OAK RIDGE National Laboratory

Update on Frontier and Early Science

Bronson Messer
Director of Science
Leadership Computing Facility
Oak Ridge National Laboratory

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

U.S. DEPARTMENT OF ENERGY

Oak Ridge National Laboratory's Journey from Petascale to Exascale

Mission: Providing world-class computational resources and specialized services for the most computationally intensive global challenges	Vision: Deliver transforming discoveries in energy technologies, materials, biology, environment, health, etc.

Steady progress of 10x per generation

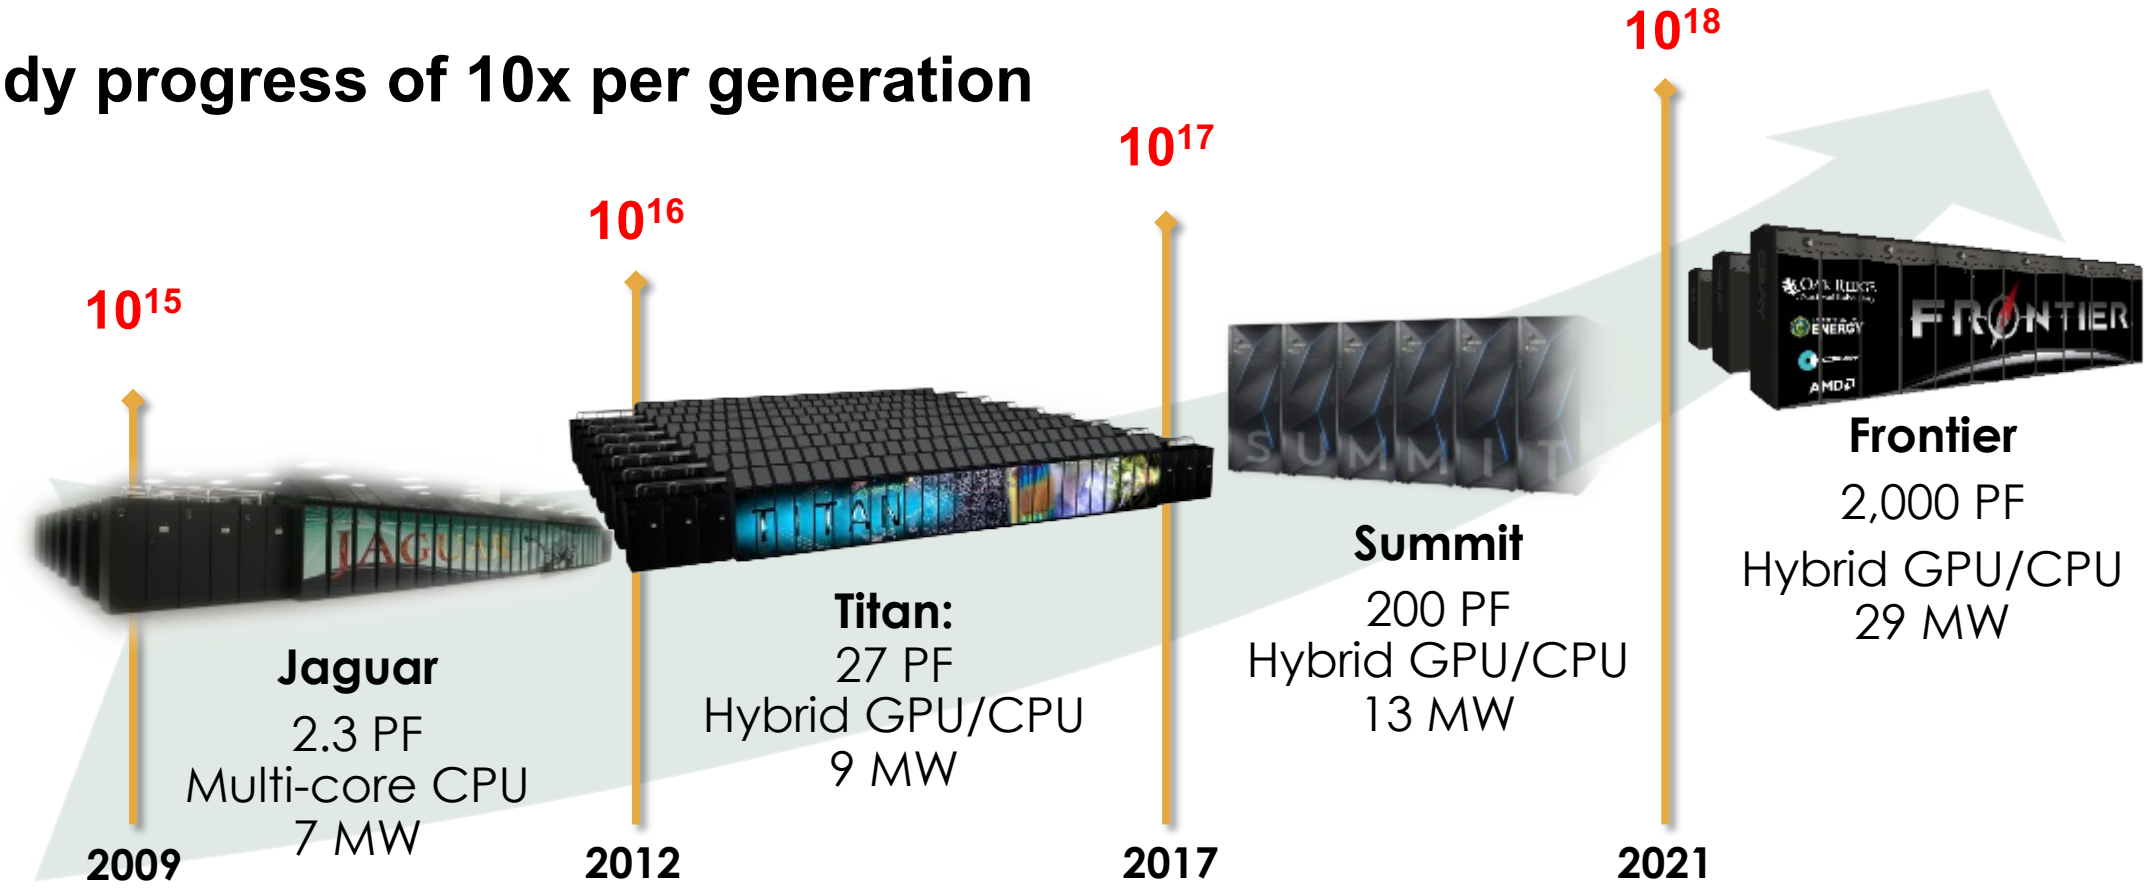

Energy-efficient computing – Frontier achieves 14.5 MW per EF

Since 2009 the biggest concern with reaching Exascale has been energy consumption

