HPI Hardware Update – January 2016

Markus Dreseler,
markus.dreseler@hpi.de

Summary

- First Intel NVDIMMs will have a capacity of 512 GB per DIMM, Microsoft working on NVM-aware file systems
- Omni-Path fabric launched
- Google shows 100M-fold speedup with a quantum computer on selected benchmarks
- Google, HPE, and Oracle work on open-source RISC-V processor
- Intel finishes acquisition of FPGA producer Altera
- Samsung produces 128GB DIMMs
- AMD ships first ARM-based server CPU

Non-Volatile Memory

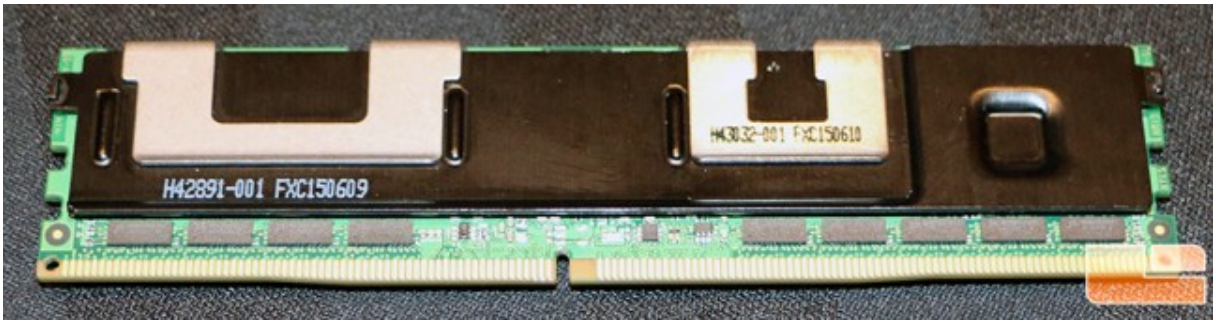

Figure 1: Mechanical sample of a 3D XPoint DIMM [NV1]

Few more details about the upcoming 3D XPoint memory have been made public. With a maximum capacity of 512 GB per DIMM (compare to 128GB for traditional memory, see "Newsflash") a dual-socket server could use up to 6 TB [NV2]. While first prototypes are "just around the corner", mass production could take another 12 to 18 months. The chips use around 100 new materials, which raises supply chain questions, and require more complicated production steps [NV3]. 3D XPoint is expected to have a total addressable market of $34B in 2020 [NV4].

Furthermore, Microsoft announced that they have an internal Windows Server 2016 version with a file system optimized for NVM. With memory-mapped files, this allows for unbuffered, zero-copy file accesses [NV5].

For as long as actual NVM is not available, Micron (also working on 3D XPoint) announced a new battery-backed DDR4 DIMM that uses supercapacitors and NAND flash to backup the data in case of a power failure. This means that the capacity is limited to that of DRAM, the hardware cost is higher than traditional DRAM (1.5x-2x), but also that the higher latencies of NVM are not coming into play [NV6].

A different approach is used by Netlist, which just announced a $22M collaboration with Samsung. Their Hypervault product combines DRAM and NAND flash, but instead of using NAND only as a place to write data to in the case of a power loss, they also use it to extend the memory's capacity. By using DRAM as a buffer for the NAND memory, they claim that it "allows non-volatile memory (NAND) to compete with DDR4 NVDIMM products and other future non-volatile memories, such as 3D XPoint" [NV7].

Intel Omni-Path / Scalable System Framework

In November 2015, Intel formally launched the Omni-Path Architecture (OPA) fabric. Hardware is not yet available. Advantages of OPA over existing fabrics include a higher port density in the switches, a peak port bandwidth of 100 Gbps, and a 17% lower latency than Infiniband [OP1]. While PCIe adapters will be sold, the future of OPA is supposed to be in the on-chip integration with the Xeon Phi (2016) and the Xeon processors (2017) [OP2]. Omni-Path is supported by a new ecosystem that, besides Intel, includes both hardware and software vendors.

OPA is path of Intel's Scalable System Framework (SSF), which they present as an overall design foundation for both data- and compute-intensive workloads. Charles Wuischpard, GM of the HPC group at Intel, explained: "[Y]ou really have to take more of a systemic view and look at memory, I/O, storage, and the software stack" [OP3].

Mellanox, producer of Infiniband hardware, claims that Omni-Path has a significantly higher CPU overhead than Infiniband [OP4].

SSF appears to be the roof to a number of hardware technologies (such as Omni-Path and Xeon Phi), software technologies (Lustre file system, cluster management), and guidelines (such as reference architectures and hardware configurations). One of the main goals is to reduce the complexity and cost of adapting HPC for both small companies and large supercomputers [OP5].

Quantum Computing

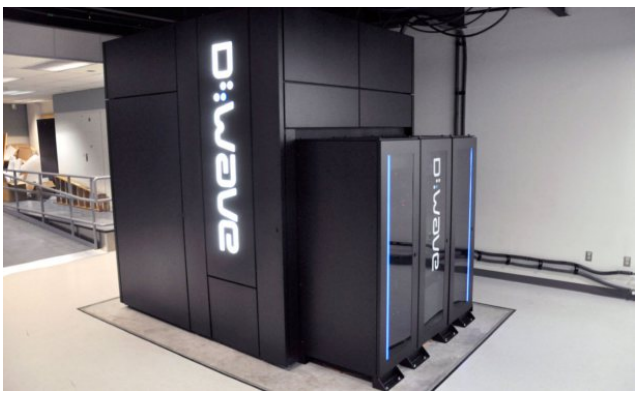

Figure 2: The D-Wave 2X

Google announced that their quantum computer, a D-Wave 2X produced by a Canadian company, outperformed traditional computers by a factor of 100M for carefully selected problems. One of the benchmarks used was an optimization of a system of equations with 1000 binary

variables [QC1]. While this is a first step, Google admits that more work is needed for "problems of practical relevance".

