OLCF Overview

20 Years of Excellence in Computational Science

OLCF

Oak Ridge Leadership Computing Facility

1992–2012

Presented by:

Jack C. Wells

Director of Science

Oak Ridge Leadership Computing Facility (OLCF)

National Center for Computational Sciences (NCCS)

OLCF East Coast Training

February 19, 2013

Knoxville

ORNL's mission

Deliver scientific discoveries and technical breakthroughs that will accelerate the development and deployment of solutions in clean energy and global security, and in doing so create economic opportunity for the nation

Science Strategy for the Future: Major Laboratory Initiatives

Deliver science using neutrons

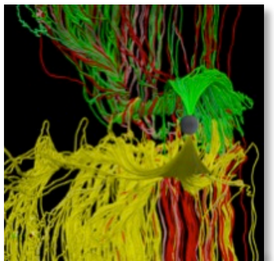

Scale computing, data infrastructure, and analytics for science

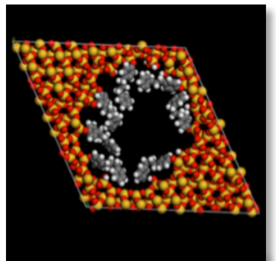

Discover and demonstrate advanced materials for energy

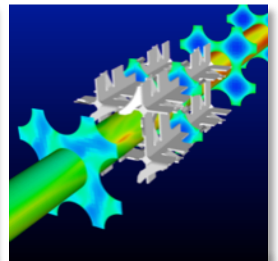

Advance scientific basis for new nuclear technologies and systems

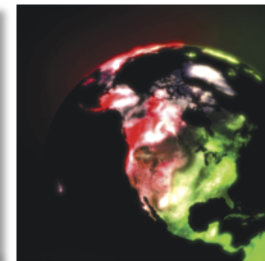

Advance understanding in biological, environmental systems, and climate change impacts science

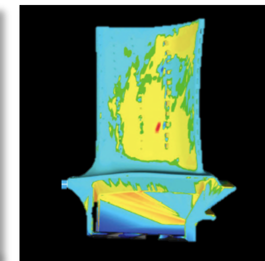

Enhance building energy efficiency, sustainable transportation, and advanced manufacturing

Solve the nation's most compelling global security challenges

Scale computing, data infrastructure, and analytics for science

Vision

- World's most capable complex for computational science: Infrastructure, staff, multiagency programs
- **Outcome:** Sustained world leadership in transformational research and scientific discovery through advanced computing

Strategy

Provide the nation's most capable computational site for advancing to the exascale and beyond	Leverage partnerships to lead U.S. development of exascale computing technology	Attract top talent; deliver outstanding user program; educate and train next generation of researchers

Big Problems Require Big Solutions

Climate Change

Energy

Healthcare

Competitiveness

High-impact science at OLCF across a broad range of disciplines.

For example in 2012:

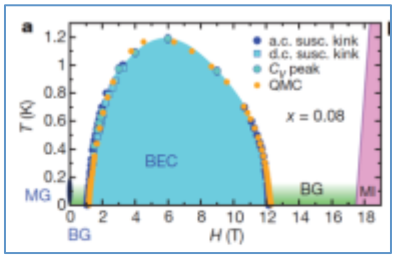

Materials: Quantum Magnets
"Bose glass and Mott glass of quasiparticles in a doped quantum magnet"
Rong Yu (Rice U.)
Nature (2012)

Nuclear Physics
"The Limits of the Nuclear Landscape"
J. Erier, (UT/ORNL)
Nature (2012)

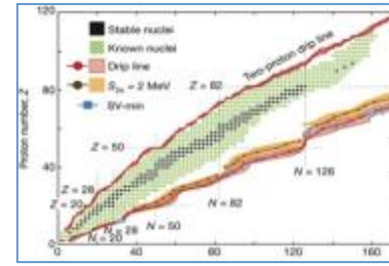

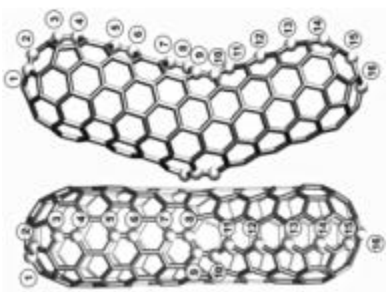

Carbon Nanomaterials
"Covalently bonded three-dimensional carbon nanotube solids via boron induced nanojunctions"
Hashim (Rice), **Scientific Reports** (2012)

Climate Prediction and Mitigation
"Relative outcomes of climate change mitigation related to global temperature versus sea-level rise"
G.A. Meehl (NCAR),
Nature Climate Change (2012)

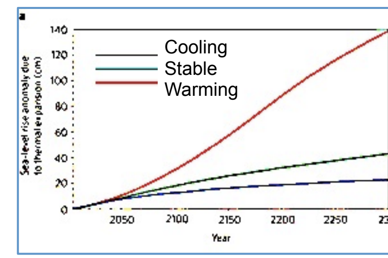

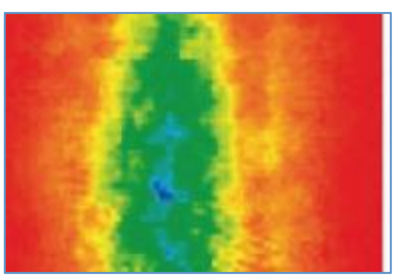

Plasma Physics:
"Dynamics of relativistic transparency and optical shuttering in expanding overdense plasmas"
S. Palaniyappan (LANL)
Nature Physics (2012)

Paleoclimate Climate Change
"Global warming preceded by increasing carbon dioxide concentrations during the last deglaciation"
J. Shakun, (Harvard/Columbia)
Nature (2012)

High-impact climate-change science:

Paleoclimate Change Understanding

- First IPCC-class CGC Model to simulate 22,000 years.
- Reproduces Earth's most recent natural global warming
- Reproduces phasing of the climate at last deglaciation
- Proves that CO_2 caused global warming at Ice Age's end

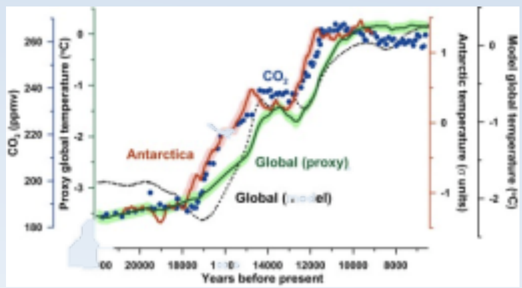

"Transient simulation of last deglaciation with new mechanism for Bølling-Allerød warming"
Zhengyu Liu, (CAS/UW-Madison)
Science (2009)
82 cites

"Global warming preceded by increasing carbon dioxide…."
J. Shakun, (Harvard/ Columbia)
Nature (2012)

Climate Predictions and Mitigation

- Ocean circulation supports hiatus periods of cooling
- Sea level will rise in the future for all mitigation scenarios – shows commitment to rising seas
- Aggressive mitigation measures affect the rate of sea-level rise

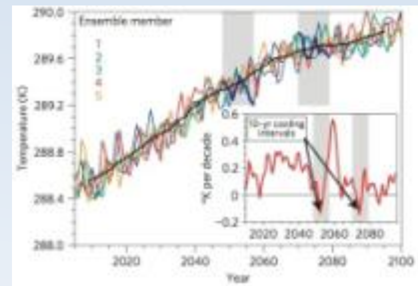

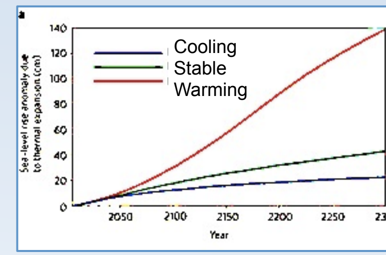

"Model-based evidence of deep-ocean heat update during surface-temperature hiatus periods"
G.A. Meehl (NCAR),
Nature Climate Change (2011)

"Relative outcomes of climate change mitigation related to global temperature versus sea-level rise"
G.A. Meehl (NCAR),
Nature Climate Change (2012)

Climate Science: Ocean circulation drove warming and deglaciation of the Southern Hemisphere

Nature, 7 February 2013

Feng He, University of Wisconsin-Madison

Objectives

- Discover what triggered beginning of last period of Southern Hemisphere warning and deglaciation
- Solve this mystery by explaining the global temperature record during the last deglaciation
- Produce a more comprehensive global temperature dataset
- Simulate earth system energy transport mechanisms

Meltwater pond in southwest Greenland. The team's simulations show that a large amount of meltwater led to a gradual slowing of the AMOC, thus triggering warming in the Southern Hemisphere.

