OLCF Year in Review

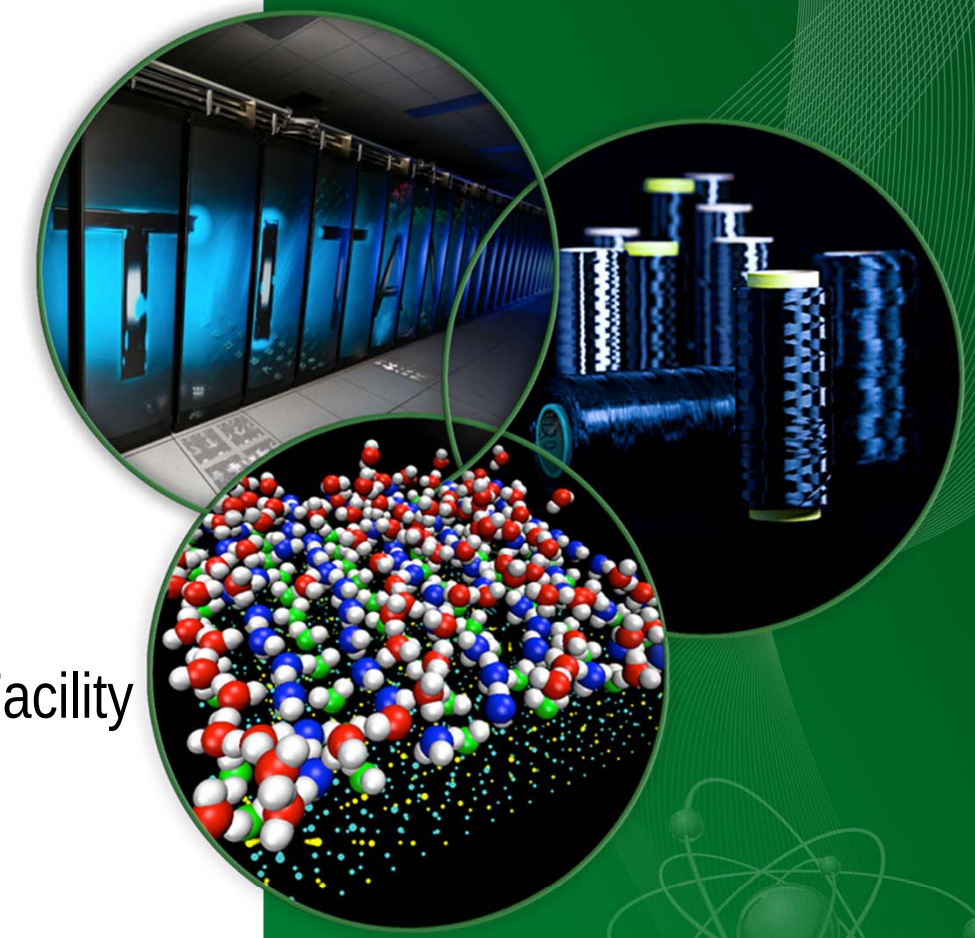

Jack Wells
Director of Science
Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory

2014 OLCF User Group Meeting
Oak Ridge
22 July 2014

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

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

ORNL has increased system performance by 1,000 times 2004-2010

Hardware scaled from single-core through dual-core to quad-core and dual-socket , 12-core SMP nodes	Scaling applications and system software was the biggest challenge

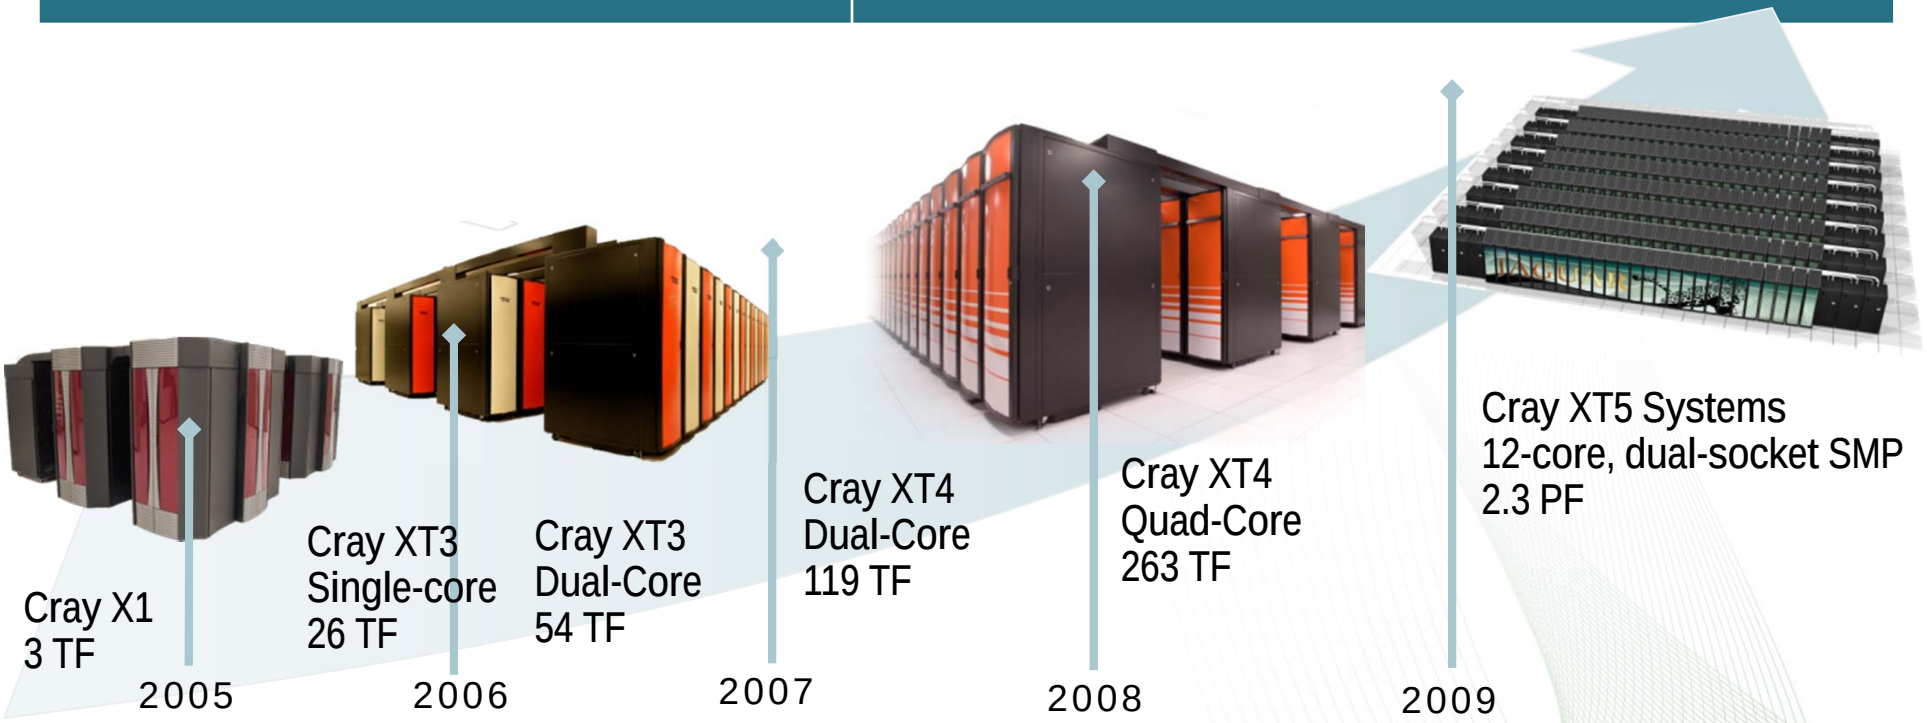

Our Science requires that we advance computational capability 1000x over the next decade.

