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

EXDCI-2

European eXtreme Data and Computing Initiative - 2

Grant Agreement Number: 800957

D5.3

Presence of Europe in HPC-HPDA standardisation and recommendations to promote European technologies

Final

Version: 1.0

Authors: JF Lavignon, TS-JFL; Marcin Ostasz, ETP4HPC; Thierry Bidot, Neovia Innovation

Date: 21/02/2020

Project and Deliverable Information Sheet

EXDCI Project	**Project Ref. №: FETHPC-800957**	
	Project Title: European eXtreme Data and Computing Initiative - 2	
	Project Web Site: http://www.exdci.eu	
	Deliverable ID: **D5.3**	
	Deliverable Nature: Report	
	Dissemination Level: PU *	**Contractual Date of Delivery:** 29 / 02 / 2020
		Actual Date of Delivery: 21/02/2020
	EC Project Officer: Evangelia Markidou	

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

Document Control Sheet

Document	**Title: Presence of Europe in HPC-HPDA standardisation and recommendations to promote European technologies**	
	ID: **D5.3**	
	Version: <1.0>	**Status:** ***Final***
	Available at: http://www.exdci.eu	
	Software Tool: Microsoft Word 2013	
	File(s): **EXDCI-2-D5.3-V1.0.docx**	
Authorship	**Written by:**	JF Lavignon, TS-JFL Marcin Ostasz, ETP4HPC Thierry Bidot, Neovia Innovation
	Contributors:	
	Reviewed by:	Elise Quentel, GENCI John Clifford, PRACE
	Approved by:	MB/TB

Document Status Sheet

Version	Date	Status	Comments
0.1	03/02/2020	Draft	To collect internal feedback
0.2	07/02/2020	Draft	For internal reviewers
1.0	21/02/2020	Final version	Integration of reviewers' comments

Document Keywords

Keywords:	PRACE, , Research Infrastructure, HPC, HPDA, Standards

Table of Contents

List of Tables

References and Applicable Documents

[1] http://www.exdci.eu
[2] http://www.prace-project.eu
[3] http://www.etp4hpc.eu

List of Acronyms and Abbreviations

Below is a list of acronyms used within the EXDCI-2 project. The acronym specific to this report are explained inside the text where they are used.

AISBL	Association Internationale Sans But Lucratif (International Non-for-Profit Association)
BDEC	Big Data and Extreme-scale Computing
BDV	Big Data Value
CoE	Centres of Excellence for Computing Applications
cPPP	contractual Public-Private Partnership
CSA	Coordination and Support Action
D	Deliverable
DG	Directorate General
DoW	Description of Work
EC	European Commission
ECMWF	European Centre for Medium-range Weather Forecasts
EESI	European Exascale Software Initiative
ENES	European Network for Earth System modelling
EPOS	European Plate Observing System
EsD	Extreme scale Demonstrators
EU	European Union
FET	Future and Emerging Technologies
FP7	Framework Programme 7
GDP	Gross Domestic Product
H2020	Horizon 2020 – The EC Research and Innovation Programme in Europe
HPC	High Performance Computing
IDC	International Data Corporation
IESP	International Exascale Software Project
INVG	Istituto Nazionale di Geofisica e Vulcanologia (National Institute of Geophysics and Volcanology)
ISV	Independent Software Vendor
IT	Information Technology
KPI	Key-Performance Indicator
M	Month
OS	Operating System
PM	Person Month
Q	Quarter
R&D	Research and Development
R&I	Research and Innovation
RFP	Request for Proposal
ROI	Return On Investment

SHAPE	SME HPC Adoption Programme in Europe
SHS	Social and Historical Sciences
SME	Small and Medium Enterprise
SRA	Strategic Research Agenda
SWOT	Strengths, Weaknesses, Opportunities and Trends
TRL	Technology Readiness Level
US	United States
WG	Working Group
WP	Work Package

Executive Summary

The domain of standards relevant for HPC is currently evolving due to two complementary trends that shape the future of HPC: the inclusion of HPC in the digital continuum and the convergence of HPC with AI and Big Data Analytics. In the HPC standard landscape this translates into a need for new standards to deal with the:

- Proliferation of new architectures that support more diverse workloads
- Emergence of new integration levels such as photonics or chiplet integration
- New memory coherent models for better performance
- New memory hierarchy interfaces
- Interoperability to support complex workflows
- Security challenges
- Support of more programming languages
- Development of meta data to facilitate identification and usage of data
- Emergence of high-level software packages supporting deep learning and data analytics.

In our analysis of the HPC standard organisations we have taken into account these trends. We have looked at four categories of standardisation organisations: hardware, HPC management software stack, programming model and application software.

The conclusion of this analysis is that Europe's presence in international standard organisations is limited to the representatives of a few projects or mainly academic organisations, which have a vested interest in developing and maintaining standards. Most of the standard organisations are controlled by US-based organisations. Also, the other two prominent HPC ecosystems, Japan and China, have been making headway in strengthening their presence and thus have a direct influence on standards. In order to keep up, Europe must invent a mechanism to ensure its continuous presence in standard management.

A set of interviews with European industrial companies interested by standard has confirmed this underinvestment. The interviewees are aware of the resulting risk and ready to be more active if they are helped. The suggested incentives include funding opportunities for consortia to work on standards, processes to facilitate the sharing of the efforts between the European stakeholders and a more connected and organised value chain to help the emergence of user communities around de facto standards.

Another important finding of the interviews is that some European players would like to see new standards to be put in place in order to allow the development of new products or new services that cannot exist within a proprietary or over fragmented environment. The main domains where such new standards are expected by European players are Big Data/AI and HPC as a service (proposed by cloud providers). In these new domains, European SMEs have ideas for innovative services or products but due to the absence of standards it is either too complex to develop them or too difficult to prove the benefits that they could bring.

This work on standards leads to two sets of recommendations. The first set targets the optimisation of the current efforts without a significant increase of resource:

- R1 Achieve a continuous survey of emerging standards, produce maps of new standard organisations relevant in the context of the digital continuum, maintain all maps up to date and disseminate this information.
- R2 Establish a network of interested people around each European standardisation expert to share the information about the standard activity.

The second, more ambitious, set of recommendations aims to make Europe a leader in this domain. These recommendations would mean additional resources that can partly be provided through the EuroHPC JU[1].

- R3: Implement a specific HPC organisation with sufficient financial and human resources, whose goal will be to develop European HPC presence in standard organisations.
- R4: Put in place a process to help consortia organise their presence in a standard organisation or write emerging standard specifications.
- R5 Establish a process for participants to apply to European funded projects for a specific funded effort to facilitate the emergence of a standard.
- R6 Identify and launch standardisation projects in strategic domains, to establish a level playing field, to favour European player's new initiatives and to encourage green IT.

The rise of HPC in the digital continuum is creating new challenges in terms of standards. Europe can be a significant player and develop a sound and dynamic environment for the emergence of new scientific, industrial and societal applications of HPC with a more active approach of standards. The proposed recommendations can push in this direction and make Europe a leader in some standard domains which foster an open, sustainable and competitive IT ecosystem.

[1] Joint Undertaking

1 Introduction

The aim of Task 5.2 of the EXDCI-2 project is to assess the position of Europe in HPC related standardisation organisations and to propose solutions to increase the weight of European players in the setting up of present or future HPC standards.

This report presents the activity of this task and the recommendations that are the results of the analysis and the discussions with European stakeholders.

In Chapter 2, we present the standard organisations that are relevant for HPC. As the HPC field is not static, we start with the analysis of the main trends that have an impact on the standardisation process in HPC which is most of the time driven by a de facto approach where a set of specifications becomes a standard adopted by the community. The analysis is segmented according to the HPC value chain with:

- Hardware oriented standards
- Management software stack oriented standards
- Programming model software oriented standards
- Application software oriented standards

In Chapter 3, we describe the main findings of this analysis of standardisation organisations. We focus on the presence of European stakeholders and on the European experts who are influential in standardisation organisation. This chapter is concluded by an analysis of the strengths and weaknesses of Europe in the standard organisations.

To complement the previous analysis of the standardisation landscape, we have conducted some interviews with European stakeholders. Our objective is to discuss the stakes and the obstacles with different European companies ranging from SMEs to large organisations. Chapter 4 presents the methodology and the main findings of the interviews.

In Chapter 5, we present the recommendations that result from the standardisation landscape analysis and the discussion with the stakeholders. These recommendations are aimed at the EuroHPC Joint Undertaking and the European Commission with the objective to help these organisations set actions that will increase the influence of Europe in HPC standard creation.

2 Presentation of standard organisations relevant in HPC

This chapter presents the most relevant organisations active in standardisation that impact the HPC domain. Covering all such organisations is out of reach, nevertheless, we expect to have selected a large and diverse enough panel to analyse the most strategic ones and to see the different situations that can arise. After an analysis of the main trends that influence the evolution of standards, the organisations are briefly presented. The organisations are segmented according to the HPC value chain:

- Hardware oriented standards
- Management software stack oriented standards
- Programming model software oriented standards
- Application software oriented standards

A template has been used to gather the information about the standardisation organisations. This template can be found in Annex 1 (Section 7.1). Some examples of standard organisation descriptions can be found in Annex 2 (Section 7.2).

2.1 Trends in standards for HPC/HPDA

At a global level, we can see two complementary trends that shape the future of HPC:

- The inclusion of HPC in the digital continuum;
- The convergence of HPC with AI and Big Data Analytics.

As explained in the Strategic Research Agenda issued by the EXDCI-2 project, the rapid proliferation of digital data generators, the unprecedented growth in the volume and diversity of the data they generate, and the intense evolution of the methods for analysing and using that data are radically reshaping the landscape of scientific computing. The most critical problems involve the logistics of wide-area, multistage workflows that will move back and forth across the computing continuum, between the multitude of distributed sensors, instruments and other devices at the network's edge, and the centralized resources of commercial clouds and HPC centres. This new paradigm is called 'The Digital Continuum'. As a consequence, HPC systems have to be more and more interconnected with other IT infrastructures (edge and cloud) and the HPC software stacks have to support complex workflows and data logistics.

The second global trend is that Big Data Analytics and Artificial Intelligence (and especially Deep Learning) applications are becoming a larger and larger part of the HPC system workload replacing more traditional simulation applications. This has important consequences for the HPC hardware and software stacks that need to propose energy efficient solutions and new features for these new types of applications.

At the hardware level, these two main trends translate in the emergence of accelerator chips and the need to support HPC in a loop. Hardware is also influenced by some new technologies such as package level integration or photonics interconnects that are developed in response to the forthcoming end of Moore's law. In a more long-term perspective, we can also expect some new materials to be used for compute, communication or storage. These new technologies might complement the silicon centric value chain rather than replace it.

At the software level, the current major evolutions are:

- Fast integration of AI: this integration goes in two directions, adaptation of the stacks to support AI applications, and integration of AI techniques into the stacks to increase HPC or simulation efficiency.
- Fast integration of data coming from sensors or separate computing or AI to increase HPC or simulation efficiency.
- Emphasis on the accessibility of data.
- Development of specific workflow tools allowing those two above evolutions through transformation of simulation software to complex workflows.
- HPC software have to serve applications ranging from edge to HPC centres. This evolution drives new features for the software stacks such as complex end-to-end workflow management, security, integration with cloud capability or interactive HPC.
- Adaptation to support the HPC hardware evolution: here the main current drivers are new data storage hierarchy and integration of heterogeneous computing options.
- Sustainability: developing new features to reduce the environmental impacts of HPC; one of the axes is to balance energy consumption and performance; another one is to extend the life time of hardware resources.

