Accelerating COMPUTING

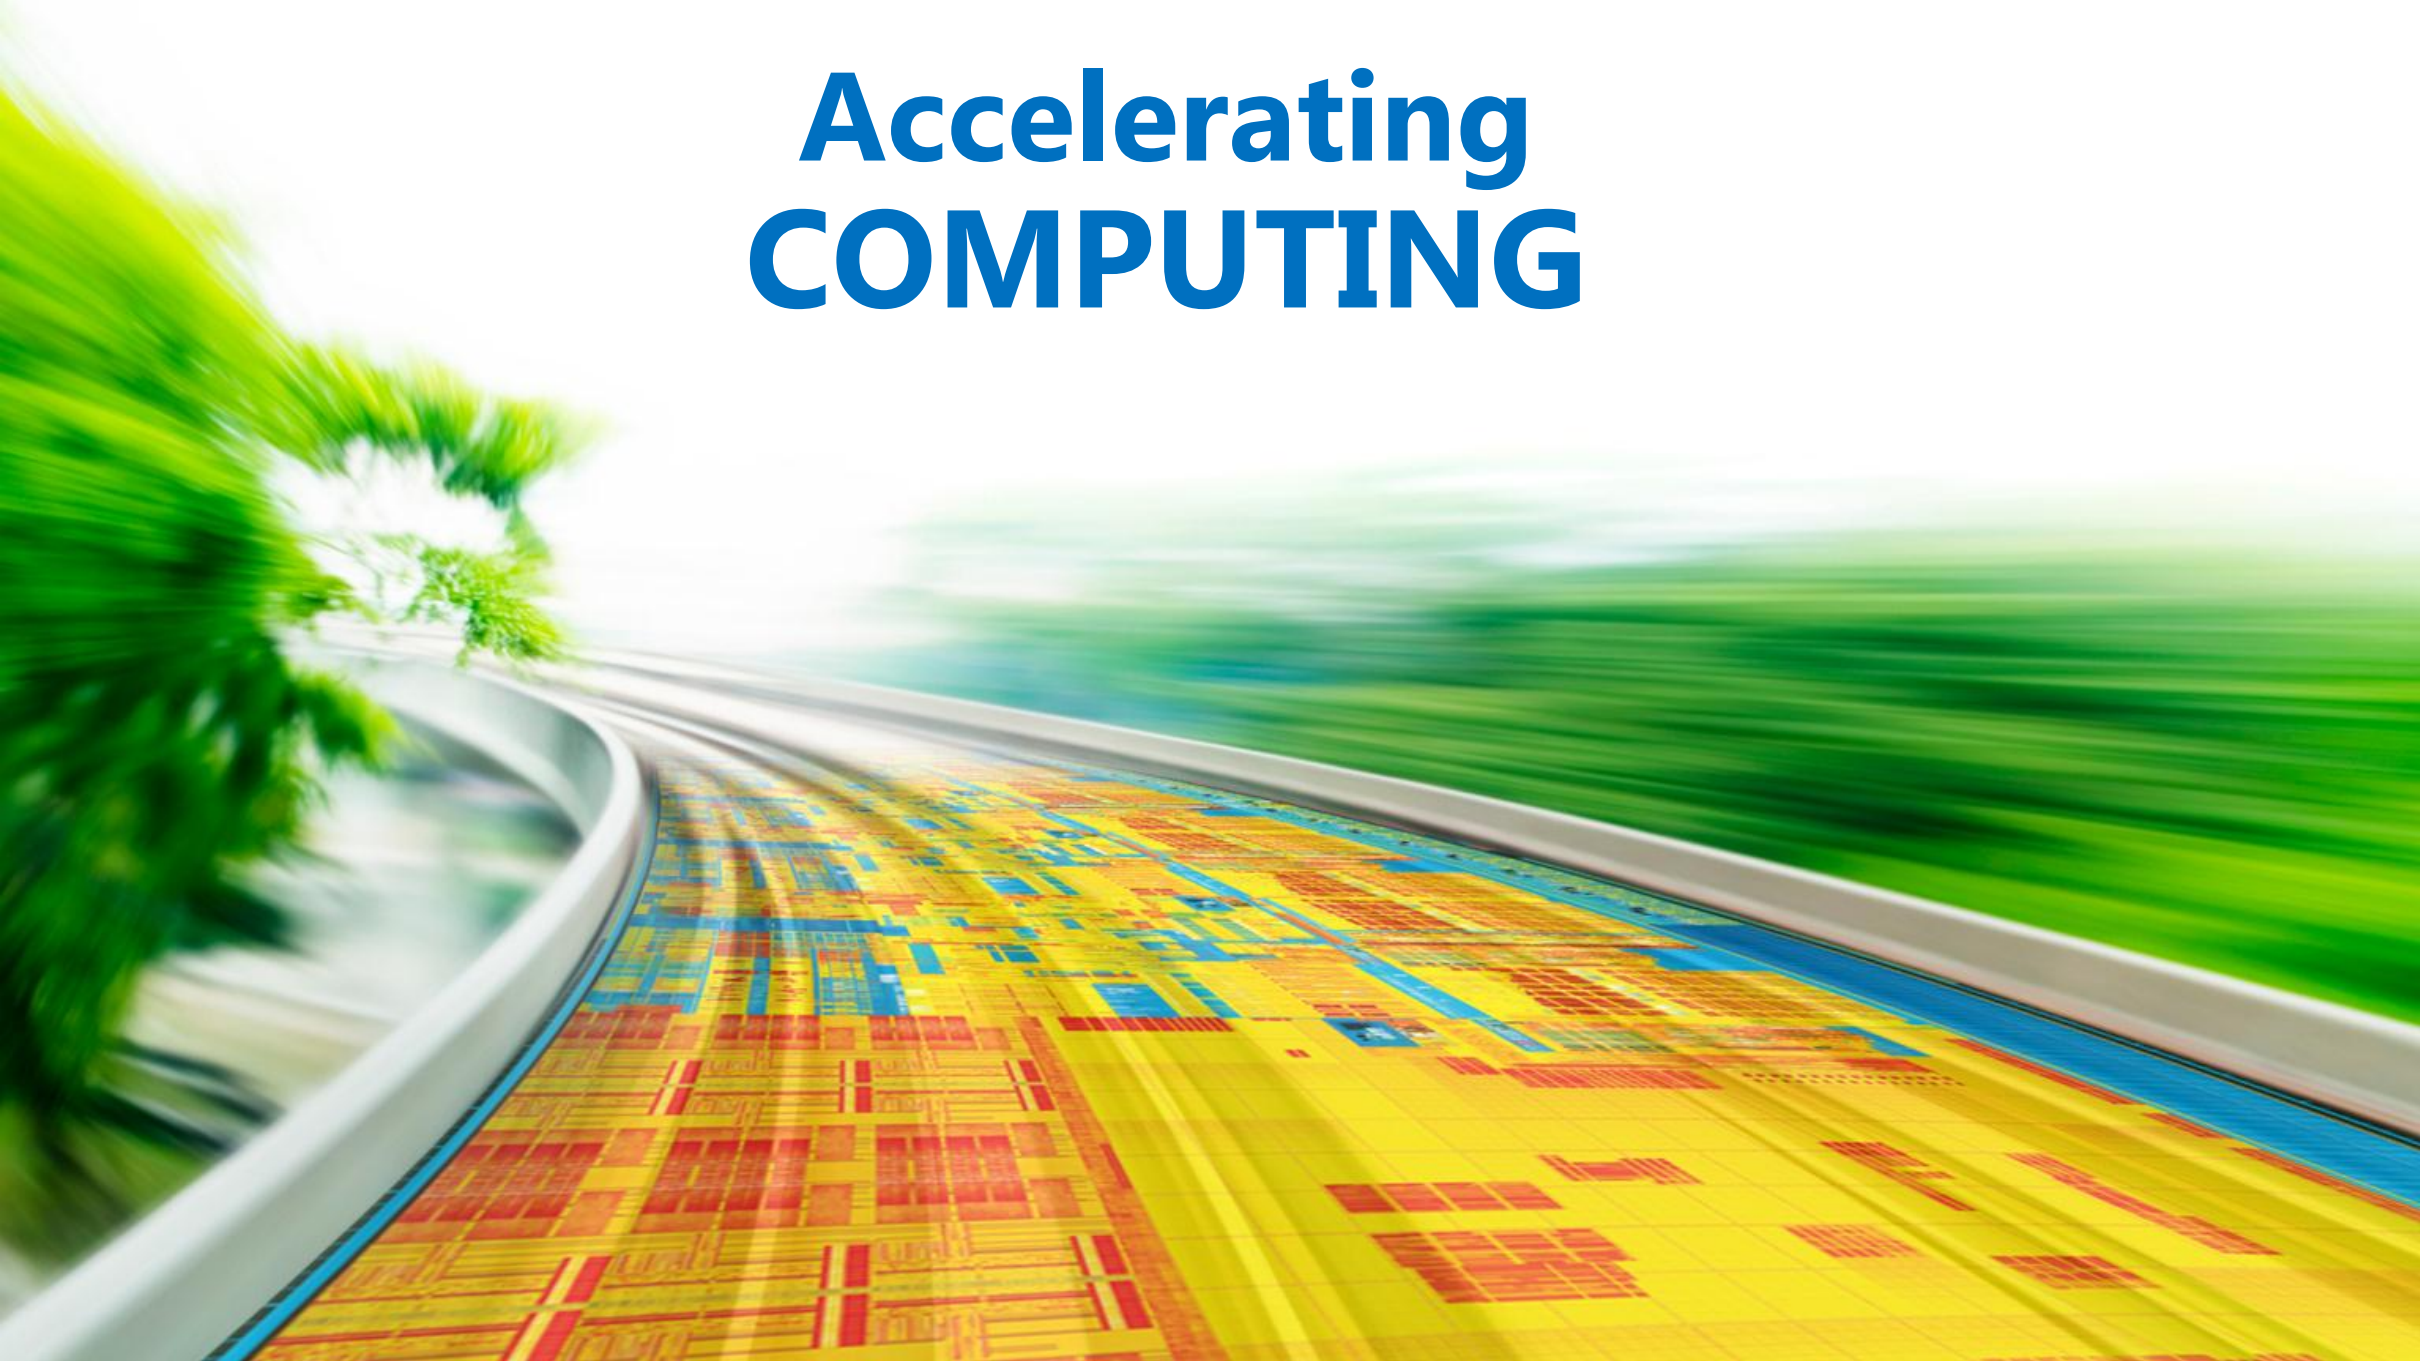

Intel Technologies for Advanced Computing

Andrey Semin
HPC Technology Manager, Sr. Staff Engineer
Intel Corporation, EMEA

June 2012
Moscow

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 PROPERTY 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 PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS.

• 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.

• Penryn, Nehalem, Westmere, 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

• 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.

• Intel, Xeon, Netburst, Core, VTune, 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.*

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.

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

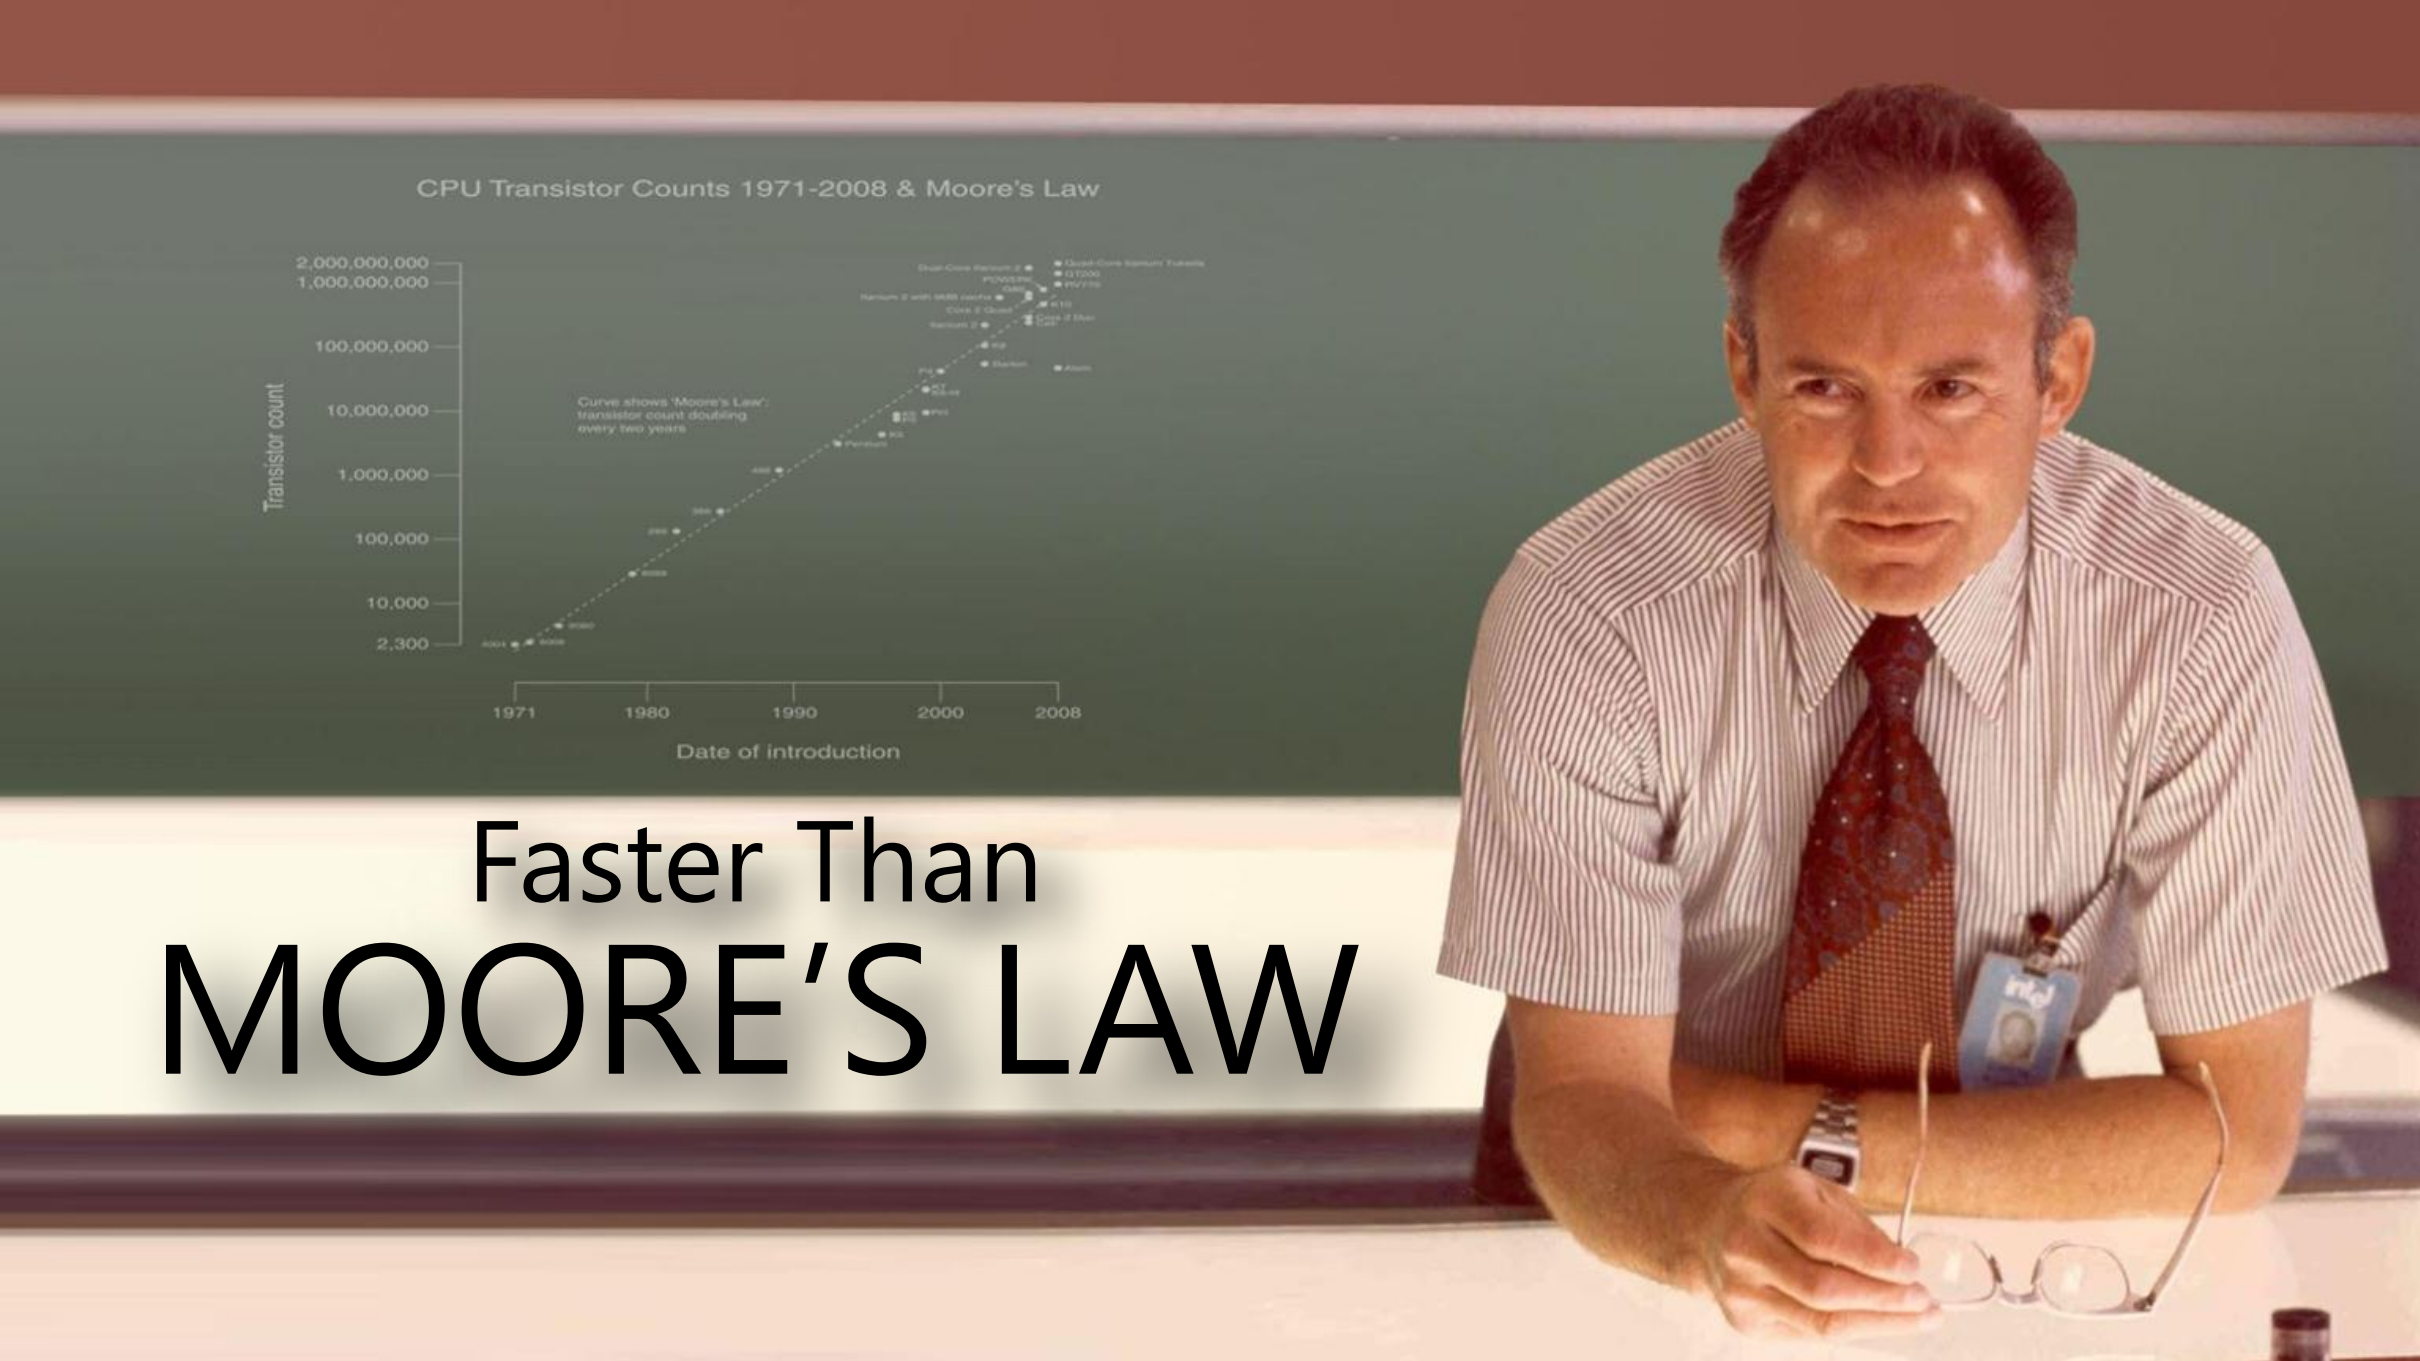
CPU Transistor Counts 1971-2008 & Moore's Law
2,000,000,000
1,000,000,000
100,000,000
10,000,000
1,000,000
100,000
10,000
2,300
Transistor count
Curve shows 'Moore's Law' - transistor count doubling every two years
1971
1980
1990
2000
2008
Date of introduction
Faster Than
MOORE'S LAW

Moore's Law at Work

1997

1 TFLOPS DP-F.P.

9298 Chips

"ASCI RED"

~2500 Square Feet

850KW Supercomputer

Source: Sandia

2012

1 TFLOPS DP-F.P.

Single node

with Xeon® Phi™ (MIC)

Intel in High-Performance Computing

Dedicated, Renowned Applications Expertise

Large Scale Clusters for Test & Optimization

Tera-Scale Research

Exa-Scale Labs

Defined HPC Application Platform

Broad Software Tools Portfolio

Industry Standards

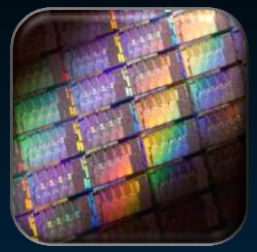

Manufacturing Process Technologies

Leading Processor Performance, Energy Efficiency

Many Integrated Core (MIC) Architecture

Platform Building Blocks

A long term commitment to the HPC market segment

Xeon Spanning a Diverse Set of Workloads

Xeon® E3 (Bromolow)	Xeon® E5 (Sandy Bridge-EP)	Xeon® E7 (Westmere-EX)	Xeon® Phi™ (MIC platform)
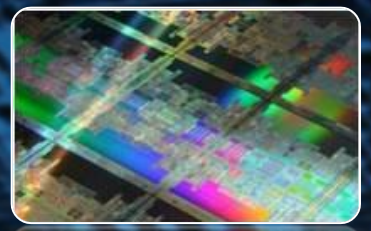		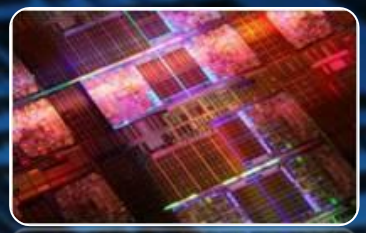	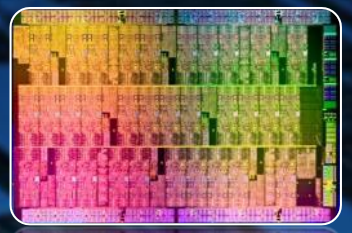
Entry (1S)	**Compute Nodes (2S)**	**SuperNodes (4+S)**	**Many-Core Computing**

