U.S. DEPARTMENT OF ENERGY

OFFICE OF SCIENCE

Advanced Scientific Computing Research Program

The View from Germantown

ALCF Getting Started Workshop
January 27-29, 2010

Betsy Riley
Advanced Scientific Computing Research
Office of Science, U.S. Department of Energy
Betsy.Riley@science.doe.gov

U.S. DEPARTMENT OF ENERGY

DOE Organizational Structure

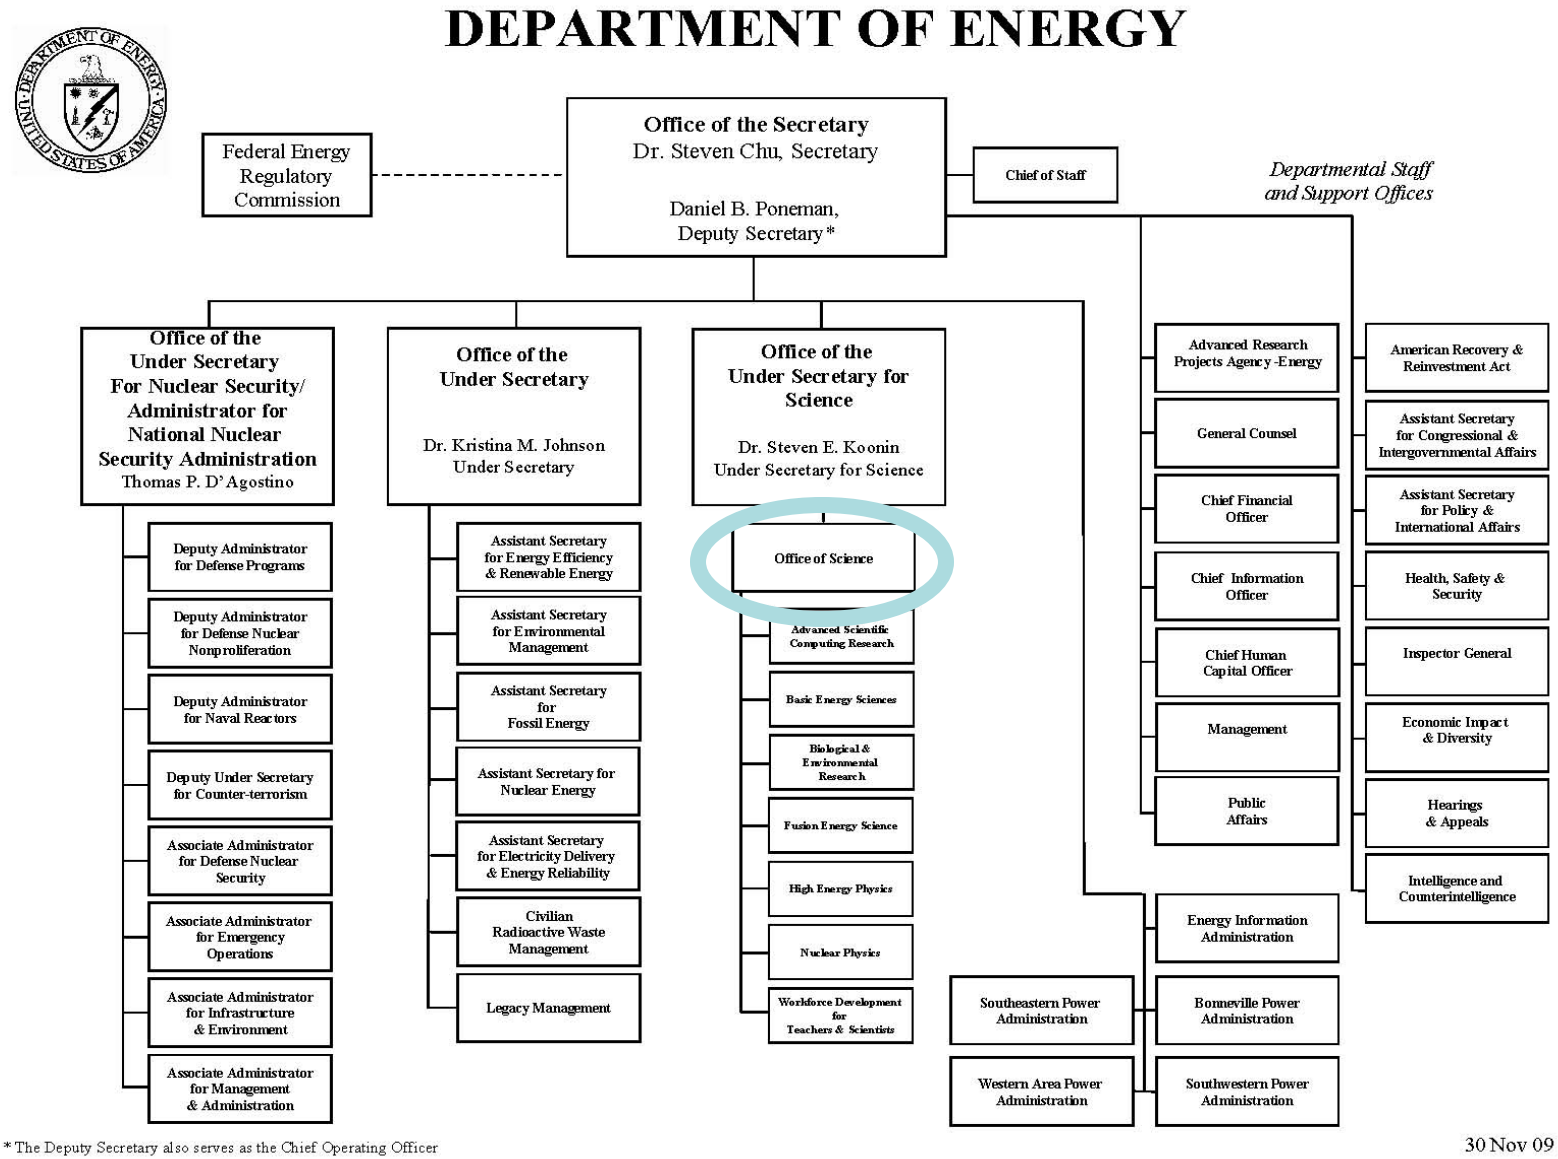

U.S. DEPARTMENT OF ENERGY

Office of Science

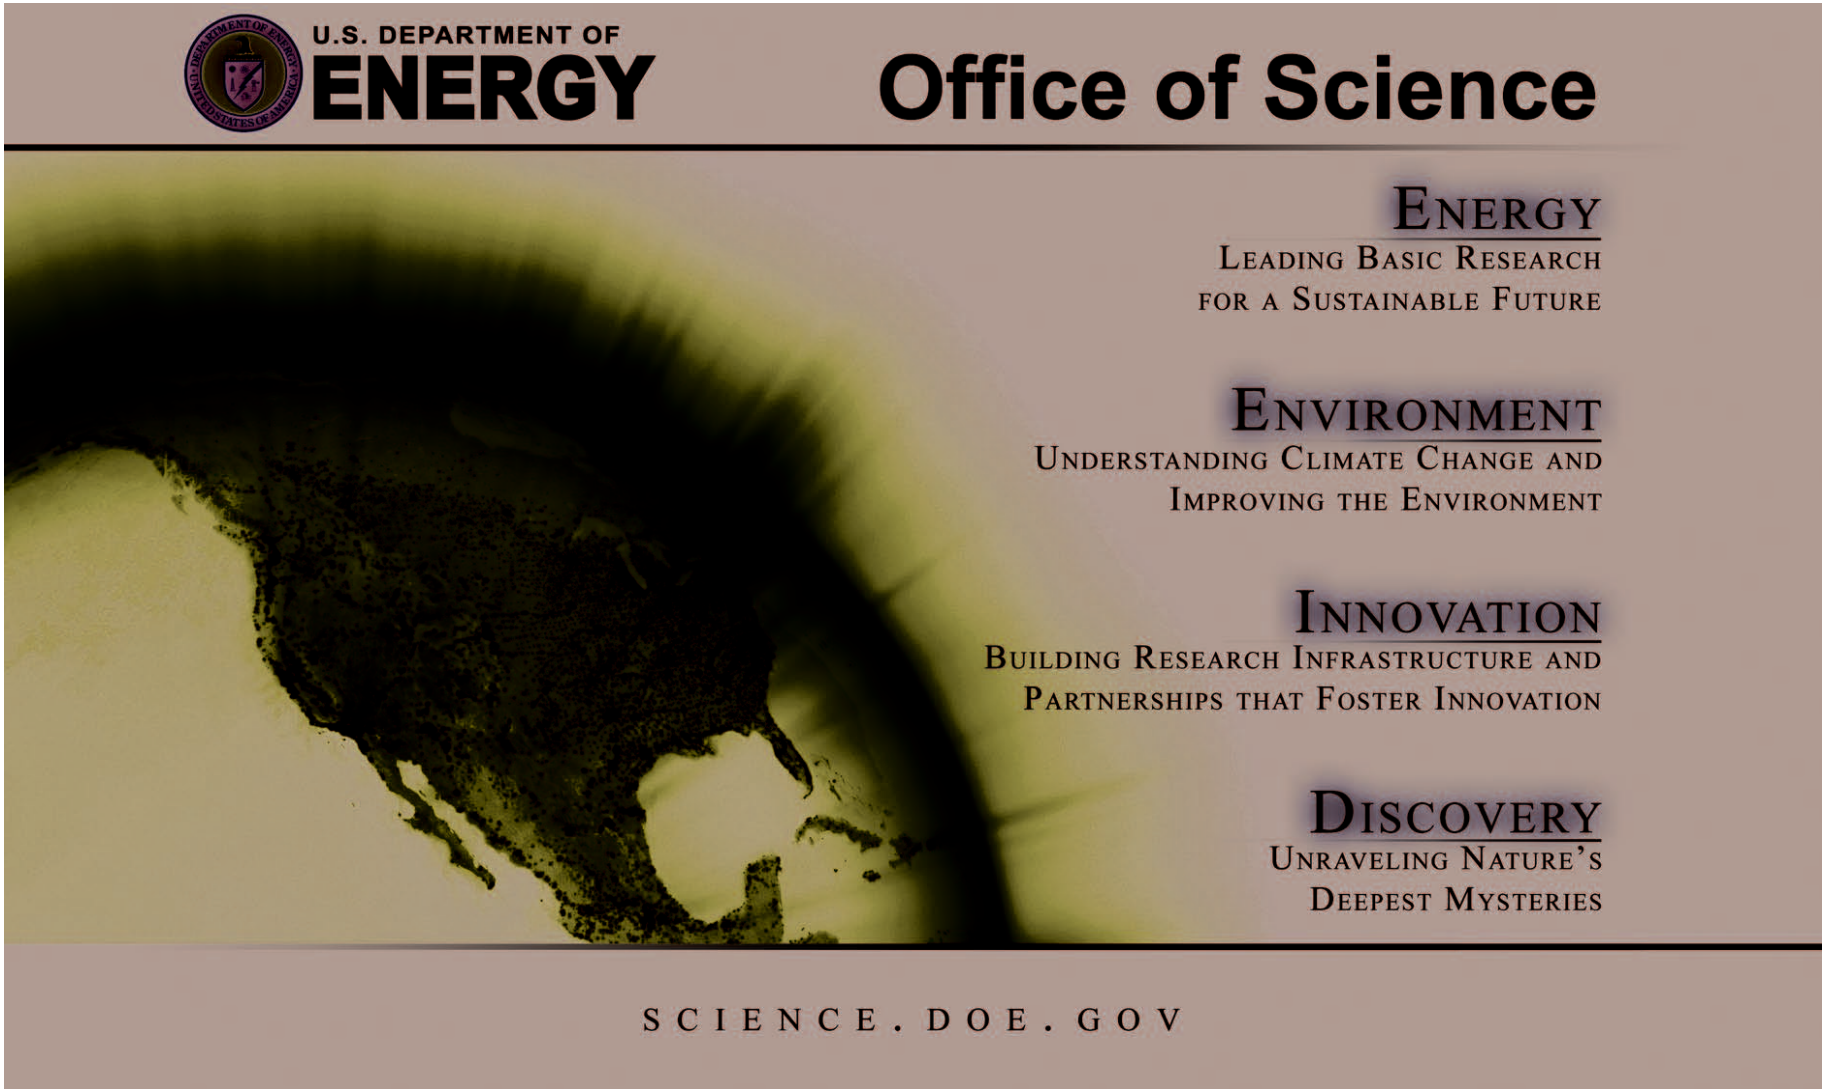

Office of Science:

Advanced Scientific Computing Research (ASCR)

