


Lecture 14

Top 500

EN 600.320/420/620

Instructor: Randal Burns


21 October 2020



Department of Computer Science, *Johns Hopkins University*

https://www.top500.org/media/filer_public/54/77/5477d858-1f1e-410b-994b-b7122cfd1d57/top500_2020_06_v2_web.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
 - Does not measure I/O to storage
 - Does measure networking between nodes (needed to succeed)



TOP500: Interpreting their Slides

https://www.top500.org/media/filer_public/54/77/5477d858-1f1e-410b-994b-b7122cfd1d57/top500_2020_06_v2_web.pdf

- Slide 10
 - An Exaflop! has been reached (at reduced precision)
 - New architectures have emerged
 - Both in Green500 and Top500
- Slide 11
 - 4 new entrants: 2 in Europe, 1 from Industry
- Slide 12
 - HPCG is an alternative benchmark than HPL (LinPack)
 - Complex gradients involve sparse patterns of a suite of modern applications



Green 500 (slide 13-14)

https://www.top500.org/media/filer_public/54/77/5477d858-1f1e-410b-994b-b7122cfd1d57/top500_2020_06_v2_web.pdf

- Two new architectures
- nVidia A100: low-power data-center tensor GPU
 - Supports multiple word sizes
 - Up to 7 GPUs collaborate on single task
 - Selene system – most powerful industrial system ever
 - AI designed – train big neural nets
- MNCORE – deep learning training accelerator
 - Optimized for matrix operators
 - Supposed to be a ‘neural processor’ don’t know what that means
- These two new architectures follow on two new ones last time
 - ZettaScaler: hierarchical with stacked ram
 - A64FX: 512 wide vector
- There is a revolution in accelerator architecture



Replacement (slide 15-17)

https://www.top500.org/media/filer_public/54/77/5477d858-1f1e-410b-994b-b7122cfd1d57/top500_2020_06_v2_web.pdf

- Lowest replacement rate ever
 - Sum lower than trend
 - 500th lower than trend
- Technology improvement pacing
 - #1 machine on Moore's law
 - With discretization effects



Countries (slide 19-23)

https://www.top500.org/media/filer_public/54/77/5477d858-1f1e-410b-994b-b7122cfd1d57/top500_2020_06_v2_web.pdf

Rise of Chinese manufacturing

- 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.
 - Triggered a commitment to self-sufficiency
 - Remain export controls in effect today



Research v. Industry (slides 23,27,29)

https://www.top500.org/media/filer_public/54/77/5477d858-1f1e-410b-994b-b7122cfd1d57/top500_2020_06_v2_web.pdf


- Diverse set of vendors
 - Prefer performance share to system share
 - In both categories
 - Different vendors in each sector
- Chip share is more confusing
 - Multiple offerings per vendor
 - Intel CPUs and Xeon Phi
 - NVidia GPU accelerators and A200 tensor
 - ARM is Fujitsu
- Split between
 - Research: HPC for science
 - Industry: training deep networks for AI
 - e.g. driverless cars were a stated goal of NVidia.



Couple Comments

From previous years presentations

- Evolution of types of systems shows Moore's law consequences
 - Death of single CPU
 - Multicore ->
 - SMP ->
 - NUMA ->
 - Distributed memory ->
 - Accelerators
 - Rise of commodity hardware




Couple Comments

From previous years presentations

- Accelerators

- Solve problems with Dennard scaling
- In response to same factors as CPU multicore (2004)



The Rise of Accelerators

- Accelerators emerged as a way to defer Dennard scaling limitations
 - 1/3 of systems, but > 50% of flops
- Almost all Nvidia GPUs

