Intel® Xeon Phi™ Coprocessor Architecture Overview

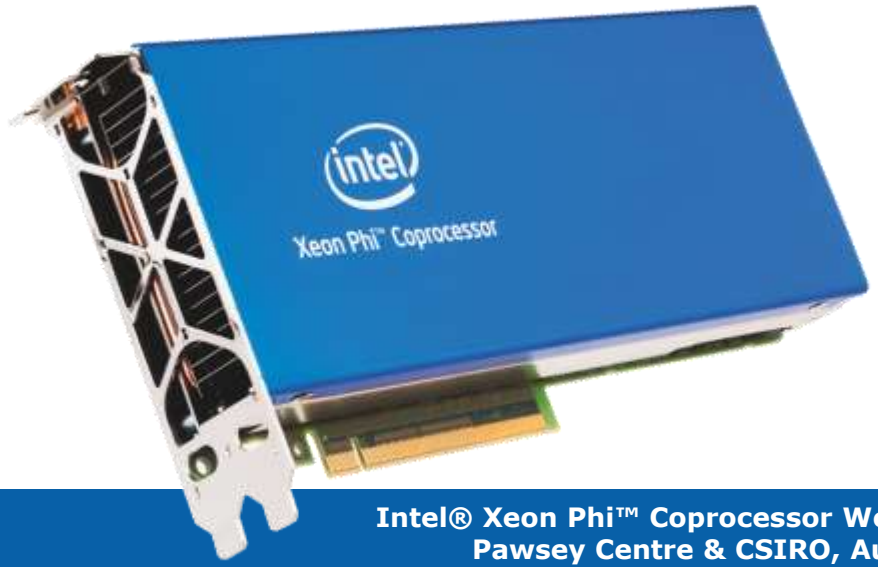

Legal Disclaimer

- INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPETY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL® PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.
- Intel may make changes to specifications and product descriptions at any time, without notice.
- All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.
- Intel, processors, chipsets, and desktop boards may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.
- Sandy Bridge and other code names featured are used internally within Intel to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user
- 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
- Intel, Core, Xeon, VTune, Cilk, Intel and Intel Sponsors of Tomorrow. and Intel Sponsors of Tomorrow. logo, and the Intel logo are trademarks of Intel Corporation in the United States and other countries.
- *Other names and brands may be claimed as the property of others.
- Copyright ©2013 Intel Corporation.
- Hyper-Threading Technology: Requires an Intel® HT Technology enabled system, check with your PC manufacturer. Performance will vary depending on the specific hardware and software used. Not available on all Intel® Core™ processors. For more information including details on which processors support HT Technology, visit http://www.intel.com/info/hyperthreading
- Intel® 64 architecture: Requires a system with a 64-bit enabled processor, chipset, BIOS and software. Performance will vary depending on the specific hardware and software you use. Consult your PC manufacturer for more information. For more information, visit http://www.intel.com/info/em64t
- Intel® Turbo Boost Technology: Requires a system with Intel® Turbo Boost Technology capability. Consult your PC manufacturer. Performance varies depending on hardware, software and system configuration. For more information, visit http://www.intel.com/technology/turboboost

Architecture Topics

- Intel® Many Integrated Core (MIC) Architecture
- Intel® Xeon Phi™ Coprocessor Overview
- Core and Vector Processing Unit
- Setting Expectations
- Performance
- Summary

Module Outline

- **Intel® Many Integrated Core (MIC) Architecture**
- Intel® Xeon Phi™ Coprocessor Overview
- Core and Vector Processing Unit
- Setting Expectations
- Performance
- Summary

Intel Architecture Multicore and Manycore
More cores. Wider vectors. Coprocessors.

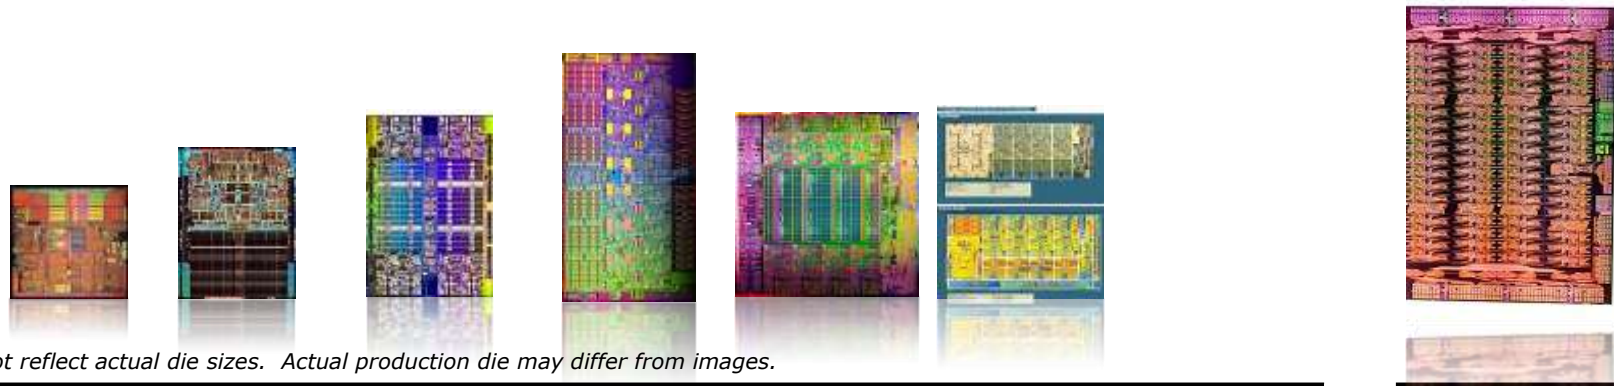

Images do not reflect actual die sizes. Actual production die may differ from images.

	Intel® Xeon® processor 64-bit	Intel® Xeon® processor 5100 series	Intel® Xeon® processor 5500 series	Intel® Xeon® processor 5600 series	Intel® Xeon® processor E5 Product Family	Intel® Xeon® processor code name Ivy Bridge	Intel® Xeon® processor code name Haswell	Intel® Xeon Phi™ Coprocessor
Core(s)	1	2	4	6	8	10	To be determined	61
Threads	2	2	8	12	16	20		244

Intel® Xeon Phi™ coprocessor extends established CPU architecture and programming concepts to highly parallel applications

Intel® Multicore Architecture

- Foundation of HPC Performance
- Suited for full scope of workloads
- Industry leading performance and performance/watt for serial & parallel workloads
- Focus on fast single core/thread performance with "moderate" number of cores

Intel® Many Integrated Core Architecture

- Performance and performance/watt optimized for highly parallelized compute workloads
- Common software tools with Intel® Xeon® architecture enabling efficient application readiness and performance tuning
- Intel® Architecture extension to Manycore
- Many cores/threads with wide SIMD

