TOP 500
The List.

Highlights of the 51st TOP500 List

ISC 2018,
Frankfurt,
June 25, 2018

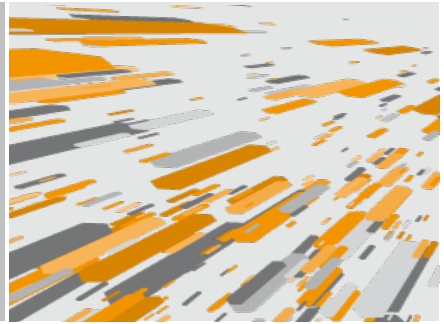

Erich
Strohmaier

ISC18 TOP500 TOPICS

- New #1
- New TOP5
- Slow-down, aging, and concentration
- 25 years of systems and sites
- China, a new twist
- Industry and GPUs
- HPCG petaflops

#	Site	Manufacturer	Computer	Country	Cores	Rmax [Pflops]	Power [MW]
1	Oak Ridge National Laboratory	IBM	Summit IBM Power System, P9 22C 3.07GHz, Mellanox EDR, NVIDIA GV100	USA	2,282,544	122.3	8.8
2	National Supercomputing Center in Wuxi	NRCPC	Sunway TaihuLight NRCPC Sunway SW26010, 260C 1.45GHz	China	10,649,600	93.0	15.4
3	Lawrence Livermore National Laboratory	IBM	Sierra IBM Power System, P9 22C 3.1GHz, Mellanox EDR, NVIDIA GV100	USA	1,572,480	71.6	
4	National University of Defense Technology	NUDT	Tianhe-2A ANUDT TH-IVB-FEP, Xeon 12C 2.2GHz, Matrix-2000	China	4,981,760	61.4	18.5
5	National Institute of Advanced Industrial Science and Technology	Fujitsu	AI Bridging Cloud Infrastructure (ABCI) PRIMERGY CX2550 M4, Xeon Gold 20C 2.4GHz, IB-EDR, NVIDIA V100	Japan	391,680	19.9	1.65
6	Swiss National Supercomputing Centre (CSCS)	Cray	Piz Daint Cray XC50, Xeon E5 12C 2.6GHz, Aries, NVIDIA Tesla P100	Switzerland	361,760	19.6	2.27
7	Oak Ridge National Laboratory	Cray	Titan Cray XK7, Opteron 16C 2.2GHz, Gemini, NVIDIA K20x	USA	560,640	17.6	8.21
8	Lawrence Livermore National Laboratory	IBM	Sequoia BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	1,572,864	17.2	7.89
9	Los Alamos NL / Sandia NL	Cray	Trinity Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries	USA	979,968	14.1	3.84
10	Lawrence Berkeley National Laboratory	Cray	Cori Cray XC40, Intel Xeons Phi 7250 68C 1.4 GHz, Aries	USA	622,336	14.0	3.94

System Overview

System Performance

- Peak performance of 200 petaflops for modeling & simulation
- Peak of 3.3 ExaOps for data analytics and artificial intelligence

Each node has

- 2 IBM POWER9 processors
- 6 NVIDIA Tesla V100 GPUs
- 608 GB of fast memory
- 1.6 TB of NVMe memory

The system includes

- 4608 nodes
- Dual-rail Mellanox EDR InfiniBand network
- 250 PB IBM Spectrum Scale
  file system transferring

#	Site	Manufacturer	Computer	Country	Cores	Rmax [Pflops]	Power [MW]
1	Oak Ridge National Laboratory	IBM	Summit IBM Power System, P9 22C 3.07GHz, Mellanox EDR, NVIDIA GV100	USA	2,282,544	122.3	8.8
2	National Supercomputing Center in Wuxi	NRCPC	Sunway TaihuLight NRCPC Sunway SW26010, 260C 1.45GHz	China	10,649,600	93.0	15.4
3	Lawrence Livermore National Laboratory	IBM	Sierra IBM Power System, P9 22C 3.1GHz, Mellanox EDR, NVIDIA GV100	USA	1,572,480	71.6	
4	National University of Defense Technology	NUDT	Tianhe-2A ANUDT TH-IVB-FEP, Xeon 12C 2.2GHz, Matrix-2000	China	4,981,760	61.4	18.5
5	National Institute of Advanced Industrial Science and Technology	Fujitsu	AI Bridging Cloud Infrastructure (ABCI) PRIMERGY CX2550 M4, Xeon Gold 20C 2.4GHz, IB-EDR, NVIDIA V100	Japan	391,680	19.9	1.65
6	Swiss National Supercomputing Centre (CSCS)	Cray	Piz Daint Cray XC50, Xeon E5 12C 2.6GHz, Aries, NVIDIA Tesla P100	Switzerland	361,760	19.6	2.27
7	Oak Ridge National Laboratory	Cray	Titan Cray XK7, Opteron 16C 2.2GHz, Gemini, NVIDIA K20x	USA	560,640	17.6	8.21
8	Lawrence Livermore National Laboratory	IBM	Sequoia BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	1,572,864	17.2	7.89
9	Los Alamos NL / Sandia NL	Cray	Trinity Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries	USA	979,968	14.1	3.84
10	Lawrence Berkeley National Laboratory	Cray	Cori Cray XC40, Intel Xeons Phi 7250 68C 1.4 GHz, Aries	USA	622,336	14.0	3.94

