Scientific Computing with Intel Xeon Phi Coprocessors

Andrey Vladimirov
Colfax International

HPC Advisory Council Stanford Conference 2015

Contents

§1 MIC Architecture, Developer's Perspective

§2 Case Studies

- Astrophysics (offload story)
- N-body simulation (offload vs native in a cluster)
- Finanical Monte Carlo (heterogeneous clustering)
- Computational fluid dynamics (legacy code)

§3 Colfax Developer Training

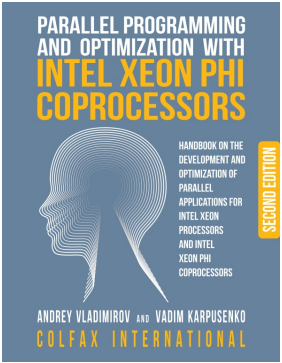

§1. MIC Architecture from Developer's Perspective

© Colfax International, 2014

Intel Xeon Phi Coprocessors and the MIC Architecture

- PCIe end-point device
- High Power efficiency
- ~ 1 TFLOP/s in DP
- Heterogeneous clustering

For highly parallel applications which reach the scaling limits on Intel Xeon processors

Examples of Solutions with the Intel MIC Architecture

Colfax's CXP7450 workstation with two Intel Xeon Phi coprocessors

xeonphi.com/workstations

Colfax's CXP9000 server with eight Intel Xeon Phi coprocessors

xeonphi.com/servers

© Colfax International, 2014

Intel Xeon Phi Coprocessors and the MIC Architecture

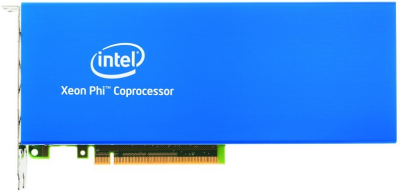

- ≤18 cores/socket ≈3 GHz
- 2-way hyper-threading
- Up to 768 GB of DDR3 RAM
- 256-bit AVX vectors

- 57 to 61 cores at ≈1 GHz
- 4 hardware threads per core
- 6–16 GB cached GDDR5 RAM
- 512-bit IMCI vectors

- C/C++/Fortran; OpenMP/MPI
- Linux OS (on host *and* on coprocessor)

Linux μOS on Intel Xeon Phi coprocessors (part of MPSS)

```
user@host% lspci | grep -i "co-processor"
06:00.0 Co-processor: Intel Corporation Xeon Phi coprocessor 3120 series (rev 20)
82:00.0 Co-processor: Intel Corporation Xeon Phi coprocessor 3120 series (rev 20)
user@host% sudo service mpss status
mpss is running
user@host% cat /etc/hosts | grep mic
172.31.1.1  host-mic0 mic0
172.31.2.1  host-mic1 mic1
user@host% ssh mic0
user@mic0% cat /proc/cpuinfo | grep proc | tail -n 3
processor : 237
processor : 238
processor : 239
user@mic0% ls /
amplxe   dev    home   lib64    oldroot  proc   sbin     sys    usr
bin      etc    lib    linuxrc  opt      root   sep3.10  tmp    var
```

Offload and Native modes

- Explicit offload mode:

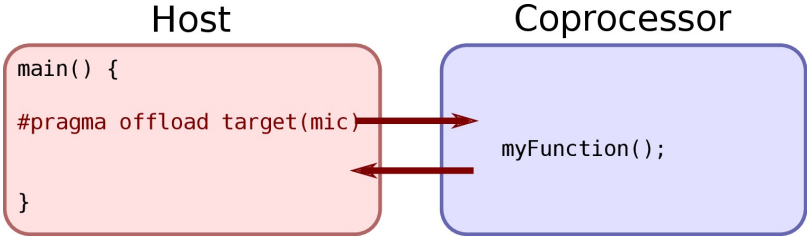

- Native mode:

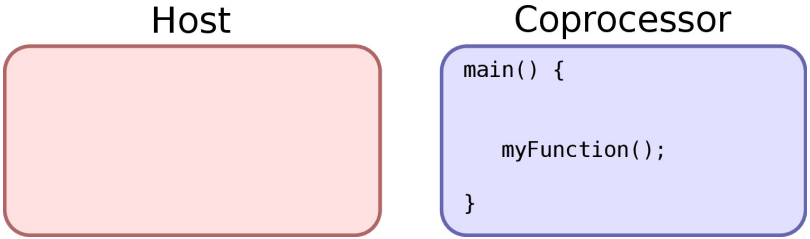

Optimization Areas

Common methods for Intel Xeon CPUs and Intel Xeon Phi coprocessors:

