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

EXDCI-2

European eXtreme Data and Computing Initiative - 2

Grant Agreement Number: 800957

D3.3
Second Report on joint brainstorming sessions among scientific and industrial users' communities
Final

Version: 1.0
Author(s): Mathis Bode (RWTH Aachen University), Lee Margetts (University of Manchester), Stéphane Requena (GENCI)
Date: 23/12/2020

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: D3.3	
	Deliverable Nature: Report	
	Dissemination Level: PU	**Contractual Date of Delivery:** 31 / December / 2020
		Actual Date of Delivery: 29 / December / 2020
	EC Project Officer: Athanasia-Charalampia Evangelinou	

* - The dissemination levels 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:** Second Report on joint brainstorming sessions among scientific and industrial users' communities	
	ID: D3.3	
	Version: 1.0	**Status:** ***Final***
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): EXDCI-2-D3.3-v1.0.docx	
Authorship	**Written by:**	Mathis Bode, RWTH Aachen Univ. Lee Margetts, University of Manchester Stéphane Requena, GENCI
	Contributors:	NAFEMS
	Reviewed by:	Maike Gilliot, ETP4HPC David Tur, HPCNow
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	10/12/2020	Draft	Before internal review
0.2	22/12/2020	Draft	After internal review
1.0	23/12/2020	Final version	Approved by TB/MB

Document Keywords

Keywords:	PRACE, Research Infrastructure, Covid-19, Industry, Science, AI.

Table of Contents

List of Figures

References and Applicable Documents

[1] http://www.exdci.eu/
[2] http://www.prace-project.eu
[3] https://arxiv.org/abs/1911.11380
[4] https://covid.bioexcel.eu
[5] https://molssi.org/
[6] https://Covid-19-hpc-consortium.org/
[7] https://prace-ri.eu/hpc-access/hpcvsvirus/
[8] https://www.state.gov/g7-science-and-technology- ministers-declaration-on-covid-19/
[9] https://www.euvsvirus.org/
[10] https://www.exscalate4cov.eu

List of Acronyms and Abbreviations

AI	Artificial Intelligence
CAE	Computational-Aided Engineering
CFD	Computational Fluid Dynamics
CoE	Centres of Excellence for Computing Applications
CPU	Central Processing Unit
D	Deliverable
E4C	EXSCALATE4CoV
EC	European Commission
EESI-2	European Exascale Software Initiative 2
EIP	European Innovation Partnership
FEM	Finite Element Model
FETHPC	Future and Emerging Technologies in High Performance Computing
GAN	Generative Adversarial Network

GPU	Graphics Processing Unit
HTC	High Throughput Computing
HPC	High Performance Computing
HPDA	High Performance Data Analytics
MS	Milestone
NAFEMS	National Agency for Finite Element Methods and Standards
PIESRGAN	Physics-Informed Enhanced Super-Resolution Generative Adversarial Network
SMEs	Small and Medium-sized Enterprises
WP	Work Package

Executive Summary

This deliverable summarizes recent information about joint brainstorming sessions between scientific and industrial users' communities. The Covid-19 pandemics changed the focus of current interests tremendously as also reflected by this deliverable. The originally planned EuroHPC Summit Week 2020 was cancelled due to Covid-19 restrictions. Instead, virtual meetings and programs fighting the pandemics have become more important.

An overview of the EXDCI-2 workshop about HPC and AI convergence is briefly given here. The workshop was held virtually with 86 participants from Centers of Excellences (CoEs), FETHPC (Future and Emerging Technologies in High Performance Computing) projects and end-user scientific and industrial communities. An example for AI methods developed in science and later used in industry is given.

The results of the updated NAFEMS (National Agency for Finite Element Methods and Standards) survey about simulations in industrial environments are presented. A focus is the difference compared to the results of the first NAFEMS survey done in 2015.

1 Introduction

Work Package 3 (WP3) "Excellence in HPC applications and usages" focuses on applications and best practice usage in the context of potential requirements towards Exascale platforms. This concerns classical High Performance Computing (HPC) applications but also High Performance Data Analytics (HPDA), High Throughput Computing (HTC), and Artificial Intelligence (AI). As ascertained during the EXDCI project, Europe is developing a significant fraction of the applications used in the world and the biggest producer of data. Therefore, this effort must be continued as science opportunities are evolving very quickly with the expected availability, in a few months/years, of Exascale supercomputers (and corresponding HPDA and AI tools) since: a) large research infrastructure designs are evolving with new capabilities of HPC, Big Data, and data-driven workflows; b) new tools and approaches are increasingly needed to take into account new technical realities; c) as users get a better understanding of the potential of HPC, Big Data, and data-driven applications, possibly combined, science goals evolve rapidly and new application domains appear.

For scientific applications, WP3 relies on the PRACE [2] Scientific Case and interacts with the PRACE user communities, and the HPC Centers of Excellence for Computing Applications (CoEs) with the intention to focus on the Exascale aspects, as well as on the influence of specific technological or algorithmic innovations. WP3 roadmaps HPC applications and usages.

Similar considerations apply to industrial applications, but in this case, it is often required to fit the application in a complex proprietary workflow, for instance a "digital twin" of a product or manufacturing setup. This requires engaging and working directly with the users or their collective organizations including European Innovation Partnerships (EIPs) and national initiatives. The roadmap results are of importance to avoid gaps in the value chain, but also to evaluate scenarios where "disruptive innovation" entails changes in the value chain or permit different entry strategies.

To fulfil this role within the EXDCI-2 project, WP3 is divided into three tasks:

- Task 3.1 – Roadmap of HPC applications and usages
- Task 3.2 – Engagement with HPC users' communities and CoEs
- Task 3.3 – Preparation of industrial codes to exascale

This deliverable entitled "D3.3 – Second Report on joint brainstorming sessions among scientific and industrial users' communities" summarizes developments and analyzes in the European scientific and industrial users' communities in 2019/2020. Originally, the EuroHPC Summit Week 2020 planned to be held in Porto played an important role for the preparation of this deliverable as it brings scientific and industrial communities together. The relevant milestones for this deliverable are:

- MS7 – 2nd brainstorming session
- MS9 – Workshop during HPC Summit Week 2020

Due to the cancelled EuroHPC Summit Week 2020, the content of this deliverable changed. Instead of focusing on brainstorming sessions during the EuroHPC Summit Week, other activities were emphasized. An example is HPC activities with respect to Covid-19, which are summarized here. Another highlight of this deliverable is the updated NAFEMS (National Agency for Finite Element Methods and Standards)/EXDCI-2 survey on "Computing Platforms for Engineering" simulations.

2 Summary of Deliverable

This delivery summarizes discussions and developments between scientific and industrial users' communities in 2020. The description is separated into three parts:

- Summary of Workshop on HPC and AI convergence (cf. Chapter 3)
- Summary of NAFEMS/EXDCI-2 survey (cf. Chapter 4)
- Summary of Covid-19 HPC activities (cf. Chapter 5)

The report finishes with some conclusions.

3 Workshop on HPC and AI convergence

WP3 organized a workshop with the target of discussing important Exascale topics among various scientific communities including participants from industry. The workshop was designed as the continuation of a series of workshops and therefore denoted as "2nd European Communities Workshop on Exascale Computing ". In order to support fruitful discussions, AI and HPC convergence was chosen as the target topic, and 12 speakers were invited to present their ideas/experiences/solutions/questions with respect to this topic.

Due to the importance of getting a detailed idea about the entire European Exascale landscape, the EXDCI-2-WP3 workshop was organized in close collaboration with Guy Lonsdale from FocusCoE. Finally, the speakers represented CoEs, FETHPC (Future and Emerging Technologies in High Performance Computing) projects, science, and industry. The agenda of the workshop is shown in Figure 1 and Figure 2.

