Perspectives on Application Benchmarking at OLCF

Bronson Messer & Jack Wells
Oak Ridge Leadership Computing Facility
National Center for Computational Sciences
Oak Ridge National Laboratory

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

Oak Ridge Leadership Computing Facility (OLCF)

Mission: Deploy and operate the computational resources required to tackle global challenges

- Providing world-leading computational and data resources and specialized services for the most computationally intensive problems
- Providing stable hardware/software path of increasing scale to maximize productive applications development
- Operate would-class user programs accessible to industry, universities, national laboratories and federal agencies, and where access is based on the merit of the proposed work.
- With our partners, deliver transforming discoveries and engineering breakthroughs across a breadth of applications.

OLCF is one of the DOE Office of Science Advanced Scientific Computing Research Facilities

Providing the Facilities – High-End and Leadership Computing

- **National Energy Research Scientific Computing Center (NERSC)** Lawrence Berkeley National Laboratory
  - Delivers high-end capacity computing to entire DOE SC research community
  - Over 7,000 users and 700 projects
- **Leadership Computing Centers at Argonne National Laboratory (ALCF) and Oak Ridge National Laboratory (OLCF)**
  - Delivers highest computational capability
    - Open to national and international researchers, including industry
    - Not constrained by existing DOE or Office of Science funding or topic areas
  - Approximately 1,000 users and 80% of resources consumed by 50-60 large projects at each center, each year

Linking it all together – Energy Sciences Network (ESnet)

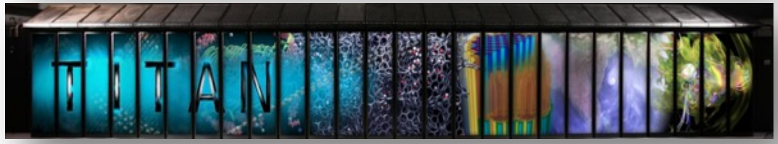

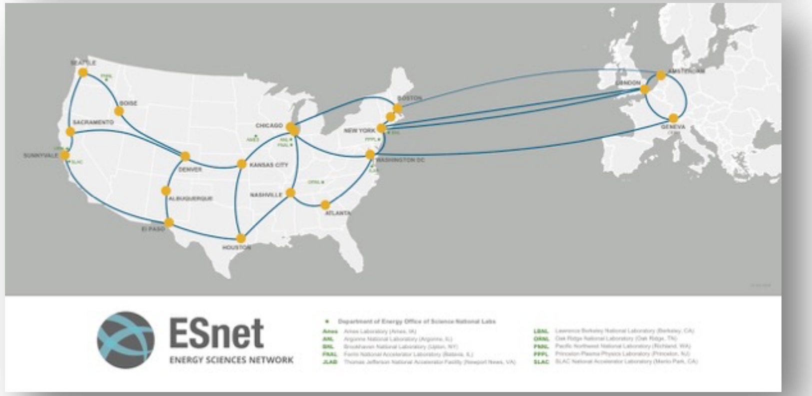

OLCF Use Cases for Benchmarking

- Program Development and Marketing
- Application Development and Performance
- Compiler Development and Performance
- Procurements
- User Program Management

Program Development and Marketing

We build supercomputers for science!

Science Accomplishments Highlights

All from 2014 INCITE Program on Titan

Cosmology

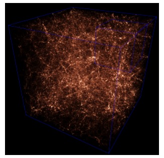

Salman Habib
Argonne National Laboratory

Habib and collaborators used its HACC Code on Titan's CPU–GPU system to conduct today's largest cosmological structure simulation at resolutions needed for modern-day galactic surveys.

K. Heitmann, 2014. *arXiv.org,* 1411.3396

Combustion

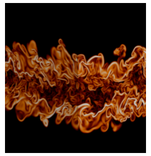

Jacqueline Chen
Sandia National Laboratory

Chen and collaborators for the first time performed direct numerical simulation of a jet flame burning dimethyl ether (DME) at new turbulence scales over space and time.

A. Bhagatwala, et al. 2014. *Proc. Combust. Inst.* **35**.

Superconducting Materials

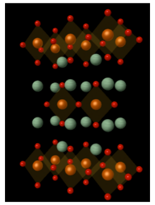

Paul Kent
ORNL

Paul Kent and collaborators performed the first ab initio simulation of a cuprate. They were also the first team to validate quantum Monte Carlo simulations for high-temperature superconductor simulations.

K. Foyevtsova, et al. 2014. *Phys. Rev. X* **4**

Molecular Science

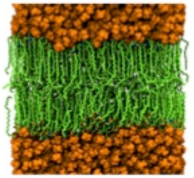

Michael Klein
Temple University

Researchers at Procter & Gamble (P&G) and Temple University delivered a comprehensive picture in full atomistic detail of the molecular properties that drive skin barrier disruption.

M. Paloncyova, et al. 2014. *Langmuir* **30**

C. M. MacDermaid, et al. 2014. *J. Chem. Phys.* **141**

Fusion

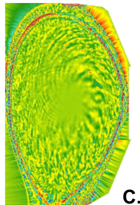

C.S. Chang
PPPL

Chang and collaborators used the XGC1 code on Titan to obtain fundamental understanding of the divertor heat-load width physics and its dependence on the plasma current in present-day tokamak devices.

