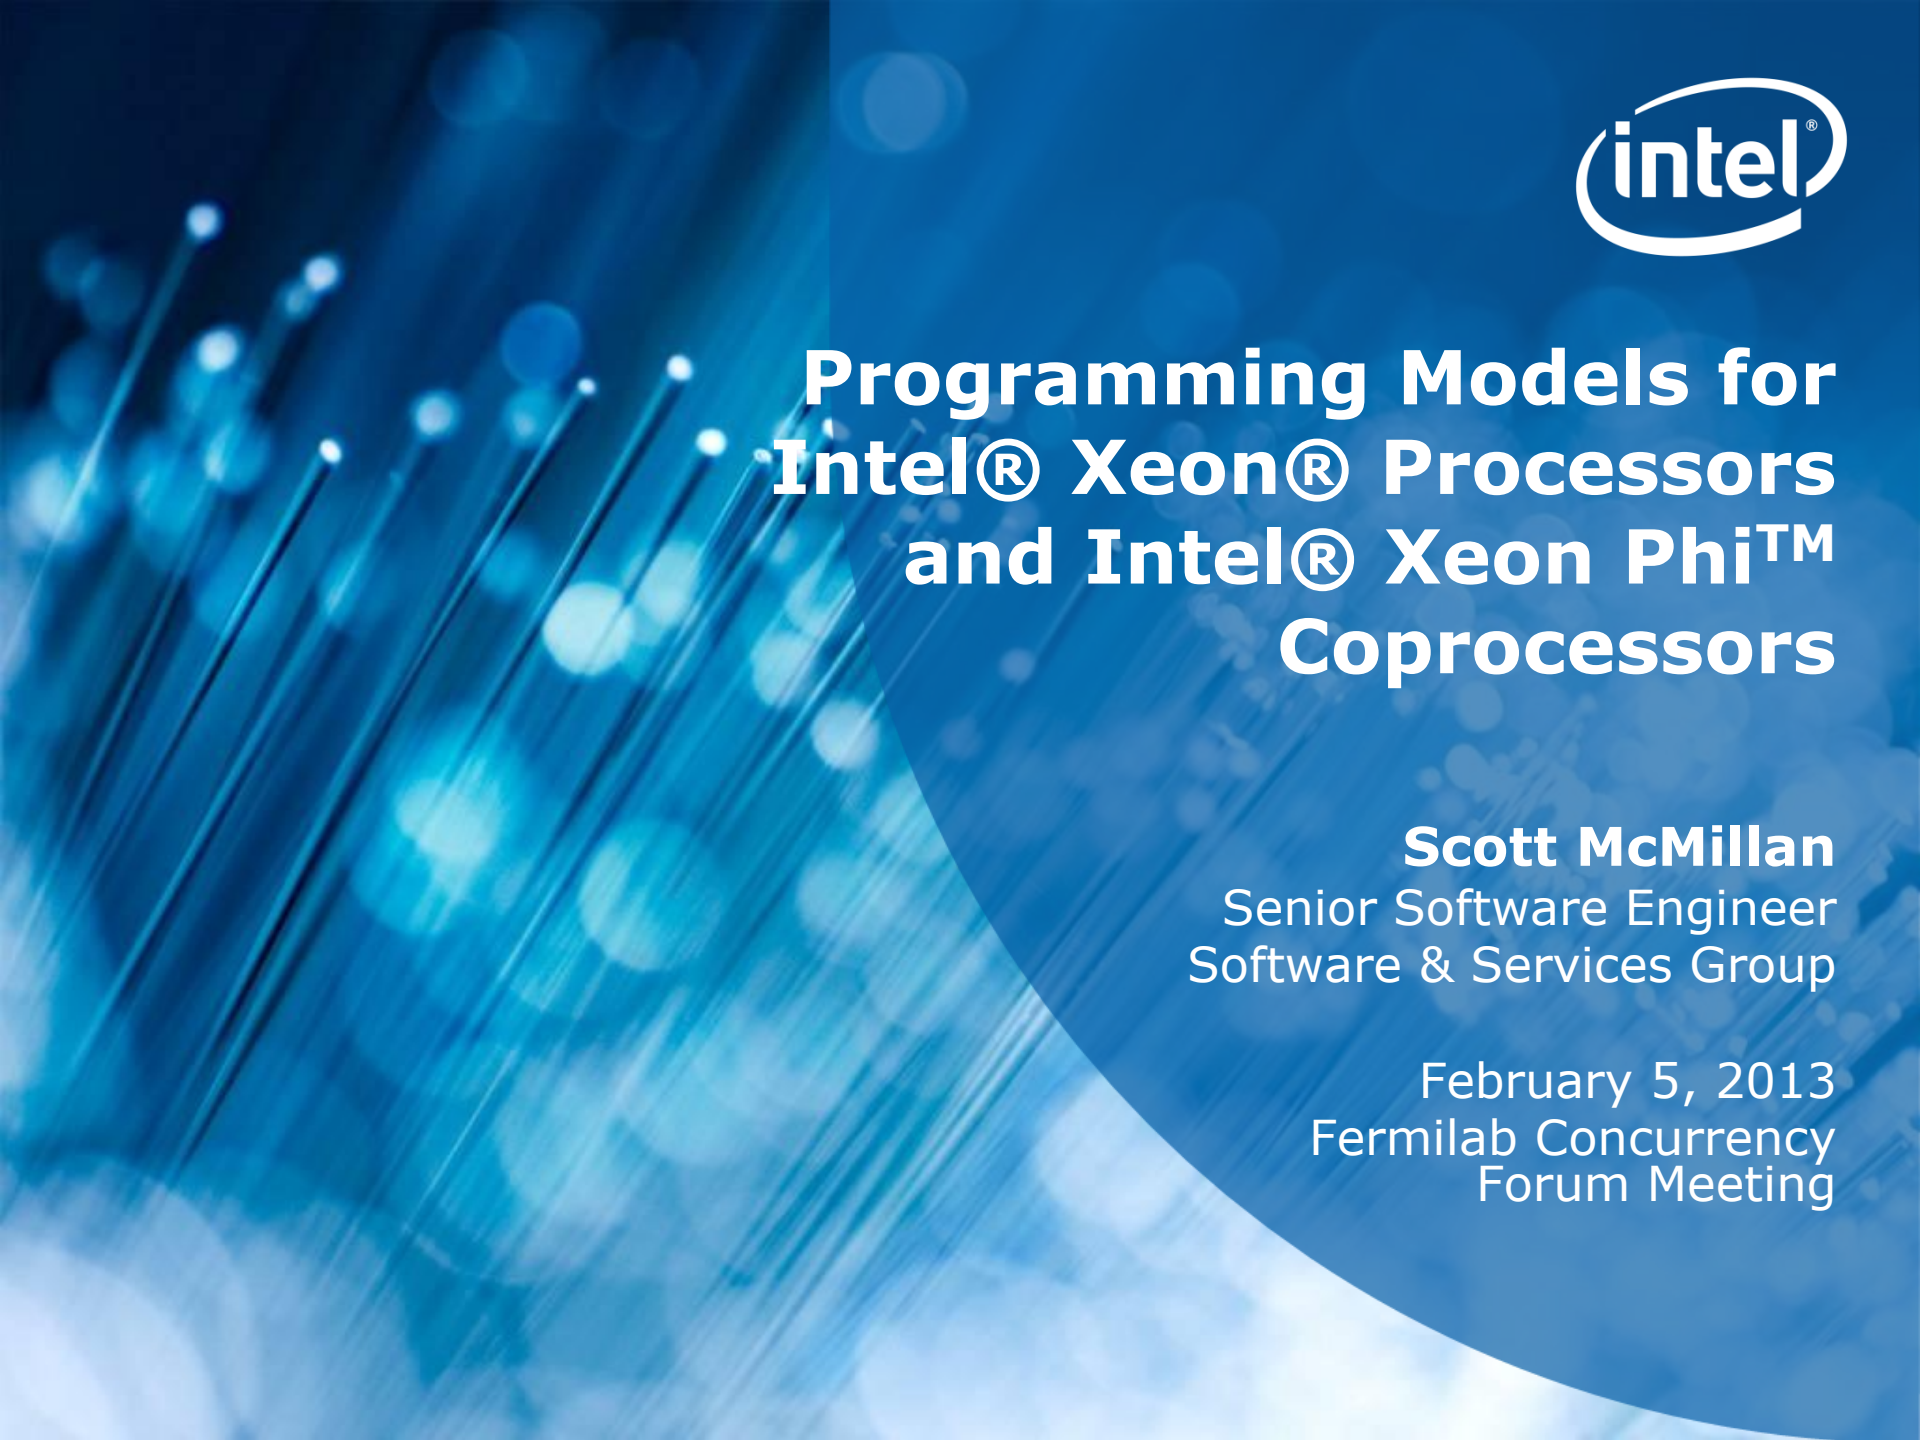

Programming Models for Intel® Xeon® Processors and Intel® Xeon Phi™ Coprocessors

Scott McMillan
Senior Software Engineer
Software & Services Group

February 5, 2013
Fermilab Concurrency Forum Meeting

Legal Disclaimer

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.

http://intel.com/software/products

Optimization Notice

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2®, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

Intel® Xeon® Processors + Intel® Xeon Phi™ Coprocessors: Complimentary Solutions for Parallel Workloads

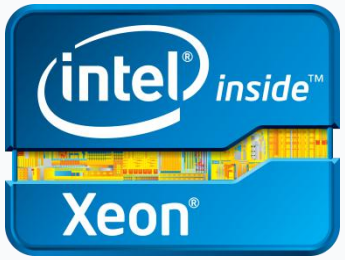

Leadership performance for the majority of server & workstation workloads

Versatile foundation to meet rapid growth in users, devices, and data

Robust energy efficiency, security, and reliability to reduce data center costs

Advanced performance for highly parallel workloads for breakthrough innovation and discovery

Based on Intel® MIC Architecture; Works synergistically with Intel® Xeon® Processors

Increased developer productivity via programming models & tools common with Intel® Xeon® Processors

Develop with Intel tools for Intel® Xeon Processor today, Scale your software investment to include Intel® Xeon Phi™ Products

Shipping January 28, 2013

Intel® Xeon Phi™ Coprocessor 5110P - $2649 RCP

Performance

Up to 1 TFLOP of double-precision (peak)[1]

8GB GDDR5
320 GB/s Bandwidth
Passive form factor at 225W TDP

Programmability

C, C++, Fortran
Intel and 3rd party tools

Intel and
3rd party tools

Applications

Memory Bandwidth / Capacity Bound workloads

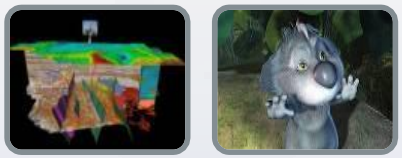

Ideal for
Molecular Modeling,
Digital Content Creation,
and Energy

