

H2020-FETHPC-2014

Coordination of the HPC strategy



EXDCI

European eXtreme Data and Computing Initiative

Grant Agreement Number: FETHPC-671558

D3.2

EXDCI Inputs to the PRACE Scientific Case

Final

Version: 1.0

Author(s): S. Requena (GENCI), Y. Fournier (EDF), H. Pitsch (Univ. Aachen), M. Bode (Univ. Aachen), G. Aloisio (CCMC), JC. André (Jca Consultance), S. Brun (CEA), S. Krieg (JSC), R. Apostolov (KTH-BioExcel), P. Coveney (UCL-CompBioMED) and D. Wright (UCL - CompBioMED)

Date: 30.11.2017

D3.2

Project and Deliverable Information Sheet

EXDCI Project	Project Ref. №: FETHPC-671558	
	Project Title: European eXtreme Data and Computing Initiative	
	Project Web Site: http://www.exdci.eu	
	Deliverable ID: D3.2	
	Deliverable Nature: Report	
	Dissemination Level: PU	Contractual Date of Delivery: 30 / 11 / 2017
		Actual Date of Delivery: 30 / 11 / 2017
EC Project Officer: Evangelia MARKIDOU		

* - 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: EXDCI Inputs to the PRACE Scientific Case	
	ID: D3.2	
	Version: 1.0	Status: Final
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): EXDCI-Deliverable-D3.2.docx	
Authorship	Written by:	S. Requena (GENCI)
	Contributors:	Y. Fournier (EDF), H. Pitsch (RWTH Aachen), M. Bode (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), P. Gibbon (JSC), E. Molinari (MaX CoE) and G. Latu (CEA).
	Reviewed by:	Catherine Inglis (EPCC) and Katia Soto (PRACE)
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1 and then	03/10/2017	Draft	Initial and internal versions
0.5	12/11/2017	Draft	Final consolidation, sent to review
1.0	30/11/2017	Final	Final version of the report

Document Keywords

Keywords:	PRACE, Research Infrastructure, applications requirements, Centre of Excellence, use cases, breakthroughs and recommendations, convergence, HPDA, AI, BDVA
------------------	--

Copyright notices

© 2017 EXDCI Consortium Partners. All rights reserved. This document is a project document of the EXDCI project. All contents are reserved by default and may not be disclosed to third parties without the written consent of the EXDCI partners, except as mandated by the European Commission contract for reviewing and dissemination purposes.

All trademarks and other rights on third party products mentioned in this document are acknowledged as own by the respective holders.

Table of Contents

1	Introduction	15
2	Scientific (and industrial) challenges and roadmaps	20
2.1	Industrial and engineering applications.....	20
2.1.1	<i>Challenges on Aeronautics and Aerospace</i>	21
2.1.2	<i>Challenges in the automotive industry</i>	24
2.1.3	Challenges in the Oil and Gas	26
2.1.4	<i>Challenges in Power generation and nuclear plants</i>	29
2.1.5	Challenges in Combustion	31
2.1.6	HPC and data requirements	32
2.1.7	Recommendations in terms of R&D	33
2.1.8	Recommendations to PRACE RI	34
2.2	Weather, Climate and Solid Earth Sciences	37
2.2.1	Introduction	37
2.2.2	<i>Challenges in Earth's climate system</i>	41
2.2.3	<i>Challenges in oceanography and marine forecasting</i>	46
2.2.4	<i>Challenges in weather and air quality</i>	48
2.2.5	<i>Challenges in Solid Earth sciences</i>	50
2.2.6	<i>Common requirements and recommendations</i>	54
2.3	Fundamental Sciences	56
2.3.1	<i>Challenges on Nuclear Physics and QCD</i>	57
2.3.2	<i>Challenges on Plasma Physics</i>	58
2.3.3	<i>Challenges on Fusion</i>	59
2.3.4	<i>Challenges on Astrophysics and Cosmology</i>	61
2.3.5	<i>Challenges on Material Sciences</i>	64
2.3.6	<i>Common requirements and recommendations</i>	66
2.4	Life Science & Health	68
2.4.1	<i>Data driven bioscience</i>	68
2.4.2	<i>Biomolecular modelling and simulation</i>	70

2.4.3	<i>Biomedical simulation</i>	72
2.4.4	<i>Common requirements and recommendations</i>	75
3	Global recommendations.....	79
3.1	Convergence between in-situ/in-transit post processing techniques and machine/deep learning methods	79
3.1.1	<i>Introduction</i>	79
3.1.2	<i>Key issues and scientific and industrial data analysis challenges</i>	80
3.1.3	<i>Toward the convergence of scientific data analysis and machine/deep learning techniques</i>	82
3.2	Development of new services toward urgent computing and link with scientific instruments	83
3.3	Development of new Centres of Excellence in Europe.....	84
3.3.1	<i>Engineering and industrial applications</i>	85
3.3.2	<i>(Open-source) software sustainability</i>	85
3.3.3	<i>High performance Data Analytics</i>	86
3.4	Development of EU-wide standard tools including scalable solvers, meshers and code couplers, and optimisation and UQ frameworks	87
3.4.1	<i>Code Couplers</i>	87
3.4.2	<i>Unified frameworks for UQ and optimisation</i>	87
3.4.3	<i>Scalable solvers</i>	88
3.5	Development of training and education	89
4	Conclusions	90
5	Annex 1 – List of WP3 experts	91

List of Figures

Figure 1 - distribution of experts across the 4 working groups	16
Figure 2 - WP3 hardware and software recommendations to ETP4HPC' SRA	17
Figure 3 - Temperature field of a jet emanating from a coaxial chevron nozzle (<i>RWTH Aachen University Aerodynamics Institute</i>)	22
Figure 4 - roadmap of CFD needs of the aeronautics industry.....	24
Figure 5 - HPC needs of Renault.....	25
Figure 6 - FWI coupled with DL workflow (copyright Schlumberger)	28
Figure 7 - evolution of the needs of TOTAL in seismic modelling	28
Figure 8 – large-scale CFD simulations using Code_Saturne using complex geometries	30
Figure 9 - combustion chamber/turbine LES of the factor test rig (F. Duchaine et al, JARAHPC'16)	31
Figure 10 - Instantaneous emission mass fraction - LES of the SGT 100 methane air burner operated at DLR. (Stopper et al CF 2013) T. Jaravel A. Felden (T. Jaravel et al PROCI 2016)	32
Figure 11 - scaling of CIAO over 28 IBM BG/Q racks	35
Figure 12 - scaling of Code_Saturne over 1.835 million threads.....	35
Figure 13 - CAM atmospheric model scaling at 1km resolution on 1.5 million cores, awarded by the 2016 Gordon Bell Prize.....	39
Figure 14 - Performance of current weather forecast models on up to 120 hours.....	40
Figure 15 - HPC and data requirements between CMIP5 and CMIP6 exercises	44

Figure 16 - comparison of the scalability of different models with a target resolution of 1km and a production of 1 year/date.....	45
Figure 17 - Global view of EPECC FET challenges.....	46
Figure 18 - Simplified illustration of the number of compute cores and power required for single 10-day model forecast (lower curves) and 50-member ensemble forecast (upper curves) as a function of model resolution, given today's model code and compute technology. From Bauer et al., 2015.....	50
Figure 19 - Expected requirements in weather forecast by ECMWF.....	50
Figure 20 - first result of lossy data compression on massive 3D datasets, leading to a 26x gain in I/O bandwidth with the GYSELA code.....	61
Figure 21 - evolution of the size (number of particles) of dark matter simulations over the time	63
Figure 22 - Full data cycle for Earth, Planetary and Universe Sciences (copyright BDEC).....	64
Figure 23 - evolution of the cost of genome sequencing vs Moore' law	69
Figure 24 - overview of XIOS (developed by IPSL), an asynchronous in-transit I/O library for climate simulations with less than 6% overhead.....	81
Figure 25 - in-situ framework developed by Aachen Univ/ for large-scale combustion simulations ...	81
Figure 26 - Identified extreme climate events using supervised (left) and semi-supervised (right) deep learning. Green = ground truth, Red = predictions (confidence > 0.8). [NIPS 2017]	83

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] EESI-2 reports: <http://www.eesi-project.eu/ressources/documentation/#EESI-2-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 and 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 and 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>
- [46] [3] J. Demmel, L. Grigori, M. Hoemmen, and J. Langou, Communication-Optimal Parallel and Sequential QR and LU Factorizations, SIAM J. Sci. Comput. Vol. 34, No. 1, pp. A206A239, 2012, available since 2008.
- [47] [4] L. Grigori, J. Demmel, and H. Xiang, Communication avoiding Gaussian elimination, Proceedings of the IEEE/ACM SuperComputing SC08 Conference, November 2008.
- [48] [5] L. Grigori, J. Demmel, and H. Xiang, CALU: a communication optimal LU factorization algorithm, SIAM J. Matrix Anal. & Appl., 32, pp. 1317-1350, 2011
- [49] <http://www.exascale.org/bdec/sites/www.exascale.org/bdec/files/whitepapers/bdec2017pathways15Nov17.pdf>
- [50] www.acare4europe.org/sites/acare4europe.org/files/document/Vision%202020_0.pdf
- [51] <https://ec.europa.eu/transport/sites/transport/files/modes/air/doc/flightpath2050.pdf>
- [52] <https://trilinos.org>
- [53] <http://sites.onera.fr/cwipi/>

List of Acronyms and Abbreviations

AI	Artificial Intelligence
AISBL	Association Internationale Sans But Lucratif (International Non-for-Profit Association)
ANL	Argonne National Labs (USA)
AMR	Adaptive Mesh Refinement
BDEC	Big Data and Extreme-scale Computing
BDVA	Big Data Value Association
BSC	Barcelona Supercomputing Center (Spain)
CAE	Computer-aided engineering
CASP	Critical Assessment of protein Structure Prediction
CFD	Computational Fluid Dynamics
CMIP	Coupled Model Intercomparison Project
CNN	Convolutional neural networks
CoE	Centres of Excellence for Computing Applications
DNS	Direct Numerical Simulation (used in CFD)

D3.2

EXDCI inputs to the PRACE Scientific Case

DoE	Design of Experiments
DoW	Description of Work
DSL	Domain Specific Language
EC	European Commission
ECI	European Cloud Initiative
ECMWF	European Centre for Medium-range Weather Forecasts
EDI	European Data Infrastructure
EESI	European Exascale Software Initiative projects
ENES	European Network for Earth System modelling
EPECC	European Programme on Extreme Computing and Climate
EPOS	European Plate Observing System
ESA	European Space Agency
EsD	Extreme Scale Demonstrators
ESRF	European Synchrotron Radiation Facility
ESS	European Spallation Source
EOSC	European Open Science Cloud
ETP4HPC	European Technology Platform for HPC
ESGF	Earth System Grid Federation
EU	European Union
FET	Future and Emerging Technologies
FP7	EC Framework Programme 7
FP16	IEEE Floating Point representation using 16 bits, called also half precision
FP32	IEEE Floating Point representation using 32 bits, called also simple precision
FP64	IEEE Floating Point representation using 64 bits, called also double precision
FPGA	Field Programmable Gate Array
GDP	Growth Domestic Product
GPU	Graphical Processing Unit
H2020	Horizon 2020 – The EC Research and Innovation Programme in Europe
HPC	High Performance Computing
HPDA	High Performance Data Analytics
HPL	High Performance Linpack (benchmark used for the top500 ranking)
IDC	International Data Corporation (now called Hyperion)

D3.2

EXDCI inputs to the PRACE Scientific Case

IESP	International Exascale Software Project
INVG	Instituto Nazionale di Geofisica e Vulcanologia (National Institute of Geophysics and Volcanology)
I/O	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
LES	Large Eddy Simulations (used in CFD)
MHD	MagnetoHydroDynamics
MIC	Intel Manycore product line
NWP	Numerical Weather Prediction
OS	Operating System
POC	Proof of Concept
QCD	QuantumChromoDynamics
RANS	Reynolds Average Number Simulations (used in CFD)
R&D	Research and Development
RDA	Research Data Alliance
ROI	Return On Investment
SDN	Software Defined Networks
SHAPE	SME HPC Adoption Programme in Europe
SHS	Social and Historical Sciences
SIMD	Single Instruction Multiple Data model
SME	Small and Medium Enterprise
SRA	Strategic Research Agenda
SSC	PRACE Scientific Steering Committee
SWOT	Strengths, Weaknesses, Opportunities and Trends
TRL	Technology Readiness Level
UQ	Uncertainties Quantification
VVUQ	Verification, Validation and Uncertainties Quantification
WCES	Weather, Climate and solid Earth Sciences
WG	Working Group
WP	Work Package

D3.2

EXDCI inputs to the PRACE Scientific Case

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 by the PRACE Scientific Steering Committee (SSC) 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 and services.

The roadmapping activity of WP3 relies on documents produced by previous European projects such as EESI and EESI-2 (European Exascale Software Initiative, [5]), by PRACE on its second edition of the PRACE Scientific Case, inputs from the nine European Centres 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

After the production by EXDCI/WP3 in November 2016 of first inputs (D3.1), this final report summarizes the most 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 (more than 50% in the field of chemistry/materials 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, beyond the traditional model of theory, experiment and simulation, 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 multi-scale/multi-physics simulations), and the public World Wide Web.

This now opens the door to an accelerated convergence between HPC/HTC/Edge computing (directly in the instrument), big data (called high performance data analytics, HPDA) and artificial intelligence (AI) techniques, in order to be able to acquire, analyse, value and make available, in a competitive timescale, large amounts of refined data in a trustable, open (in some cases) and user-friendly way. With that goal, a recent communication from the European Commission in April 2016 advocated an ambitious plan for a European Cloud Initiative providing innovative data Cloud services (European Open Science Cloud) through a unified European-wide Data Infrastructure (EDI) to researchers from academia and industry as well as public entities. This vision has then been endorsed and amplified by 13 European countries² and the EC on a project called EuroHPC, using a new pooled financial instrument (Joint

¹ as defined in EXDCI's DoA

² Luxembourg, Italy, Spain, Germany, France, Netherlands, Portugal, Slovenia, Belgium, Poland, Switzerland, Greece and Croatia

Undertaking, with an indicative budget of 4.7 to 5.2 B€) to deploy (pre-)Exascale systems into a federated pan-European HPC/data infrastructure open to research, industry and public services, and develop HPC technology in Europe and HPC testbeds in areas such as Personalised Medicine, Smart Space and Civil Protection, Smart Mobility, connected and autonomous vehicles, Industry 4.0 and Smart Manufacturing, Advanced Materials, Fintech, Smart Agrifood and Smart City Applications.

The Exascale area, which is about to rise in the coming years, will bring a lot of disruptions mainly due to a double objective to handle efficiently this deluge of data and to face heavy energy limitations on the IT side, trying to cap the power envelope to an acceptable 20 to 30MW. If previous shifts from Giga to Tera and then from Tera to Petascale were quite smooth (even if some code adaptations for jumping from vector to massively scalar parallel were needed), the transition from Petascale to Exascale aiming to sustain a performance improvement 50 to 100x on real applications (and not anymore on HPL) will be more complex.

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 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 national HPC services) and services (training/education, 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 in some cases already available), communities have started to rewrite/modernise their applications with the support of recently created European Centres 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 such as 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 such as Domain Specific Languages (DSL), frameworks like Kokkos or Raja relying on smart underlying system software layers, or new languages like Julia 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 now need to be accelerated;
- Addressing the needs of the convergence between HPC and data will require rethinking and supporting end-to-end workflows as well as a new computing chain spanning edge computing (close to the instrument) to large-scale HPC systems, using novel network services such as Software Defined Networks (SDN);
- Future (converged) HPC and data e-infrastructures will need to support both capability and capacity simulations. Communities have been engaged in major efforts in

improving scalability of the applications beyond 100 000 cores (a lot of examples are provided in this report but also in previous EESI-2 reports), but the major fraction of workloads will consist of ensemble or coupled multi-scale/multi-physics applications. This leads to reformulating previous EESI-2 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 amounts 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 a joint call for proposals, bridging teams in domain science, HPC, applied maths and DL/ML for 10 to 12 lighthouse projects assessing the potential of such approaches on different cases.
- Developing features of HPC resource 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 Centres 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 such as 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 for many reasons are not industrialised and subsequently used. 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 such as 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 such as 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 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.

- Scale at the European level the development of mandatory standard tools for scientific communities such as scalable meshers, code couplers, numerical solvers or unified frameworks supporting Uncertainty Quantification and Optimisation. Such tools are critical for ensuring that research communities will be able to use – both in capacity and capability – future Exascale systems. In all these domains Europe owns strong assets but teams are too sparse and not coordinated enough in order to reach the critical mass; structuring the applied mathematics community into a single entity like a European Technology Platform (ETP) could be a very good way to reach these goals and promote such tools more widely across European scientific and industrial communities.
- Finally, as a last but critical recommendation, foster training and education in Europe, especially in new fields addressing the convergence between HPC, HPDA and AI.

1 Introduction

This document describes the final consolidation of the activity of the 4 working groups inside WP3: “Applications roadmap toward Exascale”. The goal of this work package is to setup and manage four working groups of applications 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 Exascale 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 such as aeronautics, automotive, oil & gas, energy or finance, to name a few, are making heavy use of HPC internally or through access to research infrastructures such as PRACE.

This WG has investigated and provided updated roadmaps 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 coming from large companies as well as SMEs and academia have also been involved. Special support has also been 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 in the preparation of the PRACE Scientific Case, and in relation with the EsiWACE CoE on "Weather and climate", this group addressed issues such as new climate models, dynamical cores, and future couplers, as well as knowledge compression, increased concurrency, complex workflows for ensemble simulations, uncertainty quantification, and scientific data management.

Contributors came from EsiWACE, ENES and from ECMWF and the national meteorological services it comprises. Other contributions have been received from 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 has further considered and updated the challenges in fundamental sciences and has taken a close look at 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 has organized the panel of experts into four main areas: Systems Biology, Genomics, Neuro-informatics, Molecular Simulations and Biomedical Simulations.

The group has been analysing the impact of Life Sciences Exascale computing in the society, economy and welfare state.

In order to provide inputs to the PRACE Scientific Steering Committee (SSC) in the elaboration of the third version of the PRACE Scientific Case, planned for end 2017, a fifth working group (D3.5, chaired by S. Requena/GENCI and by representatives of PRACE Scientific Steering Committee, SSC) was in charge of the global coordination of WP3 and the consolidation of the different roadmaps.

During the two years of EXDCI a total of 43 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.

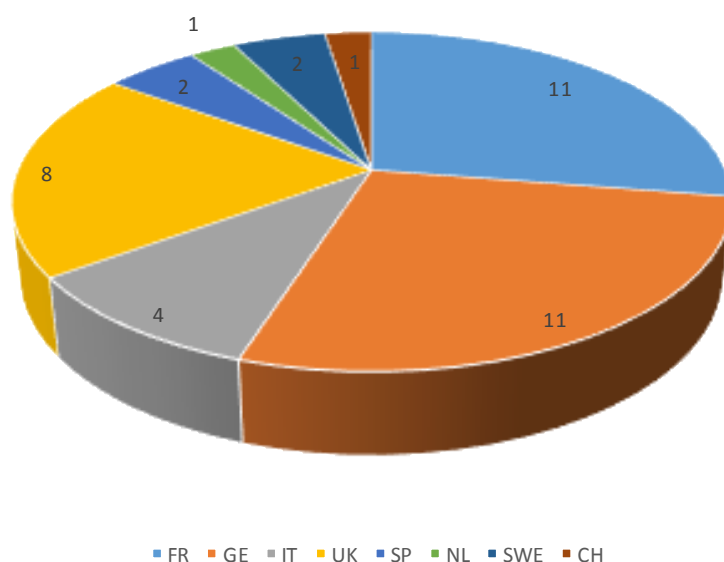


Figure 1 - distribution of experts across the 4 working groups

WP3 also had several interactions with other 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 the

two EXDCI technical meetings (organised on 21 and 22 September 2016 in Barcelona and in Bologna on 7 and 8 September 2017) where interlock sessions between experts from WP2, WP3 and WP4 allowed the exchange of current needs, visions and ways of collaboration for establishing joint roadmaps taking into account both applications, technology and cross-cutting requirements.

