The Architecture of the Intel® Xeon Phi™ Coprocessor

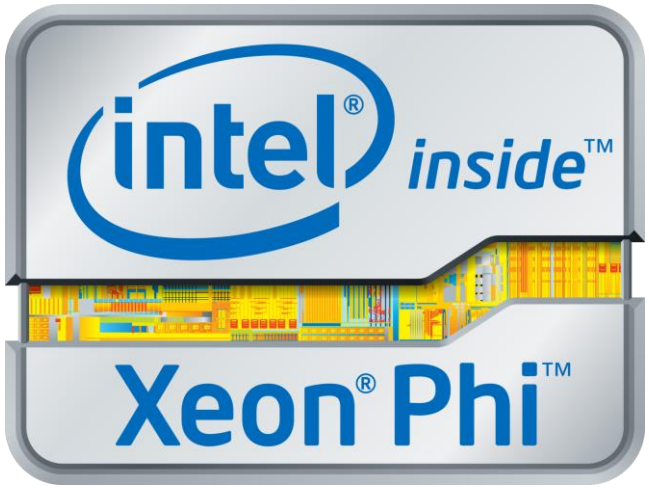

Dr.-Ing. Michael Klemm

Software and Services Group
Intel Corporation
(michael.klemm@intel.com)

Legal Disclaimer & Optimization Notice

INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS". NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

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, reference www.intel.com/software/products.

Optimization Notice

Intel's compilers 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 SSE2, SSE3, and 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. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

References

- Intel® Xeon Phi™ Coprocessor High-Performance Programming
  (J. Jeffers, J. Reinders, published by Morgan Kaufman, ISBN 978-0-12-410414-3)
- Intel® Developer Zone for the Intel® Xeon Phi™ Coprocessor
  http://software.intel.com/mic-developer
- Intel® Xeon Phi™ Coprocessor Developer's Quick Start Guide
  http://software.intel.com/en-us/articles/intel-xeon-phi-coprocessor-developers-quick-start-guide
- Intel® Xeon Phi™ Coprocessor Instruction Set Architecture Reference Manual
  http://download-software.intel.com/sites/default/files/forum/278102/327364001en.pdf

Intel® Xeon Phi™ Coprocessor (former codename: Knights Corner)

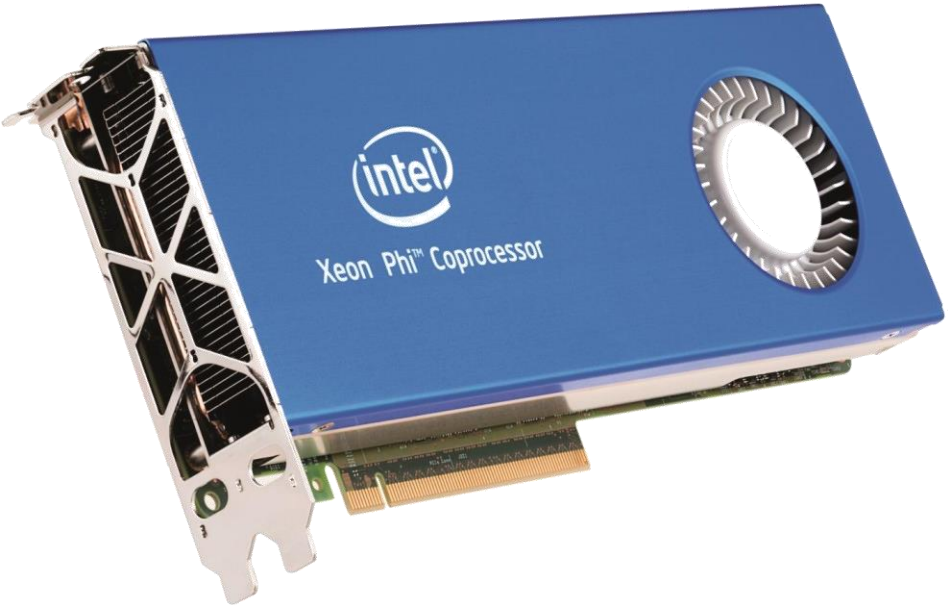

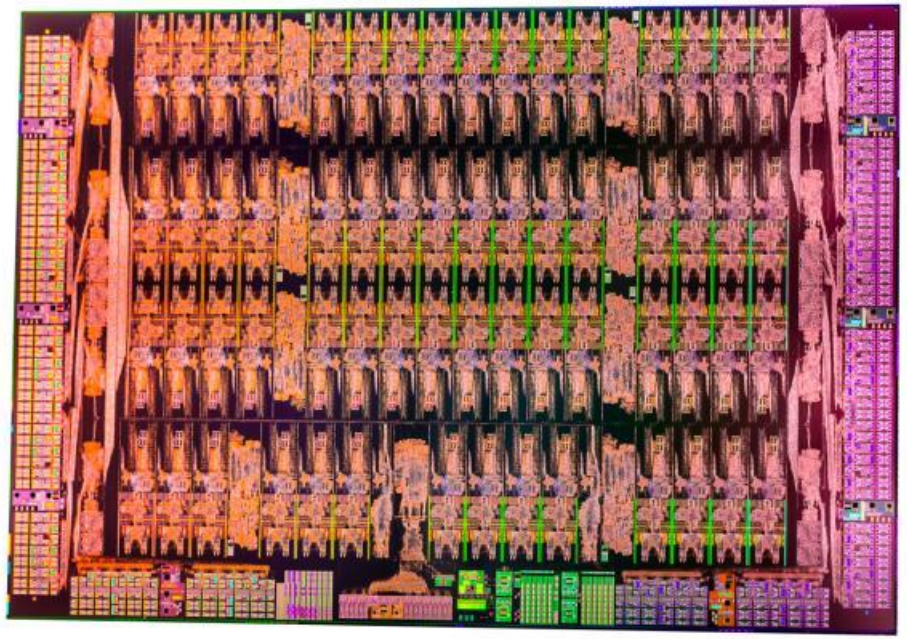

IA-based coprocessor for massively parallel applications.

Intel and Parallelism

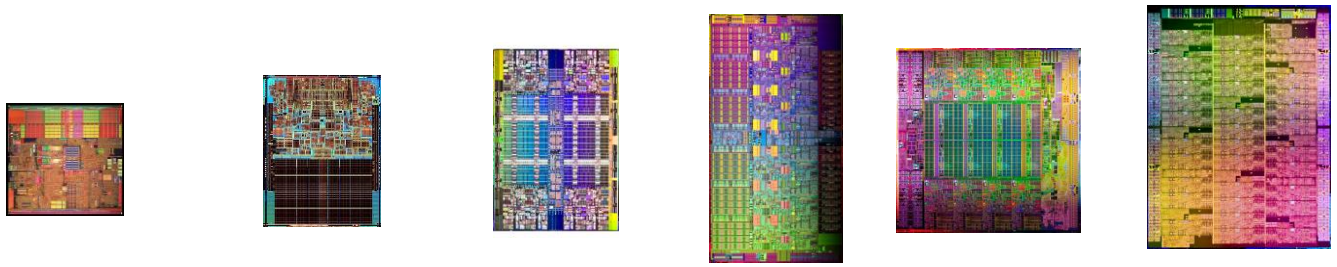

Images not intended to reflect actual die sizes

	64-bit Intel® Xeon® processor	Intel® Xeon® processor 5100 series	Intel® Xeon® processor 5500 series	Intel® Xeon® processor 5600 series	Intel® Xeon® processor E5-2600 series	Intel® Xeon Phi™ Co-processor 7120P
Frequency	3.6GHz	3.0GHz	3.2GHz	3.3GHz	2.7GHz	1238MHz
Core(s)	1	2	4	6	8	61
Thread(s)	2	2	8	12	16	244
SIMD width	128 (2 clock)	128 (1 clock)	128 (1 clock)	128 (1 clock)	256 (1 clock)	512 (1 clock)

Intel® Xeon Phi™ coprocessors extend established CPU architecture and programming concepts to highly parallel applications.

