European HPC Technology Handbook

13 November 2015

Supercomputing'15 Bird-of-a-Feather session on 19th Nov 2015

The following project details are included:

Project name ● Web address ● Partners Involved ● Contact Name ● Contact email ● Contact phone ● Project Profile ● Which areas does your project target? ● What areas of your project do you think have a potential for cross-continent synergies?

For additional information and project summaries please visit:

A list of the European Horizon 2020 projects:

http://cordis.europa.eu/search/result_en?q=contenttype%3D%27project%27%20AND%20%27FETHPC%27&p=1&num=10&srt=/project/contentUpdateDate:decreasing

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

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

Web address

http://www.allscale.eu

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

Contact Name

Thomas Fahringer

Contact email

thomas.fahringer@uibk.ac.at

Contact phone

+43 676872564410

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.

Which areas does your project target?

Programming methodologies, environments, languages and tools

What areas of your project do you think have a potential for cross-continent synergies?

Programming small to extreme scale parallel systems

Resilience management for exascale systems

Online performance analysis

Runtime systems for extreme scale HPC

ANTAREX - AutoTuning and Adaptivity appRoach for Energy efficient eXascale HPC systems

Web address

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

Partners Involved

1) Politecnico di Milano (Italy)

2) ETH Zurich (Switzerland)

3) Universidate do Porto (Portugal)

4) INRIA Rennes (France)

5) Consorzio Interuniversitario CINECA (Italy)

6) IT$Innovations (Czech Republic)

7) Dompé Farmaceutici SPA (Italy)

8) Sygic (Slovakia)

Contact Name

Prof. Cristina Silvano

Contact email

cristina.silvano@polimi.it

Contact phone

+39 02 2399 3692

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 the software 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 autotune applications for green and heterogeneous High Performance Computing (HPC) systems up to the Exascale level.

Which areas does your project target?

Programming methodologies, environments, languages and tools

APIs and system software for future extreme scale systems

What areas of your project do you think have a potential for cross-continent synergies?

The proposed approach based on Domain Specific Language the application self-adaptivity and to runtime manage and autotune applications for green and heterogeneous HPC systems has a high potential for cross-continent synergies.

Collaborative Research into Exascale Systemware, Tools and Applications (CRESTA)

Web address

http://www.cresta-project.eu

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

Contact Name

Dr Lorna Smith

Contact email

l.smith@epcc.ed.ac.uk

Contact phone

+44 131 651 3408

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 significance that 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.

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 scale systems

What areas of your project do you think have a potential for cross-continent synergies?

CRESTA's efforts to enable future applications for extreme scale systems are universal and similar efforts are being carried out across continents. Applications that can test extreme scale prototypes are rare and of interest world-wide.

Programming models for future extreme scale systems require international backing and standardization; hence cross-continent engagement is essential.

ComPat

Web address

http://www.compat-project.eu

Partners Involved

http://www.compat-project.eu/consortium/

Contact Name

Prof. Dr. Alfons G. Hoekstra

Contact email

a.g.hoekstra@uva.nl

Contact phone

n/a

Project Profile

see http://www.compat-project.eu/wp-content/uploads/2015/06/ComPat-project-overview-Sept-2015.pdf

Which areas does your project target?

New mathematical and algorithmic approaches

What areas of your project do you think have a potential for cross-continent synergies?

Overall development of the High Performance Multiscale Computing paradigm, and the development of the Multiscale Computing Patterns.

In more detail, the applications that drive the research as well as ComPat middleware and tools have lots of potential.

DEEP (Dynamical Exascale Entry Platform) & DEEP-ER (Dynamical Exascale Entry Platform - Extended Reach)

Web address

http://www.deep-

project.eu

http://www.deep-er.eu

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

Contact Name

Dr. Estela Suarez, Project Manager

Contact email

pmt@deep-

project.eu

pmt@deep-er.eu

Contact phone

n/a

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.

