HPCG on
Intel Xeon Phi 2nd Generation, Knights Landing

Alexander Kleymenov and Jongsoo Park

Intel Corporation

SC16, HPCG BoF

Outline

- **KNL results**
- Our other work related to HPCG

November 2016 HPCG Results

Rank	Site	Computer	Cores	HPL Rmax (Pflop/s)	TOP500 Rank	HPCG (Pflop/s)	Fraction of Peak
1	RIKEN Advanced Institute for Computational Science Japan	**K computer** – , SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705,024	10.510	7	0.6027	5.3%
2	NSCC / Guangzhou China	**Tianhe-2 (MilkyWay-2)** – TH-IVB-FEP Cluster, Intel Xeon 12C 2.2GHz, TH Express 2, Intel Xeon Phi 31S1P 57-core NUDT	3,120,000	33.863	2	0.5800	1.1%
3	Joint Center for Advanced High Performance Computing Japan	**Oakforest-PACS** – PRIMERGY CX600 M1, Intel Xeon Phi Processor 7250 68C 1.4GHz, Intel Omni-Path Architecture Fujitsu	557,056	13.555	6	0.3855	1.5%
4	National Supercomputing Center in Wuxi China	**Sunway TaihuLight** – Sunway MPP, SW26010 260C 1.45GHz, Sunway NRCPC	10,649,600	93.015	1	0.3712	0.3%
5	DOE/SC/LBNL/NERSC USA	**Cori** – XC40, Intel Xeon Phi 7250 68C 1.4GHz, Cray Aries Cray	632,400	13.832	5	0.3554	1.3%
6	DOE/NNSA/LLNL USA	**Sequoia** – IBM BlueGene/Q, PowerPC A2 1.6 GHz 16-core, 5D Torus IBM	1,572,864	17.173	4	0.3304	1.6%
7	DOE/SC/Oak Ridge Nat Lab USA	**Titan** – Cray XK7, Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray	560,640	17.590	3	0.3223	1.2%
8	DOE/NNSA/LANL/SNL USA	**Trinity** – Cray XC40, Intel Xeon E5-2698-V3, Aries custom Cray	301,056	8.101	10	0.1826	1.6%
9	NASA / Mountain View USA	**Pleiades** – SGI ICE X, Intel Xeon E5-2670, E5-2680V2, E5-2680V3, E5-2680V4, Infiniband FDR HPE/SGI	243,008	5.952	13	0.1752	2.5%
10	DOE/SC/Argonne National Laboratory USA	**Mira** – IBM BlueGene/Q, PowerPC A2 1.6 GHz 16-core, 5D Torus IBM	786,432	8.587	9	0.1670	1.7%

~47 GF/s per KNL

~10 GF/s per HSW

Single-Node KNL

	Perf. (GFLOP/s)
72c Xeon Phi 7290	51.3 (flat mode)
68c Xeon Phi 7250	49.4 (flat mode), 13.8 (DDR)
64c Xeon Phi 7210	46.7 (flat mode)

Cache mode provides a similar performance (~3% drop)

MCDRAM provides >3.5x performance than DDR

Easier to use than KNC

- Less reliance on software prefetching
- 2 threads per core is enough to get the best performance
- Smaller gap between SpMV using CSR and SELLPACK for U of Florida Matrix Collection

n=192 usually gives the best results. All results are measured with quad cluster mode and code from https://software.intel.com/en-us/articles/intel-mkl-benchmarks-suite

