Early Application Readiness Results on Summit

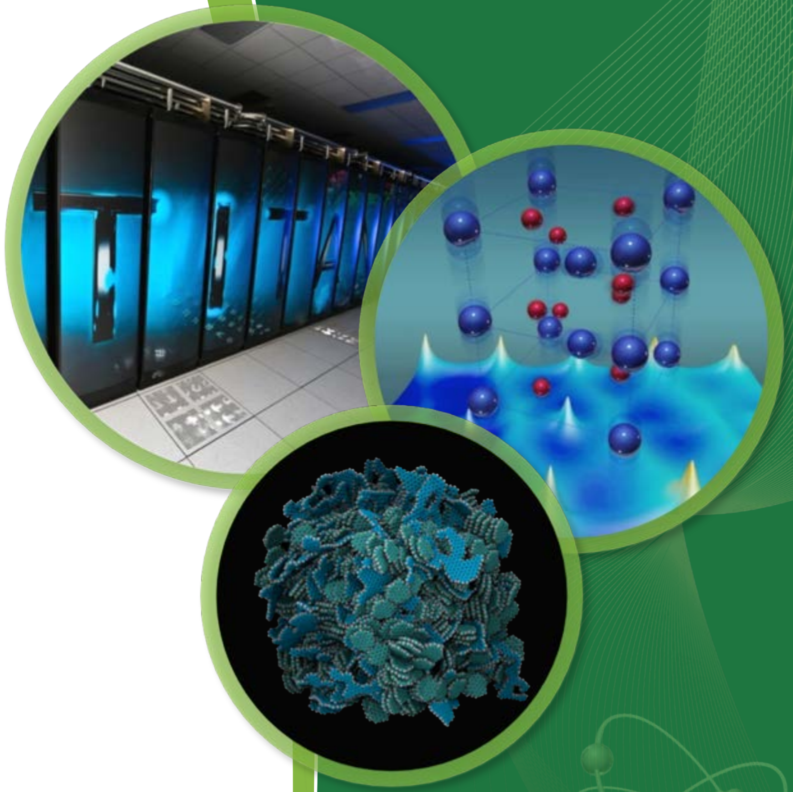

Tjerk Straatsma
Oak Ridge National Laboratory

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

Summit will replace Titan as OLCF's Leadership Supercomputer

Coming soon:

Summit, slated to be more powerful than any other existing supercomputer, is the Department of Energy's Oak Ridge National Laboratory's newest supercomputer for open science.