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

What areas of your project do you think have a potential for cross-continent synergies?

HPC core technologies: the Cluster-Booster concept can be generalised into a modular approach for HPC. Different components are combined to provide the system functionalities exploiting the best characteristics of each technology.

Prototype development: exchange of lessons learned.

Programming models and application porting: the DEEP programming model is based on standards (MPI and OpenMP) and has been designed to make application porting between platforms as seamless as possible.

Application scientists: invited to test their codes on the DEEP System. Exchange of experience in code modernisation

ECOSCALE

Web address

http://www.ecoscale.eu/

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

Contact Name

Iakovos Mavroidis

Contact email

iakovosmavro@gmail.com

Contact phone

+30 2810254899

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.

Which areas does your project target?

HPC core technologies and architectures

Programming methodologies, environments, languages and tools

What areas of your project do you think have a potential for cross-continent synergies?

The main areas of ECOSCALE are High Level Synthesis and Reconfigurable Computing. Several commercial systems provide FPGA-based supercomputing nodes, namely Maxeler's MPC series, Convey's HC-2, BeeCube's BEE4, SRC's MAPstation, and Timelogic's DeCypher. The importance of the utilization of reconfigurable computing for future HPC systems is also demonstrated by the fact that there is a special National Science Foundation (NSF) research center in the USA consisting of more than 30 industrial and academic partners focusing only on this topic; the outcome is the most powerful HPC system in the world utilizing reconfigurable technology (Novo-G). Altera and Xilinx have also shown significant interest in reconfigurable computing providing High Level Synthesis tools. It is obvious that High Level Synthesis and Reconfigurable Computing are two close areas of global interest and thus there is a great potential for cross-continental synergies.

EPiGRAM - Exascale ProGRAmming Models

Web address

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

Partners Involved

KTH Royal Institute of Technology

Technische Universität Wien

EPCC - University of Edinburgh

Cray UK

Fraunhofer

Contact Name

Stefano Markidis

Contact email

markidis@kth.se

Contact phone

+46 87907896

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 for extreme-scale computing.

Which areas does your project target?

Programming methodologies, environments, languages and tools

What areas of your project do you think have a potential for cross-continent synergies?

Our work on MPI could have cross-continent synergy with Argonne National Laboratory projects on MPICH.

*ES*CAPE: Energy-efficient Scalable Algorithms for Weather Prediction at Exascale

Web address

http://www.hpc-escape.eu

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

Contact Name

Dr Peter Bauer

Contact email

peter.bauer@ecmwf.int

Contact phone

+44 1189499080

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.

Which areas does your project target?

Programming methodologies, environments, languages and tools

New mathematical and algorithmic approaches

What areas of your project do you think have a potential for cross-continent synergies?

The fundamental algorithmic building blocks are common to most weather and climate prediction models. There is therefore a potential to share the ESCAPE outcome with operational and research centres in other countries.

EXA2CT

Web address

http://www.exa2ct.eu

Partners Involved

1 INTERUNIVERSITAIR MICRO-ELECTRONICA CENTRUM VZW IMEC Belgium

2 UNIVERSITEIT ANTWERPEN UA Belgium

3 UNIVERSITA DELLA SVIZZERA ITALIANA USI Switzerland

4 VYSOKA SKOLA BANSKA - TECHNICKA UNIVERZITA OSTRAVA VSB-TUO Czech Republic

5 INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE INRIA France

6 UNIVERSITE DE VERSAILLES SAINT-QUENTIN-EN-YVELINES. UVSQ France

7 T-SYSTEMS SOLUTIONS FOR RESEARCH GMBH TS-SFR Germany

8 FRAUNHOFER-GESELLSCHAFT ZUR FOERDERUNG DER ANGEWANDTEN FORSCHUNG E.V

9 INTEL CORPORATION SAS France

10 NUMERICAL ALGORITHMS GROUP LTD NAG United Kingdom

Contact Name

Tom Vander Aa

Contact email

tom.vanderaa@imec.be

Contact phone

+32 16 28 80 53

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.