One of the last produced results of collaboration with WP2, toward the establishment of applications community hardware and software requirements for steering the deployment of the first ESD (Extreme Scale Demonstrators) planned for 2018, was this following table:

	WG3.1 Industrial & engineering apps	WG3.2 WCES	WG3.3 Fundamental Sciences	WG3.4 Life Sciences & Health
HPC System Architecture and Components	move to GPU or others accelerators virtualisation of resources	co design, GPU/MIC move hierarchical storage	GPU, MIC, FPGA, specialised HW ok with opt. libs, memory per core (astrophysics, fusion), mem & network BW	Large width vector units, low- latency networks, high- bandwidth and large memory; fast CPU<-> accelerator transfer rates, Heterogeneous acceleration, floating-point
System Software and Management	smart runtime systems, scalable couplers, data-aware schedulers	scalable couplers, smart scheduling	scalable couplers, mix of capacity/capability	urgent computing, link with instruments, co-location of compute & data, dynamic scheduling, support for workflows
Programming Environment	use of standards, sustainability	use of DSL, solution & performance portability	task programming / runtimes, use of DSL, standards	portability, fast code driven by python interfaces
Energy and Resiliency	scalable monitoring tools, data compression, application based FT	mixed precision	reduced precision, data compression	Distributed computing techniques to handle resiliency/fault tolerance
Balance Compute, I/O and Storage Performance	in-situ/in transit post processing	in-situ/in transit post processing active storage techniques	dynamic load balancing in mesh refinement and spectral element techniques	I/O driven, in memory database, Data-focused workflows, handling lots of small files in bioinformatics
Big Data and HPC usage Models	compression of data, remote viz, UQ/optimisation, ML/DL, mix of capacity/capability	UQ, compression of data, multi-site experiments to support multi-model analysis experiments, in- memory analytics, HPC- through-the cloud, ML/DL	remote data viz, ML/DL	convergence HPC-HTC, integration of HPDA tools (Hadoop, Spark, ...) inc ML/DL Cloud Computing security / privacy
Mathematics and algorithms for extreme scale HPC systems	ultra-scalable solvers, // in time, automatic/adaptive meshers, model order reduction, meshless and particle simulations, coupling stochastic and deterministic methods	parallelisation in time, ensemble simulations, auto tuned kernels	scalable solvers, adaptive meshers, // in time, FMM and h- matrices, auto tuned kernels	multiscale/physics workflows tools, ensemble simulation, model order reduction

Figure 2 - WP3 hardware and software recommendations to ETP4HPC' SRA

WP3 together with the CoEs has been involved in the elaboration of the third version of the ETP4HPC' SRA (Strategic Research Agenda³) published in November 2017, especially reflecting a section called application requirements. The following requirements associated with scientific/industrial challenges have been issued:

	Requirement	Challenge
General Requirements		
A1	Portability / Maintainability	Development of standards for programming heterogeneous architectures, deeper memory hierarchies, and new networking concepts; DSLs with smart system software layers
A2	Automatic profiling of applications	Develop automatic lightweight and accurate profiling tools for taking into account performance,

³ <http://www.etp4hpc.eu/en/sra.html>

		power consumption and data movement of applications
A3	Analysing large amounts of data	In-situ and in-transit scalable analytics using e.g. ML/DL techniques
A4	Cohabitation of different software stacks	Allow classical HPC software to interact with Data Analytics software stacks
A5	“Separation of concerns”: refactoring applications to separate the parts that are specialized for different HW.	Encapsulate kernel functionalities in self-contained HW-specific modules that should become part of domain specific libraries, or kernel system libraries when common to other application fields. The latter should be optimised, tuned and co-developed across different domains, and co-designed with HW vendors and specialists in performance optimisation and exascale programming paradigms.
Requirements industrial / engineering applications		
IND1	Ultra scalable solvers	Communication avoiding methods, parallel in time, encapsulating space and time combination of high order schemes (like DG) with implicit or semi-implicit time stepping allowing large time steps.
Requirements for weather and climate apps		
WC1	Disruptive numerical methods	Communication avoiding methods, parallel in time, encapsulating space and time, combination of high order schemes (e.g. DG) with implicit or semi-implicit time stepping allowing large time steps. Optimized communication for advection schemes; mixed- precision methods & AI techniques
WC2	Efficient coupling of capacity and throughput computing	Perform ultra-large $O(100-1000)$ ensemble calculations implying strong I/O systems with novel storage layouts (beyond POSIX) and support to resilient (fault tolerant) end to end workflows
Requirements for biomolecular research		
BM1	To deal with millions of independent jobs	Efficiently run a huge number of jobs deal with their input and output data
BM2	To efficiently simulate the long- range electrostatics treatment.	Efficient and scalable 3D FFT or alternatives
BM3	Simulating diffusion processes within a hybrid resolution (multiscale) approach.	Adaptive load balancing schemes
Requirements for energy apps		
E1	Meteo or hydro modelling using combined elliptic/parabolic equation systems (eg: incompressible Navier-Stokes equations for wind turbines, gyrokinetic or full Maxwell equations for fusion application)	Scalable algebraic solvers (e.g. multigrid), capable of exploiting accelerator hardware (see also WC1)
E2	Efficient coupling of capacity and throughput computing	Perform ultra-large $O(100-1000)$ ensemble calculations to generated probabilistic power forecasts; batch systems and resource managers tuned for this coupling

Requirements for fundamental sciences		
FS1	Immersive (remote) visualisation	Develop scalable post processing tools able to visualise remotely in an immersive way massive amount of data
FS2	Extreme level of scalability and optimal exploitation of the available compute resources	A system software that facilitates optimisation close to the hardware to allow, e.g., for very low-latency communication
Requirements for global system science		
GSS1	High-speed data transfer of diverse data sources, e.g. data streams but also of the results of simulations in a workflow execution of different simulations.	Develop complex data management techniques to couple HPC and HPDA components to efficiently process data sets (generated and from external sources such as streams)
GSS2	Providing the necessary means to execute a workflow (of partially parallel running) simulations, with potential strong dependencies on output data of simulations within the workflow.	Workflow scheduling on single system
Requirements for material apps		
M1	Enabling high-throughput calculation, automatic storage of data, and sharing of data and workflows	Workflow and data management systems
M2	Improve basic or kernel libraries, e.g. linear algebra and FFT, to optimize for the specific sizes that are needed. This requires co-design and closer interaction with library developers.	Better parallel implementation of 3DFFT (or better HW supporting its memory/communication pattern), new iteration loop for dense Hermitian matrix eigenvectors.

As stated in the EXDCI DoA, WP3 established connections with the Centres of Excellence (CoE) as follows:

- ESiWACE⁴ the Centre of Excellence in Simulation of Weather and Climate in Europe is represented in WG3.2;
- EoCoE⁵ the Centre of Excellence in the field of Energy, E-CAM⁶ the Centre of Excellence in the field of molecular dynamics, quantum dynamics and electronics structure modelling, NoMAD⁷, a Centre of Excellence in the field of novel materials discovery and MaX⁸ in the field of materials are involved in WG3.3;
- As planned in the DoA, BioExcel⁹ (a Centre of Excellence in biomolecular simulation) and CompBioMED¹⁰ (a Centre of Excellence in biomedicine, recently created in September 2016) are now chairing WG3.4.

WP3 together with WP2 collaborated with the BDVA ETP [7] in the field of big data, and with the EuroLab-4-HPC H2020 project in order to address issues related to big data and long-term applications roadmaps.

⁴ <https://www.esiwace.eu>

⁵ <http://www.eocoe.eu>

⁶ <https://www.e-cam2020.eu>

⁷ <https://nomad-coe.eu>

⁸ <http://www.max-centre.eu>

⁹ <http://bioexcel.eu>

¹⁰ <http://www.compbiomed.eu>

With BDVA, WP3 experts worked on the establishment of several joint use cases of big data generation/analysis and exploitation in the field of scientific computing and large-scale scientific instruments. This collaboration will be continued in the context of the EXDCI-2 project.

Finally, in the field of international collaboration, WP3 experts contributed to the “Pathways to Convergence” report [49], published during the SC’17 conference (Nov 2017, Denver).

2 Scientific (and industrial) challenges and roadmaps

2.1 Industrial and engineering applications

Working group 3.1 focuses on industrial and engineering applications. 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.

Many industrial and engineering applications concern improvement of mature technologies (power plants, industrial processes, aircraft, automobiles, transportation, ...), which have been developed for a very long time. Though similarities exist with other type of applications, this has strong impacts on several aspects:

- Compared to other scientific domains, there is often less consensus or presence of “community codes”, as many scientific codes and tools have been developed by competing partners, often at first for simulation of specialized aspects, and use of these tools may in some cases be included in licensing or regulatory processes concerning the technologies which may be supported by these tools.
- For the same reasons, many tools are long-lived, and ensuring a long-life cycle, reliability and reproducibility are often major deciding factors in their development process. Often, only a fraction of development resources may be allotted to high performance computing aspects.
- In many cases, code development and expertise in a given domain are closely linked; code development represents a capitalization of domain expertise, so collaboration is often possible on some aspects of a code, but rarely on all.
- 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”.

Further increasing efficiencies or performances will require the exploitation of the potential of effects which are currently mostly neglected (e.g. unsteadiness, even smaller scales, full simulations of coupled systems, etc.). One consequence will be long runtimes for unsteady simulations and thus large amounts of data. Furthermore, coupling of several reduced sub-simulations with different resolutions need to be developed, as well as consistent numerical methods.

In the future, more multi-physics / multi-scale simulations will be needed. For example, today fluid mechanics codes generally retain simplified wall interaction modelling while interactions between flows and structures and between flows and materials interactions, in the largest meaning, are key factors in resistance and aging of systems, especially with reacting and/or hot flows. A nice way to address such simulations is code coupling allowed by massively parallel machines but this faces two difficulties: (i) coupling of existing codes, repository of large amount of know-how but developed independently by different communities; (ii) efficiency and reliability of coupling procedures.

Also, these applications usually require sensitivity analysis to many factors, including model and numeric choices in addition to operating conditions, leading to many simulations for a given application. Today, most challenging simulations (for example development of combustion instabilities, multi-cylinder internal combustion engines, ignition of a realistic gas turbine combustion chamber in aeronautics, or design of nuclear plant component with longer lifetime) require “capability-level” computational resources and can only contribute to feasibility and qualitative studies. Obtaining trusted quantitative values would require much higher computational capacity.

2.1.1 *Challenges on Aeronautics and Aerospace*

Aeronautics and Aerospace are industrial fields with a traditional affinity to HPC computing going back 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 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 report from ACARE (Advisory Council for Aeronautics Research in Europe)[50] and the Flightpath 2050 report “Europe's Vision for Aviation” [51] are expecting, by 2050, a CO₂ reduction of 75%, a NO_x reduction of 90%, a reduction of 65% in perceived aircraft noise (compared to a typical aircraft in the year 2000), an 80% reduction in the accident

rate – while air companies are expecting reliable and less kerosene-hungry planes with a 3-fold increase in traffic and 99% of flights arriving within 15 minutes 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 (called also virtual or digital twin)** 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 realisability issues can now be tackled using HPC methods and systems. The continuing effort to increase HPC methods have rendered methods such as RANS (Reynolds Average Numerical Simulation) as everyday tools in industrial applications in the past. With increasing HPC capabilities on an 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 in order to jointly investigate issues which have been examined separately in the past (e.g. combustion, soot formation, fluid-structure interaction, propulsion systems and noise generation for aeroplane propulsion systems). Research efforts have shown that entire aeroplanes 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-the-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 across several orders of magnitude. Further improvement of low-level (e.g. RANS) tools can be expected for the production environment through the advancement of non-linear interdisciplinary research in aeronautical and aerospace sciences.

As an example, in the field of noise simulation reported by the RWTH Aachen University Aerodynamics Institute, the design of new quieter planes using “chevron” nozzles requires massive DNS simulations using roughly 2 billion points in a computational mesh running on up to 180,000 cores on the Cray XC40 Hazel Hen system located at the High-Performance Computing Centre Stuttgart (HLRS).

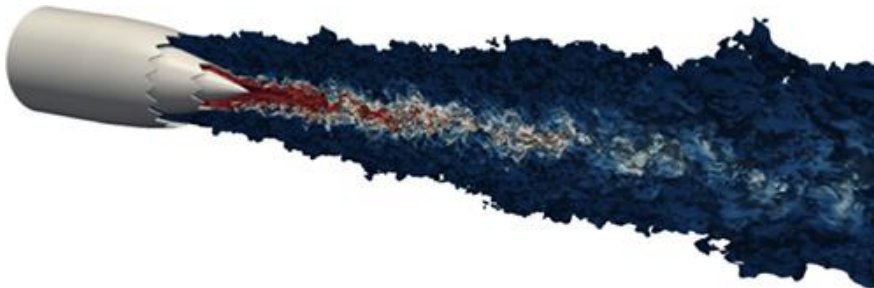


Figure 3 - Temperature field of a jet emanating from a coaxial chevron nozzle
(RWTH Aachen University Aerodynamics Institute)

In order to gain greater insights into turbulence processes and properties of particle-laden flows, the team will need somewhere between 15 and 20 billion points in its simulation mesh. With further increases in computational power, the demand for high memory bandwidth will play an

increasingly crucial role in order to scale codes to future exascale machines. The team knows that it will need faster I/O and data compression techniques to help cope with more complex machine setups. The team relies heavily on visualization resources, and with the advent of next-generation hardware, it will be important to have visualization hardware that has fast data access, including the possibility for *in situ* visualization.

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 the next figure was provided by NASA two years ago, and widely adopted by the entire aeronautics/aerospace community, is still valid and demonstrates **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.

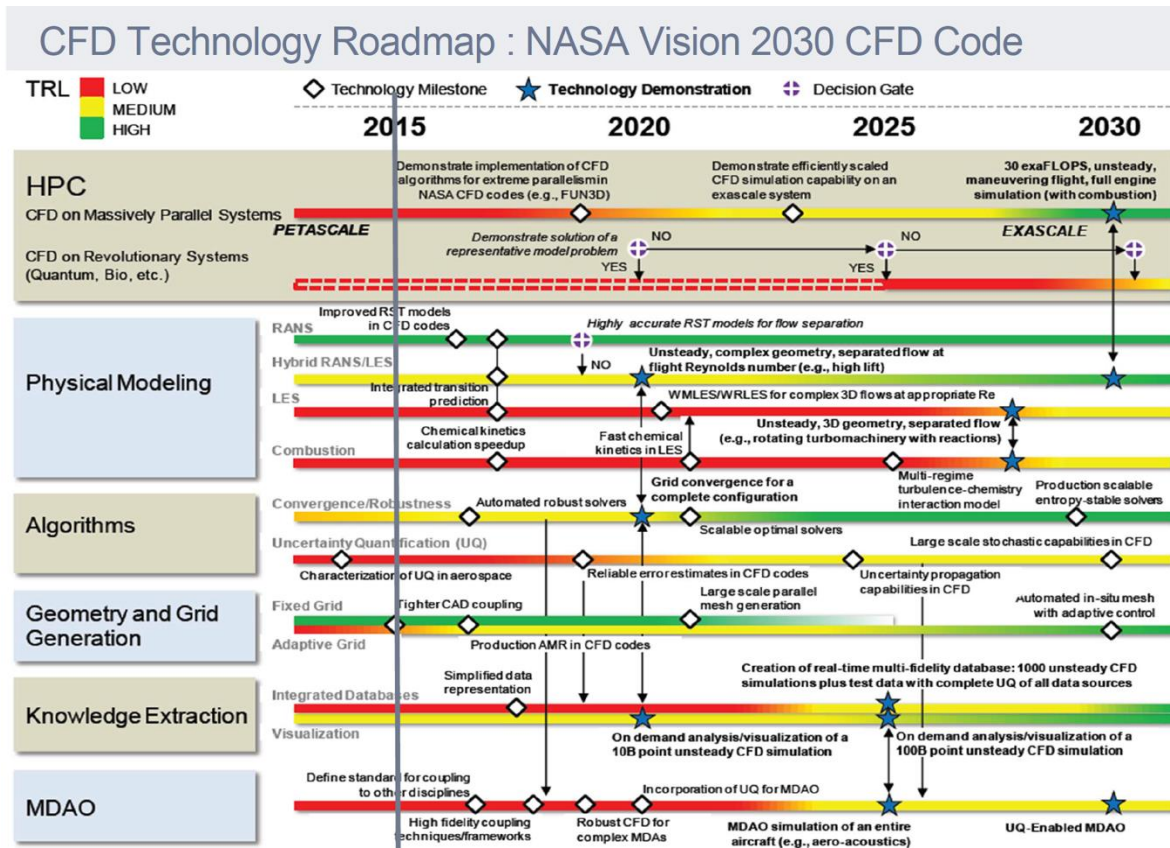


Figure 4 - roadmap of CFD needs of the aeronautics industry

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 costs several million dollars. Virtual design saves both cost and time, and improving the trust at each level of that design reduces 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-driven 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.

Another strong domain using HPC in the automotive industry is crash simulation 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 each customer the real performance of his car (and not only the data for the best-selling version as required by current regulations). The physical validation of all possible versions is not possible, so numerical simulation becomes mandatory;
- The design process needs to use systematic optimization with bigger and bigger studies. The current optimization algorithms last for 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 using potentially artificial intelligence. 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 for a lot of crash simulations. Model reduction uses all available data from crash simulation and will be seen as one of the biggest contributors to increased use of HPC (horizon: 5 years).

As a consequence, companies such as Renault-Nissan who used PRACE resources in 2013 for massive crash optimisation studies (using more than 41 million core hours), **started to re-engage in the use of HPC** and now have developed aggressive roadmap in both crash simulations, aerodynamics, noise simulation (engine, wind, road) and combustion, as seen in the following figure (values are Tflops installed at Renault).

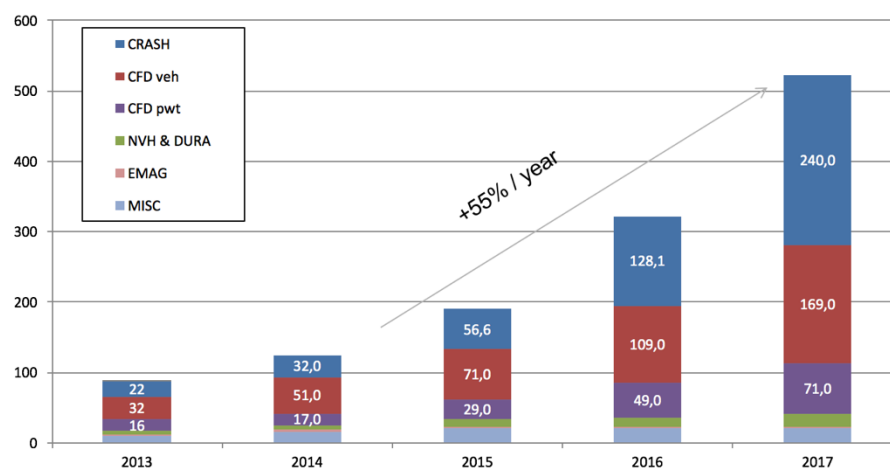


Figure 5 - HPC needs of Renault

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 (such as 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.

As an example, during the summer of 2017, Renault announced an investment of 25 M€ for a new Autonomous Driving centre called ROADS (for Renault Optimisation Autonomous Driving Simulator), to be launched in 2019, integrating a new HPC/AI simulator. This will help Renault to comply with new safety regulations obliging car makers to provide in each car 30 (instead of 12) advanced driver-assistance systems (ADAS) with an increased number of test phases.

Renault announced their intention to expand their HPC facilities from 500 TFlops now to 2 PFlops in 2019 in order to simulate billions of kilometres for each autonomous car usage (using a virtual twin approach), reducing to only a few million kilometres the final validation using a physical vehicle on roads.

The autonomous driving will come gradually with hands-on, hands-off, eyes-off and at the end even mind-off. The complexity between the first step: hands-on, to the last one: mind-off, is increasing by a factor of more than 1000 with an extremely low failure rate per hour obviously 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 controls while it could be impossible for a driver.

These challenges are illustrating a strong convergence between HPC and embedded systems (performing edge computing at the first level) 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, aeroplanes 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 shale gases) by using more and more precise seismic algorithms and more accurate reservoir modelling methods. Oil & Gas has become the second largest profitable market for HPC (just after Finance) with, according to the International Data Corporation (IDC, Hyperion now), an increase of the Compound annual growth rate (CAGR) of 9.2% between 2012 and 2017. Most of the major Oil & Gas companies or contractors own multi-Petascale HPC resources including:

TOTAL (with Pangea, a 6.7 PFlops system and in 2018 a 30 PFlops system expected), ENI,

BP (with a new 2PFlops manycore system installed 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 new 2.8 PFlops converged system addressing HPC and AI needs), CGG-Veritas or Schlumberger for the major contractors have also distributed HPC datacentres with accumulated performance beyond 10 PFlops.

Most notable is that the first non-public systems in the Top 500 lists (www.top500.org) are owned by these companies.

Such companies have 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.

With the massive increase of seismic sensors and the recording of data from heterogeneous format and sources, artificial intelligence using machine/deep learning is a promising path for classification, segmentation, pattern recognition, exploration of temporal series, ...

In Oil & Gas, this will go over domains such as seismic interpretation, real time analysis of data from wells in production, bio-stratigraphic analysis, analysis of satellite images in case of oil leakage, and forecast of production, to well planning before production into existing fields, smart forecast of production, anticipation of failures, ... from upstream to downstream.

In consequence, the rise of new hardware architectures providing AI or neuromorphic computing will provide a promising short-term perspective, while more mid-term technologies such as quantum computing are already on the radar of Oil & Gas companies. The acceleration of AI problems through quantum computing will be one of the first fields of application but porting numerical methods such as FFTs, Darcy laws for flow modelling through a porous medium (used in basin and reservoir modelling) will be a real challenge. **This will raise in a few years the question about rewriting HPC applications to be scalable beyond 300+ PFlops sustained (if possible) or starting now to think about using such applications in the future with quantum computing.**

Recently multiple research results provided by PGS and Schlumberger highlighted the interesting coupling between HPC and AI techniques for full-waveform inversion. In geological areas such as the Gulf of Mexico, however, reconstructing complex salt geobodies poses a huge challenge to FWI due to the absence of low frequencies in the data needed to resolve such features. A skilled seismic interpreter has to interpret these geobodies and manually insert them into the earth model and repeat this process several times in the earth model building workflow. Such companies investigate the use of deep learning algorithms, mainly Convolutional neural networks (CNNs), to generate useful prior models for full-waveform inversion by learning features relevant to earth model building from a seismic image.

