From Milliwatts to PFLOPS

GridKa School 2014

Karlsruhe, Germany
02 September 2014

Dr.-Ing. Michael Klemm
Intel

Legal Disclaimer

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.

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 infully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Intel product plans in this presentation do not constitute Intel plan of record product roadmaps. Please contact your Intel representative to obtain Intel's current plan of record product roadmaps.

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not uniqueto Intel microprocessors. These optimizations include SSE2, SSE3, and SSE3 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 specificinstruction sets covered by this notice.Notice revision #20110804

All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

Intel processor numbers are not a measure of performance.Processor numbers differentiate features within each processor family, not across different processor families.Go to: http://www.intel.com/products/processor_number

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.

Intel, Intel Xeon, Intel Xeon Phi, Intel Hadoop Distribution, Intel Cluster Ready, Intel OpenMP, Intel CilkPlus, Intel Threaded Buildiingblocks, Intel Cluster Studio, Intel Parallel Studio, Intel CoarrayFortran, Intel Math KernalLibrary, Intel Enterprise Edition for LustreSoftware, Intel Composer, the Intel Xeon Phi logo, the Intel Xeon logo and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

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

*Other brands and names are the property of their respective owners

Legal Disclaimers: Performance

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, Go to: http://www.intel.com/performance/resources/benchmark_limitations.htm.

Intel does not control or audit the design or implementation of third party benchmarks or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmarks are reported and confirm whether the referenced benchmarks are accurate and reflect performance of systems available for purchase.

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.

SPEC, SPECint, SPECfp, SPECrate. SPECpower, SPECjAppServer, SPECjEnterprise, SPECjbb, SPECompM, SPECompL, and SPEC MPI are trademarks of the Standard Performance Evaluation Corporation. See http://www.spec.org for more information.

TPC Benchmark is a trademark of the Transaction Processing Council. See http://www.tpc.org for more information.

SAP and SAP NetWeaver are the registered trademarks of SAP AG in Germany and in several other countries. See http://www.sap.com/benchmark for more information.

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.

*Other brands and names are the property of their respective owners

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

The Path to Discovery & Innovation

EXPERIMENT
Observation

THEORY
Mathematical Model

$$\left(\frac{\partial u_r}{\partial \theta} + \frac{u_\theta}{r}\frac{\partial u_r}{\partial \phi} - \frac{u_\theta^2 + u_\phi^2}{r}\right) = -\frac{\partial p}{\partial r} + \rho g_r$$

$$\frac{1}{\sin(\phi)^2}\frac{\partial^2 u_r}{\partial \theta^2} + \frac{1}{r^2\sin(\phi)}\frac{\partial}{\partial \phi}\left(\sin(\phi)\frac{\partial u_r}{\partial \phi}\right) - 2\frac{u_r + \frac{\partial u_\theta}{\partial \theta} + u_\phi}{r^2}$$

$$\left(\frac{\partial u_\theta}{\partial \theta} + \frac{u_\theta}{r}\frac{\partial u_\theta}{\partial \phi} + \frac{u_r u_\theta + u_\theta u_\phi \cot(\phi)}{r}\right) = -\frac{1}{r\sin(\phi)}\frac{\partial p}{\partial \theta} + \rho g_\theta$$

$$\frac{1}{\sin(\phi)^2}\frac{\partial^2 u_\theta}{\partial \theta^2} + \frac{1}{r^2\sin(\phi)}\frac{\partial}{\partial \phi}\left(\sin(\phi)\frac{\partial u_\theta}{\partial \phi}\right) + \frac{2\frac{\partial u_r}{\partial \theta} + 2\cos(\phi)}{r^2\sin(\phi)}$$

$$\left(\frac{\partial u_\phi}{\partial \theta} + \frac{u_\phi}{r}\frac{\partial u_\phi}{\partial \phi} + \frac{u_r u_\phi - u_\theta^2 \cot(\phi)}{r}\right) = -\frac{1}{r}\frac{\partial p}{\partial \phi} + \rho g_\phi$$

$$\frac{1}{\sin(\phi)^2}\frac{\partial^2 u_\phi}{\partial \theta^2} + \frac{1}{r^2\sin(\phi)}\frac{\partial}{\partial \phi}\left(\sin(\phi)\frac{\partial u_\phi}{\partial \phi}\right) + \frac{2}{r^2}\frac{\partial u_r}{\partial \phi} - \frac{u_\phi +}{r^2}$$

HPC
Numerical Simulation

Performance

it's all about Parallelism

and

Energy Efficiency

HPC IMPERATIVES

High Performance
Capabilities & Capacity

Energy Efficiency
TCO

Ease of Use
Productivity & Sustainability

Intel's Assets for HPC

Processors
Intel® Xeon® Processor

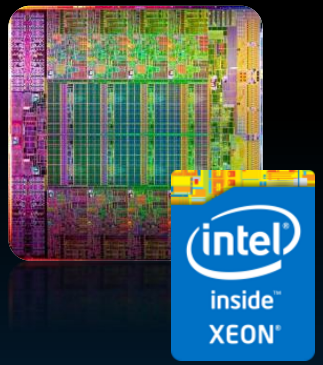

Coprocessors
Intel® Many Integrated Core

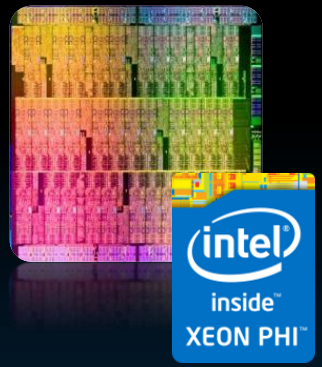

Network & Fabrics

Storage

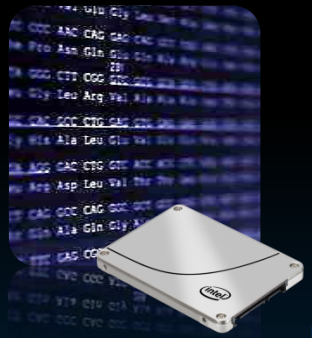

Software & Services

… and many, many application experts

*Other brands and names are the property of their respective owners

Simplicity
is the ultimate sophistication.

