H2020-FETHPC-2014

Coordination of the HPC strategy

EXDCI

European eXtreme Data and Computing Initiative

Grant Agreement Number: FETHPC-671558

D3.1
First set of recommendations and reports toward applications

Final

Version: 1.0
Author(s): S. Requena (GENCI) et al.
Date: 25.10.2016

Project and Deliverable Information Sheet

EXDCI Project	**Project Ref. №: FETHPC-671558**	
	Project Title: Europea eXtreme Data and Computing Initiative	
	Project Web Site: http://www.exdci.eu	
	Deliverable ID: **D3.**1	
	Deliverable Nature: Report	
	Dissemination Level: PU	**Contractual Date of Delivery:** 30 / 11 / 2016
		Actual Date of Delivery: 30 / 11 / 2016
	EC Project Officer: Beatrice MARQUEZ-GARRIDO	

* - The dissemination level are indicated as follows: **PU** – Public, **CO** – Confidential, only for members of the consortium (including the Commission Services) **CL** – Classified, as referred to in Commission Decision 2991/844/EC.

Document Control Sheet

Document	**Title: First set of recommendations and reports toward applications**	
	ID: **D3.**1	
	Version: 1.0	**Status:** Final
	Available at:	http://www.exdci.eu
	Software Tool:	Microsoft Word 2013
	File(s):	EXDCI-Deliverable-D3.1.docx
Authorship	**Written by:**	S. Requena (GENCI)
	Contributors:	Y. Fournier (EDF), H. Pitsch (RWTH Aachen), G. Aloisio (CCMC-Univ Salento), JC. André (JCA Consultance), S. Brun (CEA), S. Krieg (JSC), R. Apostolov (KTH-BioExcel), P. Coveney (UCL-CompBioMED) and D. Wright (UCL / CompBioMED) with the support of all the experts of the different working groups and JM. Alimi (OBSP), Y. Tourbier (Renault), P. Gibbon (JSC) and G. Staffelbach (CERFACS)
	Reviewed by:	D. Henty (EPCC) and F. Bodin (Irisa)
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	03/10/2016	Draft	Initial version
0.2	22/10/2016	Draft	Incorporated inputs from WG3.4
0.4	04/11/2016	Draft	Incorporated inputs from WG3.2 and WG3.3
0.5	10/11/2016	Draft	For internal review

1.0	19/11/2016	Final version	Inputs and comments from the reviewers

Document Keywords

Keywords:	PRACE, Research Infrastructure, applications requirements, Center of Excellence, use cases, breakthroughs and recommendations, HPDA, BDVA

Table of Contents

List of Figures

References and Applicable Documents

[1] http://www.exdci.eu
[2] http://www.prace-ri.eu
[3] http://www.etp4hpc.eu
[4] PRACE Scientific Case: http://www.prace-ri.eu/prace-the-scientific-case-for-hpc
[5] EESI2 reports: http://www.eesi-project.eu/ressources/documentation/#eesi2-deliverables
[6] https://www.eurolab4hpc.eu
[7] http://www.bdva.eu
[8] http://wssspe.researchcomputing.org.uk
[9] https://github.com/PRUNER/archer
[10] https://www.rd-alliance.org
[11] https://www.eudat.eu
[12] http://esgf.llnl.gov
[13] http://science.sciencemag.org/content/331/6018/700
[14] http://extremecomputing.labworks.org/nuclearphysics/index.stm
[15] http://www.deus-consortium.org/a-propos/dark-energy-universe-simulation-full-universe-run
[16] https://zenodo.org/collection/user-emmc
[17] http://www.fz-juelich.de/ias/jsc/EN/Expertise/High-Q-Club/_node.html
[18] http://dx.doi.org/10.1080/10618562.2016.1218481
[19] Yi Wang, Sébastien Chevrot, Vadim Monteiller, Dimitri Komatitsch, Frédéric Mouthereau, Gianreto Manatschal, Matthieu Sylvander, Jordi Diaz, Mario Ruiz, Franck Grimaud, Sébastien Benahmed, Hélène Pauchet and Roland Martin, The deep roots of the western Pyrenees revealed by full waveform inversion of teleseismic P waves, Geology, vol. 44(6)
[20] http://link.springer.com/article/10.1007/s11432-016-5588-7
[21] G.R. Mirams, C.J. Arthurs, M.O. Bernabeu, R. Bordas, J. Cooper, A. Corrias, Y. Davit, S.-J. Dunn, A.G. Fletcher, D.G. Harvey, M.E. Marsh, J.M. Osborne, P. Pathmanathan, J. Pitt-Francis, J. Southern, N. Zemzemi, D.J. Gavaghan. Chaste: An open source C++ library for computational physiology and biology. PLoS Comput. Biol. 9(3): e1002970, 2013. doi: 10.1371/journal.pcbi.1002970.
[22] Vázquez, M., Arís, R., Houzeaux, G., Aubry, R., Villar, P., Garcia-Barnés, J., Gil, D. and Carreras, F. (2011), A massively parallel computational electrophysiology model of the heart. Int. J. Numer. Meth. Biomed. Engng., 27: 1911–1929. doi: 10.1002/cnm.1443
[23] M.D. Mazzeo, P.V. Coveney, HemeLB: A high performance parallel Lattice-Boltzmann code for large scale fluid flow in complex geometries, Computer Physics Communications, Volume 178, Issue 12, 15 June 2008, Pages 894-914, ISSN 0010-4655, http://dx.doi.org/10.1016/j.cpc.2008.02.013.
[24] http://www.palabos.org
[25] http://www.continuity.ucsd.edu
[26] J. R. Hussan, P. J. Hunter, P. A. Gladding, N. Greenberg, R. Christie, A. Wu, H. Sorby, and J. D. Thomas, Bioinformatics. 2015 Apr 15; 31(8): 1331–1333. doi: 10.1093/bioinformatics/btu809
[27] http://simvascular.github.io
[28] https://blogs.scientificamerican.com/video-of-the-week/a-stunning-and-groundbreaking-simulation-of-the-human-heart/ (accessed November 2016)

[29] A. Randles, E. W. Draeger, T. Oppelstrup, L. Krauss and J. A. Gunnels, "Massively parallel models of the human circulatory system", Proceedings: SC '15 Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, Article No. 1, ACM New York, NY, USA 2015 ISBN: 978-1-4503-3723-6

[30] Advancing Regulatory Science at FDA (Online). http://www.fda.gov/regulatoryscience

[31] US Food and Drug Administration. Guidance for industry and FDA staff—Nonclinical engineering tests and recommended labelling for intravascular stents and associated delivery systems, 2010. http://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance

[32] Avicenna Coordination Support Action, "in silico Clinical Trials: How Computer Simulation will Transform the Biomedical Industry", http://avicenna-isct.org/wp-content/uploads/2016/01/AvicennaRoadmapPDF-27-01-16.pdf (accessed Nov 2016)

[33] M. Eisenstein, "Big data: The power of petabytes", Nature 527, S2–S4 (05 November 2015) doi:10.1038/527S2a

[34] http://biorxiv.org/content/early/2016/06/19/059717

[35] Horowitz S et al. Determining crystal structures through crowdsourcing and coursework. Nat Commun. 2016 Sep

[36] Marx V. Genomics in 3D and 4D. Nat Methods. 2016 Sep

[37] Norrby M et al. Molecular Rift: Virtual Reality for Drug Designers. J Chem Inf Model. 2015 Nov

[38] Hopf TA et al. Sequence co-evolution gives 3D contacts and structures of protein complexes. Elife. 2014 Sep

[39] Toth-Petroczy A et al. Structured States of Disordered Proteins from Genomic Sequences. Cell. 2016 Sep

[40] http://a16z.com/2016/10/25/bio-cs-machinelearning-medicine

[41] http://dx.doi.org/10.3233/978-1-61499-621-7-439

[42] http://dx.doi.org/10.1063/1.4952638

[43] EU Intertwine Project : www.intertwine-project.eu

[44] EU NEXTGenIO Project : www.nextgenio.eu

[45] http://predictioncenter.org

List of Acronyms and Abbreviations

AISBL	Association Internationale Sans But Lucratif (International Non-for-Profit Association)
BDEC	Big Data and Extreme-scale Computing
BDVA	Big Data Value Association
CAE	Computer-aided engineering
CASP	Critical Assessment of protein Structure Prediction
CFD	Computational Fluid Dynamics
CoE	Centres of Excellence for Computing Applications
cPPP	contractual Public-Private Partnership
DoE	Design of Experiments
DoW	Description of Work

DSL	Domain Specific Language
EC	European Commission
ECMWF	European Centre for Medium-range Weather Forecasts
EESI	European Exascale Software Initiative projects
ENES	European Network for Earth System modelling
EPOS	European Plate Observing System
EsD	Extreme Scale Demonstrators
ESGF	Earth System Grid Federation
EU	European Union
FET	Future and Emerging Technologies
FP7	Framework Programme 7
FPGA	Field Programmable Gate Array
GDP	Growth Domestic Product
GPI	Global Address Space Programming Interface
GPU	Graphical Processing Unit
H2020	Horizon 2020 – The EC Research and Innovation Programme in Europe
HPC	High Performance Computing
IDC	International Data Corporation
IESP	International Exascale Software Project
INVG	Istituto Nazionale di Geofisica e Vulcanologia (National Institute of Geophysics and Volcanology)
IO	Input/Output operations
IoT	Internet of Things
IPCC	Intergovernmental Panel on Climate Change
ISV	Independent Software Vendor
ITER	International Thermonuclear Experimental Reactor
IT	Information Technology
KPI	Key-Performance Indicator
MHD	Magneto HydroDynamics
MIC	Intel Manycore product line
NWP	Numerical Weather Prediction
OS	Operating System
POC	Proof of Concept
QCD	QuantumChromoDynamics
R&D	Research and Development
RDA	Research Data Alliance
ROI	Return On Investment
SHAPE	SME HPC Adoption Programme in Europe
SHS	Social and Historical Sciences
SME	Small and Medium Enterprise
SRA	Strategic Research Agenda

SWOT	Strengths, Weaknesses, Opportunities and Trends
TRL	Technology Readiness Level
UQ	Uncertainties Quantification
WCES	Weather, Climate and solid Earth Sciences
WG	Working Group
WP	Work Package
WP2	EXDCI work package 2 “Technological ecosystem and roadmap toward extreme and pervasive data and computing »
WP3	EXDCI work package 3 “Applications roadmaps toward Exascale”
WP4	EXDCI work package 4 “Transversal vision and strategic prospective”

Executive Summary

EXDCI's goal is to "coordinate the development and implementation of a common strategy for the European HPC Ecosystem in order to achieve its global competitiveness within the Horizon 2020 Programme"[1].

The objectives of the EXDCI work package 3 (WP3) "Applications roadmaps toward Exascale" are to provide updated roadmaps of needs and expectations of scientific applications as inputs to be used for the update of the PRACE Scientific Case in 2017 [4]. This document will support PRACE in the deployment of its (Pre)Exascale pan European HPC research infrastructure.

The roadmapping activity of WP3 will rely on documents produced by previous European projects like EESI and EESI2 (European Exascale Software Initiative, [5]), by PRACE on its second edition of the PRACE Scientific Case, inputs from the newly created nine European Centers of Excellence (CoE) and from experts involved inside the following 4 working groups (WG):

- WG3.1: Industrial and engineering applications
- WG3.2: Weather, Climate and solid Earth Sciences
- WG3.3: Fundamental sciences
- WG3.4: Life Sciences and Health

This deliverable summarizes the first important results, findings, conclusions and recommendations of these four WP3 working group reports. It also provides global and transverse recommendations.

It is important to highlight that applications are the meaning of HPC and big data, Europe is developing a major fraction of the applications used in the world (70% in the field of chemistry for example) and Europe is the biggest producer of data (thanks to a high level of equipment of large scale scientific instruments).

The nature of science is changing – new scientific discoveries and socio-economical innovation are emerging from the analysis of large amounts of complex data generated by high-throughput scientific instruments (e.g. sequencers, synchrotrons, scanners, microscopes), observational systems (e.g. telescopes, satellites, network of sensors, IoT, etc.), extreme-scale computing (for both capability based large scale 3D simulations as well as ensemble or coupled multiscale/multiphysics simulations), and the public World Wide Web.

This now opens the door to an accelerated convergence between HPC/HTC and big data in order to be able to acquire, analyse, value and make available larger amount of refined data in a trustable, open (in some case) and user-friendly way. In that goal, a recent communication from the European Commission in April 2016 advocated for an ambitious plan for a European Cloud Initiative providing innovative data Cloud services (European Open Science Cloud) through a unified European-wide Data Infrastructure to researchers from academia and industry as well as public entities.

On the other hand, the Exascale area, which is about to rise in the coming years, will bring a lot of breakthroughs mainly due to heavy energy management concerns on the IT side. If applications want to continue to benefit from improvements both in terms of sustained compute and data performance, major efforts of code modernisation will be needed. This shift will be

[1] as defined in EXDCI's DoA

successful only if it bridges even more scientific communities with HPC, applied mathematics and big data communities through co-design activities and if communities are able at the same time to address the hardware/software prototypes.

Finally, as HPC and advanced numerical simulations are now crucial for all the scientific and industrial domains, access to persistent and complementary national and world-class European HPC/data infrastructures (such as PRACE) and services (training, user support etc.) is a key element for sustaining the European competitiveness.

In this context, WP3 experts representing industrial and scientific communities are highlighting the following aspects:

- As some (pre)Exascale architectures are already known (and for some of them available), communities have started to rewrite/modernise their applications with the support of recently created European Centers of Excellence (CoE). Activities include for example optimizations for heterogeneous/manycore architectures or deep memory/storage hierarchies. This is facilitated by the existence of standards like OpenMP (and to a lesser extend OpenACC) which ease and secure the shift to novel architectures in a context where applications last for decades when HPC architectures could last years;
- Some communities started also to explore novel programming approaches like Domain Specific Languages (DSL) relying on smart underlying system software layers, for abstracting the complexity of future HPC architectures;
- As data movement will be a strong concern regarding energy efficiency and overall performance, some communities started to engage with using in-situ and in-transit post processing techniques or using novel big data approaches but these efforts need to be now accelerated;
- Future (converged) HPC and data e-infrastructures will need to support both capability and capacity simulations. Communities engaged major efforts in improving scalability of the applications (a lot of examples are provided in this report but also in previous EESI2 reports) beyond 100 000 cores, but the major fraction of workloads will consist on ensemble or coupled multi-scale/multi-physics applications.
  This leads to reformulating previous EESI2 recommendations toward the development of scalable meshing tools, ultra-scalable solvers, unified and scalable frameworks for code coupling and for the support of uncertainty quantification and optimisation.

Here are the global recommendations:

- Developing in-situ/in-transit post processing tools with extended machine/deep learning (DL/ML) features in order to detect pertinent structures in massive amount of data. This could foster closer relations between the HPC community and the Big Data community where Europe owns strong skills (in both academia and industry). WP3 is proposing to first launch joint call for proposal, bridging teams in domain science, HPC, applied math and DL/ML for concretely on 10 to 12 lighthouse projects assessing the potential of such approaches on different cases.
- Developing features of HPC resources managers for the coupling with large scale scientific/medical instruments and the development of urgent computing services by HPC infrastructures by supporting co-scheduling of resources, smart (application-based) checkpoint/restart or complex workflows, etc.

- Extend the European Centers of Excellence with the following ones:
  - Engineering and industrial applications: One of the strengths of Europe is the

number of industrial users of HPC having already strong HPC internal roadmaps (in Oil & Gas, Aeronautics, Automotive, Energy etc.) or using HPC facilities (PRACE in 4 years worked with more than 50 groups from large companies to SMEs). In order to democratise even more HPC and advanced numerical simulation this CoE could federate the European ecosystem in order to reach a critical mass for the support of European engineering applications. This CoE could start first with the support of CFD and turbulent applications and then expand its activity to other engineering domains. Sustainability of this CoE could be ensured by specific services provided to industry like user support of (open source) software, licensing or specific tailored developments.

- (Open Source) Software sustainability: European and national funding agencies are funding a lot of research projects leading to promising applications/tools which are not industrialised and subsequently used for many reasons. This is especially crucial for industry which requires software to be highly industrialised and provided with long-term and reactive user-support and training.

  This transverse CoE could interface on one side with existing vertical CoE and H2020 projects and on the other side with e-infrastructures like PRACE, in order to industrialise, promote and provide long-term support of scientific software once it reaches a given level of readiness. Sustainability of this CoE could be ensured by specific services provided to industry like user support of (open source) software or tailored industrialisation of in-house software.

- High Performance Data Analytics: due to the convergence of HPC and big data detailed previously, it is appearing mandatory to provide specific services to scientific and industrial communities toward High Performance Data Analytics (HPDA). In the same way as the POP (Performance Optimisation and Productivity) CoE is already acting for performance analysis and optimisation, this transverse CoE could work on assessing/auditing the needs of "client" communities in terms of data analytics/management and then provide solutions through Proof of Concept (PoC) and training actions based on standard approaches when possible.
  This CoE could be established in collaboration with the Big Data Value Association (BDVA) in order to foster synergies between the HPC community and the Big Data community.

1 Introduction

This document describes a first draft consolidation of the activity of the 4 working groups inside WP3: “Application roadmap toward Exascale”. The goal of this work package is to setup and manage four working groups of applicative experts to:

- investigate the application drivers for Exascale computing (meaning using machines delivering Exaflops class performance but also able to flow Exabytes of data);
- identify needs and expectations of scientific applications in the Excascale time-frame
- evaluate the economic dimension and impact on European competitiveness;
- build a European vision and suggest a roadmap as inputs for elaboration of the third version of the PRACE Scientific Case at the end of 2017.

The WP3 is chaired by Stephane Requena (GENCI) and it is organized into four representative working groups (WG), each one managed by a Chair and a Vice-Chair:

- WG 3.1: Industrial and Engineering Applications

  Chair: Yvan Fournier (EDF) and Vice Chair: Heinz Pitsch (Univ. Aachen)
  Industrial sectors like aeronautics, automotive, oil & gas, energy or finance to name a few are using heavily HPC internally or through access to research infrastructures like PRACE.
  This task will investigate and provide updated roadmaps required to support capacity and capability simulations, development of new tools dealing with load-balancing in industrial geometries with automatic/adaptive meshing, couplers, scalable solvers, frameworks for optimisation and uncertainty quantification, etc. Industrial experts will participate coming from large companies as well as SMEs and academia. A special support will be also provided by the PRACE Industrial Advisory Committee
- WG 3.2: Weather, Climate and solid Earth Sciences

  Chair: Giovanni Aloisio (Univ. Salento-CCMC) and Vice Chair: Jean Claude André (JCA Consultance)
  Starting from the activity carried out during EESI projects or into the PRACE Scientific Case, and in relation with the future CoE on "Weather and climate", this group will address issues like new climate models, dynamical cores, and future couplers, as well as knowledge compression, increased concurrency, complex workflows for ensemble simulations, uncertainty quantification, scientific data management, …
  Contributors will come from EsiWACE, ENES and from ECMWF and the national meteorological services it comprises. Other contributions are expected in the fields of air quality and oceanography, as well as from INGV and EPOS for the solid Earth Science part.
- WG 3.3: Fundamental Sciences (Chemistry, Physics)

  Chair: Allan Sacha Brun (CEA) and Vice Chair: Stefan Krieg (JSC)
  The field of fundamental sciences is covering various fields from nuclear physics, fusion to cosmology or molecular physics and quantum chemistry; this implies a diverse field of algorithmic and methodological approaches. This workgroup will further consider and update the challenges in fundamental sciences and will have a close look onto computational requirements and data management for Exascale computing.
- WG 3.4: Life science and Health

Chair: Rossen Apostolov (KTH, representing the BioExcel CoE) and Vice Chair: Peter Coveney (UCL, representing the CompBioMed CoE)
Since there is a wide range of projects requiring Exascale performance in Life Sciences, this working group will organize the panel of experts into four main areas: Systems Biology, Genomics, Neuro-informatics, Molecular Simulations and Biomedical Simulations.
The group will analyse the impact of Life Sciences Exascale computing in the society, economy and welfare state.

A fifth working group (D3.5, chaired by S. Requena/GENCI and by representatives of PRACE Scientific Steering Committee, SSC) is in charge of the global coordination of WP3 and the consolidation of the different roadmaps in order to provide inputs to the PRACE SSC for the elaboration of the third version of the PRACE Scientific Case.
During the first year of EXDCI a total of 40 additional scientific experts, from academia and industry, representing 8 European countries have been progressively enrolled by the 4 applications working groups. The full list of experts is provided as annex 1 of this document.

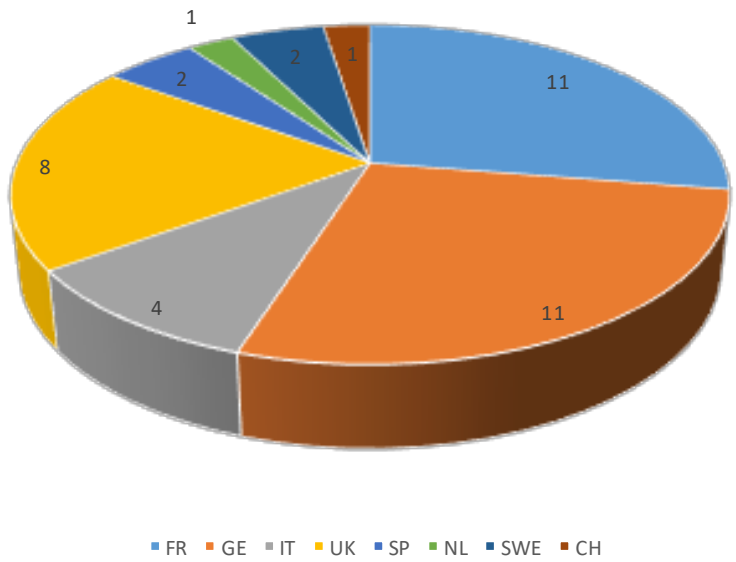

Figure 1 - distribution of experts across the 4 working groups

Each of the four working groups recruited progressively around 8 to 12 experts in their respective areas and had several meetings or telco during the first year. WP3 also had several interactions with others work packages of EXDCI including WP2 “Technological ecosystem and roadmap toward extreme and pervasive data and computing » and WP4 « Transversal vision and strategic prospective ». In that sense WP3 participated to a first EXDCI technical meeting organised on 21 and 22 September 2016 in Barcelona where interlock sessions between experts from WP2, WP3 and WP4 allowed to the exchange of current needs, visions and ways of collaboration for establishing joint roadmaps taking into account both applications, technology and cross-cutting requirements.

