

HPCG UPDATE: ISC'16

Jack Dongarra

Michael Heroux


Piotr Luszczek

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: $(np_x \times np_y \times np_z)$
- Process layout: $(n_x * np_x) \times (n_y * np_y) \times (n_z * np_z)$
- Global domain:
- Sparse matrix:
 - 27 nonzeros/row interior.
 - 8 – 18 on boundary.
 - Symmetric positive definite.



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:
 - $n_x = n_y = n_z = 104$ gives
 - $n_{\text{local}} = 1,124,864; 140,608; 17,576; 2,197$
 - 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.
- Optimized 3.0 version:
 - Vendor or site developed.
 - Used for all results (AFAWK).
 - Intel, Nvidia, IBM: Available to their customers.
- All future results require HPCG 3.0 use.
- Quick Path option makes 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-benchmark.org.
 - Mail list hpcg.benchmark@gmail.com
- HPCG & Student Cluster Competitions.
 - Used in SC15/16, ASC
 - SC15: HPCG replaced HPL, ranking matched overall cluster ranking.
- HPCG-optimized kernels going into vendor libraries.
- Next event: SC'16:
 - 80 entries ISC16
 - 61 – SC15, 42 – ISC15, 25 – SC14, 15 – ISC14

Summary

- HPCG is
 - Addressing original goals.
 - Rewarding vendor investment in features we care about.
- HPCG has traction.
 - Original goal of top 50 systems is reachable, and more.
- 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

JUNE 2016

And The Winners Are...

