Accelerating HPC

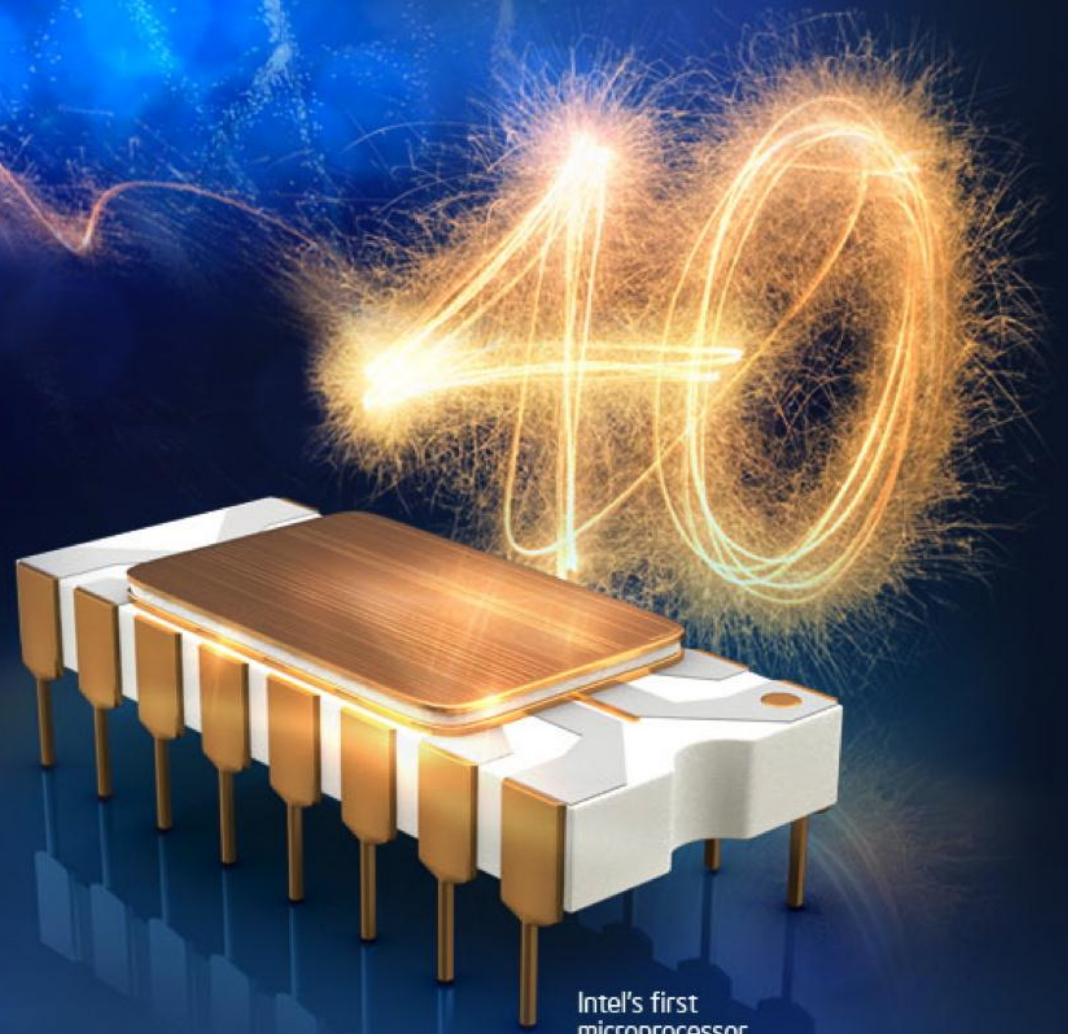

Intel's first microprocessor, circa 1971.

CERN OpenLab, Geneva
February 6th, 2012

Herbert Cornelius
Intel

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.*

Copyright © 2010 Intel Corporation.

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.*
- *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

COMPUTING 2020

(Source: IDC Directions 2010, March 2010, ICT Outlook forecast)

