ETP 4 HPC

EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

2022 ETP4HPC's SRA 5

STRATEGIC RESEARCH AGENDA FOR HIGH-PERFORMANCE COMPUTING IN EUROPE

SEPTEMBER 2022

EUROPEAN HPC RESEARCH PRIORITIES 2023 - 2027

- Quantum for HPC
- Sustainability the Next Big Thing
- HPC in the Digital Continuum
- Industrial Use of HPC
- On the path to Exascale
- Heterogeneous High-Performance Computing
- Unconventional HPC Architectures
- Centre to edge framework
- HPC for Urgent Decision Making
- Federated HPC, Cloud and Data Infrastructures
- Application co-design

ETP 4 HPC

EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

2022 ETP4HPC's SRA 5

STRATEGIC RESEARCH AGENDA FOR HIGH-PERFORMANCE COMPUTING IN EUROPE

SEPTEMBER 2022

EUROPEAN HPC RESEARCH PRIORITIES 2023 - 2027

ETP4HPC Chairman's Message

Our Association has the long-standing tradition of documenting the priorities of European research in HPC technology. We have been doing this for almost ten years (this being our fifth Strategic Research Agenda - "SRA") and we have been quite successful in that the European Commission's HPC research programmes throughout that time have reflected, or indeed referred to, our SRAs and the community of industrial stakeholders, researchers and experts we represent felt involved in the process of making Europe home to the most competitive HPC technology.

With the introduction of EuroHPC, our community has received a powerful funding and coordination mechanism and ETP4HPC is honoured to be the leading private partner of this Joint Undertaking. The delivery of this SRA is part of our commitment to provide considered and comprehensive input to the development of the EuroHPC research programme. The ETP4HPC SRA provides a strong basis for the discussions of the EuroHPC Research and Innovation Advisory Group (RIAG) in which we participate. To fulfil this goal, we work together with all our partners to ensure that the priorities of this SRA are fully representative of the research and development needs of the expanding HPC ecosystem. This document is the work of over 130 experts associated with ETP4HPC. A huge amount of expertise lies within, an asset that should be exploited fully. Our community would be delighted to receive feedback on its work and how the priorities we have identified will be reflected in the work programmes. I think this an equally important part of the process, on par with the writing of the SRA itself, and ETP4HPC would like to encourage all other stakeholders to work with us to achieve that.

This document stipulates where the funding available should be directed to achieve a competitive, sovereign and sustainable European HPC ecosystem. In today's changing world, one needs to have independent sources of critical technology at their disposal. HPC is necessary to tackle the most important and urgent issues our society is facing: climate change, reducing global dependencies in the areas of energy and new technologies, health and many others.

And to save our planet and let the future generations live in a clean and safe world, we need to make HPC, other technologies and all the areas of our lives sustainable. Huge amounts of work are still to be done in this field and this SRA is the first stepping stone on the giants' causeway toward a green HPC ecosystem. In line with this goal, a separate chapter includes our ecosystem-level recommendations on the strategy of European HPC research with an aim to deliver a sovereign and sustainable HPC ecosystem in Europe. Apart from defining the Work Programme, we would also like to contribute to the improvement of the entire funding mechanism.

I would like to thank all authors - the leaders of the working groups, members of the working groups, the office team and all other contributors - for the work and trust they have put in this process.

ETP4HPC Chairman
Jean-Pierre Panziera

Executive summary

This Strategic Research Agenda contains the priorities of European research in HPC technology as identified by ETP4HPC. We make this SRA available to the EuroHPC Joint Undertaking and in particular its RIAG (Research and Innovation Advisory Group) with a view to basing the European HPC technology research programme on the findings of this SRA. The first part of this document introduces the key technological trends on which this work builds in addition to explaining how this SRA was written and its relevance for the European HPC ecosystem. A separate chapter includes our ecosystem-level recommendations on the strategy of European HPC research with an aim to deliver a sovereign and sustainable HPC ecosystem in Europe. While the technological part of this SRA is divided into research domains and clusters, due to HPC's complexity, an emphasis is also placed on treating HPC as part of a broad technological ecosystem i.e., the Digital TransContinuum, with concepts such as Digital Twins and the inclusion of AI-based solutions marking the current trends. This SRA also highlights the importance of the industrial use of HPC and the role of the EuroHPC-funded prototypes in developing technology that contributes to Europe's technological sovereignty. The other underlying theme is that of sustainability and this SRA summarises the initial discussions held to date on this topic and lays the ground for the extensive work still to come.

Note: Vendor neutrality

Any reference to products or solutions is intended as a reference in order to present the context and not as an implied promotion. The SRA remains completely agnostic in relation to brands and it maintains the diversity of implementation options.

Table of contents

1

What has changed since the last SRA?

Reflecting on the period since our last SRA (4), one needs to realise that our lives have been significantly impacted at least in four dimensions:

- The importance of a **digital autonomy** has been demonstrated in multiple ways: disrupted supply chains of hardware components have led to severe economic consequences; the distribution of expertise and skills across multiple digital disciplines has further increased Europe's dependency on other geographies which puts Europe's industrial competitiveness at risk.
- The **security of data- and compute infrastructures** is compromised on a daily basis by cybercriminal activities.
- The **Covid-19 pandemic** has not only led to a dramatic loss of human lives, but it has also caused unprecedented economic and social disruptions.
- The increasing occurrence of environmental disasters (e.g., floods, forest fires) demonstrates the need for an expedited global action to tackle **climate change**.

Nowadays, HPC plays a critical role in all four aspects above, reflecting how essential it is for our daily lives. It is a pivotal tool in rapidly advancing research in complex disciplines such as design of vaccines and medicines[1], or modelling, monitoring and simulating natural phenomena and related human activities as intended by the Destination Earth Initiative[2]. However, there is the challenge of maintaining the use of HPC resources in continuously stable, effective and sustainable conditions and taking measurable steps towards the ability to master the necessary R&D for new-era HPC - technology and systems.

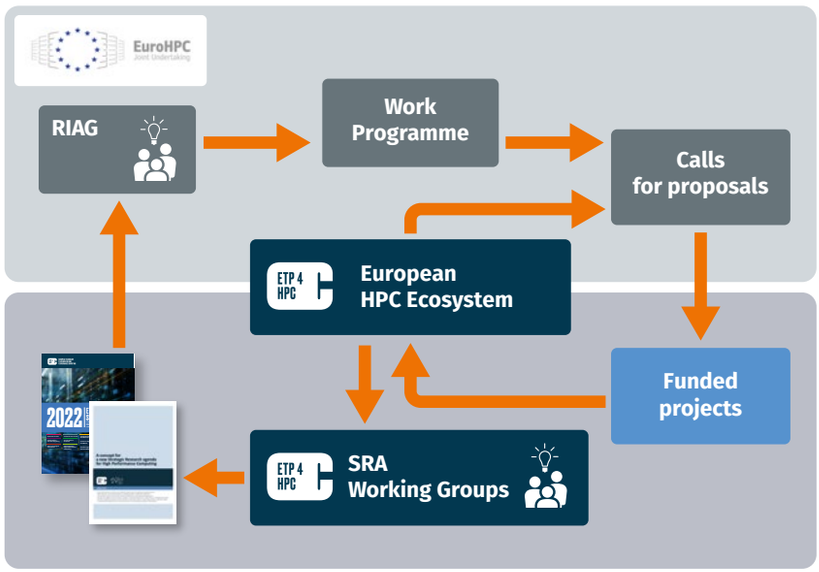

Figure 1: The role of SRA.

1. https://cordis.europa.eu/project/id/101003551
2. https://digital-strategy.ec.europa.eu/en/policies/destination-earth
3. https://ecfr.eu/europeanpower/european_sovereignty/

1.1

The role of this Strategic Research Agenda

The document feeds research and development priorities developed by the European HPC ecosystem into EuroHPC's Research and Innovation Advisory Group with an aim to define the HPC Technology research Work Programme and the calls for proposals included in it and to be launched from 2023 to 2026.

This SRA also describes the major trends in the deployment of HPC and HPDA methods and systems, driven by economic and societal needs in Europe, taking into account the changes expected in the technologies and architectures of the expanding underlying IT infrastructure. The goal is to draw a complete picture of the state of the art and the challenges for the next three to four years rather than to focus on specific technologies, implementations or solutions.

1.2

European digital sovereignty in the context of HPC

The term "European Sovereignty" has become prevalent in strategy statements concerning key technologies and also regarding the definition of research programmes. Achieving European sovereignty is a response to the complex economic interdependencies that emerged in the era of globalisation, and which resulted in multiple asymmetric dependencies that have limited Europe's freedom of action[3]. The concept of "strategic sovereignty" has been proposed in order to guide the EU through this new era of geopolitical competition. **Strategic sovereignty implies that the EU and its member states need to retain the capacity to act in a self-sufficient way as a unison of states. Promoting European digital sovereignty is a critical piece of this effort.**

European digital sovereignty naturally has an impact on the area of HPC and the most notable actions taken so far are the following:

- The creation of the EuroHPC-JU in 2018 as a legal and funding entity to promote a pan-European supercomputing infrastructure and supporting research and innovation (R&I) actions – see next Chapter for more details.

- National Competence Centres (NCCs) in the framework of EuroHPC: participating countries are tasked with establishing their single HPC National Competence Centres. The NCCs are the central points of contact for HPC and related technologies in their country. Their missions are (1) to develop and display a comprehensive and transparent map of HPC competences and institutions in their country, (2) act as a gateway for industry and academia toward national or international providers with suitable expertise or relevant projects, (3) collect HPC training offers in their country and present them in a central location together with international training offers collected by other NCCs and (4) foster the industrial uptake of HPC[4].
- There are fifteen European Centres of Excellence (CoEs) for the field of HPC application software: they bring together the European world-class knowledge and expertise in user-driven development, performance tools and programming models for HPC and co-design activities for real systems based on leading edge technologies[5].
- A framework partnership agreement targeting European low-power microprocessor technologies has been formed under the name of "EPI – the European Processor Initiative"[6] with twenty-eight partners collaborating in the conception and design of a general-purpose European low-power processor and a RISC-V based accelerator. The start-up company SiPearl[7] has been founded as a development and bring-to-market focal point for all EPI products.
- On a larger scale, in the "European Chip Act"[8], in October 2021, the EC announced the mobilisation of investments to reinforce Europe's capabilities in semiconductors to ensure future competitiveness and maintain its technological leadership and security of supply. The strategy underlying the EU Chips Act plans to mobilise more than €43 billion euros of public and private investments to finance technology leadership in research, design and manufacturing capacities up to 2030. In this context, the "Key Digital Technologies – Joint Undertaking (KDT-JU) kicked-off in late 2021 and is to become the single entity tasked with European chip development[9].

Digital Sovereignty is one of the key objectives of the EuroHPC Joint Undertaking of which ETP4HPC is one of the private members.

1.3

EuroHPC – the European HPC funding and development mechanism

The European Commission, with the support of the national research ministries, has identified HPC as a priority for Europe. It enjoys a special status among the technologies targeted by the European research programme, as its development is now managed and funded by one centralised body: **the EuroHPC Joint Undertaking (JU) which aims at the creation of world-class supercomputing in Europe**. As of writing this document, the JU consists of the European Union (represented by the Commission), selected Member States and Associated Countries that have chosen to become members of the Joint Undertaking, and three participating private partners, the European Technology Platform for High Performance Computing (ETP4HPC), the Big Data Value Association (BDVA) and the European Quantum Industry Consortium (QuIC). The EuroHPC Joint Undertaking is jointly funded with a budget of around EUR 7 billion for the period 2021-2027. The EuroHPC JU allows the European Union and the EuroHPC JU participating countries to coordinate their efforts and pool their resources. The objectives of the JU are to boost Europe's scientific excellence and industrial strength and support the digital transformation of its economy. The JU emphasises the importance of technological sovereignty. The JU deploys its strategy along the following pillars: infrastructure; federation of supercomputing services; technologies; applications; leadership in use and skills.

4. https://www.eurocc-access.eu
5. https://www.hpccoe.eu/eu-hpc-centres-of-excellence2/
6. https://www.european-processor-initiative.eu/
7. https://www.sipearl.com/index_en.php
8. https://ec.europa.eu/commission/presscorner/detail/en/qanda_22_730
9. https://www.kdt-ju.europa.eu/news/kdt-ju-become-chips-joint-undertaking

1.4

Global Challenges – the need for world-wide collaboration

The world is facing challenges at a scale that call for an increased global collaboration, in particular in the supply and then sharing of necessary resources, which include hardware and software infrastructure, data, models and expertise. For instance, from an economical and environmental point of view, it is debatable whether every region in the world should deploy local exascale infrastructure. The global challenges cannot be addressed using only local datasets. Also, the rapid nature of various events (natural disaster management, humanitarian and environmental crisis management, etc.) requires a high degree of agility in allocating resources (which is one of the applications of urgent computing) exploiting the global availability of resources. Overall, a global coordinated use, billing, and trust is required that spans across local datacentre boundaries.

Further use cases for global collaboration in HPC can be found in the WHO sustainable development goals[10]. For instance, global trade could benefit from techniques that promote a circular economy that manage the complete lifecycle of products across the global supply chain, i.e. from production, distribution, consumption, to reuse, repair and recycling of materials throughout their lifecycle. Possible solutions include the use of digital twins that represent the physical asset on the Internet in a secure way and facilitate the open exchange of resource information. For example, a product manufactured in China has a digital twin that contains all the information on the materials that it is made of. The product is subsequently shipped to Europe, where the digital twin can be used to ensure the product conforms to health and safety standards. Further examples include the use of Earth observation data to combat the detrimental effects of climate change and global access to health data without compromising a person's right to confidentiality.

Respecting local sovereignty aspirations in terms of infrastructure (local infrastructure use, possibly also local supply) and data (local data compliance & privacy rules, etc.) means that these must be accounted for in a global design that enables collaboration while satisfying such local requirements. The efficient, fair, and secure joint exploitation of local and distributed global resources introduces a number of technical challenges that are being explored in a number of use scenarios. This includes, for instance, the technical integration of infrastructure resource towards a hybrid multi-cloud model. Other dimensions include the creation and protection of local intellectual property, fair sharing according to the different levels of investments being made, all of which can profit from advances in technical solutions.

The federation of European supercomputing services is a new pillar in EuroHPC, which was introduced together with the new pillar of "international collaboration"[11]. These two pillars could serve as a foundation towards enabling a global federation of such services in the future whenever European and global actors agree to share and combine their resources. The definition of a such European federation architecture with interfaces, mechanisms for access, security and data protection, etc., would therefore profit from collaboration with global actors in order to achieve not only a European federation and autonomy but also influence the definition of a global federation.

A number of technical solutions are required to enable the realisation of such a vision. For example, this includes data and compute orchestration solutions in order to perform compute near data to meet performance and compliance requirements; zero-trust computing solutions, such as fully-homomorphic encryption and computing to ensure privacy of data and/or code; smart automated contracts, to expedite administrative aspects of resources sharing; and hybrid/multi-cloud including compliance and remote attestation solutions to guarantee efficient, secure and seamless execution/transfer/storage across hybrid and multi-cloud environments.

1.5

Ecosystem-level Recommendations

ETP4HPC's objective is a competitive, sustainable and sovereign HPC ecosystem in Europe. The technological research priorities identified in this SRA pave the path for the development of a competitive HPC technology value chain in Europe. However, these priorities require an adequate research project funding mechanism that would optimise the research effort within the work programme.

Instead of funding a wide range of diverse proposals, there should be a process for focusing the funding available on certain areas (e.g., investing in the development of domain-specific applications targeted for exascale-class systems). The Work Programme should thus favour those selected topics and allow the related projects to fail - taking risks should be encouraged. Furthermore, this lack of focus results in our inability to address complexity of various levels (HW, SW).

A flexible support tool is needed to nurture disruptions – i.e., break-through work which involves a lot of risk requires substantial, accessible funding. Also, we recommend a mechanism in which projects are object-driven and are regularly assessed based on their results and are subject to modifications if needed. We need to encourage our resources to overcome their resistance to new systems and new technologies (in particular, by industry) and project funding targeting risky and innovative topics would help reduce this barrier. We should also actively collaborate in the definition of standards and reference architectures. ■

10. https://www.who.int/europe/about-us/our-work/sustainable-development-goals
11. https://eurohpc-ju.europa.eu/system/files/2022-05/Decision%207.2022_%20Amendment%20EuroHPC%20Work%20Plan%20and%20Budget%202022.pdf

WHAT HAS CHANGED SINCE THE LAST SRA?

On the positive note, good R&D is coming out of HPC research programs. HPC Research program outcomes need to convert into viable and mature solutions quickly. Will private funding (Private Equity, Venture Capital, etc.) funding help to propel this by championing some SMEs in this area - for example HPC ISVs? However private funding supply seems inadequate. HPC projects need to have much stronger incentives and plans for commercialisation and/or usages in production environments at scale.

Also, industrial usage of HPC will make a very positive impact to European economy. Focus on commercial HPC and better engagement models (and incentives for HPC centres to engage with Industry) between HPC centres and the industry needs to exist. Large corporations will benefit financially through more collaborative approach using HPC centres rather than building everything in house. Education of using HPC as a tool is also very important here especially for start-ups trying to assess the use of HPC to optimise their manufacturing/production processes. We should educate the next generations of scientists/technologists to be knowledgeable about HPC technologies.

We need to bear in mind that HPC is not an isolated technology. This constitutes a great opportunity for our ecosystem as our products are in a position to aid other technologies. On the other hand, we do feel the consequences of the lack of a solid electronic components industry in Europe. While the alleviation of this market failure is beyond the capacities of the HPC ecosystem, it is important to emphasise its effect on HPC.

1.6

Process of writing this Strategic Research Agenda

This document is the work of over 130 experts who were either represented members of ETPHPC or were associated with our Association (see "Acknowledgements' in the final part of the document). Much input came from related activities such the TransContinuum Initiative and other exchanges within the European HPC community that either directly or indirectly affected the contents of this SRA (e.g., the "High Performance and Embedded Architecture and Compilation" (HiPEAC) project, the "Big Data Value Association" (BDVA/DAIRO) and the "Centres of Excellence for Computing Applications" (CoE) projects). In addition, feedback has been collected from twelve of the currently running R&I HPC technology projects.

Work on this SRA began in September 2021, and it was completed in July 2022. The experts working on this SRA were divided into working groups which represent the research domains of this SRA (below) and cross-cutting research clusters (below). Each of the working groups had multiple co-leaders. The working groups met in on-line meetings, and they worked to assemble their respective SRA 5 chapters. Half of the leaders of the SRA 5 Working Groups represent companies that are members of ETP4HPC.

SRA 5 Working Groups (Domains):

WG 1. System Architecture
WG 2. System Hardware Components
WG 3. System Software and Management
WG 4. Programming Environment
WG 5. IO and Storage
WG 6. Mathematical Methods & Algorithms
WG 7. Application co-design
WG 8. Centre to edge framework
WG 9. Unconventional HPC Architectures
WG 10. Quantum for HPC

SRA 5 Research Clusters:

- Sustainability
- HPC in the Digital Continuum
- HPC for Urgent Decision Making
- Federated HPC, Cloud and Data Infrastructures
- Heterogeneous High-Performance Computing

2

Projects resulting from the previous SRAs

There are a number of research projects in progress as a result of the research priorities defined in the previous SRA 4. The success of these projects proves the relevance of the findings included in that document. Here are a few examples:

- **IO-SEA**[12] (started in 2021) employs hierarchical storage including tape and specialised data nodes with non-volatile main memory to temporarily store data before handing off to more persistent hierarchical storage tiers. The IO-SEA project is based, besides others, on results which have been achieved at the end of the Sage2 project (2021), where memory style addressing of persistent storage was used to map objects from persistent layers to memory as well as Quality of Service overlaid on top of object storage. In-storage computing was also fully realised towards the end of Sage2.
- The **ADMIRE**[13] project (started in 2021) aspires to remove the barriers of multi-tier storage hierarchies by introducing appropriate interfaces and policies for managing an enhanced I/O stack. ADMIRE establishes control over the different storage levels by creating an active I/O stack that dynamically adjusts computation and storage requirements through intelligent global coordination, malleability of computation and I/O, and the scheduling of storage resources along all levels of the storage hierarchy. To achieve this, ADMIRE develops a software-defined framework based on the principles of scalable monitoring and control, separated control and data paths, and the orchestration of key system components and applications through embedded control points. ADMIRE builds on developments which have been introduced in previous European and national projects, like the ad-hoc file systems GekkoFS and Hercules, and combines them with proven technologies such as Lustre and Slurm.
- **DEEP-SEA**[14] (started in 2021) further develops critical elements of the software stack and programming environment to enable a more dynamic and malleable operation of the modular supercomputing architecture. DEEP-SEA focuses in particular on the support of heterogeneous compute resources and deep memory hierarchies, developing solutions driven by an application-requirement approach. In doing so, DEEP-SEA fully embraces the recommendations of the SRA working groups on system architecture, system software, programming environment, and application co-design. Furthermore, DEEP-SEA, as part of the SEA-projects family, works tightly together with IO-SEA to address the I/O related recommendations from the working group system software and AI and Data everywhere, as well as with RED-SEA to tackle the interconnect-related recommendations from the SRA4 working group system hardware components.
- **REGALE**[15] (started in 2021) is directly relevant to the priorities of ETP4HPC stated in SRA4. At the core of the project objectives lies the definition and implementation of an open system software architecture that will manage supercomputing resources efficiently with a primary focus on energy efficiency and effective operation under power constraints. As such, REAGLE works exactly on the intersection of ETP4HPC's «System Software and Management» research domain with the «Energy efficiency» cluster in SRA4. Additionally, REGALE also targets to assist in the development of complex, non-traditional HPC applications that involve ML pipelines and Big Data processing.
- **ACROSS**[16] (started in 2021) roots stem from well recognized research priorities described in ETP4HPC's SRA4. ACROSS focuses on integrating technological solutions for improving energy efficiency when running complex scientific and industry-oriented workflows on petascale and upcoming pre-exascale supercomputers. ACROSS aims at supporting these improvements through the large exploitation of GPUs, FPGAs, and other Machine/Deep Learning tailored architectures which allow to greatly speed up the execution of jobs and at investigating the potential benefit of neuromorphic architectures. Indeed, ACROSS has a primary focus on developing a multi-level workflow orchestration system. Aiming at supporting "AI everywhere" (a Research Cluster in SRA49, the project will provide technological solutions (e.g., FastML Engine) for workflows mixing HPC simulations, ML/DL models and HPDA. Other targeted SRA4 priorities are found in supporting data everywhere (through the use of HPDA, in-situ/in-transit DAMARIS solution) and HPC and digital continuum (use of Cloud resources). ■

12. https://iosea-project.eu
13. https://research.zdv.uni-mainz.de/research/admire/
14. https://www.deep-projects.eu
15. https://regale-project.eu/
16. https://www.across project.eu

3

Two new challenges for HPC

3.1

Sustainability – the next big thing

In December 2019, the European Commission released the Communication "The European Green Deal"[17] with an objective to implement the United Nation's 2030 Agenda and the sustainable development goals[18]. The goals of this Agenda include actions over the next fifteen years in areas of critical importance for humanity and the planet. Some of the keywords of these actions are "People", "Planet", "Prosperity" and "Peace".

Within "Planet", the objective is to protect the planet from degradation, including the application of sustainable consumption and production, sustainably managing its natural resources and taking urgent action in the area of climate change so that our planet can support the needs of the present and future generations.

On a regular basis, the UN assesses the science related to climate change and issues reports through its Intergovernmental Panel on Climate Change (IPCC)[19] . The IPCC was created to provide policymakers with regular scientific assessments on climate change, its implications and potential future risks, as well as to put forward adaptation and mitigation options. The sixth report is being developed and it is expected in September 2022.

Also, in 2015, during the UN Climate Change Conference (COP21[20]), the Paris Agreement was signed by 193 nations as a legally binding treaty to tackle climate change and its negative impacts. The agreement commits all signing countries to end the fossil fuel dependency as quickly as possible in order to keep the global temperature increase significantly under two degrees Celsius compared to the pre-industrial levels.

The overarching goal of all these actions is to tackle climate change and limit the resources consumption and global warming effects in order to "Meeting the demand of current generations without putting the demands of future generations at stake» (see chapter "Sustainability").

What role does HPC play in this context?

The extremely positive effects of the application of scientific modelling and simulation supported by HPC in terms of reducing the environmental impact are difficult to quantify (one example is the reduction of waste and energy in the automotive sector through numeric modelling of car crashes). It seems easier to tackle the negative environmental impact of using HPC compute- and data infrastructure and name the mitigation actions necessary.

A study on "Energy-efficient Cloud Computing Technologies and Policies for an Eco- friendly Cloud Market" issued by the Borderstep Institute[21] addresses the issue of the exponential growth in energy consumption due to the expansion of cloud services at a European level as well the development of voluntary and regulatory policy instruments. On-going digitisation and especially the increasing availability of cloud services are leading to a significant growth in data centre capacities. This growth is so strong that it has more than off-set the significant efficiency gains achieved at all levels (hardware, software, data centre infrastructure) and, consequently, the total energy consumption of data centres in Europe has risen.

While HPC data centre capacities constitute only a subset of the entire datacentre and cloud capacities, they are growing as well. Therefore actions need to reduce the CO_2 footprint, the use of rare material for the production of hardware components, increase the lifetime of systems, reduce electronic waste and introducing the scheme of a circular economy. There are also further issues such as HPC in a circular economy (cf the "Sustainability" research chapter).

The "Technical Research Priorities 2023 – 2026" chapter gives more details: the "Sustainability" Research Cluster addresses these challenges in a holistic way and each of the Research Domains details priorities in the context of sustainability to be taken into account in the upcoming R&D priorities.

The "Next actions" Chapter summarises the objectives of a new work-effort recently launched by ETP4HPC focusing on (1) a better understanding of the parameters and drivers for sustainability improvements, (2) generating an R&I roadmap focused on sustainability-driven research priorities and (3) acting as catalyst in an interaction between players inside and outside the HPC IT space (e.g. policy makers).

For many years, it has been a topic only addressed from a hardware point of view because of the impossibility to access real systems and production systems. More recently, the rise of the electricity cost and the requirements that datacentres make energy management a global responsibility, made this topic not only a research topic but also and foremost a requirement for the future of the planet itself. New programs must guarantee that energy efficiency is not only a keyword in the projects proposals: it must be part of the core of the projects. Energy management at the system and application levels have not been covered by many projects, and their applicability in real (production) systems has been very complicated. The reduction of the PUE (Power Usage Effectiveness) for datacentres, the system software and application optimization together will make the energy consumption adapted to real needs without over-consuming.

On top of the technical implementation and automation deployed to make infrastructures more sustainable, it is critical to empower end user to understand the ecologic impact of their decisions and educate the scientific & research community until behaviours evolve. Energy efficiency focus drives frugality, sustainability and ultimately eco-responsibility.

17. https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52019DC0640&from=EN
18. https://sdgs.un.org/2030agenda
19. https://www.ipcc.ch
20. https://unfccc.int/process-and-meetings/conferences/past-conferences/paris-climate-change-conference-november-2015/cop-21
21. https://www.borderstep.org

3.2

HPC in the Digital Continuum: a widely accepted paradigm

The previous SRA (4) introduced the concept of "TransContinuum digital infrastructures" to account for the convergence of data and compute capabilities: stimulated by massive deployments of compute and storage capabilities at the Edge, there is a need for a new "system design" to accommodate the ecosystem change to be expected in the coming decades (environmental and technological) and horizontally integrate the different actors representing the technology used in all domains of this continuum (cf the "HPC in the Digital Continuum" Research Cluster) .

As a result, in 2020, ETP4HPC created a framework for horizontal collaboration among stakeholders active in the continuum: a "TransContinuum Initiative" was established which includes eight technology-oriented associations:

- **ETP4HPC** - European technology Platform for High Performance Computing
- **BDVA/DAIRO** - Big Data Value Association
- **ECSO** - European Cyber Security Organisation
- **AIOTI** – the Alliance of Internet Of Things Innovation
- **6G IA** – 6G Infrastructure Association
- **EPoSS** – the European Technology Platform on Smart Systems Integration
- **EU-MATHS-IN** - European Service Network of Mathematics for Industry and Innovation
- **HiPEAC** - High Performance Embedded Architecture and Compilation

The scope of this collaboration was jointly formulated in a TCI Vision paper[23] :

- Identify priorities and recommendation for European R&I work programs - Jointly we will elaborate recommendations for R&D to be carried out in EU-funded work programs addressing challenges in the digital continuum. The recommendations will cover challenges in technological (hardware and software) functionality, interoperability, and APIs. New standards, best practices, methodologies and project-type related suggestions will also be generated. Applications deployed in the digital continuum are addressed wherever needed.
- Joint work with European R&I funding agencies and R&D programs (e.g., JUs, Missions) - The recommendations will be presented to EU-funding entities like Joint Undertakings (JUs) and applicable programs in the MFF 2021-2027. TCI representatives will be available for presenting and explaining the recommendations as well to discuss any possible further analysis and elaborations.

Transcontinuum Initiative (TCI)

PARTICIPATING ORGANISATIONS

A CONTINUOUS DYNAMIC WORKFLOW

Between
Smart Sensors and IOT devices

and
HPC / cloud centers

passing through
Edge platforms & subsystems

as well as
Smart Networks and Services

executing
Simulation & Modelling, Big Data Analytics and ML*

based on
Math. Methods & Algorithms incl. MSODE**

pervasively augmented by
Artificial Intelligence

protected and secured by
Cybersecurity

back to
Cyber-Physical Systems,

all based on
Data and compute platforms (hw and sw)

* ML: Machine Learning
** MSODE: Modelling, Simulation and Optimization in Data-rich Environment

Original courtesy HiPEAC

Figure 2: The TransContinuum Initiative components and members (TCI Vision)[22].

21. https://www.etp4hpc.eu/tci-vision.html
22. https://www.etp4hpc.eu/tci-vision.html

- Generate and foster an interdisciplinary network of experts - We look forward to a lively exchange of news about EC work programs, calls and related events, events of partner organisations and potentially joint activities. We will jointly analyse new industrial and scientific use cases to better understand the challenges presented. On one side, this is a pre-requisite for any R&I recommendations, on the other side it facilitates the forming of interdisciplinary consortia for upcoming calls.
- Contribute to SR(I)As and other partners' documents - Based on the results of the joint work mentioned above, contributions to the Research and Innovation Agendas or any other road mapping documents issued by participating partners will be offered.
- Contribute to the five Horizon Europe missions - One of the first pragmatic actions will be to design the contribution of the Digital Twin enabler to the Horizon Europe missions (adaptation to climate change including societal transformation, cancer, healthy oceans, seas coastal and inland waters, climate-neutral and smart cities, soil health and food.) These missions will need digital technologies to achieve their respective goals and Digital Twins should be one of the key elements.

In 2021, the TCI partners analysed a selection of Digital Twin (DT) use cases and issues two TCI white papers related to these use cases. This work was a good preparation for further similar engagement as technical consultants in important EU projects such as Destination Earth.

Digital Twins, originally introduced by NASA in 2012, have gained an enormous traction in industrial production control processes whereby an online digital replica in form of a numerical model of the actual process receives real-time information from observations of the physical system (the "physical twin") to optimise production, predict and optimise the performance during the lifetime of the physical system or process. For, example predicting malfunction and the need for maintenance ("preventive maintenance") is a key goal of the Digital Twin approach.

Classically, this concept is highly relevant for industries closely linked to digital technologies and has been heavily used in R&D since years. Recent technical breakthroughs allow to package easy-to-reuse real-time versions of it, opening up novel sensing or control concepts during operation of the physical twin - so called "Executable Digital Twins". These are key technologies applied in the Industrial Internet of Things, e.g., they enable milling robots to achieve unprecedented accuracies or measuring motor temperatures at locations not accessible by sensors.

The DT concept can equally be applied to use cases outside industrial production such as the development of a Digital Twin of the Earth System in order to enable predictions of extreme weather events with unprecedented accuracy. At its heart is the ability to constrain the model with available observations in the model initialisation step. Faithful representation of the current state of the system in a virtual environment is a prerequisite to

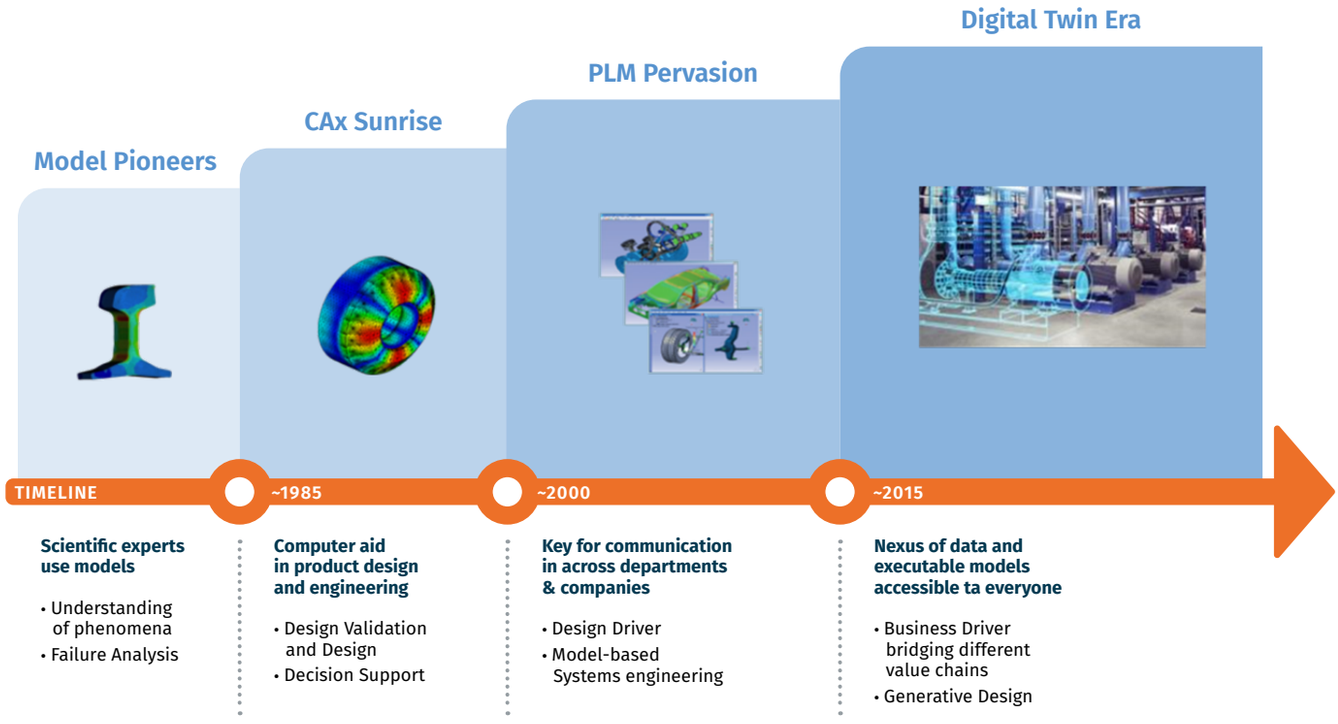

Source: HARTMANN, Dirk; VAN DER AUWERAER, Herman. Digital twins. In: Progress in Industrial Mathematics: Success Stories. Springer, Cham, 2021

Figure 3: Simulation is evolving from a troubleshooting tool to key business driver in the form of digital twins.

the skilful numerical weather prediction. Both tasks not only have an extreme scale computing and data footprint, but more crucially their implementation relies on the entire landscape of the digital continuum.

Even though High-Performance Computing (HPC) and Big Data (BD) handling are at the core of DTs, only the full integration of all components of the digital continuum will allow to exploit the full potential of both HPC and BD and generate effective DT capabilities for all use cases. Digital Twins are the next wave in simulation technologies, merging high performance simulation with big data and artificial intelligence technologies. It is predicted that companies who invest in digital twin technology will see a 30% improvement in cycle times of critical processes.

Example: Extremes Prediction use case

Natural hazards represent some of the most important socio-economic challenges our society is facing in the next decades. Natural hazards have caused over 1 million fatalities and over 3 trillion Euros economic loss worldwide in the last 20 years, and this trend is increasing given drastically rising resource demands and population growth. Apart from the impact of natural hazards on Europe itself, the increasing stress on global resources will enhance the political pressure on Europe through yet unprecedented levels of migration.

Dealing responsibly with extreme events does not only require a drastic change in the ways our society is solving its energy and population crises. It also requires a new capability of using present and future information on the Earth-system to reliably predict the occurrence and impact of such events. A breakthrough of Europe's prediction capability can be made manifest through science-technology solutions delivering yet unprecedented levels of predictive accuracy with real value for a society.

These science-technology solution includes the entire loop from (1) basic Earth-system science, fundamental to developing (2) enhanced prediction models, that need to be combined with (3) the vast range of Earth observation data ranging from advanced satellite instruments to commodity devices available through the internet of things, exploiting (4) extreme-scale computing, big-data handling, high-performance data analytics and artificial intelligence technologies, for feeding (5) impact models that translate scientific data into information close to the real users responsible for critical infrastructures, hydrology and water, energy, food and agriculture, health and disaster management. ■

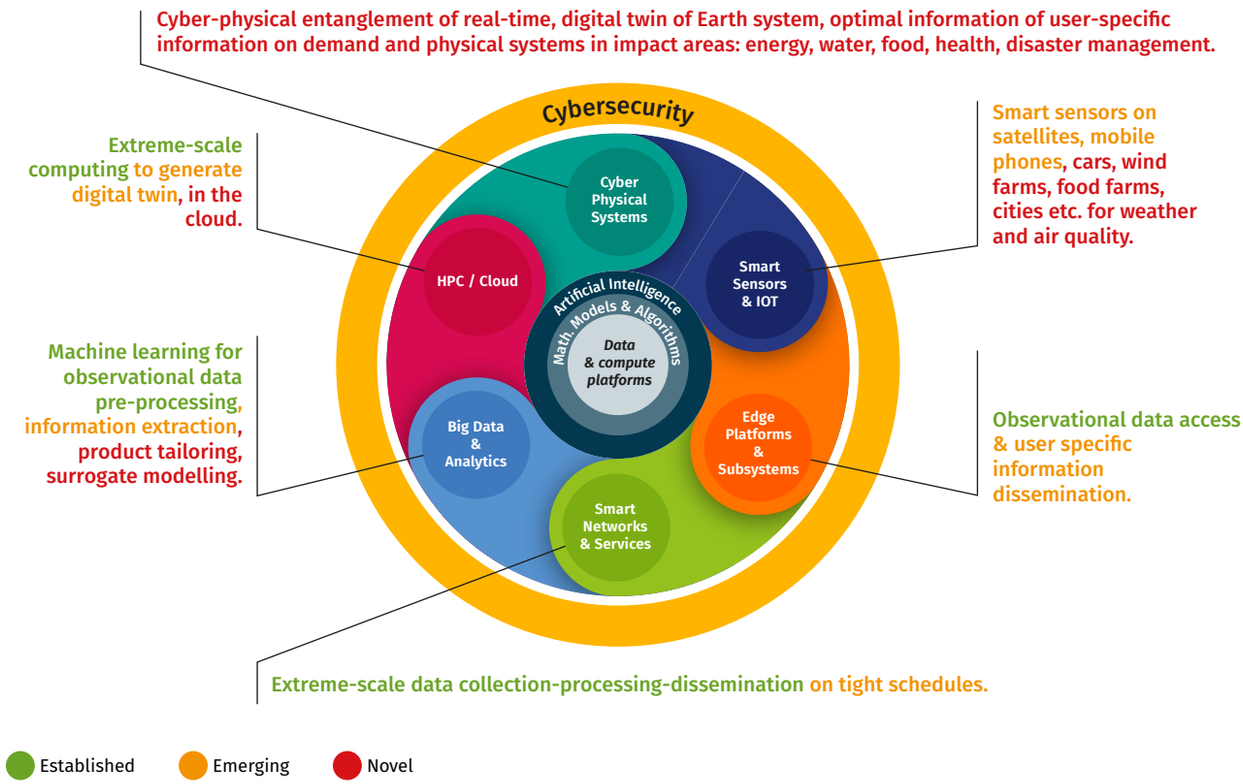

Original version courtesy HiPEAC

Figure 4: Main elements of digital continuum and relevance for extremes prediction use case including readiness of technologies (green = established in production, but not optimal in performance; yellow = research, not in production mode; red = novel and unexplored).

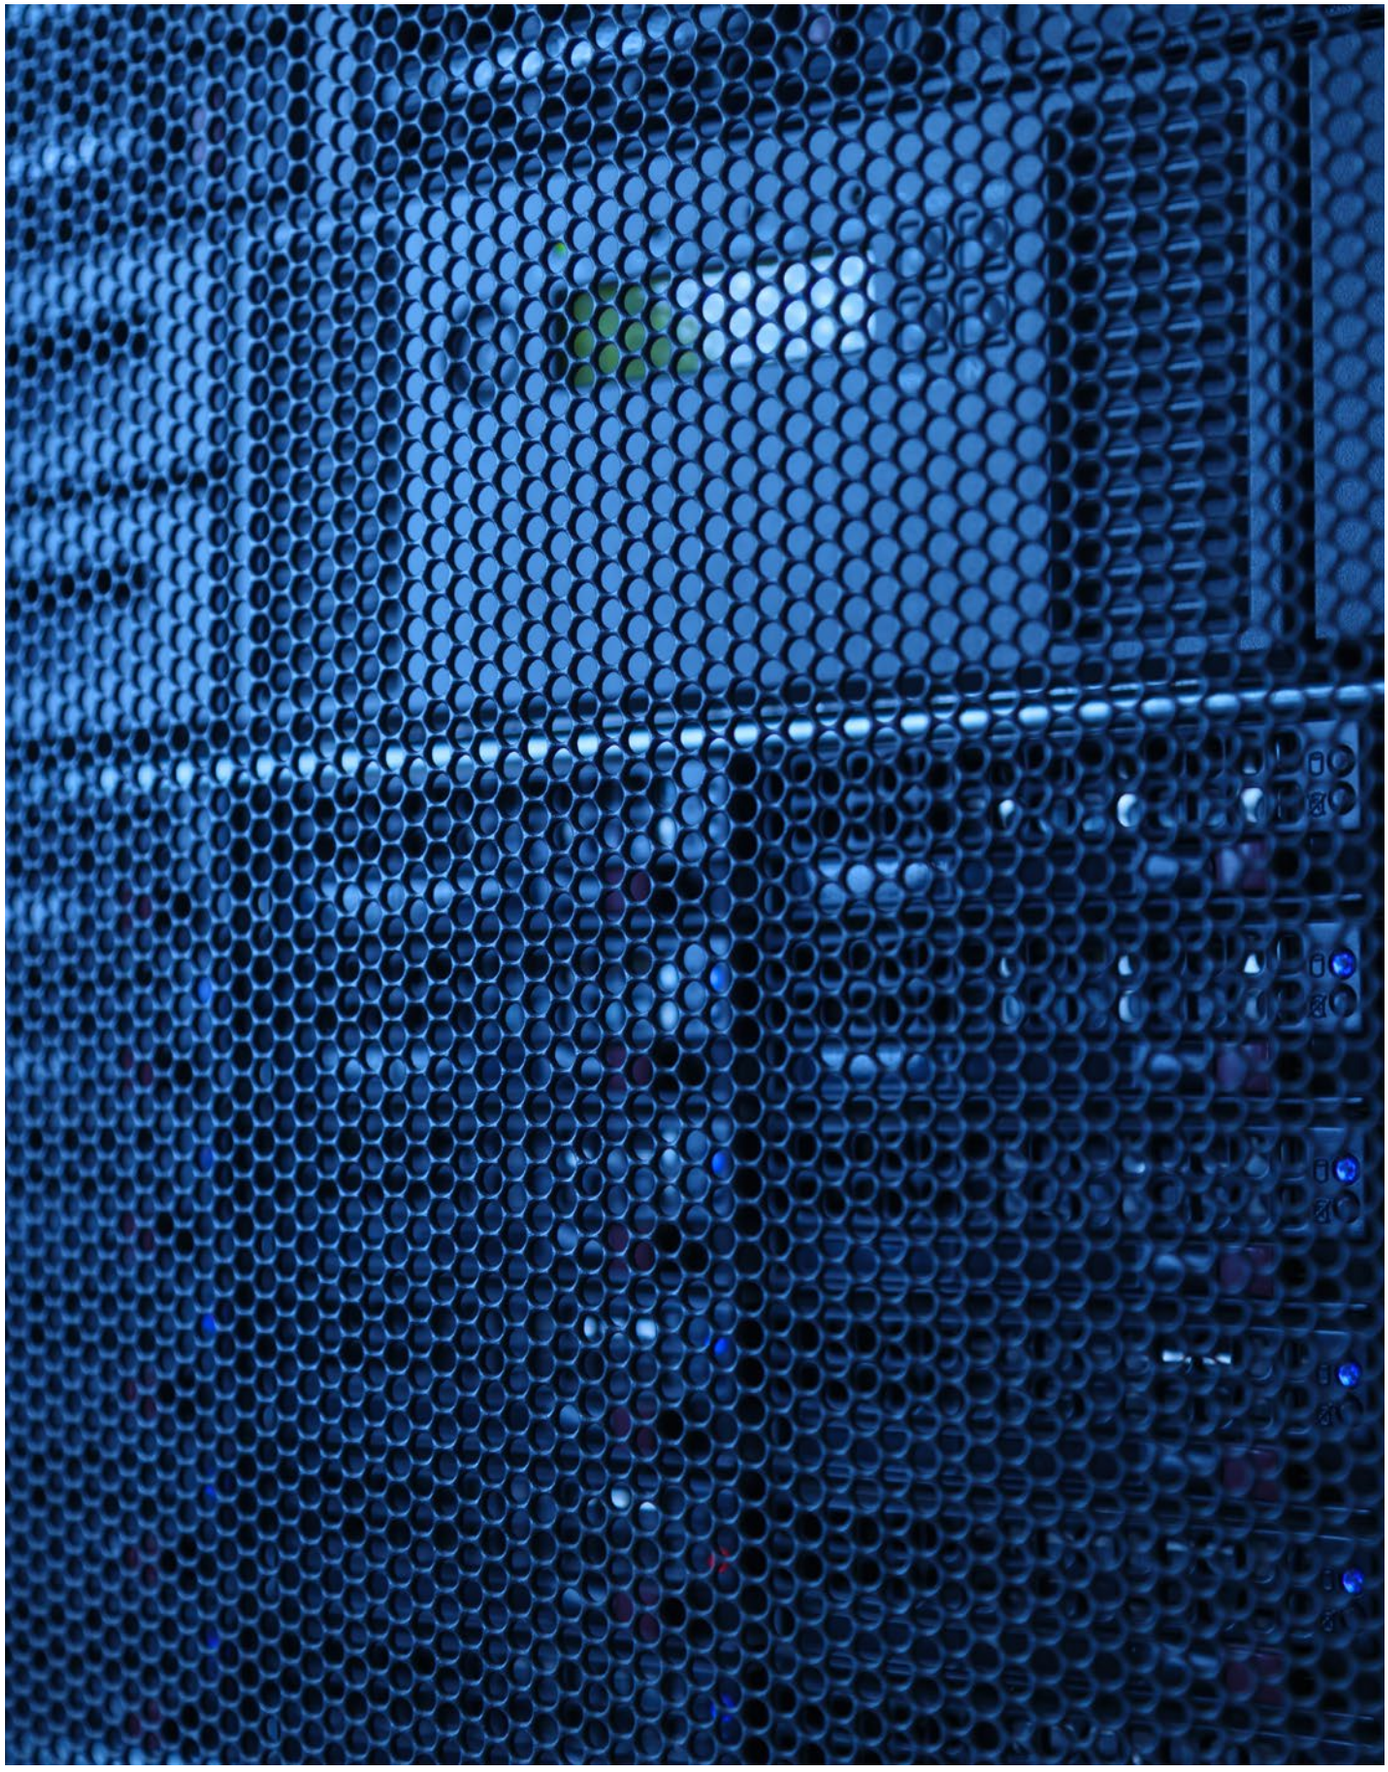

4 Industrial Use of HPC

CREDIT AND DISCLAIMER
This chapter summarises the topics raised in market studies performed by the HPC-GIG project[24], regarding industrial use of HPC resources/services in Europe, and where possible relates them to the corresponding chapters of SRA5. Courtesy of HPC-GIG.

4.1

Context – Collecting intelligence on actual & potential industrial use of HPC in Europe

The European Commission (EC) and member states have a long history of co-funding research and innovation efforts in the HPC field. Scientific use and the needs of European scientists have shaped these research and innovation activities. Industry across Europe has taken up HPC, and in some fields, HPC is now an integral element of cutting-edge product development or service provision.

The EC and EuroHPC JU have funded valuable industry-targeted innovation actions around HPC (such as the Fortissimo family of projects), yet today, no comprehensive, detailed analysis of actual HPC use across industrial and commerce in Europe is available, and knowledge about potential use cases is even more limited. There is a clear need to identify actual and potential industrial/commercial HPC users, to analyse their specific requirements in the technical, legal and business space, and to understand the currently available HPC resources and services. The HPC-GIG project is working to provide such an analysis, relying on market studies conducted by impartial, well-established third parties.

4.2

Principal HPC usage scenarios

Significant industrial use of HPC has historically started with R&D collaborations involving public HPC centres and proceeded to rely on the use of in-house HPC resources. Today, we see four principal industrial usage scenarios: usage of publicly funded HPC centres, use of remote HPC platform or service offerings owned and offered by Cloud service providers (CSPs) or private HPC providers, utilisation of private Cloud resources which are part of the customer's IT infrastructure, and of course reliance on in-house systems owned by the customer. Each of these have different characteristics regarding the key aspects of commercial use sketched below. Cloud use of HPC resources is a topic of WG8 , and it has a major influence on WG3[25].

24. The analysis reproduced here was supported by the European Commission as part of the Horizon 2020 project HPC-GIG (grant no. 824151).
25. The SRA 5 Working Groups are: WG 1 - System Architecture, WG 2 - System Hardware Components, WG 3 - System Software and Management, WG 4 - Programming Environment, WG 5 - IO and Storage, WG 6 - Mathematical Methods & Algorithms, WG 7 - Application co-design, WG 8 - Centre to edge framework, WG 9 - Unconventional HPC Architectures and WG 10 - Quantum for HPC.

4.3

Technical requirements of commercial/industrial users

Industrial HPC end-users rely on a large set of core applications, which perform simulations, optimisations, complex data analysis, and increasingly AI/ML processing. Such core applications cover a huge scope: simulation of internal and external flows (CFD), combustion and other chemical reactions, structural analysis (often using finite element codes), materials sciences, molecular dynamics, quantum effects, circuit simulations, financial risk analysis, docking and protein folding, *-omics, traffic and crowd simulations, weather, large-scale graph analytics, discrete optimisation and others. While the algorithms used in these do not differ widely from scientific applications, differences are caused by a different scale and complexity of problems, the frequent need for ensemble scaling to drive optimisations, and the role of closed-source, licensed applications. These differences are relevant for WG3, WP4, and WG6, and they can also influence the choice of system technology & design.

Often, thorough validation of the application results is required, to be performed by the application developers/vendors, or sufficiently tight bounds have to be established using uncertainty quantification (UQ) techniques. This is a topic for WG4 and WG6.

Commercial users store and control significant amounts of data with high commercial value, plus for many use cases personal data in the sense of the GDPR. Such data commonly resides in the protected IT infrastructure of the customer, and it has to be transferred onto the IT infrastructure of an HPC provider, accessed and modified via HPC applications and workflows there, and transferred back (results) or purged (input data which is no longer useful). It is critical that technical measures to protect data from unauthorised access and modification are in place throughout the complete data processing chain. In addition, effort and costs for data movement have to be kept under control. WG5 and WG8 look at these topics.

Industrial companies all operate their own IT infrastructure, which often relies on Microsoft Windows operating system and associated applications. More importantly, the industrial work environment relies on high-level frameworks and workflows, which match the business or engineering processes in place, and provide end-users with easy-to use interfaces. Such frameworks and the corresponding high-level workflows have to be supported by any HPC usage model. This is considered by WG3 and WG4.

Authentication, authorisation and accounting support to non-personal accounts may be required, and strong monitoring and security mechanisms are required for any operation that touches confidential data. Accounting and billing have to comply with the business requirements of the end users, and accommodate paying per use. WG3[26] covers these aspects.

4.4

Legal requirements of commercial/industrial users and HPC providers

In most cases, commercial companies operate with data subject to the European GDPR rules, and in almost all cases they derive value from confidential information or trade secrets. GDPR requires specific organisational assignments and processes in place across the full processing chain (for example, to implement notification requirements), and requires state-of-the art data protection measures to be in place. External HPC providers have to collaborate with the customers in assuring compliance, and for trade secrets, specific non-disclosure agreements with potential financial penalties are required.

Commercial use of HPC requires contracts including specific availability and quality of service guarantees, pricing and payment terms and rules on liability. This might be a challenge for public providers operating under particular statutory frameworks.

Co-design activities with industrial end-users will require a sufficient coverage of closed-source application codes, agreements on how to limit the exposure of trade secrets embodied in algorithms or codes, and an apportioning of intellectual property rights created during such projects. This is relevant for WG7.

26. The SRA 5 Working Groups are: WG 1 - System Architecture, WG 2 - System Hardware Components, WG 3 - System Software and Management, WG 4 - Programming Environment, WG 5 - IO and Storage, WG 6 - Mathematical Methods & Algorithms, WG 7 - Application co-design, , WG 8 - Centre to edge framework, WG 9 - Unconventional HPC Architectures and WG 10 - Quantum for HPC.

4.5

Business aspects of industrial and commercial HPC use

Providers of HPC resources or services have to protect confidential, non-personal data of industrial/commercial customers (trade secrets), usually embodying intellectual property, or related to the specific use of HPC systems. As mentioned above, breaking such agreements can expose the provider to litigation. On the technical side, encryption of data in transport and at rest will be required, potentially complemented by encryption in memory. Customers may also request additional access rights protection and detailed access logs. WG3 and WG5 can contribute here.

License conditions and pricing for closed-source, commercial applications is a key factor – ideally, license models would enable providers to amortise such costs across multiple customers or be based on actual use of applications rather than on maximum job/system sizes.

Commercial use of HPC as part of development or business processes requires availability of resources according to business needs and resulting negotiated schedules. This has to be balanced against the need of HPC providers to maximise continuous utilisation of their resources. As technical aspect, provision of priorities and pre-emption of running jobs of lower priority may be applicable here. This is one of the topics of WG3.

Regarding support of industrial users (meaning timely and effective assistance in case of malfunctions encountered by the customer which are not related to more general availability problems), key aspects are the scope of support (which frameworks and applications are covered) and the reaction time. Integration of third parties like application developers might be required, and times to resolution will be more critical than with most scientific uses.

Pricing models for HPC resources/services can range from paying per actual use to buying a reservation for resources for a specific time range up front. Combinations of up-front subscriptions or payments with additional (reduced) per-use billings are also possible. The latter helps providers to plan for the likely amount of resources required and help financing these.

Finally, commercial users are used to automatic methods of specifying and negotiating quality of service and pricing. The former refers to automation of service level agreements (SLAs), and the latter to an expectation to support marketplaces or brokering of resources/services from several providers. ■

5

On the path to Exascale – Review of EuroHPC Extreme-scale Prototypes

This chapter introduces the three prototype projects funded by EuroHPC. They aim to showcase and demonstrate the readiness of the technology derived from the European R&D projects. The importance of technology integration projects at a reasonable scale was emphasised already in SRA 1 in 2013. Later, starting in 2014, ETP4HPC continues to issue proposals for so-called "Extreme Scale Demonstrators", which at that time were not based on the processors or accelerators delivered by the EPI project. EUPEX and EUPILOT implement a complete HPC stack on top of EPI processors/accelerators and include an extensive validation and benchmarking phase. HPQS integrates quantum technology with high-end HPC systems and aims to deliver a pan-European quantum-HPC infrastructure integrating Tier-0 systems with various quantum hardware technologies.

5.1

EUPEX

EUPEX
European Pilot for Exascale

The EUPEX pilot brings together academic and commercial stakeholders to co-design a **European modular Exascale-ready pilot system**. Together, they will deploy a pilot hardware and software platform integrating the full spectrum of European technologies, and will demonstrate the readiness and scalability of these technologies, and particularly of the Modular Supercomputing Architecture (MSA), towards Exascale.

EUPEX's ambition is to **support actively the European industrial ecosystem around HPC**, as well as to **prepare applications** and users to efficiently exploit future European exascale supercomputers.

EUPEX will **leverage the best of the assets developed in previous H2020 projects** to design, build and validate the **first HPC platform integrating the full spectrum of European HPC technologies**, from architecture, processors and interconnect to system software and applications. Some examples of the European technologies integrated in EUPEX are:

- the architecture: OpenSequana, Modular Supercomputer Architecture
- processors: SiPearl's Rhea processor
- the network interconnect: Atos's BXI
- the system software: ParTec's ParaStation Modulo
- development tools: Bull energy + dynamic power optimizers, ScoreP/Scalasca, PMIx, MERI, COUNTDOWN, etc.
- applications: CINECA's LiGen, ECMWF's Integrated Forecasting System (IFS)

Modularity at the heart

The EUPEX pilot system will be **modular**, thanks to the OpenSequana-compliant hardware platform and the matching HPC software ecosystem implementing the Modular Supercomputing Architecture, so as to be able to integrate and manage efficiently a **variety of hardware modules, including upcoming architectures, and to handle heterogeneous workflows**.

European Pilot for Exascale (EUPEX) started on 1 January 2022 and will run for 4 years with a budget of up to €40.7 million provided by the EU and the Participating States of the EuroHPC JU. It is the largest of the three EuroHPC pilots funded under call EuroHPC-2020-01.

EUPEX is coordinated by Atos and involves a balanced consortium of 17 European academic and commercial stakeholders from 7 countries:

- **Atos (Bull SAS) - France**
- **Forschungszentrum Jülich GmbH - Germany**
- **CEA - France**
- **GENCI - France**
- **CINECA - Italy**
- **E4 - Italy**
- **FORTH - Greece**
- **CINI - Italy**
- **ECMWF - International**
- **IT4I – Czech Republic**
- **University of Zagreb, Faculty of Electrical Engineering and Computing - Croatia**
- **ParTec - Germany**
- **EXAPSYS - Greece**
- **INGV - Italy**
- **Goethe University Frankfurt - Germany**
- **SECO - Italy**
- **Cybeletech- France**

A true proof-of-concept

The EUPEX pilot system will be **large enough to be a proof of concept** for a modular architecture relying on European technologies, and in particular on the European processor technology (EPI), and to demonstrate the Exascale readiness of the applications selected for co-design.

Production-grade

The technical choices of the EUPEX pilot were guided by the **maturity** of the European solutions available, to guarantee that the resulting system will be truly **production-grade** and deployable in a near future.

Addressing software and applications

EUPEX is not just about building a hardware prototype: the consortium will devote as much effort to the system software stack and to the applications as to hardware.

A set of applications selected from a large variety of domains (climate and weather forecast, biology and health, remote sensing analysis, material science, astrophysics, engineering, seismology) will be optimised for the target architecture, used for benchmarking on the EUPEX Pilot, and analysed to issue recommendations for exploitation on future European Exascale systems.

European industrial sovereignty

The EUPEX consortium involves a balanced mix of established European technology suppliers and vendors, recognized research organisations and universities, European-scale computing centres, and application owners. This alliance of public and private sector stakeholders is a guarantee that the innovative research achieved by the consortium will quickly translate into European-grown industrialised solutions that will reduce Europe's dependence on foreign supercomputing technologies.

In a context where the European Union is striving to regain its digital sovereignty, turning from a simple user of computer technologies into a self-reliant player capable of designing, developing, manufacturing and marketing all the necessary technological bricks, it is an essential asset to master extreme-scale supercomputing technologies, an asset whose success will percolate through the entire computing continuum.

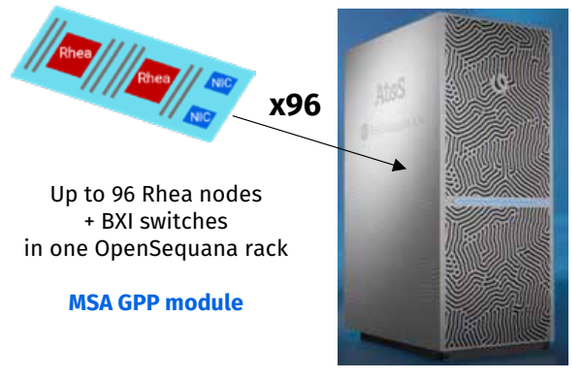

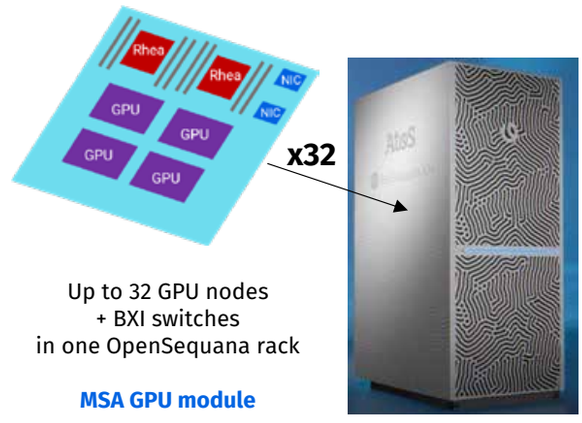

5.2

EUPilot

The *European PILOT (EUPILOT)* project's main goal is to develop and deploy an end-to-end demonstrator of HPC and ML accelerators maximising the use of European technology, open source and open standards. The effort covers the entire process from chip design and fabrication to datacentre rack deployment.

EUPILOT is a combination of accelerator designs, based on the open source RISC-V instruction set architecture (ISA), as well as High Performance Computing (HPC) and High Performance Data Analytics (HPDA) (Machine Learning (ML), Deep Learning (DL) and Artificial Intelligence (AI)) applications and system software - all of it integrated into open standard systems and deployed in advanced cooling and power delivery solutions in liquid immersion racks.

EUPILOT brings together an ensemble of partners to leverage and extend pre-existing Intellectual Property (IP), combined with new IP, that can be used as building blocks for future HPC systems, both for traditional as well as emerging application domains.

At the heart of this project are two European HPC/HPDA chips. EUPILOT will produce three tapeouts (fabricate chips), the first one being a test chip to validate the use of the 12nm silicon technology node, the second and third ones are concurrent, one containing a vector accelerator with up to 16 cores, and the other containing a machine learning and stencil accelerator with up to 8 cores. Each type of accelerator is integrated on a printed circuit board (PCB) accelerator module, as illustrated in Diagram 1.

EUPILOT's overall goal is to demonstrate European HPC accelerator technology, based on open source and open standards, with a full stack ecosystem (software, hardware, and integrated system).

Extend Open Source to Hardware for HPC

EUPILOT technology is based on the open source RISC-V ISA and will feature high performance computing (VEC) and data analytics (MLS) accelerators coupled to a high performance, low energy general purpose platform, like SiPearl's Rhea, as well as any other CPU platforms with PCIe. This is a major next step in the direction of an open European software/hardware ecosystem. EUPILOT is leading the way to the first demonstration of technology independence.

Software/hardware Co-design, System Software and HPC/HPDA Applications

EUPILOT will also demonstrate the next generation of software tools, compiler technologies, middleware, and drivers to support a co-designed RISC-V HPC and HPDA accelerator. EUPILOT will target key applications in HPC (GROMACS and EC-Earth) and HPDA (BERT, YOLO, and industrial video anomaly detection) to drive the co-design, with a flexible, coherent, and high-performance connection to a host CPU via PCIE with CXL.

System Integration Innovations

Power density and scaling continue to be a challenge for the next generation exascale systems. EUPILOT will demonstrate advanced liquid immersion cooling technology required for the power densities resulting from increased scale. It will feature integrated power modules and system management, supporting ultra-efficient rack power densities far surpassing today's solutions.

European Collaboration

EUPILOT partners will leverage their existing IP from multiple European projects such as European Processor Initiative (EPI SGA1/2), Highly Efficient, Eco-friendly Immersion Cooling for Data Centres, MareNostrum Experimental Exascale Platform (MEEP), *e*Processor, several DEEP projects, POP2 CoE, MontBlanc2020, EuroEXA, and ExaNeSt, extending their capabilities, and improving their Technical Readiness Level (TRL). EUPILOT will collaborate with the EUPEX pilot, which showcases the general-purpose processor Rhea from SiPearl.

Accelerating Exploitation with Cutting-edge Methodologies and Tech

The EUPILOT extends, enhances, develops, and delivers silicon-proven IP at higher TRL, in a leading-edge silicon geometry. This will provide a faster time-to-market, with higher potential for European exploitation.

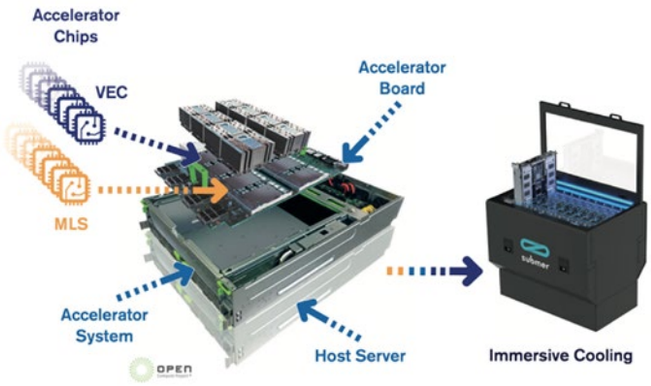

PROJECT DETAILS

https://eupilot.eu

Duration: 42 months (12/2021 - 05/2025)

Partners (19): BSC (coordinating), 2CRSI, CEA, CHALMERS, CINI, ETHZ, EXAPSYS, EXTOLL, FORTH, FZJ, Fraunhofer, KTH, LEO, SMD, SUBMER, TUBITAK, TUK, UNIBO, UNIZG-FER

The European PILOT project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No. 101034126. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Spain, Italy, Switzerland, Germany, France, Greece, Sweden, Croatia and Turkey.

5.3

HPCQS – High-Performance computer and quantum simulator hybrid

Project duration: December 1, 2021 – November 30, 2025
Budget: 12 Mio euro, provided by the EuroHPC JU and the participating member states in equal parts

The overall goal of the HPCQS project is to prepare European research, industry and society for the use and federal operation of quantum computers and simulators. HPCQS aims at developing, deploying and coordinating a European federated infrastructure at European level, integrating a quantum simulator (QS) of circa 100+ interacting quantum units in the HPC systems of the supercomputer centers Jülich Supercomputing Centre (JSC) at Forschungszentrum Jülich (FZJ) and GENCI/CEA. This federation is the first step towards a pan-European quantum-HPC infrastructure integrating Tier-0 systems with various quantum hardware technologies. The infrastructure will be accessible via the cloud to public and private European users on a non-commercial basis.

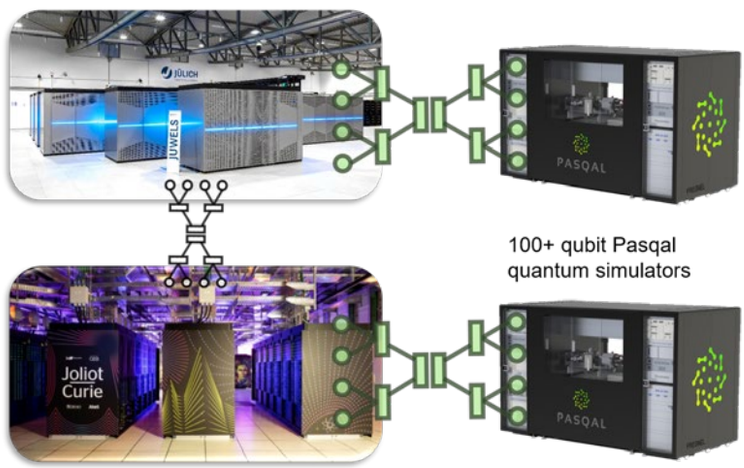

The integration challenge is met through two major technological developments within HPCQS supported by a consequent co-design approach via use cases. One important technical component is a production-ready programming environment and middleware for federated QSs based on the Quantum Learning Machine QLM® by ATOS. The programming environment is complemented by technical libraries that allow users to translate use cases into quantum programs without having to deal with low-level instructions. It will also include a set of compilers that will convert the hardware-agnostic instructions generated by the technical libraries into hardware-compliant instructions, with a special emphasis on variational algorithms, a class of hybrid quantum-classical algorithms used in many of the use cases.

A second essential technical component serves to closely integrate the QS into the high-end HPC systems. To this end, the interconnection between the classical HPC supercomputer and the QS has to be developed. HPCQS will employ the concept of the Modular Supercomputing Architecture (MSA) for the lowest latency integration of the QS. The MSA has been developed in the series of EU-funded DEEP projects and is based on ParTec's ParaStation Modulo middleware suite.

HPCQS concentrates on a few selected use cases to develop its infrastructure in the co-design process. Training material, user guides and courses will be provided for both scientific and industrial user needs and supporting the access and use of the HPCQS infrastructure.

The HPCQS infrastructure is designed to be open to new quantum hardware architectures and to new sites, making HPCQS an unprecedentedly sovereign European offering in terms of HPC/quantum coupling, architectures supported, number of sites, diversity of users and uses, and integrated services offered. Together with JUNIQ (Jülich UNified Infrastructure for Quantum computing), HPCQS is the seed for the European Quantum Computing and Simulation infrastructure (EuroQCS), an infrastructure described in the Strategic Research Agenda of the Quantum Flagship.