Ideal for memory bandwidth and memory capacity bound workloads

Stay Tuned in 2013

Intel® Xeon Phi™ Coprocessor 3100 Product Family under $2000 RCP

Performance

Up to 1 TFLOP of double-precision (peak)[1]

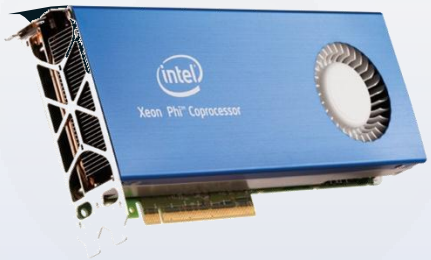

6GB GDDR5
240 GB/s Bandwidth
Active and passive form factors at 300W TDP

Programmability

C, C++, Fortran
Intel and 3rd party tools

Intel and
3rd party tools

Applications

Compute Bound workloads

Ideal for MonteCarlo, Black-Scholes, Life Sciences

Ideal for compute bound workloads

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

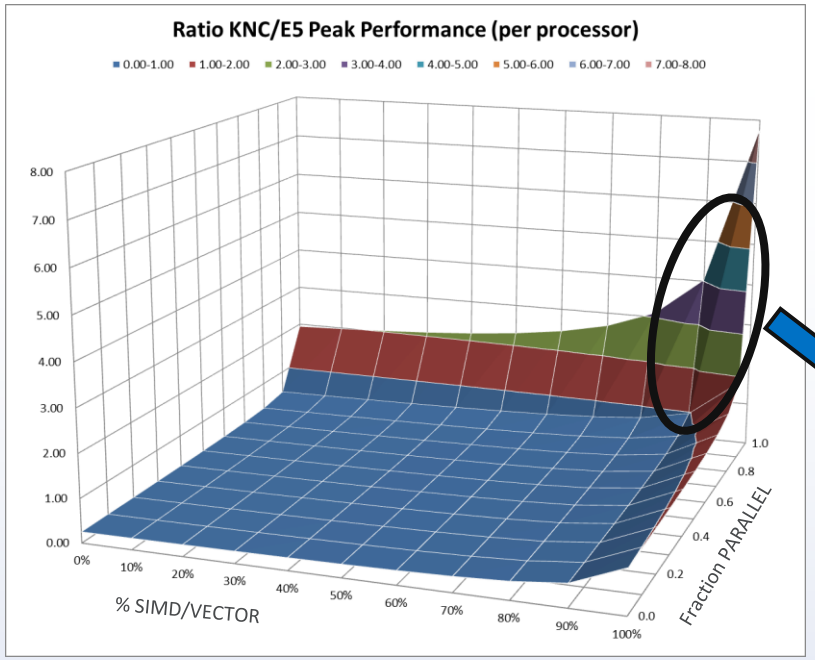

- Intel® Xeon Phi™ coprocessor accelerates highly parallel & vectorizable applications. (graph above)
- Table provides examples of such applications

Application Performance Examples

Customer	Application	Performance Increase[1] vs. 2S Xeon*
Los Alamos	Molecular Dynamics	Up to 2.52x
Acceleware	8th order isotropic variable velocity	Up to 2.05x
Jefferson Labs	Lattice QCD	Up to 2.27x
Financial Services	BlackScholes SP Monte Carlo SP	Up to 7x Up to 10.75x
Sinopec	Seismic Imaging	Up to 2.53x[2]
Sandia Labs	miniFE (Finite Element Solver)	Up to 2x[3]
Intel Labs	Ray Tracing (incoherent rays)	Up to 1.88x[4]

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

Notes:

1. 2S Xeon* vs. 1 Xeon Phi* (preproduction HW/SW & Application running 100% on coprocessor unless otherwise noted)
2. 2S Xeon* vs. 2S Xeon* + 2 Xeon Phi* (offload)
3. 8 node cluster, each node with 2S Xeon* (comparison is cluster performance with and without 1 Xeon Phi* per node) (Hetero)
4. Intel Measured Oct. 2012

Source: Customer Measured results as of October 22, 2012 Configuration Details: Please reference slide speaker notes.

