



Advancing Science in Alternative Energy and Bioengineering with Many-Core Processors

W. Michael Brown, *Intel® Corporation*[†]

Jan-Michael Carrillo, *Oak Ridge National Laboratory*^{*}

[†]*Disclaimer:* This talk includes results from my tenure at ORNL^{*} that do not necessarily reflect the views, research directions, etc. of Intel® Corporation.



Legal Disclaimers

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. 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.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. 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. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document 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.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>

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.

Intel does not control or audit the design or implementation of third party benchmarks or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmarks are reported and confirm whether the referenced benchmarks are accurate and reflect performance of systems available for purchase.

Relative performance is calculated by assigning a baseline value of 1.0 to one benchmark result, and then dividing the actual benchmark result for the baseline platform into each of the specific benchmark results of each of the other platforms, and assigning them a relative performance number that correlates with the performance improvements reported.

SPEC, SPECint, SPECfp, SPECrate, SPECpower, SPECAppServer, SPECjbb, SPECjvm, SPECWeb, SPECCompM, SPECCompL, SPEC MPI, SPECjEnterprise* are trademarks of the Standard Performance Evaluation Corporation. See <http://www.spec.org> for more information. TPC-C, TPC-H, TPC-E are trademarks of the Transaction Processing Council. See <http://www.tpc.org> for more information.

Hyper-Threading Technology requires a computer system with a processor supporting HT Technology and an HT Technology-enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. For more information including details on which processors support HT Technology, see [here](#)

Intel® Turbo Boost Technology requires a Platform with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software and overall system configuration. Check with your platform manufacturer on whether your system delivers Intel Turbo Boost Technology. For more information, see <http://www.intel.com/technology/turboboost>

No computer system can provide absolute security. Requires an enabled Intel® processor and software optimized for use of the technology. Consult your system manufacturer and/or software vendor for more information.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families: Go to: [Learn About Intel® Processor Numbers](#)

Intel product plans in this presentation do not constitute Intel plan of record product roadmaps. Please contact your Intel representative to obtain Intel's current plan of record product roadmaps.

Copyright © 2014 Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon and Intel Core are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. All dates and products specified are for planning purposes only and are subject to change without notice

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



Risk Factors

The above statements and any others in this document that refer to plans and expectations for the third quarter, the year and the future are forward-looking statements that involve a number of risks and uncertainties. Words such as “anticipates,” “expects,” “intends,” “plans,” “believes,” “seeks,” “estimates,” “may,” “will,” “should” and their variations identify forward-looking statements. Statements that refer to or are based on projections, uncertain events or assumptions also identify forward-looking statements. Many factors could affect Intel's actual results, and variances from Intel's current expectations regarding such factors could cause actual results to differ materially from those expressed in these forward-looking statements. Intel presently considers the following to be the important factors that could cause actual results to differ materially from the company's expectations. Demand could be different from Intel's expectations due to factors including changes in business and economic conditions; customer acceptance of Intel's and competitors' products; supply constraints and other disruptions affecting customers; changes in customer order patterns including order cancellations; and changes in the level of inventory at customers. Uncertainty in global economic and financial conditions poses a risk that consumers and businesses may defer purchases in response to negative financial events, which could negatively affect product demand and other related matters. Intel operates in intensely competitive industries that are characterized by a high percentage of costs that are fixed or difficult to reduce in the short term and product demand that is highly variable and difficult to forecast. Revenue and the gross margin percentage are affected by the timing of Intel product introductions and the demand for and market acceptance of Intel's products; actions taken by Intel's competitors, including product offerings and introductions, marketing programs and pricing pressures and Intel's response to such actions; and Intel's ability to respond quickly to technological developments and to incorporate new features into its products. The gross margin percentage could vary significantly from expectations based on capacity utilization; variations in inventory valuation, including variations related to the timing of qualifying products for sale; changes in revenue levels; segment product mix; the timing and execution of the manufacturing ramp and associated costs; start-up costs; excess or obsolete inventory; changes in unit costs; defects or disruptions in the supply of materials or resources; product manufacturing quality/yields; and impairments of long-lived assets, including manufacturing, assembly/test and intangible assets. Intel's results could be affected by adverse economic, social, political and physical/infrastructure conditions in countries where Intel, its customers or its suppliers operate, including military conflict and other security risks, natural disasters, infrastructure disruptions, health concerns and fluctuations in currency exchange rates. Expenses, particularly certain marketing and compensation expenses, as well as restructuring and asset impairment charges, vary depending on the level of demand for Intel's products and the level of revenue and profits. Intel's results could be affected by the timing of closing of acquisitions and divestitures. Intel's results could be affected by adverse effects associated with product defects and errata (deviations from published specifications), and by litigation or regulatory matters involving intellectual property, stockholder, consumer, antitrust, disclosure and other issues, such as the litigation and regulatory matters described in Intel's SEC reports. An unfavorable ruling could include monetary damages or an injunction prohibiting Intel from manufacturing or selling one or more products, precluding particular business practices, impacting Intel's ability to design its products, or requiring other remedies such as compulsory licensing of intellectual property. A detailed discussion of these and other factors that could affect Intel's results is included in Intel's SEC filings, including the company's most recent reports on Form 10-Q, Form 10-K and earnings release.

