MIC & OpenMP 4

TCG Micro
SSG DPD
NERSC Threading Workshop, March 2015

High-performance Parallel Computing

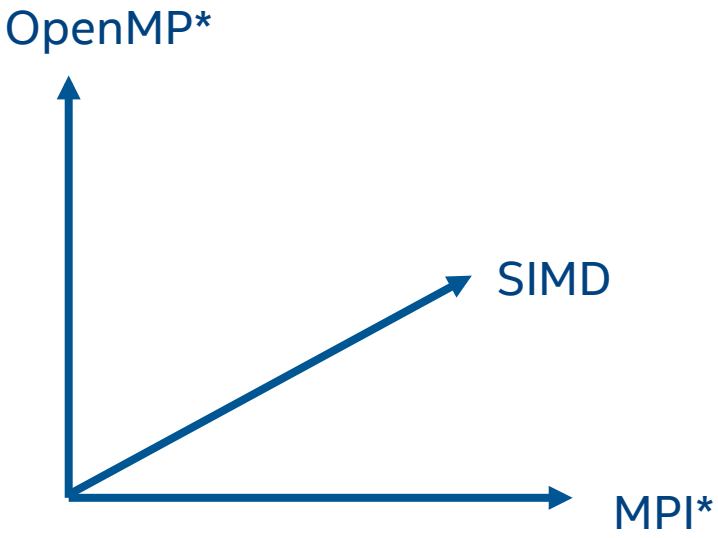

Moving data is expensive!

- Node-node
- Socket-Socket; Processor-(co)processor
- Core-core
- SIMD lanes

At each parallel level

- Find enough parallelism
- Decide the optimal granularity
- Optimize locality/data movement
- Ensure load balance
- Reduce the impact of coordination and synchronization

All the parallel units have to be coordinated with maximum overlap of data movement and computing.

Programming on 61 cores and 244 threads

	ES-2650 v2 (Ivy Bridge)	Xeon Phi 7120
Clock Rate GHz	2.40	1.24
# Cores/socket	12	61
# Threads/node (core)	48 (2)	244 (4)
# of sockets	2	1
Peak GFLOP/S	460.8	1210.24

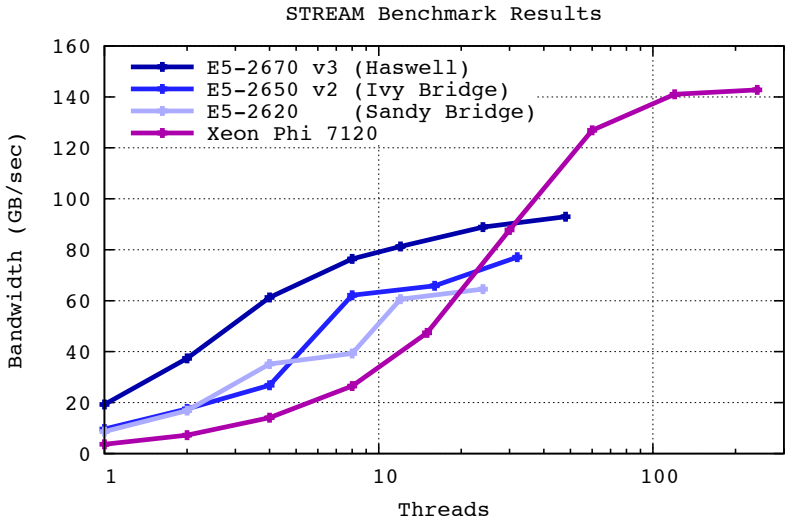

Basics of high-performance parallel computing is much more critical on Xeon™ Phi.

† STREAM bandwidth.

NERSC and community resources

Cori

https://www.nersc.gov/users/computational-systems/cori/

Trainings

https://www.nersc.gov/users/NUG/annual-meetings/nug-2015/hack-a-thon/

https://www.nersc.gov/users/training/events/OMP-vectorization-oct14/

The Intel Xeon Phi User's Group (IXPUG)

