Application Experience on a GPU-Accelerated Arm-based HPC Testbed

Wael Elwasif

Computer Science & Mathematics Division &
Oak Ridge Leadership Computing Facility (OLCF)

Oak Ridge National Laboratory

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

U.S. DEPARTMENT OF ENERGY

Evaluation Team

ORNL: Wael Elwasif, William Godoy, Nick Hagerty, J. Austin Harris, Balint Joo, Paul Kent, Damien Lebrun-Grandie, Elijah MacCarthy, Verónica G. Melesse Vergara, Bronson Messer, Ross Miller, Sarp Oral, Oscar Hernandez,

Helmholtz-Zentrum Dresden-Rossendorf: Sergei Bastrakov, Michael Bussmann, Alexander Debus, Klaus Steiniger, Jan Stephan, Rene Widera

Georgia Tech: Spencer H. Bryngelson, Henry Le Berre, Anand Radhakrishnan, Jeffrey Young

University of Basel: Florina Ciorba, Osman Simsek

Swiss National Supercomputing Center: Sebastian Keller, Jean-Guillaume Piccinali

Sandia National Lab: Christian Trott

University of Illinois: David Hardy

University of Delaware: Sunita Chandrasekaran

NVIDIA Corp: M. A. Clark, Filippo Spiga, Jeff Hammond, John E. Stone

Why Do We **Care** About ARM Architecture?

- Exploring technologies that might be viable in the next few years
  - Essentially, we're looking for technologies we might use in the next supercomputer

- "AArch64 for HPC" has been around for years
  - Now with more large-scale / leadership-class deployments
  - Astra, Fugaku, Graviton, …

- AArch64 has been on the OLCF's radar since about 2017
  - The Wombat cluster started in 2018
  - Using GPUs on AArch64 started in September of 2019
  - Added A64fx nodes in mid-2020 and NVIDIA ARM DevKit nodes in late-2021

- OLCF doesn't have any AArch64 hardware deployed in production

OAK RIDGE
National Laboratory

ARM: Wombat testbed at NCCS

Hardware:

- 4 HPE Apollo 70 nodes:
  - ThunderX2 CN9980 CPU
  - 2x NVIDIA V100 GPUs

- 16 HPE Apollo 80 nodes
  - Fujitsu A64FX CPUs

- ***8 NVIDIA ARM Developer Kit nodes***
  - ***Ampere Altra, 80 Core CPU***
  - ***2x NVIDIA A100 GPUs***

- Infiniband EDR and HDR

Website:
https://www.olcf.ornl.gov/olcf-resources/compute-systems/wombat/

Software:

- Centos 8.1

- CUDA 11

- Compilers
  - NVHPC/22.1
  - ARM HPC Compiler 22.0.2
  - LLVM 15.0
  - GCC 10.2 and 11.2

- Open MPI 4.0.5

OAK RIDGE
National Laboratory

Comparison of Summit vs Wombat Nodes

System Specs	Summit	Wombat TX2	Wombat NVDevkit
CPU	IBM POWER9™ (2 Sockets * 21 Cores / socket)	Marvell ThunderX2 (2 Sockets * 28 Cores / socket)	**Ampere Altra Q80-30 (80 Cores, 1 Socket)**
GPU	NVIDIA Volta (6 per node)	NVIDIA Volta (2 per node)	**NVIDIA A100 (2 per node)**
RAM	512 GB DDR + 96 GB HBM (16 GB per GPU)	256 GB DDR + 64 GB HBM (32 GB per GPU)	**512 GB DDR4 + 80 GB HBM (40 GB per GPU)**
On-node interconnect	NVIDIA NVLINK2 (50 GB/s) Coherent memory across the node	PCIe Gen3 (16 GB/s) No coherence across the node	**PCIe Gen4 (32 GB/s) No coherence across the node**
System Interconnect	Mellanox Dual-port EDR IB network 25 GB/s	Mellanox Single-port EDR IB network 12.5 GB/s	**2x NVIDIA BlueField-2 E-Series DPU, 200GbE/HDR single-port QSFP56**
On-node Storage	1600GB NVME	480GB SATA SSD	**2 TB SATA Spinning Disk**

