

Lecture 20

Top 500

EN 600.320/420/620

Instructor: Randal Burns


12 March 2019



Department of Computer Science, *Johns Hopkins University*

Images in lecture from
https://www.top500.org/static/media/uploads/top500_ppt_201806.pdf

TOP500



http://ornl.gov/info/ornlreview/v45_3_12/images/a04_p10_lrg.jpg


- Fastest 500 computing machines
 - According to LINPACK benchmark
- Controversial
 - Other metrics represent other problems
 - Green500: best computing power per watt
- Currently "Summit" ORNL, US
 - 122 TFLOP/s, 2.2M cores, 8.8 kW
 - IBM Power 8 + NVidia V100 GPUs
- #2 "Sunway TaihuLight" Xiwu Chine
 - 93 TFlops, 10,649,000 cores, 15,371 kW
 - All Chinese hardware



Top500 Growth (Exascale coming)

https://www.top500.org/static/media/uploads/top500_ppt_201806.pdf

PERFORMANCE DEVELOPMENT




Top500: Flops

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	DOE/SC/Oak Ridge National Laboratory United States	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM	2,282,544	122,300.0	187,659.3	8,806
2	National Supercomputing Center in Wuxi China	Sunway TaihuLight - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRCPC	10,649,600	93,014.6	125,435.9	15,371
3	DOE/NNSA/LLNL United States	Sierra - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM / NVIDIA / Mellanox	1,572,480	71,610.0	119,193.6	
4	National Super Computer Center in Guangzhou China	Tianhe-2A - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000 NUDT	4,981,760	61,444.5	100,678.7	18,482




Top500 and Moore's Law

ANNUAL PERFORMANCE INCREASE OF THE TOP500




Top 500 Evolution




Who has supercomputers?

COUNTRIES



The Rise of Accelerators

ACCELERATORS



Top 500 Conclusions

- Exaflop in reach
- Accelerators are ubiquitous
- Despite US reemergence, China dominates
 - 40% of systems
 - Chinese made hardware



Green 500

- GFlops/watt

- Power efficiency will be the limited factor in reaching exascale computing
- Also co-branded as environmentally friendly



MOST ENERGY EFFICIENT ARCHITECTURES



Computer				Rmax/ Power
Shoubou system B , ZettaScaler-2.2	Xeon 16C 1.3GHz	Infiniband EDR	PEZY-SC2	18.4
Suirens2 , 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



What's a Zettascaler??

- PEZY-SCx accelerator
 - 3 level hierarchy of computer 128 cities x 4 villages x 4 cores
 - L2 shared @ "city"
 - L1 shared @ "village"
- ExaScaler fluoro-carbon based liquid cooling
- TCI stack die RAMs