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

Roadmap to Exascale

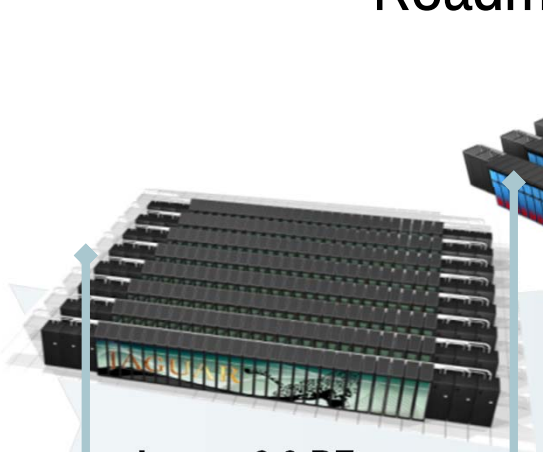

Jaguar 2.3 PF
362 TB DRAM

2010

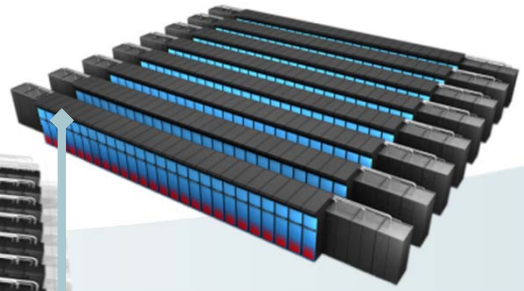

Titan 27 PF
600 TB DRAM
Hybrid GPU/CPU

2012

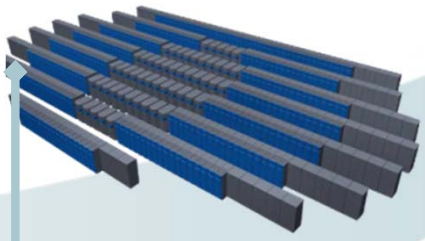

OLCF-4: 100-250 PF
4000 TB memory
> 20MW
6 day resilience

2017

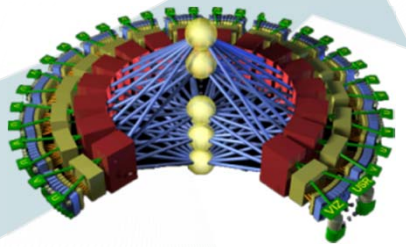

OLCF-5: 1 EF
20 MW

What are the Challenges?

2022

Science breakthroughs at the LCF:

SELECTED science and engineering advances over the period 2003 - 2013

Astrophysicists discover supernova shock-wave instability, *Astrophys. J. (2003)*
268 citations, 7/2014

Researchers solved the 2D Hubbard model and presented evidence that it predicts HTSC behavior, *Phys. Rev. Lett (2005)*
111 citations, 7/2014

First-Principles Flame Simulation Provides Crucial Information to Guide Design of Fuel-Efficient Clean Engines, *Proc. Combust. Insti.* (2007)
78 citations, 3/2014

Largest simulation of a galaxy's worth of dark matter, showed for the first time the fractal-like appearance of dark matter substructures, *Nature* (2008)
347 citations, 7/2014

World's first continuous simulation of 21,000 years of Earth's climate history, *Science (2009)*
133 citations, 7/2014

Biomass as a viable, sustainable feedstock for hydrogen production for fuel cells, *Nano Letters (2011)*
JACS. (2010)
92 & 82 citations, respectively, 7/2014

Demonstrated that three-body forces are necessary to describe the long lifetime of ^{14}C
Phys. Rev. Lett. (2011)
28 citations, 3/2014

Global Warming preceded by increasing carbon dioxide concentrations during the last deglaciation, *Nature (2012).*
95 citations, 7/2014

Calculation of the number of bound nuclei in nature, *Nature (2012), 36 citations, 3/2014* , 46 citations, 7/2014

MD simulations show selectivity filter of a trans-membrane ion channel is sterically locked open by hidden water molecules, *Nature (2013)*

SASI Model Discovered at OLCF Confirmed Observationally a Decade Later in Supernova Remnant

- In 2003, 3D simulations on OLCF's Phoenix supercomputer first predicted the standing accretion shock instability (SASI) in core-collapse SN.
- SASI is a sloshing of stellar material that destabilizes a shock wave pushing material outward during the death o a massive star, helping lead to a supernova explosion that releases all the elements found in the universe.
- In 2013, researchers mapping radiation signatures from the Cassiopeia A supernova with NASA's NuSTAR high-energy x-ray telescope array published evidence that supports the SASI model based on the distribution of Titanium-44 in Cas A.
- This is a example of simulation predicting a physical phenomenon before it is observed experimentally.
- Tony Mezzacappa of the University of Tennessee–Knoxville (UT), an author of the 2003 Phoenix paper, and 2014 INCITE project principal investigator Eric Lentz of UT are using 85 million core hours to simulate a more detailed supernova explosion on Titan. They expect to learn more about the SASI's role in generating explosion.

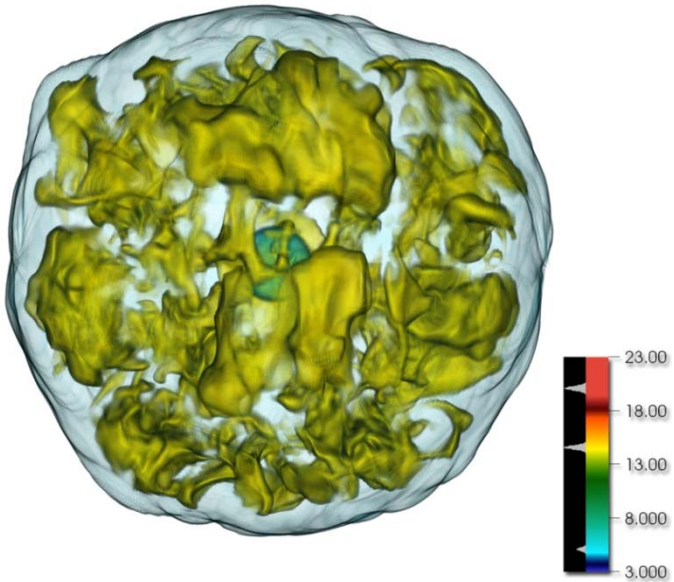

The entropy of the inner 250 kilometers of a 15 solar-mass star during a 3D simulation of a core-collapse supernova using the CHIMERA code. Large-scale distortion of the supernova shock can be seen, along with smaller-scale convection. Visualization by Mike Matheson.

High-impact science across a broad range of disciplines

For example in 2013:

Paleoclimate Science
"Northern Hemisphere forcing of Southern Hemisphere climate during the last deglaciation,"
Feng He (UW Madison), *et al*., *Nature*, February (2013)

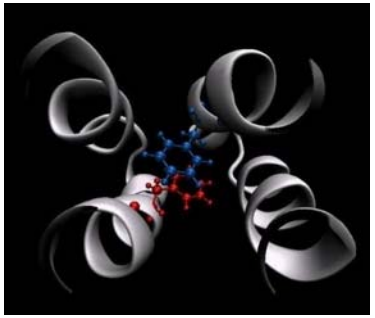

Molecular Biology
"A phenylalanine rotameric switch for signal-state control in bacterial chemoreceptors"
D. Ortega (UTK),
Nature Communications
December (2013)

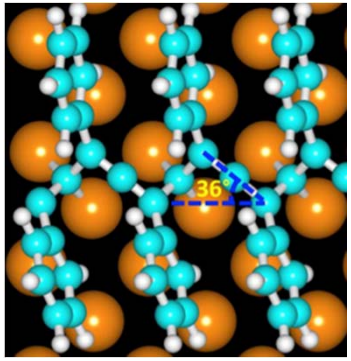

Polymer Science
"Self-Organized and Cu-Coordinated Surface Linear Polymerization"
Qing Li, B. Sumpter (ORNL),
Nature Scientific Reports.
July (2013)

