Extreme computing

The Dynamical Exascale Entry Platform (DEEP) project is developing concepts for supercomputers one thousand times more powerful than existing machines. Project Manager **Estela Suarez** describes the project's potential and discusses the obstacles her team is overcoming to achieve this feat

How did your academic career lead you to become involved with your current research endeavours?

My background lies in high-energy astrophysics. Quite early in my studies, I gained first-hand experience in scientific programming, something by which I was immediately fascinated. Over time, I established a real passion for it. In my PhD, I developed the full simulation package of an X-ray polarisation detector. Pursuing a career in the area of high-performance computing (HPC) took me to the Jülich Supercomputing Centre, Germany. I have been working here for the last three years, coordinating the Dynamical Exascale Entry Platform (DEEP) and DEEP – Extended Reach (DEEP-ER) projects.

Could you describe how a computer can be 'exascale' and detail the difficulties in building such a supercomputer?

Exascale means that future supercomputers will be 100-1,000 times more powerful than those currently. Although hard to imagine, there will be cases requiring machines with enormous computing power. We are facing great societal challenges – climate change, food safety, energy security and an ageing population. Simulating future scenarios will help us to find solutions to these challenges. However, complex issues

Enabling exascale

A collaborative project coordinated by the **Jülich Supercomputing Centre** in Germany and involving more than 15 partners from all over Europe is working to develop a computing system of unprecedented power, able to simulate some of the biggest challenges facing society

IN THE LAND of computing, speed is king. Current speedy computers enable their users to achieve data-heavy tasks unimaginable even 10 years ago. The computing performance of electronic devices is measured by FLoating-point Operations Per Second (FLOPS). The next milestone in computing's evolution is exascale computing systems capable of performing more than one exaFLOP, or 10^{18} FLOPS. This would represent a monumental achievement for computer science, as such supercomputers would have processing power tantamount to the human brain. Computing at this level would enable the simulation of several phenomena, including climate, space weather and brain networks.

The Dynamical Exascale Entry Platform (DEEP) project, led by Dr Estela Suarez, aims to make this future a reality by developing a novel, exascale-enabling, supercomputing platform. There are many challenges en route, such as the need to increase computer power by a factor of 1,000 while reducing the overall power consumption. However, by addressing the problems of programmability, scalability and energy efficiency, the DEEP project has derived novel solutions for each.

EUROPE-LED INNOVATION

DEEP – which began in December 2011 and will be completed in May 2015 – is coordinated by the Jülich Supercomputing Centre in Germany and involves partners across eight countries.

The €18.5 million project advances the concept of computer acceleration to an unprecedented level. Instead of using the usual approach of adding accelerator cards to Cluster nodes, DEEP uses an accelerator Cluster – called a Booster – to complement conventional high-performance computing (HPC) systems and increase performance. By combining this with a software stack dedicated to meeting exascale requirements, DEEP's computer architecture will enable an unparalleled level of scalability.

NOVEL ARCHITECTURE

This new type of 'Cluster-Booster Architecture' is the crux of the project's approach to HPC and represents a pioneering method of solving the challenges of exascale computing. The individual computer nodes are the building blocks of any supercomputer, and DEEP uses two kinds. "The Cluster part of the system uses classic central processing units (CPUs), while the Booster uses more powerful and energy efficient Xeon Phi accelerators," Suarez explains. "The code of an application is distributed to the appropriate nodes; the complex parts of the code run on the Cluster, and the simpler ones that can be processed quickly run on the Booster. This approach will map inherently heterogeneous applications onto exascale-ready hardware." This approach removes the central bottleneck of traditional accelerated computing.

DEEP's approach is made even more effective by directly connecting the nodes within the Booster to a high-performance interconnecting network called EXTOLL. It is the first project worldwide to achieve such a feat.

A HOLISTIC APPROACH

Although the hardware is transformational, the core innovation of the project is its holistic approach, integrating the architectural, system software and application levels. The project partners believe such co-design is the only way to benefit from exascale performance at an application level.

The team's matching software stack will make the adaptation of applications to this architecture seamless, as the programming model employed in DEEP provides a dedicated development and runtime environment. Additionally, the codes developed in this environment will run on other systems easily, as the researchers have built the applications using industry standards. Finally, the disparate characteristics of the Cluster and Booster sides makes them suited to different kinds of application codes. Parts of an application with limited scalability run well on the Cluster side, like these require more sophisticated and computer-intensive algorithms than we have at hand today. Therefore, we need to develop computers capable of performing such complex calculations.

To build exascale-ready computers, we first need to overcome a number of challenges. In DEEP, we are addressing the topics of programmability, scalability and energy efficiency. A highly scalable input/output system and resiliency are important aspects addressed in the follow-up project, DEEP-ER.

Why is Europe the hub for such R&D initiatives?