OAK RIDGE
National Laboratory

Evaluating ARM + GPU : 20 unique codes

2019 : ThunderX2 + V100

- 12 Codes
- *LAMMPS, MiniSweep, NAMD/VMD*
- BabelStream   Comet
- DCA++   Gamera
- Gromacs   LSMS
- Patatrack   SNAP
- Tea Leaf

2022 : Ampere Altra + A100

- 11 Codes
- *LAMMPS, MiniSweep, NAMD/VMD*
- Exastar   GPU-I-TASSER
- MFC   MILC
- PIConGPU   QMCPACK
- SpecHPC   SPH-EXA2

Exascale Computing Project (ECP) Codes

OAK RIDGE
National Laboratory

NVIDIA ARM DevKit Evaluation Phase 1 : HPC mod/sim

Scientific Applications:

Exastar
Stellar Astrophysics

GPU-I-TASSER
Bioinformatics

LAMMPS
Molecular Dynamics

MFC
Fluid Dynamics

MILC
QCD

NAMD
Molecular Dynamics

VMD
MD Viz.

PIConGPU
Plasma Physics

QMCPACK
Chemistry

SPH-EXA2
Hydrodynamic

Benchmarks & Mini-apps:

SpecHPC 2021
Benchmark

MiniSweep
Radiation Transport

Parallel Programming Models and Tools

Kokkos
C++ Prog. Model

Open MPI
Distributed Prog. Model

CUDA / NVHPC, NSIGHT
Prg. Env and Tools

Arm Compilers Performance Libraries, Allinea
Prog Env and Tools

UCX

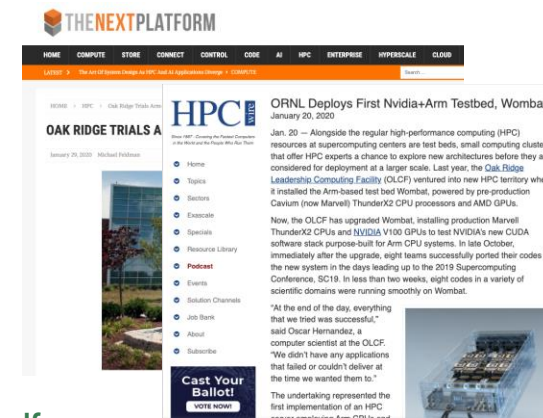OpenACC
Directives for Accelerators

OpenMP

Evaluation done by:

Evaluation Overview : "*Breadth-First study*"

2019

- Focus on software eco-system and programming environment

- Builds on the work done for the 2019 evaluation
  - https://www.olcf.ornl.gov/wp-content/uploads/2019/11/ARM_NVIDIA_APP_EVALUATION-1.pdf

- Apps were typical HPC modeling and simulation apps
  - Phase 1: 10 HPC apps from 9 institutions

- Emphasis mostly on the Ampere nodes, but some teams tested on the other nodes, too.
  - Focus **NOT** on direct comparisons of results from different node types because the hardware was too disparate
  - Some teams were interested in the SVE support on the A64fx CPU's

- *Phase 2 (ongoing): AI & ML workloads*

- *Phase 3 (planned): Focus on Bluefield DPUs*

OAK RIDGE | LEADERSHIP
National Laboratory | COMPUTING FACILITY

ExaStar (thornado)

Austin Harris, Bronson Messer
Oak Ridge National Laboratory

Application information

- Science Domain & Method
  - Computational stellar astrophysics, spectral neutrino radiation transport

- Programming Models and languages
  - Fortran and OpenACC

- HPC systems where app is used in production
 - Summit, Crusher, Frontier*, Perlmutter*, Aurora*