C. S. Chang, et al. 2014. *Proceedings of the 25th Fusion Energy Conference, IAEA*, October 13–18, 2014.

2 OLCF OAR 3/2015 Strategic Results

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

Top500.org has been a success in marketing HPC

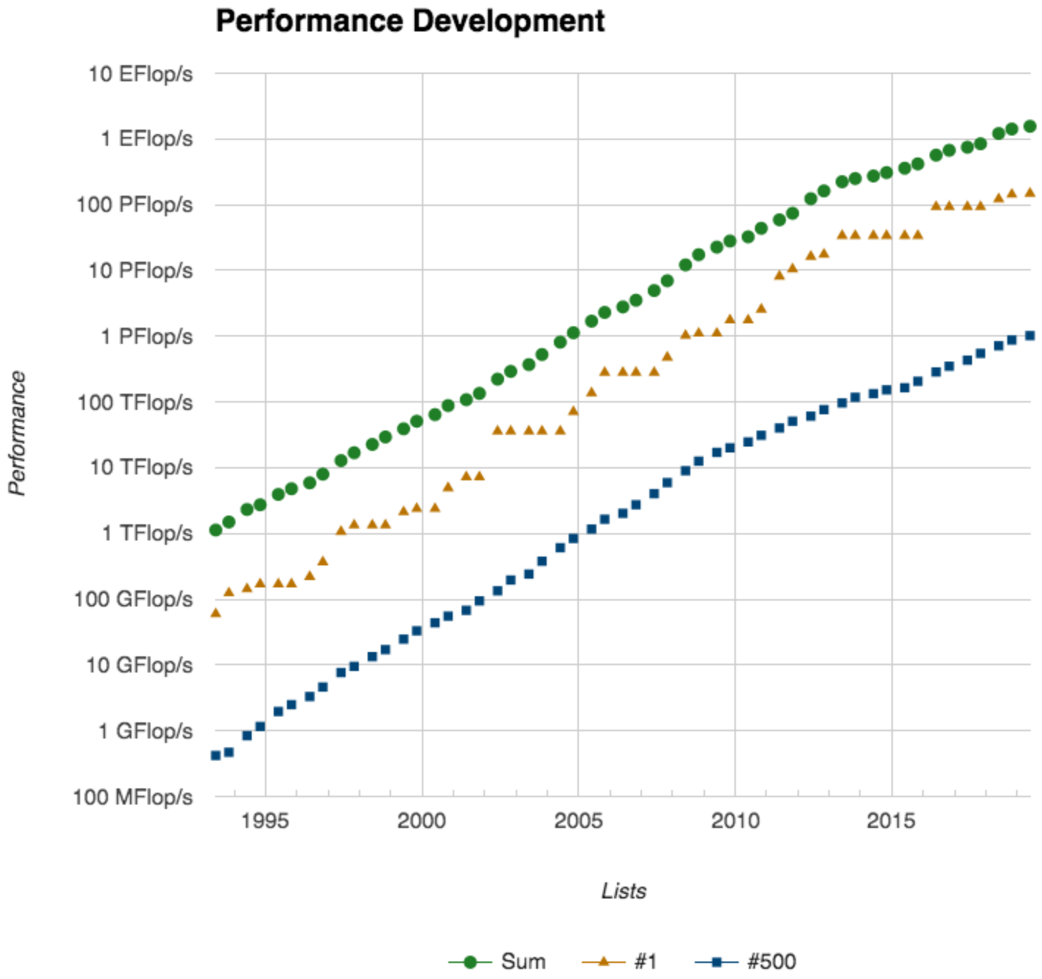

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

ORNL aspires to be among world leaders in HPC and computational sciences

- **Summit** – World's fastest computer in 2018
  200 PetaFLOPS FP64, 3.3 ExaFLOPS FP16.

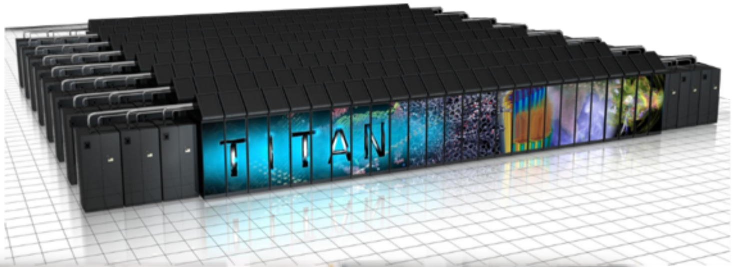

- **Titan** – World's fastest computer in 2012 and still at top 10 system in November 2018 at 27 PetaFLOPS

- **Jaguar** – World's fastest computer in 2009
  2.33 PetaFLOPS

ORNL has had a Top 10 supercomputer in every year since the Leadership Computing Facility was founded in 2005. Jaguar, Titan, and Summit are the only DOE/SC systems to be ranked #1

