

Edition 16.01.2020

Eurolab4HPC Long-Term Vision on High-Performance Computing (2nd Edition)

Editors: Theo Ungerer, Paul Carpenter



Overall Editors and Authors

Prof. Dr. Theo Ungerer, University of Augsburg
Dr. Paul Carpenter, BSC, Barcelona

Authors

Sandro Bartolini	University of Siena	Photonics, Single Source Programming Models
Luca Benini	ETH Zürich	Die Stacking and 3D-Chips
Koen Bertels	Delft University of Technology	Quantum Computing
Spyros Blanas	The Ohio State University	Data Management
Uwe Brinkschulte	KIT	Memory Hierarchy
Paul Carpenter	BSC	Overall Organization and Diverse Sections
Giovanni De Micheli	EPFL	Nanowires, Superconducting Electronics, 2D electronics
Marc Duranton	CEA LIST DACLE	Overall Comments
Babak Falsafi	EPFL	Data Centers, Cloud Computing, Heterogeneous Systems
Dietmar Fey	University of Erlangen-Nuremberg	Memristors, Resistive Computing
Said Hamdioui	Delft University of Technology	Near- and In-Memory Computing, Resistive Computing
Christian Hochberger	Technical University of Darmstadt	Memory Hierarchy, Nanotubes
Avi Mendelson	Technion	Hardware Impact
Dominik Meyer	Helmut-Schmidt-Universität Hamburg	Reconfigurable Computing
Ilia Polian	University of Stuttgart	Security and Privacy
Ulrich Rückert	University of Bielefeld	Neuromorphic Computing
Xavier Salazar	BSC	Overall Organization and Related Initiatives
Werner Schindler	Bundesamt für Sicherheit in der Informationstechnik (BSI)	Security and Privacy
Theo Ungerer	University of Augsburg	Overall Organization and Diverse Sections

Compiled by

Florian Haas

University of Augsburg

We also acknowledge the numerous people that provided valuable feedback at the roadmapping workshops at HiPEAC CSW and HPC Summit, to HiPEAC and EXDCI for hosting the workshops and Xavier Salazar for the organizational support.

Radical changes in computing are foreseen for the next decade. The US IEEE society wants to “reboot computing” and the HiPEAC Visions of 2017 and 2019 see the time to “re-invent computing”, both by challenging its basic assumptions. This document presents the second edition of the “EuroLab4HPC Long-Term Vision on High-Performance Computing” of January 2020¹, a road mapping effort within the EC CSA² EuroLab4HPC that targets potential changes in hardware, software, and applications in High-Performance Computing (HPC).

The objective of the EuroLab4HPC Vision is to provide a long-term roadmap from 2023 to 2030 for High-Performance Computing (HPC). Because of the long-term perspective and its speculative nature, the authors started with an assessment of future computing technologies that could influence HPC hardware and software. The proposal on research topics is derived from the report and discussions within the road mapping expert group. We prefer the term “vision” over “roadmap”, firstly because timings are hard to predict given the long-term perspective, and secondly because EuroLab4HPC will have no direct control over the realization of its vision.

The Big Picture

High-performance computing (HPC) typically targets scientific and engineering simulations with numerical programs mostly based on floating-point computations. We expect the continued scaling of such scientific and engineering applications to continue well beyond Exascale computers.

However, three trends are changing the landscape for high-performance computing and supercomputers. The first trend is the emergence of data analytics complementing simulation in scientific discovery. While simulation still remains a major pillar for science, there are massive volumes of scientific data that

are now gathered by sensors augmenting data from simulation available for analysis. High-Performance Data Analysis (HPDA) will complement simulation in future HPC applications.

The second trend is the emergence of cloud computing and warehouse-scale computers (also known as data centres). Data centres consist of low-cost volume processing, networking and storage servers, aiming at cost-effective data manipulation at unprecedented scales. The scale at which they host and manipulate (e.g., personal, business) data has led to fundamental breakthroughs in data analytics.

There are a myriad of challenges facing massive data analytics including management of highly distributed data sources, and tracking of data provenance, data validation, mitigating sampling bias and heterogeneity, data format diversity and integrity, integration, security, privacy, sharing, visualization, and massively parallel and distributed algorithms for incremental and/or real-time analysis.

