



# MODERN CODE PRACTICES AND INTEL<sup>®</sup> ARCHITECTURE


Part 1 of 3

*Colfax International* — [colfaxresearch.com](http://colfaxresearch.com)

September 2016

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.

- ▶ **Part 1:** Multi-threading strategies
  - Minimizing synchronization
  - Avoiding false sharing
  - Exposing parallelism
- ▶ **Part 2:** Vectorization tuning
  - Roles of compiler and developer
  - Tuning with directives
  - Data container optimization
  - Language extensions
- ▶ **Part 3:** Memory traffic control
  - Maximizing cache utilization
  - Optimizing memory bandwidth
  - Intel Xeon Phi processors: high-bandwidth memory





## **§2. INTEL ARCHITECTURE**



# COMPUTING PLATFORMS


## Intel Xeon Processor



Current: Broadwell  
Upcoming: Skylake

Multi-Core Architecture

## Intel Xeon Phi Coprocessor, 1st generation



Knights Corner (KNC)

## Intel Xeon Phi Processor, 2nd generation\*



\* socket and coprocessor versions

Knights Landing (KNL)


Intel Many Integrated Core (MIC) Architecture

# BOOTABLE INTEL XEON PHI PROCESSORS


- ▶ Bootable Host Processor
- ▶ RHEL/CentOS/SUSE/Win
- ▶ 64 cores × 4 HT, 1.3 GHz
- ▶ ≤ 384 GiB DDR4, > 90 GB/s
- ▶ 16 GiB HBM, > 400 GB/s
- ▶ PCIe bus for networking

[dap.xeonphi.com](http://dap.xeonphi.com)

Servers:





Workstations:



# FUTURE FORM-FACTORS

## KNLF: KNL with Fabric

- ▶ Fabric integrated on CPU
  - Intel OmniPath Architecture
- ▶ Socket mount processor



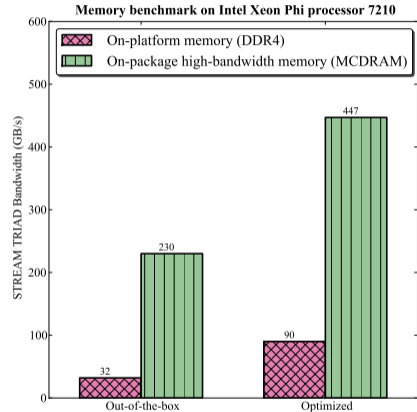

\*KNC image

## KNL Coprocessor

- ▶ PCIe add-in card
  - Requires host
- ▶ Multiple KNLs in a system




# IT TAKES GOOD SOFTWARE TO UNLOCK THE PERFORMANCE!



Details on N-body simulation in Chapter 23 of [this book](#)




# PERFORMANCE OPTIMIZATION




# KNL DIE ORGANIZATION

- ▶ Mesh interconnect relaxes data locality requirement [somewhat]
- ▶ All-to-all, quadrant or sub-numa domain communication in mesh



≤ 16 GiB on-package MCDRAM, ~ 400 GB/s





## **§3. MULTI-THREADING: WHAT CAN GO WRONG?**




# THREADS AND OPENMP

# "HELLO WORLD" OPENMP PROGRAM

```
1  #include <omp.h>
2  #include <stdio.h>
3
4  int main(){
5      // This code is executed by only 1 thread
6      const int nt=omp_get_max_threads();
7      printf("OpenMP with %d threads\n", nt);
8
9      #pragma omp parallel
10     {
11         // This code is executed in parallel
12         // by multiple threads
13         printf("Hello World from thread %d\n",
14                omp_get_thread_num());
15     }
16 }
```

- ▶ OpenMP = “Open Multi-Processing” = computing-oriented framework for shared-memory programming
- ▶ Threads – streams of instructions that share memory address space
- ▶ Distribute threads across CPU cores for parallel speedup

# CO-EXISTENCE WITH VECTORS



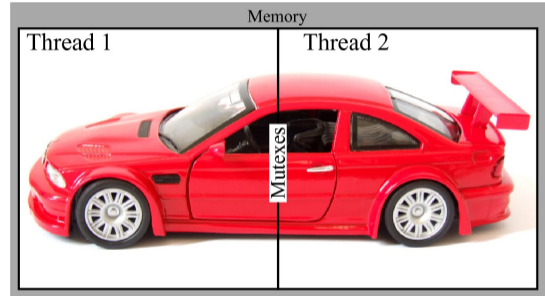
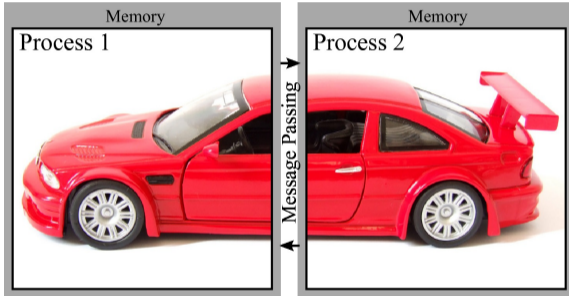
**Utilize cores:** run multiple threads/processes (MIMD)

**Utilize vectors:** each thread (process) issues vector instructions (SIMD)



# THREADS VERSUS PROCESSES


Option 1: Partitioning data set between threads/processes



Examples: computational fluid dynamics (CFD), image processing.