Optimization Notice

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 SSE3 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

Outline


- Motivation:* A subset of science problems I have been involved with that have been difficult to investigate with previous hardware/software solutions
- What we have done:* Recent hardware and software advances that have enabled new investigations
- What we are doing:* Issues and the solutions required to continue to advance HPC capabilities

3 Example Science Drivers

Liquid Crystal Biosensors

Label Free Diagnostics

- Alignment changes induced by interactions with biomolecules can be detected with optical microscopy
- *"Critical to this area is a sound understanding of the theory and modelling of liquid-crystal materials at interfaces. Computer simulations of the molecular and mesoscopic interactions between thermotropic liquid-crystal materials ... are a necessity. It is important to understand the interactions at these complicated interfaces"* - Woltman, et al., *Nature Materials**, 6, 929 (2007)
- Prior to 2013, no molecular simulations at experimentally relevant scales
 - Computational limitation
 - Extrapolations from small-scale studies difficult/dangerous due to the presence of undulative modes



Woltman, et al., *Nature Materials**, 6, 929 (2007)

Icephobic Surfaces for Wind Power

Many of the regions of the world that can benefit from wind power are in cold climates

Ice can reduce the efficiency of turbines or force shutdown

Extremely difficult to probe ice formation experimentally

Probing with simulation requires large simulations at time scales sufficient to capture rare nucleation events leading to ice formation



Organic Solar Cells

Organic photovoltaic (OPV) solar cells are promising renewable energy sources


- Low cost, high-flexibility, light weight

Morphology of the OPV active layer has a significant impact on device efficiency

- Donor regions should be large enough to absorb light efficiently, but small enough to allow charge excitations to diffuse to acceptor before recombining
- High interface-to-volume ratio for efficient exciton dissociation
- Proper donor molecular alignment to optimize the carrier mobility of the donor phase

Multi-scale problem

- Very large simulations to reach experimental scales



Hardware and Software Advances

Many-core and GPUs

Issues with power consumption, heat dissipation, and memory access latencies have forced a change in direction towards many-core processors for current and next-generation supercomputers that advance the limits of peak floating point performance

GPU-based architectures were the first affordable many-core processors generally available

ORNL* Titan came online in late 2012 as a hybrid supercomputer capable of over 17PF HPL performance.

- Cray* XK7 architecture with a Gemini* network interconnect and nodes containing a single AMD* Opteron* 6274 processor and Nvidia* Tesla K20X

Center for Accelerated Application Readiness (CAAR at ORNL*)

In order to exploit the potential performance gains on Titan, changes to the software are required

CAAR was formed to address this issue for several codes before the upgrade to Titan

- LAMMPS* (Large-scale Atomic/Molecular Massively Parallel Simulator) was the molecular dynamics code chosen for CAAR
- We implemented the 'GPU package' for LAMMPS* with new algorithms suited for running on hybrid GPGPU architectures

Brown, W.M., Wang, P. Plimpton, S.J., Tharrington, A.N. **Implementing Molecular Dynamics on Hybrid High Performance Computers - Short Range Forces.** *Computer Physics Communications**. 2011. 182: p. 898-911.