    *Planned

Isotope distributions for early and late stages
of core-collapse supernovae simulation
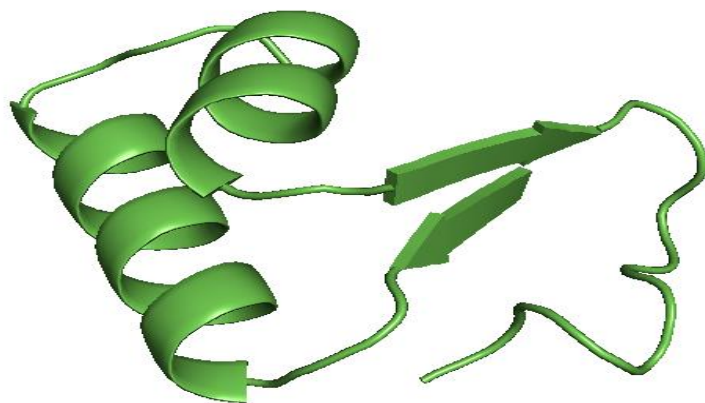
WOMBAT ARM+NVIDIA DevKit Experience

- Easy to transition Summit to Wombat (NVHPC)

- Minimal modifications of build system to use OpenBLAS instead of IBM's ESSL for CPU libraries

- Wombat speedup relative to Summit (1 rank):
  - Altra to POWER9 (multicore): 1.6x – 2.2x
  - Altra+A100 to POWER9+V100: 1.3x – 1.7x

Normalized time of bottleneck kernels.
Wall-clock times above each bar.

Elijah MacCarthy, Verónica Melesse Vergara
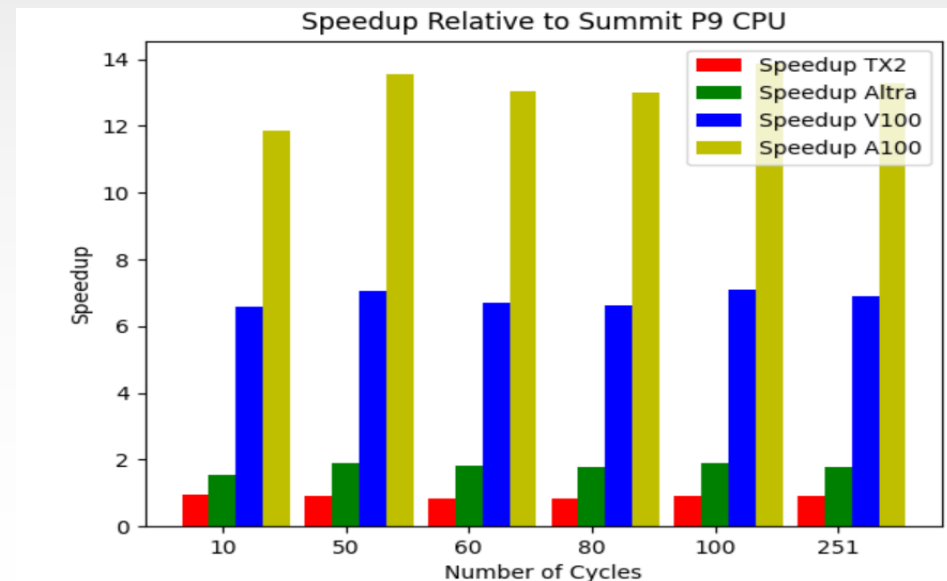Oak Ridge National Laboratory

Application information

- Science Domain & Method
  - Computational Biology, Bioinformatics
- Programming Models and languages
  - Fortran and OpenACC
- HPC systems where app is used in production
  - PSC Bridges2

Predicted structure of 1YV8 protein
from *Crambe hispanica subsp. abyssinica*

WOMBAT ARM+NVIDIA DevKit Experience

- Compiled on Wombat with updated NVHPC 22 compiler.

- Wombat's Ampere Ultra CPU recorded speed-up of 1.8x relative to Summit's P9 CPU.

