Center for Accelerated Application Readiness

Call for Proposals

Getting Applications Ready for **Summit**

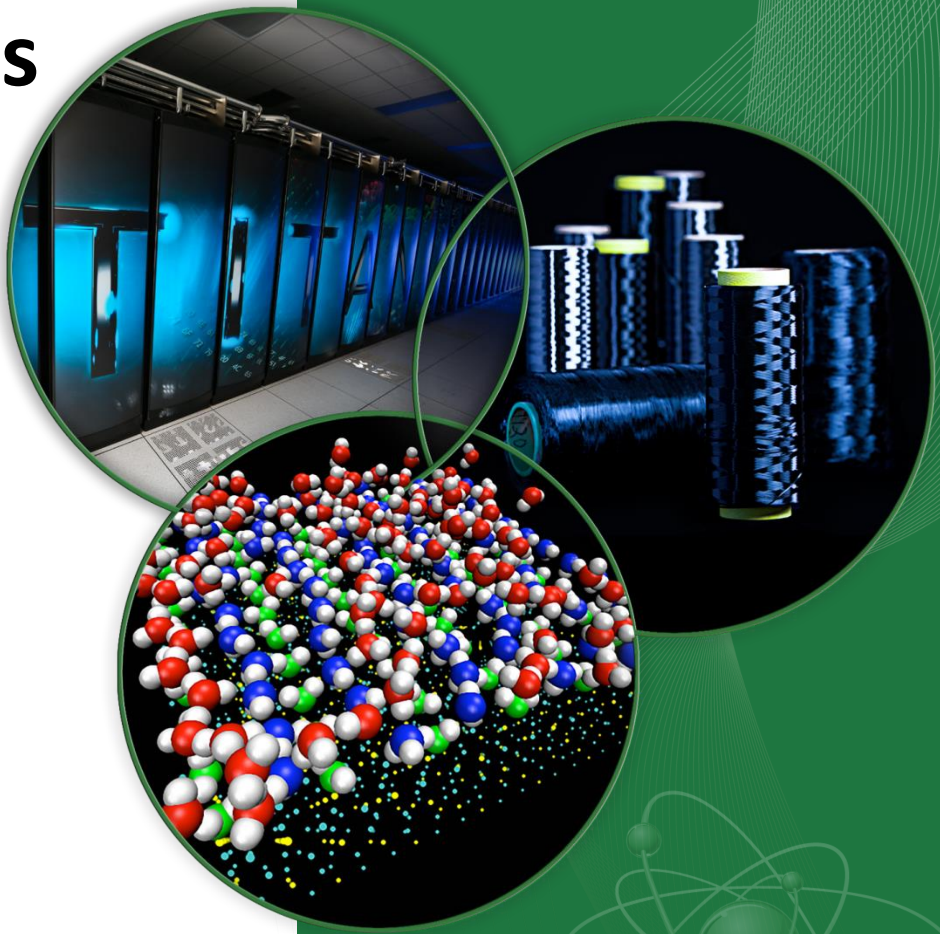

Tjerk Straatsma

OLCF Scientific Computing Group

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

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

Summit
OLCF's Next Leadership System

OLCF on the Road to Exascale

Since clock-rate scaling ended in 2003, HPC performance has been achieved through increased parallelism. Jaguar scaled to 300,000 cores.	Titan and beyond deliver hierarchical parallelism with very powerful nodes. MPI plus thread level parallelism through OpenACC or OpenMP plus vectors