For the HPC standard landscape these evolutions have several consequences:

- Proliferation of new architectures that support more diverse workloads
- Emergence of new integration levels such as photonics or chiplet integration
- Memory coherent models for better performance
- New memory hierarchy interfaces
- Interoperability to support complex workflows
- Security challenges
- Support of more programming languages
- Development of meta data to facilitate identification and usage of data
- Emergence of high-level software packages supporting deep learning and data analytics.

In our analysis of the standard organisations, we have taken into account these trends. Nevertheless, for some of them there is still a lack of maturity to have a translation in standard. This can be seen as good news, as it gives Europe more of a chance to have an influence at the global level in the adoption of these emerging technologies.

2.2 Hardware

In the domain of hardware, we can find standard oriented organisations for each of the components of the HPC systems (processor, memory, interconnect and storage) and some targeting the whole system. Our analysis has focused on (all the acronyms are explained in the text below):

- Processors: OpenPOWER, ARM, Risc-V, OpenHW Group
- Interface between chips: Gen-Z, CCIX, CLX, OpenCapi, NVlink, AIB
- Memory: JEDEC (DDR, HBM specifications)
- Interconnects: PCI-SGI, Infiniband Trade Association
- Storage: INCITS (SCSI, SATA), SNIA
- System and data centres: Open Compute, ASHREA.

Processors

The dominant computing solution in HPC is a combination of x86 processors and Nvidia GPUs[2]. These two architectures are closed with the market leaders Intel, AMD and Nvidia defining the future evolutions without a strong concern about external views.

This domination can be a problem for economical or strategic reasons and so it has generated different initiatives to provide a more diverse market. These initiatives to counterbalance the market power of x86 and Nvidia have a more open approach and try to involve the HPC community to define the future evolution of their solutions.

The first example of this trend is the OpenPOWER initiative (https://openpowerfoundation.org). Having lost the HPC processor battle to Intel, IBM has tried to sustain its POWER architecture by opening its evolutions to external partners. The OpenPOWER foundation was incorporated in 2013 with the purpose to create an ecosystem around the POWER architecture and its associated technologies. It has attracted a large number of IT players including some of the main European HPC organisations such as Bull, Cineca and FZJ[3]. In August 2019, the OpenPOWER Foundation announced that POWER ISA[4] would be under an open model – guided by the OpenPOWER Foundation within the Linux Foundation – and be made available to the open technical commons. Nevertheless, the influence of POWER architecture in HPC is still decreasing and the momentum around this architecture has not been created in the HPC community.

Another ecosystem around the ARM (https://developer.arm.com/) technology seems more successful in the HPC arena. Coming from low power and embedded processors, ARM architecture has been complemented to be able to provide the performance level requested by HPC workload. This technology developed by a UK company is now owned by Softbank a Japanese investor though most of the developments remains in Europe or US research labs. Europe has been one of the first to consider ARM technology for HPC with the Montblanc set of projects led by Atos/Bull and BSC[5]. These projects have made a strong contribution for the emergence of a complete HPC solution based on the ARM technology. The Japanese flagship program "Fugaku"[6] has also chosen this technology and ARM is becoming a highly visible HPC ecosystem. Nevertheless, the ARM company manages the evolution of the technology in a not completely open way. Even if the software stack aspects are managed by the Linaro organisation which is a member-based organisation, the evolution of the ISA is under the sole control of ARM. ARM sometime invites external companies to work on the future features (as it has done for the adoption of the vector extension with Fujitsu) but it has always the final decision. This approach and the Japanese control over ARM could be a problem for European players to invest more in this ecosystem.

As an alternative the RISC-V ecosystem (https://riscv.org) is more open. This technology which started as a Berkeley project is now managed by a foundation that decides on the evolutions and specifications. This foundation is member-based with a significant participation of European organisations. Most of these European members are nevertheless from domains other

[2] Graphical Processing Unit

[3] Forschungszentrum Jülich

[4] Instruction Set Architecture

[5] Barcelona Supercomputing Centre

[6] Name the Japanese flagship system that is currently being installed in Kobe at RIKEN and was called post K in the past.

than HPC except BSC and CEA. This is certainly due to the fact that RISC-V is not as mature for HPC as ARM. There is currently a vector extension working group that plans to release the specification at the end of 2020. So, things can evolve quite quickly and RISC-V can be considered as a candidate for future HPC systems.

In the RISC-V ecosystem, the OpenHW Group (https://www.openhwgroup.org) has just been created in 2019 for the purpose of developing an open hardware ecosystem. The first contribution comes from ETHZ (Eidgenössische Technische Hochschule Zürich) with a CORE-V IP (Intellectual Property) that implement a RISC-V architecture. More globally the aim of OpenHW Group is to establish a collaboration of hardware and software designers for the development of open-source cores, related IP, tools and software. For the time being it is very focused on the exploitation of the ETHZ PULP platform and has a strong involvement of NXP. It is too early to know if it will have a broader impact but it is a good example of the trend toward open hardware.

The European Processor Initiative (EPI) is investigating processor designs based on ARM and RISC-V architectures. This initiative could give Europe more influence on the evolutions of these two ecosystems. Some specific efforts from EPI to actively participate in the related standard oriented organisations are required to effectively ensure that the two ecosystems will evolve in the direction of European interests.

Interface between chips

This domain is quite new as in the past processors had only either internal proprietary bus or external non-coherent interfaces such as PCI (Peripheral Component Interconnect). The first attempt to open the processor domain was made by AMD with the HyperTransport Link used for NUMA multiprocessor or connection to GPU. Intel remained with a closed technology with its CSI (Common System Interface), QPI (QuickPath Interconnect) and now UPI (Ultra Path Interconnect) protocols.

Now with the development of accelerators and new memory solutions, this domain of high-speed communication with the ability to have coherent memory domains has developed with a large number of standard oriented solutions that target different design options.

One of the first, to be created in 2016, was the GEN-Z consortium (https://genzconsortium.org/) initiated by HPE following its effort to promote “the machine” concept that was an attempt to develop a disaggregated architecture with computing, networking and memory resources that can be combined according to the need of the application workload. GEN-Z is a new data-access technology designed to provide high-speed, low-latency, memory-semantic access to data and devices via direct-attached, switched or fabric topologies. The GEN-Z consortium is a membership-based organisation with very little European participation. The main target of GEN-Z is disaggregated architecture with high performance communication with memory.

At the same time, the CCIX (Cache Coherent Interconnect for Accelerators) Consortium (https://www.ccixconsortium.com/) was also created under the leadership of processor vendors such as AMD and ARM. The mission of the CCIX Consortium is to develop and promote adoption of an industry standard specification to enable coherent interconnect technologies between general-purpose processors and acceleration devices for efficient heterogeneous computing. The standard allows processors based on different instruction set architectures to extend the benefits of cache coherent, peer processing to acceleration devices including FPGAs, GPUs, network/storage adapters, intelligent networks, and custom ASICs, allowing system designers to seamlessly integrate the right combination of heterogeneous components for their specific system needs. European participation, with the exception of ARM, is also rather weak with only a few “adopters” members (the lowest level of membership) from Europe.

Nevertheless, this standard is very well suited for the communication between processors and accelerators.

A more recent initiative targeting the same technical objective is CLX (Compute Express Link) (https://www.computeexpresslink.org). CXL is an open industry standard interconnect offering high-bandwidth, low-latency connectivity between host processor and devices such as accelerators, memory buffers, and smart I/O devices. The first specifications were released in 2019 and the technology is built upon the PCI Express (PCIe) physical and electrical interface with protocols in three key areas: I/O, memory and cache coherence. Due to this positioning upon PCIe, it requires fewer specific developments and is now supported by a large amount of IT players including the main processor and server vendors. It is perhaps too early to assess the participation of European organisations but up to now it seems that they are almost completely absent from this initiative.

It is worth also mentioning the OpenCAPI consortium (Coherent Accelerator Processor Interface) (https://opencapi.org) that is pushed by IBM. It aims to create an open coherent high performance bus interface based on a new bus standard and grow the ecosystem that utilizes this interface. Organized as a membership based organisation, it has few European members which are academic organisations.

Even if it is not managed as a standard, the NVLink (https://www.nvidia.com/en-us/data-center/nvlink) is also an important technology of this domain developed by Nvidia which is by far the accelerator leader today. NVLink is a wire-based communications protocol for near-range semiconductor communications developed that can be used for data and control code transfers in processor systems between CPUs and GPUs and solely between GPUs. To design a chip implementing NVLink you need a licence from Nvidia. This shows that the interchip communication domain is still not completely open and that dominant players continue to try to control the game.

The previous technologies were designed to interconnect chips. We are now starting to see initiatives to interconnect chiplets at the interposer level. This trend is important for HPC and if Europe wants to have a place in the processor or accelerator market, it is worth actively following this domain.

One noticeable initiative is the proprietary technology developed by Intel Advanced Interface Bus pushed as a potential standard. AIB is a low-power, high-bandwidth die-to-die PHY (Physical Layer) level interface designed to allow direct chiplet-to-chiplet communication. With this initiative Intel is pushing an interposer interconnect standard that is compatible with its EMIB (Embedded Multi-Die Interconnect Bridge) a technology only available with Intel interposers. Even if Intel is providing a royalty-free license for their Advanced Interface Bus this solution developed under the DARPA (Defense Advanced Research Projects Agency) ERI (Electronic Resurgence Initiative) programme is not really open.

European initiatives in this domain of chiplet integration can benefit from the knowledge developed in integration by its big electronics RTOs (Research and Technologies Organisations) CEA/LETI, Fraunhofer and IMEC.

Memory

As a much more mature domain, the standard oriented initiatives for memory are all managed by a single organisation JEDEC (Joint Electron Device Engineering Council) (https://www.jedec.org) created in 1958. JEDEC develops open standards for the microelectronics industry, with more than 3,000 volunteers representing nearly 300 members

companies. It is accredited by ANSI (American National Standards Institute) and maintains liaisons with numerous standards bodies throughout the world.

JEDEC is organized in committees and subcommittees and organisations can apply for one of several of them. The JC-42 is in charge of the DRAM (Dynamic Random-Access Memory), NVM (Non-Volatile Memory) and HBM (High Bandwidth Memory) standards and the JC-64 of Solid State Drives (SSD).

Even if Europe does not anymore have any ambitions in developing its own solutions, it is worth following the work done on new DRAM, HBM, NVM and SSD standards to be sure that European server or processor designs will be able to integrate the latest memory technologies.

Interconnects

To interconnect computing element to the external world, PCIe (Peripheral Component Interconnect Express) has been one of the main standards established back in 2004. The evolution of this standard is managed by the Peripheral Component Interconnect Special Interest Group (PCI-SIG) (https://pcisig.com) which is an association (Oregon non-profit corporation) of around 700 industry companies committed to advancing its non-proprietary peripheral component interconnect technology. This organisation has around 40 European members which range from electronics to hardware vertical solutions. Even if this technology is not at the centre for future HPC evolutions, it will remain an enabler for HPC systems and so is important for European stakeholders.

Another important standard for interconnects is the Infiniband technology. Even if Mellanox (now acquired by Nvidia) has reached a monopoly (its competitors Voltaire, Qlogic or Intel have either merged with Mellanox or stopped their Infiniband business) the technology is driven by the InfiniBand Trade Association (IBTA) (https://www.infinibandta.org). It is chartered with maintaining and furthering the InfiniBand Architecture specification defining hardware transport protocols sufficient to support both reliable messaging (send/receive) and memory manipulation semantics (e.g. RDMA (Remote Direct Memory Access)) without software intervention in the data movement path. These transport protocols are defined to run over Ethernet (RoCE RDMA over Converged Ethernet) as well as InfiniBand fabrics. IBTA is an Oregon non-profit corporation managed by a steering committee that includes major US companies. Besides ARM there is no other European company member of IBTA. It could be a problem for Europe as Infiniband remains the main HPC interconnect technology even if European players such as Atos and Extoll have developed alternative solutions.

