Intel flirts with exascale leap in supercomputing

June 19 2012, by Nancy Owano

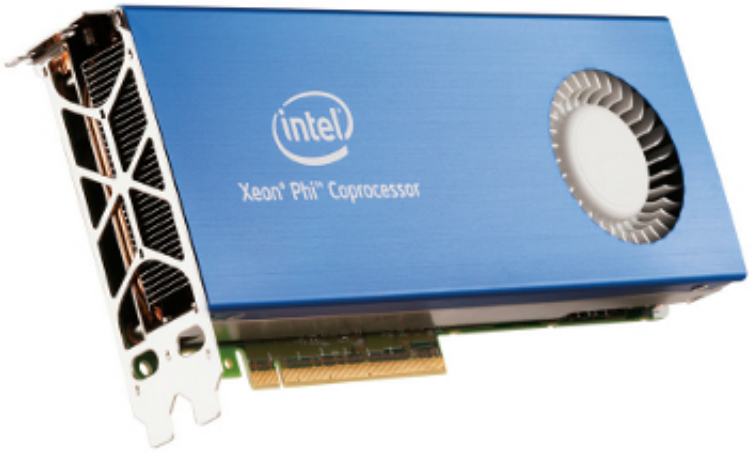

(Phys.org) -- If exascale range is the next destination post in high-performance computing then Intel has a safe ticket to ride. Intel says its new Xeon Phi line of chips is an early stepping stone toward exascale. Intel on Monday announced its high performance chip family as Xeon Phi at the International Supercomputing conference in Hamburg, Germany. Xeon Phi is now the brand name for future Intel's Many Integrated Core (MIC) architecture-based products. The announcement

was obviously made at the right venue. The ISC is a key gathering for high-performance computing, networking and storage experts.

Intel chips are used in the majority of the 500 fastest supercomputers. The Xeon processor is found in 70 percent of these top 500.

The Xeon Phi is intended as a Xeon complement, optimized for highly parallel supercomputing. Xeon Phi works as a co-processor alongside a server CPU to accelerate workloads. The coprocessor will be compatible with x86 programming models, running as an HPC-optimized, highly parallel, separate compute node with its own Linux-based operating system, independent of the operating system of the host. The Phi technology is to enable both high performance and energy efficiency when processing highly parallel applications.

The first Xeon Phi chip, codenamed Knights Corner, will be out at the end of this year. Knights Corner will have more than 50 cores and will deliver four or five gigaflops per watt. It will use 22 nanometer silicon and Intel's 3-D TriGate Transistors. The architecture is ready for cluster-based experimental deployments.

To reach exascale requires 40 to 50 gigaflops of performance-per-watt. The first Phi chip would deliver about four to five gigaflops per watt, according to Intel's John Hengeveld, director of marketing for high-performance computing.

Intel is working with other technology companies to support the new chip family. Intel has commitments from a number of computer partners, say reports, to make use of the Xeon Phi in their roadmaps. In its press announcement, Intel noted that its acquisition of Infiniband and interconnect assets from QLogic and Cray further present areas for Intel to innovate in delivering future scalable exascale-class platforms.

Supercomputers can harness the parallel-processing capabilities of graphics processors and chips like Phi to carry out complex calculations in scientific and math research. The task of building specialized chips that can execute more calculations per second while keeping power consumption in check, though, is not trivial. Intel is targeting no sooner than 2018 as the year it reaches “exascale performance.”

More information:
www.intel.com/content/www/us/e … ver-reliability.html
www.intel.com/newsroom/kits/is … 012-Presentation.pdf