggled to keep up with computational performance and have become a key bottleneck on today's largest systems. NEXTGenIO will develop a prototype computing platform that uses on-node non-volatile memory, bridging the latency gap between DRAM and disk, thus removing this bottleneck. In addition to the hardware that will be built as part of the project, NEXTGenIO will develop the software stack (from OS and runtime support to programming models and tools) that goes hand-in-hand with this new hardware architecture. Two particular focal points are a data and power aware job scheduling system, as well as an I/O workload and workflow simulator that will allow us to stress-test our hardware and software developments. We believe that the new platform that is being developed in NEXTGenIO will be capable of delivering transformational performance across high performance and data intensive computing.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

Partners that have IO requirements that currently prevent them from using traditional HPC systems.

What future partners are you looking to collaborate with?

Partners that have IO requirements that currently prevent them from using traditional HPC systems.

NLAFET - Parallel Numerical Linear Algebra for Future Extreme-Scale Systems

Website address http://www.nlafet.eu/

Contact name Prof. Bo Kågström, Umeå University, Sweden

Contact email bokg@cs.umu.se info@nlafet.eu

Contact telephone +46 -90 -786 5419

Partners involved: Umeå University, Sweden (coordinator of NLAFET), The University of Manchester, UK, Institute National de Recherche en Informatique et en Automatique, France, Science and Technology Facilities Council, UK

Which areas does your project target?

Programming methodologies, environments, languages and tools, New mathematical and algorithmic approaches
See focus topics in Project profile below.

Project profile

The aim is to enable a radical improvement in the performance and scalability of a wide range of real-world applications relying on linear algebra software, by developing novel architecture- aware algorithms and software libraries, and the supporting runtime capabilities to achieve scalable performance and resilience on heterogeneous architectures. The focus is on acritical set of fundamental linear algebra operations including direct and iterative solvers for dense and sparse linear systems of equations and eigenvalue problems. The main research objectives of NLAFET are:

(i) development of novel algorithms that expose as much parallelism as possible, exploit
(ii) heterogeneity, avoid communication bottlenecks, respond to escalating fault rates, and help meet emerging power constraints;
(iii) exploration of advanced scheduling strategies and runtime systems focusing on the extreme scale and strong scalability in multi/many core and hybrid environments;
(iv) design and evaluation of novel strategies and software support for both offline and online auto-tuning.

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 deliverables also include a sustainable set of methods and tools for cross-cutting issues such as scheduling, auto-tuning, and algorithm-based fault tolerance packaged into open source library modules.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

We already have contact and meetings with HPC HW/SW vendors and one of our work packages are focusing on four challenging applications (materials science, power systems, study of energy solutions, and data analysis in astrophysics).

NUMEXAS - Numerical Methods and Tools for Key Exascale

Website address http://www.numexas.eu

Contact name Cecilia Soriano, Project Manager

Contact email csoriano@cimne.upc.edu

Contact telephone +34 93 401 7440

Partners involved: Cimne (Centre Internacional De Metodes Numerics En Enginyeria) Spain, Csuc (Consorci De Serveis Universitars De Catalunya) Spain, Luh (Gottfried Wilhelm Leibniz Universitaet Hannover) Germany, Ntua (National Technical University Of Athens) Greece, Quantech Atz Sa Spain

Which areas does your project target?

Programming methodologies, environments, languages and tools, New mathematical and algorithmic approaches

Please list the main achievements of your project in one sentence

Is greatly improving the scalability of the linear solver used within Kratos (for the MPI case), with a highly scalable embedded solver and MPI porting of a very complex "commercial grade" problem.

Project profile

The overall aim of NUMEXAS is to develop, implement and demonstrate the next generation of numerical simulation methods to be run under exascale computing architectures. This cannot be done by just scaling currently available codes, but by implementing new algorithms for advanced numerical methods to really exploit the intrinsic capabilities of the future exascale computing infrastructures.

The specific goal of NUMEXAS is the development of numerical methods based on validated models that enable scaling to millions of cores along the complete simulation pipeline. The main outcome of NUMEXAS will be a new set of numerical methods and computer codes that will allow industries, governments and academia to routinely solve multidisciplinary large-scale class problems in applied sciences and engineering with high efficiency and the simplicity of the best nowadays user-friendly computer codes.

