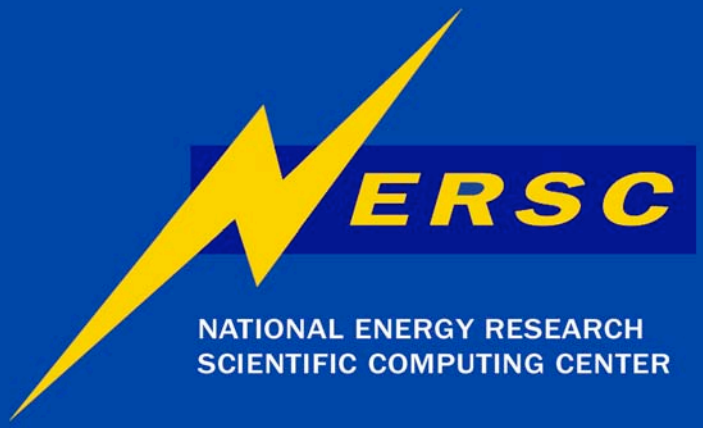

National Energy Research Scientific Computing Center (NERSC)

Horst D. Simon
NERSC Center Division, LBNL
November 15, 2005
SC05, Seattle

NERSC Mission

The mission of the National Energy Research Scientific Computing Center (NERSC) is to accelerate the pace of scientific discovery by providing high performance computing, information, data, and communications services for research sponsored by the DOE Office of Science (SC).

Office of Science
U.S. DEPARTMENT OF ENERGY

NERSC Must Address Three Trends

- The widening gap between application performance and peak performance of high-end computing systems
- The recent emergence of large, multidisciplinary computational science teams in the DOE research community
- The flood of scientific data from both simulations and experiments, and the convergence of computational simulation with experimental data collection and analysis in complex workflows

