

HPCG on Intel Architecture Update

Nov 2015

Jongsoo Park, Alexander Kleymenov,
Mikhail Smelyanskiy, and Vadim Pirogov

Intel Corporation

SC15, HPCG BoF

Outline

- **IA result updates**
- Our other work related to HPCG
- HPCG 3.0 optimizations

Single-Node Results in IA

	Single-node Perf. (GFLOP/s)	Efficiency w.r.t. STREAM
2-skt 14c Xeon E5-2697 v3 @ 2.6 GHz (HSW) ¹⁾	18.4	104%
61c Xeon Phi 7120 @ 1.24 GHz (KNC) ²⁾	21.2	72%

- 1) Results from JHPC paper. STREAM BW 108 GB/s is used for BW efficiency. 1 MPI rank per socket, 14 threads per rank (1 thread per core), KMP_AFFINITY=granularity=fine,compact,1, lexicographical GS at all MG levels, 50 CG iterations
- 2) Results from JHPC paper. STREAM BW 170 GB/s is used for BW efficiency, 12 MPI ranks, 20 threads per rank (4 threads per core), KMP_AFFINITY=granularity=fine,compact, lexicographical GS at top 2 MG levels, multi-color GS at the other levels. 51 CG iterations

Single-Node Results in IA

	Single-node Perf. (GFLOP/s)	Efficiency w.r.t. STREAM
2-skt 14c Xeon E5-2697 v3 @ 2.6 GHz (HSW) ¹⁾	18.4	104%
61c Xeon Phi 7120 @ 1.24 GHz (KNC) ²⁾	21.2	72%

1) Results from IHPC paper. STREAM BW 108 GB/s is used for BW efficiency. 1 MPI rank per socket. 14 threads

Newer Xeon processors have higher efficiency w.r.t. STREAM BW.

color 0.5 at the other levels. 512 iterations

Multi-Node Results in IA

	HPCG perf. (TFLOPS)	Per-node HPCG perf. (GFLOPS)	Efficiency w.r.t. STREAM
TH2 ¹⁾	580	37.8	50%
Shaheen II ²⁾	114	20.5	112%
2/3 of Stampde ³⁾	99	19.9	48%

- 1) Results from ISC14. STREAM BW 150 GB/s per Intel® Xeon Phi™ coprocessors. Intel® Xeon® processors were not used for computation, and their STREAM BW was excluded when computing BW efficiency.
- 2) STREAM BW 110 GB/s per node is used for BW efficiency..
- 3) Both Xeon processors and Xeon Phi coprocessors are used. 80 GB/s Xeon processor STREAM BW and 170 GB/s Xeon Phi coprocessor STREAM BW is used for BW efficiency

Multi-Node Results in IA

	HPCG perf. (TFLOPS)	Per-node HPCG perf. (GFLOPS)	Efficiency w.r.t. STREAM
TH2 ¹⁾	580	37.8	50%
Shaheen II ²⁾	114	20.5	112%
2/3 of Stampede ³⁾	99	19.9	48%

Stampede result is from only 2/3 of the full cluster. Stampede has similar per node performance to GPU clusters.

Outline

- IA result updates
- **Our other work related to HPCG**
- HPCG 3.0 optimizations

Work related to HPCG (1)

MKL 11.3 inspector-executor sparse BLAS routines

- Sparse matrix-vector multiplication (SpMV)
- Sparse triangular solver (SpTS)
- Sparse matrix-matrix multiplication (SpGEMM)
- ...

<https://software.intel.com/en-us/articles/intel-math-kernel-library-inspector-executor-sparse-blas-routines>

Work related to HPCG (2)

SpMP: Open-source library for optimized sparse matrix pre-processing

- <https://github.com/jspark1105/SpMP>
- Task dependency graph construction, BFS/RCM reordering, ...
- Can be useful for libraries implementing their own SpTS, ILU, ...
- Helped a customer wanted optimized ILU preconditioners

Work related to HPCG (3)

Optimizing algebraic multi-grid implementation in HYPRE library

- High-Performance Algebraic Multigrid Solver Optimized for Multi-Core Based Distributed Parallel Systems, Park et al., SC15
- Common optimizations: efficient lexicographical GS, matrix reordering for cache locality optimization, ...