The virtual workshop took place on Thursday, 26th November, 2020 from 2:30pm to 6:00pm. Eighty-six people attended the workshop in total, which was quite successful. The presentation covered various aspects of AI and HPC convergence, from Big Data processing and parallel data analysis, to provenance and reproducibility of simulation results. The discussions and questions showed that many European communities face similar challenges and would be thankful for more joint efforts and combined solutions. For WP3, it was important to get this feedback, to improve the collaboration with various actors of the European ecosystem.

A snapshot of some participants during the meeting is presented in Figure 3. A full summary of the workshop can be found in D3.5 "Second report on the organisation of WP3 workshops during HPC Summit Week 2020".

One example for HPC and AI convergence presented during the workshop is briefly discussed here as an example of how new developments transfer from science to industry within the European HPC ecosystem. Physics-informed deep learning methods emerged over the last few years. The idea is to use physical information, such as underlying equations, to support the accuracy of deep learning-based predictions. These physical constraints can be either enforced as part of the network architecture or the loss function. The applicability of this method is very general and ranges from control of power supplies to climate modeling.

A physics-informed deep learning method for turbulence modeling was presented by Mathis Bode during the workshop. The network is called Physics-Informed Enhanced Super-Resolution Generative Adversarial Network (PIESRGAN), sketched in Figure 4. It employs physical information of turbulence in the loss function. As GAN, it consists of a generator deep network and a discriminator deep network and features an adversarial loss term, which can be seen as feedback between both network parts. The generator is finally used as a model for turbulence closure in industrial relevant flows. It was emphasized how the adversarial and physics-informed loss terms contribute to an extrapolation capability for higher Reynolds numbers and that the model predicted mixing more accurately than classical models, such as the dynamic Smagorinsky model [3]. The model was developed as a scientific project at RWTH Aachen University and finally used to support the development of engines towards Euro 7-vehicle emission standard within the automotive industry.

European
Extreme Data
& Computing
Initiative

Virtual Workshop on AI and HPC convergence

organized by

EXDCI-2 (European eXtreme Data and Computing Initiative - Phase 2)

November 26th, 2020

Agenda November 26th, 2020

14:30	**Welcome and brief introduction**	
14:30-16:00	**Session 1 – AI and HPC Convergence**	**Chair: Mathis Bode**
14:30-14:45	HPDA for Biomolecular Research: The BioExcel Case	Rosa Badia (BSC)
14:45-15:00	Towards the next level of use of HPC in Engineering	Amgad Dessoky (HLRS)
15:00-15:15	Efficient CNN for space data classification on EuroEXA multi-FPGA platform	Iakovos Mavroidis (ICS-FORTH)
15:15-15:30	An HPC-enabled Data Science and Learning Environment for Climate Change Experiments at Scale	Donatello Elia (CMCC/ESiWACE)
15:30-15:45	Elastic Ensemble Run Data Processing with Melissa	Bruno Raffin (INRIA)
15:45-16:00	Towards the use of HPC for HEP workflows	Maria Girone (CERN OpenLab)

Figure 1 Agenda of the 2nd European Communities Workshop on Exascale Computing (1/2)

16:00-16:30	*Coffee break*	
16:30-18:00	**Session 2 – AI and HPC Convergence**	**Chair: Giovanni Aloisio**
16:30-16:45	Maestro: Orchestrating Data for HPC (and AI) Applications	Dirk Pleiter (Jülich Superc. Centre)
16:45-17:00	The Combination of Real-Time Data, HPC, and Interactive Visualization in the VESTEC project	Max Kontak (DLR)
17:00-17:15	ASPIDE: A data-oriented programming model for AI/HPC Convergence	Jesus Carretero (UC3M)
17:15-17:30	SAGE - Object Storage Platform for HPC & AI	Sai Narasimhamurthy (Seagate Technology, LLC)
17:30-17:45	Subgrid Modeling of Turbulent Reactive Multiphase Flows Using Physics-Informed Enhance Super-Resolution Generative Adversarial Networks	Mathis Bode (RWTH Aachen Univ.)
17:45-18:00	The New HPC: Accelerating AI, HPC and Visualisation	Tim Lanfear (NVIDIA)
18:00	**Wrap up and closing session**	**Stephane Requena**

Program Committee

Giovanni Aloisio (CMCC & University of Salento)	*giovanni.aloisio@unisalento.it*
Stephane Requena (GENCI)	*stephane.requena@genci.fr*
Mathis Bode (RWTH Aachen University)	*m.bode@itv.rwth-aachen.de*
Sandro Fiore (University of Trento)	*sandro.fiore@unitn.it*
Stefan Krieg (FZ-JUELICH)	*s.krieg@fz-juelich.de*
Jean-Claude André (Jca Consultance and Analysis)	*jc_andre@sfr.fr*
Guy Lonsdale (SCAPOS)	*guy.lonsdale@scapos.com*

This event is funded by EXDCI-2: the EXDCI-2 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 800957 http://www.exdci.eu

Figure 2 Agenda of the 2nd European Communities Workshop on Exascale Computing (2/2)

