

ECE 574 – Cluster Computing

Lecture 3

Vince Weaver

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

vincent.weaver@maine.edu

11am, Barrows 133

4 September 2026

Announcements

- HW#1 was due!
- A break on homeworks until next week.
- Don't forget no class Monday due to Labor Day

Top500 List – June 2026

#	Name	Country	Arch /Proc	Cores	PFLOPS Max/P	Accel	Power
1	LineShine	China	LX2 (ARM9)	13,789,440	2.2k/2.7k	?	42MW
2	El Capitan	US/LLNL	AMD EPYC4	11,340,00	1.8k/2.8k	AMD Instinct	30MW
3	Frontier	US/ORNL	AMD EPYC	9,066,176	1.4k/2.1k	AMD Instinct	25MW
4	Aurora	US/Argonne	Intel SPR	9,264,128	1.0/2.0k	Intel MAX	39MW
5	JUPITER	Germany	GH Superchip (ARM)	4,801,344	1.0/1.2k	NVIDIA GH200	16MW
6	HPC7	Italy	AMD EPYC4	3,461,472	572/861	AMD Instinct	9MW
7	Eagle	US/Microsoft	Intel SPR	2,073,600	561/846	NVD H100	?MW
8	HPC6	Italy	AMD EPYC3	3,143,520	478/607	AMD Instinct	8MW
9	SC Fugaku	Japan/Riken	ARM64	7,630,848	442/537	N/A	30MW
10	Alps	Switzerland	NVD Grace (ARM)	2,121,600	435/575	NV GH200	7MW
11	LUMI	Finland	AMD EPYC3	2,752,704	380/531	AMD Instinct	7MW
12	Leonardo	Italy	Intel ICL	1,824,768	239/304	NVD A100	7MW
13	Isambard-AI	UK	NVIDIA Grace (ARM)	1,028,160	216/3278	NVD GH200	?
14	Tuolumne	US/LLNL	AMD EPYC4	1,161,216	208/289	AMD Instinct	3MW
15	ISEG2	Netherlands	Intel SPR	718,848	202/338	NVD H200	5MW
16	MareNostrum5	Spain/BSC	Intel SPR	663,040	175/249	NVD H100	4MW
17	Discovery 6	US/ExxonMobil	NVD Grace	806,208	164/218	NVD GH200	?
18	Maximus-02	US	Intel Xeon	747,520	156/343	NVD H200	?MW
19	ABCI 3.0	Japan	Intel Xeon	479,232	145/181	NVD H200	4MW
20	NIPA	Japan	AMD Epyc	643,072	137/156	NVD B200	?MW
21	CHIE-4	Japan	Intel Xeon	662,256	135/152	NVD B200	?MW
22	Shaheen III	Saudi Arabia	NVD Grace (ARM)	574,464	123/155	NVD GH200	2MW
23	EOS SuperPodDGX	US/NVIDIA	Intel SPR	485,888	121/189	NVD H100	?MW
24	MAXIMUS-384	US	Intel Xeon	976,896	115/251	AMD Instinct	?MW
25	SSC-24	US	Intel Xeon	349,440	106/151	NVIDIA H100	?MW

Notes on the Top500 BoF Video

- What is a BOF anyway?
- Why is Exaflop a big deal?
 - They're always pushing for the next order of magnitude
 - There was an “Exaflop Challenge” in 2011 hoping for an Exaflop in 20MW by 2018. They more or less got there eventually but not by 2018.
 - There were a lot of challenges, it was GPUs that saved the day

Extra Notes on #1

- LineShine ()
- National Supercomputing Center in Shenzhen
- LingKun arch using LX2 CPUs (ARMv9 based)
 - Each LX2 CPU 304 cores running 1.5GHz
 - ARMv9 with SVE and SME vector
- NO GPU
- Memory: 1.4 PB HBM, 11 PB DDR5
- 650 PB storage
- LingQi high speed network 1.6 TB/s
- 42MW power, 52 FP64 GFLOP/W

Notes on the Top500 BoF Video 2025

- When at SC I try to attend these, was not there in 2025 (and even though I was there at 2024 there was a conflict with the time)
- Jack Dongarra, Turing award winner, previous boss of mine
- You can't tell but this presentation is in a massive auditorium

Notes on the Top500 BoF Video 2025

- Not many changes since previous year, so not very exciting video
- A lot of changes happened by June 2026 making this video obsolete
- I couldn't find any good coverage of the June 2026 announcement at ISC Hamburg

Benchmarks

- Linpack – regular matrix multiply
- HPCG – High performance conjugate gradient
- HPL-MxP – Multi-precision
 - Traditionally HPL was 64-bit fp
 - GPUs better at 32-bit
 - Modern GPUs optimized for AI like 16-bit and 8-bit better
 - It's possible to use the lower precision but still get 64-bit results

