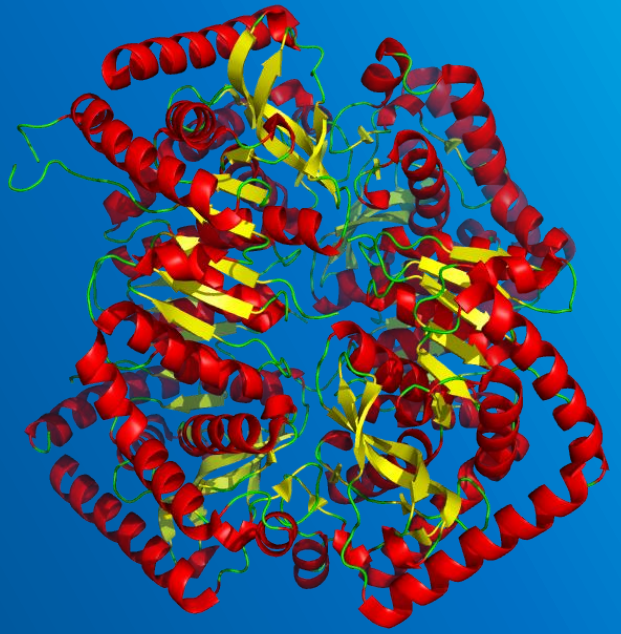

FASTER SIMULATIONS WITH LAMMPS*

Advancing Performance on Intel® Xeon® and Xeon Phi™ Processors

W. Michael Brown, Intel® Corporation
2015 LAMMPS Users Workshop and Symposium, Albuquerque, NM

Agenda

- Software challenges for HPC and why we believe the Intel® hardware and software roadmap can reduce software engineering, programming, and validation efforts
- Optimizations for many-core processors and the Intel package for LAMMPS*
- Performance results
- Current/future hardware and software efforts

Software Barrier for Modern HPC

The breakdown in Dennard scaling has led to significant changes in HPC computer architectures that achieve power-efficient performance.

Modifying legacy software for these systems can be a significant barrier to performance

- Example: Still only a subset of HPC codes can efficiently utilize multicore server processors with GPGPU accelerators at scale

The Intel roadmap seeks to address both challenges with x86-based many-core coprocessors and bootable processors that achieve performance with standard programming models

2016
Knights Landing
The processor version of the next generation Intel Xeon Phi product family
- 14 nm process
- Processor & Coprocessor
- Over 3 TF DP Peak[1]
- Up to 72 Cores
- On Package High-Bandwidth Memory
- 3x single-thread performance[2]
- Out-of-order core
- Integrated Intel® Omni-Path

2013
Knights Corner
Intel® Xeon Phi™ x100 product family
- 22 nm process
- Coprocessor
- Over 1 TF DP Peak
- Up to 61 Cores
- Up to 16GB GDDR5

FUTURE
Knights Hill
Next generation of Intel® MIC Architecture Product Line
- 10 nm process
- 2nd Generation Integrated Intel® Omni-Path
- In planning –

All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. All projections are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

[1]Over 3 Teraflops of peak theoretical double-precision performance is preliminary and based on current expecations of cores, clock frequency and floating point operations per cycle.
[2] Projected peak theoretical single-thread performance relative to 1st Generation Intel® Xeon Phi™ Coprocessor 7120P

Reducing Divergent Source Code

Acceleration with GPGPU can require different algorithms forcing an expanded source code base and further complicating optimization and validation for HPC systems

- Example: GPGPU acceleration for MD uses different algorithms and potentially different MPI* communications to efficiently handle the large number of lightweight threads in flight.

	Intel® Xeon® Processors	GPGPUs	Intel® Xeon Phi™ Coprocessors	Intel® Xeon Phi™ Bootable Processors
Performance	✓	✓	✓	✓
Electrical Power Efficiency for HPC		✓	✓	✓
Source Code Modification for Hybrid Computers Not Necessary	✓			✓
Software Modification for non-x86 Architectures Not Necessary	✓		✓	✓

The Software Challenge in Modern HPC

Example: Optimizations for GPGPU acceleration in LAMMPS* did not improve performance on CPUs

- Many production codes supporting GPU acceleration still use the CPU for many routines in order to achieve performance and must still support CPU-only