Highly Parallel Applications

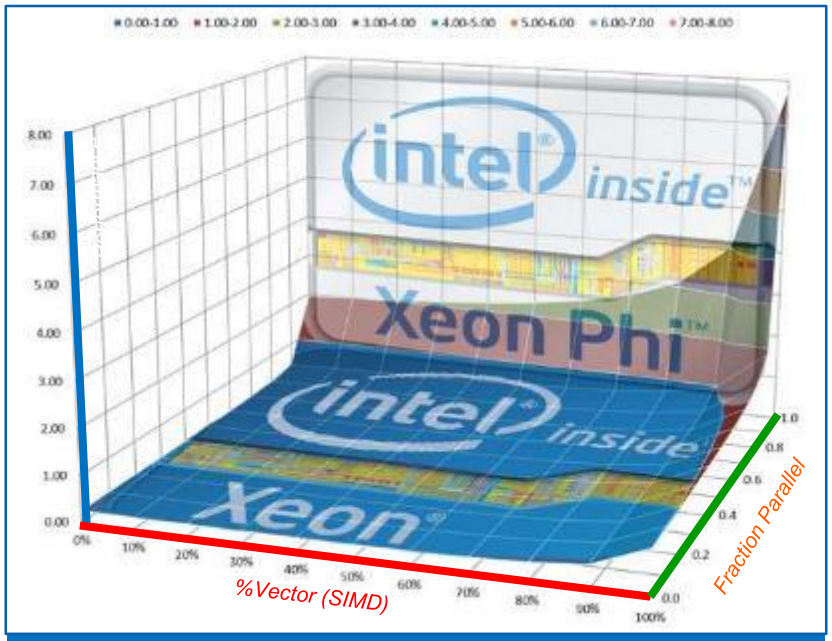

Theoretical acceleration of a highly parallel processor over a Intel® Xeon® parallel processor (<1: Intel® Xeon® faster) – For illustration only

Efficient vectorization, threading, and parallel execution drives higher performance for suitable scalable applications

Intel® Xeon Phi™ Coprocessor: Increases Application Performance up to 10x

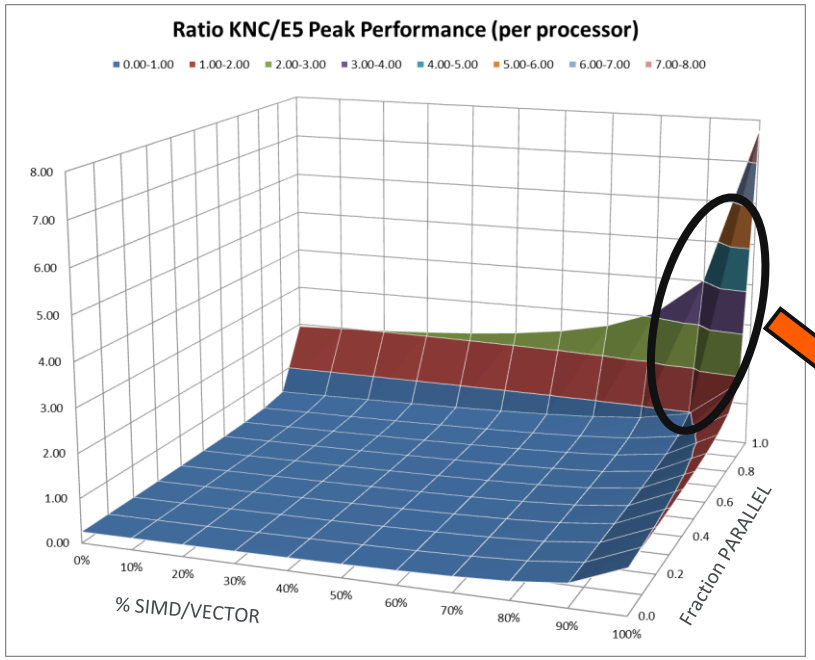

- Intel® Xeon Phi™ coprocessor accelerates highly parallel & vectorizable applications. (graph above)
- Table provides examples of such applications

Application Performance Examples

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	BlackScholes SP Monte Carlo SP	Up to 7x Up to 10.75x
Sinopec	Seismic Imaging	Up to 2.53x[2]
Sandia Labs	miniFE (Finite Element Solver)	Up to 2x[3]
Intel Labs	Ray Tracing (incoherent rays)	Up to 1.88x[4]

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

Notes:

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. 8 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

