PRACE

PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE

PRACE days 19

Executive Summary

PRACEdays19

PRACEdays19 Executive Summary

www.prace-ri.eu

Editor-in-Chief
Marjolein Oorsprong – PRACE Communications Officer

Proofreading
John Clifford, PRACE aisbl
Silke Lang, PRACE aisbl

Article Editors
Harry O'Neill, Insight Media
William Davis, Insight Media

Graphic Design
Crystal Turner, Insight Media

Published by Insight Media Group, Bristol, UK

The Partnership for Advanced Computing in Europe (PRACE) is an international non-profit association with its seat in Brussels. The PRACE Research Infrastructure provides a persistent world-class high performance computing service for scientists and researchers from academia and industry in Europe. The computer systems and their operations accessible through PRACE are provided by 5 PRACE members (BSC representing Spain, CINECA representing Italy, ETH Zurich/CSCS representing Switzerland, GCS representing Germany and GENCI representing France). The Implementation Phase of PRACE receives funding from the EU's Horizon 2020 Research and Innovation Programme (2014-2020) under grant agreement 823767.

Published: November, 2019
ISBN: 9789082942231

Foreword

Serge Bogaerts, Managing Director of PRACE, Chair of the EuroHPC Summit Week 2019 and PRACEdays19 Organisation and Programme Committee

This year, PRACE held its sixth edition of the PRACE Scientific and Industrial Conference – PRACEdays19 – as part of the EuroHPC Summit Week 2019. The EHPCSW was held in Poland in the city of Poznań, hosted by PRACE Member the Poznań Supercomputing and Networking Centre (PSNC).

This year's EuroHPC Summit Week provided a comprehensive and rich programme for the various European HPC stakeholders to gather information and discuss the latest developments in the world of HPC. The event underlined in particular the key role that PRACE and its Members play in the HPC ecosystem and gave the opportunity to show how it may support the newly-created EuroHPC Joint Undertaking. We would like to thank all of the organisers for their efforts to make the EuroHPC Summit Week an inspiring event.

More than 360 of you joined us to listen to and take part in our plenary sessions, our parallel sessions focused on specific domains, and the various discussions that took place about the pressing topics that science, industry and the broader HPC landscape face.

As always, we spent a lot of time bringing together the speakers for our plenary sessions, and this year brought us a selection of fascinating talks, each with their own take and perspective on the world of HPC at large. Marina Becoulet provided a brilliant overview of the role that HPC is playing in the fusion energy project ITER, and Julien Harou demonstrated how HPC can be used to help society better manage valuable resources such as water and energy.

Our eight parallel scientific tracks and one industrial track covered topics including physics, life sciences, the environment, materials and chemistry, and computational science. It was exciting to hear from so many of you about the pioneering work that you have been doing using some of the most powerful computers in Europe, and I'm sure these sessions will have provided inspiration for many more of you for your own research.

The User Forum provided all of the HPC users present with the chance to learn about the scientific and industrial achievements made possible by the PRACE Research Infrastructure and, more importantly, the very positive feedback of PRACE users as collected by the PRACE User Forum in their recent survey. The panel discussion similarly provided a lot of food for thought, with the conversation covering many topics including the impact of big data, and the implications of the new EuroHPC Joint Undertaking.

We were very happy to present Debora Sijacki with this year's PRACE Ada Lovelace Award, which celebrates the work of women working with HPC. Her presentation of her astrophysical studies of galaxy formation and supermassive black holes was fascinating, and provided a clear picture of how far High-Performance Computing has helped the field progress in recent years.

Organising and participating in the EuroHPC Summit Week this year has been an extremely positive and exciting experience for us all, and we hope that every one of you who attended went away with a renewed optimism about Europe's aspirations for the coming years. We would like to thank you all for attending and are looking forward to seeing you all in Porto, Portugal from 23 to 27 March 2020.

PRACEdays19

Table of Contents

Foreword 3

Opening ceremony - Welcome to Poznań 6

PRACEdays19: Summary 8

EuroHPC explained - Thomas Skordas 11

Managing resources with HPC - Julien Harou 14

Harnessing the power of fusion - Marina Becoulet 17

EDF and HPC - Ange Caruso 18

Diversity and equality in HPC 20

HPC Ecosystem Summit - Florian Berberich 23

Winner of the PRACE Ada Lovelace Award 2019 - Debora Sijacki 25

Meeting the future demands of EuroHPC - Erik Huizer 28

PRACEdays19 User Forum summary 30

Tribute to Dimitri Komatitsch 32

PRACEdays19 panel discussion summary 34

BDEC - Big Data and Extreme-scale Computing 40

Poznań, Poland, 2019

Welcome to Poznań

PRACEdays19 Opening Ceremony

PRACEdays19 began on Tuesday, 14 May in Poznań with a number of opening remarks. Hosted by the Vice-Chair of the PRACE Council **Janne Ignatius**, of CSC, Finland, these words would set the tone for the rest of the conference.

Serge Bogaerts, the Managing Director of PRACE, offered a warm welcome to all who had gathered for the opening ceremony. This year marked the advent of a changing European HPC landscape with the onset of the EuroHPC Joint Undertaking (EuroHPC JU), he said, and he expressed hope that the conference programme, organised by PRACE, EXDCI-2 and ET-P4HPC, would suitably honour the occasion.

He highlighted the fact that the programme this year was organised with the three pillars of the HPC ecosystem in mind – infrastructure, applications, and technological development. He spoke of the diverse range of activities that had been arranged to keep conference attendees interested, including keynote speeches, parallel sessions focused on specific domains, a presentation of the PRACE Scientific Case, an international session organised by BDEC, the PRACE User Forum, a panel discussion, and much more. He also mentioned the training sessions on offer to provide students with hands-on experience of working in HPC.

Next up was Norbert Meyer (Supercomputing Division Director, Poznań Supercomputing and Networking Centre), who spoke about the city and how it has contributed to research in Poland and beyond. Poznan is a city of history, he said, but also one that looks to the future in terms of technology and the arts. It is, first and foremost, an academic city with eight public universities, and its 120 000 students make up around 20

Serge Bogaerts, PRACE MD, welcomed a packed hall to Poznań, saying that this year marked the advent of a changing landscape for HPC

per cent of its population. This academic core of the city is where PSNC first began delivering its infrastructure and services and, 25 years on, the organisation is going from strength to strength. Nowadays, its main quarters are at the Polish Optical Internet Research Centre, and it provides more than just HPC. PSNC ensures that the scientific community has access to world-class IT infrastructure, delivering networks to universities and R&D institutions, at the same time as delivering HPC and more recently cloud computing.

PSNC is part of a large consortium of universities in institutions in Poland that have collectively built a national academic fibre optic network, with around 6000 kilometres of cable within Poland and a further 2500km connecting the country to major hubs throughout the rest of Europe, including CERN and the GÉANT network. This is just one of the actions that is pushing Poland to become a major European force in the world of computational sciences. Norbert Meyer concluded his speech by wishing everyone a good time in Poznań (despite the weather), and expressing his hope that everyone present would have a fruitful and productive conference.

Thomas Lippert, Chair of the PRACE Council was next to speak, providing a summary of what PRACE does and how it serves High-Performance Computing in Europe. First and foremost, PRACE provides open access to leading HPC systems for EU researchers. This access is allocated through its rigorous peer review process, which ensures that only the best science gets carried out on the best machines. A variety of computing architectures are available to meet the algorithmic needs of different groups and domains.

PRACE does not just support academic research though, said Lippert, as he highlighted the SME HPC Adoption Programme in Europe (SHAPE), which equips European SMEs with the expertise necessary to take advantage of the innovation possibilities opened up by HPC. PRACE also provides more than just access to supercomputers. It supports HPC users in other ways, such as through training by the PRACE Training Centres, Seasonal Schools, and its Summer of HPC programme for students. Support in creating and developing applications so that they are ready for use on the largest computers

Session chair: Janne Ignatius of CSC

Speaker: Norbert Meyer of PSNC

Speaker: Thomas Lippert Chair of the PRACE Council

is provided by the High-Level Support Teams, and through the domain-specific Centres of Excellence.

Thomas Lippert then provided the audience with an overview of the Tier-0 machines available to PRACE users, and some of the research highlights that have been carried out on them in recent times. For instance, the MareNostrum machine at BSC has been used to model gravitational wave signals so that they can be identified when detected from events such as merging black holes. More work on astrophysics has been occurring at SuperMUC hosted by GCS@LRZ, Leibniz Supercomputing Centre (LRZ), where researchers are trying to make sense of the data being gathered by the Event Horizon Telescope, which produced the first ever image of a black hole earlier this year. MARCONI at CINECA has been used in the study of carbon nanotubes as excitonic insulators, while researchers using Piz Daint at CSCS have significantly improved the representation of the diurnal cycle, heavy precipitation and clouds in climate simulations. At GENCI, the Curie machine has been used by Renault for crash optimisation, helping to push the competitiveness of European industry. Similarly, the Hazel Hen machine at GCS@HLRS has been used by SME *Ingenieurbüro Tobias Loose* in a SHAPE project to simulate multi-layered welding lines, providing the company not only with significant knowledge and experience in HPC, but also the high commercial benefit of better cost estimates.

Finally, the JUQUEEN machine hosted by GCS@FZJ has been instrumental in solving a long-standing riddle in particle physics of why protons and neutrons have slightly different mass. 80 years after the experimental detection of this, it has been simulated using quantum chromodynamics. These projects represent just a small slice of the excellent science being carried out on Tier-0 machines in Europe, said Thomas Lippert, and it is thanks to PRACE's Scientific Steering Committee that the organisation is able to adapt to the needs of researchers as new developments occur in the world of science.

"Researchers using Piz Daint significantly improved the representation of the diurnal cycle, heavy precipitation and clouds in climate simulations."

PRACEdays19: Summary
Bringing together the people of HPC

The sixth edition of PRACEdays took place in Poznań, Poland from 13 to 17 2019. As part of the EuroHPC Summit Week, this year's conference once again provided a carefully curated set of talks and activities about every aspect of the world of High-Performance Computing. The conference brought together all of the main players in HPC, hosted by PRACE's Polish Member, the Poznań Supercomputing and Networking Center (PSNC).

The initial welcome on Tuesday morning from Serge Bogaerts (Managing Director, PRACE) was followed by talks from Thomas Lippert (Chair of the PRACE Council) and Thomas Skordas (Director, DG CONNECT, Vice-Chair, EuroHPC Governing Board) about the shape of the European HPC landscape and the roles that both PRACE and EuroHPC will play in the coming years. Erik Huizer of GÉANT also provided an interesting insight into NRENs (European National Research and Education Networks) and how they will be essential for delivering the EuroHPC programme.

Tuesday morning's other session provided two talks from two very different areas of HPC. Marina Becoulet of the nuclear fusion project ITER gave an overview of how High-Performance Computing is being used in the project, and the variety of codes that are needed to overcome the numerous challenges posed by trying to sustainably contain and maintain a fusion reaction. Julien Harou, on the other hand, showed the audience how HPC can be used to help society better manage valuable resources such as water and energy.