1993~2012 CPUs Performance (vs. Pentium)*

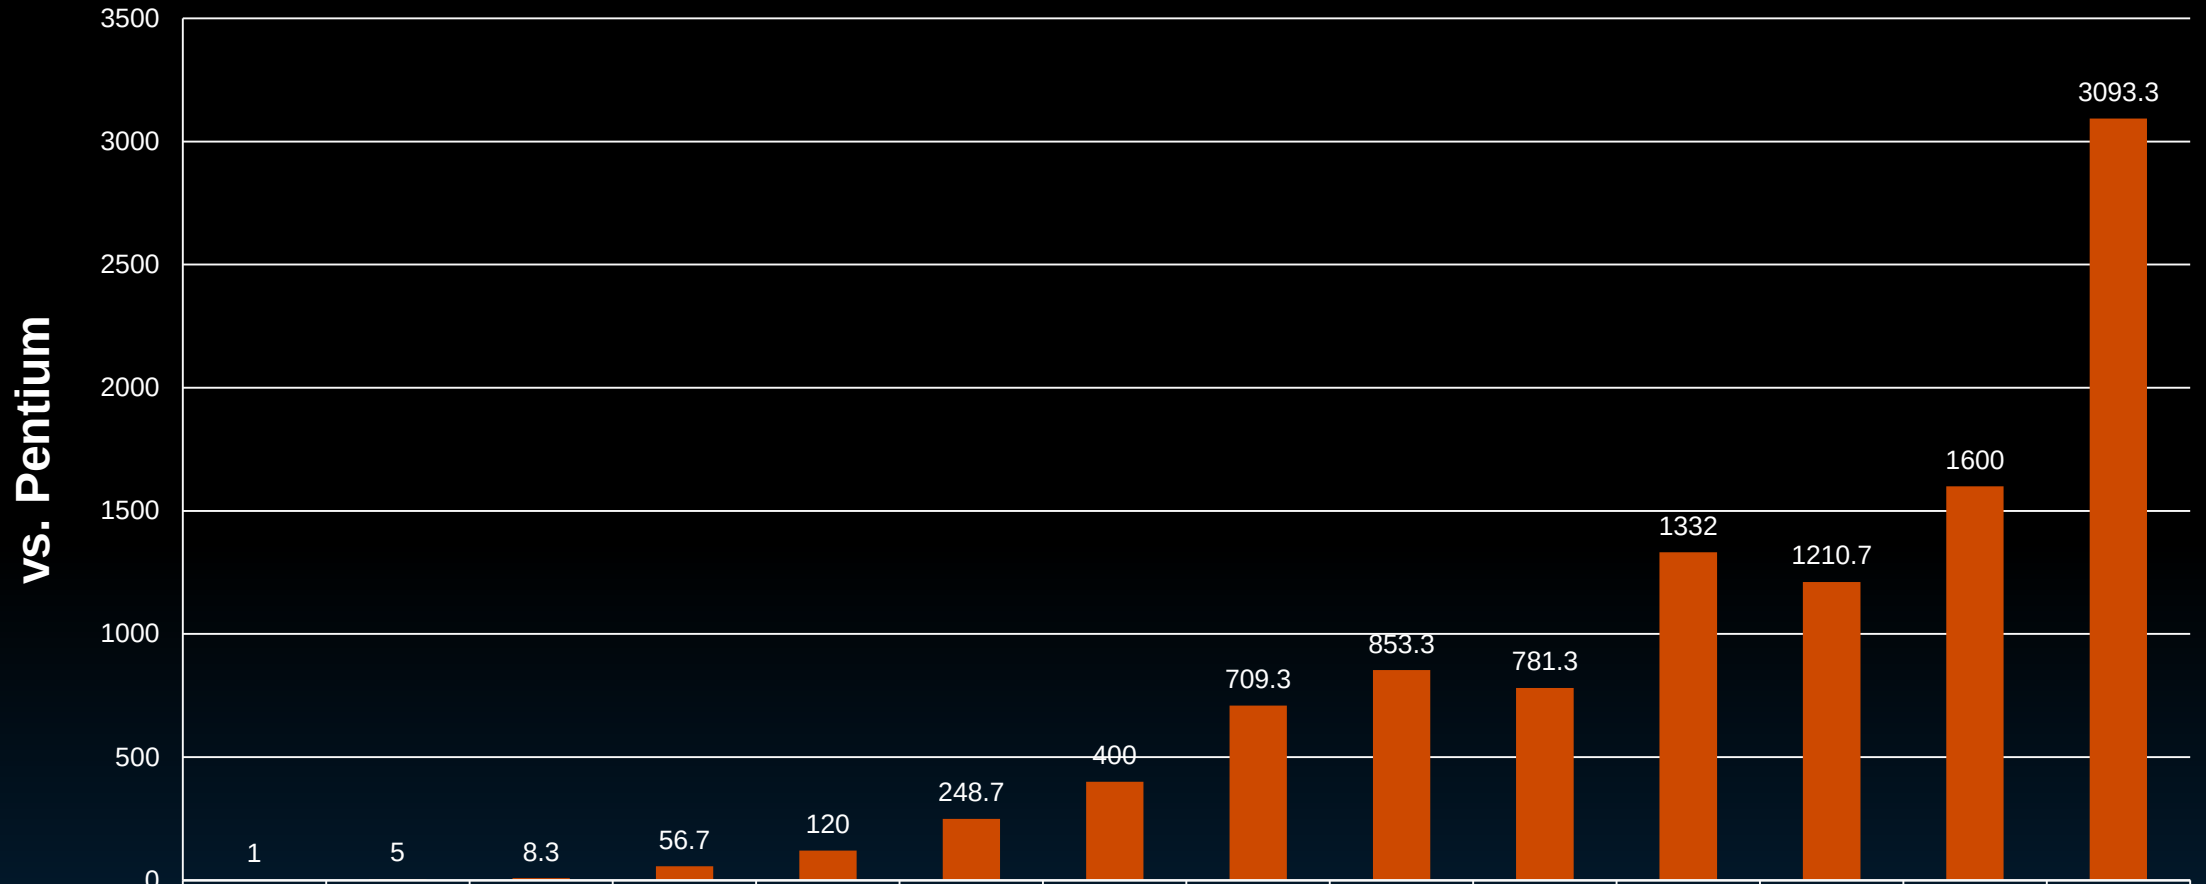

	Pentium	Pentium II Xeon	Pentium III Xeon	Pentium 4 Xeon	Pentium 4 Xeon	Xeon 5000	Xeon 5100	Xeon 5300	Xeon 5400	Nehalem-EP 5500	Westmere-EP 5600	Nehalem-EX 7500	Westmere-EX E7	Sandy Bridge E5
	1993	1997	1999	2001	2005	2006	2006	2006	2007	2008	2010	2010	2011	2012
vs.Pentium	1	5	8.3	56.7	120	248.7	400	709.3	853.3	781.3	1332	1210.7	1600	3093.3

*** Float Theoretical Peak performance**

CPU Index

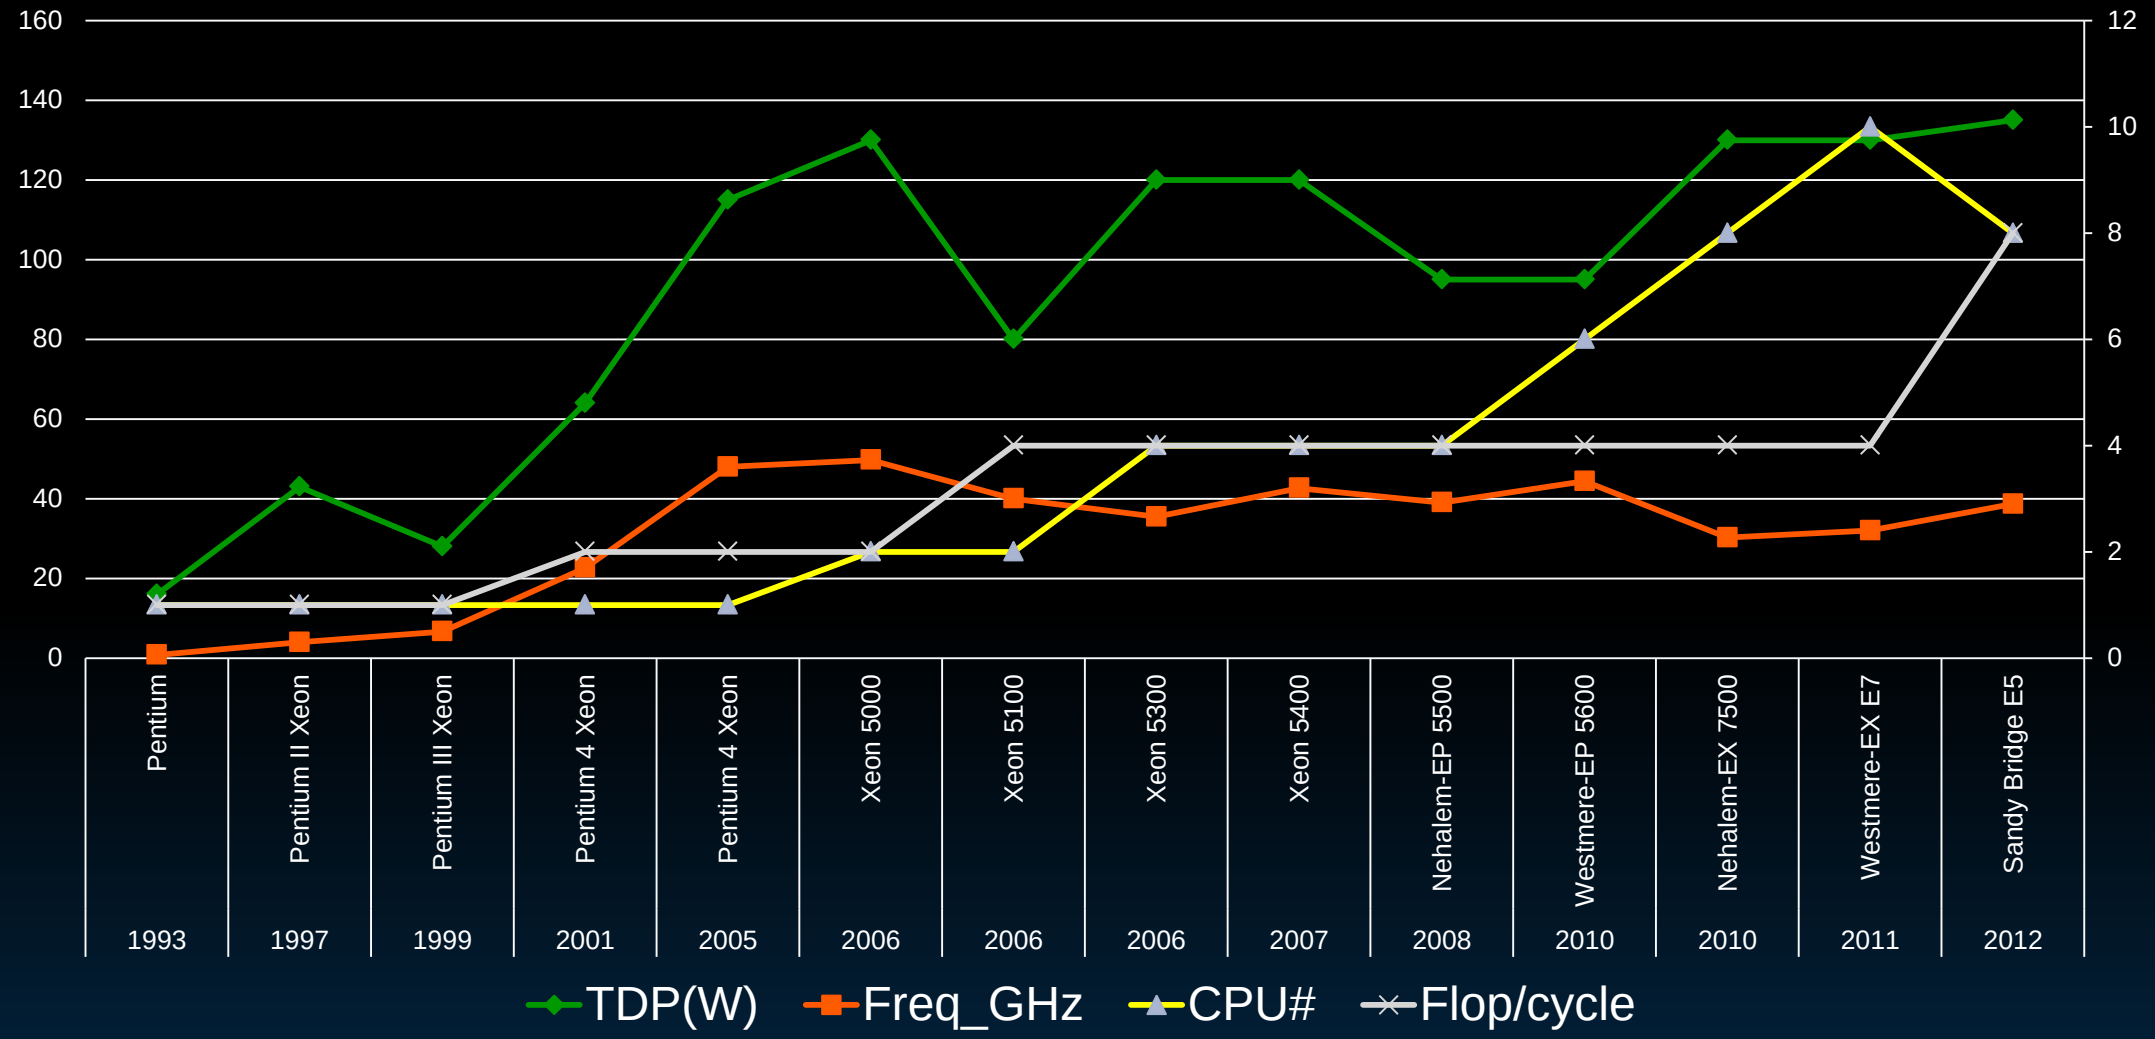

Flop/cycle, CPU# - going up
Power and Freq - fluctuating year by year

Increasing Processor Performance

Through Many-Core Technologies for Highly Parallel Workloads

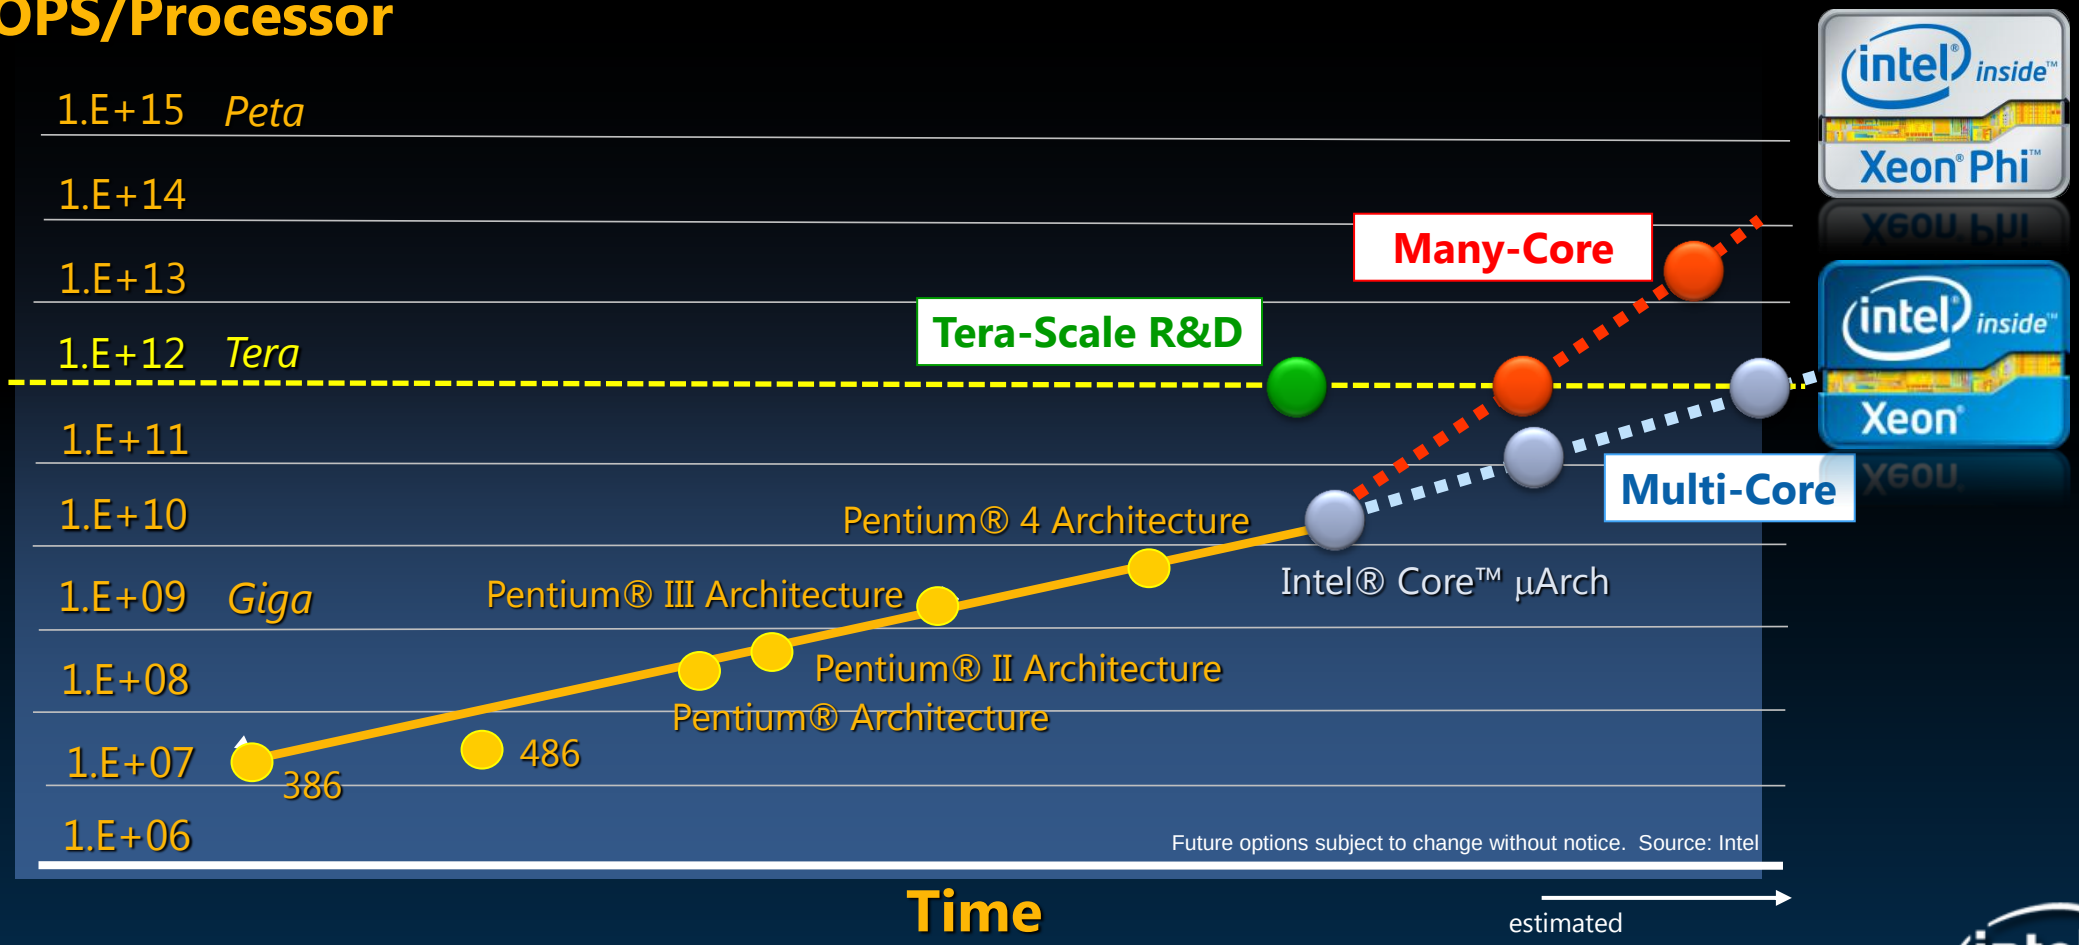

All dates, product descriptions, availability, and plans are forecasts and subject to change without notice.

IA Cores build on a Common Architecture

Scalable Performance
Energy Efficient
Microarchitecture

Highly Parallel
Energy Efficient
Architecture

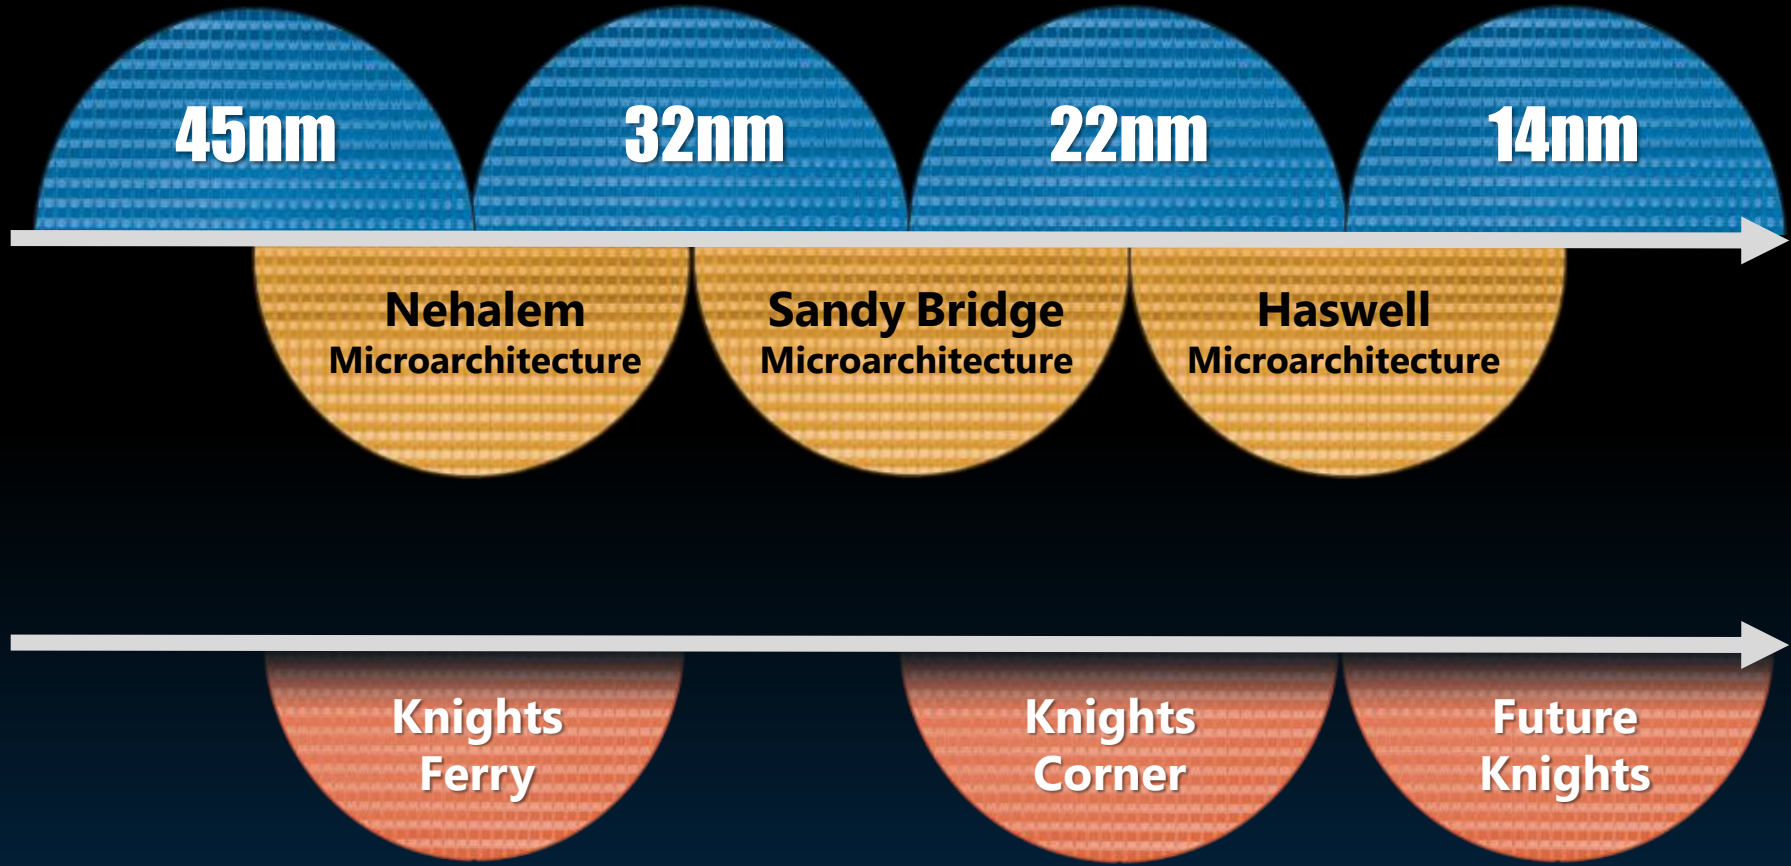

Potential future options, subject to change without notice.

