Intel®: Accelerating the Path to Exascale

Kirk Skaugen
Vice President Intel Architecture Group
General Manager Data Center Group

An Insatiable Need For Computing

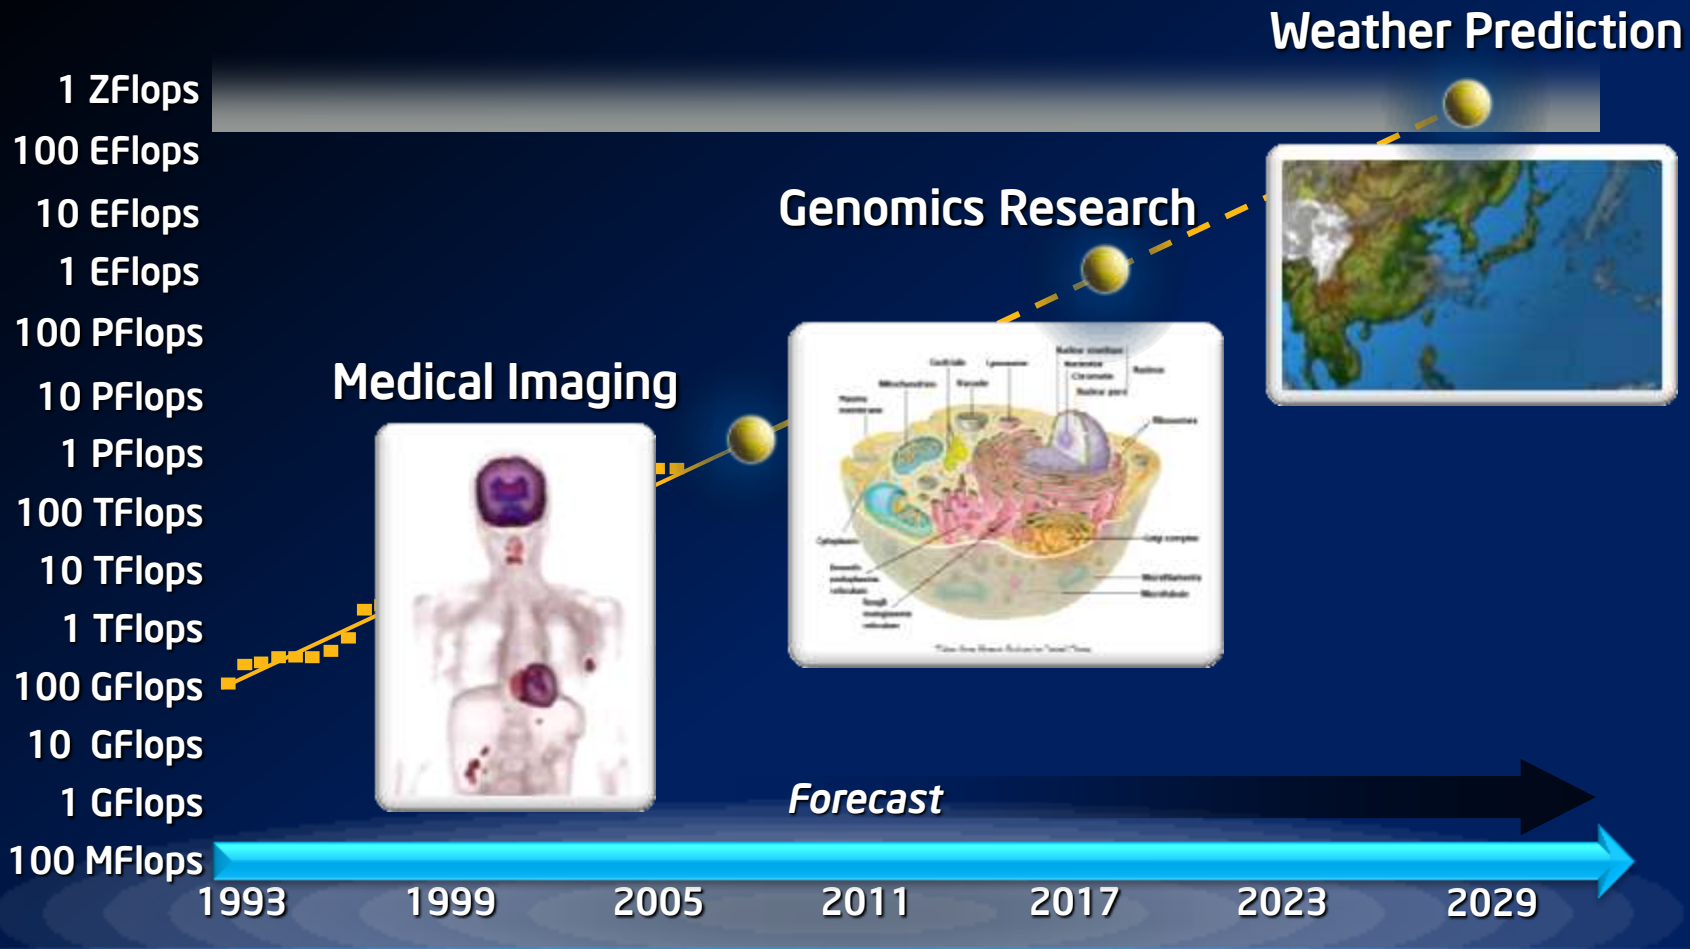

Exascale Problems Cannot Be Solved Using the Computing Power Available Today

Source: www.top500.org

intel

Exascale Answers Mankind's Challenges In…

Weather / Climate

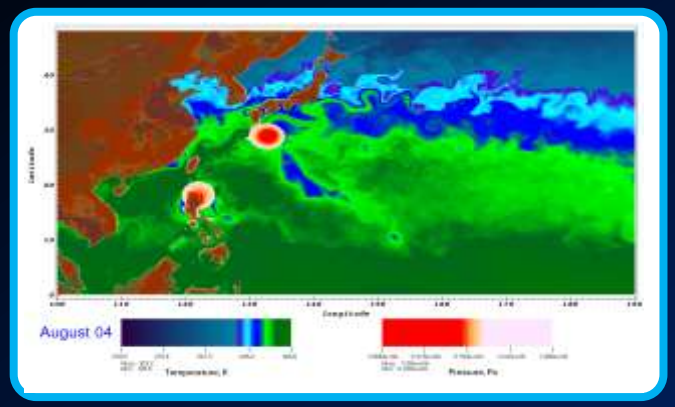

Healthcare

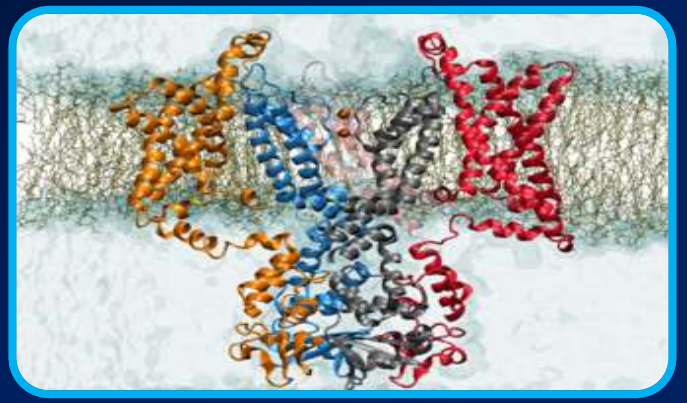

New Forms of Energy

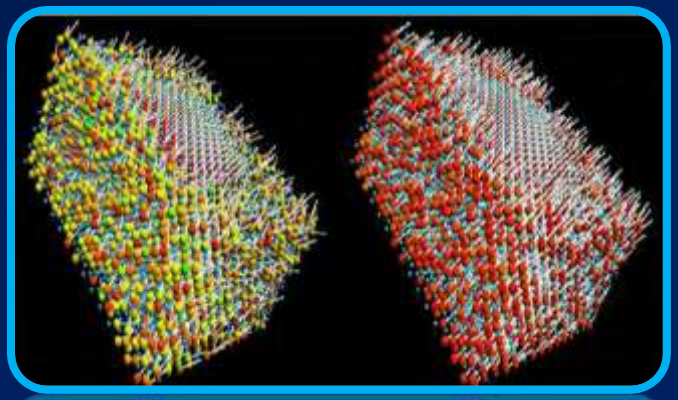

intel

We've Helped Transform Industries

