PRACE
Annual Report

2019

www.prace-ri.eu

Annual Report

PRACE Annual Report 2019
www.prace-ri.eu

Editor-in-chief
Marjolein Oorsprong – PRACE Communications Officer

Editors
Fabio Affinito, CINECA
Evangelia Athanasaki, GRNET
Sanzio Bassini, CINECA
Mauro Bianco, ETH Zurich
Florian Berberich, FZJ
Serge Bogaerts, PRACE aisbl
François Bodin, IRISA
Pascal Bouvry, University of Luxembourg
Mark Bull, UEDIN
Zuzana Cervenkova, IT4Innovations
John Clifford, PRACE aisbl
Stelios Erotokritou, CaSToRC
Jacob Finkenrath, CaSToRC
Paul Gibbon, FZJ
Sergi Girona, BSC
Troels Haugbølle, NBI, KU
Anni Jakobsson, CSC
Chris Johnson, EPCC
Earl Joseph Jordan, KTH
Zoltan Kiss, KIFÜ
Leon Kos, UL
Silke Lang, PRACE aisbl
Tiina Leiponen, CSC
Walter Lioen, SURFsara
Thomas Lippert, FZJ
Martti Louhivuori, CSC
Lee Margetts, University of Manchester
Michal Merta, IT4Innovations
Petia Minkova, NSCA
Cristiano Padrin, CINECA
Oriol Pineda, BSC
Janez Povh, UL
François Robin, CEA
Núria Lopez, ICIQ
Gabriela Scipione, CINECA
Catarina Simões, PRACE aisbl
Alan Simpson, EPCC
Raffaele Solcà, ETH Zurich
Veronica Teodor, FZJ
Volker Weinberg, LRZ
Simon Wong, ICHEC

Article Editor
Harry O'Neill, Insight Media

Project Manager
Lydia Faircloth, Insight Media

Graphic Design
Ellie Mone, Insight Media

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: June 2020 **ISBN: 9789082942255** Published by Insight Media, Bristol, UK

Welcome

Serge Bogaerts

Before I begin my reflections on 2019, I will address the "elephant in the room". I understand that many of you reading this publication in 2020 might be wondering why there is no mention of the global COVID-19 pandemic that has swept across the world. It is now of course at the forefront of all of our minds in everything we do, but as this publication is a reflection on PRACE and its surrounding ecosystem in 2019, we will not be referring to it from hereon. Needless to say, PRACE is now actively involved in international efforts within the global HPC community to lend its expertise to this ongoing crisis.

2019 has seen many developments in the world of high-performance computing. The EuroHPC Joint Undertaking is now in full swing, and inside (page 6) you will find an overview of the activities of the eight countries that will be hosting a brand new set of high-performance computers, three of which will be precursors to exascale machines and therefore at the very pinnacle of HPC technology.

Of course, to make the most of this new hardware there is a pressing need for similar developments in codes and algorithms that can push these machines to the very limits of their capabilities. In our section on forward-looking software solutions (page 47) you can find out more about some of the great work that is being done in this area, all of which was supported by PRACE.

PRACE as an organisation has taken a number of steps over the past years to strengthen its collaboration with similar bodies across the globe. The Memorandum of Understanding on information exchange that was signed along with RIST and XSEDE in 2017 was a defining moment in formalising this need for closer collaboration and, as our interview with the president of RIST Masahiro Seki (page 38) shows, all sides are keen to continue in this direction as we move forward, particularly in regards to the sharing of expertise at the human level.

As usual, this annual report reserves a space for our Scientific Steering Committee (page 5), our User Forum (page 40), and our Industrial Advisory Committee (page 12), groups that consist of highly dedicated people who provide invaluable advice to improve our services. We also include a showcase of a number of success stories from PRACE-supported projects (page 32), which highlight how important the provision of core hours on Tier-0 machines and advice from HPC experts is to the success of scientific and industrial research and innovation.

Last but not least, I would like to thank all of you who have contributed your time, knowledge, effort, and support to PRACE in 2019, and will be happy to be able to count on you again in 2020.

Serge Bogaerts, Managing Director PRACE aisbl

Contents

From the desk of

The Scientific Steering Committee

The 2019-2020 period has been very hectic for the PRACE Scientific Steering Committee. The identification of the three sites for the new pre-exascale machines and the five petascale machines in June 2019 was excellent news for all the community.

These new supercomputers bound to the EuroHPC JU will allow simulations in all fields of science and even in new emerging fields such as artificial intelligence. The challenges ahead are enormous and the need for training and adaptation of the codes to the new architectures is one of the most important concerns for the scientific community.

Such efforts, sometimes performed through jobs without a clear scientific/technical career path, are fundamental to get the best from the pre-exa and exascale systems. The challenges with the new facilities have prompted us to help define new access paths for the supercomputers.

Merging the scientific view of the SSC and the corresponding perspective of the EuroHPC JU will result in a new flexible and adaptive system to grant access to the supercomputer power in Europe, for instance suggesting excellence tracks or defining the needs for community access.

In addition, important discussions about the use of supercomputers for industrial initiatives are being conducted. The appearance of new initiatives through the extension of the Centres of Excellence in dedicated areas and the new Centres of Competence have enhanced the robustness of the computational ecosystem, but enhanced interaction between all these different stakeholders is needed to attract and train the best scientists.

The challenges with the new facilities have prompted us to help define new access paths for supercomputers.

Núria López, Chair of the PRACE Scientific Steering Committeee

The challenges are big and the needs for education are enormous, but the scientific and technical quality of the European landscape is huge and we need to make a collective effort to do our best.★

EuroHPC Hosting Sites

In June 2019, the EuroHPC Joint Understanding (JU) announced that it had selected eight sites across eight different countries to host the next generation of European supercomputers. These machines will support the development of major applications in domains such as personalised medicine, drug and material design, bio-engineering, weather forecasting, and climate change. Here, we summarise the plans for each site

Bulgaria

The petascale machine based in Sofia is interconnected via networks with terabit/sec capacity, and is available throughout the Balkan countries - Bulgaria, North Macedonia, Greece, Turkey, Serbia, Romania, and possibly Ukraine (pending a decision from the European Commission), says Professor Stoyan Markov, head of the Bulgarian National Centre for Supercomputing Applications (NCSA).

This large-scale project provides many advantages for this region, not only for researchers and scientists, but for citizens, businesses and industry. In Bulgaria, for example, the most prominent scientists from several universities and institutes of the Bulgaria Academy of Sciences participate in different workgroups with NCSA's team.

NCSA's resources are currently being used to carry out research into a range of computationally intensive topic of research, from muremical homogenisation, to supramolecular organisation, and from climate and pollution research to drug design.

The EuroHPC hosting entities open up unique chances for the European HPC ecosystem to reach smaller countries with less developed HPC infrastructures, and for scientists to receive access to world-class computers - a strategic resource for European citizens, industry and science. The benefits for all these countries will stimulate positive developments in many sectors of society and economy: health, climate change, industry, cybersecurity, energy, agriculture, public sector, pharmaceuticals, automotive industry, shipbuilding, and technology among others.

Czech Republic

In June 2019, IT4Innovations National Supercomputing Centre, located at the VSB – Technical University of Ostrava (Czech Republic), succeeded in a pan-European competition for hosting a EuroHPC petascale system, which will be acquired within the implementation of the EuroHPC Joint Undertaking.

The EuroHPC petascale system, called EURO_IT4I, to be procured by IT4Innovations, will be among the most powerful HPC systems in Europe. Nowadays, the HPC system aggregated performance would rank it 21st worldwide and 6th in Europe. When installed in 2020, its aggregated performance will be among the 50 most powerful supercomputers worldwide and 10 most powerful supercomputers in Europe. The HPC system is designed to coherently respond to the needs of its user communities, addressing complex scientific and industrial challenges, including standard numerical simulations, demanding data analysis, and artificial intelligence applications.

The IT4Innovations National Supercomputing Centre also succeeded in the call for a EuroHPC pre-exascale system, being a member of the LUMI (Large Unified Modern Infrastructure) consortium, coordinated by

Finland and participated in by Belgium, the Czech Republic, Denmark, Estonia, Iceland, Norway, Poland, Sweden, and Switzerland. The HPC system will be installed in 2020 in Kajaani, Finland,

The procurement of the EuroHPC petascale system by IT4Innovations and involvement in the LUMI EuroHPC pre-exascale system consortium is a very significant contribution by the Czech Republic to the successful implementation of the EuroHPC Joint Undertaking and, in general, to strengthening the competitiveness of the EU in the areas of digital technologies. These achievements will provide the Czech and European scientific and industrial communities with the opportunity to use one of the most powerful computing facilities for research, development, and innovation.

IT4Innovations is a longtime member of PRACE and will thus become a natural bridge between PRACE and EuroHPC, sharing its investments in the new resources with both entities and ensuring a top level of operations and support for European science and industry.

Finland

One of the EuroHPC pre-exascale supercomputers, LUMI, will be located at CSC's data center in Kajaani, Finland. The supercomputer will be hosted by CSC together with the LUMI consortium including nine European countries: Finland, Belgium, the Czech Republic, Denmark, Estonia, Norway, Poland, Sweden, and Switzerland.

This consortium provides a high-quality, cost-efficient and environmentally sustainable HPC ecosystem based on true European collaboration. Nearly all LUMI consortium countries are also members of PRACE, and this European HPC collaboration was a key element when forming the LUMI consortium. CSC has been an active player in European HPC collaboration throughout the years, starting as early as 2002 with the DEISA (Distributed European Infrastructure for Supercomputing Applications) project. Based on this collaboration that has continued over the years in several PRACE projects, the LUMI consortium was built, and eventually EuroHPC decided to place one of its world-class supercomputers in Finland.

Kimmo Koski, Managing Director of CSC said: *"EuroHPC continues the European collaboration on high-performance computing, from which European researchers have already benefitted from over the years. The investment will make CSC's data centre one of the world's largest players in the field of high-performance computing."*

Italy

On 4-5 June 2019, the 7th EuroHPC Governing Board approved funding for three pre-exascale and five petascale supercomputers.

Italy will host one of the pre-exascale systems, funded by the EU and the Ministry of Education, University and Research, conceived and managed by Cineca, the Italian HPC data centre, in collaboration with the National Institute of Nuclear Physics and SISSA, the International School of Advanced Studies.

The Italian supercomputer, Leonardo, will be one of the most powerful supercomputers in the world. Leonardo will deliver between 220 and 250 peak petaflops, spread over 5 000 compute nodes and laced together with a 200 GB/sec interconnect. The system will also include 150 PB of storage and the entire machine is expected to consume 9 MW. Leonardo will foster the convergence of HPC, AI and high-performance data analytics (HPDA).

During 2019, Austria, Slovakia, Slovenia and Hungary – all PRACE members – joined the Italian initiative, establishing the Leonardo Consortium of the Central-Eastern European Countries. The objective is to improve collaboration, foster the European HPC

The LUMI wolf, symbolising the LUMI consortium coordinated in Finland and involving the Czech Republic

ecosystem and develop common activities of joint infrastructures, research and innovation projects in the European context within EuroHPC.

Luxembourg

In the context of EuroHPC, the Luxembourg project Meluxina, which aims to install a 10 petaflop supercomputer at Luxprovide in Bissen, was selected and will obtain European co-financing (around 35% in CAPEX). The new centre should be operational by the end of 2020.

The Luxembourgian digital innovation hub, Luxinnovation, will steer the EuroHPC National Competence Centre in partnership with the University of Luxembourg and Luxprovide, providing a one-stop shop for HPC-related competencies and resources in Luxembourg. Professor Pascal Bouvry is the national delegate for PRACE, and Dr Sébastien Varrette his advisor. The University of Luxembourg is also active in the PRACE-6IP project.

Slovenia

Slovenia was selected to host one of the EuroHPC petascale systems, which will be hosted under the name VEGA at IZUM, Institute of Information Science in Maribor, Slovenia, which already has the appropriate infrastructure to host such a system. At the end of April 2020, Vega – the procurement of the purchase, supply and maintenance – will be launched, with an estimated budget of €17.2 million.

This investment into the new HPC system will be accompanied with growing activities in training and research in Slovenia. Active participation in EuroHPC competence centres and in PRACE training activities, including hosting a new PRACE training centre and PRACE autumn schools, is the main driver that will result in a higher number of more competent young people coming to work with the new system. Additionally, many national and EU funded research projects in several domains, including engineering, personalised medicine, weather forecasting, climate change, discovering new materials and drugs will

benefit from the new supercomputer. All of this will, in the short term, result in a higher uptake of HPC technology by SMEs in Slovenia and in the wider region.

Spain

The aim of the MareNostrum5 Consortium, which coordinates the efforts of Spain, Portugal, and Turkey, together with the European Commission within the EuroHPC Joint Undertaking, is to procure and deploy a world-class HPC system for the benefit of society through the provision of computing services to scientists, industry, and public administrations. The new supercomputer, with an expected aggregated sustained performance above 150 petaflops, will integrate cutting-edge technologies, following the path set by its predecessors. This will be enhanced with a high-level user support service to promote and improve the usage of the system, and with the development of prototypes to facilitate the path to exascale supercomputers. In this way, MareNostrum5 will contribute to the objectives of EuroHPC JU and of the countries who are members of the consortium.

"The design of this new supercomputer will improve the services to current HPC users while facilitating the adoption of European technologies into HPC systems," says Mateo Valero, director of the Barcelona Supercomputing Centre.

"The international consortium is another demonstration of the strong European efforts to boost the capacities of public research infrastructures." says Sergi Girona, coordinator of the MareNostrum5 Consortium.

Portugal

A supercomputer called "Deucalion", co-financed by European funds as part of the EuroHPC Joint Undertaking, is going to be acquired and installed at the Minho Advanced Computing Centre in Portugal, with operations to begin by the end of 2020.

Deucalion will be making a fundamental contribution to the development of both European and Portuguese HPC capabilities across various disciplines and at different levels. It will support and foster research and innovation into the co-design of HPC systems, from their implementation all the way to their optimisation and exploitation. As well as this, it will build capacity in the management and operation of HPC systems.

Finally, it will foster the creation, expansion and consolidation of computational engineering and innovation teams and centres on European and national high-priority scientific domains requiring either large-scale digital simulations or very large-scale data management and processing.

Fundação para a Ciência e a Tecnologia (FCT), the Portuguese public agency that supports science, technology and innovation, is set to procure a heterogeneous 10 petaflop state-of-the-art system that will be based on the x86 and the ARMv8 architectures.

Deucalion will sport a subset of GPU accelerated nodes and another of bleeding-edge HPC-oriented RISC nodes for experimentation purposes. The new system will offer a high-speed shared storage of at least 10 petabytes and be fully connected by state-of-the-art interconnect and ethernet networks. The power consumption of the whole system is expected not to exceed 1MW and to be fully supplied by renewable energy sources. It will be possible to fully run and manage Deucalion with open source software. To serve the research community, the system will have to efficiently run a large and diverse set of applications, including those devoted to fluid dynamics, molecular dynamics, material science, high energy physics, as well as emerging scalable data-driven applications for analytics and machine learning.★

Q&A

Herbert Zeisel, chair of the EuroHPC Governing Board

As chair of the EuroHPC Governing Board, **Herbert Zeisel** has a key role to play in shaping the future of high-performance computing in Europe. We spoke to him to gather his thoughts on the development of EuroHPC, its relationship with PRACE, and how he thinks HPC will impact Europe in the future.

What have been the main developments in EuroHPC in 2019?

EuroHPC has three main pillars. The first is the R&D pillar where we develop, support and fund research in the area of developing the right software and architecture for high-performance computers. We fund the use of these systems, and also have a line called the European Processor Initiative where we are developing a new family of low-power European processors for extreme-scale computing, high-performance big data, and a range of emerging applications. In this area, we have already achieved quite a lot.

The second pillar was where we had calls out to get high-performance computing machines for different national states. We had two calls, one for petaflop machines, which we bought five of. We are now on the cusp of the phase where they are beginning to run. Then we had a call for three exascale machines.

The third pillar is about supporting the communities that surround HPC. Using high-performance computers is quite complicated! It's not just that you have to use a computer to run a programme on. You have to translate the physical or chemical problem that you are working on into a code that can be run on a computer to do a simulation. Supporting people in doing this is a very important issue, and we are helping to establish communities that do this. We already have good examples of this happening, for example with the work PRACE has done, and now we have a situation where we have all of the communities coming together and working towards common goals.

Herbert Zeisel

How do you see the relationship between EuroHPC and PRACE moving forward into the future?

We have had a number of discussions with PRACE about how we will act together in the future. PRACE represents an excellent example of how to run an organisation to give the right people access to computing and provide them with support. But right now, we have not got all of the instruments needed for EuroHPC.

For example, we have a huge problem in what we are doing with providing industrial access to supercomputers, especially SMEs (of course some of the larger companies have their own supercomputers). Smaller businesses need support in this area. This is something that is not that well established at the moment.

Today, access to supercomputers by PRACE is regulated by the PRACE Council. The question is, what will be the shape of the governance between the EuroHPC Council and the PRACE Council. Right now, we are discussing many different things to do with this governance, as well as the different access lines to HPC. In the end, we want to use PRACE for the common usage of computers, because we think it is necessary to have one point of contact for all the people rather than many. However, this is not so easy because the computers we have bought only have partial open access, and a national part which the nations have for themselves.

In general, we think that we have a good agreement with PRACE about how want to move forward with access. But with governance, we are yet to have proper discussions. These discussions will be happening in 2020.

If you look to the future, there will be a lot of new business models based on data. . . High-performance computing will play a huge role in our future in this respect.

How do you envisage the role of HPC in Europe in the future?

I think if you look at the development in Europe of data infrastructure and the whole IT infrastructure, including the cloud and HPC networks, they are hugely important for supporting what we need as a society moving forward in terms of academia and industry. If you look to the future, there will be a lot of new business models based on data. Artificial intelligence is one key aspect here, but there are others as well, and all of these will need the right infrastructure – networks, clouds, data storage, and computing power. High-performance computing will play a huge role in our future in this respect.

From another perspective, we need to look at futuristic applications such as the development of new materials. How we will develop materials in the future is hugely important, but at the moment we use a lot of trial and error. We try many different things out in order to get one result, but this is very time consuming and inefficient. If we are able to use computing power to simulate what we do with our atoms and molecules so that we can create materials on the computer, we can make this process much faster and more efficient.

We have to develop more infrastructure for Europe as we move on and combine it all under one roof. This is why I think EuroHPC will need to help develop things in this way. We need to provide industry and academia with the tools to solve their problems.

We will also need to help develop middleware that allows people centralised access to distributed data. This is a very important issue. We have to see all of these issues as part of one wider common goal for Europe.★

Our thanks to Herbert Zeisel for contributing to the PRACE Annual Report 2019

From the desk of

The Industrial Advisory Committee

Looking back over 2019, there have been two main highlights regarding industry adoption of Tier-O systems in Europe. Firstly, 10% of the resources offered via PRACE open calls are reserved for projects led by a company. This so-called "industry track for open science" passed a successful pilot stage and is now a permanent feature of future calls. The second highlight is that EuroHPC promises more: to reserve 20% of its resources for firms on a pay-to-use basis.

So, does this mean we should expect EuroHPC to have a better track record of technology transfer to industry than PRACE? The past can be a useful guide, particularly major historical periods such as the Renaissance, the Age of Enlightenment and the Industrial Revolution. In digital form, let's consider the supercomputing renaissance to be the period before PRACE, when scientific artisans hand-crafted scalable applications so that they could run at scale. If the supercomputing enlightenment was the bold vision that led to the creation of PRACE, in order to share knowledge between European centres, then EuroHPC marks the start of the digital industrial revolution. Whilst there is political will, through investment in supercomputing, EuroHPC is not enough. Other aspects of the landscape need to change for the revolution to happen.

My ancestor George Margetts' experience in the 18th century Enlightenment illustrates how technology transfer used to proceed. George revolutionised shipping both through the manufacture of precision instrumentation and numerical analysis. He designed a clock that kept time for a record eight days and published a book of astronomical and parallax corrections for calculating longitude. His scientific work was paid for by various subscribers including the East India Company (the world's largest shipping company) and the Royal Navy. In our age, we rarely find this direct link between technology provider (scientist) and technology consumer (industry). Instead we rely on third party agents who invest on behalf of society, through funding bodies such as the European Commission. Direct engagement with industry is rare, but still exists. For example, the NCSA at Illinois has an industrial partner programme where firms pay a subscription to access technology and experts.

