EuroHPC Summit Week 2022 & PRACEdays22

Executive Summary

EuroHPC SUMMIT WEEK 2022 & PRACEdays22

EuroHPC Summit Week 2022 & PRACEdays22 Executive Summary

www.prace-ri.eu

Editor-in-chief

Marjolein Oorsprong – PRACE Head of Communications

Proofreading

Patricia Fredes – PRACE

Article editor

Harry O'Neill, Insight Media

William Davis, Insight Media

Graphic design

Alison Davis, Insight Media

Published by Insight Media, 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 made available 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.

ISBN: 9789464597707

Contents

Forewords

Anders Dam Jensen
Executive Director, EuroHPC JU

This year's EuroHPC Summit Week was a landmark event for many reasons, not least because it brought the European HPC community together under one roof for the first time in three years, which for my part was a very positive experience.

During the discussions in Paris, it was very inspiring to meet experts, academics, researchers, vendors, policy-makers and users who have been building the European HPC ecosystem over the years. The presentations and debates were very informative and lively. My team and I were energised after our week in Paris. We will now take the ideas and suggestions made during the Summit Week and we will make sure to carry that work forward.

PRACE has done an excellent job over the years to build this community and I am delighted to see how much has been done and how dynamic it is. The future is bright for the HPC Community in Europe and EuroHPC JU will play its part in building it further.

I have three main take-aways about this conference:

Everybody is trying to understand what activities EuroHPC JU will undertake in the years ahead. My answer is simple – our activities are based on the guidelines set out in our founding EU Regulation. Our work plans are developed with our members which are the participating states and partners which are HPC Hosting entities, researchers and industry.

Many in Paris mentioned the need to ensure that we build training programmes to develop skills in HPC that will attract a more diverse population who want to engage with supercomputing. The JU is thinking about this in partnership with PRACE and the European HPC competence centres and I hope to have proposals approved by our governing board before the end of 2022.
Another key takeaway for me was the prominent discussion we had on quantum computing. I felt that the community is really coming together to embrace the potential capabilities quantum computing has to offer. This is very exciting and will be a fascinating part of EuroHPC JU's programme.

"The future is bright for the HPC Community in Europe and EuroHPC JU will play its part in building it further."

We have spent the last few years building up the EuroHPC Joint Undertaking – lots of back-office stuff to get everything in place while restricted in our ability to interact with the community due to the pandemic. EuroHPC Summit Week 2022 in Paris gave us a renewed sense of energy for the future of HPC in Europe after two years in our virtual world.

A big thank you to PRACE colleagues for this excellent event!

Serge Bogaerts
Managing Director, PRACE aisbl

The EHPCSW / PRACEdays event has always been about providing a place where we can feel a sense of community in the world of HPC. Of course, COVID-19 has robbed us all of many face-to-face experiences over the last few years, and although we thoroughly enjoyed the challenge and novelty of last year's purely digital event, it has to be said that seeing you all together in Paris truly brought a smile to my face.

One thing I was struck by at the event was how the HPC ecosystem has developed over the last decade. Researchers from communities we would never have considered to be interested in computing are now here in droves, and the number of people attending from industry continues to grow.

Bringing all these people together for a week of talks each year is so valuable to all of us in terms of the potential for cross pollination of ideas and strategies, and I hope all of you managed to make some connections with people you may not have expected to.

Quantum computing took centre stage this year. The community working on quantum information has quietly been working away in the background for a long time, but it finally feels like the technology has come of age and there is so much to be excited about in this area.

Keynotes from Elham Kashefi, Kristel Michielsen and George Em Karniadakis all touched on how quantum computing might evolve in the coming years, and it is something that we should all be aware of and planning for.

The coming years will be critical for high-performance computing in Europe. EuroHPC will be the driving force behind how we approach this as a community in the long term, and I would like to congratulate them for all the hard work they have done in preparation for this monumental task.

"Now is the time for us all to join hands and help build Europe's HPC ecosystem towards its full potential."

It was great to have so many of their core team present at EHPCSW this year, and now is the time for us all to join hands and help build Europe's HPC ecosystem towards its full potential.

We at PRACE could not be prouder of what we have achieved since our inception. The training programmes and other forms of support for researchers that we provide, as well as our fantastic peer review process, are hallmarks of the quality of our organisation and our unerring desire to bring the best out of everyone who wishes to participate in the world of high-performance computing.

Long may it continue - and a big thank you to everyone involved in European HPC.

Modelling the climate

High-performance computing plays a huge role in many of the most pressing issues our society faces, and none more so than in our ability to understand and predict climate change. **Sylvie Joussaume**, senior scientist at CNRS, used her keynote at EHPCSW to lay out the evidence about the nature of climate change, as well as to provide an insight into international efforts to improve the climate models that we rely on.

Expert climate modeller Sylvie Joussaume opened her talk at EHPCSW 2022 with a stark message: global warming is unequivocal, and can only be explained by human influence. Although the peak of climate change denial has now passed, Joussaume's presentation of the pertinent facts still made for compelling listening.

Global surface air temperature, one of the key indicators for climate change, has risen 1.09°C over the last decade compared to the reference period between 1850-1900. Each of the last four decades has been warmer than any decade that preceded it since 1850. The trend of warming is clear.

Occurring parallel to this warming has been a change in atmospheric composition. Atmospheric CO2 has increased by 48% since 1750, a rise unprecedented in 800,000 years. More carbon on top of this has been sequestered into the ocean and land and, although this may have other adverse consequences, Joussaume remarked we would be in a far more dire situation if the Earth lacked this ability to clean the atmosphere.

To determine whether these rises in temperature and atmospheric carbon are related, climate models are used. These Earth System Models (ESMs) are the culmination of our amassed knowledge about how things work on our planet and include, amongst other factors, representations of the atmosphere (winds, humidity), oceans (energy, temperature, and salinity), the water cycle, soil, topography, vegetation, and includes natural (solar and volcanic) and anthropogenic (greenhouse gas

Sylvie Joussaume delivers her talk

emissions) forcing factors. ESMs are undoubtedly unwieldy things, stuffed full of legacy code, but to say that due diligence is done when investigating our climate is an understatement.

With these models, it is possible to simulate how the temperature of the Earth has varied since 1850. In a straight comparison between simulations using just natural forcings versus those that also include human factors, the undeniable conclusion is simple: it is impossible to explain the trends in temperature up to today without including the influence of humans.

As well as looking back, these models allow us

> *“The coupling of ice sheets and the coupling of vegetation still need improving, as well as the way we represent clouds, which is currently the main uncertainty in the models.”*

to peer into the many possible futures that await us. In a best-case scenario in which we reach zero emissions by 2050, we will end the century with a global rise in temperature of 1.5°C, in line with the ideal outcome laid out in the Paris Agreement of 2015. If we carry on as we are, however, we face the prospect of a global rise in temperature in the range of 4°C, along with which will come devastating knock-on effects such as heatwaves, droughts, floods, and extreme precipitation events.

Undoubtedly, concluded Joussaume, our best models of the Earth tell us that temperature change is highly correlated to cumulative CO_2 emissions, and now is the time to act if we do not wish to face unprecedented global disruption.

Climate models and the international landscape

Earth System Models use physical laws such as the Navier-Stokes equations for the dynamics of the atmosphere and the oceans, but also include parameterisations for processes that we do not have physical laws for, such as cloud formation and dissipation, radiation, and sub-grid scale processes. By bringing all these aspects together and making them interact, a “coupled” model is created.

Ongoing efforts by the international community to improve coupled climate models aim to build upon the pioneering work of Syukuro Manabe, the father of modern climate science and the joint winner of the 2021 Nobel Prize in Physics for his work in creating the first global climate models. The World Climate Research Programme (WCRP) guides and oversees the Coupled Model Intercomparison Project, now in its sixth phase, which runs internationally coordinated experiments to assess the performance of the various current models and how they differ in their projections of the future.

The scale of this work requires a huge common data infrastructure to provide access to all the available data. This infrastructure is supported by multiple agencies across the world, including IS-ENES (Infrastructure for the European Network for Earth System Modelling) and the Department of Energy in the US, and together they handle over 30 petabytes of data.

Challenges in climate modelling

Climate science has come along leaps and bounds in the last few decades, but many challenges remain to improve our fundamental understanding of climate systems. The WCRP recently issued a new strategy outlining four main objectives to achieve this. The aims are to improve our understanding of variability and changes in coupled systems, improve our ability to predict near-term evolution of climate systems with seasonal to decadal predictions, improve our understanding of the long-term responses of the climate, including any possible future tipping points, and to bridge the gap between climate science and society through communication efforts.