Annual Server Processor Shipments

Supercomputing in 2000

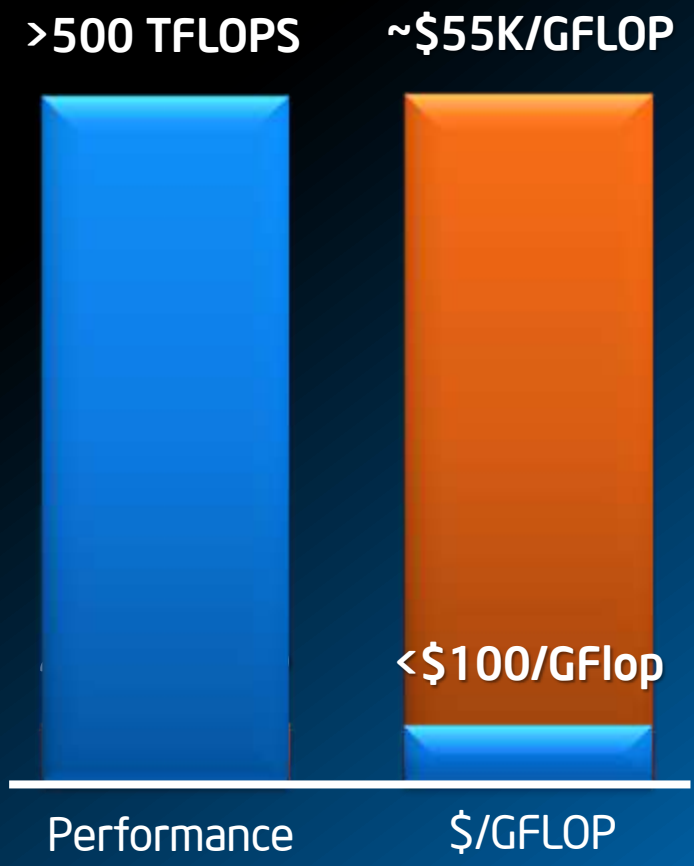

Intel Commitment To Exascale

Efficient Performance

Programming Parallelism

Extreme Scalability

Intel Exascale Commitment:
>100X Performance Of Today At
Only 2X The Power of Today's #1 System
Scaling Today's Software Model

Exascale Requirements

Petascale Machine of 2010: TFLOP of Compute

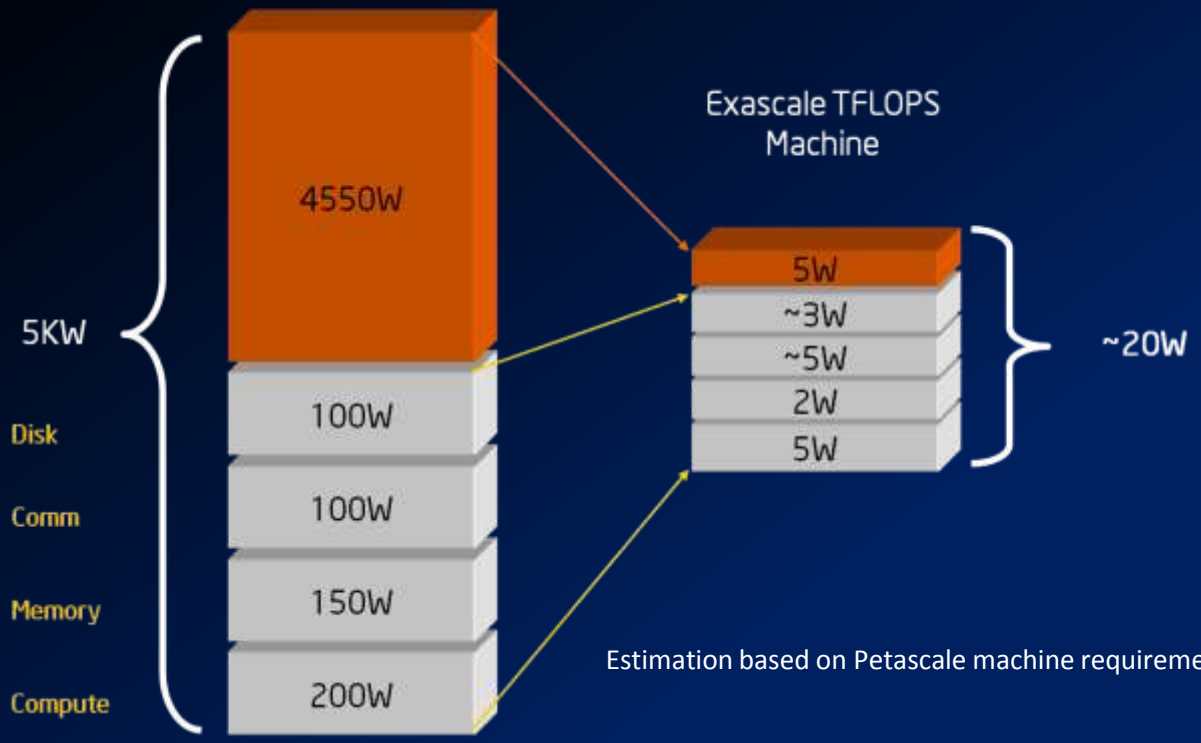

Compute 40x
Memory 75X
Comms 20x
Disk/Storage 33x
Other 900x

Estimation based on Petascale machine requirements circa 2010.