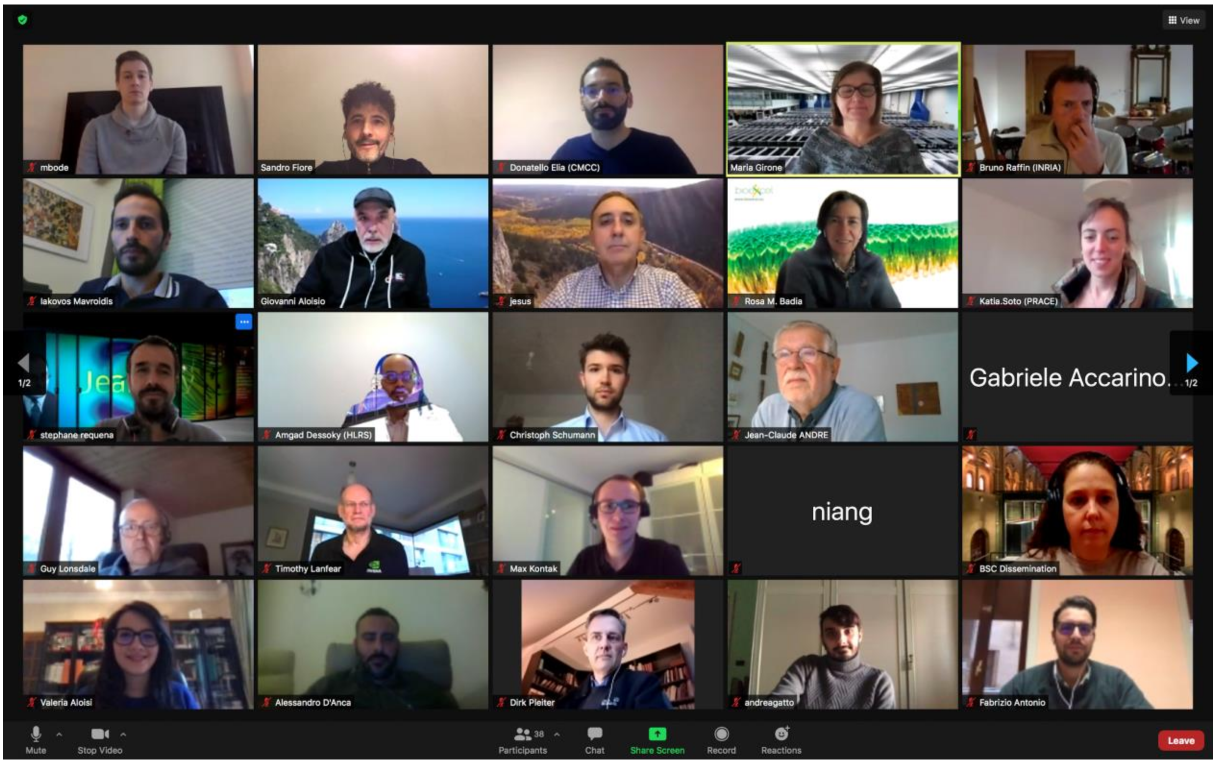

Figure 3 Overview of remote participants using the Zoom session

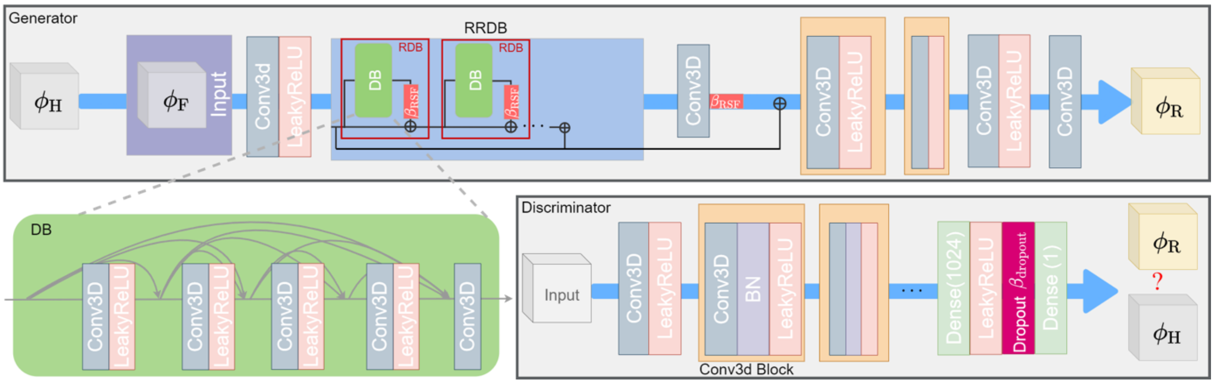

Figure 4 Sketch of PIESRGAN [3].

4 Survey of Computing Platforms for Engineering Simulation – 2015 vs 2020

The aim of this chapter is to present some early results for the recent 2020 NAFEMS/EXDCI-2 survey on "Computing Platforms for Engineering" simulations and compare them with the same survey carried out in 2014 (published in 2015). The initial survey was prepared by the NAFEMS HPC Working Group in 2015; with input from colleagues in the European Exascale Software Initiative (EESI-2), the N8 HPC Service (UK) and Teratec (France). The second survey was prepared as part of EXDCI-2 task 3.3 run by NAFEMS. The survey was advertised to NAFEMS members and the broader engineering simulation community in August 2020, and data collection ended in November 2020.

NAFEMS is an international trade association, established in the 1980s, that focuses on promoting best practice in the use of simulation in engineering. It has 1800+ institutional members across an engineering ecosystem that includes blue chip firms such as Boeing, Airbus and RollsRoyce and micro-firms consisting of less than 10 employees. NAFEMS organizes events, provides training, certifies engineers under the industry devised "Professional Simulation Engineer" scheme, and publishes guides for the practitioner. All of these activities are centrally organised by a not-for-profit SME, with input from both regional and technical working groups.

The original motivation for the 2014 survey was to find out to what extent the CAE (Computational-Aided Engineering) community is making use of HPC, the cloud and other advanced computing platforms for engineering simulation. That said, the survey was aimed at all users of engineering simulation to ensure that, as far as possible, survey responses represented the engineering simulation community as a whole. The survey was repeated using the same wording, categories, and classification in 2020.

Here, survey responses are reviewed for the following topics: (i) the largest simulations carried out for the main types of engineering analysis, and (ii) the maximum number of cores used.

4.1 Respondents

A total of 231 respondents started the survey published in 2015 and 374 commenced the survey in 2020. This was thought to be a reasonable sample of the targeted engineering simulation community when compared with response rates for NAFEMS surveys carried out in the past. Figure 5 shows the segmentation of respondents by geography and Figure 6 and Figure 7 by business area.

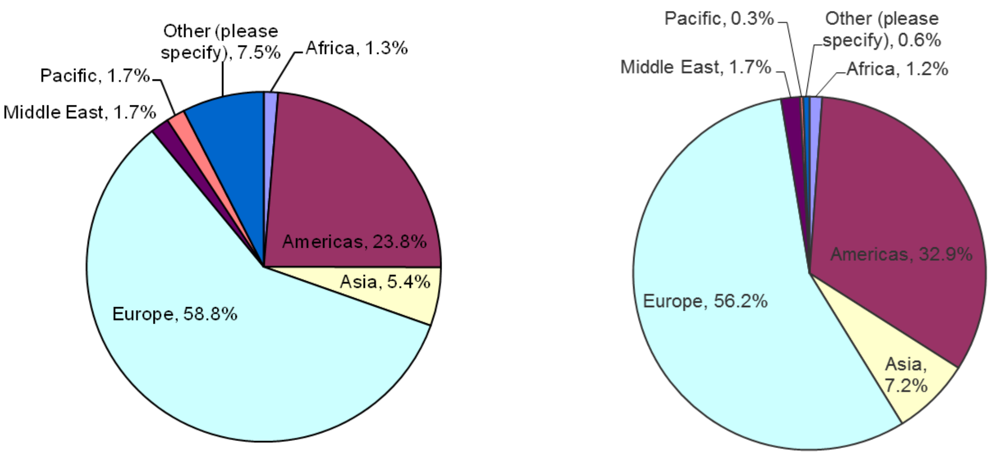

Figure 5 Respondents segmented according to location of headquarters.

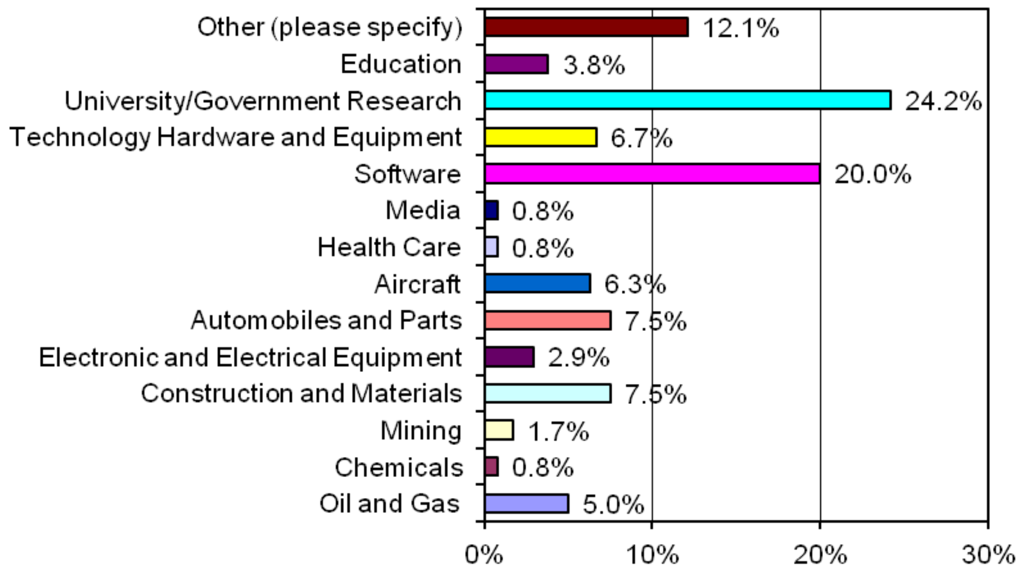

Figure 6 Respondents segmented by main business area in 2015.