Molecular Biology
"Recovery from slow inactivation in K1 channels is controlled by water molecules"
Jared Ostmeyer, et al. (U. Chicago) *Nature*, Sept. (2013)

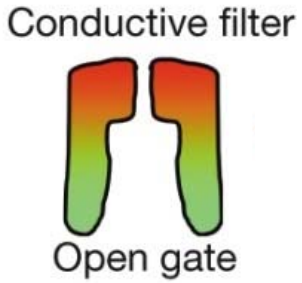

Superconductivity
"Doping dependence of spin excitations and correlations with high-temperature super-conductivity in iron pnictides,"
Meng Wang(IOP CAS Beijing),
Nature Communications.
December (2013)

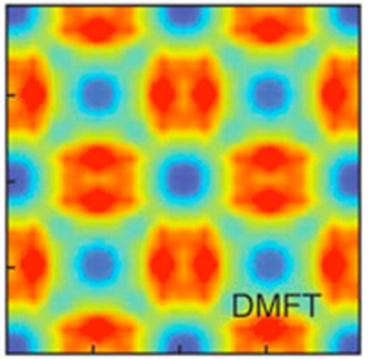

Complex Oxide Materials
"Atomically resolved spectroscopic studyof Sr2IrO4: Experiment and theory," Qing Li (ORNL), E.G. Eguiluz (UTK)
Nature Scientific Reports.
October (2013)

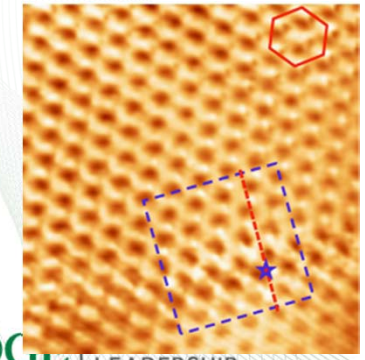

What is the Leadership Computing Facility (LCF)?

- Collaborative DOE Office of Science program at ORNL and ANL
- Mission: Provide the computational and data resources required to solve the most challenging problems.
- 2-centers/2-architectures to address diverse and growing computational needs of the scientific community
- Highly competitive user allocation programs (INCITE, ALCC).
- Projects receive 10x to 100x more resource than at other generally available centers.
- LCF centers partner with users to enable science & engineering breakthroughs (Liaisons, Catalysts).

OAK RIDGE National Laboratory

Argonne NATIONAL LABORATORY

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

Three primary ways for access to LCF
Distribution of allocable hours

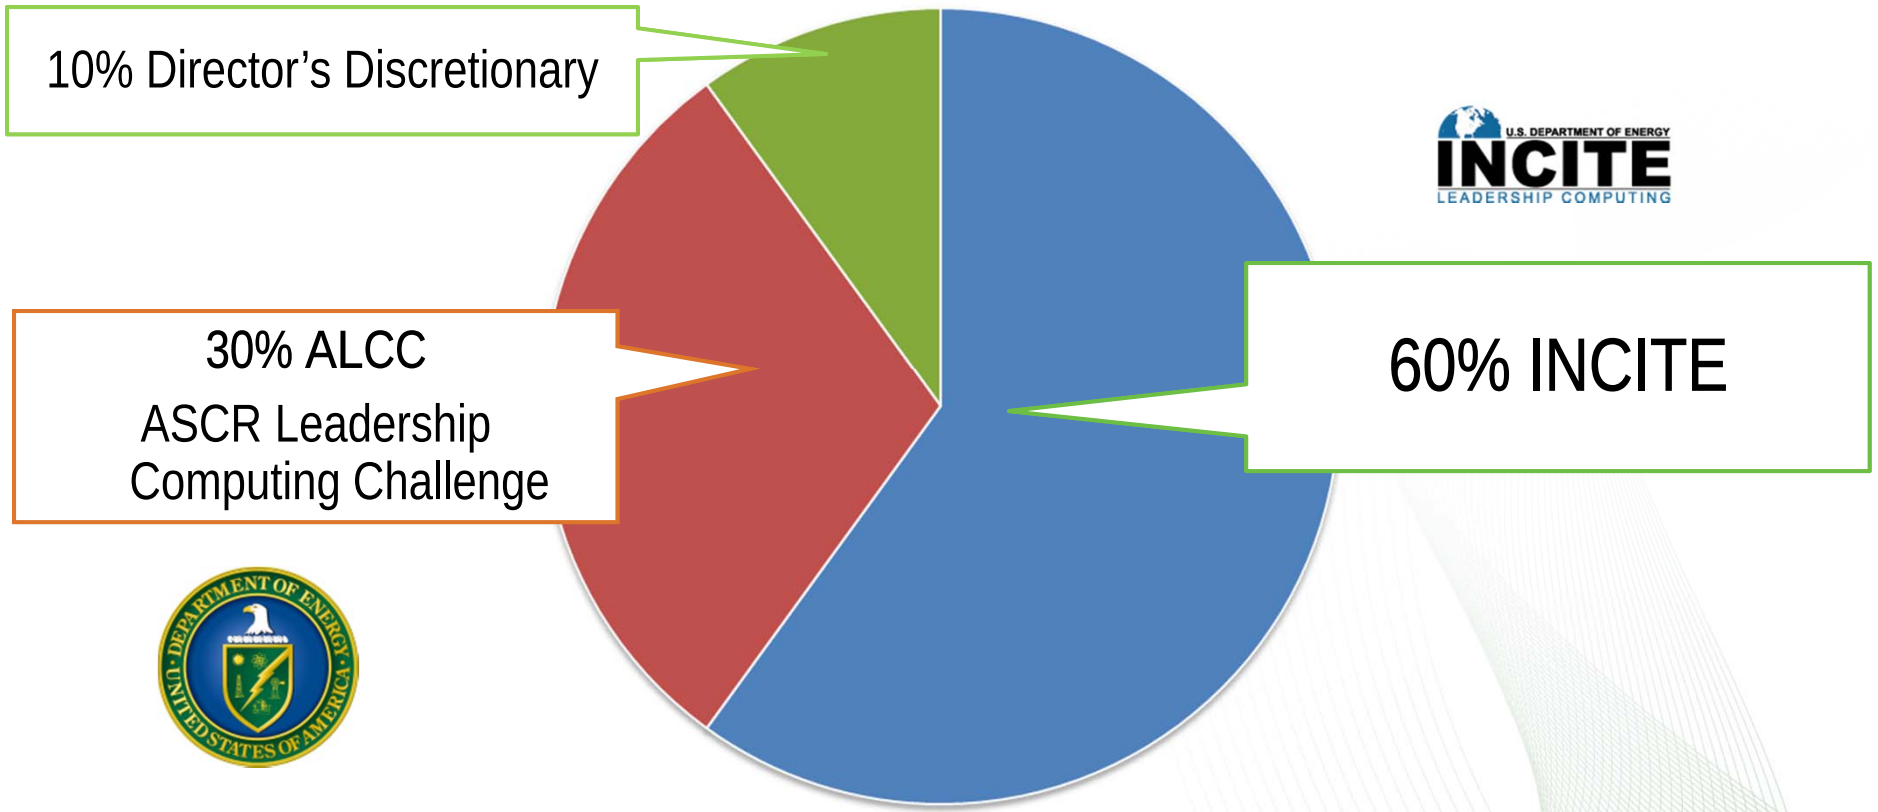

Demand for INCITE resources outstrips supply with more time requested than available – Number of projects remains flat