1. **Scalar optimization** (compiler-friendly practices)
2. **Vectorization** (must use 16- or 8-wide vectors)
3. **Multi-threading** (must scale to 100+ threads)
4. **Memory access** (streaming access or tiling)
5. **Communication** (offload, MPI traffic control)

© Colfax International, 2014

Getting Ready for the Future

- Knights Landing (KNL) – next generation of Intel MIC architecture
- 3x the performance of current generation
- Available as a stand-alone processor or as a coprocessor

Getting Ready for the Future

The best way to prepare applications for KNL is to optimize them for Intel Xeon Phi coprocessors based on KNC.

§2. Case Studies

Astrophysical Code HEATCODE: an Offload Story

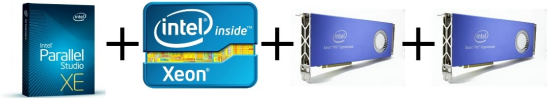

Porting to Intel® Xeon Phi™ coprocessors

- 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!

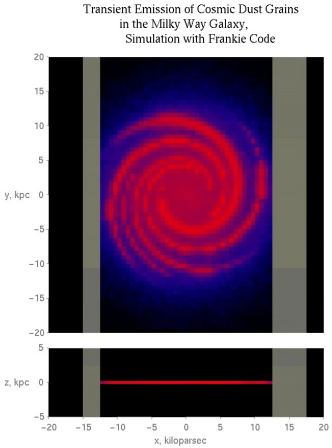

xeonphi.com/papers/heatcode

Astrophysical Code HEATCODE: an Offload Story

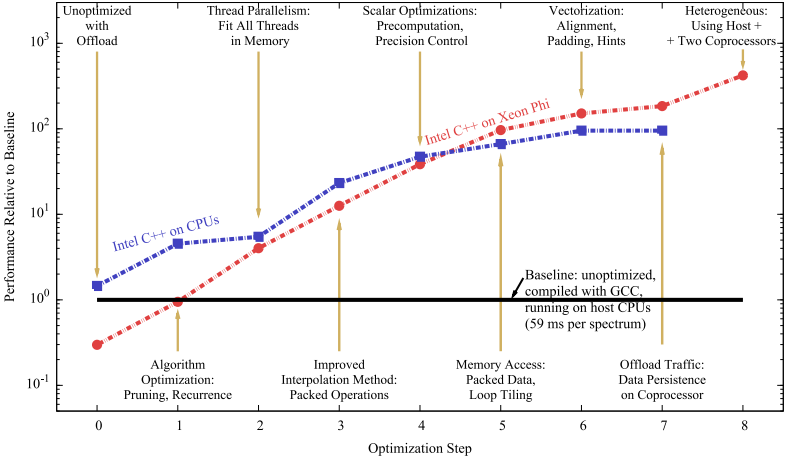

xeonphi.com/papers/heatcode

