



Preview:

Developer's Guide to Knights Landing Introduction to 2nd Generation Intel[®] Xeon Phi[™] Processors

Colfax International — [@colfaxintl](#)

April 2016 - Rev. 1.1

About This Document

This document represents the materials of a Web-based training “Introduction to 2nd Generation Intel® Xeon Phi™ Processors: Developer’s Guide to Knights Landing” developed and run by Colfax International.

© Colfax International, 2013-2016

Parallel Programming Boot Camp (1-Day) / Workshop (4-Days)



Instructor-led 1-day or 4-days training, at your office or at Colfax facility in Sunnyvale, CA

[Click here to learn more](#)

1-Day Parallel Programming Boot Camp

For software engineers and architects, providing an overview of parallel programming frameworks and optimization guidelines for multi-core CPUs (Intel® Xeon®) and many-core coprocessors (Intel® Xeon Phi™):

- Discussions about three layers of parallelism: SIMD, Threads, Cluster environment
- Tips for quick porting/development of HPC software applications
- Real-life examples of code and optimization techniques
- Hardware solution and corresponding software implementations, APIs, and frameworks

4-Days Parallel Programming Workshop

For the developer who wants to hit the ground running with the modern multi-core CPUs (Intel® Xeon®), many-core coprocessors (Intel® Xeon Phi™) and leading software development tools:

- Hardware installation
- MPSS tools and the Linux environment on the Intel® Xeon Phi™ coprocessor
- Exploring differences in serial vs. parallel programming / processing / hardware usage
- Accelerated clusters
- Optimizations of vector arithmetics, memory traffic, thread parallelism and communication
- Using the Intel® Math Kernel Library

[Register Now!](#)

colfaxresearch.com/knl-webinar/

Disclaimer

While best efforts have been used in preparing this training, Colfax International makes no representations or warranties of any kind and assumes no liabilities of any kind with respect to the accuracy or completeness of the contents and specifically disclaims any implied warranties of merchantability or fitness of use for a particular purpose. The publisher shall not be held liable or responsible to any person or entity with respect to any loss or incidental or consequential damages caused, or alleged to have been caused, directly or indirectly, by the information or programs contained herein. No warranty may be created or extended by sales representatives or written sales materials.

Resources

COLFAX RESEARCH
Log In/Out or Register

CONTRIBUTING TO INNOVATIONS IN COMPUTING

/ READ WATCH LEARN CONNECT JOIN



To search, type and hit enter

Popular

The Hands-On Tutorials (HOT) webinars: details on efficient programming for Intel architecture

The Hands-On Workshop (HOW) Series

Introduction to Intel DAAL, Part 1: Polynomial Regression with Batch Mode Computation

Research and Educational Publications

Introduction to Intel DAAL, Part 1: Polynomial Regression with Batch Mode Computation

Optimization Techniques for the Intel MIC Architecture, Part 3 of 3: False Sharing and Padding

Software Developer's Introduction to the HGST Ultrastar Archive HaaS SMR Drives

Optimization Techniques for the Intel MIC Architecture, Part 2 of 3: Strip-Mining for Vectorization

Optimization Techniques for the Intel MIC Architecture, Part 1 of 3: Multi-Threading and Parallel Reduction

Performance to Power and Performance to Cost Ratios with Intel Xeon Phi Coprocessors (and why xx Acceleration May be Enough)

Parallel Programming Book

Introduction to parallel programming, deep discussion of optimization techniques, exercises.

© 2015, Colfax International. 506 pages.