Large datacentres are fundamentally different from traditional supercomputers in their design, operation and software structures. Particularly, big data applications in data centres and cloud computing centres require different algorithms and differ significantly from traditional HPC applications such that they may not require the same computer structures.

With modern HPC platforms being increasingly built using volume servers (i.e. one server = one role), there are a number of features that are shared among warehouse-scale computers and modern HPC platforms, including dynamic resource allocation and management, high utilization, parallelization and acceleration, robustness and infrastructure costs. These shared concerns will serve as incentives for the convergence of the platforms.

There are, meanwhile, a number of ways that traditional HPC systems differ from modern warehouse-scale computers: efficient virtualization, adverse network topologies and fabrics in cloud platforms, low memory and storage bandwidth in volume servers.

¹<https://www.eurolab4hpc.eu/vision/>

²European Commission Community and Support Action

HPC customers must adapt to co-exist with cloud services; warehouse-scale computer operators must innovate technologies to support the workload and platform at the intersection of commercial and scientific computing.

It is unclear whether a convergence of HPC with big data applications will arise. Investigating hardware and software structures targeting such a convergence is of high research and commercial interest. However, some HPC applications will be executed more economically on data centres. Exascale and post-Exascale supercomputers could become a niche for HPC applications.

The third trend arises from Deep Neural Networks (DNN) for back propagation learning of complex patterns, which emerged as new technique penetrating different application areas. DNN learning requires high performance and is often run on high-performance supercomputers. GPU accelerators are seen as very effective for DNN computing by their enhancements, e.g. support for 16-bit floating-point and tensor processing units. It is widely assumed that it will be applied in future autonomous cars thus opening a very large market segment for embedded HPC. DNNs will also be applied in engineering simulations traditionally running on HPC supercomputers.

Embedded high-performance computing demands are upcoming needs. It may concern smartphones but also applications like autonomous driving, requiring on-board high-performance computers. In particular the trend from current advanced ADAS (automatic driving assistant systems) to piloted driving and to fully autonomous cars will increase on-board performance requirements and may even be coupled with high-performance supercomputers in the Cloud. The target is to develop systems that adapt more quickly to changing environments, opening the door to highly automated and autonomous transport, capable of eliminating human error in control, guidance and navigation and so leading to more safety. High-performance computing devices in cyber-physical systems will have to fulfil further non-functional requirements such as timeliness, (very) low energy consumption, security and safety. However, further applications will emerge that may be unknown today or that receive a much higher importance than expected today.

Power and thermal management is considered as highly important and will continue its preference in future. Post-Exascale computers will target more than

1 Exaflops with less than 30 MW power consumption requiring processors with a much better performance per Watt than available today. On the other side, embedded computing needs high performance with low energy consumption. The power target at the hardware level is widely the same, a high performance per Watt.

In addition to mastering the technical challenges, reducing the environmental impact of upcoming computing infrastructures is also an important matter. Reducing CO₂ emissions and overall power consumption should be pursued. A combination of hardware techniques, such as new processor cores, accelerators, memory and interconnect technologies, and software techniques for energy and power management will need to be cooperatively deployed in order to deliver energy-efficient solutions.

Because of the foreseeable end of CMOS scaling, new technologies are under development, such as, for example, Die Stacking and 3D Chip Technologies, Non-volatile Memory (NVM) Technologies, Photonics, Resistive Computing, Neuromorphic Computing, Quantum Computing, and Nanotubes. Since it is uncertain if/when some of the technologies will mature, it is hard to predict which ones will prevail.

The particular mix of technologies that achieve commercial success will strongly impact the hardware and software architectures of future HPC systems, in particular the processor logic itself, the (deeper) memory hierarchy, and new heterogeneous accelerators.

There is a clear trend towards more complex systems, which is expected to continue over the next decade. These developments will significantly increase software complexity, demanding more and more intelligence across the programming environment, including compiler, run-time and tool intelligence driven by appropriate programming models. Manual optimization of the data layout, placement, and caching will become uneconomic and time consuming, and will, in any case, soon exceed the abilities of the best human programmers.

If accurate results are not necessarily needed, another speedup could emerge from more efficient special execution units, based on analog, or even a mix between analog and digital technologies. Such developments would benefit from more advanced ways to reason about the permissible degree of inaccuracy in calculations at run time. Furthermore, new memory

technologies like memristors may allow on-chip integration, enabling tightly-coupled communication between the memory and the processing unit. With the help of memory computing algorithms, data could be pre-processed “in-” or “near-” memory.