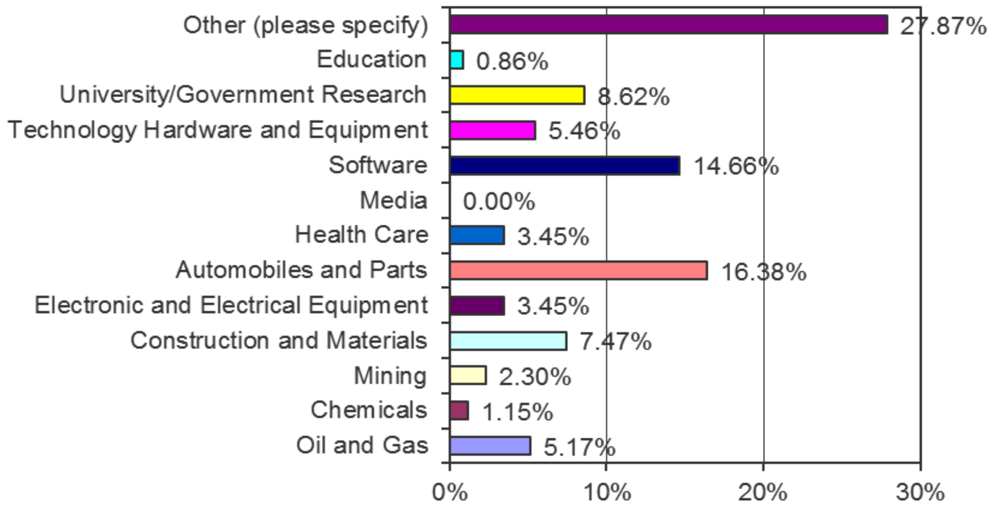

Figure 7 Respondents segmented by main business area in 2020.

In 2015, responses were dominated by firms with headquarters in Europe (59%) and the Americas (24%). In terms of business area, there was a roughly even split between engineering firms (30%), the public sector (29%) and software/hardware vendors (28%). 13% of respondents marked their business area as "other".

In 2020, responses were dominated by firms with headquarters in Europe (56%) and the Americas (32%). In terms of business area, respondents were predominantly engineering firms (47%). Respondents from the public sector (10%) and software/hardware vendors (15%) decreased substantially compared with 2015. 28% of respondents marked their business area as "other".

In 2015, there was an even split between SMEs with fewer than 250 employees (47%), and large firms with more than 250 employees (53%). The SME responses include micro-firms with fewer than 10 employees (17%), small firms with 10 to 49 employees (14%) and medium sized firms with 50 to 249 employees (16%). All percentages stated in this section are the percentages of all 231 respondents.

In 2020, 43% of respondents were SMEs and 57% were large firms. The SME responses include micro-firms with fewer than 10 employees (12%), small firms with 10 to 49 employees (13%) and medium sized firms with 50 to 249 employees (17%).

4.2 Facilities

The survey asked respondents to state how often they used (or might use) a broad range of devices in their organisation, now and in the future. The options that could be chosen were "always", "often", "sometimes" and "never". The responses for each device were ranked according to the sum of "always", "often" and "sometimes", giving a broad indication of the preferences of the engineering simulation community.

Figure 8 shows the response for current usage (2015). There is a strong bias towards the use of computer hardware hosted by the organization, with laptops and workstations being preferred to internally hosted HPC facilities. An equally strong, but negative, bias is seen for the use of externally hosted facilities. Regardless of the type of organization doing the hosting (academic, government or commercial service), usage of externally hosted facilities is low compared with desktop systems. According to the data collected, smartphones and tablets appear to be the devices used least for engineering simulation. The main change in 2020 is an increase in the use of HPC systems owned by the respondent's own organization and high memory workstations (Figure 9).

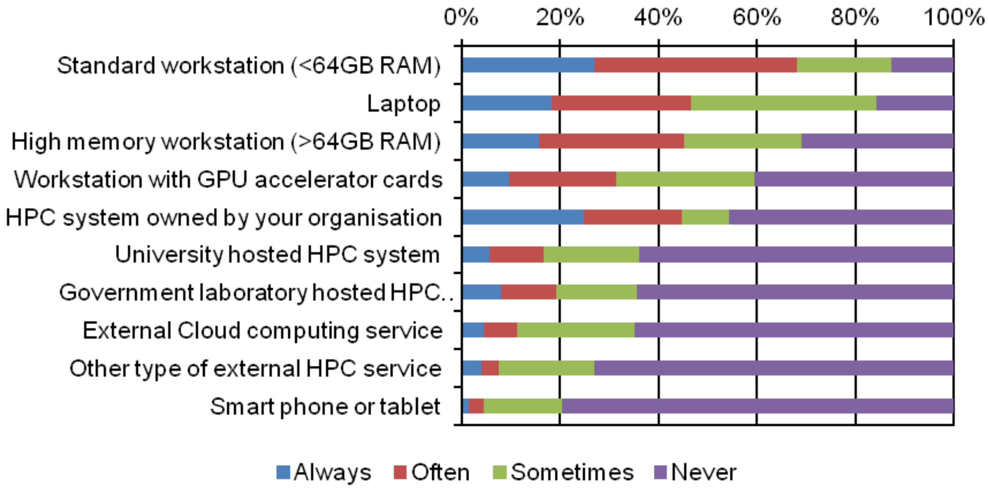

Figure 8 Facilities currently used for simulation in the organization (2015).

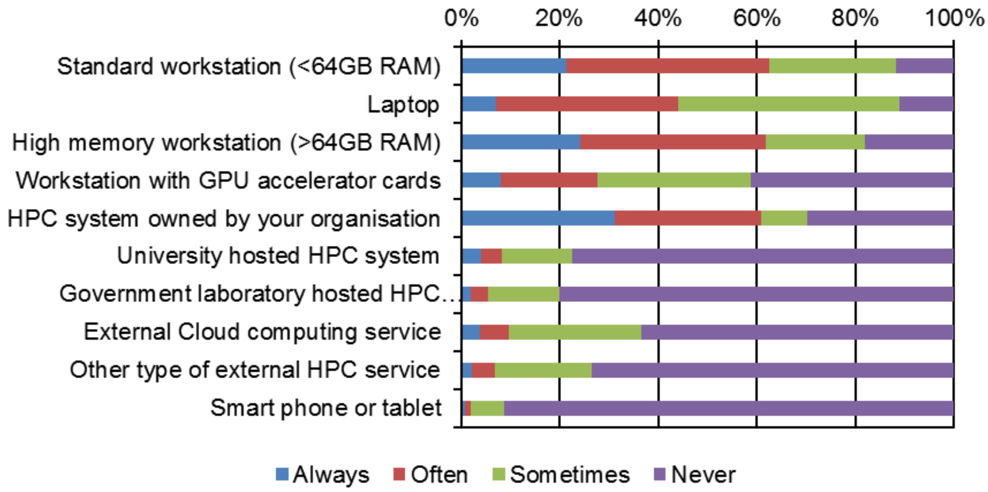

Figure 9 Facilities currently used for simulation in the organization (2020).

Figure 10 gives an indication of how respondents in 2015 viewed their use of these hardware platforms in 2020. There is little correlation between the predicted change and actual change.

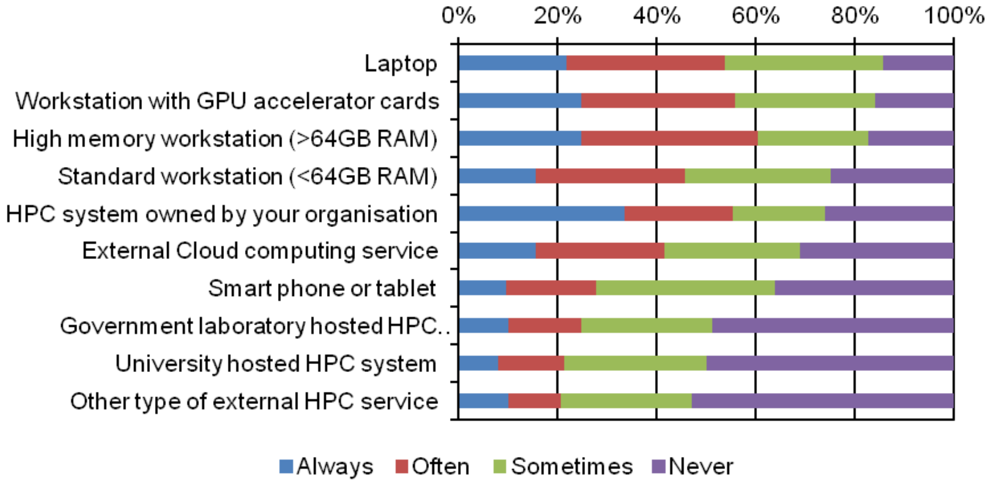

Figure 10 Facilities that may be used for simulation in the future (2015).