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

What areas of your project do you think have a potential for cross-continent synergies?

The work on solvers is disseminated in PETSc which is used world-wide. The work on GASPI has influence on the MPI standard, which is used world-wide. We are collaborating with TACC in the USA via Intel

ExaFLOW

Web address

http://exaflow-project.eu

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

Contact Name

Erwin Laure

Contact email

erwinl@pdc.kth.s

e **Contact phone**

n/a

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.

Which areas does your project target?

New mathematical and algorithmic approaches

What areas of your project do you think have a potential for cross-continent synergies?

Working on challenges and codes of global importance.

ExaHyPE - An Exascale Hyperbolic PDE Engine

Web address

http://www.exahype.eu

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

Contact Name

Michael Bader

Contact email

bader@in.tum.de

Contact phone

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

Project Profile

ExaHyPE develops a novel Exascale-ready engine to simulate large-scale problems that may be 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.

Which areas does your project target?

New mathematical and algorithmic approaches

What areas of your project do you think have a potential for cross-continent synergies?

For general areas of collaboration, see 12.

High-order numerics and parallel adaptive mesh refinement is still a computational challenge, and addressed by supercomputing groups in the US, e.g.

In terms of supercomputing architectures, ExaHyPE will concentrate on homogeneous designs, such as envisaged for NERSC's Cori system or the CORAL platform "Aurora". The ExahyPE software will be designed as a compute-bound application that fully exploits the capabilities of these machines.

ExaNeSt - European Exascale System Interconnect and Storage

Web address

http://www.exanest.eu

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)

Contact Name

Manolis G.H. Katevenis (FORTH), Coordinator

Contact email

nickpap@ics.forth.gr

Contact phone

+30 2811.39.1078

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.

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

What areas of your project do you think have a potential for cross-continent synergies?

Areas where world-wide API's are needed to take advantage of new technologies:

Applications Tuning, API's and Systems Software for: zero-copy, user-level, RDMA-based interprocessor communication, and for taking advantage of a Global Address Space.

Storage Systems and their API's and Systems Software for: distributed storage systems where NVM's are local to the compute cores.

API's and OS mechanisms for congestion mitigation and resilience in interconnection networks.

ExaNoDe

Web address

not yet defined

Partners Involved

CEA

ARM

ETH Zürich

FORTH

Fraunhofer

SCAPOS

UNIVERSITY OF MANCHESTER

BULL

VOSYS

BSC

JUELICH

KALRAY

CNRS

Contact Name

Denis DUTOIT

Contact email

denis.dutoit@cea.fr

Contact phone

+33 4 38 78 62 58

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 reently 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.

Which areas does your project target?

HPC core technologies and architectures

What areas of your project do you think have a potential for cross-continent synergies?

By decreasing development costs with 3D integration and ARM cores, ExaNoDe will enable more companies to develop application-targeted advanced computing solutions. This will have a potential for cross-continent synergies between our European project and US or Japan companies.

ExCAPE (Exascale Compound Activity Prediction Engines)

Web address

http://www.excape-h2020.eu

Partners Involved

Imec vzw (Belgium)

JANSSEN CILAG SA (Spain)

IT4I (Czech Republic)

ASTRAZENECA AB (Sweden)

UNIVERSITAT LINZ (Austria)

Aalto University (Finland)

Intel (Belgium)

IDEAconsult (Bulgaria)

Royal Holloway (UK)

Contact Name

Tom Ashby

Contact email

ashby@imec.be

Contact phone

+32-16-28-1251

Project Profile

Our project focuses on developing better machine learning algorithms for predicting biological activity 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.

Which areas does your project target?

Programming methodologies, environments, languages and tools

New mathematical and algorithmic approaches

What areas of your project do you think have a potential for cross-continent synergies?

