ECE 574 – Cluster Computing
Lecture 24

Vince Weaver

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

vincent.weaver@maine.edu

20, 25, 27 April 2023

Announcements

- Remember no final exam for this course
- Still catching up on grading homeworks and 2nd exam
- Don't forget to do online student evaluations

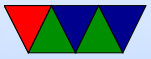

HW #7 notes

- Fine grained parallelism
- Running on the Pi-cluster
  - Test with np=7, some code failed that worked with 2
    9d4b6548fa8c6ff66602ef5993aca90f common
    seems to be not gathering in the extra lines
  - Reading from each core rather than Bcast doesn't help anything.
  - Some analysis of pi-cluster results
    Only scale up to 4.

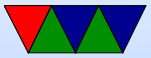

cores	load	bcast	convolve	combine	gather	tail	store	total
1	1.0	0	12.8	3.8	0.1	0	3.4	21.2
2	1.0	0.1	6.4	1.9	1.8	0	2.4	13.7
4	1.0	0.3	3.2	0.9	3.0	0	2.4	10.9
8	1.0	5.6	1.7	0.5	4.6	0	2.4	15.8
16	1.0	7.3	0.7	0.2	6.5	0	2.4	18.2
32	1.0	8.0	0.3	0.1	6.4	0	2.4	18.3
64	1.0	8.8	0.1	0.06	6.9	0	2.4	19.5

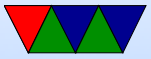

HW #8 notes

- Be careful memory copying, if copying an array of 9 ints need to copy 36 bytes (not 9)
- Also, you can pass in ints as parameters (no need to allocate space and then memcpy in. Or you could, but if you do you would use points and allocate space properly)
- Be sure you are using *unsigned char* for the image data, not signed char.
- Limits and matrix indexing
- My results with a large image

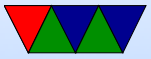

Type	Load	Copy	Convolve	Combine	Store	Total
OMPx16	1s		0.882	0.135	0.9	3.12
MPIx16	1s	0.5+1.4	0.6	0.1	1.0	4.9
Cuda	1s	0.3	0.2	0.2	1.0	3.3
OpenCL CUDA	1s	0.2	0.4	0.4	0.9	2.9
OpenCL intel	1s	0.3	0.2	0.2	0.9	3.0
OpenCL pocl	1s	0.4	0.6	0.7	1.0	3.4

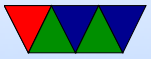

HW #9 notes

- Didn't grade yet

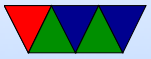

Myths and Legends in High-Perofmance Computing

- Paper by Matsuoka, Domke, Wahib, Drozd, and Hoefler
- Myth1: Quantum Computing
- Myth2: Deep Learning
- Myth3: Extreme Specialization
- Myth4: Accelerators
- Myth5: FPGAs
- Myth6: Zettascale
- Myth7: Memory/Core Ratio

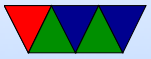

- Myth8: Diassgregation
  - Silicon Photonics

    Circuit switching vs Packet Switching
  - Compute Express Link (CXL)
  - Speed of Light Issues
- Myth9: Applications are Improving
  - Post Moore:

    Improve architecture?

    Non-CMOS?

    Abandon traditional computers?
  - Is there an Algorithmic Moore's Law?

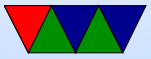

- Myth10: End of Fortran
  - What is the proper lay of abstraction for optimization? low-level virtual machine? C/C++/assembly? dataflow representation?
- Myth11: Low or Mixed Precision
  - Can break application
  - Abandon IEEE-754?
  - AI had interesting issue, fp16 faster than fp32, but did not always converge. bfloat16 had wider range, but still issue. So then 19-bit (tensorfloat-32) still faster

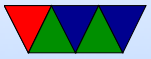

than fp32 but better behaved

- Myth12: The Cloud
  - Cloudification of Supercomputers

    many supercomputers have cloud-like features
  - Supercomputerification of Clouds

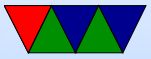

More Exascale

- Exascale Day – 10/18
- Frontier (OLCF-5)
  ORNL
  1.5 Exaflops, $600 milion
  AMD Epyc and Radeon Instinct GPUs, 30MW, 100 racks

- Aurora
  DOE (Argonne)
  Intel and Cray (now HPE)

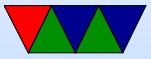

originally supposed to be 2018 with Xeon Phi
1 exaFLOP
$500 million
general scientific community. low carbon tech, subatomic particles, cancer, cosmology, solar cells
over 9000 nodes, each two Intel Sapphire Rapids CPUs, Golden Cove, 10nm, DDR5 RAM
6 Xe ponte vecchio GPUs, chiplets

- El Capitan
  LLNL (2023), NNSA
  $600 million

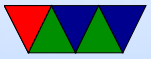

Replace Sierra (IBM Power9 + Nvidia)

2 Exaflops

Zen 4

less than 40MW

Infinity Fabric

Connecting nodes is AMD slingshot fabric, 200Gb/s, one port per CPU

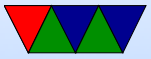

Zettascale

- Eurolab- 4 -HPC Long-Term Vision on High-Performance Computing Editors: Theo Ungerer, Paul Carpenter
- zettascale by 2030?
- convergence with big data?
- deep neural networks?
- die-stacking
- non-volatile memory
- resistive computing?
- neuromorphic computing? – try to replicate nerves in

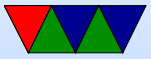

silicon

- quantum computing?
- nanotubes?
- graphene/diamond based transisitors?
- optical networks on die / Terahertz communication
- HP Labs "the Machine"

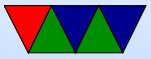

Zettascale

- Challenges
- Lines of code. 10-100 Euro per line?
- Approximate Computing
- auto-tuning
- debugging and profiling
- extreme data
- cloud, big data
  modern data centers 20MW cover 17 football fields
- exabytes of data, merge with cloud

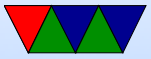

Disruptive Tech

- Moores Law continues? 1.5nm by 2030?
- DRAM to 7.7nm in 2028, 32GB/chip? Scaling DRAM below 20nm hard
- Might be stuck at 32GB unless something new happens
- DUV argon Flouride excimer lasers, 193nm
  (deep ultraviolet) excited dimer, noble gas plus reactive gas
- die stacking, chiplets
- non-volatile memory

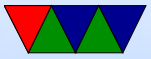

- spintronics
- memristors
- Photonics, 15ps/mm in silicon, 5ps/mm in waveguides
  stacked chips can have photonic layer
- mode-division multiplexing, free-air propogation, plasmonics
- photonic non-volatile mem, photonic computing
- memristive computing
- neuromorphic
- quantum computing
  d-wave, qubits, mili-Kelvin, new algos

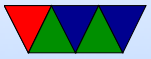

- nanowires
- graphene, how you make it, 100GHz transistor
- diamond transistors

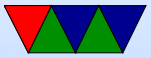

Last Notes

- Near memory / in memory computing
- power
- analog computing
- end of von neuman (memory hierarchies)
- Green computing, liquid nitrogen temps (memory story)
- System software, programming languages