

FROM PETASCALE TO EXASCALE

FROM PETASCALE TO EXASCALE



PARTNERSHIP FOR ADVANCED
COMPUTING IN EUROPE

Copyright notice

ALL RIGHTS RESERVED. This report contains material protected under International and Federal Copyright Laws and Treaties. Any unauthorized reprint or use of this material is prohibited. No part of this report may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system without express written permission from the author / publisher.

The copyrights of the content and images contained within this report remain with their original owner(s). The PRACE name and the PRACE logo are © PRACE.

All diagrams and most photographs are original to this book. See copyright notice for individual copyrights.

Printed by Bishops Printers Ltd, Spa House, Walton Road, Portsmouth, Hampshire, PO6 1TR, United Kingdom, Telephone +44 239 2334 900

Published by Insight Publishers Ltd, 12-13 King Square, Bristol, BS2 8JH, Telephone +44 117 2033 120, www.ipl.eu.com

Project Manager

Ellen Haggan

Designer

Mike Stafford

Additional designers

Thomas Hunt

Layout assistant

Simon Browne

Proofreader

Becky Freeman

Print manager

Paul Rumbold

PRACE

The PRACE Research Infrastructure is established as an international non-profit association with seat in Brussels and is named “Partnership for Advanced Computing in Europe aisbl”. It has 24 member countries (June 2012) whose representative organisations are creating a pan-European supercomputing infrastructure, providing access to computing and data management resources and services for large-scale scientific and engineering applications at the highest performance level.

www.prace-ri.eu

This work has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) in the PRACE-1IP Project under grant agreement n° RI-261557.

Lead Author

Martyn Guest – Cardiff University, UK

Secretary

Giovanni Erbacci – PRACE-1IP WP4 leader

Editor-in-chief

Marjolein Oorsprong – PRACE Communications Officer

Panel Chairs

Giovanni Aloisio – University of Salento and ENES-CMCC, Italy

Stefan Blügel – Forschungszentrum Jülich, Germany

Modesto Orozco – Institute for Research in Biomedicine, Spain

Philippe Ricoux – TOTAL, France

Andreas Schäfer – University of Regensburg, Germany

Scientific Case

Management Group

Turlough Downes – PRACE User Forum (Chair)

Richard Kenway – PRACE Scientific Steering Committee (Chair)

Thomas Lippert – PRACE-1IP Project Coordinator

Maria Ramalho – PRACE aisbl Acting Managing Director

Executive Summary

1 In 2005 and 2006, an international panel produced a White Paper entitled 'Scientific Case for Advanced Computing in Europe' that argued the case for High-Performance Computing (HPC) to support European competitiveness. The document was published by the 'HPC in Europe Taskforce' in January 2007. The initiative was instrumental to the establishment of PRACE – The Partnership for Advanced Computing in Europe – and the PRACE Research Infrastructure (PRACE-RI) in April 2010. This document represents the culmination of an initiative by PRACE to create an update of the Scientific Case and capture the current and expected future needs of the scientific communities. It has involved the PRACE Scientific Steering Committee (SSC), leading scientists from across all major user disciplines, and has been funded through the 1st Implementation Phase of PRACE – the PRACE-1IP project.

2 Five years after the publication of the scientific case, the HPC landscape in Europe has changed significantly. The PRACE Research Infrastructure is providing Tier-0 HPC services – large allocations of time on some of the most powerful computers in the world – to researchers in Europe, a global effort has been launched towards achieving exascale HPC by the end of this decade, and the importance of HPC in solving the socio-economic challenges and maintaining Europe's competitiveness has become even more evident.

3 The scope of this report is wide-ranging and captures the conclusions of five scientific areas, each derived from the work of an associated panel of experts. The five panels include those in the areas of: Weather, Climatology and solid Earth Sciences; Astrophysics, HEP and Plasma Physics; Materials Science, Chemistry and Nanoscience; Life Sciences and Medicine; and Engineering Sciences and Industrial Applications.

Communication from the EC

'High-Performance Computing: Europe's place in a Global Race' COM (2012) 45 Final. Brussels 15.02.2012

'The race for leadership in HPC systems is driven both by the need to address societal and scientific grand challenges more effectively, such as early detection and treatment of diseases like Alzheimer's, deciphering the human brain, forecasting climate evolution, or preventing and managing large-scale catastrophes, and by the needs of industry to innovate in products and services.'