HPCG

PRESENTED AT
ISC
High Performance
JUNE 21, 2016

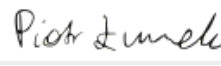
NUMBER 3

SYSTEM **Sunway
TaihuLight**
National Supercomputing
Center in Wuxi
CHINA

ACHIEVED **0.371**
Pflop/s


JACK DONGARRA


MICHAEL HEROUX


PIOTR LUSZCZEK

IN COLLABORATION WITH



SPONSORED BY



HPCG

PRESENTED AT



ISC

High Performance

JUNE 21, 2016

NUMBER 2

2

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

ACHIEVED **0.554**
Pflop/s



JACK DONGARRA



MICHAEL HEROUX



PIOTR LUSZCZEK

IN COLLABORATION WITH



SPONSORED BY



HPCG

PRESENTED AT
ISC
High Performance
JUNE 21, 2016



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

ACHIEVED **0.580**
Pflop/s


JACK DONGARRA


MICHAEL HEROUX


PIOTR LUSZCZEK

IN COLLABORATION WITH

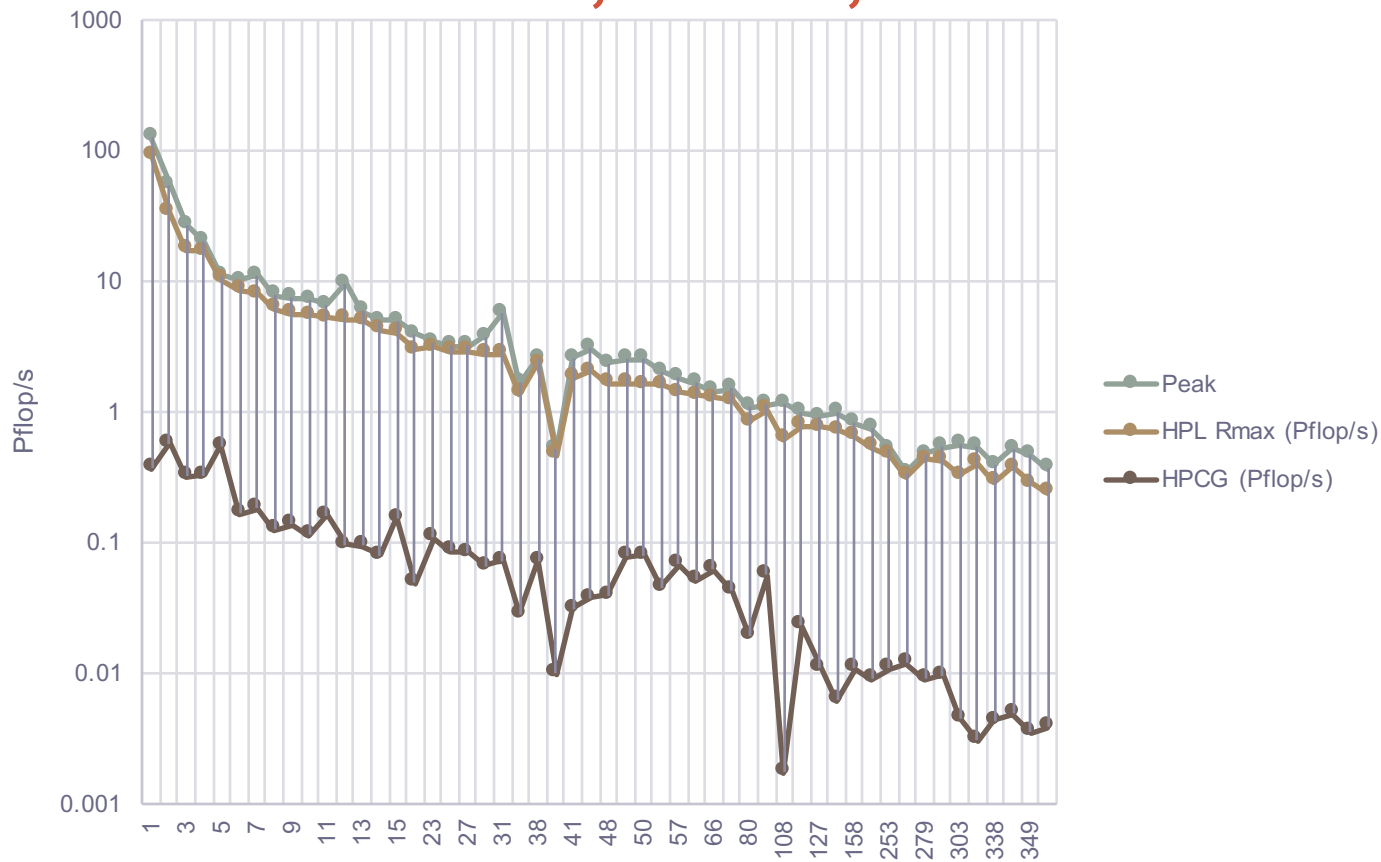


SPONSORED BY



Rank (HPL)	Site	Computer	Cores	Rmax	HPCG	HPCG/HPL	% of Peak
1 (2)	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 (5)	RIKEN Advanced Institute for Computational Science	K computer, SPARC64 VIIIfx 2.0GHz, Tofu interconnect	705,024	10.510	0.5544	5.3%	4.9%
3 (1)	National Supercomputing Center in Wuxi	Sunway TaihuLight -- SW26010, Sunway	10,649,600	93.015	0.3712	0.4%	0.3%
4 (4)	DOE/NNSA/LLNL	Sequoia - IBM BlueGene/Q	1,572,864	17.173	0.3304	1.9%	1.6%
5 (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%
6 (7)	DOE/NNSA/LANL/SNL	Trinity - Cray XC40, Intel E5-2698v3, Aries custom	301,056	8.101	0.1826	2.3%	1.6%
7 (6)	DOE/SC/Argonne National Laboratory	Mira - BlueGene/Q, Power BQC 16C 1.60GHz, Custom	786,432	8.587	0.1670	1.9%	1.7%
8 (11)	TOTAL	Pangea -- Intel Xeon E5-2670, Infiniband FDR	218592	5.283	0.1627	3.1%	2.4%
9 (15)	NASA / Mountain View	Pleiades - SGI ICE X, Intel E5-2680, E5-2680V2, E5-2680V3, Infiniband FDR	185,344	4.089	0.1555	3.8%	3.1%
10 (9)	HLRS/University of Stuttgart	Hazel Hen - Cray XC40, Intel E5-2680v3, Cray Aries	185,088	5.640	0.1380	2.4%	1.9%

Bookends: Peak, HPL, and HPCG



HPCG Highlights

- 80 Systems:
 - Up from 61 at SC'15, 42 at ISC'15, 25 at SC'14 and 15 at ISC'14.
 - All of the top 10, most of the top 20.
- Notable new entries are:
 - New #3: Sunway TaihuLight, HPL #1
 - New #4: Sequoia, HPL #4.
- Expect upgraded Trinity result (with KNL) for SC'16.
- HPCG has diverse adoption:
 - Vendor/site versions complete.
 - Cluster competitions.
 - Industry analysis.