For more information go to http://www.intel.com/performance

Synthetic Benchmark Summary (Intel® MKL) (5110P)

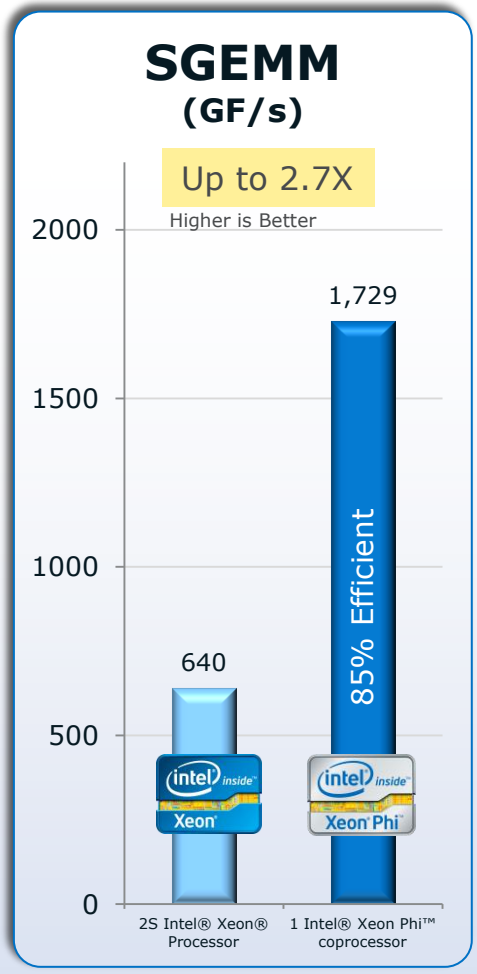

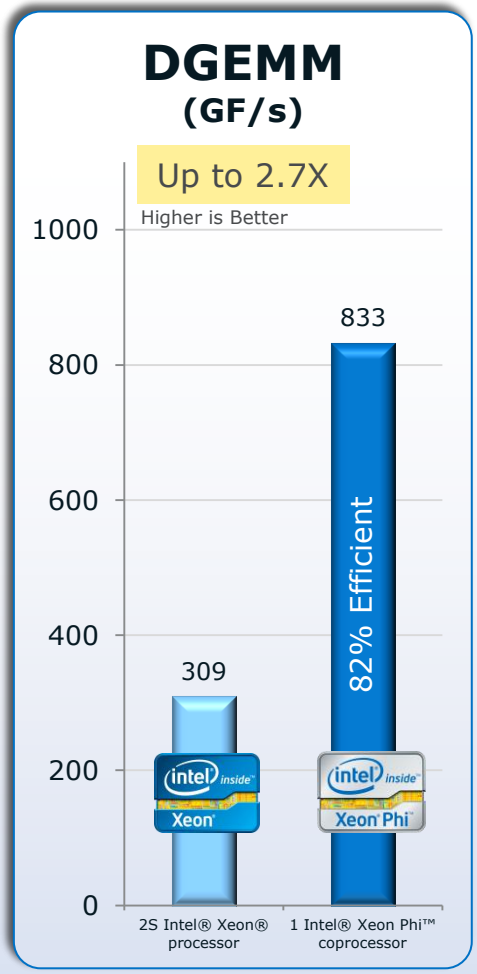

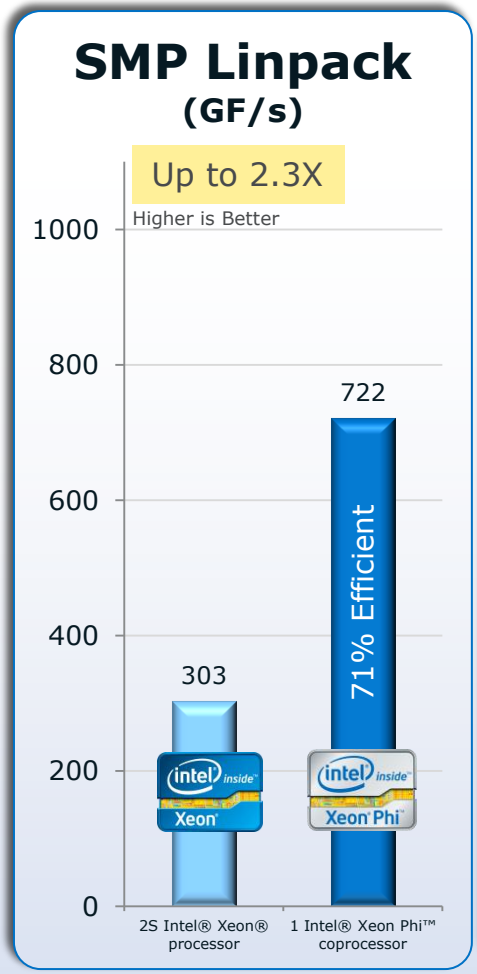

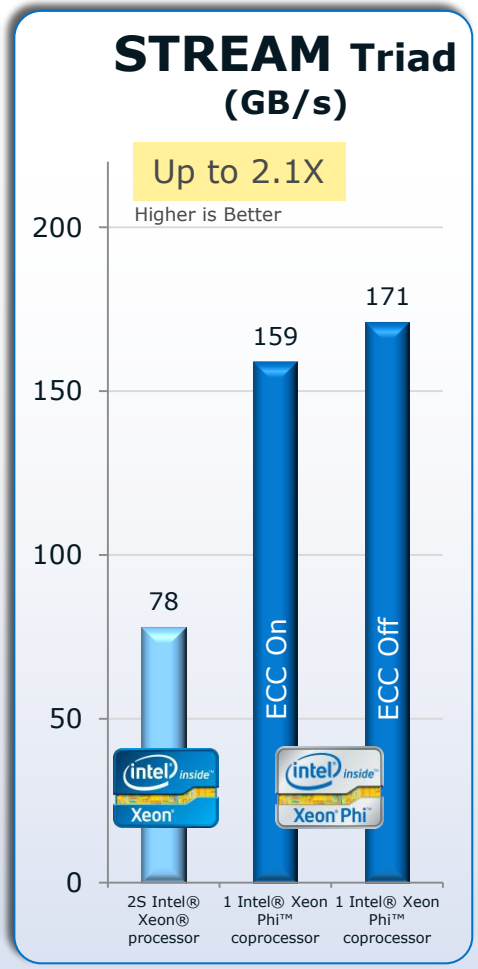

Coprocessor results: Benchmark run 100% on coprocessor, no help from Intel® Xeon® processor host (aka native)

Notes

1. **Intel® Xeon® Processor E5-2670 used for all SGEMM Matrix = 13824 x 13824 , DGEMM Matrix 7936 x 7936, SMP Linpack Matrix 30720 x 30720**
2. **Intel® Xeon Phi™ coprocessor 5110P (ECC on) with "Gold Release Candidate" SW stack SGEMM Matrix = 11264 x 11264, DGEMM Matrix 7680 x 7680, SMP Linpack Matrix 26872 x 28672**

Intel® Xeon Phi™ Product Family
based on Intel® Many Integrated Core Architecture

Optimized, Highly Parallel

Intel® Xeon Phi™ coprocessor

(Pairs with Intel® Xeon® processor host via PCIe)

Runs *Complete* Applications

IP Addressable

Open Source Linux OS

Common Source Code

Standard models of clustering

Builds on / Advances

State of the Art in Parallelism

Intel Developer tools

An SMP on a chip!

