INTEL NEWS FACT SHEET

Intel Unveils Next Generation Technologies to Advance High Performance Computing and Artificial Intelligence Capabilities

June 19, 2017 – The accelerating growth of the world's data combined with the diversity of new digital services and business models is demanding new ways to extract insight from information – leading to a blurring of traditional High Performance Computing (HPC) with data analytics and emerging artificial intelligence models. Commercial, government and academic users all face similar dilemmas: how to scale computing resources in this new paradigm, how to democratize tools and access, how to simplify and future proof solution decisions and how to fit within resource constraints. Intel continues to advance state-of-the-art technologies for enabling complex workloads to scale and generate more answers to the world's most challenging questions, quickly and efficiently.

Delivering Performance and Scalability to HPC Systems with the Next Generation Intel® Xeon® Scalable Processor

With the shift in market dynamics and rise of diverse workloads, there is a need to increase performance and scalability for HPC systems. Intel is addressing this demand with the next generation Intel® Xeon® Scalable Processor that delivers impressive performance to accelerate HPC applications. This new processor represents a significant leap forward in performance and efficiency that unlocks scalable performance for a broad range of HPC systems from the smallest clusters to the world's largest supercomputers. The new processor offers a boost to performance for demanding workloads by incorporating Intel® Advanced Vector Extensions 512 (Intel® AVX-512). With Intel AVX-512, the Intel Xeon Scalable Processor can deliver up to twice the amount of flops per clock-cycle peak compared to the previous generation[1], boosting performance for demanding computational workloads in HPC applications such as modeling and simulation, data analytics, machine learning, and visualization. Intel AVX-512, combined with improvements in cores, cache and memory, delivers up to 2.27x more performance than today's Intel Xeon processor E5 (formerly codenamed Broadwell)[2], and an increase of up to 8.2x more double precision GFLOPS/sec when compared to a four-year old Intel Xeon processor E5 family in the installed base (formerly codenamed Sandy Bridge)[3]. Intel is already seeing use in some early Top500 systems, three were included in the June 2017 Top500 list[4].

The next generation Intel Xeon Scalable Processor also integrates the Intel® Omni-Path Architecture (Intel® OPA), an end-to-end fabric solution designed and optimized for HPC, offering high MPI message rate, low latency and high bandwidth performance. The combination of the new Intel Xeon Scalable Processor with Intel OPA will enable supercomputers to achieve greater performance and cost-effective scaling to extremely large node counts. When deployed with the new Intel Xeon platform, Intel OPA offers the added benefit of a fabric that is integrated with the processor, which improves density and helps optimize platform costs.

INTEL NEWS FACT SHEET

Intel® Xeon® Scalable Processors - Line up and Availability

- Intel® Xeon® Platinum Processor
- Intel® Xeon® Gold Processor
- Intel® Xeon® Silver Processor
- Intel® Xeon® Bronze Processor

The Intel Xeon Scalable Processor will be generally available in mid-2017.

Accelerating Processing Performance for AI Solutions

To further fuel innovation for AI and deep learning solutions across a wide range of industries, Intel will offer a new extension to the Intel® Xeon Phi™ processor family with its upcoming processor code named Knights Mill (KNM).

Knights Mill is Intel's first CPU that is optimized to enable new levels of deep learning systems while reducing time to train deep learning models. By leveraging low-precision computing, KNM delivers faster computing for deep learning workloads, enabling data scientists to train deep learning models in a matter of hours as opposed to days or weeks. To achieve these performance increases, KNM is introducing new instruction sets that improve low-precision computing called Quad Fused Multiply Add (QFMA) and Quad Virtual Neural Network Instruction (QVNNI). The QFMA can up to double the amount of single precision performance KNM can deliver over that of the current Intel Xeon Phi processor and QVNNI, which reduces the precision further, bringing additional performance gains for deep learning workloads. With both QFMA and QVNNI, Knights Mill can deliver significantly higher training performance than today's Intel Xeon Phi processor, making it a targeted solution for deep learning workloads.

The new Knights Mill is expected for production in the fourth quarter of 2017.

Increasing Efficiency with Turnkey Software Defined Visualization Solution

Software defined visualization is becoming more prevalent in high performance computing and with new advancements, HPC systems are moving toward a more integrated, software-defined infrastructure. To optimize visualization for HPC, Intel is delivering an integrated turnkey solution with its SDVis Appliance (Software Defined Visualization Appliance). Based on an open standards foundation, the SDVis Appliance allows HPC users to gain benefits of a developed software ecosystem without having to construct the hardware system themselves or work with vendors to customize the configuration. As a visualization-ready solution, the SDVis Appliance comes bundled with necessary software for visualization, including SDVis Software (ParaView, VTK, VisIt, VMD), and Intel® Parallel Studio Cluster Edition. In addition, through its high performance capabilities, the SDVis Appliance offers a pre-configured solution for in-situ, post-processing, and professional rendering visualization tasks that can support much larger data sets than competing GPU-based solutions.

INTEL NEWS FACT SHEET

Intel Continues Top500 Leadership

Intel made significant advancements in the Top500 list of supercomputers including three systems based on the upcoming Intel Xeon Scalable Processor, of which the highest ranked is the No.13 system from Barcelona Supercomputing Center. Overall, 464 of the Top500 systems are based on Intel processors, including 13 systems with Intel Xeon Phi processors. Intel OPA continues to be the leading 100Gb/s fabric in the Top500 used in 38 systems; a 36 percent increase over the November 2016 Top500 list.

Intel Demonstrates a Variety of Cutting-Edge Solutions for HPC at ISC 17

- **Attala System Intel® FPGA-based solution** - Intel field programmable gate arrays (FPGAs) support an emerging storage standard known as NVMe over RoCE (Non-Volatile Memory Express over RDMA over Converged Ethernet). The Attala System Intel® FPGA-based solution enables efficient CPU access to pooled storage, and eases host configuration, while using the FPGA to offload functions. The solution also enables inline acceleration and HPC system future-proofing.
- **High Content Screening for Drug Discovery** - With Intel technology, artificial intelligence helps businesses accelerate solutions, automate operations, and gather better insights to make smarter decisions. One example is high-content imaging application that harnesses AI to improve treatments and leads to better understanding in life sciences. Intel will showcase this capability by demonstrating an Intel Xeon Phi processor-based solution including Intel's latest AI software optimizations for Caffe which can be used to significantly reduce the time to train an AI model.

Intel will also show demonstrations of our Intel Xeon Phi Developer Access Program and our latest Software Visualization appliance.

1. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Results are based on internal testing and are provided to you for informational purposes. Any differences in your system hardware, software, or configuration may affect your actual performance.
2. Baseline config: 1-Node, 2 x Intel® Xeon® Processor E5-2699 v4 on Red Hat Enterprise Linux* 7.0 kernel 3.10.0-123 using Intel® Distribution for LINPACK Benchmark, score: 1446.4 GFLOPS/s vs. estimates based on Intel internal testing on 1-Node, 2x Intel® Xeon® Scalable processor (codename Skylake-SP) system. Score: 3295.57
3. Baseline configuration: 1-Node, 2 x Intel® Xeon® Processor E5-2690 on Intel® Server Board S2600CP2 with 32 GB Total Memory on Red Hat Enterprise Linux* 6.0 (Santiago) kernel version 2.6.32-504.el6.x86_64 using Intel® Distribution for LINPACK Benchmark using 56000 problem size. Score: 366.0 GFLOPS/s vs. estimates based on Intel internal testing on 1-Node, 2x Intel® Xeon® Scalable processor family (codename Skylake-SP) system.
4. https://www.top500.org/lists/2017/06/