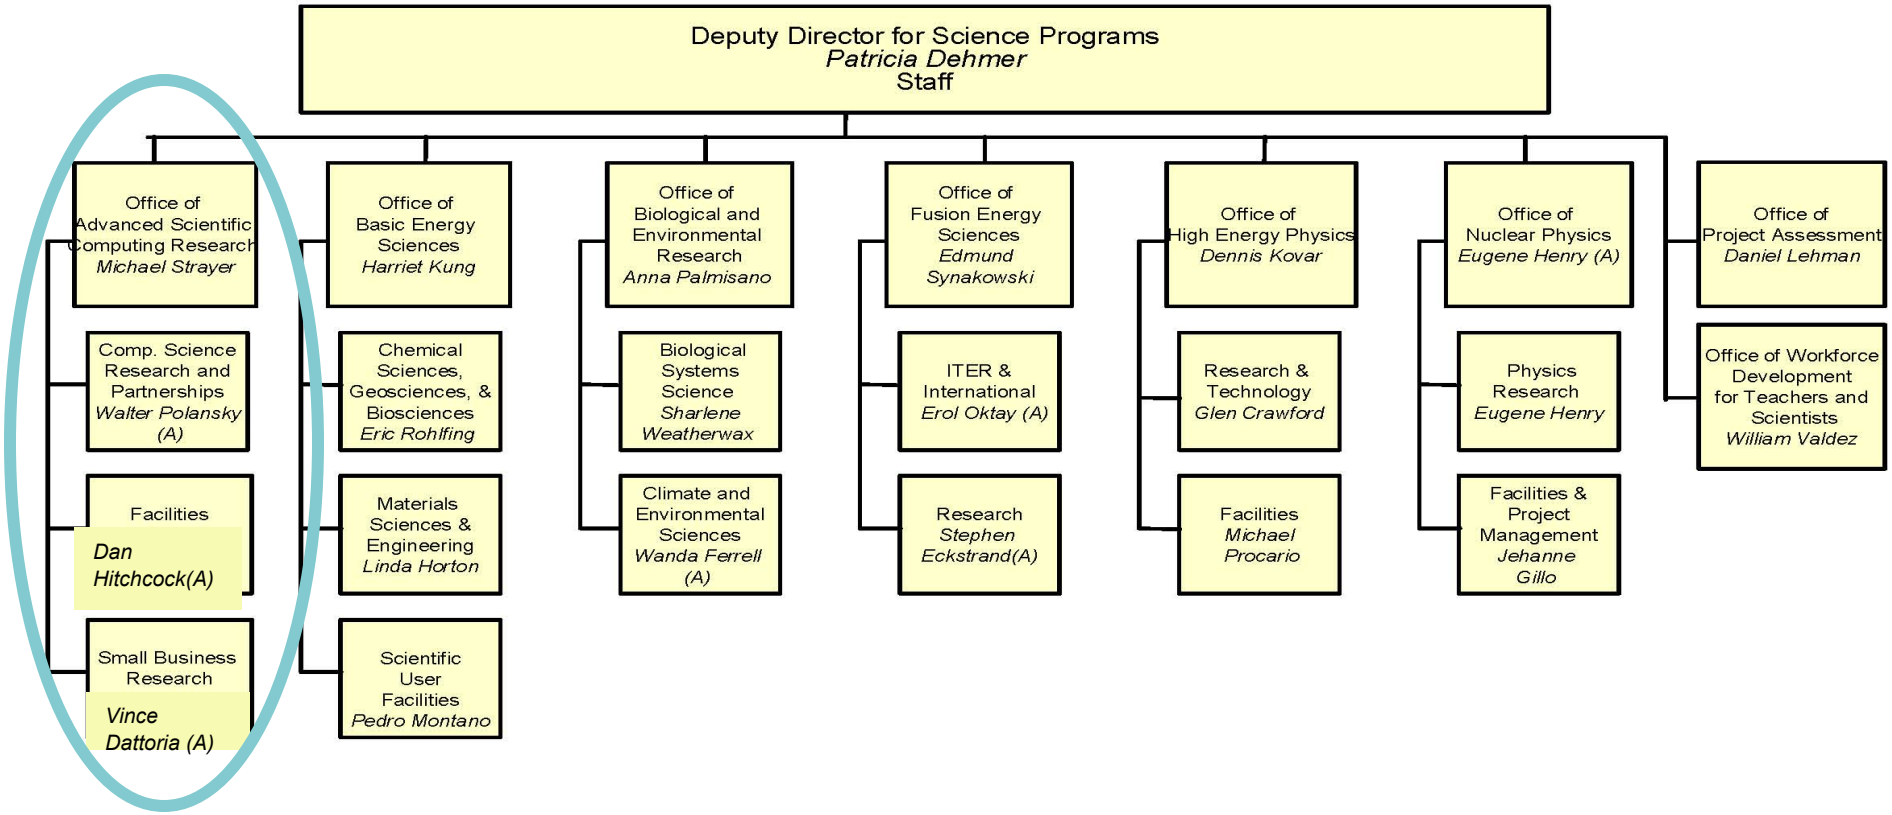

U.S. DEPARTMENT OF ENERGY

ASCR Facilities Division

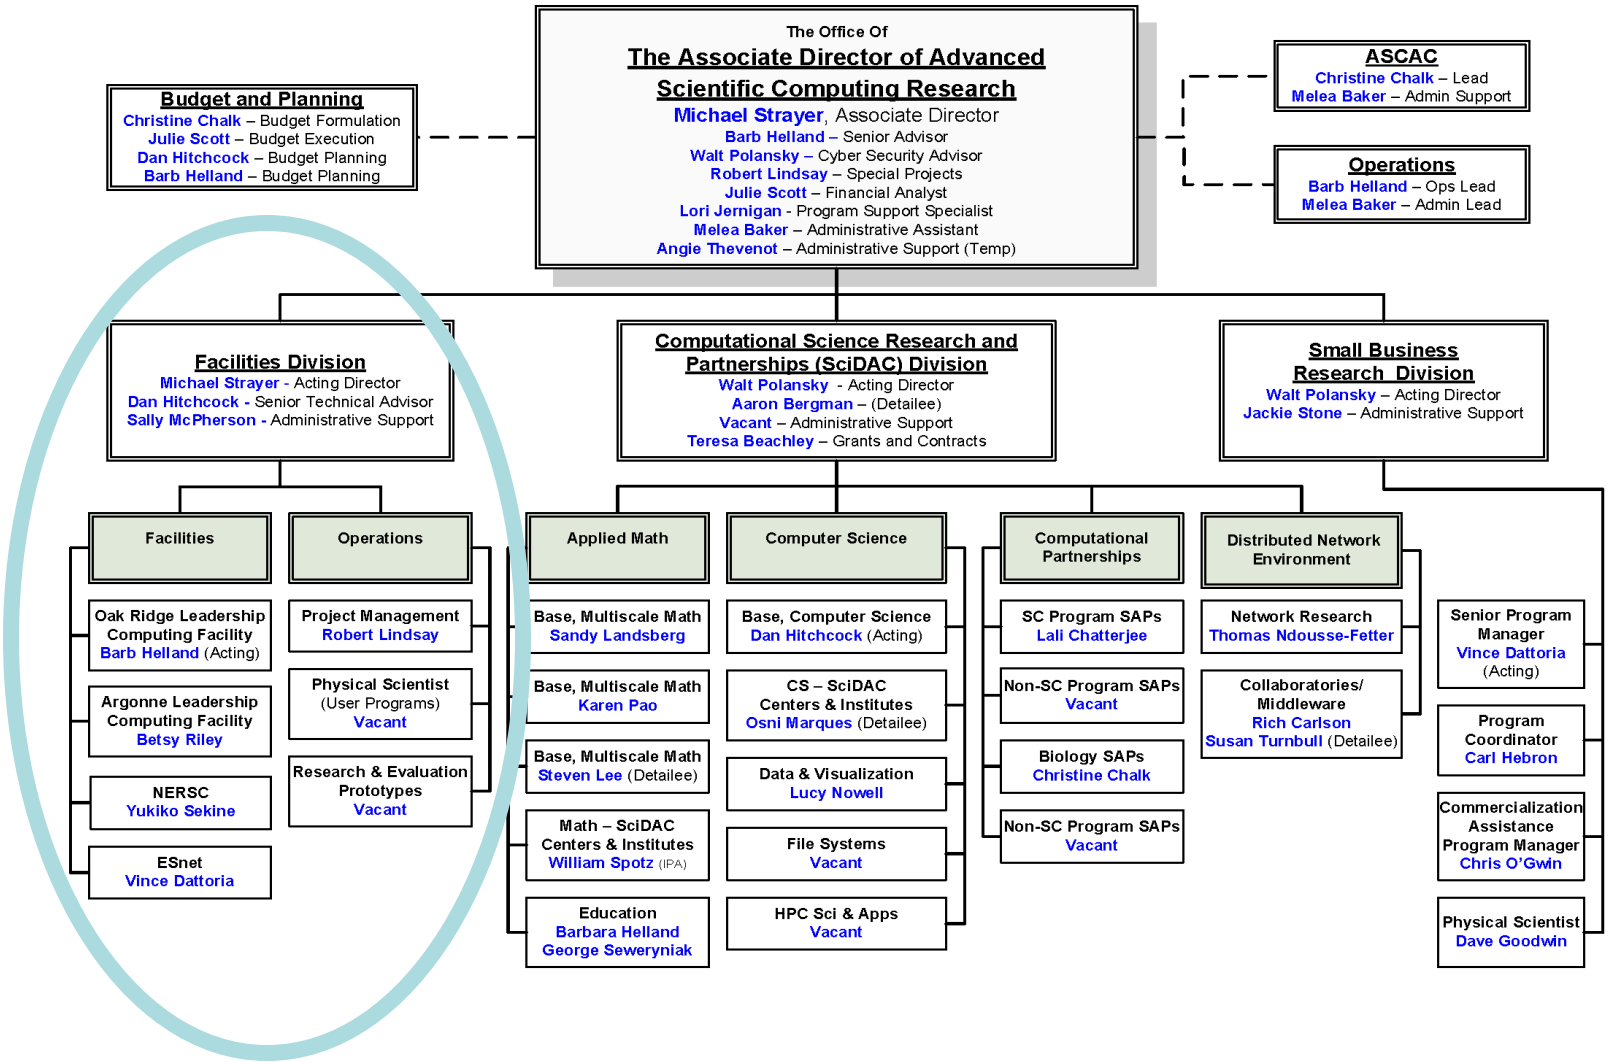

ASCR Facilities Strategy

- **Providing the Tools – High-End Computing**
  - **High-Performance Production Computing –** National Energy Research Scientific Computing Center **(NERSC)** at Lawrence Berkeley National Laboratory (LBNL)
    - Delivers high-end capacity computing to entire DOE SC research community
  - **Leadership-Class Computing –**
    Leadership Computing Facilities (LCFs)
    -Argonne National Laboratory (ANL -- ALCF)
    -Oak Ridge National Laboratory (ORNL -- OLCF)
    - Deliver highest computational capability to national and international researchers through peer-reviewed **Innovative and Novel Computational Impact on Theory and Computation (INCITE)** program (80% of resources)
- **Investing in the Future - Research and Evaluation Prototypes**
- **Linking it all together – Energy Sciences Network (ESnet)**