PRACE then presented its own recently published Scientific Case for Computing in Europe 2018-2026. Núria López and Sinéad Ryan, the incoming and outgoing chairs of the PRACE Scientific Steering Committee, provided a vision of the future of computing in Europe, giving examples of achievements in each area of science and showing how each of these domains might benefit over the coming years through further investment in European HPC infrastructure.

The afternoon session brought a number of separate scientific tracks, providing more in-

Getting together and interacting - PRACEdays19 was a lively affair

Listening, taking note and presenting ideas - the plenary sessions at EHPCSW 2019 / PRACEdays19 gave all a chance to interact

depth talks around the use of HPC in specific domains. These took place throughout the week and included tracks on physics, life sciences, earth and environment, materials and chemistry, industry, and computational science.

The HPC Ecosystem Summit was also held, which hosted discussions on how to create a coherent HPC landscape including access to HPC resources, services for users and research. The summit helped to clarify the future roles of each actor in the field and was followed by an HPC Ecosystem Plenary and a number of HPC Ecosystem Workshops on Thursday afternoon, organised by EXDCI-2. Those workshops gave an overview of the European effort and results in HPC, with the presentation of the new Strategic Research Agenda from ETP4HPC, the analysis of the rich results of the EU FET HPC projects, and the upcoming works of EuroLab4HPC.

As well as this, PRACEdays hosted its first ever Equality and Diversity in HPC session, designed to raise awareness about some of the difficulties faced by under-represented groups in computational research, as well as to inspire participants just starting out in the field. Wednesday morning brought a plenary session hosted by the Big Data and Extreme-Scale Computing (BDEC) project. Rosa M Badia of Barcelona Supercomputing Centre (BSC) spoke about ELASTIC, a software architecture for extreme-scale big data analytics in fog computing ecosystems, and how it is being applied with trams in the city of Florence to improve traffic management, allow predictive maintenance and enhance the tramway operation.

The PRACE Ada Lovelace Award for HPC is presented each year to a woman who has made an outstanding contribution and impact on HPC in Europe and the world. This year's winner was Debora Sijacki of the University of Cambridge. Debora could not receive her award personally, but she presented her research work on how galaxies are formed in a video presentation.

A number of other awards were dished out this year at the end of the conference. Michele Buzzicotti of the University of Rome Tor Vergata

"PRACEdays hosted its first ever Equality and Diversity in HPC session, designed to raise awareness about some of the difficulties faced by under-represented groups in computational research, as well as to inspire participants just starting out in the field."

won the PRACE Best Scientific Presentation Award for his presentation titled "Energy cascade in rotating turbulent flows". The PRACE Best Industrial Presentation Award went to Oriol Lehmkuhl of BSC for his presentation, "High fidelity simulation of vortex induced vibrations for flow control and energy harvesting". Finally, the PRACE Best Poster Award was given to Martin Golasowski of IT4Innovation for his poster titled "Distributed environment for traffic simulations".

Two of the main events at PRACEdays each year are the panel discussion and the User Forum open session. This year was no different, with the panel discussion on "The Future of Scientific and Industrial Computing in Europe" bringing about some enlivening and engrossing debate on various issues, including the dawn of big data and data analytics, and the best ways to improve industrial use of HPC.

The User Forum open session provided a good insight into the needs of the users, underlined by the results of a survey carried out in spring 2019 among the principal investigators of PRACE Tier-0 projects.

On a sadder note, Vadim Monteiller gave a speech about the life and research work of his good friend and colleague Dimitri Komatitsch, who sadly passed away earlier this year at the age of 48. His contribution to the field of computational seismology has been huge, and he has paved the way for young scientists working in his field.

The keynote presentations were broadcast live on YouTube

Sinéad Ryan, Trinity College Dublin and Anders Kvellstad, Imperial College London

The EuroHPC Joint Undertaking State of Play and Next Steps

The EuroHPC Joint Undertaking (EuroHPC JU) is a €1 billion joint initiative between the EU and 29 European countries that has been established to develop a world-class supercomputing infrastructure and ecosystem in Europe. After 15 months of planning in 2017-2018, the initiative is beginning to take shape, and **Thomas Skordas**, Director of the European Commission's Directorate-General for Communications Networks, Content and Technology, Directorate C "Digital Excellence and Science Infrastructure", and Vice Chair of the Governing Board of the EuroHPC JU, was keen to update everyone about the progress that has been made.

Thomas Skordas provided delegates with the latest news of the EC's collaborative efforts that drive EuroHPC JU

So what will the EuroHPC JU actually do? First of all, it will acquire and provide a world-class petascale and pre-exascale supercomputing and data infrastructure for Europe's scientific, industrial and public users, matching their demanding application requirements, before the end of 2020.

In addition, it will continue to build upon the thriving big data and High-Performance Computing (HPC) ecosystem which already exists in Europe, strengthening its HPC supply industry while pushing to maintain its status as a world leader in HPC applications. According to the website of the EuroHPC JU[1], this ecosystem development will entail "supporting an ambitious research and innovation agenda to develop and maintain a world-class HPC ecosystem, exascale and beyond, covering all scientific and industrial value chain segments, including low-power processor and middleware technologies, algorithms and code design, applications and systems, services and engineering, interconnections, know-how and skills for the next generation supercomputing era".

Providing HPC infrastructure in Europe

Mr Skordas used the first half of his talk to outline the EuroHPC JU's plans to improve Europe's HPC infrastructure, giving an overview of the HPC machines that will be acquired in 2020. These will be the first HPC machines funded and procured by the EU and 20 of the 29 countries participating in the EuroHPC JU. Three new pre-exascale machines[2] will be funded, to be hosted at CSC – IT Center for Science in Finland, Cineca in Italy, and the

Barcelona Supercomputing Centre in Spain. Each of these new machines will be around ten times more powerful than the current most powerful supercomputer in Europe (Piz Daint in Switzerland). Five more sites in Bulgaria, Czech Republic, Luxembourg, Portugal and Slovenia will host petascale machines to consolidate this new HPC infrastructure.

Overall, these new machines will help to provide a four-to-five-fold increase in the supercomputing power currently available in Europe. They will be interconnected with existing national supercomputers and will be made available throughout Europe to public and private users for developing leading scientific and industrial applications. The procurement process for the pre-exascale machines is now underway, and it is expected that the new machines will be installed and become operational by the end of 2020. The timeline for the petascale machines is slightly quicker, with these expected to be up and running by mid-2020. Access to them will be provided through a similar process to PRACE's own peer review process, and it is envisaged that over time, some (up to 20%) of these computing resources will be offered on a pay-per-use basis to support European industry for commercial purposes.

A world-class supercomputing ecosystem

The plans for the EuroHPC JU are ambitious, and they go well beyond just providing Europe with state-of-the-art HPC machines. The idea, Mr Skordas highlighted, is to support the coordinated development in Europe of new technologies with the potential for impact in many different areas. For example, the European Processor Initiative is already in the process of developing a European low-power microprocessor technology that will make Europe autonomous in this strategic field that is essential for many emerging markets, such as future driverless cars, data servers and smart phones. These processors will be used to develop and build the first exascale machines based on European hardware.

Another major target for the EuroHPC JU will be to support the development of leading applications in a large number of social, scientific and industrial domains, in order to maintain Europe's leadership in existing HPC application sectors, as well as to develop novel ones in high added-value areas such as personalised medicine, bioengineering, precision agriculture and smart cities.

Mr Skordas outlined a number of R&I calls related to the development of HPC software technologies for scientific communities, as well as calls for innovative HPC applications that will benefit industry. The EuroHPC JU will also build on the success of the HPC Centres of Excellence, helping to serve as many scientific communities as

possible by helping to update and maintain their most essential codes.

The HPC ecosystem in Europe will be developed around national HPC Competence Centres, of which there will be one in each country participating in the EuroHPC JU. These will be one-stop-shops for providing on-demand services to public and private users, providing access to tools, equipment and everything that is needed to help industry in Europe to leverage the power of HPC.

Planning for the future

It is important for everyone in Europe to understand that the scientific and industrial challenges that HPC users tackle are closely linked to the societal challenges that affect us all, such as the need to fight climate change and make advances in personalised medicine. EuroHPC will make a significant contribution to meeting these challenges both now and in the future, and is already planning for the advent of new post-exascale computing architectures that will most likely use AI-based (neuromorphic) chips and even quantum computing.

In the EU's next budgetary period (2021-2027), three European programmes will be essential for the success of the EuroHPC JU and the wider strengthening of HPC in Europe. The first of these is the Horizon Europe research framework programme, which will support R&D efforts towards new and emerging HPC architectures and applications for science and innovation. The second of these programmes is the Digital Europe programme, which will focus on building the capabilities Europe needs, such as acquiring next-generation HPC machines, supporting further the development of the national HPC competence centres and their networking, developing large-scale HPC applications, and fostering skills and training. Finally, the Connecting Europe Facility programme will provide greater connectivity between Europe's supercomputing centres.

The overriding message from Mr Skordas was that the EU's EuroHPC strategy is set to provide European leadership in HPC technology supply, to facilitate increased HPC access, and to support the development of the necessary skills to enable a wider HPC uptake by the next generation of HPC users. HPC will be one of the key critical technologies for driving the digital transformation of Europe, and the EuroHPC initiative is providing an ambitious and long-term plan for Europe to make this a reality.

References

1. https://eurohpc-ju.europa.eu/
2. EuroHPC JU Governing Board Decision June 2019. https://eurohpc-ju.europa.eu/documents/Selection_of_the_HE_for_the_Precursors_to_Exascale_Supercomputers.pdf

Disclaimer: The views expressed in the article are the sole responsibility of the author and in no way represent the view of the European Commission and its services.

Julien Harou spoke inspiringly about how HPC can and will be used for the greater good

HPC helping society manage its valuable resources better

Julien Harou wowed the PRACEdays19 audience with his plenary talk in which he demonstrated how High-Performance Computers can be used to help society better manage valuable resources such as water and energy. We spoke to him afterwards to get his opinion on a number of topics related to the real-world applications of supercomputers.

Q&A

Could you summarise how HPC will help us make the best decisions in terms of our shared resources in the future?

Firstly let's recall that human development depends on a series of networked resource systems like water, energy, food and transport. Getting these human-resource system relationships right fundamentally influences how sustainable and prosperous human civilisation can be.

HPC allows you to consider the full richness of resource system decision-making problems that involve both human interventions (like infrastructure investment or changes in management) and the uncertain supply and demand of the resources those interventions are meant to harness. The issue is there are deep uncertainties

– we don't really know many of the probabilities of future conditions because there are too many interdependent factors causing changes. Luckily, HPC allows you to use computers to help you cost-effectively think about and experiment with possible situations in such an uncertain future.

Nowadays, with information so easily available, many different constituencies have to be satisfied when big costly public decisions are made, and even with private decisions if they are high-impact. Therefore, quantifying the many different stakeholder interests and agreeing on reasonable trade-offs between them is laborious work. For large systems, you just cannot do that with a

"HPC allows you to use a computer to help you think about possible situations in a future in which you don't know what will happen."

spreadsheet on your laptop. If you want to look at multiple facets of the issue and search through many interdependent options under many plausible future scenarios, then HPC is a needed tool.