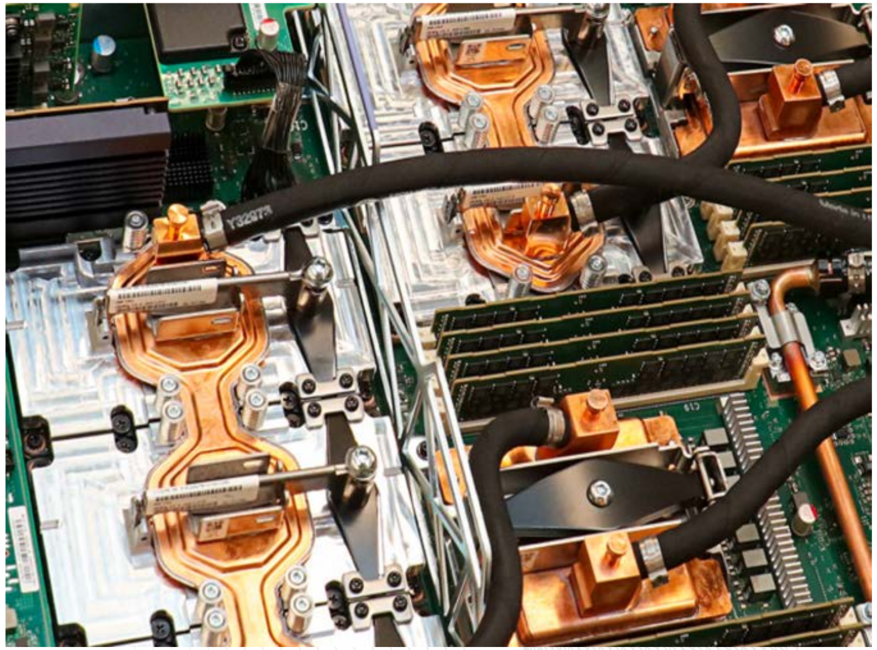

Summit Overview

Components

IBM POWER9
- 22 Cores
- 4 Threads/core
- NVLink

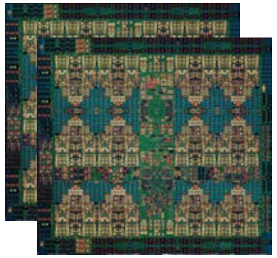

NVIDIA GV100
- 7 TF
- 16 GB @ 0.9 TB/s
- NVLink

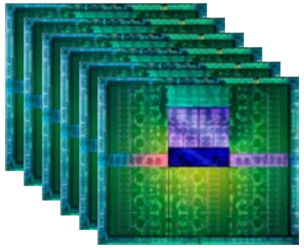

Compute Node

2 x POWER9
6 x NVIDIA GV100
NVMe-compatible PCIe 1600 GB SSD

25 GB/s EDR IB- (2 ports)
512 GB DRAM- (DDR4)
96 GB HBM- (3D Stacked)
Coherent Shared Memory

Compute Rack

18 Compute Servers
Warm water (70°F direct-cooled components)
RDHX for air-cooled components

39.7 TB Memory/rack
55 KW max power/rack

Compute System

10.2 PB Total Memory
256 compute racks
4,608 compute nodes
Mellanox EDR IB fabric
200 PFLOPS
~13 MW

GPFS File System

250 PB storage
2.5 TB/s read, 2.5 TB/s write

OpenPOWER™

Oak Ridge National Laboratory | Leadership Computing Facility

Summit Node Architecture

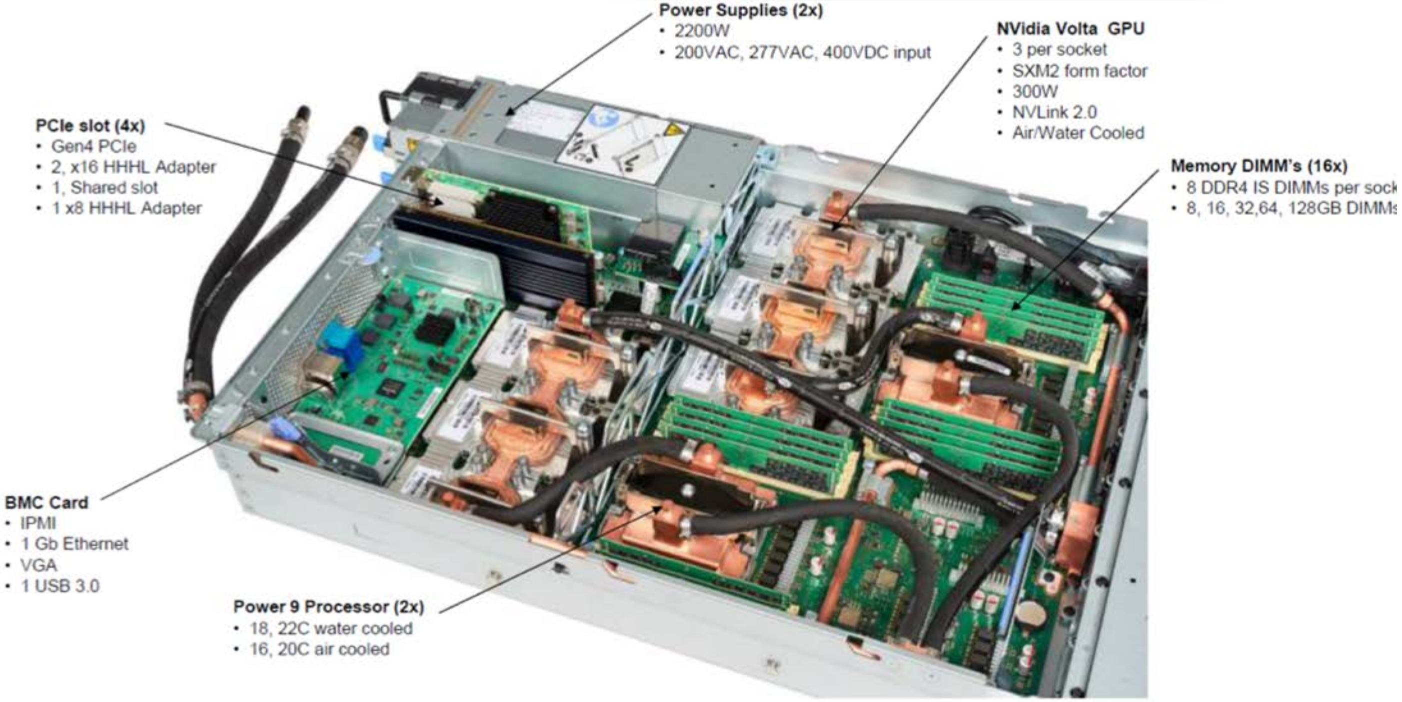

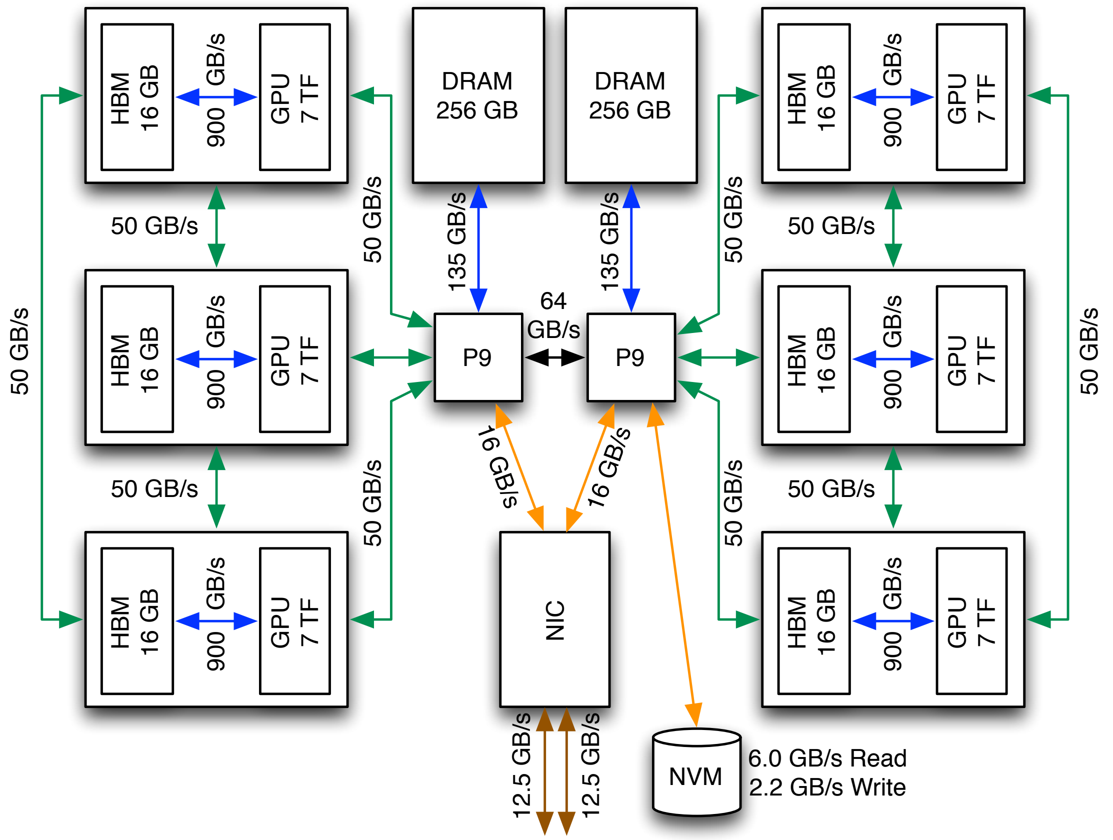

TF	42 TF (6x7 TF)
HBM	96 GB (6x16 GB)
DRAM	512 GB (2x16x16 GB)
NET	25 GB/s (2x12.5 GB/s)
MMsg/s	83

- HBM/DRAM Bus (aggregate B/W)
- NVLINK
- X-Bus (SMP)
- PCIe Gen4
- EDR IB

HBM & DRAM speeds are aggregate (Read+Write).
All other speeds (X-Bus, NVLink, PCIe, IB) are bi-directional.

A comparison of Summit and Titan

- Many fewer nodes
- Much more powerful nodes
- Much more memory per node and total system memory
- Faster interconnect
- Much higher bandwidth between CPUs and GPUs
- Much larger and faster file system

Feature	Titan	Summit
Application Performance	Baseline	5-10x Titan
Number of Nodes	18,688	4,608
Node performance	1.4 TF	42 TF
Memory per Node	32 GB DDR3 + 6 GB GDDR5	512 GB DDR4 + 96 GB HBM2
NV memory per Node	0	1600 GB
Total System Memory	710 TB	>10 PB DDR4 + HBM2 + Non-volatile
System Interconnect	Gemini (6.4 GB/s)	Dual Rail EDR-IB (25 GB/s)
Interconnect Topology	3D Torus	Non-blocking Fat Tree
Bi-Section Bandwidth	15.6 TB/s	115.2 TB/s
Processors	1 AMD Opteron™ 1 NVIDIA Kepler™	2 IBM POWER9™ 6 NVIDIA Volta™
File System	32 PB, 1 TB/s, Lustre®	250 PB, 2.5 TB/s, GPFS™
Power Consumption	9 MW	13 MW

OLCF Program to Ready Application Developers and Users

- We are preparing users through:
  - Application Readiness and Early Science through Center for Accelerated Application Readiness (CAAR)
  - Training and web-based documentation
  - Early access on SummitDev and Summit Phase I system (already accepted)
  - Access for broader user base on final, accepted Phase II system
- Goals:
  - Early science achievements,
  - Demonstrate application readiness,
  - Prepare INCITE & ALCC proposals,
  - Harden Summit for full-user operations

Early Test & Development System SummitDev

Each IBM S822LC node has:

- 2x IBM POWER8+ CPUs
  - 32x 8GB DDR4 memory (256 GB)
  - 10 cores per POWER8+, each core with 8 HW threads
- 4x NVIDIA Tesla P100 GPUs
  - NVLink 1.0 connects GPUs at 80 GB/s
  - 16 GB HBM2 memory per GPU
- 2x Mellanox EDR InfiniBand
- 1600 GB NVMe storage

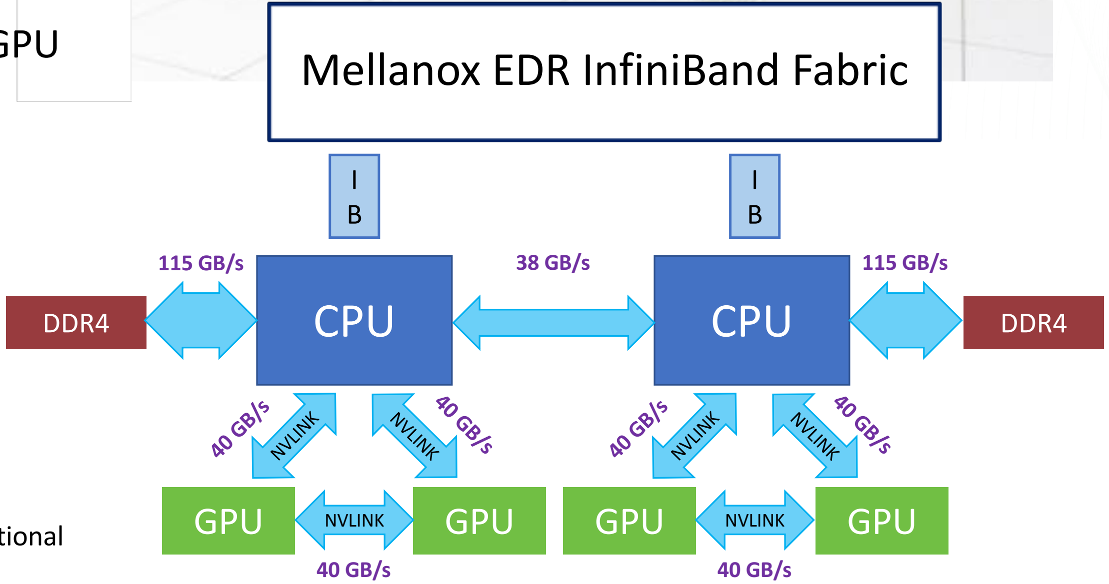

DDR4 BW is Read+Write
X-Bus and NVLINK are bi-directional

Summit Early Access System:

- 3 racks for development, each with 18 nodes
- One rack of login and support servers
- Nodes connected in a full fat-tree via EDR InfiniBand
- Liquid cooled w/ heat exchanger rack
- One additional 18-node rack is for system software testing
- Primary purpose is for code teams to prepare applications for Summit

Center for Accelerated Application Readiness (CAAR) Projects 2015 - 2018

ACME: Climate Research: Advancing Earth System Models
PI: Dr. David Bader, Lawrence Livermore National Laboratory
Science Domain: Climate Science
CAAR Liaison: Dr. Matt Norman
CSEEN Postdoc: Dr. Anikesh Pal
NESAP