# THREADS VERSUS PROCESSES

## Option 2: Sharing data set between threads/processes



Examples: particle transport simulation, machine learning (inference).



## **ISSUE 1: TOO MUCH SYNCHRONIZATION**

# RACE CONDITIONS AND UNPREDICTABLE PROGRAM BEHAVIOR

```


1  #include <omp.h>
2  #include <stdio>
3  int main() {
4      const int n = 1000;
5      int total = 0;
6      #pragma omp parallel for
7      for (int i = 0; i < n; i++) {
8          total = total + i; // Race condition
9      }
10     printf("total=%d (must be %d)\n", total,
11            ((n-1)*n)/2);
12 }

```

```

vega@lyra% icpc -o app omp-race.cc -qopenmp
vega@lyra% ./app
total=208112 (must be 499500)

```



**Race Condition!**

- Occurs when 2 or more threads access the same memory address, and at least one of these accesses is for writing

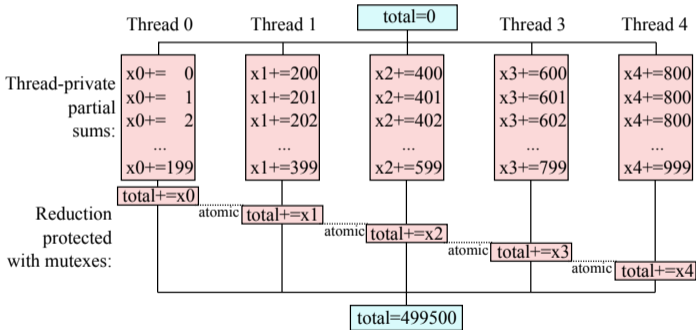
# AVOIDING RACES WITH THREAD-PRIVATE STORAGE

Correct and efficient code:

```

1  int total = 0;
2  #pragma omp parallel
3  {
4      int total_thr = 0;
5      #pragma omp for
6      for (int i=0; i<n; i++)
7          total_thr += i;
8
9      #pragma omp atomic
10     total += total_thr;
11
12 }

```



# EXAMPLE: BINNING PROBLEM


Computing a histogram ( $m \ll n$ ):

```

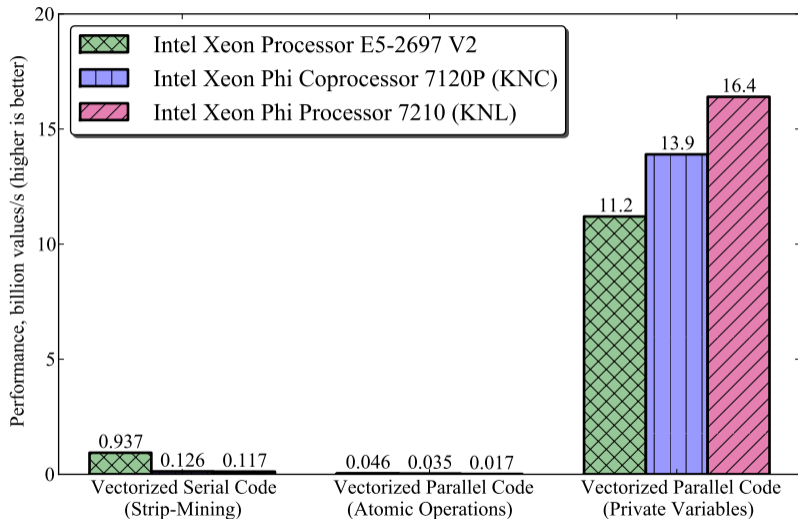
1 void Histogram(
2     // Ages, values from 0.0f to 100.0f:
3     const float* age,
4     // Size of array age, n=100000000:
5     const int n,
6     // Output: counts in groups:
7     int* const hist,
8     // Size of array hist, m=5:
9     const int m,
10    const float group_width) {
11    for (int i = 0; i < n; i++) {
12        const int j = int(age[i]/group_width);
13        hist[j]++;
14    }
15 }