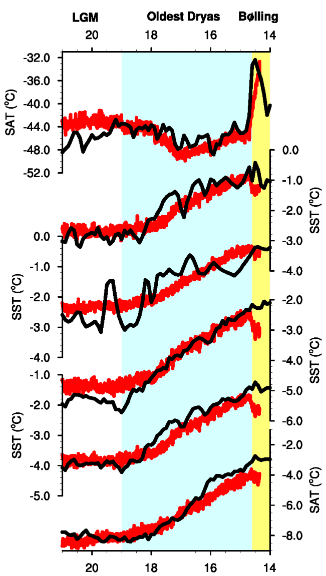

Comparison of data and model simulation. From top to bottom:

- *Greenland Surface air temperature*
- *Regional sea surface temperature in South Atlantic, Indian Ocean, South Pacific, Southern Hemisphere*
- *Antarctic region surface air temp.*

Black, data; ***red, simulation***.

LGM, ***Last Glacial Maximum***.

Accomplishments

1. Innovative reconstruction of insolation, ocean currents
   - Demonstrated how Atlantic Meridonal Ocean Circulation (AMOC) contributed to "bipolar seesaw"
2. Transient global Earth System Model simulation
   - 15,000 year simulation reproduces reconstructed global temperature response to CO_2
   - Points to CO_2, insolation, and ocean currents' role in driving global climate change over glacial cycles

"Our project could have only been done using OLCF resources, giving the computational and storage requirements." – Quarterly report

Nuclear Physics
The Limits of the Nuclear Landscape

Project PI: James Vary, Iowa State University
Allocation Program: INCITE

Science Objectives and Impact

- Identify how many protons and neutrons can be bound within a nucleus.
- Identify the nuclear drip lines that denote the end of nuclear binding.
- Project represents at step toward "designer nuclei," with uses ranging from potential cancer treatments to a better understanding of superconductivity.

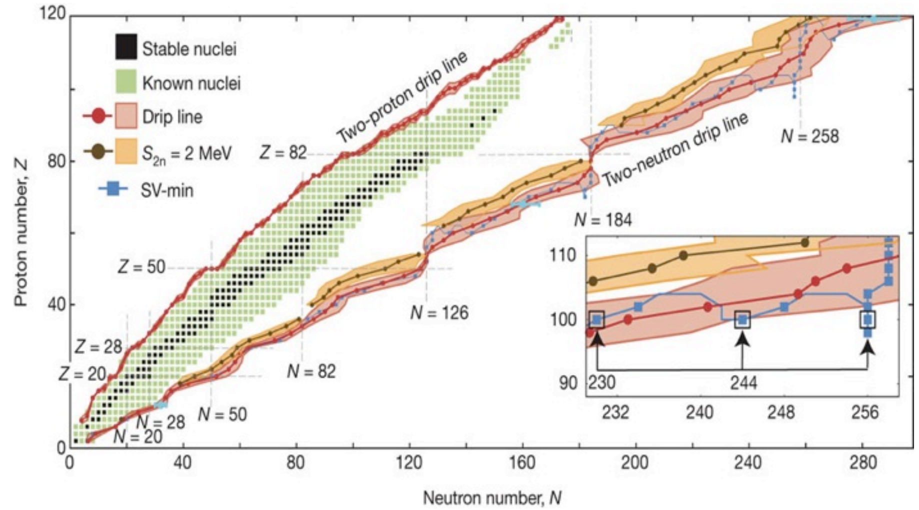

There are 767 even–even isotopes known experimentally, both stable (black squares) and radioactive (green squares). Mean drip lines and their uncertainties (red) were obtained by averaging the results of different models.

Application Performance

- Code (co-authored by OLCF staff member Hai Ah Nam) used density functional theory to solve the nuclear structure, which involved a large algebraic non-linear eigenvalue problem.
- Each set of nuclei took about two hours to calculate on 224,256-processor Jaguar system.
- Each run considered about 250,000 possible nuclear configurations.

Results (Erier, et al., Nature 2012)

- Calculations identified about 7,000 possible combinations of protons and neutrons allowed in bound nuclei with up to 120 protons.
- Several leading models of nuclear interaction shown to be largely in agreement.
- Accurate calculation of the number of bound nuclei in nature.

Publication:
J. Erier, et al. 2012. "The Limits of the Nuclear Landscape." *Nature* **486**, 509-512. doi:10.1038/nature11188

Innovation Through Industrial Partnership

Engine Cycle-to-Cycle Variation
Develop an high-performance computational strategy for modeling cyclic variation. Would be the first one-of-a-kind simulation to study CCV using 1000s of processors .

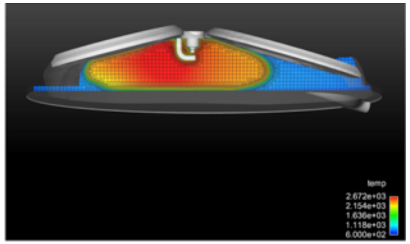

Gasoline Engine Injector Optimization
Optimization of injector hole pattern design for desired in-cylinder fuel-air mixture distributions - 4 - 40x potential improvement in workflow throughput via 100s of ensemble simulations

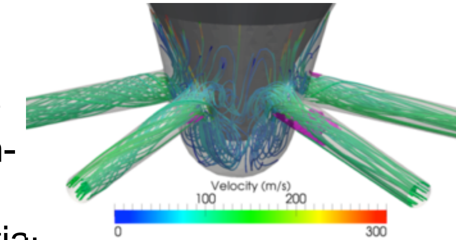

Li-ion Batteries
Discover and optimize new classes of solid inorganic Li-ion electrolytes with high ionic and low electronic conductivity, and good electrochemical stability.

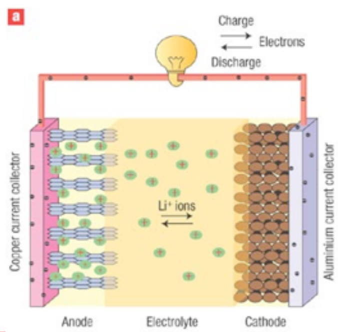

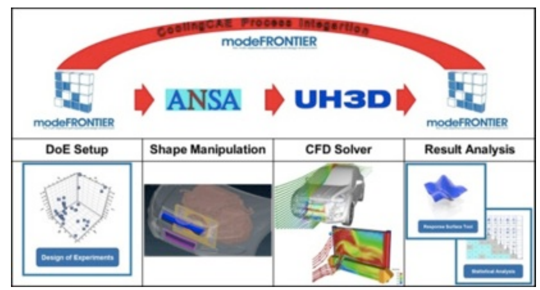

Underhood Cooling
Developed a new, efficient & automatic analytical cooling package optimization process leading to one of a kind design optimization of cooling systems

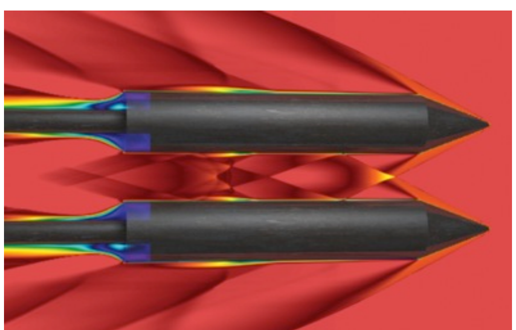

Design Innovation
Ramgen Power Systems accelerates their design of shock wave turbo compressors for carbon capture and sequestration

Turbo Machinery Efficiency
General Electric, for the first time, simulated unsteady flow in turbo machinery, opening new opportunities for design innovation and efficiency improvements.