- **ORNL pioneered GPU use in supercomputing** beginning in 2012 with Titan thru today with Frontier. Significant part of energy efficiency improvements.
- **ASCR [Fast, Design, Path] Forward vendor investments** in energy efficiency (2012-2020) further reduced the power consumption of computing chips (CPUs and GPUs)..
- **200x reduction in energy per FLOPS** from Jaguar to Frontier at ORNL
- ORNL achieves additional energy savings from using warm water cooling in Frontier (32 C).
  ORNL Data Center PUE= 1.03

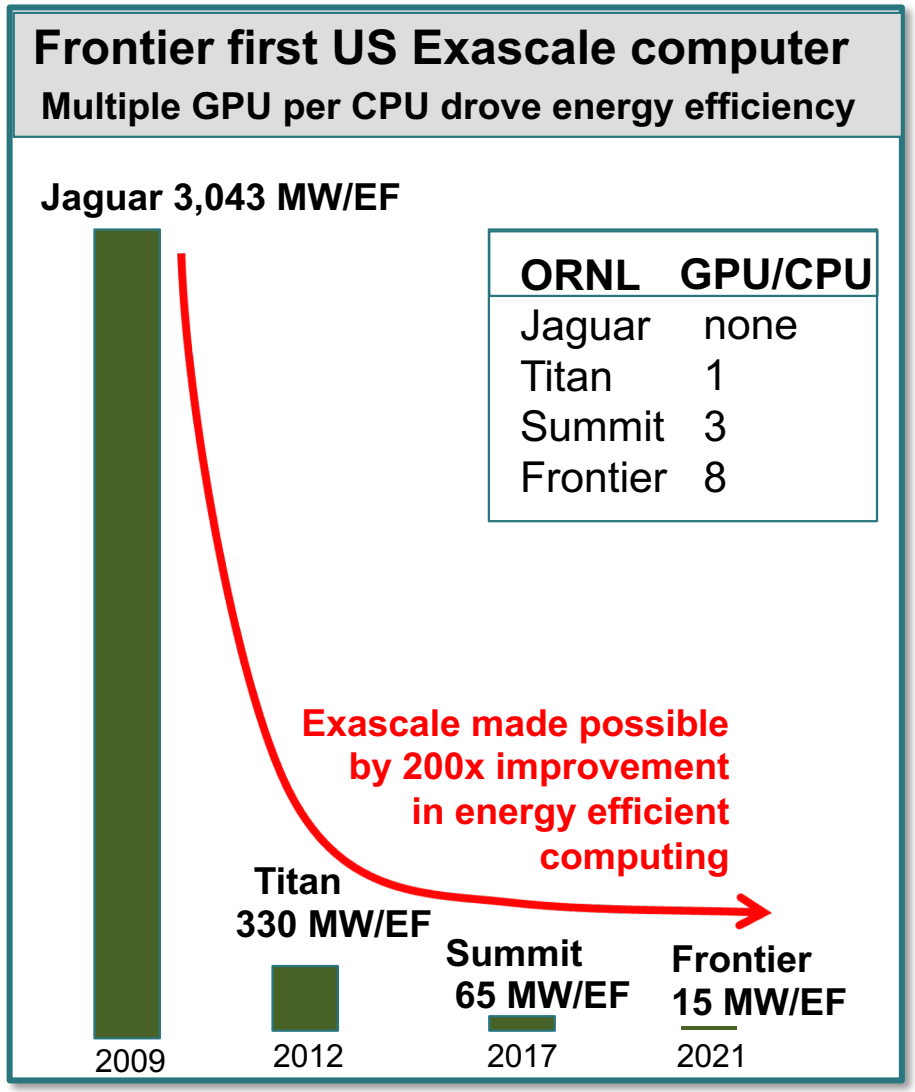

Frontier Overview

Extraordinary Engineering

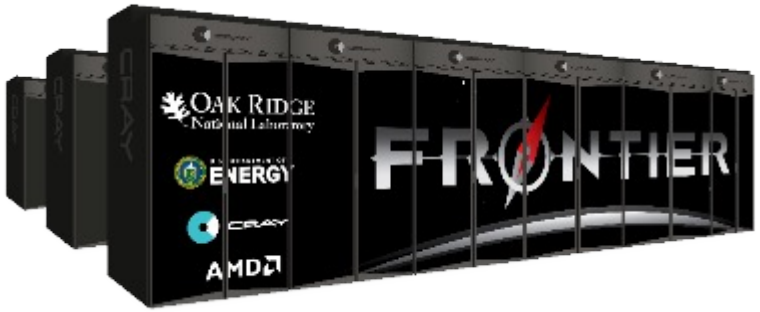

System

- 2 EF Peak DP FLOPS
- 74 compute racks
- 29 MW Power Consumption
- 9,408 nodes
- 9.2 PB memory
  (4.6 PB HBM, 4.6 PB DDR4)
- Cray Slingshot network with dragonfly topology
- 37 PB Node Local Storage
- 716 PB Center-wide storage
- 4000 ft^2 foot print