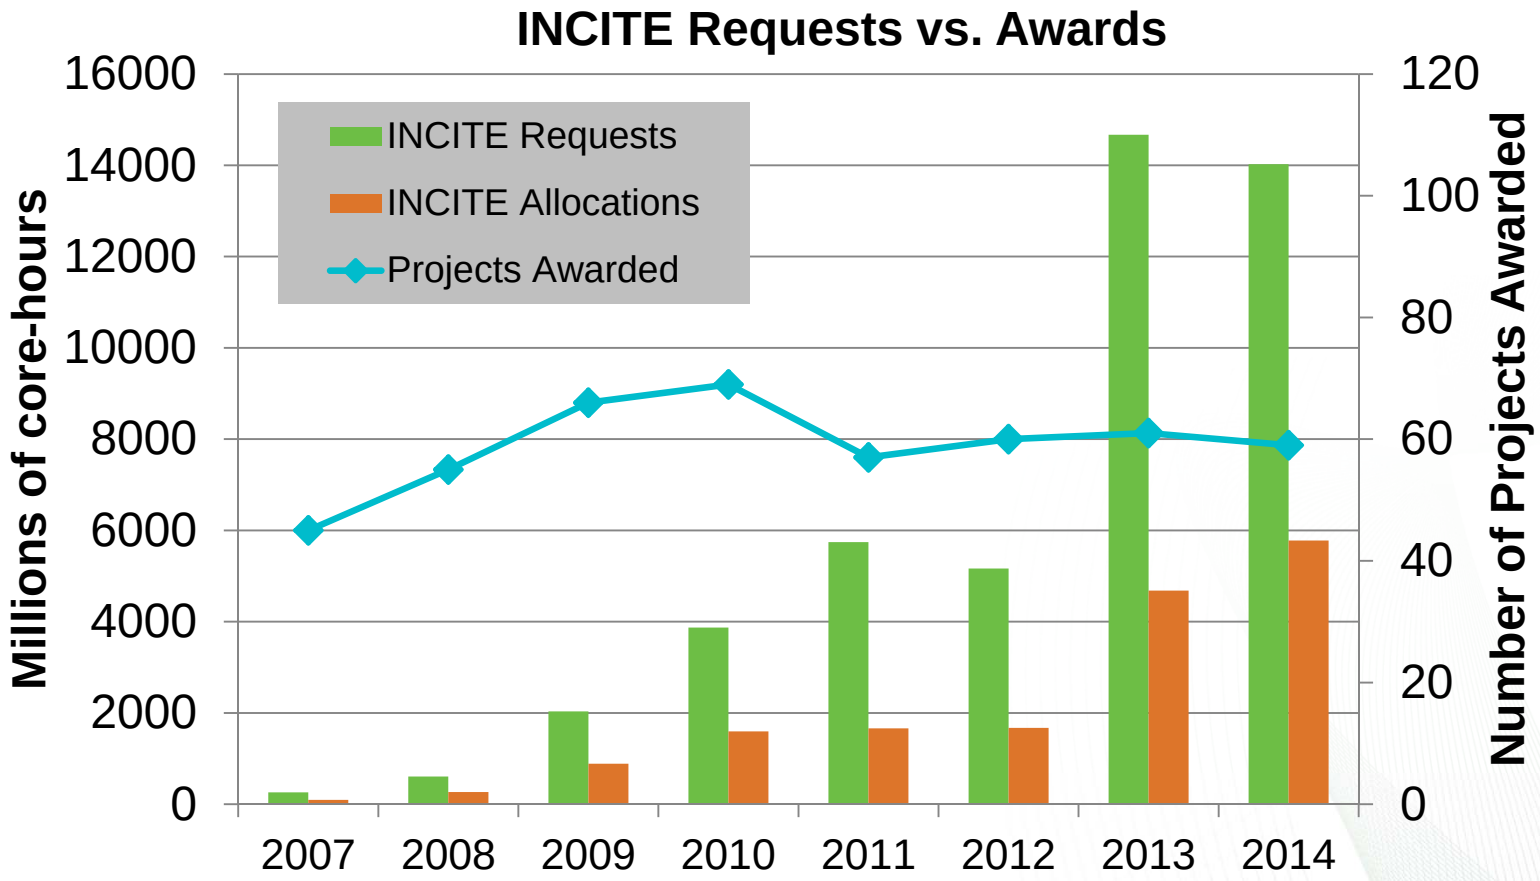

	Titan System (Cray XK7)		
Peak Performance	27.1 PF 18,688 compute nodes	24.5 PF GPU	2.6 PF CPU
System memory	710 TB total memory		
Interconnect	Gemini High Speed Interconnect	3D Torus	
Storage	Luster Filesystem	32 PB	
Archive	High-Performance Storage System (HPSS)	29 PB	
I/O Nodes	512 Service and I/O nodes		

High-impact science at OLCF: Four of Six SC13 Gordon Bell Finalists Used Titan

Peter Staar **ETH Zurich**	**Massimo Bernaschi** **ICNR-IAC Rome**	**Michael Bussmann** **HZDR - Dresden**	**Salman Habib** **Argonne**
High-Temperature Superconductivity	**Biofluidic Systems**	**Plasma Physics**	**Cosmology**
Taking a Quantum Leap in Time to Solution for Simulations of High-T_C Superconductors	20 Petaflops Simulation of Protein Suspensions in Crowding Conditions	Radiative Signatures of the Relativistic Kelvin-Helmholtz Instability	HACC: Extreme Scaling and Performance Across Diverse Architectures
Titan (15.4 PF)	Titan (20 PF)	Titan (7.2 PF)	Titan (10 PF)
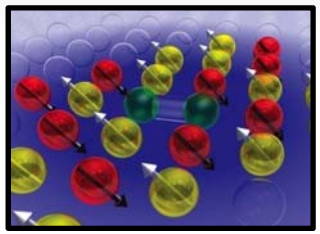	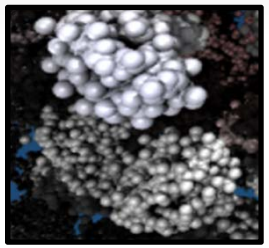	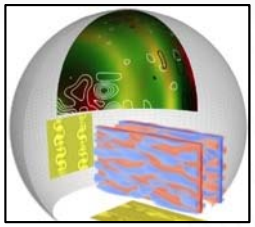	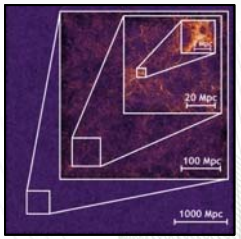

Titan Daily Node Failures:

30-Day Moving Average

Operational Highlights:
Delivered commitments to major programs

- Delivered 1.8 Billion hour commitment to the 2013 INCITE Program
- Delivered 400 Million hour commitment to 2013-2014 ALCC Program
- Exceeded OLCF's leadership computing metric and all other operational goals