Consistent Tools & Programming Models

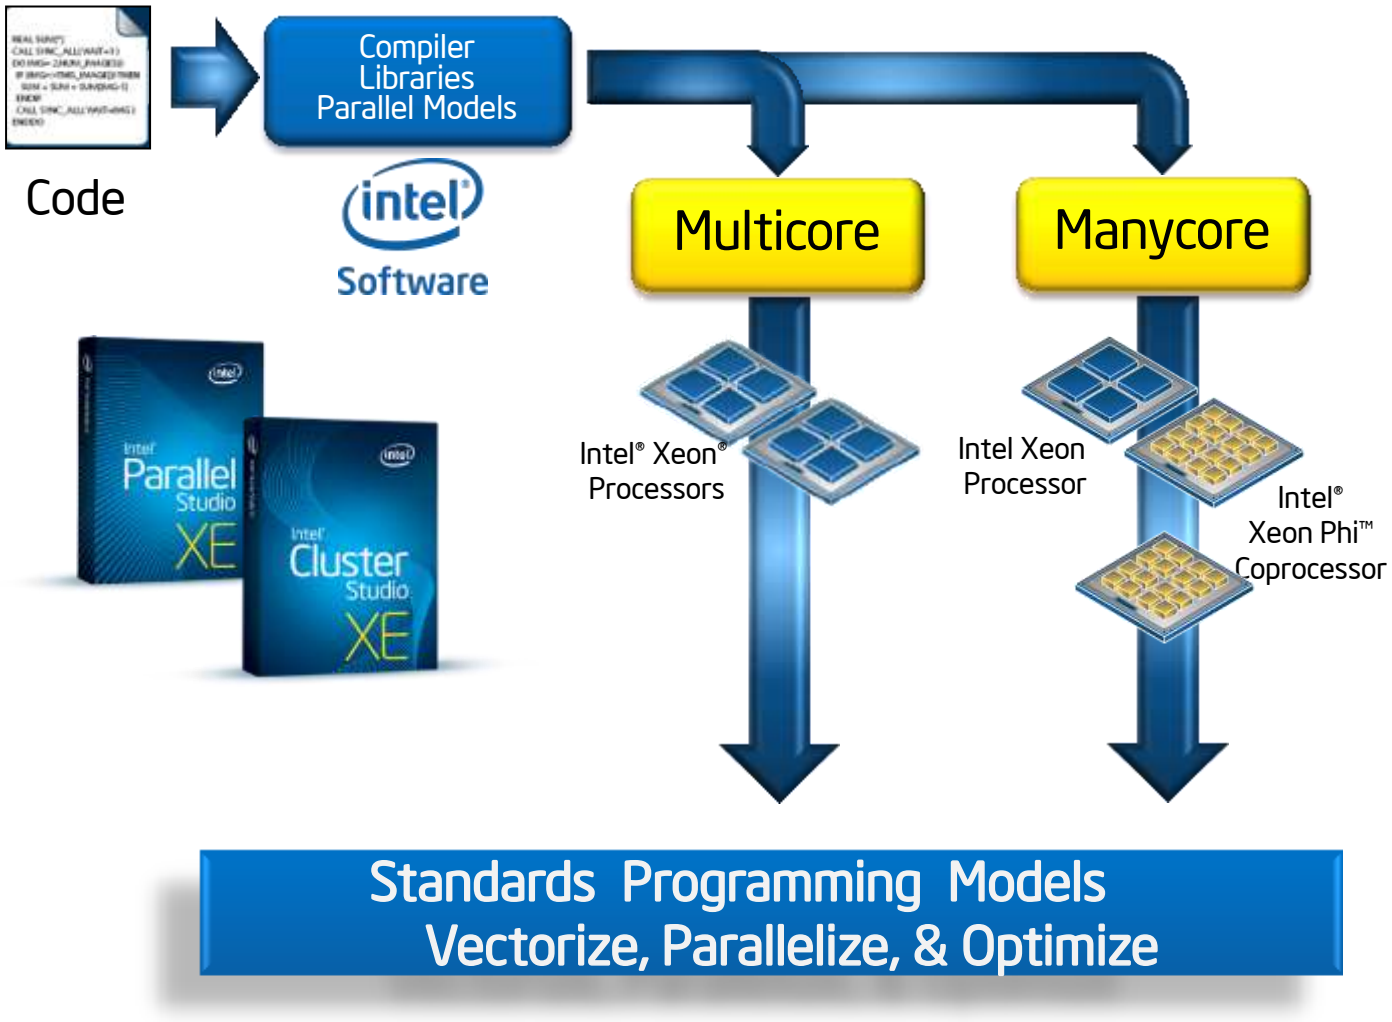

Module Outline

- Intel® Many Integrated Core (MIC) Architecture
- **Intel® Xeon Phi™ Coprocessor Overview**
- Core and Vector Processing Unit
- Setting Expectations
- Performance
- Summary

Introducing Intel® Xeon Phi™ Coprocessors

Highly-parallel Processing for Unparalleled Discovery

Groundbreaking differences

Up to 61 Intel® Architecture cores/1.1 GHz/ 244 threads

Up to 8 GB memory with up to 352 GB/s bandwidth

512-bit SIMD instructions

Linux* operating system, IP addressable

Standard programming languages and tools

Leading to Groundbreaking results

Over 1 TeraFlop/s double precision peak performance[1]

Up to 2.2x higher memory bandwidth than on an Intel® Xeon® processor E5 family-based server[2]

Up to 4x more performance per watt than with an Intel Xeon processor E5 family-based server[3]