Featured Video

Additional Reading

See Research material on vectorization in a streaming code

<http://colfaxresearch.com/?p=708>

Consulting

Colfax offers consulting services for enterprises, research help you to:

- Optimize your existing application to take advantage of parallelism, from vectors to cores to clusters and GPUs
- Future-proof your application for upcoming innovations
- Accelerate your application using coprocessor technology
- Investigate the potential system configurations that satisfy your cost, power, performance requirements.
- Take a Hands-On Workshop: a virtual workshop to explore your computing problems, software experience in advance

Episode 2.1 — Purpose of the MIC architecture

Parallel Computing in the Search for New Physics at LHC

COLFAX RESEARCH
Log In/Out or Register

CONTRIBUTING TO INNOVATIONS IN COMPUTING

/ READ WATCH LEARN CONNECT JOIN

Consulting

Colfax offers consulting services for enterprises, research help you to:

- Optimize your existing application to take advantage of parallelism, from vectors to cores to clusters and GPUs
- Future-proof your application for upcoming innovations
- Accelerate your application using coprocessor technology
- Investigate the potential system configurations that satisfy your cost, power, performance requirements.
- Take a Hands-On Workshop: a virtual workshop to explore your computing problems, software experience in advance

Episode 2.1 — Purpose of the MIC architecture

Parallel Computing in the Search for New Physics at LHC

COLFAX RESEARCH
Log In/Out or Register

CONTRIBUTING TO INNOVATIONS IN COMPUTING

/ READ WATCH LEARN CONNECT JOIN

Software Developer's Introduction to the HGST Ultrastar Archive HaaS SMR Drives

Configuration and Benchmarks of Peer-to-Peer Communication over Gigabit Ethernet and InfiniBand in a Cluster with Intel Xeon Phi Coprocessors

Parallel Computing in the Search for New Physics at LHC

COLFAX RESEARCH
Log In/Out or Register

CONTRIBUTING TO INNOVATIONS IN COMPUTING

/ READ WATCH LEARN CONNECT JOIN

Fluid Dynamics with Fortran on Intel Xeon Phi coprocessors

Interview with James Reinders: future of Intel MIC architecture, parallel programming, education

COLFAX RESEARCH
Log In/Out or Register

CONTRIBUTING TO INNOVATIONS IN COMPUTING

/ READ WATCH LEARN CONNECT JOIN

Fluid Dynamics with Fortran on Intel Xeon Phi coprocessors

Interview with James Reinders: future of Intel MIC architecture, parallel programming, education

<http://colfaxresearch.com/>
 (already registered? get \$10 off the book)

HOW Series



THE "HOW" SERIES

DEEP DIVE

WITH CODE MODERNIZATION EXPERTS

STARTS MAY 23

*10x 2-hour sessions | 24-hour 2-weeks remote access to a system | Filling up fast, register now!

Interested? Sign-up at:

colfaxresearch.com/how-series

Developer Access Program (DAP)

Can't wait to get your hands on Knights Landing?




Find out more at dap.xeonphi.com
or contact us at dap@colfax-intl.com.

§2. Intel Architecture: Today and Tomorrow

Intel Xeon Phi Processors (2nd Gen)


2nd Generation of Intel Many Integrated Core (MIC) Architecture.
Specialized platform for demanding computing applications.

- Bootable host processor or coprocessor
- 3+ TFLOP/s DP
- 6+ TFLOP/s SP
- Up to 16 GiB MCDRAM
- MCDRAM bandwidth $\approx 5x$ DDR4
- Binary compatible with Intel Xeon
- **Public disclosures**



Standard CPU Form Factor

- **Bootable Host Processor**
 - ▶ No need for “host”. OS runs on KNL processor.
 - ▶ Supports common OS.
 - ▶ No more PCIe bottleneck!
- **Direct access to ≤ 384 GiB DDR4 RAM**
 - ▶ Up to ≈ 90 GB/s DDR4 bandwidth
- **Access to PCIe Bus**



§3. Cores and Threading

Importance of Parallelism

If KNL was a steam engine...



It would have a lot of furnaces.




***Single-threaded code on KNL
is like having 1 fireman service 72 steam engines***

Implementing Multi-threading

You have to use a core to benefit from it!