40th Anniversary of the Microprocessors

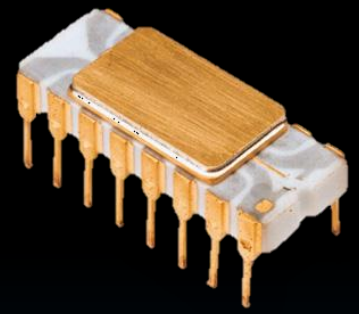

Intel® 4004 (1971)
10000nm, 2300 Transistors
740KHz

Intel® Core™ i7 (2011)
32nm, 2.27B Transistors
3.6GHz

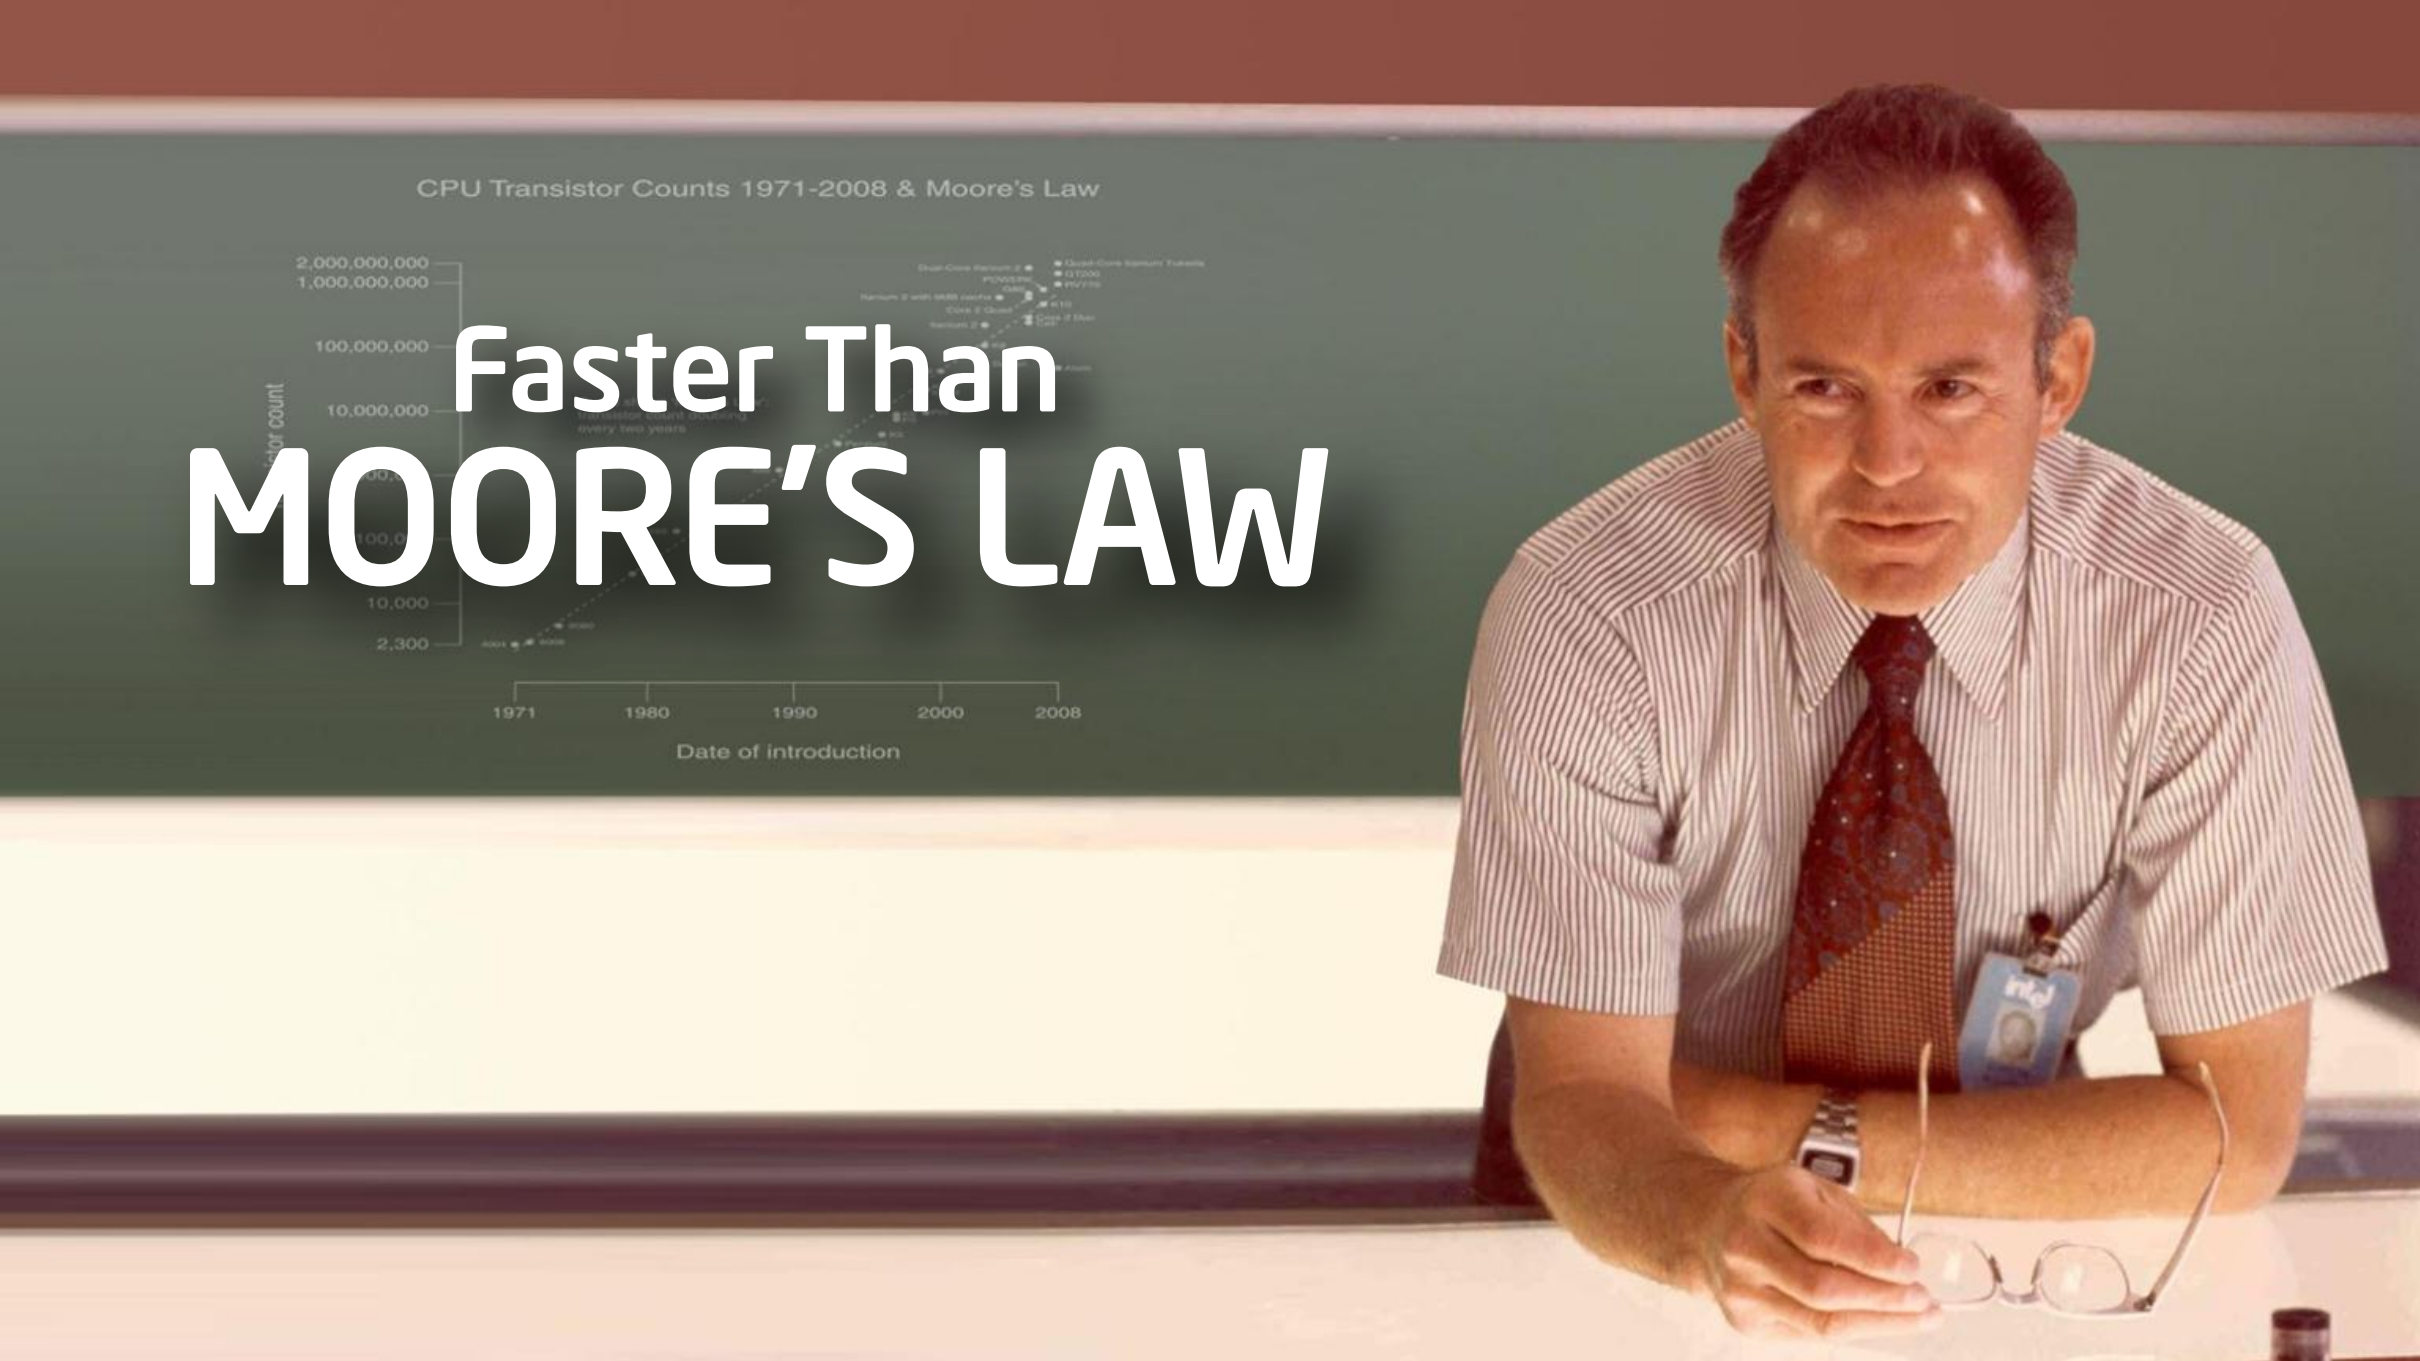

Faster Than MOORE'S LAW

Transistor Innovations Enable Technology Cadence

2003	2005	2007	2009	2011
90 nm	65 nm	45 nm	32 nm	22 nm
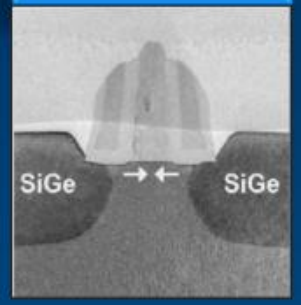	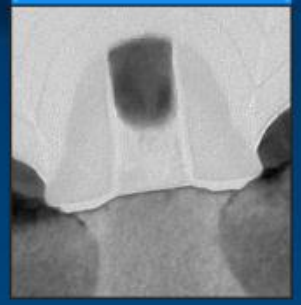	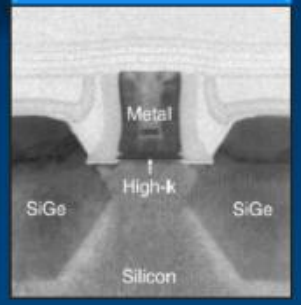	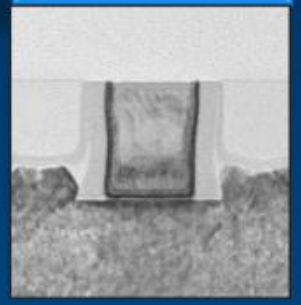	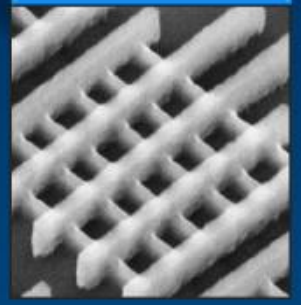
Invented SiGe Strained Silicon	2nd Gen. SiGe Strained Silicon	Invented Gate-Last High-k Metal Gate	2nd Gen. Gate-Last High-k Metal Gate	First to Implement Tri-Gate
Strained Silicon				
		High k Metal gate		
				Tri-Gate

22nm

37%
Performance Gain at Low Voltage[+]

>50%
Active Power Reduction at Constant Performance[+]