Storage

The domain of storage is also a very mature domain with a large organisation INCITS (http://www.incits.org/) (InterNational Committee for Information Technology Standards) that manages part of the domain standards. It is the central forum dedicated to creating technology standards for the next generation of innovations from cloud computing to communications, from transportation to health care technologies. In the domain of storage, its main activities are focused on SCSI(Small Computer System Interface) Storage Interfaces, Fibre Channel Interfaces and ATA (Advanced Technology Attachment) Storage Interfaces. INCITS coordinates technical standards activity between ANSI (American National Standards Institute) in the US and joint ISO/IEC (International Organisation for Standardisation International Electrotechnical Commission) committees worldwide. INCITS is very US centric.

Another storage standard organisation with activities in the domain of storage is SNIA (Storage Networking Industry Association) (https://www.snia.net/). This is a Californian corporation organized as a non-profit mutual benefit corporation under the Non-profit Mutual Benefit

Corporation Law of California. SNIA standards are primarily related to data, storage, and information management and address such challenges as interoperability, usability, and complexity. SNIA is organised with working groups dealing with the topics of importance for storage solutions such as Cloud Storage Technologies, Data Management, Data Security, Networked Storage, Next Generation Data Centres, Persistent Memory, Physical Storage and Storage Management. Is has very few European members. In the context of integration of HPC in the "Digital Continuum", European HPC stakeholders should have a more active participation in SNIA activities.

System and data centres

The system level and data centre level are interesting domains as we see some important evolutions such as the emergence of open hardware or the increasing focus on green IT.

One of the more visible standardisation efforts is the Open Compute Project (OCP) (https://www.opencompute.org). This is a collaborative community focused on redesigning hardware technology to efficiently support the growing demands on compute infrastructure. It has been created by Facebook and others with the objective to have open solutions for servers, racks, data storage or energy efficient solutions for data centres. The Open Compute Project Foundation is a non-profit, nonstock corporation based in the State of the Delaware. There is good participation of European companies around in OCP (around 30 members, of which 6 have platinum status meaning a strong participation in OCP activities). The OCP has a marketplace where vendors can present products and solutions that are implementing the standards. OCP has also developed software that can be used for the operation of a data centre. OCP is organized by projects and new projects continue to be created. Some of them are less relevant for HPC but globally OCP illustrates the trends towards open hardware and energy efficient solutions.

Regarding standards for data centres, it is important to mention the American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHREA) (https://www.ashrae.org). Even if created in America, this organisation has an international footprint. It develops a set of standards and guidelines relating to HVAC&R (heating, ventilation, air conditioning and refrigeration) systems. These standards are important for the design of HPC systems that will be compliant with regulations based upon ASHREA norms.

In the data centre space, the Data Centre Alliance (https://dca-global.org/) is more a trade association than a standard oriented organisation. Nevertheless, DCA has established Special Interest Groups (SIG) that are a forum were some of the main evolutions and interoperability issues for data centres are discussed. The topics of these SIG include security access, energy efficiency and thermal management.

2.3 Middleware

Operating Systems

Linux Foundation

Linux[7] is a family of open source Unix-like operating systems based on the Linux kernel, an operating system kernel. Linux is typically packaged in a Linux distribution. Distributions include the Linux kernel and supporting system software and libraries, many of which are

7 https://en.wikipedia.org/wiki/Linux

provided by the GNU (GNU's Not Unix) Project[8]. Linux was originally developed for personal computers based on the Intel x86 architecture, but has since been ported to more platforms than any other operating system.

Linux is the leading operating system on servers and other big iron systems such as mainframe computers, and the only OS used on TOP500 supercomputers[9] (since November 2017, having gradually eliminated all competitors). Linux also runs on embedded systems. Linux is one of the most prominent examples of free and open-source software collaboration. The source code may be used, modified and distributed—commercially or non-commercially—by anyone under the terms of its respective licenses, such as the GNU General Public License.

The Linux Foundation[10] supports the creation of sustainable open source ecosystems by providing financial and intellectual resources, infrastructure, services, events, and training. Working together, The Linux Foundation and its projects form the most ambitious and successful investment in the creation of shared technology.

The Linux Foundation hosts many of the most important open source projects in the world, including Linux. There are a few different ways to host a project with The Linux Foundation, and they are all guided by each community's requirements and goals. Most generate open source code, but projects that provide infrastructure support for open source communities, standards, or specifications also choose The Linux Foundation as their home. Linux Foundation projects meet the following requirements: 1) They use a license approved by the Open Source Initiative and 2) They allow The Linux Foundation to own community assets, like a domain or trademark, on behalf of the project community.

The Foundation is based in the US and its members, board and management are pre-dominantly US- or Asia-based (e.g. AT&T, Cisco, IBM, Intel, Fujitsu, Huawei, Microsoft, Google) but it also has a number of prominent European members – both IT providers and industrial or service companies (ABN-AMRO, Airbus, Atos, BBVA, Comarch, Deutsche Bank, Ericsson, Siemens, Total) as well as research organisations (e.g. EPFL, Budapest Technical University, University of Cambridge, GRNET). The Linux Foundation Board of Directors is comprised of 22 senior leaders from across the IT industry. However, there are no recognised European experts among the Board Members and the most influential people associated with the Foundation.

File system

Lustre

The Lustre[11] file system is an open-source, parallel file system that supports many requirements of leadership class HPC simulation environments. Born from a research project at Carnegie Mellon University, the Lustre file system has grown into a file system supporting some of the world's most powerful supercomputers. The Lustre file system provides a POSIX compliant file system interface that can scale to thousands of clients, petabytes of storage and hundreds of gigabytes per second of I/O bandwidth. The key components of the Lustre file system are the

8 https://en.wikipedia.org/wiki/GNU

9 https://www.top500.org/

10 https://www.linuxfoundation.org

11 http://lustre.org/

Metadata Servers (MDS), the Metadata Targets (MDT), Object Storage Servers (OSS), Object Server Targets (OST) and the Lustre clients.

The name Lustre is a portmanteau word derived from Linux and cluster. Lustre file system software provides high performance file systems for computer clusters ranging in size from small workgroup clusters to large-scale, multi-site clusters.

Because Lustre file systems have high performance capabilities and open licensing, it has consistently been used by at least half of the top ten, and more than 60 of the top 100 fastest supercomputers in the world.

Lustre file systems are scalable and can be part of multiple computer clusters with tens of thousands of client nodes, tens of petabytes (PB) of storage on hundreds of servers, and more than a terabyte per second (TB/s) of aggregate I/O throughput. This makes Lustre file systems a popular choice for businesses with large data centres, including those in industries such as meteorology, simulation, oil and gas, life science, rich media, and finance. The I/O performance of Lustre has widespread impact on these applications and has attracted broad attention.

Lustre file system software is available under the GNU General Public License (version 2 only). In 2014, the European Open File System (EOFS[12]) and Open Scalable File Systems, Inc (OpenSFS[13]) obtained Lustre assets from Seagate. In November 2019, OpenSFS and EOFS announced at the SC19 Lustre BOF that the Lustre trademark had been transferred to them jointly from Seagate. Since 2011 OpenSFS has been in charge of organising the annual Lustre User Group (LUG) event, traditionally held in April, for discussion and seminars on Lustre. In 2011, Lustre 2.1 was the first community release endorsed by OpenSFS. OpenSFS began direct funding of community releases in early 2012, focused on introducing new features and targeted every six months. Maintenance releases are targeted every three months. EOFS continues similar developments in Europe.

HDF

HDF supports n-dimensional datasets and each element in the dataset may itself be a complex object. **HDF5[14]** is a general purpose library and file format for storing scientific data.

HDF5 can store two primary objects: datasets and groups. A dataset is essentially a multidimensional array of data elements, and a group is a structure for organizing objects in an HDF5 file. Using these two basic objects, one can create and store almost any kind of scientific data structure, such as images, arrays of vectors, and structured and unstructured grids. You can also mix and match them in HDF5 files according to your needs.

HDF5 was created to address the data management needs of scientists and engineers working in high performance, data intensive computing environments. As a result, the HDF5 library and format emphasize storage and I/O efficiency. For instance, the HDF5 format can accommodate data in a variety of ways, such as compressed or chunked. And the library is tuned and adapted to read and write data efficiently on parallel computing systems.

NCSA (The National Center for Supercomputing Applications of the US) maintains a suite of free, open source software, including the HDF5 I/O library and several utilities. The HDF5 user

[12] https://www.eofs.eu/lustre/start

[13] https://opensfs.org/

[14] https://www.hdfgroup.org/solutions/hdf5/

community also develops and contributes software, much of it freely available. There is little commercial support for HDF5 at this time.

The HDF Group[15] is the developer of HDF5. The HDF Group based in the US is a non-profit organisation with the mission of advancing state-of-the-art open source data management technologies, ensuring long-term access to the data, and supporting its dedicated and diverse user community.

Spectrum Scale

IBM Spectrum Scale[16] is high-performance clustered file system software developed by IBM. It can be deployed in shared-disk or shared-nothing distributed parallel modes. It is used by many of the world's largest commercial companies, as well as some of the supercomputers on the Top 500 List.

Before 2015, Spectrum Scale was known as IBM General Parallel File System (GPFS).

BeeGFS

BeeGFS[17] (formerly FhGFS) is a parallel file system, developed and optimised for high-performance computing. BeeGFS includes a distributed metadata architecture for scalability and flexibility reasons. Its most important aspect is data throughput.

BeeGFS was originally developed at the Fraunhofer Center for High Performance Computing in Germany. ThinkParQ, the spin-off company of this project was founded in 2014 to maintain BeeGFS and offer professional services.

CEPH

Ceph is a next generation open-source distributed object-based storage system that is designed for massive scalability, high performance and reliability. Ceph Storage Clusters use a distributed object storage service known as the Reliable Autonomic Distributed Object Store (RADOS), which provides applications with block, object, and file system storage in a single unified space.

Ceph is backed by a robust, worldwide open source community[18] effort with broad participation from major HPC and storage vendors. The Ceph Foundation[19] project is organised as a directed fund under the Linux Foundation. It exists to enable industry members to collaborate and pool resources to support the Ceph project community. The Foundation provides an open, collaborative, and neutral home for project stakeholders to coordinate their development and community investments in the Ceph ecosystem. Among its members are three European organisations: CERN, grnet and STFC, with CERN having the Associate Member Representative position.

15 https://www.hdfgroup.org/about-us/

16 https://www.ibm.com/us-en/marketplace/scale-out-file-and-object-storage

17 https://www.beegfs.io/content/

18 https://ceph.io/

19 https://ceph.io/foundation/

HPC system stack

OpenHPC

OpenHPC[20] is a Linux Foundation Collaborative Project whose mission is to provide a reference collection of open-source HPC software components and best practices, lowering barriers to deployment, advancement, and use of modern HPC methods and tools.

Thus, the objectives of the Project are: to create a stable and flexible open source HPC software stack, validated to run on a variety of hardware platforms; increase the simplicity and to reduce the cost of deploying and managing HPC systems and the performance and efficient utilization of HPC systems; include insights and technical contributions from across the HPC ecosystem, to integrate the leading edge work in the field and make it available to the community, and to adapt to new hardware insights and new technologies addressing scalability and performance, quickly responding to technology changes; and host the infrastructure for the open source project, establishing a neutral home for community meetings, events and collaborative discussions and providing structure around the business and technical governance of the Project.