- A100s on Wombat resulted in 2x speed-up compared to NVIDIA V100s.
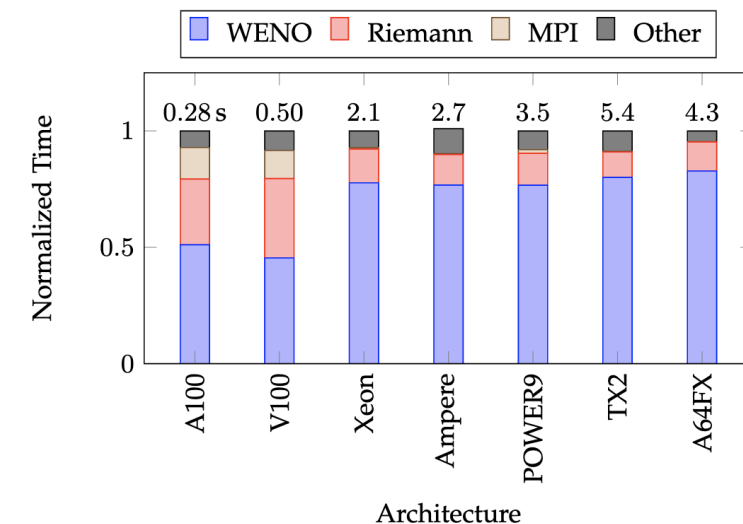
Application information

- Science Domain & Method
  - Multiphase Computational Fluid Dynamics

- Programming Models and languages
  - Fortran90, OpenACC 2.6, MPI/RDMA

- HPC systems where app is used in production
  - OLCF Summit, PSC Bridges2, SDSC Expanse

Bubbles cavitating and collapsing near a wall.
Streamlines are shown.

WOMBAT ARM+NVIDIA DevKit Experience

- Sufficient out-of-the-box support for Arm environments with NVHPC 22, GCC11

- Ampere Arm CPU competitive with latest Xeon CPUs, A100s offer expected speedups

- 2x speed-up over a Summit node with same number of GPUs

Most expensive routines. **MPI** is all communication costs.
Wall time is atop each architecture bar.

LAMMPS and Kokkos

Damien Lebrun-Grandie (ORNL)
Christian Trott (SNL)

Application information

- Science Domain & Method
  - Classical molecular dynamic code with focus on material modeling
  - Domain decomposition + shared memory parallelism

- Programming Models and languages
  - MPI+Kokkos and C++

- HPC systems where app is used in production
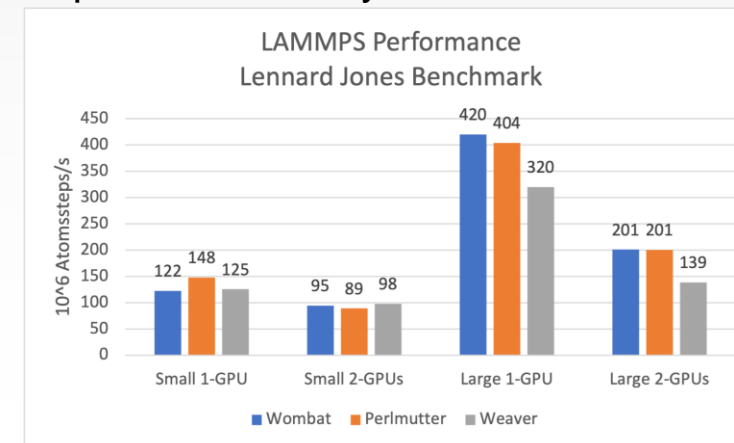 - OLCF Summit, NERSC Perlmutter, etc.