1) Machine learning component for biological activity. The Japanese HPC community has already looked at the same use case, and sharing experiences with them would be very beneficial.

2) Machine learning in general. There are many groups around the world looking at large scale machine learning, although typically in data centres. Any other groups doing ML on supercomputers would be very useful to speak to, regardless of the actual use-case.

3) HPC challenges in the pharmaceutical industry. The same challenges are being faced globally in the pharmaceutical industry, so international collaboration with other industry partners would make sense, to understand their experiences with HPC and their future plans.

4) Programming models. Machine learning (and the life sciences) do not typically use the classical programming models from HPC (MPI, OpenMP, C, Fortran etc). Discussion with other groups looking at programming models (and associated frameworks and implementations) for these domains would be useful.

EXTRA

Web address

http://www.extrahpc.eu

Partners Involved

Ghent University Belgium Telecommunications

Systems Institute Greece Imperial College

London UK

Politecnico di Milano, Italy

University of Amsterdam, The Netherlands

Rurh-University Bochum, Germany

Maxeler, UK

Synelixis, Greece

University of Cambridge, UK

Contact Name

Dirk Stroobandt

Contact email

Dirk.Stroobandt@UGent.be

Contact phone

+32 9 264 34 01

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.

Which areas does your project target?

HPC core technologies and architectures

Programming methodologies, environments, languages and tools

Reconfigurable Computing for HPC

What areas of your project do you think have a potential for cross-continent synergies?

The main concept of reconfigurable computing is studied all over the world, with a large body of work being done in the US and Canada, due to the main suppliers of FPGA technology (Xilinx and Altera) being headquarted in the US. We focus more on higher level reconfiguration aspects and HPC applications that make efficient use of reconfiguration (possibly through innovations in the reconfiguration architectures) which can also benefit other fields of applications using reconfigurable technology.

green flash

Web address

http://green-flash.lesia.obspm.fr/

Partners Involved

Observatoire de Paris

University of Durham

Microgate

PLDA

Contact Name

Damien Gratadour

Contact email

damien.gratadour@obspm.fr

Contact phone

+33 145077757

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.

Which areas does your project target?

HPC core technologies and architectures

Programming methodologies, environments, languages and tools

What areas of your project do you think have a potential for cross-continent synergies?

The baseline strategy for AO real-time control, involving MVM with a wavefront reconstructor matrix, relies extensively on dense linear algebra (DLA) both for the Real-Time box task and the supervisor module tasks. In the case of the former, the problem can be distributed efficiently over several independent nodes processing portions of the matrix with a single data concentration step at the end of the computation. While the strong constraint is on determinism and end-to-end latency of the overall, the job scheduling is relatively easy and follows a bulk synchronous parallel approach, driven by the data transfer from the sensors and the explicit synchronization by data concentration before it can be sent to the optics. The strategy in this case is to optimize the process execution on a single node and overlap communication with sensors and first processing steps. The algorithm is memory bound and the local implementation should rely on optimized utilization of the compute device's memory bus. Concerning the supervisor module, the most compute intensive task is the computation of the reconstructor matrix. It involves the inversion of a large covariance matrix either through matrix factorization of through direct inversion. In the latter case, an embarrassingly parallel problem with a numerical complexity scaling with N3, an optimized implementation should optimize the usage of computing cores rather than the memory bus. However, the underlying

algorithms require frequent synchronizing of global communications, which represents a bottleneck that limits performance. This bottleneck will be further exacerbated when concurrency will reach billions of core in Exascale systems and should be addressed as a general issue for the HPC community.