Optimizations for x86 coprocessors also improved CPU performance, but …

- Optimization still required
- What are the optimizations? →

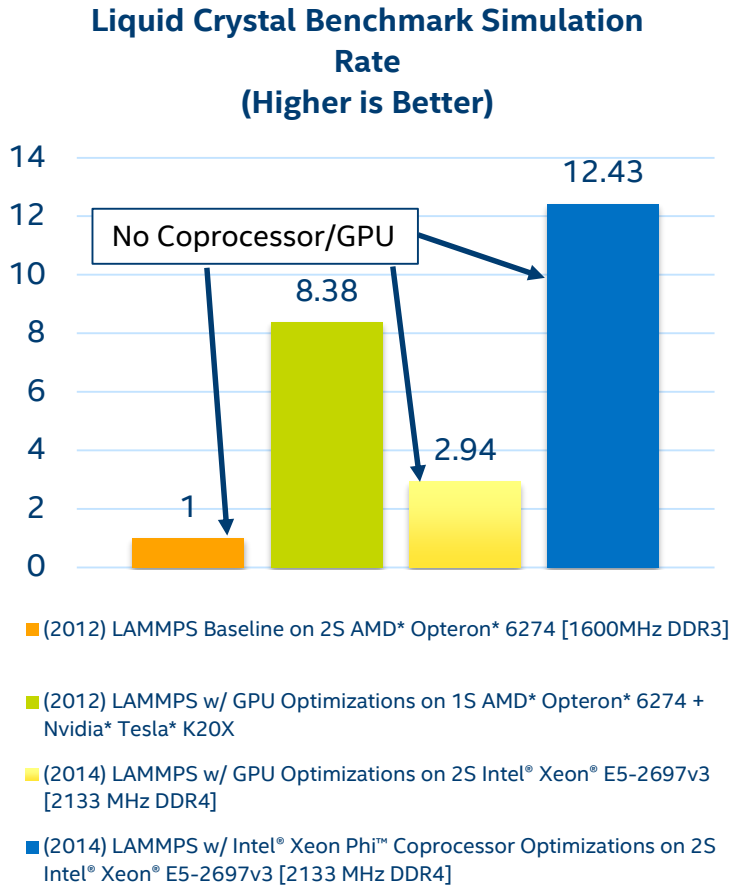

Source:
AMD* & AMD/NVIDIA* results: http://www.nvidia.com/docs/IO/122634/computational-chemistry-benchmarks.pdf
Intel Results: Intel Measured August 2014

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. See benchmark tests and configurations in the speaker notes. For more information go to http://www.intel.com/performance

Optimizing for Many-Core Processors

- The general optimization guidelines remain the same for codes that have been running on large-scale clusters and supercomputers for some time
- However, the increase in parallelism required for power efficiency with many-core can amplify bottlenecks due to:

Unvectorized code

- Use vectorization directives when compiler can't prove vectorization is safe

Inefficient data layout for vectorization and/or many cores sharing the memory subsystem

- Modify data layout for efficient (re)use of cache lines, software prefetching for random access, correctly aligned data, efficient loading of vector registers, and use hyper-threads as appropriate

Inefficient synchronization

- Avoid collective synchronization where possible

Lack of overlap of internode communications and computation

- Use non-blocking MPI* calls, task-based parallelism, etc. to minimize idle time

Suboptimal MPI task mapping of subdomains to physical cores

- Map MPI tasks to minimize the surface to volume ratio of domains mapped to nodes or NUMA nodes, not just individual MPI tasks

Intel Package for LAMMPS*

- Initially released in the 15 Aug 2014 version of LAMMPS
- Currently includes support for simulation of soft matter, biomolecules, 3-body potentials, and aspherical particles
- Single binary can be used for runs on CPUs with or without coprocessors in single, mixed, and double precision

Includes optimizations for:

- Data layout
  - More efficient use of cache lines, less vectorization overhead
- Vectorization
  - Vectorization for non-bonded force and energy calculations
- Multiple precision modes
  - Runtime selection of single, mixed, or double precision calculation
- Use of coprocessors
  - Same source code routine is called for runs on coprocessor and CPU
    - Optimize in a single source code path
    - Directives manage data transfer, offload, and synchronization
      - Ignored for compilers and/or systems that don't support coprocessors

Using the Intel Package in LAMMPS*

https://software.intel.com/en-us/articles/lammps-for-intel-xeon-phi-coprocessor

Can also use command-line options `-pk` and `-sf` without modifying input script.

Choose number of coprocessors to use on each node.

Select floating-point precision mode (**`mixed`** recommended):

{ **`single, mixed, double`** }

```
# Rhodopsin model
package intel 1 mode mixed balance -1
package omp 0
suffix intel

units           real
neigh_modify    delay 5 every 1

# ...

timestep        2.0
run 10
run             100
```

Choose fraction of work to offload to coprocessor.

`0.0` → Run optimized routines without using coprocessor.

`0.5` → Calculations for half of the atoms are on coprocessor.

`-1` → Fraction adjusted automatically by load balancer

For benchmarking with short runs, add a warm up-run so load balancer and other startup penalties are not included in time.

