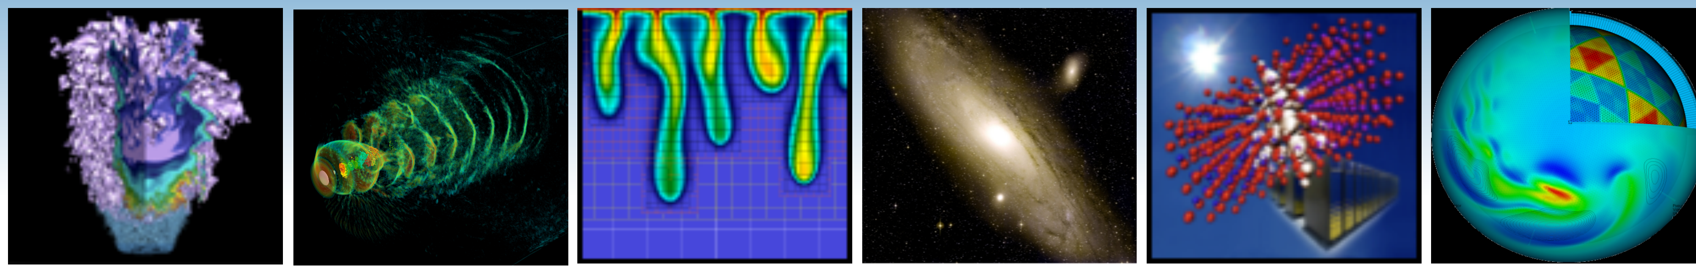

NERSC Overview

Francesca Verdier
fverdier@lbl.gov

NERSC Services Department Head
Lawrence Berkeley National Laboratory

Center for Multiscale Modeling of Atmospheric Processes meeting
January 11, 2011

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

NERSC Mission

NERSC's mission is to *accelerate the pace of scientific discovery* by providing high-performance computing, information, data, and communications services to the DOE Office of Science community.

NERSC Priority: Science First

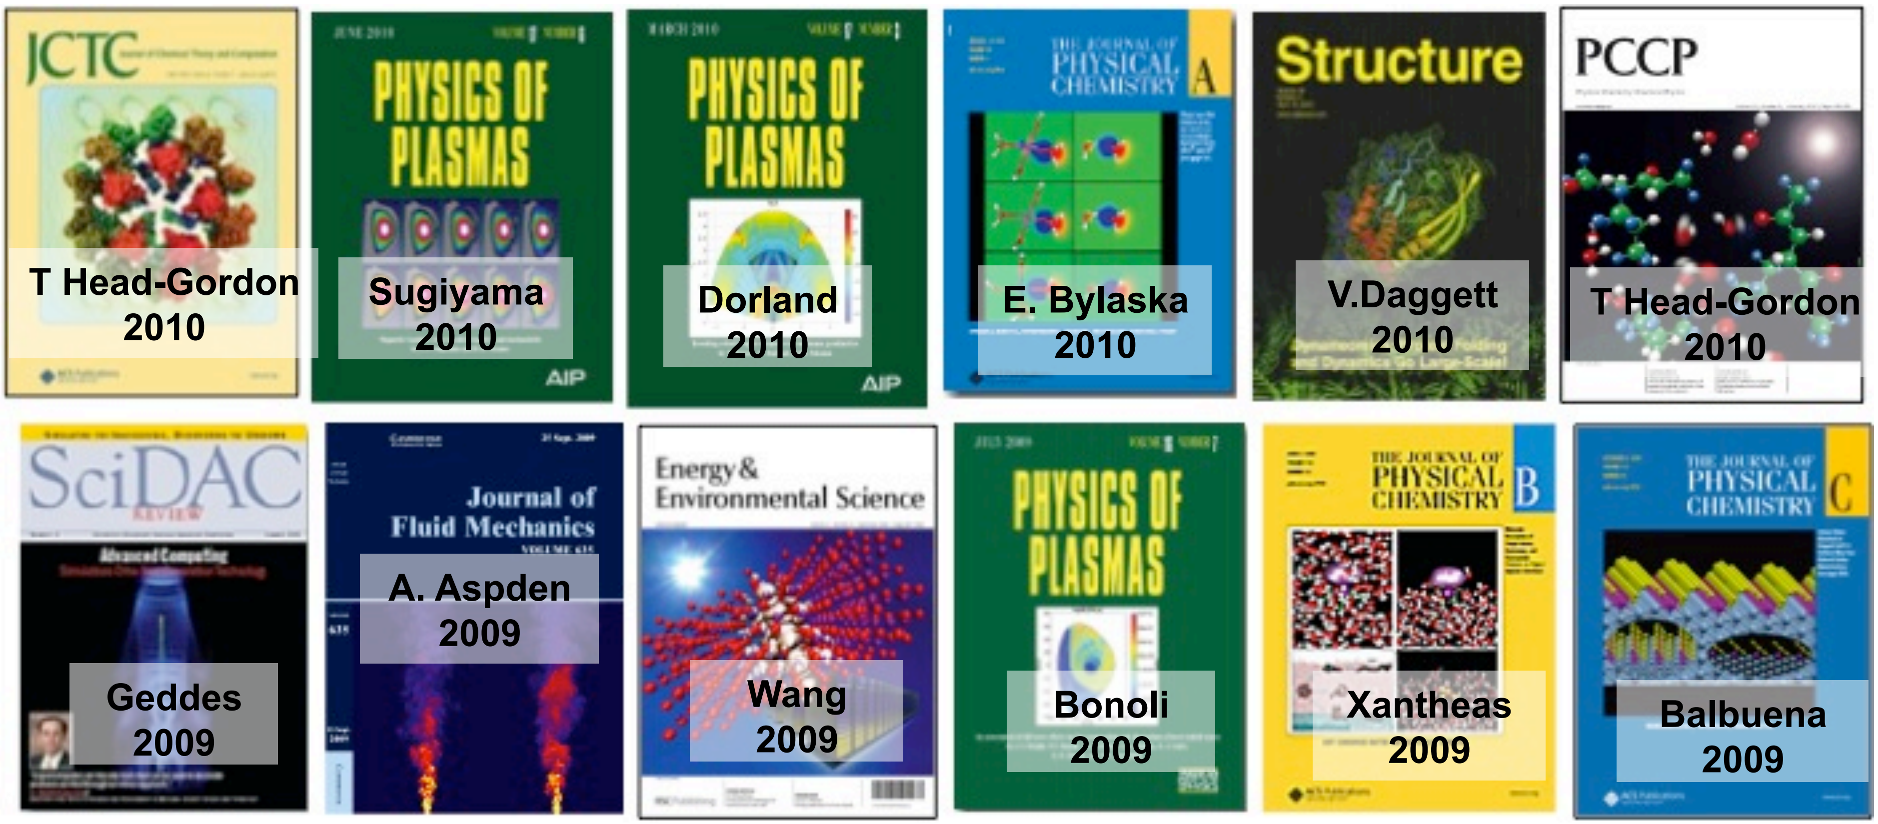

NERSC is enabling new high quality science across disciplines, with over *1,600 refereed publications* last year