As stated into the EXDCI DoA, WP3 worked to establish connection for the newly created Center of Excellence (CoE) and up to now:

- ESiWACE[2] the Centre of Excellence in Simulation of Weather and Climate in Europe is represented into WG3.2;
- EoCoE[3] the Center of Excellence in the field of Energy, E-CAM[4] the Center of Excellence in the field of molecular dynamics, quantum dynamics and electronics structure modelling and NoMAD[5], a Center of Excellence in the field of novel materials discovery are involved into WG3.3;
- As planned in the DoA, BioExcel[6] (a Center of Excellence in biomolecular simulation) and CompBioMED[7] (a Center of Excellence in biomedicine, recently created in September 2016) are now chairing WG3.4.

WP3 will continue to extend relationships with the others 3 CoE during the second year and will enrol additional experts. Following the outcome of the first EXDCI technical meeting (see D4.4 Report on first technical workshop), relations with the BDVA ETP [7] in the field of big data and with the Eurolab4HPC H2020 project will be enhanced in order to address issues related to big data and long-term applications roadmaps.

2 Scientific (and industrial) challenges and roadmaps

2.1 Industrial and engineering applications

Working group 3.1 focuses on industrial and engineering applications. The experts are Yvan Fournier (EDF, CFD, focus on HPC, pre, post and coupling aspects), Heinz Pitsch (RWTH Aachen University, CFD and combustion, focus on engine applications) Klaus Adams (TU Munich, complex flows, aerodynamics), Philippe Ricoux (Total, reservoir modelling), Norbert Kroll (DLR, aeronautics and external aerodynamics), and Denis Veynante (ECP and CNRS, combustion and turbulence). CFD is very well represented in this group, while other areas less so, but applications at least at EDF and Total are quite varied, and feedback from other areas was requested.

2.1.1 Challenges on Aeronautics and Aerospace

Aeronautics and Aerospace are industrial fields with a traditional affinity to HPC computing for more than 20 years. Tools developed by industry and in academic institutions in separate branches as well as in joint research/industrial projects on national and EU levels in the last decades have enable the community to reduce the number of experiments through predictive large-scale simulations in the fields of CFD, structural and failure analysis, material research, heat management and propulsion (combustion), etc.

One of the main benefits for the application of HPC in the design process is the increased number of case studies possible on a numerical level hand in hand with an increase in safety and reliability, and faster turn-around times in design cycles assuring competitive aerospace products coming out of the EU. An even bigger step has been undertaken in optimizing

2 https://www.esiwace.eu
3 http://www.eocoe.eu
4 https://www.e-cam2020.eu
5 https://nomad-coe.eu
6 http://bioexcel.eu
7 http://www.compbiomed.eu

efficiency (i.e. fuel consumption) as well as increasing eco-friendliness (i.e. noise of air traffic) and the health impact reduction of traffic and combustion.

These developments are strongly driven by the EC or worldwide regulations. As an example, the latest ACARE (Advisory Council for Aeronautics Research in Europe) (and the Flightpath 2050 Europe's Vision for Aviation) reports are expecting for 2050 a CO_2 reduction by 75%, a NO_x reduction by 90%, by 65% a reduction in perceived aircraft noise (compared to a typical aircraft in the year 2000), a 80% reduction in the accident rate, while air companies are expecting reliable and less kerosene hungry planes with a 3x increase of the traffic and 99% of flights within 15' of schedule!

To meet the challenges of future aircraft transportation (Greening the Aircraft), it is indispensable to be able to flight-test a virtual aircraft with all its multi-disciplinary interactions in a computer environment and to compile all of the data required for development and certification with guaranteed accuracy in a reduced time frame.

Non-linear, interdisciplinary challenges not accessible to experimental investigations due to cost or realizability issues can now be challenged with HPC methods and systems. The continuing effort to increase HPC methods have rendered methods as RANS (Reynolds Average Numerical Simulation) as everyday tools in industrial application in the past. With increasing HPC capabilities on a EU level through PRACE institutions as well as companies are capable of considering unsteady codes and issues very early in their development process manifesting the need for Exascale resources with data sizes increasing by at least two orders of magnitude.

Codes used become more and more complex as they increasingly become coupled to jointly investigate issues which have been investigated separately in the past (e.g. combustion, soot formation, fluid-structure interaction, propulsion systems and noise generation for airplane propulsion systems). Research efforts have shown that entire airplanes or propulsion units can be tackled with HPC codes and resources. The tools have to become more precise and more reliable to support EU institutions and industrial partners to be able to step forward in competitive and sensitive fields.

DNS (Direct Numerical Simulation) and LES (Large Eddy Simulation) simulations with high-order accurate numerical methods calculating at billions of nodes and producing TB of result data are state-of-art in high-resolution turbulence or combustion simulations. They still face limitations of resolution at solid boundaries, in complex configurations and/or areas with multiple phases and multiple scales that can range of several orders of magnitude. Further improvement of low-level (e.g. RANS) tools can be expected for the productions environment through the advancement of non-linear interdisciplinary research in aeronautical and aerospace sciences.

Such Digital Aircraft vision will have an impact on the use of HPC by all the aeronautics value chain, requiring the availability of leading edge HPC resources in both capacity (or farming mode) as well as capability. This will also require **strong software development efforts** for:

- Increasing the scalability of individual simulation codes such as CFD, structure, acoustics, … by working on new scalable numerical solvers;
- Introduction of non HPC technologies (CAD) in multi-disciplinary optimisation;
- Developing automatic grid generation tools for handling complex geometries;
- Next generation of couplers for handling multi scale and multi physics heterogeneous applications;
- Next generation of uncertainties/optimisation framework;
- Handling and visualisation of big data.

The CFD roadmap in Figure 2, provided by NASA two years ago, and widely adopted by the entire aeronautics/aerospace community is still valid and demonstrate **needs up to 2030 with Zetascale HPC facilities**. This roadmap highlights more precisely some issues raised by the experts of the working group toward the evolutions of the physical modelling (RANS to cover the area where the boundary layer (BL) is too thin for LES, hybrid RANS/LES for a good compromise in performance and quality, or full LES), the development of new algorithms and meshing tools, the growing need of analytics and in-situ technologies and finally the importance of the development of multi-disciplinary analysis and optimisation (MDAO) with uncertainties quantification.

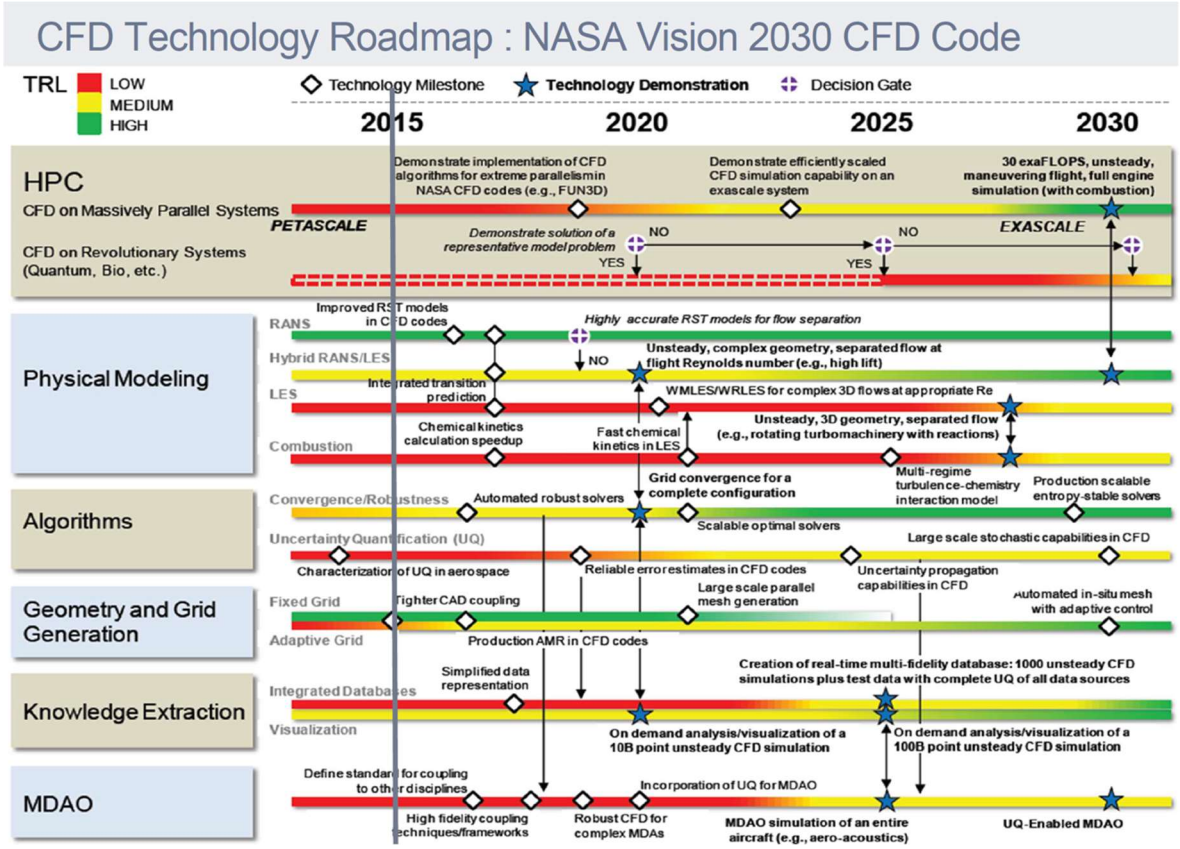

Figure 2 - roadmap of CFD needs of the aeronautics industry

2.1.2 Challenges in the automotive industry

In the automotive industry, reducing the need for prototypes has long been a goal. Avoiding late changes and fixes can help avoid delays of the market launch, and requiring tooling changes also cost several million dollars. Virtual design saves both cost and time, and improving the trust at each level of that design allows reducing the need for validation on physical models.

Finding the right balance between design and engineering priorities is essential to building a successful vehicle, and for some automotive companies, improving visualization software to a photorealistic level is a large part of attaining that goal, as it allows design teams to work more tightly with engineering.

The importance of such simulation-drive design is true especially for the external aerodynamics, where using HPC simulations, the effect of design changes on aerodynamics

may be computed and visualized with a high degree of fidelity, allowing both design and engineering teams better insight into the effect of changes.

According to announcements by Jaguar Land Rover (JLR) and its partner Exa Corp., it should be possible to **eliminate the need to build physical prototypes of a new vehicle by 2020**. This will require going much beyond aerodynamics, up to achieving full vehicle verification. In 2014, JLR reports using 36 million CPU hours, the equivalent of 7000 wind tunnel tests.

To drive this virtual engineering capability, several HPC principles are applied, namely:

- the HPC resource is strategic
- one global scheduling system for the ecosystem
  - allowing fairshare, and specific handling of urgent priority
  - everything can run anywhere based on resources

The last point may seem in conflict with choices from other domains, where minimum movement of data is sought, though quite practical for users. A remark on convergence is given in the "Job Managers" section of this document.

Another strong domain using HPC in the automotive is crash simulations where:

- The model size is growing continuously, quicker than the performance increase of individual CPU;
- The carmakers want to reduce the delay of crash simulation in order to accelerate the design process. They use more and more cores for each crash simulation, but the scalability of that kind of simulation (explicit solvers) is not perfect, HPC growth is more than linear;
- The carmakers will have to validate a lot of configurations, to give to each customer the real performance of his car (and not only the data for the best-selling version like in current regulation). The physical validation of all possible version is not possible, numerical simulation becomes mandatory;
- The design process needs to use systematic optimization with bigger and bigger studies. The current optimization algorithms last several crash simulations per design parameter, especially with robust optimization and uncertainty propagation.

All these items give an exponential growth of HPC usage, but one breakthrough in optimization field can counteract this evolution: model reduction. The main optimization technology used with big finite elements simulation is Design Of Experiments (DOE). One can find a lot of DOE based algorithms but they all have a common attribute: they use the crash simulation as a black-box (use very few information from each simulation), and then last a lot of crash simulation. Model reduction uses all available data from crash simulation and will be seen as one of the most contributors of HPC increase (horizon: 5 years).

As a consequence, companies like Renault-Nissan who used in 2013 PRACE resources for a massive crash optimisation studies, **started to re-engage into the use of HPC** and expose now aggressive roadmap in both crash simulations, aerodynamics and combustion like in the following figure (values per year are proportions of use of the HPC resources wrt 2013).

Figure 3 - HPC needs of Renault

Finally, the latest and one of the most HPC requiring challenges in the future will be **connected cars and autonomous driving**. This has been popularised by the Google car but now ALL the vendors and their Tier1 (like Dephi, Valeo, Continental and Faurecia, …) are developing these technologies.

The connected car will bridge all the on-board IT infrastructure of the car to the outside in order to provide traffic simulation, interaction with other cars, interaction with any other service for the safety or the leisure of the driver or the passengers.

The autonomous driving will come gradually with hand-on, hands-off, eyes-off and at the end even mind-off. The complexity between the first step: hand-on to the last one: mind-off is increasing by a factor of more than 1000 with obviously extremely low failure rate per hour expected, close to 10^{-9}. This is something which is already present in the aeronautics but automotive applications will require even more control and reliability since a pilot almost always has time to take back the commands while it could be impossible for a driver.

These challenges are illustrating a strong convergence between HPC and embedded systems with severe issues on real time image/signal processing, data management, security, reliability and insurance regulations.

As an example of volume of data to ingest and analyse in almost real time, airplanes generate approximately 2,5 billion Terabyte of data each year from the sensors installed in the engines and upcoming self-driving cars will generate 2 Petabyte of data every year.

2.1.3 Challenges in the Oil and Gas

Despite the drop of the price of the crude oil barrel in 2014, oil & gas companies are pushed to invest in new technologies including HPC to explore and exploit new ultra-deep offshore, or non-conventional oil fields (including shales gases) by using more and more precise seismic algorithms and more accurate reservoir modelling methods. **Oil & Gas became the second largest profitable market for HPC** (just after Finance) with according to IDC an increase of the CAGR of 9.2% between 2012 and 2017. Most of the Oil & Gas major companies or contractors **owns multi Petascale HPC resources** including TOTAL (with Pangea a 6.7 PFlops system), ENI, BP (with a new 2PFlops manycore system in 2016), Exxon, Shell and Chevron for the majors or Petrobras, Woodside Energy (Australia), Petrochina (China) for the National Oil Companies or PGS (Petroleum Geo-Services with a primary system of 5 PFlops which has been completed end of 2016 with a a new 2.8 PFlops system), CGG-Veritas or Schlumberger for the major contractors have also distributed HPC datacenters with accumulated performance beyond 10 PFlops.

Such companies developed a strong roadmap towards Exascale primarily for the development of **efficient and accurate novel seismic processing methods for exploration** (such as Reverse Time Migration, Separated Wavefield Imaging (SWIM), and Wave Equation Reflectivity inversion) and production (4D seismic coupled to reservoir modelling, uncertainties quantification, multi-scale modelling from the pore to the reservoir scale).

In this domain and due to the progress of the process of seismic acquisition (more sensors, more streamers, higher frequencies of acquisition, multi-component data, …) it is expected to see an increase of needs by one order of magnitude before 2020. Also, the integration of more physics, more complex approximations and more iterations of the methods will lead to an increase between one to 3 orders of magnitude.

The following roadmap provided by TOTAL in 2015 is still valid and representative of the Oil & Gas industry requirements in seismic exploration:

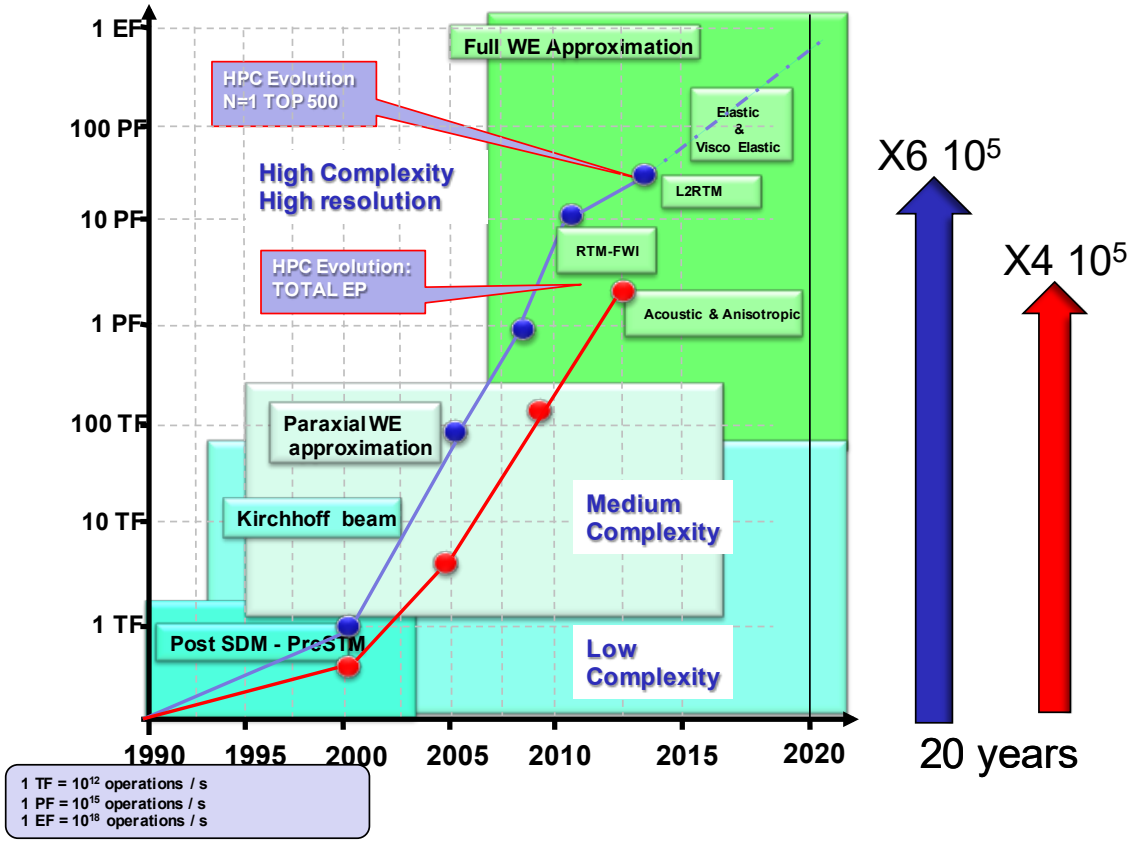

Figure 4 - evolution of the needs of TOTAL in seismic modelling

Also, some companies **expanded the use of HPC** to scale out basin and reservoir modelling or more recently to molecular dynamics and chemistry applied to processes of refining or global optimization of plants including multi-disciplinary interactions, e.g. simulating full plant lifetime and control systems, emulation of operations and risk analysis.

Finally, HPC in conjunction with high performance data analysis is starting to become also crucial in performing **real time simulations** by coupling simulations and inputs from network of sensors or seismic acquisitions when drilling a well or operating it in order to be able to better steer the process and by consequence to optimise operations and reduce costs and risks.

2.1.4 Challenges in Power generation and nuclear plants

In this industrial domain, the objectives are multiple: first improvement of safety and efficiency of the facilities (especially nuclear plants), and second optimization of maintenance operation and life span. In this field, physical experimentation, for example with nuclear plants, can be not only impractical but also unsafe. **Computer simulation**, in both the design and operational stages, **is therefore indispensable**.

In the thermal hydraulics the improvement of efficiency may typically involve mainly steady CFD calculations on complex geometries, while improvement and verification of safety may involve long transient calculations on slightly less complex geometries, and less well resolved meshes.

This will require HPC for the study of flow-induced loads (to minimize vibration and wear through fretting in components such as fuel assemblies), flow-induced deformation and de-nucleate boiling avoidance in pressure water reactors (PWR) cores, and the use of detailed simulations designed to verify and increase safety.

In order to validate such models, it will be mandatory to run quasi-DNS type calculations on subsets of the calculation domain may be necessary.

Immediate needs in the field require (unstructured) meshes in the billion-cell range. Studies in the near future could easily an order of magnitude more. **Such studies will require in the range of several hundred thousand (or millions) of cores during several weeks (involving multiple computation restarts)**.

Developers of several CFD tools seem to reach the agreement that due to CFL and time step constraints, DNS computations at very high Reynolds numbers will never be possible given the current hardware trends, as latency becomes the major issue. So LES, RANS, and approaches bridging scales will remain necessary, and increased computing power may be used more to increase geometric details and domain sizes, which requires progress not only in the codes themselves, but in CAD, meshing, and post-processing aspects.

Improving CAD, meshing, or otherwise means to handle geometric details (through some forms of coupling for example) is essential to make these studies viable.

Also, to reach very fine resolutions on detailed meshes, to reduce data movement, features such as **"in-situ" mesh refinement** or adaptation and improved coupling methods may be of great importance.

In the nuclear industry, materials science also has a very high importance, due in great part to safety demonstration requirements for materials in a neutronic flux. These studies bridge many domains, from the ab-initio to molecular dynamics and up to finite elements, **and use a variety of multiphysics codes at all scales**, from such well-known tools such as VASP to more specialized tools, for example the EDF-CEA CRESCENDO cluster dynamics code for the mesoscale dislocation dynamics parts. Such tools have traditionally required as larger or slightly larger resources than even CFD. Most approaches end themselves well to parallelisation and HPC, although a few (e.g kinetic Monte-Carlo) may be difficult to parallelise.

With other methods of power generation, such as wind or water turbines, study of deposition on blades (especially for marine turbines) can be important, requiring coupling of CFD with other models.

2.1.5 Challenges in Process Engineering

In process engineering, CFD also has a major role, in a typology of flows not familiar to those in aeronautics or nuclear engineering. Here, interaction with specific material behaviour models, coupling with lower scales, and moving meshes or domains with rotating parts bring

additional challenges, and progress in both CFD and coupling tools and methods can have a major impact.