Brown, W.M., Kohlmeyer, A. Plimpton, S.J., Tharrington, A.N. **Implementing Molecular Dynamics on Hybrid High Performance Computers - Particle-Particle Particle-Mesh.** *Computer Physics Communications**. 2012. 183: p. 449-459.

Spinoidal Dewetting in Liquid Crystal Layers


Science Team: Trung Dac Nguyen, Jan-Michael Carrillo, Mike Matheson, W. Michael Brown, (ORNL*)

Problem: How/why do defects form in liquid crystal layers?

- Not possible to study on previous Jaguar supercomputer with software available unless the machine was dedicated to the problem

Innovation: Coarse-grain simulation model with high arithmetic intensity/vector potential, GPU algorithms

Result: > 7X performance gain (1S CPU+GPU vs 2S CPU); First simulations at scale with molecular detail



Nguyen, T. D., Carrillo, J.-M. Y., Matheson, M. A., Brown, W.M., **Rupture Mechanism of liquid crystal thin films realized by large-scale molecular simulations.** *Nanoscale**. 2014. 6: 3083-96.

Icephobic Surfaces


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
- Additional concurrency for 3-body: $O(N^3)$ independent force computations vs $O(N^2)$
[for N atoms within cutoff radius]
 - Additional 2-3X Simulation speedup on XK7 node (1S CPU+GPU) vs 2S Opteron*

Result: Simulation of water freezing process now typical with hundreds of nodes.



Brown, W.M., Masako, Y. **Implementing Molecular Dynamics on Hybrid High Performance Computers – Three-Body Potentials.** *Computer Physics Communications**. 2013. 184: p. 2785-2793.

Issues with Hybrid GPU Machines

Issues with the Programming Model

Using CUDA* with C/C++ semantics introduces **multiple languages**, code paths, compiler requirements, etc. into a code base.

- Support for Fortran or x86 targets requires proprietary compilers

OpenACC* can potentially address this, but in my opinion, will still require **separate code paths** and **different algorithms** for x86 in general

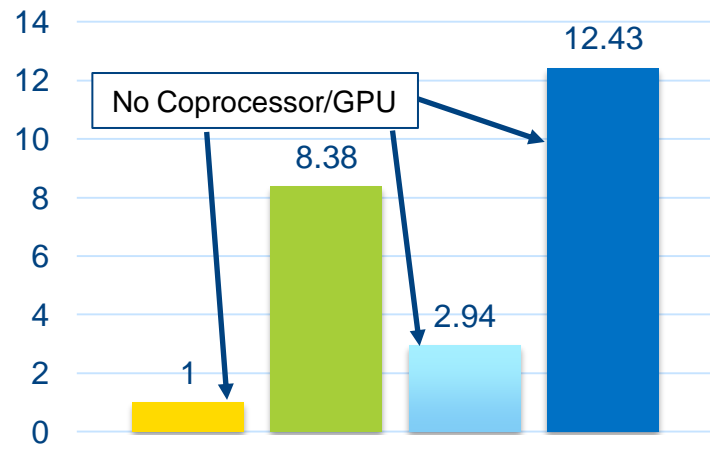
- GPUs have 10,000s of threads in flight, different performance for atomic operations, different penalties for thread synchronization, context switching, etc.
- For example, the algorithms used for molecular dynamics on the GPU would not perform well on x86 processors.
 - For the 3-body water/ice simulation presented here, we *triple* the number of force computations on the GPU with redundant computations!

Optimizing separate code paths

Optimizations for GPU do not necessarily improve performance on CPUs

- Many production codes supporting GPU acceleration still use the CPU for many routines and must still support CPU-only
- Intel® Xeon® processors are still improving the performance/socket
- 85% of all systems on TOP500* use Intel® processors (97% of new systems)
- Future Intel® many-core processors will be bootable (no coprocessor necessary)
- Difficult to manage/debug/port software with separate optimization code paths

Liquid Crystal Benchmark Simulation Rate (Higher is Better)



■ (2012) LAMMPS Baseline on 2S AMD* Opteron* 6274 [1600MHz DDR3]

■ (2012) LAMMPS w/ GPU Optimizations on 1S AMD* Opteron* 6274 + Nvidia* Tesla* K20X

■ (2014) LAMMPS w/ GPU Optimizations on 2S Intel® Xeon® E5-2697v3 [2133 MHz DDR4]