Built by HPE

Olympus rack

- 128 AMD nodes
- 8,000 lbs
- Supports 400 KW

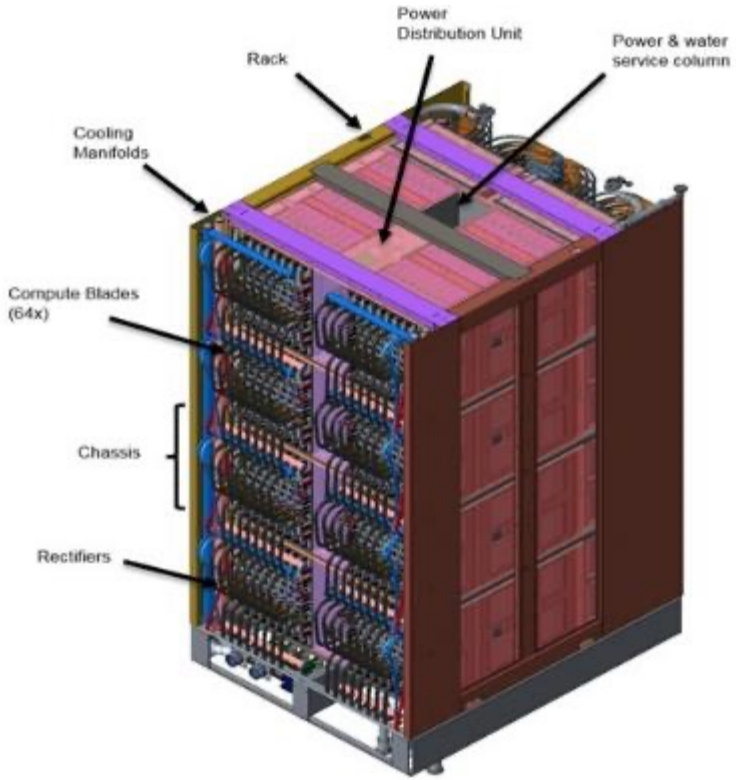

Powered by AMD

AMD node

- 1 AMD “Trento” CPU
- 4 AMD MI250X GPUs
- 512 GiB DDR4 memory on CPU
- 512 GiB HBM2e total per node
  (128 GiB HBM per GPU)
- Coherent memory across the node
- 4 TB NVM
- GPUs & CPU fully connected with AMD Infinity Fabric
- 4 Cassini NICs, 100 GB/s network BW

Compute blade

- 2 AMD nodes

All water cooled, even DIMMS and NICs

Frontier multi-tier storage system

	Capacity	Performance	
Multi-tier I/O Subsystem		**Read**	**Write**
	37 PB Node Local Storage	65.9 TB/s	62.1 TB/s
		11 Billion IOPS	
	11 PB Performance tier	9.4 TB/s	9.4 TB/s
	695 PB Capacity tier	5.2 TB/s	4.4 TB/s
	10 PB Metadata	2M Transactions per sec	

Two 2TB SSD NVM per node Local Storage (Flash)

Gazelle SSD Storage board (Performance Tier and Metadata)

Moose HDD Storage board (Capacity Tier)

During Frontier build -- the chip shortage hit in earnest!

When HPE began ordering parts, suppliers said the lead time on orders was increasing an additional 6-12 months.

60 Million parts needed for Frontier

685 Different part numbers used in Frontier

167 Frontier part numbers affected by the chip shortage
(more than 2 million parts from dozens of suppliers worldwide)

12 Part numbers blocked building the first compute cabinet

15 Part numbers shortage for AMD building all the MI200 cards for Frontier

It wasn't exotic parts like CPUs or GPUs, rather parts needed by everyone – in cars, TVs, electronics, such as voltage regulators, oscillators, power modules, etc.

Getting Frontier ready for Early Science

- As we saw with Titan and Summit, it takes a number of months to get all the hardware and software stabilized
- HPE continues testing and stabilization of Frontier and its file system
- Early Science Teams in CAAR and ECP got access to the “Crusher” Test & Development system in November 2021.
- Rest of ECP users (~800) given Crusher access January 2022.

ECP is scheduled for full Frontier access July 2022
INCITE use of Frontier scheduled for January 2023

“Crusher” TDS system

- 2 cabinets of Frontier HW
- 192 nodes
- Slingshot 11 w/ Cassini
- Same Software as Frontier

Crusher (Frontier test and development system)

- 2 cabinets, the first with 128 compute nodes and the second with 64 compute nodes, for a total of 192 compute nodes.
- Each node
  - One 64-core AMD EPYC 7A53 CPU
  - 512 GB of DDR4 memory.
  - Four AMD MI250X, each with 2 Graphics Compute Dies (GCDs) for a total of 8 GCDs per node
  - Connected with 4 HPE Slingshot 200 Gbps NICs
- Kept in rough sync with Frontier SW stack

CAAR

The **Center for Accelerated Application Readiness (CAAR)** is the primary OLCF program to achieve and demonstrate application readiness

- ***Build on the experience from the successful CAAR programs for OLCF-3 (Titan) and OLCF-4 (Summit)***
- ***CAAR project resources***
  - Dedicated collaboration with OLCF staff
  - Support and consultation from other project personnel, particularly from the Programming Environment and Tools area, and the vendor Center of Excellence
  - OLCF Postdoctoral fellows (both during application readiness and early science)
  - Allocations to available compute resources (Summit, early access systems)

Characteristics of CAAR Projects

Application	Programming languages	Scientific libraries used	I/O	Algorithms	Initial parallelization
Cholla	C++	None.	HDF5	Finite volume hydrodynamics	MPI, CUDA
NAMD	C++	FFTW (node-level)	VMD (custom)	MD, PME	CHARM++, CUDA
LSMS	F90/C++	BLAS, LAPACK, FFTW	HDF5	Dense Linear Solvers, Coupled ODE, Poisson Eq., Monte Carlo	MPI+CUDA
CoMet	C++	cuBLAS, MAGMA	None	2-way and 3-way Proportional Similarity Method and Custom Correlation Coefficient	MPI+OpenMP, CUDA
GESTS	F90	FFTW	HDF5	Fourier pseudo-spectral methods	MPI+OpenMP 4.5
NUCCOR	F90 + F2008; C	BLAS, LAPACK	HDF5	CCSD + CCSDT, Hartree-Fock, Sparse and dense linear algebra (eigensolvers)	MPI+OpenMP, CUDA
PIConGPU	C++	Alpaka, SOLLVE	ADIOS	PIC	MPI+OpenMP, CUDA/HIP/TBB thru Alpaka
LBPM	C++	Zlib	SILO, HDF5	Lattice Boltzmann methods	MPI, CUDA

Introduction & Overview of Cholla

Cholla (Computational Hydrodynamics on Parallel Architecture) is a hydrodynamics code developed to run natively on GPU.