AVERAGE SYSTEM AGE

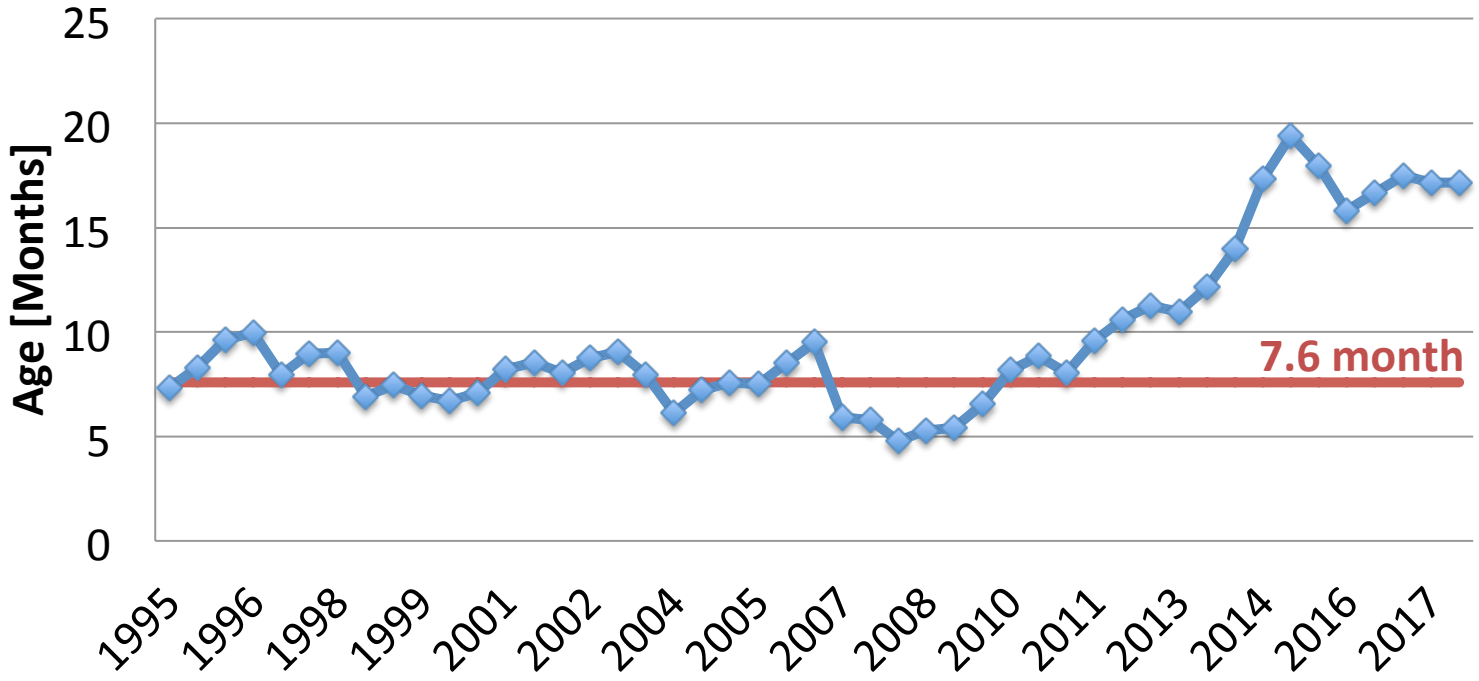

RANK AT WHICH HALF OF TOTAL PERFORMANCE IS ACCUMULATED

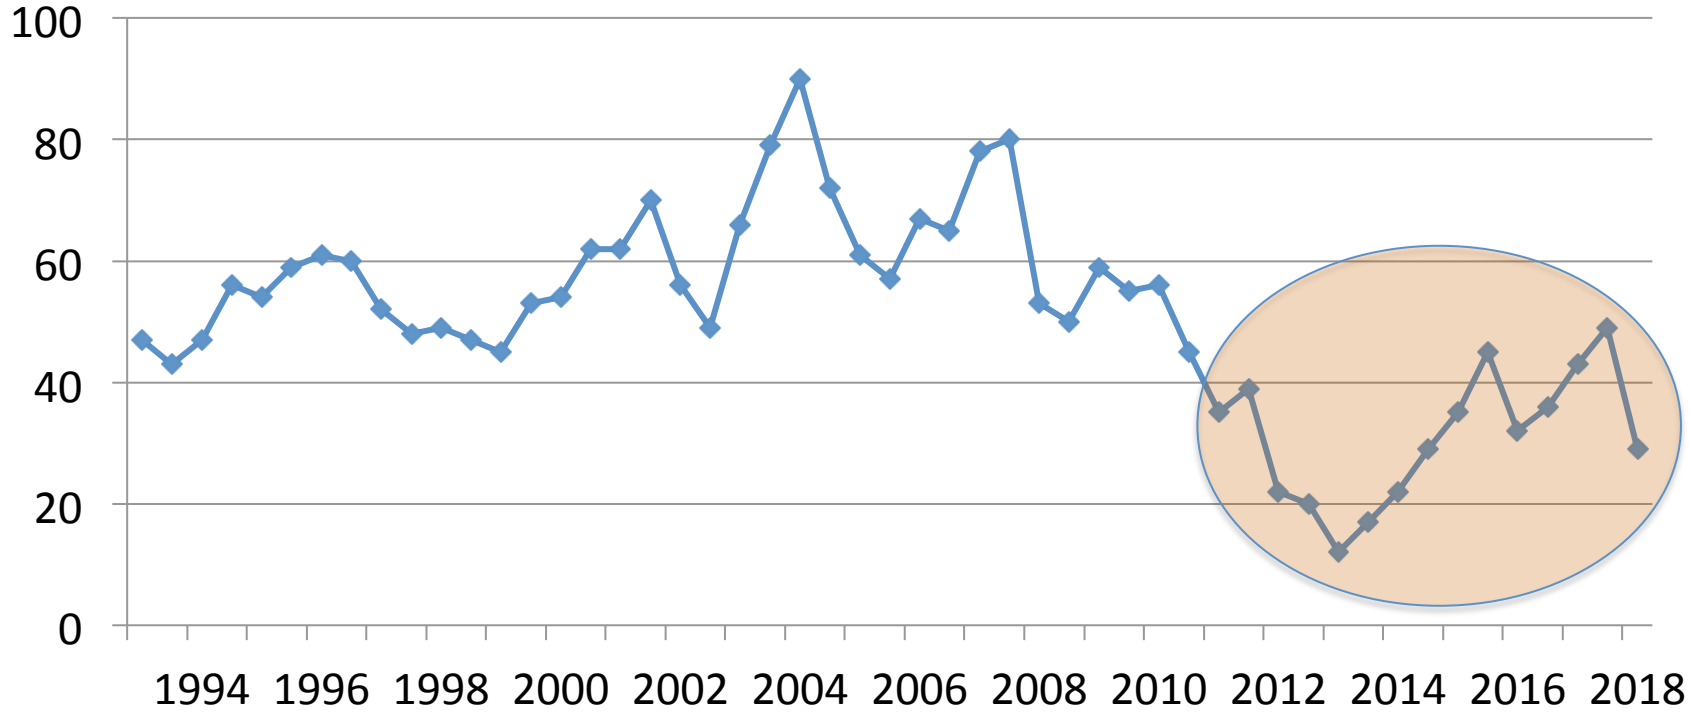

PERFORMANCE DEVELOPMENT

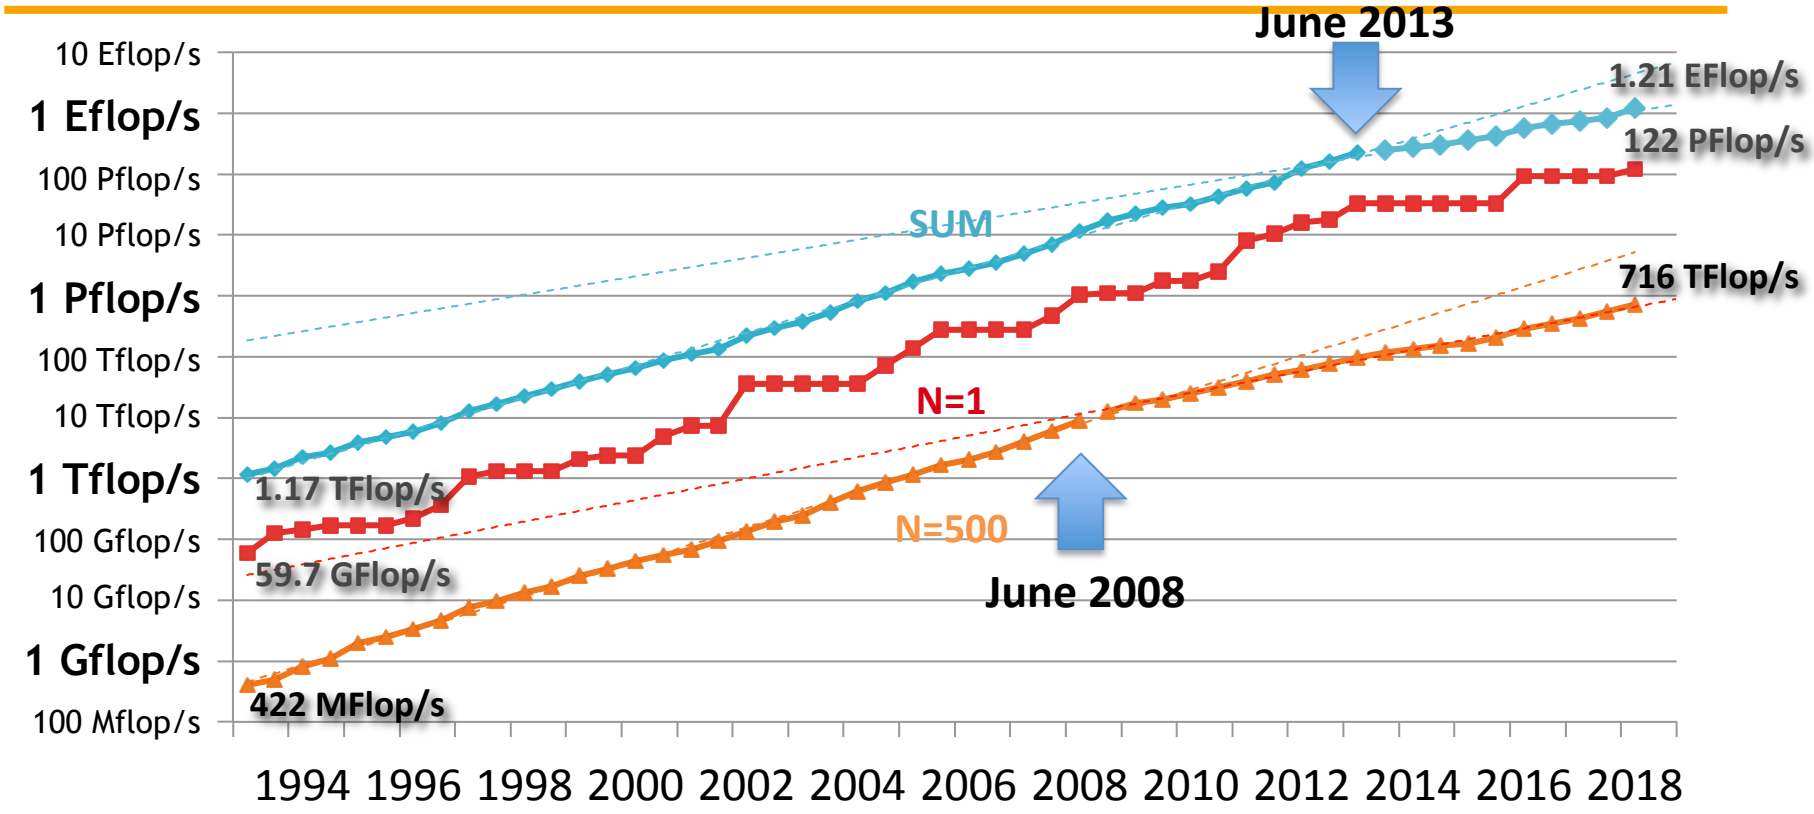

ANNUAL PERFORMANCE INCREASE OF THE TOP500

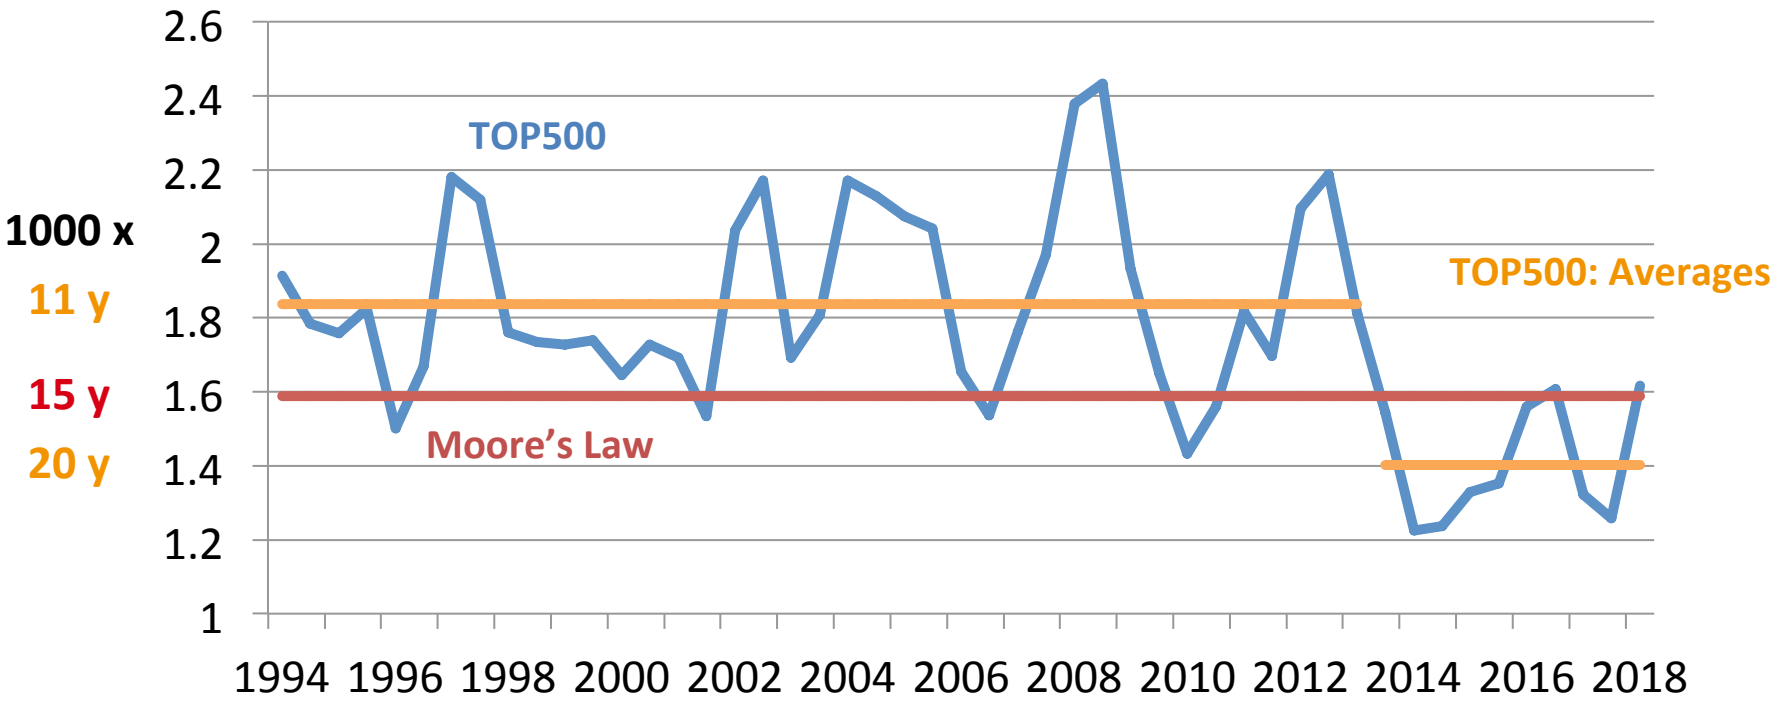

25 Years – 51 Editions

How do we integrate/aggregate over time/editions?

- System counts are focused on the low end
- Moore's Law overpowers everything performance based
- **Normalize each list by average performance**
  - HPL not Peak
  - Average performance not max (#1) or min

  - **Norm-HPL**

- Add up contributions from various lists over time
  - Each full lists contributed a total of 500
  - 51 lists together have a total weight of 25,500