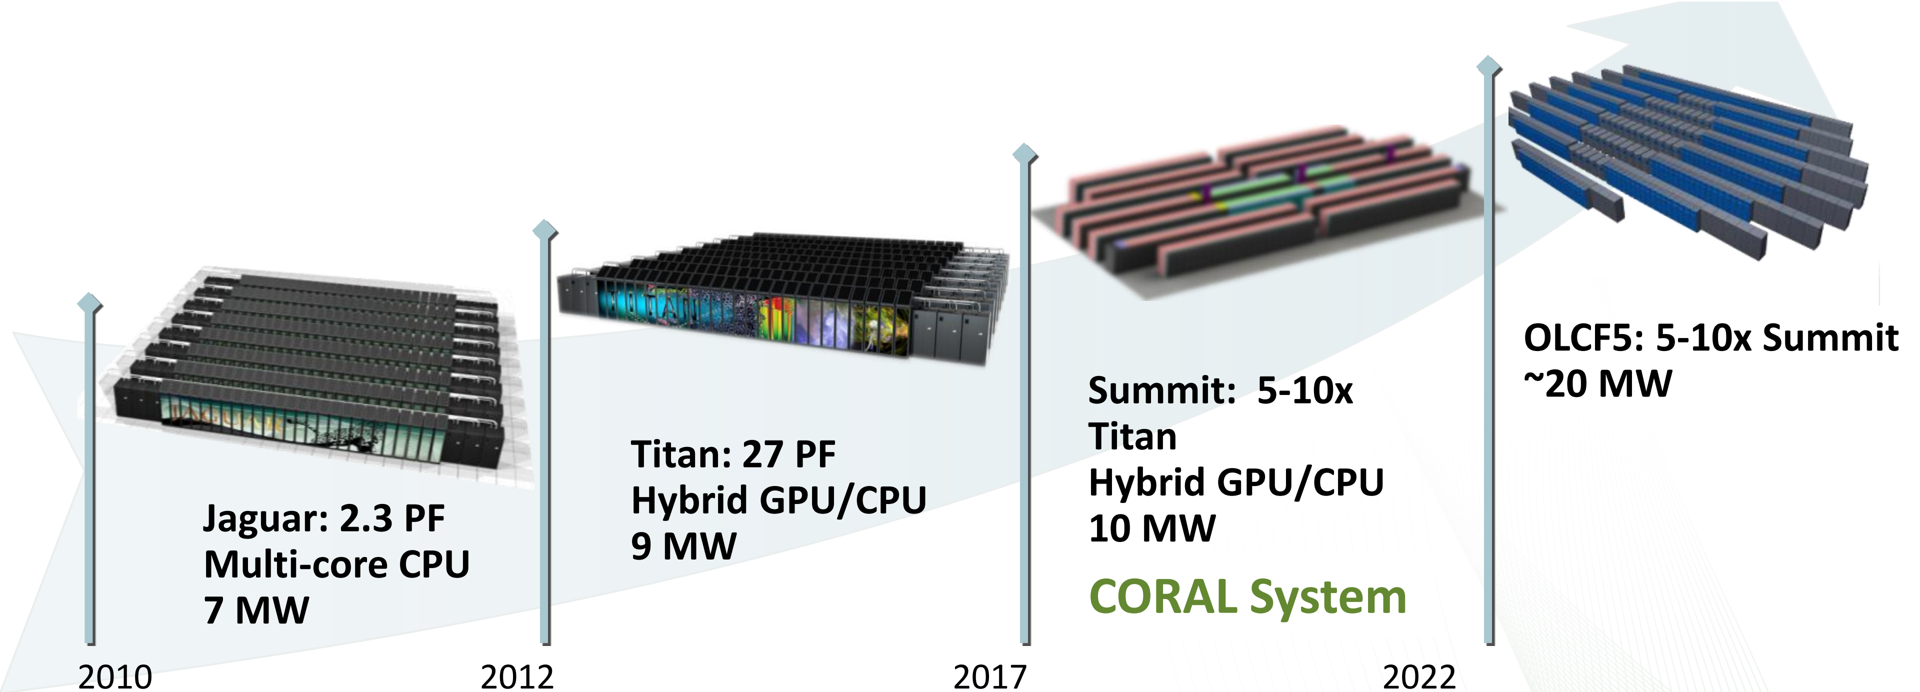

OLCF Current Systems

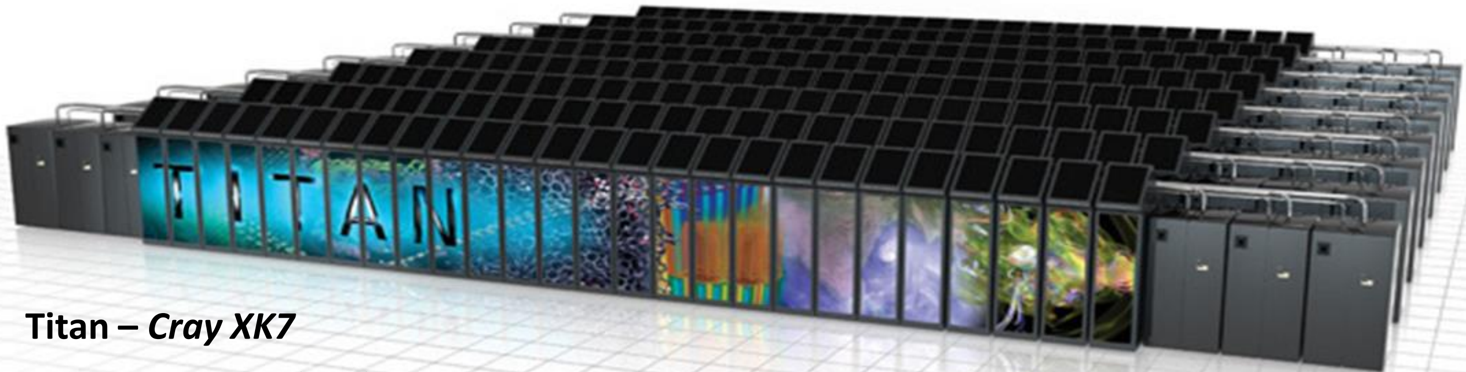

Titan – *Cray XK7*

- 18,688 nodes, each with
  - AMD™ Opteron™- 141 GF, 32 GB DDR3 memory
  - NVIDIA™ Kepler™ K20X GPU - 1,311 GF, 6 GB GDDR5 memory
  - PCIe2 link between GPU and CPU
- Cray Gemini 3-D Torus Interconnect
- 688 TB of memory
- Peak flop rate: 27 PF

Eos – *Cray XC30*

- 744 nodes, Intel Xeon E5-2670
- 48 TB of memory
- 248 TF

Rhea

- Pre- and post-processing cluster
- 512 nodes, dual 8c Xeon, 64 GB

EVEREST– Visualization Laboratory

- Stereoscopic 6x3 1920x1080 Display Wall, 30.5’ x 8.5’
- Planar 4x4 1920x1080 Display Wall
- Distributed memory Linux cluster

Storage – Spider Lustre® filesystem: 40 PB, >1 TB/s BW; HPSS archival mass storage: 240PB, 6 tape libraries

OLCF Next System: Summit

Vendor: IBM® (Prime) / NVIDIA™ / Mellanox Technologies®

At least 5X Titan's Application Performance

Approximately 3,400 nodes, each with:

- Multiple IBM POWER9 CPUs and multiple NVIDIA Tesla® GPUs using the NVIDIA Volta™ architecture
- CPUs and GPUs completely connected with high speed NVLink™
- Large coherent memory: over 512 GB (HBM + DDR4)
  - all directly addressable from the CPUs and GPUs
- An additional 800 GB of NVRAM, which can be configured as either a burst buffer or as extended memory
- over 40 TF peak performance

Dual-rail Mellanox® EDR-IB full, non-blocking fat-tree interconnect

IBM Elastic Storage (GPFS™) - 1TB/s I/O and 120 PB disk capacity.

