U.S. DEPARTMENT OF ENERGY | Office of Science

ASCR OVERVIEW

October 27, 2011
HEPAP

Daniel Hitchcock

Acting Associate Director

Advanced Scientific Computing Research

Advanced Scientific Computing Research

Delivering world leading computational and networking capabilities to extend the frontiers of science and technology

The Scientific Challenges:

- Deliver next-generation scientific applications using today's petascale computers.
- Discover, develop and deploy tomorrow's exascale computing and networking capabilities.
- Develop, in partnership with U.S. industry, next generation computing hardware and tools for science.
- Discover new applied mathematics and computer science for the ultra-low power, multicore-computing future.
- Provide technological innovations for U.S. leadership in Information Technology to advance competitiveness.

FY 2012 Highlights:

- Research in uncertainty quantification for drawing predictive results from simulation
- Co-design centers to deliver next generation scientific applications by coupling application development with formulation of computer hardware architectures and system software.
- Investments in U.S. industry to address critical challenges in hardware and technologies on the path to exascale
- Installation of a 10 petaflop low-power IBM Blue Gene/Q at the Argonne Leadership Computing Facility and a hybrid, multi-core prototype computer at the Oak Ridge Leadership Computing Facility.

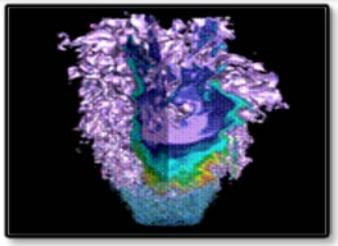
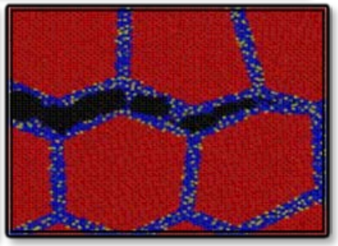
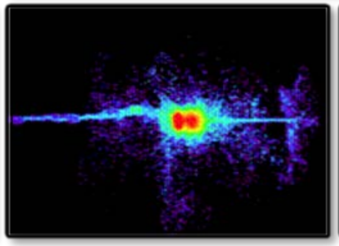
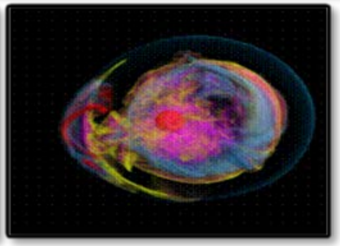

ASCR Facilities

- Providing the Facility – **High-End and Leadership Computing**
- Investing in the Future - **Research and Evaluation Prototypes**
- Linking it all together – **Energy Sciences Network (ESnet)**

