



PROGRAMMING AND OPTIMIZATION FOR INTEL[®] ARCHITECTURE

Hands-On Workshop (HOW) Series "Deep Dive"
Session 1

Colfax International — colfaxresearch.com

February 2017

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.

- ▶ **Module I. Programming**
 - 01. Intel Architecture and Modern Code – Feb 13
 - 02. Xeon Phi, Coprocessors, Omni-Path – Feb 14
- ▶ **Module II. Expresssing Parallelism**
 - 03. Expressing Parallelism with Vectors – Feb 15
 - 04. Multi-threading with OpenMP – Feb 16
 - 06. Distributed Computing, MPI – Feb 17
- ▶ **Module III. Optimization**
 - 06. Optimization Overview: N-body – Feb 20
 - 07. Scalar tuning, Vectorization – Feb 21
 - 08. Common Multi-threading Problems – Feb 22
 - 09. Multi-threading, Memory Aspect – Feb 23
 - 10. Access to Caches and Memory – Feb 24

S	M	T	W	H	F	S
			1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28				
■ — Webinar+remote access						

Course page:

colfaxresearch.com/how-17-02

- ▶ Slides
- ▶ Code
- ▶ Video
- ▶ Chat

More workshops:

colfaxresearch.com/training



GET YOUR QUESTIONS ANSWERED: CHAT



colfaxresearch.com/how-17-02

GET YOUR QUESTIONS ANSWERED: FORUMS



Forum

Colfax Cluster

Discussion of Colfax Cluster usage policies, troubleshooting.

Modern Code

Discuss with Colfax Research and colleagues any topics related to computational science, parallel programming, performance optimization and code modernization.

Developer Training

Questions about any of the Colfax trainings? Usage of training servers, experience with specific exercises, inquiries on what's inside, suggestions for future trainings - post them here.

Colfax News

Subscribe to this forum if you wish to receive updates on new papers, training events, and other news we may have. Updates here are frequent and

colfaxresearch.com/discussion

- ▶ All registrants receive an invitation from `cluster@colfaxresearch.com`
- ▶ Queue-based access to Intel Xeon E5, Intel Xeon Phi (KNC and KNL)
- ▶ Can access the cluster the entire 2 weeks of the workshop