Threading frameworks:

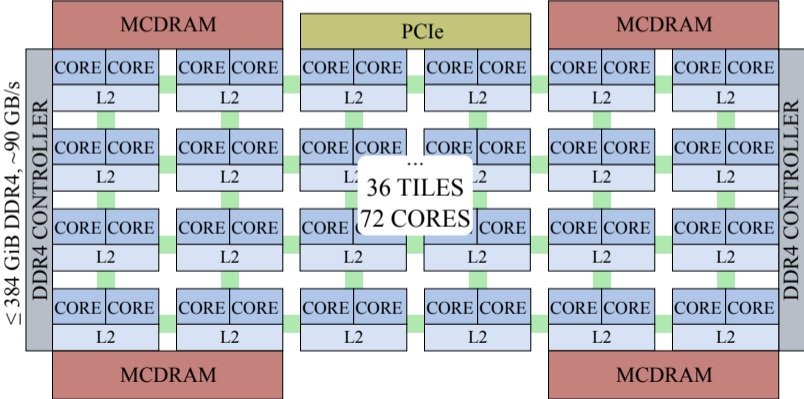
- OpenMP
- TBB
- Cilk Plus
- Pthreads
- Hybrid with MPI
- Any others supported by Xeon Processors



KNL Die Organization: Tiles

- Up to 36 tiles, each with 2 physical cores (72 total).
- Distributed L2 cache across a mesh interconnect.

≤ 16 GiB MCDRAM, ~ 400 GB/s



Affinity and cpubinfo

Neighboring threads often work on nearby/same locations in memory.

→ Use thread pinning to have them share L2 cache!

You can use cpubinfo (part of Intel[®] MPI) to find which cores share cache.

```
user@knl% cpubinfo
// ... cpubinfo output ... //
L2 1  MB (0,1,64,65,128,129,192,193)(2,3,66,67,130,131,194,195)(4,5,68, ...
```

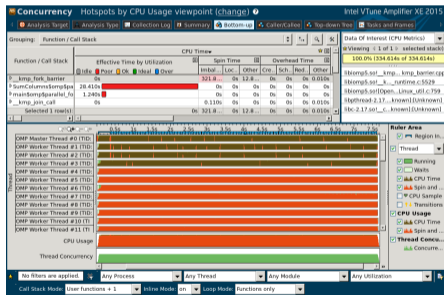
Use KMP_AFFINITY for Intel Compilers or OMP_PROC_BIND for GCC compilers.

```
user@knl% export KMP_AFFINITY=compact
user@knl% export OMP_PROC_BIND=close
```

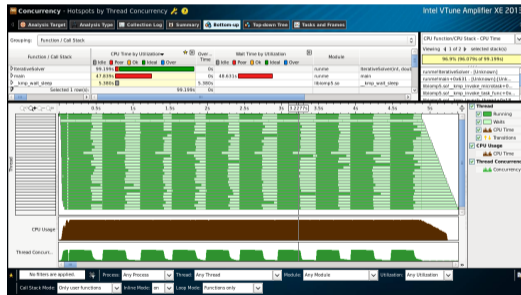
Tuning Threading

Even if you do have multi-threaded code, look out for issues such as...

- Insufficient parallelism



- Load imbalance



Learn more at: colfaxresearch.com/how-series

§4. Vectorization

Scalar Code and Our Fireman



*Not using vector instructions in KNL
is like feeding a steam engine with a spoon.*

AVX-512 Features

Knights Landing: first AVX-512 processor

- AVX-512F (Fundamentals)
 - Extension of most AVX2 instructions to 512-bit vector registers.
- AVX-512CD (Conflict Detection)
 - Efficient conflict detection (application: binning).
- AVX-512ER (Exponential and Reciprocal)
 - Transcendental function (exp, rcp and rsqrt) support.
- AVX-512PF (Prefetch)
 - Prefetch for scatter and gather.