OLCF Summit Key Software Components

- **System**
  - Linux®
  - IBM Elastic Storage (GPFS™)
  - IBM Platform Computing™ (LSF)
  - IBM Platform Cluster Manager™ (xCAT)
- **Programming Environment**
  - Compilers supporting OpenMP, OpenACC, CUDA
    - IBM XL, PGI, LLVM, GNU, NVIDIA
  - Libraries
    - IBM Engineering and Scientific Subroutine Library (ESSL)
    - FFTW, ScaLAPACK, PETSc, Trilinos, BLAS-1,-2,-3, NVBLAS
    - cuFFT, cuSPARSE, cuRAND, NPP, Thrust
  - Debugging
    - Allinea DDT, IBM Parallel Environment Runtime Edition (pdb)
    - Cuda-gdb, Cuda-memcheck, valgrind, memcheck, helgrind, stacktrace
  - Profiling
    - IBM Parallel Environment Developer Edition (HPC Toolkit)
    - VAMPIR, Tau, Open|Speedshop, nvprof, gprof, Rice HPCToolkit

Summit compared to Titan

Feature	Summit	Titan
Application Performance	5-10x Titan	Baseline
Number of Nodes	~3,400	18,688
Node performance	> 40 TF	1.4 TF
Memory per Node	>512 GB (HBM + DDR4)	38GB (GDDR5+DDR3)
NVRAM per Node	800 GB	0
Node Interconnect	NVLink (5-12x PCIe 3)	PCIe 2
System Interconnect (node injection bandwidth)	Dual Rail EDR-IB (23 GB/s)	Gemini (6.4 GB/s)
Interconnect Topology	Non-blocking Fat Tree	3D Torus
Processors	IBM POWER9 NVIDIA Volta™	AMD Opteron™ NVIDIA Kepler™
File System	120 PB, 1 TB/s, GPFS™	32 PB, 1 TB/s, Lustre®
Peak power consumption	10 MW	9 MW

Application Readiness

Two Tracks for Future Large Systems

Tianhe-2 (NUDT): TH-IVB-FEP
Intel Xeon E5-2692 12 C 2.2 GHz
TH Express-2
Intel Xeon Phi 31S1P

Titan (Cray): Cray XK7
AMD Opteron 6274 16C 2.2 GHz
Cray Gemini
NVIDIA K20x

Sequoia (IBM): BlueGene/Q
Power BQC 16C 1.6 GHz

K computer (Fujitsu)
SPARC64 VIIIfx 2.0 GHz
Tofu

Mira (IBM): BlueGene/Q
PowerPC A2 16C 1.6 GHz

Piz Daint (Cray): Cray XC30
Intel Xeon E5-2670 8C 2.6 GHz
Cray Aries
NVIDIA K20x

Edison (Cray): Cray XC30
Intel Xeon E%-2695v2 12C 2.4 GHz
Aries

Many Core

- 10's of thousands of nodes with millions of cores
- Homogeneous cores
- Multiple levels of memory – on package, DDR, and non-volatile
- Unlike prior generations, future products are likely to be self hosted

Hybrid Multi-Core

- CPU / GPU Hybrid systems
- Likely to have multiple CPUs and GPUs per node
- Small number of very fat nodes
- Expect data movement issues to be much easier than previous systems – coherent shared memory within a node
- Multiple levels of memory – on package, DDR, and non-volatile

Cori at NERSC

- Self-hosted many-core system
- Intel/Cray
- 9300 single-socket nodes
- Intel® Xeon Phi™ Knights Landing (KNL)
- 16GB HBM, 64-128 GB DDR4
- Cray Aries Interconnect
- 28 PB Lustre file system @ 430 GB/s
- Target delivery date: June, 2016

Summit at OLCF

- Hybrid CPU/GPU system
- IBM/NVIDIA
- 3400 multi-socket nodes
- POWER9/Volta
- More than 512 GB coherent memory per node
- Mellanox EDR Interconnect
- Target delivery date: 2017

ALCF-3 at ALCF

- TBA
- Target delivery date: 2017-18

Importance of Performance Portability

Application portability among NERSC, ALCF and OLCF architectures is critical concern of ASCR

- Application developers target wide range of architectures
- Maintaining multiple code version is difficult
- Porting to different architectures is time-consuming
- Many Principal Investigators have allocations on multiple resources
- Applications far outlive any computer system

Primary task is exposing parallelism and data locality

Challenge is to find the right abstraction:

- MPI + X (X=OpenMP, OpenACC)
- PGAS + X
- DSL
- …

Synergy Between Application Readiness Programs

NESAP at NERSC

NERSC Exascale Science Application Program

- Call for Proposals – June 2014
- 20 Projects selected
- Partner with Application Readiness Team and Intel IPCC
- 8 Postdoctoral Fellows

Criteria

- An application's computing usage within the DOE Office of Science
- Representation among all 6 Offices of Science
- Ability for application to produce scientific advancements
- Ability for code development and optimizations to be transferred to the broader community through libraries, algorithms, kernels or community codes
- Resources available from the application team to match NERSC/Vendor resources

CAAR at OLCF

Center for Accelerated Application Readiness

- Call for Proposals – November 2014
- 8 Projects to be selected
- Partner with Scientific Computing group and IBM/NVIDIA Center of Excellence
- 8 Postdoctoral Associates