Visceral Focus on System Power Efficiency Improvement

Scaling Programmability

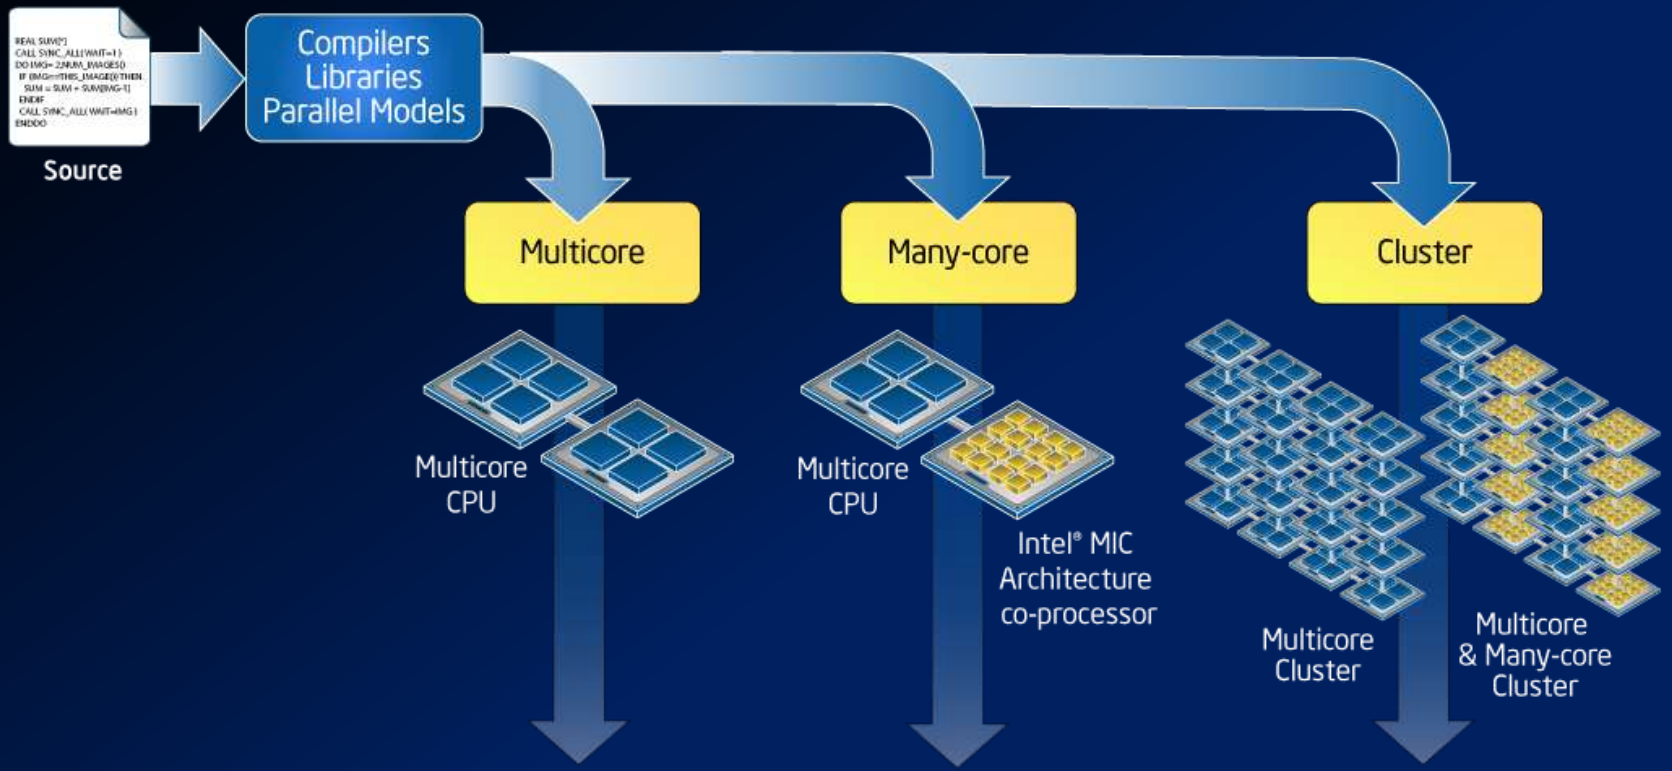

One Programming Model Democratizes Usage

…Avoid Costly Detours

intel

Process Technology Leadership

2003	2005	2007	2009	2011
90 nm	65 nm	45 nm	32 nm	22

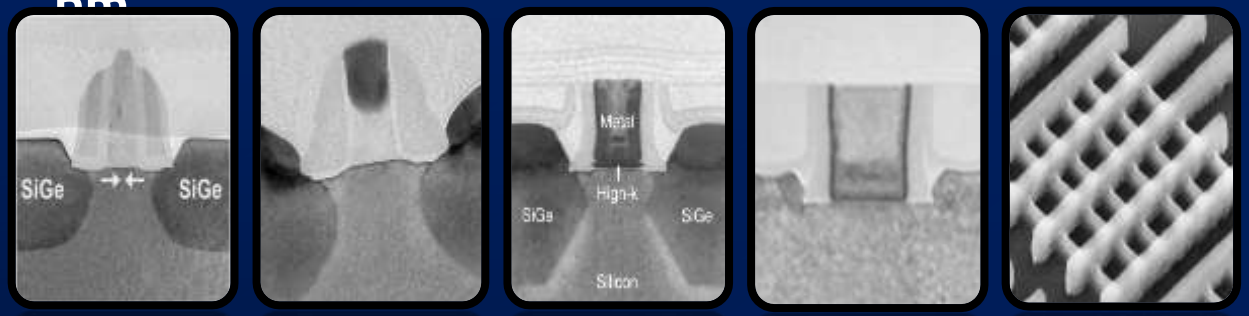

Invented SiGe Strained Silicon	2nd Gen. SiGe Strained Silicon	Invented Gate-Last High-k Metal Gate	2nd Gen. Gate-Last High-k Metal Gate	First to Implement Tri-Gate

STRAINED SILICON

HIGH-k METAL GATE

TRI-GATE

The foundation for all computing

22nm

A Revolutionary Leap in Process Technology

37%
Performance Gain at Low Voltage*

>50%
Active Power Reduction at Constant Performance*

Source: Intel
*Compared to Intel 32nm Technology

Intel Labs & HPC

Strong Research Partnerships

Industry

Government

Universities

World Class Research in HPC

Memory Stacking & Technologies

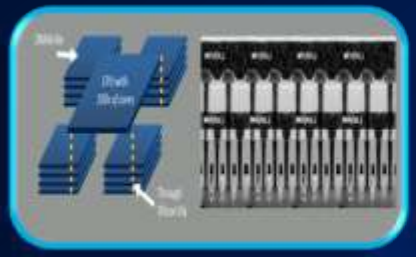

Silicon Photonics

Security

Programmability

Interconnect Technologies

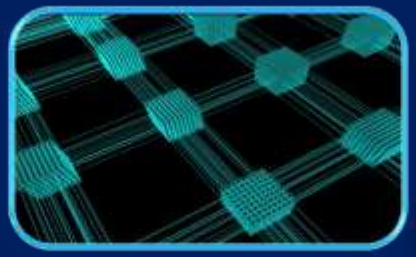

Power Reduction

Delivering Breakthrough Technologies to Fuel Innovation

Continuing The Journey: Next Intel® Xeon® Processor Codenamed Sandy Bridge-EP

Growing Performance

- Up to 8 cores per socket
- 2X FLOPS with Intel® Advanced Vector Extensions

Efficient I/O

- Integrated PCIe reduces latency and power

The Foundation of the Innovation in Science and Technology

intel

Highly Parallel Performance

Intel® Many Integrated Core (Intel® MIC) Architecture

Delivered Performance

Launching on 22nm with >50 cores to provide outstanding performance for HPC users