N-body Simulation: Offload vs Native in a Cluster

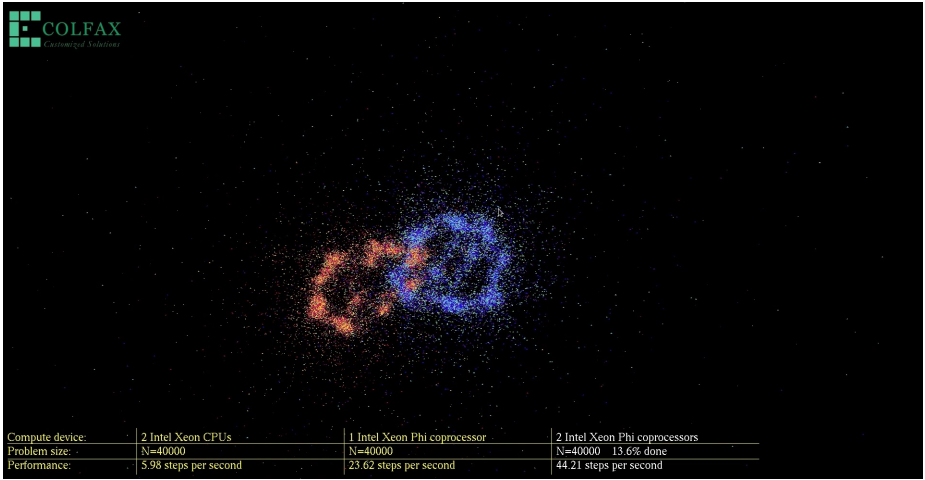

xeonphi.com/papers/nbody-basic

N-body Simulation: Offload vs Native in a Cluster

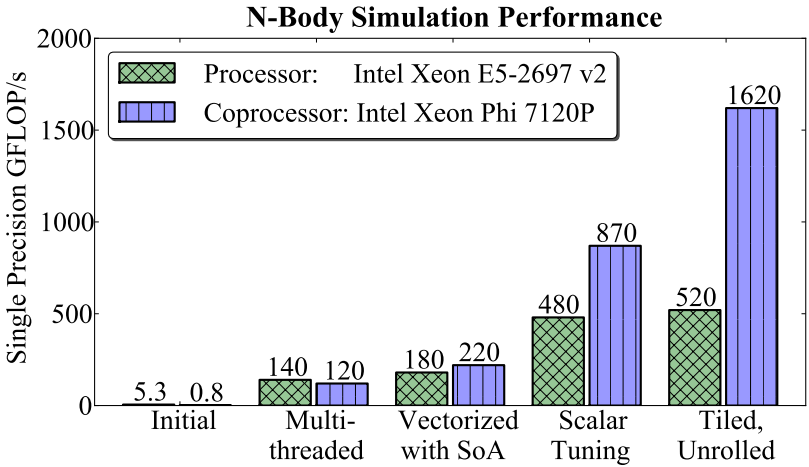

xeonphi.com/papers/sc14

N-body Simulation: Offload vs Native in a Cluster

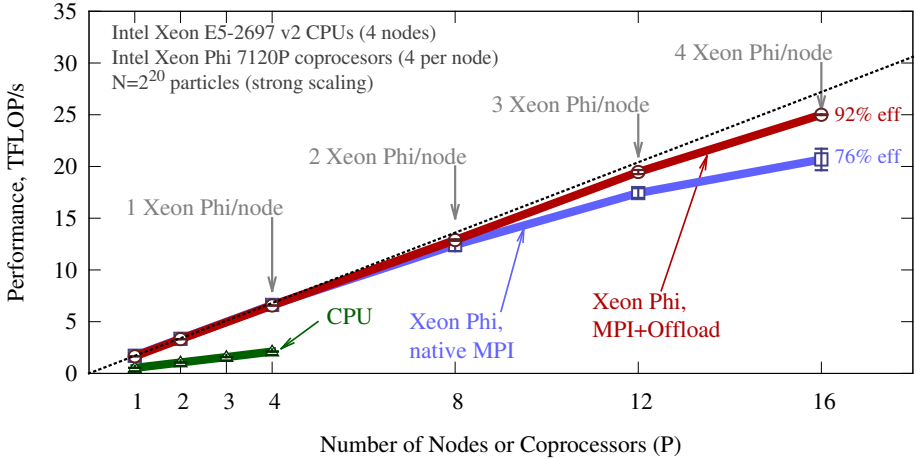

xeonphi.com/papers/sc14

Asian Option Pricing: Heterogeneous Clustering

Heterogeneous Clustering with Homogeneous Code:
Asian Option Pricing

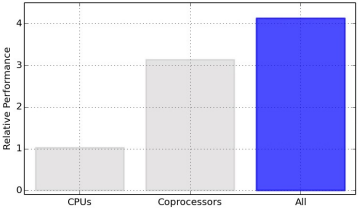

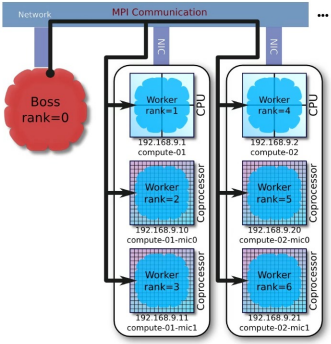

COLFAX
Customized Solutions

xeonphi.com/papers/heterogeneous

Computational Fluid Dynamics: Legacy Code

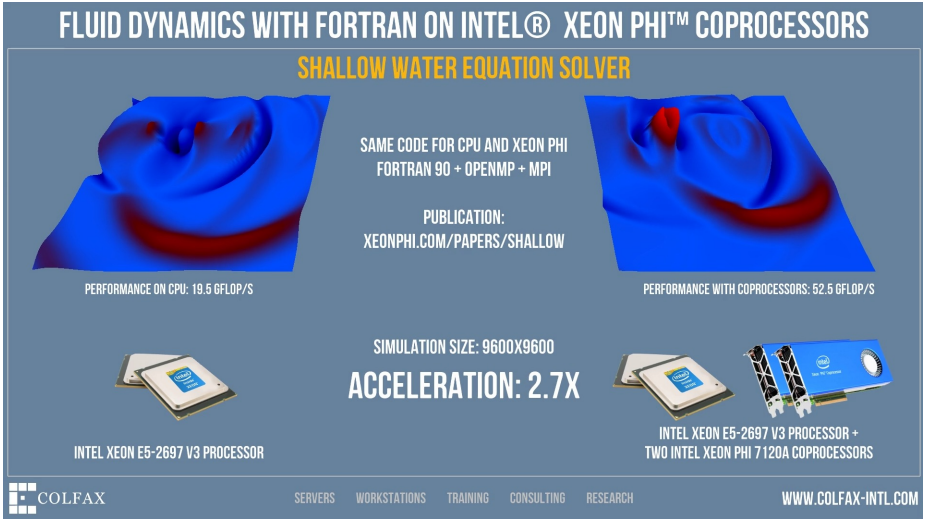

xeonphi.com/papers/shallow

§3. Colfax Developer Training

Colfax Developer Training

Intel Xeon Phi Coprocessor Programming

Future-Proofing Applications for Knights Landing (KNL)

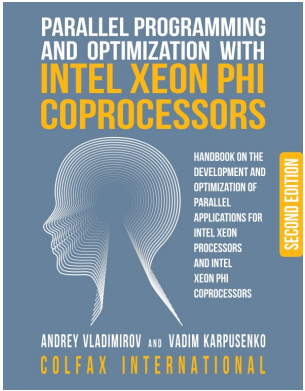

xeonphi.com/training

 © Colfax International, 2014

Free Training for HPCAC Stanford 2015 Participants

COLFAX DEVELOPER TRAINING FOR INTEL® XEON PHI™ COPROCESSOR