Criteria

- Anticipated impact on the science and engineering fields
- Importance to the user programs of the OLCF
- Feasibility to achieve scalable performance on Summit
- Anticipated opportunity to achieve performance portability for other architectures
- Algorithmic and scientific diversity of the suite of CAAR applications.
- Optimizations incorporated in master repository
- Size of the application's user base

ESP at ALCF

Early Science Program

- Call for Proposals
- 10 Projects to be selected
- Partner with Catalyst group and ALCF Vendor Center of Excellence
- Postdoctoral Appointee per project

Criteria

- Science Impact
- Computational Readiness
- Proposed science problem of appropriate scale to exercise capability of new machine
- Confidence code will be ready in time
- Project code team appropriate
- Willing partner with ALCF & vendor
- Diversity of science and numerical methods
- Samples spectrum of ALCF production apps

Synergy Between Application Readiness Programs

- Application Developer Team involvement
  - Knowledge of the application
  - Work on application in development "moving target"
  - Optimizations included in application release
- Early Science Project
  - Demonstration of application on real problems at scale
  - Shake-down on the new system hardware and software
  - Large-scale science project is strong incentive to participate
- Vendor technical support is crucial
  - Programming environment often not mature
  - Best source of information on new hardware features
- Access to multiple resources, including early hardware
- Joint training activities

- Portability is a critical concern
- Experience benefits other developers and users
  - Coverage of scientific domains
  - Coverage of algorithmic methods and programming models
- Persistent culture of application readiness
  - More computational ready applications available
  - Experience of science liaisons and catalysts for user programs
  - Synergy with libraries and tools projects