4.3 Matters of Size

One of the objectives of the survey was to get an idea of the size of the largest simulations carried out in the community and the size of the largest facilities used for those simulations. Respondents were asked to indicate the size of the largest simulation carried out in their organization, measured in terms of the number of degrees of freedom, for a list of different types of engineering analysis. For example, in a two-dimensional finite element model (FEM) of a structure, displacements in the x-direction and y-direction are the degrees of freedom in the system of equations to be solved. The size of facilities is measured in terms of the number of cores in a standard multi-core CPU. Respondents were asked not to include GPU cores in their estimates.

Figure 11and Figure 12 plot survey data for 5 bins of maximum problem size. The numbers on the scale indicate the percentage of respondents in a particular bin and add up to 100% for each simulation category. For example, adding the values along the Computational Fluid Dynamics (CFD) line, for each size of bin, gives 100% for the CFD data.

Figure 11 shows the 2015 data. The largest problems, for all types of engineering simulation, are mostly in the 100,000 to 10 million degree of freedom range. There is a strong bias towards smaller problems in data analysis, systems simulation and boundary simulation. In terms of the larger categories, there is a bias towards CFD and multiphysics for problems >10 million degrees of freedom and only CFD has a relatively strong response (10%) for problems >1 billion degrees of freedom.

Figure 12 shows the 2020 data. The main change is an increase in the 10M to 100M problem size for all types of engineering simulation. Similar to 2015, problems are still mostly in the 100,000 to 10 million degree of freedom range. The increase in 10M to 100M degrees of freedom problems appears to be at the expense of <100K and 100M to 1 billion degrees of freedom problems.

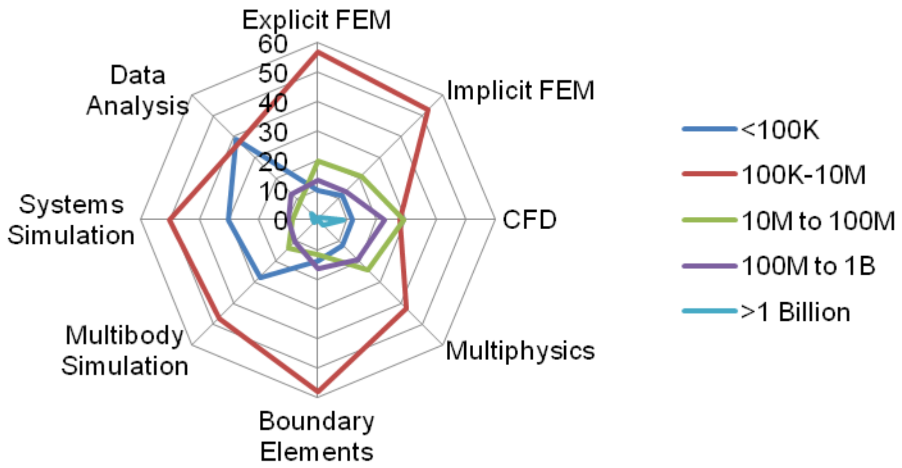

Figure 11 Maximum number of degrees of freedom by type of simulation 2015.

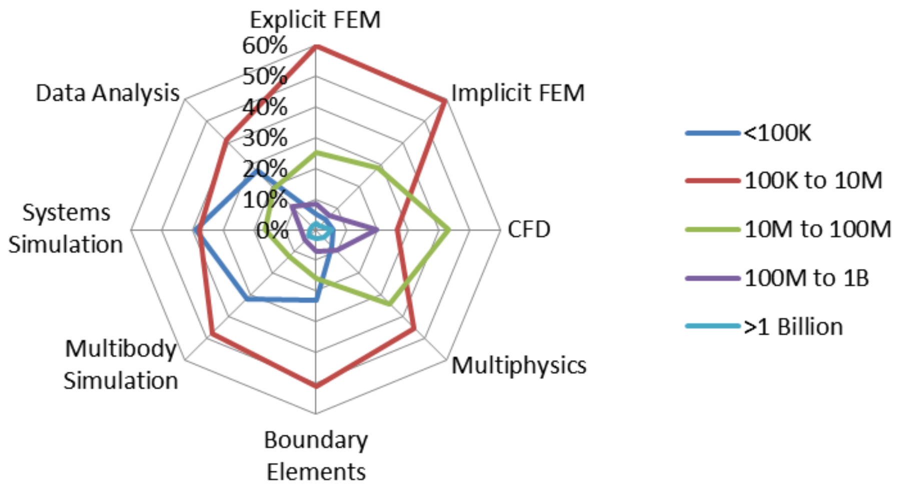

Figure 12 Maximum number of degrees of freedom by type of simulation 2020.

Figure 13 plots data for the maximum number of cores used in each of the categories of engineering simulation for 2015. The values along any line, corresponding with a particular type of simulation, add up to 100%. The figure shows a clear bias towards small core counts for data analysis, systems simulation, and multibody simulation. The ranges 9-64 cores and 65-1024 cores are fairly similar. There is a small bias towards CFD and multiphysics for 65-1024 cores and a stronger bias towards CFD for simulations using more than 1024 cores. The survey included a >8196 core bin, but only a handful of respondents ticked this box and therefore the data was added to the >1024 bin. Figure 14 shows the results for the same question. As before, the systems simulation bin is missing due to an error in setting up the survey. The main difference is the growth of the 9-64 core bin, which may coincide with the change in the number of cores in a workstation.

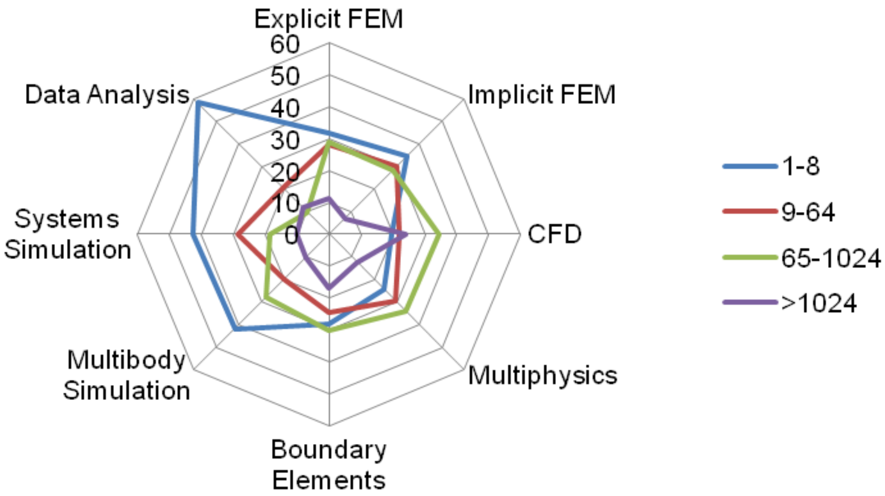

Figure 13 Number of cores used by type of simulation 2015.