11-20

Rank	Site	Computer	Cores	Rmax	HPCG	HPCG/HPL	% of Peak
11	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%
12	KAUST / Jeddah	Shaheen II - Cray XC40, Intel Haswell 2.3 GHz 16C, Cray Aries	196,608	5.537	0.1139	2.1%	1.6%
13	Japan Aerospace eXploration Agency	SORA-MA -- SPARC64 Xlfx	103,680	3.157	0.1102	3.5%	3.2%
14	Texas Advanced Computing Center/Univ. of Texas	Stampede - PowerEdge C8220, Xeon E5-2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi SE10P	522,080	5.168	0.0968	1.9%	1.0%
15	Forschungszentrum Jülich	JUQUEEN - BlueGene/Q	458,752	5.009	0.0955	1.9%	1.6%
16	Information Technology Center, Nagoya University	ITC, Nagoya - Fujitsu PRIMEHPC FX100, SPARC64 Xlfx, Tofu interconnect 2	92,160	2.910	0.0865	3.0%	2.7%
17	Leibniz Rechenzentrum	SuperMUC - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR	147,456	2.897	0.0833	2.9%	2.6%
18	DOE/NNSA/LLNL	Vulcan - IBM BlueGene/Q	393,216	4.293	0.0809	1.9%	1.6%
19	EPSRC/University of Edinburgh	ARCHER - Cray XC30, Intel Xeon E5 v2 12C 2.700GHz, Aries interconnect	118,080	1.643	0.0808	4.9%	3.2%
20	DOE/SC/LBNL/NERSC	Edison - Cray XC30, Intel Xeon E5-2695v2 12C 2.4GHz, Aries interconnect	133,824	1.655	0.0786	4.8%	3.1%

21-30

Rank	Site	Computer	Cores	Rmax	HPCG	HPCG/HPL	% of Peak
21	National Institute for Fusion Science	Plasma Simulator - Fujitsu PRIMEHPC FX100, SPARC64 Xifx, Tofu Interconnect 2	82,944	2.376	0.0732	3.1%	2.8%
22	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%	1.3%
23	Forschungszentrum Jülich	JURECA - T-Platform V-Class Cluster, Xeon E5-2680v3 12C 2.5GHz, Infiniband EDR, NVIDIA Tesla K80/K40	49,476	1.425	0.0683	4.8%	3.8%
24	HLRS/Universitaet Stuttgart	Hornet - Cray XC40, Xeon E5-2680 v3 2.5 GHz, Cray Aries	94,656	2.763	0.0661	2.4%	1.7%
25	Max-Planck-Gesellschaft MPI/IPP	iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband FDR	65,320	1.283	0.0615	4.8%	4.2%
26	CEIST / JAMSTEC	Earth Simulator - NEC SX-ACE	8,192	0.487	0.0578	11.9%	11.0%
27	Information Technology Center, The University of Tokyo	Oakleaf-FX -- SPARC64 lxfx	76,800	1.043	0.0565	5.4%	5.0%
28	CEIST / JAMSTEC	Earth Simulator -- NEC SX-ACE	8,192	0.487	0.0547	11.2%	10.4%
29	CEA/TGCC-GENCI	Curie thin nodes - Bullx B510, Xeon E5-2680 8C 2.700GHz, Infiniband QDR	77,184	1.359	0.0510	3.8%	3.1%
30	Exploration & Production - Eni S.p.A.	HPC2 - iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.8GHz, Infiniband FDR, NVIDIA K20x	62,640	3.003	0.0489	1.6%	1.2%