- Leonardo da Vinci

Transforming the Economics of HPC

Executing to Moore's Law

Predictable Silicon Track Record – well and alive at Intel.
Enabling new devices with higher performance and functionality while controlling power, cost, and size

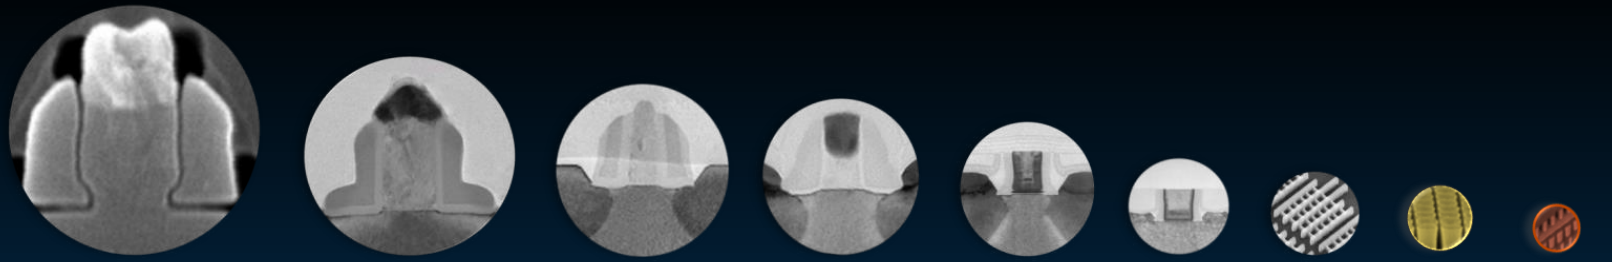

Future options subject to change without notice.

Tick-Tock Development Cycles

Integrate. Innovate.

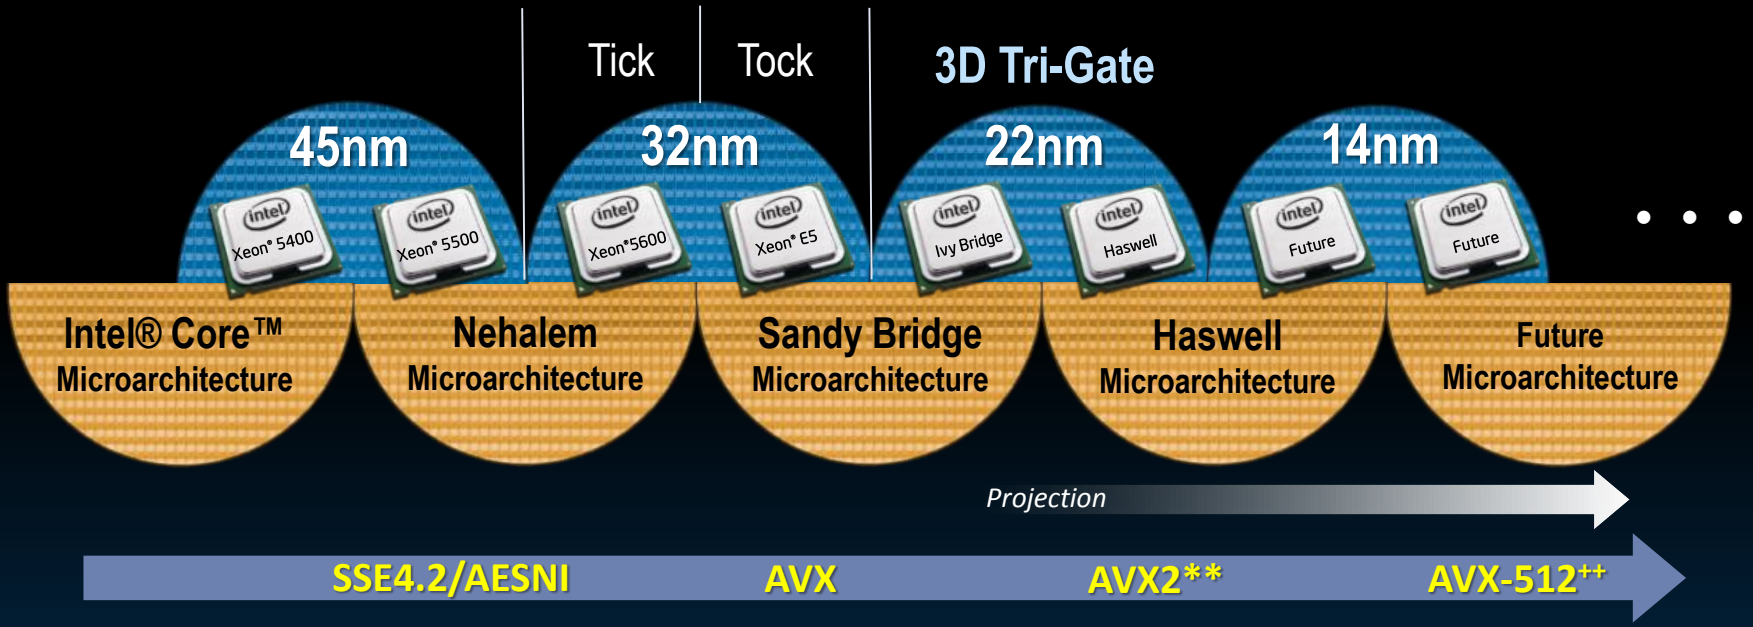

**Intel® Architecture Instruction Set Extensions Programming Reference, #319433-012A, FEBRUARY 2012

++Intel® Architecture Instruction Set Extensions Programming Reference, #319433-015, JULY 2013

Potential future options, subject to change without notice.