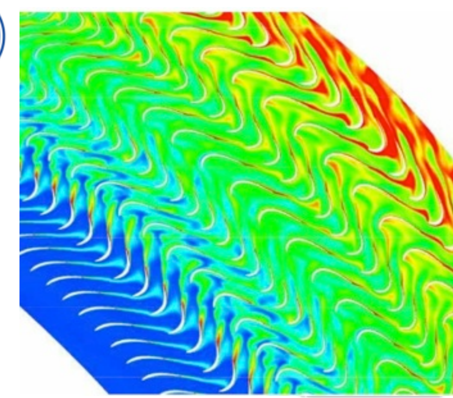

Computational Fluid Dynamics
"Smart Truck Optimization"

PI Mike Henderson,
Allocation Program: Discretionary

Science Objectives and Impact

- Apply advanced computational techniques from aerospace industry to improve fuel efficiency of Class 8 Long Haul Trucks
- New California Air Resources Board (CARB) requires 5% fuel efficiency increase on all Class 8 trucks; national drive to reduce emissions.
- Determine design of add-on parts to substantially reduce drag
- If all 1.3 million trucks operated with a passenger car's drag:
  - ✓ Save 1.5 billion gallons of diesel annually
  - ✓ Reduce 16.4 million tons CO_2 annually
  - ✓ Save $5 billion in fuel costs annually

Science Results

Unprecedented detail and accuracy of a Class 8 Tractor-Trailer aerodynamic simulation.

UT-6 Trailer UnderTray System reduces Tractor/ Trailer drag by 12%.

- Minimizes drag associated with trailer underside components
- Compresses and accelerates incoming air flow and injecting high energy air into trailer wake
- Pulls high energy, attached air flow from the top of the trailer down into trailer wake

Application Performance

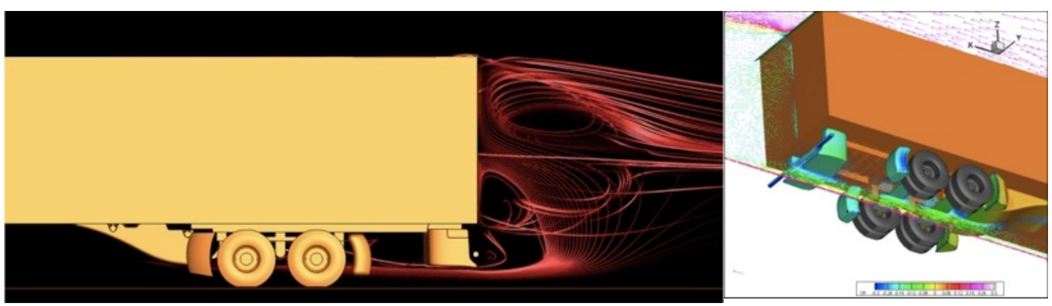

- CFD results using NASA' a FUN3D within 1% of physical test results
- Access to Jaguar reduced single run times from 15 hours to less than 2 hours
- Stage set to do full Navier-Stokes-based optimization of large trucks

Competitiveness Impact

For BMI:

- Time from conception to manufacture reduced 50% (from 3 years t0 1.5 years)
- Early-to- market advantage will result in earlier and greater revenue recognition

For BMI's customers

- Demonstrated fuel mileage improvements of 7% to 12% available 2011.
- Exceeds California CARB requirements.

Turbo Machinery Efficiency

Project PI: Holmes, GE
Allocation Program: Discretionary

Removing Steady-Flow Assumptions Helps Improve Efficiency and Reduce Noise

Science Objectives and Impact

- Design efficient turbo machinery generating less noise
  - 1% reduction in fuel consumption can save $20B
- What is the “steady-flow” assumption hiding?

Image on right:
Computational Fluid Dynamics results depicting entropy for unsteady flow through turbine with fan blades removed.

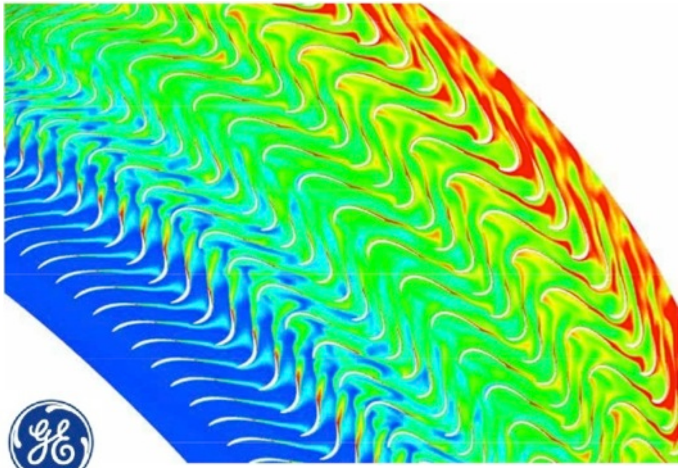

Holmes, et al., SciDAC (2011)

Application Performance

- Scaled GE’s TACOMA CFD code from ~1K cores to 80K cores
- TACOMA has good weak scaling

OLCF contributions:

- Guide GE in porting TACOMA.

Results

- GE ran their largest ever CFD simulation on Jaguar
- Unsteady flow and efficiency losses localize at end walls
- On the basis of the success of these calculations, GE Global Research has purchased a Cray XE6

> “Our two million hours on 8K processors of Jaguar has allowed us to see phenomena we have never seen before.”
>
> -- Graham Holmes, GE Global Research

Turbo Compressor Innovation

Scaling CFD code to machining titanium for a new turbo compressor in a few months

Science Objectives and Impact

- Ramgen Power Systems is developing shock wave compression turbo machinery to meet DOE goals for reducing Carbon Capture and Sequestration costs
- Complementary goal: design a gas turbine with dramatically lower costs and higher efficiency

Image on right:
3.7 billion grid cell simulation of two-body shock-wave, boundary layer interaction, showing high-resolution shock capture

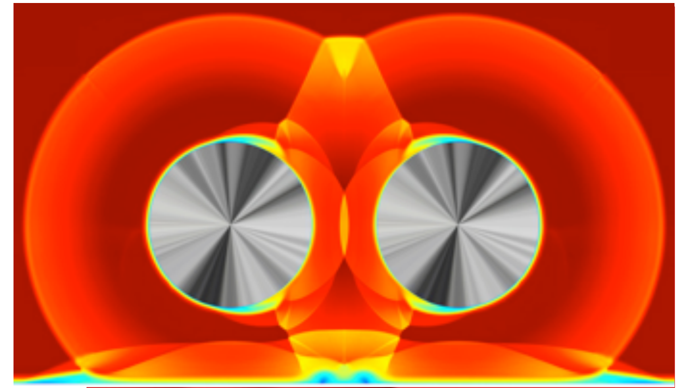

Grosvenor et al., EUCASS Conf. (2011)

Application Performance

- 50x improvement in code scalability with more efficient memory utilization
- Accelerated I/O by 10x with optimizations and ADIOS
- Intelligent use of ensembles to explore parameter space using over 240,000 cores

Results

- Advanced Ramgen's aerodynamic design process
- Observed designs with valuable new characteristics, from ensembles not possible without Jaguar
- Created a new workflow paradigm that accelerates design of compressors
- Accelerated computational design cycle for turbo machinery from months to 8 hours!

"The use of Jaguar has cut the projected time from concept to a commercial product by at least two years and the cost by over $2 million," -- Ramgen's CEO and Director Doug Jewett.

What is the Leadership Computing Facility?

- Collaborative, multi-lab, DOE/SC initiative ranked top domestic priority in ***Facilities for the Future of Science: A Twenty-Year Outlook.***
- Mission: Provide the computational and data science resources required to solve the most important scientific & engineering problems in the world.
- Highly competitive user allocation program (INCITE, ALCC).
- Projects receive 100x more hours 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

We have increased our system capability by 10,000 times since our founding in 2004

- Strong partnerships with computer designers.
- LCF users employ large portions of the machine for large fractions of time.
- We accomplished this through strong partnerships with our users to scale codes and algorithms.

