Enabling Exascale in Europe for Industry

PRACEdays15 Satellite Event by European Exascale Projects

Date: Tuesday, 26th May 2015
Location: Aviva Stadium, Dublin

Exascale research in Europe is one of the grand challenges tackled by the Seventh Framework Programme for Research and Technological Development (FP7) and becomes even more important in Horizon 2020, the next European framework programme. Started in 2007 and with a total budget of over €50 Billion, the idea is to respond to Europe's needs in terms of jobs and competitiveness, and to maintain leadership in the global knowledge economy. The "European Exascale Projects" represent the collective efforts of multiple collaborative projects funded by the European Commission to investigate and develop hardware and software for Exascale supercomputers.

One of the strategic objectives is to strengthen the scientific and technological base of European industry. Following this aim, the European Exascale projects will organize a European Exascale workshop with a particular focus on industry, in conjunction with PRACEdays15. More information about the European Exascale projects can be found at www.exascale-projects.eu.

Agenda

09:00 Welcome address by EC, Dr. Panagiotis Tsarchopoulos

09:15 Welcome & Introduction by Prof. Dr. Thomas Lippert, Head of Jülich Supercomputing Centre

09:30 Keynote: Exascale Needs & Challenges for Aeronautics Industry, Eric Chaput, Airbus

10:30 – 11:00 Coffee break

11:00 DEEP & DEEP-ER: Innovative Exascale Architectures in the light of user requirements, Gabriel Staffelbach, Estela Suarez & Marc Tchiboukdjian

11:45 Mont-Blanc: High Performance Computing from Commodity Embedded Technology, Filippo Mantovani, Yoon K Ho

12:15 – 13:15 Lunch break

13:15 CRESTA: Developing Software and Applications for Exascale Systems, Mark Parsons

13:45 EPiGRAM: Software in Support of Current and Future Space Missions, Stefano Markidis

14:15 EXA2CT: Mining Chemical Space Annotation to tackle the Phenotypic Challenge of Pharma Industry, Hugo Ceulemans

14:45 NUMEXAS: Embedded Methods for Industrial CFD Applications, Riccardo Rossi

15:15 Panel Discussion moderated by Gilad Shainer

15:45 Conclusion & Farewell

Registration: https://events.conferencepartners.ie/ei/getdemo.ei?id=259&s=_8800Q6PMD

Abstracts

Keynote: Exascale Needs & Challenges for Aeronautics Industry
Exascale computing is seen as a key enabling technology for future aircraft design to be developed and optimised in a fully multi-disciplinary way, making a wide use of design systems that provide integrated analysis and optimisation capabilities which allow for a real-time/interactive way of working. The move from RANS to unsteady Navier-Stokes simulation, (ranging from current RANS-LES to full LES) and/or Lattice Boltzmann method will significantly improve predictions of complex flow phenomena around full aircraft configurations with advanced physical modelling. For instance moving LES capability from Petascale to Exascale computing will accelerate the understanding of noise generation mechanisms and will enable the elaboration of flow control strategy for noise reduction. Multi-disciplinary analysis and design, and real time simulation of aircraft manoeuver, supported by affordable CFD-based aerodynamic and aero elastic data prediction will be a significant change of paradigm in aeronautics industry. The challenges faced by our industry at the horizon of 2025 will be presented together with the expectations on Exascale computing likely to bring operational benefits at that time.

DEEP & DEEP-ER: Innovative Exascale Architectures in the Light of User Requirements
When developing new architectures for the Exascale era, the chicken-or-egg question arises of what to work on first: new hardware or new codes actually able to fully exploit Exascale systems. In the DEEP and DEEP-ER projects we tackle this challenge by adopting a comprehensive, holistic approach. We have come up with an innovative hardware concept, called the Cluster-Booster architecture. At the same time we develop the software stack and work very closely with our application partners to thoroughly integrate all three aspects. For our pilot applications, on the one hand we optimise their codes for our system, and on the other hand we are developing the system design based on the Exascale requirements that our users have. In this session we will explain our basic concept and share two of our use cases: Our industry partner CGG will talk about seismic imaging in the oil and gas industry, and our partner CERFACS on computational fluid dynamics. These two use cases will clearly demonstrate the potential the DEEP architecture offers at Exascale, not least for industrial users.

Mont-Blanc: High Performance Computing from Commodity Embedded Technology
In this session, the coordinator of the Mont-Blanc project will present an overview and status of this European project together with RR HPC Tech Lead Specialist-Aerothermal Methods at Rolls-Royce, a member of the Industrial End-User Group. He will present their observations from the process of testing the low-energy HPC prototypes produced by the project.