https://www.ixpug.org/what-ixpug

Resources

https://software.intel.com/en-us/mic-developer

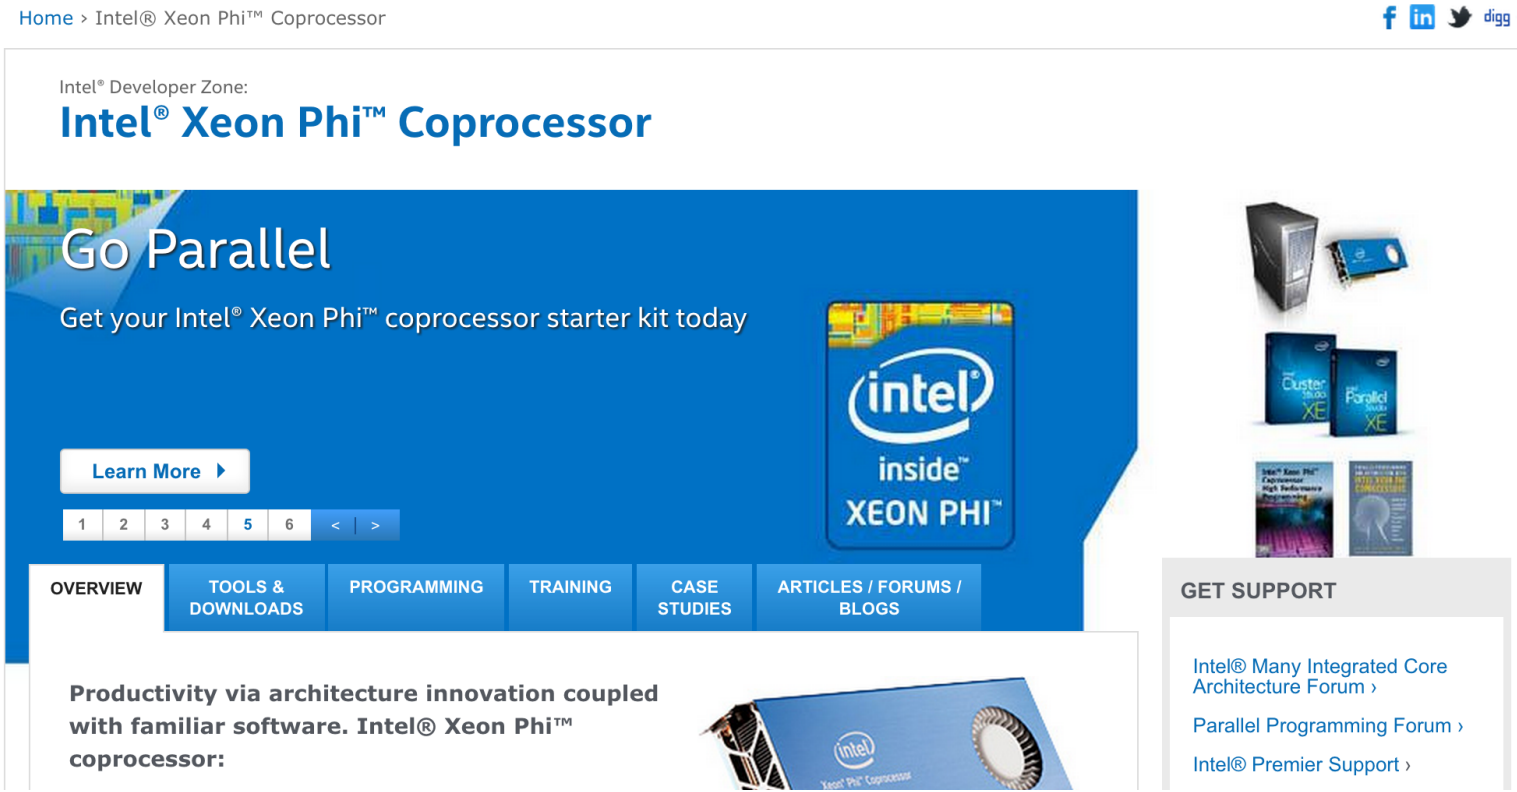

“I need a cookbook”

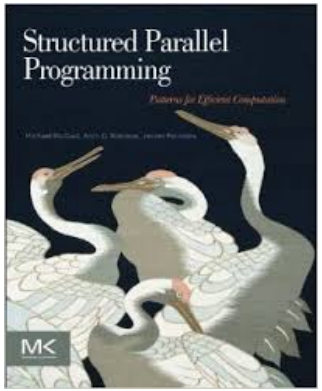

http://www.amazon.com/Structured-Parallel-Programming-Efficient-Computation/dp/0124159931

http://www.amazon.com/Intel-Xeon-Coprocessor-High-Performance-Programming/dp/0124104142

