www.hpcg-benchmark.org

HPCG UPDATE: SC'15

Jack Dongarra
Michael Heroux
Piotr Luszczek

HPCG Snapshot

- High Performance Conjugate Gradient (HPCG).
- Solves $Ax=b$, A large, sparse, b known, x computed.
- An optimized implementation of PCG contains essential computational and communication patterns that are prevalent in a variety of methods for discretization and numerical solution of PDEs

- Patterns:
  - Dense and sparse computations.
  - Dense and sparse collective.
  - Multi-scale execution of kernels via MG (truncated) V cycle.
  - Data-driven parallelism (unstructured sparse triangular solves).
- Strong verification (via spectral properties of PCG).

Model Problem Description

- Synthetic discretized 3D PDE (FEM, FVM, FDM).
- Zero Dirichlet BCs, Synthetic RHS s.t. solution = 1.
- Local domain: $(n_x \times n_y \times n_z)$
- Process layout: $(np_x \times np_y \times np_z)$
- Global domain: $(n_x * np_x) \times (n_y * np_y) \times (n_z * np_z)$
- Sparse matrix:
  - 27 nonzeros/row interior.
  - 8 – 18 on boundary.
  - Symmetric positive definite.

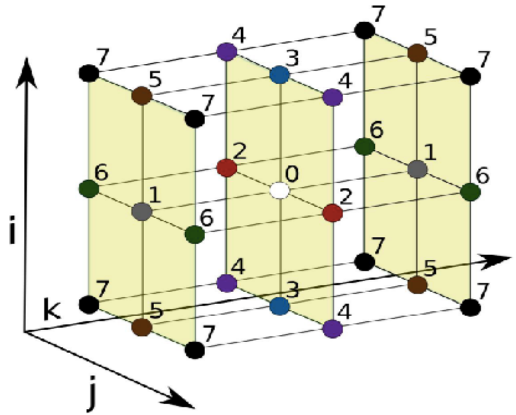

27-point stencil operator

Merits of HPCG

- Includes major communication/computational patterns.
  - Represents a minimal collection of the major patterns.
- Rewards investment in:
  - High-performance collective ops.
  - Local memory system performance.
  - Low latency cooperative threading.
- Detects/measures variances from bitwise reproducibility.
- Executes kernels at several (tunable) granularities:
  - nx = ny = nz = 104 gives
  - nlocal = 1,124,864; 140,608; 17,576; 2,197
  - ComputeSymGS with multicoloring adds one more level:
    - 8 colors.
    - Average size of color = 275.
    - Size ratio (largest:smallest): 4096
  - Provide a “natural” incentive to run a big problem.

HPL vs. HPCG: Bookends

- Some see HPL and HPCG as “bookends” of a spectrum.
  - Applications teams know where their codes lie on the spectrum.
  - Can gauge performance on a system using both HPL and HPCG numbers.

HPCG Status

HPCG 3.0 Release, Nov 11, 2015

- Available on GitHub.com
  - Using GitHub issues, pull requests, Wiki.
- Intel, Nvidia optimized 3.0 version available.
- IBM has 2.4 optimized version.
- For ISC'15, HPCG 3.0 any new results should be obtained using 3.0 unless not possible.
- Quick Path option will make this easier.

Main HPCG 3.0 Features

See http://www.hpcg-benchmark.org/software/index.html for full discussion

- Problem generation is timed.
- Memory usage counting and reporting.
- Memory bandwidth measurement and reporting
- "Quick Path" option to make obtaining results on production systems easier.
- Provides 2.4 rating and 3.0 rating in output.
- Command line option (--rt=) to specify the run time.

Other Items

- Reference version on GitHub:
  - https://github.com/hpcg-benchmark/hpcg
  - Website: hpcg-benchark.org.
  - Mail list hpcg.benchmark@gmail.com
- HPCG used in SC15 Student Cluster Competition.
- HPCG-optimized kernels going into vendor libraries.
- Next event: ISC’16:
  - 63 entries so far (42 – ISC15, 25 – SC14, 15 – ISC14)
  - Quick Path option should accelerate adoption.

Special Issue: International Journal of High Performance Computer Applications

1. Reference HPCG.
2. Intel.
3. Nvidia.
4. NUDT.
5. Riken.
6. Coming a little later: IBM.

- Discussion and results from vendor optimizations.
- Articles in final review.
- Some highlights…

Rewards investment high performance collectives.

"Edison spends only 1.9% of the total time in all-reduce while SuperMUC, Occigen, and Stampede spend 12.9%,5.9%, and 22.0%, respectively. We believe this difference primarily comes from that Edison uses a low-diameter high-radix Aries network with Dragonfly topology."

Intel HPCG Paper

Collectives futures

- “Addressing the bottleneck in collective communications will be also an important challenge as the collectives are shown to often take well above 10% of the total time. Even though high-radix Dragonfly topology considerably speedups the collectives, we envision that continued innovation in network infrastructure will be necessary due to ever increasing concurrency in high performance computing systems.”

Impact broader set of computations

"The optimizations described in this paper are not limited to the HPCG benchmark and can be also applicable to other problems and sparse solvers as exemplified by our evaluation with unstructured matrices shown in [our previous report]."

Looking toward next generation memories

“We expect challenges and opportunities laid out for HPCG in the next few years. One of the significant challenges will be effective use of emerging memory technologies and the accompanied diversification of memory hierarchy.”

Detecting FP Variations (Reproducibility)

Residual=4.25079640861055785883e-08 (0x1.6d240066fda73p-25)
Residual=4.25079640861032293954e-08 (0x1.6d240066fd910p-25)
Residual=4.25079640861079079289e-08 (0x1.6d240066fdbd3p-25)
Residual=4.25079640861054528568e-08 (0x1.6d240066fda60p-25)
Residual=4.25079640861068491377e-08 (0x1.6d240066fdb33p-25)
Residual=4.25079640861059094605e-08 (0x1.6d240066fdaa5p-25)

"The code correctly identified small variations in the residuals, caused by the network off-loading collectives. There is a small improvement in performance but the off-loading collectives introduce a small non-reproducibility."