In order to achieve the above mentioned goals, improvements are required along the whole simulation pipeline, including parallel pre-processing of analysis data and mesh generation, parallel, scalable, parallel field solvers in fluid, solid mechanics and coupled problems, optimum design parallel solvers considering uncertainties and parallel post-processing of numerical results

READEX - Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing

Website address http://www.readex.eu

Contact name Joseph Schuchart, Wolfgang E. Nagel

Contact email joseph.schuchart@tu-dresden.de wolfgang.nagel@tu-dresden.de

Contact telephone +49 351 463 31673 +49 351 463 35450

Partners involved: Technische Universität Dresden/ZIH, Norges Tekniski-Naturvitenskapelige Universitet, Vysoka Skola Banska – Technicka, Univerzita Ostrava National, University of Ireland/Galway, Intel Corporation SAS, Technische Universität München, Gesellschaft für numerische Simulation mbH

Which areas does your project target?

Programming methodologies, environments, languages and tools

Project profile

The READEX project aims at developing a tools-aided methodology for dynamic auto-tuning of HPC applications to exploit the dynamically changing resource requirements for improved energy-efficiency. It connects technologies from both ends of the computing sprectrum: the methodology will be based on the System Scenario Methodology for dynamic tuning developed in the Embedded Systems domain paired with the technology from the Periscope Tuning Framework (PTF) developed in the FP7 AutoTune project for static auto-tuning in the area of HPC.

READEX is a H2020-FETHPC-2014 project funded under topic b) Programming methodologies, environments, languages and tools and focuses entirely on the development of methodology and required software; no hardware development will be done.

SAGE - Percipient Storage for Exascale Data Centric Computing

Website address www.sagestorage.eu

Contact name Sai Narasimhamurthy, Seagate Systems UK

Contact email info@sagestorage.eu, sai.narasimhamurthy@seagate.com

Contact telephone +44(0)7584080691

Partners involved: Allinea, UK, Atos, France, CEA, France, CCFE(UK Atomic Energy Agency), UK, DFKI, Germany, Diamond Light Source , UK, Juelich (FZJ), Germany, KTH, Sweden, Seagate (Co-ordinator), UK, STFC, UK

Which areas does your project target?

HPC core technologies and architectures, Programming methodologies, environments, languages and tools APIs and system software for future extreme scale systems

Please list the main achievements of your project in one sentence

The project has completed the first phase of requirements gathering activity from the use cases for building the SAGE storage platform and completed the architecture and design of the key Software & Hardware components of SAGE in the first year.

Project profile

Worldwide data volumes are exploding and islands of storage remote from compute will not scale. We will demonstrate the first instance of intelligent data storage, uniting data processing and storage as two sides of the same rich computational model. This will enable sophisticated, intention-aware data processing to be integrated within a storage systemsinfrastructure, combined with the potential for Exabyte scale deployment in futuregenerations of extreme scale HPC systems. Enabling only the salient data to flow in and out of compute nodes, from a sea of devices spanning next generation solid state to low performance disc we enable a vision of a new model of highly efficient and effective HPCand Big Data demonstrated through the SAGE project.

Objectives

- Provide a next-generation multi-tiered object-based data storage system (hardware and enabling software) supporting future generation multi-tier persistent storage media supporting integral computational capability, within a hierarchy.

- Significantly improve overall scientific output through advancements in systemic data access performance and drastically reduced data movements.

- Provides a roadmap of technologies supporting data access for both Exascale/Exabyte and High Performance Data Analytics.

- Provide programming models, access methods and support tools validating their usability, including 'Big-Data' access and analysis methods

- Co-Designing and validating on a smaller representative system with earth sciences, meteorology, clean energy, and physics communities

- Projecting suitability for extreme scaling through simulation based on evaluation results.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

Further/continued collaboration on the codesign of the storage platform (incl. its API) and supporting ecosystem for exascale and beyond by better understanding future extreme computing/data workflows as needed by the international community.

What future partners are you looking to collaborate with?

We are looking to collaborate with big data extreme computing (BDEC / extreme HPDA) HPC workflow based projects or users which, without the SAGE system could be IO limited or unable to scale. These could be in both the HPC or Big data community. We are looking to partner with communities doing work on tools, programming models, runtimes, data analytics, extreme IO etc. in the realm of BDEC. We are also looking to partner with systems developers or integrators needing to balance compute and IO capabilities at extreme scale or providers of next generation storage device technologies.

