


# Lecture 25

## Top 500

EN 600.320/420/620

Instructor: Randal Burns


6 December 2021



Department of Computer Science, *Johns Hopkins University*

<https://www.top500.org/lists/top500/2021/11/>

# TOP500



[http://ornl.gov/info/ornlreview/v45\\_3\\_12/images/a04\\_p10\\_lrg.jpg](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
  - Does not measure I/O to storage
  - Does measure networking between nodes (needed to succeed)



Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	RIKEN Center for Computational Science Japan	<b>Supercomputer Fugaku</b> - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D Fujitsu	7,630,848	442,010.0	537,212.0	29,899
2	DOE/SC/Oak Ridge National Laboratory United States	<b>Summit</b> - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM	2,414,592	148,600.0	200,794.9	10,096
3	DOE/NNSA/LLNL United States	<b>Sierra</b> - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband IBM / NVIDIA / Mellanox	1,572,480	94,640.0	125,712.0	7,438
4	National Supercomputing Center in Wuxi China	<b>Sunway TaihuLight</b> - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway NRCPC	10,649,600	93,014.6	125,435.9	15,371
5	DOE/SC/LBNL/NERSC United States	<b>Perlmutter</b> - HPE Cray EX235n, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 SXM4 40 GB, Slingshot-10 HPE	761,856	70,870.0	93,750.0	2,589