Features:

- 1st, 2nd, 3rd order space spatial reconstruction
- Exact, Roe, HLLC Riemann Solvers
- CTU & Van Leer integrator
- Optically-thin radiative heating / cooling
- Additional (non-hydro) features:
  - self-gravity
  - particle-based gravity (for dark matter)
  - magnetic-fields (in development)

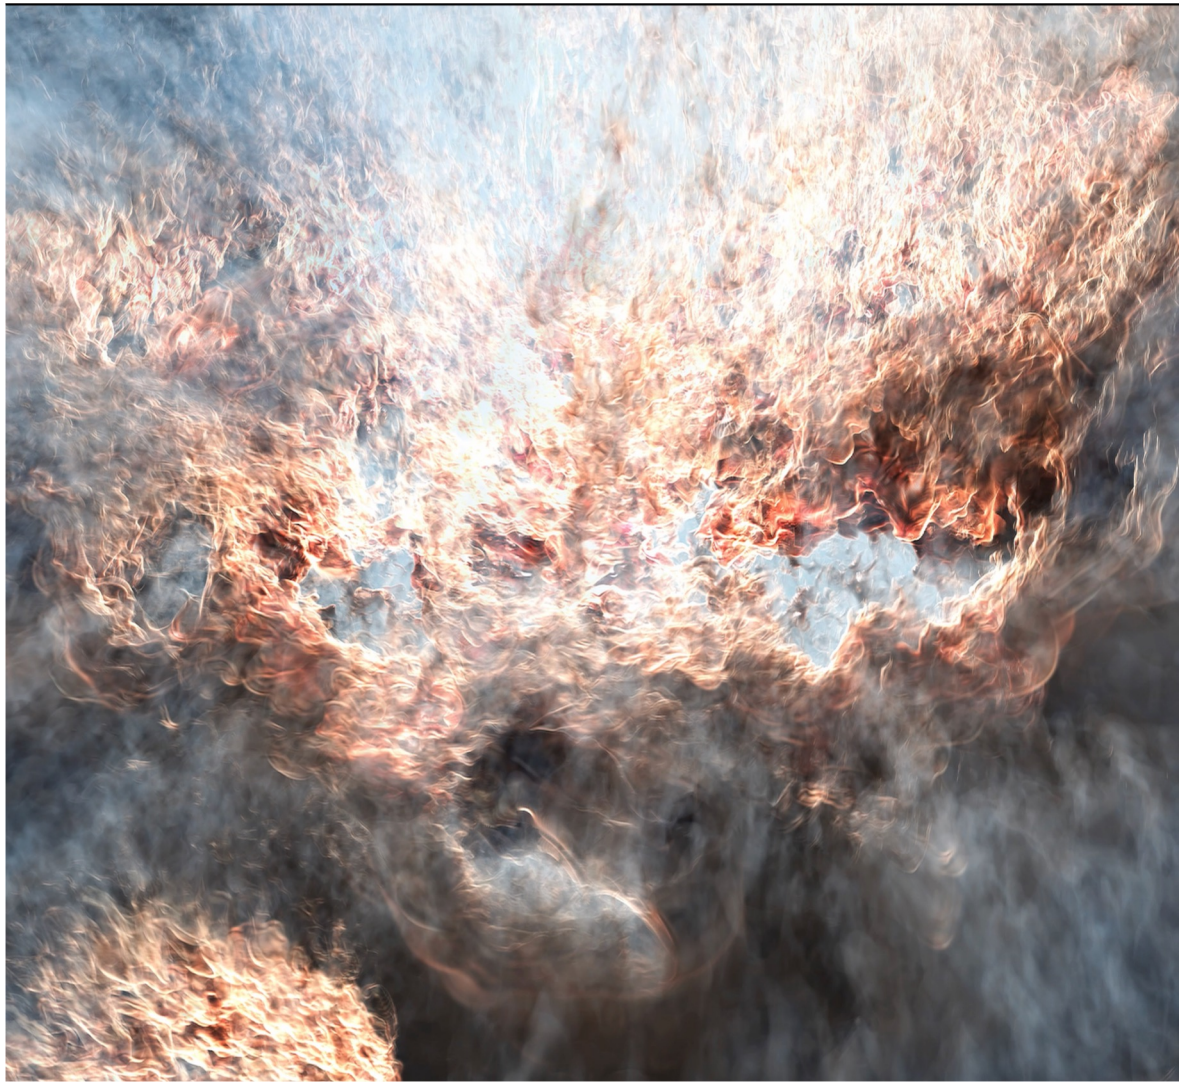

CAAR Cholla Status (Feb 2022)

- Total speedups: **~16X** on 64 GPUs on Frontier hardware (Crusher) from baseline (see plot).