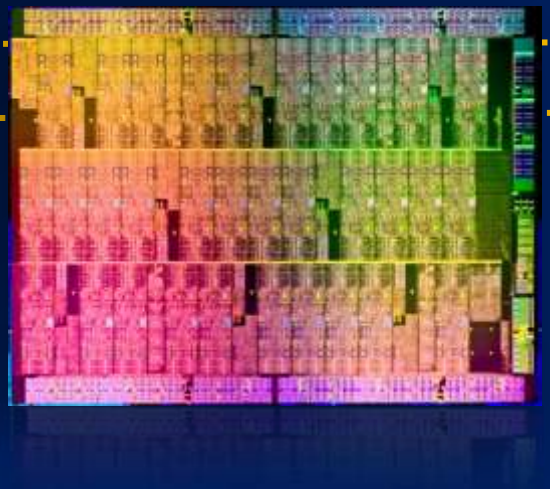

Performance Density

The compute density associated with specialty accelerators for parallel workloads

Programmability

The many benefits of broad Intel CPU programming models, techniques, and familiar x86 developer tools

A Step Forward In Dealing With Efficient Performance & Programmability

intel


lrz

Evaluating the Intel MIC Architecture

Arndt Bode
Leibniz Supercomputing Centre, Germany
with input from Iris Christadler, Alexander Heinecke and Volker Weinberg

June 2011, ISC, Hamburg

Preface

- Programming models are the key to harness the computational power of massively parallel devices.
- Obviously, Intel has realized this trend and substantially supports open standards and invests in innovative programming models.
- LRZ and TUM are using Intel hard- and software for many years and know the tool chain by heart.
- We expect: A hardware product that delivers good performance (and energy-efficiency) without loosing programmability.

Advantages of the MIC Architecture

- Is a standard x86 architecture!
- Allows many different parallel programming models like OpenMP, MPI and Intel Cilk!
- Offers standard math-libraries like Intel MKL!
- Supports whole Intel tool chain, e.g. Compiler & Debugger!

➢ ***Writing MIC-accelerated code with minimal effort and great performance***

Workloads under Investigation

- Euroben Kernels (7 dwarfs of HPC)
- Data Mining
- TifaMMy – Matrix Operations (Demo here at ISC'11!)
- Further Linear Algebra and Simulation Codes

Euroben Kernels

- Selected micro-benchmarks used in PRACE for the evaluation of accelerator hardware & new languages: http://www.prace-project.eu/documents/public-deliverables/d6-6.pdf

‹ multiplication

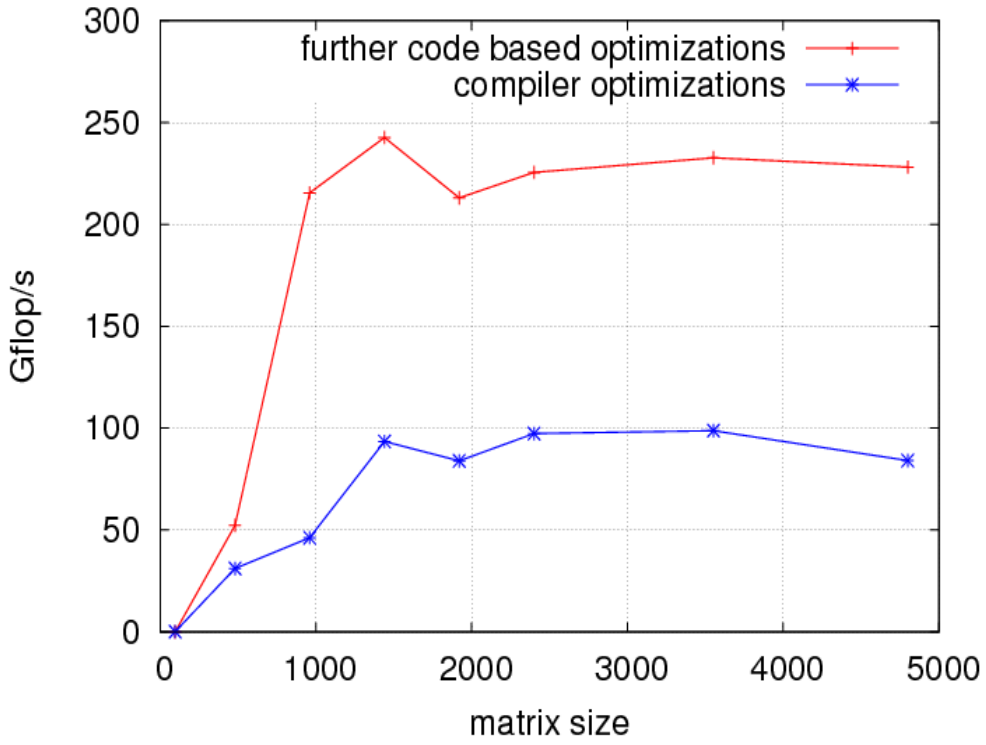

Performance evaluation of mod2am on KNF with 30 cores @1050 MHz using Intel's Offload Compiler, single precision, data transfer times excluded

Data Mining with Adaptive Sparse Grids

- Machine learning algorithm
- Learning function from a training dataset
- Important workload for classification and regression of huge datasets

➢ *MIC-Execution: Straightforward*
  - ➢ *First version within a few hours*
  - ➢ *Optimized version took **2** days*

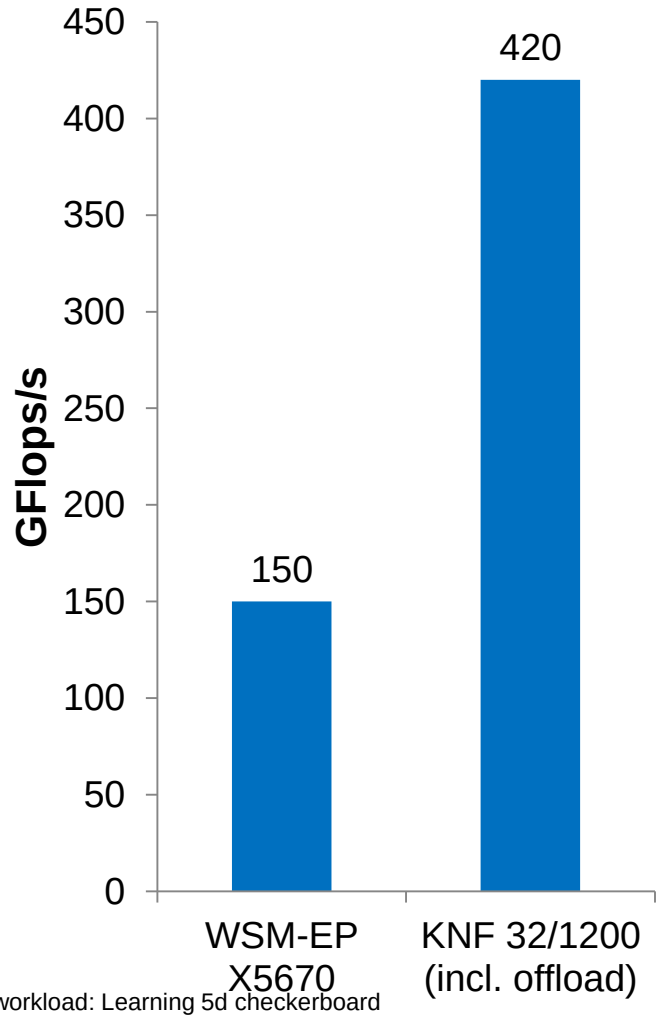

Testworkload: Learning 5d checkerboard with 262144 instances and classification accuracy of 92%

TifaMMy – Idea and Application

- TifaMMy: self-adaptive and cache-oblivious framework for matrix operations optimized on fat x86 cores
- This is done by nested recursions and vectorized kernels
  - On MIC only the kernels were changed, MIC's x86 cores are able to tackle nested recursions!