DIRAC: CAAR Oak Ridge Proposal for getting the Relativistic Quantum Chemistry Program Package DIRAC ready for SUMMIT
PI: Prof. Dr. Lucas Visscher, Free University Amsterdam, the Netherlands
Science Domain: Relativistic Quantum Chemistry
CAAR Liaisons: Dr. Dmitry Liakh, Dr. Tjerk Straatsma
CSEEN Postdoc: TBD (backfill Dr. Amelia Fitzsimmons)

FLASH: Using FLASH for Astrophysics Simulations at an Unprecedented Scale
PI: Dr. Bronson Messer, Oak Ridge National Laboratory
Science Domain: Astrophysics
CAAR Liaisons: Dr. Bronson Messer
CSEEN Postdoc: Dr. Austin Harris

GTC: Particle Turbulence Simulations for Sustainable Fusion Reactions in ITER
PI: Prof. Dr. Zhihong Lin, University of California - Irvine
Science Domain: Plasma Physics
CAAR Liaison: Dr. Wayne Joubert
NESAP

HACC: Cosmological Simulations for Large-scale Sky Surveys
PI: Dr. Salman Habib, Argonne National Laboratory
Science Domain: Cosmology
CAAR Liaison: Dr. Bronson Messer
NESAP, ESP

LS-DALTON: Large-scale Coupled-cluster Calculations of Supramolecular Wires
PI: Prof. Dr. Poul Jørgensen, Aarhus University
Science Domain: Quantum Chemistry
CAAR Liaison: Dr. Dmytro Bykov
INCITE

NAMD: Molecular Machinery of the Brain
PI: Dr. James Phillips, University of Illinois at Urbana-Champaign
Science Domain: Biophysics
CAAR Liaison: Dr. Tjerk Straatsma
NESAP

NUCCOR: Nuclear Structure and Nuclear Reactions
PI: Dr. Gaute Hagen, Oak Ridge National Laboratory
Science Domain: Nuclear Physics
CAAR Liaison: Dr. Gustav Jansen
CSEEN Postdoc: TBD (backfill Dr. Micah Schuster)

NWCHEM: Developing Coupled Cluster Methodologies for GPUs
PI: Dr. Karol Kowalski, Pacific Northwest National Laboratory
Science Domain: Computational Chemistry
CAAR Liaison: Dr. Dmitry Liakh
IBM Postdoc: Dr. David Appelhans
NESAP

QMCPACK: Materials Science Research for High-Temperature Superconductors
PI: Dr. Paul R. C. Kent, Oak Ridge National Laboratory
Science Domain: Materials Science
CAAR Liaison: Dr. Ying Wai Li
CSEEN Postdoc: Dr. Andreas Tillack

RAPTOR: Fluid Dynamics Research to Accelerate Combustion Science
PI: Dr. Joseph Oefelein, Sandia National Laboratory, Livermore
Science Domain: Engineering/Combustion
CAAR Liaison: Dr. Ramanan Sankaran
CSEEN Postdoc: TBD (backfill Dr. Kalyana Gottiparthi)

SPECFEM: Mapping the Earth's Interior Using Big Data
PI: Dr. Jeroen Tromp, Princeton University, Princeton University
Science Domain: Seismology
CAAR Liaison: Dr. Judy Hill
CSEEN Postdoc: Dr. Yangkang Chen

XGC: Multiphysics Magnetic Fusion Reactor Simulator, from Hot Core to Cold Wall
PI: Dr. CS Chang, Princeton Plasma Physics Laboratory, Princeton University
Science Domain: Plasma Physics
CAAR Liaison: Dr. Ed D'Azevedo
CSEEN Postdoc: TBD (backfill Dr. Stephen Abbott)
NESAP

FLASH

Dr. Bronson Messer, Oak Ridge National Laboratory, Oak Ridge, TN
Dr. Austin Harris, ORNL
Dr. Alan Calder, Stony Brook University, NY
Dr. Sean Couch, California Institute of Technology, CA
Dr. Petros Tzeferacos, University of Chicago, IL


Bronson Messer


Austin Harris

Domain Area: Astrophysics

Research supported by the DOE SC/NP seeks to understand the nuclear processes that have shaped the cosmos, including the origin of the elements, the evolution of stars, and the detonation of supernovae.

FLASH is a publicly available, component-based, massively parallel, adaptive mesh refinement (AMR) code that has been used on a variety of parallel platforms

In particular, measurements made at the under-construction Facility for Rare Isotope Beams (FRIB) – coupled with simulations of the late-time evolution of supernovae – will help determine how the elements from iron to uranium are created.

Targeted for CAAR

1. Hydrodynamics module performance
2. Equation of State (EOS) threading and vectorization
3. Nuclear kinetics (burn unit) threading and vectorization, including Jacobian formation and solution using GPU-enabled libraries

http://flash.uchicago.edu

FLASH Early Summit Results

- FLASH: Component-based, massively parallel, adaptive-mesh refinement code
  - Widely used in astrophysics community (>1100 publications from >600 scientists)
- CAAR work primarily concerned with increasing physical fidelity by accelerating the nuclear burning module and associated load balancing.
- Summit GPU performance **fundamentally changes the potential science impact** by enabling large-network (i.e. 160 or more nuclear species) simulations.
  - Heaviest elements in the Universe are made in neutron-rich environments – small networks are incapable of tracking these neutron-rich nuclei
  - Opens up the possibility of producing precision nucleosynthesis predictions to compare to observations
  - Provides detailed information regarding most astrophysically important nuclear reactions to be measured at FRIB

Preliminary results on Summit

GPU+CPU vs. CPU-only performance on Summit for 288-species network : **2.9x**

P9: 24.65 seconds/step
P9 + Volta: 8.5 seconds/step

288-species impossible to run on Titan

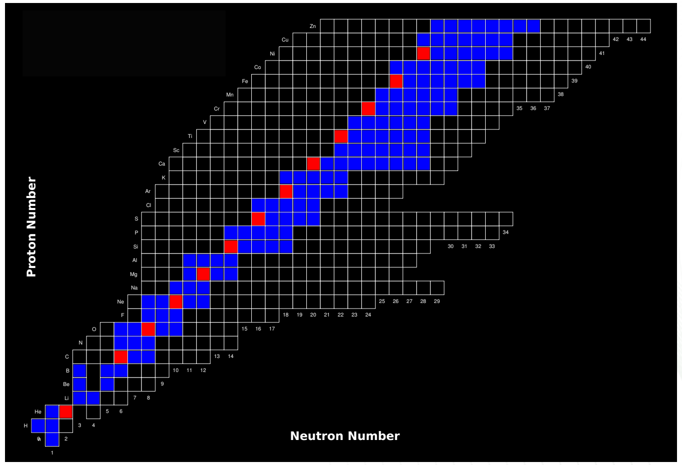

NASA, ESA, J. Hester and A. Loll (Arizona St. Univ.)

Time for 160-species (blue) run on Summit ***roughly equal*** *to 13-species "alpha" (red) network run on Titan*

>100x the computation for identical cost

GTC

Prof. Dr. ZhiHong Lin, University of California, Irvine
Dr. Wayne Joubert, ORNL
Prof. Dr. William Tang, Princeton University
Dr. Ihor Holod, University of California, Irvine
Dr. Animesh Kuley, University of California, Irvine
De. Bei Wang , Princeton University

ZhiHong Lin

Wayne Joubert

Domain Area: Plasma Physics / Fusion

Plasma simulations supporting the ITER project are a key DOE/FES focus and are required to understand the complex kinetic dynamics governing magnetic confinement properties of fusion-grade plasmas.

The Gyrokinetic Toroidal Code (GTC) is a massively parallel particle-in-cell code for first-principles, integrated simulations of burning plasma experiments such as the International Thermonuclear Experimental Reactor (ITER). GTC solves the five-dimensional (5D) gyrokinetic equation in full, global torus geometry to address kinetic turbulence issues in magnetically-confined fusion tokamaks.

The proposed CAAR project will develop the capability for direct numerical simulations of key instabilities that limit the burning plasma performance and threaten device integrity in magnetically-confined fusion systems.