Process Technology Research @ Intel

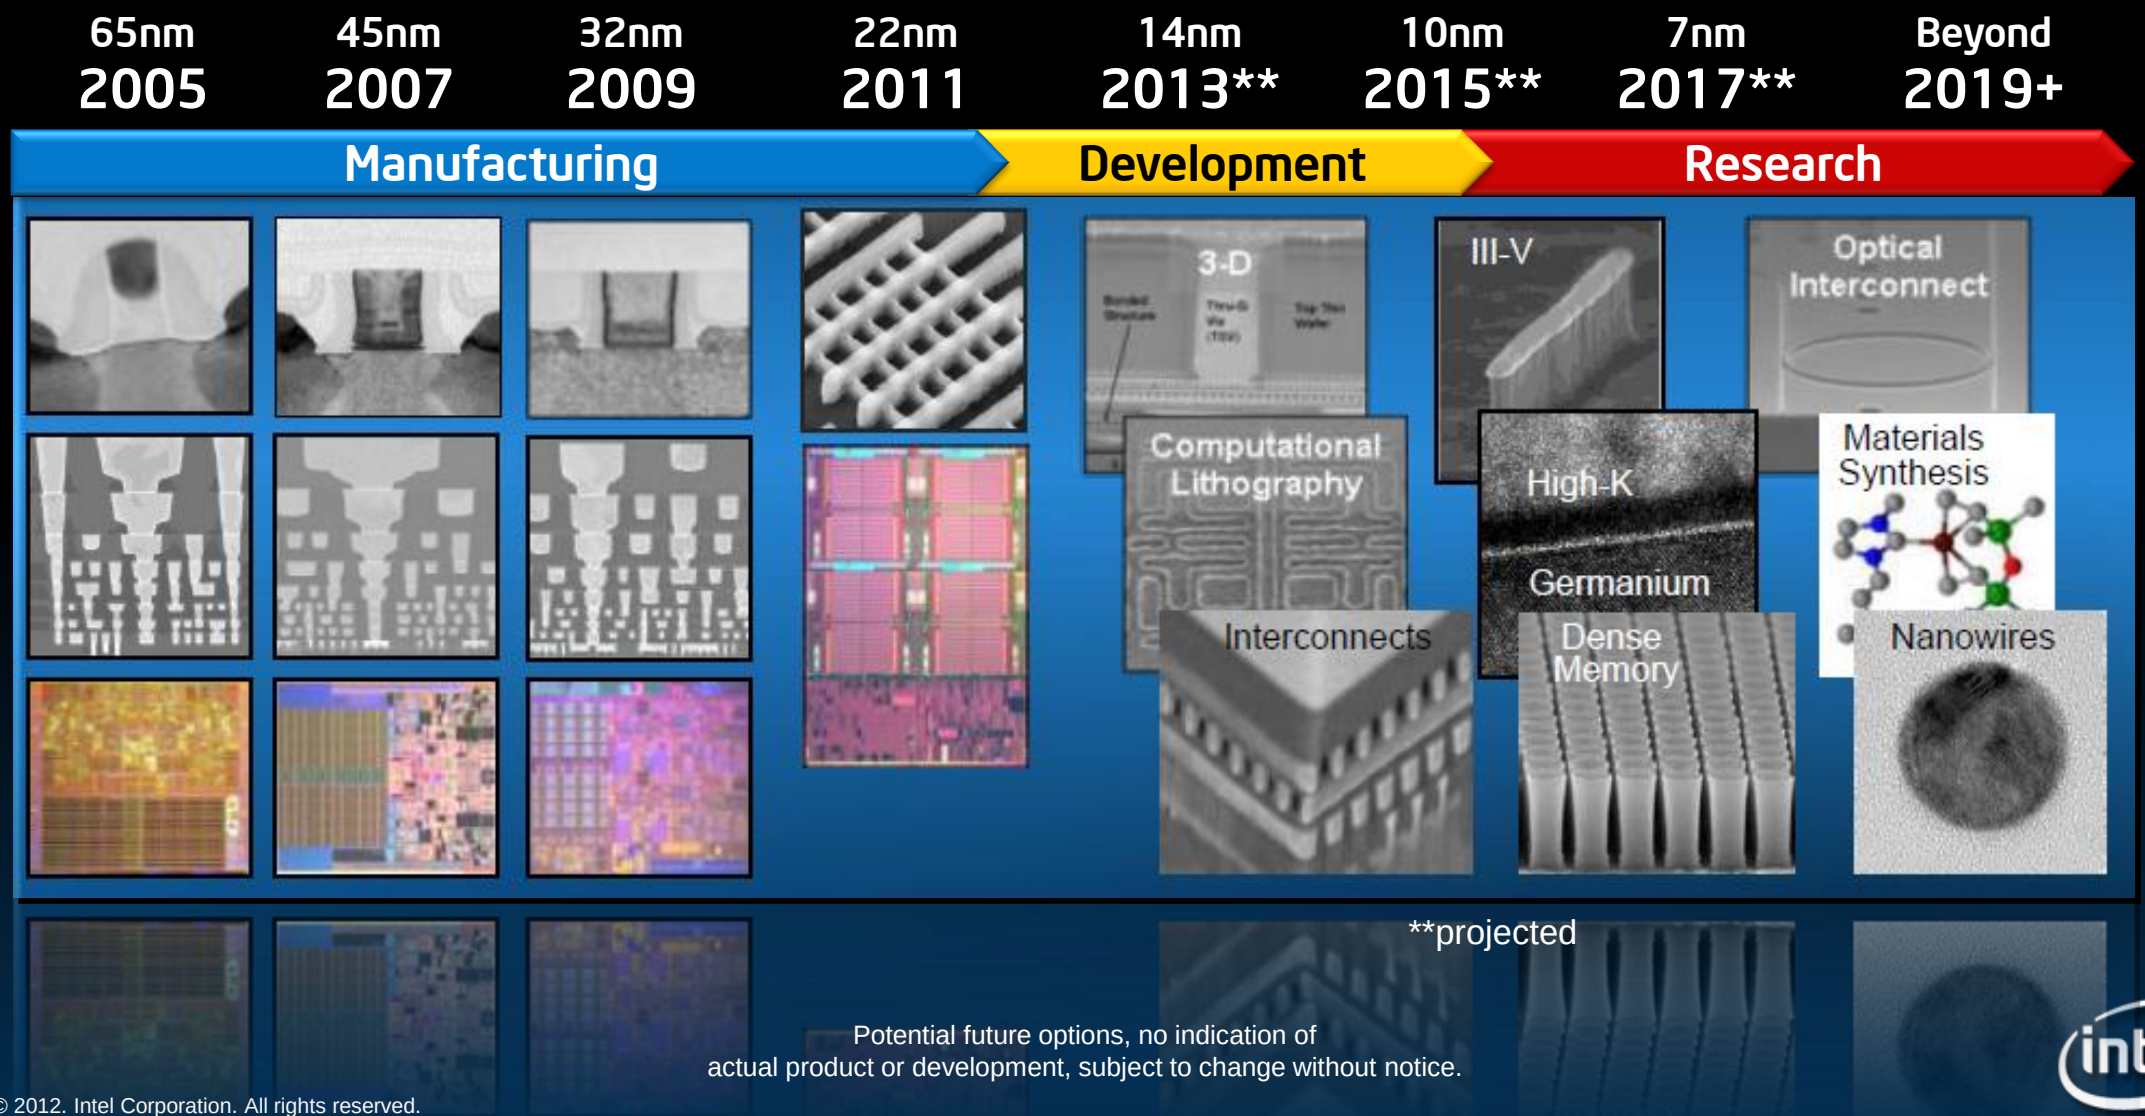

**projected

Potential future options, no indication of
actual product or development, subject to change without notice.

Intel 22nm Tri-Gate Technology

Smaller
Faster
Less Power

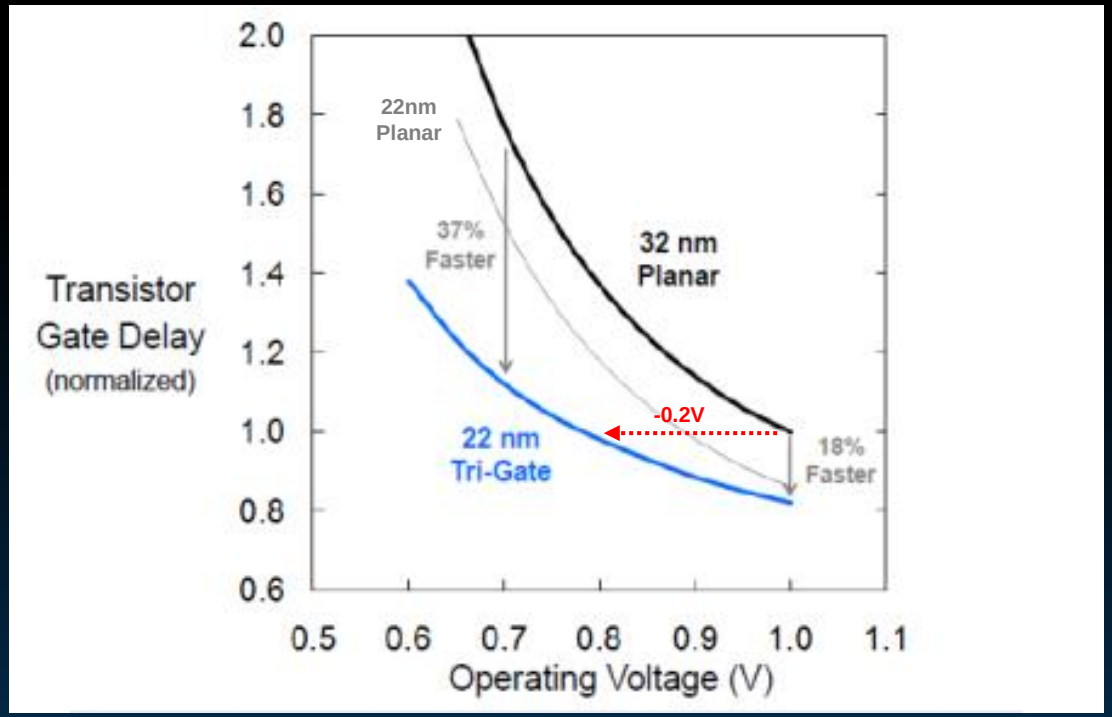

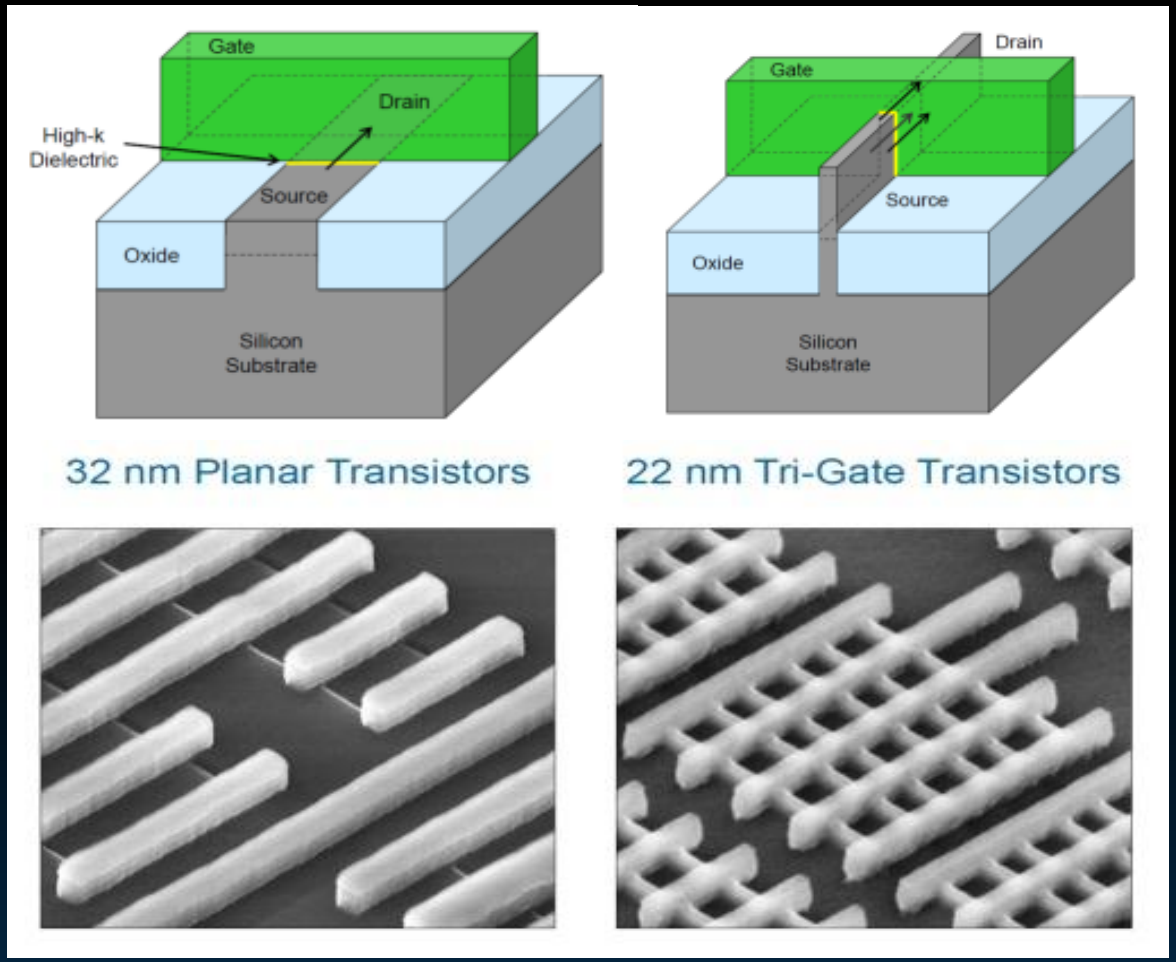

Increasing Performance and Energy Efficiency

Process Technology
22NM

Legacy
TURBO-BOOST

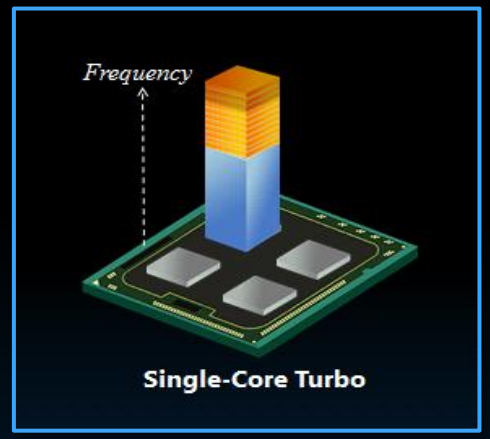

Core Architecture
AVX

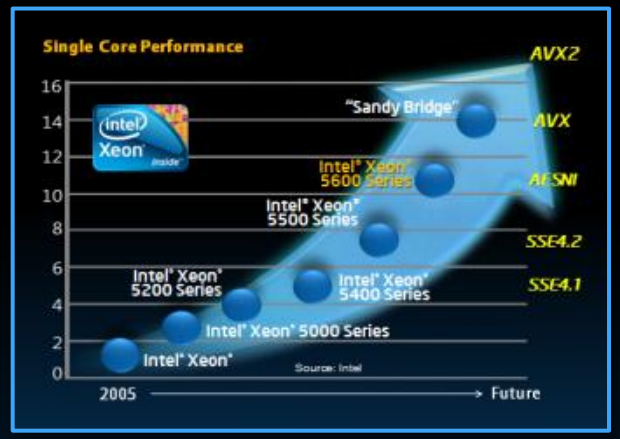

Processors
MULTI/MANY-CORE

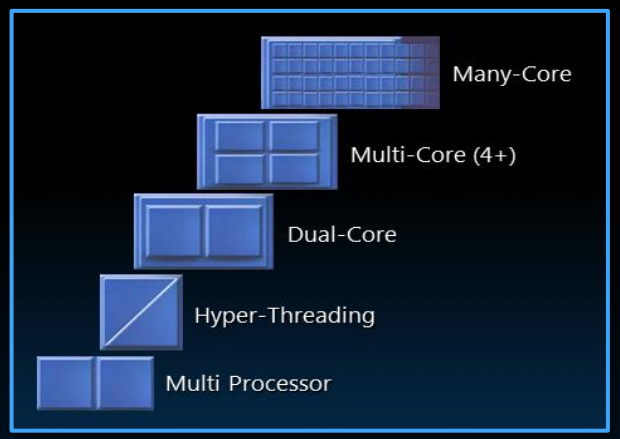

Note:
not all cores are equal !

Potential future options, subject to change without notice.

Industry Trend to Multi/Many-Core

Energy Efficient (HPC) Computing with Multi/Many-Core Processors

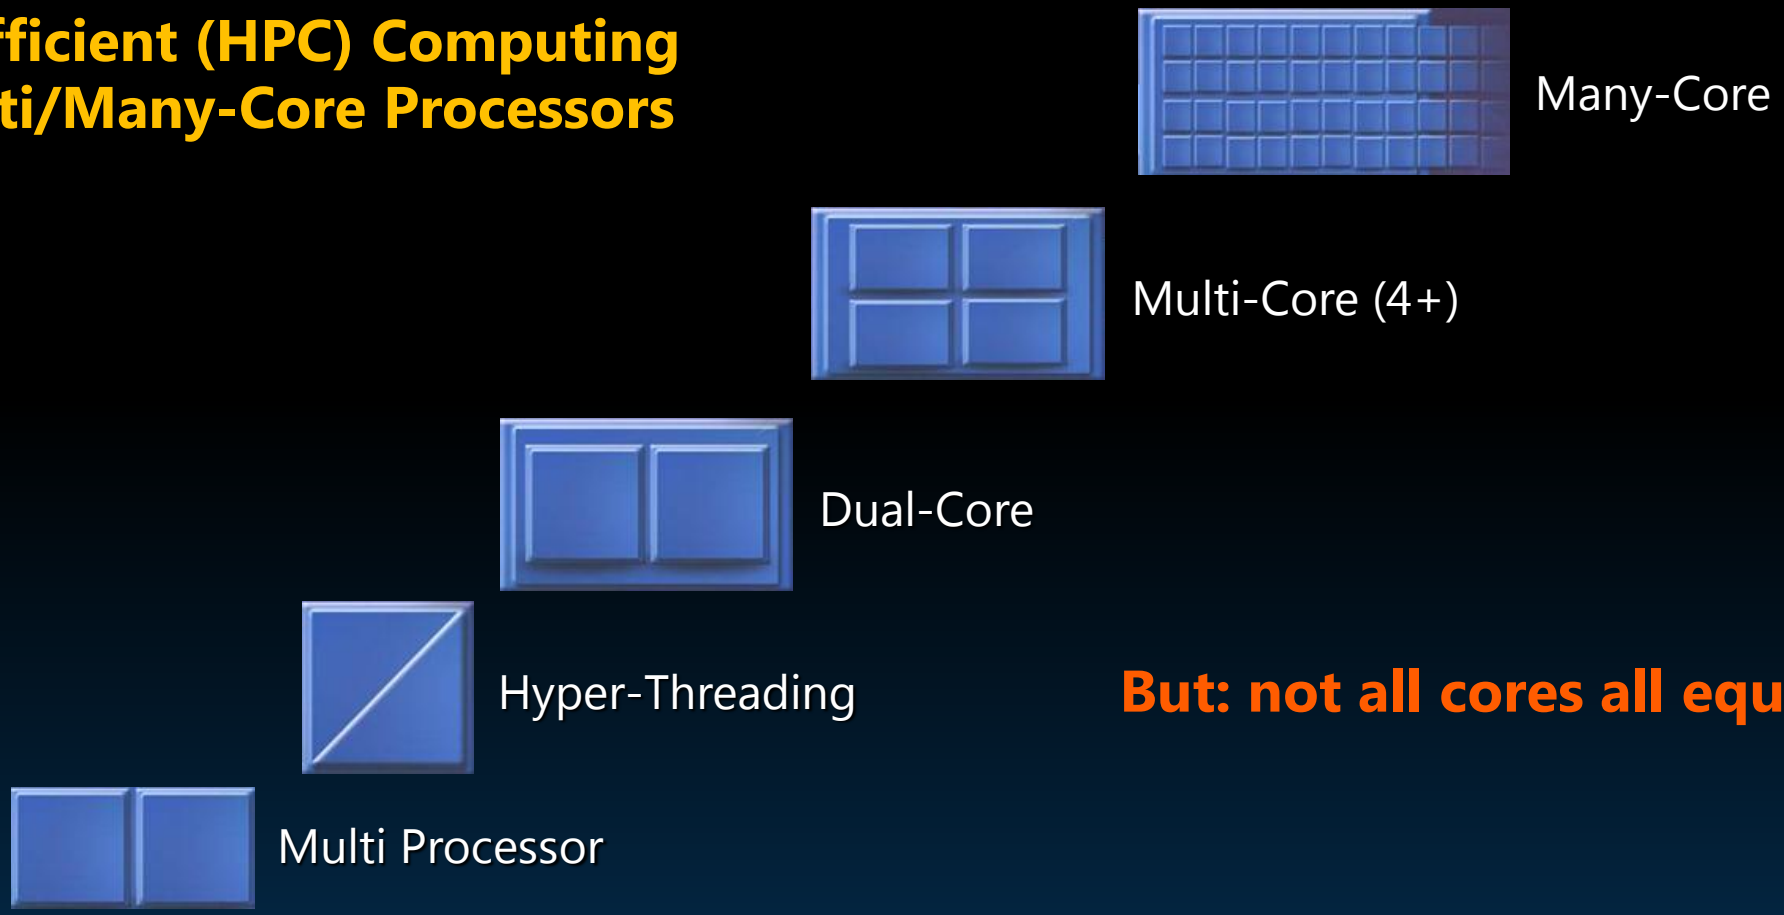

But: not all cores all equal !