Intel Compiler support for AVX-512

Intel Compiler versions ≥ 15.0 supports AVX-512 instruction set.

```
user@knl% icc -v
icc version 16.0.1 (gcc version 4.8.5 compatibility)
user@knl% icc -help
// ... truncated output ... //
-x<code>
    ...
    MIC-AVX512
    CORE-AVX512
    COMMON-AVX512
```

- `-xMIC-AVX512` : for KNL (supports F, CD, ER, PF)
- `-xCORE-AVX512` : for future Xeon (supports F, CD, DQ, BW, VL)
- `-xCOMMON-AVX512` : common to KNL and Xeon (supports F, CD)

GCC support for AVX-512

GCC \geq 4.9.1 supports AVX-512 instruction set.

```
user@knl% g++ -v
gcc version 4.9.2 (GCC)
user@knl% g++ foo.cc -mavx512f -mavx512er -mavx512cd -mavx512pf
```

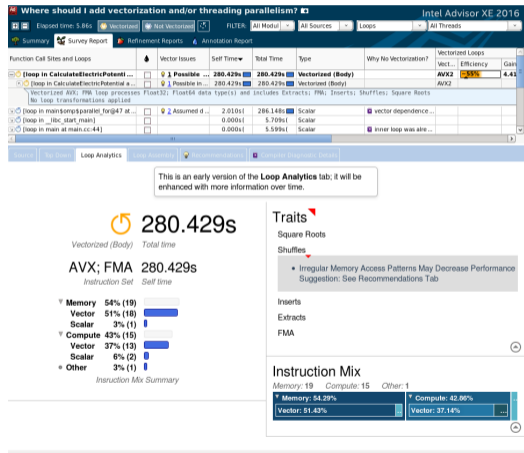
Basic automatic vectorization support: add -O2 or -O3.

```
1 // ... foo.cc ... //
2 for(int i = 0; i < n; i++)
3   B[i] = A[i] + B[i];
```

```
user@knl% g++ -s foo.cc -mavx512f -O3
user@knl% cat foo.s
...
vmovapd -16432(%rbp,%rax), %zmm0
vaddpd -8240(%rbp,%rax), %zmm0, %zmm0
vmovapd %zmm0, -8240(%rbp,%rax)
```

Performance Considerations

Even if your code is vectorized, tuning may unlock more performance.



- Providing enough parallelism.
 - ▶ More consecutive vector operations required to overcome vectorization latency.
- Loop pipelining and unrolling.
 - ▶ Double the pipeline stages to populate.
- Better vectorization patterns.
 - ▶ Avoid long latency operations with unit-stride and unmasked operations.

§5. Memory Architecture

Where's the Fuel?

Cores can't do work if the data is not there.



Source: [wikipedia](https://en.wikipedia.org/wiki/Coal_train)




KNL memory must be used efficiently to feed the cores.

MCDRAM on KNL

MCDRAM Memory Modes

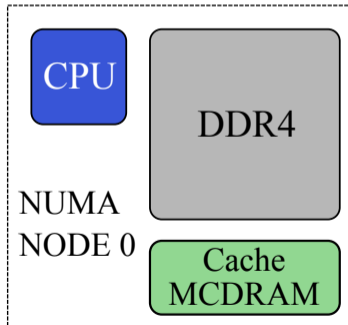
Flat Mode

- MCDRAM treated as a NUMA node
- Users control what goes to MCDRAM




Cache Mode

- MCDRAM treated as a Last Level Cache (LLC)
- MCDRAM is used automatically



Hybrid Mode

- Combination of Flat and Cache
- Ratio can be chosen in the BIOS



Numactl

- Finding information about the NUMA nodes in the system.

```
user@knl% # In Flat mode with All-to-All
user@knl% numactl -H
available: 2 nodes (0-1)
node 0 cpus:  ... all cpus ...
node 0 size: 98207 MB
node 0 free: 94798 MB
node 1 cpus:
node 1 size: 16384 MB
node 1 free: 15991 MB
```

- Binding the application to MCDRAM (Flat/Hybrid)

```
user@knl% gcc myapp.c -o runme -mavx512f -O2
user@knl% numactl --membind 1 ./runme
// ... Application running on MCDRAM ... //
```