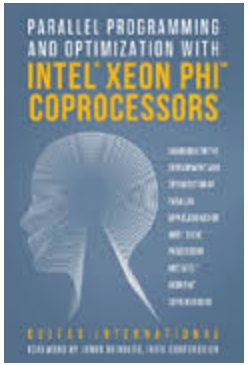

http://www.colfax-intl.com/nd/xeonphi.aspx

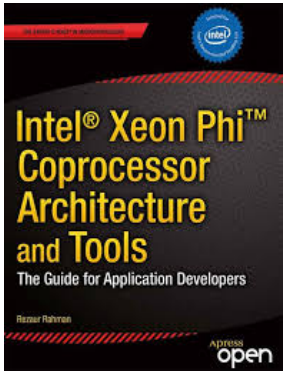

http://click.intel.com/intelr-xeon-phitm-coprocessor-architecture-and-tools-the-guide-for-application-developers.html

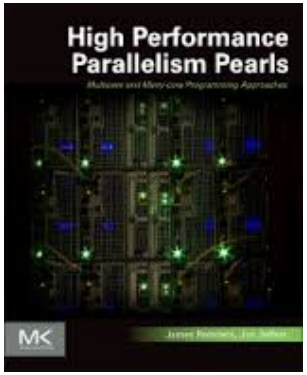

http://store.elsevier.com/High-Performance-Parallelism-Pearls/James-Reinders/isbn-9780128021187/

OpenMP*: Shared Memory Systems

(Basic) OpenMP How-to:

1. Identify loops with independent operations.
2. Insert compiler directives/pragmas.
3. Add thread synchronization and/or restructure code to avoid data races.

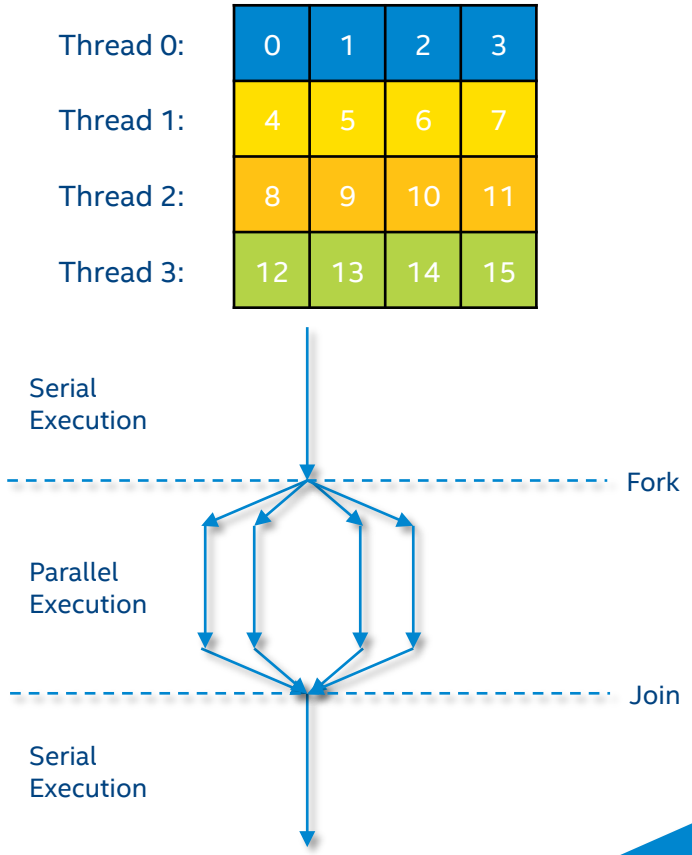

OpenMP*: Shared Memory Systems

Advantages

- Allows incremental development.
- Compiler-driven optimization.
- Support for dynamic load-balancing.

Disadvantages

- Implicit communication (“false sharing”).
- Data races are easy to introduce.
- Encourages “bolted on” parallelism.

