HPCG UPDATE: ISC'17

Jack Dongarra
Michael Heroux
Piotr Luszczek

HPCG Update: Agenda

- Mike Heroux, HPCG Update.
- Jack Dongarra, Awards & Results.
- Erich Strohmaier, HPCG Analysis.
- Lin Gan, NSCC Wuxi
- Open discussion.

HPCG Snapshot

- High Performance Conjugate Gradients (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 collectives.
  - 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.

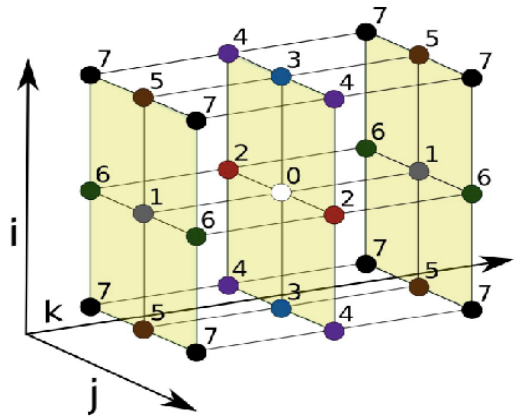

27-point stencil operator

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.
- Full performance discussion:
  - http://www.hpcg-benchmark.org -> “Performance Overview” tab.

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.
- Optimized 3.0 version:
  - Vendor or site developed.
  - Intel, Nvidia, IBM: Available to their customers.
  - Need: Contacts with ARM providers.
- All results require HPCG 3.0 use.

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
- Provides 2.4 rating and 3.0 rating in output.
- Command line option (--rt=) to specify the run time.
- ***"Quick Path" option to make obtaining results on production systems easier.***
  - Important for filling TOP500 gaps.

Other Items

- Reference version on GitHub:
  - https://github.com/hpcg-benchmark/hpcg
  - Website: hpcg-benchark.org.
  - Mail list hpcg.benchmark@gmail.com
- HPCG & Student Cluster Competitions.
  - Used in SC15/16/17, ASC
  - ISC17: First time HPCG is used for SCC.
- HPCG-optimized kernels in vendor libraries.

Summary

- HPCG is
  - Addressing original goals.
  - Rewarding vendor investment in features we care about.
- HPCG is integrated into TOP500.
  - Will work to fill in results gaps.
- Version 3.X is the final planned major version.
- Version 3.1:
  - Minor bug fixes to API.
  - Better support for heterogeneous systems.
  - More rigorous, automated policy checks.

HPCG RANKINGS
JUNE 2017

And The Winners Are…

We Have a Tie for Third Place

- **Sunway TaihuLight** – Sunway MPP, SW26010 260C 1.45GHz, Sunway
  NRCPC

- **Piz Daint** – Cray XC50, Intel Xeon E5-2690v3 12C 2.6GHz, Aries interconnect, NVIDIA Tesla P100
  Cray

- Both systems “clock” in at 0.48… Pflop/s for HPCG
  - Less than 1% difference in the performance


HPCG
JUNE 2017
PRESENTED AT
ISC High Performance
JUNE 21, 2017
NUMBER 3
SYSTEM
Sunway TaihuLight
National Supercomputing Center in Wuxi
CHINA
JACK DONGARRA
MICHAEL HEROUX
PIOTR LUSZCZEK
IN COLLABORATION WITH
ICL INNOVATIVE COMPUTING LABORATORY
THE UNIVERSITY OF TENNESSEE KNOXVILLE
Sandia National Laboratories
SPONSORED BY


HPCG
JUNE 2017
PRESENTED AT
ISC High Performance
JUNE 21, 2017
NUMBER 3
SYSTEM
Piz Daint
Swiss National Supercomputing Centre (CSCS)
SWITZERLAND
JACK DONGARRA
MICHAEL HEROUX
PIOTR LUSZCZEK
IN COLLABORATION WITH
ICL INNOVATIVE COMPUTING LABORATORY
THE UNIVERSITY OF TENNESSEE KNOXVILLE
Sandia National Laboratories
SPONSORED BY

HPCG

JUNE 2017

PRESENTED AT ISC High Performance JUNE 21, 2017

NUMBER 2

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, THE UNIVERSITY OF TENNESSEE KNOXVILLE, Sandia National Laboratories