(for illustration only)

Intelligent Processor Performance Scaling Forward

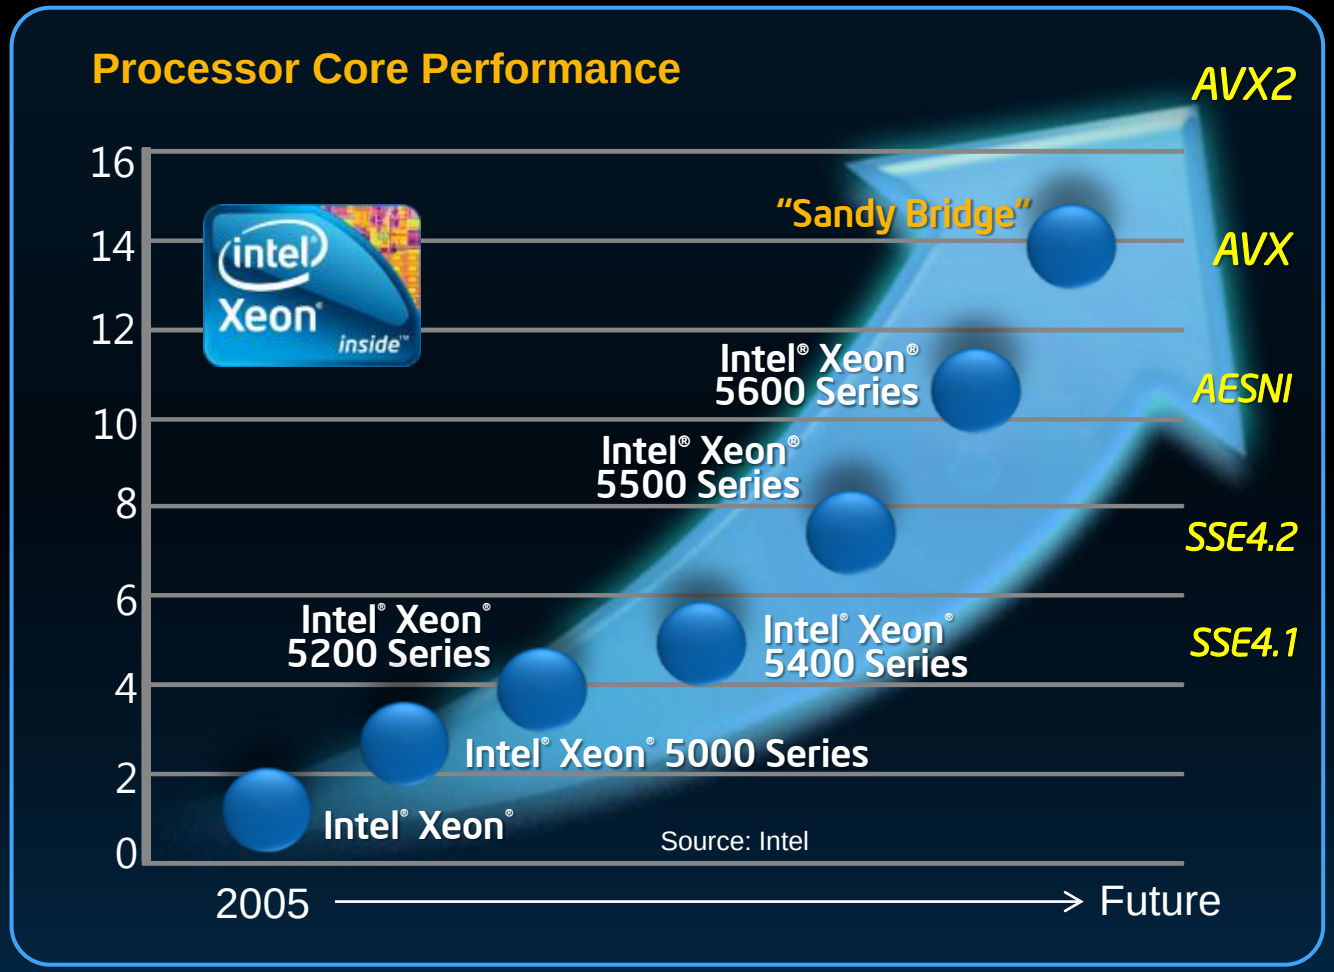

Faster Time To Productivity

- Total Application Performance
- Increased Single Thread Performance
- Increased Floating Point Performance and Bandwidth
- Irregular Data-Access
- Balanced Processor and System Architecture
- Less Complex Software Development and Support

Potential future options, subject to change without notice.

Intel® Turbo Boost Technology 2.0

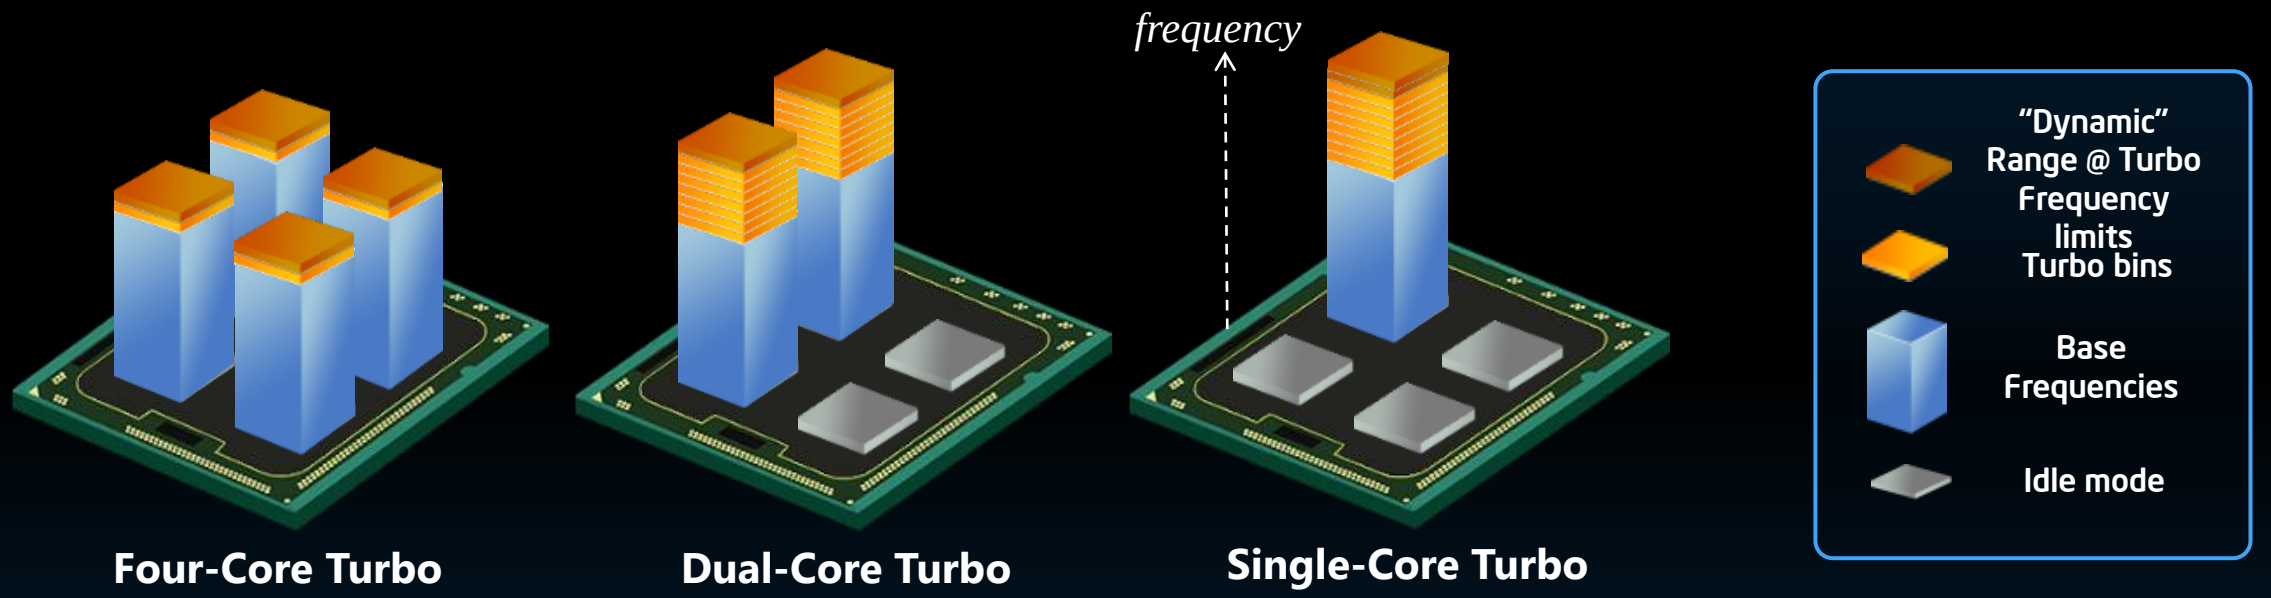

Intelligent and energy efficient performance on demand

The number of Turbo bins shown is only for illustrative purposes and is not representative of the actual number of turbo bins available.

Dual-/Quad- Socket Xeon® Processor

Sandy Bridge "Tock"

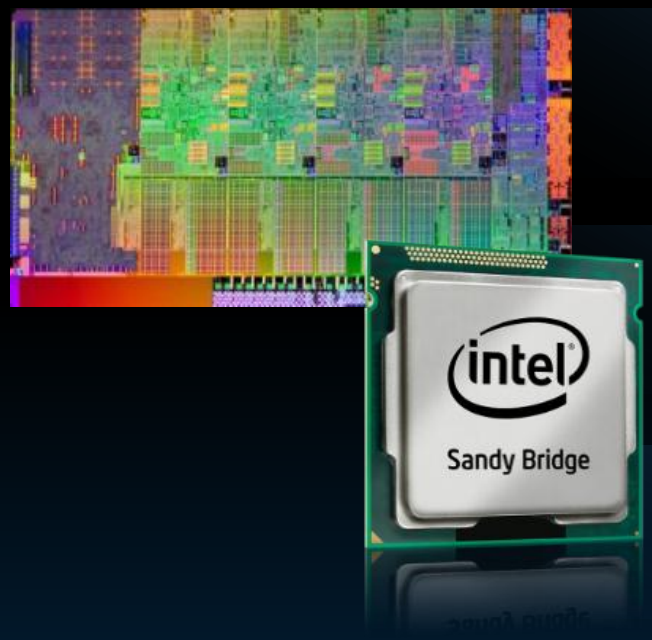

Significantly **greater performance** with higher core count, Intel® Hyper-threading and Turbo Technology

2x Flops / clock peak using new **AVX** instructions

Making **Petascale** Widely Available for Leading Science

Two Socket Platform Evolution

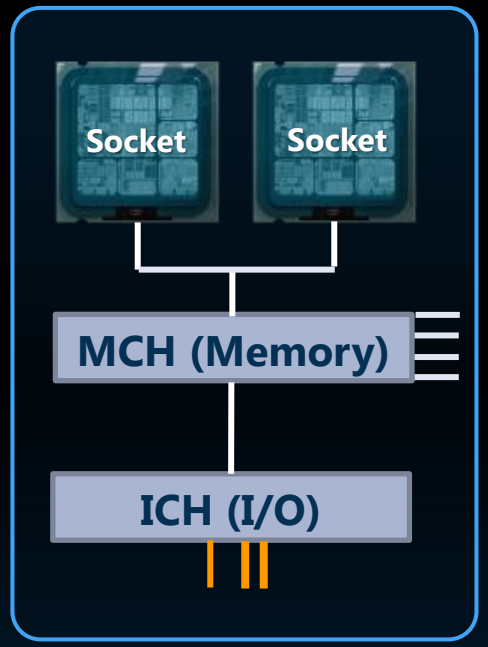

Frontsidebus

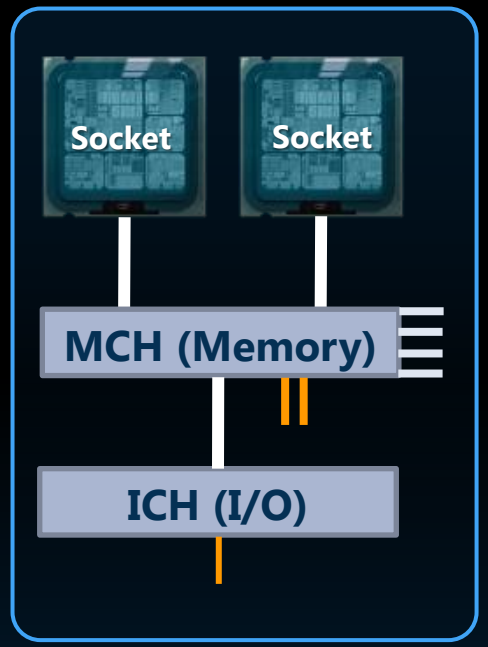

Multiple Frontsidebusses

Nehalem Architecture

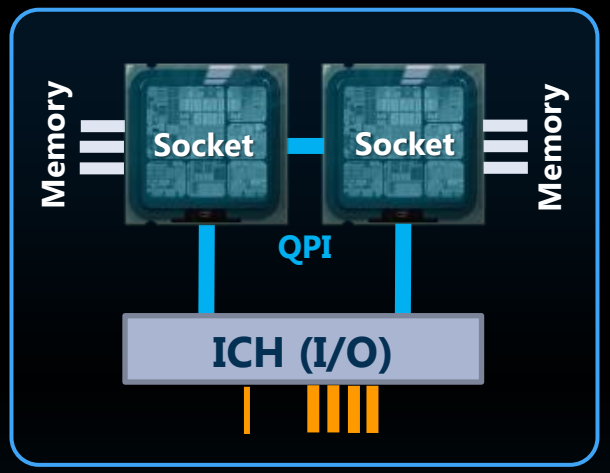

Integrated Memory Controller, QPI, PCIe

Sandy Bridge Architecture

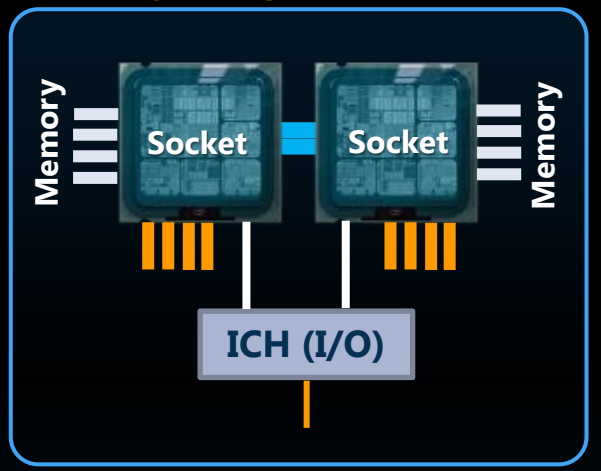

Integrated Memory Controller, QPI, Integrated I/O PCIe

Schematic overview, potential future options, subject to change without notice.

Sandy Bridge/Romley based Server Platforms

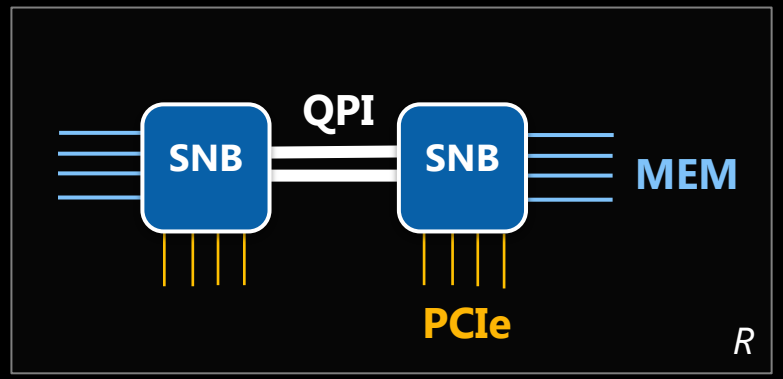

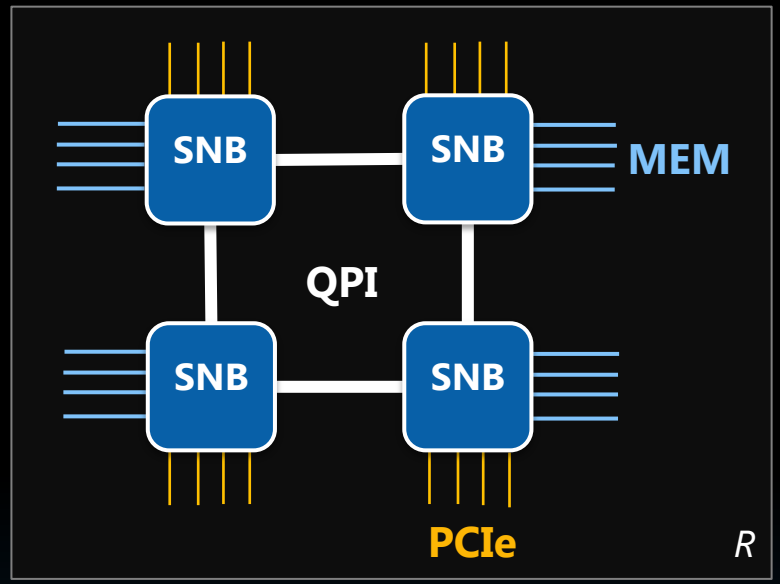

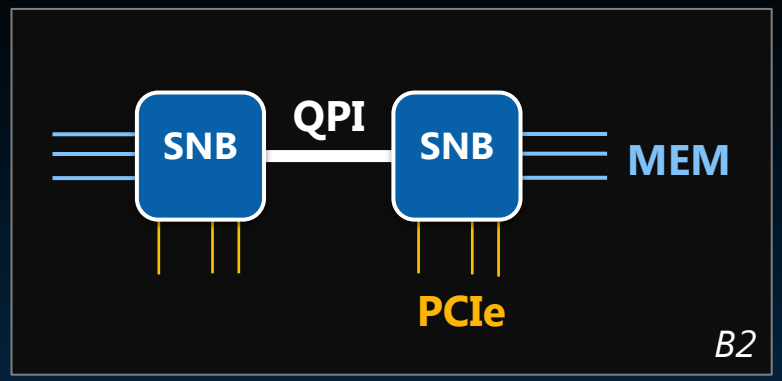

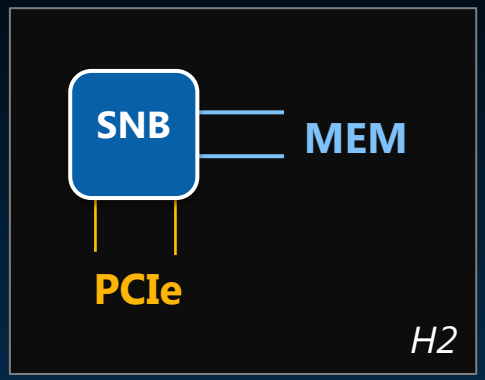

Sandy Bridge EP - "Ring Bus"

Ring bus architecture delivers an efficient bi-directional highway for platform data movement

Compared to Xeon 5500/5600:

- Lower L3 cache latency (~20%)
- Higher bandwidth between cores, L3, Memory & I/O
- Up to 8x more L3 to core bandwidth

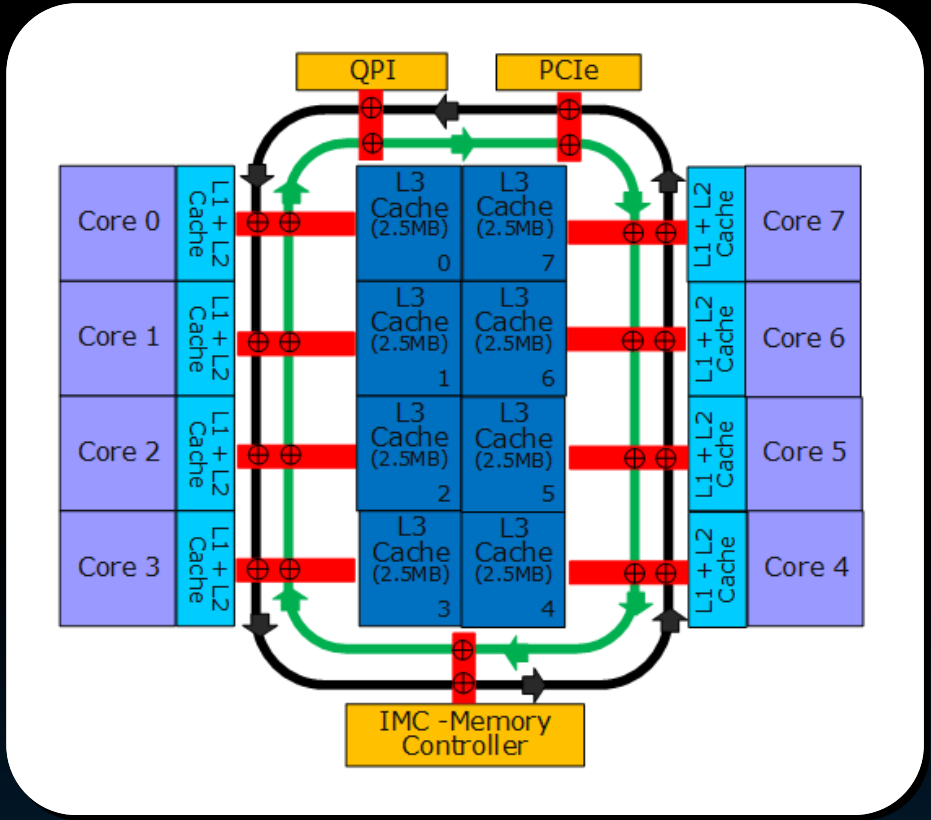

Higher performance starts with The Ring!