Dr Lee Margetts, chair of the PRACE Industrial Advisory Committee

James Dyson has set up his own college to train the digital engineers his firm needs to employ. The warning: Europe must modernise its schools, colleges and universities to teach digital science, data analytics and artificial intelligence.

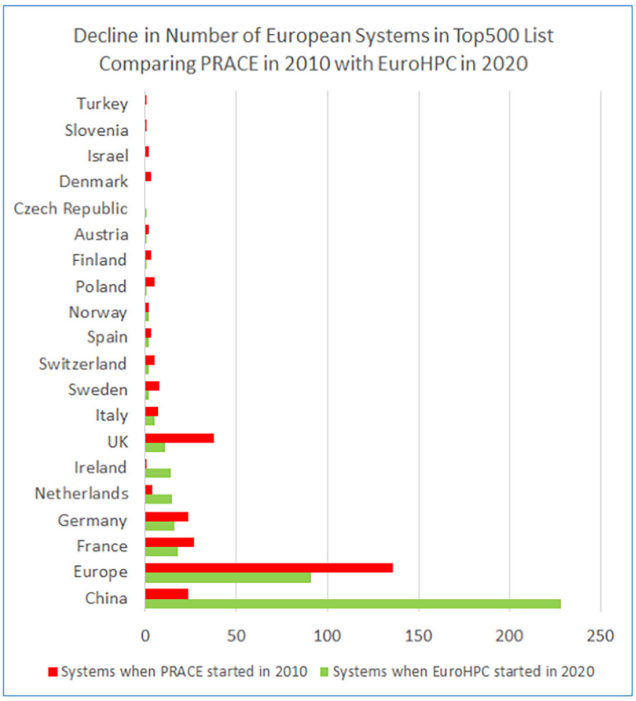

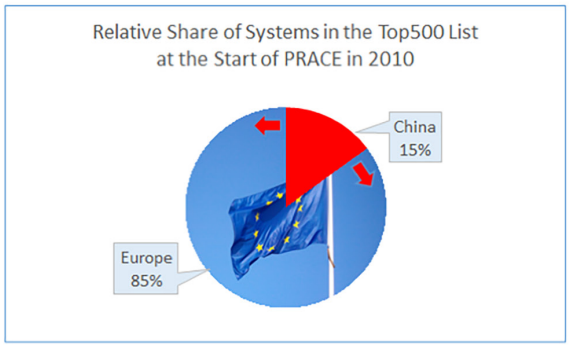

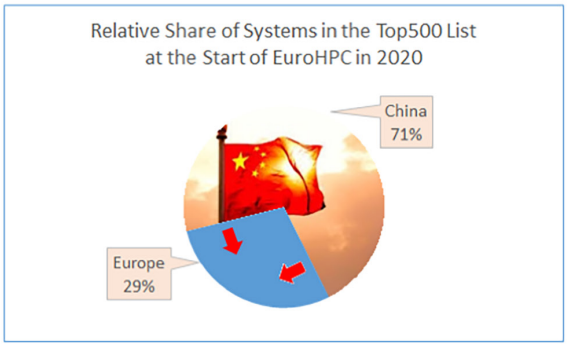

An illustration of the decline in systems in Europe compared to a rise in China

Another distinction between the analogue industrial revolution and the present concerns the provision of training. PRACE training is excellent and suits the age of supercomputing enlightenment, but it is not sufficiently scalable to support a digital industrial revolution. In the 19th century, there was an explosion of new technical colleges, set up to train the entire working population in the ideas that had emerged during the Age of Enlightenment. Donald Trump has highlighted the present-day skills crisis. Whilst the USA has millions of open positions in the digital economy, the available workforce lacks the required skills. The tension is even leading industry to reinvent education. James Dyson has set up his own college to train the digital engineers his firm needs to employ. The warning: Europe must modernise its schools, colleges and universities to teach digital science, data analytics and artificial intelligence.

I have the impression that some politicians believe that signing up to EuroHPC will transform industry in their home country. EuroHPC is of critical importance to Europe, in terms of supporting world leading use of digital technologies. Members of the PRACE Industrial Advisory Committee are enthusiastic supporters of PRACE, EuroHPC, and any other initiative aimed at digitising industry. However, a technology-driven industrial revolution relies on the technology being ubiquitous (the SMEs in the bazaar), not centralised (the EuroHPC cathedral). Data from the Top500 list of supercomputers shows that Europe is seriously underinvesting in supercomputers when compared with China (see figure which includes a subset of data from China and Europe). Investing in EuroHPC is essential, but not enough. Each European Member State must make additional local investments in new supercomputing centres, transform education and help science and industry work together for the benefit of their local industries.

If there is no additional investment in local supercomputing by European Member States, no effective engagement with industry, and no supply of digitally literate workers, I regretfully predict that there will be no digital industrial revolution in Europe. ★

‘No industry, no digital revolution’

PRACE

Key Performance Indicators

Given the evolution of computational power in the PRACE portfolio, PRACE-related statistics are becoming increasingly important to highlight the impact of PRACE on HPC-based research, HPC know-how in Europe, and European industrial uptake of HPC.

PRACE's impact on evolving research

Offer and demand of resources

Figure 1 shows the evolution of PRACE resources offered and requested in Project Access Calls. PRACE first provided HPC services in 2010 with contributions from German Tier-0 systems. France, Italy and Spain added their contributions gradually. This is reflected in the constant increase of HPC resources offered by PRACE to the scientific community until the 6th call, where a stable regime was reached. The phasing out of the initial phase of PRACE (known as PRACE 1) started in the 10th call, while the second phase of PRACE (known as PRACE 2) only started in call 14. This is reflected in the valley from the 10th call to the 13th call. With PRACE 2, a substantial increase in the amount of available resources can be seen, thanks to the renewed contributions of all the original PRACE hosting members and Switzerland as a new hosting member. The stable regime has been reached quickly, and since the 14th call an average of 1.9 billion core hours per call has been offered. The 19th call shows a high increase from this average, due to the inclusion of novel, more dense architectures into PRACE standard offer.

The demand for HPC resources has always exceeded the capacity of PRACE to provide them. The average oversubscription of PRACE calls is 3:1, reaching a 5:1 ratio during the phasing out of PRACE1. This constant interest in PRACE resources has been driving the periodic upgrades and additions of new systems that PRACE offers, specifically concerning the PRACE 2 programme. ★

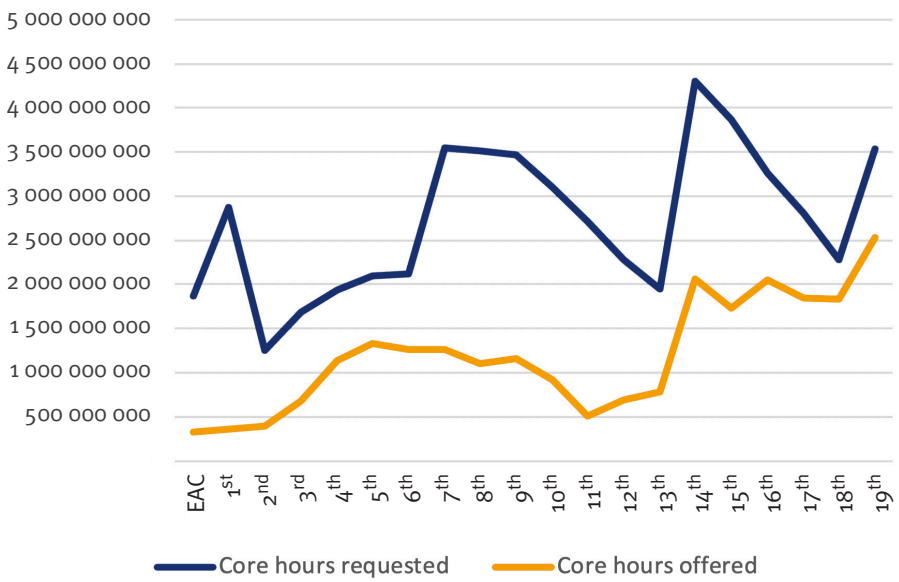

Figure 1

Core hours offered and requested in each PRACE call

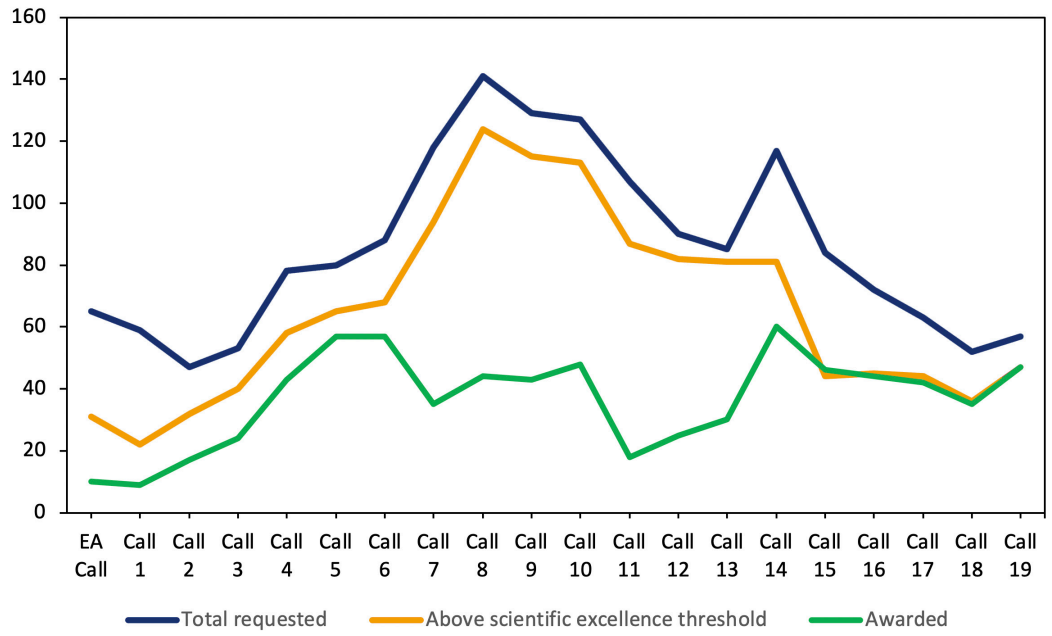

Figure 2

Total number of applications received, proposals above scientific excellence threshold, and projects awarded

Number of projects

During the initial phase of PRACE (known as PRACE 1), the number of project applications received via PRACE Calls for Proposals for Project Access exhibited a clear overall upward trend. The phasing-in of PRACE 1 naturally incited an increase in demand for Tier-0 resources. This is particularly evident up to the 8th call, with a large sustained increase between the 6th and 8th call, followed by a slight decrease (Figure 2).

A downward trend of rejected projects below the scientific excellence threshold is noted, displaying maturity of submitted proposals, in which researchers put more effort into the quality of their proposals, as a reaction to increased competition. Moreover, the evolution reflects the positive outcomes of PRACE Preparatory Access Calls (including access type C) that enable prior technical support for application and scalability tests. Figure 2 also highlights an increase in rejected projects above the scientific threshold, particularly after the 6th call. This is correlated with the increase in total applications.

During the phasing-out of PRACE 1, the number of available core hours dropped (Figure 1), and this decreased the demand, as researchers anticipated even stronger competition for the remaining resources. This trend was mitigated in the 12th and 13th call, when PRACE Hosting Members made additional core hours available during the preparation of the PRACE 2 programme, which started in the 14th call.

With the start of PRACE 2, there has been an increase in the number of projects awarded, which combined with the increase of resources that started in the 14th call (Figure 1) shows the clear success of the second phase of PRACE. In this second phase, the scientific objectives of PRACE have been updated to increase of the scope and excellence of the projects awarded. The minimum size of allocations has been increased three-fold, and the scientific threshold as well. The decrease in proposals submitted and the apparent decrease in their quality in Call 15 is a positive sign of the success of this scientific update.

Despite the competition, demand for PRACE resources remains high and all PRACE calls are oversubscribed (Figure 1), indicating that scientists consider Tier-0 access an essential asset to their work. This was also underlined by the PRACE Scientific Steering Committee (SSC) during the preparation of PRACE 2. ★

Recurring users

PRACE also keeps track of the submission of Project Access proposals by recurrent principal investigators (PIs) (Figure 3). This KPI is created by checking for each call if a PI is new to PRACE.

The ratio of first-time applicants is relatively high – roughly 50% of PIs who submitted to the two project access calls in 2019 were recurrent applicants to a PRACE Call for Proposals for Project Access. This means that half of all project proposals are submitted by new users. This indicates that PRACE is continuously attracting new PIs, while remaining an essential support for existing users. The upward trend of the ratio of recurrence is visible, particularly from the 6th call onwards, influenced by the downward trend on awarded projects, but recovered with the onset of PRACE 2.★

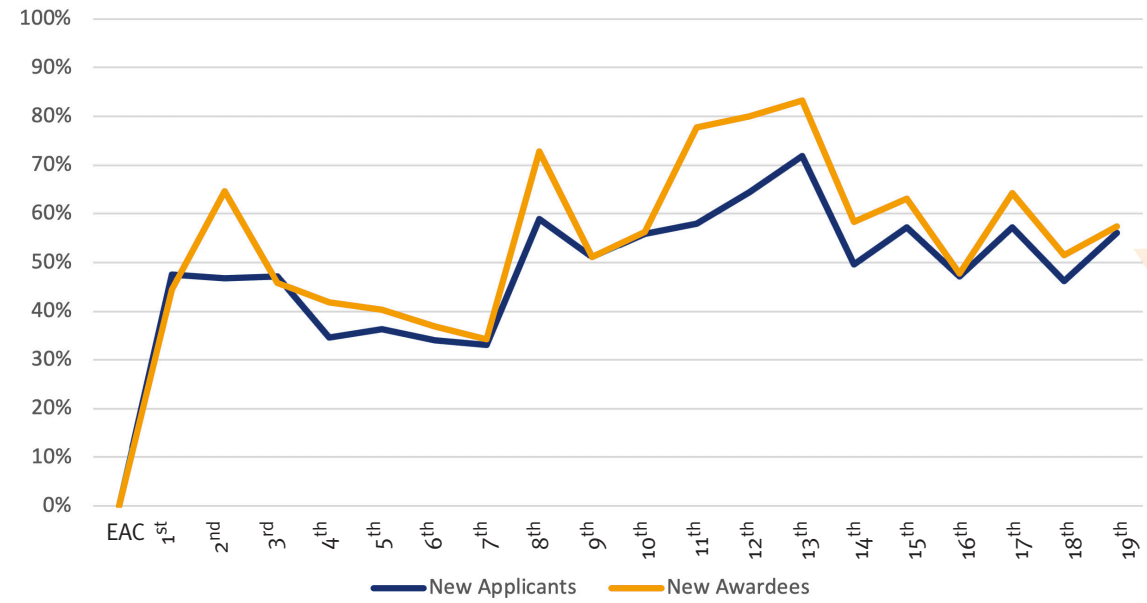

Figure 3

Ratio of new applicants and new awardees in each PRACE call

International transnational co-operation

Two-thirds (63%) of resources awarded under the early access call through to the 15th call are awarded to "foreign projects". Foreign projects are defined as projects with principal investigators (PIs) from a different country (recorded as the country of the PI's primary institution) than the machine on which the research is executed.

The ratio of awarded foreign projects has remained rather stable over time (Figure 4). This shows that the nationality of the PI's institution does not influence the chances of a project being awarded. It also demonstrates PRACE's impact in the enhancement of European and International collaboration.★

Figure 4

Ratios of awarded 'foreign' projects and resources for awarded 'foreign' projects

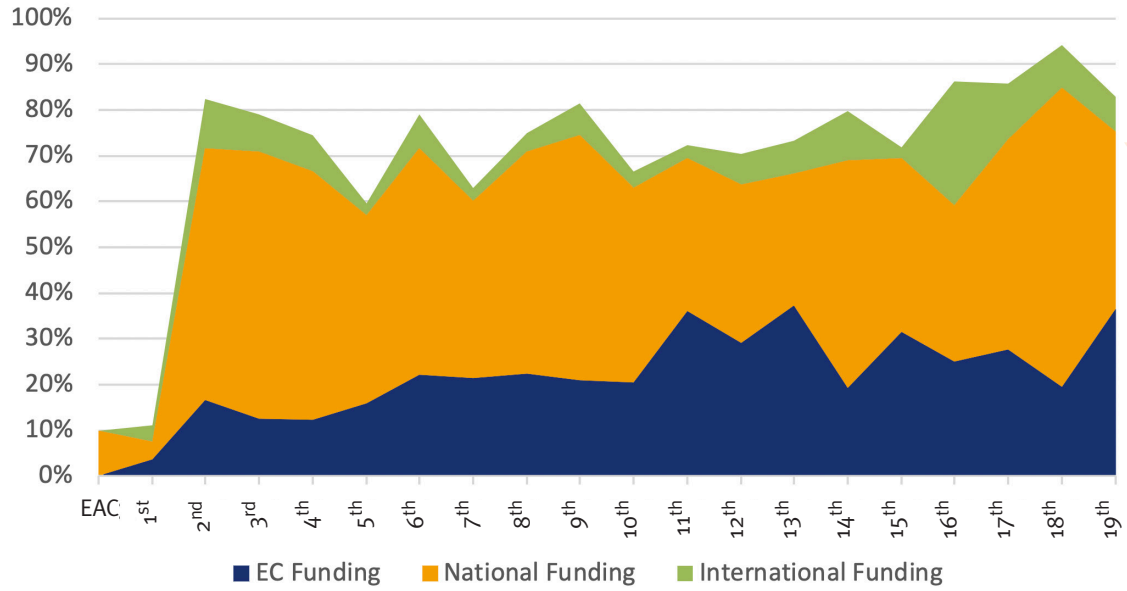

Figure 5

Ratios of awarded projects with EC, national, and international support

Co-funding

PRACE awards are normally developed within larger scientific initiatives, where HPC resources are part of the needs of the project. PRACE asks Project Access awardees to declare such synergies.

On average, 75% of PRACE users have declared that their awards are complemented with EC, national, or international funds (Figure 5).

The major fraction corresponds to national projects, which is slowly showing a downward trend since the 10th Call. EC funding shows an increasing trend, coinciding with the implementation of the H2020 programme. International funding remains low, with 8% being the average contribution.

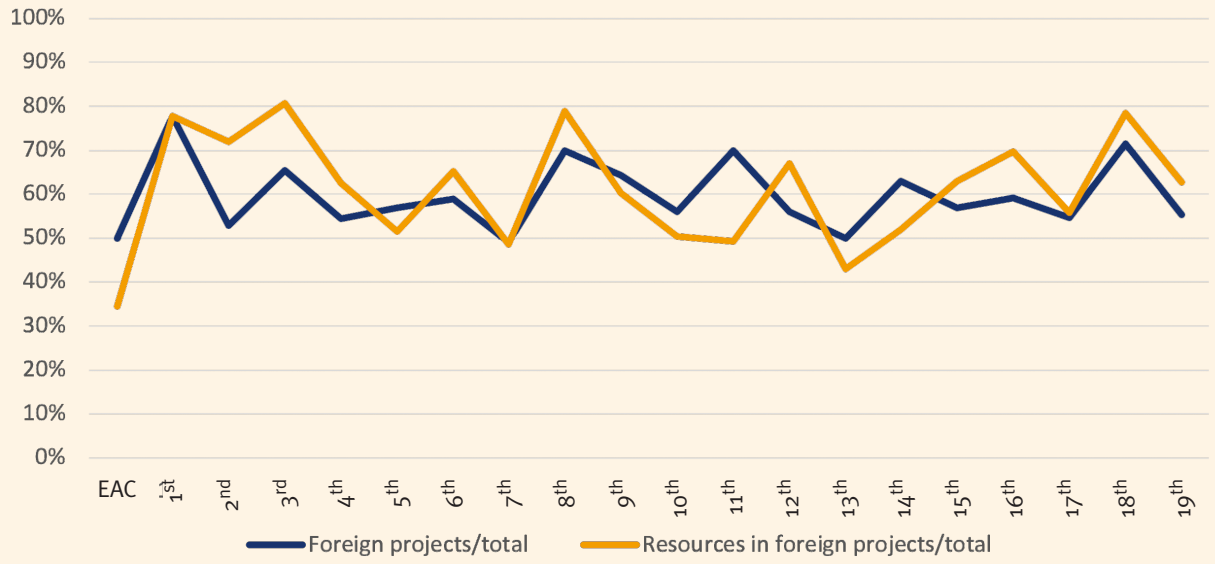

PRACE's impact on growing know-how in Europe

Since 2008, PRACE has been engaged in providing top-class education and training for computational scientists in Europe through the PRACE Training Centres (PTCs), the International HPC Summer School, and PRACE Seasonal Schools, with a clear increase of participants registered (Figure 6).