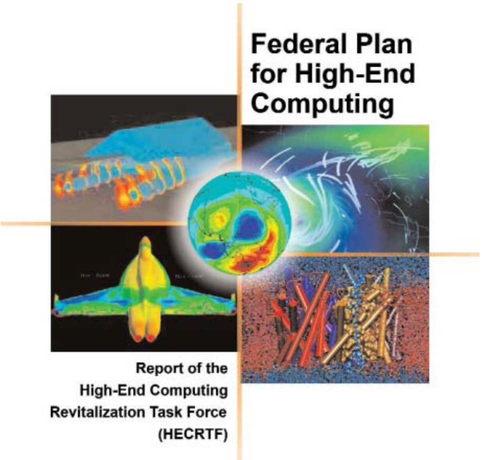

OLCF
Oak Ridge Leadership Computing Facility

NERSC

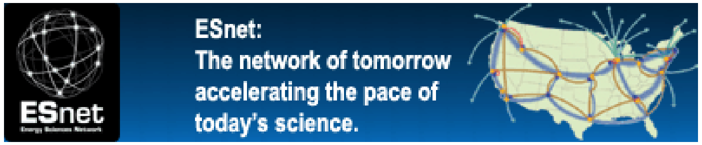

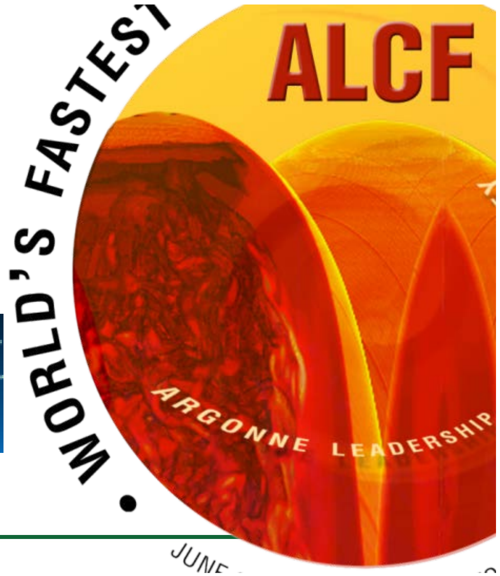

U.S. DEPARTMENT OF ENERGY | Office of Science

Energy Sciences Net

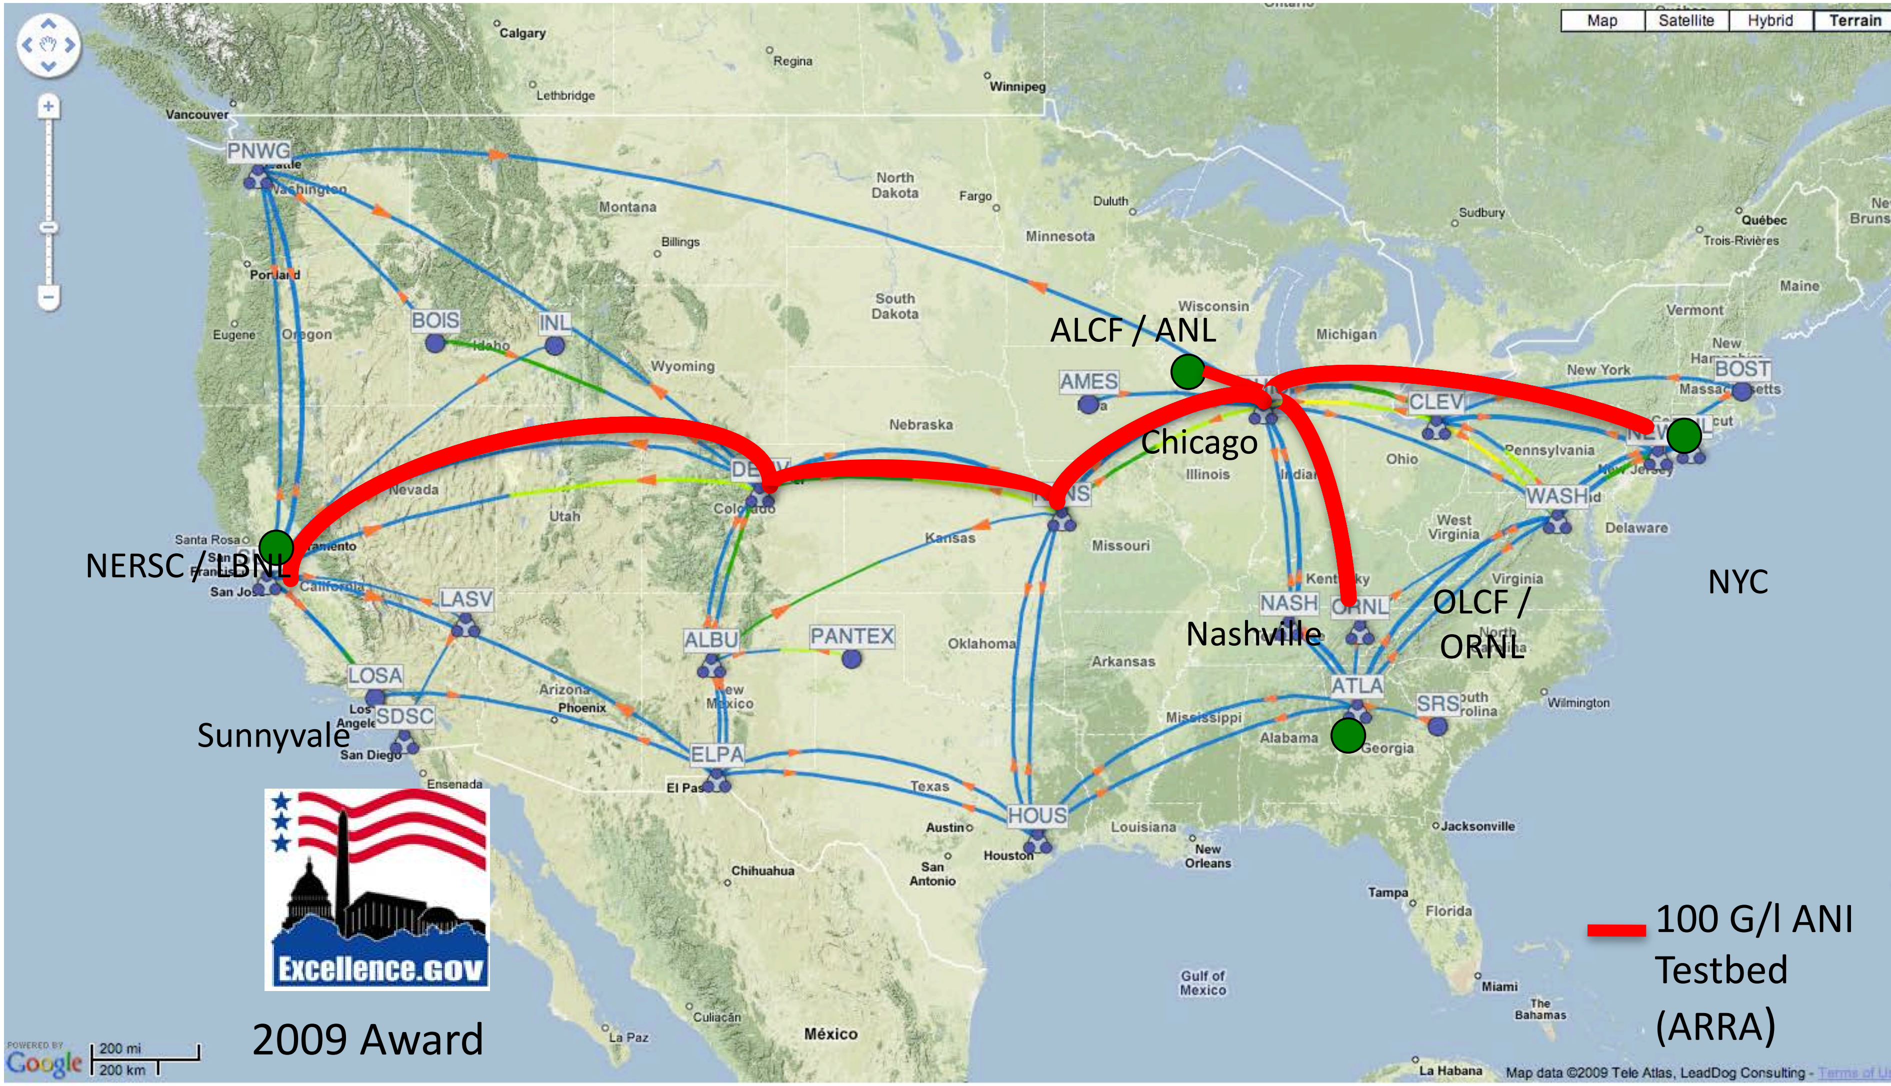

ASCR Research

Substantial innovation is needed to provide essential system and application functionality in a timeframe consistent with the anticipated availability of hardware