Application Readiness Tentative Timelines

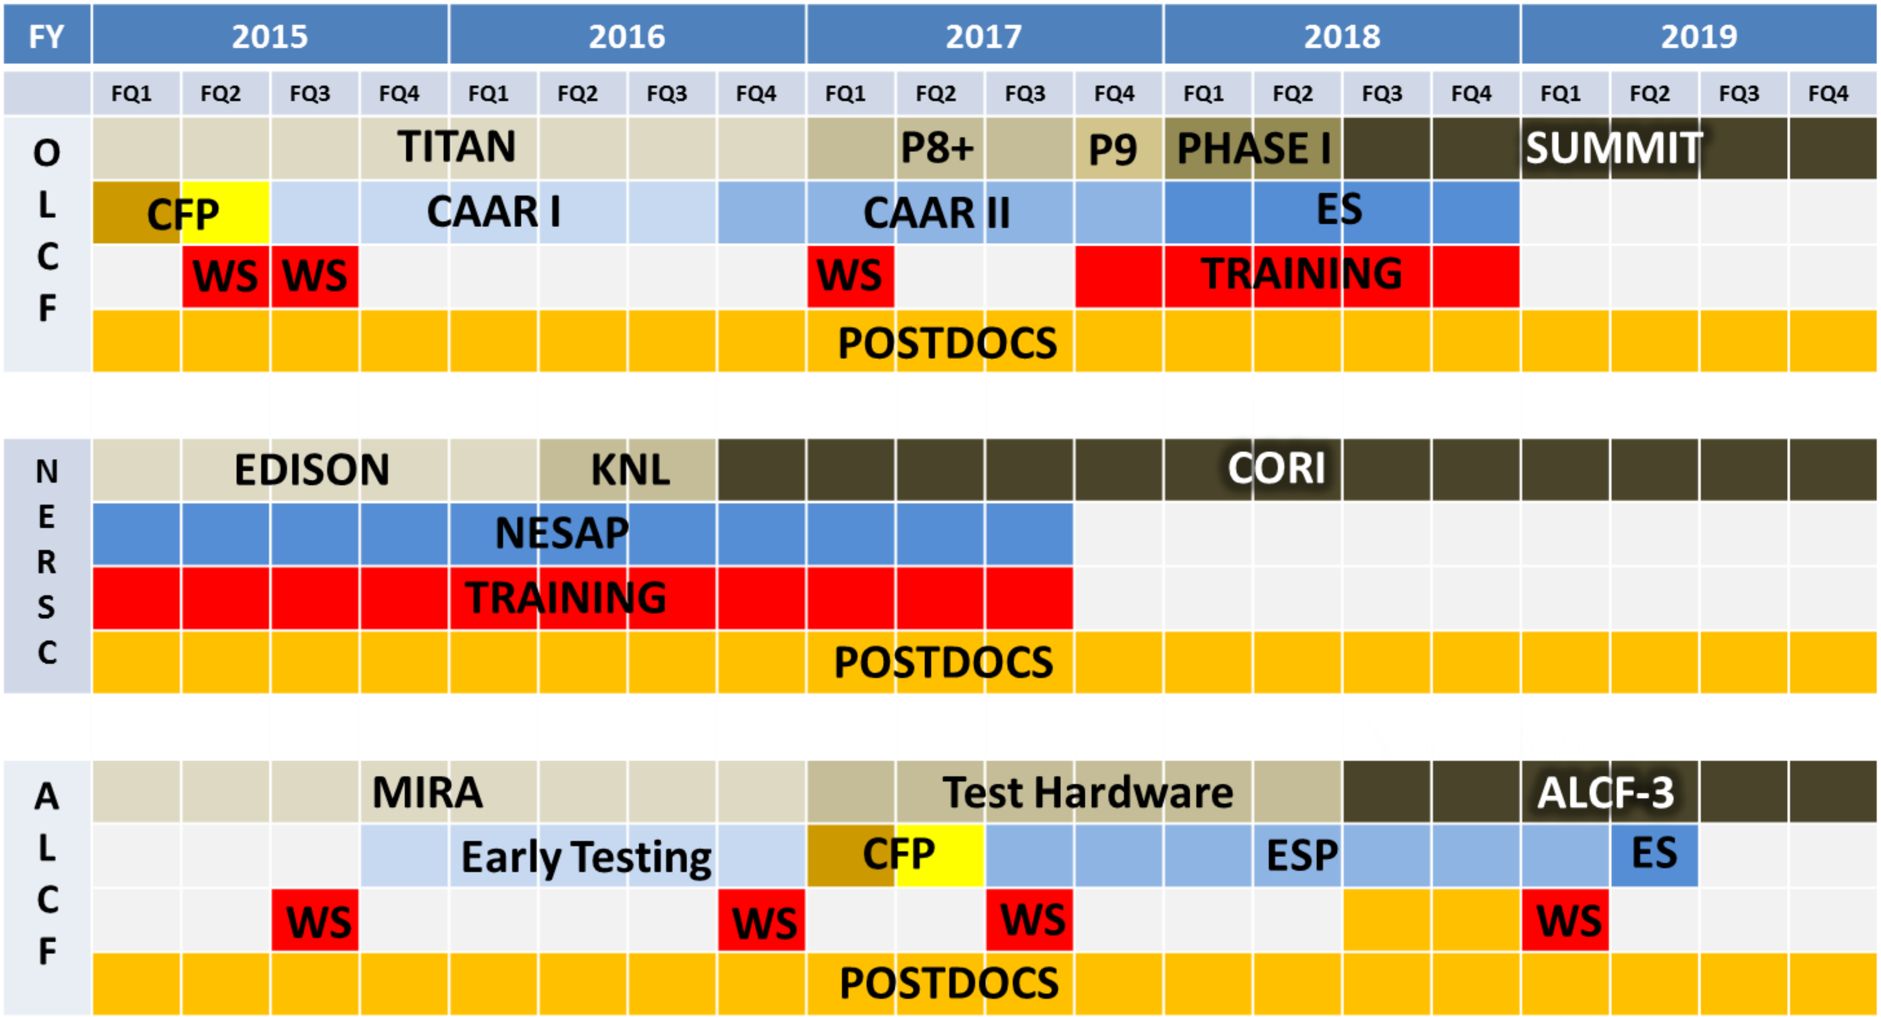

OLCF Center for Accelerated Application Readiness

WL-LSMS
Illuminating the role of material disorder, statistics, and fluctuations in nanoscale materials and systems.

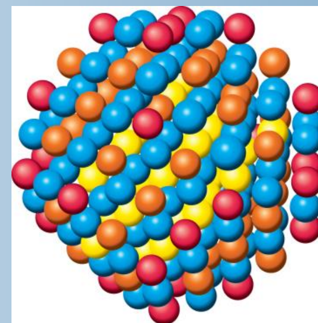

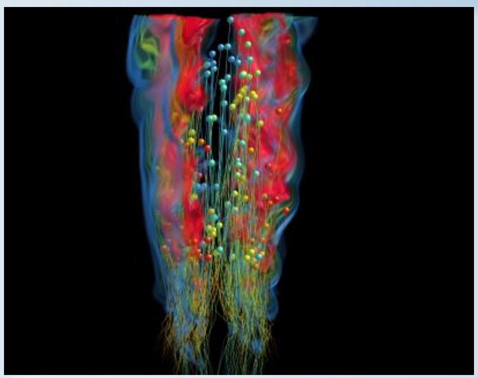

S3D
Understanding turbulent combustion through direct numerical simulation with complex chemistry.

NRDF
Radiation transport – important in astrophysics, laser fusion, combustion, atmospheric dynamics, and medical imaging – computed on AMR grids.

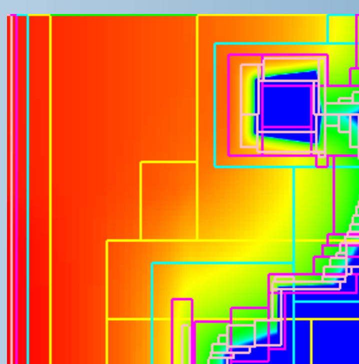

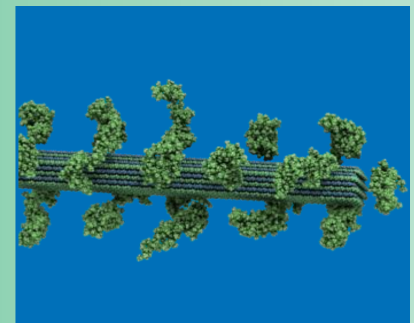

LAMMPS
A molecular description of membrane fusion, one of the most common ways for molecules to enter or exit living cells.

CAM-SE
Answering questions about specific climate change adaptation and mitigation scenarios; realistically represent features like precipitation patterns / statistics and tropical storms.

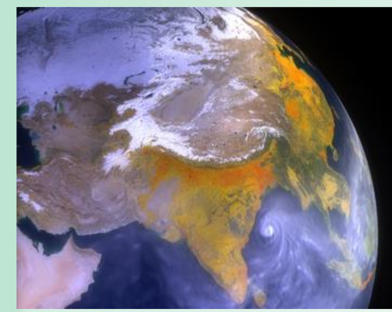

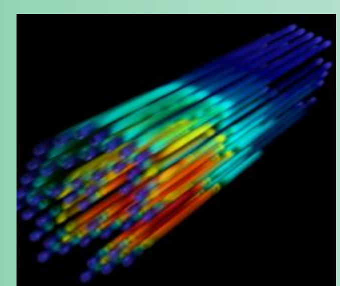

Denovo
Discrete ordinates radiation transport calculations that can be used in a variety of nuclear energy and technology applications.