Targeted for CAAR:
The GTC particle-in-cell (PIC) algorithm is the most computationally dominant component of the GTC code. As a sequel to previous work, a large part of the project's performance optimization work will thus focus on efficient multithreading of this computation for Summit.

http://phoenix.ps.uci.edu/gtc_group/

GTC: Summit Results

- GTC is a gyrokinetic toroidal fusion code for modeling fusion reactors
- Much of the code already uses GPUs effectively on Titan, e.g., particle push, using OpenACC
- Extensive development work has been done to optimize for Summit
- Code now uses the NVIDIA AmgX solver to improve performance over the previous PETSc field solver

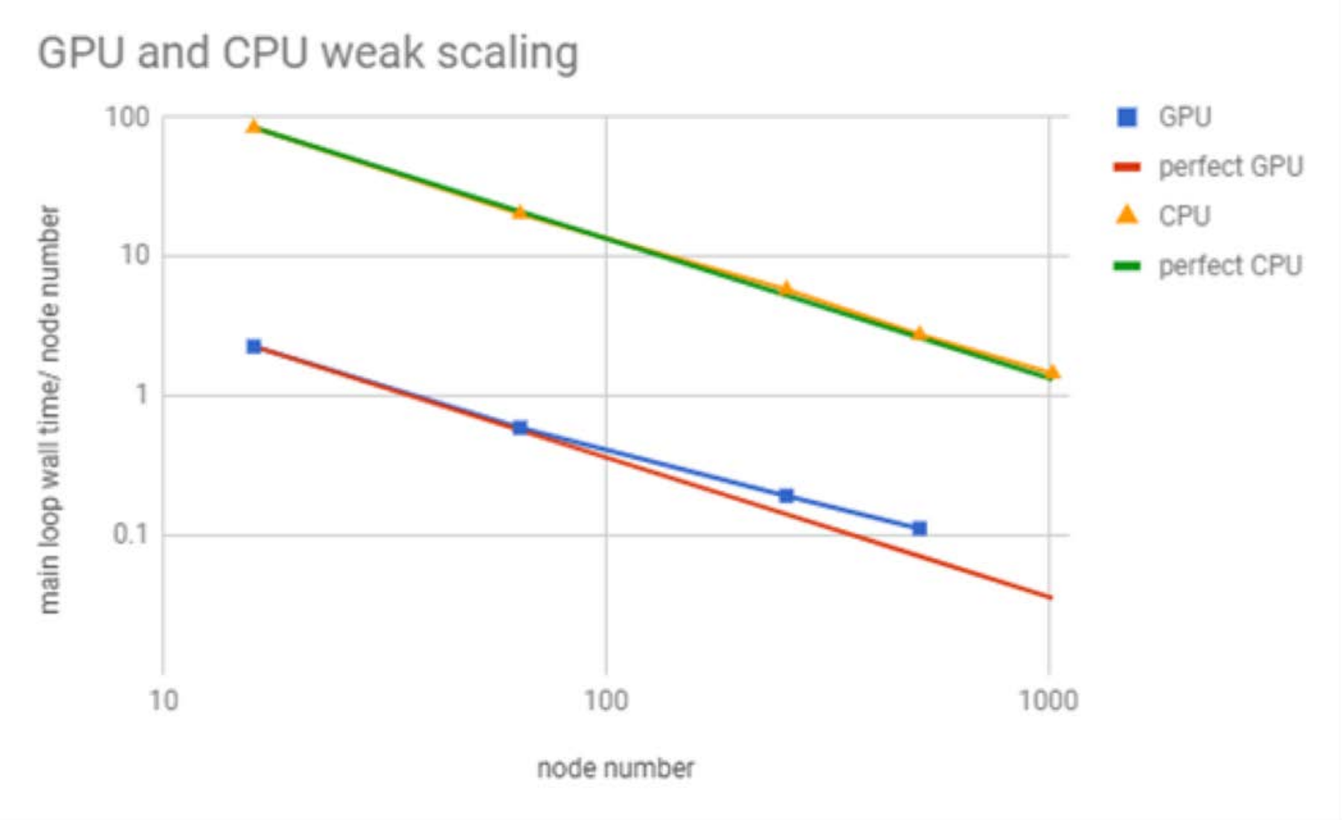

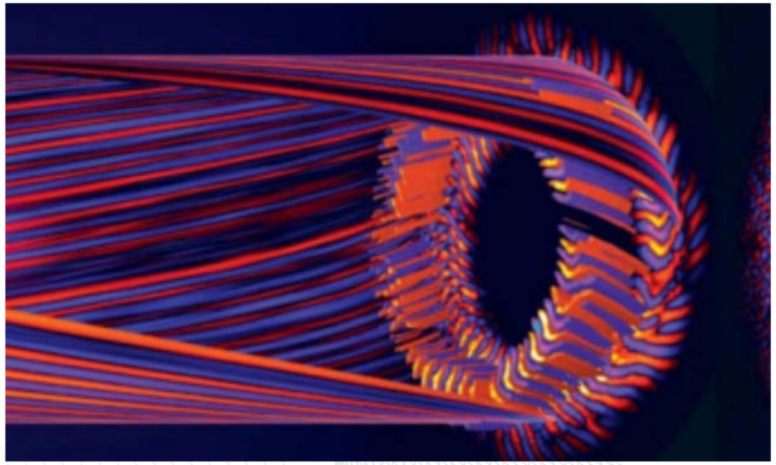

HACC

Dr. Salman Habib, Argonne National Laboratory
Dr. Hal Finkel, Mr. Nicholas Frontiere, Dr. Katrin Heitmann, Dr. Vitali Morozov,
Dr. Adrian Pope, Argonne National Laboratory
Mr. Esteban Rangel, Northwestern University
Dr. Jim Aherns, Mr. Li-Ta Lo, Dr. Chris Sewell, Los Alamos National Laboratory

Salman Habib

Bronson Messer

Domain Area: Astrophysics

Research motivated by the Cosmic Frontier Science Driver supported by the DOE Office of Science Program in High Energy Physics seeks to reveal the nature of dark matter and dark energy. The ancient light from distant galaxies allows the distribution of dark matter to be mapped and perhaps the nature of dark energy to be unraveled. Large-scale cosmological simulation is an essential tool in these investigations.

HACC is a particle-mesh cosmology code. It has been used to produce numerical surveys to guide the design of observational galaxy surveys and to develop prediction tools to help survey teams interpret large amounts of observational data.

Targeted for CAAR

1. The addition of baryonic hydrodynamics via Conservative Reproducing Kernel Smooth Particle Hydrodynamics (CRKSPH)
2. Optimization of both the Tree-PM and P^3M short-range solver on Summit
3. A new, GPU-enabled visualization and analysis framework

https://www.alcf.anl.gov/projects/next-generation-cosmology-simulations-hacc-challenges-baryons