Summit is latest DOE #1 system on Top500

143.5 PF HPL Shows math performance

2.9 PF HPCG Shows fast data movement

14.7 GF/W Shows energy efficiency

#1 on the IO-500 Shows file system performance

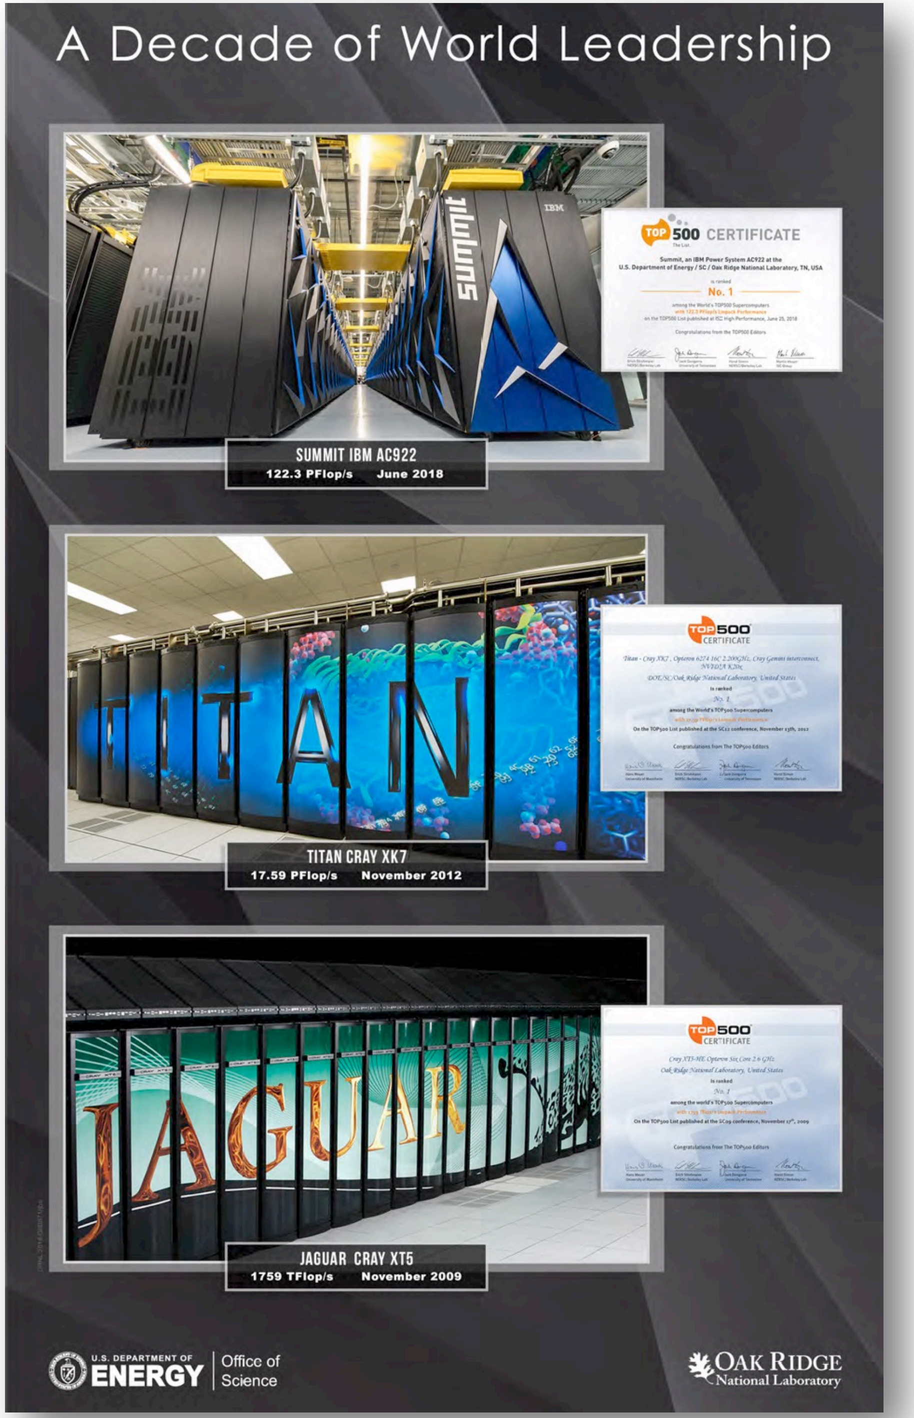