➢ *parallelization scheme employing OpenMP can be reused*

➢ *having SSE kernels, bringing code to MIC is nearly for free*

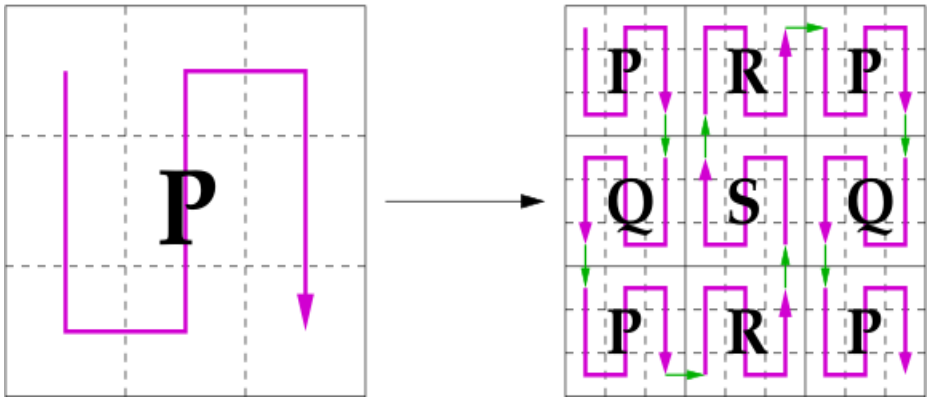

TifaMMy – Performance Matrix Multiplication

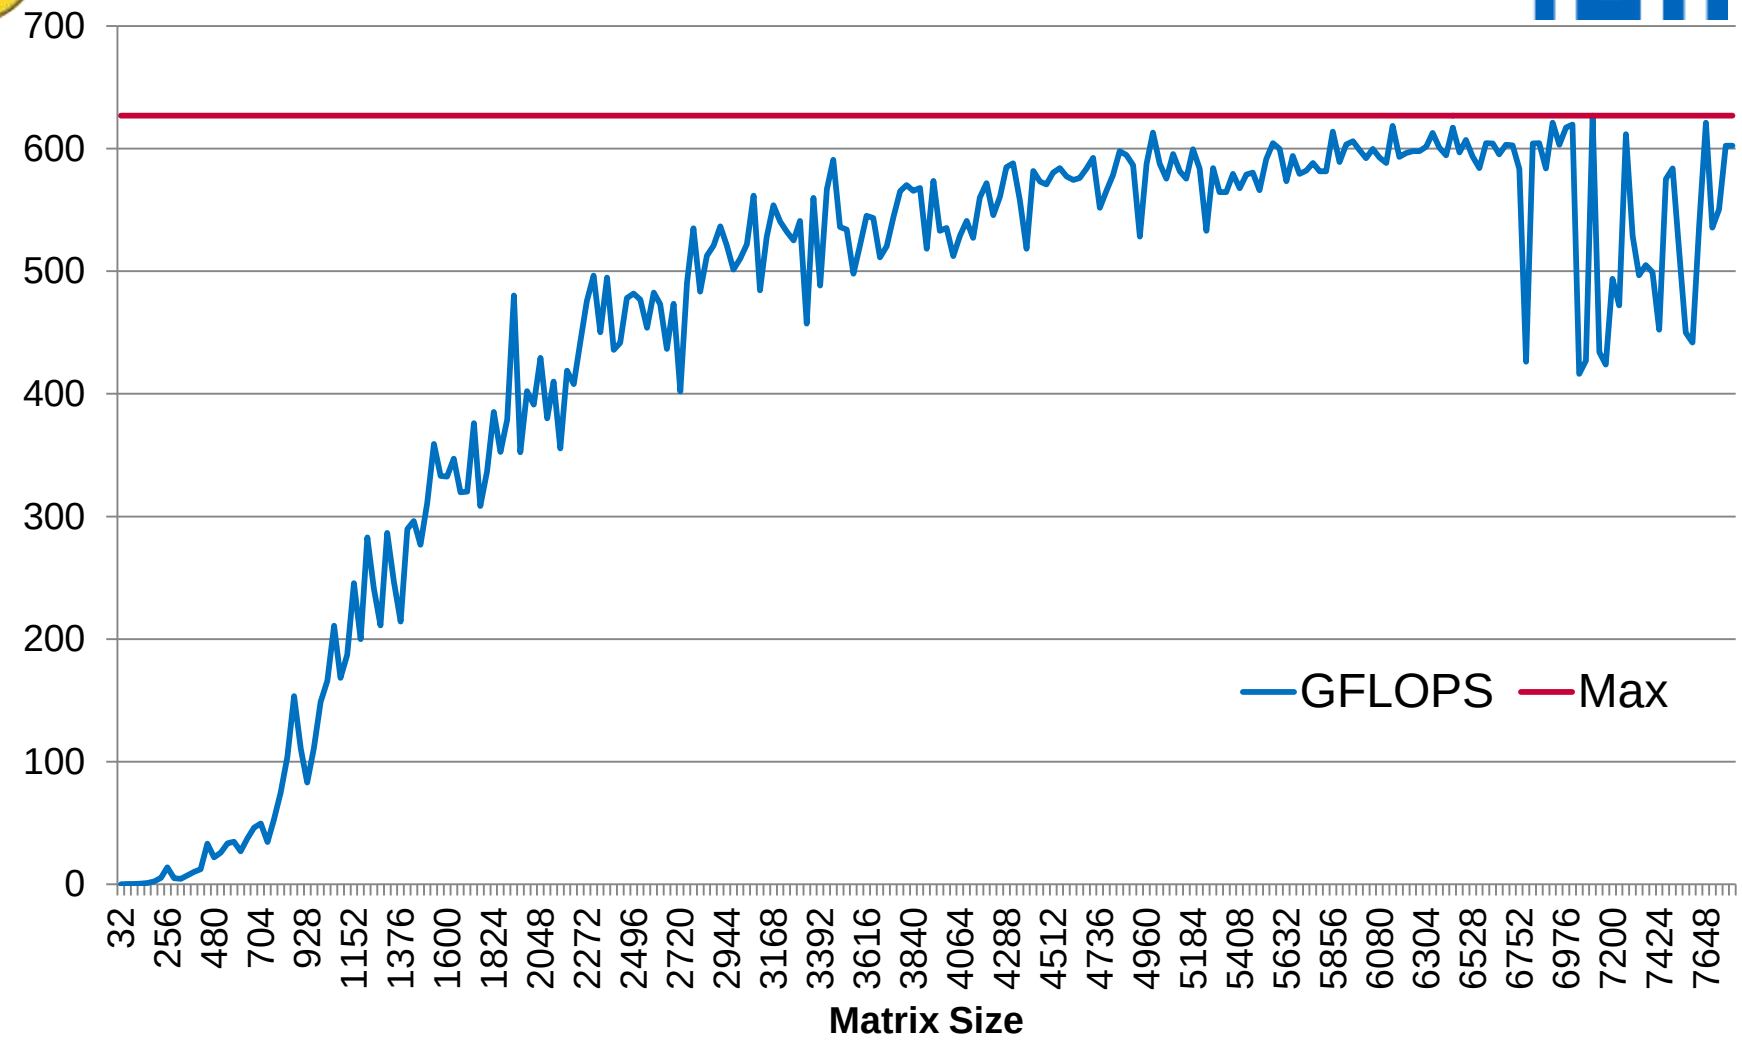

Testworkload: TifaMMy
Executed on KNF with
32 cores @1200 MHz

Advantages of the MIC Architecture

- Is a standard x86 architecture!
- Allows many different parallel programming models like OpenMP, MPI and Intel Cilk!
- Offers standard math-libraries like Intel MKL!
- Supports whole Intel tool chain, e.g. Compiler & Debugger!