MaX - Materials design at the eXascale

Website address	http://www.max-centre.eu/
Contact name	Elisa Molinari
Contact email	elisa.molinari@unimore.it
Contact telephone	N/A

Partners involved: (CNR - Modena, SISSA - Trieste, ICN2 - Barcelona, FZ - Jülich, EPFL - Lausanne) and five supercomputing centres (CINECA - Bologna, ETH/CSCS - Zürich/Lugano, FZ - Jülich, KTH - Stockholm , BSC - Barcelona), (ICTP - Trieste), E4 Computer Engineering (Italy), CloudWeavers (UK)

Project profile

Materials are crucial to scientific and technological advances and industrial competitiveness, and to tackle key societal challenges - from energy and environment to health care, information and communications, manufacturing, safety and transportation.

The current accuracy and predictive power of materials' simulations allow a paradigm shift for computational design and discovery, in which massive computing efforts can be launched to identify novel materials with improved properties and performance; behaviour of ever-increasing complexity can be addressed; sharing of data and work-flows accelerates synergies and empowers the science of big-data; and services can be provided in the form of data, codes, expertise, turnkey solutions, and a liquid market of computational resources. Europe has the human resources, track record and infrastructure to be worldwide leader in this field, and we want to create a CoE in materials' modelling, simulations, and design to endow our researchers and innovators with powerful new instruments to address the key scientific, industrial and societal challenges that require novel materials.

This CoE will be a user-focused, thematic effort supporting the needs and the vision of all our core communities: domain scientists, software scientists and vendors, end-users in industry and in academic research, and high-performance computing centres.

The proposal is structured along two core actions: (1) Community codes, their capabilities and reliability; provenance, preservation and sharing of data and work-flows; the ecosystem that integrates capabilities; and hardware support and transition to exascale architectures. (2) Integrating, training, and providing services to our core communities, while developing and implementing a model for sustainability, with the core benefit of propelling materials simulations in the practice of scientific research and industrial innovation.

NOMAD Laboratory - The Novel Materials Discovery Laboratory

Website address	https://nomad-coe.eu/
Contact name	Kylie O'Brien, Pintail Limited
Contact email	kylie.obrien@pintailservices.com
Contact telephone	+353 1 841 8046

Partners involved: Fritz Haber Institute of the Max Planck Society (Coordinator), Aalto University, Humboldt-Universität zu Berlin, King's College London, Max Planck Institute for the Structure and Dynamics of Matter, University of Barcelona, University of Cambridge, Technical University of Denmark, Barcelona Supercomputing Centre, CSC - IT Center for Science, Leibniz Supercomputing Centre, Max Planck Computing and Data Facility, Pintail Ltd

Which areas does your project target?

Computational materials science

Please list the main achievements of your project in one sentence

In the first year of the *NOMAD Laboratory CoE*, methods were established to convert heterogenous data from the major computational materials science codes to a homogeneous, code-independent format and tools were developed to intuitively access, query and visualize this data in ways that will be useful for both academic and industrial users.

Project profile

The *NOMAD Laboratory CoE* is one of a number of international centers to provide 'materials genomics' high-throughput screening. It is special because it focuses on computations but is not restricted to one or a few codes. Instead, it considers essentially all important codes from the fields of ab initio electronic-structure theory and force fields. Eight complementary research groups in computational materials science along with four high-performance computing (HPC) centers form the synergetic core of this CoE. The composition of the *NOMAD Consortium* also reflects the strong embedding of the CoE in the Psi-k, CECAM, and ETSF communities. The goal of NOMAD is to develop *Big-Data Analytics Tools* and a *Materials Encyclopedia* for materials science and engineering. This will be reinforced by advanced graphics and animation tools. In this way, NOMAD will enable a leap forward in computational materials science and 'materials genomics'. To ensure the *NOMAD Laboratory CoE* is valuable and relevant to end-users and results in lasting impact, extensive outreach to industrial and academic end users is a key element of the project, as is developing a sustainable business model for the CoE after the project ends.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

The *NOMAD Laboratory CoE* is well placed to collaborate with other international initiatives in computational materials science, including those led by members of our Scientific Advisory Committee in the US (MICCoM) and Japan (CMI^2).