Beryllium depositing on tungsten surface (SNAP Wbe)

Helium bubble forming near tungsten surface

WOMBAT ARM+NVIDIA DevKit Experience

- Kokkos already supported ARM CPUs and all current NVIDIA GPU architectures
  - Wombat: KOKKOS_ARCH="Ampere80"
  - Weaver: KOKKOS_ARCH="Power9;Volta70"
  - Perlmutter: KOKKOS_ARCH="ZEN3;Ampere80"

- Noteworthy observations on the platform and software stack
  - Latency-sensitive scenario performs worse on Wombat
  - Otherwise perform similarly on Wombat and X86

LAMMPS Performance Lennard Jones Benchmark

	Wombat	Perlmutter	Weaver
Small 1-GPU	122	148	125
Small 2-GPUs	95	89	98
Large 1-GPU	420	404	320
Large 2-GPUs	201	201	139

(10^6 Atomssteps/s)
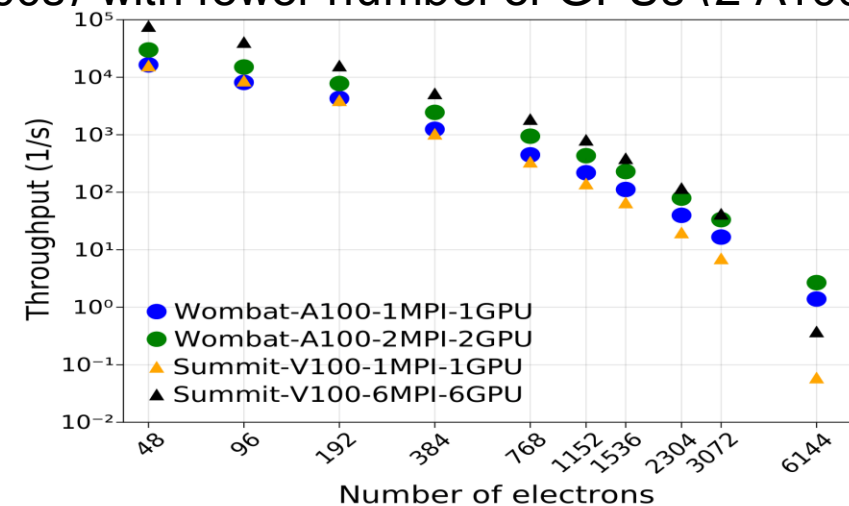
Application information

- Science Domain & Method
  - Highly-accurate Quantum Monte Carlo for predicting the quantum-mechanical properties of materials, molecules, and atoms. Open source.

- Programming Models and languages
  - C++17, OpenMP target offload, CUDA/HIP/SYCL wrappers, Python workflow automation (Nexus). Older CUDA-only implementation tested.

- HPC systems where app is used in production
  - OLCF Summit, ALCF Theta & Polaris, NERSC Cori & Perlmutter, numerous CPU clusters.

Optical Properties of Correlated Defects via QMC

WOMBAT ARM+NVIDIA DevKit Experience

- Sufficient out-of-the-box support for Arm environments with NVHPC22 for GPU

- Ampere Arm CPU competitive with latest Xeon CPUs on single threads, need to evaluate Arm performance libraries with OpenMP

- Competitive with Summit node performance (6 V100s) with fewer number of GPUs (2 A100s)

Greater scalability (higher is better) on Wombat vs Summit nodes using GPUs for NiO benchmark problem

SPH-EXA2

https://github.com/unibas-dmi-hpc/SPH-EXA

F. M. Ciorba (UniBas), S. Keller (CSCS),
J.G. Piccinali (CSCS), O. S. Simsek (UniBas)
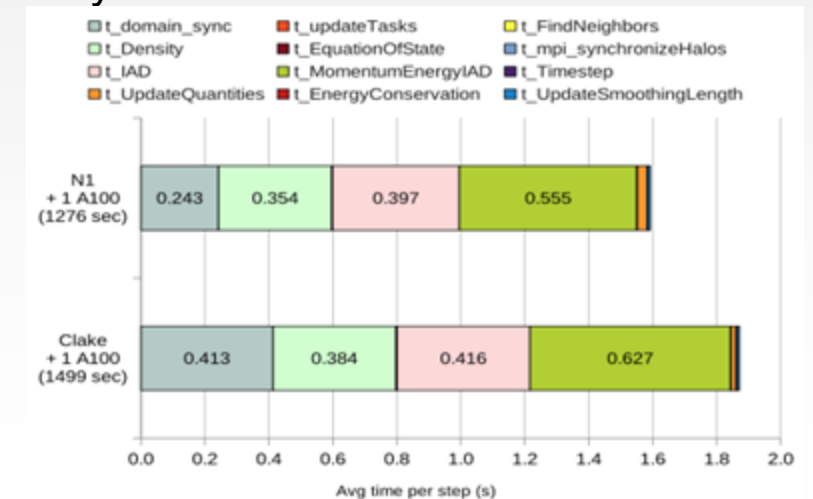
Application information

