IXPUG Workshop at HPC Asia – Taiwan *(presented remotely)*

Initial Spectrum of Exascale Applications on Aurora

TIM WILLIAMS
Deputy Director
Computational Science Division (CPS)
Argonne National Laboratory

19 February 2025

Abstract

In this talk, I will survey the set of scientific applications that are just beginning their exascale run campaigns on the Aurora supercomputer. I will discuss the work of many project teams on targeting Aurora's Intel GPU accelerators, portability approaches, and performance considerations. These applications sample a broad spectrum of scientific domains, numerical methods, and AI training and inference components.

Aurora

Intel® Data Center GPU Max Series

4th Gen Intel XEON Max Series CPU *with High Bandwidth Memory*

Platform
HPE Cray-Ex

Racks - 166
Nodes - 10,624
CPUs - 21,248
GPUs – 63,744

Interconnect
HPE Slingshot 11
Dragonfly topology with adaptive routing
Cassini NIC, 200 Gb/s (25 GB/s), 8 per node
Network Switch:
25.6 Tb/s per switch (64 200 Gb/s ports)
Links with 25 GB/s per direction

Peak FP64 Performance
≧ 2 exaFLOPS

Memory
10.9PiB of DDR @ 5.95 PB/s
1.36PiB of CPU HBM @ 30.5 PB/s
8.16PiB of GPU HBM @ 208.9 PB/s

Network
2.12 PB/s Peak Injection BW
0.69 PB/s Peak Bisection BW

Storage
230PB DAOS Capacity
31 TB/s DAOS Bandwidth

Aurora

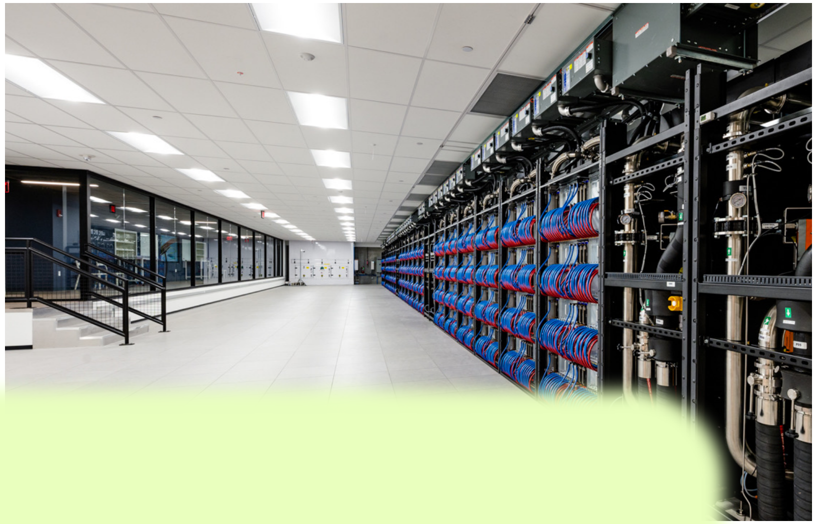

NOW IN PRODUCTION!

#3 in Top500 w/9234 nodes – HPL 1.012 exaFLOPS

#1 in Top500 w/9500 nodes – HPL-MxP 11.6 exaFLOPS

Int
Ma

4th
Se
Ban

Platform
HPE Cray-Ex

Dragonfly topology with adaptive routing
Cassini NIC, 200 Gb/s (25 GB/s), 8 per node
Network Switch:
- 25.6 Tb/s per switch (64 200 Gb/s ports)
- Links with 25 GB/s per direction