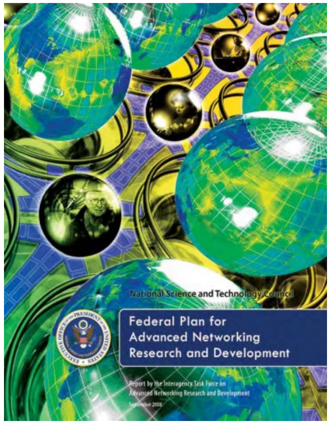

Provide forefront research knowledge and foundational tools:

- Applied Mathematics
- Computer Science
- SciDAC
- Next Generation Networking for Science

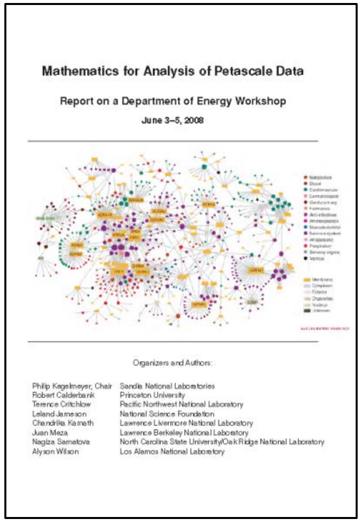

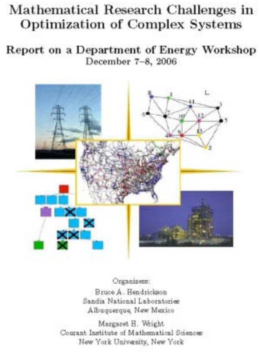

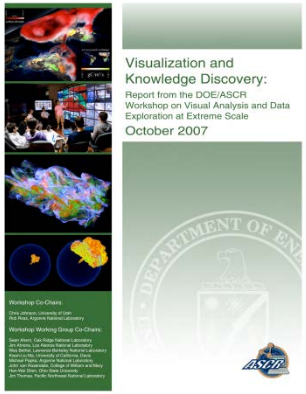

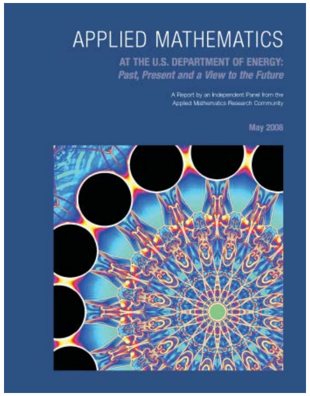

Looking Backward and Forward

SciDAC Centers for Enabling Technology and Institutes

Science Application Partnerships

Hopper
Jaguar
Intrepid

Many Core
Energy Aware

X Stack

CoDesign

SciDAC Institutes

Strategic ASCR – SC Office Partnerships

Titan
Mira
???

SciDAC Institutes

Goals & Objectives

- Deliver tools and resources to lower barriers to effectively use state-of-the-art computational systems;
- Create mechanisms to address computational grand challenges across different science application areas;
- Incorporate basic research results from Applied Mathematics and Computer Science into computational science challenge areas and demonstrate that value
- Grow the Nation's computational science research community.

Awards- Up to $13M/year over 5 years available to support 1–5 Institutes

Eligible applicants- DOE National Laboratories, Universities, Industry and other organizations

Expected outcome- Institutes that cover a significant portion of DOE computational science needs on current and emerging computational systems.

Timeline

- Solicitations opened- February 23, 2011
- Letters of Intent- March 30, 2011
- Solicitations closed- May 2, 2011
- First awards- end of FY2011

Answers to Inquiries- http://science.doe.gov/ascr/research/scidac/SciDAC3InstitutesFAQ.html

Scientific Discovery through Advanced Computing (SciDAC) Institutes – FY11 Awards

FASTMath – Frameworks, Algorithms, and Scalable Technologies for Mathematics

Topic areas: Structured & unstructured mesh tools, linear & nonlinear solvers, eigensolvers, particle methods, time integration, differential variational inequalities

QUEST – Quantification of Uncertainty in Extreme Scale Computations

Topic areas: Forward uncertainty propagation, reduced stochastic representations, inverse problems, experimental design & model validation, fault tolerance

SUPER – Institute for Sustained Performance, Energy and Resilience

Topic areas: Performance engineering (including modeling & auto-tuning), energy efficiency, resilience & optimization

FASTMath Director – Lori Diachin, LLNL	QUEST Director – Habib N. Najm, SNL	SUPER Director – Robert F. Lucas, USC
Argonne National Laboratory	Los Alamos National Laboratory	Argonne National Laboratory
Lawrence Berkeley National Lab	**Sandia National Laboratories***	Lawrence Berkeley National Lab
Lawrence Livermore National Lab*	Johns Hopkins University	Lawrence Livermore National Lab
Sandia National Laboratories	Massachusetts Institute of Technology	Oak Ridge National Laboratory
Rensselaer Polytechnic Institute	University of Southern California	University of California at San Diego
	University of Texas at Austin	University of Maryland
		University of North Carolina
		University of Oregon
		University of Southern California*
		University of Tennessee at Knoxville
		University of Utah

Strategic ASCR – SC Program Partnerships

- **Goals and Objectives**
  - Partner with SC Programs to combine the best applied mathematics, computer science, and networking with SC program expertise to enable Strategic advances in program missions.
- **Five FOAs out before September 30, 2011:**
  - Fusion Energy Science: topics in Edge Plasma Physics, Multiscale Integrated Modeling, and Materials Science
  - High Energy Physics: topics in Cosmological Frontier, Lattice Gauge Theory QCD, and Accelerator Modeling and Simulation
  - Nuclear Physics: topics in Low- and Medium-Energy Nuclear Physics and Heavy Ion Collider Physics
  - Biological and Environmental Research: topics in dynamics of atmospheres, oceans, and ice sheets that support Earth System Modeling
  - Basic Energy Sciences: topics in theoretical chemistry and materials science, especially excited states of atoms and molecules, materials electron correlation.