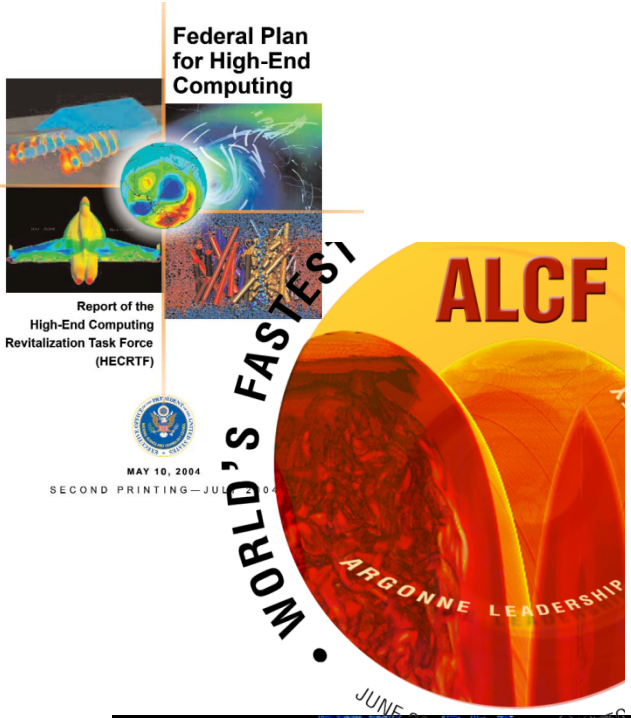

U.S. DEPARTMENT OF ENERGY

Facilities Update

Next Generation of LCFs

Mission need approved January 2009

"The upgrade of the Leadership Computing Facilities to tens of petaflops by the 2011-2013 timeframe is vital to the U.S. playing a leading role in several important international programs, including:

- ***climate science (International Panel on Climate Change),***
- ***fusion energy research (ITER), and***
- ***the Nuclear Energy Advanced Modeling and Simulation (NEAMS) program"***

- ORNL's Cray XT5 upgraded to 2.3 PF, increasing allocatable hours by 50%
- New NERSC-6 Cray System "Hopper" will be >1PF in 3Q10
- ALCF-2 upgrade "Mira" Blue Gene/Q will be 10PF in FY12

ALCF-2 Upgrade System "Mira"

- Enables key science impact:
  - Predict abrupt regional climate change
  - Design safer, cost-effective nuclear power reactors
  - Enhancement of the extraction of biofuels from biomass
  - In silico design of nano-structured storage systems
- Builds on ASCR/NNSA investment and LLNL BG/Q Sequoia competitive bid procurement

Science Enablement Program

- 15 Science Teams
- Call: 1Q10, 1st Selection: 3Q10
- Port, Optimize, New Modeling Approaches
- Early Access to Pre-Production hardware

Mira Blue Gene/Q System

- 10 Pflop/s peak
- ~800K cores, 16 per chip
- ~70PB disk, ~470 GB/s I/O bandwidth
- Power efficient, water cooled

4Q10	1Q11	2Q11	3Q11	4Q11	1Q12	2Q12	3Q12	4Q12

U.S. DEPARTMENT OF ENERGY

Innovative and Novel Computational Impact on Theory and Experiment- INCITE

- **Initiated in 2004 at NERSC**
- **Now provides Office of Science Leadership Class Computing resources to a small number of computationally intensive research projects of large scale, that can make high-impact scientific advances through the use of a large allocation of computer time and data storage**
- **Open to national and international researchers, including industry**