0.69 PB/s Peak Bisection BW

Storage
230PB DAOS Capacity
31 TB/s DAOS Bandwidth

Argonne NATIONAL LABORATORY

First-Year Science Projects on Aurora

ALCF Node-Hours

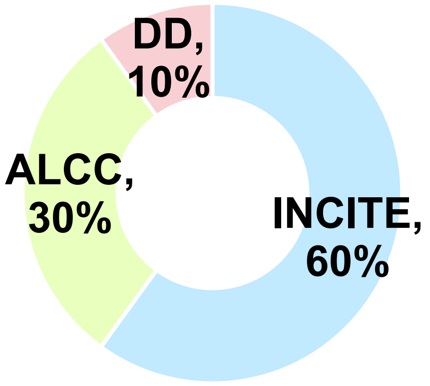

- Aurora Early Science Program (ESP)
  - 19 projects
  - 19.8 million node-hours

AURORA | EARLY SCIENCE PROGRAM

- INCITE
  - 30 projects (+ 4 porting/optimization projects)
  - 22.6 million node-hours

U.S. DEPARTMENT OF ENERGY INCITE LEADERSHIP COMPUTING

Open to researchers from academia, government laboratories, and industry, the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program is the major means by which the scientific community gains access to the nation's fastest supercomputers. The program aims to accelerate scientific discoveries and technological innovations by awarding, on a competitive basis, time on supercomputers to researchers with large-scale, computationally intensive projects that address "grand challenges" in science and engineering.

- ALCC (ASCR Leadership Computing challenge)
  - 13 projects
  - 5 million node-hours

U.S. DEPARTMENT OF ENERGY | Office of Science

People

- Readiness efforts behind 1st year projects on Aurora
  - ESP:
    - 19 project teams with over 200 multi-institutional team members
    - 16 ALCF postdocs
    - 19 ALCF Catalysts/POCs (10% effort for ~7 years)
  - ECP: 7 years, $1.8B, 2800 people
    - 15 (of 21) Applications Development projects engaged with ALCF toward Aurora
    - ALCF POC for each, via ECP Hardware & Integration funding
  - Intel-Argonne Center of Excellence
    - 5 Intel staff
- INCITE, ALCC
  - Always multi-person
  - Typically multi-institutional

Aurora Early Science Program

Science campaigns start 1/27 with workflows incorporating the 3 pillars

Simulation

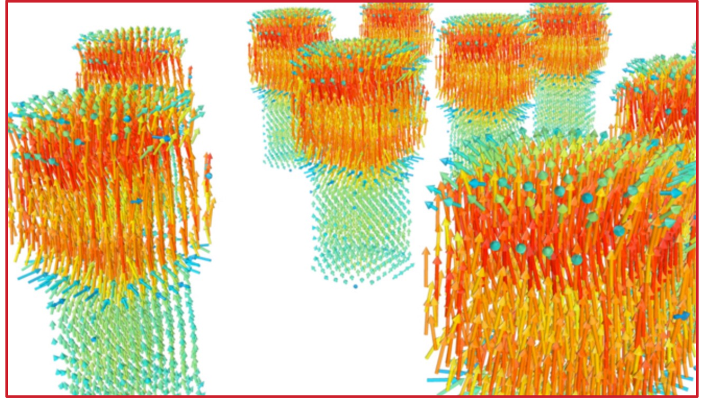

- Reactive, quantum, and classical MD for materials and biophysics
- ITER tungsten impurity ions kinetics—PIC
- Cosmological P^3M + SPH hydrodynamics
- CCSD(T) quantum chemistry
- Quantum Monte Carlo
- FEM CFD, CFD+combustion

Data

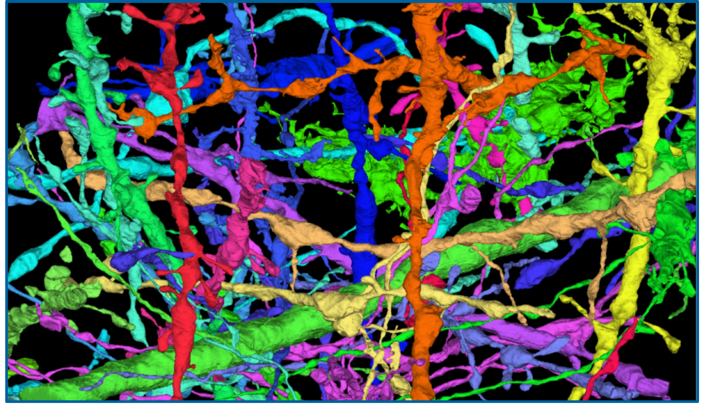

- High-throughput electron microscopy data
- In situ data from blood flow and aerodynamics simulation
- Atlas detector event Monte Carlo analysis
- Cosmological sky survey and P^3M simulation data
- Chemistry kinetics database
- Tokamak experiment data