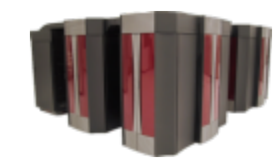

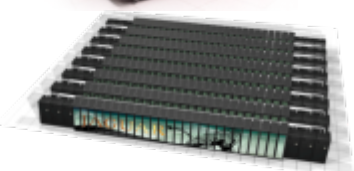

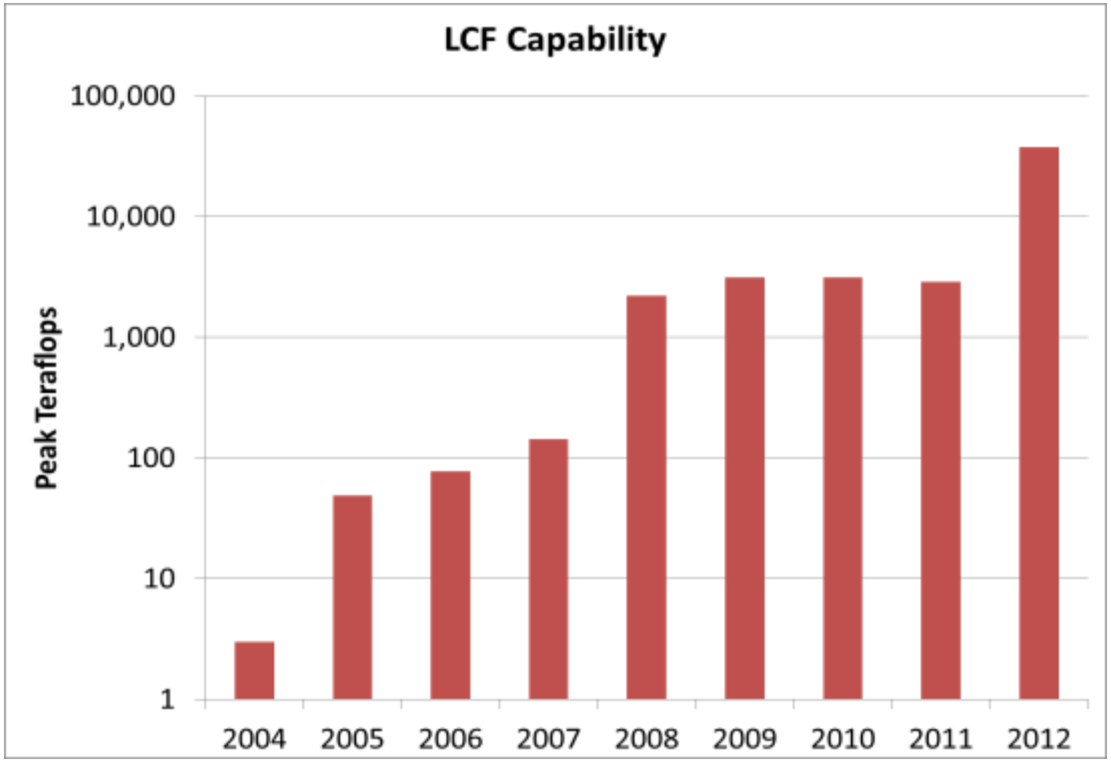

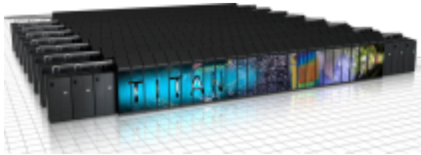

Oak Ridge National Laboratory
Argonne NATIONAL LABORATORY

Science challenges for LCF in next decade

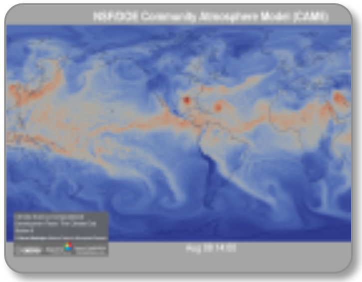

Climate Change Science
Understand the dynamic ecological and chemical evolution of the climate system with uncertainty quantification of impacts on regional and decadal scales.

Combustion Science
Increase efficiency by 25%-50% and lower emissions from internal combustion engines using advanced fuels and new, low-temperature combustion concepts.

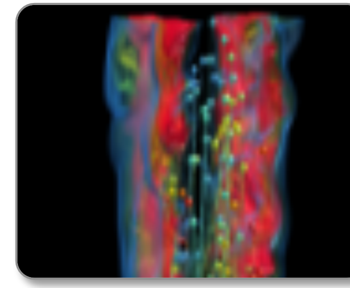

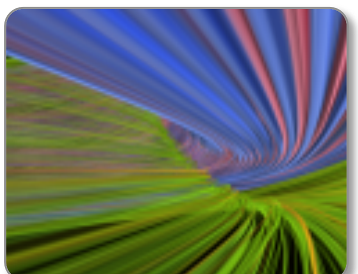

Fusion Energy/ITER
Develop predictive understanding of plasma properties, dynamics, and interactions with surrounding materials.

Biomass to Biofuels
Enhance the understanding and production of biofuels for transportation and other bio-products from biomass.

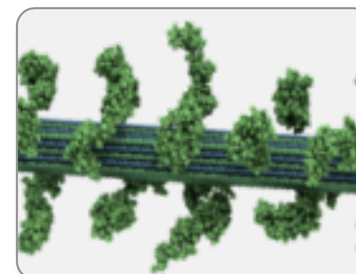

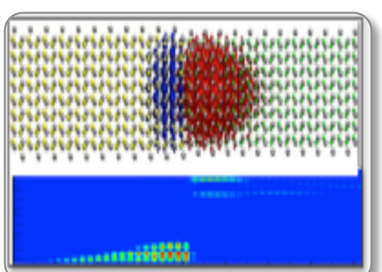

Solar Energy
Improve photovoltaic efficiency and lower cost for organic and inorganic materials.

Globally Optimized Accelerator Designs
Optimize designs as the next generations of accelerators are planned, increasingly detailed models will be needed to provide a proof of principle and a cost-effective method to design new light sources.

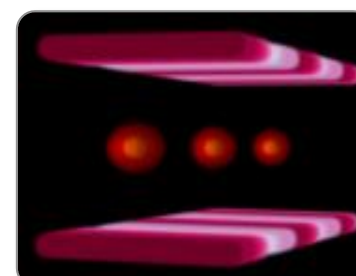

Oak Ridge National Laboratory
Argonne National Laboratory

Leadership Facility Strategy for the future
Roadmap and Timeline

Through a three-phase plan executed over the next decade, the LCF will deploy a series of ever-more-powerful, balanced, scalable, HPC and data resources to support the most challenging computational problems of the nation.

Phase 1: Procure and operate pre-Exascale systems (2016)
Phase 2: Procure and operate Exascale systems (2020)
Phase 3: Procure and operate Second generation Exascale systems (2024)

Computer System requirements for each Leadership Computing Center

	2012	2016	2020	2024
Peak FLOP/s	10-20 PF	100-200 PF	500-2000 PF	2000-4000 PF
Memory	0.5-1 PB	5-10 PB	32-64 PB	50-100 PB
Burst Buffer	N/A	500 TB	3 PB	5 PB
Storage Disk +tape	20+100 PB	100+1000 PB	1+10 EB	5+50 EB
Power & Space	6-12 MW 5,000-10,000 ft^2	15-20 MW 8,000-15,000 ft^2	20-30 MW 20,000 ft^2	25-35 MW 25,000 ft^2

Oak Ridge National Laboratory   Argonne National Laboratory

Science requires exascale computational capability in this decade

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

- Deliver transforming discoveries in climate, materials, biology, energy technologies, etc.
- Enabling investigation of otherwise inaccessible systems, from regional climate impacts to energy grid dynamics

Vision: Maximize scientific productivity and progress on largest scale computational problems

- World-class computational resources and specialized services for the most computationally intensive problems
- Stable hardware/software path of increasing scale to maximize productive applications development