2.1.6 Challenges in Combustion

Most practical combustion processes happen in the turbulent combustion regime. This requires high spatial resolution in simulations in order to resolve very small length scales which are of importance for turbulence but also accurate numerical schemes. Additionally, pollutant and emission formation can only be predicted accurately, if many different chemical species are considered. This results in trillions of degrees of freedoms in current simulations and consequently high demand for computing time and data storage.

Efficient analytical methods are required to develop models from highly-resolved DNS data for cheaper but also predictive simulation types such as Large Eddy Simulations (LES).

Additionally, technical combustion devices are often very complex, not entirely accessible by measurements.

Today, with the increase in computational resources, numerical simulation applications range from gas turbines to car engines and even rocket engine. Subdomain simulations allow for tackling highly complex systems such as combustion chamber/multi-stage turbines (see fig. 3) with increased modelling accuracy using multi-physics code coupling for fluid/solid interaction for both structure dynamics and thermal response for example.

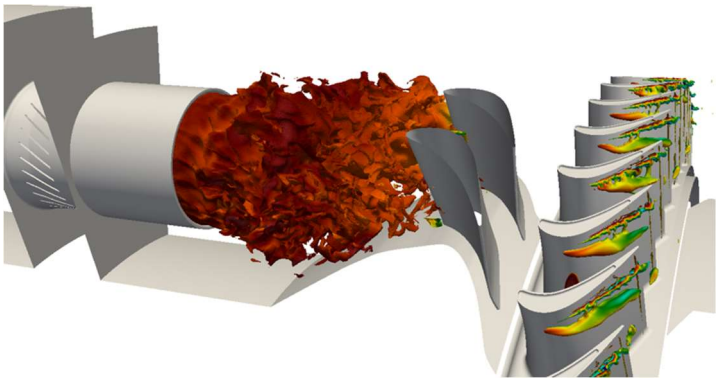

Figure 5 - combustion chamber/turbine LES of the factor test rig (F. Duchaine et al, JARAHPC'16)

The advent of Exascale computing foreshadows the birth of **fully integrated virtual engine design in the next decade** with great leaps forward. Fully resolved complex chemistry in real engines **is still out of reach from current HPC platforms**, however analytically reduced chemistry models allow today to predict global tendencies for emissions predictions with reasonable computational cost increase (x3 compared to simple models).

Additionally, by its nature, combustion involves highly intermittent phenomena which entails require long time-resolved simulations. This translates to generating large amounts of data where data mining will become crucial but also long simulation times where reliability and high level parallelisation and code computational efficiency will be critical. This will translate to current and future efforts to modernise legacy code to adapt to new and innovative hardware and software changes such as short vector vectorisation or massively manycore systems.

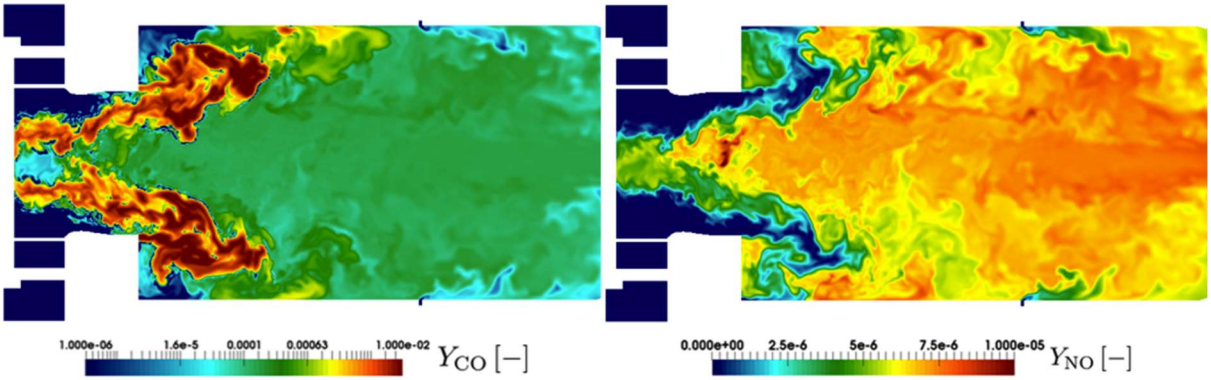

Figure 6 - Instantaneous emission mass fraction - LES of the SGT 100 methane air burner operated at DRL. (Stopper et al CF 2013) T. Jaravel A. Felden (T. Jaravel et al PROCI 2016)

2.1.7 Common requirements and recommendations

All of the sections above concern mature technologies (power plants, airplanes, …), which have been developed for a very long time. Further increasing efficiencies or performances will require to exploit potential of nowadays mostly neglected effects (e.g. unsteadiness, even smaller scales, full simulations of coupled systems, …). One consequence will be long runtimes for unsteady simulations and thus large amount of data. Furthermore, coupling of several reduced sub-simulations with different resolutions need to be developed and consistent numerical methods.

Evolving Hardware and software environments

As computing power increases faster than memory bandwidth and capacity, current hardware trends lead an increasing number of cores, with decreasing memory per core. This causes issues with single-level parallelism (such as pure MPI), as some fixed overheads (both memory and computational) are multiplied by the number of cores. These issues are mitigated with “lighter” programming models (such as OpenMP), at the cost of increasing programming effort.

Code complexity

General-purpose codes are quite complex, as they consist not only of numerical kernels, but also many other subsets, relative to many tasks, such as I/O, coupling, mesh manipulation, etc. To reduce disk I/O, it is often most efficient to link most functions into a single executable (often through an extensive set of libraries), but this means the executables are larger and more complex. Separating stages into distinct modules is not an easier path, as interfaces may then need to be very well specified and API changes avoided.

As industrial codes evolve over time, many specific models are added, and they become a repository for large amounts of “know-how”. As such, the volume of code representing even the core features of those tools may become important, and research on new algorithms is hampered by the volume of code. Still, extracting mini-apps is often quite difficult in itself, as removing too many features may lead to not being able to run at all, at least not on cases of significant complexity or size. This would tend to argue the case for simpler, more specialized codes, which could be coupled together, but mastering several tools and coupling them efficiently only shifts the difficulties from the developer to the user, which is unacceptable for most user of “industrial” tools.

So though evolution of large codes is possible, it is often slower than evolution of hardware, making it more and more difficult to “catch up”.

Co-processing, co-visualisation and in-situ visualization

The EESI2 project already strongly recommended in 2015 working on in-situ visualization. There has been significant progress in this area, with some of the codes used of developed by WG 3.1 experts now implementing some in-situ capabilities. For example, EDF's *Code_Saturne* CFD tool leverages ParaView's Catalyst library, while RWTH Aachen's CIAO code uses VisIt's libsim library. Additionally, interface libraries such as FZ Jülich's JUSITU have been developed in order to help developers enabling in-situ visualization in their codes without knowing all details about the underlying libraries (Catalyst/libsim).

In-situ visualization seems less present in non-CFD tools, and lacks maturity (at least in the case of Catalyst), but maturity has improved significantly over the last 2 years.

Some national-level projects within the EU already aim to foster progress in this area, such as the French AVIDO project (EDF, Total, UPMC-LIP6, Inria, funded by BPI France). EU-level projects could help improve the impact and dissemination of similar projects throughout Europe and the launch of first call of proposals (FETHPC-1-2017) targeting in-situ and in-transit processing is a good signal.

Other efforts around this subject go beyond just visualization but include other post-processing and analytics aspects. Some non-EU efforts such as LBNL Scalable Environment for Scientific Exploration In Situ (SENSEI) project or the ANL's GLEAN infrastructure in the US should also be watched or experimented with.

Another important issue is about the analysis of the resulting data itself. With increasing data set sizes, the analysis becomes more and more difficult and commonly developed analysis methods such as deep learning or machine learning could be used to support this process.

Software Engineering

The issues relative to software engineering are not specific to this workgroup, and other areas are expected to have similar needs.

Software Sustainability

It can be noted that some initiatives in Europe relative to software engineering in HPC areas are already active, such as WSSSPE [8] which tries to reach beyond Europe, having organized workshops at both during SC'14 and SC'15.

Most experiences shared in these workshops seem to indicate training and best practice are essential, as many researchers are domain specialists, whose initial training and education is quite poor in the software development aspects, though more than 90% of researchers surveyed use research software.

Actually, in engineering domains, codes developed for the industry may suffer somewhat less from developer turnover than some research codes, when at least part of development teams has a lower turnover than students or post-doc, but training and good practice remain essential.

Another aspect in this context would be the development of very general software libraries, not limited to certain applications (such as AMR libs). HPC experts could be employed in the long-run, while scientists do not need to know all details of the software and consequently the training effort is reduced. Due to increasing complexity, this might be necessary in any case.

Programming Languages and smart runtimes

Most engineering codes use Fortran, C, C++, with MPI, OpenMP and libraries. Will other tools reach critical mass for long-lived tools? Should the HPC community try to develop its own languages, or should it try to push mainstream languages in more HPC-friendly directions?

There are critical mass issues here, even considering international collaboration not restricted to the EU.

Vendor "behaviour" here seems short-sighted also, with major hardware vendors as well as research teams each advocating their own approach or solution. It seems some of the main reasons for the success of MPI (good collaboration on a common standard making code development much easier, with competition on the implementation, not the standard) have been forgotten, with risks of short-term technology lock-in often slowing adoption ("wait and see") more than helping it. This may be in part due to the greater complexity of today's hardware, and lack of maturity of proposed runtimes, but from an industrial code developer's point of view, the lack of a clear direction here are often an impediment to starting code modernization efforts. Fewer, better supported, and more interoperable proposed solutions would certainly encourage earlier experimentation and adoption by industrial code development teams. The recently EU funded project InterTwine is going into this direction [43].

Build & Productivity & Debug tools

Debugging tools have made good progress in recent years, and instrumentation of major compilers has improved significantly, allowing catching many errors early, but some aspects are still not well covered. Notably, debugging of OpenMP thread issues remains difficult, as vendor specific solutions exist, but open source tools are lagging. Valgrind's DRD or GCC's thread sanitizer require compiling codes with a build of GCC using specific options to avoid false positives, and the ARCHER tool [9] seems promising, but requires keeping up as a separate package with LLVM, and so build/version compatibility issues may arise easily, limiting the tools practical availability.

Job Managers

Here also, there is a plethora of job managers, and their number tends to increase. For complex workflows (especially coupling of codes), interaction with job managers (or usage of advanced job manager features) may be required, and the lack of standardization is an issue for the sustainability of complex couplings.

Given the needs of balancing ease of use and limitation of IO, systems allowing running HPC computations on all available (and compatible) resources but prioritizing recently used systems with a form of "cached" data would be very interesting. Different batch systems already handle some data dependency and transfer aspects, but all in their specific way. As configuring such complex priorities, and combining them with complex code usage schemes (especially when using multiple coupled codes), standardization of job manager usage would be essential here. In that end, the EU project NEXTGenIO [44] is looking at this with "data-aware" schedulers implemented under SLURM.

Coupling

The EESI2 project mentioned a need for strong and unified at the European scale code couplers. As of today, most code couplers are restricted to relatively loose coupling algorithms (with explicit schemes at the time step level).

Today, there are several solutions for code coupling, which acceptable parallel performance relative to today's codes, but which will require more work (or a new generation of tools) for increased parallelism or stronger coupling schemes.

The US has some strong libraries, such as ANL's MOAB, or the DataTransferKit coupling library initiated in the CASCL project (and now part or Trilinos), but these libraries still seem quite complex to use but the EU also has some well-established libraries, such as PALM coupler from CERFACS. EDF and CEA also have parallel coupling libraries, MEDCoupling (part of

the SALOME platform project), and PLE (Parallel Location and Exchange, part of the *Code_Saturne* CFD tool, but also used at least by BSC). These tools are each useful in their respective usages, and have limited overlap, covering different subsets or levels of coupling requirements, but increased collaboration could improve all those tools. More specifically, the PLE and MedCoupling developers have iterated over specs/requirements in the past, and work together on other subjects, and OpenPALM's CWIPI tool is based on an ancestor of EDF's PLE library, its main developer having worked with EDF in the past. With some developers from these teams already having some collaboration experience, a well-defined common effort around a multi-layered spec allowing better interoperability and coordinated developments could lead to a strong EU coupling presence, with a good balance between ease of use and performance.

Again here, the experts stressed the urgent need to develop at the European scale a unified and scalable code coupling environment, federating the efforts of multiple teams across Europe, in order to address the need of multi-scale and multi-physics simulations.

2.2 Weather, Climate and Solid Earth Sciences

WG3.2 Working Group is dealing with scientific and society-oriented applications in the fields related to earth sciences, ranging from the fluid envelops to the planet interior. WG3.2 may be called WCES, standing for "Weather, Climate and solid Earth Sciences".

The scientific questions at stake include the dynamics of both the atmosphere and the ocean, at all spatial and temporal scales, with the many interactions between the various scales which are significant from an energetic and/or transport points of view. This includes progressing in the numerical simulation at increased spatial resolution, so that the effects of unresolved scales can be made as small as possible. This includes also achieving a proper description of transfer through the various media, both in terms of turbulent and wave processes. In particular, propagation of waves directly induced by man-made forcing is critical for assessing the physical properties of these medias, which are otherwise very difficult, if not impossible, to document with in situ techniques.

Besides being aimed at allowing progress in the description of the physics and dynamics of the different media at stake, and also in the properties resulting from their interactions and couplings, WCES is concerned with progress in such critically important issues of applied and societal interests like:

- meteorology and forecasting weather at all scales, from local to global, from now-casting (less than a few hours ahead) to medium-range (a couple of weeks up to a month);
- climate dynamics, climate change and its impacts, including numerical simulation of future climate toward the end of the present century, including possible changes in the extreme events regime, assessment of relative importance of natural climate variability as compared to anthropogenic climate change;
- understanding and predicting seismic phenomena;
- exploring through wave propagation the structure the upper kilometres of the earth crest for supporting gas and petroleum industry.

From a computing point a view, addressing the above scientific issues requires to be able (i) to perform numerical simulation at higher resolution, (ii) to interface the models with the wealth of observational data, (iii) to efficiently collect data produced by various numerical models and

exchange them between a large number of groups, (iv) to efficiently adapt the current computational models to the new generation of high performance platforms.

The long-standing issues concerning the evolution of future architectures of computers and how these new architectures will/should influence the development of next-generation of weather and climate models is now a bit clearer: future machines will be based on manycore and hybrid processors, with memory levels which will require modellers to **rethink, re-engineer and recode very substantial parts of their codes** in order to increase the computation per byte of accessed data. Moreover, new algorithms will be needed in order to take advantage of the new possibilities of the hardware. The possible use of GPUs or FPGA's for climate modelling has been explored recently and the speed of computations of realistic models can be already significantly improved by limiting the data movement overhead or achieving a realistic load balance in the heterogeneous computation. On the other side, manycore (MIC) or latest x686 architectures could be efficiently exploited by improving the code vectorization. Another interesting possibility to reduce computing time could be to use mixed precision (single precision instead of double-precision computations), by taking into account the effect on the accuracy of physical model results.

The necessary refactoring of numerical codes is given a lot of attention and is stirring a number of discussions. Energy efficiency is an increasingly demanded requirement and a very real necessity to achieve future goals, where the waste of computational resources is not anymore a possibility. Various options are still opened, *e.g.*, about the way to deal with concurrency in coupled modelling and/or to decouple the description of physical processes (chemistry, radiative transfer, interaction with biosphere, …) from the simulation of the "dynamics", *i.e.* of the flow itself. Another issue is the question of possibly separating the science issues from the HPC optimization: is this separation of concerns feasible? How far could Domain Specific Languages (DSL's) help in this respect? Are there other promising avenues for achieving this separation of concerns? Which are the main benefits in terms of code development and computational performance deriving from separation of concern- approaches? Could code generation benefit from shared standards, as supported by ESMF (Earth System Modelling Framework) in the US? The number of questions to explore is quite large, the gap between actual and future models has never been so important. Each particular modelling group cannot face all these difficulties and challenges and solve them efficiently by itself. The **entire community must then look for coordination and sharing of tasks**, a necessary condition for preparing efficiently to the forthcoming exascale era.

On the data side, the **Exabyte challenge is recognized as being even more crucial than the Exaflops one**. Some ideas have found their way and are already implemented and used with success, *e.g.*, XIOS. The community nevertheless feels it is just at the start of addressing the data issues: better ways to efficiently process data have still to be explored in a much more detailed and complete manner, and this will require more exchange, and possibly more organization, between the groups. From a climate data analysis perspective, a different approach based on high-performance/big data analytics "tools" exploiting server-side analysis capabilities, should to be explored. This is especially true for the climate change community, considering that ESGF (which has been serving the Coupled Model Inter-comparison Project Phase 5 experiment, providing access to 2.5PB of data for the IPCC AR5) is mainly a large, federated, data sharing infrastructure. Considering CMIP6 (the next major international climate exercise just started in 2016), the challenge will be even stronger considering that the expected number of published datasets will be around 20-30 times bigger than CMIP5 in 2019/2020.

The **collaboration among vendors, computational and climate scientists** (in a co-design approach) should help putting together expertise at different levels. It is seen as an efficient way to (re)-design climate models able to benefit to a large extend from the next Exascale systems. In this respect, a series of meetings between climate and weather scientists and vendors is planned for late-November/early-December and for mid-January in Hamburg. The ESiWACE CoE is further working toward the active participation of weather and climate modelling for the definition and preparation of ETP4HPC' future Extreme-scale demonstrators (EsD's).

From a climate data analysis perspective, a paradigm shift must be envisioned in the current ESGF infrastructure (Earth System Grid Federation). A **different approach based on data-intensive facilities running high-performance analytics frameworks jointly with server-side analysis capabilities, should be explored.** *Data intensive facilities* (representing a counterpart to the HPC eco-system for climate models simulations) close to the different storage hierarchies will be needed to address high-performance scientific data management. *Server-side approaches* will intrinsically and drastically reduce data movement. Download will only relate to the final results of an analysis. Such an approach would strongly reduce the amount of data downloaded on the client side as well as the complexity related to the analysis software to be installed on client machines. Yet, as server-side approaches will benefit from data locality, the geographic datasets distribution will require specific tools or frameworks to orchestrate multi-site experiments to support multi-model analysis experiments. So, end-to-end workflow systems will be needed to handle such kind of aspects. Finally, server-side approaches will also require a strong effort on interoperability and standard interfaces as well as governance and organisational issues solving in order to build highly interoperable tools and environments for climate data analysis. In this regard, the Research Data Alliance (RDA) [10], EUDAT [11] and ESGF [12], amongst others, are already working on these topics at different levels with valuable contributions on big data analytics, array-databases, persistent identifiers, standard interfaces for server-side processing, workflow tools etc. **Providing more compatible data services between these and PRACE would be a very valuable help to application scientists**.

The following graph issued from an article called "Climate Data Challenges in the 21st Century" published in Science [13], highlight the evolution of the volume of climate data coming from instruments (satellite, radar, …), large scale simulations and in-situ/other, and the growing importance of computational data coming from both capabilities simulations and capacity simulations (ensemble).

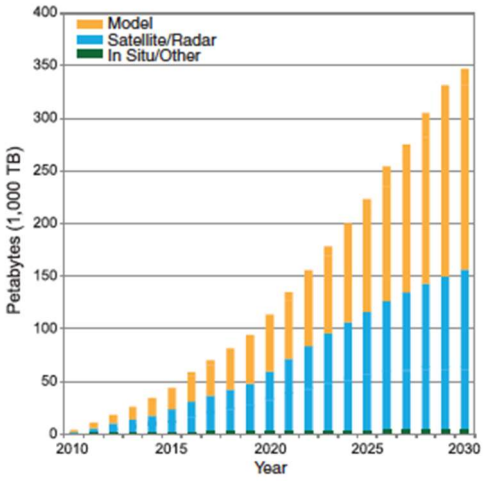

Figure 7 - evolution of data requirements in climate

In such a landscape **joining HPC and big data with Cloud technologies** could help on deploying, in a flexible and automated manner, analytics "tools" as containers or virtual machines, thus enabling highly dynamic scenarios in both private clouds and cluster environments. It should nevertheless be recommended that these efforts be prepared with caution, and that the participation of the modelling communities be obtained beforehand, in order to avoid any waste of time and resources.

2.2.1 Challenges in numerical weather forecast and meteorology

Numerical weather prediction (NWP) will undergo fundamental scientific evolution in the next decade because current atmosphere-focused models will evolve towards Earth-system models with much enhanced detail of processes in the ocean, on land, sea-ice and snow, and regarding atmospheric composition. This evolution brings NWP models closer to the physical complexity already covered in climate prediction models, but at much higher spatial resolution and using ensembles for predicting forecast uncertainty. However, NWP production schedules will remain tight as the socio-economic value of forecasts is strictly tied to the timely availability of forecast data. Given the increase in model complexity and much enhanced user communities feeding NWP model output into downstream services, for, *e.g.*, regional predictions, hydrology, air quality and marine services, data volumes will also substantially increase. This challenge requires investment in efficient post-processing and data dissemination tools.

Currently, global medium-range operational NWP models run on O(10,000) cores to complete forecasts within ~1 hour and produce about 1 PB of data per week. These numbers are expected to increase by **2-3 orders of magnitude in the next 10 years due** to the increase in model complexity, spatial resolution, and ensemble size. The following figure highlights mid-term needs of ECMWF in terms of model resolution, number of computing cores required and associated power consumption, for both single executions of NWP models and ensemble simulations (coupling multiple single executions).