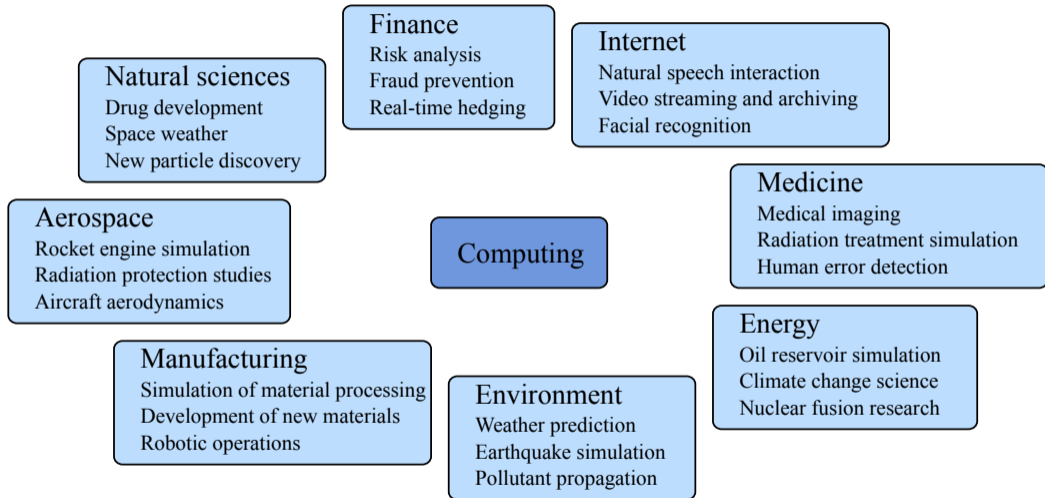


§2. MODERN CHALLENGES



AREAS OF APPLICATION

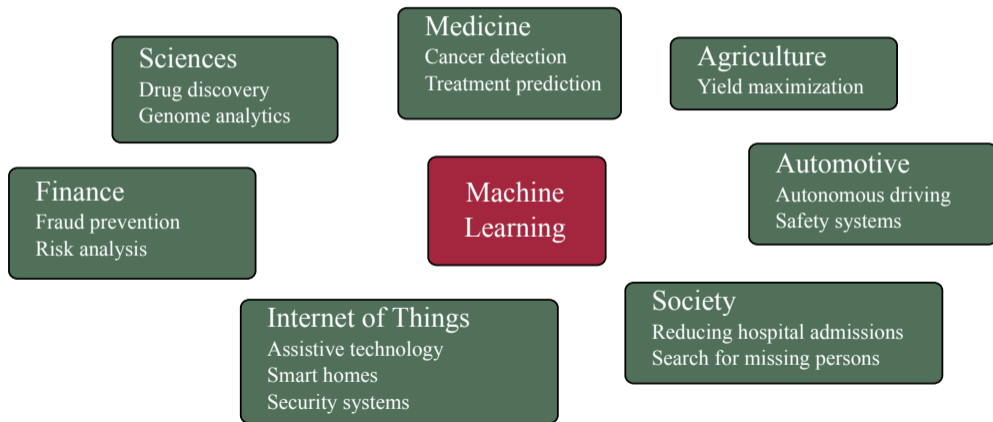
TRADITIONAL COMPUTING APPLICATIONS



ARTIFICIAL INTELLIGENCE AND MACHINE LEARNING


AI = reasoning, perception, action and adaptation by machines

ML = data-driven AI, intelligence without explicit programming







NEED FOR SPEED



WHAT IS YOUR PERFORMANCE METRIC?

- ▶ Performance per System
- ▶ Performance per Watt
- ▶ Performance/Cost Ratio



See [this paper](#) for details




§3. MODERN ARCHITECTURE



HOW PROCESSORS GET FASTER


40 YEARS OF MICROPROCESSOR DATA



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
 New plot and data collected for 2010-2015 by K. Rupp

Source: karlrupp.net

POWER WALL



Liquid nitrogen for CPU
overclocking

youtu.be/WZr0W_g0dqk



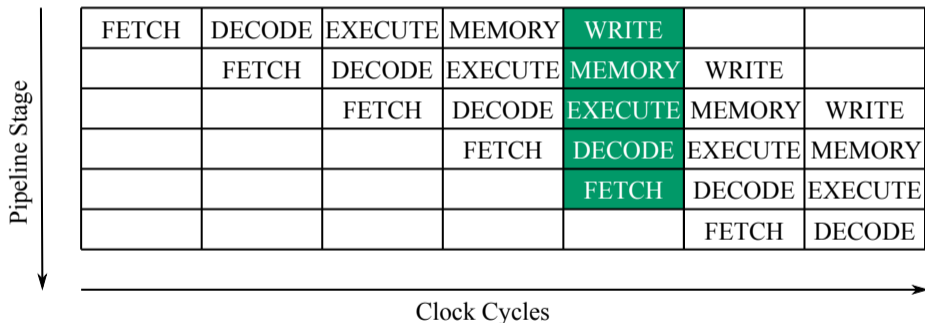
Microsoft's project Natick:
immersed datacenter

news.microsoft.com/natick

Cooling solutions for high clock speeds are not practical or expensive

INSTRUCTION-LEVEL PARALLELISM (ILP) WALL: PIPELINING

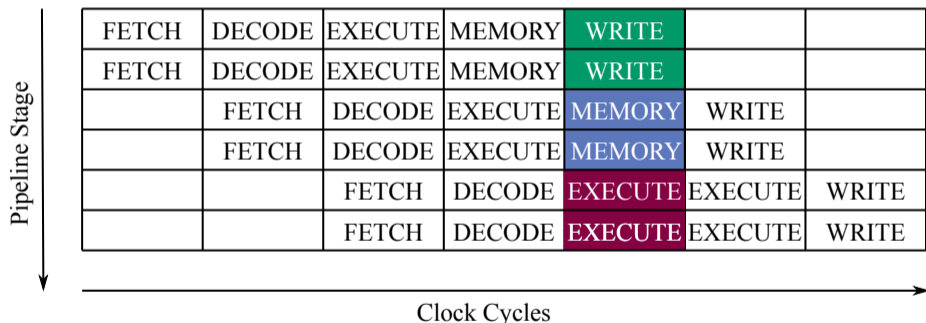
Pipelining – replication of hardware to run different stages of different instruction streams at the same time