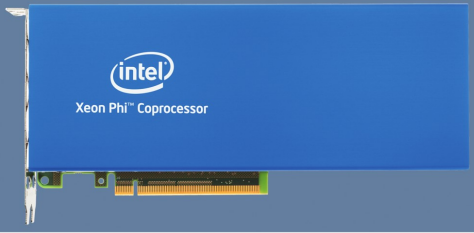

FOR A LIMITED TIME, COLFAX INTERNATIONAL IS OFFERING FREE 1-DAY AND 4-DAY CLASSES! CLASS SEATS ARE LIMITED.
SIGN UP AT: XEONPHI.COM/HPCAC2015

CDT 101: 1-DAY LECTURE SERIES

PROGRAMMING MODELS, PARALLEL FRAMEWORKS AND OPTIMIZATION PRACTICES FOR INTEL XEON PHI COPROCESSORS. LECTURE-ONLY TRAINING. PREREQUISITE FOR CDT102.

CDT 102: 1-DAY HANDS-ON LABS

HANDS-ON EXERCISES ON PROGRAMMING INTEL XEON PHI COPROCESSORS. TUTORIALS ON EXPRESSING PARALLELISM AND OPTIMIZING PERFORMANCE WITH A FOCUS ON METHODS APPLICABLE TO MULTI-CORE CPUS AS WELL AS INTEL MIC ARCHITECTURE. PREREQUISITE: CDT 101.

CDT 401: 4-DAY WORKSHOP

IN-DEPTH DISCUSSION AND EXERCISES ON PROGRAMMING INTEL XEON PHI COPROCESSORS, INCLUDING HETEROGENEOUS CLUSTERING WITH NATIVE AND OFFLOAD METHODS. NUMEROUS HANDS-ON EXAMPLES ON TUNING THE PERFORMANCE OF SCALAR MATH, VECTORIZATION, MULTI-THREADING, MEMORY TRAFFIC AND COMMUNICATION.

COLFAX INTERNATIONAL | WWW.COLFAX-INTL.COM | 1.408.730.2275

xeonphi.com/hpcac2015