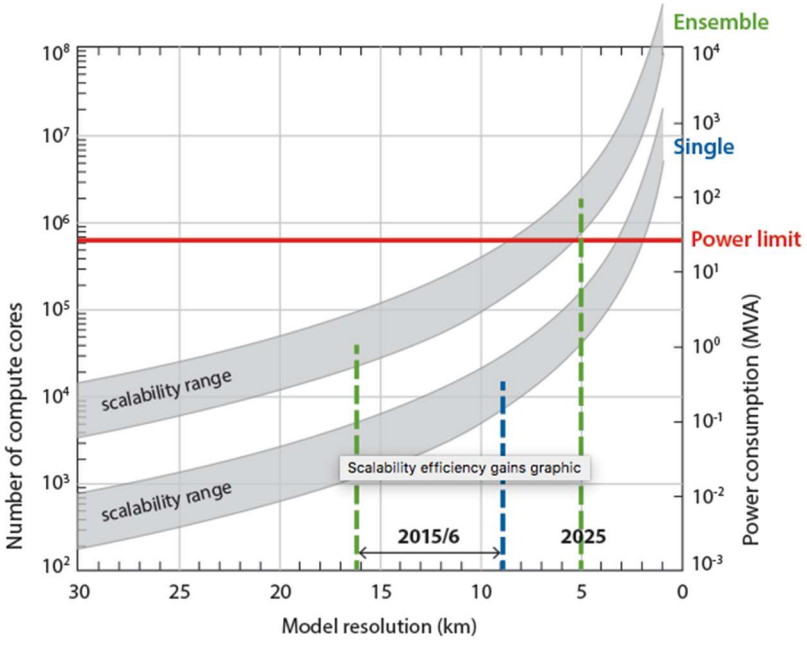

Figure 8 – Long term roadmap of ECMWF needs addressed into its Scalability Programme

Important for NWP is that this growth needs to be addressed along the entire forecast production chain, namely from observational data dissemination and pre-processing, the generation of initial conditions using complex mathematical algorithms and the forecast model, the forecast production itself, to the post-processing of model output and product dissemination.

The key areas of development are currently:

- *Numerical model components and algorithms*. The focus here is on dynamical core developments and data assimilation algorithms, where a key requirement is to maintain scientific flexibility with respect to grid-mesh choices, advection schemes, physical parameterization schemes, solvers, and between variational and ensemble Kalman filter techniques. Approaches to exploit concurrency between physical processes and model components is investigated as well;
- *Work flow management* targeting a reduction of data processing in the time critical path, continuous streaming and processing as data is being produced, parallelization of processing across fields and products, and resilience;
- *Programming models* supporting the separation of concerns between physical science code and highly optimized libraries. Enhanced parallelism and computing intensity requires a combination of existing programming models (MPI, OpenMP, GPI) with domain specific languages and hardware specific compiler directive or languages. Given the complexity of NWP models, flexibility and portability are key;
- *Novel hardware options* promising more (energy) efficient processing. Currently MIC and GPU processor types are being tested but efficiency gains need to be traded off against code adaptation effort and portability.

At present, these developments are pursued in H2020 projects such as ESCAPE, NextGenIO, InterTwine and others where co design actions are planned with the hardware and software vendors.

Furthermore, the recently established Center of Excellence ESiWACE offers many opportunities to increase significantly the convergence between NWP and climate modelling. Up to now only the Met Office (UK) and Météo-France (FR) were able to coordinate in-house developments between these two domains, as these establishments are tasked with both applications and are consequently interested in using as much as possible common software elements. The fact that most European national meteorological services are now working within ESiWACE in closer relationship with all European climate modelling centres will increase the exchange of information, will draw attention toward solutions originating from one domain and of likely interest to the other (as it could especially be the case for data issues), and will hopefully lead to common developments.

2.2.2 Challenges in climate oceanography

Progress in ocean science is intricately linked to the computing power available because of the need for high resolution in ocean models and data assimilation systems. Furthermore, ocean models need to be coupled to wind waves, tidal motion and marine biogeochemical numerical models (leading altogether to about 100 state variables). The latter are empirical numerical models and optimization algorithms for empirical parameter values should be developed for different parts of the world oceans. Operational oceanography is a new and rapidly growing sector, providing key assessments for coastal water quality, fisheries and marine ecosystems, offshore industry, military operations, maritime transport, etc.

The advent of satellite measurements of sea level (altimetry) and of the global Argo array of profiling floats have led to major breakthroughs by overcoming the historically sparse data coverage in the ocean. However, ocean biogeochemistry is still largely subsampled and

coupling of ocean with hydrological models is at its infancy so that large uncertainties are associated with ocean biogeochemical cycles. Ocean wind waves and sub-mesoscales drive variations in the upper ocean vertical motion that affects nutrient supply and thus ocean biota and they are not yet resolved by the global as well as the regional models. Another key concern is that ocean general circulation models do not consider yet tidal forcing so that they do not resolve the mixing processes induced by internal tides that are not possible to parameterize. Coupled ocean currents and wave modelling has greatly advanced at the theoretical level in the past 15 years but no large-scale model is still coupled to surface waves and only few coastal models are. These key concerns can only be addressed by building upon recent advances in ocean modelling to include more processes and multivariate data assimilation systems, new high-resolution and accurate numerical models.

Four main challenges have been identified in the field of oceanography:

- Challenge #1: High-resolution ocean circulation and data assimilative models;
- Challenge #2: Understanding and monitoring the carbon cycle in marine ecosystems;
- Challenge #3: Increasing the lead time of ocean predictions, from weeks to decades;
- Challenge #4: Predictions and climate impacts at the coastal scales.

Running ocean/sea-ice circulation models that can resolve the **whole spectrum of ocean dynamics** (*e.g.*, at resolution of about 1/100°) is at **present beyond any computing capability**, except in very local areas and over short periods. Reaching kilometric resolution at global scale demands indeed a thousand-fold increase in computer power. Adding a full carbon cycle model increases the computational cost and storage capacity by a factor of approximately 5. The highest resolution global ocean/sea-ice circulation model presently used in Europe for research and operational forecasting (*e.g.*, ORCA12 used in CMEMS) uses a grid resolution of 1/12° (*i.e.* 5 to 10 km). A single 50-year long run of this model requires an available peak computational power of 25 Teraflops for a period of 2 months, which represents about 2.5 % of the annual peak power available on a Petaflops computer. The most commonly used eddying ocean models (for operational forecasting, seasonal and climate prediction, ocean climate variability studies) use a grid resolution of about 1/4° (10 to 25 km) and require, for a 50-year long run, a peak computational power of 10 Teraflops for a period of 20 days. In the near future, the scientific objectives of the oceanographic community will require, on the one hand, performing series of multi-decadal experiments with O (1/12°) models or ensemble runs of O(50) members with O(¼°) models, or increasing eddying ocean model complexity with a full carbon cycle. Operational oceanography, on the other hand, urgently needs to develop higher resolution products for its Marine Core Services, and global models with a grid resolution of, *e.g.*, 1/24° (resolution of 5 km at low latitudes and up to 2 km at high latitudes) need to be developed in the **medium-term**, asking for an increase of computational power by a factor of 10 or more. At the same time, we need to make significant progress in data assimilation: accurate ocean simulations require accurate initial conditions, accurate forcing fields and accurate calibration of model parameters. These requirements are even more demanding than modelling in terms of computational power. Moreover, it would be important to include grid refinements, down to 1 km-100 m scales to explicitly resolve the dynamics of specific oceanic regions (coasts) that are critical for the global oceanic circulation. Altogether, the above requirements call for peak computational resources of **500 to 1000 Teraflops available for periods of months** in the medium-term. This can be obtained only with O (10-100) Petaflops computers coupled to very large storage facilities O (10-100 Petabytes) that will store the simulation outputs over long periods (at least O (5 y)) for subsequent studies. In the **long-term, Exascale** systems are needed to provide the computing power to move towards very high resolution models (of about 1/100°, i.e. 1 km), which also take into account processes and

factors (such as coastlines, eddies, tides and surface wind waves), essentially unresolved in current climate models.

2.2.3 Challenges in Solid Earth sciences

Computational challenges in Solid Earth sciences span a wide range of scales and disciplines and address fundamental problems in understanding the Solid Earth evolution and structure and its interaction with the Ocean and Atmosphere. Solid Earth sciences have significant scientific and social implications, playing a central role nowadays in natural hazard mitigation (seismic, volcanic, tsunami and landslides), in energy and resource exploration, in environmental protection (containment of underground wastes and carbon sequestration), and in national security (nuclear test monitoring and treaty verification). In the realm of seismic hazard mitigation alone, it is important to recall that, despite continuous progress in building codes, one critical remaining step is the ability to forecast the earthquake ground motion to which a structure will be exposed during its lifetime. Another direction is to develop capabilities to make this type of prediction in **near real-time for Earthquake Early Warning systems**. In terms of seismic imaging, going to full-waveform inversion and to high frequencies and thus to much increased resolution still remains a formidable challenge, both from a high-performance computing and from a big data point of view. Building capabilities for analysis of huge datasets collected by Solid Earth observation systems is another challenge for this discipline. All these areas of expertise require increased computing and data analysis capabilities in order to provide breakthrough science. A programme of **provision of European leadership class computational resources (like PRACE) will make it increasingly possible** to address the issues of resolution, complexity, duration, confidence and certainty, and to explicitly resolve phenomena that were previously simply parameterised.

Recent promising extensions and developments in terms of application fields of solid earth sciences include applying similar imaging technologies to industrial non-destructive testing.

2.2.4 Common requirements and recommendations

Access to data – generated by large simulations, large instruments and observational systems – is **changing the ways in which we think** about and address problems in science and societal research, including climate change and environmental hazard and risk. Shared data is catalysing change in global collaboration and in businesses. **A user-driven sustainable data and-compute European wide e-infrastructure needs to be application-oriented**, easily accessible, open and agile so that it can continuously adapt to changes in technology and research practice. Providing data adjacent to the HPC does indeed avoid data traffic and save energy, besides providing high speed IO between and storage and compute systems.

Case studies must, therefore, include the relevant science and technology expertise while identifying gaps in missions, resources and capabilities, and also involving appropriate business models to explore how to sustain long-term stewardship. A crucial innovation to be explored is the agile strategies that will allow to intimately couple research thinking with technical innovations. These strategies need to be deployed in a multi-organizational context in order to discover solutions and ways forward that closely match domain researchers' requirements, which evolve rapidly as the potential of new capabilities and data are appreciated. Eventually, this will change the focus from HPC production environments and community data services,

where middleware standards, security procedures and connectivity dominated, to research-focused scenarios.

Having large memory per node, possibly resulting in fewer nodes with higher memory so as to remain economically feasible, would facilitate a lot of the computation in ways that a number of geoscientists don't have available currently on most machines. A target of 2-5TB per node would be appropriate.

Further, making the large computational resources practical to use for a large number of application scientists may require easier interfaces/languages to be developed. Easier access to training workshops for young application scientists, to become specialists not only in their own science domain or HPC but both, should be provided. Funds to establish collaborations between scientists and computer scientists or programmers, are also thought to be very helpful steps.

2.3 Fundamental Sciences

WP3.3 is in charge of assessing strategic development in high performance computing and how to prepare the scientific community to Exascale computing for the vast topics of fundamental sciences, covering statistical physics, physical chemistry, plasmas or astrophysics applications, lattice QCD, particle physics, materials, molecular dynamics etc. Given the broadness of topics many different techniques and degree of parallelism of the applications have been developed and implemented on either a particle-based, fluid-based or hybrid approach.

With the emergence of the new architectures towards Exascale computing, the software and algorithms need to be adapted, which may require substantial code developments. Locality in memory and reduced communication patterns are keys to succeed efficient applications on these new architectures. With the increase in computation power, we will also be able to increase the number of physical processes that can be simulate and therefore improve the realism of our current simulations and increase the resolutions ($10\ 000^3$ particles for example and more).

Mini applications and demonstrators on test machines will pave the way to find the best strategies, e.g.:

- Deep changes in the algorithms such as a switch from Godunov MUSCL/Hancock schemes to Discontinous Galerkin (DG) in compressible fluid dynamics;
- Deep changes in the software such as a switch to task programming.

2.3.1 Challenges on Nuclear Physics and QCD

Answering questions related to our field of research have mostly a societal impact on our understanding of how elements are formed within our universe, particularly elements relevant to the formation of carbon-based life as we know it. Research on nuclear reactions could potentially impact industrial applications related to storage of nuclear waste.

The goal by 2021 and beyond is to perform most lattice calculations of hadronic systems at or near the physical pion mass, with lattices representing physical volumes of $(4\ \mathrm{fm})^3$ and larger. To achieve robust signals from these types of calculations, the scale and of the problem must be **increased by at least a 1000-fold compared to today's calculations**, and most likely larger.

We are not aware of any roadmap that has been developed in Europe, but a roadmap has been formulated in the US [14] and an update to the roadmap is in the works. For NLEFT, one challenge will be to find the edges of nuclear stability, both for neutron- and proton-rich nuclei. This requires extreme computing resources (at least a factor 100 better than used so far) to deal with the extreme small energy gaps in these systems.

In terms of software, the available simulation codes for simulations in Lattice QCD (LQCD) are highly advanced. Low-level optimized implementation exists for any large-scale architecture, be it XeonPhi (e.g. QPhiX) or GPU (QUDA) based architectures or the still existing Blue Gene's. Optimized communications libraries exist as well, which utilize the low-level interfaces to the hardware to cut down latencies and to optimally use the hardware capabilities. These software packages are maintained by a broad community that is and will be willing to invest into optimized software also in the future.

Algorithms: Multi-Grid methods are available for most of the different discretisation of the theory of QCD. Important algorithmic developments include:

- Optimally computing matrix traces in so-called disconnected diagrams;
- The development of multi-level techniques to improve signal/noise in future calculations (removing an existing exponential noise problem in complex calculations);
- Development and verification of simulation methods to be used at finite chemical potential which causes a "sign-problem", i.e. the Boltzmann weight becomes complex (stochastic Langevin/Lefschetz thimbles).

In terms of problem size, a moderate increase in the problem sizes to 256^3x512 is conceivable, but will not be required for all use cases. However, the requirements for precision required for important questions such as the QCD contributions to g-2 are dramatically larger as for quantities computed thus far and, hence, even if the problem size does not increase, the resource requirements can still increase significantly.

In terms of infrastructure, the availability of prototype hardware has proven to be essential for the software readiness and software quality as described above. The field maintains close ties with hardware vendors to ensure access to upcoming architectures, which typically are made available as a small testing installation at a national or European wide HPC sites.

Lattice QCD applications have made good use of all available architectures. In almost all cases, the codes are bandwidth bound, where the exact bandwidth requirements could be improved by recent algorithmic developments, at the price of a larger memory footprint (e.g. Multi-Level methods using smoothers with improved data locality). Still communication and memory bandwidth are the primary concern of LQCD implementations. Here, the stagnation in communication bandwidth in upcoming architectures is an important concern.

Access to adequate HPC resources provided by national or the pan European HPC infrastructure (PRACE) is one roadblock. Lattice calculations have poor signal-to-noise ratios, and therefore ample statistics are required for robust signals, which in turn require large HPC resources. Further, many systems of interest have an inherent sign problem due to an imaginary term within the Lagrangian. No adequate algorithm to date has been developed to deal with such an imaginary term, although some partial answers can be given. Contraction algorithms grow factorially in complexity, and therefore clever algorithms need to be developed in conjunction with future architectures to alleviate this problem.

Both NLEFT and LQCD have benefited from utilising current capability machines, and in the case of LQCD, have even helped develop such machines (e.g. QCDOC, IBM BG series, etc.). The relatively simple stencil structure of these calculations means that they are readily adaptable to different processors and architectures. In the case of LQCD, with its lower memory footprint, it has been highly successful in utilising accelerator GPU systems that are typically memory bound. Vectorisation has been successfully adapted for the new Xeon Phi processors for LQCD calculations, showing comparable or better performance than accelerator systems on calculations that are I/O bound.

2.3.2 Challenges on Plasma Physics

Plasma physics of either cold or hot plasmas is more and more common in our technological society (non-thermal plasma for food process, plasma torch, space plasmas, etc) and any improvement in our modelling of such complex systems will have scientific, industrial and economic impact. Compact laser-based particle accelerators and x-ray sources may soon become important tools in medicine (ion beam cancer treatment) and in-situ imaging of biological systems respectively.

High-energy astrophysics is evolving in the direction of hybrid codes requiring fluid and particles algorithms with mutual feedback to be developed. Problem sizes will probably double in the next 5 years provided CPU power is adequate. This requires more than **50 times the actual available computational power**. It is desirable for numerical methods to achieve better than 2nd order accuracy to maximise the efficiency of future-generation codes. **Adaptive mesh refinement also plays a crucial role in this framework.**

In a timeframe of 2022, being able to model the gyroradius scale for ions and electrons while following the global scale dynamics is one of the main goal of plasma physics. They share with astrophysics and fusion the difficulty of bridging large range of scales both in space and time.

In terms of hardware/software environments, Magneto-Hydro-Dynamics (MHD) applications do not do well on BlueGene-like architecture relying on the employment of lots of cores with small computing power. In the future, Intel-phi technology with AVX SIMD performance will be most suited for this kind of applications. By contrast, explicit PIC codes have been rather successful at exploiting the IBM BlueGene line: a number of the major community codes have demonstrated scalability up to 0.5M cores, albeit for relatively homogeneous systems. GPU requires immense re-coding and may not be worth it unless a new common graphic accelerator language is becoming a world-wide standard as MPI did 2 decades ago or the integration inside OpenMP of GPU acceleration is correctly handled.

Exascale computation will produce very large data arrays and reduction techniques will be a must. One possibility is to have reduced floating point precision (2 bytes or less). The second one is to write data at lower resolution.

In the field of laser-matter interaction the commissioning of 3 large-scale facilities under the European Light Infrastructure (ELI) project will likely drive a huge increase in demand for heroic Particle-in-Cell simulations. Fully electromagnetic, kinetic PIC models are essential for predicting and interpreting the outcome of experimental campaigns to advance laser-driven particle accelerators and short-wavelength light sources. Currently 3D, multi-billion particle simulations are routine and are already capable of matching experimental conditions for some laser-electron schemes. However, ion acceleration schemes that rely on denser material are still woefully under-resolved, and numerical results are often too optimistic regarding the beam

properties. **To achieve quantitative predictive power in this case, at least 10-100x more particles would probably be necessary.**

For laser-plasma interactions there are multiple challenges. For electron acceleration, spatial and temporal resolution must be a fraction of the laser wavelength (1 micron) and period (~ 1fs) respectively; yet the simulation times must follow an entire acceleration stage of at least several centimetres, ideally up to a metre, thus requiring the integration of Maxwell's equations over millions of timesteps. This leads to issues with numerical dispersion and artefacts in the propagation behaviour. Radiation sources based on laser wakefield acceleration in principle require sub-nanometre spatial resolution, e.g 1000x better than the current state of the art. For ion acceleration and some laser-fusion (fast ignitor) schemes a pressing issue is the need to model the high-density, opaque plasma accurately. This is currently addressed using some kind of hybrid or coupled model technique in order to simultaneously treat the collision-less laser heating region and the denser, colder material in which the hot electrons are transported. Many of these studies suffer from poor statistics of the interesting (accelerated) populations of the distribution function and would therefore benefit from order-or-magnitude increases in simulation size.

2.3.3 Challenges on Fusion

Energy production by nuclear fusion would be characterized by the abundance of the primary fuel, a very low carbon footprint, and modest radioactivity problems (compared to fission). It is therefore essential to pursue research in this field at the EU and world level. This need is materialized by the ITER project, whose organization involves all the world economic powers and whose goal is to produce a prototype of fusion reactor which demonstrates the feasibility of nuclear fusion as an energy source.

Magnetic confinement fusion aims to achieve controllable thermonuclear fusion by confining a plasma (composed mainly by deuterium and tritium) in a suitable magnetic configuration. This plasma needs to be sufficiently dense and must be heated to sufficiently high temperatures in order to obtain enough thermonuclear power output to make the scheme worthwhile for energy production.

The challenges are enormous. On the technological side one needs materials and engineering solutions able to withstand the high-energy throughput and mechanical stresses. On the scientific side one needs to understand the plasma behaviour in a complex geometry, but in this respect the challenges are common to most of plasma physics, e.g. in the astrophysics context.

Plasma modelling relies on fluid models, such as MHD, when applicable, or more and more often on kinetic models that incorporate the small scale non-collisional physics that plays a key role in phenomena such as plasma instabilities and turbulence which are at the heart of the transport processes in fusion reactors. The mathematical problems take the form of partial differential equations (PDE) in 3 spatial variables, in the case of MHD or other fluid models, and in five or six phase space variables in the case of the kinetic equations. The HPC challenges, especially in the kinetic case, can be extreme, but, overall, not very different from other fields such as astrophysics, fluid dynamics, weather forecasting. The challenge stems from the wide range of spatial and time scales to be covered, from meters and seconds scale at the machine scale down to millimetres and milliseconds or microseconds for the small-scale processes governed by non-collisional physics. All this translates in mesh sizes of several thousand in at least two spatial variables (and in the hundred range in a third spatial variable), several hundred

in the velocity variables, and extremely long time integrations encompassing thousands of time units (taken on the basis of the fastest relevant process).

The goal by 2022 is to **model a full tokamak device with ion and electrons coupling using gyrokinetic approach** (so called full-f distribution function) in **support of ITER initiative/reactor device**. By 2022, the electrons will certainly still be heavier than in reality but the mass ratio between ion and electron will become more realistic as time and supercomputers progress.

The biggest fusion codes are still based on the MPI/OpenMP model, although small scale projects to test the use of accelerators are under way. A key aspect in limiting the adaptation to new architectures is the manpower factor. Because of the **limited human resources** one needs to compromise between scientific production achievable with older and not optimal codes and the necessity to adapt, at some point, the software tools. This is a common problem these days.

Experience shows that the broad range of codes and algorithms needed to model the different plasma phenomena require the availability of a correspondingly broad range of computer architectures. It is simply too risky and scientifically limiting to invest only in a given class of codes in view of the Exascale. However, one element is in common to most of the codes: the need to have sufficient memory per core. The current trend of putting inadequate memory per core will produce a memory bottleneck for an Exascale size computer.

The biggest simulation carried out so far by our group involves a data structure of 2TB to describe the plasma state (distribution function). However, this is not saved and instead one produces a series of sections in lower dimension as the simulation results. They do not seem to pose particular big data problems. The situation would change if, for any reason, such as testing a particular theory, one wished to study the self-correlation of the distribution function at different times. To this end, techniques of in-situ lossy data compression are envisioned.