What future partners are you looking to collaborate with?

Computational Materials Science Centers, other Supercomputing Centers for the calculation of missing data of materials properties.

E-CAM - a European HPC Centre of Excellence that supports Software, Training and Discussion in Simulation and Modelling.

Website address	www.e-cam2020.eu
Contact name	Luke Dury, UCD and Dublin Institute for Advanced Studies
Contact email	ld@cp.dias.ie
Contact telephone	+353-1-6621333 Ext 321

Partners involved:	18 partners (including 3 PRACE centres) from the extended CECAM network of research institutes involved in material science simulations.

Which areas does your project target?

All areas of computational material science from the electronic to the supra-molecular scale with an emphasis on software refactoring and human capacity building.

Please list the main achievements of your project in one sentence

E-CAM is only in its first year, but has already had its first significant engagement with Industrial partners through a successful scoping workshop in Mainz and has started its training activities.

Project profile

All areas of computational material science from the electronic to the supra-molecular scale with an emphasis on software refactoring and human capacity building, training and software development.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

E-CAM seeks to collaborate in three main areas: advanced training; algorithmic innovation and software; and in industrial outreach.

What future partners are you looking to collaborate with?

Industrial partners seeking to use advanced material science simulations; algorithmic innovators and software developers interesting in material simulations; and groups interested in the human capital side of supercomputing, including gender and status issues.

EoCoE - Energie Oriented Center of Excellence

Website address	http://www.maisondelasimulation.fr/
Contact name	Edouard Audit CEA Maison de la Simulation
Contact email	edouard.audit@cea.fr
Contact telephone	N/A

Partners involved: Forschungszentrum Julich Gmbh, Germany, Max-Planck-Gesellschaft Zur Forderung Der Wissenschaften Ev, Germany, Agenzia Nazionale Per Le Nuove Tecnologie, L'energia E Lo Sviluppo Economico Sostenibile, Italy, Centre Europeen De Recherche Et De Formation Avancee En Calcul Scientifique, France, Instytut Chemii Bioorganicznej Polskiej Akademii Nauk, Poland, Universita Degli Studi Di Trento, Italy, Fraunhofer Gesellschaft Zur Foerderung Der Angewandten Forschung E.V., Germany, University of Bath,United Kingdom, The Cyprus Institute, Cyprus, Consiglio Nazionale Delle Ricerche, Italy,Universite Libre de Bruxelles, Belgium, Barcelona Supercomputing Center - Centro Nacional de Supercomputacion, Spain

Project profile

The aim of the present proposal is to establish an Energy Oriented Centre of Excellence for computing applications, (EoCoE). EoCoE (pronounced “Echo”) will use the prodigious potential offered by the ever-growing computing infrastructure to foster and accelerate the European transition to a reliable and low carbon energy supply. To achieve this goal, we believe that the present revolution in hardware technology calls for a similar paradigm change in the way application codes are designed. EoCoE will assist the energy transition via targeted support to four renewable energy pillars: Meteo, Materials, Water and Fusion, each with a heavy reliance on numerical modelling. These four pillars will be anchored within a strong transversal multidisciplinary basis providing high-end expertise in applied mathematics and HPC.

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; a prerequisite for the long-term sustainability of EoCoE and also ensuring that it is deeply integrated in the overall European strategy for HPC. The primary goal of EoCoE is to create a new, long lasting and sustainable community around computational energy science. At the same time, EoCoE is committed to deliver high-impact results within the first three years. It will resolve current bottlenecks in application codes, leading to new modelling capabilities and scientific advances among the four user communities; it will develop cutting-edge mathematical and numerical methods, and tools to foster the usage of Exascale computing. Dedicated services for laboratories and industries will be established to leverage this expertise and to foster an ecosystem around HPC for energy. EoCoE will give birth to new collaborations and working methods and will encourage widely spread best practices.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

Several applications within the EoCoE are and will continue to be developed within international collaborations. Adding new features to applications and enable application for a broader range of architectures.

