Ref. Ares(2020)1228323 - 27/02/2020

H2020-FETHPC-3-2017 - Exascale HPC ecosystem development

EXDCI-2

European eXtreme Data and Computing Initiative - 2

Grant Agreement Number: 800957

D3.1
Roadmap of HPC applications and usages

Final

Version: 1.0
Author(s): Catherine Lambert (CERFACS), Jean Claude André (JCA Consultance), Stéphane Requena (GENCI), and Mathis Bode (RWTH Aachen University)
Date: 26/02/2020

Project and Deliverable Information Sheet

EXDCI Project	**Project Ref. №: FETHPC-800957**	
	Project Title: European eXtreme Data and Computing Initiative - 2	
	Project Web Site: http://www.exdci.eu	
	Deliverable ID: < **D3.1**>	
	Deliverable Nature: <Report>	
	Dissemination Level: PU	**Contractual Date of Delivery:** 28 / 02 / 2020
		Actual Date of Delivery: 26 / 02 / 2020
	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: Roadmap of HPC applications and usages**	
	ID: D3.2	
	Version: <**1.0**>	**Status:** ***Final***
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): D3.1-v1.0.docx	
Authorship	**Written by:**	Catherine Lambert (CERFACS), Jean-Claude Andre (JCA Consultance), Stéphane Requena (GENCI), and Mathis Bode (RWTH Aachen University)
	Contributors:	Andre Folch (BSC, ChEESE) Josep de la Puente (BSC, ChEESE), Jorge Macías (UMA, ChEESE), Emily Lumley (UCL, CompBioMed), Rossen Apostolov (KTH, BioExcel), Sacha Brun (CEA), Corentin Lapeyre (CERFACS, EXCELLERAT), Thierry Poinsot (IMFT/CERFACS, EXCELLERAT), and Gabriel Staffelbach (CERFACS, EXCELLERAT)
	Reviewed by:	Maike Gilliot (ETP4HPC) and Pascale Rossé-Laurent (ATOS)
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	10/02/2020	Draft	Before internal review
0.2	21/02/2020	Draft	After internal reviews
1.0	26/02/2020	Final version	Approved by TB/MB

Document Keywords

Keywords:	PRACE, ETP4HPC, Research Infrastructure, applications, societal challenges, convergence HPC/AI, artificial intelligence, data analytics, workflows, urgent computing, large scale instruments analysis, edge computing

Table of Contents

1 Introduction 2
2 Recent changes in hardware and methods 3
2.1 Hardware evolution and trends 3
2.2 Machine learning 5
2.3 Revisiting applications and industrial roadmaps: an ever-continuing exercise 6
2.4 How the present deliverable was prepared 6
3 Toward the convergence of HPC/HPDA and AI riding a new digital continuum 8
3.1 Exascale: the convergence of HPC/HPDA and AI 8
3.1.1 Why HPC needs AI? 11
3.1.2 Why AI needs HPC? 13
3.2 Towards a digital continuum from the edge to the tape 17
3.2.1 AQMO: an edge-to-HPC digital continuum framework for air quality 19
3.2.2 PHIDIAS: an edge-to-HPC digital continuum framework for earth system sciences 20
3.3 Leading to new services like urgent computing 21
4 Some updates of application roadmaps 24
4.1 Application example for HPC/AI convergence 24
4.2 Weather and climate applications 25
4.2.1 Weather prediction 26
4.2.2 Climate modeling 27
4.3 High-energy particle physics, astrophysics, and plasma physics 28
4.3.1 High-energy particle physics applications 28
4.3.2 Astrophysics 29
4.3.3 Plasma physics 30
4.3.4 Some recommendations 31
4.4 Biosciences 31
4.4.1 Biomolecular sciences 31
4.4.2 At the larger scale, a growing number of issues concerns human health 35
4.5 Combustion 39
4.6 Chemistry and material sciences 42
4.7 Social Sciences 43
4.8 Engineering and industrial applications 47
5 Applications and co-design 52
6 The main messages 54

List of Figures

Figure 1 Increase in computing power as a function of time of the ten most powerful supercomputers* 3
Figure 2 Mean electrical consumption as a function of time of the ten most powerful supercomputers. Same remark as in Figure 1 concerning November 2017.......... 4
Figure 3 Number of cores, including accelerators, as a function of time of the ten most powerful supercomputers. Same remark as in Figure 1 concerning November 2017. 5
Figure 4 Overview of CMIP5/6 data management (courtesy of CMIP) 9
Figure 5 SKA project overview (courtesy of SKA).......... 9
Figure 6 Picture of GENCI's Jean Zay converged system at IDRIS (courtesy of C. Frésillon) 11
Figure 7 Picture of Google v3 TPU Pod and card with direct liquid cooling (courtesy of Google)..... 11
Figure 8 Inria's MELISSA post-processing framework (courtesy of B. Raffin).......... 12
Figure 9 Two distinct eras of compute usage in training AI systems (source [11]).......... 14
Figure 10 ResNET-50 pipeline (courtesy of M. Riedel) 15
Figure 11 Optimization of ResNET 50 done at JSC. ResNET is a leading convolutional neural network that is trained on more than a million images from the ImageNet database. The network is 50 layers deep and can classify images into 1000 object categories, such as keyboard, mouse, pencil, and many animals. Its computational complexity of training relies on the fact that it has around 25.6 millions of trainable parameters which make RESNET-50 very suitable for parallelization via distributed training on multiple GPUs (courtesy of M. Riedel). 15
Figure 12 The PlantNet system (courtesy of A. Joly) 16
Figure 13 Scalability of the training of the 390,000 classes PlantNet model on x86 systems (courtesy of A. Joly).......... 16
Figure 14 A new digital continuum: scheme (source BDEC2). 17
Figure 15 A new digital continuum: table (source BDEC2). 18
Figure 16 AQMO continuum overview (courtesy of F. Bodin) 19
Figure 17 Synthetic view and components of the main workflow of the PHIDIAS project (courtesy of B. Dintrans) 21
Figure 18 Example of a complete end-to-end workflow of real time tsunami alert system (courtesy of A. Floch) 23
Figure 19 DL network used for subgrid model development. Data generation and training were performed on supercomputers (courtesy of M. Bode). 25
Figure 20 Progress in the degrees of freedom of the European Centre for Medium-range Weather Forecasts (ECMWF) operational global atmosphere model in comparison to Moore's law. The ambitious goal of reaching an operational 1 km horizontal resolution with 180 levels in 2030 requires a faster progress. The numbers indicate the average grid-point distance in kilometers and the corresponding spectral resolution and levels used (source [15]) 26
Figure 21 Test simulation of climate with high-resolution coupled climate model (courtesy of M. Castrillo).......... 28
Figure 22 Space-time view of the re-ionization process. Vertical axis shows a region of space extending over 12 Mpc ("Megaparsec"), while the horizontal axis covers the first 109 years of the universe. Hor regions are shown in red, dark regions correspond to ionized zones, and the variation in green color refers to the gas density (courtesy of D. Aubert). 30
Figure 23 Coupling of the HPC/AI re-ionization code using auto-encoders for the radiative transport module (courtesy of D. Aubert).......... 30
Figure 24 Ensemble simulations in GROMACS provide close to 100-fold reduction in time-to-solution for opening 20 base pairs on Piz Daint CPUs + P100 GPUs at Centro Svizzero di Calcolo Scientifico, Switzerland (courtesy of R. Apostolov) 33
Figure 25 High-throughput hybrid-QM/MM modeling of biomolecular systems. A related, previous, implementation using the same distributed computing model achieved linear scaling until 40,000 cores on an older Cray CX40. Much higher scalability is expected with the new codebases and redesigned implementation (courtesy of R. Apostolov). 34
Figure 26 Atomistic model of the chromatophore of the purple bacteria (source [17]).......... 34

Figure 27 Intervening scales and organizational levels in computational biomedicine, from health care at the top to cell and molecular scales at the bottom (source [18]) 35
Figure 28 Fully-coupled fluid-electro-mechanical simulation of a heart. The sequence shows the systole process of a third degree atrio-ventricular block and the action of a trans-catheter intra-ventricular pacemaker (the small red cylinder close to the heart apex). Tissue is colored by electrical activity and blood shows the so-called Q-criterion, which depicts blood flow vortices evolution (courtesy of A. Santiago).. 37
Figure 29 Overview of a combined MD/QM simulation and ML approach From left to right (source [19]): (a) MD data generation will produce output files size of several petabytes; (b) ML may then replace QM to predict dihedral energies given a neural network trained with QM simulations; (c) data augmentation by MD will then improve the description of some processes (protein-ligand binding poses for a set of protein-ligand pairs with unknown binding mode; binding affinity data for a set of resolved protein-ligand complex structures of unknown affinities). 38
Figure 30 Overview of brain cytoarchitecture simulation/learning (courtesy of C. Poupon)............... 39
Figure 31 Example of coupled computation on HPC architectures: two codes are coupled to compute the combustion chamber on one side and the turbine on the other side (source CERFACS)............... 41
Figure 32 Coupling of AVBP and a neural network for computing the total surface of the flame front (source CERFACS).. 41
Figure 33 Yearly edition for 2019 of the One Internet Minute (source [21])..................................... 44
Figure 34 Global framework of scalable social sciences framework for pandemic spreading (courtesy of Barcelona Supercomputing Center) 46
Figure 35 Base pillars of the CoE EXCELLERAT (source EXCELLERAT) 48
Figure 36 Seismic processing of TOTAL (courtesy of H. Calandra) as published inside previous EXDCI document 49
Figure 37 Early use of quantum computing at TOTAL (courtesy of H. Calandra) 50
Figure 38 1 billion cells multiphasic industrial scale fluidized bed reactor CFD simulation (courtesy of H. Néau) 51
Figure 39 First quantum experiments of EDF R&D UK for optimal battery operations (courtesy of S. Tanguy) 51

References and Applicable Documents

[1] https://exdci.eu/sites/default/files/public/files/d3.2.pdf
[2] http://www.exdci.eu
[3] http://www.prace-project.eu
[4] http://www.prace-ri.eu/third-scientific-case
[5] http://www.etp4hpc.eu
[6] https://www.etp4hpc.eu/pujades/files/SRA%203.pdf
[7] https://exdci.eu/resources/sra
[8] https://www.eurolab4hpc.eu/media/public/vision/D2-2--final-vision.pdf
[9] https://arxiv.org/abs/1702.07675
[10] https://github.com/melissa-sa/melissa/wiki/1-Home
[11] https://openai.com/blog/ai-and-compute/
[12] https://www.hpcwire.com/2018/08/14/cern-incorporates-ai-into-physics-based-simulations/
[13] https://arxiv.org/abs/1910.00928

[14] https://arxiv.org/abs/1911.11380

[15] T.C. Schulthess et al., Reflecting on the goal and baseline for Exascale computing: A roadmap based on weather and climate simulations. Computing in Science and Engineering, 30-41, 2019

[16] J. Kates-Harbeck et al., Predicting disruptive instabilities in controlled fusion plasmas through deep learning. Nature, 568, 526-543, 2019

[17] https://www.sciencesetavenir.fr/fondamental/biologie-cellulaire/136-millions-d-atomes-pour-simuler-la-photosynthese_140381

[18] www.compbiomed.eu/wp-content/uploads/2019/11/D2.4-Report-on-the-Impact-of-Modeling-and-Simulation-within-Biomedical-Research-as-enabled-by-CompBioMed_V1.0_reduce.pdf

[19] A. Pérez et al., Simulations meet Machine Learning in Structural Biology. Current Opinion in Structural Biology, 49:139-144, 2018

[20] https://scikit-learn.org

[21] https://www.garyfox.co/what-happens-internet-minute-2019

[22] https://www.bsc.es/discover-bsc/organisation/scientific-structure/hpc-modelling-and-simulation-societal-challenges

List of Acronyms and Abbreviations

AI	Artificial Intelligence
AMD	Age-related Macular Degeneration
API	Application Programming Interface
BDEC	Big Data and Extreme-scale Computing
CFD	Computational Fluids Dynamics
CINES	Centre Informatique National de l'Enseignement Supérieur
CMIP	Coupled Model Intercomparison Project
CoE	Centers of Excellence
CPU	Central Processing Unit (=processor)
CSA	Coordination and Support Action
D	Deliverable
DFT	Density Functional Theory
DIAS	Copernicus Data and Information Access Services
DL	Deep Learning
DoE	Department of Energy
DSL	Domain Specific Language
EC	European Commission
ECMWF	European Centre for Medium-range Weather Forecasts
ECP	Exascale computing project
EOSC	European Open Science Cloud
ERC	European Research Council
ERCC	Emergency Response Coordination Center
ETP4HPC	European Technology Platform for High Performance Computing

EXDCI	European Extreme Data & Computing Initiative
EU	European Union
EC	European Commission
FAIR	Findable Accessible Interoperable and Reusable (data standard)
FET	Future and Emerging Technologies
FFT	Fast Fourier Transform
FTRT	Faster-Than-Real-Time
GAN	Generic Adversarial Network
GDP	Growth Domestic Product
GENCI	Grand équipement national de calcul intensif
GPGPUs	General Purpose computing on GPUs
GPU	Graphical Processing Units (hardware accelerator)
H2020	Horizon 2020 – The EC Research and Innovation Programme in Europe
HBP	Human Brain Project EU flagship
HPC	High Performance Computing
HPDA	High Performance Data Analytics
HTC	High Throughput Computing
IaaS	Infrastructure as a Service
IDC	International Data Corporation
ILP	Instruction Level Parallelism
INCITE	Innovative and Novel Computational Impact on Theory and Experiment
INRA	Institut national de la recherche agronomique
I/O	Input/Output
IPCC	Intergovernmental Panel on Climate Change
IT	Information Technology
ITER	International Thermonuclear Experimental Reactor
IoT	Internet of Things
JU	Joint Undertaking
LES	Large-Eddy Simulation
LHC	Large Hadron Collider
LQCD	Lattice QCD
LUMI	Large Unified Modern Infrastructure
MD	Molecular Dynamics
MELISSA	Modular External Library for In-Situ Statistical Analysis
MHD	Magneto-HydroDynamics
ML	Machine Learning
MM	Molecular Mechanics
Mpc	Megaparsec
MRI	Magnetic Resonance Imaging
NVMe	NVM express
PRACE	Partnership for Advanced Computing in Europe

QCD	Quantum ChromoDynamics
QLM	Quantum Learning Machine (simulator from ATOS)
QM	Quantum Mechanics
R&D	Research and Development
RDMA	Remote Direct Memory Access
RNN	Recurrent Neural Network
RoCE	RDMA over Converged Ethernet
SLA	Service Level Agreement
SKA	Square Kilometer Array (radio telescope project)
SME	Small and Medium Enterprise
SRA	Strategic Research Agenda
TGCC	Très Grand Centre de calcul
TPU	Tensor Processing Unit
US	United States
WAN	Wide Area Network
XAI	Explainable AI
XML	Extensible Markup Language

Executive Summary

A number of new developments have changed the high performance computing (HPC) landscape over the last five years or even less.

On the hardware side, the new computing architectures are increasingly relying on accelerators, while the more traditional CPUs represents only 25% of the processors available in the actual top supercomputers now. While this evolution allows for the necessary limitation of the energy consumption of the machines, corresponding adaptation of legacy codes, both in the scientific and industrial contexts, is one of the main challenges facing them. Another important architecture evolution should also be emphasized: HPC centers are less and less isolated but fully integrated inside a global cyber-infrastructure ensuring a digital continuum of the data, from the place where they are generated (large-scale instruments, Internet of Things, IoT) to the place where they are finally stored or archived after having been processed.

On the methodological side, artificial intelligence (AI), or more precisely machine learning (ML) and deep learning (DL), is modifying the way scientific problems are now being attacked. On the one hand ML is contributing to more traditional HPC by offering methods for either performing more efficiently some parts of the calculations (parameterizations schemes, solver preconditioners …), or identifying features which would otherwise be very difficult to identify, or providing innovative methods to process the wealth of date produced by these calculations. On the other hand, this physics-based HPC production of important datasets contributes to ML efficiency, while physics-based HPC it is also able to constrain neural network approaches.

This convergence between HPC and AI/ML is a further incentive for the actual development of the so-called converged infrastructures.

Besides the longstanding need for more computing power, e.g. access to Exascale resources, the new game-changer is leading to some modification for the roadmaps in a number of application domains, and roadmaps have been consequently modified and updated to account for these new opportunities, especially as compared to the preceding exercise conducted under the earlier EXDCI project (see "EXDCI inputs to the PRACE Scientific Case" [1]). This is the case, among others, for weather and climate, for high-energy particle physics, astrophysics and plasma physics, for biosciences and neurosciences, from the molecular level up to the medicine level, for combustion, for material sciences, for social sciences, and for engineering and industrial applications. For each of this domains, new opportunities resulting from the new hybrid approaches are given and discussed. The development of urgent computing is also benefiting from converged HPC/AI approaches, and applications are now prepared for real-time decision making. These concerns are, among others, natural hazards (earthquake, tsunamis, …), biological hazards (propagation of pandemic, …), industrial damages and accidents, and (cyber) terrorism.

Co-design between vendors and application developers in Europe is taking place at a rather slow pace, at least as compared to what happens in other countries, e.g. Japan.

Finally, a series of recommendations are proposed: 1) increase the support for the development and use of hybrid methods (modeling, resource infrastructures, initial and life-long training); 2) keep-on inserting HPC facilities in a global cyber-infrastructure; and 3) sustain and increase co-design with application developers, among others to facilitate the transition to converged hardware and software infrastructures.

1 Introduction

In mid-2018, the former EXDCI(-1) ("European Extreme Data & Computing Initiative") project [2] published, as its deliverable 3.2, a quite complete overview of high performance computing (HPC) roadmaps of both scientific and industrial applications entitled "EXDCI inputs to the PRACE Scientific Case" [1]. This document included two main parts, the first one dealing with an analysis of scientific challenges in various scientific disciplines and industrial applications, the second one concerned with a set of quite high-level recommendations. The coverage of scientific disciplines was quite extensive, including numerical weather forecasting and meteorology, climate, oceanography, solid earth sciences, nuclear physics and QCD ("Quantum ChromoDynamics") plasma physics, fusion, astrophysics and cosmology, material sciences, data-driven bioscience, molecular simulation, and biomedical simulation. It should nevertheless be noted here that social sciences were not at all addressed. The range of industrial and engineering applications was also large: aeronautics and aerospace, automotive industry, oil and gas, power generation and nuclear plants, process engineering, and combustion.

A set of three main global recommendations were proposed:

- Convergence between in-situ/in-transit post-processing techniques and machine learning (ML)/deep learning (DL) methods;
- Development of new services toward urgent computing and link with scientific instruments;
- And development of new Centers of Excellence (CoE) in Europe (Engineering and industrial applications, (Open-source) software sustainability, high performance data analytics, HPDA).

The purpose of the present deliverable is to update this earlier analysis. For this, some new trends, which extend or even partially challenge the analysis from 1.5 years ago, are mentioned in Chapter 2. The new topics of convergence of HPC/HPDA and artificial intelligence (AI) as well as the digital continuum, which were not discussed in the earlier version are summarized in Chapter 3. The application cases are updated in Chapter 4 and co-design is addressed in Chapter 5.

2 Recent changes in hardware and methods

Before considering the evolution of scientific and industrial roadmaps, it is necessary to recap the two main changes, which influenced the HPC approach, either in the hardware domain or with the disruptive development of ML in a number of applications.

2.1 Hardware evolution and trends

It is now well known that the increase in computing power cannot come from an increase of the core frequency above a few GHz, as this would lead to an unaffordable energy consumption. In fact, this has been experienced since about 2006, clock frequencies are not increasing much (if at all not going down). The consequence is that speed ups require more parallelism at all level: there is indeed such a possibility to gain performance by reducing the transistors size to include on the same surface of silicium many computing cores and then by using multi-level parallelism, from the many-core processors to the many-sockets nodes and to many-nodes supercomputers. But the number of transistors in a dense integrated circuit, which has been doubling roughly every two years for a long time, needs now slightly longer than three years for doubling, while the physics of making transistors suggests reaching the limit is near. Electrical consumption, although significantly reduced, is still an issue. Let us recall that, approximatively 12 years ago, the 1 MW consumption has been reached for 100 TeraFlop/s computers, as can be seen in Figure 1 and Figure 2.

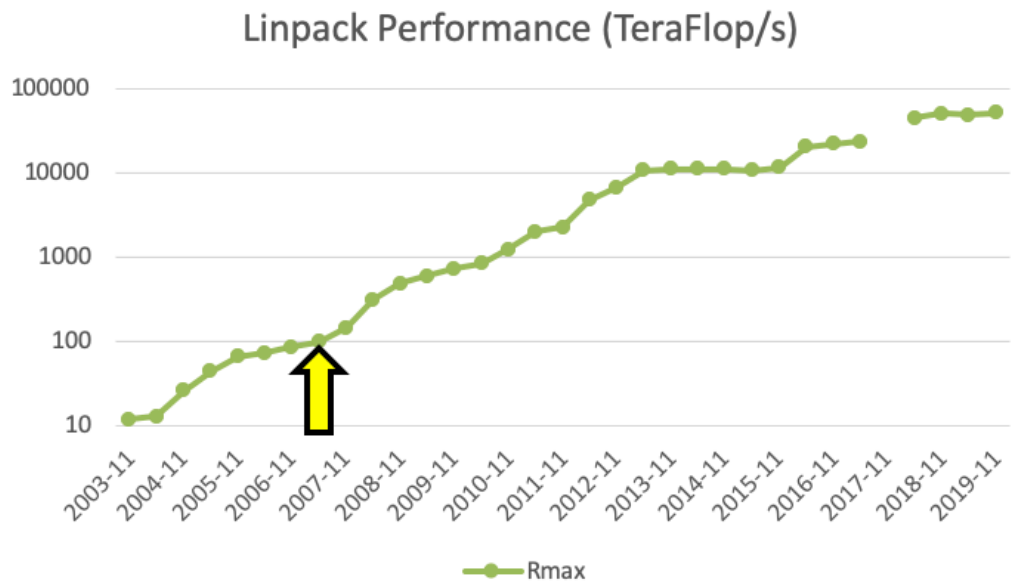

Figure 1 Increase in computing power as a function of time of the ten most powerful supercomputers*

*Data for November 2017 do not appear as, for this only date, the Japanese “Gyoukou” machine, then ranked #4 with 19.136 PetaFlop/s, had a reduced electrical consumption of 1.35 MW, due to its 19,840,000 accelerators out of a total of 19,860,000 cores. Including it here would have led to an incorrect appreciation of the global trends.