Six PTCs were first established, and these are located at:

- Barcelona Supercomputing Centre (Spain)
- CINECA – Consortio Interuniversitario (Italy)
- CSC – IT Center for Science Ltd. (Finland)
- EPCC at the University of Edinburgh (UK)
- Gauss Centre for Supercomputing (Germany)
- Maison de la Simulation (France)

After the rapid increase between 2010 and 2012, a plateau is evident since 2012. As this indicated that the maximum capacity of PRACE training offerings had been reached, four new PRACE Training Centers were opened in 2017.

The new training centres are located at:

- IT4Innovations (Czech Republic)
- GRNET (Greece)
- ICHEC (Ireland)
- SURFSara (The Netherlands)

PTC training events, seasonal schools and the International HPC Summer School are offered free-of-charge to eligible participants.★

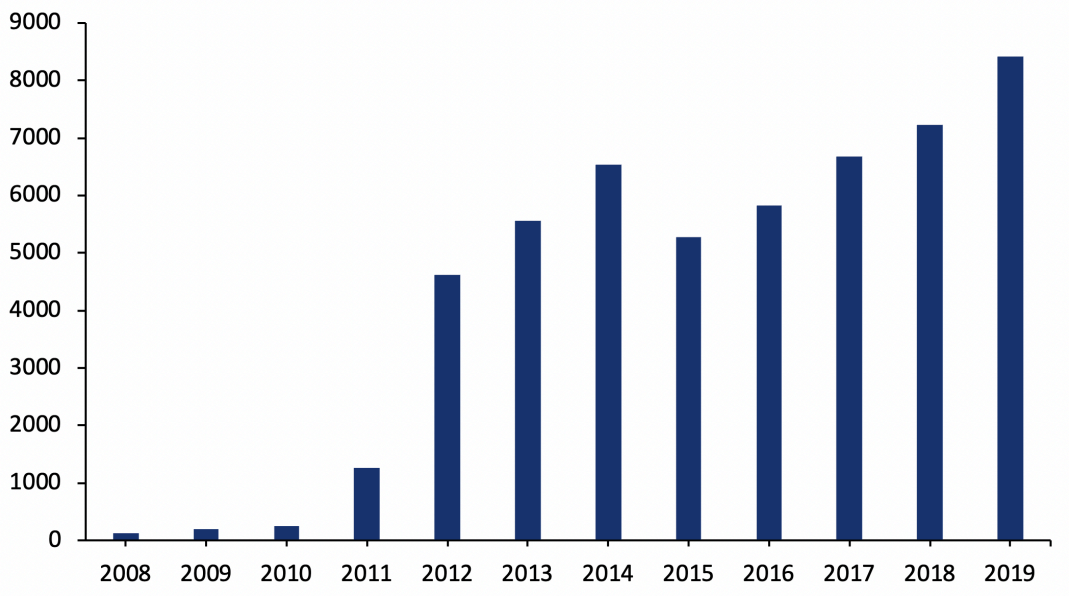

Figure 6

Number of person-days registered at PRACE training days between 2008 and 2019

Between August 2008 and December 2019, PRACE provided close to 52 000 participant-days of training through attendance-based courses, with an upward attendance trend. PRACE courses were attended by over 16 400 individuals. This shows the effectiveness of PRACE in attracting, training, and retaining competences.

In 2019 the number of participants attending PTCs courses was 2 618 (359 with non-academia affiliation). 86% of participants attending PTCs trainings days have an academic affiliation, illustrating the impact of such event on research and scientific communities, in particular for early stage researchers and PhD students.

A clear difference of attendance is observed between the first and second semester of 2019. As observed in Figure 7, the total number of attendees registered in the first semester (Q1 and Q2) is significantly higher than during the second semester (Q3 and Q4). This indicates that the bulk of the training offered occurs in the first semester, with a notable drop in attendance during Q3 which corresponds with the summer and winter vacation periods. ★

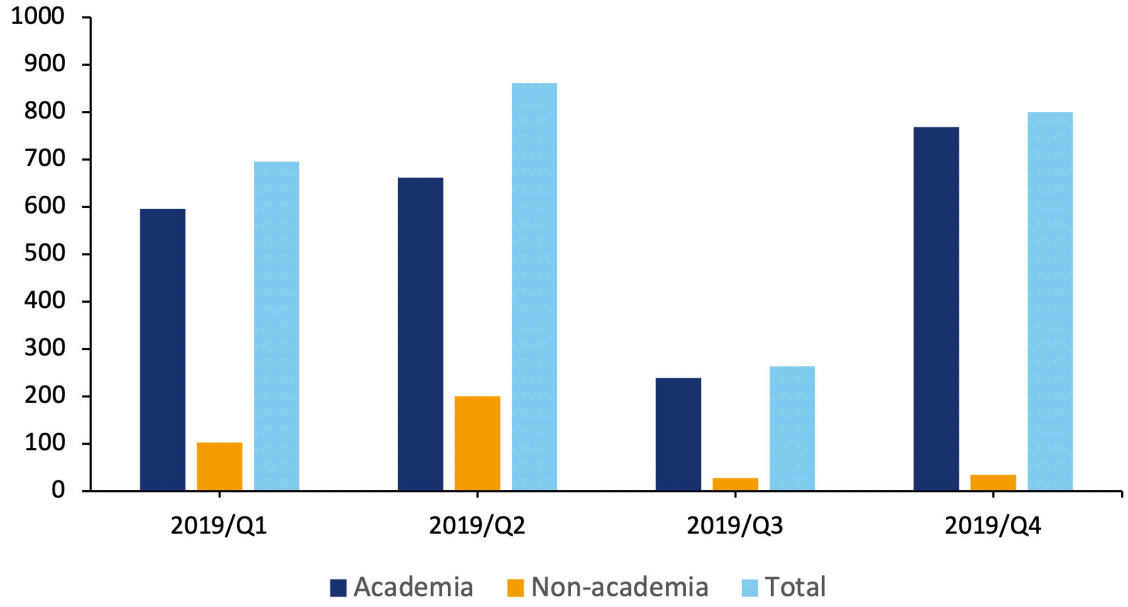

Figure 7

Number of person-days registered at PRACE training days in 2019

PRACE's impact on attracting the industrial sector

Industrial visitors of the PRACE booth at ISC and SC

The interest of industry in PRACE at high-level international events has remained high over the past years (Figure 8). The total number of industrial participants showing interest in PRACE during the two main HPC events (Supercomputing (SC) in the USA, and the International Supercomputing Conference (ISC) in Germany) between 2008 and 2019 was 2 058 individuals.

More than half of the companies that visited the PRACE booth at ISC19 and SC19 were first-time visitors. This indicates that industrial interest in PRACE is growing on both sides of the Atlantic. ★

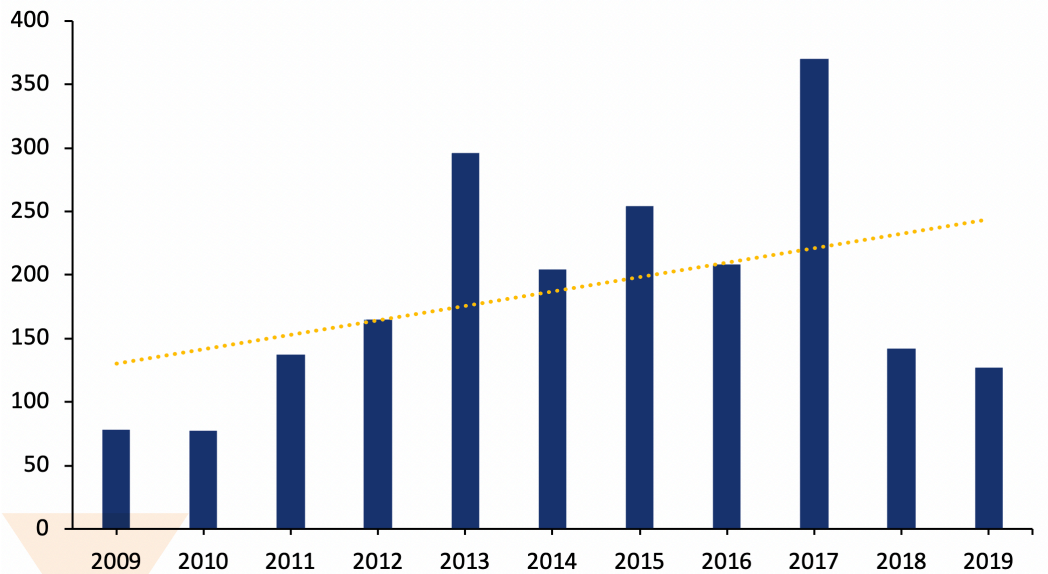

Figure 8

Number of industrial attendees that made contact with the PRACE booth at ISC and SC; and related trend line

Industrial participants in PTCs

The average participation of industry in PTC trainings is 15.42% between 2012 and 2019 (12.9% in 2017). The increasing interest from industry in participating in HPC training is visible in Figure 9. More than 359 industrial participants were trained by PRACE in 2017.

Eligible industrial participants enjoy the same service as academic trainees and can attend PTC courses free of charge.★

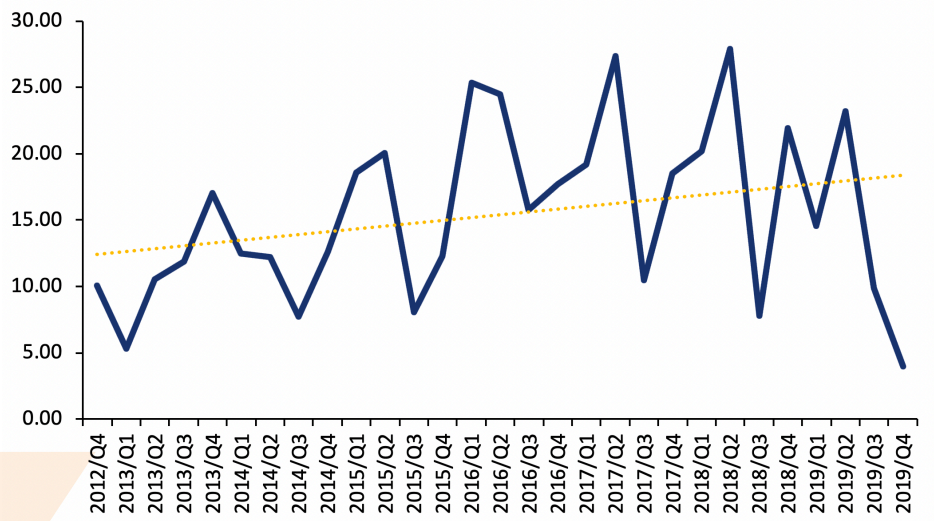

Figure 9 *Industrial participation in PTC training days, and related trend line*

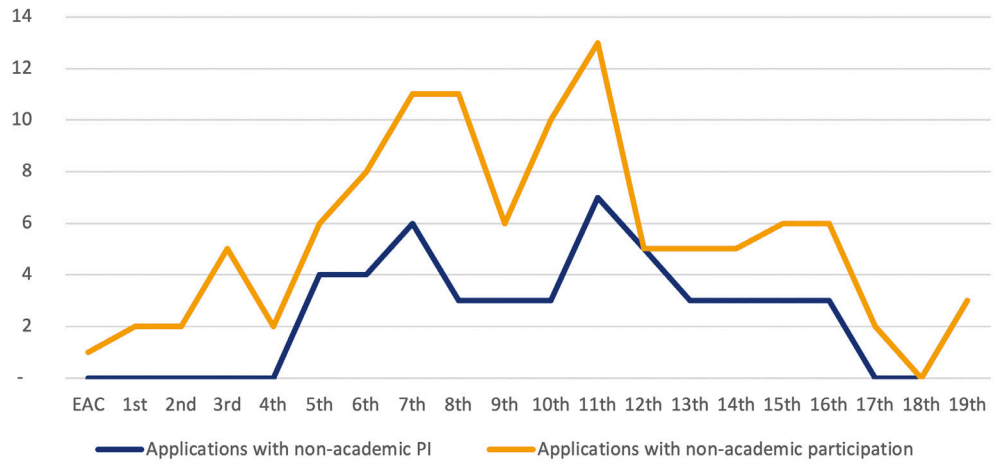

Figure 10 *Industry participation in PRACE*

Non-academic use of PRACE HPC resources

PRACE opened its calls for proposals to non-academic applications mid-2012. This can take the form of a project led by a principal investigator coming from a private company, or a researcher from industry collaborating in an academia-led project. The number of applications with non-academic participation can be seen in Figure 10. Up to call 18, applications with a non-academic PI have competed for PRACE resources with applications with an academic PI. This was changed in call 19 where an "Industry Access Pilot" which offered principal investigators from industry the possibility to apply for single-year access to a special industry track which prioritised 10% of the total resources available. This will continue in future calls and we expect an increasing non-academic use of PRACE HPC resources in future calls. Through the SHAPE programme, started in 2013, PRACE has received 73 applications so far, of which a total of 55 have been awarded both PRACE HPC resources and, more importantly, support and know-how in the PRACE centres.★

PRACE aisbl structure

This article lists the members of each body of PRACE aisbl. Our thanks goes out to all those who have supported our organisation during 2019. More information on the bodies of PRACE aisbl can be found on our website:

prace-ri.eu/about/organisation/

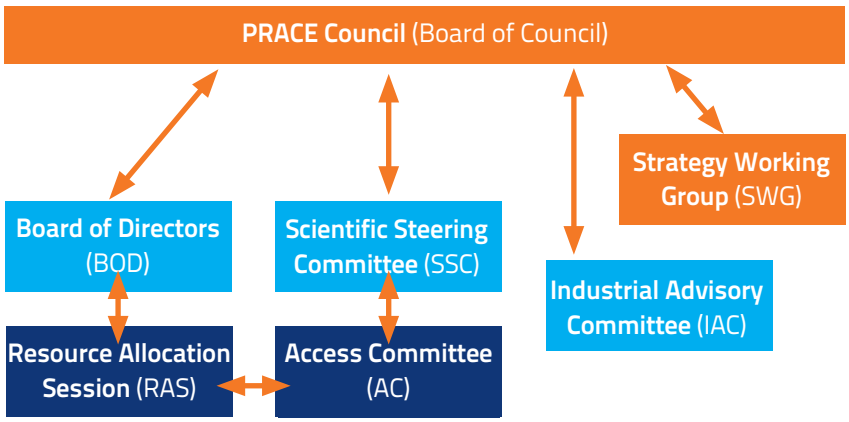

PRACE Member Organisations and delegates to the PRACE Council

The Council is the deliberative body of PRACE aisbl and decides on all matters of the association. It is composed of one representative from each member. The Board of the council is elected by the Council from among the delegates representing the Members.

Board of the PRACE Council			
Thomas Lippert	Council Chair	Forschungszentrum Jülich	Germany
Janne Ignatius	Council Vice-Chair	CSC-Tieteen tietotekniikan keskus Oy (IT Center for Science Ltd.)	Finland
Florian Berberich	Council Secretary	Forschungszentrum Jülich	Germany
Serge Bogaerts	Managing Director	PRACE aisbl	Belgium
Observers			
Zlatan Car	Observer	University of Rijeka	Croatia
Zigmunds Zitmanis	Observer	Riga Technical University	Latvia
Delegates			
Christoph Dellago	Delegate	ACOnet Association	Austria
Philippe Geuzaine	Delegate	Direction générale opérationnelle de l'Économie, de l'Emploi et de la Recherche – Service Public de Wallonie (DGO6-SPW)	Belgium
Ivan Dimov	Delegate	Executive Agency "Electronic Communication Networks and Information Systems"	Bulgaria

Constantia Alexandrou	Delegate	The Cyprus Institute	Cyprus
Vit Vondrák	Delegate	VŠB-Technical University of Ostrava	Czech Republic
René Belsø	Delegate	Danish e-Infrastructure Cooperation (DeIC)	Denmark
Philippe Lavocat	Delegate	Grand Equipement National De Calcul Intensif (GENCI)	France
Dieter Kranzlmueller	Delegate	GAUSS Centre for Supercomputing (GCS) e.V.	Germany
Ognjen Prnjat	Delegate	Greek Research And Technology Network Sa (GRNET	Greece
Tamas Maray	Delegate	KIFÜ Governmental Information Technology Development Agency	Hungary
Jean-Christophe Desplat	Delegate	National University of Ireland - Irish Centre for High-End Computing (ICHEC)	Ireland
Haim Taitelbaum	Delegate	Inter University Computation Center - IUCC	Israel
Gabriella Scipione	Delegate	CINECA Consorzio Interuniversitario	Italy
Pascal Bouvry	Delegate	Université du Luxembourg	Luxembourg
Peter Michielse	Delegate	SURFsara	The Netherlands
Gunnar Bøe	Delegate	UNINETT Sigma2 AS	Norway
Norbert Meyer	Delegate	Instytut Chemii Bioorganiczej Pan - Institute of Bioorganic Chemistry	Poland
Pedro Alberto	Delegate	University of Coimbra	Portugal
Jozef Noga	Delegate	Computing Centre of the Slovak Academy of Sciences	Slovakia
Marko Bonač	Delegate	ARNES	Slovenia
Sergi Girona	Delegate	Barcelona Supercomputing Center (BSC) - Centro Nacional de Supercomputacion (CNS)	Spain
Hanifeh Khayyeri	Delegate	Swedish Research Council	Sweden
Thomas Schulthess	Delegate	ETH Zurich/CSCS	Switzerland
Ertugrul Karacuha	Delegate	National Center for High Performance Computing of Turkey - Ulusal Yüksek Başarımlı Hesaplama Merkezi (UYBHM)	Turkey
Richard Gunn	Delegate	The Engineering and Physical Sciences Research Council (EPSRC)	United Kingdom

PRACE Strategy Working Group (SWG)

The Strategy Working Group (SWG) is a sub-committee of the PRACE Council. It counts 14 seats, including one for each hosting member, and three for representatives of the general partners.

Thomas Lippert	PRACE Council Chair
Janne Ignatius	PRACE Council Vice-Chair and delegate of Finland
Florian Berberich	PRACE Council Secretary
Gabriella Scipione	Delegate of Italy
Dieter Kranzlmueller	Delegate of Germany
Philippe Lavocat	Delegate of France
Sergi Girona	Delegate of Spain
Thomas Schulthess	Delegate of Switzerland
Richard Gunn	Delegate of United Kingdom
Peter Michielse	Delegate of the Netherlands
Rene Belsø	Delegate of Denmark
Serge Bogaerts	Managing Director
Núria López	Chair of the Scientific Steering Committee (SSC)
Lee Margetts	Chair of Industrial Advisory Committee (IAC)

PRACE Board of Directors (BoD)

The Board of Directors (BoD) is the executive body of the association and is generally responsible for managing and representing the association.

Florian Berberich	Member for Germany
Serge Bogaerts	Chair and Managing Director
Maria Grazia Giuffreda	Member for Switzerland
Oriol Pineda	Member for Spain
Philippe Segers	Member for France
Núria López	Member as Chair of the Scientific Steering Committee (SSC)
Debora Testi	Member for Italy

Scientific Steering Committee (SSC)

The Scientific Steering Committee (SSC) is composed of leading researchers from Europe. They provide advice and guidance on all matters of a scientific and technical nature that may influence the scientific work carried out using the association's resources.