This relates to the issue of complexity. Imagine you're interested in large-scale railway systems for example, you have to look at many interdependent and connected elements, along with the complexity of how humans interact with the infrastructure services. Then you realise weather plays a big role, and train prices etc. The tendency is to have to consider several factors when designing real world systems. This leads to large simulations which need large computers. This is especially true if decision-makers are aware there are lots of choices, and they want to ensure they select a good design. In that case you need to connect the exploratory simulation tools to search engines (optimisation), and that's computationally intensive. Finding credible balances of benefits if you are working with several stakeholder groups inevitably means you will need to demonstrate that all good options have been tested. Computers are needed for that; there is no other way to do it really.

To do the work that needs to be done in the future, are we going to need much more computing power? Where is this going to come from?

I don't really know. I am just an HPC user so I can't really answer that accurately. In our group at the University of Manchester we use what is available today, and progressively we are realising it is not really enough. We may be able to do more with less, and that is part of this new science, but it is clear that if decision-makers looking at major infrastructure investments want to consider uncertainty, multiple interests, and the complexity of real-world behaviours and interdependencies, then this points to larger and more simulations, i.e. more computing.

For example, if you are investing in a big piece of infrastructure costing billions, you want to make sure the decision you make will perform acceptably over at least 100 plausible futures. But interacting with stakeholders may reveal that really 1000 scenarios are needed to acceptably sample future conditions given the uncertainty, so that will be the driver for more computing – assessing how today's investments might work under tomorrow's possible conditions. For example, if you are building a motorway, you have to think about what technology people will use in 20 years. Surely that should impact today's motorway designs? That example is quite engineering-centric, but when you consider for example interconnected water-energy-food system changes, then social, environmental and economic

dimensions become relevant and their uncertainties must be considered.

With today's computing, analysts can already offer valuable insights and recommendations that are helpful for decision makers today. Already decision-making under deep uncertainty assessments are more advanced today than they were a decade ago. It's been remarkable to see how quickly society can integrate such new advanced approaches and tools. Adoption is mostly driven by forward-thinking organisations like utilities or certain well-informed government actors. In the future such advanced system-scale investment assessments will become standard, a bit like cost-benefit analysis was last century.

We cannot predict how people are going to behave, but people's behaviour has a huge impact on the things we design. Can we ever predict behaviour?

That's a good point, and a tough one. Human choice is the one thing we will probably never be able to simulate very well. That being said, with classical methods from economics and social simulation we might understand just enough about motivations and behaviours such that the integrated intervention assessments are useful.

Many academics today are trying to achieve a new multi-dimensional engineering-economic science, where people and institutions are not reduced to single objective optimisers, where system-scale cumulative impacts of interventions are assessed, and where the trade-offs implied by today's decisions are quantified even if we don't know the future. Despite knowing we'll never fully understand humans or artificially simulate choices accurately, we still need to make decisions today about how to build the future, so this is a strong driver to use the best science we can and get on with it.

Who will have or should have access to HPC?

That's a big deal. Imagine in 5-10 years when the use of system-scale development intervention assessments will have become a routine way for wealthy decision-maker and stakeholder groups. to deliver better resource system services. At that point we are talking about tangible social benefits being generated on a regular basis by standardised planning and governance practices. It would be unfortunate if that were only accessible to wealthy countries.

FutureDAMS.org, one of Manchester's flagship development research projects, aims to make that type of computer-assisted systems analytics available to partners in Africa and Asia designing complex resource systems. We are working with our partners to enable them to become advanced users of HPC for development intervention decision-making. If we could find a way to make these approaches available to more countries, that would be great.

Access to supercomputing as a global equity issue is just appearing on the horizon. If making good decisions is enabled by system-scale analytical approaches that require supercomputing, then who has access to HPC becomes an equity issue. I think HPC service providers should take on some of the responsibility of ensuring that advanced computer support decision-making methods are available across social, gender and country lines. Being able to access decision-making services through the web will certainly help democratise them.

Are you optimistic about a future that will be shaped by supercomputers?

Yes, definitely. I didn't use HPC when I finished my own PhD 12 years ago, but now everyone in our group uses it. The new AI-based algorithms allow for efficient and effective searching over large, non-linear systems. It's become very real and practical now. The demand for interaction with system analysts to evaluate human-resource system changes is growing. Several academics I know are having trouble meeting the growing demands from society to deploy these services. These new types of services are enabled by high-performance computing, so it's a real HPC success story.

Computer models won't make better decisions themselves for a long-time, and perhaps never, but they can assist stakeholder coalitions to assess changes to human-resource systems. The real world is a very costly place to fail in. Bad decisions cost social, political and financial capital. If we can develop better insights and understanding by evaluating changes to essential resource services virtually, then we'll likely make less bad decisions in the real world. Ultimately the benefit of system-scale human-resource system design is socio-economic progress, and the cost of not doing will be measured in lost human opportunity.

Unlimited energy

Harnessing the power of fusion can bring about a new era of safe and almost-limitless energy for humanity. **Marina Becoulet** of the Atomic Energy Commission used her plenary talk to describe some of the challenges of first principle modelling and the role of HPC in the field of magnetic fusion research

Marina Becoulet of the AEC spoke about harnessing the power of fusion

In Greek mythology, the Titan Prometheus defied the gods by stealing fire and giving it to humans, an act that enabled progress and civilisation. Marina Becoulet, a senior research physicist for the Atomic Energy Commission (ACE), described to the audience how those working on the ITER Tokamak, the most complex machine ever built, are aiming to do something similar for humanity by providing it with the almost limitless power of fusion.

Fusion, the nuclear reaction that powers the Sun and the stars, is a potential source of safe, non-carbon emitting and virtually limitless energy. Harnessing fusion's power is the goal of ITER, which has been designed as the key experimental step between today's fusion research machines and tomorrow's fusion power plants. In stars, the most common fuel for fusion is hydrogen, and gravity creates the conditions needed for fusion energy production such as plasma confinement and extreme density and temperature. However, these conditions are extremely difficult to reproduce on Earth.

The goal of scientists working in the field of fusion is to develop a clean, safe and limitless energy source. The most developed concept so-far involves magnetically confined plasma heated by external sources (like RF waves or high-energy ion beams) to create conditions for hydrogen isotopes (deuterium and tritium) to fuse and release energy.

A new era in magnetic fusion started in 2007 with the construction of the experimental device ITER in the frame of a large international partnership involving the EU, the US, China, Russia, Japan, India and Korea. Now, magnetic fusion research is facing challenges in the development of high confinement plasma scenarios combined with a high level of fusion plasma control. In this respect, theory and predictive modelling play a crucial role for ITER and the possible future reactor DEMO based on magnetic fusion.

The biggest challenge in terms of HPC resources and numerical methods is so-called first principles modelling. The HPC challenges for this highly non-linear physics, involving 3D complex geometry and kinetic effects, can be extreme. The difficulty stems from the wide range of spatial and time scales that need to be covered. This complexity of magnetic fusion first principle modelling has resulted in the development of new numerical methods adapted to these highly non-linear problems. These demand large amounts of memory, a high level of parallelisation and long real-time modelling that can reach the scale of months.

Becoulet used her time in front of the PRACEdays19 audience to describe some of the specific challenges of modelling in fusion, and gave some examples of recent progress in the understanding of phenomena such as the small-scale instabilities responsible for turbulent transport and plasma confinement, and large-scale magneto hydrodynamic instabilities which define the operational limits for any given plasma scenario.

HPC will play an increasingly important role in our future

Ange Caruso of R&D, Électricité de France (EDF) spoke in one of the plenary sessions about how the company is using High-Performance Computing to improve many aspects of their business. We caught up with him afterwards to discuss this further.

Can you explain EDF's motivation for investing so much in HPC?

Our motivation is simple. We would like to improve the safety of our processes involved in the production of energy, mainly in the complex and dangerous domain of nuclear energy. The use of HPC allows us to guarantee safety to obtain margin in terms of performance in producing energy and in the cost of that production. Improving safety, performance and cost is important for EDF.

HPC helps us do this by providing a deeper understanding of the phenomena that exist inside our power plants and helps us modify their design to optimise their performance. HPC enables us to find the best ways to do these vital tasks.

EDF is not only using HPC to optimise the production of energy, however. HPC is also being used to help us understand the introduction of electricity generated by renewable sources into the electricity networks. Introducing renewables into the grid can lead to destabilisation and impact on the quality of electricity, and so we are able to counter this by getting a better understanding of how the grid will react. It helps us anticipate where the problems will arise more effectively by the use of simulations. If the simulation is good, then we can then go on to develop the physical solutions, which is far more cost-effective and efficient as we only develop solutions we know will work.

The same thing goes for optimisation. The storage of electricity is extremely complex. We can do it, but we must still maintain a balance between production and consumption. Production is very diversified – nuclear, solar, coal, etc. – and we are always looking to optimise quality, efficiency and cost. We want to know if it is more efficient and cost-effective to use nuclear rather than renewables at a particular time, for example, but we cannot

Q&A

know for sure in advance. Simulations can help us decide what is the best way to produce electricity at a particular time to optimise performance and cost, and to maintain equilibrium between production and consumption – a balanced grid. Simulations using HPC help us make these decisions in advance, like whether to build a nuclear power plant or plan for the introduction of renewable energy sources. But we can also do simulations in real time – and balancing the grid is a real-time challenge.

What do you see as the main challenges and opportunities we are facing in terms of the development and use of HPC?

The increasing application of artificial intelligence is having a major impact and will continue to do so, but if I had to identify one challenge that we need to address, then combining simulation with artificial intelligence is the big one. The challenge is how to introduce artificial intelligence algorithms into simulations to improve them. At the same time, we need to improve the artificial intelligence algorithms

Ange Caruso of EDF presented an optimistic future of HPC being used as an essential tool by industry

to deliver better simulations. The problem we have is that we have two communities – those doing physical simulations and those working on artificial intelligence algorithms.

Trying to get these two communities to overlap is very difficult. We have technical problems, of course, but we have to change the way these two groups of people think – when a physical simulation is being planned, we must think now what artificial intelligence can do to help. The two ways of thinking need to be more integrated. This is starting to happen more, but it is still a big technical challenge. We have to develop new methodologies, new algorithms to bring artificial intelligence into simulations, and this challenge is one of change management in the way people think and work.

What is the future for HPC and how will EDF be involved?

Today, HPC is a revolution, but in years to come it will become more and more classical. Like Windows was a revolution in computing some years ago and now we all know it and use it, HPC will go the same way as the power of computing facilities keeps increasing.

We will, of course, reach a limit to the power of the machines due to physics, but in the future HPC will be seen as an everyday tool we use without thinking, to simulate and analyse. The challenges in the future will not be using HPC, but using disruptive technologies such as quantum computers.

As a utility, EDF takes the new opportunities offered by more powerful computers and we follow this development and use what is available. But the revolution in HPC is not only fuelled by the users but also by the manufacturers of these machines.

Every year, for example, we discuss with all manufacturers the integration of their roadmaps into ours. We see this as necessary if we are to complete our roadmap. We cannot do that without integrating HPC and using the very best on offer.