SPONSORED BY


HPCG
JUNE 2017
PRESENTED AT
ISC High Performance
JUNE 21, 2017
NUMBER 1
SYSTEM
K computer
RIKEN Advanced Institute for Computational Science
JAPAN
ACHIEVED
0.603
Pflop/s
JACK DONGARRA
MICHAEL HEROUX
PIOTR LUSZCZEK
IN COLLABORATION WITH
ICL INNOVATIVE COMPUTING LABORATORY
THE UNIVERSITY OF TENNESSEE KNOXVILLE
Sandia National Laboratories
SPONSORED BY

HPCG Results, Jun 2017, top-10

Rank	Site	Computer	Cores	Rmax Pflops	HPCG Pflops	HPCG /HPL	% of Peak
1	RIKEN Advanced Institute for Computational Science, Japan	**K computer**, SPARC64 VIIIfx 2.0GHz, Tofu interconnect	705,024	10.5	0.60	5.7%	5.3%
2	NSCC / Guangzhou China	**Tianhe-2** NUDT, Xeon 12C 2.2GHz + Intel Xeon Phi 57C + Custom	3,120,000	33.8	0.58	1.7%	1.1%
3	National Supercomputing Center in Wuxi China	**Sunway TaihuLight** – Sunway MPP, SW26010 260C 1.45GHz, Sunway, NRCPC	10,649,600	93.0	0.48	0.5%	0.4%
3	Swiss National Supercomputing Centre (CSCS) Switzerland	**Piz Daint** – Cray XC50, Xeon E5-2690v3 12C 2.6GHz, Aries interconnect, NVIDIA P100, Cray	361,760	19.6	0.48	2.4%	1.9%
5	Joint Center for Advanced HPC Japan	**Oakforest-PACS** – PRIMERGY CX600 M1, Intel Xeon Phi 7250 68C 1.4GHz, Intel OmniPath, Fujitsu	557,056	24.9	0.39	2.8%	1.5%
6	DOE/SC/LBNL/NERSC USA	**Cori** – XC40, Intel Xeon Phi 7250 68C 1.4GHz, Cray Aries, Cray	632,400	13.8	0.36	2.6%	1.3%
7	DOE/NNSA/LLNL USA	**Sequoia** – IBM BlueGene/Q, PowerPC A2 16C 1.6GHz, 5D Torus, IBM	1,572,864	17.2	0.33	1.9%	1.6%
8	DOE/SC/Oak Ridge National Lab	**Titan** - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x	560,640	17.6	0.32	1.8%	1.2%
9	DOE/NNSA/LANL/SNL	**Trinity** - Cray XC40, Intel E5-2698v3, Aries custom, Cray	301,056	8.10	0.18	2.3%	1.6%
10	NASA / Mountain View	**Pleiades** - SGI ICE X, Intel E5-2680, E5-2680v2, E5-2680v3, E5-2680v4, Infiniband FDR, HPE	243,008	6.0	0.17	2.9%	2.5%