The global climate is a hugely multiscale system, with events happening from the scale of metres up to thousands of kilometres, as well as from seconds up to millenia. Current climate models generally have spatial resolutions of around 100-200km, and although ongoing work aims to reduce this down to the kilometre scale, much would still be missing. Better spatial resolution would improve our understanding of regional patterns and our ability to predict phenomena such as tropical cyclones, but with a doubling of resolution comes an eight-fold increase in computing time. When running models on high-performance computers, compromises must therefore always be made between the resolution, complexity, duration, and ensemble size of simulations depending on the goal.

Speaking after her keynote, Joussaume explained how heterogenous architectures might bring some difficulties to the climate modelling community. "Climate models require a lot of internal exchange of data between points on the grid, which also require exchanges between nodes" she said. "Most models still run on CPUs, but a number of groups have now started work towards porting their models for use on GPUs, while others have decided to completely redesign their models using a separation of concerns approach to allow more flexibility". Dealing with increasing amounts of model data also calls for new approaches. "This will require to combine on-the-fly processing, near-data processing, compression, and data reduction".

Joussaume also highlighted the fact that there are still several processes that are not modelled as well as they could be. "The coupling of ice sheets and the coupling of vegetation still need improving, as well as the way we represent clouds, which is currently the main uncertainty in the models. We need to keep improving so that we can understand extreme events such as the heat domes recently seen in Canada, a type of phenomenon that we still know little about."

"As climate modellers we must not forget our role outside of climate change, which is to improve our understanding."

The use of AI in climate modelling is likely to increase in the coming years. Although it has already been widely used for analysis of results, it is now being investigated for developing or replacing parameterisations as well.

"There are ongoing developments to use AI to represent subgrid-scale processes and factors like cloud formation," says Joussaume. "However, there are still questions around how much it can be used if only based on the observations we have so far, especially under climate conditions differing from today. As well as this, it is possible that in the future we may be able to harness AI to accelerate some parts of our code."

A bright and busy future

The coming years will be defining in addressing the mitigation of climate change, said Joussaume, but we will also need to look at the question of adaption. "As climate change occurs, we need to work out how to adapt to it. We need to use the information from our models to do this, and part of our work in the future will be devising ways of presenting this information in a way that makes it more usable for this adaptation."

"Climate models are the core of our information for climate change mitigation and adaptation", concluded Joussaume, "but as climate modellers we must also not forget our role outside of climate change, which is to improve our understanding of the climate in general."

There is little doubt that climate change remains at the forefront of climate science, but to address this, a wide range of model configurations will be needed. High resolution simulations with shorter time scales that give highly detailed regional information can be complemented by lower resolution runs with larger ensembles that address issues around uncertainties and understanding.

High-performance computing is tightly intertwined with the way these models are built and run, and it will be important for the climate science community to continue working to ensure they can use the newest machines to their full capacity.

Quantum computing
in industry

This year's EHPCSW felt like the year that quantum computing took to the mainstream, and this was no better demonstrated than during the session on industrial quantum computing, where various companies took to the stage to highlight their initiatives and integration of quantum computing strategies.

Airbus – Gerd Büttner

The first presentation came from Gerd Büttner of Airbus. Airbus has long been exploring quantum technologies to enable their ambitious aim to be the first company to offer a hydrogen-powered zero-emission aircraft by 2035. This will require them changing the way they do almost everything, for instance considering how to store hydrogen on a plane. Quantum computing is also being used to explore challenges around military aircraft, space systems, connected intelligence, and unmanned aerial systems.

Airbus is looking at three areas of quantum technologies: computing, sensing, and communications. Büttner's talk focused on computing. The Airbus Quantum Computing Challenge set out five problem statements: aircraft climb optimisation, computational fluid dynamics, quantum neural networks for solving partial differential equations, wingbox design optimisation, and aircraft loading optimisation. Teams competed to find the best solutions to these issues, and the successful teams continue to work with Airbus in various forms.

EDF – Marc Porcheron

EDF aims to improve internal skills in quantum computing, identify the relevant use-cases for EDF businesses, experiment with the available technologies through proof of concept, and enlighten its operational divisions on the disruptive applications of quantum technologies.

Current use-cases under study by EDF include: the production, transport, and distribution of electricity; the safety of installations; energy management; electric mobility; renewable energy development; and the protection of data and expertise. In these use cases, the quantum technologies they are exploring are: combinatorial optimisation; quantum machine learning and neural networks; quantum walks; cloud quantum computing; quantum sensors; and quantum simulations for studying materials. Overall, they are seeking potential advantages in terms of performance and energy consumption.

Porcheron delved deeper into the example of quantum optimisation for electric mobility. Vehicle-to-grid technology considers electric vehicles not only as a power demand for the grid but also as a power resource. Using electric vehicles in this way, however, involves many difficult combinatorial optimisation problems, for instance minimising the completion time of charging tasks with priorities, and optimising a huge number of electric vehicle charges and discharges to contribute to the stability of the frequency of the grid. They have already had encouraging results using a quantum approximate optimisation algorithm on real instances of these problems, but of limited size.

"Quantum machine learning would impact areas from target assessment to the clinic, for instance allowing lung cancer to be identified very early."

BMW – Elvira Shishenina

Next up was host of the session, Elvira Shishenina, to present the situation of BMW. They began their journey into exploring quantum technologies back in 2017, and now have more than 50 use-cases being investigated, such as: quantum chemistry to develop new materials for energy storage; optimisation for various industrial processes; machine learning for automated driving, automated quality assessment, and generative modelling of vehicle components; and numerical simulations for structural analysis, crash simulations and more.

Shishenina outlined BMW's strategy for prioritising different areas of research into quantum technologies based on the hardware requirements (qubits required) and software requirements. Quantum chemistry applications appears as low-hanging fruit based on these factors, but it is important to be aware of the technical and business challenges that remain.

CA-CIB – Didier M'Tamon

Quantum computing has a place in the future of finance. This was the basis of the talk from Didier M'Tamon of CA-CIB. When the company began their exploration of quantum technology, they were unsure about the maturity of the technology, but were attracted by the fact that others such as Amazon were investing in it. They began by building knowledge about quantum computing, and M'Tamon himself took a course at MIT on the subject.

Banks do a lot of computations, including those to do with financial product pricing and portfolio optimisation. Accuracy and speed are critical in the sector. CA-CIB wanted to be ahead of the curve in order to be able to compete. Speaking to a number of companies, they understood that at present, noisy intermediate quantum computers were what were available now, which might be able to solve some classes of problems relevant to CA-CIB.

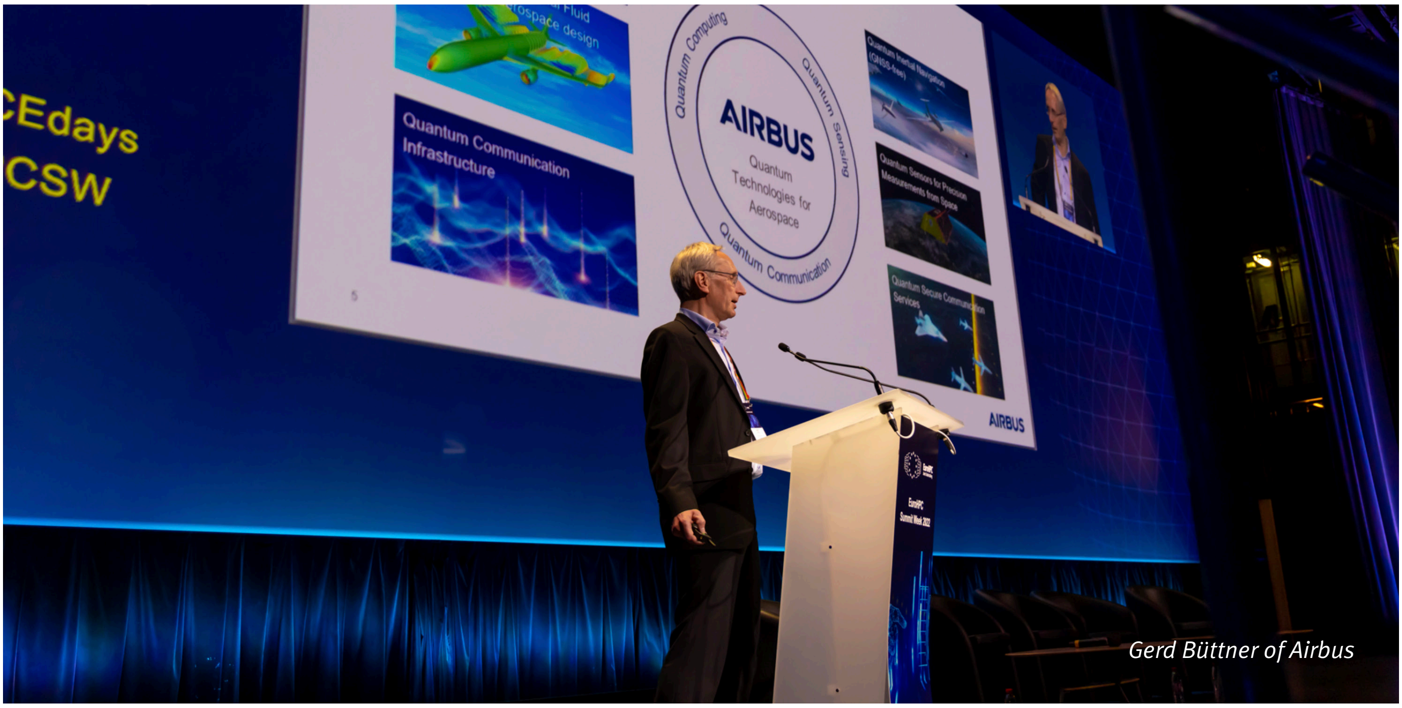

Gerd Büttner of Airbus

They have focused on quantum-inspired problems with tensor networks in order to accelerate the pricing of financial instruments, and quantum machine learning for credit scoring i.e. the ability to predict whether a client will default. Because so few of their clients actually default, it is very difficult to have a good model to predict defaults using classical computing. CA-CIB have looked to quantum computing to try and achieve better performance in this area.