INTERTWinE (Programming Model INTERoperability ToWards Exascale)

Web address

http://www.intertwine-

project.eu/ **Partners Involved**

EPCC, University of Edinburgh, UK (Lead)

Barcelona Supercomputing Center, Spain

Kungliga Tekniska Hoegskolan (KTH), Sweden

Inria, France

Fraunhofer, Germany

Deutsches Zentrum für Luft und Raumfahrt,

Germany TS-SFR Solutions for Research, Germany

Universitat Jaume I de Castellon, Spain

University of Manchester, UK

Contact Name

Mark Bull (Project Leader) / George Beckett (Project Manager)

Contact email

m.bull@epcc.ed.ac.uk

george.beckett@ed.ac.uk

Contact phone

George Beckett : +44 131 651 3577

Project Profile

The INTERTWinE project addresses the problem of programming model design and implementation for the Exascale. It is funded through the European Commission Horizon 2020 program, FET-HPC 2014 Research and Innovation Action.

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 will be required.

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.

Although there remains room for improvement in individual programming models and their implementations, the main challenges lie in interoperability between APIs. It is this interoperability, both at the specification level and at the implementation level, which we seek to address and to substantially improve.

INTERTWinE brings together the principal European organisations driving the evolution of programming models and their implementations. We will focus on seven key programming APIs: MPI, GASPI, OpenMP, OmpSs, StarPU, QUARK and PaRSEC, each of which has a project partner with extensive experience in API design and implementation.

Interoperability requirements and evaluation of implementations will be driven by a set of kernels and applications, each of which has a project partner with a major role in their development.

INTERTWinE will implement a co-design cycle, by feeding advances in API design and implementation into the applications and kernels, thereby driving new requirements and hence further advances.

Which areas does your project target?

Programming methodologies, environments, languages and tools

What areas of your project do you think have a potential for cross-continent synergies?

- Engagement with (international) standards bodies in the area of programming models for HPC applications (GASPI, MPI, OpenMP).

- Collaboration with researchers in other regions, including:

    - University of Tennessee, USA - Jack Dongarra is a project investigator.
    - Argonne National Labs, USA:

    - Pavan Balaji is collaborating with project partner BSC on GLES? (funded G8 project).

- Marcus Neil actively researching better interoperability between the two most popular HPC programming models (that is, MPI and OpenMP).

Mont-Blanc

Web address

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

Partners Involved

BSC

Bull/Atos ARM

LRZ

Juelich

Allinea

CEA HLRS

Inria

Universidad de Cantabria (Santander)

GENCI

University of Bristol ETH

Zurich

CRNS (Centre national de la Recherche Scientifique) / LIRMM AVL

Institute for Scientific Computing of KF Univ. Graz

Université de Versailles Saint Quentin

Contact Names

Filippo Mantovani, Etienne Walter

Contact email

filippo.mantovani@bsc.es, etienne.walter@atos.net

Contact phone

+34 934137716, +33 130807496

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-efficient embedded 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 bit architecture
- 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 focussed 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.

Which areas does your project target?

HPC core technologies and architectures

Programming methodologies, environments, languages and tools

What areas of your project do you think have a potential for cross-continent synergies?

We consider the following topics as powerful tools and starting points for future collaborations (with no geographical limitations):

- To access the computing platforms (ARM-based prototypes) made available to partners and to a selected pool of End-Users of the Mont-Blanc project.
- To adopt and contribute to the system software stack developed within the Mont-Blanc project over the years, that is needed for future ARM based Exascale systems
- To organize training and workshop promoting the interchange of experiences in area related to the project (prototyping, ARM-based computing, innovative programming models, etc.).

Network for Sustainable Ultrascale Computing (NESUS)