All of this creates remarkable density

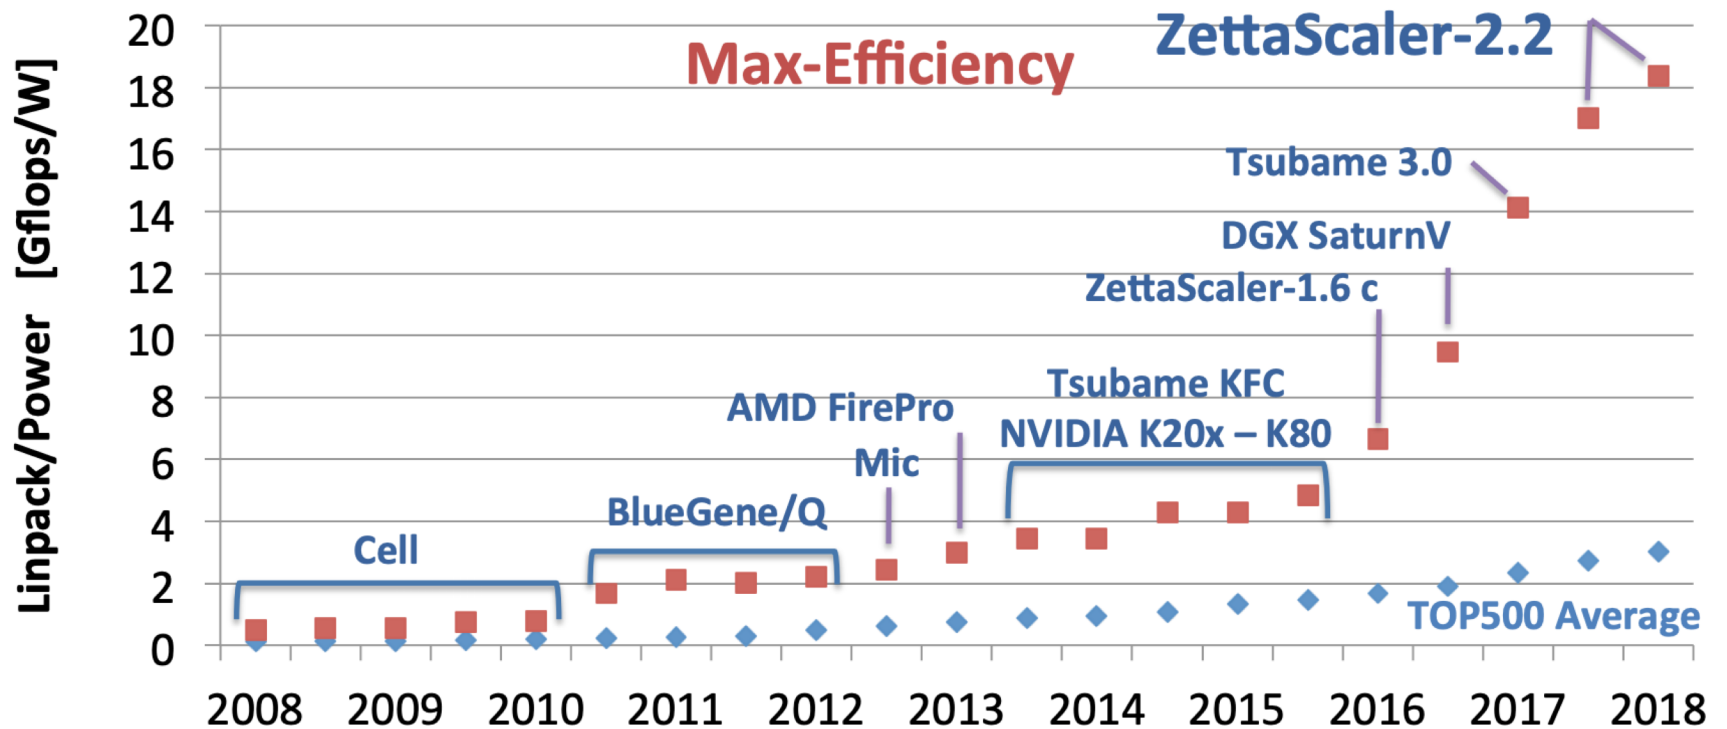


Green Trends Architectures

The
GREEN
500

ENERGY EFFICIENCY

TOP 500




What else is going on?

- Convergence between TOP500 and Green 500
 - Power and cooling density limits overall performance
 - Can't make a fast machine that's not efficient
- T#1 = G#5

