

HPC and AI at OLCF

Balint Joo (OLCF)

2nd Workshop on AI at the EIC

College of William and Mary, Williamsburg VA, Oct 12, 2022

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




U.S. DEPARTMENT OF
ENERGY

Who we are and what we do

- Oak Ridge Leadership Facility (OLCF) was founded
 - ***“to accelerate scientific discovery and engineering progress by providing world-leading computational performance and advanced data infrastructure”***
- OLCF is home to two world leading supercomputer systems
 - Summit: #1 on Top500 in June 2018, #4 in 2022
 - Frontier: #1 on Top500 this year and officially the first Exaflop system
- OLCF also hosts facilities for Data Archival (HPSS) and for Scientific Workflows (SLATE)



OLCF Summit



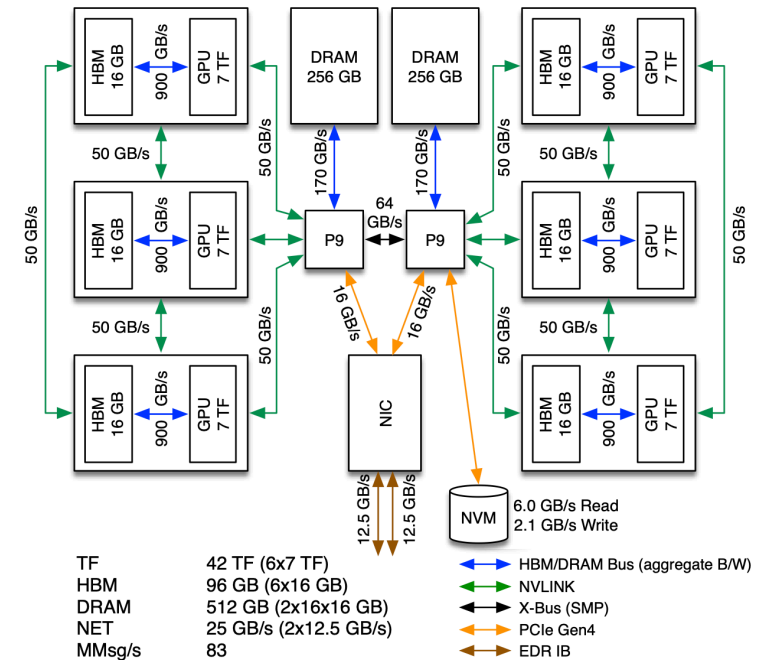
OLCF Frontier

OLCF Summit

- 4,608 nodes, each with:
 - 2 IBM Power9 CPUs
 - 6 NVIDIA Tesla V100 GPUs
 - 512 GB DDR4 memory, 96 GB HBM2 in GPUs
 - 1.6 TB of NV Memory (Burst-buffer)
- Dual-port Mellanox EDR InfiBand network
- 250 PB IBM File-system with 2.5TB/s transfer rate
- Peak Performance
 - 200 PFLOPs (FP64) – modeling & simulation
 - 3.3 ExaOps (FP16) -- analytics and AI
- 2018 ACM Gordon Bell Prize: ExaOps achieved in Climate related AI application [paper](#), [press release](#)