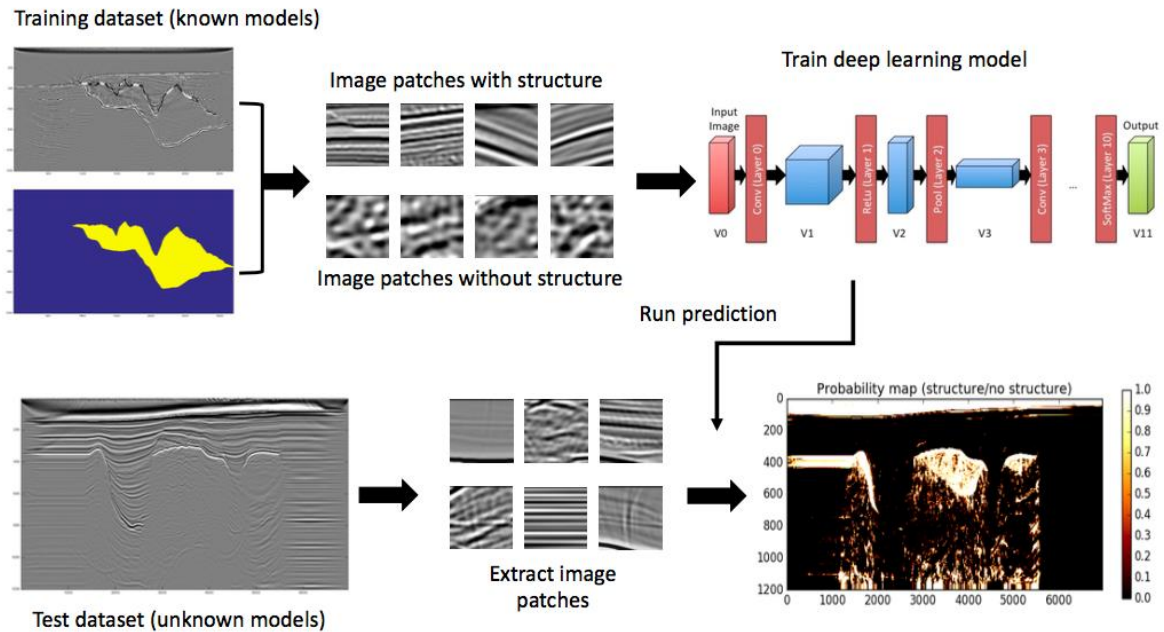


Figure 6 - FWI coupled with DL workflow (copyright Schlumberger)

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

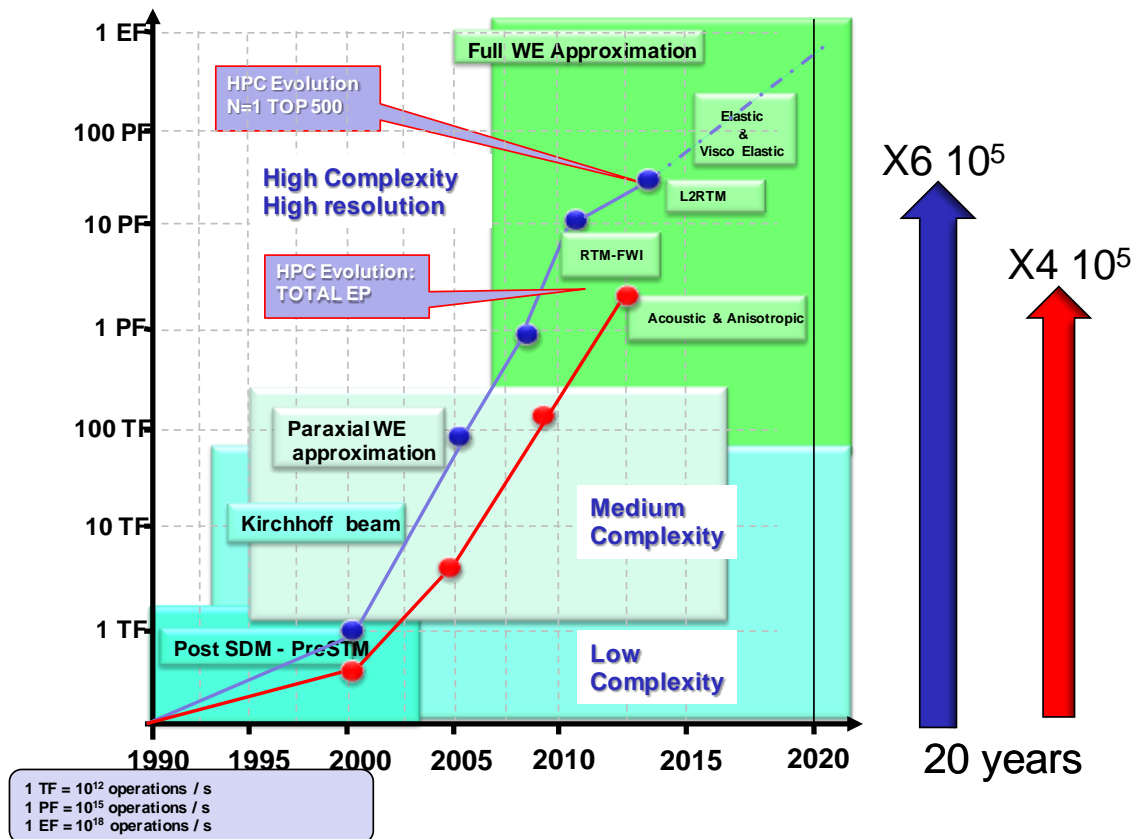


Figure 7 - evolution of the needs of TOTAL in seismic modelling

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 also starting to become crucial in performing **real time simulations** by coupling simulations and inputs from networks of sensors or seismic acquisitions when drilling a well or operating it in order to be able to better steer the process and in consequence to optimise operations and reduce costs and risks.

TOTAL is also starting to tackle the merging of data-driven and physics-driven modelling to optimize production of mature hydrocarbon reservoirs. The approach is based on using **simplified physical models, combined with neural networks or more classical machine learning approaches**. This is intended to allow testing of the impact of many parameters and scenarios to guide optimization of the reservoir, increasing the reliability of simplified models through the use of data. This is one example of possible HPC-Big data convergence, and will require new research.

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 reactor (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.

Immediate needs in the field require (unstructured) meshes in the billion-cell range. Studies in the near future could easily increase by 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).**

For renewables, though safety aspects are less of an issue, optimizing production means placement requires sensitivity studies with very high computational requirements, and optimizing maintenance aspects also requires high fidelity CFD computations.

Developers of several CFD tools seem to reach the agreement that due to CFL convergence criteria (Courant–Friedrichs–Lewy) 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 other means to handle geometric details (through some form 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 multi-physics codes at all scales**, from 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 larger, or slightly larger, resources than even CFD. Most approaches lend 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.

Recent large-scale industrial computations in CFD toward a full virtual twin nuclear reactor have required more than 1 billion unstructured cells. One example is a study of pressurized water reactor guide tubes by EDF, using the *Code_Saturne* CFD using LES with more than one billion cells, requiring a run over several weeks on 120,000 cores.

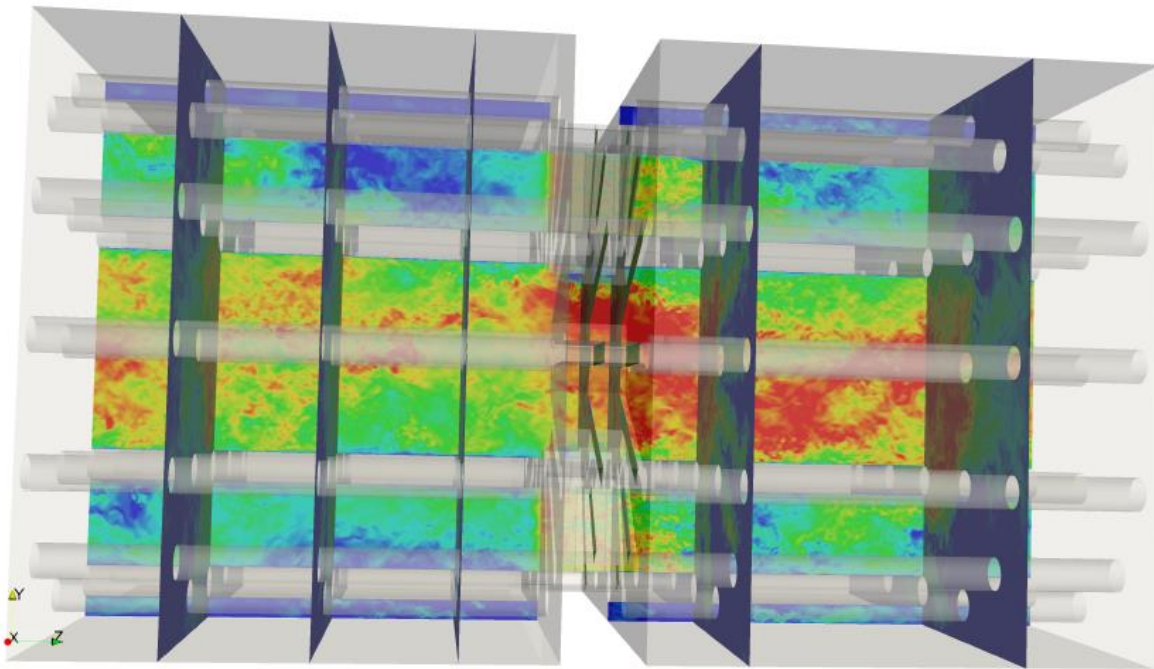


Figure 8 – large-scale CFD simulations using Code_Saturne using complex geometries

Computations on 1.6 billion cells have also been run by STFC Daresbury and EDF Energy UK using RANS models, for the study of sleeve leakage gas impact on a British AGR (Advanced Gas-Cooled Reactor) fuel assembly. These studies required a complex mesh, built of many sub-meshes, which are joined in parallel [18]. Note that for studies using more precise turbulence and heat flux models, meshes beyond 5 billion cells would be necessary.

2.1.5 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 formations can only be predicted accurately when detailed chemical schemes are considered¹¹. 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, and not entirely accessible by measurements.

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

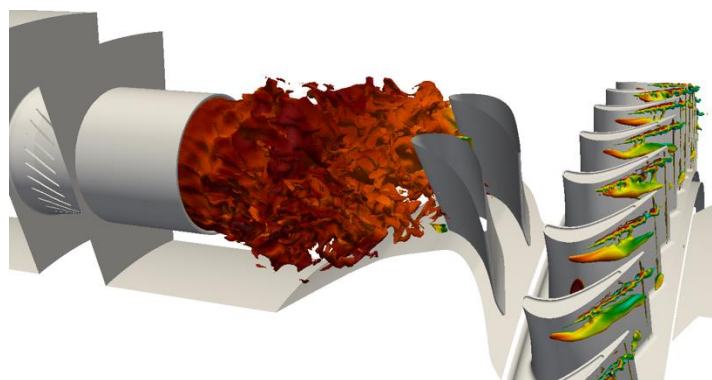


Figure 9 - combustion chamber/turbine LES of the factor test rig (F. Duchaine et al, JARAHPC'16)

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 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

¹¹ Note that usual pollutants (carbon monoxide, nitrogen oxides, ...) are present, but dangerous for health and environment, in very small amounts (tens or hundreds ppm) compared to main species (few to tens of per cents). Their prediction then requires precision (chemical schemes, modelling, numerical schemes).

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.

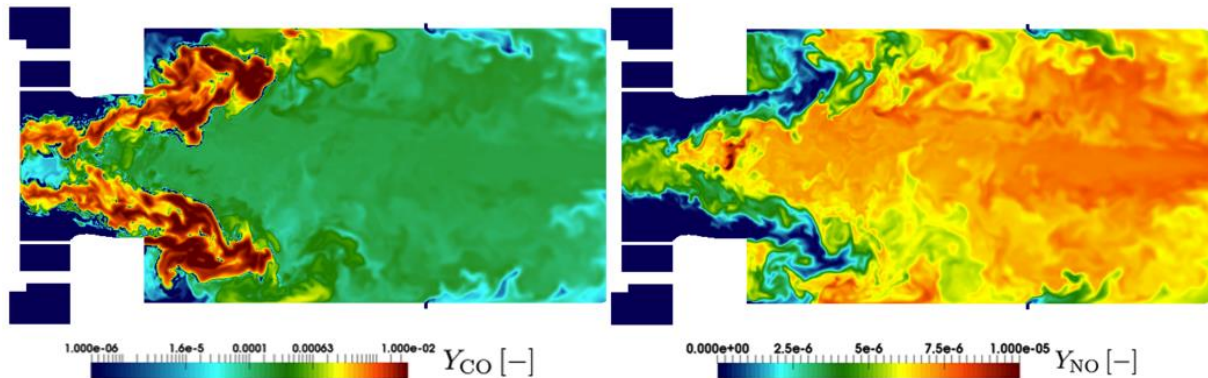


Figure 10 - Instantaneous emission mass fraction - LES of the SGT 100 methane air burner operated at DLR. (Stopper et al CF 2013) T. Jaravel A. Felden (T. Jaravel et al PROCI 2016)

Detailed combustion simulations are generally performed for aeronautic burners (see illustrations above) or internal combustion car engines. Both are typically about ten centimetres in size.

A challenge concerns systems with a large relative length scale, so while it is difficult to simulate situations such as industrial furnaces (tens of metres long), fires (forest fires may run over kilometres and have an atmospheric impact for thousands of kilometres), and deflagration-to-detonation transition following reactant leakages in plants, small devices with high Reynolds numbers may be even more challenging.

2.1.6 HPC and data requirements

As some (pre)Exascale architectures are already known (and in some cases already available), communities have started to rewrite/modernise their applications with the support of recently created European Centres 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 such as OpenMP (and to a lesser extent OpenACC) which ease and secure the shift to novel architectures in a context where applications last for decades when HPC architectures could last years; but although porting codes is usually not difficult, obtaining good and portable performance on these systems often can be difficult. Several factors make optimization difficult: one being adaptation of algorithmic choices to the hardware, as mostly compute-bound algorithms benefit from new architectures, while memory-bound ones do not so much; another being that for many codes of a certain size, optimizing a significant portion of the code may take longer than the lifetime of the hardware, which is often less than 5 years, compared to the 20 or more years for codes.

Some communities started also to explore novel programming approaches such as **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 now need to be accelerated.

Future (converged) HPC and data e-infrastructures will need to **support both capability and capacity simulations**. Communities have been engaged in major efforts in improving scalability of the applications beyond 100 000 cores (a lot of examples are provided in this report but also in previous EESI-2 reports), EESI-2 but the major fraction of workloads will consist of ensemble or coupled multi-scale/multi-physics applications.

This leads to reformulating previous EESI-2 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.

2.1.7 *Recommendations in terms of R&D*

Code couplers

The EESI-2 project mentioned a need for code couplers that are strong and unified at the European scale. 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, with 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 of Trilinos [52]), but these libraries still seem quite complex to use. However, 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 specifications/requirements in the past, and work together on other subjects, and OpenPALM's CWIPI [53] 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 some multi-layered specifications 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.

In-situ analytics

In-situ/transit techniques take 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, 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.

Also, although in-situ tools are increasingly being used, lack of standardisation increases the cost of experimenting with different toolsets. While standardisation opportunities may be

limited at this level of maturity, at least encouraging broad initiatives in this direction could be helpful.

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.

There has been significant progress in the in-situ visualization area, with some of the codes used or 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 FZ Jülich's JUSITU library (developed in order to help developers enabling in-situ visualization in their codes without knowing all details about the underlying libraries), allowing both VisIt's libsim and ParaView's Catalyst libraries. This approach still leads to a more complex software stack, for which developers of engineering tools need to evaluate not just the lifecycle in installation of the included tools, but also those of the wrapper library, so it does not replace standardization.

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.

2.1.8 *Recommendations to PRACE RI*

Experts from WG3.1 expect that PRACE could participate in strengthening the European HPC industrial ecosystem in several related aspects:

Promotion of highest scaling codes

Promoting the current highest-scaling codes may help potential users choose applications with good scaling. Examples of such promotion include Jülich's High Q club, based on codes scaling on the full JUQUEEN 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 anyway). 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.

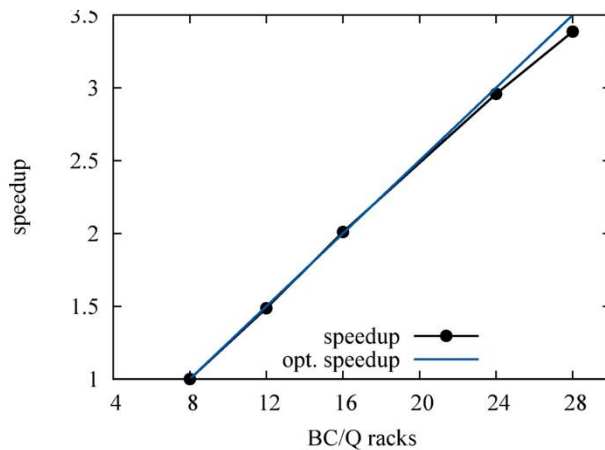


Figure 11 - scaling of CIAO over 28 IBM BG/Q racks

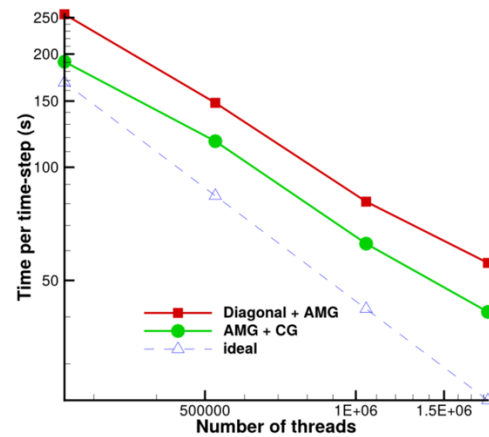


Figure 12 - scaling of Code_Saturne over 1.835 million threads

In the US in September 2016, the Department of Energy (DOE)'s Advanced Manufacturing Office announced up to \$3 million in available funding for manufacturers to use high-performance computing resources at the DOE 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 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. DoE funded in early 2017 a total of 13 projects for this third round of funding including projects from large groups like Ford, Samsung, GE, Applied Materials and also from SME like Arconic (casting), 7AC Technologies (air conditioning) or Sierra Energy (gasification). Selected projects received up to \$300,000 to support access to supercomputers and experts at the partnering national labs.

Promotion of sustainable software

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 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 low in terms of software development aspects, though more than 90% of researchers surveyed use research software.

However, in engineering domains, codes developed for industry may suffer somewhat less from developer turnover than some research codes, since it is likely that at least part of the development teams will have a lower turnover than groups of students or post-docs. Nevertheless, 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 would not need to know all details of the software, and consequently the training effort would be reduced. Due to increasing complexity, this approach might become necessary in any case.

Build & Productivity & Debug tools

Debugging tools have made good progress in recent years, and instrumentation of major compilers has improved significantly, allowing many errors to be caught 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 behind. 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 together with LLVM, and so build/version compatibility issues may arise easily, limiting the tool's practical availability.

Programming Languages and smart runtimes

Most engineering codes use Fortran, C or C++, together 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 is 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 EU-funded FETHPC project INTERTWinE is working in this direction [43].

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 HPC computations to be run 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. When 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. To that end, the EU project NEXTGenIO [44] is looking at this with "data-aware" schedulers implemented under SLURM.

Requiring unification of job managers across PRACE RI machines could provide impetus towards such standardization.

Databases and repositories

In engineering, huge data volumes are often generated by simulations, processed for a given objective and then lost, while some could be re-used for other analysis, sometimes by other

teams. This is especially true for, but not restricted to, DNS databases. It would be interesting to think about archives and availability to scientific communities of existing data.

In this regard, it seems scientific communities, such as the climate community, are much more advanced than engineering communities, where many databases are considered proprietary and not shared. This results in the fact that even when agreements for sharing exist, the infrastructures and standardization that would enable such sharing do not.

2.2 Weather, Climate and Solid Earth Sciences

2.2.1 Introduction

Weather, Climatology and solid Earth Sciences (WCES) encompass a wide range of disciplines from the study of the atmosphere, the oceans and the biosphere to issues related to the solid part of the planet. These Earth system sciences address many important societal issues, from weather prediction to air quality, ocean prediction and climate change to natural hazards such as meteorological, seismic, volcanic and tsunami hazards.

The **losses associated with natural catastrophes** are in the range from 100 to 400 B€ per year, among which 85% of the events are hydro-meteorological ones¹². The year 2017 will unfortunately surpass these figures due to the consecutive very intense tropical cyclones, Harvey and Irma, each of which caused more than 100B€ damages, as well as the South-East Asia flooding. Furthermore, other events such as heat waves and droughts are estimated to lead to a 9-10% reduction in the worldwide annual cereal production, *i.e.* losses of about 70 B€ per year¹³. Still more worrying, climate change is very likely to lead to an intensification and/or increase of these types of weather extremes¹⁴.

Furthermore, the major earthquake which struck Haiti in January 2010 killed more than 230,000 people and injured another 300,000. Altogether it led to the displacement of 1.5 million people. The total damages are estimated to be close to 10 B€.