To achieve these goals, HPCQS brings together leading quantum and supercomputer experts from science and industry from six European countries. HPCQS will become an incubator for practical quantum HPC hybrid computing that is unique in the world. The five participating European HPC centers JSC at FZJ (Germany), GENCI/CEA (France), Barcelona Supercomputing Center (Spain), CINECA (Italy) and NUIG-ICHEC (Ireland) collaborate closely with the technological partners Atos (France), ParTec (Germany), FlySight (Italy), ParityQC (Austria), and the research partners CEA (France), CNRS (France), Inria (France), Centrale Supélec (France), Sorbonne Université (France), CNR (Italy), the University of Innsbruck (Austria), Fraunhofer IAF (Germany) and the SME partner Eurice (Germany). Prof. Dr. Kristel Michielsen from JSC coordinates HPCQS. ■

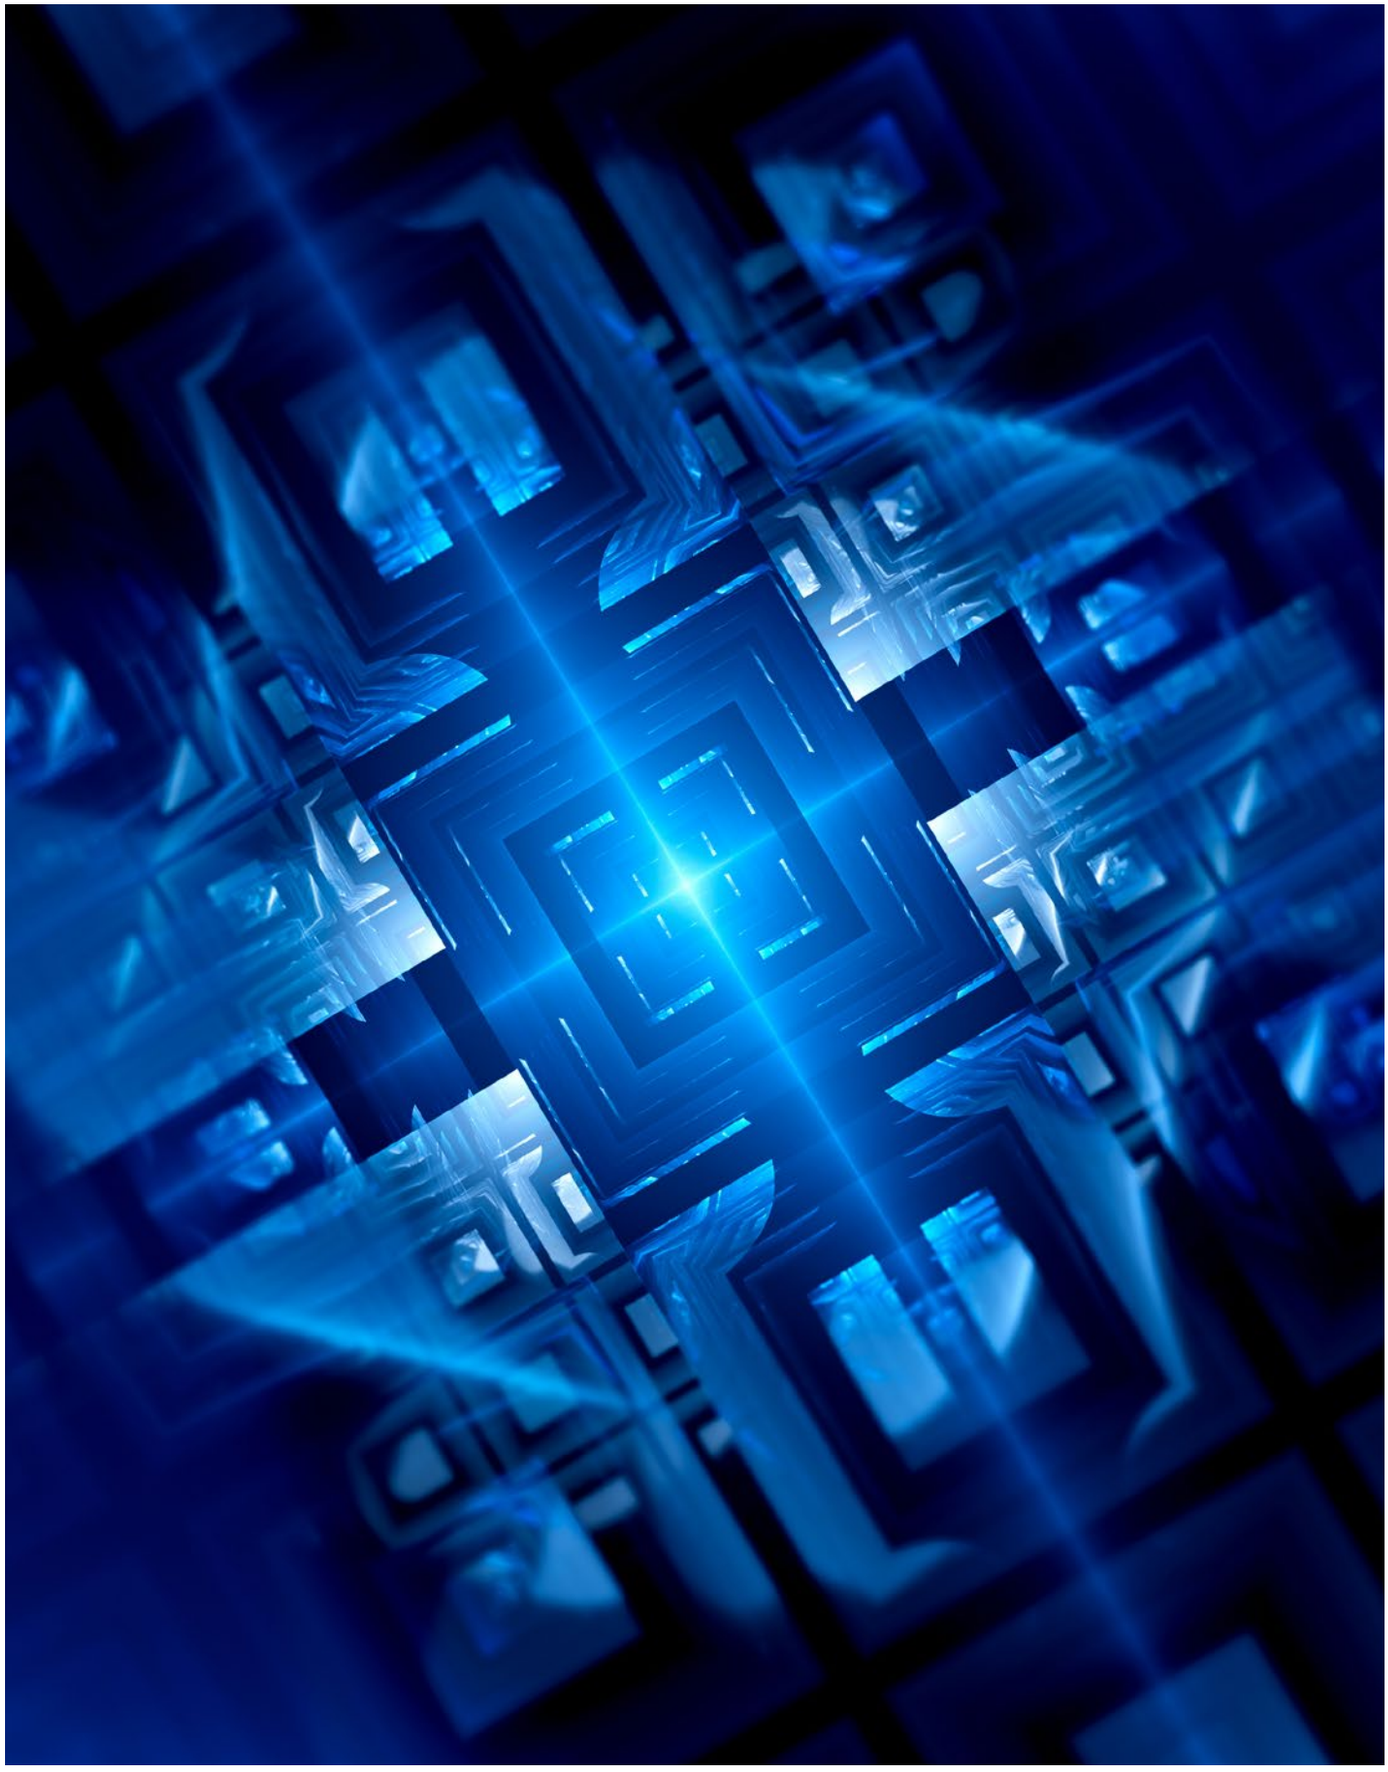

6

Technical Research Priorities 2023 – 2026

Our model of research priorities has two dimensions: **Research Domains and Research Clusters.**

Research Domains describe the essential layers and elements of a HPC functional stack. While "System Architecture" applies a holistic view of all elements of the stack, "System Hardware Components" covers the lowest hardware level. "Unconventional HPC Architectures" addresses the deployment of new, non-mainstream components to gain higher efficiencies and exploit the full system-level acceleration potential with specific workloads. "System Software and Management" focuses on all aspects of operating and managing the underlying hardware facilities. "Programming Environment" adds a user's (programmers) view to using the HPC hardware and system software infrastructure. "I/O and Storage" covers the space of feeding the compute nodes with data and storing data. "Mathematical Methods and Algorithms" look at the backbone of any level of software used in HPC systems (not to be confused with the algorithms used in the application layer). "Application Co-design" offers an application writer's input into the needs for the implementation of next generation HPC infrastructure all together. "Centre-to-Edge Framework" covers all aspects of HPC functionalities used in complex multi-tiered workflows. Finally, "Quantum for HPC" describes the HPC aspects of using Quantum - based technology.

Research Clusters represent cross-cutting "themes", which capture a set of research priorities of high interest for the next generation of HPC infrastructure. They are shown as vertical elements as the impact of most of them cuts through the majority or all of the research domains. Under the heading of each "Cluster", several aspects with a high degree of similarity are grouped (or 'clustered') together. Each cluster theme is being discussed very lively in the HPC community these days. All clusters were published as separate ETP4HPC white papers and are now included in this SRA to facilitate the definition of research priorities.

In our model, the research areas identified are grouped according to two dimensions: horizontal "Research Domains" and vertical ''Research Clusters'' (*See figure 5*).

Each domain description is divided into three parts:

1. "Research trends and current state of the art": introduces the domain by presenting the status quo and today's research areas and priorities.
2. "Challenges for 2023 – 2026": outline the most important problems and limitations that need R&I-attention in that period.
3. "Intersections with Sustainability": given this SRA's special focus on the role of HPC in the context of Sustainability, this part describes the specific technical characteristics common to both the research domain and the cluster "Sustainability". It illustrates the areas where each domain needs to feed solutions to the problems and challenges mentioned in this cluster.

The intersections with the other four clusters receive special attention in the respective sub-chapters "Challenges for 2023 – 2026".

Figure 5: The model of Research Clusters and Research Domains

6.1
Research Clusters

6.1.1
Sustainability

INTRODUCTION

While the term 'sustainability' is used in various interpretations depending on the domain in question to which the concept of sustainability is being applied, within the high-performance computing (HPC) world, it is mostly related to CPU energy efficiency, which is also measured in FLOPS per Watt. However, supercomputers are not only used to run simulation codes for a deeper insight into processes of the real world, but they are also increasingly used for data analytics, especially the field generally described as artificial intelligence.

This not only introduces different algorithms and architectures but also highlights a development that has been present in the classical HPC world for a while now: the ever-increasing amount of data moved around in a distributed world with an enormous impact on energy consumption. This data is not necessarily created by supercomputers - a lot of it comes from devices (e.g., mobile phones) at the edge of the network.

The carbon footprint of data centres during operation is only one parameter in an overall assessment of the environmental impact of computing and equally important are the CO_2 produced during production of the infrastructure as well as the huge demand for scarce natural resources.

Thus, the term "sustainability" includes much more than just straightforward energy efficiency during the operation of a data centre. Following the recommendations of GeSI, an international organisation dedicated to ICT sustainability[1], we advocate for a holistic perspective that takes into account everything that is part of the overall life cycle of systems, including not just hardware but also software, infrastructure, and, in particular, workflows and the humans working on and with these elements. It also needs to address the cost of runtime, manufacturing and disposal.

Of special interest to the HPC community is the identification of R&D priorities to further reduce the CO_2 footprint associated with designing, developing, producing, and operating HPC infrastructure with its specific characteristics, including workflows and HPC/HPDA applications.

KEY INSIGHTS

- There is a significant lack of an overall perspective on the ecological footprint of the computing ecosystem.
- The emphasis on optimisation of performance per energy of HPC systems (Watts per FLOPS) is too narrow.
- The slowing down of Moore's law might help to extend the duration of the use of high-end computing systems.
- Investments into optimisation and better workflows will lead to a more efficient use of HPC.
- The amount of data is growing rapidly, which has a great impact, economically as well as ecologically.

KEY RECOMMENDATIONS

- Sustainability, addressed in a holistic perspective, needs to gain significantly more focus in the design of HPC IT infrastructure and its operation in HPC centres and as an element of the digital continuum.
- The Life-Cycle Analysis (LCA) method used mostly for data- and cloud centres should be applied to the specific parameters of HPC infrastructure and HPC workloads. This should help understand and quantify options for further reducing the environmental impacts.
- An emphasis should be placed on the optimisation of the overall solution instead of only looking at single elements. To support this, it is important to:
  - establish opportunities to bring all stakeholders together and
  - increase the awareness of the roles they play.
- Develop the eco-system towards building modular systems facilitating a longer use of subsystems and operating them in a mix of technology generations.
- Because of CPU performance not increasing as much as in the past,
  - code developers should tune their codes better to systems and systems architectures since they will be available for a longer time; this includes but is not limited to the energy efficiency of the code,
  - workflows shall be developed further to be more efficient if systems are to be available longer,
  - at procurement, buyers should request modular upgrade options, an extension of support including the availability of security updates so that they will be able to use the hardware longer.
- Reduce the ecological cost of data volume and the movement of data
- Continue to improve the power efficiency of data centre infrastructure and supercomputers.

1. https://www.gesi.org/research/ict-sector-guidance-built-on-the-ghg-protocol-product-life-cycle-accounting-and-reporting-standard

DEFINITION OF THE TERM "SUSTAINABILITY"

It may well be that "Sustainability" is one of the most flexibly used words nowadays. This immediately leads to the question: how is sustainability defined in the context of this white paper?

A very general definition of sustainability that shows how broad the interpretability of this term can be is the following:

> *«Meeting the demand of the current generations without putting the demands of the future generations at risk».*

That is: we can sustain ourselves and our descendants only with practices that do no future harm.

Sustainability is often broken into three intertwined categories: social sustainability, economic sustainability and environmental sustainability. Together, these three forms of sustainability are known as the «three pillars of sustainability[2].» They provide a framework for applying a solutions-oriented approach to complicated sustainability issues. It is generally considered that the convergence of positive impacts on these three dimensions makes a given development more sustainable.

Another classification is the notion of **weak** and **strong sustainability**[3]:

- **Weak sustainability** is built on the neoclassical economic approach, which perceives the natural resources as superabundant or substitutable by another kind of capital. Weak sustainability also considers the economic system as a closed and isolated system that can grow infinitely.

  The natural capital can be substitutable by other capital forms and the environmental degradation can be compensated with other economic activities. In this sense, the CO_2 market, for example the degradation resulting from the gas emissions, can be compensated by payments, and that is sufficient to mitigate the global warming, promote the green revolution, and continue with the global economic agenda.

  Proponents of weak sustainability postulate that future generations will have best technologies, more information and more capital available to address their environmental problems so that we do not have to invest too much effort today to avoid such unknown problems.

- **Strong Sustainability**
  Proponents of this approach recognise weak sustainability as an essential first step in the right direction but only one step is insufficient. It is about intergenerational equity and justice.

  Why can future generations not solve their problems? Because as we all already know, future conditions depend on the present generation's actions. Today, the global economic activities have reached a scale that is big enough to undermine the welfare of the future generations. Some scientific studies point out the long-term and the irreversible effects of environmental pollution.

  Strong sustainability postulates that there are some critical natural capital resources and services that are non-substitutable through other forms of capital and such resources and services must be preserved; for example, the ozone layer and the biological diversity.

  Unlike weak sustainability, strong sustainability puts an emphasis on ecological scale over economic gains. This implies that nature has a right to exist and that it has been borrowed and should be passed on from one generation to the next intact in its original form[4].

Additionally, sustainability is often related to the concept of resilience, which can be defined as the ability of a system to absorb a disturbance (fault) while remaining viable or operational.

It becomes obvious, that "sustainability" can be addressed in such a diverse way that defining a single working solution and interpretation is a challenge even if we focus only on a specific discipline like HPC.

Sustainability in an ecological sense is in the focus of this chapter: it is about the amount of resources that are used to build and operate the systems as well as the amount of e-waste generated.

SUSTAINABILITY IN THE CONTEXT OF HPC

With the growing need for HPC-based high precision simulation and modelling to solve various social and environmental global issues, the negative impact of using HPC infrastructure on the climate becomes a matter of concern. The **energy consumption** of data centres is rising above what should reasonably be supplied:

The report "Energy-efficient Cloud Computing Technologies and Policies for an Eco-friendly Cloud Market" issued by the Borderstep Institut[5] states that "The analysis and modelling of the future energy demand of data centres across the EU Member States shows that the energy consumption of data centres in the EU28 increased from 53.9 TWh/a to 76.8 TWh/a between 2010 and 2018. This means that in 2018, data centres accounted for 2.7% of the electricity demand in the EU28. This growth is so strong that it has more than off-set the significant efficiency gains achieved at all levels (hardware, software, data centre infrastructure) and consequently, the total energy consumption of data centres in Europe has risen. Compared to 2018, the energy consumption of data centres is expected to increase by 21% to 92.6 TWh/a by 2025."

Of course, the share of HPC data centres is a subset of data centres in general, the growth rates are expected to be in the same percentage range though. Not only the CO_2 emitted during operation of the infrastructure is relevant but also the huge demand for

2. https://www.treehugger.com/what-are-the-three-pillars-of-sustainability-5189295
3. https://smartcity.pharosnavigator.com/static/content/en/1032/Strong-Sustainability-vs-Weak-Sustainability.html
4. https://en.wikipedia.org/wiki/Weak_and_strong_sustainability
5. https://www.borderstep.de

natural resources (where scarcity can no longer be overlooked) and the large carbon footprint of the production of ICT equipment needs to be dealt with to be able to sustain these infrastructures.

Improving the manufacturing process is therefore a key issue and it offers probably (like many other industrial processes) a high potential for improvement with respect to the overall environmental cost. The energy used during the development and production of HPC systems is often called "grey energy": it is commonly accepted that nowadays building the components, delivering them and assembling them represents a significant part of the ecological footprint. Depending on how much the power used during the operating phase of the system is decarbonated this can be up to 80% of the total ecological footprint (see[6]); however, on the other end of the scale it could also go down to less than 50% if power during the use of the system is CO_2-intensive[7].

It is also noteworthy that only 6–12 percent of the energy consumed by a large HPC system is devoted to active computational processes. The remainder is allocated to cooling and maintaining the consecutive chains of redundant fail-safes to prevent costly downtime[8]. Optimising the combination of reliability and resiliency so as to minimise redundancies is also a key factor for reducing the environmental footprint of operations.

The share of the end of the life cycle of HPC systems (i.e. decommissioning) in the ecological footprint is to a certain extent already defined during the development and building period, and should be included in this grey energy; e-waste should be avoided and has to be considered from the beginning.

Hence, in a «weak sustainability» approach, extending the lifetime of the hardware (e.g., preventive maintenance and circular usage, (i.e., recycling HW from supercomputers to equip general purpose data centres) could be a first step forward reducing the environmental impact.

With respect to **the architecture of HPC systems**, energy efficiency has been greatly improved (x10 over the last ten years) but at the expense of greater system heterogeneity with the generalisation of accelerators (e.g., GPU). The use of accelerators is also promoted by the adoption of AI technologies. This tendency has contributed to adding complexity to the development process and to the extension and variety of the required software stack. Furthermore, this makes systems less usable as well as more difficult to render applications sustainable (i.e., running efficiently on multiple generations and varieties of systems). Besides this, the current state of the art in AI training uses power-hungry algorithms.

In addition, **system software, applications and workflows** have a significant influence on the CO_2 footprint of HPC systems during operation:

While the development of new codes is of course important, the elements which are essential for a longer lifetime of the codes are easily neglected. If the reusability of a code can be improved, the efficiency of the code development of other programmers will benefit significantly (in the commercial as well as in the academic world). Closely related to this is the transfer from research into productive and commercial use. A holistic approach here would have a great potential to improve the efficient use of the systems, thus increasing sustainability.

The more a system is used in a product development context (i.e., frequently and embedded in a business process and environment with a high demand for reliability), the more important are its **workflows**. Companies invest a lot of time and money to improve R&D workflows because they allow them to make better use of the scarce resources of an experienced workforce. If a process can be automated, engineers can focus much better on the real problem at hand. Every time the system environment changes though, these workflows need to be adapted; depending on the nature of the change this can also be costly, time consuming and could have carbon footprint implications of its own. If system developers (hardware and software) better take into account what changes for the productive use of the system arise when they change a technical aspect, the overall efficiency will benefit.

Another aspect is the **cost of adapting and porting** software/applications/workflows to new generations of infrastructures. There is a trade-off to be made between highly optimised and tuned applications/workflows (which would require less hardware and energy) and the usability of the infrastructures/systems for an extended range of standard applications. This is also relevant for the other two dimensions of sustainability (social and economic): the usability of system software for a wider range of applications (e.g., industry so as to increase its positive economic impact or academia to increase its social impact) and the dissemination of the results (to address both training of students/staff and also be complemented by other communities).

A big roadblock in creating synergies between the data and compute elements is **cybersecurity**. On the one hand, we need to connect data and HPC infrastructures to deploy efficient large-scale application workflows. On the other hand, each infrastructure tends to shield itself from external attacks making the life of application developers more difficult; this either makes it more expensive or leads to less efficient applications. Thus, designing system software with security AND energy efficiency in mind besides the traditional set of parameters like performance is required.

6. Thomas Taro Lennerfors, Per Fors, and Jolanda van Rooijen. "Sustainable ICT: A Critique from the Perspective of World Systems Theory". In: ICT and Society. Ed. by Kai Kimppa et al. Berlin, Heidelberg: Springer Berlin Hei-delberg, 2014, pp. 57–68. isbn: 978- 3-662-44208-1.9
7. U. Gupta et al., «Chasing Carbon: The Elusive Environmental Footprint of Computing», in IEEE Micro, doi: 10.1109/MM.2022.3163226 (2022). https://arxiv.org/pdf/2011.02839.pdf
8. Monserrate, Steven Gonzalez. 2022. "The Cloud Is Material: On the Environmental Impacts of Computation and Data Storage." MIT Case Studies in Social and Ethical Responsibilities of Computing, no. Winter 2022 (January). https://doi.org/10.21428/2c646de5.031d4553

While many driving forces push towards new uses of the infrastructure, **legacy codes** must not be neglected. They are the basic blocks of many new applications but were designed for different, less complex system architectures. In their original form, they of course do not take advantage of modern systems features facilitating a more power efficient execution. Although they represent a very important scientific capital and therefore must not be dilapidated, studies show that the modernising effort is far from being painless (see EXDCI[9]).

The **amount of data is growing steadily** (e.g.[10]), increasing data locality issues (if there is more data, it is more likely that the data that is needed is stored somewhere else) and requiring larger storage facilities.

The complexity is in part due to the increasing number of user access endpoints (PC, phones, tablets, etc.) as well as more data producers (e.g., IoT) that generate a lot of data movement, and makes the global efficiency challenging to quantify. In this context, optimising the usage of the infrastructure in a holistic manner requires new tools and the ability to provide meaningful feedback to users. This addresses all three dimensions of sustainability:

- minimise the environmental impact
- increase the added value of the data product (positive economic impact)
- enhance user experience and provide new services (positive social impact)

Another orthogonal consideration in the context of policies relates to the **environmental-related taxations** as well as the geopolitical situation. The current economy of the HPC ecosystem can be strongly impacted if new carbon or environmental taxes were to increase. The importance of reusing hardware or having long lasting systems could become the only solution to achieve cost effective solutions.

TOWARDS MORE SUSTAINABLE HPC-INFRASTRUCTURES

In this section, we provide a set of recommendations to help improve the sustainability of HPC infrastructures (in a broad context).

BUILD MODULAR AND OPTIMISE THE WHOLE SOLUTION, NOT ONLY ITS PIECES

Sustainability must be an important factor when it comes to the way to design new solutions with a limited budget of natural resources and energy consumption. As energy gets greener, the production step will be the dominating issue regarding the environment. The production phase includes all activities from natural resources extraction to manufacturing and deployment of products.

Modularity and flexibility are to be implemented at multiple levels of the data centre with the cooling and power infrastructure, the rack, the blades, the servers, the network and the storage. However, it is difficult to innovate in terms of energy efficiency without introducing new hardware components. It has also to be taken into account that on the very high-end new technology has to be introduced more frequently and might be more disruptive; however, for those environments, where HPC has become a commodity and where the focus is more on capacity than on capability, increased modularity might not only bring ecological but also economical benefits.

Building modular systems aims at extending the exploitation time of systems and sub-systems. Completely new hardware design approaches are required allowing for partial upgrades which facilitates a circular economy and helps to reduce e-waste. Accelerator-based architectures can be a way to consider gaining performance while keeping the central parts of the systems unchanged. It should also be noted that this path of reasoning is in line with the end of Moore's Law, which forces system designers to look for performance and energy efficiency by using specialised hardware.

This approach has been known for a long time in the area of embedded system design. At the core of this approach are advances in technology such as silicon photonics connections that, thanks to their short latency, may help to build more modular machines (for instance by separating the memory from the compute part). At the edge of the continuum, the key factor is reconfigurability in order to adapt over time to new application needs.

Extending the lifetime of a system also requires increasing its resilience to faulty hardware or other system failures. This is by no means a new domain. However, an increase of the number of components (within the HPC system itself as well as all other parts of the workflow), along with a longer lifetime necessitates reviewing and improving the resilience and fallback mechanisms. In particular, this challenges the use of the latest technological node (7nm or less) that increases the faults caused by electromigration[11] (e.g., this is why a processor or a memory DIMM can only have a limited lifespan).

9. https://exdci.eu/
10. Petrie, R., S. Denvil, S. Ames, G. Levavasseur, S. Fiore, C. Allen, F. Antonio, et al. 2021. "Coordinating an Op-erational Data Distribution Network for Cmip6 Data." Geoscientific Model Development 14 (1): 629–44. https://doi.org/10.5194/gmd-14-629-2021.
11. https://www.sciencedirect.com/topics/engineering/electromigration

It is important to note that the full application environment needs to be included into the design of solutions. As more applications integrate external data sources and processing, we need to consider large-scale workflows (i.e., application workflows distributed over multiple infrastructures, some compute oriented, others data oriented). Different steps in these workflows ask for different hardware and thus today's HPC systems are often composed of different modules optimised for certain stages in these workflows (see for instance[12]). Concerning rack and server modularity, initiatives such as the Open Compute Project, OCP[13] are a good illustration of this trend, which is expected to become stronger in the future.

These changes which aim at capitalising on existing hardware beyond the current five years are not risk-free in terms of stability and maintainability of the supercomputer. A European initiative aims to explore a modular and sustainable architecture at the supercomputer level (e.g. Modular System Architecture: Modular Systems Architecture, MSA[14]).

All these considerations require not only new system design approaches but also to re-assess the rationale behind "programmed obsolescence". As mentioned above, the paradigm of a circular economy can also be applied to HPC infrastructure: for example, some hardware may already be obsolete for supercomputers but it may be reused in smaller, less demanding use scenarios. This calls for a new business model that develops around this circular process. Services are needed that refurbish second-hand hardware and fine tune the ecosystem to its new use.

In the context of **modernising applications,** it should be noted that scientific domains rely significantly on legacy codes with a significant impact on power and resource efficiency. Application evolution is rather a composition with the integration of extensions than a complete redesign (almost nobody starts from scratch). The approach "AI for code", i.e., integrating an AI phase before or during simulation, thus avoiding the re-write of large amounts of code, might help to increase the overall efficiency of the code at runtime.

The functional evolution of major applications seems to be their use in the "computing continuum"; this is defined as a fluid ecosystem where distributed resources and services are programmatically aggregated on demand to support emerging data-driven application workflows. This covers the fact that new applications integrate sources of data, complementary hardware resources such as storage pools with shared databases or scientific data lakes, visualisation pools or different supercomputers with specific capabilities.

The fundamental issue of the continuum is to provide the ability to deploy complex end-to-end application workflows.

From a sustainability point of view, the major issue for applications is the hardware heterogeneity that requires specific portability efforts and an evolution of the software stack. While the trend towards heterogeneity is certainly helpful in many aspects, it also makes the task of programming these systems and using them efficiently much more complicated. Often, the combination of different programming models is required and selecting suitable technologies for certain tasks or even parts of an algorithm is difficult although it could be cost effective.

For sustainability purposes, it could be valuable to use virtualisation solutions or container technologies. High Performance Computing (HPC) workloads have been traditionally run on bare-metal, non-virtualised. Virtualisation was often seen as an additional layer that leads to performance degradation. Performance studies have shown that virtualisation in many cases has a minimal impact on HPC application performance. In general, the method of virtualising physical system resources like processors, memory, network and I/O, and using them by a multitude of programs at the same time increases the utilisation of the physical system and might avoid to upgrade/enlarge it.

REDUCE THE COST OF DATA VOLUME AND MOVEMENTS

Reducing the cost of data volume and movements is another key factor in achieving more sustainable systems. The exponential growth of data volume[15] is accompanied by an exponential need of storage hardware. These storage components such as SSDs or HDDs require a lot of resources and energy in the production phase, out-consuming the operation phase (see Dell analysis[16]). Also, there are storage components such as tape cartridges that require much less resources to produce and use. Smarter workflows that can predict their use of data could take better advantage of high-latency media such as tape, by prefetching data from tape to online storage.

With the introduction of AI in the supercomputing domain and new practices such as continuum computing, this exponential factor is not only relevant for data centres.

Discarding raw data is intrinsically difficult, because future progress in data analytics and AI may lead to new discoveries. Therefore, addressing the data issue does not only require improvement in technologies, but also in data management policies and modelling capabilities. This is necessary in order to 1) reduce the effort of keeping the data and 2) improve the ability to extract key information from large datasets.

12. https://www.phidias-hpc.eu/
13. https://www.opencompute.org
14. https://zenodo.org/record/6508394#.YotFwi-21GM
15. PRACE scientific case https://prace-ri.eu/about/scientific-case/, 2018
16. https://corporate.delltechnologies.com/content/dam/digitalassets/active/en/unauth/data-sheets/products/servers/lca_poweredge_r740.pdf

A rather different angle to address this issue would be: if the value of stored data can be increased, the overall cost calculation of moving data around must also be recomputed (especially if the value depends on the location). This could be achieved by processing/grouping/sorting and including provenance in terms of both, data and workflow into the meta data and bringing "compute closer to data" instead of "data to compute".

Another fairly new issue related to data volume is the energy required for deep learning techniques[17]. Improving the efficiency of AI training algorithms may lead to more sustainable approaches. Some scientific domains such as particle physics, astrophysics, weather forecasting or the Square-Kilometer-Array project (SKA[18]), are pioneers in the research of data workflow and data management.

In order to strengthen data sustainability, the scientific community also takes care of obsolescence of data format and type. Most of scientific data is stored in databases or files with poor efficiency. Translation or migration to new data models and storage formats will be required (which can be made easier by an increased use of open data formats[19]). This might also be helpful for the cases when the value of the data is increased.

CONTINUE TO IMPROVE THE POWER EFFICIENCY OF DATA CENTRES AND SUPERCOMPUTERS

The main driver of emissions in the use phase is the electricity that is consumed by the servers, networks and storage. The carbon intensity of this electricity (usually a mix of different sources) depends on geography and the forms of energy that are used to generate electricity and this is also true for electricity costs.

Carbon intensity is a parameter which can reflect the share of renewable energy used during the operational phase. During the production phase, it might be more difficult to address the carbon footprint.

Supercomputers and data centres have similar issues: they shall offer the highest availability and efficiency in a power envelope that cannot increase beyond certain limits: in the 1990s, exascale international initiatives targeted the energy consumption of an exascale system in a power envelope of 20MW. It was and still is a huge challenge to design a system with this performance/Watt efficiency: first, each hardware component of a supercomputer had to be optimised and then the entire software stack. This target is still in use today, but we can assume that the first exascale system will consume between 30 and 40MW.

In the race of power efficiency several best practice examples in energy-efficient data centres and server rooms include a variety of approaches, in public as well as in private procurement. The most common approaches are[20]:

- More efficient cooling systems
- Heat reuse, e.g., for district heating
- Virtualisation of software, optimal use of server capacity
- Use of renewable energy to supply data centres
- Construction of data centres in regions with a cold climate

Further significant improvements in energy efficiency of supercomputers are mandatory despite two significant hurdles:

Infrastructure efficiency and the PUE of data centres improved over the past few years. The remaining energy-efficiency potentials are getting smaller as technology is moving closer towards the physical limits. At the same time, unlocking the remaining potentials is getting more complicated.

The "rebound effect" has to be taken into account: A direct rebound effect refers to the increased level or frequency of use of products by consumers as a result of the improved levels of energy-efficiency and the lower costs per use. Thus, the more efficient data centres and supercomputers become, the more they tend to scale up and grow in capacity. The efficiency gains are mostly lost or even overcompensated.

USE SUSTAINABILITY AS A GUIDE FOR DESIGNING NEW SOLUTIONS AND ROADMAPS

The development of new R&D roadmaps and technical solutions seems to be a constant in the world of HPC. What needs to change is that sustainability is established as a key element in this process, as a crucial guide to design. It has already led to efforts to reduce data centre energy demand during operations.

To be fully transparent, the sustainability aspects must be analysed at solution level, because improvements in one area and to one impact can affect sustainability of other components or services. Overall, a holistic perspective is mandatory, trade-offs seem to be inevitable!

For example, it does not make sense to extend the hardware lifecycle in regards of environmental costs of its production without evaluating the impact of obsolescence:

- Technical materials obsolescence (fragile components, lack of spare parts, without repair capabilities or the lack of compatibility),
- Software obsolescence (software incompatibility, version and power problems)
- Cultural obsolescence, lack of expertise or education capabilities to maintain those technologies...

17. Emma Strubell, Ananya Ganesh, and Andrew McCallum. "Energy and Policy Considerations for Deep Learning in NLP." In: ACL (1). Ed. by Anna Korhonen, David R. Traum, and Lluís Màrquez. Association for Computational Linguistics, 2019, pp. 3645–3650. isbn: 978-1-950737-48-2. url: http://dblp.uni-trier.de/db/conf/acl/acl2019-1.html#StrubellGM19.
18. https://www.skatelescope.org/the-ska-project/
19. https://www.opencompute.org
20. https://www.borderstep.de/wp-content/uploads/2020/11/KK0320210ENN.en_.pdf

Sustainability must be a way to design new solutions with a limited budget of natural resources and energy consumption. Environmental impact has to be considered carefully but it may be acceptable if

1. it is non-destructive and
2. it provides significant benefits for the other dimensions of sustainability (society and economy).

A good example for this is our ability to design more efficient solutions on top of existing ones. If they were not built, our long-term progress might be hampered. This is a difficult trade-off of course, but the general flow of developments could well be as important as looking at negative immediate environmental impacts. Compromising the environment in the short term though is only acceptable if the longer-term environmental benefits of technical advancements are quantified and their implementation is planned for, yielding an overall positive balance.

On a different level, it might also make sense to think about whether there is a bridge to be made between sustainability and the European strategy of digital sovereignty. The latter includes more technology manufactured and produced locally which probably helps us regain leverage on the manufacturing capacity with a lower CO_2 footprint, besides the benefit of hosting and taking control of our data.

The discussion becomes even more complex taking the "Enabling Effect" into account: The use of HPC technology enables new discoveries and insights into materials, processes, human social behaviour which can lead to overall savings of CO_2. One example is the aero-dynamic modelling of entire planes vs. a single wing only. The outcome can be the potential for further fuel-savings. Only exascale supercomputers will allow for this size of a modelling problem. So, a "bigger supercomputer" not only generates more CO_2 on the downside, but it might also help reduce it (or even more) in a totally different and unrelated fields of our daily lives.

CONCLUSION

Sustainability definitely is an issue today, almost everywhere. And IT, especially high-performance computing and data analytics/artificial intelligence, is too big and too important to not be in the focus of the activities aiming to improve sustainability in our personal but even more important in our business lives.

However, we have also found that what is done today is not enough. Currently, it looks more like a set of individual actions, each addressing a single issue. However, in order to be able to really make significant progress, it is important to take a holistic perspective. Only then will we be able to uncover the hidden treasures, those that can only be lifted, if we address the challenge beyond isolated silos. The entire life cycle of HPC IT equipment, as shown in Figure 6, needs to be characterised, measured and quantified. Recently processes and tools emerge which help with these evaluations and assessments[21].

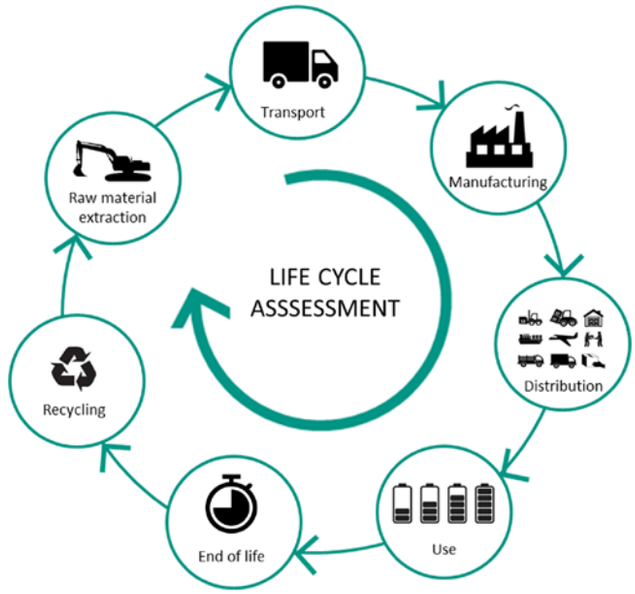

Figure 6: The stages of the life cycle of IT equipment (courtesy CINECA)

In terms of increasing the speed of action towards advancements in sustainability, we need to make sure that the exchange between all relevant players (e.g., technology providers, software developers, users, and policy makers) is vital, open, integrative and driven by the same goals. This interaction, if done with focus, will provide faster progress and optimise the whole ecosystem.

21. https://www.gesi.org/research/tool-evaluating-the-carbon-reducing-impacts-of-ict

We identified topics such as energy efficiency, the cost of storing and moving data, the carbon footprint during production and the concept of a circular economy; even longer lifetimes might be achievable. And, on the same token, we need better tools to measure energy usage, energy efficiency, the overall carbon footprint, etc. But all of this should be seen as bricks of a building, where each of them plays a specific, evenly important role; and only if they are all put together in the right way, we will be able to achieve the best possible level of sustainability. This includes that we develop and refine methods to measure the ecological impact so that it can be included in the TCO computations and thus help to improve the optimisation.

Therefore, our recommendations address this bigger picture and we think that more effort should be put into convincing the community that it is worth to apply an end-to-end approach at sustainability, without neglecting the work on all the bricks that are necessary to achieve more aggressive targets (which, of course, are moving over time!). With the extensive growth in energy demand of the ICT sector and the scarcity of resources needed to produce the equipment and the huge expanding carbon footprint, it is essential we follow these recommendations rather today than tomorrow. Even more so when we acknowledge that the sustainability roadmap is fully consistent with the emerging digital sovereignty policies in Europe.

MORE REFERENCES ON SUSTAINABILITY

Asch, Mark, et al. «Cybercosm: New Foundations for a Converged Science Data Ecosystem.» *arXiv preprint arXiv:2105.10680* (2021). https://arxiv.org/abs/2105.10680

Deelman E, Mandal A, Jiang M, Sakellariou R. The role of machine learning in scientific workflows. The International Journal of High Performance Computing Applications. 2019;33(6):1128-1139. doi:10.1177/1094342019852127

Huerta, E.A., Khan, A., Davis, E. *et al.* Convergence of artificial intelligence and high performance computing on NSF-supported cyberinfrastructure. *J Big Data* 7, 88 (2020). https://doi.org/10.1186/s40537-020-00361-2

Kasim, Muhammad F. et al. "Building high accuracy emulators for scientific simulations with deep neural architecture search." Machine Learning: Science and Technology 3 (2021): n. pag. https://arxiv.org/pdf/2001.08055.pdf

A. S. Szalay, «From simulations to interactive numerical laboratories,» Proceedings of the Winter Simulation Conference 2014, 2014, pp. 875-886, doi: 10.1109/WSC.2014.7019948.

https://www.hpcwire.com/off-the-wire/using-ai-and-supercomputers-researchers-remove-one-of-the-biggest-roadblocks-in-astrophysics/

6.1.2

HPC in the Digital Continuum

Supporting integrated applications across the Edge-Cloud-Supercomputer layers to address critical scientific, engineering and societal problems

INTRODUCTION

Modern use cases such as autonomous vehicles, digital twins, smart buildings and precision agriculture, greatly increase the complexity of application workflows. They typically combine physics-based simulations, analysis of large data volumes and machine learning and require a hybrid execution infrastructure: edge devices create streams of input data, which are processed by data analytics and machine learning applications in the Cloud, and simulations on large, specialised HPC systems provide insights into and prediction of future system state. From these results, additional steps create and communicate output data across the infrastructure levels, and for some use cases, control devices or cyber-physical systems in the real world are controlled (as in the case of smart factories). All of these steps pose different requirements for the best suited execution platforms, and they need to be connected in an efficient and secure way. This assembly is called the Computing Continuum (CC) [1][1]. It raises challenges at multiple levels: at the application level, innovative algorithms are needed to bridge simulations, machine learning and data-driven analytics; at the middleware level, adequate tools must enable efficient deployment, scheduling and orchestration of the workflow components across the whole distributed infrastructure; and, finally, a capable resource management system must allocate a suitable set of components of the infrastructure to run the application workflow, preferably in a dynamic and adaptive way, taking into account the specific capabilities of each component of the underlying heterogeneous infrastructure.

To address the challenges, we foresee an increasing need for integrated software ecosystems which combine current "island" solutions and bridge the gaps between them. These ecosystems must facilitate the full lifecycle of CC use cases, including initial modelling, programming, deployment, execution, optimisation, as well as monitoring and control. It will be important to ensure

KEY INSIGHTS

- **There is a clear trend to combine numerical computations, large-scale data analytics and AI techniques to improve the results and efficiency of traditional HPC use cases, and to advance new use cases in fields such as autonomous vehicles, digital twins, smart buildings/towns etc. Such use cases are typically implemented as complex workflows and will require the coordinated use of supercomputers, cloud data centres and edge-processing devices.**
- **Today, separate, efficient software ecosystems [2][2] exist for the management of computation, communication and data on supercomputer facilities, cloud infrastructures or edge-based systems. Yet, these address the specific requirements of their infrastructure layer and typically fail to smoothly interoperate and cooperate. The same is true for the different machine learning (in particular Deep Learning) ecosystems.**
- **Complex workflow orchestration across the whole continuum leads to challenges at multiple levels: application/algorithmic level (programming paradigms), middleware level (deployment, execution, scheduling, monitoring, data storage and transfer, processing and analysis) and resource management level.**
- **CC workflows will likely have specific hardware requirements e.g. for on-the-fly encryption, efficient and low latency communication, transparent compression, automatic brokerage, in-network data processing and generally for energy efficient computation and communication as large scale systems will have significant energy footprints[3,4].**
- **Co-design and cooperation among experts in the different areas involved (e.g. HPC, data analytics, AI/Deep Learning, cybersecurity, mobile communication) will be a key prerequisite for building Computing Continuum hardware and software and support complex application workflows in an effective, efficient and secure way. The TransContinuum Initiative (TCI)[5] is providing such a cross-domain cooperation framework.**
- **HPC datacentres today feel the pressure of Cloud computing vendors like AWS and MS Azure[6]. It will be attractive for small to medium scale HPC workloads to leverage Cloud interfaces and instantly available Cloud resources which offer almost unlimited scalability and compute throughput.**
- **HPC centres will have to adopt such Cloud interfaces, and they are poised to gain additional customers from CC use cases and from the deep learning and data analytics communities.**

1. D. Balouek-Thomert, E. Gibert Renart , A. R. Zamani, A. Simonet and M. Parashar, "Towards a computing continuum: Enabling edge-to-cloud integration for data-driven workflows," The International Journal of High Performance Computing Applications, vol. 33, no. 6, 2019.
2. G. Nguyen, S. Dlugolinsky , M. Bobák, V. Tran, A. López García, I. Heredia, P. Malík and L. Hluchý, "Machine Learning and Deep Learning frameworks and libraries for large-scale data mining: a survey," Artificial Intelligence Review, vol. 52, p. 77–124, 2019.
3. O. Peckham, "15 Years Later, the Green500 Continues Its Push for Energy Efficiency as a First-Order Concern in HPC," HPCwire, no. https://www.hpcwire.com/2021/07/15/15-years-later-the-green500-continues-its-push-for-energy-efficiency-as-a-first-order-concern-in-hpc/ , 17 07 2021.
4. E. Masamet, A. Shehabi, N. Lei, S. Smith and J. Koomey, "Recalibrating global data center energy use estimates," Science, vol. 367, no. 6481, pp. 984-986, 2020.
5. https://www.etp4hpc.eu/transcontinuum-initiative.html
6. https://hyperionresearch.com/wp-content/uploads/2021/07/Hyperion-Research-HPC-Briefing-Slides-During-ISC-June-2021.b.pdf

KEY RECOMMENDATIONS

- **Applications/algorithms: develop programming models and systems which integrate HPC, AI/ML and data analytics processing and facilitate hybrid applications such as AI-enabled simulations. This is not a CC requirement by itself, yet it will help in programming the steps in a CC workflow.**
- **Data storage, transfer and processing: unified abstractions must enable interoperable data storage and processing across the continuum, and facilitate data analytics at all levels (HPC center, cloud, edge). In addition, automated data placement and transformations in transfer should be supported.**
- **Workflow programming: programming models and systems are required which can support large, dynamically evolving workflows, with workflow steps combining HPC, data analytics and AI/ML processing.**
- **Workflow deployment, orchestration and monitoring: deployment and orchestration of workflows must be seamless and dynamically adapt to the load and the condition of the distributed infrastructure. Fine-grained monitoring of running workflows and reproducibility of results (to a degree adequate for the actual application) are important. This could, for instance, be achieved by extending recent advances in the Cloud Computing field.**
- **Software interoperability and composability: this concerns the established HPC, data analytics and AI/ML stacks for which interoperable and composable implementations will be required .**
- **Authentication, authorisation and accounting: these have to be interoperable across all layers of a CC infrastructure, and be combined with a secure, pay per use billing mechanism.**
- **Data security: interoperable mechanisms for data encryption and transfer, data access control and monitoring will be needed.**
- **Management of heterogeneous systems: a CC infrastructure will include a wide variety of heterogeneous compute, storage & communication systems, including accelerators like GPGPUs, FPGAs and others. It will be critical to find an integrated way to manage these resources in an efficient and secure manner.**
- **HPC data centres must rethink their offerings, adapt to new usage models and customer requirements and evolve their business models accordingly. De-facto cloud standard interfaces APIs for storage and computation must be considered.**
- **HPC data centres must address the increasing heterogeneity of computing hardware with new approaches for resource management and user management.**
- **Energy efficiency: the energy consumption of CC use cases must be reduced as far as possible, for instance by the development and use of highly efficient hardware, the minimisation of data movement and communication between systems and layers, and the adoption of novel algorithmic approaches, such as mixed precision computing.**

adequate reproducibility of workflow results and to find ways for creating and managing trust when sharing systems, software and data. All of these will in turn require novel or improved hardware capabilities. This white paper provides an initial discussion of the gaps. Our objective is to accelerate progress in both hardware and software infrastructures to build CC use cases, with the ultimate goals of accelerating scientific discovery, improving timeliness, quality and sustainability of engineering artefacts, and supporting decisions in complex and potentially urgent situations.

WHAT IS THE COMPUTING CONTINUUM AND WHY IS IT IMPORTANT

There is an extremely rapid proliferation of digital devices generating valuable data in many fields of application (science, meteorology, autonomous vehicles, industry 3.0/4.0, social media, etc.). The combination of physics-based simulation (classic HPC) and data-driven modelling (using machine learning techniques) has been shown to create impressive results, besides improving time and energy required[7,8]. Large scale distributed supercomputer/Cloud/Edge infrastructures demonstrate how compute and data throughput can be scaled[9]. Given the above, it now seems possible to create new data-driven use cases which address important scientific, commercial and societal challenges in novel ways, by combining simulation, analytics and learning. Such use cases will make full use of CC infrastructures, and to a large degree depend on their existence.

7. S. Dewitte, J. P. Cornelis, R. Müller and A. Munteanu, "Artificial Intelligence Revolutionises Weather Forecast, Climate Monitoring and Decadal Prediction," Remote Sensing, vol. 13, no. 16, p. 3209, 2021.
8. M. Feldman, "HPC in 2020: AI is no longer an experiment," The Next Platform, no. https://www.nextplatform.com/2020/01/09/hpc-in-2020-ai-is-no-longer-an-experiment/ , 09 01 2020.
9. Google Cloud, "Helping researchers at CERN to analyze powerful data and uncover the secrets of our universe," [Online]. Available: https://cloud.google.com/customers/cern .

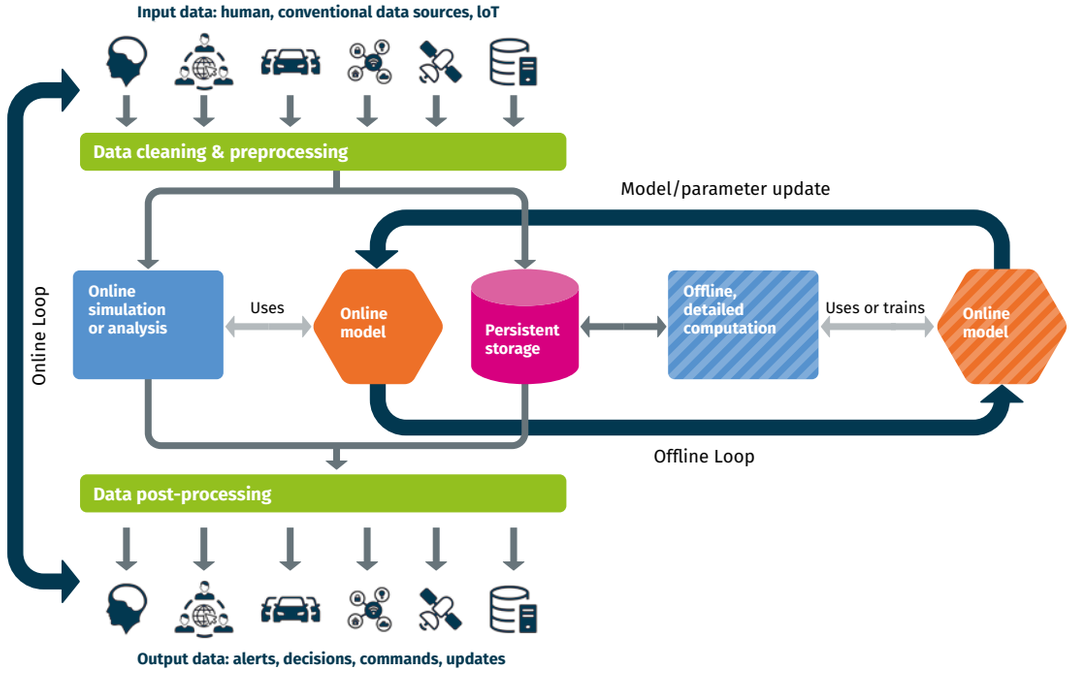

A USE CASE IN PRECISION AGRICULTURE

Modern greenhouses have sophisticated ways to control the vegetable environment to increase their production, based on models of plant growth. However, they require more and more parameters to be set by the grower (e.g., over 200 in a standard soil-less glasshouse for tomatoes)[10]. To be relevant, these models need to be adapted to the location of the farm, the needs of the species and of the cultivar, their potential in yield and quality (dry matter and sugar content). A promising approach consists in combining models of plant and climate development with actual monitored data processed through machine learning algorithms, to accurately model the greenhouse costs and its production potential. To introduce near-real time data assimilation of climate and plant observation to correct the simulations scenarios of plant needs, data inputs will be computed by a model (Digital Twin) executed in the cloud; then fusion with recent data dynamics and projections will be done by the models deployed on the edge.

A major challenge in such a scenario is that the complexity of the system represented (indoor climate, greenhouse management systems, plants interacting and reacting to uncertain weather) can make the simulation results drift significantly from the mon-

10. This use case is proposed by CYBELETECH https://www.cybeletech.com/en/home/ in collaboration with Ctifl https://www.ctifl.fr.

itored data dynamics (energy failure, climatic extremes, changes in grower production strategy...). When this occurs, the models need to be corrected according to the observed data dynamics (calibration of the models, learning of machine learning algorithms). When such drifts are detected for the Digital Twin, the process should automatically launch corrections of part of the models, which in turn calls for on-demand, dynamic resource allocations in the cloud or at the edge Edge. For example, if a breakdown in greenhouse climate management occurs and is fixed during the day, the quasi-real-time forecast and decisions on the Edge should be reviewed while the global strategy computed mainly in the cloud stays valid. In contrast, on-demand Cloud simulations need to be deployed when significant changes in recent data dynamics is detected.

Overall, this use case calls for the deployment and the execution of a distributed, cross platform workflow from the edge to the cloud, with the potential of the unpredictable data dynamics triggering on-demand, computationally intensive simulations in the cloud and related processing at the edge.

The use case scales by including more greenhouses, growers and regions.

A USE CASE IN SMART BUILDINGS & CITIES

Buildings and infrastructure are pivotal in the socio-economic transition towards a sustainable and climate-neutral economy. People spend most of their life in buildings, which are responsible for approximately 50% of the global energy consumption across their whole life-cycle. Optimizing their energy consumption is therefore highly important to ensure sustainable development.

Additionally, more and more sensors are installed over built-up areas to form so-called smart city environments. This way real time air pollution monitoring, traffic and citizen dynamics, road conditions and other events can be monitored and analyzed. Based on that data and on detailed meteorological forecasts, prescriptive measurements can be taken e.g. to control building lighting and heating/cooling/ventilation or the traffic flow to reduce energy use and air pollution, trigger warnings or take action with predictive maintenance.

To address this goal, based on historical operation data, machine-learning techniques are largely deployed. This kind of data-driven analytics can be carried out on a distributed platform. Real-time sensor data will be used as an input,integrating Edge processing for data preprocessing with Cloud-based analytics. Further computation-bound simulations will be used for

AFRISOL[11]: a sensorised smart building in Madrid (this is a pilot building for the project investigating the use case described here)

11. https://www.ciemat.es/portal.do?IDM=61&NM=2&identificador=92

predicting effects. To enable high-resolution computations, access to an HPC facility may be required. Thus data movement and data-driven calculations will continuously be performed on the hybrid platform.

A key aspect is that the input data can vary in frequency, relevance, amount of available data, etc. This can lead to variations in data provisioning, that will induce variations of the computational/processing workloads; this requires dynamic reallocation of processing tasks and resources across the CC; at the same time, this will result in a more/less intensive computation for learning behavioural models, and, lastly, to a modification in the forecast and hindcast calculation due to better match actually observed data. This variability does require support for seamless data processing across the CC and also smart, dynamic scheduling, as well as allocation of processing tasks, to react to data variations.

EARTH SYSTEM DIGITAL TWIN

The "Destination Earth" initiative[12] is creating detailed, digital replicas of the Earth, and will model a wide range of natural systems (including atmosphere, oceans and rivers, the cryosphere and subterranean water) at a very high and hitherto unattained level of detail and precision, using proven simulation methods in combination with AI/ML techniques. Besides its obvious value for all Earth-related scientific activity, this digital twin will enable detailed investigations of the full range of climate change effects, guide the definition of effective mitigation measures and support data-driven decisions by political and administrative bodies. In addition, the platform will be used to forecast environmental extremes and disasters (like hurricanes, wildfires or floods) and enable effective responses to limit their natural, economic and societal impact.

The system relies on a large, distributed network of sensors providing real-time input data, which includes space assets, meteorological measurement stations, sensors in planes and ships, buildings and in perspective also sensors in mobile devices. The data coming in is of a very diverse nature - it includes time-series of scalar and vector data (like temperature, wind direction and speeds) all across the globe (and up to the tropopause), digital multi-spectral images and videos, plus potentially acoustic and seismic data. In addition, much of the incoming data has to be calibrated and assimilated with the more controlled traditional meteorological data. This task clearly requires significant data processing on the edge and in cloud layers, combined with a massive and reliable storage system, since no byte should ever be left behind.

The simulation of the natural phenomena considered based on the input data is a classical HPC task, and proven technology does exist in this area. However, the scale of the simulations is unheard of due to a significant increase in resolution/detail, and capturing the interactions between processes such as weather, ocean dynamics, and hydrology is very challenging. All simulation results will be annotated and stored forever, to serve as inputs for science, basis for policy decisions, or as a prerequisite for improving and calibrating models. The amount of data generated here will likely dwarf the input data. To reduce equipment costs, as well as time and energy to solution, advanced AI/ML techniques will be used. The natural location for the simulation activities are supercomputer centres, working together with large Cloud-based storage centres. The AI/ML models used will be trained on suitable systems in the CC, and regular validation and if necessary re-training will be performed.

To explore effects of human activities, be it the long-term impact of different economic and societal measures on climate change or the short-term effects of disaster relief, reduced-order models will be created and kept up-to-date. Such models can run for inference in smaller, Cloud-based data centers, or at the edge, close to, for instance, a flooded area. The initial creation and update of such models will likely happen in other parts of the CC. The number of end users of such models, or of the simulation data, will be very large, so effective data distribution and replication techniques will need to be implemented, and secure authenticated access has to be provided using scalable and intuitive Cloud mechanisms.

Taken together, Destination Earth will exercise all CC infrastructure and software layers, integrating data acquisition and assimilation, simulation of multiple scenarios, creation of forecasts, creation and calibration of both forecast and reduced order "decision" models, reliable storage of all input and simulation data, and handling of unforeseen user requests with highest priority (such as for disaster relief). The Destination Earth infrastructure will be mission critical, requiring special attention to ensure high availability, and it will need to handle tens of thousands of users from science, government, and industry.

AUTONOMOUS VEHICLE MANAGEMENT

Autonomous, self-driving vehicles need to rely on powerful edge computing support for their operation. On-board sensors deliver a continuous, high-bandwidth stream of data about the vehicle's environment (through videos or radar/lidar), neighboring vehicles might share their data, and internal sensors reflect the condition of the vehicle itself. All this information is processed under hard real-time requirements by AI/ML based software components, resulting in specific steering decisions which will protect the safety of the vehicle and its environment while proceeding on a course to the eventual target. Due to the real-time schedule, such AI/ML inference has to happen on an Edge system, be it local systems in the vehicle or in a close-by location (near to a rail track/road).

The AI/ML models used have to work for a truly overwhelming number of complex scenarios, and it is clear that regular updates of such models will be necessary. Anticipating all conditions a priori (or even after a long test driving campaign) seems unfeasible. In particular, accidents or "near misses" will no doubt happen, and "safe fail" mechanisms will be triggered if the main AI/ML

12. https://digital-strategy.ec.europa.eu/en/library/destination-earth

engine cannot reach an optimal decision. In these cases, the underlying models must be updated, using the actual incident data, to prevent accident scenarios from being repeated, and to cover corner cases by the main AI/ML engine.

For a large fleet of autonomous vehicles in day-to-day operation, a CC use case can be derived: a large number of edge devices send incident reports and large-volume data to large HPC installations, which will update or re-train the AI/ML models used by the edge systems, and regularly download updates to all managed edge systems.

The system will process private/protected data, and for a multi-tenant scenario, confidentiality about AI/ML models and vehicle-specific data has to be guaranteed. Should severe accidents happen, their data has to be fed into a high-priority training and update operation, whereas incidents relating to possible optimisations (such as avoiding triggering safe-fail) might conceivably be batched together.

The recently disclosed supercomputing installation of Tesla[13] shows the scale of HPC platform needed for the initial creation of models for self-driving automobiles by a single vendor. Our CC use case scales by the number of client edge systems/vehicles, the number of different versions (we would anticipate that such a CC infrastructure would be used across vendors), and the aggregated usage time (or mileage) of the vehicles. The resolution and complexity of environment data will play a role in scaling, too. This clearly shows that leading-edge HPC capacity will be needed for re-training, in addition to a highly capable data transmission and aggregation system handling many millions of endpoints, ways to prioritize urgent updates, storage systems able to keep all transmitted vehicle and environment data, and a scheduling/orchestration system which will safeguard the creation and distribution of urgent updates.

WHERE ARE WE NOW?

Today's landscape: an archipelago of disconnected solutions. Today's software approaches available to address the needs of CC use cases consist of separate software stacks optimised for different goals, specific to the target infrastructure (supercomputer, cloud datacenter, edge devices, respectively). There is no simple solution making it possible to deploy and orchestrate a combination of consistent interoperable components across the full continuum.

Moreover, the existing stacks are often specialised to be applied in their prevalent use cases; for example, combining an ab-initio QCD simulation with an AI engine for assessing the resulting, macroscopic characteristics of materials characteristics is extremely hard, even if the entire code was to run on a single system. Below, we list the major hurdles to overcome:

- Multitudes of software development stacks are tailored to specific use cases, with no guarantee of interoperability and composability between them. This greatly impedes application software development for integrated CC use cases.
- Across HPC, data analytics and AI, the existing software stacks have very different requirements for their execution infrastructure, and cannot be run efficiently on a single, homogeneous system.

13. O. Peckham, "Ahead of 'Dojo,' Tesla Reveals Its Massive Precursor Supercomputer," HPCwire, no. https://www.hpcwire.com/2021/06/22/ahead-of-dojo-tesla-reveals-its-massive-precursor-supercomputer/, 22 06 2021.

- The different compute, storage and communication systems of a complex CC installation will belong to different owners and be operated according to different rules:
  - Common abstractions of AAA (Authentication, Authorisation and Accounting) plus resource usage policies and pricing are needed; lack of this has limited the uptake of grid computing in the past.
  - Similarly, compatibility and interoperability across all parts of a CC infrastructure must be assured; this includes data formats, communication, security/encryption, data processing paradigms, etc. Suitable standards and their wide-spread adoption will play a very important role here.
  - Large-scale heterogeneity has to be managed in an effective and efficient way. This again cuts across compute, storage and communication systems, and the scheduling/orchestration middleware has to optimize the mapping of workflows with regards to performance and energy use.
- Security of platforms, communication links and data has to be assured, and specific rules & restrictions for the access to private or confidential data have to be enforced.
- On the hardware and systems side, the basic building blocks are common between the HPC, data analytics and AI/ML field, yet the system architectures and configurations are tailored to their respective usage area and do differ significantly. Flexible and efficient operation of CC infrastructures will require either reconfigurability at the system level or very agile mapping of processing steps to different systems.

BUILDING INTEGRATED SOFTWARE ECOSYSTEMS FOR THE CONTINUUM: CHALLENGES

Moving from an archipelago of disconnected software and system solutions towards integrated ecosystems requires addressing the aforementioned hurdles. This leads to challenges at multiple levels.

APPLICATION-LEVEL CHALLENGES

At the application level, traditional physics-based simulations (traditionally executed on HPC systems) need to smoothly cooperate with data-driven, learning-based analytics and prediction engines (typically run on clouds). Programming the workflow at the highest level requires the ability to consistently combine all these components in a unified framework. This requires innovative algorithmics and flexible programming models and supporting environments, which also safeguard performance and energy-efficiency.

Composability (the ability to combine multiple programming models or software stacks for a single application with defined rules) will be needed. Workflows could then combine the use of different robust programming models to enhance usability and achieve efficiency using standardized interfaces within and between workflow steps.

STORING AND PROCESSING DATA ACROSS THE CONTINUUM

Mastering the data flows from edge devices to clouds and supercomputers requires flexible, **shared data abstractions and unified mechanisms for data storage, to support distributed processing and analytics** across the whole CC infrastructure. This further increases the relevance of the classic Big Data challenges (e.g., Volume, Velocity, Variety), which can be projected in specific ways on the continuum:

Coping with Extreme Volume. Hybrid workflows running across the CC might need to process both synthetic data generated by simulations and real, sensor-originated data. Therefore, data volumes are virtually infinite.

The storage infrastructure needs to support the access and processing of “cold”, historical data and “hot”, real-time data (possibly streaming data), which further accentuates the volume pressure. Data lifecycle management becomes a key aspect which must consider technical, legal and application requirements and restrictions.

Coping with Extreme Velocity. The data processing ecosystem typically needs to combine processing of historical data with real-time data processing (e.g., in-situ/in-transit processing on HPC systems and cloud-based or edge-level stream-based processing) in a unified way. Unifying data processing approaches in a common software ecosystem becomes a challenge. Focusing on energy efficiency will require data movements to be kept to a minimum.

Coping with Extreme Variety. Unified data storage abstractions and systems enabling efficient data sharing to enable distributed processing and analytics across the Computing Continuum will have to overcome the extreme variety challenge: 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”. This requires:

- Interoperability of the data exchange formats.
- «Semantic interoperability” through shared ontologies understandable at all levels.
- Storage interfaces should match the needs of the application, e.g. by supporting object or key-value stores.

MANAGING COMPUTATION ACROSS THE CONTINUUM

At the middleware level, a major challenge is to facilitate the development of tools for seamless deployment, orchestration, scheduling, and execution of complex workflows across hybrid, heterogeneous CC infrastructures. This in turn calls for:

- Dynamic scheduling and orchestration of workflows which evolve at runtime, optimizing time and energy to solution on a dynamically evolving system.
- Support of a wide variety of processors, accelerators, storage devices and systems, and communication systems. A CC infrastructure will be deeply heterogeneous, and very likely so even within each of its layers.

- Support seamless deployment and migration of workflows or workflow steps, maybe aided by containerisation methods.
- Monitoring systems which provide the data required by adaptive, dynamic schedules to efficiently run workflows and workflow steps; recent advances in the support of micro-services on Cloud data centres may be relevant here.
- Definition and automatic derivation of performance models, which would enable a priori scheduling decision to be taken with a high degree of confidence.
- HPC centres should offer APIs to be integrated into workflows and for monitoring.

MANAGING DYNAMIC WORKFLOWS WITH AD-HOC LOAD VARIATION

In complex workflows, combining simulation and real-time data analytics, there are situations when one must react to certain events, depending on data contents or depending on interactive requests. This typically happens when on-demand simulations are triggered with real-time constraints for rapid decision making. The use cases described above exhibit a need for efficient management of dynamic load variations across the continuum.

In complex workflows executed over a CC infrastructure, changes in data characteristics and dynamic changes in the infrastructure create the necessity to dynamically adapt the mapping of the workflow onto the infrastructure and swiftly reconsider on which resources the different workflow steps are to be executed. In some applications, such as disaster warning and response[14], this need may appear when parts of the infrastructure suddenly become unavailable. T his requires efficient coupling between Cloud-oriented dynamic orchestrators and traditional batch-based resource management systems, as a step towards more integrated software approaches to dynamic resource management across the continuum. It might even be necessary to question the batch-oriented job scheduling in HPC systems.

AI-RELATED CHALLENGES

AI-powered workflows running on HPC-enabled infrastructure are gaining momentum. Their potential execution on hybrid infrastructures create new challenges:

- Optimising resource usage for AI workflows requires strategies to improve resource utilisation by automatically rescheduling jobs to best suited hardware option which in fact maximises the energy efficiency and reduces unnecessary allocation of hardware.
- The new heterogeneity of use cases and hardware must be taken into account;
- The deep learning software stacks must be supported (python dependencies handling, containerisation);
- Ad-hoc training and inference runs with tight timing constraints must be supported (urgent and interactive computing);

CYBERSECURITY CHALLENGES

To effectively leverage a CC infrastructure, significant challenges relate to cybersecurity[15], when it comes to federated authentication, authorisation and accounting, monitoring, resource allocations, encryption, user insulation, container certification, etc. These important considerations are currently impeding the deployment of large scale workflows. This is especially true for handling GDPR-related data. HPC centres that consider being involved in workflows handling such data must make sure to provide all tools necessary to address regulatory requirements.

COOPERATION CHALLENGES

Addressing the aforementioned challenges requires the collaboration of several expert communities (HPC, Big Data, AI, cybersecurity, IoT, 5G, etc.). Establishing commonly agreed, shared goals and priorities, as well as a common vocabulary and common roadmaps appears as a key strategic objective. Achieving this goal will enable strong synergies and well-founded, sound formulation of relevant and useful research directions. This is precisely the core motivation underlying the TransContinuum Initiative (TCI)[16], whose efforts are building the framework for such a cooperation. The open directions discussed in this paper on how to build an integrated ecosystem to leverage the CC are contributing to this effort.

CONCLUSIONS

In an increasing number of areas we are witnessing the emergence of complex workflows combining simulations, data analytics and learning, running on hybrid infrastructures where supercomputers are connected to cloud data centres and edge devices. To address the requirements of such workflows, we foresee a growing need for integrated software ecosystems which should build on state-of-the-art "island" solutions to bridge the gaps between them. This paper discusses the associated challenges and formulates recommendations to explore specific research topics which will contribute to this direction.

14. K. Fauvel, D. Balouek-Thomert, D. Melgar, P. Silva, A. Simonet, G. Antoniu, A. Costan, V. Masson, M. Parashar, I. Rodero and A. Termier, "A Distributed Multi-Sensor Machine Learning Approach to Earthquake Early Warning," in 34th AAAI Conference on Artificial Intelligence, Outstanding Paper Award - Special Track for Social Impact, New York, https://ojs.aaai.org//index.php/AAAI/article/view/5376 , 2020.
15. M. Asch, F. Bodin, M. Beck, T. Moore, M. Taufer, M. Swany and J.-P. Vilotte, "Cybercosm: New Foundations for a Converged Science Data Ecosystem," https://arxiv.org/abs/2105.10680v3 , 2021.
16. https://www.etp4hpc.eu/transcontinuum-initiative.html

6.1.3

HPC for urgent decision-making

Exploiting the full range of HPC capabilities to promptly gain insights and support decisions in a dynamic data-driven environment.

INTRODUCTION

Emerging use cases from incident response planning and broad-scope European initiatives (e.g. Destination Earth[1,2], European Green Deal and Digital Package[3]) are expected to require federated, distributed infrastructures combining computing and data platforms. These will provide elasticity enabling users to build applications and integrate data for thematic specialisation and decision support, within ever shortening response time windows.

For prompt and, in particular, for urgent decision support, the conventional usage modes of HPC centres is not adequate: these rely on relatively long-term arrangements for time-scheduled exclusive use of HPC resources, and enforce well-established yet time-consuming policies for granting access. In urgent decision support scenarios, managers or members of incident response teams must initiate processing and control the resources required based on their real-time judgement on how a complex situation evolves over time. This circle of clients is distinct from the regular users of HPC centres, and they must interact with HPC workflows on-demand and in real-time, while engaging significant HPC and data processing resources in or across HPC centres.

This white paper considers the technical implications of supporting urgent decisions through establishing flexible usage modes for computing, analytics and AI/ML-based applications using HPC and large, dynamic assets.

The target decision support use cases will involve ensembles of jobs, data-staging to support workflows, and interactions with services/facilities external to HPC systems/centres. Our analysis identifies the need for flexible and interactive access to HPC resources, particularly in the context of dynamic workflows processing large datasets. This poses several technical and organisational challenges: short-notice secure access to HPC and data resources, dynamic resource allocation and scheduling, coordination of resource managers, support for data-intensive workflow (including data staging on node-local storage), preemption of already running workloads and interactive steering of simulations. Federation of services and resources across multiple sites will help to increase availability, provide elasticity for time-varying resource needs and enable leverage of data locality.

KEY INSIGHTS

- **The need to give prompt access to users and entities for time-sensitive decision support in response to unexpected events does not fit the common access/use policies of HPC centres well. Moreover, such users rely on HPC workflows that are part of data-driven processes and might not have a long-term and direct connections to supercomputing centres.**
- **Use cases for insight extraction and decision support in short timescales require scheduling of resources in reaction to urgent demands, and is complicated by high dynamicity and diversity in the units of work to be managed; these include ensembles of jobs, data-staging to support workflows, as well as interactions with services or facilities external to the HPC systems/centres.**
- **Model-driven simulation and real-time interaction with decision makers requires interactive-style access to HPC resources, in the context of dynamic workflows processing large datasets. This leads to technical and organisational challenges. Pragmatically, we could consider maintaining both classical HPC clusters for batch-style use of computational and storage resources, together with clusters configured for dynamic workflows and interactive usage modes. For urgent decision making (such as evacuation and rescue operations planning coordination), the HPC resource requirements can be quite substantial; therefore, a static partitioning might turn out to be ineffective.**
- **Specific technical challenges arise in several areas: short-notice secure access to federated HPC and Cloud resources, dynamic resource allocation and scheduling (including large-scale use of accelerators), coordination of resource managers (including elastic resource allocation, unified authentication and authorisation, and usage accounting/monitoring), support for data-intensive workflows plus data staging on node-local storage, and increased interactivity (including near real-time steering of simulations and dynamic preemption of already running jobs as needed to free-up resources).**
- **The current state-of-the-art at supercomputer centres fails to adequately support the objective of providing prompt decision support.**

1. Destination Earth (DestinE) initiative. https://ec.europa.eu/digital-single-market/en/destination-earth-destine
2. Destination Earth: Use Cases Analysis, JRC Technical Report JRC122456, 2020. https://publications.jrc.ec.europa.eu/repository/handle/JRC122456
3. https://ec.europa.eu/info/strategy/priorities-2019-2024/european-green-deal_en

KEY RECOMMENDATIONS

- **Advance the state-of-the-art in short-notice secure access to federated HPC resources, system software and management to address challenges in resource allocation and scheduling, data-intensive workflow support (including data staging on node-local storage), and increased interactivity (including near real-time computational steering and preemption of jobs).**
- **Introduce enhancements and extensions to the relevant HPC programming models and APIs which will support expressing and managing complex and dynamic units of work, including ensembles of jobs, data-staging to support workflows, and interactions with services/facilities external to the HPC system.**
- **Initiate research and innovation actions, in a timeframe of 3-4 years, towards enabling interactive-style usage of HPC resources for near real-time workflows in support of insight extraction and decision making.**
- **Use Digital Twins to mirror various facets of complex events or systems, and develop predictive models to predetermine which line of action would be the best to take when a similar situation is evolving in real-time. The Digital Twins would run HPC simulations without the need for fast, interactive turnaround (i.e. offline operation).**
- **Continuously evaluate and improve reference system architectures and interoperability standards to support the effective operation of HPC environments for urgent decision making.**

This white paper considers the implications of flexible and convergent usage modes for data assets and HPC infrastructure to support modelling/simulation software, as well as analytics and AI/ML-based applications. Emerging use cases from broad-scope initiatives (e.g. Destination Earth, European Green Deal and Digital Package) are expected to be built on top of federated platforms with some degree of resource elasticity and supporting the combined use of data and infrastructure, where users can build applications and integrate their data for thematic specialisation for the purposes of decision support in short timescales, including urgent response to evolving complex situations.