'Industry has a dual role in high-end computing: firstly, supplying systems, technologies and software services for HPC; and secondly, using HPC to innovate in products, processes and services. Both are important in making Europe more competitive. Especially for SMEs, access to HPC, modelling, simulation, product prototyping services and consulting is important to remain competitive. This Action Plan advocates for a dual approach: strengthening both the industrial demand and supply of HPC.'

4 The position of HPC has evolved since 2007 from a technology crucial to the academic research community to a point where it is acknowledged as central in pursuing 'Europe's place in a Global Race'. The extract from the communication from the Commission to the European Parliament, in the panel, above, pays testimony to this position.

5 HPC is currently undergoing a major change as the next generation of computing systems ('exascale systems') is being developed for 2020. These new systems pose numerous challenges, from a 100-fold reduction of energy consumption to the development of programming models for computers that host millions of computing elements, while addressing the data challenge presented by the storage and integration of both observational and simulation/modelling data. These challenges cannot be met by mere extrapolation but require radical innovation in several computing technologies. This offers opportunities to industrial and academic players in the EU to reposition themselves in the field.

6 All of the panels contributing to this report are convinced that the competitiveness of European science and industry will be jeopardised if sufficiently capable computers are not made available, together with the associated infrastructure and skilled people necessary to maximise their exploitation. The panels have listed multiple areas at risk in concluding that access to high-performance computers in the exascale range is of the utmost importance. Thus, in aerospace, considerable changes in the development processes will lead to a significant reduction in development times while at the same time including more and more disciplines in the early design phases to find an overall optimum for the aircraft configuration. This will enable the European aircraft industry to keep a leading role in worldwide competition, facing both an old challenge, i.e. competing with the USA, and a new, rapidly emerging one – keeping an innovation advantage over China. However, while aerospace can afford its own HPC provision, it may not have the capability to exploit exascale if similar systems are not available to academia for training and software development. In a similar vein, the lack of high-performance computers appropriate for life sciences research will displace R&D activities to the USA, China or Japan, putting European leadership in this field at risk.

7 Providing scientists and engineers with ongoing access to computers of leadership class must be recognised as an essential strategic priority in Europe: there is a compelling need for a continued European commitment to exploit the most powerful computers. Such resources are likely to remain extremely expensive and require significant expertise to procure, deploy and utilise efficiently; some fields even require research for specialised and optimised hardware. The panel stresses that these resources should continue to be reserved for the most exigent computational tasks of high potential value. It is clear that the computational resource pyramid must remain persistent and compelling at all levels, including national centres, access and data grids. The active involvement of the European Community along with appropriate member states remains critical in maintaining a world-leading supercomputer infrastructure in the European ecosystem. Europe must foster excellence and cooperation in order to gain the full benefits of exascale computing for science, engineering and industry in the European Research Area.

“ Providing scientists and engineers with ongoing access to computers of leadership class must be recognised as an essential strategic priority in Europe: there is a compelling need for a continued European commitment to exploit the most powerful computers. ”

Key statements

By way of summary, we present below key statements from the Commission and from our thematic panels that emphasise the essential role of a sustainable top-level infrastructure



'High-Performance Computing (HPC) is critical for industries that rely on precision and speed, such as automotive and aviation, and the health sector. Access to rapid simulations carried out by ever-improving supercomputers can be the difference between life and death; between new jobs and profits or bankruptcy.'

Communication from the EC: 'High-Performance Computing: Europe's place in a Global Race' COM (2012) 45 Final. Brussels 15.02.2012

Weather, Climatology and solid Earth Sciences (WCES)

"In the last decade, our understanding of **Climate Change** has increased, as has the societal need for pull through to advice and policy. However, whilst there is great confidence in the fact that climate change is happening, there remain uncertainties. While there is uncertainty about the levels of greenhouse gas emissions and aerosols likely to be emitted, there are perhaps more significantly uncertainties on the degree of warming and the likely impacts. Increasing the capability and comprehensiveness of 'whole Earth system' models that represent in ever-increasing realism and detail, scenarios for our future climate is the only way to reduce these latter uncertainties."

"A programme of provision of Leadership-class computational resources will make it increasingly possible in **solid Earth Sciences** to address

the issues of resolution, complexity, duration, confidence and certainty. Challenges have significant scientific and social implications, playing today a central role in natural hazard mitigation, treaty verification for nuclear weapons, and increased discovery of economically recoverable petroleum resources and monitoring of waste disposal."

