JSC's Horizon 2020

Projects for Designing Future HPC Technologies

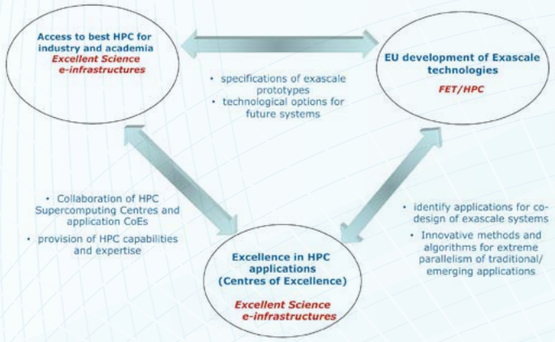

In 2012 the European Commission (EC) adopted a dedicated HPC strategy [1] which formulates as one objective the independent access to HPC technologies and systems for the EU (the pillars of the strategy implementation are shown in Fig. 1). Within the new program for research and innovation, Horizon 2020, the EC started funding 19 new projects. The call for these projects had been formulated on the basis of the first Strategic Research Agenda [2] of the European Technology Platform for HPC (ETP4HPC). JSC successfully joined the efforts of two consortia which aim for the development of future HPC core technologies and architectures, namely ExaNoDe (European Exascale Processor & Memory Node Design) and SAGE (Percipient StorAGe for Exascale data centric computing), which are coordinated by CEA and Seagate, respectively. The goal of the ExaNoDe project is to design a high-performance, heterogeneous compute element based on the chiplet concept and Unimem memory architecture previously explored in the EUROSERVER project [3]. This memory architecture aims for an elastic allocation of memory resources to different coherence islands by routing load/store operations between different chiplets. While ExaNoDe focuses on the design of future compute nodes, SAGE has the objective of providing a next-generation multi-tiered data storage that integrates computing capabilities. The project addresses two important exascale challenges: Today's disk-based storage architectures, which are highly cost-efficient for providing large storage capacity, will not be able to scale bandwidth as compute performance increases. Hierarchical storage architectures comprising high-bandwidth non-volatile memory devices will allow to mitigate this problem. The second challenge is the need for minimizing data movement as it is expensive in terms of energy consumption. SAGE's approach to this challenge is to integrate compute capabilities into the storage hierarchy, i.e. move data processing capabilities to where the data is. These new exascale projects are meant to be a first step within Horizon 2020 towards a European ecosystem for HPC capable of providing exascale class solutions. In a few years the results of these projects are expected to be integrated in extreme-scale demonstrator systems. This will be the litmus test for these development projects as they will have to proof that they can deliver technology which is ready for addressing large-scale computational challenges.

References

[1] http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2012:0045:FIN:EN:PDF

[2] http://www.etp4hpc.eu/strategy/strategic-research-agenda/

[3] http://dx.doi.org/10.1109/DSD.2014.15

• Dirk Pleiter

Jülich Supercomputing Centre (JSC), Germany

contact:
Dirk Pleiter,
d.pleiter@fz-juelich.de

Future Supercomputers for Brain Research: Pre-Commercial Procurement entered final Phase

Human Brain Project

By developing and expanding the use of information technology, the Human Brain Project (HBP) [1] wants to open new opportunities for brain research. The goal of this European project is to enable a multi-level, integrated understanding of brain structure and function. Particular challenging is the enablement of large-scale simulations of brain models as todays HPC architectures do not meet their requirements. This includes both, the need for extremely large memory footprint and interactive supercomputing. For realistic network sizes, e.g., the amount of data generated during a simulation becomes too large for being written to an external storage system and thus new memory technologies have to be integrated. Furthermore, the complexity of the simulations requires interactive steering. To ensure that suitable solutions for realizing HBP's future High-Performance Analytics and Computing Platform will exist, the project published in April 2014