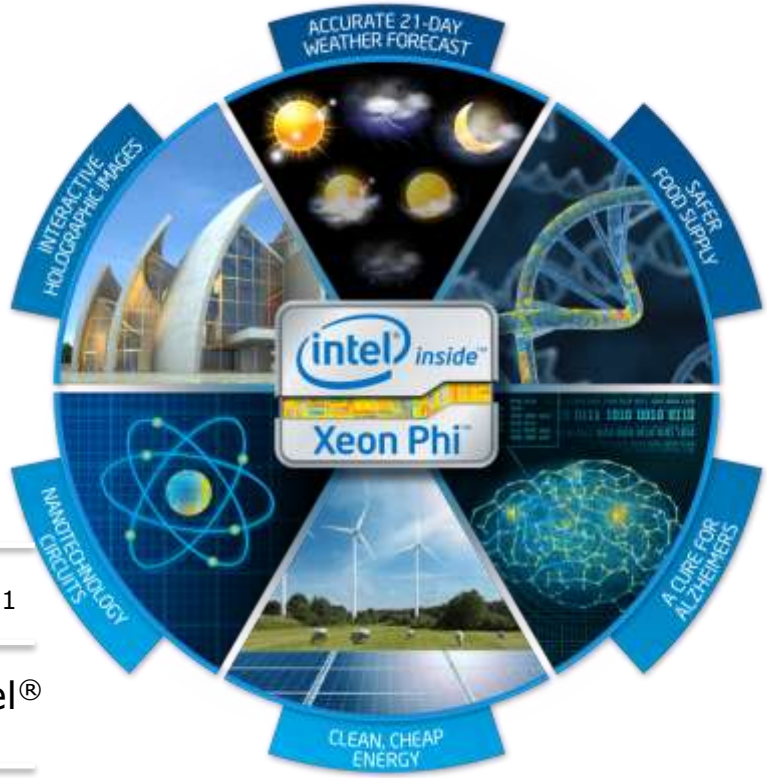

Intel® Xeon Phi™ Architecture Overview

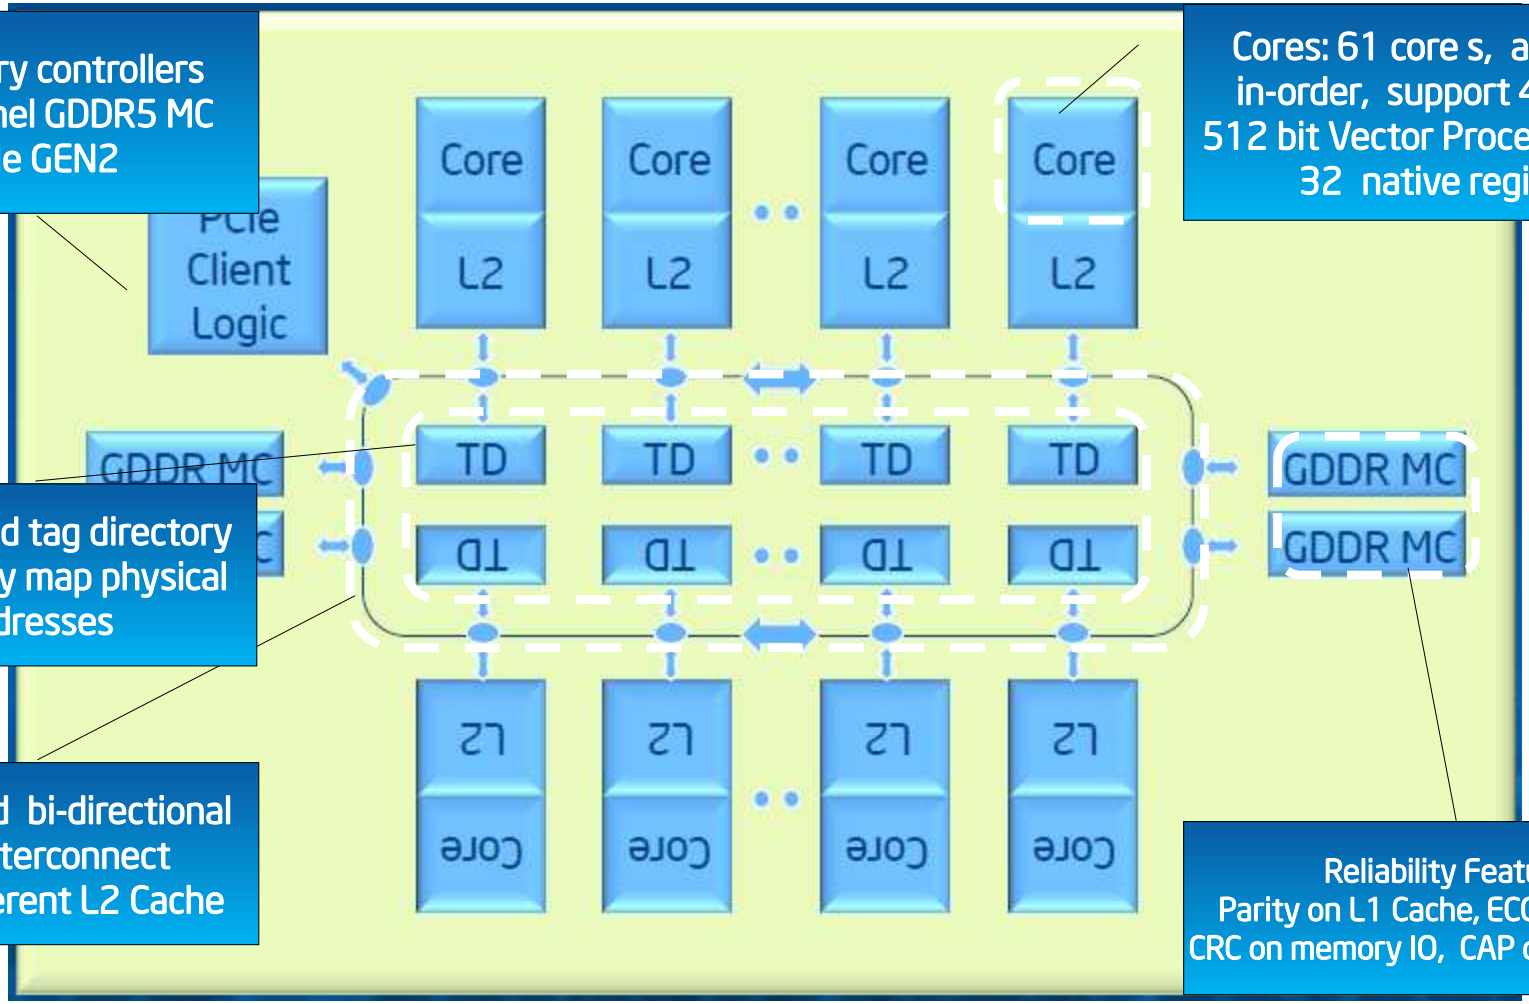

Core Architecture Overview

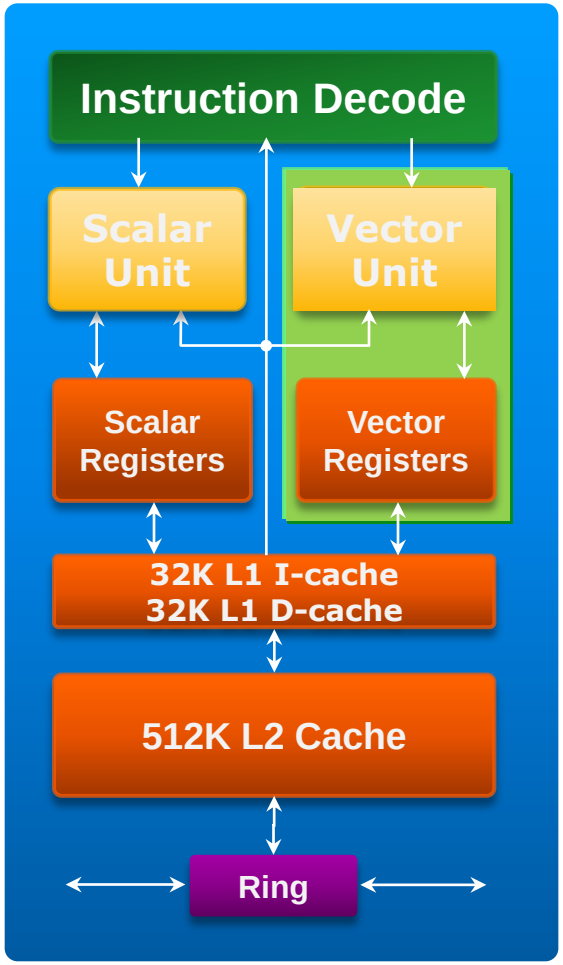

- 60+ in-order, low-power Intel® Architecture cores in a ring interconnect
- Two pipelines
  - Scalar Unit based on Pentium® processors
  - Dual issue with scalar instructions
  - Pipelined one-per-clock scalar throughput
- SIMD Vector Processing Engine
- 4 hardware threads per core
  - 4 clock latency, hidden by round-robin scheduling of threads
  - Cannot issue back-to-back inst in same thread
- Coherent 512 KB L2 Cache per core

Core Architecture Overview

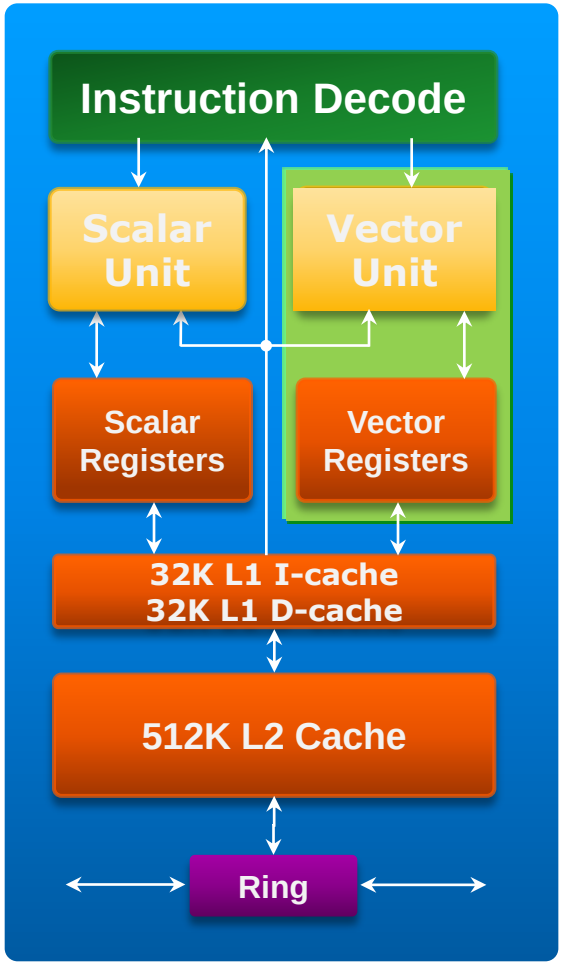

- 2 issue (1 scalar/1 vector)
- 2 cycle decoder: no back-to-back cycle issue from the same context (thread)
- Most vec instructions have 4 clock latency
- At least two HW contexts (thread/proc) to fully utilize the core

Core Architecture Overview

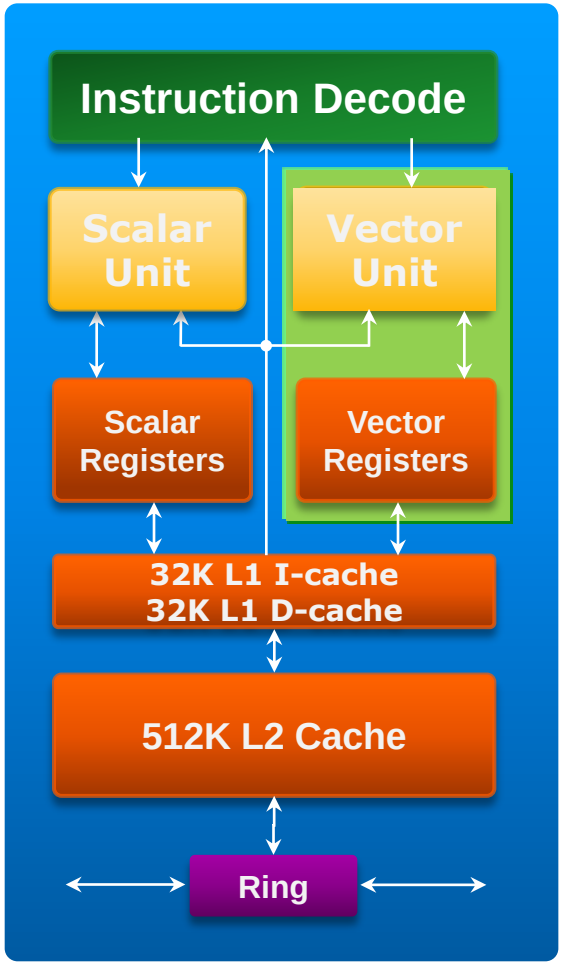

- Performance Monitoring Unit
  - 4 event select register
  - 4 performance counters
  - Shared among the 4 HW threads
  - Programmable via model specific registers (MSR) using RDMSR/WRMSR