Roche – Agnes Meyder

Roche is the world's leading pharmaceutical company, and has had a quantum computing team for four years. They began by asking how quantum computing can support drug development, communicating with IBM, Microsoft, Google and Intel to gauge the technology available. They also funded Masters and PhD programmes to further scope out the landscape, and collaborated with other industrial users.

Their use cases cover a broad range of research. For chemistry simulations, accurate binding affinity estimations would strongly influence efficiency in lead identification and lead optimisation. In terms of optimisation, rapid identification of minima would strongly influence the whole pipeline of pRED and many other Roche units. Finally, quantum machine learning would impact areas from target assessment to the clinic, for instance allowing lung cancer to be identified very early. These use cases have been explored using real quantum hardware.

Roche has now established baselines that are relevant for the pharmaceutical industry with specific use cases that are ready to be used, so that they can better compare improvements in quantum computing and decide whether it is game-changing or not. They also found that translating problems onto quantum hardware through quantum algorithms is a difficult task, and that the algorithms that will scale well are yet to emerge.

Thales – Frédéric Barbaresco and David Sadek

Aerospace company Thales are currently setting up their roadmap for developing quantum algorithms and identifying their major use-cases. They are pursuing quantum technology because it has the potential to tackle complex problems, offer higher resolution, and reduce energy consumption. Augmented engineering and augmented operations are what they hope to achieve.

After Frédéric Barbaresco and David Sadek outlined a number of the use cases and the different types of algorithms and quantum computers being used, Sadek highlighted one of the key issues, which is the need for people who possess true expertise in quantum algorithms.

It is vitally important that education and training of people in this domain occurs, and to be cautious of the fact that those truly talented in this field will be rare and that demand for them globally will be huge, so brain-drain in Europe could be an issue.

Quantum computing has the potential to solve a myriad of computational challenges that will enable science and industry to gain a foothold over complex systems that have previously seemed unconquerable. **Professor Kristel Michielsen**, quantum computing pioneer and head of the Quantum Information Processing group at Jülich Supercomputing Centre, described in her keynote talk at EHPCSW 2022 how three classes of simulations currently carried out on HPC systems may well be revolutionised by the advent of quantum computing.

The potential of quantum computing

Quantum simulations can be used in scientific fields ranging from elementary particle physics to quantum biology, as well as industrial challenges such as battery design and photovoltaics. Optimisation problems are another area where quantum computing can potentially excel, helping to plan flights, improve the health sector and provide a crucial advantage to those working in finance. Finally, quantum machine learning algorithms have the potential to help turbo-charge data classification, helping to identify specific terrain types such as water and vegetation from satellite remote sensing data, as well as improving image-based medical diagnosis.

Michielsen is coordinator of the High-Performance Computer and Quantum Simulator hybrid (HPCQS) project which, backed by EuroHPC, aims to accelerate the use of quantum computing in Europe through the deployment of quantum simulators and a hybrid HPC-quantum infrastructure that will boost the computing speed of existing classical supercomputers. In reaction to the influx of newly developed quantum computing devices in recent years, in 2017 Kristel Michielsen and Thomas Lippert created a new metric, the Quantum Technology Readiness Level (QTRL) that aims to provide an idea of where each technology lies on a scale of one to nine, ranging from the initial design phase all the way up to real world use.

Two main types of quantum computers exist. The first, digital quantum computing devices, allow one to control qubits individually, but the lack of error correction in these devices means they are currently rated at a QTRL 4-5. The second type are known as analogue quantum computing devices, in which the qubits are controlled together as an ensemble.

Examples of such devices are quantum simulators and quantum annealers, the latter of which have been given a QTRL 8 rating due to processor sizes doubling every two years and increasing connectivity of qubits meaning that they are almost ready for real world use in solving optimisation problems.

It is obvious that huge opportunities lie in the quantum computing space, but the real challenge will be coming up with applications. "Everyone here at EHPCSW needs to be thinking about how to develop prototype applications and use cases for these analogue and digital quantum computing devices," said Michielsen. "Working with quantum computers requires you to learn a new way of thinking. We should therefore all invest some time in analysing our software codes to see whether there are some parts – even very small parts that do not really contribute to the production outcome – that can be run on quantum computers."

> ***"We should all invest some time in analysing our software codes to see whether there are some parts that can be run on quantum computers."***

One of the long-term visions of HPCQS is to forge a link between high-performance computing and quantum computing. HPC systems are already used to emulate quantum computers to gain a better understanding of how they operate and ultimately improve their design. HPC systems can therefore play an important role in the benchmarking of quantum computers, with the largest supercomputers able to simulate quantum computers made up of around 50 qubits. But for the next step – the creation of hybrid HPC-quantum systems – Michielsen believes that every aspect of the existing HPC ecosystem has a role to play.

Professor Kristel Michielsen

HPC users must develop new algorithms and software for the integrated systems. A new quantum computing and simulator full software stack will need to be created, and the development of integrated systems will need to follow a codesign approach alongside the use cases.

During her talk at EHPCSW, Michielsen outlined an example of a hybrid algorithm known as a quantum approximate optimisation algorithm. She demonstrated how it can be used to solve the tail assignment problem – an optimisation problem that aims to minimise costs when assigning aircraft to various flights, while also considering other operational constraints such as airport curfews and maintenance. "We took this algorithm and executed the quantum parts on the JUWELS booster at the Jülich Supercomputing Centre," explained Michielsen. "What we are working on for the future is to be able to run the quantum parts on real quantum circuits in a hybrid modular supercomputing architecture."

Michielsen is heavily involved in developing JUNIQ (Jülich UNified Infrastructure for Quantum computing), a quantum computing user facility that is accessible for scientists from academia and industry in Europe. It installs, operates, and provides access to various types of quantum computers at different levels of technological maturity. It also develops algorithms and prototype applications in collaboration with users, all with the goal of creating an ecosystem for modular quantum-HPC hybrid computing to flourish.

Asked whether she believes Europe is well-placed to be a global leader in quantum computing, Michielsen gave a positive outlook. "In Europe there is a lot of expertise in the fundamental knowledge of quantum physics, more than the US, for instance," she said. "We now need to back up this expertise through projects like HPCQS to ensure that we start integrating quantum computing devices into our HPC infrastructure and begin building the expertise to use them."

"Quantum computing has the potential to provide benefits in many ways, not just in terms of helping to solve complex problems from science and industry. For instance, the amount of energy they consume is tiny in comparison to giant HPC systems. Of course, efforts are already being made to reduce energy consumption in the world of HPC, but we should not ignore quantum computing in these conversations."

In her keynote address **Professor Elham Kashefi** of Sorbonne University addressed the need to the question of the verification and testing of quantum computing as its ties with HPC become ever closer

The challenge of verifying quantum computing

The past five years have seen the communities of quantum computing and high-performance computing draw closer together as the promises of quantum computing come closer to being realised. Traditional methods for solving optimisation problems and those involving big data, for example, will likely be swept aside once quantum approaches begin to flourish, but it will take strong collaboration between the two communities for all of this to be a success. This collaboration has already begun, for example with the HPCQS project, which is bringing together top people from quantum and HPC to work together on quantum-HPC hybrid computing.

Several physical features differentiate quantum and classical computing. For example, measuring a bit in classical computing will reveal the value of the bit: 0 or 1, while measuring a qubit will change the state of the computation. Similarly, information from a bit can be copied to other bits, but a qubit cannot be copied. These are fundamental features that separate the two types of computing at the physical level, and although they might seem almost like bugs in the system, they are actually features that can be taken advantage of.

Professor Elham Kashefi of Sorbonne University spoke at this year's EHPCSW to highlight a less discussed topic in the quantum space – the verification and testing of quantum computing. "A simplistic view is that quantum computing merely represents the next generation of high-performance computing, which we can easily swap in. This is not the case." she said. "Although a lot of the same expertise can be drawn through to quantum computing, it is fundamentally different to classical computing, so questions remain about verification, benchmarking, testing, and even just comparing two different quantum devices."