  - **Σ Norm-HPL**

No. 1's – HPL R_max

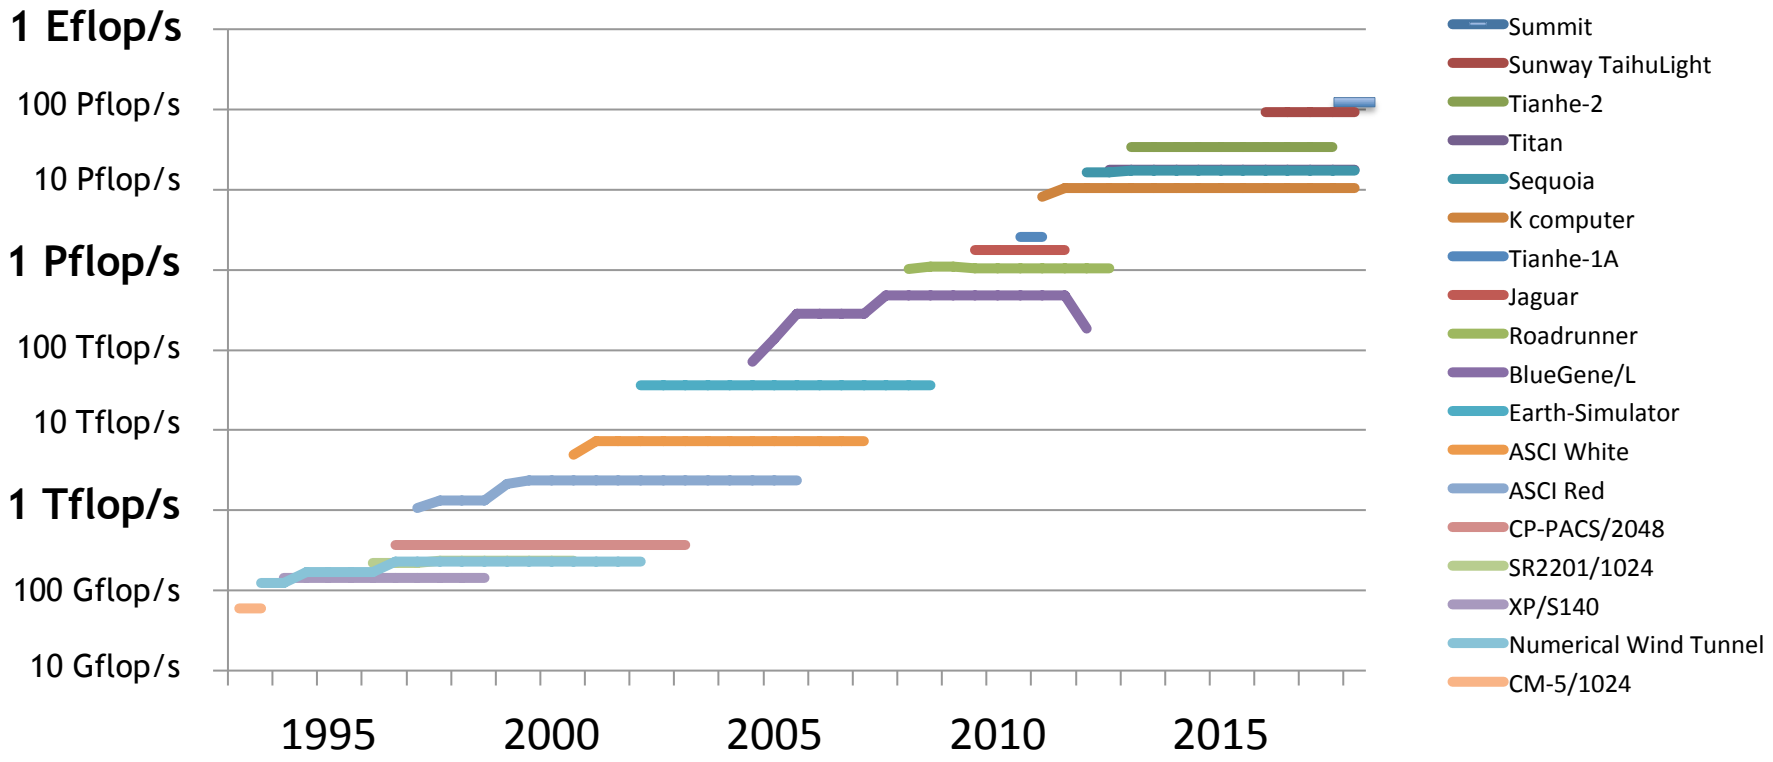

1's – Accumulated Norm-HPL

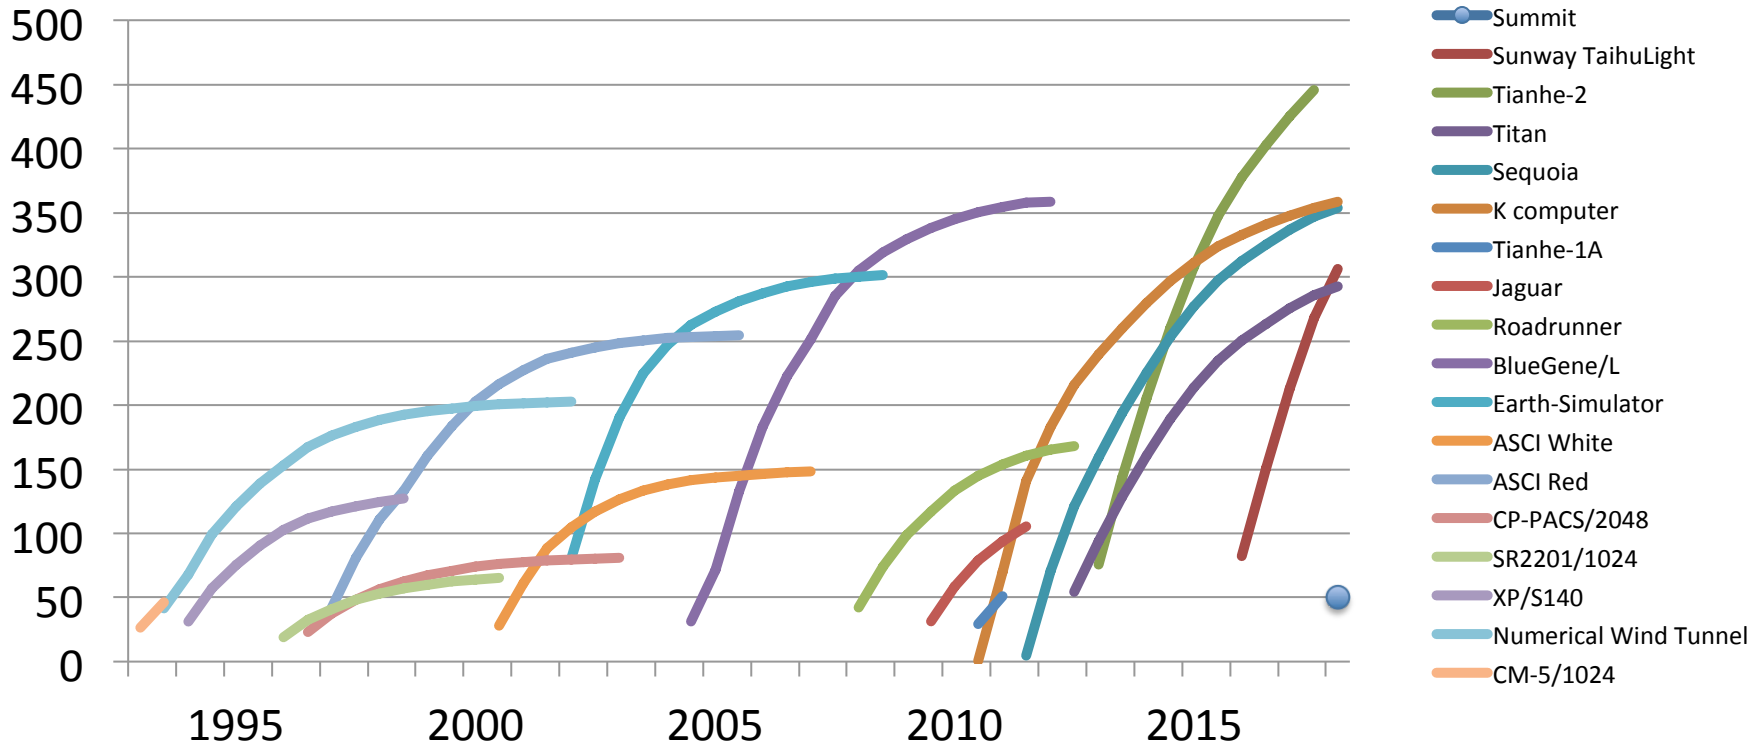

Dominant Sites

#	Site	Country	Norm-HPL
1	LLNL	USA	1,504
2	LANL	USA	816
3	ORNL	USA	795
4	SNL	USA	658
5	NSCC Guangzhou	China	474
6	RIKEN AICS	Japan	361
7	NASA/Ames	USA	356
8	FZ Jülich	Germany	339
9	JAMSTEC	Japan	325
10	NERSC / LBNL	USA	310
11	ANL	USA	305