Knights Corner Architecture Overview – Cache

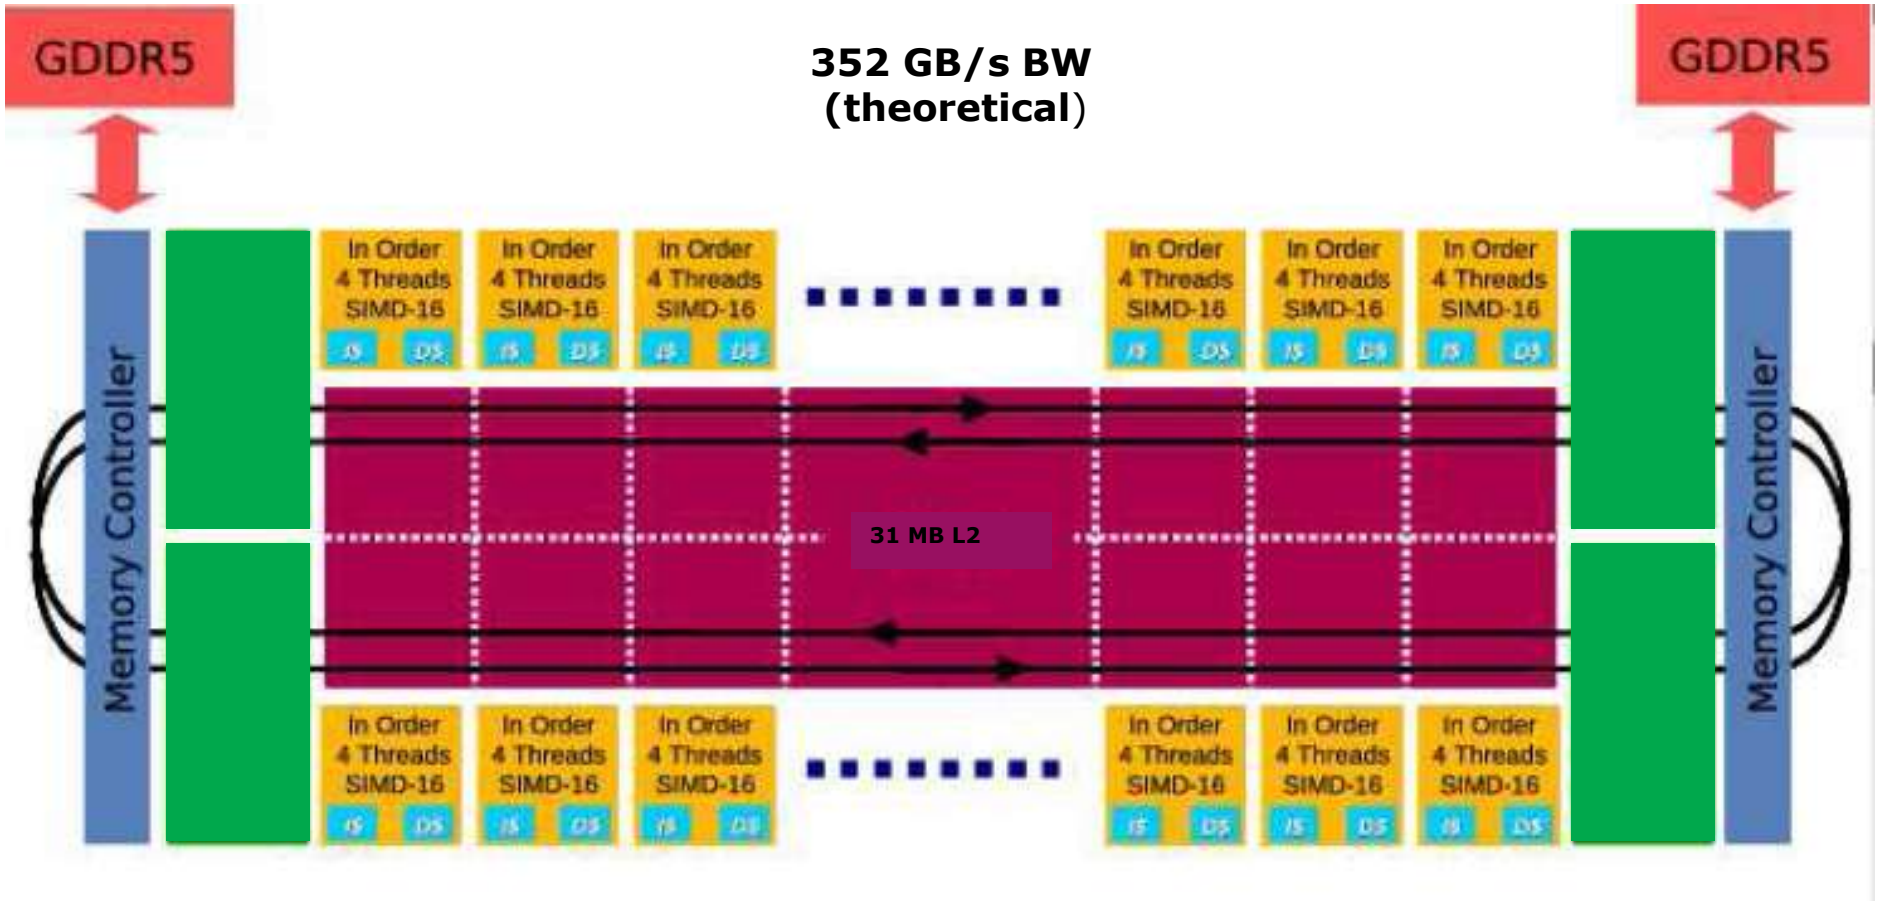

Knights Corner Architecture Overview – Cache

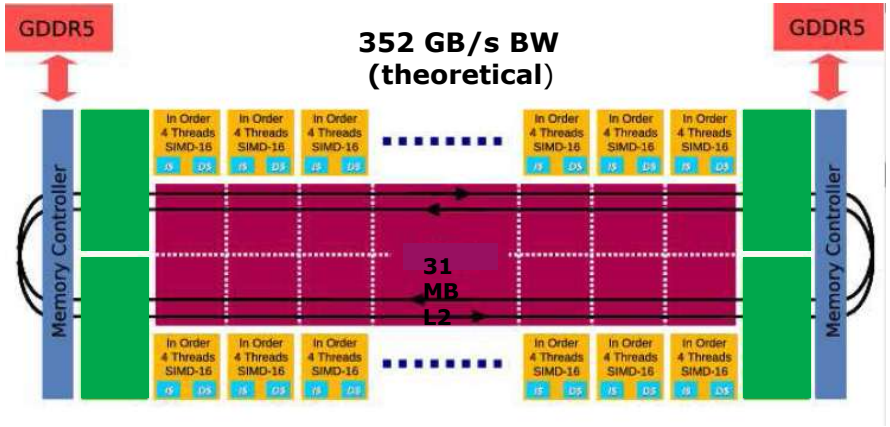

- L1 cache
  - 32K I-cache per core
  - 32K D-cache per core
  - 8 way associative
  - 64B cache line
  - 3 cycle access
  - Up to 8 outstanding requests
  - Fully coherent (MESI)

- L2 cache
  - 512K Unified per core
  - 8 way assoc
  - Inclusive
  - 31M total across 62 cores
  - 11 cycle best access
  - Up to 32 outstanding requests
  - Streaming HW prefetcher
  - Fully coherent

Knights Corner Architecture Overview – Cache

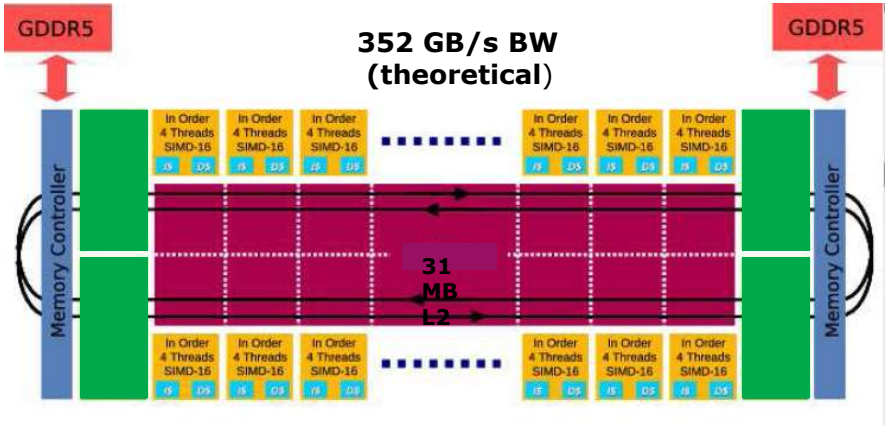

- Alignment
  - Based upon number of elements, element size, and vector load and store instruction
  - 64B alignment for 4B (float) data elements for a 16 to 16 vector load

- Memory
  - 8GB GDDR5
  - 8 Memory controllers, 16 GDDR5 channels, up to 5.5 GT/s
  - 300 ns access
  - Aggregate 352 GB/s peak memory bandwidth
  - ECC
- PCI Express*
  - Gen2 (Client) x16 per direction

Module Outline

- Intel® Many Integrated Core (MIC) Architecture
- Intel® Xeon Phi™ Coprocessor Overview
- **Core and Vector Processing Unit**
- Setting Expectations
- Performance
- Summary