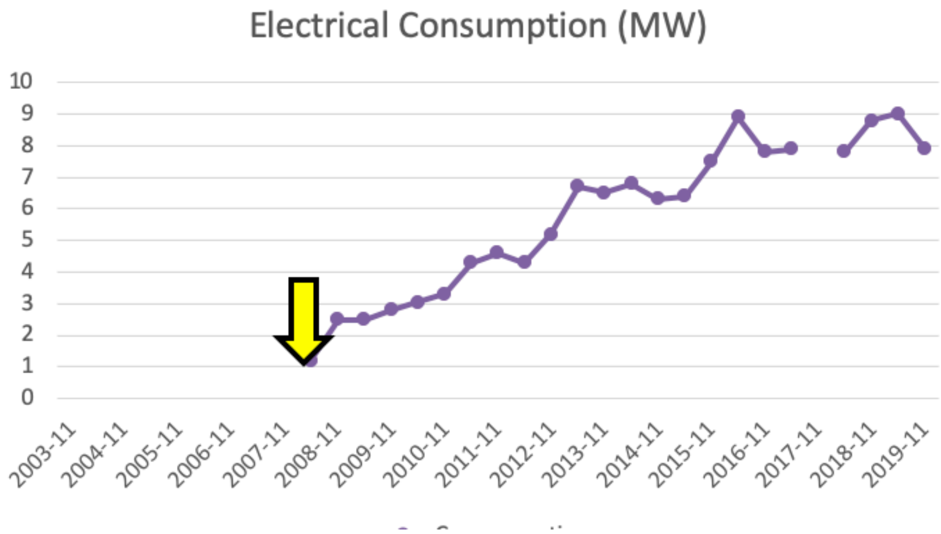

Figure 2 Mean electrical consumption as a function of time of the ten most powerful supercomputers. Same remark as in Figure 1 concerning November 2017.

A way to reduce the increase in electrical consumption is to use specialized computing cores, requiring less energy and then serving as accelerators to the more traditional cores. Such GPU accelerators are now used for an ever-growing number of applications, although they require to develop new software tools.

Altogether, and for the most recent period, i.e. the last five years, the computing power of the mean of ten best-supercomputers-at-a-time has been multiplied by a factor close to five, from 11 PetaFlop/s to slightly more than 50 PetaFlop/s, linked to an increase in the number of cores by a factor slightly less than three, from 800,000 to almost 2,200,000. During the same period, the number of accelerators, which started to be part of supercomputers around 2011, increased by a factor of almost five, and they now account for about 75% of the total number of cores (see Figure 3), allowing for an almost levelling of the electrical consumption.

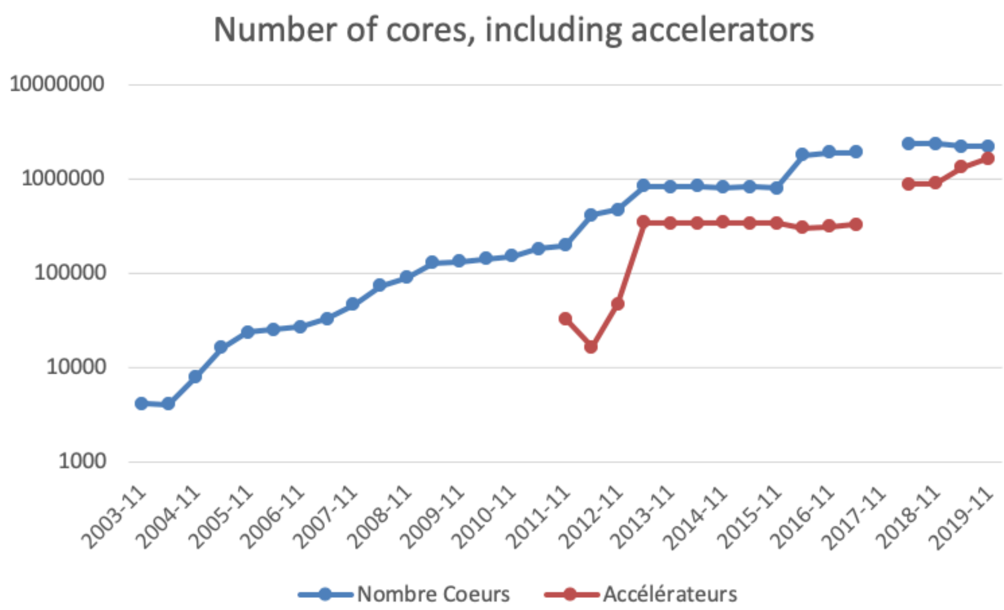

Figure 3 Number of cores, including accelerators, as a function of time of the ten most powerful supercomputers. Same remark as in Figure 1 concerning November 2017.

Straightforward extrapolation of these trends would call for computers with much more parallelism and customized hardware, but being significantly slower than earlier predicted from the so-called Moore's law. The importance of power consumption, responsible for these trends, will remain in the forthcoming years as a crucial and unavoidable constraint. Future computers will have to provide 1 ExaFlop/s with 20 to, at most, 30 MW power consumption, therefore calling for a massive use of low-consumption processors. With present technologies this implies a very high proportion of accelerators, and improved interconnects and software.

On the longer term, new technologies presently under development should provide alternative and more efficient ways to increase the computing power without increasing the energy consumption: die stacking and 3D chip, nonvolatile memory, photonics, resistive computing, neuromorphic computing, quantum computing, nanotubes, graphene, and diamond-based transistors. It is however too early to predict which technologies will mature first, and when. As a consequence, roadmaps for applications should take as an input the continuation of actual hardware technologies for the next few years.

2.2 Machine learning

ML, and especially DL, first appeared as domain more or less independent from model-based HPC, both for the software stack and type of computers being used. It has been however obvious for more than a couple of years now that ML/DL is indeed synergetic with HPC. Most of the applications are indeed able to benefit from the complementary approach, with strong influence on roadmaps in a number of application domains. Some applications are given in more detail in Chapter 4.

2.3 Revisiting applications and industrial roadmaps: an ever-continuing exercise

Both the trends in hardware and architectures over the next- or longer-term futures, and the modification of HPC paradigm due to the ever-growing importance of dealing with massive data and taking advantage of ML, make it necessary and timely to review the earlier roadmaps for a number of scientific applications and industrial and engineering usages. This exercise should however recognize and take advantage of a few recent analysis conducted by other organizations.

In the field of applications, the new vision document by PRACE ("Partnership for Advanced Computing in Europe") [3], "The scientific case for computing in Europe 2018-2026", published in October 2018 [4], reviews the scientific and technical challenges and potential breakthroughs which are facing the various applications domains: climate, weather and earth sciences; life science and improvement of human health; energy; infrastructure and manufacturing for mankind; future materials, from molecules to machines. It also underlines the need for adapting the approaches to take care of complexity and massive data as well as of architectures of next-generation computers. It concludes with recommendations for infrastructures, both from the compute and storage points of view and from operations and environments. The BDEC ("Big Data and Extreme-scale Computing") paper, "Big data and extreme-scale computing: Pathways to convergence-toward a shaping strategy for a future software and data ecosystem for scientific inquiry", also published in 2018, is a thorough analysis of the scientific potential of hybriding the traditional HPC approach with HPDA and ML, as well as the constraints and difficulties on the road to converged approaches. It makes a number of recommendations for both decentralized edge and peripherical ecosystems on the one hand and centralized facilities on the other hand.

For issues closer to hardware and architectures we have to consider three recent studies of importance. Firstly, the ETP4HPC ("European Technology Platform for High Performance Computing") [5] "Strategic Research Agenda" (SRA) is continuously revised, its third version (SRA-3) has been published at the end of 2017 [6], and the fourth version (SRA-4) should be available in the end of 2020 [7]. SRA is aimed mostly at providing contextual guidance for research in hardware and software, for business and for EU. It concentrates on technological aspects and keeps an updated view of actual context and vision for the future. Exascale developments are central to its work. Secondly, the document "Long-Term Vision on High-Performance Computing" [8] published by the European Commission (EC) coordination and support action (CSA) "Eurolab-4-HPC" in 2018 outlines a long-term vision for excellence in European HPC research, with a timescale beyond Exascale computers, i.e. a time span of approximately 2023-2030. In this document issues concerned with disruptive technologies are dealt with, as well as their consequences with respect to software and programming environments. Finally, EuroHPC strategy documents are in the process of being finalized and published. They should emphasize the new avenues for European HPC.

2.4 How the present deliverable was prepared

It has already been said that a number of recent prospective documents, offering complementary visions for hardware, software, methods, applications, … were available. If relevant for the present deliverable, visions and recommendations from these documents have been taken into account. Other sources of information have been considered as well.

Firstly, most of the CoEs were contacted and, when available, the recommendations and conclusions taken from their prospective documents have been included. As some of these documents will be available only in 2022, when their actual contracts with the EC will end, it was decided to discuss directly with the CoEs to learn about their roadmaps. For scientific domains which are not covered by CoEs it was decided to contact experts in such fields, to collect their vision documents and/or invite them to produce ad hoc considerations for contributing to the present deliverable. This was in particular the case for social sciences and astrophysics nuclear fusion.

3 Toward the convergence of HPC/HPDA and AI riding a new digital continuum

In the following part of the document, one will use either AI or ML/DL, as AI is most commonly used in the literature at stake, but as it covers only ML/DL for almost all of the developments referred to below.

As already said, the EXDCI project issued its last roadmap document on academia and industrial applications mid-2018, as such elements are not varying so much in 18 months the following elements are much more new elements or a gap analysis.

3.1 Exascale: the convergence of HPC/HPDA and AI

HPC, the use of supercomputers for accelerating numerical simulation, is a **strategic tool for a competitive science, fostering innovation and supporting public decision making**. In science, after having being used since more than 30 years in climate research, numerical weather prediction, astrophysics or chemistry, HPC is now irrigating all scientific fields from biology, life sciences and health, high fidelity combustion, materials sciences to social sciences and humanities. In industry, HPC is widely used in oil and gas exploration, aeronautics, automotive and finance, and becomes now crucial for ensuring personalized medicine, developing nanotechnologies or enabling developing renewable energies. Finally, HPC is becoming a tool of growing importance for supporting public decision making by allocating urgent computing resources in case of natural risk events (earthquakes, thunderstorms, flooding, …), industrial risks, biological risks or (cyber)terrorism risks.

While Moore’s and Dennard’s laws and the shift from vector to scalar parallel processing allowed relatively smooth transitions every 10 to 12 years from Gigascale[1] (1985), Terascale[2] (1997) and finally to Petascale[3] (2008), the upcoming transition from Petascale to Exascale[4] (expected between 2019 and 2023) is going to raise many issues.

The first one is related, as already indicated, to the **limitation of the power consumption** to an acceptable 20 to 30MW for 1 ExaFlop/s system leading to dense massively parallel hardware architectures including heterogeneous manycore compute nodes, and complex memory and storage hierarchies. Benefiting from these architectures will require profound changes/rewriting of HPC applications in order to adapt them to new (and standard) programming models (including domain specific language, DSL, and interpreted languages), extract more parallelism, use dynamically available computing/data resources, maximize data locality and reuse for avoiding data movement and thus save energy, deal with hardware and software resiliency, … As a consequence, instead of focusing like in the past on peak performance of HPC systems, HPC agencies and centers are now expecting future systems to improve drastically by 50x to 100x sustained performance on real applications, considering co-design between scientific communities, HPC vendors and HPC centers as the only way to develop simultaneously efficient HPC architectures and applications.

The second issue is related to the **convergence between HPC and big data workloads** due to the deluge of data coming from next generation scientific instruments (satellites, (radio-) telescopes, accelerators, microscopes, sequencers, Internet of Things, IoT, …) and from large-scale simulations (massive 3D simulations, multi-scale and multi-physics coupled simulations, ensemble/optimization studies, uncertainties quantification, …). Exploiting and valuing such

[1] one GigaFlop/s = 1 billion floating point operations per second
[2] one Tera/Flop/s = 1 thousand billion floating point operations per second
[3] one PetaFlop/s = 1 million billion floating point operations per second
[4] one ExaFlop/s = 1 billion billion floating point operations per second

amount of structured or unstructured data in a reasonable and competitive time is not anymore possible for human beings, leading to the rise of HPDA supported by new data assimilation/interpretation/extraction/prediction techniques benefiting from AI. As example, Figure 4 states the increase of final (refined) data generated by successive CMIP5 (2010-2012, fifth Coupled Model Intercomparison Project) and CMIP6 (2016-2018) climate exercises. For CMIP7 planned to be performed around 2022-2024, a 30x increase in the size of the data, resulting from higher resolutions and number of multi-physics ensembles, is expected by climate communities.

Status CMIP5 data archive:
- **1.8 PB for 59000 data sets stored in 4.3 Mio Files in 23 ESGF data nodes**
- **CMIP5 data is about 50 times CMIP3**

Extrapolation to CMIP6:
- **CMIP6 has a more complex experiment structure than CMIP5.**
- **Expectations: more models, finer spatial resolution and larger ensembles**
- **Factor of 20: 36 PB in 86 Mio Files**
- **Factor of 50: 90 PB in 215 Mio Files**

Figure 4 Overview of CMIP5/6 data management (courtesy of CMIP)

Another example shown in Figure 5 is coming from large-scale instruments where next generation projects like SKA (square kilometer array), a radio telescope project, are going to generate up to 500 TB/s of data at the level of the edge and up to 4 EB of refined data per year.

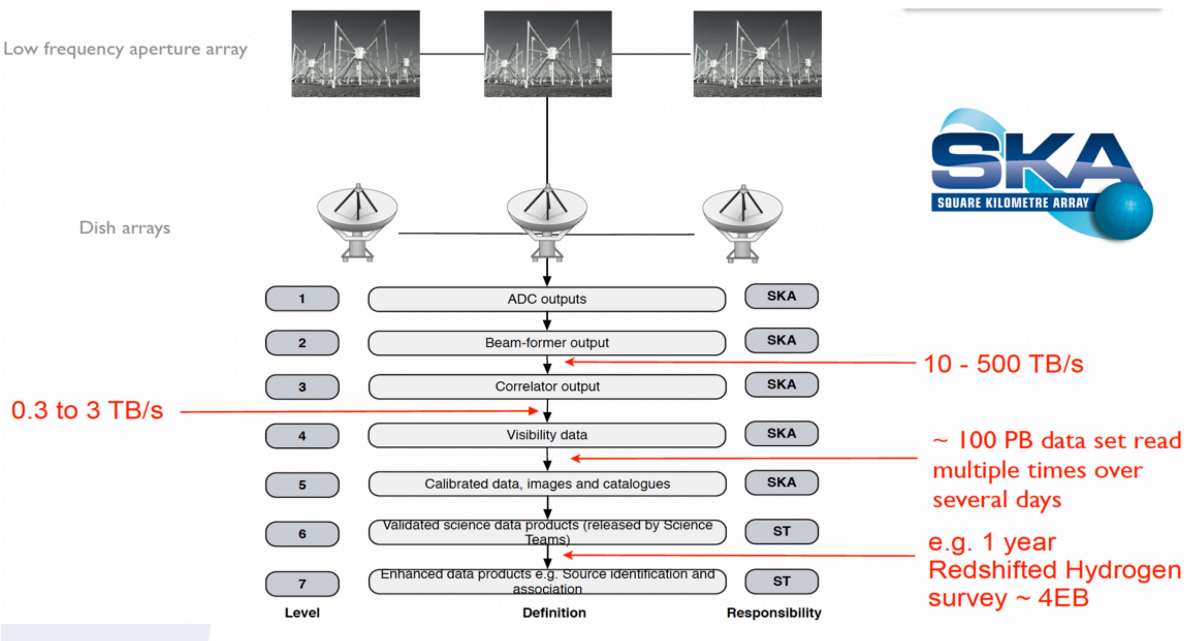

Figure 5 SKA project overview (courtesy of SKA)

By consequence, Exascale will be a new paradigm, not only technological but also organizational in the way to **rethink data workflows from end to end, federate converged HPC/data infrastructures through high bandwidth network services, develop new skills in HPC, data management and AI and foster new insights from scientific and industrial**

disciplines toward data discovery. As a consequence, we see since few years the rise of converged systems able to accommodate the needs of AI-augmented numerical simulation and the scale-out of AI models using HPC.

In the US such systems like IBM Summit at Oak Ridge National Laboratory are providing a peak performance of 200 PetaFlop/s using more than 27,000 NVIDIA V100 GPUs across more than 4600 Power9 based compute nodes. Originally designed to accommodate the needs of traditional massive HPC workloads, Summit is now running HPC, AI workloads at scale, and real time analytics. The next generation of US Department of Energy (DoE) systems, providing Exacale class performance for 2021/22 will be converged systems (HPC+AI) deployed at Argonne National Laboratory (A21, Intel/Cray system of 1.3 ExaFlop/s peak end of 2021), Oak Ridge National Laboratory (Frontier Cray/AMD system of 1.5 ExaFlop/s peak end of 2021) and Lawrence Livermore National Laboratory (El Capitan Cray system of 1.5 ExaFlop/s peak, end of 2022).

In Japan the Post-K computer (now known also as Fugaku) at RIKEN will also address starting 2021 converged workloads with more than 150,000 high-performance CPUs, ARM based A64FX scalar processors developed by Fujitsu.

In Europe, the three pre-Exascale systems ongoing purchased by the EuroHPC joint undertaking (JU) together with the three hosting entities (Barcelona Supercomputing Center for MareNostrum 5, CINECA for Leonardo, and CSC for the Large Unified Modern Infrastructure, LUMI, consortium) will be based on hybrid nodes (CPU+GPU) addressing potentially converged workloads. These systems with a performance >150 PetaFlop/s peak will be available for European researchers starting 2021.

In France finally, following the announcement of the President mid-2018 about the French AI strategy (called "AI For Humanity"), Grand équipement national de calcul intensif (GENCI) have been asked to make available for French AI research community (academia and industry) leading edge HPC resources. This led beginning of 2019 to the installation of Jean Zay at IDRIS at Centre national de la recherche scientifique of an Hewlett Packard Enterprise converged system of 16 PetaFlop/s based on a mix of scalar (more than 60,000 Intel cores) and converged nodes (scalar nodes accelerated with 4 to 8 NVIDIA V100 32 GB GPU), a tiered fast storage infrastructure and a converged software stack including AI-native schedulers (Kubernettes), the support of secured containers technologies and Jupyter notebooks (see Figure 6).

Beyond the traditional HPC world, it is also interesting to see that the **world of hyperscalers (clouds providers) are also moving to converged architectures** by integrating components from the HPC world in terms of integration (notion of dense pods), cooling (direct liquid cooling on Google's latest tensor processing unit, TPU, see Figure 7), interconnects (adoption of low latency remote direct memory access, RDMA, over converged ethernet, RoCE, deployment of InfiniBand networks at Amazon Web Services) and tiered storage technologies (use of NVM express, NVMe, over fabrics, use of Lustre parallel file system, …).

Beyond the convergence of hardware and middleware architectures, it is also interesting to see growing convergence of the usage between HPC and AI, with the moto *HPC needs AI and AI needs HPC*.

Figure 6 Picture of GENCI's Jean Zay converged system at IDRIS (courtesy of C. Frésillon)

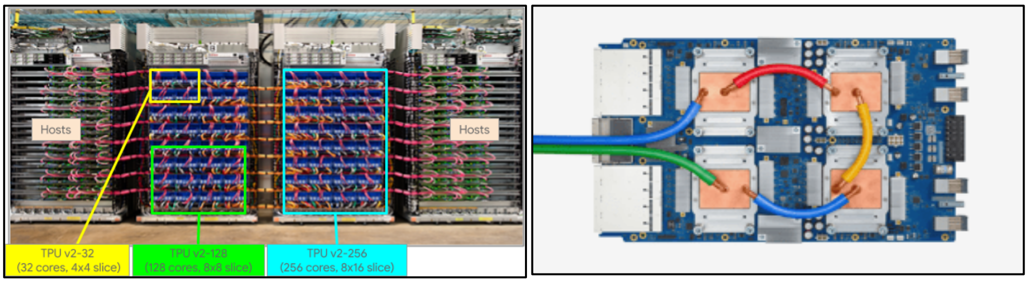

Figure 7 Picture of Google v3 TPU Pod and card with direct liquid cooling (courtesy of Google)

3.1.1 *Why HPC needs AI?*

The first reason is that since the 1990s, researchers use HPC/high throughput computing (HTC) to post-process the data coming from large scale instruments (satellites, telescopes, microscopes, network of sensors, sequencers, …) and the amount of data is now so huge that traditional HPDA techniques and **humans in the loop are not enough to process and value data in a competitive time**. The amount of data that will be generated by various sources will reach 175 ZB (175×10^{21} Bytes) by 2025, growing from 33 ZB in 2018 according to Hyperion (formerly International Data Corporation, IDC). AI techniques such as DL gained recently a lot of at attention in the field of medical imaging (breast cancer detection, detection of diabetic retinopathy, age-related macular degeneration, AMD, and glaucoma, …), large-scale instruments (for denoising and detecting gravitational wave [9], AI-guided post-processing of the next CERN's large hadron collider (LHC) experiments, …) and natural risk prevention (for example using AI, such as recurrent neural networks, RNNs, for listening the earth background seismic noise in order to anticipate the rise of earthquakes). This implies, for example, to adapt HPC operations for accommodating to stream based access of data coming from instruments with traditional batch access in order to avoid when possible to store raw data and post-process it on the fly and the support of end to end workflows (see specific section later).

The second reason is the rise of **AI-augmented numerical simulation**. Many experts are not considering that AI will replace numerical simulation but much more will accelerate it by coupling learnt models with existing simulation codes for multi-scale/multi-physics simulations, being able to interpolate/extrapolate missing information in molecular dynamics

(MD) simulations or as input for numerical models, accelerating the convergence of iterative methods by providing refined inputs, being integrated into optimization and uncertainties quantification studies or in the management of meshes (pre-processing, re-meshing, …).

The third reason is related to the post-processing of data coming from massive simulations running on supercomputers. In order to save time and energy, it is preferable to take benefit of data locality for performing **AI-guide in-situ or in-transit post-processing of the data and only store pertinent inferred data on disk**. A lot of research is done in Europe, as an example the MELISSA ("Modular External Library for In-Situ Statistical Analysis") framework [10] (see Figure 8) from Inria aims to propose standard and portable AI-guided in-situ post-processing for massive ensemble statistical analysis with concrete results on energy applications with EDF (France). Such frameworks will also evolve in the future to AI-driven computational steering of the simulations, allowing numerical simulations to converge faster end thus optimize time and energy-to-solution.