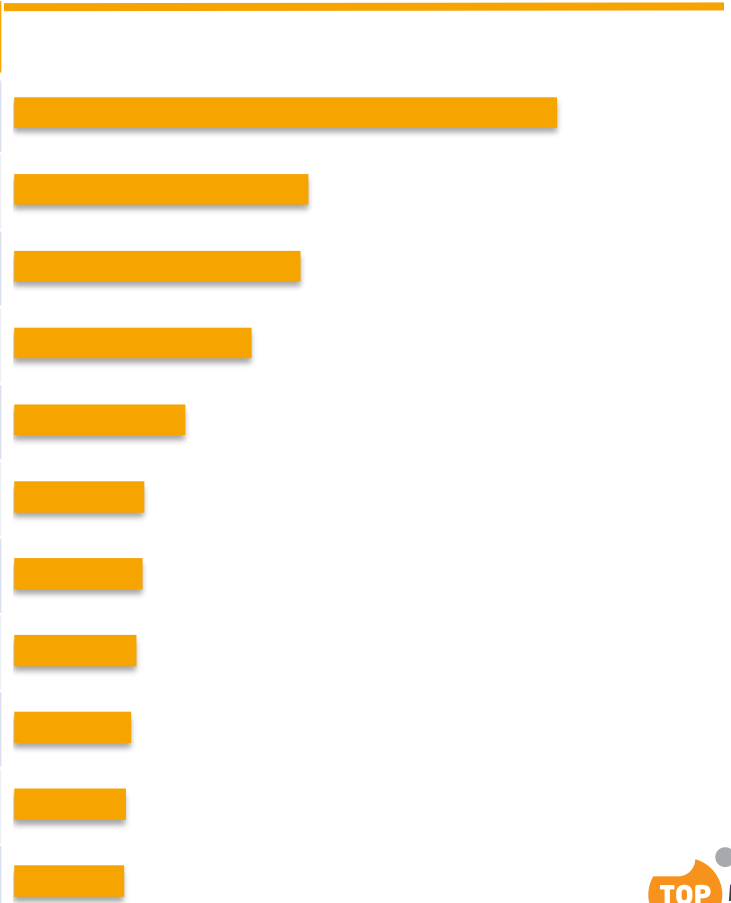

COUNTRIES

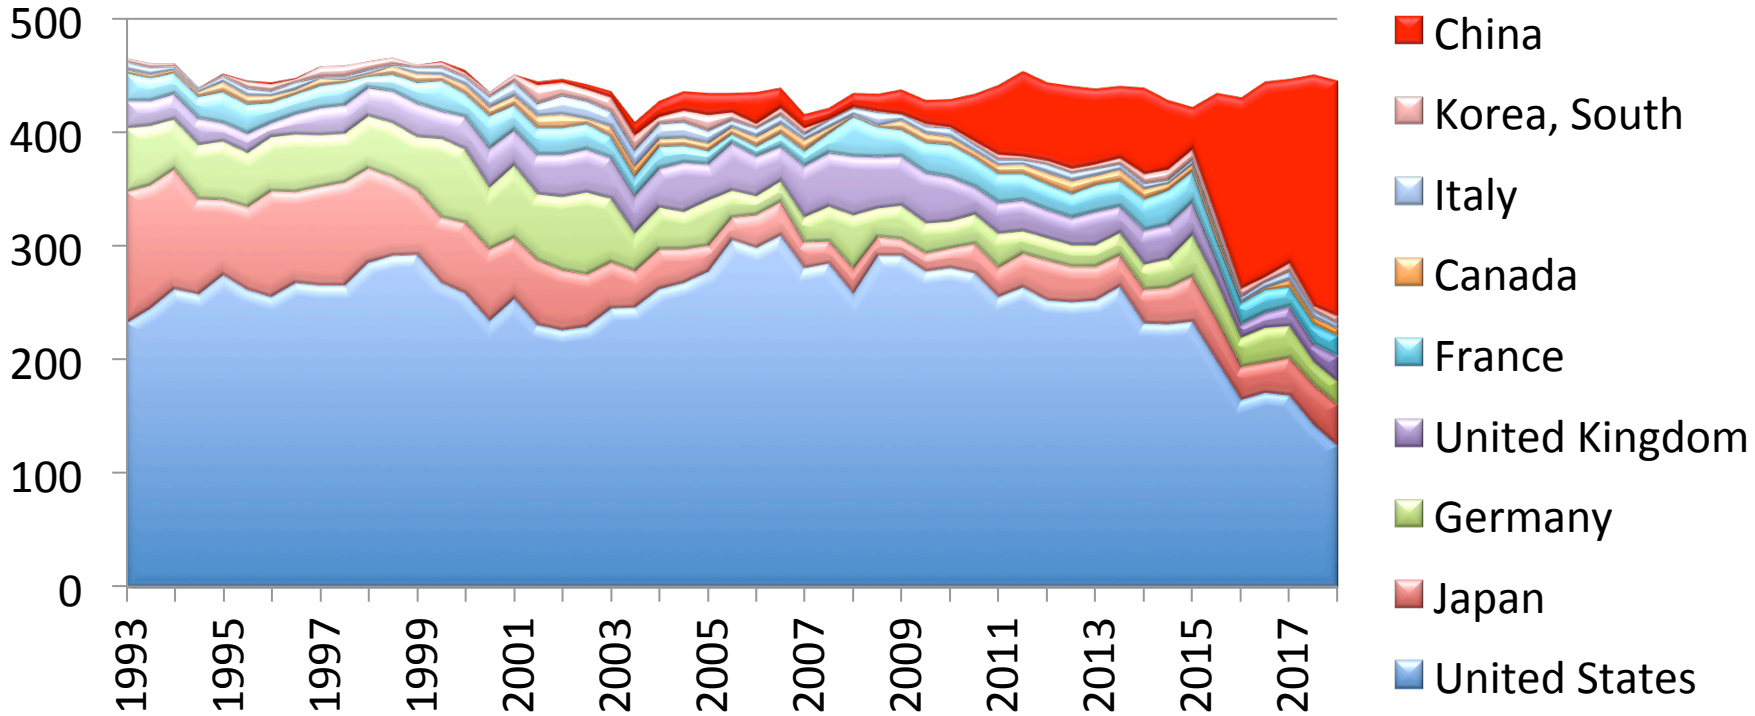
500
400
300
200
100
0
1993
1995
1997
1999
2001
2003
2005
2007
2009
2011
2013
2015
2017
China
Korea, South
Italy
Canada
France
United Kingdom
Germany
Japan
United States

COUNTRIES / SYSTEM SHARE

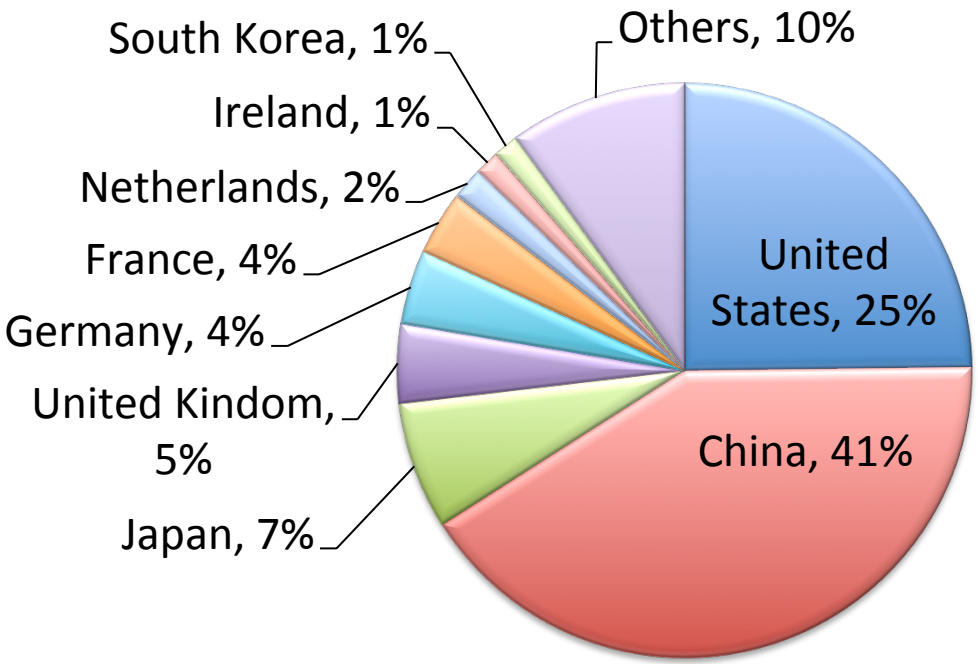

PRODUCERS

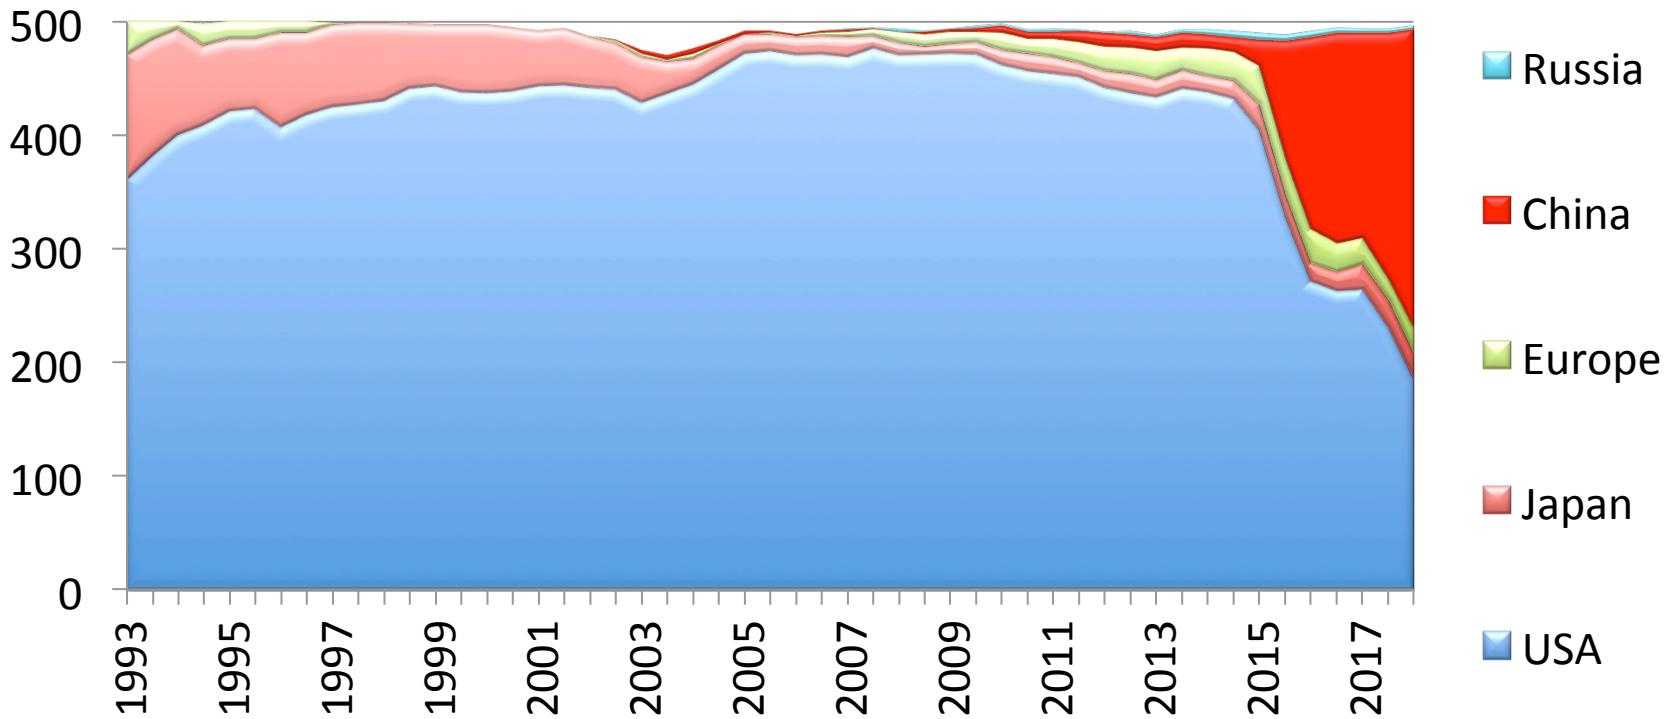
500
400
300
200
100
0
1993
1995
1997
1999
2001
2003
2005
2007
2009
2011
2013
2015
2017
Russia
China
Europe
Japan
USA