Source: Intel
[+]Compared to Intel 32nm Technology

intel

14nm

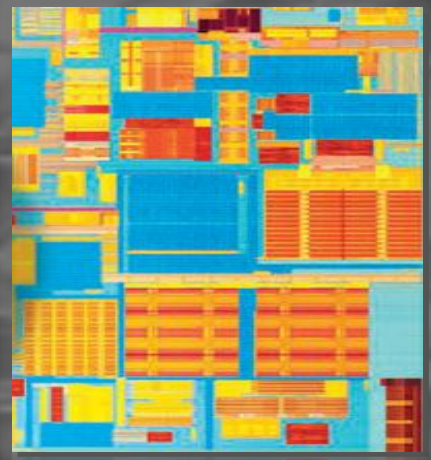

Well on track

Intel in High-Performance Computing

Dedicated, Renowned Applications Expertise

Large Scale Clusters for Test & Optimization

Tera-Scale Research

Exa-Scale Labs

intel Cluster Ready

Defined HPC Application Platform

intel Software

Broad Software Tools Portfolio

Industry Standards

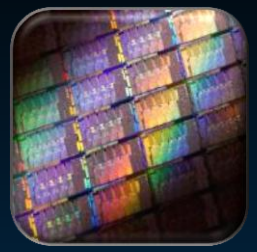

Manufacturing Process Technologies

Leading Processor Performance, Energy Efficiency

Many Integrated Core (MIC) Architecture

Platform Building Blocks

A long term commitment to the HPC market segment

Increasing Processor Performance

Through Many-Core Technologies for Highly Parallel Workloads

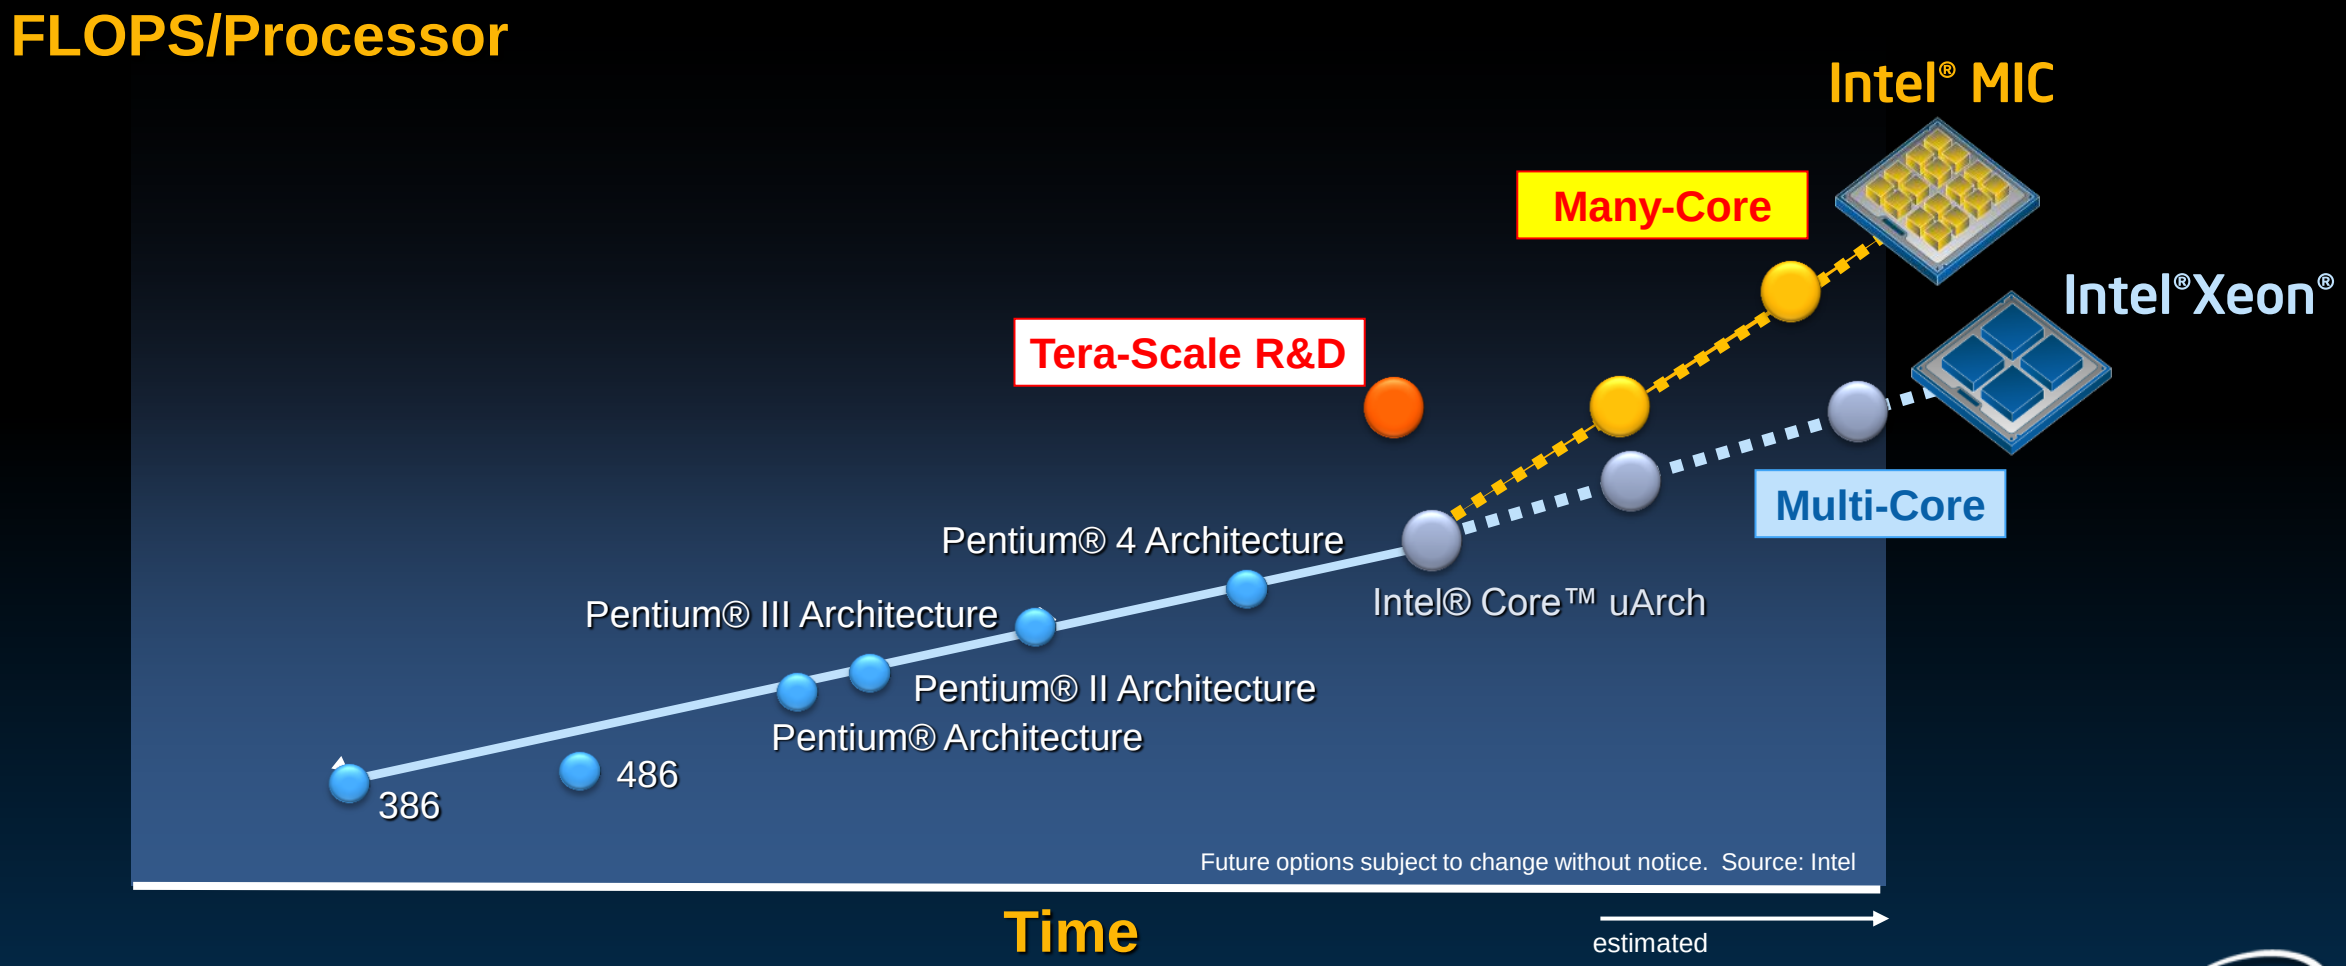

All dates, product descriptions, features, availability, and plans are forecasts and subject to change without notice.

Tick-Tock Development Cycles

Integrate. Innovate.

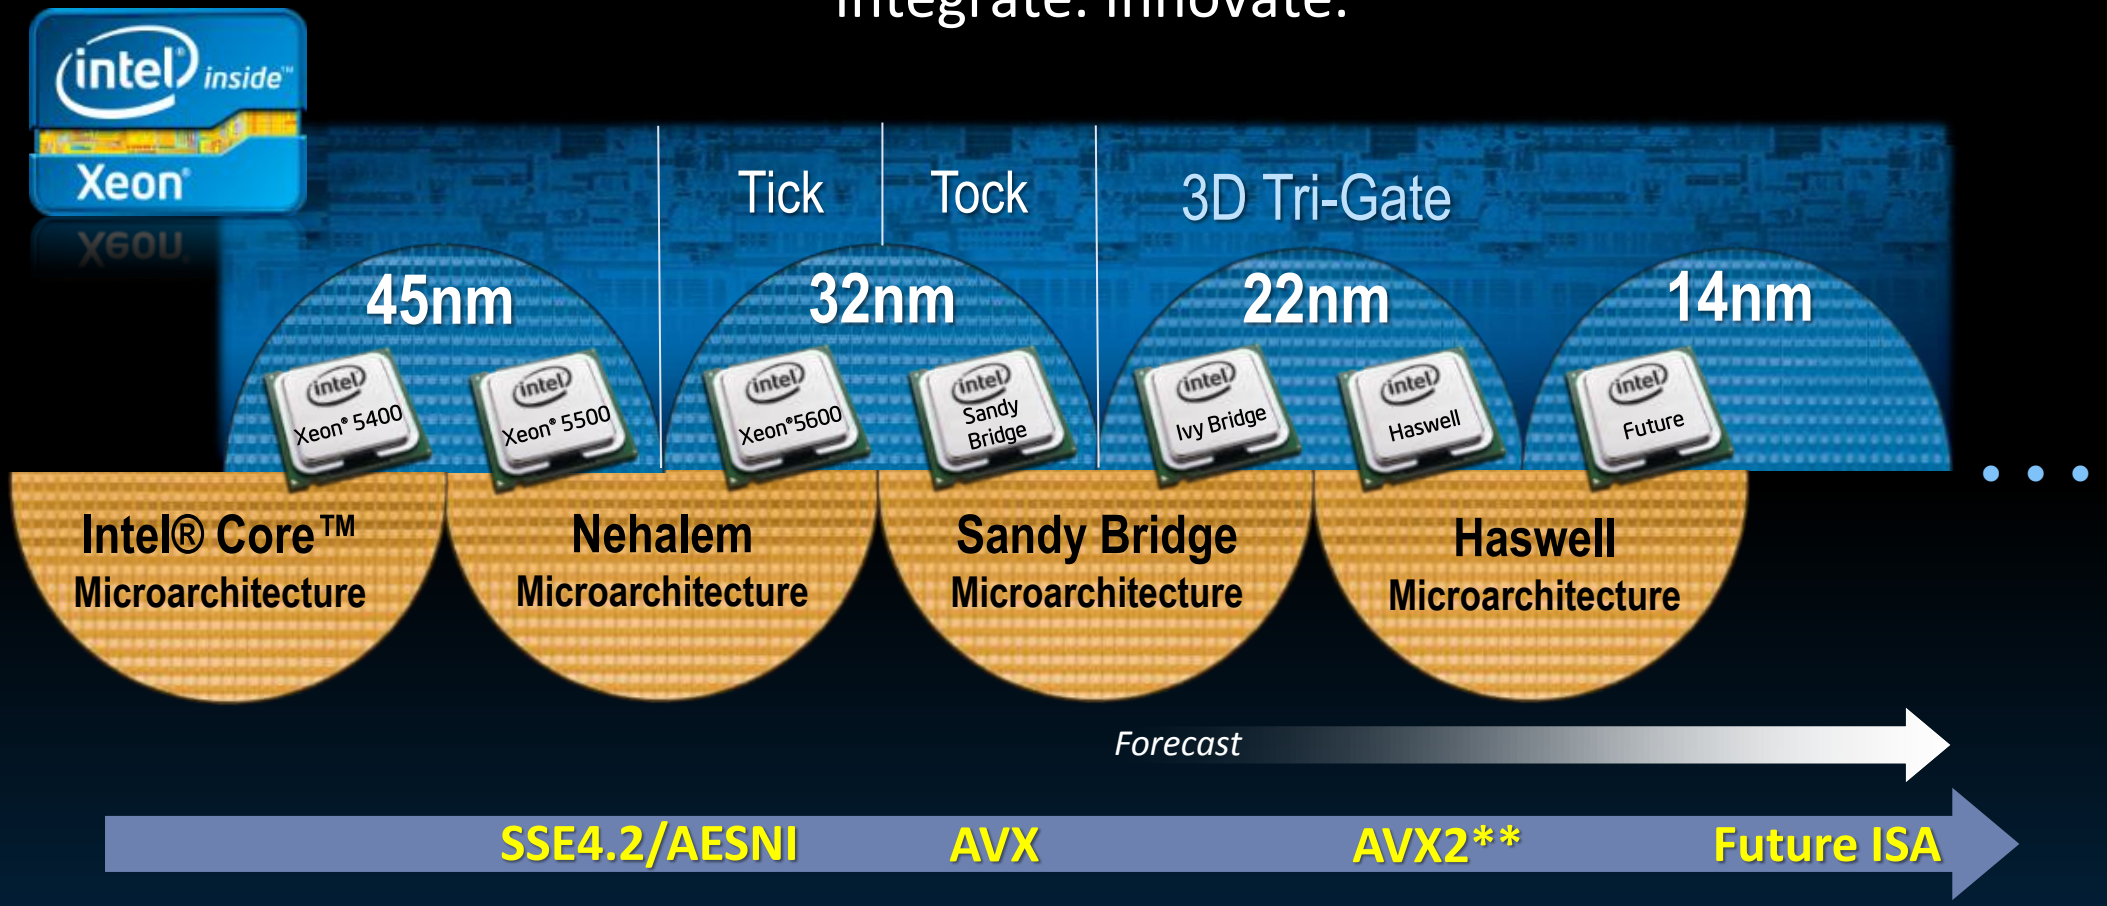

***Intel® Advanced Vector Extensions Programming Reference, Ref. #319433-011, JUNE 2011*

Potential future options, subject to change without notice.

Intel® Xeon® E5 Series Processor

Codename Sandy Bridge-EP

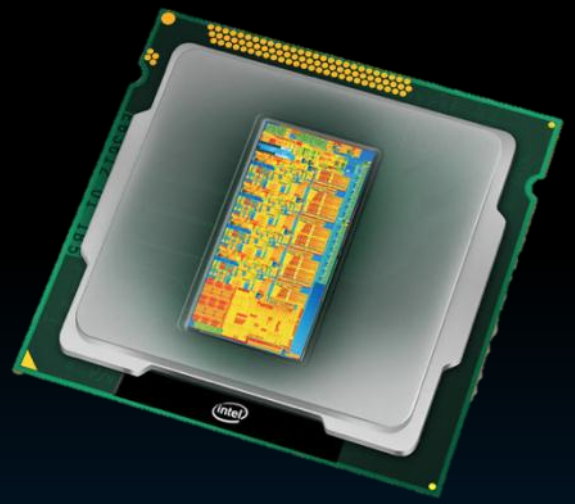

152 GFLOPS DP-F.P. per processor on HPL with 91% efficiency (8C, 2.6GHz)

AVX Instructionset (256-bit SIMD F.P.)
Enhanced Hyper-Threading and Turbo-Technology
Large on-die L3-Cache
First server processor chip with integrated PCI Express* 3.0 I/O

Early-performance benchmarks**:

- up to 2.1x raw FLOPS (Linpack)
- up to 70% percent more performance using real-HPC workloads
- compared to the previous generation of Intel Xeon 5600 series processors (Westmere)

Balanced compute/memory, fast single core/thread multi-core architecture

** Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. 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.

Increased Application Performance

Intel® Xeon® Processor E5 Family

Relative Geometric Mean Scores by segment
Actual performance will vary by workload.
Early Performance Benchmarks.
Higher is better.