From MILLIWATTS to TERAFLOPS

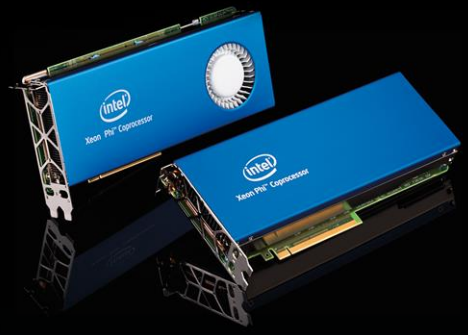

Smartphones with Intel® Inside

Intel® Xeon® Processors

Intel® Many Integrated Core Architecture

Energy Efficient

Driving Innovation and Integration

Enabled by Leading Edge Process Technologies

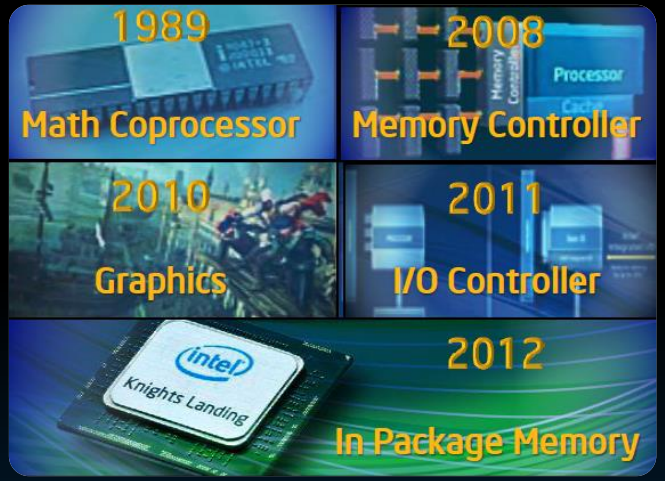

Integrated Today

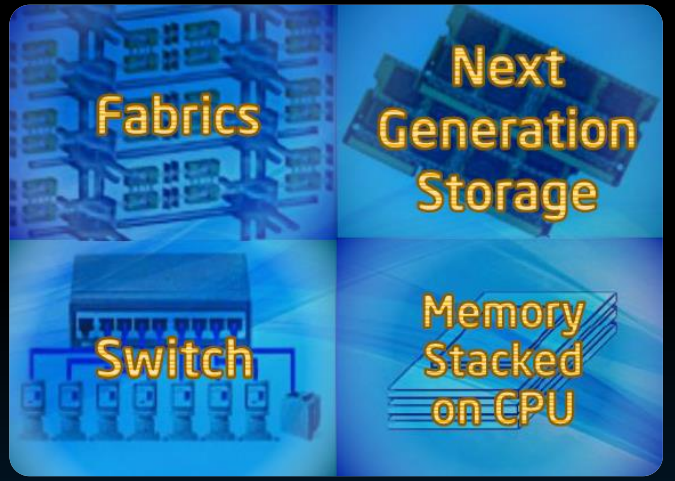

Coming in the Future

SYSTEM LEVEL BENEFITS IN COST, POWER, DENSITY, SCALABILITY & PERFORMANCE

The Magic of Integration

Moore's Law at Work & Architecture Innovations

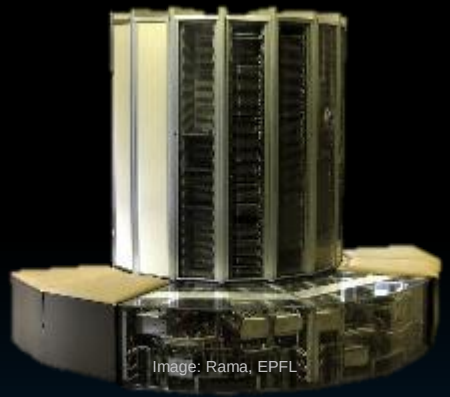

6666x

1970s

150 MFLOPS

CRAY-1

2013

1000000 MFLOPS

Intel® Xeon Phi™

#1 TOP500 June 2014

33 PFLOPS HPL

54 PFLOPS Peak

32000 Intel® Xeon® E5v2 Processors

48000 Intel® Xeon Phi™ Coprocessors

Increasing Processor Performance

Through Many-Core Technologies for Highly Parallel Workloads

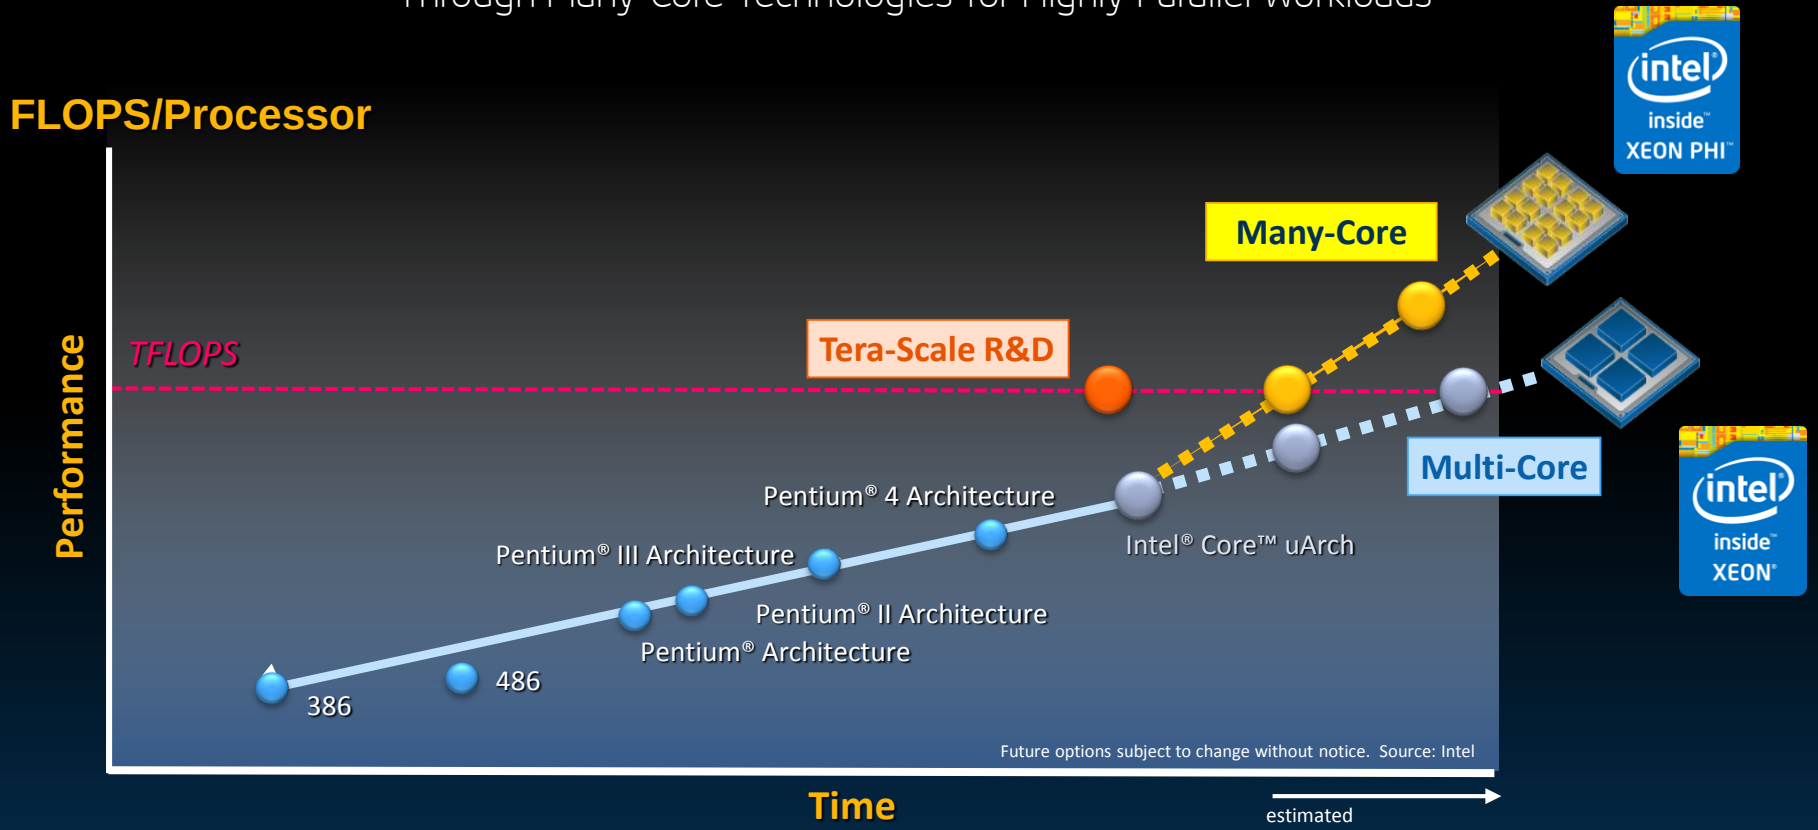

For illustration only. All dates, product descriptions, features, availability, and plans are forecasts and subject to change without notice.