■ (2014) LAMMPS w/ Intel® Xeon Phi™ Coprocessor Optimizations on 2S Intel® Xeon® E5-2697v3 [2133 MHz DDR4]

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>

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

Intel® Package for LAMMPS

Motivation

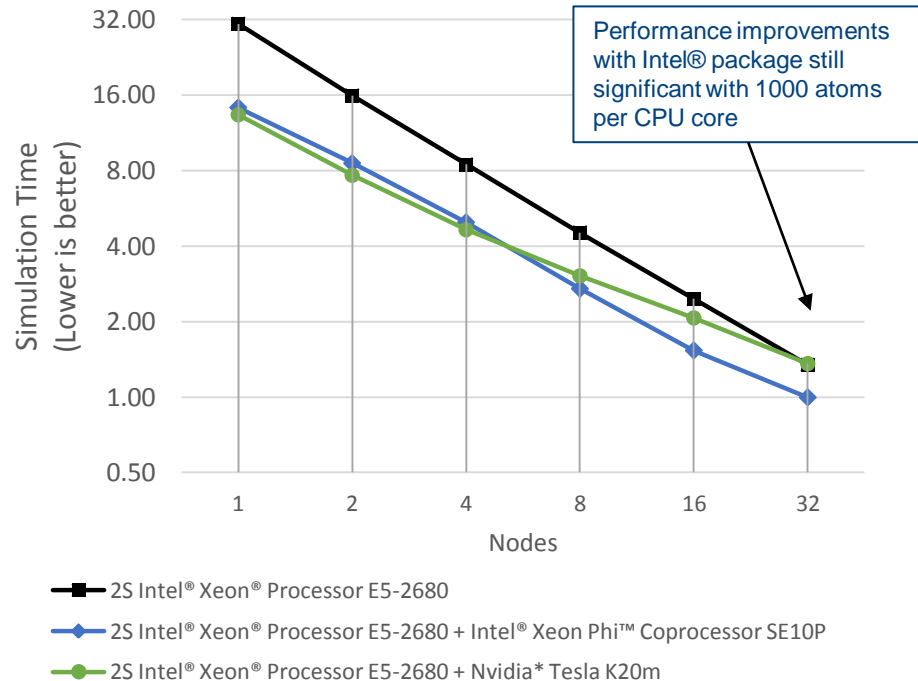
Provide a workspace for demonstrating code modernization strategies with vectorization for different precision modes, OpenMP*, and MPI*-3

Support computation offload to Intel® coprocessors with performance that is comparable or better than GPU performance

Demonstrate portable performance with a single code path using standard MD algorithms with C++/OpenMP*

Provide routines that allow the community to easily add new functionality by example

LAMMPS* Rhodopsin Protein Benchmark 512K Atoms (TACC* Stampede)



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

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>

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

Current Optimizations in Intel Package

Data alignment

Support for mixed and single precision modes in addition to double

SIMD directives to allow compiler vectorization for routines with data dependencies

Modifications to conditional branches to better support compiler vectorization for both coprocessors and Intel® Advanced Vector Extensions (Intel® AVX)

Neighbor-list padding to prevent execution of vector remainder code

Offload directives to manage data allocation, transfer, and concurrent computation on the coprocessor

Advantages of Intel® Package vs GPU Package (1)

- Same code for routines run on the CPU and coprocessor (with or without offload)
 - Optimizations for Intel® Xeon Phi™ coprocessors resulted in faster performance on Intel® Xeon® processors (up to 3.5X)
 - GPU package uses **different algorithms** and different code/language
- Support for both ‘newton’ settings allows for more flexibility for new force-fields
- Improved flexibility for heterogeneous calculations
 - Intel® coprocessor offload not limited to 16 MPI* tasks on CPU (CUDA*-MPS limitation)
 - Intel® package supports OpenMP* with multiple threads on the CPU (GPU package does not use OpenMP)
 - MPI* tasks sharing coprocessor are able to get exclusive core affinity

Advantages of Intel® Package vs GPU Package (2)

- More options for overlap of MPI* communications and computation
- Build process is simpler and does not require building a separate library for coprocessor routines
 - One compiler/Makefile for everything
- Precision mode (single, mixed, or double) can be switched at run-time without rebuilding
- Package written in standard C++ with OpenMP*
 - Offload directives used for the coprocessor