Vendor improvement: Intel 4X

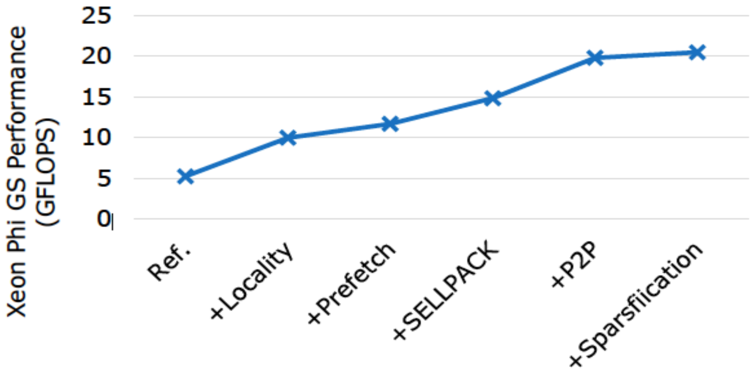

Fig. 5: The impact of optimizations on the Xeon Phi performance of SymGS parallelized with task scheduling.

- `Ref.`: the reference implementation ran with 240 MPI ranks
- `+Locality`: storage layout optimization for locality (Section IV-A1)
- `+Prefetch`: software prefetches
- `+SELLPACK`: vectorization-friendly matrix storage format [43]
- `+P2P`: point-to-point synchronization instead of barriers
- `+Sparsification`: eliminating unnecessary synchronization [10]

Summary

- HPCG is
  - Addressing original goals.
  - Rewarding vendor investment in features we care about.
- HPCG has traction.
  - Original goal of top 50 systems seems reachable, and more.
- Biggest challenge (my bias):
  - Pre-mature conclusions based on incomplete analysis of reference version.
  - IJHPCA papers should dispel these concerns.
- Version 3.X is the final planned major version.
- HPL and HPCG make a nice set of bookends.
  - Anyone got a (wood) router?

HPCG RANKINGS NOVEMBER 2015

And The Winners Are…

HPCG

PRESENTED AT SC15 Austin, TX | hpc transforms. NOVEMBER 18, 2015

NUMBER 3

SYSTEM **Titan**
DOE
Oak Ridge National Laboratory
USA

ACHIEVED **0.322** Pflop/s

JACK DONGARRA

MICHAEL HEROUX

PIOTR LUSZCZEK

IN COLLABORATION WITH ICL INNOVATIVE COMPUTING LABORATORY, Sandia National Laboratories

