

Breakthrough Science at NERSC

Richard Gerber

NERSC User Services Group

Cray Technical Workshop, Isle of Palms, SC

February 25, 2009





Outline

- **Overview of NERSC**
- **NERSC's Cray XT4 - Franklin**
- **Franklin 2008-2009 Usage**
- **Breakthrough Science**



NERSC

NERSC

The *National Energy Research Scientific Computing Center*

NERSC's Mission

Accelerate the pace of scientific discovery for all DOE Office of Science research

NERSC Provides High Performance

Computing, information, data, and communications services

NERSC 2009 Configuration

Large-Scale Computing System

Franklin: Cray XT4

- 9,660 nodes
- 38,640 cores
- 79 TB Memory
- 400+ TB shared disk
- SeaStar 2 3D torus interconnect

95% of NERSC
compute capacity



Clusters



Bassi

- IBM Power5 (888 cores)

Jacquard

- LNXI Opteron (712 cores)

PDSF (HEP/NP)

- Linux cluster (~1K cores)

NERSC Global Filesystem (NGF)



230 TB; 5.5 GB/s

HPSS Archival Storage

- 44 PB capacity
- 10 Sun robots
- 130 TB disk cache



Analytics / Visualization Davinci SGI Altix



NERSC's Cray XT4

- **Franklin**
 - 102 Cabinets in 17 rows
 - 9,660 nodes (38,640 cores)
 - 79 TBs Aggregate Memory (4 x 2GB DIMMs per node)
- **Interconnect: Cray SeaStar2, 3D Torus**
 - >6 TB/s Bisection Bandwidth
 - >7 GB/s Link Bandwidth
- **Shared Disk: 400+ TBs**
- **Network Connections**
 - 24 x 10 Gbps + 16 x 1 Gbps
 - 60 x 4 Gbps Fibre Channel
- **Performance:**
 - Sustained application: 38 Tflops
 - Peak performance: 355 Tflops
 - Linpack: 266 Tflop/sec






NERSC is the Production Facility for all of DOE SC

- **NERSC serves a large and diverse community**
 - ~ 3,000 users
 - ~ 400 projects
 - ~ 500 codes
 - ~ 100 institutions (primarily DOE labs, US Univ)
- **DOE allocates NERSC resources across all SC research areas**
 - High impact through large awards
 - Broad impact across science domains


2009 Allocations by DOE Office

Office	XT4 hours	Num projects
Adv Sci Comp Res	12.1 M	46
Bio Environ Res	35.4 M	62
Basic Energy Sci	57.2 M	163
Fusion Energy Sci	32.6 M	60
High Energy Phy	30.5 M	32
Nuclear Phy	18.0 M	31

Percent of Allocation




















Percent of Projects



- ASCR
- BER
- BES
- FES
- HEP
- NP

08/09 Franklin Use by Science Field



- | | | |
|---|--|---|
|  <i>Materials Sciences</i> |  <i>Climate Research</i> |  <i>Fusion Energy</i> |
|  <i>Lattice Gauge Theory</i> |  <i>Chemistry</i> |  <i>Combustion</i> |
|  <i>Accelerator Physics</i> |  <i>Astrophysics</i> |  <i>Other</i> |
|  <i>Life Sciences</i> |  <i>Nuclear Physics</i> |  <i>Applied Math</i> |
|  <i>Geosciences</i> |  <i>Computer Sciences</i> |  <i>Environmental Sciences</i> |
|  <i>Engineering</i> | | |



Scientific Publications

- **1,487 refereed publications based on research conducted using NERSC in mid-2007 to mid-2008**
- **Examples**
 - "Improving PEM Fuel Cell Catalyst Activity and Durability Using Nitrogen-Doped Carbon Supports," Y. Zhou, R. Pasquarelli, J. Berry, D. Ginley, T. Holme, R. O'Hayre, *Nature Materials*
 - Genomics of cellulosic biofuels, Rubin, EM, *Nature*, 454 (7206): 841-845 AUG 14 2008
 - Questioning the existence of a unique ground-state structure for Si clusters, W. Hellmann, R. G. Hennig, S. Goedecker, C. J. Umrigar, B. Delley and T. Lenosky. *Phys. Rev. B* 75, 085411



Solving Real-World Problems Using Computing

- **DOE and NERSC have research on all aspects of the energy & climate change problems**
 - **Energy sources: solar, bio, nuclear,...**
 - **Energy production: catalysts**
 - **Energy efficient engines,...**
 - **Carbon sequestration**
 - **Climate change**



Breakthrough Science




Climate Change

How do you validate climate models?