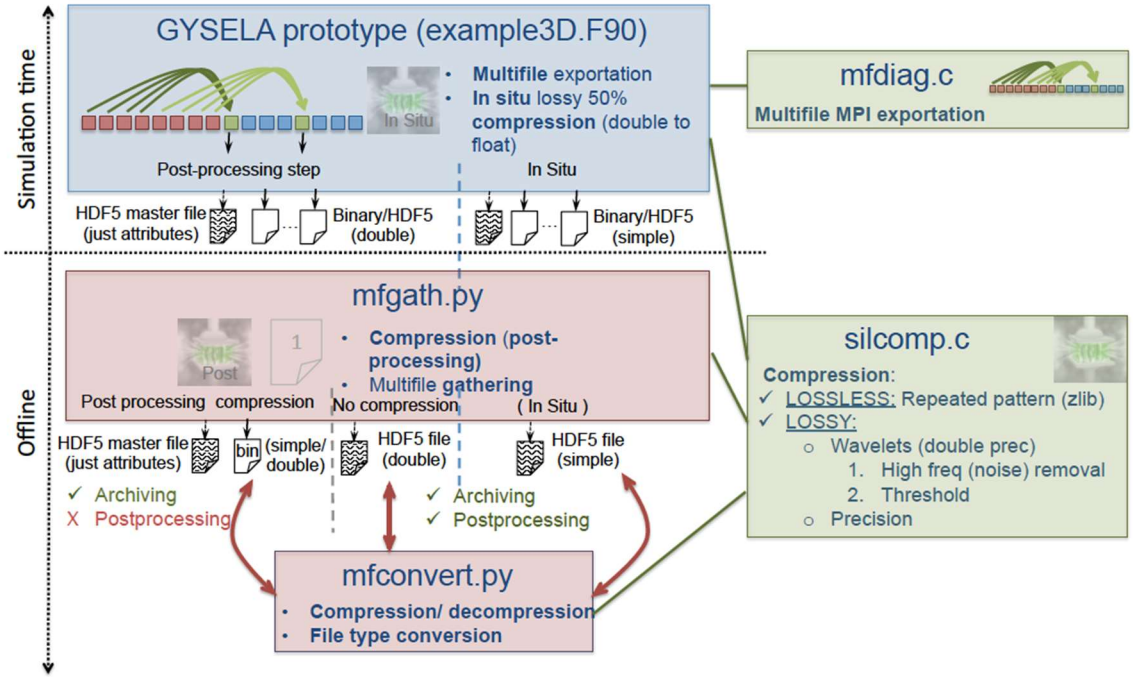

Figure 9 - first result of lossy data compression on massive 3D datasets, leasing to a 26x gain in I/O bandwidth with the GYSELA code

2.3.4 Challenges on Astrophysics and Cosmology

Astrophysics and Cosmology seek to answer fundamental questions such AS how the universe was born, where does life come from, how were the Earth and our solar system formed, are we alone? Any improvement in answering such questions will have deep impact on our society such as the recent proof of existence of the gravitational waves, highlighted in theory by Einstein one century ago.

From the industrial/economical view point, answering such difficult questions requires novel detectors and multi-physics carrier (not only light of various energy, but also neutrinos or gravitational waves). This implies advanced technical development to be able to detect such carriers.

Further the Earth living around an active star, the Sun, it is directly influenced by its storms and intense magnetic activity. Space weather provides simulations and tools to prevent such activity from impacting our technological society. Recent nationwide report in US and Europe to develop a coherent view of taking into consideration large solar impacting events (with hundreds of billions of euros of damage cost if a Carrington like event would strike the Earth), have been issued. There is thus a direct impact of space hazards (including asteroids) on Earth.

In astrophysics and geophysics, modelling the complex turbulent medium occurring in stars, planets and the interstellar medium is key. Turbulence is often coupled in these multi-scales systems with gravity, rotation, magnetism, radiative transfers, nuclear and chemical reactions etc) making their modelling extremely time consuming and difficult. They share with other field the needs to model global dynamics organization while resolving small scale phenomena, often in non-equilibrium state or subject to a large range of fluid and plasmas instabilities. Key problems like the origin of the Earth and Solar magnetism and 11-yr cycle for the latter, the formation of stars and planets, how accreting disk behave, how radiation field ionized the Universe, or galaxies form and evolve to create the Hubble sequence, or how stars blow up and enrich the interstellar medium are very demanding in term of computing resources and physical description. Linking all these scales together is also a tremendous challenge or the year to come and require at least Exascale-class supercomputers.

In astrophysical fluid dynamics (and its geophysical counterparts), high performance simulations are devoted to the understanding of multi-scales, multi physics systems such as the interstellar medium, convection and turbulence in stars and planets, dynamo action and magnetized (low plasma beta) dynamical systems, global instabilities, disk accretion. The goal in 2022 is to bridge as much as possible the large scale global dynamics with the small-scale reconnection and dissipative scales, sharing with turbulence/fluid dynamics problem the need of having larger and larger Reynolds numbers (or its equivalent Rayleigh, Elsasser, Taylor (or Ekman), etc…). It is realistic to believe that by **2022 at least 4 order of magnitude larger in each dimension** will be performed on a regular basis in order to do a systematic parameter space exploration, with the most extreme grand challenge simulations reaching about 4.5 to 5 order of scale difference in each direction.

The largest simulations in astrophysics and geophysics to model either geo or solar dynamo, interstellar medium structuring or galaxy mergers required huge resolution to resolve the global and small scale dynamics. This translates into hundreds of terabytes of data per simulations. **Huge data management of 4-D structures (space + time)** are required to understand the complex nonlinear physics and feedback among the various scales/objects/processes. **Immersive (remote) data visualisation is required to identify key structures.**

In solar and geophysics data assimilation of the large amount of observational data are becoming more and more common and require multi-level parallelism to have ensemble approach while computing highly resolve physical models that are driven by the assimilation procedure to yield more realistic solutions.

Overall large memory per core (4 GB/core or beyond) and high bandwidth must be accessible as the global nature of many of the astrophysical problems make it difficult to run on slow/low memory systems.

In cosmology the trend is to both increase the resolution of the FUR (Full Universe Resolution) simulations and to better take into account gravitational physics laws for being able to 1) simulate bigger volumes for rebuilding the redshift space with higher resolutions and 2) take consider non-linear collapse phenomena of dark matter. Finally, a third objective is the possibility to take into account several fluids like baryonic gases and hydrodynamics into n-body simulations for understanding the formation of macro-stars.

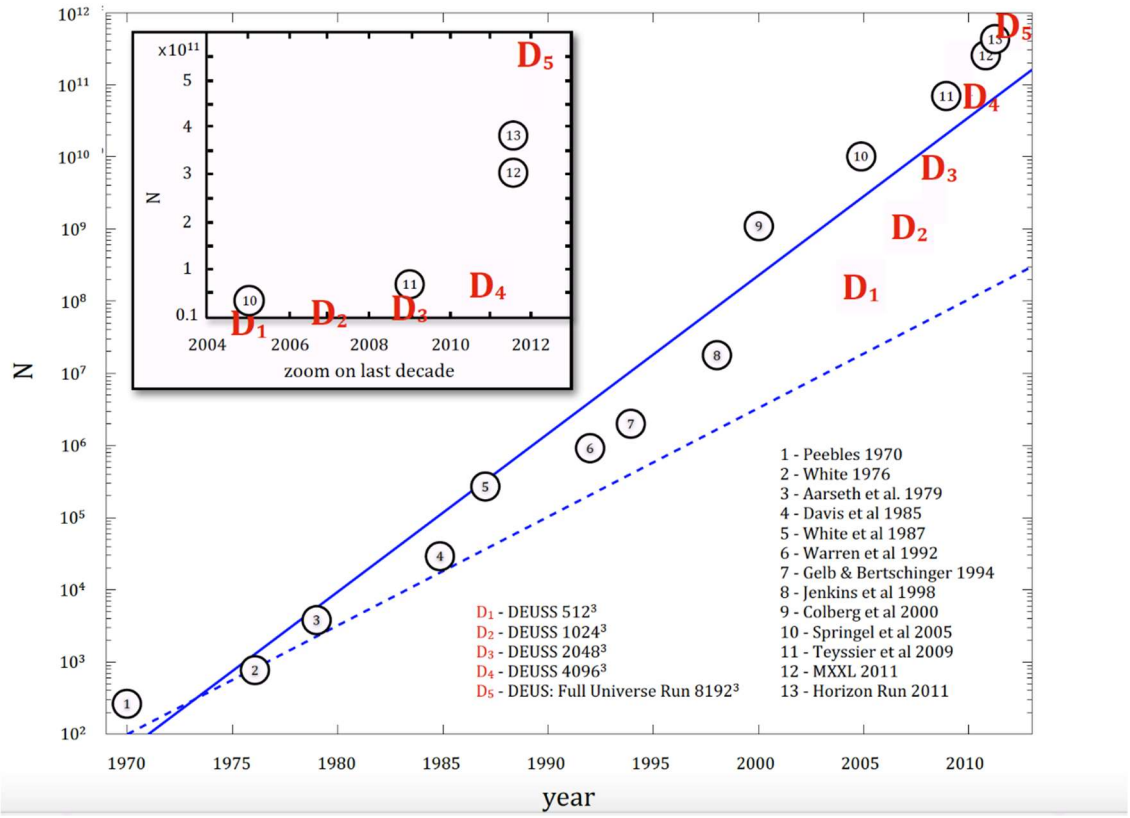

Figure 10 - evolution of the size (number of particles) of dark matter simulations over the time

This will lead to upcoming **multi-scale multi-physics dark matter simulations with a resolution of 32384^3 particles** (instead of 8192^3 or 16384^3 now) able to access to the formation of solar halos of a size of 10^{11} solar masses. These simulations are very important in the context of the launch in 2019 of the EUCLID mission to map the geometry of the dark energy by ESA, which could represent 75% of the Universe.

Such simulations will generate massive amount of data (for example already 150 PB of rough data for a 8192^3 FUR simulation performed in 2012 by Observatoire de Paris [15]) which will require the development of on the fly in-situ post processing techniques coupled with deep learning methods for being able to follow for example only pertinent turbulent structures.

2.3.5 Challenges on Material Sciences

Industry already does use atomistic simulations – particularly DFT based ones as they are predictive and easy to use. However, the amount of use is modest – even very big companies have only a few people working in this area – at first sight this is difficult to reconcile with an annual publication rate exceeding 30,000. We feel that this is because industry challenges tend to be very complex but our atomistic simulations are predominantly applied to rather simple systems – thus there is a disconnect between academic usage and industrial requirements. We suppose that the E-CAM CoE is trying to address this problem as is, to some extent, the MAX CoE. There are reports available on the economic impact of modelling [16].

The field of quantum mechanical modelling has expanded enormously over the last two decades primarily driven by the increasing use of density functional theory (DFT). Remarkably, the rate of publication of papers that use DFT now exceeds 30,000 papers per year and is continuing to grow. What is also remarkable is that on the materials side of DFT the field has been utterly dominated by Europe. However, this leadership is potentially under severe pressure at present (see later).

By 2021, then a new generation of post-density functional methods (DFT) will be available and will become as usable as DFT. While there will continue to be a role for DFT calculations these new techniques will offer the possibility of truly predicting physical and chemical properties to an accuracy of the order of 0.1% or better compared to the typical DFT value of 1-2%. **Some of these new methods may scale to very large numbers of cores but on the whole materials and chemistry methods do not scale to enormous core counts**. However, the computational power is still required to do multiple instances – e.g. vibrational modes at every wave-vector in order to calculate free energies. One area of advance will be high quality calculations at finite temperature – for instance to predict phase diagrams. **High throughput methods to generate large volumes of materials data are becoming more established and by 2021 there will be very large amount of materials data that could help enormously in designing improved materials** etc if we find effective methods for extracting meaning and insight from this data.

Given that quantum simulations and, indeed, most materials/chemical simulations methods will not (and probably do not need) to run on entire PRACE level machines and the use of such machines will primarily be through multiple instances, parallel time etc. then scale up of most of the science we would want to do does not need particular resources beyond machines routinely available. We think there would be interest to see whether we could use efficient Petascale algorithms once these were developed in order to parallelise further but this is not really a roadblock to our progress. Of the materials CoEs, then MAX is probably the one closest to this area of work.

Large simulations (100,000 atoms) will be performed probably based on QM/QM (quantum Mechanics) or QM/MM (Molecular Modelling) first in the field of biology.

These multiscale approaches will be also developed with the paradigm of the complexity reduction. The idea of complexity reduction is to consider few freedom degrees for some part of the systems. Because we call do full QM calculations for large systems, then we can control the complexity reduction. An illustration of that is to consider the multipole of solvent molecules far from a protein.

By 2022, if asked to identify a ‘Grand Challenge’ for the field it would be to **compute entire finite temperature phase diagrams of materials ‘ab initio’** and, furthermore, identify all competing structures within a finite free energy window of the stable phase at each temperature

pressure. Knowledge of this would provide an invaluable resource for explaining materials' behaviour.

2.3.6 Common requirements and recommendations

As the spatial resolution of the I/O datasets is expected to continue to grow, the experts consider that the use of adaptive meshes generated by scalable meshing tools will be generalised as well as smart data management/analysis tools.

In most of the cases, simulations will involve coupling of multi-scale and multi-physics components so the development of ultra-scalable solvers as already recommended by the EESI2 projects is still one of the key priority.

In order to achieve performance portability across the different new architectures (with multilevel parallelism and/or hybridation) toward Exascale computing, different innovative methods with general purpose application are being pursued such as:

- Task programming with a smart runtime scheduler e.g. starPU;
- Implementation with directives e.g. pragmas with OpenMP or OpenACC;
- External libraries e.g. Kokkos.

Finally, **disruptive approaches like parallel-in-time methods are interesting** and subject of on-going studies, but presently face serious obstacles due to the requirements of the simulation process (area conservation, reversibility). Some potential use of this approach has been reported in astrophysics and material sciences, where in this last field a lot of independent calculations to sample phase space (for instance in crystal structure prediction). Multi-level parallelism has been widely exploited in the plane wave codes such that, for instance, CASTEP can now use 24,000 cores. While this is an order of magnitude higher than what one might have expected – particularly given the steady fall in memory or I/O bandwidth of parallel machines compared to their compute performance (and we are very bandwidth dependent) this gives a good indication of how far away we are from using a Exascale resource for a single calculation!

The experts are proposing the 6 following recommendations:

- The need to have a Center of Excellence of reference for the field of Astrophysics and Cosmology. It is surprising that it does not exist given the leading role of Europe in those fields of research. Geophysics and plasma physics could join as well in a new or existing centre which will make place to turbulence, gravity, magnetism and nonlinear phenomena in complex fluid/plasma systems.

- One should assess what is crucial to be developed in EU with support of EC and national funding agencies and what could be leveraged by collaborating with countries outside EU. In the same vein, it would be ideal to set up a long-range plan, similar to that done in the US, related to the coupling of various areas of sciences and HPC. Within this long-range plan, the roadmap for creating and utilising HPC at the Exascale and beyond could be made explicit, as well as a stronger coupling between the various science disciplines (due to similar algorithm methods, exchange of ideas, etc.).

- Europe has to keep investing in the development and procurement of HPC architectures, as is done in the Exascale research projects such as DEEP and Montblanc, but generalizing it to all field of fundamental sciences.

At the beginning of this decade Europe seemed able to keep up with the growth of the other major players, but now the gap seems increasing especially after the decision to renew PRACE took so long. Looking at the recent TOP 10 HPC list, Europe comes in at #8 and #9, behind China, US, and Japan, so Europe is potentially trailing, as for instance in the US the CORAL initiative for Exascale computing is already in place.

- HPC and fundamental sciences is not just about clock cycles – endless examples show that being clever, developing better methods etc. buys you a factor of 100 against every factor of 10 gained through Moore's Law. Despite the endless evidence on this point, every country in the world invests far too little in software and expert computational scientists and far too much in big computers. So, yes we need HPC but if the budgets were better balanced with greater investment in software Europe would achieve far more than our rivals even though they have more clock cycles.
- The likely trend of smaller memory per core needs to be avoided as much as possible
- Although deep learning and big data approaches are quite popular nowadays, we need to be careful in using them for the analysis in HPC simulations in fundamental sciences. We highlight the fact that in many scientific communities, a simulation has very little value if the result is not understood in terms of the key physical ingredients and reduced models that can highlight these key processes. Deep learning is a powerful tool e.g. for classifications or model reductions but will never give the physical understanding that might (or not) be hidden behind the classification or reduced model, so it can't be the alpha and omega in our field of research.

2.4 Life Science & Health

The life and medical sciences community is highly diverse and tackles problems with hugely varying length and time scales, ranging from the interactions of the atomic building blocks of living cells to those of human populations. The study of such systems encompasses data-centric approaches in fields such as genomics as well as theoretical models based on physical and chemical understanding. This combination of targets and approaches creates a highly diverse set of computational requirements. In order to capture the needs of the community we have focussed on three key areas in life sciences and health that have specific HPC requirements: data driven bioscience (including genomics and informatics), molecular simulation and biomedical simulation. Our expert panel brings together members of two EU-wide Centres of Excellence (CoEs), BioExcel and CompBioMed, designed to support academia and industry in using high-performance (HPC) and high-throughput computing (HTC). The focus of BioExcel is exclusively on biomolecular simulation, whereas CompBioMed aims to establish modelling and simulation at all scales as an integral part of clinical decision making.

2.4.1 Data driven bioscience

The development of high throughput experimental techniques has revolutionized the way science is performed in many areas of biology. The vast amount of data produced presents new challenges in terms of both its management and interpretation. The best-established example of the increased volumes of biological data is genomics, where the throughput of next-generation sequencing techniques is increasing much faster than Moore's law.

Since 2000, the cost to sequence a whole human genome has continued to collapse. From $3.7 billion, it dropped to $10 million in 2006, and to $5,000 in 2012. In 2014 it costed $1,000. To date, the rate of the decline has outpaced Moore's Law by three to four times. As shown below, at either the historic rate of decline or Moore's Law, the cost to sequence a human genome will fall below $100 in the next five years.

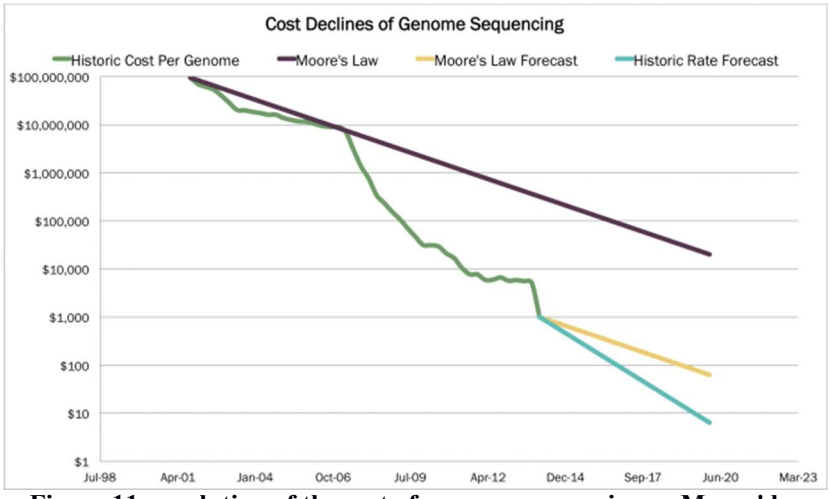

Figure 11 - evolution of the cost of genome sequencing vs Moore' law

With trillions of genomes waiting to be sequenced, both human and otherwise, **the genomic revolution is in its infancy**.
Variation in the genetic makeup of each patient influence their response to clinical interventions, in both positive and negative ways (such as increasing toxic side effects). **Understanding of these factors should pave the way for stratified or personalised medicine**, in which treatments are tailored to groups of similar patients or even the individual, to replace the outdated concept that a single drug is the solution for all members of the population. These advances are set to accelerate as sequencing projects are extended to entire populations, enabling more complex investigations such as linkage studies.
The **key requirements in this field computationally hinge on data management** (including issues of **confidentiality** and **privacy**). Increasing data sizes will make increased demands for increasing algorithmic efficiency and the use of distributed data architectures (such as Hadoop or Spark). Genomics is the canonical example of a Big Data application, requiring resources which enable flexible usage models and where I/O performance is of equal or greater importance than processing power. The use of single GPUs or a small number of cores over a constrained timeframe lends itself perfectly to the use of cloud resources.

Whilst genomics has been a trailblazer, it is by no means unique in processing large volumes of data. From investigations to identify the functional elements of the genome (such as the ENCODE project[8]) to understanding the structure of the brain, it is increasingly common in the life sciences to handle large and heterogeneous datasets. For example, typical datasets derived from high resolution volumetric brain scans are of the order of 10 TB in size. **Image processing and batch machine vision of such a dataset will usually use hundreds of thousands of cores**. This will often be only one step in a multi-stage workflow potentially involving further compute intensive steps using tens of thousands of cores. Similar scenarios are increasingly common in other fields too, such as the processing of data from cryo-electron microscopy (Cryo-EM) experiments. In many such instances, a key concern is the accessibility of data produced by an instrument, in such circumstances **co-location of compute and storage is highly desirable**.

[8] https://www.genome.gov/26525202/encode-pilot-project

2.4.2 Molecular simulation

In the biomolecular area, detailed computer simulations add an essential extra dimension to be investigation of biological molecules by allowing functionally relevant motions to be visualised. This is transforming the fundamental understanding of biochemistry and molecular evolution, and has practical applications such as the design of new drugs. Such models can also be used for repositioning and targeting therapies for precision medicine, through rapid and accurate assessment of drug efficacy in specific disease cases. Molecular simulation techniques vary in resolution from those which incorporate quantum mechanical detail to those which abstract much of the atomic detail (such as Brownian dynamics) in order to investigate interactions in larger systems and on longer timescales. The grand challenges facing the field are to successfully **model larger scale biological systems**, such as organelles or whole cells, and the **generation of accurate, reproducible quantitative results** (such as drug binding affinities and kinetics) in addition to qualitative insight.

The most common level of theory employed at present is atomistic molecular simulation, which is well suited to exploiting the current scale of computational resources whilst investigating single macromolecules. Typical job sizes range between tens and thousands of cores, depending on the size of the system of interest. Simple aggregation of a greater number of cores on ever larger machines can enable the study of larger systems but will not favour the study of the longer timescales necessary to study some phenomena (such as folding and conformational changes). However, much recent work has suggested that for many properties of interest more efficient sampling can be gained through the use of ensembles of multiple shorter simulations. As the number of atoms in the systems of interest moves into the millions, as required to simulate sub-cellular compartments or viruses, it is likely that multi-scale approaches involving will become increasingly important.