VENDORS / SYSTEM SHARE

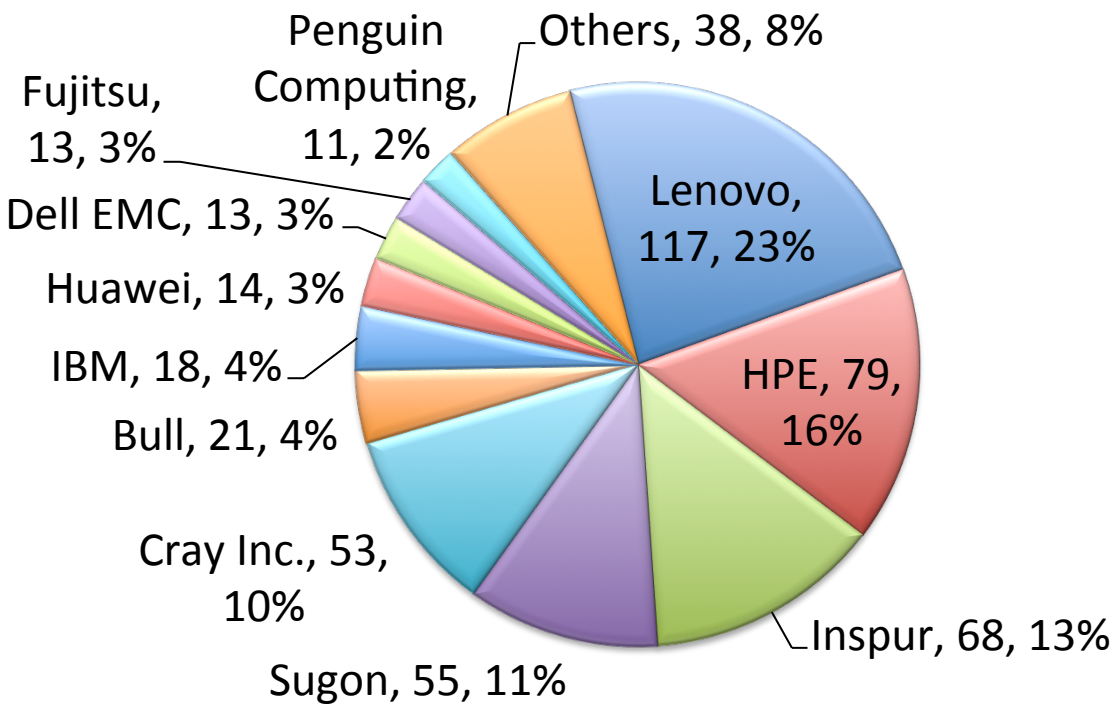

# of systems, % of 500

VENDORS / PERFORMANCE SHARE

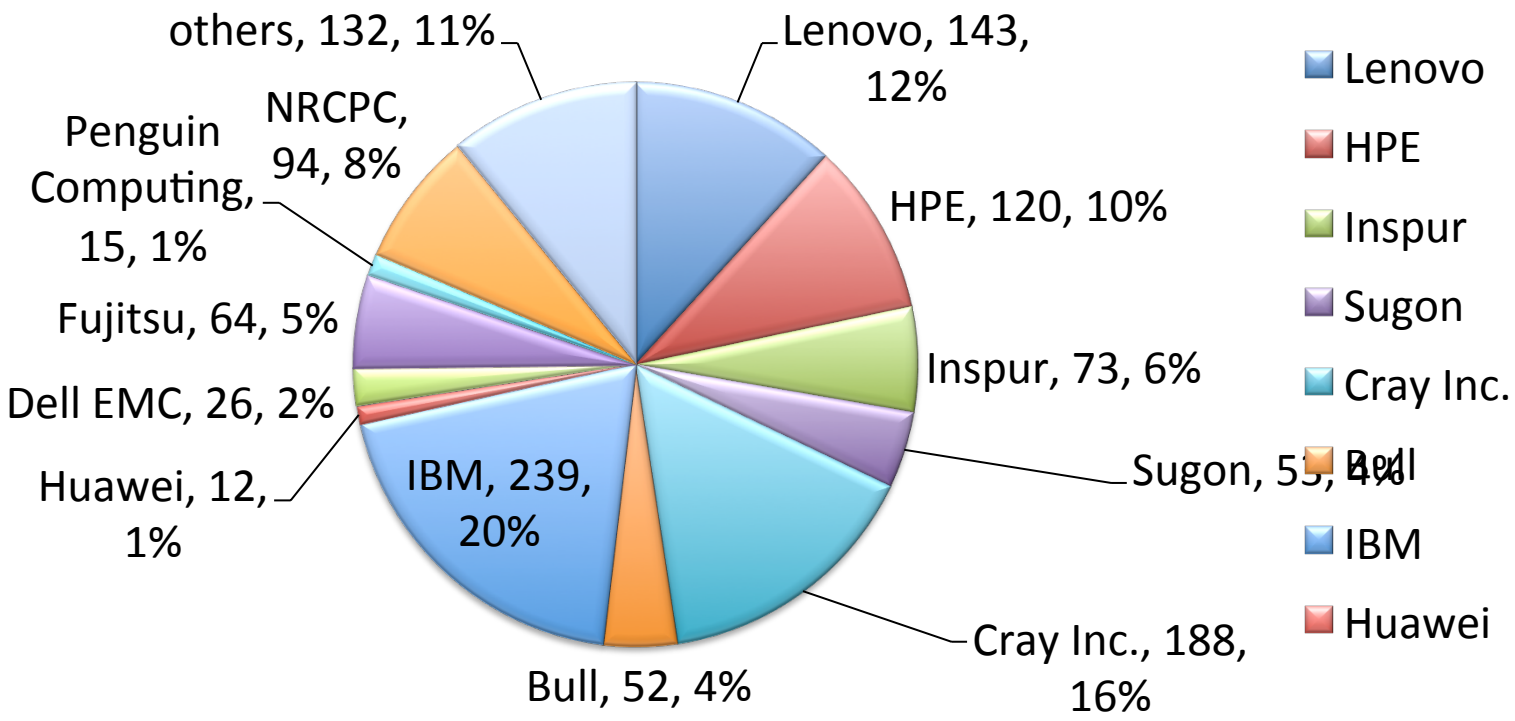

Sum of Pflop/s, % of whole list

ACCELERATORS

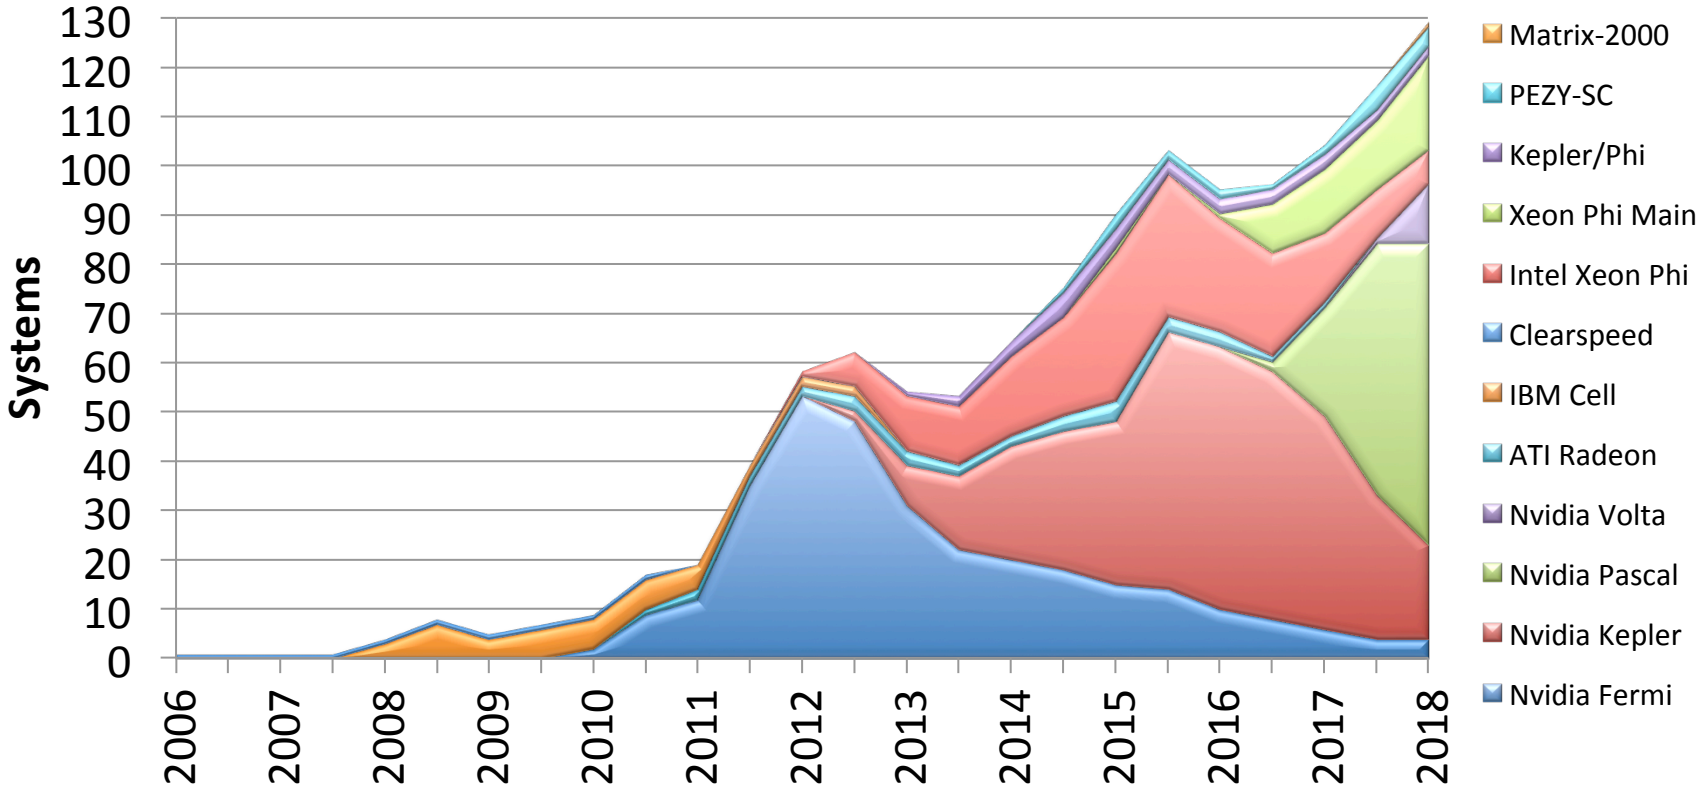

PERFORMANCE SHARE OF ACCELERATORS

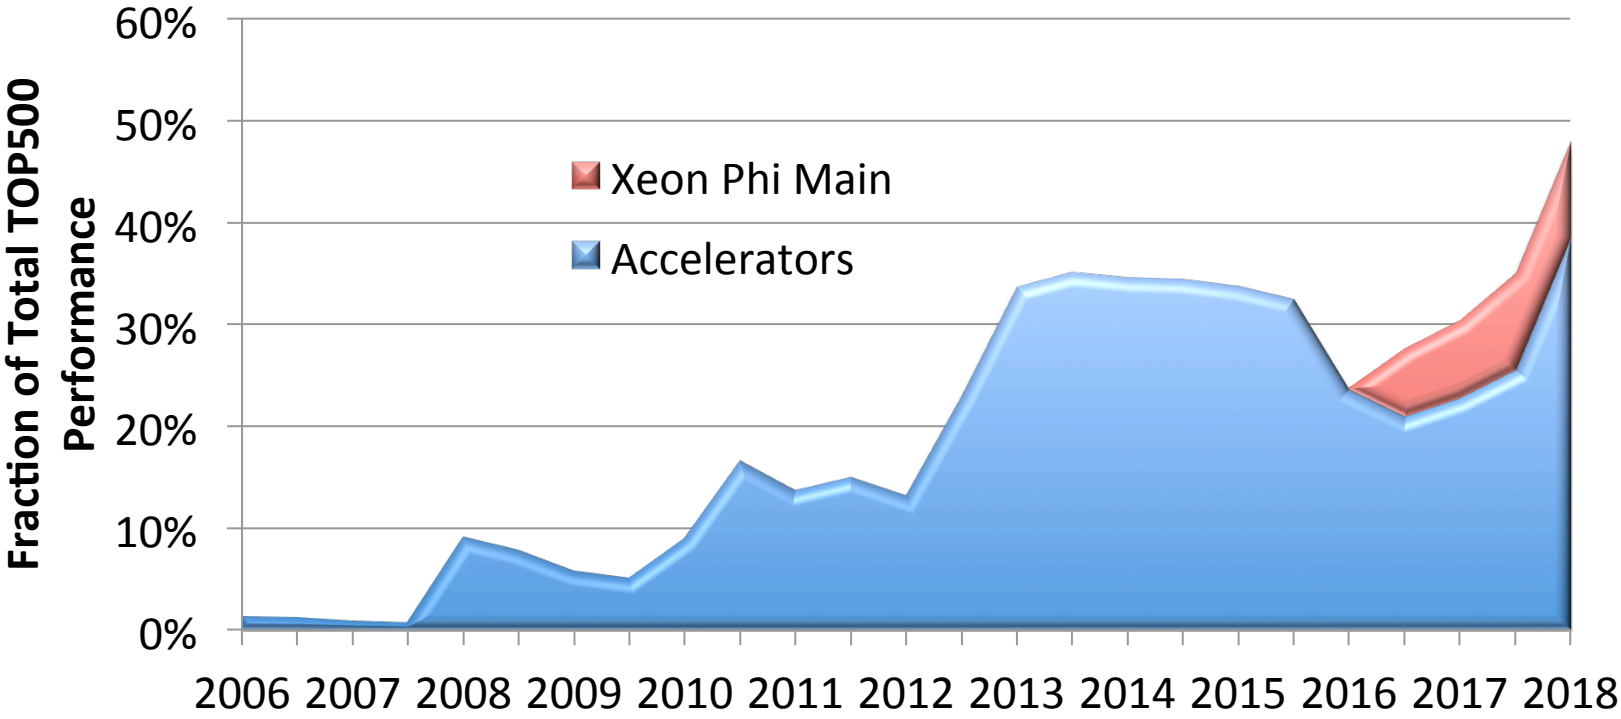

MOST ENERGY EFFICIENT ARCHITECTURES

Computer				Rmax/ Power
Shoubou system B, ZettaScaler-2.2	Xeon 16C 1.3GHz	Infiniband EDR	PEZY-SC2	**18.4**
Suiren2, ZettaScaler-2.2	Xeon 16C 1.3GHz	Infiniband EDR	PEZY-SC2	**16.8**
Sakura, ZettaScaler-2.2	Xeon 8C 2.3GHz	Infiniband EDR	PEZY-SC2	**16.7**
DGX Saturn V, NVIDIA DGX-1 Volta36	Xeon 20C 2.2GHz	Infiniband EDR	Tesla V100	**15.1***
Summit, IBM Power System	Power9 22C 3.07GHz	Infiniband EDR	Volta GV100	**13.9**
Tsubame 3.0, SGI ICE XA	Xeon 14C 2.4GHz	Intel Omni-Path	Tesla P100 SXM2	**13.7***
AIST AI Cloud, NEC 4U-8GPU	Xeon 10C 1.8GHz	Infiniband EDR	Tesla P100 SXM2	**12.7**
AI Bridging Cloud Infrastructure (ABCI), Fujitsu PRIMERGY, NVIDIA Tesla V100	Xeon Gold 20C 2.4GHz	Infiniband EDR	Tesla V100 SXM2	**12.1**
MareNostrum P9 CTE, IBM Power System	Power9 22C 3.1GHz	Infiniband EDR	Tesla V100	**11.9**
Wilkes-2, Dell C4130	Xeon 12C 2.2GHz	Infiniband EDR	Tesla P100	**10.4**

* Efficiency based on Power optimized HPL runs of equal size to TOP500 run.

[Gflops/Watt]