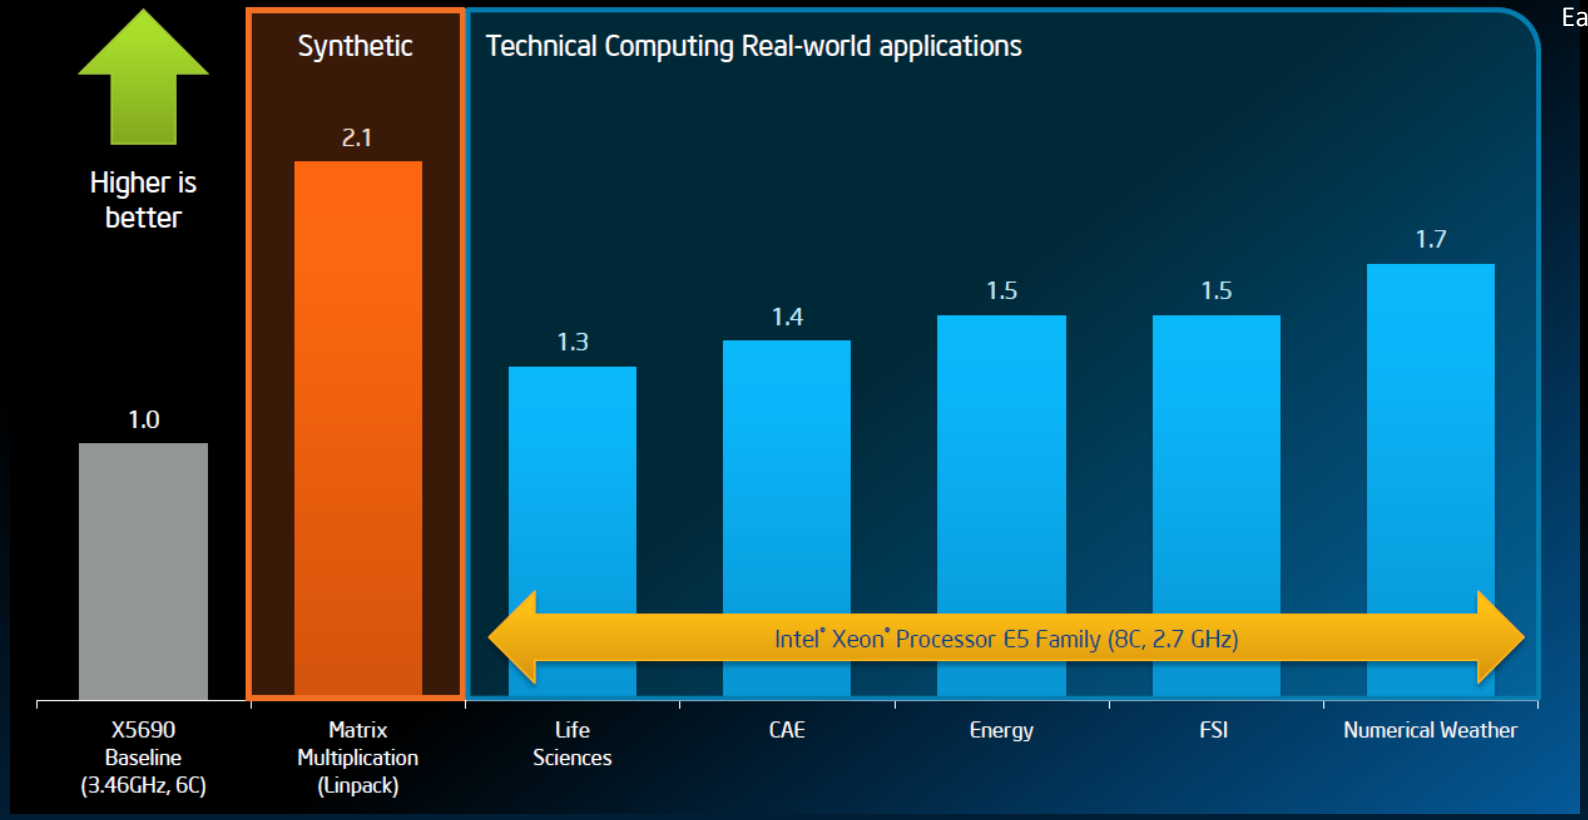

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may 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 October 2011, See backup for configuration details. For more information go to http://www.intel.com/performance . Any difference in system hardware or software design or configuration may affect actual performance.

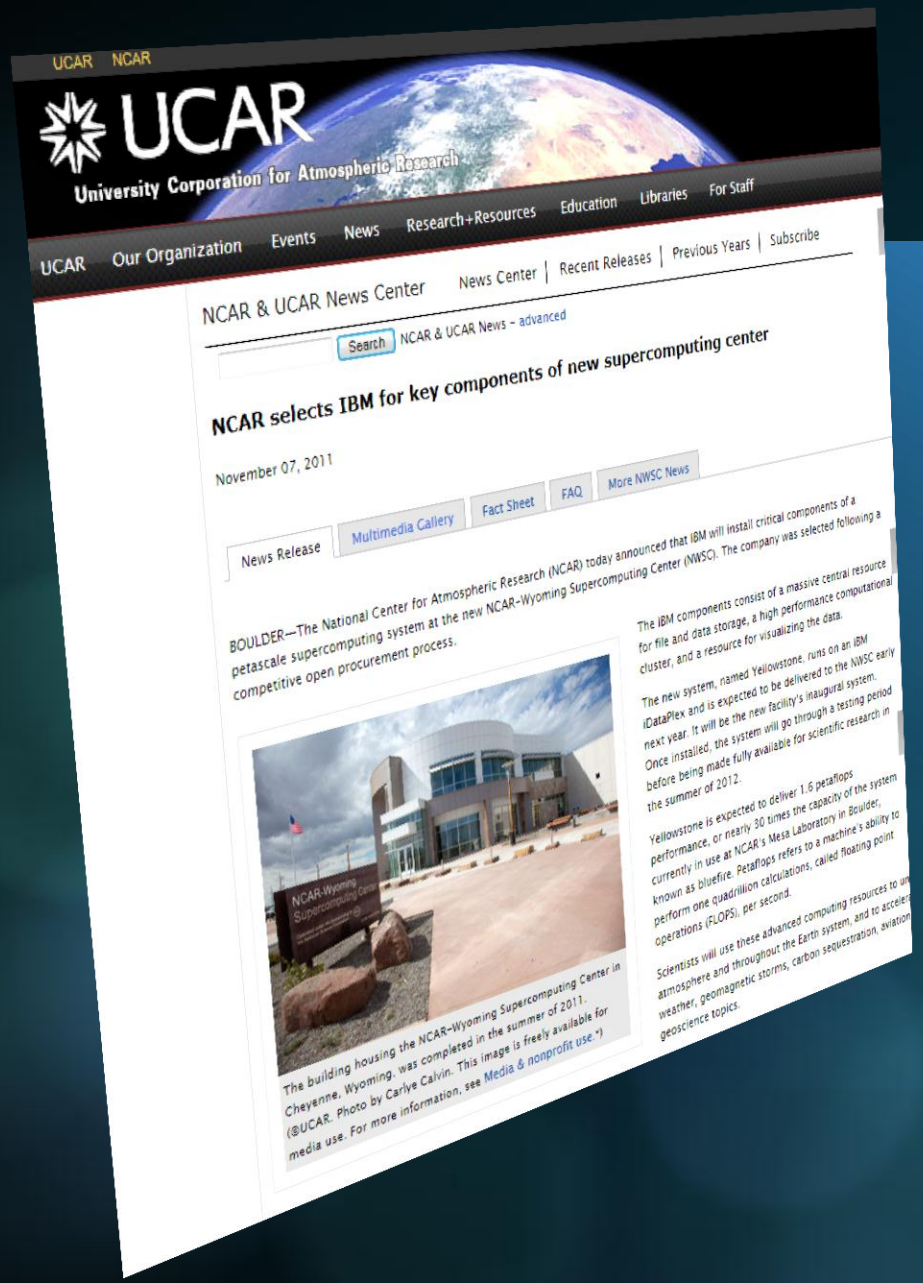

"Yellowstone"

- 1.6 PFLOPS peak
- 74592 processor cores
- 149 TB memory
- 17 PB storage
- Deployment scheduled in H1'2012
- IBM* iDataPlex* System
- Intel® Xeon® E5 (Sandy Bridge-EP) processors
- Mellanox* FDR InfiniBand (56Gb/s) cluster fabric

INTEL® MANY INTEGRATED CORE (MIC) ARCHITECTURE

GENERAL PURPOSE ENERGY EFFICIENT TFLOPS PERFORMANCE FOR HIGHLY PARALLEL WORKLOAD

USING COMMON X86 STANDARD PROGRAMMING MODELS AND SW-TOOLS

Intel's Multi-Core and Many-Core Engines

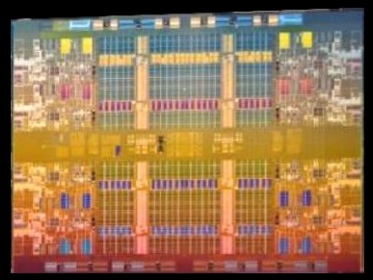

Multi-core Intel® Xeon® processor up to 3.6 GHz

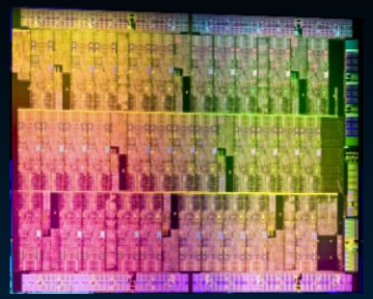

Many Integrated Cores at 1+ GHz

Intel® Xeon® Processor

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

Intel® MIC Architecture

- 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-sizes not to scale)

High Performance and Energy Efficiency

for highly Parallel Workloads

high F.P. performance (VPU/SIMD)