So, what is the problem with quantum verification? As quantum computers have crept up to the point of supremacy over classical computers, excitement has built around their capabilities. The problem is that now this threshold has been passed, there is no way of verifying if the answers that the quantum devices give are correct, as there are no classical machines powerful enough to fulfil this function.

The important thing to understand, Kashefi says, is that to a certain extent, small and large classical computers are not that different from each other. "Despite the obvious differences in architecture, fundamentally the calculations and the robust behaviour in a home computer and a high-performance computer are the same. That is not the case with quantum computers. With quantum computers, understanding one 10 qubit system does not mean that you understand another 10 qubit system. That is the reason quantum computers are not already widely used. All the big companies like Google and IBM have built 50 qubit machines, but the issue of scaling remains."

Kashefi believes the answer lies in a kind of bootstrapping. By starting with a small quantum computer that is entirely simulated and characterised using classical methods, techniques can then be employed to certify bigger

and bigger computers. By embedding a small quantum computer in a larger one, it is possible to “obfuscate” from this device and essentially change the behaviour of the larger machine so that the noise of the larger machine behaves like the noise of the smaller machine. “By starting with this minimum trust assumption, we can then work upwards to find out things about the larger quantum devices,” she says.

Despite the many announcements that have been made about the achievement of quantum supremacy, it has yet to be scientifically proven. This is why Kashefi stresses the need for due diligence to be done in order to make sure that quantum devices are performing in the way that they are supposed to. “A lot of challenges lie ahead for quantum computing in this area, but at the same time we should not panic, as the solutions are there. Huge communities of theoreticians from error mitigation, error characterization, quantum fault tolerance, and more are developing toolkits. But it is important that industries who hope to use quantum computers are aware that the challenges of benchmarking need to be overcome before we can move forward.

> ***“We should not think that quantum computing can slot in and improve what HPC already does.”***

“What I can say, though, is that we have a clear roadmap for testing in the future. In the same way that there is a roadmap for building bigger and bigger quantum computers from an engineering perspective, there is already a plan for testing these machines as they emerge.”

For Kashefi, speaking at EHPCSW represented an important coming together of two communities that will need to help and work with each other in the near future. “Coming from the quantum community, I have come to this crowd of HPC users to tell them about the challenges that we are facing, and ensure that we have a good dialogue from the start. The HPC community has so much knowledge about parallel programming, and we in the quantum community are bringing this huge capacity for computation, but we still need to work out the best ways for these things to work in synergy.

“We need to work out how to translate the problems of end users in a way that quantum computing and HPC can interact effectively. It might be the case that quantum computing will be very good for specific tasks as a subroutine for HPC. What we should not do is think that quantum computing can just slot in and improve everything that HPC already does. We will need to think carefully how we address the many issues around quantum computing, such as benchmarking and verification, and build a roadmap together. I am very keen as a representative of quantum computing to raise these issues and see what kind of interaction comes out of it.”

As quantum computing continues its march towards the mainstream of computer science, it is only through proper communication that the challenges that lie ahead can be addressed so that its full potential can be realised. Kashefi’s appearance at EHPCSW represents the beginning of a paradigm shift in the HPC community that will happen through collaboration and cooperation between the two communities.

European Master of Science programme for HPC

EUMaster4HPC, the first pan-European Master programme in HPC, is to be launched in September to welcome students hoping to advance their careers using high-performance computers. **Professor Pascal Bouvry**, head of the Computer Science department at the University of Luxembourg, is coordinating the launch of the programme and spoke in Paris about how it has been put together and the importance of industry in its conception.

As part of the EuroHPC strategy to provide clear career pathways into the rapidly expanding field of HPC, a consortium of European universities, HPC centres and industry partners has been selected to design and run the first ever pan-European HPC master's programme. The consortium including the 8 awarding universities led by the University of Luxembourg and, from September this year, will be offering specially designed courses to enhance opportunities in HPC.

Importantly, the European Master of Science programme for HPC is being designed in collaboration between academia and research and supercomputing centres along with industrial and SME partners. It has two main objectives – to educate students in how to use current and future HPC machines and HPC-related technologies, and to educate experts on how best to link HPC activities with industry to drive opportunities for industry to leverage the benefits HPC offers. As such, students will be able to go on to do a PhD in their topic or move on into industry.

The organisations involved in this € 7 million project are currently working together to build the body of knowledge necessary for the programme to be relevant to industry as well as developing materials for lectures and exploring new ways of teaching. This involves teachers and academics working alongside HPC administrators and architects, data scientists, expert users and application developers all focused on the training of highly skilled graduate students.

The curriculum being developed will be modular in its structure so it can be integrated where necessary into other computer science education programmes and new or existing master's programmes. This modularity is also important in terms of lifelong learning pathways for those already working in the industry. People will be able to re-use the modules and put them into a lifelong learning activity. Meanwhile, by bringing good examples from real life into the teaching, the hope is students will become more engaged and interested and will then provide their real-life feedback, which will in turn help improve the lectures – a learning loop.

Essentially, the programme's overall goal is to strengthen mobility between universities, research centres and industry by training the next generation of HPC experts in Europe and make sure industry has a highly-skilled workforce available to develop new machines and their applications.

Pascal Bouvry, Professor and head of the Computer Science department at the University of Luxembourg, spoke at a plenary session in Paris, introducing the programme alongside representatives from all the participating universities. Speaking after this session, he made it clear how important the ties with industry are at this development phase of the project.

"A vital part of the development of the programme is building links with industry," he says. "We don't want a purely academic-focused programme and then have to look for how to apply it in industry. We want industry involved from the start, which is why we are working with ETP4HPC to help define the curriculum, help with exchanges into industry for students, provide teachers from industry and organise training activities."

Bouvry is also part of the board od directors of LuxProvide, the National supercomputer HPC organisation in Luxembourg, as well as being a private company selling computing time, storage and

solutions. LuxProvide is just one of the industrial and SME partners involved in the development of the master's programme and Bouvry uses it as an example of how this type of partnership is working: "We are providing teachers to the programme, with a view to sharing an understanding of industry needs, from space, industry 4.0, health, fintech, manufacturing, process industries and so on.

Of course, the HPC ecosystem is ever changing and Bouvry is careful to acknowledge this when he speaks of how the European Master of Science programme for HPC is being put together. "What we are trying to do is foresee what will be on the market 10 years from now," he says. "But this world is amazing, and we discover new things every day - the question is how we guide this research, so we need to look at the short, medium and long-terms with things like the new machines, quantum computing and how things that are experimental now may emerge into real applications in the future.

"The idea is not only to guide the research but to do it to ensure that it is research with impact."

"So, the idea is not only to guide the research but to do it to ensure that it is research with impact," he continues. "A few years ago, people were talking about excellence in research, now we are talking about excellence with impact."

As well as bringing industry and academia together in these efforts, the project is also bringing standardisation bodies into the picture as they also play an important role, for example in terms of ensuring interoperability. "By bringing all these actors together, not only does it change the way we see and do things for our students," says Bouvry, "but it provides them with a smoother transition and integration into real life and so delivers more impact for them and for industry."

Bouvry is well aware of the challenges he and his team now face in promoting the programme and ensuring the engagement with industry on which it relies. "This is truly a challenge," he says. "If you go into industry

Professor Pascal Bouvry

now and mention HPC, they will look at you and go 'what are those three letters?'".

"But HPC and supercomputing power is a real game changer so what we need to do now is start democratising this," he continues. "We need more talent, we need more people able to understand the potential of this technology and what it can do, who can go to the decision makers and explain how business can really be boosted using HPC. And we need those on the inside, who can be the bridge between decisions in industry to making it happen using HPC. We need to be able to translate the value.

"Up to now, a lot of effort has gone into putting a lot of beautiful equations on the board. This does not talk to a CEO."

The programme is now being actively promoted by all the partners involved. This activity is targeting those studying now and efforts are also focused on expanding the network of universities participating. "What is important is that we have a joint programme where everyone offering this knows what is involved and knows what needs to be taught to master HPC," says Bouvry.

"Industry hires people, not a degree or a curriculum,"

> ***"Industry hires people, not a degree or a curriculum."***

he continues. "So, we need to attract the right people, with the right mindset, those ready to be mobile, move countries, work in a new culture. This way the right people will be able to find the right areas of study for them."

Bouvry has serious ambitions for this programme as well. These involve extending it to other countries, like the US and Japan, and involving other universities as well as looking to bring more students in through additional funding from Erasmus.

"I hope people start copying us and doing what we are doing, particularly in relation to lifelong learning," he says. "I like working in academia, of course, but when you meet people who have been working for 10 years, I like to attract them back to learning, provide them with all the latest information about the new technologies and inspire them to learn more.