Compare simulations to observations going back as far as you can.

But what if the observed data is incomplete and unreliable?





Surface Input Reanalysis for Climate Applications

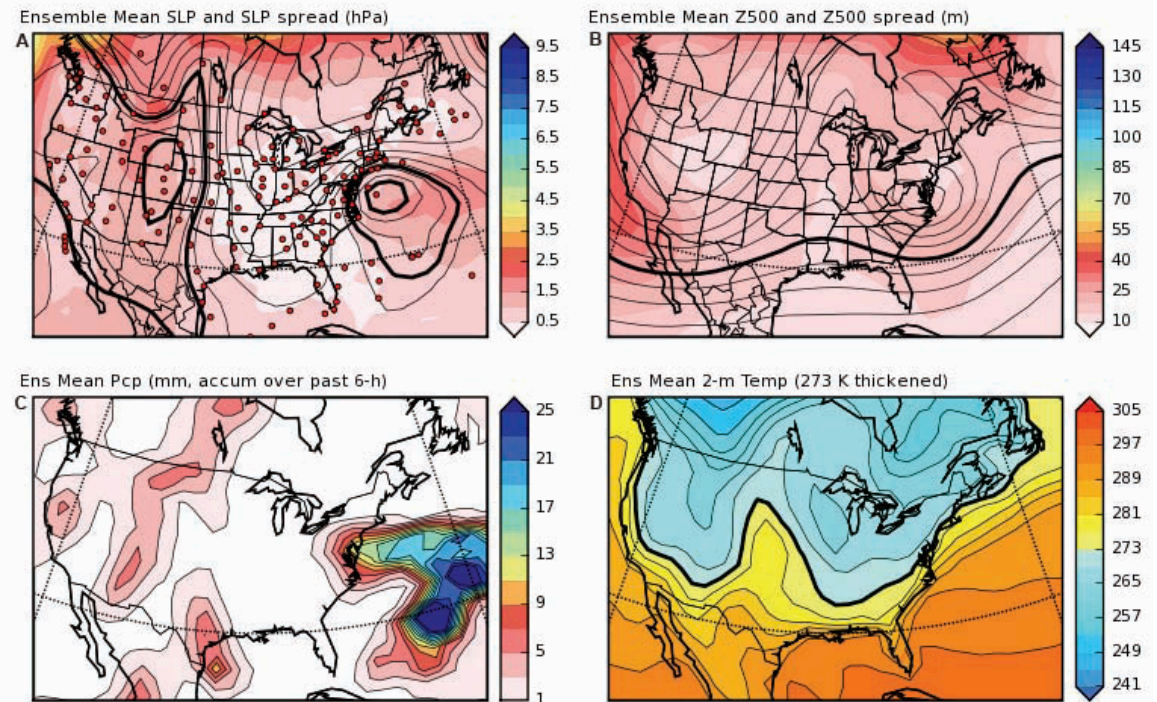
- **The SIRCA project, lead by Gil Compo (U. Colo) seeks to create the first complete 3D data set of lower atmospheric pressure from 1871-2009.**
- **Input: error-prone manual sea-level measurements in the N. Hemisphere and random account of storms, etc.**
- **Apply a recursive filter and a weather prediction model to produce 2 degree (0.5 degree hires), 6-hour, 3D datasets**
- **To date: First complete dataset of atmospheric observations for 1908 to 1948, when modern observations began.**

Validating Climate Models

Data will also enable investigation of

- 1878 El Nino of the Century
- Dust Bowl of the 1930s
- Arctic warming of the 1920s to 1940s