Science-Driven Computing Strategy 2006 -2010

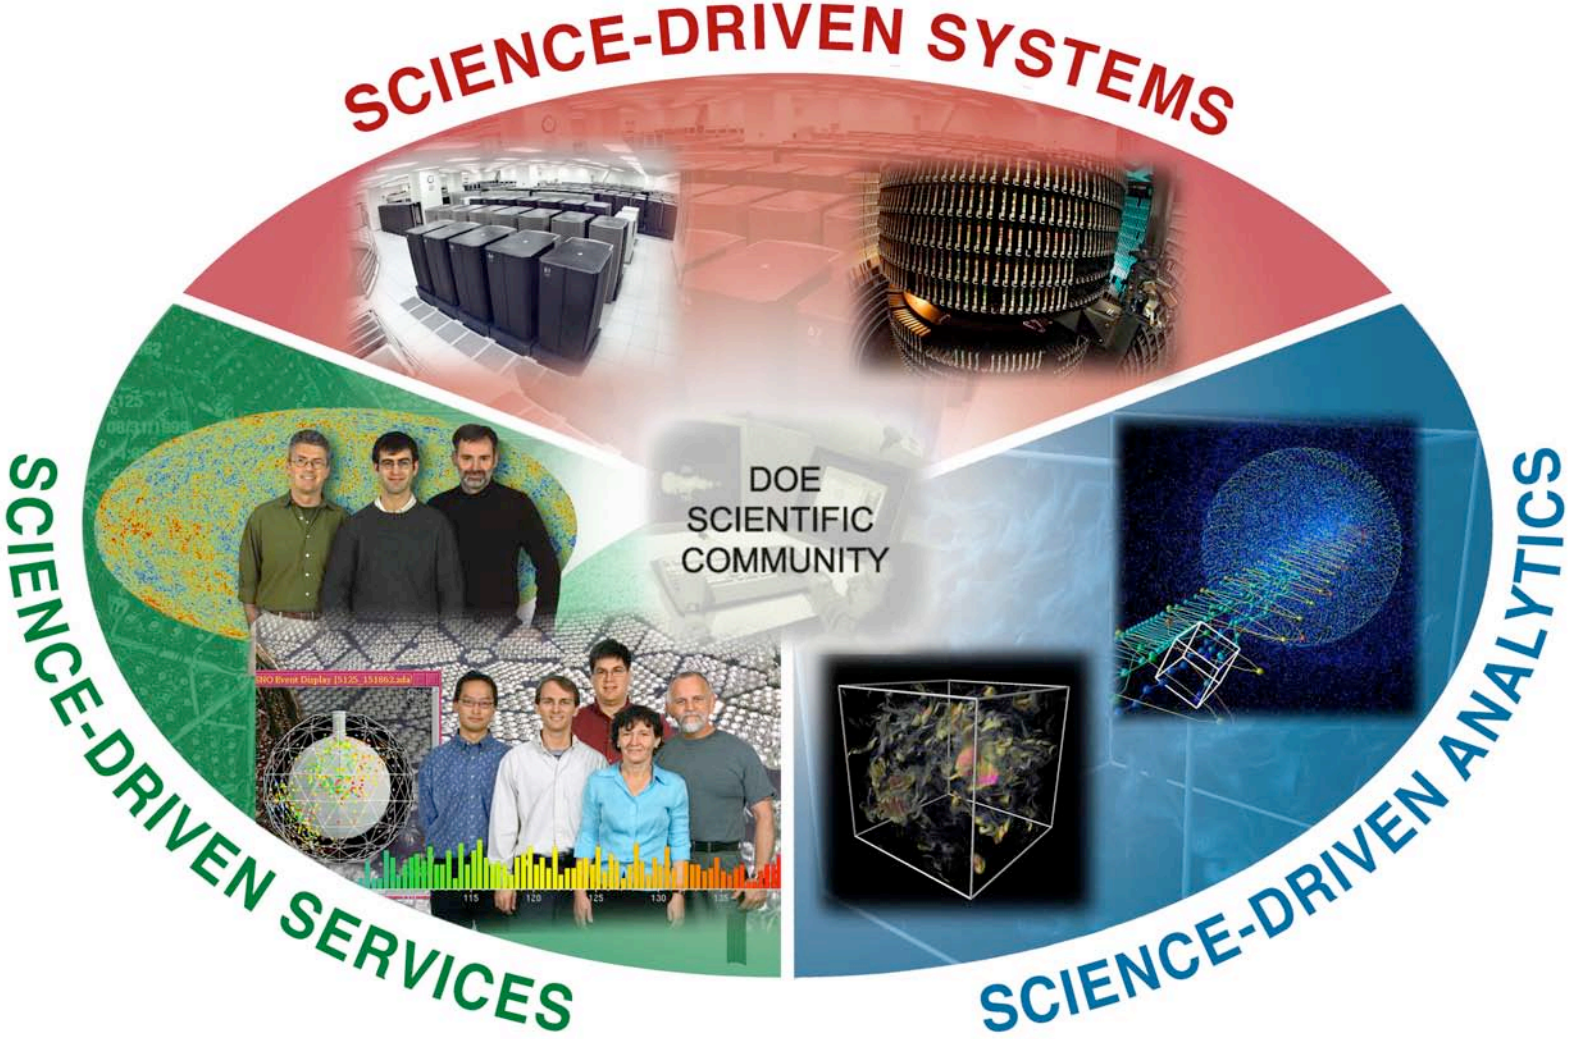