NERSC is the Primary Computing Center for DOE Office of Science

- **NERSC serves a large population**
  - **Approximately 4000 users, 500 projects, 500 codes**
- **Focus on "unique" resources**
  - **Expert consulting and other services**
  - **High end computing & storage systems**
- **NERSC is known for:**
  - **Excellent services**
  - **Large and diverse user workload**

> "…NERSC, which is widely regarded as the best managed High Performance Computing Center in the world."
>
> *– Pat Dehmer, 2010*
> *DOE Deputy Director for Science*

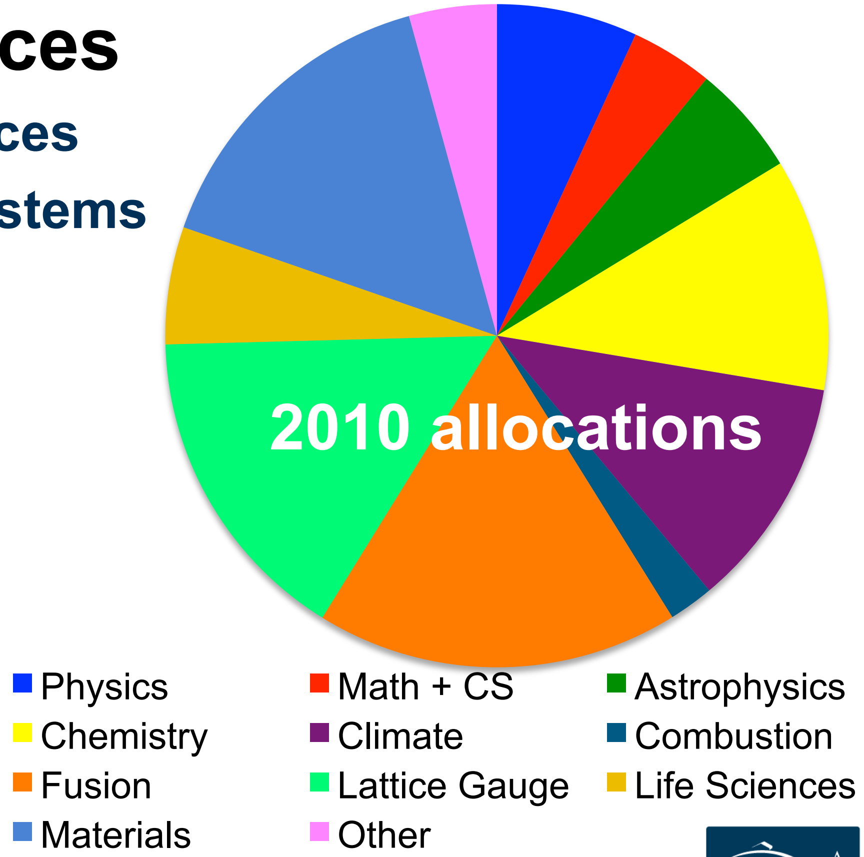

NERSC

NERSC Systems

Large-Scale Computing Systems

Franklin (NERSC-5): Cray XT4
- 9,532 compute nodes; 38,128 cores
- ~25 Tflop/s on applications; 356 Tflop/s peak


Hopper (NERSC-6): Cray XE6
- Phase 1: Cray XT5, 668 nodes, 5344 cores
- Phase 2: Cray XE6, 6392 cores, 153408 cores
- 1.27 Pflop/s peak


Clusters

140 Tflops total

Carver
- IBM iDataplex cluster

PDSF (HEP/NP)
- ~1K core throughput cluster

Magellan Cloud testbed
- IBM iDataplex cluster

GenePool (JGI)
- ~5K core throughput cluster

NERSC Global Filesystem (NGF)

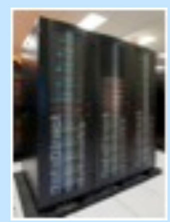

Uses IBM's GPFS
- 1.5 PB capacity
- 5.5 GB/s of bandwidth

HPSS Archival Storage

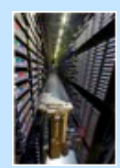

- 40 PB capacity
- 4 Tape libraries
- 150 TB disk cache

Analytics

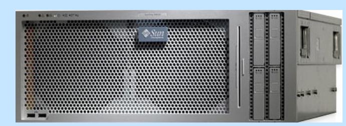

Euclid
(512 GB shared memory)

Dirac GPU testbed (48 nodes)

User Services

- **Consultants provide broad front line support to users**
- **Expertise in:**
  - **Material Science, Chemistry, Astrophysics and Climate codes**
  - **Math and I/O libraries**
  - **Compilers**
- **Communicate with users via:**
  - **Trouble ticket system/email**
  - **Phone calls**
- **Web documentation**
- **Software support**
- **Training**

> *"The quality of the technical staff is outstanding. They are competent, professional, and they can answer questions ranging from the trivial to the complex"*
>
> *NERSC User Survey*

Number of User Tickets Created Each Year

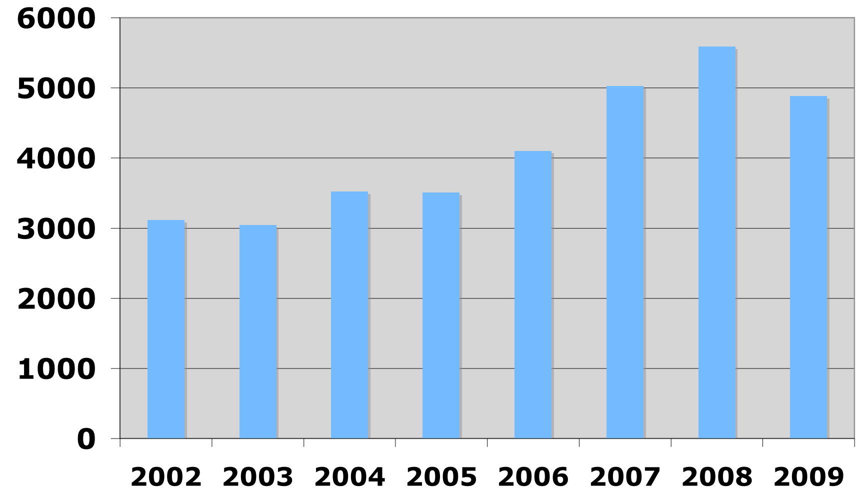

User Satisfaction with NERSC Consulting

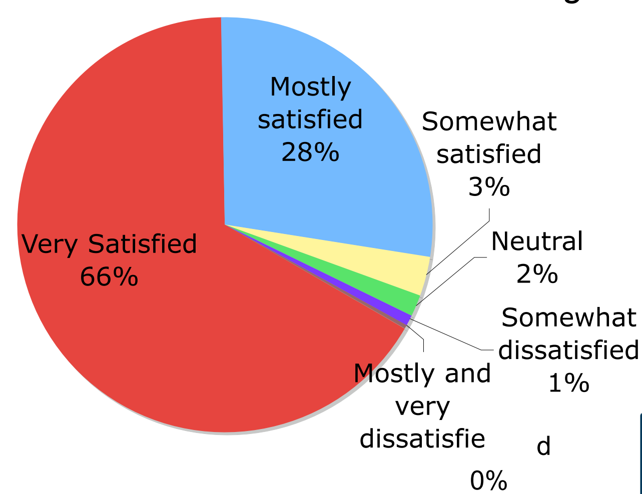

Getting an Allocation at NERSC

- **80% of NERSC time awarded by DOE Office of Science**
  - Annual request cycle; proposals due each August
  - Most of the time awarded to projects with a DOE Office of Science research grant
  - Awards may be made to other "DOE relevant" projects
  - Awards may be made "out of cycle" depending on reserves
  - http://www.nersc.gov/nusers/accounts/allocations/ercap/help.php
- **10% awarded via ASCR Leadership Computing Challenge**
  - http://www.er.doe.gov/ascr/Facilities/ALCC.html
- **10% of the time awarded by NERSC via NISE**
  - NERSC Initiative for Scientific Exploration
  - http://www.nersc.gov/nusers/accounts/NISE.php