Diversity and equality in HPC

PRACEdays19 hosted its first ever Equality and Diversity in HPC session. It was hosted by **Debora Testi** of CINECA and was designed to raise awareness about the difficulties faced by under-represented groups in computational research, as well as to inspire participants to think more about the problem.

A whole afternoon at PRACEdays19 was dedicated to the important issue of equality and diversity in HPC. Chaired by Debora Testi of CINECA, the session not only focussed on the presence and contribution of women in the various scientific domains, but tried to examine the notion of all generally under-represented groups and how barriers to their participation can be removed.

It was a lively discussion, which included many personal stories of both prejudice being faced and obstacles overcome, mainly by the women in the audience. One conclusion was that many of the procedures put in place to support the inclusion of under-represented groups are often perceived to be undermining the competence of those they seek to help and this is an issue that should be addressed immediately.

Debora Testi opened the session with an account of the situation she faced at the University of Bologna in early 1990, where out of 350 people specialising in ICT, just 30 were women, a situation made worse in the second year of study when more women left than men. The scientific domains do vary, with some having greater female participation than others. Meanwhile, women are less likely than men to continue studies to doctoral level and are still under-represented in the field of Information and Communication Technology, engineering, manufacturing and construction, while the number of women in all STEM subjects is not growing as fast as it is with men. Efforts to improve this are, however, starting to pay off. The percentage of women applying to Tier-0 calls, for example, was below 20%, but is improving and now stands at over 50%. Women in HPC in CINECA stands at 20% out of 6 500 users. Figures do vary depending on the HPC domain. Bioinformatics is faring well, but engineering is still under-represented.

Testi concluded that despite all these statistics from all over the world being available, there are not many solutions – and none are proven to be effective. Policy change takes time to make an impact. Change is happening for the younger generation, but it is a slow process. Specific examples include better conditions for women with small children and improvements are being introduced by HPC centres in universities, but more could always be done.

Núia López and Rosa M Badia

"Women are less likely than men to continue studies to doctoral level while the number of women in all STEM subjects is not growing as fast as it is with men."

Núria López of ICIQ followed by providing an academic perspective of women in STEM, acknowledging that diversity is much more than talking about women, although a women-focused perspective does highlight common issues.

One particular problem is the number of women in senior academic positions, which is not good, with a common issue being that women scientists are often lost over time. According to a L'Oreal Foundation report, there are often regional and cultural influences although a major contributory factor is an unconscious bias in those who make decisions on hiring. Gender stereotypes emerge early and influence children's education, both in terms of how boys and girls perceive themselves, but also in how they are treated.

Núria López pointed out that it does not have to be like this and there are examples of things changing. The issue is being addressed in many universities with professors being taught how to recognise bias and deal with it. Smart games designed for boys and girls under six make a difference, but often things

Debora Testi, Member of the PRACE BoD; CINECA

change as children get older and stereotypes are reinforced. Núria López concluded by calling for more exposure of HPC in schools.

Rosa M Badia presented examples from the Barcelona Supercomputing Centre of how the organisation is proactively promoting diversity across the organisation, in particular its organisation of the conference Advances in Computational Biology (www.iscb.org/advcompbio2019). The conference aims to foster collaboration amongst women scientists focusing on systems biology, artificial intelligence and HPC applied to biology. Although it is open to all, "for a change, the organisers and the presenters will all be women," announced Badia.

"The objective here is to draw attention to the existence of a gender bias that negatively impacts on women scientists and to visualise and promote research done by women," she continued. Topics at the conference, which takes place in Barcelona on 28 and 29 November 2019, include learning from biological sequences, when computational biology meets medicine and how machines are speeding up research.

Badia also spoke about the initiatives for young researchers at BSC, using the example of Bioinfo4Women as a channel for supporting outstanding young female bioinformaticians that aims to create an international network of women working in bioinformatics to amplify the impact they have. BSC offers mentoring for women working in this field as well as training plans. Finally, Badia placed this work into context in terms of the diversity of the people at BSC. While the centre is certainly diverse in terms of the nationalities of the people it employs, with staff coming from 48 countries in total, there is still some way to go in terms of the gender balance, where only 24 per cent of researchers are women – 128 to 405 in favour of men. Countering this, BSC has established a Gender and Diversity Equality Plan.

"It is not left to the under-represented who must come forward to share the issues they face, but for all involved to be proactive in finding out what issues everyone faces"

Lee Margetts spoke about his experiences at the University of Manchester, where there is a strong culture of improving equality and diversity through policies, practices, action plans and culture. The university is very active in promoting diversity and equality through training sessions for staff, including bias training, interviewing, and recognising inequality. The university has a directorate for equality and diversity, so the direction of travel towards a more representative place of learning and work, comes from the top and stems from the UK's Equality Act. The act defines nine areas of discrimination as unlawful – and gender is only one of these.

Margetts believes that UK universities take equality and representing diversity very seriously. On the one hand, it can be costly if they do not, while on the other, the legal systems in place encourage individuals to consider ethical issues carefully.

There is still room for improvement, however, and organisations in the UK still need to consider their weaknesses – this is what happens in Manchester and from this exercise a six-point policy has been developed through which it is not left to the under-represented who must come forward to share the issues they face, but for all involved in the organisation to be proactive in finding out what issues everyone faces. Groups, workshops, staff consultations, student and union bodies are all on board with the university's mission and it makes the place more attractive to work, study and live in.

In terms of gender, Manchester has roughly a 50/50 split, while professors keep records of interviews and attend training sessions to recognise bias in terms of gender, sexuality, disability and other areas of representation. Margetts concluded by saying that the pathway to equality and diversity is not through activism, it is through policy, practice, action plans and culture.

Quotes from the floor

Following the talks about equality and diversity in HPC, a number of delegates attending the workshop spoke of their own experiences and of what they believe needs to happen if we are to continue improving how under-represented groups are better served by our institutions working in this sector. Here is a flavour of the views expressed.

"Feeling out of place is not necessarily a gender issue. As a student from South America I feel like an imposter when studying in Europe. You need to feel sure of yourself and keep going."

"Be a colleague. Support colleagues. Make sure that the first voice to ask a question is a woman as if you don't they will be less likely to ask others. Small steps."

"In Brazil there is a positive bias towards black people. Why? The answer is that if you give people a chance and let them in, then people get used to mixing together and it becomes the norm. This has been beneficial in Brazil – correcting the unconscious bias."

"I worked with a person with a severe disability. He couldn't hear or see. When he was in the lab, the lab was very tidy as people made sure there were no hazards in his way. The message here is don't be scared to help if people need it. He was teaching us."

"Bias is not just gender – moving from one country to another presents challenges."

"What can each of us do? Become aware of your bias and be an ally of equality."

"Through training, I now recognise the unconscious bias."

"Positive bias is detrimental. We should not be obliged to employ women as it corrals women into groups and doesn't integrate them."

"We don't want to make ghettos of women in science. Groups of women talking about women in science are a waste of time."

"Choose where to work. If you are working in a company that is biased to a particular group, leave."

"Inclusion is much more fun. And the more you have the easier it gets."

The HPC Ecosystem Summit

The recently founded EuroHPC Joint Undertaking will permit the EU and participating countries to coordinate their efforts and share resources, with the objective of deploying in Europe a world-class supercomputing infrastructure and a competitive innovation ecosystem in supercomputing technologies, applications and skills. At this year's EuroHPC Summit Week, the HPC Ecosystem Summit was held to discuss the creation of a coherent HPC landscape including access to HPC resources, services for users and research. We spoke to the chair of the session **Florian Berberich** of Jülich Supercomputing Centre, Germany, about how he sees the development of this ecosystem panning out in the future.

Could you describe the European landscape of HPC as it exists at the moment and the different actors in the ecosystem?

In Europe, the HPC strategy is based on the three pillars of technology, infrastructure and applications. The technology pillar is obviously aimed at developing European technology, and includes actions such as the European Processor Initiative (EPI), which is developing European microprocessors for future supercomputers. The Future Emerging Technology Projects for High Performance Computing (FETHPC) projects and the European Technology Platform for High Performance Computing (ETP4HPC) are also a big part of this pillar.

The infrastructure pillar is composed of PRACE and GÉANT. PRACE provides computing resources to users, while GÉANT provides high-speed networks between the computing sites and between the users. The various national research and education networks (NRENs) are also involved in this infrastructure pillar.

Finally there is the application pillar, which is mainly composed of all the Centres of Excellence, of which there are now eight. There is a project called FocusCoE that coordinates these, pushing for joint training, joint trademarks of excellence in Europe, and other things. This is the current landscape.

Besides the three classical pillars, data infrastructure will become more and more important in order to handle the growing amount of data from large-scale scientific infrastructures such as CERN and SKA. An HPC infrastructure tailored to treat these large amounts of

Florian Berberich chaired the important HPC Ecosystem Summit

data will also be indispensable for machine learning and artificial intelligence. The data infrastructure is involved in part of all three pillars and is provided jointly by PRACE and EUDAT, which share and preserve data across borders and disciplines, through a common model and service infrastructure for managing data.

Last year, the EC created a new structure, the EuroHPC Joint Undertaking, which will get funding from the EC and EU Member States. EuroHPC will procure and prepare pre-exascale systems, exascale systems, and beyond. This process has started already. But it will also be involved in various other activities, and will cover all three of these pillars in some respect.

With the formation of EuroHPC, where does PRACE fit into the changing landscape?

PRACE was founded in 2010, and preparation for it began all the way back in 2004, so it has an enormous amount of experience to draw from in providing computing resources to industrial and academic researchers. This experience in providing HPC access based on scientific excellence through a rigorous peer review process is at the core of PRACE's activities. The difference between EuroHPC and PRACE is that EuroHPC will actually own the computers, but the way they are used needs still to be discussed and decided. It is very likely that at least part of the resources will be distributed also by a pan-European peer review process. So PRACE has a lot of expertise to offer in that respect.

PRACE also provides a wide range of training facilities and resources, and these will eventually synergise with EuroHPC. User support is essential for ensuring that the machines that exist in Europe are used properly.

What do you see as the main obstacles for HPC in Europe in the future?

The main obstacle I see is getting people to talk to each other and to collaborate, but overall I think the situation is quite good. Eventually, we want to create a single offer for HPC users that combines all of the activities and resources that Europe has to offer in one place, but this will not happen overnight. I think the most important thing is that we are aware of the individual initiatives and what they provide, and then to coordinate them. I can definitely understand that entities such as the Centres of Excellence want their own identity and trademark, and this can be achieved by discussing and agreeing on how to collaborate. Coordination and talking will help us to quickly create a situation where the user can get hold of all actors, including EuroHPC, in one place. For the moment, users need to know where to go, and this can be difficult because there are so many different projects. I think if we can arrange to put all of it in one place in a universal portal, it would solve some of the issues we face.

Access to training and support is high on the agenda of the EC, but users have experienced difficulties in finding it, partly because of lack of awareness. How can this be addressed?

We acknowledge that training is a key asset of the ecosystem, because it will not help to have HPC systems in the first 10 places of the Top 500 list without having the people that can use them effectively. If we don't have the experts able to make the codes and maintain them, these systems are useless. In order to get these experts, we have to increase engagement and training. There are many stakeholders working in this area, including the European Open Science Cloud and the Centres of Excellence, so it would be great if we could make things less confusing by having all of these resources in one place. We need a training portal that each actor can feed its own training into in a joint database, and also extract from.