Office of Science
U.S. DEPARTMENT OF ENERGY

NERSC System Architecture April 2005

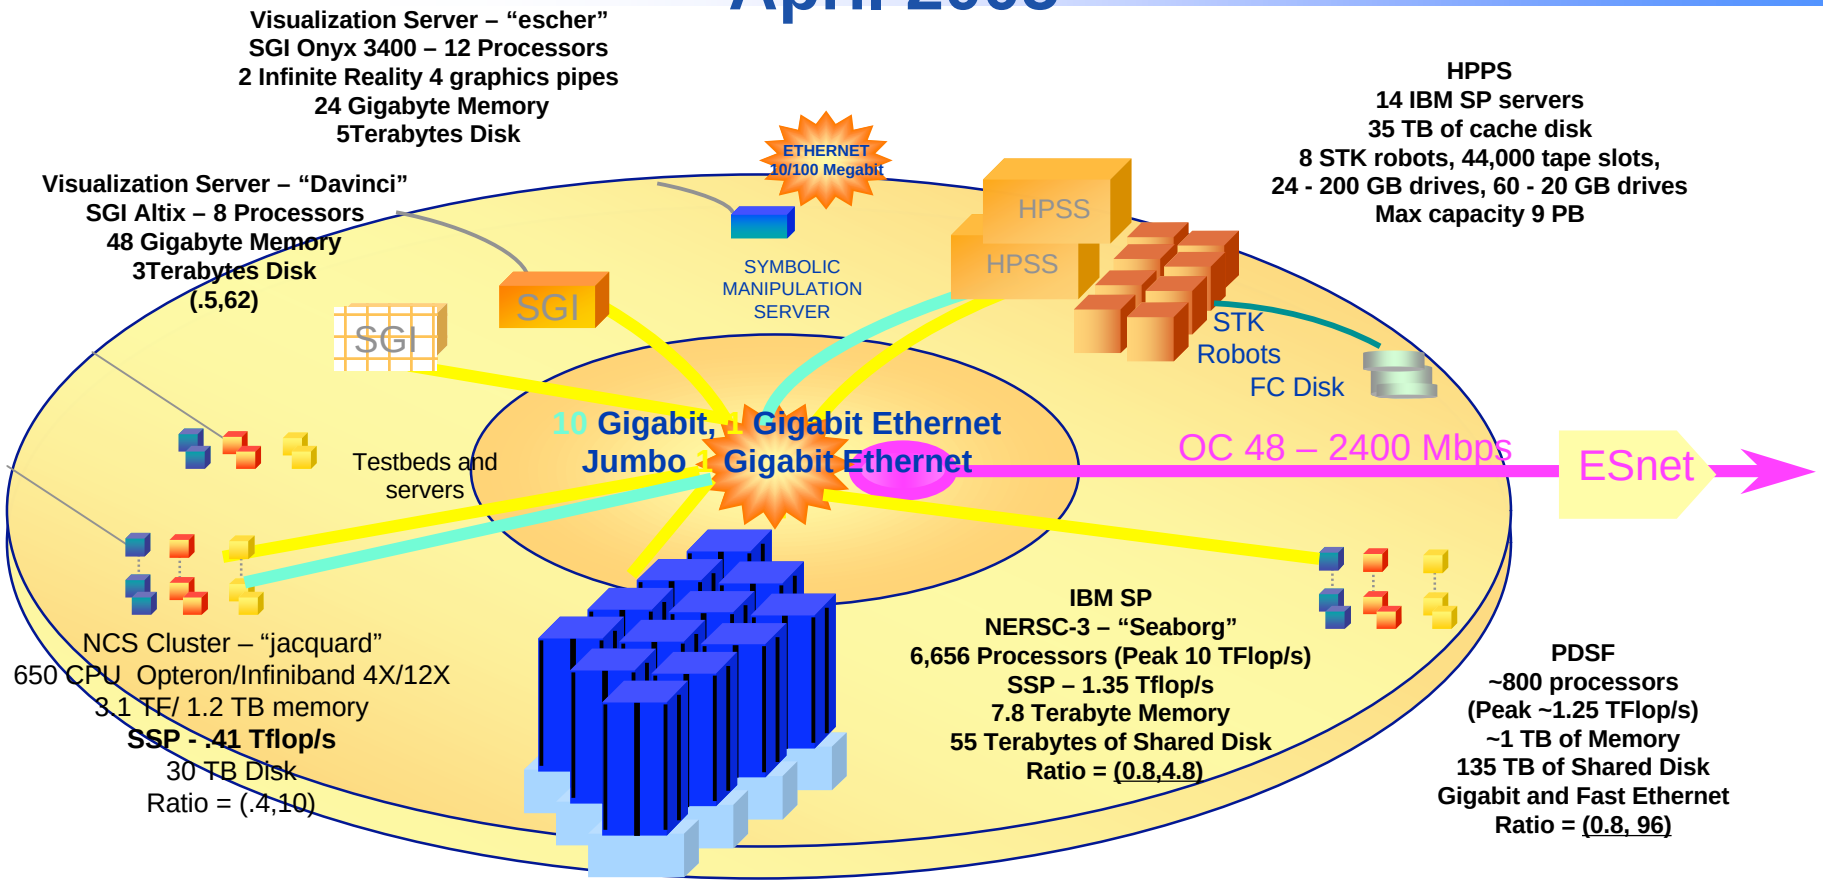

Ratio = (RAM Bytes per Flop, Disk Bytes per Flop)