OLCF Center for Accelerated Application Readiness

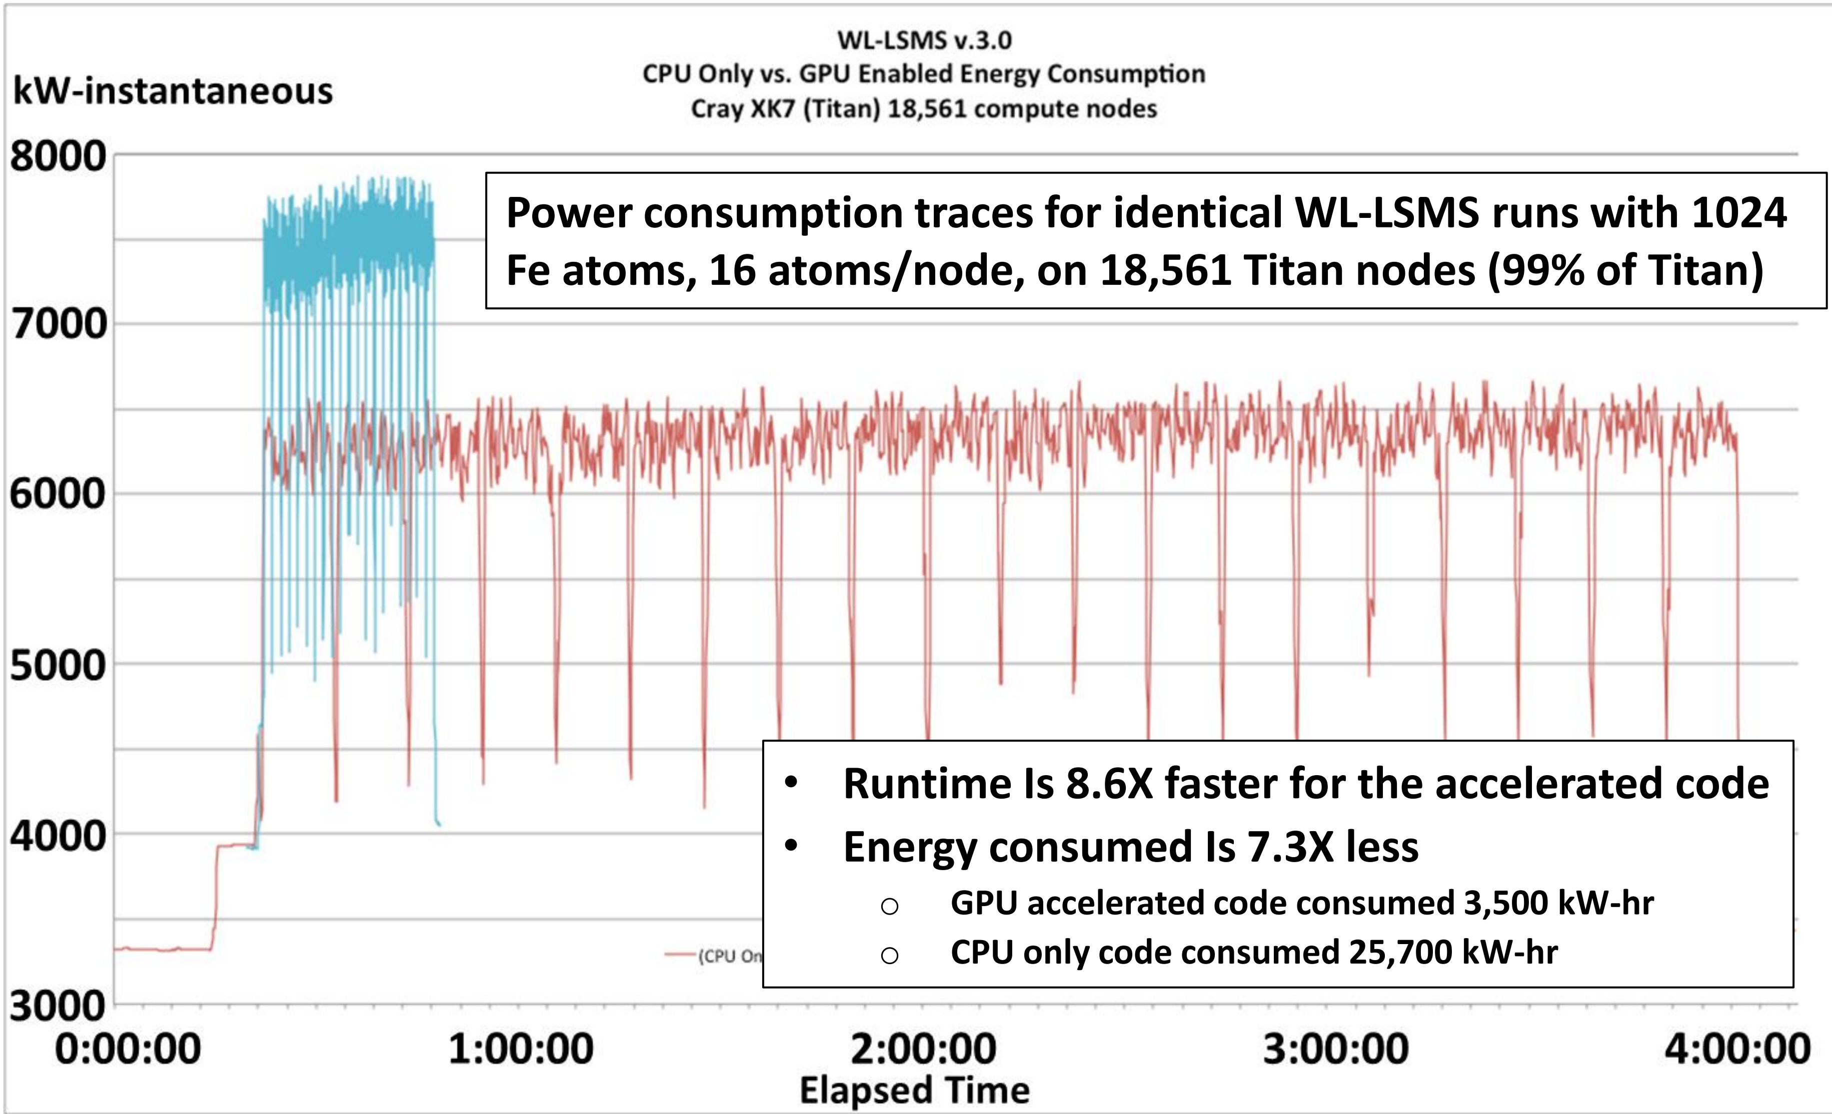

- Runtime Is 8.6X faster for the accelerated code
- Energy consumed Is 7.3X less
  - GPU accelerated code consumed 3,500 kW-hr
  - CPU only code consumed 25,700 kW-hr

CAAR

Application Readiness Partnership Projects

CAAR Projects Overview

- Call for Proposals for eight Partnership Projects
- Partnership between
  - Application Development Team
  - OLCF Scientific Computing group
  - IBM/NVIDIA Center of Excellence
- Application Readiness phase for restructuring and optimization
- Early Science phase for grand-challenge scientific campaign
- OLCF Postdoctoral Associate per project
- Extensive training on hardware and software development environment
- Portability is critical concern
- Coordination with NERSC NESAP and ALCF ESP programs
- Allocations on Titan, early delivery systems and Summit
- Allocations on NERSC and ALCF

CAAR Partnership Responsibilities

- *Develop and execute a* ***Technical plan*** for application porting and performance improvement, developed and executed with reviewable milestones
- *Develop and work according to a* ***Management plan*** with clear description of responsibilities of the CAAR team
- *Develop and execute an* ***Early Science project*** for compelling scientific grand-challenge campaign
- *Assign an* ***Application Scientist*** to carry out the Early Science campaign together with the CAAR team
- *Provide* ***Documentation*** for semi-annual reviews of achieved milestones, and intermediate and final reports

CAAR Partnership Resources

- The core development team of the application, with a stated level of effort
- An ORNL Scientific Computing staff member, who will partner with the core application development team
- A full-time postdoctoral fellow, located and mentored at the OLCF
- Technical support from the IBM/NVidia Center of Excellence
- Allocation of resources on Titan
- Allocation of resources at ALCF & NERSC
- Access to early delivery systems and the *Summit*
- Allocation of compute resources on the full *Summit* system for the Early Science campaign

CAAR Partnership Activities

1. **Common training of all Application Readiness teams**
   a. Architecture and performance portability
   b. Avoidance of duplicate efforts

2. **Application Readiness Technical Plan Development and Execution**
   a. ***Code analysis & benchmarking*** to understand application characteristics: code structure, code suitability for architecture port, algorithm structure, data structures and data movement patterns, code execution characteristics ("hot spots" or "flat" execution profile)
