OAK RIDGE National Laboratory

EXASCALE COMPUTING PROJECT

Exascale Panel – Introductory Slides

Bálint Joó (joob AT ornl.gov) – *Oak Ridge Leadership Computing Facility, Oak Ridge, USA*

Fri. Aug 4, 2023 – Fermilab, Batavia, IL, U.S.A

ORNL is managed by UT-Battelle LLC for the US Department of Energy

U.S. DEPARTMENT OF ENERGY

About me

- Born in Debrecen, Hungary
- Moved to Aberdeen, Scotland at age 12
- University education at Edinburgh University
  - B. Sc. (Hons) Physics & Computer Science (1996)
  - EPCC Summer Scholarship Programme (1996)
  - Ph. D. Theoretical Physics (2000)
- Post Docs at
  - University of Kentucky (2000)
  - University of Edinburgh at Columbia Univ. (2000-2002) working on QCDOC
  - University of Edinburgh (back in Edinburgh) (2002-2005)
- Staff Scientist at Jefferson Lab (2005-2020)
  - Working on Lattice QCD software and algorithms
- Oak Ridge National Laboratory (2020-present)
  - Group Leader for Advanced Computing in Nuclear Particle and Astrophysics in the Science Engagement Section
  - Exascale Computing Project (HI - Application Integration at the Facilities)
  - Preparing Lattice QCD Code for Frontier
  - **and now: taking an interest in other areas: medical imaging & HemeLB**









images from Google – may be subject to copyright

Oak Ridge Leadership Computing Facility Mission

from the OLCF web pages

- "The Oak Ridge Leadership Computing Facility (OLCF) was established at Oak Ridge National Laboratory to accelerate scientific discovery and engineering progress by providing world-leading computational performance and advanced data infrastructure."
- "The Oak Ridge Leadership Computing Facility is charged with helping researchers solve some of the world's most challenging scientific problems with a combination of world-class high-performance computing (HPC) resources and world-class expertise in scientific computing."

Our Systems

- Frontier – Exascale System (1.194 EF/s in HPL)
  - HPE/Cray EX system, 9408 nodes
  - 1 HPC Optimized 3rd Gen EPYC CPU (Trento)/node
  - 4 AMD MI250X Instinct GPUs (8 GCDs)/node
  - Cray Slingshot 11 Interconnect
  - 21 MW of power, liquid cooled (~6000 gal/min)
  - #1 on June 2023 Top 500 List
  - Now in EARLY production for INCITE, ECP
- Summit (148.6 PF/s in HPL)
  - IBM AC922 system, 4600 Compute nodes
  - 2xPower 9 CPUs/node
  - 6 NVIDIA Volta V100/node
  - Infiniband Interconnect
  - 13 MW of power, liquid cooled
  - #5 on June 2023 Top 500 list (#1 in 2018)

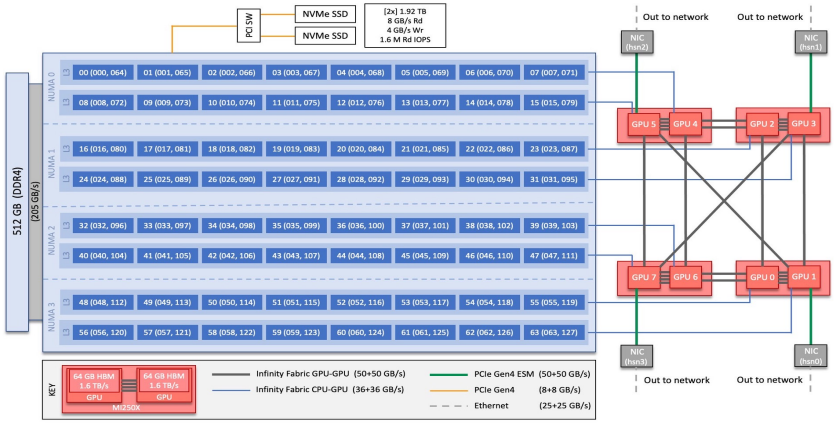

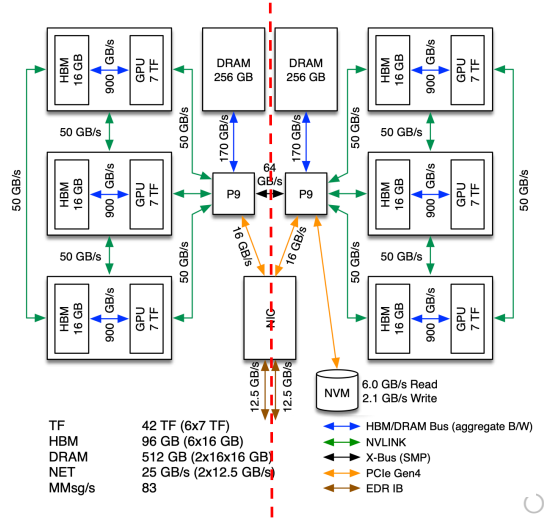

EXASCALE COMPUTING PROJECT

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

Come and Compute with Us!!!