Intel® Xeon Phi™ Coprocessor Becomes a Network Node

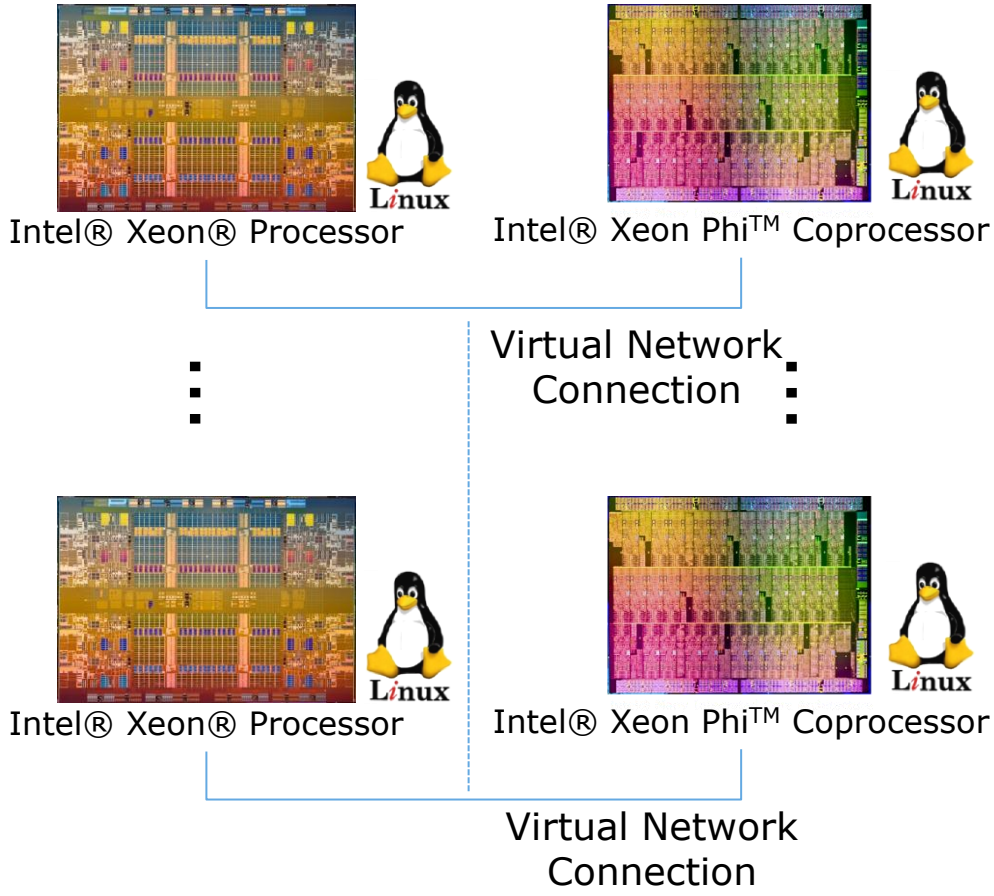

Intel® MIC Architecture + Linux enables IP addressability

Spectrum of Programming Models and Mindsets

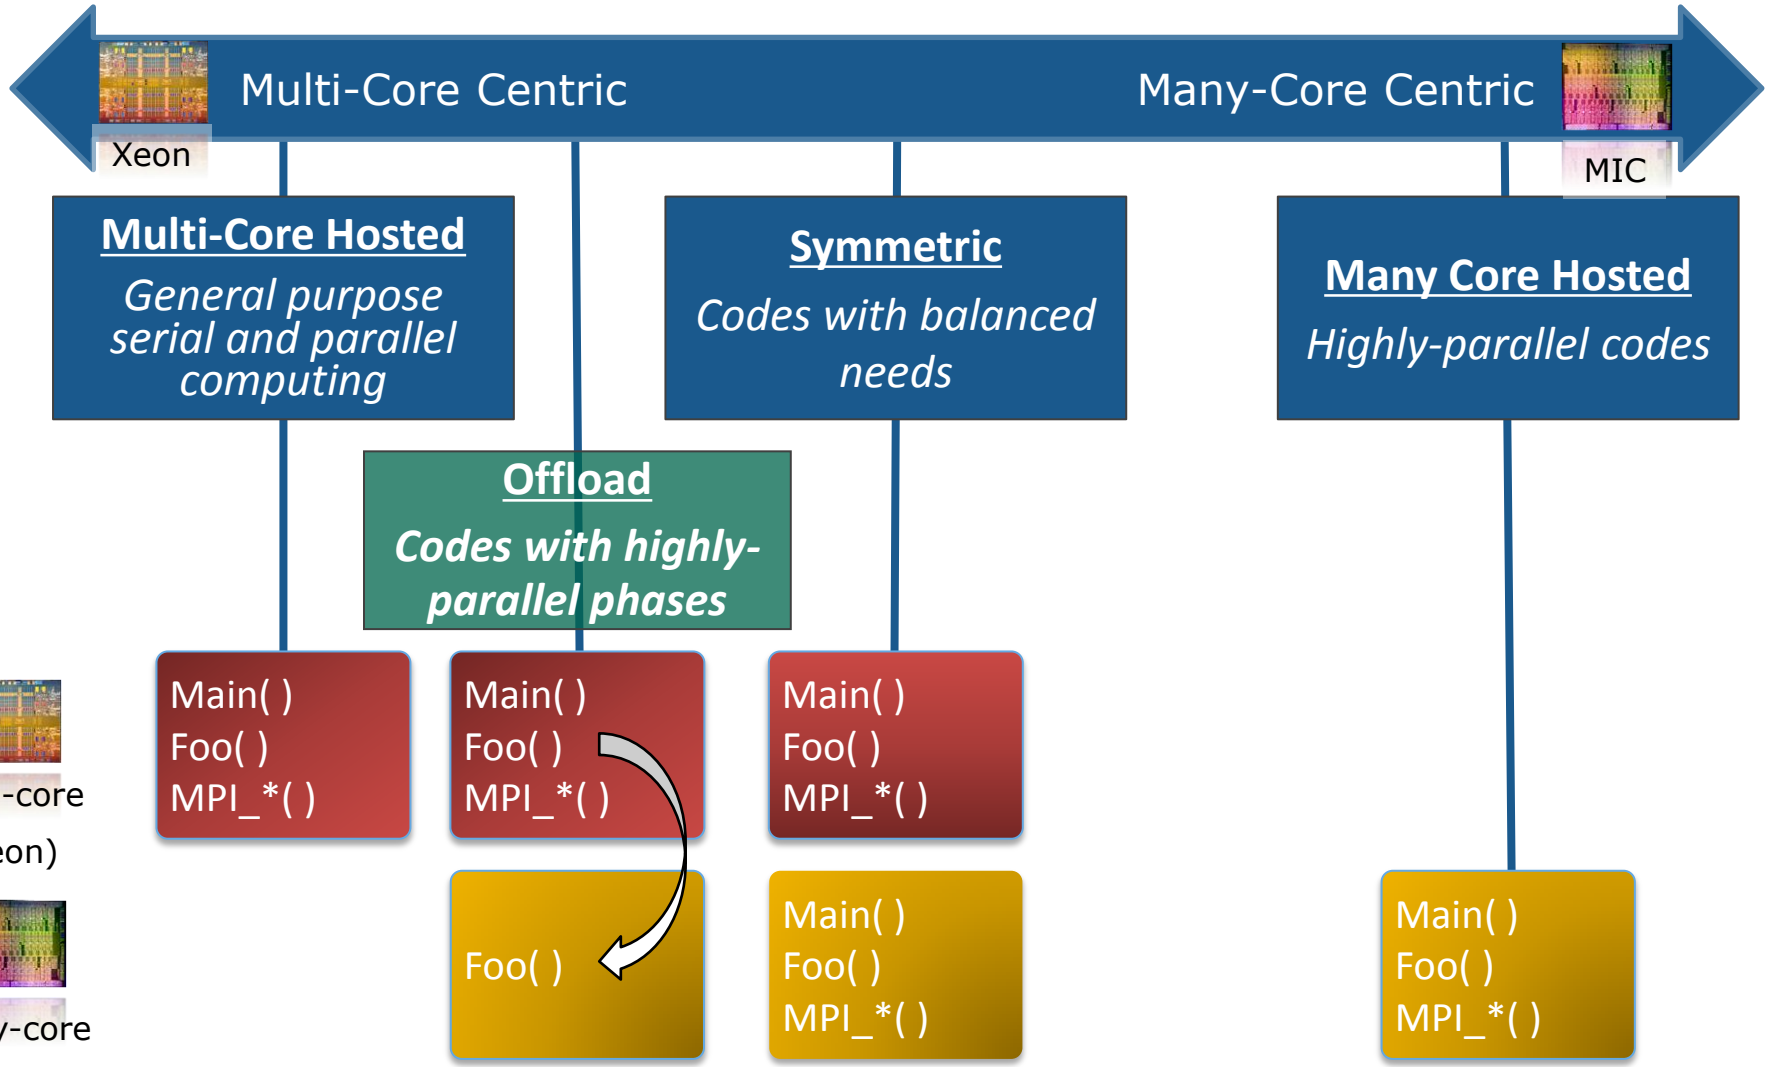

Range of models to meet application needs

Go Parallel with High Performance Math Kernel Library

Intel® Math Kernel Library (Intel® MKL)


```
void foo()  /*  Intel® Math Kernel Library */
{
      float *A, *B, *C; /* Matrices */
      sgemm(&transa, &transb, &N, &N, &N, &alpha, A, &N, B, &N, &beta, C, &N);
}
```