"This is particularly true with HPC," he continues., "When people see what HPC can do, they see beyond the boxes. They have a real understanding of how this potential can be applied in their worlds and they are so excited by this potential. So, what we need to do is to provide them with the training to learn more about this potential, and then reward them for taking on this new knowledge for the good of their industry."

Debating the needs of end users

In a lively parallel session that took place on the Industry Day at the EuroHPC Summit Week, speakers representing industry, SMEs, academia and advanced user groups, debated the requirements of those seeking to use the power of high-performance computing and what is needed to ensure they get the most out of this potential. Consensus was soon reached on the need for training, expert support for business and the fact that there is no one-size-fits-all approach that will work across our wide and varied industry landscape.

SMEs first time users from the NCCs perspective

Bastian Koller from CASTIEL/EuroCC, focused on the industrial challenges faced by HPC, highlighting the fact that business comes in many shapes and sizes, so there is no one HPC strategy that fits all. Many don't know about modelling, simulation, AI or data analytics and how they can benefit from them. They are also put off by the initial cost and nervous of the potential ROI, while a skills gap also exists between business and HPC and the associated technology.

Trust is also an issue for business in terms of investing in HPC. This is not something you can buy, and business needs support in building this trust. Openness, patience and the ability to translate complex science into the language of business are required to build this trust, alongside simple, reliable access to expert support and help.

There are actions that could help business understand the opportunities HPC offers and enable them to make the investment of money and time. These include facilitating SME access to the best and most appropriate systems for their needs and support in finding the right partners for them.

He offered a three-tier approach to this:

- Find SMEs who can benefit from HPC – networking, raising awareness of HPC and what it can do
- Convince SMEs – explain how they can use HPC, what it can do and the benefits
- Make it simple – offer training and support for workers involved

User needs from the SME perspective

Guy Lonsdale spoke about his work with Scapos, a Fraunhofer spin-off company which provides expert project management support to collaborative projects, as well as strategic planning for HPC through the FF4EuroHPC project.

Bastian Koller from CASTIEL/EuroCC opening the session

The key focus for any collaborative effort between HPC and business should be addressing the business-oriented challenges. FF4EuroHPC seeks to explain what the experiments do and the results they deliver for participating SMEs, for example modelling and simulation, data analytics, machine learning and AI, and the general issues they face when seeking to transform business processes. One of the key aspects of this work by FF4EuroHPC is to pass on success stories to business so others can follow, helping them go from experiment level to SME innovation.
This support is based on the work of 42 ongoing projects which will provide these success stories and showcase a range of methods and

The audience took part in the interactive workshop examining user requirements for HPC

services that business will find useful. These will focus on the end user needs and not the technology or the machine used – finding the partners and software vendors to support a business case and providing the support the business will need to be able to adapt and integrate them into their specific workflows and so use the technology successfully.

Lonsdale was keen to emphasise the point that the successful adoption of HPC by industry is not going to be about the machines being used, but the needs of the business and how the machines can meet those needs.

The needs of advanced user communities

Elisa Molinari, representing Centres of Excellence (CoEs) spoke about the needs of advanced user communities and CoEs, which provide support to industry and academia using large sets of applications. This is a highly connected ecosystem that feeds into the needs of SMEs but one that, in turn, has urgent needs that can be met by HPC centres and consortia and from the EuroHPC community as a whole.

A key, critical issue that needs support, is the software stack, where bottlenecks occur due to languages and programming models. Standards are key to portability but are still lacking particularly for GPUs, while support for code development and complex workflows is also required by advanced users and developers. The scheduling of early access to systems is also required along with system-specific support for porting codes to new systems.
Molinari also mentioned access to detailed statistics from HPC centres and consortia being required, along with more efficient data handling and provision of computer time for scaling tasks.

The objective of all this support is to interface CoEs and competence centres in order to improve and engage the pool of advanced users, which in turn feeds into the competence centres and CoEs to provide the support needed.

HPC and SMEs

Jeanette Nilsson, of NCC in Sweden, spoke about what SMEs need from HPC, starting her session by posing the question that if SMEs make their money by making things and selling services, why do they need HPC in the first place? And the answer she put forward is that if SMEs understood the potential benefits HPC can deliver, then they would soon understand how HPC could help them.

Engagement with business by the HPC community in Europe will be key to enabling this understanding and so it is necessary now to identify people who understand the possibilities. This could be done by reaching out to business through their own networks and through competence centres where strategies for how SMEs can benefit from European investment into HPC can be developed.

It will also be important to build trust and develop networks that enable business to enhance what they already do through HPC by offering help and support. Successful industrial and public sector applications of HPC use should also be highlighted so that examples of what is possible and real benefits can speak for themselves.

User needs panel session

Following the opening remarks by the speakers, a roundtable session took place during which the panel focused on the main challenges facing the HPC community in terms of expanding how HPC is more widely used by industry and SMEs.

Consensus emerged that businesses don't have time to spend on HPC, so require mentoring and the support of experts. The main challenges faced here are that people do not know that opportunities with HPC exist, so they need expert help applying what does exist to their particular needs. This highlights the issue that the pool of people able to bridge the gap between computing science and industry needs enlarging and more experts are needed in the system which will come through the offer of better opportunities for women and attracting scientists from other disciplines. The EUMaster4HPC programme (see page 18) was seen as a positive step in this direction, and this programme is also key to addressing the need for lifelong learning, with better training essential for those already in work.

The notion that aspects of working with HPC in an academic sense do not apply in the industry context was also a view expressed on the panel. It is important to view HPC opportunities in the commercial context, that using HPC should solve a business problem. Meanwhile, the idea that SMEs are interested in whether something has been done before and has been peer reviewed, for example, is not important – solving the problem they have or improving what they are doing is all that matters.

Co-designing systems that use HPC in industry was also discussed, with the idea that these systems should mirror what those working in industry already know being essential for smooth and successful integration. Again, communication about the benefits at the start of this process is essential, while a clear vision of how the systems will work and what it will deliver from the outset is also vital. Software developers, consultants and expert users should work with industry on this co-design process.

George Em Karniadakis, Professor of Applied Mathematics at Brown University, provided a memorable remote keynote on data-driven computing and in particular how physics-informed neural networks (PINNs), an idea he first introduced in a paper of 2017, will be used in high-performance computing to transform many areas of research in the future.

Transforming research using PINNs

In 2017, George Em Karniadakis published a paper introducing the idea of physics-informed neural networks (PINNs), a neural network architecture that integrates data and physics in a seamless way. "Our idea was to encode physical laws into neural networks," he says. "In science and engineering, we do not have the huge amounts of data that an entity like Facebook has, but what we do have are physical laws that can provide us with insight into our domains. The concept of PINNs allows for neural networks to be shaped according to these physical laws using a kind of semi-supervised learning approach."

Since the publication of his paper, the topic has exploded in popularity and was covered in a talk by Jensen Huang, CEO of NVIDIA, in his talk at SC19, where he talked about how the company's AI developments could help to improve these types of code. NVIDIA now has a whole division working on this topic, which demonstrates the huge potential they see in PINNs.

"In most realistic situations in science and engineering, you will know part of the physics and also have some data available," says Karniadakis. "These are the situations that PINNs are most useful. In a PINN, an extension of the standard neural network is used to enforce the given physical laws using automatic differentiation, the same technique that is widely used to backpropagate errors in neural networks. The simplicity of our method will allow many teams to reduce the size of their codes from thousands of lines down to hundreds. By defining the differential equations of physical laws using automatic differentiation, PINNS eliminate what I like to call the tyranny of mesh generation. Defining a mesh of a gearbox, an engine or a wing is a hugely time-consuming task – entire companies specialise in this work, but our method gets rid of this problem."

"The simplicity of our method will allow many teams to reduce the size of their codes from thousands of lines down to hundreds."

At EHPCSW22, Karniadakis used his keynote address to showcase several ways in which PINNs can be used. One particularly compelling case involved flow visualisations of brain aneurysms. Although these visualisations are already useful to brain surgeons, Karniadakis showed how it was possible to extract the velocities and pressures from these visualisations using a PINN-based framework called hidden fluid mechanics that encodes the Navier-Stokes equations into a neural network. Using this method, his group was able to provide accurate information about the stresses being placed on the aneurysmal tissue, which can be used to predict the point and time of rupture.

Despite the impressive array of applications that Karniadakis demonstrated for PINNs, he also sought to temper expectations about this method in the short term. "PINNs are not a panacea," he says, "and the reason for this is that, compared to classical methods, they are powerful but computationally expensive. We are currently reliant on advances in hardware made by NVIDIA. Of course, to apply our methods for real world applications, we need to scale up the method so that it can work efficiently with parallel computing."

