

ECE 571 – Advanced Microprocessor-Based Design Lecture 30

Vince Weaver

<https://web.eece.maine.edu/~vweaver>

vincent.weaver@maine.edu

22 November 2024

Announcements

- HW11 will be a GPU reading (after Thanksgiving)
- Project status update due on the 25th
Note! This requires related work (see handout)
Three references if working alone, six if in group



Some SC24 notes

- Show pictures
- Competitive, 21% acceptance rate
- SCInet infrastructure, 8.14TB/s
- Student cluster competition
- Keynote. 17,500 people. NASA talk.
- Notes on Quantum talks. There is hardware? Not sure how to use it? Different hardware (lots of N2 tanks?) TUM (Reinaldo). Shulzs Quantum Valley. Good at brute force problems, traveling salesman, prime



factoring. How to integrate with SC setup. Do you need Quantum Physicists as Sysadmins?

- How do you stop people running bitcoin miners on your hardware. Talk to trying to watch which programs are run and stop if signature doesn't match typical HPC programs. Many issues with this including eventually people can modify code to look more similar
- IPv6, use there. Use in clusters, things like Lustre, RDMA
- RISC-V accelerators. Coming in a few years?
Dave Ditzel (Transmeta) was there with fancy RISC-V



accelerator plans but not shipping until 2027? People concerned about FORTRAN support

- ARM Firmware... everything chiplets these days
- Showroom, lots of CLI (?) RAM, lots of cooling
- Green500, didn't realize it was so hard to get measurements, most still aren't full system, many just 20% and extrapolating

Level1/Level2/Level3 reporting rules. Most of top computers in Europe. 72 GFLOP/W to hit top



Notes on the Top500 BoF 2024

- Not released in time to have us watch
- Didn't see in person, was at poster session
- El Capitan at LLNL new #1
 - 1.742 EFLOPs
 - 11,039,616 cores (CPU+GPU)
 - AMD 4th gen EPYC (24 cores)
 - AMD Instinct MI300A
 - Cray Slingshot 11 network
 - #18 on GREEN500 (58.89 GFLOPs/W)



Need in the 70s to hit top (I was at that presentation)



Notes on the Top500 BoF Video – 2023

- What is a BOF anyway?
- Why is Exaflop a big deal?



Top500 BoF Video – Aurora

- Aside, long history, supposed to be 200 PFLOPS in 2018 with lots of Xeon Phi
- Intel Sapphire Rapids, Intel having problems with chips
- Intel Xeon MAX (Ponte Vecchio) GPU, had to turn to TSMC to get parts made
Roughly 2x A100 perf, not as fast as H100
- Only part of it running, goal is 2 Exaflops (mostly from GPUs)
- 160 Racks, 10,624 Nodes, 21,248 CPUs, 63,744 GPUs



- HPE Slingshot-11 interconnect (formerly Cray)

<https://www.nextplatform.com/2022/01/31/crays-slingshot-interconnect-is-at-the-heart-of-hpes-hpc>

a lot of history on interconnects

- Dragonfly Topology?
- 10.9 PB of DDR5 RAM (512 GB/CPU?)
- 1.36 PB of CPU HBM
- 8.16 PB GPU HBM (100GB/GPU?)
- Storage 230 PB



Top500 BoF Video – MS Eagle

- Azure HPC
- 14,400 H100 GPUs
- Infiniband Quantum 7
- 1800 servers
- Any customer can use
- Ubuntu
- Generative AI



Top500 BoF Video – General

- HPCG results different as usual. Frontier only 16 PFLOPS
- HPL-MXP – Piotr working on. Mixed precision. Estimates can get 10x performance
Complicated coding. Worth it? Old days would just wait, Moore's Law.
Now maybe it is
- Green 500. Henri, 65 GFlops/W
- Less churn, computers on list for longer time and used



longer before getting rid of

- Unlikely to hit 10 Exaflops by end of decade



Notes on the Top500 BoF Video – 2022

- Frontier
 - 74 racks, 9.2PB RAM (half HB, half DDR4)
 - 90 miles of network cable
 - \$600 million
 - Quiet, water cooled. Warm water (32C)
 - Trouble getting more than 600PFLOPs, turned out to be linear-time thing in Cray MPI library
 - 3 hours to run Linpack, nodes keep failing when try to do run



- Non-linpack results, HPGC Frontier #2 only 14 PFLOPs
- Green 500, GFLOPs/W. Frontier much lower. Top was machine with first NVIDIA H100 (Hopper)
- Systems appearing more slowly on list, aging more before dropping off



ARM 2024 Chips

- *ARM Unveils 2024 CPU Core Designs, Cortex X925, A725, and A520: Arm v9.2 Redefined for 3nm* by Gavin Bonshor
- What is 9.2? How different from v8?
 - Add various instructions
 - v9.3 hinted conditional branches, NMI, memcpy instructions?
- Move to 3nm
- What is CSS (compute subsystem?)



- Enhanced security, memory tagging (MTE), confidential compute arch (CCA)
 - CCA – ability to hide data from root/superuser, especially useful in virtual machine hosting environments
- CSS - scale across different form factors
- Big/middle/little cores
- Speeds greater than 3.6GHz
- Cortex X925 – wider pipeline, improved branch prediction, enhanced OoO. Boost IPC
- 3nm – more transistors in area
- “Holistic view of RTL”



- Kleidi AI
- Kleidi CV (computer vision)
- Cortex-X925
 - 10-wide decode
 - Double instruction window
 - Reduce stalls
 - 2x I1 icache bandwidth, 2x I1 TLB size
 - better branch predictor
 - Good power
 - Compare IPC against PCs?
- Cortex-A725



- (builds on A720)
- 1MB L2 cache
- DVFS, half-slice power modes (?)
- Cortex-A520
 - Previously existing, shrunk to 3nm
 - 15% energy saving



Quick ARM Versions Reference

- ARMv7: Cortex A5, A7, A8, A9, A15, A17
- ARMv8: A34, A35, A53, A57, A72, A73
- ARMv8.1: A55, A75, A76, A77, A78, X1, N1
- ARMv9.0: 510, 710, 715, X2, X3, A520, A720, X925, AWS graviton, NVIDIA Grace
- ARMv9.2: Apple M4



Apple Chips

- Essentially Apple Press Release?
- Second-gen 3nm
- 2x-faster ray tracing
- Thunderbolt 5
- 2x faster neural engine (16 core)




Apple M4

- 10-core CPU, 4 perf and 6 efficiency (taking Intel terminology)
 - Perf: improved branch prediction, 10-wide ins decode, 40% larger reorder
 - eff: improved branch predictor, lower latency fp
- 1.8x faster than M1 (note comparison)
- 10-core GPU, 2x faster than M1



Apple M4 Pro

- 14-core CPU, 10perf+4eff, 1.9x faster M1, 2.1x faster graphics
- 64GB fast unified memory, 273GB/s memory (for comparison DDR5 is 64BG/s)



Apple M4 Max

- 16-core CPU, 12perf+4eff, 2.2x faster M1, 2.5x faster graphics
- 128GB fast unified memory, 546GB/s memory