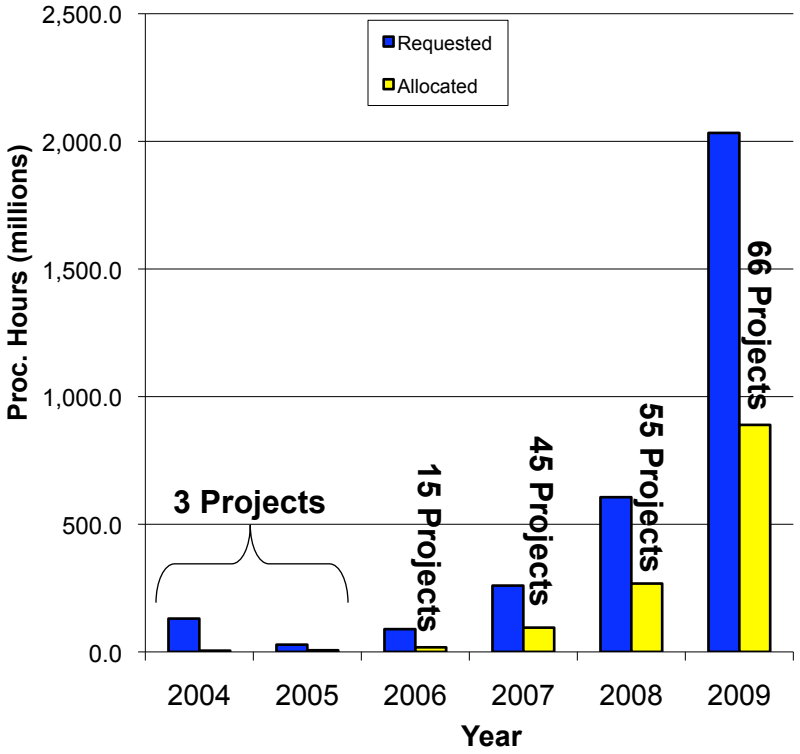

- **No requirement of DOE Office of Science funding**
- **Peer and computational readiness reviewed**

U.S. DEPARTMENT OF ENERGY

ALCF: IBM Blue Gene/P—557 TFs

Intrepid: 40,960 quad-core compute nodes (163,840 proc.),80TB of memory

Expectations for INCITE

- **From ALCF:**
  - Sign ALCF User Agreements
  - Adhere to ALCF Cyber Security and other computing policies
  - Work with the Catalyst, Performance Engineering, and Viz & Data Analysis groups to take advantage of the IBM Blue Gene architecture
  - Don't wait to start computing
- **From DOE:**
  - Share your scientific accomplishments!

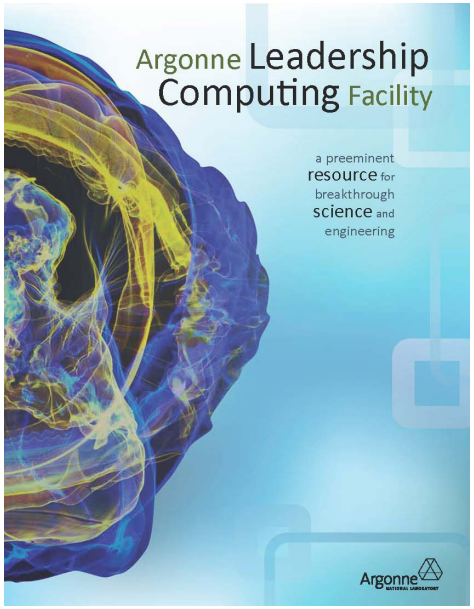

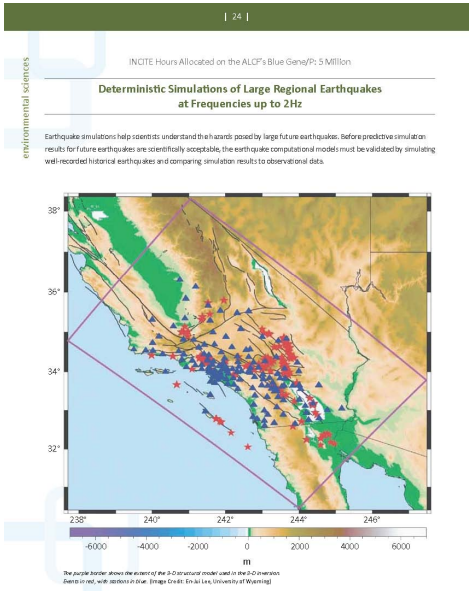

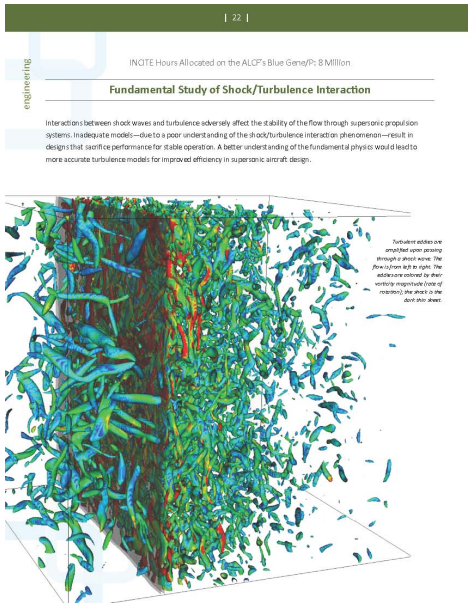

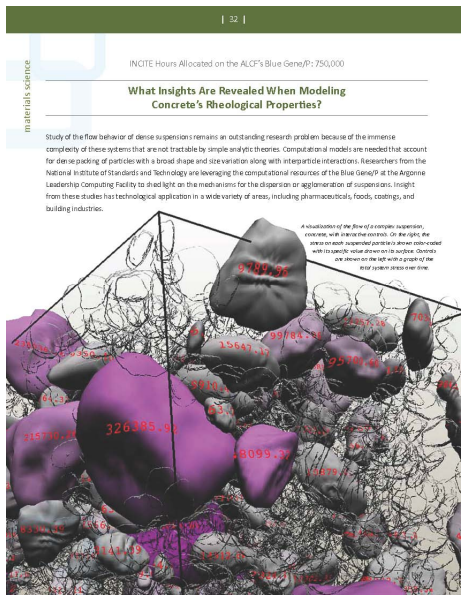

U.S. DEPARTMENT OF ENERGY

Thousand Atom Nanostructures

DD

Lin-Wang Wang
LBL

Science

- **Design better materials for products including solar cells**
- ***Ab initio* electronic structure calculations**
- **Lin-Wang Wang, B. Lee, H. Shan, Z. Zhao, J. Meza, E. Strohmaier, D. Bailey, "Linear Scaling Divide-and-conquer Electronic Structure Calculations for Thousand Atom Nanostructures," SC08, to appear.**