Bookends: Peak and HPL Only

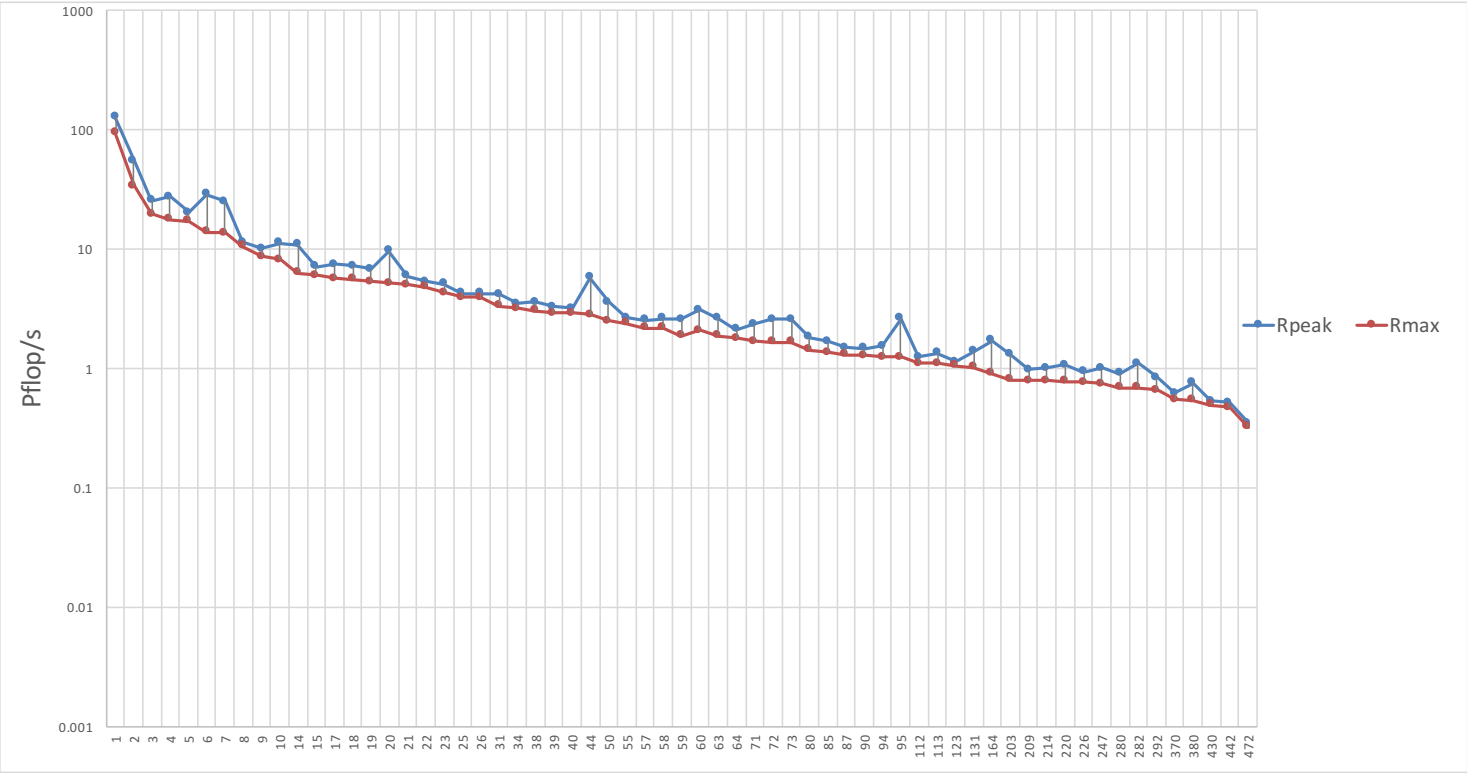

Bookends: Peak, HPL, and HPCG

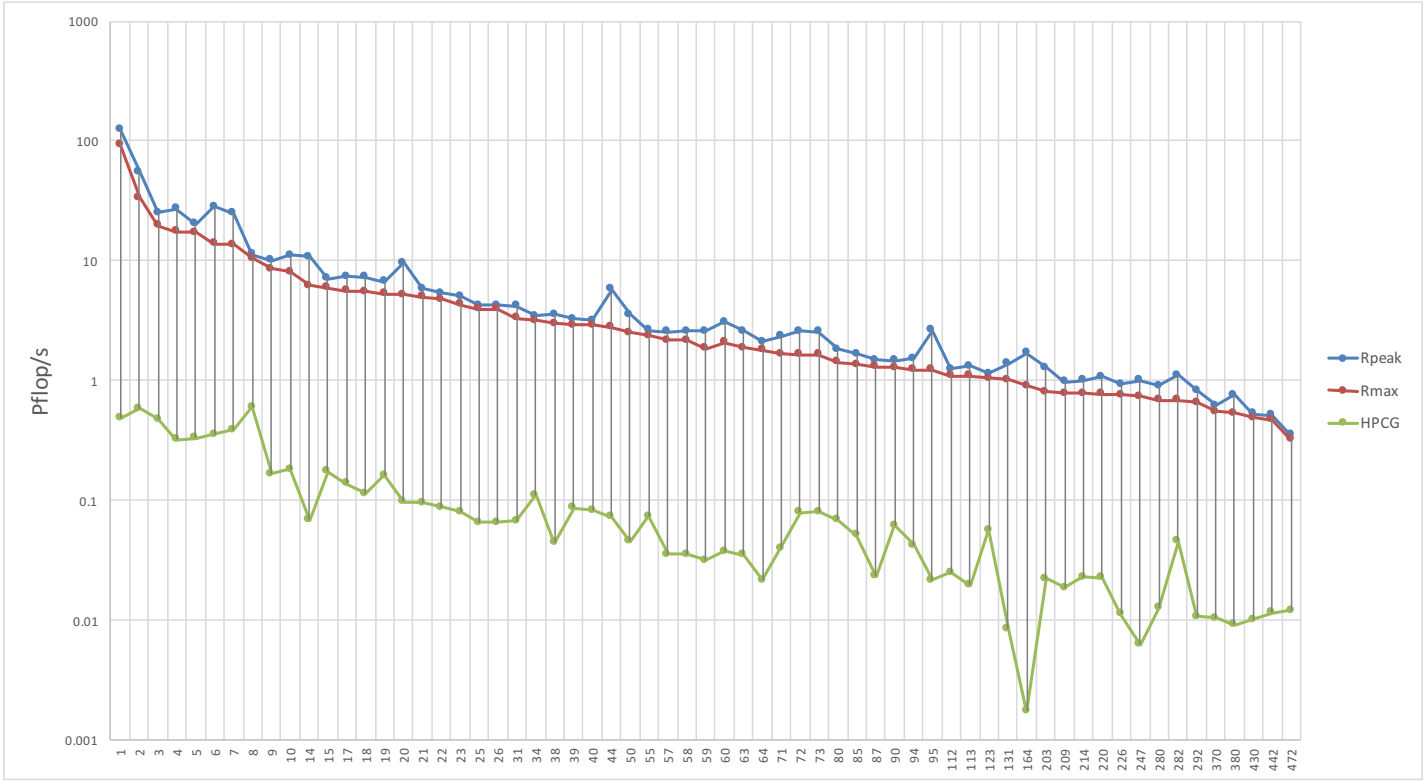

HPCG Lists over Time