- Science Domain & Method

  – Smoothed Particle Hydrodynamics and gravity for Astrophysics and Cosmology

- Programming Models and Languages

  – C++ code, parallelized with MPI and OpenMP, and accelerated with CUDA and HIP

- HPC systems where app is used in production

  – Piz Daint, LUMI-G

High resolution subsonic turbulence

WOMBAT ARM+NVIDIA DevKit Experience

- Straightforward porting w/ GCC, ARM, NVCC

- Performance experience and comparisons

  – CPU-only execution performance is limited on ARM CPUs compared to Intel Xeon CPUs

  – ARM Ampere CPUs deliver competitive performance on the GPU-accelerated version

- HPC tools for ARM are quite mature and useful for future ARM systems

Execution times of SPH-EXA2 kernels on ARM N1 and Intel Clake CPUs with NVidia GPUs

PIConGPU

S. Bastrakov[1], M. Bussmann[1], S. Chandrasekaran[2],
A. Debus[1], K. Steiniger[1], J. Stephan[1], R. Widera[1]
[1]Helmholtz-Zentrum Dresden-Rossendorf, [2]University of Delaware
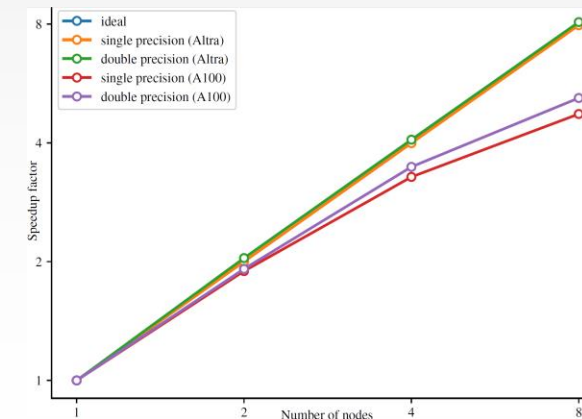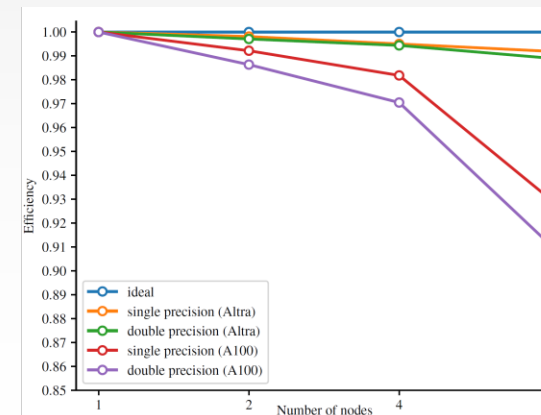
Application information

- Science Domain & Method
    - Radiation physics & particle-in-cell simulations

- Programming Models and languages
    - C++ & *alpaka* abstraction layer over compute APIs (CUDA, HIP, SYCL, OpenMP, …)

- HPC systems where PIConGPU is used for production
    - ORNL Frontier, Summit, Crusher, NERSC's Perlmutter, JSC JUWELS, TUD Taurus

Still image from an uncut LWFA simulation video using Summit and 48 V100s using ISAAC 1.5

WOMBAT ARM+NVIDIA DevKit Experience

- ARM CPUs and NVIDIA GPUs seamlessly supported by *alpaka* (using the OpenMP back-end) → no porting effort required