We are developing an exascale ready supercomputing architecture because we want to help tackle the aforementioned societal challenges. This requires IT experts and scientists to work together extremely closely. Europe provides excellent infrastructures to do so. It is also renowned for its fundamental research in several areas that will require exascale computing, for instance, medical engineering or environmental studies. Also, energy efficiency is an important challenge at the exascale. European researchers have a lot of experience and expertise in this field, and this is supported by politics and society. Hence, support for this element of our research is much higher in Europe than it might be in Asia or the US.

What excites you about its potential?

We follow a holistic approach in our project: the topics include hardware, system software, tools, energy efficiency, resiliency and applications. It presents an architectural concept, which addresses the most important challenges we face when working towards exascale. This makes it a challenging but very exciting project. Integrating all these aspects into medium-sized prototypes constitutes the first step on the path to building an exascale production system that really works. This is not only important within the HPC world, but also for society at large, since technology is driving forward European industry and research, increasing their future competitiveness.

What stage are you at in your investigation? When do you hope to be operational?

The project has been running for around two-and-a-half years. We've created the first hardware prototype and expect the full-sized prototype with the software environment up and running to be ready by the end of this year.

while its highly scalable code benefits from the Xeon Phi accelerators of the Booster. This enables users to run both parts concurrently and speed up the whole application.

ECO EXASCALE

Energy efficiency is a key problem for exascale-computing. Although HPC necessarily requires large amounts of energy, the project is working to minimise energy consumption on all fronts. DEEP will take advantage of novel many-core chip technologies and software-aided cooling technologies to improve current cluster energy efficiency by an order of magnitude. The team has created energy efficient infrastructures, cooling the DEEP system directly with warm water using a cold plate solution, and it employs energy aware system management with power monitoring capabilities.

In addition to ensuring the system is environmentally friendly, its real-world applicability plays a central role: "Simply because a future exascale machine will be a thousand times faster than current machines does not mean the applications will run as quickly. Therefore, we are testing our concept by running scientific applications on the DEEP systems, and allowing time to optimise the applications to the hardware architecture," Suarez elucidates. The team is running six applications, one of which simulates brain function and could help scientists to develop new treatments for Alzheimer's disease.

MEETING SOCIETAL CHALLENGES

DEEP, together with its successor project Dynamic Exascale Entry Platform – Extended Reach (DEEP-ER), is confronting all the challenges facing exascale, and finding solutions to important societal challenges in the process. The project team is well on its way to achieving its goal of building a prototype hardware and software supercomputing system with the potential to reach 0.5 petaFLOPS, with the actual machine that will come from it capable of 100 petaFLOPS.

DEEP-ER

Dynamical Exascale Entry Platform – Extended Reach (DEEP-ER) started in October 2013. It will extend the Cluster-Booster Architecture of DEEP. DEEP-ER enables parallel inputs/outputs to be highly scalable and improves high system resiliency, two of the most significant issues in exascale computing. The prototype from the project will be used to build a supercomputer capable of 500 PetaFLOPS.

INTELLIGENCE

DEEP & DEEP-ER

DYNAMICAL EXASCALE ENTRY PLATFORM & DYNAMICAL EXASCALE ENTRY PLATFORM – EXTENDED REACH

OBJECTIVES

- To develop a novel, exascale-enabling supercomputing architecture with a matching SW stack and a set of optimised grand-challenge simulation applications
- To extend the Cluster-Booster Architecture with a highly scalable, efficient, easy-to-use input/output system and resiliency mechanisms

KEY COLLABORATORS

Barcelona Supercomputing Center (BSC) • **Eurotech** • **Intel** • **Forchungszentrum Jülich** • **Leibniz Supercomputing Centre** (LRZ) • **ParTec** • **University of Heidelberg**

PARTNERS

Astron, The Netherlands • **BSC**, Spain • **CINECA**, Italy • **Eurotech**, Italy • **Forchungszentrum Jülich**, Germany • **Fraunhofer Institute for Industrial Mathematics ITWM**, Germany • **Inria**, France • **Intel**, Germany • **KU Leuven**, Belgium • **LRZ**, Germany • **ParTec Cluster Competence Center**, Germany • **University of Regensburg**, Germany • **Xyratex Technology Limited**, UK • **École polytechnique fédérale de Lausanne (EPFL)**, Switzerland • **Mellanox Technologies**, Israel • **CERFACS**, France • **CGG**, France • **The Cyprus Institute**, Cyprus • **University of Heidelberg**, Germany

FUNDING

EU Seventh Framework Programme (FP7)

CONTACT

Estela Suarez
Project Manager

Jülich Supercomputing Centre
Leo-Brandt-Straße
52428 Jülich
Germany

E e.suarez@fz-juelich.de

Sabrina Eisenreich
PR Manager Project

T +49 89 35831 8887
E sabrina.eisenreich@lrz.de

www.deep-project.eu
www.deep-er.eu

ESTELA SUAREZ is Project Manager of both DEEP & DEEP-ER. She works at Jülich Supercomputing Centre. She holds a PhD in Physics and has collaborated on high-energy astrophysics projects at the University of Geneva, Switzerland.