Marina Bécoulet		Plasma physics, fusion	France
Luke Drury		Plasma physics, fusion	Ireland
Claudia Filippi		Electronic structure, multiscale modelling	The Netherlands
Laura Grigori		Numerical mathematics, HPC	Germany
Christian Holm		Statistical physics	Germany
Heiner Igel		Seismology	Germany
Erik Lindahl		Life sciences	Sweden
Núria López	Chair	Computational chemistry	Spain
Aimee Morgans		Mechanical engineering	United Kingdom
Ignacio Pagonabarraga		Computational physics	Spain
Mike Payne		Computational physics	United Kingdom
Silvia Picozzi		Materials	Italy
Matej Praprotnik		Chemistry, multiscal modelling	Slovenia
Sinéad Ryan	Vice-Chair	Particle physics, mathematics	Ireland

Financial Oversight & Risk Assessment Committee (FORAC)

FORAC provides a high-level and visible forum for monitoring standards of internal control and propriety, economy, effectiveness. It also evaluates the extent to which systems and procedures help PRACE objectives to be met.

Mark Parsons	Chair	United Kingdom
Edouard Brunel		France
Claus-Axel Müller		Germany

Industrial Advisory Committee (IAC)

The Industrial Advisory Committee (IAC) is composed of European industry representatives (both from multi-nationals and SMEs) representing 11 industrial sectors.

Enric Gibert	Vice-Chair	Life sciences, pharmaceuticals	Spain
Tomi Ilijaš		Engineering, manufacturing, SMEs	Slovenia
Lee Margetts	Chair	ISV	United Kingdom
Alain Martin		Energy	France
Marc Morere		Aeronautics, aerospace	France
Martin Winter		Materials, chemistry	France
Observers (with no voting rights)			
Maike Gilliot		HPC vendors (ETP4HPC)	France

User Forum (UF)

The PRACE User Forum was set up in December 2011 through an initiative of PRACE itself. It is not a body of the association, but an independent entity where PRACE users can discuss their experiences and express their future needs as well as provide feedback on the current services and resources of the PRACE HPC Research Infrastructure.

Marc Baaden		CNRS	France
Carmen Domene		King's College London	United Kingdom
Turlough Downes		Dublin City University	Ireland
Stefano Fabris		CNR-IOM DEMOCRITOS	Italy
Derek Groen	Vice-Chair	Brunell University London	United Kingdom
Troels Haugbølle	Chair	University of Copenhagen	Denmark
Koen Hillewaert		Centre de recherche en Aéronautique (Cenaero)	Belgium
William Sellers		University of Manchester	United Kingdom
Gabriel Staffelbach		Cerfacs	France
Maria-Teresa Parra		University of Valladolid	Spain
Jorge Vieira		Instituto Superior Técnico	Spain
Gustavo Yepes		Universidad Autonoma de Madrid (UAM)	Spain

Access Committee (AC)
The Access Committee is responsible for the scientific assessment of PRACE proposals. They produce a final ranked list of proposals and suggested allocations. The AC is composed of researchers experienced in areas of science, engineering and supercomputing. Members are proposed by the SSC and approved by the Council.

Takayuki Aoki		Fluid dynamics, engineering	Japan
Edouard Audit		Universe	France
Marc Baaden	Vice-Chair	Biology, chemistry	France
Peter Bauer		Earth system sciences	United Kingdom
George Biros		Computational Engineering	United States
Hans-Joachim Bungartz		Mathematics, computer science	Germany
Giovanni Ciccotti		Structure of matter, biological systems, molecular dynamics	Italy
Georges-Henri Cottet		Turbulence, engineering	France
Christine Davies		High-energy physics	United Kingdom
Luigi Del Debbio	Chair	Particle physics	United Kingdom
Stefan Goedecker		Chemical sciences, materials	Switzerland
Simone Hochgreb		Engineering	United Kingdom
Maria Paola Lombardo		High Energy Physics	Italy
Fernando Martin Garcia		Chemistry, materials	Spain
Victor Milman		Materials (with industrial applications)	United Kingdom
Nadia Pinardi		Atmospheric physics, oceanography	Italy
Luciano Rezzolla		Astrophysics	Germany
Friederike Schmid		Chemistry, polymers	Germany
Spencer Sherwin		Computational Engineering	United Kingdom
Alexandre Tkatchenko		Chemistry, materials	Luxembourg
Dietrich Wolf		Computational and statistical physics	Germany

PRACE Members

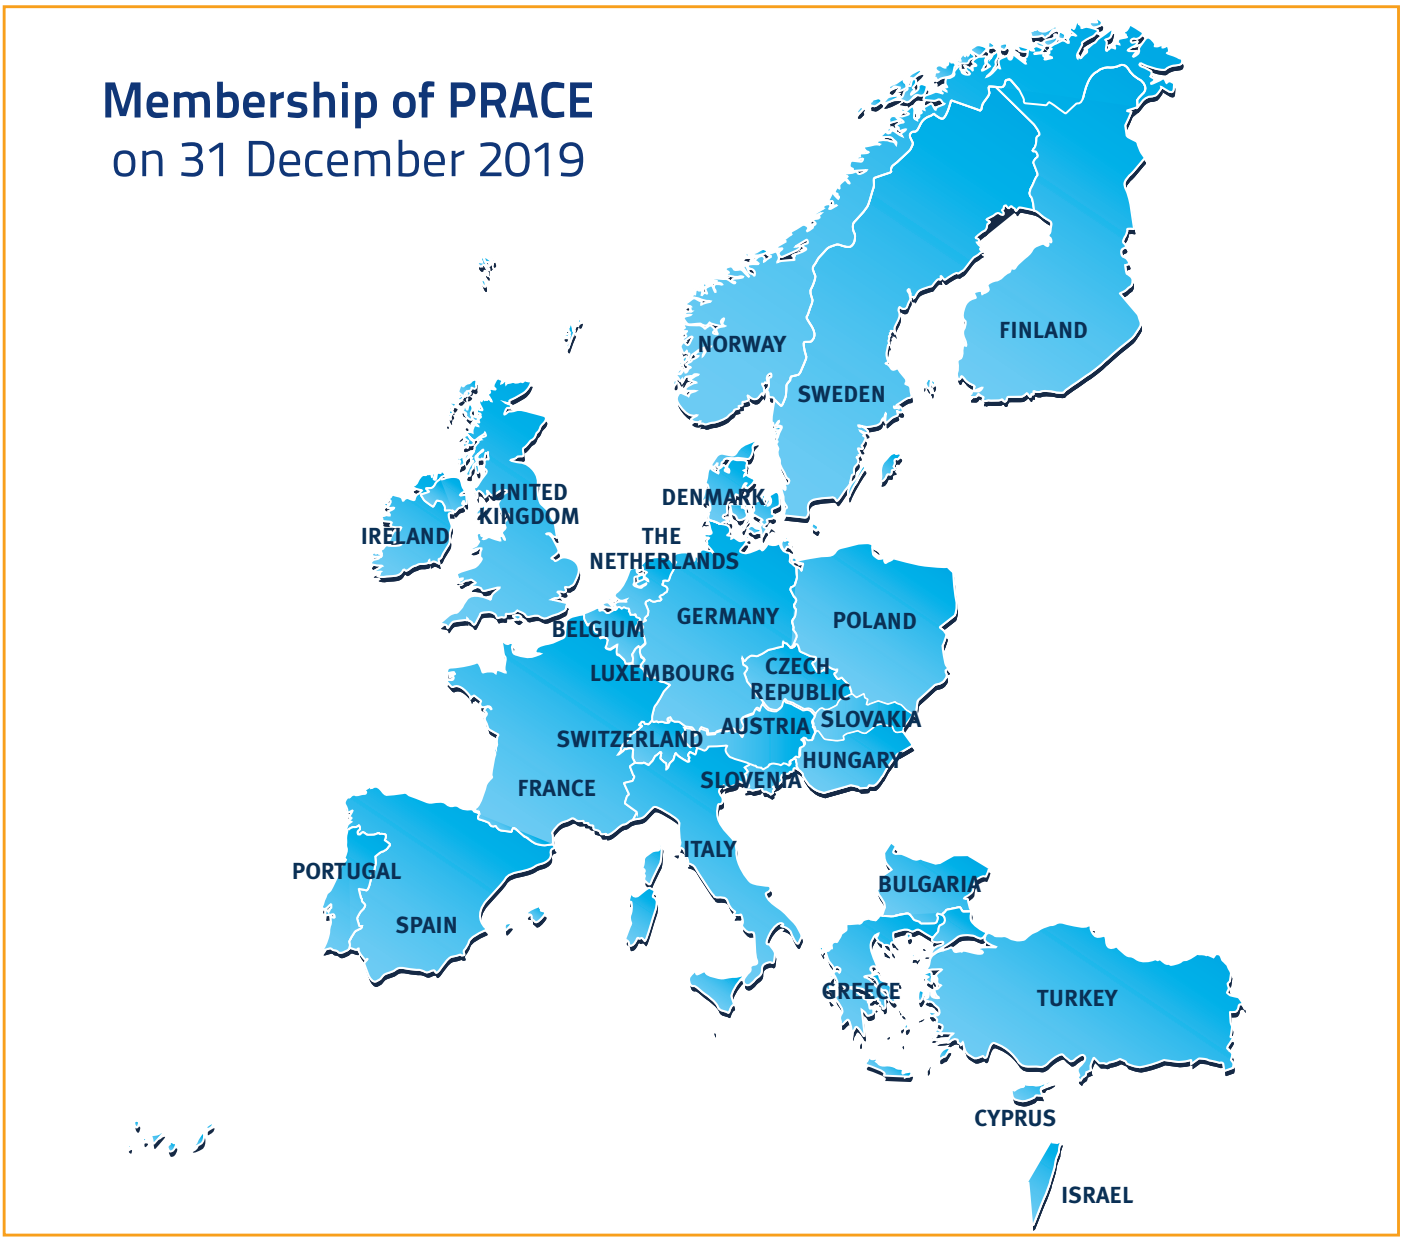

AUSTRIA	ACONET Association – Austrian Academic Computer Network www.aco.net
BELGIUM	Direction générale opérationnelle de l'Économie, de l'Emploi et de la Recherche – Service Public de Wallonie (DGO6-SPW) https://recherche-technologie.wallonie.be

Members & Partners

BULGARIA	Ministery of Transport, Information Technology and Communications **www.mtitc.government.bg**
CYPRUS	CaSToRC – Computation-based Science and Technology Research Center, The Cyprus Institute **www.cyi.ac.cy**
CZECH REPUBLIC	IT4I, VŠB – Technical University of Ostrava **https://it4i.cz/lang=eng; https://vsb.cz/en**
DENMARK	DeIC – Danish e-Infrastructure Cooperation **www.deic.dk**
FINLAND	CSC – IT Center for Science Ltd. **www.csc.fi**
FRANCE	GENCI – Grand Equipement National de Calcul Intensif **http://genci.fr/en**
GERMANY	GCS – GAUSS Centre for Supercomputing e.V **www.gauss-centre.eu**
GREECE	GRNET – Greek Research and Technology Network S.A. **https://grnet.gr/en/company**
HUNGARY	KIFÜ – Kormányzati Informatikai Fejlesztési Ügynökség **http://kifu.gov.hu/**
IRELAND	ICHEC – Irish Centre for High-End Computing **www.ichec.ie**
ISRAEL	IUCC – Inter-University Computation Center **https://iucc.ac.il/en**
ITALY	CINECA – Consorzio Interuniversitario **https://.cineca.it/en**

LUXEMBOURG	University of Luxembourg **https://wwwen.uni.lu**
THE NETHERLANDS	SURFsara **https://www.surf.nl/en**
NORWAY	SIGMA – UNINETT Sigma AS – The Norwegian Metacenter for Computational Science **www.sigma2.no**
POLAND	PSNC – Instytut Chemii Bioorganicznej Pan – Institute of Bioorganic Chemistry – Poznan Supercomputing and Networking Center **http://man.poznan.pl/online/en**
PORTUGAL	Universidade de Coimbra **http://uc.pt/en**
SLOVAKIA	Computing Center of the Slovak Academy of Science **https://vs.sav.sk/?lang=en**
SLOVENIA	ARNES **http://arnes.splet.arnes.si/en**
SPAIN	BSC – Barcelona Supercomputing Center – Centro Nacional de Supercomputación **www.bsc.es**
SWEDEN	Vetenskapsrådet – Swedish Research Council **https://www.vr.se.english.html**
SWITZERLAND	CSCS, ETH – Eidgenössische Technische Hochschule Zürich – Swiss Federal Institute of Technology, Zürich **https://ethz.ch/eng.html**
TURKEY	Istanbul Technical University **https://global.itu.edu.tr/**
UNITED KINGDOM	EPSRC – The Engineering and Physical Sciences Research Council **https://epsrc.ukri.org**

PRACE Systems

This article lists the systems that PRACE provides access to. The systems are listed in alphabetical order, and the information reflects their status on 31 December 2019. Updates and upgrades of these systems are published here: **www.prace-ri.eu/prace-resources/**

Hazel Hen, GCS@ HLRS, Germany

Hazel Hen is the Cray XC40 system (upgrade of Hornet system) and is designed for sustained application performance and highly scalable applications. It delivers a peak performance of 7.42 petaflops. This new system is composed of 7 712 compute nodes with a total of 185 088 Intel Haswell E5-2680 v3 compute cores. Hazel Hen features 965 terabytes of main memory and a total of 11 petabytes of storage capacity spread over 32 additional cabinets containing more than 8 300 disk drives. The input-/output rates are +/- 350 gigabytes per second.★

JOLIOT Curie, GENCI@CEA, France

JOLIOT CURIE of GENCI, located in France at the Très Grand Centre de Calcul (TGCC) operated by CEA near Paris. JOLIOT CURIE is an Atos/BULL Sequana system X1000 based on a balanced architecture (compute, memory, network and I/O) with 2 compute partitions:

SKL (standard x86)

- 1 656 compute nodes, each with Intel Skylake 8168 24-core 2.7 GHz dual processors, for a total of 79 488 cores and 6.86 petaflop/s peak performance
- 192 GB of DDR4 memory per node – (4GB/core)
- InfiniBand EDR interconnect

KNL (manycore x86)

- 828 Intel KNL 720 nodes each with a 1.4 GHz 68-core processor and 16 GB of MCDRAM for total peak performance of 2.52 PFlops
- 96 GB of DDR4 memory / node
- BULL BXI high speed interconnect

25 additional nodes for post processing and remote visualisation, access to a 500 GB/s multi level Lustre filesystem.★

JUWELS, GCS@FZJ, Germany

The Jülich Wizard for European Leadership Science (JUWELS) is the successor of JUQUEEN and represents a milestone on the road to a new generation of ultraflexible modular supercomputers that can carry out a broader range of tasks – from big data applications right up to compute-intensive simulations. With its first module alone, JUWELS qualified as the fastest supercomputer in Germany.

The cluster module, which was supplied in spring 2018 by French IT company Atos in co-operation with software specialists at German enterprise ParTec, is equipped with Intel Xeon 24-core Skylake CPUs and excels with its versatility and ease of use. It has a theoretical peak performance of 12 petaflop/s. The nodes are connected to a Mellanox InfiniBand high-speed network. Another unique feature of the module is its ultra-energy-efficient warm-water cooling system.★

MARCONI, CINECA, Italy

CINECA's Tier-0 system named MARCONI provides access to PRACE users since July 2016. The MARCONI system is equipped with the new Intel Xeon processors and it has two different partitions:

- Marconi – Broadwell (A1 partition) consists of ~7 Lenovo NeXtScale racks with 72 nodes per rack. Each node contains 2 Broadwell processors that have 18 cores and 128 GB of DDR4 RAM.
- Marconi – KNL (A2 partition) was deployed at the end of 2016 and consists of 3 600 Intel server nodes integrated by Lenovo. Each node contains an Intel Knights Landing processor with 68 cores, 16 GB of MCDRAM and 96 GB of DDR4 RAM.

The entire system is connected via the Intel OmniPath network. The total computational power of the Marconi system is in excess of 20 petaflops.★

MareNostrum, BSC, Spain

MareNostrum is based on Intel's general purpose Xeon E5 processors with 2.1 GHz (two CPUs with 24 cores each per node), 2 GB/core and 240 GB of local SSD disk acting as local/tmp. There are 48 racks, each with 72 compute nodes, giving a total of 3 456 nodes. Just over 200 of the nodes have 8GB/core. All nodes are interconnected through an Intel Omni-Path 100Gbits/s network, with a non-blocking fat tree network topology. MareNostrum has a peak performance of 11.14 petaflops.★

Piz Daint, ETH Zurich/CSCS, Switzerland

The Piz Daint supercomputer is a Cray XC50 system and the flagship system at CSCS – Swiss National Supercomputing Centre, Lugano.

Piz Daint is a hybrid Cray XC50 system with 4 400 nodes. The compute nodes are equipped with an Intel® Xeon® E5-2690 v3 processors with 2.60GHz (12 cores, 64GB RAM) and NVIDIA® Tesla® P100 16GB. The nodes are connected by the "Aries" proprietary interconnect from Cray, with a dragonfly network topology.★

SuperMUC, GCS@ LRZ, Germany

SuperMUC-NG is the Tier-0 supercomputer at the Leibniz Supercomputing Centre (Leibniz-Rechenzentrum, LRZ) of the Bavarian Academy of Sciences and Humanities in Garching near Munich, Germany. It provides resources to PRACE via the German Gauss Centre for Supercomputing (GCS).

SuperMUC-NG consists of 6 336 thin nodes (96 GB each) and 144 fat nodes (768 GB each), equipped with Intel Skylake processors. Each node has 48 cores. All 311 040 compute cores together, connected by an Intel OmniPath Interconnect Network with a fat tree network topology, deliver a peak performance of 26.9 PFlop/s.

The parallel filesystem (IBM Spectrum Scale, GPFS) has a capacity of 50 PByte with 500 GByte/s I/O bandwidth.

For long term data storage, 20 PByte capacity with 70 GByte/s bandwidth are available. The programming environment is Linux (SLES12 SP3), Intel Parallel Studio and OpenHPC. An OpenStack Compute Cloud is attached to SuperMUC-NG.

SuperMUC-NG is cooled with hot water at temperatures up to 50 centigrade. The heat removal efficiency is 97%.

An energy aware scheduling system further assists in saving energy. Adsorption chillers reuse the waste heat to generate cooling for other components.

The LINPACK performance of SuperMUC-NG was measured to be 19.5 PFlop/s, positioning SuperMUC-NG as number eight on the November 2018 world's TOP500 list of supercomputers.★

PRACE

Participation in projects

PRACE is striving to establish an HPC ecosystem at European and international level. Part of this work involves cooperating with other EU-funded scientific, research, industrial and HPC-related projects. The following is a list of projects PRACE participated in during 2019.

eInfraCentral

eInfraCentral's mission is to ensure that, by 2020, a broader and more varied set of users (including industry) discovers and accesses the existing and developing e-infrastructure capacity. A common approach to defining and monitoring e-infrastructures services will increase the uptake and enhance understanding of where improvements can be made in delivering e-infrastructure services.

Project duration: 30 months from 01 January 2017 through 30 June 2019

Funding from the EC: €1 499 037.50 (100%)

Budget for PRACE – Funded: €81 372.50 (100%)

Role of PRACE aisbl: Project Partner

www.einfracentral.eu

Project Consortium

- GEANT VERENIGING
- JNP STRATIGIKI KAI EPICHIRISIAKI SYMVOULEFTIKI IKE
- PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE AISBL
- ETHNIKO KAI KAPODISTRIAKO PANEPISTIMIO ATHINON
- GOTTFRIED WILHELM LEIBNIZ UNIVERSITAET HANNOVER
- EUROPEAN FUTURE INNOVATION SYSTEM CENTRE
- CONSIGLIO NAZIONALE DELLE RICERCHE
- GEANT LIMITED
- STICHTING EGI
- THE UNIVERSITY OF EDINBURGH

ELITRANS

While the implementation of the "Extreme Light Infrastructure" ELI is near completion in the Czech Republic, Hungary and Romania, its remaining challenge is to create the necessary conditions for its future operation as a single, distributed international laser user facility of pan-European dimension. The goal of the ELITRANS project is to complement the final implementation stage of ELI by preparing and undertaking the transformation from three legally (but not conceptually) independent construction projects, towards operation as a single international legal entity, the ELI European Research Infrastructure Consortium (ELI-ERIC).