Only so many pipeline stages, possible conflicts

INSTRUCTION-LEVEL PARALLELISM (ILP) WALL: SUPERSCALAR EXECUTION


Superscalar Execution – hardware checks for independence of operations, pipelines multiple instructions in a cycle



Automatic search for independent instructions requires extra resources

MEMORY WALL: OUT-OF-ORDER EXECUTION

Out-of-order Execution – hardware re-orders instructions in a stream to minimize latencies

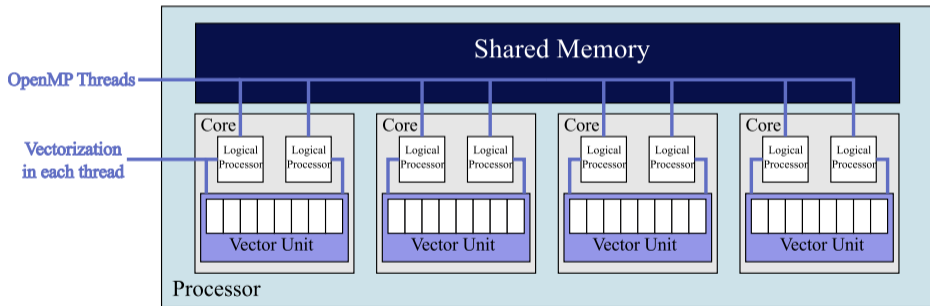


CPU performance grows faster than RAM bandwidth, OOE can't fill gap

PARALLELISM

CORES – multiple instructions on multiple data elements (MIMD)

VECTORS – single instruction on multiple data elements (SIMD)



Unbounded growth opportunity, but **not automatic**

PARALLELISM IS THE PATH FORWARD

- ▶ Clock speed has hit the power wall
- ▶ Automatic parallelism has hit the ILP wall
- ▶ Out-of-order execution cannot overcome the memory wall

The Show Must Go On

Hardware keeps evolving through parallelism.
Software must catch up!



INTEL ARCHITECTURE

INTEL COMPUTING PLATFORMS

General-Purpose Processors

Intel® Xeon®
Intel® Core™
Intel® Atom™, ...



Specialized Processors

Intel® Xeon Phi™
processors
and coprocessors



Computing Accelerators

Intel® VCA (x86)
Intel® Nervana™ Platform
Intel® DLIA™ (FPGAs)



Network Interconnects


Intel® Omni-Path™
Architecture



INTEL XEON CPU: PURPOSE AND SPECIFICATIONS

General-purpose platform for demanding computing applications.

- ▶ Up to ~ 1 TFLOP/s in DP*
- ▶ Up to ~ 2 TFLOP/s in SP*
- ▶ Up to 3072 GiB DDR4 RAM*
- ▶ ~ 154 GB/s bandwidth*
- ▶ Hardware-rich: forgiving of sub-optimal code





* 2-way Intel Xeon processor, Skylake architecture, top-of-the-line (e.g., E5-2699 V4)

INTEL XEON PHI PROCESSORS (1ST GEN)

Specialized platform for demanding computing applications.



- ▶ PCIe add-in card
- ▶ ~ 1.2 TFLOP/s in DP
- ▶ ~ 2.4 TFLOP/s in SP
- ▶ Up to 16 GiB GDDR5 RAM
- ▶ ~ 176 GB/s bandwidth
- ▶ Heterogeneous clustering
- ▶ Runs special Linux distribution



INTEL XEON PHI PROCESSORS (2ND GEN)


Specialized platform for demanding computing applications.