# Top500 Growth (Exascale coming)

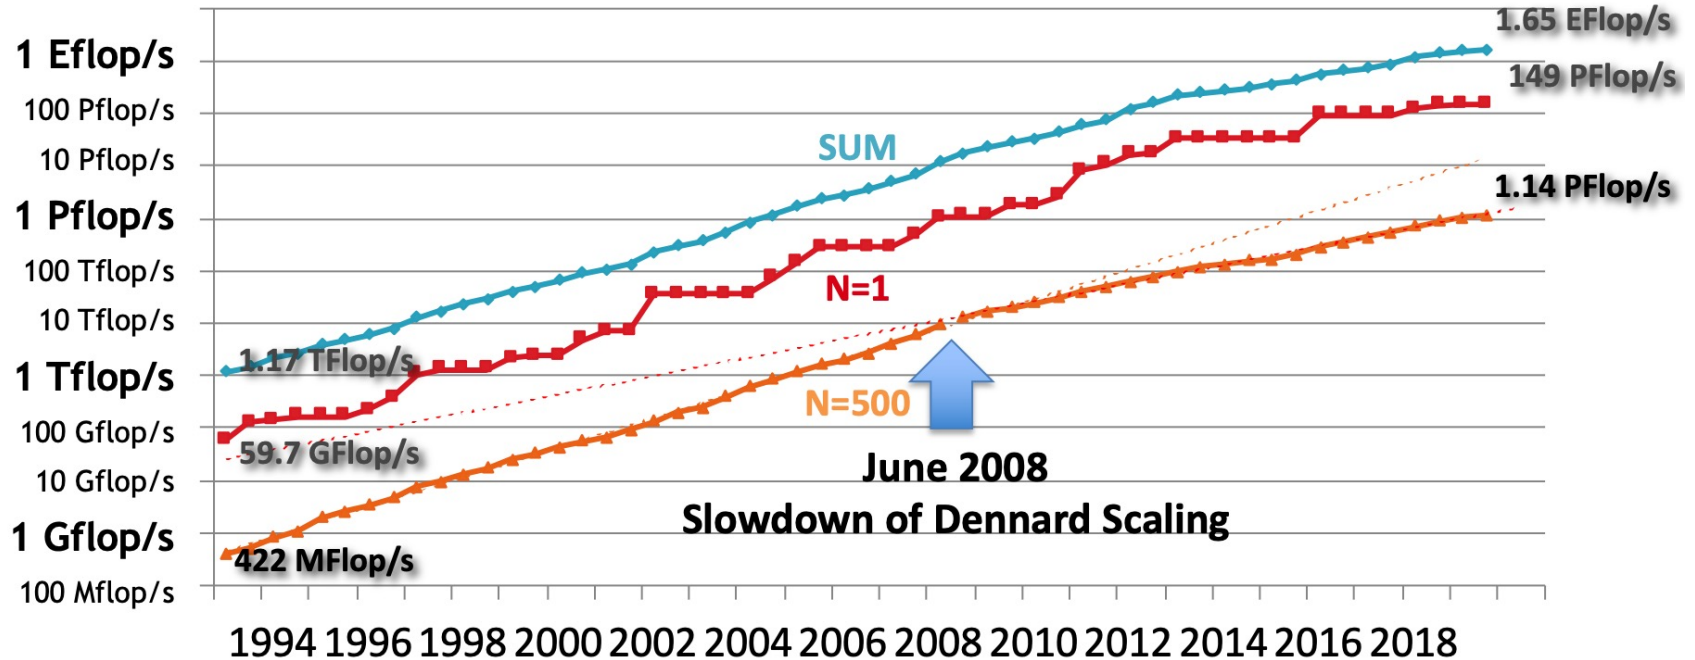


- Tracking Moore's law pretty well
- Top computer always represents a large fraction of total FLOPS/s



# Top500 Growth (Exascale coming)

[https://www.top500.org/static/media/uploads/top500\\_ppt\\_201911.pdf](https://www.top500.org/static/media/uploads/top500_ppt_201911.pdf)




- Slide from 2 years ago connects to Dennard scaling



# CPUs

- Not accelerators, which hold most of the FLOPs, since 2012)
- Proprietary = Sunway, Fujitsu




# Newish Processors

- Fujitsu A64FX (2019): ARM for HPC
  - 48 compute cores
  - 2 or 4 assistant cores
- Sunway SW26010 (2018)
  - RISC 260-core manycore



# Accelerators/Co-processors





# Newish Accelerators


- PezySC (2018)
  - Hierarchy of 128 cities (shared L2) x 4 villages (shared L1) x 4 cores
  - TCI stack die RAMs
- MN-Core (2021): deep-learning, neural processor
- Matrix 2000: 128 core many-core
  - Replacement for Intel Xeon Phi on Tianhe-2
- NVidia A100 (2020): tensor cores added to GPU

Accelerators are the solution to Dennard scaling.