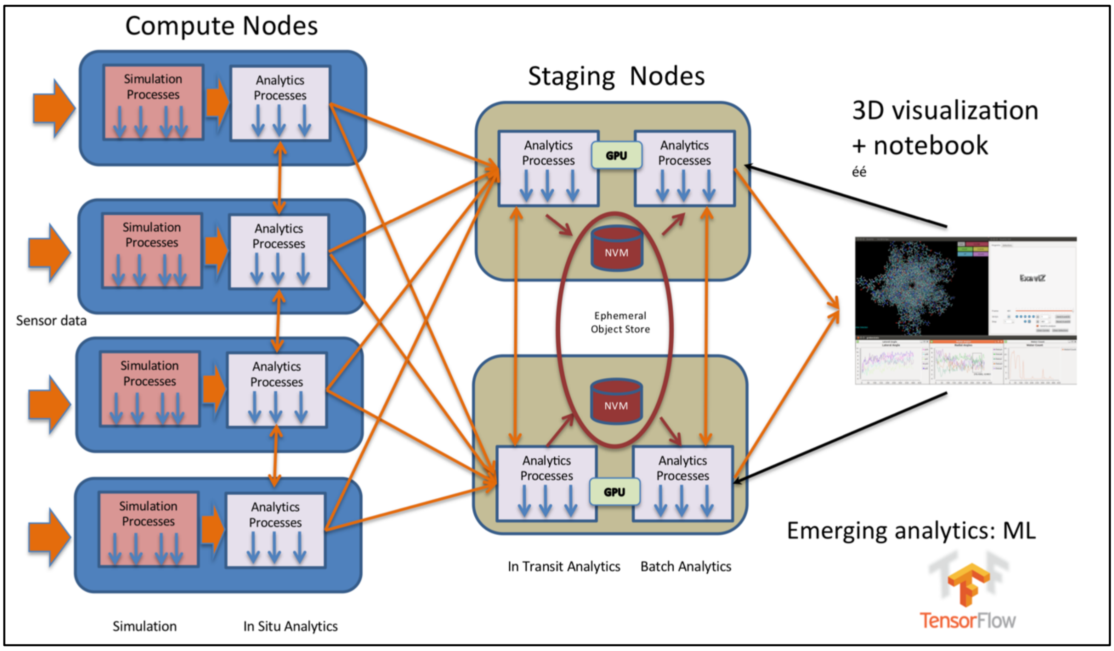

Figure 8 Inria's MELISSA post-processing framework (courtesy of B. Raffin)

Finally, and for sure not the least, the last reason is to use AI for **better exploiting HPC systems and more globally computing centers**. According to Hyperion, by 2020, the demands of next-generation applications and new IT architectures will force 55% of enterprises to either update existing data centers or deploy new ones. As datacenters architectures are becoming more complex and denser, there is also a major stake in a context of limited human workforces in optimizing the global production of heterogenous workloads while reducing energy footprint. In that context the use of AI for assisting system administrators is growing from data center infrastructure management (cooling, power), IT environment (storage, networks, security) to HPC systems.

As a major example for HPC systems, **AI could be used for augmenting the capabilities of the resource scheduler for**:

- performing smart preventive self-healing of the components,

- identifying automatically the running applications and optimize their placement knowing their profile (in terms of communication patterns, input/output (I/O) operations and of course computing requirements associated to the datasets used allowing to adjust the frequency of the processors),
- being able to application-based checkpoint/restart on the fly for having a more elastic production and thus accommodate to stream or urgent computing,
- deferring the I/O traffic when congestion periods are occurring, …

Based on the permanent information collected on running infrastructures, AI could foster the development of predictive on the fly maintenance avoiding costly and lengthily regular shutdown of HPC systems for major maintenances.

Ultimately, based on this fine grain profiling of information, AI could be used for developing next generation of wizard tools able in a post mortem way to provide tailored feedback to system administrators as well as end user about how their applications behave and what is needed to improve them.

3.1.2 *Why AI needs HPC?*

First, HPC and numerical simulations are generating a **lot of data which can be used for training AI models.** This often includes data which is already labeled and thus suitable for massively supervised DL. HPC is also a very good tool for rapidly generating missing data in the field of active learning for improving the training of AI models.

As an example, recent CMIP6 climate experiments performed in France generated around 14 PB of structured and labeled data, this data-lake is closely integrated with the Jean Zay converged system at IDRIS and will become a natural playground for fostering the growing synergies between AI and numerical simulation communities. While scale and variety (in order to avoid bias) of data matter for an efficient training of the models, HPC systems could be very useful by design (large amount of compute and storage resources) for generating large synthetics datasets for training or testing AI models.

The second reason, and the most important one, is related to the **compute requirements of AI while training models especially deep neural networks**. The high computational demand of fields like DL has contributed to emphasize the need of high performance hardware and software (e.g. general purpose computing on GPUs, GPGPUs), and therefore, put HPC technologies in the foreground. Figure 9 issued by OpenAI illustrates clearly this exponential growth of the compute usage for AI training since the inception of the first DL models (Perceptron). The availability of massive datasets (e.g. smart phone user data or social media usage) and the development of new AI models lead now during the so-called modern-area of AI to the doubling of the compute power needed every 3.4 months, exploding by far the well unowned micro-electronics Moore’s law.

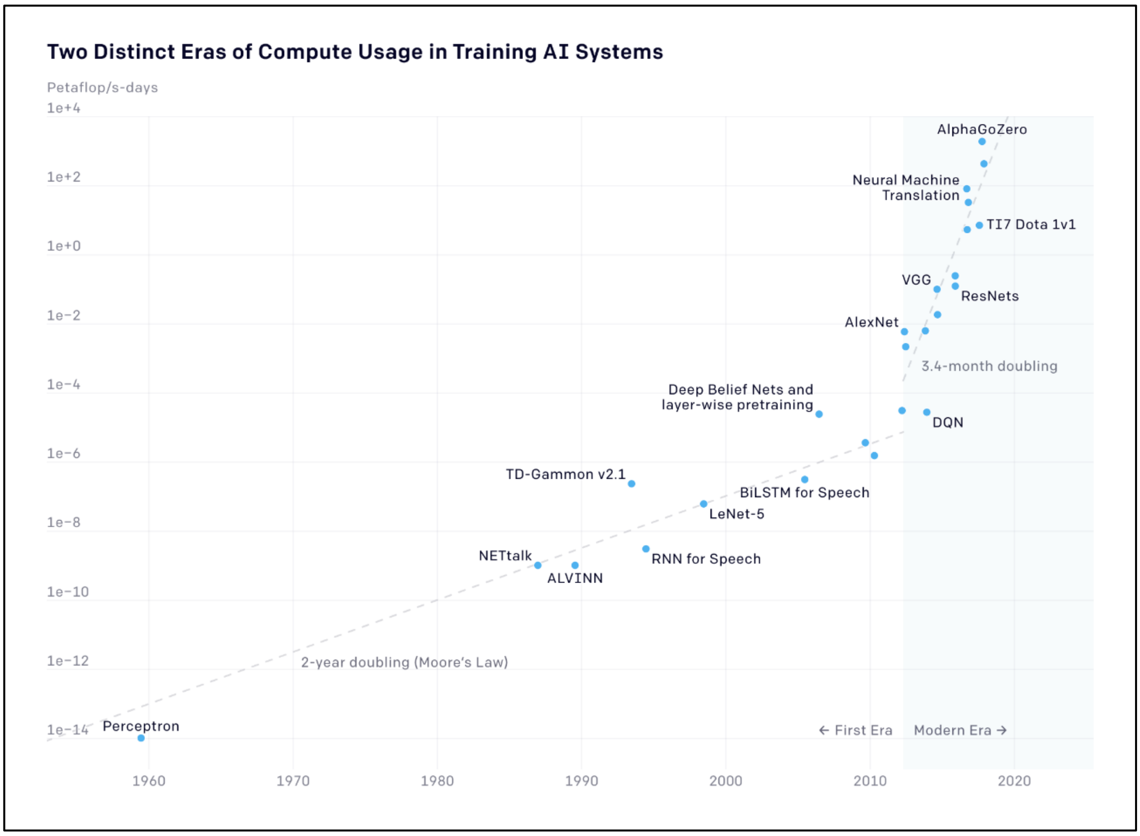

Figure 9 Two distinct eras of compute usage in training AI systems (source [11])

Beyond the raw access to HPC resources, the AI communities may benefit from the synergies with the HPC communities for optimizing (in terms of performance and energy efficiency) and scaling out the learning of AI models taking into account the multiple levels of parallelism of such process:

- parallelization of hyper parameter search (network type and topology, learning rate, hidden units, convolution kernel width, …) as one of the most expensive tasks in the process of training;
- parallelization of the different datasets across multiple batches;
- domain-decomposition and parallelization of the forward/backward gradient computation;
- intra chip ILP by using vectors, tensors (like NVIDIA tensorcores) for parallel local convolutions.

As an example, a research team from the Jülich Supercomputing Centre (Germany) accelerated the training of RESNET-50 (see Figure 10) on up to 96 GPU of the JUWELS tier-0 system but using the Horovod parallel framework and optimizing the collective operations needed while exchanging computed gradients at each phase of the training (see Figure 11). Such optimization will allow a faster and more efficient scaling of the training process, the use of more complex/deep/dynamic networks.

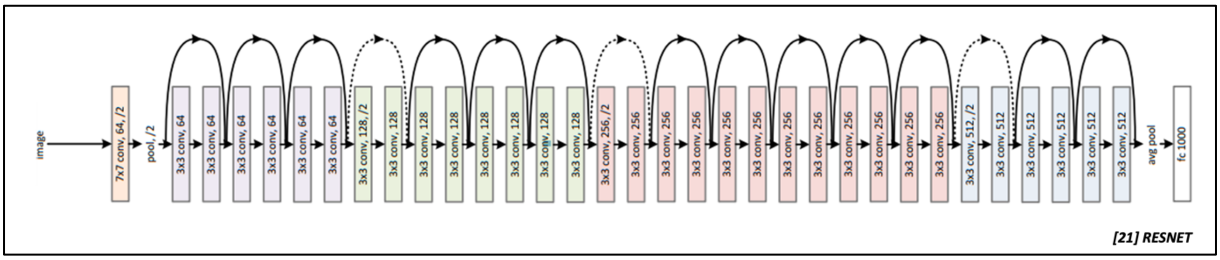

Figure 10 ResNET-50 pipeline (courtesy of M. Riedel)

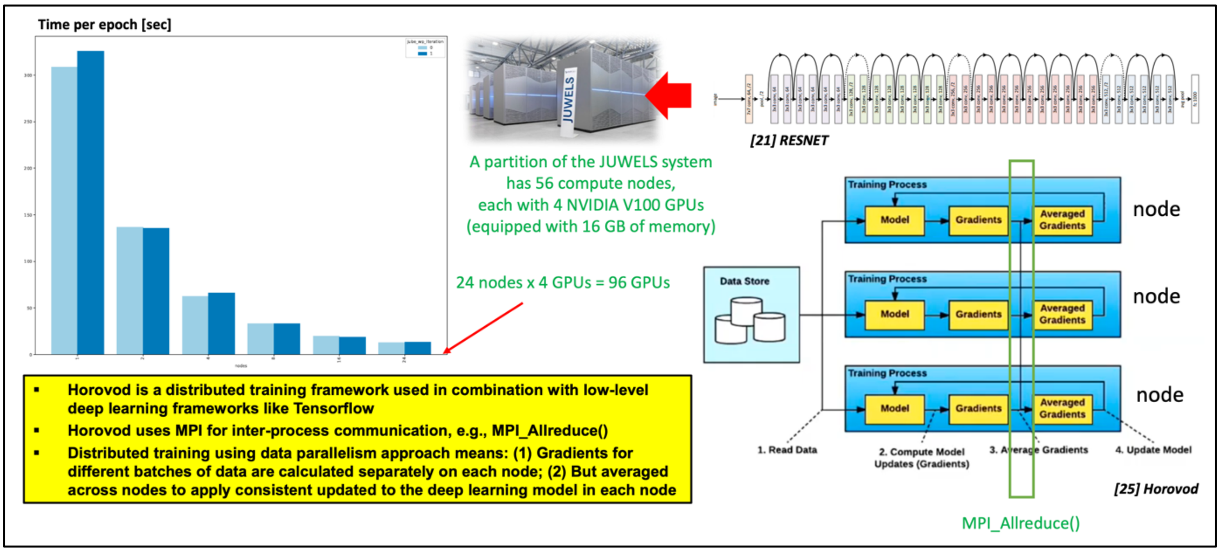

Figure 11 Optimization of ResNET 50 done at JSC. ResNET is a leading convolutional neural network that is trained on more than a million images from the ImageNet database. The network is 50 layers deep and can classify images into 1000 object categories, such as keyboard, mouse, pencil, and many animals. Its computational complexity of training relies on the fact that it has around 25.6 millions of trainable parameters which make RESNET-50 very suitable for parallelization via distributed training on multiple GPUs (courtesy of M. Riedel).

As another recent example, in 2018 a team from Inria, IRD, and the Institut national de la recherche agronomique (INRA) worked with Intel, SurfSARA, and GENCI on tier-1 and tier-0 HPC systems to scale out a plant recognition AI model (PlantNet, Figure 12) on up to 400,000 classes which was at this time a first ever. Previously, the mobile application was using only 17,000 classes for the inference of the photos submitted by users (more than 100,000 users/day) so the challenge was to expand it to the biggest plant and flora available database with 390,000 species around the world.

The framework used was Caffe using an optimized version provided by Intel while using GENCI's x86 facilities (Irene, the French tier-0 at Très Grand Centre de calcul, TGCC, at CEA and Occigen at the Centre Informatique National de l'Enseignement Supérieur, CINES). Multiples issues have been addressed including the memory scalability of the *protobuf* library and finally it has been possible to train a model with 390,000 classes in a few hours, which was clearly to reachable with current PlantNet HPC facilities. Figure 13 illustrate the good scalability on up to 1024 nodes (around 50,000 cores) even if tests have been conducted on up to 1320 nodes (63,000 cores). This example clearly emphasizes the impact of using HPC systems for the training of large-scale AI models.

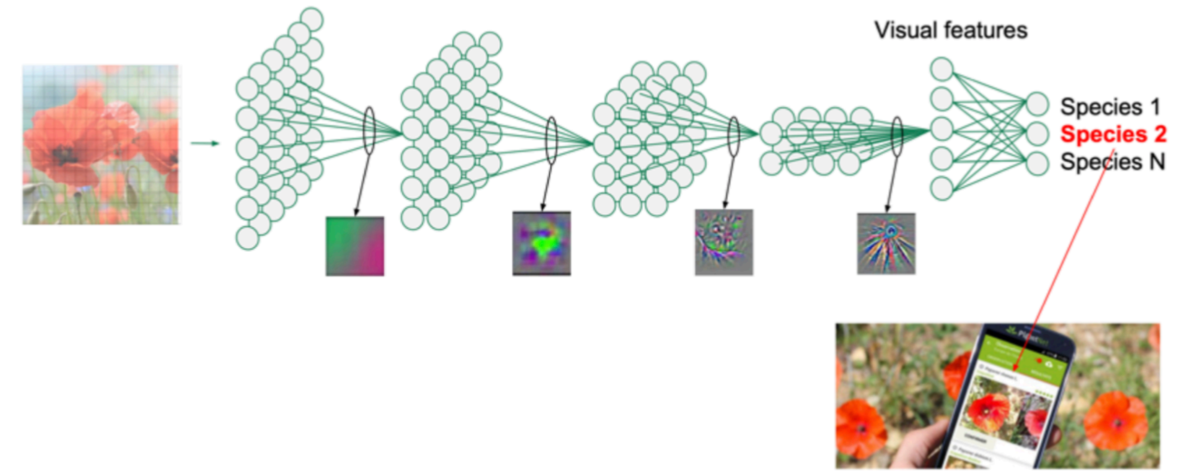

Figure 12 The PlantNet system (courtesy of A. Joly)

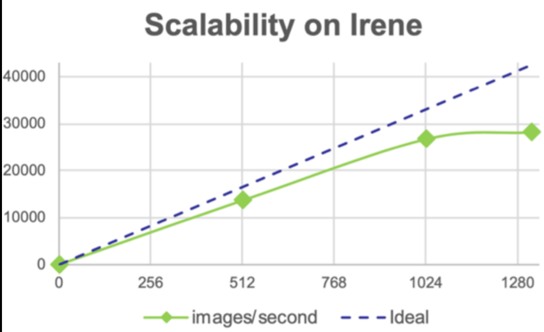

Nodes	global batch size	images/second	Scaling efficiency
2	64	64.5	**100%**
512	16384	13650	**83%**
1024	32768	26670	**81%**
1320	**42240**	**28160**	**66%**

Figure 13 Scalability of the training of the 390,000 classes PlantNet model on x86 systems (courtesy of A. Joly)

The use of **HPC will also foster the development and the use of automatic tools** like AutoML, AutoAI for auto-tuning of the models or the **development of generative adversarial network (GAN)** approaches like in high energy physics. Scientists from CERN are replacing Monte Carlo methods with GANs, an unsupervised learning technique where two neural networks compete to produce more accurate results, in order to accelerate their ability to study collision data from high-granularity calorimeters that measure particle energy. In this case, the AI code is able to produce similar results as the numerically intensive code, but uses only a fraction of the compute cycles, thereby enabling both bigger and faster models [12].

Finally, as the holy grail of the AI community, HPC could help in **developing the next generation of AI called XAI (eXplainable AI)** for providing more traceability, trust, and explicability on decisions taken by AI models, a major stake for the adoption and the trust on such tools by societies and be compliant with legal, regulatory and moral frameworks. Such approaches could be based on so-called hybrid AI connectionist and symbolic methods.

3.2 Towards a digital continuum from the edge to the tape

The second disruption is coming from the fact that HPC centers are less and less isolated but **fully integrated inside a global cyber-infrastructure ensuring a digital continuum of the data**, from the place where its generated like large scale instruments (satellites, telescopes, microscopes, sequencers, network of sensors, …) and IoT to the place where the final (and refined or useful) data is stored (or archived). In Figure 14 and Figure 15, issued from EXDCI-1 project/BDEC2, data is generated by instruments/devices at the edge, then pre-processed/inferred locally thanks to a small amount of (energy efficient and low cost) computing power, allowing to reduce the burden of data to be sent across communication networks and in some case enforce security/privacy (local processing only or encryption of data) for sensible data.

After this first stage of acquisition/processing at the edge level, data is now sent to a hierarchy of processing levels of (re)processing called fog computing thanks to upcoming 5G/Sigfox/LoRa and traditional wide area networks (WANs), until it reaches more traditional, larger-scale HPC systems such as supercomputers or HPC-enabled clouds. After this final stage of processing, data is archived and made available to wide communities using open data and FAIR principles transforming previous computing centers into (computing+data) centers. In some cases, either this post-processed data or synthetic data generated by simulations run on supercomputers are jointly analyzed with real-time data monitored on edge systems, as part of a global process allowing for continuous refinement and self-improvement of simulation models (example of real time digital twins or real time steering of large-scale scientific instruments or oil & gas major seismic acquisitions).

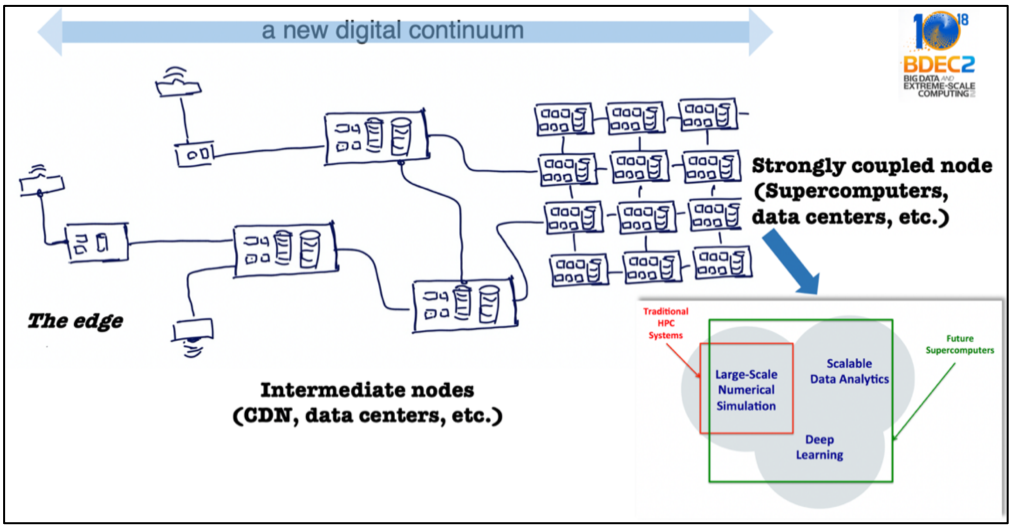

Figure 14 A new digital continuum: scheme (source BDEC2).

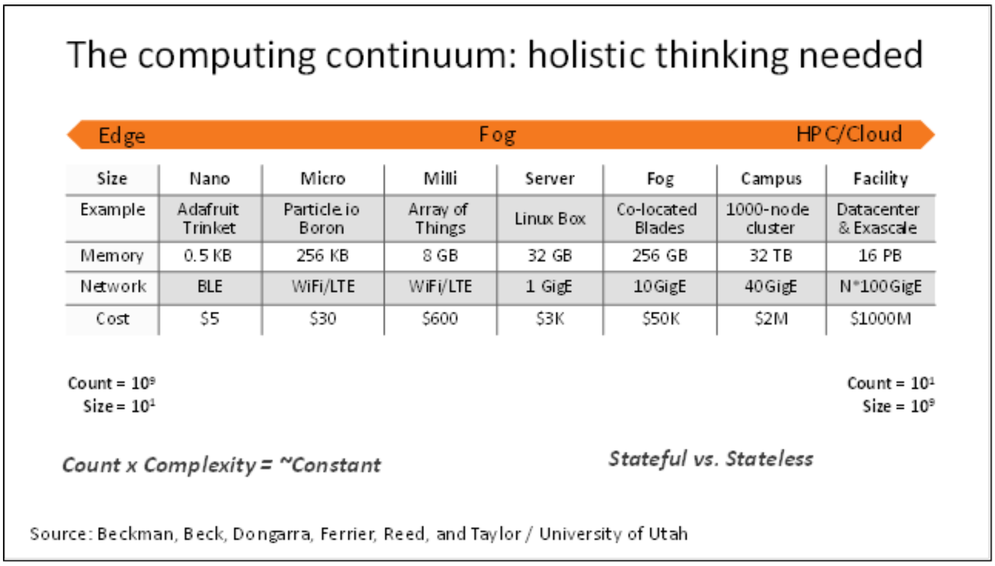

The computing continuum: holistic thinking needed

Edge — Fog — HPC/Cloud

Size	Nano	Micro	Milli	Server	Fog	Campus	Facility
Example	Adafruit Trinket	Particle.io Boron	Array of Things	Linux Box	Co-located Blades	1000-node cluster	Datacenter & Exascale
Memory	0.5 KB	256 KB	8 GB	32 GB	256 GB	32 TB	16 PB
Network	BLE	WiFi/LTE	WiFi/LTE	1 GigE	10GigE	40GigE	N*100GigE
Cost	$5	$30	$600	$3K	$50K	$2M	$1000M

Count = 10^9
Size = 10^1

Count = 10^1
Size = 10^9

Count x Complexity = ~Constant

Stateful vs. Stateless

Source: Beckman, Beck, Dongarra, Ferrier, Reed, and Taylor / University of Utah

Figure 15 A new digital continuum: table (source BDEC2).