„Big Core“ – „Small Core“

Different Optimization Points
Common Programming Models
and Architectural Elements

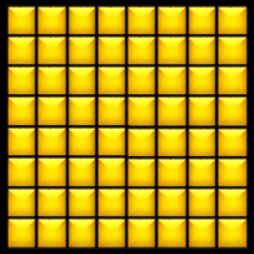

Intel® Xeon® Processor

Simply aggregating more cores generation after generation is not sufficient

Performance per core/thread must increase each generation, be as fast as possible

Power envelopes should stay flat or go down each generation

Balanced platform (Memory, I/O, Compute)

Cores, Threads, Caches, SIMD

Intel® Xeon Phi™ Coprocessor

Optimized for highest compute per watt

Willing to trade performance per core/thread for aggregate performance

Power envelopes should also stay flat or go down every generation

Optimized for highly parallel workloads

Cores, Threads, Caches, SIMD

For illustration only

Intel SIMD Evolution

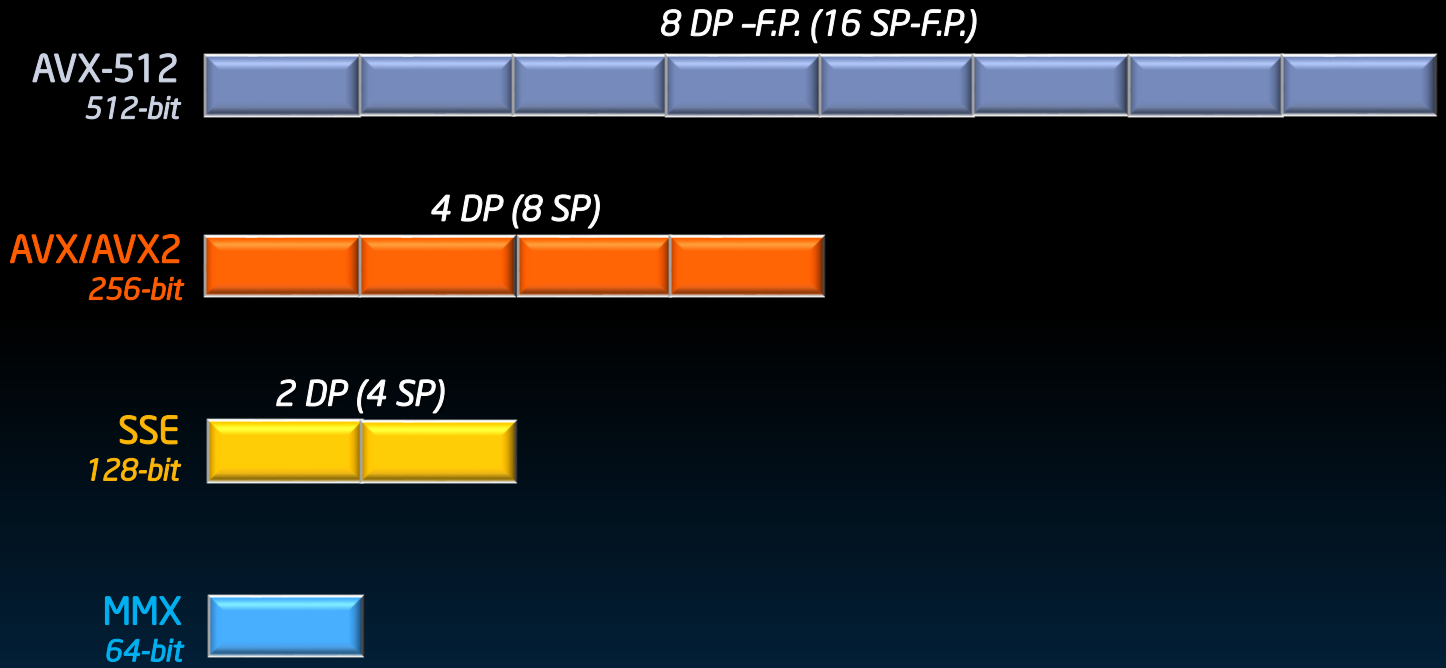

Potential future options and features subject to change without notice.

Intel Roadmap to Exascale

1.000.000.000.000.000

Intel's Exascale Goal:
Reach Exascale by ~2020 with Intel technologies including Intel® Xeon Phi™ Coprocessors

Intel® Xeon Phi™ Product Family
Key ingredient in Intel Exascale Roadmap:

- Programmability
- Power efficiency
- Scalability
- Resiliency

Future options subject to change without notice.

Common Programming Models & Software Tools

Common Intel® architecture enables applications to run across the full spectrum of Intel® Xeon® family based servers so programmers don't have to "start over".

Use the same development tools you used for Intel® Xeon® processors, such as Intel® Cluster Studio XE and Intel® Parallel Studio XE

*Other brands and names are the property of their respective owners

Intel® Xeon® E5 Processor Family

Foundation of HPC Performance
suited for full scope of workloads

Industry leading performance and
performance/watt
for serial & parallel workloads

General purpose with focus on
fast single core/thread performance
with “moderate” number of cores

www.intel.com/xeon

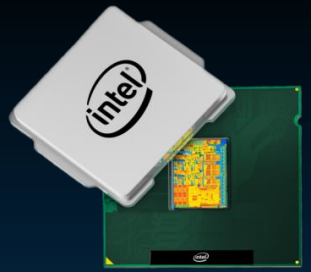

www.intel.com/xeonphi

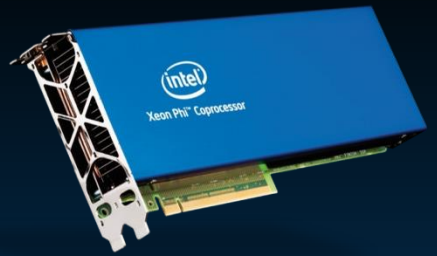

Intel® Xeon Phi™ Coprocessor

Up to 61 Cores, 244 Threads
512-bit SIMD instructions
>1TFLOPS DP-F.P. peak
Up to 16GB GDDR5 Memory, 352 GB/s
PCIe* x16
Up to 300W TDP (card)

22nm with the world's first
3-D Tri-Gate transistors
Linux* operating system
IP addressable native node
Common x86/IA
Programming Models and SW-Tools

