



PRODUCT BRIEF

The Intel® Xeon Phi™ Coprocessor 5110P

Highly-parallel Processing for Unparalleled Discovery



Breakthrough Performance for Your Highly-Parallel Applications

Extracting extreme performance from highly-parallel applications just got easier—much easier. The Intel® Xeon Phi™ coprocessor 5110P, based on Intel® Many Integrated Core (MIC) architecture, complements the industry-leading performance and energy-efficiency of the Intel® Xeon® processor E5 family to enable dramatic performance gains for some of today's most demanding applications. You can now achieve optimized performance for even your most highly-parallel technical computing workloads, while maintaining a unified hardware and software environment.¹

Intel® Xeon Phi™ Coprocessor 5110P

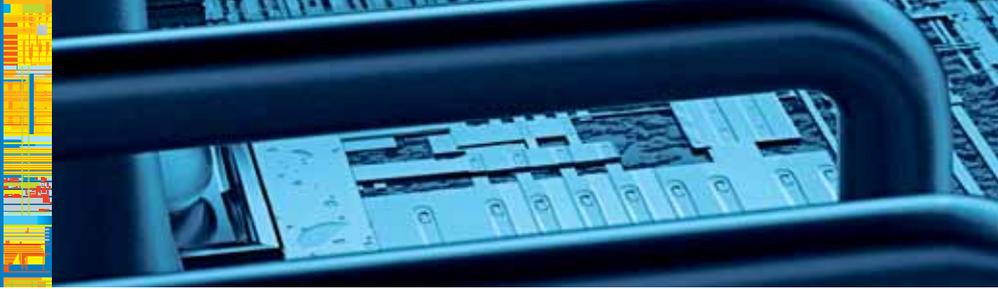
Key Specifications:

- Up to 1 teraflops double-precision performance^{1,3}
- Exceptional performance-per-watt for highly parallel workloads
- Single programming model for all your code
- Flexible usage models to maximize your investment

"Moving a code to Intel Xeon Phi might involve sitting down and adding a couple lines of directives that takes a few minutes. Moving a code to a GPU is a project."²

—Dan Stanzione, Deputy Director at Texas Advanced Computing Center

The Intel Xeon Phi Coprocessor



Intel Xeon Phi Coprocessor 5110P Specifications

Ideal for:

- Highly parallel applications using over 100 threads
- Memory bandwidth-bound applications
- Applications with extensive vector use

Key Specifications:

- 60 cores/1.053 GHz/240 threads
- Up to 1 teraflops double-precision performance³
- 8 GB memory and 320 GB/s bandwidth
- Standard PCIe* x16 form factor
- Linux* operating system, IP addressable
- Supported by the latest Intel® software development products
- 512-bit wide vector engine
- 32 KB L1 I/D cache, 512 KB L2 cache (per core)
- 8 GB GDDR5 memory (up to 320 GB/s)
- 225W TDP
- X16 PCIe form factor (requires IA host)
- Host OS: Red Hat Enterprise Linux 6.x, SuSE Linux 12+

A Single Programming Model for All Your Code

A wide assortment of programming languages, models, and tools support Intel architecture and all of them can be used with both Intel Xeon processors and Intel Xeon Phi coprocessors. Applications that run on one processor family will run on the other. This uniformity can greatly reduce the complexity of developing, optimizing, and maintaining your software code. Existing applications will need to be tuned and recompiled for parallelism to maximize throughput, but your developers won't need to rethink the entire problem and they won't need to master new tools and proprietary programming models. Instead, they can reuse existing code and maintain a common code base using familiar tools and methods.

Code can be optimized just once for both Intel Xeon processors and Intel Xeon Phi coprocessors. The same techniques—such as scaling applications to many cores and threads, blocking data for hierarchical memory and caches, and effective use of SIMD—deliver optimal performance for both processor and coprocessor families. The investment you make in parallelizing your code will deliver benefits across the full range of computing environments.

Even Higher Efficiency for Parallel Processing

While the Intel Xeon processor E5 family remains the preferred choice for the majority of applications, Intel Xeon Phi coprocessors provide more efficient performance for highly-parallel applications. They include many more and smaller cores, many more threads, and wider vector units. The high degree of 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. You can use Intel Xeon processors and Intel Xeon Phi coprocessors together to optimize performance for almost any workload. Because both processor and coprocessor support the same software code and programming models, your developers won't have to reinvent the wheel to deliver optimized performance.