NERSC Startup Allocations

- **NERSC also makes Startup award decisions**
  - Startup projects should demonstrate need for HPC resources
  - Startup awards are for a small amount of time (15,000 hours)
  - Can move to "production status" and get more time with concurrence of DOE
  - Award decisions within 1-3 weeks of applying
  - If new to NERSC must first request a NERSC username:
    - https://nim.nersc.gov/newpi.php
  - If new to NERSC must fill out the policy form:
    - http://www.nersc.gov/nusers/accounts/usage.pdf
- **To apply:**
  - Go to NERSC home page: www.nersc.gov
  - Click "Apply for an Allocation" under About NERSC

DOE Office of Science Computing Facilities

NERSC at LBNL

- 1000+ users,100+ projects
- Allocations:
  - 80% DOE program manager control
  - 10% ASCR Leadership Computing Challenge
  - 10% NERSC reserve
- Science includes all of DOE Office of Science
- Machines procured competitively

LCFs at ORNL and ANL

- 100+ users 10+ projects
- Allocations:
  - 60% ANL/ORNL managed INCITE process
  - 30% ACSR Leadership Computing Challenge
  - 10% LCF reserve
- Science limited to largest scale; no limit to DOE/SC
- Machines procured through partnerships

NERSC-6 Hopper Phase 2

Hopper phase 2 currently available to all users

- 1.25 PFlop/s peak performance
- Over 1 billion annual core-hours
- Gemini high performance resilient interconnect
- Two 12-core AMD Magny-Cours chips per node

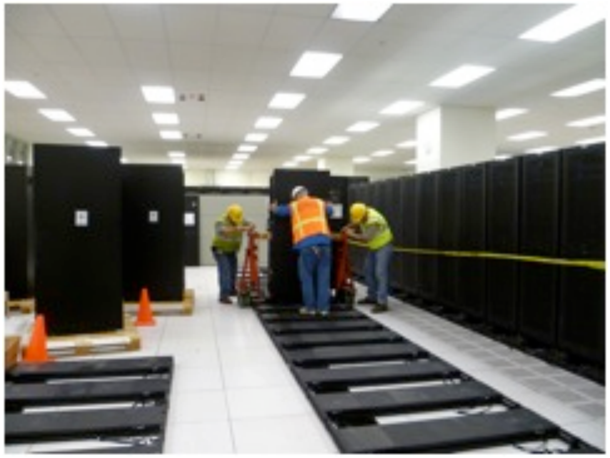

Hopper installation, August 2010

NERSC/Cray Center of Excellence

- Programming Models for Multicore systems
- Ensures effective use of new 24-core nodes

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Carver and Magellan Clusters