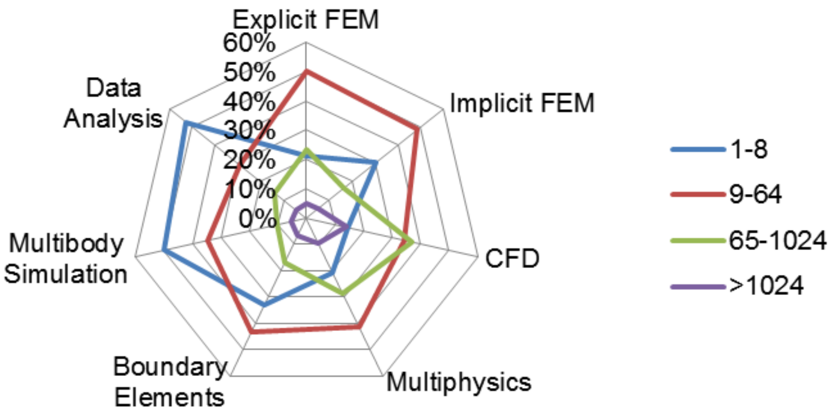

Figure 14 Number of cores used by type of simulation 2020.

4.4 Barriers to Use a Larger Number of Cores

Respondents were asked to comment on a list of possible barriers to using more cores for engineering simulation; both from the point of view of facilities provided by their own organization and with respect to access to externally hosted services, such as public cloud services or facilities provided by universities.

For each of the potential barriers listed, respondents were asked to indicate whether each was "not a barrier", was a "weak barrier" or was a "strong barrier". The data collected has been weighted by giving each response a score of 0 for "not a barrier", 1 for a "weak barrier" and 2 for a "strong barrier". An average score was determined by dividing the total score for each barrier by the number of respondents. A score of less than 1 indicates that the average response is somewhere between "not a barrier" and a "weak barrier". A score greater than 1 indicates the average is somewhere between a "weak barrier" and a "strong barrier". The maximum possible score using this methodology is 2.00, which would be given if all respondents marked "strong barrier" as their response. The results are shown in Table 1 and Table 2. Notably none of the items had a score close to zero, thus all the items on the lists were considered as barriers to some extent. The highest scoring barriers are cost related.

Table 1 Barriers to using more cores in own organization.

	Barrier to uptake	**2020**	**2015**
1	Prohibitive licensing cost of the software	1.41	0.95
2	Purchase cost of the hardware	1.32	1.13
3	Total cost of ownership of the hardware	1.21	1.13
4	Difficulty automating workflows between different packages	0.84	0.80
5	Difficulty in managing increasing volumes of data	0.83	0.84
6	Poor performance of the software	0.74	0.85
7	Decision makers not convinced of the business benefits	0.73	0.71
8	Existing facilities oversubscribed (too small for demand)	0.70	0.75
9	Difficulty scheduling usage of facilities	0.60	0.63
10	Staff lack the necessary skills	0.60	0.70

Table 2 Barriers to using more cores on externally hosted platforms.

	Barrier to uptake	**2020**	**2015**
1	Total cost of service (compute cycles & software licenses)	1.38	1.28
2	Concerns about data security and intellectual property	1.28	1.32
3	Difficulty in managing increasing volumes of data	0.99	1.11
4	Decision makers not convinced of the business benefits	0.92	0.86
5	Difficulty in automating workflows between different packages	0.90	0.82
6	Staff lack the necessary skills	0.74	0.63
7	Poor performance of the software	0.70	0.74
8	Difficulty scheduling usage of facilities	0.66	0.77

4.5 Discussion

The survey provides some useful insight into the use of a range of computing platforms for engineering simulation. However, the results provide just two data points in a rapidly evolving era of computing. The 2015 survey asked respondents to reply on the behalf of their organization and there was much discussion as to whether this focus was correct, particularly on social forums such as LinkedIn. Given that an individual respondent will have a much better idea of what they do than what their peers do, the 2020 survey asked for the individual's usage rather than their assessment of use in their organization.

In terms of facilities, there is a danger that the survey results will be out of date quite quickly. The survey asked respondents to comment on the number of cores used in their simulations for a "typical" multi-core processor (Intel Xeon, AMD Opteron or IBM Power series). The use of accelerators, such as GPUs was not considered. Future surveys may have to deal with a broader range of hardware; for example, desktops, HPC clusters and cloud computing platforms comprising "standard" processors packaged together with co-processors, accelerators and other technologies that the hardware vendors may bring to market. Comparing the computational power of these systems may require a new effort from NAFEMS on benchmarking; for computational speed rather than simulation accuracy.

The responses to the "barriers to using more cores" question indicate that there may be a suppressed demand for more computational power in the engineering community. There are no surprises with hardware costs, licensing and data protection being important issues. However, care needs to be taken in interpreting the data. These responses do not represent a ranking of the barriers most cited by the respondents, but the respondents' ranking of a list of barriers proposed by the NAFEMS HPC Working Group.

Finally, the author acknowledges that further processing of the survey data may lead to further important insights. For example, separating the responses of industrial and academic users may indicate whether there is a strong difference between industrial and academic choices regarding the facilities each uses and the size of problem solved. There may also be enough data to assess whether there are significant differences between the responses received from Europe and the Americas.

4.6 Implications

The survey responses show that laptops and workstations are the preferred hardware platform for engineering simulation in both 2015 and 2020 and this may limit problem size. Comparing the two surveys, there is a noticeable increase in the 10M to 100M degree of freedom problem size in 2020 and a reduction in the 100K to 10M bin. An increase in the number of cores used is also seen in 2020, but the 9-64 core bin which has grown probably corresponds to increased capability in workstations.

Usage of HPC facilities is seen across all types of engineering simulation, but core counts (65-1024 cores) are low compared with the capability of the largest HPC systems (~100 million cores). The CFD community runs the largest problems (1 billion degrees of freedom) on the largest number of cores (~1024), highlighting maturity in the use of advanced hardware platforms in this area. Comparing the results published in 2015 with the 2020 survey, there is an increase in the use of HPC systems owned by the respondents' own organization. Use of externally hosted systems such as cloud services appears unchanged.

Considering the respondents' own institution, the most significant barriers to using more cores for simulation were hardware and software licensing costs. The scores have increased in 2020 compared with the results published in 2015. This may reflect an increase in respondents looking at increasing their use of HPC and being more aware of the costs. As in 2015, concerns about data security and intellectual property still appear to be the main issue in 2020 holding firms back from using the externally hosted cloud facilities.

5 HPC Support to Covid-19 Researchers in Europe

Since the rise in Europe of the Covid-19 pandemic and especially since the start of the first lockdown measures in Europe, many HPC research infrastructures at the national level or at the European level (PRACE) have decided to setup fast-track calls with priority access (urgent computing) toward researchers engaged against Covid-19.

In that sense applications received have been evaluated scientifically and technically in few days allowing research teams from academia and industry to rapidly access to leading edge supercomputers.

Such measures include access to compute facilities on national (Tier1) or European (Tier0) systems as well as storage facilities and dedicated user support for helping sometimes user communities not used to access to such level of supercomputers to run their simulations.

5.1 PRACE Covid-19 Fast Track

The PRACE HPC research infrastructure setup in 24 March 2020 a fast-track access to HPC resources (#HPCvsVIRUS) with the active support of the following members of the PRACE Scientific Steering Committee (SSC):

- Luigi Del Debbio (Chair of PRACE Access Committee), University of Edinburgh, UK
- Marc Baaden (Vice-Chair of PRACE Access Committee), CNRS, Université de Paris, UPR 9080, Laboratoire de Biochimie Théorique, France
- Matej Praprotnik (Chair of PRACE Scientific Steering Committee), National Institute of Chemistry & Faculty of Mathematics and Physics, University of Ljubljana, Slovenia
- Laura Grigori (Vice-Chair of PRACE Scientific Steering Committee), Inria and Sorbonne University, France
- Núria López (Former Chair of PRACE Scientific Steering Committee), Institute of Chemical Research of Catalonia, ICIQ, The Barcelona Institute of Science and Technology, BIST, Spain

At date a total of 83 Covid-19 proposals were received by PRACE, from 21 different countries leading to **30 proposals awarded** (for a total of more than **500 million core hours allocated**), 49 rejected and 4 redirected to national resources by PRACE.