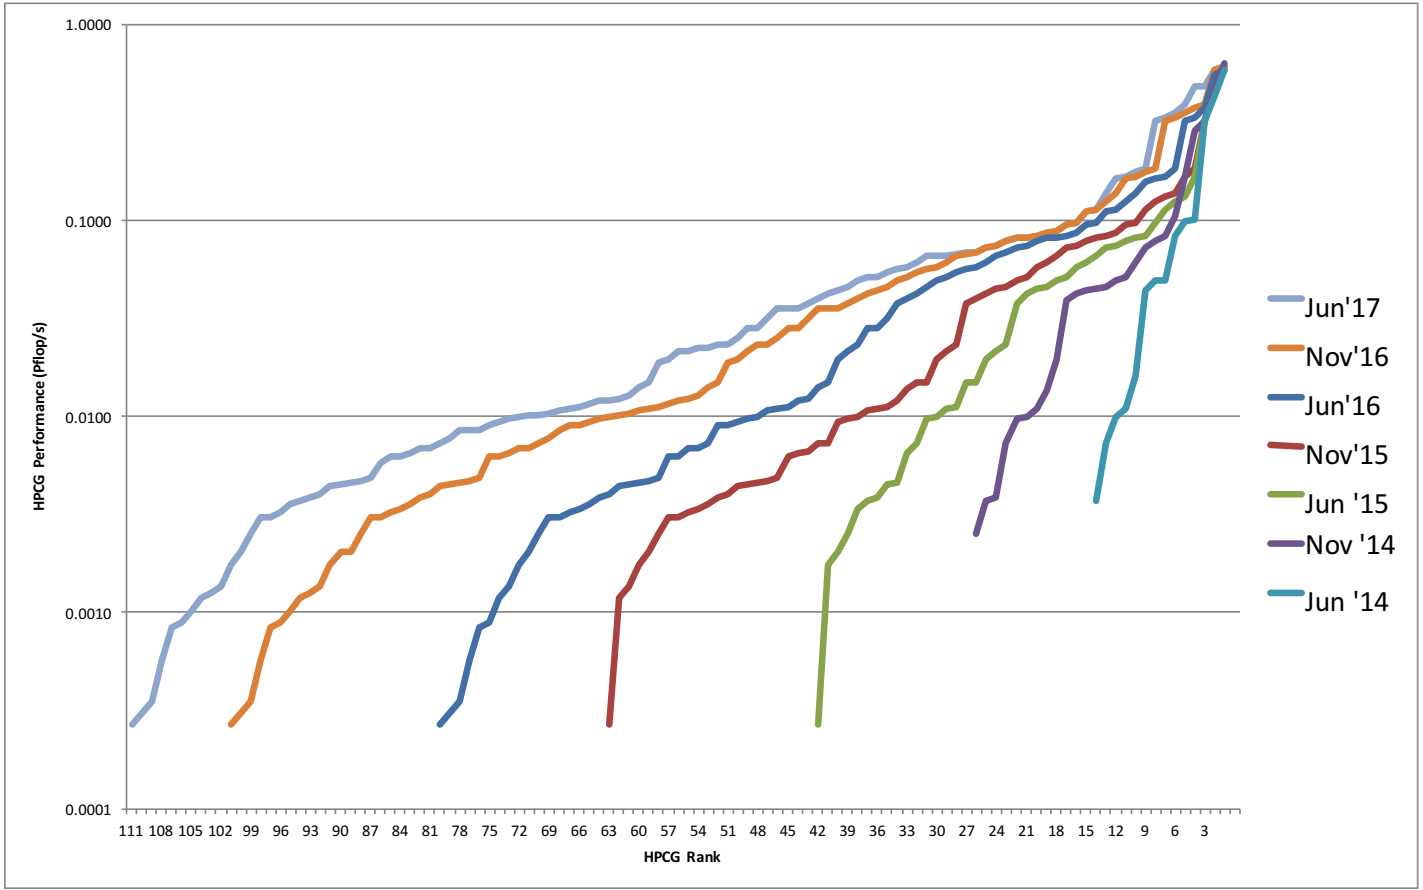

Performance by Region

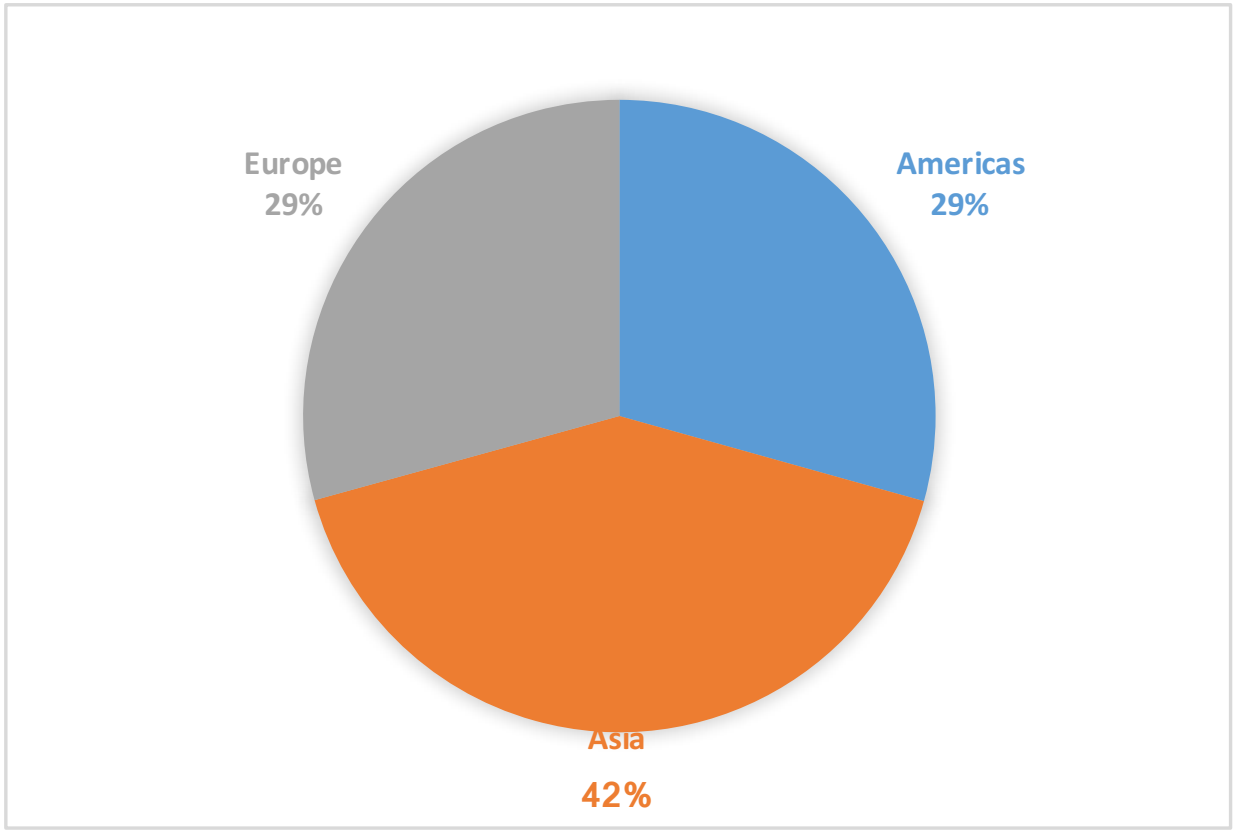

Top500
Americas: 34%
Asia: 41%
Europe:17%