As stated inside the forth version of ETP4HPC's SRA, ensuring such digital continuum across a well-connected heterogeneous cyber-infrastructure will need to address major challenges in:

- **ensuring a proper and secured data logistics** with unified data storage abstractions and systems enabling efficient data sharing across the digital continuum with federated resources: data have to be exchanged from edge devices to HPC-class machines, therefore the data should be presented in a coherent and easy to use form for all machines in the "continuum";
- unified real-time data processing techniques favoring the joint use of HPC-originated approaches such as in-situ/in transit processing with stream-based processing techniques now common in big data analytics frameworks;
- **interoperability and standardization of the different software stacks** from the edge to the HPC center either for HPC, HPDA, and AI workloads, in that goal secured containerization could be a solution;
- (near to) real-time computation to generate near to real-time decisions: as most cyber physical systems will have to close the loop between sensing-computing-acting by reacting on the physical world, the time constraints of the physical world will drive the maximum execution time of the part sensing-computing-acting;
- cohabitation of processing **in stream mode together with classical batch access mode** using smart resource schedulers: the real world generates data continuously, and due to the expected real-time reaction (see above), this implies continuous processing of streams of data and elastic allocation of resources;
- enforced privacy and data security: As HPC systems will be open to "untrusted" data and accesses from outside, the requirements of security and ethics such as privacy should be enforced;
- **dynamic end to end application workflows management** (allowing coupling of simulation, databases and data streams, data analytics and visualization that interact together in real time), in that sense increased interaction between the HPC, data analytics and AI communities is critical;
- as these challenges will also have impact on the mindset of people operating these new generation of HPC systems education and training of new skills in data management and increased user support.

Some examples can be given of the use of such continuum for a couple of applications. These new services are pilot implementation of new business model called "HPC as a service", that will be analyzed as new way to access the futures European Exascale HPC facilities.

3.2.1 *AQMO: an edge-to-HPC digital continuum framework for air quality*

Air quality improvement is a major challenge for most metropolises. Proposing efficient policies to address this challenge requires solving two issues. The first one is to perform air quality measurement with a thorough temporal and spatial coverage. The second one is to understand the dispersion of the pollution as well as being able to analyze "*what-if*" scenario. The measurement issue is addressed using more sensors while the second one related to the use of HPC numerical simulations. Of course, the two issues are intimately entangled. The measurements provide the basis to elaborate the model inputs and validation while the numerical dispersion model gives insight in how the pollution reaches the citizens.

The AQMO project provides an **end-to-end urban platform that extends current practices in air quality measurements**. Figure 16 shows an overview of the digital continuum design to implement the platform. This **continuum integrates edge technology, cloud facilities and supercomputers**. It is intended to provide citizens, local authorities, scientific organizations and private companies with new open-data and innovative services based on computing simulation (HPC and edge-computing/IoT).

To implement air quality analysis in a cost-effective manner in a wide area, the local transportation bus network is embedding mobile sensors. In the case of measurements for catastrophic event, use of drones is explored. This strategy allows to use fewer but more accurate sensors. The edge computing part of the continuum aims at two main functions. The first one is the storage of the collected data in order to manage the intermittent communications issue. The second one is to respect citizen privacy. Indeed, the platform makes use of cameras to detect if a bus is stuck behind a truck or another vehicle (which tamper with the pollution measurement). An AI engine is implemented at the edge and only the image analysis is sent back, no images are stored. The resource continuum is logically assembled using a global workflow management technology and is design to support new sensors as well as complex numerical models in a routine usage in the context of a smart-city.

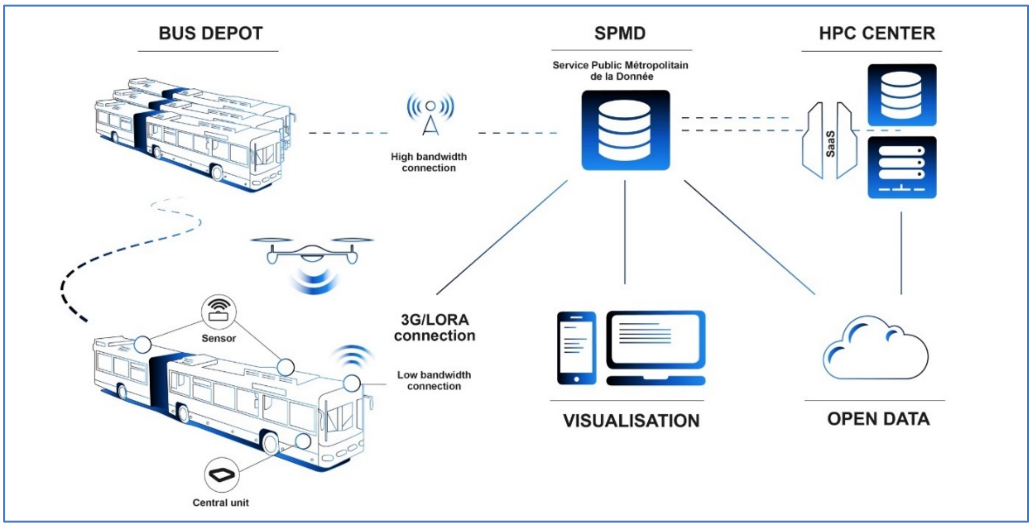

Figure 16 AQMO continuum overview (courtesy of F. Bodin)

3.2.2 *PHIDIAS: an edge-to-HPC digital continuum framework for earth system sciences*

The European project PHIDIAS will address the development and concrete realization of a set interdisciplinary services and tools based on HPC for earth system sciences to exploit large datasets of public European interest provided by satellite observation of earth, and provide fair access to these processed datasets and value added services (from "standard" data processing applied to big data heterogeneous datasets, to more advanced services such as AI or HPC on demand, called "urgent computing") through large data storage capability and high bandwidth network across Europe (see Figure 17).

The solutions which will be tested and validated leverage existing HPC and data management capacities and are suitable *by-design* to spin-off across other scientific fields. In particular PHIDIAS will develop and offer a catalogue allowing users to discover and access data, but also relevant open source software, public application programming interfaces (APIs) and also interactive processing services. This catalogue will implement interoperable services for the discovery, access and processing of the data, and be connected to other major data repositories such as the European Data Portal, GEOSS, NextGEOSS, and the European Open Science Cloud (EOSC). PHIDIAS will also implement an end-user web common interactive processing service based on notebook and datacube technologies. This, coupled with powerful toolboxes, such as the Orfeo Toolbox, and AI methods, such as convolutional neural networks, will allow new users to easily (since the data are pre-processed in the datacube) have access to the HPC capacities and develop new algorithms. This may lead to commercial applications, and new decision-support tools for public authorities that could be executed on the DIAS ("Copernicus Data and Information Access Services"), and that will be evaluated as part of the economic equation for a sustainable business model for the Industrialization of these proof of concept services. Finally, PHIDIAS will also directly contribute to the development of EOSC by industrializing its HPC/HPDA/AI processes and algorithms, as well as making its services available and accessible through the EOSC portal, in compliance with the EOSC rules of participation.

This will not be business as usual for HPC, where resources are mainly allocated through a Peer Review process. This new usage, demonstrating the capacity to offer operational capacity for near real time link to the instrument (edge) for data processing, but also on demand access to computational resource, up to urgent computing automatically triggered in case of automatic detection, on the fly, of natural hazards, will require adaptation of the data flow of the HPC and data centers involved within this project across Europe, to offer Infrastructure As a Service services (IaaS). This will require complex architecture design, addressing challenges on end-to-end security and SLA needs for these new usage and best practice acquired from project partners with demonstrated skills on spatial data logistics and HPC management. By doing this, PHIDIAS not only aims to leverage the large amount of valuable data generated by these research infrastructures but also to increase usage and ease access to HPC and data storage capacities which represent the foundations of the emerging European data infrastructure that is need to handle huge societal that our society are facing, such as climate change and biodiversity preservation.

PHIDIAS will address three concrete use cases:

- in oceanography: improve the use of cloud-based services for marine data management, service and processing, considering the DIAS and the EOSC challenges. This includes improving long term stewardship of data (storage, capacity, archive, metadata enhancement, …), mix of in-situ data with other data sources especially satellite data for cross validation or resolution improvement, ... requiring new data storage and processing capacities;

- an intelligent screening of satellite data for detection and identification of anomalous atmospheric composition events in order to improve accuracy and genericity. This will include an HPC/AI based smart filtering service in order to 1) detect relevant data within the huge volume of measurements, 2) identify and assess the pertinent information, and 3) qualify the targeted data for dedicated exploitation by users and provide access to filtered data and corresponding metadata;
- big data earth observation by enhancing the toolchains scalability for environmental monitoring from the end-users needs of Theia land data centre network and geoinformation for sustainable development national satellite data infrastructure for environmental research and environmental management and policy.

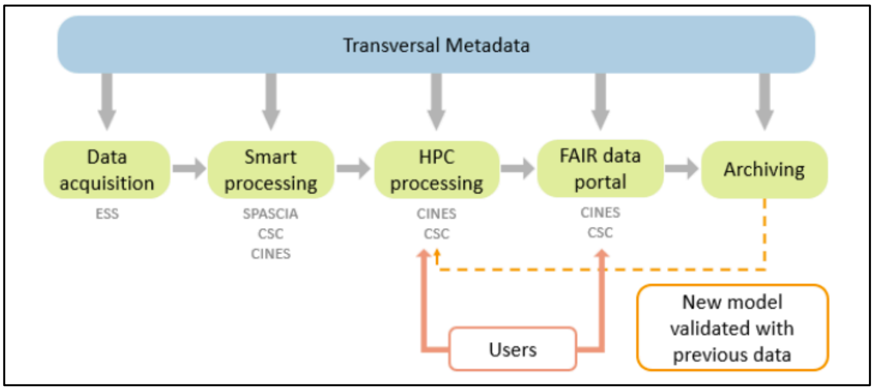

Figure 17 Synthetic view and components of the main workflow of the PHIDIAS project (courtesy of B. Dintrans)

3.3 Leading to new services like urgent computing

After being widely used since decades for science and innovation, converged HPC+AI systems inside a fully connected digital continuum are now also used for (real-time) decision making for public and private bodies. In a case of natural hazard (like earthquake, tsunamis, flooding, …), biological hazard (propagation of pandemic like the ongoing 219-nCoV coronavirus), industrial damage (rupture of a dam, incident on an industrial plant, …) or (cyber)terrorism, HPC systems could help when properly parametrized (including installation, regular tests in red button mode of the full alert workflow) used by allocating HPC and storage facilities on the fly. Such use modes are developed by the ChEESE CoE in Europe.

Early warning of natural hazards is essential for an effective mitigation of related socio-economic and environmental impacts, particularly in case of tsunamis, earthquakes and volcanoes. Prompt reaction to these scenarios requires of computing infrastructures, complicated data workflows, and engagement with stakeholders formally involved in emergency management (e.g. the European Emergency Response Coordination Center, ERCC) through shared protocols and policies.

However, urgent supercomputing during the aforementioned phenomena involves complexities at many levels. From a technical perspective, the frameworks required to manage the data intake, analysis, simulation pre-processing, execution (often on tier-0 systems), and final post-processing are very complex. On the other hand, access policies to HPC clusters that account for the time constraints imposed by urgent simulations are radically different from current resource access at public supercomputers.

The CoE ChEESE is preparing flagship HPC applications in earth sciences and natural hazards for the Exascale era. The CoE has a strong commitment to build services for industry and public governance bodies (e.g. civil protection), including **urgent computing end-to-end workflows of data coming from networks of sensors to HPC resources s for earthquakes and tsunamis**. A series of pilot demonstrators on urgent seismic simulations, Faster-Than-Real-Time (FTRT) Tsunami simulations, and high-resolution volcanic ash dispersal forecast are being developed to test future Exascale urgent computing workflows supporting contingency plans for seismic, volcanic and tsunami events. On the other hand, building a prepared society resilient to natural catastrophes requires the capability of managing in a probabilistic framework the complexity of the natural phenomena and the large uncertainty associated with their development. This requires performing large *ensembles* of accurate scenario simulations to reproduce the complex physics of the natural systems and the wide variability of initial and boundary conditions.

This is the probabilistic hazard assessment, which ChEESE will implement in another series of pilot demonstrators to answer questions like:

- How big will the next volcanic eruption be at Jan Mayen in Norway? How will that potentially impact the population and the air traffic in Europe, considering the statistical variability of wind intensity and directions?
- What is the probability of ground acceleration exceeding a given threshold at a nuclear plant site, or for a critical infrastructure in Europe? What is the probability that a large earthquake in the Mediterranean produces a tsunami wave higher than one meter in the Marseille harbor?

Another example is the tsunami service providers, providing tsunami warnings in the framework of the systems coordinated by IOC/UNESCO worldwide, and other national tsunami warning centers, are striving to complement, or replace, decision matrices and pre-calculated tsunami scenario databases with FTRT tsunami simulations. The aim is to increase the accuracy of tsunami forecast by assimilating the largest possible amount of data in quasi real time, and performing simulations in a few minutes wall-clock time, possibly including the coastal inundation stage. This strategy of direct real time computation, that could seem unfeasible a decade ago, it is now foreseeable thanks to the astonishingly recent increase in the computational power and bandwidth evolution of modern GPUs.

This is "Early Warning" and, for that, ChEESE will make use of the most efficient HPC architectures, software and workflows to demonstrate prototype early warning systems for tsunamis. It is depicted in Figure 18.

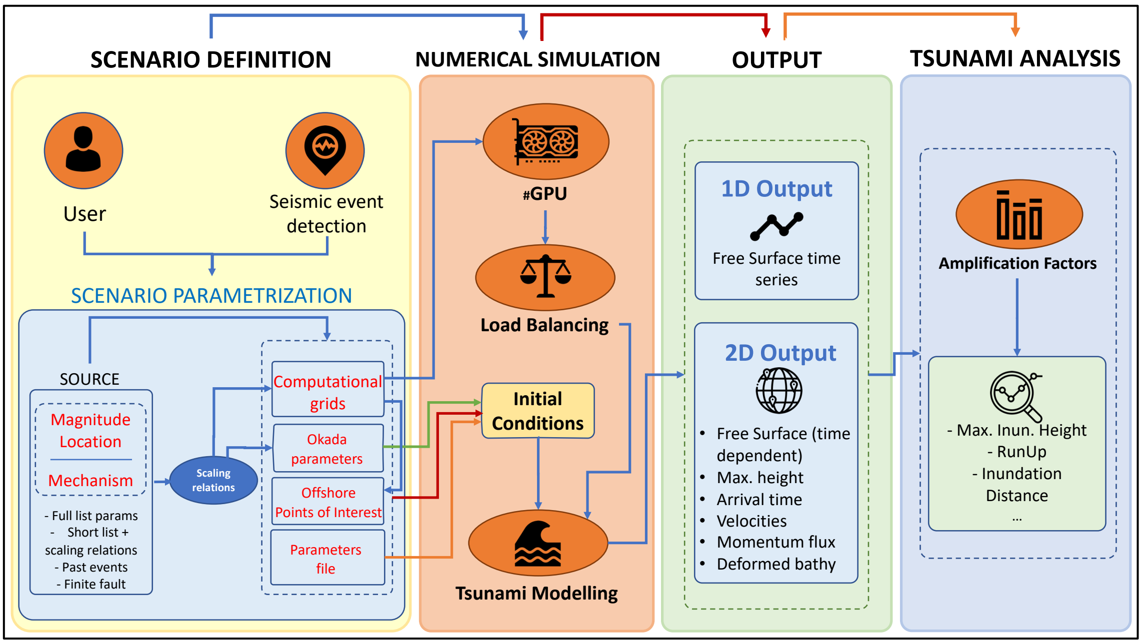

Figure 18 Example of a complete end-to-end workflow of real time tsunami alert system (courtesy of A. Floch)

4 Some updates of application roadmaps

Before describing in a more detailed manner how different scientific disciplines and industrial and engineering applications are considering their future, it should first be recognized that there are now three main categories of roadmaps. Depending upon the relative importance of the underlying physical laws, if any, *vs*. the information content of a wealth of data relative to the field, if any, the various challenges can be addressed in different ways:

- need for always increasing compute power (Exascale and beyond) for simulations based mostly on physical laws.
- need for using ML, more precisely computer DL, to extract knowledge from data of various sources. Most of the time the data come from earlier collections or from experiments performed on the physical system at stake;
- need for developing new *hybrid* methods, to take the most advantage of both approaches.

The simultaneous use of HPC for solving the physical laws and of ML, based on either experimental or computer-generated data, for replacing some parts of the HPC calculation is indeed a way to reduce "time to solution", among other possible benefits. It seems that such hybrid methods, although still in their infancy for a large number of them, are developing quite rapidly. Examples, most of them detailed below, concern meteorology and climate, nuclear fusion, biosciences and human health, combustion, material sciences, social sciences, … It should nevertheless be emphasized that the convergence between HPC and ML is only at its beginning as far as applications are concerned. For most of the application developers, their roadmaps are in a large part still dominated by the need to access to significantly increased computer resources, in a "more-or-less traditional" HPC manner. The convergence with ML is included in a slightly more modest way by most of the application developers, e.g. to document which parts of the calculations could benefit from such a convergence. This is underlined below, for each of the application domain at stake, by using special characters (*italics*) to recognize the parts of the roadmaps benefiting from this convergence. We should indeed recall here that the aim of the present deliverable is to report about actual roadmaps produced in the various application domains, even if these roadmaps are not yet taking full advantage of the foreseen convergence described above. Another specific example, which also summarizes some of the developments described in Chapter 3, is given in the next section.

Despite the above differences, it should already be underlined that all the above approaches do require increased HPC-compute and data-processing power to address the challenges, as they refer to systems of ever-increasing complexity and multi-disciplinarily nature.

4.1 Application example for HPC/AI convergence

Subgrid models are required for many predictive simulations. One way to develop such models is to perform fully-resolved simulations for generating "truth" data. This data can then be analyzed with ML/DL or used for training ML/DL models.

Recently, German researchers presented different reconstruction- and regression-based data-driven approaches [13]. They used HPC to generate large datasets, which were then fed into a developed DL framework. Training the DL networks was only possible by using again supercomputers as the amount of training data was very large showing the convergence of HPC and AI.

More precisely, Figure 19 shows the generator and discriminator parts of the used DL network used by the German researchers. As the network is used to model underlying physical processes, the researchers emphasize the need for developing physics-informed networks [14].

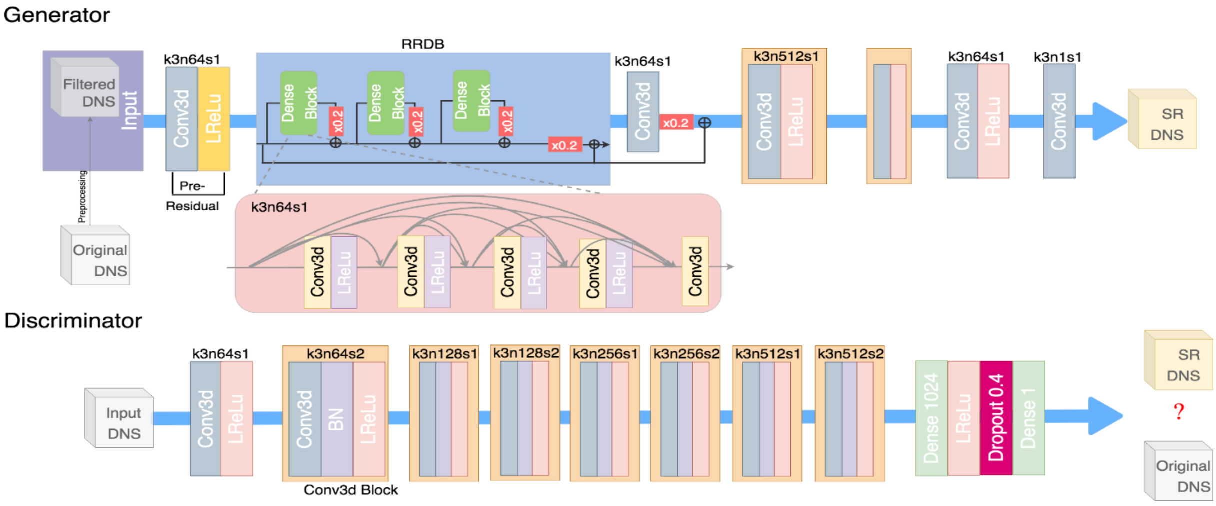

Figure 19 DL network used for subgrid model development. Data generation and training were performed on supercomputers (courtesy of M. Bode).

4.2 Weather and climate applications

Roadmaps for weather and climate share a number of common issues and challenges, among which the need to reach higher spatial resolution is a long-standing concern. High spatial resolution is indeed crucial for predicting weather within the next few days as local meteorological phenomena affect almost all human activities and therefore have to be predicted at the scale where these activities take place, i.e. locally. Kilometer, or even sub-kilometer, scale resolution weather predictions are presently within reach and this will drive many developments over the next five to ten years. High spatial resolution is also required for climate modeling: The Intergovernmental Panel for Climate Change (IPCC) had concluded in its fifth assessment report that detection and attribution studies focused on extreme events were constrained by model resolution. Recent studies have shown that enhancing the horizontal resolution of models is seen to significantly affect aspects of large-scale circulation as well as improve the simulation of small-scale processes and extremes when compared to earlier, lower-resolution models.

There are also specific challenges for each of these domains. For example, meteorology has to interface with an extremely large number of users, so that weather models have to include the generation and management of data bases so that the users can interface on an operational basis with their own models. The consequences are developed below. On the other hand, climate simulations, although they are thriving toward kilometer-scale resolution on the longer term, will still need to rely on sub-grid scale parameterizations for a large number of years. *They are now methods relying on ML, applied to either experimental or computer-generated data, which may facilitate addressing the problem of scale-dependency of parameterizations. This is also developed below.*

At the same time such physical improvements of models are undertaken, it is also of utmost importance to adapt codes to the new architectures, with particular concern to accelerators. The

European state-of-the-art in this field is quite diverse, with a number of weather or climate models being still well adapted to classical CPUs, while initiatives to transform the codes for an efficient use of GPUs has shown clear success. It is likely that this trend will continue to develop and, quite likely, to accelerate.

4.2.1 *Weather prediction*

Predicting the weather at higher spatial and temporal resolution with longer lead times comes firstly as a result of improved model initialization. Such improvement is driven by more accurate assimilating methods, the wealth of data coming each day from in-situ networks, satellite-borne instruments and, soon to come, connected sensors on all type of personal devices (cell phones, cars, …). Observational data are delivered from actual systems at a rate close to one billion per day, but this rate will keep regularly increasing due to more numerous, more refined satellite-borne sensors and to new connected sensors. Furthermore, the numerical weather prediction model must be run very fast, so that it delivers its results for the forecast to be delivered in time to the public and all types of users. Typically forecasting the next 24 hours should not take more than a few wall-clock minutes.