Intel® Xeon Phi™ Coprocessor Microarchitecture Overview

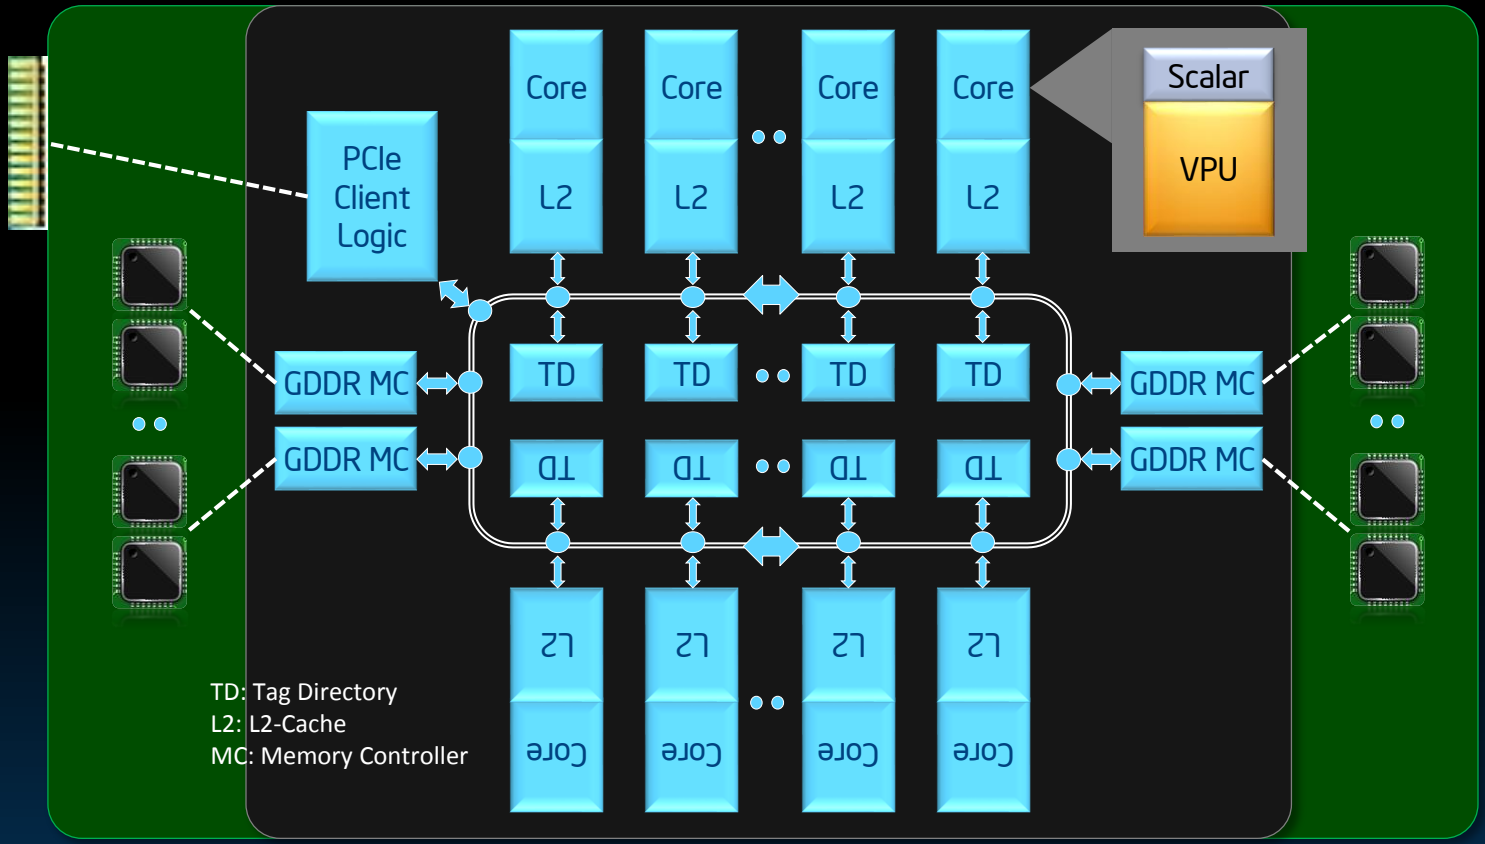

For illustration only.

Intel® Xeon Phi™ Coprocessor

Codename: Knights Corner - It is so much more

Restricted Architectures

DSP
FPGA
GPU
ASIC

Run restricted code

Custom HW Acceleration

Supercomputer on a chip

Operate as a compute node
Run a full OS (Linux*)
Run MPI
Run OpenMP*
Run x86 native code
Run offloaded code