SPONSORED BY DEPARTMENT OF ENERGY · UNITED STATES OF AMERICA

HPCG

PRESENTED AT SC15 Austin, TX | hpc transforms. NOVEMBER 18, 2015

NUMBER 2

SYSTEM **K computer**
RIKEN Advanced Institute
for Computational Science
JAPAN

ACHIEVED **0.461** Pflop/s

JACK DONGARRA

MICHAEL HEROUX

PIOTR LUSZCZEK

IN COLLABORATION WITH ICL INNOVATIVE COMPUTING LABORATORY

Sandia National Laboratories

SPONSORED BY DEPARTMENT OF ENERGY UNITED STATES OF AMERICA

HPCG

PRESENTED AT SC15 Austin, TX | hpc transforms. NOVEMBER 18, 2015

NUMBER 1

SYSTEM **Tianhe-2**
National Super Computer Center in Guangzhou
CHINA

ACHIEVED **0.580** Pflop/s

JACK DONGARRA

MICHAEL HEROUX

PIOTR LUSZCZEK

IN COLLABORATION WITH ICL INNOVATIVE COMPUTING LABORATORY

Sandia National Laboratories

SPONSORED BY DEPARTMENT OF ENERGY UNITED STATES OF AMERICA

Comparison Peak, HPL

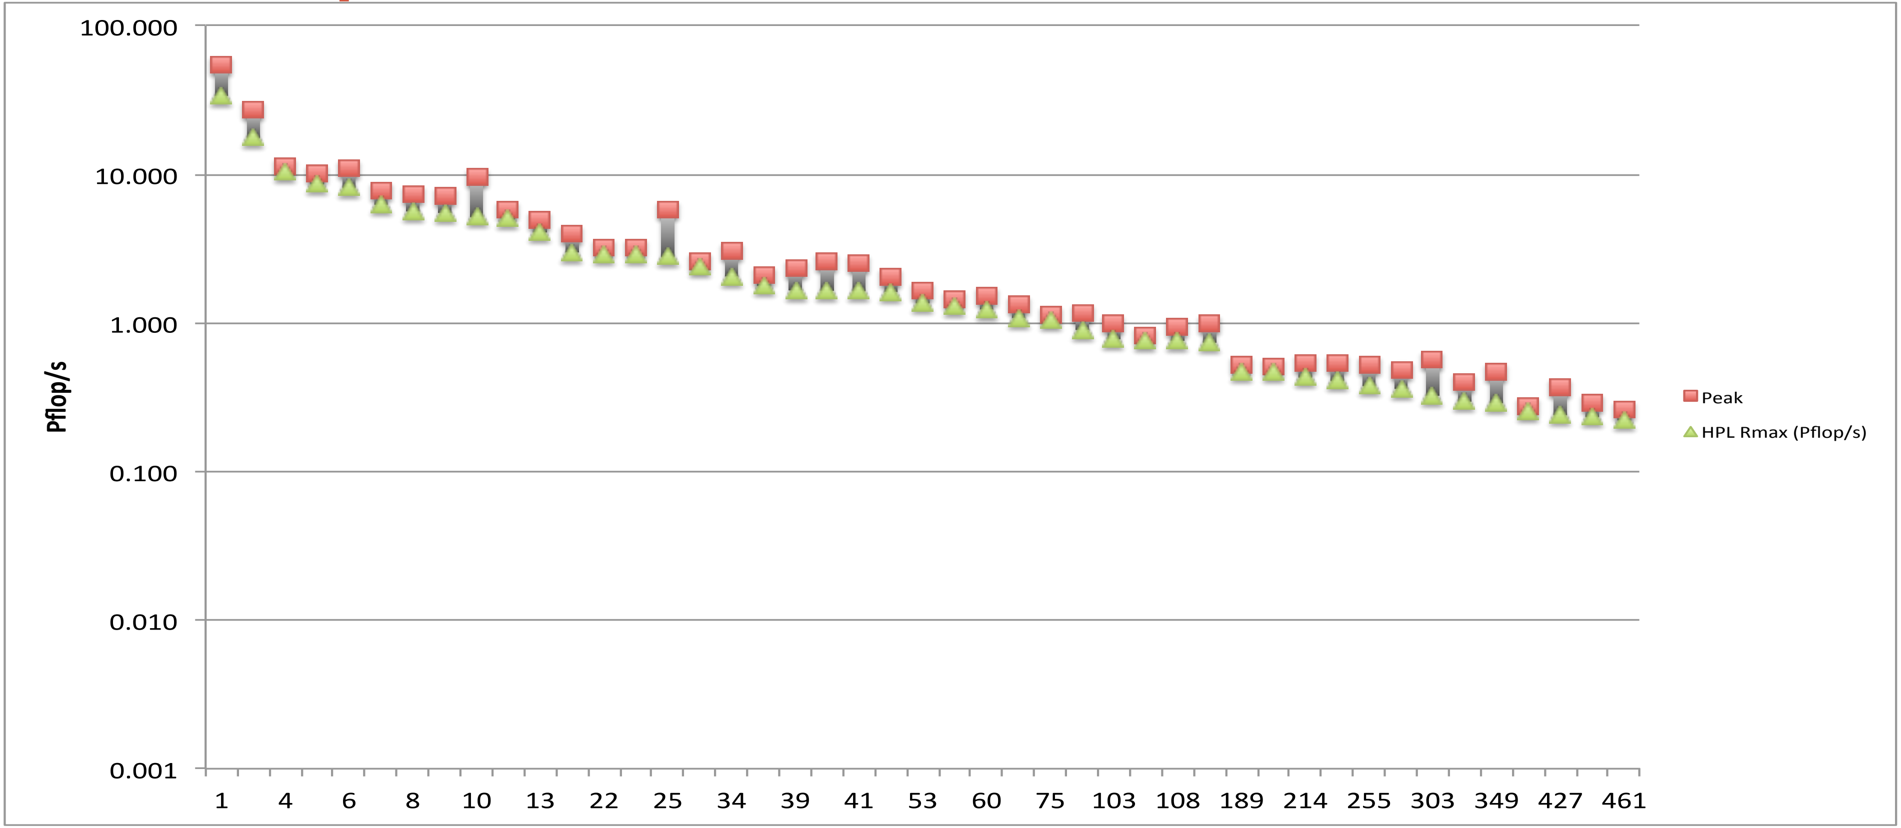

Comparison Peak, HPL, & HPCG

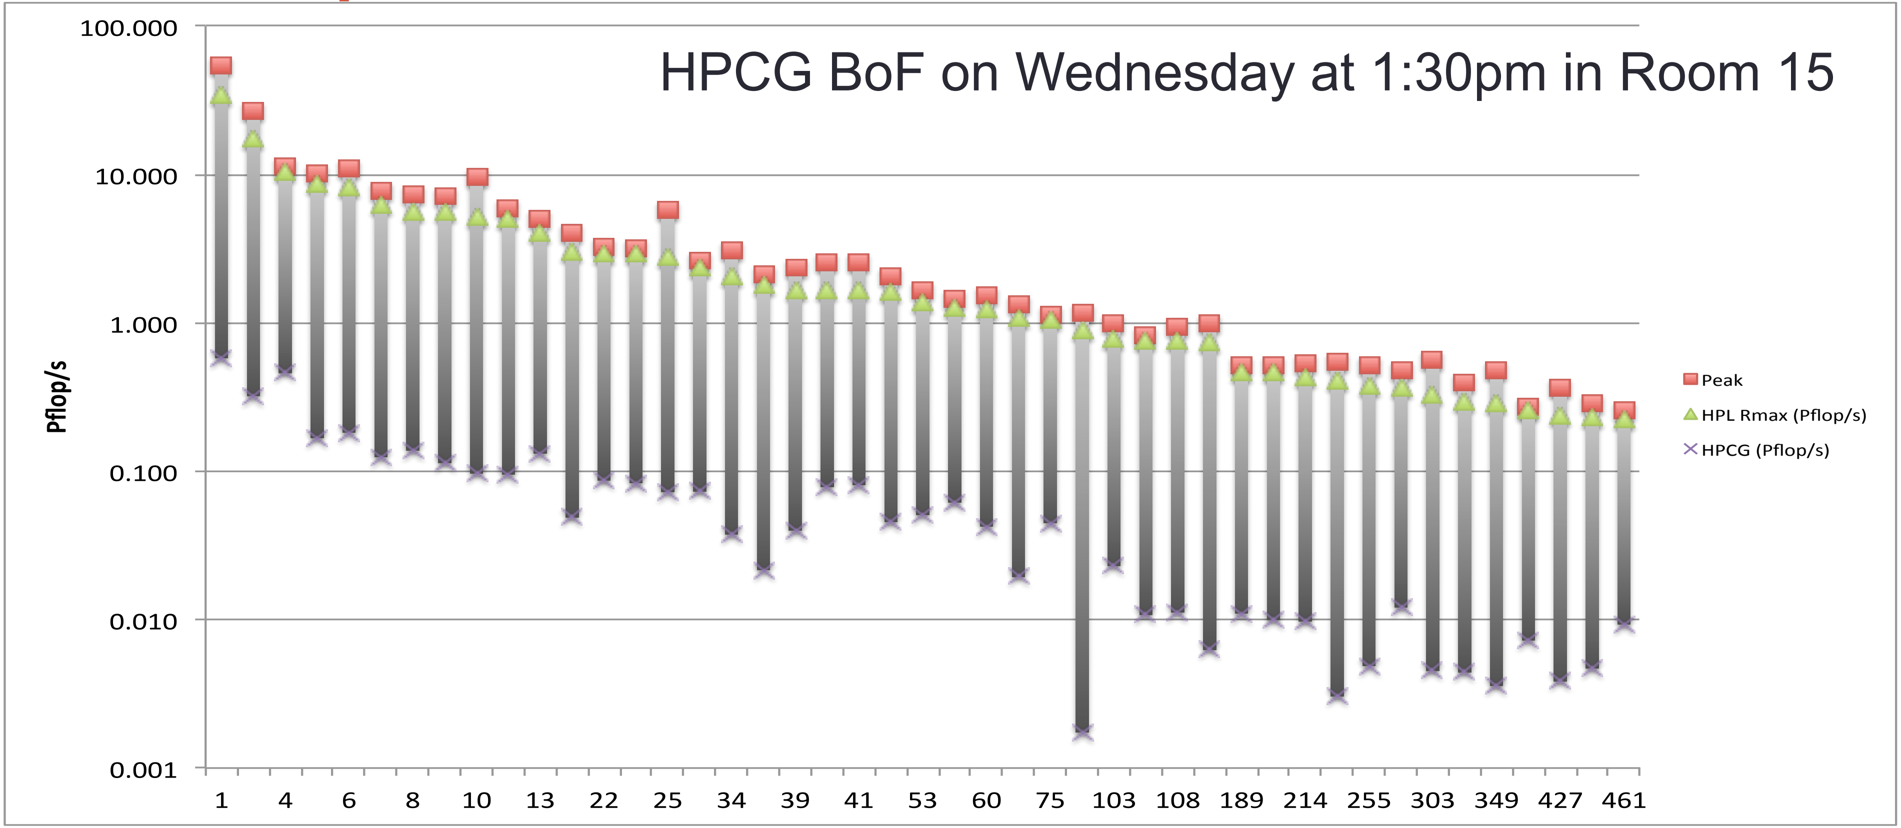

HPCG Results, Nov 2015, 1-10

Rank	Site	Computer	Cores	Rmax	HPCG	HPCG /HPL	% of Peak
1	NSCC / Guangzhou	Tianhe-2 NUDT, Xeon 12C 2.2GHz + Intel Xeon Phi 57C + Custom	3,120,000	33.863	0.5800	1.7%	1.1%
2	RIKEN Advanced Institute for Computational Science	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect	705,024	10.510	0.4608	4.4%	4.1%
3	DOE/SC/Oak Ridge Nat Lab	Titan - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x	560,640	17.590	0.3223	1.8%	1.2%
4	DOE/NNSA/LANL/SNL	Trinity - Cray XC40, Intel E5-2698v3, Aries custom	301,056	8.1009	0.1826	2.3%	1.6%
5	DOE/SC/Argonne National Laboratory	Mira - BlueGene/Q, Power BQC 16C 1.60GHz, Custom	786,432	8.587	0.1670	1.9%	1.7%
6	HLRS/University of Stuttgart	Hazel Hen - Cray XC40, Intel E5-2680v3, Infiniband FDR	185,088	5.640	0.1380	2.4%	1.9%
7	NASA / Mountain View	Pleiades - SGI ICE X, Intel E5-2680, E5-2680V2, E5-2680V3, Infiniband FDR	186,288	4.089	0.1319	3.2%	2.7%
8	Swiss National Supercomputing Centre (CSCS)	Piz Daint - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x	115,984	6.271	0.1246	2.0%	1.6%
