Computational Working Group Cross Group Meeting

Chris Jones & Adam Lyon

17 April 2019

Areas Covered by Computational Working Group

- Applications
- Storage
- CPU

Challenge: Huge data rates & volume

- Unprecedented data taking in upcoming experiments
  - HL-LHC: ~1.5 EB/year data just from the detector
  - DUNE: ~ 30-60 PB/year data just from the detector
- Processed and simulated data grow proportionally
- Stretching limits of
  - data storage
  - data movement
  - data processing

Challenge: Data Complexity

- Computation needed to process a unit of data are increasing
- E.g. CMS in HL-LHC
  - 200 overlapping events
    - compared to ~50 in last running period
    - Processing time grows quadratically with number of overlapping events
  - More complex sub-detector components
    - Require time-consuming algorithms to get most from the new detectors

Challenge: Heterogenous Computing

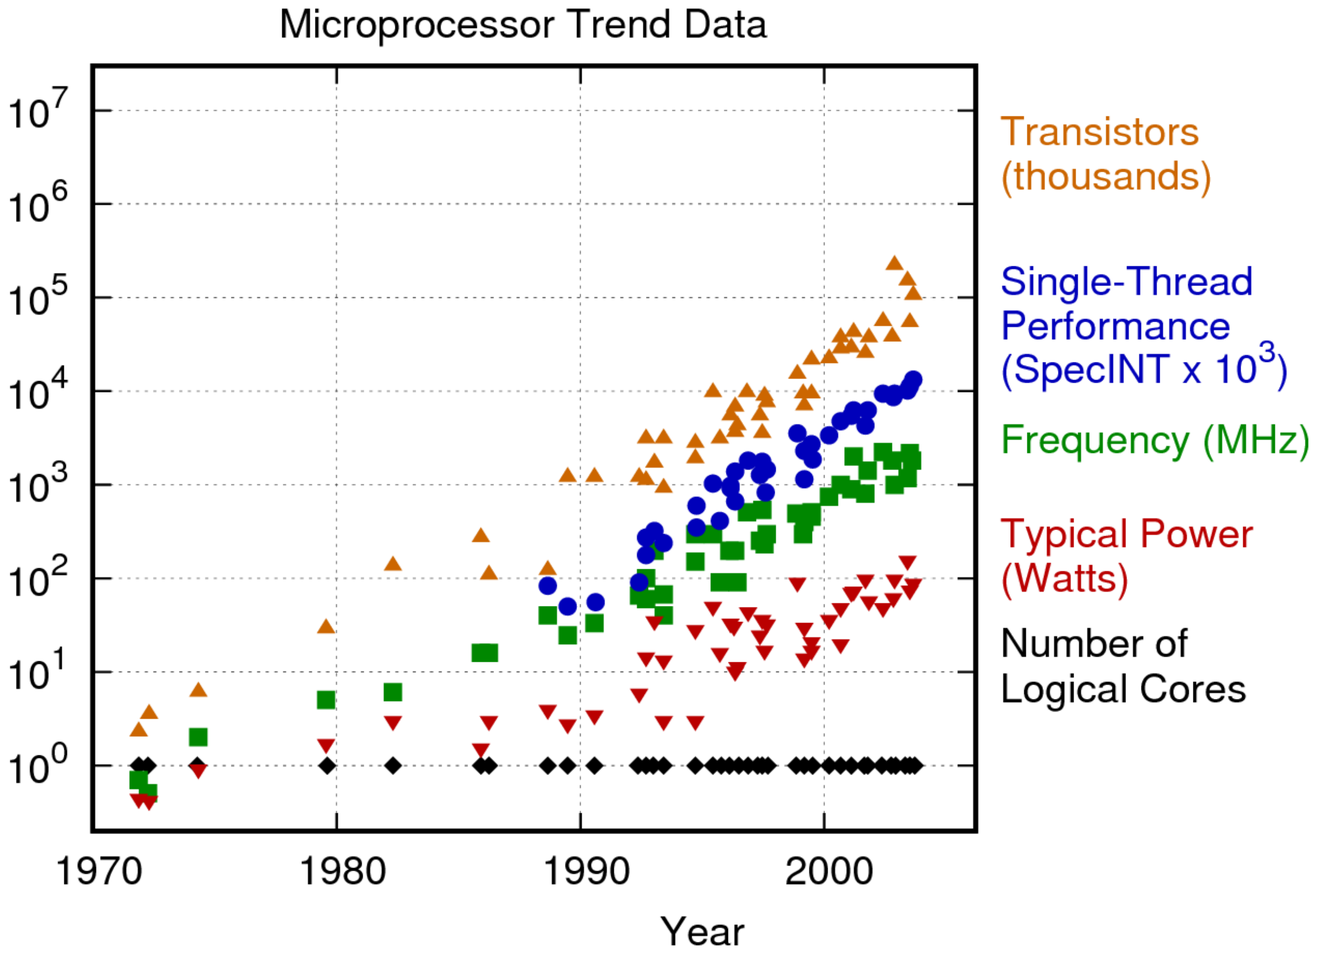

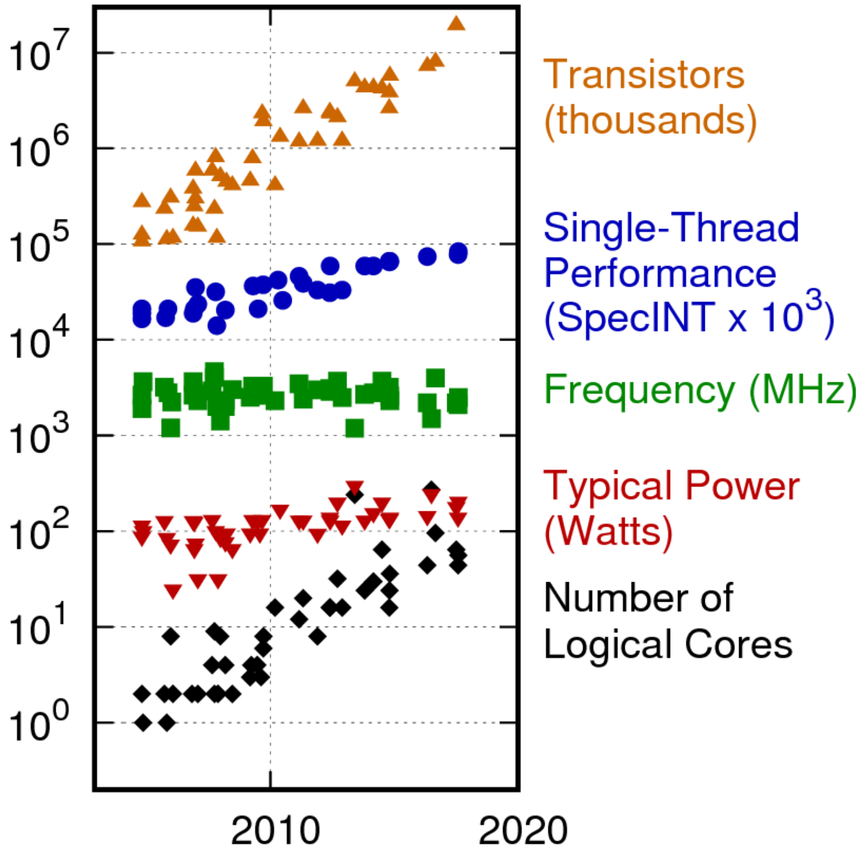

Fermilab

Challenge: Heterogenous Computing

- Future: multi-core, limited power/core, limited memory/core, memory bandwidth increasingly limiting
- The old days are not coming back
- The DOE is spending $2B on new “Exascale” machines…
  - OLCF: Summit — IBM CPUs & 27K NVIDIA Volta GPUs (#1 supercomputer in the world)
  - NERSC: Perlmutter — AMD CPUs & NVIDIA Tensor GPUs (2020)
  - ALCF: Aurora — Intel CPUs & Intel Xe GPUs (2021) — first Exascale machine
- Notice a pattern above? GPUs are winners. (Intel KNL is a big loser)

- These machines offer massive computing capacity … much more than what we’re used to
- How do we use these machines efficiently?
- Run inefficient code in containers and adiabatically shift to HPC style coding? Will the HPC centers allow that or will this be a disruptive change
- GPUs will be everywhere … can we use them?

You?

- What have you heard here that worries you?

- What is special computing-wise about your future experiment/project?