The Future is about Energy Efficient Computing

- **At $1M per MW, energy costs are substantial**
- **1 petaflop in 2010 will use 3 MW**
- **1 exaflop in 2018 at 200 MW with "usual" scaling**
- **1 exaflop in 2018 at 20 MW is target**

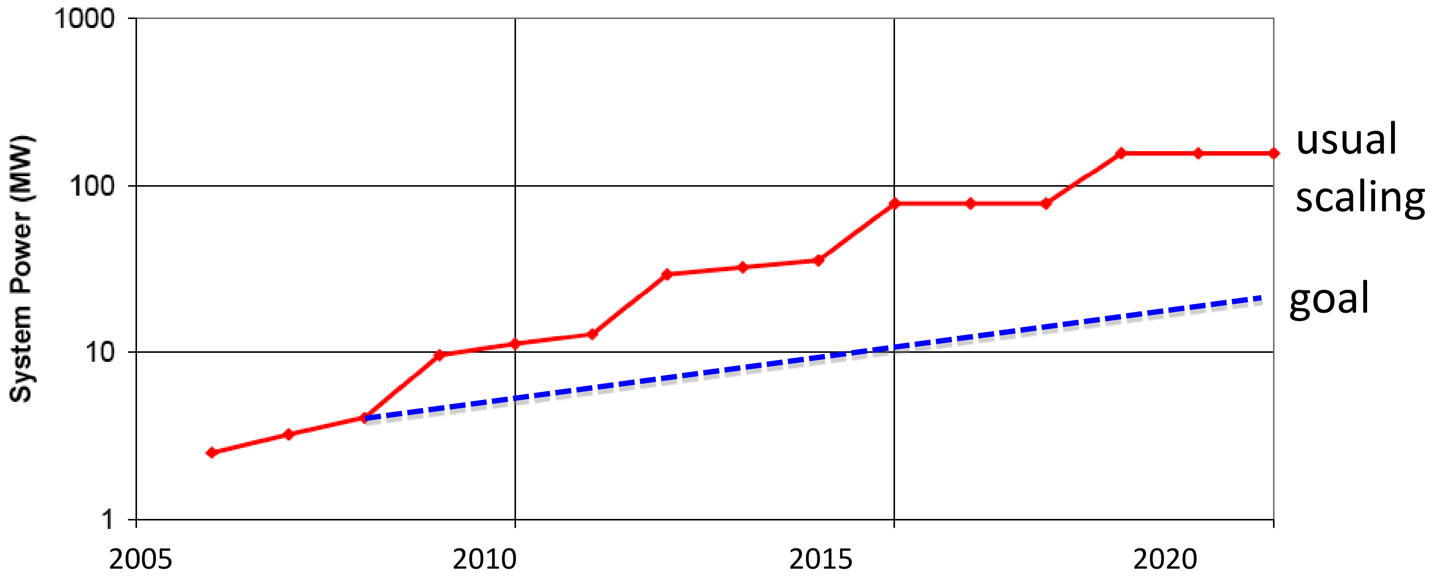

The Fundamental Issue: Where does the Energy (and Time) Go?