Strong Capability Usage on Titan

Percentage of INCITE Time running leadership jobs

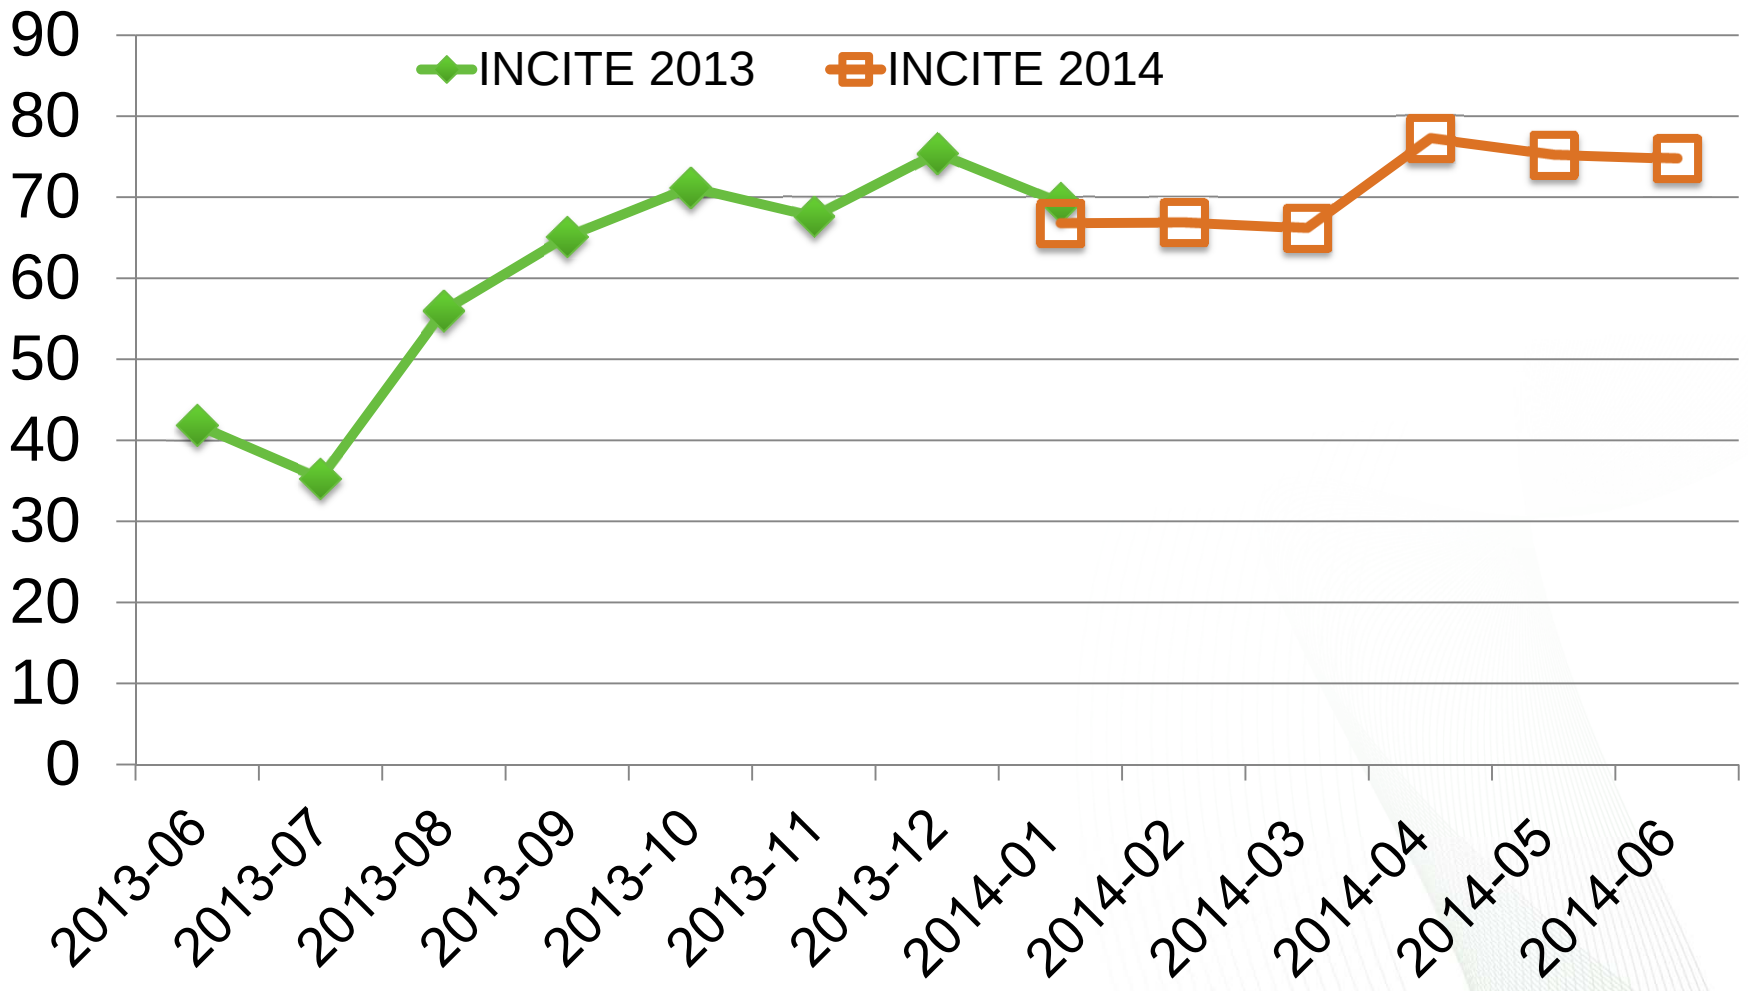

Leadership job uses 3,750 nodes or more of Titan's compute nodes, 20% of resource