Building and Running w/ Intel Package

Install with '`make yes-user-intel`' in `src` directory.

Example Makefiles in `src/MAKE/OPTIONS/Makefile.intel*`

Run LAMMPS as usual

```
> export OMP_NUM_THREADS=1
> mpirun -n 48 -perhost 24 ./lmp_intel_phi -in
in.intel.rhodo

> export OMP_NUM_THREADS=2
> mpirun -n 24 -perhost 12 ./lmp_intel_phi -in
in.intel.rhodo
```

When using offload to a coprocessor, better to have Intel® Hyper-threading Technology enabled for threads using non-blocking offload calls.

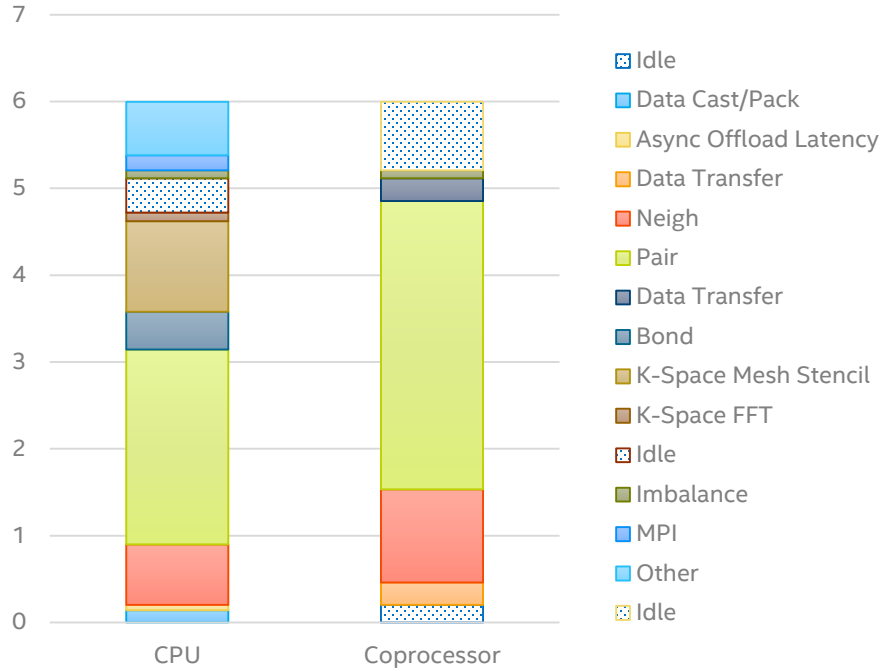

Concurrent calculations and data movement for Rhodopsin

PERFORMANCE RESULTS

Organic Solar Cells

Science Team: Jan-Michael Y Carrillo, Rajeev Kumar, Monojoy Goswami, S. Michael Kilbey II, Bobby G Sumpter (ORNL/UT*)*

Problem: Predictive simulation of active layer morphology and molecular alignment based on blend composition

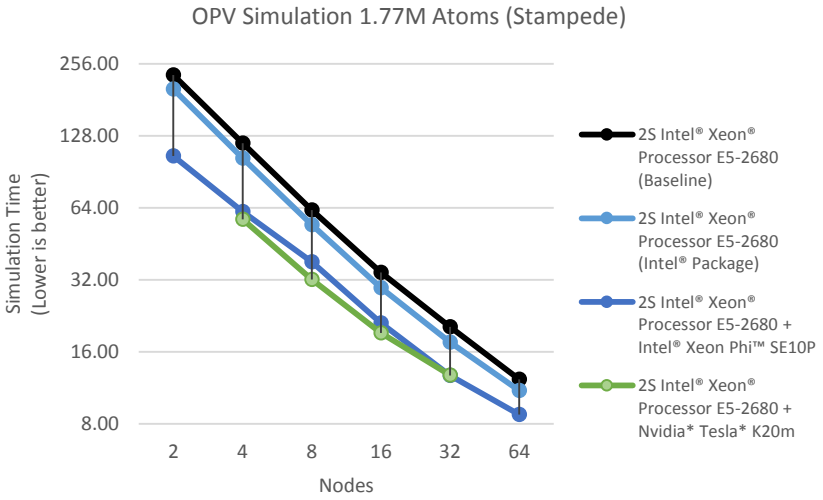

Result: Up to 2.2X performance improvement on Stampede* (1.9X with use of a coprocessor)

- *Simulations include all of the statistics and I/O (about 10% of run time) from the production runs*

- Significant potential for advanced multiscale simulation models with coprocessors…