Figure 1 illustrates an indicative flow for decision support workflows, relying on model-based simulation and data analysis and expected to be triggered either as part of regular proactive preparation for incident response or as an urgent reactive response to unexpected events. In this white paper, we focus on decision-making support based on model execution and data processing in latency-sensitive response to evolving situations. We assume that model development and the evolution of the corresponding codebase are carried out offline, although with interactions with the data and code assets on the decision support environment. The assessment of model quality and applicabili-

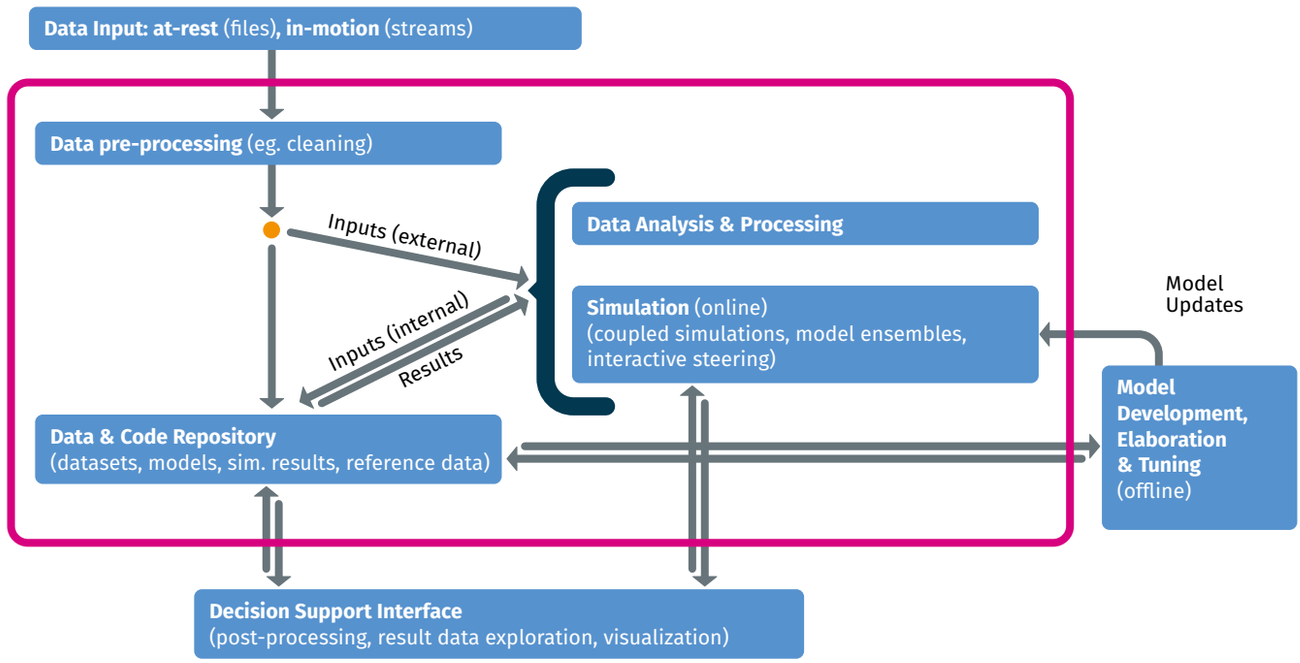

Figure 7: Decision support flow based on model-based simulation and data analysis. This white paper focuses on model execution and data processing in latency-sensitive response to evolving situations. Model development and the evolution of the corresponding codebase are assumed to be carried out offline, although with interactions with the data and code assets on the decision support environment

ty of results is also assumed to be carried out offline. Likewise, the decision support interfaces per-se (including data processing for exploration and visualisation) are outside the scope of this white paper. We focus on the requirements regarding the access to HPC and data processing resources, particularly HPC platform access, data access and workload management in latency-sensitive use cases.

We mention in passing that this type of decision support flow would also facilitate data-driven workflows in several scientific disciplines (e.g. Earth Observation, Astrophysics, Biomedicine), with a strong focus on exploratory data processing . However, such scientific workflows usually lack the urgency of having to respond within hard deadlines, with unforeseeable start times and with the involvement of non-registered decision makers.

CONTEXT AND USE CASES

This section considers requirements coming from a set of use cases that illustrate the expectations in terms of urgent computing from existing technologies[4]. The overall objective of the Destination Earth Initiative (DestinE) is to develop a "Digital Twin of the Earth"[5]- i.e. a digital modelling platform to "visualise, monitor and forecast natural and human activity on the planet in support of sustainable development". Accessible and interoperable data, combined with digital infrastructure and AI solutions, facilitate evidence-based decisions and expand the capacity to understand and tackle environmental challenges.

In the DestinE initiative[6], use cases that require urgent reactions are related to **disasters like earthquakes, tsunamis, forest fires, storms**. However other important events must be monitored such as floods, radiological incidents, epidemic outbreaks, but also astrophysical events ("space weather") as particle eruptions (flares) emanating from the Sun and impacting the Earth's magnetosphere with potentially dire disruptions of satellite operations[7] or telecommunications including the internet.

Such extreme events require a rapid response, to prevent or at least minimise human injuries or loss of life, and damage to infrastructure or property. In the worst case, the response would be in the form of data-driven management of disaster relief. The decision makers (persons or organisations) will require accurate information as quickly as possible to choose the best course of action. As clearly stated in the DestinE announcement[8], there are different classes of stakeholders: civil protection agencies, scientists and expert users, public authorities and humanitarian aid organisations. As a further example from the DestinE set of use cases, extreme weather constitutes "the second most likely global risk to the economy" and failure of climate change mitigation and adaptation is seen as "the highest global risk in terms of impact on the economy". Thus, DestinE use case #1 - "Accelerating weather-related disaster risk management" focuses on the urgent need to "improve disaster risk management with timely information on exposure of human settlements and fine scale hazard and risk information accelerating weather-related disaster risk management".

At the same time, there exists a large number of use cases that do not require strictly "urgent" reactions but rather either **rapid analysis or "medium-term" predictable responses**[9,10,11]] for ever increasing (spatial and temporal) resolution. As an example, for Climate Modelling scientists, public authorities and private sector entities continuously simulate advanced and well-validated high-resolution Earth System Models (ESMs), which are the primary tools for making future projections of global climate change. An ESM is a coupled climate model that also explicitly models the movement of carbon through the entire earth system, with consideration of physical, chemical and biological processes[12]. They are linking such projected changes to allowable carbon emissions commensurate with staying below a given warming target. However, they also guide short-term and mid-term actions affecting agriculture, forest environment, and day water management, from experts, policy officers, local and global authorities. A particularly acute challenge is the sheer volume of the data resulting from weather and climate simulations with a high spatial resolution[13].

Advances in technology over the past years have opened up **many new opportunities in aiding urgent decision-makers**. In the Big Data domain, a huge amount of data is collected from numerous sources, like satellites, instruments, IoT sensors, and even social media sources such as Twitter[14]. There is a tremendous wealth of information in all this data if it could be mined, aggregated,

4. F. Løvholt, S. Lorito, J. Macias, M. Volpe, J. Selva and S. Gibbons, «Urgent Tsunami Computing,» 2019 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC), 2019, pp. 45-50, doi: 10.1109/UrgentHPC49580.2019.00011.
5. Destination Earth (DestinE) initiative. https://ec.europa.eu/digital-single-market/en/destination-earth-destine
6. Destination Earth (DestinE) initiative. https://ec.europa.eu/digital-single-market/en/destination-earth-destine
7. R. Kube et al., «Near real-time analysis of big fusion data on HPC systems,» 2020 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC), 2020, pp. 55-63, doi: 10.1109/UrgentHPC51945.2020.00012.
8. Destination Earth (DestinE) initiative.https://ec.europa.eu/digital-single-market/en/destination-earth-destine
9. R. Kube et al., «Near real-time analysis of big fusion data on HPC systems,» 2020 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC), 2020, pp. 55-63, doi: 10.1109/UrgentHPC51945.2020.00012.
10. A. Kremin, S. Bailey, J. Guy, T. Kisner and K. Zhang, «Rapid Processing of Astronomical Data for the Dark Energy Spectroscopic Instrument,» 2020 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC), 2020, pp. 1-9, doi: 10.1109/UrgentHPC51945.2020.00006.
11. Jiang, M., Bu, C., Zeng, J. et al. Applications and challenges of high performance computing in genomics. CCF Trans. HPC (2021). https://doi.org/10.1007/s42514-021-00081-w
12. Earth System Modeling Framework : https://earthsystemmodeling.org/
13. T. C. Schulthess, P. Bauer, N. Wedi, O. Fuhrer, T. Hoefler and C. Schär, «Reflecting on the Goal and Baseline for Exascale Computing: A Roadmap Based on Weather and Climate Simulations,» in Computing in Science & Engineering, vol. 21, no. 1, pp. 30-41, 1 Jan.-Feb. 2019, doi: 10.1109/MCSE.2018.2888788.
14. CISCO 2020, Global Network Trends Report, Tech. rep., CISCO. URL https://www.cisco.com/c/dam/m/en_us/solutions/enterprise-networks/networking-report/files GLBL-ENG_NB-06_0_NA_RPT_PDF_MOFU-no-NetworkingTrendsReport-NB_rpten018612_5.pdf

processed, analysed, visualised and compared with numerical simulations[15]. The use of data (both observed and simulated) is crucial in a prompt decision-making environment. In particular, there are substantial benefits in facilitating the creation and access to data that are Findable, Accessible, Interoperable and Reusable without or with minimal human intervention (FAIR principles[16]), giving computational systems the ability to find, access, process, and reuse data. Moreover, following the Digital Twin approach, forecasting models could be continuously calibrated and validated.

Data comes from primary sources (e.g. physical measurement streams) and from processing pipelines (e.g. output of filtering steps or models) that may include extensive simulations, or ensembles of interrelated simulations. Incident response scenarios introduce additional constraints due to the use (and generation) of sensitive datasets, governed by stricter access and usage policies. Data annotation and metadata standards as well as established data curation processes are central to ensure that findings meet high standards of credibility and reproducibility. Metadata provides context and provenance to raw data and methods; therefore, they are essential to both discovery and validation. A good example is the meteorology field, which uses data and metadata formats standardised by the World Meteorological Organisation (WMO).

With the rapidly increasing power of HPC machines and the advent of Exascale computing, critical processes can be modelled and simulated much faster than in the past, making it possible to work on short- and mid-term predictions and analyses together with long-term studies. Examples are **complex natural phenomena** such climate change, ocean environment and sea level rising. The efficient and effective exploitation of large-scale computing capabilities is fundamental to support prompt data-driven decision making, as it is to obtain new scientific insights. A combined use of post-processing and visualisation serves to reduce, analyse and explore large volumes of data resulting from simulations. Modelling and data processing codes must be able to efficiently use heterogeneous systems that combine general-purpose processors with accelerators, high-speed networks and multiple classes of memory/storage devices, that represent the state-of-the-art and future trends in supercomputing resources (for an example, see the directions set by the EuroHPC Joint Undertaking[17]).

The opportunity of using cutting edge computing systems is not just a way to reduce the time-to-solution, but the sole viable approach to process datasets of the size and complexity expected in the type of use cases considered[18,19]. The observed data can be used as input for a numerical simulation that must then be post-processed and delivered to the decision makers or domain experts so they can take appropriate actions in near real-time. Moreover, they must be compared with theoretical simulations to enable domain experts to understand nuanced predictions, as well as shape experiments more efficiently (see Figure 1). The time scale for the analysis and simulations to be completed depends on the specific use case. However, the ability to intervene and modify predetermined, batch processing schedules to secure the resources needed ad hoc is essential for generating actionable results within the tight time windows for events whose occurrence cannot be predicted a-priori.

USING CUTTING EDGE COMPUTING SYSTEMS IS NOT JUST A WAY TO REDUCE THE TIME-TO-SOLUTION, BUT THE SOLE VIABLE APPROACH TO PROCESS DATASETS OF THE SIZE AND COMPLEXITY EXPECTED

Data exploration and visualisation are further crucial aspects of the decision-making workflow, relying on the automated analysis and determination of possible actions (plus their consequences). Data exploration is the essential initial step in data analysis, where experts and policymakers explore a large data set in an unstructured way to uncover initial patterns, characteristics, and points of interest. Data exploration can use a combination of manual methods and automated tools to generate artefacts for decision support, such as data visualisations, charts, and initial reports. Data visualisation assists experts and policymakers in reaching a clear understanding of the information, by giving it visual context through maps or graphs. It is possible to identify different scenarios where data exploration is implemented for urgent decision making - for example: automatic decision based on pre-defined or AI driven approaches[20,21] where visualisation is used as a means for post-mortem analysis of the outcomes resulting from human-driven decision making where exploration is necessary to confine available data and focus the decisions on relevant information. As in AI models, there is the concern of explainability[22], i.e. ensuring that humans can adequately un-

15. Asch M, Moore T, Badia R, et al. Big data and extreme-scale computing: Pathways to Convergence-Toward a shaping strategy for a future software and data ecosystem for scientific inquiry. The International Journal of High Performance Computing Applications. 2018;32(4):435-479. doi:10.1177/1094342018778123
16. Wilkinson MD, Dumontier M, Aalbersberg IJ, et al.. The FAIR Guiding Principles for scientific data management and stewardship. Sci Data. 2016 Mar 15;3:160018. doi: 10.1038/sdata.2016.18. Erratum in: Sci Data. 2019 Mar 19;6(1):6. PMID: 26978244; PMCID: PMC4792175.
17. https://eurohpc-ju.europa.eu
18. CISCO 2020, Global Network Trends Report, Tech. rep., CISCO. URL https://www.cisco.com/c/dam/m/en_us/solutions/enterprise-networks/networking-report/files/GLBL-ENG_NB-06_0_NA_RPT_PDF_MOFU-no-NetworkingTrendsReport-NB_rpten018612_5.pdf
19. F. Løvholt, S. Lorito, J. Macias, M. Volpe, J. Selva and S. Gibbons, «Urgent Tsunami Computing,» 2019 IEEE/ACM HPC for Urgent Decision Making (UrgentHPC), 2019, pp. 45-50, doi: 10.1109/UrgentHPC49580.2019.00011.
20. Tzachor, A., Whittlestone, J., Sundaram, L. et al. Artificial intelligence in a crisis needs ethics with urgency. Nat Mach Intell 2, 365–366 (2020). https://doi.org/10.1038/s42256-020-0195-0
21. Chen, N., Liu, W., Bai, R. et al. Application of computational intelligence technologies in emergency management: a literature review. Artif Intell Rev 52, 2131–2168 (2019). https://doi.org/10.1007/s10462-017-9589-8
22. R. Roscher, B. Bohn, M. F. Duarte and J. Garcke, «Explainable Machine Learning for Scientific Insights and Discoveries,» in IEEE Access, vol. 8, pp. 42200-42216, 2020, doi: 10.1109/ACCESS.2020.2976199.

derstand and oversee algorithmic processes, and establishing ways for explainable decision-making technologies to be used in combination with domain knowledge from the application areas. Finally, HPC is part of an overarching framework of organisational procedures in support of urgent decision making; therefore, HPC resource acquisition and usage protocols should be regularly reviewed and rehearsed to ensure effective response to emergencies. However, human-driven analysis and exploration steps, and protocols for orchestrating emergency response, are outside the scope of this white paper.

IMPACT ON TECHNOLOGIES AND REQUIREMENTS

The most widely adopted and reliable method for supporting prompt decision making is the use of dedicated resources as in the case of the Japanese earthquake early warning system[23] or operational weather forecasting organisations (such as the European Centre for Medium-Range Weather Forecasts, Meteo France, UK Met Office, DWD). This is a cost-efficient way to ensure the required response times, in particular for the use cases where continuous computing is required, for example, weather forecasting or ocean and climate modelling. However, setting up dedicated resources for each domain-specific urgent computation, constantly ready and waiting to run forecasts for disasters which may never come, is economically unviable, in particular when computations take place rarely and require a huge amount of resources.

Additionally, having a fixed set of HPC resources for disaster response could also cause problems should multiple disasters occur simultaneously, as there would be no capacity to scale to larger computing resources if required. For most cases, it is therefore prudent to use time on existing international, national or regional HPC systems, or to combine dedicated resources with existing ones into a flexible and scalable environment. Rapidly engaging such resources poses significant challenges:

- Preempting already running, lower priority jobs to free required resources (rapid resource availability for urgent computing)
- Provisioning the complete workflow needed to react to an urgent (real-time or rapid analysis) event
- Performing data acquisition, execute simulations on HPC machines, run data analytics and presenting the results to decision-makers (e.g. public authorities, scientists, domain experts)
- Adapting the set of resources used and the workflow as the situation evolves in unpredictable ways.

Existing HPC infrastructures are not commonly used for urgent

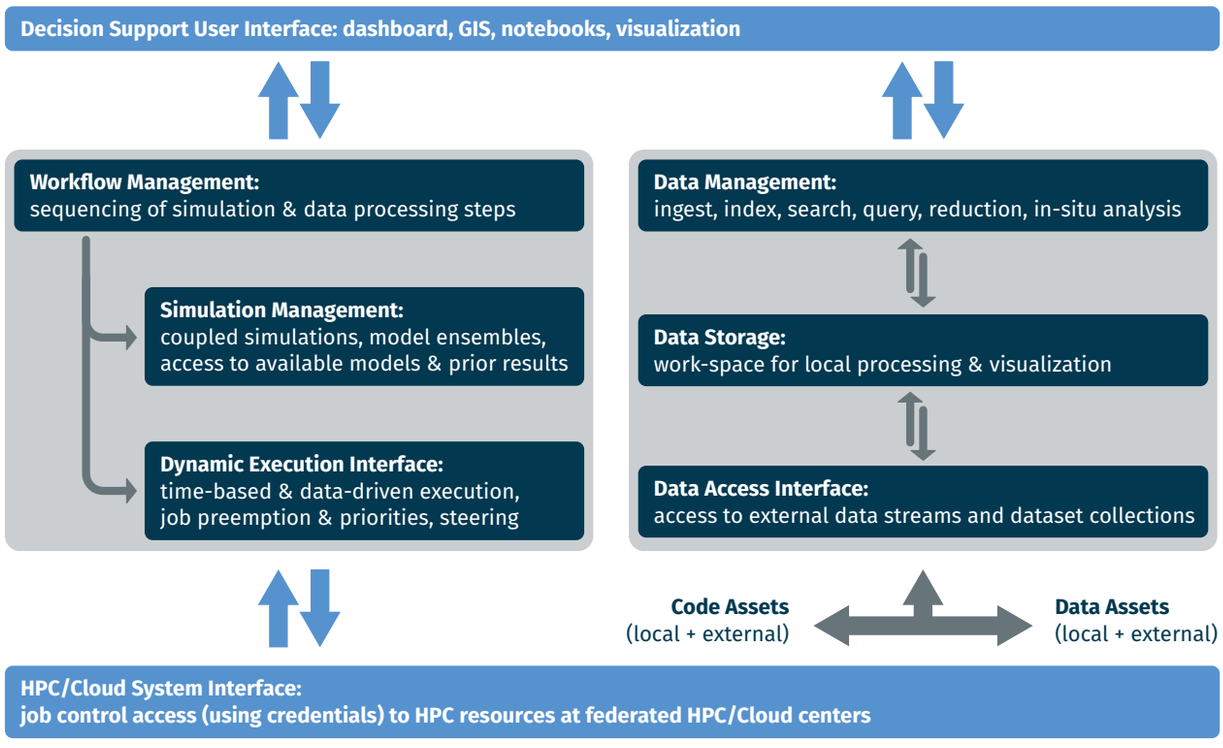

Figure 8: Architectural overview of an HPC-based urgent decision support environment.

23. E.Yamasaki, 2012, What We Can Learn From Japan's Early Earthquake Warning System, Momentum: Volume 1: Issue 1, Article 2.

computing, since rapidly mobilising them faces a few (but crucial) technological challenges. For example, the usual approach of reserving resources well in advance does not meet the requirements of urgent computing. Preemptive scheduling is an effective alternative to swiftly obtain resources on HPC infrastructures when an urgent need arises unexpectedly. This requires a change in the mode of operation, and in addition the use of novel technologies: for example, checkpointing any running applications that are preempted in a non-disrupting and efficient way becomes mandatory. Furthermore, challenges related to pre-emption in the context of high-performance interconnects and accelerated architectures need to be solved and resource occupation by preempted jobs has to be minimised.

CHALLENGES IN ADVANCING HPC FOR URGENT DECISION-MAKING

Figure 2 illustrates a high-level view of the architecture of an HPC-based system designed for urgent decision support. The main capabilities as expected are simulation and data management, but with additional features arising from the need to utilise distributed resources and data assets. The user interface can be part of a web portal for HPC resources.

The following list summarises ideas and challenges to consider in advancing the current state of the art for HPC-based urgent decision support:

- HPC systems tend to be optimised for job throughput so there could be very long (hours or even days) wait times in the batch queue until the resources for urgent computing become available. It is therefore necessary to **improve the management of job priorities and of job preemption, in the context of coarse-grain units of work**.
- Priority management and preemption are necessary for prompt response to incoming workloads associated with decision support workflows. Moreover, appropriate system software and application level **checkpointing and restart procedures** must be implemented to enable any preempted jobs to restart without losing significant work. Extensive usage of accelerators in HPC workflows is essential for prompt response; however, it exacerbates the complexity of state capture and restoration (in particular, in the context of checkpointing protocols).
- HPC systems are usually unique environments with specific interconnects and accelerators, as well as software sets customised to particular domains and use cases. Running an optimised code requires machine-specific porting. Deployment can be expedited with the use of (hardware platform-specific) **container technologies** that facilitate handling of software dependencies and portability.
- CPU and storage resources are unlikely to be co-located in urgent computing use cases. Input data must be accessible from the HPC systems inside supercomputing centres machines for processing, and at the same time results need to be stored and made accessible to decision makers. **Advanced data management and caching services** should be implemented, able to interoperate with modern data storage and management architectures such as data lakes, and support interaction with large-scale data repositories of both raw and transformed data. A key problem is the staging of required data onto high-performance parallel filesystems or node-local storage.
- Opening up HPC centres for utilising their resources on short-notice and in the context of dynamic workflows will exacerbate cybersecurity challenges due to the increased attack surface. Protection and validation of system and data integrity become particularly acute challenges. **Data integrity validation**, including reliable tracking of provenance, becomes a particularly acute challenge, throughout the entire lifecycle of large-scale data sets. Executing workflows involving diverse computing and data resources entails adding data transfers and storage solutions to the decision support workflow, which can make it more difficult to ensure that data have not been corrupted in transit or at rest. When data integrity is not preserved, computations can fail and result in increased computational cost due to reruns, or worse, results can be corrupted in a manner not apparent to data analysts, and produce invalid results with severe consequences to the real-world impact of the decision-making process.
- **Simplified access to HPC resources** and the results of the urgent computing analysis is of course mandatory: the ability to obtain context-relevant information quickly and easily allows decision makers to act and respond to findings swiftly. Science platforms[24] are a promising modern approach to interact with data, resources and software, and could provide a useful starting point. APIs and common standards should be defined to develop community/service specific platforms, including standards to facilitate interaction with different workload managers.
- **Federated HPC environments** involving supercomputing centres all over Europe can be instrumental to optimise resource usage for different urgent actions. This usage mode of HPC resources requires the following:
  - Ability to pick the proper supercomputer architectures and systems according to the code and workflow to be executed
  - Identification of cross-centre policies (rather than a large set of disparate HPC centre-specific policies) to allow the use of resources that are rare and extremely expensive, including authorisation, authentication and accounting policies
  - Definition of policies for storage access and data persistence.

The idea of a Federated HPC environment opens a new set of challenges, including unified authentication and authorisation infrastructure (AAI) for user management, elastic and resilient resource management, and the definition of an innovative resource information system able to monitor supercomputer health status and load, and eventually predict resource status in advance.

24. Strategic Research and Innovation Agenda of the European Open Science Cloud (EOSC), Feb. 2021. https://www.eosc.eu/sites/default/files/EOSC-SRIA-V1.0_15Feb2021.pdf

A common challenge present in a wide range of use cases[25,26,27] is the increasing dynamicity and diversity in the definition of what units of work have to be managed, including ensembles of jobs, data-staging to support workflows, and interactions with services/facilities external to the HPC system. Our analysis of salient features of these use cases identifies the need for advances towards **more interactive access to HPC resources**, particularly in the context of dynamic workflows over large datasets. Several challenges, both technical and organisational, have to be considered. A pragmatic approach is to maintain both classical HPC clusters for batch-style use of computational and storage resources, together with clusters more tuned for dynamic workflows and interactive usage modes. HPC centres supporting prompt decision making could provision (i) one set of resources for use by predictive simulations not directly involved in the decision making in near real-time, (ii) a specialised set of resources for interactive-style use by decision makers (e.g. visualisation and data exploration) , and (iii) one set of dedicated or at least preemptable resources for actual processing within tight time constraints. Nevertheless, we should keep in mind that the HPC resource requirements for urgent decision support can be quite substantial; therefore, a static partitioning might turn out to be ineffective, necessitating a set of technical improvements along the lines suggested in this white paper.

Major challenges arise in the following areas: dynamic resource allocation and scheduling, coordination of resource managers, short-notice secure access to federated HPC and Cloud resources, data-intensive workflow support (including data staging on node-local storage), increased interactivity (including preemption and simulation steering). Additionally, cybersecurity concerns need to be considered.

CONCLUSIONS

The current state-of-the-art does not adequately cover the use of HPC environments for prompt decision support. In particular, applying HPC in a wide range of use cases for gaining insight and supporting decisions in short time scales causes increased dynamicity and diversity in the units of work which have be managed. This includes ensembles of jobs, data-staging to support workflows, and interactions with services/facilities external to the HPC systems or centres. Model-based simulation, the real-time adaptation of interlinked computational models of evolving complex processes and distributed data-driven scientific collaborations require advances towards more interactive access to HPC resources and support for preempting running jobs. Challenges in the technical and organisational area have to be addressed, all in the context of dynamic workflows over large datasets.

Specific technical challenges arise in the following areas: dynamic resource allocation and scheduling, preemption techniques with transparent and efficient job checkpointing and restart, coordination of resource managers, short-notice secure access to federated HPC resources, data-intensive workflow support (including data staging on node-local storage), interactive style and near real-time use (including preemption and simulation steering). We have outlined a set of short- to medium-term recommendations for R&D actions which would enable the effective use of HPC resources for extracting insights and supporting decisions in short time scales.

25. Destination Earth (DestinE) initiative. https://ec.europa.eu/digital-single-market/en/destination-earth-destine
26. Jiang, M., Bu, C., Zeng, J. et al. Applications and challenges of high performance computing in genomics. CCF Trans. HPC (2021). https://doi.org/10.1007/s42514-021-00081-w
27. CISCO 2020, Global Network Trends Report, Tech. rep., CISCO. URL https://www.cisco.com/c/dam/m/en_us/solutions/enterprise-networks/networking-report/files/GLBL-ENG_NB-06_0_NA_RPT_PDF_MOFU-no-NetworkingTrendsReport-NB_rpten018612_5.pdf

6.1.4

Federated HPC, cloud and data infrastructures

An increasing interest is observed in making a diversity of compute and storage resources, which are geographic spread, available in a federated manner. A common services layer can facilitate easier access, more elasticity as well as lower response times, and improved utilisation of the underlying resources. In this white paper, current trends are analysed both from an infrastructure provider as well as an end-user perspective. Here the focus is on federated e-infrastructures that among others include high-performance computing (HPC) systems as compute resources. Two initiatives, namely Fenix and GAIA-X, are presented as illustrative examples. Based on a more detailed exploration of selected topical areas related to federated e-infrastructures, various R&I challenges are identified and recommendations for further efforts formulated.

INTRODUCTION

Federating geographically spread HPC resources as well as storage resources have been an attractive vision for a long time. It matches, for instance, the needs of various distributed research communities, allows for more flexible access to more resources and access to a broader diversity of computing resources[1]. In this white paper, we discuss and describe trends that drive the realisation of federated e-infrastructures that include HPC resources. The documentation of selected initiatives show-case these trends. We furthermore explore different topical areas related to federated e-infrastructures in more detail to identify different R&I challenges, which are the basis for our recommendations. Throughout this white paper, we will focus on federated e-infrastructures that include HPC resources and that require tight integration and close proximity of computing and storage resources.

Many efforts have been invested in the past on federating HPC resources as a means of realising e-infrastructures that allow addressing large-scale scientific challenges. Early examples are the following projects that all started in the early 2000s: TeraGrid in the US[2], DEISA in Europe[3] and NAREGI in Japan[4]. These projects had in common that they were based on Grid concepts and technologies, which had been developed mainly by academic organisations with rather limited commercial uptake. With the advent of commercial cloud providers that deployed federated compute and storage resources this situation fundamentally changed. Technologies required for federating compute and storage resources as well as services are now used at a much wider scale. Furthermore, they are driven by competing large-scale commercial entities, which increases the probability of different solutions being available.

One of the key assumptions made in this white paper is that infrastructures consisting of stand-alone HPC systems with attached high-performance storage resources will increasingly be replaced by e-infrastructures based HPC and cloud compute as well as storage resources that are made available through a set of federated services. In this context, federation means that the services are integrated in (one or more) access, identity and resource management mechanisms such that users can flexibly leverage different services, even if they are offered by different organisations at geographically different locations. These infrastructures should be realised without creating a risk for vendor lock-ins.

In the context of this white paper, it is important to note that the term "cloud" is used in different contexts. In the following, we will use this term to refer to technologies that are primarily used for realising public cloud infrastructures. We will not consider cloud business models although they may start to play an increasing role in future e-infrastructures that include HPC-based services.

TRENDS

In the past, the primary role of HPC centres was to operate one or more supercomputers and to provide access to this resource. In the future, we expect these centres to transform to providers of an *e-infrastructure services layer*, where the latter is based on different types of underlying compute and storage resources. While this will continue to include supercomputers, an increasing number of centres have also deployed cloud-type resources, i.e. on-premise private cloud instances. Most of the new EuroHPC pre-exascale and petascale systems do include such instances. These services started to become harmonised and federated across multiple sites, e.g. in the context of the Fenix initiative[5]. The trends sketched here have meanwhile become part of the future EuroHPC strategy, as EuroHPC added "HPC Federation and Services" as one of the five pillars of future activities[6].

1. Examples are the bioinformatics community ELIXIR https://elixir-europe.org/platforms/compute or various astronomy and particle physics communities such as John D. Swinbank et al. (ESCAPE), "D5.3 - Performance Assessment of Initial Science Platform Prototype", 2021 (https://projectescape.eu/sites/default/files/ESCAPE-D5_3.pdf).
2. Charles Catlett et al., "TeraGrid: Analysis of Organization, System Architecture, and Middleware Enabling New Types of Applications", in: L. Grandinetti (Ed.), "High-Performance Computing and Grids in Action", IOS Press, 2008.
3. Wolfgang Gentzsch et al., "DEISA - Distributed European Infrastructure for Supercomputing Applications", J Grid Computing, 2011 (DOI: 10.1007/s10723-011-9183-2)
4. Satoshi Matsuoka et al., "Japanese Computational Grid Research Project: NAREGI", Proceedings of the IEEE, Vol. 93, Issue 3, 2005 (DOI: 10.1109/JPROC.2004.842748)
5. https://www.fenix-ri.eu/
6. European Commission, "Equipping Europe for world-class High-Performance Computing in the next decade", SWD(2020) 179 final, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52020SC0179&rid=9.

With these architectural changes, HPC centres respond to changing user needs as well as those of emerging new science and engineering domains, which do need HPC resources for their research. A first trend to highlight are efforts towards establishing domain-specific *platform services layers* that facilitate collaborative research within geographically dispersed communities. There are various opportunities for creating benefits through such higher-level services (see also[7]):

- Web Portals hiding complexity: The existence of user-friendly web-based interfaces can improve efficiency and increase the use of the infrastructures. Science Gateways[8] are also used in HPC to increase the usability of modelling of data and simulations on top of complex computing infrastructures.
- The realisation of complex workflows: Realising workflows that combine simulations, big data processing and interactive steering can quickly become challenging due to the complexity resulting from the use of different, sometimes incompatible, tools and frameworks.
- Interactivity for enhanced execution steering: The choice of parameters and the tuning of the frameworks constitute a difficult task, which also impacts the application performance.
- Legacy applications transformation: The use of clouds and virtualization approaches enables the transformation of legacy applications for traditional computing infrastructures to a SaaS model (Software-as-a-Service).
- FAIR data management[9]: Platform services allow for integration with global data infrastructures and connect to data lakes that will grow to exabyte capacities.

Domain-specific platform services are being implemented among others by projects like the Human Brain Project[10] or different ESFRIs like those organised in the European Science Cluster of Astronomy & Particle physics ESFRI research infrastructures (ESCAPE)[11,12] or the ESFRI for life-science information Elixir[13].

There are several benefits that can be achieved by federating e-infrastructure services. Firstly, research communities can deploy domain-specific services without becoming dependent on a single resources provider. Secondly, enabling data sharing and exploiting data locality become easier in a federated infrastructure. Data does not have to be collected in a single location but can remain close to where it has been produced while compute resources can be allocated near the data. Furthermore, federated infrastructures can help to improve the availability of services by geo-replication of services and data. If one site becomes unavailable then both services and data remain accessible at other sites. Finally, federated infrastructures can provide more flexibility to serve a diversity of communities as more resources are available for allocation and a larger variety of services can be offered through different resource providers.

SELECTED INITIATIVES

There are a large number of initiatives throughout Europe towards the federation of e-infrastructure services. We introduce in the following two initiatives as illustrative examples that have a connection to HPC and involve various ETP4PHC members.

FENIX

Fenix[14,15] is a collaboration of HPC centres working on the harmonisation and federation of their offerings of e-infrastructure services with the goal of supporting a variety of science and engineering communities. This service portfolio's distinguishing characteristic is that different types of data repositories, scalable supercomputing systems, and private cloud instances are in close proximity and thus well connected and integrated. The different sites that act as e-infrastructure services providers are interconnected via a high-speed network. A key feature of the Fenix approach is the separation of concern with HPC centres being in the role of providing a set of federated e-infrastructure services that support user communities that deploy their domain-specific services. A schematic view is provided in *Figure 1*.

Fenix has the ambition of serving in a sustainable manner relevant science and engineering domains that strongly benefit from diverse e-infrastructure services for their collaborative research. For being able to scale to a larger number of such domains, Fenix focuses on a consolidated portfolio of services. To stay aligned with the needs of current and upcoming science and engineering domains, Fenix governance foresees a representation of these domains such that they can drive the evolution of the e-infrastructure services portfolio.

7. Marco A. S. Netto et al., "HPC cloud for scientific and business applications: taxonomy, vision, and research challenges", ACM Computing Surveys (CSUR) 51 (1), 1-29, 2019 (DOI: 10.1145/3150224)
8. S. Gesing, "Science Gateways in HPC: Usability Meets Efficiency and Effectiveness" In: Modelling and Simulation in HPC and Cloud Systems, Book Series «Studies in Big Data» Print ISBN: 978-3-319-73766-9, Electronic ISBN: 978-3-319-73767-6, Springer International Publishing, 2018
9. https://www.go-fair.org/fair-principles
10. See, e.g., EBRAINS Simulation services offering (https://ebrains.eu/services#category2)
11. https://projectescape.eu/
12. Rizart Dona, Riccardo Di Maria (ESCAPE project), "The ESCAPE Data Lake: The machinery behind testing, monitoring and supporting a unified federated storage infrastructure of the exabyte-scale", CHEP 2021, EPJ Web of Conferences, 2021 (DOI: 10.1051/epjconf/202125102060)
13. https://elixir-europe.org/
14. Fenix, "Fenix Strategy Document", 2021 (https://fenix-ri.eu/sites/default/files/public/file-uploads/Fenix%20Strategy%20Document%20202111111-public.pdf)
15. Sadaf Alam et al., "Archival Data Repository Services to Enable HPC and Cloud Workflows in a Federated Research e-Infrastructure", IEEE/ACM International Workshop on Interoperability of Supercomputing and Cloud Technologies, 2020 (DOI: 10.1109/SuperCompCloud51944.2020.00012)

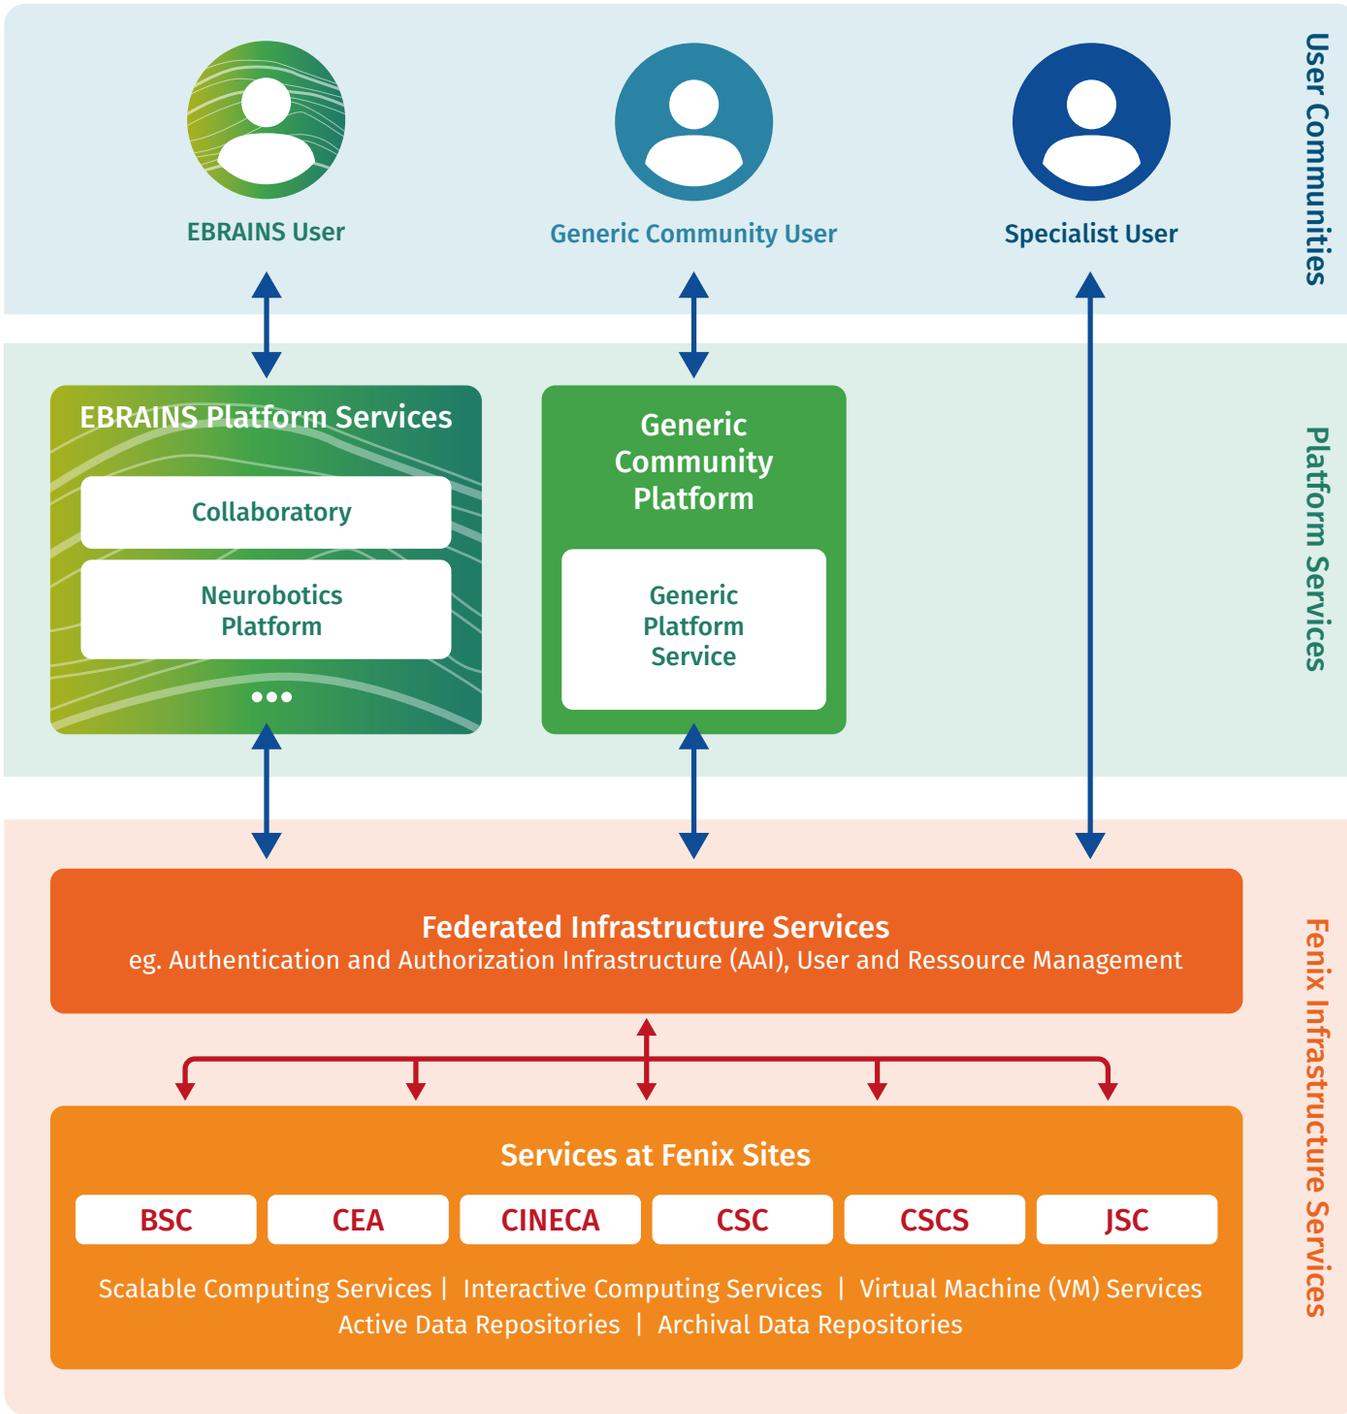

Figure 9: Schematic overview of Fenix's layered approach for support users from different science and engineering domains (source: Fenix)

Fenix members leverage national, European and international funding programs to realise the compute, storage and network resources sustaining the e-infrastructure services. The coherence of the approach is ensured through a clear governance model and close technical collaboration at various levels. Furthermore, the Fenix partners share the responsibility for the operation of federation-level services and the integration at each of the sites.

GAIA-X

The backdrop of Gaia-X can be traced back to late 2019. There had been a general concern in Europe for many years that big global hyperscalers (like Amazon or Alibaba) outside of Europe handle the data of European citizens and organisations. Hence, there existed a political will to create something in Europe that breaks European dependency on these big non-European hyperscalers. Thus arose Gaia-X, which was a Franco-German initiative to create a European data infrastructure and data economy.

In Europe, there are many challenges to creating an indigenous data economy. Firstly, there are decentralised data processing locations throughout Europe, which are not consolidated. These locations all use their own stacks. Interoperability and flow of data between these locations are not seamless. There is an absence of widely accepted APIs to share data and insufficient clarity about applicable jurisdiction. Also, there are sector-specific data spaces (e.g. automotive, financial, pharma) and sharing data across sectors is extremely hard.

The Gaia-X solution envisioned a sharable data ecosystem between various data spaces all relying on a federated infrastructure ecosystem to store data. The federated infrastructure ecosystem consists of existing network and interconnection providers (e.g., telcos), Cloud Service Providers (CSPs), sector-specific clouds, edge locations and HPC. The data spaces and the infrastructure ecosystem are all orchestrated by the Gaia-X Federation Services, which deals with issues of identity and trust management, provision and maintenance of a federated catalogue of infrastructure providers, and compliance management. It is envisioned that smart third party services (AI, IoT, etc) could be built on top of these data spaces. The key European data spaces identified are agriculture, energy, health, industry 4.0, mobility, public sector, smart living and finance.

More details can be found in the Gaia-X architecture document[16].

Some of the key considerations of the Gaia-X Infrastructure ecosystem are:

- Establishing portability and interoperability amongst the different infrastructure providers (HPC, edge, CSPs, etc);
- Providing the ability for finding, combining and connecting services from participating providers.

16. Gaia-X, “Gaia-X Architecture Document”, 21.09 release, 2021 (https://www.gaia-x.eu/sites/default/files/2021-10/Gaia-X_Architecture_Document_2109.pdf)

At the time of writing this white paper, Gaia-X has been fully set up as an AISBL organisation with a fully functioning organisational structure - Policy & Rules Committee, Data Spaces Business Committee and a Technical Committee consisting of various working groups looking into different aspects from the perspective of Users, Providers, Architecture, etc. Various Gaia-X hackathons are being organised to develop proofs-of-concepts supporting Gaia-X and the first version of the fully-functioning Federated Services is expected to become available in 2022.

HPC as a part of a Gaia-X infrastructure ecosystem is a subject of current discussions. HPC sites in Europe may become resource and service providers within Gaia-X. Some of the issues that need to be addressed are how the HPC services describe themselves within Gaia-X, what are the possible economic models for their engagement, how they work vis-a-vis other initiatives such as the European Open Science Cloud, etc. One possibility is the focus on industrial usage of HPC, since Gaia-X seems to be heading in the direction of being Europe's "Industrial Federated Cloud", considering very strong participation from the Industrial actors. The focus on HPC within Gaia-X creates the opportunity for making HPC services available to a wider user community as well as to drive the development of solutions for federated services based on HPC resources.

TOPICAL AREAS

In this section a few topics are considered that are relevant for realising future federated e-infrastructures including HPC resources.

AUTHENTICATION AND AUTHORISATION INFRASTRUCTURE

Federation of e-infrastructure services requires integration in a common Authentication and Authorisation Infrastructures (AAI) such that a user with a single identity as well as one (or few) credentials can authenticate with different services.

A typical AAI architecture for federated e-infrastructures foresees, on the one hand, one or more identity providers and on the other hand a set of service providers, which are geographically distributed in case of federated e-infrastructures considered here. Different initiatives, including Fenix or EOSC[17], foresee such an architecture as they follow the architecture proposed by the AARC project[18] with important elements being realised by eduGAIN[19] and MyAccessID[20]. One key challenge for realising such an architecture is the establishment of suitable trust relations between identity providers and service providers. This is particularly relevant for sensitive services, which includes any service that provides direct access to HPC resources. For such services, a high level of assurance is required, while other services (e.g. web portals) may require only a low level.

For a federated e-infrastructure, authentication mechanisms need to be harmonised and support multiple authentication mechanisms to facilitate user-friendly access to different services. While authentication to HPC resources today typically relies on authentication using SSH, different technologies like OpenID Connect (OIDC)[21] are used for cloud-type services. At the same time, security will have to be pushed to a higher level, e.g. by replacing today's dominant single-factor by multi-factor authentication mechanisms. In the context of a federated e-infrastructure, there are additional security concerns as a single compromised identity would put more service providers at risk. Today many services are accessible using a single-factor for authentication. There is an emerging consensus that soon multi-factor authentication will be required, which is likely to become mandatory due to European or local regulations.

Also, authorisation mechanisms need to be established at the level of a set of federated services. While this may generically be solved through attribute services as foreseen in the AARC architecture, in practice interoperability of different solutions is typically at best at an early stage.

INTEGRATION OF HPC- AND CLOUD-BASED COMPUTE AND DATA SERVICES

For multiple reasons, services based on HPC and cloud instances are still difficult to integrate. While direct access to HPC resources requires strict control, cloud instances are typically more openly accessible. Another reason is the diversity of how HPC resources are operated at different sites. This makes, for instance, the deployment of web-based portal services that facilitate the spawning of HPC jobs difficult and tedious due to the required customisation. Such kind of functionality is expected to become increasingly important in the context of community-specific platform services layers (see Section 2), the use of HPC for urgent decision making (see[22]) or the realisation of use cases identified by the Transcontinuum Initiative[23].

Similarly, federating HPC-based services and data services remains challenging as integration is still often proprietary. An increasing number of data platforms are created and realised through lake concepts, but implementing workflows that consume data available on such platforms or publish data produced on the HPC system often requires manual steps. They can often not be automated as handling of the necessary credentials is not supported.

RESOURCE MANAGEMENT, ALLOCATION AND ACCOUNTING

HPC and cloud infrastructures tended to adopt different approaches to resource management. In the area of HPC, resource

17. European Commission, Directorate-General for Research and Innovation, Vaghetti, D., Kanellopoulos, C., Johansson, L., et al., "EOSC Authentication and Authorization Infrastructure (AAI): report from the EOSC Executive Board Working Group (WG) Architecture AAI Task Force (TF)", Publications Office, 2021 (https://data.europa.eu/doi/10.2777/8702)
18. Nicolas Liampotis et al., "AARC Blueprint Architecture 2019", 2019 (DOI: 10.5281/zenodo.3672785)
19. https://edugain.org/
20. https://wiki.geant.org/display/MyAccessID/MyAccessID+Home
21. https://openid.net/connect/
22. Manolis Marazakis et al., "HPC for urgent decision making", ETP4HPC White Paper, 2022 (DOI: 10.5281/zenodo.6107362)
23. https://www.etp4hpc.eu/tci-use-cases.html

managers compromise on response times to maximise resource utilisation without oversubscribing compute resources; they are allowed to significantly delay start-up of jobs. In the area of cloud, the focus is rather on short response times and elastic use of the available hardware resources. Another difference is that HPC resources are typically assumed to be concurrently used by a smaller number of long-running large-scale jobs (in the extreme case a single job), while cloud computing resources may be used by a very large number of small microservices, which may run for only short periods of time, so that set-up and tear-down latencies become critical. Both HPC and cloud computing resource management is becoming more complex when underlying hardware resources become more heterogeneous (for instance, when part of the servers host compute accelerators such as GPUs).

For an e-infrastructure comprising HPC and cloud resources this means that different types of resource managers will co-exist. Additionally, in a federated e-infrastructure coordination between resource managers at different sites may be required to ensure timely availability of resources. This is an area ripe for further collaboration between the developers and providers of services based on HPC and Cloud resources and technologies. Further research and innovation is needed here, for example on how to orchestrate different types of resource managers to optimally use the available HPC or cloud infrastructure, while supporting co-allocation of different types of resources. Innovations in this area will also help in the seamless transition to HPC cloud bursting - where HPC centres can leverage cloud services on demand when existing HPC infrastructure has to deal with demand increases that cannot be met locally.

Another challenge in the context of federated e-infrastructures is centralised allocation of resources and accounting of used resources. Such a capability would on the one hand allow to allocate resources that are made available by different organisations to single projects such that, for instance, a project could have different types of computing resources at different locations. On the other hand, members of such projects would be able to monitor consumption of these resources at a central location without going through a process of collecting this information from different resource providers using different proprietary interfaces. Realising such capabilities would, in particular, require standardised and protected interfaces and mechanisms for distributing information on resource allocation and consumption. One example for such a centralised resource allocation and accounting service is the Fenix FURMS service[24].

TRUST, SECURITY AND DATA COMPLIANCE

Federated e-infrastructures integrate geographically distributed resources that are operated in different security domains such that data will be transferred over organisational boundaries. To use such e-infrastructures for workflows, for which confidentiality and security are particularly critical (e.g. to protect personal data or trade secrets), requires reconsideration of existing security measures, a more active security management targeting harmonised security levels at different sites as well as suitable mechanisms for establishing trust. While some aspects need to be addressed at a policy level others require suitable technical solutions.

As an example, we consider the need for an end-to-end secure set-up for workflow execution. Secure and trustworthy execution of HPC workloads is in its infancy while cloud system operations often rely on virtualisation to provide this, from node attestation to verifying workload attestation. The HPC community efficiency concerns prohibit such costly isolation procedures or the use of trusted execution environments (TEE)[25]. This will become more challenging within a federated infrastructure with multiple and heterogeneous sets of computing resources with different hardware support for TEEs. Future middleware solutions, in concert with the scheduling system, will need to address this challenge. Furthermore, trustable deployment of workload and other software components must be possible. These may be deployed using containers within an e-infrastructure supporting available security guidelines[26] and container encryption. Finally, the trust may be enhanced by putting mechanisms in place for auditing distributed workflow execution. These can be based on distributed solutions for generating audit trails using blockchain technologies (see, e.g.,[27]).

Another example concerns data compliance. While data is transferred over organisational and possibly state boundaries, the complexity of maintaining compliance with regulatory, organisational, or contractual data handling requirements becomes more complex. Data management solutions may support enforcing compliance, for instance by tagging data repositories and services depending on the met requirements (see, e.g.,[28]).

DISTRIBUTED INFRASTRUCTURE MONITORING

One key challenge for federated infrastructures is to monitor (and operate) them in such a way that in particular end-user services function seamlessly and with very high availability. Any malfunction or performance degradation needs to be identified in a central manner within a short period of time. Continuous testing and monitoring mechanisms must be organised in such a manner that difficult to understand error patterns, which are the result of a complex interplay between different services, can be

24. https://github.com/unity-idm/furms
25. For recent studies on trusted environments in the context of scientific computing workloads, see Ayaz Akram et al., "Performance Analysis of Scientific Computing Workloads on General Purpose TEEs", IPDPS, 2021 (DOI: 10.1109/IPDPS49936.2021.00115) and Sean Peisert, "Security Trustworthy Scientific Computing", Communications of the ACM, 2021 (DOI: 10.1145/3457191)
26. Murugiah Souppaya, John Morello, Karen Scarfone, "Application Container Security Guide", NIST Special Publication 800-190, 2017 (DOI: 10.6028/NIST.SP.800-190)
27. Abdullah Al-Mamun et al., "HPChain: An MPI-Based Blockchain Framework for Data Fidelity in High-Performance Computing Systems", Supercomputing, 2019 (https://sc19.supercomputing.org/proceedings/tech_poster/poster_files/rpost106s2-file3.pdf)
28. Martin Henze et al., "Practical Data Compliance for Cloud Storage", IEEE International Conference on Cloud Engineering (IC2E), 2017 (DOI: 10.1109/IC2E.2017.32)
29. As an example we refer here to the ESCAPE data lake testing infrastructures: Rizart Dona, Riccardo Di Maria (ESCAPE project), "The ESCAPE Data Lake: The machinery behind testing, monitoring and supporting a unified federated storage infrastructure of the exabyte-scale", CHEP 2021, EPJ Web of Conferences, 2021 (DOI: 10.1051/epjconf/202125102060).

identified[29]. Today, realising monitoring of distributed infrastructure components still requires a significant amount of customisation due to a lack of interfaces for collecting the necessary information with suitable access control.

Monitoring is expected to become more important as distributed infrastructures need to be able to support workflows with particular service quality demands, for instance in terms of service availability. A prominent example is numerical weather prediction, where particular calculations need to have been completed within pre-defined (short) periods of time as they otherwise would become obsolete. Today typically dedicated resources are provided for weather services, but in future also general-purpose resources provided through, e.g., the EuroHPC Joint Undertaking are expected to be used[30]. Monitoring mechanisms therefore need to support compliance with service-level agreements.

HARMONISATION OF SERVICE PROVISIONING

To ensure coherence within a federate e-infrastructure, harmonisation of the service provisioning is an important aspect. This can be supported by promoting the use of automatized service deployment using tools like Ansible[31], Puppet[32], or Terraform[33]. Sharing the relevant configuration files will not only improve harmonisation but also improve commoditization of the services.

Another area concerns the harmonisation of software environments within a distributed e-infrastructure. Different strategies started to be explored and implemented. The most popular approach is the use of containers, an OS-level virtualisation technique that facilitates deployment of customised environments within different host environments. It has become a standard technology in the cloud environment and is meanwhile commonly supported on HPC systems. To avoid customisation for individual users and user groups, tools like EasyBuild[34] and Spack[35] can be used as they allow sharing recipes for software installation that can be used on rather different systems. Finally, there are efforts for creating unified software environments that can be distributed to different systems of different types (see, for instance,[36]).

USABILITY BY END-USERS

Federation of different types of e-infrastructure services open many new opportunities for applications that require using a combination of services based on HPC, cloud computing and various storage resources. The lack of usability is one of the main drawbacks that both applications and tool developers are currently facing. Improving usability deserves significant efforts in order to leverage the opportunities offered by these e-infrastructures and to increase their effective use.

With regards to new research and innovation efforts in terms of usability, the following are some of the most promising approaches:

1. Establish DevOps for federated e-infrastructures: DevOps has been introduced as a term to refer to sets of common practices that bring together software development and infrastructure operations[37]. Such concepts have been applied in other areas to facilitate the process of development while ensuring the deployed code fits with the characteristics of infrastructure. In the area of AI-based services, the term AIOps was introduced to refer to engineers capable of efficiently and effectively building and operating online services that use AI and techniques[38]. Translating this philosophy to federated e-infrastructures including HPC will contribute to the usability and improvement of solutions tailored to these e-infrastructures, and therefore help to increase their popularity.
2. Services and resource selection support: Within distributed e-infrastructures the process of deciding on the choice of services and resources can become very complex when taking into account, for instance, cloud costs and benefits trade-offs as well as data movement policies to increase data locality.
3. Intelligent Automation for knowledge and service work: The term Intelligent Automation has been introduced to describe application of AI in ways that can learn, adapt and improve over time to automate tasks that as of today are performed by humans[39]. Examples are decision support and expert systems or recommendation agents. This approach could benefit from the aforementioned AIOps.
4. Creation of value-added cloud services integrated with HPC: A usable approach is characterised by providing value to the consumers. The virtualization of HPC resources and its deployment in a cloud-based approach, on the one hand, and the extension of data centre capabilities with value-added services, on the other hand, will provide usability, scalability, elasticity and will decrease the total cost of ownership (TCO).

30. European Commission, “Equipping Europe for world-class High-Performance Computing in the next decade”, SWD(2020) 179 final, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52020SC0179&rid=9.
31. https://www.ansible.com
32. https://www.puppet.com
33. https://www.terraform.io
34. https://docs.easybuild.io/en/latest/
35. https://spack.readthedocs.io/en/latest/
36. Maxime Boissonneault et al., “Providing a Unified Software Environment for Canada’s National Advanced Computing Centers”, PEARC’19, 2019 (DOI 10.1145/3332186.3332210)
37. Frédéric Desbiens, “EdgeOps for IoT: A Vision For Edge”, Connected World, 2021 (https://connectedworld.com/edgeops-for-iot-a-vision-for-edge/)
38. Y. Dang, Q. Lin and P. Huang, «AIOps: Real-World Challenges and Research Innovations”, IEEE/ACM 41st International Conference on Software Engineering, 2019 (DOI: 10.1109/ICSE-Companion.2019.00023)
39. C. Coombs, D. Hislop, S. K. Taneva and S. Barnard, “The strategic impacts of Intelligent Automation for knowledge and service work: An interdisciplinary review”, The Journal of Strategic Information Systems, 2020 (DOI: 10.1016/j.jsis.2020.101600)

RECOMMENDATIONS FOR R&I EFFORTS

On the basis of the analysis of the different topical areas the following recommendations for further research and innovation efforts by ETP4HPC members, in particular in the context of R&I actions supported by the EuroHPC Joint Undertaking:

1. The integration of HPC systems and HPC-based services within the AAI of a federated e-infrastructure should become a commodity while at the same time different and sufficiently strong authentication mechanisms should be supported, e.g. through multi-factor authentication.
2. The integration of different types of services based on secured HPC and data resources and other services that can be made more openly available, e.g. web-based portal services, needs to be improved. Specific examples are web-based APIs to resource managers, support of data management involving HPC systems and data platforms.
3. Improve support of centralised resource allocation and resource consumption monitoring within a federated e-infrastructure with different types of resources provided by different organisations.
4. Seek collaboration between HPC and cloud communities for research and innovation on orchestration of different types of resource management systems within a federated e-infrastructure.
5. Improve support for end-to-end secure setup for workflow execution, mechanisms for trustable deployment of workload and other software components as well as solutions that help to enforce data compliance.
6. Establish standardised interfaces for collecting monitoring information with suitable access control within a distributed environment.
7. Promote usability by end-users by establishing DevOps for federated e-infrastructures, developing services and resource selection support tools and exploring Intelligent Automation for knowledge and service work.

SUMMARY AND CONCLUSIONS

In this white paper, we discussed general trends towards integration of HPC systems within federated e-infrastructures and presented two initiatives that are working on establishing such e-infrastructures. While such e-infrastructures have been established and are in use, there are still many areas where research and innovation efforts are needed for commoditization, improved capabilities and usability. We explored some of these areas in more detail and made a number of specific recommendations for research and innovation efforts in the context of ETP4HPC.

GLOSSARY AND LIST OF ACRONYMS

AAI	Authentication and authorization infrastructure
AI	Artificial intelligence
API	Application programming interface
CSP	Cloud service providers
DevOps	A practice that combines software development (Dev) and IT operations (Ops)
EOSC	European Open Science Cloud (https://eosc-portal.eu/)
ESFRI	European Strategy Forum on Research Infrastructures
HPC	High-performance computing
IoT	Internet-of-things
Level of assurance	Degree of confidence in the claimed identity of a person
ML	Machine learning
OIDC	OpenID Connect (https://openid.net/connect/)
OS	Operating System
TEE	Trusted execution environment

6.1.5

Heterogeneous High Performance Computing

Heterogeneity is here to stay: Challenges and Opportunities in HPC.

INTRODUCTION

Modern HPC systems are becoming increasingly heterogeneous, affecting all components of HPC systems, from the processing units, through memory hierarchies and network components to storage systems. This trend is on the one hand due to the need to build larger, yet more energy efficient systems, and on the other hand it is caused by the need to optimise (parts of the) systems for certain workloads. In fact, it is not only the systems themselves that are becoming more heterogeneous, but also scientific and industrial applications are increasingly combining different technologies into complex workflows, including simulation, data analytics, visualisation, and artificial intelligence/machine learning. Different steps in these workflows call for different hardware and thus today's HPC systems are often composed of different modules optimised to suit certain stages in these workflows.

While the trend towards heterogeneity is certainly helpful in many aspects, it makes the task of programming these systems and using them efficiently much more complicated. Often, a combination of different programming models is required and selecting suitable technologies for certain tasks or even parts of an algorithm is difficult. Novel methods might be needed for heterogeneous components or be only facilitated by them. And this trend is continuing, with new technologies around the corner that will further increase heterogeneity, e.g. neuromorphic or quantum accelerators, in-memory-computing, and other non-von-Neumann approaches.

In this paper, we present an overview of the different levels of heterogeneity we find in HPC technologies and provide recommendations for research directions to help deal with the challenges they pose. We also point out opportunities that particularly applications can profit from by exploiting these technologies. Research efforts will be needed over the full spectrum, from system architecture, compilers and programming models/languages, to runtime systems, algorithms and novel mathematical approaches.

KEY INSIGHTS

- **Heterogeneity is here to stay, and the upcoming disruptive technologies (e.g. neuromorphic, quantum, processing in memory) will only increase it.**
- **Dynamic orchestration and management of heterogeneous components are key for effective resource utilisation.**
- **Applications need appropriate tools (standardised programming models; smart compilers, runtime and workflow systems; debugging and performance tools) to master the complexity of heterogeneous systems.**
- **Novel heterogeneous systems provide opportunities for novel workflows, novel mathematical formulations and new application features.**
- **Heterogeneity cannot be tackled in an independent manner but coordinated efforts at different levels need to be combined (e.g. compilers and tools and runtime systems)**

KEY RECOMMENDATIONS

- **The challenge of heterogeneity must be faced as a community effort, sharing best practises to combat complexity.**
- **Heterogeneity has to be tackled at all levels of the stack: from the system architecture, (choosing an appropriate distribution of the resources) passing by the middleware and software stack (to manage and orchestrate the resources), up to the application codes (adapting them to use the various resources).**
- **Integration projects hardening and combining results covering all levels of the HPC ecosystem are needed**
- **Interoperability and exchangeability of components (hardware and software) should be improved and become part of the overall design**

Today's HPC systems exhibit heterogeneity everywhere - from the compute and storage subsystems, through integrated systems, to systems software and even application software. In this Chapter, we provide an overview of the current situation, point out some potential future directions and highlight future research directions.

PROCESSING TECHNOLOGIES

Multiple different processing technologies are today combined in HPC systems in order to achieve the required performance and functionality at an affordable power envelope. The most frequently used components are general purpose processors (CPUs) and graphic cards (GPUs), but others are also used (e.g. field programmable gate arrays (FPGAs), domain-specific accelerators) and, on experimental basis, even more exotic ones like neuromorphic and quantum devices. This hardware diversity comes at the price of a more complex software and programming environment, and the need to adapt the applications to exploit the full computational power of heterogeneous HPC systems.

CPUS

Arguably, **CPUs** are currently the best devices in terms of programmability, with a long history of usage and very wide and stable software support that facilitates running applications across different hardware generations and even with different CPU architectures porting applications is typically quite straight-forward (e.g. from x86 to ARM CPUs). However, modern CPUs have become much more complex than they used to be in the past. The end of Dennard's law[1] forced CPUs to become multi-core, for which cache coherency protocols and parallel programming techniques are required. Following the RISC architecture principle, the CPU pipelines are now deeper and apply multiple optimisation strategies for which specific hardware elements inside the CPU have been added. Examples include arithmetic-logic units to execute different kinds of operations, memory prefetchers, and registers for vector operations with progressively larger and even variable vector lengths. All these added capabilities bring a higher peak performance on the CPU but require specific optimisations such as data alignment and vectorisation in order to actually benefit from them. Some of these optimisations can be achieved by compilers and runtime systems, but often applications also need to be modified to best profit from them. The CPUs used in HPC today have become very heterogeneous themselves and this trend is likely to continue. A particular type of CPUs are **many-core devices**, which were designed as CPUs containing a large number of relatively weak cores, originally conceived as accelerators, i.e. PCIe-attached cards that depend on a stronger host-CPU to boot and communicate through the network. However, the distinction between multi-core CPUs and many-core devices has almost diluted, since the latter became more and more autonomous and the former today can contain over 100 cores; and with the end of Intel's Xeon PHI[2] products, there are currently no wide-spread many-core devices on the market.

GRAPHICS PROCESSING UNITS (GPUS)

Graphics Processing Units (GPUs) were originally designed for the gaming and film industries to speed-up graphics operations, but are used today also as computing devices by arithmetic-intensive HPC and Artificial Intelligence (AI) applications. High-end GPUs contain a very large number of arithmetic-logic units supporting various precisions (double, single, half precision) and specific tensor cores used to perform very efficiently vector and matrix operations widely used in machine- and deep learning algorithms. Furthermore, they contain high-bandwidth memory (HBM) technologies that are very interesting for memory-bound applications. To use GPUs, accelerator programming models such as OpenACC[3] and OpenCL[4] are required, some of them vendor specific such as CUDA or HIP, which hampers application portability and particularly performance portability across platforms. There are ongoing efforts to include vendor-independent GPU support in other programming environments, e.g. OpenMP[5] or even C++, but so far, the best performance is still achieved using the vendor specific environments. Nevertheless, their many execution units and the high memory bandwidth make GPUs very computationally powerful and energy efficient devices, reaching peak performances an order of magnitude higher than CPUs.

FIELD PROGRAMMABLE GATE ARRAYS (FPGAS)

In Field Programmable Gate Arrays (FPGAs), the logic gates can be set up to execute the specific operations of the problem to be solved, which allows designing a computing device specifically for each use case so that the maximum performance and energy efficiency can be achieved. This configuration of the system before the user code is launched induces some overhead, but it becomes negligible if the total application runtime (eventually over various executions) is large enough. The caveat is that the environment for the pre-configuration and FPGA programming using low-level approaches like VHDL or high-level synthesis (HLS) is still not familiar to application developers. OpenCL is moderately higher level, but the kernels still need to be managed from the host side. High-level programming environments are possible for a more comfortable approach[6], and the FPGA can be used to accelerate the management of kernels in addition to the kernels themselves but achieving maximum performance still requires some low-level programming and tuning between different hardware generations, specific device configurations, and technologies. This makes achieving performance portability among

1. R. H. Dennard, F. H. Gaensslen, H.-N. Yu, V. L. Rideout, E. Bassous and A. R. LeBlanc, "Design of ion-implanted MOSFET's with very small physical dimensions," IEEE Journal of Solid-State Circuits, vol. 9, no. 5, pp. 256-268, October 1974.
2. "Intel discontinues Xeon Phi 7200 series Knights Landing coprocessor cards," [Online]. Available: https://www.anandtech.com/show/11769/intel-discontinues-xeon-phi-7200-series-knights-landing-coprocessor-cards.
3. https://www.openacc.org/
4. https://www.khronos.org/registry/cl/specs/opencl-1.1.pdf
5. https://www.openmp.org/
6. https://pm.bsc.es/ompss-at-fpga

different devices and generations of devices even more challenging than for different CPU/GPU architectures and generations.

DISRUPTIVE APPROACHES

Disruptive approaches such as neuromorphic and quantum technologies are also being explored. For example, quantum technologies have been shown to solve selected optimisation problems much faster and more efficiently than traditional von-Neumann systems. Similarly, neuromorphic computers are well suited for pattern recognition or neural networks. Until now, these technologies have proven their capabilities on individual use cases tuned to exploit them. Acceptance by the wider community still is hampered by a yet rather immature and device-specific software and programming environments, but progress is being done in all fronts to widen their use and applicability.

This diversity of processing technologies is posing new challenges to system integrators and application programmers alike. One challenge is memory management, where most approaches currently do not share the same memory and data placement needs to be managed explicitly by the application. Recently, approaches for cache coherent GPU/CPU memory systems have been announced, their efficiency needs however still to be proven. Not all accelerators or even more so disruptive approaches will allow the implementation of such coherency, though. In addition, application programmers have to deal with different, sometimes vendor-specific, programming environments (like OpenMP, MPI, Cuda, HiP, Sycl, OpenACC, OpenCL, etc.) to be able to exploit the different components. In addition, application developers need to decide, which platform or platform mix is best suited for their applications and part of an application should be executed on what platform. In turn, system integrators need to decide what mix of technologies is best suited for the expected application workload. To overcome these difficulties, more research is needed in standardised programming environments (with associated intelligent compilers and runtime systems) as well as modelling approaches to predict the potential performance of applications on certain technologies.

MEMORY TECHNOLOGIES

While the compute performance of CPUs, GPUs, etc. has traditionally captured most of the attention in HPC, e.g. in the TOP500 ranking, the performance of many HPC applications is actually constrained by the memory system, either due to latency/bandwidth (e.g. HPCG) or capacity (genomics, fluid dynamics). For more than five decades, HPC memory systems have followed the same basic design: a von Neumann architecture where the CPU controls a passive memory hierarchy. Nowadays, most HPC systems use a cluster architecture, in which each node has one or two sockets and its own memory hierarchy comprising L1, L2 and L3 caches and DRAM. However, the memory hierarchy is getting deeper with the introduction of further layers, particularly of High Bandwidth Memory (HBM)[7] and non-volatile memory DIMMs (NVDIMMs)[8].

Several innovations in the memory subsystem organisation and underlying memory technologies are already in production (e.g. HBM, Non-Volatile Memory (NVM), disaggregated memories) or under active industrial and academic development (e.g. Processing in Memory (PIM)[9,10,11]). These systems provide opportunities to improve performance and reduce power consumption, but effectively and productively exploiting them presents many challenges:

EXPLICITLY MANAGED MEMORY

Complex memory systems are generally architected in hardware not as a cache hierarchy, but on an equal ranking, e.g. Optane[12] is connected on a normal DDR interface[13]. The data placement and migration among multiple types of memories can be done in various places, for example as a transparent cache by the memory controller (DRAM is an inclusive cache for Optane in “memory mode”), transparently by the OS (at a page granularity), or by the runtime system or specific support library (at a page or object granularity, or even by the compiler. Alternatively, they may be exposed to the application either by exposing the specific memory types and topology or using a generalised interface in which memory is requested in terms of capabilities (high bandwidth, high capacity, etc.). Some APIs already exist: memkind[14], SICM (Simplified Interface to Complex Memory)[15] and some European projects are pioneering work in this direction, e.g. DEEP-SEA[16].

7. “High Bandwidth Memory (HBM) DRAM, White Paper,” JEDEC Solid State Technology Association, 2013.
8. “JEDEC Publishes DDR4 NVDIMM-P Bus Protocol Standard,” February 2021. [Online]. Available: https://www.jedec.org/news/pressreleases/jedec-publishes-ddr4-nvdimm-p-bus-protocol-standard.
9. K. Wang, K. Angstadt, C. Bo, N. Brunelle, E. Sadredini, T. Tracy, J. Wadden, M. Stan and K. Skadron, “An Overview of Micron’s Automata Processor,” in Proceedings of the Eleventh IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis (CODES), 2016.
10. Y.-C. Kwon, S. H. Lee, J. Lee, S.-H. Kwon, J. M. Ryu, J.-P. Son, O. Seongil, H.-S. Yu, H. Lee, S. Y. Kim, Y. Cho, J. G. Kim, J. Choi, H.-S. Shin, J. Kim, B. Phuah, H. Kim, M. J. Song, A. Choi, D. Kim, S. Kim, E.-B. Kim, D. Wang, S. Kang, Y. Ro, S. Seo, J. Song, J. Youn, K. Sohn and N. S. Kim, “A 20nm 6GB Function-In-Memory DRAM, Based on HBM2 with a 1.2TFLOPS Programmable Computing Unit Using Bank-Level Parallelism, for Machine Learning Applications,” in Proceedings of IEEE International Solid-State Circuits Conference (ISSCC), 2021.
11. F. Devaux, “The True Processing in Memory Accelerator,” in IEEE Hot Chips Symposium (HCS), 2019.
12. “Intel Optane Memory - Responsive Memory, Accelerated Performance,” [Online]. Available: https://www.intel.com/content/www/us/en/products/details/memory-storage/optane-memory.html .
13. “Intel at last announces Optane memory: DDR4 that never forgets.,” May 2018. [Online]. Available: https://arstechnica.com/gadgets/2018/05/intel-finally-announces-ddr4-memory-made-from-persistent-3d-xpoint .
14. http://memkind.github.io/memkind/
15. https://github.com/lanl/SICM
16. DEEP-SEA Project. Horizon 2020 grant agreement 955606, https://www.deep-projects.eu/

PROCESSING IN-MEMORY

Decades after being initially explored in the 1970s, Processing in Memory (PIM) is currently experiencing a renaissance. PIM moves part of the computation to the memory devices, thereby addressing the mismatch between the von Neumann architecture and the requirements of important data-centric applications[17]. The interest in PIM has grown dramatically over the last years. The recent white paper from ETP4HPC[18] advises that wide acceptance of PIM in high-performance computing depends on our ability to create an ecosystem in which a number of PIM approaches can be designed and evaluated, leading to the selection of potential winners and adoption by system architects and end users.