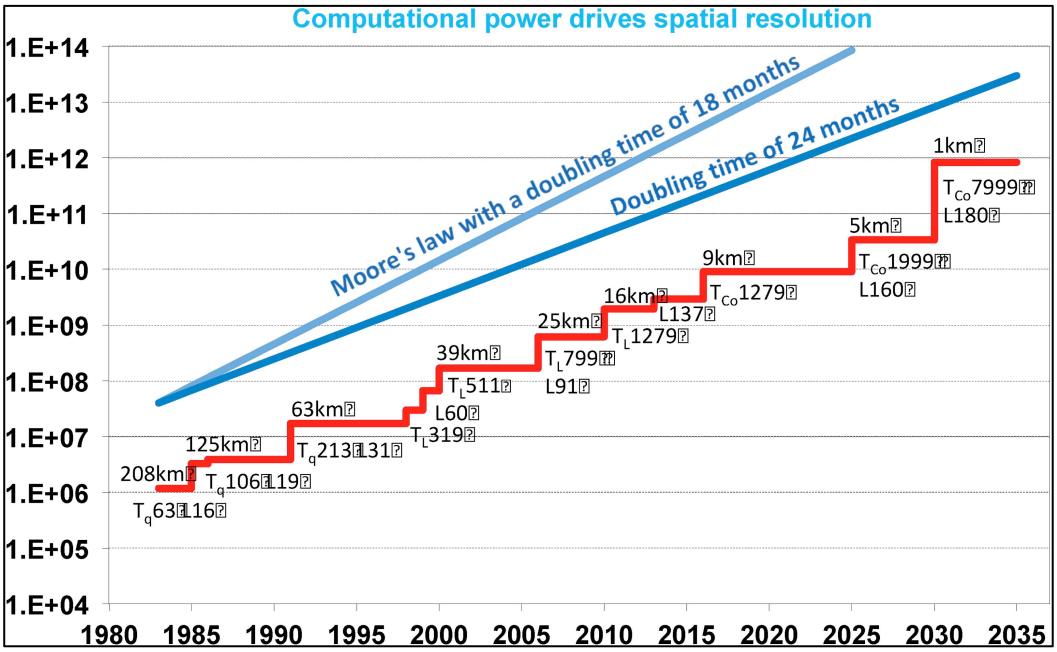

Figure 20 Progress in the degrees of freedom of the European Centre for Medium-range Weather Forecasts (ECMWF) operational global atmosphere model in comparison to Moore's law. The ambitious goal of reaching an operational 1 km horizontal resolution with 180 levels in 2030 requires a faster progress. The numbers indicate the average grid-point distance in kilometers and the corresponding spectral resolution and levels used (source [15])

The horizontal spatial scale of the most advanced *global* numerical weather prediction models is now about 10 km (see Figure 20), although less-demanding *limited-area* models are run at the 1 km horizontal scale. There is whatsoever always scientific incentive to run at still improved spatial resolutions as a number of atmospheric phenomena, and especially extreme weather events, have important features at sub-kilometer scale. It should also be emphasized that the value of weather simulations comes not only from high-resolution models, but also from two other factors:

- the number of realizations which can be performed during the forecasting time. Atmosphere is indeed a deterministic chaotic system and consequently, despite improved assimilation methods used for providing the initial state of the forecast, uncertainties present in the smaller scales of the initial state lead, over a few days, to a progressive decorrelation of atmospheric trajectories. Characterizing the rate at which this scatter develops is still the only way to assess the decreasing value of the forecast over the forecast range. Present numerical weather prediction models may use as much as 100 realizations for producing a single forecast;
- the model used for the prediction system must also be enriched for better physics, better interactions between the atmosphere and its interfaces (soils and vegetation, upper ocean, atmospheric chemistry, …). While the early (hydrostatic 2D) models had only three prognostic variables (the two horizontal components of the wind and the temperature), actual non-hydrostatic 3D models include a number of additional variables (vertical wind, pressure, soil moisture, sea-surface temperature, …). Furthermore, subgrid-scale phenomena are parameterized with schemes of increased sophistication and realism.

All this results in ever-increasing compute power. There are indeed 12 computers reserved for meteorological forecasting within the first 150 entries to the last Top500 list.

A further new trend is to allow users of the forecasts to interact directly with the prediction system, in order for them to extract information at full scale. This would mean allowing users to actively shape these prediction systems, reconfiguring the "value chain" as a collaborative space. Such an opening up of the prediction systems would not only accelerate the uptake of information by existing communities but also create opportunities for innovation by new application communities. Such developments would require a comprehensive application programming interface to extreme-scale computing applications.

4.2.2 *Climate modeling*

Simulating the climate and its changes under anthropic influence shares, as already said, a number of challenges with numerical prediction: need for improved spatial resolution, need for larger ensemble of realizations, and need for improved physical description of the system. But for each of these challenges, there are specific "climate features".

The long-term goal is still to develop and run global coupled climate models at the 1 km horizontal scale for two main reasons:

1. kilometer-scale phenomena are of crucial importance for driving climate processes and understanding greenhouse effect and its further change, in particular convection, *i.e.* the vertical transfer of heat within the atmosphere through convective cells, and the resulting development of clouds;
2. a number of extreme meteorological events, either in actual or future climate are kilometer-scale events. One main difference between simulation of weather and simulation of climate is nevertheless the length of the simulated time. While the counts are in days or weeks, up to a few months, for weather forecasting, they are in decades or centuries, up to millennia, for climate. Such an increase of the length of the simulations by about four orders of magnitude has a number of consequences. Even accounting for the fact that, from a wall-clock perspective, a climate simulation can be realized in months, as compared to hours for weather, therefore reducing the gap by three orders of magnitude, the actual rate of annual increase in compute power shows it will not be possible before years to run coupled climate simulations, as they should, at the same spatial resolution as weather simulations. As an example, the most advanced high-resolution coupled climate models are still in the demonstrator phase under the

ESiWACE CoE, with a resolution of 9 km in the ocean and 16 km in the atmosphere (see Figure 21). This means that climate modeling has also to develop alternative strategies.

There is indeed a developing trend to emulate and derive much-improved parameterization schemes, therefore making simulations with actual spatial resolution more realistic. ML is key to such developments: data coming either from in-situ observations or from very-high resolution models of specific phenomena (e.g. convection, cloud microphysics and cloud formation, surface vegetation, ...) are used as well-documented data bases for deriving, from the climate model variables themselves, the requested output from subgrid-scale phenomena. Such a strategy would also further allow for automatic scale-dependency of sub-grid parameterization schemes, a long-standing issue in climate modeling. This type of hybridization between physically-based laws and data-based information is likely to become more and more popular for climate modeling.

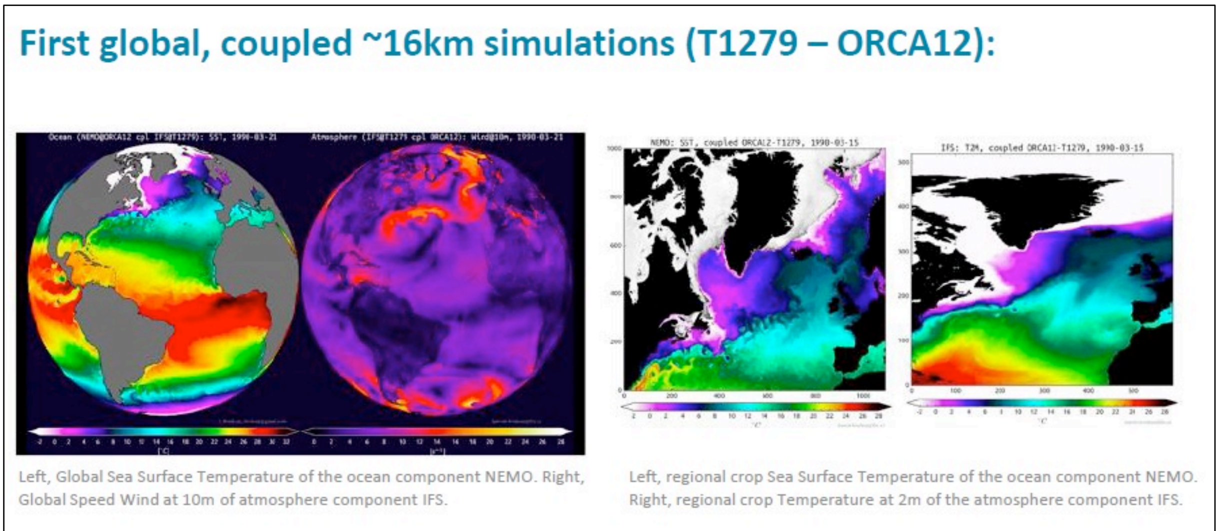

Figure 21 Test simulation of climate with high-resolution coupled climate model (courtesy of M. Castrillo)

4.3 High-energy particle physics, astrophysics, and plasma physics

These fundamental sciences address deep questions for humanity and society about the nature and origin of our Universe and the matter it contains. Plasma physics can also have unparalleled industrial and economic impact within the next 20 years by demonstrating of the feasibility of a clean and unlimited energy source based on nuclear fusion. In last few decades, all three disciplines have undergone a dramatic change and accelerated progress due to the development of powerful observation and experimental facilities and in tandem the development of HPC infrastructure.

4.3.1 *High-energy particle physics applications*

Such applications seek to explain the nature of matter in the universe. The standard model of particle physics is the subject of stringent experimental and theoretical research at the LHC at CERN. A significant component of these tests is the requirement for precise, robust calculations of strong interaction effects. This can be done through QCD, for which one has to turn to discretization of space-time on a four-dimensional hypercube, the so-called Lattice QCD

(LQCD), for achieving useful numerical simulations. The available simulation codes for LQCD are highly advanced and perform well on many HPC architectures and computing paradigms. The simulations can be described briefly as the solution of a very high-dimensional integral, the QCD path integral, by Monte Carlo methods, followed by the evaluation and averaging of observables on the stochastically generated ensembles. This requires repeated calculations of the determinant of a large sparse matrix and the inversion of a high-dimensional matrices. There has been enormous progress in the last ten years, but the extension of the LQCD approach to new problems will require at least a 1000-fold increase in computing power. The LQCD community has been highly successful up-to-now in making use of GPUs, but present multi-level methods will require larger memory footprint and hence communication bandwidth, which is a primary concern for future LQCD implementations.

4.3.2 *Astrophysics*

Any breakthroughs or improvements in the answers to fundamental questions in **astrophysics** (e.g. origin and evolution of the universe, dark matter and dark energy, solar physics, …) will have deep effect on society, an example being the recent proof of existence of gravitational waves. Space weather directly impacts many activities and its predictive modeling is crucial for a technologically-dependent society. Very large amount of data from observational facilities, one typical example being the SKA telescope, need to feed numerical simulations, for both proper interpretation and discrimination between models. Such high-performance simulations use combined fluid dynamics and kinetic approaches.

There are at least three scientific drivers for accessing more powerful HPC infrastructures:

1. increase by orders of magnitude the Reynolds numbers in simulations;
2. increase the resolution of the full universe resolution simulations and better take into account gravitational physics;
3. consider several fluids like baryonic gases and hydrodynamics into n-body simulations. The complexity of astrophysics and cosmology modeling and comparison with observations leads to hundreds of terabytes of data per simulations. Huge data management of 4D structures are required to understand the complex nonlinear physics and feedback among the various scales/objects/processes. Immersive (remote) data visualization is required to identify key structures. It should be mentioned that hybrid methods, associating physics-based modeling and DL, are presently developed for more efficient simulations. As an example, for the re-ionization which affected the cosmic gas 10^9 years after the big-bang due to the propagation of radiation from the early astrophysical sources is simulated by splitting the calculations in two parts. DL using data issued from simulation of radiative hydrodynamics is providing re-ionization processes, which can then be incorporated in non-radiative, and consequently much cheaper, simulation. The results, shown in Figure 22, demonstrate the value of this hybrid method.

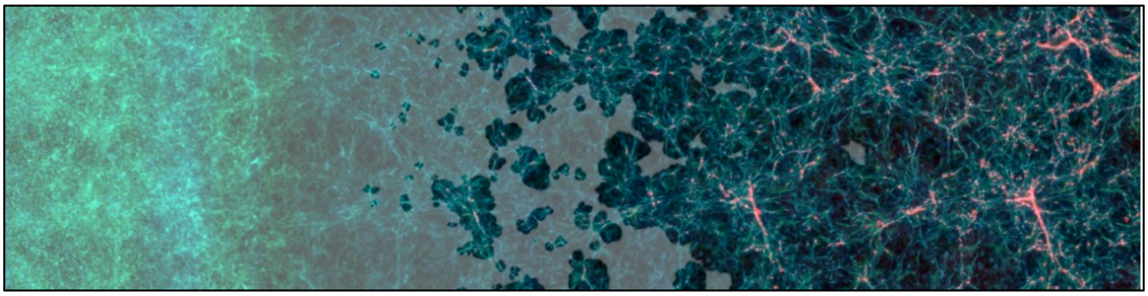

Figure 22 Space-time view of the re-ionization process. Vertical axis shows a region of space extending over 12 Mpc ("Megaparsec"), while the horizontal axis covers the first 109 years of the universe. Hor regions are shown in red, dark regions correspond to ionized zones, and the variation in green color refers to the gas density (courtesy of D. Aubert).

Some early experiments of coupling numerical methods and ML are ongoing in the field of the study of the re-ionization process where costly operations of radiative transport are replaced by a learnt model and coupled with the rest of the simulation, i.e. gravitation and hydrodynamics.

The team of D. Aubert from Observatoire de Strasbourg is currently implementing such approaches based on auto-encoders on the Jean Zay system at IDRIS on up to 256 GPUs applied to cubes of 128 Mpc on meshes of 1024^3 (see Figure 23).

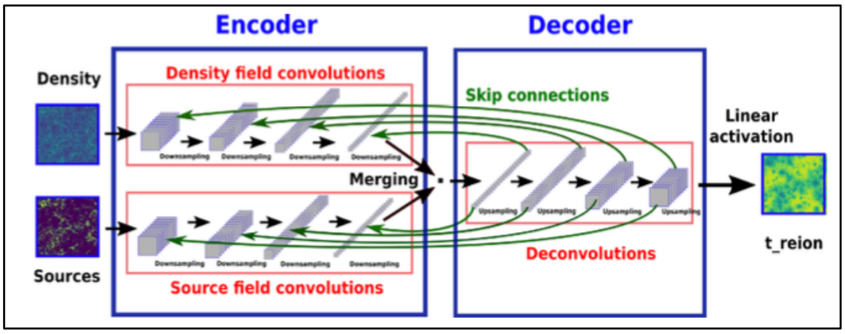

Figure 23 Coupling of the HPC/AI re-ionization code using auto-encoders for the radiative transport module (courtesy of D. Aubert)

4.3.3 *Plasma physics*

Plasmas physics has a very wide range of technological applications, including the destruction of toxic materials, modification material surfaces (coating), food processing, plasma torch, and cancer treatments, not to mention the demonstration of the feasibility of an energy source based on nuclear fusion. Theory and predictive modeling play indeed crucial roles for ITER ("International Thermonuclear Experimental Reactor"). The most challenging in terms of HPC resources is the so-called first principles modeling codes that are concentrating on the detailed understanding of, e.g., the many instabilities which lead to operational limits. These challenges are not very different in nature from the ones faced in other fields such as astrophysics, fluid dynamics and weather forecasting: they stem from the wide range of spatial and time scales to be covered, here from meters and seconds at the machine macro-scale down to millimeters and microseconds magneto-hydrodynamics (MHD) and turbulence microphysics. Nearest future developments in actual gyro-kinetic fusion codes are concerned with the modeling of both ion and electron dynamics. In the next ten years, MHD first principles modeling should approach the realistic plasma parameters for ITER as well as further developing their level of complexity and integration by including plasma-wall interactions. The final goal of such modeling is the “numerical tokamak”, with increased predictive capability for ITER. However, for the next ten years, one can expect only more modest achievements, such as in global MHD-electromagnetic turbulence models with plasma core-edge-surface integration. *The need to increase the spatial resolution, as required for capturing the many instabilities that occur in fusion plasmas and may be detrimental to ITER and tokamaks, could be somewhat alleviated by using hybrid methods between “traditional” HPC and ML. It indeed appears that data from earlier tokamaks (e.g. the Joint European Torus) can be used to train a method to predict the appearance of disruptive instabilities when simulating the plasmas in the ITER machine [16]. As for the re-ionization problem mentioned above, the rationale for such a hybrid method is very close in principle to the one already described for weather and climate simulations.*

4.3.4 *Some recommendations*

Some recommendations are emerging for preparing to Exascale. Due to small-scale nonlinearities and instabilities, many codes are using implicit algorithms, and hence require inversion of large (often ill-conditioned) matrices leading to large memory footprint: having sufficient memory per core remains then essential. Both communication and memory bandwidth remain then a primary concern. The current trend is however to offer new computing architectures with limited memory-per-core. Adapted numerical methods (especially in matrices inversion) should then be rapidly developed to adapt to such future Exascale architectures. Strong collaborative groups or consortia which combine different specialists in physics, applied mathematics and computer sciences should be strongly encouraged. It may, however, take some time before new methods become available, quite often too long with respect to the existing deadlines needed for urgent issues and modeling-based decisions, for example in ITER design and construction. This leads to the recommendation of keeping a good balance between urgently-needed scientific production, achievable only with existing numerical tools, and development of advanced numerical methods and codes for new HPC architectures. This calls consequently for the availability of a correspondingly broad range of computer architectures. Without such a balance, a memory bottleneck for Exascale computers would limit the progress in the simulation of nuclear fusion plasmas and astrophysics, like in a number of other large non-linear problems.

4.4 Biosciences

Life science, biochemistry, biology and medicine are vast areas of extreme societal importance.

4.4.1 *Biomolecular sciences*

Such applications fundamentally are at the basis of all these broad areas, ultimately rely on advancing the computational state-of-the-art applied in the domain. It is already possible to simulate systems with millions of atoms such as viral capsids, fusing vesicles, drug delivery systems or other multi-molecular assemblies. With the upcoming Exascale compute capability systems, this limit will be pushed-up, and molecular modeling will be able to tackle a much wider scale of bio- and non-bio systems. At the same time, some of the greatest challenges in the field are how to make predictions on longer timescales and enable high-throughput structural studies.

Understanding structure and motion in biomolecules is a fundamental challenge. MD simulations have become a universal tool to understand both how molecules interact and move, for instance how pumps transport molecules across cell membranes, how ion channels respond to pharmaceutical compounds, how virus capsids are assembled, or how complex organs such as the skin are formed from lipid building blocks. In addition to using standalone simulation engines (such as GROMACS developed by the BioExcel CoE), the field has made substantial progress on advanced techniques for devising workflows and ensemble simulations for studies of large, complex systems. Those latter approaches are extremely important for reaching Exascale performance, while integrating multiple tools with powerful core applications dramatically improves the productivity of researchers and shorten the time-to-solution.

Integrative structural biology is a rapidly evolving field. Dissecting the processes at atomic detail is invaluable, but there is no single technique that can provide all the answers. Researchers increasingly rely on coupling computational modeling with experimental data. E.g. one of the most popular applications for integrative modeling, HADDOCK (see below) is supported by BioExcel CoE.

Modeling reactions, catalysis, photochemistry, and detailed binding are relatively little studied in the life science domain. As the amount of structural and biological data has grown, and given the complexity of the molecules at stake, the facing challenges cannot be handled neither with classical methods nor with pure quantum mechanics (QM). An attractive solution to this problem is to employ hybrid and multi-scale methods combining QM and molecular mechanics (MM).

Drug discovery and optimization is becoming ever more crucial for the pharmaceutical industries. Therapeutically used peptides and proteins are engineered to increase their affinity towards the corresponding binding partners. Routine use of integrative modeling via docking is a powerful approach to model the interaction of drugs and compounds with their targets and highly accurate molecular dynamics-based free energy calculations for large scale mutation scans in proteins, DNA, and small organic molecules to improve their binding affinity and stability. Biomarkers are often used in different stages of the drug development processes or as diagnostics. Since the biomarker-design process is similar to drug development, molecular simulation methods are an integral part of this process.

GROMACS is designed mainly for simulations of proteins, lipids, and nucleic acids. Its future development will focus on a number of key user-requested features, in particular adding multiple time stepping, driving simulations with external data, facilitating generation of input data, and better free energy tools. Performance optimization for additional new platforms will also continue together with efforts targeting significantly improved scaling by using ensemble parallelism, as well as co-design projects to port the code to new accelerator platforms (see Figure 24).

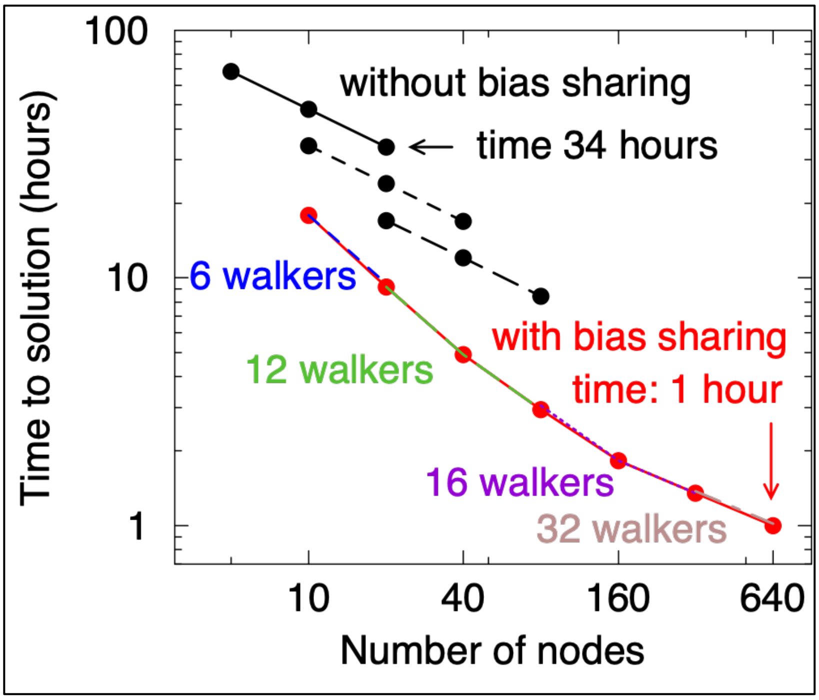

Figure 24 Ensemble simulations in GROMACS provide close to 100-fold reduction in time-to-solution for opening 20 base pairs on Piz Daint CPUs + P100 GPUs at Centro Svizzero di Calcolo Scientifico, Switzerland (courtesy of R. Apostolov)

Compatible with the GROMACS code, the PMX package is used as an automated procedure to generate hybrid structures and topologies for the amino acid mutations in all commonly used force fields. The development of the PMX package will follow two main directions. Firstly, for successful future usability, PMX will be fully rewritten in Python3. Secondly, it will incorporate frequently user-requested features, the addition of new mutation libraries and creation of workflows for free energy calculation setup and analysis.