9	KAUST / Jeda	Shaheen II - Cray XC40, Intel Haswell 2.3 GHz 16C, Cray Aries	196,608	5.537	0.1139	2.1%	1.6%
10	Texas Advanced Computing Center/Univ. of Texas	Stampede - PowerEdge C8220, Xeon E5-2680 8C 2.7GHz, Infiniband, Phi SE10P	522,080	5.168	0.0968	1.9%	1.0%

HPCG Results, Nov 2015, 11-20

Rank	Site	Computer	Cores	Rmax	HPCG	HPCG /HPL	% of Peak
11	Forschungszentrum Jülich	JUQUEEN - BlueGene/Q	458,752	5.0089	0.0955	1.9%	1.9%
12	Information Technology Center, Nagoya University	ITC, Nagoya - Fujitsu PRIMEHPC FX100	92,160	2.91	0.0865	3.0%	3.0%
13	Leibniz Rechenzentrum	SuperMUC - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR	147,456	2.897	0.0833	2.9%	2.9%
14	EPSRC/University of Edinburgh	ARCHER - Cray XC30, Intel Xeon E5 v2 12C 2.700GHz, Aries interconnect	118,080	1.643	0.0808	4.9%	4.9%
15	DOE/SC/LBNL/NERSC	Edison - Cray XC30, Intel Xeon E5-2695v2 12C 2.4GHz, Aries interconnect	133,824	1.655	0.0786	4.8%	4.8%
16	National Institute for Fusion Science	Plasma Simulator - Fujitsu PRIMEHPC FX100, SPARC64 Xifx, Custom	82,944	2.376	0.0732	3.1%	3.1%
17	GSIC Center, Tokyo Institute of Technology	TSUBAME 2.5 - Cluster Platform SL390s G7, Xeon X5670 6C 2.93GHz, Infiniband QDR, NVIDIA K20x	76,032	2.785	0.0725	2.6%	2.6%
18	HLRS/Universitaet Stuttgart	Hornet - Cray XC40, Xeon E5-2680 v3 2.5 GHz, Cray Aries	94,656	2.763	0.0683	2.4%	2.4%
19	Max-Planck-Gesellschaft MPI/IPP	iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband	65,320	1.283	0.0661	4.8%	4.8%
20	CEIST / JAMSTEC	Earth Simulator - NEC SX-ACE	8,192	0.487	0.0615	11.9%	11.9%

HPCG Highlights

- 63 Systems:
  - Up from 42 at ISC'15, 25 at SC'14 and 15 at ISC'14.
  - Most entries from the very top of the TOP500 list.
- New supercomputers (also coming to TOP500) are:
  - New #4: DOE Trinity (Haswell-only), HPL #6.
  - New #6: HLRS "Hazel Hen", HPL #8.
- Strong showing from Japan and NEC SX machines:
  - Achieve over 10% of peak performance with HPCG