If there was one call for action to come from your talks this year at the conference, what would it be?

To be open and share information. I think this is what I have learned during this conference. We know that all of these initiatives exist but we don't know what they are all doing. The readiness to collaborate I think is not an issue. We just have to collect this information and put it in one place. EuroHPC is an enormous opportunity. The European Commission recognised the need for an EU-level policy in HPC to optimise national and European investments, addressing the entire HPC ecosystem, and in 2012 published the communication 'High Performance Computing: Europe's place in a Global Race'. Since then the European Commission has increased its investment in HPC significantly and supported the three HPC pillars with various projects and initiatives. Thomas Skordas' talk showed this, and we now know that many billions of euros are being spent on this. It's a great opportunity for all of us, and we now have the chance to create a really sound HPC infrastructure.

Winner of the PRACE Ada Lovelace Award 2019: Debora Sijacki

The PRACE Ada Lovelace Award is awarded each year to a female scientist making an outstanding contribution and impact on HPC in Europe and the world, and serves as an inspiration for women who are at the start of their scientific careers. This year's award went to the University of Cambridge's **Debora Sijacki** for her astrophysical studies of galaxy formation and supermassive black holes through the use of high performance computers. Although unable to receive the award in person, she recorded a presentation to provide a summary of her work to the PRACEdays19 audience.

Debora Sijacki began her video presentation with a small tribute to Ada Lovelace, who achieved an extraordinary mental leap from calculation to computation. In Lovelace's own words, "the analytical engine [a proposed steam-powered programmable computer designed by Charles Babbage in 1837] might act upon other things beside numbers, where objects found whose mutual fundamental relations could be expressed by those of the abstract science of operations, and which should also be susceptible of adaptions to the action of the operating notation and mechanism of the engine". Sijacki argued that this notion still holds up in encapsulating what scientists today are trying to achieve with High-Performance Computers.

Debora Sijacki and Ada Lovelace

Moving on to her own work, which is looking towards the future of creating exascale simulations of our universe, Sijacki told the audience how, unlike observational astronomers who use telescopes to look up at the sky, HPC users working in the field of universe science are using supercomputers to explore the sky to try to understand the fundamental rules of nature.

HPC users working in the field of universe sciences are today addressing some of the biggest questions that remain in astrophysics. How did our universe form and evolve? What does it look like on a large scale? What are galaxies, why do they exist, and how are they distributed throughout the universe? More exotic questions are also being asked: why do supermassive black holes exist, and what relationship do they have with the surrounding stars and the host galaxies that they are embedded in? This is the area where Sijacki has excelled in answering with her own work. Simulations on our most powerful computers are an essential tool for astrophysicists.

Although the equations that we use to describe the fundamental laws of physics do not in themselves look overly-complicated, when one tries to use them to describe a sample of the universe, they become completely unmanageable to solve analytically, with strong non-linear effects occurring. This is the reason high performance computers are so important, which can produce simulated galaxies or patches of the universe that researchers can then analyse in great detail.

The universe is, in fact, rather strange in that it began with quite simple conditions of almost uniform temperature and density, but quickly developed into an unimaginably complex structure. The universe's composition (as we understand it) is also strange – most of the universe is dark i.e. most of the universe that reacts to gravity is dark. Only a small component of all matter is the ordinary matter made of baryons that we know so much about. The biggest segment of

the universe consists of dark energy, which drives its accelerated expansion.

Excellent observational data now exists that gives us precise initial conditions of what the universe was like near its inception. This can be used in simulations of the universe, evolving these conditions as a function of cosmic time, assuming that gravity is the fundamental force acting upon everything. One famous depiction of this was the Millennium Simulation, which used more than 10 billion particles to trace the evolution of dark matter distribution in a cubic region of the universe over two billion light-years in length. But in this type of approximation, the contribution from baryons in the form of gas and stars is neglected. When we look up at the sky at night, we do not see the dark matter that dominates the universe. What we see instead is gas and stars, and estimates put the number of hydrogen atoms that make up this visible part of the universe at around 10^{80}.

As well as the huge number of particles, a vast range of spatial scales are involved, from the scale of the whole universe, down through the galaxies and its components and finally the stars. Then, of course, the supermassive black holes that are now thought to exist at the centre of most galaxies have to be taken into account. So, to simulate this extremely complex visible component of the universe, clever discretisation approximations need to be adopted which can reduce the number of particles to a manageable amount, which is currently around 10^{12}.

State-of-the-art simulations that incorporate both dark matter and visible matter are currently working in the range of 100 megaparsecs, running on around 10 000 CPUs with well parallelised codes. These simulations use around 10-50 million CPU hours, creating petabytes of data that then need to be analysed to understand what is happening.

Sijacki has recently been involved in the Illustris project (www.illustris-project.org), an ongoing series of astrophysical simulations run by an international collaboration of scientists that aim to elucidate the processes of galaxy formation and evolution in the universe. One of the main successes of this project has been the accurate reproduction of galaxies as they are seen in the observed universe. Sijacki showed the audience a comparison between visualisations from the project and pictures of the real universe which demonstrated how accurate these simulations are becoming. A long standing problem in this field has