Memkind Library and hbwmalloc

Manual allocation to MCDRAM possible with hbwmalloc and Memkind Library.

```
1 #include <hbwmalloc.h>
2 const int n = 1<<10;
3 // Allocation to MCDRAM
4 double* A = (double*) hbw_malloc(sizeof(double)*n);
5 // No replacement for _mm_malloc. Use posix_memalign
6 double* B;
7 int ret = hbw_posix_memalign((void*) B, 64, sizeof(double)*n);
8 .....
9 // Free with hbw_free
10 hbw_free(A); hbw_free(b);
```


Fortran Allocations.

```
1 REAL, ALLOCATABLE :: A(:)
2 !DEC$ ATTRIBUTES FASTMEM :: A
3 ALLOCATE (A(1:1024))
```

Cluster Modes on KNL


Cluster Modes: All-to-All

No affinity between the distributed Tag Directory (TD) and memory.




Cluster Modes: Quadrant/Hemisphere

Tag Directory (TD) and memory reside in the same quadrant.



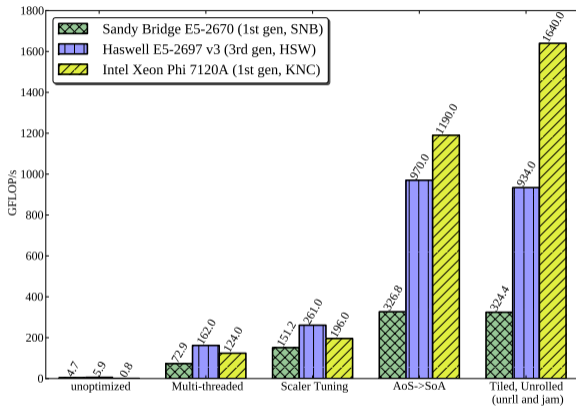
Cluster Modes: SNC-4/SNC-2

Will appear as 4 NUMA nodes (similar to quad-socket system).



§6. Importance of Code Optimization

N-body Simulation on Intel Architecture



Get Ready

The best way to prepare your code for KNL is to optimize for KNC

[Learn More](#)

HOW Series



THE "HOW" SERIES

DEEP DIVE

WITH CODE MODERNIZATION EXPERTS

STARTS MAY 23

*10x 2-hour sessions | 24-hour 2-weeks remote access to a system | Filling up fast, register now!

Interested? Sign-up at:

colfaxresearch.com/how-series


Textbook

ISBN: 978-0-9885234-0-1 (508 pages, Electronic or Print)

Parallel Programming and Optimization with Intel® Xeon Phi™ Coproprocessors

Handbook on the Development and
Optimization of Parallel Applications
for Intel® Xeon® Processors
and Intel® Xeon Phi™ Coprocessors

© Colfax International, 2015



<http://xeonphi.com/book>

§7. Closing Words

COLFAX RESEARCH

CONTRIBUTING TO INNOVATIONS IN COMPUTING

[Log In/Out](#) or [Register](#)

/ READ
/ WATCH
/ LEARN
/ CONNECT
/ JOIN



Introduction to Intel DAAL, Part 1: Polynomial Regression with Batch Mode Computation

To search, type and hit enter

Popular

The Hands-On Tutorials (HOT) webinars: details on efficient programming for Intel architecture

The Hands-On Workshop (HOW) Series

Introduction to Intel DAAL, Part 1: Polynomial Regression with Batch Mode Computation

Parallel Programming Book

Introduction to parallel programming, deep discussion of optimization techniques, exercises.
© 2015, Colfax International. 506 pages.

Featured Video

See Research material on vectorization in a learning video

<http://colfaxresearch.com/?p=708>

Research and Educational Publications

Introduction to Intel DAAL, Part 1: Polynomial Regression with Batch Mode Computation

Optimization Techniques for the Intel MIC Architecture, Part 3 of 3: False Sharing and Padding

Software Developer's Introduction to the HGST Ultrastar Archive HaaS SMR Drives


Optimization Techniques for the Intel MIC Architecture, Part 2 of 3: Strip-Mining for Vectorization



Optimization Techniques for the Intel MIC Architecture, Part 1 of 3: Multi-Threading and Parallel Reduction