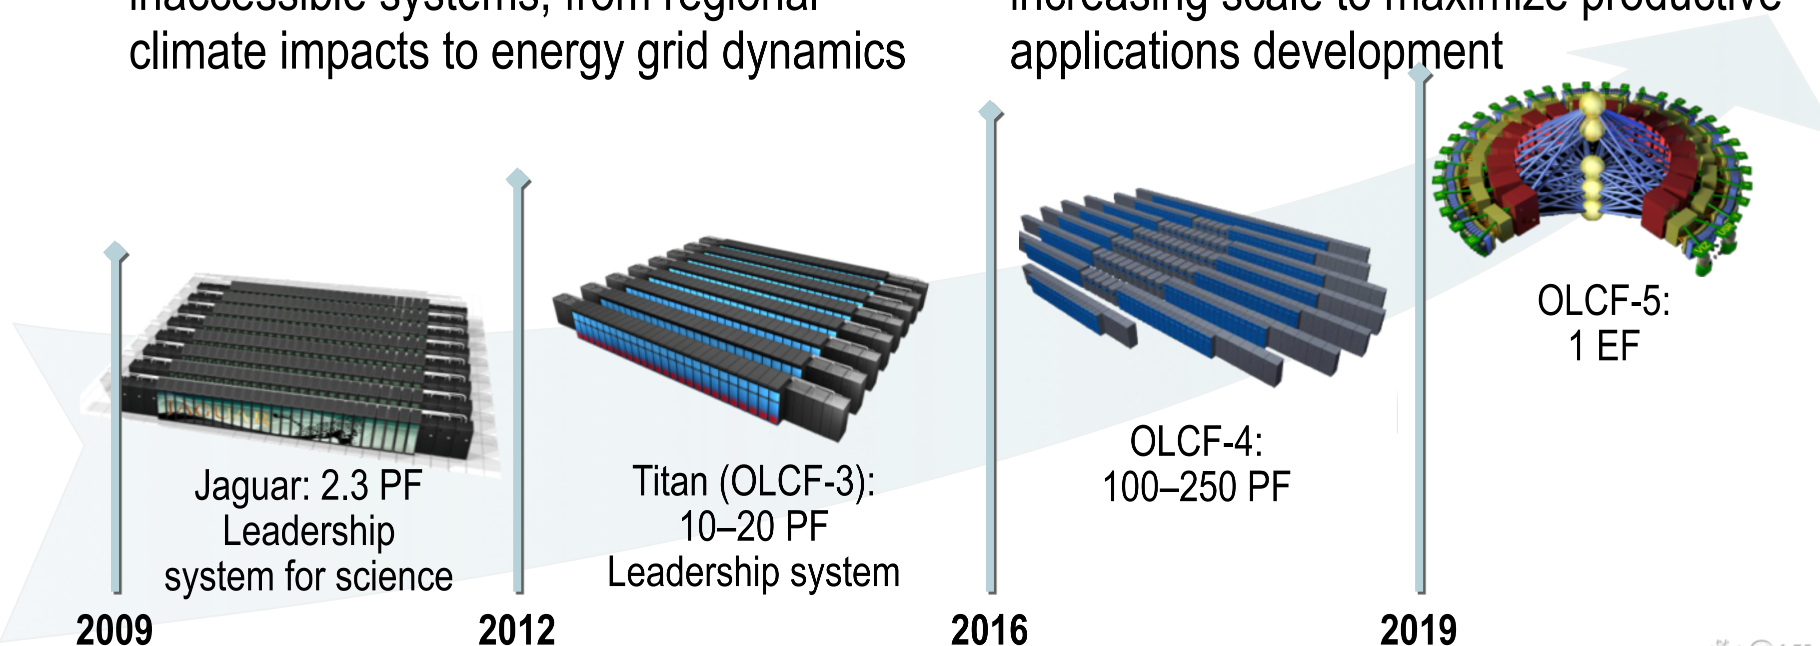

Slide courtesy of Jeff Nichols

Architectural Trends – No more free lunch

- Moore's Law continues (green line) but the CPU clock rates stopped increasing in 2003 (dark blue line) due to power constraints. (light blue line)
- Power is capped by heat dissipation and $$$
- Performance increases have been coming through increased parallelism

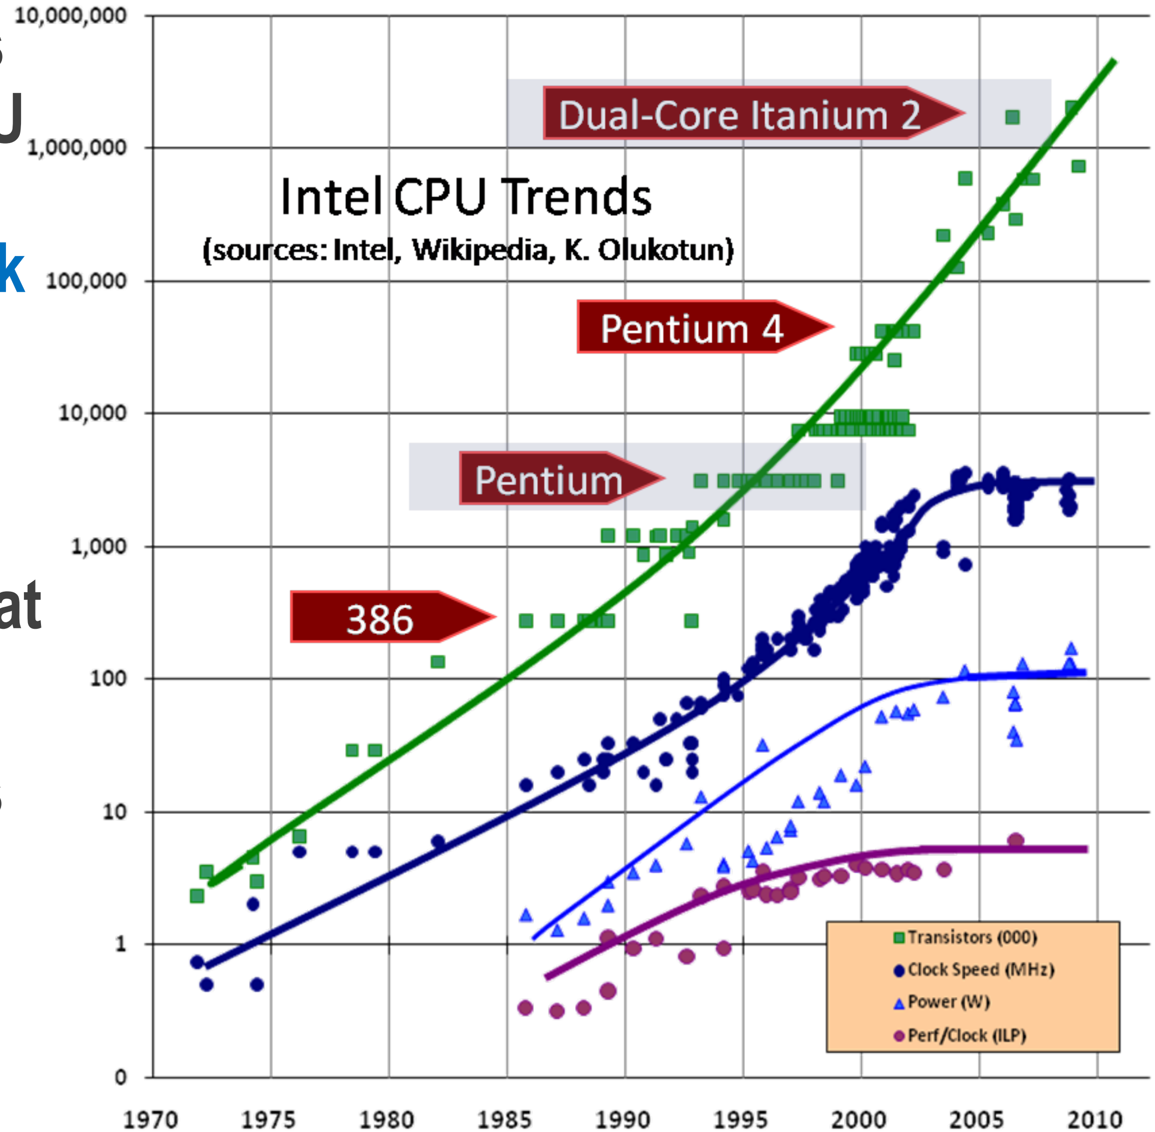

Herb Sutter: Dr. Dobb's Journal:
http://www.gotw.ca/publications/concurrency-ddj.htm