Many diverse approaches are under development for PIM, differing in where the computation takes place (inside the memory array or periphery, or near the memory but outside it)[19] and in the type of operations supported, as well as in terms of cost, maturity, etc.. Any modifications required to existing codes should be as small and local as possible, and should be done in a performance-portable (if at all possible) and vendor independent way controlled in the long term by open standards. It is currently not clear what will be the eventual programming model to describe the code to run on a PIM system and how to control data placement. Overall, we can conclude that only coordinated innovations and co-design across the whole stack can make PIM a reality in production.

As with the heterogeneous computing technologies discussed above, a key challenge is to decide which mix of technologies works best for what applications and what data should be placed where and when. This requires more research to decide at what level(s) of the stack to manage data placement/migration, etc., on what basis (heuristics, historical data, sampling during the execution, user annotations etc.), subject to which metrics, and using which measurement/analysis tools—and there is still room to try various approaches. The best choice of memory for performance is not always obvious, e.g. if the memory bandwidth is low, then MCDRAM has higher latency than DDR (the bandwidth–latency curves cross), and thus worse performance. Standardisation in APIs and programming models will be needed eventually, but not too soon as the underlying technology is still evolving fast.

STORAGE TECHNOLOGIES

The storage system is typically attached to the supercomputer and made accessible via a file-system. The different storage technologies are organised in a hierarchical/pyramidal approach with the fastest (but smaller and more expensive) devices closest to the computer and the largest capacity (cheaper but slower) furthest away. This concept extends the principle of the memory hierarchy into the storage region, with non-volatile memories (which can be configured as memory or storage) sitting somewhere in the middle of both areas.

The heterogeneous storage hierarchy-"tiers" consist of byte addressable NVRAM at the very top followed by block based Flash storage (exposed through Solid State Drives), High performance Hard Disk subsystem (eg: Serial Attached SCSI, SAS) and slower archival grade drives (eg: Shingled Magnetic Recording Drives). NVRAM components are typically present in Compute nodes and they are either exposed as byte addressable DIMMs or block addressable NVMe devices. NVRAM and Flash tiers typically could act as "Burst Buffers" hiding the latency of lower performance tiers below them. All tiers are exposed to parallel file system storage services or object based storage services. Applications can do I/O directly on the tiers or through files and objects mapped to the NVRAM tiers and memory. Parallel file systems/ object storage could also work closely with support from other software infrastructures such as Hierarchical Storage Managers (HSMs), which can help to move data to the right performance tier and make them available to workflows at the right time with appropriate performance.

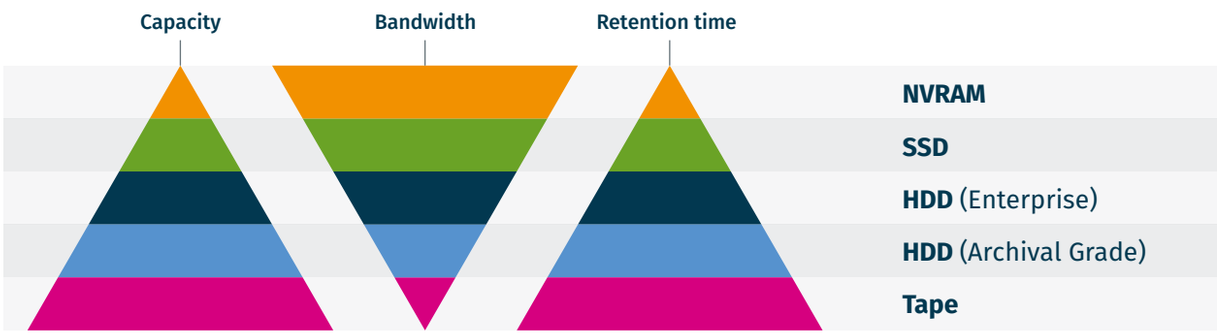

Figure 10: Storage Hierarchy

17. S. Ghose, A. Boroumand, J. S. Kim, J. Gomez-Luna and O. Mutlu, "Processing-in-memory: A workload-driven perspective," IBM Journal of Research and Development, vol. 63, no. 6, pp. 3:1-3:19, November-December 2019.
18. P. Radojković, P. Carpenter, P. Esmaili-Dokht, R. Cimadomo, H.-P. Charles, A. Sebastian and P. Amato, "Processing in Memory: The Tipping Point," ETP4HPC, 2021.
19. P. Siegl, R. Buchty and M. Berekovic, "Data-centric computing frontiers: A survey on processing-in-memory," in Proceedings of the Second International Symposium on Memory Systems, 2016.

One of the concepts explored for example by the SAGE/Sage2[20] projects in the EU is "Global Memory Abstraction" where-in data objects are mapped from lower tiers to higher tiers. The objects could be mapped to NVRAM or memory - which enables applications to work with objects in memory at very high speeds and low latencies. These objects can be "drained" back to lower tiers at a later time. The concept also raises the possibilities to consolidate NVRAM pools distributed across multiple compute nodes. SAGE and Sage2 projects have explored hierarchical storage systems with all the known tiers - NVRAM, SSD, HDD & Tape - all managed as Objects by Object storage infrastructure software.

It is also to be borne in mind that tape has always provided a very low cost persistent storage resource, primarily used for long term archival. There are exabytes of data in tape across various scientific communities worldwide. However, frequently accessing data from tape multiple times, when needed by running applications, is always a challenge. Bringing tape into the fold of hierarchical tiers for running workflows is also actively being explored. In that there are software interfaces to data in tape from object stores hence providing the ability to actively work with Tape data as objects like in any higher level tier. This is explored in the IO-SEA EuroHPC project[21].

Furthermore, heterogeneous computing resources (GPUs, FPGAs, etc) very close to data can act to provide "in-storage computing"[22] capability. Whilst standard processing cores can be leveraged for in-storage computing, GPUs, FPGAs, etc. raise interesting new possibilities in storage system design. All this will help to reduce traffic from storage systems to compute nodes helping to eventually improve the time to solution. However, it needs to be noted that inclusion of some of these technologies as part of storage systems is also driven by economic considerations. In fact the role of the DPU (Data Processing Unit)[23], pushed by some of the vendors, from the perspective of in-storage-computing is also something that needs to be explored.

In summary, heterogeneity in storage technologies exposes the same challenges as the heterogeneous memory technologies: where to place data and when to migrate data through the different tiers. More research is needed on tools that help with these decisions or are able to automatically take these decisions.

FULL SYSTEM ARCHITECTURES

High compute and data management performance can be achieved in a cost and energy efficient manner by wisely combining the diverse computing, memory, and storage devices mentioned above. The aim of a system architect is to choose complementary technologies, which potentiate each other's strengths and compensate for each other's weaknesses under a given cost-target. The system-level architecture is then defined by the choice of technologies and how they are combined and connected with each other.

Cluster computers traditionally interconnected general-purpose CPUs via a high-speed network, and attached to an external disk-based storage. Accelerators, and in particular GPUs, started to play a role around 2010, due to the increasing power-needs of CPUs and the relatively high energy efficiency of accelerators. The typical integration consists of building one cluster node as a CPU with one or more PCIe-attached accelerators, and then interconnecting these "heterogeneous nodes" with each other via a high-speed network. Typically, the communication between the accelerators (within and between nodes) has to be managed by the host CPU, creating a communication bottleneck at the CPU-GPU connection. However, recent GPU technologies support (vendor specific) high performance interconnects to directly communicate between the GPUs in the node. The "GPU-islands" created in this way deliver a huge computational power per node and move the communication bottleneck to the inter-node interface, although new developments like GPU-direct communications might reduce the problem to some extent. An approach for multi-node FPGA-to-FPGA communication using shared memory, and without involving the CPU, has been developed by the EuroEXA project[24].

A particular approach to heterogeneous system integration targeting diverse application portfolios is the so-called Modular Supercomputer Architecture (MSA)[25] developed in the DEEP projects[26]. It combines several clusters of potentially large size (called "modules"), each one tuned to best match the needs of a certain class of applications. For instance, a homogeneous CPU-based cluster that delivers high single-thread performance for low/medium scalable applications can be attached to an accelerator-based system that delivers energy efficient performance for highly scalable codes, and to a third CPU+accelerator cluster with large volumes of high-bandwidth memory for data-analytics codes. The architecture allows even the integration of more disruptive technologies such as neuromorphic or quantum computers. A federated network connects the module-specific interconnects, while an optimised resource manager enables assembling arbitrary combinations of these resources according to the application workload requirements. The goal of the MSA-approach is to enable an efficient orchestration of heterogeneous resources at the service of very diverse application profiles, so that each code can run on a near-optimal combination of resources and achieve excellent performance, increasing throughput and efficiency of use for the whole system.

20. S. Narasimhamurthy, N. Danilov, S. Wu, G. Umanesan, S. Markidis, S. Rivas-Gomez, I. B. Peng and S. De Witt, "Sage: percipient storage for exascale data centric computing," Parallel Computing, no. 81, 2019.
21. https://iosea-project.eu
22. Z. Ruan, T. He and J. Cong, "Designing In-Storage Computing System for Emerging High-Performance Drive," in Proceedings of the 2019 USENIX Annual Technical Conference, Renton, WA, USA, 2019.
23. "What's a DPU data processing unit," [Online]. Available: https://blogs.nvidia.com/blog/2020/05/20/whats-a-dpu-data-processing-unit/
24. EuroEXA project. Horizon 2020 grant agreement number 754337, https://euroexa.eu/
25. E. Suarez and T. Lippert, "Modular Supercomputing Architecture: from idea to production," in Contemporary High Performance Computing: from Petascale toward Exascale, vol. 3, Ed. Jeffrey S. Vetter, CRC Press, 2019, pp. 223-251.
26. DEEP projects. Horizon 2020 grant agreements 287530, 610476, 754303 and 955606., https://www.deep-projects.eu/

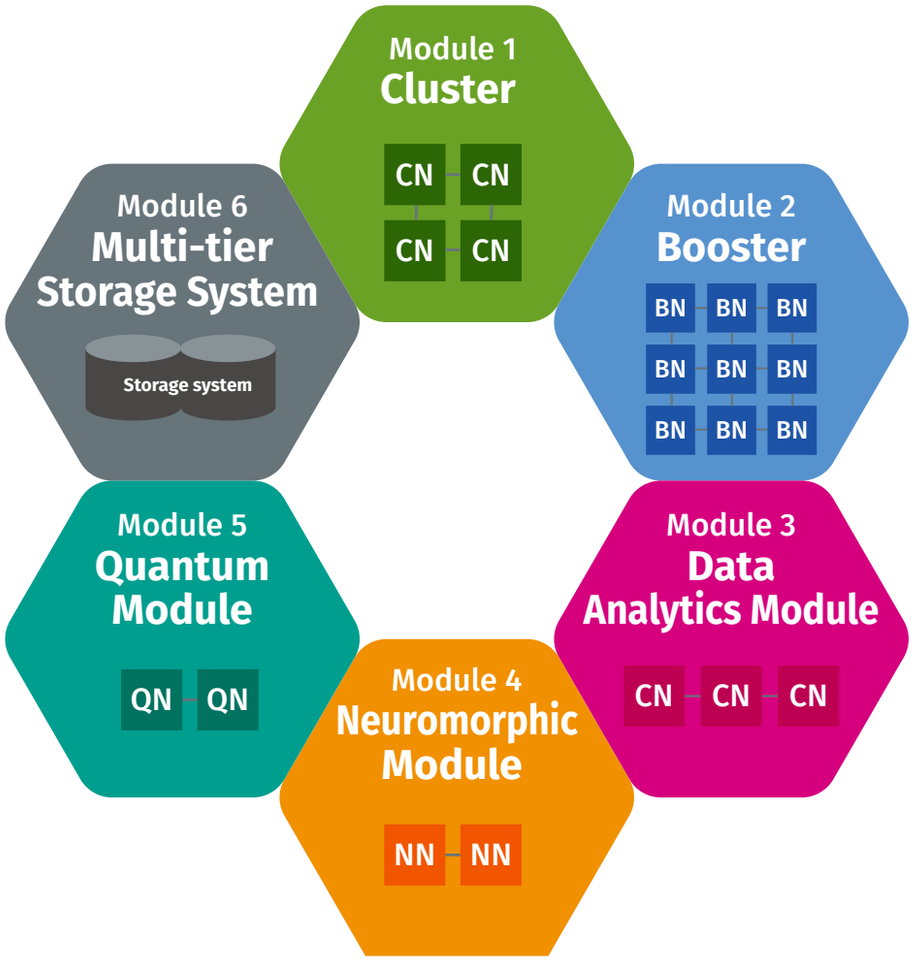

Figure 11: Modular Computing Systems

Some novel architectures attempt to turn around the traditional approach that is built around the computing elements by putting a pool of globally accessible memory at the centre of the system. Such "memory-centric" architectures[27,28,29] promise a better energy efficiency through minimisation of data movement, but (when built at large scale) require advanced resilience mechanisms to deal with typically high failure rates on memory-hardware.

Continuing on this path of disaggregation, the network fabric can take a more central stage in not only facilitating data exchange, but also computation during data transport. While this is not a new idea -- some network technologies have been integrating in-fabric support of atomic RDMA operations and (simple) MPI collective operations for a decade -- the trend continues to gain traction, with active compute elements in switches and programmable NICs becoming a feature in multiple vendor's upcoming products[30,31]. The challenge is accessing such features, e.g. making them available both in existing programming paradigms like MPI and through domain-specific toolkits (e.g., deep learning frameworks, workflow couplers). They also need to become accessible to high-end users ready to invest into offloading to the network, much like offloading to GPU accelerators has become commonplace.

27. J. Schmidt, "Network Attached Memory, Chapter 4 of the PhD Thesis: «Accelerating Checkpoint/Restart Application Performance in Large-Scale Systems with Network Attached Memory», Ruprecht-Karls University Heidelberg - Fakultaet fuer Mathematik und Informatik," http://archiv.ub.uni-heidelberg.de/volltextserver/23800/1/dissertation_juri_schmidt_publish.pdf
28. P. Faraboschi, K. Keeton, T. Marsland and D. Milojicic, "Beyond Processor-centric Operating Systems," in 15th Workshop on Hot Topics in Operating Systems (HotOS XV), 2015.
29. Rigo et al., "Paving the way towards a highly energy-efficient and highly integrated compute node for the Exascale revolution: the ExaNoDe approach," in Euromicro Symposium on Digital System Design, DSD 2017, 2017.
30. e.g., https://www.arubanetworks.com/pensando-announcement/ and https://www.mellanox.com/products/BlueField-SmartNIC-Ethernet
31. S. Schweitzer, "Panel: SmartNIC or DPU, Who Wins?," in 2020 IEEE Symposium on High-Performance Interconnects (HOTI), see https://technologyevangelist.co/2020/08/25/smartnics-vs-dpus/, 2020.

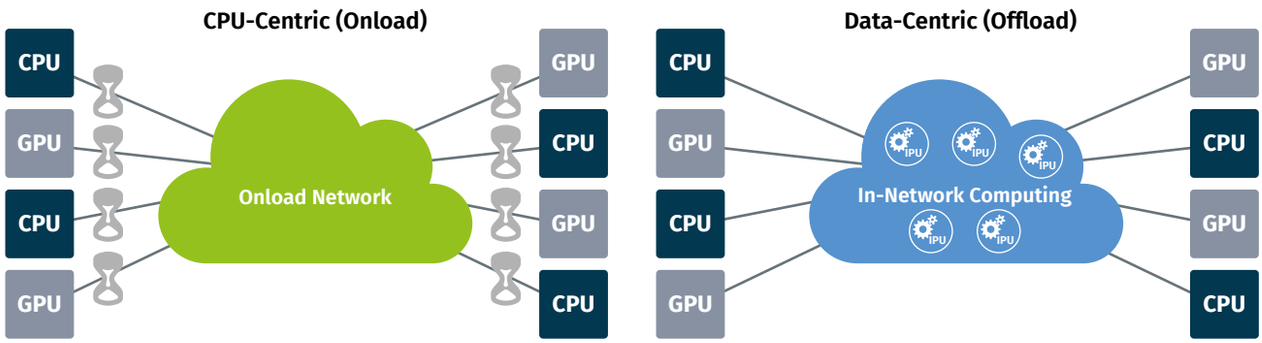

Figure 12: - CPU-centric vs. Data-centric Computing[32]

Not only resulting from heterogeneity in the HPC system, but clearly complicated by it, is the question of how to ensure increasing security requirements in such environments. While initially driven by special-purpose systems and then the need to separate potentially adversary customer's use of the same resources in a cloud computing environment, many -- but still not all -- components used in an HPC system can (or should be able to) provide some level of trusted computing capability: CPU and memory protection standards exist, some devices incorporate appropriate hardware roots-of-trust, and standards like SPIFFE are able to attest trust hierarchically[33]. However, in an HPC context there is still a lot of work left to do: UNIX style user or project separation will not be sufficient to segregate tenants on an exascale system shared by many different user communities; scheduling systems need to be aware and tied into the security design; high performance networks need to isolate traffic of jobs to both ensure guaranteed performance characteristics, but also shield users from crosstalk of other jobs, as such behaviour could otherwise permit DoS or side channel attacks. Finally, there is a trustworthy computing view expressed by HPC system operators: contrary to many cloud providers, HPC centres do care about the kind of workload run by their users, and they may want (or need) to prohibit usage of their resources for certain purposes. The question of how to attest user workloads as compliant with system operators' requirements is an entirely separate, but technologically intertwined, area of research.

Combining different heterogeneous components as discussed above into a full HPC system results in combinatorial effects in their complexity. It is a huge challenge to design systems such that they can be used efficiently by the expected workloads, particularly, when the workload is very heterogeneous. Modular systems as discussed above can help, deciding according to the user portfolio how much weight a particular module should get, and what connectivity is required within and between modules. To deal with these challenges, integrated projects that cover all levels of the HPC ecosystem are needed. Also, interoperability and exchangeability of components, both hardware and software, should be easier to give system designers and users, alike, more flexibility.

APPLICATIONS AND SYSTEM SOFTWARE

On top of the heterogeneous hardware described above various system and application software stacks are being executed. These stacks typically have to be aware of the underlying hardware heterogeneity to various extents in order to optimally exploit the available hardware capabilities.

SCHEDULING AND RESOURCE MANAGEMENT

A core component of the system software is the scheduling and resource management system. Traditional usage of HPC systems centres around a batch scheduling system such as Slurm or PBS, which exposes the resources of the system either by partitioning them so that groups of mostly identical nodes are visible to the user, or by tagging nodes with feature labels, so that users can specify the desired kinds of features by selectors. These are feasible models if variability between nodes is small, if all resources to be allocated are encapsulated in the same nodes, and job requirements are homogenous across all components: Users request resources matching their job's needs, and the scheduler can make a decision about how to allocate resources.

Such a model is well suited to a mostly uniform set of resources, and jobs that have a constant resource requirement throughout their lifetime and across distributed parts (say, MPI ranks). It is not well suited to heterogenous jobs, or complicated workflows, where a resource matching over time may need to be performed. Solving this problem by workflow languages is not a sustainable approach: there are literally hundreds of workflow management tools, often used only by a small community, and many do not

32. M. Malms, M. Ostasz, M. Gilliot, P. Bernier-Bruna, L. Cargemel, E. Suarez, H. Cornelius, M. Duranton, B. Koren, P. Rosse-Laurent, M. S. Pérez-Hernández, M. Marazakis, G. Lonsdale, P. Carpenter, G. Antoniu , S. Narasimhamurthy, A. Brinkman, D. Pleiter, A. Tate, A. Wierse, J. Krueger, H.-C. Hoppe and E. Laure, "ETP4HPC's Strategic Research Agenda for High-Performance Computing in Europe 4," Zenodo, 2020.
33. https://spiffe.io/

take HPC execution realms into account. Mapping scheduling decisions to one or more particular target architectures and HPC execution realms, taking into account batch schedulers with their own scheduling logic, makes holistic scheduling a far-off target. More research in these directions as well as fewer (and more broadly applicable) workflow languages will be needed.

A further complication arises if dynamically created or provisioned resources are considered. It is becoming a common requirement to allocate temporary storage space to a job (e.g., in NVRAM, as a dynamically created file system[34,35], or by creating an object store namespace[36]), and possibly also perform stage-in and/or stage-out of data, in preparation for efficient access during the job runtime, but also possibly in order to access encrypted data temporarily in near-compute storage. While prologue and epilogue tooling can be used to perform these operations, schedulers are not advanced enough to perform this task sufficiently: Data preparation likely needs far fewer resources than a large-scale compute job, but may itself be well suited to be a distributed job. Deciding (and the accounting of) the effort spent is not a user or admin task, but should be part of the scheduling decisions, taking into account detailed knowledge about system resources, current utilisation, expected performance gains, and co-scheduling options with other jobs.

Finally, applications -- whether at user level, at middleware level, or in standard HPC libraries -- are missing system introspection infrastructure. While tools like hwloc[37] and netloc[38] often are available to provide an overview of components, there is a definite lack of higher-level portable resource description tooling that also properly reflects the restrictions placed on availability and performance characteristics by the scheduling decisions.

COMPILERS AND RUNTIME SYSTEMS

Another important set of system software components are compilers and runtime systems. Despite new generations of HPC hardware becoming available at steady rates over the last decades, application developers (and so, users) could rely on the four cornerstones of compiler (Fortran, C, C++), MPI, PFS and scientific libraries (like fftw, BLAS, LAPACK) being available and tuned to the current target system. This has started to change with GPUs bringing their own programming paradigms such as CUDA[39] and HIP[40] into the mix, which cannot always be abstracted entirely by models like OpenACC, OpenMP, and OpenCL[41]. The situation worsens with novel accelerators entering the landscape, from programmable hardware like FPGAs to custom-built devices that serve best as a backend to a particular programming paradigm in a specialised application domain, such as TensorFlow. With heterogeneity becoming mainstream, relying on the compiler to make the final (and static) decisions about code transformations, pattern detection and substitution by library functions, as well as data access optimisation becomes less and less realistic

This indicates potential for the rise of additional middlewares, that is, runtimes that expose a higher level of abstraction over the actual machine architecture. It can be argued that POSIX and MPI, and standard ScaLAPACK functions are effectively such middlewares, but data centric abstractions like Kokkos[42], Raja[43], ADIOS2[44], XIOS[45], and Maestro[46], task based execution runtimes[47], object stores[48], and deep learning frameworks are becoming equally important. The interplay (and compatibility or interoperability) of all of these is posing a major software engineering and maintenance challenge. Standardisation of the system software as well as middleware layers will be an important aspect for ensuring uptake and maintainability.

Finally, software containers like Singularity[49], Docker[50], Podman[51], and Sarus[52] cannot be ignored in this context. While they solve the dependency issue, and relieve the end user from complicated software installation burden, they create new questions for efficient system operations and use: Avoiding useless rebuilds, but rebuilding containers to best match the target platform, defining, maintaining, and mapping compatible APIs between host system and containers transparently to the user to ensure full

34. F. Tessier, M. Martinasso, M. Chesi, M. Klein and M. Gila, “Dynamic Provisioning of Storage Resources: A Case Study with Burst Buffers,” in 2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), 2020.
35. “BeeOND™: BeeGFS On Demand,” [Online]. Available: https://www.beegfs.io/wiki/BeeOND
36. https://docs.daos.io/
37. https://www.open-mpi.org/projects/hwloc/
38. https://www.open-mpi.org/projects/netloc/
39. https://developer.nvidia.com/cuda-zone
40. https://rocmdocs.amd.com/en/latest/Programming_Guides/HIP-GUIDE.html
41. S. Memeti, L. Li, S. Pllana, J. Kołodziej and C. Kessler, “Benchmarking OpenCL, OpenACC, OpenMP, and CUDA: Programming Productivity, Performance, and Energy Consumption,” in ARMS-CC ‘17. Proceedings of the 2017 Workshop on Adaptive Resource Management and Scheduling for Cloud Computing, July 2017.
42. https://kokkos.org/
43. https://computing.llnl.gov/projects/raja-managing-application-portability-next-generation-platforms
44. https://csmd.ornl.gov/software/adios2
45. http://forge.ipsl.jussieu.fr/ioserver/wiki
46. https://www.maestro-data.eu/
47. O. Aumage, P. Carpenter and S. Benkner, “Task-Based Performance Portability in HPC,” Zenodo, 2021.
48. for instance DAOS (https://docs.daos.io/), CORTX (https://www.seagate.com/de/de/products/storage/object-storage-software/), MinIO (https://min.io/), and to some degree Ceph (https://ceph.io/).
49. https://singularity.hpcng.org/
50. https://www.docker.com/
51. https://podman.io/
52. https://sarus.readthedocs.io/

performance, and -- in more complicated form -- providing resource awareness to the container content from the host system cannot be considered 'solved' at this time. Secure and trustworthy execution of HPC workloads is in its infancy -- while Cloud system operations often rely on virtualisation to provide this, the HPC community efficiency concerns prohibit such costly isolation procedures. Future middlewares, in concert with the scheduling system, will need to address this challenge.

Since GPUs became a wide-spread component of large HPC systems, almost all applications have to deal with heterogeneity in computing and memory technology. While programming environments, compilers and runtime systems discussed above are intended to help improve efficiency it is ultimately the responsibility of the application programmers to exploit the heterogeneous hardware to the best possible extent. Consequently, much effort has gone into porting and optimising applications for GPU systems in recent years - and many of these efforts show great results. With further increased heterogeneity in many system components as discussed above these efforts need to continue and even intensify. Support in performance tools and models for all forms of heterogeneity are also important to enable this. Examples of ongoing efforts include the US Exascale Computing Project (ECP)[53] and the European Centres of Excellence in High Performance Computing[54].

In addition, applications are increasingly becoming heterogeneous themselves. Instead of single, large-scale simulations in many application domains different kinds of workflows are becoming a dominant usage model. This ranges from ensemble systems where the same application is executed multiple times with different parameters to complex workflows combining simulation, data analysis, and AI methods. Indeed, the increasing use of AI methods is drastically changing the way supercomputers are being used. In tasks like pre- and post-processing, application steering and in-situ analysis, AI methods play an increasingly important role. However, AI methods have hardware requirements that are often not compatible with the requirements of the application. Therefore, heterogeneous systems, where different parts of the system are configured in a modular way, offer an excellent platform for the execution of these modern workflows. However, efficient workflow systems (including scheduling etc.) are required.

RECOMMENDED R&D IN THE NEXT 2-4 YEARS

To tackle the challenges posed by heterogeneity in HPC as discussed above, research on all levels is needed. This starts with system architecture, which has to deliver a concept to efficiently combine and connect the variety of resources, including also opportunities to integrate new and upcoming disruptive technologies. For an effective use and share of resources, the resource management software needs to organise and orchestrate heterogeneous resources taking into account side constraints like thermal management and energy efficiency, as well as interactive usage and malleability of applications. More research is needed in these areas to provide the scheduler and resource manager with the necessary flexibility and dynamic functionality, potentially also applying AI methods to extract insight from both the system monitoring and application performance data. Also, efficient execution of complex workflows and large ensembles requires more attention, including analysis tools with introspection capabilities to better match application tasks to adequate hardware components.

The drastic changes in storage technologies ask for better support for the continuum from on-node to siloed storage, disaggregated storage, object store, and in-storage compute.

Security and trustworthy aspects are becoming critical also in HPC: Confidential Computing in HPC has traditionally been implemented via on-premises single-tenant systems. The on-going democratisation of HPC through cloud-like multi-tenant systems necessitates a secure solution to off-site job execution on "foreign"-hosted high performance resources. Cloud vendors currently utilise hypervisor-based virtualisation, however this is expected to be a performance inhibitor for HPC. Instead, we believe an end-to-end protocol for remote workload execution protected from in-storage, in-flight, and in-execution access by third parties through encryption and trust and identity federation is becoming important, and will also become a requirement to integrate with Gaia-X.

On the compiler side, more efficient vectorisation and parallelisation, potentially exploiting polyhedral compilation and just-in-time code generation are needed. In addition, compiler support, possibly in conjunction with appropriate runtime systems, for managing heterogeneous memory systems will be needed. Advanced runtime systems will need to support the placement of tasks on processing and memory technologies, explicit management of heterogeneous memories (not necessarily in a hierarchy), data transfers, data locality aware abstractions, malleability, and coupling of application tasks.

More research is also needed on programming models such that common models for GPUs, FPGAs, etc. can be provided. To increase the efficiency of programmers, high-level parallel programming models and DSLs can provide a more intuitive way of programming. In addition, compilers and runtime systems need to ensure interoperability/composability of different components (e.g. in-situ tasks) and ideally performance portability to the largest possible extent.

Tools also need to be adapted for heterogeneity. This includes particularly performance analysis tools, monitoring tools, and tools to identify the best resources for each step of the workflow and accordingly execute them.

Finally, applications have to exploit novel methods and algorithms enabled by heterogeneous components. This also

53. https://www.exascaleproject.org/
54. https://www.hpccoe.eu/eu-hpc-centres-of-excellence2/

includes novel formulations of existing models, using e.g. mixed precision arithmetic or non-IEEE data formats. Complex application workflows will be enabled via system software and tools, but the last-mile optimisations are likely to require intervention of the application developers.

All the above research areas must be well aligned and coordinated, in a concerted and coherent approach towards integration, support, management, and use of heterogeneous compute resources. This could be achieved by dedicated coordinated integration projects hardening and combining results covering all levels of the HPC ecosystem.

CONCLUSIONS

Heterogeneity in HPC started with the introductions of GPUs more than a decade ago. Since then, the trend towards heterogeneous hardware has continued and increased. Heterogeneity is now present in all system components, from computing elements, memory and storage technologies, to full system designs. This heterogeneity offers great opportunities for a more efficient use of HPC resources, but also poses severe challenges. To overcome these challenges, increased research efforts are needed on all levels of the HPC ecosystem. It is crucial that these efforts are perfectly coordinated, addressing hardware heterogeneity in a holistic and coherent manner.

6.2

Research Domains

HPC system architectures have been evolving over the last decades, with the long-term trend based on cluster computing approaches that combine standard off-the-shelf (increasingly heterogeneous) hardware. In the following sections first current research trends are described, followed by the challenges on the SRA-5 period. The chapter concludes with a description of the intersection between system architecture and the research clusters "Sustainability and Usability".

6.2.1

System architecture

RESEARCH TRENDS, CURRENT STATE OF THE ART AND FUTURE EVOLUTIONS

Architecting an HPC system requires balance between system performance and cost in a multi-dimensional trade-off analysis, evaluating computational power, communication and I/O, memory bandwidth, energy consumption, software support, usability, resiliency, purchasing and operating costs, and maintenance. The optimal set of parameters will depend heavily on the portfolio of applications to be used on the installed system. Therefore, it is crucial to include a comprehensive co-design strategy early in the process of designing the HPC system architecture, addressing the requirements of the future users.

HETEROGENEOUS ARCHITECTURES

One of the most difficult challenges for HPC system architects is to achieve the highest compute performance within an affordable cost and power envelope (for Exascale, a maximum power envelope of around 20 MW has been quoted). Reaching this target with homogeneous, clustered architectures based entirely on general-purpose CPUs is hardly possible. Therefore, most HPC systems include other compute elements (so-called accelerators), such as General-Purpose Graphics Processing Units (GPGPUs), which materially increase performance and/or energy efficiency. Accelerators can be integrated within a node via I/O links, which support remote memory access operations or provide a common, coherent memory system. As an alternative, accelerators can be connected via an interconnect and placed external to the nodes. Accelerator components are discussed in more detail in the WG2 (System Hardware Components).

The diversity in the hardware landscape has grown over time and heterogeneity has also entered the CPU package. We find processors using different Instruction Set Architectures (ISA), containing a diversity of execution units (SIMD, vector extensions, tensor processing, supporting different data types and/or variable precision etc.), and even CPUs that contain accelerators on-package. Also, a wider variety of accelerator devices is available, from general purpose (GPGPUs, many-core processors, FPGAs, Tensor Processing Units (TPUs)), through domain-specific accelerators (e.g,. signal processors, digital annealers), right up to disruptive technologies such as quantum or neuromorphic devices. Details on the individual computing elements can be found in WG2 (System Hardware Components), WG9 (Unconventional HPC Architectures, and WG10 (Quantum for HPC)). Here, we discuss approaches and means to combine all these technologies to create highly performant, energy-efficient and sustainable system architectures (see also White-paper on heterogeneous HPC).

As said above, traditional HPC systems were homogeneous clusters of general-purpose CPUs, until limits on power consumption forced the adoption of more efficient processing elements in order to further increase performance. Acceleration devices were first integrated within the node using I/O interfaces (typically one or two CPUs plus one or several GPGPUs), clustering the heterogeneous, accelerated nodes to scale up the machine. This static association of different resource types inside each node poses some constraints when assigning compute elements to applications. In response, disaggregated architectures (e.g., the modular supercomputing architecture[1]) have been implemented. Resources are no longer statically assigned; they are dissociated from the nodes to create pools of resources (e.g., accelerator pools), which can be dynamically reserved and allocated to each application. Here, efficient scheduling and resource management becomes crucial. Disaggregation can suffer from higher latency on the network side, but current network developments (e.g., Darpa PIPES) based on optical interconnects might help to address this issue. Meanwhile, applications are partitioned in a coarse granular manner for disaggregated architectures, limiting frequent communication to local subspaces (or modules), with inter-module communication occurring at the interfaces between models (e.g., in a multi-physics applications) or workflow steps. This is a similar approach as applied generally when optimising message-passing applications to minimise non-local communication.

All the examples mentioned above (homogeneous, node-heterogeneous, and disaggregated) consider the powerful compute resources (either CPUs or accelerators) as the central elements in the system architecture. Alternative strategies with different approaches have been proposed, such as memory-centric architectures, that creates centralised memory pools shared between compute units to minimise data movement. A variation of this concept keeps a CPU-centric design but attaches memory devices to the network,

1. E. Suarez, N. Eicker, Th. Lippert, "Modular Supercomputing Architecture: from idea to production", Chapter 9 in Contemporary High Performance Computing: from Petascale toward Exascale, Volume 3, pp 223-251, Ed. Jeffrey S. Vetter, CRC Press. (2019) [ISBN 9781138487079], http://hdl.handle.net/2128/22212

giving them some limited computation capabilities[2]. Network-centric approaches, on the other hand, put the interconnect fabric at the centre of the architecture design: instead of just transferring data between the nodes, network cards or switches are enabled to execute some operations, leaving only the most computationally demanding actions for the "strong" CPUs and accelerators. A further evolution of this scheme envisages network devices (referred to as dPUs or iPUs) to also handle scheduling and orchestration of compute tasks. The next section covers this concept in more detail.

DPU/IPU CENTRIC

To run more complex workloads and deploy ever more powerful architectures, HPC data centres grow bigger and bigger, and lessons can be learnt from operators of hyperscale data centres (mainly Cloud Service Providers or CSPs) which succeed in maximising utilisation and delivered performance by separating infrastructure from client applications. The latest step here is the introduction of "Data Processing Units" (dPUs) or "Infrastructure Processing Units" (iPUs)[3], which extend smart network interface controllers (NICS) to take over processing and computing tasks from CPUs or other accelerators, and most importantly, to control the infrastructure, including scheduling and orchestration of user workloads and services.

The objective is to improve end-to-end application performance, enable increased data centre throughput and improve overall energy efficiency. Consequently, we anticipate that dPU/iPUs will join CPUs and GPUs (or similar accelerators) as one of the three pillars of the data-centric accelerated computing model.

dPU/iPUs support a strict separation of infrastructure management and client applications. Infrastructure management tasks are moved off CPUs, making these 100% available for user tasks and minimising latency and jitter. iPU/dPUs possess scaleable accelerators that enable them to execute storage, network, security, micro-services and workload management functions. With the latest dPU/iPUs, traffic between CPUs, accelerators, and storage now happens on the "fast path", achieving low latency, high throughput of data and maximum flexibility. In addition, the distributed dPU/iPU infrastructure monitors the status and load of all datacenter components in detail and can make data-driven decisions on efficient placement of client tasks.

We can certainly anticipate increasing deployments of dPU/iPUs in the hyperscale data centres of CSPs, and a concurrent rapid evolution of hardware and software. Supercomputing faces analogous challenges with the management of their heterogeneous infrastructure in terms of optimising throughput and energy consumption. Ergo, as we move to pre-Exascale and Exascale architectures to run ever more complex, data-intensive workloads, dPU/iPUs could enable an infrastructure-optimised execution environment. The technology is being applied to AI and data analytics applications today, and its applicability to typical HPC workloads has to be investigated.

Summing up, a successful adaptation of iPU/dPU technology could provide significant benefits in overall utilisation of supercomputing resources, the adaptability to dynamic conditions, the turn-around times of user workloads and data centre throughput, and the overall energy efficiency.

SYSTEM-LEVEL PACKAGING & INTEGRATION

HPC systems are large enough to allow specific designs at the data centre, rack, or motherboard level. These designs address some of the core HPC KPIs: performance, density, and energy efficiency.

The HPC community has encouraged the use of advanced cooling techniques to achieve a greener environment and reduce costs [see Section 1.4]. One of the targets is to remove air conditioning systems and save their associated electricity consumption. The most widely adopted cooling technology today is direct liquid cooling. Even immersion cooling is sometimes applied. Additionally, new studies and proofs of concept underway aim to demonstrate that the use of solar farms producing green hydrogen could make a data centre carbon-free. Producing and storing energy and capping the power of the entire system both aim at making HPC more environmentally friendly.

Density of compute power (Flop/s) per square metre is one of the main factors of merit for racks designed for HPC. Maximising this metric requires tightly integrating and cooling the most powerful computing elements (CPU, GPU, AI accelerators, etc.), which are also those with highest power consumption (500W+). As a result, the size of HPC racks is increasing to provide more power to compute boards and dissipate more heat.

Increasing the density of the system, besides requiring less floor space in the customer's data centre, has a big impact on the HPC network as it can reduce the lengths of cables. Today three categories of cables are used: passive copper cable, active copper cable, and optical cables. Their respective costs are significantly different: shorter distances can be covered with less expensive copper cables; for long distances, standard optics and silicon photonics are now the two technologies available. The number of ports per switch increases (now in the 64-port range), as does the speed of the links (now in the 400Gbps range). Consequently, cables and connectors are constantly evolving (e.g., high density Octal Small Form Factor Pluggable (OSFP) cables).

The HPC motherboard integrates the computing element (CPU or accelerator) and external interfaces. The main interfaces are PCIe / CXL (gen5, 1x to 16x), network such as Infiniband, Omnipath, Ethernet, or Slingshot (400 Gbps, x4), memory modules (DDR5), and coherent links (similar to PCIe gen5 but with a different number of lanes). Numerous companion peripherals can be connected to the PCIe / CXL depending on the configuration chosen (disk, flash module, CXL memory extension, etc.). The increased speed and width of the interface leads to larger sockets,

2. J. Schmidt, "Network Attached Memory", Chapter 4 of the PhD Thesis: "Accelerating Checkpoint/Restart Application Performance in Large-Scale Systems with Network Attached Memory", Ruprecht-Karls University Heidelberg - Fakultaet fuer Mathematik und Informatik. [Online: http://archiv.ub.uni-heidelberg.de/volltextserver/23800/1/dissertation_juri_schmidt_publish.pdf]
3. NVIDIA uses the dPU term, and the iPU name was proposed by Intel.

more layers on the printed circuit boards (PCBs), new connectors, and finally materials.

Multi-die packages with different chiplet components (CPU cores, HBM, NIC, etc.) allow embedding in today's products part of the component which yesterday was still part of the motherboard. With this change comes new trade-offs and constraints. As an example, the integration of HBM DRAM will provide a large bandwidth but with a smaller capacity than DDR DRAM, and will add a new constraint in terms of maximum case temperature and power dissipation. Other candidates for the integration of components such as the network interface are discussed in WG2 (System Hardware Components).

ENERGY EFFICIENCY AND SUSTAINABILITY

State-of-the-art research in the energy efficient HPC systems shows that the Exascale era will not be limited by the performance growth of systems, but by their power consumption. A reasonable maximum power envelope for Exascale systems is 20-30 MW, which is today the power available to many large data centres. Since the available power is limited, maximising energy efficiency is a critical aspect when designing HPC systems and components, and when configuring HPC data centres. It is also critical to limit the operating costs of HPC systems, which are in large part due to the electricity costs.

Trivially, the quest for energy efficiency has to start with the system components; processing elements (in particular CPUs) have been a focal point for energy analysis and optimisation, yet other system elements, such as memory and interconnects, will need to be optimised as well. Details of this can be found in WG2 (System Hardware Components). Non-conventional architectures (as described in WG9 (Unconventional HPC Architectures), and WG10 (Quantum for HPC)) will also play a role here.

System architecture can improve energy efficiency by tailoring the system characteristics to the needs of applications, ensuring that all parts of a system are used at their best operating points, and no elements are left idling while drawing significant power. Due to the variety of workloads a typical HPC centre is supporting, and their dynamic nature, flexible and adaptive ways to partition and apportion resources are required.

In addition, system architecture can optimise the efficiency of the power supply and cooling sub-systems. Examples are the proposed use of single voltage supplies and of highly efficient, dynamically managed power converters close to power-hungry components, and the use of advanced, liquid cooling approaches.

Finally, the design and configuration of HPC data centres plays an important role. The ratio of total energy consumed by a data centre vs. the energy used by the HPC systems it operates is referred to as Power Usage Effectiveness (PUE). Today's average PUE of a data centre is between 1.3 and 1.5, and future generations should target an average PUE of <1.1, which is only possible with highly efficient power supply and cooling systems and equipment carefully designed for energy efficiency.

The power density of an Exascale system will require more than 100KW per rack, which is at the limit of what today's air-cooling systems are capable of dissipating. This forces data centres to use direct liquid cooling technology, in particular warm water cooling (typically 30°-50° C), which is currently the preferred approach for handling such high-power density and reaching target levels of PUE.

An efficient monitoring infrastructure capable of analysing and correlating different performance metrics is a key requirement for achieving gains in energy efficiency via adaptive management of systems and installations. In particular, data centre infrastructure management (DCIM) and HPC monitoring systems should strictly cooperate to optimise the overall efficiency of the data centre orchestrating HPC systems with the cooling and electrical systems of the facility. DCIM and monitoring systems should leverage AI technologies to reach higher levels of efficiency in terms of power consumption and computing performance.

SYSTEM ADMINISTRATION

HPC, AI and Big Data applications often run on a general-applicability HPC infrastructure, built as a cluster of powerful high-end machines, interlinked with high-bandwidth low-latency networks. The compute cluster is commonly augmented with hardware accelerators (co-processors, GPUs or FPGAs) and a fast, large capacity parallel file system. All equipment and services are set up and tuned by systems administrators.

While high-end, the infrastructure is typically still based on commodity hardware to leverage technology developed for cloud service providers and common data centres.

Nonetheless, HPC systems have special needs which system administration software needs to address. Open-Source HPC management software like xCat or Warewulf can replace costly proprietary solutions and help democratise HPC especially for Tier 2 and smaller systems. However, careful considerations have to be made as outlined in the Open Source Approach Chapter.

System administration software must support a wide range of functionality [cf. WG3 System-SW and Management]. User management is typically done by LDAP or other directory services. The deployment of bare metal servers must be supported - even if those servers do not have disks for operating systems (OS). Accelerators must be included in the management and monitoring software and interconnects such as InfiniBand or Slingshot need to be supported in addition to Ethernet. RedFish is an open hardware configuration and monitoring standard implemented by more and more hardware vendors. However, there still needs to be refinement and standardisation to support automated system administration. For an effective vendor-independent control of the hardware, these interfaces must be accessible in-band or out-of-band and documented in detail. This enables local management controllers or central management units to increase efficiency by evaluating sensor data, coordinating, and dynamically controlling the heterogeneous hardware. The data can be used by the different layers of the whole system, up to the user

application layer to accomplish different targets, such as energy or cost efficiency, reliability or availability.

For exascale systems, a scalable solution which supports tens of thousands of machines is essential.

MEMORY AND STORAGE HIERARCHY

The trend towards memory heterogeneity and the blurring of distinction between memory and storage continues. Deep memory and storage hierarchies are built using a variety of technologies, allowing optimisation for either higher bandwidth (on the upper layers of the hierarchy), or higher capacity and lower price (on the lower layers). Additionally, technology which reduces the data management and movement load on CPUs and accelerators is becoming available.

DDR4 will be phased out in favour of DDR5, offering memory bandwidth increases of more than 30%. In 2022 mainstream accelerators and CPUs will utilise HBM2 memory. HBM3 memory with bandwidth of up to 819 GiB/s per package is in active development. Non-volatile Storage Class Memories (SCM) used both as memory (NVDIMM) or as a middle-layer-storage in front of solid state/NVMe drives or traditional hard disks are utilised today in production systems and serve Big Data analytics and Artificial Intelligence workloads.

Compute node internal in-memory computing devices can reduce data movement operations, as well as time- and energy costs for these operations (cf. WG2 Hardware Components). Finally, with performance increases in interconnects, the concept of memory pools which can be dynamically subdivided and associated with accelerators and CPUs can become reality. This approach that matches the disaggregation concept described in Section 1.1 would address the common over-provisioning of memory for the worst case (i.e., most memory hungry workload) and thus contribute to cost and energy savings.

NETWORK AND INTERCONNECT

The tremendous increase of the number of computing elements in supercomputing systems brings new challenges in terms of interconnect. In particular, high bandwidth and low latency interconnect such as Mellanox InfiniBand, Intel Omni-Path and Bull eXascale Interconnect used in today's supercomputers are growing in performance, but not as fast as the computing capacity. They are becoming the main bottlenecks in the exascale era. In particular, communication latency must be in order of nanoseconds to effectively scale on millions of processors moving exabytes of data.

Another important aspect of network communication is power consumption. While over the last five decades powerful supercomputers could be built without major concerns on energy consumption, with the advent of exascale systems this is no longer sustainable. It is estimated that today's data centres use between 10% and 40% of total power consumption only to move data, but that percentage is expected to grow in the next few years due to the data-volume to be processed. Also, the movement of data across data centres is expected to increase in the next few years. This is particularly true for European HPC data centres where users will share exabytes of data. Next generation data centres and HPC systems should focus on technologies that can efficiently move data among computing resources. Photonic die-to-die interconnection is a good candidate to solve this issue thanks to its ability to directly route light instead of electronic signals, which drastically reduces latency and power consumption, improving the overall bandwidth.

The technologies able to optimise latency and bandwidth at all levels of communication, from die-to-die to data centre communications, should be the main target when aiming at further increasing the peak performance.

SYSTEM LEVEL MODELLING AND SIMULATION

End users across industry and academia are consistently concerned with qualitative and quantitative evaluation and modelling of HPC systems to optimise utilisation and capacity, and ensure satisfactory application performance. This includes benchmarking and simulation based upon state-of-the-art theories, tools and practice. The global HPC ecosystem recognises that the coverage of the term "performance" has broadened to include power consumption (energy efficiency) and reliability, and that performance modelling can be delivered both through analytical methods and simulations. As HPC systems increase in complexity and heterogeneity, it becomes more difficult to realistically predict their expected behaviour; the same holds true for end-user applications, the performance of which gets harder and harder to approximate by analytical methods.

Aside from vendor solutions, simulation environments which could provide detailed insight into system and application behaviour for thousands of computing elements and the associated infrastructure are not available. Research and development is required[4] to evaluate new approaches, which would involve use of accelerators and techniques like ML/AI, to provide such a capability for systems approaching pre-Exascale or even Exascale levels. Truly precise simulations and predictions will not be feasible at that scale - yet reduced-accuracy systems would still be useful, as long as their limitations are understood. Similarly, ways to create end-to-end performance models for applications based on system metrics applicable to realistic-size, heterogeneous HPC systems would be extremely valuable.

It is worth noting that 20% or more of compute capacity in a high-performance computing system is wasted due to failures and recovery responses (cf. Fault tolerance below). Consequently, fault tolerance in exascale systems becomes a critical requirement, where poor design can result in significant TCO overheads.

FAULT TOLERANCE

As HPC system densities, scale and power requirements increase, fault tolerance assessments frameworks will become increasingly critical in maintaining maximum utilisation and resilience.

4. Early work in Europe, for instance on the message-passing simulator Dimemas and the CPU simulator Sniper could be taken up and continued; in the US, the concepts underpinning the Aspen project do look promising.

Consequently, traditional assessment frameworks consisting of periodic checkpoint/restart (CPR) policies and protocols may no longer be best practice.

We need to enhance these frameworks to support high levels of resilience that are critical for continued operation of extreme scale HPC. At exascale level, this requires management of the impacts that directly result in decreased reliability through increased hardware and software complexity and cost risk of unexpected issues (e.g., manufacturing errors/faults, dirty power, and accelerated component wear and tear).

Exponential growth in core counts of HPC systems carries the potential to reduce fault-free compute cycles, and dramatically increases the impact of hard and soft faults. Resilience, redundancy, availability and serviceability features include all the components required to keep systems functional with maximum Mean Time Before Failure (MBTF). The ongoing transition to exascale constitutes an excellent opportunity to establish future objectives, requirements and standards for resilience and fault tolerance, that may translate into significant cost savings.

Consequently, more economical resiliency strategies should be implemented to replace CPR. Decisions and compromises need to be made on whether to focus exclusively on frameworks and APIs for fault-tolerance and fault management.

Metrics and standards for measuring, improving, and enforcing effective fault-tolerance of an extensive range of HPC architectures and applications will be essential in the technology roadmap and procurement strategy for Horizon Europe. This can be supported by an assessment framework that addresses five key categories: procurement, preferred system validation, design of future systems, application code optimisation and redesign, and data intensive workflow scheduling.

FEDERATED COMPUTING IN HETEROGENEOUS HPC, CLOUD AND DATA INFRASTRUCTURES

The traditional way of using scientific HPC systems is for end-users to apply for accounts and quota at HPC centres one at a time, and explicitly select the systems an application or workflow should run on, providing the centre-specific credentials. Federated computing promises to enable users to access any system from a pool of federated systems (across multiple centres), and moreover have an application or workflow automatically run on an available system, governed by service-level agreements. For instance, with a "fastest result" SLA, a federated computing installation would direct an application to run on the system in the pool which will provide the results fastest.

Federated computing does require HPC centres to agree on a common scheme to authenticate the identity of users, to operate a unified method of administering quotas and pay-per-use (if applicable) based on these "federated identities", and to provide higher-level discovery, reservation and meta-scheduling services across the participating centres and systems. It also requires fully automated ways to redirect execution of applications and workflow steps between centres, and to make all required data available. Data, network and large instrument resources have to be supported in addition to "just" compute.

From the above discussion, it is clear that there are no specific requirements of Federated Computing towards the system architecture, except that all systems have to be able to efficiently support the authentication and authorisation methods used, that data protection and encryption have to work for the federated identities, and that sufficient monitoring data is provided to eventually support a "meta-scheduling" scheme.

ARCHITECTURE IN THE DIGITAL CONTINUUM

Over the last two decades, our life has become increasingly dependent on the integration, processing and analysis of large amounts of information in different formats and from different sources. HPC is expanding beyond its traditional modelling and simulation applications and it is now directly interacting with us in many aspects of our life, e.g. autonomous cars, localised weather forecasts, cloud computing, AI, smart cities and Industry 4.0.

The notion of digital continuum denotes the set of computing resources and software applications acting together and cooperating in order to complete a complex task/workload involving multiple systems and multiple data streams from multiple sources. Examples include the management of a fleet of autonomous cars, or the Destination Earth effort. The end-to-end infrastructure required to implement a digital continuum use case must enable: i) value extraction value from large volumes of data quickly and reliably; ii) reduction of overall costs by minimising the amount and frequency of data transferred over a network; iii) reduction/minimisation of security vulnerabilities by retaining data locally for privacy protection; and iv) placement and orchestration of constituent tasks to minimise the overall energy required to achieve the use case results.

Therefore, applications will not be built around a monolithic piece of code running on a single platform, but will span across various resources (each performing a specific part of the computing, storage, AI workload) distributed in multiple locations. Some parts of the workloads, such as modelling, simulation and AI training, likely will run best in a centralised data centre, and other parts of the workload, such as face recognition, intelligent traffic lights and physical security, are best suited for the edge to enable local real-time data-driven decisions.

The architecture of the digital continuum infrastructure is a multivariable optimisation problem that includes considerations about the need to react to mission-critical requests and application performance, the size and frequency of the data, latency-sensitiveness and SLA, cost and security, data sovereignty, and communication bandwidth. HPC data centres play a key role in the digital continuum by providing the HPC support to the infrastructure by means of specialised HW and SW components performing at best a specific part of the workload. The technology vendors' and System Integrators' role is to design components and equipment enabling the end-to-end physical and

logical exploitation of the data/information across the resources.

The SRA working group "Centre-to-Edge" (WG8) takes a detailed look at the technical challenges involved to realise the Digital Continuum.

HPC FOR URGENT DECISION MAKING

The system architecture of HPC systems must support priority management and pre-emption in combination with checkpointing to allow a prompt response to incoming urgent workloads without sacrificing a significant amount of work. With the current level of heterogeneity in hardware and software architecture, support for container technology is mandatory to reduce portability issues. As input data must be accessible from within the HPC system and output data must be accessible by decision makers outside of the system, advanced data management and caching services should be implemented. To address current data sizes a minimum in- and outbound i In conclusion it is not realistic to design every HPC system for urgent decision making, as some needed features may be out of scope for some centres. Thus, one should designate a number of HPC centres as targets for urgent decision-making workloads and design their systems from the ground up with the needed abilities in mind. Internet bandwidth of significantly more than 100GBit/s is advisable.

CHALLENGES FOR 2023-2026

In all the topics described in the section above, challenges will need to be addressed in the time frame 2023-2026. Here, we summarise the most important and urgent aspects to be tackled:

Efficient orchestration, management and use of heterogeneous resources: In the next few years Exaflop-capable computers will be deployed, composed of diverse compute, memory and interconnecting technologies. Research and development is required at all levels of the HPC ecosystem to make these heterogeneous systems usable for the HPC user community with high efficiency. The system architecture has to combine and physically integrate heterogeneous elements at high density and energy efficiency, but most important will be support for a strong, Europe-driven heterogeneous computing software initiative. Monitoring infrastructures and cluster management software have to combine a vast amount of multi-scale sensor data and extract the information required for power-efficient operation of the system. The management software must enable the scalable installation of software packages and updates that include device specific elements. Resource manager and scheduler software must be adapted to enable the dynamic, adaptive and malleable reservation, allocation, management and use of the resources case-by-case, for jobs and fine-grain workflow steps. Long term, fast, low-overhead scheduling and orchestration mechanisms should also cover novel system components such as smart NICs (dPUs, iPUs). Programming models and tools for performance analysis and modelling have to be made composable (interchangeable from the application perspective) so that the best suitable model is employed for each hardware component, using vendor-specific models when necessary, but keeping performance portability as the ultimate target. Finally, the application codes themselves have to be adaptable for the use of heterogeneous systems, either by abstracting them behind high-level programming models, or by including Domain Specific Languages (DSL) approaches with back-ends for each of the necessary technologies.

Integration of disruptive technologies: the challenges mentioned above for heterogeneous HPC grow exponentially when integrating technologies such as neuromorphic or quantum devices, which display different operational modes and software stacks than the standard HPC devices. Research and development is required to enable interoperability of traditional compute devices and disruptive technologies in future HPC systems. Challenges will exist on the hardware side (e.g. quantum computers have much stricter operational requirements than normal HPC systems, including near-zero temperatures and vibration isolation from the environment), on the system software (e.g. enabling offloading operations and exchange data between standard and disruptive devices), and on the application side (e.g. extracting the parts of the workflow that can profit from a neuromorphic device and adapting the code to exploit this kind of hardware). Successfully implemented, the integration of disruptive devices on HPC systems will bring a broader user-base which should foster their acceptance and adoption by the wider HPC user community.

Modelling and predicting performance, power consumption and behaviour of highly heterogeneous and complex systems: Limitations of traditional architectures have resulted in the emergence of heterogeneous architectures that combine general purpose CPUs with accelerators such as GPUs and FPGAs. The major challenge of highly heterogeneous architectures is the (usually significant) effort required to port existing applications to these new system architectures, especially since the performance bottlenecks of HPC applications may vary according to the accelerator's technical features. Moreover, the choice of a particular accelerator may not be readily apparent. Performance models can provide valuable information for the selection of the accelerator as well as for the design of new architectures. To predict speedup and compare different architectures of accelerators, the following steps are typically implemented: i) identify sections of existing source code that are potential candidates to run on accelerators. HPC applications' computational and data access patterns can be expressed by a small set of commonly occurring functions (e.g., reduction, transpose, matrix operations, FFT, stencil and the like); ii) identify the impact of moving data between accelerators and CPUs. The cost of transferring data can become the performance bottleneck; iii) gather HW and SW features, performance, statistics and counters for an accelerator's architecture; vi) develop a model (digital twin) of the accelerator based on the technical features and components to simulate the behaviour of the accelerator; v) run and fine tune the model. Modelling tools should also cover aspects such as energy efficiency and allow predictions of both runtime and energy use.

Monitoring: As we embark on a new decade of technology transformation, the principal drivers of infrastructure management (from component level up to data centre environments) will be smart controls and monitoring using ML, full infrastructure automation, maximising utilisation, system performance, resiliency and energy efficiency. These should enable monitoring of actual aggregated system energy use per job and workflow step, including data transfer and storage. Going forward the deployment of an ever-increasing number of sensors and data points will require application of ML algorithms to harvest and learn legacy monitoring data. This will enable HPC environments to predict overall resiliency and performance, and in turn dynamically adjust data driven predictive management models. Smart monitoring and cross-platform controls are being deployed to integrate with the Board Management System; modern simulation and modelling systems automatically update themselves and schedule these actions in the most efficient manner to minimise disruption. This contributes to directly improving operational efficiency and further minimises service management intervention. Similarly, AI/ML and automation technologies already have a significant role in enabling sustainability, by improving utilisation of environmental data and fine-tuning of mechanical, electrical and cooling infrastructure. This in turn improves energy efficiency and supports carbon footprint reduction.

In particular, exascale HPC systems that will employ hundreds of thousands of computing nodes will create new challenges in data management for monitoring systems. An Exascale HPC system will be the first source of big data that will need to be processed to tune and optimise the system. A cloud environment to support monitoring systems will be needed to enable exascale HPC to process such vast amounts of data, thereby closing the loop of cooperation between typical HPC systems and cloud environments.

Resiliency: System resilience is one of the hardest Exascale requirements, particularly due to its cross-layer nature. The European HPC community, however, currently lacks a strong concentrated research effort in resilience, which makes resilience one of the greatest challenges of the EU HPC initiative. Ensuring the resiliency of large-scale HPC systems is complex and requires research and engineering effort for the analysis, development and evaluation of reliability features. Additionally, resilience is a vertical problem that needs holistic solutions. For all these reasons, we advocate that HPC system resilience is properly represented in future research and innovation programmes. Resiliency is crucial to ensure the success of the extreme-scale HPC systems with hundreds-thousands components. Further challenges arise from the interplay between resiliency and energy consumption: improving resilience often relies on redundancy (replication and/or checkpointing, rollback and recovery), which consumes extra energy.

Energy Efficiency: Minimising the energy requirements of operating HPC data centres is the key challenge which supercomputer facilities are facing. Solutions will have to come from all levels, starting from semiconductor design and materials/production, via component architecture and integration, system architecture and data centre set-up. Challenges at the system architecture level concern efficient and adaptive configuration of resources to best match running applications, integration of highly efficient power supply/delivery and cooling methods, and providing the system monitoring information required for a near-optimal management of the computer systems and the data centre as a whole.

For data centre design and operation, the trend towards more and more performant, yet also “hotter” packages (which is a corollary of high-density integration) and nodes will be forcing a transition from air-cooling to more efficient liquid-cooling systems, since the installed power per rack can no longer be handled otherwise. This also raises the bar for power delivery and supply systems. Raising the temperature of the cooling liquid can result in substantial energy savings, since it enables free adiabatic cooling for > 95% of the year in northern regions such as the UK, Germany or Scandinavia. The challenge here is to handle the increased heat load per rack mentioned above.

Overall, the challenge to European data centres is to target a PUE lower than 1.1. Besides warm liquid cooling as discussed above, a data centre must consider the entire system stack, from applications to HW subsystems and the cooling and electrical facility support. Tools to improve the energy efficiency of the HPC systems play an important role in the data centre stack, since they can leverage system and power management knobs of the system and the application workloads in order to optimise overall efficiency for both performance and energy. For this reason, achieving multi-layer orchestration of the whole facility is crucial to reach high energy efficiency, and it should be supported by R&D activities from data centre designers, system and technology providers, and developers of system and application software.

Security: The objective of product security can be summarised as designing trustful products that show no known vulnerability, prevent or deny attacks at customer or supply chain level, and most importantly, allow recovering to a safe state in case unknown vulnerabilities are exploited against them. Regulations around critical infrastructure protection evolve fast and will most likely propagate in one way or another to HPC systems, which face increasing security demands. Security starts from the platform where hardware protected root of trust, signed firmware and software, and controlled access allow a safe delivery and upgrade of an HPC system. Another objective is to protect the data of each user in a shared infrastructure, and to provide HPC mechanisms protecting against denial of service. This will induce specific developments in all system layers from hardware components to operating systems and networks. Additionally, some HPC centres may require additional attention to share their infrastructure between different companies working on sensible data, defence-related applications, and users processing medical or personnel data, to name just a few extremely security-concerned areas. However, not only these areas, but all the HPC

community in general is demanding growing levels of security in HPC systems and centres. Keeping up with them will require new security approaches to be applied holistically at all layers of the HPC environment, from the hardware, through the software, up to the application execution, and from the HPC centre to the edge.

Federation: As discussed in section 1.8, federated computing in itself does not pose significant system architecture challenges per se. Systems would need to efficiently support whatever authentication & authorisation methods chosen, data encryption and protection would need to work on the level of federated identities, and reservation and meta-scheduling of compute, data and network resources is required. These challenges mainly concern the "System Software" (WG3), with the potential exception of system management, which would need to provide sufficient functionality for the higher-level layers.

Open collaboration: Today, we are missing an open standard strategy to strengthen industry-wide research into hardware and software. Greater collaboration across users, data centres operators, HPC vendors and integrators and across scientific communities is needed to confront related issues. In particular, machine learning research clearly has a major influence on future hardware developments. It will be increasingly important for user communities and HPC vendors to agree on a long-term strategic vision for collaboration and co-design that encompasses the whole technology stack. This could be done building on existing relationships, developing broader strategic partnerships that, in turn, form connected networks that blend academic expertise with industry, government, vendors, integrators and other collaborators to research, catalyse, apply, and drive research and development in HPC. Such partnerships can provide important benchmarks for driving breakthroughs in architecture design into production. They can also provide a vehicle for sharing confidential, commercially sensitive and proprietary data on technical subject matters such as system performance and resiliency.

Co-Design has been a successful methodology for fostering collaborations between application scientists, software developers and vendors for many years. The process can be interpreted as identifying the full range of elements that are limiting performance across the technology stack, and assessing whether it makes technical and economic sense to develop improved architectures. It will become increasingly important to deploy co-design strategies to address requirements for greater energy efficiency, resiliency and reliability that are brought about in the deployment of future exascale systems. Where co-design encourages systematic and agile processes for rapid product development or selection, not all organisations using HPC have the investment or resources to engage at this level. They may be content to source HPC products off the shelf, rather than embark on a complex co-design process to identify elements that are limiting performance. To mitigate this, we can bring together the full range of core technology disciplines to develop architectures that support solutions to difficult methodological and real-world challenges. Co-design can yield positive technological transformation across all technology towers that will support cross-domain Horizon Europe projects. This includes bringing together different research disciplines, scientists and engineers to deliver future bleeding edge technologies that no single organisation or vendor can do alone. Ultimately, co-design should be preferred over reliance on technology vendors to deliver future architectures.

INTERSECTION WITH "SUSTAINABILITY" RESEARCH CLUSTER

Sustainability is a global challenge that is tackled by system architecture from various complementary perspectives, ensuring that systems are as energy efficient as possible, and as durable as possible.

Energy efficiency has to be considered at all levels of the HPC centre, from the system to its environment, and the associated challenges have been described already in Section 2 of this Chapter. At component level the most efficient elements (processor, memory, interconnect, etc.) need to be combined in a manner such that their operation consumes the least possible amount of energy and enables assignment of the most appropriate resources for each application. This brings as consequence the need for a comprehensive, multi-scale monitoring system from the component to the infrastructure level, in which the data is accumulated and correlated with the runs of the individual jobs and applications and is fed back to the system-administration to reduce the overall energy consumption at each point of time.

Durability, understood as long-term usage of the materials, often requires compromises on energy efficiency, since the latest computing elements are generally the most energy efficient ones. Operators perform a trade-off between the costs (both for the site and for the environment) of continuing the operation of an existing HPC system vs. exchanging it for a modern, more energy-efficient one. Currently, this leads to a typical lifetime of about 5 years for HPC systems. System-architecture concepts are needed, which enable efficient, longer-term operation of at least parts of the HPC infrastructure and renewing some of its elements without simultaneously being forced to exchange all the rest. Resiliency also plays an important role for durability: more resilient elements can be operated over a longer lifespan. However, higher component resilience comes at a price, with more expensive materials and manufacturing processes.

Disaggregation of resources, as proposed in the modular supercomputing architecture, tackles energy efficiency and durability by introducing a new level in the system hierarchy, between the node and the system-level: the module level. The benefits are manyfold: resources can be shared more efficiently between users, which maximises their usage and therefore the scientific outcome of an HPC-system; modules can be sized so that only energy-efficient resources are scaled up, while the number of power-hungry CPUs is kept low; and outdated modules can be substituted by more modern ones keeping the rest of the system in operation.

6.2.2

System hardware components

RESEARCH TRENDS, CURRENT STATE OF THE ART AND FUTURE EVOLUTIONS

INTRODUCTION

Processor architecture evolves hand in hand with semiconductor technology. At the turn of the century, as transistors became smaller, the performance of processors was improved principally through frequency increases at each new technology node. The majority of processor architectures were single core. Power consumption remained under control as the current draw for higher frequency was compensated with the reduction of the supply voltage (Dennard's scaling).

Around 2006, Dennard scaling broke down as clock frequency scaling hit a practical ceiling, prompting processor architectures to move to many-core to improve performance, and implement low power design techniques to keep dynamic power consumption reasonable. For roughly the following decade the growth in processor performance was driven by increasing core counts.

Around 2016, the processors reached a power consumption envelope limit due to increasing leakage current. Gains on the energy efficiency of compute nodes were achieved by a move to heterogeneous many core designs, using acceleration to boost performance beyond what general purpose processors could support in an equivalent power envelope.

Today the majority of the HPC systems, especially at the top end of the performance scale, use GPUs or other types of accelerator/coprocessor technology (like the Intel Xeon Phi) , and this trend is growing. Future systems will require further hardware/software co-design to meet their power and performance goals, as exemplified by Fugaku[1] at RIKEN. Thus, there is a trend towards greater heterogeneity.

Looking forward, we foresee the cost wall incurred by the slowing of Moore's law and the increasing cost of fabrication processes. Modularity at package level can be achieved by composing a chip with smaller interconnected and heterogeneous dies, also referred to as SiP, or System in Package. This chiplet technology will enable reuse of the costly silicon in different products, and even to extend to markets beyond HPC, extending the business case for development of top performance silicon.

This many-core, heterogeneous decade is marked by the race to Exascale computing. Other trends driving growth are the proliferation of AI and data centric processing, and more recently the political change in perspective, from global outsourcing to the investment in sovereign core competencies. New features of processors, such as support for virtualisation, for crypto, compression, new data types, expanded vector processing support, matrix multiple units, security (like SGX and others), will allow new usage of HPC such as "cloudification" and federation.

From early designs, computer systems have maintained a strict separation between CPU and memory, forming the von Neumann bottleneck. This bottleneck, present from small scale embedded systems to the high-performance supercomputers, is especially detrimental to data-centric applications, which poorly match the assumptions underlying a modern memory hierarchy: large data volumes, low operational intensity, low spatial and temporal locality, and unpredictable data access patterns. These characteristics arise in important application domains including data analytics (big data and database workloads), machine learning, artificial intelligence, deep learning, bioinformatics (genome analysis), and scientific computing. New memory technologies, like non-volatile memories with persistent storage, low latency, fast access and large capacity will lead to new memory hierarchies and new ways to address information, like object addressing (exemplified by DAOS). Finally, the cost of data movement is increasing, which impacts the global efficiency of HPC systems. This can be addressed by moving data and compute closer; proposed near term solutions include the use of data management accelerators (DPU or IPU), or moving (simple) computation near or in memory.

New hardware features will enable new capabilities, but it will be up to the system architects and software to take advantage of them. A new era of massive scale computing is arriving, with associated communication and storage needs. HPC systems will be the lighthouse solutions for the new system architectures and associated SW stacks.

The following paragraphs will describe the state of the art in computing elements (processors, accelerators), infrastructure and supporting components (IPU, DPU), memory systems, intra and extra node interconnect, and multi-die device architecture.

COMPUTING ELEMENTS

Processors:

Most high performance processors are designed for cloud and servers, and a few key elements merit discussion for applications in HPC; specifically Instruction Set Architecture, data types and vector processing, security, and virtualisation.

- Instruction Set Architecture (ISA): various ISAs are considered for server processors today, notably x86, Arm[2], PowerPC[3] and RISC-V[4] (but still strongly dominated by x86 architecture). The instruction set is the basis for a processor design. From a pure hardware perspective, it can be argued that any ISA can serve for any type of processor, it is just a question of design.

1. Fugaku (number 2 in TOP 500 of June 2022, and number 1 previously) was the culmination of the collaboration between Fujitsu and Arm in defining Arm´s Scalable Vector Extensions (SVE) standard; and thus required feedback from the software to drive the development of the hardware SVE specs. Additionally, the proprietary TofuD interconnect was also developed in a hw/sw co-design process. The latest Exascale class US supercomputers have done a similar hw/sw co-design process
2. AWS Graviton line of processors, Graviton3 being the latest release in the family of processors designed by Amazon
3. IBM Power10 is used in IBM servers
4. SiFive's Performance P650 RISC-V processor core is aimed at high-end servers

However, as an ISA matures, legacy hardware and software are developed and standardised. For server processors the SW ecosystem is particularly important - server end users expect to pull SW from repositories without having to patch or directly modify the code. Development of HW standards and the associated software ecosystem takes years, money, and community mindshare. As an example, if we consider that Calxeda was founded in 2008 and released the first Arm server in 2012, it has taken >10 years for Arm servers to become viable alternatives with 5% global market share[5]. As we reach 2022 there are 3 ISAs generally recognised as viable server processor alternatives with concrete market presence: x86, Power, and Arm. RISC-V, an open-source ISA, is a promising new architecture with significant traction in the embedded and acceleration applications where the software ecosystem is encapsulated and more easily managed. Given the enthusiasm for RISC-V server processors (geopolitical factors, free licensing model, large momentum behind RISC-V investments), and the fact that experience with Arm can be used as a template for new ecosystem development, it is reasonable to assume that RISC-V server processor ecosystem will develop more quickly than the Arm ecosystem did. Discussions in various HPC forums and consortia suggest a timeframe of 5-10 years for the ecosystem to become self-sustaining. In the current top500 list, processor architectures are based on x86 (AMD EPYC and Intel Xeon) accelerated with GPUs, IBM Power (also mainly GPU accelerated), and more recently Arm (the Fugaku system based on processors from Fujitsu is the only Arm-based architecture in the top 10).