In the field of a collaboration with the BioExcel CoE [4] and the MolSSI US initiative [5], specific data storage and interactive data processing services have been also provided through the use of the ICEI Fenix "Platform for Computational Molecular Data Exchange".

Among all the supported projects it's interesting to notice about the variety of scientific domains and stakes including:

- epidemiological studies (modeling of the spread of the disease across cities and countries) and effects of lockdown measures for advising public authorities;
- massive screening of candidate molecules (including repositioning techniques);
- studies of the structure and screening results of targets in the virus such as 3CLpro, Spike, nsp1, RNA-dependent RNA polymerase…
- modeling of social distancing including high resolutions for tracking infected droplets (depending to size, wind, location, …) and use of different shape of masks;
- assessment of possible side effects of some molecules used against Covid-19.

5.2 Europe-US Collaboration

Through the pre-existing relationship between PRACE and XSEDE (NSF national HPC infrastructure, USA), further collaborations to bring HPC into the fight against Covid-19 are being developed during weekly meetings of representatives from both organizations.

Regular Zoom webinars for all project leaders from US (HPC Consortium) and PRACE to discuss ideas, exchange data are organized. The number of participants is around 70 by webinar. Participants are shortly presenting their projects. The aim is to get collaboration from both sides of the border, to avoid that participants do the same simulations and to share the data in an open way.

Collaboration between researchers will also be facilitated through consolidated lists of projects and their corresponding contacts on a common website. The platform and future activities will be highlighted and promoted.

Visitors to the PRACE website can find a link to the Covid-19 HPC Consortium [6] via the "Collaborations" section on a new page within the PRACE website that lists the main activities PRACE undertakes to support the fight against COVID-19 [7].

On that same PRACE webpage, a reference is made to the G7 statement "*A Shared Vision for Science and Technology in Responding to the Pandemic, Protecting Human Health, and Promoting Social and Economic Recovery" and mentioned PRACE in the paragraph "Strengthen the use of high- performance computing for Covid-19 response*." [8]

Thanks to a PRACE-6IP project effort, PRACE has offered resources to the EUvsVirus Hackathon [9] organized beginning of June 2020 by the European Commission.

Contacts with other regions of the world, such as Latin America (RedCLARA and SCALAC), have been made to develop collaborations.

5.3 The EXSCALATE4CoV European Project

In response of the European Commission call H2020-SC1-PHE-CORONAVIRUS-2020: Advancing knowledge for the clinical and public health response to the 2019-nCoV epidemic the EXSCALATE4CoV (E4C) project [10] involves a close collaboration among three of largest supercomputing centres in the EU (CINECA in Italy, the Barcelona supercomputing centre in Spain and the supercomputing centre of Forschungszentrum Jülich in Germany) along with a pharmaceutical company (Dompé), and several universities and research Institutes.

The E4C team uses supercomputers to perform molecular simulations and in silico, i.e. using special programs or algorithms, biochemical and phenotypic screening of existing drugs against SARS-CoV-2. This approach allows the fast analysis of simulation results and reduces the time required to discover new therapeutic agents. Indeed, the EXSCALATE (EXaSCale smArt pLatform Against paThogEns) platform permits exascale virtual screening and therefore the evaluation of billions of molecules against several targets within a few weeks. This is particularly useful for pandemic viruses such as coronavirus, where the immediate identification of effective treatments is of the utmost importance.

The E4C project involves 18 partners, led by Dompé Farmaceutici SpA an Italian Pharma company with a total budget close to 3M€ over 18 months.

Among the first results several docking simulations were performed to define and optimize the machine learning and virtual screening protocols to use on SARS-CoV-2 proteins. The performances of the virtual screening strategies were assessed by evaluating their capacity to correctly rank molecules, which are endowed with antiviral activity, and in particular, with a known effect against SARS-CoV proteins, considering the lack of known actives on SARS-CoV-2 proteins.

The tuned and validated virtual screening protocols were used to screen a repurposing library, containing the set of safe in man drugs, commercialized or under active development in clinical phases, and a set of known bioactives in particular preclinical compounds identified as "CoV Inhibitors" (> 12000 drugs). The most promising drugs, identified from docking studies as potentially active against SARS-CoV-2 proteins, were selected to be tested in biochemical and phenotypic assays.

5.4 Examples of National Covid-19 Initiatives

In France

At the end of march 2020, GENCI the French HPC agency decided to set up fast-track access to its HPC facilities located on the 3 national centers (TGCC for CEA, IDRIS for CNRS and CINES for French Universities). Like PRACE this provision of computing resources comes with access to storage facilities and dedicated user support for helping sometimes research teams which are not used to supercomputers to take advantage of these resources.

GENCI as one of the 5 hosting members of PRACE has only involved into PRACE #HPCvsVIRUS fast-track by making available the Joliot-Curie Tier-0 system at TGCC.

At date GENCI supported close to 40 Covid-19 projects from French and European researchers from academia and industry.

Like in PRACE such 40 supported projects are covering a large span of scientific domains from epidemiology (in support of the French national health authorities), massive screening of candidate drugs (with one of the largest worldwide studies conducted with 1.5 billion molecules using 15 million core hours on Occigen @ CINES), biomolecular research to understand the mechanisms of the virus infection, bioinformatics research to understand mutations/evolution, large CFD studies to model the spread of Covid-19-infected droplets under various conditions (size of droplets, wind, distance, size/shape of masks, location in train/plane, air conditioning, …) or use of AI for fast analysis of computed tomography scans of lungs.

Figure 15 and Figure 16 show numerical simulations performed in the field of CFD studies, quantum chemistry for understanding mechanisms of virus infection or AI analysis for computed tomography scans. Such projects involve research teams from academia as well as industry (from large groups like Safran to SMEs like Owkin) and also hospitals.

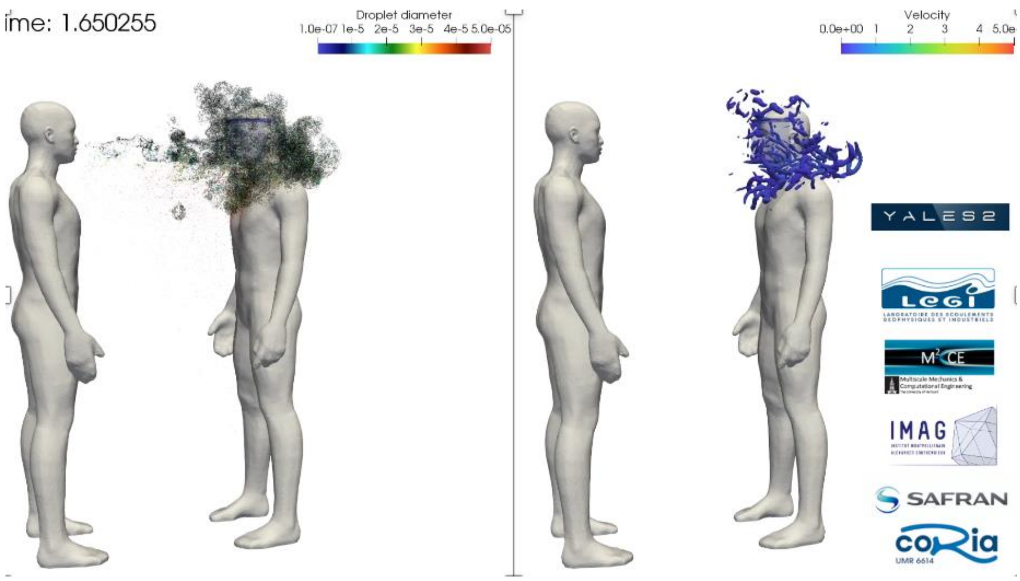

Figure 15 CFD simulations of droplets with the YALES2 code.