11 M hours (2008-09)
5,000 cores max
2,000 core median



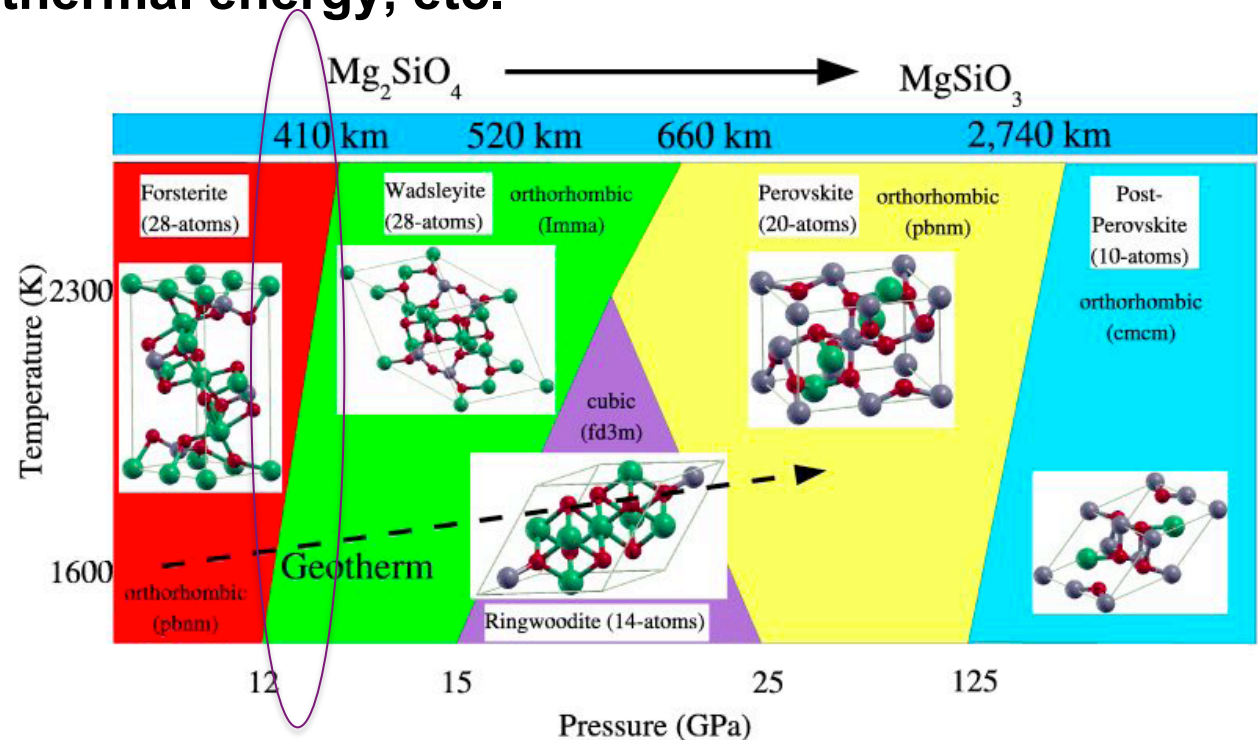
Sea level pressures with color showing uncertainty (a&b); precipitation (c); temperature (d). Dots indicate measurements locations (a).

Materials Science/Geophysics

Kevin Driver, John Wilkins (Ohio State)

Understanding how seismic waves propagate through silicates and reflect from boundaries is important for understanding sub-surface structures - implications for locating oil, carbon sequestration, geothermal energy, etc.

A seismic discontinuity at 410 km depth is attributed to a phase transition in magnesium silicate





Quantum Monte Carlo Geophysics

Using Quantum Monte Carlo techniques to study the Mg_2SiO_4 system, the transition pressure is in great agreement with experiment, compared to DFT which is 50% in error.


Present studies include effects of Fe, which makes up about 10% of material in the transition zone.

QMC calculations will accurately determine the equation of state of magnesium-iron-silicate, providing information on elastic behavior, structural transition pressure, and bulk sound velocity.