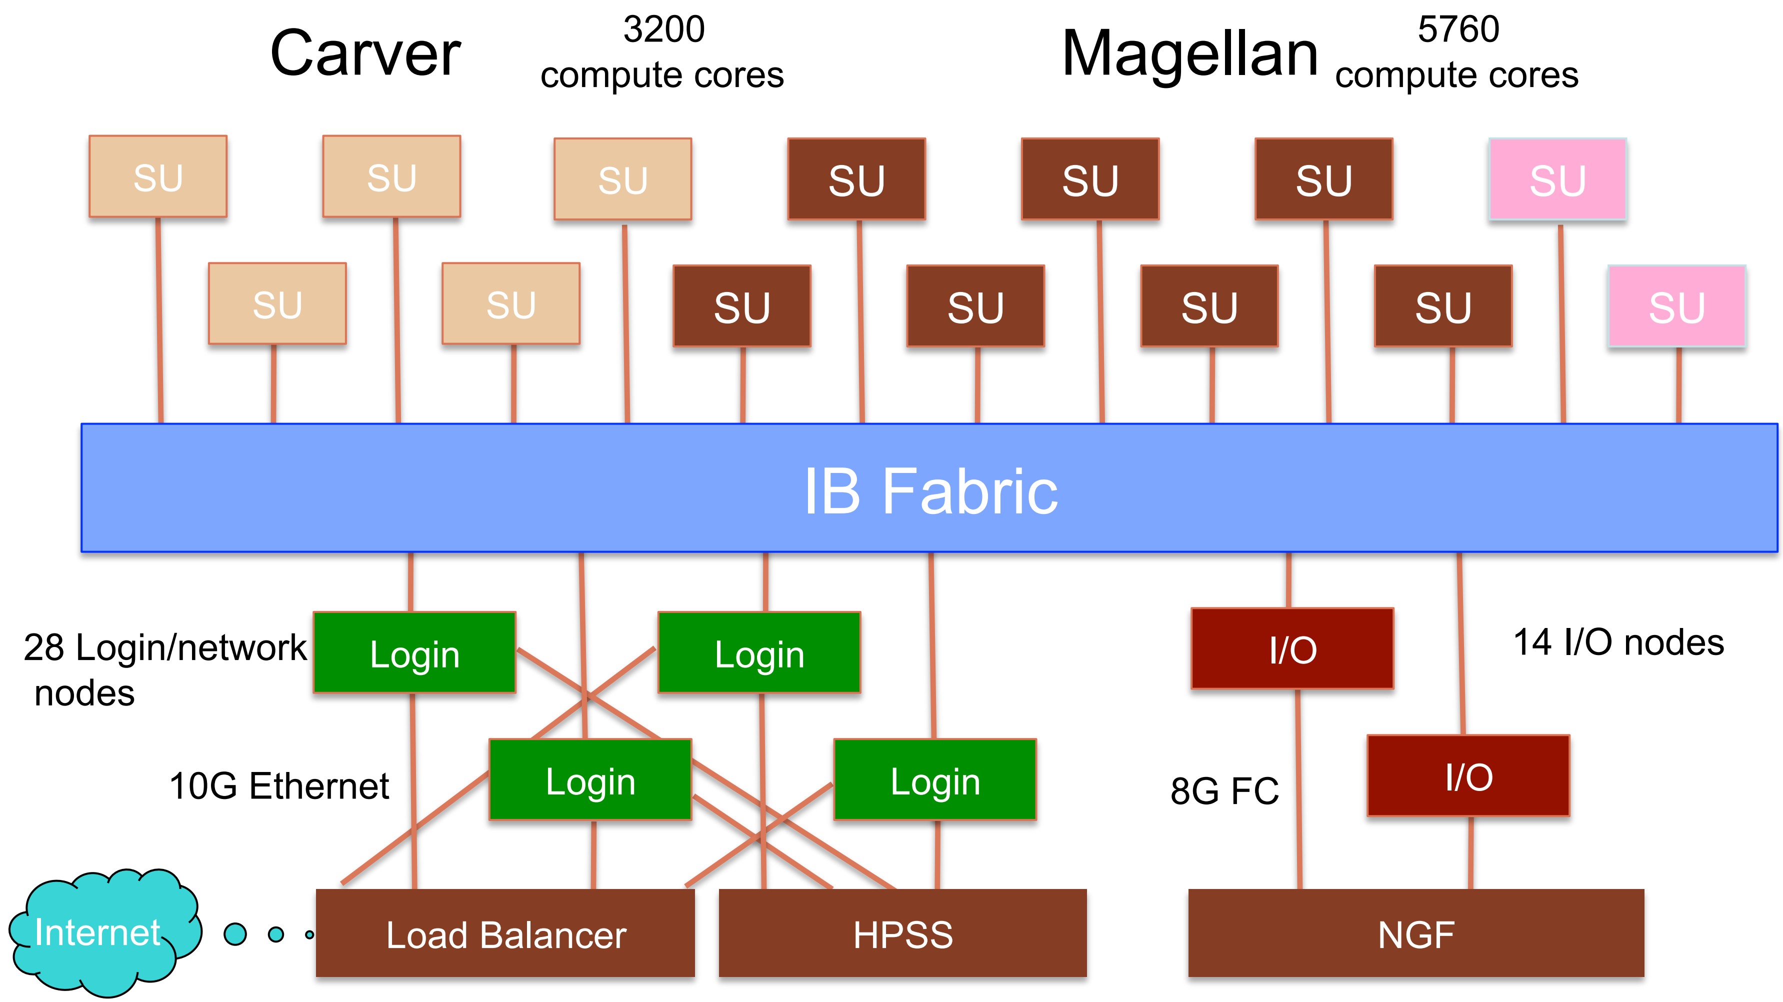

GPU Testbed

- **Installed "Dirac" GPU testbed**
  - About100 users so far
  - Popular with SciDAC-E postdocs
- **Example: Q-Chem Routine**
  - Impressive single node speedups relative to 1 core on CPU
  - Highly variable with input structure

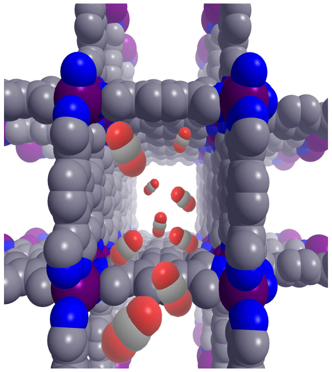

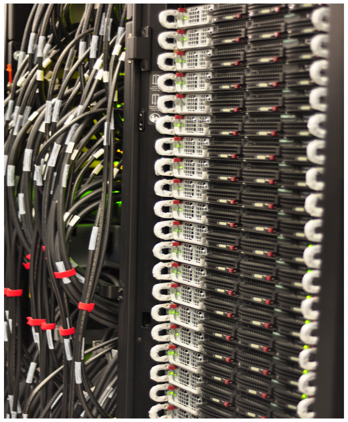

Fermi GPU Racks - NERSC

Blue: CPU 1 thread

Red: CPU 8 threads

Green: GPU

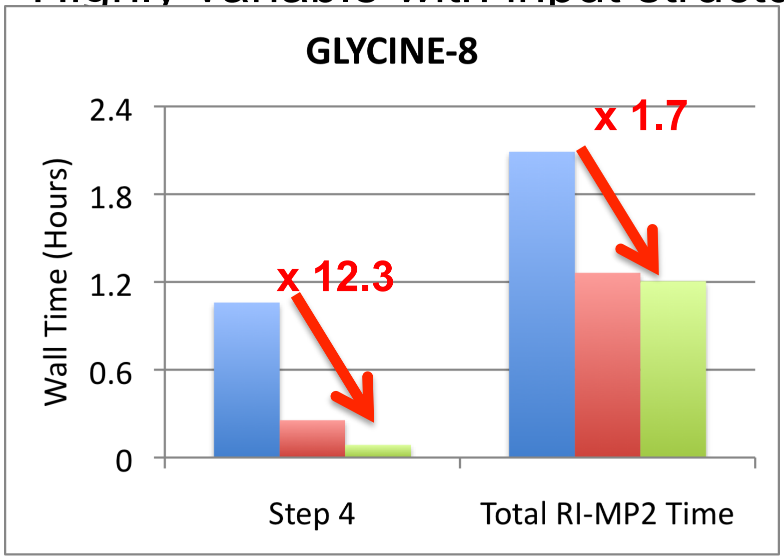

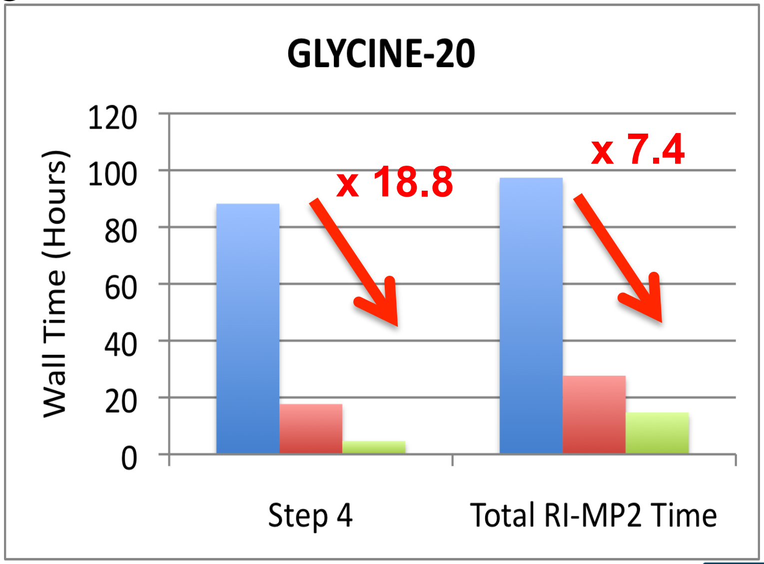

Jihan Kim (SciDAC-E NERSC Postdoc)