Multi-Node KNL

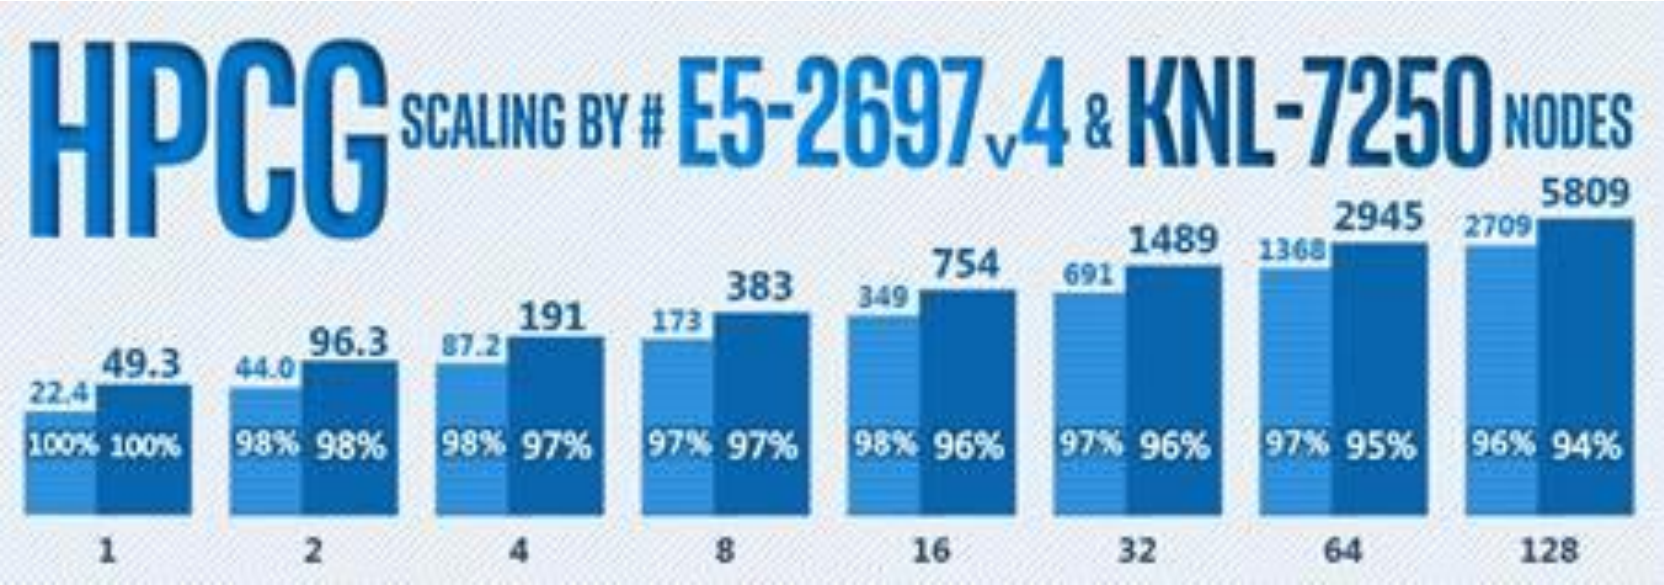

Each node in flat/quad mode connected with Omni-Path fabric (OPA)

Outline

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

Related Work (1) – Library

MKL inspector-executor sparse BLAS routines

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

SpMP open source library (https://github.com/jspark1105/SpMP)

BFS/RCM reordering, task graph construction of SpTrSv and ILU, …

Optimizing AMG in HYPRE library

Included from HYPRE 2.11.0

Related Work (2) – Compiler

Automating Wavefront Parallelization for Sparse Matrix Computations, Venkat et al., SC'16

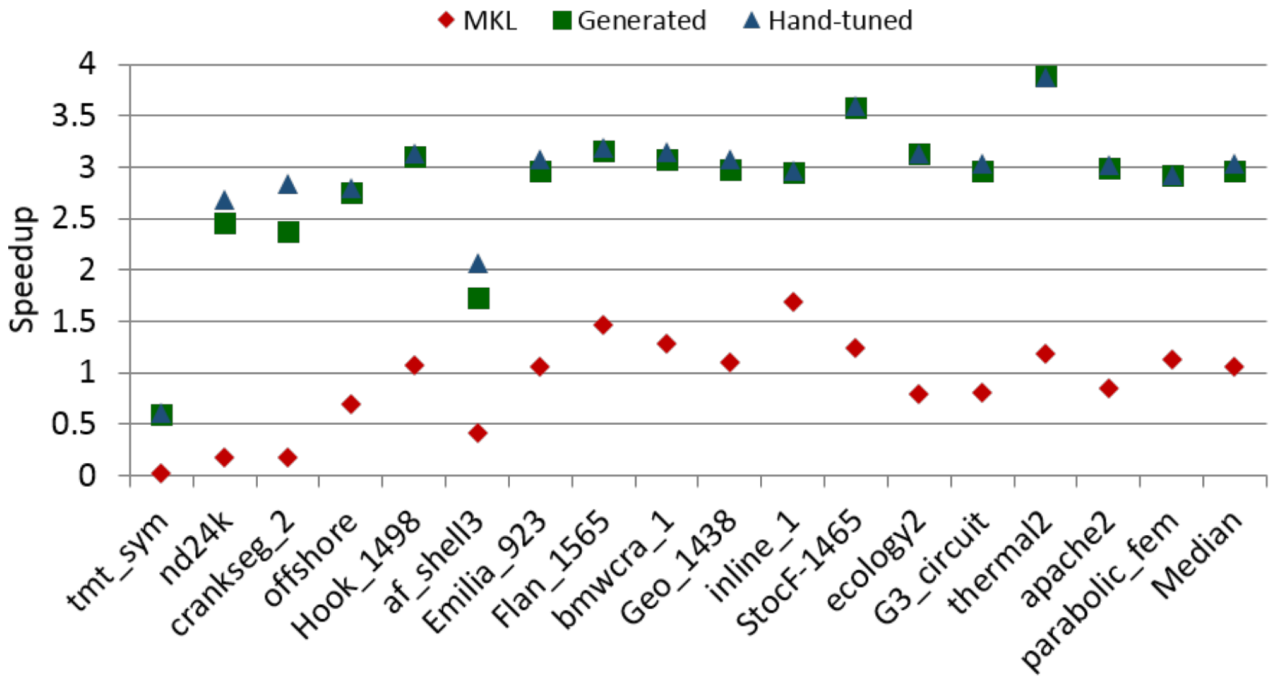

Fig. 8. Speedup of Parallel PCG over Sequential PCG.