$\Sigma \{$ 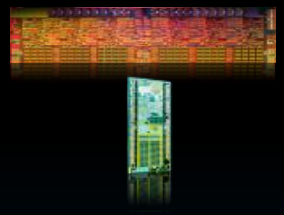 $\}$ =

small extreme energy efficient core

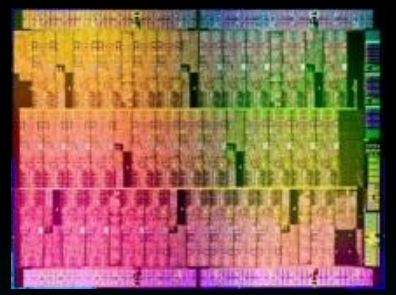

Many-core cache coherent shared memory SMP

Intel® Many Integrated Core (MIC) Architecture

The Newest Addition to the Intel Server Family.
Industry's First General Purpose Many Core Co-Processor Architecture

The “Knights” Family

Future Knights Products

Knights Corner

1st Intel® MIC product

22nm process

>50 Intel Architecture cores

TFLOPS of Performance

Energy Efficient

Offload Co-Processor and

Native Linux* Node Programming

Knights Ferry

Development Platform

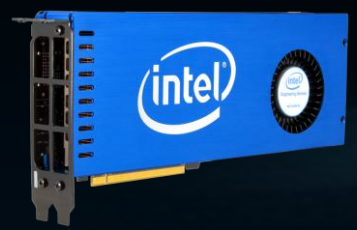

“Programmed like a computer”

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

Software Development and System Environment

HPCwire Editors' Choice Awards 2011

Same Comprehensive Set of SW Tools: Application Source Code Builds with a Compiler Switch

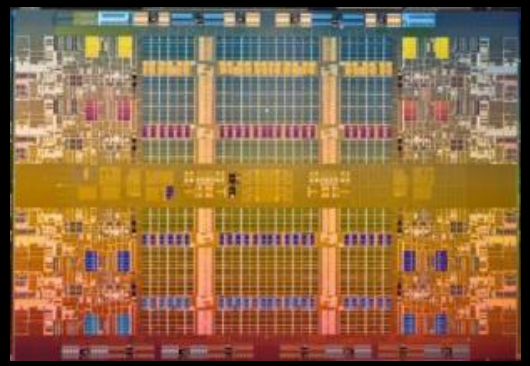

Intel® Xeon® Processor

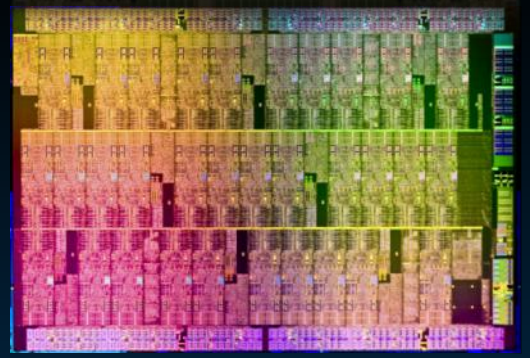

Intel® Many Integrated Core Architecture

(die-sizes not to scale)

Linux*

Established HPC Operating System

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 Message Passing

Co-Array Fortran

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

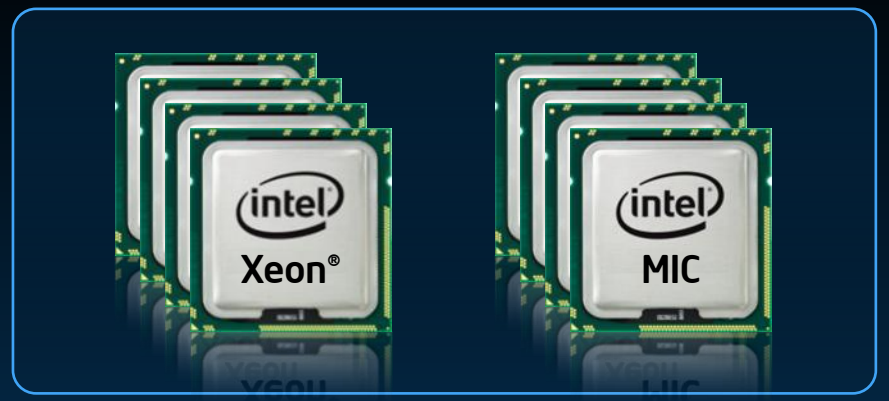

Potential future options subject to change without notice.

Single Source Code

Compiler
Libraries
Parallel Models

intel Software

Intel Parallel Studio XE

Intel Cluster Studio XE

Multicore

Many-core

Cluster

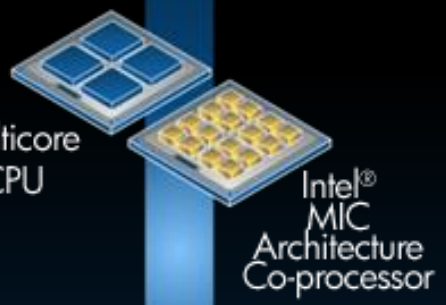

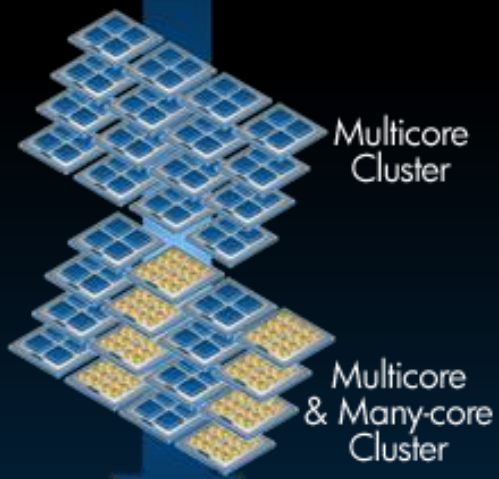

Eliminate Need for Dual Programming Software Architecture

For illustration only, potential future options subject to change without notice.

intel

Moore's Law at Work

1997

1 TFLOPS DP-F.P.

9298 Chips

"ASCI RED"

~2500 Square Feet

850KW Supercomputer

Source: Sandia

2011

1 TFLOPS DP-F.P.

Single Chip (MIC)

SC'11 November 2011

1 TFLOPS DP-F.P.

(Early Silicon Demonstration)

MIC: Knights Corner

- In 22nm process technology
- >50 cores/die energy efficient
- 512 bit SIMD instructions
- Early Si delivers 1TFLOPS sustained on DGEMM
- Runs Linux
- Can be
  - a native network node (ssh in …)
  - used as an offload co-processor
- Common x86 programming models, techniques and tools
- Targeted by Intel compilers and SW-tools
- 3rd party software being enabled

CERN OpenLab is working closely with Intel on the evaluation of Intel® MIC Architecture