Intel® Many Integrated Core Architecture

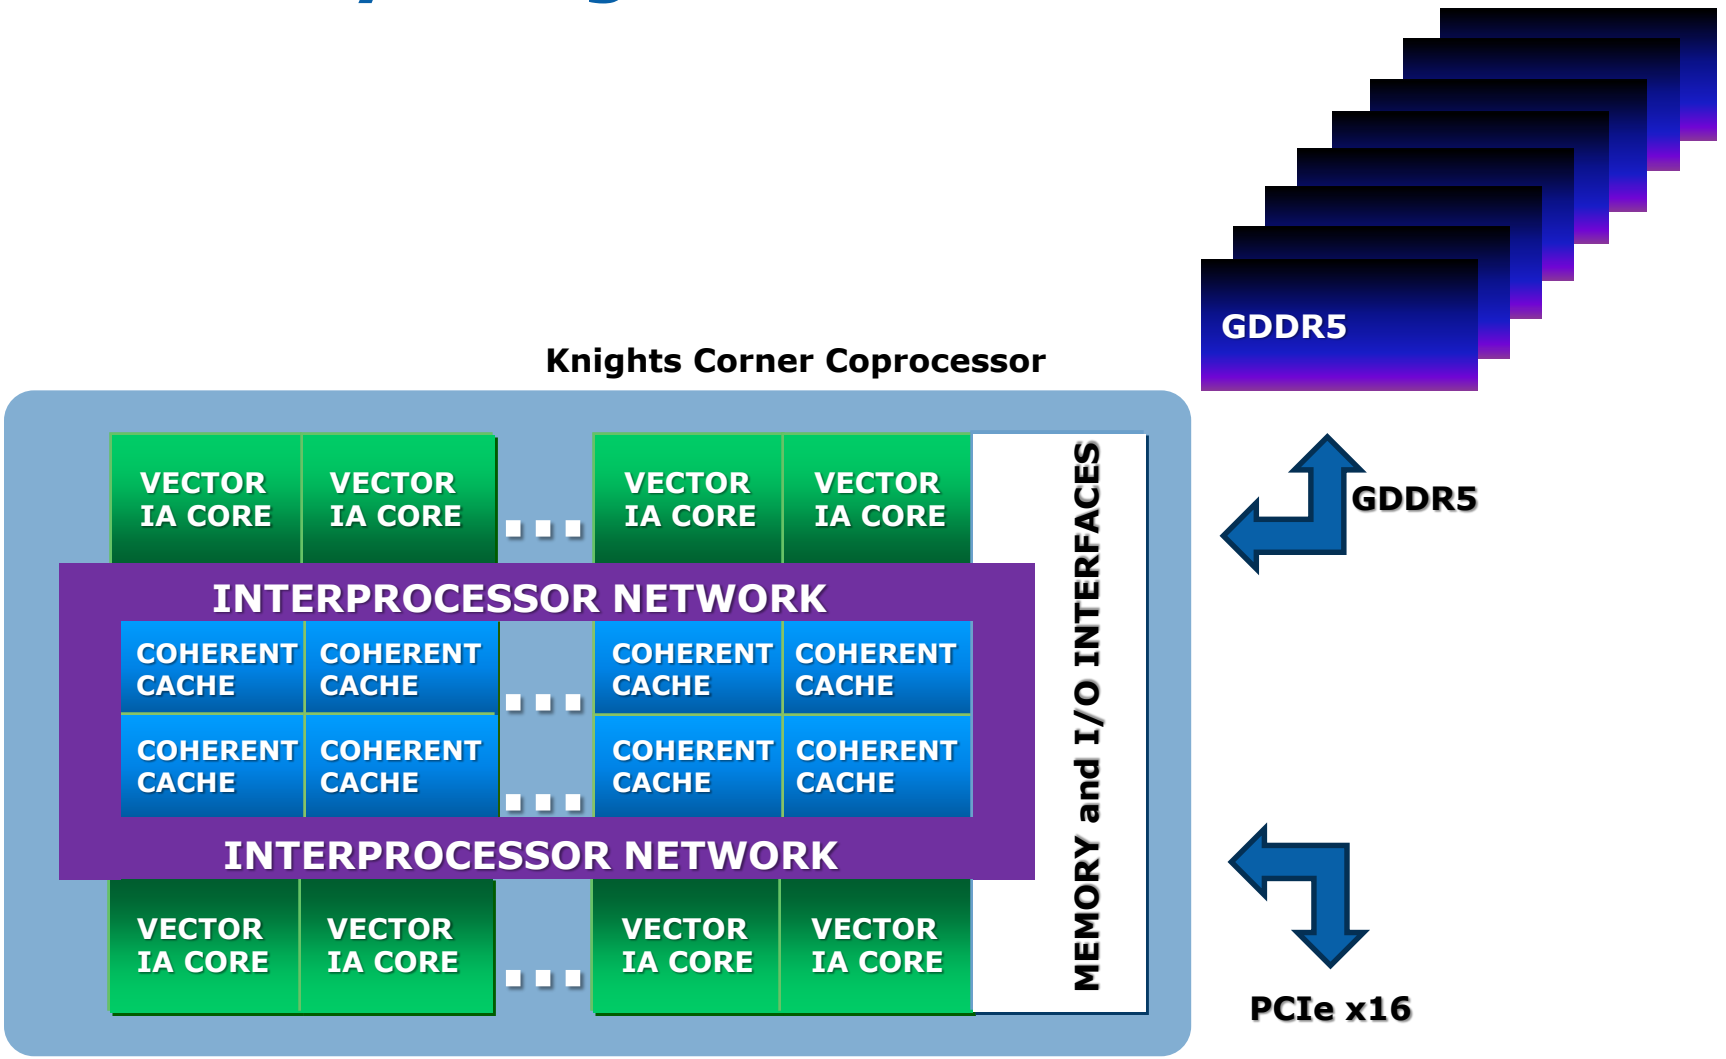

Intel® Many Integrated Core Architecture

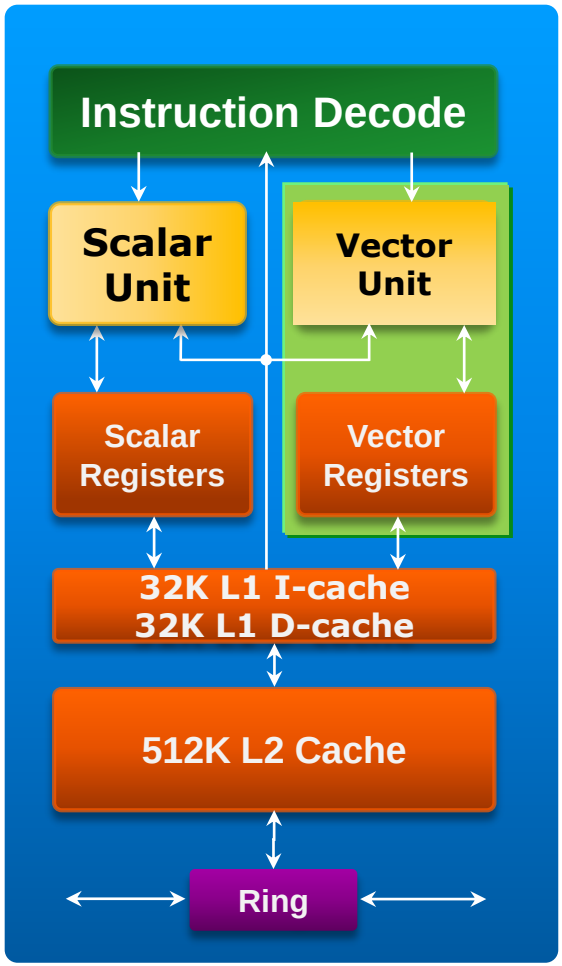

- 60 in-order cores
- 4 hardware threads per core
- Two pipelines
  - Pentium® processor family-based scalar units
  - Fully-coherent L1 and L2 caches
  - 64-bit addressing
- All new vector unit
  - 512-bit SIMD Instructions (not Intel® SSE, MMX™, or Intel® AVX)
  - 32 512-bit wide vector registers
    - Hold 16 singles or 8 doubles per register
  - Pipelined one-per-clock throughput
    - 4 clock latency, hidden by round-robin scheduling of threads
  - Dual issue with scalar instructions

Architecture of an Intel® Xeon Phi™ Coprocessor

- Cache
  - 32 KB L1 / 512 KB L2 per core
  - Fully coherent
- Core Communication
  - Bi-directional ring buffer
  - Up to 16 GB GDDR5 shared by all cores
- PCIe*
  - Gen2
  - 16 channels