Project duration: 42 months from 01 September 2015 through 28 February 2019

Funding from the EC: €3 395 383.75 (100%)

Budget for PRACE – Funded: €108 203 (100%)

Role of PRACE aisbl: Project Partner

https://eli-trans.eu

Project Consortium

- INSTITUTUL NATIONAL DE CERCETARE -DEZVOLTARE PENTRU FIZICA SI INGINERIE NUCLEARA "HORIA HULUBEI"
- PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE AISBL
- ELI-HU KUTASI ES FEJLESZTESI NONPROFIT KOZHASZNU KORLATOLT FELELOSSEGU TARSASAG
- COORDINATOR ASSOCIATION INTERNATIONALE EXTREME-LIGHT-INFRASTRUCTURE DELIVERY CONSORTIUM
- KARLSRUHER INSTITUT FUER TECHNOLOGIE
- FYZIKALNI USTAV AV CR V.V.I
- STICHTING EGI

EOSCpilot

The European Open Science Cloud will offer 1.7 million European researchers and 70 million professionals in science and technology a virtual environment with open and seamless services for storage, management, analysis and re-use of research data by federating existing scientific data infrastructures. The EOSCpilot project has been funded to support the first phase in its development.

Project duration: 28 months from 01 Jan 2017 to 30 April 2019

Budget for PRACE – Funded: €80 625 (100%)

Funding from the EC: €9 953 067.50 (100%)

Role of PRACE aisbl: Project Partner

https://eoscpilot.eu

Project Consortium

- GEANT VERENIGING
- ISTITUTO NAZIONALE DI GEOFISICA E VULCANOLOGIA
- SVERIGES METEOROLOGISKA OCH HYDROLOGISKA INSTITUT
- PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE AISBL
- CSC-TIETEEN TIETOTEKNIIKAN KESKUS OY
- THE OPEN UNIVERSITY
- TRUST-IT SERVICES LIMITED
- STICHTING NEDERLANDSE WETENSCHAPPELIJK ONDERZOEK INSTITUTEN
- KARLSRUHER INSTITUT FUER TECHNOLOGIE
- KONINKLIJKE NEDERLANDSE AKADEMIE VAN WETENSCHAPPEN – KNAW
- EUROPEAN SPALLATION SOURCE ERIC
- UNITED KINGDOM ATOMIC ENERGY AUTHORITY
- FRAUNHOFER GESELLSCHAFT ZUR FOERDERUNG DER ANGEWANDTEN FORSCHUNG E.V.
- FUNDACIO CENTRE DE REGULACIO GENOMICA
- ISTITUTO NAZIONALE DI ASTROFISICA
- STICHTING LIBER
- TECHNISCHE UNIVERSITEIT DELFT
- GEORG-AUGUST-UNIVERSITAT GOTTINGENSTIFTUNG OFFENTLICHEN RECHTS
- KONINKLIJK NEDERLANDS METEOROLOGISCH INSTITUUT-KNMI
- VSI DARBO EIGOS
- CINECA CONSORZIO INTERUNIVERSITARIO
- EUROPEAN MOLECULAR BIOLOGY LABORATORY
- BIOBANKS AND BIOMOLECULAR RESOURCES RESEARCH INFRASTRUCTURE CONSORTIUM(BBMRI-ERIC)
- ATHINA-EREVNITIKO KENTRO KAINOTOMIAS STIS TECHNOLOGIES TIS PLIROFORIAS
- TON EPIKOINONION KAI TIS GNOSIS
- CONSIGLIO NAZIONALE DELLE RICERCHE
- MAX-PLANCK-GESELLSCHAFT ZUR FORDERUNG DER WISSENSCHAFTEN EV
- CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE CNRS
- EUROPEAN X-RAY FREE-ELECTRON LASERFACILITY GMBH
- COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES
- JISC LBG
- ISTITUTO NAZIONALE DI FISICA NUCLEARE
- SURFSARA BV
- ACADEMISCH ZIEKENHUIS LEIDEN
- STIFTUNG DEUTSCHES ELEKTRONEN-SYNCHROTRON DESY
- THE UNIVERSITY OF LIVERPOOL
- UNITED KINGDOM RESEARCH AND INNOVATION – **COORDINATOR**
- STICHTING EGI
- ECRIN EUROPEAN CLINICAL RESEARCH INFRASTRUCTURE NETWORK
- THE UNIVERSITY OF MANCHESTER
- BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION
- STICHTING NETHERLANDS ESCIENCE CENTER
- INTEGRATED CARBON OBSERVATION SYSTEM EUROPEAN RESEARCH INFRASTRUCTURECONSORTIUM
- PIN SOC.CONS. A R.L. - SERVIZI DIDATTICI E SCIENTIFICI PER L UNIVERSITA DI FIRENZE
- AGENCIA ESTATAL CONSEJO SUPERIOR DE INVESTIGACIONES CIENTIFICAS
- THE UNIVERSITY OF EDINBURGH

HPC-GIG

The European Cloud Initiative implementation encompasses different elements and will require a solid governance setup. A cornerstone of this governance is a joint undertaking, which will become active in 2019. HPC-GIG proposes intelligence gathering services from the European HPC stakeholders for this future HPC Governance - HPCG. It will facilitate a timely start of its operations, transparently and non-intrusively leveraging the HPC community know-how, organised with enough flexibility to align project activities with the programme schedule and needs during its ramp-up phase, and deliver timely and useful input to the HPCG.

Project duration: 18 months from 01 September 2018 through 29 February 2020

Funding from the EC: €900 000 (95%)

Budget for PRACE – Funded: €400.432 (89%)

Role of PRACE aisbl: Co-ordinator

Project Consortium

- PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE AISBL – COORDINATOR
- GEANT VERENIGING
- EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

Members & Partners

FocusCoE

FocusCoE will contribute to the success of the EU HPC ecosystem and the EuroHPC Initiative by supporting the EU HPC Centres of Excellence to more effectively fulfil their role within the ecosystem, ensuring that extreme scale applications result in tangible benefits for addressing scientific, industrial or societal challenges. It will do this by creating an effective platform for the centres to coordinate strategic directions and collaboration, and will provide support services for the centres in relation to both industrial outreach and promotion of their services and competences.

Project duration: 36 months from 01 December 2018 through 30 November 2021

Funding from the EC: €1 997 921.25 (100%)

Budget for PRACE – Funded: €16 437.50 (100%)

Role of PRACE aisbl: Project Partner

www.focus-coe.eu

Project Consortium

- UNIVERSITY COLLEGE LONDON
- AGENZIA NAZIONALE PER LE NUOVE TECNOLOGIE
- L'ENERGIA E LO SVILUPPO ECONOMICO SOSTENIBILE
- KUNGLIGA TEKNISKA HOEGSKOLAN
- PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE AISBL
- TERATEC
- SCAPOS AG – COORDINATOR
- FORSCHUNGSZENTRUM JULICH GMBH
- NATIONAL UNIVERSITY OF IRELAND GALWAY
- BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION
- UNIVERSITAET STUTTGART
- COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES

PRACE-5IP

The objectives of PRACE-5IP are to build on and seamlessly continue the successes of PRACE and start new innovative and collaborative activities proposed by the consortium. The activities are designed to increase Europe's research and innovation potential, especially through seamless and efficient Tier-0 services and a pan-European HPC ecosystem.

Project duration: 28 months from 01 Jan 2017 to 30 April 2019
Funding from the EC: €15 000 000.00 (93%)
Budget for PRACE – Funded: €1 123 215 (85%)
Role of PRACE aisbl: Project Partner

www.prace-ri.eu/prace-5ip

Project Consortium

- PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE AISBL
- CSC-TIETEEN TIETOTEKNIIKAN KESKUS OY
- FORSCHUNGSZENTRUM JULICH GMBH – COORDINATOR
- ASSOCIATION "NATIONAL CENTRE FOR SUPERCOMPUTING APPLICATIONS
- CENTRE DE RECHERCHE EN AERONAUTIQUE ASBL – CENAERO
- UPPSALA UNIVERSITET
- GRAND EQUIPEMENT NATIONAL DE CALCUL INTENSIF
- GAUSS CENTRE FOR SUPERCOMPUTING (GCS)
- KORMANYZATI INFORMATIKAI FEJLESZTESI UGYNOKSEG
- VYSOKA SKOLA BANSKA - TECHNICKA UNIVERZITA OSTRAVA
- CINECA CONSORZIO INTERUNIVERSITARIO
- UNIVERSIDADE DE COIMBRA
- ETHNIKO DIKTYO EREVNAS TECHNOLOGIAS AE
- INSTYTUT CHEMII BIOORGANICZNEJ POLSKIEJ AKADEMII NAUK
- UNINETT SIGMA2 AS
- THE CYPRUS INSTITUTE
- SURFSARA BV
- CENTRUM SPOLOCNYCH CINNOSTI SLOVENSKEJ AKADEMIE VIED
- ISTANBUL TEKNIK UNIVERSITESI
- MACHBA - INTERUNIVERSITY COMPUTATION CENTER
- UNIVERZA V LJUBLJANI
- KOBENHAVNS UNIVERSITET
- NATIONAL UNIVERSITY OF IRELAND GALWAY
- BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION
- EIDGENOESSISCHE TECHNISCHE HOCHSCHULE ZUERICH
- THE UNIVERSITY OF EDINBURGH

EXDCI-2

The project EXDCI-2 builds upon the success of EXDCI and will continue the coordination of the HPC ecosystem with important enhancements to better address the convergence of big data, cloud and HPC. Its main goals are to develop a competitive European HPC exascale strategy by supporting the implementation of a common European HPC strategy, and to coordinate the stakeholder community for exascale HPC in Europe through joint community structuring and synchronisation.

Project duration: 30 months from 01 March 2018 through 31 August 2020

Funding from the EC: €2 440 000.00 (94%)

Budget for PRACE – Funded: €864 065 (100%)

Role of PRACE aisbl: Co-ordinator

https://exdci.eu

Project Consortium

- PARTNERSHIP FOR ADVANCED COMPUTING IN EUROPE AISBL – COORDINATOR
- UNIVERSITE DE RENNES I
- UNIVERSITE DE PICARDIE JULES VERNE
- EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

PRACE-6IP

The PRACE-6IP project continues and extends the scope of the PRACE-5IP project. The main objectives are: preparing for exascale computing through the development of forward-looking software solutions; assisting the development of PRACE 2; increasing Europe's research and innovation through seamless and efficient Tier-0 services and a pan-European HPC ecosystem including national capabilities; promoting take-up by industry and special offers to SMEs; and collaborating with other stakeholders at European and international level on future architectures, training, application, support and policies.

Project duration: 28 months from 01 Jan 2017 to 30 April 2019
Funding from the EC: €9 953 067.50 (100%)

Budget for PRACE – Funded: €80 625 (100%)
Role of PRACE aisbl: Project Partner

https://prace-ri.eu/about/ip-projects/#PRACE6IP

Project Consortium

- FORSCHUNGSZENTRUM JULICH GMBH - CO-ORDINATOR
- GAUSS CENTRE FOR SUPERCOMPUTING (GCS)
- GRAND EQUIPEMENT NATIONAL DE CALCUL INTENSIF (GENCI)
- THE UNIVERSITY OF EDINBURGH (UEDIN)
- BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION (BSC)
- CSC-TIETEEN TIETOTEKNIIKAN KESKUS OY (CSC)
- EIDGENOESSISCHE TECHNISCHE HOCHSCHULE ZUERICH (ETH Zürich)
- SURFSARA BV (SURFsara)
- UPPSALA UNIVERSITET (SNIC-UU)
- CINECA CONSORZIO INTERUNIVERSITARIO (CINECA)
- INSTYTUT CHEMII BIOORGANICZNEJ POLSKIEJ AKADEMII NAUK (PSNC)
- UNINETT SIGMA2 AS (SIGMA2)
- NATIONAL INFRASTRUCTURES FOR RESEARCH AND TECHNOLOGY S.A. (GRNET)
- UNIVERSIDADE DE COIMBRA (UC-LCA)
- NATIONAL UNIVERSITY OF IRELAND GALWAY (NUI Galway)
- ISTANBUL TEKNIK UNIVERSITESI (UHEM)
- THE CYPRUS INSTITUTE (CaSToRC)
- ASSOCIATION "NATIONAL CENTRE FOR SUPERCOMPUTING APPLICATIONS (NCSA)
- VYSOKA SKOLA BANSKA - TECHNICKA UNIVERZITA OSTRAVA (IT4I)
- KORMANYZATI INFORMATIKAI FEJLESZTESI UGYNOKSEG (KIFU)
- KOBENHAVNS UNIVERSITET (UCPH)
- MACHBA - INTERUNIVERSITY COMPUTATION CENTER (IUCC)
- PPRACE AISBL (PRACE)
- UNIVERZA V LJUBLJANI (UL)
- CENTRUM SPOLOCNYCH CINNOSTI SLOVENSKEJ AKADEMIE VIED (CCSAS)
- UNIVERSITEIT ANTWERPEN (UANTWERPEN)
- UNIVERSITE DU LUXEMBOURG (UdL)
- TECHNISCHE UNIVERSITAET WIEN (TU WIEN)
- GEANT VERENIGING (GEANT)
- EUDAT OY (EUDAT)

European Extreme Data & Computing Initiative 2 (EXDCI-2)

The European Extreme Data & Computing Initiative 2 (EXDCI-2) supports multiple aspects of the HPC ecosystem, such as the organisation of conferences, HPC projects, analysis of innovations within Centres of Excellence, road mapping, calls for proposals, and SME technology development.

During 2019, EXDCI-2 created a new strategic research agenda for HPC technologies. This work has involved more than 80 experts from our HPC ecosystem. The main finding is the emergence of new applications deployed on the "digital continuum", ranging from sensors to HPC systems, all of which need to be supported by both hardware and software technologies.

To help in the development of these new technologies for this digital continuum, active co-operation with other ecosystems, namely big data, cyber security, the Internet of Things and 5G, has been developed. Other collaborations with the photonics and electronics communities have also been undertaken to build a European technology value chain from basic technologies up to HPC systems.

The work to establish a “Roadmap of HPC applications and usages” confirms that on the methodological side, Artificial Intelligence (AI), or more precisely Machine Learning (ML) and Deep Learning (DL), are changing the approach to scientific problems. On the one hand ML is contributing to more traditional HPC by offering methods for either performing more efficiently some parts of the calculations (including parameterisation schemes and solver preconditioners), or identifying features which would otherwise be very difficult to identify, or providing innovative methods to process the wealth of data produced by these calculations.

On the other hand, this physics-based HPC production of important datasets contributes to ML efficiency, while physics-based HPC is also able to constrain neural network approaches. This convergence between HPC and AI/ML is a further incentive for the actual development of the so-called converged infrastructures.

The European HPC research projects have been monitored with the objective to increase the re-use of results between them and the impact they can have on future products. A spin-off project has been selected to experiment how to accelerate the take-off of an innovation created by a research project.

The impact of the HPC European research programme has been evaluated through a set of KPIs and a new viscosity concept has been developed to characterise legacy HPC software evolution. In the advent of a new generation of HPC systems, the question is whether legacy softwares are still relevant to their owners, and if yes, how these codes can be rewritten to efficiently run on future HPC systems. In this case, it is important to identify the effort that would be required.

This work brought insight on the European portfolio of legacy codes through the analysis of a survey that has been distributed amongst European code developers. Between November 2019 and January 2020, the survey was distributed to many contacts in industry and in academia. Based on the survey responses, the conclusions of these findings are that despite the investment already in Centres of Excellence (CoE), more funding needs to be dedicated to code porting.

At international level, EXDCI-2 has enabled the presence of European experts in the Big Data Extreme Computing (BDEC) series of workshops. Europe has organised one of these events during the EuroHPC Summit Week (EHPCSW) 2019 in Poznań, and continues to lead the reflection with the USA on the future of HPC.

In 2019, EXDCI-2 has also conducted an analysis of the European position in HPC related standard organisations. Some recommendations have been proposed to foster this position and to make Europe one of the leaders in HPC new standards.

Also in 2019, EXDCI-2 has been present in a set of well-known international events such as SC19, ISC19, the EHPCSW 2019 and the Teratec forum. ★

From the desk of

The PRACE Council Chair

The year 2019 was a year of many profound changes for the European HPC ecosystem, in particular following the establishment of the EuroHPC Joint Undertaking on 28 October 2018. In 2020/21, the EuroHPC JU will procure and operate five petascale and three pre-exascale systems, which will promote Europe's place in the global HPC race. PRACE has influenced the significance of these new and exciting developments, incorporated them early in its planning, and started to shape its role in the future landscape of European HPC. The PRACE position paper, which was published in 2018, provides a solid frame of reference for our common future.

The planning for the future co-operation between EuroHPC and PRACE was carried out as a main activity for 2019 and will continue at full speed in 2020. Several very successful coordination meetings between the chairs of the EuroHPC Joint Undertaking and PRACE took place in order to discuss and agree on the roles of EuroHPC and PRACE in the new European HPC ecosystem.

One of the main topics was the development of new access channels in order to include EuroHPC systems in the PRACE award process for the upcoming EuroHPC JU systems. For the specific requirements defined by the EuroHPC JU for access to the new systems for scientific and industrial research, PRACE developed appropriate access channels. Together with the Brussels office of PRACE, the Scientific Steering Committee of PRACE, and the PRACE-6IP project, a comprehensive scheme of adapted access channels was proposed and positively received by the EuroHPC JU.

The EuroHPC Joint Undertaking is running at full speed, and the earliest installations of their systems from the EuroHPC Joint Undertaking are expected by the end of 2020. In order to avoid a gap in the provision of computing resources to the European Research Communities, PRACE has decided to extend the current funding agreement called "PRACE 2" by providing resources until March 2022. With this extension, PRACE 2 will allocate computer resources from the hosting members France, Germany, Italy, Spain, and Switzerland for the PRACE Project Access Calls 20, 21 and 22.

Prof Dr Dr Thomas Lippert, Chair of the PRACE Council

I sincerely thank you all for your hard work in developing PRACE into a European HPC infrastructure at the service of science. Much has already been achieved, but there is still much work to be done.

Last year, a pilot project to increase the industrial acceptance of Tier-0 high-performance computers was also carried out as part of the 19th PRACE Call for Project Proposals for Project Access. This was also continued in the 20th Call. At the end of 2019, PRACE aisbl recruited an industry liaison officer with the specific task of reaching out to industrial communities and sectors and making them aware of the benefits of HPC.

During the International Supercomputing Conference (SC19) in November 2019, PRACE was awarded the HPCwire Readers' and Editors' Choice Award in the category Workforce Diversity Leadership for its PRACE Ada Lovelace Award for HPC. This demonstrates that PRACE is at the forefront of promoting diversity and inclusivity in HPC and in fields beyond.

I sincerely thank you all for your contributions and hard work over the past year in developing PRACE into a European HPC infrastructure at the service of science. Much has already been achieved, but there is still much work to be done. I look forward to seeing how the new systems of the EuroHPC Joint Undertaking will provide valuable computing time for the European research communities distributed and supported by PRACE. ★

Q&A

Dr Masahiro Seki, President of RIST

Dr Masahiro Seki is president of RIST, the Research Organisation for Information Science and Technology in Japan. We spoke to him about the objectives of RIST, its international collaborations with PRACE and XSEDE, the new Fugaku supercomputer, and the future of HPC.

Please tell us about RIST is and what its objectives are.

Before RIST existed, its predecessor, the Nuclear Energy Data Centre (NEDAC), was founded in 1981 to promote the development and utilisation of nuclear computer codes. The scope of work was then expanded to incorporate the support of integrated computational science including the support of the Earth Science Projects, and these were all consolidated into RIST – the Research Organisation for Information Science and Technology – in 1995

RIST is a general incorporated foundation that aids the development and utilisation of computational science and technology in order to support our highly information-oriented society.

The RIST Kobe Centre coordinates HPCI. HPCI – the High-Performance Computing Infrastructure of Japan – was established in 2012 as a national HPC infrastructure, connecting the flagship systems and the systems of major universities and national laboratories with a high-speed academic network. To accelerate the use of HPC in Japan and help it contribute to society, HPCI provides world-class computing resources with a variety of systems through open calls. RIST is also responsible for user selection, resource allocation and user support for these systems.