   b. ***Develop parallelization and optimization approach*** to determine the algorithms and code components to port, how to map algorithmic parallelism to architectural features, how to manage data locality and motion
   c. ***Decide on programming model*** such as compiler directives, libraries, explicit coding models
   d. ***Execute technical plan***– benchmarking, code rewrite or refactor, porting and testing, managing portability, managing inclusion in main code repository

3. **Development and Execution of and Early Science Project**, i.e., challenging science problem that demonstrates the performance and scientific impact of the developed application port

CAAR Timeline

1. November 2014: Call for CAAR applications
2. **February 20, 2015: CAAR proposal deadline**
3. March 2015: Selection of CAAR application teams
4. March 2015: CAAR application training workshop
5. April 2015: CAAR application teams start
6. June 2016: CAAR project review
7. September 2017: Call for Early Science projects
8. November 2017: Selection Early Science projects
9. December 2017: Early Science projects start
10. June 2019: Early Science project ends

Future Computational Scientists for Energy, Environment and National Security – Training Program

ASCR facilities host Distinguished Postdoctoral Associates programs with the objective of training the next generation of computational scientists

These programs provide:

1. Challenging scientific campaigns in critical science mission areas
2. Experience in using ASCR computing capabilities
3. Training in software development and engineering practices for current and future massively parallel computer architectures

Central to achieving these goals is access to leadership computing resources, availability of computational domain scientists to provide adequate mentoring and guidance, and facilities' association with universities with strong computational and computer science programs.

OLCF Distinguished Postdoctoral Associates Program

At the OLCF, **eight** Distinguished Postdoctoral Associate positions are available immediately for candidates interested in and capable of performing leading-edge computational science research and development.

Priorities include the development of methodologies and their efficient massively parallel implementation on current state-of-the-art accelerated computer architectures, as well as their application on large scientific challenge problems.