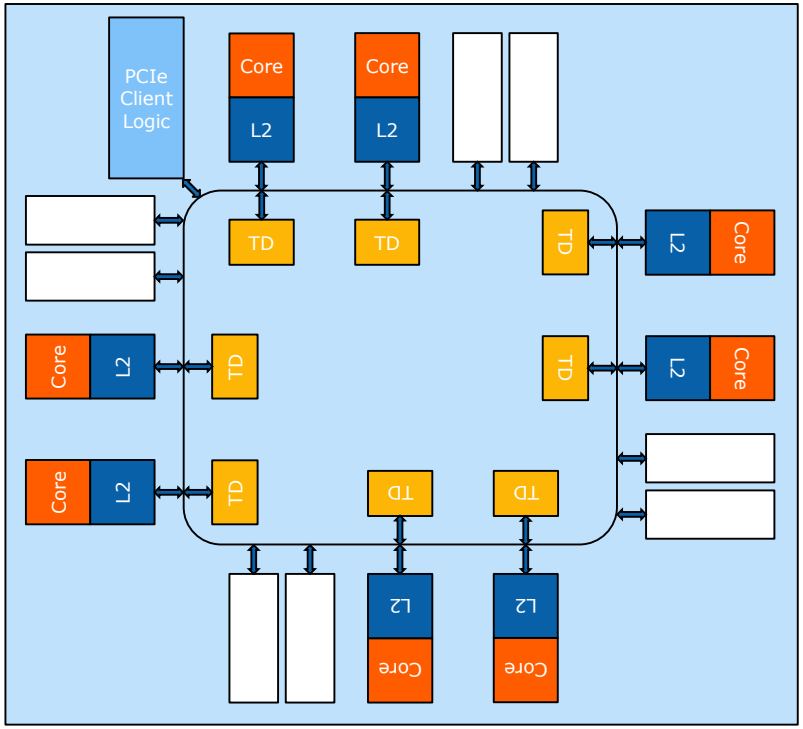

Intel MIC Architecture Overview
Vector Processing Unit and ISA

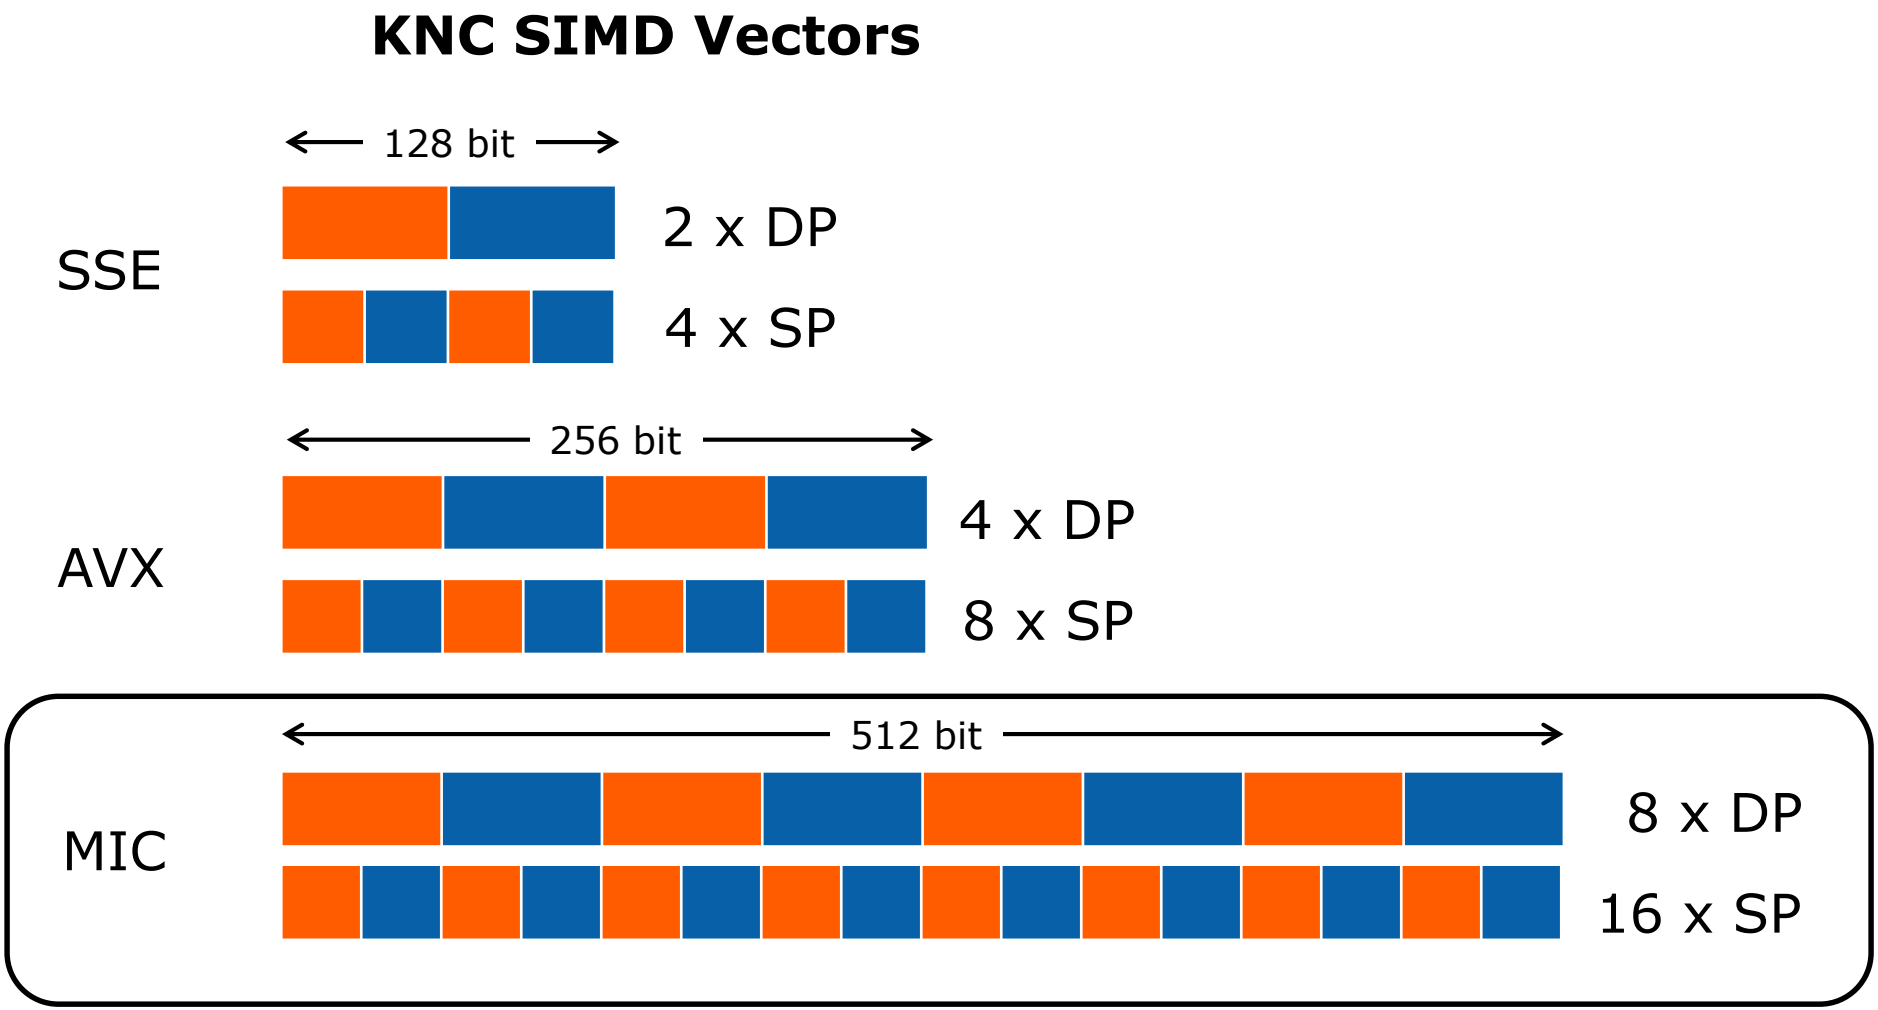

Intel MIC Architecture Overview
Vector Processing Unit and ISA

KNC SIMD Vectors Basic Arithmetic

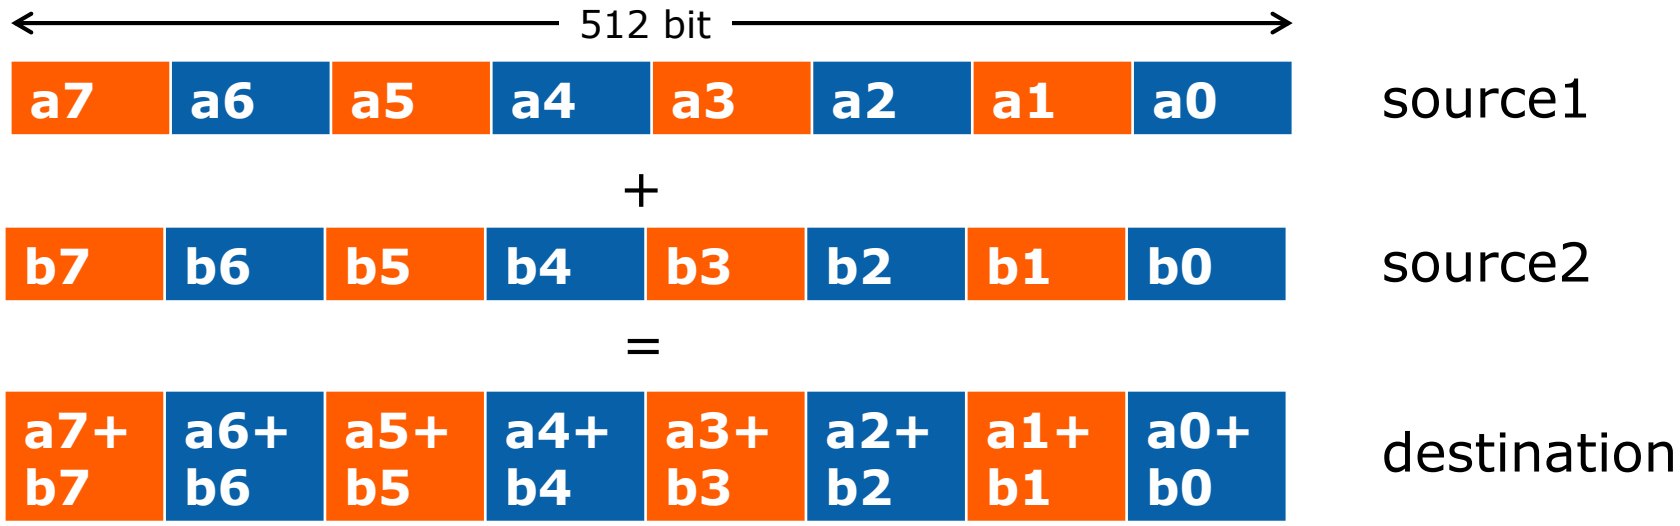

Basic arithmetic SIMD instruction usage is trivial and identical to SSE or AVX.