Learning

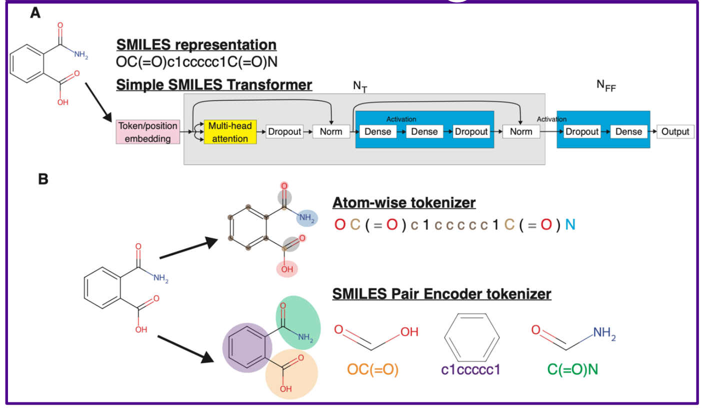

- Simple SMILES transformer for drug design
- Generative model for lattice QCD
- Material band gap predictor
- Tokamak disruption predictor
- Flood fill model for connectomics
- MLP turbulence closure model

Argonne NATIONAL LABORATORY

Computational Science Methods in Y1 Aurora Projects

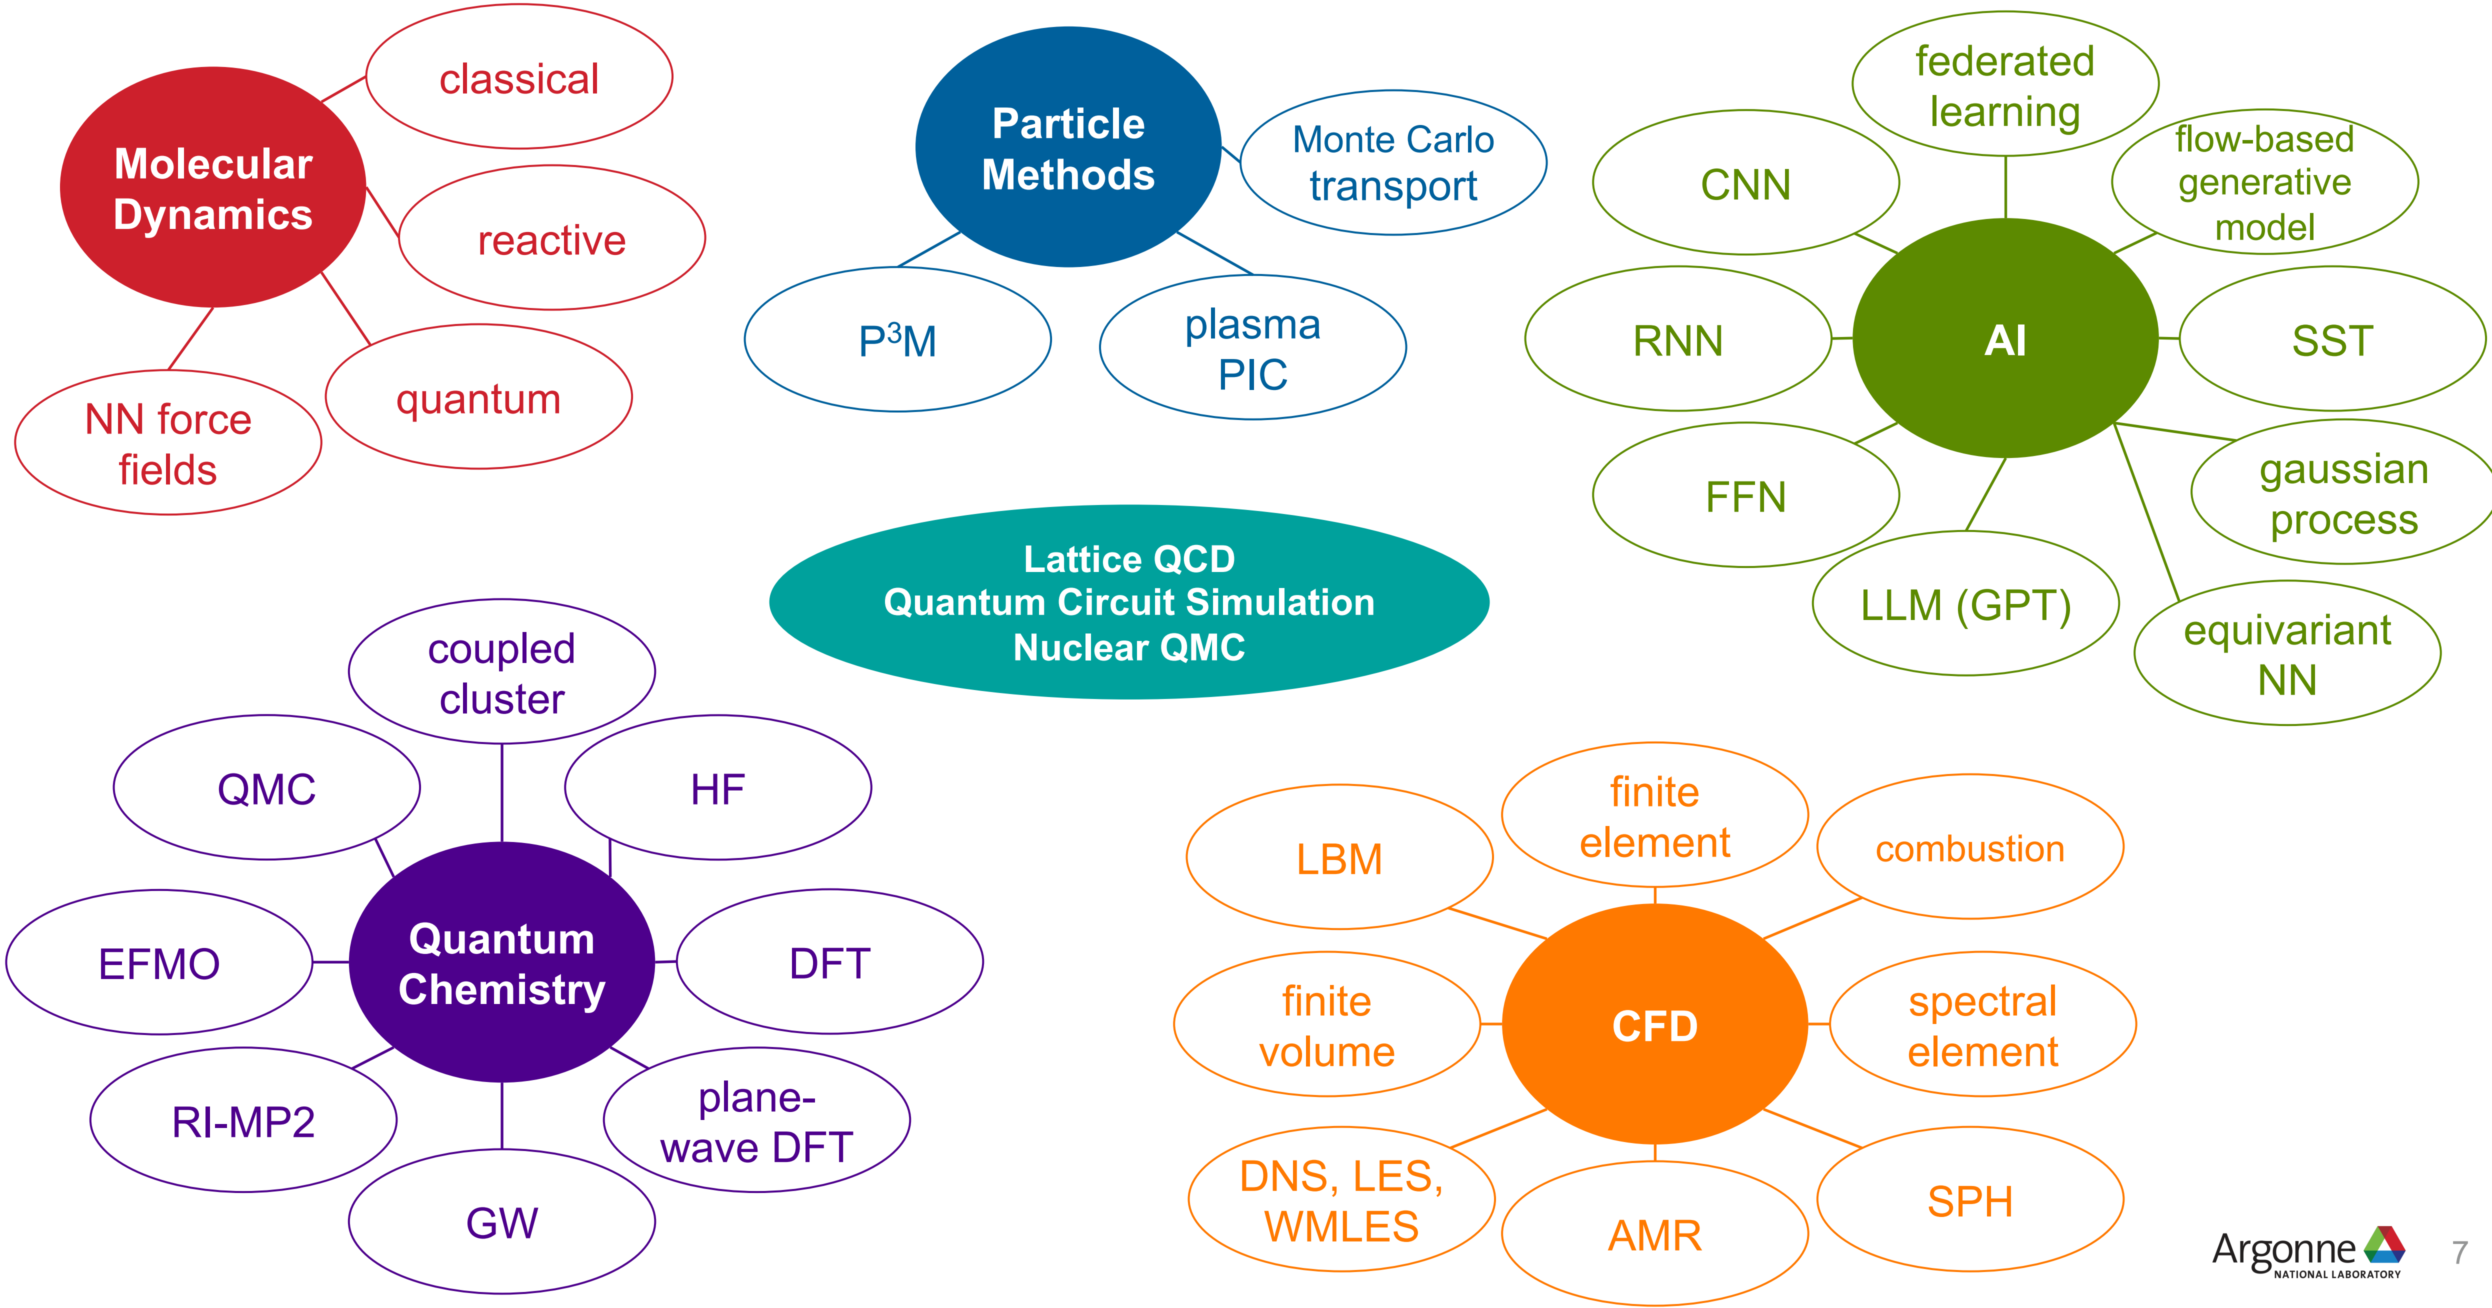

Argonne NATIONAL LABORATORY

Science Domains in Y1 Aurora Projects

- Aerodynamics
- Combustion
- Thermal hydraulics for fusion & fission
- Turbulent Rayleigh-Benard convection
- Buoyancy-driven turbulence