Web address

http://www.nesus.eu

Partners Involved

"Technische Universität Wien University of Innsbruck IMEC University of Mons (UMONS) University of Sarajevo International University of Sarajevo Institute for Parallel Processing Rudjer Boskovic Institute University of Cyprus SIX Research Center Technical University of Denmark University of Tartu Abo Akademi University CSC – IT Center of Science Ltd. INRIA Centre National de la Recherche Scientifique University Paul Sabatier University Ss. Cyril and Methodius "St. Paul the Apostle" University University Bayreuth Technische Universität Chemnitz Johannes Gutenberg-Universität Mainz Aristotle University of Thessaloniki Harokopio University of Athens University of Macedonia University of Ioannina Foundation for Research and Technology – Hellas (FORTH) MTA SZTAKI Budapest University of Technology and Economics University College Dublin University College Cork Università della Calabria CNR-ICAR Universita di Modena e Reggio Emilia University of Torino Vilnius Gediminas Technical University University of Luxembourg University of Malta University of Donja Gorica University of Amsterdam University of Bergen Norwegian University of Science and Technology Oslo and Akerrshus University College University of Tromso Czestochowa University of Technology Poznan Supercomputing and Networking Center INESC ID INESC TEC University Politehnica of Bucarest West University of Timisoara University of Nis Comenius University in Bratislava University of Zllina Jožef Stefan Institute University of Ljubljana La Laguna University Universidad Jaime I University of A Coruña University Carlos III of Madrid Barcelona Supercomputing Center Universidad de Murcia Universidad Católica San Antonio de Murcia Universidad de Valladolid University of Extremadura CIEMAT KTH Royal Institute of Technology Uppsala University University of Neuchatel University of Applied Sciences of Western Switzerland Istanbul Technical University University of Cambridge University of Derby University of Manchester University of York Polytechnic University of Tirana National Academy of Sciences of the Republic of Armenia Research and Educational Networking Association of Moldova Moscow State University Lviv Polytechnic National University The University of Sydney St Francis Xavier university Universidad de los Andes Indraprastha Institute of Information Technology (IIIT) Centro de investigación y Estudios Avanzados del Instituto Politécnico Nacional (CINVESTAV) Northwestern University "

Contact Name

Jesus Carretero

Contact email

jesus.carretero@uc3m.es

Contact phone

Project Profile

+34 916249458

Which areas does your project target?

NESUS is ICT COST action. (IC1305).

The objective is to promote cooperation and networking to create a research community focused on sustainability in extreme scale/ultrascale computing systems. Understanding sustainability in a holistic way, not only energy efficiency.

The goal of the NESUS Action is to establish an open European research network targeting sustainable solutions for ultrascale computing aiming at cross fertilization among HPC, large scale distributed systems, and big data management. The network will contribute to glue disparate researchers working across different areas and provide a meeting ground for researchers in these separate areas to exchange ideas, to identify synergies, and to pursue common activities in research topics such as sustainable software solutions (applications and system software stack), data management, energy efficiency, and resilience. Some of the most active research groups of the world in this area are members of this proposal. This Action will increase the value of these groups at the European-level by reducing duplication of efforts and providing a more holistic view to all researchers, it will promote the leadership of Europe, and it will increase their impact on science, economy, and society.

What areas of your project do you think have a potential for cross-continent synergies?

Programming methodologies, environments, languages and tools

APIs and system software for future extreme scale systems

New mathematical and algorithmic approaches

Sustainability in extreme scale systems

What tools and mechanisms do you think could be used to facilitate the cross-continent synergies your project could produce?

NESUS has working groups related to: Programming models and runtimes; Resilience of applications and runtime environments; Sustainable data management; Energy efficiency; and Applications. Research on sustainable programming and execution models in the context of rapidly changing underlying computing architecture is a major goal in NESUS. The idea is to explore synergies among emerging programming models and run-times from HPC, distributed systems, and big data management communities that can be portable and (almost) transparent to application programmers.

This research line is very important in Exascale if we want them to be adopted. We have to provide programming models supporting different kind of underlying systems and helping to reduce the effort to redesign existing applications.

Data management techniques is also a major area having potential for cross-continental synergies, provided that nowadays data management is a major challenge to reach Extreme scale. The goal of NESUS is to contribute to the evolution of the storage I/O stack towards higher-levels of scalability and sustainability by enhancing data sharing/integration (globalization of data), improving the programmability of data management and analysis, and by providing adaptivity to manage data workload and uncertainty.