Intel® Xeon Phi™ Coprocessor

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

Flexible Execution Models

Optimized Performance for all Workloads

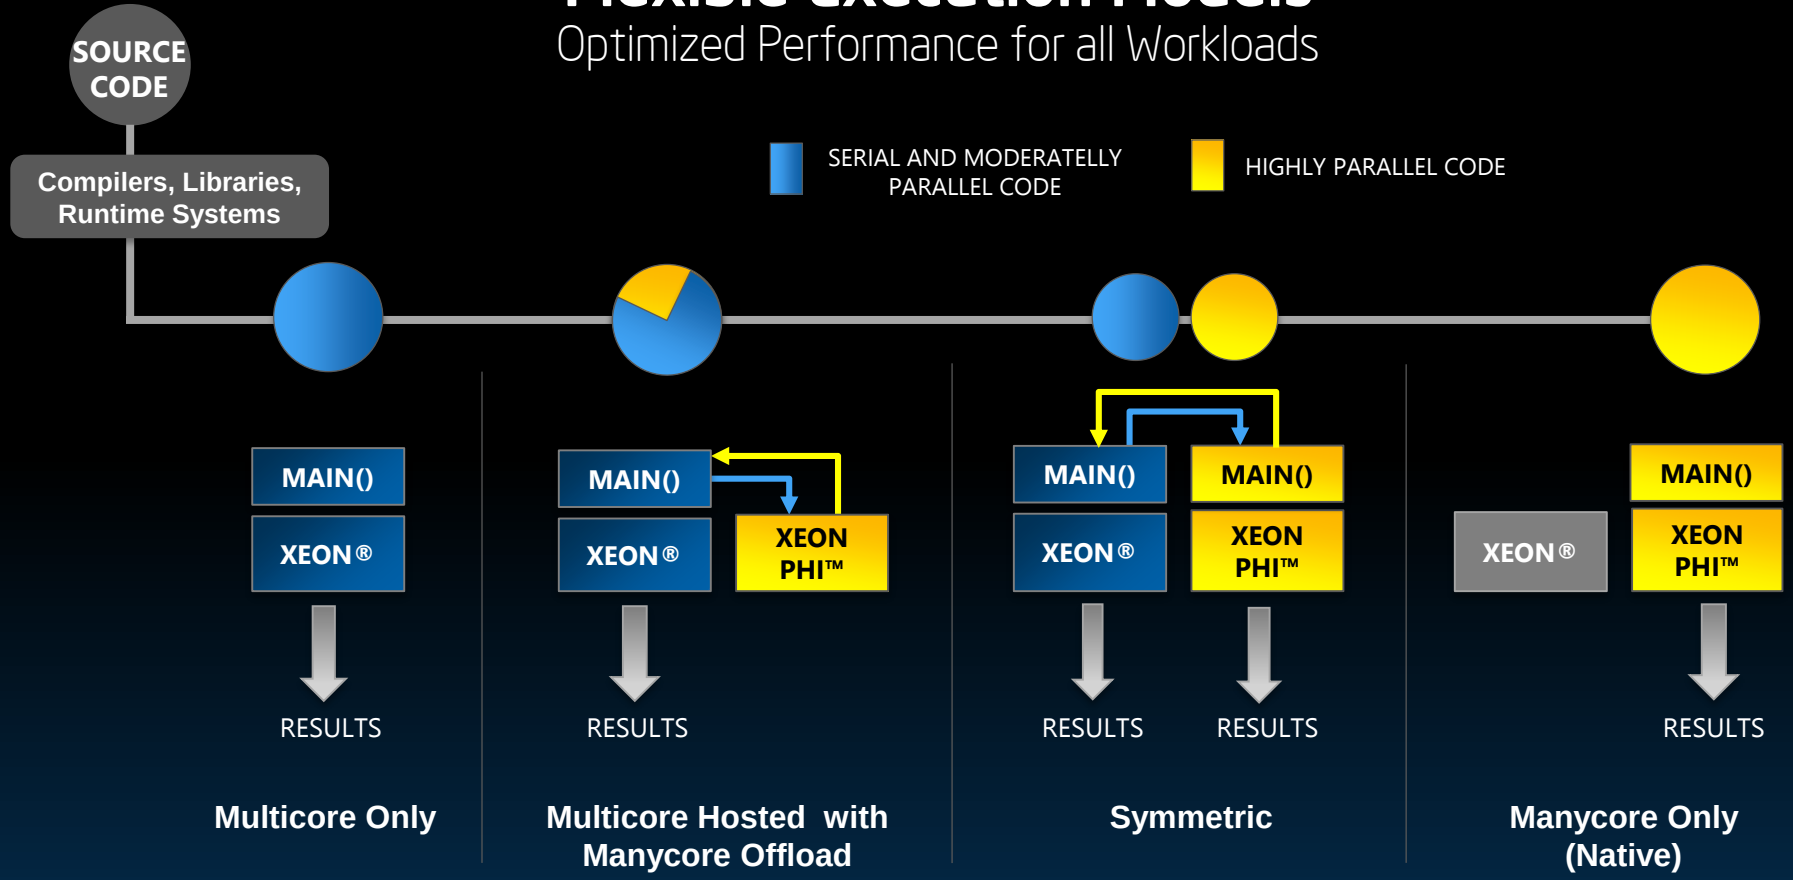

Manycore Processors – Example HPC Use Cases

1 MPI process
64 Threads each

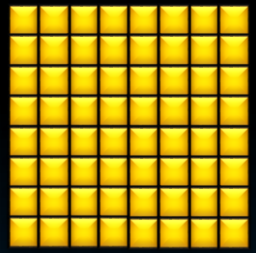

4 MPI processes
16 Threads each

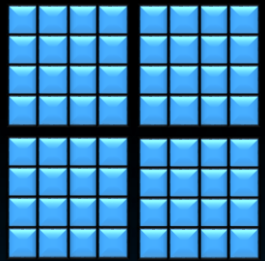

16 MPI processes
4 Threads each

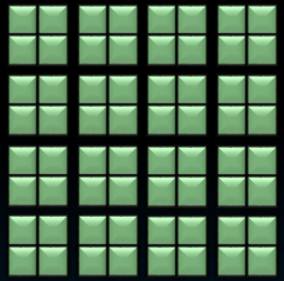

64 MPI processes
1 Threads each

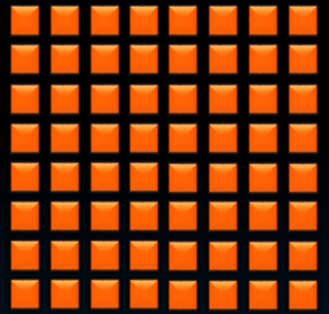

For illustration only.

Highly Parallel Applications

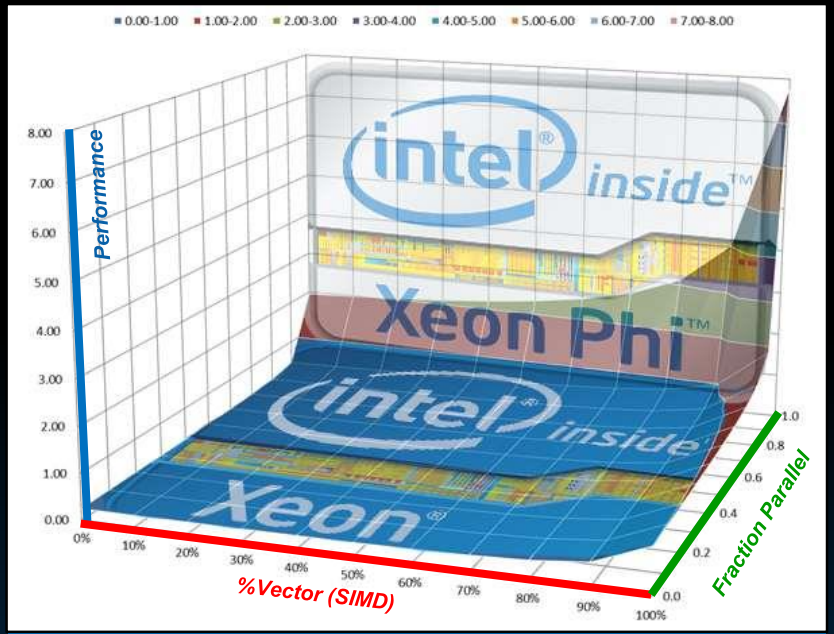

Theoretical acceleration of a highly parallel processor over a Intel® Xeon® parallel processor (<1: Intel® Xeon® faster) – For illustration only

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

Parallel Programming for Intel® Architecture (IA)

NODES	Use Intel® MPI, Co-Array Fortran
CORES	Use threads directly or e.g. via OpenMP*, pthreads Use tasking, Intel® TBB / Cilk™ Plus
VECTORS	Intrinsics, auto-vectorization, vector-libraries Language extensions for vector programming (SIMD)
BLOCKING	Use caches to hide memory latency Organize memory access for data reuse
DATA LAYOUT	Structure of arrays facilitates vector loads / stores, unit stride Align data for vector accesses

Parallel programming to utilize the hardware resources, in an abstracted and portable way

More Cores. Wider Vectors. Performance Delivered.