The exposure to the climate-change induced risks will increase, as too will the societal vulnerability to all types of catastrophes, consequently leading to increasing human and financial costs. It is then of outmost importance to be able to benefit from reliable, accurate and detailed information about either forecasted or actual events in order to prevent huge humanitarian and economic impacts. Improved preparedness and avoided damage are essential.

High performance computing (HPC) is playing a key role, and is in fact a preferred tool, for predicting the possible developments and trajectories of such events, and for anticipating their intensities, whenever they are predictable phenomena on the daily to weekly scale. In all cases, and once the events have struck, HPC is key for providing, through the so-called "**urgent computing**", crucial information to plan and organize protection and mitigation actions to prevent further damages.

Research in the WCES fields is not only knowledge-based and of very-high socio-economic relevance, but it is also of key importance for Europe for preparation of the EU policy on environment at large, *e.g.*, understanding the likely impact of the natural environment on infrastructure, economy and society, enabling informed investment decisions, developing civil

¹² <https://www.munichre.com/en/reinsurance/business/non-life/natcatservice/index.html>

¹³ Lesk, C., Rowhani, P., & Ramankutty, N. (2016). Influence of extreme weather disasters on global crop production. *Nature*, 529, 84-87.

¹⁴ <http://www.ipcc.ch/report/srex/>

protection capabilities and supporting the Copernicus programme¹⁵ to deliver information on environment and security.

Climate Change

Climate models play a crucial role to understand and predict climate variability and change and to influence advice and policy, as demonstrated by the elaboration of the Paris Agreement. However, while there is great confidence in the fact that climate change is happening, there remain many uncertainties, especially at regional scale, to inform adaptation policy. Increasing the comprehensiveness of ‘whole Earth system’ models requires more and more Earth system components to be included. Detailed scenarios for future climate can only be achieved by simultaneously increasing the resolution of the models. A further challenge is to provide more robust predictions of regional climate change from decadal to centennial timescales to underpin local adaptation policies. A dual track approach should be taken involving multi-member multi-model comparisons at the current leading-edge model resolution (about 25 km, limited to a few decades) alongside the longer-term aim to develop a global convective resolving model (down to 1 km resolution). This requires a coordinated set of experiments and multi-year access to HPC platforms. Furthermore, as a large number of climate simulations extend over long periods of wall-time and involve coordinated experiments between many international groups, the stability in time of the target HPC platforms is of key importance. Issues will need to be resolved relating to mass data storage and the dissemination of model outputs for analysis to a wide-ranging community of scientists over a long period.

At the international level, the European climate modelling teams belong presently to the leading groups, with, *e.g.*, particular reference to the high-resolution simulations. As examples of this leadership, one can mention:

- success of projects such as "Upscale", run on PRACE facilities, where the crucial importance of high-resolution was clearly shown for the simulation of extreme events such as tropical cyclones¹⁶;
- the fact that European teams are the focal points for high-resolution in international inter-comparisons such as CMIP.

It is nevertheless to be kept in mind that this forefront role could be endangered if European teams were not be able to access facilities capable of multi-petaflops, if not exascale, performance, as this is, for example, already the case for Chinese teams using the TaihuLight computer¹⁷.

¹⁵ <http://marine.copernicus.eu/>

¹⁶ Strachan et al. (2013). *J. Climate*, **26**, 133-152

¹⁷ The Chinese "CESM Project" was recently able to run preliminary global simulations at 1km resolution over one million cores and 24,000 MPI processes

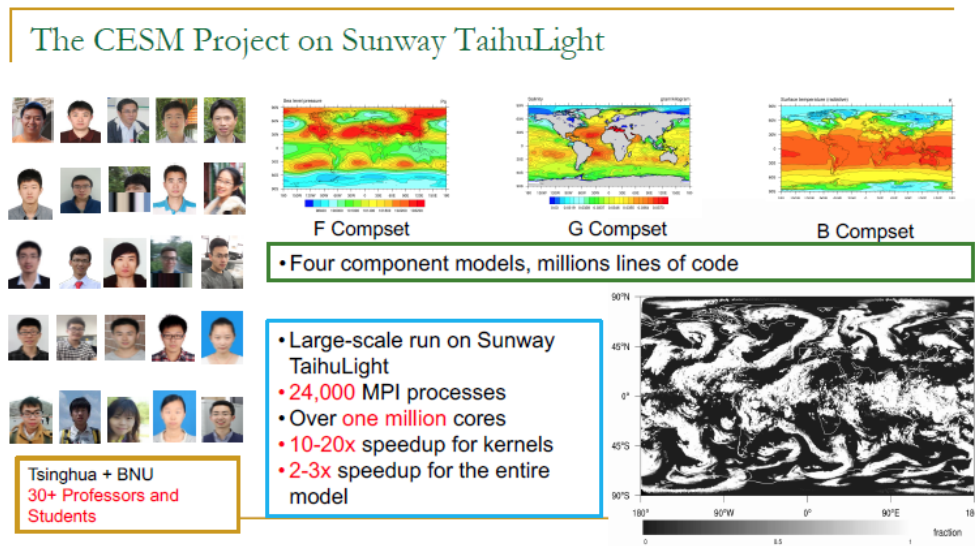


Figure 13 - CAM atmospheric model scaling at 1km resolution on 1.5 million cores, awarded by the 2016 Gordon Bell Prize

Oceanography and Marine Forecasting

Improvements in the understanding of the ocean circulation and its associated biogeochemical cycles is critical to properly assess weather and climate variability, future climate change and related impacts, *e.g.*, sea-ice extension, coastal sea level and inundation risks, ecosystem status and productivity, ocean acidification ... Beyond climate, oceanography and marine forecasting are necessary to assess and maintain the wealth of ocean services provided to society, *e.g.*, supply of food and energy, transport of goods ..., some of them being under ever-increasing stress. The stressors and the impacts must be evaluated and regulated in order to preserve the ocean's integrity and resources. Society must also protect itself against marine natural hazards such as storm surges and their trends in view of climate change. Marine safety concerns are becoming more acute as the coastal population and maritime activities continue to grow. More accurate ocean models are the key tools to assess and predict how the different components of the ocean (physical, biogeochemical, sea-ice) evolve and interact. They are also the building stone for operational oceanography, which is now delivering short lead-time forecasts using complex data assimilation systems that include all *in situ* and satellite observations. The main perspective is to increase the number and realism of physical processes represented in the models, as well as their spatial resolution: the vision is to have unstructured grid resolving the 100 m coastal scales up to the open ocean eddy field, of the order of 1 km, coupled to comprehensive data assimilation. This will provide science-based information for the development of a sustainable blue economy. Ice sheets modelling must be actively pursued, for its own sake, for its importance in driving climate, and also for its importance with respect to, *e.g.*, maritime transportation.

Europe has now reached a favourable position at the international level. Thanks to long-standing efforts of European research teams for converging toward common tools, operational oceanography, which started from the early MERCATOR-Océan and MyOcean projects, is now fully mature: in late 2014, the European Commission established its "Copernicus Marine Environment Monitoring Service" (CMEMS)¹⁸. This provides a strong incentive for research

¹⁸ <http://marine.copernicus.eu/>

teams to use, further improve and validate the numerical ocean model at stake, and to interface it with the wealth of data streaming from satellites, floats and more conventional observations. This unique concentrated effort is key to the European success.

Meteorology, Hydrology and Air Quality

Weather and flood events with high socio-economic and environmental impacts may be infrequent, but the consequences of them occurring can be catastrophic to those societies and systems that are affected. There is, of course, a link to climate prediction and climate change impacts, if severe meteorological and hydrological events are to become more frequent and/or extreme. Predicting these low frequency, high impact events a few days in advance – with enough certainty and early warning to allow practical mitigation decisions to be taken – remains difficult. Understanding and predicting the quality of air at the Earth's surface is an applied scientific area of increasing relevance. Poor air quality can cause major environmental and health problems affecting both industrialised and developing countries around the world (*e.g.*, adverse effects on flora and fauna, and respiratory diseases, especially in sensitive people). Advanced real-time forecasting systems are basic and necessary tools for allowing early warning advice to populations and practical mitigation strategies in case of air pollution crisis.

At the international level, it is fair to say that European meteorology is now at the forefront. For example, if one considers the accuracy of predicted trajectories for the Irma tropical cyclone, it has been shown recently that ECMWF¹⁹ has produced the most accurate ones, as early as 5 days in advance, see figure below²⁰. Of course, some errors are still affecting the forecasts and further improvements are needed. These improvements will concern both the meteorological models themselves, with more detailed physics and thermodynamics to be accounted for, and the smaller grid size at which they can be solved, *i.e.*, access to more powerful computers.

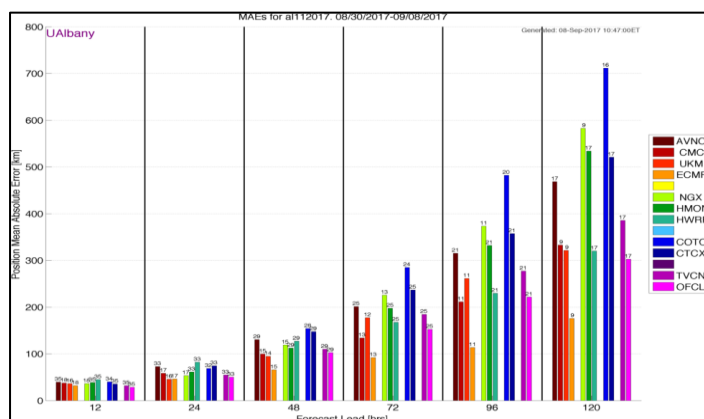


Figure 14 - Performance of current weather forecast models on up to 120 hours

Solid Earth Sciences

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 societal implications, playing a central role in natural hazard mitigation (seismic, volcanic, tsunami and landslides), hydrocarbon and energy resource exploration, containment of underground wastes and carbon sequestration, and national security (nuclear test monitoring and treaty verification). In the

¹⁹ European Centre for Medium Range Weather Forecast

²⁰ <https://arstechnica.com/science/2017/09/us-forecast-models-have-been-pretty-terrible-during-hurricane-irma/>

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. In terms of seismic imaging, going to high frequencies and thus to much increased resolution still remains a formidable challenge. 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 capability to address the issues of resolution, complexity, duration, confidence and certainty, and to explicitly resolve phenomena that were previously parameterised. Each of these challenges represents an increase by a factor of at least 100 over individual national facilities currently available. “Urgent computing” is also of increasing importance when there is a seismic event, so as to provide assistance in organizing relief operations.

A common feature of almost all of the actual WCES computational models is that they have a strong legacy, being developed through time, most often over several years, using mathematical methods or computational implementations that require adaptation to exascale architectures or even for actual massive high-performance platforms. As these computational models are almost always very large (hundreds of thousands or millions of code lines) and as most of them represent hundreds of person-years of development, it is quite difficult to change them at once. It is then of great importance to deeply re-think algorithms, in order to optimize and improve their computational performance and to obtain a real scalability of these models on complex parallel architectures.

2.2.2 *Challenges in Earth’s climate system*

Fundamental questions facing climate change research can be summarised in four key challenges, all related to the need for better qualification and quantification of uncertainty of the predictive simulations. All these questions are strongly linked to the amount of computing power and data storage capacity available since they ask for increased model resolution, large numbers of experiments, increased complexity of Earth system models, and longer simulation periods. It is also important to perform coordinated ensembles of simulations, using independently developed models, to ensure the robustness of the model results. Such coordinated multi model activities are carried out within the framework of the WCRP and serve as a basis for IPCC Assessments but considerably more could be done within the European context.

Challenge #1: The need for very high-resolution models to better understand, quantify and predict extreme events and to better assess the impact of climate change on society and economy on the regional scale

Currently, global climate models have typical grid spacing of 100km and are limited in their capacity to represent processes such as clouds, orography effects, small-scale hydrology, etc. The latest generation of models, under development or just starting to be used, have grid spacing in the 10–50 km range, and there is evidence that a number of important climate processes are better represented at this resolution (e.g., ENSO, water and energy transports, blocking, tropical storm location and frequency, etc.). A priority is to continue the development of coupled models at such high resolution and use them in multi-member, multi-model inter-comparisons focused on key climate processes. In weather forecasting applications, much higher-resolution, convection-resolving limited-domain models are now being used operationally. However, these

models cannot be run globally for climate, which requires much longer simulations because of the prohibitive cost of associated computing resources and limits in model scalability. **The climate community's first longer-term 'grand challenge' is therefore to develop global climate models that resolve convective scale motions (nominally around 1 km horizontal resolution).** Convection-resolving (mesh size of ~1km) global models are expected to improve our predictions and understanding of the effect of global warming on high impact weather events on seasonal, decadal and century timescales. Another issue is related to the simulation of regional-scale climate features, of crucial importance for assessing impacts on society and economic activities (farming, fisheries, health, transportation, etc.) and for which improved regional models, embedded in global climate models, are necessary. Such regional climate models, which are becoming more recently real Regional Climate System Models (*i.e.* including aerosols, biogeochemistry, carbon cycle, etc. for some of them), currently running at 10–50 km, are also calling for spatial resolution of a few kilometres.

Recently 20 leading European scientists, who formulated a proposal to establish a European Programme on Extreme Computing and Climate (EPECC²¹), further underlined this. The current Centre of Excellence in Simulations for Weather and Climate, ESiWACE, is aimed at the implementation and operation of current and future models that will address this HPC-challenge.

Developing such very high resolutions will require developing scalable and more efficient dynamical cores and improving physical parameterisations. Increasing model **resolution down to 1 km requires increases by factors of at least 100 to 1,000 in computing power compared to the current state**, *i.e.* in the multi-petascale to exascale range toward the end of the decade.

Challenge #2: The need to further account for the complexity of the Earth's climate system

Earth system models already include representations of biogeochemical processes and of subsystems of major importance such as the carbon cycle. These representations are of key importance to assess the sensitivity of predictions to biological and chemistry processes, like those involving land surfaces, oceans as well as greenhouse gas (GHG) reactions. Some key processes, such as the indirect effect of aerosols, remain a challenge to correctly reproduce. In addition to the value of being able to predict changes in vegetation and atmospheric composition, it turns out that biogeochemical processes can have quite a marked effect on the magnitude of climate change. For example, the European modelling groups were the first to show in coupled Earth system models that the global carbon cycle could accelerate forthcoming climate change. The carbon cycle is also intertwined with other biogeochemical cycles, such as the nitrogen cycle, so other matter cycles need to be included. Moreover, other processes, such as GHG reactions or aerosol related processes and their indirect effect on clouds or interactive vegetation, still need to be better accounted for.

It should be noted that including the representation of biogeochemical cycles using different biochemical tracers and aerosols typically increases time by a factor of between 5 and 20. **An increase of computing power by a factor of 5 to 20 is then required to better account for the complexity of the system.**

Challenge #3: The need to quantify uncertainty

Future projections of climate change are uncertain for two main reasons: (i) future forcing by greenhouse gases and aerosols is uncertain, as it depends on climate policy and societal

²¹ <https://ec.europa.eu/futurium/en/content/flagship-european-programme-extreme-computing-and-climate-0>

behaviour; (ii) models are inherently imperfect, owing to insufficient representation of physical processes, so that natural and anthropogenic components of future climate change remain difficult to correctly predict.

A wide range of underlying scientific issues needs to be solved to reduce these uncertainties: (i) assess the predictability of climate on a range of timescales; (ii) estimate the range of uncertainty that can be fully represented using the models currently available; (iii) measure the sensitivity of climate and determine how much current uncertainty can be reduced by a better account of the major feedbacks (*e.g.*, clouds, atmospheric chemistry, the carbon cycle).

The consensus approach is that the uncertainty can be estimated by combining multi-model multi-member experiments, due to multi-model partial reduction of uncertainties linked to parameterization of physical processes, and due to multi-member ensembles that help in identifying preferred system time-trajectories. Increasing the European contribution to international multi-model experiments allows investigation of the sensitivity of results to model parameters.

Computing requirements scale of course directly with the number of ensemble members, and at the same time the number of members required to keep the same signal-to-noise ratio increases as the spatial resolution increases. **Ensemble experiments are therefore computationally expensive –a factor of 10 to 100 for each experiment.**

Challenge #4: The need to investigate the possibility of climate surprises

Due inherently to its non-linearity and resulting chaotic behaviour, earth system changes could be abrupt, surprising and unmanageable, as evidenced by the occurrence of such rapid changes in the past, *e.g.*, for the Atlantic thermohaline circulation (THC). It is then crucial to determine if there are thresholds in the greenhouse gas concentrations over which such abrupt THC changes could occur again. Surprises may also arise from ice sheet collapse and large amounts of fresh water in the ocean. Some key questions arise on the possibility to model glacial–interglacial cycles, including changes in carbon cycle and major ice sheets. Observational evidence from past climates will help calibrating complex climate models and allows studying the pace of climate changes as it was punctuated by rapid events.

Investigating the possibility for such rapid changes requires longer runs, both for the past and the future, with medium- to high-resolution models and various degrees of complexity. **An increase of computing power by a factor of 10 to 1,000** is then required which, however, cannot result only from an increased number of cores; it also requires an increase of power for each individual core (*e.g.*, many-core nodes, accelerators).

Roadmap

Both physical and software infrastructures are required to solve the above issues and computing power is a strong constraint, with both capability and capacity computing being important.

Capability is needed given the long timescales every coupled model configuration needs to spin up to a stable state. Higher-resolution simulations also strongly benefit from capability. However, multi-member ensemble runs are a typical **capacity** problem. For example, participation in CMIP6 requires running of a high number of simulations, and certainly asks for capacity as all of these runs must be considered part of the same experiment. There is furthermore a 20 to 30-fold increase in the amount of data generated during CMIP6 with respect to CMIP5, explaining why the data issue is quite often the key bottleneck limiting the overall performance.

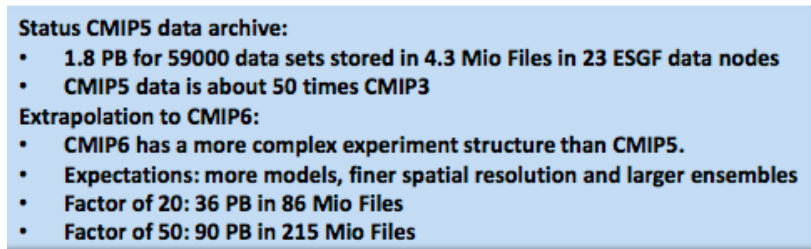


Figure 15 - HPC and data requirements between CMIP5 and CMIP6 exercises

Capacity-demanding ensemble-type runs with high-resolution models are generally done most efficiently on central systems. There are however applications for which distributed systems would provide good performance, but these cases generally depend on models with very good portability and with relatively low input/output volumes, criteria that are not fulfilled by most earth system models. While ESiWACE²² is aimed at contributing to solving these challenges with its integrated teams of climate and weather scientists, applied mathematicians, computer scientists and computer architects, the very-next computing systems will still need to provide both capability and capacity with a good balance between computer power, storage system size and read/write efficiency. On the longer term, parallelisation-in-time methods, when available, might help reduce this gap between capability and capacity requirements.

Climate models fundamentally scale with difficulty because the problems at stake are connected, physically and algorithmically. Requiring significant communication results in an increasing overhead with increasing domain decomposition. In the last few years, the climate community has gained experience in using Tier-0 PRACE machines (*e.g.*, UPSCALE, PULSATION, THROL, SPHINX). However, most climate models are still executed (*e.g.*, for IPCC scenarios) on Tier-1 national, sometimes on Tier-2 dedicated machines, which are adequately tailored with a good balance between computing performance, bandwidth to storage, and storage capacity. Lack of scalability of the multi-component (*e.g.*, atmosphere, ocean, land, coupler, etc.) climate models is just one reason why they have not been run very often yet on Tier-0 machines. Another key reason for this limited use is that running a coordinated set of experiments requires multiyear access to a stable platform (with respect to hardware and middleware), because of the high cost of loss of bit level reproducibility. **A suitable strategy for European earth system modelling is needed to exploit efficiently the entire European ‘computing ecosystem’.**

In summary, climate modelling requires:

- optimisation and adaptation of the mathematical algorithms to achieve better scalability on the new hardware and architectures;
- computing platforms offering both capability and capacity and access requirements including these two aspects;
- multi-year access so that the simulations can be carried out on the same Tier-0 machine with the same environment as the ones used during the model porting and validating phase;
- the possibility of providing multi-year multi-modelling access to European modelling groups, so that they can investigate scientific questions through targeted multi model experiments (*e.g.*, CMIP6);
- appropriate mass storage and dissemination mechanisms for the model output data;
- -high-performance data analytics, *i.e.* systems with server-side capabilities;

²² Centre of Excellence in Simulation of Weather and Climate in Europe

- high-performance I/O middleware as solutions to accelerate I/O to scientific formats;
- end-to-end workflow mechanisms for orchestration of analysis experiments and processing chains at the data centre level;
- outputs of the coordinated set of experiments easily available for analysis to a wide-ranging community of scientists over a long period.

With present technologies and with the actual state of models, the scalability is shown in the figure below. Performances keep improving, for, among others, the IFS model operated by ECMWF for numerical weather forecasting and by many other centres for climate modelling.