- New data types and vector processing: CPUs are increasing system energy-efficiency by moving accelerator-instructions into the CPU. As an example, large vector instructions are included in CPUs to process more data with fewer instructions. The Fujitsu Arm-based A64FX processor for the Fugaku (Post-K) machine has 48 compute cores plus 4 assistant cores and SVE SIMD 512-bit double precision vector extension. Intel processors have advanced 512-bit double precision vector extension. AMD processors support 256-bit vector width. Processor architectures also support single precision and bfloat16 for AI and deep-learning acceleration. Right sizing the vector acceleration, using types optimised to the applications (like 16-bit Bfloat for AI loads) or with variable precision (for solver and high precision arithmetic) will lead to more efficient solutions, but they will also need good software support, for example through optimised libraries. Likewise, matrix multiply units are being incorporated into newer CPUs. A Single instruction can now be used for manipulating an entire matrix, accelerating the capabilities of the current SIMD instructions like AVX-512 and SVE.
- Security: HPC systems are increasingly exposed to attacks, and system security is moving from a mentality of "secure room" to secure system. Processors and systems are required to support increasing security functionality such as encryption and "secure enclaves" (hardware-level isolation and memory encryption that isolates application code and data based on privileges). These are built into CPUs from Intel, AMD and various Arm processor providers. Though not included in classical HPC applications, these will be needed in HPC machines addressing new application domains and connected environments such as federated datacentres or hybrid solutions incorporating HPC "in the cloud".
- Virtualisation: virtualisation relies on hardware features found in all current server CPUs that allow a single processor to act as if it was multiple individual CPUs. In servers, this allows job instances to be isolated to multiple cores, to a single core, or even to support multiple instances on a single core while abstracting the hardware layer to the application layer. Virtualisation is used in cloud computing but it is still relatively new to HPC[6]. Where virtualisation is discussed, it is typically containers, which allow an application to run in its best environment (libraries, APIs) even if the environment is not currently installed in the system. Virtualisation (and paravirtualisation) is a challenge for HPC for two main reasons: added SW reduces performance due to overhead and memory footprint, and applications are typically spread over a large number of cores. Singularity is the most widely used software solution implementing virtualisation for HPC.

Accelerators:

Accelerators such as GPUs have been introduced in HPC based on their high energy efficiency and huge peak performance compared to general purpose processors. Eight of the top ten systems in the Top500 utilize GPUs or some other form of accelerator. Ten of the ten top systems in the Green500 of June 2022 use GPUs. This trend shows that the underlying processing technologies for HPC are trending toward heterogeneity, and many upcoming supercomputers are targeting a mix of accelerators and classical CPUs. With the introduction of data processing accelerators (DPUs - Data Processing Units or sometimes called IPUs Infrastructure Processing Units), these accelerators are not only helping the processor in heavy regular operations, but also in data management. There are also examples of accelerators being pushed to the storage devices[7]. Thus, accelerators can be distributed to various system components and utilised with data inflight. We can expect that more accelerators will be proven in HPC systems and then integrated in various non-HPC systems and placed near the storage, network and general compute components. Due to the new usages of HPC (data intensive, AI), the systems might need to be more open to the outside world, with the related increasing

5. Max Smolaks, "Omdia: Arm-based Server CPUs Experience Record Demand", Dec 13, 2021, https://www.datacenterknowledge.com/hardware/omdia-Arm-based-server-cpus-experience-record-demand
6. It is still early days for containers in HPC. Docker still has issues because it requires sudo access. Singularity is preferred, but this is an open research and devOps question. Lost in performance and memory footprint are another reason why virtualization is not being used in HPC. Furthermore, things like Hyperthreading is disabled to maximize performance and to provide consistent high performance.
7. https://www.anandtech.com/show/16245/xilinx-and-samsung-launch-smartssd-computational-storage-drive

problems of security, access control, etc. These issues will also have to be considered at the hardware level, and accelerators like DPUs and IPUs can ease the use of these new applications in HPC machines. Of course, software must be available to insure efficient use of accelerators. How caches are distributed and intertwined with communication between cores and accelerators has a drastic impact on performance and is an active system architecture topic.

- GPU: Today, GPUs are a mature technology and are a good match for vector or SIMD-kind operations, while CPUs execute most of the scalar and sequential operations. Accelerators like GPUs improves efficiency for given workloads, while this comes at the cost of more complex programming environments and memory management. GPUs are also difficult to optimise compared to CPUs, lacking the tools and transparency to enable easy code tuning. This is a challenge for the programming tools, which should ideally hide the underlying hardware characteristics. Autovectorisation is available on compilers, while pragma-based programming models, such as OpenMP (and derivatives) and OpenACC help portability of programs and help the programmers, which still should have knowledge about the underlying architecture to design efficient codes. DPU or IPU (see "Infrastructure computing" section below) cannot necessarily be considered as custom-purpose accelerator technologies. However, they offload many communication and infrastructure related tasks from the CPU. This may include virtual switching, encryption/decryption of data at rest and in motion, searching and content inspection related regex pattern matching offload (regex processing).
- TPU or more domain specific accelerators: Artificial Intelligence leads to the design of new accelerators to boost performance for this kind of heavy loads. It was pioneered by Google with its line of TPU (Tensor Processing Units) chips (they are on their 4th generation, which are on average 2.5x better performing than the previous generation[8]), followed by Amazon (AWS Inferentia[9]) and Microsoft (using Graphcore accelerators, but planning to develop their own chips as well). They generally handle tensors[10] (hence their name) that are efficient for in Deep Learning algorithms (results of development frameworks such as Tensorflow, developed by Google, are typically run on Tensor processors). Those accelerators also support data types adapted to Deep Learning loads, such as Bfloat (Brain Floating Point), a 16-bit floating point representation with 8 bit exponent bits. We can expect that such domain specific accelerators will be integrated in HPC machines if AI workloads usage increases in the HPC domain. Need will be driven by learning, with top performing neural networks such as "Transformers" require a tremendous amount of computing (3.14×10^{23} Flops for GPT3) and might still increase in the near future, together with the amount of data required for "learning" (nearly 500 billion tokens for GPT3-175B).
- FPGA/CGRA: The heterogeneity of accelerators might also include FPGAs in the future, which are very efficient when dealing with workloads that can be executed in a spatial (parallel) way, with a relatively limited complexity in data representation (operating at the bit level rather than using double precision floating-point representation). CGRA (Coarse Grain Reconfigurable Arrays) are similar to FPGAs but with reconfigurability at coarser grain. They require good software support including HLS (High Level Synthesis) tools that allow configuration of the FPGA with "classical" programming languages such as C, avoiding the need for hardware-oriented languages such as Verilog or VHDL, which eases use for software programmers. However, integration of the SW environment with the rest of the SW tools needs improvement, and it is not clear if FPGAs or new CGRAs will soon emerge as efficient accelerators in the HPC domain.
- Other acceleration technologies: specific accelerators can be integrated in HPC machines for some special purpose accelerations, for example, to speed up randomised algorithms at a very large scale, LightOn's Optical Processing Unit (OPU) uses photonics to compute those specific functions and was integrated in the French "Jean Zay" HPC machine[11]. Lightmatter is proposing optical computing modules that are more energy efficient than GPUs for particular tasks. Adding computing near or in storage (In Memory Computing/IMC or Near Memory Computing/NMC) can also be considered as new architectures for accelerators.

INFRASTRUCTURE IS COMPUTING TOO…

To advance the operation, resource management, disaggregation, security, and effective utilisation of computing, memory and storage of large-scale infrastructures (HPC / Hyperscale computing), Data Processing Unit (DPU) / Infrastructure Processing Unit (IPU) have been introduced, initially in servers and cloud, but they might also appear in the future in HPC systems.

DPU/IPU technology is an evolution of "Smart NIC" technology, which combines traditional Network Interface Controller (NIC) with Field Programmable Gate Array (FPGA) technology on a single NIC card. The FPGA technology is used to extend the NIC's features for implementing data-path accelerators, such as cryptographic functions, Open Virtual Switching (OVS), packet header processing, packet inspection, load balancing and packet steering.

The increased demand for infrastructure related custom-purpose processing and storage disaggregation led to the use of

8. Naveen Kumar, "Google breaks AI performance records in MLPerf with world's fastest training supercomputer", 2020, https://cloud.google.com/blog/products/ai-machine-learning/google-breaks-ai-performance-records-in-mlperf-with-worlds-fastest- training-supercomputer
9. Sébastien Stormacq, "Majority of Alexa Now Running on Faster, More Cost-Effective Amazon EC2 Inf1 Instances", 2020, https://aws.amazon.com/fr/blogs/aws/majority-of-alexa-now-running-on-faster-more-cost- effective-amazon-ec2-inf1-instances/
10. A rank 0 tensor is a scalar, a rank 1 is a vector, a rank 2 is a matrix, etc.
11. https://www.hpcwire.com/off-the-wire/lighton-photonic-co-processor-integrated-into-european-ai-supercomputer/

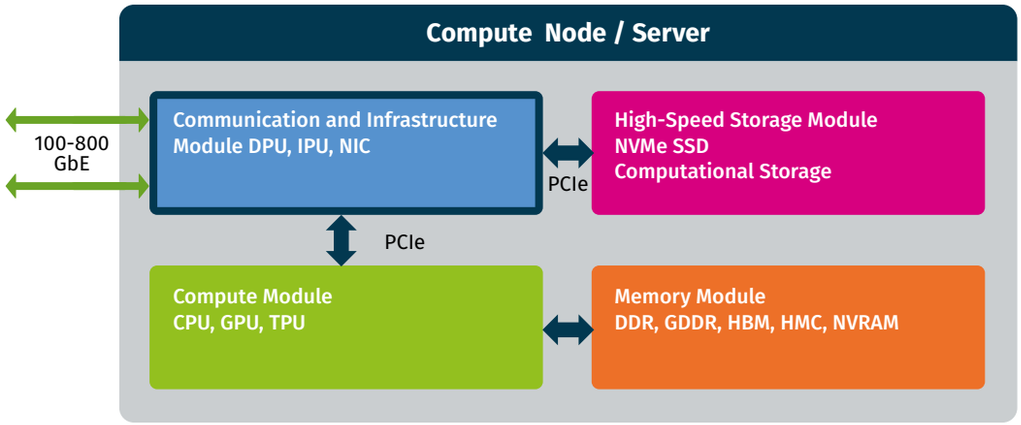

Figure 13: DPU/IPC in a compute node

embedded general-purpose and custom-purpose processor cores, including high-performance Arm-based CPU cores and micro-engines.

Figure 13 illustrates the emerging role of DPU technologies within Hyperscale Computing and HPC compute nodes.

STATE-OF-THE-ART IPU/DPU TECHNOLOGIES:

First introduced by FPGA-based Smart NICs, general infrastructure related functions have been mostly implemented in the form of programmable network-data-path and inline/lookaside offload accelerators. The focus has been the increase of programmability of accelerated data-path processing and Arm-based general-purpose processing. Various DPU architectures offer network path programmability, RDMA, virtualisation support, storage disaggregation, and range of security offloads. Several commercial products are available, such as Broadcom Stingray DPU, Intel IPU, Fungible DPU, Marvell Octeon 10 DPU, Nvidia BlueField 3 DPU, Pensando Elba DPU, Xilinx's Alveo SN1000 SmartNIC.

MEMORY SYSTEMS[12]

Society's production and use of data is exploding. This enormous wealth of data is not sitting idle, but is actively analysed to create models (learning), which are then used to analyse data (inference). Thus, among the factors that may limit future compute, we have to consider memory and storage capacity along with memory bandwidth. Memory bandwidth is particularly critical as domain-specific accelerators are introduced in common applications. To optimise power consumption and avoid thermal issues, future memories are trending to increase bandwidth while also lowering energy per bit.

Datacentres are trending to heterogeneous compute by including accelerator-based compute, with dedicated hardware to offload functions typically done by the CPU. Heterogenous compute is the original demand driver for high performance memory in the datacentre.

From early designs, computer systems have maintained a strict

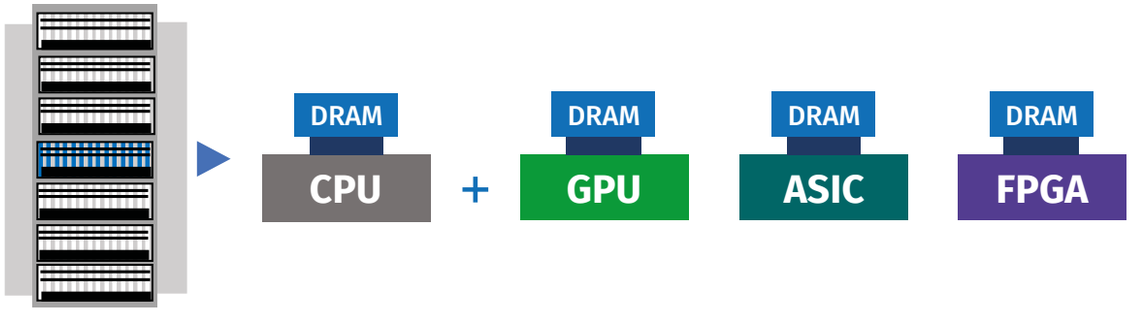

Figure 14: Heterogeneous computing

12. "Data is to this century what oil was to the last one: a driver of growth and change" [The Economist https://www.economist.com/briefing/2017/05/06/data-is-giving-rise-to-a-new-economy]
"DRAM is the oxygen that feeds the AI and VR/AR bonfire, and data from NAND storage the fuel" [Bill Joy, founder of Sun Microsystems]

separation between CPU and memory, forming the von Neumann bottleneck. This bottleneck is especially detrimental to data-centric applications that poorly match the assumptions underlying a modern memory hierarchy, i.e., that have large data volumes, low operational intensity, low spatial and temporal locality, and unpredictable data access patterns. Many of these characteristics arise in important HPC algorithms such as sorting, filtering, hashing, associative memories, pointer-chasing, array and matrix address management, graph operations and time series analysis.

Memory systems are major contributors to the deployment and operational costs of large-scale high-performance computing (HPC) clusters, as well as one of the most important design parameters that significantly affect system performance. For decades, DRAM DIMMs have been the dominant building blocks for main memory systems in server and HPC market. However, this technology is running into scaling limitations, and may not meet the bandwidth and capacity needs of next-generation systems.

Memory bandwidth is one of the main HPC applications performance bottlenecks. In contradiction, processing power of HPC platforms is increasing at a higher pace than the memory bandwidth, causing the ratio between the memory bandwidth and the floating-point performance (Byte per Flop ratio) to decrease. 3D-stacking technology and HBM devices lighten the memory bandwidth bottleneck by providing N-fold higher bandwidths than traditional DIMMs. Currently HBMs are the only viable memory solutions (in terms of acceptable Byte per Flop ratio) for chips with large compute power, such as GPUs and high-end CPUs with vector accelerators. However, although HBM pushes back the memory bandwidth problem, it has limitations. First, they are expensive, with a cost per capacity of roughly 3X DDR. Second, capacity is much lower than traditional memory such as DDR. Thus HBM, where used, will likely be part of a hybrid memory solution[13].

Memory capacity requirements are driven by the algorithms processing large data volumes that cannot be easily partitioned amongst different HPC nodes. These algorithms are widely used in important application domains, especially data analytics (big data and database workloads), machine learning, artificial intelligence, deep learning, and bioinformatics (genome analysis). The community is already trying to meet these capacity requirements with the node-level deployment of storage-class memories (SCM), such as Intel-Micron 3DXP technology. These memories have lower bandwidth and latency performance relative to DRAM DIMMs and HBM, but provide orders of magnitude higher capacity, data persistence and better energy efficiency with zero stand-by power. They are byte-addressable and come in a DIMM form factor enabling easy integration into the existing SW and HW ecosystems.

The idea of heterogeneous memory systems that would bring the best of different memory components is tempting and drives extensive research. Indeed, extending the standard DRAM DIMMs with high-bandwidth low-latency HBMs, and high-capacity, energy-efficient and persistent storage-class memories would provide all the desirable memory-systems features. In these systems, however, good performance requires efficient data allocation and migration between different memory segments. Data management requires profound application profiling, and up to now, no automatic algorithms—whether in the hardware, compiler, or runtime environment—can provide out-of-the-box performance for legacy codes. Instead, efficient use of the advanced memory organisation is still the responsibility of the programmer, which has significant impact on code development cost[14]. Many new workloads are memory-bound on traditional hardware, with a byte/flop ratio of 10 or higher. This requires an increasing bandwidth between processing (including accelerators) and storage.

Caches and scratchpads memories are increasing in size, are integrated “closer” to the CPU or accelerators, and, for example, unified memory architectures allow the CPU and GPU or accelerators to share a common pool of memory. With the GPU and CPU using a common pool of memory, the CPU does not have to transport data - it just tells the GPU to get data from a memory address and do its processing. Once the GPU is done, it releases data back to the CPU, removing the need to transfer the processed data from GPU memory to system memory[15].

	HBM	HBM2 / HBM2E (Current)	HBM3
Max Pin Transfer Rate	1 Gbps	3.2 Gbps	6.4 Gbps
Max Capacity/stack		16GB	24GB
Max Bandwidth		410 GBps	819.2 GBps

13. But, for example, this is not the case in Fujitsu A64FX
14. Newburn 2015; Cantalupo et al. 2015; Jeffers et al. 2016
15. For example, the Apple M1 architecture is an example of unified memory architecture and explains its performances.

The introduction of HBM has been a radical innovation relative to DRAM. The HBM memories, composed of stacked dies, provide high bandwidth relative to standard DDR (SK Hynix announced 819.2 GBps for HBM3) but with a limited capacity (up to 24 GB per stack in HBM3).

DDR5 and GDDR6 memories are available today, while DDR6, GDDR6+ and GDDR7 are under development (targeting 2025). DDR6 speeds will be of around 12,800Mbps on JEDEC modules. But there is often still a compromise between memory capacity and latency.

Byte addressable persistent memories allow new abstraction and, for example, more structures using key-value than address-data. Byte-addressable NVRAM (such as Intel's 3D XPoint) will contribute to the performance increase if used with an optimal data placement into the different memories according to their specificities. This leads to new data management software such as DAOS that is used in the Aurora supercomputer[16].

The community has been working to address the von Neumann bottleneck for decades. For example, high memory latency can often be mitigated using latency-hiding techniques, such as out-of-order execution, branch prediction, data prefetching, up to three levels of cache memory and two levels of main memory, data buffering, and more. These techniques have been successful in improving performance to a degree, but the downside has been increased complexity, high development and per-unit costs, and high power consumption. At this point, additional incremental opportunities to improve performance are close to be exhausted and it is time to look afresh at a different approach to the problem, the tight integration of data and processing. This approach is referred to as Processing in Memory (PIM)[17].

PIM is approaching a tipping point for three reasons. Firstly, it avoids the von Neumann bottleneck, a fundamental limitation to the effective use of computing systems for a large range of important data-centric applications. Secondly, it matches the community's requirement for efficient application acceleration by reducing the amount of expensive data transfers and, for some applications, it already demonstrates orders of magnitude improvements in performance and energy. Thirdly, it has reached a high technological readiness confirmed by various industrial prototypes and products. However, its capabilities in term of complex processing are still limited.

INTRA AND EXTRA NODE INTERCONNECT

In half of the Top10 of the TOP500 systems InfiniBand interconnect is used (generally provided by Mellanox in EDR or HDR versions). Cray has its Aries and now Slingshot interconnects technology and Omni-path - OPA – is also used but Intel will not further develop its OPA200 technology and it will rather develop Ethernet-based technologies following the acquisition of Barefoot Networks. Fujitsu systems use the Tofu interconnect. The Enhanced Data Rate (HDR) InfiniBand frameworks is evolving towards a 400 Gbps version (XDR), an upgrade over the current 200 Gbps, utilising 100Gbps Serdes technology. Europe has its own developments with Atos-Bull's BXI and EXTOLL.

Disaggregation over the network of compute and storage resources drives more bandwidth requirement from the network and requires disaggregation protocols (e.g., Non-Volatile Memory Express over Fabrics, NVMeoF).

"In-network" computing is a new method to accelerate HPC and AI applications by moving portions of the HPC application to be run "in the network". This approach was introduced in InfiniBand EDR and HDR generations and today it is used to accelerate MPI collective operations and distributed AI training workloads. The potential to process the data on the fly in the network will enable new use cases and potentially new HPC acceleration frontiers in the future.

There is a large range of interconnects. Many open and proprietary high-speed server processor interfaces have been introduced in the last decade. The classic PCIe interface is still the main server workhorse interface and continues to increase in bandwidth:

Peripheral Component Interconnect Express, officially abbreviated as PCIe or PCI-e, is a high-speed serial computer expansion bus standard. The performance is defined by its generation and by its number of lanes (the number of simultaneous sending and receiving lines of data). PCI Express bandwidth is increasing more rapidly than ever, while Gen 5 is already being used and Gen 6 is in development. Version 5.0 allows 32 GT/s (×1: 3.94 GB/s, and ×16: 63 GB/s), while the latest version, Version 6.0 performs 64 GT/s (×1: 7.56 GB/s, and ×16: 121 GB/s).

However, multiple open interfaces have been proposed by different consortiums to improve performance for various applications. There have been 4 principal open interfaces for interconnection at the board level (component to component, or component to adapter) CCIX, GenZ, OpenCAPI, CXL, this is trending toward consolidation into one: CXL. CXL (Compute Express Link), an open standard for high-speed CPU to CPU, CPU-to-device and CPU-to-memory connections. CXL is the first server interface that enables a common standard bus across Intel, AMD and Arm platforms. CXL provides an abstracted interface that can enable: (i) the attachment of accelerators to CPU, to other accelerators, and to memory with coherency; (ii) "easy" insertion of non-volatile memories with standard and non-standard interfaces; (iii) compute redistribution in the system across many sub-systems (including memory).

More recently, there are initiatives for open or standardised die-to-die interfaces (see next paragraph for more details on multi-die architectures). In the Open Compute Project (OCP), Open Domain Specific Architecture (ODSA) has the objective of defining readily available, integrated end-to-end SiP (system-in-a-package) level modelling and development flows and tools including functional, electrical, mechanical, thermal, testability,

16. https://www.alcf.anl.gov/events/daos-next-generation-data-management-exascale
17. PIM is a generic term that covers any form of integration of processing capabilities into the memory system, of which there are many approaches; e.g. computation in memory arrays, peripheral circuits, and logic layers and dies integrated into the memory devices.

and manufacturing; common chiplet data exchange formats to share chiplet information between 3rd parties; common chiplet model formats including interfaces; and shared standards for integration of Known Good Die (KGD). Specifically for the die-to-die interface, ODSA has initiatives to address different layers of the interface stack including a new parallel interface standard, Bunch of Wires (BoW).

A new die-to-die interface standard was announced in early 2022, UCIe: an open specification that defines the interconnect between chiplets within a package with the goal of enabling an open chiplet ecosystem and ubiquitous interconnect at the package level. While it is difficult to predict the acceptance and growth of individual standards, UCIe has been founded by some of the largest component, foundry, and ASAT providers. AMD, Arm, Advanced Semiconductor Engineering (ASE), Google Cloud, Intel Corporation, Meta, Microsoft, Qualcomm, Samsung, and TSMC are dedicated to this open industry standard organisation to promote and further develop the technology, and to establish a global ecosystem supporting chiplet design. The value of the UCIe standard is that it can leverage the PCIe and CXL protocol layers (it is not confined to these, but they are included in the first version of the standard) while implementing power efficient phy layers. This offers an open standard for a D2D interface that can readily extend to board-level interfaces.

Open interfaces have the advantage of working across multiple vendors/ISAs/Protocols. There have also been a number of proprietary interfaces created for interconnection between specific compute elements, typically restricted to the compute elements of a single company. This may interconnect between chiplets, external to the package, or both. Because of their restrictive nature, these proprietary interfaces typically provide higher performance than open interfaces. The downside is that they promote vendor lock. Examples of proprietary interfaces include Infinity fabric (AMD), NVLink (Nvidia), and Xe link (Intel).

MULTI-DIE DEVICE ARCHITECTURE

As die geometries shrink and data sets grow, data movement has become a significant contributor to the system energy budget: moving data takes time and energy (orders of magnitude difference between an on-die transfer and transfers between chips on boards, or even worse between boards). Reducing the distance[18] between compute nodes (general purpose processors and accelerators), networking, memory and storage (persistent and non-persistent) will increase efficiency of nodes.

Modularity and composability are also important requirements for modern systems: composable nodes (comprising CPU, accelerators, DDR or HBM memories, persistent storage, network interface, interlinked by (a) (coherent) switch(es)) allow tuning of efficiency across different workloads; composable racks (with variable ratios between compute and storage) support flexible use of system resources across workloads without having to equip each node for the max configuration.

As discussed above, compute gains are increasingly provided by the combination of acceleration and general-purpose processing. Devices can be placed on the same motherboard or even across boards and connected through high speed I/O. As gains due to silicon process technology slow (Dennard scaling[19], Moore's law[20]), processor architectures are targeting other dimensions to increase performance. One trend many suppliers are embracing is chiplet (sometimes called dielet) architectures, with silicon ICs that are co-packaged and interconnected. 2.5D interconnection, the most prevalent chiplet interconnect today, is made through an interposer. 3D interconnection is based on chip stacking. Whatever the interconnect between chiplets, the idea of ganging chiplets in a package provides the possibility to shrink functionality that typically sits on a motherboard or across boards, providing associated gains in size, power consumption, performance (bandwidth and latency), and ultimately in cost as volumes increase. Interposers allow high bandwidth connections between chiplets, with low latency. Photonic interposers are under development in research organisations and could further improve the bandwidth between chips.

Integrating different dies (chiplets) in the same package allows cost reduction (multiple small chiplets have a higher yield than an equivalent solution on a large chip) and facilitate diversity (the same chiplets can be used in different combinations to adapt to different product needs or even different markets, by changing the computing/memory ratio for example). It also reduces the length of interconnect, and therefore also the power lost[21] in connection, and increases the number of wires compared to a PCB, therefore increasing the local bandwidth. Chiplets made from different technologies can be combined to optimise cost vs functionality (e.g. logic, memories and analogue devices for power converters). Lowering overall interconnect power consumption also drives the industry towards packaging the optics close to the electronic interconnect, which is called "Optical co-packaging". This new way to build systems can be done with either VCSEL technology or Silicon photonics.

18. Emerging "compute near or in memory" architectures, where compute is moved to data instead of vice versa, is a good example of this, though this is not mature enough to be used in near-term production systems.
19. Dennard's scaling: This "law" establishes the relationships between a number of physical characteristics when a uniform reduction in size is applied to a standard MOS transistor, including power density (which remains constant), speed and supply voltage. In fact, the improvements resulting from size reduction had already slowed or disappeared completely by 2005.
20. Moore's "law" was, in fact, an observation made by G. Moore, while he was still based at Fairchild and before CMOS device technology was even in production. He observed that the density of components on a chip doubled roughly every year; in 1975, a more complete version of the paper stated that this doubling occurred every three years. This statement is more of a business model than a law, and the only true law derived from the physics of MOS transistors was described by R. Dennard in 1974.
21. Reducing power consumption has the multiplying effect of reducing the cost of cooling, which is considered at the system integration level. Most HW components allow the choice between air, cold plates or direct liquid cooling, or even immersion. The choice of cooling for large compute elements depends on the desired PUE and the desired performance point of the component.

3D stacking, and the use of chiplets and interposers is increasing, together with 2D and 3D packaging and all major silicon providers are using it. AMD was one of the firsts to develops devices using chiplets. For its Rome generation of Epyc processors, AMD uses 8 7nm chiplets (for a 64 cores processor) arranged around a 14nm interconnect and memory interface die. The European Processor Initiative (EPI) is also using the interposer and chiplets approach in their Common Platform. (cf. "Technology sourcing").

Using EMIB[22] and interposers, Intel's Ponte Vecchio chip complex has chiplets that employ five different process nodes of manufacturing across Intel and TSMC, a total of 47 chiplets, and over 100 billion transistors in the aggregate across those tiles.

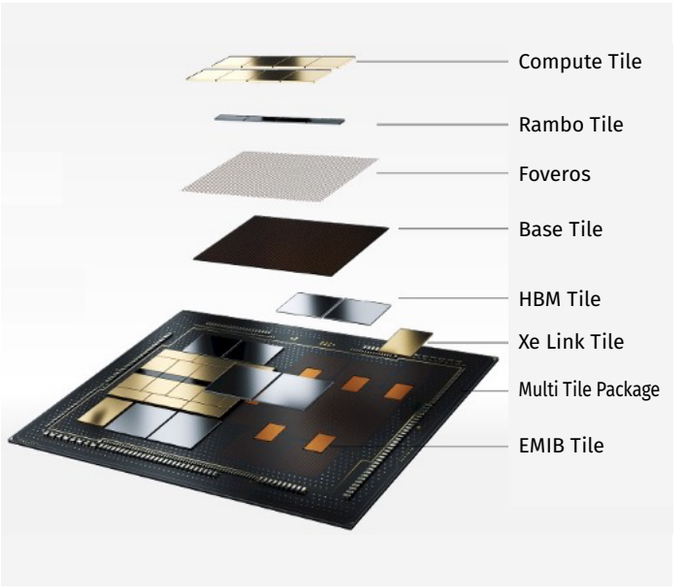

Figure 15: The structure of the Intel Ponte Vecchio chip.

(from https://www.nextplatform.com/2021/08/24/intels-ponte-vecchio-gpu-better-not-be-a-bridge-too-far/)

CHALLENGES FOR 2023-2026

As previously explained, in the coming years, there will be major increases in the amount of data to be processed. The communication capabilities need to support the increasing processing power provided by accelerators and their need for fast data access, otherwise their efficiency will be limited. The solutions to cope with the amount of data transfers and their power consumption are:

1) to limit the length of data transfers such as by unifying memory access of processors and accelerators and integrating them physically near, by using chiplets on (photonic) interposers, with the associated programming models that can transparently take benefit of this heterogeneity of computing elements,
2) to design a better memory hierarchy comprising complementary components: HBM (low latency, high bandwidth), DRAM DIMMs (low latency, mid-size capacity), NVM DIMMs (high capacity, persistency). This will also require software support for clever placement of data in the different categories of memories.
3) to move the computation close to memory. Near-memory processing (digital processing in DRAM periphery or logic layer) is already included in HMC specs, as well as recent products from Samsung (HBM PIM and AxDIMMs) and Hynix (Computational Memory Solution--CMS) and UPMEM (Data Processors - DPUs) .

These solutions will not only improve the power consumption of data movement, but will lead to significant gains in the overall system performance and energy efficiency. Their adoption in production, however, requires coordinated innovations across the whole stack.

In heterogeneous memory systems good performance requires efficient data allocation and migration between different memory components. This requires profound application profiling, ideally supported with automatic algorithms to be deployed in hardware, in compiler, or in runtime environments.

For example, wide adoption of processing in memory depends on our ability to create an ecosystem in which PIM approaches can be designed and evaluated, leading to the selection of potential winners and their adoption by HPC system architects and end users. Creating the PIM ecosystem will require innovations and co-design across the whole stack, from technology, hardware, system, software and programming environment, to updating of algorithms and applications.

Slowing of Moore's law (and the related increase in cost per technology node improvement) is forcing innovation in areas beyond semiconductor fabrication. We are in an expansion phase with innovation in packaging technologies, interconnect, accelerated components, and new compute paradigms such as neuromorphic and quantum, and wafer scale architectures. In addition, the Moore's law slow-down will favour new devices such as Non-Vol-

22. Embedded Multi-Die Interconnect Bridge

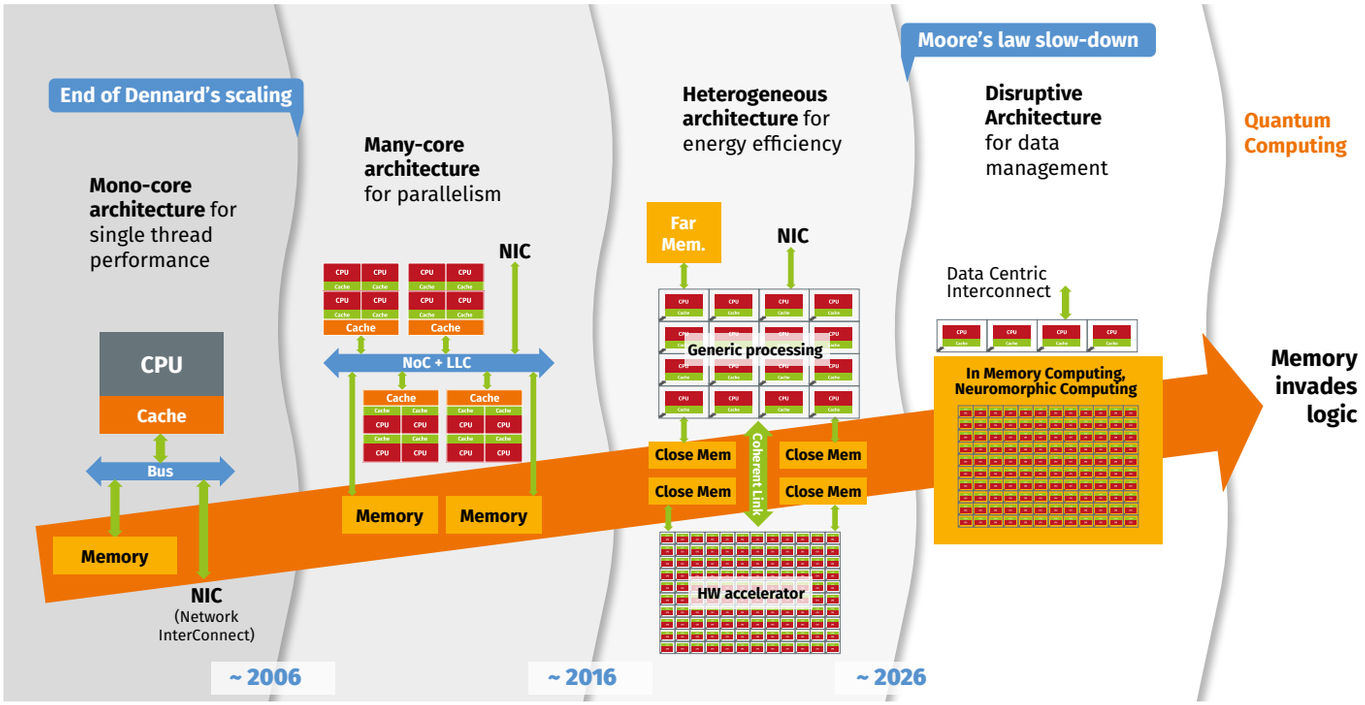

Figure 16: Processor Architecture Evolution since the turn of the century (from D. Dutoit, CEA).

atile Memories, which would also be useful in new computing paradigms such as neuromorphic or quantum computing (see the Quantum for HPC Research Domain).

To benefit from the increase of performance/better efficiency provided by the new hardware components, it is of paramount importance that they are well utilised in the system architecture and well supported by the software and tools.

Computing performance and energy efficiency at component level, as shown in figure 4, are increasingly accomplished through specialisation. Accelerated compute elements, incorporated as blocks, chiplets, or on-board components, increase the compute efficiency for targeted workloads.

SiP architectures are expected to move to a marketplace model, leading to larger and more specialised compute solutions. At the same time, we are seeing an increase in research investment for complex SoC with unified memory (e.g. like in the M1 from Apple), In Memory Computing and Near Memory computing. Either separately or combined, these will drive to a higher ratio of compute to data movement.

Silicon Photonics have been an "emerging technology" for datacentre for a number of years, but have yet to break through as a mainstream technology. The value of a photonic interconnect is low latency for a given distance (vs electrical connections). The challenge is the cost of implementing optical (including acquisition/manufacturing costs, power per bit, reliability). Silicon photonics offer the promise of disaggregated system solutions. Once a signal is converted to optical, the maximum spacing between components that have latency or bandwidth dependencies can increase.

INCREASING COMPUTING PERFORMANCE

Increasing computing performance and energy efficiency will require a complete system (hardware and software) view and will be achieved by combining several axes at the same time, i.e.:

- Increasing the number of cores per chip,
- Using aggressive power saving techniques, dynamic power management, including resource management,
- Supporting the most efficient data type for each type of computing, therefore having hardware with variable precision/formats, adjusting the floating point size. Higher precision computing could accelerate the convergence of algorithms and avoid numerical drift. Other applications, like Deep Learning, efficiently benefit from smaller data representation, such as bfloat16 or event float8,
- Using interposers and chiplets to increase diversity in designs, reduce costs and increase efficiency by closely coupling chiplets made of technologies optimised for each purpose,
- Adding more efficient accelerators and smoothly integrating them in the programming environment,
- Developing 3D stacking to have efficient interconnection between processing and memories, therefore increasing the bandwidth between processing and memory and decreasing the "distance" of the data movements,

- Increasing the (local) bandwidth and having an efficient memory hierarchy (HBM2, DRAM, NVM, etc.). Again, the goal for these new memory structures will be to reduce the data movements. The goal is to have a better ratio compute/bandwidth/storage.
- Support new addressing schemes such as byte addressing and key-value (associative access).

In parallel to these hardware-oriented challenges, the corresponding software challenges will be to make these increases of parallelism, heterogeneity, new memory hierarchy increasingly transparent to the users and efficiently and automatically used by tools (compilers, software stack, e.g., using AI techniques).

To address all these challenges, the HPC community therefore has to:

- Build software development vehicles, prototypes and commercial products that will be used by early adopters to develop and evaluate system software and applications.
- Manage the new hardware components with advanced firmware, OS and system software.
- Develop target applications and kernels with enhanced Application Programming Interfaces (APIs) and programming models supported with the new hardware emulators, compilers, runtimes and libraries.
- Provide users with profiling and analysis tools capable of detecting application phases and kernels that are suitable for offloading to the various accelerators (GPU, TPU, PIM, DPU, ...)
- Evaluate the new hardware components innovations with recognised benchmarks, simulators and HW platforms. The evaluation metrics should include performance, power, energy and total cost of ownership, but also system usability, resilience, security and economic viability, including an analysis of the market size and production cost.

CONVERGENCE (OF HPC, DATACENTRE, "CLOUDIFICATION")

The **convergence (of HPC, datacentre, "cloudification"),** i.e., systems able to support traditional HPC workloads (simulation), Big Data analytics and Deep Learning workloads will be mandatory. That will have several consequences from components specifications point of view, such as increasing the communication/computing ratio, moving processing near data, more modularity and composability, diversification of accelerators techniques, dealing with more diverse types of data; opening to the "continuum" will involve more security concerns for HPC.

Indeed, the peak communication/computing ratio (bytes per flop) should be improved. The objective should be to provide the maximum possible byte/flop ratio from the hardware side and (since it is difficult and expensive) modify the applications to be able to live with a lower byte/flop ratio that they ideally would like to have. Applications should also avoid moving data and keep computation near or in memory.

The move towards distributed systems and locating processing closer to where the data is generated is required to improve performance and efficiency. Communication bandwidth is up to 10 TB/s on a node, while it falls to several hundred of Gb/s at the interconnect level. That requirement will push towards architectures that have accelerators or systems where "Processing In Memory" or "Processing Near Memory" will be realised.

Modularity and composability are also important topics to be more efficient for the various workloads, which need various compute/memory/communication ratios. Dynamicity in this domain is a must but hard to achieve in an efficient way.

Concerning Federated Cloud/HPC & Data infrastructures, the focus should be made on more heterogeneous accelerators and node specifications: different ratio of memories, perhaps more data types (from high precision floating point for simulation, to 16 bit bfloat for AI learning, to 8 bit integer for AI inference). Here also, more flexibility in the system architecture can be achieved by multiple compositions of the same key elements (chiplets) on a common infrastructure (interposer), or by having a range of SoCs differing by the various resources that are embedded according to a particular use case.

There will also be an increasing diversity of accelerators, from GPU to Deep Learning, graph processing and other specialised function-oriented ones, including reconfigurable architectures based on FPGAs or "Coarse Grained Reconfigurable Architectures" (CGRA). This will also have an impact on the data types: the data representations will be more diverse, from bit, to bytes, to integer, to bfloat16, to float, to double precision, and should be dynamically adapted to the current workload. For example, we need to consider completely different floating-point representations such as "POSIT" or "Universal Numbers" (Unums)[23]. In the longer term, accelerators and systems should also support new computing paradigms, where information is not necessarily binary coded, such as neuromorphic accelerators ("spike" coding) and quantum ("qbits").

HPC SYSTEMS WILL BE MORE OPEN TO THE OUTSIDE

New uses (such as interactive accesses) will lead to new cybersecurity threats which will also deeply affect the HPC computer industry by creating a requirement for more robust computer systems. Securing various firmware components inside the platform which ensure secure boot and secure firmware updates will become a major requirement. Special features of the hardware should support protection against attack (secure enclaves, activity monitoring, etc...). Moreover, "encrypt everything" approaches to both system memory and network traffic are emerging. Security devices, such as integrity monitoring, should be encrypted (in hardware, to ensure efficiency) at least at the interface with the HPC centres. Secure boot, and other trusted solutions might also migrate in the core of HPC centres.

23. For the definition of Universal Numbers, see: https://en.wikipedia.org/wiki/Unum_(number_format)#POSIT

RESILIENCE

Resilience is a major requirement for system hardware components. Hot swap and monitoring of health of components are important features to ensure resilient operation, together with architecture features, e.g., redundancy and task migration[24]. Building resilient systems is one of the challenges of HPC in Europe. For this reason, the EU HPC community should carefully allocate the available resources and expertise to address the most important requirements for resilience.

The European HPC resilience initiative (https://www.resilientHPC.eu/) aims to spearhead a Europe-wide discussion on HPC system resilience and to help the European HPC community to define best practices. Its recently published blueprint Towards Resilient EU HPC Systems[25] analyses a wide range of state-of-the-art resilience mechanisms and recommends the most effective approaches to employ in large-scale HPC systems. The guidelines are useful in the allocation of available resources, as well as guiding researchers and research funding towards the enhancement of resilience approaches with the highest priority and utility. The resilience approaches are prioritized based on three main principles:

1. The resilience features implemented in HPC systems should assure that the failure rate of the system is below an acceptable threshold, representative of the technology, system size and target application.
2. Given the high cost of the uncorrected errors, if hardware errors occur frequently, they should be corrected, and corrected at low cost, preferably (if possible) in hardware.
3. Overheating is one of the main causes of unreliable device functioning. Production HPC systems should prevent overheating while balancing power/energy and performance.

The main outcome of this document is the recommendation of the following resilience features as “MUST HAVE” in production HPC systems:

- ECC in main memory,
- Memory demand and patrol scrub,
- Memory address parity protection,
- Error detection in CPU caches and registers,
- Error detection in the intra-node interconnect,
- Packet retry in the intra-node interconnect,
- Reporting corrected errors to the BIOS or OS (system software requirement),
- Memory thermal throttling,
- Dynamic voltage and frequency scaling for CPUs, FPGAs and ASICs,
- Over-temperature shutdown mechanism for FPGAs,
- ECC in FPGA on-chip data memories as well as in configuration memories.

Development and implementation of the remaining state-of-the-art resilience features should only be done based on an additional cost–benefit analysis. These types of analysis depend on the targeted workload and require significantly more work on modelling, simulation and measuring resilience.

This resilience blueprint focuses on the resilience features of HPC nodes, covering the CPU, memory and intra-node interconnect, as well as emerging FPGA-based hardware accelerators. We would strongly motivate the follow-up work to update the analysis and expand scope to include interactions among multiple levels of the software stack and to cover other types of accelerators, as well as networks and storage.

IMPACT OF ARTIFICIAL INTELLIGENCE

Artificial Intelligence will have several impacts on systems components: Firstly, it will drive the development of new hardware accelerators tuned for AI workloads (efficiently supporting tensors, Bfloat data representation, etc), secondly, as it is very data intensive (mainly during the learning phase of Deep Learning Neural Networks), it will drive new system architectures more suited for this kind of data intensive load (see above). And besides inducing new workloads for HPC, it will also be used to design better HPC systems, both software and hardware. For example, it can help in selecting the precision required for computation, or being an “oracle” to predict optimum combinations.

It is especially expected that new AI techniques will help for designing more efficient SoC, by finding the best solutions in exploring the configuration space of architectures, down to having a better routing of chips[26]. Even if this will mainly impact the EDA market and increase productivity of chip makers, this will also decrease the cost for more specific chips for HPC, and then increase the heterogeneity and efficiency of HPC systems.

AI workloads will have an impact on system hardware components in the following ways:

- Compute engines should support a diversity of objects and data types (small matrices, “tensors”) and data types (e.g., 16 bit floating point, fp16, bfloat16). Accelerators are currently developed to improve the efficiency of Deep Learning applications (both during the learning stage – mainly done today on GPUs – and during the inference stage). Neuromorphic accelerators are still in the research phase.
- The memory per node should be increased to reduce the communication overhead due to access to data (storage) during learning phases (Deep Learning).

In the case of “AI workloads”, the growing amount of data will require an increase in the ratio of communication to processing.

24. See for example: https://resilienthpc.eu/sites/default/files/pdf/Blueprint2020__Towards-Resilient-EU-HPC-Systems.pdf
25. "Towards Resilient EU HPC Systems: A Blueprint". European HPC Resilience Initiative. White paper. Online: https://resilientHPC.eu/blueprint2020. April 2020.
26. A graph placement methodology for fast chip design - https://www.nature.com/articles/s41586-021-03544-w

This will have an impact on the memory hierarchy and the local storage size.

AI and HPC workloads are blending together and "in network" acceleration of these workloads is gaining more momentum and will become yet another way to accelerate computing, in spite of the end of Dennard's scaling and the forthcoming end of Moore's law.

OPEN SOURCE HARDWARE

See the dedicated chapter in this SRA.

INTERSECTION WITH "SUSTAINABILITY" RESEARCH CLUSTER

Making semiconductors requires rare earth materials and very pure materials that have a negative ecological impact for their extraction and refining. The main improvements will therefore come from the process technologies and foundries, which should offer technologies that are more sustainable and use less materials. The fabrication process should also limit its energy and water consumption. Recycling is also a key element, but it also requires a lot of energy, and a complete life cycle analysis of components should highlight where the focus should be in terms of improving sustainability. Sustainability and usability cost of building and shipping components is evaluated today, but this is an area that is expected to advance over the next decade, partly due to sovereignty and geopolitical issues which is expected to redistribute manufacturing of components back to a worldwide scope, and partly due to the widening social awareness of carbon expenditure.

From an architectural point of view, energy efficiency and efficient use of the hardware components in the system are the main contribution areas to sustainability and usability. At the component level, several of the phenomena discussed in this chapter will contribute to sustainability and usability, including increasing compute density, distributed processing through specialisation with accelerated compute elements, and with increasing the lifetime of the components (though this is principally a board level concern rather than just component level).

Increasing compute density takes several forms. The first is the proliferation of accelerated compute components. GPUs, specialised AI processors, DPU/IPU all offer a high performance/watt for the workloads that they offload. As long as systems are well architected (SW and HW) to get a high use of installed accelerators, they will see efficiency gains. Chiplet architectures offer another avenue for higher performance density. Collapsing elements from the board level to the chip level lowers the power needed for interconnects and helps to reduce overall system size. Chiplets can also mix processor types, furthering the range of efficient processing solutions. Furthermore, chiplets reduce the cost of redesign, since specific elements can be reused in multiple generations while other chiplets are upgraded. Finally, the individual chiplets can be designed in different process nodes, optimising design and design cost for the various blocks. There is a potential downside in yield losses, since chiplet solutions can be forced to discard the entire SiP due to one or few flawed components, but this is expected to be counteracted by improvements in Known Good Die (KGD) manufacturing techniques, which will be required to implement cost effective chiplet marketplaces.

Disaggregated architectures are also expected to add to usability and sustainability over the next decade, since pooling components will alleviate component redundancy in the system. While this is primarily a system-level phenomenon, component development will be instrumental in enabling disaggregation, such as inclusion of flexible interfaces such as CXL, and the addition of silicon photonics in chiplet architectures.

Other component innovations, such as in memory computing, will improve memory usage by lowering the cost of data movement, and the associated energy.

6.2.3

System software and management

RESEARCH TRENDS AND CURRENT STATE OF THE ART

Power efficiency, scalability, and heterogeneity support still drive the system software landscape at Exascale for 2023-2026. The convergence of Simulation (HPC), Big Data (HPDA) and AI in the same IT continuum can be represented by means of an analogy with an *hourglass*, with heterogeneous HPC/HPDA/AI applications (and their corresponding software stacks) as the top part, and heterogeneous HPC/Cloud infrastructures as the bottom part. The system software resides in the waist of the hourglass, representing the meeting point among applications and infrastructure, and enabling their convergence.

The combination of multiple workloads coming from AI and data analytics and conventional scientific HPC applications on the same supercomputing infrastructure is now a reality that is reinforced by a new criterion: the carbon footprint, requiring to weigh the carbon emissions due to the production of digital equipment against the carbon emissions of their exploitation.

First studies of the carbon footprint of computing infrastructure and applications computing practice have put emphasis on the sustainability of solutions and on the critical impact of data footprint. Requirements listed in SRA4 remain valid[1]. System software solutions need to:

- Support applications diversity and an expanding computing scope, in particular meeting the requirements of the convergence of Simulation (HPC), Big Data (HPDA) and AI processing as part of the same IT continuum.
- Master the complexity of optimised or specialised hardware and software combinations, with environments and tools supporting dynamic and flexible execution models[2].
- Offer smart tools to assist in the development, deployment, optimisation and control of the efficiency of application workflows over heterogeneous hardware architectures..
- Support mixed generations of hardware and their interoperability. This need will become a challenge for runtimes, and applications portability.
- Requirements due to carbon footprint must be reinforced, including investigating the carbon footprint of data (research leading to refinement of data organisation).

Simulations often require large amount of computations, are often run on a general-applicability HPC infrastructure, built as a cluster of powerful high-end machines, interlinked with high-bandwidth low-latency networks. The compute cluster is commonly augmented with hardware accelerators (co-processors, GPUs or FPGAs) and a large-capacity and fast parallel file system. All equipment and services are set up and tuned by systems administrators. In big-data analytics, the focus is rather on the storage and access to data and data processing is often performed on Big Data infrastructure customized for the problem at hand. Those infrastructures offer specific data stores and are often installed in a more or less self-service way on a public or private Cloud, typically built on top of commodity hardware (although hardware customization is becoming more popular). It is our understanding that the HPC world realises that there is more to data storage than just files and that self-service ideas i.e. minimising system administrator intervention for resource provisioning (e.g. by offering advanced policy-driven resource provisioning functionality through a HPC gateway/portal) are attractive to users. In the meantime, the big-data world realises that co-processors and fast networks can really speed up analytics. All Cloud providers now offer HPC services and currently HPC centres are looking to add Cloud-based technologies to their offerings. In this setting, we are considering convergence of HPC/HPDA/AI workloads as a major milestone in the evolution of system software infrastructure and tools.

Additional technical challenges have emerged. It is becoming more pressing to consider the exposure of supercomputers to security risks. The convergence of workload HPC/HPDA/AI and the development of hybrid computing (Cloud +supercomputer, edge+ supercomputer) exposes supercomputers to very serious security risks, which cannot be ignored.

The computing continuum and the sharing of supercomputing resources by new types of applications are critical points for security, with more entry points to verified computing resources. The flexibility of applications containers execution could be also a way to execute verified applications. The software stacks which run on computer resources are not strictly built and maintained by supercomputer administrators anymore, but rather must become more evolvable and composable. We do not have today a process to verify through a “software factory” process that a CVE (i.e. a publicly disclosed cybersecurity flaw, listed in the Common Vulnerabilities and Exposures database) is indeed mitigated.

Furthermore, the integration of Quantum Processing Units (QPUs) into the existing HPC landscape is a challenge for the system software of future HPC systems. Only a commonly shared Resource Management (RM) and scheduling infrastructure with the existing von Neumann-based HPC enables usage scenarios going beyond pure analogue or digital quantum processing.

Such quantum-hybrid HPC codes demand low-latency data exchange between the QPU and the traditional part of the HPC system. The Modular Supercomputing Architecture (MSA[3]) provides a natural way of realising this kind of tight integration by considering the QPU as an additional module of the system. However, with QPUs still being scarce novel resource management

1. cf. the “Sustainability” Research Cluster
2. cf. the “Heterogenous High-Performance Computing” Research Custer
3. Suarez, E., Eicker, N., & Lippert, T. (2019). Modular Supercomputing Architecture: From Idea to Production. Contemporary High Performance Computing. [https://juser.fz-juelich.de/record/862856]

approaches are needed to enable their efficient utilisation by multiple users. Additionally, the system software has to be prepared for supporting even tighter integration models, e.g., as quantum technology evolves, QPUs might even be directly integrated into the nodes and finally on the chip.

CHALLENGES FOR 2023-2026

CONVERGENCE OF SIMULATION (HPC), BIG DATA (HPDA) AND AI IN THE SAME IT CONTINUUM

Converged HPC/HPDA/AI workloads have different characteristics from pure HPC workloads and therefore demand additional features on the systems they run on. With compute intensive workloads currently running on HPC clusters, primary concerns are close-to-the-metal performance and efficient use of high-end/dedicated hardware in an environment where users are granted exclusive, albeit time-limited, access to resources. With data intensive (Big Data, AI) workloads, currently running mostly on Cloud systems, primary concerns are the instant and elastic availability of resources and fault tolerance, in a multi-tenant environment and with sufficient flexibility to select between a "self-service" operating mode or rely on a ready-made software stack. Converged systems should be able to cope with wide-ranging workload diversity and AI-inspired solutions could be used for efficiently managing the complexity introduced. It is important to be able to reconfigure data centre resources, building new virtual infrastructures out of existing building blocks dynamically: elastic reconfiguration and efficient scheduling are potential solutions. New technologies such as AI methods and AI-optimised hardware or commodity-device observations and IoT data streaming offer new potential for scientific methodologies and workflows. Some of the toughest technical challenges will depend on understanding and modelling the data and workflows in the underlying multi-owner, and multi-tenant IT infrastructure. The data and computing continuum is a disruption for present application development putting a major emphasis on a data-aware execution flow and security across the full application workflow.

HPC systems and particularly converged HPC/HPDA/AI systems will need to provide system software support for advanced security mechanisms to satisfy increasingly pressing requirements for trustworthy processing of sensitive datasets. This support includes assistance from OS for putting in place several essential security mechanisms enabling isolation and verifiably trusted execution. Particular concerns arise from the need to process sensitive data sets (such as Personally Identifiable Information and data under intellectual property restrictions) , necessitating protection against loss of integrity and confidentiality.

At the architecture and system levels, computing together with data concerns will drive supercomputer design. Near-memory computing, including processing-in-memory (PIM), integrated at different locations of the architecture will have a great impact, and necessitates advances in system software support. Kernel-bypass networking is a promising technique to address the throughput and system stack scalability issues but we still lack standardised interfaces and protocol implementations for applications to benefit from them. Programmable packet processing is an emerging technique for applications that can offload their request processing in part or in full to the NIC. OS disaggregation will offer new capabilities for applications and runtime-specific optimisations.

Even if the operating system and hardware and smart I/O services are expected to enhance data access and data sharing, major efficiency improvements will depend on application redesign to minimise data movement over the multiple steps of a computation or the multiple steps of the workflow. The real challenge of convergence appears to lie in integrating flexibility with heterogeneity. System architects and application programmers need to re-think the way that information is accessed, shared and stored. A more flexible science-technology co-design flow with fast turn-around of innovation at the interface between science applications, engineering and computational science is clearly needed. Adaptation can be very intrusive, influencing the way in which applications are designed and deviating from the classic first-principles-driven science and linear-workflow approach.

Support capabilities such as workflow and dataflow deployment and orchestration, data location and logistics and dynamic resource allocation (compute, network, storage) are complex issues that will require progressive enhancements in order to address major challenges such as security, efficiency, programmability and reproducibility.

Resource management is challenging for any computing setup. However, in a computing continuum environment[4,5], where edge devices, heterogeneous nodes and both cloud and HPC resources are integrated, this management involves additional difficulties. Indeed, the main challenges to face are:

- There is a lack of a global architecture of the complete environment and scheduling decisions that should be done in a distributed manner.
- The wide variety of compute, storage and communication systems that take part of a computing continuum infrastructure makes more complex the efficient and secure integration and management of these resources. Interoperability is also an issue for the seamless integration of HPC and cloud technologies[6].
- Edge devices and intermediate nodes (fog computing) could be limited in processing and memory resources, and these limitations affect the appropriate allocation of resources to the different parts of the computing continuum infrastructure.
- It is vital to enable local computing on the edge, minimising the data movement between layers and the energy consumption as far as possible, but keeping adequate quality of service.

4. "Towards Integrated Hardware/Software Ecosystems for the Edge-Cloud-HPC Continuum", G. Antoniu, P. Valduriez, H-C Hoppe, J. Krügger, ETP4HPC White paper, 2021.
5. cf. the "HPC in the Digital Continuum" Research Cluster
6. cf. The "Federated HPC, Cloud and Data Infractrucrture Research Cluster"

- Workflows should be deployed and/or migrated between different layers of the computing continuum environment, and the scheduling and orchestration systems have to map the workflows adequately in the complex infrastructure, minimising the energy used and optimising the performance.

EFFICIENCY OF THE COMBINATION OF IT INFRASTRUCTURE AND APPLICATIONS EXECUTION ENVIRONMENTS

To minimise power consumption, the hardware will be increasingly heterogeneous, and it will use processors for computation and orchestration of the dataflows and diverse accelerators, such as FPGAs and GPUs or their derivatives for Deep Learning. At the global level, electricity power sizing and its associated cost ultimately limit the size of the machine: higher energy efficiency allows a more powerful system for the same cost. There are several options to increase the efficiency of the machine and in practice, they should all be combined:

- "Adequate/ appropriate" computing - The idea is to adapt the accuracy of the operation to the needs. For example, the learning process in Deep Learning does not really need double precision floating point operations and GPUs are directly supporting half-precision (float16), which is enough while decreasing the size and energy required. Some operations do not even need to be exact, so operators can be simplified while being "good enough" for the requirements. On the other hand, floating point representation can induce errors in iterative computing and new formats (e.g. UNUM) can help in solving effects such as numerical instability. A similar approach (called "model quantisation") is being applied in the AI domain.
- Application specific hardware is more efficient in terms of FLOPS/Watt or Ops/Watt than the general purpose option because computing resources are tuned to the application class and their control is more efficient. For example, in terms of throughput, GPUs are more efficient than general purpose processors, yet their compute capabilities are more limited (SIMD instead of MIMD execution) and, as a result, programming can be more difficult in the general case. Reprogramming capabilities or dedicated software optimisation will need to be designed. It will be a challenge not only from the programmer's point of view but also from the point of view of system managers to combine different accelerators into a unified programming model supported by a dynamic and elastic resource management infrastructure.
- In situ/in transit processing - In situ processing is a more efficient alternative, allowing data visualisation, curation, structuring or analysis to happen online, i.e. as data is generated by the simulations, thus reducing the volume of refined data to be stored and in consequence saving energy. Big Data management approaches include in situ processing capabilities that are of particular interest for addressing this challenge, i.e. by bringing the computation to where data is located. With ever increasing emphasis on data processing in converged HPC/HPDA/AI infrastructures, there is a pressing need to improve data streaming support at the network and OS levels.
- Decarbonation is an emerging topic which puts emphasis on sustainability of hardware and software for HPC. Providers of hardware, software, or data will need to publish the carbon footprint of production , execution, and recycling for each product offering. For that we need tools to be able to fix the carbon footprint, and we need also to rethink the life requirements, and the Carbon ROI of solutions. The slowdown in the growth of CPU clock speed, and corresponding core performance, is likely to incentivise extending the useful life of HPC system deployments. Operating different generations of hardware may bring interoperability challenges, which need to be balanced with the potential savings from avoiding, or at least deferring, the effort of rewriting application codes for performance tuning.
- Data-aware system policy: Heterogeneity of computing resources and the large amount of data required by applications have a very bad impact on the locality data/processing capabilities. The energy cost, combined with the carbon footprint, of data movement is a great challenge that needs to be considered in the co-design of applications and in workflow orchestration.

AVAILABILITY OF TOOLS FOR DYNAMIC AND FLEXIBLE EXECUTION MODELS

New software stack integration and compliance capabilities will be necessary to support applications portability over heterogeneous infrastructures. Software-defined infrastructure solutions offer advanced policy-driven data and resource management capabilities for managing storage and compute resources respectively. The current state-of-the-art does not yet adequately cover the evolution of HPC environments towards the goal of prompt decision support[7]. Regardless of whether the computing is HPC or Big Data, launching jobs with high resource requirements will require efficient support to reduce job launch latency, monitor job progress and resource consumption in real time and handle runtime nodes and other failures.

We need to consider the technical implications of flexible usage modes of HPC infrastructure, scientific modelling software, analytics and AI/ML-based applications in combination with data assets for the purposes of decision making support. For prompt, or even urgent, decision support, the conventional usage model of HPC centres, which relies on relatively long-term arrangements to enable time-scheduled use to shared resources based on well-established access policies, is not adequate. In urgent decision support scenarios, stakeholders distinct from the conventional users of HPC centres, particularly the members of incident response teams, need to interact with HPC workflows on a demand-based basis, i.e. have the capability to initiate and control the actual processing schedule, while engaging signifi-

7. "HPC for urgent decision making", ETP4HPC Whitepaper, 2022.

cant HPC and data processing resources (both hosted in a HPC centre, but possibly external as well) based on their real-time judgement on how a complex situation evolves over time.

In particular, applying HPC in a wide range of use-cases for insight extraction and decision support in short time-scales brings about increased dynamicity/diversity in the definition of what is the unit-of-work to be managed, including ensembles of jobs, data-staging to support workflows, and interactions with services/facilities external to the HPC system. Model-based simulation, the real-time adaptation of interlinked computational models of evolving complex processes and distributed data-driven scientific collaborations necessitate advances towards more interactive access to HPC resources and pre-emption support, in the context of dynamic workflows over large datasets that require consideration of several challenges, with both technical and organisational aspects. Specific technical challenges arise in the following areas: dynamic resource allocation and scheduling, pre-emption techniques with fast job checkpointing and restart, coordination of resource managers, short-notice secure access to federated HPC resources, data-intensive workflow support (including data staging on node-local storage), increased interactivity (including pre-emption and simulation steering). We have outlined a set of short- to medium-term recommendations for R&D actions towards improving the effective use of HPC resources for insight extraction and decision support in short time scales.

Public clouds are attractive for several use cases where users benefit by providing flexibility, scalability, pay-per-use, and removing initial investment for their HPC workloads. However, they are either low performing, expensive with unpredictable costs, or hard to optimise for the increasing diversity and complexity of many domain-specific, data-intensive HPC and AI/DL workloads. On the other hand, private clouds (on-premises or off-premises delivered by a managed services provider with HPC expertise) can be highly customised with legacy integration to provide significant HPC capabilities while enabling cloud end-user experience and agility. Therefore, a trending usage mode for HPC is to shift towards *hybrid HPC*, leveraging cloud-based options to augment on-premises HPC capabilities, i.e. extending traditional on-premises HPC systems with flexible private cloud (off-premises) infrastructure. A data-aware governance policy is essential in such a deployment, with emphasis on data sovereignty and lifecycle, as well as controls for operational costs.

Matching hardware resource capabilities with applications-oriented environments is a great challenge, requiring the evolution of multiple tools. Even if virtualisation abstractions help to deploy applications over multiple architectures, significant tools evolution will be required to address variability of resources. Application development for this level of complex architecture will require new programming APIs, new run time combinations and tools that offer an abstraction layer which hides a part of this complexity and guarantee application portability. Moreover, we need tools to analyse, profile, trace and predict efficiency of flexible execution environments. Embedded AI and analytics methods will be helpful to master the complexity of development and deployment of that new style of applications.

Moreover, AI integration in HPC/HPDA application workflows (including specific libraries, tensor data types, data ingestion, visualisation, continuum processing) increases the pressure for flexible and effective support for integration of diverse platform capabilities in application workflows. Examples include system software integration packages, mixed-precision arithmetic support, workflow and orchestration capabilities, integration of workflow control with resource management, front-end persistent integration services (e.g. for Spark and TensorFlow).

Efficient integration of virtualisation or container approaches would improve the ease of use, efficiency and resilience of systems. Virtualisation/containerisation must, on the applications side, fulfil their requirements of portability and reproducibility by allowing the definition of encapsulated and customised software stacks, the efficient creation/termination of those software environments on-demand, the seamless and efficient use of HPC fabrics (eg. GPUs, Infiniband), and the secure use of the underlying platform (e.g. rootless, sandboxing). On the infrastructure side, virtualisation/containerisation must feature complete isolation of the applications in a multi-tenant environment, allow agile and fine-grain dynamic resource provisioning to orchestrate resource sharing in those environments, and integrate HPC (e.g. gang scheduling, affinity, pre-emption, topology-awareness, checkpoint/restore) and Cloud (e.g. autoscaling, elasticity, migration) scheduling and resource management techniques, while providing fault tolerance, energy efficiency, and scalability.

To allow arbitration between different users and applications in the current resource management tools, some features will need to be re-thought in terms of the global workflow: the allocation rules, data provisioning and dataflow management policy or engine. This will raise new challenges in terms of resiliency, security and reproducibility of simulation.

Reproducibility will be a major challenge in the next decade. Integration in applications workflow of capabilities to capture contextual information and provenance information during application execution with limited scalability or efficiency impact is therefore an important topic. In order to enable the reproducibility, the use of virtualisation will be again essential to provide a portable environment where experiments can be controlled and mimicked.

Research should target mechanisms for adaptive and dynamic scheduling, management and use of heterogeneous system components to achieve energy efficiency and resilience, while meeting application performance requirements. HPC as a Service has had an impact on the management solutions of the Exascale supercomputers, with an emerging trend towards hybrid infrastructures that include both on-premises and off-premises elements. Supercomputers are increasingly expected to be accessible from the Cloud and be compliant with the Cloud in terms of system management criteria and practices, while still keeping

high performance and scalability as major objectives. Significant challenges lie in the coordination of application workflows, resource management and data management cycle, utilising the combination of HPC and Cloud resources. A meta-OS/meta-orchestration approach, which manages both HPC and Cloud orchestrators and seamlessly enables the combined use of HPC and Cloud resources, is essential to reach the required levels of functionality and compliance.

With an increasing evolution of HPC/HPDA/AI infrastructure towards more dynamic and elastic resource provisioning and management, infrastructure owners/operators are increasingly expected to pay attention to alignment with development tools and standards. It is worth stressing that there are still many differences in terms of tools, protocols and philosophy of usage between the HPC and Cloud communities. This creates a gap, which will have to be bridged for the desired convergence between both disciplines to become a reality. As technology progresses, applications and computing infrastructures become more complex and heterogeneous. As such, modern scientific applications are reflected into multi-domain (i.e. AI, HPC, Big-Data) converged workflows, where different processing stages are executed by a combination of diverse frameworks/stacks (e.g., Spark, MPI, CUDA) and infrastructures (Cloud and HPC). Gluing together such diverse ecosystems is not sustainable in the long-term, since even more and more applications benefit from the cross-stack hybridisation; therefore, a more holistic approach is required. This poses significant challenges to the software management layer, with a focus on the Resource and Job Management Software (RJMS), which is responsible for provisioning the computing, storage and networking resources on one side, and for properly scheduling the execution of jobs on the other hand (orchestration). While Cloud computing relies on a large basis of automation tools to provision resources in an elastic manner, Supercomputing systems traditionally rely on batch schedulers where multiple queues allow the management of different execution platforms (e.g., CPUs, CPUs + GPUs, CPUs + FPGAs, etc.) and priorities.

Although current-generation batch schedulers support the convergence with Cloud environments, they are limited by simplistic models used to express resources and the way these resources are taken ahead of time and relinquished. Furthermore, despite their sophisticated heuristic priority functions that allow to guarantee fair access to requested resources by all users, certain limitations still apply. First, these functions are set at the beginning and hardly adapt over time and as workload changes. Second, I/O-awareness is (in general) poorly taken into account or totally missed, leading to degraded performance on modern architectures which embed a growing number of I/O layers (i.e., Burst Buffers, NVRAM). Similarly to I/O heterogeneity, access to heterogeneous nodes and accelerators may suffer the stiffness of traditional job schedulers, in particular when multiple types of nodes are required by the same job.

Moreover, current generation batch schedulers do not provide any workflow-specific mechanism beyond a way to express dependencies among jobs. This can lead to long waiting times incurred by chained jobs. Alternatives can be identified in those strategies that submit large jobs (also referred to as pilot jobs) to the queuing system, acquiring the maximum required resources for the entire runtime. As such, jobs experience a shorter turnaround when compared to chained jobs, but computing resources get higher idle periods. There are challenges to improve RJMS to better support future exascale machines and hybrid workflows. I/O-awareness scheduling is necessary to exploit the full advantages offered by modern I/O and storage hierarchies (which fill the bandwidth gap between the computing nodes and the parallel filesystems). Therefore, improving scheduling algorithms to optimise computing and I/O operations in a coordinated manner are of primary importance (e.g. limiting the number of I/O jobs running concurrently). Furthermore, full utilisation of the available resources can no longer be based solely on simple information provided ahead of time by the users (which tend to overestimate the required resources); the integration of more accurate models is deemed necessary. Machine learning paves the way for extracting the required information in a much more affordable way, allowing the RJMS to rely on a larger knowledge base for making scheduling decisions.

Hierarchical scheduling approaches are seen as means for tackling the scalability problem of applications (higher levels of the hierarchy) while allowing the scheduling algorithm to get advantage of the specificity of the various application domains being part of the workflow at the lower levels. While some notable examples already exist (e.g., Flux), the majority of the currently available RJMS are flat. An additional point can be found in the way workflows are described and presented to the RJMS. Workflows are commonly described as Directed Acyclic Graphs (DAGs), in which vertices correspond to the tasks to be executed and the edges are the control/data dependencies among these tasks. Generally, tasks must be mapped to jobs either using the piloting strategy or the chaining strategy described above. However, DAGs are not always sufficient to express the complexity of modern applications that may require features like feedback loops. Moreover, workflow tools are typically assuming that the underlying infrastructure is always available. In fact, even if they can interface with HPC job schedulers, no optimisation of the resources allocation is made to ensure, for instance, that resources needed at the same time by two different steps in the workflow, are allocated together. In general, HPC-oriented workflow management tools require gaining more expressivity in terms of handling different levels of parallelism (DAG level, in-job MPI level, node-level) and mapping them to underlying computational resources in an efficient way (working around limitations of current job schedulers), while keeping the execution level and the resources description and scheduling separated.

Moreover, looking at a future perspective, the combination and optimisation of different technological solutions (also already existing) will provide the capability of composing the right set of (virtual) resources needed by a certain workflow, leaving the RJMS

the responsibility to find an optimal mapping with the underlying physical ones. In this perspective, also the capability of the RJMS to support an elastic provisioning of such physical resources will allow moving beyond the traditional (old-style) queuing systems, and will set the point for a seamless integration of HPC and Cloud based ecosystems.

The convergence of traditional HPC and HPDA/AI results has more complex usage modes and workloads as a result. Especially for large-scale HPC systems, this necessitates a modularisation of the runtime environment using standardised interfaces (e.g., PMIx) for information exchange between the various components. This will enable the decoupling of key functionalities, e.g, RM, scheduling, and monitoring/accounting. Only this way, the Resource Management (RM) and scheduling infrastructure can be tailored to the requirements of the supercomputing centre by customising and even replacing individual components of the system software stack.

Finally, HPC systems and particularly converged HPC/HPDA/AI systems will need to provide system software support for advanced security mechanisms to satisfy increasingly pressing requirements for trustworthy processing of sensitive datasets. This support includes assistance from OS for putting in place several essential security mechanisms enabling isolation and verifiably trusted execution. Particular concerns arise from the need to process sensitive data sets (such as Personally Identifiable Information and data under intellectual property restrictions), necessitating protection against loss of integrity and confidentiality.

INTERSECTION WITH "SUSTAINABILITY" RESEARCH CLUSTER

Building modular systems aims at extending the lifetime of some system parts exploitation and at recycling. This requires completely new hardware designs in order to make future HPC systems more modular, thus reducing e-waste and allowing for partial upgrades. Accelerator-based architectures can be a way to consider gaining performance while keeping the central parts of the systems. Increasing the lifetime of a system also requires increasing its resilience to faulty hardware or other system failures. This is by no means a new domain; but an increase of the number of components (within the HPC system itself, and all other parts of the workflow), along with a longer lifetime necessitates to review and improve resilience and fallback mechanisms. From a sustainability point of view the major issue for applications is the hardware heterogeneity that requires specific portability efforts and an evolution of the software stack. While the trend towards heterogeneity is certainly helpful in many aspects, it also makes the task of programming these systems and using them efficiently much more complicated. Often, the combination of different programming models is required and selecting suitable technologies for certain tasks or even parts of an algorithm is difficult and could be cost effective. For sustainability purposes, it could be valuable to use virtualisation solutions; container technologies and microservices are particularly interesting to address these issues. Moreover, efficient and timely metrics collection and low-level resource monitoring APIs will continue to be crucial for overall infrastructure effectiveness and efficiency. Examples include augmenting job-level accounting with profiling of power consumption, power-aware job resource allocation via extensions to workload managers, and power control and power saving modes for increasingly heterogeneous resources.