Office of Science
U.S. DEPARTMENT OF ENERGY

2007

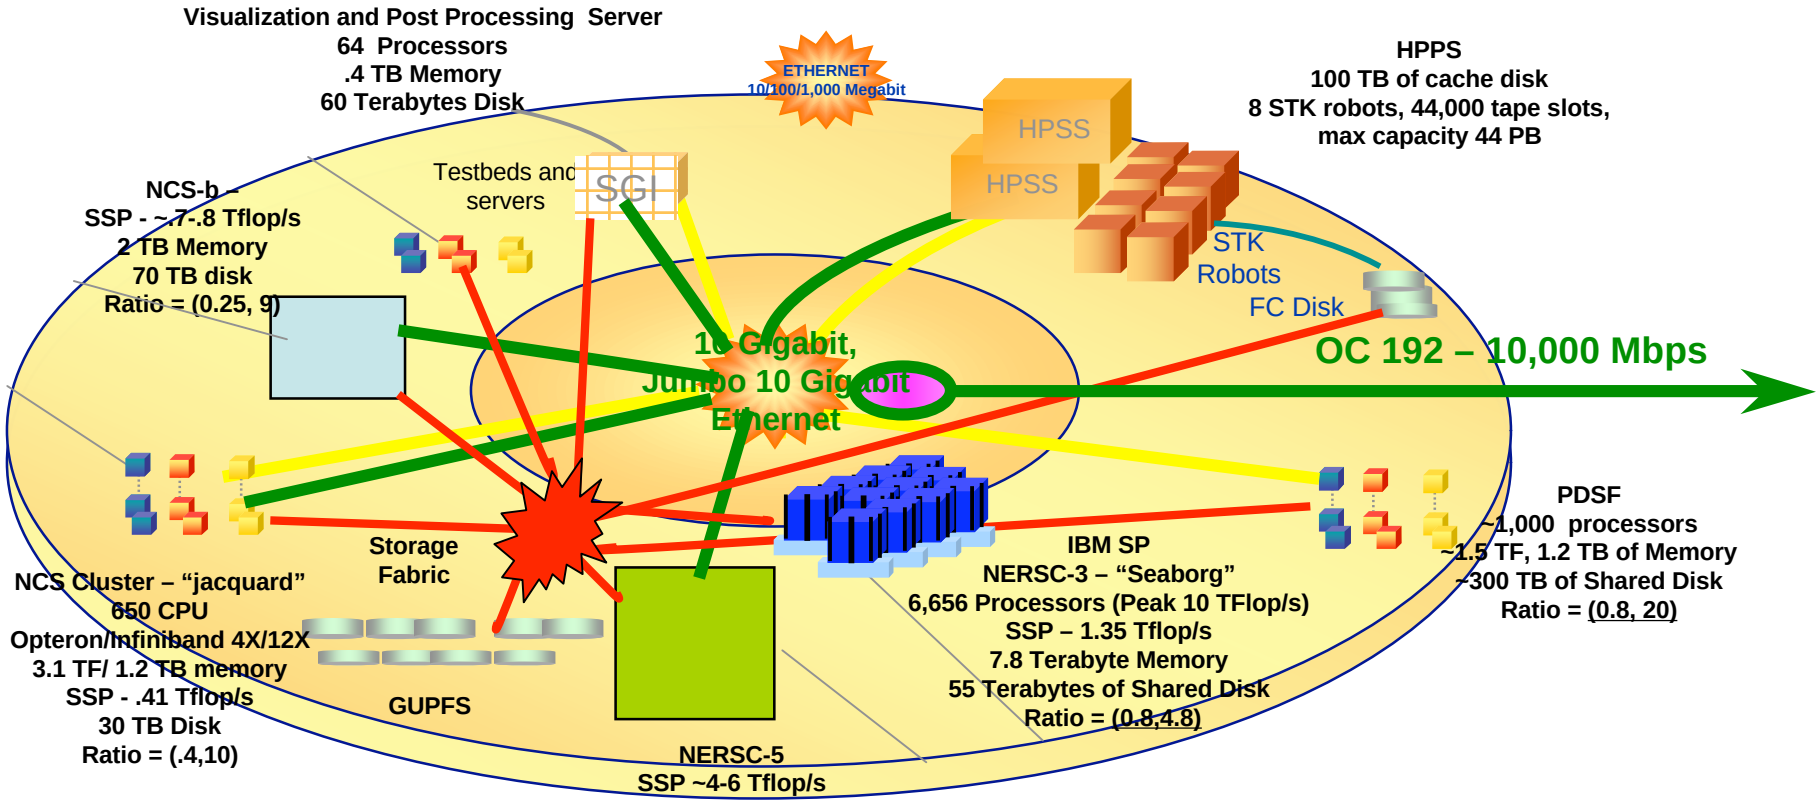

Ratio = (RAM Bytes per Flop, Disk Bytes per Flop)