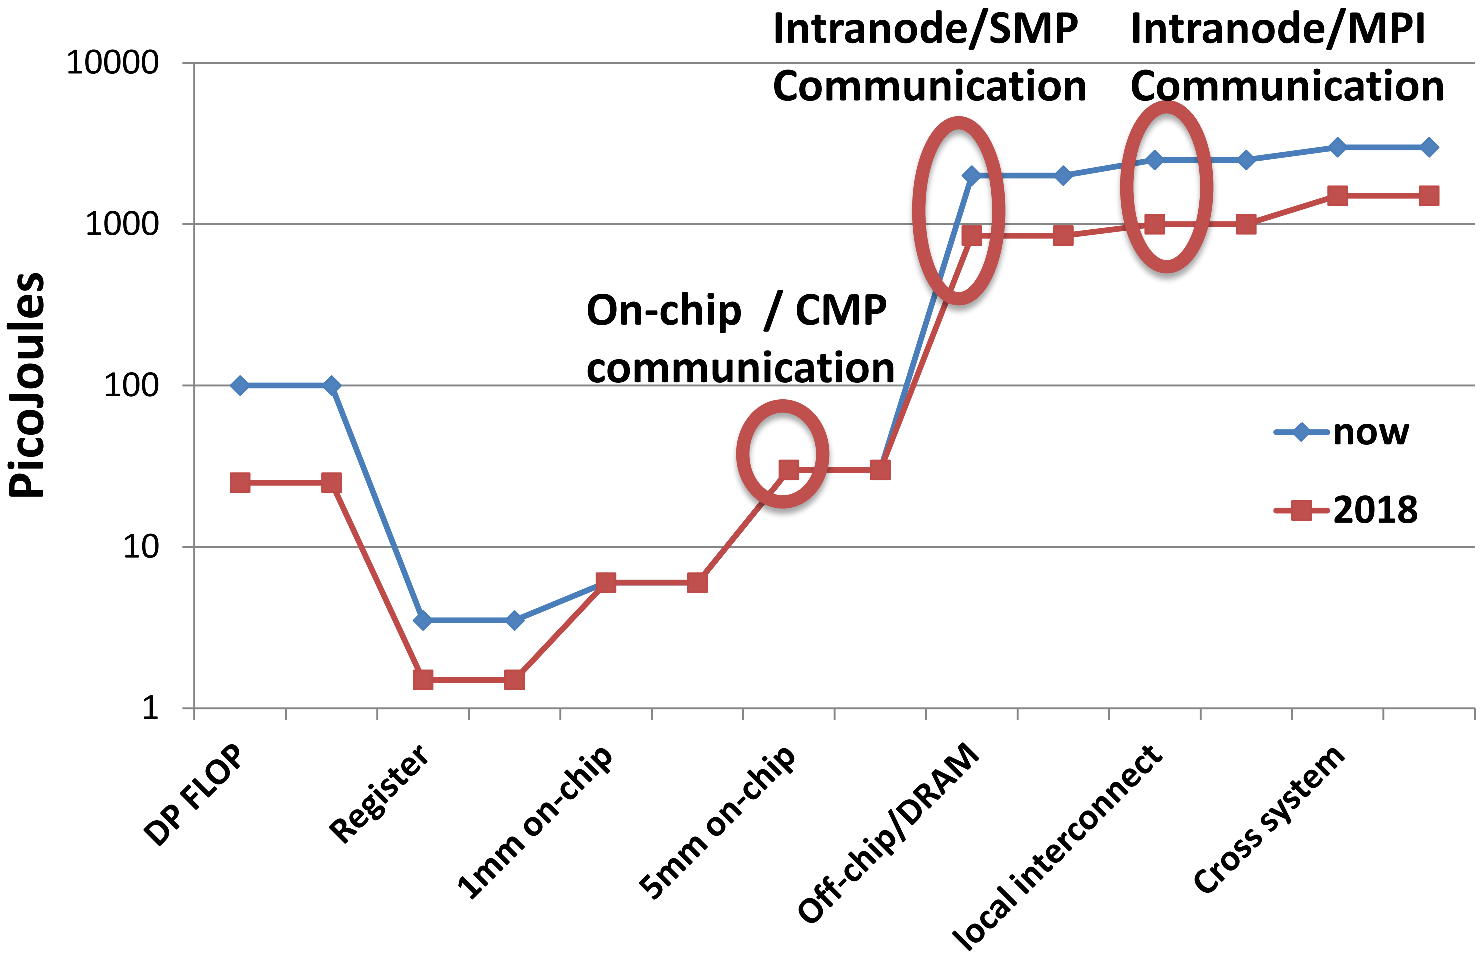

Memory Technology: Bandwidth costs power

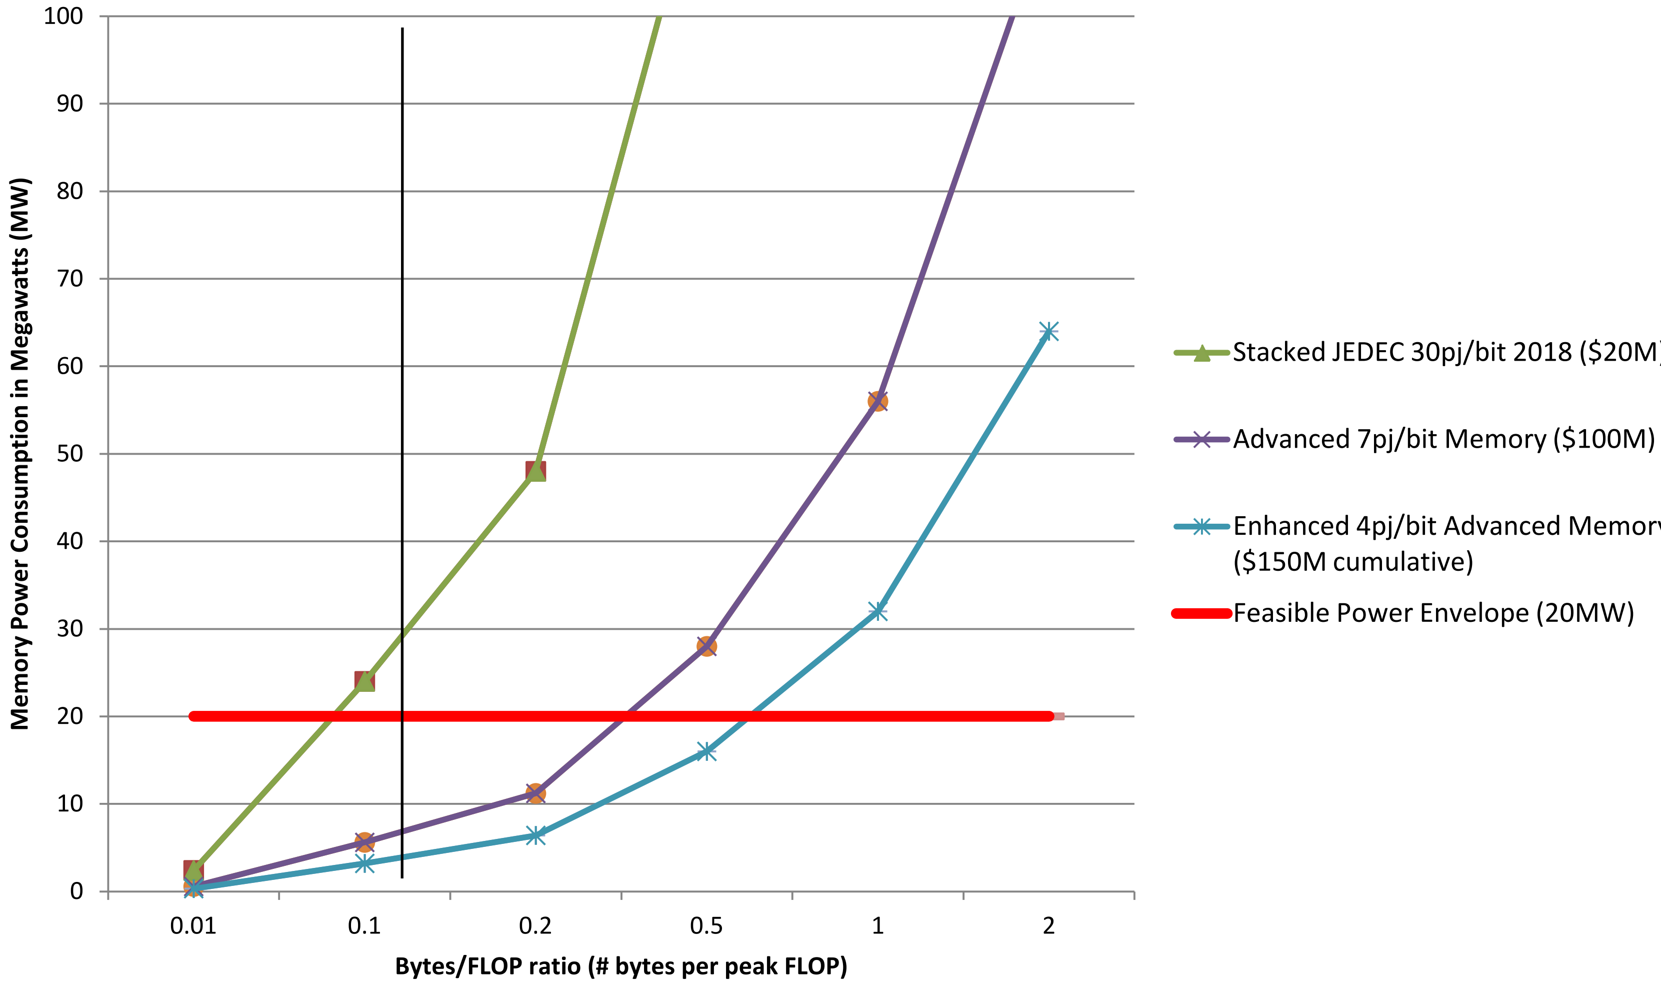

Approaches

- **Locality, Locality, Locality!**
- **Billion Way Concurrency;**
- **Uncertainty Quantification (UQ) including hardware variability;**
- **Flops free data movement expensive so:**
  - Remap multi-physics to put as much work per location on same die;
  - Include embedded UQ to increase concurrency;
  - Include data analysis if you can for more concurrency
  - Trigger output to only move important data off machine;
  - Reformulate to trade flops for memory use.
- **Wise use of silicon area**

Co-Design

Goals & Objectives

- **Understand how to allocate complexity between hardware, systems software, libraries, and applications;**
- **Modify application designs at all levels;**
- **Understand reformulating as well as reimplementing tradeoffs;**
- **Explore uncertainty quantification, in line data analysis, and resilience in applications;**
- **Co-adapt applications to new programming models and perhaps languages;**
- **Impact of massive multithreaded nodes and new ultra-lightweight operating systems.**

Awards- June 2011

Expected outcome- Understanding, Guidance for Future Applications, Application Readiness

Three Exascale Co-Design Centers Awarded