- Software development contributed to ~5X speedups on Summit (blue bars on the plot). Major highlights:
  - Made hydro grid fully GPU resident
  - Exploited GPU-aware MPI
  - Ported gravity solver to GPU
  - Ported particle solver to GPU

- Hardware improvements from Summit to Crusher: ~3X speedups

- Pending: Scaling up to the full Frontier

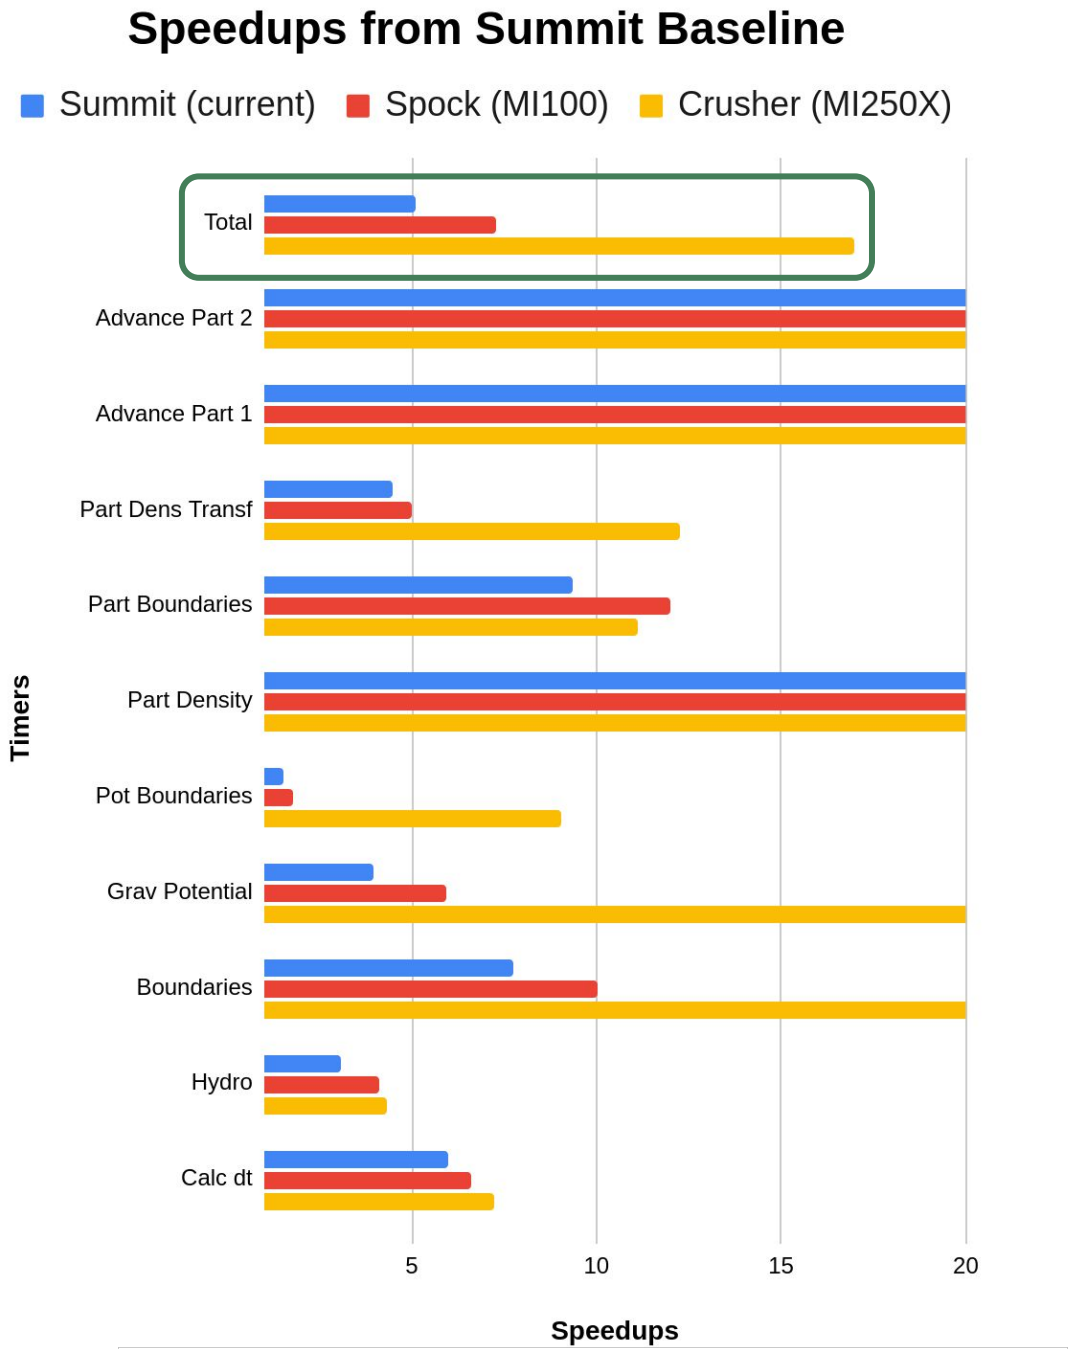

CoMet Application Overview

- CoMet is a data analytics application used to find epistatic and pleiotropic relationships in large genomic datasets
- Applications include human health (opioid addiction, Alzheimer's disease, etc.), bioenergy (poplar, switchgrass), others
- Peak performance is over 2 ExaOps mixed precision on Summit
- Gordon Bell Prize Winner 2018

https://commons.wikimedia.org/wiki/File:DNA_con_GGN.jpg

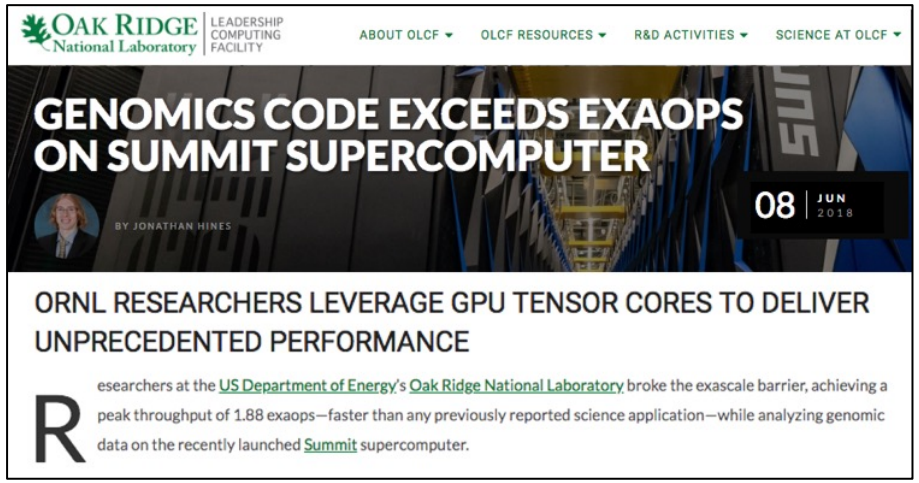

CoMet CAAR update: Crusher, Frontier

- 3-way CCC method has been scaled to 3,210 nodes of **Frontier**
- Mixed precision GEMM from rocBLAS is achieving over 157 TF per GCD
- Trendline indicates achieving FOM target on full Frontier