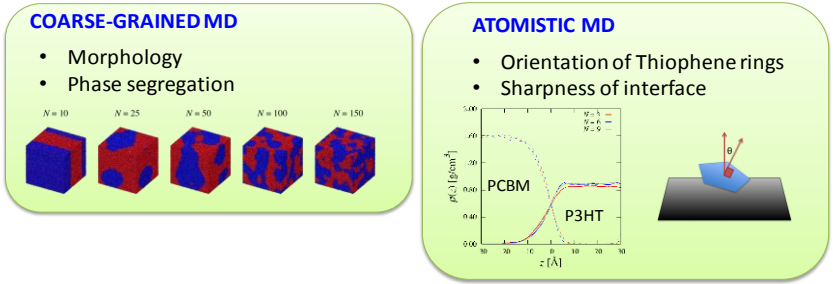

Stampede Configuration: HT Off, 32GB DDR3-1600MHz, PCIe2 (GPU), PCIe2 (Intel® Coprocessor), FDR 56 Gb/s, MPSS 3.3, MVAPICH 2.0b

Source: Michael Brown

Hydrocarbon Lubricants

Science Team: Mishra Biswajit, Nitin Gavhane, Foram M. Thakker, A.R. De Kraker (Shell India Markets Private Limited)*

Problem: Study thermodynamic, transport, and rheological properties of hydrocarbon mixtures (typically 30-carbon atoms long) used as lubricants

- CHARMM* force field, 11Å/13Å cutoff
- 2Å skin, PPPM 1e-4
- NPT ensemble, 165K Atoms

Result: *Up to 2.3X simulation rate on Stampede* (1.6X from use of a coprocessor)*

- *Simulations include all of the statistics and I/O from the production runs*

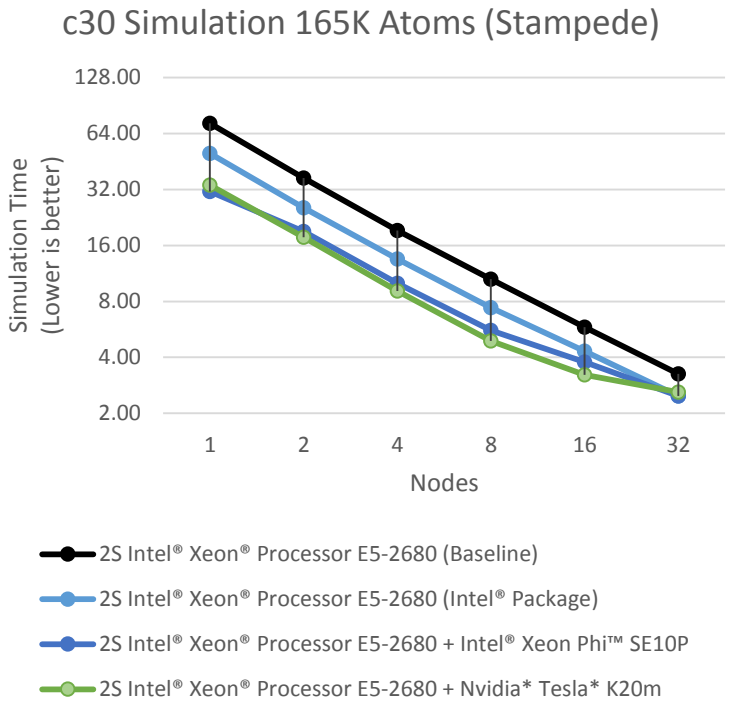

Source: Michael Brown

Water Simulation w/ Stillinger-Weber Potential

Science Team: Yamada Masako, et. al (GE Global Research)*

Problem: Probe mechanism for water droplet freezing on a surface with molecular detail

Innovation: New 3-body potential for water simulation (E. B. Moore, V. Molinero, *Nature** 479 (2011) 506–508)

- Coarse-grain model, eliminate all-to-all communications for electrostatics, larger timestep
  - > 100X Simulation rate speedup from model

Result: Stampede* simulations up to 5.16X faster, (1.5X with coprocessor)