HPCG Lists over Time

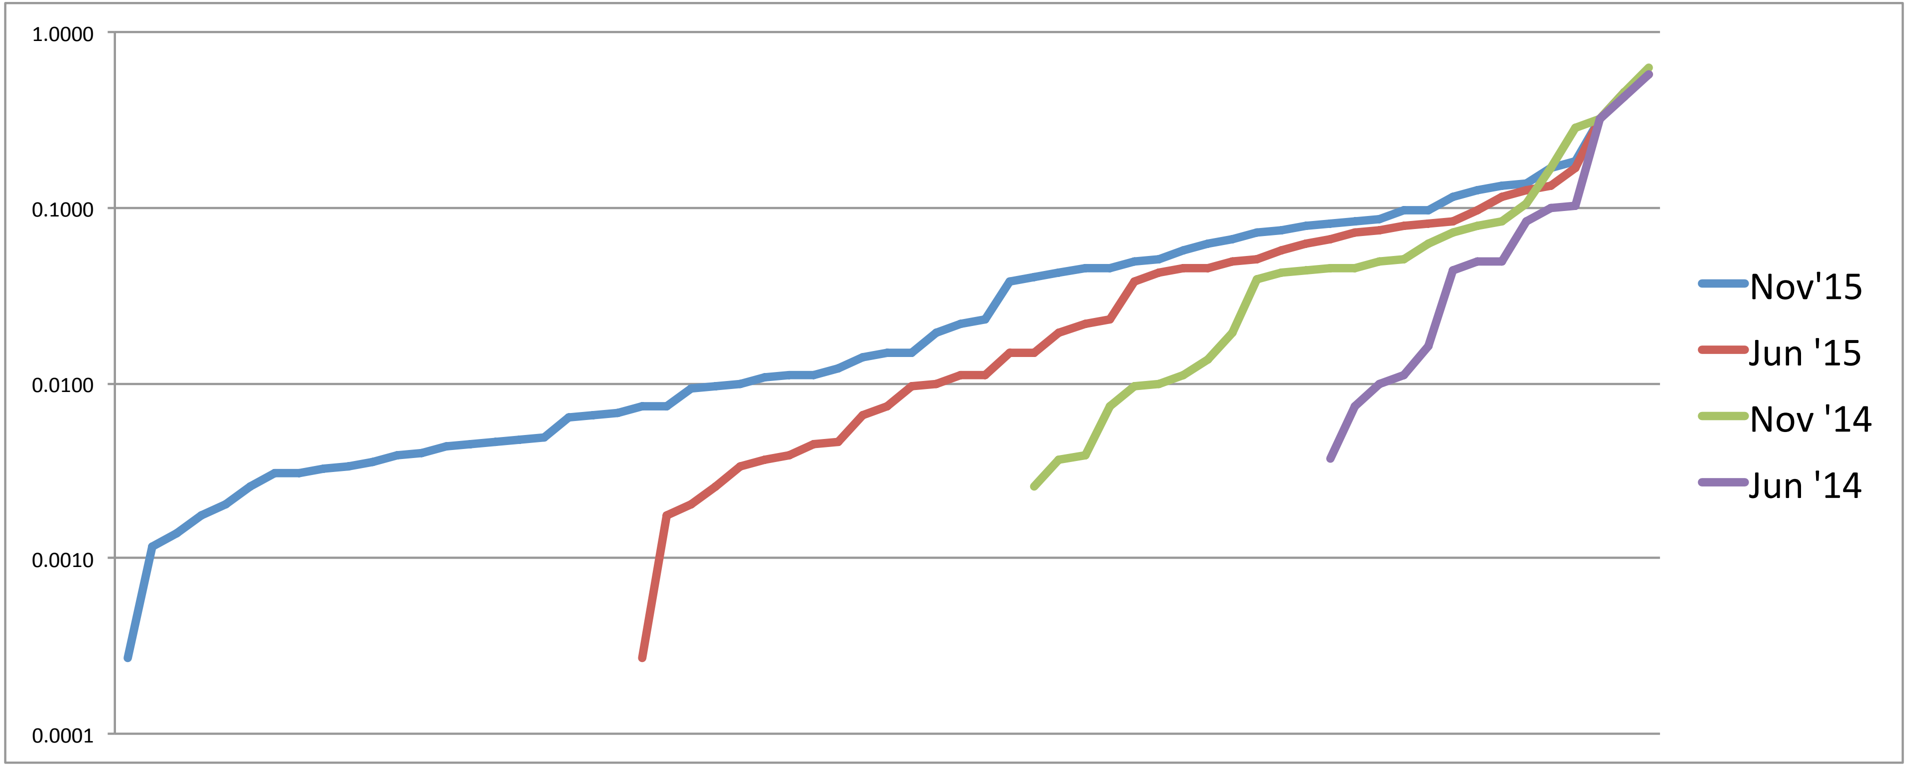

Performance by Region

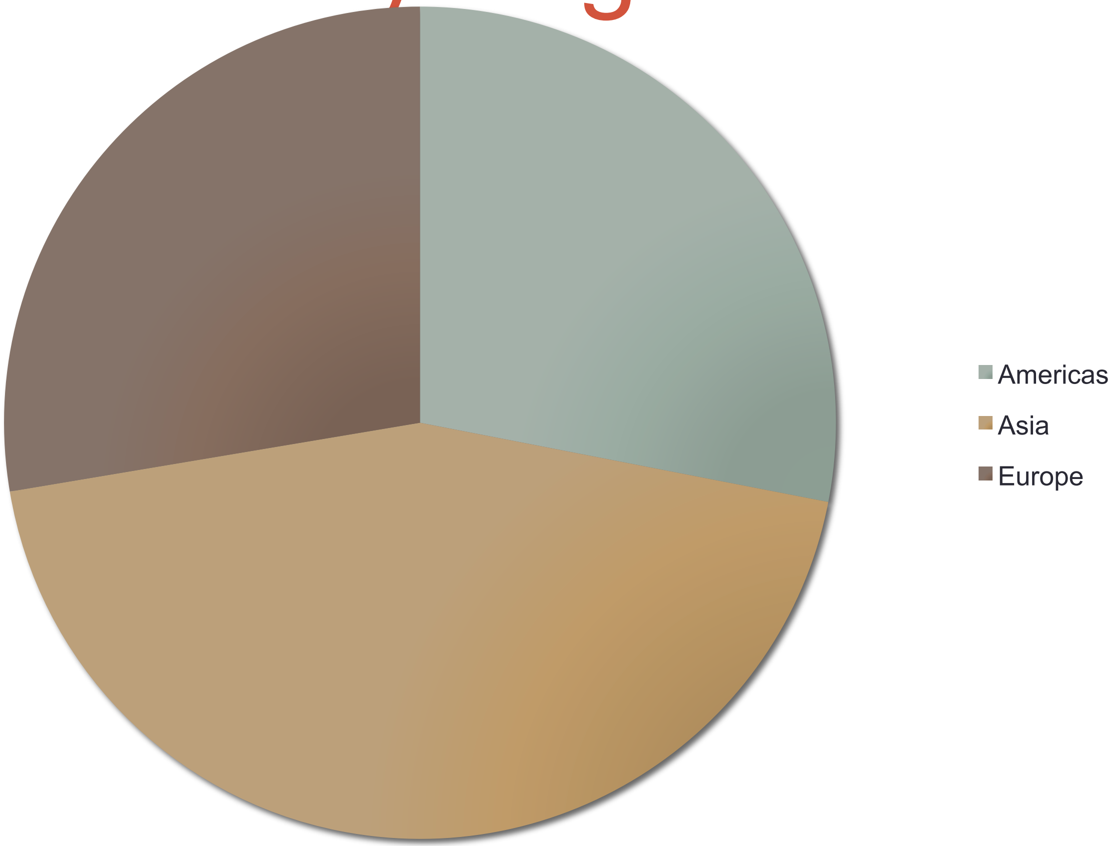

Performance by Country