ECP ExaSky Project: First HACC Runs on Summit under CAAR Application Readiness

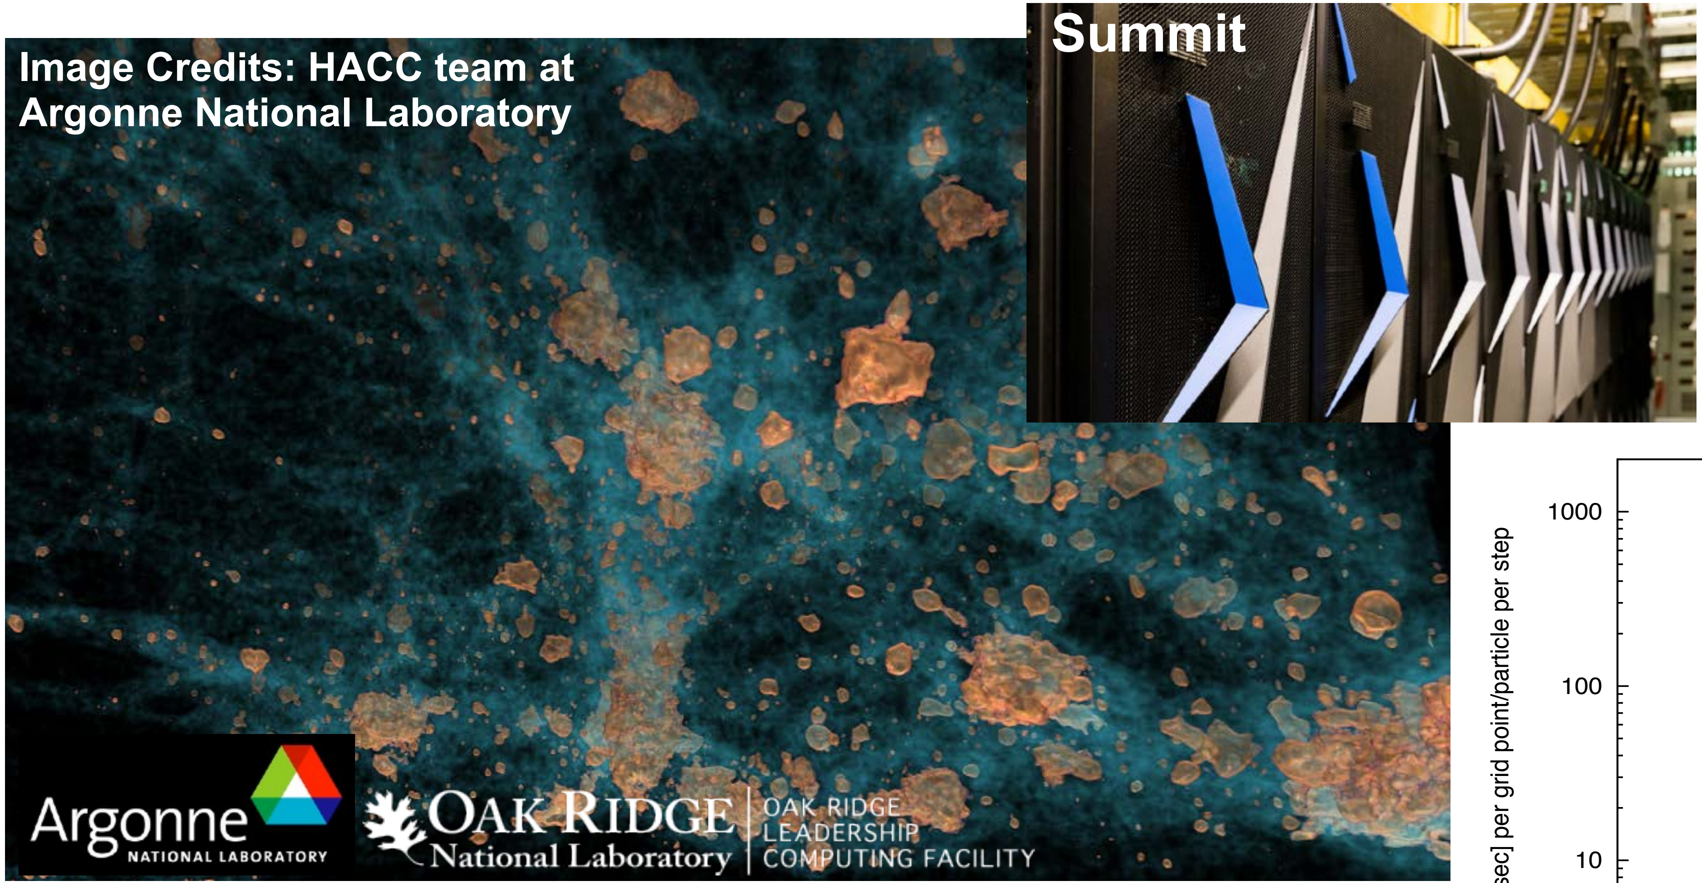

The HACC (Hardware/Hybrid Accelerated Cosmology Code) is being developed under ASCR's Exascale Computing Project to carry out simulations for next-generation cosmological surveys sponsored by DOE HEP

Matter distribution in the universe from a HACC test run on Summit at ORNL; this work is part of an OLCF Center for Accelerated Application Readiness (CAAR) supported project

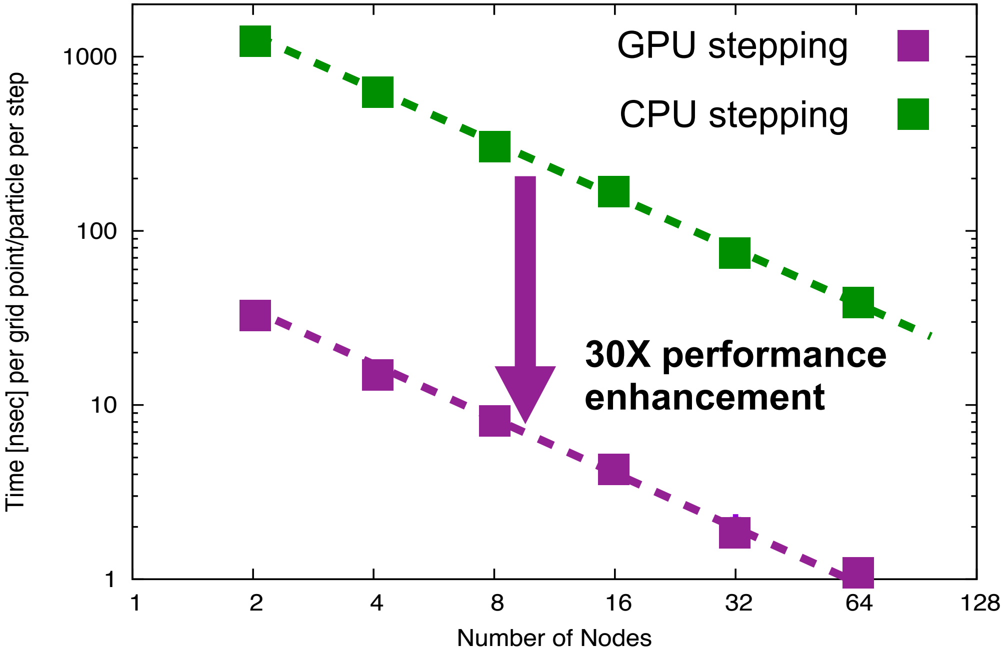

Demonstration of GPU acceleration on Summit with ExaSky's extreme-scale HACC simulation framework

Argonne NATIONAL LABORATORY · Los Alamos NATIONAL LABORATORY EST.1943 · BERKELEY LAB · ECP EXASCALE COMPUTING PROJECT

QMCPACK


Paul Kent


Ying Wai Li


Andreas Tillack

Dr. Paul R. C. Kent, ORNL
Dr. Ying Wai Li, Dr. Andreas Tillack, Dr. F. A. Reboredo and Dr. J. T. Krogel, ORNL

Domain Area: Materials Science

QMCPACK is an open source, ab initio quantum Monte Carlo code for the study and prediction of materials properties by stochastically solving quantum many-body problems using variational MC and more accurate but computationally demanding diffusion MC. It allows for highly accurate, predictive ab initio calculations of larger and more complex systems that are limited or inaccessible by conventional methods such as density functional theory (DFT).

Algorithmic improvements and enhanced predictive capacity targeted by CAAR will accelerate the understanding of a wide range of materials, e.g., strongly correlated systems and transition elements, which are essential for materials design for energy capture, storage and conversion, as well as high temperature superconductors.

Targeted for CAAR:
1. Orbital evaluation and memory usage
2. Slater determinant evaluation: improve numerical stability of current (less computationally intensive) updating scheme
3. QMC energy evaluation: optimize and explore extra level(s) of parallelism to improve time-to-solution ratio
4. Load balancing: reduce synchronizations and global collection operations to maximize performance

http://qmcpack.org/

Oak Ridge National Laboratory | Leadership Computing Facility

QMCPACK Early Summit Results

- QMCPACK: Accurate quantum mechanics-based simulation of materials, including high-temperature superconductors.
- QMCPACK runs correctly and with good initial performance on up to 64 Summit nodes.
- Runs use the latest version from GitHub without modification.
- Improvements to both Power CPU and Volta GPU performance are expected via optimization and new algorithms.
- A Summit node is 50-times faster than a Titan node for this problem, indicating a ~3.7x increase in the complexity of materials (electron count) computable in the same wall time as Titan.

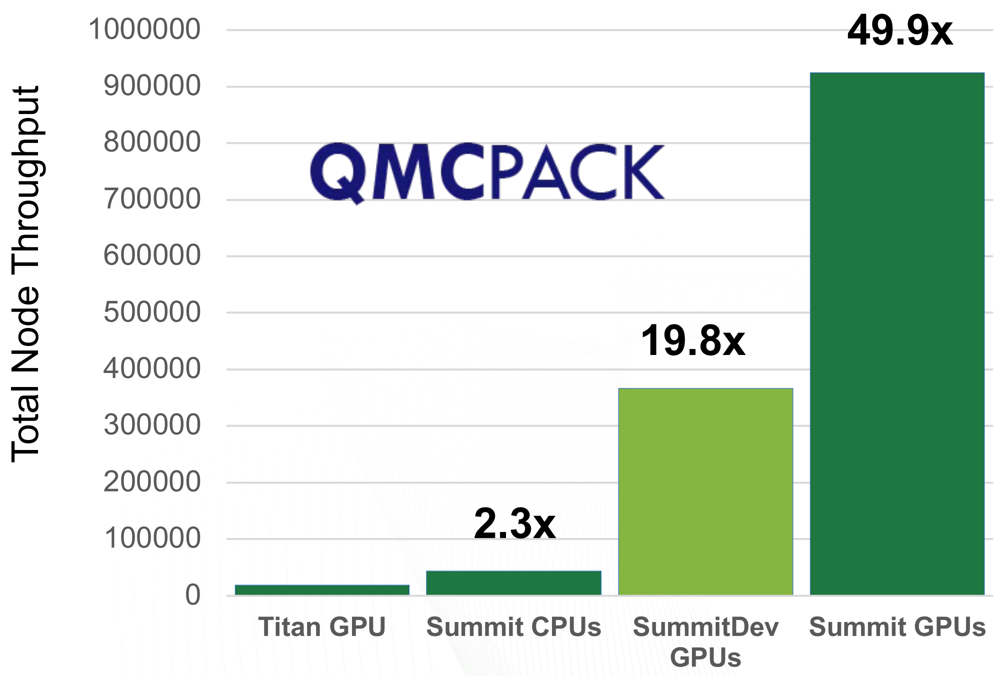

Initial data by Andreas Tillack/OLCF: QMCPACK v3.4.0 NiO 128-atom cell. Power CPU reference uses optimal 2 MPI tasks, 42 OpenMP threads each and optimized "SoA" version.

