NEXTGenIO Project: Addressing HPC and Data Analytics I/O Challenges

The NEXTGenIO project[1] develops an innovative HW/SW architecture that fully leverages the Intel® 3D XPoint™ technology to provide unmatched I/O performance, scalability and energy efficiency, and supports complex HPC and high performance data analytics (HPDA) workloads and workflows. The project has started October 1st, 2015, and did establish a virtuous co-design circle involving key European HPC operators, end-users and technology providers, preparing the NEXTGenIO system architecture design phase in the first half of 2016. An integrated system prototype will be built by Fujitsu, and a thorough testing and evaluation process will ensure its functionality and performance. The prototype will consist of Intel® Xeon® CPUs, include the announced NVDIMM variant of 3D XPoint storage-class memory, and be complemented with a full systemware stack including Intel® Lustre*-based parallel I/O, a data object store, and an advanced scheduling/orchestration engine that accommodates data location.

Today's practice in high-end HPC systems is to employ data storage separate from the main computer system. These I/O subsystems often struggle to keep up with the ever increasing compute performance and in particular the sharply rising degree of parallelism, given that they need to bridge the gap to comparatively slow disk storage. The result is poor I/O performance, which impacts application runtimes and thus science throughput. As we move into the domain of extreme parallelism at the Exascale, and as large-scale data analytics emerge, we need to address the I/O bottleneck challenge if such systems are to deliver acceptable performance and efficiency for their application user communities.

The 3D XPoint memory technology developed by Intel and Micron will close the performance gap between DDR RAM and storage class memory, and in addition a NVDIMM version will become available that provides fine-granular access to large amounts of persistent memory, enabling completely new use cases. Figure 1 (left) illustrates the basic organization of 3D XPoint, with memory cells being addressed by row and column signal lines. Figure 1 (right) shows the place of this new technology in a memory and storage hierarchy. Availability of first 3D XPoint products is tentatively announced for late 2016/early 2017.

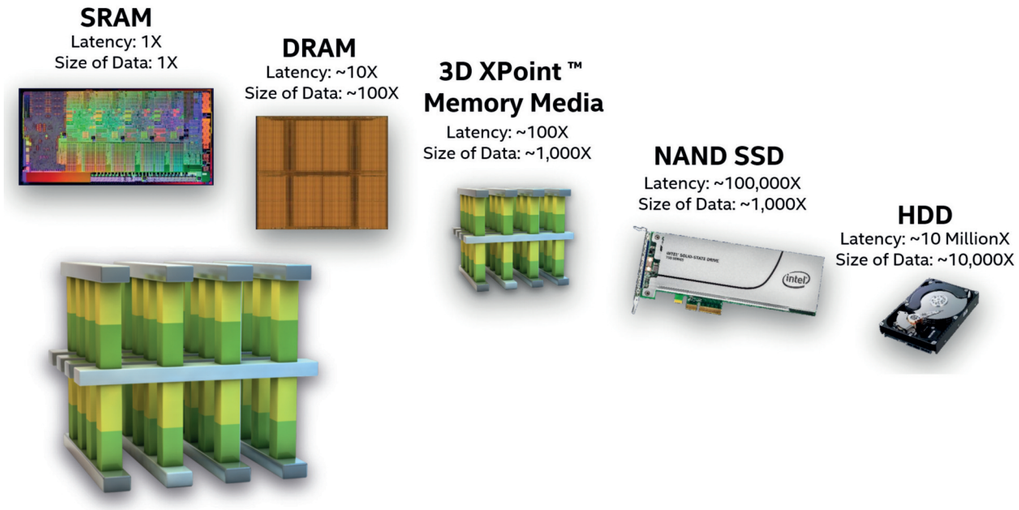

Figure 1: the lower left part illustrates the basic organization of Intel 3D XPoint memory devices, with a combination of row and column signal lines addressing each cell. The upper part shows the place to the new memory technology in an extended memory hierarchy.

The NEXTGenIO project is funded by the European Commission under grant agreement no. 671591.

The NEXTGenIO project revolves around exploring the use of 3D XPoint technology and associated systemware developments through a co-design process with three end-user partners: Edinburgh Parallel Computing Centre (EPCC, also coordinating the project) as a high-end academic HPC service provider, ECMWF as a globally leading numerical weather centre and Arctur as a commercial on-demand HPC service provider mainly for industrial and commercial users. A key output of NEXTGenIO will be a prototype system built by Intel and Fujitsu, based on the new technologies.

NEXTGenIO will also develop an I/O workload simulator to allow improvements in performance to be directly measured on the new system for a wide range of research configurations. Systemware developed in the project will include performance analysis tools, advanced schedulers that take into account data locality and energy efficiency, optimized programming models, and APIs and drivers for optimal use of the new I/O hierarchy.

In the year of 2015, the end-users in NEXTGenIO have started the co-design cycle by first defining a set of abstract usage scenarios to be supported by the new architecture. From these, specific use cases were derived, and in turn specific qualitative and quantitative requirements were constructed. Each requirement was labeled according to its criticality, and in the next step, the feasibility of meeting each requirement was established in a dialog with the system architecture experts. This annotated list then defines the usage scenarios and use cases that can be achieved in the project, and feedback from the system architects was used to improve the specificity of requirements. Figure 2 illustrates the co-design cycle.

In 2016, a first detailed system architecture will be defined, and design for the prototype HW and SW systems will progress. In addition, the test and evaluation methodology will be defined, including a precise list of workloads and input data sets, the metrics to be recorded and the tools and procedures used to do so, and baseline measurements with today's systems at EPCC and ECMWF will be established.

Full information on NEXTGenIO can be found at the project website: www.nextgenio.eu

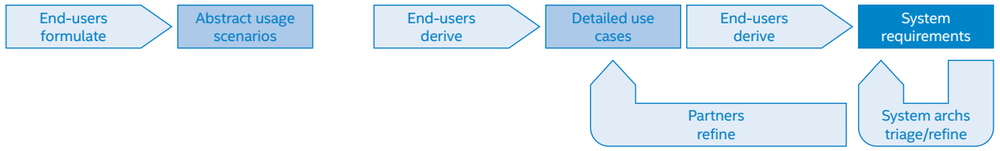

Figure 2: NEXTGenIO has started the co-design process in 2015 by a definition of usage scenarios and use cases and the systematic collection and evaluation of requirements. The results will direct the work on HW and SW architecture definition and design in 2016.