What are your views on the RIST-XSEDE-PRACE collaboration?

As you know, we have agreed a memorandum of understanding on information exchange for promoting the publica utilisation of HPC, and this has been very fruitful for us. We have just started a new programme, a joint call for interoperable support services

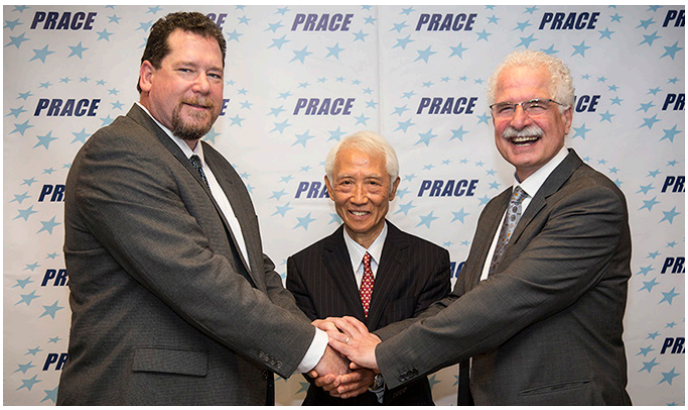

Dr Seki with John Towns, project director for XSEDE and Anwar Osseyran, former Chair of the PRACE Council.

between RIST, XSEDE and PRACE. Fugaku, a supercomputer installed at RIKEN Center for Computational Science in Kobe, will begin public operation in 2021. After that, we will try to make Fugaku resources available through a RIST joint call. The amount of allocations will not be that large, but we will try to make those resources available. I am expecting that the use of Fugaku will bring about an increase in applications for the joint call. I hope that our co-operation can contribute to enhancing international collaboration in the field of science and technology.

Right now, our information exchange takes place only in a limited manner. In the future, I think if we exchange people – sending our colleagues to PRACE and bringing some of the researchers from Europe to us – we can really help to strengthen our collaboration. In particular, I think that the exchange of young scientists would be very fruitful for us.

Could you tell us more about the Fugaku project?

In fact, RIST has not taken any responsibility in the development and construction of Fugaku, but I will give you a rough outline of the project. The project goal is to develop the next generation of supercomputers in Japan, the successors to the K computers. It also aims to develop a wide range of applications that will resolve high-priority social and scientific issues.

The development of Fugaku started in 2014. Over 400 racks were installed at RIKEN, and it will start official operation in 2021. The first open call for the use of Fugaku will start this autumn. The total budget, including for application projects, has been approximately one billion US dollars.

We see the Fugaku machine as a game-changer due to the amount of new technology that has been used for it. It is ranked as the world's greenest supercomputer, and uses ARM CPUs, which will allow a wide range of applications to be run on it and will make it useful in the fields of data science and artificial intelligence. 384 of the CPUs developed by Fujitsu can fit into one rack, so the Fugaku machine has just 10 racks for every 900 that the K machines had.

There have been three key players in this development. The RIKEN Centre for Computational Science has led the development of Fugaku. Fujitsu Ltd. was selected as the development and manufacturing vendor of Fugaku. Finally, RIST was appointed to be the organisation responsible for the facilities use promotion, user selection, resource allocation, user support, publication management, and other relevant activities of Fugaku.

What do you see as the main challenges for the world of high-performance computing in the near future? How can we meet these challenges?

One of the key issues for HPC in the future will be the fusion of HPC, artificial intelligence, and data science. This includes the application of AI and data science technologies to HPC, which will enhance the capabilities and performance of HPC simulations. On the other hand, the application of HPC technologies to AI and data science will accelerate those fields due to the huge amounts of data generated by HPC simulations.

We see the Fugaku machine as a game-changer due to the amount of new technology that has been used for it. It is ranked as the world's greenest supercomputer

HPC currently plays an important role in quantum computing R&D, and I expect this to continue in the future. This will include quantum algorithm research and the enrichment of quantum computing through HPC platforms.

Japan is adopting the field of HPC as a strategic subject in Japanese science and technology basic plan. This means the government will support this field. Moving forward, I think enhancing user support capabilities, especially human resources, will be of great importance. Scientists, although extremely intelligent and competent within their disciplines, are not necessarily experts in using computers. That is why it is so important to surround them with skilled support staff who can help them use supercomputers to their full potential.

In coordination with Japan's Ministry of Education, Culture, Sports, Science and Technology (MEXT), RIST is now planning to have a new shared use framework for Japanese HPCI. A new resource allocation category which promotes and accelerates international collaboration may be included in this framework.

In the future, we expect continuous support from PRACE and XSEDE, and close communication with them for international computing resources, human resources, and data sharing activities. As I mentioned before, we would like to examine the possibility of a human resource exchange programme between these organisations. Through these activities, we would like to contribute to the global HPC community. ★

Our thanks to Dr Seki for contributing to the PRACE Annual Report 2019

From the desk of

The PRACE User Forum

The PRACE User Forum represents the users of PRACE computational resources and provides a communication channel between the users of the PRACE research infrastructure and the computational centres. Our concerns are generic issues, such as problems in the review process or during the allocation of resources, and the provision by PRACE for new use-cases that are currently not catered for. We also give a voice to the users of high-performance computing in the debate about the development of HPC in Europe, which has become ever more important with the transition to exascale and the advent of EuroHPC.

PRACEdays19

In 2019 the User Forum held its seventh open session in Poznań during PRACEdays19 as part of the EuroHPC Summit Week. The User Forum was happy to see the strong uptake from users and their contribution to the EuroHPC Summit Week giving them the possibility to interact with peers from different disciplines, and showcasing the main product of PRACE: the many excellent scientific and industrial results produced through access to world-leading HPC infrastructure in Europe.

Surveying Tier-0 PRACE users

In 2019 the most important activity of the User Forum was a thorough survey of present and past PRACE Tier-0 users to understand HPC user requirements and their view of PRACE and HPC in Europe in general. This gathered responses from more than 50 current and former Tier-0 PIs representing all scientific domains. The survey highlighted a number of interesting trends that are particularly important given the transition to accelerated pre-exascale systems in Europe under the EuroHPC joint undertaking. 80% of the research groups are using "in-house" codes developed by the group, underscoring the importance of technical insight when using the largest HPC infrastructures in Europe. This includes all of the largest user group answering the survey – those that consume more than 100 million core hours per year. The results also showed that only 36% of the users have access to codes that are adapted for running on either many-cores, such as Blue-Gene and

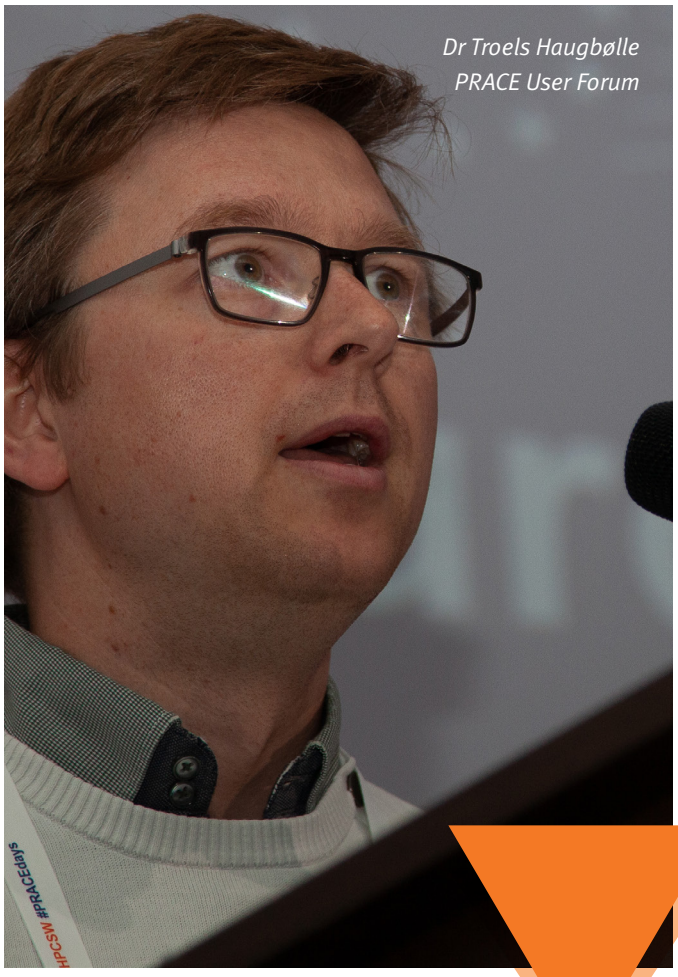

Dr Troels Haugbølle
PRACE User Forum

In the User Forum, we hope the same process, based on scientific excellence and technical readiness, will be used for access to future pre-exascale EuroHPC systems. We strongly believe this will result in access to the most excellent scientific projects.

Xeon Phi, or accelerators, such as GPUs. At present, the typical PRACE Tier-0 user is using a code adapted to traditional multi-core platforms. Another highlight from the survey was the result on workflows. While the majority of users mostly carry out traditional large-scale modelling, there is a growing demand for data intensive science, post processing, data portals, and long-term storage. This is something the EU is now addressing with the Fenix research infrastructure.

Over the years, the PRACE peer review process is by far the most important topic that the User Forum has addressed. We have forwarded and consolidated suggestions by the user community for improving the peer review process, and lobbied for a transparent peer review. We are therefore very happy to see the general praise given by the users for the PRACE application process that is viewed by the large majority of users as fair, impartial, and with the right balance between reward and workload. In particular, PRACE should be applauded for implementing a transparent refereeing process with the possibility of responding to reviewers' comments by the applicants. In the User Forum, we hope the same process, based on scientific excellence and technical readiness, will be used for access to future pre-exascale EuroHPC systems. We strongly believe this will result in giving access to the most excellent scientific projects.

Overall, the survey shows how appreciated and unique the PRACE user infrastructure is, as highlighted by one of the comments submitted in the survey: "I think in Europe there is no alternative to PRACE if one need an allocation of size in the order of 100M core hours."

In Europe there is no alternative to PRACE if one need an allocation of size in the order of 100M core hours

EuroHPC and green accelerated computing

Three next generation pre-exascale system will be installed in Europe in early 2021 under the auspices of EuroHPC. The main users of EuroHPC will be the current Tier-0 PRACE users that include some of the best European scientists, and the groups who have applications that can run efficiently on the current largest machines. All three installations will obtain the majority of the performance using accelerators such as GPUs or FPGAs. This reflects a general paradigm shift in high-performance computing internationally towards a greener HPC infrastructure, and was the only possibility given the combined economic and performance envelope required by EuroHPC.

Running efficiently on a pre-exascale machine is a formidable challenge, and many user groups across Europe are right now developing new codes and rewriting old ones, to meet the challenge of a 100 million-level concurrency.

The user survey showed that even at the echelon of computational science in Europe a minority of the user groups are ready to efficiently exploit accelerators. Therefore, it is important to develop not only hardware and middleware, but also support new approaches to end-user applications. Otherwise, Europe risks sub-optimal use of its new world-class hardware by a smaller community than the current PRACE machines.

Europe has a long tradition of international leadership in HPC applications. To maintain this, it is critical that the support programmes and EU calls by EuroHPC, for example, support the development of scalable parallel applications, with a particular emphasis on accelerated computing. As also shown by the survey, most of the leading computational scientists in Europe use in-house developed codes; the Formula 1 racecars of the HPC racetracks. Therefore, in the User Forum we believe it is absolutely critical that the EU and EuroHPC not only support development of middleware and standard packages, but also allocate resources for smaller groups to optimise and modernise existing applications. This will stimulate the creation of a diverse ecosystem of exascale-ready applications and is a critical requisite for having a large enough user base for the future pre-exascale installations that very soon will be operational in Europe. ★

Supporting theoretical science

Laser-matter interactions at ultra-high intensities


Jorge Vieira

Jorge Vieira is an assistant professor and researcher at the Institute for Plasmas and Nuclear Fusion at IST in Lisbon. He is currently leading a PRACE-supported project that aims to explore how high intensity lasers with exotic properties can be used to accelerate particles and generate radiation.

The project we are currently working on is called "OptiMom - Optical angular momentum in laser-matter interactions at ultra-high intensities". It is concerned with the use of intense lasers in connection with plasmas to investigate new regimes that could lead to accelerated particles and radiation with new properties. The distinctive aspect of this work is that we are investigating lasers with exotic properties. One of these properties, known as orbital angular momentum, can cause matter to spin and create magnetic fields that can be used for radiation generation.

These lasers are fairly well studied at low intensities, but their interactions at very high intensities are not well known. As soon as a laser with such high intensity interacts with matter, it will instantly ionise it into plasma, a state where electrons become unbound from the nucleus forming a soup of electrons and ions.

We have access to amazing software, but without supercomputers we would not be able to make meaningful advances

We are investigating how these lasers interact with matter at very high intensity and their potential applications for accelerating particles and generating radiation. These processes are extremely complicated. When you think back to the physics problems that you used to solve at school, they were always very well defined in terms of the physics and maths, and allowed you to calculate a straightforward answer. However, in research where you are looking at real problems, it is never really possible to solve things in such an exact manner. The problems we are looking at are highly non-linear, and relativity plays a crucial role given the laser intensities we consider. Besides, we need to capture the motion of millions of particles at the same time, as well as understanding physics at a huge range of scales simultaneously, from the level of millimetres to below the microscale. This means that it is very hard to build an exact theoretical answer. Instead, we rely on simulations to capture both the motion of all the particles self-consistently with the electric and magnetic fields that are created by this motion.

These simulations cannot be run on laptops as the memory required to track that many particles is too great, and the calculations would take far too long. Access to supercomputers is crucial as it allows us to divide the problem between several computing nodes so that each one can address a smaller, more manageable problem and provide results within hours. The answers given by these simulations provide points of data that allow us to develop predictive models for these new areas of physics.

The project is halfway through now, and so far we have obtained some results that point towards a previously unrecognised mechanism that generalises the usual temporally coherent x-ray emission processes in free electron lasers. This could open new research pathways for X-ray emissions in compact plasma accelerators. In general, our experience has been extremely positive. Without access to HPC, a theory simulation team like ours would not be able to have the impact that we do – we would be like an experimental team without a lab. We have access to amazing software, but without supercomputers to run it on we would not be able to make meaningful advances and develop theoretical models for making predictive theories about the complex problems we are investigating.★

PRACE industry access

Corner separation predictions using LES

Dimitrios Papadogiannis, a specialist in high fidelity simulations for Safran Tech, is currently leading a project that is investigating corner separation in aeronautical engine compressors. Provided with an allocation through the PRACE Industry Access Track, his project is simulating this phenomenon at a high level of accuracy, which will in the future help to improve less computationally expensive models.

My main employer is Safran, an international group that works on aeronautics and defence. Specifically, I work for Safran Tech, which does applied research and facilitates knowledge transfer between academia and industry. I work in the modelling and simulation department, specialising in high fidelity simulations of rotating components of engines. The project I am currently leading is called "CornerLES – Corner Separation Predictions using Large Eddy Simulations", which was awarded 20.4 million core hours on Joliot Curie-AMD@GENCI@CEA, France. This PRACE allocation, which was given as part of the PRACE Industry Access Track, will allow us to analyse in unprecedented detail the corner separation that occurs in compressors of aeronautical engines.

Compressors are vital components in aircraft engines – the first half of most engines consists of various stages of compression. They are essentially blades, some spinning and some non-spinning, that increase the pressure of the air before coming into the combustion chamber. They are critical components as they impact on the efficiency of the engine and also control the stability of the engine. The main instabilities that can hinder the safe operation of engines come from the compressor components, so it is a vital part for all aeronautical engines. Corner separation is a phenomenon that occurs in compressors at the junction of the blade and the hub. Normally the job of the blade is to guide the flow, directing it to a specific angle that will help compression. But at this junction, different pressure gradients can lead to what is called flow separation, where the flow no longer follows the form of the blade as intended. This causes it to separate from the normal path and creates a zone where a lot of undesirable things can happen. When you have a flow not acting like you want it to, it creates aerodynamic losses and results in reduced efficiency of the component, which is very undesirable. Another problem with corner separation is that if it grows in size, it can lead to instabilities that completely perturb the operation of the compressor as a whole, causing what is known as a rotating stall.

Corner separation is particularly tricky to investigate because it changes shape and size through time, making it very hard to capture in a simulation. The industry standard is to use relatively affordable simulations that use a lot of modelling assumptions about important turbulence phenomena and are steady state, meaning that they don't take into account unsteadiness over time. We are applying high-fidelity CFD methods – large eddy simulations – to get much more accurate predictions of what is going on in corner separation and how it potentially grows over time to cause instabilities that can hinder the performance of the compressor. Simulations such as these produce terabytes of data. In this project, we had a numerical domain that had over 1.1 billion degrees of freedom. Therefore, a second part will be to hand over our data to other partners in a Horizon 2020 project who will use it to improve the more affordable models that are commonly used in industry. In this way, our work will benefit many others.

The project has been a great experience for us. We are collaborating with CERFACS in Toulouse, who are world-leading in the field of large eddy simulations and were responsible for developing the code we used. Interacting with them on a regular basis has been a very fruitful experience and has accelerated our efforts to create a real impact on these problems we are investigating. We have also received invaluable support from the supercomputing centre. From the moment we received the allocation we have worked with them a lot, and their exceptional input has been a vital part of our success. so far.★

Supporting climate research

Predicting atmospheric growth rate of CO_2

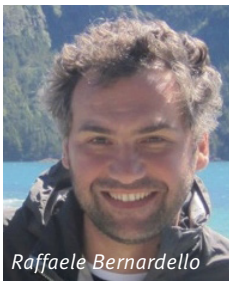
Raffaele Bernardello

Raffaele Bernardello, a research scientist at the Barcelona Supercomputing Centre (BSC), is currently leading a PRACE-supported project – TOPSyCled – that is looking to establish a system for predicting the atmospheric growth rate of CO_2. His work will help to improve negotiations in future climate summits and is made possible by the workflow manager tool developed in-house at BSC.

We are part of a Horizon 2020 project called 4C, which is trying to gain better knowledge about the global carbon cycle and the interactions it has with climate. Part of this project is helping to improve databases based on observations, while the other part involves using climate models. This is where our PRACE-supported project, TOPSyCled, comes in.

We are trying to establish a system for predicting the atmospheric growth rate of CO_2. The reason this is important is that although the overall trend is clearly going upwards, growth rates year on year are extremely variable. This is mostly due to natural variability in the location of carbon. Carbon doesn't just exist in the atmosphere – it is also present in the ocean and in the land, and depending on natural climate oscillations such as El Niño, the amount held in each changes constantly.

When you frame this in the context of the Paris Agreement, this natural variability of atmospheric carbon can be damaging in terms of countries sticking to their CO_2 emission reduction commitments. In 2023, a global stocktake will occur where the emission reductions are analysed and then compared with actual observations of atmospheric CO_2 growth rate. If natural variability causes CO_2 growth rate to be lower than expected, policymakers might decide that their actions taken were too severe. Conversely, if the natural variability causes CO_2 growth rate to be higher than expected, policymakers might decide that the measures are not working, or even begin to lose trust that other signees of the agreement are doing what they say they have done. This would make further co-operation around climate more difficult in the future. The aim of this project is to make a way to predict this natural variability in order to inform the stocktake so that we can see the real effect of our efforts to reduce emissions and help improve future negotiations.

How do we do this? With climate prediction, the physics used in the models is the same as that used for weather prediction. However, they are used in a different way. Weather predictions depend strongly on knowing the initial state of the weather at the beginning. If you then have a system that models the physics sufficiently well, you can predict the evolution of the system over time. This works up to a point, but limitations due to imperfect observations and our inability to fully represent variables such as turbulence in simulations mean that weather prediction is only reliable up to around one week in the future.