Performance by Country

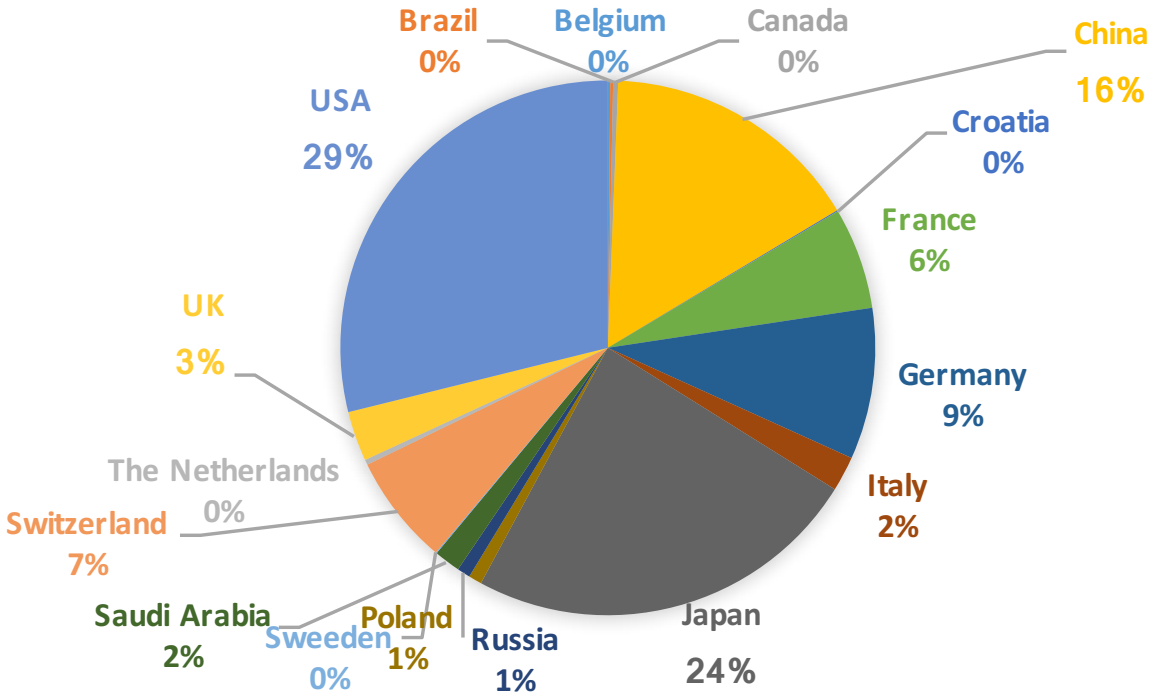

Top500
USA: 34%
China: 34%
Japan: 8%
Germany: 5%
France: 3%

Performance by Network Type

Top500