Addressing the audience of HPC users, Karniadakis outlined how PINNs can be scaled to work on large parallel computing architectures. "With these kinds of algorithms, you interrogate the system many times with your iteration loop, so you need to be efficient. Scaling neural networks on parallel computers means you have to take into account the fact that you will

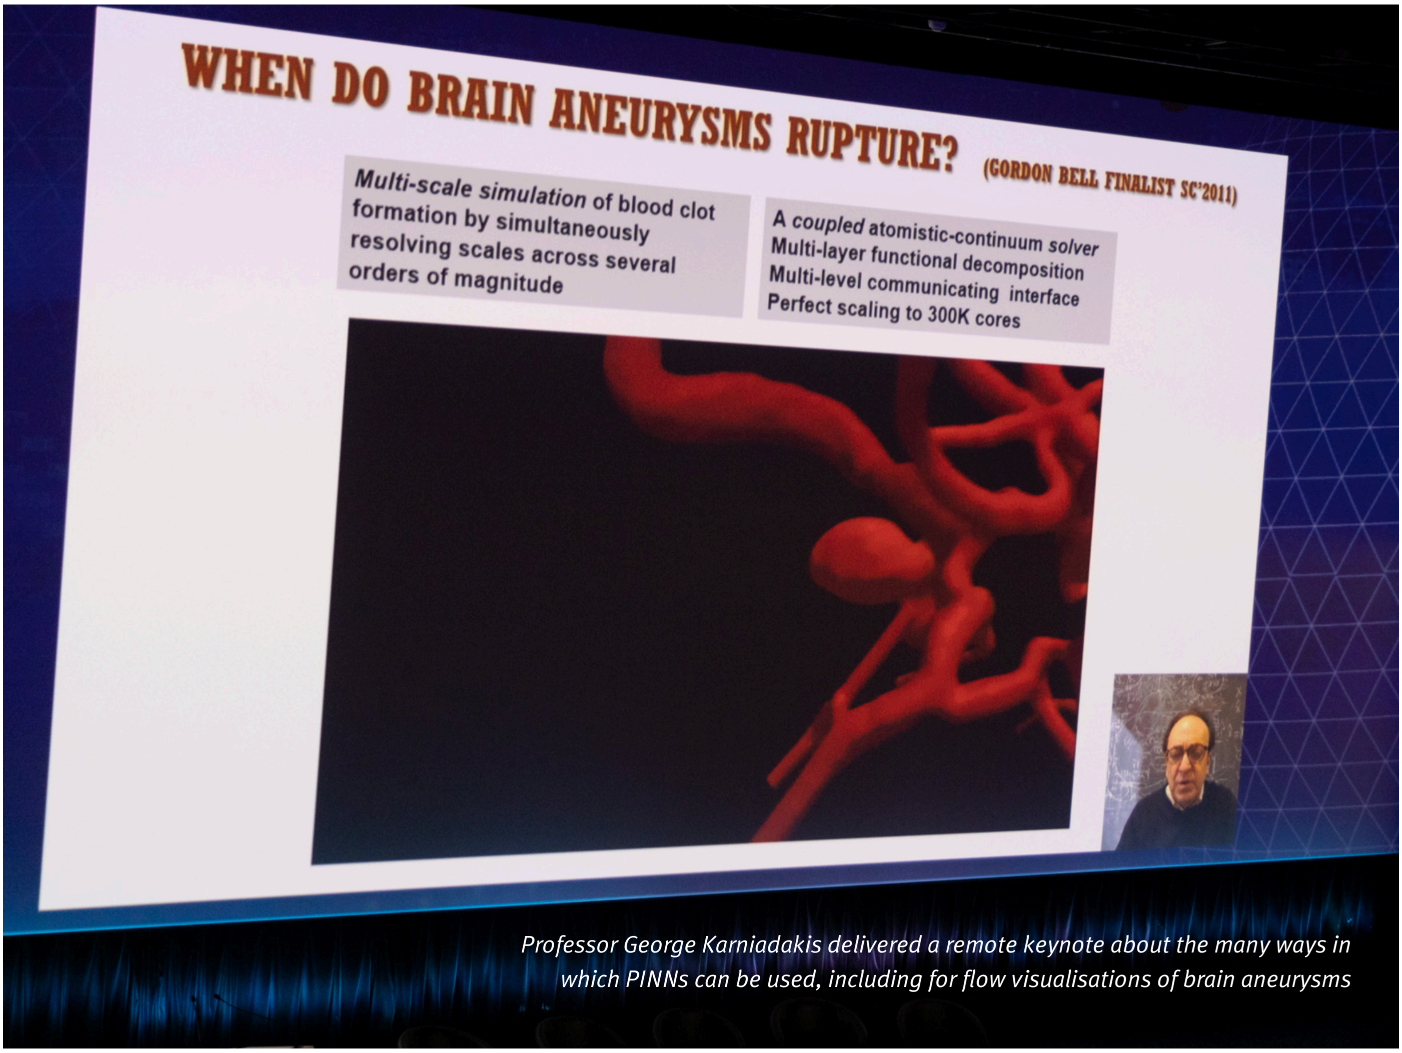

Professor George Karniadakis delivered a remote keynote about the many ways in which PINNs can be used, including for flow visualisations of brain aneurysms

be using a combination of CPUs and GPUs. In terms of data parallelisation, you must split the data so that each processor is working on a different part. This is not entirely scalable, so we use the concept of domain decomposition, well-known in the world of HPC, to parallelise spacetime, assigning hypercubes of spacetime to each processor. Working like this, we can assign different neural networks on the separate subdomains."

PINNs have the potential to be used in many domains of research and have already been successfully deployed in the fields of fluid mechanics, solid mechanics, and geophysics. Big industrial players are also getting involved, including Bosch and NVIDIA, the latter of which has developed a tensor flow implementation that is free to use for academia. "There are now many smaller start-ups also working with PINNS," continues Karniadakis. "One company I am working with called Anailytica is using them for renewable energy design and optimisation. What we can take from all this is that PINNs are agnostic in terms of their application – all you have to do is dial in the relevant physics into the neural network. If you do it fast, accurately, and use any data you have, it will work well.

"PINNs are agnostic in terms of their application – all you have to do is dial in the relevant physics into the neural network."

"In this new paradigm, we like to say that no data is left behind. For example, I worked on a project for Boston Harbour about ocean acidification using satellite data and in situ measurements using pH sensors. On the one hand, we had top of the range pH sensors worth around $ 10 000 each, and on the other we had the $ 100 sensors being deployed by students. On top of this, we also had satellite data, sometimes uncalibrated. This mixture of high and low fidelity data would in the past have been difficult to work with, but our neural networks can be trained using multi-fidelity and multi-modality data."

In the world of science and technology, it is the simple ideas that are often the most revolutionary. It is this simplicity that give Karniadakis' PINNs the potential to transform many areas of research in the future, and the world of HPC will surely be heavily involved in their uptake into the mainstream of computational science.

Celebrating HPC success

Every year PRACE honours scientists working with HPC on their research with awards that celebrate their achievements and bring their work to wider audience. In Paris, five scientists took to the stage to receive their prizes, including **Dr Marija Vranic**, who was awarded the prestigious Ada Lovelace Award 2022 for her work on plasma.

An important part of the EuroHPC Summit Week 2022 was the awards ceremony that took place to recognise the outstanding work done by scientists the previous year. Five awards were presented, the first being the prestigious PRACE Ada Lovelace Award, awarded annually to a female scientist who makes an outstanding contribution to and impact on HPC in Europe and the world, and who serves as a role model for women who are at the start of their scientific careers.

The award is named after the Countess of Lovelace, a British mathematician who lived in the 19^{th} century and worked with Charles Babbage on the machine they called the Analytical Engine — one of the first precursors of computers. Many historians regard Ada Lovelace's contribution to this mechanical calculator as the very first algorithm — and herself as the first person to be rightly called a programmer.

The 2022 award was presented to Dr Marija Vranic, Invited Professor at the Instituto Superior Técnico of the University of Lisbon, for her work developing pioneering techniques for representing quantum effects in extreme plasmas. She is also actively involved in improving the visibility of women in physics and HPC — for instance, by reducing the gender imbalance of invited speakers at scientific conferences as well as being part of the Diversity and Gender balance group at IST, and mentoring a group of graduate students at IST to form a Women in Physics group, set up to provide female scientists with a place to meet, discuss and get advice.

Dr Vranic accepted the award from Dr Laura Grigori, Chair of the PRACE Scientific Steering Committee, who said: "Dr Marija Vranic's contributions, achieved partly through PRACE-supported projects, are crucial to advance our understanding of plasmas in extreme conditions".

Dr Marija Vranic receives the Ada Lovelace Award 2022 from Dr Laura Grigori, Chair of the PRACE Scientific Steering Committee

Dr Vranic presented her work in a keynote address during a plenary session of the Week. Read more about this on page 28.

As well as providing this detailed overview of this important work, Dr Vranic was also keen to use the honour to highlight the important role women play in HPC and science. "With this award, PRACE raises awareness for women in HPC and science, which is extremely important," she said.