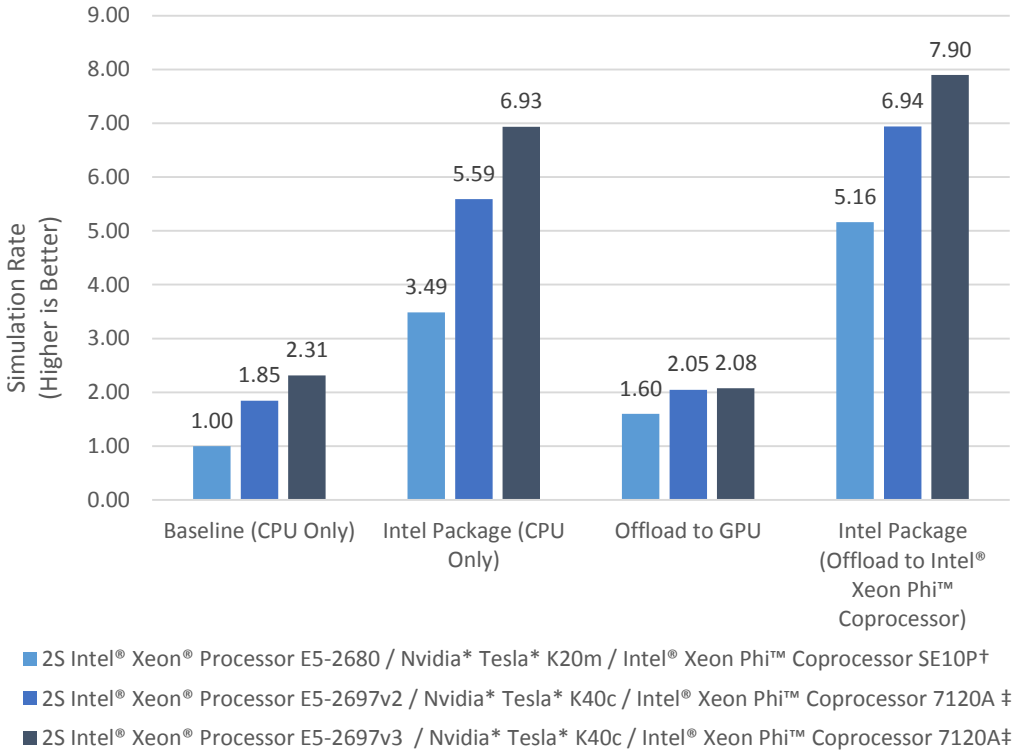

Source: Intel Measured June 2015

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. See benchmark tests and configurations in the speaker notes. For more information go to http://www.intel.com/performance

