CHEP'12 Report

Don Holmgren

SCF Monthly Meeting

July 11, 2012

Outline

- GPUs
- High speed networks (online only)
- Multicore
- Storage
- Misc
- Recommended Talks

GPUs

- Primarily mentioned in online & event processing tracks
  - Exceptions: Giele *et al.* (event generation), LQCD (two posters), a few plenaries
  - ALICE: GPUs for tracklet finding in their TPC
    - 3X speedup compared to their 6-core Xeons
    - Xeon vs GPU results don't precisely match, due in part to non-associative arithmetic
    - Per-GPU add-on cost less than the cost of the CPUs (more discussion below)
    - Their code/workflow is structured to allow CPUs to do other work while the GPUs calculate tracklets, so a bigger win than just the tracklet speed-up

GPUs

- PANDA (contribid=353)
  - Studied track finding, comparing GPU to FPGA on a helix tracking algorithm
  - GPU is 200x faster than CPU
  - CPU is 30% faster than FPGA
- Edinburgh – optimizing TMVA ("Toolkit for Multivariate Analysis") (contribid=492)
  - Used GPUs in neural network training
  - Since event-based parallelism won't work (training is sequential), instead did neuron-based parallelism
  - For small networks, CPUs are faster, but for large/complex networks, GPUs win. Increasing the number of hidden layers doesn't increase GPU run time
  - Also trained multiple networks simultaneously on a GPU

GPUs

- Other than a passing mention of ATI GPUs in one (maybe two?) plenaries, all used NVIDIA
- Also, in the talks I attended and the posters I read, I only saw CUDA implementions
  - No OpenCL or OpenACC
- ALICE uses gaming cards, noting they are cheaper than Tesla cards (and their Xeon processors)
  - They don't care about double precision, ECC (!?), warranties, but only price and performance
  - I confirmed that they see hardware failures, but the cards are so cheap they didn't mind just doing replacements

High Speed Networks (DAQ only)

- Particularly in talks about upgrades, many mentions of 10/40 gigE and Infiniband
- Interesting study of TCP (IPoIB) vs uDAPL in the context of the "DAT collaborative" (www.datcollaborative.org)
  - Andrea Petrucci (contribid 282)
  - For CMS upgrade, in their XDAQ distributed framework they use DAT "pluggable Peer-Transports", e.g. SOAP/HTML, I_2O (TCP/Myrinet/FIFO)
  - Application code doesn't change moving from Ethernet to Myrinet to Infiniband
  - ptuDAPL is their new peer-transport for Infiniband and iWARP
  - Lots of data showing DAPL in general better than TCP (and IPoIB) for event builder traffic
- Myrinet to be replaced as the tunnel-to-surface DAQ network in the CMS upgrade (by Ethernet or Infiniband)

Multicore (EP only)

- Several high level talks (including two plenaries) covered in depth the need to go parallel within individual events rather than across events
- Nice talk by Christopher Jones (contribid=194) on using Apple's *libdispatch* for a threaded framework
- Noticed a few talks that investigated Intel Threading Building Blocks
  - Some static code analysis to look for parallelization strategies

Storage

- I didn't hear much about hardware technologies, just standard bits in plenaries about growth of disks vs smaller growth in disk speeds.
- The storage plenary (Peters, contribid=603) had a nice discussion of HEP vs other big data (e.g. Google, Yahoo)
  - Data volumes are comparable (HEP smaller)
  - Object counts are 3+ orders of magnitude smaller in HEP
  - Can full POSIX filesystems survive at huge metadata counts?
    - Scaling up may require moving from POSIX IO to cloud-like (e.g. S3 API) APIs
  - RAID vs RAIN (Redundant Array of Inexpensive Nodes)
- Heard several folks talking about moving away from CASTOR HSM (to what – my notes aren't good here, but I can dig more if others are interested)

Miscellaneous

- Expert systems
  - LHC experiments have done a lot with expert systems to assist shifters (one was CLIPS based, another a Perl-based framework)
- Data taking efficiencies
  - Very interesting, lightly-attended talks by CDF and D0 giving DAQ summaries
  - CDF's final efficiency was about 84%, vs high 90's for LHC experiments
- Configuration management – heard "Puppet" several times

Recommended Talks

- The summaries were (as usual) very good – kudos to Adam for having working links in his Event Processing Track Summary to the actual talks
- "Computing Technology Future" plenary (Lennart Johnsson) had lots of good info. Much was about power consumption:
  - ICT (information and communications technology) alone produces CO_2 emissions equivalent to the entire aviation industry. Direct emissions of internet + ICT = 2-3% of worlds emissions. ICT growth rate is the highest of any sector (4 to 6 years doubling time). A small server generates the CO_2 of an SUV getting 15 mpg.
  - The costs to power and cool a server over a 3 year lifetime now exceed the initial cost of the server
  - Energy efficiency (Flops/KWHr) is doubling more slowly (19 mos) than Top500 performance (#1: 13.6 mos, #500: 12.9 mos)
- "A Reflection on Software Engineering in HEP" plenary (Federico Carminati) is well worth looking through (thesis is that HEP developed agile/extreme programming as a software engineering method but never realized it)