vaddps, vsubps, vmulps, …
vaddpd, vsubpd, vmulpd, …

Intel MIC Architecture Overview
Vector Processing Unit and ISA

KNC SIMD Fused Multiply and Add/Subtract

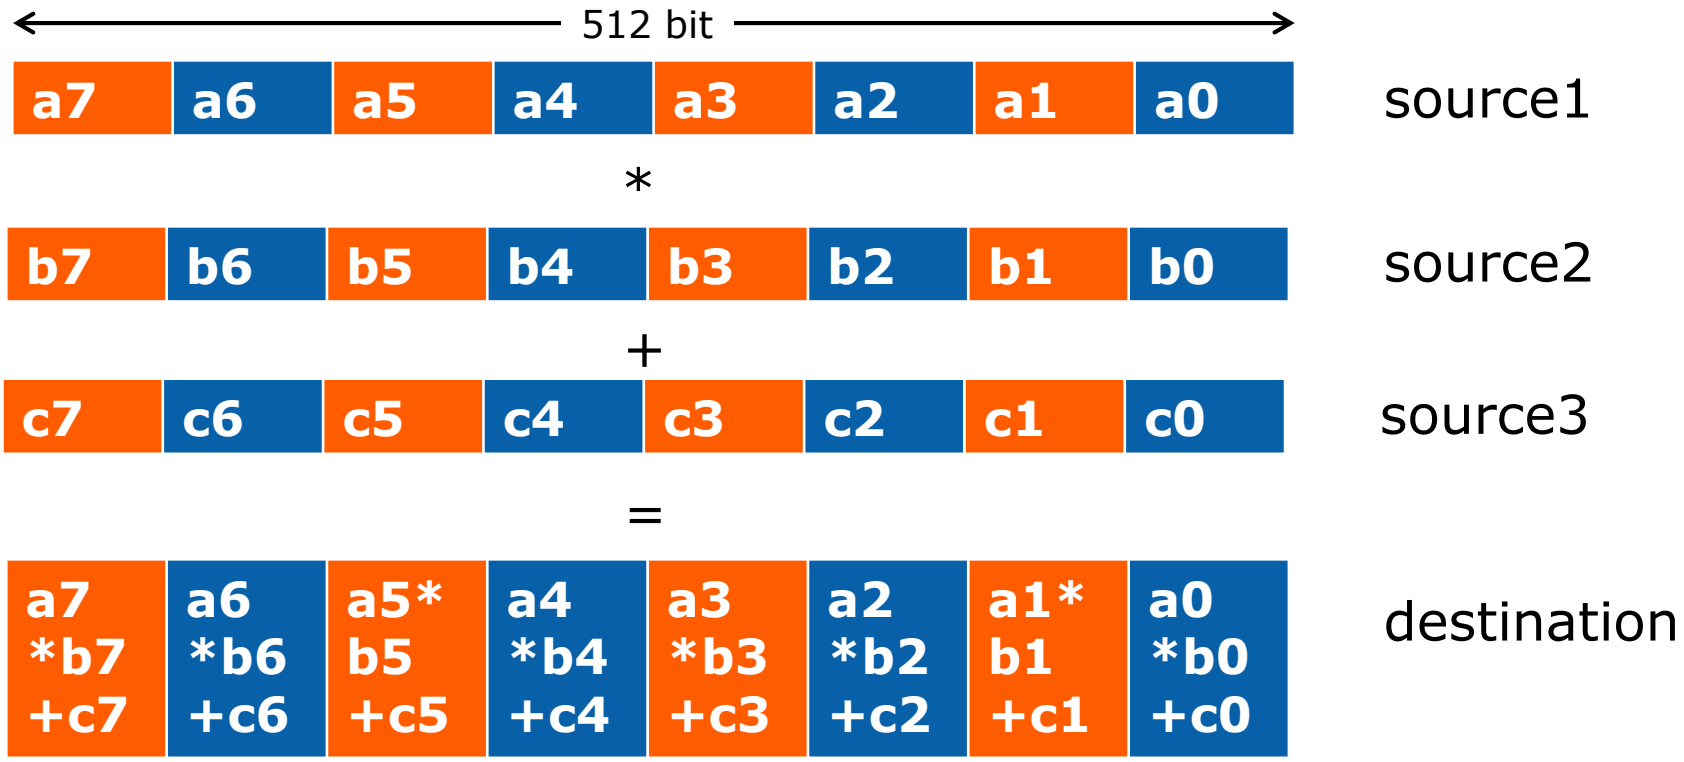

vfmadd213ps source1,source2,source3

Intel MIC Architecture Overview
Vector Processing Unit and ISA

KNC SIMD Vectors Masking

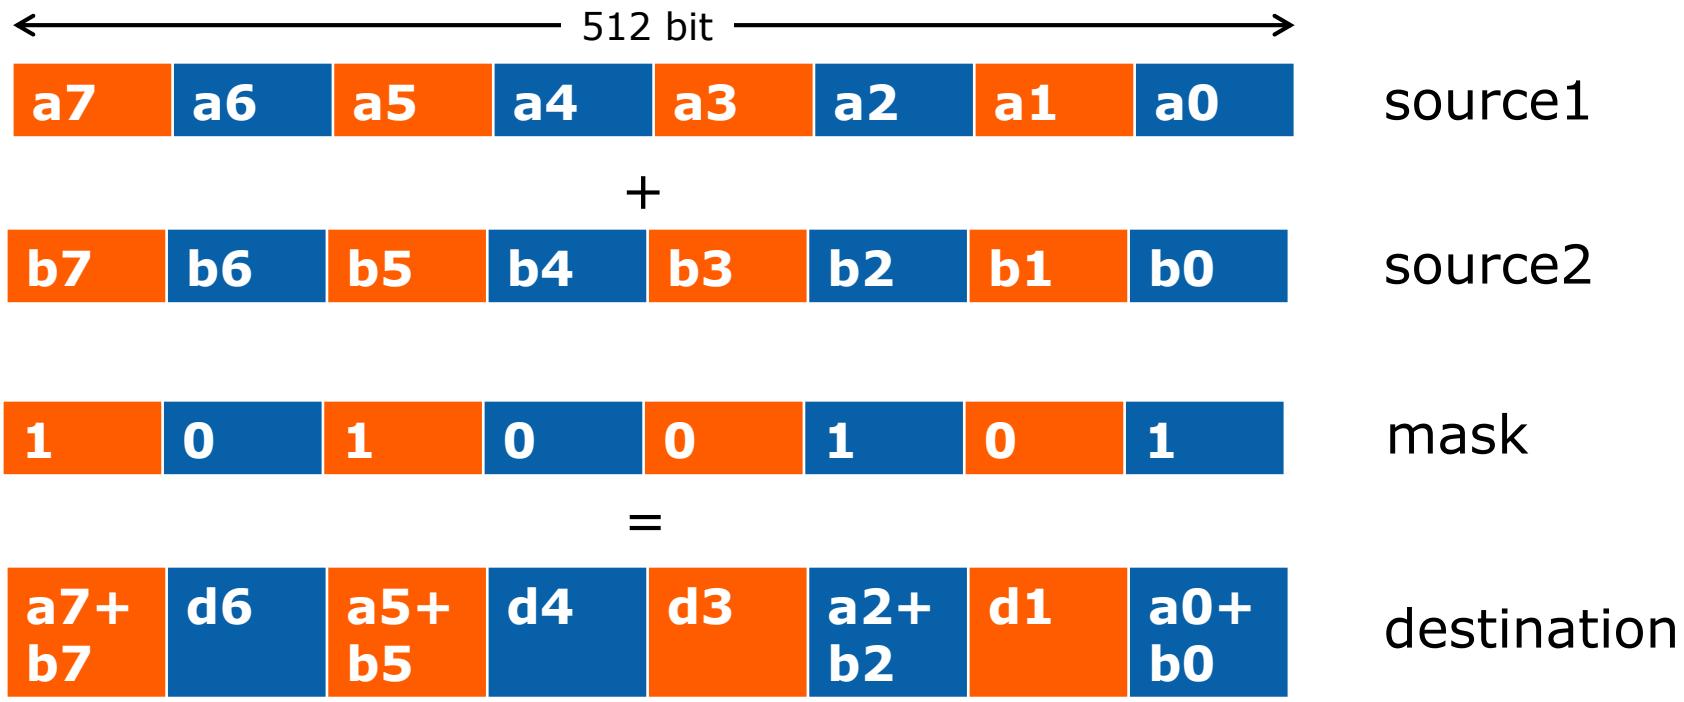

vaddps zmm0{k1}, zmm1, zmm2

Masking allows non-destructive writing to the destination (unlike AVX). Every Knight's Corner instruction has write masking.

Intel MIC Architecture Overview
Vector Processing Unit and ISA

KNC SIMD Vector Swizzling

Swizzling is the modification of the last source. One can easily envision it as creating a modified copy for the following operation.

Example:

vmovapd zmm1, zmm0{dacb}

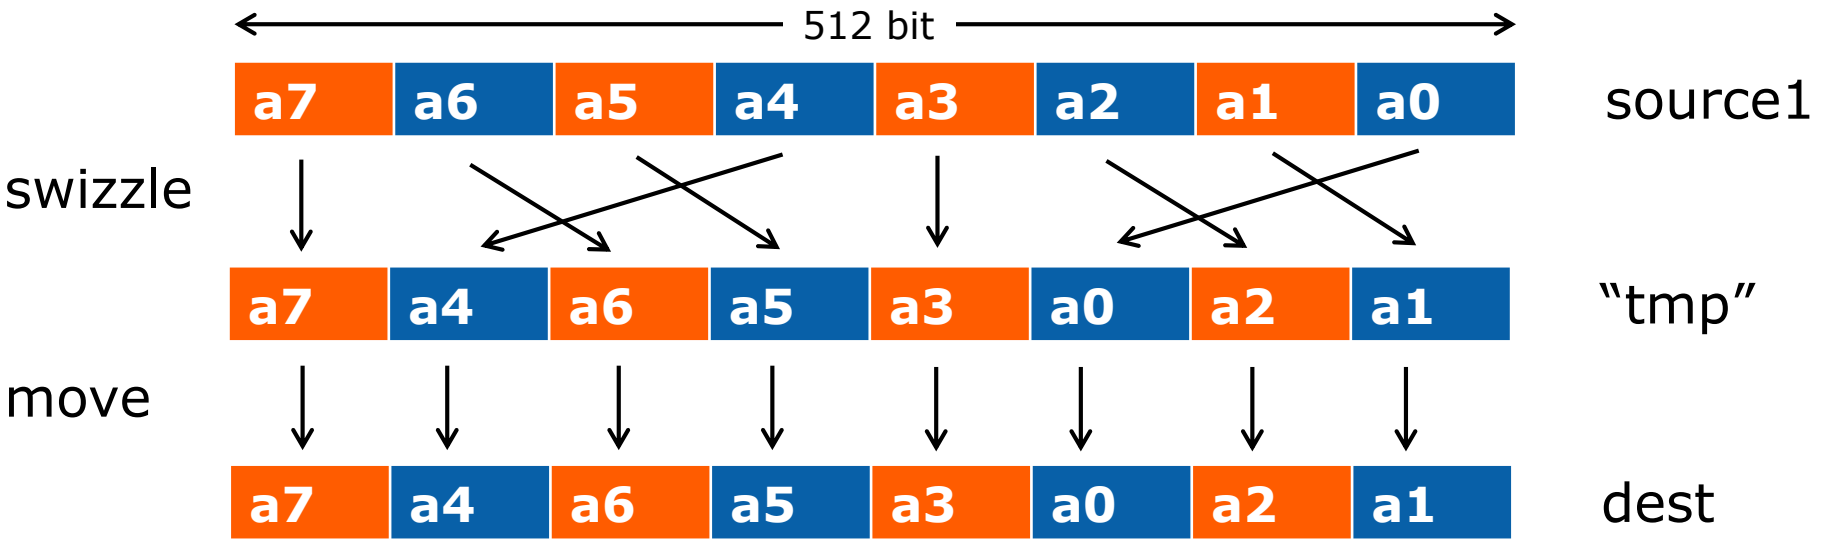

Knights Corner Architecture Overview
Software Architecture

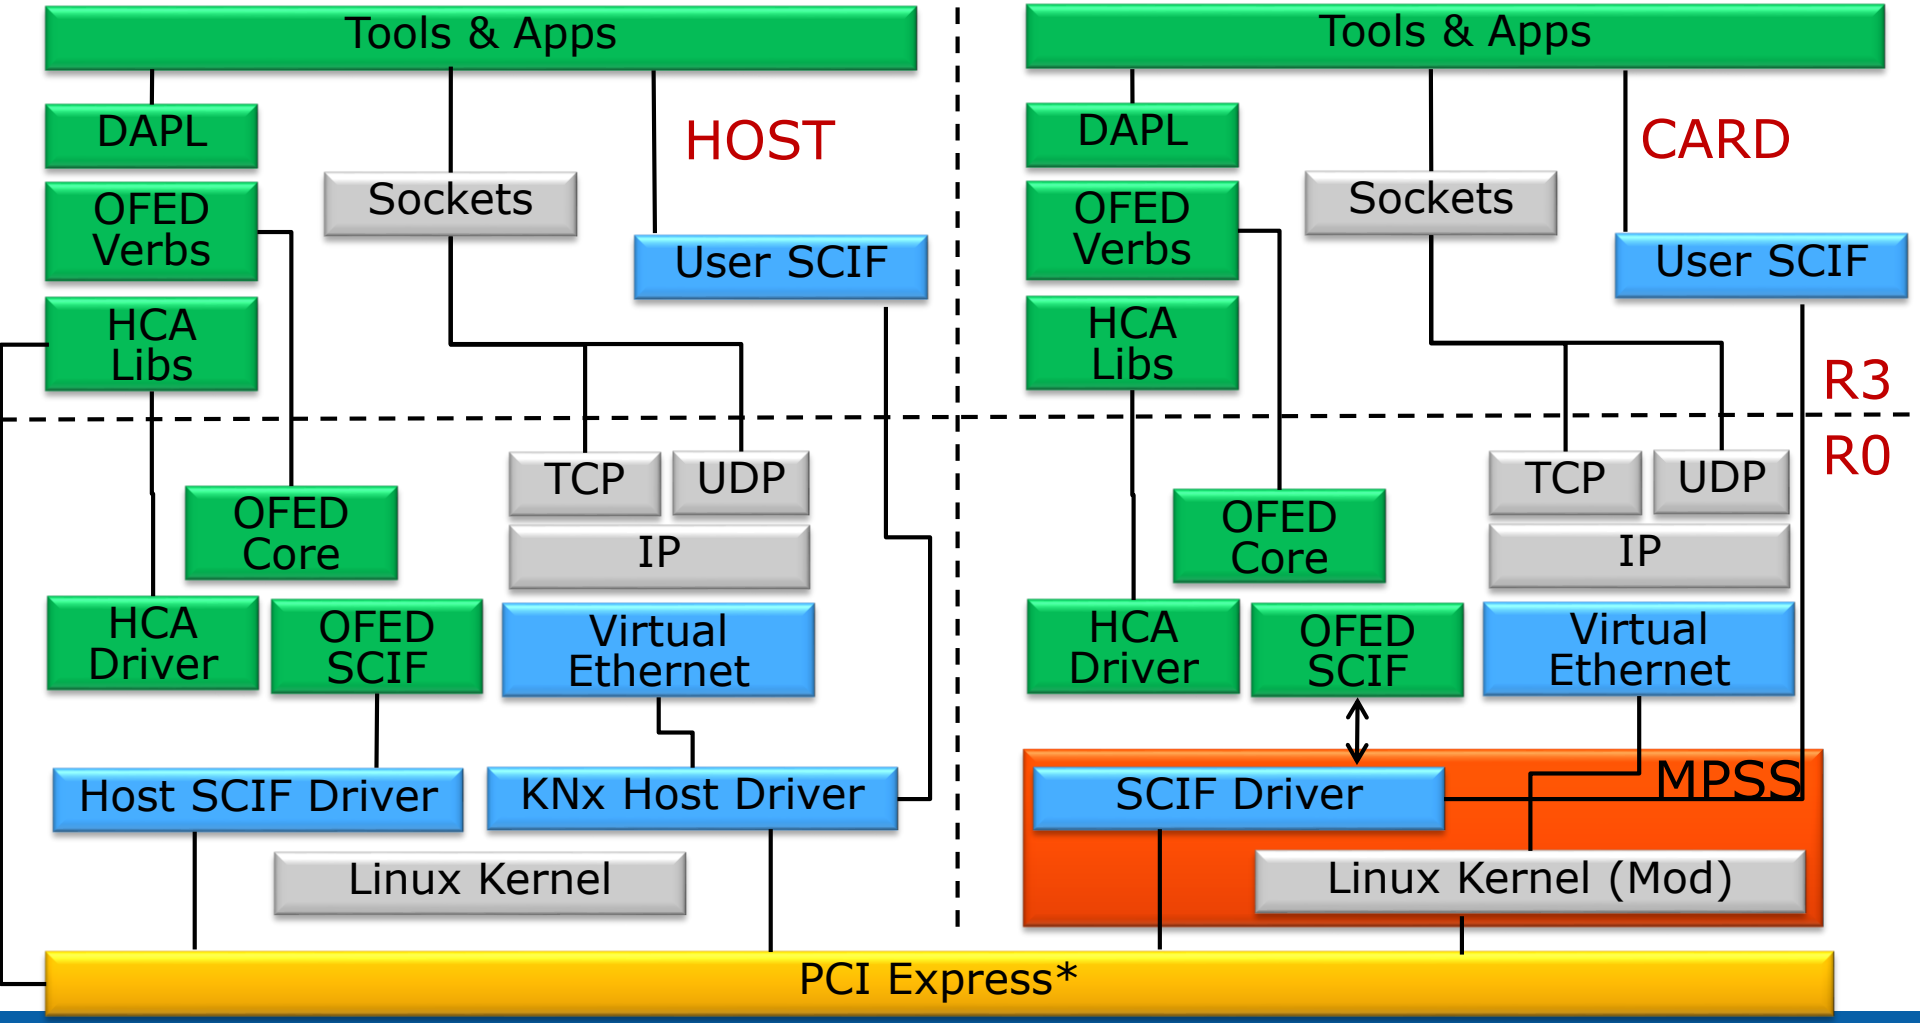

Spectrum of Programming Models