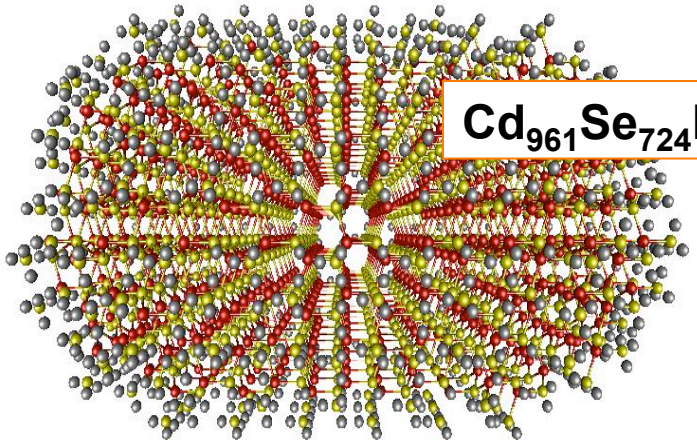

Dipole moment calculated on 2633 atom quantum rod

Methods and Challenges

- **Novel divide & conquer approach to solve DFT but reducing $O(n^3)$ to O(n)**
  - Many months to 30 hours
  - Direct DFT impractica
- **Mapping critical**
  - Linear scaling to 160K cores and a 10% improvement in

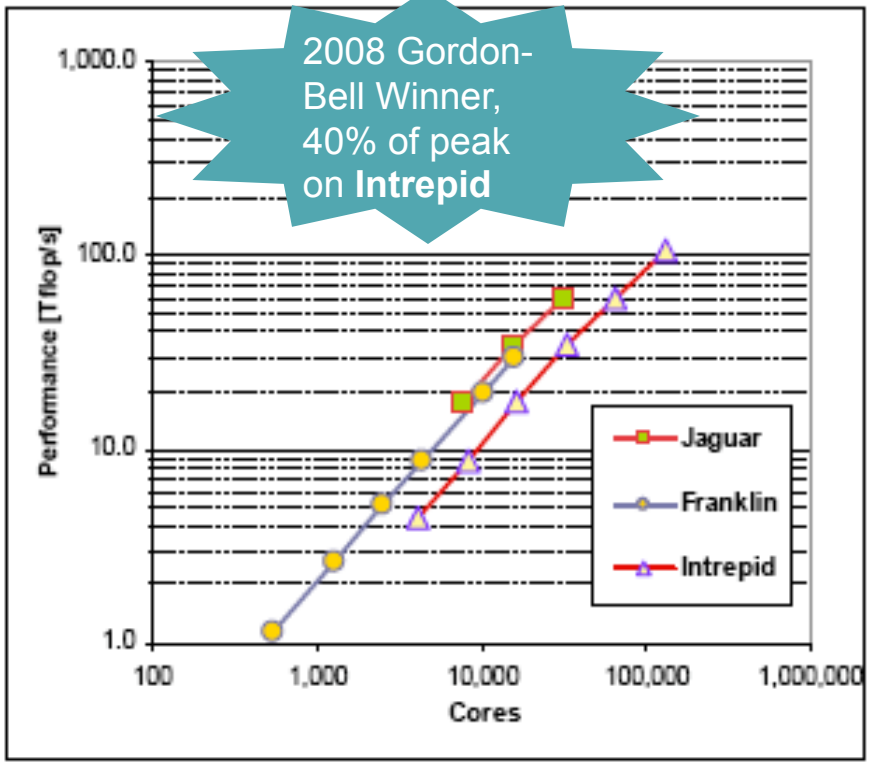

Identifying Potential Drug Targets

Michael Wilde
ANL

Science

- **Reduce dead ends in antibiotics and anticancer drugs with DOCK5 and DOCK6**
  - 9 enzymatic proteins in core metabolism of bacteria and humans screened against 15,351 natural compounds and existing drugs
  - Study correlations and re-prioritize proteins for further study
- **Able to complete 21.43 CPU-years of analysis in 2.01 wall-hours**

•MCS Computer Science teams, using discretionary allocation, to facilitate science
•Participated in submitting 2 INCITE proposals

Methods and Challenges

- **Port of framework, Falkon, to manage run**
- **Falkon requires non-standard BG/P kernel (ZeptoOS)**
- **Huge demand on I/O system as each core is controlling multiple files**
- **118,000 cores were used running nearly one million tasks**