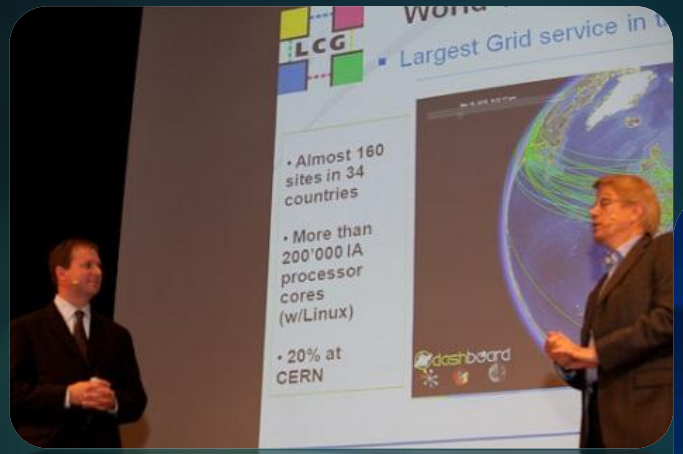

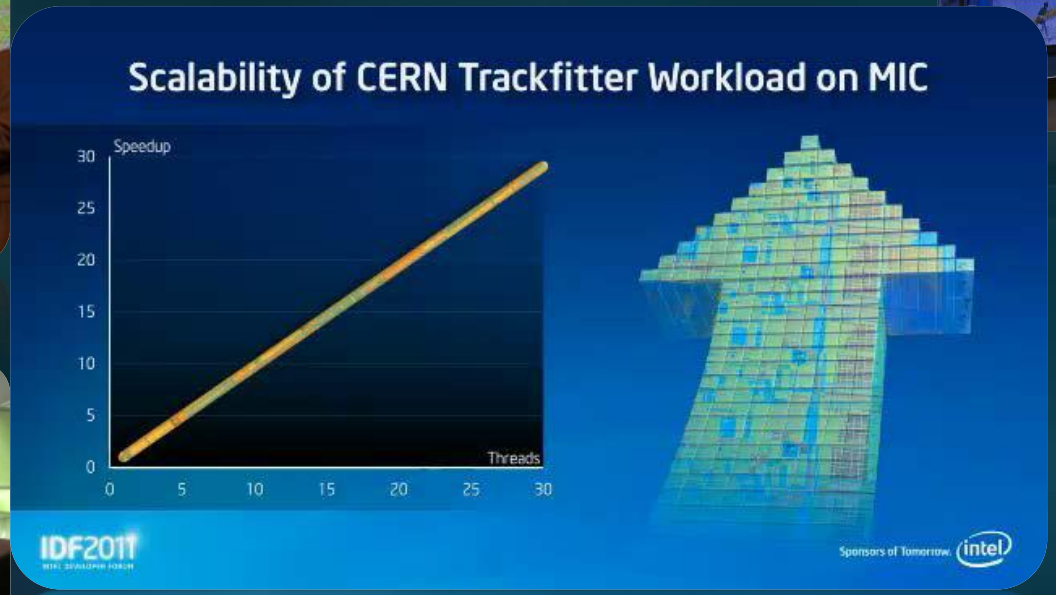

Courtesy CERN OpenLab: Sverre Jarp, Andrzej Nowak, Julien Leduc

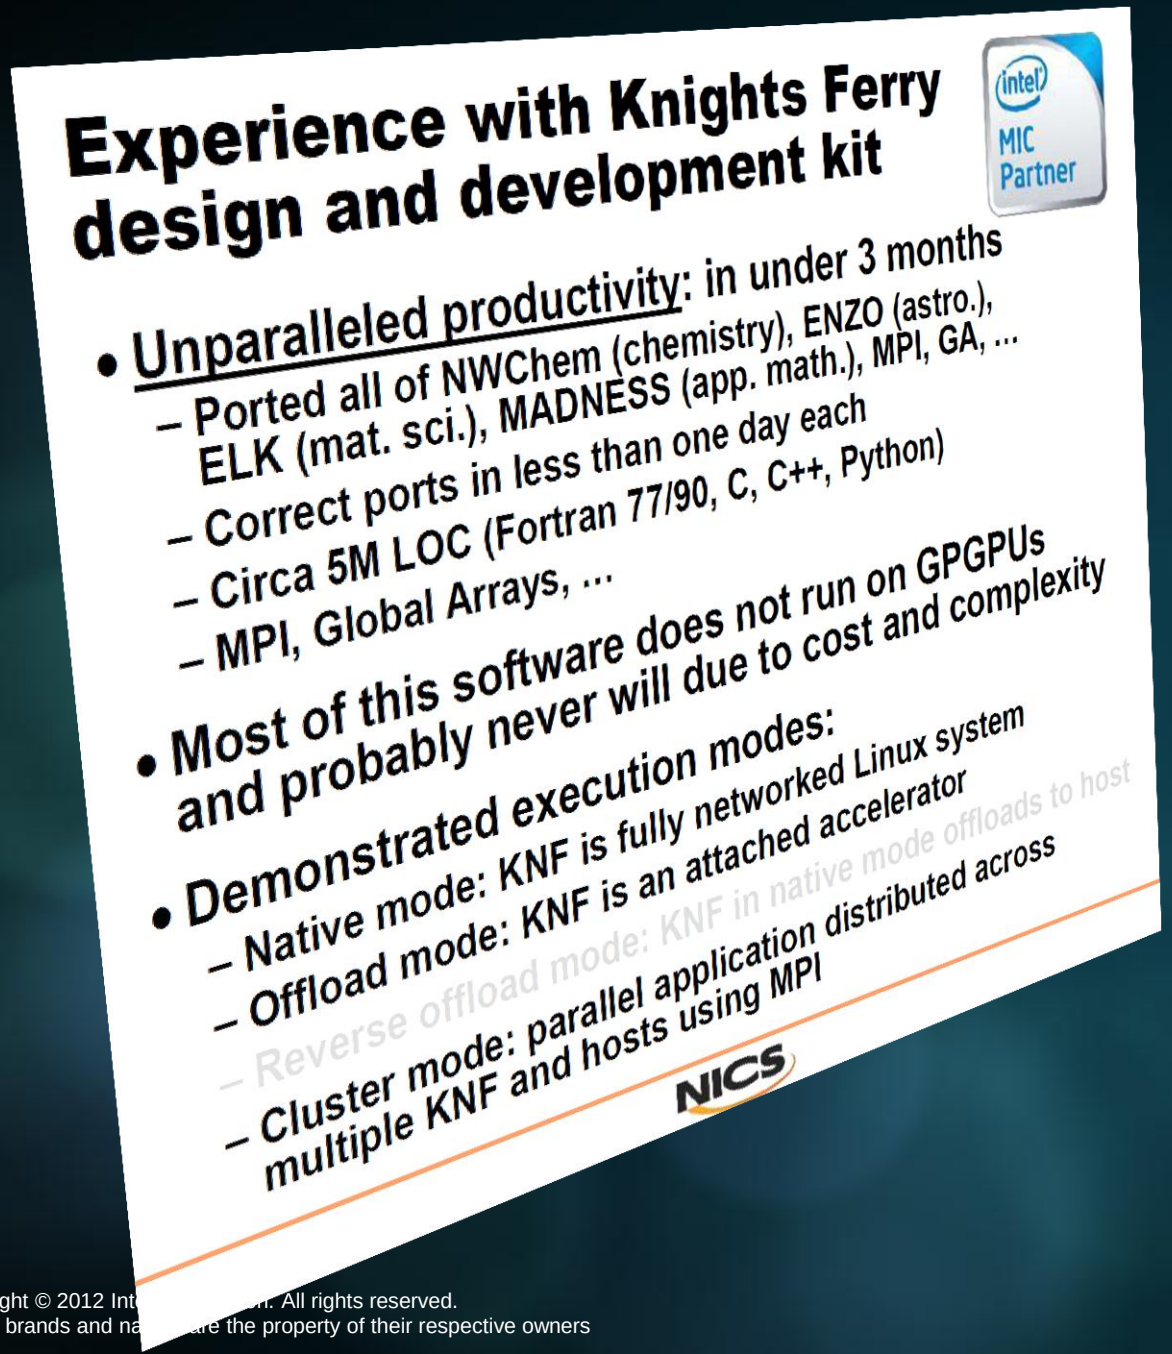

National Institute for Computational Sciences

Joint Institute for Computational Sciences
University of Tennessee & ORNL

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" (22.09.2011)

- 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

Example: Theoretical Speed-Up with 8 Cores and 8-way SIMD

Applying Amdahl's Law Twice

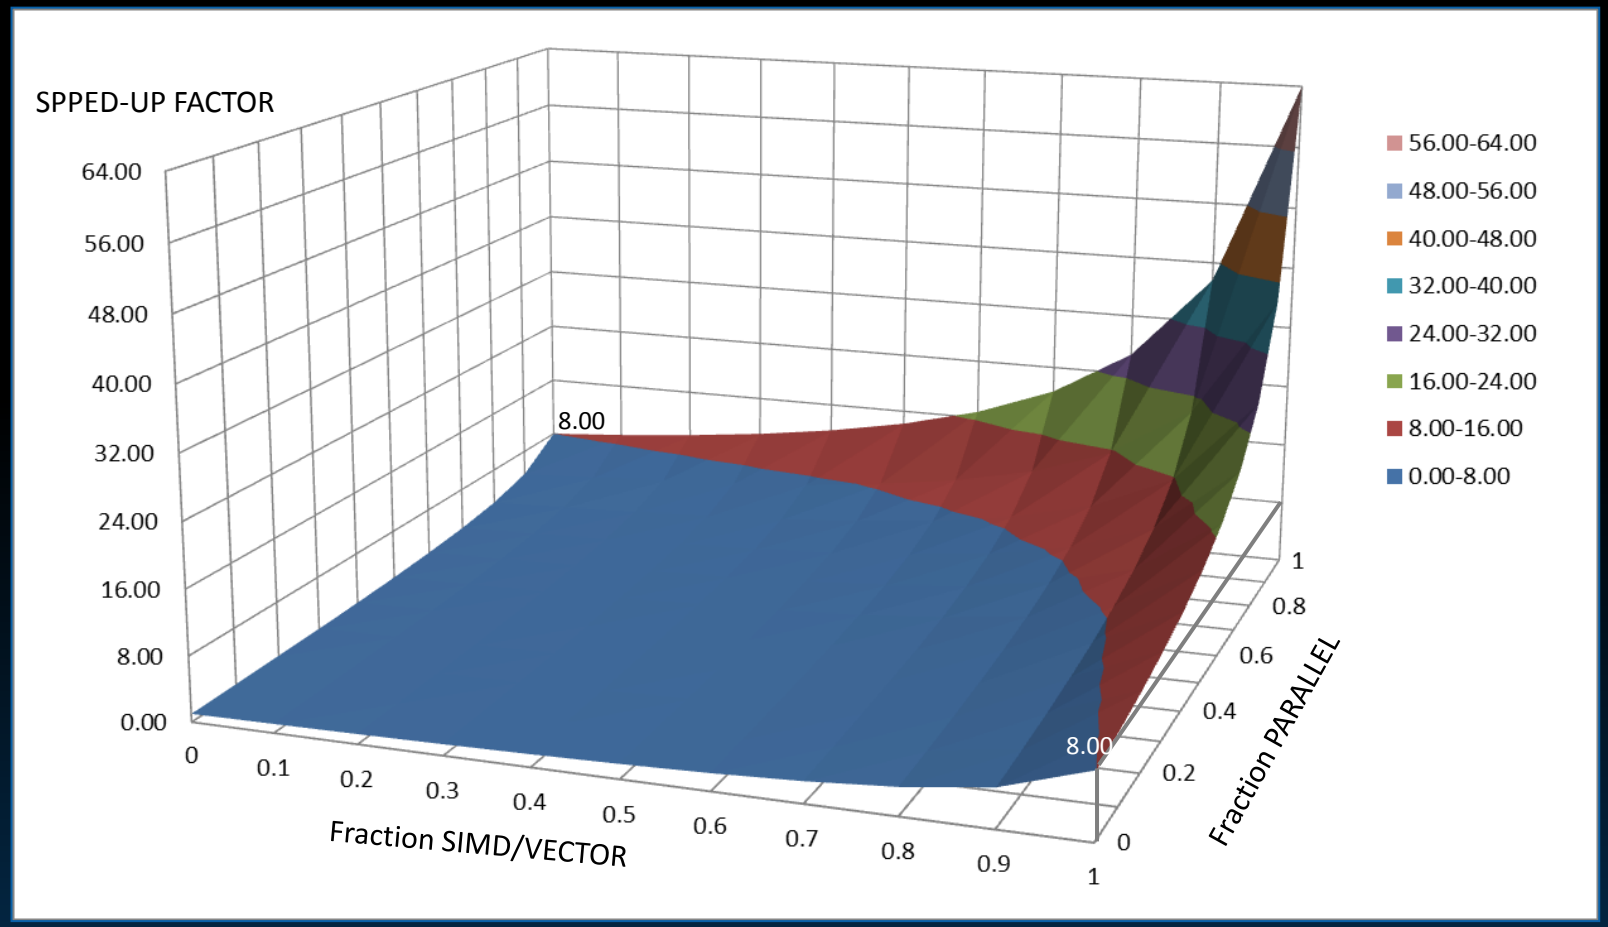

“If you are not scared …

your dreams are not big enough”

Assume Exascale Computing at 20MW …

New Forms of Energy

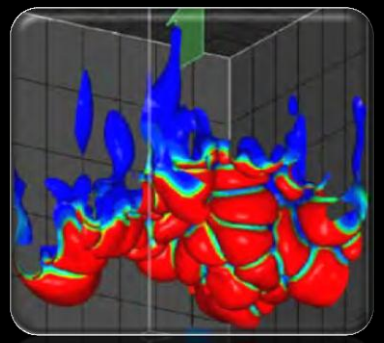

Ecological Sustainability

Space Exploration

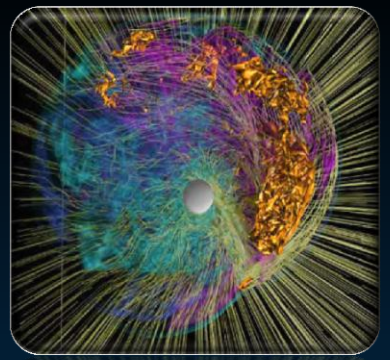

Medical Innovation

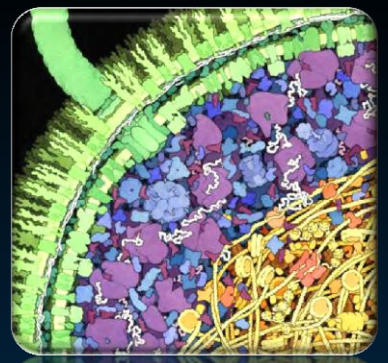

And many others….

Data Center Sized Exascale System

Lower Volume
Higher Cost

20MW

Rack Sized Petascale System

"Mainstream"

20KW

Embedded Terascale System

Higher Volume
Lower Cost

20W

For illustration and concept only.

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 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

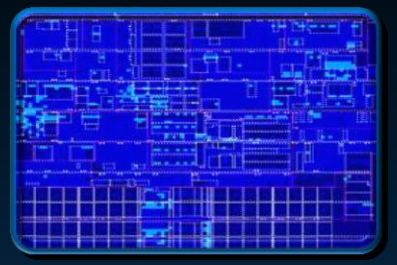

Memory & Storage

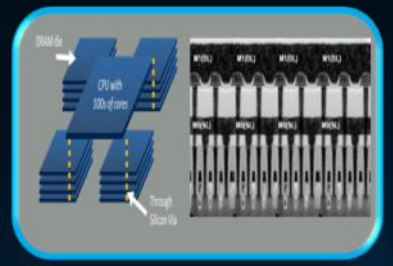

Interconnect

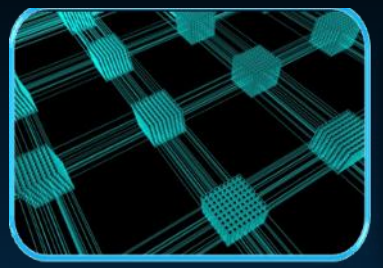

For illustration and concept only.