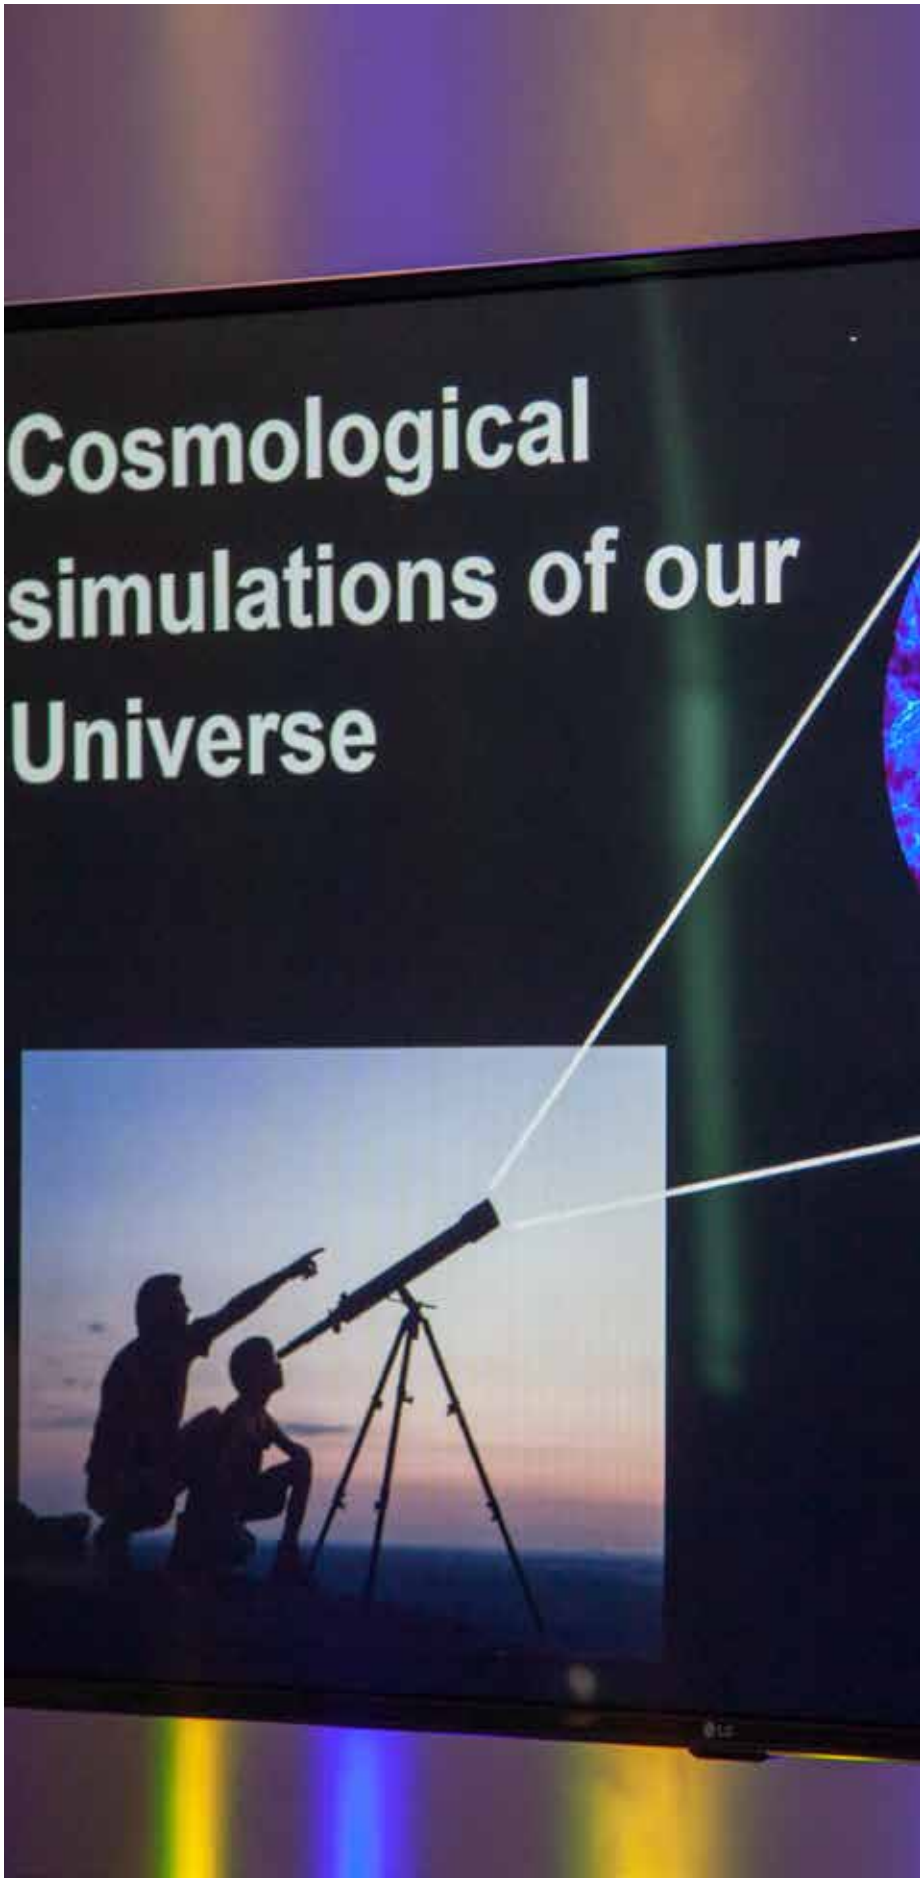

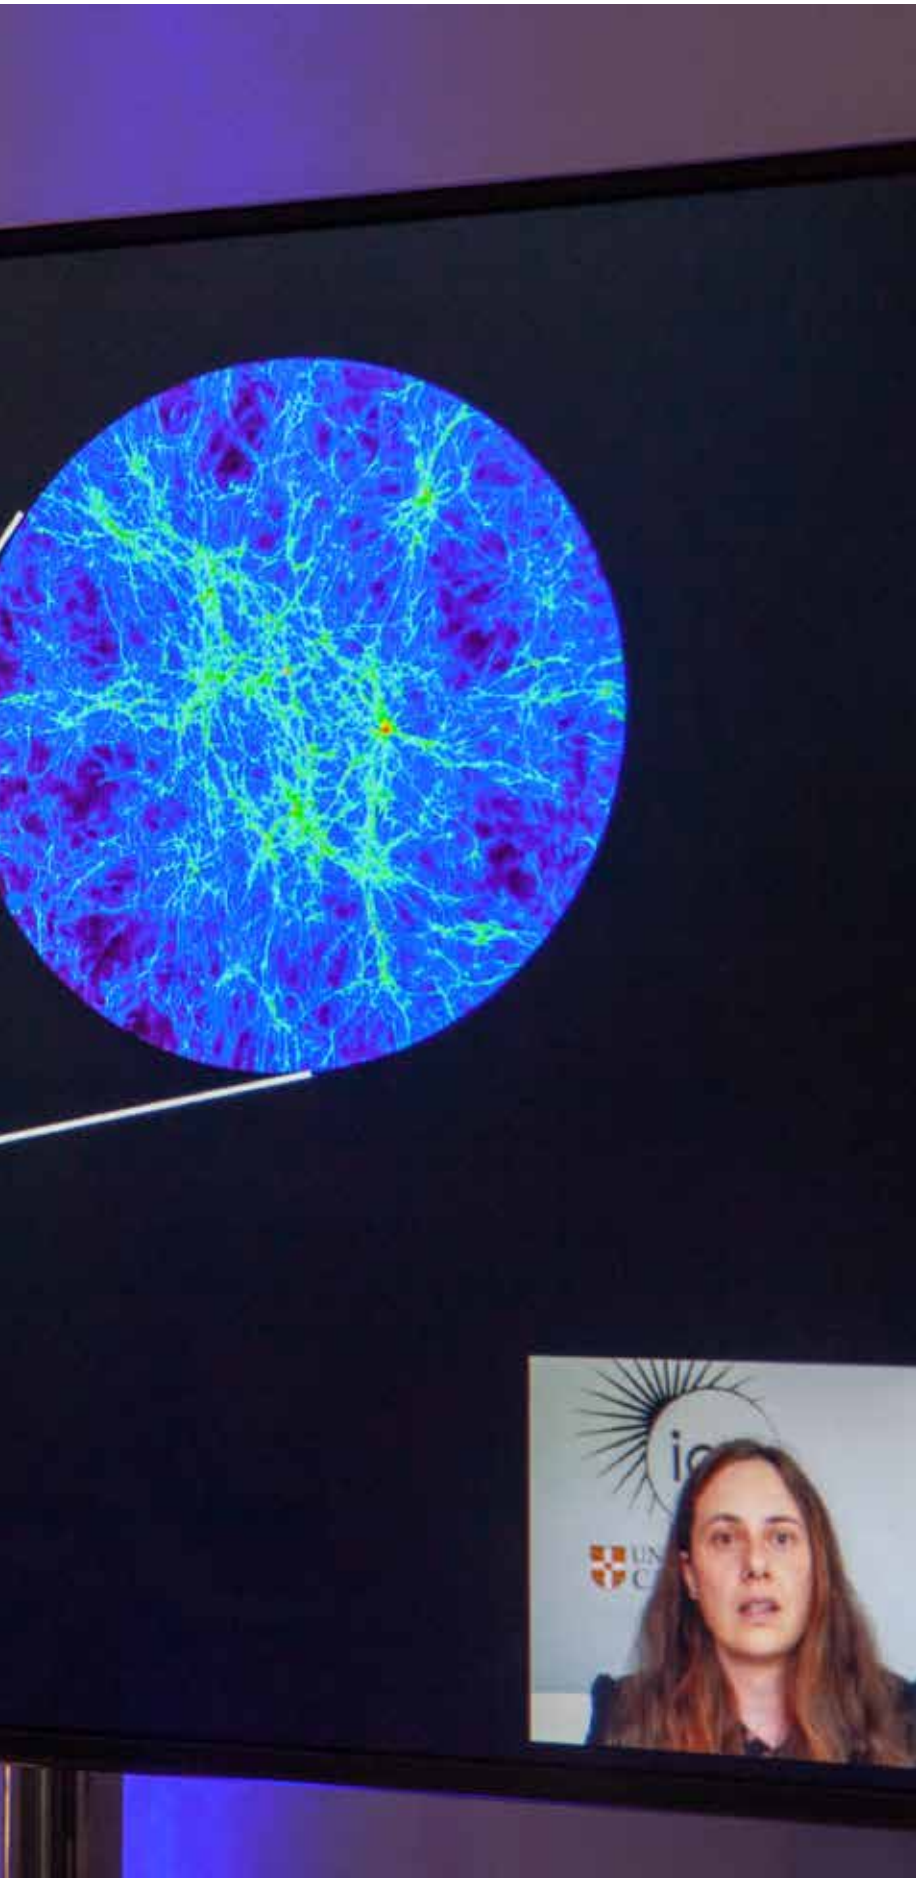

been reproducing the diversity of galaxy morphologies that are seen in the observed universe. The Illustris project is the first to have produced elliptical, disk-shaped, and irregular galaxies, and has also provided a surprising explanation for why this diversity of shapes exists. Sijacki explained how it is now thought that the supermassive black holes that exist in the cores of galaxies are a key factor for their morphology. A bigger stellar mass in a galaxy is correlated with a larger supermassive black hole at the core, and the energy from this black hole in turn tends to produce an elliptical morphology in the galaxy.

"A main success has been the accurate reproduction of galaxies as they are seen in the observed universe."

More recent developments in Sijacki's research involve the production of a new algorithm, alongside her former PhD student Michael Curtis, which allows them to zoom into the central regions of galaxies where the supermassive black holes reside and increase the mass and spatial resolution by many orders of magnitude. This makes it possible to capture the physical processes that happen in these regions far more accurately, and provides a bridge between some of the many scales that are involved in studying the universe.

These algorithmic developments have allowed Sijacki and her colleagues to begin investigating the jets that are sometimes seen being ejected from accreting supermassive black holes, known as active galactic nuclei. These jets can only be seen in simulations with the newly improved mass and spatial resolution, and their physical features such as vorticity, velocity and density can now be studied in this way and compared with real observations of jets. It is thought that the energy transfer from these jets to the surrounding medium could be an important heating source for these regions.

Ultimately, the goal for those working on simulations of the universe will be to breach all of the scales of the universe within a single simulation. Sijacki's work has demonstrated that with enough focus on algorithmic development and with powerful enough machines, this goal will be within our grasp in the years to come.

Meeting the future demands of the EuroHPC programme

The European National Research and Education Networks (NRENs) and GÉANT have together built a network and set of services to meet the demands of HPC. This will be further developed to meet the future demands of the EuroHPC programme by delivering a quality and targeted set of services that distinguishes NRENs from commercial Internet providers. We spoke to **Erik Huizer** at PRACEdays19 about how these actions will be implemented, and heard his thoughts about the future of HPC in Europe.

Q&A

Why are NRENs so essential for HPC to be successful in Europe?

The purpose of my talk at the event was to show that networks are an essential part of making sure that the EuroHPC project will succeed as access to HPC is also essential to this success. We have the unique situation in Europe where several countries don't build their own supercomputers but invest in supercomputers in other countries so they can use those instead. For this to work effectively, we need to provide access for researchers to be able to use those machines from their own desks, wherever they are based.

Research networks like GÉANT and the NRENs in Europe differ from commercial networks and are essential for researchers to have access to HPC. We have done performance tests that show that traffic patterns for researchers are significantly different to those of normal industrial and civil use of networks. Our networks are fully optimised for research traffic patterns and they deliver exactly as promised. Commercial networks are usually oversold – you buy a 10GB network link and you only get 1gb or less if you start to measure it. That is not what happens in the research network world.

But we do a lot more than networking. We also provide federated identity systems that work worldwide that allow researchers to log into various services all over Europe using their own institute ID. This identity system has now been adopted by PRACE to make sure that researchers can access HPC and auxiliary equipment in computing centres.

GÉANT
Networks • Services • People

Can you summarise GÉANT's role in EuroHPC and how the collaboration works?

Our role is to make sure, together with the national research and education networks that make up our membership, that there is optimal and sufficient connectivity between the computing centres and from the HPC centres to the rest of the research institutes all over Europe.

How are you able to do that?

We already have a network in place that connects all universities and research institutes in 40 countries in Europe, so more countries than are in the EU. We are currently involved in the process of upgrading our network, which will take place over four years. The upgrade will ensure that the network is based on fibre infrastructure so that we have optimum connectivity options by maximising the bandwidth and minimising latency. We can even have access to wavelengths on the fibre that will mean we can separate out channels that are specifically dedicated to connecting supercomputing centres.

What did you consider to be the highlights of the PRACEdays19 event?

I thought that the talk by the EC's Thomas Skordas on the next steps for EuroHPC was particularly interesting. We have been talking about EuroHPC for some time now and we are getting closer to fulfilling the promise, but I certainly had some questions about the funding of it all in Europe by the Member

Erik Huizer of GÉANT presented a clear vision of how the future demands of EuroHPC will be met through collaborative networks

States and I felt that Thomas Skordas certainly addressed and answered these questions.

What else needs to be addressed in the world of HPC in Europe?

There are a couple of things I would like to see addressed in the next year. Firstly, I think we need to get clarity together with members of PRACE to define what the expected needs are for networking between exascale and pre-exascale computers. We are capable of dealing with almost anything, but it would help to have the requirements agreed and down on paper.

Secondly, PRACE and GÉANT have signed a memorandum of understanding on a framework called the Wise Trust and Identity Framework to secure collaboration between infrastructures, and

"We need to get clarity to define what the expected needs are for networking between exascale and pre-exascale computers."

I would like to see that implemented. It will mean that we will have delivered a blueprint architecture, which has been discussed with the HPC community, on how to operate our federated identity system and use that for authorisation and authentication in the HPC world. Having signed the memorandum, I think it's time to start implementation and we should be able to have everything in place within a year.

Finally, security is another issue we need to start working on with PRACE. We have talked a lot about secure access, which is, of course, important, but security is another big challenge we face. With cyber threats being so real, there is a great need for GÉANT, the NRENs and PRACE to work together to make sure that whatever we provide in terms of services to researchers is as secure as possible.

The PRACE User Forum

The PRACE User Forum is an independent entity where PRACE users can discuss their experiences and express their future needs as well as give feedback on the current services and resources provided by PRACE. The aim is to provide an effective mechanism through which the Tier-0 user community can give feedback to PRACE. At PRACEdays19, the PRACE User Forum held a session where they presented a number of points of interest about the feedback they receive, as well as providing an open session for users to voice their concerns and give suggestions in the presence of PRACE representatives.