- Modification of build system required to use *armclang* compiler

- Weak and strong scaling working almost ideal for Ampere Altra CPUs, stronger MPI impact with GPUs due to GPU computation speed

Nick Hagerty, Verónica Melesse Vergara
ORNL

Application information

- Science Domain & Method
  - Portable benchmarking of real-world applications

- Programming Models and languages
  - C, C++, and Fortran (NVHPC only)
  - OpenACC offload via NVHPC 22.1
  - OpenMP TGT via LLVM 15.0.0, NVHPC 22.1

- HPC systems where app is used in production
  - Summit (ORNL), Frontera (Texas Advanced Computing Center), Hemera (Helmholtz-Zentrum Dresden – Rossendorf)
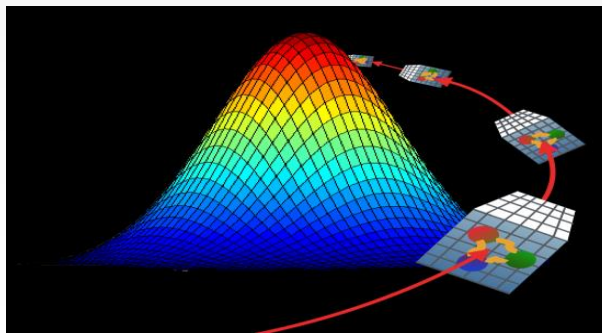
WOMBAT ARM+NVIDIA DevKit Experience

- Availability of both LLVM and NVHPC on both Wombat and Summit eased transition

- Wombat and Summit exhibit consistent OpenMP performance trends

Performance relative to OpenACC
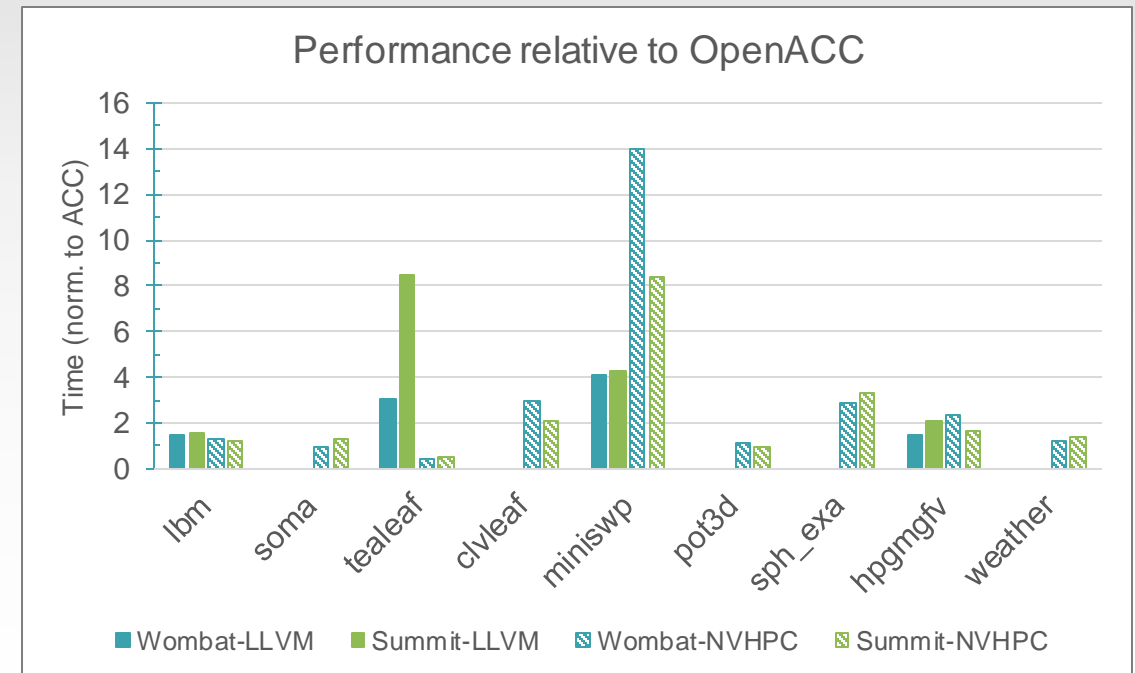
Application information

- Science Domain & Method
  - Lattice Quantum Chromodynamics (LQCD)
  - Dominated by iterative linear solve (CG) of discretized PDE (stencil) on 4-d grid
- Programming Models and languages
  - Utilizes QUDA library for GPU acceleration
  - C (MILC) and CUDA C++ (QUDA)
- HPC systems where app is used in production
  - OLCF Summit, NERSC Perlmutter, etc.