Power is THE problem

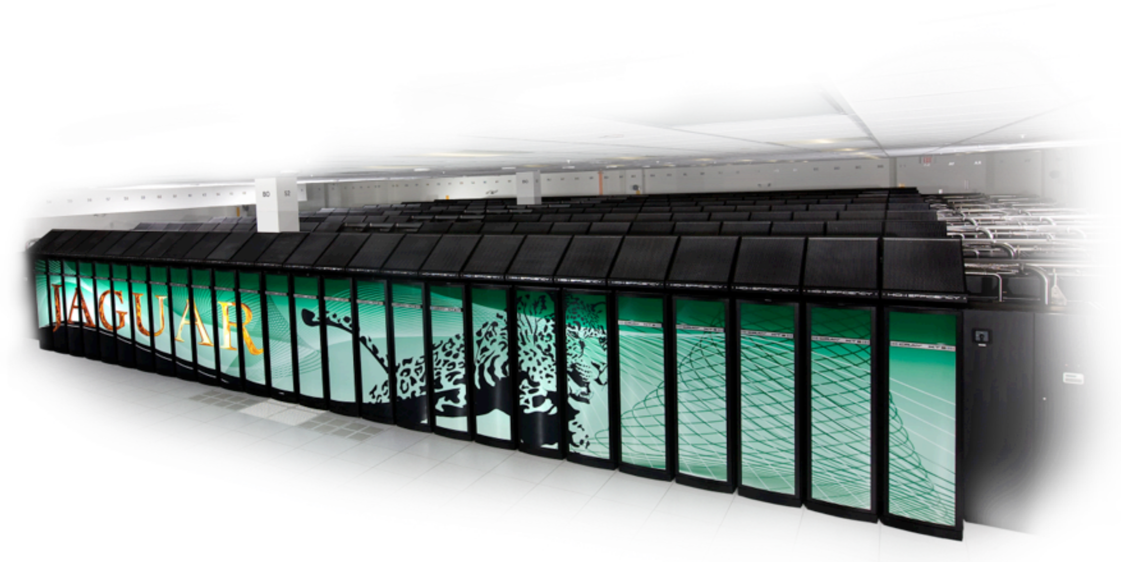

Power consumption of 2.3 PF Jaguar:
7 megawatts, equivalent to that of a small city (5,000 homes)

Using traditional CPUs is not economically feasible

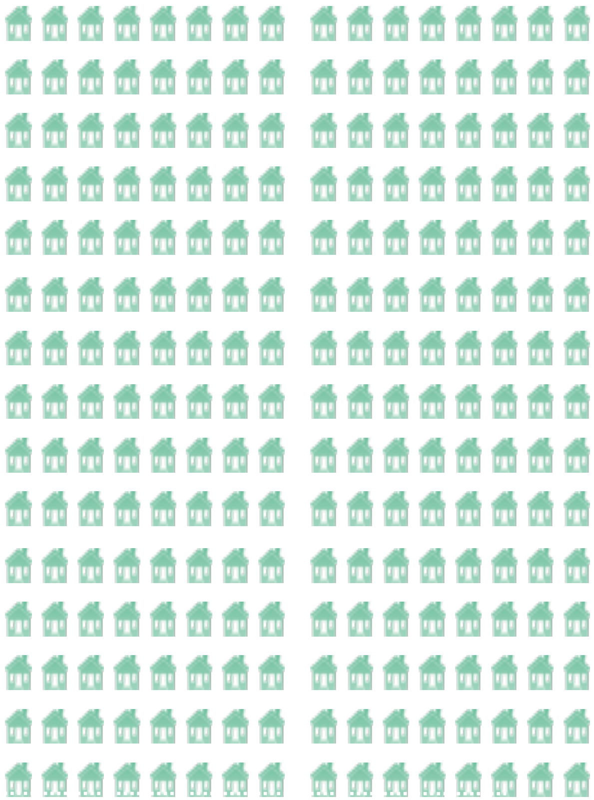

20 PF system:
30 megawatts (30,000 homes)

Why GPUs?
High performance and power efficiency on path to exascale

- Hierarchical parallelism improves scalability of applications
- Expose more parallelism through code refactoring and source code directives
  - Doubles performance of many codes
- Heterogeneous multicore processor architecture: Using right type of processor for each task
- Data locality: Keep data near processing
  - GPU has high bandwidth to local memory for rapid access
  - GPU has large internal cache
- Explicit data management: Explicitly manage data movement between CPU and GPU memories

CPU

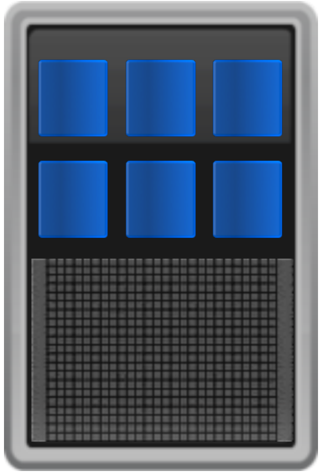

- Optimized for sequential tasks

GPU Accelerator

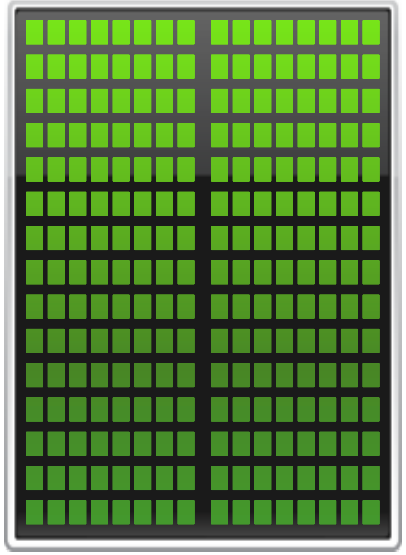

- Optimized for many simultaneous tasks
- 10× performance per socket
- 5× more energy-efficient systems