Energy efficiency is another area with high potential for cross-continent synergies. NESUS is contributing to develop techniques for monitoring and energy analysis of large scale infrastructures and to propose new holistic models of energy consumption for ultrascale systems. The goal is to provide those models and information to the several layers of the system to help in the design of energy aware software components that can allow users to specify energy issues in their applications.

Finally, there is also potential for cooperation in the applications area, as a NESUS goal is to study the impact of application requirements on the ultrascale system design. It is mandatory to evaluate the needs of the HPC applications concerning scalability, programmability, portability and resilience, trying to identify computational patterns for expressing the applications at a higher level of abstraction for leveraging programming for ultrascale systems. The final goal is to reduce design and programming efforts, while building portable programs that can live longer.

NEXTGenIO

Web address

http://www.nextgenio.eu

Partners Involved

EPCC

Intel

Fujitsu

BSC

TU Dresden

ECWMF

Allinea

Arctur

Contact Name

Michele Weiland

Contact email

m.weiland@epcc.ed.ac.uk

Contact phone

+44 (0)131 651 3580

Project Profile

The overall objective of the Next Generation I/O Project (NEXTGenIO) is to design and

prototype a new, scalable, high-performance, energy efficient computing platform, based on non-volatile memory technologies, designed to address the challenge of delivering scalable I/O performance to applications at the Exascale. These hardware and systemware developments will be coupled to a co-design approach driven by the needs of some of today's most demanding HPC applications. By meeting this overall objective, NEXTGenIO will solve a key part of the Exascale challenge and enable HPC and Big Data applications to overcome the limitations of today's HPC I/O subsystems.

Today most high-end HPC systems employ data storage separate from the main system and the I/O subsystem often struggles to deal with the degree of parallelism present. As we move into the domain of extreme parallelism at the Exascale we need to address I/O if such systems are to deliver appropriate performance and efficiency for their application user communities.

The NEXTGenIO project will explore the use of non-volatile memory technologies and associated systemware developments through a co-design process with three 'end-user' partners: a high-end academic HPC service provider, a global numerical weather centre and a commercial on-demand HPC service provider. These partners will develop a set of I/O workload simulators to allow quantitative improvements in I/O performance to be directly measured on the new system in a variety of research configurations. Systemware software developed in the project will include performance analysis tools, improved job schedulers that take into account data locality and energy efficiency, optimised programming models, and APIs and drivers for optimal use of the new I/O hierarchy.

The project will deliver immediately exploitable hardware and software results and show how to deliver high performance I/O at the Exascale.

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

What areas of your project do you think have a potential for cross-continent synergies?

All areas of the project have that potential.

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

Web address

N/A - NLAFET web will be developed after kick-off meeting (see below).

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

Contact Name

NLAFET Coordinator: Bo Kågström

Professor of Numerical Analysis and Parallel Computing

Director of High Performance Computing Center North (HPC2N)

Contact email

bokg@cs.umu.se

Contact phone

+46 90 786 5419

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 a critical 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 heterogeneity, avoid communication bottlenecks, respond to escalating fault rates, and help

meet emerging power constraints;

(ii) exploration of advanced scheduling strategies and runtime systems focusing on the extreme scale and strong scalability in multi/manycore and hybrid environments;

(iii)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.

Which areas does your project target?

Programming methodologies, environments, languages and tools

New mathematical and algorithmic approaches

See focus topics in Project profile above.

What areas of your project do you think have a potential for cross-continent synergies?

NLAFET is a co-design effort for designing, prototyping, and deploying new linear algebra software libraries

- Exploration of new algorithms

- Investigation of advanced scheduling strategies

- Investigation of advanced auto-tuning strategies

- Validation of results in challenging scientific and engineering HPC applications

- Open source library software

NLAFET will undertake and encourage different types of collaborations with different stakeholders, including library software users, academia, as well as HW and SW vendors. These collaborations could certainly lead to potential cross-continent synergies.