Implicit automatic offloading requires no code changes, simply link with the offload MKL Library

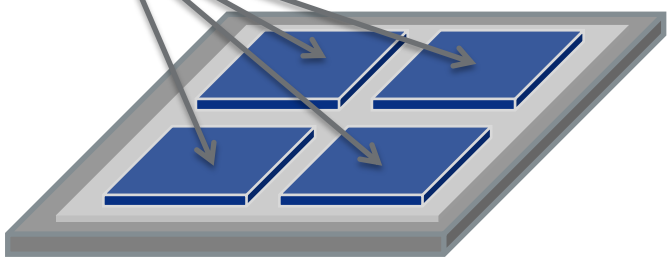

Intel® Xeon® processor

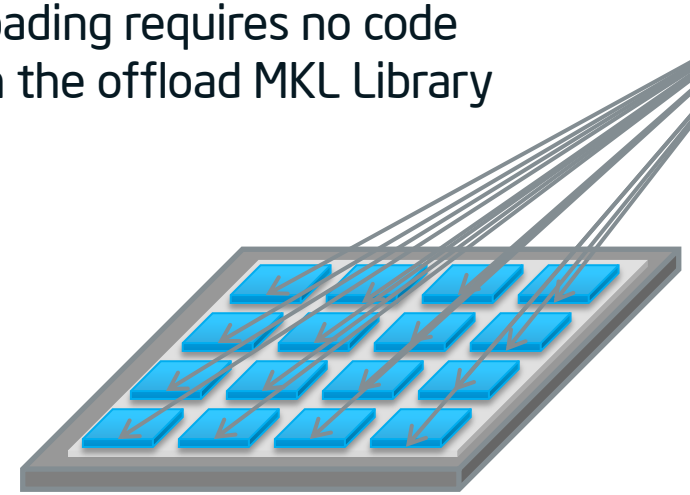

Intel® Xeon Phi™ coprocessor

Intel High Performance Math Kernel Library is Applicable to Multicore and Many-core Programming

Go Parallel with OpenMP*

Intel® C/C++ and Fortran Compilers

(C Example)


```
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);     }
```

One Line Change to Offload to the Intel® Xeon Phi™ coprocessor

Intel® Xeon® processor

Intel® Xeon Phi™ coprocessor

OpenMP* is Applicable to Multicore and Many-core Programming

Go Parallel with Message Passing Interface (MPI)

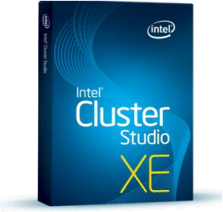

Intel® MPI Library

Extend your cluster solutions to the Intel® Xeon Phi™ coprocessor

- E.g., Intel Xeon Phi™ coprocessor in every node of the cluster using Intel® MPI and Intel® Threading Building Blocks and/or Intel® Cilk™ Plus on nodes
- Same model as an Intel® Xeon processor based cluster .