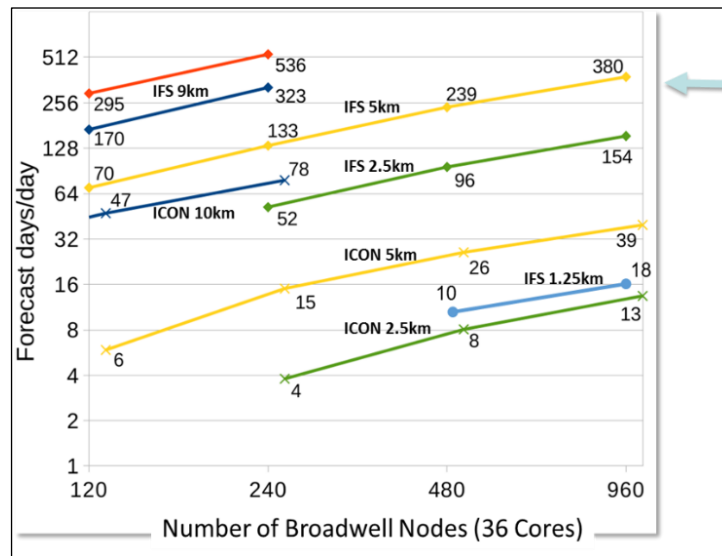


Figure 16 - comparison of the scalability of different models with a target resolution of 1km and a production of 1 year/date

These requirements will be central to the usefulness of the PRACE infrastructure for the climate community, the reason for ENES²³ collaboration with PRACE aimed at fostering its use of Tier-0 machines. PRACE operates general-purpose machines, while climate has specific requirements as described above. Those would be better met by a dedicated world-class machine. Such a facility would also allow the production of ensembles of very high-resolution simulations, necessary to develop and provision future climate services.

Besides using PRACE and its Tier-0 computers to address the complex issues behind climate science, the community developed recently the EPECC initiative ("European Programme on Extreme Computing and Climate"²⁴).

EPECC is an ambitious pan-European flagship project to implement a wholly new class of models for the analysis and prediction of weather and climate extremes, and their impacts, in a changing climate. Through a more fundamental and accurate representation of all components of the Earth system - made imaginable by the emergence of exascale computing - these models will greatly enhance the capacity to predict extremes, their underlying drivers, and eventual impacts. EPECC follows an initiative of world leading weather and climate centres in Europe

²³ European Network for Earth System Modelling; www.enes.org

²⁴ <https://ec.europa.eu/futurium/en/content/fet-flagships> and <https://ec.europa.eu/futurium/en/content/flagship-european-programme-extreme-computing-and-climate>

gathered within ENES. It is supported by the largest European high-performance computing and data handling facilities, it is coordinated by ECMWF, and it works in partnership with key European science and management consortia. The topic "Energy, environment and climate change" was selected by the EC following the submission of ideas for Flagships. The selection process starts with a so-called preparatory action, which allows 5-6 separate proposals to advance and define the blueprint for the respective Flagships. The call for proposals for the preparatory action was issued in late October with a submission deadline in February 2018. A writing team has been established for preparing the EPECC proposal, including sections on technology, science, business cases, downstream applications, and co-funding.

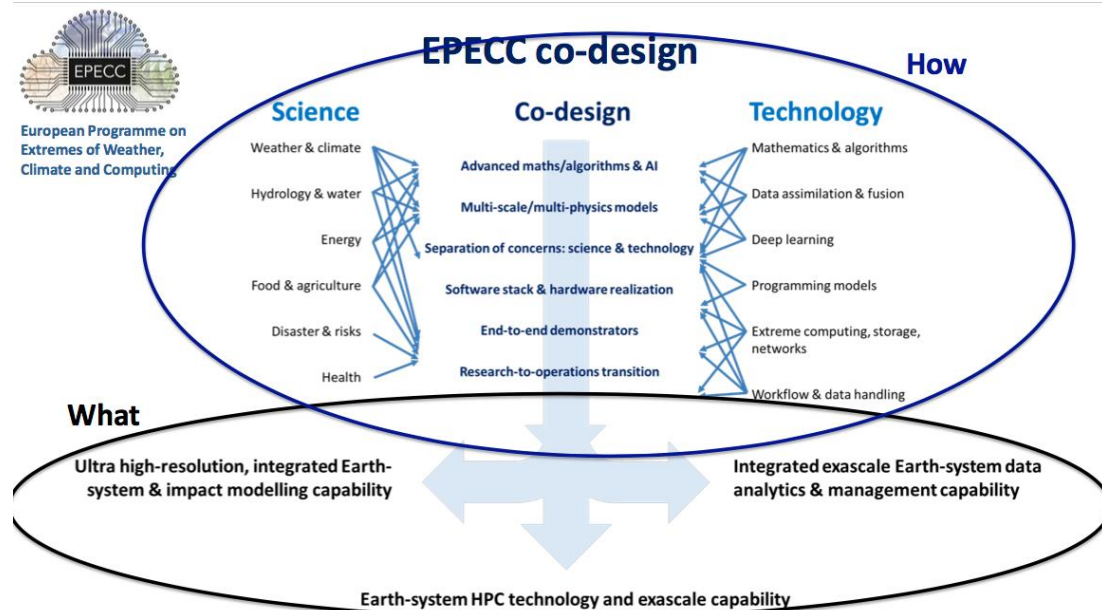


Figure 17 - Global view of EPECC FET challenges

2.2.3 Challenges in oceanography and marine forecasting

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 processes need to be coupled to wind waves, tidal motion and marine biogeochemistry. Variations in the upper ocean vertical motions affect nutrient supply and ocean biota, but they are not yet resolved by the actual models. No large-scale model is yet coupled to surface waves and only few coastal models are. Altimetry measurements and global Argo array of profilers have led to major breakthroughs, but biogeochemistry is still largely subsampled. Uncertainties also arise from coupling with hydrological models, still in its infancy. On the other hand, 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.

Multi-peta- to exa-flops capabilities would greatly help resolve four major issues:

Challenge #1: The need for high-resolution ocean circulation and data assimilative models

Major ocean currents are critically governed by small-scale topographic features, including the coastlines, by energetic eddies, tides and surface wind waves. Present eddy-permitting

O(10km) grid models have now begun to capture eddy processes in the subtropics and mid-latitudes (with strong effects, *e.g.*, on the Gulf Stream), but much higher resolution is needed to achieve comparable progress in the subpolar and polar oceans. A number of groups are beginning to operate eddy-resolving models under the HighResMIP CMIP6 protocol, but coupling with wind waves and tidal motion will increase the computing time by several orders of magnitude. Developing data assimilation in eddying ocean models is another challenge, as it requires more efficient use of all the satellite altimetry data, ARGO floats, drifters, and other *in situ* observations. The reconstruction of the past state of the ocean via high-resolution models and new data assimilation schemes should be done for at least the past 130 years, where observations have been collected.

Challenge #2: Understanding and monitoring the carbon cycle in marine ecosystems

Another great challenge is to understand the evolution of marine ecosystems and their sensitivity to a changing environment (*e.g.*, reduction of ocean pH due to increasing atmospheric CO₂ content threatens corals and other marine organisms). Accurate ocean biogeochemical models of these interactions would greatly improve understanding, monitoring and forecasting of the marine carbon cycle and support the preservation of the ocean ecosystems. A key issue is to reproduce as accurately as possible the biomass primary production, as well as the structure of the food web that re-cycles the organic carbon in the water column. High resolution is essential, possibly using nested regional models (grid refinement up to 1/100 degree) embedded within larger-scale models.

Challenge #3: Increasing the lead time of ocean predictions, from weeks to decades

Climate physical mechanisms associated with natural decadal and longer time scale variability remain unclear. Although ocean predictions at mesoscale resolution and weekly time scales are now available for some regions, key issues are the verification of retrospective predictions that are possible only for a handful of “instances” of decadal change (*e.g.*, mid-1990s warming of the northern North Atlantic). Model climate is indeed systematically different from the observed climate, and numerical models initialized with an ocean state that is “close” to observations will drift to their preferred equilibrium state over the same time-horizons of relevance for decadal predictions. Seasonal, inter-annual and decadal predictions are an enormous challenge since it is not yet clear what is the major source of uncertainty and how to deal with it. High-resolution ocean-atmosphere coupled models are required here.

Challenge#4 Predictions and climate impacts at the coastal scales

Simulating, reconstructing and predicting the coastal ocean is possible, timely and needed. The coastal ocean influences the transport of materials that fertilizes the deep ocean; it controls the dynamics of the large-scale currents (Gulf Stream, Kuroshio) and it shapes the major upwelling areas. The present-day strategy for modelling the coastal ocean in the global ocean models is to use two-way nested higher-resolution models in areas of particular deficiencies by the coarse large-scale models (*i.e.*, the eastern sites of the world ocean basins). The future challenge is to use unstructured grid models to resolve the shelf-coastal dynamics with few hundred metres resolution, while keeping a resolution of a few kilometres in the open ocean areas.

Roadmap

Running models that can resolve the spectrum of ocean dynamics down to the submesoscale (*e.g.*, 1/100°) is at present beyond any computing capability, except in very local areas and over short periods, as reaching kilometric resolution at global scale would demand a thousand-fold increase in computer power. **Adding a full carbon cycle model increases the computational cost and storage capacity by a factor of O(5) i.e 5 times more.** Adding wind

waves to the circulation model probably doubles the computing time of any model.

The highest-resolution global model presently used in Europe for research and operational forecasting uses a grid resolution of $1/12^\circ$ (*i.e.*, 5 to 10 km). A single 50-year long run requires an available crest computational power of 25 Teraflops for a period of 2 months. Present scientific objectives require performing either series of multi-decadal experiments with such models, or ensemble runs of $O(50)$ members with coarser $O(1/4^\circ)$ models, or increasing eddying ocean model complexity with a full carbon cycle. Moreover, it would be important to include grid refinements, down to 1 km-100 m scales, to explicitly resolve the coastal dynamics that are critical for the global circulation. Operational oceanography, on the other hand, urgently needs to develop higher-resolution (*e.g.*, $1/24^\circ$) products for its Marine Core Services, asking for an increase of computational power by a factor of at least 10. At the same time, significant progress in data assimilation is needed and this is even more demanding in terms of computational power.

Altogether, the above requirements call for crest computational resources of 500 to 1000 Teraflops available for periods of months. **This can be obtained only with $O(10-100)$ petaflops computers coupled to very large storage facilities of $O(10-100)$ Petabytes** that will store the simulation outputs over periods of a few years for subsequent studies.

2.2.4 *Challenges in weather and air quality*

The mitigation of high-impact weather achievable by having a more accurate and timely weather nowcasting system requires that the comprehensiveness of the various atmospheric models will have to be enhanced dramatically, *e.g.*, a horizontal resolution increased to about 1 km to resolve convection explicitly. Probabilistic approach will have to be taken to arrive at meaningful warning scenarios. Similarly, one of the most far-reaching developments results from the enhanced capabilities in air quality modelling forecasts. The operational air quality forecast systems require high spatial resolution, significant computational efforts and a large volume of input data. The availability of increased computational power and the possibility of accessing scattered data online with the help of a cloud infrastructure, coupled with advances in the computational structure of the models, now enable their use in real time air quality forecasting. These challenges are currently being tackled, albeit on a much smaller scale. Today it is only possible to solve the problems for a limited number of variables in a limited area. It is known in principle what needs to be done in future, with computing power about three orders of magnitude larger than today.

Fundamental questions facing weather and air quality research can be summarised in three key challenges.

Challenge #1: The need for very high-resolution atmospheric models and associated forecasting system

The resolution of models will have to be increased in both time and space to be able to resolve explicitly physical processes which today are still parameterised. To arrive at meaningful results, this work also entails the evaluation of the error growth due to the uncertainty of the initial and boundary conditions. Furthermore, the data gathered by new, high-resolution observing systems, either space- or ground-based, need to be assimilated using new techniques such as 4DVAR²⁵. Furthermore, the I/O rates of the applications will be in the order of 5 GB/s

²⁵ Four-dimensional variational method. This method consists in combining all types of observations at different times (from satellite, radiosondes, aircraft and surface networks) with background fields from a previous forecast to create a starting point for a new forecast.

for the duration of the runs, resulting in files of more than 6 TB. In order to test these ideas, a preoperational trial of a complete end-to-end system needs to be carried out, checking whether the integrated system consisting of data acquisition, data assimilation, forecast run and product generation can be handled in a sufficiently small time-period to allow future operational deployment.

Challenge #2: The need for high-resolution assimilated air quality forecast and cloud–aerosol–radiation interaction models

Reliable, high-resolution (*i.e.*, 1 km) air-quality forecasting system for Europe becomes essential in informing and alerting the population. The accuracy of such a system depends dramatically on the accuracy of emission data. It is thus timely to develop a system using the 4DVAR technique. 4DVAR is computationally very expensive and requires a lot of memory. An important unresolved question is the role of aerosols in modifying clouds, precipitation and thermal atmospheric structure, as, *e.g.*, dust generates complex feedbacks within atmospheric processes. Other aerosols may also act as cloud condensation nuclei and affect precipitation processes. In current generation models, these processes are still highly simplified.

Challenge #3: The need for developing a pan-European short-range weather and air quality modelling systems

Further improvements in these systems are needed, which all require more computing power. Among these one can cite: (i) representation of clouds and small-scale processes; (ii) coupling between the biosphere, atmosphere, aerosols and clouds; (iii) chemical reactions in the environment; (iv) estimation and computation of different sources of emissions. Furthermore, short-range and very-high-resolution models need observational data and subsequent data assimilation. **This increases the computing requirements by a factor of at least 10.**

Roadmap

The expected growth of computing resource will allow a gradual increase in complexity and resolution of the various models and will narrow the gap between what is possible and what is required. The work will be carried out by the already-existing collaboration between European National Meteorological Services, the European Centre for Medium Range Weather, collaborating universities, and scientific research centres.

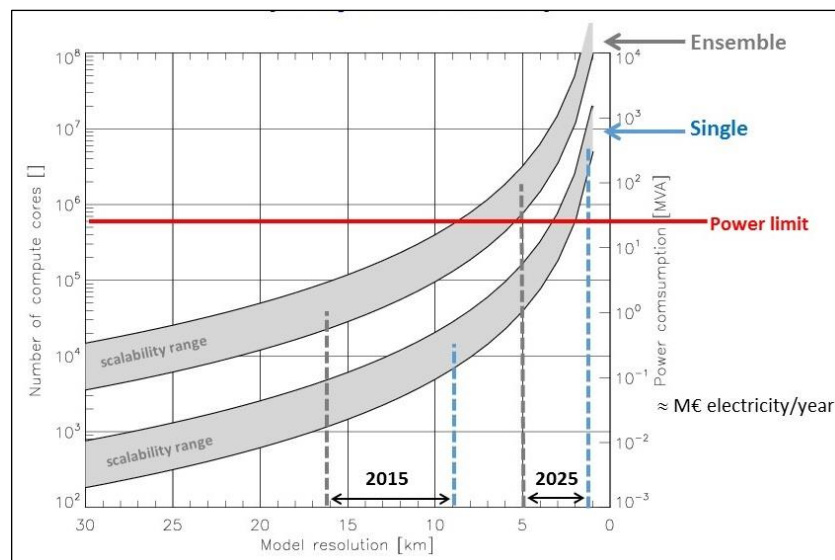


Figure 18 - Simplified illustration of the number of compute cores and power required for single 10-day model forecast (lower curves) and 50-member ensemble forecast (upper curves) as a function of model resolution, given today's model code and compute technology. From Bauer et al., 2015²⁶.

Integrated air-quality forecasts for the area of all EU Member States at the highest feasible resolution (*i.e.*, 1 km) require extensive computing resources, typically of the order of 100–300 Tflop/s sustained performance. The most relevant initiative in this field is the Global Monitoring for Environment and Security (GMES) project, a joint initiative between EU and ESA to strengthen the acquisition and integration of high-quality environmental, geographical and socio-economic data, that will help improve policymaking from local to global level.

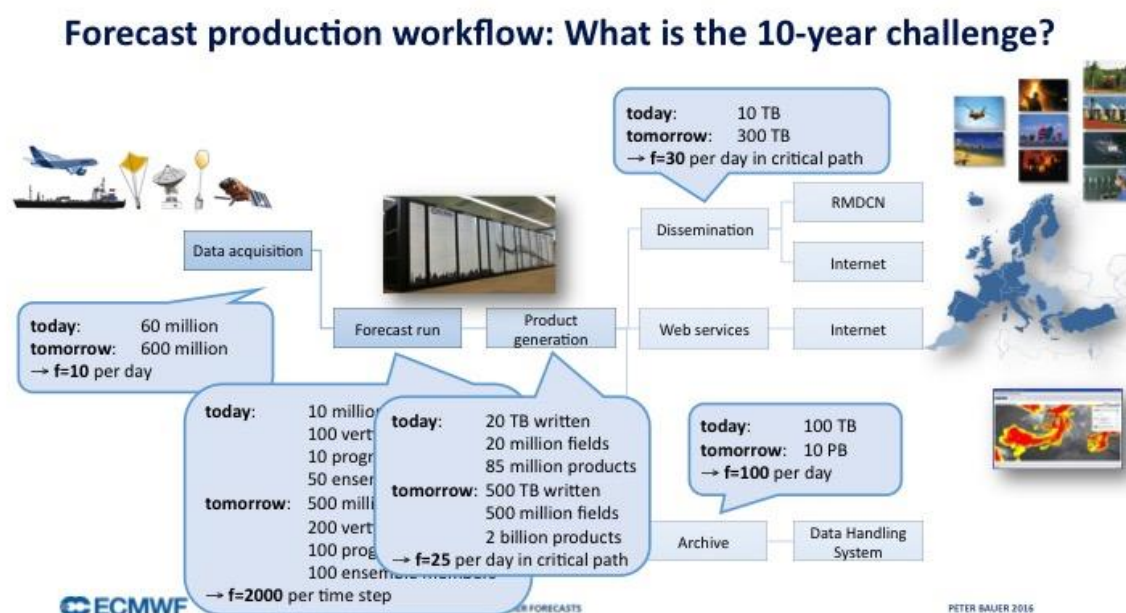


Figure 19 - Expected requirements in weather forecast by ECMWF

2.2.5 Challenges in Solid Earth sciences

A key issue is to better quantify uncertainties, and estimate the probability of extreme events through simulation. For some problems, the underlying physics is adequately understood and the main limitation is the amount of computing and data capabilities available. For other problems, a new level of computing and data capabilities is required to advance understanding, *e.g.*, for example modelling and simulating earthquake dynamics rupturing processes together with high frequency radiation in heterogeneous earth, or imaging based on high-frequency body waves.

The solid-earth community is preparing itself for massive use of supercomputers by the current (re)-organisation of some of the communities through large-scale projects, including the ESFRI project EPOS²⁷, the FP7 Infrastructure project VERCE²⁸, the Marie Curie ITN initiative QUEST²⁹, and other initiatives such as SHARE, TOPOEurope, TOPOMod and MEMoVOLC.

Five major challenges are outlined below.

²⁶ Bauer, P., Thorpe, A., & Brunet, G. (2015). The quiet revolution of numerical weather prediction. *Nature*, 525, 47-55

²⁷ www.epos.eu.org

²⁸ www.verce.eu

²⁹ www.quest.itn.org

Challenge #1: Seismic cycle, earthquake physics, and ground motion simulation for the seismic hazard mitigation

In the absence of deterministic prediction methods, the forecasting of earthquake ground motion based on simulation of scenarios is one of the most promising tools to mitigate earthquake-related hazard. The probabilistic model of earthquake recurrence has however been recently challenged with the discovery of transient deformation processes called “slow earthquakes”. Better understanding of these phenomena is now needed, which requires systematic intensive analysis of large volumes of data recorded by seismic and geodetic instruments, *i.e.*, TBytes today, continuing to increase at a very rapid pace. Another important issue is to improve actual physical understanding of earthquake rupture processes and seismicity, requiring large-scale simulations of these phenomena together with data assimilation and inversion. Coupling the modelling of earthquakes together with high-frequency radiation in heterogeneous media will remain a grand challenge problem even with the next generation of computers. There is also an urgent need to enhance earthquake simulations and to improve model realism by incorporating more fundamental physics related to subsurface soil behaviour into the simulations.

High-resolution models are also required to develop and assess fast operational analysis tools for quasi-real-time seismology and early warning systems. Accurate simulations must span from metres near the earthquake source to hundreds of kilometres across the entire region, and from hundredths of a second to capture the greatest impact on buildings to hundreds of seconds for the full event.

Altogether the goal is to:

- extend by a factor of 10 the spatial dimensions of the models;
- increase the highest resolved frequency above 5 Hz (for structural engineering purposes), implying a 64-fold increase in computational size;
- move to more realistic soil behaviours, at least a 2 orders of magnitude increase in computational complexity;
- incorporate a new physics-based dynamic rupture component at 100 m resolution for near-field risk assessment, implying at least 1 order of magnitude increase in computation;
- invert for both the earthquake source and the geological parameters, which necessitates repeated resolutions of the forward problem, *i.e.* an increase of 1 to 2 orders of magnitude in computations;
- perform stochastic modelling of seismic events and wave propagation for quantifying uncertainties and exploring earthquake scenarios, which implies a 10–50 times increase in computation.

Challenge #2: High-resolution imaging techniques

The capacity to accurately image the earth’s subsurface is one of the challenging problems that have even more important economic applications in terms of resource management (identification of new energy reservoirs and storage sites as well as their monitoring through time), and understanding the formation and evolution of earth system.