"There is a fundamental need in **Oceanography and Marine Forecasting** to build and efficiently operate the most accurate ocean models. Improved understanding of ocean circulation and biogeochemistry is critical to assess properly climate variability and future climate change and related impacts on e.g. ocean acidification, coastal sea level, marine life, polar sea-ice cover, etc."

Astrophysics, High-Energy Physics and Plasma Physics

"Astrophysics, high-energy physics and plasma physics have, in recent years, shared a dramatic change in the role of theory for scientific discovery. In all three fields, new experiments become ever more costly, require increasingly long time scales and aim at the investigation of more and more subtle effects. Consequently, theory is faced with two types of demands: Precision of theory predictions has to be increased to the point that it is better than the experimental one. This is a most demanding requirement as the latter can be expected to increase by further orders of magnitude until 2020. In parallel, the need to explore model spaces of much larger extent than previously investigated also became apparent. For example:

- In *astrophysics* determination of the nature of dark energy and dark matter requires a detailed comparison of predictions from large classes of cosmological models with data from the new satellites and ground based detectors which will be deployed until 2020.
- In *high-energy physics*, one of the tasks is to explore many possible extensions of the Standard Model to such a degree, that even minute deviations between experimental data and Standard Model predictions can serve as smoking guns for a specific realization of New Physics.
- In *plasma physics*, one of the tasks is to understand the physics observed at ITER at such a high level that substantially more efficient fusion reactors could be reliably designed based on theoretical simulations which explore a large range of options."

Materials Science, Chemistry and Nanoscience

*"Computational materials science, chemistry and nanoscience is concerned with the complex interplay of the myriads of atoms in a solid or a liquid, thereby producing a continuous stream of new and unexpected phenomena and forms of matter, characterised by an extreme range of length, time, energy, entropy and entanglement scales. The target of this science is to design materials ranging from level of a single atom up to the macroscopic scale, and unravel phenomena and design processes from electronic reaction times in the femtosecond range up to geological periods. Computational material science, chemistry and nanoscience stand in close interaction with the neighbouring disciplines of biology and medicine, as well as the geo-sciences, and affect wide fields of the engineering sciences. A large and diverse computational community that views as critical assets the conceptualization, development and implementation of algorithms and tools for cutting edge HPC will achieve this goal. These tools are used to great benefit in other communities such as **Medicine and Life Sciences** and **Engineering Science and Industrial Applications**."*

*"The advance from petascale to exascale computing will change the paradigm of **computational materials science and chemistry**. The move to petascale is broadening this paradigm – to an integrated engine that determines the pace in a design continuum from the discovery of a fundamental physical effect, a process, a molecule or a material, to materials design, systems engineering, processing and manufacturing activities, and finally to the deployment in technology, where multiple scientific disciplines converge. Exascale computing will significantly accelerate the innovation, availability and deployment of advanced materials and chemical agents and foster the development of new devices. These developments will profoundly affect society and the quality of life, through new capabilities in dealing with the great challenges of knowledge and information, sustained welfare, clean energy, health, etc."*

Life Sciences and Medicine

“In **Life Sciences and Medicine**, Eflop/s capabilities will allow the use of more accurate formalisms (more accurate energy calculations for example) and enable molecular simulation for high-throughput applications (e.g. study of larger number of systems). Molecular simulation is a key tool for computer-aided drug design. The lack of high performance computers appropriate for this research will displace R&D activities to the USA, China or Japan, putting European leadership in this field at risk. Appropriate exascale resources could revolutionize the simulation of biomolecules, allowing molecular simulators to decipher the atomistic clues to the functioning of living organisms.

Biomedical Simulation will reduce costs, time to market and animal experimentation. In the medium to long term, simulation will have a major impact on public health, providing insights into the cause of diseases and allowing the development of new diagnostic tools and treatments. It is expected that understanding the basic mechanisms of cognition, memory, perception etc. will allow the development of completely new forms of energy efficient computation and robotics. The potential long-term social and economic impact is immense.”



“While Exaflop machines are essential for specific areas of Life Sciences (e.g. brain simulation), and higher computational power will enable significantly increased accuracy for current modelling studies, some extremely important fields in Life Science will be mainly limited by throughput and data management.”