But it is also possible that new hardware developments reduce software complexity. New materials could be used to run processors at much higher frequencies than currently possible, and with that, may even enable a significant increase in the performance of single-threaded programs.

Optical networks on die and Terahertz-based connections may eliminate the need for preserving locality since the access time to local storage may not be as significant in future as it is today. Such advancements will lead to storage-class memory, which features similar speed, addressability and cost as DRAM combined with the non-volatility of storage. In the context of HPC, such memory may reduce the cost of checkpointing or eliminate it entirely.

The adoption of neuromorphic, resistive and/or quantum computing as new accelerators may have a dramatic effect on the system software and programming models. It is currently unclear whether it will be sufficient to offload tasks, as on GPUs, or whether more dramatic changes will be needed. By 2030, disruptive technologies may have forced the introduction of new and currently unknown abstractions that are very different from today. Such new programming abstractions may include domain-specific languages that provide greater opportunities for automatic optimization. Automatic optimization requires advanced techniques in the compiler and runtime system. We also need ways to express non-functional properties of software in order to trade various metrics: performance vs. energy, or accuracy vs. cost, both of which may become more relevant with near threshold, approximate computing or accelerators.

Nevertheless, today’s abstractions will continue to evolve incrementally and will continue to be used well beyond 2030, since scientific codebases have very long lifetimes, on the order of decades.

Execution environments will increase in complexity requiring more intelligence, e.g., to manage, analyse and debug millions of parallel threads running on heterogeneous hardware with a diversity of accelerators, while dynamically adapting to failures and performance variability. Spotting anomalous behavior may be viewed as a big data problem, requiring techniques

from data mining, clustering and structure detection. This requires an evolution of the incumbent standards such as OpenMP to provide higher-level abstractions. An important question is whether and to what degree these fundamental abstractions may be impacted by disruptive technologies.

The Work Needed

As new technologies require major changes across the stack, a vertical funding approach is needed, from applications and software systems through to new hardware architectures and potentially down to the enabling technologies. We see HP Lab’s memory-driven computing architecture “The Machine” as an exemplary project that proposes a low-latency NVM (Non-Volatile Memory) based memory connected by photonics to processor cores. Projects could be based on multiple new technologies and similarly explore hardware and software structures and potential applications. Required research will be interdisciplinary. Stakeholders will come from academic and industrial research.

The Opportunity

The opportunity may be development of competitive new hardware/software technologies based on upcoming new technologies to advantageous position European industry for the future. Target areas could be High-Performance Computing and Embedded High-Performance devices. The drawback could be that the chosen base technology may not be prevailing but be replaced by a different technology. For this reason, efforts should be made to ensure that aspects of the developed hardware architectures, system architectures and software systems could also be applied to alternative prevailing technologies. For instance, several NVM technologies will bring up new memory devices that are several magnitudes faster than current Flash technology and the developed system structures may easily be adapted to the specific prevailing technologies, even if the project has chosen a different NVM technology as basis.

EC Funding Proposals

The Eurolab4HPC vision recommends the following funding opportunities for topics beyond Horizon 2020

(ICT):

- Convergence of HPC and HPDA:
 - Data Science, Cloud computing and HPC: Big Data meets HPC
 - Inter-operability and integration
 - Limitations of clouds for HPC
 - Edge Computing: local computation for processing near sensors
- Impact of new NVMs:
 - Memory hierarchies based on new NVMs
 - Near- and in-memory processing: pre- and post-processing in (non-volatile) memory
 - HPC system software based on new memory hierarchies
 - Impact on checkpointing and resiliency
- Programmability:
 - Hide new memory layers and HW accelerators from users by abstractions and intelligent programming environments
 - Monitoring of a trillion threads
 - Algorithm-based fault tolerance techniques within the application as well as moving fault detection burden to the library, e.g. fault-tolerant message-passing library
- Green ICT and Energy
 - Integration of cooling and electrical subsystem
 - Supercomputer as a whole system for Green ICT

As remarked above, projects should be interdisciplinary, from applications and software systems through hardware architectures and, where relevant, enabling hardware technologies.