CRESTA: Developing Software and Applications for Exascale Systems
The CRESTA project was one of three complementary Exascale software projects funded by the European Commission. The recently completed project employed a novel approach to Exascale system co-design, which focused on the use of a small set of representative applications to inform and guide software and systemware developments. The methodology was designed to identify where problem areas exist in applications and to use that knowledge to consider different solutions to those problems, which inform software and hardware, advances. Using this approach, CRESTA has delivered on all of its outputs, producing a set of Exascale focused systemware and applications.

EPiGRAM: Software in Support of Current and Future Space Missions
During the preparation of NASA and ESA space missions, several simulations of different scenarios in space are carried out on HPC systems. These large scale simulations allow scientists to plan the space missions and to investigate possible phenomena of interest. In this talk, we present the new software developed by the EPiGRAM project to increase the scalability of these codes, the performance of the I/O activities and the amount of useful data for analysis. The impact of the EPiGRAM software on the current NASA Magnetospheric Multiscale Mission (MMS) and on the proposed ESA THOR mission (http://thor.irfu.se/) is discussed.

NUMEXAS: Embedded Methods for Industrial CFD Applications
A problem of paramount importance in the simulations of real engineering problems is the construction of a suitable discretization. It is widely acknowledged that the meshing step required to obtain a suitable geometry may take 90% of the time needed to obtain an engineering result. The objective of our work is to develop a technology to embed "dirty" geometries within a background mesh, which is then adapted to fit the requirements of the simulation. The technique employed results in a methodology, which is both robust and scalable on modern HPC hardware.

EXA2CT: Mining Chemical Space Annotation to tackle the Phenotypic Challenge of Pharma Industry
The trajectory from a biological concept to a drug available to patients is expensive and typically spans over a decade. Drug discovery starts by mapping a disease mapped to a scalable experiment in a test tube. This enables the screening libraries of chemicals for hits or active compounds, from which chemical starting points or leads are selected. These leads are then taken through a cascade of follow-up assays and animal models to optimize their potency on the intended protein targets implicated in disease, while controlling their activity on undesired targets associated with side effects. Finally, the compound is transferred to drug development, where the candidate drugs are tested in human subjects in three subsequent clinical phases. Still, the vast majority of candidates that enter drug development do not make it through to approval. One current trend to mitigate the high attrition rate is to do the initial screening in more complex, so-called phenotypic assays, which are believed to emulate the disease much better than biochemical assays, and that do not rely on the limiting knowledge of which targets are critical for effect. The phenotypic approach, however, presents challenges of their own: their throughput is lower, implying a need for more compact libraries. Secondly, many of the existing compound optimization processes require knowledge of the target. Both of these challenges can be addressed by improving the industries capabilities to predict the activities of chemical on not just the intended protein target, but on as many proteins and processes as possible. To this end, we propose scaled-up machine learning approaches that can mine extensive but heterogeneous information on biological activities of chemicals that is accessible to the industry, to learn to predict it comprehensively. Moreover, we believe computational approaches enable us to extract much more relevant primary information for these exercises from industry standard screens; for instance, by more extensive image analysis, feature selection and machine learning microscopy based screens. Finally, progress is being made in not only formulating predictions, but also quantifying the reliability of predictions, not just for each model for a certain target, but even for individual prediction of a given chemical at a given concentration on a given target.

Speakers

Eric Chaput, Airbus - Flight-Physics Capability Strategy

Dr. Eric Chaput joined Airbus in 1992 after a Ph.D. in Energetics and Optimisation, Post-doctoral positions in Experimental and Numerical Simulation at University of Poitiers, and six years' experience at Airbus Defence & Space working for ARIANE and HERMES programmes. He became subsequently CFD Research Manager, before managing Aerodynamics Methods and in 2004, Senior Manager of Flight-Physics Methods. He is currently the leader of Airbus Flight-Physics capability strategy and a Senior Expert in Aerodynamics Flow Simulation Methods. He has long experience and interest in HPC, driving within the HPC Steering Board the needs and investment for Airbus Engineering, and for more than 15 years member of the Management Board of CERFACS, a research organization and Center of Competence in HPC serving a wide range of industrial sectors.

Estela Suarez, Jülich Supercomputing Centre

Dr. Estela Suarez is the project manager for DEEP & DEEP-ER, two European funded Exascale research endeavors. She works at Juelich Supercomputing Center in Germany and holds a PhD in Physics from the university of Geneva. Already early on Estela has engaged intensively in scientific simulations and computing. Her passion for this research area made Estela follow a career in an HPC environment and enter the Exascale world.