➢ ***Pre-release MIC-accelerated code for a typical scientific workload (e.g. Data Mining, TifaMMy) can reach up to 50% of peak performance!***
Visit demo here at ISC'11!

"SGI understands the significance of inter-processor communications, power, density and usability when architecting for exascale. Intel has made the leap towards exaflop computing with the introduction of Intel® Many Integrated Core (MIC) architecture. Future Intel® MIC products will satisfy all four of these priorities, especially with their expected ten times increase in compute density coupled with their familiar X86 programming environment."

Dr. Eng Lim Goh, SGI CTO

Intel® MIC Architecture : Needed for Exascale

Exaflop by 2018

125x compute power

- 25x : Moore's Law
- 5x : remains

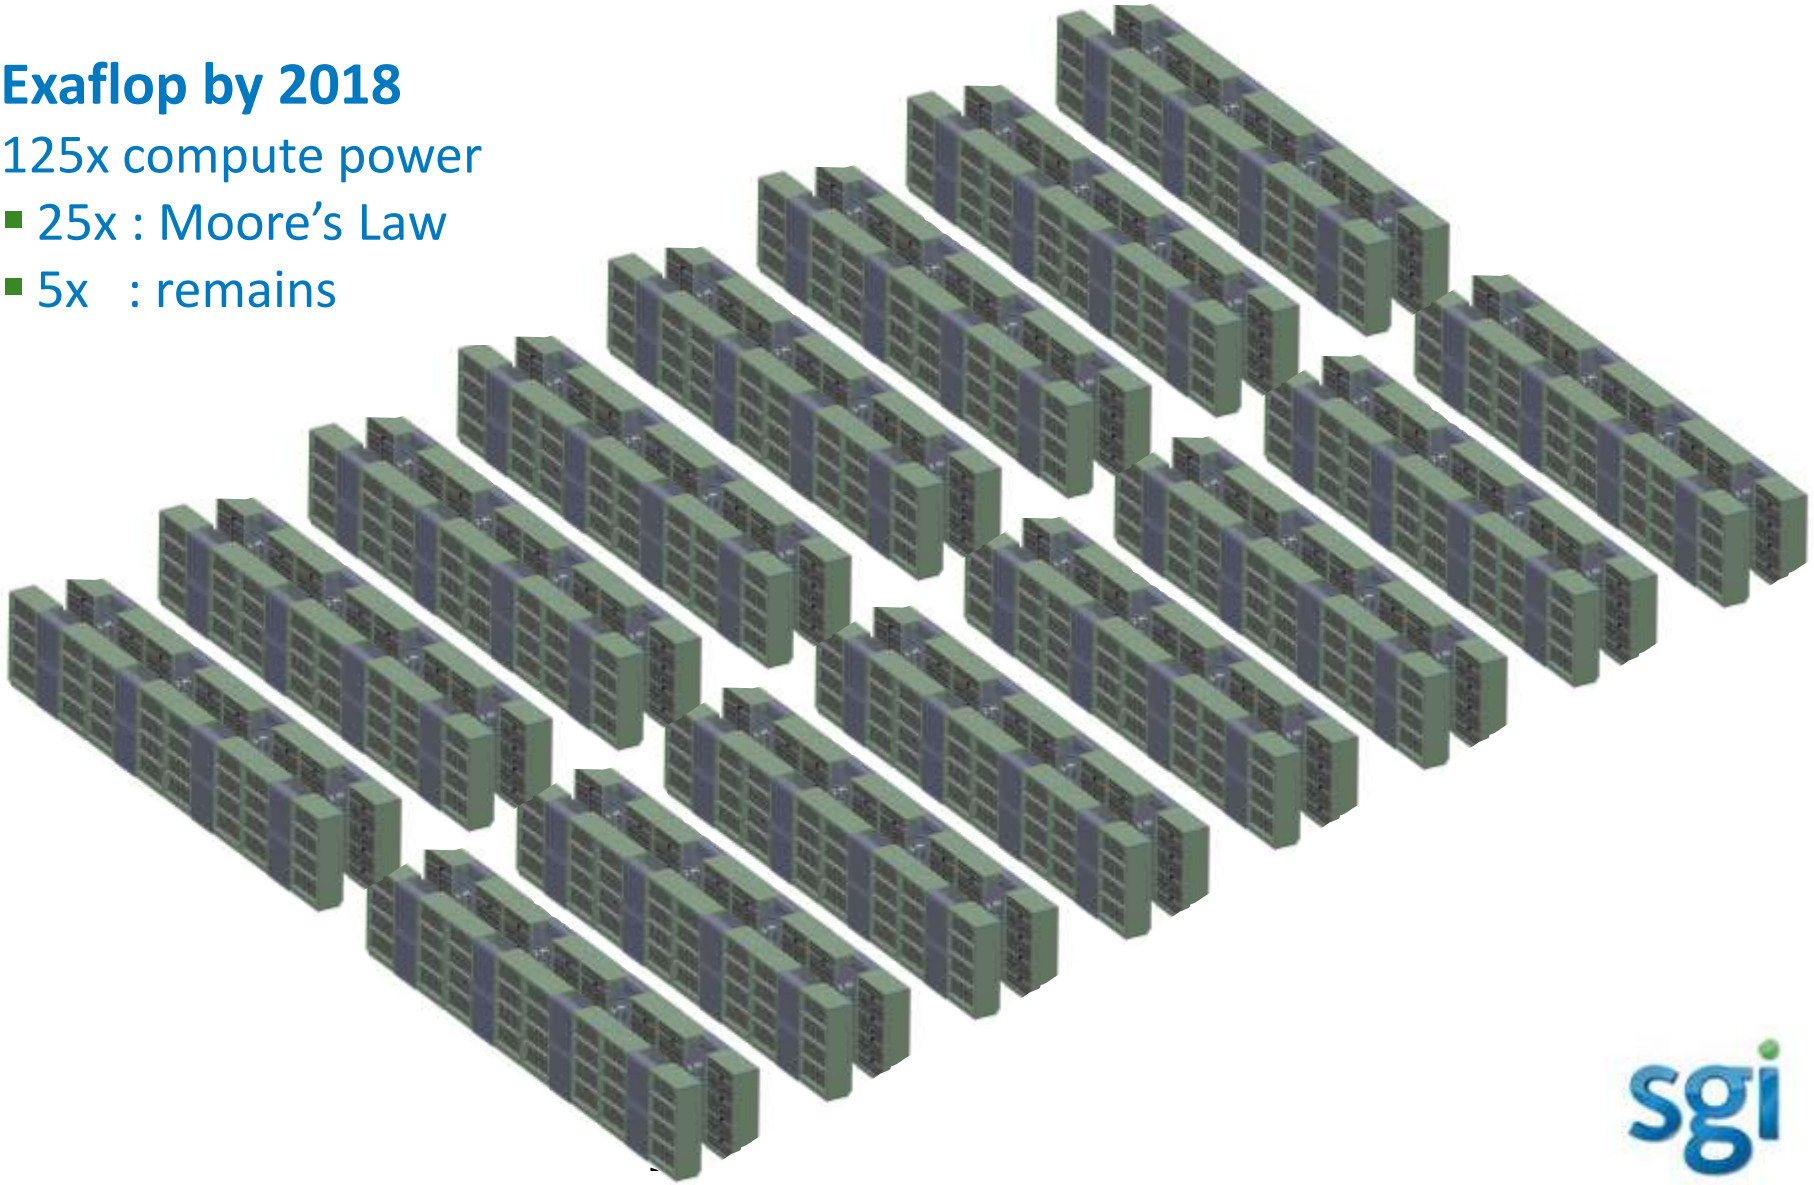

Intel® MIC Architecture : Familiar x86 Programming

Intel® Xeon® processor

Intel® MIC co-processor

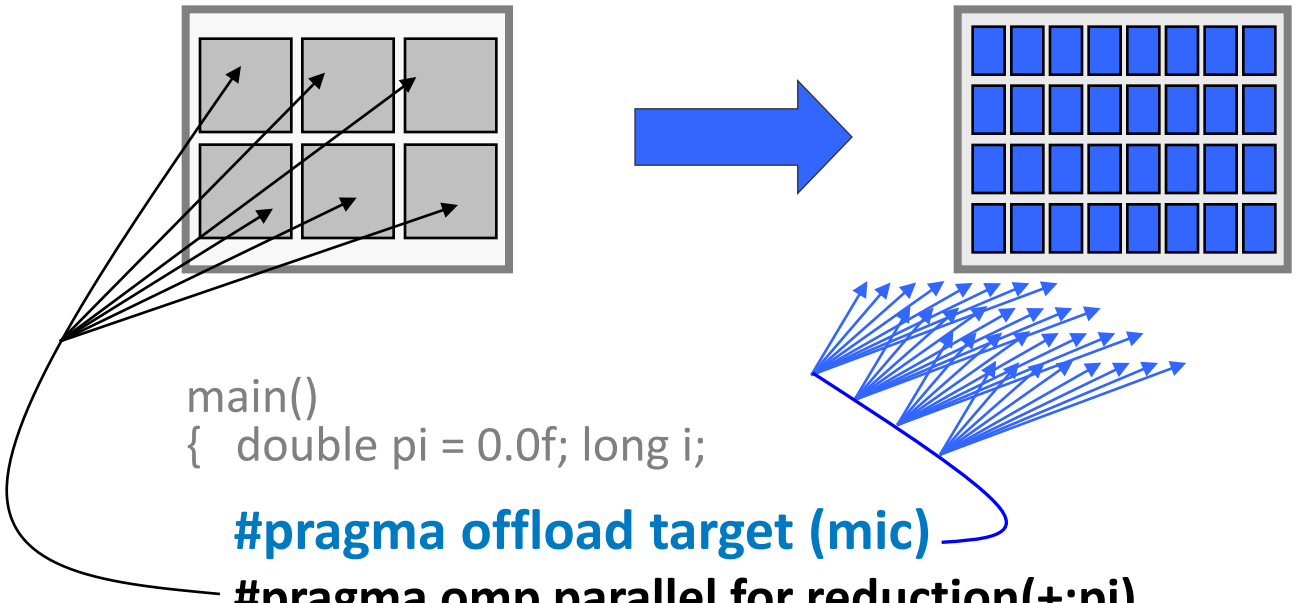

```
main()
{   double pi = 0.0f; long i;

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

   printf("pi = %f\n",pi/N);    }
```


sgi

MIC On Track: ISC Demonstrations[1]

Hybrid LU Factorizarion

Up to 772 GFLOP

Leverages compute power of both Intel® Xeon® CPUs and Intel® MIC
Delivers optimal performance by dynamically balancing large and small matrix Computations between Intel® Xeon® and Intel® MIC

Hybrid Computing – SGEMM with Intel® MKL

1+ TFLOP

High performing SGEMM with just 18 lines of code – common between Intel® Xeon® CPUs and Knights Ferry
Uses Intel® MKL in current version of Alpha stack/tools on Knights Ferry