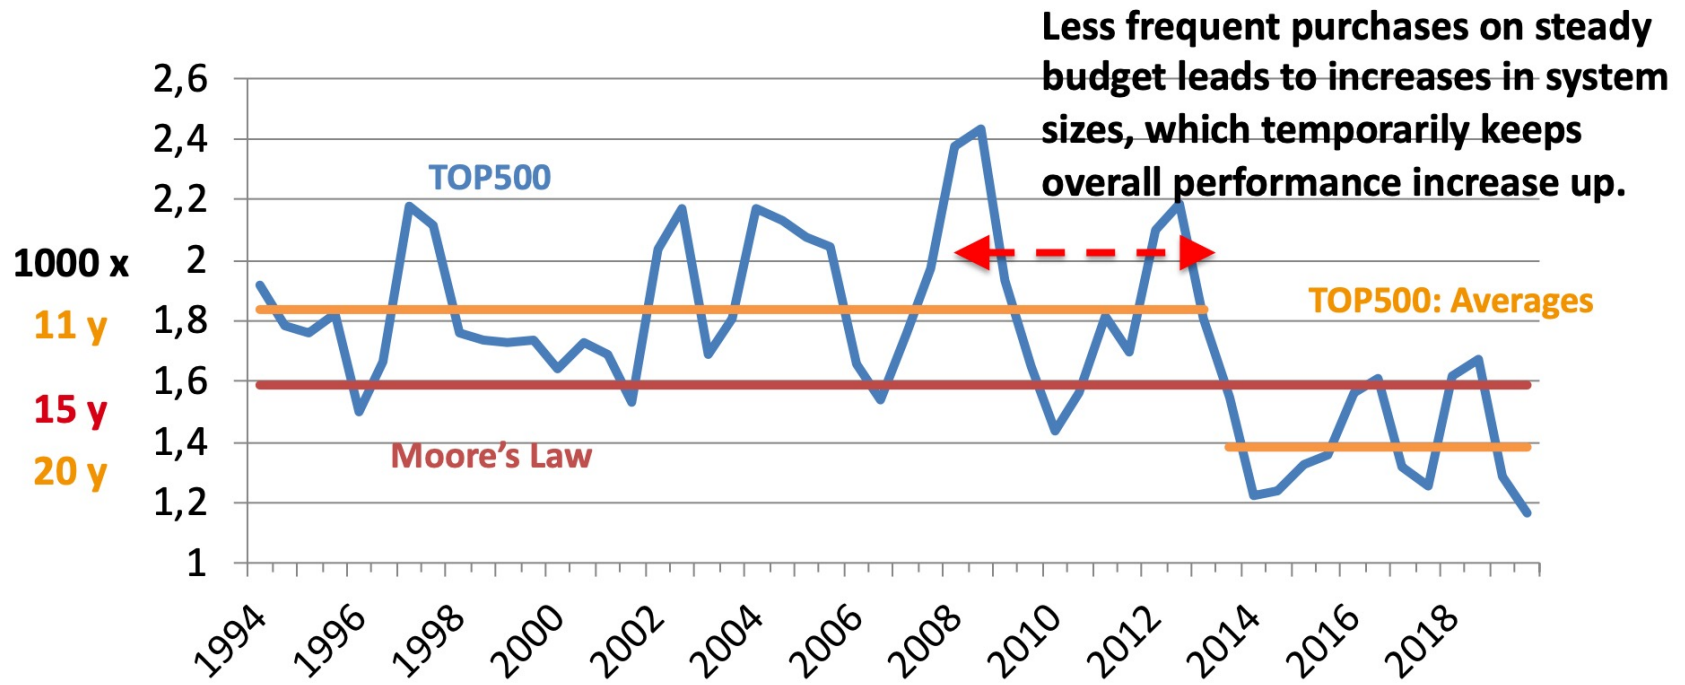
# TOP 500 Average Age (2019)

- Increase in age starts in 2009
  - When Dennard scaling slows
- Computers will have longer lives!




# Top500 and Moore's Law (2019)

- Exceeded Moore's law for decades. Why?
  - Now, closely following Dennard scaling limits on Moore's Law
- TOP 500 has concentrated performance in few systems





# The Rise of China

- Fraction of machines by count



# China (2)

- Compute Fraction
- Producer



# Discussion

- Diversity narrowed to US only by 2009
  - These are the system vendors/integrators
  - Representing a global (mostly Asia) supply chain
- China committed to vertical intergration
  - From chip to HPC, the whole system
  - First all Chinese supercomputer 2016

But it's more complex:

In 2015, US Government rejected an export license for Intel to send CPUs and coprocessors boards.



# Green 500


- GFlops/watt: dominated by accelerators

Rank	TOP500 Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Power Efficiency (GFlops/watts)
1	301	<b>MN-3</b> - MN-Core Server, Xeon Platinum 8260M 24C 2.4GHz, Preferred Networks MN-Core, MN-Core DirectConnect, Preferred Networks Preferred Networks Japan	1,664	2,181.2	55	39.379
2	291	<b>SSC-21 Scalable Module</b> - Apollo 6500 Gen10 plus, AMD EPYC 7543 32C 2.8GHz, NVIDIA A100 80GB, Infiniband HDR200, HPE Samsung Electronics South Korea	16,704	2,274.1	103	33.983
3	295	<b>Tethys</b> - NVIDIA DGX A100 Liquid Cooled Prototype, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100 80GB, Infiniband HDR, Nvidia NVIDIA Corporation United States	19,840	2,255.0	72	31.538



# Green Trends (as of 2019)

- Nvidia has caught back up
- MN-Core: new architecture
  - Prototype? For next Top 500 machine





# Final Comments

- Convergence between Green 500 and Top 500?
  - Low power leads to more compute density
  - Practical limits on how much power to a single computer
    - 29MW at Fugaku is amazing
- But, that's not what we're seeing
  - New architectures built/prototyped at smaller scales
  - New processors/accelerators coming out at great rates