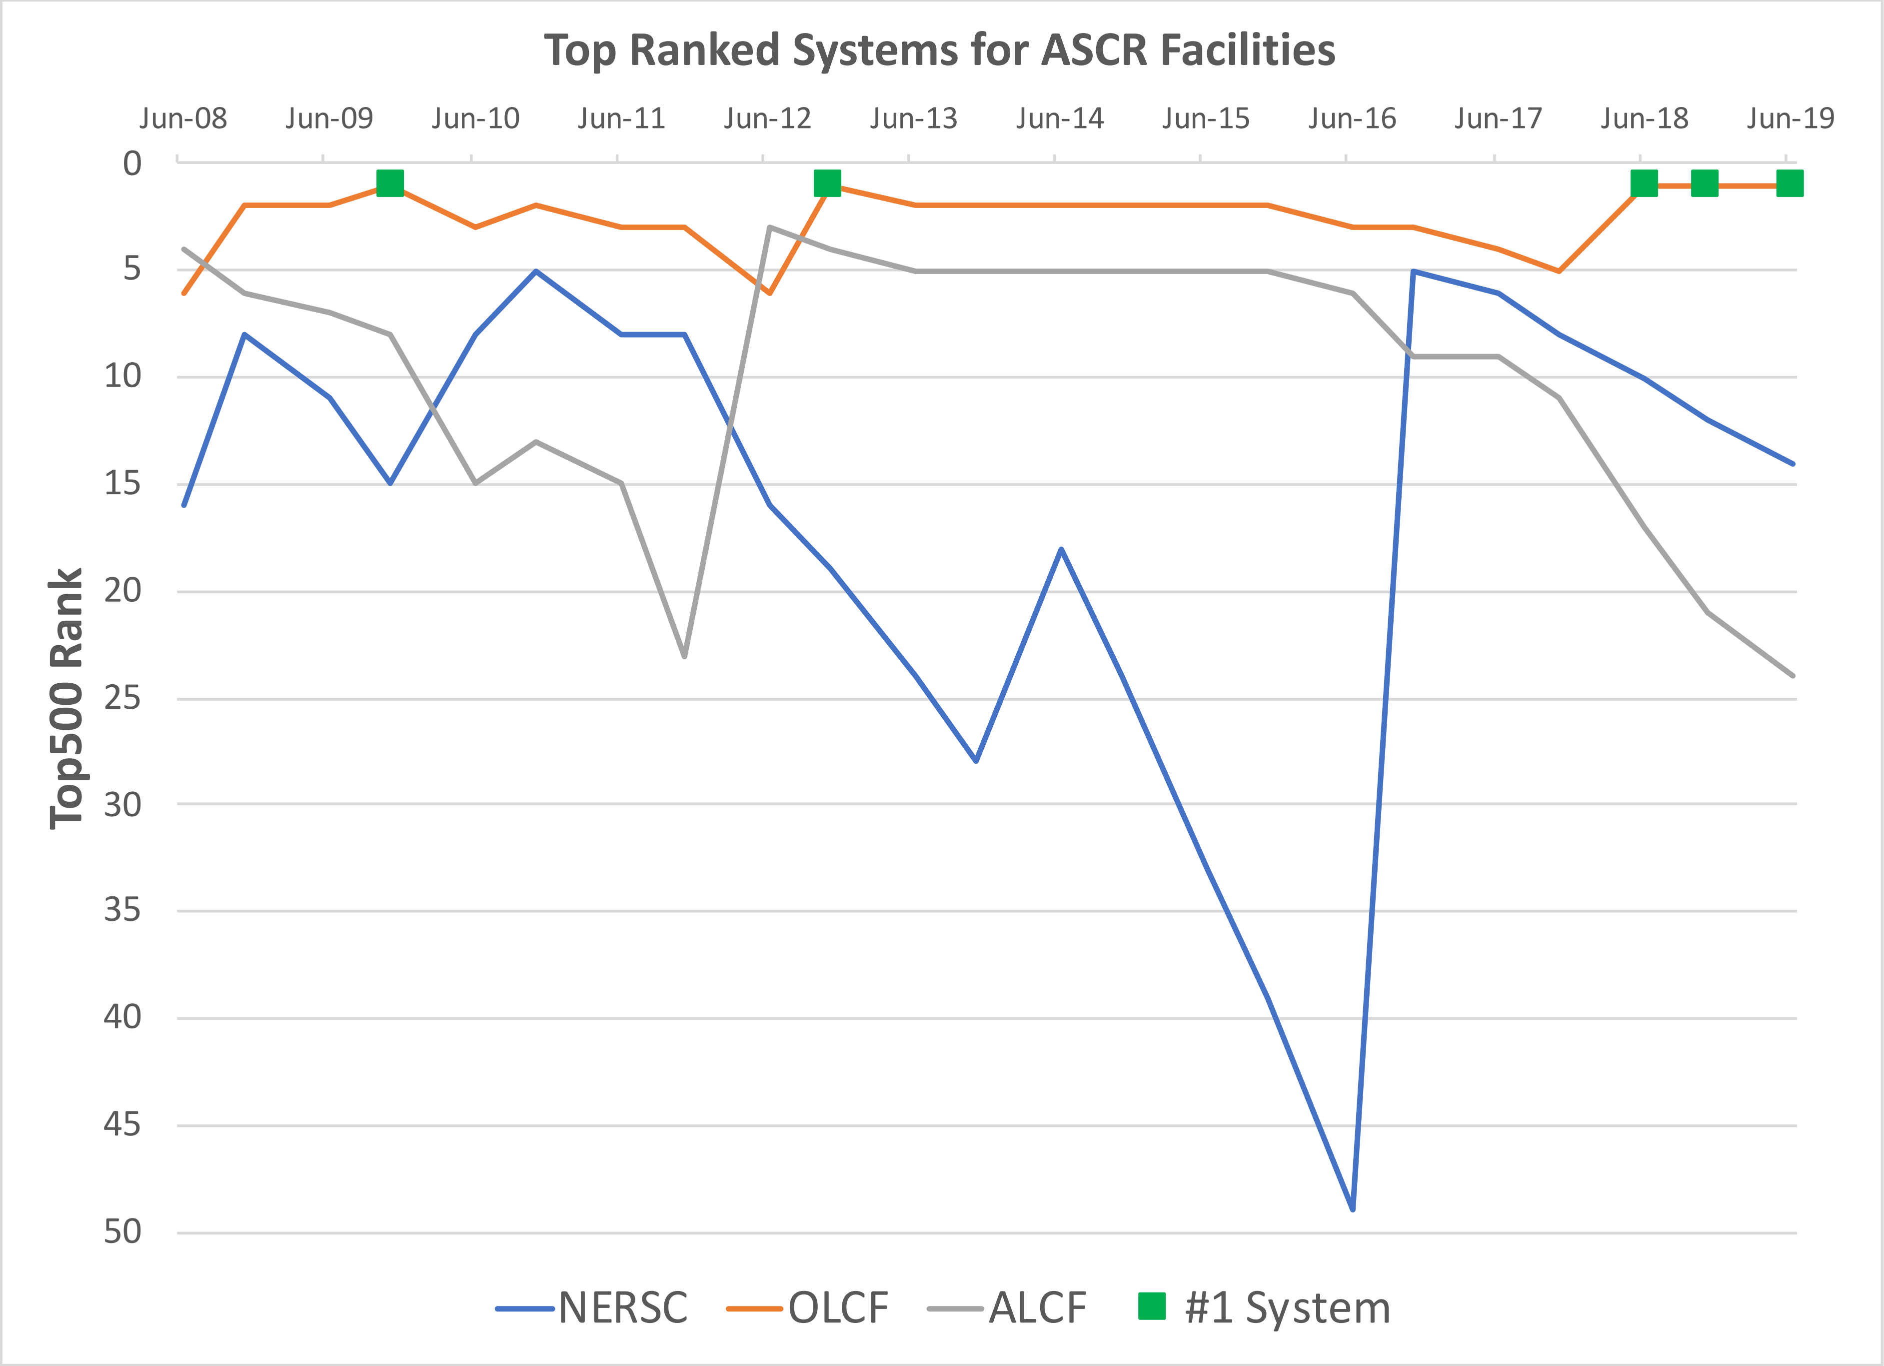
Top Ranked Systems for ASCR Facilities
Jun-08
Jun-09
Jun-10
Jun-11
Jun-12
Jun-13
Jun-14
Jun-15
Jun-16
Jun-17
Jun-18
Jun-19
Top500 Rank
0
5
10
15
20
25
30
35
40
45
50
NERSC
OLCF
ALCF
#1 System