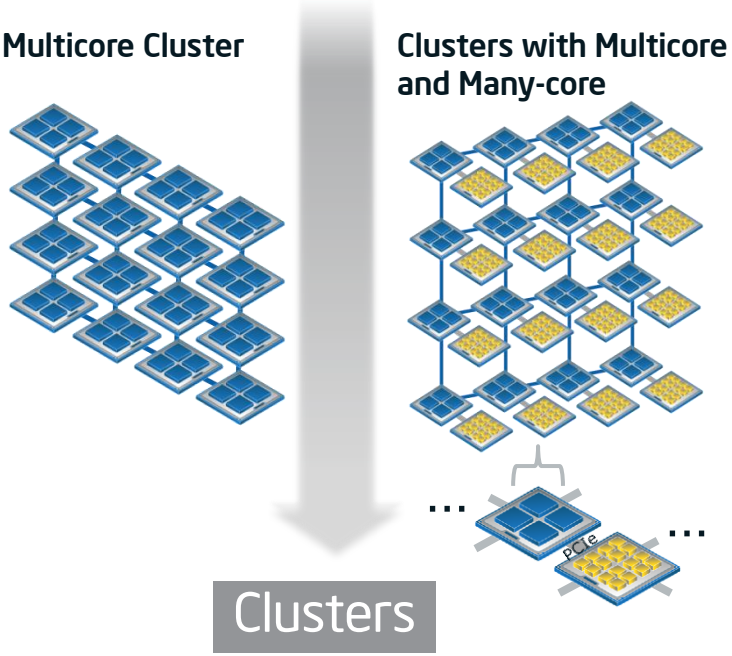

Intel is a leading vendor of MPI implementations and tools

Learn more at http://intel.com/go/mpi

MPI is applicable to Multicore and Many-core Programming

Improving Load Balance: Real World Case

Collapsed data per node and MIC card

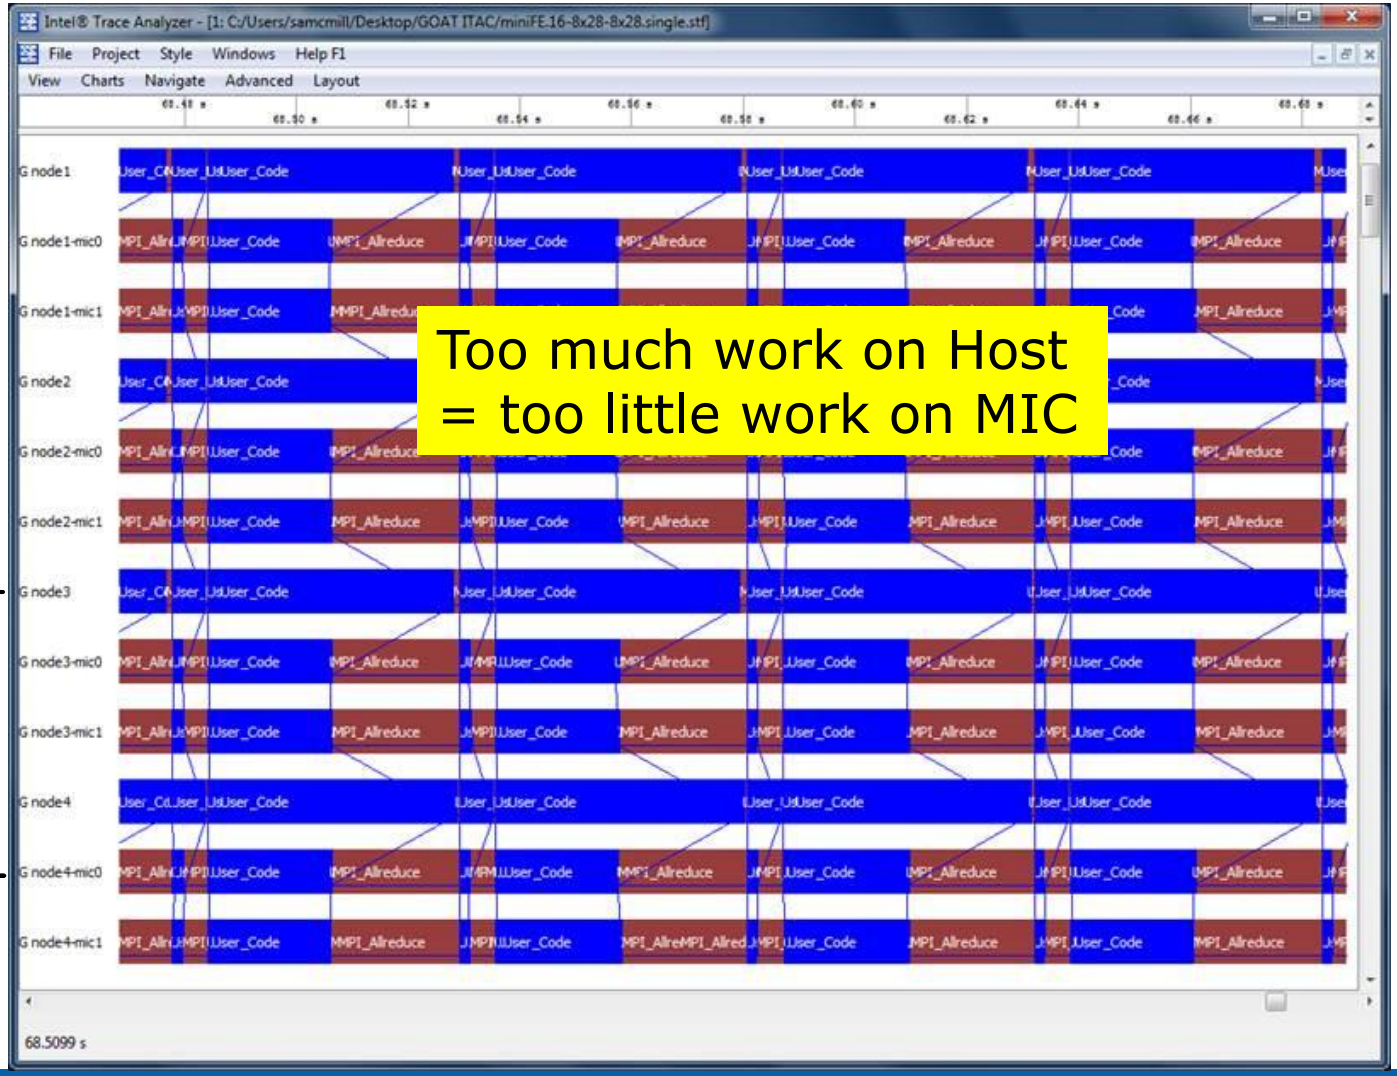

Host
16 MPI procs x
1 OpenMP thread

MIC
8 MPI procs x
28 OpenMP threads

Improving Load Balance: Real World Case

Collapsed data per node and MIC card

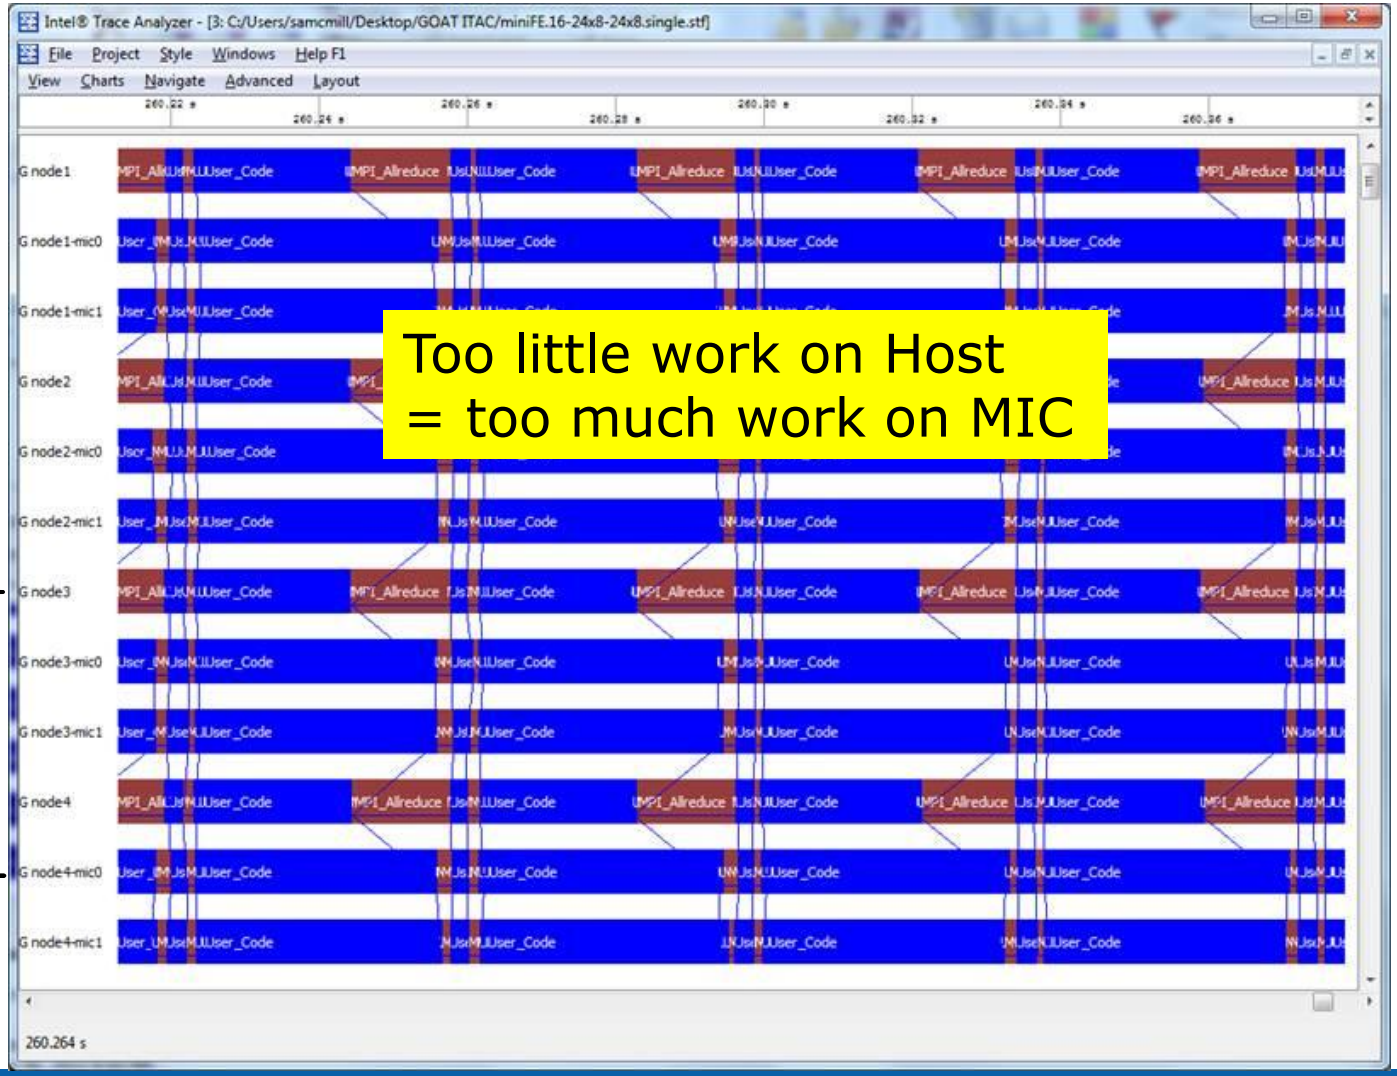

Host
16 MPI procs x
1 OpenMP thread