Knights Corner Architecture Overview
Vector Processing Unit and ISA

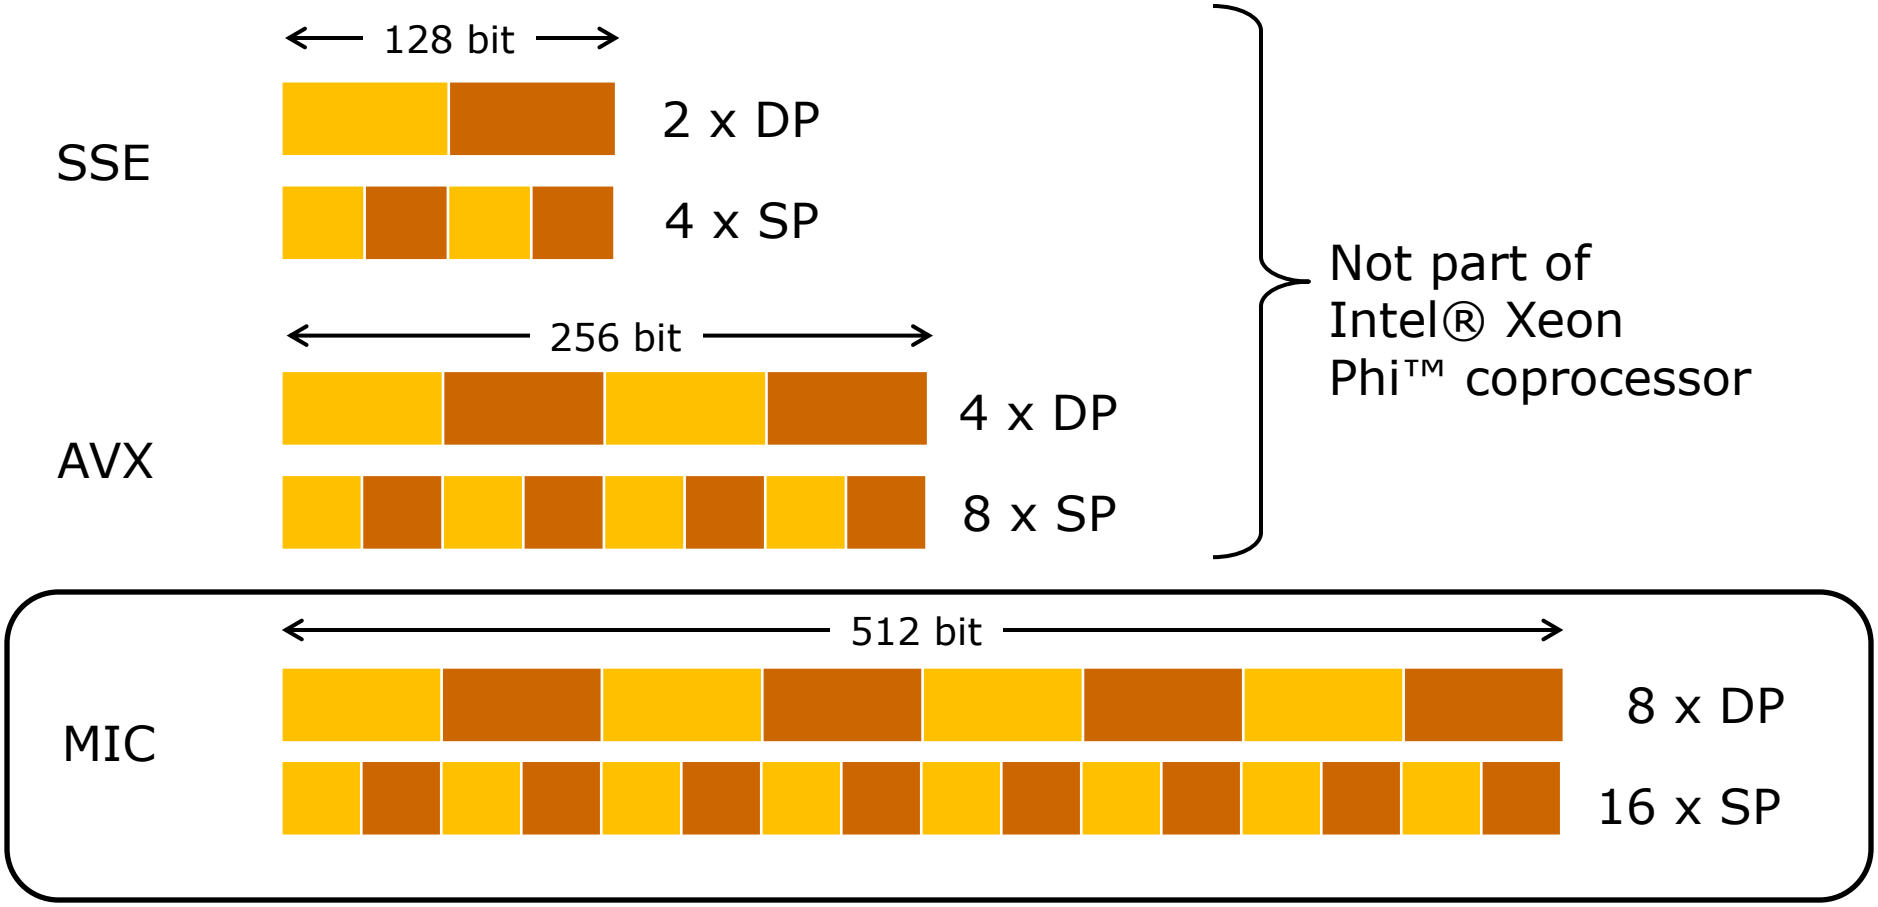

Vector Processing Unit Extends the Scalar IA Core

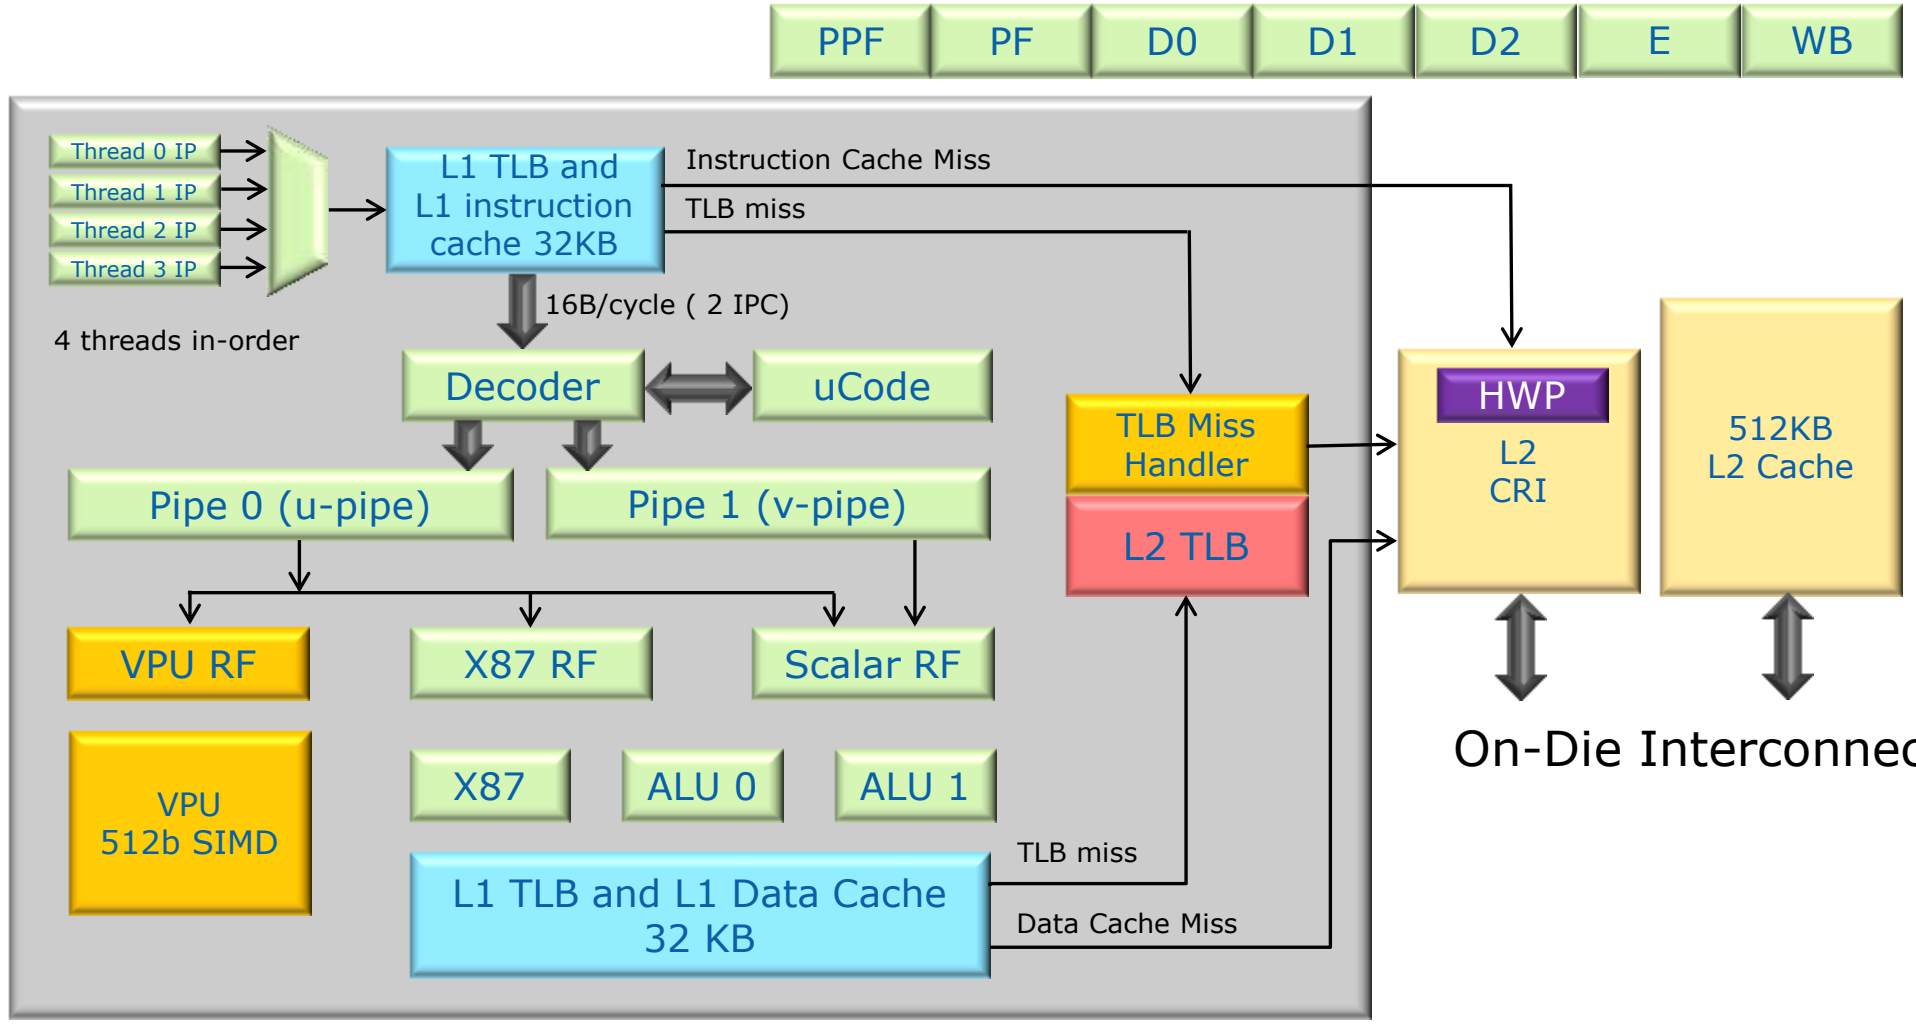

Vector Processing Unit and Intel® Initial Many Core Instructions (Intel® IMCI)

- Vector Processing Unit Execute Intel IMCI
- 512-bit Vector Execution Engine
  - 16 lanes of 32-bit single precision and integer operations
  - 8 lanes of 64-bit double precision and integer operations
  - 32 512-bit general purpose vector registers in 4 thread
  - 8 16-bit mask registers in 4 thread for predicated execution
- Read/Write
  - One vector length (512-bits) per cycle from/to Vector Registers
  - One operand can be from memory
- IEEE 754 Standard Compliance
  - 4 rounding models, even, 0, +∞, -∞
  - Hardware support for SP/DP denormal handling
  - Sets status register VXCSR flags but not hardware traps

Vector Instruction Performance

- VPU contains 16 SP ALUs, 8 DP ALUs
- Most VPU instructions have a latency of 4 cycles and TPT 1 cycle
  - Load/Store/Scatter have 7-cycle latency
  - Convert/Shuffle have 6-cycle latency
- VPU instruction are issued in u-pipe
- Certain instructions can go to v-pipe also
  - Vector Mask, Vector Store, Vector Packstore, Vector Prefetch, Scalar

21

Module Outline

- Intel® Many Integrated Core (MIC) Architecture
- Intel® Xeon Phi™ Coprocessor Overview
- Core and Vector Processing Unit
- Software Ecosystem
- Performance
- Summary

Programming Considerations - Setting Expectations (1)

- Getting full performance from the Intel® MIC architecture requires both a high degree of parallelism *and* vectorization
  - Not all code can be written this way
  - Not all programs make sense on this architecture

- The Intel® Xeon® processor is not an Intel® Xeon® processor
  - It specializes in running highly parallel and vectorized code