Recent progress in seismic acquisition from dense networks of sensors and further data analysis makes it possible to extract new information from fine structures associated with strongly diffracted waves. The challenge is nowadays to fully exploit these fine details, exploring late arriving signals associated with strongly and multiply diffracted waves. This opens new perspectives in very complex geological settings, *e.g.*, possible ascents and descents of magmas

within volcanoes. Accurately imaging seismic rupture evolution on complex faulting systems proceeds in similar fashion.

Only adjoint-based methods are tractable to date for local analysis, and back projection is the mathematical tool for image reconstruction. Going to semi-local analysis will require simulated annealing genetic algorithms, leading to a drastic increase in computer resources that cannot yet be foreseen. Because thousands of forward problems should be achieved in an iterative optimisation scheme related to the number of sources and receivers, one must investigate techniques for solving these forward problems in a combined way efficiently. Other issues are related to reflection and diffraction in heterogeneous media, diffraction by rough topography, at very high frequencies (10–40 Hz), for which complex attenuation is expected, and bridging between deterministic estimations and probabilistic approaches.

Challenge #3: Structure and dynamics of the Earth's interior

Seismology is the only method to probe the earth's interior from its surface to its inner core. Improving the quality of 3D tomographic images of the earth's interior, with a resolution lower than tens of kilometres, using the continuously-increasing data sets of broadband seismological records, is nowadays essential to improve core–mantle dynamical models as well as the knowledge of the earth system's physics. So far, only a small fraction of the information contained in broadband seismograms is used for this purpose, while broadband seismological data volumes are still increasing at a faster rate than computational power. It must however be emphasized that the community is well organized with respect to the data problem, through the federation of digital seismic networks and to the notable presence of the European Integrated Data Archive(s) (i.e. ORFEUS and EIDA) which, irrespective of the specific archive to which the data request is submitted, provides data contained in all the federated archives. Recent advances in high performance computing have facilitated 3D simulations of seismic wave propagation at unprecedented resolution and accuracy. The realm of PFlop/s computing opens the door to full waveform tomographic inversions to significantly enhance the resolution of the earth's interior image. This is a grand challenge problem due to the large number of mesh-dependent model parameters and of wave propagation simulations required in the inversion procedure.

Solid-earth internal dynamical processes often take place at scales of tens of millions of years. In order to understand these systems, simulations must then be carried out concurrently with observations. Mantle convection provides the driving force behind plate tectonics as well as the other geological processes. Realistic models of thermomechanical mantle convection in 3D spherical geometry are required to better assimilate mineral physics and seismology information into the deep earth dynamics. The short time scale dynamic behaviour will also serve as the monitor for stress build up that loads seismically-active regions. Mantle convection drives plate tectonics and the earth thermal evolution. Mantle convection is dominated by slow viscous creep, involving a timescale of hundreds of millions of years. The Rayleigh number is of the order of 10^7 , inducing a time-dependent dynamics and small convective scales comparable to the size of the domain. Three-dimensional numerical simulations of mantle convection are currently performed both in Cartesian and spherical shell geometries. Those including both melt-induced compositional differentiation and self-consistent plate-like behaviour have so far been performed only in two dimensions or in small 3D Cartesian geometries. One computational challenge is thus to resolve convective features smaller than 100 km over a spherical domain of depth 2,900 km and circumference of 40,000 km. Another challenging issue is the resolution of the rapid spatial variations of the physical properties involved: viscosity strongly depends on temperature (by six orders of magnitude), pressure and stress (by two orders of magnitude with depth). Furthermore, how plate tectonics arises from

mantle convection is an outstanding issue. Incorporating full physics in high-resolution spherical shell models is nowadays a challenging problem that can be addressed only in the realm of multi-Pflop/s computing.

Challenge #4: Generation of the Earth's magnetic field

The extremely involved magneto-hydrodynamic of the core dynamics and the associated external magnetic field, and its reversals, are essential to understand, because of its relations with plate motions and also because of many analogies with astrophysical processes. Simulations are crucial to make progress in this field. Although dynamically self-consistent models of the geodynamo have simulated magnetic fields that may appear quite similar to the geomagnetic field, none is able to run in an earth-like parameter regime because of the considerable spatial resolution that is required.

Simulation tools are already implemented in parallel, but much higher resolution is required to compare with the conditions observed in nature. In addition, many realisations are necessary to create stable results for highly nonlinear processes with strong dependence on initial and boundary conditions. A grand challenge for the next generation of geodynamo models is to produce simulations with thermal and viscous (eddy) diffusivities set no larger than the actual magnetic diffusivity of the earth's fluid core, while using the core's dimension, mass, rotation rate and heat flow. Another challenge is to develop new highly-parallel, adjoint-based assimilation numerical methods.

Challenge #5: Interaction between solid earth, ocean, and atmosphere

Some of the most energetic and hazardous processes in the solid-earth field (*e.g.*, seismogenesis, volcanism) occur very close to the surface and generate strong oceanic (tsunamis) and atmospheric perturbations (ash plumes from volcanic eruptions). On the other hand, oceanic and atmospheric waves are strongly coupled to the solid earth and generate the so-called “seismic noise”. This type of seismic signals has been shown recently to be of efficient use for imaging of the earth's interior. Such inter-disciplinary studies have both observational and modelling components: huge volumes of continuous observations are analysed and large-scale and high-resolution numerical models help understand the physics of these coupled interfaces.

Roadmap

Nowadays, simulations of seismic wave propagation in small geological volumes have reached the Tflop/s plateau production and state-of-the-art cases are rising to the realm of Pflop/s computing. However, these simulations remain limited (see above, Challenge #1). Enhancing the resolution and the physics, performing inversion of seismic parameters, and quantifying the uncertainties will push these simulations into the realm of Pflop/s computing. **A key requirement will be the provision of large-scale end-to-end data cyber infrastructures to handle, analyse and visualise PBytes of simulated data for storage.** Visualisation of very large data sets will be a related important challenging problem.

In general, codes used in our solid Earth community are memory bound rather than compute bound. For example, arithmetic intensities of current seismic wave propagation algorithms are around 2. This means that most currently used codes profit more from increased memory bandwidths in newer hardware systems than additional floating-point capabilities in a single compute node. New techniques will have to be developed which address current algorithms in order to increase their arithmetic intensities in order to fully benefit from future hardware developments.

Formulation in the *time domain* (e.g., wave propagation in time) allows handling 3D imaging problems at the expense of computer time. For simulations on models of 100 km x 100 km x 25 km, with 10 Hz content, sustained performances range between 10 Tflop/s and 100 Tflop/s. Moving to more powerful resources will increase the size of the box and/or the maximum frequency.

The 3D imaging problem in the *frequency domain* remains a challenge nowadays both for computer resources and for numerical algorithms due to the lack of efficient large-scale parallel direct algebraic solvers. The realm of multi Pflop/s computing is providing the possibility of performing such seismic imaging. It requires access to a large memory/processor ratio, efficient algorithms for direct decomposition of very large matrices, and optimised parallel and sequential IO's. Achieving load balancing between processors will be a challenge. Unfortunately, resorting to iterative methods dims the interest of a frequency formulation compared to a time domain formulation.

Current advances in data acquisitions have increased data volume several folds, and processing methods at high resolution requires an increase in the computational effort beyond current resources. Large data volumes and complex mathematical algorithms make seismic data processing a very compute and I/O intensive task that requires high performance computers with a large amount of memory.

Global simulation of body wave phases is nowadays a Pflop/s challenge problem. This requires global wave simulations, at periods of 1 second or less, and space resolution at wavelengths of tens of kilometres, in 3D earth models including high-resolution crustal models, topography and bathymetry together with rotation and ellipticity. Current front-end 3D global seismology simulations run at wavelengths of tens of kilometres and typical periods down to 1–2 seconds, on $O(10^5)$ or more cores with a sustained performance of about 200 Tflop/s. The next-generation simulations will require hundreds of TBytes of memory and 1 Pflop/s sustained performance. Another great challenge will be to go for the adjoint-based inversion of complete waveforms using these 3D wave-propagation simulation models. This will lead to at least one order magnitude increase in computational requirements.

Massive access to Pflop/s computing will allow European researchers to investigate the mechanism of the dynamos and understand their physical principle. Yet the parameters available on such resources are still out by a factor of 1 million from actual geophysical values. Recent progress indicates that an earth-like solution could be reached only by decreasing the relevant parameter controlling viscous effects by a factor of 1,000. Constructing such an earth-like numerical dynamo model is therefore only realistic in the realm of Pflop/s computing. This will also require PBytes of storage to describe the 4D (time and space) magneto hydrodynamic solution.

Another relatively new computational challenge is the processing of large datasets (see all the above challenges). To date, most advanced examples of this type processed up to ~100 Tbytes of raw data. The next step will be to build applications using Pbytes of records. The mentioned analysis of data is often based on cross-correlations between different parts of records and involves consequently computing huge numbers of Fourier transforms and cross products. These types of computation can become very dependent on I/O rates and are very efficiently parallelized with architectures such as GPUs.

2.2.6 Common requirements and recommendations

All the computational requirements described above have one common feature: the urgent need for access to very large computational resources does not stem from only a single aspect, because currently available compute power restricts the applications at stake in several ways. They all require simultaneously higher resolutions, a more sophisticated representation of processes, and ensemble methods to quantify uncertainty. The need to improve these multiple aspects implies very high computational requirements. The requirements are typically a factor of 1,000 above what can be run today on the top European computational facilities. In absolute terms, the performance requirements of these applications are of the order of 1 Pflop/s *sustained*, with some of the applications having even higher longer term requirements. As a large number of such applications are concerned (see the many challenges described above), the *total sustained* performance which would be necessary is then in the range of 10+ to 100 Pflop/s. The ratio between sustained and peak performance varies from application to application; in the past, typical factors were 1:10 for scalar architectures and 1:3 for vector-processor based systems. Due to the massive parallelism present in the envisaged exascale systems, such performance ratios could be sustained only if the application codes were modified to deal with such parallelism. The peak performance requirement is therefore in the exascale range, when considering a scalar architecture³⁰. Newer mathematical algorithms might also be required.

Due to the high internal communication requirements and the continuous need to modify and enhance the model codes, a general-purpose computing system that offers excellent communication bandwidth and low latencies between all processors is required. To ensure efficient utilisation of the system, an I/O subsystem that supports very high transfer rates is essential.

On the data side, substantial amounts of online disk storage (at least 1+ PBytes today and 10+ for the 2020 timeframe) is essential. Such online storage needs to be complemented by local offline storage (at least 10+ PBytes), to enable inputs and outputs to be stored up to 12 months. A possible long-term storage strategy would be for each community to develop its own federated database joining data from multiple sites, while preserving at the same time local autonomy. As stated above this is already the case for climate and solid earth. Depending on the community, these archives would hold between 20 and 100 PBytes of data. To implement a federated data archive system, high-speed network links between the European resources and the larger of the national facilities are a fundamental requirement. Such a strong link with national facilities would enable the bulk of the pre- and post-processing to be carried out at these facilities. Equally, visualisation and analysis of model outputs would be possible through these network links.

Requirements for communities organization and development

For applications to run efficiently and productively on future multi-Pflop/s and exascale architectures, software will have to be re-designed and adapted, in particular with respect to the extreme parallelism. Software refactoring, algorithms re-engineering and new data analytics are crucial to achieve these goals. Work in these directions has already started for adaptation to the present sustained Pflop/s computers, but a lot more efforts are needed. Furthermore, ways will have to be implemented to cope with the expected high failure rate of the components.

³⁰ It should however be underlined that for applications such as medium-resolution paleo-climate models, strong scaling would be required to an extent that seems not to be possible. These applications would only benefit from increased per-core performance (*e.g.*, powerful new vector type architectures)

These trends are not limited to the high end of supercomputing, but apply to all tiers of the HPC ecosystem. The compute nodes architecture follows current technology trends and will integrate thousands of cores, including accelerators. The effort involved in software re-factoring is substantial and requires a deep knowledge of the algorithms, of the codes and, in many cases, of the physics at stake. Therefore, a successful practice is to place these activities in community projects where substantial code development activities are common and thus long-term software development activities can be sustained.

Supercomputer centres, with their profound expertise of computing architectures and programming models, have to recast their service activities in order to support, guide and enable scientific program developers and researchers in refactoring codes and re-engineering algorithms, influencing the development process at its root. These services should be provided for longer periods without the necessity for users to change their codes. The resulting codes should be adapted to both Tier-1 and Tier-0 machines, and it will be the particular requirements of the users of these community codes that determine whether applying for Tier-1 or Tier-0 resources.

Due to the complexity of these novel HPC environments and the relevance of the scientific challenges, interdisciplinary teams and training programmes will be strongly required. Training programmes will allow WCES scientists to improve their HPC background as well as to establish stronger links with the HPC community. In this respect, funding specific actions to support training activities, summer/winter schools, intra-European fellowships as well as international incoming and outgoing fellowships will play a strategic role in preparing new scientists with a stronger and more interdisciplinary background. As almost all applications are multi-physics, and given the increased complexity of the component models and of future exascale computing platforms, a lot of resources should be devoted to the technical aspects of coupling and coupler-development teams should be reinforced.

2.3 Fundamental Sciences

WP3.3 is in charge of assessing strategic development in high performance computing and how to prepare the scientific community for Exascale computing for the vast topics of fundamental sciences, covering statistical physics, physical chemistry, plasma 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 key for efficient applications to succeed 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 simulated 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 Discontinuous 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 $(6 \text{ fm})^3$ and larger. To achieve robust signals from these types of calculations, the scale 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, arXiv:1603.09303] 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 Genes. Optimized communications libraries exist as well, which utilize the low-level interfaces to the hardware to cut down latencies and to use the hardware capabilities optimally. These software packages are maintained by a broad community that is, and will continue to be, willing to invest in optimized software 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).
- Propagator contractions exhibit factorial growth in computational complexity. New algorithms need to be developed in conjunction with future architectures to alleviate this problem.

In terms of problem size, a moderate increase in the problem sizes to $256^3 \times 512$ is conceivable, but will not be required for all use cases. However, the requirements for precision for important questions such as the QCD contributions to $g-2$ are dramatically larger than 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 national or European-wide HPC sites. In this context, the recent initiative of ETP4HPC toward ESD is a real progress for our communities.

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. As more complex observables are coming under consideration, a shift of the computational efforts from scalable to less scalable algorithms may occur. Here, “modular” HPC systems such as the EU funded DEEP architecture offer a promising solution to address potentially arising bottlenecks.

Access to adequate HPC resources provided by national HPC infrastructures or the pan-European PRACE HPC infrastructure 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. 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.**

Within A goal that has been set out to for 2022 is to be able to model the gyroradius scale for ions and electrons while following the global scale dynamics is one of the main goals of plasma physics. They share with astrophysics and fusion the difficulty of bridging large ranges 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 application. In 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 becomes 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 ($\sim 1\text{fs}$) 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-of-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 being addressed 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 metres 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.

Leading European gyrokinetics codes such as Gysela5D or ORB5 or US codes such as XCG are now optimised or rewritten to (pre)Exascale architecture for ensuring increased scalability between 500k to 2M cores, extended MPI decomposition over the 5 dimensions, use of AMR methods, optimised memory consumption to 1 GB/core, I/O requirements of 200 to 800 TB per simulation written at least at 200 GB/s using multi-level hierarchical storage with application based checkpoint/restart and use of in-situ post processing.

An especially time-urgent and very challenging problem facing the development of a fusion energy reactor is the need to deal reliably with large-scale major disruptions in magnetically-confined tokamak systems such as the EUROfusion Joint European Torus (JET) today, and the burning plasma ITER device in the near future. The JET team has led the development of supervised Machine Learning Support Vector Machine (SVM) predictive methods which have achieved over 80% predictive capability for disruptions 30 milliseconds (ms) prior to such damaging events, far exceeding current HPC first principles approaches. However, ITER will require around 95% or better predictive accuracy with less than 5% false positives at least 30 ms before disruptions. Very encouraging advances in the development and deployment of deep learning recurrent neural nets have recently been demonstrated by results obtained by the Princeton University/PPPL team which have already exceeded that of SVM methods i.e., better than 90% predictive accuracy with 5% false positives at least 30 ms before the occurrence of JET disruptions. Moreover, scalability studies of the associated Deep Learning Code—FRNN (Fusion Recurrent Neural Net) Code – first to 200 GPUs on Princeton University’s “Tiger” Cluster and subsequently to 6000 GPUs on ORNLs Titan Supercomputer – provide great encouragement that sufficiently rapid training of higher physics fidelity classifiers is now feasible.

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 the GYSELA group (at CEA) 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.

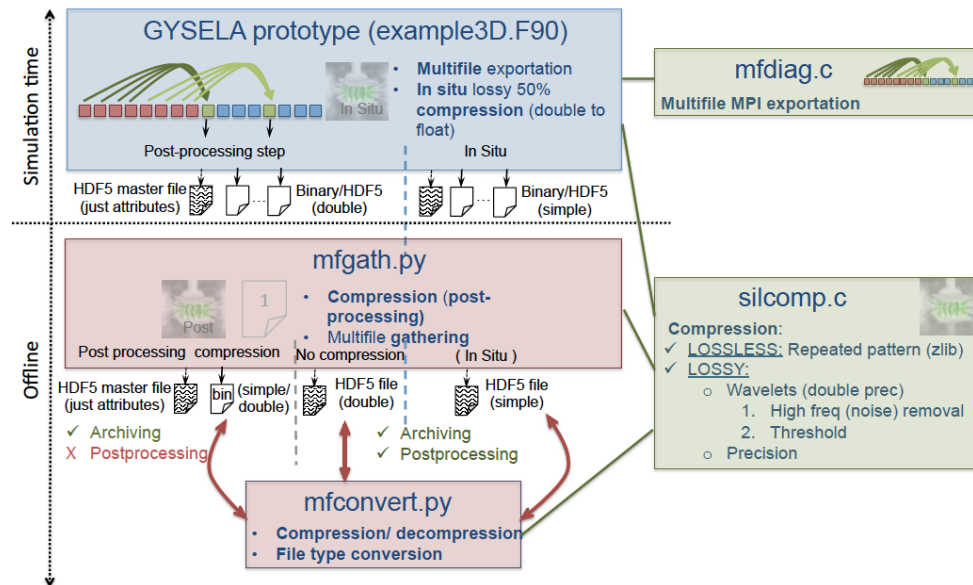


Figure 20 - first result of lossy data compression on massive 3D datasets, leading to a 26x gain in I/O bandwidth with the GYSELA code

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 carriers (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. Study of space weather provides simulations and tools to prevent such activity from impacting our technological society. Recent nationwide reports 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, a massive solar storm in 1859, 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-scale

systems with gravity, rotation, magnetism, radiative transfers, nuclear and chemical reactions etc making their modelling extremely time consuming and difficult. They share with other fields the need to model global dynamics organization while resolving small scale phenomena, often in non-equilibrium state or subject to a large range of fluid and plasma instabilities. Many key problems are very demanding in term of computing resources and physical description. These include problems such as: the origin of the Earth and solar magnetism and its 11-year cycle, the formation of stars and planets, how accreting disks behave, how radiation fields ionized the Universe, how galaxies form and evolve to create the Hubble sequence, and how stars blow up and enrich the interstellar medium. Linking all these scales together is also a tremendous challenge for the years 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-scale, 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 problems 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 orders 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 dynamos, 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 simulation. **Huge data management of 4-D structures (space + time)** are required in order 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 is becoming more and more common and requires multi-level parallelism to have an ensemble approach while computing highly-resolved 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 makes 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 in order to be able to 1) simulate bigger volumes for rebuilding the redshift space with higher resolutions and 2) consider non-linear collapse phenomena of dark matter. Finally, a third objective is the possibility to take into account several fluids such as baryonic gases and hydrodynamics into n-body simulations for understanding the formation of macro-stars.