Rhodopsin Protein Benchmark

Standard LAMMPS Benchmark Scaled to 512K Atoms

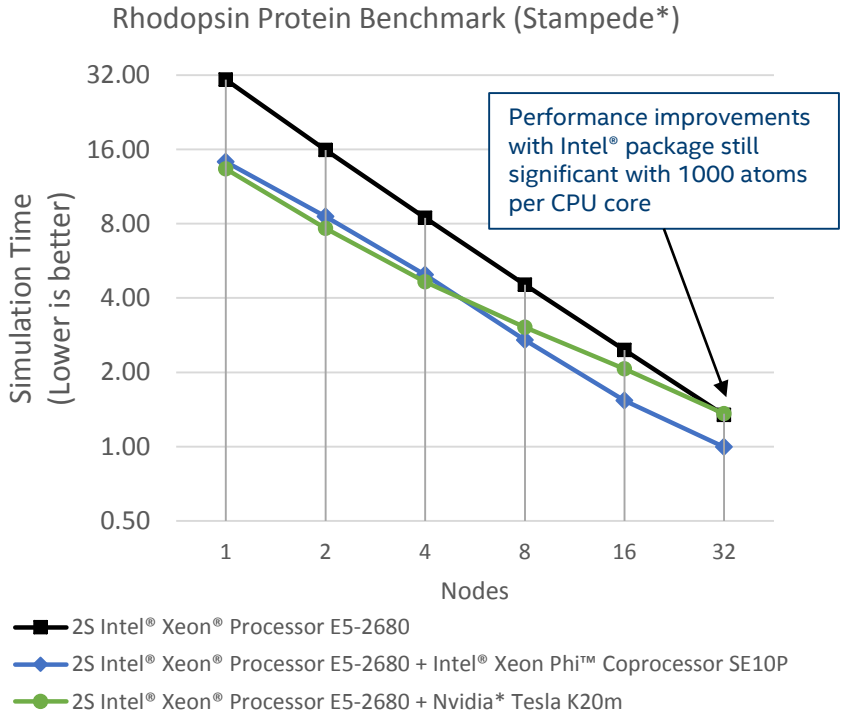

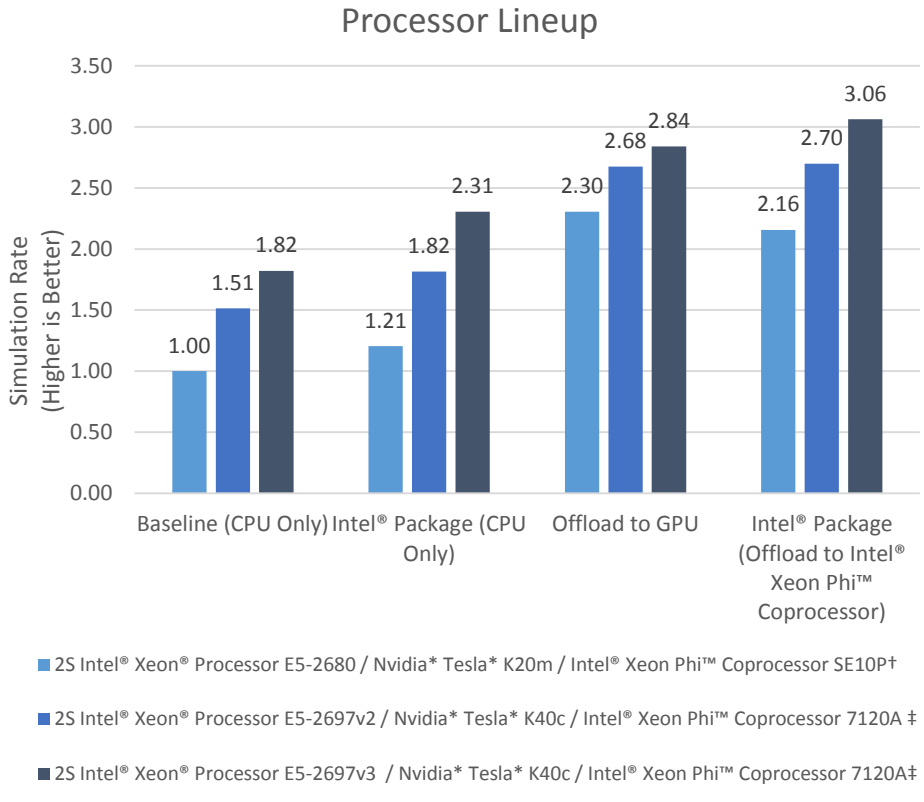

Source: Intel Measured August 2014

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. See benchmark tests and configurations in the speaker notes. For more information go to http://www.intel.com/performance

Liquid Crystal Benchmark

Included with LAMMPS, Aspherical coarse-grain liquid crystal mesogens, 524K Particles*