- Wind energy

- Core collapse supernovae
- Binary neutron star mergers
- Core collapse supernovae
- Black hole accretion
- Neutron star remnants

- MOF design
- Carbon capture
- Transition metal chemistry
- Heterogeneous catalysis
- Heterogeneous reactions
- Biofuels

- Membrane transport proteins
- Predictive molecular epdimiology
- Blood flow & transport
- Cancer drug design
- Heteropolymer design
- Brain connectomics
- Pathogenesis

- Fusion radiation transport
- Fusion plasma kinetics
- Tokamak control

- Functional quantum materials
- Fusion device materials
- Multiscale materials modeling
- Quantum materials
- Materials science

- Neuroscience foundation model
- Science foundation model
- Federated learning
- Quantum computing
- AI assistant for simulation

- QCD
- Cosmology
- Nuclear structure and reactions
- Collider data analysis

- Climate modeling
- Runoff inundation

Aurora Performance and Expectations

- DOE Exascale Computing Project (ECP)
  - Achieve 50X improvement in figure-of-merit (FOM) w.r.t. pre-ECP baseline (2017 systems)
  - Aurora: 3 projects measured this:
    - EXAALT – 89X on 1024 nodes
    - ExaSMR – 84X on 512 nodes
    - ExaSky – 277X on 4096 nodes

- Aurora Early Science Program (ESP)
  - Demonstrate INCITE-level computational readiness
    1. Scale efficiently to 20% of full system
    2. Make use of GPUs
    3. Ready to run science in 1-3 months
  - Internal review: 16/19 projects met goal in Fall 2024

Aurora Performance and Expectations

- Argonne Applications Working Group performance tracked single-GPU performance figures of merit across >25 ECP and ESP applications for several years
  - Expectations
    - Comparable performance to AMD MI-250X (Frontier GPU)
    - Better performance than NVIDIA A100 (Polaris GPU)