Engineering Science and Industrial Applications


“All of us experience the effects of HPC in our day-to-day lives, although in many cases we are unaware of that impact. We travel in cars and aeroplanes designed using modelling and simulation applications run on HPC systems so that they are efficient and safe. HPC is essential for ensuring that our energy needs are met. Finding and recovering fossil fuels require engineering analysis that only HPC can deliver. Nuclear power generation also relies heavily on HPC to ensure that it is safe and reliable. In the coming years HPC will have an even greater impact as more products and services come to rely on it.”

“The **automotive industry** is actively pursuing important goals that need exaflop computing capability or greater. Examples include (i) vehicles that will operate for 250,000 kilometres on average without the need for repair; this would provide substantial savings for automotive companies by enabling the vehicles to operate through the end of the typical warranty period at minimal cost to the automakers, and (ii) Insurance companies require full-body crash analysis that includes simulation of soft tissue damage; today’s “crash dummies” are inadequate for this purpose.”

“The impact of **computer simulation in aircraft design** has been significant and continues to grow. Numerical simulation allows the development of highly optimized

designs and reduced development risks and costs. Boeing, for example, exploited HPC in order to reduce drastically the number of real prototypes from 77 physical prototype wings for the 757 aircraft, to only 11 prototype wings for the 787 “Dreamliner” plane. HPC usage saved the company billions of Euros.”

“In addition to the automotive and aeronautics examples above, many areas within the engineering sciences – waves seismic equation inversion, engine combustion (chemical and multiphysics combustion), and turbulence demand highly scalable or so called “hero applications” to deliver long-term social and economic impact.”



Key Recommendations

In pointing to the compelling need for a continued European commitment to exploit leadership class computers, the scientific panels have considered the infrastructure requirements that must underpin this commitment, and present their considerations as part of the review of computational needs. This considers both the vital components of the computational infrastructure, and the user support functions that must be provided to realize the full benefit of that infrastructure. This review has led to a set of key recommendations deemed vital in shaping the future provision of resources, recommendations that are justified in full in the Scientific Case and outlined below

1

The need for HPC infrastructure at the European level

The scientific progress that has been achieved using HPC since the “Scientific Case for Advanced Computing in Europe” was published in 2007, the growing range of disciplines that now depend on HPC, and the technical challenges of exascale architectures make a compelling case for continued investment in HPC at the European level. **Europe should continue to provide a world-leading HPC infrastructure to scientists in academia and industry, for research that cannot be done any other way, through peer review based solely on excellence.** This infrastructure should also address the need for centres to test the maturity of future Exascale codes and to validate HPC exascale software ecosystem components developed in the EU or elsewhere.

2

Leadership and management

The development of Europe’s HPC infrastructure, its operation and access mechanisms must be driven by the needs of science, industry and society to conduct world-leading research. This public-sector investment must be a source of innovation at the leading edge of technology development and this requires user-centric governance. **Leadership and management of HPC infrastructure at the Europe level should be a partnership between users and providers.**

3

A long-term commitment to Europe-level HPC

Major experiments depend on HPC for analysis and interpretation of data, including simulation of models to try to match observation to theory, and support research programmes extending over 10-20 year timeframes. Some applications require access to stable hardware and system software for 3-5 years. Data typically need to be accessed over long periods and require a persistent infrastructure. Investment in new software must realise benefits over at least 10 years, with the lifetime of major software packages being substantially longer. **A commitment to Europe-level HPC infrastructure over several decades is required to provide researchers with a planning horizon of 10-20 years and a rolling 5-year specific technology upgrade roadmap.**

4

Algorithms, software and tools

Most applications targeting Tier-0 machines require some degree of rewriting to expose more parallelism and many face severe strong-scaling challenges if they are effectively to progress to exascale, as is demanded by their science goals. There is an on-going need for support for software maintenance, tools to manage and optimise workflows across the infrastructure, and visualisation. Support for the development and

There is grave concern about HPC skills shortages across all research areas and, particularly, in industry. The need is for people with both domain and computing expertise. The problems are both insufficient supply and low retention, because of poor career development opportunities for those supporting academic research.

Europe's long-term competitiveness depends on people with skills to exploit its HPC infrastructure. It must provide on-going training programmes, to keep pace with the rapid evolution of the science, methods and technologies, and must put in place more attractive career structures for software developers to retain their skills in universities and associated institutions.

Notes