Application Development and Performance
Center for Accelerated Application Readiness (CAAR)
Goals and Anticipated Outcomes

- Scalable, accelerated scientific applications at the start of OLCF-5 operations
- Experience translated to robust training program and Best Practices documentation
- Close collaboration with Programming Environment and Tools Team
- Further hardening of the OLCF-5 system at scale with a broader set of applications now included in the acceptance test
- Communications
  - Best Practices in conference papers
  - Reports to ASCR
- Build OLCF staff expertise to enable a smooth transition and effective support of user programs on OLCF-5

OLCF-5 CAAR Project Selection Criteria

Category	Description
Science	• Compelling scientific vision alignment with Nation's science needs • Broad coverage of science domains
Implementation (models & algorithms)	• Broad coverage of relevant programming models, environment, languages, implementations • Broad coverage of relevant algorithms and data structures
Development Plan	• Feasibility: measure of success is "Figure of Merit" compared to Summit • Clear challenge problem for execution on Frontier
Development Team	• Commitment from development team • Plan for integration with other active development directions • OLCF liaison domain-specific skills and expertise with the application • Engagement with Vendor Center of Excellence

CAAR Partnerships for Frontier Announced

https://www.olcf.ornl.gov/caar/Frontier-CAAR/

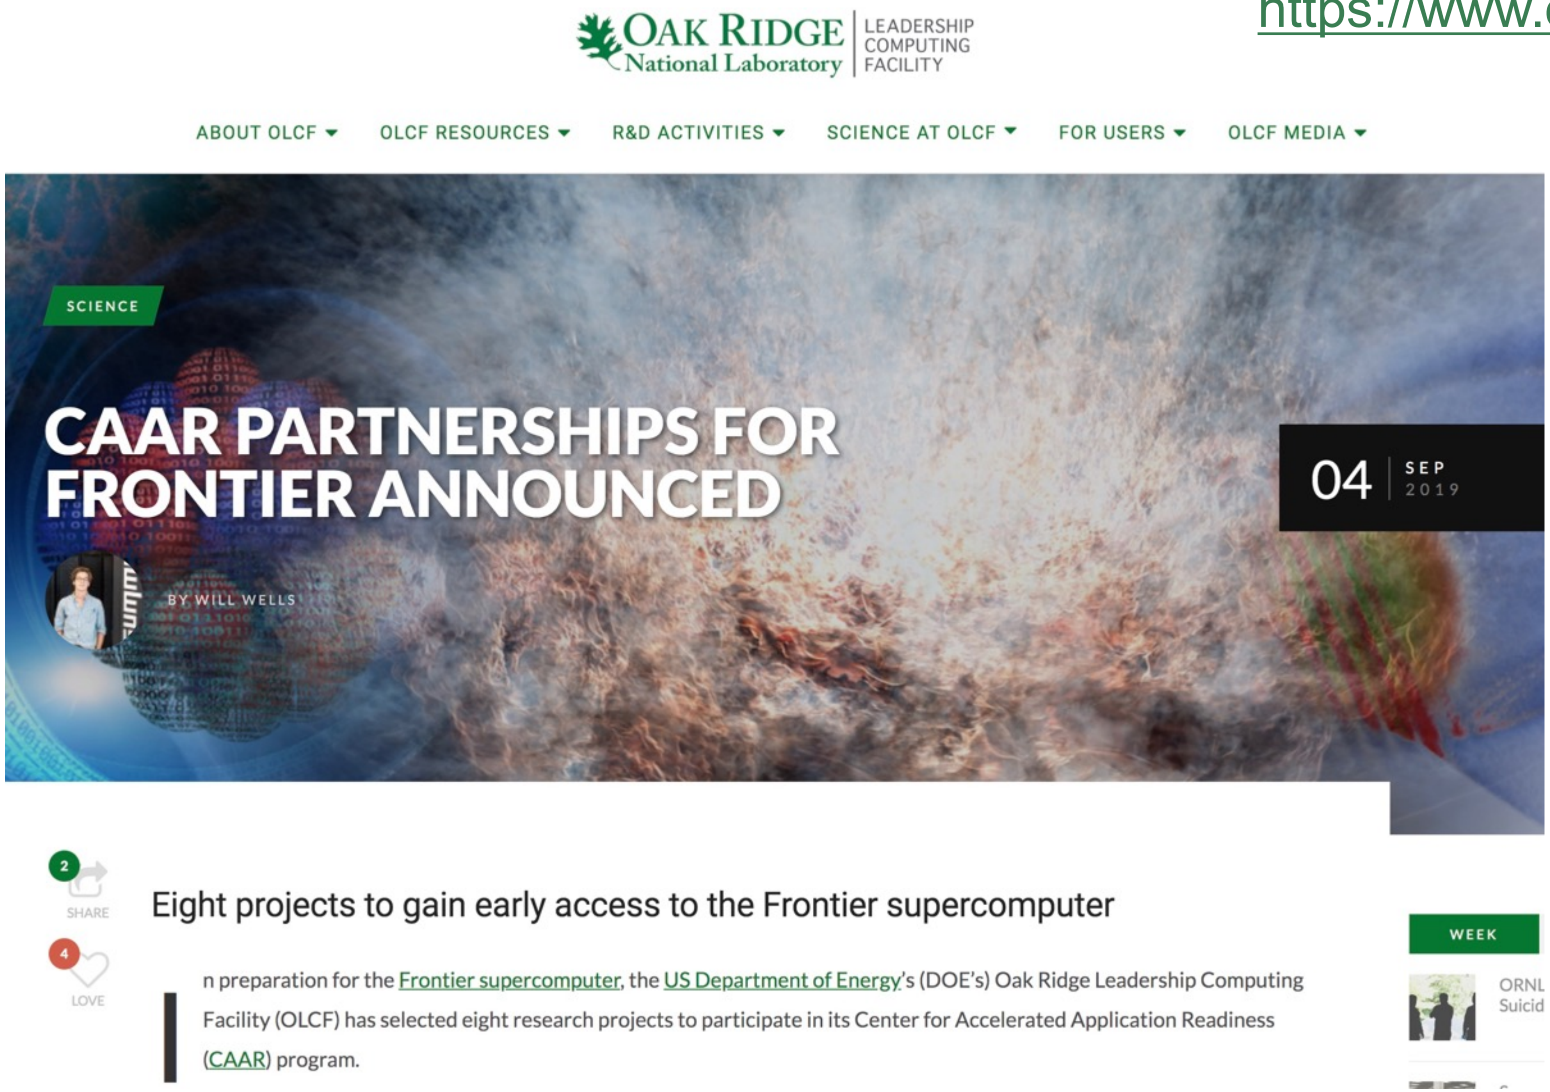

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

ABOUT OLCF OLCF RESOURCES R&D ACTIVITIES SCIENCE AT OLCF FOR USERS OLCF MEDIA

SCIENCE

CAAR PARTNERSHIPS FOR FRONTIER ANNOUNCED

BY WILL WELLS

04 SEP 2019

2 SHARE

4 LOVE

Eight projects to gain early access to the Frontier supercomputer

In preparation for the Frontier supercomputer, the US Department of Energy's (DOE's) Oak Ridge Leadership Computing Facility (OLCF) has selected eight research projects to participate in its Center for Accelerated Application Readiness (CAAR) program.

WEEK

ORNL Suicid

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

PIConGPU

Prof. Sunita Chandrasekaran,
University of Delaware

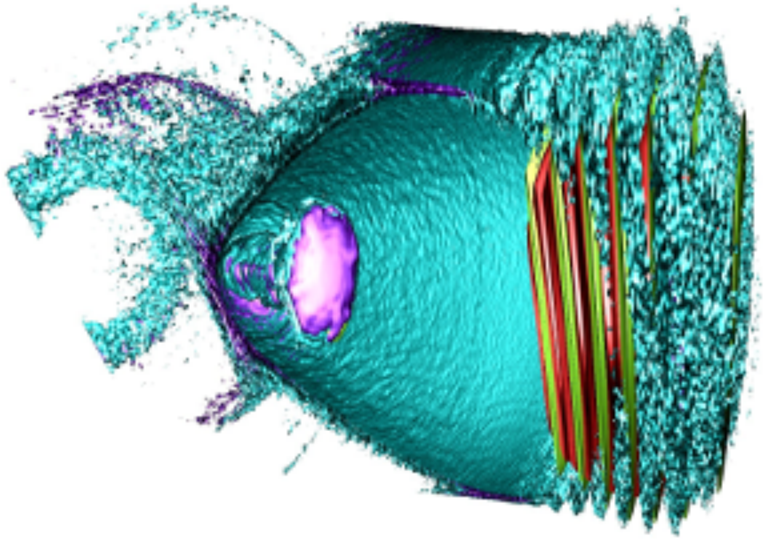

Domain Area: Plasma Physics

PIConGPU is an extremely scalable, heterogeneous, fully relativistic particle-in-cell (PIC) C++ code available as open-source project on GitHub. PIConGPU provides a modern simulation framework for laser-plasma physics and laser- matter interactions suitable for production-quality runs on state-of-the-art supercomputers driven by manycore accelerators as well as traditional architectures with a single, maintainable code base.