6.2.4

Programming environment

RESEARCH TRENDS AND CURRENT STATE OF THE ART

There is a clear trend to combine numerical computations, large-scale data analytics and AI techniques to improve the results and efficiency of traditional HPC use cases, and to advance new use cases in fields such as autonomous vehicles, digital twins, smart buildings/towns, etc. Such use cases are typically implemented as complex workflows and will require the coordinated use of supercomputers, cloud data centres and edge-processing devices. It becomes critical to effectively support the productive development of scalable, dynamic, efficient and effective high-performance applications across the extended HPC landscape (the Digital Continuum). This results in the following strategic research directions:

- Develop programming models and systems that integrate HPC, AI/ML and data analytics and facilitate hybrid applications such as AI-enabled simulations. Enabling effective application development and deployment at extreme scales requires high-productivity and performance-oriented programming environments that abstract away the details of the hardware and execution environment. This is also true when targeting the expansion of HPC workloads from simulation-centric applications to applications from Cloud, analytics or AI arenas while maintaining high efficiency.
- The adaptation of existing programming models (or the creation of new ones) and their associated runtime systems and compilers requires co-design and efficient interaction with aspects covered by the other SRA domains, such as the system software, scalable underlying data I/O, storage and system architecture.
- There should be strong interoperability throughout the programming environment across the digital continuum, including compiler tools and runtime systems, debuggers and performance tools and system software, linking information from these tools to the programming model and source code.

An aspect of key importance is to establish the acceptance and adoption of the programming models and application programming interfaces (APIs) in new and existing software systems used in industrial and scientific production scenarios. This requires an emphasis on long-term reliable and robust support of the programming models, APIs and the related runtime system functions and tools, including via formal and de facto standardisation.

Modelling and simulation (well established in industrial and scientific computing) benefits from a relatively long past history of HPC use. Big Data analytics ("BDA" that includes learning-based analytics) and Artificial Intelligence (AI, in particular the Deep Learning variant of Machine Learning) both experience very rapid development and deployment of programming frameworks and associated languages, tools and software systems. A convergence of HPC and BDA/AI offers great opportunities and significant potential to identify commonalities in the software stacks[1] (e.g., common storage and compute abstractions) and common approaches to provision "cross-area" programming environment components (e.g. unified data processing techniques and common models for tensors). While the definition of unified APIs may appear as a long-term goal, possible approaches to address this convergence in the shorter term include composability and the definition of interoperability-oriented APIs.

The programming environment should support the development of complex application workflows, including those expected to run across the HPC-related Digital Continuum, in particular to support "HPC in the loop" scenarios, in situ data analysis and visualisation. Such scenarios require support for dynamic resource management of hybrid workflows adapted for high-performance execution as well as support for dynamic resource availability (malleability). Both have implications throughout the programming environment and system software, encompassing the communications library, runtime, and performance/debugging tools. Together, they will enable the coupling of simulation, databases and data streams, data analytics, AI/ML training or inference, and visualisation that interact together in real time. For instance, results from intermediate data analysis steps performed during the execution of a numerical simulation application code should be able to trigger detailed (or refined) further simulation steps. In other scenarios, to improve the quality of decision making based on data analytics, on-demand HPC simulations or AI predictions can be necessary.

CHALLENGES FOR 2023–2026

As has been discussed above, the use of HPC technologies is evolving from dedicated data centres to encompass the Digital Continuum spanning edge-to-fog-to-cloud/HPC centre deployment and workflows involving modelling and simulation, data analytics, and machine-learning/AI components. The HPC programming environment must necessarily follow that evolution. Nevertheless, the key programming environment challenges presented in the previous SRA (reflected in the following four subsections) essentially retain their relevance and importance; those key challenges permeate through the digital computing continuum.

Irrespective of the target deployment of an application in this new HPC scope, there is a crucial need to support the evolution of HPC applications by providing high-productivity and performance-oriented programming environments. Improved productivity for application developers can be addressed by a reduction in the programming complexity through advancements throughout the programming model and system software stack. An approach is to explore the convergence of the (different) programming models and languages traditionally used by the areas of HPC, Big Data analytics and ML/AI. Potential approaches for this include increased intelligence throughout the programming environment and higher-level abstractions allowing separation of core algorithmic issues from implementation and optimisation concerns.

1. https://www.bdva.eu/sites/default/files/bigdata_and_hpc_FINAL_16Nov18.pdf

While scalability is an attribute that is obviously needed for exascale computing, it is actually a technology attribute required across the Digital Continuum. Scalability requires the programming environment to be developed in a co-design activity at multiple levels, including with the developers of the computing systems and digital continuum infrastructure and with the development of highly relevant applications .

Energy efficiency and support for resilience are two technology characteristics that require close collaboration with the system software level so that the application programmer can realise optimal deployment of production codes. The successful adoption by industrial and scientific production codes necessitates the establishment of formal or de-facto standardisation (including interoperability and composability across the communities and technical areas arising in the foreseen digital continuum).

HIGHER-PRODUCTIVITY PROGRAMMING FOR COMPLEX HPC-ENABLED WORKFLOWS RUNNING ACROSS THE DIGITAL CONTINUUM

This topic covers approaches targeting increased productivity of application development, including legacy codes, through complexity reduction and includes in particular research into application/domain frameworks, extreme heterogeneity abstraction, as well as support for dynamic workflows. These should facilitate code readability, maintainability, portability and performance portability. Over the last few decades, the HPC research community has not reached a consensus on metrics for the quantification of coding productivity, and those metrics defined for other areas are unlikely to be directly applicable to high-productivity parallel programming. Given the many factors involved, including those based on subjective perception such as code expressiveness, research is necessary in this field. We expect advancements to be based on solid qualitative analyses along with quantitative metrics.

The separation of algorithmic expression vs. implementation concerns is a key approach to reduce hardware dependencies and address dynamicity, and thus to decrease programming complexity for the application developer. The aim is to provide an abstraction of the computational algorithms (which should typically be hardware-neutral) separated from the underlying system-specific optimisations, including data structures and runtime system implementation. In addition, programming environments will need to to support the generation of portable applications/services (multi-target, containerised, etc.) across an increasingly diverse set of systems and installations, due in particular to the emergence of the computing continuum.

An intrinsic aspect of the high-productivity programming approach is that specific hardware (and related runtime software) features such as accelerators, in-memory/in-storage or near-memory/near-storage processing, heterogeneous memory systems, or even data distribution across network interconnects, would be supported in a way that is vendor neutral and transparent to the application developer. Other aspects of complexity, such as dynamic resource availability, resource sharing, and load balancing should similarly be addressed in a transparent way. Approaches include the use of meta-programming, high-productivity languages (such as Python), directive-based parallel programming, inclusion of parallel algorithms and data structures in standard libraries, and domain-specific languages (DSLs), particularly those built upon a general-purpose framework. Similarly, the approach should facilitate the use of auto-/self-tuning libraries by applications (including for legacy applications).

At the workflow level, there is a need for application-independent dynamic workflow systems, adapted for high-performance execution, that enable the integration of simulation and modelling with data analytics and AI. Such workflows are expected to be composed of HPC simulations, data analytics (at the input, interleaved with computation, or at the output), AI/ML training or inference steps, visualisation and output to persistent storage/databases or to data streams.

EFFECTIVE INTERACTION FROM THE APPLICATION WORKFLOW TO THE UNDERLYING SOFTWARE INFRASTRUCTURE

The effective interaction between applications and the runtime system also requires use of appropriate APIs or directives for the applications (or high-level application environments) to transfer information (application metadata) between the application and the computing system thereby realising computational schemes that best exploit the system for the targeted metrics (performance, energy consumption/capping, etc). For the latter, key aspects include data layout, data movement, dynamic load balancing, resiliency and malleability (the ability to dynamically adapt to, and request, changing resources and application needs).

In complex workflows, combining simulation and real-time data analytics, there are situations when the application must react to certain events, depending on data contents or depending on interactive requests. This typically happens when on-demand simulations are triggered with real-time constraints for rapid decision making. Changes in data characteristics and dynamic changes in the infrastructure create the need to dynamically reassign resources and remap the workflow onto the infrastructure and swiftly reconsider on which resources the different workflow steps are to be executed. In some applications, such as disaster warning and response, this need may appear when parts of the infrastructure suddenly become unavailable. For the targeted distributed, heterogeneous and dynamic compute and data systems, effective workload forecasting and scheduling of the dynamic workloads will be essential at all levels of granularity (from fine grain tasks to coarse grain jobs). The ability for running processes to respond to changes in resource availability requires modifications to the system software, communications library (typically MPI), runtime system and programming model. Fine-grained monitoring of running workflows and reproducibility of results (to a degree adequate for the actual application) are important. This could, for instance, be achieved by extending recent advances in the Cloud Computing field.

Existing workflow models and environments from Cloud and general data centre computing were not developed with effective and efficient support for parallel, HPC-style applications and supercomputers. It will be important to find ways of integrating such applications and systems with existing, commonly accepted ways to define and run workflows. This requires efficient coupling between Cloud-oriented dynamic orchestrators and traditional HPC batch-based resource management systems, as a step towards more integrated software approaches to dynamic resource management across the continuum. It might even be necessary to develop alternatives to the batch-oriented job scheduling in HPC systems, which nevertheless fulfil the high-performance targets of such scheduling. In addition, the handling of persistent objects, potentially distributed across storage and computational systems, is needed to support the future modes of use of application workflows, e.g. object re use across "work sessions".

The design of suitable abstractions at the application level, and programming model development to support those abstractions, is needed in order to enable the runtime system to realise optimisations. The information transfer (from API through to the runtime system) needs to support flexible runtime hierarchies occurring in dynamic and heterogeneous systems, supported by malleable resource management approaches and systems for storage and computation. In addition, favouring the collection and processing of provenance metadata through the API exposed to the users should facilitate effective interaction with the runtime system. The operating systems will need the appropriate hooks for the runtime system to be flexible, efficient and effective. This includes runtime system - operating system collaboration, such as, e.g., proactive data placement (runtime) combined with transparent reactive data movement (OS) to favour near-memory computing.

Potential approaches for high-productivity and performance-oriented programming environments include exploiting increased intelligence arising from machine learning throughout the programming environment and its underlying runtime and system software, to support smart and efficient resource usage.

INTEROPERABILITY, COMPOSABILITY AND STANDARDISATION TO SUPPORT COMPLEX WORKFLOWS

Efforts made to achieve and improve interoperability among programming models, communications libraries (e.g. MPI), and resulting frameworks or environments should be continued with a particular emphasis on establishing migration paths and improving performance for legacy codes in established industrial and scientific HPC application communities.

To efficiently integrate simulations, data analytics and learning, thereby ensuring a high interoperability for data processing, is an important step, which should ultimately lead to the definition of unified APIs for managing data globally across the continuum. Such unified APIs should facilitate the design and implementation of extremely scalable data processing architectures combining traditional Big Data processing (batch- and stream-based) with HPC-inspired data processing (in situ, in transit). Data access raises similar interoperability challenges, as data models and data access APIs are currently very heterogeneous (byte-level, object-level, structured data APIs, etc.). As stream processing gains momentum, byte- level access to storage is needed more and more to support manipulating data items with fine granularity. In addition, AI/ML training or inference steps will need to be integrated.

At the application level, traditional physics-based simulations (typically executed on HPC systems) need to smoothly cooperate with data-driven, learning-based analytics and prediction engines (typically executed on clouds). Programming the workflow at the highest level requires the ability to consistently combine all these components into a unified framework. Composability (the ability to combine multiple programming models or software stacks for a single application with defined rules) will be needed. Where composability involves multiple "components" (including the runtime system), they must cooperate among themselves and with the system software to efficiently exploit the shared physical resources. Composability between programming models/languages and higher-order frameworks is a particular challenge, as is the ability to handle applications aimed for deployment across the edge–fog–cloud–data centre continuum. One relevant aspect for composability and interoperability across such a large digital continuum regards the semantics and the management level of data consistency. For instance, data consistency is often managed at storage level on cloud-based Big Data storage systems, but substantial improvements in performance and energy efficiency are available if data consistency is exposed at the programming level in HPC systems. An important challenge is to reconcile these aspects as HPC becomes a piece in a larger digital continuum.

While existing standards need to be upheld by new developments, filtering innovations into those standards, new standards (formal or de facto) addressing in particular composability and the broadened HPC targets are important to ensure take-up by industrial and scientific production applications.

Finally, complex dynamic workflows combining HPC simulations and analytics are expected to be deployed in a possibly broad hybrid environment across the digital continuum (including edge/fog devices, clouds and supercomputers or a subset of them). In this new context, interoperability (including storage abstractions and processing techniques), composability and standardisation become critical for the design of programming frameworks and of their supporting tools for data storage and processing, computation and analytics across such hybrid infrastructures. This includes unified real-time data processing techniques favouring 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.

PERFORMANCE ANALYTICS, DEBUGGING & PROGRAM CORRECTNESS

With the evolution of computing systems expected to encompass the whole continuum from embedded HPC to HPC clouds and exascale computing, it is clear that the amount of data produced will be enormous, while only few useful tool sets are available to make sense of the data. Thus, the scalability of these tools is a specific concern. While significant progress has been made in the research area of automated data centre monitoring/surveillance/maintenance techniques, there is a need to develop things further for the HPC context (and for HPC within the Digital Continuum). Other performance analysis themes that have been concerns in the past, have increased in importance in the context of extreme scale computing and dynamic application workflows deployed in the digital continuum:

- Intelligent performance and energy analysis tools.
- Mapping of information, relating to debugging state or performance behaviour, across the (multiple layers of) source codes, to present a view in terms of the original application source code in a high-level language.
- Integrated and user-friendly tools allowing users to collect, analyse and trace debugging state or performance behaviour for both system- and application-levels.
- Provision of information relating data access in the application source code to run-time data layout and transfers.

Moreover, debugger technology is needed which can fully support applications that have been developed on and for dynamic, heterogeneous computing systems, using both current and non-conventional programming models, languages and APIs, and deployed on the full range of target systems within the digital computing continuum.

In addition, open challenges remain for program correctness. Current compilers and runtimes already perform many analyses and checks to warn application users about (potential) issues on sequential applications. This kind of support should be extended and enhanced to also cover issues related to the parallel/distributed execution of an application (e.g., data-race detection in parallel applications).

INTERSECTION WITH "SUSTAINABILITY" RESEARCH CLUSTER

R&D on programming environments can support sustainability objectives in three ways: first, enabling applications and workflows to achieve high performance and efficiency on deployed HPC systems directly contributes to reducing the size of systems needed and the energy used by them. Second, coming up with intuitive and easy-to-use programming paradigms and environments will reduce the effort required for creating such efficient applications. Third, keeping these paradigms and the programmer-visible interfaces stable over extended periods of time and across multiple generations of HPC systems reduces the effort to be invested in maintaining applications and workflows, and enables long-term use of systems even with newer applications.

Combining these three directions, focused and sustained R&D effort should be directed towards programming environments that separate the high-level representation of an application, to be created by the code developer driven by a stable high-level system architecture, from the manifold system and component-specific transformations and optimisations required to achieve high performance and efficiency. This field covers high-efficiency and domain-specific languages, as well as static and dynamic compilation, transformation and optimisation technologies transparent to the higher-level languages. A key request from system architecture R&D is the definition of an architecture concept which can remain stable and at the same time covers the plethora of existing (data and compute) processing elements.

In this field, the equivalent of “system modularity” is the composability and interoperability of different, high-level languages and their implementations. Achieving this is of high importance to allow use of emerging processing technologies across a large body of applications and workflows.

6.2.5

IO and storage

RESEARCH TRENDS AND CURRENT STATE OF THE ART

The research trends for HPC storage systems continue to be based on the requirement to integrate classical HPC simulations, big data, and AI in a single environment. Traditional HPC storage systems have been designed to support huge sequential data accesses per process in a parallel I/O environment, while deep learning and machine learning applications require random accesses to the backend storage systems, especially during the training phases. Additionally, most big data and AI applications are built using machine learning frameworks like Spark and Hadoop for big data, and PyTorch, TensorFlow, or Caffe for AI. The efficient use of these frameworks in HPC environments requires solutions such as specialised hardware and tailored storage plugins within these frameworks, e.g., by adapting their security mechanisms to fit the security policies of HPC storage.

The growing heterogeneity of systems and the increasing performance gap between storage and compute pose additional challenges. Accelerators and co-processors are becoming more diverse and especially the number of different GPU-types with subtly different programming models increases the number of non-POSIX storage accesses between the CPU and the GPU memories. Additionally, classical processors are increasing their core count and the wide-spread application of accelerators is driving up computational throughput. As a result, the overall storage bandwidth requirements per node are growing substantially.

These two trends of converging HPC, big data, and AI as well as growing performance requirements have not shown any signs of abating since the previous SRA. The interplay between different storage workloads, the complexity of their access patterns, and the performance requirements of modern systems are therefore stressing traditional HPC storage solutions beyond their limits. Building unified, robust storage frameworks which can efficiently serve all three scenarios under the constraints of ever-increasing bandwidth and capacity requirements therefore remains one of the main tasks within the time frame of SRA 5.

The requirement list is extended by the orthogonal need to couple data from many different sources. Huge scientific experiments, e.g., from the domains of particle and astro-physics, generate data at a speed and volume which mandate distributed storage and processing. The Internet of Things (IoT) requires collecting huge data sets consisting of small individual data chunks and taking decisions in near-real time. These two use cases require very different approaches to adapt the system infrastructure and the connection between HPC centres and the outside world. The full HPC system therefore consists of the HPC data centre, intermediate caching layers, and the storage infrastructure very close to the devices and experiments themselves. These distributed data sources need new concepts to securely couple data from multiple sites without putting a too high burden on the HPC users as also discussed in WG8, which is looking at such distributed infrastructures .

New technologies help to build faster and more scalable storage solutions. The first generation of non-volatile main memory (NVMM) has been recently commercialised and is now being integrated into NVMe SSDs and in the memory bus to provide different performance tiers. NVMM is less expensive than DRAM and much faster than flash-based SSDs, so that the storage hierarchy can be extended to use tape and magnetic disks for huge, cold data, NVMe-devices using either flash or NVMM for hot data, while NVMM inside the memory bus will build an ultra-fast tier, which blurs the boundary between memory and storage.

These new storage technologies can take advantage of new hardware architectures like on-node cache coherency using Compute Express Link (CXL) and Gen-Z technologies, which, again, have to be integrated with HPC storage systems. Additionally, data-processing units (DPUs), which are designed to offload data movement oriented tasks within a data centre, inside smart network interface cards or switches can help to move additional processing burden from the CPU to the network itself. This offload scheme can help to improve throughput and reduce latencies, while also decreasing jitter that impacts HPC simulations.

Next, we provide a retrospective on technologies that we had been identified in previous SRAs that have now gone into production.

RETROSPECTIVE

TECHNOLOGIES COMMERCIALISED FROM SRA3

We had introduced object storage as a key trend in HPC since the early SRAs, especially SRA3. Object storage, which originated from Cloud infrastructures, does away with the hierarchical organisation of data as files and directories and provides a flat namespace on top of which any data model or protocol can be overlaid. Object storage also changes the I/O paradigm such that there are no more strict consistency needs in the face of multiple I/O operations to storage. We are now in a position to report commercial successes on that front after the technology was investigated in R&D projects just after the SRA3 timeframe.

The SAGE project (www.sagestorage.eu) looked at an object store called CORTX Motr (formerly called Mero) that arose out of co-design with the HPC community. Motr was unique in that it had its inception in HPC rather than in the Cloud. The SAGE project ran between 2015-2018 and was followed by the Sage2 project which ended in 2021. The Motr object store is now fully open sourced and there are now products in the market based on Motr[1].

Also, non-volatile main memory (NVMM) technologies which were introduced and discussed in SRA3 were used in the NextGenIO project (http://www.nextgenio.eu/). NVMM is commercially widely available in the market, including form factors suitable to be integrated as main memory.

1. https://www.seagate.com/in/en/products/storage/object-storage-solutions/lyve-drive-rack/

SRA4 TECHNOLOGIES USED IN R&I PROJECTS

The IO-SEA[2] and the ADMIRE[3] projects both started in 2021 and are based on many of the concepts that were discussed in SRA4. Please refer to the "Projects resulting from the previous SRAs" chapter.

In the following, we discuss the challenges in the 2022 - 2026 timeframe in more detail in the area of storage and I/O.

CHALLENGES FOR 2023 - 2026

FEDERATION

HPC data centres and Cloud come ever closer

There are many initiatives and plans in Europe that will eventually bring HPC data centres and Cloud infrastructures a lot closer. As an example, the GAIA-X[4] project that kicked off in late 2020 and aims to create a federated cloud infrastructure for Europe. GAIA-X aims to consolidate Edge Locations, Cloud Service Providers, and HPC data centres - that are ready to play by the rules of the GAIA-X federation. GAIA-X enables different industry verticals (Automotive, Finance, Health Care, etc) to share data seamlessly through the GAIA-X infrastructure layer. Also, the European Commission aims to achieve better federation between Cloud and Edge through SIMPL middleware, which should also be highly relevant for HPC data centres. Furthermore, there are various initiatives such as FENIX[6] that aim to consolidate data repositories (amongst others) into a common federation for use by cutting edge scientific use cases such as the Human Brain Project - enabling the "cloudification" of HPC data infrastructures. We also note that EuroHPC clearly recognises "HPC Federation and Services" as a priority[7]. With HPC and Cloud infrastructures coming together through such federation mechanisms, there are various research topics that need to be addressed such as:

- Identity and Access Management
- Resource and Job Management
- Trust and Security
- Distributed Infrastructure Monitoring
- Usability by End Users

Various applications will run with different IO patterns and Cloud interfaces like the Amazon Storage-API S3 (which includes the Swift interface from OpenStack and the Rados Gateway interface from Ceph[8]) will become more wide-spread. HPC environments therefore need to include new APIs, so that workflows can access the cloud and HPC environments interchangeably. S3, for example, already nicely integrates with applications that will populate caches or working data sets, an action that can be managed via "get" and "put" requests, while a non-optimised use of this interface can easily lead to new challenges. S3 nevertheless has many drawbacks that do not make it the perfect *de facto* interface for HPC. The size of the records, e.g., only fits well if data exchanges are kept small and each transfer can be reduced to a small number of elementary requests. On the other hand, applications capable of producing partial results that will never change later (for example some checkpoint / restart files) could be adapted to work with S3.

Indeed, this challenge intersects with the Research Clusters "HPC in the Digital Continuum", which is discussed earlier in the report.

Distributed Data Management

The Cloud paradigm promises that data can be processed independent of its current location in the best suited data centres. Transferring this paradigm to HPC requires the introduction of federated data and compute management across HPC centre boundaries (including national boundaries) and across multiple industry and academic verticals, while adhering to governance rules for data sharing and to provide wide-spread APIs.

Spreading data across several locations has to consider network technologies. Those considerations include the classical "bandwidth" and "latency" metrics, but they should also focus on the robustness of the network and the power required to perform data transfers. It is necessary to find trade-offs between these conflicting optimization criteria, which drive scheduling and orchestration SW systems. One approach is to use an efficient caching infrastructure, which can be coupled with in-data or near-data processing. This would enable an application to submit a "compute request" to an intelligent and distributed storage system, which pre-processes the requested data at suitable points in the distributed infrastructure and only sends back pre-computed and filtered results. Such an approach directly ties in with the Edge Computing paradigm. This caching infrastructure also has to distinguish between moving data and moving metadata to make data locatable and manageable. Keeping metadata only at one location simplifies the overall storage design, yet it can induce long latencies which reduce storage performance.

An even more challenging aspect is data federation across multiple industry verticals (because of the different data formats, models and rules for data sharing) , which requires widely used interfaces and which has to take data ownership, access permissions, and regulations into account. The federation of infrastructures, e.g., implies that data could reside in countries which may have different rules concerning data sharing and data handling, while at the same time the federation for European users has to follow the European GDPR regulations. Fortunately, HPC can build on experience from Cloud computing, where middle

2. https://iosea-project.eu/
3. https://www.admire-eurohpc.eu/
4. https://www.gaia-x.eu/
5. https://digital-strategy.ec.europa.eu/en/news/simpl-cloud-edge-federations-and-data-spaces-made-simple
6. https://fenix-ri.eu/about-fenix
7. European Commission, "Equipping Europe for world-class High Performance Computing in the next decade", SWD(2020) 179 final, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52020SC0179&rid=9.
8. https://access.redhat.com/documentation/en-us/red_hat_ceph_storage/5/html/object_gateway_guide/the-ceph-object-gateway_rgw

layers sitting between use cases and infrastructure such as the GAIA-X federation services[9] and SIMPL are under development. These frameworks and corresponding APIs have to be integrated with existing HPC data management solutions.

This challenge intersects with the Research Cluster "HPC for Urgent Decision making" where actionable intelligence is needed close to the data sources (aka, within the Edge), which has to be coupled with intelligent data caching mechanisms (and associated processing) within the federation for timely decision making.

Storage infrastructure co-design with DPUs, FPGAs, GPUs, and new processor designs

The huge scale of the market for Cloud Computing is leading to the introduction of new efficient hardware technologies, which bears great potential for take-up in HPC centres. Examples of new hardware technologies with an impact on storage infrastructures are data processing units (DPUs), field-programmable gate arrays (FPGAs), and GPUs[10], while new CPU designs also provide opportunities to increase the efficiency of the HPC storage stack.

The most promising architectural components to increase storage performance, and at the same time improve energy efficiency in HPC centres, are DPUs. They are tightly coupled with networking interfaces and dedicated accelerators which allow to offload tasks from compute nodes and storage servers into smart network interface cards or even smart switches, leveraging multi-core CPUs, GPUs, or FPGA elements.. First research results indicate that the distributed nature of HPC storage can benefit from including DPUs in the storage stack. Especially highest-performance file systems, like ad-hoc file systems, require huge processing capabilities to ensure data security and consistency. The corresponding overheads can lead to interference with the HPC applications when being used in standard architectures, while DPUs promise to relieve CPUs from these tasks. Additionally, DPUs inside smart switches can cache small intermediate results, steer, pre-process, and combine data streams of file systems, and serve as metadata storage. While the potential benefits of such smart networking cards and switches have been (partly) investigated in the last decade, its adaption inside HPC has just become possible by the introduction of newer and faster hardware. One caveat to note is that these cards tend to be Ethernet based, whilst HPC is almost exclusively Infiniband or similar.

GPUs and FPGAs can extend the application of DPUs in two directions. Firstly, FPGAs are an integral part of many DPU architectures and enable hardware acceleration matching the specific requirements of workloads to be deployed on the fly. Secondly, the trend to pre-process data at its source requires the integration of energy-efficient hardware, so that in-storage processing can reduce energy consumption by offering more energy-efficient hard- and software, while at the same time sending less data over the network. The successful integration of GPUs and FPGAs into HPC storage stacks is nevertheless still in its infancy.

There also have been many new developments in the area of processing to support storage infrastructure software within HPC. HPC storage nodes typically run on x86 cores[11]. However, we now start to see the co-design of storage software with ARM-based processors in European projects availing some of the benefits they offer such as energy and cost efficiency. Also, the ARM processor architecture and instruction set is being adapted to work with NVDIMMs, potentially negating an advantage x86 CPUs have. For the longer term, RISC-V is an emerging CPU architecture that the storage and I/O software can explore to co-design with and understand some of the benefits of.

This challenge intersects with the Research Cluster "Heterogeneous High Performance Computing".

Non-Volatile Main Memories (NVMM)

Non-Volatile Main Memories (NVMM) were already discussed in the last SRAs, long before becoming available in the market, and existing projects both in EuroHPC and FETHPC are already adapting NVMM to be used as underlying block devices for parallel file systems and as memory extensions in HPC infrastructures. The first NVMM devices have recently been commercialised, and in the last few years a plethora of research has been devoted to understanding their behaviour and limitations under different access patterns.

NVMM technology is less expensive than DRAM and much faster than flash, which enables the manipulation of larger-than-memory datasets without incurring high I/O performance penalties. Also, the maximum realistic NVMM capacities are significantly larger than DRAM capacities. More importantly, NVMM devices are addressable at the granularity of 64 byte cachelines, which enables direct (and efficient) load/store accesses without moving data to DRAM, thus enabling in-place execution. This opens the door to new kinds of storage systems able to take advantage of NVMM for in-storage processing, thus twice improving data locality: on the one hand, by avoiding data transfers between the storage system and the application; on the other hand, by avoiding data copies between layers within the storage system itself. This scenario is also a perfect fit for ad-hoc storage systems, which can prevent competition for compute resources within the store.

In addition, NVMM also offers the opportunity to be used in the context of resilience. Reliability is ever more important at extreme scales where the number of components is so high that failures are common and not the exception. As data stored in NVMM is non-volatile by definition, it is possible to exploit this characteristic in order to develop new efficient mechanisms for fault tolerance. Depending on the failure model and the type of application and software stack, one can leverage the remaining data from the NVMM after a failure to avoid a global coordinated restart, which usually requires heavy I/O and recomputation. Thus, it is important to make progress in the context of NVMM

9. https://www.gaia-x.eu/what-is-gaia-x/federation-services
10. This WG mainly covers the usage and exploitation of hardware by the I/O stack, middleware and applications – but not the chips and architecture of these
11. Note that x86 processor co-design is more driven by the Cloud Service Providers and not HPC

techniques to support resilience for extreme scale numerical simulations.

Therefore, many questions still need to be addressed in terms of the optimal usage of NVMM in the I/O stack and what kinds of paradigms are needed to best extract value. Their dual mode usage as both block device and extension of memory raises interesting prospects. In the area of file and object storage, they can potentially be used not just as a high-performance burst buffer tier, but also as a pool that stores specific latency and performance sensitive data sets for, e.g., object metadata, while the less latency-sensitive data can be stored elsewhere. It is also possible to map data from lower tiers into the extended address space enabled by NVDIMMs or mix and match different storage stacks working together, including NVMM, NVMe SSDs, and HDDs to offer the best performance for applications.

Hierarchical storage systems that mix and match NVMM along with other pools of storage to provide the required "points" on the performance/cost curve hence continue to be a major trend for upcoming I/O intensive HPC and AI applications.

Last but not the least, RISC-V support for NVMMs is also an important milestone to achieve.

POSIX issues & Object storage

POSIX filesystems are ubiquitous i n HPC systems. They come with a familiar, structured namespace and permission models which are used by a large body of existing code. They also provide consistency and correctness guarantees which support various use cases but impose significant performance constraints on highly parallel, networked environments. Ways to overcome the drawbacks of POSIX have been discussed for a long-time inside the HPC community, while first progress has only been achieved within the last few years.

Within the HPDA and C loud ecosystems, significant scale and performance benefits have been achieved by relaxing the consistency model and using different approaches for namespacing and permissions. Some of the non-POSIX approaches involve objects rather than file storage (eg: DAOS and CORTX Motr - both object stores targeted towards HPC for scratch storage). There has also been a proliferation of POSIX-like filesystems, which provide the structured namespace and APIs that can be used by existing software without changes, without providing the full range of consistency guarantees.

In an HPC environment, various workflows have very different requirements on the storage backend. Some workflows have demonstrated significant gains by reworking their I/O approaches to entirely avoid POSIX interfaces, but the efforts involved can be significant. Libraries, middleware, and standards will be required which applications can use for I/O in a manner that abstracts the underlying technologies and make the benefits of new technology more widely available. Domain-specific adaptor layers will become essential to map the required semantics for different data uses to the behaviour of the underlying technology, to maximise the performance benefit that can be achieved. Many workflow components are not I/O critical, and it is important to provide POSIX-like, or trivially adaptable, interfaces to minimise the migration effort for these parts.

It has to be noted that object storage continues to strongly evolve in the area of HPC. Object storage can be used as a very flexible, highly performant, scalable base layer on top of which any other protocol and data access semantic can be overlaid. This is applicable for HPC along with AI workflows. Object storage with any of the "plug and play" higher level protocols is suitable for both structured and unstructured data storage. While work on them has begun in earnest, many of the required features and properties (QoS, In-Storage compute, Telemetry etc.) still have to be implemented for existing object stores.

Software Infrastructures to work with AI/DL (Artificial Intelligence/Deep Learning) emerge

With the emergence of AI/DL, mechanisms are needed to easily interface the storage system to these workflows, which need to ingest and work with very large volumes of data often with irregular, unpredictable access patterns for learning and inference. Popular AI based frameworks are now available to develop AI/DL applications (e.g., Keras, pyTorch, TensorFlow and Caffe). Interfacing them with advanced HPC storage stacks however poses challenges as they typically work only with legacy POSIX style interfaces. These frameworks need to be adapted to take advantage of latest developments in HPC such as object storage, which will be a very natural substrate for AI/DL applications to work with. Storage connectors to work with various types of backend storage substrates need to be built.

Interoperability with "cold" data (& dealing with experimental data)

An enormous amount of value is tied up in large data sets held by various scientific initiatives and organisations. Much of this data is stored in 'cold' storage which is typically a tape archive or the lower tiers of a hierarchical storage system. The latency of accesses is typically very large. The use of these large data sets in complex HPC workflows is hindered by the ability to access these data sets in a timely manner. There are also cost constraints with respect to data access that will become more pressing if these increasingly larger data sets are stored in a Cloud.

Middleware and libraries are being developed to integrate workflows and complex storage hierarchies, but these are largely focussed on more efficient use of higher/faster tiers (NVRAM, Burst Buffers, etc). The existence of large "cold" datasets provides different kinds of challenges for workflow execution. Some of the associated issues are:

- The workflows need to be adapted to deal with these cold data sets resulting in suboptimal use of resources (energy, compute cycles, etc.) through various pre-staging and related methods.
- Users must know which data they need and they have to be able to express this correctly in a workflow rather than being forced to copying entire data sets back and forth

Data identification, data labelling, data indexing, and data retrieval become challenging as data sets grow larger and older - which will hence become more important for more efficient pre-staging of data to higher tiers. To remain both relevant and usable, data sets need to be managed. The most effective frameworks for working with deep archives of cold data involve domain-specific solutions. Software frameworks and management tools are needed to build domain-specific interfaces to cold data archives for the new data sets that are being generated today. These tools will need to coordinate data identification and scheduling of data movement as needed by workflow components. In conjunction, this will allow more complex, data-dependent workflows to deliver value by using more data, more efficiently. Best practises for handling such cold data also needs to be put in place.

Large scientific experiments are already handling volumes of data in the range of hundreds of Petabytes, which will soon reach Exabytes in the upcoming years. High Energy Physics experiments were among the first, but more and more scientific communities are now seeing similar challenges. To mitigate the rise in storage costs, scientific communities have started projects to increase the use of less expensive storage, e.g., tapes.

An example project in this field is the "Data Carousel"[12] at CERN, which is an orchestration between the workflow management system, the data management system, and the tape services. It enables a bulk processing campaign, with input data resident on tape, to be executed by staging and promptly processing a sliding window of a fraction of the input data onto a buffer disk such that only a small percentage of the data are pinned on disk at any one time. In order to promptly process staged data and to improve turnaround time, a deeper level of integration between the Workflow Management and Data Management systems is needed: data are dynamically transformed and delivered to let computing resources process data on time, decoupling data pre-processing, delivery, and main processing in each workflow and allowing them to run asynchronously[13].

Understandability & Machine Introspection for Storage and I/O

Storage infrastructure failures at Exascale are going to be a norm rather than an exception. It will be important for data centre operators to obtain detailed insights into the storage and I/O subsystem and take preventive actions in case of issues developing. The telemetry subsystem needs to help provide a systematic and detailed view of the state of the system, including storage and I/O which is still not possible with existing mechanisms, for example of manually going through unstructured log messages. Smarter and better handling of overall system telemetry has been described in earlier WG sections. The storage system forms an integral part of this.

Further, there has been recent work on real-time system monitoring both for I/O as well as performance prediction. However, most of the data generated by system monitoring is used offline instead of online. This is because current systems are rigid and not flexible enough to adapt in real-time when the system behaviour changes. It is important to design real-time information about the load of the storage system as well as API points to access real-time I/O counters and interact with them (e.g., current load, expected load). This would give us the ability to dynamically adapt heavy I/O tasks, such as checkpointing, in order to execute them at the best moment to access the storage depending on the load.

INTERSECTION WITH "SUSTAINABILITY" RESEARCH CLUSTER

Sustainability issues are highly relevant to the HPC sector. Most visibly, large HPC facilities use enormous amounts of electricity and operate hardware whose constituent materials and manufacturing processes have significant environmental footprints. On the other hand, a large proportion of HPC computational resources are directly used in the service of activities with a positive environmental impact. HPC-based climate simulations are probably the most visible, but other examples include research into novel materials, development of more energy efficient systems and products, understanding the complex behaviour of renewable energy systems, and earth-system and weather forecasting that are used for decision making. As a result, as a sector HPC may already be environmentally positive in aggregate.

Most visibly, enormous progress has been made to improve the energy efficiency of computing systems due to both architectural improvements and more efficient cooling - albeit driven by commercial rather than sustainability pressures. But specific support should be given to the development of in-situ data processing approaches, and other strategies to reduce data movement. As datasets grow larger, and longer-lived, and storage hierarchies grow deeper and more heterogeneous, attention should be given to keeping 'cold' data on storage media that do not consume power when not being accessed. Reduction of the purchasing costs of solid-state storage, increasing its longevity, and smarter usage of hierarchical storage (tape and disk) in light of energy efficiency needs to be better assessed

Attention should also be paid to overall lifecycle costs, including re-use and recycling of hardware at end-of-life.

12. https://cds.cern.ch/record/2693664?ln=en
13. https://doi.org/10.1051/epjconf/202024504035 https://doi.org/10.1051/epjconf/202125102006

6.2.6

Mathematical methods & algorithms

RESEARCH TRENDS AND CURRENT STATE OF THE ART

The shape of future HPC architectures strongly depends on the evolution of key technologies that are available for integration, e.g. CPUs with an increasing number of cores and wide SIMD pipelines, high-bandwidth memory, or accelerators like GPUs that feature an even more extreme level of parallelism. These trends have a strong impact on the quest for mathematical methods and algorithms that enable efficient exploitation of these HPC architectures. For both CPUs and GPUs, a trend to improved support of lower-precision arithmetics and new instructions supporting tensor operations could be observed. This observation can, e.g., be exploited by mixed-precision algorithms. Also, the integration of different memory technologies leading to more complex memory hierarchies will shape these future HPC architectures and thus a need for better understanding of how to leverage data locality properties and optimise data placement.

While the speed at which performance improvements are achieved through new CMOS technologies is reducing, there are good arguments that suggest that despite this slow-down, significant improvements can be achieved above the CMOS level[1]. In particular, there is plenty of room for improvement at the top of the computing-technology stack, including software, algorithms and hardware architecture. In particular, new mathematical methods and algorithms are important ingredients in ensuring efficient usage of future architectures and technologies, and improved energy efficiency.

One example is the development of algorithms that support mixed-precision calculations. They allow the exploitation of new CPU and GPU technologies that feature a significantly higher throughput of reduced-precision arithmetic operations, e.g. based on floating-point formats like bfloat16, compared to double-precision arithmetics. Significant reductions of both time-to-solution and energy-to-solution could be demonstrated for specific cases using mixed-precision iterative solvers[2]. Integration of multi-precision capabilities into numerical libraries is making progress (see[3] for an overview). Another area of active research is the development of algorithms and implementation of numerical libraries for unconventional computing solutions based on data-flow architectures (typically realised using FPGAs), less attested processing devices, which may require more robust algorithms, and approaches such as in-memory or in-network processing. These solutions typically operate on data in place and thus are potentially much more energy-efficient as unnecessary data movement is avoided.

In the context of this SRA the focus is on the interplay between HPC architectures and technologies on the one hand and mathematical methods and algorithms on the other hand. The proposed research directions are to be seen in the context of a broader effort in strengthening mathematical research in Europe[4].

CHALLENGES FOR 2023 - 2026

ROBUST METHODS AND ALGORITHMS ENABLING EXTREME SCALABILITY

With each new generation of HPC systems we continue to observe further increase in terms of number of floating-point operations per clock cycle. This ever increasing amount of parallelism is exposed at different levels ranging from instruction, device, node to system level. For instance, on modern GPUs, a single instruction is able to perform 2048 16-bit floating-point operations, while full systems like Fugaku comprise almost 160,000 nodes. This trend requires further research on algorithms that can exploit the increasing concurrency and **enable extreme scalability of algorithms**. Several strategies like parallel-in-time integration[5] but also hybrid stochastic/deterministic algorithms[6] continue to be promising.

The costs of many algorithms both in terms of time and energy are driven by the costs of data transport within modern HPC architectures. This ranges from the costs of moving data between different levels of the memory hierarchies to transport of data at system level through a high-performance network. With memory hierarchies becoming more complex and diverse and new network topologies becoming popular (e.g. Dragonfly-like topologies) there is an ongoing need for improved understanding of lower bounds for the amount of data, which needs to be transported for specific algorithmic approaches in addition to computational complexity[7].

Another strategy for reducing communication and synchronisation costs is to introduce asynchronous algorithms. These have been explored for iterative methods showing promising reductions in terms of time-to-solution. These asynchronous algorithms remove the requirement to wait starting a next iteration until all data from a previous iteration has been received[8]. Hence, communication latencies can be hidden and increasingly expensive global synchronisation steps can be removed.

1. C. E. Leiserson et al., “There's plenty of room at the Top: What will drive computer performance after Moore's law?”, Science, 2020 (DOI: 10.1126/science.aam9744)
2. Azam Haidar et al., “Harnessing GPU Tensor Cores for Fast FP16 Arithmetic to Speed up Mixed-Precision Iterative Refinement Solvers”, SC18 proceedings, 2018 (DOI: 10.1109/SC.2018.00050)
3. Ahmad Abdelfattah et al. (ECP project), "A Survey of Numerical Methods Utilizing Mixed Precision Arithmetic", technical report, 2020 (arXiv:2007.06674 [cs.MS])
4. See, e.g., the report Mathematics for Europe (https://ec.europa.eu/futurium/en/system/files/ged/finalreport_maths.pdf) that documents the outcome of a European-wide consulting process.
5. For a list of references related to parallel-in-time integration see https://parallel-in-time.org/references/index.html.
6. Recent work includes Diego Davila et al., “On Monte Carlo Hybrid Methods for Linear Algebra”, ScaIA'16, 2016 (DOI: 10.1109/ScalA.2016.015).
7. For recent work related to parallelization of matrix-matrix multiplications and matrix factorisation see, e.g., Grzegorz Kwasniewski et al., “Red-blue pebbling revisited: near optimal parallel matrix-matrix multiplication”, SC'19, 2019 (DOI: 10.1145/3295500.3356181) Grzegorz Kwasniewski et al., “On the parallel I/O optimality of linear algebra kernels: near-optimal matrix factorizations”, SC'21, 2021 (DOI: 10.1145/3458817.3476167).
8. For a recent evaluation of asynchronous Schwarz solvers on GPUs see Praktik Nayak et al., “Evaluating asynchronous Schwarz solvers on GPUs”, The International Journal of High Performance Computing Applications, 35(3), pp. 226–236, 2021 (DOI: 10.1177/1094342020946814).

Exploiting increasingly more parallel and therefore complex systems will furthermore require **fault tolerance at the algorithmic level**. Also in this context asynchronous solvers are of interest as the time for recovering from faults may be large. Both the increase in the number of components integrated into high-end HPC systems, which may fail, the growing risks of undetected silent errors as well as the use of unreliable computing devices (e.g. noisy in-memory processing technologies) drive the need for enhanced robustness in terms of improved fault tolerance. There is a need for improving error detecting algorithms, error aware algorithms as error oblivious algorithms[9]. Error detecting algorithms are algorithms that can detect errors that are not detected by the underlying system hardware and software layers, e.g. algorithms that allow reconstructing a state before a node failure occurred[10]. Error aware algorithms can apply correction or recovery techniques in case errors are detected, e.g. using error correction codes or restart of iterative solvers based on earlier, uncorrupted results. Finally, error oblivious algorithms can recover from errors without explicit detection of these errors.

Enhanced robustness of algorithms may also be required when leveraging new capabilities of computing devices to **perform floating-point operations with reduced precision** compared to the typically used double-precision arithmetics. While an increasing number of multi-precision algorithms have been developed (see[11] for an overview), more algorithms are needed and the existing ones, which are still at different levels of maturity, need to be improved and made available in production-quality libraries. Furthermore, there is still a lack of quantitative evaluation of the benefits for specific use cases.

ALGORITHMS AND MACHINE LEARNING

Fast developments in data-driven algorithms are leading to new and largely unexplored opportunities for HPC, especially for problems at the exascale, by combining Machine Learning techniques with more traditional workflows. In various scientific and industrial examples, **surrogate models** have shown lots of potential for greatly improving time- and energy-to-solution of computational-intensive simulations, while at the same time maintaining good accuracy. Surrogate models are simplified models that mimic the behaviour of a model, which is expensive to simulate, reasonably good but that are much cheaper to evaluate. In those cases where some prior knowledge is available, refinements such as **physics-informed neural networks** proved to be even more accurate and robust. Both types of models are also relevant in the wider scope of the digital continuum[12].

Despite these initial successes, further developments are needed before being able to expand the range of applications of such algorithms to more complex, real-world scenarios, especially for fields like weather prediction where the physical system presents chaotic behaviour: new training methods should be envisioned to eliminate training instabilities frequently occurring in these models, often related to the multi-scale and multi-physics nature of the problem addressed; new techniques should be developed to automate the search for the most accurate Neural Network architecture; more fundamentally, a better theoretical understanding of the mathematical limitations and capabilities of these new models should be pursued. While weather prediction employs data assimilation for state and parameter estimation, both data assimilation and machine learning are inverse methods based on a common Bayesian framework[13]. Integrating the two approaches to learning from data could leverage their respective advantages[14,15] .

With the increasing importance of these kinds of algorithms, it will also become crucial to create a set of commonly accepted and comprehensive benchmarks and metrics to evaluate them, along the lines of what has traditionally been done for other Artificial Intelligence models.

Many algorithms involve tunable parameters that can have a critical impact on the performance of these algorithms. As exploration of the parameter space to identify optimal parameter choice can easily become rather challenging, efforts started to leverage **machine learning and other automatised optimisation algorithms for optimising algorithmic parameters**[16].

In combinatorial optimisation, performance depends fundamentally on the instance. For a given complexity, in general, problems are NP-complete and often formalised as mixed-integer linear programming (MILP), there are easy and difficult instances for the solvers. This situation presents three major difficulties, which are opportunities for the use of machine learning:

First, the execution time fundamentally depends on the instance and thus can vary a lot. Guaranteed polynomial-time approximation algorithms are an answer to this difficulty. However, guarantees are often difficult to obtain, ad-hoc and the algorithms, by construction, tend to stick to these guarantees, even for easy instances. Even if machine learning does not generally offer the possibility of guaranteeing a priori the quality of the results, it offers the possibility to build algorithms whose complexity can be fixed a priori (albeit at the cost of an unknown approximation error).

9. Emmanuel Agullo et al., “Resiliency in Numerical Algorithm Design for Extreme Scale Simulations”, 2020 (arXiv:2010.13342 [cs.DC])
10. Carlos Pachajoa et al., “A Generic Strategy for Node-Failure Resilience for Certain Iterative Linear Algebra Methods”, IEEE/ACM 10th Workshop on Fault Tolerance for HPC at eXtreme Scale (FTXS), 2020 (DOI: 10.1109/FTXS51974.2020.00010)
11. Ahmad Abdelfattah et al. (ECP project), "A Survey of Numerical Methods Utilizing Mixed Precision Arithmetic", technical report, 2020 (arXiv:2007.06674 [cs.MS])
12. cf. the Digital Continuum Research Cluster
13. A. J. Geer, “Learning earth system models from observations: machine learning or data assimilation?”, Philosophical Transactions of the Royal Society A, 2021 (DOI: 10.1098/rsta.2020.0089)
14. J. Brajard et al. "Combining data assimilation and machine learning to infer unresolved scale parametrization", Philosophical Transactions of the Royal Society A, 2021 (DOI: 10.1098/rsta.2020.0086)
15. Alban Farchi et al., "A comparison of combined data assimilation and machine learning methods for offline and online model error correction", Journal of Computational Science, 2021 (DOI: 10.1016/j.jocs.2021.101468)
16. See, e.g., Harshitha Menon et al., "Auto-tuning Parameter Choices in HPC Applications using Bayesian Optimization", IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2020 (DOI: 10.1109/IPDPS47924.2020.00090)

Second, it is often possible to specialise optimisation algorithms according to a characterisation of the instances. For example, there are specialised algorithms that are very efficient for solving the Euclidean Travelling Salesman Problem compared to the classical Travelling Salesman Problem even if the complexities are the same. Related to this, there are many heuristics for solving problems like the Boolean satisfiability problem (SAT). But the question of choosing the best heuristic (and its hyper-parameters) for a given instance is a difficult and crucial problem. In a general case, identifying which latent parameters[17] of a set of instances make them easy and using that set to solve a hard problem is typically a task for which in some cases machine learning has been shown to be successful. For example, scheduling, logistics in general, robotics, and molecular dynamics are promising fields to exploit these latent features.

Finally, the formulation of the optimisation problems, the optimisation strategies (how to choose variables, constraints and cuts in Branch and Bound (B&B), Branch and Cut (B&C), Cutting and Packaging (C&P) strategies) and the choice of the auxiliary problems are of heuristic nature because the underlying problems are in general of exponential complexity. For these strategies, one relies heavily on experts. Here too, the success of ML and in particular of RL for game strategies such as Go is promising. This topic is also related to the issues around hyper-parameter selection and AutoML. Covariance matrix adaptation evolution strategy (CMA-ES) for numerical continuous optimisation is an example where the adaptation of the algorithm is done at runtime according to the history of the search.

Among the open questions, we can mention (I) the design of networks adapted to these optimisation problems and (II) the problem of validation for problem sizes for which we do not have satisfactory solutions in particular.

SYSTEM RESOURCE MANAGEMENT AND CO-SCHEDULING

As larger and larger systems are available to the industrial and scientific community, new challenges arise in the context of resource orchestration. Actually, this community provides several use-cases that require the definition of custom workflows, each workflow consisting of multiple tasks with specific dependencies, usually modelled through Directed Acyclic Graphs (DAGs). Current HPC architectures' standard scheduling policies do not always allow for the efficient submission of the jobs belonging to complex workflows. Therefore, there is a need for **workflows-aware scheduling policies for HPC resources**. Workflow-awareness means that the scheduler is aware of the relations of the different tasks that belong to a workflow and is able to benefit from this awareness, e.g. in terms of concurrent execution of independent tasks or of minimising data transport needs in case of one task producing data that is consumed by another task. Realising workflow-aware schedulers that are able to achieve efficient utilisation of available resources of HPC systems while adhering to policy constraints strongly relies on suitable algorithms for solving complex optimisation problems.

While optimising the scheduling behaviour can easily be justified by higher throughput, lower energy-to-solution, or more interactive usability, dynamically tunable scheduling mechanisms are crucial to successfully employ HPC for urgent decision making[18], for instance by permitting guaranteed schedule makespans.

Among the main challenges concerning resulting in optimisation problems, we can highlight the following ones:

- The definition of the objective could be an issue by itself, it is a multi-objective problem where the requirements of the interested party should be balanced, one could target minimisation of energy consumption, time-to-solution, data movement, budget as well as maximisation of resources allocation (e.g. memory, CPU usages) and performances. A proper scheduling policy should take into account which relevance to assign to each objective.
- The scheduling policy should also consider that the access to the HPC resources today is managed by resource management systems, such as PBS and SLURM. This is of particular interest in the workflows-focused context, where multiple jobs submissions are associated with a single workflow. Each job, potentially requiring heterogeneous resources[19], goes through the queues before beginning its execution. Two trivial allocation policies could be either to submit one job at a time once all the preceding jobs in the corresponding BAG are completed, thus accumulating queueing time each time, or to submit one single job aggregating the whole workflow, thus requiring a huge amount of resources that probably would not be necessary for all the tasks. This second option increases both resource wasting and the queueing time to provide such an amount of resources. Between these two naïve policies, many other levels of granularity in jobs aggregation could be investigated, thus eliminating some redundant queue time and allocating just the needed resources. Interaction of jobs with the queue system for dynamically expanding or reducing the resource allocation is an advanced mechanism that could increase the malleability of the workflow.
- Another difficulty arises from the considerable stochastic component inherent in the problem, proper scheduling addressing the target objective function should be able to somehow anticipate the consequences of its decisions. In this context, despite some previous investigation through state-of-the-art statistical and ML models, the queue systems modelling still represents a challenge. Reliable predictions could highly improve the efficiency of the scheduling policies, for instance, in the case of the multi-jobs workflows, a good model able to predict the queuing time for each job could be useful to find the optimal time to pre-submit the aggregated tasks in the workflows. Although this is a very difficult task that requires the collection

17. The term "latent parameter" refers to variables that cannot be observed.
18. cf. the Urgent Decision-Making Research Cluster
19. cf. the Heterogeneous HPC Research Cluster

of enough representative datasets and good modelling of the temporal component.

Another type of scheduling problem occurs in the context of deeper memory and hierarchical storage architectures. These architectures integrate different memory and storage technologies to realise tiers with different capacity and performance characteristics. This approach typically allows creating tiers that provide more capacity at a reduced performance or, vice versa, fast memory or storage tiers with reduced capacity. Exploiting such architectures requires **algorithms for data-aware scheduling and smart data placement**. There had recently been various efforts to develop algorithms for scheduling in the context of such architectures, e.g. for scheduling HPC jobs' data on processors with two types of memory[20] or storage systems with two tiers[21]. However, these have only been explored and evaluated for simple algorithms or synthetic workload distributions.

Today the realisation of data-aware scheduling and smart data placement algorithms suffers from the lack of information available about the application's data objects, indeed most of the scheduling policies that deal with data management are focused more on data locality or dimensions than on data I/O needs.

Future data-aware programming frameworks and middleware architectures are expected to be able to make efficient scheduling decisions, taking into account more specific characterisation of the HPC jobs' data. This opens up an avenue for HPC to be at the forefront of solving these scheduling problems, which exist similarly in the wider digital continuum[22].

Performing computations across multiple federated HPC systems and/or federated data storage (cf the corresponding Research Cluster - Federated HPC Cloud Infrastructure), or with workflows spanning HPC, cloud and edge systems creates another level of scheduling complexity, parameterised in a similar fashion, but with widely different parameter values, fault tolerance requirements, and possibly no central algorithmic coordination feasible.

Data consistency within a distributed system is a crucial topic and ensuring consistency can significantly impact performance. **Consensus-based algorithms** are expected to become more important. This includes the Paxos family of protocols, which in the past has been used for distributed storage systems[23]. Consensus-based algorithms are also explored in the context of Cloud computing[24]. In the context of increasingly heterogeneous and more distributed computing systems (cf the Heterogenous HPC Research Cluster), where in the vision of a digital continuum HPC[25] will extend to the edge, it can be expected that finding consensus becomes more critical as consensus needs to be reached among participants that are less reliably connected and where connections may be restricted or affected by large latencies. Moreover, different situations will require matching to the right consensus approach for increased consistency or to avoid performance issues.

ALGORITHMS AND SYSTEM ARCHITECTURE CO-DESIGN

The different types of memory, which are integrated in current HPC systems, contribute significantly to the overall costs of these systems. A **better understanding of how capacity and performance of the memory sub-system relates to performance of algorithms** would help in the co-design of future HPC architectures, in particular in view of heterogeneous HPC systems (cf the Heterogenous HPC Research Cluster). Such analysis can, e.g., contribute to a better understanding of the required memory capacity[26]. Conversely, better understanding of upcoming memory sub-system architectures helps selecting suitable algorithms or their implementations. One example is graph-traversing algorithms, where it has been shown that different approaches feature different performance depending on whether memory latencies are small or memories are optimised for high-bandwidth streaming access[27]. Another example is the use of algorithms that feature memory locality properties that can be exploited in the context of deeper memory hierarchies, e.g. by means of cache-oblivious algorithms[28].

Algorithms can play an important role in **reducing energy costs for solving particular numerical tasks**. There are different opportunities resulting from the properties of the underlying hardware technologies. For instance, costs of moving data are since long known to require more energy than the actual computations. Therefore, algorithms minimising data transfer, often called communication-avoiding algorithms, can help to reduce energy-to-solution. Another opportunity arises from reduced-precision arithmetics, which do not only reduce the amount of transferred data, but also the consumed energy per arithmetic operation. This can be pushed further leveraging new instructions for matrix operations, for which the opportunity for very significant reduction of error costs in combination

20. Anne Benoit et al., "A performance model to execute workflows on high-bandwidth-memory architectures", ICPP 2018 (DOI: 10.1145/3225058.3225110)
21. Leah E. Lackner et al., "Efficient Job Scheduling for Clusters with Shared Tiered Storage", CCGRID 2019 (DOI: 10.1109/CCGRID.2019.00046)
22. cf. the Digital Continuum Research Cluster
23. See e.g., Juan Rao et al., "Using Paxos to build a scalable, consistent, and highly available datastore", Proceedings of the VLDB Endowment, 2011 (DOI: 10.14778/1938545.1938549)
24. See e.g., J. Lim, T. Suh, J. Gil, H. Yu, "Scalable and leaderless Byzantine consensus in cloud computing environments", Inf Syst Front 16, 19–34, 2014 (DOI: 10.1007/s10796-013-9460-7)
25. cf. the Digital Continuum Research Cluster
26. Darko Zivanovic et al., "Main Memory in HPC: Do We Need More or Could We Live with Less?", ACM Transactions on Architecture and Code Optimization, 2017 (DOI: 10.1145/3023362)
27. See e.g., Zhen Peng et al., "GraphPhi: Efficient Parallel Graph Processing on Emerging Throughput-oriented Architectures", PACT'18, 2018 (DOI: 10.1145/3243176.3243205)
28. Neil A. Butcher et al., "Cache Oblivious Strategies to Exploit Multi-Level Memory on Manycore Systems", IEEE/ACM Workshop on Memory Centric High Performance Computing (MCHPC), 2020 (DOI: 10.1109/MCHPC51950.2020.00011)

with mixed-precision algorithms have been demonstrated[29]. Together with new algorithms, there is a need for a quantitative assessment of the benefits for real-life problems. With increasing heterogeneity of the systems[30], algorithms adaptively taking advantage of the different energy usage profiles of the system components will be required to leverage this potential.

There is a growing interest in using HPC systems for processing sensitive data, e.g. in the context of genomics, precision medicine or brain research[31]. This requires the ability to deploy workloads on HPC systems with minimised trust requirements for the end user towards the HPC service provider. This results in the need for **mathematical methods for enhancing security and data protection**. One example is data encryption methods that allow to protect access to the data such that the overhead costs for accessing encrypted data remains small and the scalability of the applications are not affected (see, e.g.[32], and[33] for efforts on using encryption in the context of HPC). While robust cryptographic methods are known, their integration in the infrastructure and the various components of the relevant software stacks is often lacking. In the context of urgent decision making, federation of workloads (cf the corresponding Research Clusters) to HPC service providers will profit from such enhanced security features, as will trustworthy computation on possibly strategically important data.

VERTICAL INTEGRATION AND VALIDATION OF MATHEMATICAL METHODS AND ALGORITHMS

Efforts are required to ensure that any of the developed mathematical methods and algorithms can be efficiently implemented on different suitable types of HPC architectures and can easily be used by a broad user community. Supporting **performance portability of algorithms within emerging diversity of architectures** requires full vertical integration of new mathematical methods and algorithms should ensure that efficient exploitation of the hardware resources is ensured at all relevant levels, from applications to compilers, to middleware software layers and finally the hardware level. This topic is becoming more urgent with the emerging broader diversity of processing technologies used for HPC, not only in general, but even in the same heterogeneous HPC system[34]. The challenge of performance portability across vastly different hardware architectures may require mathematical algorithms to be formulated using suitable higher-level representation from which implementations can be derived according to the system architecture. In the context of the digital continuum[35] with its wide range of hardwares this is of paramount importance.

INTERSECTION WITH "SUSTAINABILITY" RESEARCH CLUSTER

Sustainability of HPC systems is a multifaceted topic, as it encompasses not only the operational aspects, like energy consumption, and resource lifecycle of the components, but also extends into the integration – both as a facility, and as a tool – into a circular economy. These different aspects are associated with different mathematical as well as algorithmic requirements, which we try to highlight in the following.

Transition towards a decarbonised economy will require the **development of new or improved technologies** like new materials for photo-voltaic or improved windmills where the use of HPC is becoming increasingly important. In particular, computational fluid dynamics challenges, e.g. in the context of optimising windmills, require extreme scalability of algorithms. These computations will run for long periods of time and require in parts exascale computing capabilities, which are systems with a level of complexity where the mean time between failures becomes increasingly shorter. Enabling such simulations will therefore strongly benefit from support of fault tolerance at the algorithmic level.

Minimising **energy usage of computations** is the most obvious and often primary aspect discussed when considering sustainability in HPC. Starting from microarchitectural considerations in the hardware design, classic optimisation passes in the compilers, as well as in operating system scheduling have a large influence on the power consumption of the application. Part of the research efforts proposed in this SRA should go into research on reducing energy costs for solving particular numerical tasks. For some classes of algorithms, e.g. those based on iterative solver schemes, an already successfully applied strategy is to reduce power consumption by performing floating-point operations with reduced precision. To facilitate such computations without compromising on the required precision of the final results requires suitable algorithms.

Yet another strategy takes into account that a significant fraction of the energy may be consumed by the different components even if they are not in use. Therefore, **energy efficient scheduling** becomes a prominent and user-visible requirement at the level of concurrent execution of parts of a complex algorithm on heterogeneous hardware with complex memory hierarchies and presence of compute accelerators, or in a distributed compute environment. Here, algorithms for data-aware scheduling and smart data placement can help not only to improve performance but also to reduce energy consumption. Some of the scheduling decisions may well be made at compile time, leading to variants

29. Azam Haidar et al., "Harnessing GPU Tensor Cores for Fast FP16 Arithmetic to Speed up Mixed-Precision Iterative Refinement Solvers", SC18 proceedings, 2018 (DOI: 10.1109/SC.2018.00050)
30. cf. the Heterogenous HPC Research Cluster
31. Michael Schirner et al., ""Brain Modelling as a Service: The Virtual Brain on EBRAINS", 2021 (arXiv:2102.05888)
32. Andre Brinkmann et al., "Secure Genome Processing in Public Cloud and HPC Environments", IWSG, 2017 (http:/ /ceur-ws.org/Vol-2363/paper7.pdf)
33. Leigh Lapworth, ""Parallel encryption of input and output data for HPC applications", The International Journal of High Performance Computing Applications, 2021 (DOI: 10.1177/10943420211016516)
34. cf. the Heterogenous HPC Research Cluster
35. cf. the Digital Continuum Research Cluster

of job shop scheduling problems but more often similar to vehicle routing problems, since data movement and code execution need to be considered simultaneously. More often, though, a large part of the decisions needs to be made online and would profit from middlewares that implement robust online scheduling methods with strong quality guarantees as they are required for workflows-aware scheduling policies for HPC resources.

This pushes the challenge of scheduling individual applications to scheduling problems that arise at system level. But now balancing the distribution of concurrent applications on a large system may directly lead to more even heat distribution and thus lower cooling requirements. Algorithms that can take system-wide or even site-wide monitoring feedback into consideration will be required for such **site-wide energy aware operations**. Only with such abilities will the HPC community systems be enabled to let their **data centres become active energy grid participants**, able to adjust their workload to absorb extra power in the grid, or reduce consumption upon external request. Such adaptivity of workloads has been highlighted as a crucial prerequisite for improving the sustainability of various industries, as they make targeted consumption of renewable energy possible. The HPC industry is no exception here. Coordination algorithms will need to both deal with unit commitment models common in energy system modelling, but also control theory. It has been shown recently[36] that solving such problems at European scale will require and can make successful use of HPC resources itself, using large-scale mixed integer nonlinear optimisation techniques. Building on this, HPC can be used to improve energy network operations and planning, in particular to properly integrate the variability and decentralised nature of renewables, a core requirement of sustainability at a societal level.

HPC also shares all the sustainability concerns of the IT industry in general: **re-purposing, refurbishing, and re-manufacturing equipment** used in the high-end HPC systems when their end-of-service for the initial purpose has been reached can extend the useful lifetime of many components, and reduce resource consumption as well as waste production. Common wisdom is that the largest influence on sustainability in this context can be taken in the design phase: Different system components can be designed for different useful lifespans, and to be independently serviceable or reclaimable, avoiding the swap-all approach to repairs.

Finally, **digital twins** of physical and cyber-physical models, from natural resource usage and tracking, design and manufacturing processes, require HPC technology, algorithms, and are already influencing the design of current and future systems, as the requirements of the platforms for such workflows differ from “classic HPC” systems. We expect to see HPC-based digital twins being used to predict and track economic value flows and track sustainability indices, but also to verify claims by market participants about them, which will require extensive data analysis and simulation.

36. https://www.plan4res.eu/

6.2.7

Application co-design

RESEARCH TRENDS AND CURRENT STATE OF THE ART

Sustaining excellence and European world-leadership in HPC applications is key for European science, industry (incl. SMEs) and the public sector. There is a breadth of applications in the fundamental, applied and social sciences where computing plays a pivotal role[1]:

- Simulations are critical in Climate, Weather, and Earth Sciences. Exascale resources will enable sub-kilometre scale resolutions, a more realistic representation of all Earth-system components, better mathematical models, and larger ensembles of simulations for uncertainty quantification. This will extend the reliability of forecasts to the extent needed for the mitigation and adaptation to climate change at global, regional and national levels, in particular with respect to extreme events. In addition, coupling of mesoscale and microscale models will enhance the accuracy of pollutant propagation in cities, mid-term power generation in wind farms, etc as well as "nowcasting" predictions requiring short term HPC access. In analogy to weather and climate prediction, much enhanced simulation capabilities of solid Earth physics from higher spatial resolution and seismic frequencies down to 10Hz will enable a break-through in the detection and prediction of the precursors of volcanic eruptions and earthquakes, and their impact on infrastructures. High resolution interferometry from satellite data can also help detecting small displacements of natural as well as artificial objects. A prediction capability at this level of detail is crucial for a wide range of societal impact sectors for food and agriculture, energy, water management, natural hazard response and mitigation, and finance and insurance.
- High-end computing capabilities are becoming increasingly important for life sciences, medicine, and bioinformatics and will have tremendous impact, e.g. for enabling personalised medicine. Researchers are already able to rapidly identify genetic disease variants, and it will become possible to identify diseases that are caused by combinations of variants and design treatments tailored both to the patient and state of the disease. Structural biology will increasingly rely on computational tools, allowing researchers to predict how the flexibility and motion of molecules influences function and disease. Deep learning techniques will provide more specific diagnoses and treatment plans than human doctors, making medical imaging one of the largest future computing users. At anatomical scales, organ level simulations present both a challenge and an opportunity: the Virtual Patient modelling for precision medicine. It is the ultimate example of supercomputer usage for multiphysics/multiscale modelling problems. Tumours can be simulated by following the fate of thousands, millions or even billions of cells as their internal circuitry of signalling molecules respond to each other and their environment. This approach can be used to predict the effect of targeted therapies, which is particularly useful to triage potential combination therapies based on the patient's molecular profile. At organ level, modelling is done by tightly coupling different physics (e.g. fluid, tissue, electrophysiology, chemical reactions, heat, transport of large bodies, particles or species) with contributions from different temporal and spatial scales (cells, tissue, organ, system). To make things more complex, these problems present issues such as patient variability and comorbidities in complex geometries, with extremely difficult validation. A strategy to address these two issues is to run problems in a virtual patients population. All these things together make the use of supercomputers a decisive factor.
- For Energy applications, simulations are of importance to improve the efficiency of hydropower, wind turbines, nuclear power, and, not least, batteries and high-voltage cables to enable transmission and storage. Likewise, simulations are essential for the discovery and optimisation of renewable forms of energy, their storage and distribution. The oil and gas industries are moving to full waveform inversion combined with neural networks for accurate detection. Exascale resources will make this technique feasible, allowing more accurate predictions of reservoirs and, although still fossil fuels, oil and gas have much reduced CO^2 emissions and air pollution compared to those produced by coal, which is still the dominant source of energy in the world. Accurate magnetohydrodynamic simulations of plasma are critical for fusion energy, as in the ITER project.
- Computing is already used widely in Engineering & Manufacturing. Engineering applications based on fluid dynamics, combined with orders-of-magnitude-faster resources, will enable direct numerical simulations of the governing equations of fluid mechanics with better accuracy, leading to improved designs and thus significantly better fuel efficiency e.g. for cars and airplanes, while also helping us understand phenomena such as cavitation, flow separation and pollutant formation. New data-driven approaches will enable scientists in academia and industry to integrate all aspects of design in models, use information from internet-of-things sensors, include uncertainty quantification in predictions, and consider the entire life-cycle of a product rather than merely its manufacture (digital twins).
- Chemistry & 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 handle systems, properties and processes of increased complexity, and drive towards extreme accuracy. These methods are being complemented both with multi-scale models and data-driven approaches using high-throughput and deep

1. See for instance: The Scientific Case for Computing in Europe (period 2018-2026), PRACE

learning to predict properties of materials and accelerate discovery. This will enable researchers to address the grand challenge of designing and manufacturing all aspects of a new material from scratch, ushering in a new era of targeted manufacturing.

- Global Challenges, like urban air pollution, financial crises, behaviours in social networks, pandemics (like COVID-19), food and water shortages, management of disasters or conflicts, and migration streams require complex simulations. These are often, using agent-based systems with a huge number of elements, and specific coupling with other simulations (such as weather and climate predictions and other CFD simulations in 3D models representing digital twins of real locations). These simulations can be used by policy makers in order to understand how to solve existing global issues or how to mitigate their negative effects. Therefore, it is necessary to advance research on agent-based simulations, coupling technologies and executing ensemble methods, as well as leveraging state-of-the-art HPDA and AI techniques. All of this significantly contributes to the retrieved models, their understanding, and the further extraction of knowledge for scientific-based decision making.
- A world class European computational infrastructure will expand the frontiers of fundamental sciences like physics and astronomy, supporting and complementing experiments. Researchers will be able to simulate the formation of galaxies, neutron stars and black holes, predict how solar eruptions influence electronics, and model properties of elementary particles. This will explain the source of gamma-ray bursts in the universe, advance our understanding of general relativity, and help us advance understanding of the fundamental structure of matter by means of simulating the theory of strong interactions called quantum chromodynamics. This fundamental research itself leads to advances in the state-of-the-art of scientific computing and helps attract new generations to science, technology, engineering and mathematics.

Throughout this spectrum of applications, handling of complex data plays an increasingly important role. Existing research fields are beginning to use deep learning to generate knowledge directly from data instead of first formulating models of the process, so next-generation infrastructure must be able to handle these applications with dramatically increased data storage and I/O bandwidth capabilities. This will be needed for e.g. autonomous driving, Industry 4.0, Cyber Security and the Internet of Things, but will also enable the application of computing across a whole range of non-traditional areas including the humanities, social sciences, epidemiology, finance, promoting healthy living, determining return-on-investments for infrastructure by considering behavioural patterns and, not least, in helping to develop society and secure democracy.

As the previous topics show, industry can be found almost everywhere when it comes to large scale simulation science. While in many cases, especially when SMEs are using simulation, data analytics, and AI, industry does not yet require the fastest systems available, it is absolutely crucial that industry has a direct link to those who use these systems and to their technologies in order a) to prepare for their own use as soon as they have a need for this kind of performance and b) to be able to run exceptionally complex tasks, not every day but for particularly large challenges when this compute power is needed to open up new possibilities. If industrial companies are not able to seamlessly connect to the highest-end simulations, data analytics, and AI technologies, they will suffer severe disadvantages compared to those that do. This also requires lowering the entrance bar, through approaches like “simulation as a service” and support from dedicated HPC teams.

CHALLENGES FOR 2023-2026

The computing requirements of these scientific and industrial applications drive technological development and application developers must respond in turn by porting, adapting (including the use of new methods and algorithms), and optimising application codes for the new technologies. Knowledge of and sufficient access to these new technologies for application developers will assist in ensuring that full scale systems can be used in the most efficient manner. An effective collaboration, including co-design processes, between application and technology developers is crucial for a successful HPC ecosystem to ensure technology is relevant for applications and can be fully exploited when available. It must also be recognised that not all application maintainers possess the knowledge to independently and effectively conduct the necessary code modifications for new technology. Specialist knowledge and support will also be needed to assess the level of revision needed to move applications which are successful at small-scales to large-scale HPC infrastructures. Application co-design will involve providing support to application codes and tools as they move through these processes. In addition, the challenge of transferring knowledge between computational and scientific specialists extends beyond porting and optimisation of applications and should include training existing and upcoming researchers in skills critical to fundamental HPC usage. This is even more essential in research fields that have a relatively immature level of uptake and competence in HPC usage. Incorporating all of these issues into a wider approach will foster an ecosystem of skills and expertise in Europe by retaining scientific talent and enhancing the competitiveness of industries on the global market.