Intel® Xeon Phi™ Coprocessor 5110P Overview

FEATURES	BENEFITS
Intel® Many Integrated Cores (MIC) architecture	Large numbers of cores and threads enable efficient execution of hundreds of simultaneous tasks to dramatically boost aggregate performance for highly-parallel applications.
Familiar Intel® architecture programming model	Enables broad reuse of existing code. Also allows developers to use familiar tools and methods and to maintain a common code base for Intel® Xeon® processors and Intel® Xeon Phi™ coprocessors.
Linux* hosting capability	Can operate as a dependent coprocessor or an independent server node to enable flexible usage models and optimized support for diverse hardware and software environments.
IP Addressable	Supports standard clustering models for simple integration into clustered environments.
Intel 22 nm technology with 3-D Tri-Gate transistors	Provides exceptional compute density and energy efficiency.
Up to 8 coprocessors per host server (requires one PCIe* slot per coprocessor)	Enables simple scaling of highly-parallel execution resources to deliver desired performance levels.

Table 1.

Better Performance, More Flexibility

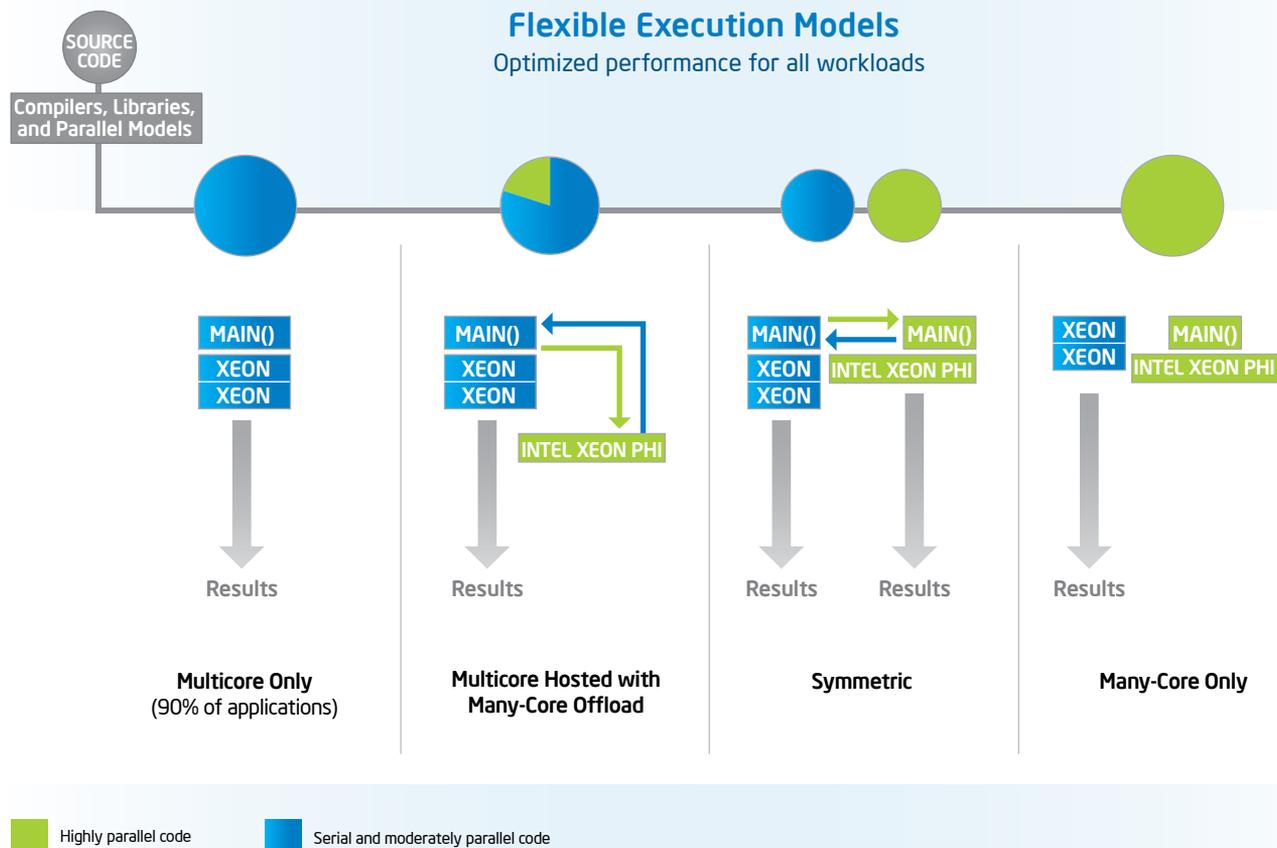


Figure 1. Computing model comparison.