"I am personally painfully aware of the gender imbalance in HPC and physics — in plasma physics especially," she added. "I grew up and studied in Serbia, where women are more represented in

"I grew up in Serbia, where women are more represented in science. It was a shock to learn that there are much fewer female physicists in most other European countries."

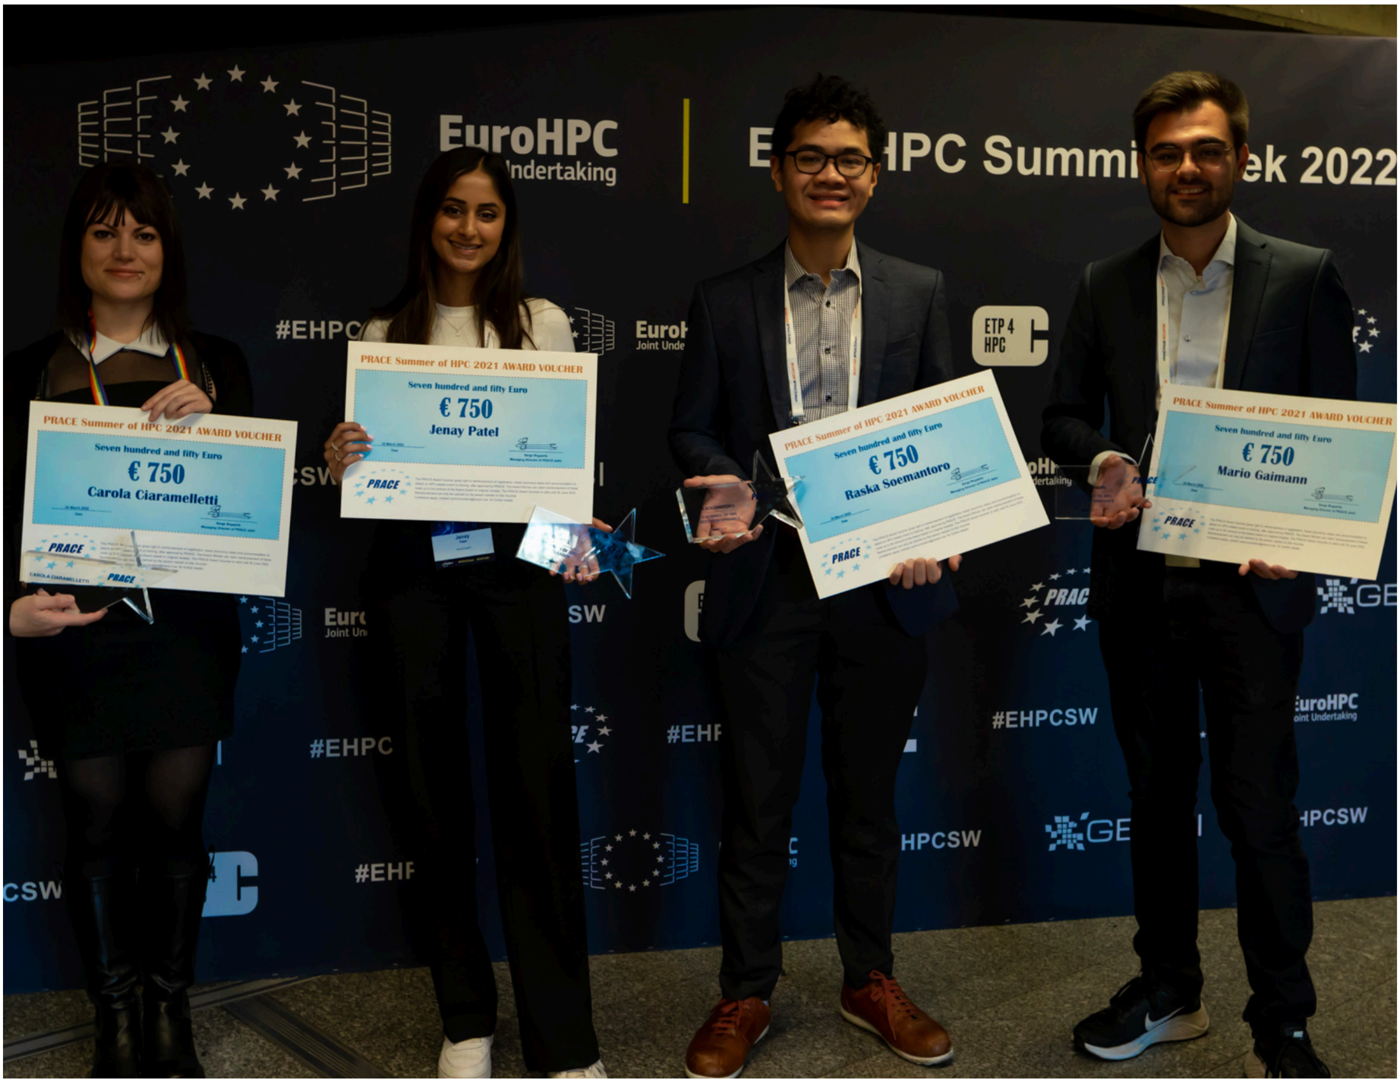

Four young scientists honoured for their work at the ninth PRACE summer of HPC, 2021. They are, from left, Carola Ciaramelletti and Jenay Patel, who received HPC Ambassador Awards along with Raska Soemantoro and Mario Gaimann, who received Best Performance awards

science. So, for me, it was a shock to learn that there are much fewer female physicists in most other European countries."

Four other young scientists received awards at the ceremony, honouring their successes at the ninth PRACE Summer of HPC 2021. The event, which was held remotely on this occasion due to COVID restrictions, allows students from all scientific disciplines to spend eight weeks working on a dedicated HPC project, while being supported and mentored by a tutor from an PRACE HPC centre.

The PRACE Summer of HPC 2021 Best Performance award was presented to Mario Gaimann and Raska Soemantoro for their collaboration on their *Automated Classification for Mapping submarine structures by Artificial Intelligence strategies* project.

The two young scientists developed an automated tool to recognise seabed structures. They worked hard researching the various AI methods available to them and then deciding which methods were best to use. They also researched, identified and learned various technologies and tools to use for their project, with the outcomes far exceeding initial expectations.

They will now aim to publish their results through funding via the HPC-Europa3 scheme.
The PRACE Summer of HPC 2021 HPC Ambassador Award was also a double honour, this year being presented at the ceremony to Carola Ciaramelletti and Jenay Patel, for their work on the IT4I project Molecular Dynamics on Quantum Computers.

The successful collaboration on the project involved the two young scientists blogging, keeping readers up to date with their research throughout the summer, explaining complex science in a way that enabled wide understanding by members of the non-academic community. This is an important quality for HPC ambassadors.
As well as writing about their work, and producing exceptional results, Ciaramelletti and Patel also made a project video that also communicated their work in an accessible and engaging way.

Q&A

Dr Marija Vranic

winner of the PRACE Ada Lovelace Award 2022

Dr Marija Vranic, Invited Professor at the Instituto Superior Técnico of the University of Lisbon, is the winner of the 2022 PRACE Ada Lovelace Award for HPC for her outstanding impact on HPC in Europe. Dr Vranic has not only pioneered techniques for representing quantum effects in extreme plasmas, but has also been actively involved in improving the visibility of women in physics and HPC — for instance, by reducing the gender imbalance of invited speakers at scientific conferences.

The award was presented to her at EuroHPC Summit Week 2022 / PRACEdays22, where she also gave a keynote speech entitled "Extreme plasma on a supercomputer". We caught up with her to talk about her work, the importance of providing visibility to women who have succeeded in the field of HPC, and the future of the European HPC ecosystem.

Could you tell us about your research?

I work in extreme plasma physics, which aims to improve our understanding of some of the most intense environments in the universe, such as those found around neutron stars. The combination of extreme fields with matter in these places means there are a lot of particles undergoing extreme acceleration and radiation, so there are a lot of things going on at the same time!

My work is focused on laser-matter interactions. We have very intense lasers nowadays and we will have even more intense lasers in the future, and the fact that we have them is owed to an amplification scheme invented by Donna Strickland and Gérard Mourou that won them the Nobel Prize for Physics. Thanks to their work, the intensity of lasers available has increased by many orders of magnitude over the past 20 years, and this will likely increase by another three orders of magnitude in the next 10 years with facilities that have been planned.

Thanks to these lasers, we can recreate extreme environments like those found around neutron stars in the lab by firing the lasers at matter. The thing that makes extreme environments hard to study is that there are many things going on at the same time and at various scales, which is why the use of supercomputers is essential to our work. The size of the system and the scale of the processes involved are sometimes very different, so you need very high resolutions to capture the relevant physics. Luckily for our type of research, we can decompose the problem efficiently and use parallel computers to solve it.

Dr Vranic giving her keynote: "Extreme plasma on a supercomputer"

What made you pursue this line of research?