MIC
24 MPI procs x
8 OpenMP threads

Improving Load Balance: Real World Case

Collapsed data per node and MIC card

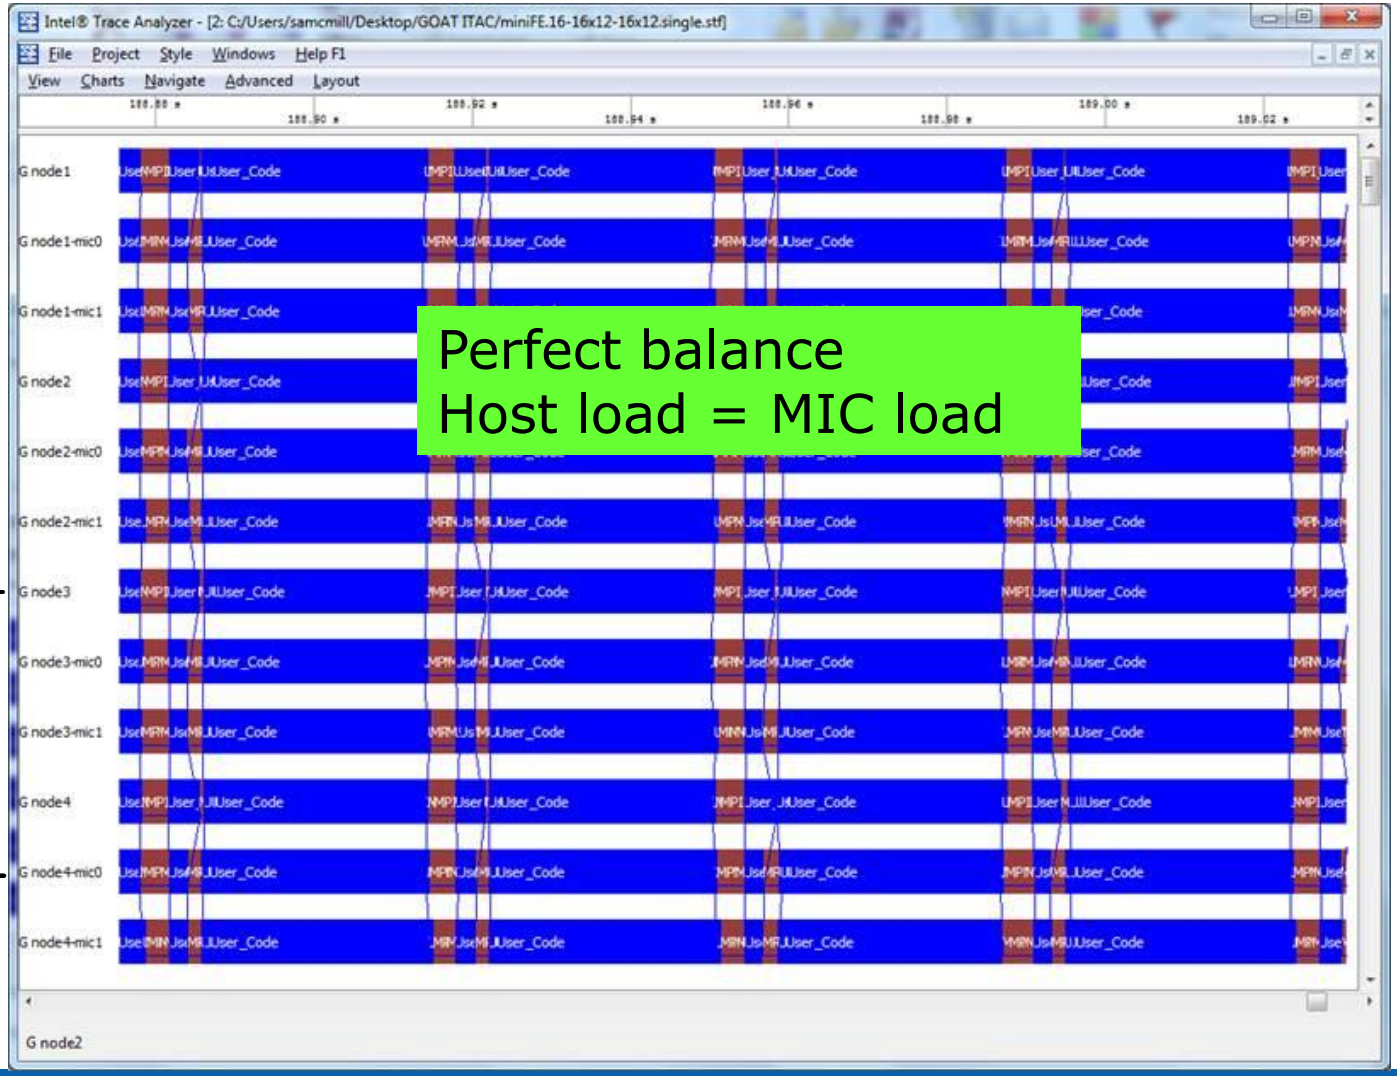

Host
16 MPI procs x
1 OpenMP thread

MIC
16 MPI procs x
12 OpenMP thrds

Intel® Xeon® Phi™ Product Family: Game Changer for HPC Performance & Programmability

Source

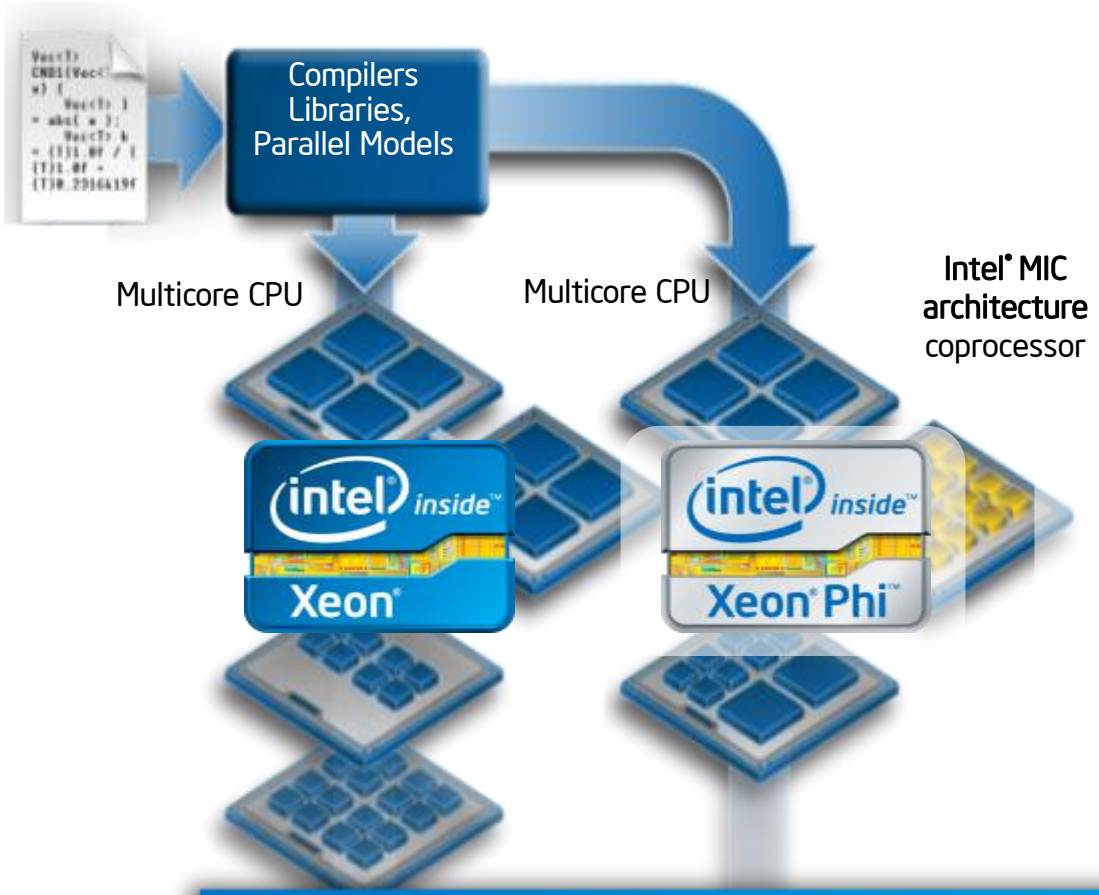

Common with Intel® Xeon® Processors

- Languages
- C, C++, Fortran compilers
- Intel developer tools and libraries
- Coding and optimization techniques
- Ecosystem support