31-40

Rank	Site	Computer	Cores	Rmax	HPCG	HPCG/HPL	% of Peak
31	Grand Equipement National de Calcul Intensif - Centre Informatique National de l'Enseignement Superieur (GENCI-CINES)	Occigen Bullx B720, Xeon E5-2690v3 12C 2.600GHz, InfiniBand FDR	50,544	1.629	0.0455	2.8%	2.2%
32	International Fusion Energy Research Centre (IFERC), EU(F4E)- Japan Broader Approach collaboration	Helios Bullx B510, Xeon E5-2680 8C 2.700GHz, Infiniband QDR	70,560	1.237	0.0426	3.4%	2.8%
33	Cyfronet	Prometheus - HP ProLiant Intel E5-2680v3, Infiniband FDR	55,728	1.670	0.0399	2.4%	1.7%
34	Lvliang/National University of Defense Technology	Tianhe-2 Lvliang - Intel Xeon E5-2692v2 12C, TH Express-2, Intel Xeon Phi 31S1P	174,720	2.071	0.0376	1.8%	1.2%
35	Moscow State University / Research Computing Center	Lomonosov 2 - Intel Xeon E5-2680V2, Infiniband FDR, NVIDIA K40	37,120	1.849	0.0315	1.7%	1.2%
36	DKRZ - Deutsches Klimarechenzentrum	Mistral -- Intel Xeon E5-2695v4, Infiniband FDR	19,200	1.371	0.0283	2.1%	1.7%
37	Cyberscience Center, Tohoku University	Cyberscience Center, Tohoku University -- NEC SX-ACE	4,096	0.246	0.0279	11.3%	10.7%
38	Stanford University / Palo Alto	Xstream - Dual Intel E5-2680V2, 8-way NVIDIA K80, Infiniband FDR	237,120	0.781	0.0230	2.9%	2.3%
39	CINECA	Fermi - IBM BlueGene/Q	163,840	1.789	0.0216	1.2%	1.0%
40	SURFsara, Amsterdam	Cartesius2 bullx B720, dual socket Intel Xeon E5-2690 v3, Infiniband FDR	25,920	0.848	0.0195	2.3%	1.8%

41-50

Rank	Site	Computer	Cores	Rmax	HPCG	HPCG/HPL	% of Peak
41	Cyberscience Center/ Tohoku University	NEC SX-ACE 4C+IXS	2,048	0.123	0.0150	12.2%	11.4%
42	Cybermedia Center, Osaka University	Osaka U ACE -- NEC SX-ACE	2,048	0.123	0.0142	11.5%	10.8%
43	SGI	SGI ICE X -- Intel Xeon E5-2690v4, Infiniband EDR	16,128	0.602	0.0122	2.0%	1.8%
44	LNCC	Santos Dumont, Bullx Intel E5-2695v2, Infiniband FDR	17,616	0.321	0.0121	3.8%	3.5%
45	Intel	Endeavor - Intel Cluster, Dual Intel Xeon E5-2697v3 14C 2.700GHz, Infiniband FDR, Intel Xeon Phi 7120P	51,392	0.759	0.0112	1.5%	1.2%
46	Meteo France	Beaufix - Bullx DLC B710 Blades, Intel Xeon E5-2697v2 12C 2.7GHz, Infiniband FDR	24,192	0.469	0.0110	2.3%	2.1%
47	Saint Petersburg Polytechnic University	Polytechnic - RSC Tornado Intel E52697v3, Infiniband FDR	17,444	0.658	0.0108	1.6%	1.3%
48	Meteo France	Prolix - Bullx DLC B710 Blades, Intel Xeon E5-2697v2 12C 2.7GHz, Infiniband FDR	23,760	0.465	0.0100	2.1%	1.9%
49	Bull Angers	Manny Bullx B720, Xeon E5-2690v3 12C 2.600GHz, InfiniBand FDR	12,960	0.430	0.0097	2.3%	1.8%
50	University Heidelberg and University Mannheim	bwForCluster - Intel E5-2630v3, Infiniband QDR	7,552	0.241	0.0093	3.9%	3.2%

51-60

Rank	Site	Computer	Cores	Rmax	HPCG	HPCG/HPL	% of Peak
51	Michigan State University	Laconia -- Intel Xeon E5-2680v4, Infiniband EDR FDR	1,008,760	0.536	0.0091	1.7%	1.2%
52	University of Duisburg-Essen	magnitUDE -- Intel Xeon E5-2650v4, Intel OmniPath	12	0.437	0.0090	2.1%	1.9%
53	CALMIP / University of Toulouse	EOS - Bullx DLC B710 Blades, Intel Xeon E5-2680v2 10C 2.8GHz, Infiniband FDR	12,240	0.255	0.0073	2.8%	2.6%
54	Christian-Albrechts-Universitaet zu Kiel	NEC SX-ACE -- NEC SX-ACE	1,024	0.062	0.0068	11.1%	10.5%
55	GSIC Center, Tokyo Institute of Technology	TSUBAME-KFC/DL -- Intel Xeon E5-2620-V2, Infiniband FDR	2,720	0.273	0.0068	2.5%	1.6%
56	University of Tuebingen	BinAC -- Intel Xeon E5-2680v4, Infiniband FDR	4,800	0.209	0.0063	3.0%	2.2%
57	The Institute of Atmospheric Physics, Chinese Academy of Sciences	Earth System Numerical Simulator-1 - Intel E5-2680-V3, Infiniband FDR	24,912	0.738	0.0063	0.8%	0.6%
58	Joint Supercomputer Center RAS	MVS-10P - Intel E5-2690, Infiniband FDR, Xeon Phi SE10X	2,992	0.376	0.0049	1.3%	0.9%
59	University of Rijeka	Bura - Bullx Intel E5-2690v3, Infiniband FDR	5,952	0.234	0.0047	2.0%	1.6%
60	CINECA	Galileo - Dual Intel E5-2630 v3 2.4 GHz, Infiniband QDR, Dual NVIDIA K80	2,720		0.0046		1.9%