XGC

Dr. CS Chang, Princeton Plasma Physics Laboratory, Princeton, NJ
Dr. Pat Worley, Dr. Ed D'Azevedo and Dr. Scott Klasky, Oak Ridge National Laboratory
Dr. Seung-Hoe Ku, Dr. Jianying Land, and Dr Stephane Ethier, PPPL
Dr. Mark Adams, LBNL

CS Chang

Ed D'Azevedo

Domain Area: Plasma Physics

Plasma Physics is a DOE/FES research program focused on advance research on magnetically confined plasmas to develop the science needed to create a sustainable fusion energy source.

XGC is a widely used advanced whole-device tokamak gyrokinetic code for modeling edge plasma at first-principles level. XGC is a 5D gyrokinetic particle-in-cell (PIC) with finite element that can model tokamak plasma in realistic divertor magnetic geometry with neutral particle recycling at material wall surface.

Whole device integrated modeling for ITER to obtain understanding of the edge pedestal boundary physics and disruption events are Tier 1 priority recommendation from DOE Fusion Energy Science Advisory Committee. Acceleration via the CAAR effort will greatly improve XGC's ability to model these critical areas of interest.

Targeted for CAAR:
- Acceleration of computational kernels (particle push, collision physics, charge deposition) for GPU.
- Innovative use of NVRAM for   check-point/restart, in-situ data analysis and data reduction.
- Reorganization of data structures to take advantage of deep memory hierarchies.

http://epsi.pppl.gov/computing/xgc-1

XGC Current Status

- Extended weak scaling to 1024 nodes for the GPU+CPU runs for the collisionless turbulence case.
- With collisions, the GPU-accelerated code [(CPU+GPU)/CPU} has a speedup of almost 8 at 1024 nodes.

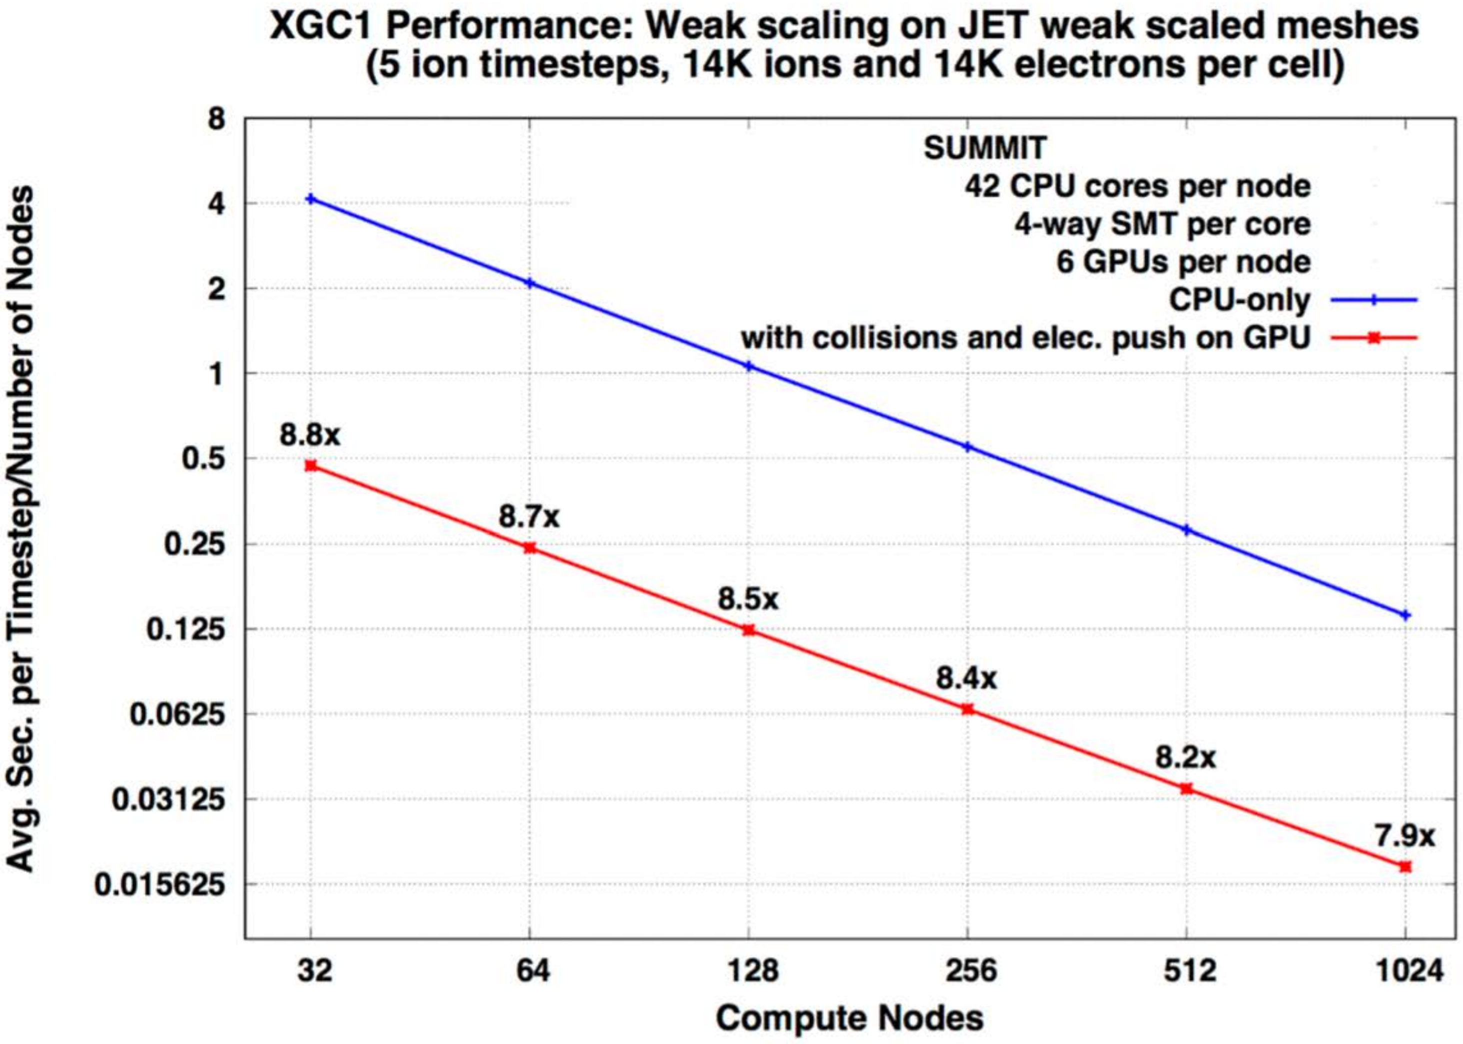

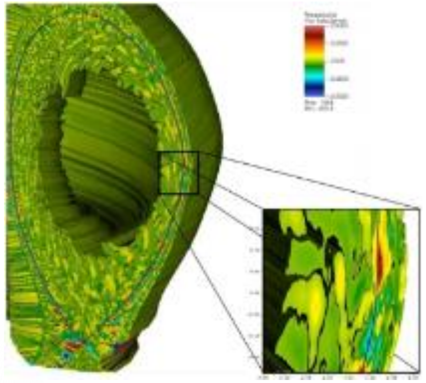

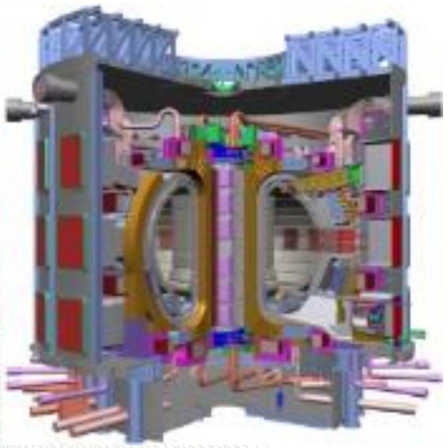

Summit Early Science Program

1. Call for Proposals
   a. To be issued December 2017
   b. Open to CAAR teams and others who ported applications for Summit
   c. Rolling acceptance with a deadline in CY18Q2
2. Criteria
   a. Meeting CAAR scalability and performance targets
   b. Mission relevant scientific campaign
3. Goals
   a. Science achievements and publications early
   b. Demonstration of CAAR and other applications
   c. Transition of applications & projects to INCITE and/or ALCC programs
   d. Harden system for production and availability using important, mission-relevant applications

Summit Early Science Program (ESP)

We put out a Call for Proposals in December 2017

- Resulted in 62 Letters of Intent (LOI) received by year's end.
  - 27 are from PIs at universities
  - 32 are from PIs at national laboratories or research institutions (DOE, NASA)
  - 14 are CAAR project-related LOIs
  - 27 have had past INCITE allocations
  - 9 have had past ALCC allocations
  - 15 have connections to the US DOE Exascale Computing Project
  - 9 are AI or deep learning-related
- Proposals are due June 1, 2018
- ESP projects will gain full access to Summit for Early Science later this year

Center for Accelerated Application Readiness – Success Metrics

- ***Scalability***: Applications should demonstrate reduced time to solution (for strong scaling benchmarks) or time to solution divided by the number of nodes used (for weak scaling benchmarks) to 20% or more of the full Summit machine, N_{20}. This is also known as the capability metric.