Performance to Power and Performance to Cost Ratios with Intel Xeon Phi Coprocessors (and why xx Acceleration May be Enough)

/ Events
/ Demonstrations
/ Careers

Consulting


Share

Colfax offers consulting services for enterprises, research help you to:

- Optimize your existing application to take advantage of parallelism, from vectors to cores to clusters and GPUs
- Future-proof your application for upcoming innovations
- Accelerate your application using coprocessor technology
- Investigate the potential system configurations that satisfy your cost, power, and performance requirements.
- Help you design the architecture of a virtual environment to enhance your computing performance

Episode 2.1 — Purpose of the MIC architecture

Download this video on YouTube

[http://www.youtube.com/watch?v=...](#)

or this video episode 2.1 we will introduce Intel Xeon Phi coprocessors based on the Intel Many Integrated Core, or MIC, architecture and will take a look at the architecture of the Intel Xeon Phi coprocessor.

Intel Xeon Phi Architecture

Download this video on YouTube

Software Developer's Introduction to the HGST Ultrastar Archive HaaS SMR Drives

Download this video on YouTube



In this paper we will discuss the new HGST Ultrastar Archive HaaS SMR drives, a different feature set which offers storage capacities of 10 TB and beyond. These drives have high density and large capacities. These drives are well suited for large "data archive" applications, such as data archive applications, the data backup and recovery solutions.

For more information on storage and data management, visit our website and download our white paper on storage and data management solutions.

Configuration and Benchmarks of Peer-to-Peer Communication over Gigabit Ethernet and InfiniBand in a Cluster with Intel Xeon Phi Coprocessors

Download this video on YouTube

Intel Xeon Phi coprocessors are designed for high performance computing applications. They are designed to be used in clusters with other Intel Xeon Phi coprocessors. This paper describes the configuration and benchmarks of a cluster of Intel Xeon Phi coprocessors connected via Gigabit Ethernet and InfiniBand. The paper also describes the performance of the cluster in terms of throughput and latency.


Parallel Computing in the Search for New Physics at LHC

Download this video on YouTube

Intel Xeon Phi coprocessors are designed for high performance computing applications. They are designed to be used in clusters with other Intel Xeon Phi coprocessors. This paper describes the configuration and benchmarks of a cluster of Intel Xeon Phi coprocessors connected via Gigabit Ethernet and InfiniBand. The paper also describes the performance of the cluster in terms of throughput and latency.

Fluid Dynamics with Fortran on Intel Xeon Phi coprocessors

Download this video on YouTube




In this paper we will discuss the Intel Xeon Phi coprocessors and the application of the Fortran language on the Intel Xeon Phi coprocessors. The paper also describes the performance of the Intel Xeon Phi coprocessors in terms of throughput and latency.

Interview with James Reinders: future of Intel MIC architecture, parallel programming, education

Download this video on YouTube

In this video we will see a conversation with James Reinders, the Director and Chief Architect of Intel Xeon Phi. We discuss the future of the parallel programming and how the performance can be improved through parallel programming and how the performance can be improved through parallel programming and how the performance can be improved through parallel programming.



Parallel Computing in the Search for New Physics at LHC

Download this video on YouTube

Intel Xeon Phi coprocessors are designed for high performance computing applications. They are designed to be used in clusters with other Intel Xeon Phi coprocessors. This paper describes the configuration and benchmarks of a cluster of Intel Xeon Phi coprocessors connected via Gigabit Ethernet and InfiniBand. The paper also describes the performance of the cluster in terms of throughput and latency.

<http://colfaxresearch.com/>
 (already registered? get \$10 off the book)

Developer Access Program (DAP)

Can't wait to get your hands on Knights Landing?



Find out more at dap.xeonphi.com
or contact us at dap@colfax-intl.com.

Bottom Line

KNL is a **highly-parallel** successor of KNC, with improvements in both **performance** and **ease-of-use**.



KEEP CALM
AND
GO PARALLEL

Thank you for Tuning In!