Intel® Advanced Vector Extensions (Intel® AVX)

2X Vector Width → 256-bit vector extension to SSE

- Intel® AVX extends all 16 XMM registers to 256 bits
- Intel® AVX works on either
  - The whole 256-bits
  - The lower 128-bits (like existing SSE instructions)
    - A drop-in replacement for all existing scalar/128-bit SSE instructions
- The new state extends/overlays SSE
- The lower part (bits 0-127) of the YMM registers is mapped onto XMM registers
- Intel® microarchitecture (Sandy Bridge) targets a full-performance first implementation
  - 256-bit Multiply, Add and Shuffle engines (2X today)
  - 2nd load port
  - New Operations to enhance vectorization
    - Broadcasts
    - Masked load & store

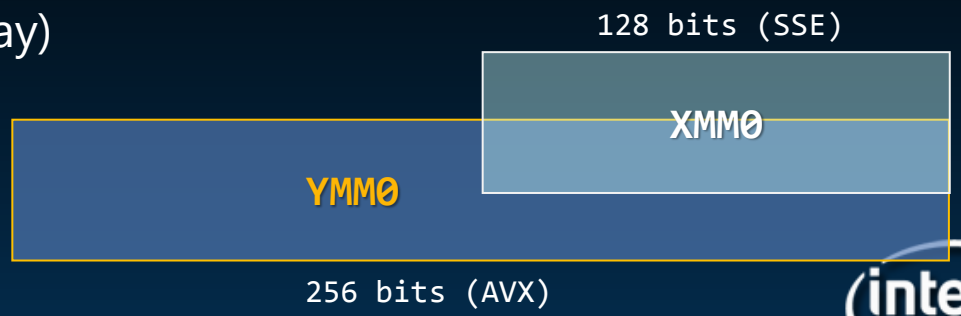

Nehalem Core Micro-architecture

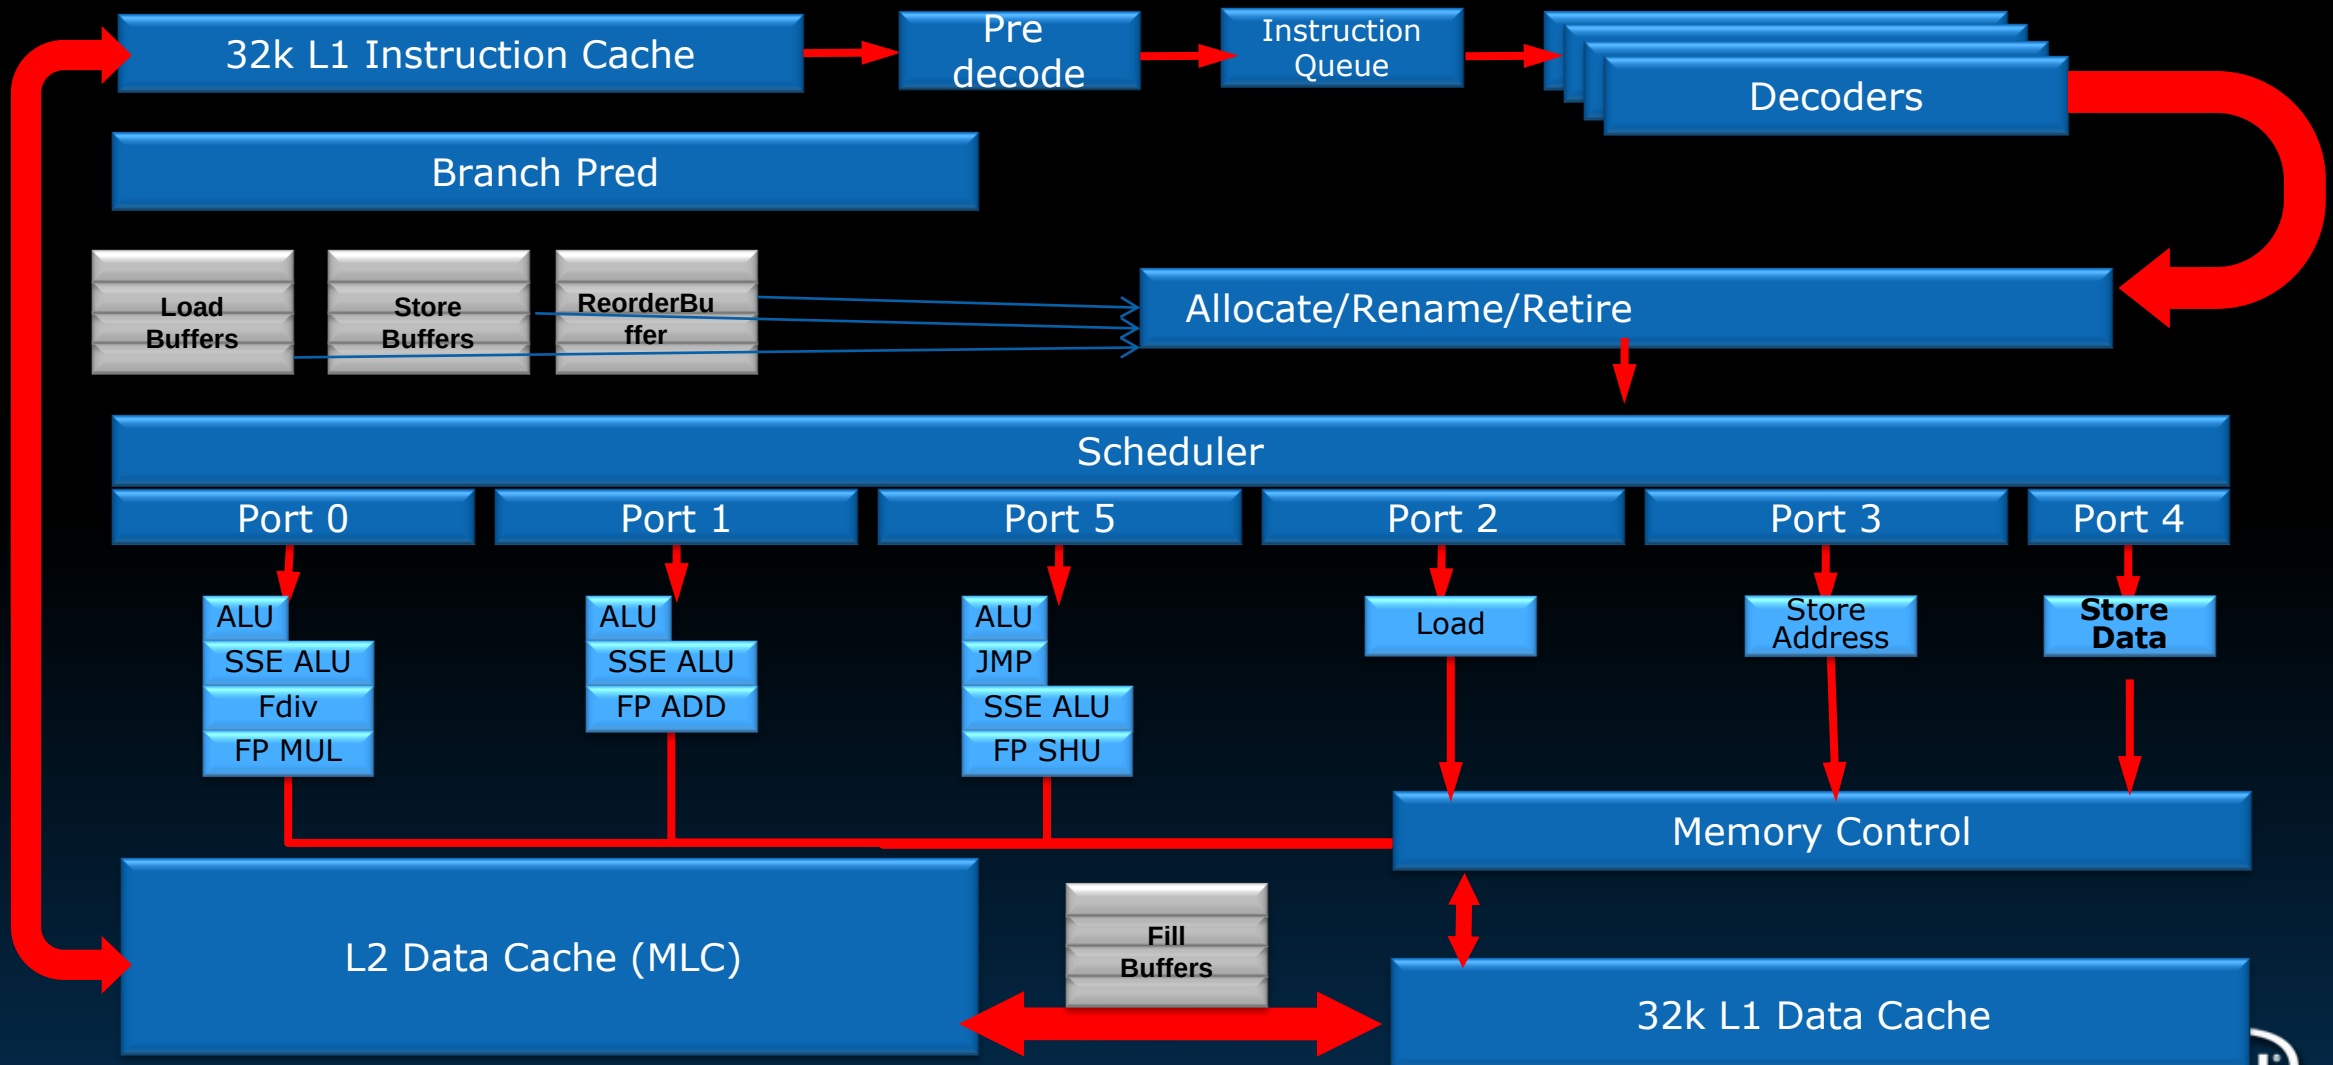

Nehalem - Intel® Core 2 Microarchitecture

SSE = Streaming SIMD Extentions

Sandy Bridge Core Micro-architecture

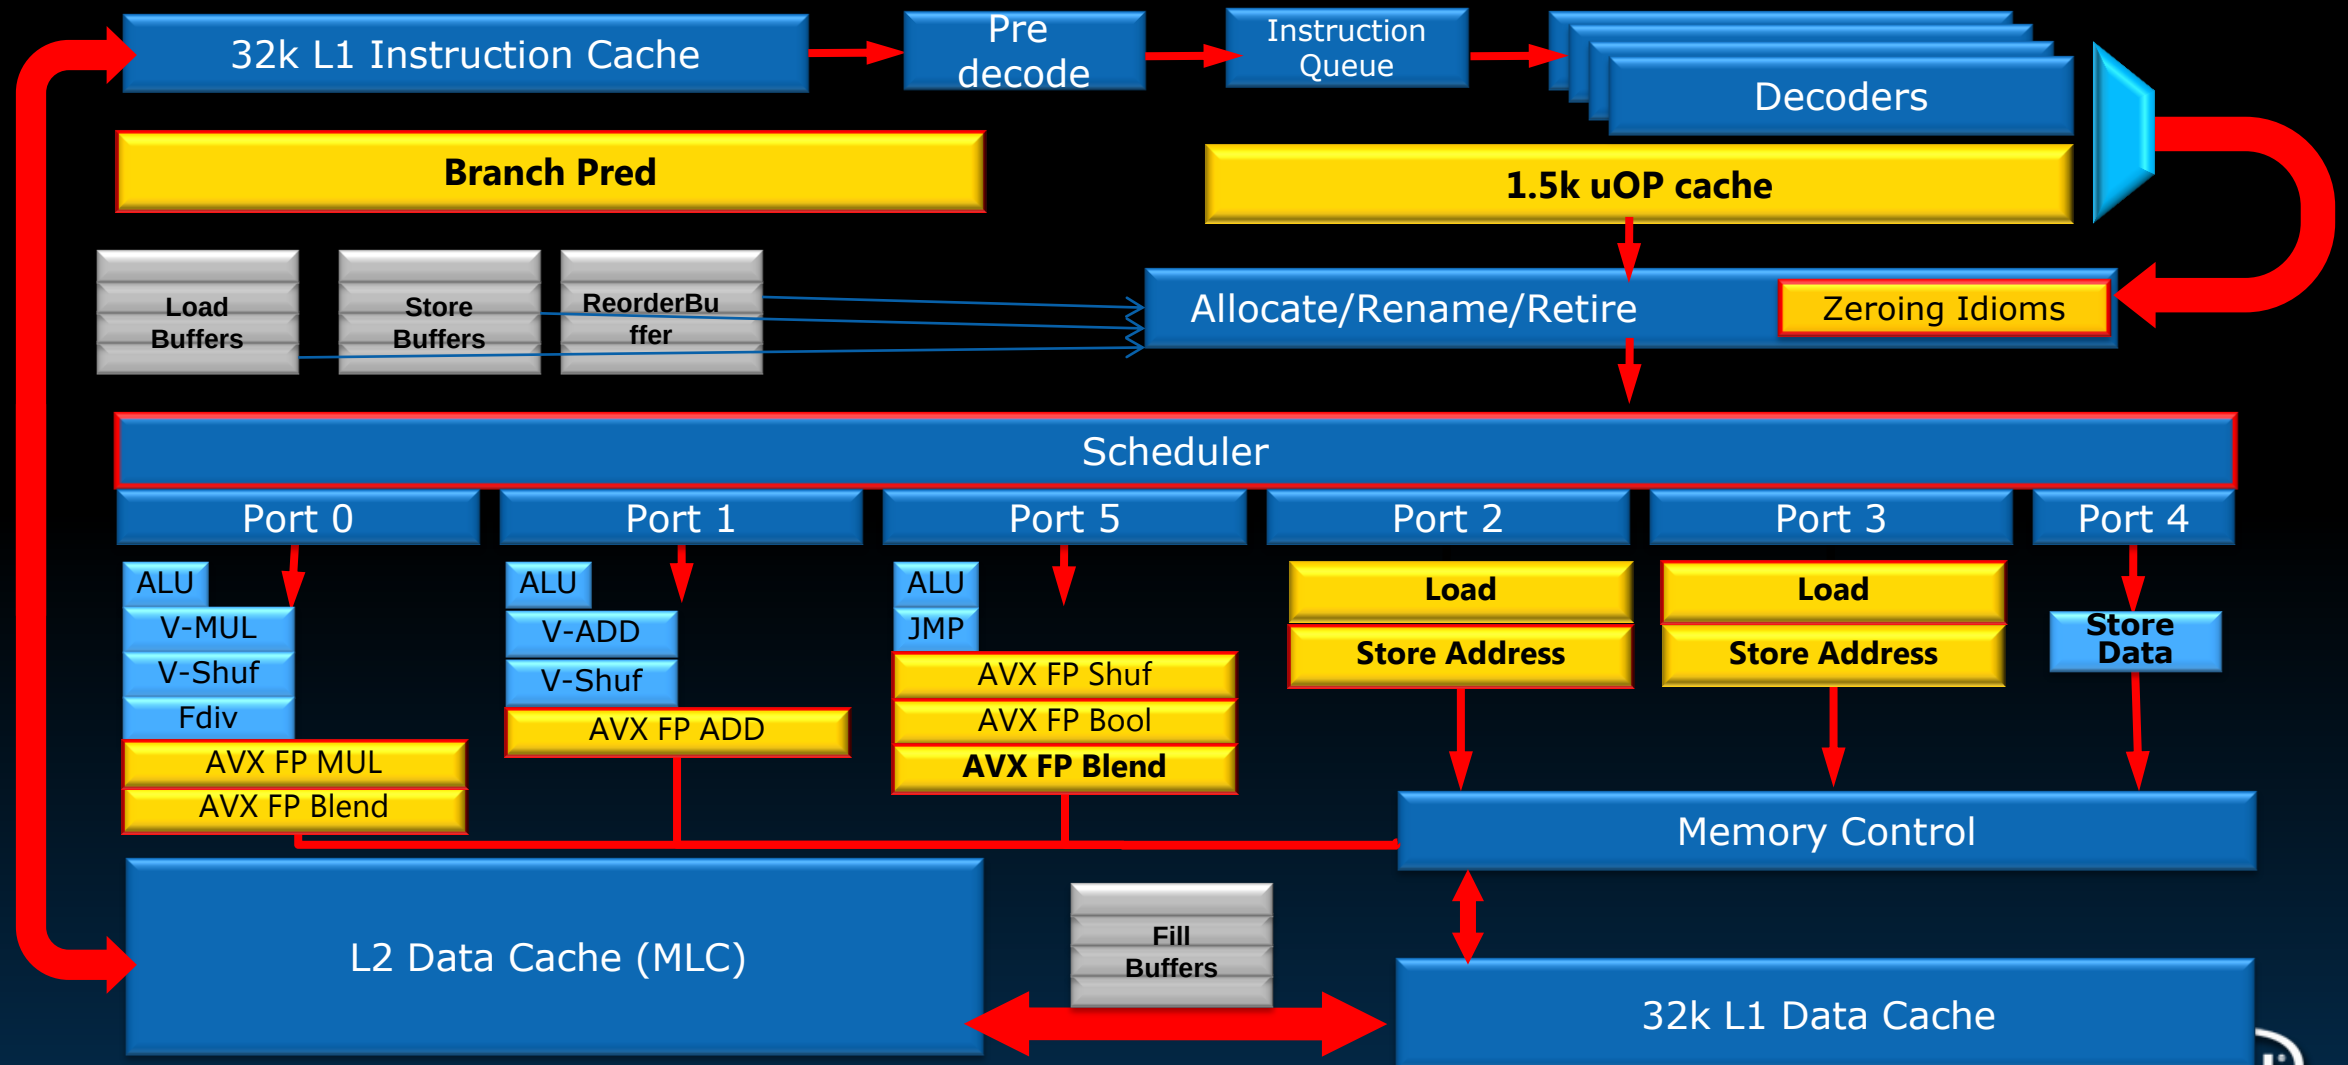

Sandy Bridge - Intel® Next Generation Microarchitecture

AVX= Intel® Advanced Vector Extensions (Intel® AVX)

Integrated PCIe: Latency Benefit

Xeon® 5600

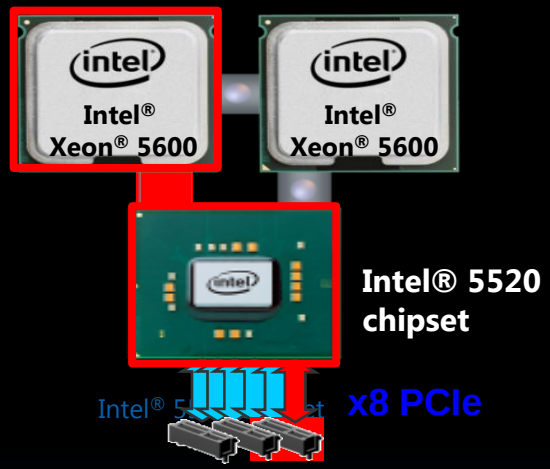

PCIe Latency
340ns

Xeon® E5 (Sandy Bridge)

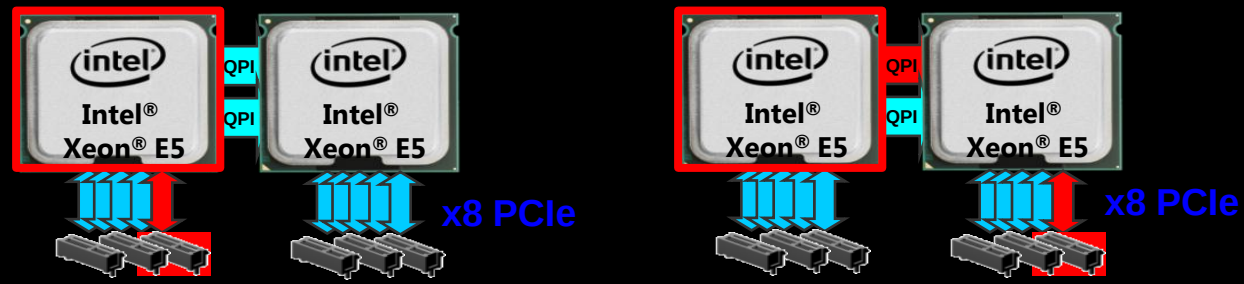

Local PCI Latency
255ns

25% reduction

Remote PCI Latency
320ns

5% reduction

Up to 25% reduction in PCIe latency with E5-2400/2600 Processors

Lower is better
Source: Intel internal measurements Oct 2011

Early Sandy Bridge-EP Performance

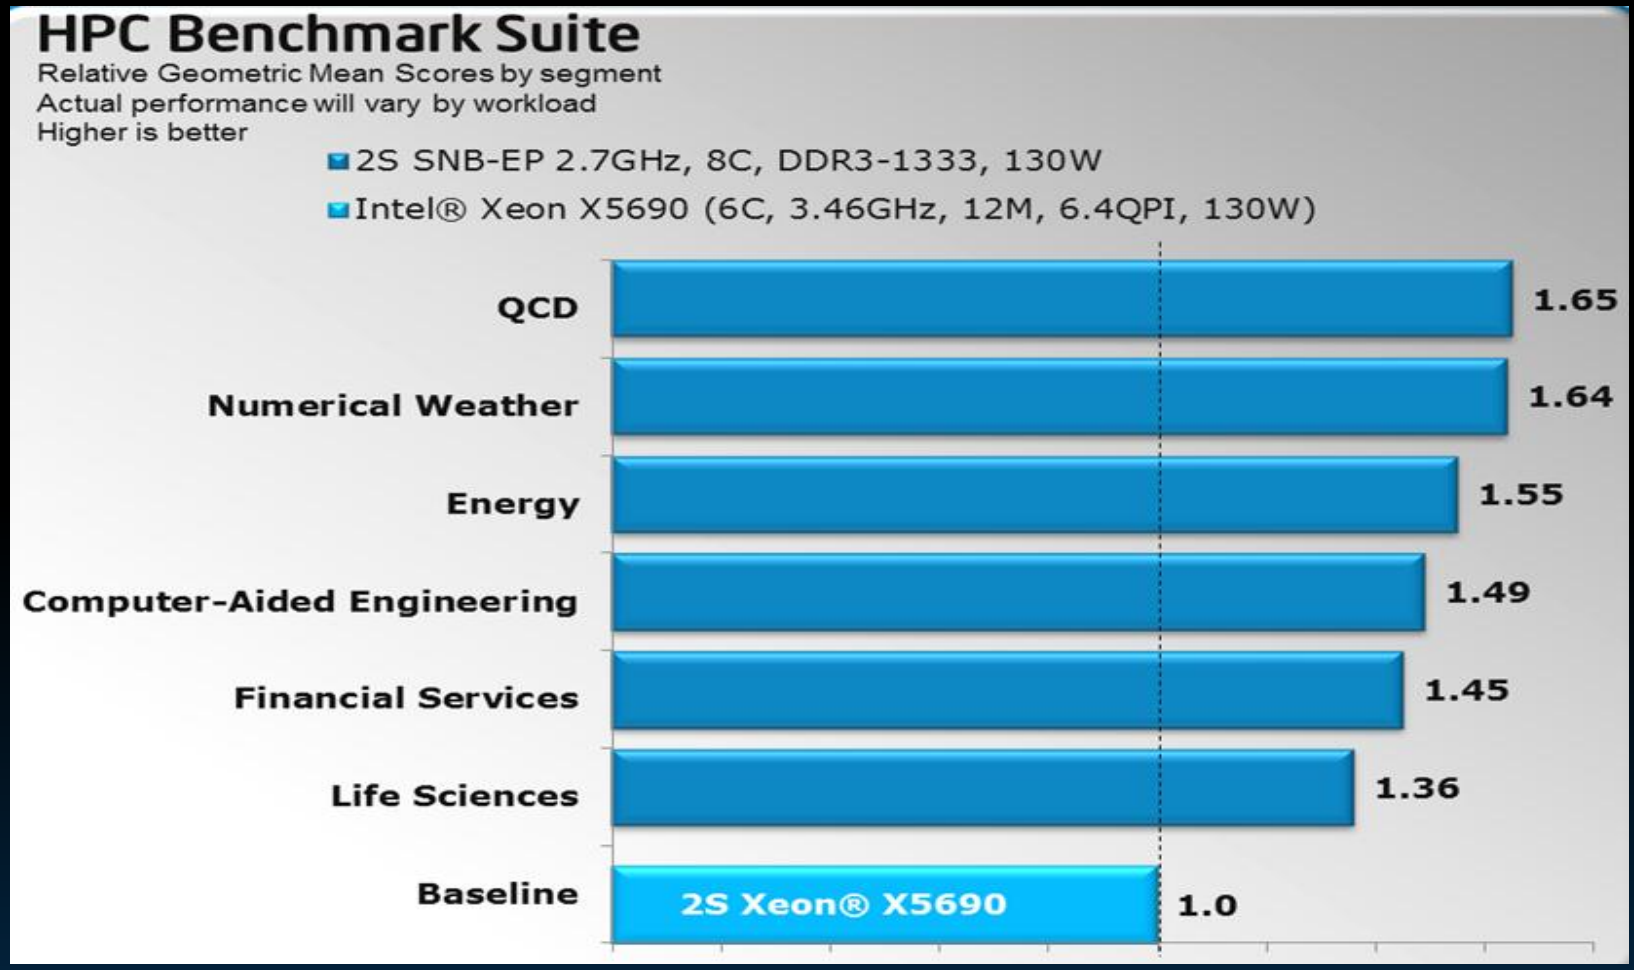

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may 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. Configurations: Intel Internal measurements February 2011, See backup for configuration details. For more information go to http://www.intel.com/performance . Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance. Copyright © 2011, Intel Corporation.