"Unparalleled productivity… most of this software does not run on a GPU" - Robert Harrison, NICS, ORNL

"R. Harrison, "Opportunities and Challenges Posed by Exascale Computing - ORNL's Plans and Perspectives", National Institute of Computational Sciences, Nov 2011"

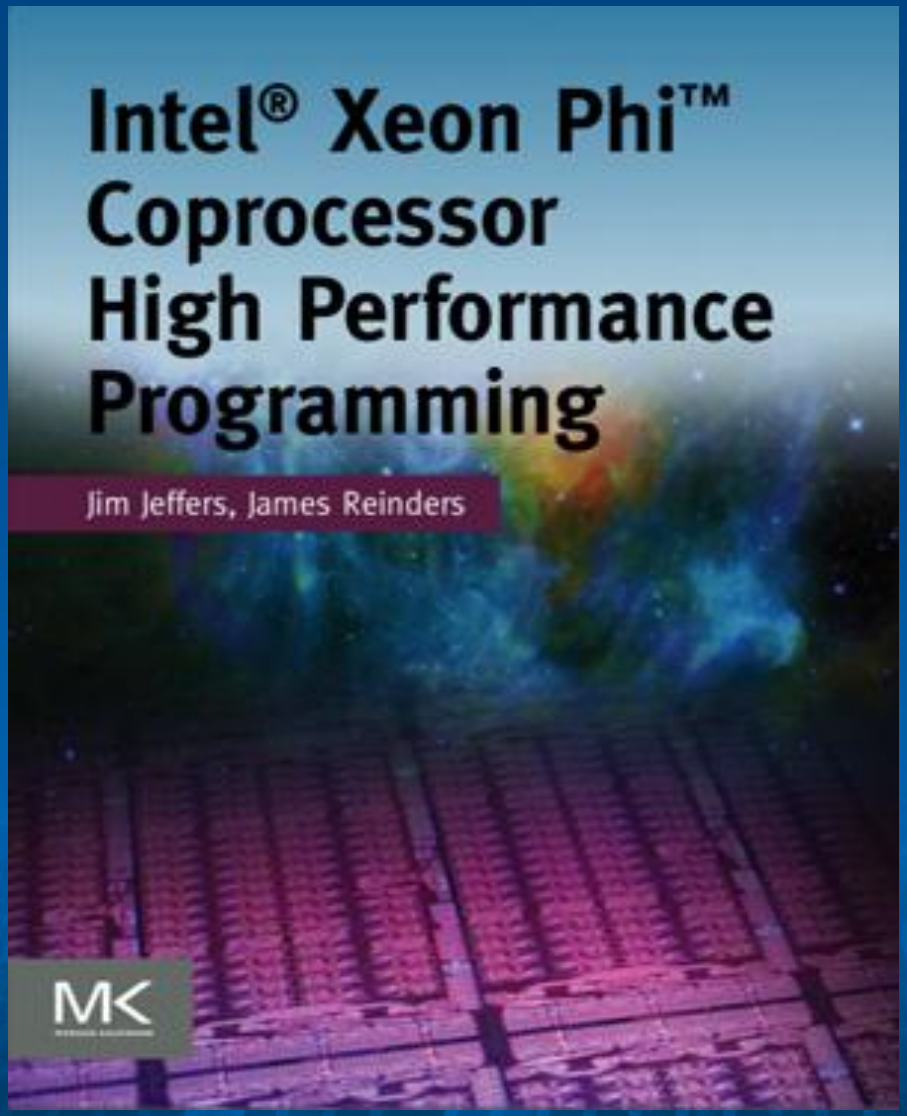

Available in February 2013.

~450 pages completely focused on Intel Xeon Phi coprocessors.

It all comes down to PARALLEL PROGRAMMING !
(applicable to processors and Intel® Xeon Phi™ coprocessor)

(c) 2013, publisher: Morgan Kaufmann (ISBN 978-0-124-10414-3)

Introduction to High Performance Application Development for Multicore and Manycore-Live webinar- 2 Day Series

Abstract: This two day webinar series introduces developers to the world of multicore and manycore computing with Intel® Xeon processors and Intel® Xeon Phi™ coprocessors. Expert technical teams at Intel discuss development tools, programming models, vectorization and execution models that will get development efforts powered up to get the best out of high performance applications and platforms.

When: Day 1 – Feb 26th & Day 2 – Feb 27th

Where: Online

Who: High Performance Application Developers

Agenda for the Days (Must Register for Each Day)

Feb 26th - Live Webinar Day 1 (Pacific Time)	
6:45am – 7:00am	Welcome and Introduction to Developing Applications for Intel® Xeon and Intel® Xeon Phi processors and coprocessors
7:00am – 8:00am	Introduction to Intel® Xeon Phi™ coprocessor hardware and software architecture: native and offload execution basics
8:00am – 9:30am	Compilation for Intel® Xeon Phi™ coprocessor: vectorization, programming models, alignment, pre-fetch, & more
9:30am – 10:00am	Debugging on Intel® Xeon Phi™ coprocessor: using The GNU Project Debugger (GDB)

REGISTER NOW FREE!

Day 1 - https://www1.gotomeeting.com/register/366181513

Feb 27th Live Webinar Day 2 (Pacific Time)	
7:00am – 8:00am	Intel® Math Kernel Library (Intel® MKL) on the Intel® Xeon Phi™ coprocessor
8:00am – 9:00am	Message Passing Interface (MPI) on Intel® Xeon Phi™ coprocessor: special considerations for MPI on Intel Xeon Phi and Intel® Trace Analyzer & Collector
9:00am – 9:50am	Performance analysis and events: Intel® VTune Amplifier introduction, GUI and command line, setup and collection, hot spots, bandwidth, events & more
9:50am – 10:00am	Attendee Q&A, wrap-up

REGISTER NOW FREE!

Day 2 - https://www1.gotomeeting.com/register/241666904

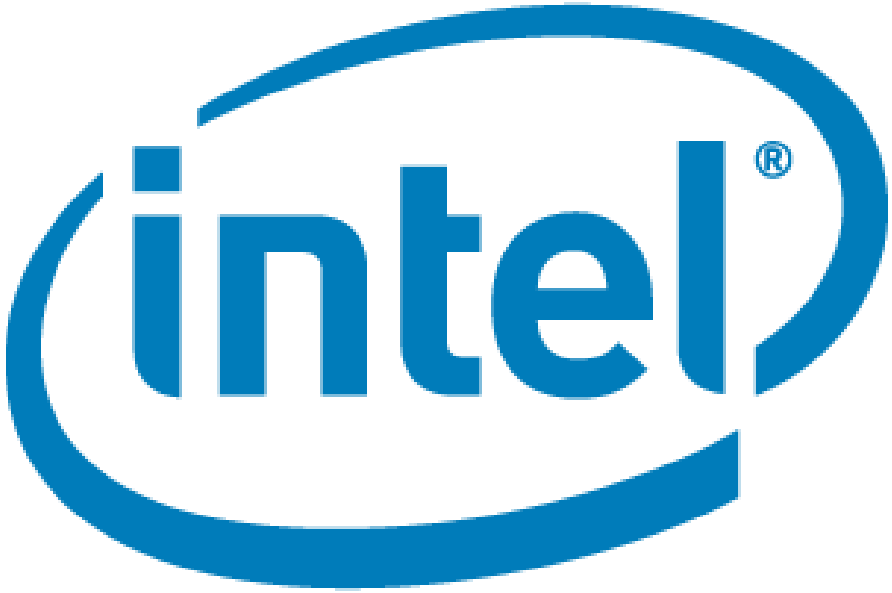
intel


intel