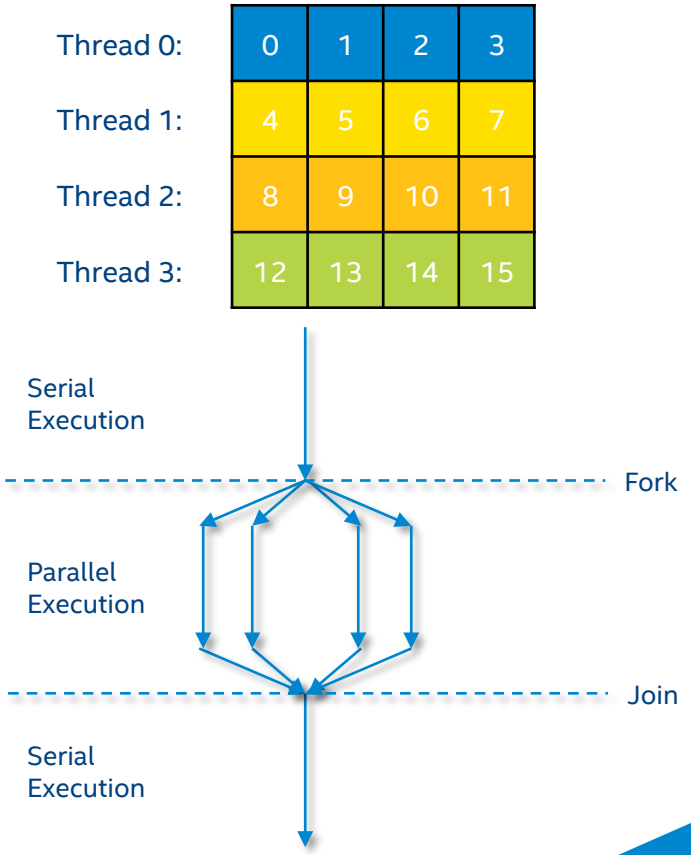

Intel® Xeon Phi™ Coprocessor: Programming Models

Native

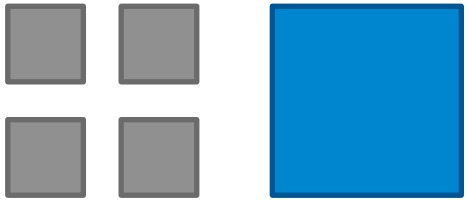

MPI*:
Ranks started on device.

OpenMP*:
Threads spawned per rank.

SIMD:
Vector loops run by threads.

Offload

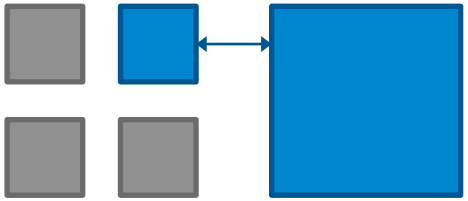

MPI:
Ranks started on host.

OpenMP:
Threads spawned in code sections offloaded to device.

SIMD:
Vector loops run by threads.

Symmetric

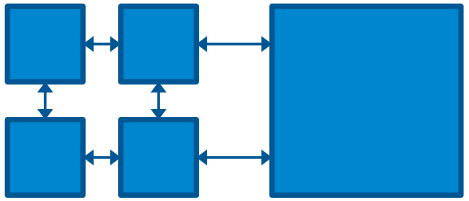

MPI:
Ranks started on host/device.

OpenMP:
Threads spawned per rank.

SIMD:
Vector loops run by threads.

Advanced Offload Use

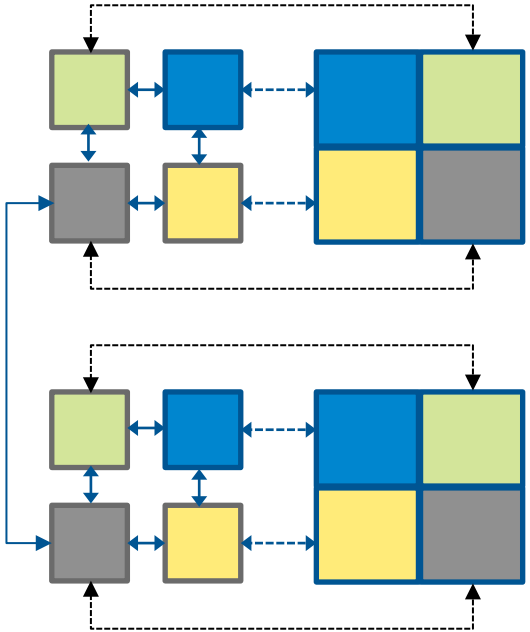

- Host employs MPI*
- Offload to a part of a device
- Asynchronous computations and communication on the host as usual

Same with host OpenMP* threads

Skipped MPI paths

OpenMP API

- De-facto standard, OpenMP 4.0 out since July 2013
- API for C/C++ and Fortran for shared-memory parallel programming
- Based on directives (pragmas in C/C++)
- Portable across vendors and platforms
- Supports various types of parallelism