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

2- to 4-Socket Scaling Performance Summary

Intel® Xeon® Processor E5-4600 Product Family

Top-bin 2S Intel® Xeon® processor E5-2690 (8C, 2.9 GHz) vs. Top-bin 4S Intel® Xeon® processor E5-46xx (8C, 2.7 GHz)

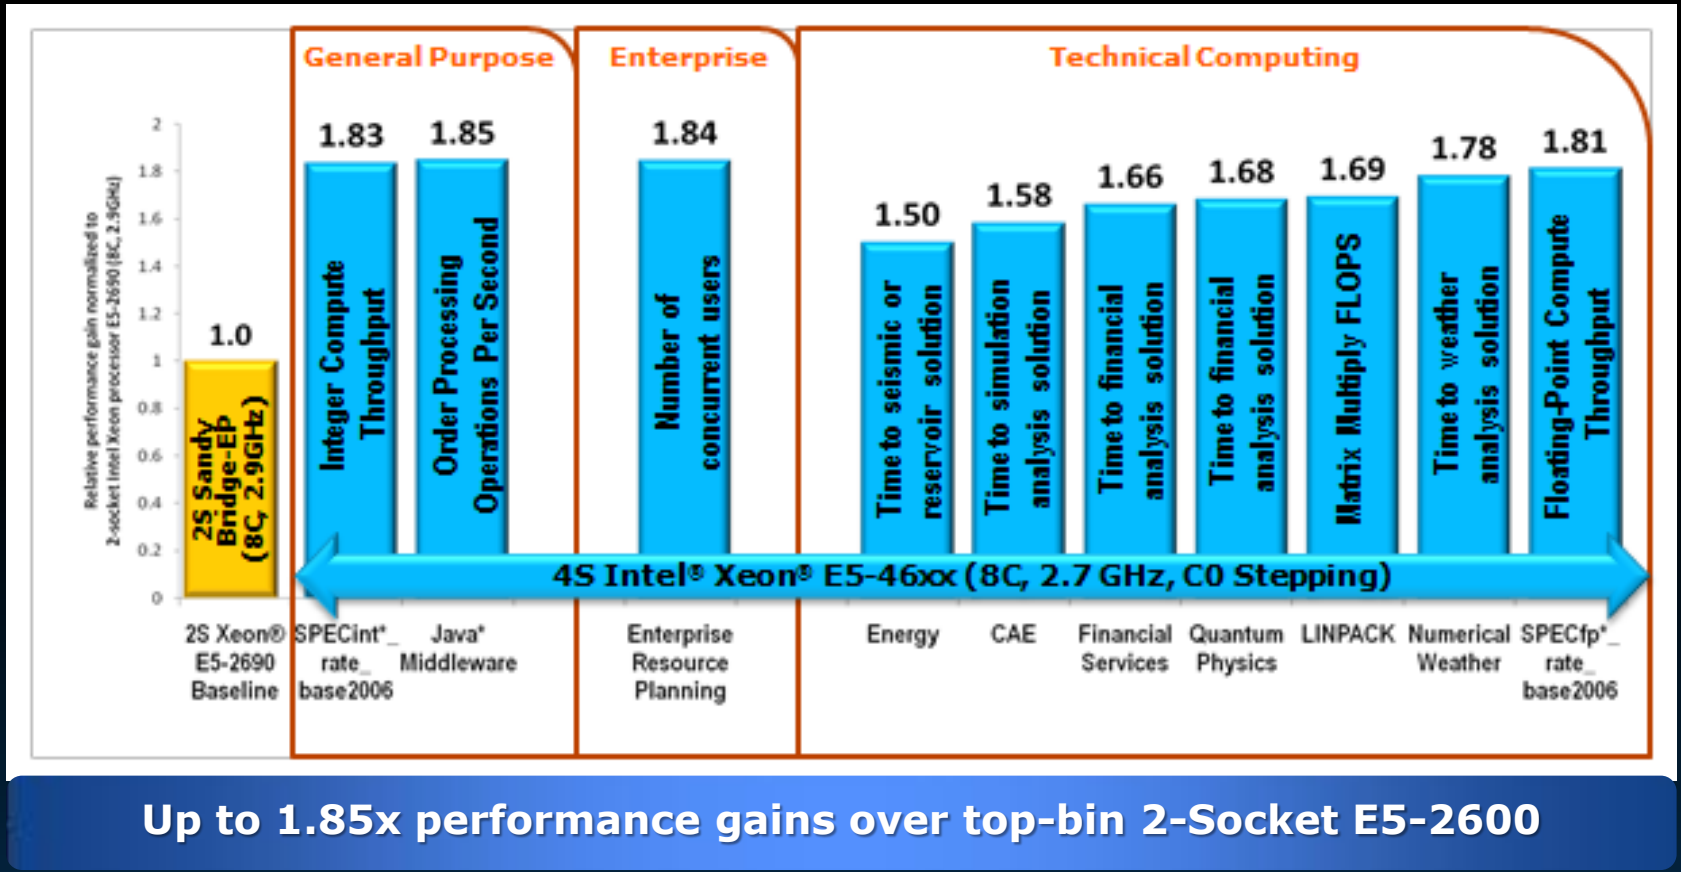

Up to 1.85x performance gains over top-bin 2-Socket E5-2600

Results that have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance. 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.

Intel® Xeon®

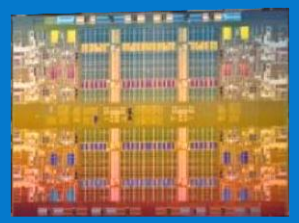

Foundation of HPC Performance
Suited for full scope of workloads

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

Focus on fast single core/thread performance with “moderate” number of cores

Intel® Xeon® Phi™

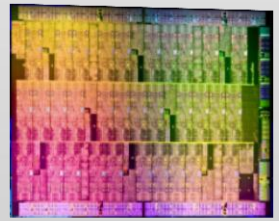

Performance and performance/watt optimized for highly parallelized compute intensive workloads

Common software tools with Xeon enabling efficient application readiness and performance tuning

IA extension to Many-Core

Lots of cores/threads with wide SIMD

[die size not to scale]

MIC: Knights Corner – the 1st Intel® Xeon® Phi™

- 22nm process, Production in 2012
- > 50 cores
- SIMD instructions
- 8GB+ of GDDR5 RAM
- Early Si delivers 1TFLOPS sustained on DGEMM and 1TFLOPS HPL in a node
- Runs Linux, IP addressable, common source code with CPU

Knights Corner node >1TF HPL

Intel® Xeon® Phi™ (Intel® MIC architecture) hits another performance milestone on the road to launch:

<u>1 Teraflop HPL (Linpack)!</u>

System configuration:

The demonstration system features dual Intel® Xeon® E5 processors, 64GB of DDR3 memory, and 1 Knights Corner coprocessor

>1 TFLOPs Linpack (HPL) in a node

C/C++, FORTRAN

OpenMP, MPI, …

Same Comprehensive Set of SW Tools

Application Source Code Builds with a Compiler Switch

Established HPC Operating System

Intel® Xeon®

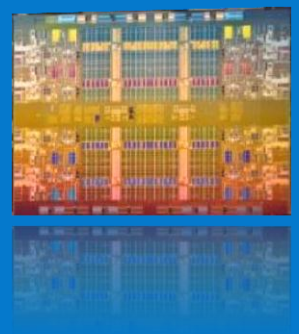

Intel® Xeon® Phi™

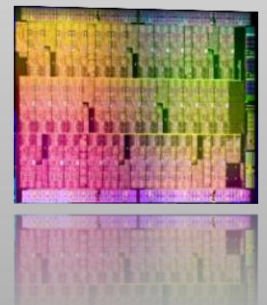

[die size not to scale]

A Very Simple Arithmetic Example

using IEEE 64-bit DP-F.P.

X_1	X_2	X_3	X_4	X_5	$SUM(X_1:X_5)$
1.00E+21	-1.00E+21	17	-10	130	137.00 ✔

Source: Ulrich Kulisch, *Computer Arithmetic and Validity,* de Gruyter Studies in Mathematics 33 (2008), p. 250

A Very Simple Arithmetic Example

using IEEE 64-bit DP-F.P.

X_1	X_2	X_3	X_4	X_5	SUM(X_1:X_5)	
1.00E+21	17	-10	130	-1.00E+21	0.00	✗✗
1.00E+21	-1.00E+21	17	-10	130	137.00	✓

Source: Ulrich Kulisch, *Computer Arithmetic and Validity,* de Gruyter Studies in Mathematics 33 (2008), p. 250

A Very Simple Arithmetic Example

using IEEE 64-bit DP-F.P.

X_1	X_2	X_3	X_4	X_5	SUM(X_1:X_5)	
1.00E+21	17	-10	130	-1.00E+21	0.00	✗ ✗
1.00E+21	-10	-1.00E+21	130	17	147.00	✗
1.00E+21	-1.00E+21	17	-10	130	137.00	✓

Source: Ulrich Kulisch, *Computer Arithmetic and Validity,* de Gruyter Studies in Mathematics 33 (2008), p. 250

A Very Simple Arithmetic Example

using IEEE 64-bit DP-F.P.

X_1	X_2	X_3	X_4	X_5	SUM(X_1:X_5)	
1.00E+21	17	-10	130	-1.00E+21	0.00	✗✗
1.00E+21	-10	-1.00E+21	130	17	147.00	✗
1.00E+21	-1.00E+21	17	-10	130	137.00	✓
1.00E+21	17	130	-1.00E+21	-10	-10.00	✗✗✗

Source: Ulrich Kulisch, *Computer Arithmetic and Validity,* de Gruyter Studies in Mathematics 33 (2008), p. 250

A Very Simple Arithmetic Example

using IEEE 64-bit DP-F.P.

X_1	X_2	X_3	X_4	X_5	SUM(X_1:X_5)	
1.00E+21	17	-10	130	-1.00E+21	0.00	✗✗
1.00E+21	-10	130	-1.00E+21	17	17.00	✗✗
1.00E+21	17	-1.00E+21	-10	130	120.00	✗
1.00E+21	-10	-1.00E+21	130	17	147.00	✗
1.00E+21	-1.00E+21	17	-10	130	137.00	✓
1.00E+21	17	130	-1.00E+21	-10	-10.00	✗✗✗

Source: Ulrich Kulisch, *Computer Arithmetic and Validity,* de Gruyter Studies in Mathematics 33 (2008), p. 250

"Results can be satisfactory, inaccurate or completely wrong. Neither the computation itself nor the computed result indicate which one of the three cases has occurred."

History: Software
How old is HPC software?

Code	Area	Date
NASTRAN	Structures	1968
Spice	E-Cad	1972
Pam-Crash	Structures	1978
UKMO Unified Model	Weather	1990
PETSc	Solvers	1991
LAPACK	Solvers	1992
NWCHEM	Chemistry	1995
WRF	Weather	2000

- **Code lasts much longer than hardware**
- **We must support old code on new hardware**

History: Software
How old are languages?

Language	Date
Fortran	1966 (FORTRAN 66)
C	1978 (K&R)
C++	1985 (C++ Programming Language)
MPI	1994
OpenMP	1997

- Languages that work have a long life
  - Investment in code
  - Investment in brain-cells
- All have open specifications & many implementations
  - Formal standards (C, C++, Fortran)
  - Open industry standards (MPI, OpenMP)
- We must support old languages on new hardware

Heterogeneous Programming

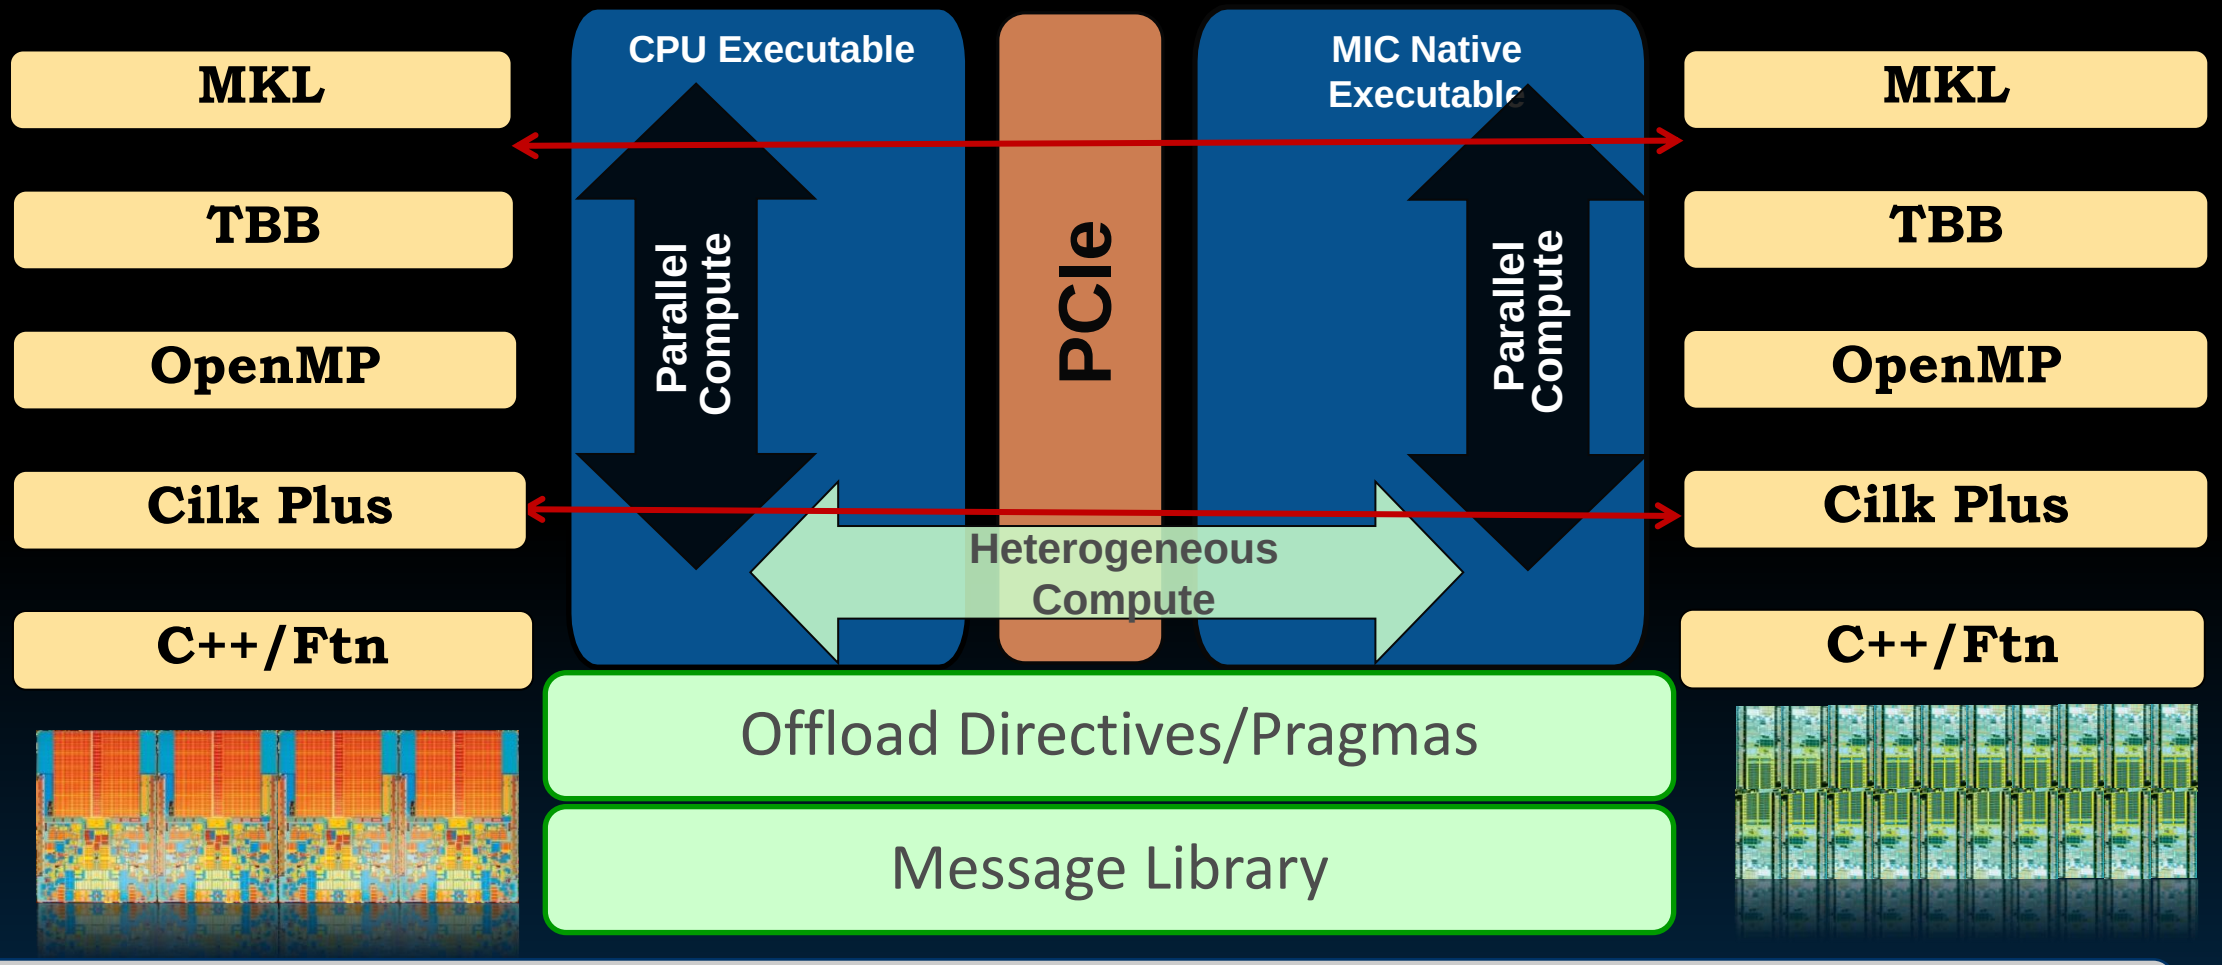

Parallel programming is the same on Intel® MIC and CPU

Offload Directives

	C/C++ Syntax	Semantics
New offload pragma	`#pragma offload <clauses> <statement>`	Execute next statement on Intel® MIC (which could be an OpenMP parallel construct)
Compile function for CPU and MIC	`__attribute__ ((target (MIC)))`	Compile function for CPU and Intel MIC target

	Fortran Syntax	Semantics
New offload directive	`!dir$ omp offload <clauses>`	Execute next OpenMP parallel construct on Intel® MIC
	`!dir$ offload <clauses>`	Execute next statement (function call) on Intel MIC
Compile function for CPU and Intel MIC	`!dir$ attributes offload:<MIC> :: <rtn-name>`	Compile function for CPU and Intel MIC target

Offload Directives (contd.)

Variables restricted to scalars, structs, arrays and pointers to scalars/structs/arrays

Clauses / Modifiers	Syntax	Semantics
Target specification	`target ( name [: ] )`	Where to run construct
Inputs	`in (var-list modifiers`$_{opt}$`)`	Copy CPU to target
Outputs	`out (var-list modifiers`$_{opt}$`)`	Copy target to CPU
Inputs & outputs	`inout (var-list modifiers`$_{opt}$`)`	Copy both ways
Non-copied data	`nocopy (var-list modifiers`$_{opt}$`)`	Data is local to target
Modifiers		
Specify pointer length	`length (element-count-expr)`	Copy that many pointer elements
Control pointer memory allocation	`alloc_if ( condition )` `free_if ( condition )`	Allocate/free new block of memory for pointer if condition is TRUE

Offload Directive Examples: OpenMP, Intel® Cilk™ Plus

C/C++ OpenMP

```
#pragma offload target (mic)
#pragma omp parallel for
  reduction(+:pi)
 for (i=0; i<count; i++)  {
   float t = (float)((i+0.5)/count);
   pi += 4.0/(1.0+t*t);
 }
 pi /= count;
```

Fortran OpenMP

```
!dir$ omp offload target(mic)
!$omp parallel do
   do i=1,10
     A(i) = B(i) * C(i)
   enddo
```

C/C++ Cilk

```
#pragma offload target(mic)
_Cilk_for (int i=0; i<count;  i++)
{
    a[i] = b[i] * c + d;
}
/*          (in a forthcoming update)       */
```

Intel Parallel & HPC Programming

Intel® FORTRAN Compiler

Fortran Language

Intel® C/C++ Compiler

C/C++ Language

Intel® Parallel Building Blocks

Intel® Cilk Plus

C/C++ Language Extensions to simply Parallelism

Open sourced, Also an Intel product

Intel® Threading Building Blocks

Widely used C++ Template Library for Parallelism

Open sourced, Also an Intel product

Established Standards

MPI

PGAS

Co-Array FTN

OpenMP*

OpenCL*

Domain-Specific Libs

Intel® Integrated Performance Primitives (IPP)

Intel® Math Kernel Library (MKL)

Research and Exploration

Intel® Concurrent Collections

Offload Extensions

Intel® Array Building Blocks

Intel® SPMD Parallel Compiler

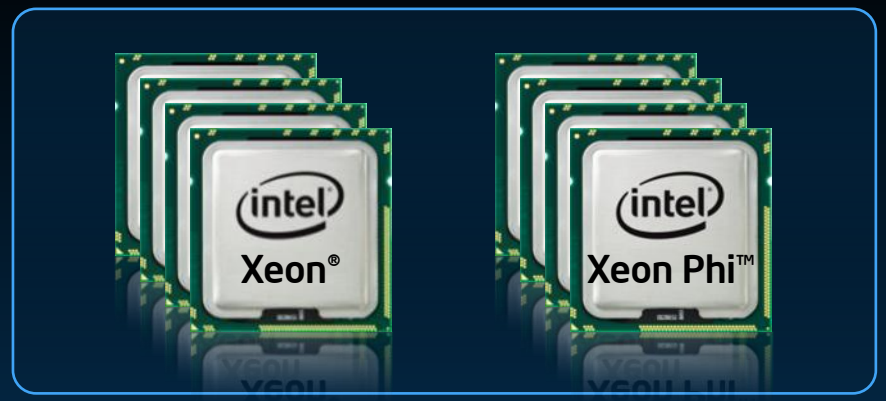

Potential future options subject to change without notice.