7.4 TFLOP SGEMM in a node

7.4 TFLOP

Simultaneous execution of SGEMM on 8 Knights Ferry cards to deliver 7.4 TFLOPS in 1 4U server

[1] Refer to backup material for system configurations

Optimized MIC Software Development Platform Performance

MIC On Track: ISC Demonstrations #2[1]

Forschungszentrum Juelich	SMMP Protein Folding	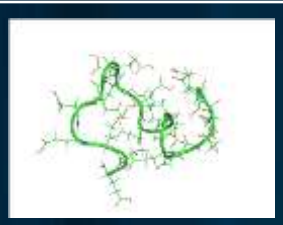	Simulates the folding process of proteins to reach their final shape after they are produced by a cell
KISTI	Molecular Dynamics		Empirical-potential molecular dynamics, widely used for simulating nano-materials including carbon nanotube, graphene, fullerene, and silicon surfaces
LRZ	TifaMMy Matrix Multiplication	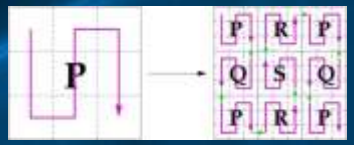	Cache-oblivious implementation of matrix-matrix multiply which uses a recursive scheme to partition input data for computation and parallelization
CERN	Core Scaling of Intel® MIC Architecture	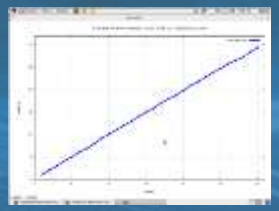	Benchmark kernel extracted from the CBM/ALICE HLT software development for collider experiments. It estimates real trajectories from imprecise measurements

[1] Refer to backup material for system configurations

Programmability For HPC Applications

MIC Partners at International Supercomputing 2011

Showcasing Applications

Showcasing Platforms with Knights Ferry

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

How Intel® Delivers its Commitments:

Intel Exascale Commitment:
>100X Performance Of Today At
Only 2X The Power Of Today's #1
Scaling Today's Software Model