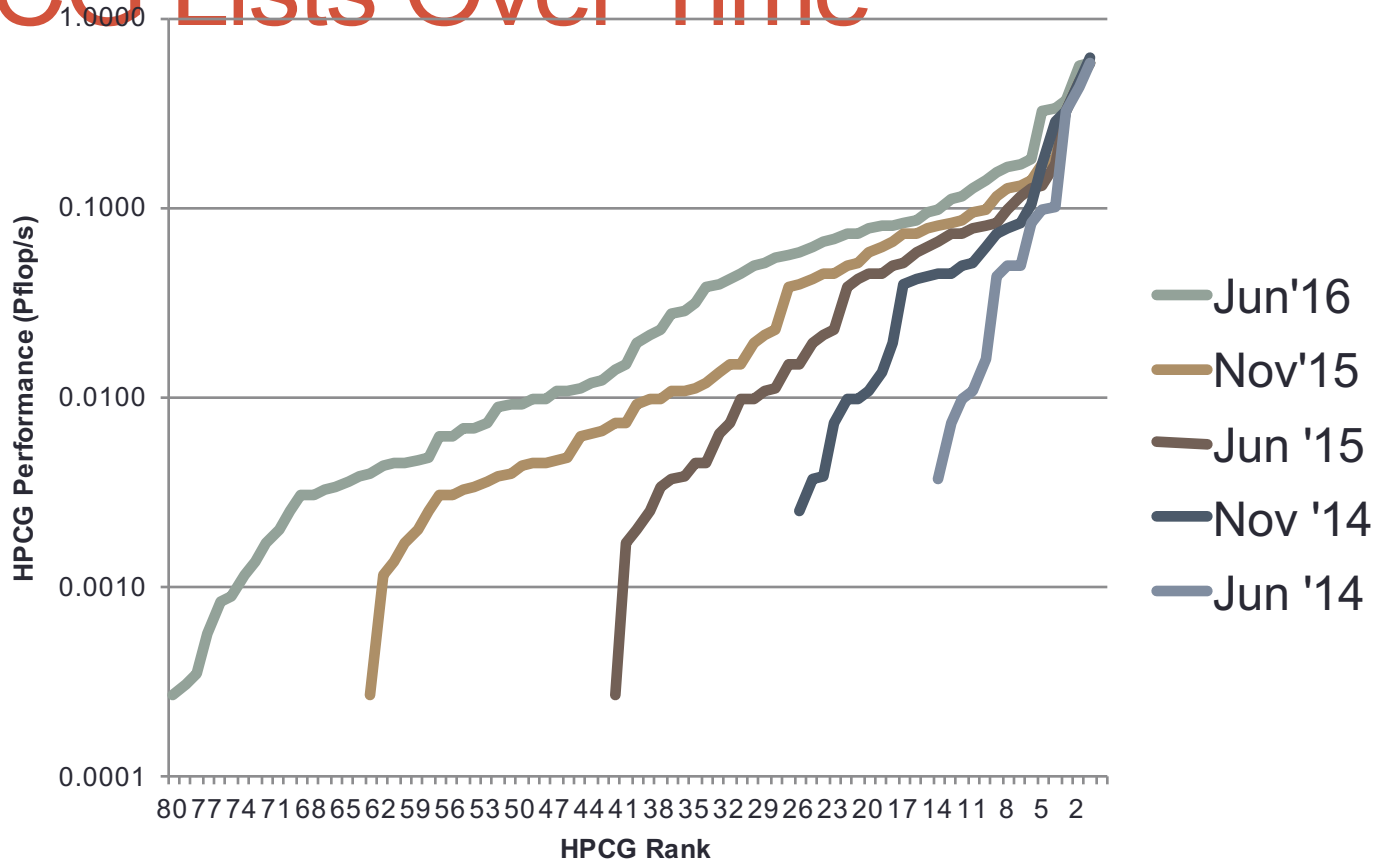
61-70

Rank	Site	Computer	Cores	Rmax	HPCG	HPCG/HPL	% of Peak
61	NSC / Linköping	Bifrost - ASUS, Intel Xeon E5-2640v3 8C 2.6GHz, Intel Truescale Infiniband QDR	10,256	0.326	0.0045	1.4%	0.8%
62	Shanghai Supercomputer Center	Magic Cube II - Intel E5-2680-V3, Infiniband EDR	9,960	0.296	0.0044	1.5%	1.1%
63	Max-Planck-Institut für Mikrostrukturphysik	Cruncher - Intel E5-2680-V3, Intel Truescale Infiniband QDR	12	0.112	0.0040	3.6%	2.8%
64	Cambridge University	Wilkes - Dell T620 Cluster, Intel Xeon E5-2630v2 6C 2.600GHz, Infiniband FDR, NVIDIA K20	5,120	0.240	0.0039	1.6%	1.0%
65	Chelyabinsk	RSC Tornado SUSU, Intel X5680, Infiniband QDR, Xeon Phi SE10X	4,032	0.288	0.0036	1.2%	0.8%
66	CINECA	Galileo - Dual Intel E5-2630 v3 2.4 GHz, Infiniband QDR, Dual Intel Xeon Phi 7120P	13,600		0.0034		1.5%
67	Atos Angers	Sid - Bullx Intel E5-2680v3, InfiniBand FDR	4,224	0.129	0.0032	2.5%	2.0%
68	St. Petersburg Polytechnic University	Polytechnic RSC PetaStream - Intel E5-2650 v2, Infiniband FDR, Xeon Phi 5120D	232	0.170	0.0031	1.8%	1.2%
69	Supercomputing Center of Chinese Academy of Sciences	Era-2 - Intel E5-2680-V3, Infiniband FDR, Xeon Phi + NVIDIA K20	13560	0.407	0.0030	0.7%	0.6%
70	SURF Sara	Cartesius - Bullx B515 cluster, Intel Xeon E5-2450v2 8C 2.5GHz, InfiniBand 4x FDR, Nvidia K40m	3,036	0.154	0.0025	1.7%	1.2%