The center is specifically looking for candidates with strong computational expertise in the following scientific areas:

Astrophysics, Biophysics, Chemistry, Climate Science, Combustion, Fusion Energy Science, Materials Science, and Nuclear Physics.

For information about the positions or to apply, visit:
https://www.olcf.ornl.gov/summit/olcf-distinguished-postdoctoral-associates-program/

OLCF Scientific Computing

Ramanan Sankaran, Mike Matheson, George Ostrouchov, Duane Rosenberg, Valentine Anantharaj, Bronson Messer, Mark Berrill, Matt Norman, Ed D'Azevedo, Norbert Podhorski, Wayne Joubert, JJ Chai (postdoc, now in CSM), Judy Hill, Mark Fahey, Hai Ah Nam, Jamison Daniel, Dmitry Liakh, Supada Loasooksathit (postdoc), Markus Eisenbach, Arnold Tharrington, Ying Wai Li, Mingyang Chen (postdoc), Peyton Ticknor, Tjerk Straatsma, Dave Pugmire and Jan-Michael Carrillo (postdoc, now at SNS).

CAAR
Proposal Elements

Summit Application Readiness Partnerships

CAAR Proposals - Overview

Proposals should include the following:

1. ***Name and description of the software application (2)***
2. ***Current degree of parallelization and porting plan for the application (6)***
3. ***Description of scientific libraries used (1)***
4. ***Description of the existing software engineering and distribution (1)***
5. ***Plan for inclusion of code modifications (1)***
6. ***Summary of the use of the application (1)***
7. ***Summary of the expected impact (1)***
8. ***Vision for a scientific campaign (1)***
9. ***Composition of the proposing team (1)***

Estimate of number of pages in parentheses

1. Description of the application

Proposals should include the following:

1. ***Name of the application***
2. ***Domain science field(s) for which application is important***
3. ***Number of active developers of the application***
4. ***Description of modules:***
   a) ***Methods implemented***
   b) ***Size of the code***
   c) ***Interdependencies between modules***
   d) ***Modules that will be targeted in Application Readiness project***
   e) ***Description of the languages used (FORTRAN, C, C++)***
5. ***Description of programming style/model***
6. ***List of architectures on which the application currently runs***

Deadline for proposal submittal to caar@ccs.ornl.gov is February 20, 2015

2. Current degree of parallelization and porting plan

Proposals should include the following:

1. ***Description of the parallel programming model***
   a) ***Threading: OpenMP, TBB, …***
   b) ***MPI, MPI+OpenMP, MPI+OpenACC, …***
   c) ***PGAS: GA, UPC, …***
   d) ***DSL***
2. ***Benchmarking data for representative calculations illustrating scalability***
3. ***Description of performance bottlenecks in the code***
4. ***Description of porting and optimization challenges***
5. ***Description of potential porting and optimization strategies***
6. ***Description of architecture and performance portability strategy***

Deadline for proposal submittal to caar@ccs.ornl.gov is February 20, 2015

3. Scientific Libraries

Proposals should include the following:

1. ***Approach in the application to the use of third-party libraries***
2. ***Description of performance critical scientific libraries used in the application***
3. ***Description of implemented methods that do/may exist in available libraries***

Deadline for proposal submittal to caar@ccs.ornl.gov is February 20, 2015

4. Software Engineering and Distribution

Proposals should include the following:

1. ***Version control***
2. ***Software engineering practices***
3. ***Team development practices***
4. ***Testing practices***
5. ***Documentation for users***
6. ***Documentation for developers***
7. ***Software licensing***
8. ***Software distribution mechanism and schedule***
9. ***User training activities and tutorials***

Deadline for proposal submittal to caar@ccs.ornl.gov is February 20, 2015

5. Including Modifications into Repository

Proposals should include the following:

1. ***Description of mechanism to include CAAR modified code into main repository***

Deadline for proposal submittal to caar@ccs.ornl.gov is February 20, 2015

6. Summary of Application Usage

Proposals should include the following:

1. ***Description of the use of the application by the broader scientific community***
2. ***Estimate of the current user base***

Deadline for proposal submittal to caar@ccs.ornl.gov is February 20, 2015

7. Summary of Impact

Proposals should include the following:

1. ***Summary of the expected impact of the CAAR ported version of the application***

Deadline for proposal submittal to caar@ccs.ornl.gov is February 20, 2015

8. Scientific Campaign

Proposals should include the following:

1. ***Vision for a scientific campaign for the Early Science phase of the CAAR project***

Deadline for proposal submittal to caar@ccs.ornl.gov is February 20, 2015

9. Composition Proposing Team

Proposals should include the following:

1. ***For the members of the proposing software development team list:***
   a) ***Name, title, and institution***
   b) ***Level of experience with development of the application***
   c) ***Level of experience with HPC systems use and software development***
   d) ***Level of effort committed to the CAAR partnership***

2. ***Curriculum Vitae from each team member should be appended to the proposal***

Deadline for proposal submittal to caar@ccs.ornl.gov is February 20, 2015

CAAR Selection Criteria

- *Composition, experience and commitment of application development team*
- *Assessment of anticipated scientific impact of ported application*
- *Assessment of porting feasibility based on provided porting plan and benchmarks*
- *Application user base*
- *Compelling vision of an Early Science project*

- *Technical domain expertise of OLCF liaison*
- *Technical expertise of IBM/NVIDIA Center of Excellence*
- *Coverage of science domains in CAAR portfolio and support for DOE and US mission*
- *Coverage of algorithms, programming approaches, languages, data models*

- *Consultation with NERSC and ALCF*
- *Consultation with DOE Office of Advanced Scientific Computing Research*

Questions & Discussion