Levels of Parallelism in OpenMP 4.0

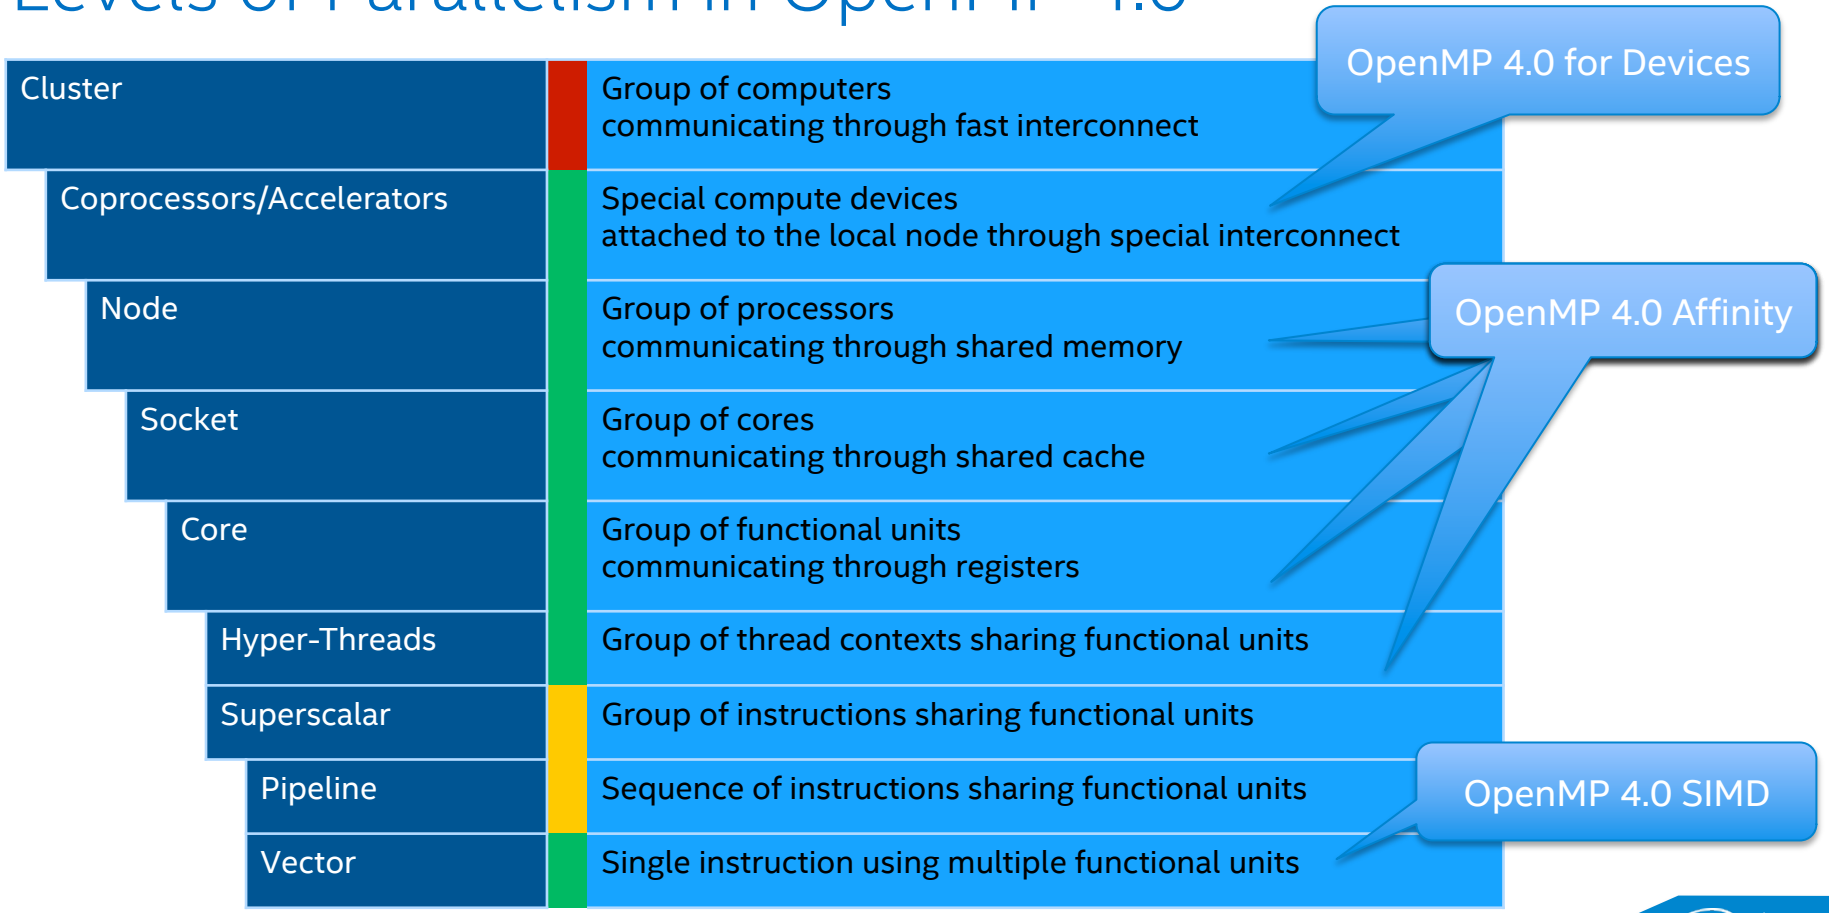

NUMA is here to Stay…

- (Almost) all multi-socket compute servers are NUMA systems
  - Different access latencies for different memory locations