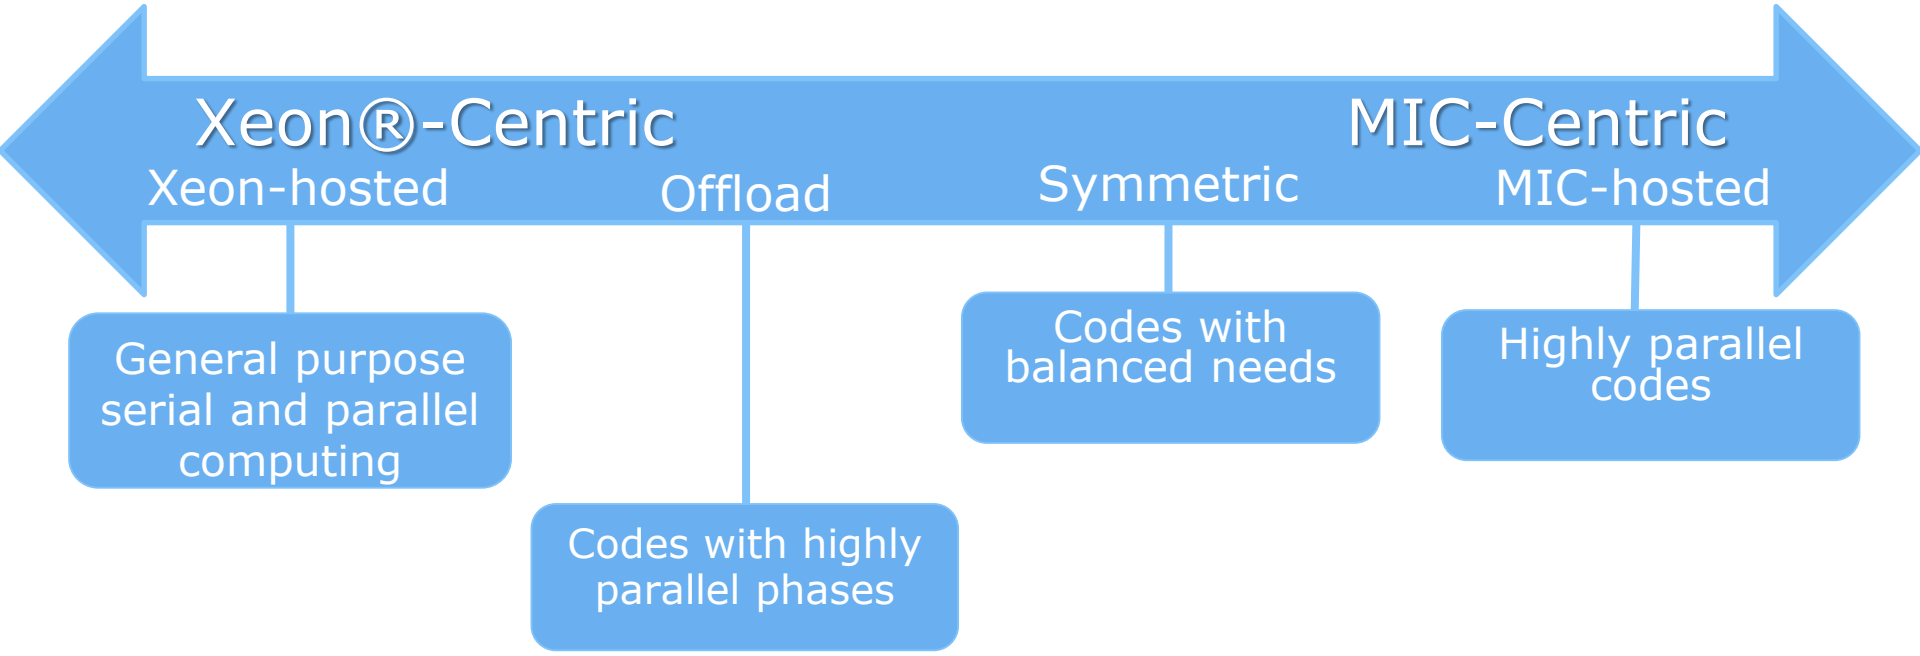

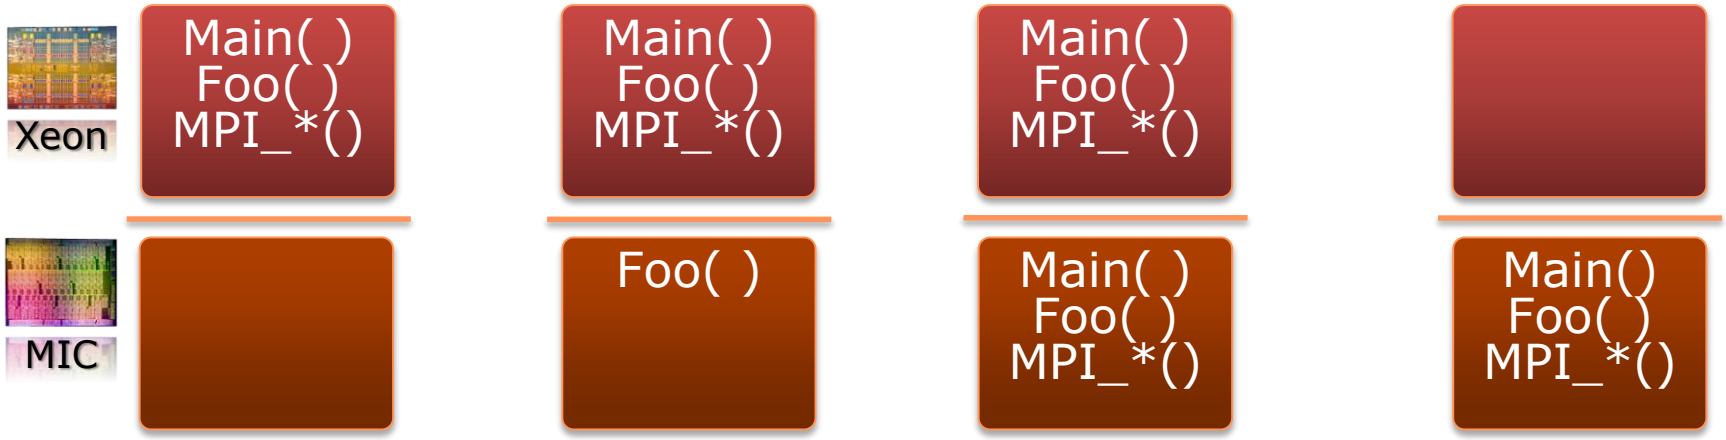

Heterogeneous Programming

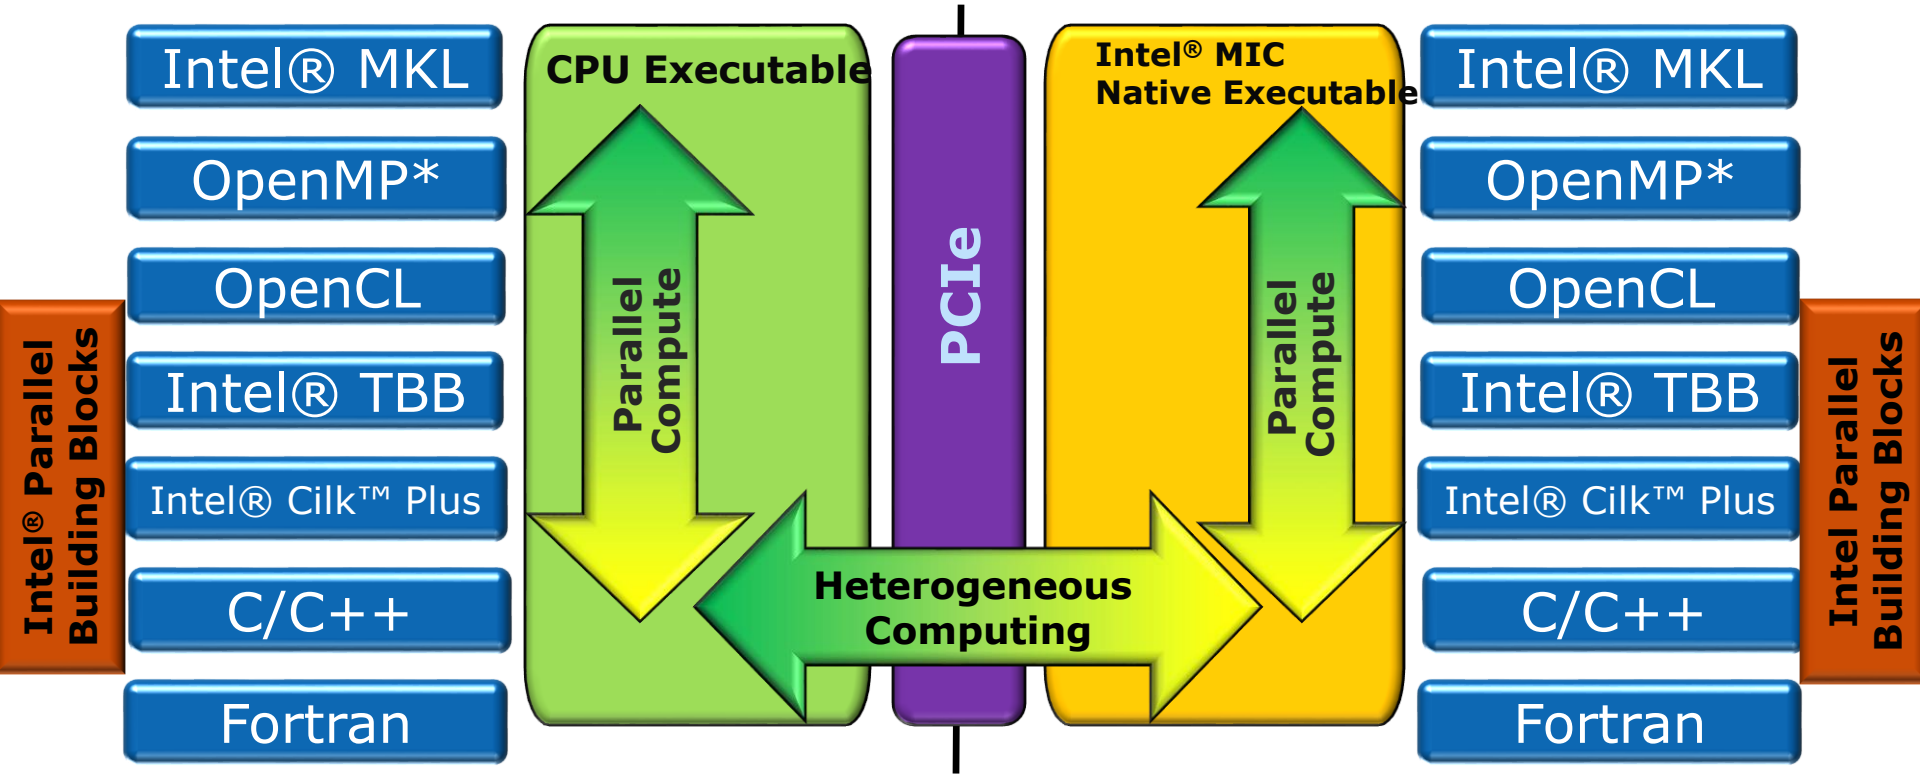

Native Programming for Intel Xeon Phi

IA benefit: your code, your choice!