ORNL's "Titan" Hybrid System: World's Most Powerful Computer

#1 TOP500 SUPERCOMPUTER SITES

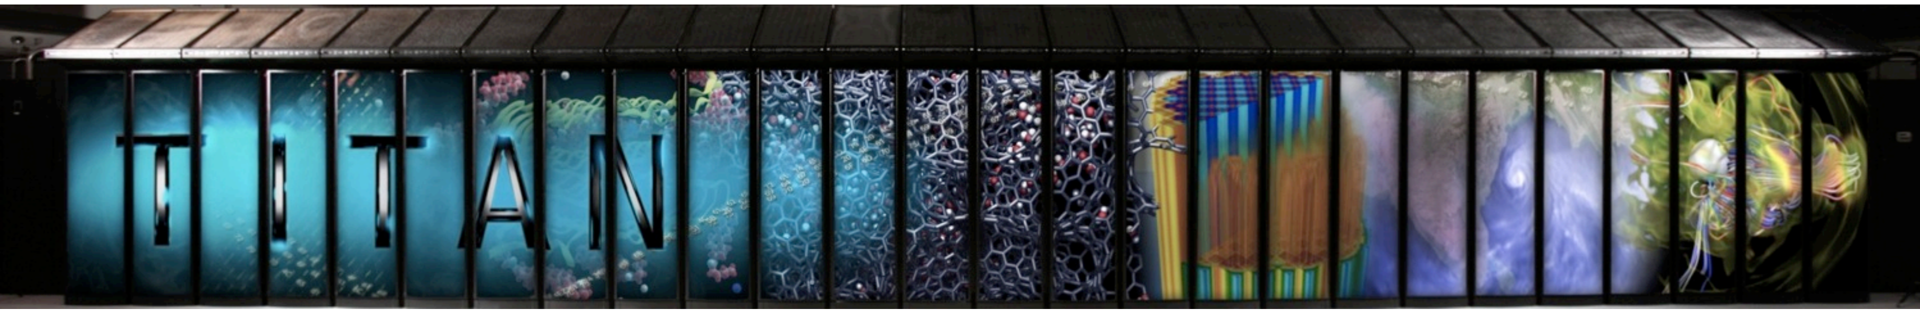

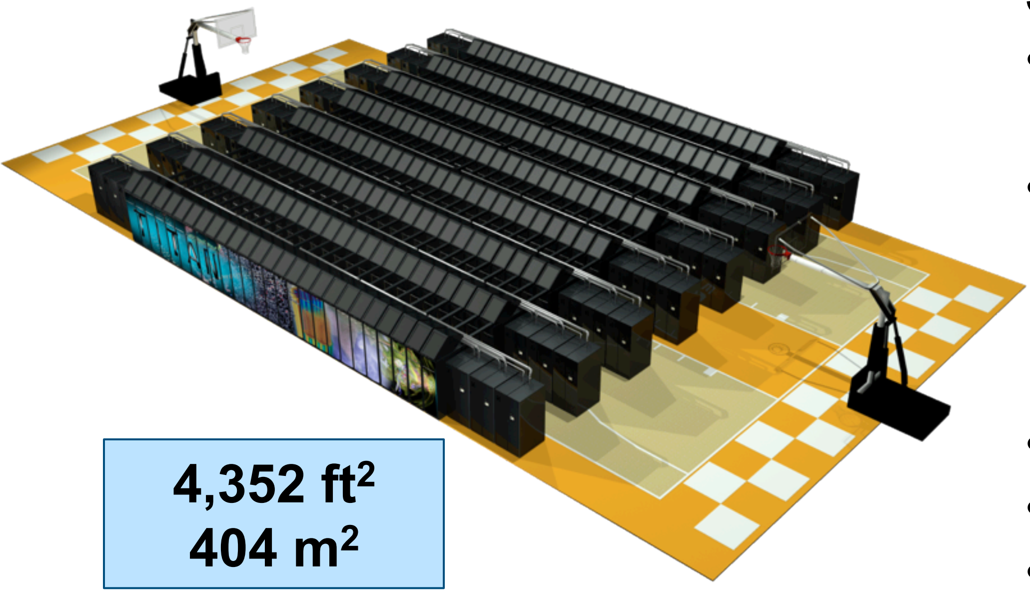

SYSTEM SPECIFICATIONS:
- Peak performance of 27.1 PF
  - 24.5 GPU + 2.6 CPU
- 18,688 Compute Nodes each with:
  - 16-Core AMD Opteron CPU
  - NVIDIA Tesla "K20x" GPU
  - 32 + 6 GB memory
- 512 Service and I/O nodes
- 200 Cabinets
- 710 TB total system memory
- Cray Gemini 3D Torus Interconnect
- 8.9 MW peak power

OAK RIDGE National Laboratory

Cray XK7 Compute Node

XK6 Compute Node Characteristics
AMD Opteron 6200 Interlagos 16 core processor @ 2.2GHz
Tesla M2090 @ 665 GF with 6GB GDDR5 memory
Host Memory 32GB 1600 MHz DDR3
Gemini High Speed Interconnect
Upgradeable to NVIDIA's next generation KEPLER processor in 2012
Four compute nodes per XK6 blade. 24 blades per rack

Hybrid Programming Model

- **On Jaguar, with 299,008 cores, we were seeing the limits of a single level of MPI scaling for most applications**
- **To take advantage of the vastly larger parallelism in Titan, users need to use hierarchical parallelism in their codes**
  - Distributed memory: **MPI**, SHMEM, PGAS
  - Node Local: **OpenMP**, Pthreads, local MPI communicators
  - Within threads: Vector constructs on GPU, libraries, **OpenACC**
- ***These are the same types of constructs needed on all multi-PFLOPS computers to scale to the full size of the systems!***

Early science applications on Titan

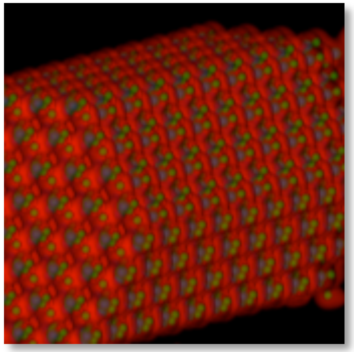

Materials science (WL-LSMS)
Role of material disorder, statistics, and fluctuations in nanoscale materials and systems

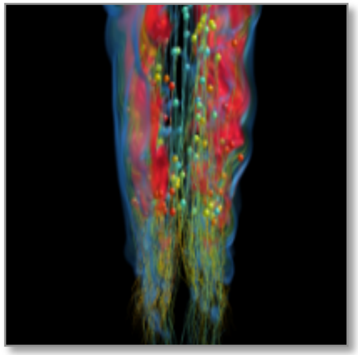

Combustion (S3D)
Combustion simulations to enable the next generation of diesel/ biofuels to burn more efficiently

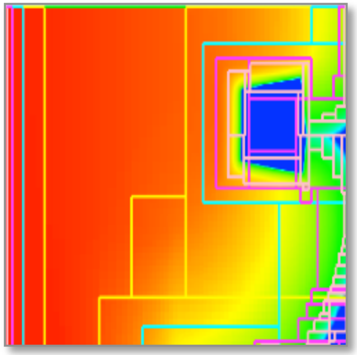

Astrophysics (NRDF)
AMR radiation transport – critical to astrophysics, laser fusion, combustion, atmospheric dynamics, and medical imaging

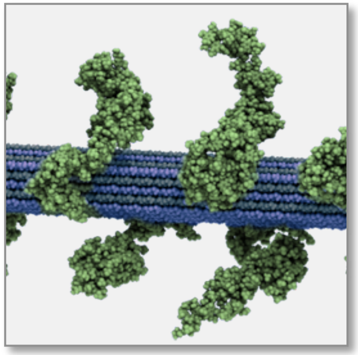

Biofuels (LAMMPS)
A multiple capability molecular dynamics code

Climate change (CAM-SE)
Answer questions about specific climate change adaptation and mitigation scenarios; realistically represent features like precipitation patterns/statistics and tropical storms

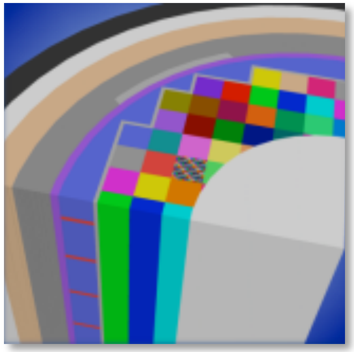

Nuclear energy (Denovo)
Unprecedented high-fidelity radiation transport calculations that can be used in a variety of nuclear energy and technology applications

How effective are GPUs on scalable applications?

OLCF-3 early science codes: Very early performance measurements

Application	Performance ratio, XK7 (with K20x) vs XE6	Cray XK7: K20x GPU plus AMD 6274 CPU Cray XE6: Dual AMD 6274 and no GPU
S3D	1.8	• Turbulent combustion • 6% of Jaguar workload
Denovo sweep	3.8	• Sweep kernel of 3D neutron transport for nuclear reactors • 2% of Jaguar workload
LAMMPS	7.4 (mixed precision)	• High-performance molecular dynamics • 1% of Jaguar workload
WL-LSMS	3.8	• Statistical mechanics of magnetic materials • 2% of Jaguar workload • 2009 Gordon Bell Winner
CAM-SE	1.8 (estimate)	• Community atmosphere model • 1% of Jaguar workload

DOE Computational Facilities Allocation Policy for Leadership Facilities

Primary Objective:

- *"Provide substantial allocations to the open science community through an peered process for a small number of high-impact scientific research projects"*