How and When to Choose Next Big System

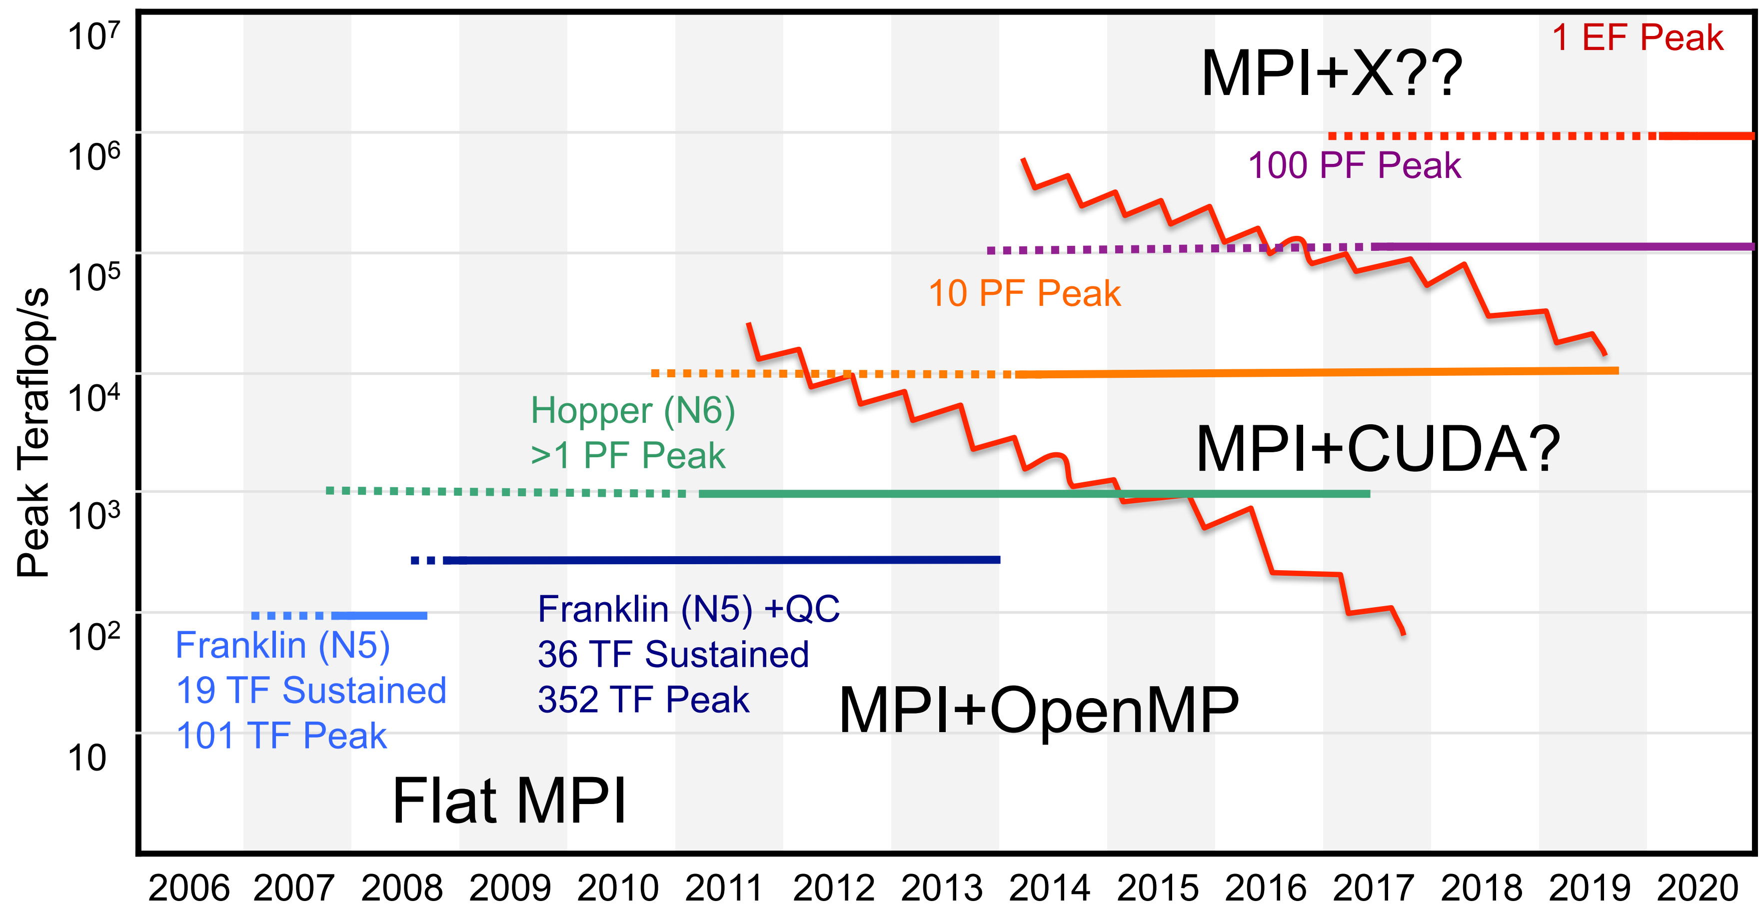

Want to avoid two paradigm disruptions on road to Exa-scale

Data Driven Science

- **Scientific data sets are growing exponentially**
  - **Ability to generate data is exceeding our ability to store and analyze**
  - **Simulation systems and some observational devices grow in capability with Moore's Law**
- **Petabyte (PB) data sets will soon be common:**
  - ***Climate modeling:* estimates of the next IPCC data is in 10s of petabytes**
  - ***Genome:* JGI alone will have .5 petabyte of data this year and double each year**
  - ***Particle physics*: LHC is projected to produce 16 petabytes of data per year**
  - ***Astrophysics*: LSST and others will produce 5 petabytes/year**
- **Create scientific communities with "Science Gateways" to data**

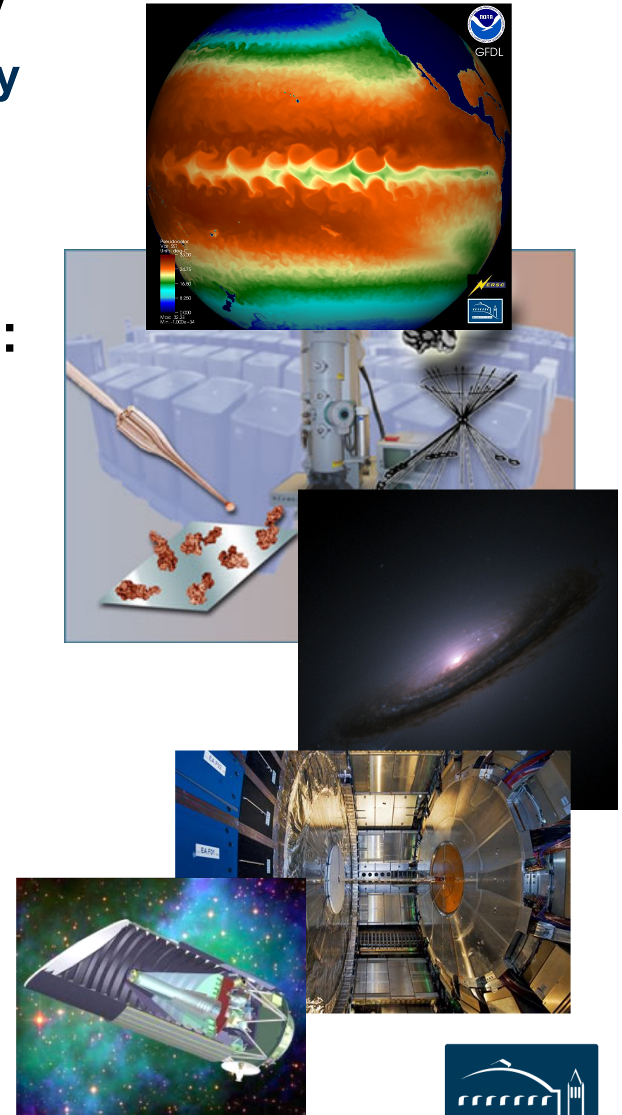

Science Gateways at NERSC

- **Create scientific communities around data sets**
  - Models for sharing vs. privacy differ across communities
  - Accessible by broad community for exploration, scientific discovery, and validation of results
  - Value of data also varies: observations may be irreplaceable

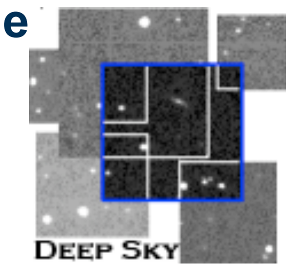

- **A *science gateway* is a set of hardware and software that provides data/services remotely**
  - Deep Sky – "Google-Maps" of astronomical image data
  - Gauge Connection – Access QCD Lattice data sets
  - PyDap: Interactive selection of 20th century climate data

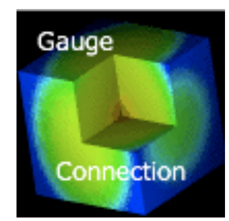

- **NERSC provides building blocks for science on the web**
  - Remote data analysis, databases, job submission

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Visualization Support

Petascale visualization: Demonstrate visualization scaling to unprecedented concurrency levels by ingesting and processing unprecedentedly large datasets.