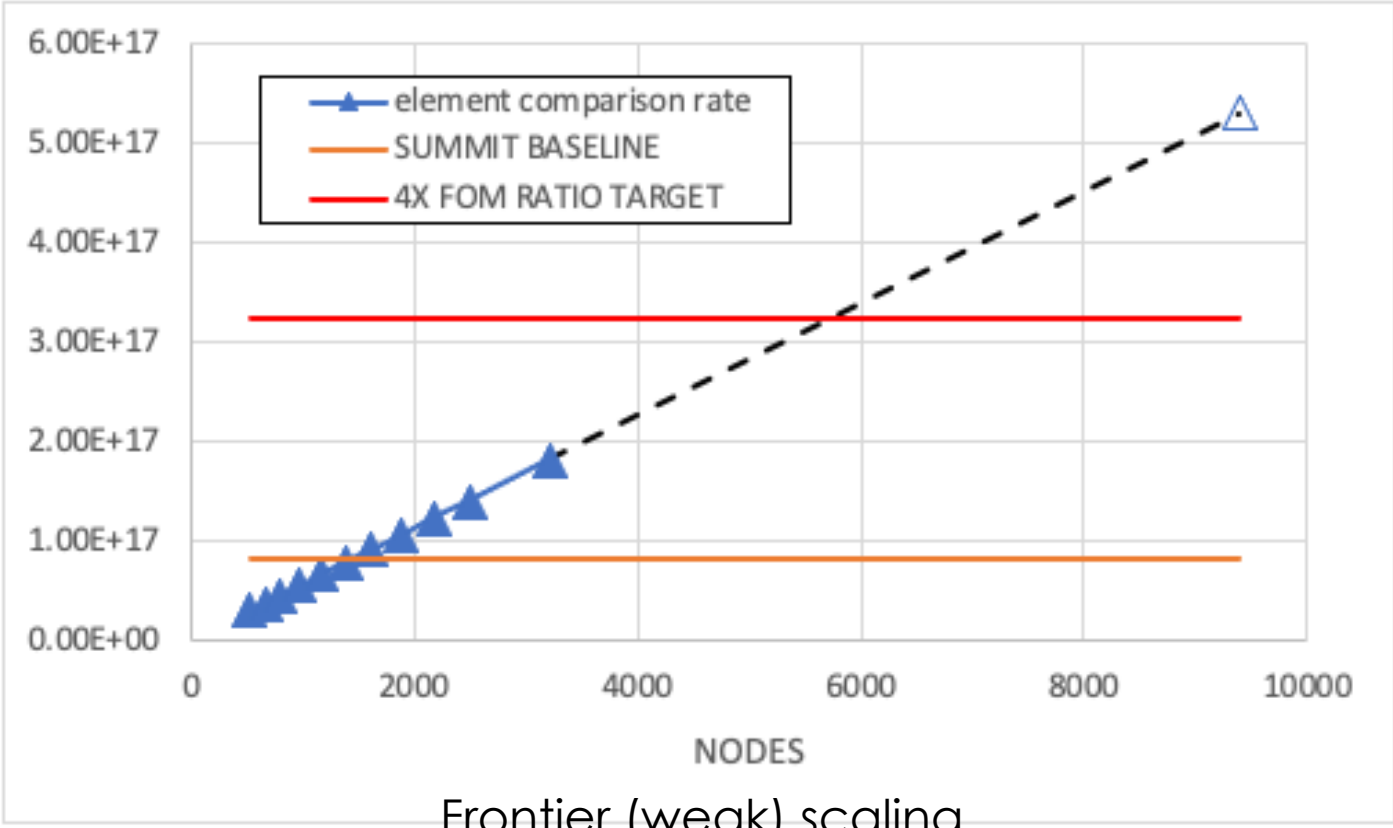

Frontier (weak) scaling

Introduction/Overview of GESTS PSDNS

- A Fourier-pseudo-spectral-based direct numerical simulation code to study fundamental behavior of turbulent flows.
- Written in Fortran
  - CUDA-Fortran on Summit
  - rocFFT, OpenMP, HIP functions on Crusher
- Extensive use of FFTs
  - cuFFT or rocFFT for primary computations on GPUs
- Majority of runtime spent in communication
  - Asynchronous FFT ops and MPI_AllToAll to mitigate bottleneck