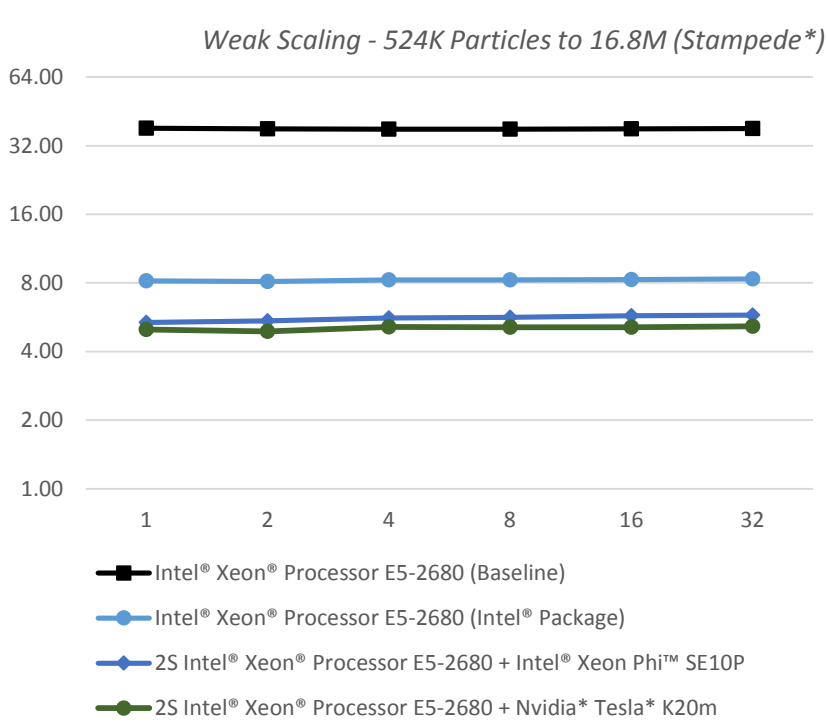

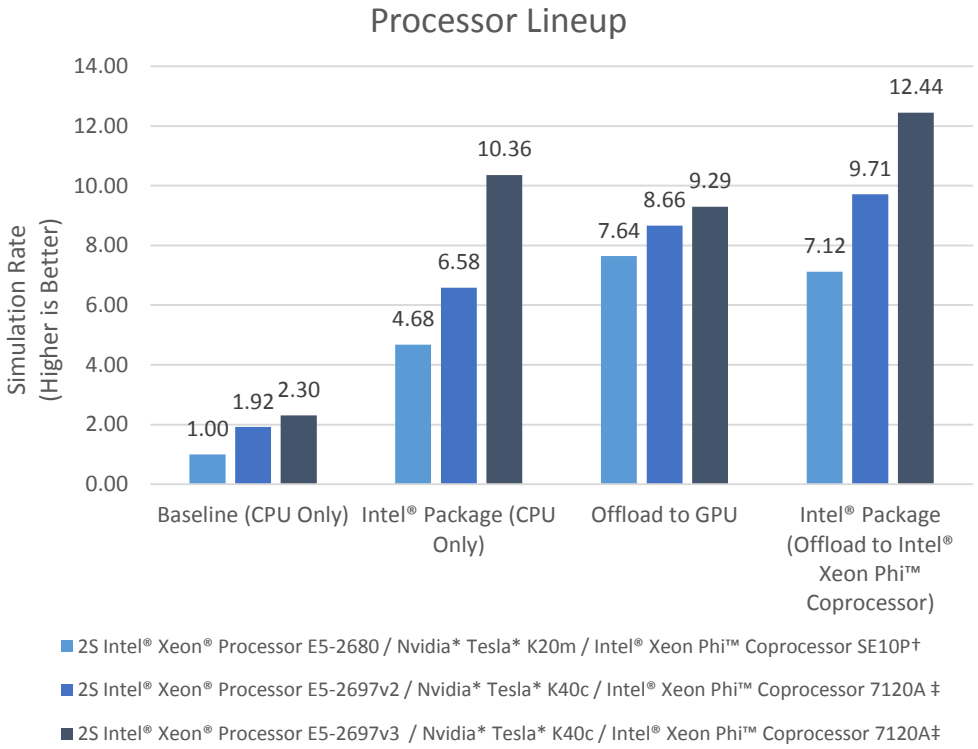

Source: Intel Measured August 2014

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. See benchmark tests and configurations in the speaker notes. For more information go to http://www.intel.com/performance

CURRENT/FUTURE EFFORTS

Intel® Xeon Phi™ (Knights Landing) Processors

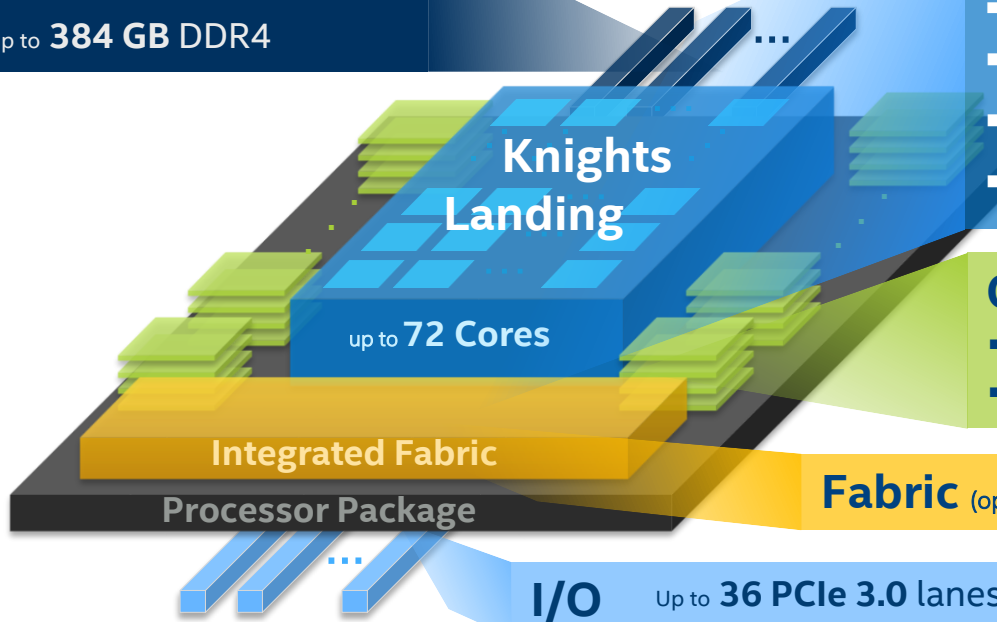

Platform Memory

Up to **384 GB** DDR4

Compute

- Intel® Xeon® Processor Binary-Compatible
- **3+ TF**LOPS[1], **3X ST**[2] (single-thread) perf. vs KNC
- **2D Mesh** Architecture
- **Out-of-Order** Cores

On-Package Memory

- Up to **16 GB** *at launch*
- Over **5x STREAM**[3] *vs. DDR4 at launch*

Fabric (optional)

- **1st** Intel processor to integrate

I/O Up to **36 PCIe 3.0** lanes

>50 System Providers Expected[4]