Educating new generations of HPC experts and computational scientists is crucial for a healthy HPC environment. In addition, such people need rewarding employment conditions with clear and structured career paths to retain skilled personnel and avoid a brain drain to other parts of the world. Unfortunately, within the traditional academic systems that exist in most European countries this is still difficult for computational experts in certain disciplines and even more so for interdisciplinary experts.

Access to increased computational power and to applications exploiting the features of the system remains crucial to enable more detailed and large-scale (compute intensive) modelling and simulations. At the same time, new approaches like data-driven computing, High Performance Data Analytics, and AI, and their convergence with classical HPC enable new opportunities and require new capabilities. This includes efficient access to large amounts of data with low latencies and high-bandwidth and support for new and large workflows and ensembles encompassing orders of magnitude more active tasks or computational jobs than today. Additionally, some simulations require coupling with others, since their results need to be synchronised, requiring further development in coupling techniques (from data sharing to message passing) and adequate resources access and allocation mechanisms.

Many applications are severely limited by memory bandwidth or communication latency, and as the throughput of floating-point operations has increased faster than the data transport capabilities, even many traditionally floating-point bound applications are now highly memory sensitive. Intra-node and inter-node communication require drastically reduced latencies and high-speed networks, particularly for algorithms based on fast iterations of short tasks so that they can achieve significantly-improved performance and strong scaling. In this respect, the process of co-designing the hardware with the applications in mind can already be seen with the addition of High Bandwidth Memory (HBM) interfaces not only to accelerators such as GPUs but also on general purpose processors such as the Fujitsu A64FX processor and in upcoming x86 variants such as the Intel Sapphire Rapids CPU. The impact that this has on application performance has already been demonstrated with the Fugaku system that is based on the A64FX processor and where significant real application speed-ups were obtained on a wide range of scientific codes. Another area that has seen a large degree of co-design between hardware and applications is reduced precision. Modern GPU and CPU processors now offer native FP16 support that mostly targets matrix-matrix operations found in machine learning workloads but there are already a number of applications such as in weather and climate that are adapting parts of their model to exploit reduced precision with promising results both in terms of speed-up and accuracy.

Storage and I/O requirements are expected to grow even faster than compute needs, with much larger data sets being used e.g. for data-driven research and machine learning. This is not limited to the amount of storage, but data-heavy applications will also need exceptionally high-bandwidth parallel file systems, and/or advanced data caching solutions on each node. This increase in storage and I/O resources must be coupled with provisioning of a large-scale end-to-end data e-infrastructure to collect, handle, analyse, visualise, and disseminate petabytes to exabytes of data. In addition, some applications like genome analyses or AI applications require ancillary data like reference datasets or underlying models. Making them available on fast storage with pre-configured access mechanisms can significantly speed-up such applications.

Long-term maintenance and portability of codes (both in terms of performance portability as well as compiling/building the codes on new architectures) are other important factors, requiring standardised, open, and supported programming environments and APIs, including container technologies, supporting a wide range of different hardware technologies. Investing in tools automating and tracking deployed software stacks can help to sort out the challenging combination of codes and architectures. In addition, the software engineering practices like unit testing, continuous integration and deployment need to improve. These efforts, typically performed by the development teams, need to be complemented with community benchmarking activities, helping to select the right software for certain environments and ensuring the validity of the results. A particular challenge is the uncertainty about future hardware developments. Porting established codes is a major undertaking and with uncertainties about the long term future of certain hardware technologies, there is often a reluctance to engage in these expensive endeavours. Performance portability frameworks can shield application developers from hardware changes to some extent, but long term, stable hardware roadmaps are of equal importance.

Many of today's applications are made up of millions of lines of code. Analysing these monolithic codes to identify expensive computational kernels to port to future architectures can be difficult. Creating application dwarfs to simplify the execution and interaction with the algorithmics can be a solution to enhance this task. As an example, in Earth Sciences, dwarfs represent functional units in the forecasting model, such as an advection or a physics parameterisation scheme, which also come with specific computational patterns for processor memory access and data communication. To ensure a smooth user experience, dwarfs should be distributed with an accurate user guide and examples of input files or namelists. In addition, modularisation of codes and the application of modern software engineering best practises are important.

While in the longer-term computing leadership will require the development of alternatives to the current technologies, including quantum, data-flow, neuromorphic or RNA computing, there is consensus that even today the fundamental mathematical and computer science algorithms that are needed to meet the requirement of leadership science and industrial competitiveness are not in place. The energy-efficient, application-oriented next-generation computing platforms therefore require an ambitious programme of algorithm development integrated with the co-design/co-development of the overall infrastructure and sustained over longer timescales than the usual 3 to 5 years funding cycles. This is particularly true for quantum computing, where in Europe and globally a vibrant research environment is being built rapidly. World-leading technology and applications innovation in this field requires significant investment in algorithmic and application development, ensuring that the anticipated quan-

tum computers will be useful for the European Research Area. This must include research funding and support for fundamental new algorithms and quantum information theory, and for robustness, reproducibility, data & I/O and the convergence with classical computing.

All these approaches require co-design activities involving architectures, OS, communication libraries, workload management and end-user applications to achieve the intended results.

INTERSECTION WITH "SUSTAINABILITY" RESEARCH CLUSTER

Application co-design has various interfaces to sustainability, particularly when it comes to more efficient usage of resources:

- Improved applications will usually have better performance in various ways, but any of these actually delivers a result at least as good in a better time. This immediately reduces the usage of resources like for example energy.
- Improved applications will have a better effect on their user's group, thus making simulations more attractive and thereby contributing to sustainability (e.g. obviously in the fields of energy efficiency, global challenges or in Climate, Weather, and Earth Sciences)
- Being able to support the porting of codes to new (and usually more efficient and faster) platforms, will help sustainability in two ways: directly by being able to run more efficiently and indirectly by reducing the need to develop new software for new platforms.
- The improvement regarding the handling of the increasingly large amounts of data will lead to better efficiency and less resource usage.

Another angle is the sustainability of application development itself. Particularly open-source (or otherwise freely available) software, which is the majority of scientific software, requires institutional funding to sustain the efforts. Often, research funds are being redirected to support application development, however this often is not optimal, since making software sustainable is mostly not exactly a research task; in addition this leads to an unsteady stream of funding and thus a critical lack of continuity in the further development. Much has improved over the past years with research software engineers accepted as a new profession in research environments and funding dedicated to software development. Yet, all these efforts are still somewhat in their infancy and need to be stabilised and increased.

In general it can be stated that many activities around application co-design will have an impact on sustainability in some way. It therefore definitely makes sense to improve the communication between application co-designers and people focussing on sustainability making sustainability a central component in any co-design activity.

6.2.8

Centre-to-edge framework

RESEARCH TRENDS AND CURRENT STATE OF THE ART

HPC calculations are increasingly integrated as core elements of large computing and data workflows and infrastructures which enable advanced scientific discovery and engineering of superior artefacts, provide monitoring, prediction and control of highly complex man-made systems, or support far-reaching, goal oriented decision making (like towards mitigating natural disasters and establishing a truly sustainable economy). Such infrastructures regularly involve a huge amount of sensor or actuator devices, Edge or Fog layers which perform computations and reduce sensor/actuator data volumes, and central computing and data storage centres. Such infrastructures also involve ML/AI and Big Data techniques, and are referred to as "Centre-to-Edge Frameworks" (CtE).

EXAMPLES OF CTE USE CASES

Large instruments like particle accelerators or optical and radio telescopes have significantly advanced scientific discovery and the associated expense has led to the establishment of international and increasingly global collaborations. With new, geographically distributed instruments such as the Square Kilometre Array (SKA), global HW/SW infrastructures are being build that receive the instruments' enormous data rates, pre-process and clean the data, and make it available to scientists worldwide, which in turn use HPC and ML/AI techniques to derive knowledge and discoveries. The SKA involves several data concentration layers, which collect, correlate and filter the data coming from the antenna arrays, uses distributed data archival sites and a network of large HPC resources for data analytics, simulation and discovery.

In industry, factories are continuously growing in size (for instance in automobile and chip production), and multiple sites across the globe have to be closely coordinated. Digital twins mirror the components of a real factory, monitor their state in detail and control factory operation to optimise efficiency and ensure health & safety. Increasingly, simulation and prediction methods guide the automated operation of factories. The growing need for flexible production (small lot sizes or on-demand) requires pre-planned, automated operation and retooling.

The EC-funded Destination Earth initiative takes the Digital Twin concept to a new level: it models and simulates the earth's atmosphere, oceans, river systems, flora/fauna, and selected human activities. For this, it uses novel, micro- and mesoscale simulation and prediction algorithms, and provides fine-grained forecasts for disaster prediction and management, complemented by less detailed long range predictions. The latter enable making informed decisions affecting society as a whole on the way towards an eco-friendly, sustainable economy. The system will also model the impact of human activities on the Earth, giving decision makers access to the information to reach near-optimal decisions. Data will come from a huge number of traditional (satellites, observation stations) and less conventional (like mobile phones and other IoT devices) sources. Observation and simulation data will be archived for use by scientists, operational and political decision makers.

Widespread deployment of interconnected communication systems and of autonomous vehicles requires a global-scale monitoring, fault detection and deployment infrastructure, which ensures safe operation of the fleet of deployed devices/vehicles. It creates and rolls out updates of control models (for instance in reaction to failures and accidents) and thus benefits customers and society as a whole. Since autonomous vehicle operation will involve highly complex ML/AI-based control layers, the improvement of such models will require significant HPC-class compute capacity. An example is Tesla's current Dojo supercomputer installation.

Communication in social networks involves billions of users and an extremely high data rate and volume. Rising concerns about quickly recognising and countering "fake news" and the need to identify criminal or terrorist activities as early as possible lead to large-scale automated surveillance of social media communication, increasingly based on advanced graph analytics methods coupled with ML/AI. Most graph analytics algorithms use sparse linear algebra techniques and require very significant computational throughput for operations on sparse data structures. Protection of data has to be guaranteed (due to GDPR), and the surveillance schemes are run by governments or by social network providers. Similar, smaller-scale use cases exist in cyber security (like for large enterprise networks) and financial transactions (fraud detection).

HPC IN THE LOOP

A key element in all CtE use cases is the integration of potentially extreme-scale HPC resources in a geographically distributed pipeline (like for the large scientific instruments), or in a closed control loop (like for instance for digital twins), with huge data streams providing input data and, depending on the use case, high-volume output streams controlling a very large set of devices or distributing large-volume updates.

To support HPC in the loop use cases, HPC resources have to be continuously available and with high reliability, at short notice (like for disaster recovery), or at precise, pre-planned times. The amount of computation required likely depends on the data rates and often on the exact data itself. Most often it cannot be exactly predicted with state-of-the art methods. This raises a major challenge regarding the efficient operation of HPC resources, which today relies on batch scheduling of jobs with prescribed and fixed resource requirements and durations, with a near-optimal, fixed execution schedule built at job entry time. HPC resource management, orchestration technology and parallel programming environments need to evolve to provide the required flexibility to run continuous or regularly recurring computations with varying resource allocations without keeping an excessive number

of nodes dormant (to be able to meet peak demand). In addition, quick reaction to incoming, high priority HPC jobs is essential to meet the near real-time requirements of certain CtE scenarios (like for instance disaster recovery).

CtE use cases are based on continuous high-volume data streams, at least for the input data, while current HPC installations are optimised for a stage-in/compute/stage-out operation, with highly tuned parallel file systems receiving input data before the HPC calculation starts and making results data available for download after the job or job step has ended. While support for more interactive use of HPC is growing, this usually pertains to control of applications based on visualisation of results, and does not address the input stream problem, or provide prompt access to the full output data stream. The heart of the challenge is to find ways to prevent starvation of the fast HPC processors (which parallel file systems succeed in) while avoiding the time and energy effort of staging and copying data.

FEDERATED USE OF HPC RESOURCES

Federated use of HPC resources and centres is important for many CtE use cases to provide redundancy in case of system failures or overloads, and to provide the required end-to-end data and compute throughput by use of multiple centres. Today, local user ids and access privileges are still prevalent in HPC, with scientific or engineering users having to maintain multiple accounts, credentials and resource allotments across different systems. Efforts to provide federated access based on a single set of credentials and fungible resource grants date back to the last century (e.g. Unicore), yet little sustained progress has been made. Recent, trans-domain efforts in the science community, like the Fenix research infrastructure, are making progress in establishing a modern software architecture with leading HPC centres as supporters. A key factor here is the adoption of modern mechanisms for identity management, authentication and authorisation, and standardised access procedures (including operational and legal aspects). Compared to earlier efforts, there has been significant progress in the virtualisation field, with approaches like containers or serverless computing addressing the federation of the many technical differences between systems.

Still, CtE use cases require a tight integration of IoT, Edge, Fog, Cloud (AI and data analytics) and conventional HPC resources, preferably in a commonly accepted reference architecture. Gaia-X could play a beneficial role here, and the recent EC initiative on funding a European resource and data federation middleware named SIMPL looks like a path to a solution. Gravity of data will play a significant role here, and common data federation interfaces and infrastructures will be required.

LARGE, DISTRIBUTED WORKFLOWS INVOLVING HPC RESOURCES

The use of workflows, that is mostly graphs of tasks connected by timing or explicit data dependencies, has become commonplace in many uses of HPC resources. Often, these workflows consist of repetitions of three distinct phases: stage-in and preparation of input data, the main HPC computation on that data, and processing, visualisation and stage-out of results. CtE use cases, in comparison, will use much more complex workflows which involve the full chain of data generation by sensors, concentration and pre-processing in an edge layer and potentially on several central Cloud resources, with results being streamed into HPC compute steps and results data being distributed to many recipients via multiple Cloud and Edge/Fog layers. Also, feedback loops will occur, with results data shaping steps like model improvement and AI training. Such workflows must run reliably and with high efficiency across vast distributed infrastructures, and for mission-critical use cases, resiliency against failures of any infrastructure component has to be assured. One way to achieve the latter would be to support quick and transparent failover to alternative infrastructure components.

For urgent decision-making scenarios, which could be seen as simple precursors of large CtE use cases like Destination Earth, complex, distributed workflows have been developed and, more importantly, prototypes for their integration with HPC centres were created (such as in the VESTEC and eFlows4HPC projects). This work must be extended to cover the genuine requirements of CtE frameworks, in particular the integration of very different disciplines like HPC, AI/ML and data analytics, and the support for multi-layer, potentially globally distributed infrastructures.

DISTRIBUTED SCHEDULING AND ORCHESTRATION

A key requirement of CtE use cases is the ability to run workflows with steps being executed at many different infrastructure elements (sites) in an organised and reliable way, with the potential need to meet near real-time deadlines. As discussed in section 1.2, the batch scheduling schemes prevalent with today's HPC centres are not well suited, yet there is a much more difficult problem to address: the various resource providers or centres making up a CtE infrastructure will all use their own, local resource management and orchestration systems, and CtE frameworks must provide a higher-level scheduler/orchestrator that decides where to place workflow steps in order to meet execution deadlines. Since the workload of each step can and in reality will vary, and since we cannot assume a fully dedicated infrastructure, scheduling and orchestration has to dynamically adapt to changes in workload of the use case and in condition/availability of the infrastructure.

Cloud service providers have demonstrated how scheduling and orchestration of large, dynamic workloads can be handled with high efficiency and reliability for huge data centres or for hyper-connected ensembles of these. HPC computation steps are of a different nature (e.g. requiring large numbers of co-located nodes), yet it should still be possible to learn from the current practice of the Hyper-scalers at least for the close monitoring of millions of systems and adaptive orchestration of services driven by this data.

Higher-level schedulers (often called Meta-schedulers) have been developed and used in the HPC field in the past, with one example being the Globus Gridway engine or ATOS Yorc. For the CtE

use cases, meta-scheduling approaches have to be extended to support large numbers of nodes of disproportionate capabilities (like IoT devices, edge/fog gateways, general-purpose and AI compute units, and accelerated HPC resources). It will also be necessary to cover the orchestration aspects (including on-demand provisioning of optimised implementations and builds of required software), and to double down on optimising work placement to achieve highest energy efficiency.

PROGRAMMING MODELS

Programming models for different disciplines (HPC, ML/AI in its deep learning and conventional varieties, data analytics, and maybe even Cloud-specific models geared towards "serverless computing" like Amazon Lambda and others) must be composed in efficient and seamless ways (i.e., minimising data transformations and transport) into either workflow steps (analogous to linking modules into a single executable) or sequences of these steps (analogous to multiple distinct executables). The former will usually be more efficient than the latter, since it will run all operations in a single memory space on the same data. On the other hand, it will require agreement on APIs and data formats, whereas the latter can accommodate substantial differences in execution systems, environments and data formats, with additional data transfer and transformation steps requiring additional energy.

Clearly, it will not be possible to a priori mandate which of the two approaches should be taken for any set of computational disciplines and programming models. The decision will depend on the exact circumstances, including the intended execution and interconnection systems, their system software, the exact use case and the relative importance of efficiency (i.e., low energy use) vs. aspects like availability of a combined solution or its ease of programming. Generally, it will be critical to define APIs and data formats in a reusable way and avoid fragmentation caused by overlapping ad-hoc, point solutions.

The working groups on programming environments and mathematics & algorithms are key contributors.

SERVERLESS PROGRAMMING MODELS INVESTIGATION

Modern "serverless" programming has been mentioned above – in the Cloud Service Provider world, its implementation as execution graphs of microservices has been gaining significant momentum, and Cloud data centre monitoring, orchestration and scheduling techniques fully support this concept, achieving very high efficiency for Cloud workloads. Such installations also increasingly rely on smart NICs, which take over key parts of the scheduling and orchestration operations. HPC workloads are of course very different from typical Cloud use cases – it should, however, be investigated whether microservices or the system software functionality supporting them (extremely detailed system monitoring and fine-grain scheduling/orchestration decisions) could drive up efficiency of HPC workloads beyond what is possible with conventional approaches. CtE use cases would be a good field for this kind of exploratory research, since they integrate the use of Cloud or Edge resources already.

CHALLENGES FOR 2023-2026

CLOUD-LIKE FEDERATED ACCESS TO HPC RESOURCES & CENTRES

Key challenges to be addressed 2023-2024 include federated handling of identities, authentication/authorisation and resource/budget allotments. In effect, a single identity (which can be tied to a natural person or to a legal construct like a set of customers, a project or a company) must be able to access and use all compute and data components in a CtE infrastructure using the same authentication procedure, which must be able to be automatically completed to support lights out operation. Resource allotments or budgets must apply to the CtE infrastructure as such and be fungible across its physical and organisational components. Alternatively, billing on-demand is an option, as long as it is able to use a single ledger and source of payment. The full set of resources and functionalities must be covered: computation, data storage and retrieval, data transformations, use of network communications etc. Where possible, existing mechanisms from the Cloud area should be adopted or adapted.

For 2024-2026, focus should shift to federation of the capabilities of each infrastructure component or service provider (covering compute, storage and networking) – instead of addressing a named site, users would request certain functionality and SLA conditions, and a broker layer would identify potential takers. Conditions of access, including pricing and performance deadlines must be expressed in SLAs that can be automatically processed and supervised.

It is expected that interfaces to resources or providers would be based on a service formulation and be supported by tailored user interfaces emphasising ease of use and user productivity. Methods for data access and transmission into/out of HPC centres have to be adapted to work with data streams, and ways to make streaming data available to HPC processors that are more efficient than copying the data stream into parallel file systems have to be found.

The European Commission has started a significant initiative to fund the implementation of a European resource and data federation middleware named SIMPL (https://digital-strategy.ec.europa.eu/en/news/simpl-cloud-edge-federations-and-data-spaces-made-simple). At the time of writing, the requirements analysis and top-level architecture had been finished, and fully funded calls for consortia to complete the design of SIMPL and implement it are indicated for autumn of 2022. The final implementation of SIMPL shall be at a high TRL and suitable for actual, operational use. Support for HPC is indicated in the architecture, without further details given. Coordination with the SIMPL effort looks prudent, since it would leverage the very considerable amount of investments by the EC into a European Cloud middleware.

The second challenge will need coordination with WG3 (System Software), which is responsible for HPC resource management. In addition, mechanisms for optimised access to streaming data are in the WG5 (I/O and Storage) area of competency.

HPC-IN-THE-LOOP USE CASES

The main challenge covers the integration of on-demand or interactive use of shared HPC resources managed by a single organisation (such as an HPC centre). Objectives are to extend the traditional batch-style scheduling systems to support priority driven on-demand usage of parts or all of the managed resources, with an additional requirement to enable interactive use of such resources by humans or by other steps of an CtE workflow. Key elements of this challenge are to uphold the high utilisation rates and efficiency of traditional batch systems, while at the same time offering fast redeployment of resources to on-demand usage. Integrating mechanisms for storing and reinstating state of preempted batch jobs would be required to avoid excessive duplication of work.

In addition, mechanisms for accommodating time-varying workloads should be prototyped and integrated. CtE use cases will mostly not have a fixed effort profile – being able to acquire and release additional resources in an adaptive way will be necessary. As long as performance SLAs are satisfied, this could to some extent be implemented using effective co-scheduling and sharing of resources among applications.

There will be the need to closely coordinate with R&D work of WG 3 (System Software) around extensions to resource management and scheduling layers to accommodate rapid and malleable/elastic provisioning of HPC resources. Also, support for fast pre-emption of running jobs and preservation/restoration of their state requires support from WG3 for software interfaces and WG 2 (System Architecture) for hardware support in preserving system state.

RELIABLE AND EFFICIENT DEPLOYMENT AND EXECUTION OF DISTRIBUTED WORKFLOWS

This challenge will require the development and demonstration of effective hierarchical scheduling and orchestration layers which liaise with advanced local scheduling and orchestration systems to place and execute steps in large-scale distributed workflows. In detail, this will entail development of interfaces to inquire status of the local systems, in particular approximate cost and execution deadlines for workflow steps, investigation of scalable scheduling algorithms and implementation of suitable candidates that are able to cope with the scale of upcoming CtE use cases, and creation of interfaces for the local deployment execution and monitoring of workflow steps. A hierarchical scheduling stack could work with a combination of pre-computed schedules and ad-hoc decisions, as long as it can adapt its decisions to the variability of specific CtE use cases and infrastructure availability. Local, "lower-level" scheduling and orchestration layers will need to evolve to provide the functions required by the hierarchical layers.

An important element is the ability to predict at which time a workflow step under consideration will be run on the suitable infrastructure components, and what energy it is likely to consume. Reliability of these estimates will determine the overall reliability and efficiency of the meta scheduler in a CtE deployment.

Part of this work will depend on results from 2.2 above, and there will be the need to closely coordinate with R&D work of WG 3. Selection and extension of workflow models and interfaces have to be coordinated with WG 4 (Programming Models).

SIMULATION AND ANALYSIS CAPABILITIES

The amount of energy used by the execution of a CtE use case and the capital costs for its realisation will to a large degree depend on the architecture and setup of the infrastructure and on the success of the hierarchical scheduler to place work in a near-optimal way. This challenge comprises ways to size and simulate a CtE infrastructure given a suitably detailed description of its function and data characteristics, as well as innovative ways to measure and predict the end-to-end amount of energy a workflow step will take on the suitable infrastructure components. End-to-end here means that energy used for computation, data access and communication is added up.

The simulation capability will support the initial architecture, design and sizing of a CtE infrastructure; the measurement capabilities will serve to inform the simulator and to build a database of ground truth which can then support the predictor in forecasting likely execution time and energy used. Both measurement and prediction will interface with the local and hierarchical scheduling layers, and they should in addition archive all relevant data to accommodate the simulator and off-line analysis of software efficiency and system operation. Modelling data-dependent calculation paths and pre-emption by the OS or other applications presents a major challenge here.

The proposed actions need close collaboration with WG 3 (System Software) for scheduling aspects and SW interfaces and WG 2 (System Architecture) for component system simulation and measurement/monitoring functionality.

SERVERLESS COMPUTING, MICROSERVICES AND HIGHLY AGILE SCHEDULING/ORCHESTRATION

As discussed in the last paragraph of 1.6, micro-services are emerging as a central way to program Cloud use cases, and HW and system SW solutions are becoming available which promise to further raise the efficiency of operating huge Cloud data centres based on the micro-services paradigm.

A deep investigation should be conducted which looks into potential uses of the micro-service paradigm to implement HPC use cases, and conversely into whether and how the new smart NICs / Data Processing Units / Infrastructure Processing Units and the associated fine-grained monitoring, scheduling and orchestration system software can benefit conventional HPC applications.

This activity is in the area of WG 4 (Programming Environments); the investigation actions proposed will however not require significant contributions from WG 4.

INTERSECTION WITH "SUSTAINABILITY" RESEARCH CLUSTER

Minimising the amount of energy used for the operation of a CtE use case is a key factor contributing to its sustainability. Energy is determined by the duration for which the CtE infrastructure components are used and their power consumption. For ordinary HPC applications, energy to solution can be addressed by increasing performance (as long as the increase in power consumption is lower) or by driving down energy consumption (as long as performance decreases are lower). For CtE use cases, the situation is made more complex by the potential near real-time deadlines (which require a certain minimum performance), and by the option to run on a dedicated infrastructure, which should be built out in a way for which deadlines are just reached without keeping components idling for significant periods. Mechanisms that support “dark silicon” and enable fast on/off switching of components (including cooling units) would be a significant factor in maximising the energy efficiency of dedicated infrastructures while still providing a safety buffer in case of raises in computation or data requirements.

A second key factor is the placement of work, in particular for the case of shared infrastructures. If the hierarchical scheduler/orchestrator is aware of the load of all infrastructure components and the energy cost of placing a workflow step (which partly depends on data location) , it can minimise overall energy used by reaching smart placement decisions. Judicious use of Edge/Fog devices vs. central Cloud or HPC resources will be an important area to cover.

Other factors influencing sustainability are capital costs for acquiring and installing the required systems, and for developing and maintaining the necessary software layers and components. A balance will need to be found between energy efficiency, affordability (i.e., capital costs), and longevity of the infrastructure and the use case implementation. Upgradeability of HW and SW are important elements in extending the useful life of the system, which influences the drain on potentially rare materials and depreciation costs.

The simulation capability described in challenge 2.4 will be an important element of designing CtE infrastructures to be efficient, affordable and thus sustainable.

6.2.9

Unconventional HPC architectures

RESEARCH TRENDS AND CURRENT STATE OF THE ART

As CMOS scaling is facing increasing challenges to deliver performance gains by conventional means, a range of new unconventional architectures emerge.

SHORT INTRODUCTION

Moore's Law, which stated that "*the complexity for minimum component costs has increased at a rate of roughly a factor of two per year* ", is slowing down due to the enormous cost of developing new process generations along with feature sizes approximating silicon interatomic distances. With the end of Dennard scaling (i.e; the rather constant power density across technology nodes, and the increase of operating frequency with each technology node) more than 15 years ago, using more transistors to build increasingly parallel architectures was a key focus. Now, other ways need to be found to deliver further advances in performance. This can be achieved by a combination of innovations at all levels: technology (3D stacking, using physics carry out computations, etc.), architecture (e.g., specialisation, computing in/near memory, dataflow), software, algorithms and new ways to represent information (e.g., neuromorphic – coding information "in time" with "spikes", quantum, mixed precision). A closely related challenge is the energy required to move data, which can be orders of magnitude higher than the energy of the computation itself.

NOVEL HPC ARCHITECTURES

The concepts of standardisation and specialisation have long been recognised as competing forces driving the development of electronics and computers. Standardisation has the benefit of making the technology easier to use and deploy, while specialisation can offer greater efficiency. The concept of a general-purpose processor in itself is arguably highly generic and standardised, and in the last decade we observed a dominance of x86 architectures in desktop computers and data centres, and ARM architectures in mobile devices.

However, in recent years one could observe a trend towards more diverse and specialised architectures. Many desktop CPUs now include GPUs and mobile CPUs move towards completely integrated system-on-chip architectures. They offer different cores for light and more intense workloads as is the case in the ARM big.LITTLE[1] architecture. Other examples include neural accelerators, such as the ones used in Apple's mobile and desktop processors. Embedded systems include specialised accelerators as well, for data manipulation (DSPs, FFT, encoders/decoders, encryption, etc). While these examples are specific to desktop and mobile, we also observe a focus on new architectures in HPC and data centres.

For example, GPUs are now frequently used for a range of non-graphics tasks (General-Purpose GPU or GPGPU), and they are the most common form of accelerator used in HPC. They are commonly used to accelerate compute intensive workloads which exhibit regular parallelism and require floating-point computations and high throughput. In the following we describe some of the most prominent examples of emerging unconventional HPC architectures.

FPGA-based implementation of unconventional HPC architectures

One very promising approach for accelerating a wide range of compute-intensive applications is to use reconfigurable hardware such as field-programmable gate arrays (FPGAs). FPGAs operate at much lower clock frequencies than most CPUs but they allow to create accelerators with massive amounts of parallelism. The accelerator is not programmed using serialised instructions from a fixed and limited instruction-set-architecture (ISA); instead, it is created by spatially mapping the components of the algorithm to a large number of reconfigurable logic, memory and interconnect resources inside the FPGA. Therefore, the generated hardware is highly customised to the unique nature of the algorithm that needs to be accelerated. Often, dataflow models are used to program FPGA accelerators, which avoid the costly movement of data in and out of the multi-tier memory hierarchy of a CPU. In a dataflow model, computations are executed based on data availability. For suitable algorithms, this results in higher performance than multi-core CPUs or GPUs but also much higher energy efficiency.

As a result, FPGAs have started to become competitors to many-core accelerators, such as GPUs and vector processors, in the HPC ecosystem. For more than a decade, there has been extensive research in using FPGAs for application acceleration, and in particular in applications with irregular memory access patterns, and applications that may not need full double-precision floating point formats and instead benefit from reduced precision without loss of accuracy, where significant improvements over GPUs can be realised. For example, AI inference with low precision (even down to 1 bit) can be carried out very efficiently on FPGAs. However, a major obstacle to the widespread commercial adoption of FPGAs in HPC is the higher programming effort and the limited maturity of the development tools, which discourage application developers from utilising them. Historically, FPGAs are used as reprogrammable hardware blocks for special purpose applications and logic emulation, and more recently, as standalone SoCs (e.g. in mobile communication), so the tools and programming environments are usually optimised for building a single monolithic application for the FPGA by following an electrical engineering approach. These programming tools are very difficult to use for software developers and the challenges of using these tools are compounded by the fact that a developed solution is typically limited to a particular device type from one FPGA vendor and porting to another device requires special skills and reengineering efforts.

The lack of software-oriented development tools and the limitations in application portability have been a major challenge for the wider uptake of FPGAs in the HPC field. Another factor has

1. https://www.arm.com/technologies/big-little

been the lack of commercial, off-the-shelf FPGA accelerator cards with surrounding software infrastructure and management tools which is expected for an HPC environment and the higher price of FPGA cards compared to GPUs. Finally, FPGAs are typically used as accelerators on the PCIe bus, where it is inefficient to accelerate fine-grain tasks. Acceleration of larger tasks or batches can be highly efficient but requires a streaming model of computation. All of these challenges have often relegated FPGAs to more niche applications in HPC.

More recently, we saw the introduction of several commercial FPGA solutions oriented for HPC. In 2017, Amazon Web Services (AWS) introduced EC2 F1 FPGA instances, which made FPGAs available to a wide community without the need to purchase special purpose hardware. Xilinx launched its Alveo accelerator cards in 2018, the first datacentre-oriented FPGA solution backed by a major vendor. Several hyperscalers, such as Alibaba, Baidu, Microsoft, Huawei, Nimbix and Tencent also started recently to expose FPGAs to application developers in their datacenter infrastructure or to provide accelerated application-as-a-service, relieving users from the complex FPGA development processes. FPGA programming tools have improved and diversified with various concepts to make them more appealing to traditional software developers. Both Intel and Xilinx have backed OpenCL based programming models which provide developers a commonly used programming model and enables some degree of code portability between different accelerator technologies.

However, the overall suitability of the OpenCL-based programming model for FPGAs remains contested and highly optimised code will no-longer be portable. Other tools have focused on High-Level Synthesis (HLS) where developers can build FPGA applications from commonly used software languages such as C and C++. While this is a more appealing model for software developers, it is not suitable for simply importing existing code and it requires careful code optimisation to create a high-performing FPGA application. Carrying out such optimisations also requires a hardware-oriented mindset and consideration of chip resources. Domain Specific Languages or languages with a specific programming paradigm can help using FPGA in this context. For example, a dataflow-oriented model to programming FPGAs can be used, such as introduced in Maxeler's MaxCompiler[2], leading to a dataflow model that is conceptually closer to an optimal FPGA accelerator architecture and lets developers reason about important design aspects such as bandwidth and parallelism while delegating lower-level optimisations to the compiler. The underlying FPGA technology has also advanced in recent years to include ultra-density on-chip memory, enable dramatic improvements in frequency, adopt 2.5D/3D stacking technology to integrate High Bandwidth Memory (HBM) and to integrate AI-oriented processing units in a network-on-a-chip architecture (e.g. in Xilinx Versal ACAP).

Another approach to optimising FPGA architectures for HPC applications is the move towards more coarse-grain reconfigurable architectures (CGRAs)[3]. While CGRAs have not been commercially successful within the original concept of changing the device architecture to more coarse grain programmable units such as byte-wide interconnect and programmable ALUs and memory tiles, one could argue that modern FPGAs include various coarser grain blocks in addition to the conventional fine-grain fabric. One example is the introduction of a mesh-like CGRA structure of a programmable processing array in Xilinx Versal ACAP devices. While the elements themselves are fairly general-purpose, they are mostly promoted by Xilinx as targeting deep learning and telecommunication applications. However, an adequate software stack is needed to leverage the algorithm-hardware co-design of such a heterogenous fabric with fundamentally different programming models. One could also interpret Groq's AI-oriented Tensor Streaming Processor (TSP), which uses a static dataflow model, as an evolution of the FPGA concept with much coarser grain vector and matrix processing units and wider interconnect (more in section AI chips).

Neuromorphic computing architectures

Neuromorphic or brain-inspired computing describes a field of unconventional computing research that takes certain aspects of what is known about the brain (where knowledge is far from complete) and applies it to novel computing models using various combinations of novel device technologies, analogue circuit approaches, and digital technologies. A common feature of all these approaches is the use of some form of event-based processing, inspired by the use of spikes, which are purely events, as a primary means of representing and communicating information in the brain. The information is essentially coded in time (frequency, coincidence, timestamping) and not in space, like bits. Some neuromorphic research also considers neural models without event-based processing. Neuromorphic systems can be seen as having one or several of the following characteristics:

- Using information encoded in time;
- Using a sparse representation of information (i.e. variation/derivative of signals, event-based);
- Using physical phenomenon to carry computation (e.g., Ohm's law for product, Kirchhoff's law or accumulation of charges for adding, optics for non-linear functions, etc.);
- Using new materials for storing or simultaneously storing and computing (see just above), like the ones used in non-volatile memories (MRAM, RRAM, Spintronic, ...).

While a strict definition of neuromorphic computing is still debated[4], ranging from "very strict high-fidelity mimicking of neuroscience principles" to "very vague high-level loosely

2. R. Dimond, M. J. Flynn, O. Mencer and O. Pell, "MAXware: Acceleration in HPC," 2008 IEEE Hot Chips 20 Symposium (HCS), 2008, pp. 1-26, doi: 10.1109/HOTCHIPS.2008.7476552.
3. A. Podobas, K. Sano and S. Matsuoka, "A Survey on Coarse-Grained Reconfigurable Architectures From a Performance Perspective," in IEEE Access, vol. 8, pp. 146719-146743, 2020, doi: 10.1109/ACCESS.2020.3012084.
4. "2022 roadmap on neuromorphic computing and engineering", Dennis Valbjørn Christensen, et al., accepted manuscript, Journal of Neuromorphic Computing and Engineering, IOP Science.

brain-inspired principles", there is a "wide consensus that neuromorphic computing should at least encompass some time-, event-, or data-driven computation. In this sense, systems like spiking neural networks (SNN), sometimes referred to as the third generation of neural networks, are strongly representative."

Neuromorphic systems will operate as accelerators for specific classes of applications alongside a general-purpose host system. The host system will be responsible for configuring the neuromorphic subsystem - setting up the network topology and parameters, and for converting the data representations and handling data I/O streams. The neuromorphic system will handle the event-based computation itself.

Current developments in neuromorphic computing include architectures based upon analogue circuits, in sub-threshold (as per the original work by Carver Mead) or above threshold (e.g., the Heidelberg BrainScaleS) modes of operation, parameterised digital engines (e.g., IBM TrueNorth, Intel Loihi) and many-core programmable systems (e.g., the Manchester SpiNNaker). There are also developments in novel device technologies, such as memristors, for neuromorphic systems which, when they are mature, will transform the density and energy efficiency of those systems. It should be noted these neuromorphic accelerators chips differ from the new emerging class of AI accelerator chips that carry out neural network computation on architectures that are optimised for this class of applications (efficient sparse matrix operations, tensor operations, convolutions, etc.). These architectures are described in section (AI chips).

In many respects neuromorphic accelerators will play a similar role to accelerators for artificial neural networks (ANNs) but offer improved energy-efficiency due to their sparse event-based model of computation. Compared with conventional ANNs, spiking neural networks (SNNs) have greater capacity to process spatio-temporal data streams, though algorithms for such processing are less developed than those for conventional classification.

Analog computing for AI and neuromorphic applications

In the quest for reduced power consumption of AI workloads, a renaissance of alternative computing technologies can be observed. Computing based directly on the analog quantities of electronic circuits, like electric charges and currents, could in principle deliver the same AI performance at a lower energy and price point compared to the usage of binary numbers, since they require far less transistors to perform the same operation[5,6].

In addition, digital implementations depend on recent process nodes to achieve state-of-the-art power efficiency. Thereby, they incur high NRE costs. Analog realizations can be as efficient in much older technologies, thus allowing for a significant cost reduction in certain applications.

The most common analog architectures accelerate inference in deep convolutional neural networks. They implement the matrix-vector multiplication (MVM), which is the most costly operation in a deep neural network, by a predominantly analog circuit. By utilizing a processing-in-memory structure, the number of data movements is also reduced. They operate by directly multiplying their memory content with the incoming activations, usually by pulse-width modulation. The results of these multiplications are summed-up as currents or charges on shared output lines.

Analog-to-digital converters are used at the output of the analog processing block to restore the partial sums to the digital domain for further processing and the subsequent transfer to the next layers. Analog computing can be combined with novel memory technologies based on memristive, magnetic or phase change materials to store the weight matrix in a dense, power-efficient and persistent way[7]. A first demonstration has already been made of a fully integrated analog in-memory computing core with phase-change memory on 14nm CMOS[8]. There are also novel implementations of floating-gate and ferro-electric transistors as memory cells[9].

Recently, a second analog computing model has been demonstrated successfully, using light instead of electrical signals, for instance as shown by LightOn or LightMatter. Photonic computing uses a substrate similar to microelectronics with integrated waveguides and optical modulators[10,11]. Although at the current state of the technology only small networks can be implemented, their operational speed in the Gigahertz range makes them unique for applications requiring extremely low latencies.

AI chips

In the last few years, a range of new architectures have emerged that are targeted specifically at accelerating AI-related computations. Demand for AI computations is driven by many applications including search, advertisement placement, personalization and data analytics. The widespread introduction of services with natural language processing or image recognition have significantly increased the performance requirements for AI computations where

5. S. Yu, H. Jiang, S. Huang, X. Peng, and A. Lu, "Compute-in-Memory Chips for Deep Learning: Recent Trends and Prospects," IEEE Circuits and Systems Magazine, vol. 21, no. 3, pp. 31–56, 2021, doi: 10.1109/MCAS.2021.3092533.
6. A. Sebastian et al., "Memory devices and applications for in-memory computing", Nature Nanotechnology, 15, 529-544, 2020
7. Y. Xi et al., "In-memory Learning with Analog Resistive Switching Memory: A Review and Perspective," Proc. IEEE, vol. 109, no. 1, pp. 14–42, Jan. 2021, doi: 10.1109/JPROC.2020.3004543.
8. R. Khaddam-Aljameh et al., "HERMES-Core--A 1.59-TOPS/mm² PCM on 14-nm CMOS In-Memory Compute Core Using 300-ps/LSB Linearized CCO-Based ADCs", IEEE Journal of Solid-State Circuits, 2022
9. M. Halter et al., "Back-End, CMOS-Compatible Ferroelectric Field-Effect Transistor for Synaptic Weights," ACS Appl. Mater. Interfaces, vol. 12, no. 15, pp. 17725–17732, Apr. 2020, doi: 10.1021/acsami.0c00877.
10. G. Wetzstein et al., "Inference in artificial intelligence with deep optics and photonics," Nature, vol. 588, no. 7836, pp. 39–47, Dec. 2020, doi: 10.1038/s41586-020-2973-6.
11. B. J. Shastri et al., "Photonics for artificial intelligence and neuromorphic computing," Nat. Photonics, vol. 15, no. 2, pp. 102–114, Jan. 2021, doi: 10.1038/s41566-020-00754-y.

conventional CPUs or GPUs were no longer cost effective. Further applications of AI include cyber security and fraud detection, self-driving vehicles, industrial, healthcare and HPC applications.

The focus of AI architectures is typically on inference; training is only carried out once and can be done with very high effort. Inference is run at a very large scale, often with strict performance requirements, therefore requiring special architectures that can carry out the computations quickly and efficiently. Nevertheless, the acceleration of training is also gaining attention recently. Broadly speaking, AI chips target tensor operations (in particular matrix multiply) in combination with mixed or lower precision operations. Both of these aspects are relevant in AI (in particular inference) applications and focusing on lower precision matrix multiplication makes these chips more efficient for AI applications than GPUs which are more vector oriented.

In 2016, Google introduced the Tensor Processing Unit (TPU). TPUs are focused on matrix multiplication operations and support lower precision operations. The first version TPU only supported Int8/Int 16 and were used for inference, but since the second version, they are also optimized for learning and support a new 16 bit floating point format. TPUs are very efficient for computing models with very large batch sizes and they are programmed through Google's TensorFlow framework. TPUs are mostly used to support Google's internal services with limited access for other vendors to the TPU hardware. Groq recently launched an AI chip called the Tensor Streaming Processor (TSP), a massively parallel architecture of matrix and vector units interleaved with memory units that execute in a static dataflow model[12]. This provides abundant data parallelism and the static execution model eliminates all caches, branch predictors and arbiters. As a result, the chip offers very efficient processing with deterministic behaviour. The TSP is also very efficient at small batch sizes and computes with low latency which is important for real time applications. Furthermore, the TSP offers fast chip-to-chip interconnect which enables the efficient scaling of computational models. Graphcore is offering the Intelligence Processing Unit (IPU), a massively parallel architecture with many small processor cores and local memory, which is efficient at executing fine-grain threads with irregular behaviour[13]. The IPU is conceptually closer to a conventional multi-core processor but provides more parallelism than CPUs and is more efficient at irregular workloads than GPUs. Cerebras has launched the Wafer Scale Engine (WSE), a single, wafer-scale processor that includes compute, memory and interconnect[14]. Processor cores are optimised for sparse linear algebra. The entire architecture is optimised on system level rather than adding PCIe form factor cards to conventional servers. They released the CS-2 system featuring a very large number of cores, on-chip SRAM, and bandwidth.

To address the pervasive need for efficient AI computations in many applications, several conventional architectures have also been extended to enable efficient inference. Intel Xeon Ice Lake processors contain an instruction set extension called DL Boost that supports efficient multiply-accumulate operations and lower precision 16-bit floating point operations. IBM Power10 includes the new matrix-multiply assist instructions executed by four engines per core to accelerate inferencing operations. Nvidia Ampere GPUs have been extended with tensor cores that also support 16-bit floating point operations. Xilinx Versal FPGAs contain AI engines that provide simple vector processors in a VLIW architecture.

Processing in memory - PIM

Processing in memory (PIM) is a generic term that covers any form of integration of processing capabilities into the memory system, of which there are many approaches, e.g., computation in memory arrays, peripheral circuits, and logic layers and dies integrated into the memory devices. PIM has been explored by a large number of academic studies and it has already been demonstrated in various industrial prototypes and products, initially targeting near-memory computing, with DRAM and SRAM, and extending to non-volatile memory technologies of PCM, RRAM, and Flash. Analog in-memory computing can be seen as a form of PIM, where some computations are performed in the memory arrays themselves.

Architectures, storage, and networking

HPC system architecture is evolving from traditional monolithic servers towards servers that can be dynamically composed from a precise set of resources to meet workload requirements. Such resources are envisioned to be disaggregated into pools of various types of CPUs, accelerators, (persistent) memory, storage, etc, accessible via high-performance fabric and orchestrated by an intelligent infrastructure. This will allow maximisation of usage of such resources while powering down unused ones.

This evolution in HPC architecture is needed as a result of the more complex workflows that are increasingly implemented in HPC applications. As an example, they can include large knowledge graphs (memory capacity and bandwidth intensive)[15], training of and inference on AI models (large datasets and compute intensity), and the execution of classic simulation codes (memory or compute bound) . The technologies that will enable such an evolution include the introduction of high-performance coherent fabrics to flexibly wire such composable systems together. Promising steps are the introduction of standards such as CXL, with fast signaling protocols such as PCIe Gen 6 that will provide the latency and bandwidth needed for disaggregation of system resources. In terms of I/O technologies, the continued increase in speed is expected to come from the adoption of new modulation techniques (e.g. frequency domain) and the pairing of short low-latency electrical links together with co-packaged optics.

12. D. Abts et al., "Think Fast: A Tensor Streaming Processor (TSP) for Accelerating Deep Learning Workloads," 2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA), 2020, pp. 145-158, doi: 10.1109/ISCA45697.2020.00023.
13. https://doi.org/10.48550/arXiv.1912.03413
14. 10.1109/MM.2021.3112025
15. Hagleitner Ch. et al., "Heterogeneous Computing Systems for Complex Scientific Discovery Workflows", Feb 2021.

UNCONVENTIONAL MODELS OF COMPUTATION

Approximate computing

Approximate computing encompasses diverse approaches such as variable precision models, novel number representations and stochastic rounding. These techniques are enabled by the unique features of novel architectures (e.g., low precision and integer modes in AI chips and completely flexible arithmetic on FPGAs). The vast majority of current HPC applications use either single or double-precision floating point formats but the actual numerical requirements and properties of algorithms are often not explored in greater detail as conventional CPU architectures simply offer a binary choice between single- and double-precision formats. Double-precision arithmetic is often seen as "correct" even though this number representation is also an approximation in a strict sense and not immune from introducing numerical problems. New unconventional architectures offer a much wider choice of number formats such as low-precision floating point, integer or fixed-point numbers, and using them effectively requires a careful analysis of the numerical behaviour in the application development and porting process. It requires detailed knowledge of how much error is acceptable at the output of an algorithm or solver, and then tracing backwards through the computation what value ranges and precision are needed at the various stages in order to produce suitably accurate output numbers. There is typically no general and analytic solution available to this problem. Instead, the optimisation process often requires experimentation and detailed instrumentation of the code, and the derived solution can be limited to the context of a particular input dataset type or algorithm configuration. Figure x shows an example of numerical application profiling. (F1)

Stochastic rounding describes a rounding process where the result is randomly rounded up or down with a probability that is a function of the residue being rounded. It has been shown to yield more accurate results over long computational sequences than any of the conventional (deterministic) rounding algorithms such as round to nearest, and has the potential to be embedded at low cost into fixed- and floating-point arithmetic units. If the random process is generated by a high-quality pseudo-random number generator (PRNG) the computation can be repeatable by using the standard technique of controlling the PRNG seed. The improved accuracy can be traded against the use of reduced precision operands, leading to reduced energy and memory requirements.

Another aspect of approximate computing is to use approximation circuits instead of conventional arithmetic units, or to use arithmetic units that are somewhat faulty or operating at lower voltages thus producing a certain number of bit errors. Individual results may be wrong, but if the application is based on statistical models, then occasional incorrect results may be acceptable as long as statistical properties of the algorithm are maintained.

Dataflow computing

Conventional processors function by carrying out a sequence of instructions. This model is very flexible but also inherently sequential. Over the decade, many architectural innovations have been developed to boost performance of a sequential CPU such as superscalar execution, out-of-order execution, branch prediction, Single Instruction Multiple Data (SIMD) extensions, and finally multi-core CPUs. All of these make modern CPUs very complex where a lot of silicon real estate is needed to keep a small execution unit running at high speeds.

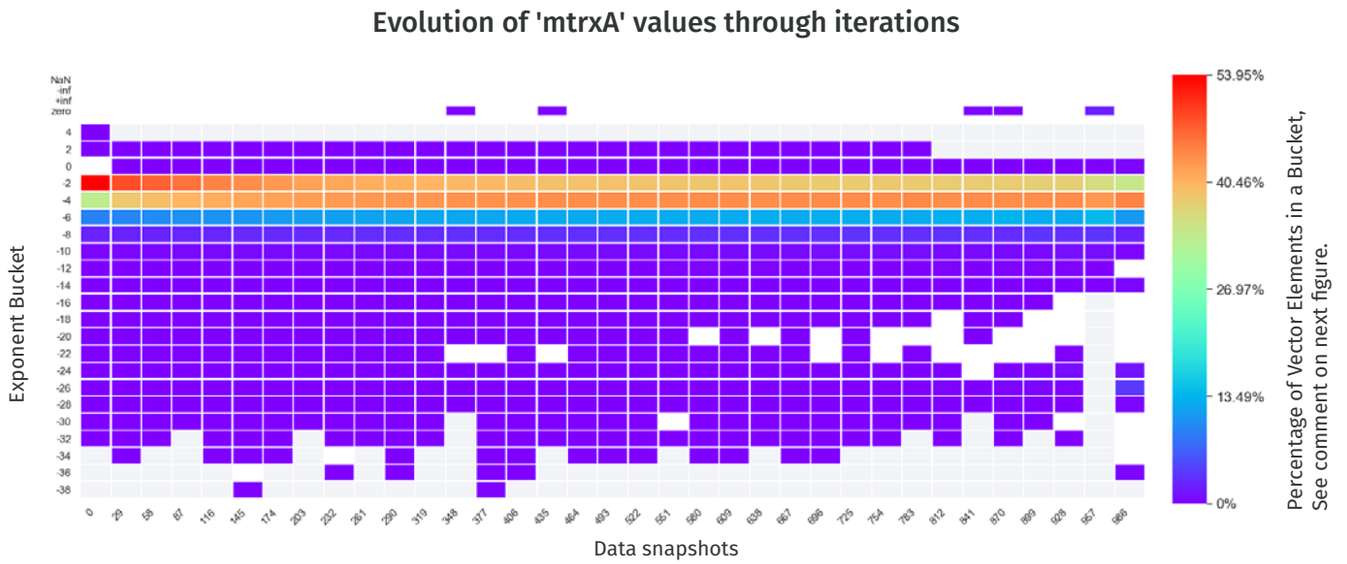

Figure 17: Numerical analysis with Maxeler's value profiling tool.

Dataflow architectures fundamentally differ in that they do not rely on control flow or program counters for execution; instead operations are carried out by making input arguments available to instructions. A whole range of different architectures have been developed under the umbrella term of dataflow and they can be typically classified as dynamic or static dataflow. Dynamic dataflow architectures result in non-deterministic behaviour and require large content addressable memories (CAMs). To date, dynamic dataflow machines have seen limited applications in practice although out-of-order execution can be seen as a very limited form of dataflow where sequential execution can be broken and rearranged depending on the availability of input data. Event-based processing as used in neuromorphic computing can also be seen as a special case of dataflow. Static dataflow is often applied in the context of a systolic array or processing network architecture. Here, data flows from memory over an array of operators where computations are performed when data arrives at the inputs. The outputs are sent forward to the next operator; hence computations are performed by data flowing through the system. This model of computation is deterministic and it eliminates the need for caches as data arrives at the inputs of an operator at a predetermined time. A significant advantage of a deterministic model of computation in static dataflow is that performance and system behaviour is completely predictable, simplifying the design and optimisation process[16]. Furthermore, it minimises data transfers to mostly small local movements and mostly avoids inefficient transfers to off-chip memory which is one the major bottlenecks in conventional CPU architectures. Maxeler Technologies has applied the model of static dataflow to programming FPGAs where a Java-like dataflow language called MaxJ is used to build deep pipelines that are then mapped to the FPGA. The dataflow model is also used in the Groq TSP where tensors flow over the execution units of the chip, resulting in deterministic and low latency operations.

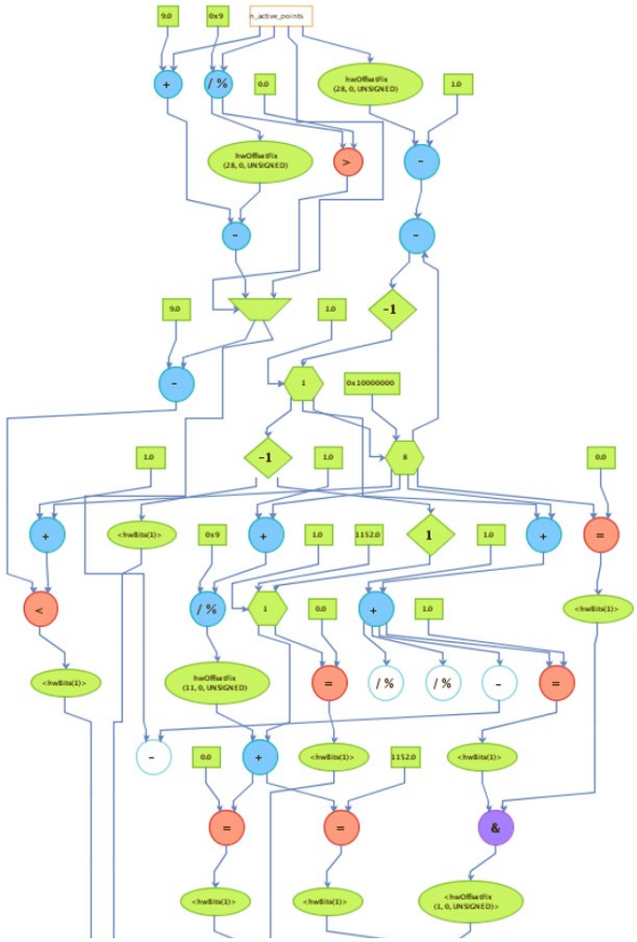

Figure 18: Example of a static dataflow graph. Computations are carried out whenever input data is available and partial results flow through the graphprofiling tool.

CHALLENGES FOR 2023-2026

New unconventional HPC architectures, through some form of specialisation, achieve more efficient and faster computations. A range of new unconventional HPC architectures are currently emerging. While these architectures and their underlying requirements are naturally diverse, AI emerges as a technology that drives the development of both novel architectures and computational techniques due to its dissemination and computing requirements. The ever-increasing demand of AI applications requires more efficient computation of data centric algorithms. Furthermore, minimising data movement and improving data locality plays an important role in achieving high performance while limiting power dissipation and this is reflected in both architectures and programming models. New models of computation emerge that differ from those used in conventional CPU architectures, or models that are purely focussed on achieving performance through parallelisation. Finally, the challenges of developing, porting and maintaining applications for these new architectures are increasingly relevant.

NOVEL HPC ARCHITECTURES

FPGA-based implementation of unconventional HPC architectures

While significant advances towards using FPGAs in HPC have been made, the programming challenge of using productive programming models that are high-level, maintain efficiency of the developed application and support code portability remains one of the central aspects. On the system level, further work is required to better manage FPGAs through an Operating System, execute an application in-parallel on multiple FPGAs, run multiple applications on the same FPGA, and provide support for resource management in a multi-user environment and virtualisation support.

16. Nils Voss, Bastiaan Kwaadgras, Oskar Mencer, Wayne Luk, and Georgi Gaydadjiev. 2021. On Predictable Reconfigurable System Design. ACM Trans. Archit. Code Optim. 18, 2, Article 17 (June 2021), 28 pages. DOI:https://doi.org/10.1145/3436995

Neuromorphic integration with HPC

The next stage of efforts will be channelised at integrating neuromorphic architectures into the HPC ecosystem or "computing continuum" with use cases e.g. in health, agriculture, digital twins & sustainable energy. With the emergence of new architectures, scientific computing workflows are also going through radical transformations. In these transformations Neuromorphic architectures could be the main driver for energy efficient AI/ML processing at edge and at data centers. The importance and need for specialised computational architectures to cater to emerging hybrid workloads and achieve optimal energy efficiency have been highlighted in "Compute-in-Memory Chips for Deep Learning: Recent Trends and Prospect"[17].

The following scenarios are envisioned for future of neuromorphic architectures:

1. Intelligent edge co-processors for distributed cross-platform edge-cloud-HPC workflows
   a. AI inference & training at edge. The data movement is minimised.
2. Datacenter co-processors / accelerators for machine learning training & inference
   a. Energy efficient AI/ML training & inference at scale
      I. Inference for HPC-AI hybrid workloads
      II. Training for AI algorithm (Spiking neural network, backpropagation)

Programmability is still a challenge for these chips, for which the tuning to the application is not done by programming, but by a "training" phase, done on device or off device (on a GPU for example). In terms of computability, there is an equivalent between space and time coding and there are tools that allow to convert Neural Networks trained "with bytes" to "spiking" representation. "Classical" training methods such as back propagation can also be performed with time encoded information ("Spikes"). University level training programs and courses must take the leading role to introduce the concept of neuromorphic computing. Also at the same time HPC centers have to facilitate workshops / training sessions to map current algorithms onto neuromorphic devices.

To facilitate development of neuromorphic applications, various frameworks have been introduced. For instance, Intel provides a software framework called "Lava"[18]. Lava is a modular, extensible, open-source software framework for developing neuro-inspired applications and mapping them to neuromorphic hardware. Lava is platform-agnostic so that applications can be prototyped on conventional CPUs/GPUs and deployed to heterogeneous system architectures spanning both conventional processors as well as a range of neuromorphic chips such as Intel's Loihi[19].

There are several other software frameworks for neuromorphic applications, for example NEST[20], Nengo[21], EBRAINS[22], N2D2[23], or others like NEURON, pyGeNN and the meta-framework pyNN.

AI chips

AI processors are typically programmed through established machine learning frameworks such as PyTorch or TensorFlow. Some AI processors also offer lower-level programming models which can be used to develop non-AI applications on these platforms. In a similar way to how the emergence of GPGPU architectures has led to non-graphics applications being considered for GPU acceleration, there is now interest to port applications from HPC and other fields onto AI chips, leveraging the efficient matrix-compute capabilities of these architectures. This typically requires the algorithms to be restructured such that they can be expressed in a machine learning framework, use the arithmetic precision of the accelerators or some other tensor-level API, and therefore presents similar challenges to other unconventional architectures in that developers need to restructure code and learn new development frameworks.

Processing in memory - PIM

The interest in PIM has grown dramatically over the last years. In 2020, the EuroLab4HPC Vision[24] already included a survey of the approaches for near- and in-memory processing and it recommended funding in this direction. In the same year, ETP4HPC's Strategic Research Agenda[25] called near and in-memory architectures the "ultimate option" to improve energy efficiency, but correctly saw the approach as not mature enough for short-term adoption in production systems. A recent white paper from ETP4HPC[26] advises that wide acceptance of PIM in high-performance computing depends on our ability to create an ecosystem in which a number of PIM approaches can be designed and evaluated, leading to the selection of potential winners and adoption by system architects and end users. The study identifies the main PIM challenges and provides guidelines to prioritise the research effort and funding, encompassing development and evaluation platforms, system software to manage the hardware, APIs and programming models to program it, profiling and analysis tools to analyse it, and benchmarks, simulators and platforms to eval-

17. S. Yu, H. Jiang, S. Huang, X. Peng and A. Lu, "Compute-in-Memory Chips for Deep Learning: Recent Trends and Prospects," IEEE Circuits and Systems Magazine, vol. 21, no. 3, p. 31–56, 2021.
18. https://github.com/lava-nc/lava
19. https://www.hpcwire.com/2021/09/30/intel-unveils-loihi-2-its-second-generation-neuromorphic-chip/
20. https://www.nest-simulator.org/
21. https://www.nengo.ai/
22. https://ebrains.eu/
23. https://github.com/CEA-LIST/N2D2
24. Eurolab4HPC Long-Term Vision on High-Performance Computing (2nd Edition), 2020.
25. ETP4HPC's SRA 4, "Strategic Research Agenda for High-performance Computing in Europe," 2020.
26. P. Radojković, P. M. Carpenter, P. Esmaili-Dokht, R. Cimadomo, H. P. Charles, A. Sebastian, P. Amato. "Processing in Memory: The Tipping Point". European Technology Platform for High-Performance Computing (ETP4HPC). White Paper. July 2021.

uate the overall benefits. Overall, it can be concluded that only coordinated innovations and co-design across the whole stack can make PIM a reality in production.

Architectures, storage, and networking

New challenges will arise to ensure that pools of resources such as persistent memory offer a similar level of reliability and availability as with traditional fabric attached storage. This will necessitate the exploration of new approaches to perform certain types of computing tasks close to memory, or in the fabric itself. Similarly, this will apply to storage to offload the fabric, exploiting techniques such as active storage or processing in storage.

From an ecosystem point-of-view, such composable systems based on a standardised fabric will enable the creation of multi-vendor systems, avoiding lock-in to proprietary fabrics, as well as allow independent updates of components without the need to synchronise all release dates. Furthermore, these more complex HPC workflows are expected to seamlessly take advantage of resources available on-premises as well as across various locations from external cloud providers, for instance to benefit from temporary spill-over capacity or access to highly-specialized and/or bleeding-edge accelerators that aren't available elsewhere. Again, orchestration by an intelligent infrastructure will need to extend to such scenarios[27].

UNCONVENTIONAL MODELS OF COMPUTATION

AI-inspired models for HPC

We also observe a convergence of AI and HPC technologies. AI is very compute intensive and requires sophisticated HPC infrastructure while HPC applications can also benefit from AI algorithms. Fundamentally, HPC applications attempt to model some form of real word process and in many cases (e.g. image or time series analysis) this model can be obtained by feeding training data into an AI algorithm. Practical examples of this are the identification of cancer from a medical imaging process or the discovery of events in a high-energy physics experiment. However, at the moment there is a strong experimental component in the application of AI algorithms to HPC algorithms. There is no common approach to pick the appropriate model and hyper parameters that can achieve fast convergence during training, and new mathematical frameworks will be needed to guide HPC developers in integrating AI into their applications. Another challenge lies in the black box behaviour of AI where the algorithm is known to work with a certain statistical level of accuracy but the behaviour of specific outputs (in particular bad outputs) cannot be understood in the context of the deployed AI model. This can be solved through new explainable AI techniques that allow the behaviour of an output to be traced back and understood through the model. Explainable AI is of particular importance for the greater acceptance of AI in the medical field. Finally, AI and HPC applications are typically built around different programming frameworks and data scientists and HPC developers often have very different backgrounds and skill sets. The convergence of AI and HPC will require efforts to bring together the different technologies and communities. At the same time, AI may also be able to address some of the development and technological challenges by applying AI to the coding process.

PROGRAMMING MODELS, PORTING AND CODE MAINTENANCE

Unconventional architectures typically require specific programming models and tool chains. The introduction of new tool chains often results in training needs for developers who have to become proficient in using a new programming language or model, in particular if the model leverages a different style of execution which requires different thinking when developing applications. For example, using a multi-threading library is relatively easy for code that naturally presents various blocks that can operate independently, but tweaking complex code for efficient parallel execution can be very challenging and requires the developer to fundamentally re-examine the existing code. Similar challenges are faced and are even increasing when programming unconventional architectures. In particular, porting and maintaining large existing code bases is challenging. HPC applications are often developed over many years by a team of developers, resulting in large and complex code bases. A need to port this code to a new architecture is often seen as a significant obstacle by the developer team. Furthermore, the porting may not be just a one-time effort as highly specialised architectures have limited code portability to newer generation devices. The same code may still be usable but may require re-optimisation in order to fully leverage the performance of the new device, as opposed to legacy x86 code which typically does not need constant re-optimisation and can simply benefit from faster processors. This creates a significant challenge in terms of code porting and maintenance which is not sufficiently addressed at the moment.

Developing applications for systems with unconventional architecture can be far from trivial and significantly more complex than conventional HPC programming. This is mainly due to the high degree of specialisation where a set of very specific application and data properties have to be captured in efficient computational solutions. To prevent repeating the errors from the past, e.g., the initial lack of software and tools for the Cell processor, the failure of, for example, Larabee, Tabula and other products, we have to invest early into creating better software tools as well building vivid designer communities around the new, unconventional acceleration technologies. These designer communities can be organised by computing technology, by application domains or ideally by the right combination of the above two. Repositories with design examples, template architectures, application centric methodologies and active discussion boards will serve as the driving vehicle of this process. Considering the challenges around porting and developing full system applications for unconventional architectures, this is going to be crucial. Universities and research centres will play a leading role by involving their undergraduate and graduate students who will contribute at dif-

27. cf. The System Software Research Domain (WG3)

ferent levels. There will be plenty of opportunities for high-impact research and engineering contributions that should naturally attract many European universities to actively participate.

Moreover, these communities will also streamline the discussions between the domain experts, who are the final users of the developed systems, and the computer engineers and researchers working on specific challenges at all different system levels. If organised correctly, such vertical cross-disciplinary discussions will also provide valuable insights for fine tuning of the underlying hardware technologies, computational models and abstractions. The repositories with e.g., template designs, Q&A discussions and various documentation, should be open source.

A complementary approach could be to better structure the code into modules that have defined computing characteristics (that are exposed in a systematic way), and therefore only these parts should be optimised to benefit from the new accelerators that fit with the characteristics. From a system point of view, the work of orchestrating those different accelerators should be further explored and optimised, both at compile time and at run-time.

These challenges of code porting and development are examples that can also be addressed by the introduction of AI-based approaches, such as AI for code. Emerging benchmark datasets such as Project Codenet, and tools such as OpenAi Codex or Deepmind's AlphaCode are a first step in this direction. More generally, methodologies to develop software-hardware platforms must shift towards AI. A leading example is already showing the benefit of ML for the IC EDA tool chain[28]. Developing AI-based tool chains to address the problem of finding the best HPC architecture and associated code bases is clearly one of the most promising approaches for unconventional HPC systems.

CONCLUSIONS

With the end of Moore's Law and the end of "free" performance improvements in sight, the industry has shifted from following a parallelism-oriented approach of pursuing further performance gains to embracing specialised and unconventional architectures as a paradigm to improve both performance and energy efficiency. A specialisation trend has already been observed in mainstream CPU architectures from mobile and desktop to data centre. FPGA accelerators are now emerging in the data centre and offer an extreme amount of specialisation in terms of compute architecture and numerical choices, but are burdened by non-standard programming models and code porting challenges. New AI chips, developed to address the processing needs of data-driven applications, are highly efficient at lower-precision, matrix-oriented computations and are now also being considered for HPC applications. Neuromorphic architectures, based on spiking models, can play a similar role in AI but could potentially be more energy efficient because of their inherent sparsity and potential simpler non-CMOS implementations. Analogue implementations could also help to drive down development cost and power consumption. Processing in memory architectures address the IO bottleneck problem of conventional von-Neuman architecture and enable simple operations close to the data. An overarching theme across all of these architectures is that they reduce data movements because these (in particular off-chip IO) are increasingly becoming a limiting factor from both a performance and energy perspective.

These new architectures are accompanied by new models of computation such as computing with custom-precision and different number formats which can deliver good-enough results at higher efficiency. Dataflow computing models benefit from minimal data movements as well as deterministic execution. AI algorithms can help to derive models for HPC applications from large amounts of data. All of these technologies provide unique benefits but they can also pose challenges because they are non-standard. In order to harness these technologies effectively we need support for better tools for software development and porting as well as increased efforts for community building, training and education.

INTERSECTION WITH "SUSTAINABILITY" RESEARCH CLUSTER

Clearly, there is the expectation that unconventional approaches presented here, as well as many others that are being explored, will also contribute to lower energy consumption when applied to HPC. The adoption of such unconventional approaches means an increase in specialisation and hence heterogeneity in the infrastructure. This poses new problems, some of them are covered in more detail in the preceding chapters, as well as others that will require attention too:

First and foremost, an increased specialisation of accelerators is expected to exacerbate issues related to sustainability unless targeted efforts are undertaken towards code portability, hardware and software modularity, etc. This is covered to some extent in the preceding chapters.

As a consequence, there will be a growing need to find out how and where to deploy increasingly complex and changing HPC workloads, also in an energy-efficient manner. In particular, this includes the exploitation of HPC resources that are made available also from various cloud providers, each with potentially unique options exploiting the unconventional architectures and models presented here. This means finding the best fit of a provider and configuration of its resources at a given time for a desired workload, without doing an exhaustive exploration, using automated approaches with predictive and/or search-based methods. This is an area of active research that will benefit the overall sustainability of unconventional approaches in on-premise as well as cloud environments.

28. Mirhoseini, A., Goldie, A., Yazgan, M. et al. A graph placement methodology for fast chip design. Nature 594, 207–212 (2021). https://doi.org/10.1038/s41586-021-03544-w

6.2.10

Quantum for HPC

RESEARCH TRENDS AND CURRENT STATE OF THE ART

Quantum computing describes a disruptive way of computing based on the principles of quantum mechanics with the potential to accelerate a selection of applications. Being a disruptive computing technology, the integration of quantum computers and simulators (QCS) in high-performance computing (HPC) systems is one of the challenges for 2023 - 2026, as indicated in WG 1 (System Architecture). From an HPC perspective QCS will be accelerators to HPC clusters. While the EU Quantum Flagship encompasses qubit development efforts and creates QCS technologies, the EuroHPC JU is meant to acquire and deploy those technologies, stemming from the Flagship or other efforts, which have reached a sufficient maturity, to support their deployment and their integration in HPC environments and to prepare software tools, libraries and applications.

QCS devices need to be integrated in HPC systems:

- on a system level, where QCS technologies need to be integrated in HPC clusters;
- on a programming level where disruptive ways to program the devices mean that a full hardware-software stack needs to be built;
- on the application level where QCS promise disruptive changes in the complexity of some applications so that compute intensive or intractable problems for HPC might become tractable in the future.

Currently we see an era of quantum enablement in which no quantum advantage for an industrially relevant application has yet been achieved. The race to scalable error-corrected quantum computer hardware is still open and we will see multiple QCS technologies in use, at least for an intermediate stage of development.

Three application areas are commonly attributed where quantum advantage is likely to be first demonstrated: optimisation, quantum chemistry and quantum machine learning. Benchmarks (application-oriented) are needed to assess the properties and capabilities of the different technologies. Many quantum algorithms, such as variational algorithms, are hybrid and divide the problem into a classical and a quantum part. Therefore, an effective integration of quantum and classical computers is important.

There are several options for the integration of HPC and QCS, and these may coexist in the medium term, until applications with quantum advantage and workflows between classical and quantum computers have been developed. New challenges concerning the hardware-software stack for the deployment and the emulation of QCS will arise for larger quantum systems.

Education is the cornerstone on which new generations of quantum aware scientists will be trained since quantum computing requires a new approach (or mindset) to application development. This activity is essential for the long-lasting growth of quantum computer usage and is further detailed in the Education and Training chapter of this document.

CHALLENGES FOR 2023-2026

APPLICATION-CENTRIC BENCHMARKING

Despite various initiatives, no general and widely adopted benchmark suite is currently available. Beyond gate or circuit level protocols, application-centric benchmarks are necessary since benchmarking at the algorithm level tests *ipso facto* the combination of the algorithm, the software stack and the technology. However, at the current stage of technology and quantum algorithm development, it is not clear which algorithms will show quantum advantage and are of special interest. Likely candidates will come from the prominent application areas such as optimisation, quantum chemistry and quantum machine learning. Therefore, the characteristics of the most important benchmarks are not yet known and, at the current stage it is important to promote the development of a broad set of benchmarks in close cooperation with the application communities, so that we do not prematurely lock ourselves into a particular technology or algorithm.

Progress in application development and experimentation should form the basis for the definition of benchmarks, i.e., benchmarks will play an important role in the hardware-software co-design, as defined in the application co-design section. Such a co-design approach is essential. It is also important to develop a widely acknowledged ranking system that is mature and flexible enough to accommodate future technological developments.

SOFTWARE-HARDWARE STACK

In the current era of developing and deploying QCS, most attention is drawn to the development of the underlying hardware technology. To use the hardware, a dedicated software stack is necessary to bring-up and run QCS (including calibration, characterisation, optimal control, instruction set architecture, device simulation), to operate the systems (including control electronics, scheduling, authentication, recalibration, recharacterisation) and to translate the user's code to run on the systems (from high level languages and algorithms to quantum circuit tablature which can be executed on the hardware). Although most parts of the software stack exist today it becomes clear that most software components can still be significantly improved. Here standardization can be key in order to allow a fixed interface design between the software components. Open APIs, widely adopted by the quantum computer industry, will facilitate a modular approach to the implementation of the software-hardware stack. It will allow different parts of the quantum solution to evolve independently, using best of breed from different vendors and preventing technology lock-in. In addition, new challenges will occur with larger QCS, e.g., optimal mapping of an algorithm on realistic QCS is an optimisation problem in itself.