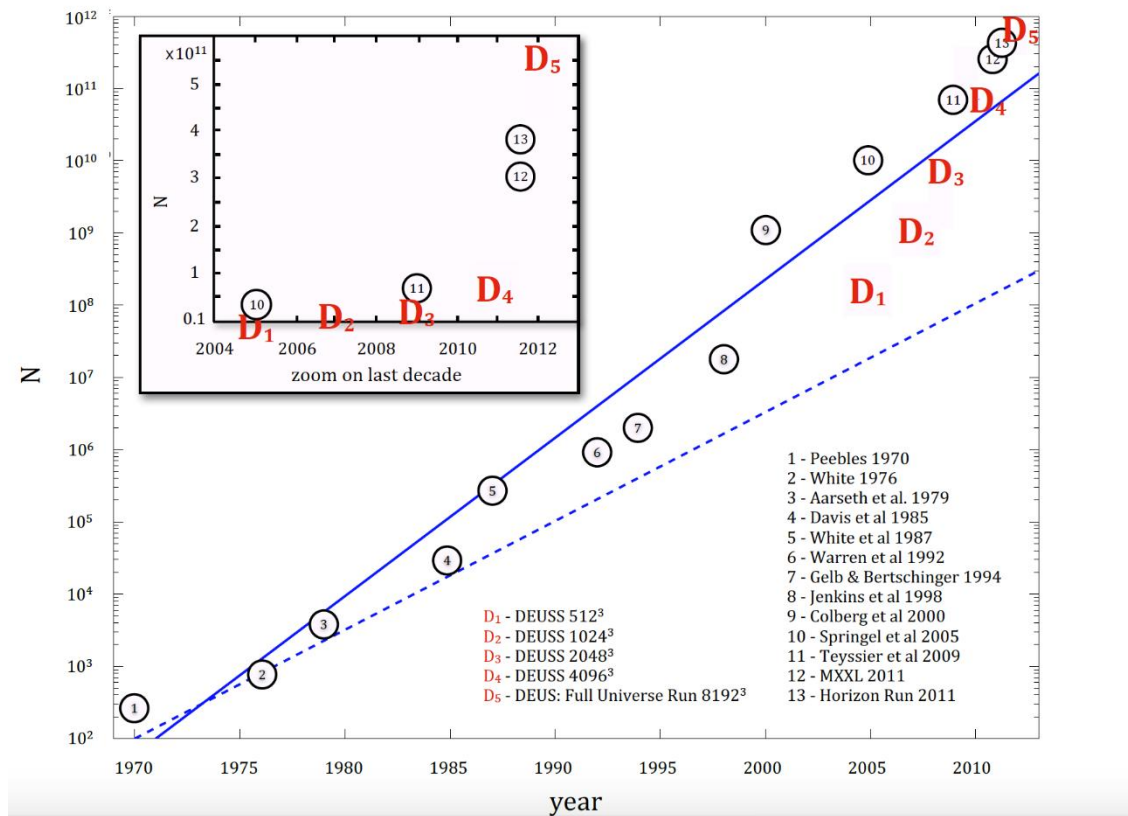


Figure 21 - evolution of the size (number of particles) of dark matter simulations over the time

This will lead to upcoming **multi-scale multi-physics dark matter simulations with a resolution of 32384³ particles** (instead of 8192³ or 16384³ now), able to access 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 ESA's EUCLID mission, which aims to map the geometry of the dark energy, which could represent 75% of the Universe.

Such simulations are generating massive amounts of data (for example already 150 PB of rough data for a 8192³ 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, to be able to follow, for example, only pertinent turbulent structures. Recently, in 2017, multiple teams based in Switzerland and China performed massive dark matter simulations over 2 and 10 trillion particles respectively, using more than 4000 GPUs on PizDaint (CSCS, Switzerland) and 1.5 million cores on the Sunway TaihuLight system (Wuxi, China), using fine-tuned n-body simulation codes (AMR codes, mixed precision, optimisation of memory footprint, load balancing, ...).

Astrophysics (or Beyond Earth and Universe Science) communities are well organised and federated at the national and international levels around more and more complex space missions (e.g. EUCLID, LISA, INTEGRAL, SVOM, ATHENA, InSight, COPERNICUS etc.), large instruments (e.g. LSST, CTA, SKA) and observation systems (e.g. KM3NET, ASTROMEV, LIGO/VIRGO, EarthCube, EPOS-IP, etc.) often run by international inter-governmental consortia and agencies. The communities place a great premium on internationally distributed and federated data resources for archiving and distributing data and have pioneered the prevailing philosophies of globally shared and open data with internationally approved

standards for data, metadata and exchange, including provenance and interoperability, together with a growing commitment to open science.

Building federated computing and data analysis platforms featuring persistent storage capability to capture and aggregate the large volumes and the diversity of data (events, time series, images) from the distributed archive resources associated with the observation systems, and from large numerical simulations (virtual instruments), and to accelerate the full path of data use, can best be accomplished by engaging a research-driven strategy aligned upon (see following figure) :

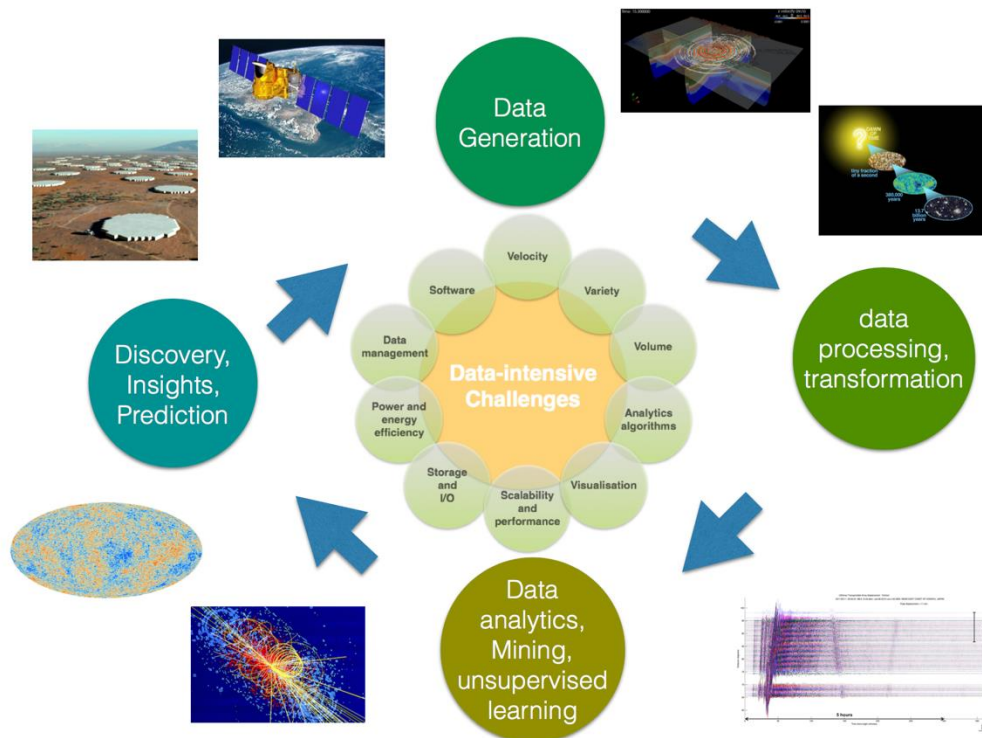


Figure 22 - Full data cycle for Earth, Planetary and Universe Sciences (copyright BDEC)

This new strategy driven by the deluge of data generated either by instruments (sensors, satellite, telescopes, ...) or by massive numerical simulation (advanced physics-based stochastic simulations of multi-physics and multi-scale natural systems) is leading to the development of **innovative big-data-driven methods** accelerating the full path of data use, with **data-stream end-to-end workflows orchestrating data processing and physics-based or data-driven statistical analysis methods**, in a Bayesian inference framework featuring machine/deep learning (ML/DL).

2.3.5 Challenges on Material Sciences

Europe is recognized to be leading in the development of scientific codes for first principles simulations in materials science, i.e. applications that allow predictive simulations of materials and their properties from the laws of quantum physics and chemistry, with no need of previous knowledge and empirical parameters.

Today, the use of these codes is still limited in terms of the complexity/size of the structures to be studied and of the availability of affordable approximations to treat many-body interactions.

The exascale perspective can boost their massive use in two main directions: on the one side, the accurate prediction, understanding and design of realistic complex materials, nanostructures and devices; on the other side, the high-throughput screening of vast classes of materials with desired properties.

The main expected scientific impact is the discovery of new materials, as well as novel phenomena, functionalities and processes, e.g. in synergy with the new experimental spectroscopy/microscopy facilities deployed in Europe. Theoretical and computational spectroscopies will also be possible through the adoption of a more accurate treatment of interactions at the exascale.

Major challenges that will be tackled include, for example:

- the development of high-performance materials for energy harvesting and storage, for water filtration, or for efficient and environmentally sustainable production processes;
- the simulation and design of the next generation electronics, where nanoscale features require radically new approaches with respect to current standards: full microscopic and quantum description of materials, interfaces, devices, as well as implementations of quantum information technologies.

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).

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 papers. Experts from this working group 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. 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].

By 2021, a new generation of post-DFT methods 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.

In the high-throughput perspective, we consider PRACE the (pre-)exascale engine that will enable the production of a huge mass of materials science data from quantum simulations.

To focus on data aspects of relevance for EDI (European Data Challenges), experts envisage major additional challenges. Disk storage will not be the main point, although we will continue to need a storage 5-10 times the size of the volatile memory of the allocated partition (see above). The key will be its integration in a multi-tier hierarchy (NVM+SSD+HardDisk+Tape), where the first tier, NVM, serves for near-online or interactive analysis of data that do not need longer term storage (and can be recalculated if needed). A federation of data and instruments throughout different sites would best allow for analytics and AI inference from data. Such federation process should be developed for distributed and federated DBs, as well as virtual machines/containers and analytics/AI tools. Making containers available for access throughout the European HPC systems will be important to open HPC to a broader pool of users including industries and SMEs.

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 multi-scale 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 can 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.

First principle materials science applications, typically used in most HPC centres (involved in MaX, eCam, Nomad first, but in EoCoE, Bioexcel, Compbiomed as well) are based on community or proprietary codes maintained by groups of core developers that are themselves most often users of their applications. Most of the users running these codes on HPC systems are not developers themselves, and looking to the future, they would better benefit from the provisioning of these applications “as a service”, in a “separation of concerns” approach. Their “time-to-science” could be reduced if they could be shielded from the specific details of this or that HPC system (e.g., they do not care about the presence of GPUs if they do not have to take care of the required code modifications).

To make this “separation of concerns” possible, in the future (exascale or post-exascale), the provisioning of these applications should be combined with the provisioning of optimized “HPC containers” (the specific technology may change over time) featuring these applications and the whole software stack involved in complex workflows. One can imagine a material science library of HPC containers, maintained by the CoEs/PRACE/EDI specialists, for all the Tier-0 and Tier-1 EU HPC systems, where containerized workloads have been enabled, on which the scientific and industrial users may rely. The same library could be used on different HPC systems as well, provided they have deployed a compatible container technology.

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 EESI-2 projects is still one of the key priorities.

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 such as 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 an Exascale resource for a single calculation!

The experts are proposing the following 6 recommendations:

- The need to have a Centre 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 focus on turbulence, gravity, magnetism and nonlinear phenomena in complex fluid/plasma systems.
- One should assess what is crucial to be developed in the EU with the support of the EC and national funding agencies and what could be leveraged by collaborating with countries outside the 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 fields 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 to be increasing, especially after the decision to renew PRACE took so long. Looking at the recent TOP 10 HPC list, Europe comes

only at #3, 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 innovative and developing better methods buys you a factor of 100 against every factor of 10 gained through Moore’s Law. Every country in the world invests far too little in software and expert computational scientists and far too much in big computers.
- 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 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 should not 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.

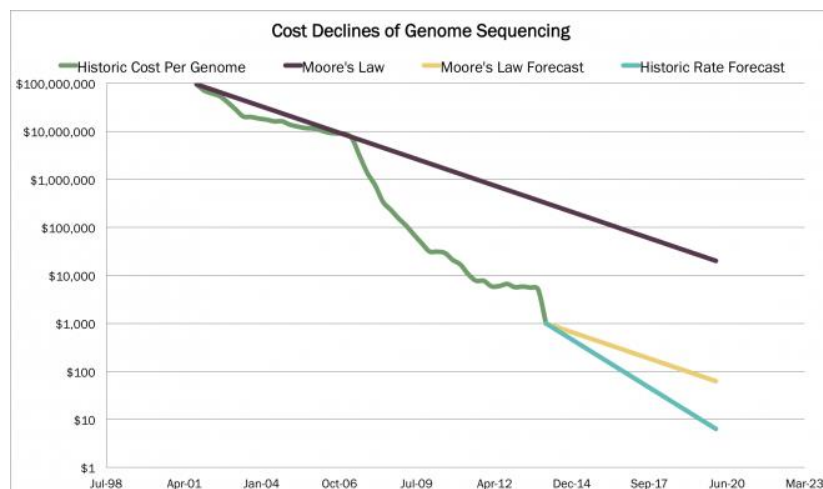


Figure 23 - evolution of the cost of genome sequencing vs Moore' law

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 improved 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 project³¹) 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

³¹ <https://www.genome.gov/26525202/encode-pilot-project>

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**.

2.4.2 *Biomolecular modelling and simulation*

In the biomolecular area, detailed computer simulations add an essential extra dimension to the 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 will become increasingly important.

Coupling coarse-grained molecular dynamics, Brownian dynamics or even higher-level models (reaction-diffusion models for example) in **multi-scale/multi-physics simulations will require both the availability of flexible, hierarchical resources but also hierarchical and concurrent multi-scale 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 **are now mandatory to develop**. 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).

The human genome has revealed over 20,000 expressed proteins, the proteome, which are the workhorses of our life, performing all kind of critical functions in our cells. These proteins undergo further all kinds of post-translational modifications that affect and control their fate and function, increasing greatly the complexity of the proteome. Proteins however do not work in isolation, but by interacting with each other and with other biomolecules, form a complex network called the interactome. In humans, this interactome consists of hundreds of thousands of protein-protein (and other molecules) complexes. It is a dynamical network whose wiring might change as a function of the cell cycle and other regulatory events such as post-translational modifications. Miscommunication in this complex network can be the origin of diseases, which is why it is important to understand how this network functions at atomic level. Thus, this requires adding the 3D structural dimension to it. Considering the size and complexity of this network, it is clear that experimental methods alone will not be able to provide all answers. This is where biomolecular modelling, and in particular docking, can play a crucial and complementary role. There are two main challenges requiring exascale computing:

1. Adding the structural dimension to the hundreds of thousands of interactions.
2. Given the knowledge of the proteome, predicting what the interactome will look like.

These require **hundreds of millions of docking runs generating hundreds of exabytes of data.**

Molecular dynamics is a powerful tool that can provide insight in molecular processes in atomistic detail. Due to short time steps (femtoseconds) compared to the time scales of biologically relevant transitions, which are on the microsecond to second range, an enormous number of integration steps is required. The computational cost is dominated by the non-bonded pair interactions, which can be computed efficiently on modern SIMD or GPU hardware. Simulating biomolecules is a strong scaling problem, because the system size is fixed (and larger systems come with larger time scales). A simulation of a **typically sized system of 200,000 atoms takes about 10 exaflop per microsecond of simulation time**. A single simulation is not an exascale problem. But nearly all problems of interest involve calculating distributions or free energies and how these vary with e.g. different ligands and/or protein mutations, which requires more sampling and increases the cost by orders of magnitude. **In current large-scale studies the number of combinations times the simulation time needed per case reaches a zettaflop/s.** There is parallelization at several different levels, for a typical case listed from tighter to weaker coupling: 1) Within the simulation of a single copy of the system, 2) when computing a free-energy for one molecule and 3) when computing free-energy differences between ligands or mutations. These parallelization options multiply, producing the potential for massive parallelization.

Exascale computing, combined with specific theoretical developments in *ab initio* QM/MM codes (such as CPMD³²), in particular a coherent approach for a quantum grand-canonical statistical ensemble in a QM/MM scheme, will allow a totally novel class of biophysical

³² <http://www.cpmid.org>

problems at the first principle level to be discovered, including the investigation of realistic large biochemical systems (up to millions of atoms) involving electron and proton transfer, quantum electron dynamics, excited state dynamics and quantum nuclear effects. First principles QM/MM is currently one of the most powerful methods to investigate enzymatic reactions and photobiophysical processes. It is also extremely useful in many other problems, such as, for instance, the investigation of metalloproteins (more than one third of the human genome). The importance of the method is acknowledged by the recent Nobel Prize in Chemistry in 2013 for multi-scale modelling as well as by the myriad of successful applications in the field. However, a significant fraction of quantum processes faces theoretical challenges associated with solvent/solute diffusion. Examples include proton transfer and diffusion mechanisms in catalytic sites of enzymes. In these cases, imposing constraints on the solvent is not a physically sound approach. A fully satisfactory theoretical framework is instead the implementation of the so-called Hamiltonian adaptive multi-scale (QM/MM) scheme. Here, the description of the diffusive particle (QM versus MM) is updated as they diffuse away from the quantum region. In a transition region, between the QM and the MM ones, the diffusive particles continuously change character, becoming partially QM and partially MM. These schemes currently face challenges in correctly describing the structures at the boundary between the two regions. The difficulty lies of course in the complex QM potential energy expression, with a many-body expansion that contains higher order terms. Solving this long-standing problem will allow microcanonical adaptive first principle QM/MM simulations. This in turn will allow predicting rigorously the energetics of a myriad of processes involving water or ion diffusion.

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 biomolecular 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 at the level of 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 coupling 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 of 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 (around 4 cores) and can require longer run times (several days);
- Multi-scale/multi-physics simulations typically involve a 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.

Europe is a hub which has fostered a range of cardiovascular simulation codes that have been gearing towards the use of HPC. The EU needs to provide the computational platform and expertise to create HPC multi-physics codes to address the gap between gene expression and the onset of illnesses. Internationally, both the USA and Japan are highly advanced in this field. The UT Heart, from the University of Tokyo, Japan, is the most detailed, complete HPC, multi-scale, multi-physics simulations code to understand the cardiovascular system to date. The EU has fallen behind the USA, which now supports the use of simulations for non-invasive diagnostic applications through their main regulatory body, the FDA. Future developments in the field will require not only investment in scientific expertise but also the fostering of collaborations with regulatory bodies throughout the EU.

Mechanobiology is a rising star in fundamental life sciences. Most biology has assumed cells to be best represented as chemical machines, totally neglecting the other, vital, biophysical

aspects. For example, if you mildly deform the nucleus of a cell, changes in DNA folding can induce large changes in the proteins transcribed from the same genetic code. Considering every cell is attached to some tissue that deforms or some fluid flow, neglecting these mechanisms misses many processes fundamental to understanding healthy function and disease. **Europe leads the world in multi-scale modelling of the skeleton, and in the development of some of the methodological aspects of simulation of muscular skeletal systems such as the elastic registration.** However, leading researchers from USA has a better track record of scaling up modelling, and is able to study problems large enough to require Tier-1 or even Tier-0 resources.

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 increased 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*

Data intensity: New methods are needed for researchers to deal with increased volumes of data coming from advancing experimental methods. Often HPC resources and experimental facilities are not co-located and TBs of data must be transferred. This should be viewed in the context of the needs of machine learning approaches as well as those of simulation.

Data security: Biomedical applications require the storage, transfer and processing of private and sensitive information. For further progress to be made both technical and policy work is required to allow researchers to use data for research legally and safely. Similar concerns also impact industrial collaborations, although in this instance much progress could be made through adoption of standards such as ISO security certification at HPC or public cloud centres.

Complex Workflows: Many biomedical applications consist of many distinct stages (such as image processing, mesh building, mechanical simulation, etc) which must be connected in order to model a given system. This requires continuing developments in workflow tools that facilitate this, eg. by providing convenient tools to ensure input and output formats match, units are checked, etc.

Heterogeneity of codes: Related to the above point many biomedical simulation projects involve many different codes all of which frequently have different computational requirements (for example some codes benefit greatly from GPGPUs). Frameworks that help to facilitate running such applications and potentially migrating them between resources (i.e. HPC and cloud) are needed.

Urgent Computing: Use of simulation in clinical settings will require the ability to demand computation as and when needed with a guaranteed turnaround time for jobs.

Validation, Verification and Uncertainty Quantification: In order to produce actionable predictions for use in medical situations it is key that uncertainties are well understood. This is rarely, if ever, the case for complex workflows. Frameworks to facilitate VVUQ are needed and expert statistical knowledge must be leveraged to ensure validity.

Resources that can be leveraged

There is considerable scope for making use of the knowledge and infrastructure present at large experimental facilities such as ESRF/ILL (Grenoble, France) and ESS (Lund, Sweden), as well as Diamond/ISIS (Diamond Light Source and the ISIS neutron and muon facilities) in the UK. Options might include co-location of HPC or provision of high speed data links to appropriate resources.

A number of requirements need to be satisfied in order to address the challenges described above. The list can be quite long and detailed but some of the main aspects are listed below:

- Dealing with millions of independent jobs
- Efficiently simulating the long-range electrostatics treatment
- Allowing strong scaling of single simulations
- Computing differences in free energies between ligands or mutations
- Allowing microcanonical adaptive first principle QM/MM simulations

Moving towards post-petascale computing, hardware systems and software stacks need to address a number of requirements, as detailed in the table below:

Exascale aspects	Requirements
HPC System Architectures and Components	large width vector units, low-latency networks; high-bandwidth memory; fast transfer rates between CPUs<->accelerators; heterogeneous acceleration, floating-point
System Software and Management	dynamic (task) scheduling, support for adaptive scheduling of workflows
Programming Environments	standardization, portability, task parallelism, fast code driven by e.g. Python interfaces, implementations accessible also at sub-Exascale levels
Energy and resiliency	distributed computing techniques to handle resiliency/fault tolerance
Balance Compute, I/O and Storage Performance	post-processing on the fly, data-focused workflows, handling lots of small files in bioinformatics
Big Data and HPC Usage Models	proximity of data generation and analysis/visualization resources, workflows, machine learning for analyzing simulation data, high-throughput sampling, efficient HPC and HTC job scheduling
Mathematics and Algorithms for extreme scale HPC systems	multi-scale algorithms, task-parallel algorithms, electrostatics solvers, ensemble sampling & clustering theory, ensemble simulations

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 datasets which are often highly sensitive. The former can be further broken down into the needs for hierarchical hardware and a multi-scale modelling software stack. This 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)

Multi-scale 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 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, and 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 in terms of 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 Centre of Excellence where the molecular dynamics simulations engine GROMACS (www.gromacs.org) currently supports the latest high-end GPU processors, Xeon Phi processors, 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) 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 Exaflop/s 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.