Committed roadmap now and in the future

Flexible, open and scalable programming models

Collaborating with others to ensure the exascale future

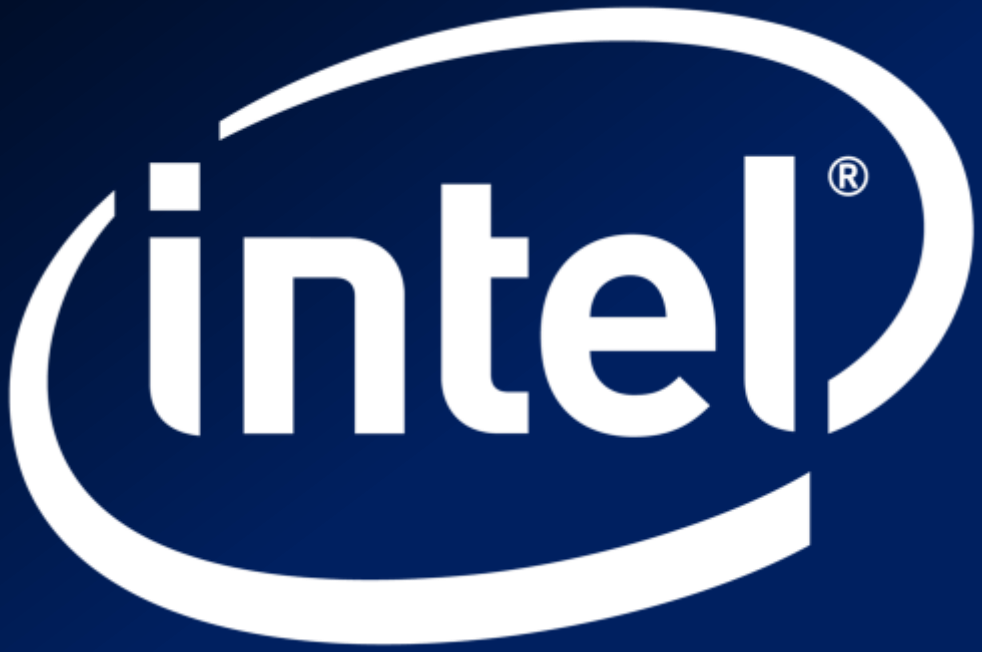
intel


intel

System Configuration
7 TFLOPS SGEMM in a node

- Colfax Model: CXT8000 Server w/Intel® 5520 chipset and 4 PLX PEX8647 Gen 2 PCIe switches
- Intel Alpha level software (Intel® Compilers, drivers etc.)

HW specifications	
8 x KNF	D0 Si @1.2GHz, 2GB GDDR5@3.6GT/s
Host	Colfax CXT8000: 2 socket platform with 2 Intel® Xeon® processor X5690 (3.46GHz, 6 cores, 12MB L3 cache) with 24GB DDR3 @1333MHz, Dual Intel® 5520 IOH, OS RHEL 6.0

KNF SW Stack
Larrabee kernel driver ver. 1.6.197
Flash Image/uOS: 1.0.0.1137/1.0.0.1137-EXT-HPC
Offload compiler (w/data xfer): Composer XE for MIC 0.043
Native compiler (w/o data xfer): Version Alpha Build 20110518

System Configuration
Hybrid Computing with Intel® MKL

- Knights Ferry Software Development Platform (Shady Cove)
- Intel Alpha level software (Intel® Compilers, Intel® MKL, drivers etc.)

HW specifications	
1 x KNF	D0 Si @1.2GHz, 2GB GDDR5@3.6GT/s
Host	Shady Cove 2 socket platform with 2 Intel® Xeon® processor X5680 (3.33GHz, 6 cores, 12MB L3 cache) with 24GB DDR3@1333MHz, single Intel® 5520 IOH, OS: RHEL 6.0

SW specifications	
MKL4KNF	MKL KNF.b2 build 20110518
MKL	10.3.3

KNF SW Stack
Larrabee kernel driver ver. 1.6.197
Flash Image/uOS: 1.0.0.1137/1.0.0.1137-EXT-HPC
Offload compiler (w/data xfer): Intel® Composer XE for MIC 0.043
Native compiler (w/o data xfer): Version Alpha Build 20110518

System Configuration
Hybrid Computing LU Factorization

- Knights Ferry Software Development Platform (Shady Cove)
- Intel Alpha level software (Intel® Compilers, drivers etc.)

HW specifications	
1 x KNF	D0 Si @1.2GHz, 2GB GDDR5@3.6GT/s
Host	Shady Cove 2 socket platform with 2 Intel® Xeon® processor X5680 (3.33GHz, 6 cores, 12MB L3 Cache) with 24GB DDR3@1333MHz, single Intel® 5520 IOH,

KNF SW Stack
Larrabee kernel driver ver. 1.6.197
Flash Image/uOS: 1.0.0.1137/1.0.0.1137-EXT-HPC
Offload compiler (w/data xfer): Intel® Composer XE for MIC 0.043
Native compiler (w/o data xfer): Version Alpha Build 20110518

System Configuration
KISTI Molecular Dynamics

- Dell Precision Workstation
- Intel Alpha level software (Intel® Compilers, drivers etc.)

HW specifications	
1 x KNF	C0 Si @1.2GHz, 2GB GDDR5@3.0GT/s
Host	Dell Precision Workstation 1 socket platform with 1 Intel® Xeon® processor X5620 (4 cores, 2.4GHz, 12MB L3 cache) with 24GB DDR3@1333MHz, single Intel® 5520 IOH, OS: RHEL 6.0

KNF SW Stack
Larrabee kernel driver ver. 1.6.197
Flash Image/uOS: 1.0.0.1137/1.0.0.1137-EXT-HPC
Offload compiler (w/data xfer): Intel® Composer XE for MIC 0.043
Native compiler (w/o data xfer): Version Alpha Build 20110518

System Configuration
CERN openlab: Core Scaling of Intel® MIC Architecture

- SGI H4002 System
- Intel Alpha level software (Intel® Compilers, drivers etc.)

HW specifications	
1 x KNF	C0 Si @1.2GHz, 2GB GDDR5@3.0GT/s
Host	SGI H4002 2 socket platform with 2 Intel® Xeon® processor X5690 (6 cores, 3.46GHz, 12MB L3 cache) with 24GB DDR3@1333MHz, single Intel® 5520 IOH, OS: RHEL 6.0

KNF SW Stack
Larrabee kernel driver ver. 1.6.197
Flash Image/uOS: 1.0.0.1137/1.0.0.1137-EXT-HPC
Offload compiler (w/data xfer): Intel® Composer XE for MIC 0.043
Native compiler (w/o data xfer): Version Alpha Build 20110518

System Configuration
LRZ: TifaMMy Matrix Multiplication

- Knights Ferry Software Development Platform (Shady Cove)
- Intel Alpha level software (Intel® Compilers, drivers etc.)

HW specifications	
1 x KNF	C0 Si @1.2GHz, 2GB GDDR5@3.0GT/s
Host	Shady Cove 2 socket platform with 2 Intel® Xeon® processor X5680 (3.33GHz, 6 cores, 12MB L3 Cache) with 24GB DDR3@1333MHz, single Intel® 5520 IOH, OS: RHEL 6.0

KNF SW Stack
Larrabee kernel driver ver. 1.6.197
Flash Image/uOS: 1.0.0.1137/1.0.0.1137-EXT-HPC
Offload compiler (w/data xfer): Intel® Composer XE for MIC 0.043
Native compiler (w/o data xfer): Version Alpha Build 20110518

System Configuration
FZ Jülich: SMMP Protein Folding

- Knights Ferry Software Development Platform (Shady Cove)
- Intel Alpha level software (Intel® Compilers, drivers etc.)

HW specifications	
1 x KNF	C0 Si @1.2GHz, 2GB GDDR5@3.0GT/s
Host	Shady Cove 2 socket platform with 2 Intel® Xeon® processor X5680 (3.33GHz, 6 cores, 12MB L3 Cache) with 24GB DDR3@1333MHz, single Intel® 5520 IOH, OS: RHEL 6.0

KNF SW Stack
Larrabee kernel driver ver. 1.6.197
Flash Image/uOS: 1.0.0.1137/1.0.0.1137-EXT-HPC
Offload compiler (w/data xfer): Intel® Composer XE for MIC 0.043
Native compiler (w/o data xfer): Version Alpha Build 20110518