71-80

Rank	Site	Computer	Cores	Rmax	HPCG	HPCG/HPL	% of Peak
71	CINECA	Galileo - Dual Intel E5-2630 v3 2.4 GHz, Infiniband QDR	6,400		0.0020		1.6%
72	Moscow State University / Research Computing Center	Lomonosov - Intel Xeon X5570/X5670/E5630 2.93/2.53 GHz, PowerXCell 8i Infiniband QDR, Dual NVIDIA Fermi 2070	78,660	0.617	0.0017	0.3%	0.2%
73	IT Services Provider	Aquarius - Intel Xeon E5-2640-V3, Infiniband QDR	8	0.034	0.0014	4.0%	3.2%
74	Joint Supercomputer Center RAS	RSC PetaStream - Intel E5-2667 v2, Infiniband FDR, Intel Xeon Phi 7120D	3,904	0.054	0.0012	2.2%	1.5%
75	Yaqingjie Street 30	hbemc_2016A -- Intel E5-2680v3, Infiniband FDR	2,304		0.0009		
76	Hefei City, Anhui Province	YUJING -- Intel Xeon E5-2680v3, custom	1,440	0.001	0.0008		
77	No. 180 Wusidong Road. Baoding City, Hebei Province, P.R.C	KunYu -- Intel Xeon E5-2680v3, Infiniband FDR	960	0.001	0.0006		
78	hongguancun Software Park II, No. 10 West Dongbeiwang Road, Haidian District, Beijing 100193, China	CSRC -- Intel Xeon E5-2680v3, Infiniband FDR	528	0.000	0.0004		
79	18, Xueyuan Road, Haidian District, Beijing, China	geo -- Intel Xeon E5-2680v3, Infiniband FDR	12	0.000	0.0003		
80	CINECA	Pico - Dual Intel Xeon E5-2670v2 2.5 GHz, Gigabit Ethernet	1,200		0.0003		1.1%