Outline

- IA result updates
- Our other work related to HPCG
- **HPCG 3.0 optimizations**

HPCG 3.0 Optimizations in newly timed routines, `GenerateProblem` and `SetupHalo`

Separate interior and boundary points

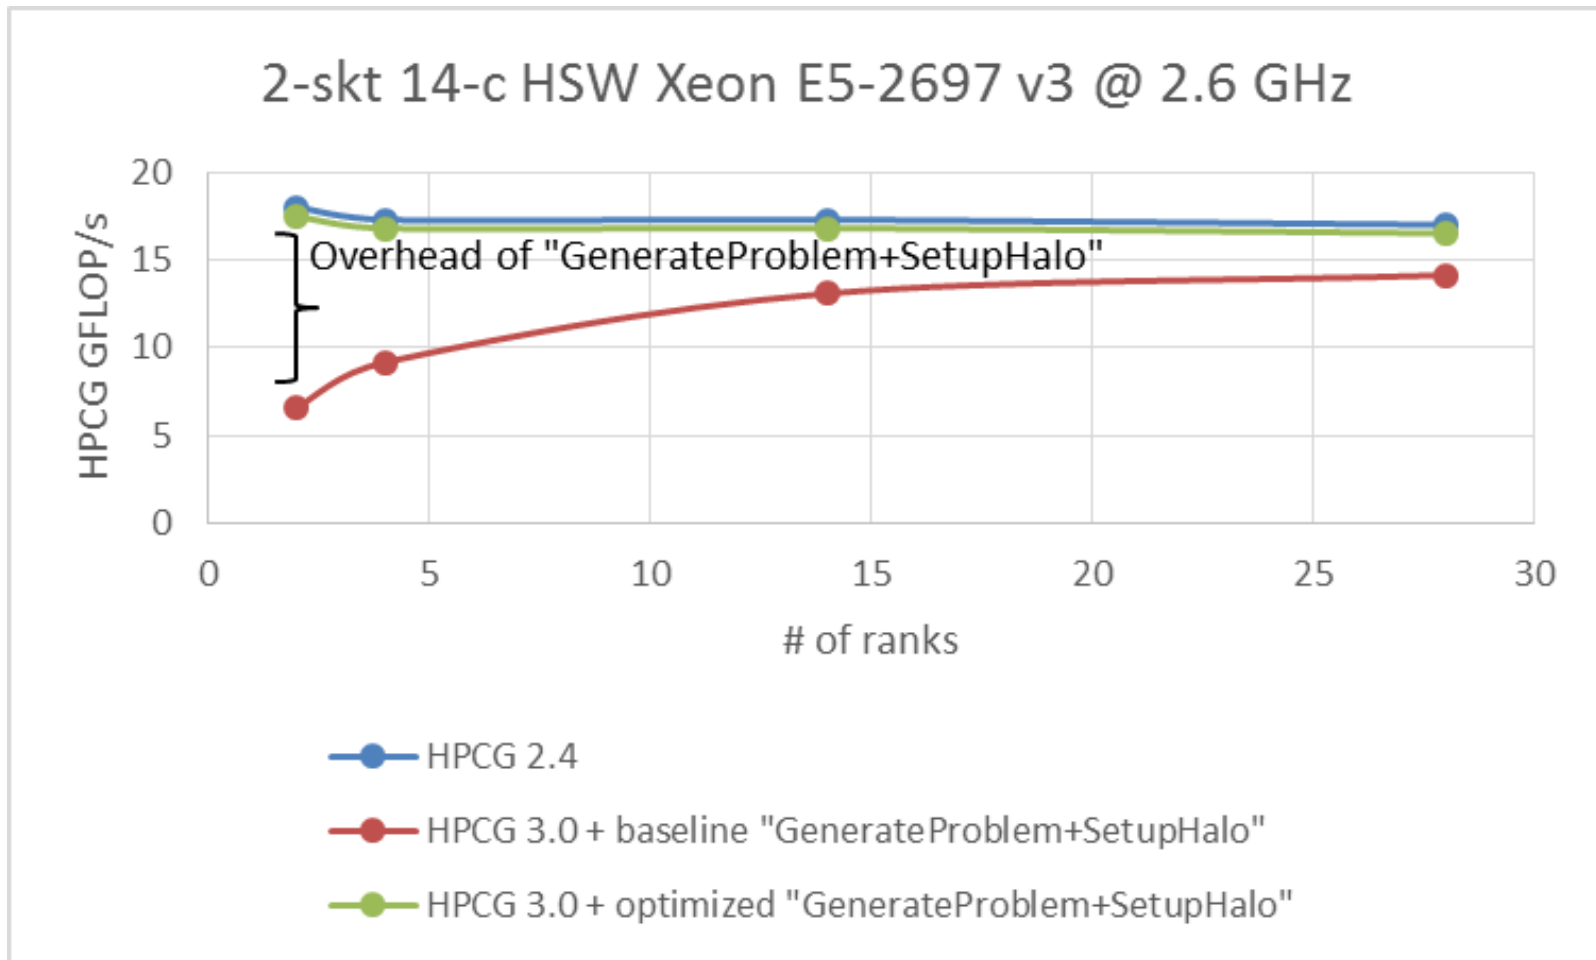
- Handling interior points (common case) is easily parallelized