Infiniband: 30%
Custom: 53%
GigE: 17%

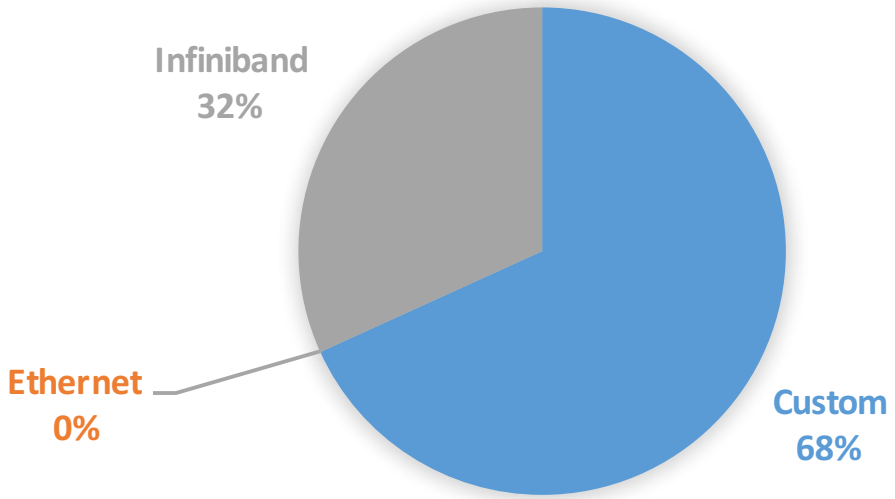

Performance by Network Details

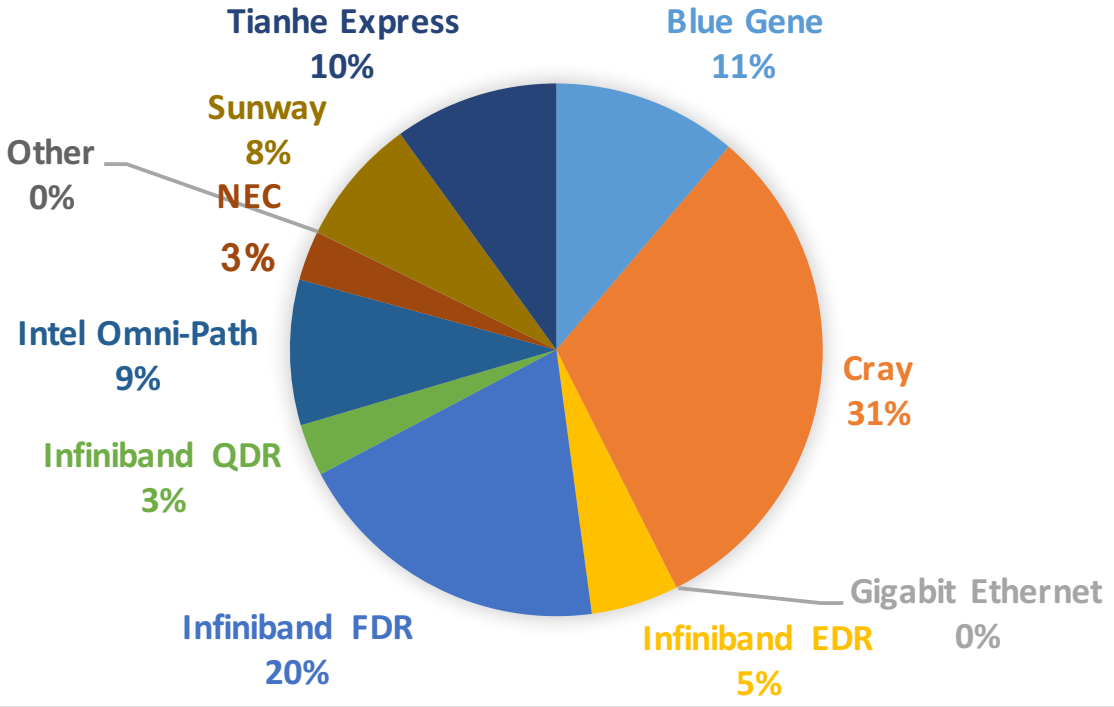