  - New vector instruction set and 512-bit wide registers
  - Not optimized for processing serial code

Programming Considerations - Setting Expectations (2)

Thread setup and comm overhead

Off load

- Coprocessor comes with 8 GB of memory
  - Only ~7 GB is available to your program
    - The other ~1GB is used for data transfer and is accessible to your coprocessor code as buffers.
- Very short (low-latency) tasks not optimal for offload to the coprocessor
  - Costs that you need to amortize to make it worthwhile:
    - Code and data transfer
    - Process/thread creation
  - Fastest data transfers currently require careful data alignment
- This architecture is not optimized for serial performance

Intel® MIC Programming Considerations – This is not a GPU

- Very different memory architectures
  - The Intel MIC Architecture is not optimized for concurrent out-of-cache random memory accesses by large numbers of threads (GPUs are)
  - The Intel MIC Architecture has a "traditional" coherent-cache architecture
  - GPUs have a memory architecture specialized for localized "shared memory" processing

- "Threads" and "cores" mean something very different - GPU versions of these are limited and lighter-weight

- Each architecture (host CPU, Intel MIC Architecture, or GPU) is really good at some things, and not others
  - Because the Intel MIC Architecture is similar to the Intel® Xeon® processor, probably your best choice for further accelerating highly parallel and vectorized code developed on Intel Xeon processor

Module Outline

- Intel® Many Integrated Core (MIC) Architecture
- Intel® Xeon Phi™ Coprocessor Overview
- Core and Vector Processing Unit
- Setting Expectations
- **Performance**
- Summary

Theoretical Maximum

(Intel® Xeon® processor E5-2670 vs. Intel® Xeon Phi™ coprocessor 5110P & SE10P/X)