Maria Grazia Giuffreda, Turlough Downes, Maria Teresa Parra Santos, Troels Haugbølle

The PRACE User Forum was streamed live. Watch a recording here:
www.youtube.com/watch?v=WUUckLkocoo

Troels Haugbølle

Oriol Pineda

Maria Grazia Giuffreda

After an initial talk by Maria Grazia Giuffreda, in which she described the finer details of the peer review process that PRACE uses to select the best applications based on scientific excellence, Troels Haugbølle and Turlough Downes presented the results of a survey that the User Forum has carried out of past and present Tier-0 users to find out their thoughts and opinions about PRACE and the rest of the HPC landscape. Anyone who has used PRACE resources is considered to be a member of the forum, and around 50 responses were received from users from a large variety of disciplines.

What is a typical PRACE user?

The first questions in the survey aimed to decipher what the typical PRACE HPC user is like. 90% of those who responded were academics, and these came from a large scattering of disciplines, with the largest segment being those from universe sciences, which made up 30% of those who responded to the survey. No industrial users responded.

The next question in the survey asked users how many CPU hours they used per year, with the largest group responding that they used somewhere between 10-100 million per year. Users were also asked about the type of code they use. The greatest proportion of users work with codes developed in-house, including the fifteen largest groups who responded, while only a small percentage use community codes. Downes noted that while some people might find this concerning, as in-house codes are more likely to contain errors, the diversity of codes used in Europe is actually a strength because it means that research results are being validated by addressing the same problems with different algorithms.

Concerning computing architectures, the majority of users are still working with multi-core processors, but there is a growing number of groups, one third of whom are able to run on many-cores such as GPUs and KNL. The users who responded worked with a large spread of the different machines available across Europe.

The application process

Following on from questions about the users themselves, the survey then moved on to finding out about what the users thought about the application process for access to PRACE resources. As Haugbølle pointed out, this is one of the main

"We need to prepare more users to run on the incoming pre-exascale systems."

focal points of the User Forum's work, as the process needs to be rigorous enough to select the people who can do the best science, but not so burdensome that it puts people off applying. Fortunately, 62% of users responded that PRACE applications require a similar amount of work to other grants, and there was a similar response about the writing of the technical justification, although the results suggested that users found gathering the information for this section quite difficult. Comments from users recommended that more specific guidelines were given for answering this section.

Users found writing the scientific justification easier than the technical one, and also did not find it a challenge to answer the referee comments. This is probably due to the average PRACE user being well-versed in writing scientific papers, and also because they likely enjoy writing about their domains of expertise!

The survey also questioned users about what they thought of the frequency and size of the calls that PRACE has to offer. At the moment, PRACE has two large calls per year for large one-year allocations, and the majority of those surveyed agreed that they liked it this way, although there was probably some bias here because those being surveyed were the ones who had succeeded. Haugbølle noted that smaller allocations are probably better served by applying for time on Tier-1 national systems. A number of people also said that they would like to be able to apply for a machine architecture (x86 multi-core, GPUs etc) instead of the specific machine, to make the allocation process easier.

Infrastructure and experience

62% of those surveyed replied that they felt that their workflow was very well supported. In terms of the types of workflow that people are doing, the majority of users answered that they were doing large-scale parallel simulations, although of course most PRACE machines are geared towards this. There was, however, a marked upturn in those doing more data-oriented science, as well as those asking for long-term storage, an issue which PRACE is in the process of addressing. Happily, when users were asked whether they would apply to PRACE again, the answer was overwhelmingly positive, with 70% saying that they definitely would.

Finally, participants were asked whether they feel PRACE is still relevant, or if it has been superseded by other HPC facilities. There was a strong response that it has not been outshone, and still remains an important part of most of their research careers. The small amount of people who said that their needs were better covered by national facilities tended to come from larger countries with stronger national HPC infrastructures.

Lessons learned

The presentation finished with a selection of general comments from the users who were surveyed. One suggested that users should make sure to collaborate with PRACE's High-Level Support Teams to help with the creation of new algorithms and implementations that will not only advance their field, but also maintain Europe's world-leading excellence in the production of applications. Another user said: "PRACE's three-year access is crucial for sustained, long-term, ambitious projects from large groups".

Haugbølle finished the session with a summary of the lessons learned from the survey. The general outlook is that PRACE has created a fantastic mechanism for grading and selecting applications, and remains a unique and important infrastructure for European researchers.

Some final thoughts from those on the Programme Committee who run the PRACE User Forum were also raised. There is a need to prepare more users to run on the incoming pre-exascale systems, which will require a diversity of users. When we look at the landscape of community codes, it becomes apparent that most of them have risen organically, starting out as codes developed by small teams of researchers.

To create widely-used codes such as GROMACS that will work with next-generation machines, PRACE will need to have input from as many people as possible. Algorithmic development will also flourish if more users are encouraged to use the established mechanism of working with High-Level Support Teams to develop codes used on Tier-0 systems.

In memoriam - Dimitri Komatitsch

Dimitri Komatitsch, a visionary computational geophysicist and well-beloved member of the PRACE community, passed away unexpectedly on 21 January 2019 at the age of 48. His colleague and friend, **Vadim Monteiller**, spoke about Dimitri's vital contributions towards the use of HPC in seismology and beyond.

After obtaining an engineering degree from Telecom ParisTech together with a Master's degree in Electronics from Université Pierre-et-Marie Curie, Dimitri Komatitsch joined the Institut de Physique du Globe de Paris, where he worked in Albert Tarantola's group on geophysical tomography. As part of his PhD thesis, he developed the spectral-element method and its application to seismology issues, under the supervision of Jean-Pierre Vilotte.

Vadim Monteiller's presentation of Dimitri Komatisch's achievements

After a post-doctoral fellowship at the Department of Earth and Planetary Sciences at Harvard University, Komatitsch and his colleague Jeroen Tromp moved to the Department of Geological and Planetary Sciences at Caltech in Pasadena, USA. There, he became a senior research fellow in scientific computing and geophysics. The SPECFEM codes (2D and 3D versions), used world-wide for seismic and acoustic imaging in the academic and industrial fields, are the results of his fruitful collaboration with Tromp.

In 2004 Dimitri went back to France and became a professor at the University of Pau at the Laboratory of Modelling and Imaging in Geosciences, where he would eventually become director from 2007 to 2010. In 2011, he became a professor at the University of Toulouse at the Laboratory of Geosciences and Environment, and then a Research Director at CNRS at the Laboratory of Mechanics and Acoustics in Marseille.

At CNRS, Dimitri continued his extensive work at the interface of applied mathematics, geophysics and High-Performance Computing on adjoint/inverse imaging problems based on acoustic waves. He collaborated with a large number of high-level scientists, always with great motivation and enthusiasm, on seismology-related topics, on GPU computing for seismic wave propagation, and on the optimisation of high-order finite-element codes on SMP machines. With colleagues at LMA, he extended his enthusiasm and curiosity into researching various non-geophysical applications, such as underwater acoustics, non-destructive testing, and medical (ultrasonic) imaging.

Dimitri Komatitsch 1970-2019

Dimitri was a member of the Institut Universitaire de France from 2007 to 2012. He was also a member of the Scientific Steering Committee of PRACE. He received numerous awards including the Gordon Bell Super Computing Award in 2003, the Atos-Bull Award in 2010, the 3rd Bull Joseph Fourier Award in 2009, and the Simone and Cino del Duca Foundation/IUF/Académie des Sciences Award in 2013 for his work on the numerical simulation of wave field propagation on new High-Performance Computing architectures. He was a member of the American Geophysical Union, the Society of Exploration Geophysicists (SEG), and the French Acoustical Society (SFA). He served the SEG as an Associate Editor for the journal Geophysics.

There is unanimous agreement among the PRACE community that Dimitri's work has paved the way for computational seismology towards High-Performance Computing, which has opened up a whole new branch of careers for young geophysicists.

Panel discussion: The future of scientific and industrial computing in Europe

The two pillars of scientific experiments and theories have been the seeds of our understanding of the world for centuries. The appearance of computers and their implementation in efficient numerical techniques has provoked the expansion of simulations as a new way to understand, guide and design new experiments, models, materials, and processes that can better support the technical and economic development of our societies.

Simulations are now used in all areas of science and humanities and have been the main use of supercomputers world-wide. However, the large amounts of data generated by simulations together with those derived from instruments now call for a different era. The use of data science and statistical learning has impacted our daily lives almost without us noticing, and will impact the way we use supercomputers.

Exascale computers will need to answer societal needs for efficiency in the use of resources, sustainable growth, and security, and have to address crucial problems like energy supply, health, and active ageing. The combination of simulations and data will certainly change the way scientific and industrial computing is performed in Europe and the rest of the world. This will challenge our capacities to educate the next generation, develop newer and efficient algorithms and promote smooth technology transfer.

As in previous years, PRACEdays19 brought together a panel of experts to discuss these issues in front of its attendees, providing a platform to begin discussions about what the future of High-Performance Computing might bring.

The panel

Moderator
Guy Lonsdale
SCAPOS

Thomas Lippert
Chair of the PRACE Council

Vít Vondrák
IT4Innovations, VŠB

Lee Margetts
Chair of the PRACE Industrial Advisory Committee

Zoe Cournia
Biomedical Research Foundation, Academy of Athens, Greece

Núria López
Chair of the PRACE Scientific Steering Committee

The age of data

Moderator Guy Lonsdale began the discussion by asking the panel about how the shift from modelling and simulations to data analysis and artificial intelligence has affected their research communities and facilities, and how this will affect the next generation of researchers in their field.

As Chair of the PRACE Council, Thomas Lippert stated that he was confident that PRACE is well positioned to meet the challenges of this new era, with various computer architectures to meet the changing needs of its users. There are complex workflow issues with this new kind of computing, with increasing amounts or research using data-intensive calculations and analytics and AI alongside simulations. Experts are convinced this will require major software development, and with workflows with latency issues due to concurrency of codes, architectures will need to become more heterogeneous.

Vít Vondrák of IT4Innovations replied by saying that researchers in the Czech Republic were increasingly requesting to carry out data science-driven work. Numerical simulations still dominate, but they often now have machine learning techniques embedded within them. The HPC community needs to respond to this demand by designing systems that can handle it.

"From the point of view of software vendors, all of the software on the market needs to be rewritten as a consequence of increased use of data driven analytics."

Chair of the PRACE Industrial Advisory Committee Lee Margetts reprised his role from last year as the voice of industry in the PRACEdays panel discussion. He told the audience how he initially joined the committee as a representative of the NAFEMS trade association, which represents those engaged in the use of engineering simulations. From the point of view of software vendors, he said, all of the software on the market needs to be rewritten as a consequence of increased use of data-driven analytics. He pointed out that companies like EDF are moving away from traditional commercial packages and investing more in open source, which brings the advantage of continually updated software to meet innovations in algorithms and the hybridisation of hardware.