[1]Over 3 Teraflops of peak theoretical double-precision performance is preliminary and based on current expectations of cores, clock frequency and floating point operations per cycle. FLOPS = cores x clock frequency x floating-point operations per second per cycle. .
[2]Projected peak theoretical single-thread performance relative to 1st Generation Intel® Xeon Phi™ Coprocessor 7120P (formerly codenamed Knights Corner).
[3]Projected result based on internal Intel analysis of STREAM benchmark using a Knights Landing processor with 16GB of ultra high-bandwidth versus DDR4 memory only with all channels populated.
[4] Intel internal estimate

First *Public* Knights Landing Systems

Cori

"NERSC, Cray, Intel Partner on **Next-gen Extreme-Scale Computing System**... named Cori"
Scientific Computing – Apr 30, 2014

"... our goal is to enable performance that is **portable** across systems and will be **sustained** in future supercomputing architectures."
Scientific Computing – Apr 30, 2014

Trinity

"... **multipetaflop supercomputer called Trinity**"
PCWorld – Jul 10, 2014

"...powered by future Intel Xeon and Intel Xeon Phi processors, will deliver great application performance for a wide set of codes while the **binary compatibility between the processors will allow the NNSA to reuse existing codes.**"
insideHPC – Jul 10, 2014

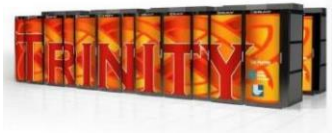

Theta

"Intel is working in-cooperation with Cray to deliver **two supercomputers for Argonne National Laboratory, named as Aurora and Theta**" *wccftech.com – Apr 9, 2015*

"Theta Supecomputer – **8.5 Pflops With Knights Landing in 2016**" *wccftech.com – Apr 9, 2015*

Future LAMMPS* Efforts

- Vectorization in other code regions, AVX512 conflict detection
- Memory prefetch optimizations for x200 family Intel® Xeon Phi™ processors
- Efficient non-blocking PPPM communications
- Support for additional simulation models

Intel® Parallel Computing Centers (IPCC) Community focused on LAMMPS

- Chinese Academy of Sciences – Computer Network Information Center*
- RWTH Aachen University*

Open-Source Contributions

- Shell India Markets Private Limited*
- ***Any contributors welcome, discussion at the workshop or contact via e-mail***

Additional References

Additional details on LAMMPS optimizations in the Intel Package

- Brown, W.M., Carrillo, J.-M.Y., Gavhane, N., Thakkar, F.M., Plimpton, S.J. **Optimizing Legacy Molecular Dynamics Software with Directive-Based Offload**. *Computer Physics Communications*. 2015. 195: p. 95-101.

  Preprint: https://sites.google.com/site/wmbrown85/lammps_mic1.pdf

- Optimization Webinar:

  https://software.intel.com/sites/default/files/managed/72/19/Optimizing-LAMMPS-for-Intel-Xeon-Phi-Coprocessors-slides.pdf

Summary

We continue to see significant performance improvements for molecular dynamics simulation in LAMMPS* with hardware and software advances

- 3X-12X simulation rates compared to several years ago

Intel is working to provide solutions for HPC intended to allow software developers to achieve performance with reduced programming effort on future power-efficient processors

- Standard programming models such as MPI* and OpenMP*
- A community effort is developing to continue to optimize LAMMPS performance for current and future processors
- The Intel package allows a workspace for experimenting with optimization options that can ultimately be used to inform changes important for the default code base
  - All optimizations are tested thoroughly before release allowing scientists to immediately realize benefits from the effort

Acknowledgements

Research by the teams presented here made possible by:

NSF TACC* Stampede Project: ACI-1134872

The researchers acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin* for providing HPC resources that have contributed to the research results reported within this presentation. URL: http://www.tacc.utexas.edu*

NSF UT* Beacon Project

This material is based upon work supported by the National Science Foundation under Grant Number 1137097 and by the University of Tennessee through the Beacon Project. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation* or the University of Tennessee*.*

Legal Notices and Disclaimers

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY RELATING TO SALE AND/OR USE OF INTEL PRODUCTS, INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT, OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life-saving, life-sustaining, critical control or safety systems, or in nuclear facility applications.

Intel products may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel may make changes to dates, specifications, product descriptions, and plans referenced in this document at any time, without notice.

This document may contain information on products in the design phase of development. The information herein is subject to change without notice. Do not finalize a design with this information.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Intel Corporation or its subsidiaries in the United States and other countries may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

Wireless connectivity and some features may require you to purchase additional software, services or external hardware.

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, visit Intel Performance Benchmark Limitations

Intel, the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Other names and brands may be claimed as the property of others.

Legal Disclaimer & Optimization Notice

INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS”. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

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

Thanks!