HADDOCK is a computational docking program that follow a data-driven strategy. It has been applied successfully and reliably to a plethora of biomolecular systems, and more than 60 structures solved via HADDOCK docking have been deposited in the protein data bank. HADDOCK's development focuses on the continuous improvement of the web user's interface, and the design of a new modular version, HADDOCK3. The web server is being concomitantly developed alongside the local HADDOCK version, so that new features are readily available with the user-friendly interface. The modular design of HADDOCK3 will provide flexibility for workflow creation and a straightforward framework to integrate new tools as well as more efficient execution.

Upcoming Exascale compute resources give a tremendous opportunity to explore in ever greater detail biochemical reactions such as enzyme function. In order to take full advantage of those capabilities, it is necessary to improve the interfaces between classical MM/MD engines such as GROMACS and QM codes such as CP2K, an open-source, highly parallelized density functional theory (DFT) engine with large user base (see Figure 25). Coupling of GROMACS with CP2K will allow multi-scale simulations, where part of system undergoing a chemical reaction is treated at the QM level (e.g. with DFT), while the rest, and much larger, part of the system is described with a classical molecular mechanics force field. The updated roadmap reflects the transition to using CP2K as a QM engine due to its larger user-base and advanced compute capabilities.

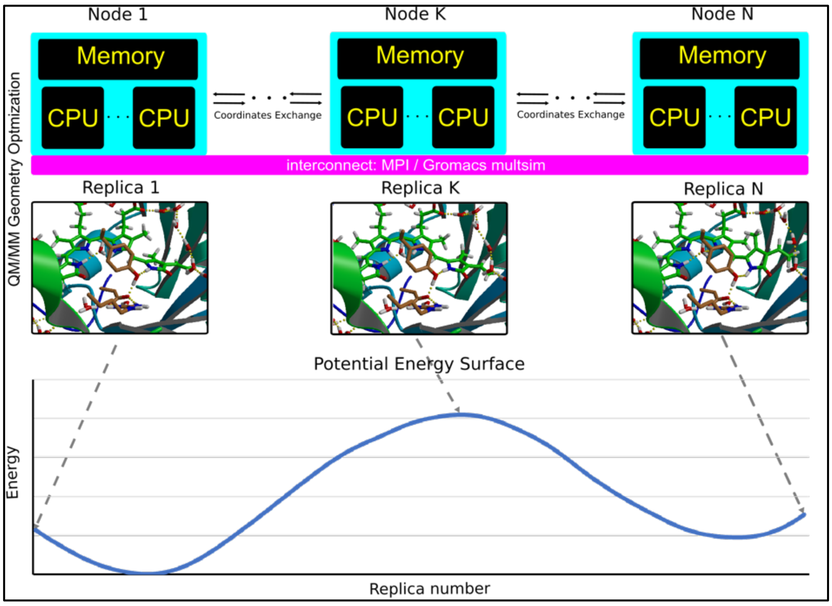

Figure 25 High-throughput hybrid-QM/MM modeling of biomolecular systems. A related, previous, implementation using the same distributed computing model achieved linear scaling until 40,000 cores on an older Cray CX40. Much higher scalability is expected with the new codebases and redesigned implementation (courtesy of R. Apostolov).

Still in the field of molecular processes, but turning now to vegetation problems, very recent new advances concern the simulation of processes involved in the capture and transformation of photon energy into chemical energy by the purple bacteria, in a way that is similar, although different in some respect, to the photosynthetic plant activity (see Figure 26). The purple bacteria is relatively simple but very efficient to produce its metabolic energy from the only few photons which reach it at the bottom of the lakes where it develops. The processes at stake are nevertheless very complex, as their simulation did require the consideration of 136 millions of atoms, three order of magnitude more than in standard simulations, with a time step of 1 msec. This requires multi PetaFlop/s computers.

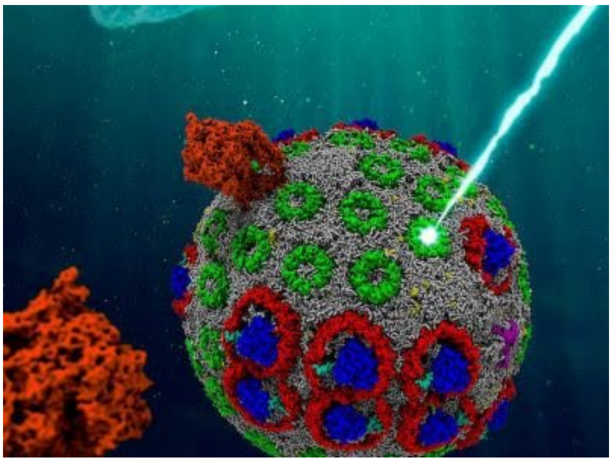

Figure 26 Atomistic model of the chromatophore of the purple bacteria (source [17])

4.4.2 *At the larger scale, a growing number of issues concerns human health*

Human-based computer models and simulations are a fundamental asset of biomedical research. They augment experimental and clinical research through enabling detailed mechanistic and systematic investigations. Owing to a large body of research across biomedicine, their credibility has expanded beyond academia, with vigorous activity also in regulatory and industrial settings. Thus, human in silico trials are now becoming a central paradigm, for example, in the development of medical therapies.

Biomedical research is particularly sensitive to computational tools and their efficiency on large systems because, quoting PRACE, "Life science is one of the fastest growing users of HPC both in Europe and worldwide, with a wide range of uses from chemistry, bioinformatics, and structural biology to diagnosis and treatments in clinical settings." The topics at stake may concern, among others, neuro-musculoskeletal treatments, in silico heart assays on HPC, simulation of blood flow through a stent (or other flow diverting device) inserted in a patient's brain, or in silico drug trials in populations of human cardiac cell. They all require HPC modeling for predicting safe and efficient medical interventions and treatments. Other examples are schematically represented in Figure 27.

This wide range of scales extends from the organization of the health-care system to the molecular level, where the deepest roots of the genome reside. Compared to other complex multiscale systems, in biomedicine the lowest scales are largely responsible for emergent properties at the highest scales. Conversely, upper scales feedback to lower ones, creating a non-linear coupling loop. This does not mean that a complete system's understanding is required to simulate all scales, but it shows that neglecting the effects of any single scale upon the others may lead to erroneous predictions. Considering that the final target is human healthcare, it is imperative to arrive at the correct prediction.

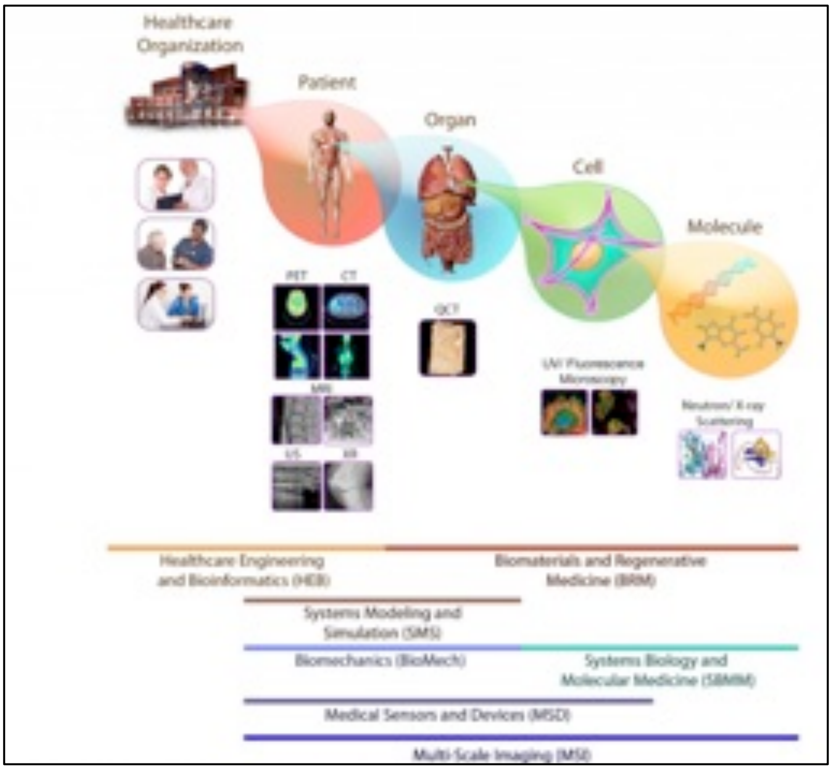

Figure 27 Intervening scales and organizational levels in computational biomedicine, from health care at the top to cell and molecular scales at the bottom (source [18])

Cardiovascular diseases, having a direct effect on the function of the heart itself, be it on the electrophysiology, mechanics or blood flow, and disorders in the arteries, be it aneurisms in the abdominal aorta or in intracranial arteries, or stenosis in carotid or coronary arteries, accounts for half of sudden deaths in Europe, calling for urgent and challenging improvements in patient risk stratification and prediction of clinical intervention. Human cardiac physiology is one of the most advanced areas in physiological modeling and simulation with many areas of medicine relying on a molecular understanding of the underlying human biology. Indeed, the pharmaceutical industry's success has been largely underpinned by such knowledge. Its business model is being seriously challenged today, with rapidly increasing sums of money invested in an attempt to maintain an acceptable pipeline of patentable products. However, the central premise of drug production, namely that one can hope to produce "blockbuster" one-size-fits-all drugs for the entire global population, has proven impractical; most drugs that have been developed for specific disease treatments only apply to subsets of the population. Instead, the industry now needs to think in terms of multiple drugs, which address any specific disease case, using stratification (based primarily on gene sequencing) as a first step along the way to ultimate personalized drug selection and treatment. Due to advances in gene sequencing, the basic patient specific data are now available, allowing the development of personalized drug treatment to be initiated. Next-generation of drug discovery, relying on computational predictive models with the ability to test millions of compounds in silico, will provide accurate and precise results and reduce the number of experimental tests needed on the design process.

Biomedical systems are, however, difficult to simulate. Simulating physiology requires complex computational models that, to be accurate, require very fine space and time discretization. The models are usually non-linear, increasing the need for sophisticated solver strategies, which are always computationally expensive. Compounding this, space and time scales range very widely in these systems, each scale being solved using a different model, increasing the complexity through a strong multi-scale and multi-physics character. The choice of medical intervention to save a person's life requires the bringing together of substantial quantities of data together with the performance of multi-dimensional simulations before the event in question occurs. Such forms of calculation are among the most demanding, as they need to be done rapidly, accurately, precisely and reliably. Moreover, they must include the quantification of the uncertainties associated with them. All these systems are multi-scale in nature, as their accuracy and reliability depend on the correct representation of processes taking place on several length and time scales. Only now, when approaching the Exascale era, can one expect to be able to tackle such problems effectively and, eventually, in a routine manner.

Containerization is explored in several scenarios. At the molecular level, an example concerns Acellera, a CompBioMed CoE consortium partner. Its "In Silico Binding Analysis" service makes extensive use of Amazon's EC2 cloud computing platform, for fulfilling the need to meet the pressing deadlines of customers. At the other end Alya, the multi-physics code developed by Barcelona Supercomputing Center, Spain, solves biomedical applications at the cell, tissue, and organ level (see Figure 28).

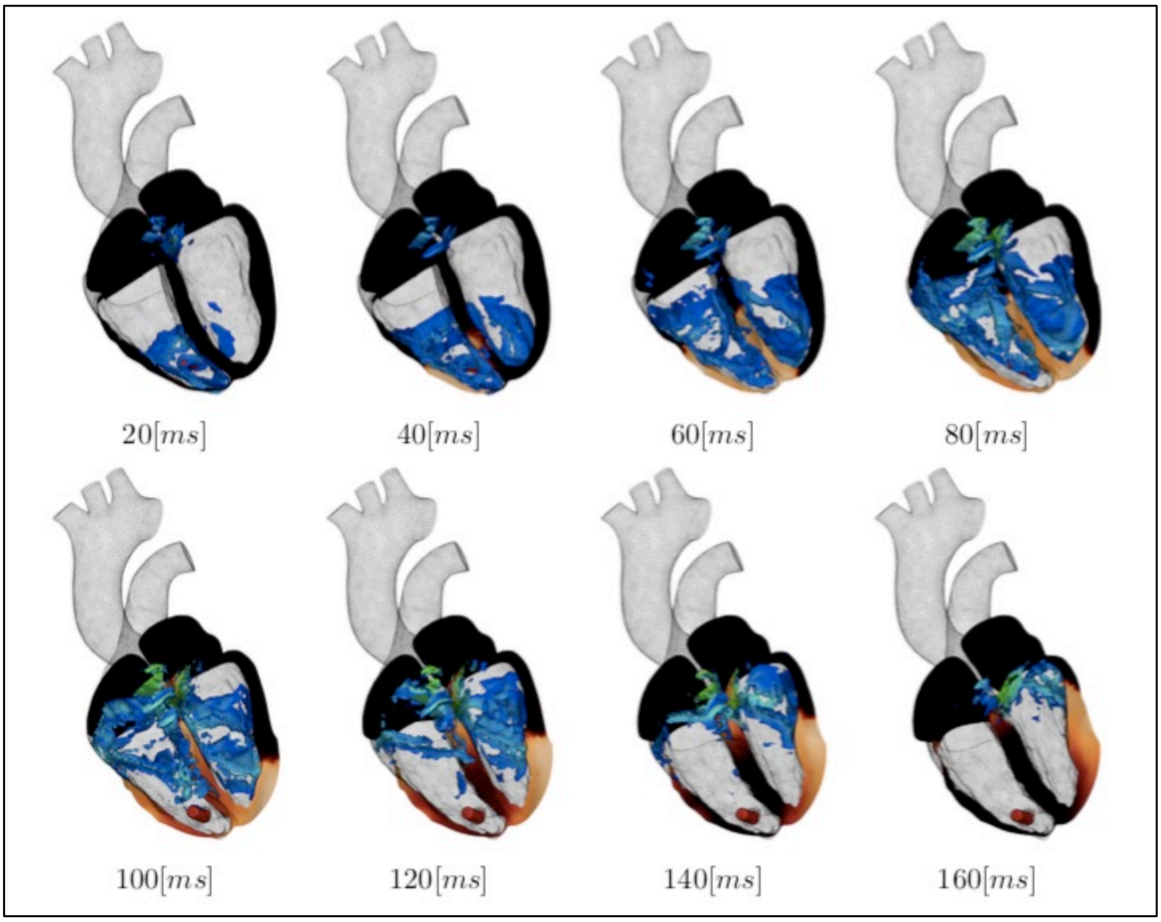

Figure 28 Fully-coupled fluid-electro-mechanical simulation of a heart. The sequence shows the systole process of a third degree atrio-ventricular block and the action of a trans-catheter intra-ventricular pacemaker (the small red cylinder close to the heart apex). Tissue is colored by electrical activity and blood shows the so-called Q-criterion, which depicts blood flow vortices evolution (courtesy of A. Santiago).

A powerful way to attack multi-scale simulations deployed in HPC cloud environments is with data analysis techniques of all kinds, including genetic algorithms, ML, and AI. In their most complex form, these techniques can combine high- and low-resolution simulations with experiments of all kinds, to obtain the correct input parameters, to assess sensitivities of inputs or to create surrogate predictive models. Examples concern, e.g., simulation of aggregation of blood platelets, for which five parameters, related to the deposition process, can be tuned by using comparisons between in vitro experiments and simulations. Hybrid methods, where HPC simulations are combined with ML, are developing quite rapidly and are expected to be able to solve both the accuracy and time-to-prediction problem by learning predictive models using expensive simulation data, a kind of post-processing method (see Figure 29). This would nevertheless present more difficulties, such as I/O, storage and data analysis. The synergies between classical, quantum simulations and ML methods, such as artificial neural networks, have the potential to drastically reshape the way predictions are made in computational structural biology and drug discovery.

One should finally underline that visualization is still of great importance for all the above problems at stake, as simulation results are often best represented in the more familiar time-dependent three-dimensional form.

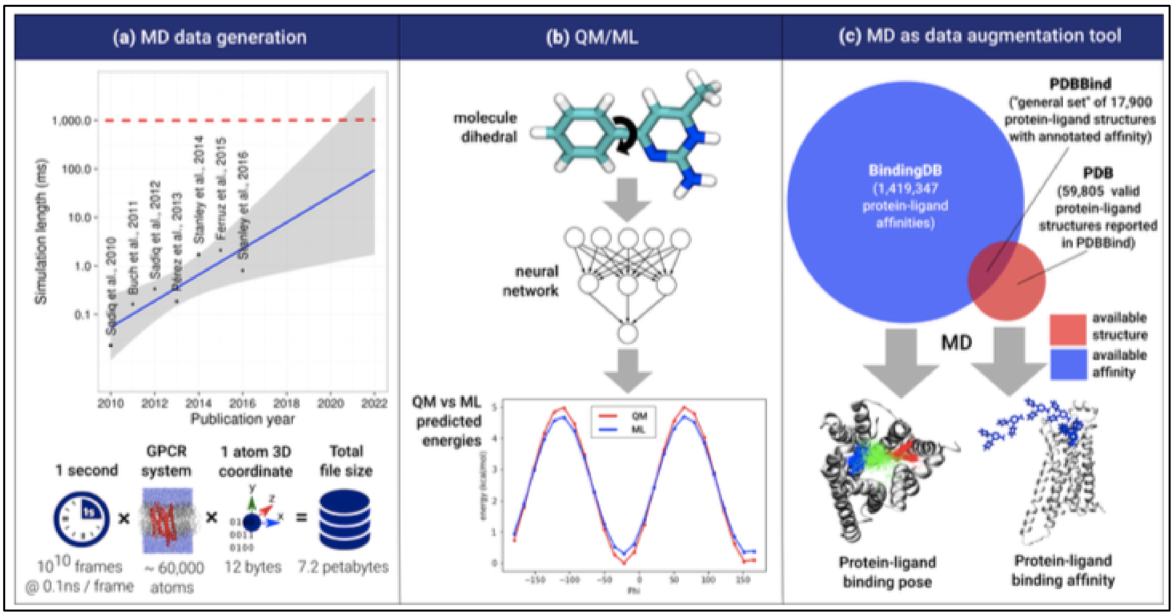

Figure 29 Overview of a combined MD/QM simulation and ML approach From left to right (source [19]): (a) MD data generation will produce output files size of several petabytes; (b) ML may then replace QM to predict dihedral energies given a neural network trained with QM simulations; (c) data augmentation by MD will then improve the description of some processes (protein-ligand binding poses for a set of protein-ligand pairs with unknown binding mode; binding affinity data for a set of resolved protein-ligand complex structures of unknown affinities).

4.3.3 The human brain

Understanding how the human brain is structured and how it functions remains one of the major challenges facing neurosciences. Medical imaging, which has experienced a tremendous boom since the mid-1970s, has greatly contributed to the improvement of knowledge in this field by allowing in vivo and non-invasive exploration, unlike exploration methods based essentially on ex vivo studies of anatomical parts developed by the major neuroanatomists of the early twentieth century. And yet, medical imaging still does not map the cytoarchitecture (understanding which zone of the cortex is responsible for vision, motricity, audition, emotion, cognition, …) of the cerebral cortex. *This project performed in the field of the FET Human Brain Project by researchers of CEA's Neurospin team using the Jean Zay converged HPC system of GENCI at IDRIS, aims to develop new approaches combining numerical simulations and ML to eventually decode in vivo the cellular organization of the cerebral cortex from magnetic resonance imaging data coming from two state-of-the-art 3 T and 7 T whole body Magnetic Resonance Imaging (MRI) scans and three preclinical MRI scans at 7 T, 11.7 T, and 17 T.*

The Neurospin team developed a simulation code called MEDUSA which will be used in the context of cerebral cortex studies for allowing to perform successively the three simulation steps: 1) creation of a virtual fabric 2) Monte-Carlo simulation of the diffusion process and 3) simulation of the MRI signal weighted in diffusion.

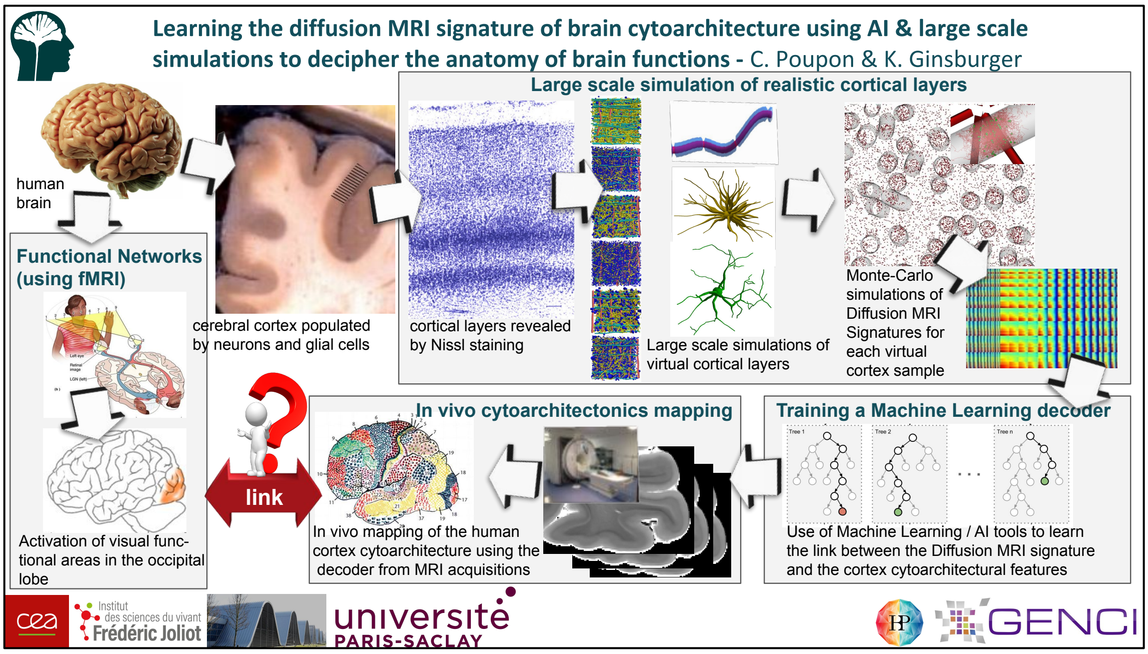

Figure 30 Overview of brain cytoarchitecture simulation/learning (courtesy of C. Poupon)

At the end of the large-scale simulations to be conducted on the supercomputer provided, the Neurospin microstruture team will have a dictionary consisting of $0.5x10^6$ elements, each element corresponding to a pair of vectors, the first vector containing the thirty or so parameters characteristic of the cell geometry of the virtual tissue corresponding to the element in question and the second vector containing its specific MRI signature.