Intel® Parallel Studio XE and Intel® Cluster Studio XE

More Cores

Multicore (8+)

Manycore (60+)

Wider SIMD/Vector

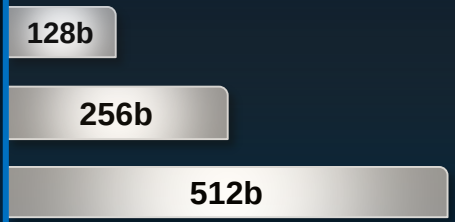

Scaling Performance Efficiently

Serial

Task & Data Parallel

Distributed

Industry Leading Software Tools

High-Performance from advanced compilers

Comprehensive libraries

Parallel programming models

Insightful analysis tools

OpenMP 4.0 Heterogeneous Programming

```
#pragma omp target data device(0) map(alloc:tmp[:N]) map(to:input[:N]) map(from:res)
  {
#pragma omp target device(0)
#pragma omp parallel for
    for (i=0; i<N; i++)
      tmp[i] = some_computation(input[i], i);

    do_some_other_stuff_on_host();

#pragma omp target device(0)
#pragma omp parallel for reduction(+:res)
    for (i=0; i<N; i++)
      res += final_computation(tmp[i], i)
  }
```

Case Study: NWChem CCSD(T) Method

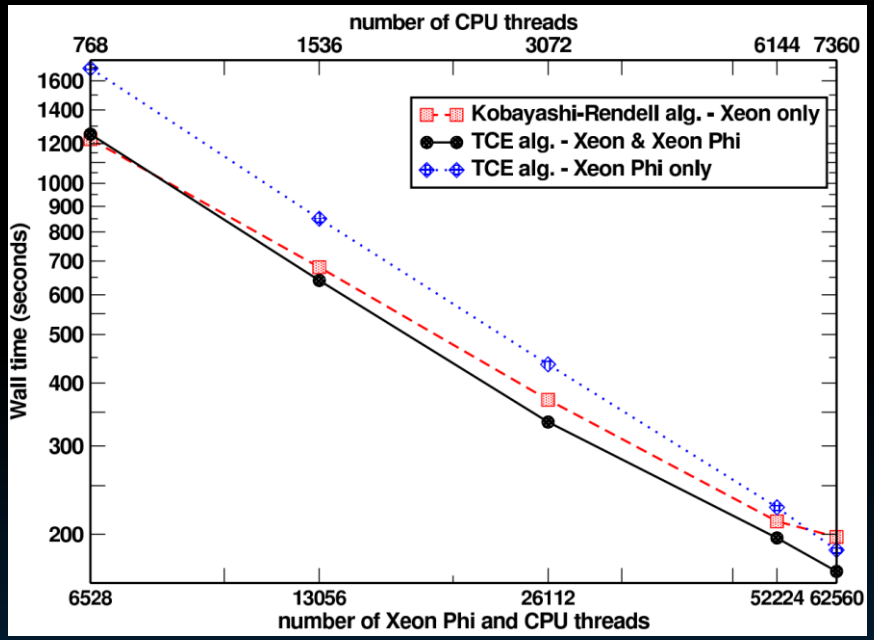

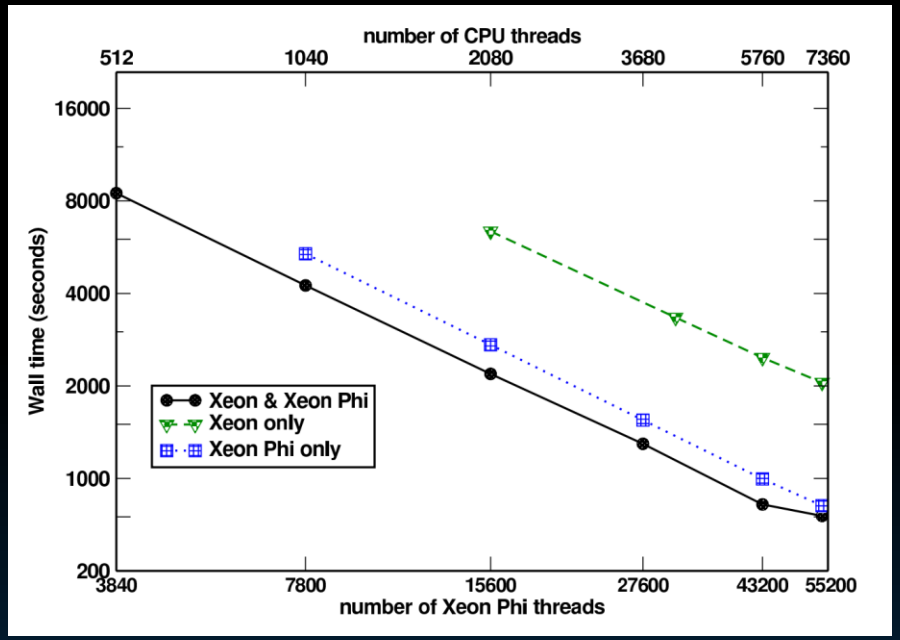

Edoardo Apra, Michael Klemm, and Karol Kowalski. *Efficient Implementation of Many-body Quantum Chemical Methods on the Intel Xeon Phi Coprocessor*. In Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, New Orleans, LA, November 2014. To appear.

Performance tests 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. System configuration: Atipa Visione vf442 server with two Intel Xeon E5-2670 8-core processors at 2.6 GHz (128 GB DDR3 with 1333 MHz, Scientific Linux release 6.5) and Intel C600 IOH, two Intel Xeon Phi coprocessors 5110P (GDDR5 with 3.6 GT/sec, driver v3.1.2-1, flash image/micro OS 2.1.02.0390, Intel Composer XE 14.0.1.106).

Next Intel® Xeon Phi™ Processor
Codename: Knights Landing

Designed using Intel's cutting-edge
14nm process

Not bound by "offloading" bottlenecks
Standalone CPU
or PCIe Coprocessor

Leadership compute & memory bandwidth
Integrated
On-Package Memory

Next Intel® Xeon Phi™ Processor

Codename: Knights Landing

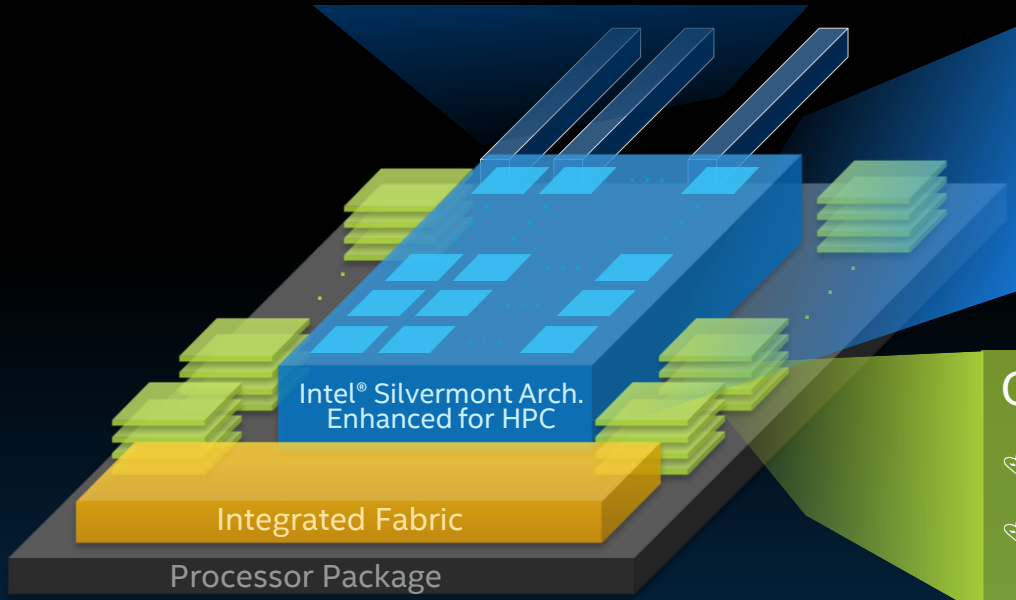

Compute: Energy-efficient IA cores[2]