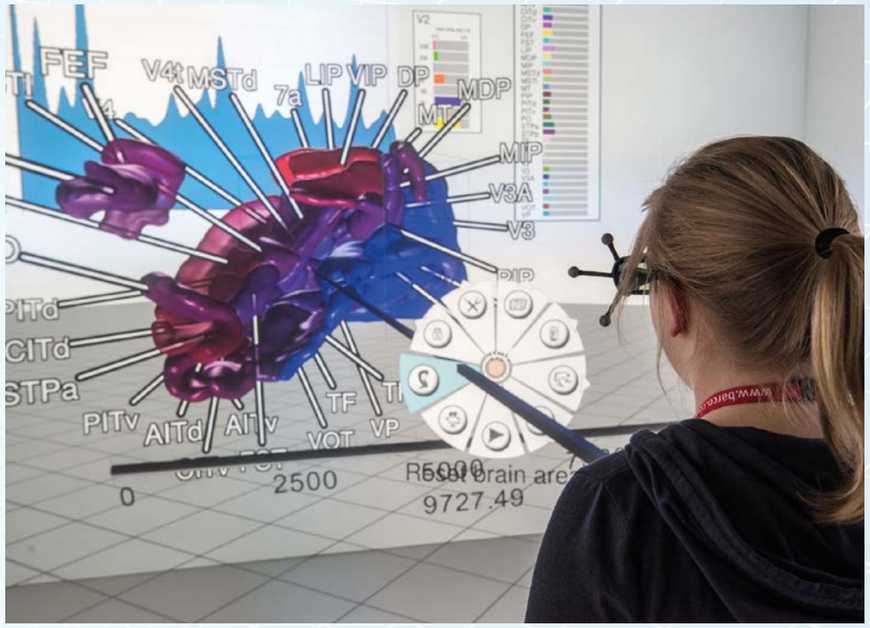

Figure 1: NEST [5], an interactive analysis tool for neural activity data, being used in the aixCAVE at RWTH Aachen University (Source: Virtual Reality Group, RWTH Aachen University).

a tender for a pre-commercial procurement (PCP). PCP [2] is an instrument promoted by the European Commission (EC) to foster innovation through public procurement. It allows for procuring research and development services to enable development of new solutions which would otherwise likely not be available. A comparison with the already announced pre-exascale systems in the US, like Summit at ORNL [3] and Aurora at ANL [4], confirm additional research and development efforts will be required to realize the planned HBP Platform. This concerns in particular the integration of dense memory technologies, scalable visualization (see Fig. 1 for a visualization use case) as well as dynamic management of resources required for interactive access to the systems. By design a PCP is organized as a multi-phase, competitive process. During Phase I the suppliers had the task for sketching-out a design meeting the different challenges, which then had been refined in Phase II. At the end of that phase, three competitors, namely Cray, a consortium comprising Dell and the German SMEs Extoll and ParTec as well as a consortium consisting of IBM and NVIDIA, presented their design specifications. An expert committee evaluated these solutions in July 2015 and recommended awarding Cray and IBM-NVIDIA with contract for Phase III. These contractors now have the task of implementing the proposed solutions and demonstrate their technological readiness on pilot systems that will be installed in 2016. A PCP is a still a new instrument which needs to be carefully designed to balance goals, timing and available budget. But the efforts within the HBP demonstrate that PCP can be a suitable instrument to drive development of future supercomputers.

References

[1] https://www.humanbrainproject.eu

[2] http://ec.europa.eu/digital-agenda/en/pre-commercial-procurement

[3] https://www.olcf.ornl.gov/summit

[4] http://aurora.alcf.anl.gov

[5] http://dx.doi.org/10.1109/BioVis.2013.6664348

contact:
Dirk Pleiter,
d.pleiter@fz-juelich.de

• Dirk Pleiter

Jülich Supercomputing Centre (JSC), Germany

The Vector Computer NEC SX-ACE

In the eighties of the last century High Performance Computing was a synonym for the usage of so-called vector supercomputers, machines that used pipelining in the functional units as well as in data path for the acceleration of numerical codes. Memory bandwidth was high in comparison to the peak floating-point performance. These machines from Cray (later split in Cray Research, Cray Computer Corporation, Super Computer Systems Incorporated), CDC, IBM, Fujitsu, Hitachi and NEC had low chip integration density, but a very sophisticated packaging, which made these technologies expensive. Shared memory parallelism was early integrated and led to OpenMP as standardized parallel model. In the nineties large scale parallel distributed memory computers based on relative inexpensive integrated processors gained traction, the so called “killer micros”, together with message passing as parallel model. Indeed these machines replaced the vector supercomputers, but not as fast as expected. One reason for surviving was that vector machines delivered predictable reliable performance for vectorized codes. Vectorization as parallel paradigm can be handled easyer automatically by compilers in contrast to other kinds of parallelization. The other reason was that they came up as parallel machines (Earth Simulator in 2002, HLRS in 2005). Nevertheless, vector machines seem not to play a major/any role today. But is that true? A closer look to modern processors shows that all of these have SIMD support and are to be used as vector machines. Suppressing vectorization in the code by inappropriate programming (many procedure calls, usage of array of structures instead of structure of arrays, recursive loops, short