Flexible Usage Models to Maximize Your Investment

The Intel Xeon Phi coprocessor is designed to provide the highest level of flexibility in conjunction with Intel Xeon processor-based systems and clusters. It can operate under the host server's operating system (OS), in which case the OS and application run on the Intel Xeon processors and highly-parallel code segments are off-loaded to the Intel Xeon Phi coprocessor to accelerate performance. However, unlike a basic accelerator, the Intel Xeon Phi coprocessor can also function as an independent server node with its own Linux OS and IP address. In this scenario, it can run applications independently and it can off-load serial code segments (or moderately parallel code segments) to the host system so they perform more quickly and efficiently. The flexibility provided by these new and exclusive usage models is illustrated in Figure 1.

A single Intel Xeon Phi coprocessor provides up to 60 cores and 240 threads and can deliver up to a teraflop of double-precision performance for targeted applications. These coprocessors are manufactured using Intel's industry-leading 22 nm technology to provide exceptional compute density and energy efficiency and they feature the world's first 3-D Tri-Gate transistors. The Intel Xeon Phi coprocessor 5110P comes in a standard PCIe x16 form factor that can be added to a supported Intel® Xeon® processor-based server. Up to eight cards can be used with a single two-socket host server to scale parallel processing capability.

Get Started Today!

The Intel® Xeon Phi™ coprocessor can dramatically accelerate performance for your highly-parallel applications to help you push the boundaries of innovation and scientific discovery—without requiring your software developers to reinvent the wheel.

Driving Supercomputing to New Heights

Some of today's most successful high performance computing centers are already using Intel Xeon Phi coprocessors to deliver massive new parallel computing capability.

- The Texas Advanced Computing Center will soon launch a 10 petaflop supercomputer that will include thousands of Intel Xeon Phi coprocessors.⁴
- The DEEP project is using the Intel Xeon processor E5 family and Intel Xeon Phi coprocessors as the foundation for "an exascale-enabling supercomputing platform."⁵
- Intel built a small prototype cluster using Intel Xeon Phi coprocessors that ranked as the 150th largest supercomputer in the world.⁶



Is Intel® Xeon Phi™ Coprocessor Right for Me?

While a majority of applications will achieve maximum performance running on Intel Xeon processors, certain highly-parallel applications will benefit dramatically by using Intel Xeon Phi coprocessors. To qualify as highly-parallel, an application must scale well to over one-hundred threads, and either make extensive use of vectors or efficiently use more local memory bandwidth than is available on an Intel Xeon processor. Intel has many software tools to help optimize your code and determine whether your application is best suited for Intel Xeon Phi coprocessors. Learn more at intel.com/software/products.

Learn more at intel.com/xeonphi

¹ 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. For more information go to <http://www.intel.com/performance>

² hpcwire.com/hpcwire/2011-04-21/tacc_steps_up_to_the_mic.html

³ Claim based on calculated theoretical peak double precision performance capability for a single coprocessor. 16 DP flops/clock/core * 60 cores * 1.053 GHz = 1.01088 Tflops.

⁴ Read the Intel press release at: http://newsroom.intel.com/community/intel_newsroom/blog/2011/09/22/chip-shot-intel-xeon-processors-intel-mic-co-processors-to-power-10-petaflop-supercomputer

⁵ Get more information at: http://www.deep-project.eu/deep-project/EN/Project/Hardware/_node.html;jsessionid=508843AD062718950894B545578887A2

⁶ As measured by the June 2012 Top 500 list (www.top500.org). Get more information at: <http://blogs.intel.com/technology/2012/06/intel-xeon-phi-coprocessors-accelerate-discovery-and-innovation/>

Copyright © 2012 Intel Corporation. All rights reserved. Intel, the Intel logo, the Intel Inside logo, Xeon, and Xeon Phi are trademarks of Intel Corporation in the U.S. and other countries.

* Other names and brands may be claimed as the property of others.