POWER EFFICIENCY

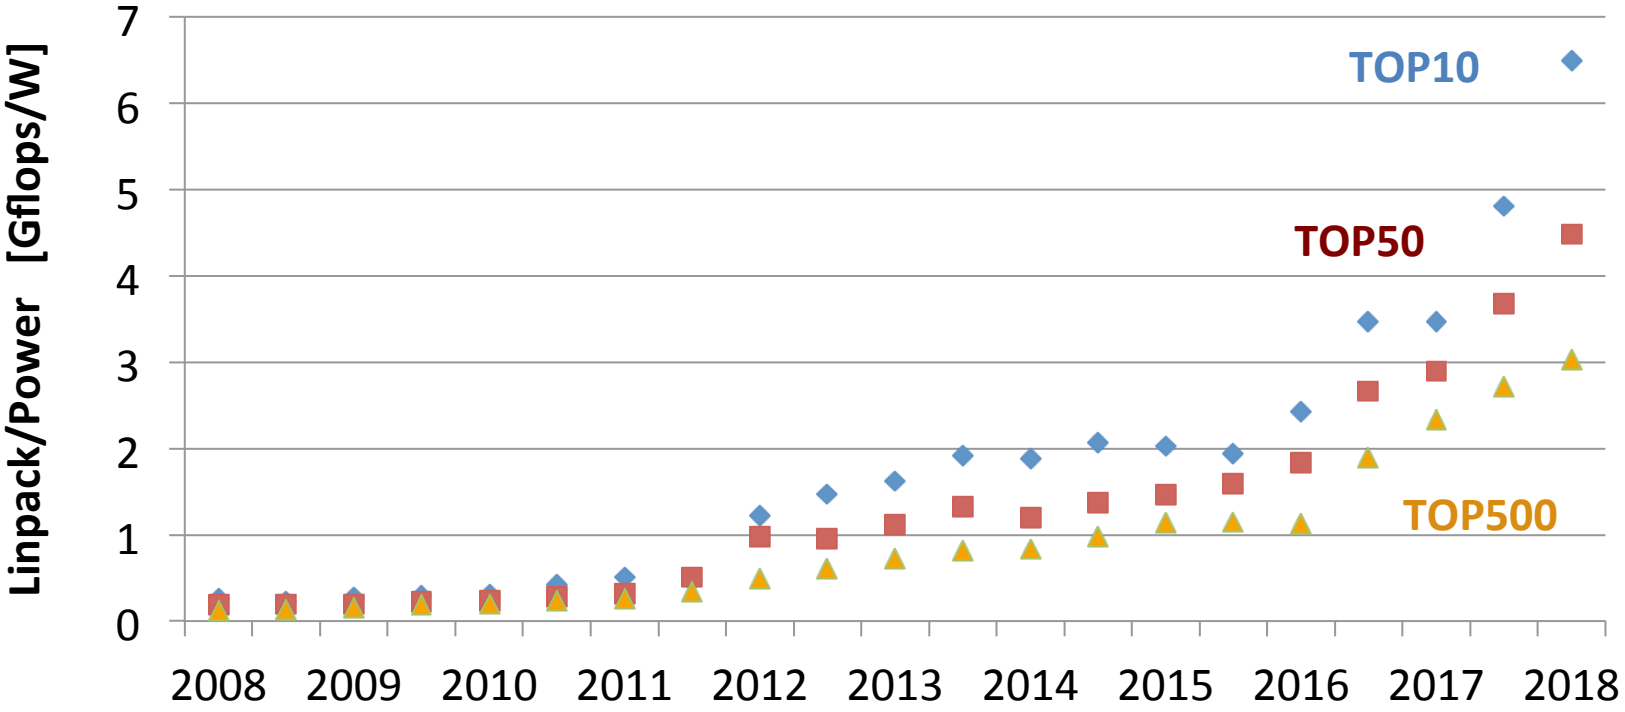

ENERGY EFFICIENCY

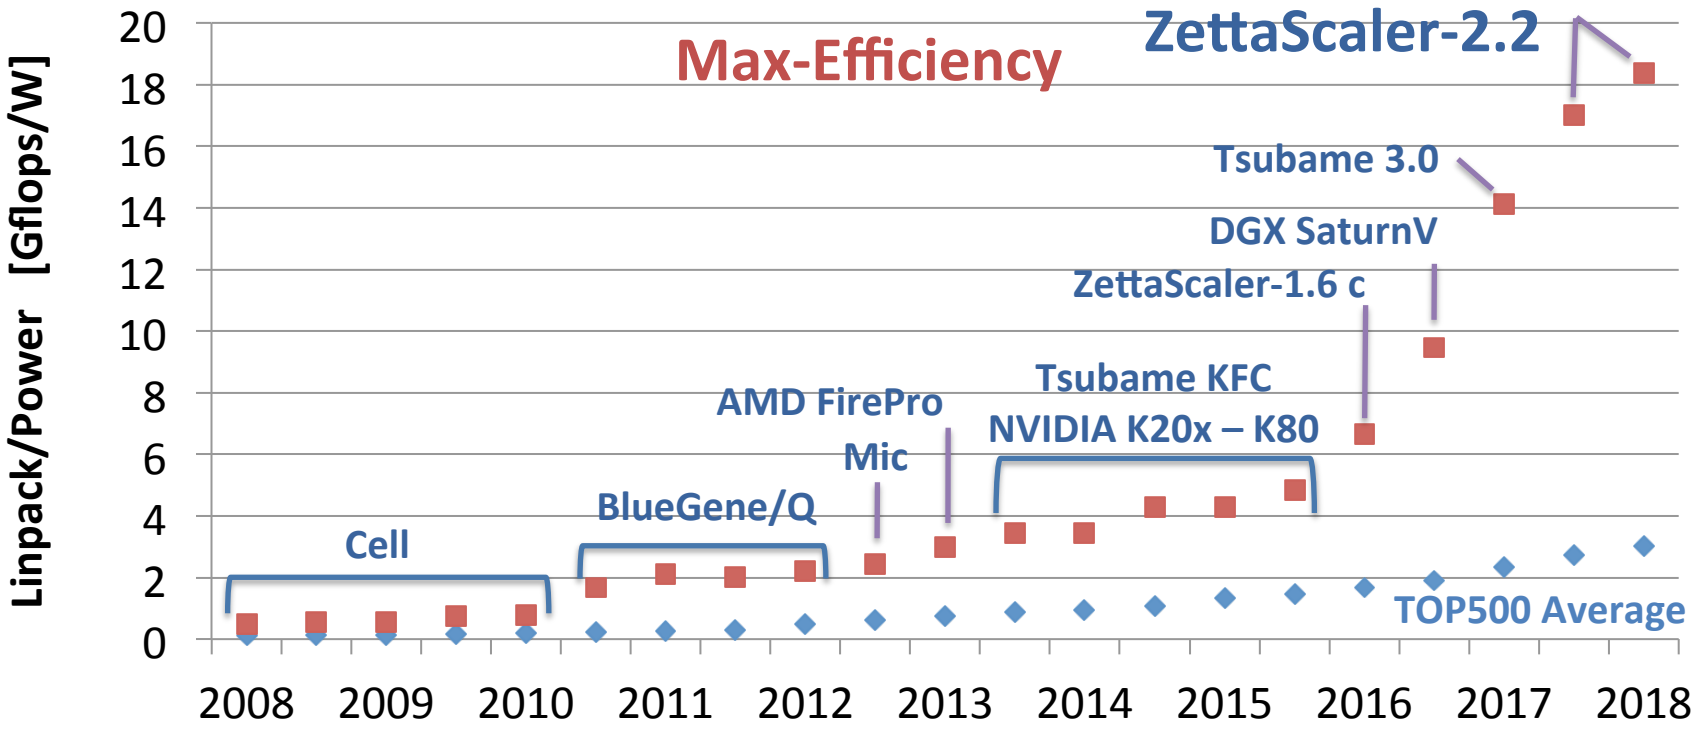

#	T	Site	Manufacturer	Computer	Country	HPCG [Pflop/s]	Rmax [Pflop/s]	HPCG/ Peak	HPCG/ HPL
1	1	Oak Ridge National Laboratory	IBM	Summit IBM Power System, P9 22C 3.07 GHz, Volta GV100, EDR	USA	2.9258	122.3	1.6%	2.4%
2	3	Lawrence Livermore National Laboratory	IBM	Sierra IBM Power System, P9 22C 3.1 GHz, Volta GV100, EDR	USA	1.7957	71.6	1.5%	2.5%
3	16	RIKEN Advanced Institute for Computational Science	Fujitsu	K Computer SPARC64 VIIIfx 2.0GHz, Tofu Interconnect	Japan	0.6027	10.5	5.3%	5.7%
4	9	Los Alamos NL / Sandia NL	Cray	Trinity Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Aries	USA	0.5461	14.1	1.2%	3.9%
5	6	Swiss National Supercomputing Centre (CSCS)	Cray	Piz Daint Cray XC50, Xeon E5 12C 2.6GHz, Aries, NVIDIA Tesla P100	Switzerland	0.4864	19.6	1.9%	2.5%
6	2	National Supercomputing Center in Wuxi	NRCPC	Sunway TaihuLight NRCPC Sunway SW26010, 260C 1.45GHz	China	0.4808	93.0	0.4%	0.5%
7	12	JCAHPC Joint Center for Advanced HPC	Fujitsu	Oakforest-PACS PRIMERGY CX1640 M1, Intel Xeons Phi 7250 68C 1.4 GHz, OmniPath	Japan	0.3855	13.6	1.5%	2.8%
8	10	Lawrence Berkeley National Laboratory	Cray	Cori Cray XC40, Intel Xeons Phi 7250 68C 1.4 GHz, Aries	USA	0.3554	14.0	1.3%	2.5%
9	14	Commissariat a l'Energie Atomique (CEA)	Bull	Tera-1000-2 Bull Sequana X1000, Intel Xeon Phi 7250 68C 1.4 GHz, Bull BXI 1.2	France	0.3338	12.0	1.4%	2.8%
10	8	Lawrence Livermore National Laboratory	IBM	Sequoia BlueGene/Q, Power BQC 16C 1.6GHz, Custom	USA	0.3304	17.2	1.6%	1.9%

ISC18 TOP500 HIGHLIGHTS

- ORNL's Summit is new #1 (IBM, NVIDIA, Mellanox).
- Four 'new' system in the TOP5! (Summit, Sierra, Tianhe-2A, ABCI)
- Slow-down in performance growth since 2013 goes hand in hand with
  - Longer system usage (~2x) and
  - Concentration of capabilities at the top (relatively larger top systems)
- Lenovo is first Chinese manufacturer to sell systems in numbers outside of China (everywhere) (China: 20, USA: 21, ROW: 23).
- Accelerated system get finally adopted by industrial users (25% of new systems in November + June).
- Summit and Sierra are the first systems to achieve over 1 Pflop/s on HPCG.