- ▶ Socket version or coprocessor
- ▶ 64-72 cores × 4 HT at 1.3-1.5 GHz
- ▶ 3+ TFLOP/s in DP (FMA)
- ▶ 6+ TFLOP/s in SP (FMA)
- ▶ ≤ 384 GiB DDR4 (> 90 GB/s)
- ▶ 16 GiB HBM (MCDRAM, > 400 GB/s)
- ▶ Binary-compatible with Xeon
- ▶ Common OS
(RHEL/CentOS/SUSE/Windows)





- C/C++/Fortran
- Linux/Windows
- ≤ 3 TiB DDR4
- ≤ 44 cores (2-way)
- ≈ 3 GHz
- 2 HT/core
- 256-bit AVX




- C/C++/Fortran
- Special Linux
- ≤ 16 GiB GDDR5
- 57-61 cores
- ≈ 1.2 GHz
- 4 HW THR/core
- 512-bit IMCI



- C/C++/Fortran
- Linux
- MCDRAM+DDR4
- 64-72 cores
- 1.3-1.5 GHz
- 4 HT/core
- 512-bit AVX-512

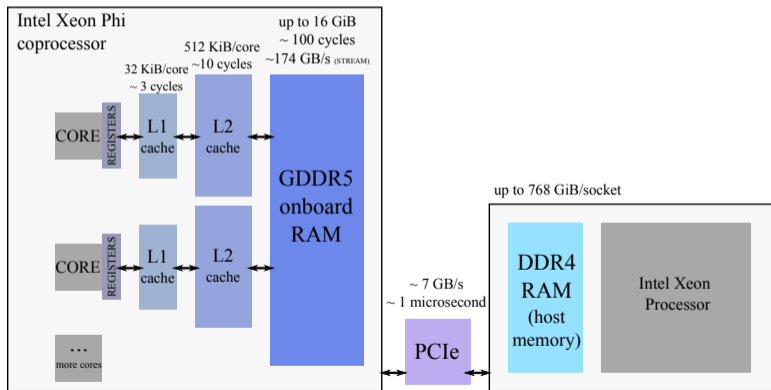

INTEL XEON CPU: MEMORY ORGANIZATION

- ▶ Hierarchical cache structure
- ▶ Two-way processors have NUMA architecture





KNC MEMORY ORGANIZATION

- ▶ Direct access to ≤ 16 GiB of cached GDDR5 memory on board
- ▶ No access to system DDR4, connected to host via PCIe



KNL MEMORY ORGANIZATION (BOOTABLE)

- ▶ On-package high-bandwidth memory (HBM) – MCDRAM
- ▶ Optimized for arithmetic performance and bandwidth (not latency)