Hence, successfully porting PIConGPU via CAAR will immediately make more large-scale relativistic laser-matter interaction challenges accessible on FRONTIER. We expect the strongest impact for compact sources of ion beams for radiation therapy of cancer, fundamental studies of warm-dense matter and lab-astrophysics, studied using high-power laser systems.

Ronnie Chatterjee

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

Application Readiness: Community Effort

- Readiness applications will be drawn from CAAR, ECP engagement applications, as well as INCITE and ALCC projects on Summit
- CAAR provides the primary risk mitigation strategy for meeting the application readiness KPP
- CAAR is also the vanguard for the broader application readiness ecosystem and for future science
  - Development of training and documentation
  - Knowledge development for staff
  - Improvements to the software stack robustness and performance

Application Readiness

Oak Ridge National Laboratory | Leadership Computing Facility

CAAR

U.S. Department of Energy INCITE Leadership Computing

ECP Exascale Computing Project

ALCC ASCR Leadership Computing Challenge

Compiler Development and Performance

Summit
OpenMP -- XL 16.1.0
OpenACC -- PGI 18.3

Titan
OpenMP - CCE 8.7.0
OpenACC - PGI 18.4

Ratio = OpenACC time / OpenMP time

SPEC ACCEL 1.2 Benchmark	Summit	Titan
Stencil (C)	1.12	0.67
LBM (C)	1.77	1.95
MRI-Q (C)	0.9	0.92
MD (Fortran)	0.14	
EP (Fortran)	0.27	1.07
CLVRLEAF (C,Fortran)	0.88	
CG (C)	0.32	1.12
SEISMIC (C)	1.49	
SP (Fortran)	1.27	
SP (C)	0.21	0.54
MinGhost (C,Fortran)	2.18	
LBDC (Fortran)	2.45	
Swim (Fortran)	1.18	
BT (C)	0.12	0.7

OpenMP Enabling HPC since 1997

OpenMP is Better

OpenACC is Better

OpenACC More Science, Less Programming

Update from Swen Boehm, Swaroop Pophale, Veronica G. Vergara Larrea, and Oscar Hernandez. "Evaluating Performance Portability of Accelerator Programming Models using SPEC ACCEL 1.2 Benchmarks" P3MA, ISC'2018

Snapshot of default compilers September 2018

Disclaimer –research use of SPEC non-compliant results

Status of OpenACC and OpenMP compilers on Summit

	Open ACC						Open MP					
			GNU		PGI				GNU		XL	
	Benchmark	Reference time	Runtime	Spec Score	Runtime	Spec Score	Benchmark	Reference time	Runtime	Spec Score	Runtime	Spec Score
ostencil	303.ostencil	145			12.1		503.postencil	109			10.2	
olbm	304.olbm	455			36.3		504.polbm	122			19.9	
omriq	314.omriq	956			35.5		514.pomriq	621			45.5	
md	350.md	252			9.28		550.pmd	241			21.2	
palm	351.palm	370			117		551.ppalm	544			203	
ep	352.ep	530			45.8		552.pep	231			179	
clvrleaf	353.clvrleaf	445			35.9		553.pclvrleaf	1145			55.5	
cg	354.cg	408			31.2		554.pcg	333			72.8	
seismic	355.seismic	370			26		555.pseismic	282			45.8	
sp	356.sp	276			21.6		556.psp	818			29.3	
csp	357.csp	270			19.5		557.pcsp	859			92.4	
miniGhost	359.miniGhost	369			35.8		559.pmniGhost	397			41.5	
ilbdc	360.ilbdc	367			27.3		560.pilbdc	653			30.5	
swim	363.swim	230			34.2		563.pswim	159			28	
bt	370.bt	223			9.37		570.pbt	780			75.7	

Unofficial SPEC HPG result – Academic use

Acknowledgment: Swen Boehm (ORNL)

Disclaimer –research use of SPEC non-compliant results

CORAL Procurements

Objective - Procure 3 leadership computers to be sited at Argonne, Oak Ridge and Lawrence Livermore in 2017.

Current DOE Leadership Computers

Titan (ORNL)
2012 - 2017

Sequoia (LLNL)
2012 - 2017

Mira (ANL)
2012 - 2017

Leadership Computers RFP requested >100 PF, 2 GB/core main memory, local NVRAM, and science performance 4x-8x Titan or Sequoia

Approach

- Competitive process - one RFP (issued by LLNL) leading to 2 R&D contracts and 3 computer procurement contracts
- For risk reduction and to meet a broad set of requirements, 2 architectural paths were selected and Oak Ridge and Argonne must choose different architectures
- Multi-year Lab-Awardee relationship to co-design computers
- Both R&D contracts jointly managed by the 3 Labs
- Each lab manages and negotiates its own computer procurement contract, and may exercise options to meet their specific needs
- Understanding that long procurement lead-time may impact architectural characteristics and designs of procured computers

Oak Ridge National Laboratory · Argonne National Laboratory · Lawrence Livermore National Laboratory