Zoe Cournia, winner of the PRACE Ada Lovelace Award in 2016 and representing those working in the life sciences, spoke about the huge accumulation of data in her field, brought about by the digitisation of health records, genomics, next-generation sequencing and the elucidation of protein structures through computational biology. This mass of data is dictating new architectures, networks, and workflows for the archiving, analysis and collection of data. A new kind

of expert is now emerging, she said, with many Master's programmes catering for people who work at the intersection of computer science and medicine/natural sciences. The community needs to leverage the power of these new scientists, who can help in gaining tremendous computational speedup. Chair of the PRACE Scientific Steering Committee Núria López told a similar story about her field, materials science, where there is also a huge amount of data being produced. A lot of work needs to be done to upgrade codes to meet these new demands, and this will require collaboration between all of the groups in PRACE and EuroHPC. But she was also keen to stress that this new way of computing is exciting, and will bring about possibilities that we cannot imagine at present. This will be the role that the hybrid experts mentioned by Zoe Cournia will play in the future, so scientific and industrial communities need to adapt and integrate them.

The dawn of exascale

The next question was about the upcoming era of exascale computing. Guy Lonsdale asked the panel what their communities hoped they would be able to do with exascale computers and what would be the strategic challenges of achieving these goals. Núria López spoke about the recently published Scientific Case for Computing in Europe 2018 – 2026 (www.prace-ri.eu/third-scientific-case/), which outlines the aspirations of each of the main HPC domains for exascale computing. There will be increased integration of experiments with HPC, she said, citing the recently produced image of a black hole. For many domains, their codes will be able to achieve higher precision and accuracy.

The field of computational molecular biophysics examines the structure and function of biomolecules and tries to create new drugs that can act upon them, said Zoe Cournia. Right now, it is possible to simulate a 200 million atom virus, which is already huge, but the dawn of exascale computing will allow researchers to simulate these entire systems using quantum chemistry, bringing higher accuracy and spatiotemporal resolution and allowing for a much more predictive approach to drug creation, which will reduce the time and cost of bringing drugs to the market. She also mentioned that the EU is at present dependent on non-EU technologies.

This will change with actions such as the European Processor Initiative, which will bring about growth of technology from within Europe. Investment in software for the exascale era will be crucial for maintaining Europe's leadership in HPC applications.

"Investment in software for the exascale era will be crucial for maintaining Europe's leadership in HPC applications."

Lee Margetts highlighted the plenary talk from Julien Harou, who demonstrated that he could use exascale computing right now with his code to address socioeconomic problems with water resource management. Engineering companies are increasingly looking at the use of digital twins, which will be used to design the next generation of many products. The challenge for exascale will be how to deal with multi-physics, multiscale workflows. At present, we have the mind-set of scaling a code up to 100 000 cores, whereas what industry needs is heterogeneous applications.

Vít Vondrák answered the question by saying that, when it comes down to it, most communities will be looking for exascale to bring greater precision to their work. Some codes will have to be rewritten, though, and more training will be needed for this.

Thomas Lippert then offered an exciting vision of the future with exascale computing, telling the audience

how it will not only bring greater precision to existing fields of research, but also open up new fields that have previously been unable to use HPC. Nuclear chemistry, for example, will be able to explore fusion and lean nuclear energy, which can only be done through simulations on exascale computers. Weather prediction models will also be able to augment traditional modelling with data driven modelling using things such as satellite imaging and leading to predictions valid on the time scale of months.

HPC in industry

The panel were then asked whether they saw a divergence in requirements in academia and industry towards extreme scale computing, and what can be done to transfer knowledge more effectively from academia to industry. Núria López stated that it was perhaps the case that industries themselves are less interested in exascale, with some of the larger companies buying hardware themselves and doing things in-house. But academia and industry need to work together to make sure that technology and the knowledge gathered from it is transferred between the two. Vít Vondrák expressed a similar sentiment, but also mentioned that we must focus on SMEs, for which we need to make it easier to access hardware, knowledge and training. Developing HPC as a service is a must, so that they do not have to spend huge amounts of time learning how to use it.

"The EU needs to find out what industry needs, what citizens need, and to work towards providing it."

Academia is well-educated in the use of HPC, but Zoe Cournia said that industry is lagging behind. There is the well-known European paradox of the perceived failure in Europe to translate scientific advances into marketable innovations, and this is because of a lack of synergies between academia and industry. Europe needs to look across the Atlantic for inspiration. In the USA, calls are dictated by the Ministry of Defence, the Ministry of Energy, the Ministry of Commerce and others, who are in constant correspondence with the needs of their industries which are then passed on to the research community.

The EU needs to find out what industry needs, what citizens need, and to work towards providing it. Cournia herself is working with EuroHPC to document the

needs of industry, and believes that stronger industrial investment in HPC can be a huge driver for the European economy. Thomas Lippert spoke about the difficulty that start-ups have of even considering what they can do with exascale machines, due to the fact that it is so unlikely that they will get to use them.

Europe needs to provide the best entrepreneurial minds and ingenious innovators with access to the most powerful computers, but it needs the political will to make this happen. It will be PRACE and EuroHPC's responsibility to advocate for this.

The EC recognises the needs for industry to increase its uptake of technology, and has acted upon this with the formation of 36 separate European Technology Platforms, each supporting specific industries, including pig breeding and power station suppliers.

But, Lee Margetts argued, there seems to be an issue with the language being used, as many of these industries do not see HPC as an avenue for improving the use of technology. To improve industrial uptake of HPC, we need someone to be talking to these platforms and explaining that HPC is actually technology.

Continuing on from the previous question, Guy Lonsdale brought up EuroHPC's pledge to reserve 20% of the use of its computers to pay-for-access by industry. Lee Margetts expressed that he did not understand the motivation behind this pledge.

The Industrial Advisory Committee's liaisons with industrial leaders have always shown that what industry values is working relationships with scientists, not time on computers. The value proposition for industry is not computing time – it is the relationship with the academic HPC community.

Zoe Cournia stated that she believed the motivation was to bring some sustainability to the infrastructure by generating some revenue without distorting the market), but largely agreed with Margetts, as did Núria López, who added that these relationships are especially important for SMEs, where the input from an outside expert can hugely expand the overall knowledge of the company.

Thomas Lippert offered a different opinion, arguing that the 20% figure could be changed in the future, and that it was a good thing that more computing time was being made available to industry.

Data analytics: a threat to the HPC community?

The session finished with a discussion of the possibility of machine learning and data analytics proving so disruptive that they replace modelling and simulation and change the face HPC as we know it. Thomas Lippert said that these fields could be viewed as a threat by those who refuse to embrace them. Wonderful things can be done through a combination of exascale computing and machine learning, and we need to think about how they can grow together.

Vít Vondrák said that he thought that there was sometimes too much emphasis on machine learning and artificial intelligence. If numerical simulations can solve something easily, then machine learning is not really needed. But it is useful to use machine learning when standard simulations are not capable of providing answers, so it is a case of using the right technique for the right problem.

Lee Margetts said that the HPC community often sees its future as a march towards a big machine in a box. But, when we get it, will people actually be doing things in a way that does not require it? The production of data on experimental devices such as synchrotron light sources is sometimes so big that it cannot be stored, and has to be processed on the fly using edge computing. So, could the envisioned exascale computer of the future, rather than being a cathedral of HPC, actually become a mausoleum? Zoe Cournia spoke of her role as editor of a computational chemistry journal. She said that it is often the case that papers that use artificial intelligence to obtain results are rejected by reviewers because they do not understand why the results are correct. It can be a battle to get these papers through, she said, but using HPDA and machine learning is the future. There are massive amounts of data that we can now exploit with artificial intelligence to get meaningful results and overcome big data bottlenecks, but we have to accept the fact that often we may not understand the physical origin of the results.

It was up to Núria López to finish with a positive note, saying that if we invest in new ways of analysing the results of artificial intelligence work, we can retrospectively retrieve the understanding of why they are what they are. All of them agreed in the power and the need for training.

BDEC: Where big data and HPC join forces

The Big Data and Extreme-scale Computing (BDEC) plenary provided a European and international view on the convergence of big data and HPC. The keynote speeches emphasised the challenges and expectations, and potential solutions for the shift towards this new era.

Rosa M Badia of Barcelona Supercomputing Centre

In the past five years, the United States, the European Union, Japan and China have each moved aggressively to develop their own plans for achieving exascale computing in the next decade. But while these exascale initiatives have understandably focused on the big challenges of exascale for hardware and software architecture, the relatively recent emergence of the phenomenon of big data in a wide variety of scientific fields represents a shift that is transforming the entire research landscape on which all plans for exascale computing must play out.

Building on the work of the IESP, the Big Data and Extreme-scale Computing (BDEC) community has staged a series of workshops that have aimed to map out and account for the ways in which the major issues associated with big data intersect with, impinge upon, and potentially change, the national (and international) plans that are now being laid for achieving exascale computing. The latest workshop was held in Poznań during the EuroHPC Summit Week 2019, the first part of which took place as a plenary session at PRACEdays19.

The first keynote of the plenary was given by David Keyes, director of the Extreme Computing Research Centre (ECRC) at the King Abdullah University of Science and Technology. His talk, titled "The convergence of big data and large-scale simulation: Leveraging the continuum", argued that, just as the convergence of theory and experiment in the pre-computational era launched modern science, the convergence of simulation and big data in the exascale computational era has the potential to give humanity predictive tools to overcome our greatest natural and technological challenges. Using edge-computing to provide on-the-fly analysis of the vast swathes of data that we now produce will be essential for this new era. Machine learning will also play a big role, both in the development of new applications and for enabling the previously mentioned edge-computing. BDEC's recent report "Pathways to convergence"[1], highlighted a number of challenges and potential actions that can bring about a new era of HPC and big data working in tandem. Science and engineering may be minority users of big data at present, but they can become leaders in the community by harnessing high performance computing and being pathfinders for other applications. In the end, said David Keyes, the marriage between the fields will undoubtedly happen naturally over time, as the benefits for those working in big data and HPC will be too big to ignore.

> ***"The convergence of simulation and big data in the exascale computational era has the potential to give humanity predictive tools to overcome our greatest natural and technological challenges."***

Rosa M Badia of the Barcelona Supercomputing Centre provided the second keynote, providing a real-world example of how the worlds of big data and HPC can join forces. In her presentation, titled "Workflows for continuum computing platforms (from edge to HP computing)", Badia talked about the ELASTIC project, which is developing a software architecture that will enable smart systems to satisfy the performance requirements of extreme-scale analytics workloads.

ELASTIC's software architecture will be tested in the city of Florence, Italy, equipping tram vehicles with V2X communication and a variety of detectors aimed at improving traffic management, allowing predictive maintenance and enhancing the tramway operation. This project uses a complex infrastructure with sensors and instruments in the edge that generate data and computation that will be done across the compute continuum, from edge to cloud. Executing applications on such distributed environments is complex, especially if real-time constraints need to be considered as well.

In conclusion, the BDEC workshop in Poznań emphasized the challenges and expectations, and potential solutions for the "Digital continuum" on which the next generation cyber-infrastructure platforms will rely, dealing with heterogeneity of data and computing resources: from the edge all the way to HPC centres. The main objective is to enable transnational research communities in a wide range of disciplines to converge on a common "digital continuum platform" (DCP), a next generation network computing platform for creating distributed services in a world permeated by devices and saturated by digital data.

Reference

1. Asch, et al. (2018) Big data and extreme-scale computing: Pathways to Convergence-Toward a shaping strategy for a future software and data ecosystem for scientific inquiry. The International Journal of High Performance Computing Applications, 32(4), 435–479.

Executive Summary

www.prace-ri.eu