ESiWACE - Excellence in Simulation of Weather and Climate in Europe

Website address	http://fieg@dkrz.de
Contact name	Dr. Kerstin Feig
Contact email	fieg@dkrz.de
Contact telephone	N/A

Partners involved: Deutsches Klimarechenzentrum GmbH (DE), Coordinator, European Centre for Medium-Range Weather Forecasts (UK), Centre National de la Recherche Scientifique-Institut Pierre Simon Laplace (FR), Max-Planck-Institut für Meteorologie (DE), Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (FR), Barcelona Supercomputing Center (ES), Science and Technology Facilities Council (UK), Met Office (UK), University of Reading (UK), Sveriges Meteorologiska och Hydrologiska Institut (SE), National University of Ireland Galway - Irish Centre for High End Computing (IE), Fondazione Centro Euro-Mediterraneo sui Cambiamenti Climatici (IT), Deutscher Wetterdienst (DE), Seagate Systems UK Limited (UK), BULL/ATOS (FR), Allinea (UK)

Project profile

ESiWACE will substantially improve efficiency and productivity of numerical weather and climate simulation on high-performance computing platforms by supporting the end-to-end workflow of global Earth system modelling in HPC environment. This will be obtained by improving and supporting (1) scalability of models, tools and data management on state-of-the-art supercomputer systems (2) Usability of models and tools throughout the European HPC eco-system, and (3) the Exploitability of the huge amount of resulting data. We will develop solutions for cross-cutting HPC challenges particular to the weather and climate domain. This will range from the development of specific software products to the deployment of user facing services for both, computing and storage.

ESiWACE leverages two established European networks, namely (1) the European Network for Earth System modelling, representing the European climate modelling community and (2) the world leading European Centre for Medium-Range Weather Forecasts. The governance structure that defines the services to be provided will be driven by the European weather and climate science community.

Weather and climate computing have always been one of the key drivers for HPC development, with domain specific scientific and technical requirements that stretch the capability and capacity of existing software and hardware to the limits. By developing solutions for Europe and at European scale, ESiWACE will directly impact on the competitiveness of the European HPC industry by engendering new products, providing opportunities for exploitation beyond the project itself, and by enhancing the skills base of staff in both industry and academia. ESiWACE will be at once thematic, as it focuses on the HPC application domain of climate and weather modeling, transversal, as it covers several aspects of computational science, and challenge-driven, as climate and weather predictability represents a major societal issue.

CoeGSS - Centre of Excellence for Global Systems Science

Website address	http://coegss.eu
Contact name	Michael Gienger, High Performance Computing Center Stuttgart
Contact email	gienger@hlrs.de
Contact telephone	+49 711 685 63824

Partners involved: University of Potsdam, Germany; Global Climate Forum, Germany; University of Stuttgart; HLRS, Germany; Poznań Supercomputing and Networking Center, Poland; Institute for Scientific Interchange, Italy; IMT School for Advanced Studies Lucca, Italy; Chalmers Tekniska Hoegskola AB, Sweden; Atos Spain SA, Spain; TOP-IX Consortium, Italy; The CoSMo Company, France; DIALOGIK, Germany; CSP scarl, Italy

Which areas does your project target?

Agent-based modelling; New HPC community; High Performance Data Analytics

Please list the main achievements of your project in one sentence

CoeGSS addresses an entirely new community for the HPC domain

Project profile

The main outcomes of the CoeGSS project are the development of the technology and the environment for successful collaboration between the stakeholders dealing with global challenges and the High Performance Computing institutions that provide the mandatory capabilities to address those complex challenges at the required scale.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

Collaborations are of particular interest within the Global Systems Science domain (political stakeholders, industrial stakeholders, decision makers, ...).

What future partners are you looking to collaborate with?

Data providers, agent-based modelling experts and emerging technology providers

BioExcel - CoE for Computational Biomolecular Research

Website address http://bioexcel.eu

Contact name Rossen Apostolov,
KTH Royal Institute of Technology

Contact email rossen@kth.se

Contact telephone +46-700-242-545

Partners involved: KTH Royal Institute of Technology, University of Utrecht, University of Edinburgh/EPCC, Barcelona Supercomputing Center (BSC), University of Manchester, EMBL-EBI, Institute for Research in Biomedicine (IRB), Forward Technologies, Ian Harrow Consulting, Jülich Supercomputing Center, Max-Planck Institute for Biophysical Chemistry

Which areas does your project target?