**8 M hrs (2008-09)
4K-5K cores max
2K typical**

Fusion Energy

- Fusion holds the promise of abundant, clean energy
- Problem: How to confine hot plasma needed for fusion reaction
- In tokamak reactors (e.g. ITER), strong magnetic fields confine a torus-shaped fusion plasma
- Effects at the outer edge of the confined plasma are of critical importance: e.g., interactions with container walls





Center for Plasma Edge Simulation

PI: C.S. Chang, NYU


- **ITER requires a substantial core plasma rotation in order to stabilize the dangerous resistive wall modes.**
- **No known external mechanism known to drive this rotation.**
- **Experimental evidence suggests that a spontaneous rotation in the edge could propagate into the core plasma by an ion turbulence process to generate enough rotation speed in ITER.**
- **A spontaneous co-current rotation in the tokamak edge plasma has been identified using the code XGC1.**

8.6 M hrs (2008-09) 16K cores max 10K cores median

Combustion

New combustion systems based on lean premixed combustion have the potential for dramatically reducing pollutant emissions in power generation and transportation.

Lean premixed combustion - esp. with H₂ - is highly susceptible to a variety of combustion instabilities, making robust, reliable systems difficult to design.





Combining a low Mach number formulation with adaptive mesh refinement provides a significant increase in computational efficiency compared to traditional uniform grid approaches to solving the full compressible flow equations


PI: John Bell, LBNL

Low-Swirl Burners

- **Discovered in 1991 at LBNL.**
- **May dramatically reduce pollutants by using special “lean premixed” fuels in power generation and transportation.**
- **Now being developed for fuel-flexible, near-zero-emission gas turbines (2007 R&D 100 Award)**



1" burner (5 kW, 17 KBtu/hr)



28" burner (44 MW, 150 MBtu/hr)

Low-Swirl Burner Simulation



PI: John Bell, LBNL

- Numerical simulation of an ultra-lean premixed hydrogen flame in a laboratory-scale low-swirl burner.
- Interaction of turbulence and chemistry.
- Method captures the hydrogen flame cell structures (lower right).

16,000 cores max
2,000 core typical

10 M Hours (2008-2009)

J B Bell, R K Cheng, M S Day, V E Beckner and
M J Lijewski, Journal of Physics: Conference
Series 125 (2008) 012027



Example Projects

Lattice Gauge Theory

Quantum Chromodynamics with three flavors of dynamical quarks - Toussaint, U. Arizona

Climate Research

Surface Input Reanalysis for Climate Applications (SIRCA) 1850-2011 - Compo, U. Colo.

Combustion

Interaction of Turbulence and Chemistry in Lean Premixed Laboratory Flames - Bell, LBL

Fusion Energy

Center for Plasma Edge Simulation: SciDAC FSP Prototype Center - Lee, PPPL

Materials Sciences

Modeling Dynamically and Spatially Complex Materials - Wilkins, OSU

Accelerator Physics

Particle simulation of laser wakefield particle acceleration - Geddes, LBNL

Chemistry

Large Eddy Simulation of Turbulence-Chemistry Interactions in Reacting Multiphase Flows - Oefelin, Sandia

Astrophysics

Computational Astrophysics Consortium - Woosley, UCSC



Example Projects

Life Sciences

Advanced Theoretical Models to Characterize the Alzheimers Amyloid Beta Peptide - Head-Gordon, UC-Berkeley

Applied Math

Applied Partial Differential Equations Center - Bell, LBNL

Geosciences

Large Scale 3D Geophysical Inversion & Imaging - Newman, LBNL

Computer Sciences

Development and Test of an IO API for the Global Cloud Resolving Model - Schuchardt, PNL

Scott?

Environmental Science

Impact of vegetation on turbulence over complex terrain: a wind energy perspective - Patton, UCAR

Summary

- **Science continues to thrive at NERSC on Franklin.**
 - Still fun, too.
- **Franklin is being used to address fundamental science and issues of urgent importance in energy sciences and applications.**



Acknowledgements

- **Kathy Yelick, NERSC Director**
- **Harvey Wasserman, NERSC SDSA**
- **Francesca Verdier, NERSC Services Department Head**