Implications: Visualization and analysis of Petascale datasets requires the I/O, memory, compute, and interconnect speeds of Petascale systems.

Accomplishments: Ran VisIt SW on 16K and 32K cores of Franklin.

- First-ever visualization of two *trillion* zone problem (TBs per scalar); data loaded in parallel.
- Petascale visualization

Plots show 'inverse flux factor,' the ratio of neutrino intensity to neutrino flux, from an ORNL 3D supernova simulation using CHIMERA.

b

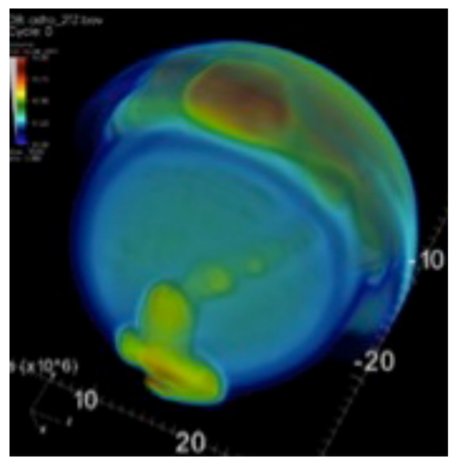

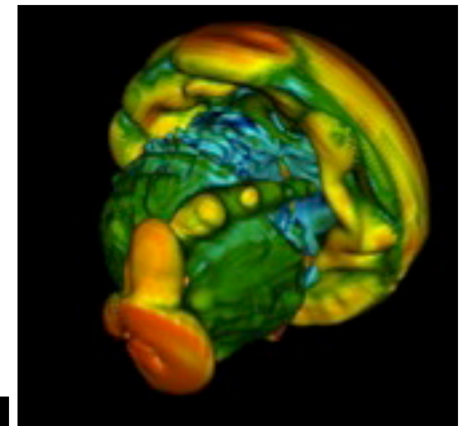

a

Isocontours (a) and volume rendering (b) of two trillion zones on 32K cores of Franklin.

Sample Scientific Accomplishments at NERSC

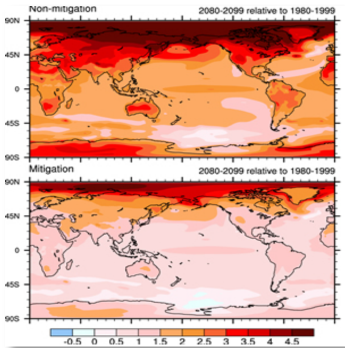

Climate

Studies show that global warming can still be diminished if society cuts emissions of greenhouse gases.
(Warren Washington, NCAR)

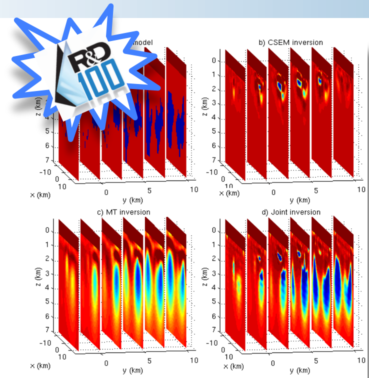

Energy Resources

Award-winning software uses massively-parallel supercomputing to map hydrocarbon reservoirs at unprecedented levels of detail.
(Greg Newman, LBNL)

Fusion Energy

A new class of non-linear plasma instability has been discovered that may constrain design of the ITER device.
(Linda Sugiyama, MIT)

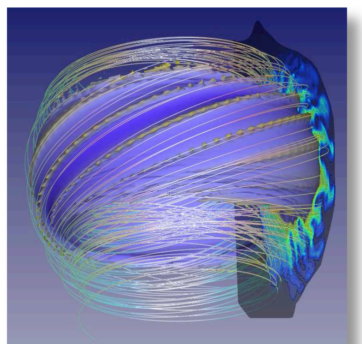

Combustion

Adaptive Mesh Refinement allows simulation of a fuel-flexible low-swirl burner that is orders of magnitude larger & more detailed than traditional reacting flow simulations allow.
(John Bell, LBNL)

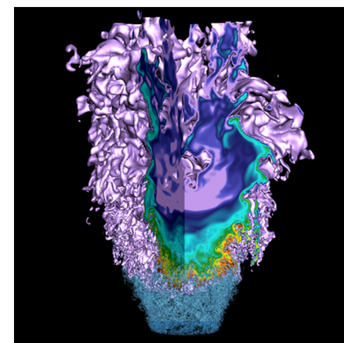

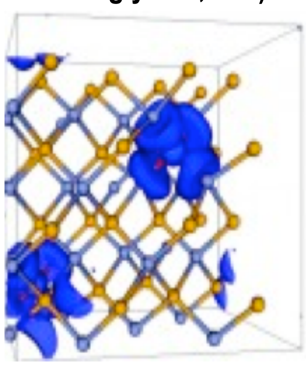

Materials

Electronic structure calculations suggest a range of inexpensive, abundant, non-toxic materials that can produce electricity from heat.
(Jeffrey Grossman, MIT)

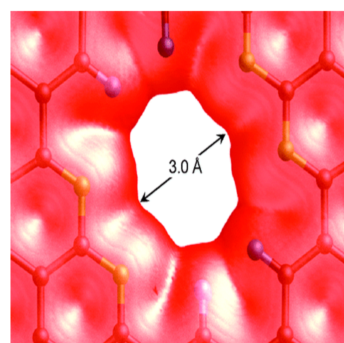

Nano Science

Using a NERSC NISE grant researchers discovered that Graphene may be the ultimate gas membrane, allowing inexpensive industrial gas production.
(De-en Jiang, ORNL)

U.S. DEPARTMENT OF ENERGY | Office of Science

Cloud-Resolving Climate Model

Objective: Climate models that fully resolve key convective processes in clouds; ultimate goal is 1-km resolution.

Implications: Major transformation in climate/weather prediction, likely to be standard soon, just barely feasible now.

Accomplishments: Developed a coupled atmosphere-ocean-land model based on geodesic grids.

- Multigrid solver scales perfectly on 20k cores of Franklin using grid with 167M elements.

NERSC:

- NERSC/LBNL played key role in developing critical I/O code & Viz infrastructure to enable analysis of ensemble runs and icosohedral grid.