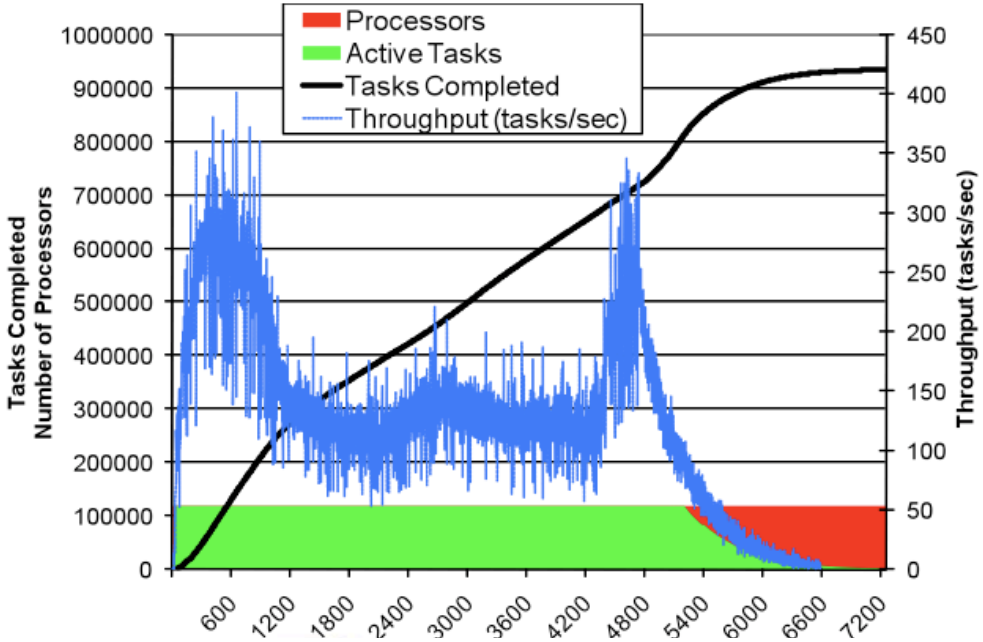

U.S. DEPARTMENT OF ENERGY

Gating Mechanism of Membrane Proteins

Benoit Roux
ANL, University of Chicago

Science

- Understand how proteins work so we can alter them to change their function: micromachines
- Validated the atomic models of Kv1.2 and first to calculate the gating charge in the two functional states

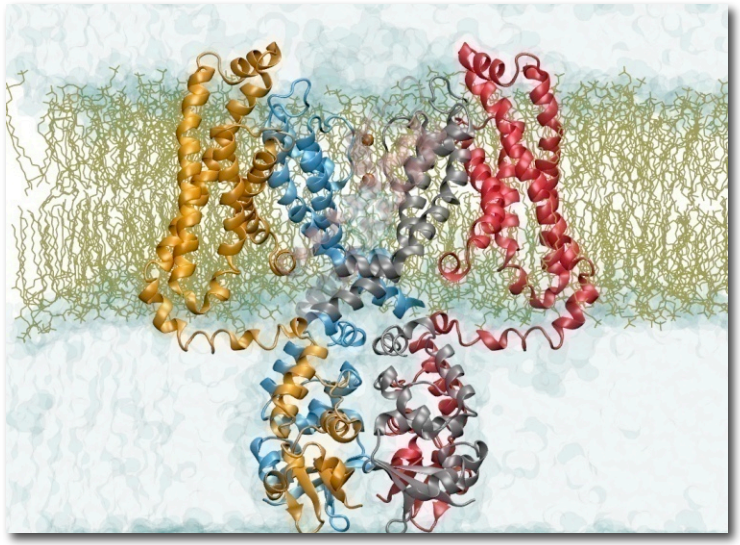

Methods and Challenges

- NAMD with periodicity and particle-mesh Ewald method

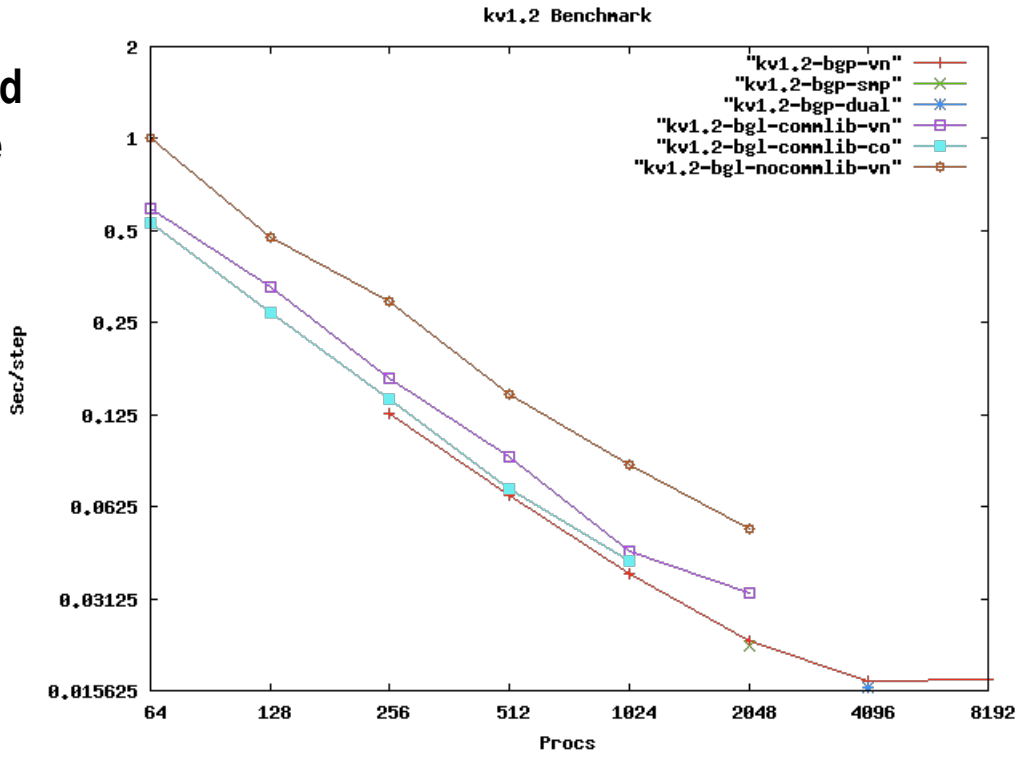

kv1.2 Benchmark (352K atoms)

15-20% gain over BG/L customized ("commlib") version.

What are YOUR Accomplishments?