Potential Path to Exascale

Intel® MIC Architecture

- Aggressively drive FLOPS/Watt to stay within 20MW Exascale target
- Scale memory & interconnect bandwidth through innovation
- Highly parallel, programmable IA
- Large local storage solutions

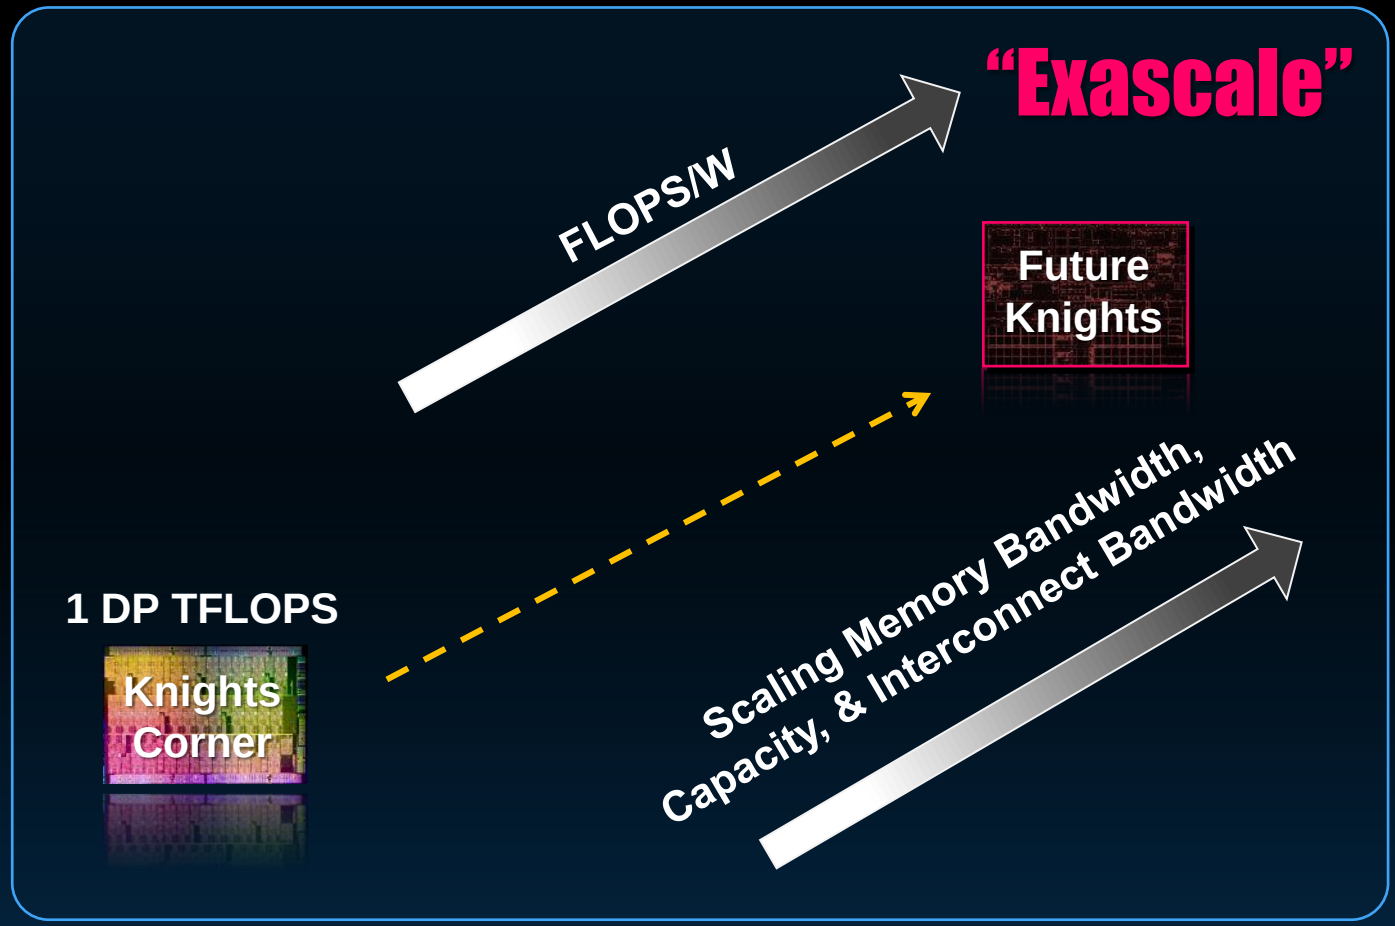

Potential future options; plans, dates and features are subject to change without notice.

Intel TeraScale Research Areas

MANY-CORE COMPUTING

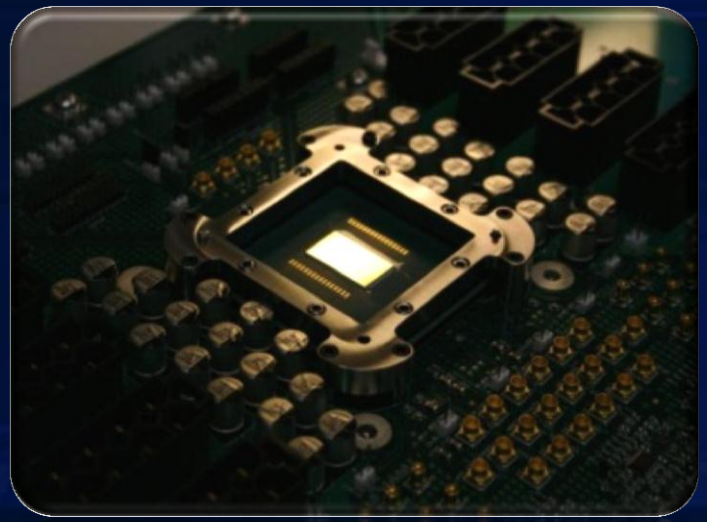

Teraflops
of computing power

3D STACKED MEMORY

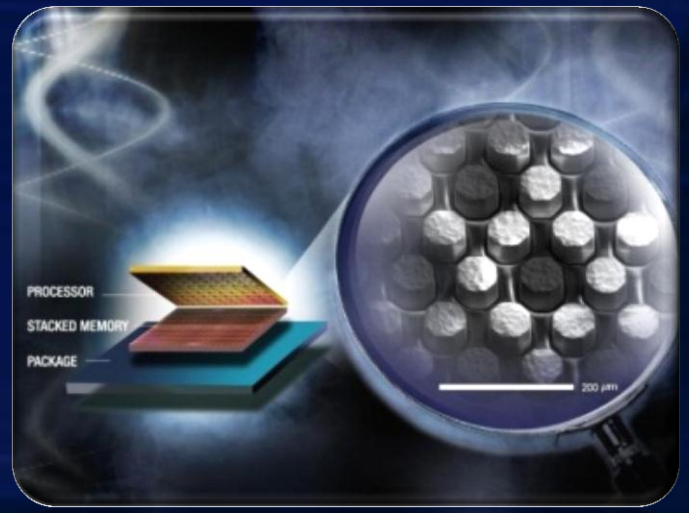

Terabytes
of memory bandwidth

SILICON PHOTONICS

Terabits
of I/O throughput

For illustration only. Future vision, does not represent real products.