Intel Development Tools for HPC

Leading developer tools for performance on nodes and clusters

Advanced Performance

C++ and Fortran Compilers, MKL/IPP Libraries & Analysis Tools for Windows*, Linux* developers on
IA multi-core node

Distributed Performance

MPI Cluster Tools, with C++ and Fortran Compiler and MKL Libraries, and analysis tools for Windows*, Linux* developers on
IA clusters

Optimized Intel Libraries

Math	Power, Root	Trig	Hyper	Tounding	Exp, Log Special
Add	Pow	Cos	Cosh	Floor	Exp
Sub	Powx	Sin	Sinh	Ceil	Expm1
Div	Pow2o3	SinCos	Tanh	Round	Ln
Sqr	Pow3o2	Cis	Asinh	Trunc	Log10
Mul	Sqrt	Tan	Acosh	Rint	Log1p
Conj	Cbrt	Acos	Atanh	NearbyInt	Erf
MulByConj	InvSqrt	Asin	-	Modf	Erfc
Abs	InvCbrt	Atan			ErfInv
Inv	Hypot	Atan			

Random-Number Generators
Pseudo-random
Multiplicative Congruential 59-bit
Multiplicative Congruential 31-bit
Multiple Recursive
Feedback shift register
Wichman-Hill
Mersenne Twister 19937
Mersenne Twister 2203
Quasi-random
Sobol
Niederreiter

Probability Distributions	
Continuous	Discrete
Uniform	Uniform
Gaussian	UniformBits
GaussianMV	Bernoulli
Exponential	Geometric
Laplace	Binomial
Weibull	Hypergeometric
Cauchy	Poisson PTPE
Rayleigh	Poisson Norm
Lognormal	Poisson V
Gumbel	Negative Binomial
Gamma	–
Beta	–

Intel® MKL

Math Kernel Library

- Science, Engineering and Financial applications oriented
- Incl. BLAS, LAPACK, ScaLAPACK, Sparse Solvers, Fast Fourier Transforms, Vector Math

Intel® IPP

Integrated Performance Primitives

- Multimedia, Data Processing, and Communications applications oriented
- Cryptography and String Processing

TACC TEXAS ADVANCED COMPUTING CENTER
Powering Discoveries That Change The World
THE UNIVERSITY OF TEXAS AT AUSTIN

Home About Resources User Services Research & Development Partnerships Education & Outreach News

"Stampede's" Comprehensive Capabilities to Bolster U.S. Open Science Computational Resources

Texas Advanced Computing Center's new supercomputer to enable traditional HPC, data-intensive computing, and scientific visualization for nation's scientists

AUSTIN, Texas—The Texas Advanced Computing Center (TACC) at The University of Texas at Austin today announced that it will deploy and support a world-class supercomputer with comprehensive computing and visualization capabilities for the open science community, as part of the National Science Foundation's (NSF) "eXtreme Digital" (XD) program.

The new system, called Stampede, will be built by TACC in partnership with Dell and Intel to support for four years the nation's scientists in addressing the most challenging scientific and engineering problems. NSF is providing $27.5 million immediately and Stampede is expected to be up and running in January 2013. The estimated investment will be more than $50 million over four years, the Stampede project may be renewed in 2017, which would enable four additional years of open science research on a successor system.

"Stampede will be one of the most powerful systems in the world and will be uniquely comprehensive in its technological capabilities," said TACC Director Jay Boisseau. "Many researchers will leverage Stampede not only for massive computational calculations, but for all of their scientific computing, including visualization, data analysis, and data-intensive computing. We expect the Stampede system to be a model for supporting petascale simulation-based science and data-driven science."

When Stampede is deployed in 2013, it will be the most powerful system in the NSF XD environment, currently the most advanced, comprehensive, and robust collection of integrated digital resources and services enabling open science research in the world. As a critical part of XD, the Extreme Science and Engineering Discovery Environment (XSEDE) consortium comprising more than a dozen universities and two research laboratories, has now replaced the TeraGrid as the integrating fabric for the bulk of the NSF's high- end digital resources. Researchers from any U.S. open science institution can apply to use Stampede for a variety of novel scientific and educational activities through the XSEDE project.

News
Press Releases
Feature Stories
Events
TACC "In The News"
Case Studies
Annual Research Report
Subscribe

Quick Links
About TACC
Staff

Contact
Faith Singer-Villalobos
Public Relations
faith@tacc.utexas.edu

Watch a video of TACC Director Jay Boisseau discussing Stampede.

Stampede System Overview

"Stampede"

- 10 PFLOPS peak
- 272 TB memory
- 14 PB storage
- Deployment scheduled in 2013
- DELL System
- Intel® Xeon® E5 (Sandy Bridge-EP) processors
- Intel® MIC (Knights Corner) co-processors
- FDR InfiniBand (56Gb/s) cluster fabric

“If you are not scared …
your dreams are not big enough”

Intel's Plans For Exascale

Efficient Performance

Programming Parallelism

Extreme Scalability

Intel Exascale Plans for 2018+:
>100X Performance of today at
only 2X the Power of today's #1 System
Scaling today's (and future) Software Models …

All dates, data and figures are preliminary and are subject to change without any notice

Architecting for ExaScale

Needs a Multi-Disciplinary Approach

Power Management

Parallel Software

Reliability & Resiliency

Microprocessor

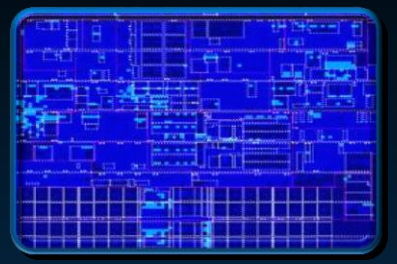

Memory & Storage

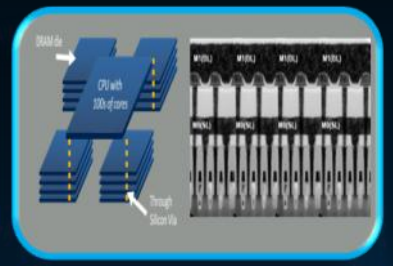

Interconnect

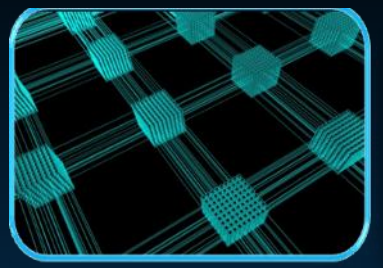

For illustration and concept only.

HPC Platform Power Consumption

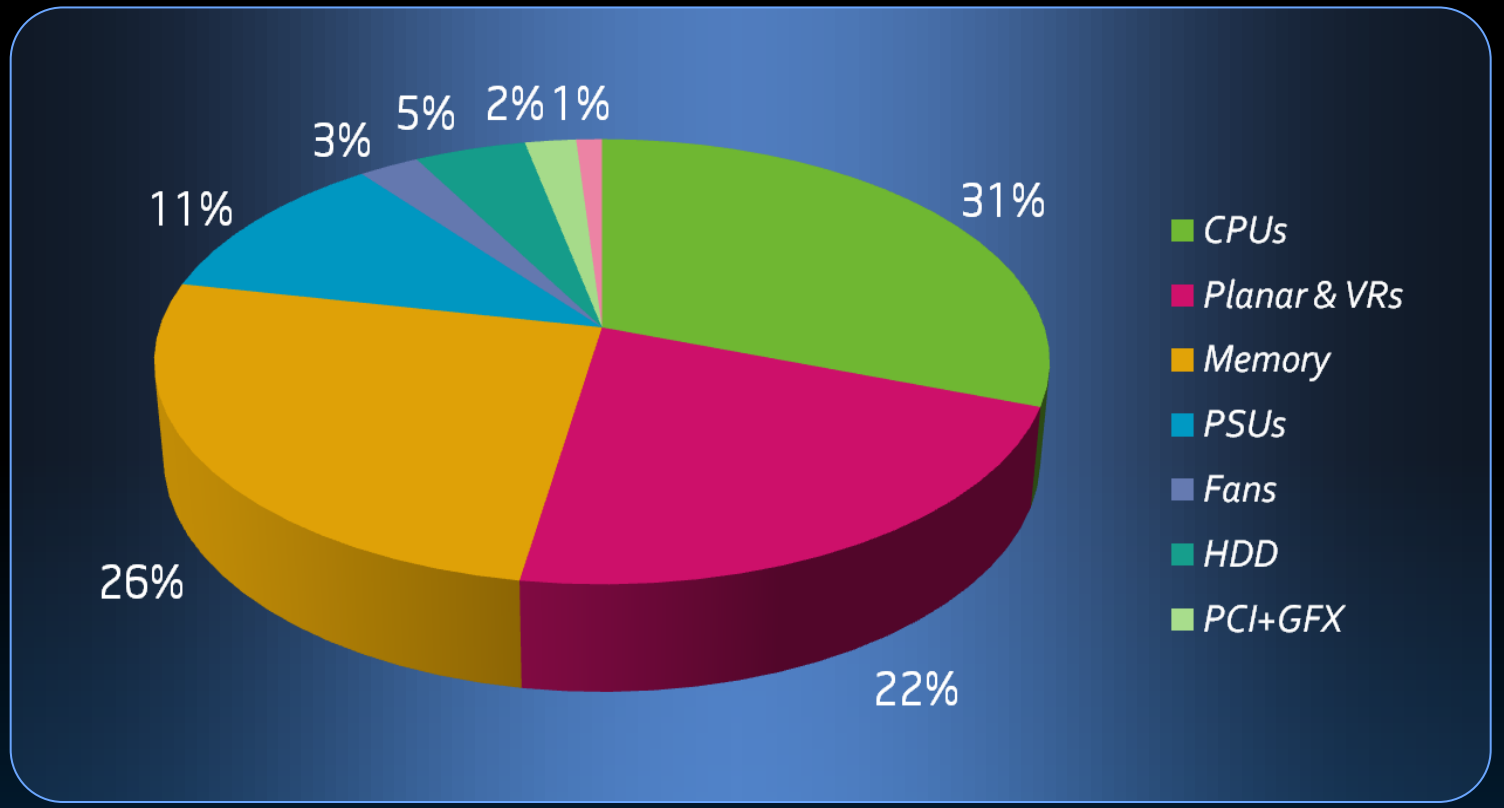

Data from P3 Jet Power Calculator, V2.0, DP 80W Nehalem
Memory – 48GB (12 x 4GB DIMMs) Single Power Supply Unit @ 230Vac

Need a platform view of power consumption: CPU, Memory and VR, etc.

Reduce Memory and Communication Power

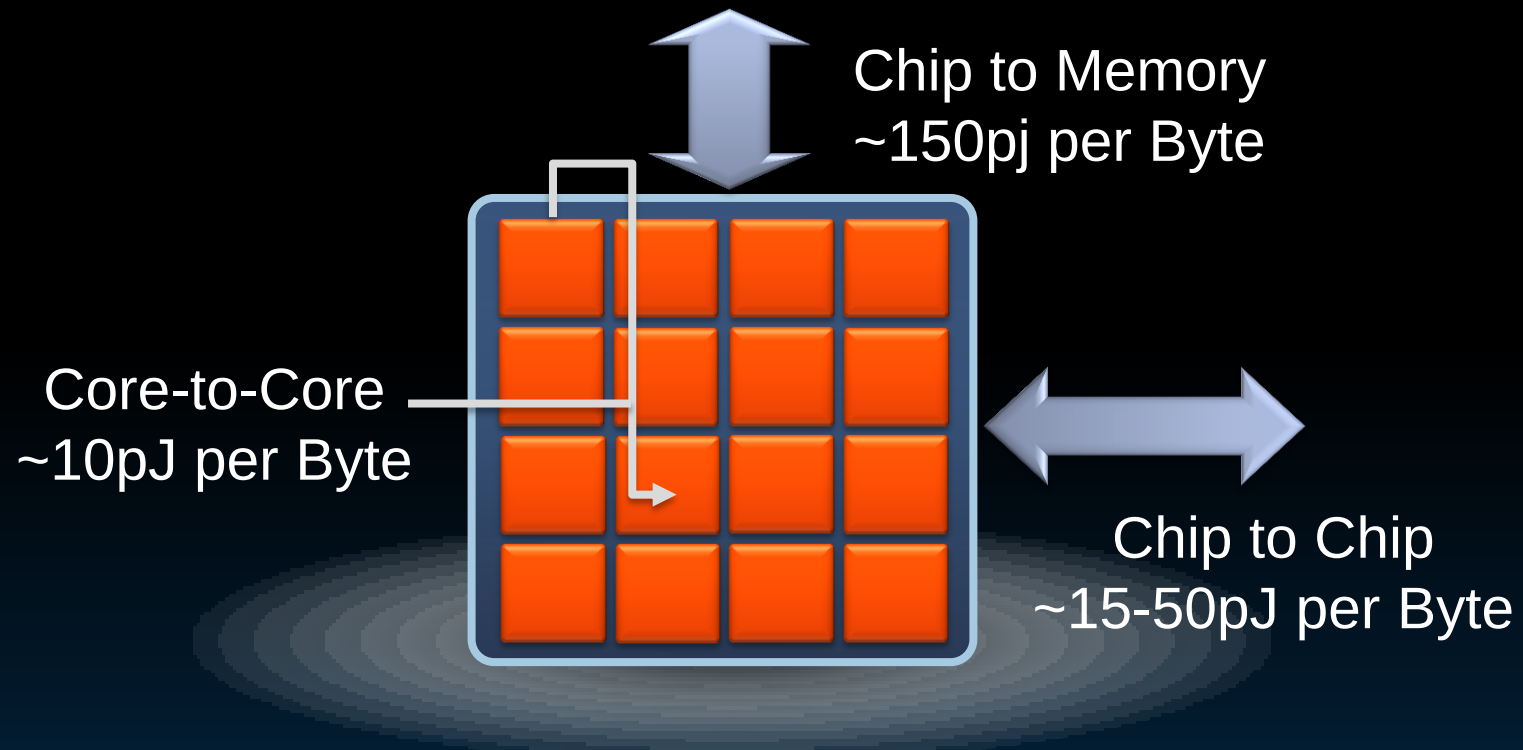

Data movement is expensive

Source: Intel

"Business As Usual" Isn't Trending Towards Exascale

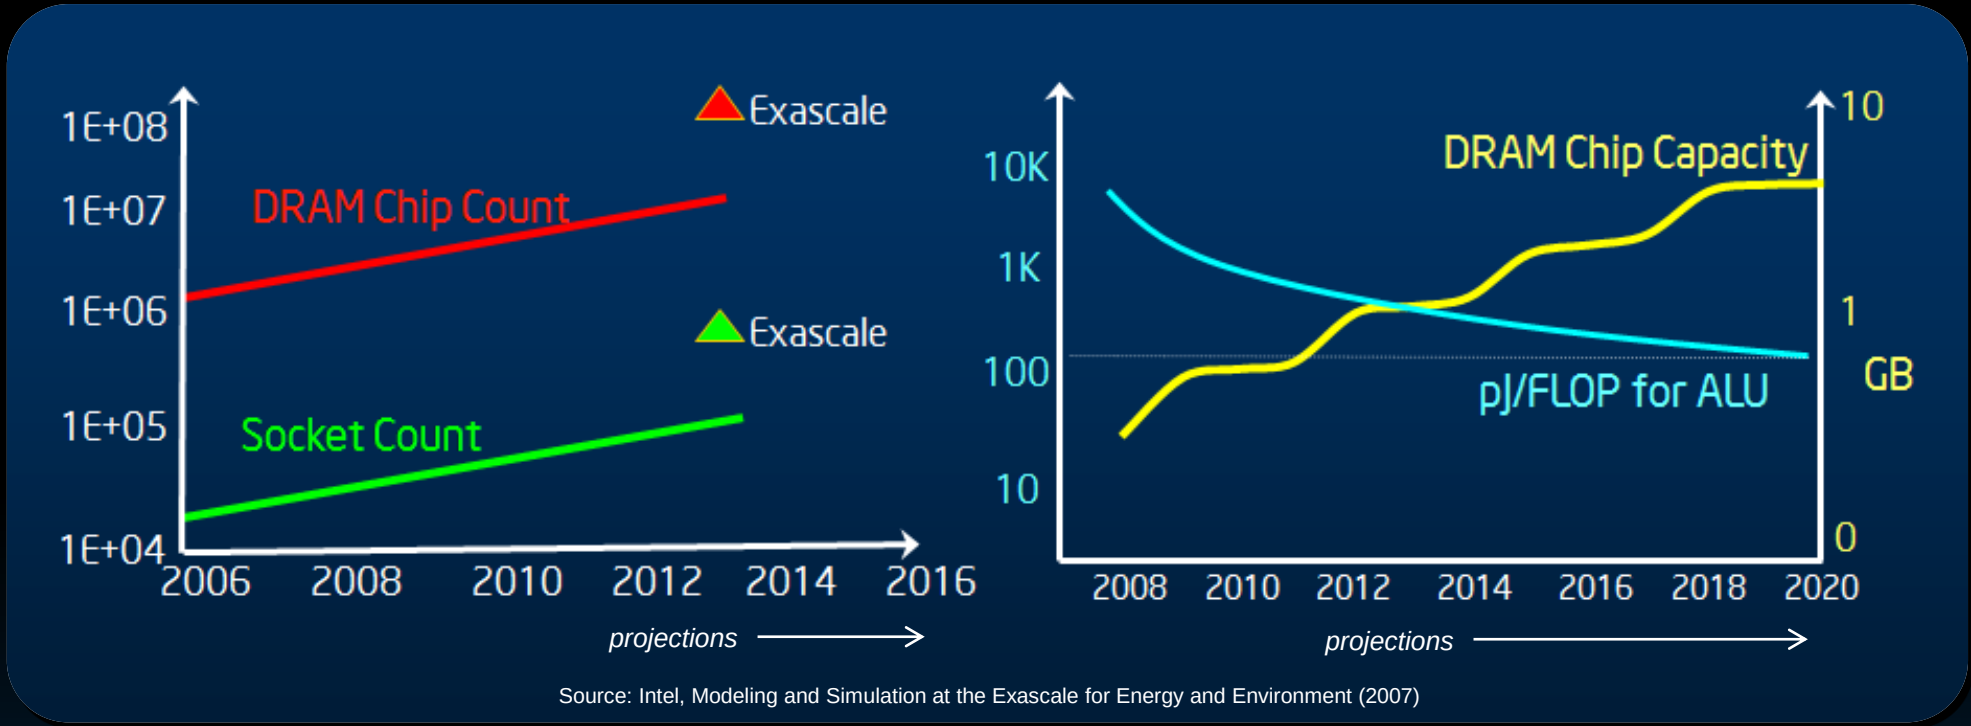

Source: Intel, Modeling and Simulation at the Exascale for Energy and Environment (2007)

- **Core count increase trend falls short of the exascale requirements**
- **Energy/Op for ALU operations alone is too high**
- **DRAM power is too high and Capacity/BW is too low**

→ Current Natural Progression of Technology is Insufficient

→ A Paradigm Shift in All Areas of Computing is Needed

Intel TeraScale Research Areas

MANY-CORE COMPUTING

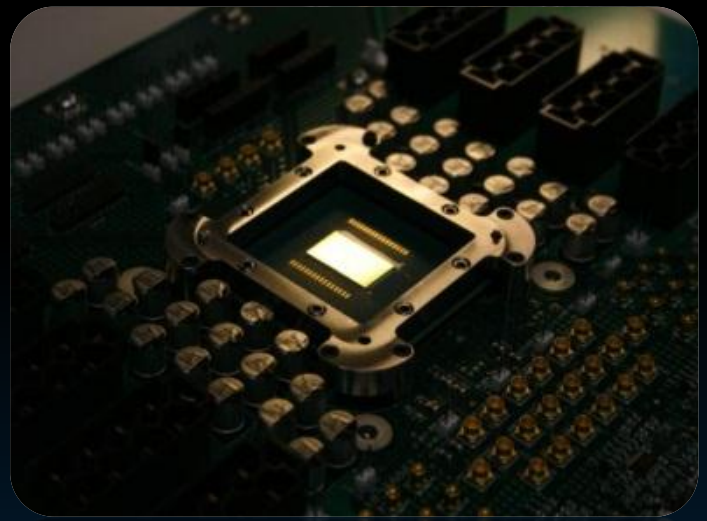

Teraflops of computing power

3D STACKED MEMORY

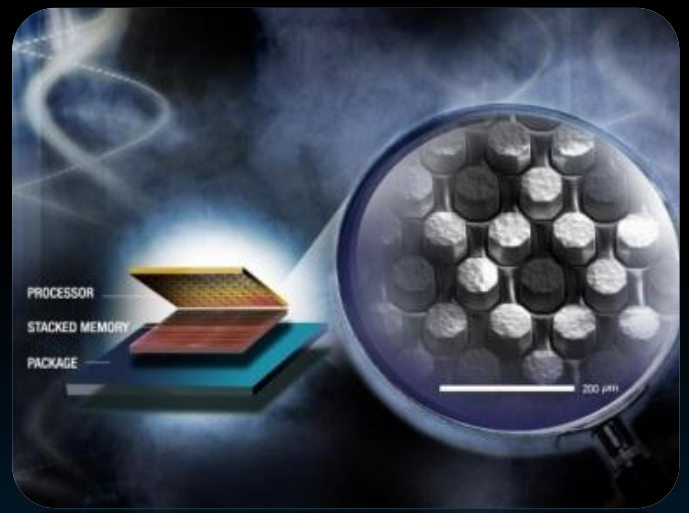

Terabytes of memory bandwidth

SILICON PHOTONICS

Terabits of I/O throughput

Future vision, does not represent real products.