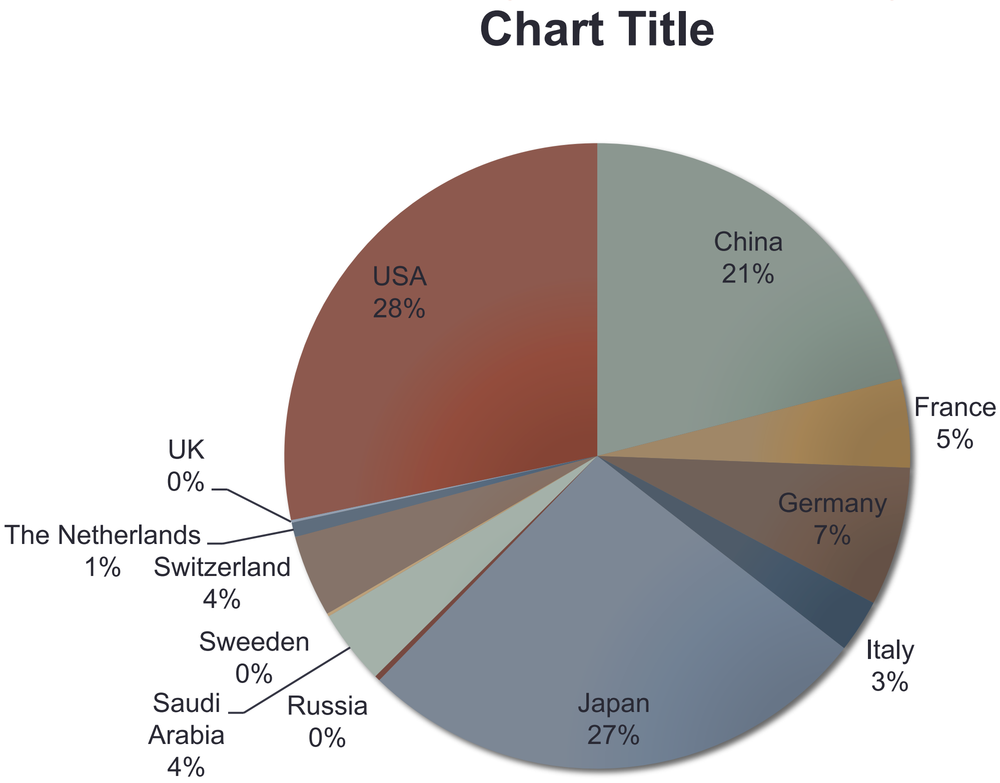

Performance by Network

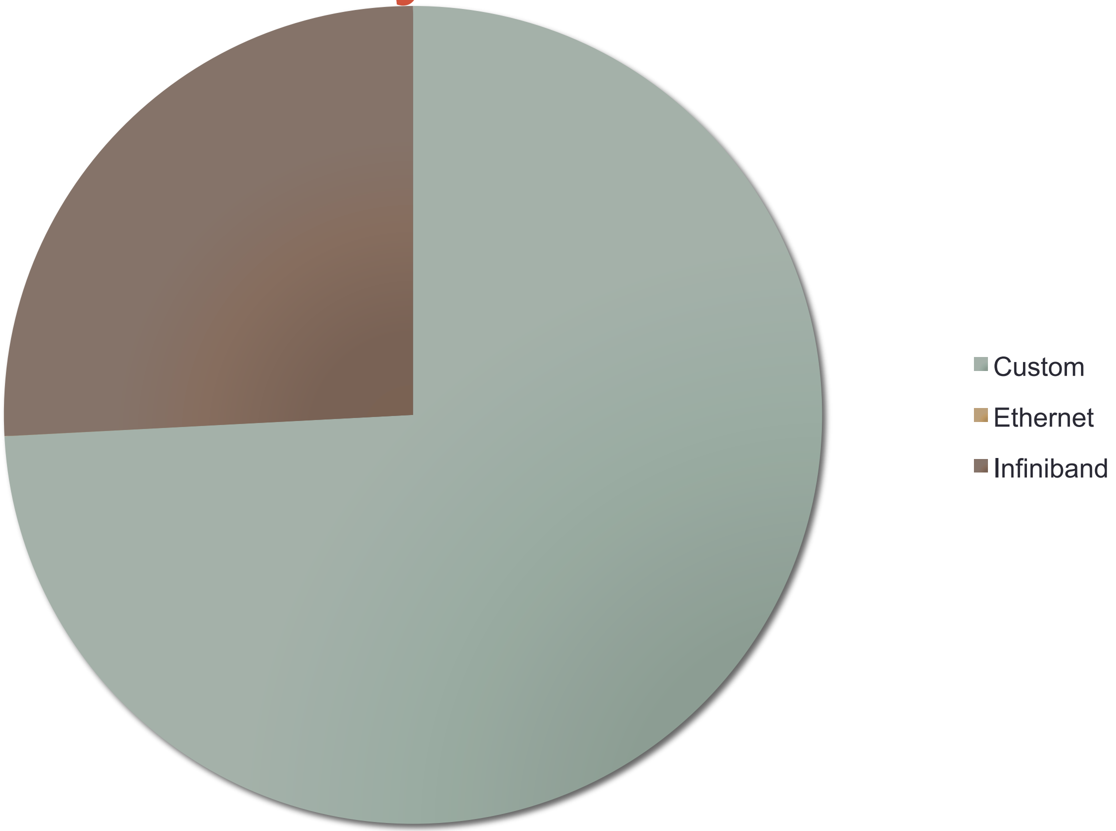

Performance by Network Details