  - Different bandwidth observed for different memory locations
- Example: Intel® Xeon E5-2600v2 Series processor

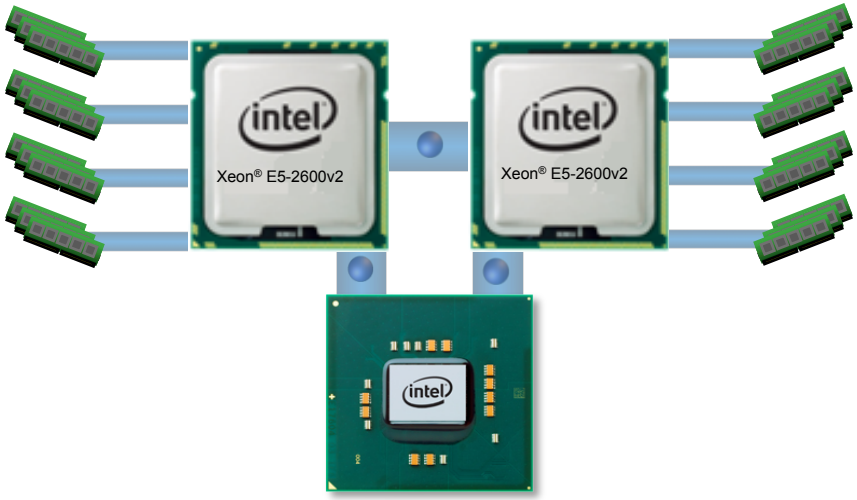

Thread Affinity – Why It Matters?