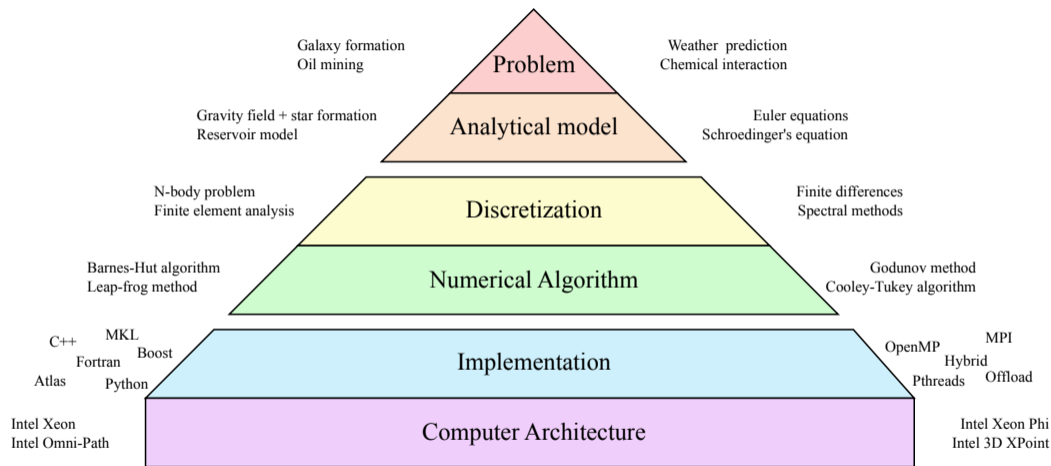





§4. MODERN CODE



OPTIMIZATION AND FUTURE-PROOFING

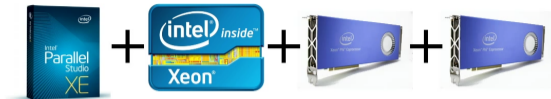






MOTIVATING EXAMPLES

ASTROPHYSICAL CODE HEATCODE: AN OFFLOAD STORY

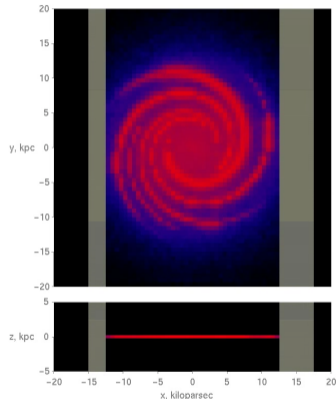


Porting to Intel® Xeon Phi™ coprocessors

- We ported Frankie code using explicit offload model
- Same code & optimization methods for Xeon Phi™
- Simultaneous calculations on CPU and coprocessors with automatic load balancing was easy to implement
- With two Intel® Xeon Phi™ coprocessors, performance for high-res calculations is 3.2x better than with two Intel® Xeon® E5 processors alone.
- **RESULT:** estimated target project calculation time is now 2 weeks (down from 6+ years)

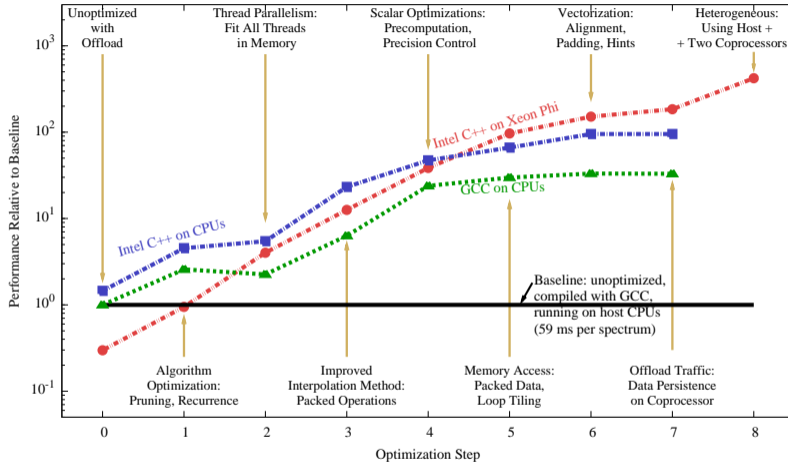
Goal achieved!

Transient Emission of Cosmic Dust Grains
in the Milky Way Galaxy,
Simulation with Frankie Code



<http://xeonphi.com/papers/heatcode>

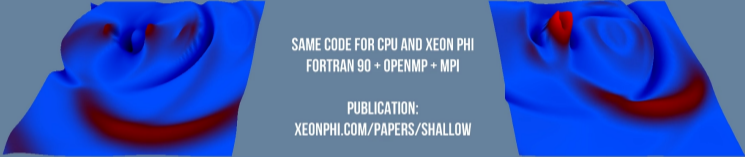
ASTROPHYSICAL CODE HEATCODE: AN OFFLOAD STORY



<http://xeonphi.com/papers/heatcode>

FLUID DYNAMICS WITH FORTRAN ON INTEL® XEON PHI™ COPROCESSORS

SHALLOW WATER EQUATION SOLVER



SAME CODE FOR CPU AND XEON PHI
FORTRAN 90 + OPENMP + MPI


PUBLICATION:
XEONPHI.COM/PAPERS/SHALLOW

PERFORMANCE ON CPU: 19.5 GFLOP/S


PERFORMANCE WITH COPROCESSORS: 52.5 GFLOP/S

SIMULATION SIZE: 9600X9600

ACCELERATION: 2.7X



INTEL XEON E5-2697 V3 PROCESSOR



INTEL XEON E5-2697 V3 PROCESSOR +
TWO INTEL XEON PHI 7120A COPROCESSORS

COLFAX

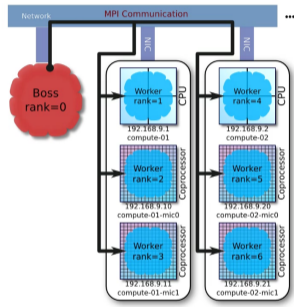
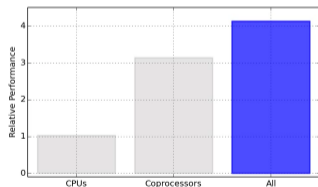
SERVERS WORKSTATIONS TRAINING CONSULTING RESEARCH