Coupling coarse grained molecular dynamics, Brownian dynamics or even higher level models (reaction-diffusion models for example) in **multiscale/multiphysics simulations will require both the availability of flexible, hierarchical resources but also hierarchical and concurrent multiscale modelling frameworks and workflow tools**. Such use cases are better served by heterogeneous machines, with hierarchical capabilities in terms of the number of cores, amount of memory, memory access bandwidth and inter-core communication (in contrast to a ‘flat’ machine with peak processing power). Even for single scale simulations workflow tools that automate the running of analyses and potentially even reactively respawn simulations based on the results of them. Such automation would remove the need for the costly (in terms of time and effort) cycles of simulation, transfer of data to a local resource for analysis and then human involved decision making as to whether more simulation is needed. These advances would build on simulation and machine learning tools which are already installed on most HPC resources but require both innovative combination of these tools and more flexible modes of running than are permitted by queueing rules at present. The combination of resources and policy would ideally allow combination of single core simulations or analyses with the ability to run larger scale simulations for extended periods of time (up to weeks in extreme cases).

Whilst the use cases of the community are varied, biomolecular simulations remain very much dependent on floating-point performance and will benefit from the wide availability of compute resources with GPUs or other accelerators. Both traditional HPC and cloud are key for different types of users. In particular interactions with the pharmaceutical industry are frequently enabled

by commercial cloud offerings where security concerns have been addressed and ongoing costs are minimized.

2.4.3 Biomedical simulation

Biomedical simulation encompasses a wide range of computational approaches seeking to provide computational models of cells, tissues and organs in order to aid in clinical decision making and treatment development. In addition to seeking to describe complex living systems this field interacts directly with clinicians and clinical scientists which creates additional concerns surrounding turnaround times, as well as privacy and confidentiality issues. *In silico* models are a source of innovation and are therefore of great interest to industries such as pharmaceuticals and medical device manufacturers (including imaging instrument manufactures).

HPC is capable of enhancing each of these sectors and underpinning emerging sectors, such as those concerned with *e*-health and personalised medicine. Biomedical simulations can complement advances in bimolecular simulation, for example predicting toxicity in drug development. They can also provide added value to medical device measurement data, for example as acquired by various imaging modalities. In the medium to long term, biomedical simulations could lead to *in silico* clinical trials for use in the development and regulatory evaluation of medical products, devices, or interventions. They offer the prospect of significantly reduced costs, time to market and animal experimentation.

Modelling of entire systems in the detail necessary to incorporate patient specific information requires the ability to perform many simulations with high memory requirements, large numbers of parallel tasks and/or long run times. Even with relevant resources available, discovery will be limited by computability. A key step in the development of productive models, particularly those required to generate results on the timescales needed to influence clinical decisions, will be the ability to derive reduced-order models which capture the complex physics of patient-specific anatomy and boundary conditions from solution of a representative population of compute intensive simulations.

In many areas the first step towards building biomedical models is the creation of systems based models integrating data from multiple levels of biological organization. Examples include the understanding of brain-wide connectivity or gene and transcription networks. These cases often share many of the same concerns as the field we have described as data driven biosciences above but additionally run predictive simulations. This again leads to the need for better methods for the couple of data collection from experiments and analysis and modelling.

The precise computational requirements of any given simulation are largely problem dependent. Patient-specific biomedical simulations generally involve computational fluid dynamics (CFD) and/or finite element analysis (FEA). Differential equation (both ODE and PDE) solvers are also used to represent cellular behaviour in multi-scale simulations. An idea of the variety of requirements is presented below:

- CFD analyses typically utilise around 16 cores to provide analysis of local vascular regions. Run-times vary in the range 4 to 48 hours for typical jobs. Memory requirements are of the order 10s of GB;
- FEA typically requires more memory (order 100s GB), benefit less from parallelisation (~ 4 cores) and can require longer run times (several days);
- Multiscale/multiphysics simulations typically involve combination of CFD and/or FEA with other simulation approaches, often requiring iterative simulation between codes at each time-step of the solution. This imposes the same individual requirements as for the cases above, but increases run-times by an order of magnitude (to days/weeks);

- Tissue-scale modelling, thanks to high-throughput experimental techniques, may require the solution of non-linear finite element models with up to a billion degrees of freedom; tissue adaptation simulations might require that these models are coupled to a bone remodelling model, so that may have to run hundreds of times.

A similar need for easy to use tools for the running and management of ensembles of simulations are required here as was articulated for biomolecular simulation. These are necessary not just to gain statistical power in results but also for sensitivity analyses and parameter scans that are vital for model reduction and refinement.

Additionally, Big Data applications impact biomedical simulation through the relationship between population-based models of health, disease and response to intervention and patient-specific analysis. This includes the use of healthcare databases to define anatomical atlases and generate representative models to allow patient-specific simulation when not all model parameters are captured during the clinical workup. Treatment stratification for individuals is informed through integration of clinical outcome data from healthcare records with appropriate simulation metrics.

HPC facilities play an important role in all of the scientific challenges identified above. In cases where reduced-order models can be effectively used to deliver simulation support for end-users it is likely that iterative refinement of the reduced-order model will result in continued requirements for HPC resource. Cloud compute provides a solution for rapid development and deployment of new software solutions, particularly where there is a focus on graphics/user interaction or a requirement for specific operating system support. In order to support the use of mature simulation cases in the healthcare domain on demand delivery of HPC-based simulation will be required. In order for this to become routine this provision will have to reach higher TRLs than is currently available. In particular, this use case raises challenges in ensuring the availability of the simulation service and ease of use for end-users without a technical computing background. Cloud solutions offer a bridge between HPC resource and end-users with limited technical computing background. Accelerators (in particular GPUs) are critical for many of the floating-point intensive codes used in predictive modelling.

Whilst increase computational resources are clearly necessary for conducting the ambitious agenda of biomedical simulation, novel algorithms for model order reduction are likely to be of at least equal importance. Advances in this area would provide enormous benefits for applications where the reduced model is required to quickly provide information to the end–user (e.g. a clinician during a patient visit). In particular, for applications where models are integrated within the treatment process (e.g. image-processing) the aim is to deliver real-time patient-specific analysis.

In order for biomedical simulation to achieve its ambition of integration with clinical workflows, **specialised software stacks** are required to simplify the delivery of *in silico* medicine services to end users. This involves a significant commercial effort, and associated support in non-technical areas such as **ethico-legal** aspects, supported by access models that ensure **secure data digestion and transmission**.

2.4.4 Common requirements and recommendations

Whilst the overall landscape of life sciences presents a hugely varied set of demands for HPC and cloud resources several key issues are found across the field. Two main issues unite much of the field; the need to couple understanding at different resolutions or levels of model and the centrality of large dataset which are often highly sensitive. The former can be further broken

down into the needs for hierarchical hardware and a multiscale modelling software stack. The encompasses both a need for **an integrated data infrastructure and data security policies**.

Further to these requirements there is a general need for the understanding that there will be large amounts of work just keeping current high-performance simulation engines working well on next-generation hardware.

Multiple classes of HPC resources

The present state of the art in many fields of the life sciences consists of the integration of models with highly diverse computational requirements. Often the best way to make progress is believed to involve coupling simulations and/or analyses with very different computational requirements. Broadly this results in the need for HPC resources that concurrently support three classes of parallel jobs:

- **Massively parallel**: tightly coupled fast connectivity between tens or hundreds of thousands of cores (e.g. Organ modelling, molecular dynamics of organelles)
- **Embarrassingly parallel**: no connectivity between cores but independent use of 1,000-10,000 cores simultaneously (parameter scans of network models, post-processing free energy calculations from molecular dynamics, some CFD methodologies, reaction diffusion)
- **Annoyingly parallel**: ensembles of hundreds to thousands of independent/low-frequency intercommunicating simulations each of which requires 100-1000 tightly coupled cores (ensemble MD simulations for free energy and molecular kinetics determination)

Multiscale software stack

In order to maximize the utility of hardware supporting the requirements we have outlined it will be necessary to have a complementary software stack which facilitates the execution of complex workflows in which components have heterogeneous computational requirements. Such software is likely to need some level of co-design with HPC centres to ensure that queueing policies and technical capabilities do not conflict. A common problem at present is that HPC resources are usually encourage use of the massively parallel regime, at the expense of the ability to conduct innovative science with more mature technology using fewer tightly coupled cores.

Integrated data infrastructure

A key requirement in many fields of the life sciences is a data infrastructure that prioritises not just the storage of large datasets but their accessibility to a variety of computational resources. It is also vital that data storage is not the only priority as in many cases metadata curation, transfer and dissemination are of equal importance. Large-scale datasets need specific services to make them searchable, viewable, accessible, analysable without relocating them. Data access is increasingly becoming a bottleneck, in particular when combining simulations with information coming from high throughput experiments.

In many fields the establishment of a viable data infrastructure needs to be supported by the development of metadata and file format standards. These are pre-requisites for creating the knowledge and data management and sharing tools that underpin both large scale academic collaborations and interactions with industrial partners.

Data security policies

Data security is particular critical for protecting data in private collaborations, and in order to protect human data governed by specific data use agreements. In particular, the data derived directly from healthcare records to inform patient-specific analysis (e.g. imaging, ECG, blood tests, gait analysis data) results in specific requirements for:

- control of access to data (role management);
- tracking of data access (governance);
- anonymization and mechanisms to maintain linked data;
- security requirements (including ISO certification).

End users of any cloud or HPC service will need to provide details of the specific arrangements in place on any infrastructure to secure ethical approval to store data on the infrastructure for computational analysis. This information should be provided as standard in a relevant form by the infrastructure provider.

Porting software to new platforms and improvements in scalability

Rapid hardware development requires big efforts into porting existing software to new accelerator systems in order to take advantage of increased capabilities. Core developers of major codes are working extensively in this direction, e.g. in the scope of the BioExcel Center of Excellence where molecular dynamics simulations engine GROMACS (www.gromacs.org) currently supports the latest high-end GPU processors, Xeon PHIs and OpenCL platforms (e.g. AMD GPUs) and is expected to considerably improve its performance through task-based parallelism; CPMD (a major code for hybrid-QMMM enzymatic studies, www.cpmd.org) is expected to have a dramatic increase in performance based on a new coupling interface for the back-end engines; and one of the most important codes for Cryo-EM image post-processing RELION has been recently ported to GPUs giving great increase in performance [34].

Making new ensemble-simulations methods available to non-experts

A significant number of important biological questions can be addressed in terms of a massive number of ensemble simulation studies as opposed to a single Exaflops job. Such simulations are already being widely used for screening of potential drug targets, studies of pathways for structural conformational changes, dynamic pathways e.g. ion transport, macromolecular formation and reactions, protein/DNA/ligand/saccharides interaction studies etc. There is a great value in utilizing automation systems for such large-scale studies and there are major efforts in devising efficient workflow solutions that allow a much larger and non-expert user base to take advantage of those. Some examples currently under development within BioExcel CoE include workflows for harvesting and processing of genomics data, large-scale free energy simulations of protein mutants, biomolecular recognition and virtual screening of pharmacologically active compounds. In the same direction CompBioMed is developing a "giant workflow" simulator that was successfully used for drug target screening studies.

Exascale orchestration frameworks

Efficient utilization of large Exascale resources via ensemble simulation studies will not be possible without powerful frameworks for orchestration of such level of executions. Notable efforts in the area are managers such as COMPSs and Copernicus, which are also actively developed in the BioExcel CoE

Machine learning

Outside EU, we witnessed the raise of deep learning in contacts prediction as seen in latest CASP experiments [45], which noticeably lead to performance improvements in the last round, whose results will be presented in December at the Gaeta meeting.

USA specifically is investing heavily in startups working on machine learning techniques e.g. in the areas of drug design [40]. Hardware manufacturers such as NVidia are also working on specialized software stacks (https://developer.nvidia.com/deep-learning)

3 Overview of recent accomplishments and breakthroughs in Europe and outside

This section describes for several scientific domains covered by WP3 recent major accomplishments and breakthroughs performed by European teams, using national or European (PRACE) or outside EU (US DoE, K-Computer in Japan, systems in China etc) HPC and data facilities.

3.1 Industrial and engineering applications

3.1.1 Promotion of highest scaling codes

Promoting current highest scaling codes may help potential users choose applications with good scaling. Examples of such promotion include Jülich's High Q club [17], based on codes scaling on the full JUQEEN machine (preferably including multi-threading capability). This list currently contains 27 applications, of which slightly less than half may be applicable to engineering (the others being more "science" oriented). The list includes both CIAO and *Code_Saturne*, in whose development experts from WP 3.1 also participated. This type of list is not exhaustive, and describes mostly weak scaling only on some types of current architectures (mostly IBM BlueGene/Q in this case). Readiness of the listed codes for hybrid architectures is not measured here (and may be different for Xeon phi or GPU anyways). Strong scaling or "baseline" performance of codes is not measured either.

Specific optimisations on these 2 codes have been performed in order to improve the MPI/OpenMP parallelisation, scalability of the numerical solvers, memory consumption and I/O and MPI behaviour.

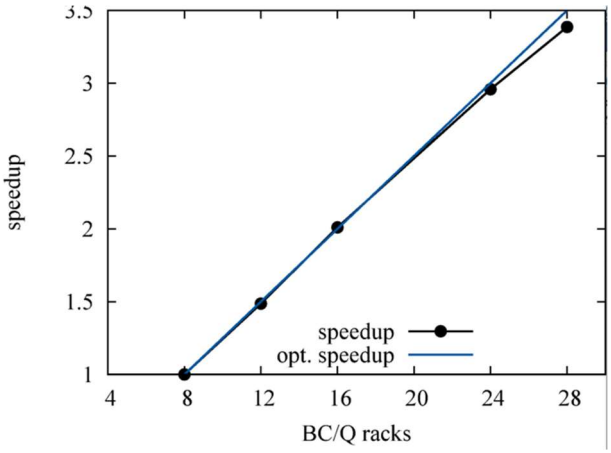

Figure 12 - scaling of CIAO over 28 IBM BG/Q racks

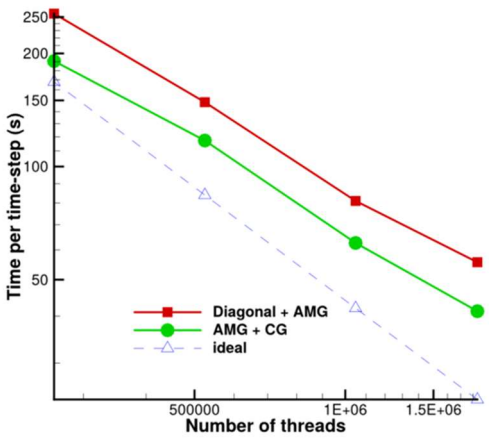

Figure 13 - scaling of Code_Saturne over 1.835 million threads

In the US in September 2016, the Energy Department's Advanced Manufacturing Office announced up to $3 million in available funding for manufacturers to use high-performance computing resources at the Department's national laboratories to tackle major manufacturing

challenges. The High-Performance Computing for Manufacturing (HPC4Mfg) program enables innovation in U.S. manufacturing through the adoption of high performance computing (HPC) to advance applied science and technology in manufacturing, with an aim of increasing energy efficiency, advancing clean energy technology, and reducing energy's impact on the environment. The Energy Department plans to select 8-10 projects for this third round of funding and seeks qualified industry partners to participate in a one-year term collaborative project. Selected projects will receive up to $300,000 to support access to supercomputers and experts at the partnering national labs.

3.1.2 In-situ visualisations

Some in-situ visualisation solutions have recently been shown to scale to large machines. For example, CIAO's in-situ visualisation, based on VisIt, has shown its ability to run on the full JUQUEEN machine. *Code_Saturne*'s in-situ visualisation is based on ParaView/Catalyst, which has not been tested at such scale (due in part to more difficult installation of ParaView on Blue Gene compute nodes), but has shown good behaviour on several hundred cores on Intel Xeon/Infiniband clusters.

3.1.3 Large scale computations

Recent large scale industrial computations in CFD have required more than 1 billion unstructured cells. One example is studies of sleeve leakage gas impact on a British AGR (Advanced Gas cooled Reactor) fuel assembly, done by STFC Daresbury and EDF Energy UK. These studies using the *Code_Saturne* CFD tool have shown that the effect of the ribs might be limited to a region close to the pins and therefore it is possible that the two streams of cold and hot gas do not mix completely before exiting the assembly. These studies required a complex mesh, built of many sub-meshes, which are joined in parallel [18].

Illustrations of the matching computations are shown here. Note that for studies using more precise turbulence and heat flux models, meshes beyond 5 billion cells would be necessary.

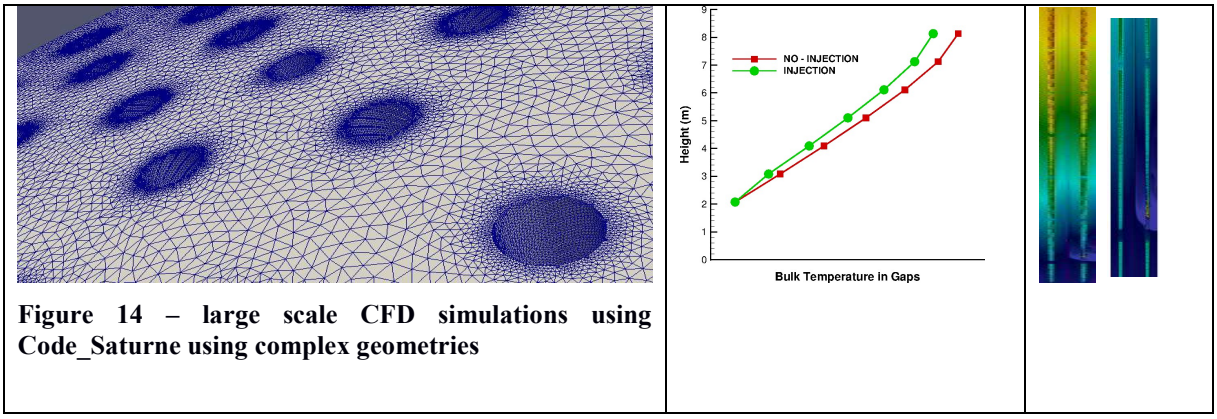

Figure 14 – large scale CFD simulations using Code_Saturne using complex geometries

3.2 Weather, Climate and Solid Earth Sciences

Climate modelling groups involved in HiResClim I & II PRACE projects (i.e. Cerfacs, BSC, KNMI and SMHI) carried out high-resolution simulations for the final purpose of improving our capability to perform initialized climate predictions on time scales from months to a decade. For example, CERFACS, using a coupled model with horizontal resolution of 0.25° and 75

vertical levels for the ocean and 50 kilometres and 31 vertical levels (T359L31) for the atmosphere, showed that decadal predictions have significant skill at 2-5 years lead time over the North Atlantic and Indian oceans. In particular, the decadal simulations captured the 1995 North Atlantic observed temperature shift. This accomplishment was made possible thanks to the 86 million core-hours allocated in the framework of these PRACE allocations on Barcelona Supercomputing Centre (BSC) MareNostrum IBM platform, from January 2013 to November 2014.

ECMWF has recently started testing the computability of global 1km atmosphere model simulations on its own HPC infrastructure for the first Extreme-scale Demonstrator. The simulations confirmed that such a simulation requires approximately 150 times the employed (35,000) core allocation to complete in time if perfect scaling is assumed. However, this does not include Earth-system model capability, which will increase the requirement even further. These requirements will be explored with the second Extreme-scale Demonstrator, where BSC and SMHI will test EC-Earth using a global 10km atmosphere model and a 1/12-degree ocean model. Further, ECMWF is starting high-resolution seasonal forecasting system integrations in 2016 following the award of 60 million core hours on the new NCAR machine in the US, which is being installed in Cheyenne. These simulations will serve to better understand tele-connections between the tropics and mid latitudes, and their dependence on model spatial resolution.

A team forward by researchers from UK-Met demonstrated by simulations significant advances in predicting up to one year ahead the phases of the North Atlantic Oscillation (NAO), which drives European and North American winter variability. The NAO – a large-scale gradient in air pressure measured between low pressure around Iceland and high pressure around the Azores – is the primary driver of winter climate variability for Europe. Understanding and predicting the changes in the North Atlantic Oscillation could have significant economic benefits including potential boosts in climate services for a range of sectors including transport, energy, water management and the insurance industry. This discovery was made possible thanks to increased Met Office supercomputer capacity (A Cray XC40 of 16 Pflops when fully installed for a total budget of £97 million in a new high performance computing facility). This allowed the team to increase the resolution of the climate model and test the retrospective skill of their forecasts over the 35-year period from 1980 onwards.

The Met Office long-range outlook, made in real-time one month ahead of winter, gave good advice ahead of last winter's flooding.

In the field of solid earth sciences, a team formed by researchers from JAMSTEC (Japan), King Abdullah University of Science & Technology (Saudi Arabia), Laboratory of Mechanics and Acoustics (France) and Princeton University, (ISA), performed in 2016 a high-performance simulations of global seismic wave propagation with an unprecedented accuracy of 1.2 s seismic period for a realistic three-dimensional Earth model using the spectral element method on the K computer (Japan). These seismic simulations used a total of 665.2 billion grid points and resolved 1.8 trillion degrees of freedom. To realize these large-scale computations, the team optimized a widely-used community software code (SPECFEM3D) to efficiently address code vectorisation, cache efficiency and all hardware parallelization, especially thread-level parallelization to solve the bottleneck of memory usage for coarse-grained parallelization.

The new code exhibits excellent strong scaling for the time stepping loop, that is, parallel efficiency on 82,134 nodes relative to 36,504 nodes is 99.54%. Sustained performance of these computations on the K computer is 1.24 petaflops, which is 11.84% of its peak performance.

The obtained seismograms with an accuracy of 1.2 s for the entire globe should help us to better understand rupture mechanisms of devastating earthquakes.

Very accurate simulations using data coming from the Tohoku major earthquake, (magnitude 9.0) which happened in March 11, 2011 and which lead to a disastrous tsunami in the Fukushima region of Japan have been performed.