On the opposite side, with climate projections over timescales of around 100 years, the initial state of the weather does not matter, as the memory of these initial conditions is lost over time. What matters when trying to accurately create climate projections far into the future are boundary conditions such as

Effect of sustained reductions in CO_2 emissions

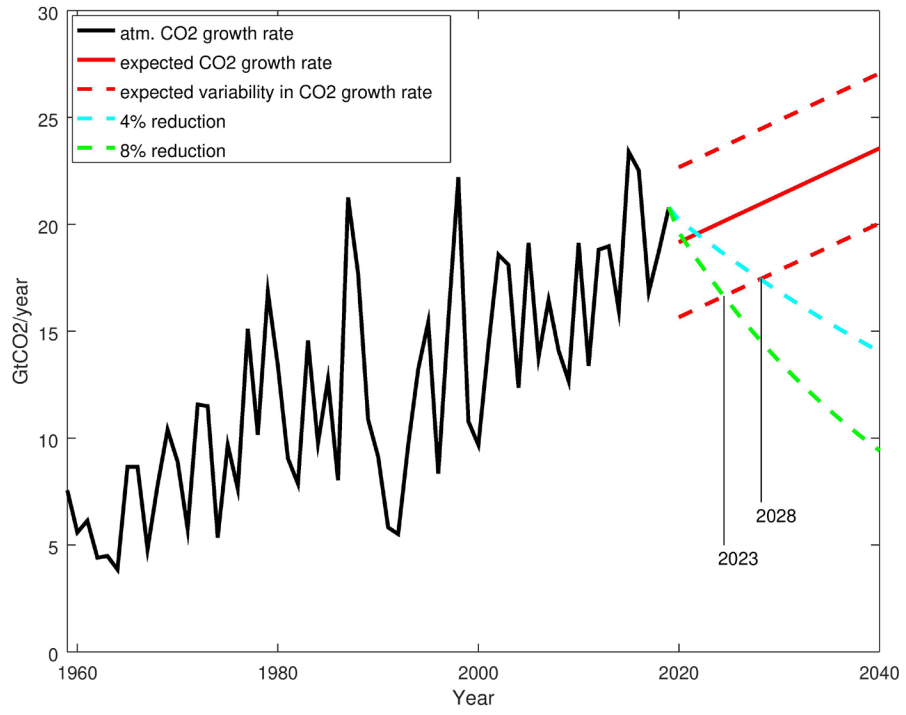

solar activity and greenhouse gases. The predictions made are also very different. Weather predictions will tell you that it will rain in a specific place at a specific time, whereas climate projections use more generalised terms, for example saying that certain extreme weather events will be more likely, or that average rainfall in a region will be lower.

In the middle of these two ends of the climate modelling spectrum you have what is known as seasonal to decadal predictions. With these, both the initial conditions and the boundary conditions are important. Models for predicting the future at these timescales are first tested by seeing if they can predict the past. Initial conditions from past observations are used as a starting point. But, to account for the imperfectness of these observations, many simulations are run and the initial conditions are perturbed slightly for each one. An average of this ensemble of simulations is then taken and compared with what the actual observations were later on to assess the accuracy of the model.

Weather predictions will tell you that it will rain in a specific place at a specific time, whereas climate projections use more generalised terms, for example saying that certain extreme weather events will be more likely

In the TOPSyCled project, we are applying the techniques used for seasonal to decadal climate predictions to the carbon cycle. We use the same climate models but also include a representation of the carbon cycle that is coupled to the physics of the climate model. We are currently working on testing a model over five years from 30 different starting dates in the past, each of which will have 15 different perturbations. We therefore rely heavily on PRACE resources, as carrying out this many simulations is extremely computationally expensive.

One tool that is extremely important to us is the workflow manager software called Autosubmit. This tool is developed and maintained in-house here at the Barcelona Supercomputing Centre and allows us to submit all of the simulations in our ensembles very easily. Submitting each of these simulations manually would be very time consuming, as would gathering all of the data generated.

The workflow manager automatically splits all of the different pieces of our process and compiles the script for each one, then works out the most efficient way to do all of this work in the parallelised MareNostrum architecture. Without this tool developed by BSC, our work would be impossible.★

PRACE-6IP project

Supporting research infrastructure

In May 2019, the 6th PRACE Implementation Phase (PRACE-6IP) project began, supporting the implementation of PRACE aisbl and initiating new innovative and collaborative activities, such as:

- Preparing for exascale computing by developing forward-looking software solutions
- Assisting the development of PRACE 2 to increase Europe's research and innovation This will be achieved through seamless and efficient Tier-0 services, and a pan-European HPC ecosystem which includes national capabilities, promoting take-up by industry, and special offers to SMEs
- Collaborating with other stakeholders at European and international level on future architectures, training, application, support, and policies

The PRACE-6IP project is a joint effort, where 26 PRACE countries contribute and work together to deliver the best support for PRACE users and the implementation of the PRACE research infrastructure. So far, the PRACE IP projects have received €121 million of funding from the EC. The project partners supported the projects with €161 million.

On the following pages, we present highlights of the PRACE-6IP project from 2019

Forward-looking software solutions projects

Performance-portable linear algebra

The goal of the "Performance-portable linear algebra" project is to implement distributed eigenvalue solvers for dense matrices. The project contributions are divided in two parts:

- implementation of a direct eigensolver based on MPI and HPX libraries
- improvement of the performance of the iterative eigensolver implemented in the ChASE Library

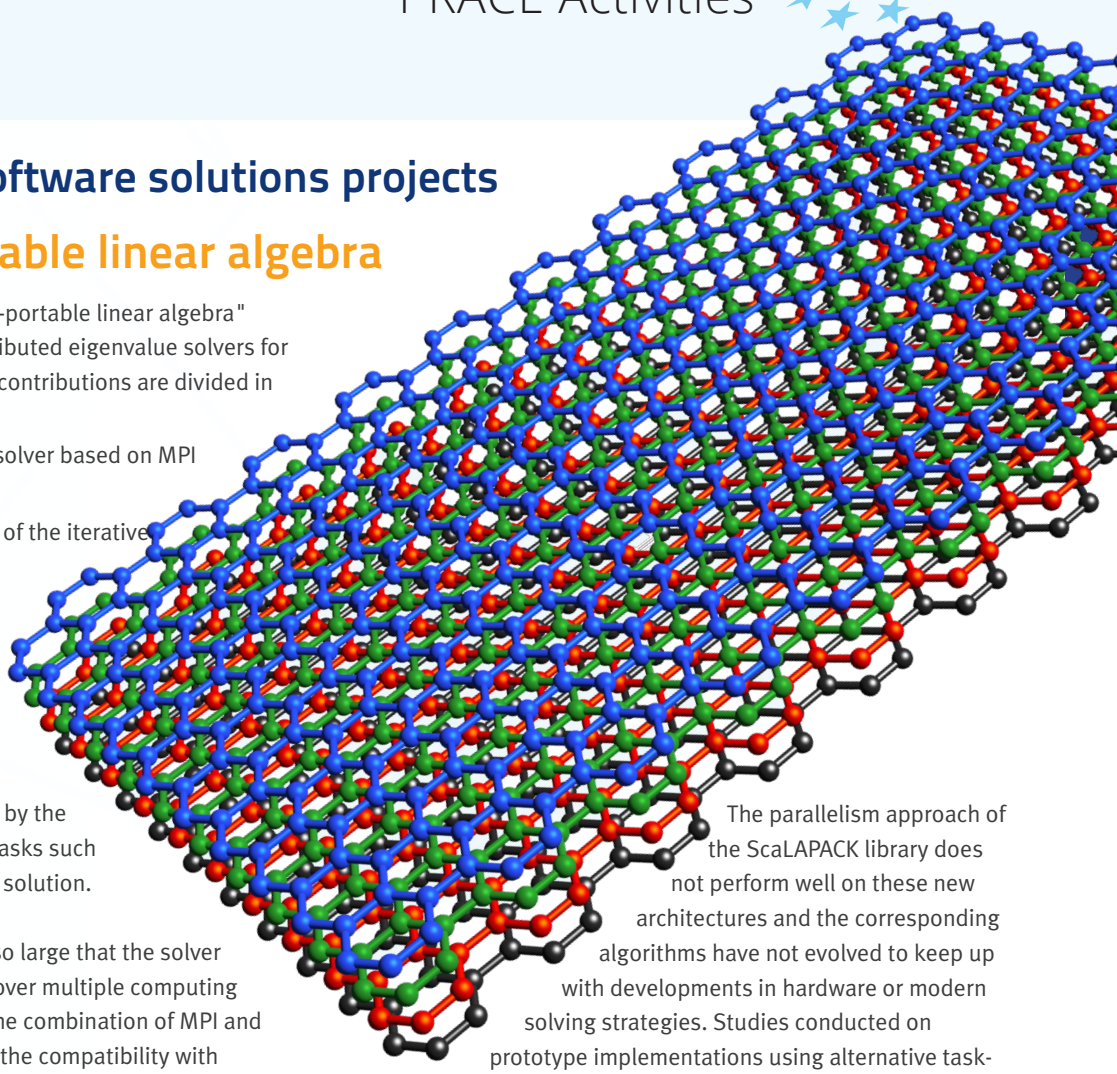

The eigenvalue solver is used in many scientific applications. A typical example is to be found in materials science (e.g. Quantum ESPRESSO, SIRIUS and CP2K) where the computational time is dominated by the execution of complex linear algebra tasks such as the Hermitian eigenvalue problem solution.

The size of the operands involved is so large that the solver of choice has to necessarily operate over multiple computing nodes. The reason for the choice of the combination of MPI and HPX is that MPI is required to ensure the compatibility with the existing applications. As a task-based framework, our choice naturally fell on HPX since it is a C++ Standard Library for concurrency and parallelism. HPX implements all of the corresponding facilities as defined by the C++ Standard and, additionally, functionalities proposed as part of the ongoing C++ standardisation process. Moreover, the combination of HPX libraries with scientific HPX based codes further reduces the need of synchronisation points.

In current applications, each invocation of a library routine introduces a barrier before the function call (all the rest of the computation on the current node has to be finished before calling the routine and all the computation of the routine has to be completed before returning). Through use of the C++ future concept, HPX further reduces the need of such barriers, since the runtime can track the task dependencies of the full application. Currently these applications use ScaLAPACK, which is the de-facto standard library for distributed linear algebra and was developed in 1995, when the architectures of supercomputers were based on nodes which had a single CPU core. Since then, the node architecture has evolved and multi-socketed nodes, multi-core CPUs, GPU accelerators and many-core CPUs have been introduced. The parallelism approach of the ScaLAPACK library does not perform well on these new architectures and the corresponding algorithms have not evolved to keep up with developments in hardware or modern solving strategies. Studies conducted on prototype implementations using alternative task-based runtime libraries show that significant performance improvements are possible. For example, using such task-based systems decreases the time to solution of the Cholesky factorisation by a factor of between 20-50%.

An alternative strategy in the development of an eigensolver is to leverage on well-known and well-established iterative algorithms such as subspace iteration. A modern example of such an algorithm has recently been implemented in the Chebyshev Accelerated Subspace iteration Eigensolver (ChASE) library. When tackling sequences of Hermitian eigenproblems, as they often appear in electronics structure codes, ChASE takes advantage of the distinctive features connecting adjacent problems in a sequence. This project aims to improve some of the key kernels of the ChASE library (e.g. skinny QR decomposition), and extend its support to generalised Hermitian eigenproblems without the necessity of Cholesky factorisations, and include partial singular value decomposition computations. Instrumental to the goals above is the need to simplify the transition of application software to new modern libraries. This will be achieved through the development of an interface which permits an easy transition from ScaLAPACK to different libraries. ★

NB-LIB – Performance-portable library for N-body force calculations at the exascale

A large number of scientific applications use particle interactions (e.g. molecular dynamics, Monte Carlo or multiscale simulations in life sciences or materials), and several smaller codes or combinations of codes have unique features. However, while computers have become more specialised, many codes are not optimised for GPUs or other accelerators and it is increasingly hard to achieve parallelisation. This will make these codes increasingly difficult to use on next-generation exascale systems.

The goal of the NonBonded-LIBrary (NB-LIB) project is to address these issues with a library of cutting-edge performance non-bonded interactions that also delivers excellent performance for other types of interactions, and has a parallelisation framework that can be used by arbitrary applications. This way, future acceleration, porting, and library features will benefit all applications. This is achieved by separating the lower-level parts of the GROMACS code into a C++ API of non-bonded and other force routines to enable reuse of widely used, highly portable, and performant HPC code. This will benefit applications in multiple domains and leverage exascale optimisation underway in GROMACS.

While NB-LIB targets performance with an eye towards the exascale, it also aims for ease of system setup and modularity of run modes. This will allow rapid prototyping of complex workflows without compromising on performance. NB-LIB allows users to programmatically define molecular topologies, facilitating building complex molecules with user-defined force-field parameters. In addition, NB-LIB is designed to allow modular simulation setup and execution, which can give the ability to run multiple simulations of the same system with different parameters in a highly performant way while not requiring any understanding of GROMACS internals. This will allow the user the flexibility to manage data and execution pipelines. ★

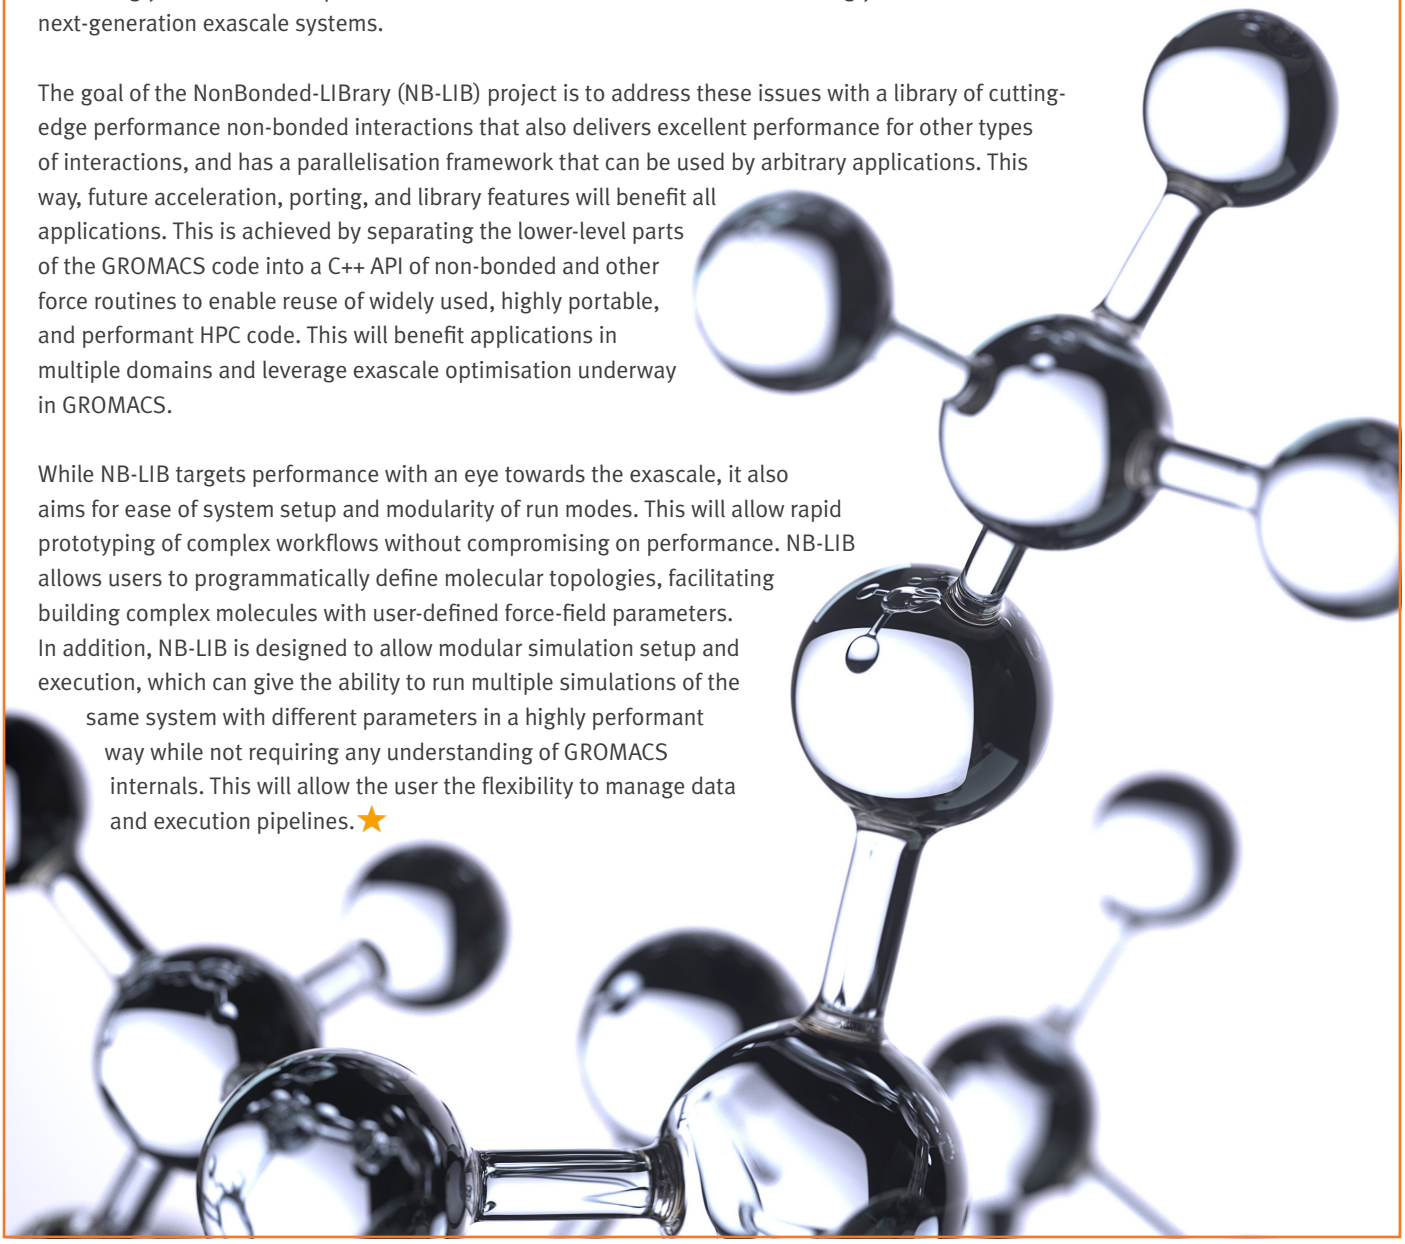

LoSync – Synchronisation reducing programming techniques and runtime support

The LoSync project aims to improve the scalability of applications by removing unnecessary synchronisation, serialisation and full realising opportunities, to overlap calculation and communication. To do this, modern features of well-standardised APIs have been used to ensure portability and relevance. These techniques include:

- **Using OpenMP/OmpSs tasks with data dependency clauses.** This not only includes expressing computation as tasks, but also communication, by wrapping MPI or GASPI library calls inside tasks. The Task-Aware MPI (TAMPI) and Task-Aware GASPI (TAGASPI) interoperability libraries developed by Barcelona Supercomputing Centre have been used to make this as efficient as possible.
- **MPI single-sided communication.** Recent developments in MPI libraries have significantly improved the performance of single-sided communication to the point where its benefits can be realised in real applications.
- **GASPI single-sided (put-notify) communication.** This is a lightweight alternative to MPI single-sided communication which interoperates well with MPI, and offers different synchronisation semantics which can help remove serialisation constraints.

The project is initially implementing these techniques in small kernels and mini-apps, with the aim of moving to key kernels of larger applications later in the project. In addition, development is being carried out on runtime library implementation to support this work.

