News Fact Sheet

Intel® Xeon Phi™ Coprocessor: Pushing the Limits of Discovery

Nov. 12, 2012 – Increasingly, society's most pressing challenges require the greatest computational complexity. From climate modeling to pharmaceutical research to national security, High Performance Computing (HPC), driven by parallelism, is helping the industry push the boundaries of competitiveness and discovery.

Today, Intel Corporation is helping push discovery forward by announcing the Intel® Xeon Phi™ coprocessor, a product that will help deliver exponential levels of performance for highly parallel workloads. Based on the Intel® Many Integrated Core (Intel® MIC) architecture, the Intel Xeon Phi coprocessor complements the existing Intel® Xeon® processor E5-2600/4600 product families to deliver unprecedented performance and efficiency for highly parallel applications that will enable breakthroughs in manufacturing, life sciences, energy and other areas. Intel Xeon Phi coprocessors also signal Intel's commitment to reaching Exascale computing (thousandfold increase over Petascale) by 2018.

Shared Programming Models Across All Your Code

A wide assortment of programming languages, models and tools support Intel® architecture and they can be used with both Intel Xeon processors and Intel Xeon Phi coprocessors. This uniformity can greatly reduce the complexity of developing, optimizing and maintaining a software code. Applications still need to use parallel programming for any parallel hardware, but with Intel's approach developers can rely on common model, tools and experiences without needing to switch to master new tools and proprietary programming models. The investment made in parallelizing the code will deliver benefits across a broad range of computing environments.

Higher Efficiency for Parallel Processing

While the Intel Xeon processor E5 family remains the best choice for the majority of applications, Intel Xeon Phi coprocessors provide more efficient performance for highly parallel applications. The emphasis on many more and smaller cores, many more threads, and wider vector units offers a high degree of parallelism. More parallelism compensates for the lower speed of each individual core to deliver higher aggregate performance for workloads that can be subdivided into a sufficiently large number of simultaneous tasks.

Intel Xeon Phi Coprocessors: Key Features

The Intel Xeon Phi coprocessor is an innovation milestone that contains several capabilities that will help accelerate HPC opportunities.

- Memory fit for HPC:

The Intel Xeon Phi coprocessor boasts up to 8 GB of memory capacity (7.75GB with ECC enabled), with a peak memory bandwidth of 352 Gigabytes/second. This level of bandwidth ensures that data can be quickly aggregated and analyzed.

- Cache where it matters:
  With up to 30.5 MB of L2 cache, Intel Xeon Phi coprocessor is able to deliver the right level of compute power for highly parallel workloads.

- Wider data parallelism via an advanced 512-bit wide vector engine

- Supported by many software tool vendors as part of standard tools

- Supported Host Operating Systems: Red Hat Enterprise Linux 6.x or SuSE Linux 12+

The First Intel® Xeon Phi™ Coprocessor Products

SKU #	Form Factor, Thermal	Peak Double Precision	Cores	Clock Speed (GHz)	GDDR5 Memory Speeds (GT/s)	Peak Memory BW	Memory Capacity (GB)	L2 Cache (MB)	Coprocessor (Board) TDP (Watts)	Process
SE10P (special edition)	PCIe Card, Passively Cooled	1073 GF	61	1.1	5.5	352	8	30.5	300	22nm
SE10X (special edition)	PCIe Card, No Thermal Solution	1073 GF	61	1.1	5.5	352	8	30.5	300	
5110P	PCIe Card, Passively Cooled	1011 GF	60	1.053	5	320	8	30	225	
3100 Series	PCIe Card, Actively Cooled	>1 TF	Disclosed at 3100 series launch (H1'13)		5	240	6	>25	300	
	PCIe Card, Passively Cooled	> 1 TF			5	240	6	>25	300	

Performance

The performance results released here by Intel provide a conservative approach to benchmarking by only comparing well-optimized processors performance with well-optimized coprocessor performance. This avoids overstating the value of the coprocessor that would happen if unoptimized processor performances were used. Also, two processors are used to compare to one coprocessor, this in an effort to compare equal power utilization to avoid overstating the coprocessor while allowing it to consume more power than what it is compared against.

Synthetic Benchmark Summary

Benchmark	Result
SGEMM	Up to 2.9x improvement
DGEMM	Up to 2.8x improvement
SMP Linpack	Up to 2.6x improvement
STREAM Triad	Up to 2.2x improvement

*Benchmarks were measured utilizing a single Intel® Xeon Phi™ coprocessor vs. a 2-socket Intel® Xeon® processor E5-2670-based server. Only parallel, threaded and vectorized codes were compared. Actual Intel Math Kernel Library (MKL) was used for SGEMM, DGEMM and SMP Linpack, not a custom benchmarking library.

Application Performance Results

Customer	Application	Performance Increase[1] vs. 2S Xeon*
Los Alamos	Molecular Dynamics	Up to 2.52x
Acceleware	8th order isotropic variable velocity	Up to 2.05x
Jefferson Labs	Lattice QCD	Up to 2.27x
Financial Services	Black-Scholes SP Monte Carlo SP	Up to 10.75x Up to 8.92x
Sandia Labs	MiniFF (Finite Element Solver)	Up to 1.7x[3]
Sinopec	Seismic Imaging	Up to 2.53x[2]

About Intel

Intel (NASDAQ: INTC) is a world leader in computing innovation. The company designs and builds the essential technologies that serve as the foundation for the world's computing devices. Additional information about Intel is available at newsroom.intel.com and blogs.intel.com.

– 30 –

CONTACT: Radoslaw Walczyk
Intel Corporation
408-765-0012
radoslaw.walczyk@intel.com

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Relative performance is calculated by assigning a baseline value of 1.0 to one benchmark result, and then dividing the actual benchmark result for the baseline platform into each of the specific benchmark results of each of the other platforms, and assigning them a relative performance number that correlates with the performance improvements reported.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations.

Synthetic Benchmark Notes

1. Intel® Xeon® Processor E5-2670 used for all SGEMM Matrix = 13824 x 13824, DGEMM Matrix 7936 x 7936, SMP Linpack Matrix 30720 x 30720
2. Intel® Xeon Phi™ coprocessor SE10P (ECC on) with "Gold Release Candidate" SW stack SGEMM Matrix = 15360 x 15360, DGEMM Matrix 7680 x 7680, SMP Linpack Matrix 26872 x 28672

Application Performance Notes

* Xeon = Intel® Xeon® processor;
* Xeon Phi = Intel® Xeon Phi™ coprocessor

1. 2S Xeon* vs. 1 Xeon Phi* (preproduction HW/SW & Application running 100% on coprocessor unless otherwise noted)
2. 2S Xeon* vs. 2S Xeon* + 2 Xeon Phi* (offload)
3. 4 node cluster, each node with 2S Xeon* (comparison is cluster performance with and without 1 Xeon Phi* per node) (Hetero)
4. Intel Measured Oct. 2012