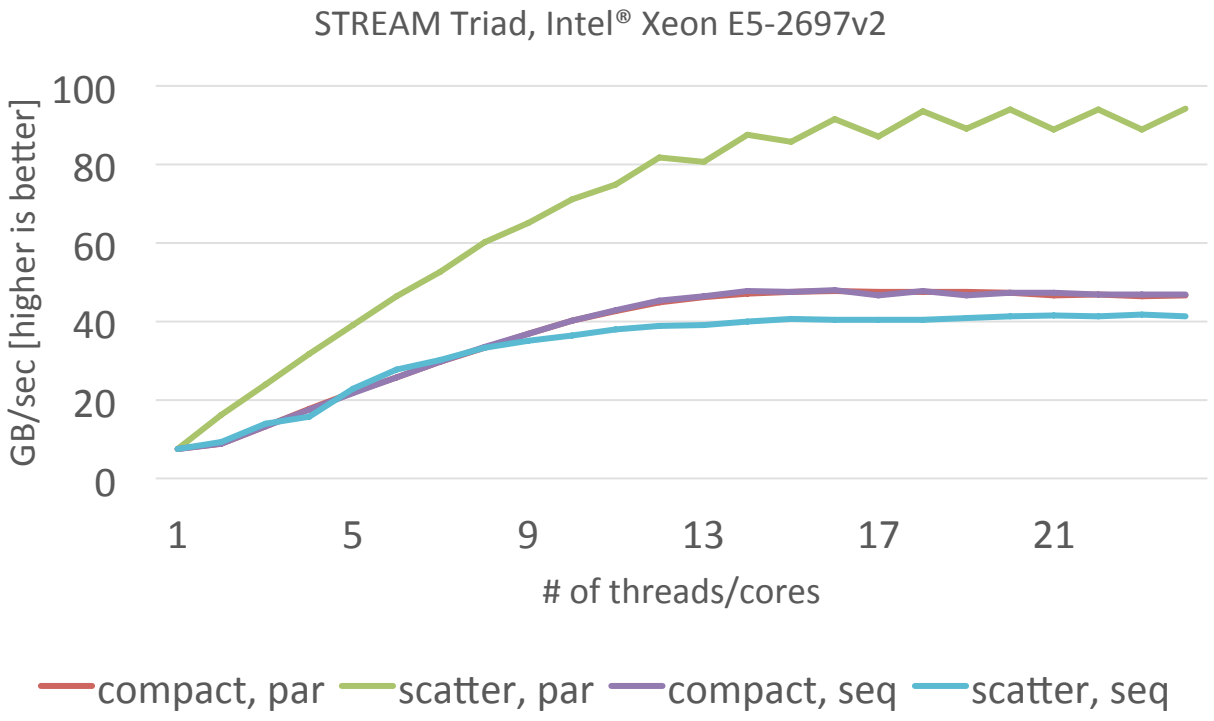

Thread Affinity – Processor Binding

- Binding strategies depends on machine and the app
- Putting threads far, e.g. on different packages
  - (May) improve the aggregated memory bandwidth
  - (May) improve the combined cache size
  - (May) decrease performance of synchronization constructs
- Putting threads close together, e.g. on a core sharing cache
  - (May) improve performance of synchronization constructs
  - (May) decrease the available memory bandwidth and cache size (per thread)
- Affinity is critical to MPI.

Thread Affinity in OpenMP* 4.0

OpenMP 4.0 introduces the concept of places…

- set of threads running on one or more processors
- can be defined by the user
- pre-defined places available:
  - threads: one place per hyper-thread
  - cores: one place exists per physical core
  - sockets: one place per processor package

… and affinity policies…

- spread: spread OpenMP threads evenly among the places
- close: pack OpenMP threads near master thread
- master: collocate OpenMP thread with master thread

… and means to control these settings

- Environment variables OMP_PLACES and OMP_PROC_BIND
- clause proc_bind for parallel regions

Thread Affinity Example

Example (Intel® Xeon Phi™ Coprocessor): Distribute outer region, keep inner regions close

```
OMP_PLACES=cores(8); OMP_NUM_THREADS=4,4;OMP_PROC_BIND=spread,close
```

Alternatively,

```
#pragma omp parallel proc_bind(spread)
  # pragma omp parallel proc_bind(close)
```

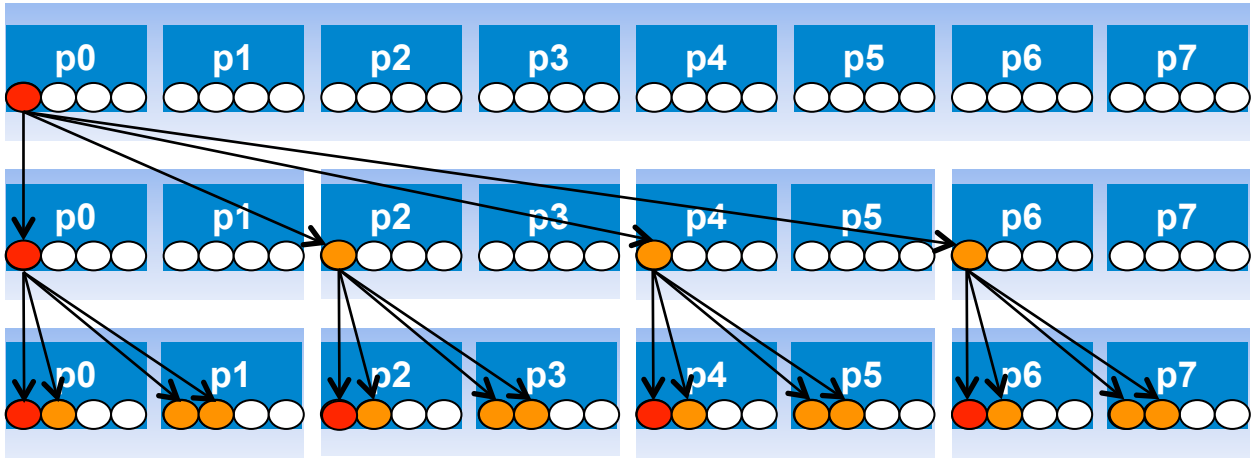

Vectorization: SIMD Machines

OpenMP* 4.0 SIMD How-to:

1. Identify loops with independent operations.
2. Insert compiler directives/pragmas.
3. Follow compiler's recommendations to enable/improve generated vector code.