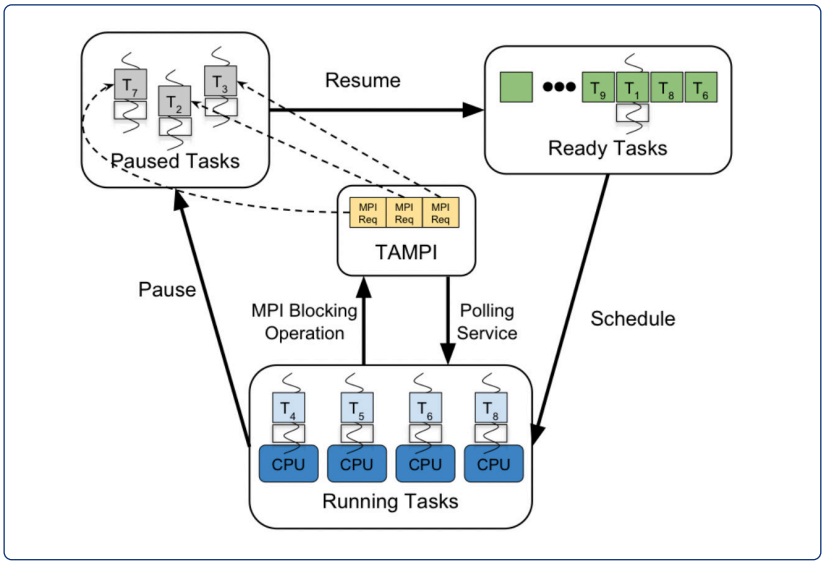

Diagram illustrating how the TAMPI library, used in the LoSync project, interacts with the OmpSs-2 runtime to integrate tasks and blocking MPI calls

This includes:

- Continuing the development of the TAMPI and TAGASPI interoperability libraries to support interaction of MPI and GASPI with OpenMP/OmpSs tasks with dependencies.
- Exploring extensions to the OpenMP tasking model to support task dependencies on external events, task-nesting, fine-grained dependencies, weak dependencies and early release of dependencies, to avoid artificial synchronisation and serialisation effects.
- Making use of performance analysis tools and techniques to identify optimisation targets in real applications where these techniques can be most beneficially applied.

The project is engaged with the relevant standards bodies for MPI, OpenMP and GASPI, to track relevant upcoming features, propose new features where necessary, and to produce prototype implementation to test these features. ★

MoPHA – Modernisation of plasma simulation codes for heterogeneous pre-exascale architectures

Numerical simulations are absolutely essential to address central open questions in plasma physics, from fusion energy to space weather. The understanding of the fundamental physical processes involved in plasma turbulence is of key importance, but simulations that provide such insights require enormous computational efforts. To be able to adequately and efficiently utilise the next generation of supercomputers in the pre-exascale and exascale era, plasma simulation codes need to be pushed to the next structural level with respect to scalability and portability for heterogeneous HPC architectures.

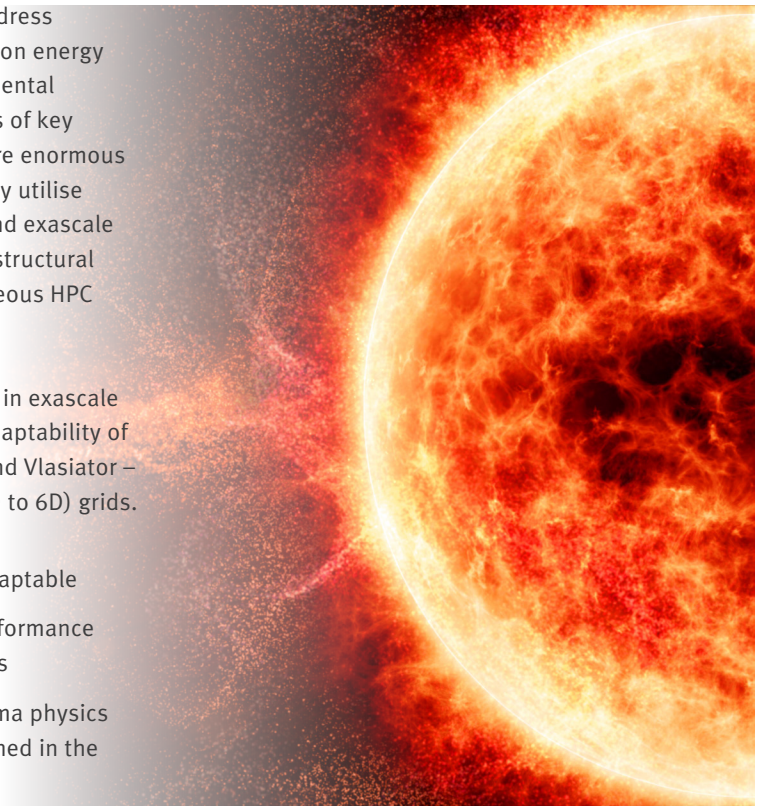

In order to prepare for the technical and scientific challenges in exascale computing, the project aims to improve the scalability and adaptability of three widely used plasma physics codes – ELMFIRE, GENE, and Vlasiator – which share the feature that they use higher-dimensional (3D to 6D) grids. The project consists of three main tasks:

- Refactor the codes to make them more accessible and adaptable
- Explore and implement task-based parallelisation for performance portability and scalability on heterogeneous architectures
- Knowledge transfer to encourage code re-use in the plasma physics community (as well as beyond) and to share lessons learned in the project.★

FEM/BEM based domain decomposition solvers

The aim of this project is to extend the domain decomposition library ESPRESO to support highly scalable solution of problems in complex domains using finite (and possibly boundary) element tearing and interconnecting (FETI/BETI) non-overlapping domain decomposition methods, thus enabling solution of large-scale sound scattering and harmonic analysis problems. The current implementation in ESPRESO focuses on the elasticity or the heat transfer problems for which it features excellent parallel and numerical scalabilities.

The implementation for complex problems (such as solution of the Helmholtz equation) will be based on the FETI-H/FETI-DPH (FETI-Helmholtz/FETI-Helmholtz dual-primal) approach where the regularisation is done using the complex interface mass matrix and the preconditioning is based on the plane wave deflation. The goal is to provide a modern, modular and portable code written in C++, parallelised on all possible levels, and capable of utilising the most powerful supercomputers. The development will take into account the heterogeneous nature of current and future supercomputers and when finished, the solver will be incorporated into a “solver as a service” platform at IT4Innovations, thus enabling scientists and engineers with zero experience in HPC to leverage the power of supercomputers.★

LyNcs – Linear algebra, Krylov-subspace methods, and multi-grid solvers for the discovery of new physics

Sparse linear systems arise naturally in a diverse range of applications in science and engineering. The sparsity is well suited for distributed computing architectures via domain decomposition of the problem space, which is why such sparse solvers have traditionally served as prototypes for HPC applications. However, on modern and upcoming architectures, especially in exascale machines, the efficient solution of such problems will require disruptive approaches to parallelism, beyond simply domain decomposition of the problem space.

The project, "Linear algebra, Krylov methods, and multi-grid API and library support for the discovery of new physics" (LyNcs), will address this challenge, pooling together software development efforts across Europe to provide communities with the next generation of parallel libraries for solving sparse linear systems at the exascale. LyNcs is led by CaSToRC of The Cyprus Institute, which is joining forces with partners from INRIA in France and LRZ in Germany.

Within LyNcs, cutting-edge sparse linear solver algorithms will be implemented, new Krylov and block Krylov solvers will be developed and prototyped, and existing parallel codes that implement a range of preconditioners for these solvers, such as multi-grid, will be optimised. The improvements proposed span all levels of the scientific application software stack, from the basic Sparse BLAS library to full-fledged simulation codes that include libraries like Fabulous and lattice QCD community solver libraries including QUDA and DD AMG. At the lowest level, the efficient sparse matrix support software librsb, its APIs and adapter libraries are being developed, pursuing a tighter integration with the aforementioned packages.

This library development and optimisation activity will be accompanied by a coherent effort in designing, implementing, documenting, and maintaining an Application Programmer Interface (API), enabling various scientific user communities to build full-fledged scientific applications on top of these libraries. In particular, communities are being directly supported by linking to major application codes running on PRACE Tier-0 systems, including tmLQCD for lattice QCD, HORSE for computational electromagnetics, and A-VCI for computational chemistry.★

PicKeX – Particle kinetic codes for exascale plasma simulation

Particle-in-cell (PIC) codes have now become an essential part of the modelling toolkit for many areas of plasma physics, whether for modelling particle acceleration with high-power lasers, or to understand detailed dynamics and transport processes near the edge – or scrape-off layer (SOL) – of magnetised plasma confinement vessels. This popular algorithm relies on a highly versatile, robust, finite-difference discretisation of the Vlasov equation for the particle distribution function in coordinate and velocity space. State-of-the-art three-dimensional PIC simulations involve up to 1012 particles on 106 cores, which generally requires careful management of the memory access and particle book-keeping to implement efficiently.

The PicKeX project focusses on two important community codes: EPOCH, a fully relativistic, electromagnetic model, and BIT1, a sophisticated PIC/Monte-Carlo model, both of which are heavily used in the laser-plasma and magnetic fusion communities respectively, but which both need heavy refactoring work to enable them to run effectively on future PRACE Tier-0 systems.

The project partners at the Jülich Supercomputing Centre and University of Ljubljana are exploring and implementing advanced algorithmic techniques such as task-based programming models, and dynamic load-balancing based on space-filling curves, to achieve this goal. A first public release of these performance-enhanced versions is planned for early 2020.★

GHEX – Performance-portable communication layer for grid applications

GHEX provides C++ library (and C/FORTRAN bindings) for performing halo-exchanges in applications based on domain-decomposition. Scientific applications are often defined on computational grids. Depending on the problem being solved, the grids can be Cartesian, with and without regular spacing or curvilinear, they can represent true 2D and 3D geometries, or they can be block-structured, in which each block exhibits regularities but their connections are general or fully unstructured meshes. To run at scale, these grids must be partitioned and distributed across address spaces, and to guarantee correctness, some information at the boundaries of the partitions, usually referred to as halos, need to be exchanged at the appropriate time during the execution.

Because of the topological differences between the individual grids and how they are represented in data structures, each individual application defines and stores the information about the halos differently. The consequence of this is that the halo-exchanges are typically custom developed for each application, and are usually based on standard communication primitives such as MPI. This makes applications hard to port to new architectures and programming paradigms, such as accelerator-based computers and/or with massive multi-threading and multi-tasking execution. GHEX strives to provide a clean and performance-portable implementation of halo-exchange operations for the most important grid types. This is achieved independently of the particular applications characteristics, even though priority will be given to grids used by the strategic scientific collaborators in the weather and climate domain, and in solar physics.

The ambitious objectives of GHEX are to enable high-resolution whole-globe climate simulations on grids with less than 1 km spatial resolution, and solar atmosphere simulations at unprecedented scales on next-generation computer architectures by being integrated in existing applications, since GHEX does not interfere with other runtime systems.

There are two key ingredients of GHEX: abstract communication end-points, and separation between the communication topology and the grid/mesh topology. The former allows the use of GHEX in multi-threaded and multi-tasking environments, while the latter allows multiple grid types to be handled by a common communication infrastructure. The separation of concerns between halo-exchange specification and implementation also allows GHEX to work seamlessly on different architectures, while the future-based interfaces allow for applications to take advantage of latency-hiding techniques.

GHEX has been developed in a collaboration between the Swiss National Supercomputing Centre (CSCS), and the University Centre for Information Technology (USIT, University of Oslo, Norway).★

Summer of HPC 2019

The Summer of HPC 2019 programme provided summer placements for 25 undergraduates and Master's students at HPC centres across Europe. The students had a memorable study week in Bologna, Italy, before embarking on a wide range of projects, from programming a Raspberry Pi-based compute cluster to deploying deep learning inference "on the edge". The award ceremony for the best projects was organised on the 6 November 2019 at the PRACE offices in Brussels.★

HPC skills for European competitiveness

A skilled, diverse HPC user/developer community is vital for improving the research and economic competitiveness of Europe. PRACE continues to maintain an extensive training programme aimed at researchers and application developers, from both academia and industry, on the productive use of world-class HPC resources. A significant part of the PRACE training programme consists of face-to-face courses with hands-on sessions, provided by a network of 14 PRACE Training Centres, complemented by seasonal schools and other courses organised on an on-demand basis with collaborators. More recent efforts have seen the offering of PRACE MOOCs and online learning materials that have broadened the scope and reach of PRACE training activities.

PRACE Training Centres network expanded

PRACE manages a network of PRACE Training Centres (PTCs) around Europe that delivered a total of 104 courses, represexnting 272 training days, over 2019. Courses are provided free-of-charge for European academic and industrial HPC users. A total of 2 618 participants have attended these courses. Feedback on the quality of PTC courses has also remained high (8.5 out of 10). In late 2019, the PTC network has expanded to a total of 14 members with the establishment of four new PTCs represented by project partners from Austria, Belgium, Slovenia and Sweden. PTCs offer a diverse array of courses from traditional HPC topics such as parallel programming (e.g. MPI, OpenMP), performance engineering tools and techniques, and accelerators, to courses that span the use of HPC technology for big data analysis (e.g. Apache Spark) as well as machine and deep learning. The PTCs often collaborate with other projects such as the European HPC Centres of Excellence. One such example is the past collaborations with the MaX and E-CAM Centres of Excellence in delivering joint training courses aimed at material science as well as molecular and atomistic simulations.★

PRACE MOOCs

PRACE´s MOOC offering grew into four courses during 2019. The latest course, “Python in High Performance Computing”, opened in September 2019 on the Future Learn learning platform. More information about PRACE MOOCs can be found at futurelearn.com/partners/prace ★

PRACE Seasonal Schools and the International HPC Summer School in 2019

PRACE-5IP and PRACE-6IP organised the following seasonal schools in 2019:

- The PRACE Winter School, “Introduction to Machine Learning for Scientists”, was held from 4 to 8 March 2019 at the University of Leuven, Belgium.
- The PRACE/BioExcel Seasonal School 2019, “HPC for Life Sciences”, took place from 10 to 13 June 2019 in Stockholm, Sweden. It was held at the PDC Center for High Performance Computing. The school was organised jointly between the PRACE-6IP project and BioExcel, the leading European Centre of Excellence for Computational Biomolecular Research.
- The PRACE Autumn School 2019, “Big Data and HPC”, was held from 17 to 20 September 2019 in Ljubljana, Slovenia and organised by the University of Ljubljana.

The 10th International HPC Summer School took place from 7 to 12 July 2019 in Kobe, Japan, hosted by the RIKEN Centre for Computational Science (RIKEN R-CCS). A total of 80 scholars from Canada, Europe, Japan and the USA participated in the event. This continues a long tradition of collaboration between PRACE in Europe and international partners from Canada (SciNet HPC Consortium), Japan (RIKEN R-CCS) and the USA (XSEDE).★

The 10th European workshop on HPC infrastructures

The European Workshop on HPC Infrastructures is an annual event that brings together specialists in HPC centre design and operation to discuss the latest trends and technologies for the infrastructure of supercomputing centres.

The 10th edition of the workshop was hosted by the Poznań Supercomputing and Networking Center (PSNC) and was held from 20 to 23 May 2019 in Poznań, Poland. CEA, CSC and LRZ together with BSC, PDC-KTH and PSNC collaborated in the programme committee of the workshop. The workshop attracted 74 participants from Europe, America, Australia, and Asia. Among participating parties, there were experts from the HPC facility management of PRACE-related datacentres, large government-funded non-European HPC sites, commercial datacentre providers, companies aiming to provide HPC relevant energy-efficient technologies for cooling and heat reuse, and experts from HPC equipment vendor labs. The workshop covered a wide range of topics relevant to HPC infrastructure management including cooling technologies, provision of electricity, overviews of standards and certification, management and monitoring tools, and future technologies. The workshop also included a guided tour of the PSNC computer centre and its facilities. A PRACE closed session was held as usual at the end of the workshop, providing an opportunity for exchanges between experts from the assembled PRACE sites.

The 10th European Workshop on HPC Infrastructures, held in May 2019

The next workshop, hosted by the Leibniz Supercomputing Centre (LRZ), will take place in Garching near Munich in May 2020. As usual, participation is by invitation only. However, people from PRACE sites interested in infrastructure related topics can contact the PRACE Work Package 5 (HPC Planning and Commissioning) leader to get an invitation.

The European HPC Ecosystem and the HPC in Europe Portal

During 2019, PRACE-6IP has engaged in the analysis of the European HPC ecosystem, in order to identify the services, competences, boundaries and overlaps of the various European HPC stakeholders. This analysis has included a survey where key players have self-declared their position regarding HPC policy, technology, computing services, training, support and research. The results of the analysis were presented in two networking events, first during the EuroHPC Summit Week 2019 in Poznań (Poland), and later during the SC19 conference in Denver (USA). Both events raised high interest and fruitful discussions among the 40 to 60 participants.

The outcome of the ecosystem analysis has been used to design the new HPC in Europe portal, a framework that will centralise the many European HPC services on offer in a single place. The portal will be structured around six major services (HPC access, training and events, support, applications, technology, and resources) and will have a special focus on the major target audiences for these HPC services (researchers, industry, students, communities, funding bodies and the general public). The advances in the portal will be presented in a webinar early 2020, and the final release of the portal is expected by the end of 2020 (to be hosted at www.hpc-in-europe.eu).

Operational and new services for the HPC ecosystem

The operational services of the ecosystem of Tier-0 and Tier-1 HPC machines in Europe are maintained daily and updated regularly to provide the best solutions to scientific and industrial communities. PRACE-6IP provides a day-to-day maintenance for eight Tier-0 and 18 Tier-1 systems with a helpdesk and persistent monitoring of resources and services. The HPC ecosystem evolves continuously with the regular integration of new systems and the natural decommission of obsolete systems, to guarantee high-tech solutions to researchers.

The operational services offered by PRACE are divided into several categories:

- Network services
- Data services
- Compute services
- AAA (authentication, authorisation, and accounting)
- Operational security services
- User services
- Monitoring services

In 2019, PRACE started a review of these services with the aim of providing more user friendly and forward-looking services for the future. A first step was the plan for four new services that have been tested already within PRACE-5IP and have been chosen for continuation within PRACE-6IP: urgent computing, in-situ visualisation, light-weight virtualisation services, and prototypal services for data analytics. Concerning urgent computing, by the end of 2019 PRACE-6IP had started deploying a scientific code called Tsunami developed within CoE CheeSE to PRACE Tier-1 systems that have strong GPU sub-clusters.

Regarding the in-situ visualisation, three use-cases related to catalyst software stacks were instrumented and tested: Migale, PLUMEFD, and OpenFoam. The light-weight virtualisation services are nearly mature (tests of different container platforms were performed), and PRACE-6IP is providing support for creating different containerised solutions (like Spark). Finally, the plan for future mandatory (core) data services Tensorflow, PyTorch, and Horovod, and for the optional data services Keras, Caffe, Caffe2, Scikitlearn, Spark, and Jupyter, was started. All these services are also considered as candidates for new future-oriented services for the updated PRACE Service Catalogue.

In the last months of 2019, the new collaboration with FENIX-RI suggested the development of a PRACE AAI (Authentication Authorisation Infrastructure) to allow easier access to resources and services. A plan will be ready before the end of PRACE-6IP. In addition to the collaboration with other EU projects, in the second half of 2019 PRACE-6IP began a collaboration with EU industries to evaluate new services towards open R&D.★

PRACE-6IP applications enabling and support

PRACE-6IP WP7 (applications enabling and support) continues to support European researchers to ensure that a broad range of important applications can effectively exploit current and future PRACE systems (both Tier-0 and Tier-1). Academic applications are supported via PRACE Preparatory Access (Tier-0) and the Tier-1 DECI (Distributed European Computing Initiative) programme, and industrial applications are supported through the SHAPE (SME HPC Adoption) Programme in Europe. WP7 also collaborates with the new PRACE High-Level Support Teams (HLSTs) to help maximise the scientific output of the Tier-0 systems. In addition to these applications enabling services, WP7 also supports European HPC research communities through the provision of benchmarks and best practice guides. Technical results are disseminated in white papers.

In PRACE-6IP, the major new focus areas are:

- DECI: the DECI programme has been re-introduced into the PRACE IP projects. DECI applicants will now benefit from enabling assistance, which can take the form of helping users get set up and running on the chosen machine, or in some cases could allow some limited work on users' codes. DECI Call 16 will close in 2020.
- SHAPE+ Pilot: SHAPE has already helped SMEs from 13 countries across Europe. To encourage applications from an even wider range of countries, a SHAPE+ pilot is being introduced to support the involvement of new countries.
- Collaboration with HLSTs: There are now HLSTs associated with each of the PRACE Tier-0 systems. WP7 will collaborate with these teams on projects that prepare applications for exascale, allowing more partners to contribute to the scientific output from Tier-0.★

N : PRACE AISBL
A : Troonstraat 98, rue du Trône
1050 Brussels, BELGIUM
P : +32 2 613 09 22
W : www.prace-ri-eu