Hybrid workflows may require multiple interactions between the classical and quantum parts. If the quantum computing system is a shared component, then algorithms might have to wait a significantly long time (each iteration) to be scheduled on the quantum processing unit (QPU) of the QCS. For example, to reduce the waiting time at every iteration of the algorithm, new programming models are emerging that bundle the classical and quantum parts of an algorithm and execute them as a tightly integrated loop, without having to queue for the QPU each time.

SYSTEM INTEGRATION

We may not know the eventual form of hybrid quantum-classical applications, but in the current era it is clear that algorithms need to be hybrid to use the current QCS. The overall efficiency is dependent on reducing the latency between the quantum and classical steps. The tight and loose integration models are the two ways that can be used to address this low latency requirement.

In the loose integration model, the QCS are not co-located with the HPC system. Latency-critical applications will require that sufficiently powerful standalone servers are directly connected to the QCS. These servers enable a tight, low-latency feedback loop between classical and quantum parts of a specific algorithm or application, and connect QCS and the HPC system over a high-speed long-distance network connection. Especially for QCS in an experimental stage, this distributed approach enables the set-up of classical and quantum resources in locations that are well-suited for each of them.

In the tight integration model, low-latency communications are key between the QCS and the HPC system. Hardware and software requirements for the QPUs of the QCS can be somehow combined within the hardware and software requirements for the CPUs and/or GPUs of the co-located HPC system, to achieve a single tightly integrated system, in which the different hardware components are connected via a federated high-speed network. Such a combination is provided, for example, by the Modular Supercomputer Architecture (MSA). In cases where the technological features of the QCS technologies are sufficiently mature, this co-location approach may be appropriate. For some QCS hardware technologies, it may even be possible in the long run to integrate QPUs directly into or close to the classical chips, as it is currently the case for GPUs. The latter tight integration model will require multiple and significant advances in the technology currently used.

EMULATION OF QUANTUM COMPUTERS WITH HPC SYSTEMS

For emulating large QPUs, supercomputers must be used because of their extensive storage and parallel computing power. On the one hand, a QPU can be emulated to obtain a system of fully connected qubits with infinite coherence time and perfect fidelity, without errors in the application of quantum gates, or on the other hand, a realistic quantum computer can be emulated. This makes emulators an irreplaceable tool for designing, analysing and benchmarking quantum algorithms, not to mention the ability of thoroughly studying new qubit technologies. An advantage is that the emulated circuit can be traced to observe, for debugging purposes, the state of the circuit at any point of its execution, which is not possible on the QPU itself.

The vast majority of available emulators, even if they are open source, have been developed by the quantum computer manufacturers themselves. Also, many of them do not support HPC implementations such as distributed computing and/or GPU acceleration. The development of a (open source) fully parallel accelerated emulator, probably using the latest technologies such as NVM for expanded memory capacity, is the key to a fast adoption of quantum technologies.

INTERSECTION WITH "SUSTAINABILITY" RESEARCH CLUSTER

Quantum computing is an emerging technology and needs to mature in order to reap the benefits it promises. The actual energy savings by an algorithm/application with quantum advantage cannot be quantified yet. However, it becomes apparent that the energy consumption of most quantum technologies in use are orders of magnitude smaller than for classical systems (10-25kW). Once quantum advantage is reached, the speedup would lead to additional savings in energy consumption. Beyond energy consumption, it is important to mention that most technologies rely on rare materials and cryogenics. This is a malus in terms of sustainability. How the different factors will add up when quantum computing becomes widely available and used is not clear yet.

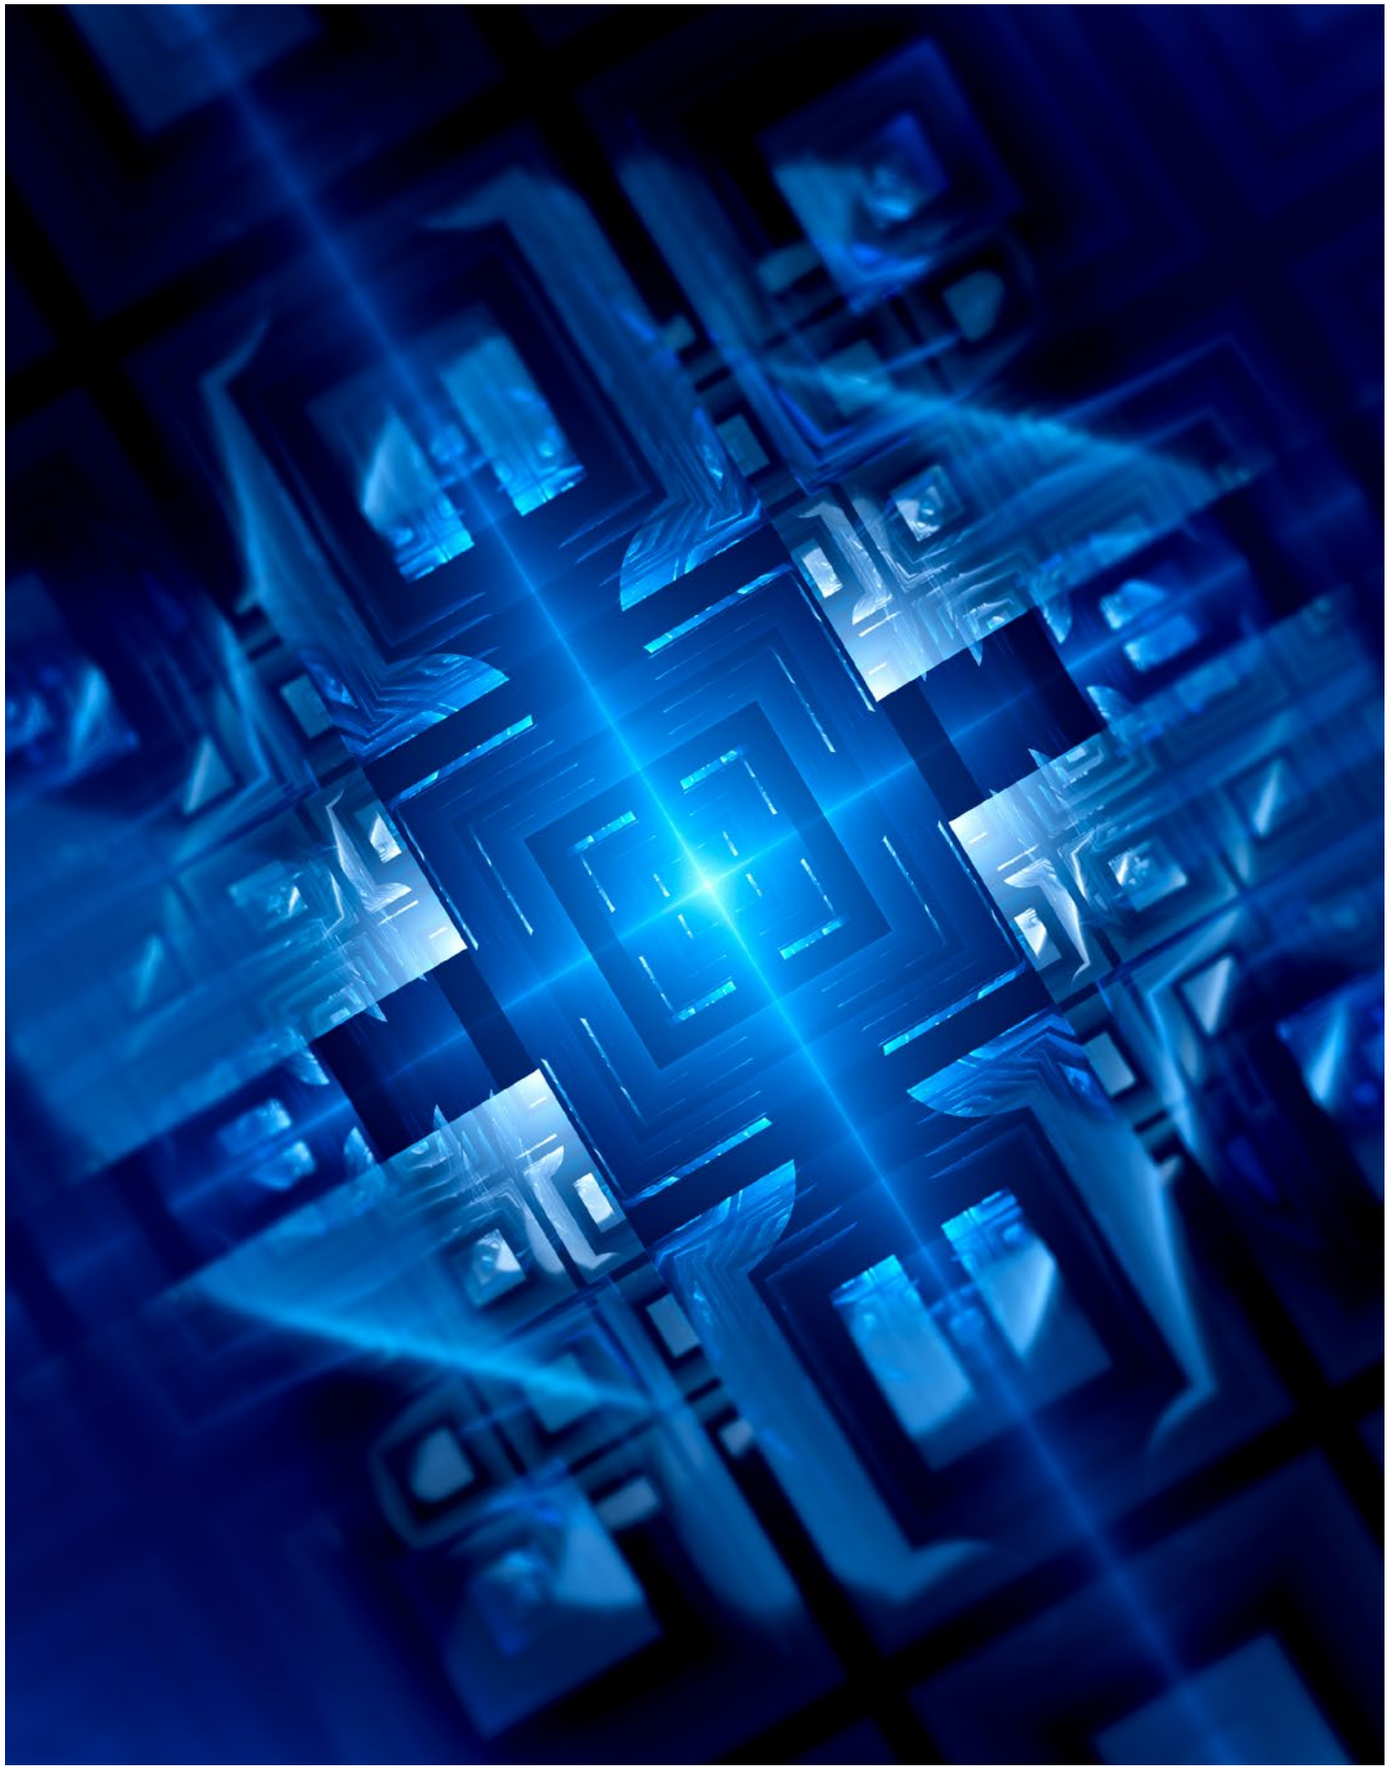

7

Open-source approach: a balanced view

7.1
Motivation

The main motivation for the use of open source is reduced acquisition costs, increased independence and control, increased flexibility, faster innovations, faster and better uptake and access by users and industry and building communities with "crowd intelligence".

Several examples exist where a mix of open source and closed source (proprietary) have provided good solutions. In addition, there are often pros and cons for the choice of open-source solutions. A decision on whether to use open-source elements should take into account a wide range of factors.

Using an open-source approach for building hardware (e.g., processors, accelerators) has gained a significant momentum over the last 2 years, worldwide. China, for example, views open source as an industrial policy tool and important part of its push for technological autonomy. Notably open source can be used as a sovereignty tool providing an alternative to licensing IPs.

The development of a strong European open- source ecosystem will drive competitiveness, enable greater and more agile innovation and give greater economic efficiency.

7.2
Definitions[1]

The term "Open Source" refers to solutions that can be modified and shared because their design and implementation are publicly accessible. Open Source projects, products or initiatives embrace and celebrate principles of open exchange, collaborative participation, transparency, meritocracy (lower risk of vendor locked solutions), and community-oriented development.

"***Open Source Software (OSS)*** *refers to software that is designed, developed, tested, and can be inspected, modified or enhanced through public collaboration, and distributed with the idea that it may be shared with others, ensuring an open (future) collaboration.*"

*"**Open Source Hardware (OSH)** refers to the design specifications of a physical object (e.g. electronic or computer hardware) which are licensed in such a way that said object can be built from design information, created, studied, modified, improved and distributed by anyone. Such information, made available for public use, can include documentation, schematic diagrams, construction details, parts lists and logic designs.*"

Software/Hardware designs and inventions are subject to copyright, patent and trademark law. Open Source uses these intellectual property laws creatively to make the respective designs publicly accessible.

7.3
Scientific software

There is a long tradition of using open-source software in the scientific community. The reasons for this are manyfold:

First and foremost, open-source software typically allows inspection of the source SW, which is important for the reproducibility of scientific results, for independent validation of the algorithms, and finding/fixing issues with the methodology or bugs in the software. Appropriate in-house knowledge provided, access to the source code also allows scientific users to adapt the software to their specific scientific problem by developing new features or tuning existing features. If such new developments are to be fed back to the main release, the user community benefits from new features being added by many contributors. Finally, open-source software is typically free of charge, which benefits the academic community where it is often difficult to obtain budgets for licensing fees, while personnel to support and develop the software is more readily available.

The downside of using open-source software in the scientific community is that if the users does not have development and support resources, it is often difficult to get required features into the software or obtain needed support. For large, widely used packages this can be mitigated by support provided by the user community and training offered by core developers and companies (sometimes even spin-offs from the original academic environment) providing professional support (and in some cases additional "premium" closed-source features).

With this professional support, open-source software is increasingly used by industry, who often require this level of support and may not do in house software development. This "freemium" business model is sometimes used by ISVs who open-source some of their products to get wider development exposure. Similarly, there is a trend in large industries with sufficient software development skills to contribute to open-source software.

A key issue for open- source scientific software is sustaining the required software engineering efforts. It takes a long time to build a well-functioning community with the required critical mass around a software package, and without a dedicated core team (be it from academia or industry as discussed above) it is often impossible to establish such a community. In addition, a well-functioning open-source project requires dedicated effort to manage the evolution of the software. This includes experts to decide which contributions should be added to a given release, quality control (code review, testing), and documentation. This essential effort it is often difficult to maintain due to lack of funding/manpower.

In some cases, scientific software is distributed through a mixed model, where a small license fee is charged but code can be freely modified and changed. The income though this model can then be used for the management effort described above. Still, the number of uses is in many cases not large enough to obtain and sustain a larger "core" team.

1. Further material at: opensource.com/resources/what-open-source and opensource.com/resources/what-open-hardware

In general, licensing is a key issue for the uptake of open-source software. Depending on the licensing scheme applied, adding open-source components to commercial products and workflows may result in difficult situations where legal experts are needed to understand the implications. It may also be difficult to provide the necessary assurance that license infringement is not a risk.

7.4

System software and management

Software in a supercomputing facility comes from a variety of sources: commercial vendors, open-source communities, as well as in-house development and adaptation. System software and management has benefited substantially by the availability and long-term evolution of relevant open-source components. With increasing adoption of open-source software, including use and contributions by major system vendors, the availability of trained personnel has increased, which increases the confidence in successful long-term deployments.

Here are some prominent examples of functionality for system management software:

- Job submission and control, including resource management and scheduling: Slurm, Torque, Flux
- Node and system-level health monitoring: Nagios, Ganglia, Prometheus
- Configuration automation and change management: Puppet, Ansible
- Low-level system initialisation and management: OpenBIOS, OpenBMC.

Turning to system software, prominent examples of open-source components include:

- Operating system: Linux kernel and distributions
- Filesystem: Lustre, BeeGFS
- Interconnects: OpenFabrics.

The Linux kernel stands on its own as an example of an open-source community with a very broad base of users and supporters, and a stable governance structure relying on a widely distributed hierarchy of long-term maintainers of subsystems, which are themselves directly or indirectly supported by stakeholders who are consistently investing in the kernel's sustainment and evolution. However, due to limited funding for system software and management in HPC, the sustainment burden rests with a small set of beneficiaries. In addition, fixing bugs and vulnerabilities is based on the surveillance done by the open-source community, which does not guarantee a high level of security relative to proprietary software. Thus, a number of concerns need to be carefully considered throughout the lifecycle including planning, development, integration, and sustainment phases:

- Licensing, with standardised well-reviewed licences used as models of best practice
- Open-source project governance frameworks for structured collaboration
- Support models, including available upgrade paths, in some cases oriented to reduce vulnerabilities
- Well-defined and well-documented APIs in a cohesive software stack, to simplify interoperability and exchangeability of software components.

The success of open-source software ultimately relies on community building centred around effective deployments to solve real-world challenges, and timely evolution to cover emerging new requirements. The existence and growth of a vibrant ecosystem is a strong indicator of the future prospects of open-source software projects. Such ecosystems compete for the engagement and retention of contributors and supporters.

On a closing note, open-source software encourages an increased level of attention to trustworthiness, starting from transparency of origin for critical software components. A functional definition of trustworthiness[2] in this context includes conformance to a checklist of criteria promoting transparency, accountability and enabling independent assessments. Such a structured framework then enables stakeholders to analyse comparative risks and make reasoned risk-benefit and resource-allocation decisions.

7.5

Open source-hardware

There is currently a large momentum, mainly in academia, around open-source hardware and especially around solutions based on open instruction sets like RISC-V. For academic use cases, RISC-V offers a licence-free, fully modifiable Instruction Set Architecture (ISA), ideal for experimenting with designs targeting various use cases. The recent attempted acquisition of ARM by Nvidia (which finally did not happen) and the supply chain problems impacting silicon delivery that arrived with the Covid pandemic have each contributed to the increase in calls for open HW for production silicon.

Open-source hardware could be a step towards European independence in the semiconductor and specifically the microprocessor domain. Open-source hardware is attractive due to the promise of free IP (no licensing fee) and the flexibility to modify cores in any direction desired. However, it poses a number of concerns that must be addressed if this approach is to be used in production silicon, and particularly for server microprocessors. The principal challenges include ecosystem, liability, and fragmentation:

2. "Creating a framework for supply chain trust in hardware and software", A Report of the Lawfare Institute's Trusted Hardware and Software Working Group, May 2022. Document URL: https://s3.documentcloud.org/documents/21831749/creating-a-framework-for-supply-chain-trust-in-hardware-and-software.pdf

- **End-use Ecosystem**: When used in research projects, embedded or acceleration applications, microprocessor or microcontroller software is generally encapsulated or even available as best effort. However, for a server processor, the majority of software is selected by the end customer and pulled from either open source or customer repositories. For end customers, the lack of support of one required SW element typically eliminates a processor from consideration due to the cost (time, effort, and money) of SW migration.

  Development of ecosystem support for a server processor ISA takes years, the involvement of many institutions, and considerable monetary investment. If we consider the development of the Arm server processor ecosystem as starting from the founding of Calxeda in 2008, it took 10-12 years for the ecosystem to develop to the point where it was self-sustaining, with many institutions contributing time and money to the ecosystem development (and many aspiring Arm based server processor companies giving up along the way). The ecosystem challenge is not insurmountable. RISC-V can follow a similar trajectory to Arm, perhaps in less time by following the template been established by Arm, but the ecosystem will be a significant hurdle to widespread competitive deployment of production server processors over the next decade.

  In addition to the end user ecosystem, the development ecosystem for open-source hardware can present challenges. This is particularly true for design verification. The RISC-V ecosystem is struggling to keep pace with rapid innovation and customization, which is increasing the amount of verification work required for each design. The historical assumption is that verification represents 60% to 80% or more of SoC project effort in terms of cost and time for a mature, mainstream processor IP core. Traditionally, both functional and performance verification of x86, MIPS, or Arm is done by companies like Intel, AMD, Arm and Arm architecture licensees, and MIPS. It is proprietary and homegrown, with processor companies using significant resources to develop internal proprietary solutions.

  RISC-V, meanwhile, has opened the door for companies from around the globe to modify the source code, which adds a whole new level of verification challenges for the companies making the modifications and for those implementing the modified cores in their devices.

- **Liability**: while the open-source IP is attractive due to lack of ISA licensing fees, it poses a risk for commercial silicon suppliers. ISAs such as Arm offer indemnity as part of the licensing agreement. If a patent troll decides to attack a successful microprocessor startup, Arm becomes involved as determined by the Service Level Agreements in the licensing contract. They bring their legal team and their war chests of patents, along with a large set of customers who also have war chests of patents and large legal teams (e.g. Qualcomm, NVIDIA, AWS or Microsoft). Open ISAs such as RISC-V cannot offer this same level of support. Open architectures can address this through consortiums, asking the consortium members to band together, pooling patents and legal funds, but until this is validated for a given consortium there is risk that patent trolls will find loopholes that could lead to disaster for young companies.

- **Fragmentation:** open hardware, as evidenced by RISC-V, offers a modular design where the ISA comes with a base set of functions, and an array of extensions that can be added. Extensions may be public and then ratified by the community, but in some cases companies or engineers, in order to protect competitive advantage, may not wish to publish extensions. The mix of proprietary and open extensions can lead to fragmentation, meaning that SW that runs on one companies' processor does not run on another. Arm avoids fragmentation by enforcing strict ISA guidelines, which must be validated by Arm, and creating a Base Architecture (initial called SBSA, now called ServerReady) verification model that can be run on Arm servers to guarantee a level of uniformity and cross platform SW support, and thus SW compatibility. Open architectures can develop base architecture standards and uniformity practices, and standardized verification practices, but this is more difficult for an open community than for a single company. The net result is that it will take time to solve the segmentation challenge.

Aside from the above listed elements, the cost savings offered by open-source hardware bears further scrutiny. For research projects, where the goal is to modify IP blocks and to explore architectures, open-source hardware can offer substantial savings. The other end of the spectrum is industrial companies developing large, complex processors: as mentioned above, the ecosystem of open-source hardware goes well beyond the processor core IP blocks:

- Even with an open-source ISA (cost free license), there are significant development costs: design, implementation, and tools to translate the implementation (Verilog or VHDL) are increasingly expensive at competitive design nodes. Once the design is complete, the GDS2 file is then sent to the foundry. This file is highly dependent on foundry specific libraries (PDK) that are generally proprietary. Mainstream CAD tools are proprietary and expensive.

- Regarding the design, the netlist of a processor core is not sufficient to make a chip - you need bus interfaces, memory interfaces, interconnect, Network on Chip, and other uncore elements. Some dedicated blocks, such as the physical interface to high-speed memories are complex to design and are technology-dependent with a short list of providers, making them high cost either to develop or to acquire.

- And once you have a GDS2 file, it needs to be transformed into real silicon by the foundry, which is expensive in advanced technology nodes (which in turn are needed for High Performance processors), then packaging and testing. The development of all the low-level software and libraries specific to the chips

are also large investments[3] in time and effort for commercial solutions.

This process of developing products with open-source hardware is far more complex than in open-source software where you can nearly clone the repository, compile the source code and have the final usable software (and iterate in a matter of minutes if you found and correct a bug).

Over the past several months, there are several working groups trying to explore what can be done in the field of open-source hardware in Europe. As an example, the report: "Recommendations on Open-Source HW & SW, incl. RISC-V, for European Sovereignty"[4] tries to set the scene and introduces all elements required to start an active ecosystem around open-source hardware in Europe. The developments in the field of HPC should be synergetic to the global European action in this context.

All investment-, roadmap- and prioritisation decisions in the context of open-source hardware components as part of the European Chip Act are expected to be handled by the KDT JU. EuroHPC should contribute by delivering design requirements specific to the HPC systems market.

At this point, the maturity of open-source hardware solutions is not yet at the level reached by HPC mainstream processor providers, and it will take significant time and investments to be competitive versus mainstream HPC processors.

Again, embedded processors and accelerators sidestep many of these issues as they work in a closed SW ecosystem, and so offer an excellent starting point for lower risk open HW solutions.

3. On the final bill, the licensing cost of commercial cores, such as ARM, is often below 10%
4. https://www.inside-association.eu/post/securing-european-sovereignty-key-recommendations-for-open-source-hardware-and-software.

8

Building and retaining skills and competence

8.1

Education and training: a lifelong journey

Little has changed over the last 100 years in how people are trained. The education system was first modelled in order to fulfil the need of providing skilled people to support the industrial revolution and the growing government's administration. That model assumed that people could have lifelong jobs (or move on to similar jobs) for which they could be trained once-off at a young age through academic-style lessons and domain-specific books. It also implied that domain-specific knowledge was independent from other domains and only changed slowly. Nowadays, the evolution of science and technology goes at an extreme speed. Domains are highly interconnected. We all have the whole knowledge developed in the previous millennia at our fingertips thanks to the Internet and mobile devices. However, relying only on the current model of education will not support the growth of the European economy in the long run.

Education and training should be a lifelong and continuous process. The process begins with the traditional, academic-style 'transmissive' approach where the knowledge accumulated up to now is passed on to the next generations. Then it should continue lifelong with formal/informal updates and refreshes covering the technological advances in scientific and technological domains. It should also teach people to think more in term of 'solution-finding' than of 'problem-solving' (e.g., human-computer interactions, critical thinking, creativity, what-if-analysis, teamwork, decision-making, cognitive flexibility, entrepreneurial mindset, IPR management). This process will create a workforce with a broad base of general supporting knowledge and skills. This base will be supplemented with up-to-date knowledge and skills covering a wide spectrum of domains, and - first and foremost - allow to create new knowledge and to innovate. Specifically in the case of HPC, society is increasingly dependent on compute-intensive applications (whether used for AI, weather forecasting, drug design, smart cities). Graduate profiles are needed to support the development and growth of a truly European ecosystem and for driving the digital transformation.

Skills are part of a holistic concept of competency, involving the mobilisation of knowledge, attitudes and values to meet complex demands. To remain competitive, workers in IT in general and HPC in particular will need to enter in the job market well-trained and acquire new skills continually. This requires some flexibility, a positive attitude towards lifelong learning and a fair amount of curiosity.

In the area of HPC (in the wider sense of techniques and methods for efficient intensive data/compute processing in the digital continuum), we see three main axes to guide reflection and action:

1. Educate and train people to leverage and harness existing technologies

Traditional, academic education should provide the initial training enabling the knowhow for exploiting the current technologies, giving people a foundation of methods to continue learning during their career, on the job. For this latter purpose, there is a whole complex of formal and informal training by companies and organisations, sometimes even providing certifications in specific skills. This is the bulk of further education after the initial training. Nowadays, several opportunities for self-directed or incidental learning are also available on the internet in the form of short movies, animations, games, demos, etc. Academia should complement their traditional curricula with non-formal learning material (videos, MOOCs, etc). Organisations should allow their employees to take time-off to upgrade/update their skills.

2. Educate people to create, design and develop new technologies and solutions

The new jobs will not only require sound domain-based skills and a continuous refresh/update but also problem-solving skills, creativity, entrepreneurship, and a think-out-of-the-box approach. The process of designing a new technology or an innovative solution includes much more than just coming up with a bright idea. Possible limitations and/or constraints on the design must be considered; safety and social impacts are important; etc. Teaching such skills is difficult to frame in an academic curriculum or in a video. However, academics and organisations should nurture them in their lessons/workforce and teach how to devise a not-yet-existing solution to the problem it is intended to solve.

3. Capitalise and transfer existing tacit knowledge to new generations

The transfer of tacit knowledge requires extensive personal contact, regular interaction, and trust. Usually, tacit knowledge is embedded (and sometimes invisible...) in organisations' processes and procedures, as a result of several years of continuous improvements and generation-over-generation of staff working side-by-side. This is not the kind of knowledge that can be transmitted in a classroom or via online meetings. Organisations should develop and implement processes to organise and disseminate the tacit knowledge of their experts, which, without such processes, can be diluted and lost. While there are no general recipes for the best way to transfer it, showing your work seems to work in most cases - making one's knowledge visible via shadowing and do-it-together with the intention of transferring the tacit knowledge embedded in how work gets done. By sharing enough information about how you get the work done, others can begin to sense, recognise and acquire the tacit knowledge behind the accomplishment[1].

In conjunction with these three axes of actions, we have identified a few specific domains in which we believe the training will play a particularly crucial role, and which we want to highlight.

1. Show Your Work, Bozarth (2014)

8.1.1

Entrepreneurship for successful commercialisation of HPC

Usage of HPC keeps broadening (from genuine numerical simulation to AI, from edge to cloud). As such, the need for bringing innovative commercial approaches with the definition and bundling of next generation services is required. In that context, training students and professionals in terms of business innovation, standardisation, and IP rights becomes a must.

New generation of HPC *technopreneurs* shall be familiar with standardisation bodies and related technical workgroups on key topics. They will need to understand the way standardisation is evolving, and master IPR (patents, copyrights), including the impact of choosing the right open source models, new platforms, technologies and applications (e.g. quantum computing, AI, etc.); and understand the way these different elements will help developing next generation of solutions and services.

8.1.2

HPC in industry requires continuous training of the workforce

With the Europeanisation and the globalisation, industries are more than before concerned by High Performance Computing. For the last two decades, more and more industries have been using HPC for shortening time to market, prototyping or increasing the quality and innovation in the product or production chain. More recently, industries are facing new challenges to digitalise their production sites to another level (digital twin). This reinforces the need for massive computation power to train artificial intelligence algorithms in order to optimise and prevent work-chain and workflow from failure or maintenance.

The need of numerical simulations and artificial intelligence at scale is no more just a need but a must have to maintain a competitive advantage.

The fundamental challenge for companies that need to resort to HPC is to find employees with enough knowledge and competences in the domain as the number of fields covered by HPC are becoming broader and broader. Therefore, training and education must respond quickly and adapt to follow the market moves and the technologies.

Industries are looking for qualified people who have enough basic blocs (education) to be trained quickly on new technologies through webinars, online courses, or even on site with instructor-led training. Those conditions are mandatory for companies because they allow them to apply agile methods within their IT, R&D and engineering departments.

It became clear that shorter and specific training sessions are demanded by industries to offer more flexibility to the employees, and also to cover a broader range of topics.

Let us note that what is stated in this section mostly stands for HPC users in research and academia as well. Industry, however, has different tactics and approaches when it comes to HR management.

8.1.3

Structured train-the-trainers programmes are necessary to keep up with demand

A fundamental challenge in today's HPC landscape is the extreme diversification and broadening of the applications that are now considered - in fact, the definition of HPC itself is undergoing a tremendous shift. Whereas before, knowledge of MPI parallelisation and one programming language out of three (C/C++/Fortran) would be sufficient for years at a time, requirements today include heterogeneous hardware/accelerators, AI methodologies, and emerging technologies such as quantum and neuromorphic computing.

With the IT workforce being under heavy pressure by a general trend towards digitalisation in the society, it becomes harder to find and retain qualified personnel to provide adequate training in the relevant HPC areas. The need to "train the trainers" has thus become a continuous quest for all providers of HPC training. Together with the considerable increase in the number of fields covered, more structured training programs are necessary that allow scientists and engineers to sustainably build up their qualifications. At the same time, it is essential to remain flexible regarding the offered training opportunities, to allow quick incorporation of emerging topics in the course programme.

8.2

SMEs are different from large companies

Large companies can afford a time-consuming periodic training for their workforce. Employees of SMEs often do not take advantage of continuing education because they or their supervisors do find it too expensive or time-consuming. This also applies to cheaper offers, or even to offers that are free of charge, because the working time loss the training implies also has an important impact.

Usually, the structures of SMEs are completely different from the large industries. In SMEs training is seldom anchored in the organisation. Therefore, the awareness of the importance of the topic is often not present nor understood. In addition, it is important that information on training offers is provided, e.g., via newsletters, social media, etc., as this information is too often lacking.

Survey results showed that shorter training formats are needed, as especially in SMEs, there is little time for periodic training. SMEs primarily demanded the following formats: workshops and seminars (1-2 days), web-based training modules as well as webinars and virtual classroom ones.

It became clear that there is a high demand for learning methods that are independent of time and place. SMEs can especially benefit from blended learning concepts as they can cover a broad range of different topics.

Virtual events and self-paced training will continue to play an important role after the Corona pandemic, as they offer learners more flexibility and thus also save costs. However, face-to-face events are still considered important, especially for networking and informal exchanges. SMEs are often particularly dependent on local offers. At the same time, they particularly benefit from local networking, both with competence partners and with other participants in their regions. Enabling ad initiating innovation awards could also help industry identify and find skills tailored to their needs, while building cooperation between industry and education and training institutions. Therefore, as the transition towards more and more on-line practices progresses, this informal networking, which is too often missing, should be offered and encouraged between the training sessions.

Finally, yet importantly, education and training is an international challenge. We are at risk that Europe is very much inwards looking, while there is a lot of stimulation from what is going on in other places in the world. Europe needs to attract the best people worldwide – training and education is one of the key possibilities to achieve this, in addition to deploying and exploiting state-of-the-art HPC infrastructures and resources, and supporting vibrant academic and industrial R&D in HPC. ■

9 Operational recommendations

The following two sections are intended to make a few operational recommendations to EuroHPC on various implementation options for future work programmes.

9.1

Research projects implementation options

The R&I processes deployed in industries and some national R&D initiatives could mirror the following structured end-to-end approach: a staged sequence of projects with a specific end-goal and intermediate checkpoints governed by sets of KPIs (Figure 6). It should be noted that projects can be started at any of the following stages, depending on the entry - TRL and maturity of the proposed solution.

Mapped onto the next 7 years Multi-annual Financial Framework, Horizon Europe, this approach could work as follows[1]:

1. Proof-of-concept, exploratory projects

- Their scope is to validate high potential proposals, quantify their technical and business potential and risks, and engage critical skills, partners and stakeholders. Also, this early validation of new concepts is the main intent of the 'working agile' style of conducting projects in many R&D organisations.
- A potentially large number of small (budget) projects with a relative short duration (ca. 12 – 24 months).
- Small and focussed projects with low number of contributors
- The expected Technology Readiness Level (TRL) is expected to be low to medium
- 100% funding

All projects will be evaluated against the first set of KPIs and only those (few) passing the test will have a chance to enter the next stage:

2. Research and innovation projects

- Their scope is to produce solutions to specific problems, generate IP and demonstrate the solution in a well- defined, limited context.
- A medium number of medium-size (budget) projects with a duration long enough to produce tangible results.
- In most cases medium and in some cases high TRL. While most projects will only be able to reach a medium TRL, specific research projects (most likely in the software area) solving important problems can be expected to exit with a high TRL. Wherever advisable, projects need to be set up with a high degree of collaboration to avoid the generation of isolated "result-islands".
- 50% funding for large companies, 100% funding for SMEs and academic partners

All projects will be evaluated against the second set of KPIs focussing on the technical feasibility, robustness, efficiency and usability will have a chance to enter the next stage:

3. Large scale integration/ demonstration projects

- The scope is the integration of high TRL technology (hardware to software) into complete solutions in a pre-commercial, problem solving and application-driven environment.
- The number of these projects is low, the KPI-driven entry criteria are very demanding, the funding envelope is larger than in the case of the other two project types.
- Small set of partners, all indispensable to reach the completed operational prototype
- The funding focusses on material and component costs or license fees, personal costs are covered only for SMEs and potential academic contributors

4. Incubator projects based on a Venture Capital approach to promote and test the market opportunities and market readiness of fully developed and tested, very high TRL product-ready solutions

- Independently of the cascade of three project types mentioned above, the forming of start-up companies introducing new technology, methods, tools or business ideas involving HPC should be supported in the same way as strengthening "young" SMEs in getting their businesses off the ground.
- This project type should facilitate the market-entry of Start-ups and market-expansion of SMEs and thus help their future business investments in Europe.

1. This proposal focusses on technology development.

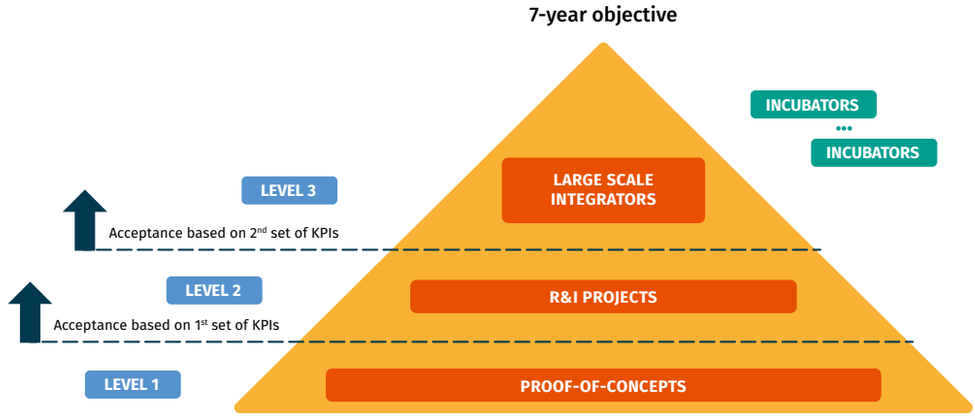

Figure 19: Conceptual structure of project types flow.

This staged approach from level 1 to 3 in Figure 19 applies mainly to projects pursuing new ideas or concepts with no preceding work in earlier projects. Given that there are numerous project results available today, future R&D projects using any of these results (potentially targeting a higher TRL level) could enter at level 2 or 3, depending on the alignment with the entry criteria.

And finally, a few challenges reported by current projects:

- Gap funding: projects undergoing an acceptance-evaluation for potentially entering the next higher level need to be kept intact on a low activity level with operational consortia, prepared to enter the next phase.
- Access to physical HPC resources should be part of calls issued by EuroHPC. It is nearly impossible to demonstrate for projects type 1, 2 and 3 the validity of the results at scale w/o access to large scale HPC compute resources.
- Raising the TRL level to TRL 8 or 9 requires multiple validations by typical potential users (SMEs and Companies in general), but today's restrictions on usage of real-life data are a real barrier to validation and adoption. Reasons are multiple, but these barriers need to be removed to progress realistically to TRL9.

9.2

Managing the next 7-year work plan, work programmes and calls

As pointed out in the previous section, a different approach would be required to increase the effectiveness the work programmes and calls. Instead of the bottom-up method, a programmatic, top-down, active program-management should be instantiated:

- First, a new long-term objective for HPC related R&D in the years following the availability of Exascale systems is required. What is the one, major goal all work programmes and calls should be aligned to?
- Second, the full spectrum of projects outlined above should be exploited, including a strict execution of the stages outlined in Figure 6.
- Active management of the entire portfolio-coverage of R&D topics and priorities should be enforced. In case certain areas of the priorities do not find enough attention and interest, these "white spots" must be taken care of. There are multiple options for "filling the holes" (e.g., through special incentives or relaxed acceptance criteria for projects).
- Facilitate the transition to commercialisation. Most European HPC hardware or software vendors in Europe are SMEs. In line with EuroHPC's objective to strengthen the indigenous industry, supporting start-ups and SMEs should be of crucial importance. It needs to be carefully evaluated whether existing instruments could be utilised or new ones (e.g., tailored to the growing needs of HPC functionality in the Digital Continuum) should be put in place.

The suggestions made in this chapter should be considered as possible implementation scenarios that would address the fragmentation seen in the past and achieve the streamlining needed on the path to a robust European HPC technology provisioning capability. ■

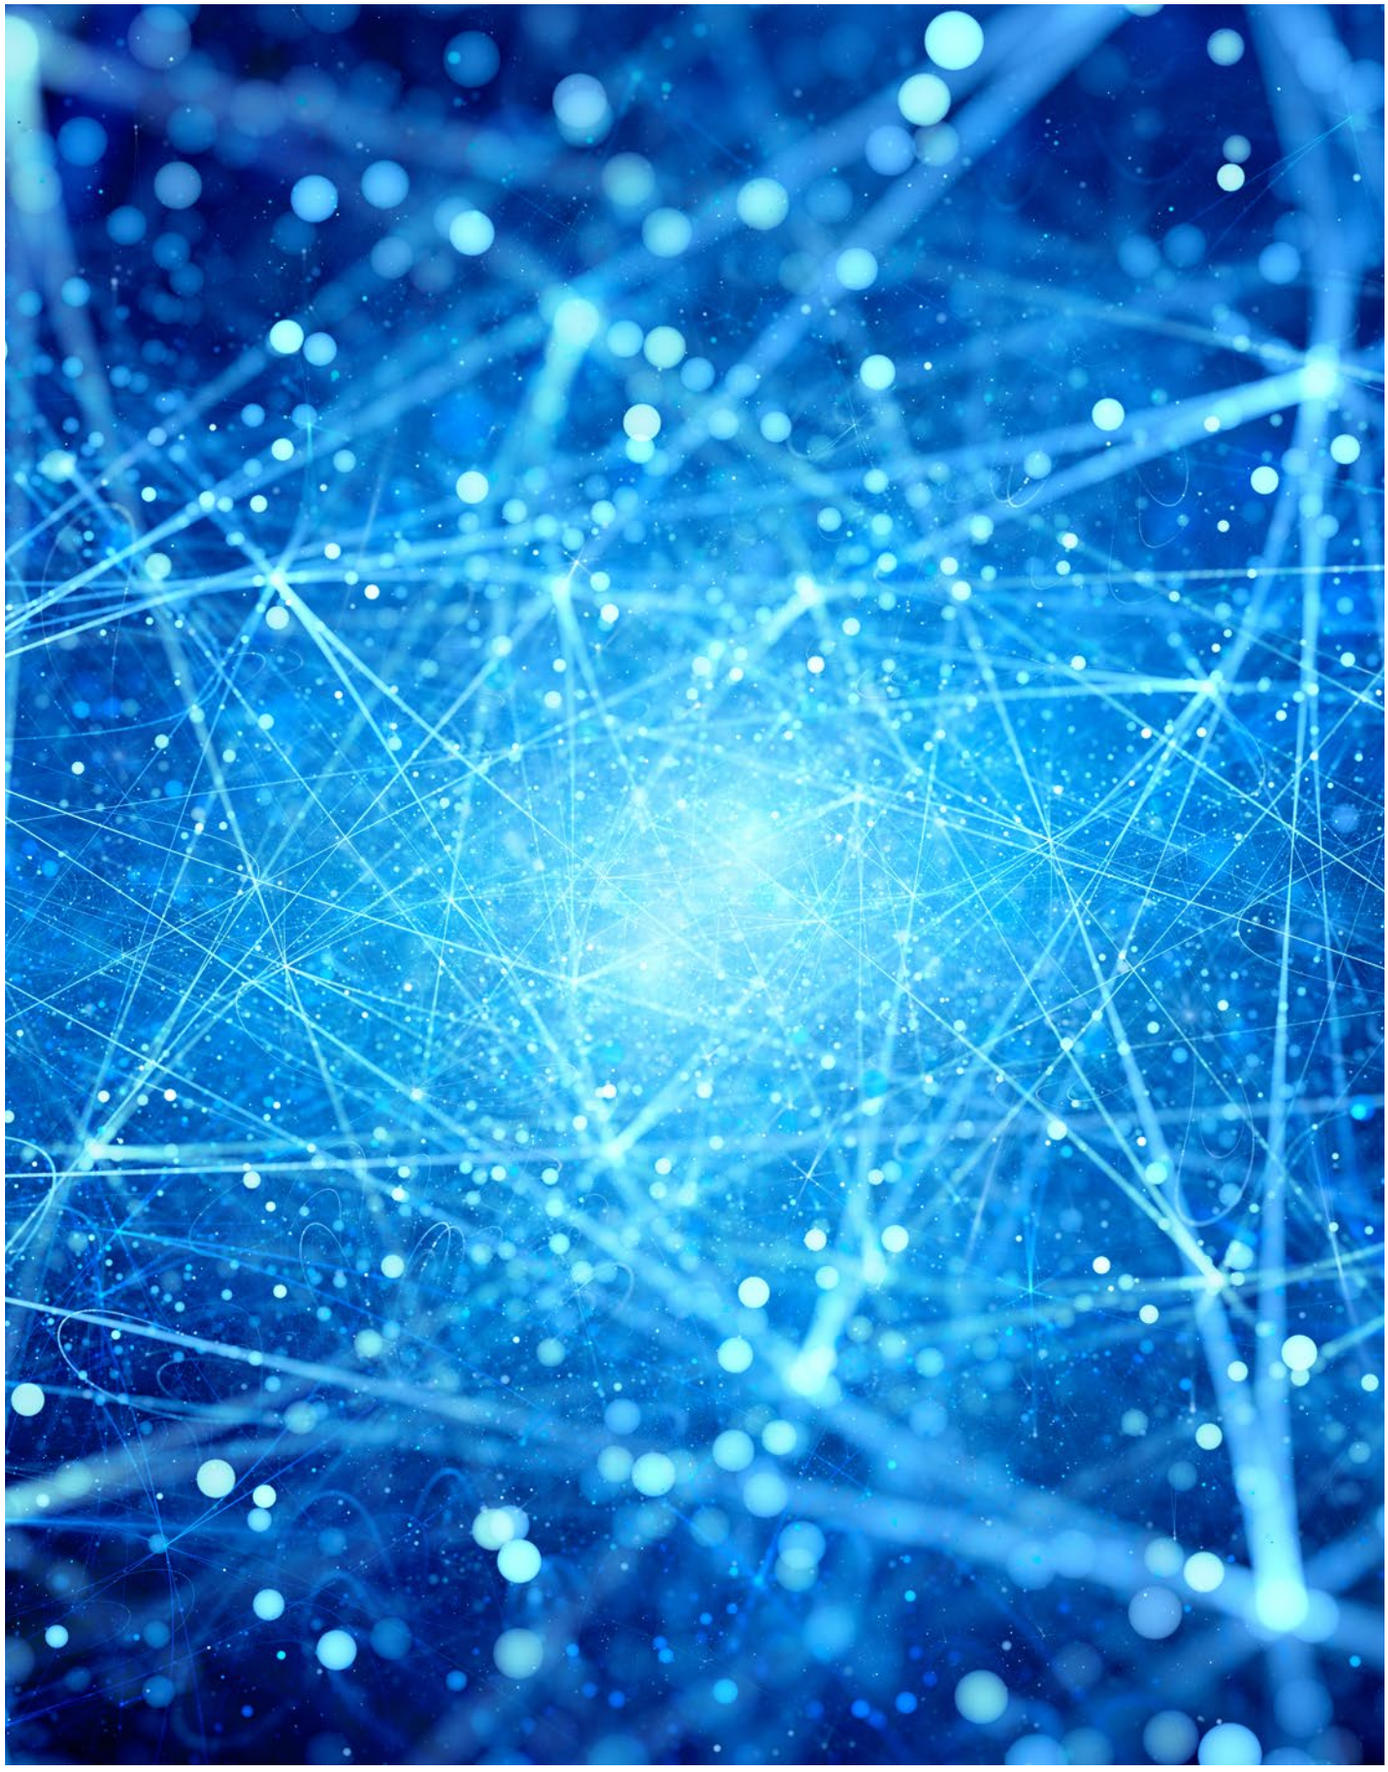

10 Conclusions

10.1
Sustainability

The challenge for the upcoming years is to efficiently deploy extreme-scale HPC infrastructures in a sustainable, environmentally compatible way and continue the R&D efforts to address diverse demands posed by application domains such as urgent-computing or modelling in the digital continuum. In the light of numerous initiatives and political directions issued by the EC and the EU member states, the HPC community is delivering its fair share. On the one hand, a significant number of HPC use cases aids the understanding and modelling of human-generated environmental impacts, e.g. the "Destination Earth" Initiative. Thus, HPC is an indispensable means to reach the environmental objectives formulated in e.g., in the Green Deal[27] communication of the EC.

On the reverse side of the coin, HPC-centres use IT resources as any other data centre. All concerns as e.g. the CO_2 footprint of the development and production of the infrastructure and during the use and operation of the centre as well as the use of rare material apply. While HPC centre resources are only a subset of all data centre/cloud resources worldwide, any possible improvements in the design & use of HPC system in data centres and, increasingly at the edge, should be undertaken.

Similarly to the "The TransContinuum Initiative" which was launched on the occasion of the previous SRA, ETP4HPC has now initiated an effort of a similar calibre targeting the issue of Sustainability: a new working group has been initiated with the following objectives:

- Describe the main factors influencing "Sustainability in HPC" in a holistic approach:
  - CO_2 – footprint of R&D, productions, operation and recycling of HPC systems
  - Approach for tackling system life-cycle extensions, re-use, circular economy, etc
  - Technical challenges and suggested solutions for efficiency improvements in the HPC stack
  - Challenges and efficiency improvements on the application level
  - Status of sustainability requirements in HPC procurement policies and hardware supply chain certification
  - Showcases of "green data centres", heat – re-use and limiting factors for further distribution

  This holistic view will be produced in close collaboration with think tanks that operate in the environmental research domain such as "The Shift Project"[28] or the "Borderstep Institute"[29] as well as other stakeholders in the European IT ecosystem like the partners in TCI (see chapter 3.2).
- Generate an R&I roadmap focused on sustainability driven research priorities listed in each of the Research Domain chapters.
- Act as catalyst for interaction between players inside and outside the HPC IT space (e.g., policy makers)

10.2
Other aspects

For several years, HPC has been occupying an increasingly prominent place among the solutions influencing our daily lives: it supports scientific applications in many crucial sectors such as drug development, climate modelling, materials research and many others. A new concept of Digital Twins is now the foundation of many applications in scientific and industrial use scenarios[30]. Furthermore, core HPC technologies and methodologies are being used to enable concurrent processing to permeate all levels of the digital computing continuum. The pace of innovation on all levels of the HPC stack seems to be increasing together with the growing demand for efficient, large-scale modelling, simulation and AI-centric projects. AI based methods are well integrated by now into the software-layers managing IT infrastructures.

Optimisation to suit different use scenarios drives an ever-increasing heterogeneity in system architectures culminating in the adoption of quantum-based accelerators which makes the effective management of these structures even more challenging.

In the context of boosting the European sovereignty, it will be important for EuroHPC to apply effective, holistic methods in the definition of R&I calls, manage the spectrum of technical focus areas across the calls and ensure that funding is applied effectively on successful projects including their transition to commercially viable solutions. ■

27. https://ec.europa.eu/info/strategy/priorities-2019-2024/european-green-deal/
28. https://theshiftproject.org/en/home/
29. https://www.borderstep.de/
30. https://zenodo.org/record/5470479#.YjBBvC9XZaQ and: https://zenodo.org/record/4683451#.YjBB5y9XZaQ

11

Appendices

11.1
Glossary

ACID - Atomicity, Consistency, Isolation, Durability

AI - Artificial Intelligence

AIOTI - Alliance for the Internet of Things Innovation

AISBL - Association Internationale Sans But Lucratif (International Non-for-Profit Association)

ALU - Arithmetic Logic Unit

AMBA - Advanced Microcontroller Bus Architecture

API - Application Programming Interface

AQMO - Air Quality and Mobility Project

ASIC - Application-Specific Integrated Circuit

AWS - Amazon Web Services

BD - Big Data

BDA - Big Data Analytics

BDC - Backup Domain Controller

BDEC - Big Data and Extreme-Scale Computing

BDV - Big Data Value

BDVA - Big Data Value Association

BSC - Barcelona Supercomputing Center

CBRAM - Conductive Bridge RAM (Random Access Memory)

CCIX - Cache Coherent Interconnect for Accelerators

CEA - Commissariat à l'énergie atomique et aux énergies alternatives

CERN - Conseil Européen pour la Recherche Nucléaire

CGRA - Coarse Grain Reconfigurable Arrays or Coarse Grained Reconfigurable Architectures

ChEESE - Centre of Excellence In Solid Earth

CIFAR - Canadian Institute For Advanced Research

CMOS - Complementary Metal-Oxide-Semiconductor

CoE - Centre of Excellence (for Computing Applications)

cPPP - contractual Public-Private Partnership

CPS - Cyber- Physical System

CPU - Central Processing Unit

CSA - Configurable Spatial Accelerator

CSA - Coordination and Support Action

CSP - Cloud Service Providers

CXL - Compute Express Link

D - Deliverable

DB - Database

DC - Direct Current

DCP - Digital Continuum Platform

DDR - Double Data Rate

DG - Directorate General

DL - Deep Learning

DOE - Department of Energy

DoW - Description of Work

DPU - Data Processing Unit

DSL - Domain Specific Language

E2E - end-to-end

EC - European Commission

ECA - electric, connected, autonomous/automated

ECC - Error-correcting code

ECG - electrocardiogram

ECMWF - European Centre for Medium-range Weather Forecasts

ECSO - European Organisation for Cyber Security

EESI - European Exascale Software Initiative

EIT - European Institute of Innovation & Technology

EMIB - Embedded Multi-Die Interconnect Bridge

ENES - European Network for Earth System modelling

EOSC - European Open Science Cloud

EPI - European Processor Initiative

EPOS - European Plate Observing System

EsD - Extreme-Scale Demonstrators

EU - European Union

EXDCI - European Extreme Data and Computing Initiative

FaaS - Function-as-a-Service

FET - Future and Emerging Technologies

FFT - Fast Fourier Transformation

FLOP - floating point operations

FORTH-ICS - Foundation for Research and Technology - Hellas

FP - Floating Point

FP - Framework Programme

FPGA - Field Programmable Gate Array

FTRT - Faster Than Real Time

FUSE - Filesystem in Userspace

GDP - Growth Domestic Product

GDPR - (EU) General Data Protection Regulation
GENCI - Grand Equipement National de Calcul Intensif
GNU - GNU's Not Unix
GPFS - General Parallel File System
GPGPU - general-purpose GPU
GPU - Graphics Processing Unit
H2020 - Horizon 2020 - The EC Research and Innovation Programme in Europe
HBM - High-Bandwidth Memory
HDD - Hard Disk Drive
HDR - Enhanced Data Rate
HiPEAC - High Performance Embedded Architectures and Compilers
HMC - Hybrid Memory Cube
HPC - High-Performance Computing
HPCG - High-Performance Conjugate Gradients
HPDA - High-Performance Data Analytics
HTC - High-Throughput Computing
HW - hardware
I/O - Input/Output
ICT - Information and communications technology
IDC - International Data Corporation
IESP - International Exascale Software Project
IIoT - Industrial Internet of Things
IMC - In-Memory Computing
IMC/PIM - In Memory Computing; Processor In Memory
INVG - Istituto Nazionale di Geofisica e Vulcanologia (Italian National Institute of Geophysics and Volcanology)
IOC/UNESCO - Intergovernmental Oceanographic Commission of UNESCO (The United Nations Educational, Scientific and Cultural Organization)
IoT - Internet of Things
IoV - Internet of Vehicles
IRISA - Research Institute of Computer Science and Random Systems
ISA - Instruction Set Architecture
ISV - Independent Software Vendor
IT - Information Technology
JGU - Johannes Gutenberg-Universität Mainz
JU - Joint Undertaking
KPI - Key Performance Indicator
KTH - Kungliga Tekniska Högskolan
LAN - Local Area Network
LDAP - Lightweight Directory Access Protocol
LDAP - Lightweight Directory Access Protocol
LHC - Large Hadron Collider
LIGO - Laser Interferometer Gravitational-Wave Observatory
LLVM - Low Level Virtual Machine (please note this acronym has officially been removed to avoid confusion)
M - Month
MB - Mega Byte
MC/PIM - In Memory Computing; Processing In Memory
MFF - Multiannual Financial Framework
MIMD - Multiple Instruction, Multiple Data
ML - Machine Learning
MPI - Message Passing Interface
MRAM - Magnetic RAM
MSA - Modular Supercomputing Architecture
MW - megawatt
NAG - Numerical Algorithms Group
NAND - NOT-AND
NGI - Next Generation Internet
NIC - Network Interface Controller
NN - Neural Network
NOSQL - non SQL
NUDT - National University of Defense Technology
NV-DIMM - non-volatile dual in-line memory module
Nvlink - a wire-based communications protocol serial multi-lane near-range communication link developed by Nvidia
NVM - Non-Volatile Memory
NVMe - NVM Express
NVMeoF - NVMe over Fabrics
OCP - Open Compute Project
OPA - Omni-path Architecture
OpenACC - Open Accelerators
OpenCAPI - Open Coherent Accelerator Processor Interface
OpenMP - Open Multi-Processing
OS - Operating System
OSH - Open Source Hardware
OSS - Open Source Software
OTA - over-the-air

OxRAM - Oxide-Based Resistive Memory
PaaS - Platform-as-a-Service
PB - petabyte
PCIe - Peripheral Component Interconnect Express
PCM - Phase Change Memory,
PCRAM - Phase Change RAM
PGA - Program Global Area
PII - Personally Identifiable Information
PM - Person Month
PMIx - Process Management for Exascale environments
POSIT – a hardware friendly version of Unums
POSIX - Portable Operating System Interface
PRACE - Partnership for Advanced Computing in Europe
PUE - Power Usage Effectiveness
Q - Quarter
QLM - Quantum Learning Machine
QoS - Quality of Service
R&D - Research and Development
R&I - Research and Innovation
RAM - Random-Access Memory
RAS - Reliability, availability and serviceability
RDMA - Remote Direct Memory Access
RFP - Request for Proposal
RIAG - Research and Innovation Advisory Group
RISC-V - Reduced Instruction Set Computer - five
RNA computing -
RoCE - RDMA over Converged Ethernet
ROI - Return On Investment
RRAM or ReRAM - metal oxide Resistive Random-Access Memory
SaaS - Software-as-a-Service
SC - Supercomputing Conference
SCM - Storage Class Memory
SDG - sustainable development goals
SerDes - Serialiser/Deserialiser
SHAPES - ME HPC Adoption Programme in Europe
SHS - Social and Historical Sciences
SICOS - SImulation, COmputing and Storage
SIMD - single instruction, multiple data
SINTEF - Stiftelsen for industriell og teknisk forskning
SiPh - Silicon Photonics
SKA - Square Kilometre Array
SLA - Service Level Agreement
SME - Small and Medium-sized Enterprise
SQL - Structured Query Language
SRA - Strategic Research Agenda
SSD - Solid-State Drive
SSH - Secure Shell
STT-RAM - Spin-transfer torque magnetic random-access memory
SW - software
TCO - Total Cost of Ownership
TOPS - tera operations per second
TPU - tensor processing unit
TRL - Technology Readiness Level
TSMC - Taiwan Semiconductor Manufacturing Company
TSP - Tsunami Service Providers
UAV - Unmanned Aerial Vehicle
UI - User Interface
UMA - University of Málaga
UN - United Nations
Unums - Universal Numbers -
UPS - Uninterrupted Power Supply
US - The United States of America
V2E - vehicle-to-everything
VCSEL - Vertical-Cavity Surface-Emitting Laser
VPU - Vision Processing Unit
vs. - versus
WG - Working Group
WLAN - Wireless LAN
WP - Work Package
XDR - External Data Representation
YE - Year End
ZB - zettabyte

11.2

Acknowledgements

ETP4HPC would like to express its gratitude to everyone involved in the writing of this Strategic Research Agenda. This document is the collective work of ETP4HPC Members and experts representing other organisations and projects, whose input and advice allow this document to tackle all components of the Digital Continuum. The writing of this SRA was supervised by our SRA editorial team: Michael Malms (ETP4HPC), Marcin Ostasz (ETP4HPC), Maike Gilliot (CEA & ETP4HPC) and Pascale Bernier-Bruna (Atos and ETP4HPC). Special thanks are in order for the Leaders of the Working Groups of this SRA as most of the work that built this document has taken place within the following teams:

- System Architecture led by Laurent Cargemel (Atos), Estela Suarez (Jülich Supercomputing Centre) and Nico Mittenzwey (Megware)
- System Hardware Components led by Marc Duranton (CEA & HiPEAC), Sakir Sezer (NVIDIA) and Craig Prunty (Sipearl)
- System Software & Management led by Pascale Rosse-Laurent (Atos), María S. Pérez-Hernández (Universidad Politécnica de Madrid & BDVA) and Manolis Marazakis (FORTH-ICS)
- Programming Environment led by Guy Londsdale (Scapos), Paul Carpenter (BSC) and Gabriel Antoniu (Inria & BDVA)
- I/O & Storage led by Sai Narasimhamurthy (Seagate) and André Brinkman (Universität Mainz – JGU)
- Mathematical Methods & Algorithms led by Dirk Pleiter (KTH) and Utz-Uwe Haus (HPE)
- Centre-to-Edge Framework led by Jens Krueger (Fraunhofer) and Hans-Christian Hoppe (Partec)
- Application co-design led by Erwin Laure (MPCDF) and Andreas Wierse (SICOS)
- Quantum for HPC led by Valeria Bartsch (Fraunhofer), Kristel Michelsen (FZJ), Cyril Allouche (Atos)
- Unconventional HPC Architectures led by Tobias Becker (Maxeler) and Robert Haas (IBM)

The core and heart of this SRA, the technical contents of the chapters of the Working Groups, was collaboratively written by the members of ETP4HPC. Special thanks are in order to those members of ETP4HPC who took part in the arduous process of defining the respective research priorities within our SRA Working Groups.

11.2.1

List of all SRA 5 contributors

ETP4HPC OFFICE

- Michael Malms - Chief Editor
- Marcin Ostasz
- Maike Gilliot
- Pascale Bernier-Bruna
- Jean-Pierre Panziera

WORKING GROUPS - RESEARCH DOMAINS

SRA 5 WORKING GROUP (WG)		Full Name	ORGANISATION
1 - System Architecture	Leader	Estela Suarez	JSC
	Leader	Nico Mittenzwey	MEGWARE
	Leader	Laurent Cargemel	Atos
		Alessandro Russo	Leonardo
		Andy Forrester	HypeAccelerator Solutions Limited
		Carlo Cavazzoni	Leonardo
		Craig Prunty	SiPearl
		Daniele Cesarini	Cineca
		David Tur	Do IT Now
		Fabrizio Magugliani	E4 Computer Engineering SpA
		Marc Casas	BSC
		Rene Oertel	MEGWARE Computer Vertrieb und Service GmbH
		Sergio Sánchez	Vicomtech
		Thierry Porcher	Do IT Now
2 - System Hardware Components	Leader	Marc Duranton	CEA / HiPEAC
	Leader	Sakir Sezer	NVIDIA
	Leader	Craig Prunty	SIPEARL
		Alessandro Russo	Leonardo
		Dirk Pleiter	KTH
		Luigi Capone	Leonardo
		Marc Casas	BSC
		Marco Cicala	E4 Computer Engineering SpA
		Osman Unsal	BSC
		Paolo Amato	Micron
		Petar Radojkovic	BSC
		Rod Evans	NVIDIA
		Thierry Porcher	Do IT Now
		Xavier Martorell	BSC
3 - System Software & Management	Leader	Manolis Marazakis	FORTH
	Leader	Maria Perez	UPM / BDVA
	Leader	Pascale Rossé-Laurent	Atos
		Alberto Miranda	BSC
		Alberto Scionti	LINKS Foundation
		Alessandro Russo	Leonardo
		David Tur	Do IT Now
		Denis Maggi	Vicomtech
		Georgios Goumas	ICCS
		Ina Schmitz	ParTec
		Jordi Guitart	BSC
		Julita Corbalan	BSC
		Michele Martone	LRZ
		Nico Mittenzwey	MEGWARE
		Nicolo Magini	Leonardo
		Paolo Viviani	LINKS Foundation
		Rene Oertel	MEGWARE Computer Vertrieb und Service GmbH
		Sebastien Varrette	University of Luxembourg
		Thierry Porcher	Do IT Now
		Thomas Moschny	Par-Tec

SRA 5 WORKING GROUP (WG)		Full Name	ORGANISATION
4 - Programming Environment	Leader	Guy Lonsdale	SCAPOS
	Leader	Paul Carpenter	BSC
	Leader	Gabriel Antoniu	INRIA (BDVA)
		Alexander Costan	INSA Rennes/Inria
		Ani Anciaux Sedrakian	IFPEN
		Antonio Peña	BSC
		Antonio Sciarappa	Leonardo
		Christian Perez	INRIA
		Francesco Iannone	ENEA
		Jose Gracia	HLRS
		Leonardo Arturo Bautista Gomez	BSC
		Luigi Capone	Leonardo
		Miguel Vasquez	BSC
		Olivier Marsden	ECMWF
		Paolo Viviani	LINKS Foundation
		Patrick Carribault	CEA
		Sebastien Varrette	University of Luxembourg
		Vicenc Beltran	BSC
		Xavier Martorell	BSC
5 - I/O & Storage	Leader	Andre Brinkmann	JGU - Mainz
	Leader	Sai Narasimhamurthy	Seagate
		Anna Queralt	BSC
		Jean-Thomas Acquaviva	DDN Storage
		Jesús Carretero Pérez	UC3M
		Leonardo Arturo Bautista Gomez	BSC
		Nicolo Magini	Leonardo Labs
		Paolo Amato	Micron
		Philippe Deniel	CEA
		Ramon Nou	BSC
		Tiago Quintino	ECMWF
6 - Mathematics & Algorithms	Leader	Dirk Pleiter	KTH
	Leader	Utz-Uwe Haus	HPE
		Adrian Tate	NAG
		Ani Anciaux Sedrakian	IFPEN
		Antonio Sciarappa	Leonardo Labs
		Chiara Vercellino	LINKS Foundation
		Dario Garcia-Gasulla	BSC
		Giovanni Samaey	KU Leuven
		Luigi Capone	Leonardo
		Marcin Chrust	ECMWF
		Maximilian Behr	Northern Data AG
		Olivier Beaumont	INRIA
		Ricard Borrell	BSC
		Ulrich Ruede	CERFACS/FAU
		Ward Melis	KU Leuven

APPENDICES

SRA 5 WORKING GROUP (WG)		Full Name	ORGANISATION
7 - Application co-design	Leader	Erwin Laure	**MPCDF**
	Leader	Andreas Wierse	**SICOS**
		Bruno Raffin	**INRIA**
		Carlo Cavazzoni	**Leonardo**
		Guillaume Houzeaux	**BSC**
		Ioan Hadade	**ECMWF**
		Kim Serradell Maronda	**BSC**
		Luigi Capone	**Leonardo**
		Miguel Vasquez	**BSC**
		Ricard Borrell	**BSC**
		Sabri Pllana	
		Sinead Ryan	**Trinity College Dublin**
		Vicence Beltran	**BSC**
8 - Centre-to-edge-framework	Leader	Hans-Christian Hoppe	**ParTec**
	Leader	Jens Krueger	**FRAUNHOFER**
		Alberto Scionti	**LINKS Foundation**
		Ander Garcia	**Vicomtech**
		Anna Queralt	**BSC**
		Benjamin Depardon	**UCit**
		Craig Prunty	**SIPEARL**
		Daniela Ghezzi	**Leonardo S.p.a.**
		Daniele Piccarozzi	**Arm**
		David Carrera	**BSC**
		Philippe Bricard	**UCit**
		Thierry Goubier	**CEA**
		Venkatesh Kannan	**ICHEC**
9 - Quantum for HPC	Leader	Valeria Bartsch	**FRAUNHOFER**
	Leader	Cyril Allouche	**Atos**
	Leader	Kristel Michelsen	**JSC**
		Andrea Scarabosio	**Links Foundation**
		Artur Garcia	**BSC**
		Chayma Bouazza	**PASQAL**
		Daniele Dragoni	**Leonardo**
		Daniele Gregori	**E4 Computer Engineering SpA**
		Daniele Ottaviani	**Cineca**
		David Bowden	**Dell Technologies**
		David Tur	**Do IT Now**
		Dennis Hoppe	**High-Performance Computing Center Stuttgart**
		Fabrizio Magugliani	**E4 Computer Engineering SpA**
		Filippo Palombi	**ENEA**
		Giacomo Vitali	**LINKS Foundation**
		Guillaume Colin de Verdière	**CEA**
		Jean-Philippe Nominé	**CEA**
		Leonardo Arturo Bautista Gomez	**BSC**
		Mikael Johansson	**CSC**
		Olivier Terzo	**LINKS Foundation**
		Osman Unsal	**BSC**
		Vekatesh Kannan	**ICHEC**

SRA 5 WORKING GROUP (WG)		Full Name	ORGANISATION
10 - Non-conventional HPC Architectures	Leader	Tobias Becker	**MAXELER**
	Leader	Robert Haas	**IBM Research**
		Angeliki Pantazi	**IBM Research Zurich**
		Cosimo Gianfreda	**E4 Computer Engineering SpA**
		Johannes Schemmel	**KIP-HD / Heidelberg University**
		Ken O'Brien	**Xilinx**
		Marco Cesena	**Leonardo S.p.A.**
		Osman Unsal	**BSC**
		Paolo Amato	**Micron**
		Paolo Palazzari	**ENEA**
		Paul Carpenter	**BSC**
		Petar Radojkovic	**BSC**
		Sagar Dolas	**SURF, Netherlands**
		Sebastien Varrette	**University of Luxembourg**
		Steve Furber	**University of Manchester**
		Tobias Becker	**Maxeler Technologies**
		Xavier Martorell	**BSC**
		Yannis Papaefstathiou	**Exascale Performance Systems - EXAPSYS P.C.**
		Yulia Sandamirskaya	**Intel**

SRA 5 Research Clusters and other chapters

FEDERATED HPC, CLOUD AND DATA INFRASTRUCTURES
- Utz-Uwe Haus (HPE)
- Sai Narasimhamurthy (Seagate)
- Maria S. Perez (UPM)
- Dirk Pleiter (KTH)
- Andreas Wierse (SICOS)

HETEROGENEOUS HIGH PERFORMANCE COMPUTING
- Paul Carpenter (BSC)
- Utz-Uwe Haus (HPE)
- Erwin Laure (MPCDF)
- Sai Narasimhamurthy (Seagate Systems)
- Estela Suarez (Forschungszentrum Jülich)

HPC FOR URGENT DECISION-MAKING
- Manolis Marazakis (FORTH)
- Marc Duranton (CEA)
- Dirk Pleiter (KTH)
- Giuliano Taffoni (INAF)
- Hans-Christian Hoppe (Scapos AG)

TOWARDS INTEGRATED HARDWARE/ SOFTWARE ECOSYSTEMS FOR THE EDGE-CLOUD-HPC CONTINUUM
- Gabriel Antoniu (Inria)
- Patrick Valduriez (Inria)
- Hans-Christian Hoppe (SCAPOS)
- Jens Krüger (Fraunhofer ITWM)

SUSTAINABILITY
- Andreas Wierse (Sicos BW)
- François Bodin (Inria)
- Sagar Dolas (SURF)
- Damien Gratadour (Université Paris Diderot)
- Michael Malms (ETP4HPC)
- Leonieke Mevus (SURF)
- Pascale Rossé-Laurent (Atos)

HPC TECHNOLOGY PROJECTS
- EUPEX - Jean-Robert Bacou (Atos)
- EUPilot - Carlos Puchol (BSC)
- HPCQS - Kristel Michielsen (JSC)

EDUCATION AND TRAINING
- Tiina Leiponen (CSC)
- Carlo Cavazzoni (Leonardo)
- Ivan Spisso (Leonardo)
- Maike Gilliot (ETP4HPC)
- Hans-Christian Hoppe (Scapos)
- Guy Lonsdale (ETP4HPC Steering Board / Fraunhofer / Scapos / FocusCoE)
- Fabrizio Magugliani (ETP4HPC Steering Board / E4 Computer Engineering)
- Jean-Philippe Nominé (ETP4HPC Steering Board / CEA)
- Andreas Wierse (SICOS)
- Pascal Bouvry (University of Luxembourg)
- Daniela Posch (SICOS)
- Gilles Civario (Intel)
- Henri Calandra (TotalEnergies & PRACE IAC)
- Vincent Galinier (Airbus & PRACE IAC)
- Xavier Vigouroux (Atos)
- Alban Rousset (LuxProvide)
- Michael Schlottke-Lakemper (HLRS)
- Carolina Berucci (Leonardo)
- Erwin Laure, MPCDF

OTHER CONTRIBUTORS
- Carolina Berucci (Leonardo)

ETP 4 HPC

EUROPEAN TECHNOLOGY PLATFORM FOR HIGH PERFORMANCE COMPUTING

Contact ETP4HPC
contact@etp4hpc.eu
www.etp4hpc.eu

@etp4hpc

Text
ETP4HPC

ETP4HPC Chairman
Jean-Pierre Panziera

Editorial Team
Michael Malms (Chief editor)
Marcin Ostasz
Maike Gilliot
Pascale Bernier-Bruna

Graphic design and layout
Antoine Maiffret (www.maiffret.net)

October 2022

Contact ETP4HPC
contact@etp4hpc.eu
www.etp4hpc.eu

@etp4hpc