Performance by Processor Type

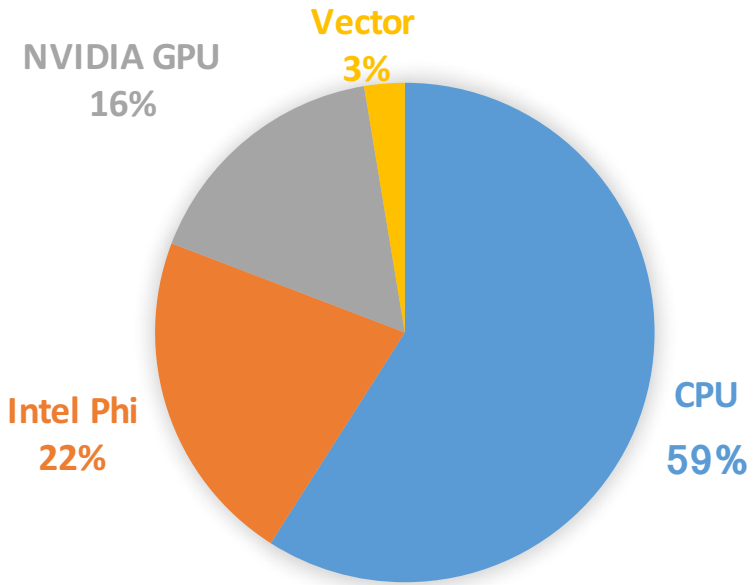

Performance by Processor Details

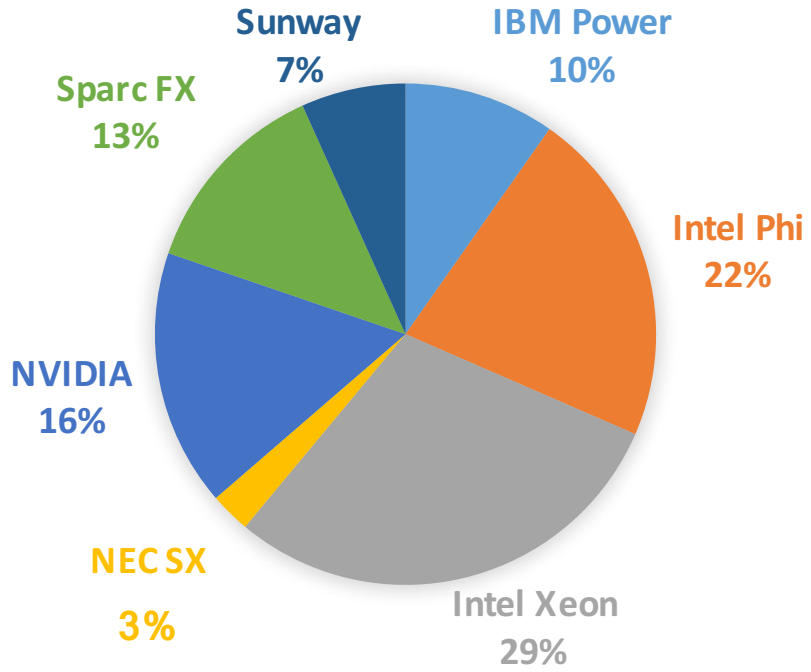

Performance by Segment