WWW.COLFAX-INTL.COM

<http://xeonphi.com/papers/shallow>

ASIAN OPTION PRICING: HETEROGENEOUS CLUSTERING

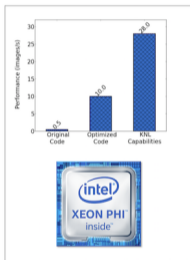
Heterogeneous Clustering with Homogeneous Code:
Asian Option Pricing



<http://xeonphi.com/papers/heterogeneous>

MACHINE LEARNING: OPTIMIZED MIDDLEWARE

INTEL® XEON PHI™ PROCESSORS — MACHINE LEARNING



NEURALTALK2 — OPEN SOURCE IMAGE TAGGING CODE (KARPATY & FEI-FEI, STANFORD)

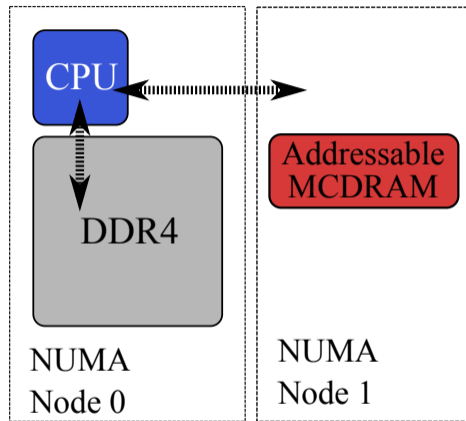
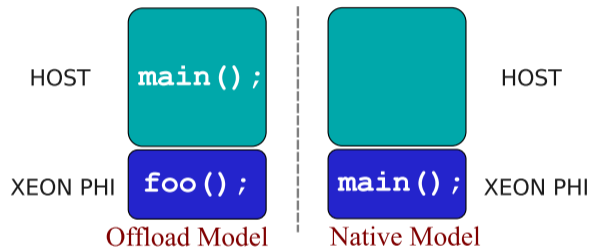


<http://colfaxresearch.com/isc16-neuraltalk>



WHAT YOU ARE GOING TO LEARN

HETEROGENEOUS AND NUMA ARCHITECTURES



Session 2: handling memory organization in Intel Xeon Phi processors

DATA PARALLELISM AND VECTOR INSTRUCTIONS

Vectors – form of SIMD architecture (Single Instruction Multiple Data).

Scalar Instructions

$$\begin{array}{r}
 4 + 1 = 5 \\
 0 + 3 = 3 \\
 -2 + 8 = 6 \\
 9 + -7 = 2
 \end{array}$$

Vector Instructions


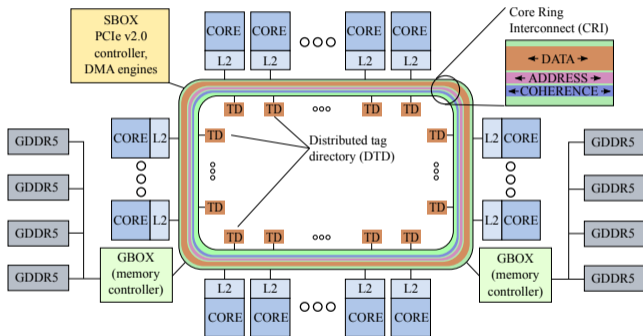
$$\begin{array}{r}
 4 \quad 1 \quad 5 \\
 0 \quad 3 \quad 3 \\
 -2 \quad 8 \quad 6 \\
 9 \quad -7 \quad 2
 \end{array}
 + =$$

↑
Vector Length
↓

Session 3: automatic vectorization with Intel compilers

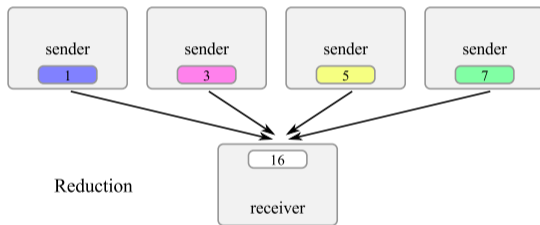
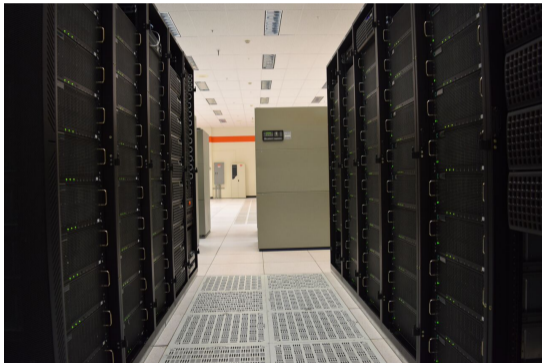
TASK PARALLELISM AND CORES