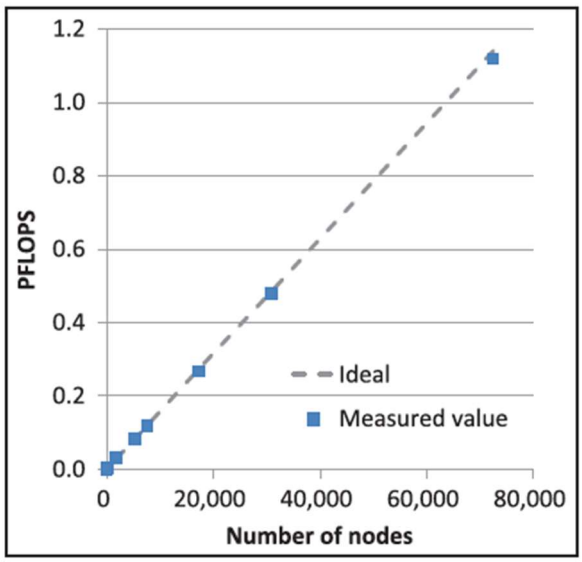

Figure 15 – Strong scaling of SPECFEM3D_GLOBE on the K-Computer over 80 000 nodes

Also, recent progress has been made in both seismic modelling and seismic imaging towards going to high frequencies, in particular for body waves. For instance, Wang et al (2016) [19] used high-frequency body-wave full waveform inversion to drastically enhance tomographic images of complex geological structures underneath mountain ranges. Both resorted to the SPECFEM3D open-source software package.

In China, the Sunway TaihuLight supercomputer became in June 2016 the fastest supercomputer in the top500 with a peak performance of 125 PFlops (93 PFlops sustained). During the presentation of this system, made with domestic components (manycore SW26010 processors with 260 cores, scalable interconnect, … and a software stack) a first set of enabled scientific and industrial applications has been detailed [20]. Over these 6 applications which have been all submitted to the 2016 edition of the Gordon Bell price, three are related to Weather, Climate and solid Earth Sciences:

- A community atmospheric model (called CAM) has been refactored, ported on the system using OpenACC and scaled out to 1 536 000 cores;
- A fully-implicit non-hydrostatic dynamic solver for cloud-resolving atmospheric simulations has been scaled out to more than 10 million cores with more than 1 PFlops sustained performance. This project received the Gordon Bell award during the SC'16 conference in November 2016 for its outstanding achievements in terms of scalability across the full TaihuLight system;
- Ultra-high resolution surface wave numerical simulation (addressing energetic motions of oceans) using the MASNUM model at 1/60º global resolution. The results show a sustained performance on the TaihuLight supercomputer of more than 30 PFlops over 8 519 680 cores (close to the full scale of the system). In order to reach this level of performance specific optimisation in terms of grid decomposition, MPI asynchronous operations have been performed. The report does not indicate however which amount of data has been generated and how it has been managed by the system.

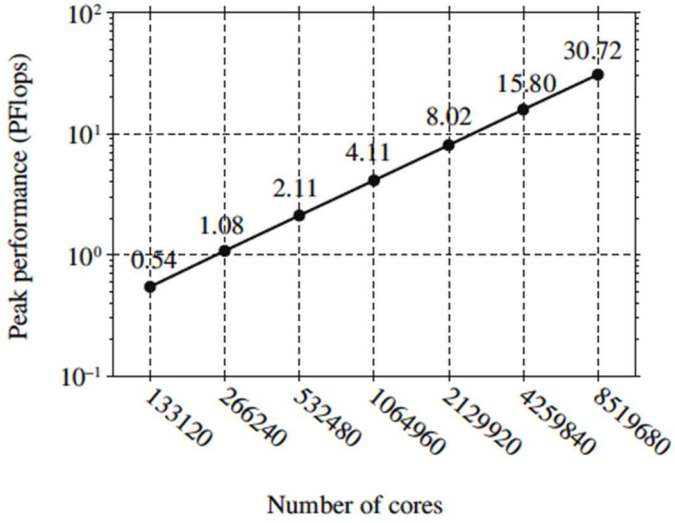

Figure 16 - scaling of the MASNUM wave model over the full TaihuLight supercomputer

3.3 Fundamental Sciences

In the field of atomistic simulations for the design of new materials, Chinese representatives reported in June 2016 with the announcement of the Sunway TaihuLight supercomputer that one application related to the design of silicon nanowires through molecular dynamics simulations. Here thanks to specific optimisation like atom ordering, SIMD vectorization, global memory access reduction, and efficient use of local memory and nonlinear mathematical functions, a sustained performance of 6.62 PFlops have been reached with systems composed by 3.22 billion atoms.

In the field of space weather, the Vlasiator project leaded by Minna Palmroth (Finnish Meteorological Institute, Helsinki) is aiming to enable the world's first global runs of the near-Earth space using a hybrid-Vlasov approach at highest resolutions so far. The domain where space weather phenomena occured is huge compared to any other computer simulation; an accurate and self-consistent simulation needs to cover a part of the solar wind, the entire near-Earth region and also preferably be coupled to the ionized upper atmosphere that is called the ionosphere. The space weather domain is also vastly complex: the physical phenomena may occur in spatial scales from a few kilometres to hundreds of thousands of kilometres, and in temporal scales from milliseconds to years. This sets a demanding topic for computer simulations, which need to be accurate both in small and large scales and be executed for at least a few hours, preferably days or weeks. This project received multiple allocations of up to 24 million core hours by PRACE on the Hornet and HazelHen systems provided by HLRS (Germany) and also support from PRACE experts in optimising the code regarding load balancing issues.

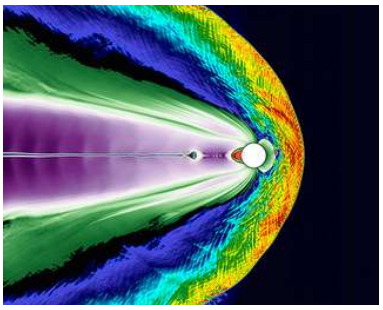

Figure 17 - Solar wind flows from the simulation from the right. Earth's magnetic domain – the magnetosphere – is like a rock in the stream, and a shock forms to encompass the magnetosphere. The color-coding shows the density of plasma, indicating that the shock packs plasma in front of the magnetosphere.

Simulations on more than 100 000 cores have been performed and Vlasiator results show rich plasma phenomena that are much more complicated than has been thought earlier. Many phenomena that can be seen in the runs have only been observed by local spacecraft measurements before, but the mechanisms of explaining the phenomena have not been placed in a larger context.

In the field of fundamental sciences, an international team with researchers from Univ. Wuppertal / Julich Forschungszentrum (Germany), Univ. Budapest (Hungary), Univ. of Aix-Marseille (France) and Univ. of Southampton (UK) performed in 2016 the first ever calculation of the neutron-proton mass difference through ab-initio calculations. The existence and stability of atoms rely on the fact that neutrons are more massive than protons. The measured mass difference is only 0.14% of the average of the two masses. For the calculations, the team developed a new class of simulation techniques combining the laws of quantum chromodynamics (QCD) with those of quantum electrodynamics (QED) in order to precisely determine the effects of electromagnetic interactions. By controlling all error sources, the scientists successfully demonstrated how finely tuned the forces of nature are. Massive simulations were performed on JUQUEEN (JSC, Germany), SuperMUC (LRZ, Germany) and Turing (GENCI/IDRIS, France) using PRACE and national allocations.
The code used is called dynQCD, it has been carefully optimized for the IBM BlueGene/Q architecture, including assembler-level serial code optimizations, low-level communications based on SPI, and threading using pthreads. It reaches an efficiency of close to 40% of machine peak, with perfect weak scaling at this efficiency: all critical communication operations are nearest-neighbour only and have been implemented in an asynchronous fashion, i.e. no global synchronization is required. The following figure highlights the speedup of dynQCD on up to 1.835 million threads on JUQUEEN, super linear speedups are due to exceptional ability to exploit caches as problem size per thread decreases

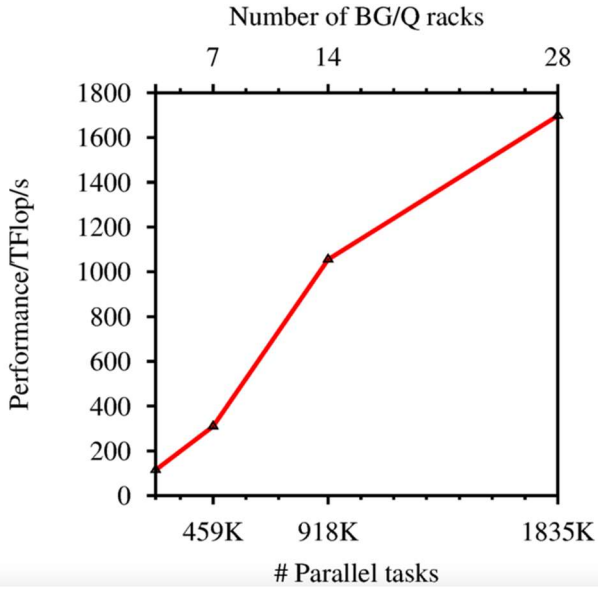

Figure 18 - scaling of dynQCD over 1.835 million threads on a IBM BG/Q system

The next figure shows a weak scaling of the newly implemented "Dual Tree Traversal" (DTT) algorithm as implemented in PEPC.

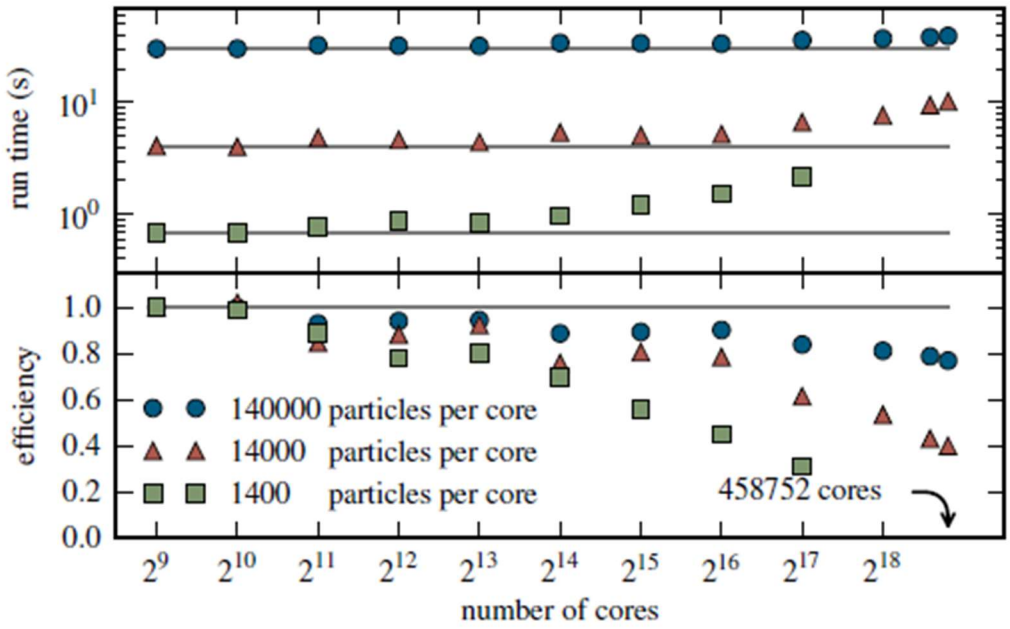

Figure 19 - weak scaling of the PEPC application

Compared to the scaling behaviour of the more conventional BH (Barnes-Hut) algorithm, the DTT (Dual Tree Transverse) shows a more pronounced fall-off of efficiency for the medium sized problem, but performs similarly for large problem sizes, achieving a parallel efficiency of about 0.8 on the whole machine. In terms of total wall time, the DTT outperforms the BH algorithm for large problem sizes, due to its O (N) complexity [41].
This code has recently been applied to the Kelvin-Helmholtz instability of plasma-vacuum interfaces in the kinetic regime - Steinbusch, Gibbon & Sydora, Phys. Plasmas (2016) [42].

Finally, in the field of the design of new materials, the US Department of Energy (DOE) plans to invest $16 million over the next four years in supercomputer technology that will accelerate

the design of new materials by combining "theoretical and experimental efforts to create new validated codes." The new program will focus on software development that eventually may run on Exascale machines. This announcement is part of the U.S. Materials Genome Initiative launched in 2015.

3.4 Life Science & Health

In the area of Cryo-EM (single-particle electron cryo-microscopy) image post-processing RELION has been recently ported to GPUs giving great increase in performance speed thanks to reduction of memory requirements on GPUs to fit widely available hardware, use of single precision arithmetic and increased parallelisation in order to benefit from all level of parallelism available on the algorithms used.

In the field of computational structural biology, some recent approaches have emerged with for instance crowdsourcing projects (FoldIt, EteRNA) [35], chromatin modelling (TADbit, BACH, AutoChrom3D, HSA etc.) [36], the use of immersion and gaming 3D engines in molecular visualisation [37] and also the use of co-evolutionary restraints derived from multiple sequence alignments to study both structured and disordered proteins and their complexes [38,39].

In the field of genomics and personalised medicine, a European research team headed by Professor Peter Coveney, Centre for Computational Science at the University College London (UCL), achieved a major success in the field of personalised medicine: the team were able to prove that by using genomic data of individual patients they are able to predict whether a specific standard drug for the treatment of breast cancer will help or not. By using up to 250 000 cores of the SuperMUC system at LRZ (Germany) during 37 hours in the field of a specific Extreme Scaling Workshop, they studied 50 drugs and candidate drugs to determine how they bind with protein targets in a range of disease cases, in order to rank their potency for drug development and for drug selection in clinical decision making. If this approach evolved into a sustained procedure, it would mark an important advance for personalised medicine, that is medicine designed with one specific patient in mind so that it works efficiently, and without side effects.

In the US, precision medicine has been a hotspot of innovation for a while now, and the interdisciplinary work of medical researchers and computer scientists is bearing fruit in clinical settings already. In fact, the U.S. government has invested $210 billion over the past few years through its Precision Medicine initiative, and earlier 2016 announced the Cancer Moonshot initiative in an effort to improve treatments and ultimately cure cancer.

Biomolecular simulation

In the past two to three years' simulation of bio-macromolecules has made huge strides in both the ability to study large systems and obtaining the sampling necessary to correctly capture important processes such as protein folding and ligand binding. There is a vibrant community involved in molecular simulation of proteins and other bio-macromolecules in the EU. However, the sheer scale of computational resources available in the US means that most of the largest and longest simulations are conducted there.

The biggest challenge facing biomedical science generally is that of reproducibility. Many recent studies have shown the difficulties in recapitulating previous results taken from the literature. This is a pressing concern in biomedical modelling where not only is it scientifically vital but also a pre-requisite for the field to be taken seriously as a tool in the clinical environment. Given this context the greatest benefit of an Exascale machine would be in enabling the running of existing models in statistically significant numbers to enable validation and uncertainty prediction. In the recent years, the meeting of computational requirements in the molecular dynamics feel has had a transformative effect on both the ability to make reliable quantitative predictions and evaluations and the range of systems that can be investigated. It is to be hoped that a similar effect will be seen as Exascale resources allow modellers to run ensembles of organ and tissue level models which test model correctness, parameter sensitivity as well as the inter-individual variability of personal physiology.

Beyond the atomistic scale the most advanced field of biomedical modelling is at the level of organs and particularly the cardiovascular system. In this field, mature modelling tools and environments are available, developed both within the EU (for example CHASTE [21] and Alya Red [22] for cardiac and HemeLB [23] and Palabos [24] for vascular systems) and outside (Continuity 6 from UCSD in the US [25] and ICMA developed in Auckland, New Zealand [26] for cardiac and SimVascular developed at Stanford for vascular modelling [27]). Recent advances have seen truly spectacular simulations of the entire heart using multiscale and multiphysics modelling techniques by the SCLS (Supercomputational Life Science) in Japan [28] and the entire systemic arterial tree (using 1,572,864 cores of a Blue Gene/Q machine) [29]. However, these simulations are not yet mature enough to provide genuine scientific insight and the field still faces considerable challenges in terms of both achieving biomedically relevant spatial and temporal scales, and scaling up to perform large numbers of simulations covering different initial conditions.

It should also be noted that work on these systems is inherently multiscale and increasing requires innovative ways to combine simulations whose computational characteristics are highly heterogeneous.

These requirements point to the pressing need for standardized tools to transform data from clinical imaging modalities into the types of three dimensional representations used within simulations reliably and accurately with minimal human intervention. In some cases the meshing and segmentation required can become a task requiring HPC itself and may need to be optimized for the target computation platform. The production of such tools would represent the first step towards producing routine workflows which would facilitate reproducible investigations into organ function and personalized simulation of clinical interventions. The first step in this direction are currently being under taken by the CompBioMed centre of excellence.

Even in the most advanced areas of biomedical modelling important challenges remain to be overcome before clinical translation becomes a reality. *In silico* trials are however receiving increasing attention from regulatory agencies due to its potential to significantly reduce the requirements for manufactured prototypes and animal experiments in phase 0 and 1 trials prior to human testing. In the US, the Food and Drug Administration (FDA) have initiated pipelines which will in future make it possible for *in silico* trials to be used as part of the evaluation of medical devices and drugs [30, 31]. In the EU context, the Avicenna project has produced a roadmap *in silico* clinical trials [32]. This activity is being promoted and advanced within CompBioMed.

In the last five years, the grand challenges that the Virtual Physiological Human (VPH) research posed were rarely defined in term of computational challenges; thus, very few (if any) accomplishments in the field do have an HPC narrative attached to it.

This is in part a chicken and egg problem, that hopefully the recent launch of the Centre of Excellence on HPC in computational medicine, will help to solve. For a long while, the HPC environments provided by PRACE or other similar initiatives were not designed to host VPH workflows, for a number of reasons including i) the complexity of software stack that frequently involved also commercial codes; ii) the need for confidentiality linked to clinical data; iii) a balanced offer for problems that are not only CPU-bounded, but also frequently I/O, storage or memory-bounded etc. As a result, developments were not driven by a major "computational ambition". It is expected that close relations between CoEs and PRACE will allow to implement now operational modes for innovative computational medicine workflows.

4 Global recommendations

Beyond the scientific challenges and specific recommendations presented by each working group, all the WP3 experts also worked in preparing global recommendations addressing some transverse issues including how to improve post-processing and the creation of new CoE in Europe.

4.1 Convergence between in-situ/in-transit post processing techniques and machine/deep learning methods

4.1.1 Introduction

The nature of science is changing – new scientific discoveries and socio-economical innovation are emerging from the analysis of large amounts of complex data generated by high-throughput scientific instruments (sequencers, synchrotrons, scanners, microscopes, …), observational systems (telescopes, satellites, network of sensors, …), extreme-scale computing (for both capability based large scale 3D simulations as well as ensemble or coupled multiscale/multiphysics simulations), and public World Wide Web.

In many domains – such as astronomy, physics, earth sciences, environmental sciences, genomics, biomolecular research, health sciences, financial, engineering, and social sciences, etc. – our ability to acquire and generate data is starting to outpace largely our ability to manage, explore, analyse, and valorise them both technically and socially, leading to the development of a new field called High Performance Data Analytics.

In order to be able to extract the wealth of information hidden in those data, and to valorise the infrastructures that generate them, new radical and holistic end-to-end data management approaches are needed.

Among many recent reports addressing such vision, the EESI2 project issued in 2015 toward the European Commission and national agencies several recommendations for the development of in-situ/in-transit post processing frameworks as well as identification of turbulent flow features in massively parallel Exascale simulations.

A joint position paper, elaborated by WP3 experts and presented during the last session of the BDEC (Big Data and Extreme Computing) initiative in June 2016, takes these

recommendations as input and propose to **increase their scope by adding machine/deep learning capabilities for the development of real cognitive tools serving an accelerated science**. It advocates urgently funding agencies for a joint call for proposals over 12 pilots, each cross-fertilizing experts of domain science and engineering (combustion, astrophysics, climate, life sciences, …), machine/deep learning as well as HPC experts and centres.

This call will concretely allow the assessment of these new technologies inside the scientific methodology given that experts considers that, as correlation is not reason, deep learning methods won't replace a true understanding of the key physical ingredients of scientific problems.

4.1.2 Key issues and scientific and industrial data analysis challenges

With regard to climate science, deep-learning techniques on large-scale datasets can provide a solid and advanced tool/methodology for understanding climate extremes (i.e. heat waves, tropical storms, and cyclones) detecting events, patterns, and trends as well studying their location, intensity, and frequency. Challenges to face in this domain relate to the large amount of data to analyse, the different scales/resolutions, the complexity of the deep/machine learning algorithms and techniques integrated into a high performance scientific data management eco-system.

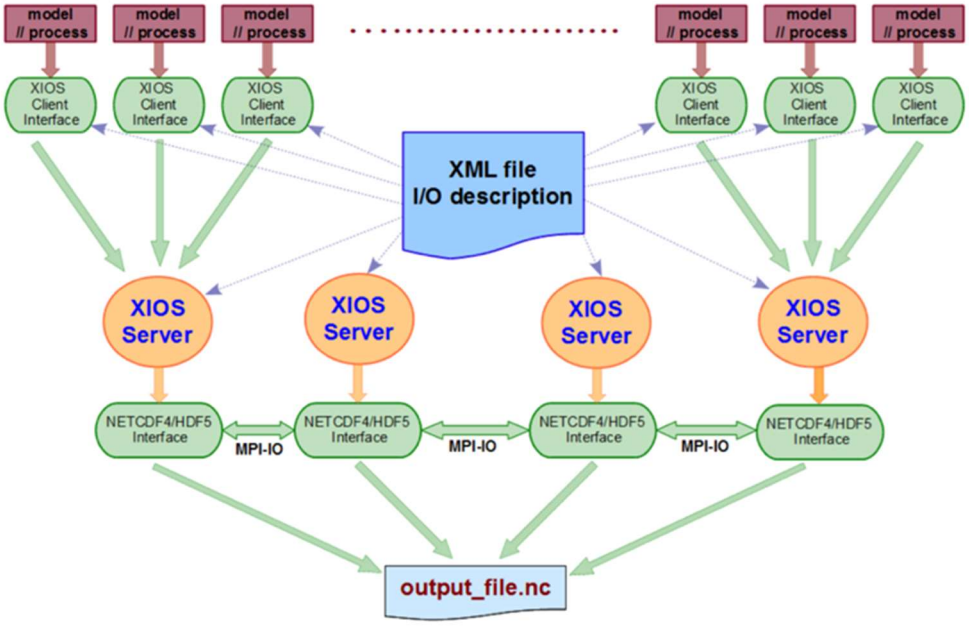

Figure 20 - overview of XIOS (developed by IPSL), an asynchronous in transit I/O library for climate simulations with less than 6% overhead

Numerical simulations play a fundamental role in cosmology today in the understanding of the origin and nature of the dark components of the universe as dark matter and dark energy and their influence on the formation of cosmic structures. In recent years, the first numerical simulations of the full observable universe with dark energy were performed generating huge amounts of data (> PB) that must be analysed and valued rapidly. But the understanding of processes in cosmology requires to study the structure and dynamics of numerous physical observables as the velocity fields, the gravitational potential, the deflection gravitational field,

the gravitational wave field, etc. The use of outcome of deep learning technology tool would allow high progress, but will require strong development of new techniques to be able to handle large volumes of multi-dimensional data.