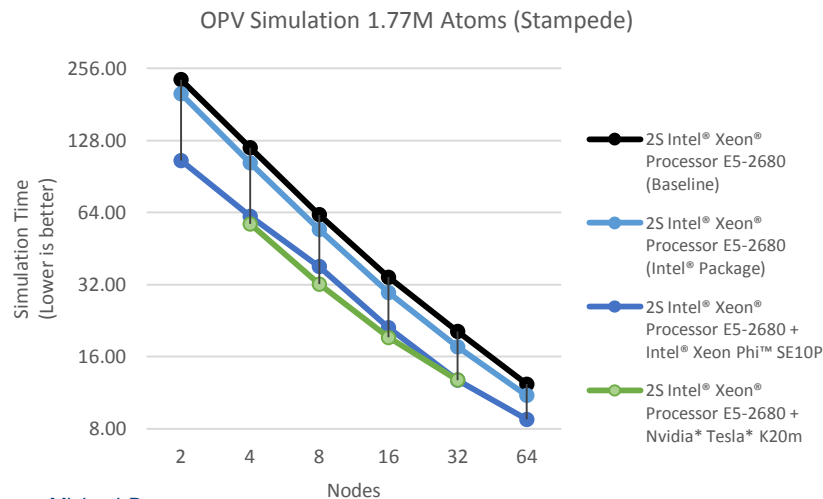
Organic Solar Cells

Science Team: Jan-Michael Y Carrillo, Rajeev Kumar, Monojoy Goswami, S. Michael Kilbey II, Bobby G Sumpter

(ORNL*/UT*)

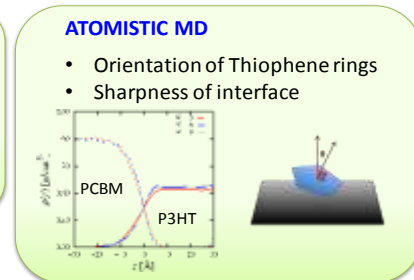

Problem: Predictive simulation of active layer morphology and molecular alignment based on blend composition

Innovation: Code Modernization/HW



Result: With code modernization and advanced HPC resources, we have been the first to perform simulations at scales that match experiment.

- Left: Up to 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...



Carrillo, J.-M. Y., Kumar, R., Goswami, G., Sumpter, B., Brown, W.M., **New Insights into Dynamics and Morphology of P3HT:PCBM Active Layers in Bulk Heterojunctions.** *Physical Chemistry Chemical Physics*. 2013. 15: p. 17873-17882.


Source: Michael Brown

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>

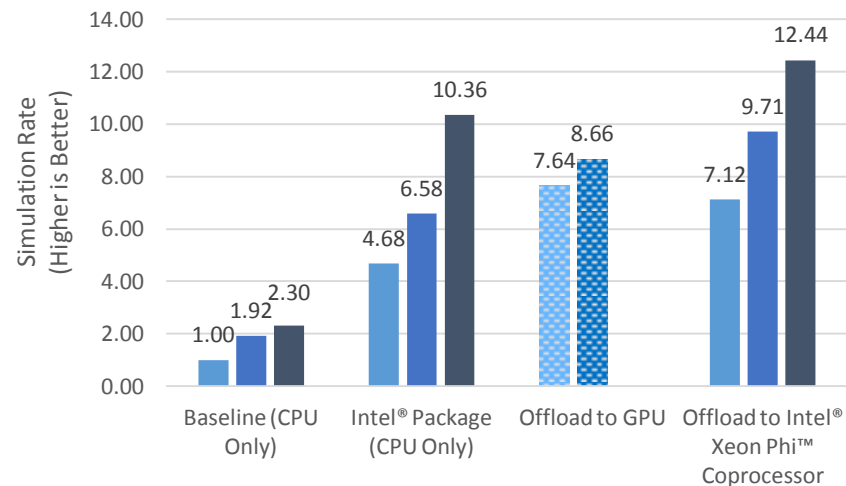
* Other names and brands may be claimed as the property of others.