Modelling and simulations of biological systems at quantum, atomistic and macromolecular assemblies up to a cellular level

Please list the main achievements of your project in one sentence

Improved the performance, scalability and functionality of major codes for biomolecular modelling – GROMACS, HADDOCK and CPMD; and devised efficient workflows environments for automated large-scale modelling, simulation and analysis with associated data integration.

Project Profile

BioExcel is a dedicated Center of Excellence that works on 1) Improvement of the performance, efficiency, and scalability of selected software packages of high importance for biomolecular Life Science research on next-generation HPC systems, 2) Improvement of the usability of ICT technologies for biomolecular researchers, both in academia and industry by devising efficient workflow environments with associated data handling and integration of simulation and analysis codes and 3) Promotion of best practices and training of end users to make best use of both software and computational infrastructure.

Please summarise, in one sentence, the potential areas of international collaboration of your project.

Biomolecular modelling and simulations; development of workflow environments for biomolecular studies; exascale approaches for HPC and HTC; co-design; data management

What future partners are you looking to collaborate with?

Experts in HPC/HTC techniques for development of exa-scale software capabilities; User communities of the codes by the CoE; Industrial users (pharma, chemical and food industries) of the codes

CompBioMed – CoE for Computational Biomedicine

Website address	http://www.compbiomed.eu
Contact name	Prof. Peter Coveney, UCL
Contact email	p.v.coveney@ucl.ac.uk
Contact telephone	N/A

Partners involved: Avicenna Alliance, Birmingham City University, Brunel University, Convergence Pharma, GSK, DNA Nexus, Rutgers University, Leibniz Supercomputing Centre, Science Museum, University of Leeds, VPH Institute, Zayed University, Abu Dhabi, Heidelberg Institute for Theoretical Studies, The Hartree Centre

Project profile

Computational Methods, based on human biology, are now reaching maturity in the biomedical domain, rendering predictive models of health and disease increasingly relevant to clinical practice by providing a personalized aspect to treatment. Computer based modelling and simulation is well established in the physical sciences and engineering, where the use of high performance computing (HPC) is now routine.

CompBioMed is a user-driven Centre of Excellence in Computational Biomedicine, 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.

Objectives are to: coalesce the burgeoning HPC user community within the biomedical research field, promote innovation in the field of computational biomedical modelling and simulation, train future generations of scientists within the field of computational biomedicine, use best practice Software Carpentry tools and techniques to develop and sustain existing community codes, engage with a range of industries across the entire healthcare value chain, and engage closely with medical professionals to promote the tools, techniques as well as access mechanisms developed within our Centre.

POP - CoE for Performance Optimisation and Productivity

Website address	https://pop-coe.eu/
Contact name	Prof. Dr. Jesus Labata
Contact email	jesus.labarta@bsc.es
Contact telephone	N/A

Partners involved: Barcelona Supercomputer Center, High Performance Computing Center Stuttgart (HLRS), Jülich Supercomputing Centre (JSC), Numerical Algorithms Group Ltd, Rheinisch-Westfälische Technische Hochschule Aachen, TERATEC

Project profile

High performance Computing is becoming a fundamental tool for the progress of science and engineering and as such for economic competitiveness. 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 when the need to get further performance and efficiency is perceived, code developers may not have sufficient insight on its detailed causes for addressing the problem properly. This may lead to blind attempts to restructure codes and consequent lack of efficiency.

The objective of POP is to operate a Centre of Excellence in Computing Applications in the area of Performance Optimisation and Productivity. POP will offer the service of precisely assessing the performance of computing application of any sort, from a few hundred to many thousand processors. Also, POP will show its Customers the issues affecting the performance of their code and the most optimal way to alleviate them.

POP will target code owners and users from all domains, including infrastructure operators, academic and industrial users. The estimated population of such applications in Europe is 1500 and within the project lifetime POP has the ambition of serving over 150 such codes. The Added Value of POP's services is the savings generated in the operation and use of a code, which will result in a significant Return on Investment (fixing a code costs less than running it below its optimal levels) by employing best-in-class services and release capacity for resolving other priority issues. POP will be a best-in-class centre. By bringing together the European world-class expertise in the area and combining excellent academic resources with a practical, hand-on approach, it will improve the access to computing applications, thus allowing European researchers and industry to be more competitive.