2009

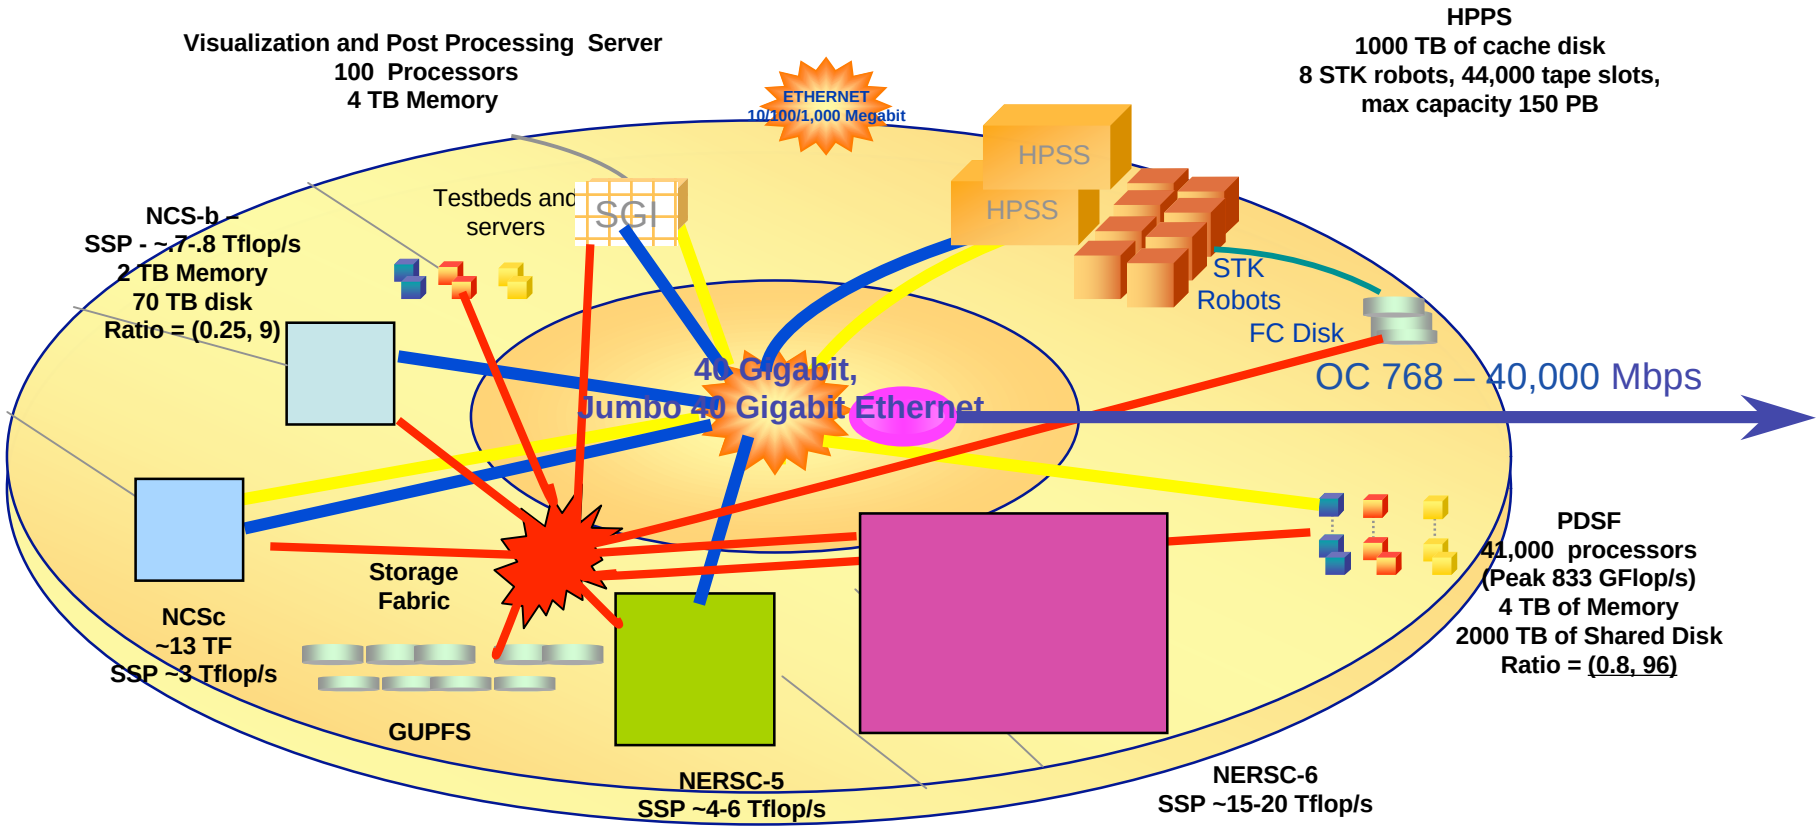

Ratio = (RAM Bytes per Flop, Disk Bytes per Flop)

Usage by Discipline (2004)