Other things

- Green 500 – Europeans dominate
- Student Cluster Competition – 10kW
Their winners would have been on the list somewhat recently
- DOE and Europe AI initiatives
- Big News: Top500 going to be run by ACM/SIGHPC

List Update

- 94 new systems (just not really in top 10)
- JUPITER at Jülich. Managed 1.0 Exaflops exactly
- Not as many Blackwell GPUs yet
 - Ampere, Hopper / Lovelace / Blackwell
- Why not as many AI/datacenter results?
 - Only Microsoft
 - Expensive and time consuming to run, better use to sell compute to customers
 - They might not want to reveal what equipment they are using

DGEMM with Tensor Cores (2024)

- Using AI hardware to do HPC (subsidized? like gaming was before?)
- Recovering 64bit precision from FP16
- If using FP32 no benefit going Hopper to Blackwell

Some Supercomputing (SC24) notes

- November 2024, in Atlanta Georgia (show pictures)
- Large academic / industry conference, 17,500 people
- Competitive conference, only 21% paper acceptance rate
- Even workshops and poster sessions reject submissions
 - Attended Security workshop. Only recently have people started about caring about security in computer systems.
 - ProTools workshop, Willow and I had a paper there on Hybrid CPU performance measurement with PAPI

SC24 – continued

- Large vendor showfloor with lots of companies, hardware and software vendors, universities and research centers
 - SCInet infrastructure, 8.14TB/s
 - Student cluster competition
 - Lots of CLX RAM, Quantum, cooling/rack solutions

(SC24) talks

- Keynote. NASA talk.
- Attended various talks
 - Quantum Computing BoF
There is hardware? Not sure how to use it?
Different hardware (lots of N2 tanks?) TUM (Reinaldo). Schulzs Quantum Valley. Good at brute force problems, traveling salesman, prime factoring. How to integrate with SC setup. Do you need Quantum Physicists as Sysadmins?
 - Talk on detecting inappropriate use of HPC resources
How do you stop people running bitcoin miners on your hardware. Talk to trying to watch which

(SC24) more talks

- IPv6 in HPC. 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
- 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 (note: m1 laptop 6 GFLOP/w, Pi5 3.6 GFLOP/W, haswell-ep 2.1 GFLOP/W)

Top500 List Notes

- Rmax vs Rpeak – Rmax is max measured, Rpeak is theoretical best
- HPL Linpack
 - Embarrassingly parallel linear algebra
 - Solves a (random) dense linear system in double precision (64 bits) arithmetic
- HP Conjugate gradient benchmark
 - More realistic? Does more memory access, more I/O bound.
 - #1 on HPCG usually not same as #1 HPL
- Green 500

My Lab's top Computer List

- <https://web.eece.maine.edu/~vweaver/group/machines.html>
- Haswell-EP (we'll use for homeworks)
436 GFLOPs, 2.3GFLOP/W, would have been #1 in 1996
- Raspberry Pi 4 – 13 GFLOPs, 2.0 GFLOP/W, #10 in 1993

Top supercomputer Architecture

- Commodity or custom
- Architecture: x86? SPARC? Power? ARM embedded vs high-speed?
- Memory (how much? 1GB per core?)
- Accelerator cards / Heterogeneous Systems

Top supercomputer Storage

- How much?
- Large hadron collider one petabyte of data every day
- Shared? If each node wants same data, do you need to replicate it, have a network filesystem, copy it around with jobs, etc?
- Cluster filesystems?

Top supercomputer – Power/Cooling

- Reliability. How long can it stay up without crashing?
Can you checkpoint/restart jobs?
Sequoia MTBF 1 day.
Blue Waters 2 nodes failure per day.
Titan MTBF less than 1 day
- Cost of Power over lifetime can be more than that of hardware cost
- Power comparison: small town? 1MW around 200 - 1000 homes? (this varies)
- Does Power include cooling
- Big river nearby?

Top supercomputer – Network

- How fast? Latency?
- Interconnect? (torus, cube, hypercube, etc)
- Ethernet? InfiniBand? Custom?

Top supercomputer – Software

- Operating System
 - Linux? Custom?
 - If just doing FP, do you need overhead of an OS?
- Job submission software, Authentication
- Tools – software that can help you find performance problems

Top supercomputer – Applications

- Software – how to program?
- Too hard to program can doom you.
- A lot of interest in the Cell processor. Great performance if programmed well, but hard to do.

Top supercomputer – Running Linpack

- How long does it take to run LINPACK?
- How much money does it cost to run LINPACK?
- Is it worth it?