OpenHPC is a community effort and there are a variety of ways to engage with the project, either as an end-user or developer. Some example participation opportunities include:

- install OpenHPC packages locally and provide feedback to the community
- host a local mirror of the OpenHPC repositories
- suggest additional components for selection; better yet, help with integrating and testing new components
- host hardware for build and test infrastructure
- share site knowledge through development of customized usage/install recipes
- participate in working groups to help establish community conventions and future directions
- register an OpenHPC cluster - https://openhpc.community/development/register-your-openhpc-system/
- join the conversations on OpenHPC mailing lists or checkout the OpenHPC code repository.

The project involves some prominent European research organisations (e.g. BSC, Cineca, CEA, Genci and LRZ) as well as non-European companies (ARM, CRAY, Fujitsu, Intel, Rd Hat).

Open Container Initiative

A container is a standard unit of software that packages up code and all its dependencies so the application runs quickly and reliably from one computing environment to another. A container image is a lightweight, standalone, executable package of software that includes everything needed to run an application: code, runtime, system tools, system libraries and settings.

The Open Container Initiative[21] (OCI) is another Linux Foundation project. It is a lightweight, open governance structure (project), formed for the express purpose of creating

20 https://openhpc.community/

21 https://www.opencontainers.org/

open industry standards around container formats and runtime. It involves the main leaders of the container industry

The Open Container Initiative provides an open source technical community within which industry participants may easily contribute to building a vendor-neutral, portable and open specification and runtime that deliver on the promise of containers as a source of application portability backed by a certification programme. The Open Container Initiative does not seek to be a marketing organisation, defines a full stack or solution requirements, and strives to avoid standardizing technical areas undergoing innovation and debate.

While the mission of OCI is not to create a complete stack, the format and runtime will:

- Provide a container format and runtime specification that enables portability across compliant runtimes
- Provide a robust stand-alone runtime that can directly consume the specification and run a container

This initiative has no European members and there are no prominent European members on its board or among its technical experts. The main non-European members are: Amazon Web Services, IBM, Intel, Microsoft, Facebook, Fujitsu, Cisco, Dell, Oracle, Red Hat.

Tools

Slurm

The Slurm Workload Manager (formerly known as Simple Linux Utility for Resource Management or SLURM), or Slurm, is a free and open-source job scheduler for Linux and Unix-like kernels, used by many of the world's supercomputers and computer clusters.

It provides three key functions:

- allocating exclusive and/or non-exclusive access to resources (computer nodes) to users for some duration of time so they can perform work,
- providing a framework for starting, executing, and monitoring work (typically a parallel job such as MPI) on a set of allocated nodes, and
- arbitrating contention for resources by managing a queue of pending jobs.

Slurm is the workload manager on about 60% of the TOP500 supercomputers[22]. Slurm uses a best fit algorithm based on Hilbert curve scheduling or fat tree network topology in order to optimise the locality of task assignments on parallel computers.

Slurm began development as a collaborative effort primarily by Lawrence Livermore National Laboratory, SchedMD, Linux NetworX, Hewlett-Packard, and Groupe Bull as a Free Software resource manager. Slurm is available under a GNU General Public License.

In 2010, the developers of Slurm founded SchedMD, which maintains the canonical source, provides development, level 3 commercial support and training services. Commercial support is also available from Bright Computing, Bull, Cray, and Science + Computing.

Slurm contributors include a number of European organisations and experts, for example: Barcelona Supercomputing Center, Bull/Atos, CEA, Jülich Supercomputing Center, Swiss National Supercomputing Centre, KTH and Trinity College Dublin.

22 https://en.wikipedia.org/wiki/Slurm_Workload_Manager

OpenStack

OpenStack is a free and open-source software platform for cloud computing, mostly deployed as infrastructure-as-a-service (IaaS), whereby virtual servers and other resources are made available to customers. The software platform consists of interrelated components that control diverse, multi-vendor hardware pools of processing, storage, and networking resources throughout a data centre. Users either manage it through a web-based dashboard, through command-line tools, or through RESTful web services.

OpenStack software controls large pools of compute, storage, and networking resources throughout a datacentre, managed through a dashboard or via the OpenStack API. OpenStack works with popular enterprise and open source technologies making it ideal for heterogeneous infrastructure. Hundreds of the world's largest brands rely on OpenStack to run their businesses every day, reducing costs and helping them move faster. OpenStack has a strong ecosystem, and users seeking commercial support can choose from different OpenStack-powered products and services in the marketplace.

OpenStack began in 2010 as a joint project of Rackspace Hosting and NASA. It is managed by the OpenStack Foundation[23], a non-profit corporate entity established to promote OpenStack software and its community. The OpenStack Foundation promotes the global development, distribution and adoption of open infrastructure with more than 500 companies, 105,000 community members from 187 countries around the world. The OpenStack Foundation was founded in September 2012 to provide an independent home for the OpenStack cloud operating system, which has since become one of the largest and most diverse open source projects in history.

The goal of the OpenStack Foundation is to serve developers, users, and the entire open infrastructure ecosystem by providing a set of shared resources to build community, facilitate collaboration and support integration of open source technologies. Primary activities include organising large-scale test infrastructure, community management, and bringing together more than 20,000 open infrastructure enthusiasts each year at global events including the Open Infrastructure Summit (formerly the OpenStack Summit).

The Board of Directors includes European organisations and companies such as: Ericsson AB - Ericsson Software Technology, Deutsche Telekom T-Systems and CERN.

2.4 Language

HPC specific

MPI (Message Passing Interface) – the Open MPI Project

Message Passing Interface is a standardised and portable message-passing standard designed by a group of researchers from academia and industry to function on a wide variety of parallel computing architectures. The standard defines the syntax and semantics of a core of library routines useful to a wide range of users writing portable message-passing programs in C, C++, and Fortran. There are several well-tested and efficient implementations of MPI, many of which are open-source or in the public domain. These have fostered the development of a parallel

23 https://www.openstack.org/foundation/

software industry, and encouraged the development of portable and scalable large-scale parallel applications.

The **Open MPI[24] Project** is an open source Message Passing Interface implementation that is developed and maintained by a consortium of academic, research, and industry partners. Open MPI is therefore able to combine the expertise, technologies, and resources from all across the High Performance Computing community in order to build the best MPI library available. Open MPI offers advantages for system and software vendors, application developers and computer science researchers.

Participation in the Open MPI project is both formal and informal. Informal participation is encouraged by working with the Open MPI team in an ad hoc manner, such as (but not limited to) engaging in discussions on mailing lists, reporting testing results, providing feature requests and bug reports, filing pull requests for code contributions, etc. Informal participants are not granted commit privileges to the Open MPI code base repositories, and do not have voting privileges on administrative matters. Individuals or organisations can request formal participation in the Open MPI Community as Members.

The project has the following working groups:

- Hardware Locality
- Network Locality
- MPI Testing Tool
- Open MPI User Docs
- Open Tool for Parameter Optimization
- And also mailing lists

This project has a number of European members and contributors: Atos, Center for Information Services and High Performance Computing (ZIH), Technische Universitaet Dresden, Computer Architecture Group, Technische Universtaet Chemnitz, HFT Stuttgart, University of Applied Science, HLRS, Hochschule Esslingen, University of Applied Sciences, Inria, The French National Institute for Research in Computer Science and Control and Friedrich-Schiller-Universitat Jena. The main non-European members (and contributors) are: Amazon Web Services, Inc., Nvidia, ARM, Facebook, Fujitsu, IBM, Intel, Mellanox and Oracle.

MPI Forum

The **MPI Forum**[25] is the standardisation forum for the Message Passing Interface (MPI). It is a self-organized effort that manages the meetings where the specifications are discussed then adopted. Europe has a strong position within the MPI Forum. The chair of the meetings and some of the working group leaders come from European organisations.

OpenMP

OpenMP is an application programming interface (API) that supports multi-platform shared memory multiprocessing programming in C, C++, and Fortran, on many platforms, instruction set architectures and operating systems, including Solaris, AIX, HP-UX, Linux, macOS, and Windows. Jointly defined by a group of major computer hardware and software vendors and

24 https://www.open-mpi.org/

25 https://www.mpi-forum.org/

major parallel computing user facilities, the OpenMP API is a portable, scalable model that gives shared-memory parallel programmers a simple and flexible interface for developing parallel applications on platforms ranging from embedded systems and accelerator devices to multicore systems and shared-memory systems.

OpenMP[26] ARB (Architecture Review Boards)

The OpenMP ARB (Architecture Review Boards) mission is to standardise directive-based multi-language high-level parallelism that is performant, productive and portable. The OpenMP ARB owns the OpenMP brand, oversees the OpenMP specification and produces and approves new versions of the specification.

The ARB is composed of permanent and auxiliary members. Permanent members are vendors who have a long-term interest in creating products for OpenMP. Auxiliary members are normally organisations with an interest in the standard but do not create or sell OpenMP products.

The Project runs a number of discussion forums, provides tool, use cases and training material.

The main European members are: BSC, Bristol University, EPCC, Inria, LRZ and RWTH. The main non-European members are: ARM, AMD, ASC/Lawrence Livermore National Laboratory, CRAY, IBM, Fujitsu and Intel.

The main European experts working on this initiative are: Xavier Martorell – BSC, Simon McIntosh-Smith – Bristol University, Volker Weinberg – LRZ and Mark Bull – EPCC.

GASPI Forum

GASPI[27] is a Partitioned Global Address Space (PGAS) approach offering the developer of parallel applications an abstract shared address space which simplifies the programming task and at the same time facilitates: data-locality, thread-based programming and asynchronous communication. This approach has been developed by Fraunhofer (see Section 4.4 for a discussion about the key success factors). The standard is managed by the GASPI Forum which has been founded by Fraunhofer ITWM and SCAI, the TU Dresden ZIH, the University Erlangen RRZE, the KTH Royal Institute of Technology, the Leibniz Supercomputing Center LRZ, the TU Munich LRR, IMEC and the European Centre for Medium-Range Weather Forecasts (ECWMF).

For accelerator

OpenACC

OpenACC[28] is a user-driven directive-based performance-portable parallel programming model designed for scientists and engineers interested in porting their codes to a wide-variety of heterogeneous HPC hardware platforms and architectures with significantly less programming effort than required with a low-level model.

26 https://www.openmp.org/ and https://www.openmp.org/about/about-us/

27 Global Address Space Programming Interface http://www.gaspi.de/

28 https://www.openacc.org/

Members of the OpenACC organisation include commercial, government, and academic organisations who value access to benefits that OpenACC has to offer including:

- Drive the OpenACC specification by setting directions and development priorities, as well as brainstorming new features with the technical team
- Influence the overall development of the organisation and the user community
- Opportunity to collaborate with fellow members and build relationships with OpenACC experts and technical developers that shape the specification

OpenACC compilers, profilers and debuggers are designed and available to download from multiple vendors and academic organisations.

The main European members are: Helmholtz Zentrum Dresden Rossendorf, Total and EPCC. The main non-European members are: CRAY, Nvidia, AMD and Oak Ridge National Laboratory.

The management includes one European person: Secretary - Guido Juckeland, Helmholtz-Zentrum Dresden-Rossendorf (HZDR).

Interpreted language

Julia

Julia[29] is a high-level, high-performance, dynamic programming language. While it is a general purpose language and can be used to write any application, many of its features are well-suited for high-performance numerical analysis and computational science.

Distinctive aspects of Julia's design include a type system with parametric polymorphism in a dynamic programming language; with multiple dispatch as its core programming paradigm. Julia supports concurrent, (composable) parallel and distributed computing (with or without MPI and/or the built-in corresponding to "OpenMP-style" threads), and direct calling of C and Fortran libraries without glue code. A just-in-time compiler that is referred to as "just-ahead-of-time" in the Julia community is used.