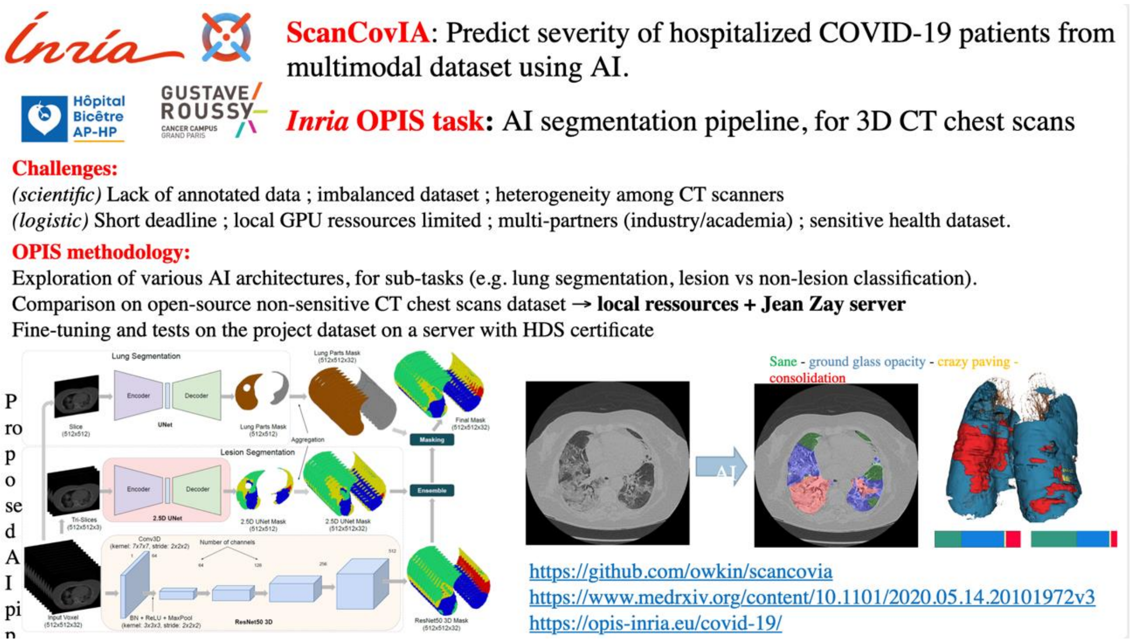

Figure 16 Overview of the ScanCovIA for AI guided computed tomography scans.

In Germany

In the same way the Gauss Centre for Supercomputing (GCS), the alliance of Germany's three national supercomputing centres has setup on March 16, a dedicated access to HPC facilities located at the High-Performance Computing Center Stuttgart (HLRS), Jülich Supercomputing Centre (JSC), and Leibniz Supercomputing Centre (LRZ).

In addition to this effort, GCS–a hosting member of the PRACE–also joined PRACE's initiative pursuing the identical purpose on March 24.

At date GCS supported more than 20 projects from Germany and Europe in many domains including Systemic epidemiological analysis of the Covid-19 epidemic, massive screening of candidate molecules, modeling Covid-19 spatio-temporal dynamics, targeting the interface of the Covid-19 spike protein with the ACE2 receptor or biomechanic simulations for quantification of the ventilation/perfusion ratio in Covid-19 patients. An example is shown in Figure 17.

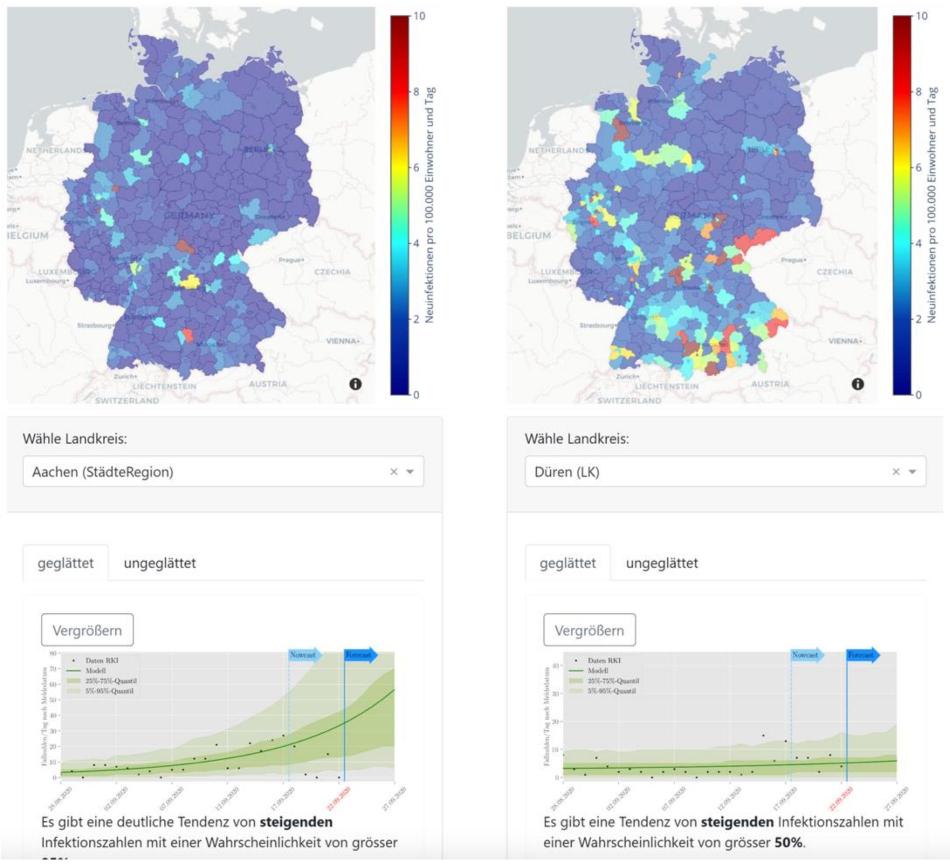

Figure 17 Daily updated corona analysis for each German district: reported data (top left) and estimation of current real case numbers (nowcast, top right) as well as forecasts (bottom) - Universität Osnabrück and Forschungszentrum Jülich.

6 Conclusions and Acknowledgements

This deliverable summarizes recent development of the scientific and industrial HPC landscape in Europe. Due to the Covid-19 pandemics, the focus of research shifted and simulations with respect to Covid-19 became a lot of attention. The HPC programs by PRACE and the EC with E4COV established for fighting the pandemics show that Europe is able to react quickly with a lot of resources.

Within Covid-19 research but also other areas of interest, AI plays a key role. Convergence between HPC and AI is happening in both directions: HPC simulations are improved by AI and AI uses the HPC system for training at scale of networks. This was also emphasized by the organized workshop focusing on HPC and AI convergence.

The update of the 2015 NAFEMS survey shows relatively slow developments, and raises the question whether the current HPC support of the EC is fitted to overcome current limitations and how issues related to the cost of software licenses are addressed which represent one of the biggest limitations for engaging users of commercial applications at scale.

Overall, all three parts of this deliverable show the adaptability of the European HPC landscape and emphasize its performance. Recent topics, such as AI, HPDA, and HTC, are integrated into HPC workflows. Importantly, that happens not only within the scientific community but these methods also emerge in industrial use cases. This full-spectrum development is also reflected by the NAFEMS survey showing stronger usage in the industrial users' community. All of these aspects are also visible as part of the fight against Covid-19 supported by the European HPC landscape.

Finally, the authors would like to acknowledge the work carried out by the NAFEMS High Performance Computing Working Group, both in designing the survey and analyzing the results. We also acknowledge support from Teratec, France and the EESI-2 project. The latter was funded by the EC under the 7th Framework Programme.