IDF 2011

Near Threshold Voltage (NTV) Operation

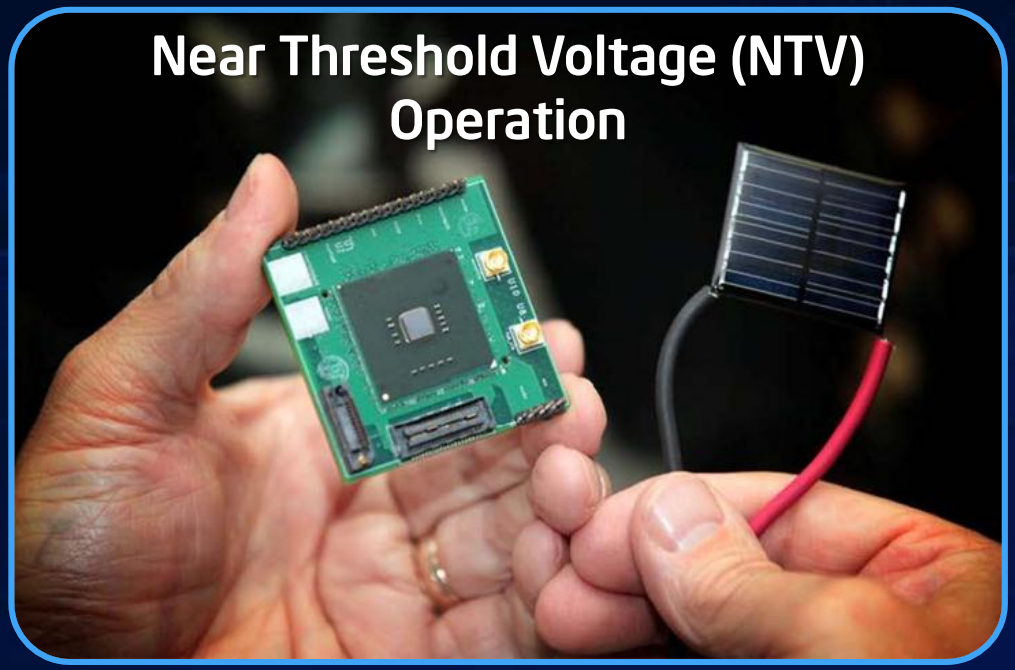

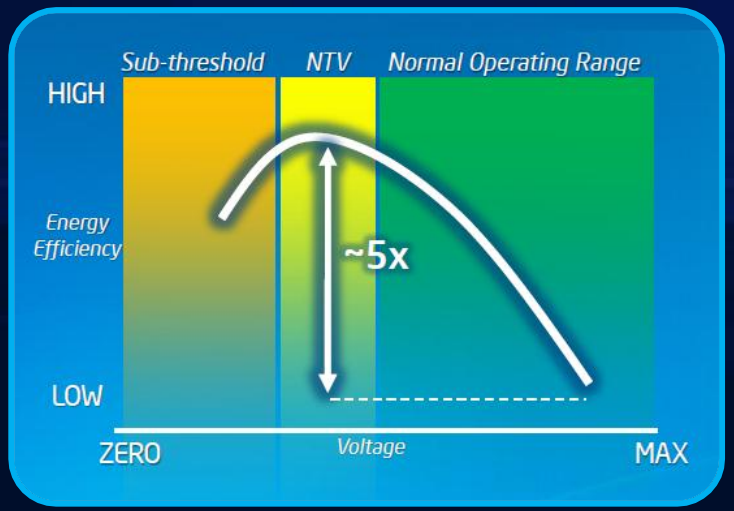

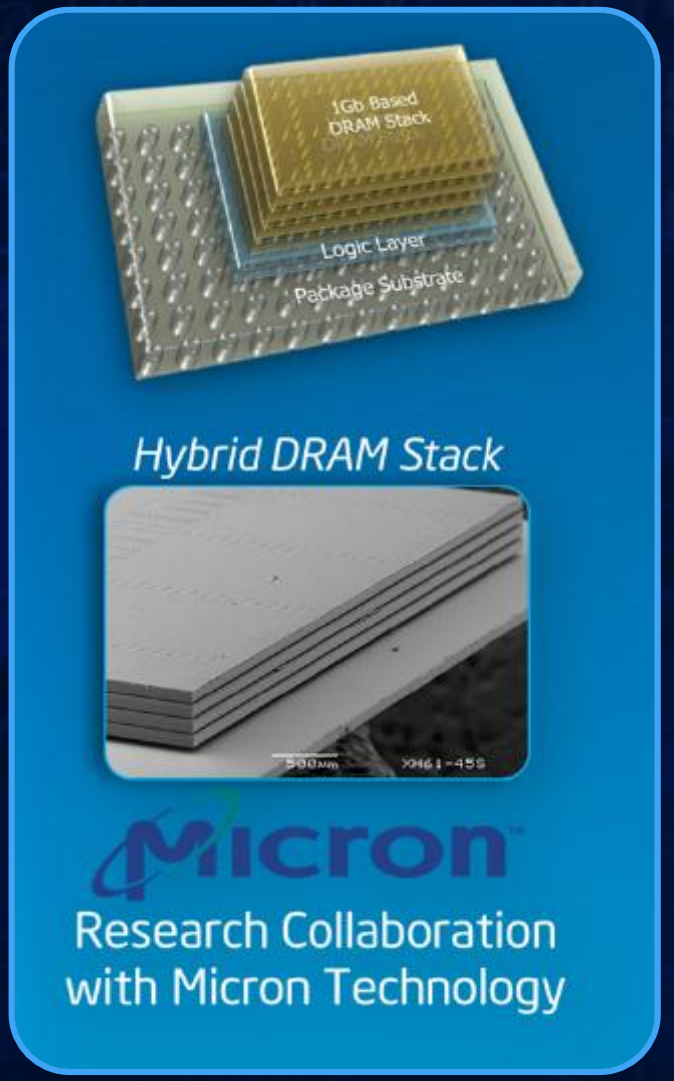

128GB/s (1Tb/s) Bandwidth
7x better energy efficiency than DDR3

Intel and Micron's Joint-Development Venture

IM Flash Technologies (IMFT)

128 Gbit (16GB) NAND in 20nm

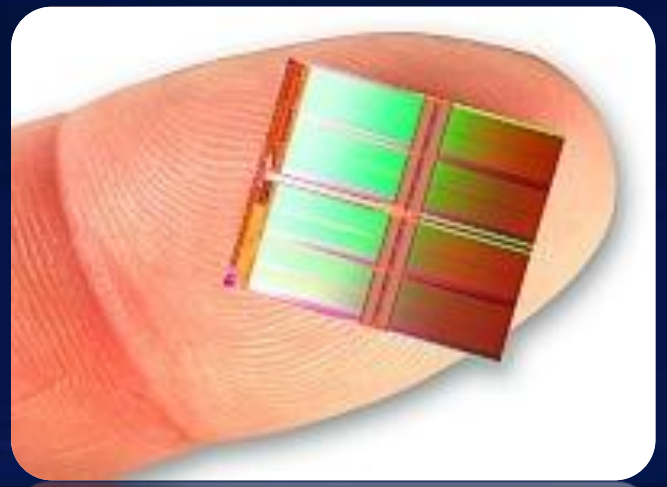

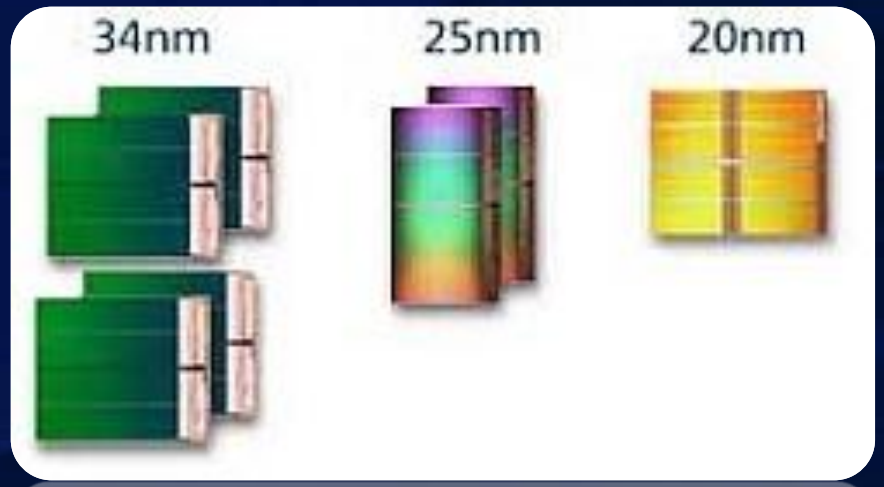

Pictures: Micron

Silicon Photonics

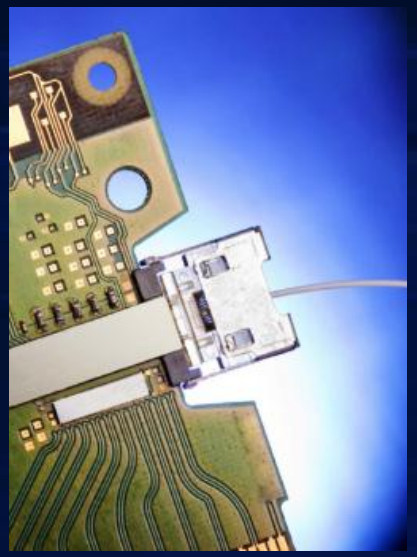

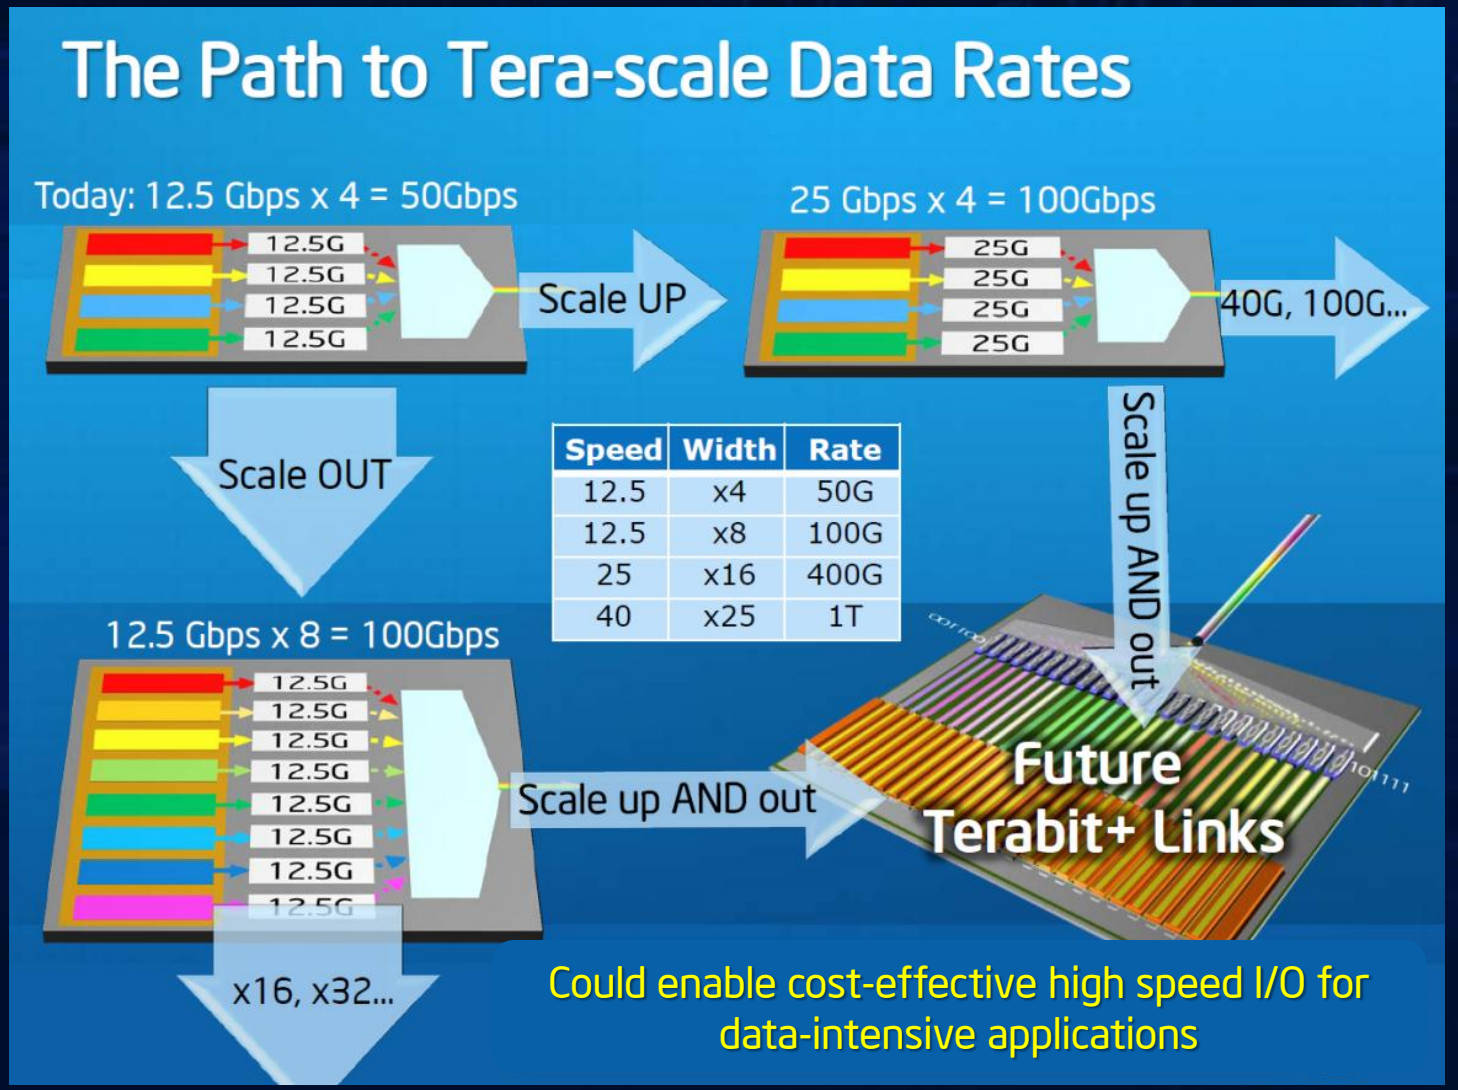

“One last thing …”

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."

Thank You.