Julia is a fast tool as it was designed from the beginning for high performance. Julia programs compile efficient native code for multiple platforms via LLVM. Julia is dynamically-typed, feels like a scripting language, and has good support for interactive use. Julia has a rich language of descriptive datatypes, and type declarations can be used to clarify and solidify programs. Julia uses multiple dispatch as a paradigm, making it easy to express many object-oriented and functional programming patterns. It provides asynchronous I/O, debugging, logging, profiling, a package manager, and more. Julia has high level syntax, making it an accessible language for programmers from any background or experience level. The Julia microbenchmarks can be browsed to get a feel for the language.

Julia is provided under the MIT license, free for everyone to use. All source code is publicly viewable on GitHub[30].

The Julia community is global, with meetups all over the world and resources in a variety of different languages. A number of them take place in Europe:

- Barcelona Julia Meetup

29 https://julialang.org/

30 https://github.com/JuliaLang/julia/blob/master/LICENSE.md

- Berlin Julia Users Group
- JuliaCPH (Copenhagen)
- Dublin Julia Users Group
- JuliaStanbul (Istanbul)
- London Julia User Group
- BeNeLux Julia User Group
- Vienna Julia Meetup
- Warszawskie Forum Julia
- Zurich Julia Users Group
- Julia Nantes User Group (France)
- Julia Paris User Group

The Julia Seasons of Contributions (JSoC) are the seasonal programmes for funding and/or mentoring students and other developers to contribute to the open source ecosystem.

Domain specific language

A domain-specific language (DSL) is a computer language specialised to a particular application domain. This is in contrast to a general-purpose language (GPL), which is broadly applicable across domains. There are a wide variety of DSLs, ranging from widely used languages for common domains, such as HTML for web pages, down to languages used by only one or a few pieces of software. Domain Specific Languages are a promising approach to hide the complexity of HPC systems, apply domain specific optimizations and boost programmer's productivity.

DSL is an important trend for the future. We have started to see some efforts for the emergence of such languages as standards. Nevertheless, this field is still maturing.

2.5 Software packages

We have considered three software categories, the boundaries between them not being always so sharp.

The first category is related to high performance data analytics (HPDA) with Scikit-learn, Apache Spark and Microsoft Distributed Machine Learning Toolkit.

The second one is related to Machine Learning (ML) with Tensorflow, Caffé, PyTorch, Microsoft Cognitive Toolkit, Apache MXNet and Open Neural Network Exchange.

The last one is related to two specific topics, huge data movement and management with dCache and weather simulation and analysis workflows where the volume and rate of data lead to performance issues with Earth System Data Middleware.

All these pieces of software are available in an open source format.

Participation in these organizations is done almost only through software contributions. All the software is accessible in open source format and the usual way to participate is to contribute through the software itself. High level contributors have a higher weight and may in time evolve through the organisations' influence structures.

HPDA

Three main software have been considered. Two of them, issued from academic organisations, are quite active and widely used. One is European led by INRIA, the second one is American led by U. of California. It must be highlighted that quite some non-European organisations are involved in the European project. The third one has been developed by Microsoft.

Scikit-learn

Scikit-learn (https://scikit-learn.org/stable/) is a free open source software machine learning Python library for data mining and data analysis started in 2007 as a Google Summer of Code project. Scikit-learn is used on a regular basis by more than half a million of people in the world.

Since September 2018 the Scikit-learn @ INRIA foundation has had the mission of fostering the development of Scikit-learn and supporting the community that builds it.

The project and the foundation are supported by European members INRIA, AXA and BNP Paribas and by non-European members Microsoft, BCG Gamma, Fujitsu, Intel, Nvidia, Data Iku, Anaconda, Columbia University, Alfred P. Sloan Foundation and the University of Sydney.

The project seems to be very active presently.

A lot of European experts are involved including Jérémie du Boisberranger (U. de Versailles-Saint Quentin en Yvelines), Loïc Estève (INRIA), Alexandre Gramfort (INRIA), Olivier Grisel (INRIA), Guillaume Lemaitre (Scikit-learn @ INRIA foundation), Jan Hendrik Metzen (Bosch Center for Artificial Intelligence (BCAI)), Vlad Niculae (Institut de Télécommunication de Lisbonne), Bertrand Thirion (INRIA), Joris Van den Bossche (Paris-Saclay Center for Data Science), Gael Varoquaux (INRIA), Noelle Varoquaux (INRIA)and Roman Yurchak (Symerio).

Apache Spark

Apache Spark (https://spark.apache.org/)is a unified open-source analytics engine for large-scale data processing. It is led by U. of California.

The project is supported by a huge number of organisations most of which are non-European.

The project seems to be very active presently. Few European experts seem to be involved.

Microsoft Distributed Machine Learning Toolkit

The goal of Microsoft Distributed Machine Learning Toolkit (DMTK) (http://www.dmtk.io/) is to make machine learning tasks on big data highly scalable, efficient, and flexible.

DMTK is a platform designed for distributed machine learning and is not focused on deep learning.

The Distributed Machine Learning Toolkit has been developed and supported by the Artificial Intelligence Group of Microsoft Research Asia.

The project does not seem to be very active presently. Few European organisations or experts seem to be involved.

Machine Learning

This field is dominated by American organisations and most major US digital economy leaders such as Google, Facebook, Microsoft and Amazon are involved in at least one project. Except Caffé, all projects seem to be very active and widely used. Europe is poorly represented.

The Open Neural Network Exchange is a transverse private US initiative to enable AI developers to use models with a variety of frameworks, tools, runtimes, and compilers. It is supported by Facebook, Microsoft and Amazon among others.

TensorFlow

TensorFlow (https://www.tensorflow.org/) is an open source library to create machine learning models (deep learning) for desktop, mobile, web, and cloud.

TensorFlow is developed by Google Brain Team.

The project seems to be very active presently. Few European experts seem to be involved.

Caffé

Caffé (https://caffe.berkeleyvision.org/) is a deep learning open source framework developed by U. of California Berkeley AI Research and by community contributors.

The project does not seem to be very active presently, its successor is PyTorch. Few European organizations or experts seem to be involved.

PyTorch

PyTorch (https://pytorch.org/)is an open source machine learning library developed by Facebook.

The project seems to be very active presently. Few European organizations and experts seem to be involved.

Microsoft Cognitive Toolkit

The Microsoft Cognitive Toolkit (CNTK) (https://docs.microsoft.com/en-us/cognitive-toolkit/) is an open-source toolkit for commercial-grade distributed deep learning.

The software is developed by Microsoft Research. The project seems to be very active presently. Few European organizations or experts seem to be involved.

MXNet

Apache MXNet (https://mxnet.apache.org/) is an open-source deep learning software framework, used to train, and deploy deep neural networks.

MXNet is supported by public cloud providers including Amazon Web Services and Microsoft Azure. Currently, MXNet is supported by 60 organisations listed on the project web site including Intel, Baidu, Microsoft and Wolfram Research, and research institutions such as Carnegie Mellon, MIT, the University of Washington, and the Hong Kong University of Science and Technology.

The project seems to be very active presently. Few European organizations or experts seem to be involved.

Open Neural Network Exchange

The Open Neural Network Exchange (ONNX) (https://onnx.ai/) is an open format built to represent machine learning models. ONNX defines a common set of operators - the building blocks of machine learning and deep learning models - and a common file format to enable AI developers to use models with a variety of frameworks, tools, runtimes, and compilers.

ONNX is developed by Facebook, Microsoft and Amazon and is supported by a large set of organisations including Alibaba, Baidu, The MathWork and in Europe, Siemens.

The project seems to be very active presently. Few European experts seem to be involved.

Data oriented framework

dCache

dCache (https://www.dcache.org/) is an open source system for storing and retrieving huge amounts of data, distributed among a large number of heterogeneous server nodes, under a single virtual filesystem tree with a variety of standard access methods.

dCache is a joint venture between three organisations two of them being European, the Deutsches Elektronen-Synchrotron DESY and the Nordic Data Grid Facility NDGF and one of them being American, the Fermi National Accelerator Laboratory FNAL.

dCache is widely used by CERN. The project seems to be very active presently. Quite some European experts are involved including Patrick Fuhrmann of the Deutsches Elektronen-Synchrotron DESY and Tigran Mkrtchyan of the Deutsches Elektronen-Synchrotron DESY.

Earth System Data Middleware

The open source Earth System Data Middleware (ESDM) (https://github.com/ESiWACE/esdm) aims at deployment in both climate and weather simulation and analysis workflows where the volume and rate of data lead to performance and data management issues with traditional approaches.

It has been developed under the ESiWACE European Centre of Excellence funded from the European Union's Horizon 2020 Research and Innovation Programme. Organizations involved in the development are U. of Reading, Centro Euro- Mediterranean sui Cambiamenti Climatici (CMCC), SEAGATE and Deutsches Klimarechenzentrum.

The project seems to be very active presently. Some European experts (all of them being from ESiWACE) are involved including Luciana Pedro, Jakob Lüttgau, Julian Kunkel, Bryan Lawrence, Alessandro d'Anca, Cosimo Palazzo, Sandro Fiore, Paola Nassisi, Alessandra Nuzzo, Marco Chiarelli, Maria Mirto, Giuseppe Congiu and Huang Hua.

It is worth mentioning other initiatives in the field of Open Data where we see different organisations such as DOI (Digital Object Identifier https://www.doi.org/), RDA (Research Data Alliance https://rd-alliance.org/), GO FAIR (Findable, Accessible, Interoperable and Reusable data https://www.go-fair.org/) or WDS (World Data System http://www.icsu-

wds.org/) aiming at proposing approach to locate and reuse data. However, as these approaches are not HPC specific we have not included them in the analysis of the European position in HPC relevant standard organisations.

3 Standard organisation analysis

3.1 Weight of European organisations and presence of European experts

To present an overall view of the importance of HPC standard oriented organisations for the future of HPC, the presence of European organisations and the existence of European influential experts we have put the different information gathered during our survey in a table, with a colour code.

	strategic position for future HPC			
minor role				
important				
strategic				
		presence of European organizations		
	weak			
	medium			
	important			
			participation of academic European experts	participation of industrial European experts
		negligeable		
		intermediate		
		strong		

Hardware oriented standard organisations

organisation name	objective	strategic position for future HPC	presence of European organizations	partipation of academic European experts	partipation of Industrial European experts
OpenPOWER	development of an ecosystem around POWER architecture and the related technologies				
ARM and Linaro	development of an ecosystem around ARM architecture and the related technologies				
RISC-V foundation	development of an ecosystem around RISC-V architecture and the related technologies				
OpenHW Group	development of open hardware solution based on RISC-V				
GEN-Z	provide high-speed, low-latency, memory-semantic access to data and devices via direct-attached or network topologies				
CCIX	adoption of an industry standard specification to enable coherent interconnect technologies between general-purpose processors and acceleration devices				
CLX	standard interconnect offering high-bandwidth, low-latency connectivity between host processor and devices such as accelerators, memory buffers, and smart I/O devices				
OpenCAPI	open coherent high performance bus interface				
JEDEC	memory technology standard (DRAM, HDM, NVM, SSD)				
PCI-SIG	non-proprietary peripheral component interconnect technology				
Infiniband TA	InfiniBand Architecture specification defining hardware transport protocols sufficient to support both reliable messaging and memory manipulation semantics				
INCITS	storage technology standards as SCSI, ATA, Fiber channel,				
SNIA	standards are primarily related to data, storage, and information management				
Open Compute	collaborative community focused on redesigning hardware technology to efficiently support the growing demands on compute infrastructure				
ASHREA	heating, ventilation, air conditioning and refrigeration standards				

Table 1: Position of Europe in hardware oriented standard organisations

HPC software stack

organisation name	objective	strategic position for future HPC	presence of European organizations	partipation of academic European experts	partipation of Industrial European experts
Linux Foundation	supports the creation of sustainable open source ecosystems by providing financial and intellectual resources, infrastructure, services, events, and training.				
Lustre	an open-source, parallel file system that supports many requirements of leadership class HPC simulation environments.				
HDF	supports n-dimensional datasets and each element in the dataset may itself be a complex object. HDF5 is a general purpose library and file format for storing scientific data.				
Spectrum Scale	high-performance clustered file system software developed by IBM. It can be deployed in shared-disk or shared-nothing distributed parallel modes.				
BeeGFS	(formerly FhGFS) is a parallel file system, developed and optimised for high-performance computing.				
CEPH	a next generation open-source distributed object-based storage system that is designed for massive scalability, high performance and reliability.				
OpenHPC	A Linux Foundation Collaborative Project whose mission is to provide a reference collection of open-source HPC software components and best practices, lowering barriers to deployment, advancement, and use of modern HPC methods and tools.				
Open Container Initiative	another Linux Foundation project. It is a lightweight, open governance structure (project), formed for the express purpose of creating open industry standards around container formats and runtime. It involved the main leaders of the container industry				
Slurm	a free and open-source job scheduler for Linux and Unix-like kernels, used by many of the world's supercomputers and computer clusters.				
OpenStack	a free and open-source software platform for cloud computing, mostly deployed as infrastructure-as-a-service (IaaS), whereby virtual servers and other resources are made available to customers.				