OLCF Summit

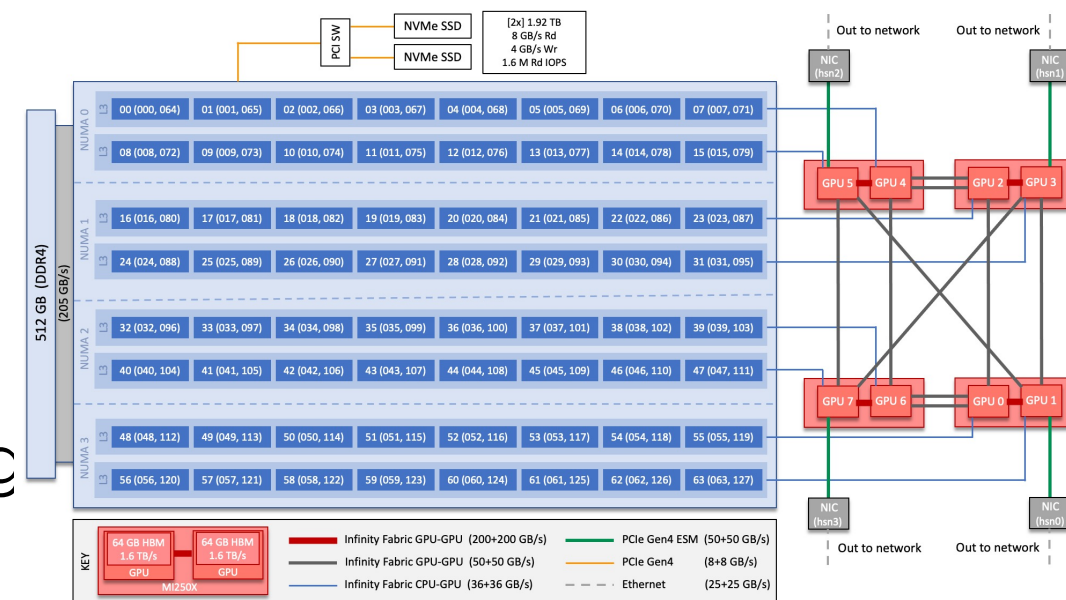



Summit Node Diagram

Open slide master to edit

OLCF Frontier

- 9408 Compute Nodes each with
 - 64 Core AMD EPYC CPU (HPC Optimized)
 - 4 AMD MI250X GPUs (2 GCDs each)
 - 64 GB HBM2E memory per GCD
 - 512 GB DDDR4 memory on the hosts
 - 2x 1.92 TB NVMe SSDs (burst buffer)
- HPE/Cray Slingshot network
- Currently the Alpine Filesystem (from Summit) is mounted.
- System is currently in Acceptance Testing



Frontier Node Diagram

Access to OLCF Systems

- ~60% of time on systems is devoted to the [INCITE](#) program
 - to solve problems which can only be solved by the worlds most powerful computer systems.
 - Typically application period is in the April-June timeframe
 - Anyone can submit an application (even from outside the US)
 - Proposals undergo 2 reviews:
 - Computational readiness; Can your code use the system, and run at scale?
 - Science Review: independent international review panel
 - Successful teams are allocated a liaison (like me) to help if they need
 - see: www.doeleadershipcomputing.org for more details
- ~30% of time on systems is devoted to ALCC projects
 - this program is allocated by ASCR to support DOE mission priorities and those of other US federal agencies
 - Typically runs July-July, application deadline tends to be around September
 - See [this link](#) for details
- ~10% of time is allocated at the Director's Discretion (DD projects)
 - e.g. porting codes and workflows, testing ideas, optimization, ready code for INCITE proposals
 - requests reviewed by Resource Utilization Council (RUC)
 - [see here for more information](#)

Science Engagement

- Mission: ***“In partnership with our users, we accelerate scientific discovery using our extensive domain and computing expertise”***
- A wealth of experience in utilizing GPU Accelerators in
 - Modeling and simulation: Lattice QCD, Many Body NP, Astrophysics, Climate, Engineering, Fusion
 - AI / ML Techniques: GRETA Detector, Astrophysics parameter searches, Drug design, Climate simulations
- Science Engagement staff can
 - help user partners port codes, optimize workflows, assist in code restructurings, advocate for users within OLCF and more

Highlight: Surrogate modeling of subgrid-scale effects in atmospheric flows: A deep learned approach using high-resolution simulation data*

Murali Gopalakrishnan Meena¹, Matt Norman¹,
David Hall², and Mike Pritchard^{2,3}