Top500
Industry: 50%
Research: 21%
Academic 19%
Govern: 7.6%

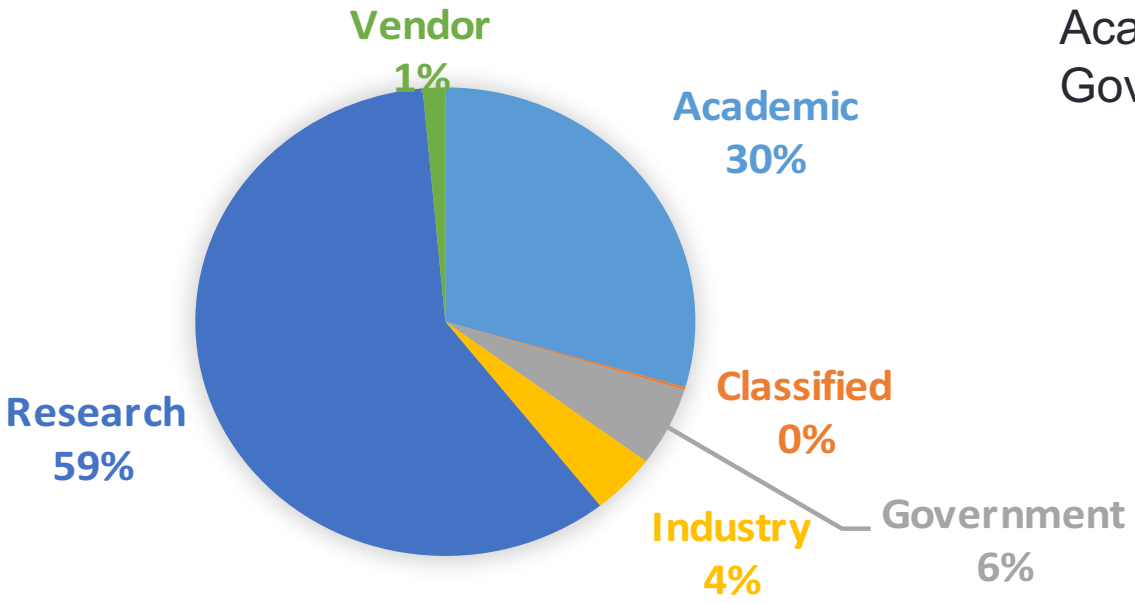

- https://www.top500.org/statistics/sublist/