Marc Tchiboukdjian, CGG

Dr. Marc Tchiboukdjian currently works as IT Architect for CGG, a fully integrated geoscience company providing leading geological, geophysical and reservoir capabilities to the oil and gas industry. He holds a PhD from the University of Grenoble and has been active in the field of Exascale research for the last four years. Within the DEEP project, Marc is working on mapping seismic imaging algorithms on the DEEP architecture and evaluating their performance.

Gabriel Staffelbach, CERFACS

Dr. Gabriel Staffelbach is a senior researcher at Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique (CERFACS). He has been working on numerical simulation of combustion and high performance computing since 2002 and is an active user of most HPC systems available to the scientific community via both the PRACE and INCITE programs.

Stefano Markidis, KTH Royal Institute of Technology

Stefano Markidis is Assistant Professor in High Performance Computing at the KTH Royal Institute of Technology. He is a recipient of the 2005 R&D100 award and author of more than 50 articles in peer-reviewed articles. His research interests include large scale simulations for space physics applications.

Mark Parsons, EPCC, the University of Edinburgh

Prof. Mark Parsons joined EPCC, the supercomputing centre at The University of Edinburgh, in 1994 as a software developer working on several industrial contracts following a PhD in Particle Physics undertaken on the LEP accelerator at CERN in Geneva. In 1997 he became the Centre's Commercial Manager and subsequently its Commercial Director. Today he is EPCC's Executive Director (Research and Commercialisation) and also the Associate Dean for e-Research at Edinburgh. He has many interests in distributed computing ranging from its industrial use to the provision of pan-European HPC services through the PRACE Research Infrastructure. His research interests include highly distributed data intensive computing and novel hardware design.

Riccardo Rossi, Universitat Politècnica de Catalunya (UPC)

Dr. Riccardo Rossi, holds a PhD in Civil Engineering from the Technical University of Catalonia (UPC) and is Senior Researcher at CIMNE and tenure-track lecturer at UPC BarcelonaTech. He has extensive experience in the field of Computational Solid and Fluid Dynamics and in the solution of Fluid-Structure Interaction problems, using both body fitted and embedded approaches. He is one of the authors of the multiphysics code KRATOS and author of 36 JCR papers and some 50 conference presentations in the field, including a plenary lecture. Dr. Rossi is also a member of the executive committee of SEMNI and has contributed to the organization of ECCOMAS conferences.

Filippo Mantovani, Barcelona Supercomputing Center

Filippo Mantovani is a postdoctoral research associate of the Heterogeneous Architectures group at the Barcelona Supercomputing Center (BSC). He graduated in Mathematics and holds a PhD in Computer Science from University of Ferrara in Italy. He has been a scientific associate at the DESY laboratory in Zeuthen, Germany, and at the University of Regensburg, Germany. He spent most of his scientific career in computational physics, computer architecture and high performance computing, contributing to the Janus, QPACE and QPACE2 projects. He joined BSC's Mont-Blanc project in 2013, becoming recently technical coordinator of the project.

Hugo Ceulemans, Janssen

Hugo holds an M.D., an M.Sc. in Bioinformatics and a Ph.D. in Molecular Biology from the University of Leuven, and did postdoctoral fellowships in molecular and computational phosphatase biology at the University of Leuven and in structural bioinformatics at the EMBL in Heidelberg. He joined Janssen in 2008 as a computational biologist supporting the Infectious Diseases and Vaccines franchise with models that predict the clinical efficacy of multi-drug regimens in HIV patients given viral sequences. Over the years, his responsibilities extended to cover additional computational approaches and all disease franchises in Janssen. Three years ago, these activities were consolidated in a new Computational Systems Biology unit, which now offers the analysis and integration of sets of chemical, biochemical, omics, phenotypic and clinical data and the formalization of drug discovery knowledge in predictive quantitative models. Mining the extensive, but heterogeneous annotation of the various biological effects of millions of chemicals is one of the major activities of the unit.

Gilad Shainer, HPC Advisory Council Chairman

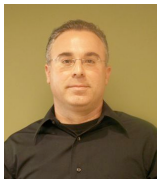

Gilad Shainer is an HPC evangelist who focuses on high-performance computing, high-speed interconnects, leading-edge technologies and performance characterizations. Mr. Shainer holds an M.Sc. degree (2001, Cum Laude) and a B.Sc. degree (1998, Cum Laude) in Electrical Engineering from the Technion Institute of Technology in Israel. He also holds patents in the field of high-speed networking.