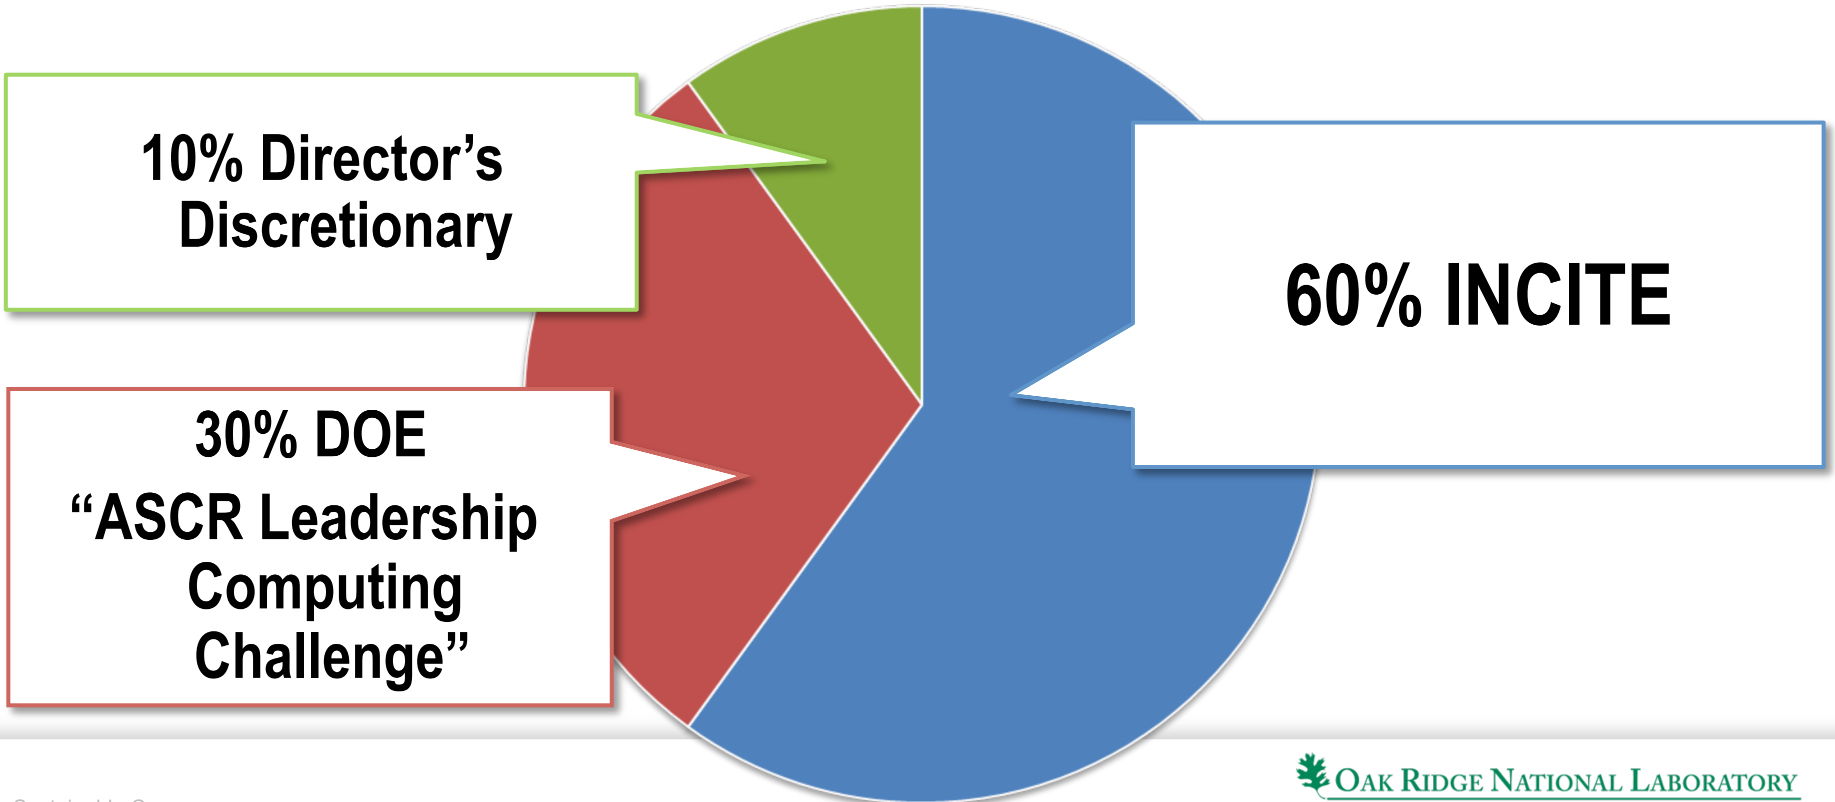

OLCF allocation programs
Selecting applications of national importance

	60% INCITE		30% ALCC		10% Director's Discretionary	
Mission	High-risk, high-payoff science that requires LCF-scale resources		High-risk, high-payoff science aligned with DOE mission		Strategic LCF goals	
Call	1x/year – (Closes June)		1x/year – (Closes February)		Rolling	
Duration	1-3 years, yearly renewal		1 year		3m,6m,1 year	
Typical Size	30 - 40 projects	20M - 100M core-hours/yr.	5 - 10 projects	1M – 75M core-hours/yr.	100s of projects	10K – 1M core-hours
Review Process	Scientific Peer-Review	Computational Readiness	Scientific Peer-Review	Computational Readiness	Strategic impact and feasibility	
Managed by	INCITE management committee (ALCF & OLCF)		DOE Office of Science		OLCF management	
Availability	Open to all scientific researchers and organizations including industry					

2014 INCITE Call for Proposals

- Planning Request for Information (RFI)
- Call opens April, 2013. Closes June, 2013
- Expect to allocate more than 5 billion core-hours
- Expect 3X oversubscription
- Awards to be announced in November for CY 2014
- Average award to exceed 50 million core-hours

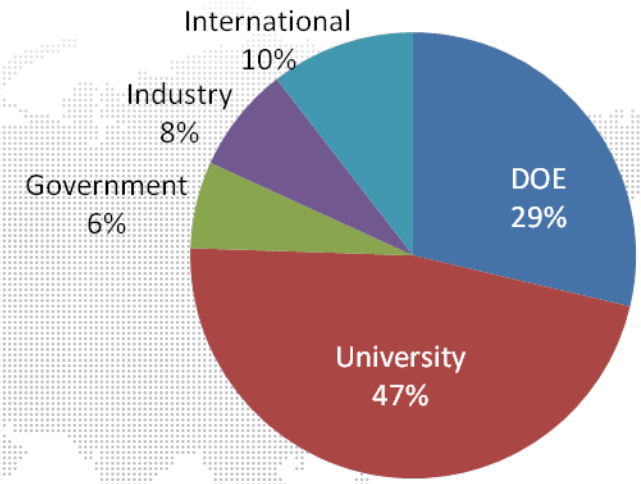

Number of 2013 INCITE Submittals by PI Affiliation

Reaching out to Researchers

Nearly 50% of the non-renewal proposals are by new PIs.

Contact information
Julia C. White, INCITE Manager
whitejc@DOEleadershipcomputing.org

Leadership Metric and Scheduling Policy

OLCF has a mandate to be used for large, *leadership-class* (aka *capability*) jobs.

OLCF implements queue policies that enable large jobs to run in a timely fashion.

- Priority in queue is set by the time a job has been waiting relative to other jobs in the queue.
- We use several factors to modify the *apparent* time a job has been waiting. These factors include:
  - The job's processor core request size.
  - The queue to which the job is submitted.

Leadership Usage Metric:

35% of the CPU time used on the system will be accumulated by jobs using 20% or more of the available nodes.

OLCF Quarterly Reports

Principal Investigators of current LCF projects must submit quarterly progress reports. The quarterly reports are essential as the OLCF must track and report the use of the center's resources.

- Achievements
- Milestones
- Report Publications
- Important Center Feedback

Conclusions

- Leadership computing is for the critically important problems that need the most powerful compute and data infrastructure
- Our user resources are in high demand and are effectively used.
- Computer system performance increases through parallelism
  - Clock speed trend flat to slower over coming years
  - Applications must utilize all inherent parallelism
- Accelerated, hybrid-multicore computing solutions are performing very well on real, complex scientific applications.
- OLCF resources are available through open, peer-reviewed allocation mechanisms.

Questions?