Pre-Commercial Procurement of the Human Brain Project Completed

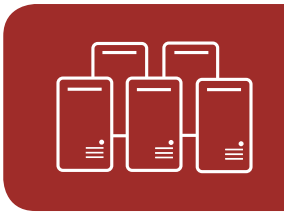

One of the key goals of the Human Brain Project (HBP) [hbp] is the creation of a research infrastructure for neuroscience, which also comprises high-performance computing (HPC) and high-performance data analytics (HPDA) systems. This area of research comes with new and challenging requirements, and thus the HBP also needs to work on enabling new technologies and architectures.

In 2014, the HBP therefore launched a pre-commercial procurement (PCP) of research and development services, which was successfully completed in January 2017. During a wrap-up workshop at JSC in March, the contractors of the last phase of the PCP, Cray and a consortium of IBM and NVIDIA, and researchers from the HBP discussed the resulting solutions and their evaluation.

One of the technical goals of this PCP was to facilitate new approaches to the integration of dense memory technologies. New solutions should enable global access to distributed dense memory. IBM expanded previous work on a Distributed Storage Access (DSA) layer to enable access to distributed dense memory within a global address space through an RDMA-type interface. On a pilot system, which was deployed at JSC in late summer 2016, it could be demonstrated that this interface allows to achieve both, extremely high bandwidth as well as a high IOPS rate. The competitor Cray proposed a solution based on a commodity technology, namely Ceph. So far, this technology has hardly been used for HPC, and its potential for software-defined storage architectures based on dense memory has not yet been exploited.

For many of the use cases of the HBP, visualization plays a key role and was thus chosen as a second focus of the PCP. The goal was to enable scalable visualization capabilities that are tightly integrated into large-scale HPC and HPDA systems. A big challenge was with the NVIDIA team to enable the complex visualization software stacks of the HBP on GPU-accelerated POWER servers. The new OpenPOWER HPC architectures do rely heavily on the compute performance provided by GPUs and thus are particular suitable for large-scale visualization.

The third technical focus area was dynamic resource management. The HBP sees this as a requirement to improve the utilization of future HPC architectures and to support the HBP's complex workflows. In this context, new features allowing to change the resources available to a running job have been added to resource managers, and mechanisms for resizing job sizes have been designed and implemented.

The PCP is a quite new instrument in Europe for working with commercial operators and promoting the development of innovative solutions. Within the HBP, its use could be successfully demonstrated for enabling new HPC capabilities. The solutions developed within the PCP will become available to a broader

The two pilot systems, developed by Cray (left) and IBM-NVIDIA (right), are designed to meet the specific requirements of neuroscience applications. Both systems are connected to an high-performance storage system from IBM/Lenovo (middle).

community and can already be exploited by HBP scientists through the two pilot systems, JURON and JULIA, which have been deployed at JSC as part of this project.

References

[hbp] Human Brain Project:

https://www.humanbrainproject.eu

[hbp-infra] B. Orth et al.:

"Towards an European infrastructure for brain research." Inside, Autumn 2016.

[dsa] B. Metzler et al.:

"Direct Storage-class Memory Access". https://www.openfabrics.org/images/eventpresos/workshops2014/DevWorkshop/presos/Monday/pdf/09_DirectStorage.pdf

Written by Dirk Pleiter and Boris Orth
(Jülich Supercomputing Centre (JSC), Germany)

Contact: Dirk Pleiter, d.pleiter@fz-juelich.de, Boris Orth, b.orth@fz-juelich.de