Increasing Requests for GPUs on Titan

Percentage of INCITE Time requesting GPUs

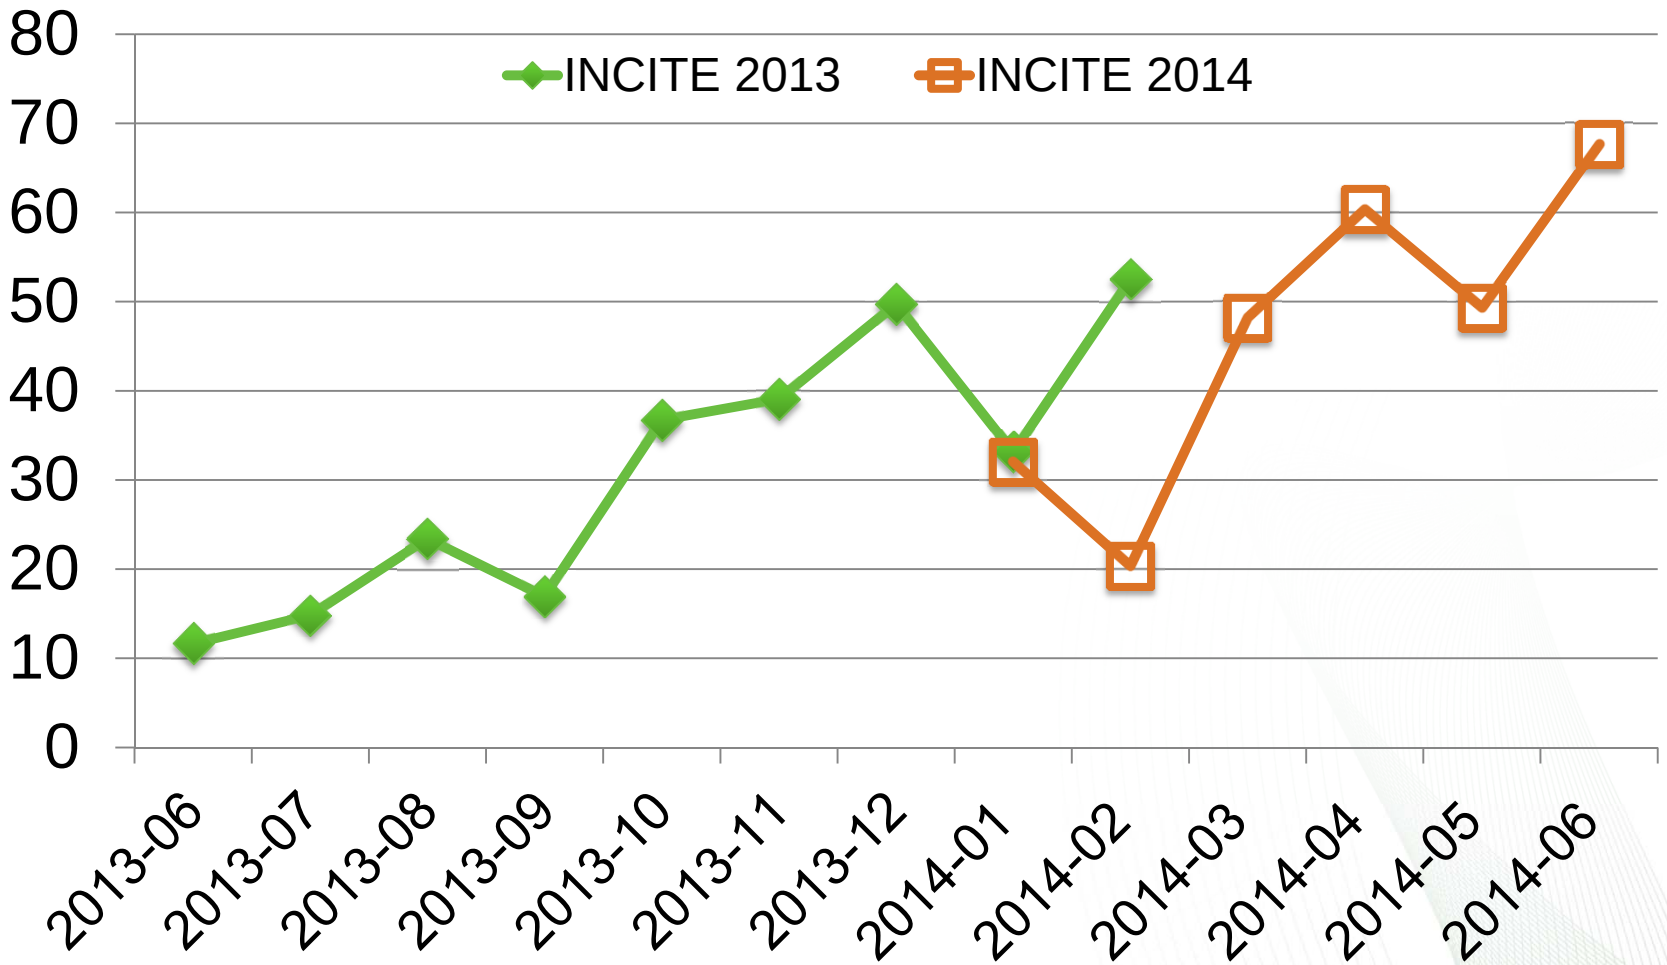

Early Science Challenges for Titan

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

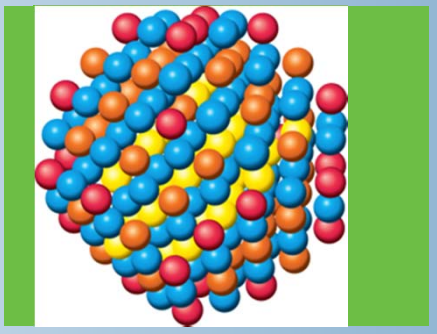

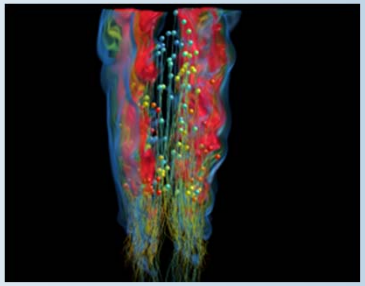

S3D
Understanding turbulent combustion through direct numerical simulation with complex chemistry.
.

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

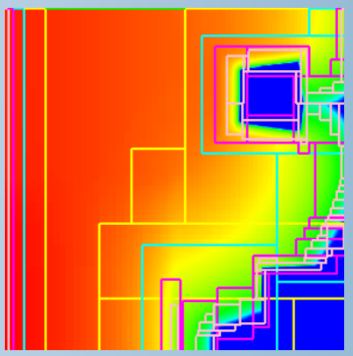

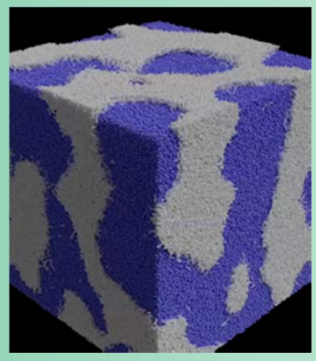

LAMMPS
A molecular dynamics simulation of organic polymers for applications in organic photovoltaic heterojunctions , de-wetting phenomena and biosensor applications

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

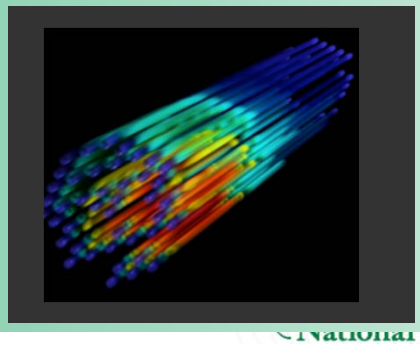

Denovo
Discrete ordinates radiation transport calculations that can be used in a variety of nuclear energy and technology applications.

Effectiveness of GPU Acceleration?

OLCF-3 Early Science Codes -- Performance on Titan XK7

Application	Cray XK7 vs. Cray XE6 Performance Ratio*
LAMMPS* Molecular dynamics	**7.4**
S3D Turbulent combustion	**2.2**
Denovo 3D neutron transport for nuclear reactors	**3.8**
WL-LSMS Statistical mechanics of magnetic materials	**3.8**

Titan: Cray XK7 (Kepler GPU plus AMD 16-core Opteron CPU)
Cray XE6: (2x AMD 16-core Opteron CPUs)
*Performance depends strongly on specific problem size chosen

Additional Applications from Community Efforts

Current Performance Measurements on Titan

Application	Cray XK7 vs. Cray XE6 Performance Ratio*
AWP-ODC Seismology	**2.1**
DCA++ Condensed Matter Physics	**4.4**
QMCPACK Electronic structure	**2.0**
RMG (DFT – real-space, multigrid) Electronic Structure	**2.0**
XGC1 Plasma Physics for Fusion Energy R&D	**2.0**

Titan: Cray XK7 (Kepler GPU plus AMD 16-core Opteron CPU)
Cray XE6: (2x AMD 16-core Opteron CPUs)

OLCF Training & Workshops

TOPIC	
Lattice QCD Computational Science Workshop	April 2013
INCITE Proposal Writing Webinars	2013 (2), 2014 (2)
Crash Course in Supercomputing	June 2013, June 2014
Introduction to HPC using GPUs	June 2013
OpenACC Tutorials	July 2013
Processing and Analysis of Very Large Data Sets	August 2013
ADIOS Code Sprint	August 2013
Compiler Optimization Webinar	January 2014
Getting Started at OLCF	2013, 2014
Allinea DDT Training	April 2014
Compiler Directives Lunch Webinars	Many
Paraview Workshop	May 2014
Joint Facilities Forum on Data-Intensive Computing	June 2014

Boosting Bioenergy and Overcoming Recalcitrance

Molecular Dynamics Simulations

2013 INCITE Program
PI: Jeremy Smith
University of Tennessee & ORNL
Used Hours: 72,167,000

Science Objectives and Impact

- Improve our understanding of lignin-cellulose interactions at the molecular level in order to overcome biomass recalcitrance and improve the efficiency of biofuel production, thereby reducing the cost of ethanol
- Ethanol, which is carbon-neutral and domestically produced, is the primary renewable substitute for gasoline

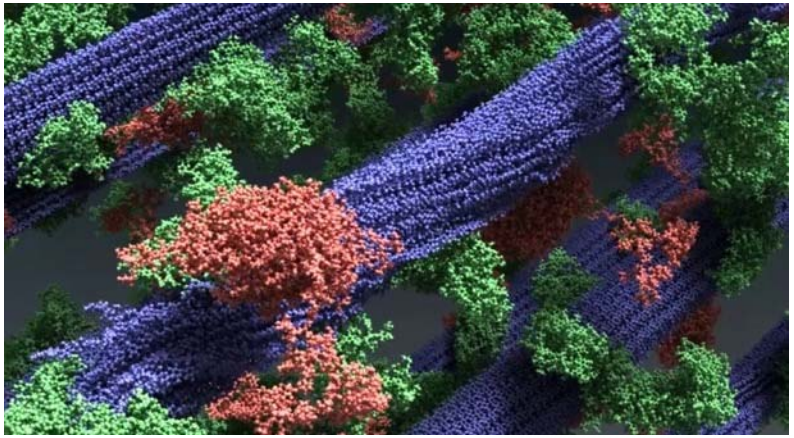

A cellulase enzyme (pink) hydrolyzing a cellulose (blue) strand despite the presence of lignin aggregates (green) on the cellulose surface. Vizualization by M. Matheson (ORNL)