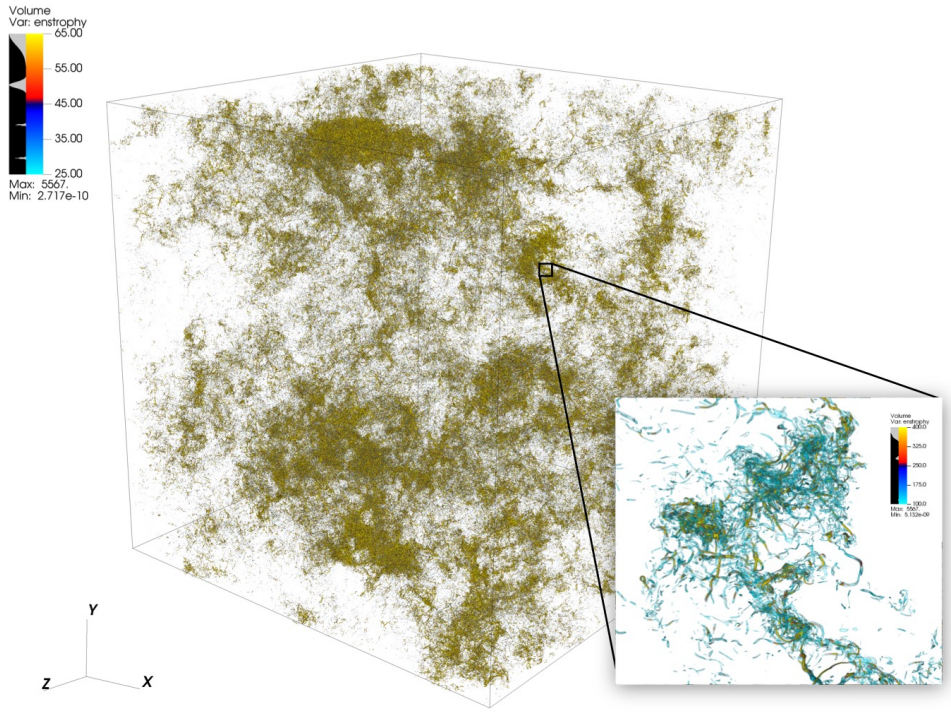

Enstrophy

GESTS: Highlights of work to date

- On Crusher
  - Spent last 6 months developing a 2-D decomposition version of the FFT3D and PSDNS codes for better flexibility with both the number of MPI ranks and memory management
  - Running the 2048^3, 4096^3, and 8192^3 problems on 1, 8, and 64 nodes respectively
  - 16384^3 and 32768^3 problems are projected to run on 512 and 4096 Frontier nodes
  - **Running the 32768^3 problem on 4096 Frontier nodes would satisfy GESTS problem size objective on less than half of Frontier**
- Current issues
  - Slower-than-expected MPI comms between multiple nodes
  - Slower-than-expected C2C forward transforms for lengths > 4096
    - Filed OLCF Ticket and coordinating the investigation with Alessandro Fanfarillo (AMD)
- OpenMP Offloading used primarily for data movement/management
  - Slow OpenMP Offload performance while packing/unpacking the send/recv buffers

Latest results from Crusher - GESTS

FFT-Slab code:

- Fortran 95 + iso_c_bindings
- 1-D decomposition
- 3-D forward/inverse FFTs
- GPU memory intensive
- MPI comm intensive
- OpenMP offloading used to manage data movement and various scaling operations

Test Problem on single node:

- 1024x1024x1024 grid points
- 3 variables per grid point
- 4 MPI ranks with either:
  - 1 V100 GPU per MPI rank
  - 1 MI250x GCD per MPI rank
- Both cases used cuda/hip launch blocking to ensure correct answers
- Strategic use of cuda/hip device synchronization is being tested for better performance w/o explicit launch blocking

Host-based MPI Comms (seconds/step)	V100 (Summit)	MI250x (Crusher)	Speedup
Host-based MPI All-to-All	1.3838	0.9651	1.43x
Data copy to/from device	1.3941	1.0872	1.28x
Pack/Unpack on GPUs	0.1989	0.0963	2.07x
Other computations on GPUs	0.2070	0.1270	1.63x
FFTs on GPUs	0.1145	0.0232	4.94x
Total wall time per step	3,2983	2.2988	1.43x

GPU-based MPI Comms (seconds/step)	V100 (Summit)	MI250x (Crusher)	Speedup
GPU-based MPI All-to-All	1.8187	0.1318	13.80x
Pack/Unpack on GPUs	0.1983	0.0962	2.06x
Other computations on GPUs	0.2066	0.1271	1.63x
FFTs on GPUs	0.1147	0.0232	4.94x
Total wall time per step	2.3383	0.3783	6.18x

Introduction/Overview of NUCCOR

NUCCOR (NUclear Coupled-Cluster Oak Ridge) is a nuclear many-body code that solves the Schrödinger equation approximately for many nucleons (protons and neutrons) using both two- and three-body interactions as input. Many different approximations are implemented following the coupled-cluster hierarchy of methods. In coupled-cluster theory, additional precision is achieved by introducing higher-order corrections to the many-body wavefunction, which increases both memory usage and computational complexity.