PI: D. Randall, Colo. St

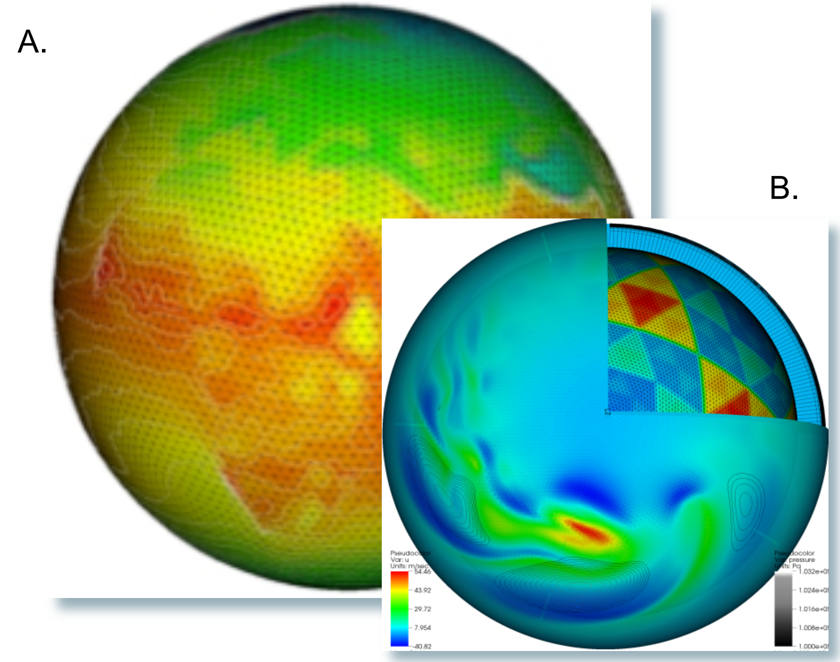

A. Surface temperature showing geodesic grid.
B. Composite plot showing several variables: wind velocity (surface pseudocolor plot), pressure (b/w contour lines), and a cut-away view of the geodesic grid.

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Mitigating Global Climate Change

Objective: **Determine if global warming can still be diminished if society cuts emissions of greenhouse gases.**

Implications: **Provide policymakers with appropriate research so they can make informed decisions to avoid the worst impacts of climate change.**

Accomplishments: **CCSM used at NERSC, ORNL, ANL, & NCAR to study a century of climate conditions, two CO_2 scenarios.**

- **70% cut in emissions would save arctic ice, reduce sea level rise.**

NERSC

- **Newer studies at NERSC include ~20,000-yr CCSM3 T42 studies of catastrophic change in Atlantic Meridional overturning circulation.**

W. Washington (NCAR)

Simulations show how average surface air temperatures could rise if greenhouse gas emissions continue to climb at current rates (top), or if emissions are cut by 70% (bottom).

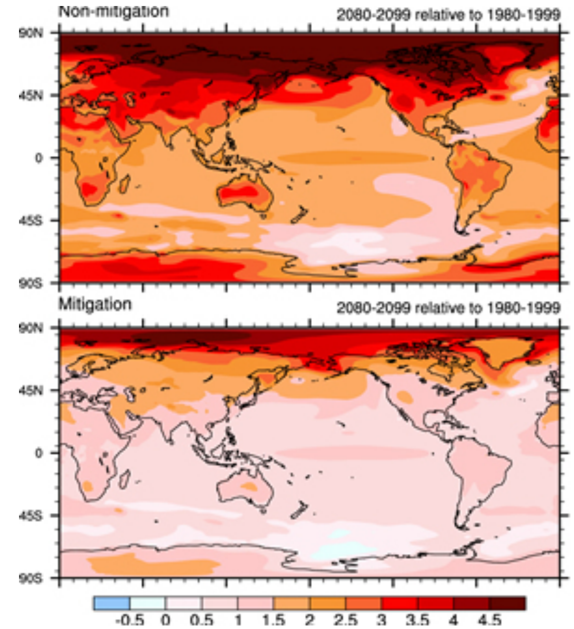

Temperatures rise by <2°C across nearly all populated areas if emissions are cut but unchecked emissions could lead to warming of >3°C in those areas.

Geophys. Res. Lett. 36, 08703 (2009)

U.S. DEPARTMENT OF ENERGY | Office of Science | BER

BERKELEY LAB

Global Warming Potential from Halocarbon Greenhouse Gases

Objective: **First-ever use of a global chemistry-transport model to estimate indirect global warming potential (GWP) of halocarbon gasses in the atmosphere.**

Implications: **Halocarbons destroy ozone; ozone can both warm and cool the earth. It is critical to understand the net influence.**

Accomplishments: **Confirms the significant importance of indirect effects on climate.**

- **Shows why GWP is a useful measure of relative climate impact; accounts for the atmospheric lifetime of greenhouse gases.**

NERSC: **Uses GFDL "Mozart" (Model of Ozone and Related Chemical Tracers) code; typically 96 cores.**

D. Wuebbles (U. Illinois)

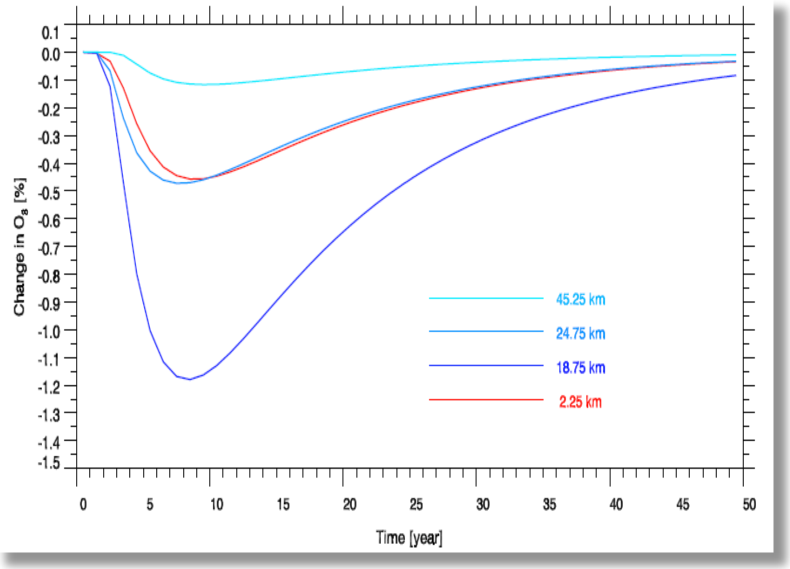

Calculated percent change in Ozone concentration as a function of time in years at four different elevations.

Atmos. Chem. Phys., 9, 8719–8733, 2009

20th Century Climate Reanalysis

Objective: Use an Ensemble Kalman filter to reconstruct global weather conditions in six-hour intervals from 1871 to the present.

Implications: Validate tools for future projections by successfully recreating – and explaining – climate anomalies of the past.

Accomplishments: First complete database of 3-D global weather maps for the 19th to 21st centuries.