Verification, Validation, Uncertainties Quantification

In many fields, a significant issue is a lack of resource (computational and financial) for reproducibility and sensitivity analysis. For simulation to provide actionable outputs (both

clinical and industrial) it is clear that simpler tools are needed to conduct replication studies and to vary parameters as well as provide meaningful error estimates.

Cross disciplinary research

The grand challenges faced by biomedical computation will require collaborations which bring together many fields of expertise. It is vital that joined-up training is provided which allows researchers from different fields to work together on common projects.

Machine learning

Outside the EU, we have witnessed the rise of deep learning in contacts prediction as seen in the latest CASP (Critical Assessment of protein Structure Prediction) experiments [45], which noticeably led to performance improvements in the last round, the results of which will be presented in December 2017.

The USA specifically is investing heavily in start-ups 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>)

Training and education

There is a need to train experts of theoretical modelling as well as experts in HPC simulations to address the above challenges. At the same time, it is necessary to raise awareness among the medical and pharmaceutical community of the technologies that could support their research. Bridging together the medical/pharmaceutical community with the modelling and simulations community, not for solving specific issues but to address the grand challenge of predictive medicine, will be a crucial step for future development.

In this respect, PRACE has to act as a facilitator and an enabler of the creation of the community of HPC centres in Europe, which should become more than just the HPC infrastructure frameworks. The centres have to become domain-specific knowledge hubs aggregating academic and private research institutions.

3 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, foster the creation of new CoEs in Europe, stimulate the development of scalable standard tools (solvers, meshers, couplers, ...), and increase training and education in new fields.

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

3.1.1 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 multi-scale/multi-physics 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 largely outpace 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 EESI-2 project issued several recommendations in 2015 to the European Commission and national agencies 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 proposes 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 consider that, as correlation is not reason, deep learning methods will not replace a true understanding of the key physical components of scientific problems.

3.1.2 *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 ecosystem.

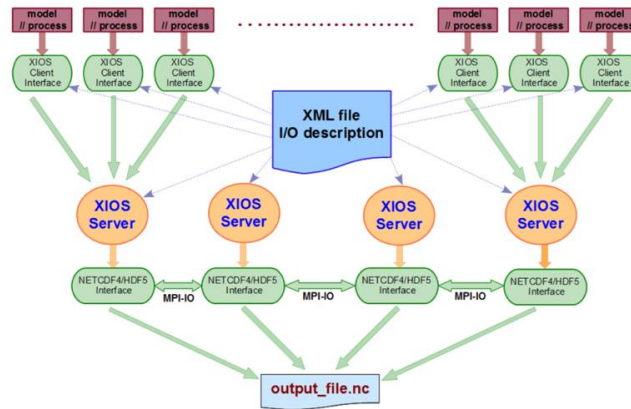


Figure 24 - overview of XIOS (developed by IPSL), an asynchronous in-transit I/O library for climate simulations with less than 6% overhead

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 the study of the structure and dynamics of numerous physical observables such as the velocity fields, the gravitational potential, the deflection gravitational field, the gravitational wave field, etc. The use of the outcomes of deep learning technology tools 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 standard “on the fly” 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.

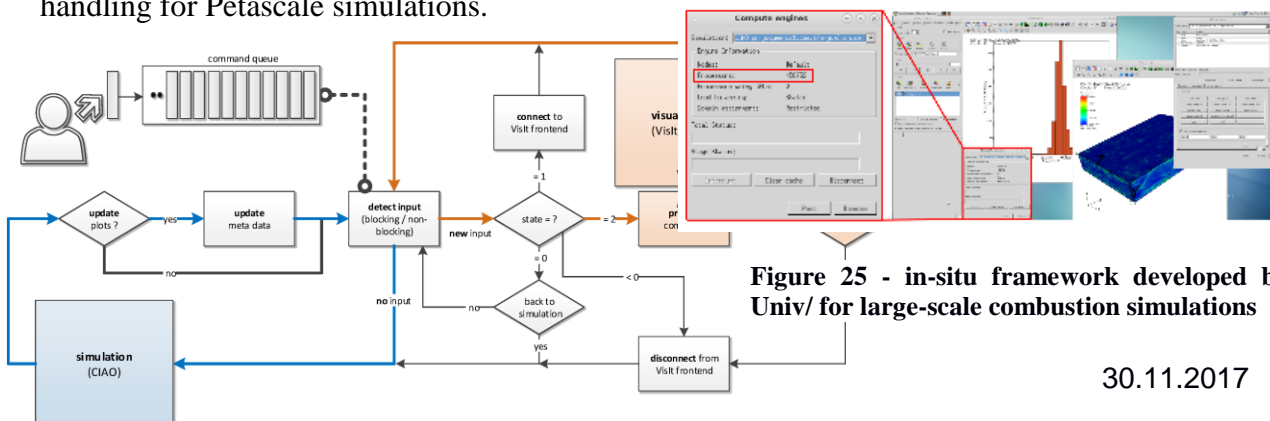


Figure 25 - in-situ framework developed by Aachen Univ/ for large-scale combustion simulations

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.

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

In 2015, inside the set of recommendations issued by the EESI-2 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 multi-physics/multi-scale simulations become impossible to process in a competitive time using 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 are benefiting 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. They 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 also requires a number of other components to be implemented, including: 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 have become more widely adopted since the 2000s, with machine learning techniques driven by the 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 such as Google, Amazon, Facebook, nVIDIA, and Microsoft have developed machine learning frameworks such as Torch, SPARK, Mahout, TensorFlow, DMTK, Shogun, Caffe, Theano, Scikit-Learn, which are now also beginning to be used in life sciences and particle physics.

Europe possesses a strong expertise in these new techniques, both in academia (with Inria and CNRS research teams in France, BSC in Spain, Technical University of Munich in Germany, and the Alan Turing Institute in UK), as well as in industry (with the Facebook FAIR laboratory and the Sony CSL laboratory in Paris and the IBM R&D lab in Zurich). **It is 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” [49] and the work conveyed by EXDCI WP3, **the European Commission and the national funding bodies were asked to organise an urgent joint call for proposals 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 proposals could be organised jointly with BDVA and could lead to data-intensive software demonstrators to be used on the upcoming Extreme Scale Demonstrators, to be 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 fund up to 12 pilot projects to create a bridge between the skills of experts in domain science from many scientific and industrial fields on the one hand, and those of deep/learning and HPC experts on the other.

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

As an example, it is interesting to note the recent breakthrough obtained in mid-2017 by Intel and NERSC on the Cori Machine (a 29 PFlops KNL system) using 9600 nodes (15 PF sustained) with supervised and unsupervised Deep learning methods for detecting on-the-fly extreme climate events:

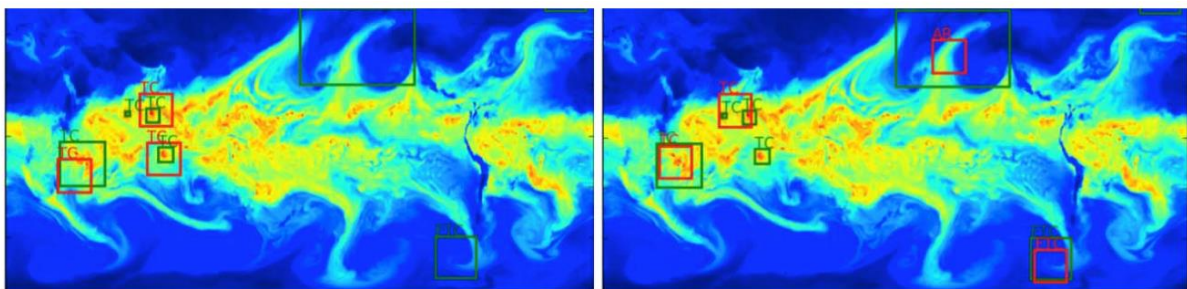


Figure 26 - Identified extreme climate events using supervised (left) and semi-supervised (right) deep learning. Green = ground truth, Red = predictions (confidence > 0.8). [NIPS 2017]

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

In the context of an explosion of the volume of data generated by large-scale instruments (telescope, satellite, network of sensors, sequencers, IoT, etc.), the use of supercomputers becomes mandatory to be able to post-process and value in a competitive time the information acquired. In the same way, being able to couple medical devices located in hospitals (scanners, microscopes, etc.) with supercomputers, or being able to allocate on the fly a partition of a supercomputer for urgent computing decision-making, in the case of major event (such as earthquakes, hurricanes, pandemic propagation, (cyber)terrorism) is also now critical for both doctors and public bodies.

Developing these new services for national or European HPC infrastructures (such as PRACE) is now a must, leading to an evolution of the operational procedures of the centres on 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 in EESI-2 reports and some actions have been implemented inside the FETHPC-02-2017 call for proposals (which closed in September 2017) but WP3 experts are insisting in the importance of implementing such services.

3.3 Development of new Centres of Excellence in Europe

In 2015 the European Commission selected 8 (increased to 9 in mid-2016) Centres of Excellence (CoEs) for computing applications. These aims to help strengthen Europe's existing leadership in HPC applications, covering important areas such as renewable energy, materials modelling and design, molecular and atomic modelling, climate change, Global System science, bio-molecular research, and tools to improve HPC applications performance. They also aim to 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 the 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 the first EESI project proposed some years ago (talking about Co Design Centres instead of Centres 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 proposals toward applications enabling. In this call 22 proposals involving 45 research teams benefited from close to \$40 million³³ 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 calls for proposals will be issued 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 fields of: Online Data Analysis and Reduction at the Exascale

³³ <https://www.hpcwire.com/2016/09/07/exascale-computing-project-awards-39-8m-22-projects/>

(CODAR); Block-structured AMR numerical methods (developing a new framework called AMReX, in support of the 22 applications); Efficient Exascale Discretisation (CEED, including work on scalable meshers, discretisation models, ...); and Particles Applications (CoPA, acting as a centralised clearing-house 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.

3.3.1 *Engineering and industrial applications*

Regarding the domains already covered by the first set of 9 Centres 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, on the one hand, either a strong established roadmap toward Peta and Exascale (e.g. TOTAL, Repsol, ENI and Shell in the Oil & Gas sector; EDF in the field of energy; Porsche, Renault and JLR in the automotive sector; and Airbus, Safran and BAE in aeronautics), or, on the other hand, a potential need for 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 expand its activity to other engineering domains. Sustainability of this CoE should be ensured by specific services provided to industry-like user support of (open-source) software, licensing or specific tailored developments.

3.3.2 *(Open-source) software sustainability*

The US National Science Foundation (NSF) launched the Software Infrastructure for Sustained Innovation (S2I2) Program in 2010³⁴ to fund software research at multiple scales, with the aim of transforming innovation in research and education into sustainable software resources that are integral to cyberinfrastructure.

In July 2016, the NSF announced 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.

Among the missions of these institutes, one is related to the industrialisation, promotion and long-term support of open-source applications developed by the NSF.

Similar missions also exist inside US DOE-funded research laboratories, and have led to the adoption many years ago by the whole scientific community of packages such as BLAS, LAPACK, PETSc, HDF5.

³⁴ https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503489

In Europe, a few isolated initiatives of this type exist, the most visible being the UK Software Sustainability Institute³⁵, founded to support the UK's research software community - a community that includes the majority of the UK's computational science researchers. The SSI is federating a network of UK applications developers and users, providing training and software evaluation services, issuing best practice guides, etc.

The European Commission and the national funding agencies are supporting the development of a vast ecosystem of scientific applications developed by researchers. Even when they reach a given level of maturity, these applications are not properly industrialised due to a number of factors, including: lack of time (industrialising applications is 5 to 10 times longer than developing it), lack of skills, lack of interest, lack of career objectives/recognition of the developers, high turnover of people, rules, structures, etc. The result of this is a massive loss of potential usage by wider research communities and industry.

Regarding the use of scientific software by industry, many reports issued by EESI-2, PlanetHPC (<https://www.epcc.ed.ac.uk/planethpc>) projects or by PRACE have highlighted the fact that industry will use research software if the following criteria are met: it reaches a high level of industrialisation (or TRL, Technology Readiness Level); industry is aware of its existence; training is available; and user support is provided.

Despite the launch of the 9 CoE in 2015 and 2016, the service of software sustainability is not always present in the main missions of these centres and by definition this service could be leveraged by a transverse action across the CoE. This remark also applies to other H2020 funded research projects.

The recommendation of the WP3 experts is to set up such a transverse CoE in order to interface with the vertical thematic CoE, with the aim of industrialising developed (open-source) scientific software once it reaches a given state of TRL. This CoE should also have a strong liaison with national HPC research infrastructures and European ones (e.g. PRACE) in order to work on training of best practices, availability and 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 such as user support of (open-source) software or tailored industrialisation of in-house software.

3.3.3 *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 the context of a strong convergence between HPC and big data, generated from large-scale instruments (e.g. sequencers, satellite, telescopes, accelerators, network of sensors, and IoT) or from supercomputers (massive 3D complex simulations, coupled multi-scale multi-physics 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).**

In 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

³⁵ <https://www.software.ac.uk>

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.

3.4 Development of EU-wide standard tools including scalable solvers, meshers and code couplers, and optimisation and UQ frameworks

These recommendations encompass 3 previous EESI-2 recommendations based on the finding that the next generation of HPC applications able to scale up to Exascale machines – both in capacity and capability – will require the development of standard transverse tools providing scalable meshers, scalable solvers, code couplers and unified frameworks for UQ and optimisation studies.

3.4.1 *Code Couplers*

As stated by many working groups, efficient and scalable code couplers is a key issue for addressing large-scale multi-physics and multi-scale complex simulations that will be the prime driver for Exascale.

Despite some existing specific initiatives (in climate, energy, combustion...), there is a crucial need to develop **new and common European-wide coupling methodologies** and tools in order to support major scientific challenges in research (evolution of the climate, astrophysics and materials) and engineering (combustion, catalysis, energy, ...).

This could bring the following features:

- the development of standard coupling API in order to enable interoperability;
- scalable performance on new (preExascale) HPC architectures and programming models;
- take into account the performance of localization process (needed for the weight and address calculations required for the interpolation of the data between the components meshes) at the beginning of the coupled computations, and optimization of this process for geometrical or mesh changes during the simulations as well as of communication performances between the coupled models have to be increase by use of asynchronous processes, hide of operations by computations, intelligent search algorithms...;
- the implementation of new methods avoiding centralization of data and preferring distribution, even for high order interpolation methods.

3.4.2 *Unified frameworks for UQ and optimisation*

In the field of large-scale scientific simulations of complex phenomena involving multi-scale and multi-physics models, the use of massive datasets acquired from large-scale instruments (network of sensors, (radio) telescopes, satellites, etc.), or the analysis of large amounts of data generated by computer simulations, can give rise to associated uncertainties in the numerical simulation process. These uncertainties can arise from different sources, e.g.:

- Lack of knowledge on a physical parameter (epistemic uncertainty),
- Parameter with a random nature (aleatory uncertainty),

- Uncertainty related to the model (model error, too simplified model),
- Uncertainty related to the numerical errors (numerical errors of the model, to the input and output data ...).

Taking into account these uncertainties is essential for the acceptance of numerical simulation for decision making. These uncertainties must be integrated in the verification and validation process of the simulation codes. This process is now commonly called VVUQ (Verification, Validation and Uncertainty Quantification). Verification consists of checking that the equations underlying the code are correctly solved. Validation is the stage during which the predictive capability of the numerical model is checked against experimental data or a reference model. Uncertainty quantification consists of defining the uncertainties that taint the output of the simulation code.

The rise of Exascale, with future simulations spanning over billions of threads executed on many-core heterogeneous devices, dealing with complex storage hierarchies and software stack, will reduce dramatically the deterministic nature of simulations, models and results generated. In order to be able to take into account such upcoming technologies and continue to trust the results obtained (with regards to the reduction of experiments) it has become mandatory to develop solid VVUQ methodologies and tools, and Europe has a strong card to play in this field, e.g. OpenTurns³⁶ and URANIE³⁷, while there is international competition in the USA from DAKOTA³⁸ and the PSUADE³⁹ Uncertainty Quantification Project by Lawrence Livermore National Labs.

The recommendations go to the **provision of a unified European UQ framework** allowing a fast development (without re-inventing the wheel) of new scalable UQ methods (based for example on Surrogate models and reduced basis models) and their large availability on HPC infrastructures.

The **same concern applies to optimisation frameworks** as well.

3.4.3 Scalable solvers

It has often been said that in parallel computing there are three key factors: data movement, data movement and data movement. The cost of moving data is consistently recognized as the biggest obstacle to Exascale computing, both within and between computing units (processors).

It is now necessary to **go beyond traditional communication overlapping** methods, to develop new communication avoiding **auto tuned solvers**, **maximising the use of overprovisioned FP capacities** (FP64, FP32, FP16 with mixed precision) and **memory hierarchies** of new compute nodes.

European researchers co-invented communication avoiding algorithms for dense and sparse linear algebra, e.g. [46, 47, 48]. These algorithms are able to provably minimize communication and are based on novel numerical algorithms and techniques. This leads to the design of a new generation of algorithms that reduce the number of communication and synchronization instances to a minimum, and hence drastically reduce the communication cost with respect to classic algorithms.

³⁶ <http://www.openturns.org>

³⁷ <https://sourceforge.net/projects/uranie/?SetFreedomCookie>

³⁸ <https://dakota.sandia.gov/release-notes-headings/uncertainty-quantification-uq>

³⁹ <https://computation.llnl.gov/projects/psuade-uncertainty-quantification>

In addition, European researchers have been extensively exploring hierarchical algorithms based on compression techniques as H-matrices and fast multipole methods, which are very relevant in this context. Hierarchical algorithms are inherently capable of reducing communication as well as synchronization at all levels of the on-chip and off-chip network system.

Here the recommendation goes to **coordinate the multiple groups working** on many different aspects and in different algorithmic areas toward **new EU-wide numerical/data mining libraries**. This coordination could be done by a **new ETP (European Technology Platform) in Applied Maths able to reach a critical mass** by gathering together all the sparse research teams in Europe and ensuring a wide diffusion of these new tools across science and industry.

3.5 Development of training and education

Major disruptions are expected due to the transition from Petascale to Exascale. There is therefore a need to develop both initial training and permanent training in the programming and use of (pre)Exascale machines driven by complex hardware (including manycore heterogeneous computing, hierarchical memory/storage), new programming paradigms/tools (standard ones “à la” MPI+X, DSL, or new tools coming from the data analysis side including python, the use of Jupyter notebooks, ...) and the convergence with HPC. **All of this means that a major effort in training in the area of HPDA and AI is required across Europe.**

PRACE, both at the European level and at the national level, is already providing a rich curriculum in terms of programming models, optimisation on Tier-0 architectures, new programming models, uncertainties quantification, etc., and the CoE have also started to provide HPC training which is more focused on applications enabling. **These joint efforts need to be amplified in order to meet the requirements of future developers.**

4 Conclusions

During the execution of the EXDCI project, WP3 worked actively on setting up 4 working groups and involving an active network of more than 40 experts across Europe in many scientific and industrial disciplines. A strong link with the European Centres of Excellence has been also established, especially with 6 of them (ESIWACE, E-CAM, EoCoE, MaX, BioExcel and CompBioMED).

The 4 working groups produced an extensive set of elements describing scientific and technical challenges and highlighting specific and global recommendations. This final report will be sent to the PRACE Scientific Steering Committee (SSC) in order to produce a third version of the PRACE Scientific Case in early 2018.

Also, WP3 has worked with ETP4HPC (within WP2) on the elaboration of the new version of the Strategic Research Agenda (SRA) and with BDVA in the setting of use cases related to generation/use of big data for computational science.

As the Exascale era is approaching, with a convergence between HPC, HPDA and AI and the strong requirement to provide HPC systems 50 to 100x more efficient on real applications, within a constrained energy envelope, experts from WP3 made several global recommendations toward addressing in Europe this convergence from an application point of view, developing centers of excellence, developing EU-wide standard tools (meshers, solvers, UQ and optimisation tools, ...) and finally accelerating efforts toward training and education.

Europe owns a lot of assets in terms of application development, use of HPC by industry (big companies as well as SME), skills in applied maths for AI, development of a pan European HPC infrastructure or generation of data from large scale instruments or simulations. With **more coordinated efforts and funding** Europe could be ideally positioned in the converged Exascale coopetition with USA, China or Japan.

5 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 Modelling, oil & gas
Norbert Kroll	DLR	Aero, External aerodynamics
Denis Veynante	Ecole Centrale Paris, CNRS	Combustion, turbulence
Henri Calandra	Total	Numerical Methods and HPC
Yves Tourbier	Renault	Expert in crash modelling
Gabriel Staffelbach	CERFACS	Combustion and HPC

Weather, Climate and solid Earth Sciences

Name	Institution	Area of Expertise
Giovanni Aloisio	Univ Salento - CCMC	Exascale Computing
Jean-Claude André	JCA Consultance	Atmosphere
Mario Acosta	BSC	
Peter Bauer	ECMWF	Exascale Computing
Reinhard Rudich	MPG	Meteorology
Rosa Filgueira	BGS	Atmosphere
Sandro Fiore	Univ Salento - CCMC	Earth Sci.
Dimitri Komatitsch	CNRS/LMA	Exascale Computing

Nadia Pinardi	Univ. Bologna	Earth Sci.
Nikolai Shapiro	IPGP	Ocean
Sophie Valcke	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	Torino 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	Multi-scale 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

D3.2**EXDCI inputs to the PRACE Scientific Case**

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