- ***Accelerated Performance***: Applications should demonstrate a performance improvement of a factor of two or better by using all six GPUs compared to using both CPUs only, with a job that runs on 20% of the full Summit machine.

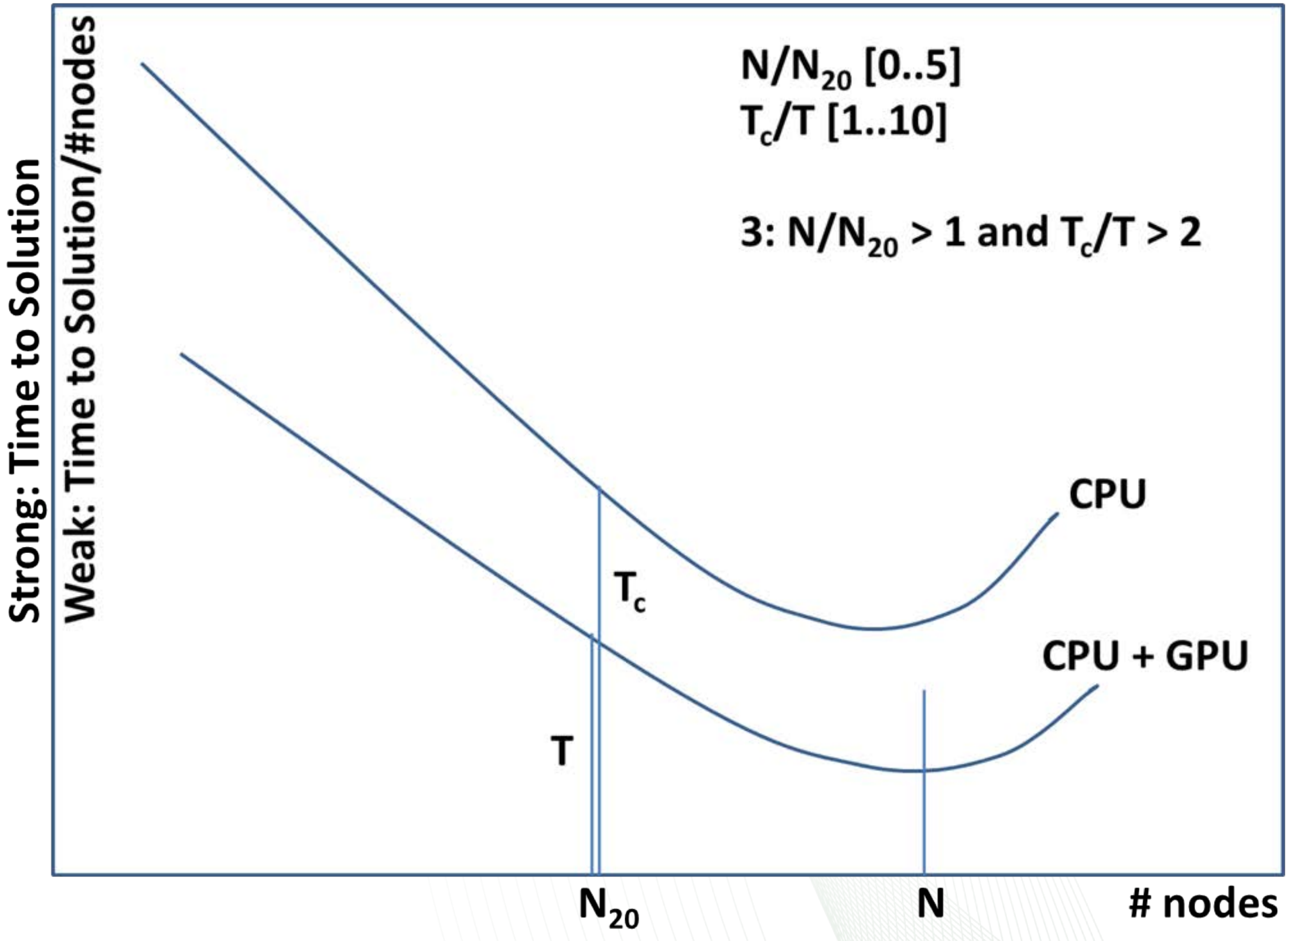

CoMet

Dan Jacobson, Wayne Joubert (ORNL)

Dan Jacobson

Wayne Joubert

Domain area: bioinformatics & biophysics

Summit allows us to:

- Discover co-evolutionary relationships across a population of genomes at an unprecedented scale
- Predict putative protein complexes which will be used in conjunction with protein-protein docking to construct the molecular machinery of the cell

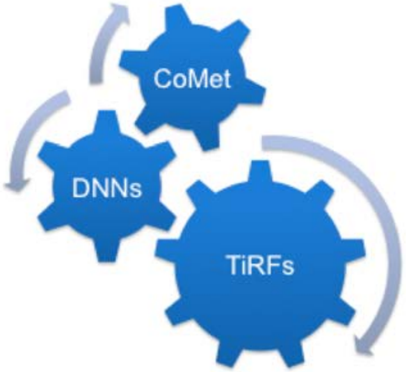

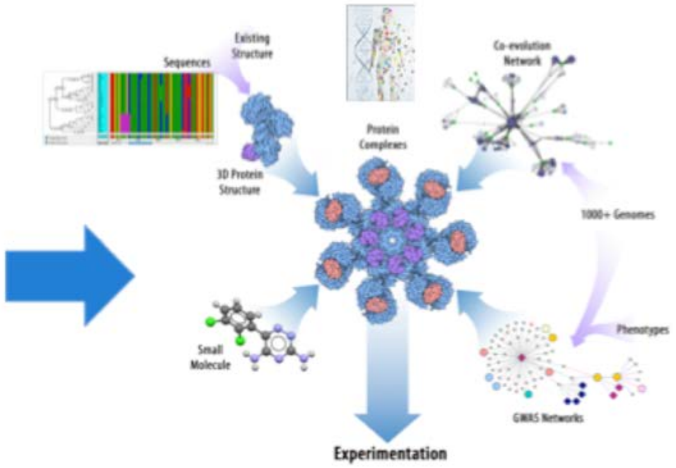

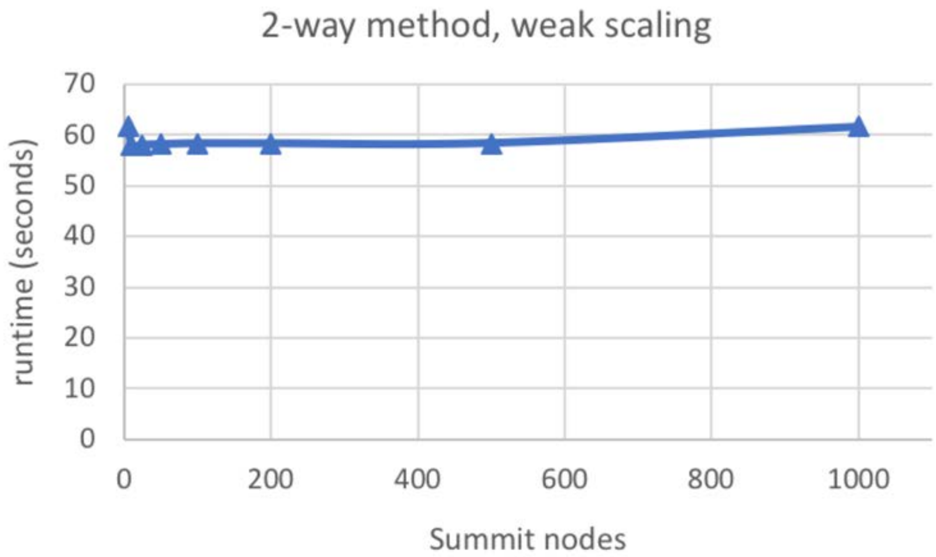

Near-ideal weak scaling up to 1000 Summit nodes

Preliminary results on Summit

- CoMet runs **38.6X** faster on a Summit node than a Titan node
- The Summit GPU version is **119.2X** faster than the CPU-only version
- Code gives **near-ideal** weak scaling performance
- We estimate **92 quadrillion** element comparisons per second possible on full Summit system – **10X** what is possible on Titan
- This is **> 30,000X** more than other, competing codes, enabling new science that was not previously possible