Performance & OLCF Contribution

- GROMACS application has been adapted to run on the GPU-accelerated Titan system: Project can handle much larger systems—30 million atoms, compared to 3 million atoms on Jaguar
- The OLCF's Mike Matheson provided visualization services

Science Results

- Discovered amorphous cellulose is easier to break down because it associates less with lignin
- This is because the less-organized cellulose interacts strongly with water, making it less available to interact with lignin
- Biofuel production may potentially become more efficient by manipulating cellulose crystallinity so as to reduce lignin precipitation onto cellulose

The Bleeding 'Edge' of Fusion

2013 INCITE Program
PI: CS Chang
Princeton Plasma Physics Laboratory
Hours used: 102 million

Science Objectives and Impact

- Efficiency of ITER is expected to be determined by edge plasma transport
- First-principles simulation of edge physics on Titan can increase understanding of ITER
- Predictions validated and new insight through analysis of DIII-D results
- Success with DIII-D gives confidence for larger JET- and ITER-scale campaigns in 2014

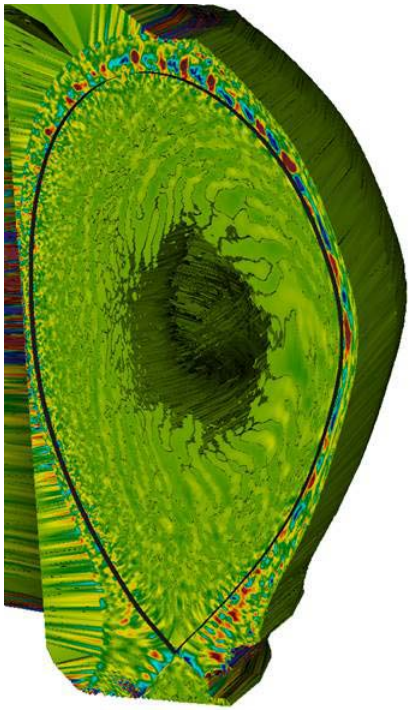

This figure shows the turbulence front from the plasma edge in the DIII-D reactor being spread inward in multiscale interaction with the evolving background profile under the central heat source. Eventually, the whole volume becomes turbulent, with the spatial turbulence amplitude distribution being just enough to produce the outward heat transport to expel the centrally deposited heat to the edge.
- Visualization by David Pugmire ORNL)

Performance & OLCF Contribution

- OLCF provided the necessary HPC power; Up to 90% of Titan's max capability has been utilized for the study
- OLCF liaison contributed to algorithmic improvement, which made the XGC1 code use GPUs and CPUs efficiently, with linear scalability to maximal size of Titan
- OLCF personnel provided the Vis support

Science Results

- Simulated nonlinear coherent turbulence structures (called "blobs") in the plasma edge of DIII-D reactor
- Identified the momentum source and its inward transport process
- Predicted the divertor heat load distribution

All for the first time in first-principles calculations of a tokamak reactor

Non-Icing Surfaces for Cold Climate Wind Turbines

Molecular Dynamics Simulations

ALCC Program
M. Yamada. GE Global Research
Allocated Hours: 40 Million
Used Hours: 38 Million

Science Objectives and Impact

- Understand microscopic mechanism of water droplets freezing on surfaces
- Determine efficacy of non-icing surfaces at different operation temperatures

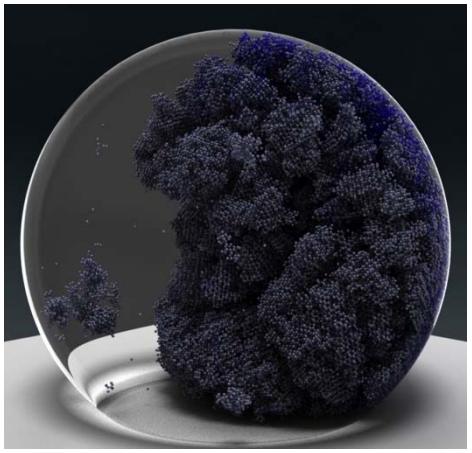

Location of ice nucleation varies dependent on temperature and contact angles. Visualization by M. Matheson (ORNL)

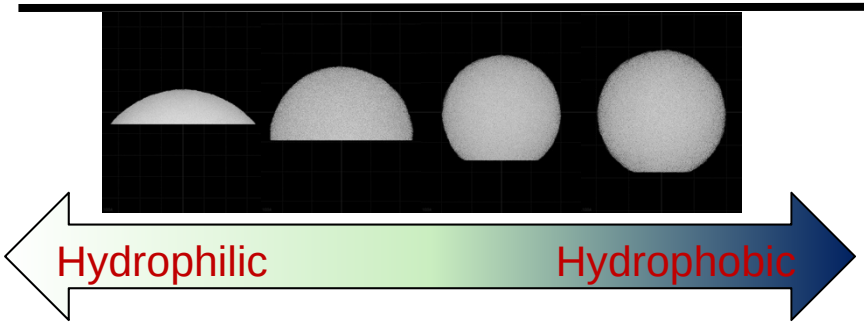

Performance Achievements

- 5X speed-up from GPU acceleration
- Achieved factor 40X speed-up from new interaction potential for water

Science Results

Replicated GE's experimental results:

- High-contact-angle surfaces delay the onset of nucleation
- The delay becomes less pronounced at lower temperatures

Competitiveness Impact

- Shift from energy intensive active ice-mitigation strategies (which can require up to 25% of the total rated output) to passive non-icing surfaces

Westinghouse-CASL Simulation of Reactor Start Up – Quarter-Core Zero Power Physics Test

Goals

- Compare fidelity and performance of Shift against Keno, SP_N, and S_N (Denovo)
- Generate high-fidelity neutronics solution for code comparison of solutions for predicting reactor startup and physics testing

Execution

- AP1000 model created and results generated for reactor criticality, rod worth, and reactivity coefficients
- Identical VERA Input models used for Shift, SP_N, and S_N
  - dramatically simpler than KENO-VI input model

Results

- Some of the largest Monte Carlo calculations ever performed (1 trillion particles) have been completed
  - runs use 230,000 cores of Titan or more
- Excellent agreement with KENO-VI
- Extremely fine-mesh S_N calculations, which leverage Titan's GPU accelerators, are under way