In industry and especially in oil & gas (reservoir simulation) or in high fidelity combustion/multiphase simulations applied to turbines or engines, the need of in-situ data processing is already a current and critical problem at the Petascale era for several applications with huge amounts of raw data to manage. Typical data rates in the order of 350 TB/run in reservoir modelling or 1 PB/30 min wall clock time in high fidelity combustion/multiphase will not even allow for the analysis of dynamic behaviour.

Current “on the fly” standard post-processing and analysis tools with features extraction applied to reservoir modelling or combustion/multiphase simulations are based on the use of different methods linking statistics, volume rendering, and topological data analysis. In that case coupling methods with new development of learning methods are certainly promising, in particular topology methods and the hierarchical deep learning. They could be for example used to find correlations between the evolution of single droplets and turbulent length scales, which is still a challenge in high-fidelity multiphase simulations.

A first joint project between RWTH Aachen University and FZ Jülich, which focused on the development of large-scale smart in-situ visualization/post-processing techniques in the field of multiphase simulations, showed significant improvement in terms of I/O speed and data handling for Petascale simulations.

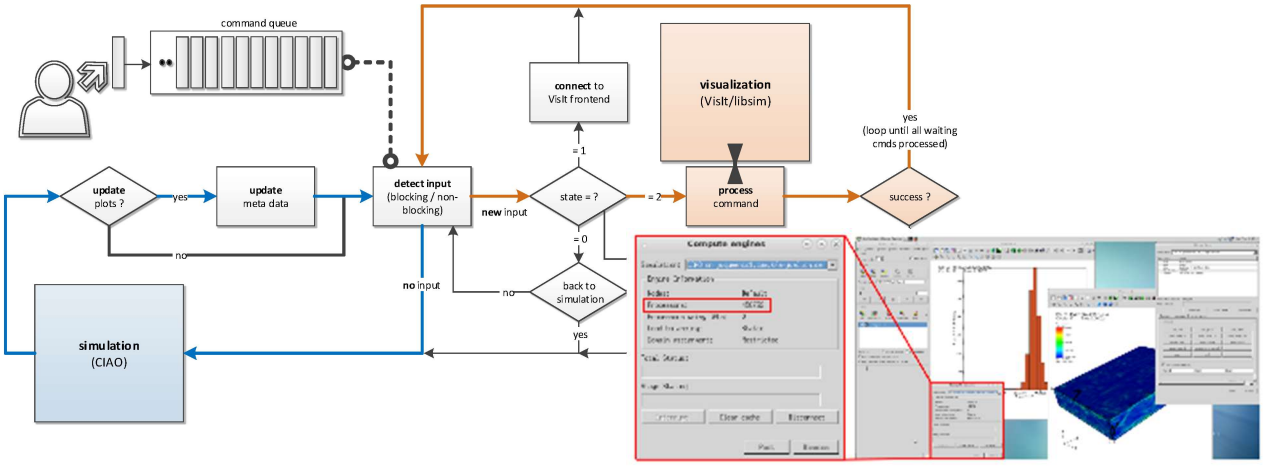

Figure 21 - insitu framework developed by Aachen Univ/ for large scale combustion simulations

Computational study and design of molecules and materials on the atomistic scale is essential in the chemical, pharmaceutical and materials sciences and industries. It requires rigorous, unbiased, and accurate theoretical treatment. While numerical approximations to the many-electron problem are available, their enormous computational cost requires HPC to overcome current limits in terms of system size, simulation length, and size of databases in high-throughput screening. Concomitant with the increasing availability of big databases for chemical compounds, crystal structures, and now electronic structure calculations themselves, machine learning models are being developed that interpolate between ab initio electronic structure calculations to accurately predict properties of new similar systems or to analyse high-volume data from simulations to uncover "hidden correlations" to gain new physical insights.

Finally, in the area of Life Sciences research, molecular simulation is a powerful tool to gain understanding of the structure and dynamic function of the basic building blocks of living organisms such as proteins, DNA, lipids, small molecules and up to the level of single cells. High-end compute infrastructures and highly scalable and efficient software packages are already capable of generating immense amounts of data e.g. by performing long time-scale simulations of multi-million particle systems or massive ensemble simulations of medium sized ones. Similarly, cryo-electron microscopy (cryo-em) methods have improved tremendously and are capable of elucidating the structure of large macromolecular complexes but their efficiency depends on the fast analysis of the terabytes of image data produced daily by the microscopes. Yet the tools for post-processing and analysis still lag in capabilities. Fostering research in the area will require development of the necessary software stack for exploitation of deep learning methods for analysis of multi-dimensional data in particular relevant to areas such as computational drug design, protein structure and function, cryo-em image processing etc.

4.1.3 Toward the convergence of scientific data analysis and machine/deep learning techniques

In 2015, inside the set of recommendations issued by the EESI2 project, one-third were related to smart data analysis of scientific applications. Based on the rationale that the deluge of data generated by large scale or ensemble/coupled multiphysics/multiscale simulations become impossible in a competitive time to process by current techniques, the recommendations were proposing to develop at the European scale frameworks for in-situ/in-transit data analysis as well as providing to such tools the possibility to identify on the fly pertinent (turbulent) structures.

In-situ/transit technique allows to benefit from data locality, over the different memory hierarchies, just after the data is computed for performing real-time and non-intrusive post-processing of the raw data and thus reduce I/O overheads and optimize energy by storing only refined data. Such efficient post-processing could lead in as reverse loop to a new class of efficient computational steering techniques, again able to reduce both time and energy to solution.

Implementing such on-the-fly post processing tools need also to be able to implement high-order low pass and high pass filters, data mining features of reduction, cross-correlation, pattern/structure/field lines reconstruction/detection, ordering, partitioning, compression of data as well as trajectory based flow feature tracking.

Artificial intelligence methods become more widely adopted since the 2000s with machine learning techniques driven by rise of big data, and more recently by deep learning techniques driven by massively parallel computational hardware and new algorithms using multi-layer neural networks. Such models by using successive computational layers that process data in a hierarchical fashion by applying on each step convolutional layers (filters) have been widely adopted in image, video, sound and speech processing. Scientific communities as well as companies like Google, Amazon, Facebook, nVIDIA, Microsoft … developed machine learning frameworks like Torch, SPARK, Mahout, TensorFlow, DMTK, Shogun, Caffe, Theano, Scikit-Learn, that now start to be used as well in Life Sciences or particle physics.

Europe possesses a strong expertise in these new techniques, both from academia with Inria or CNRS research teams in France, BSC in Spain, Technical University of Munich in Germany, or the Alan Turing Institute in UK, as well as in industry with the Facebook FAIR laboratory and the Sony CSL laboratory in Paris or the IBM R&D lab in Zurich. **It's now time to bridge**

this community with the HPC/scientific computing community in order to develop new massive data analytics techniques.

In the context of the BDEC conference about "Pathways to convergence" and the work conveyed by EXDCI WP3**, it is asked to the European Commission and the national funding bodies to organise an urgent joint call for proposal toward the convergence of scientific data analysis and machine/deep learning.** Following the first discussions with BDVA and ETP4HPC during the first EXDCI technical meeting in Barcelona, this call for proposal could be organised jointly with BDVA and could lead to data-intensive software demonstrators to be used on the upcoming Extreme Scale Demonstrators, made available in 2018/2019 by ETP4HPC.

By benefiting from the expertise of European teams in these fields, the objectives of this call, with around 3M€ funding, could be to bridge on up to 12 pilot projects from many scientific and industrial fields the skills of experts in domain science, deep/learning and HPC experts.

This mandatory cross-fertilisation of expertise could allow concretely during 2 or 3 years for each project to develop modern in-situ/in-transit post-processing techniques and assess the potential of machine/deep learning techniques for pertinent features detection in turbulent fluids, seismic processing, medical imaging, etc.

4.2 Development of new services toward urgent computing and link with scientific instruments

In a context of an explosion of the volume of data generated by large scale instruments (telescope, satellite, network of sensors, sequencers, …, IoT) the use of supercomputers becomes mandatory for being able to post process and value in a competitive time the information acquired. In a same way, being able to couple medical devices located into hospitals (scanners, microscopes, …) with supercomputers or being able to allocate on the fly a partition of a supercomputer for urgent computing decision making, in a case of major event (earthquakes, hurricanes, … pandemic propagation, …, (cyber)terrorism) is also now critical for both doctors and public bodies.
Developing these new services for European (like PRACE) or national HPC infrastructures is now a must, leading to an evolution of the operational procedures of the centres in one side but also to the efficient support by resource managers of complex workflows, co-scheduling of resources, smart (application-based) automatic checkpoint/restart and other features.
This recommendation was already mentioned into EESI2 reports and some actions have been implemented inside the FETHPC-02-2017 call for proposals (opening in April 2017) but WP3 experts are insisting in the importance of implementing such services.

4.3 Development of new Centers of Excellence in Europe

In 2015 the European Commission selected 8 (and then 9 mid 2016) Centres of Excellence (CoEs) for computing applications following the recent call under e-Infrastructures. They will help strengthen Europe's existing leadership in HPC (high-performance computing) applications and cover important areas like renewable energy, materials modelling and design, molecular and atomic modelling, climate change, Global System science, and bio-molecular research, and tools to improve HPC applications performance. They will also foster the

development of structured European scientific communities, reaching the critical mass in order to address the development of a new generation of HPC applications, able to take benefit from the (pre)Exascale systems.

A global funding of 42M€ has been awarded to 9 projects with between 4 to 5M€ over 3 to 4 years for each CoE. A second call of proposals is expected in the first semester of 2018 and will aim to consolidate successful CoEs and launch new ones in missing segments.

Experts from WP3 consider that this second round is highly welcomed, even if the budget available is below what EESI1 proposed few years ago (while talking about Co Design Centers instead of Centers of Excellence), both in terms of budget (more in the range of 50M€ per CoE) and duration (10 years instead of 4 years).

It is also important to notice that in the field of the US National Strategic Computing Initiative (NSCI) launched in august 2015 with the goal to restore US sovereignty in technology, infrastructure and applications (involving DoE, DoD, NSF and several agencies with a budget of 3 billion $ during 10 years), a major project called ECP (Exascale Computing Project) lead by DoE issued a first call for proposal toward applications enabling. In this call 22 proposals involving 45 research teams benefited from close to $40 million[9] budget for developing Exascale applications by 2023 in the field of fusion, renewable energies, combustion, high energy physics, nuclear physics, materials and cosmology. It is expected that more call for proposals will be available on a regular basis.

On November 11, 2016 the ECP again announced a new investment of $48 million to establish 4 new Co-Design Centers in the field of Online Data Analysis and Reduction at the Exascale (CODAR), Block-structured AMR numerical methods (developing a new framework called AMReX, in support to the 22 applications announced few months ago), Efficient Exascale Discretisation (CEED, including work on scalable meshers, discretisation models, …) and Particles Applications (CoPA, acting as a centralised clearinghouse for particle-based simulations by providing Exascale software platforms).

The following proposals are the result of a cross discussion between the experts of WP3 and also representatives from EXDCI WP2 and WP4 as well as CoE and BDVA ETP.

4.3.1 Engineering and industrial applications

Regarding the domains already covered by the first set of 9 Centers of Excellence the experts recommend to complete it with a new one related to engineering and industrial applications.
Europe is recognized as very strong in terms of industrial companies having in one side either an established strong roadmap toward Peta and Exascale (like TOTAL, Repsol, ENI, Shell in the Oil&Gas, EDF in field of energy, Porsche, Renault, JLR in the field of automotive, Airbus, Safran or BAE in aeronautics, etc.) or on the other side potential needs of advanced numerical simulation and HPC (from large companies to SMEs).
These companies as well as all the academia research teams involved in engineering applications (CFD, CAE, acoustics, electromagnetics, neutronics, etc.) could benefit from the establishment of a joint CoE federating the European ecosystem in order to reach a critical mass for the support of European engineering applications.
This CoE could start first with the support of CFD and turbulent applications and then expend its activity to other engineering domains. Sustainability of this CoE should be ensured by

[9] https://www.hpcwire.com/2016/09/07/exascale-computing-project-awards-39-8m-22-projects/

specific services provided to industry-like user support of (open source) software, licensing or specific tailored developments.

4.3.2 (Open-source) software sustainability

In July 2016, the US National Science Foundation (NSF) announcement the commitment of a budget of $35 million over 5 years to improve scientific software by the establishment of 2 new Scientific Software Innovation Institutes (S2I2).
NSF's Software Infrastructure for Sustained Innovation (S2I2) program, have been launched in 2010 to fund software research at multiple scales and aim to transform innovation in research and education into sustainable software resources that are integral to cyberinfrastructure.
Among the missions of these institutes, one is related to the industrialisation, the promotion and the long-term support of open source application developed by NSF.
Similar missions are also existing inside US DoE funded research laboratories, leading to the adoption by all the scientific community since decades of packages like BLAS, LAPACK, PETSc, HDF5 or others.

In Europe, few isolated similar initiatives are existing, the most visible one is the UK Software Sustainability Institute[10], founded to support the UK's research software community - a community that includes the majority of UK's researchers. The SSI is federating a network of UK applications developers and users, providing training and software evaluation services, issuing best practises guides, etc.

The European Commission and the national funding agencies are supporting the development of a vast ecosystem of scientific applications developed by researchers. When reaching a given level of maturity these applications are not properly industrialised because of lack of time (industrialising applications is 5 to 10 times longer than developing it), skills, interest, career objectives/recognition of the developers, turnover of people, rules, structures, etc. leading to a massive loss of potential usage by wider research communities and industry.
Regarding the use of scientific software by industry, many reports issued by the EESI2 or PlanetHPC (https://www.epcc.ed.ac.uk/planethpc) projects or by PRACE highlighted the fact that industry will use research software if it reaches a high level of industrialisation (or TRL, Technology Readiness Level), if industry is aware about their existence, trained and if there is a user support provided.

Despite the launch of the 9 CoE by EC in 2015, this service is not always present into the main missions of these centres and by definition this service could be leveraged by a transverse action across the CoE. Same remark applies to other H2020 funded research projects.

The recommendation of the WP3 experts is to setup a such transverse CoE in order to interface with the vertical thematic CoE for industrialising developed (open source) scientific software once it reaches a given level of TRL. This CoE could be also in a strong liaison with European (PRACE) and national HPC research infrastructures in order to work on training of best practises, availability and the dissemination of the software as well as the long-term first level user support.

Sustainability of this CoE should be ensured by specific services provided to industry like user support of (open source) software or tailored industrialisation of in-house software.

[10] https://www.software.ac.uk

4.3.3 High performance Data Analytics

This recommendation is an extension of the previous one related to the convergence of in-situ and machine/deep learning. In a context of a strong convergence between HPC and big data, generated from large scale instruments (sequencers, satellite, telescopes, accelerators, network of sensors, … and IoT) of from supercomputers (massive 3D complex simulations, coupled multi-scale multiphysics simulations, ensemble, uncertainties/optimisation studies, etc.), it is appearing mandatory to **provide specific services to scientific and industrial communities toward High Performance Data Analytics** (HPDA).
On the same way as the POP (Performance Optimisation and Productivity) CoE is already working for performance analysis and optimisation, this transverse CoE could work on assessing/auditing the needs of "client" communities in terms of data analytics/management and then provide solutions through Proofs of Concept (PoC) and training actions based on standard approaches when possible.
This CoE could be established in collaboration with the Big Data Value Association (BDVA) in order to foster synergies between the HPC community and the Big Data community.

5 Potential collaborations outside EU

Many libraries used in HPC contexts already combine both EU and non-EU teams or libraries, with most collaborations occurring between the EU and the USA. We can mention just a few of many examples:

- Most major scalable linear algebra libraries are US-based (PETSc, HYPRE, parts of Trilinos), with links to other major packages which may be non-us-based (PETSc can call MUMPS, developed by several French labs, or ViennaCL, developed at TU Wien); Some EU H2020 recent projects, such as NLAFET, also include some collaboration with US-based teams;
- In a similar fashion, several libraries, such as PETSc or Trilino's ZOLTAN partitioning library can call INRIA's PTScotch partitioning library as well as the US-based parMETIS library;
- US-based runtimes such as OpenMPI have strong links with the hwloc library, developed mostly by INRIA (France), …

For well-maintained and long-lived HPC oriented libraries such as those mentioned above, maintaining the international collaboration (or recognition of useful libraries) should be encouraged.

Many tools use visualization based on the VTK toolkit (US), and on the ParaView and VisIt tools (based on VTK) on which most in-situ visualisations today are built, which are predominantly US-based (though Kitware, which develops VTK, has offices and collaborations in France). For engineering codes, improving these libraries through feedback and collaboration is probably more important, useful, and realistic than developing a full-fledged competitor. EU labs can also work on individual VTK extensions or plugins which may become essential parts of these libraries.

We must not forget that as the HPC community is rather small, reaching critical mass for many tools often means collaborating around existing packages rather the developing complete competing packages. This does not mean the EU cannot be leaders in many areas, as with

software tools becoming larger and more complex, many major packages are assembled of many modules, and taking or keeping leadership in some of those modules is often possible. The competitive advantages may seem less immediate than with EU-only developments, but the bigger the EU teams share of work (often through non-EU restricted international collaborations) on state-of-the art international software tools, the more reactive those same teams may be when supporting EU-only projects.

Some domains seem less well covered (or well covered with less well known/established software) in the USA than in the EU, **so major projects there could provide leadership to Europe**: Two of these domains are code coupling, as detailed in the corresponding section and the development of a unified scalable Uncertainty Quantification/Optimisation framework. These 2 actions were already recommended by the EESI2 project (Toward European scalable code couplers and VVUQ (Verification, Validation and Uncertainty Quantification) and optimisation).

6 Conclusion

During this first year, WP3 worked actively on setting up 4 working groups and involving an active network of 40 experts across Europe in many scientific and industrial disciplines. A strong link with the newly created European Centers of Excellence has been also established, especially with 5 of them (ESIWACE, E-CAM, EoCoE, BioExcel and CompBioMED).
The 4 working groups produced a first set of reports describing scientific and technical challenges and highlighting specific and global recommendations.
For the next year, WP3 will enrol few more experts, collaborate with the Eurolab_4_HPC project, continue to establish connections with the other CoE, but also will extend relationships with WP2 in the field of the Extreme Scale Demonstrators, and the update of synchronised update of the ETP4HPC' Strategic Research Agenda (SRA) and the PRACE Scientific Case.

Annex 1 – List of WP3 experts

Industrial and engineering applications

Name	Institution	Area of Expertise
Yvan Fournier	EDF	CFD (focus on HPC, pre, post and coupling aspects, unstructured, semi-implicit)
Heinz Pitsch	RWTH Aachen University	CFD, combustion
Mathis Bode	RWTH Aachen University	CFD, combustion
Klaus Adams	TU Munich	Complex flows, low speed aerodynamics, multiphase flows,
Christian Stemmer	TU Munich	Complex flows, low speed aerodynamics, multiphase flows
Philippe Ricoux	Total	Reservoir Modeling, oil & gas
Norbert Kroll	DLR	Aero, External aerodynamics
Denis Veynante	Ecole Centrale Paris, CNRS	Combustion, turbulence

Weather, Climate and solid Earth Sciences

Name	Institution	Area of Expertise
ALOISIO Giovanni	Univ Salento - CCMC	Exascale Computing
ANDRE Jean-Claude	JCA Consultance	Atmosphere
ACOSTA Mario	BSC	
BAUER Peter	ECMWF	Exascale Computing
BUDICH	MPG	Meteorology
FILGUEIRA Rosa	BGS	Atmosphere
FIORE Sandro	Univ Salento - CCMC	Earth Sci.
KOMATITSCH Dimitri	CNRS/LMA	Exascale Computing
PINARDI Nadia	Univ. Bologna	Earth Sci.
SHAPIRO Nikolai	IPGP	Ocean
VALCKE Sophie	CERFACS	Earth Sci.

Fundamental Sciences

Name	Institution	Area of Expertise
Stefan Krieg	JSC	High Energy Physics
Allan Sacha Brun	CEA - Saclay	Astrophysics
Thierry Deutsch	CEA - Grenoble	Quantum Chemistry
Pascal Tremblin	CEA-Saclay	Physics/EoCoE
Maurizio Ottaviani	CEA-Cadarache	Fusion
Volker Springel	Garching, MPI Astrophysik	Astrophysics
Mike Payne	University of Cambridge	Quantum Chemistry/ECAM
Ulf Meißner	University of Bonn/ FZ Juelich	Hadron/Nuclear Physics
Chrisitan Zeitnitz	Wuppertal U.	Experimental HEP (LHC)
Paul Gibbon	KU Leuven/FZ Juelich	Laser Plasma Interaction/EoCoE
Andrea Mignone	Turino U.	computational astrophysics

Life Science & Health

Name	Institution	Area of Expertise
Rossen Apostolov	KTH Royal Institute of Technology	Molecular Simulations
David Wright	UCL	Molecular Simulations
Peter Coveney	UCL	Multiscale biological modelling
Erik Lindhal	KTH Royal Institute of Technology	Molecular biophysics and simulations
Modesto Orozco	IRB Insitute for Research in Biomedicine	Biomolecular simulations and Personal Medicine
Paolo Carloni	Juelich SCC	Electronic structure and enzymatic reactions
Sean Hill	INCF	Neuro-informatics
Marco Viceconti	University of Sheffield	Biomechanics
Blanca Rodriguez	University of Oxford	Cardiovascular modelling
Andrea Townsend-Nicholson	UCL	Cell signalling and molecular pharmacology
Kashif Sadiq	Heidelberg Institute for Theoretical Studies	Biomolecular simulations and Personal Medicine