GronOR


Remco Havenith

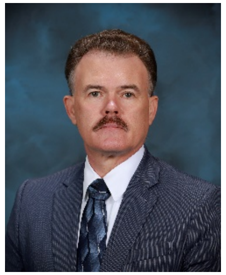
Tjerk Straatsma

Dr. Remco W. A. Havenith, University of Groningen
Dr. Tjerk P. Straatsma, Oak Ridge National Laboratory
Prof. Dr. Ria Broer, Drs. R. K. Kathir, Drs. Luis E. Aguilar Suarez, University of Groning
Drs. Maria A. Izquierdo Morelos, University of Valencia
Drs. Meilani Wibowo, University of Pisa

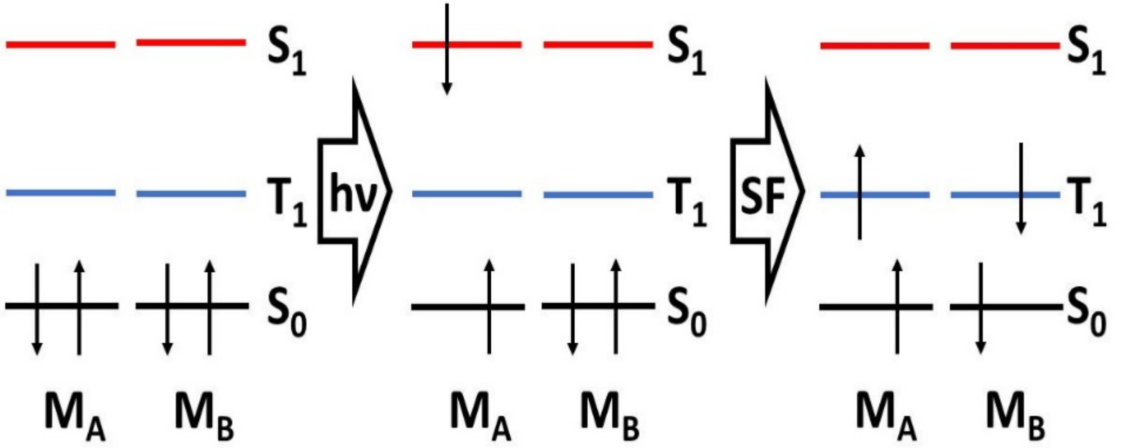

Domain Area: Chemical Physics

Research motivated by the need for more efficient photovoltaic materials for use in solar cell devices to convert solar energy into electricity. In particular materials that are able generate multiple electron-hole pairs upon absorption of a single photon have the potential to achieve the needed enhanced efficiencies. Singlet fission is an example of a process by which multiple excitons are generated from a single high-energy photon without significant thermal energy loss.

This work uses computational techniques to study the factors that govern singlet fission, through the calculation of the electronic coupling of the lowest photo-excited spin singlet state and the double triplet states that couple into a singlet state. The method of choice for the calculation of these couplings is non-orthogonal configuration interaction.

http://www.gronor.org

GronOR Results on Summit

GronOR Performance: Naphthalene Molecules, CAS(4,4)/CAS(8,8)

- Benchmark a naphthalene dimer single element of the full 4x4 Hamiltonian
- This single element requires evaluating 112,867,800 matrix elements!!!
- Estimate for the full 4x4 Hamiltonian is 1.1 billion matrix elements

- ***Near-linear accelerated strong scaling benchmarks on Summit for 64 to 1024 nodes***
- ***GPU accelerated speedup is a factor of 6.8 comparing six MPI+GPU and 28 MPI+CPU-only***

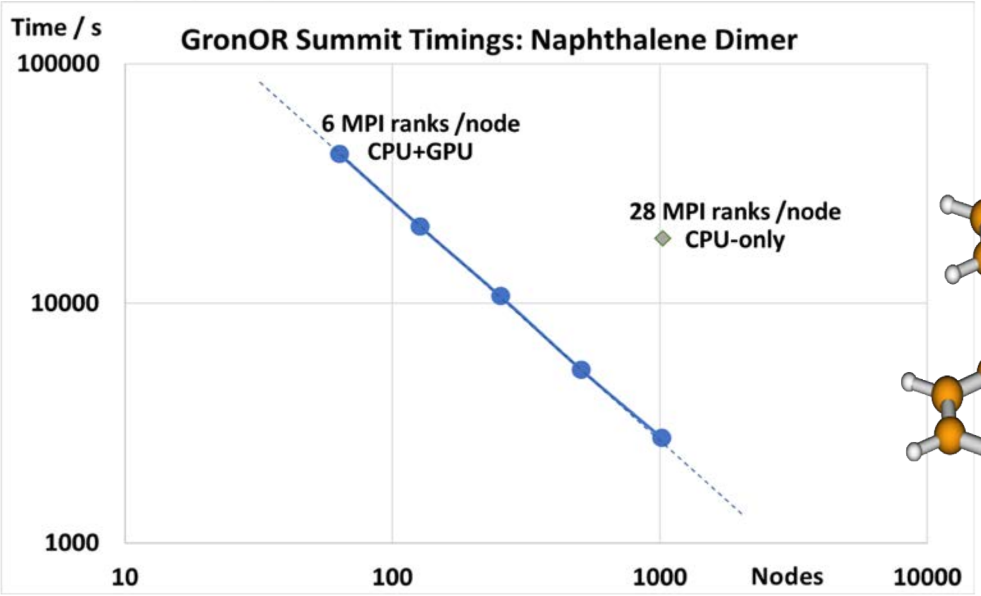

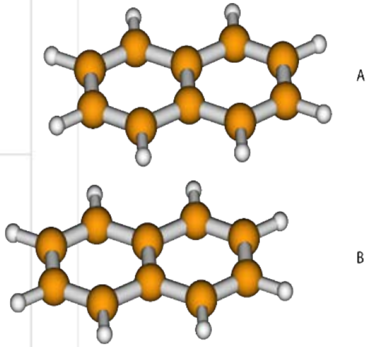

Summit Nodes	GPU-accelerated		Non-accelerated	
	Active Processes	Wall Clock Time / s	Active Processes	Wall Clock Time / s
64	379	41,851.022	-	-
128	763	20,829.594	-	-
256	1,531	10,651.503	-	-
512	3,067	5,258.905	-	-
1,024	6,139	2,734.167	28,672	18,622.110

Application Readiness for Frontier

The Center for Accelerated Application Readiness will continue with Frontier

- Start as soon as the vendor for Frontier is announced
- Application Readiness
  - Call for proposals
  - Partnership between developers, OLCF staff, and vendor Center of Excellence
  - Goals: diversity in science domain
algorithmic motifs
programming environments
- Early Science period
  - Call for proposals
  - Goals demonstrate early scientific results
hardening of hardware and software
develop experience with users and OLCF staff

Acknowledgment

- CAAR Principal Investigators and their teams: David Bader (ACME), Lucas Visscher (DIRAC), Zhihong Lin (GTC), Salman Habib (HACC), Poul Jørgensen (LSDALTON), Jim Phillips (NAMD), Gaute Hagen (NUCCOR), Karol Kowalski (NWCHEM), Paul Kent (QMCPACK), Joe Oefelein (RAPTOR), Jeroen Tromp (SPECFEM), and CS Chang (XGC)
- Director Discretionary Project CHM120 (GronOR): Remco Havenith and Ria Broer
- Early Science LoI and GB teams
- OLCF Staff: Dmytro Bykov, Ed D'Azevedo, Judy Hill, Gustav Jansen, Wayne Joubert, Ying Wai Li, Dmitry Liakh, Bronson Messer, Matt Norman, and Ramanan Sankaran
- CSEEN postdocs: Stephen Abbott, Ashleigh Barnes, Yangkang Chen, Kalyana Gottiparthi, Austin Harris, Amelia Fitzsimmons, Anikesh Pal, Thom Papatheodore, Micah Schuster, and Andreas Tillack
- Center of Excellence Staff: Stephen Abbott, David Appelhans, Jeff Larkin, Eric Luo, Matt Niemerg, Bob Walkup, Peng Wang, and many others

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.

Questions?

Tjerk Straatsma
str@ornl.gov