- Did you know that ~60% of our cycles are available via the DOE INCITE program?
  - Anyone can apply, worldwide!
  - Successive applicants get liaisons like me to help them reach their science goals.
  - The deadline for INCITE proposals for 2023 has passed... but there is always next year
  - visit: https://www.doeleadershipcomputing.org/
- Up to ~25% of our cycles go to the ASCR Leadership Computing Challenge (ALCC):
  - visit: https://science.osti.gov/ascr/Facilities/Accessing-ASCR-Facilities/ALCC

(My own) ponderings about the future (but not my ideas)

- Moore's law is slowing
  - longer lifetimes for systems before they cease to be cost effective?
  - potential rise of custom accelerators for various tasks (implemented via chiplet tech?)
  - what if your application can't use them? -> will you be sent off to use the cloud?
- Deteriorating balance between FLOPS/network/memory -> more communication avoidance (Peter's and Evans' talks at this conference)
- High performance infrastructure for instruments & data processing?
  - less focus on a single "big machine" more on data & data flow? (IRI initiative)
- Improved ISO C++ support for accelerators
  - support by other accelerator vendors than just NVIDIA
  - ISO C++ evolution:
    - parallelism & concurrency (sender/receiver model)?
    - linear algebra (already being discussed)
  - Standard evolution is steady but can be slow moving... CUDA/HIP/SYCL/OpenMP as well as Kokkos/Raja etc will not go away anytime soon.
- Increased hardware complexity => increased likelihood of failures.
  - defensive coding, defensive running w. checkpointing

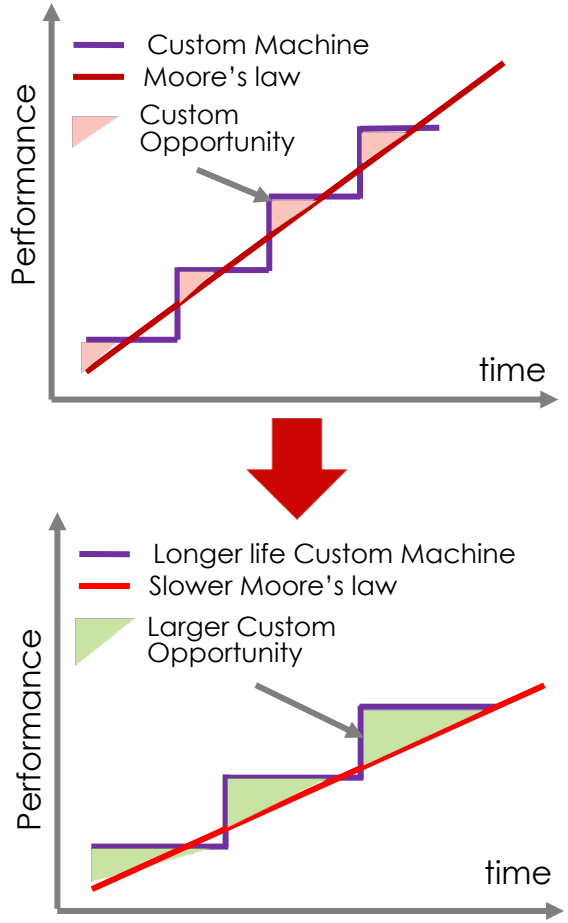

NB: I saw this discussed at the ECP AM in the panel: " What's coming next for HPC architectures?" J. Shalf, D. Reed, K. Yellick, N. Thompson. I think it came from D. Reed

Acknowledgement

- This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.
- This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.
- Balint Joo gratefully acknowledges funding under the US DOE SciDAC-5 program.

- And a gentle reminder in case it is needed:
  - Users should acknowledge the OLCF in all publications and presentations that speak to work performed on OLCF resources... (see above)
  - When discussing your research efforts that have been supported by ECP resources, we respectfully ask you to use the following acknowledgment text in your slide presentations ... (see above)