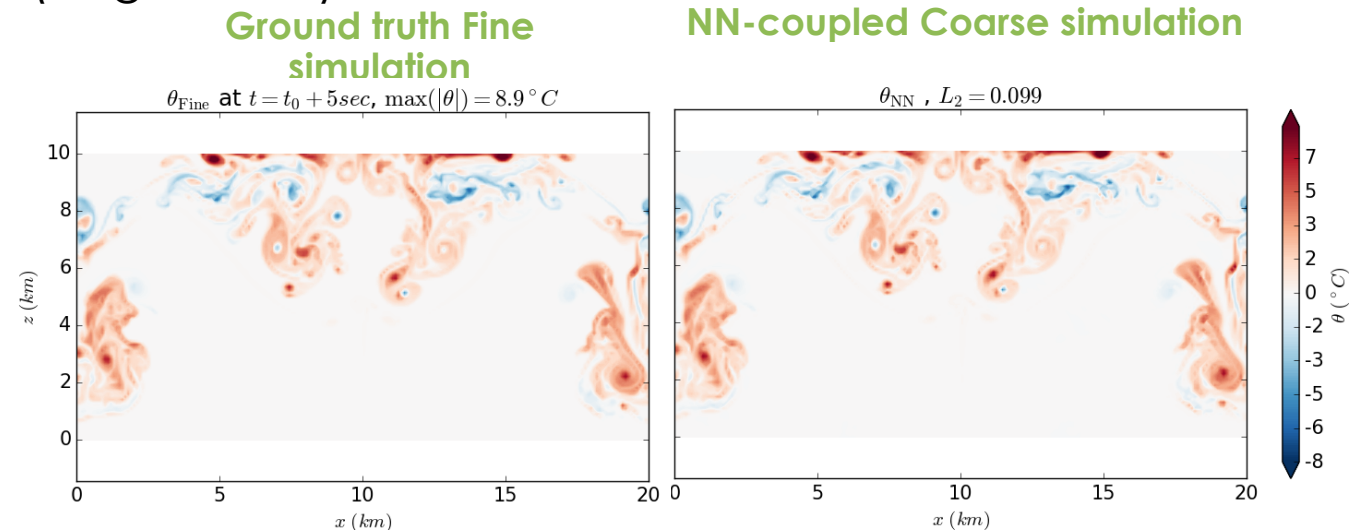
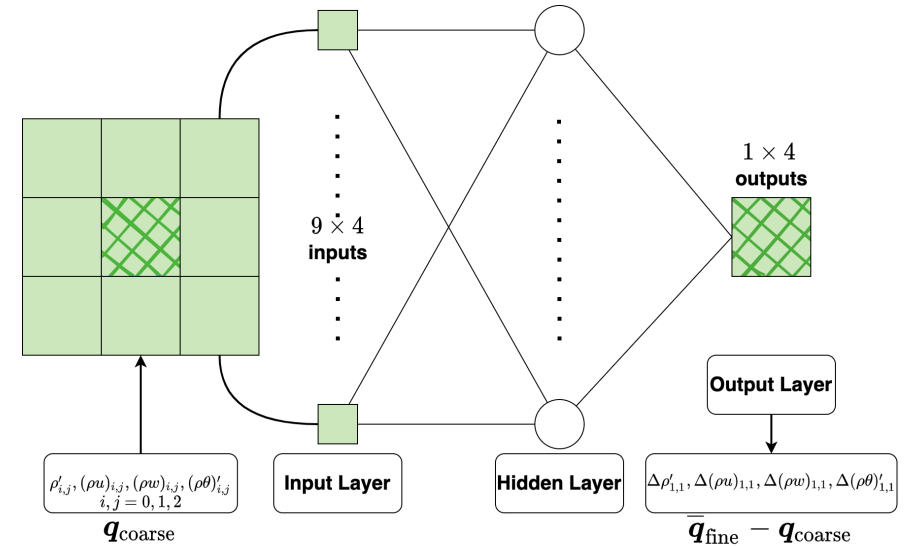
¹OLCF; ²NVIDIA; ³University of California, Irvine

Slide Courtesy of Murali Gopalakrishnan Meena (OLCF)

- Model subgrid-scale effects (dissipative & anti-dissipative) in dry, stratified turbulence for idealized atmospheric flow simulations
- Use neural networks (NN) to model full state difference between coarse & high-resolution simulations (5x grid ratio)

$$\mathbf{q}_{\text{coarse}} \xrightarrow{\text{NN}} \Delta \mathbf{q} = \bar{\mathbf{q}}_{\text{fine}} - \mathbf{q}_{\text{coarse}}$$

- ResNet model predicts accurate short-time emulations of NN-coupled coarse simulations at both laminar & turbulent regimes
- Pytorch C++ API coupled to solver using Performance Portable API (YAKL)



Future plans

- Planning has begun for OLCF-6, the next system after Frontier
 - OLCF-6 will likely include an infrastructure to facilitate analysis of data from instruments and experiments, both on and off ORNL campus (IRI)
- Several demonstration projects through ORNL INTERSECT e.g.:
 - Open Federated Architecture For the Lab of the Future
 - Ion Trap AI/ML Quantum Performance & Routine Optimization (ITAQPRO)
- We are always on the lookout for new partners & applications
 - We can develop new approaches through Directors Discretionary allocations

Summary

- OLCF operates world leading supercomputer systems
- These systems are accelerated, “AI/ML capable”, machines
- We are actively looking to form engagements with partners
- We can help with
 - Preparing software (porting, tuning, workflow deployment, etc.)
 - Training: Online/In person events, GPU Hackathons, Summer Schools, etc.
 - User Advocacy: facilitate strategic discussions and planning
- Contact
 - Balint Joo, Nuclear, Particle and Astrophysics Group: [joob AT ORNL.gov](mailto:joob@ornl.gov)
 - Tom Beck, Section Lead, Science Engagement: [becktl AT ornl.gov](mailto:becktl@ornl.gov)
 - Bronson Messer: Director of Science: [bronson AT ornl.gov](mailto:bronson@ornl.gov)

Acknowledgements:

- 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.