Table 2: Position of Europe in HPC software stack oriented standard organisations

Language

organisation name	objective	strategic position for future HPC	presence of European organizations	partipation of academic European experts	partipation of Industrial European experts
Open MPI Project	an open source Message Passing Interface implementation that is developed and maintained by a consortium of academic, research, and industry partners.				
OpenMP ARB (Architecture Review Boards)	Its mission is to standardise directive-based multi-language high-level parallelism that is performant, productive and portable.				
GASPI Forum	API to support a PGAS approach in HPC languages				
OpenACC	a user-driven directive-based performance-portable parallel programming model designed for scientists and engineers interested in porting their codes to a wide-variety of heterogeneous HPC hardware platforms and architectures with significantly less programming effort than required with a low-level model.				
Julia	a high-level, high-performance, dynamic programming language.				
Domain Specific Language (DSL)	a computer language specialised to a particular application domain.				

Table 3: Position of Europe in HPC language oriented standard organisations

Software packages

organisation name	objective	strategic position for future HPC	presence of European organizations	partipation of academic European experts	partipation of Industrial European experts
Scikit-learn	data mining and data analysis open source				
Spark	large-scale data processing open source				
Microsoft Distributed Machine Learning Toolkit	open source to make machine learning tasks on big data highly scalable, efficient, and flexible				
TensorFlow	open source to create machine learning models				
Caffé	deep learning open source framework				
PyTorch	open source machine learning library				
Microsoft Cognitive Toolkit	open-source toolkit for commercial-grade distributed deep learning				
MXNet	open-source deep learning software framework				
Open Neural Network Exchange	An open format built to represent machine learning models and a common file format to enable AI developers to use models with a variety of frameworks, tools, runtimes, and compilers				
dCache	open source system for storing and retrieving huge amounts of data				
Earth System Data Middleware	open source to manage weather simulation and analysis workflows where the volume and rate of data lead to performance and data management issues with traditional approaches				

Table 4 Position of Europe in software package oriented standard organisations

3.2 Synthesis

In the tables above, the column 'Strategic position for future HPC' reflects the views of the authors of this document and should not be considered as the only assessment available.

In all organisations, Europe's presence lags behind that of the other leading regions. With the exception of a few cases (BeeGFS and Scikit-learn), Europe's presence in all strategic international standard organisations is rather weak.

Also, Europe's presence is limited to the representatives of a few projects or mainly academic organisations, which have a vested interest in developing and maintaining standards. It also seems that the European HPC players are not as active as other, non-HPC IT companies from Europe. This is in contrast with the other regions as in the case of most organisations, it is companies (of US origin mainly) that lead their work.

All of the standard organisations are controlled by US-based organisations. Also, the other two prominent HPC ecosystems, Japan and China, have been making headway strengthening their presence and thus have a direct influence on standards. In particular, Chinese companies are very active in the hardware area. This also applies to the organisations controlling AI-related standards, which are US-dominated.

In conclusion, in order to keep pace, Europe must invent a mechanism to ensure its continuous presence in standard management.

4 Interviews

In this chapter we describe the actions undertaken to ascertain the position of the European players regarding this issue of standardisation. At first the idea was to organise an online survey targeting the ETP4HPC members. Nevertheless, as the experience of this kind of survey for other topics has been quite disappointing due to the small number of answers and because we were not looking for quantitative inputs but rather qualitative feedbacks, we decided to organize a set of interviews with pre-selected organisations.

The following sections present first the objectives behind the interview, second how it has been handled and third the main findings.

4.1 Objectives

The first objective was to know in what domains there are stakes for the European organisations in HPC related standards. The list of standard related organisations that was used for the analysis described in Chapter 3 has been the basis of the discussion. Nevertheless, new suggestions from the interviewees were welcome and looked for.

The second objective was to understand the current level of effort made by the interviewees to influence standards or to participate in standard oriented organisations. Related to this topic, we investigated also the level of management support for this standardisation effort. The way the standardisation effort is organized was also among the questions. We also look for issues that are currently slowing down the European players presence in standard organisations.

The third objective was to discover what would be the requests of the European players for the future in terms of standardisation. We investigate the standardisation topics that could be important for European players as well as the measures to boost the European activity. We tried to be as open as possible to get suggestions that are new and innovative.

4.2 Execution

To fulfill the objectives, we created an interview template (see Annex 7.3). Nevertheless, this template was just the basis for the discussion and we tried to conduct the interview in a very open way to get as many new ideas as possible from the interviewees.

Instead of trying to have a lot of short interviews, we decided to select a small panel and to conduct long interviews (between 1 and 2 hours each). This was possible as most of the interviewees were known contacts willing to invest time in the topic.

Despite this small number of interviews, we think that we got a good coverage with the panel in terms of range of standard topics, of different size of stakeholders and of approaches to standardisation. The coverage is illustrated by the following table:

Type of stakeholders	Organisations	Hardware standards coverage	Middleware standards coverage	Programming language standards coverage	Software packages standards coverage
Large company	Atos	X	X	X	X
Large company	Seagate	X	X	X	
SME	E4	X	X		
SME	UCiT		X	X	X
SME	Submer	X	X		
University	BSC	X	X	X	X

Table 5: Coverage of standards by the interviewees

After each interview, a document was created by the EXDCI-2 team and sent to the interviewee for review. After a potential feedback, a final version of the interview was issued. An example of such a document can be found in Annex 7.4.

4.3 Findings

The first finding is that almost all the interviewees consider that they do not have a standardisation effort at the right level with a clear under investment. This is the case both for large companies and SMEs. They would like to be more present in more standard oriented organisations. Most of the time, this would serve acquiring expertise onthe related developments and ensuring that their offering remains compatible with the tomorrow standards. Sometimes, they would also like to push their technologies in the standardisation process (mostly in software related topics). The main reason for this under-investment is a lack of resources of the European players that could be allocated to this activity. Right now, there is no process to share the effort and to cooperate in order to decrease the cost for each player. The connection with academic experts is not developed and even when some academic experts are influential in the standard related organisations, this has not a lot of impact on the position of the European industry. It also has been mentioned that good experts in standard related activities have specific profiles that are not easy to find in Europe. Some of the interviewees lack this kind of profiles to sustain their efforts. The cost of the fees is also mentioned especially by the SMEs as an obstacle for more efforts on standard.

When it comes to propose solutions to change this negative situation, some suggestions have been made:

- Put in place a process for several stakeholders to share the effort to be present in a given standard organisation; as most of the time, the activity in standard related organisations is pre-competitive, European players are ready to cooperate and to exchange information between them;
- Set up a financial support for strategic topics when this will be of benefit to several European players (in domains such as interconnect networks, HPC related OS features, new programming models,etc);

- Set up groups of interest for European experts to discuss and to select the ideas to be pushed in standard organisations
- Share information about who is involved in standard efforts; a kind of "who's who" for standards
- Set up a dialogue between technology providers (proposing new solutions that can become de facto standards) and the European application communities to initiate a user base for the new solutions that could become de facto standards
- Push several HPC FET projects to team up in creating the critical mass to initiate the virtual circle between applications and standard API.

Another important finding is that some European players would like to see new standards to be put in place in order to allow the development of new products or new services that can not exist within a proprietary or too fragmented environment. The main domains where such new standards are expected by European players are Big Data/AI and HPC as a service (proposed by cloud providers). In these new domains, European SMEs have ideas for innovative services or products but due to the absence of standards it is either too complex to develop them or too difficult to prove the benefits that they bring.

The following topics have been suggested as good candidates for standards:

- Benchmark for AI workloads
- API to get the energy consumption data of different cluster components (processors, GPUs, servers, network adaptors, network switches)
- API to collect data movement information in a cluster
- Standards to provide SLA for storage
- Standard API to deploy cluster as a service on different cloud providers or computing centre infrastructures
- API for orchestration in the digital continuum
- API for FPGA programming
- API for serverless computing
- API for IO in the different HPC programming models (MPI, OpenMP, PGAS)
- Specifications for chiplet integration defined by European players

Most of the interviewees are positive about the ability of European players to organise themselves and to successfully develop new de facto standards. Nevertheless, some organisations have to take the lead. Different organisation names such as ETP4HPC, CLAIRE, PRACE or EuroHPC are mentioned for being in charge of this first move that will initiate the effort in which the European players would be happy to contribute.

4.4 Other discussions with stakeholders

Beside the interviews, we had some additional discussions with European stakeholders related to standard issues. They confirmed the findings described above. One thing that can be added come from a discussion with Fraunhofer-Institute ITWM on the experience that they have in promoting their work on GASPI[31] as a standard.

GASPI (Global Address Space Programming Interface) is a PGAS (Partitioned Global Address Space) language API. The first work on the topic started in 2005 and now GASPI is a well acknowledged API that is used by the HPC community and continues to be enhanced in some

[31] http://www.gaspi.de/gaspi/

European projects. From the discussion with ITWM, it appears that several conditions have been key to make this effort a de facto standard with a user community:

- At the beginning, a disruption technology related to hardware support of RDMA which allows for efficient movement of data between distributed nodes
- Some hardware solution (such as Infiniband) to provide the acceleration
- An implementation (GPI and now GPI-2) of the GASPI standard that can be used to demonstrate the benefit of the API
- An industrial application that shows the advantages of the new standard (in the case of GASPI the "killer" application was in the seismic domain for the oil and gas industry)
- The organisation of a consortium to promote the standard (the GASPI forum).

This example shows that a European organisation can be successful in establishing a new de facto standard that is then used by the HPC community. There are of course several conditions to be met, but you can quite often in the evolution of HPC find such a context that is favourable to the emergence of a de facto standard: new performance increase coming from innovative solutions, good implementation, applications that demonstrate benefits and a good choice of partners to push the first standard.

5 Recommendations

In this section, we present recommendations that could help Europe to attain a better position and to exert more influence in the emergence and setting of standards applicable in the HPC systems. Each recommendation below is accompanied by a short discussion on a potential way to implement it. The implementation options proposed here are not the only ones available and other possibilities may exist.

These recommendations have been derived from the analysis above which is focused on HPC. Nevertheless, most of these conclusions could be applied to other IT fields.