Exascale Co-Design Center for Materials in Extreme Environments (ExMatEx)

Director: Timothy Germann (LANL)

Center for Exascale Simulation of Advanced Reactors (CESAR)

Director: Robert Rosner (ANL)

Combustion Exascale Co-Design Center (CECDC)

Director: Jacqueline Chen (SNL)

	ExMatEx (Germann)	CESAR (Rosner)	CECDC (Chen)
National Labs	LANL	ANL	SNL
	LLNL	PNNL	LBNL
	SNL	LANL	LANL
	ORNL	ORNL	ORNL
		LLNL	LLNL
			NREL
University & Industry Partners	Stanford	Studsvik	Stanford
	CalTech	TAMU	GA Tech
		Rice	Rutgers
		U Chicago	UT Austin
		IBM	Utah
		TerraPower	
		General Atomic	
		Areva	

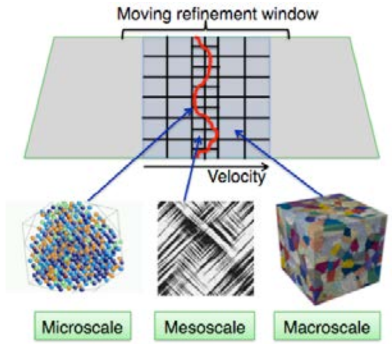

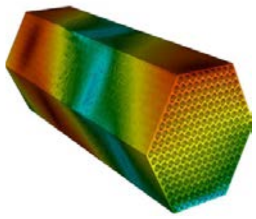

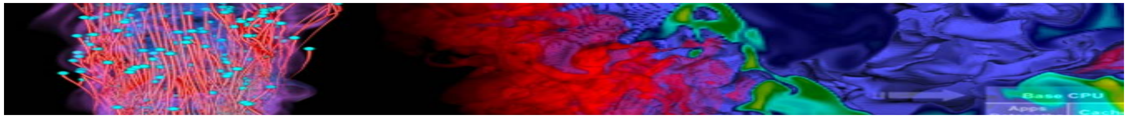

Future of Data Driven Science

- All of these hardware trends impact data driven science (in many cases more than compute intensive);
- Data from instruments still on 18-24 month doubling because detectors on CMOS feature size path;
- 100 gigabit per second per lambda networks on horizon;
- Disk read and write rates will fall further behind processors and memory;
- Significant hardware infrastructure needed to support this which probably will not be replicated at users' home institution (i.e. launching a petabyte file transfer at a users laptop is not friendly)