CORAL (I) Results

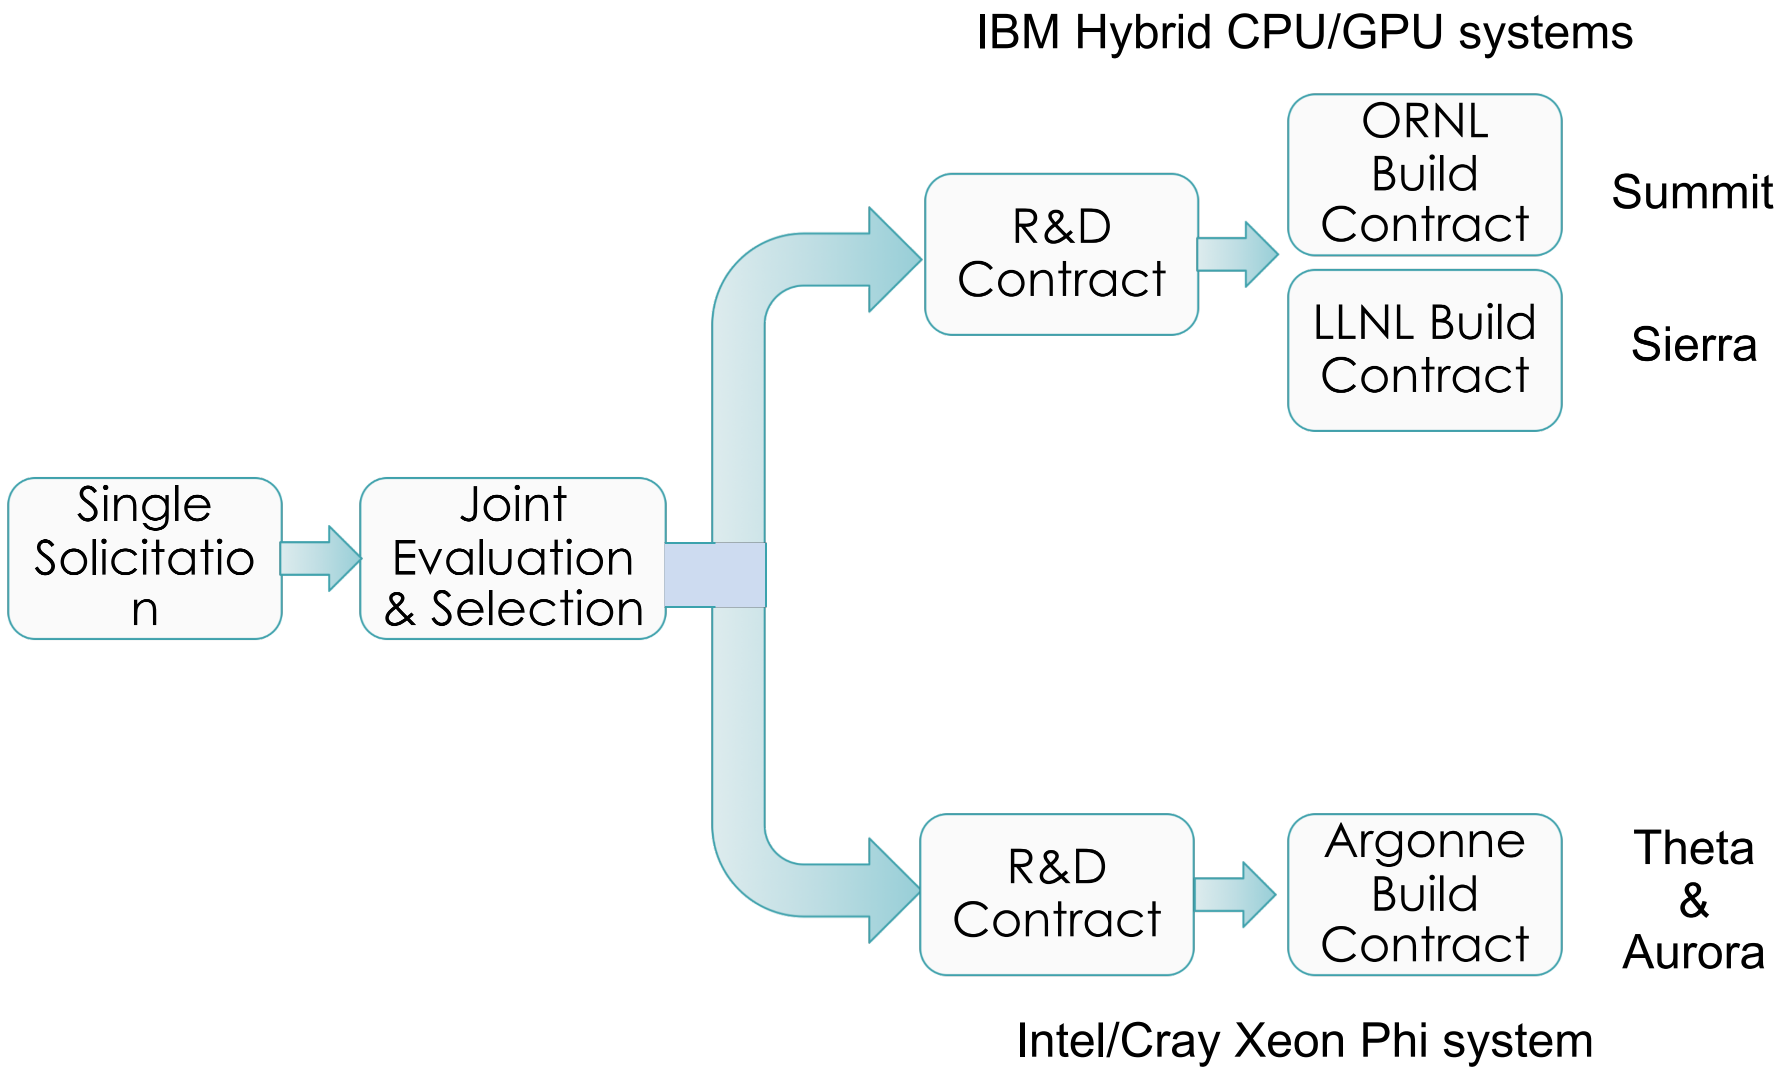

Wants and constraints

- CORAL benchmarks should
  - Span the breadth of the NNSA (LLNL) workload
  - Span the time-dependent(!) and much broader space of LCF workloads
  - Span co-spaces of algorithms, implementations, and use cases
  - Provide adequate drivers for system SW and library development
- CORAL benchmarks must
  - ..not be so numerous that vendors cannot provide sophisticated analyses on O(weeks) time scale
    - Significant challenge to cover/span the breadth of concerns, while not being onerous on vendors.
  - …not encumber application developers with 24-7 support responsibilities during those weeks
  - …use proxies for NNSA apps

CORAL Benchmark Codes

- Scalable Science Benchmarks: LSMS, QBOX, HACC, Nekbone
- Throughput Benchmarks: CAM-SE, UMT2013, AMG2013, MCB, QMCPACK, NAMD, LULESH, SNAP, miniFE,
- Data-Centric Benchmarks: Graph500, Integer Sort, Hash, SPECint2006 “peak”
  - last-minute addition, in hindsight: hurried
- Skeleton Benchmarks
- Microkernel Benchmarks

https://asc.llnl.gov/CORAL-benchmarks/

CORAL-2 Benchmark Codes

Benchmarks

- Scalable Science Benchmarks: HACC, Nekbone, QMCPACK, LAMMPS
- Throughput Benchmarks: AMG, Kripke, Quicksilver, PENNANT
- Data Science and Deep Learning Benchmarks:
  - Big Data Analytic Suite
    - [Schmidt, et al., "Defining Big Data Analytics Benchmarks for Next Generation Supercomputers," https://arxiv.org/abs/1811.02287]
  - Deep Learning Suite
- Skeleton Benchmarks
- Microkernel Benchmarks

https://asc.llnl.gov/coral-2-benchmarks/

User Program Management:
INCITE System Capability Metric (ISCM)

- Goals:
  - Develop a metric that more accurately reflects system capability for the execution of science applications
  - (Potentially) use metric for "currency unit" in user-program allocations.
- Use a modified version of the HPC Challenge benchmark suite (https://icl.utk.edu/hpcc/) to build a measure we call the INCITE System Capability Metric (ISCM).
  - Use geometric mean of the value for each benchmark component.
  - Measures capability for a machine to execute the motifs in suite.
- Focus on systems allocated under the INCITE program.
- Use application benchmarking to validate and support ISCM metric.