We have categorised the recommendations into two sets. We start with proposals that can be put in place with few additional resources and that will help optimise the current situation. Then, we present recommendations that can improve the position of Europe and that would need more efforts and some funding to be effective. In this category, the proposed actions can make Europe a leader in some standardisation efforts. If standardisation is to be considered a priority by Europe, the two categories of recommendations should be implemented.

5.1 Achieving a more efficient organisation

The analysis performed during this EXDCI-2 activity has gathered information about standard organisations and the presence of European organisations in these bodies. This is valuable information that will be shared with the community through the EXDCI-2 web site. Nevertheless, this information is dynamic and can change over time: for example, changes in participation at the level of experts or working groups, the emergence of new standardisation efforts, etc. It would be interesting to have a continuously updated landscape of these standard organisations.

With the emergence of the digital continuum and the convergence of AI/BD/HPC, the coverage of all relevant fields needs to involve different communities. The European Technology Platforms (ETPs) (or their equivalent or sub-bodies) in the different domains (IoT, Big Data, Artificial Intelligence, HPC) could be the right place to undertake this effort to maintain a map of the relevant standard organisations. We believe that a small effort (2PM[32]/Year/ETP) could be enough to perform this activity.

R1 Achieve a continuous survey of emerging standards, produce maps of new standard organisations relevant in the context of the digital continuum, maintain all the maps up to date and disseminate this information.

As seen in Chapter 3, we have already some European experts active in some standard organisations. With little additional effort, we believe that their impact can be increased. If some contacts between these experts and a broader set of European players are organised before and after the important standardisation decisions a win-win situation can ensue.

If the stakes of standardisation meetings are discussed before between the experts and other stake-holders, on the one hand more European organisations will be aware of them and on the other hand the experts can speak during the meeting in the name of more organisations than

[32] Person Month

their own and have more weight. Similarly, after the meetings some information about the decisions and the new agenda items can be shared by the experts.

Some may object that the organisation paying for the expert will not want to share the information about the standardisation activity. We believe that most of the time this activity is a pre-competitive one and that the organisation can see the benefits of having more weight by representing the European ecosystem and in being able to be in a "receiving" mode in the case of other standard organisations (being informed by other experts without paying for the active participation). Of course, the agreement of the expert is mandatory to implement this kind of network around her/him.

These networks around the standardisation experts can be organized by the relevant ETPs. It can be implemented by specific mailing lists and teleconferences prior to and after important meetings. Ideally, a person should be designated for each standard to animate the network and possibly act as an ambassador. In terms of resource, this action is costlier as it means time from the people involved in the networks and some processes and manpower to be put in place by the ETPs. However, most of the costs will be self-funded by the participants.

R2 Establish a network of interested people around each European standardisation expert to share the information about the standard activity.

5.2 Achieving a more active European ecosystem: making Europe a standard leader

We would like to propose recommendations for a more active European ecosystem in the domain of HPC standardisation. The main idea is to be proactive in developing the European presence in such activities.

The first two recommendations have the same objective of having more European experts in standard organisations. The first is a top-down approach, while the second one a bottom-up approach.

Europe can have an organisation in charge of selecting European champions/delegates/representatives to join the standards bodies considered as strategic and represent not only their own mother organisations but the whole community.

This organisation can be under the framework of EuroHPC or ETPs. Anyhow, this organisation should work closely with EuroHPC and all HPC research projects and organisations. It should also interact with others similar projects and organisations such as the StandICT project to benefit from their network and experience.

The two main axes to be considered are: first identify the upstream European novel technologies and topics which could be leverage in a standard mode. Second, for strategic standard organisations, incentivise European organisations and experts to compensate at least partly their involvement in standard organisations.

R3: Implement a specific HPC organisation with sufficient financial and human resources, whose goal will be to develop a European HPC presence in standard organisations.

The initiative can also directly come from the stake-holders. A group of European organisations can join efforts in order either to participate in an existing standard organisation or to create a

new initiative aiming at the emergence of a standard. To support and encourage such an initiative a funding mechanism can be established.

This mechanism must be a light weight process that is continuously available. For example, it can be implemented through cascade funding with a counter for interested organisations to send a proposal for setting up a standardisation effort. The consortium eligibility rules can be broad (for example at least three organisations from at least two Member States). The funding can cover the cost of the technical work for either participating in the standard activities or writing the emerging standard specifications.

R4: Put in place a process to help consortia organise their presence in a standard organisation or write emerging standard specifications.

To allow Europe be more active in standards, one should also look at leveraging the EC funded projects. Some of the projects develop technologies that could become a standard. Nevertheless, because not anticipated from the beginning, most of the time not enough resources are available inside the project to be effective in pushing the technology as a standard. Participants of these projects with good standard candidates can be allowed to receive specific funding for this activity and to gather the good players for a success story (see the GASPI example mentioned in Section 4.4).

Again, the mechanism must be a light weight process that is continuously available. It can be implemented through cascade funding with a counter for interested projects to send a proposal for setting up the standardisation effort. The consortium can differ from the project one with only a subset of the project partners complemented by additional partners in order to gather the organisations well suited for this standardisation effort.

R5 Establish a process for participants to European funded projects to apply for a specific funded effort to facilitate the emergence of a standard.

In the course of this task and also more globally during the EXDCI-2 project, some topics have emerged as strategic for Europe to become an HPC leader. Several good ideas have also emerged from the interviews conducted during this task (see Section 4.3). They can be developed by the stake-holders thanks to the implementation of R4. However, in some domains, we see some urgency to act proactively with a direct initiative to work on the emergence of standards. At least four topics are candidates for direct actions:

- Integration of chiplets: as discussed in Section 2 and as seen in EXDCI-2 Work Package 2, HPC is going in the direction of heterogeneous nodes with different kinds of accelerators tightly integrated. If European players want to get a place in this market, they require a standard for the integration of the chips that they will develop.
- Provisioning of storage and computing resources: with the development of the digital continuum, the developers of applications require a standard to be able to build and deploy their workflows using different software and hardware.
- Energy consumption and management: to minimise the environmental impact of HPC applications, it is mandatory to have a standard way to gather information about energy consumption and then to implement optimization actions through a standard interface.
- Workflow deployment: with the development of smart twins, data sources and AI capabilities, application developers would benefit from a standard to build and deploy complex integrated software workflows.

In these three domains, the existence of standards will greatly help European players to develop their activities. As already mentioned, due to the dynamic nature of HPC, additional domains can appear.

R6 Identify and launch standardisation projects in strategic domains, to establish a level playing field, to favour European players' new initiatives and to encourage green IT.

These four recommendations can make Europe a leader in the field of HPC standards. The R3 is the more structural one and the three others could be seen as actions that will be under the responsibility of the organisation resulting from implementing this R3.

5.3 Conclusion

The rise of HPC in the digital continuum is creating new challenges in terms of standards. Europe can be a significant player and develops a sound and dynamic environment to facilitate the emergence of new scientific, industrial and societal applications of HPC with a more active approach to standards. The proposed recommendations can push in this direction and make Europe a leader in some standard domains which foster an open, sustainable and competitive IT ecosystem.

6 Acknowledgments

We would like to thank the people who, through discussions with the authors, have provided valuable information and comments. Specifically, we thank Laurent Cargemel, Antonio Pena, Fabrizio Magugliani, Sai Narasimhamurthy, Daniele Rispoli, Philippe Bricard, Benjamin Depardon, Valeria Bartsch, François Bodin, Gabriel Staffelbach, Mark Ash and Stephane Requena.

7 Annex

Annex list:

1. Standard organisation analysis template
2. Examples of standard organisation description
    - RISC-V foundation
    - Linux Foundation
    - Scikit-learn
3. Interview template
4. One example of interview data

7.1 Annex 1 : Template for standard organisation description

name organization

7.2 Annex 2: Examples of standard organisation descriptions

RISC V Foundation organisation

Organisation description

Objective

The RISC-V Foundation comprises more than 325 members building the first open, collaborative community of software and hardware innovators powering innovation at the edge forward. Born in academia and research, the RISC-V ISA delivers a new level of free, extensible software and hardware freedom on architecture, paving the way for the next 50 years of computing design and innovation.

Legal form

The RISC-V Foundation, a non-profit corporation controlled by its members, directs the future development and drives the adoption of the RISC-V ISA.

In November 2018, the RISC-V Foundation announced a joint collaboration with the Linux Foundation. As part of this collaboration, the Linux Foundation will also provide an influx of resources for the RISC-V ecosystem, such as training programs, infrastructure tools, as well as community outreach, marketing and legal expertise.

Address- Web site

https://riscv.org/

Membership

Condition

Members of the RISC-V Foundation have access to and participate in the development of the RISC-V ISA specifications and related HW / SW ecosystem.

Several membership levels:

- The Platinum US$25,000 per year. Platinum members are eligible for Board seat elections and to Chair Foundation Technical Committees, Marketing Committees and Task Groups.
- Gold US$10,000. Gold members are eligible to Chair Foundation Technical Committees, Marketing Committees and Task Groups.
- The Silver US$5,000. All organisations have one vote per open position in Board elections.
- Intended for non-profit research labs and Universities, the Auditor membership level is without membership fees
- The Individual membership level (non-voting), is in place for individual contributors unaffiliated with any particular company or organisation. There are no fees associated with this membership.

Main European members

Platinum: NXP, Thales

Gold: BAE, Nokia

Silver: BSC, CEA, ETHZ, Greenwave Technologies, Infineon, KU Leuven, Raspberry PI

Auditor: Forth, INRIA, Uppsala University

Main non-European members

See board member list

Working process

Mode of organisation

Work groups are established in different domains:

Main working groups

Special Interest Group dedicated to collaboration of HPC interests. Contact email hpc@riscv.org.

Current work

Current version is

The RISC-V Instruction Set Manual

Volume I: Unprivileged ISA

Document Version 20190608-Base-Ratified

June 8, 2019

People

Direction and Staff

The Foundation has a Board of Directors comprising seven representatives from Bluespec, Inc.; Google; Microsemi; NVIDIA; NXP; University of California, Berkeley; and Western Digital.

Main European experts

BSC team

Luca Bellini

Main non European influence people

Calista Redmond, CEO of the RISC-V Foundation

Board members

- Krste Asanovic – Chairman of the Board, professor in the EECS Department at the University of California
- Zvonimir Z. Bandić is a research staff member and senior director of Next Generation Platform Technologies at Western Digital Corporation
- Charlie is CEO of Bluespec, Inc., a provider of high--level tools and IP for ASIC and FPGA design
- Rob Oshana is vice president of software engineering R&D for NXP Microcontrollers; Chairman of the Board for the OpenHW Group
- David Patterson as UC Berkeley's Computer Science Division chair
- Frans Sijstermans is a vice president of engineering at NVIDIA
- Ted Speers Microsemi's SoC Group

Linux Foundation

Linux Foundation - Organisation description

Objective

The Linux Foundation supports the creation of sustainable open source ecosystems by providing financial and intellectual resources, infrastructure, services, events, and training. Working together, The Linux Foundation and its projects form the most ambitious and successful investment in the creation of shared technology.

Legal form

THE LINUX FOUNDATION - an Oregon non-profit mutual benefit corporation, Effective as of May 18, 2018

Address- Web site

https://www.linuxfoundation.org/

Membership

Condition

Corporate Membership:

- Platinum Members
- Gold Members
- Silver Members
- and one class of non-voting Affiliate participants, called Associates.

Individual Supporters (USD 49 - free for Students)

Main European members

ABN-AMRO (a Dutch bank)

Airbus

Atos

BBVA

Comarch

Deutsche Bank

Ericsson

Siemens

Total

Universities: EPFL, Budapest Technical University, University of Cambridge

GRNET

Main non-European members

AT&T

Cisco

IBM

Intel

Fujitsu

Huawei

Microsoft

Google

Hitachi

NEC

Oracle

Qualcomm

Samsung

VMWare

Working process

Mode of organisation

Open Source Project Hosting - The Linux Foundation hosts many of the most important open source projects in the world, including Linux.