This dictionary will then be used to train in a supervised way a ML algorithm of the ExtraTrees type available in the scikit-learn toolkit [20].

Once trained, the ExtraTrees will be used to decode the cytoarchitecture of the cerebral cortex from MRI signatures that will be measured in vivo in humans on a standard clinical MRI device (3 T or 7 T MRI equipping the Neurospin department), and thus produce maps corresponding to the thirty or so characteristic parameters of individual scale cytoarchitecture that are currently inaccessible in vivo.

4.5 Combustion

Combustion provides more than 90% of the world energy and, even if renewable energies are developed at a high speed, combustion also keeps growing because of the growing world energy need. Optimizing combustion systems is therefore a critical issue and relies more and more on simulations which can be performed only on HPC systems because of the complexity of the phenomena to compute.

The simulation of turbulent reacting flows has been a main application field for HPC over the last 20 years and has led to multiple transformative results as observed in multiple European, such as Excellerat, or US projects (ASCI for example). The field where HPC tools are the most advanced is probably aerospace: today, for example, many companies developing engines for propulsion rely on HPC codes, which are used daily by design engineers and run routinely on 5,000 cores. This production use of HPC resources by industry has been made possible by aggressive academic developments and demonstrations of the corresponding codes over the largest machines (PRACE or Innovative and Novel Computational Impact on Theory and Experiment, INCITE, type) up to half-a-million cores. These computations correspond to real combustion systems and include the full complexity of the physics which is required to describe combustion: turbulence, chemical kinetics, radiation, heat transfer to walls, acoustics. The corresponding codes are developed in large teams, mostly through laboratory associations. These codes are also used in many EU projects through Marie Curie or European Research Council (ERC) programs.

The scientific difficulty associated to the simulation of turbulent reacting flow is the well-known multi-scale multi-physics nature of the problem. This is a general statement, true for many fields but critical in combustion: in many combustors, the smallest chemical time scale is of the order of a nanosecond while the computation must last a few seconds to tackle many problems, such as ignition. Similarly, the spatial smallest scales can be of the order or a micrometer, while the whole engine may be several meters long. This is a challenge which cannot be fully addressed by brute force methods even today. Solving it requires both the largest hardware available and massive software developments. The simulation codes used for reacting flows are produced by large teams and incorporate multiple levels of numerical analysis and physics. One million lines of code is a typical order of magnitude for such tools. Obviously, adapting these codes to rapidly changing architectures is a difficulty, so that successful combustion codes are developed today by composite teams gathering physics experts but also an important fraction of HPC experts. This co-design structure is well known in the HPC field but it is critical for combustion where the models evolve very fast: for example, the need to predict soot and NO_x emissions or the interest in hydrogen combustion (see European projects such as Helmeth or the ERC advanced project SCIROCCO) in the last years has led to major modifications in combustion codes physics and makes previous codes obsolete. Making these codes also compatible with new hardware at the same time is difficult.

Another evolution of combustion codes is linked to the integration of these codes with simulation tools for surrounding systems. For example, in a gas turbine engine, combustion codes are now applied to the combustion chamber but must also be coupled to the compressor and the turbine simulations. In terms of HPC, this leads to problems similar to ocean/atmosphere coupling for climate simulations: large codes must be coupled and run efficiently at the same time in a strongly coupled mode. For example, Figure 31 shows a computation where a combustion chamber computation is coupled to the simulation of the first stage of the turbine (which is rotating): at each time step, the combustion code outlet must be coupled to the turbine inlet. This requires massive message passing and interpolations from one grid to another. In certain cases, the combustion and the turbine codes might actually be running on different machines and even different architectures (a CPU system for the combustion code and a GPU one for the turbine for example).

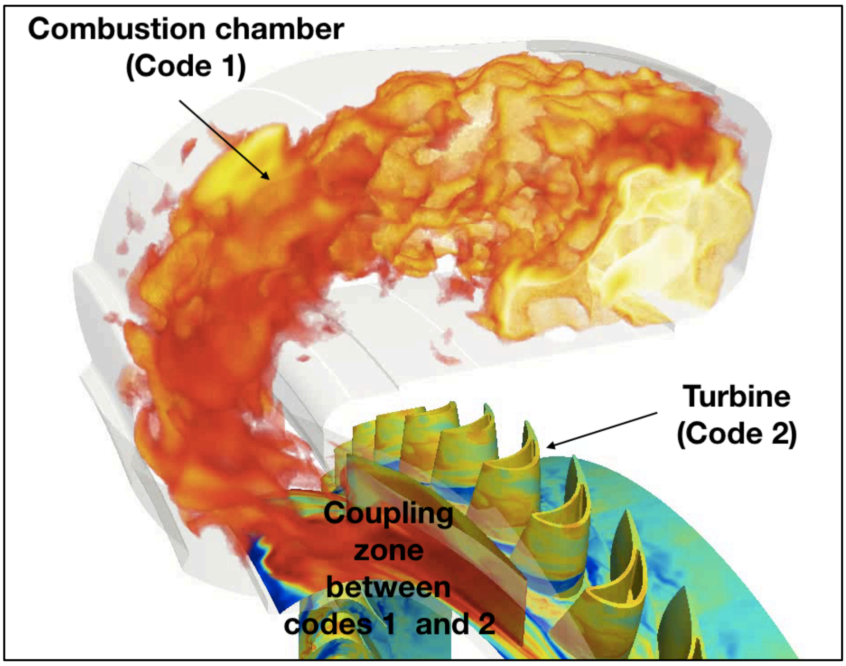

Figure 31 Example of coupled computation on HPC architectures: two codes are coupled to compute the combustion chamber on one side and the turbine on the other side (source CERFACS)

Combustion is another fascinating field of application for investigations of the ML potential in the domain of partial differential equations. The Navier-Stokes equations, which control the motion of turbulent reacting flows, require closure terms when they are averaged. Such closure schemes are statistical models which can be tackled with ML techniques. This leads to a reformulation of LES ("Large-Eddy Simulation") techniques for reacting flows, which offers great potential. This also imply new research into CPU/GPU hybrid computing, since most computational fluid dynamics (CFD) codes run today on CPUs while most ML tools live on GPUs architectures. This is another field where combustion is at the forefront of HPC, ongoing research is done on the Jean Zay system at IDRIS by CERFACS for large-scale LES of a slot jet flame coupled with a neural network-based wrinkling model with successful first co-simulations up to 2816 CPUs and 256 GPUs (see Figure 32).

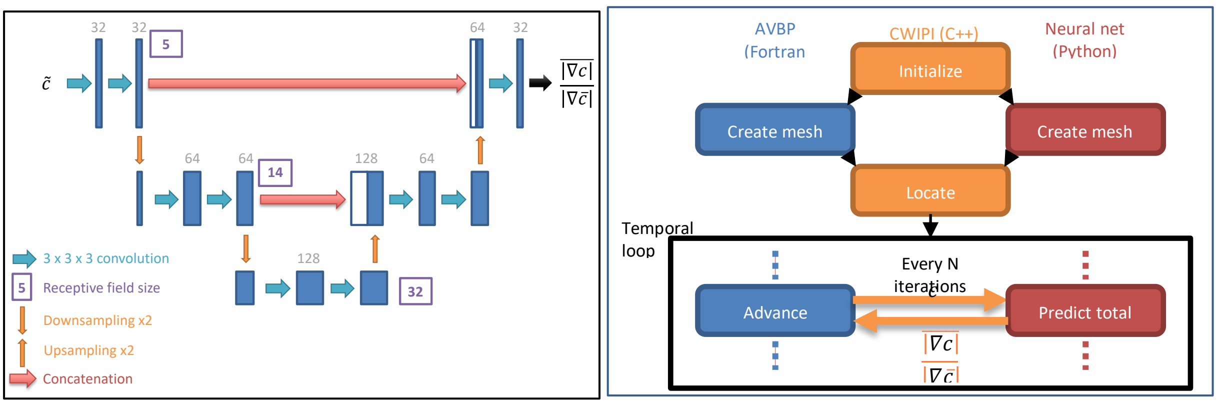

Figure 32 Coupling of AVBP and a neural network for computing the total surface of the flame front (source CERFACS)

To conclude, the European numerical combustion community is extremely strong and well connected to the HPC world. Multiple combustion codes are used to develop and benchmark new architectures because they are identified as good tools to stress machines while also attracting industrial users towards HPC systems. It is very important for the EU to propose adequate tools to extend these collaborations: combustion codes require very high computing powers and also offer direct connections with industry and important societal applications in the energy domain.

4.6 Chemistry and material sciences

Chemistry and materials science will remain one of the largest users of computing, with industry increasingly relying on simulation to design, for example, catalysts, lubricants, polymers, liquid crystals, and also materials for solar cells and batteries. Electronic structure-based methods and molecular dynamics will access systems, properties and processes of increased complexity, and towards extreme accuracy. These are being complemented both with multi-scale models and data-driven approaches using high-throughput and DL to predict properties of materials and accelerate discovery. This will enable researchers to fulfil the grand challenge of designing and manufacturing all aspects of a new material from scratch, which will usher in a new era of targeted manufacturing.

As stated into the PRACE Scientific Case the identification, development, and exploitation of new classes of materials is also key for European industry and competitiveness. To illustrate a handful of examples:

- The use of ab-initio, Born-Oppenheimer MD to understand the catalysis of petroleum cracking and other important reactions on transition metal surfaces (J. Matthey, BASF).
- The use of mesoscale simulation methods in the prediction of phase diagrams of multicomponent surfactant and polymer mixtures heavily marketed by companies such as Unilever and Procter & Gamble. This will require rapid chemical potential calculation by particle insertion methods.
- The prediction and modification of crystal habit by simulation, in particular attachment energies and entropies. This is used by major pharmaceutical companies such as Pfizer in the prediction of drug solubility and delivery.
- Calculating accurate relative free energy changes as drugs are transported from solution to the active site of proteins, as well as their binding affinity. The relative binding constant of two drugs can be calculated from the free energy cycle as one drug is transformed into another in both environments, and with Exascale computational resources these methods will be able to use even more accurate models and forcefields to replace the current more approximate docking methods.
- The prediction of lubrication and friction coefficients between two solid surfaces including ionic liquids in factory machines at extreme loads (>1 GPa), physiological lubrication of joints by polymers at loads of 7 MPa, and tertiary oil recovery (British Petroleum).
- The growth of gold nanowire using kinetic Monte Carlo calculations combined with the computational fluid dynamics of spray jets (Merck solutions).
- Mesoscale simulation of the effect of electric fields on director reorientation and orientational phase transitions in liquid crystal phases in the production of new flat screens (M. Global).
- Simulation of polymer melt mixtures to create channel structures for use in membranes for osmotic water purification (Fuji Films).
- The use of quantitative structure-activity relationship and ML techniques for the design of novel antimicrobial peptides for cleaning solutions (Unilever).
- The prediction of mechanical properties of co-block polymer melts to create new and efficient tyre composites (MD and mesoscale dynamics) for Michelin.

While data has always played a pivotal role in material design and discovery, its importance as a major driver for innovation is increasingly appreciated, as evidenced by the many current efforts worldwide in material data such as the NOMAD Laboratory and MaX CoE, independent initiatives such as ioChemDB, and the US DoE HPC for materials initiative.

The MaX CoE is enabling materials modeling, simulations, discovery and design at the frontiers of the current and future HPC, HTC, and HPDA technologies. MaX is working on scaling out to new pre-Exascale converged (HPC and AI) architectures a set of leading-edge European applications including complementary open-source codes: Quantum ESPRESSO, SIESTA, YAMBO, FLEUR, CP2k, and BigDFT. It contributes to the development of AIIDA: a Python-based materials informatics cloud based framework to manage, store, share, and disseminate the workload of high-throughput computational efforts, while providing an ecosystem for materials simulations where codes are automatically optimized on the relevant hardware platforms, and complex scientific end-to-end workflows involving different codes and datasets can be seamlessly implemented and shared.

The NOMAD CoE develops materials encyclopedia and big data analytics tools for materials science and engineering. This will be reinforced by advanced graphics and animation tools.

New ML technologies are providing substantial added value by enabling the extraction of more information from these new data projects and from existing databases – the community is increasingly generating knowledge directly from the databases. Furthermore, the transferability of different concepts through suitable metadata enhances (i) the identification of materials and compounds with desired properties and (ii) the transferability and conversion of the gained molecular understanding into mesoscopic models and ultimately devices, following the concepts developed in the European Materials Modelling Council and their "materials MOdelling DAta". The improvement of codes and their integration plays an important role in this progress as has been recognized through the Virtual Materials Market Place, an initiative that demonstrates materials research is predictive enough for wide industrial application, and that there is industrial demand for materials design using both HPC and AI.

The next generation of researchers will need to master a wide range of methodologies, ranging from identification of effective potentials starting from ab initio studies to the behavior and response of systems with complex composition over long length and time scales. A concerted effort is needed to educate a new generation of computational material scientists and direct resources towards software development in addition to the ongoing investment in computational infrastructure.

4.7 Social Sciences

Social sciences (or Humanities) are very wide with multiple definitions depending on the countries. By example, stating in Wikipedia "Social science is the branch of science devoted to the study of human societies and the relationships among individuals within those societies. The term was formerly used to refer to the field of sociology, the original 'science of society', established in the 19th century. In addition to sociology, it now encompasses a wide array of academic disciplines, including anthropology, archaeology, architecture, economics, human geography, linguistics, media studies, musicology, political science, psychology, and social history". In recent years social sciences are also encompassing new domains like finance/insurance, risk management, … and have been more and more mixed with others traditional sciences leading for example to neurosciences where both psychology and biology allow the study of perception, cognition, attention, emotion, intelligence, subjective experiences, motivation, brain functioning, and personality.

Like “traditional or hard sciences”, social sciences have in common to deal with a recent explosion of the amount of data to process, cross and correlate due to the availability of more and more of heterogeneous data sources (only think about social media, see Figure 33), the growing number legal and ethics frameworks around the collection and the handling of data and the associated social and economic stakes to process, value and expose data.

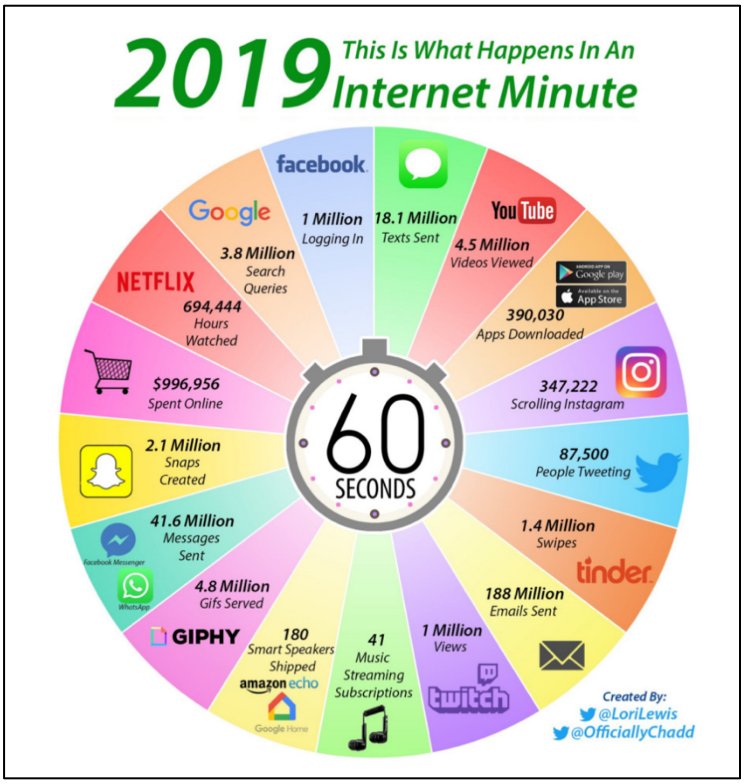

Figure 33 Yearly edition for 2019 of the One Internet Minute (source [21])

In order to ingest and process such amount of data, HPC/HTC became during the 1990s mandatory tools paving the path to HPDA in order to browse huge datasets and databases. *Since few years AI, or more precisely ML, are now feasible due to (i) the existence of a large amount of data that represents a high-potential source of valuable insights and (ii) advances in the underlying hardware and software ecosystem, which provide the computational performance to train the models associated to these ML systems.*

As a first example of the HPC for social sciences at the European level, the GSS CoE, working on the development of an HPC-based framework based on DSL to generate customized synthetic populations for GSS applications, worked on three concrete pilot studies:

- the modeling of smoking habits and tobacco epidemics to create a synthetic population by integrating large and heterogeneous data sources that will describe the prevalence of health habits in Europe and explore their expected trends. This goal will be achieved by integrating high-resolution demographic information, official population statistics, and dynamic models of social contagion. Taking advantage of HPC and developing novel algorithms to simulate the dynamics of the tobacco epidemic, the resulting system will be a powerful tool in the hands of policymakers to evaluate the impact of health programs and to increase their efficiency;
- evolution of the global car fleet and its emissions taking into account mobility factors/profile of people, congestion of infrastructures, pollution, noise, … and perform

optimization and uncertainties quantification studies on in order to answer to decision makers question like “What are possible effects of different decisions on the deployment of charging infrastructure on the uptake of electric vehicles in Germany up to 2035?”;
- and the two-way relation between transport infrastructure decisions (and associated positive impacts such public-transports development and negative as congestion and pollution) and price mechanisms, particularly concerning real-estate applied to the Paris area.

This work has been renewed and extended (in terms of partners involved) recently into the Hidalgo CoE developing novel methods, algorithms and software for HPC and HPDA to accurately model and simulate the complex processes, which arise in connection with major global challenges including three new pilots in:

- studying impacts of migration of people using scalable agent-based modeling taking into account precipitation/climate data, conflicts in the world, telecommunication data, … to enable simulations on a large scale to accurately forecast where displaced people coming from various conflict regions of the world, will eventually arrive to find safety;
- social networks, understanding and modeling the significant influence on social and economic behavior of their use, identify false/malicious messages, which intend to change the behavior of a substantial number of users, develop countermeasures on the algorithmic level in order to prevent the spread of such false messages on a large scale and finally develop a highly scalable simulation framework for such stochastic processes in real-world networks, in order to be able to analyze and predict the impact of these processes on the society;
- and urban air pollution pilot by developing an HPC framework for simulating the air flow in cities by taking into account real 3D geographical information of the city, applying highly accurate CFD simulation on a highly resolved mesh (1-2 m resolution at street level) and using weather forecasts and reanalysis data as boundary conditions.

In the field of scalable agent-based modeling research groups like at Barcelona Supercomputing Center, Spain [22], work on exploring new HPC solutions for social modeling such as data storage, load balancing, optimization of fine-grain simulations, and the use of DSL, besides developing a generic framework for simulation agent-based social dynamics to study the implications of this methodology in application areas such as tuberculosis epidemics or territorial urban planning (see Figure 34).

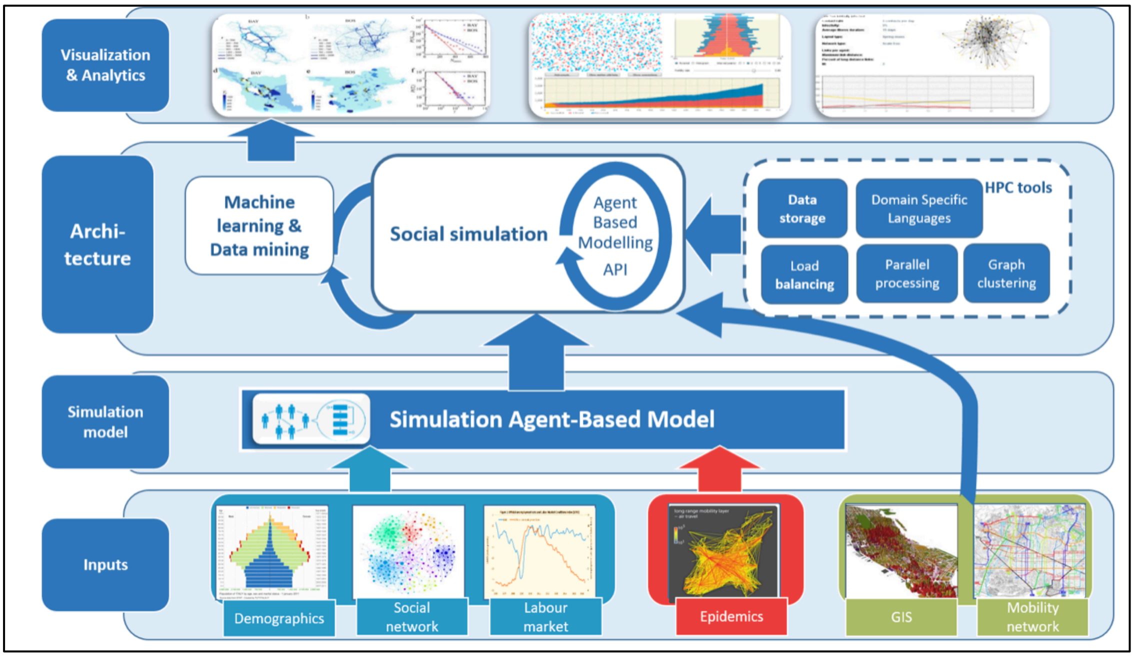

Figure 34 Global framework of scalable social sciences framework for pandemic spreading (courtesy of Barcelona Supercomputing Center)

In the field of enabling the complex analysis of large-scale digital collections a British team worked to scope out how the best HPC facilities can be used to facilitate the needs of researchers in the humanities. They analyzed, using HPC-enabled data-processing frameworks, more than 60,000 digitized books covering fiction and non-fiction publications from the 17th, 18th, and 19th centuries (representing close to 224 TB of compressed ALTO XML files), identified major barriers that humanities researchers are facing (fragmentation of communities, resources and tools, lack of interoperability, complexity and incompleteness of heterogeneous cultural heritage datasets or lack of technical skills) and worked on large-scale analysis of two case studies:

- the history of medicine in the UK, focusing on exploring issues around the spread of diseases, and the research questions were how does the occurrence of diseases (like cholera, whooping cough, consumption, and measles) in published literature compare to known epidemics in the 19th century;
- the history of images, tracking the spread of new images thanks to the development of new printing technologies between 1750 and 1850. This enabled the team to observe the dominance of full page and very small images (<15% of the page) between the 1750s and 1810s, after which time, driven by novel deployment of woodcuts and lithographs in books, the range of figure sizes diversified.

In the field of analysis of social media, a research team in Grenoble (France) used tier-2 and recently tier-1 HPC facilities to analyze in almost real time trends and behavior of potential electors just before the vote. The team analyzed using k-means methods the twitter activity (tweets, retweets, use of specific hashtags like #EE2014 or #Europeennes2014) before the 2014 European elections in order to form group of users who tended to do the same actions during the campaign.