Cores implement MIMD (Multiple Instruction Multiple Data) arch





Session 4: multi-threading with OpenMP

Clusters form distributed-memory systems with network interconnects



Session 5: Message Passing Interface (MPI)

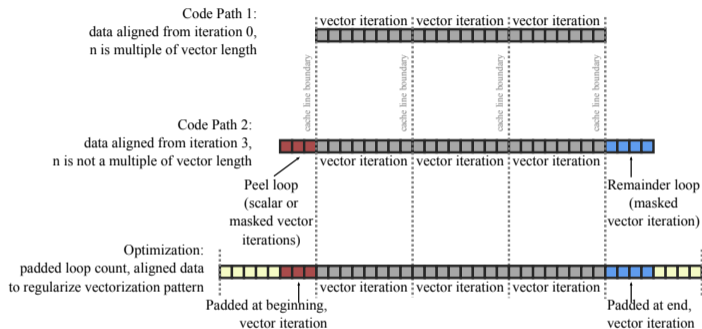
OPTIMIZATION OVERVIEW



Session 6: optimization overview, case study



SCALAR TUNING, OPTIMIZATION OF VECTORIZATION

```
for (i = 0; i < n; i++) A[i] = ...
```






Session 7: precision control, regularizing vectorization patterns

COMMON ISSUES IN MULTI-THREADING




Session 8: minimizing synchronization, avoiding false sharing, strip-mining for parallelism

MULTI-THREADING, MEMORY ASPECT



Session 9: thread affinity, NUMA locality, scheduling



Tiling

1	2	3	4
5	6	7	8
9	10	11	12
13	14	15	16

Cache-Oblivious Recursion

1	3	9	11
2	4	10	12
5	7	13	15
6	8	14	16

Session 10: loop transformations for locality, bandwidth secrets



§5. HANDS-ON DEMONSTRATION

ACCESS THE COLFAX CLUSTER

[Home](#)[Learn](#)[Connect](#)[Program](#)[Compute](#)[Log Out](#)

Welcome to Colfax Cluster!

Learn




what to expect on
the Colfax Cluster

Connect



from your home
computer to the
cluster

Program




using modern code
practices

Compute



with cluster job
management tools

JOB MANAGEMENT IN THE CLUSTER



```
[u111@c005 ~]% git clone https://github.com/ColfaxResearch/HOW-Series-Labs.git
Cloning into 'HOW-Series-Labs'...
[u111@c005 ~]% ls HOW-Series-Labs/*/
HOW-Series-Labs/2/:
2.01-native-basic      2.03-offload-basic      2.05-shared-virtual-memory-basic
2.02-native-MPI       2.04-offload-asynchronous  2.06-shared-virtual-memory-complex-

HOW-Series-Labs/3/:
3.01-vectorization    3.03-OpenMP-reduction    3.05-Cilk-Plus-basics      3.07-Cilk-Plu
3.02-OpenMP-basics    3.04-OpenMP-tasks        3.06-Cilk-Plus-reducers    3.08-MPI-basi

HOW-Series-Labs/4/:
4.01-overview-nbody          4.07-threading-affinity
4.02-vectorization-data-structures-coulomb  4.08-memory-tiling-matrix_x_vector
4.03-vectorization-tuning-lu-decomposition  4.09-memory-loop-fusion-statistics
4.04-threading-misc-histogram                4.10-offload-double-buffering-dgem
...
```

REVIEW AND WHAT'S NEXT

- ▶ Computers are getting faster through parallelism and specialization
- ▶ Intel Xeon E5 product family – general-purpose parallel processors
- ▶ Intel Xeon Phi product family – specialized parallel processors
- ▶ Coprocessor – either offload device or an additional compute node

Next session: details of Intel Xeon Phi processor and coprocessor programming.