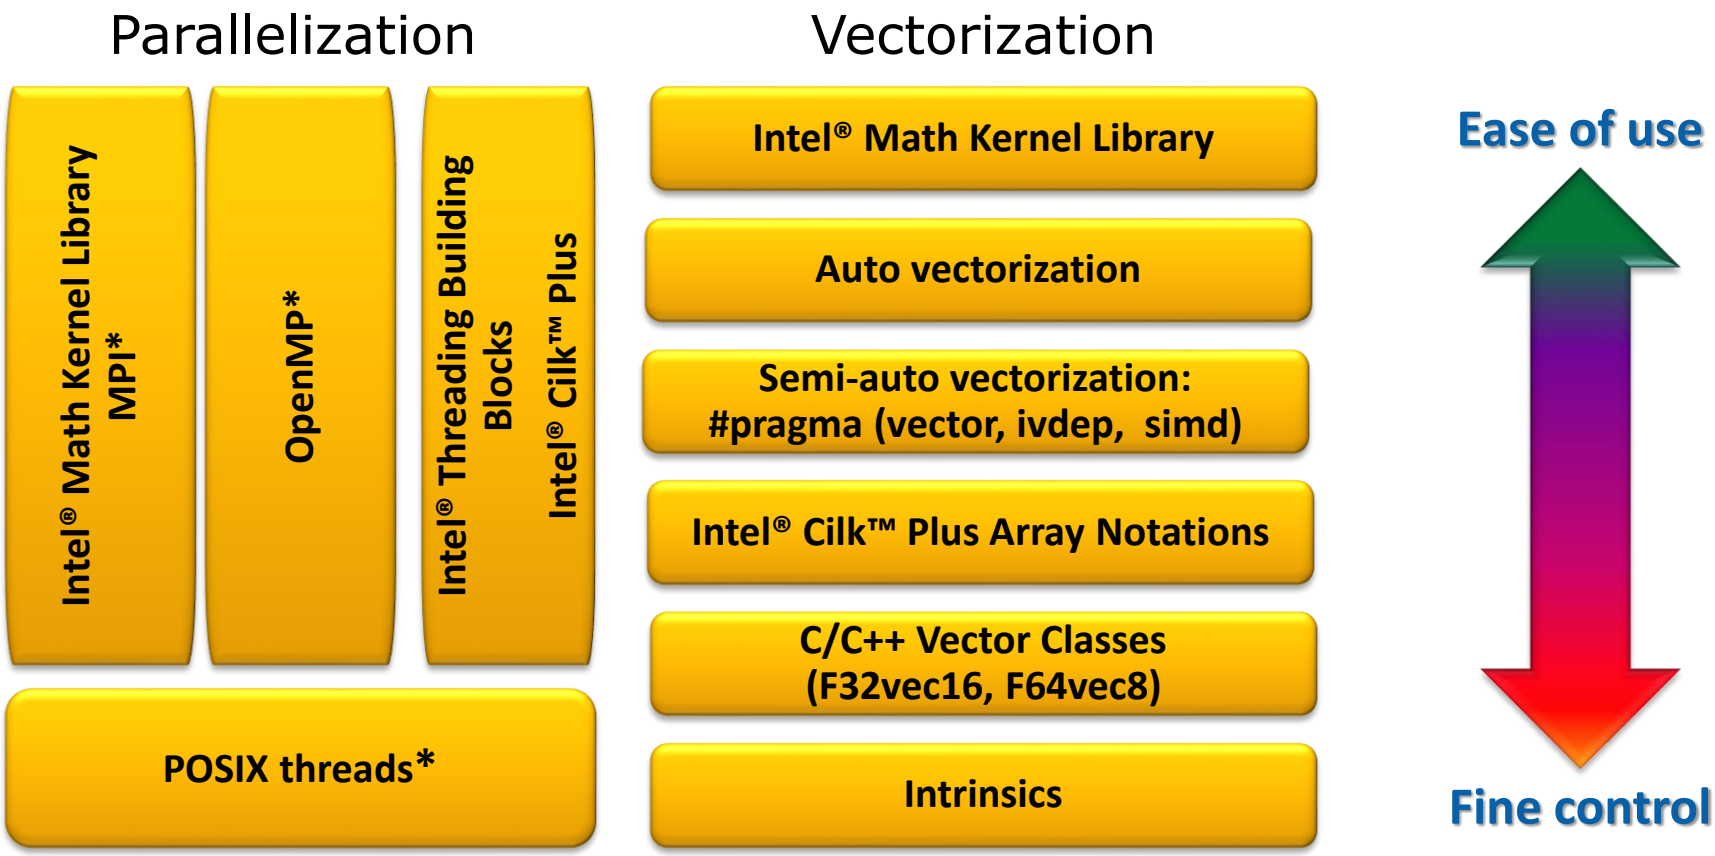

Intel® Xeon Phi™ Coprocessors Capabilities

Restrictive architectures

It's a supercomputer on a chip

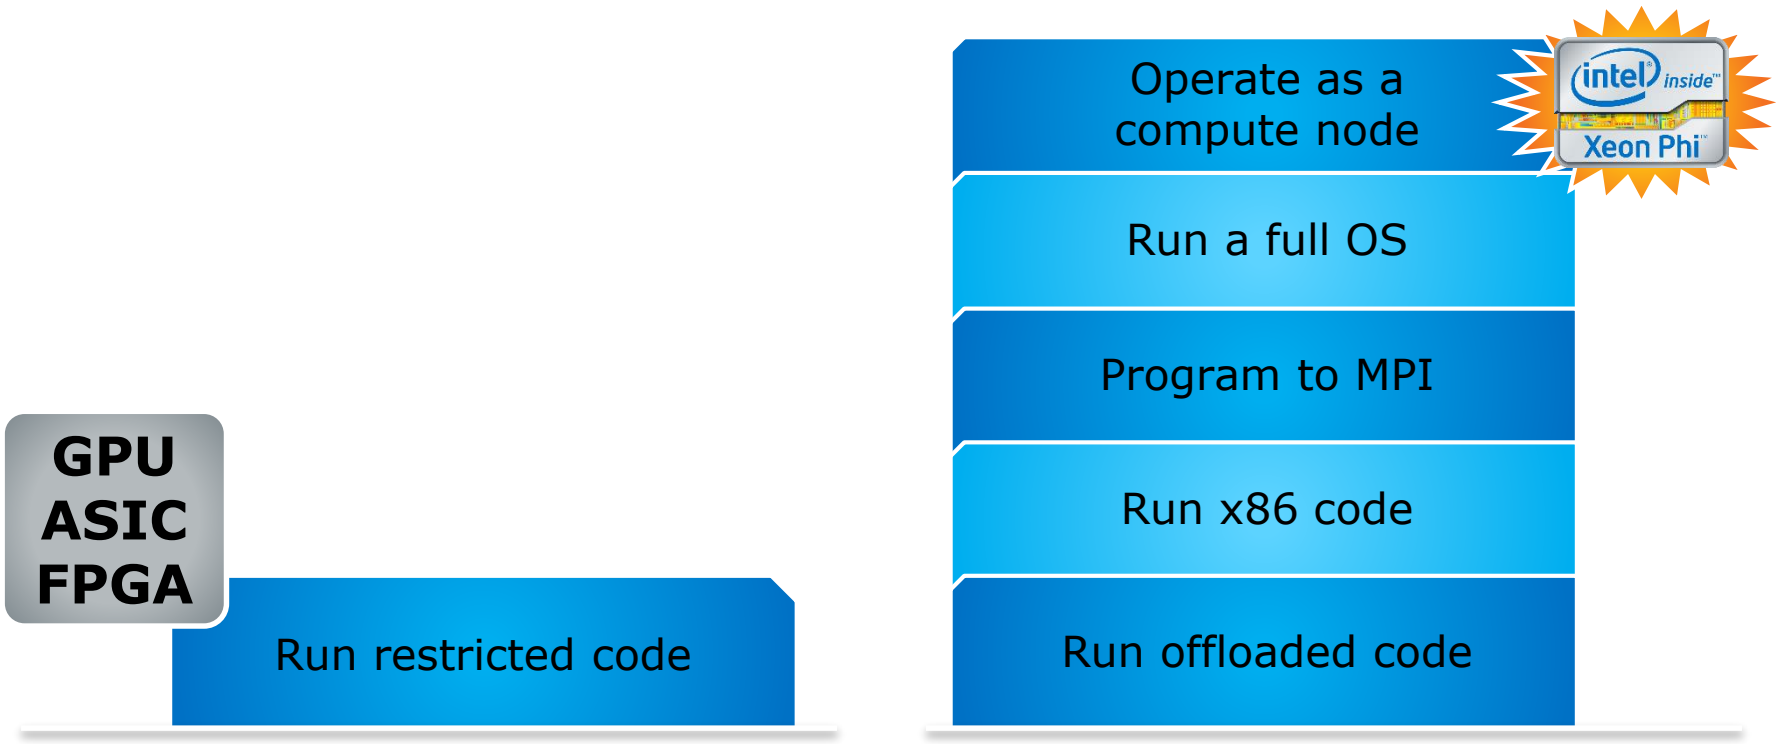

Custom HW Acceleration

Intel® Xeon Phi™ Coprocessor

Restrictive architectures limit the ability for applications to use arbitrary nested parallelism, functions calls and threading models