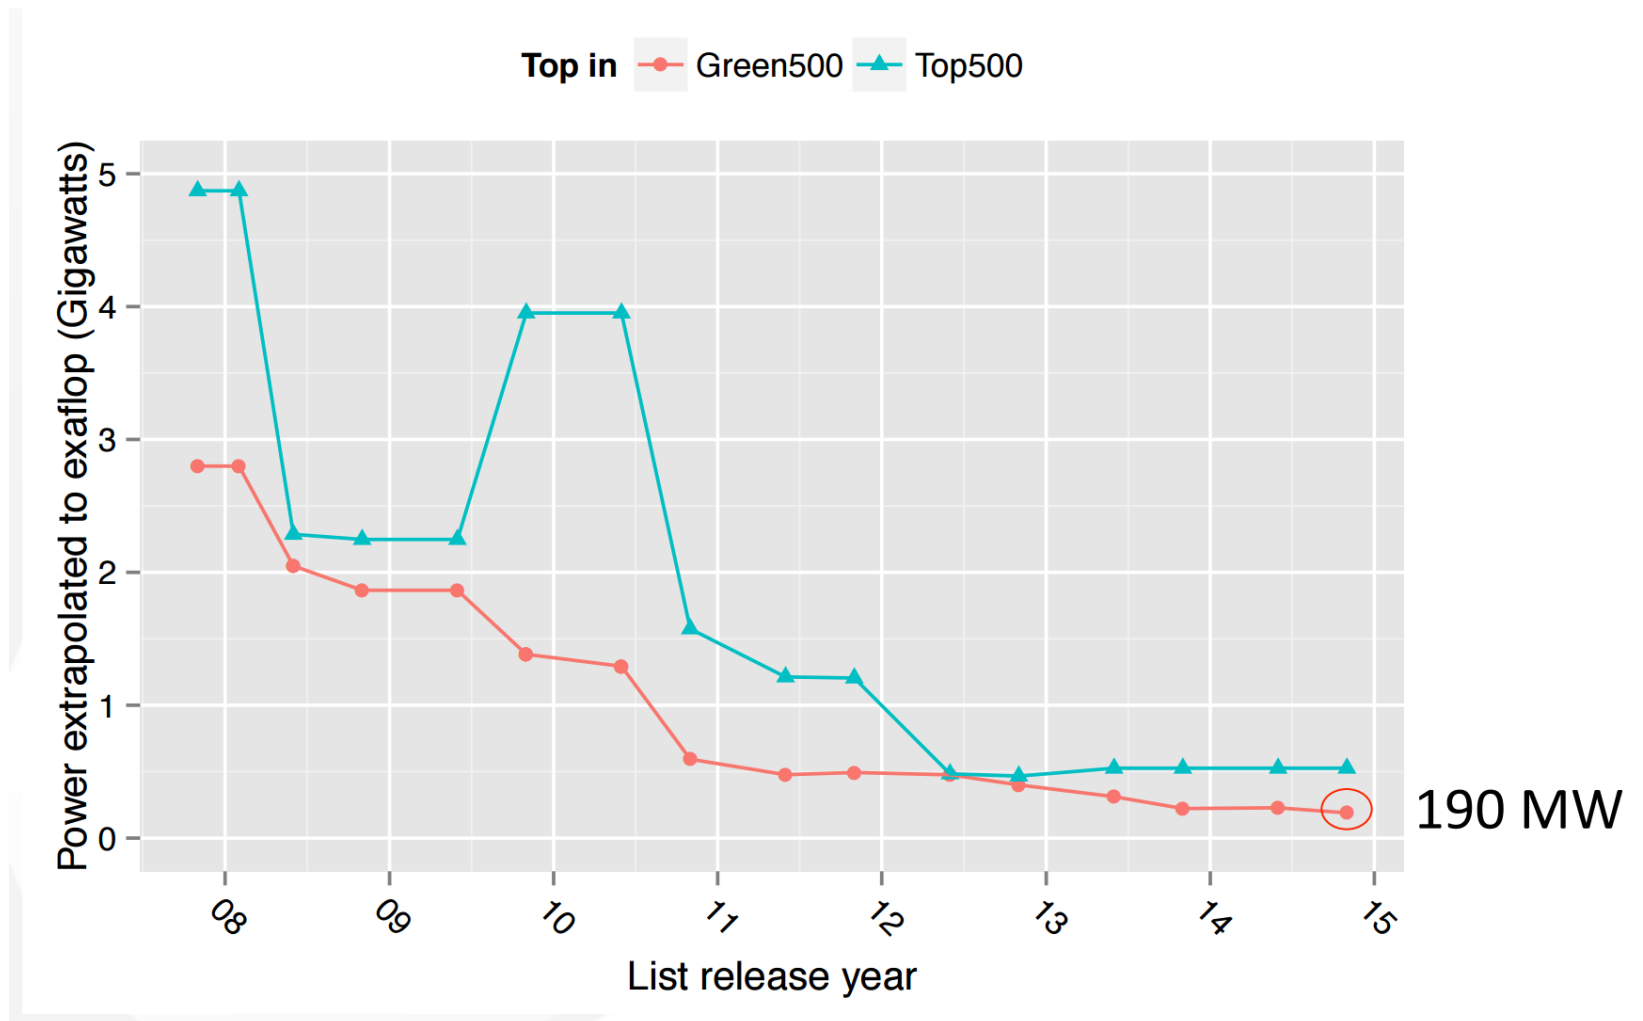


Green Trends Efficiency

POWER EFFICIENCY



Exascale at 25 MW?



The Green500 BoF, SC|14, Nov. 2014
POC: info@green500.org



Final Comments

- Multiple Chinese manufacturers
- Lenovo doing substantial business outside China
 - 21 in US, 23 in rest of world
- Slow down in performance growth
 - More focus at top machines
 - Longer lifetime on TOP500 (2x)
 - I'm not totally convinced that this is not (death of) Moore's law related?