ASCR-BES Workshop on Data

"Data and Communications in Basic Energy Sciences: Creating a Pathway for Scientific Discovery"

October 24-25, 2011 Bethesda, MD

Goals & Objectives

- **Identify and review the status, successes, and shortcomings of current data (including analysis and visualization) and communication pathways for scientific discovery in the basic energy sciences;**
- **Ascertain the knowledge, methods and tools needed to mitigate present and projected data and communication shortcomings;**
- **Consider opportunities and challenges related to data and communications with the combination of techniques (with different data streams) in single experiments;**
- **Identify research areas in data and communications needed to underpin advances in the basic energy sciences in the next ten years;**
- **Create the foundation for information exchanges and collaborations among ASCR and BES supported researchers, BES scientific user facilities and ASCR computing and networking facilities.**

Co-Chairs

- **Peter Nugent, NERSC**
- **J. Michael Simonson, SNS**

ASCR at a Glance

Office of Advanced Scientific Computing Research

Associate Director – Daniel Hitchcock (acting)
Phone: 301-903-7486
E-mail: Daniel.Hitchcock@science.doe.gov

Research

Division Director – William Harrod
Phone: 301-903-5800
E-mail: William.Harrod@science.doe.gov

Facilities

Division Director – Daniel Hitchcock
Phone: 301-903-9958
E-mail: Daniel.Hitchcock@science.doe.gov

Relevant Websites

ASCR: science.energy.gov/ascr/
ASCR Workshops and Conferences:
science.energy.gov/ascr/news-and-resources/workshops-and-conferences/
SciDAC: www.scidac.gov
INCITE: science.energy.gov/ascr/facilities/incite/
Exascale Software: www.exascale.org
DOE Grants and Contracts info: science.doe.gov/grants/

Background

Computational Science Graduate Fellowship

www.krellinst.org/csgf

20 YEARS BUILDING A COMMUNITY OF LEADERS
DOE CSGF 1991-2011

- **Funded by the Department of Energy's Office of Science and National Nuclear Security Administration, the DOE CSGF trains scientists to meet U.S. workforce needs and helps to create a nationwide interdisciplinary community.**
- **Some Talks from 2011 Fellows Conference**
  - *"The Materials Genome: An online database for the design of new materials for clean energy and beyond"*
  - *"Understanding electromagnetic fluctuations in microstructured geometries"*
  - *"Computational science meets materials chemistry: a bilingual investigation of surface effects in nanoscale systems"*

NERSC

Hopper in production April 2011

- **1.25 PFlop/s peak performance Cray XE6**
- **Over a billion core-hours to science per year**
- **Gemini high performance resilient interconnect**

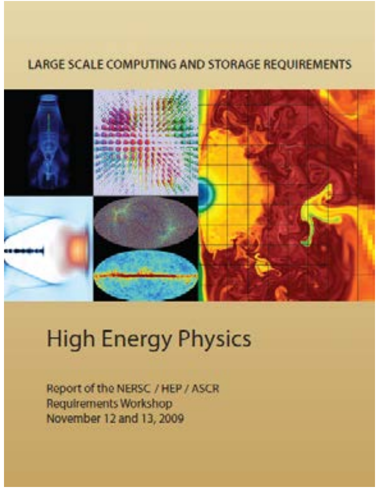

- **HEP NERSC Requirements Workshop :**
  - www.nersc.gov/assets/HPC-Requirements-for-Science/NERSC-HEP-WorkshopReport.pdf

INCITE

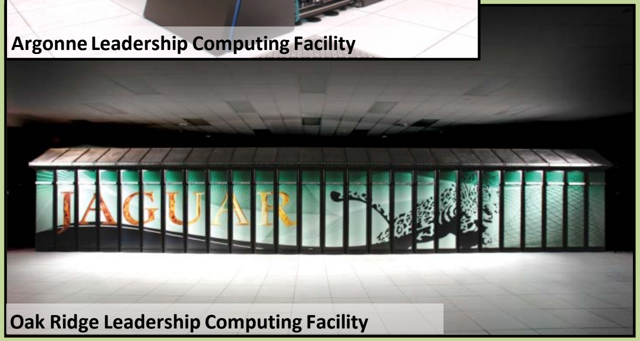

In FY 2012, the Argonne LCF will be upgraded with a 10 petaflop IBM Blue Gene/Q. The Oak Ridge LCF will continue site preparations for a system expected in FY 2013 that will be 5-10 times more capable than the Cray XT-5.

- The Cray XT5 ("Jaguar") at ORNL and the IBM Blue Gene/P ("Intrepid") at ANL will provide ~2.3 billion processor hours in FY12 to address science and engineering problems that defy traditional methods of theory and experiment and that require the most advanced computational power.
- Peer reviewed projects are chosen to advance science, speed innovation, and strengthen industrial competitiveness.
- Demand for these machines has grown each year, requiring upgrades of both.