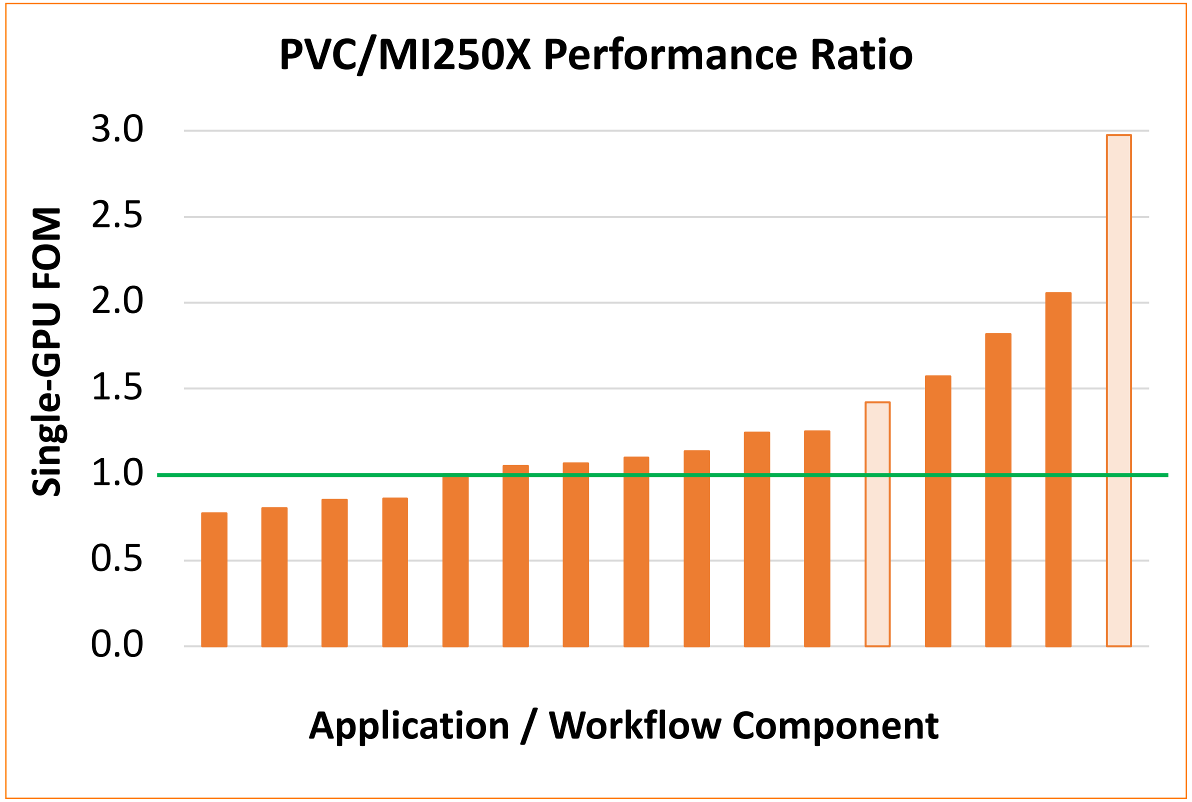

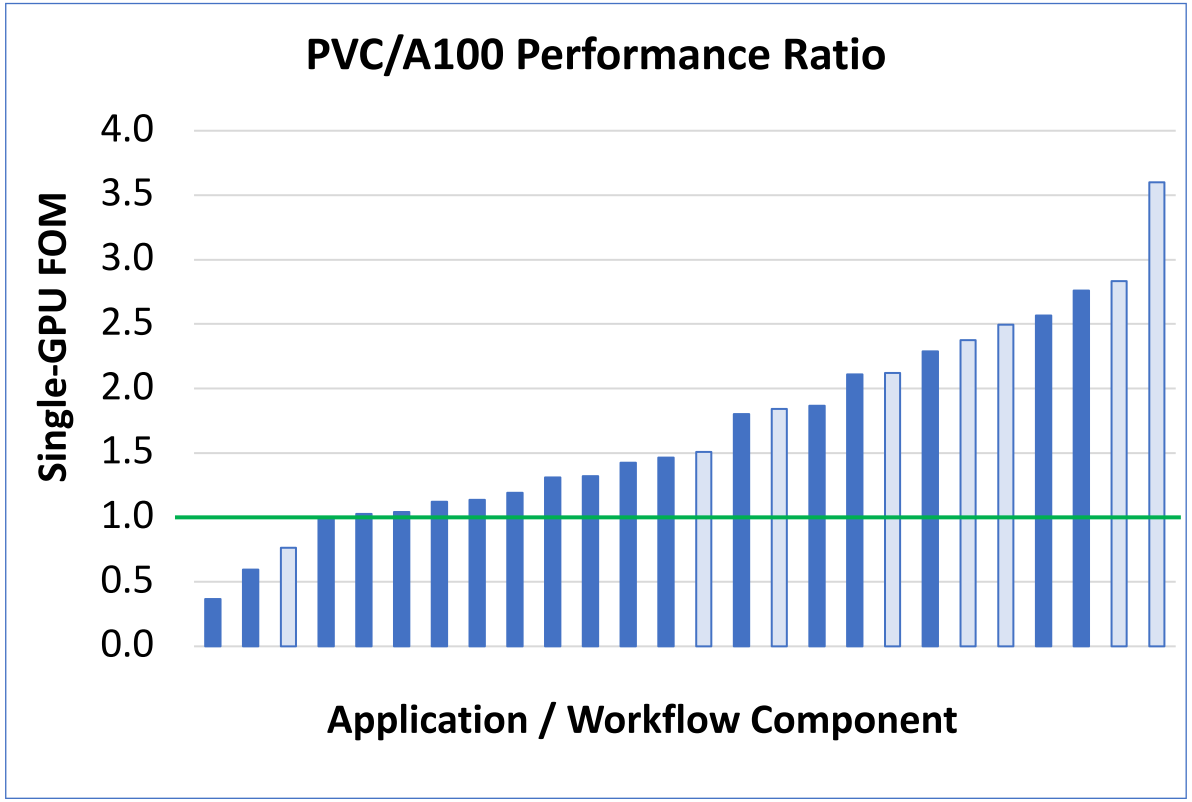

Aurora Programming Models

- **Simulation**
  - Intel oneAPI compilers
  - oneMKL, Aurora MPICH
  - PVC GPU offload
    - OpenMP
    - SYCL
    - HIP (chipStar)
    - CUDA
      - Intel Syclomatic
      - chipStar
    - Kokkos, Raja, libCEED, AMREx, OCCA,PETSc

- **Data & Learning**
  - Python – Numba, NumPy, etc.
  - Deep learning frameworks
    - PyTorch, TensorFlow, JAX, DeepSpeed, Horovod
  - Machine Learning
    - oneDAL, scikit-learn, XGBoost, *etc.*
    - Optimized communication: oneCCL

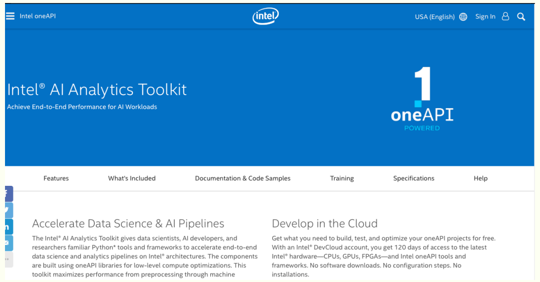

Aurora Programming Models

- **Simulation**
  - Intel oneAPI compilers