Multi/Many-Core Architecture Research

*Intel **Tera-Scale** Computing Research Program: www.intel.com/go/terascale*

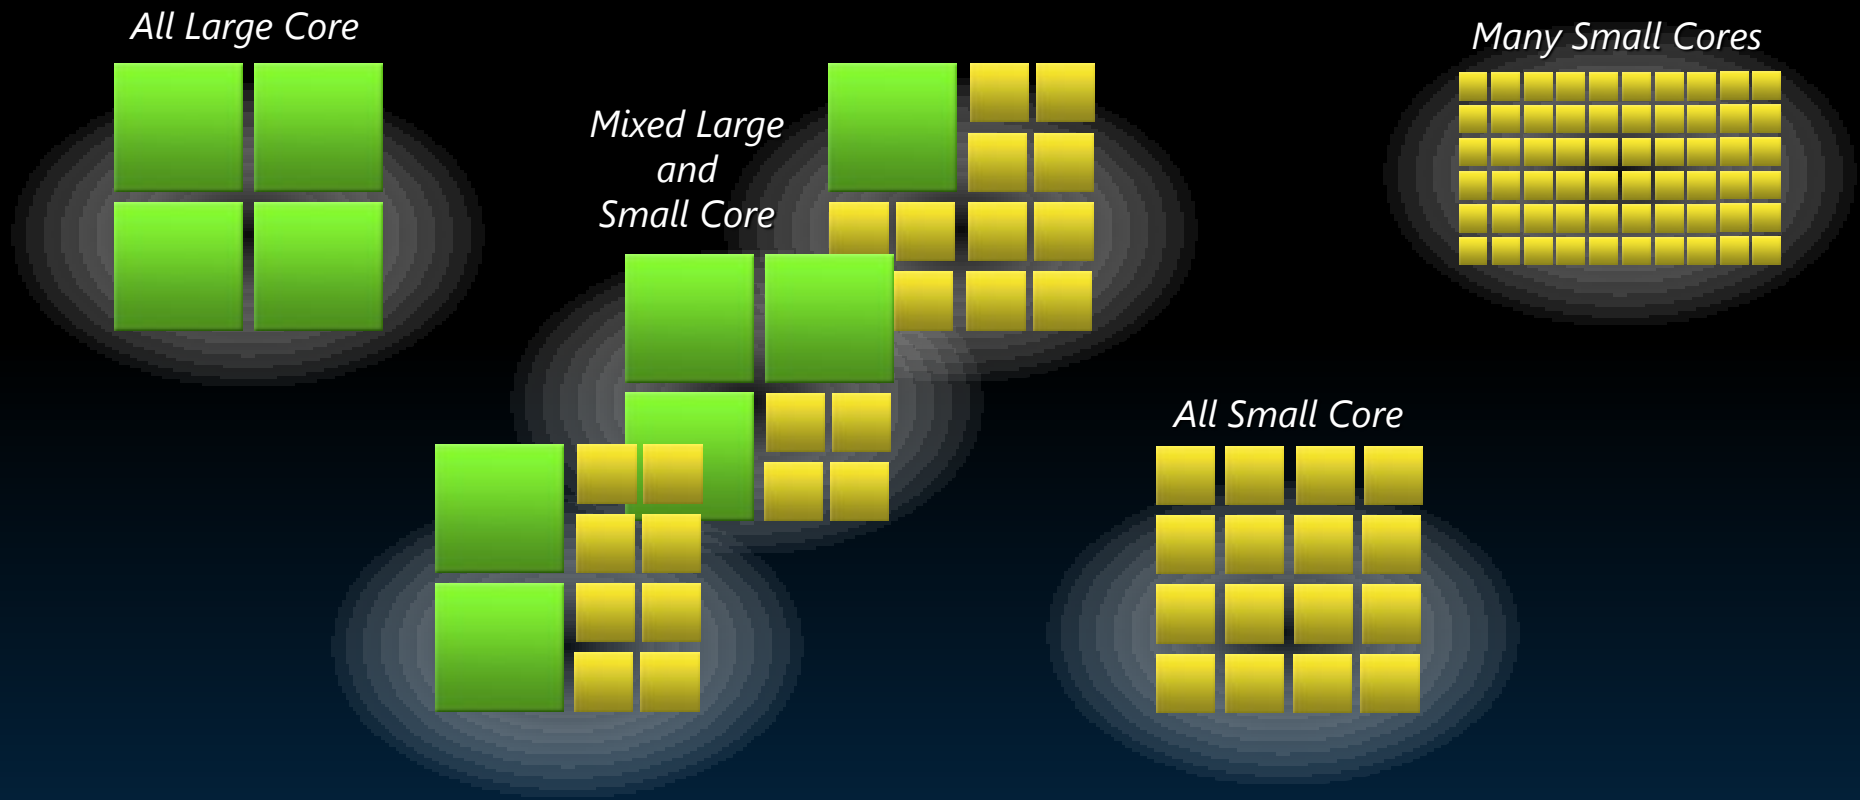

A (single) core might also have multiple HW-Threads

Note: some the above pictures don't represent any current or future products

A Theoretic Example

Processor A
100 thin cores

Processor B
90 thin cores
1 fat core, 2x faster than a thin core

Assume 10% of the application is serial (90% parallel)

What are the overall performance improvements/speed-ups?

Applying Amdahl's Law:

9.2x faster

16.7x faster

Source: A View of the Parallel Computing Landscape, Communications of the ACM, October 2009, Vol. 52, No. 10 (pp.56-67)

On-chip Interconnect Challenges and Research

Bandwidth, Link Bandwidth and Power

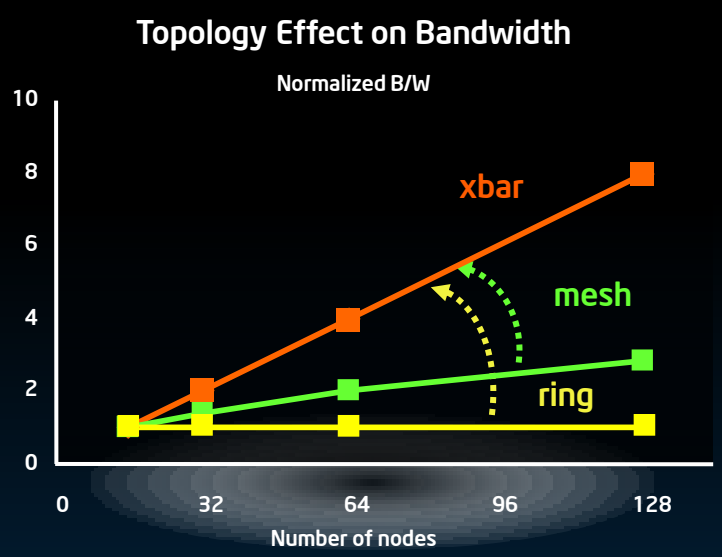

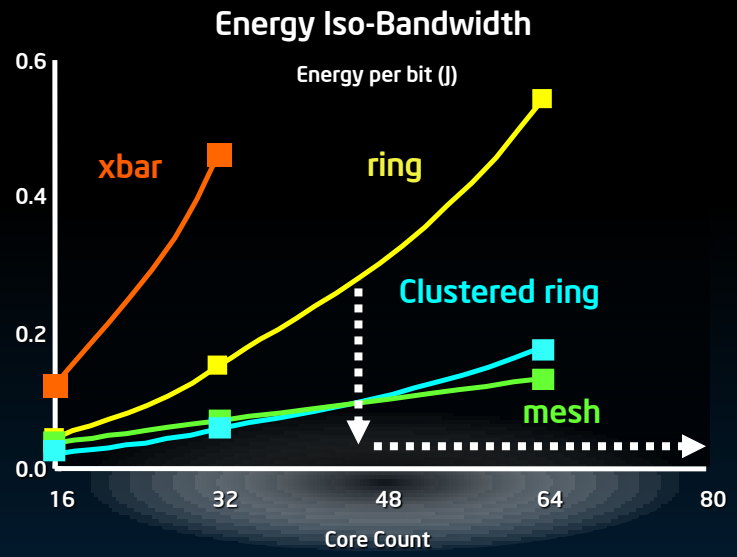

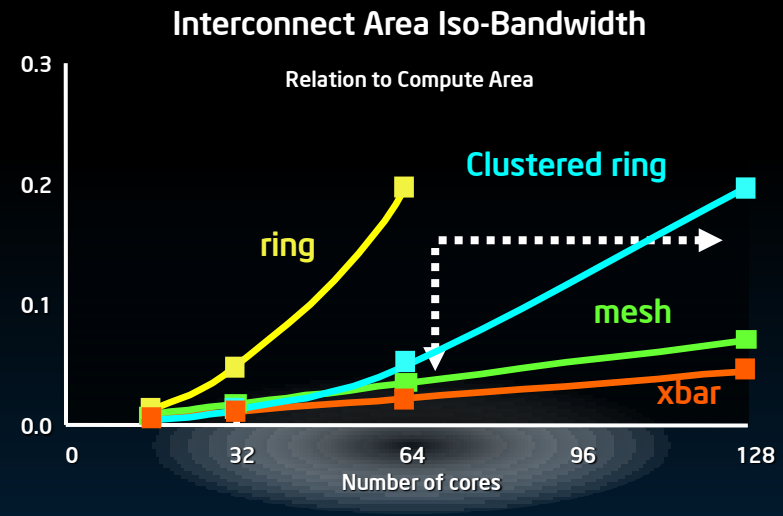

Potential future options, no indication of
actual product or development, subject to change without notice.

Memory and CPU Package Architectures for addressing Bandwidth Challenges - Research

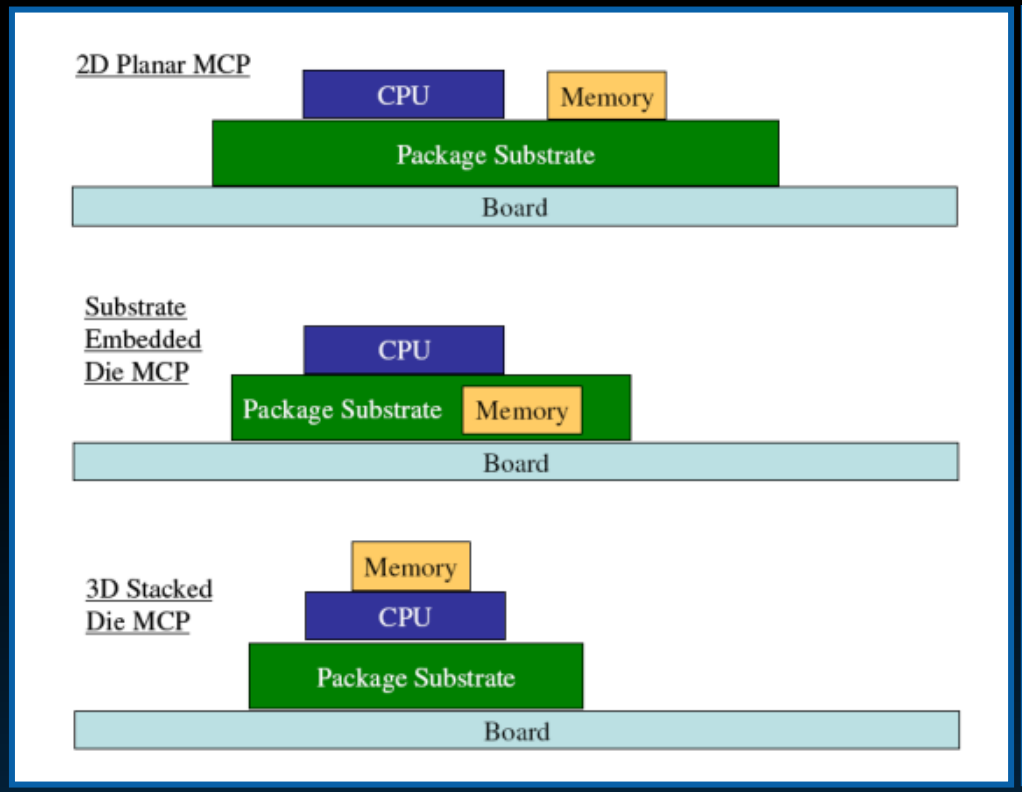

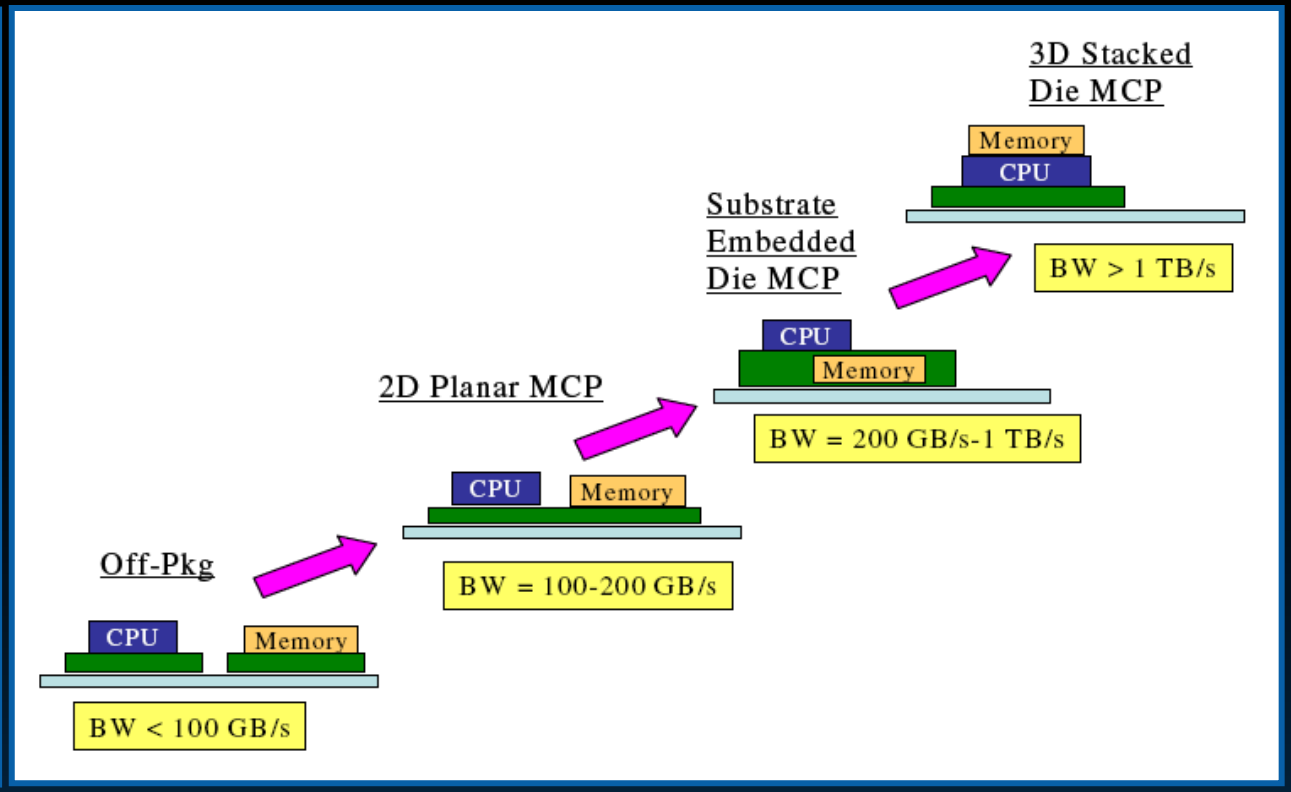

Package Technology to Address the Memory Bandwidth Challenge for Tera-scale Computing, Intel Technology Journal, Volume 11, Issue 3, 2007

Potential future options, no indication of
actual product or development, subject to change without notice.

Hybrid Memory Cube: Experimental DRAM

Highest Performance and most Energy Efficient DRAM in the Industry

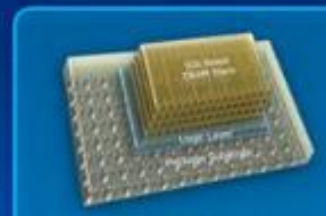

Hybrid DRAM Stack

Micron

Research Collaboration with Micron Technology

Lowest ever energy per bit (~8pJ per bit)

7x better energy-efficiency than today's DDR3

128GBps (>1 Terabit per second) bandwidth

Highest ever bandwidth to a single DRAM device

Technology	VDD	BW GB/s	Power (W)	mW/GB/s	pJ/bit
SDRAM PC133 1GB ECC Module	3.3	1.1	7.7	7226	903.3
DDR3-1333 4GB ECC Module	1.5	10.7	4.6	432	54.0
HMC Gen1 512MB Cube	1.2	128.0	8.0	62	7.78

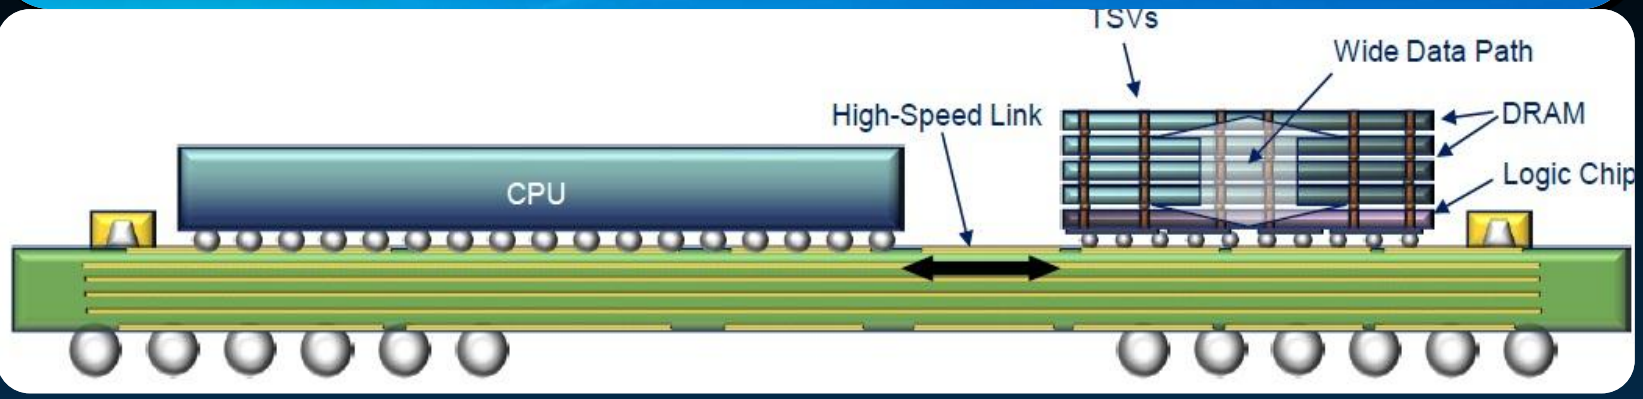

http://www.micron.com/innovations/hmc.html

whatif.intel.com

Access innovations … *in the formative stage*

Designing New Capabilities

- Adobe* Premiere® Pro/Premiere® Elements Encoder plug-in using Intel® Media SDK and Intel® Quick Sync Video Technology **New!**
- Intel® OpenCL SDK **New! Updated 2/1/2011**
- Intel Advisor Lite Now Part of Intel® Parallel Studio
- Intel® Web APIs **New!**
- Intel® Energy Checker SDK **Rev 2.0 Release**
- Intel® SOA Expressway XSLT 2.0 Processor
- Smoke - Game Technology Demo **Rev 1.2 Released**
- Isolated Execution
- Intel® Direct Ethernet Transport
- Intel® Software Development Emulator

Creating Concurrent Code

- **New!** Intel® Cilk Plus Software Development Kit
- Intel® Cilk++ Software Development Kit
- Intel® Concurrent Collections for C++ **Rev 0.6 Released**
- Intel® C/C++ STM Compiler, Prototype Edition **Rev 4.0 Released**

Math Libraries

- Intel® Cluster Poisson Solver Library
- Intel® Adaptive Spike-Based Solver
- Intel® Ordinary Differential Equations Solver Library

Performance Tuning

- Intel® Software Autotuning Tool **New!**
- Intel® Software Tuning Agent
- Intel® Architecture Code Analyzer
- Intel® Performance Tuning Utility **4.0 Update 3 Released**
- Intel® Platform Modeling with Machine Learning

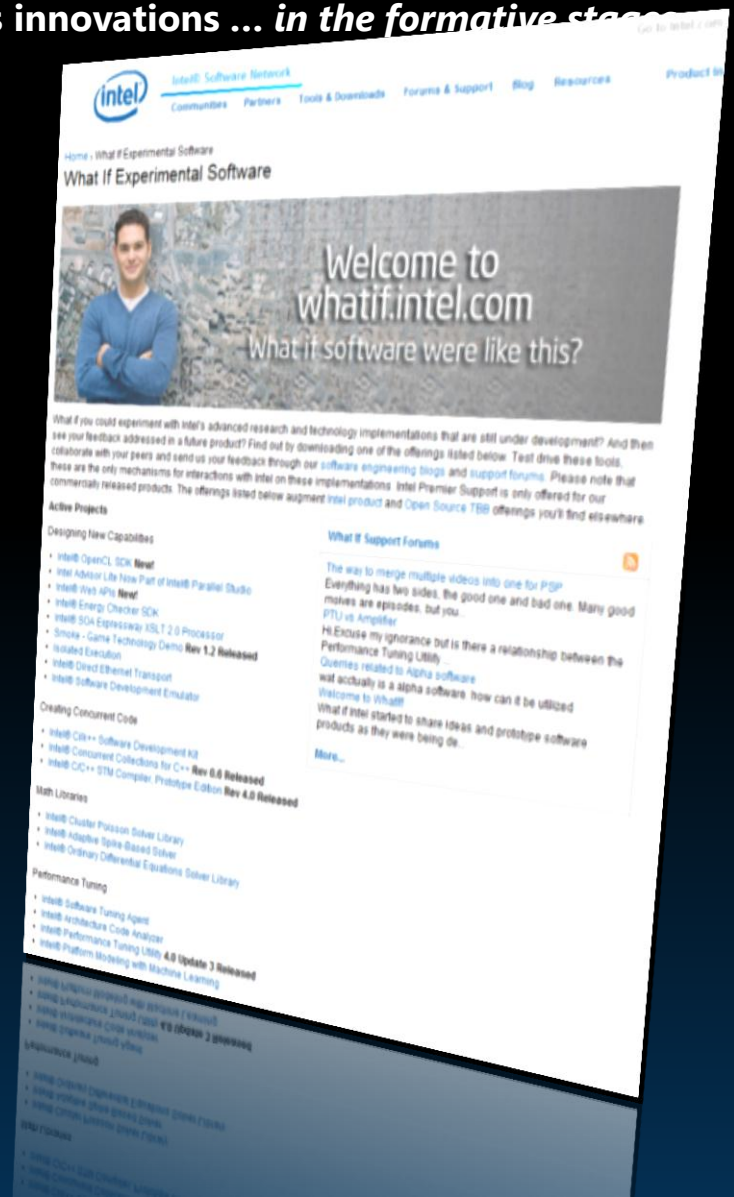

Summary

- Moore's law moving forward
- Xeon delivers the best architecture for diverse workload
- Addition of Intel® Xeon® Phi™ offers new opportunities for highly parallel applications
- Investment into software is the most precious asset


Intel 22nm 3-D Tri-Gate Transistor Technology
Unprecedented combination of performance and energy efficiency

Thank You.

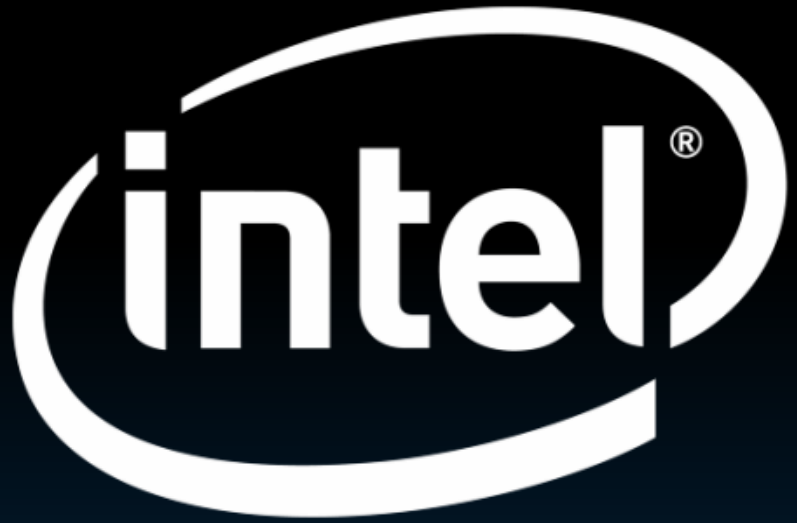
intel