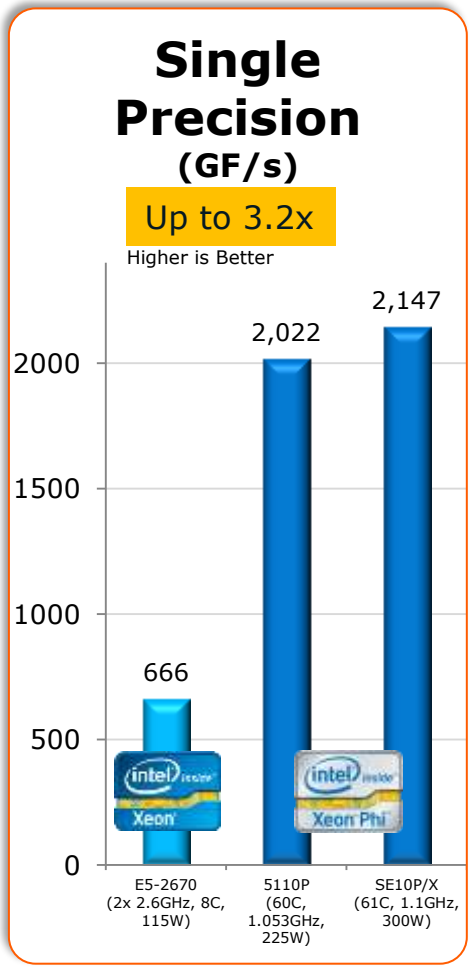

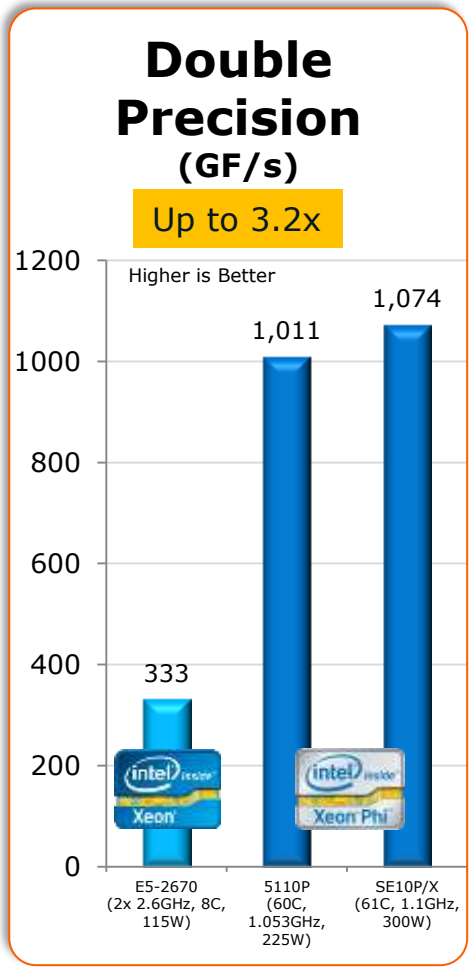

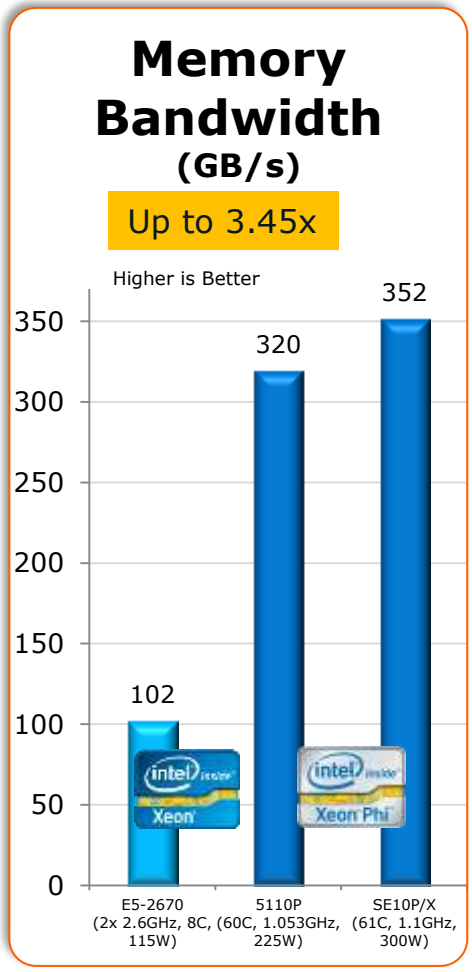

Source: Intel as of October 17, 2012 Configuration Details: Please reference slide speaker notes.
For more information go to http://www.intel.com/performance

Synthetic Benchmark Summary

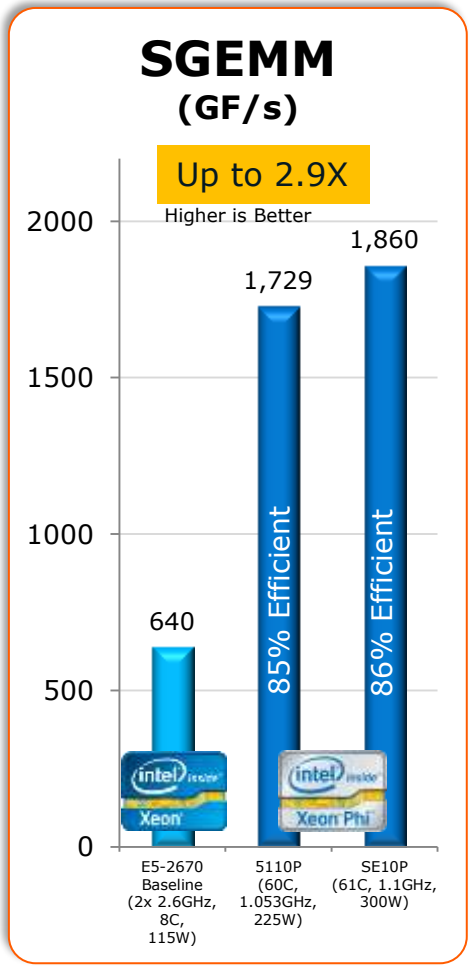

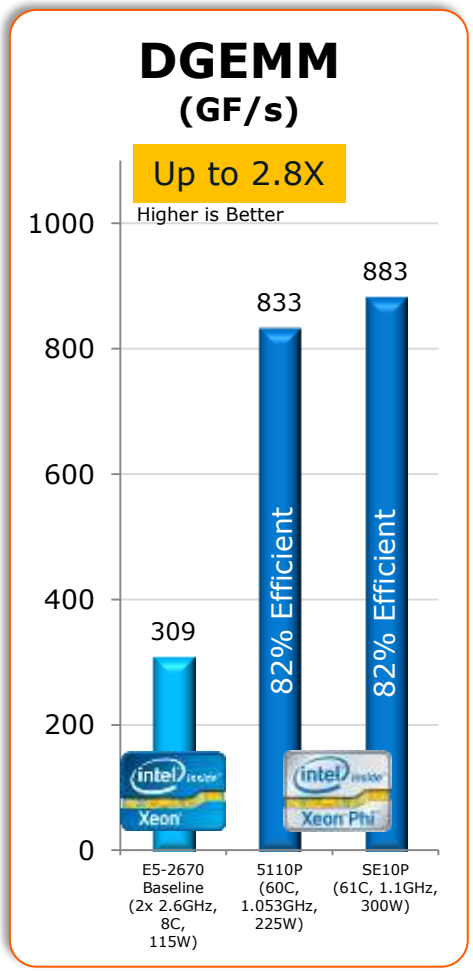

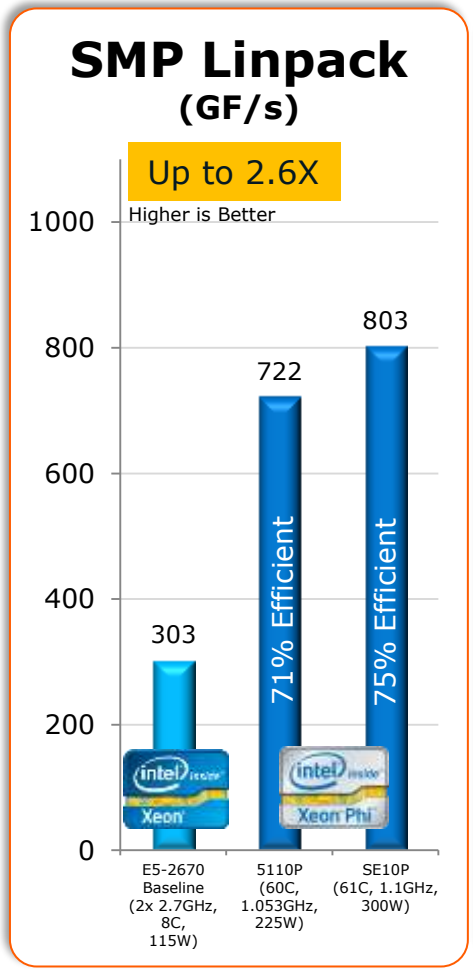

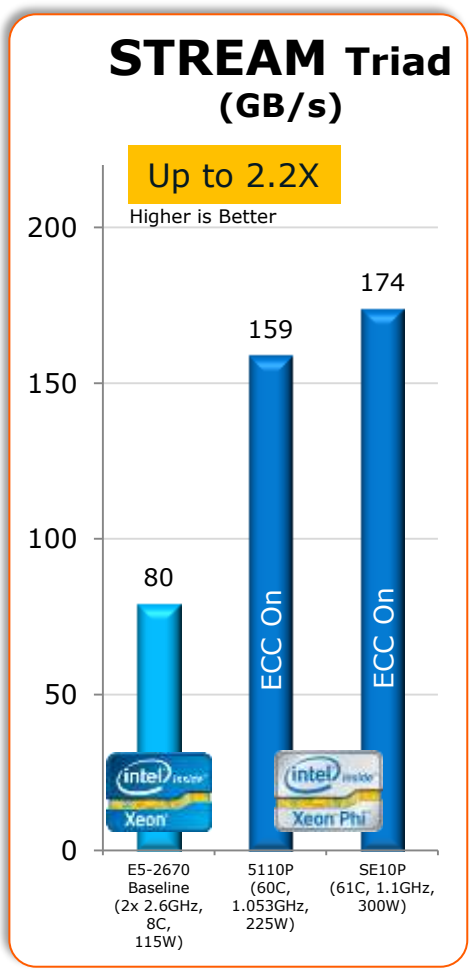

Coprocessor results: Benchmark runs 100% on coprocessor, no help from Intel® Xeon® processor host (aka native). For more information go to http://www.intel.com/performance