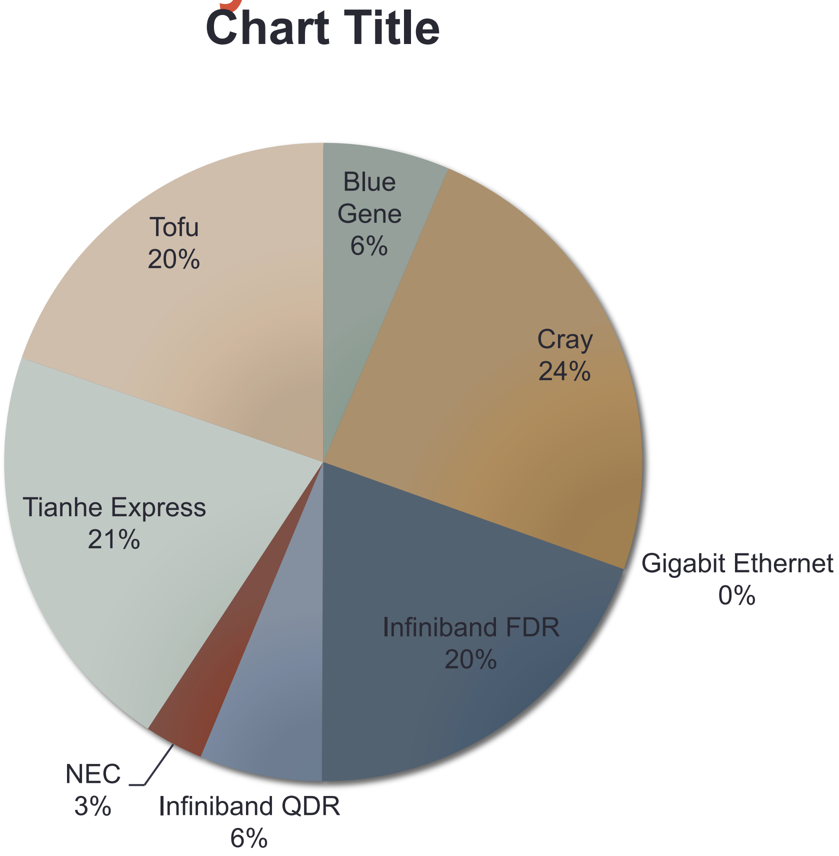

Performance by Processor

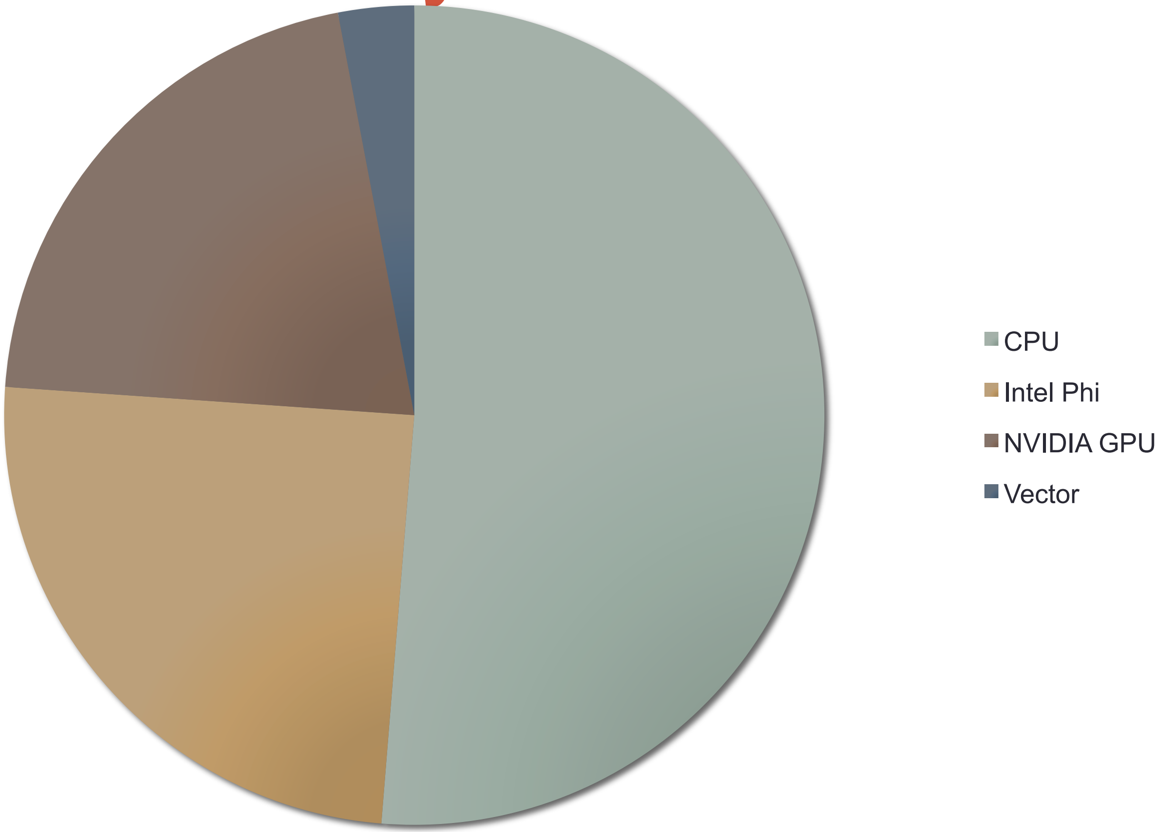

Performance by Processor Detail