Lattice QCD gauge generation codes like the MILC benchmark, sample the strong force field distribution in the vacuum, by generating one gauge field after the other using a process of Hybrid Monte Carlo updates.

WOMBAT ARM+NVIDIA DevKit Experience

- Both MILC and QUDA compiled out of the box on Wombat with no changes required
- Benchmarked with NERSC MILC generation tests
- Compared to AMD Rome + NVIDIA A100 system
  - Wombat's Ampere Ultra CPU faster than AMD ROME for running MILC CPU code
  - Single GPU performance identical
  - Scaling limited by lack of p2p on Wombat
  - Stencil with/without communication:

D – double
S – single
H – half
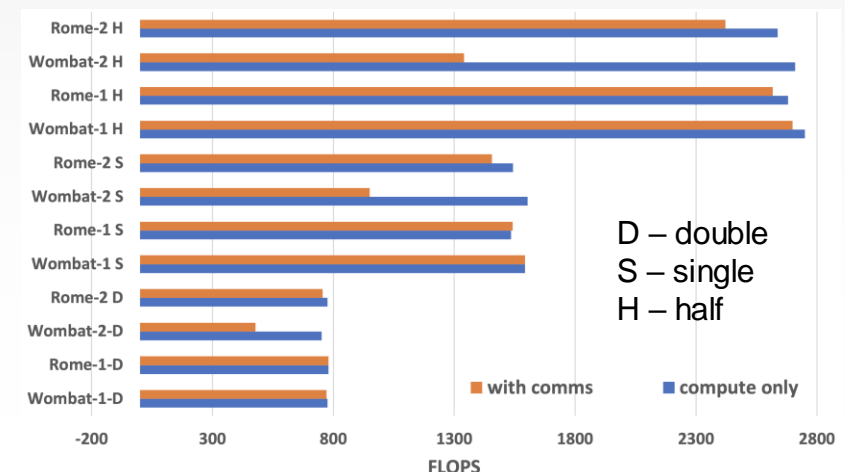
Observations: What Worked, What Didn't, What Are The Gaps

- Good news: Every team was able to build and run. Most teams required no special effort.
  - Most apps have **not** been optimized for every last FLOP of performance, though. (NAMD & VMD are exceptions – they are heavily optimized!)

- The software ecosystem has matured significantly since 2019
  - CUDA on AArch64 is officially supported. (The 2019 work used an early beta.)
  - OpenACC/OpenMP offloading works
  - More mature Spack package management on Arm

- Profiling tools were not exercised much but were used (NSIGHT, etc)
  - More effort needed to fully evaluate platform readiness.
  - Optimization not major focus of this effort

OAK RIDGE
National Laboratory | LEADERSHIP COMPUTING FACILITY

Observations: What Worked, What Didn't, What Are The Gaps

- ## Containers:
  - Not needed for this phase, more important for AI/ML in Phase 2

- ## GPU Direct, CPU-GPU Latency
  - Impact on kernels with heavy communication and data transfer dependency.

- ## Offloading programming models
  - Cuda, OpenACC (NVHPC), OpenMP (LLVM)
  - Offloading using C++ std library: need production app

- ## DPU's (in HPC) yet to be explored
  - Use case co-design - focus of Phase 3

OAK RIDGE | LEADERSHIP
National Laboratory | COMPUTING FACILITY

Thank You

OAK RIDGE
National Laboratory | LEADERSHIP COMPUTING FACILITY