R.D. Budiardja, W. Joubert, J. A. Harris, A. Tillack, T. L. Papatheodore, "ISCM: Towards a Comprehensive Metric For Comparative Evaluation of Leadership-Class System Capability for Scientific Applications" (unpublished, 2019)

ISCM Components

1. HPL
2. DGEMM
3. ML/HGEMM
4. FFT
5. STREAM
6. RandomAccess
7. Communication Bandwidth
8. Communication Latency
9. PTRANS

R.D. Budiardja, W. Joubert, J. A. Harris, A. Tillack, T. L. Papatheodore, “ISCM: Towards a Comprehensive Metric For Comparative Evaluation of Leadership-Class System Capability for Scientific Applications” (unpublished, 2019)

ISCM Benchmarking by Applications

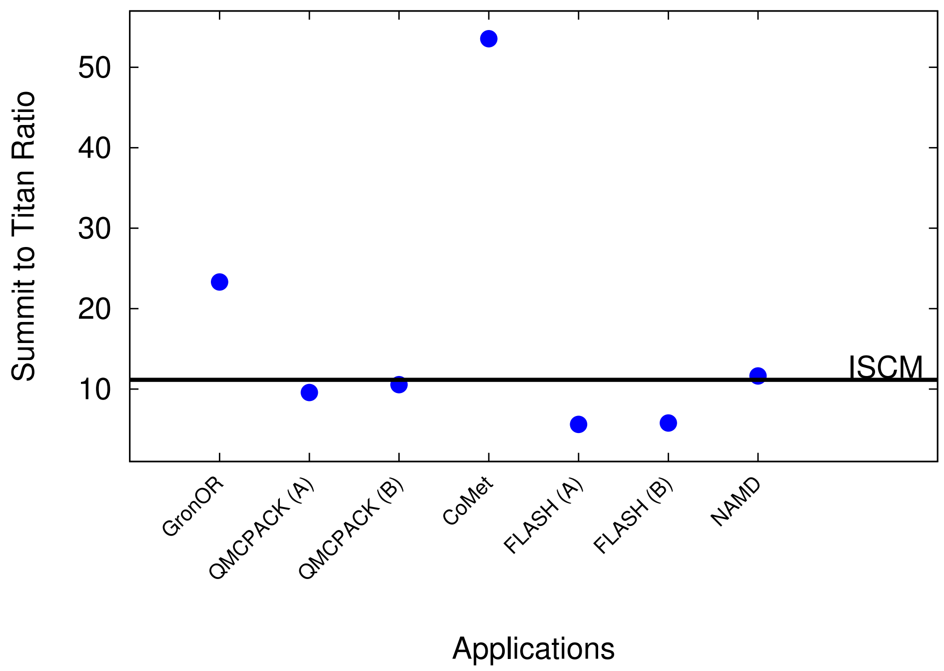

Figure 2: The performance ratio of several applications run on Summit vs. on Titan as compared to the ratio of the ISCM values of the two systems plotted as solid line.

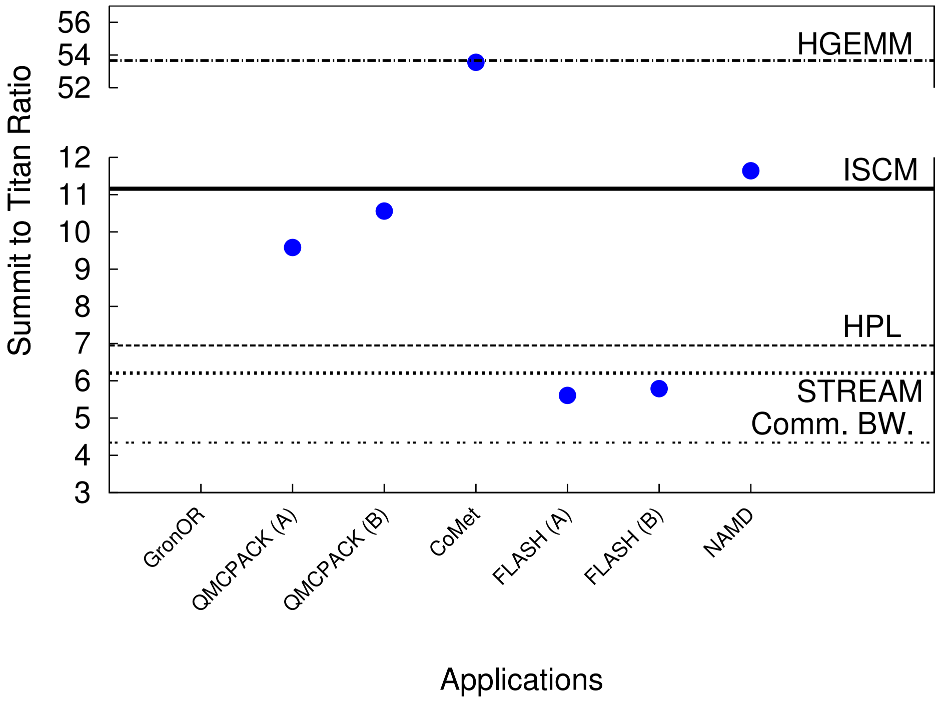

Figure 3: Same as Figure 2 with the addition of the ratio of several benchmark components plotted as dot-dashed (`HGEMM`), dashed `HPL`, dotted (`STREAM`), and dash-dashed (`Communication Bandwidth`).

R.D. Budiardja, W. Joubert, J. A. Harris, A. Tillack, T. L. Papatheodore, “ISCM: Towards a Comprehensive Metric For Comparative Evaluation of Leadership-Class System Capability for Scientific Applications” (unpublished, 2019)

Summary & Discussion

- OLCF is engaged in a variety of mission-critical activities that require application and motif benchmarking.
- Flexibility is necessary in accomplishing activities.
  - "Different horses for difference courses".
- Sustainability and maintainability are key problems to address.
- ORNL participation in SPECHPG provide real value to many mission-critical functions.