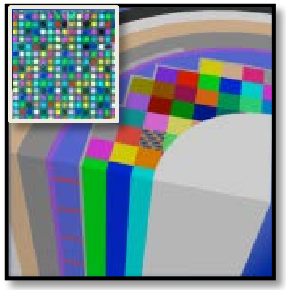
Nuclear Reactor Simulation

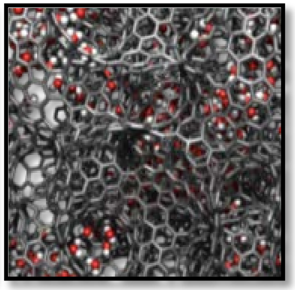
Energy Storage Materials

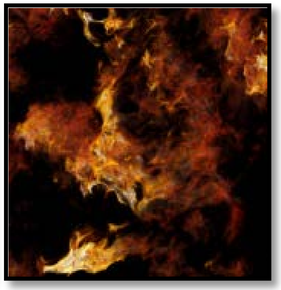
Turbulence

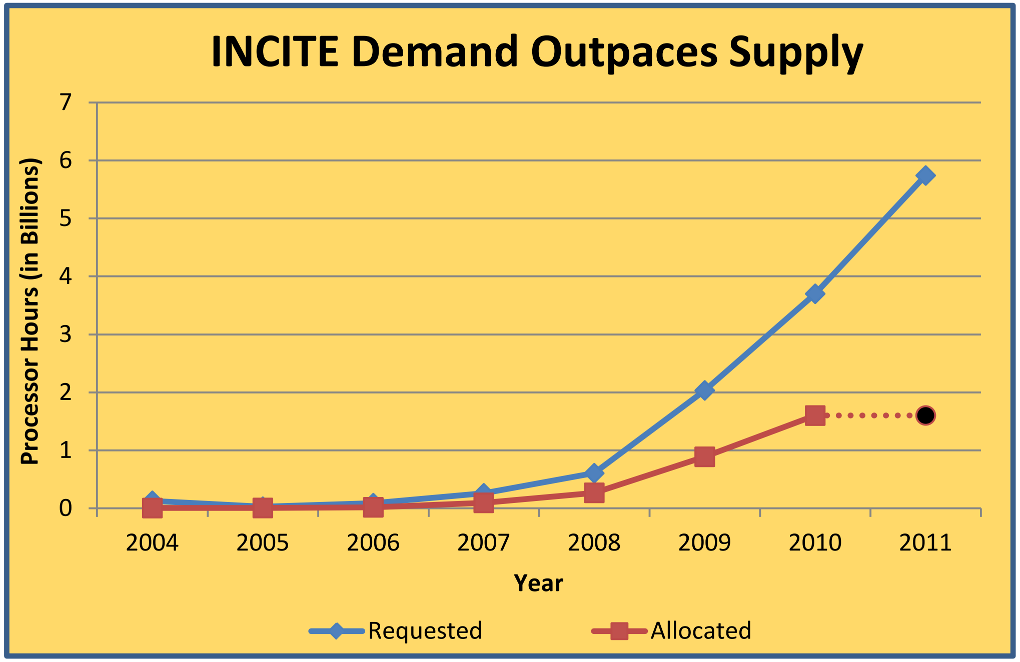

Argonne Leadership Computing Facility

ALCF - now

- more CPUs/node (x 4)
- faster CPUs (x 60)
- faster storage (x 2.7)
- more storage (x 3.6)
- more RAM/CPU (x 2)

ALCF - 2013

Intrepid

Blue Gene/P—peak **557 TF**
40 racks
40,960 nodes (quad core)
163,840 processors (3.4 GF Peak)
80 TB RAM
8 PB storage capacity
88 GB/s storage rate
1.2 MW power
Air-cooled

Mira

Blue Gene/Q—peak **10 PF**
48 racks -- *just one row more than Intrepid*
49,152 nodes (16 core)
786,432 processors (205 GF peak)
786 TB RAM
28.8 PB storage capacity
240 GB/s storage rate
Estimated 3-8 MW projected power
Water-cooled

~20 times the computing power in only 20% more space

Oak Ridge Leadership Computing Facility

- OLCF-3
  - Builds on DARPA HPCS investment in Cray
  - Prepare applications for path to Exascale computing
  - CD-1 approved, December 2009
  - Lehman CD2-3b Review scheduled for August, 2010

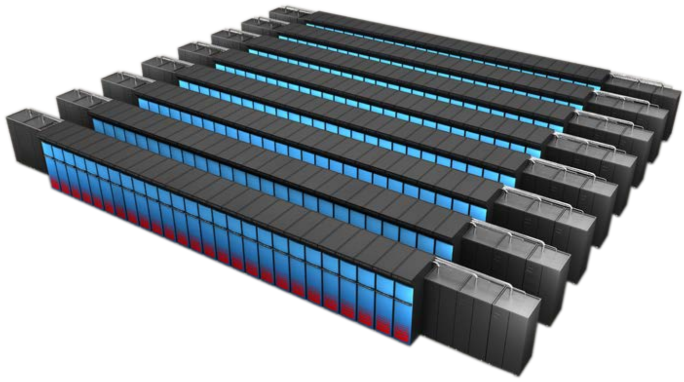

Science Application Readiness

- **Application Readiness Review August, 2010**
- **Science Teams**
- **Port, Optimize, New Modeling Approaches**
- **Early Access to development machine**

Cray-NVIDIA System

- **~10-20 Pflop/s peak**
- **Hybrid x86/accelerator architecture**
- **Liquid cooled**
- **Highly power efficient**

ESnet

- HEP 2009 Requirements Workshop
  - www.es.net/assets/Papers-and-Publications/HEP-Net-Req-Workshop-2009-Final-Report.pdf
- Electronic Collaboration Tools
  – www.ecs.es.net
- Network Performance Tools
  – fasterdata.es.net
- Virtual Circuits
  – www.es.net/services/virtual-circuits-oscars/