The results are disputed as researchers believe that the code used for the conventional machine was not optimal [QC2]. Additionally, the used hardware is a so-called adiabatic quantum computer, not a general-purpose one. These computers are limited to a small set of optimization problems [QC3], which are still of interest to Google as there are potential applications in pattern recognition and machine learning [QC4].

Google itself is also working on its own quantum computer [QC4]. Additionally, Microsoft and Rambus announced a collaboration looking at new "high-bandwidth, power-efficient memory architectures" for future quantum computers [QC5].

Newsflash

- Google, HPE, and Oracle, and others founded a trade group for the development of the open-source RISC-V processor architecture [NF1]. At the moment, it is mainly used in academics and only a single commercial consumer product exists.
- Intel finished its $16.7B acquisition of Altera, producer of field-programmable gate arrays (FPGAs) [NF2]. The new division within Intel will work closely with the data center (DCG) and IoT groups. Part of the strategy is to have one third of cloud nodes use FPGAs by 2020, accelerating tasks such as compression for big data or encryption [NF3]. For this, FPGAs will work closely with the CPUs. The first Xeon CPUs copackaged with FPGAs are to be released in 2016, allegedly boosting processor speed by 30%-50%. In the future, they are to be integrated with processors on the same die [NF3]. A major competitor is Xilinix, the biggest producer of FPGAs, which is currently more advanced in the development of integrated CPUs [NF4].
- Samsung has started mass-producing 128GB DDR4 RDIMMs, allowing a 96-slot server to hold 12.2 TB [NF5]. They claim a throughput of 2400 Mb/s. This number appears odd, as other DDR4 modules have a peak transfer rate starting at 12.8 GB/s. Most likely, the press release is incorrect and refers to 2400 MT/s (Million Transfers / second), resulting in a peak transfer rate of 19.2 GB/s. We have not yet received an answer from Samsung. To achieve the capacity and bandwidth, a technique called through-silicon via (TSV) is used in which multiple wafers are stacked and connected with a vertical connection. This allows Samsung to place 144 1GB DRAM chips on a single module [NF6].
- AMD has announced an ARM-based server processor. The A1100 series uses the reduced ARM instruction set (RISC) to allow for an inexpensive, energy-efficient processor. With eight cores and a core frequency of up to 2 GHz, even the fastest version does not

compete with the Intel-dominated x86 architecture. Instead, AMD focuses the network- and storage-heavy scale-out market, storage, and web-facing applications [NF7, NF8].

References

[NV1] http://www.legitreviews.com/intel-shows-off-512gb-optane-drive-with-3d-xpoint-memory-that-fits-in-ddr4-slot_176826
[NV2] http://www.tweaktown.com/news/49408/intels-3d-point-technology-enables-up-6tb-system-memory/index.html
[NV3] http://hexus.net/tech/news/industry/89780-3d-xpoint-memory-chip-samples-just-around-corner/
[NV4] http://files.shareholder.com/downloads/INTC/1048463185x0x862746/B63C5999-58CD-411E-A2C6-F9D9DF343CAA/2015_InvestorMeeting_Diane_Bryant_WEB.pdf
[NV5] http://www.tomshardware.com/news/intel-3d-xpoint-picture-nvdimm,30890.html
[NV6] http://www.nextplatform.com/2015/12/07/nvdimm-cant-wait-for-3d-xpoint-cant-rely-on-dram-alone/
[NV7] http://www.tomsitpro.com/articles/samsung-netlist-3d-xpoint-nvdimm,1-3048.html

[OP1] http://www.hpcwire.com/2015/12/14/dawn-of-a-new-era-in-high-performance-computing/
[OP2] http://insidehpc.com/2015/11/video-intel-debuts-omni-path-at-sc15-2/
[OP3] http://www.nextplatform.com/2015/11/16/intel-rounds-out-scalable-systems-with-omni-path/
[OP4] http://insidehpc.com/2016/01/infiniband-enables-intelligent-networks/
[OP5] http://www.hpcwire.com/2015/12/07/intels-next-generation-of-high-performance-computing-architecture/

[QC1] http://googleresearch.blogspot.co.uk/2015/12/when-can-quantum-annealing-win.html
[QC2] http://www.technologyreview.com/news/544276/google-says-it-has-proved-its-controversial-quantum-computer-really-works/
[QC3] http://www.eetimes.com/document.asp?doc_id=1326592
[QC4] http://www.technologyreview.com/news/544421/googles-quantum-dream-machine/
[QC5] http://www.hpcwire.com/2016/01/06/microsoft-rambus-collaborate-on-quantum-computing/

[NF1] http://hexus.net/tech/news/cpu/89195-google-hpe-oracle-back-risc-v-open-source-arm-alternative/
[NF2] http://newsroom.intel.com/community/intel_newsroom/blog/2015/12/28/intel-completes-acquisition-of-altera

[NF3] http://marketrealist.com/2015/12/intels-strategy-integrate-alteras-fpga-technology/
[NF4] http://www.fool.com/investing/general/2016/01/04/what-will-intel-do-now-that-it-owns-altera.aspx
[NF5] http://www.theregister.co.uk/2015/11/26/time_for_a_new_ram_cram_plan_as_128gb_ddr4_dimms_land/
[NF6] http://thememoryguy.com/samsungs-colossal-128gb-dimm/
[NF7] http://hexus.net/tech/news/cpu/89741-amd-announces-arm-based-opteron-a1100-processor/
[NF8] http://www.theregister.co.uk/2016/01/14/amd_arm_seattle_launch/