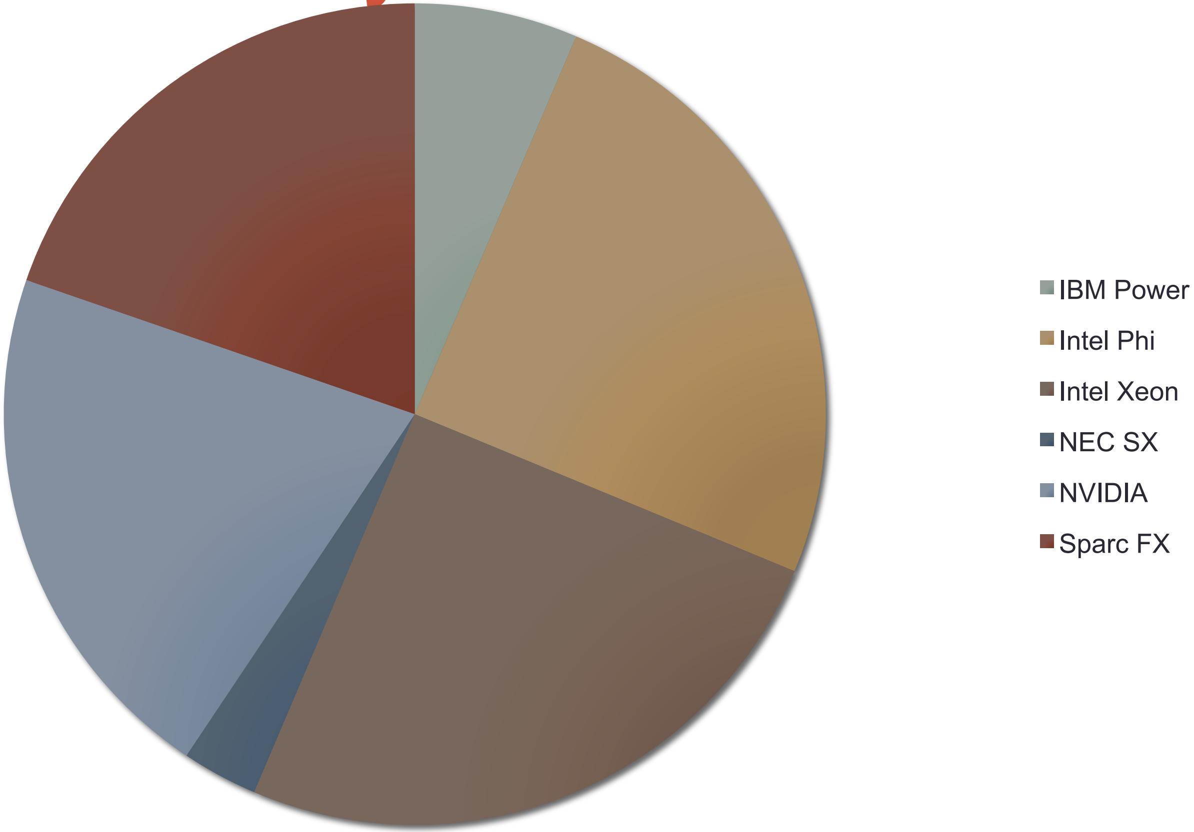

Performance by Segment

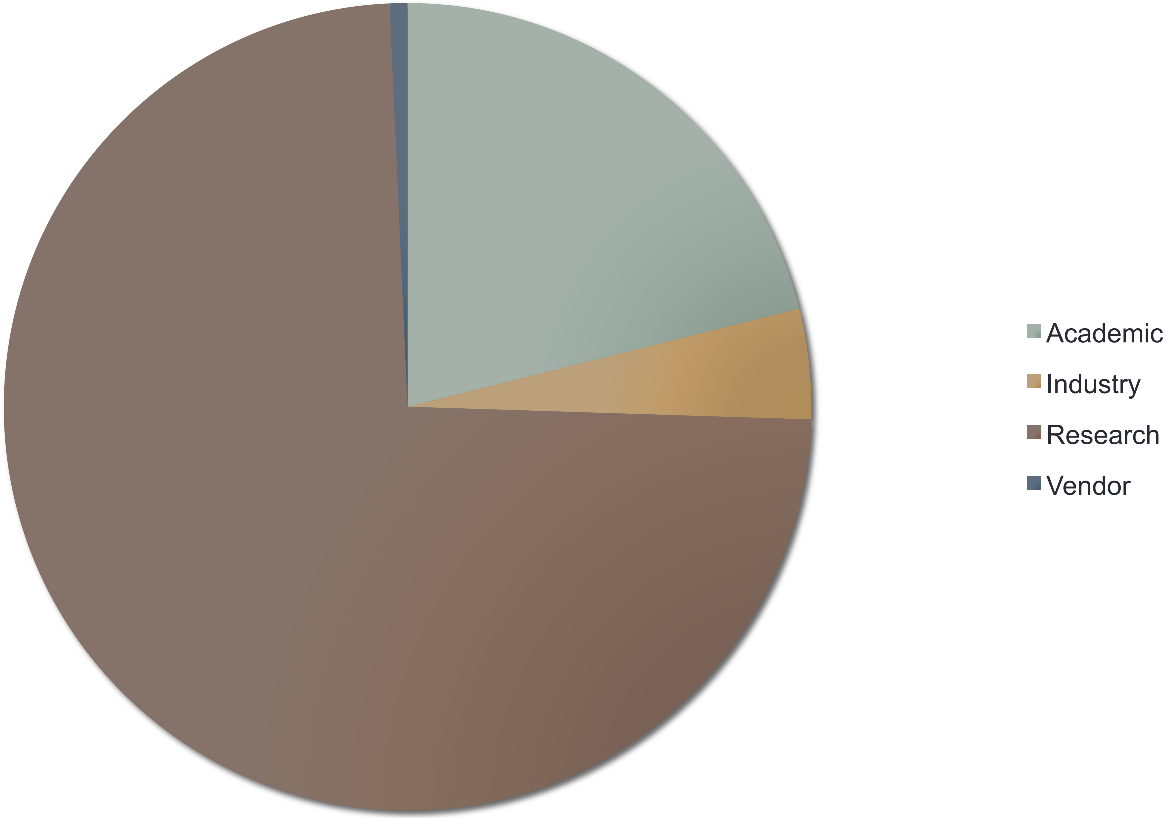

Performance by Integrator