```

- ▶ Vector dependence in `hist[j]++`
- ▶ Strip-mine or use conflict detection



# USING REDUCTION INSTEAD OF SYNCHRONIZATION






## **ISSUE 2: FALSE SHARING**

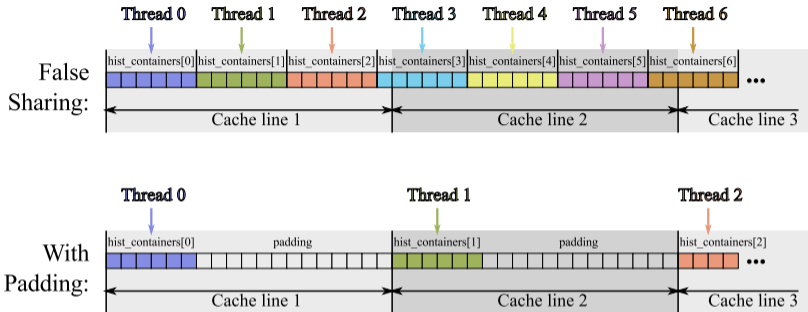


# FALSE SHARING. DATA PADDING AND PRIVATE VARIABLES



- ▶ Occurs when 2 or more threads access the same cache line, and at least one of the accesses is for writing
- ▶ Caused by *coherent caches*
- ▶ Cache line is 64-byte wide (in modern Intel architectures)

# PADDING TO AVOID FALSE SHARING




```

1 // Padding to avoid sharing a cache line between threads
2 const int paddingBytes = 64;
3 const int paddingElements = paddingBytes / sizeof(int);
4 const int mPadded = m + (paddingElements - m % paddingElements);
5 int hist_containers[nThreads][mPadded]; // New container

```

# PADDING TO AVOID FALSE SHARING







## **ISSUE 3: INSUFFICIENT PARALLELISM**

# INSUFFICIENT PARALLELISM

## Analysis in Intel VTune Amplifier XE



- ▶ Occurs when there are not enough iterations or parallel work-items exposed to the parallel loop in OpenMP.

## EXAMPLE: DEALING WITH INSUFFICIENT PARALLELISM

$$S_i = \sum_{j=0}^n M_{ij}, i = 0 \dots m. \quad (1)$$


- ▶  $m=4$  is small, smaller than the number of threads in the system
- ▶  $n \approx 10^8$  is large enough so that matrix does not fit into cache

```
1 void sum_unoptimized(const int m, const int n, long* M, long* s){
2   #pragma omp parallel for
3     for (int i=0; i<m; i++) { // m=4
4       long total=0;
5       #pragma vector aligned
6         for (int j=0; j<n; j++) // n=100000000
7           total+=M[i*n+j];
8       s[i]=total; }}
```

# EXPOSING PARALLELISM: STRIP-MINING AND LOOP COLLAPSE


```
1 void sum_stripmine(const int m, const int n, long* M, long* s){
2     const int STRIP=1024;
3     assert(n%STRIP==0);
4     s[0:m]=0;
5     #pragma omp parallel
6     {
7         long total[m];    total[0:m]=0;
8     #pragma omp for collapse(2) schedule(guided)
9         for (int i=0; i<m; i++)
10            for (int jj=0; jj<n; jj+=STRIP)
11    #pragma vector aligned
12                for (int j=jj; j<jj+STRIP; j++)
13                    total[i]+=M[i*n+j];
14        for (int i=0; i<m; i++)                // Reduction
15    #pragma omp atomic
16            s[i]+=total[i];
17    } }
```

# EXPOSING PARALLELISM: STRIP-MINING AND LOOP COLLAPSE






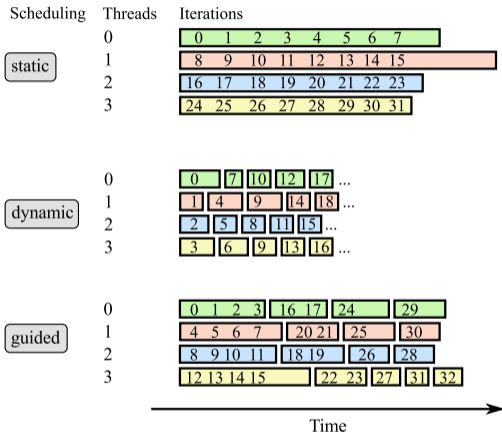
# DEALING WITH INSUFFICIENT PARALLELISM






## **OTHER OPENMP CONTROLS**

# LOOP SCHEDULING MODES IN OPENMP



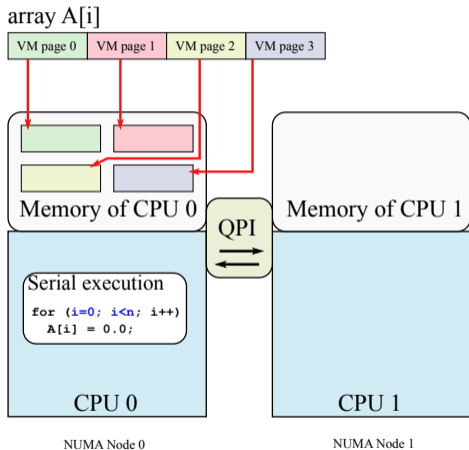
# WHAT IS THREAD AFFINITY

- ▶ OpenMP threads may migrate between cores
- ▶ Forbid migration — improve locality — increase performance
- ▶ Affinity patterns “scatter” and “compact” may improve cache sharing, relieve thread contention