- *Microarchitecture enhanced for HPC[3]*
- ***3x*** *Single Thread Performance vs Knights Corner[4]*
- *Intel Xeon Processor Binary Compatible[5]*

On-Package Memory:

- *up to 16GB at launch*
- ***1/3x*** *the Space[6]*
- *5X Bandwidth vs DDR4[7]*
- ***5x*** *Power Efficiency[6]*

Jointly Developed with Micron Technology


Journey to Exascale

Assume Exascale Computing at 20MW …

New Forms of Energy

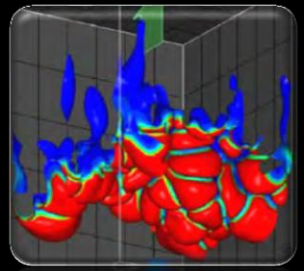

Ecological Sustainability

Space Exploration

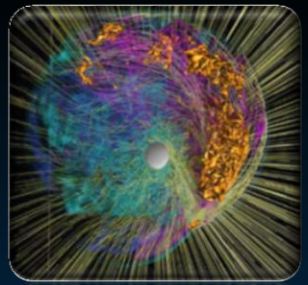

Medical Innovation

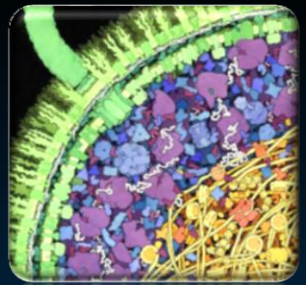

And many others….

1 Exaflop
20MW

1 Petaflop
20KW

Today's #25 system in a rack!

1 Teraflop
20W

100x the performance of today's phone at the same power

1 Gigaflop
20mW

For illustration and concept only.

HPC:
The Path to Exascale

Processors
Intel® Xeon® Processor

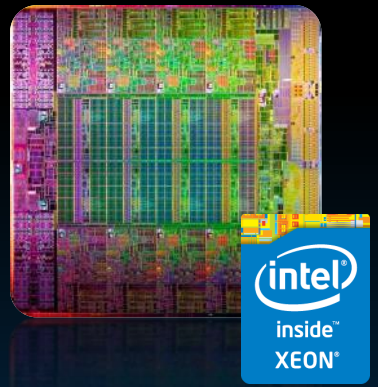

Coprocessor
Intel® Many Integrated Core

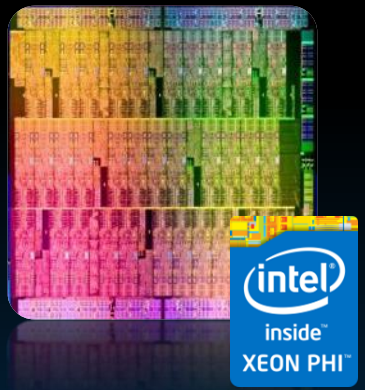

Fabrics

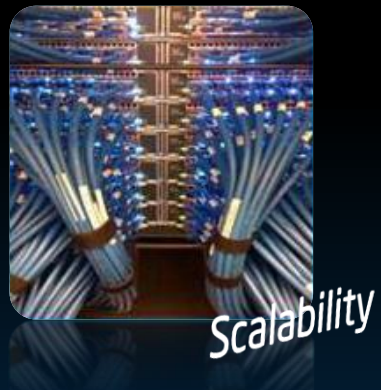

Software

HPC:
The Path to Exascale (cont.)

Memory & Storage

Networking

Reliability & Resiliency

Power Management

Intel TeraScale Research Areas

MANY-CORE COMPUTING

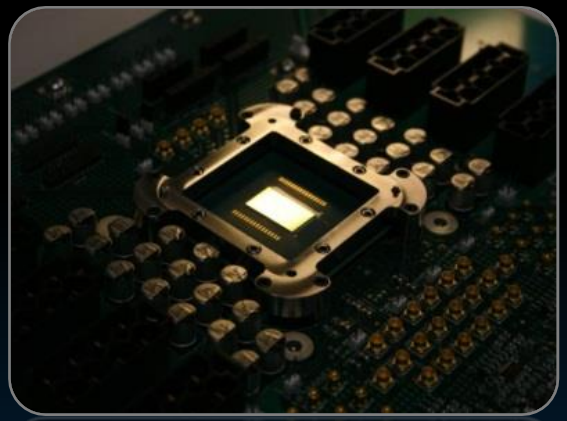

Teraflops
of computing power

STACKED MEMORY

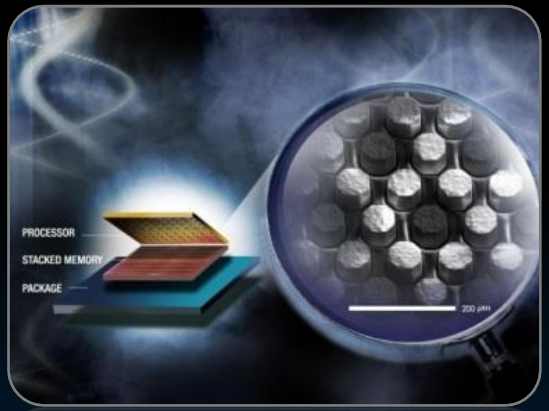

Terabytes
of memory bandwidth

SILICON PHOTONICS

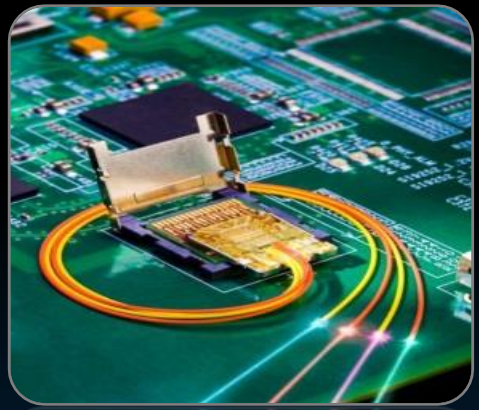

Terabits
of I/O throughput

Future vision, does not represent real products.

The Power of Solutions: Big Data Example

Sort 1TB of Data:

>4 Hours

Sort 1TB of Data:

7 MINUTES

*Other brands and names are the property of their respective owners

The Power of Platform Solutions

TeraSort for 1TB sort:

>4 hour process time

Xeon 5600
HDD
1GbE

Upgrade processor ~50% reduction

Upgrade to SSD ~80% reduction

Upgrade to 10GbE ~50% reduction

Intel distribution ~40% reduction

hadoop

<7 minutes


PERFORMANCE
it's all about SCALABLE parallelism
THINK PARALLEL!
INTEL
INNOVATION & LEADERSHIP
FOR THE ROAD AHEAD

Thank You.