Intel® Xeon Phi™ Coprocessor vs. Intel® Xeon® Processor

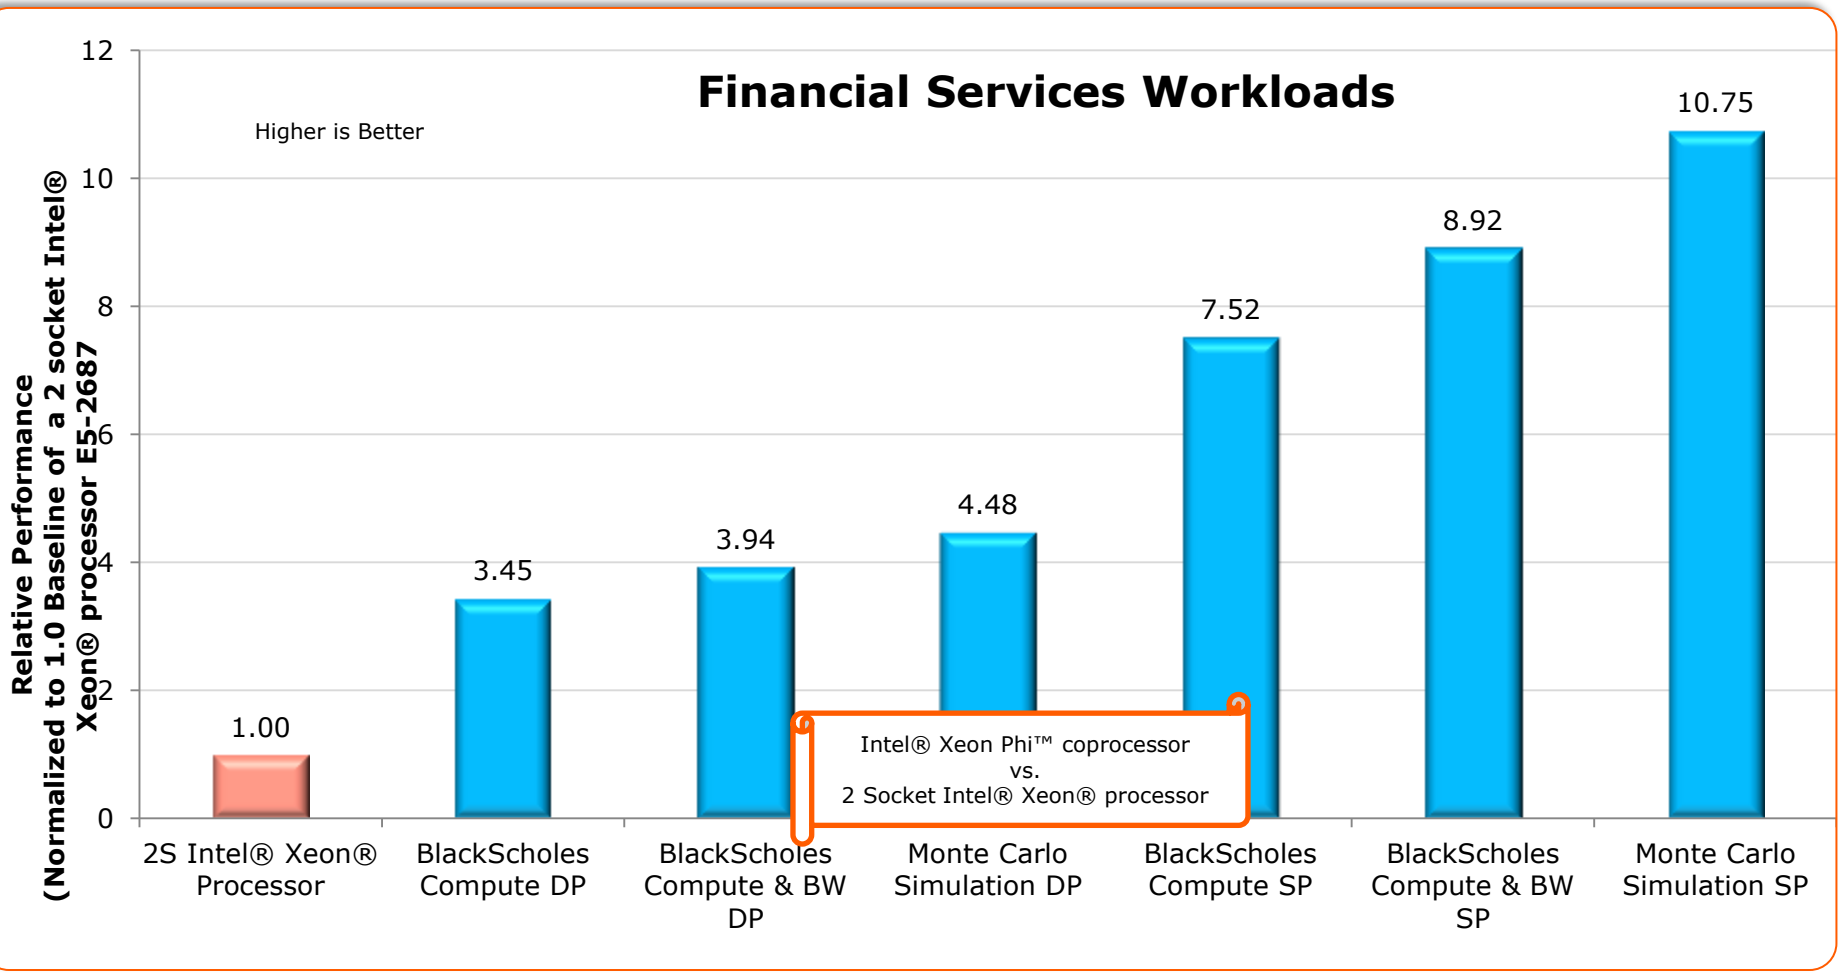

Coprocessor results: Benchmark runs 100% on coprocessor, no help from Intel® Xeon® processor host (aka native). For more information go to http://www.intel.com/performance

Notes:

1. 2 X Intel® Xeon® Processor E5-2670 (2.6GHz, 8C, 115W)
2. Intel® Xeon Phi™ coprocessor SE10 (ECC on) with pre-production SW stack

Higher SP results are due to certain Single Precision transcendental functions in the Intel® Xeon Phi™ coprocessor that are not present in the Intel® Xeon® processor

Module Outline

- Intel® Many Integrated Core (MIC) Architecture
- Intel® Xeon Phi™ Coprocessor Overview
- Core and Vector Processing Unit
- Setting Expectations
- Performance
- **Summary**

Architecture Summary

- Intel® Many Integrated Core (MIC) Architecture
- Intel® Xeon Phi™ Coprocessor Overview
- Core and Vector Processing Unit
- Setting Expectations
- Performance
- Summary

Optimization Notice

Optimization Notice

Intel® compilers, associated libraries and associated development tools may include or utilize options that optimize for instruction sets that are available in both Intel® and non-Intel microprocessors (for example SIMD instruction sets), but do not optimize equally for non-Intel microprocessors. In addition, certain compiler options for Intel compilers, including some that are not specific to Intel micro-architecture, are reserved for Intel microprocessors. For a detailed description of Intel compiler options, including the instruction sets and specific microprocessors they implicate, please refer to the "Intel® Compiler User and Reference Guides" under "Compiler Options." Many library routines that are part of Intel® compiler products are more highly optimized for Intel microprocessors than for other microprocessors. While the compilers and libraries in Intel® compiler products offer optimizations for both Intel and Intel-compatible microprocessors, depending on the options you select, your code and other factors, you likely will get extra performance on Intel microprocessors.

Intel® compilers, associated libraries and associated development tools may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include Intel® Streaming SIMD Extensions 2 (Intel® SSE2), Intel® Streaming SIMD Extensions 3 (Intel® SSE3), and Supplemental Streaming SIMD Extensions 3 (Intel® SSSE3) instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors.

While Intel believes our compilers and libraries are excellent choices to assist in obtaining the best performance on Intel® and non-Intel microprocessors, Intel recommends that you evaluate other compilers and libraries to determine which best meet your requirements. We hope to win your business by striving to offer the best performance of any compiler or library; please let us know if you find we do not.

Notice revision #20101101

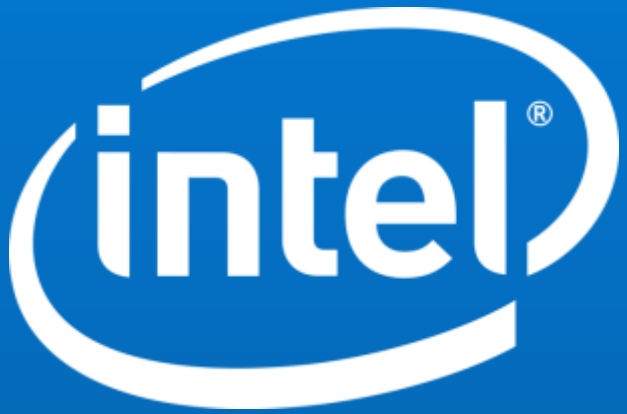
intel