HPCG Lists Over Time



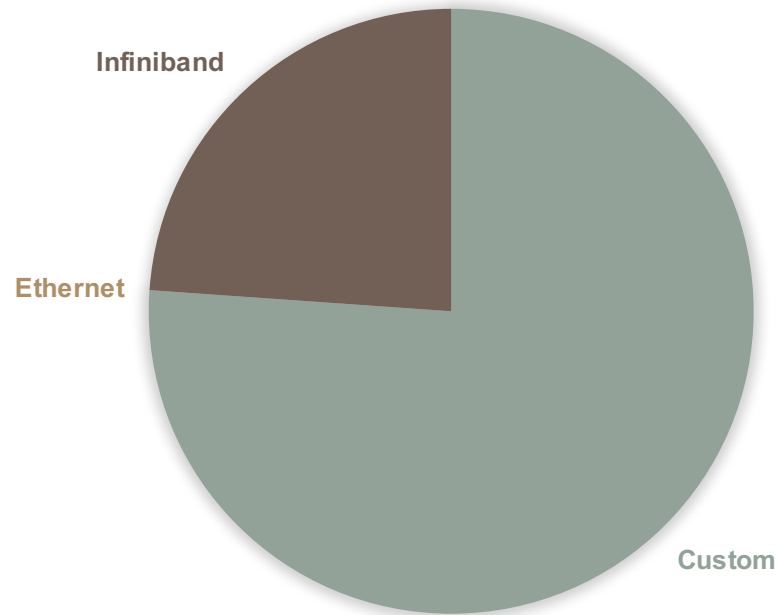
Performance by Region



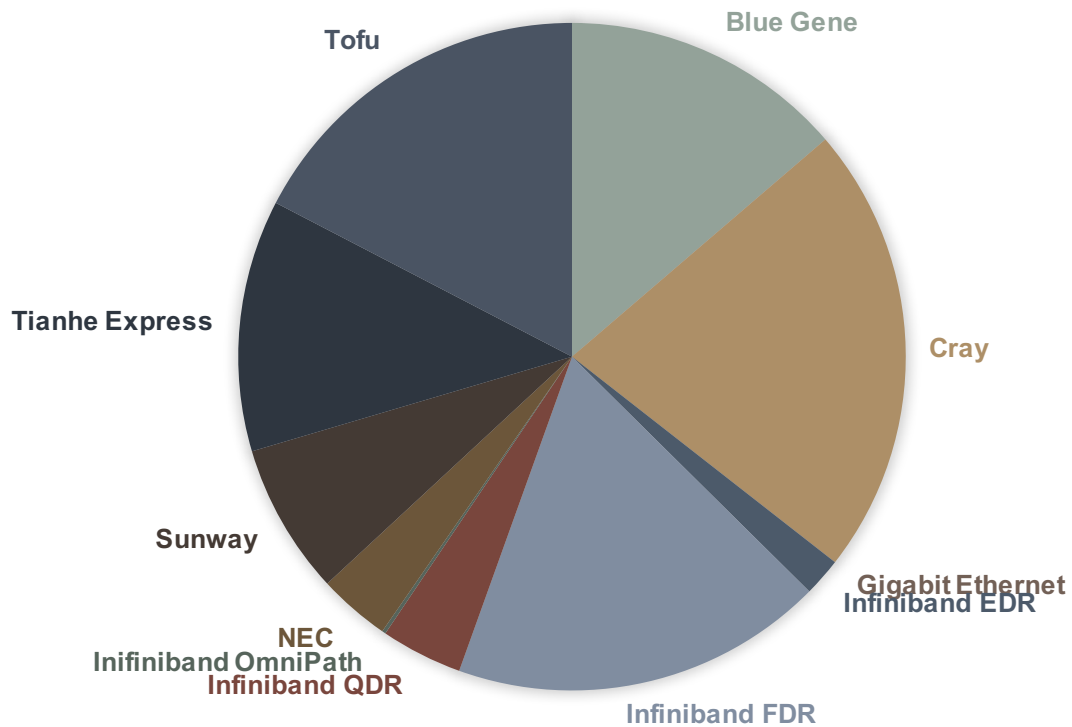
Performance by Country




Performance by Network Type




Performance by Network Details




Performance by Processor Type



Performance by Processor Detail



Performance by Segment



Performance by Integrator