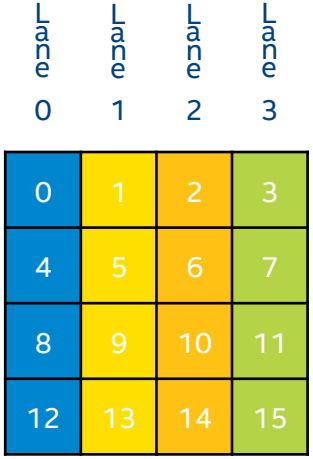

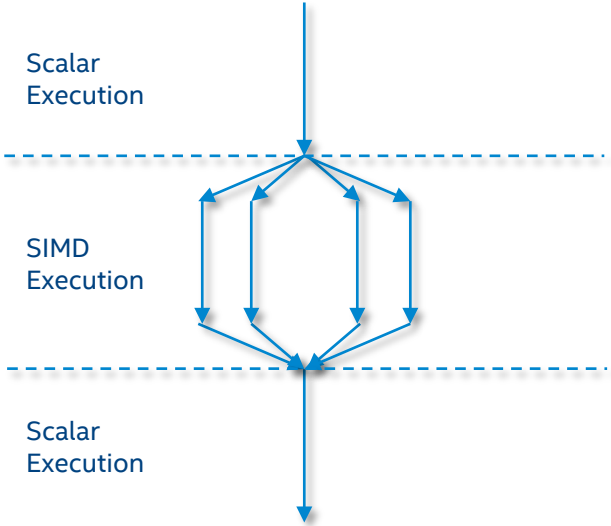

Vectorization: SIMD Machines

Advantages

- Allows incremental development.
- Compiler-driven optimization.
- Lanes are always synchronized.

Disadvantages

- Data structure/alignment requirements.
- Knowledge of ISA required to ensure efficient scalar to vector conversion.

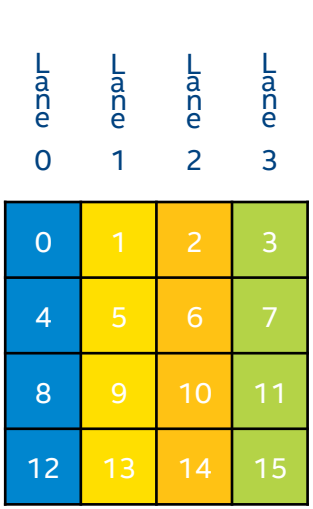

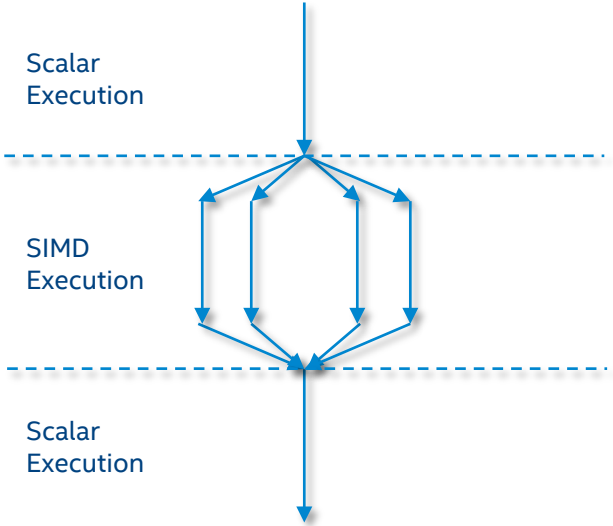

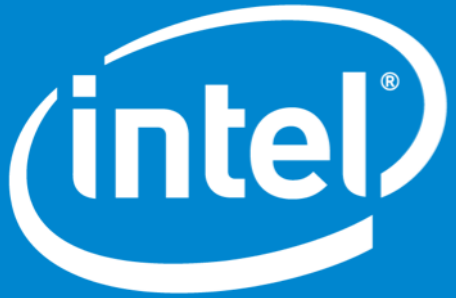

Legal Disclaimers

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

Knights Landing and other code names featured are used internally within Intel to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Intel to use code names in advertising, promotion or marketing of any product or services and any such use of Intel's internal code names is at the sole risk of the user

Intel, Look Inside, Xeon, Intel Xeon Phi, Pentium, Cilk, VTune and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

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

Legal Disclaimers

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

Legal Disclaimers

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. For more information go to http://www.intel.com/performance.

Intel® Advanced Vector Extensions (Intel® AVX)* provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at http://www.intel.com/go/turbo.

Estimated Results Benchmark Disclaimer:
Results have been estimated based on internal Intel analysis and are provided for informational purposes only. Any difference in system hardware or software design or configuration may affect actual performance.

Software Source Code Disclaimer:
Any software source code reprinted in this document is furnished under a software license and may only be used or copied in accordance with the terms of that license.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Legal Disclaimers

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.

Rev. 7/17/13