AP1000 pin powers

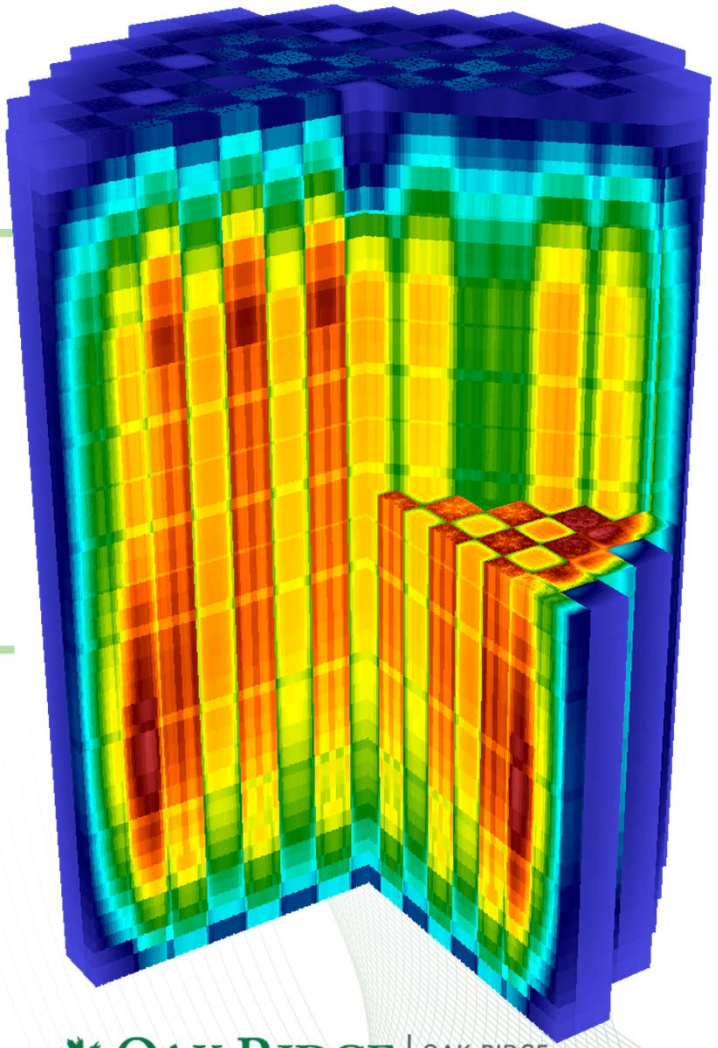

Number of New Industry Projects Launched at OLCF by Calendar Year

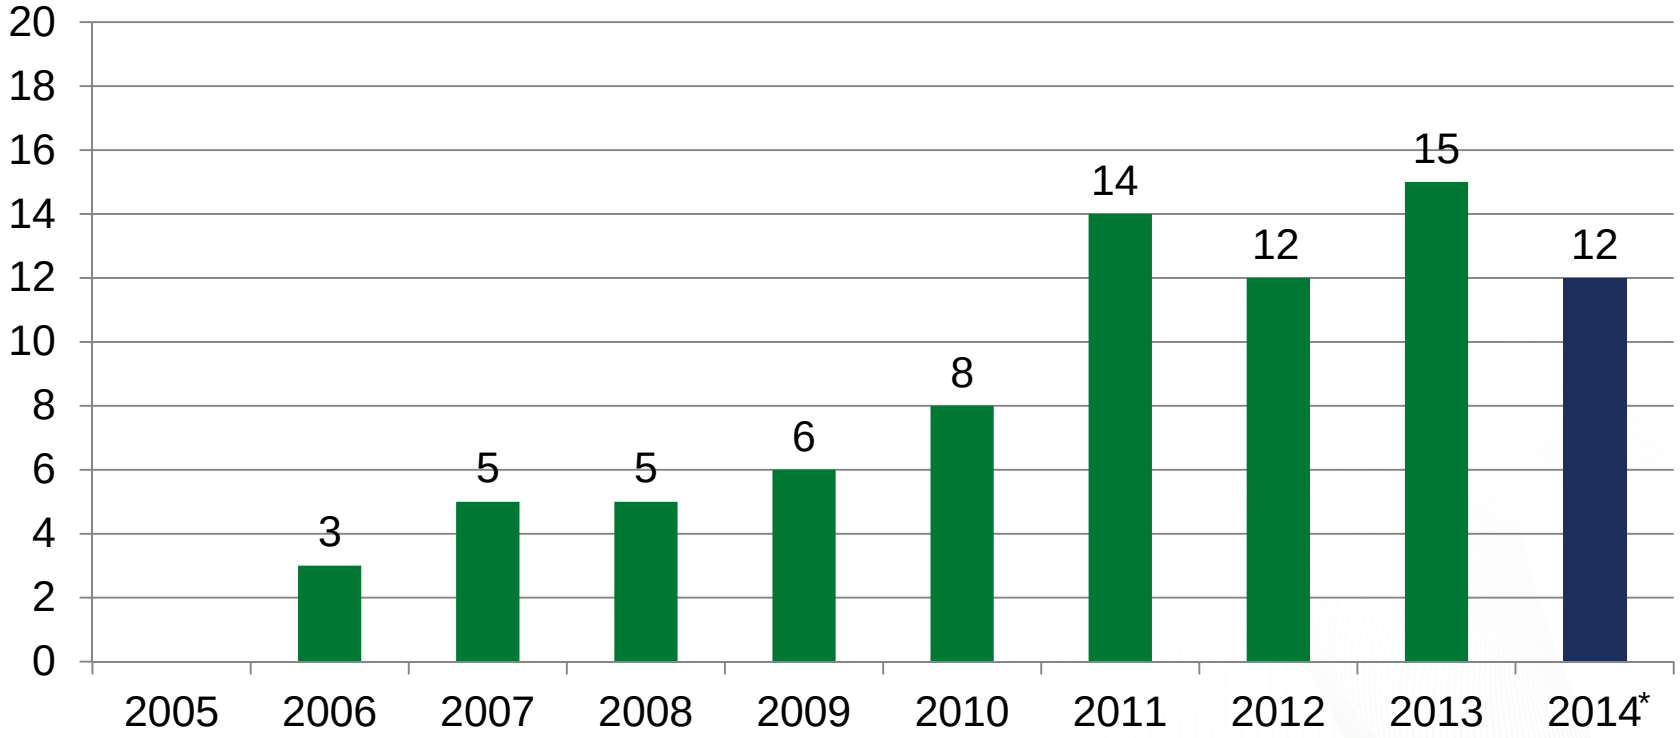

Current as of June 2014

* Year in progress

ACCEL
Accelerating Competitiveness through Computational ExcelLence

OAK RIDGE National Laboratory | OAK RIDGE LEADERSHIP COMPUTING FACILITY

PanDA Tool Provides Titan with High-Tech Scheduler

- Researchers with the ATLAS experiment in Europe have been integrating its scheduling and analysis tool, PanDa, with Titan.
- Workflow includes 1.8 million computing jobs each day distributed among 100 or so computing centers spread across the globe.
- PanDA's ability to efficiently match available computing time with high-priority tasks holds great promise for Titan.
- Team developers redesigned parts of the PanDA system on Titan responsible for job submission on remote sites and gave PanDA new capability to collect information about unused worker nodes on Titan.
- Deployment of the tool could lead to a higher utilization of available hours on Titan.

PanDA manages all of ATLAS's data tasks from a server located at CERN, the European Organization for Nuclear Research.

ADIOS wins big

OLCF project named one of top 100 tech products

- The OLCF's Scott Klasky received an R&D 100 Award for his development of the Adaptable I/O System for Big Data, or ADIOS.
- Developed by ORNL, Georgia Institute of Technology, Rutgers University, and North Carolina State University, ADIOS is a portable, scalable, easy-to-use software framework conceived to solve "big data" problems.
- Klasky began developing ADIOS in 2008, and now the middleware assists numerous applications on Titan.
- ADIOS has demonstrated impressive I/O performance results on leadership class machines and clusters, sometimes showing an improvement of more than 1,000 times over well-known parallel file formats.
- The research was funded by DOE's Oak Ridge Leadership Computing Facility, the office of Advanced Scientific Computing Research, the Office of Fusion Energy Science, and the NSF.

"ADIOS represents a grassroots effort to allow data–intensive computations and experiments to scale to current and future-generation computing resources. This award is the result of years of hard work from our team members and belief from our sponsors." – OLCF Staff Member and R&D 100 Award Winner Scott Klasky

Over Arching Questions for the scientific agenda for this OLCF user meeting

(0) What are your achievements on Titan?

(1) What problem do you want to solve in 2018 on a computer that is 5X to 10X more capable than Titan?

- What do we need to do to get ready for this opportunity?

(2) How (effectively) are you integrating compute and data requirements at OLCF to achieve breakthrough science?

- How can we integrate data and compute more effectively?

Acknowledgements

OLCF-3 Vendor Partners: Cray, AMD, NVIDIA, CAPS, Allinea

OLCF Users: CS Chang (PPPL),
Jeremy Smith(UT/ORNL),
Tom Evans/John Turner (ORNL),
Masako Yamanda (GE Global Research)

OLCF Staff: Mike Matheson (ORNL)
Dave Pugmire (ORNL) for visualizations

DOE ASCR Leadership Computing Program, ALCF, OLCF

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?

WellsJC@ornl.gov

Contact us at
http://olcf.ornl.gov
http://jobs.ornl.gov

OAK RIDGE National Laboratory