NUMEXAS

Web address

http://www.numexas.eu

Partners Involved

1-CIMNE (CENTRE INTERNACIONAL DE METODES NUMERICS EN ENGINYERIA) SPAIN

2-CSUC (CONSORCI DE SERVEIS UNIVERSITARS DE CATALUNYA) SPAIN

3-LUH (GOTTFRIED WILHELM LEIBNIZ UNIVERSITAET HANNOVER) GERMANY

4-NTUA (NATIONAL TECHNICAL UNIVERSITY OF ATHENS) GREECE

5-QUANTECH ATZ SA SPAIN

Contact Name

Cecilia Soriano - Project Manager

Contact email

csoriano@cimne.upc.edu

Contact phone

+34 93 401 7440

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

Which areas does your project target?

Programming methodologies, environments, languages and tools

New mathematical and algorithmic approaches

What areas of your project do you think have a potential for cross-continent synergies?

New Numerical methods and programming algorithms that constitute the next generation of numerical simulation techniques that are scalable to millions of cores so that exascale class problems can be solved routinely. The goal is the development and implementation of new numerical simulation techniques amenable to scalability to millions of cores along the complete simulation pipeline for a variety of large-scale multidisciplinary problems in applied sciences and engineering: parallel pre-processing and grid generation, parallel structured/unstructured field solvers of high order, parallel optimum design solvers considering uncertainties and parallel in-solver visualization and feature extraction

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

Web address

http://www.readex.eu

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

Contact Name

Joseph Schuchart

Wolfgang E. Nagel

Contact email

joseph.schuchart@tu-dresden.de

wolfgang.nagel@tu-dresden.de

Contact phone

+49 351 463 31673

+49 351 463 35450

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.

Which areas does your project target?

Programming methodologies, environments, languages and tools

What areas of your project do you think have a potential for cross-continent synergies?

There are groups working on auto-tuning in the US. While none of them is considering the System Scenario Methodology as a basis, some collaboration on the topic of tuning strategies and tuning parameters can provide viable synergies. Some of these groups have been invited to the workshop on code auto-tuning to be held at CGO 2016. The meeting will also be used to explore possible synergies.

SAGE (Percipient StorAGE for Exascale Data Centric Computing)

Web address

http://www.sagestorage.eu

Partners Involved

1. Allinea, UK

2. Atos, France

3. CEA, France

4. CCFE(UK Atomic Energy Agency), UK

5. DFKI, Germany

6. Diamond Light Source , UK

7. Juelich (FZJ), Germany

8. KTH, Sweden

9. Seagate (Co-ordinator), UK

10. STFC, UK

Contact Name

Sai Narasimhamurthy

Contact email

sai.narasimhamurthy@seagate.com

Contact phone

+44(0)7584080691

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 systems infrastructure, combined with the potential for Exabyte scale deployment in future generations 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 HPC and 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.

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

What areas of your project do you think have a potential for cross-continent synergies?

SAGE targets a range of applications & data intensive use cases for co-design (Bio-informatics, Space Weather, Satellite Data Processing, Visualization, Fusion Energy Use Cases, Synchrotron Use Cases, etc).

SAGE aims for a new storage architectures, with outstanding performance and new features which could be exploited by a broad range of applications. Within this project the set of applications will be analysed with respect to their I/O requirements. The methodologies developed in SAGE could also be applied to selected applications from the DoE CoEs.

Fusion energy use cases in SAGE will target the ITER fusion experiment, expected to have very wide international reach and impact (with seven partners. China, EU, India, Japan, Korea, Russia & the US).

The programming model areas will look at innovations in MPI & PGAS to bring about the effective utilization of NVRAM. These innovations can be looked at by the MPI consortium.

The object storage APIs developed in the project has the potential for international standardisation as an Exascale Object storage API.

The international data analytics community can take note of innovations in complex data analytics workflows and NVRAM usage with Apache Flink in SAGE. Apache Flink will pursue international cross-continent collaborations in this area.