Halo setup with boundary points

- **Thread-private `std::unordered_map`**
(most duplicates eliminated here due to locality)
- Concatenate thread private maps into an array
- `std::sort` and `std::unique`


Similar patterns in HYPRE AMG library

HPCG 3.0 on HSW




Setups are the same as in "Optimizations in High-Performance Conjugate Gradient Benchmark for IA-based Multi and Many-core Processors" in JHPC, Table II used for single node experiments.

HPCG 3.0 on HSW



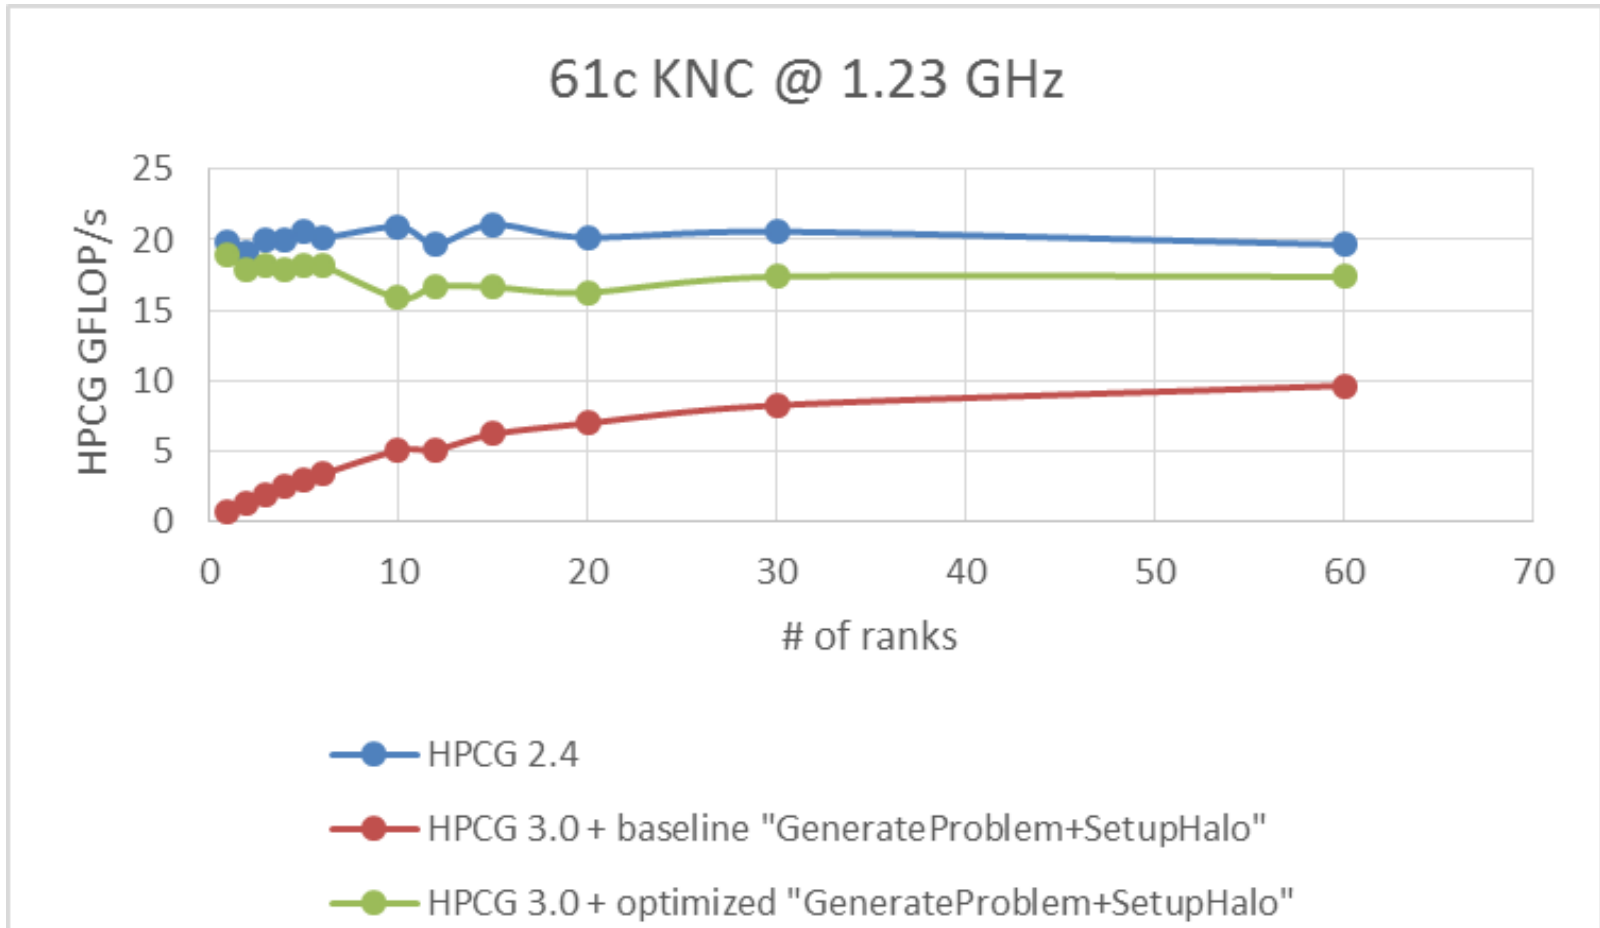
Overhead reduced to <math><3\%</math>

HPCG 3.0 on Knights Corner



Setups are the same as in "Optimizations in High-Performance Conjugate Gradient Benchmark for IA-based Multi and Many-core Processors" in JHPC, Table II used for single node experiments.

HPCG 3.0 on Knights Corner



Overhead reduced to $<4\%$

Baseline: more ranks \rightarrow smaller overhead (no omp par.)

Optimized: more ranks \rightarrow more overhead (more boundaries)

Notice 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, or life sustaining applications. Intel may make changes to specifications, product descriptions, and plans at any time, without notice.

All products, dates, and figures are preliminary for planning purposes and are subject to change 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.

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.

The Intel products discussed herein 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.

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 by visiting Intel's website at <http://www.intel.com>.

Intel® Itanium®, Intel® Xeon®, Xeon Phi™, Pentium®, Intel SpeedStep® and Intel NetBurst® , Intel®, and VTune are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Copyright © 2014, Intel Corporation. All rights reserved.

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

Notice and Disclaimers Continued ...

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'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

Q&A