And finally, at the European level again, the DARIAH-EU project (aiming to deploy thematic services for social sciences), EUDAT, and EGI collaborated to propose the DARIAH Science Gateway over the EOSC-Hub initiative. The DARIAH gateway is a platform that provides access to various digital applications and services for the arts & humanities researchers. The applications made available via the DARIAH gateway are:

- simple Semantic Search Engine: a semantic search engine which allows researchers to search for content in more than 100 languages within the Sci-GaIa e-infrastructure knowledge base, one of the largest existing databases;
- parallel Semantic Search Engine: a parallelized version of the Semantic Search Engine that enables simultaneously search across multiple platforms;
- DBO@Cloud: a cloud-based repository made of a 100-years old collection of Bavarian dialects. The datasets are provided by the Austrian Academy of Science.

4.8 Engineering and industrial applications

Engineering applications will be among the first exploiting Exascale, not only in academia but also industry. In fact, the industrial engineering field is the field with the highest Exascale potential. The European engineering industry consists of 130,000 companies of diverse size. Overall, these companies employ over 10.3 million people, with high levels of qualifications and skills. Together they generate an annual output of around 1840 billion Euros and about 1/3 of all exports from the EU. The European engineering industry plays a key role in realizing the goal of increasing the industrial production value above 20% of the gross domestic product (GDP) by 2025. To achieve this aim and meet the challenges of the fourth wave of industrialization, it is essential to support European engineering companies in their use of HPC and HPDA, thus increasing European industrial competitiveness. As massively parallel HPC systems have developed over the past decade, the engineering sector has come to face quite specific modeling and simulation challenges.

The CoE EXCELLERAT is working on a set of key European reference applications (namely Nek5000, Alya, AVBP, Fluidity, FEniCS, Flucs), which were selected as key representatives of challenges faced by broad parts of the engineering domain (especially towards Exascale). These applications cover three main topics of engineering: automotive, combustion, and aeronautics. Within the frame of the project, EXCELLERAT (see Figure 35) will prove the applicability of the results to other HPC engineering applications than the six chosen. Work on these applications will start from the pre-analysis and then enter a phase of further evolution, by applying mechanisms for maintenance, optimization and scaling where needed and performing tests and developments on current and future architectures, namely the Exascale Demonstrators, Pre-Exascale and Exascale machines, which will be available in the frame of EuroHPC and beyond. This also includes the validation of codes and their quality assurance, always in close alignment with the respective stakeholders, coupled to a focus on co-design with hardware and software vendors, to get the maximum performance of codes. A focus will be also to enable new engineering design capabilities on the basis of data analytics for improved product design.

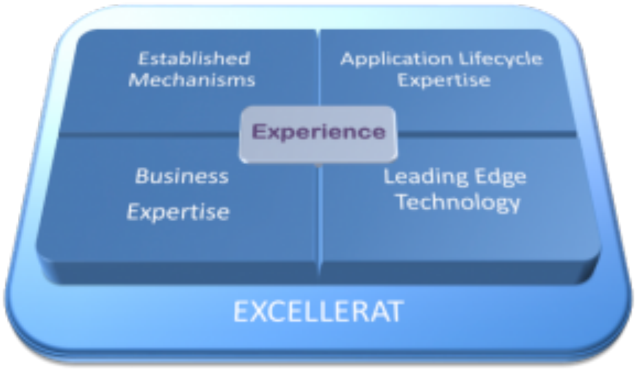

Figure 35 Base pillars of the CoE EXCELLERAT (source EXCELLERAT)

EXCELLERAT will enable enhanced engineering applications and support the community. The efficient use of HPC resources requires skillful personnel that is not only trained in how to efficiently use certain applications, but also have a good understanding of HPC hardware and technologies. Furthermore, the complexity involved in HPC requires great care in how the workflows to solve a given problem are set up, from choosing the right application and hardware mix to solving the questions of how and where to store and potentially visualize the resulting data. There is also a great need of transferring academic results to industry in an efficient and timely way to ensure leadership in the field. EXCELLERAT provides a comprehensive overview of HPC training course programs in the HPC centers comprising EXCELLERAT.

In Europe, **a lot of companies (mainly large groups but also sometimes Small and Medium Enterprises, SMEs) are using HPC in their daily business and some of them are even exhibiting Exacale roadmaps**.

In the field of Oil & Gas companies like TOTAL, BP, Shell, or ENI as well as contractors like Schlumberger or CGG Veritas are investing in large-scale own HPC facilities for high resolution seismic processing (and on a less extend reservoir simulation). TOTAL just deployed Pangea III an IBM OpenPOWER hybrid system of 26 PetaFlop/s and ENI just announced a major investment on an DELL EMC hybrid HPC system of 52 PetaFlop/s called HPC5, comprised of 1820 compute nodes, each with four NVIDIA V100 GPUs.

This system will be used for seismic processing (reverse time migration as well as first attempts of full wave equations modeling) as well as next generation reservoir modeling (4D seismic coupled to reservoir modeling, uncertainties quantification, multi-scale modeling from the pore to the reservoir scale) using a new software called ECHELON able to scale up to thousands of GPUs and used in capacity mode for uncertainty quantification of models, in order to optimize the production of oil fields during their lifetime.

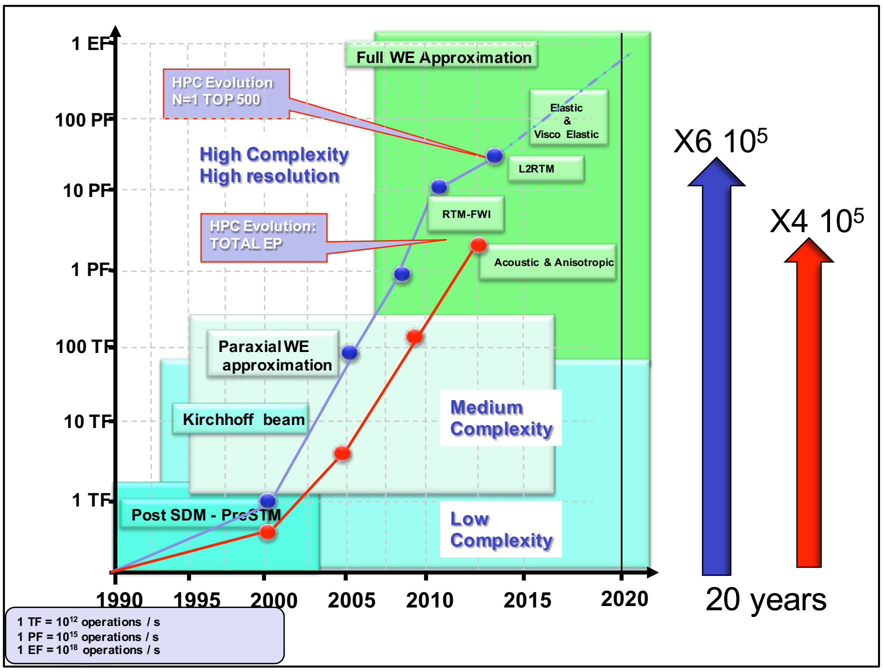

Figure 36 Seismic processing of TOTAL (courtesy of H. Calandra) as published inside previous EXDCI document

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 fast Fourier transformations (FFTs), Darcy's laws for flow modeling through a porous medium (used in basin and reservoir modeling) will be a real challenge. This will raise in a few years the question about rewriting HPC applications to be scalable beyond 300+ PetaFlop/s sustained (if possible) or starting now to think about using such applications in the future with quantum computing.

Beyond traditional computing some companies like TOTAL started to invest in quantum computing using ATOS quantum learning machine (QLM) simulators on top of existing quantum hardware in order to assess or develop new algorithms that could be used so for molecular/material chemistry simulation (developing new catalysis process, lubricants or next gen batteries) or optimization of fleets (tankers, electric cars, …) or energy grids and on a more long-term seismic processing and CFD.

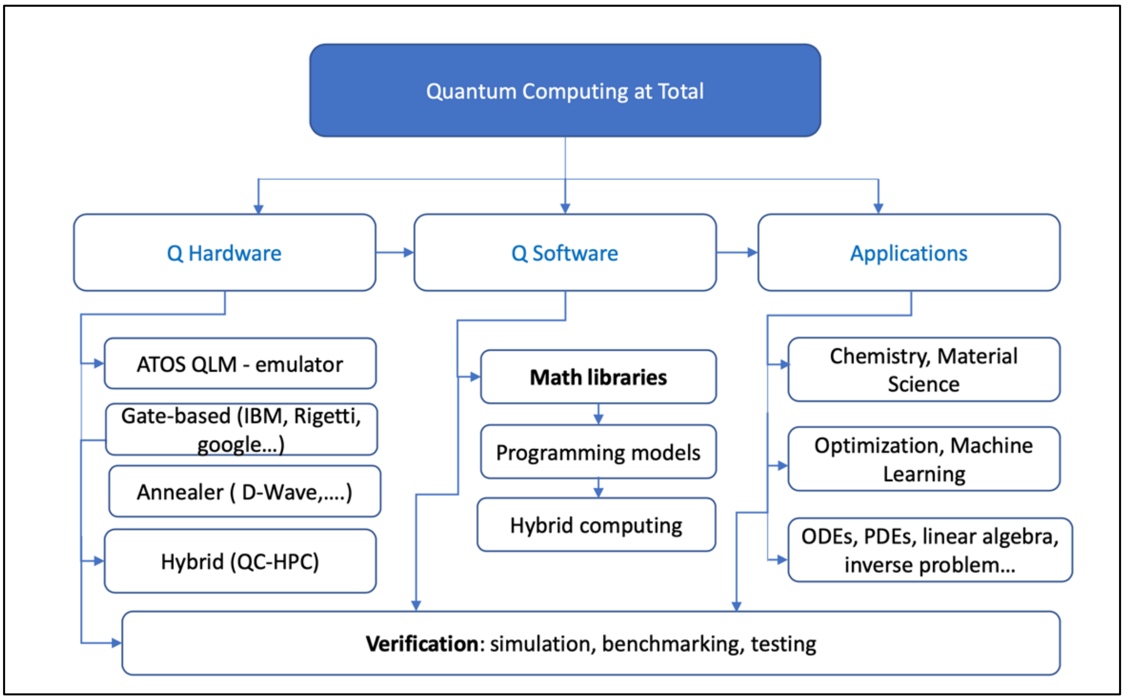

Figure 37 Early use of quantum computing at TOTAL (courtesy of H. Calandra)

In the field of energy, the objectives of using large-scale HPC (and now HPC and AI facilities) 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 cores, and the use of detailed simulations designed to verify and increase safety.

In the field of 3D unsteady CFD simulations, EDF R&D and IMFT optimized and scale out the NETPTUNE_CFD code applied to the simulation of fluidized-bed reactors in which combustion takes place at relatively low temperature have the main advantage to minimize combustion pollutants. Such reactors are used in many industrial applications, especially in solid treatment applications where energy may be supplied by direct combustion of fossil fuels inside the bed itself. Natural gas is the least polluting fossil fuel and, when burnt at low temperatures, it involves lower pollutant emissions, especially NO and NO2. Understanding and mastered natural gas combustion process in fluidized beds is thus of great interest with respect to environmental issues.

They are performing large scale multi-phasic simulations on up to eight billion cells industrial cases (a first ever in that field) on up to 61,000 cores of the Jean Zay system at IDRIS after initial successful proof of concept performed on CALMIP (Toulouse) regional facilities (see Figure 38).

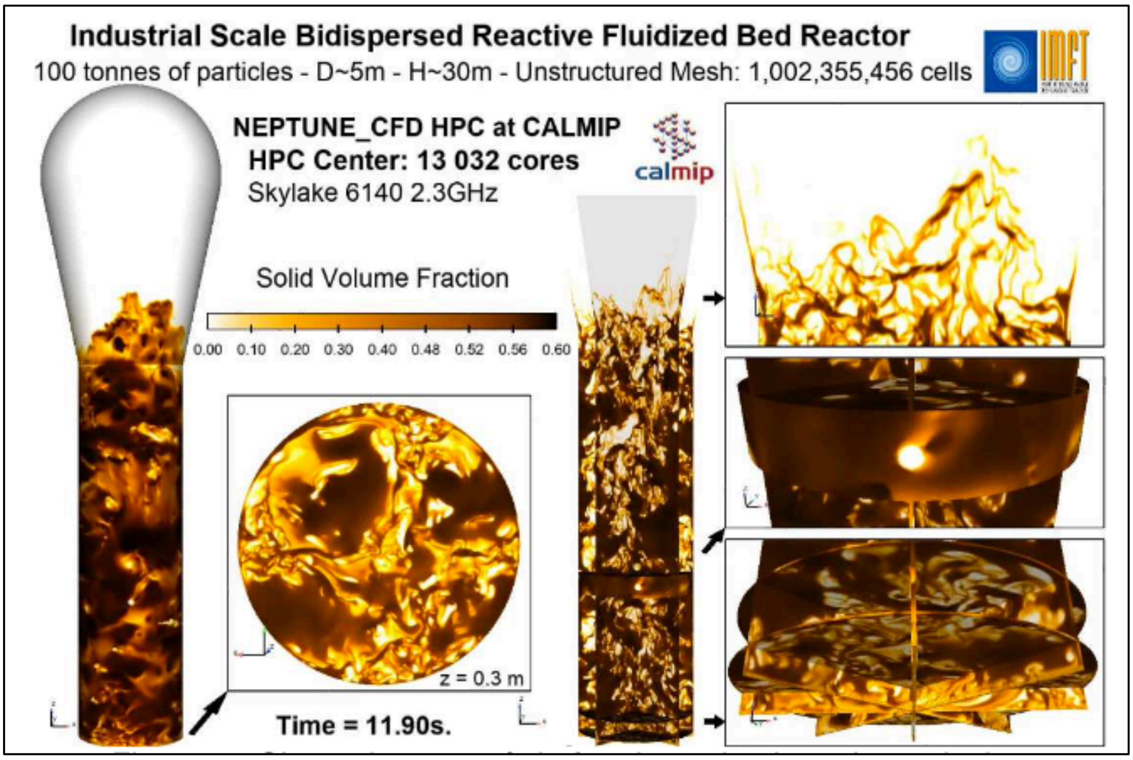

Figure 38 1 billion cells multiphasic industrial scale fluidized bed reactor CFD simulation (courtesy of H. Néau)

Beyond HPC, energy companies like EDF are also starting to look at quantum computing especially in the field of material chemistry. For integrate higher volume of renewables EDF UK plans to install large batteries on national scale power grids. But optimizing the investment and the life time operation of these storage systems is a problem whose complexity exceeds the capabilities of classical computing (see Figure 39).

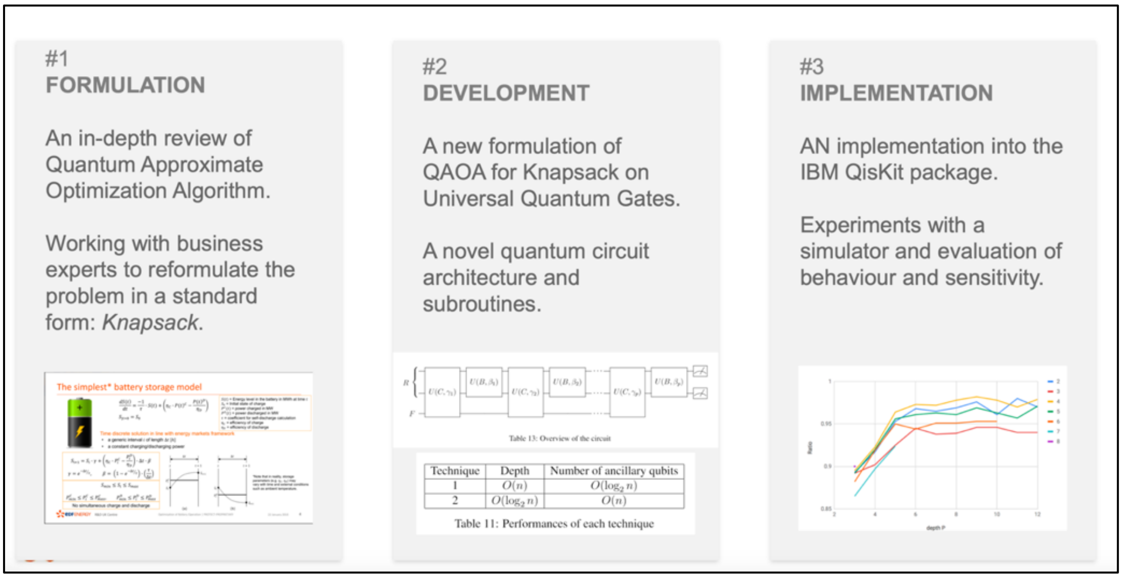

Figure 39 First quantum experiments of EDF R&D UK for optimal battery operations (courtesy of S. Tanguy)

5 Applications and co-design

The word "co-design" can be understood in different ways. Quite often it is used by vendors to refer to the simultaneous development of hardware (processors, memories, networks, architectures, …) and software (compilers, libraries, middleware, …). It is also used to refer to synergetic developments of compute infrastructure and scientific or industrial applications. We shall refer here to this later interpretation. The way applications will be able, from the computing point of view, to efficiently face their challenges will depend for a significant part upon the hardware and middleware they will be able to use in the coming years. As hardware and middleware will experience a number of important changes, co-design between applications and computing environment (machines and tools) is seen in many countries as a necessary and fruitful step. There are basically two ways to organize co-design: either by adapting the computing environment to the requested features of the main building bricks of a number of applications; or by constructing applications, or bricks necessary to a number of them, adapted to the computing environment.

The first approach is not largely developed. There are very few initiatives of this type, the most recent one being developed in Japan, to plan for the post-Riken computer Fugaku. The design of the system, processors and interconnect, is influenced by the applications needs, and is seen as a key to make it efficient and high-performance. More precisely nine target applications are considered, coming from very different fields of activity, i.e., personalized medicine, drug discovery, fundamental laws of the universe, innovative design and production processes, high-performance materials, innovative and clean energy systems, energy production, conversion and storage, meteorology and climatology, and earthquakes and tsunamis. Choices for the Fugaku system are then made so as to increase as much as possible the performance of these various applications. With respect to the preceding K-computer system (appeared in the second part of 2011) the **gain in performance for eight of the nine target applications ranges from 25 to 125**, depending on the particular application being considered. Such an increase in performance has to be compared with the rate of increase of the computer performance described above, *i.e.* a factor of (only) five over the last five years, extrapolating to a factor of 25 over ten years. This shows that co-designing the compute system with respect to the applications performance brings clearly a number of benefits. The second option for co-design is the most commonly used. In this approach, key elements of applications are identified (under such names as building bricks, kernels, dwarfs, mini-apps, motifs, …) and these key elements are then optimized on given compute infrastructures. The philosophy is that refactoring large established codes is a major effort, and that sharing software pieces between applications avoids duplication efforts. Differently from the Japanese approach, it seems then that the hardware development is first, and that the applications motifs are then optimized on this hardware. Among such co-design initiatives one can refer to the one developed under the actual Exascale computing project (ECP) initiative in the US.

ECP's focus is on delivering, not later than 2023, "capable" Exascale systems, and meaning that hardware, software, applications, platforms and facilities are co-designed to deliver sustained performance. It is aimed at solving science problems 50 times faster as compared to actual Petascale scale systems. 24 different applications have been selected, from which different motifs are extracted, and five co-design centers have been established to target crosscutting algorithmic methods that capture the most common patterns of computation and communication for future efficient Exascale applications. These co-design centers address such issues as the growing disparity between simulation needs and I/O rates that makes performing offline analysis infeasible; particle dynamics in a variety of contexts (*e.g.* MD, hydrodynamics, particle-in-cell …), solving sparse matrices and graph operations, new block-structured adaptive mesh refinement framework, development of finite element discretization libraries, … It has already been recognized that optimizing on accelerator processors is the main difficulty, requiring that the design of the code is done with awareness of this type of hardware. ECP sees it to be far more important than choosing the "right" programming model.

Co-design approach as such is not as much developed in Europe. There are a number of initiatives in different countries, but no real common initiative at the European scale to share developments. The ten CoEs, nine of them being specialized in a given range of applications, are exchanging their views and needs under the newly created concerted support action called "FocusCoE". It indeed appears that a large number of issues are shared between them, e.g. load balancing, programming models for Exascale, performance, portability of codes, standardization of programming models, dynamic (task) scheduling, scalable solvers, data flow, in-situ data analysis and I/O, ensemble runs, implementation of co-design and technology integration, post-processing on the fly, data-focused workflows, use of large width vector units, use of heterogeneous architectures and strong memory hierarchy, visualization and data processing/analysis, workflows combining HPC simulations with associated data management and analytics capabilities. They have developed expertise in a number of these fields, but, on the one hand, it does not seem that adequate mechanisms are in place to support the efficient sharing of this expertise between applications and to save from duplication efforts. On the other hand, the co-design efforts with application developers would also need to specifically address the issues concerned with the adaptation to new, converged Exascale architectures, as this is seen as a crucial and necessary step for remodeled application codes in the coming years.

6 The main messages

It may be useful to recap here, in a few words, some of the main messages which are developed in the above chapters.

For most of the application domains the roadmaps are now at a turning point. There is always a crucial need for simulating more precisely the phenomena at stake, e.g. either by increasing the resolution, or by including more detailed physics in the model, or by running more instances of the models to estimate uncertainties. Exascale simulations are obviously on the roadmap to many developments. *But there is also pressing opportunities to interface physics-driven modeling and simulations with complementary information produced by ML (DL), taking advantage of the ever-increasing amounts of available data, coming either from observations or from more detailed, generally off-line simulations. The development of the corresponding hybrid approaches, between "traditional" HPC and ML, cover many different aspects:*

- *increasing importance of* ***hybrid-modeling approaches****, either by using ML techniques to solve more efficiently parts of HPC models (parameterization of subgrid scale phenomena, solvers and preconditioners, ...), or by developing model-based (physics-based) ML;*
- *need for more-and-more resource infrastructures allowing at the same time efficient numerical simulation of physical phenomena and treatment of massive data, calling in turn for resources where different types of processors are associated, e.g. CPUs for HPC and GPUs or other types of accelerators for* ***converged HPC / ML workloads****;*
- *and, last but not least,* ***support hybrid training*** *so that application developer teams can address all aspects of these new methods.*

What has just been said concerning the evolution of resource infrastructures must also be combined with the fact that HPC facilities, either as concentrated centers or of cloud-types, are more and more integrated inside a global cyber-infrastructure, from places where the data are being produced to the place where they are used, stored and archived.

Although this is a long-standing issue, a bit out-of-scope of the present report, such improved simulation methods will require more detailed validation, calling in term for sophisticated post-processing in relation with massive validation data.

Another message would concern the co-design process between hardware and software developers on the one hand, and application developers on the other hand. Such co-design appears to be less developed within Europe as compared with the USA and, even more, with Japan. Addressing and supporting co-design issues in this way would largely facilitate efficient use of Exascale converged facilities for a number of applications.