LAMMPS performance with current processors

LAMMPS Rhodopsin Protein Benchmark 512K Atoms



LAMMPS Liquid Crystal Benchmark 524K Atoms



- 2S Intel® Xeon® Processor E5-2680 / Nvidia* Tesla* K20m / Intel® Xeon Phi™ Coprocessor SE10P†
- 2S Intel® Xeon® Processor E5-2697v2 / Nvidia* Tesla* K40c / Intel® Xeon Phi™ Coprocessor 7120A ‡
- 2S Intel® Xeon® Processor E5-2697v3 / HW Not Available / Intel® Xeon Phi™ Coprocessor 7120A‡

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>

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

Issues/Improvements

Available/Ongoing/Future Improvements

Compiler vectorization

- Significant advances in compiler vectorization using the Intel® Composer XE 2015.
- *New:* Vector variant functions that allow explicit vector coding for small subroutines
- Advanced compiler optimization reports with Intel® Composer XE 2015 and runtime analysis with Intel® VTune™ Amplifier XE

MPI Performance in symmetric modes

- Significant performance improvements with Intel® MPI* 5.
- Future many-core chips will have bootable options and also options for integrated fabric

2nd half '15
1st commercial systems

3+ TFLOPS¹
In One Package
Parallel Performance & Density

Designed using Intel's cutting-edge
14nm Transistor
Technology

Not bound by "offloading" bottlenecks
Standalone CPU or
PCIe Coprocessor

Common instruction set architecture
Intel® Advanced Vector
Extensions 512


Unveiling Details of Knights Landing

(Next Generation Intel® Xeon Phi™ Products)

Platform Memory: DDR4 Bandwidth and Capacity Comparable to Intel® Xeon® Processors

Compute: Energy-efficient IA cores²

- Microarchitecture enhanced for HPC³
- **3X Single Thread Performance** vs Knights Corner⁴
- Intel Xeon Processor Binary Compatible⁵



On-Package Memory:

- up to **16GB** at launch
 - **1/3X the Space**⁶
 - **5X Bandwidth** vs DDR4⁷
 - **5X Power Efficiency**⁶
- Jointly Developed with Micron Technology

All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. ¹Over 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. ²Modified version of Intel® Silvermont microarchitecture currently found in Intel® Atom™ processors. ³Modifications include AVX512 and 4 threads/core support. ⁴Projected peak theoretical single-thread performance relative to 1st Generation Intel® Xeon Phi™ Coprocessor 7120P (formerly codenamed Knights Corner). ⁵Binary Compatible with Intel Xeon processors using Haswell Instruction Set (except TSX). ⁶Projected results based on internal Intel analysis of Knights Landing memory vs Knights Corner (GDDR5). ⁷Projected 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.

Summary

Summary

We have demonstrated that important research for a variety of science problems continues to benefit from hardware and software advances in HPC.

We believe that the Intel® path forward for HPC architecture and software solutions is superior


- We are dedicated to providing solutions that allow for portable performance across a range of different processors for both conventional HPC centers and those on a path to exa-scale.

We have demonstrated, in a large production code, that optimizations for Intel® Xeon Phi™ coprocessors also improve performance on Intel® Xeon® processors and that the same routine can be used for efficient computations on both

Summary (2)

The innovations here also included reconsideration of the simulation model

- Important to realize that there is often a choice in the model used, and that in the past, minimizing computation was important to performance
- It may be the case that you can exploit additional arithmetic intensity and concurrency to increase the accuracy, time-step, etc. in order to advance research capabilities
- This has already been the case for simulation in materials problems



Computational Aspects of Many-body Potentials, S. J. Plimpton and A. P. Thompson, *MRS Bulletin**, 37, 513-521 (2012).

Acknowledgements

Research by the teams presented here made possible by:

DOE Early Science and ALCC Programs

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory*, which is supported by the Office of Science of the U.S. Department of Energy* under Contract No. DE-AC05-00OR22725.*

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 paper. 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*.*

Getting Access to Intel® Xeon Phi™ HPC Systems

NSF XSEDE*:

TACC* Stampede, LSU* SuperMIC

<https://www.xsede.org/allocations>

NSF UT* Beacon:

<https://www.nics.tennessee.edu/computing-resources/beacon/allocations>

DOE NERSC* Users (Babbage):

<https://www.nersc.gov/users/computational-systems/testbeds/babbage/>

Purdue* Conte (Purchase available):

<https://www.rcac.purdue.edu/compute/conte/>

Thanks!