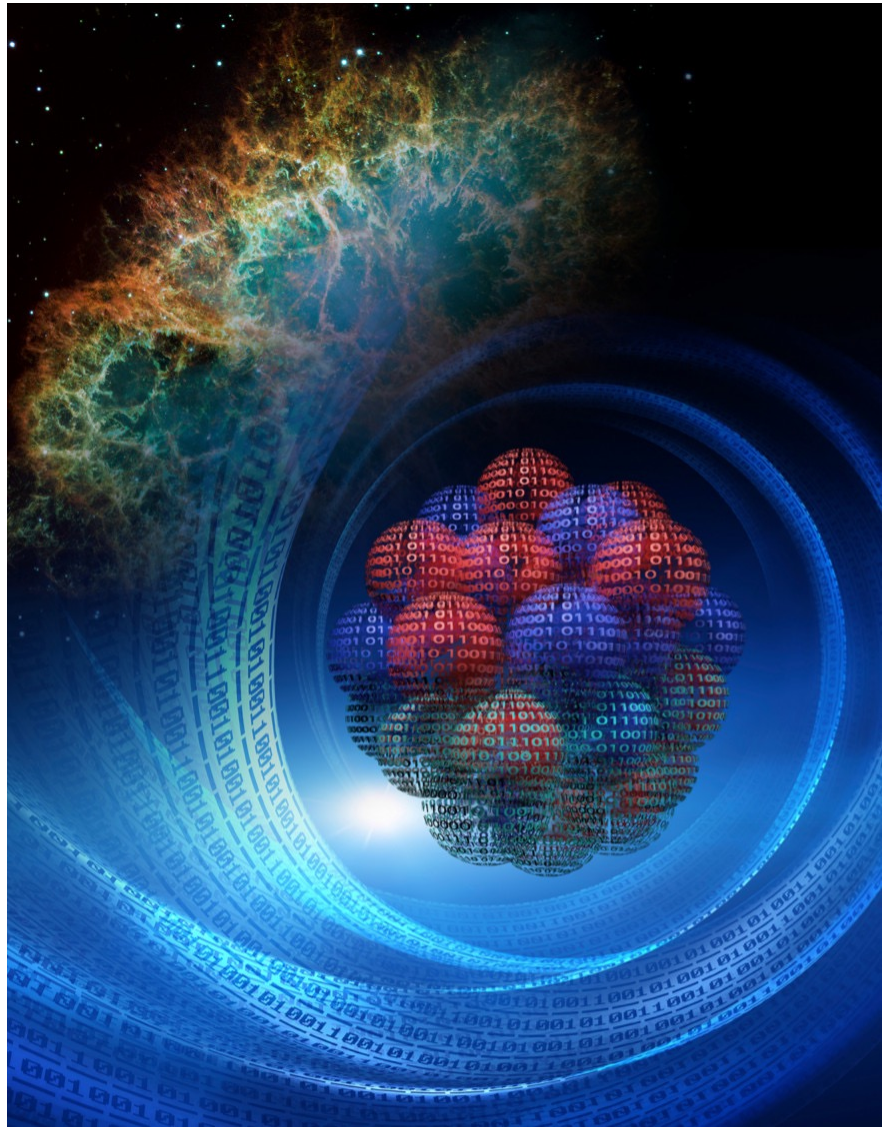

NuCCOR performance update

The main computational tasks are to perform distributed, sparse tensor contractions, followed by a global data reduction on a small (<1GB) dataset. Due to the fundamental symmetries in atomic nuclei, the sparse tensors display specific patterns that we exploit to express a distributed sparse contraction as a series of local, dense contractions. Thus, our performance is completely dependent on the dense tensor contraction kernels displayed below and MPI collective operations.

- Plot shows performance gains of tensor contraction tests using a single GCD of the MI250 on Crusher vs. the V100 on Summit.
  - No bar means that the V100 was not capable of performing the run for the given data size.
- Good performance gain across the board, especially for double precision kernels (marked by R64 and C128).
- From left to right, the kernel tests results in more complicated, non-contiguous memory access pattern.

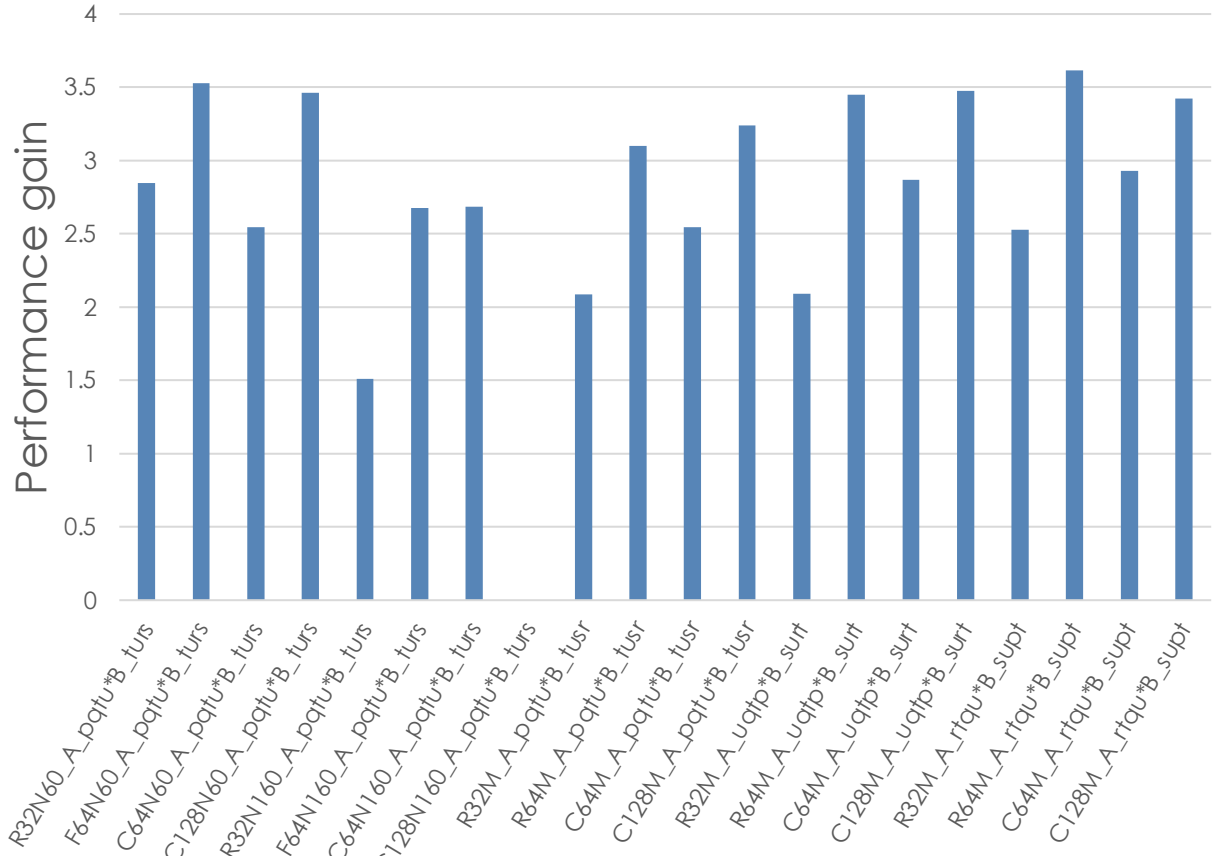

Issue motifs at the moment

- Register pressure and spills
- Unsafe FP atomics (effective workaround exists and has been documented in Crusher guide)
- Lack of information obtainable with performance tools
  - Crusher patched this week with new counter info

Questions?