12-core Xeon E5-2695 v2, ILU0 pre-conditioner, speedups include inspection overhead time

Related Work (3) – Script Language

Sparso: Context-driven Optimizations of Sparse Linear Algebra, Rong et al., PACT'16,
https://github.com/IntelLabs/Sparso

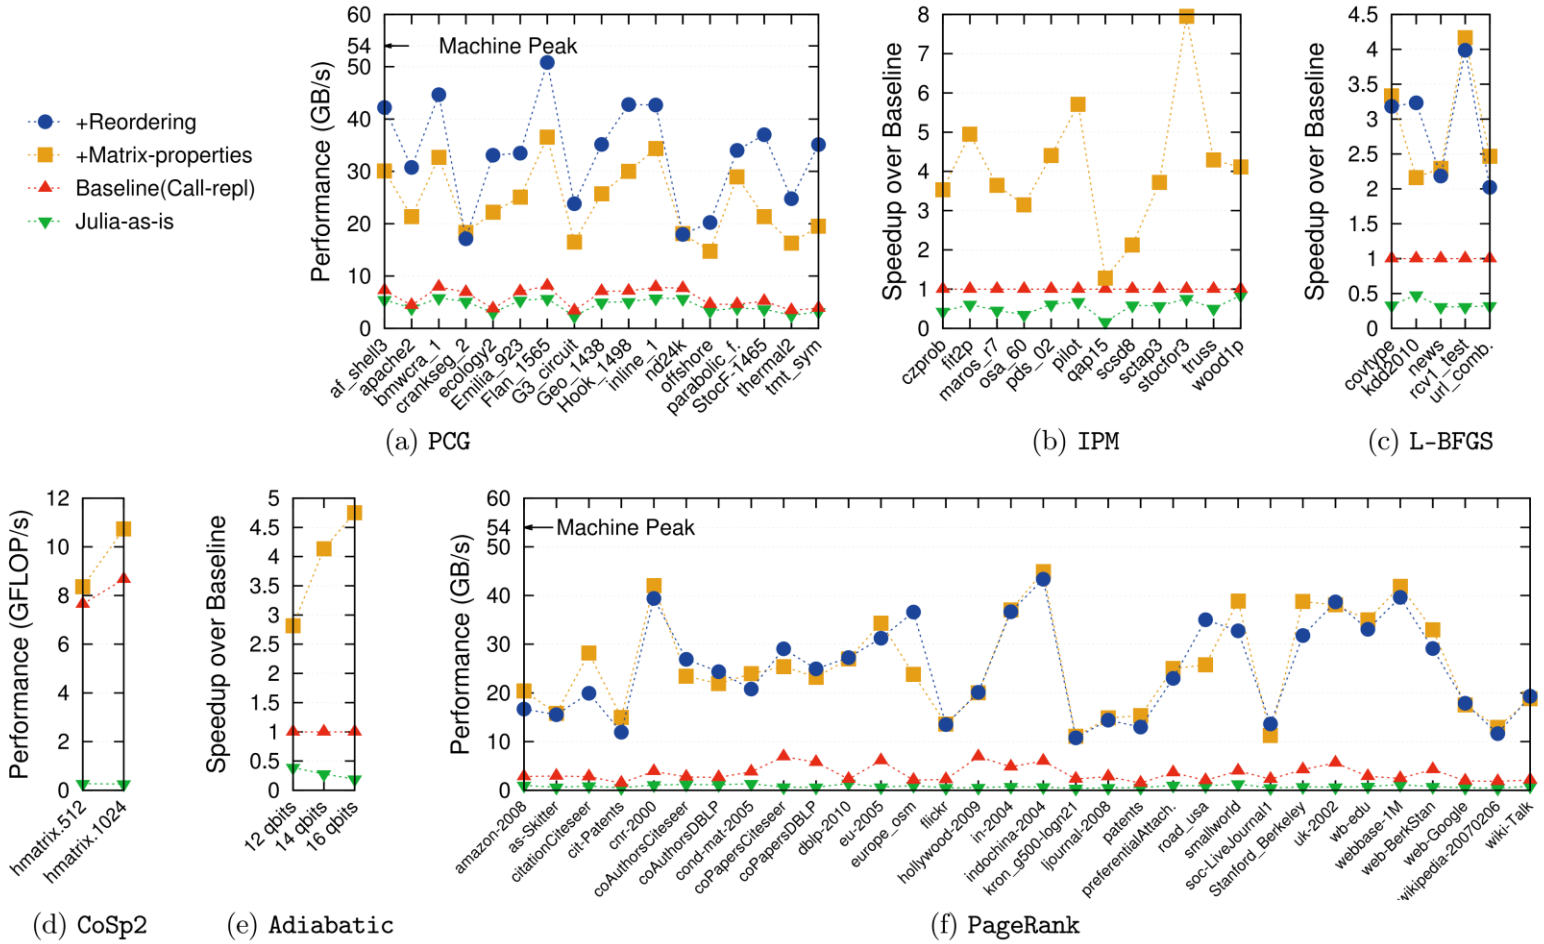

14-core Xeon E5-2697 v3,
Julia with Sparso package

Notice and Disclaimers

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