During my third year of studies, I did an internship abroad at the University of Belfast. They had just started building a big laser system called TARANIS, so it was a very exciting time. I had seen something new at the time called proton radiography, where you accelerate protons using a laser and then use those protons to take a picture of another laser propagating at 90 degrees. This allows you to see events happening at a femtosecond frame-rate, so you can literally take a picture of something happening at the speed of light. I found this fascinating, and I was able to go into the lab and play around with it a bit.

At the same time, I knew even earlier that that I wanted to do computational physics, specifically theory and modelling, so in the end when I decided to do a PhD I chose a group that was strong in this.

What role has PRACE played in your career?

PRACE has been extremely important for two reasons. The first is for having these regular competitive calls. Every year you know there is a deadline before which you have to plan your next project, which gives you a bit of a routine in terms of setting goals.

Second, and most importantly, PRACE gives you access to HPC wherever you are from in Europe. You don't have to come from a rich country where you have your own supercomputers next door. All you need is a great idea and the ability to demonstrate

"PRACE gives you access to HPC wherever you are from in Europe. You don't have to be from a rich country where you have your own supercomputers next door."

that you know how to use these resources. And even if you don't know how to use them or if your code is not yet ready to use, PRACE gives you training and the opportunity to build your code to a point where it is ready for use on HPC. Most of my research has been based on PRACE allocations, so it has been very important for my career as my university does not have the resources to even maintain a Tier-0 machine like MareNostrum.

What does winning the Ada Lovelace Award for HPC mean to you?

First of all I find it remarkable and an honour to be associated with the name of Ada Lovelace. I have been following this prize since PRACE started giving it. I think it is a great initiative to showcase role models for future generations, because we have very few women in computational science. The situation in HPC is even worse than physics – often you will find you are the only woman at a conference! I think it's important to show the younger generation of women that they can do this, and they should not be dissuaded to pursue what they want to do just because it doesn't fit the current pattern.

What more can be done to encourage women into computational science and HPC?

I think early exposure is very important. My story is that I did an internship when I was in high school – it was a computational physics science project that made me realise that I wanted to pursue it.

This project gave me the opportunity to take part in the full cycle of research, from a project proposal, through writing a paper and ending by presenting at a conference. This allowed me to gain a full view of the scientific method, which you do not really get from learning physics at school. We also had a strong programming course at our school in Belgrade which I enjoyed a lot. I was torn between going into physics or computational science, so in the end I chose computational physics which combines both.

What do you foresee for the future of the European HPC ecosystem?

I would say there is a twofold answer to this question. One more fun but maybe less secure direction is to move towards new technologies like quantum computing. I have students here at the conference and one of them had a small project on quantum computing. I believe investing some time in research focusing on quantum algorithms, even though the technology is not ready to run them, may pay off.

On the other hand, I think it is important is try and set standards, because HPC changes a lot from year to year, and codes become obsolete quickly. If people are doing research, it can take long a long time to come up with the right solution, and then if these codes become obsolete within a year it can feel like wasted effort.

For researchers to stay on board, it would help if we could prolong the lifetime of the codes we make. If we change architectures, we should try to find ways to port the existing platforms to the new architecture without big losses. I think this should receive special support in terms of funding, because it's not researchers who are necessarily going to come up with the best solutions. Maybe we need a designated computer science or engineering action for this.

Panel discussion

Emerging applications, models and implementation

The last decade has seen a range of HPC-oriented applications, computing models and implementations being established. More recently, some more potentially disruptive ones such as quantum algorithms have also started to appear. This year's EHPCSW panel discussion looked to explore these emerging advances from the various perspectives of the assembled panel, as well as from the members of the audience.

Guy Lonsdale once again took his place as moderator of the annual discussion, which over the last few years has without fail brought about some lively debate about the future of HPC. Talks began around what the panel thought were likely to be the most disruptive methodologies and application areas in the coming years. Ada Lovelace Award winner Marija Vranic began by mentioning a topic that had been widely discussed all week: quantum computing.

The quantum information community has been working on quantum algorithms for many years, but it is only in the last few that most domains of research have begun to take real notice in their potential. Much work still needs to be done to make them more effective, said Marija Vranic, and there is a need to develop them further and expand the number of problems they can solve.

Quantum expert Kristel Michielsen suggested that although quantum algorithms will clearly have a huge impact in the long run, it is the fields of science that have yet to truly leverage the power of HPC that will see the greatest returns in the near future. She gave the example of neuroscience, which has only truly begun to explore HPC with the advent of the Human Brain Project. It appears that neuroscience benefits more from data driven computing, such as image analysis, than simulations, so machine learning and AI will help to drive this field forwards. Finding other areas of science that can similarly benefit from HPC can be just as disruptive as any new type of algorithm, argued Michielsen.

Asked whether disruptive applications will drive computing models and the design of future HPC models, Michielsen answered that it is more the case that everything has an influence on everything else. AI and data driven science have influenced simulations, and we have even started to see the emergence of "quantum inspired" classical simulations and methods that do not require the use of quantum computers.

Of course, it is important to note that quantum computing will not necessarily take over the solving of entire problems; instead, certain parts of the problems will be more suited to quantum computing, whereas other parts will be better suited to classical computing. Kristel Michielsen's work on the HPCQS project which is developing hybrid quantum-classical computers will be a driving force behind this discovery phase, but for the moment it appears that there will still very much be a place for large parallel computers for a lot of research.

Harnessing the power of quantum computers will require a big investment into creating expertise in using them, and it was agreed that this should begin at the level of computer science degrees at university. Young researchers will be ones to use quantum computers to their full potential, and so it is important that they are given the chance to experiment and play with them at the early stages of their careers. On top of this, it is important to ensure that industry is part of the quantum ecosystem from the start, introducing real use cases which can be worked on and refined by quantum experts.

Marie-Alice Foujols of CNRS brought the climate modelling perspective to the discussions. Climate models are legendary for their capacity to bring various types of models and algorithms together, but this does make it difficult to quickly and entirely transform the ways in which they work. Foujols did mention that machine learning was starting to make

The panel

Guy Lonsdale (host)
FocusCoE

Marie-Alice Foujols
Climate modeling and HPC expert, CNRS, Institut Pierre Simon Laplace, Paris, France

Kristel Michielsen
Head of RG 'Quantum Information Processing' and JUNIQ
Jülich Supercomputer Centre

Marija Vranic
Instituto Superior Tecnico, Universidade de Lisboa

The panel chaired by Guy Lonsdale at the lecturn, with speakers (from left) Marie-Alice Foujols, Marija Vranic and Kristel Michielsen

inroads into their work, with many PhD students working with these methods. Tracking of cyclones from models has been made possible using machine learning, and it is also being used for some processes in the global models.

Creating new codes is costly and time-consuming. Marija Vranic said that the idea that a code could be worked on for many years and then become obsolete very quickly due to the introduction of new hardware architectures could have a demotivating effect on those putting the time in to create them. It is therefore important that, moving forwards, the community looks for ways to make algorithms more easily portable from one architecture to another, or at least to provide support when this needs to be done.

An issue that HPC has always suffered from is addressing how to bridge the gap between the knowledge held by computer experts, domain-specific scientists, and industrial researchers, each group which brings very different skills and expertise to the table. One member of the audience, a computer scientist, described how he spent at least a year and a half working with industrial partners before he truly understood the problems they were trying to solve, and suggested that training people with dual competencies from the beginning would be of great help. He asked the panel how they thought we should approach this issue in the future.

Kristel Michielsen described the setup they have at the Jülich Supercomputing Centre (JSC) to try and help this situation. Groups of scientists from specific fields work in what are known as Simulation and Data Laboratories, where they spend half of their time carrying out research and the rest working with researchers from outside to help them port their codes to the JSC systems. Then, to help industrial users, there is the Industry Relations Team, which employs researchers to work in a consultative manner with industrial researchers to help them use their codes with JSC machines.

Another topic that thread itself through the week of talks at EHPCSW22 was how to improve the energy efficiency of HPC. Marija Vranic joked that the best solution from a plasma physics perspective was to work out how to carry out efficient nuclear fusion. Failing this, she suggested that writing code as efficiently as possible was a good first step. Marie-Alice Foujols agreed, saying the climate modelling community always had this in mind, and that one of the simplest steps towards reducing energy use was sharing results within the community.

Kristel Michielsen suggested the issue ranged across building efficient hardware, writing efficient software, and using the right resources. Quantum computing is very energy efficient and so will surely have a role to play – another good reason for every researcher out there to consider what parts quantum computing could help with in their codes.

Herbert Zeisel of EuroHPC was in the audience, and asked the panel how they thought energy efficient code writing could be incentivised? They suggested a good first step would be providing information about how much energy each simulation uses, as well as providing support to improve energy efficiency in the same way support is already provided to help parallelise code.

Executive Summary

www.prace-ri.eu