There are a few different ways to host a project with The Linux Foundation, and they're all guided by each community's requirements and goals. Some communities choose to raise funding, but many do not. Several projects have dedicated staff, while others are driven by individual contributions of time and energy. Many projects have thousands of developers, whereas others start with a few and want to attract more. Most generate open source code, but projects that provide infrastructure support for open source communities, standards, or specifications also choose The Linux Foundation as their home.

Minimum Requirements For Projects

Linux Foundation projects meet the following requirements:

- They use a license approved by the Open Source Initiative.

- They allow The Linux Foundation to own community assets, like a domain or trademark, on behalf of the project community.
- At least one Linux Foundation member sponsors each project.
- They have open, neutral governance. Anyone who follows the project's contributions terms and supports the "do-ocracy" form of governance can participate. Membership is never required to participate in a Linux Foundation project technical community.

Main working groups

Work takes place in projects hosted by the Foundation:

https://www.linuxfoundation.org/projects/

Current work

All projects within the remit of the organisation:

https://www.linuxfoundation.org/projects/

People

Direction and Staff

The Linux Foundation Board of Directors is comprised of 22 senior leaders from across the IT industry. Board members represent Linux Foundation members and the Linux developer community, and set the strategic direction for the organisation.

https://www.linuxfoundation.org/about/board-members/

The Linux Foundation leadership team includes experts in business, open source, and ecosystem development. Team members help open source communities build sustainable ecosystems that advance software development and introduce the principles of shared R&D across industries.

https://www.linuxfoundation.org/about/leadership/

Fellows:

The technologists creating the most important software in the world shouldn't work for any one company. The Linux Foundation provides a neutral place for them to advance their work

https://www.linuxfoundation.org/about/linux-foundation-fellows/

Main European experts

None

Main non-European influential people

Chris Mason, Facebook

Chris Rice, AT&T

David Ward, Cisco

Daniel Park, Samsung

Imad Sousou, Intel

Jessica Murillo, IBM

John Gossman, Microsoft

Sarah Novotny, Google

Wim Coekaerts, Oracle

Scikit-learn

Organisation description

Objective

Scikit-learn is a free open source software machine learning Python library for data mining and data analysis started in 2007 as a Google Summer of Code project by David Cournapeau. Scikit-learn is used on a regular basis by more than half a million of people in the world.

Legal form

New BSD licence.

Address- Web site

https://scikit-learn.org/stable/

Membership

Condition

The open-source project is fully open.

In Sept 2018 the Scikit-learn @ INRIA foundation is implemented with the mission to foster the development of Scikit-learn and support the community that builds it.

Main European members

The project is supported by INRIA, AXA, BNP Paribas.

The foundation is supported by INRIA, AXA, BNP Paribas.

Main non-European members

The project is supported by Microsoft, BCG Gamma, Fujitsu, Intel, Nvidia, Data Iku, Anaconda, Columbia University, Alfred P. Sloan Foundation, the University of Sydney.

The foundation is supported by Microsoft, BCG Gamma, Fujitsu, Intel, Nvidia, Data Iku.

Working process

Mode of organisation

The software is under the control of a technical committee.

The Scikit-learn @ INRIA foundation is managed by a technical committee under an advisory board.

Main working groups

Current work

Release 0.22.1 January 2, 2020

People

Direction and Staff

The Technical Committee of scikit-learn consists of Alexandre Gramfort, Olivier Grisel, Andreas Müller, Joel Nothman, Hanmin Qin, Gaël Varoquaux and Roman Yurchak.

The Consulting Committee of scikit-learn foundation consists of Alexandre Gramfort, Olivier Grisel, Andreas Müller, Joel Nothman, , Gaël Varoquaux. The foundation team consists of Jérémie du Boisberranger, Olivier Grisel et Guillaume Lemaître.

Main European experts

Jérémie du Boisberranger	U. de Versailles-Saint Quentin en Yvelines
Loïc Estève	INRIA
Alexandre Gramfort	INRIA
Olivier Grisel	INRIA
Guillaume Lemaitre	Scikit-learn @ INRIA foundation
Jan Hendrik Metzen	Bosch Center for Artificial Intelligence (BCAI)
Vlad Niculae	Institut de Télécommunication de Lisbonne
Bertrand Thirion	INRIA
Joris Van den Bossche	Paris-Saclay Center for Data Science
Gael Varoquaux	INRIA
Nelle Varoquaux	INRIA
Roman Yurchak	Symerio

Main non-European influence people

Tom Dupré la Tour	Helen Wills Neuroscience Institute at UC Berkeley
Yaroslav Halchenko	Center for Open Neuroscience
Nicolas Hug	Columbia University
Adrin Jalali	Anaconda, Inc.
Thomas J Fan	Columbia University in the City of New York
Andreas Mueller	U. Columbia
Joel Nothman	U. of Sidney
Hanmin Qin	U. Pekin

7.3 Annex 3: Interview template

Guide for interviews on standards

Organisation

Name:

Type:

Interviewee:

Interviewee's position:

Standard focus

The standards of which domains are of interest to you?

1. Hardware
    a. Processor: Risc-V, OpenPower, ARM
    b. Interface between chips: CCIX, Gen-Z, CLX, OpenCapi, NVlink
    c. Interconnect: PCI, Infiniband
    d. Memory: JEDEC, HBM
    e. Storage: SCSI, SATA, SNIA
    f. System: Open Compute, Data Centre Alliance, ASHREA
2. Middleware
    a. OS: Linux foundation
    b. File system: Lustre, HDF, Spectrum Scale (new GPFS name), BeeGFS, CEPH
    c. Stack: OpenHPC, open container initiative
    d. Tools: Slurm, OpenStack
3. Language
    a. HPC: MPI, OpenMP
    b. Accelerator: OpenACC
    c. Interpreted: Julia
    d. Domain Specific Language (DSL)
4. Software
    a. HPDA: ScikitLearn, DMTK (Microsoft), Apache Spark,
    b. AI: Tensorflow (Google), Caffé, PyTorch, Microsoft Cognitive Toolkit (CNTK), Apache MXNet
    c. Dcache, Maestro, ESDM (Earth System Data Middleware) (scientific experiment software)

Current action

Are you involved in any standard organisations?

Are you in contact with anyone involved in standard definition and maintenance?

Yes

No

What are the reasons for your actions in relation to standards?

- Keeping track of a standard and their evolution
- Trying to push your innovation as a standard
- Other:

What are the obstacles to your standard-related activities?

Outlook

What are the domains where the emergence of standard is of interest to you?

What can help you better handle standard challenges?

Do you think pooling efforts with other projects would help?

Your dream about standards?

7.4 Annex 4: Examples of interviews

UCiT interviews on standards[33]

Organisation

Name: UCit

Type: an SME, software and service provider with three family of activities:

1/ developing software, allowing administrators to optimise their HPC infrastructure, collaboration with clients in new product development

2/service offering – helping clients to plan/manage the workload of their distributed HPC infrastructure (e.g. keeping it in the cloud or going for a private provider), helping mostly industrial clients (automotive) to embark on a cloud journey; tools to deploy real HPC clusters on e.g. AWS (with elasticity, nodes' communications, etc.); benchmarking of infrastructure using the AWS cloud

3/ reselling, distributing and installing NICE software (its visualization and portal) – as a way to contact new customers

Interviewee: Philippe Bricard and Benjamin Depardon

Interviewee's position: CEO and CTO (respectively)

Standard focus

The standards of which domains are of interest to you?

5. Hardware
    a. Processor: Risc-V, OpenPower, ARM
    b. Interface between chips: CCIX, Gen-Z, CLX, OpenCapi, NVlink
    c. Interconnect: PCI, Infiniband
    d. Memory: JEDEC, HBM

Comments in relation to Hardware Standards:

Objective: cluster design so that it best suits the applications, to get info from the HW on network/power consumption, also on CPU level – what is consumed; and there is no standards for that – for info format/storage, depending on the HW provider, you need specific tools, this is where standardisation would be important, whatever the HW is, to be able to analyse a cluster; we would like to add the power consumption/data movement, including a time-stamp, we are building a database on that, from the compute side, we could provide clients with the cost,

[33] Standard means technical norms or technical specifications that can be defined either by classical standard organisations, by consortium or accepted as 'de facto standard. "Standard organisations" are organisations that aim to establish standard through either consensus or market dominance.

energy would be more representative of an installation, even for big centres; we'd need something global, without using a lot of resources and without impacting what you are doing.

6. Middleware
    a. OS: Linux foundation
    b. File system: Lustre, HDF
    c. Stack: OpenHPC, open container initiative
    d. Tools: Slurm

Comments in relation to Middleware Standards:

There are two kinds of activities:

1. The deployment of a cluster - and we would like to see some kind of standard there, giving the possibility of a virtual cloud with bits of real HPC infrastructure, that could become a virtual computer on which a cloud HPC system would run (a virtualization of the cloud - it would give more flexibility to the user), a standard here could lead to containerization, etc., this is flexibility for the user to use various HPC infrastructures, there is no good adaption of common API (SGI and Slurm at a time but no strong adoption)
2. The use of such an infrastructure – the objective is to have the ability to use different clusters depending of the needs of the applications each vendor has different tools and on different levels, there is no standard/commonalities, there has been some efforts to standardize the description of jobs/schedulers, but there is no common tools now;

In Europe, there is enough diversity to provide enough flexibility; we could start working on some de-facto standard/pilot, could be an ETP4HPC Working Group, looking at some open source software, this could happen within PRACE as well – to propose something that is common to all systems, it would be a very operational approach to be sure to cover correctly the usage needs; conflict on one side using one of the public cloud vendor offering and so vendors development have to be followed; the way people share stuff is the more important aspect; an IT view of the world is to have commonalities, cloud vendors have their own roadmap, we need that intermediary level to speak to them, the user is an edge device in that paradigm.

7. Language
    a. HPC: MPI, OpenMP
    b. Accelerator: OpenACC
    c. Interpreted: Julia
    d. Domain Specific Language (DSL)

Only academically, not a priority at the moment,

8. Software
    a. HPDA: ScikitLearn, DMTK (Microsoft), Apache Spark,
    b. AI: Tensorflow (Google), Caffé, PyTorch, Microsoft Cognitive Toolkit (CNTK), Apache MXNet

It is on the radar, HPC is changing because it needs to be complimented with more and more statistics (ML a part of that), the value is between a real life use case/sensor and HPC is promising – this could be the fuel of the future HPC, if there was a good monitoring strategy (i.e. a standard) to identify some use cases to develop such standards, if the time could be shortened to identify such use cases. It will be important for parametric simulations (multi petaflops versus exaflops). We could start from use cases and then more aggressive approach and to act quickly. HPC can become more visible by the decision maker if this coupling is achieved. It would generate a new efficiency for company.

Current action

Are you involved in any standard organisations?

No, no resource available, but interested in joining the effort

Are you in contact with anyone involved in standard definition and maintenance?

Yes

No

What are the reasons for your actions in relation to standards?

- Keeping track of a standard and their evolution
- Trying to push your innovation as a standard
- Other:

The existence of standards could help our solutions to reach a broader market and bring more value to clients.

What is your current effort in this area?

What are the obstacles to your standard-related activities?

Outlook

What are the domains where the emergence of standard is of interest to you?

What can help you better handle standard challenges?

ETP4HPC would be a great proxy, organizing working groups, with simple mechanisms,

Do you think pooling efforts with other projects would help?

They could take part in a meeting but have no resources to continue such an effort

Your dream about standards?

See above the standard section