- Provide missing information about the conditions in which extreme climate events occurred.
- Reproduced 1922 Knickerbocker storm, comprehensive description of 1918 El Niño
- Data can be used to validate climate and weather models

G. Compo (U. Colorado)

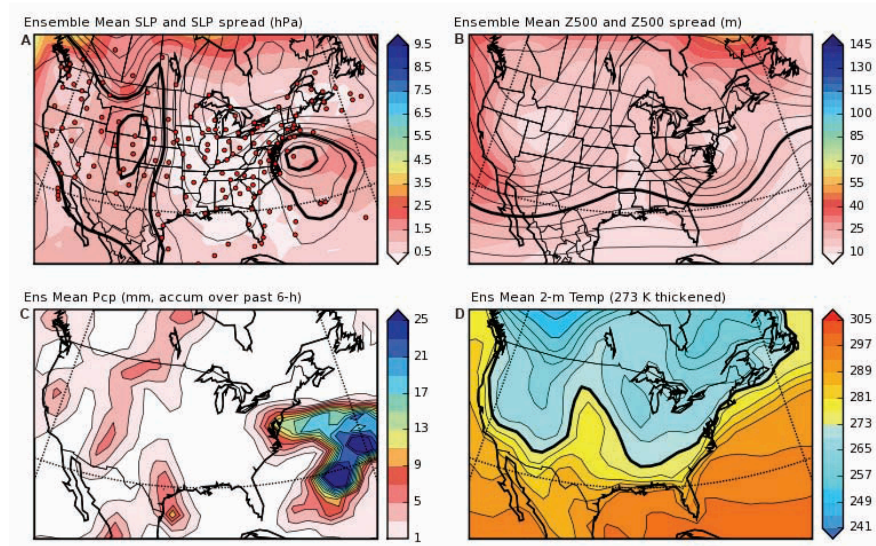

Sea level pressures with color showing uncertainty (a&b); precipitation (c); temperature (d). Dots indicate measurements locations (a).

Monthly Weather Review Vol 137(6) 2009;
Bull. Am. Meteorological Soc. (2009)

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Eddy-Resolving Ocean Model

Objective: **Understand deep ocean circulation and its response to an altered atmospheric composition.**

Implications: **Improved knowledge of CO_2 sequestration in the deep ocean and oceanic flows is crucial for understanding global climate change.**

Accomplishments: **First ocean model to resolve mesoscale flows over a wide range of parameters such as wind speed and surface temperature.**

- **Shows how dynamics of the Southern Ocean remotely control strength of meridional overturning (also known as the great ocean conveyor-belt).**

NERSC:

- **Completed over 15,000 simulation years using 1.6 M processor core hours, typically using 1,024 cores.**

P. Cessi, C. Wolfe, Scripps

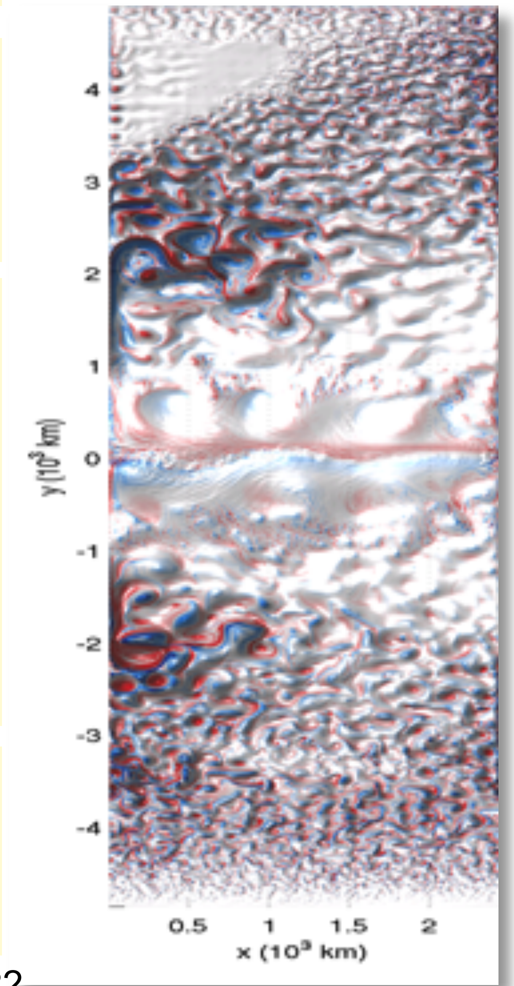

A simulation capturing eddy behavior in the Southern Ocean. A key feature is the abundance of eddies away from the equator which is shown in the center of the image at y = 0.

J. Phys. Oceanography (2008)

DOE-NOAA/GFDL Collaboration

Objective: Early look at issues involved with resolving mesoscale features in atmospheric and ocean circulations.

Implications: Provide near-term insight into regional climate change; inform the design of international modeling campaigns aimed at addressing this.

Accomplishments: Developed global models with atmosphere resolution ≅ 5 km; ocean resolution 10 - 20 km;

- Based on Flexible Modeling System (FMS) w/ tri-polar or cubed-sphere grids.
- Experiments generate 1 - 4 TB / sim. yr.
- Simulation output from Franklin loaded directly into VisIt for viz., analytics.
- Franklin can accommodate large per-core memory needs.

NERSC:

- Significant NERSC visualization support.

V. Balaji, C. Kerr GFDL

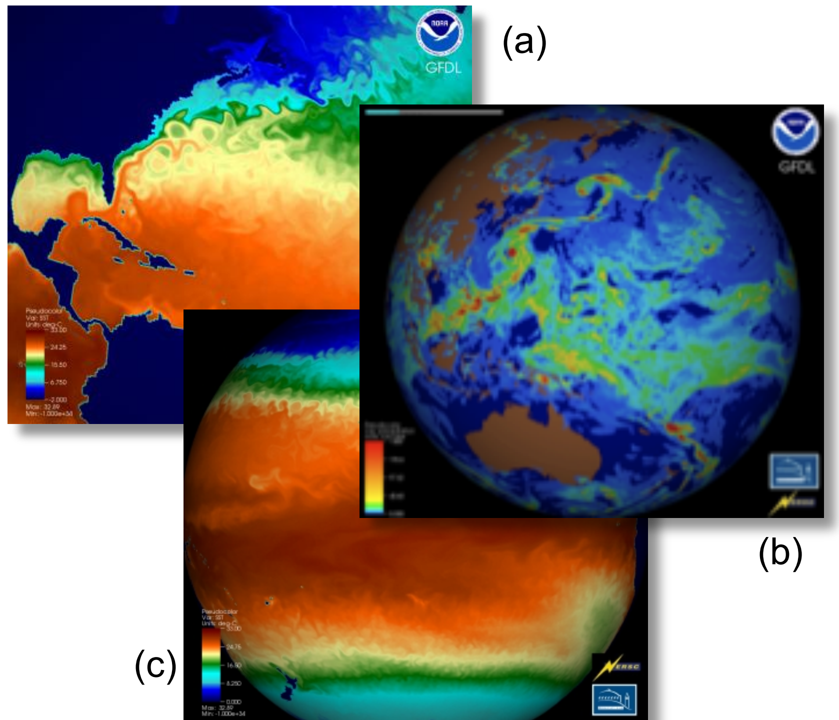

NERSC Analytics Team visualizations of GFDL-generated data: (a) CM2.4 sea surface temperature for North Atlantic Gulf Stream; (b) Pacific surface precipitation (c360 model); (c) Pacific surface temperature. Datasets provided by Chris Kerr (NOAA/GFDL)