# FIRST-TOUCH ALLOCATION POLICY


## Poor First-Touch Allocation



## Good First-Touch Allocation



# NESTED PARALLELISM WITH OPENMP



```

1  #pragma omp parallel
2  {
3  #pragma omp parallel
4    {
5      // ...
6    }
7  }

```

- ▶ Tune granularity of parallelism
- ▶ Improve resource sharing in NUMA systems



**LEARN MORE**



THE "HOW" SERIES

# DEEP DIVE

WITH CODE MODERNIZATION EXPERTS

\*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](https://colfaxresearch.com/how-series)




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

## Parallel Programming and Optimization with Intel® Xeon Phi™ Coprorocessors

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>

**COLFAX RESEARCH**  
CONTRIBUTING TO INNOVATIONS IN COMPUTING
Log In/Out or Register

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**

**Parallel Programming Book**

Introduction to parallel programming, deep discussion of optimization techniques, exercises. © 2015, Colfax International, 508 pages.

**Featured Video**

See Research material on vectorization in a training video

**Research and Educational Publications**

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

**Software Developer's Introduction to the HGST Ultrastar Archive H800 SMR Drives**

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

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

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

**Performance to Power and Performance to Cost Ratios with Intel Xeon Phi Coprocessors (and why ix Acceleration May Be Enough)**

**Consulting**

Colfax offers consulting services for enterprises, research help you:


- Optimize your existing application to take advantage of parallelism, from vectors to cores to clusters and beyond
- Future-proof your application for upcoming innovations
- Accelerate your application using coprocessor tech
- Investigate the potential system configurations that satisfy your cost, power, performance requirements.
- Take a clean slate to develop a novel approach to reduce your computing pro

**Episode 2.1 — Purpose of the MIC architecture**

**Parallel Computing in the Search for New Physics at LHC**

**COLFAX RESEARCH**  
CONTRIBUTING TO INNOVATIONS IN COMPUTING
Log In/Out or Register

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**

**Parallel Programming Book**

Introduction to parallel programming, deep discussion of optimization techniques, exercises. © 2015, Colfax International, 508 pages.

**Featured Video**

See Research material on vectorization in a training video

**Research and Educational Publications**

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

**Software Developer's Introduction to the HGST Ultrastar Archive H800 SMR Drives**

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

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

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

**Performance to Power and Performance to Cost Ratios with Intel Xeon Phi Coprocessors (and why ix Acceleration May Be Enough)**

**Consulting**

Colfax offers consulting services for enterprises, research help you:

- Optimize your existing application to take advantage of parallelism, from vectors to cores to clusters and beyond
- Future-proof your application for upcoming innovations
- Accelerate your application using coprocessor tech
- Investigate the potential system configurations that satisfy your cost, power, performance requirements.
- Take a clean slate to develop a novel approach to reduce your computing pro

**Episode 2.1 — Purpose of the MIC architecture**

**Parallel Computing in the Search for New Physics at LHC**

**COLFAX RESEARCH**  
CONTRIBUTING TO INNOVATIONS IN COMPUTING
Log In/Out or Register

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**

**Parallel Programming Book**

Introduction to parallel programming, deep discussion of optimization techniques, exercises. © 2015, Colfax International, 508 pages.

**Featured Video**

See Research material on vectorization in a training video

**Research and Educational Publications**

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

**Software Developer's Introduction to the HGST Ultrastar Archive H800 SMR Drives**

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

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

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

**Performance to Power and Performance to Cost Ratios with Intel Xeon Phi Coprocessors (and why ix Acceleration May Be Enough)**

**Software Developer's Introduction to the HGST Ultrastar Archive H800 SMR Drives**

**Fluid Dynamics with Fortran on Intel Xeon Phi coprocessors**

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

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

**Parallel Computing in the Search for New Physics at LHC**

http://colfaxresearch.com/

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



Find out more at [dap.xeonphi.com](http://dap.xeonphi.com)  
or contact us at [dap@colfax-intl.com](mailto:dap@colfax-intl.com)