  - oneMKL, Aurora MPICH
  - PVC GPU offload
    - OpenMP
    - SYCL
    - HIP (chipStar)
    - CUDA
      - Intel Syclomatic
      - chipStar
    - Kokkos, Raja, libCEED, AMREx, OCCA, PETSc

- **Data & Learning**
  - Python – Numba, NumPy, etc.
  - Deep learning frameworks
    - PyTorch, TensorFlow, JAX, DeepSpeed, Horovod
  - Machine Learning
    - oneDAL, scikit-learn, XGBoost, *etc.*
    - Optimized communication: oneCCL

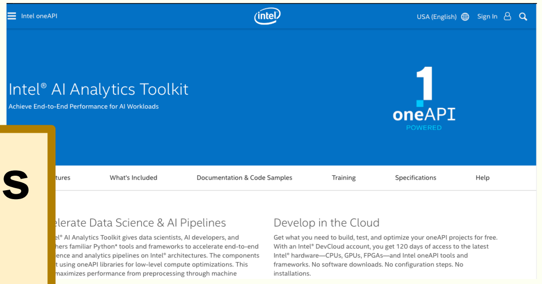

Performance portability across Intel, AMD, NVIDIA GPUs

- Pick a portability layer (Kokkos, SYCL, OpenMP 5, your own library, TensorFlow, PyTorch, ...)
- Work with implementers of layer on GPUs

Programming Model Choices by Y1 Aurora Projects