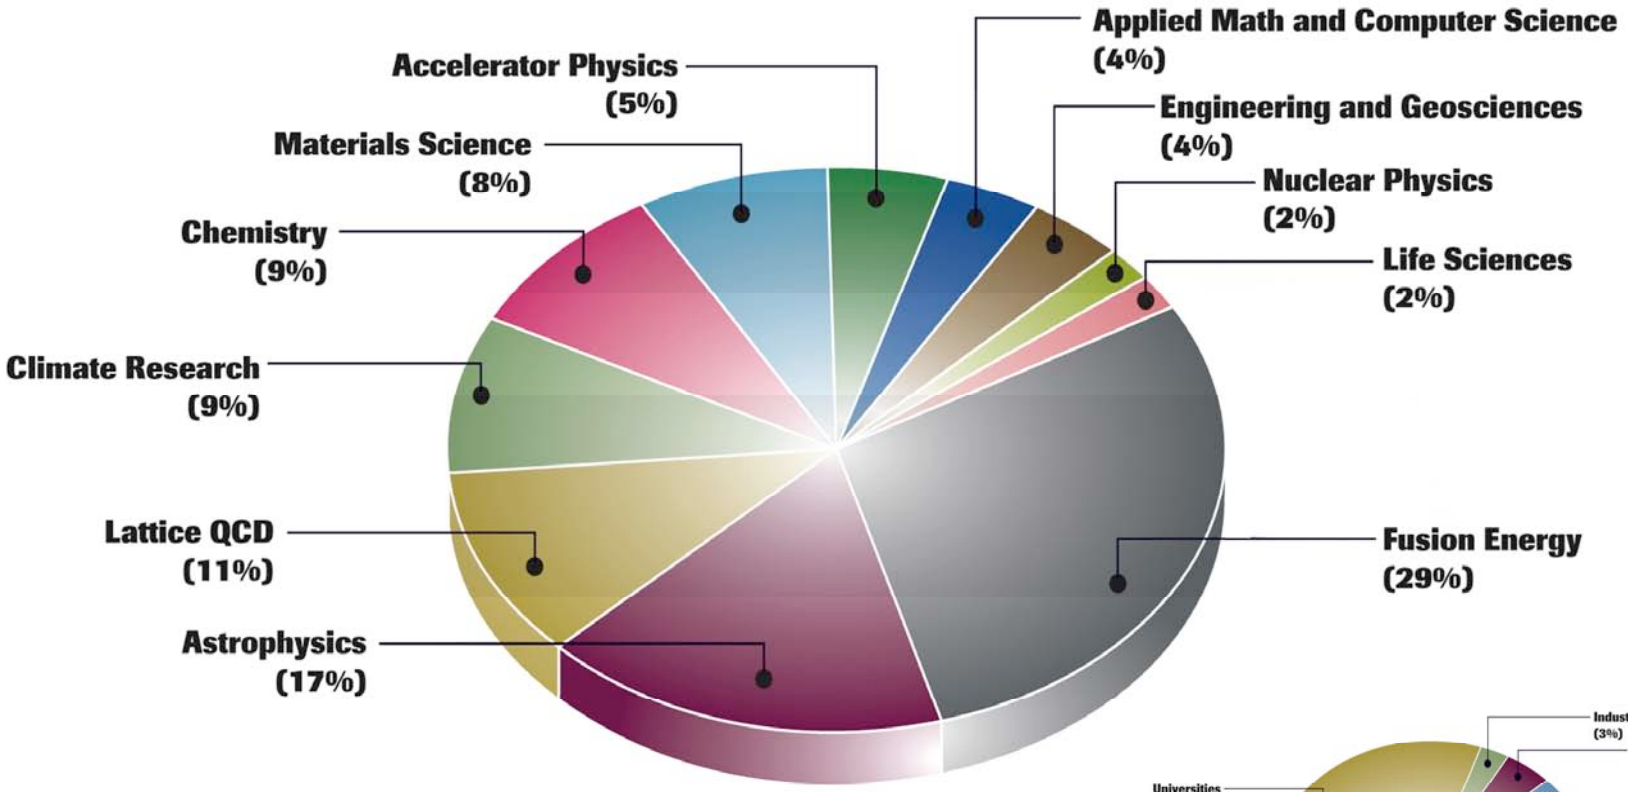

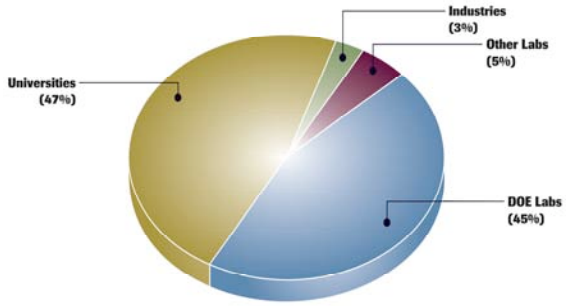

Office of Science
U.S. DEPARTMENT OF ENERGY

Overview INCITE 2004:

- **Summary**
  - Quantum Chemistry, Turbulent Flows & Super Novae
  - All projects successfully achieved their science goals
  - Late usage lead to Fall crunch on seaborg
  - All projects presented work in progress at SC04
- **Science Output**
  - “The allocation and queue priority under INCITE were a tremendous boon to our productivity. We would not have been able to address the systems studied in the INCITE project without the resources made possible in this program.” Prof. W.A. Lester
  - “The allocation allowed us for obtaining a large number of unique results that otherwise could not be obtained. This sheer size of the allocation makes INCITE a truly unique scientific resource.” T. Plewa
  - “Recognition via our INCITE Award definitely raised our profile in both the fluid-dynamics and supercomputing communities. The former is reflected in a flourishing of collaborative interest from others in the field. The data generated by our INCITE work is now available online as a resource to fluid-dynamicists ” – Prof. P.K. Yeung

Office of Science
U.S. DEPARTMENT OF ENERGY

INCITE 2004: Photosynthesis

- MPI tuning: 15-40% less MPI time
- Code scaling: from 256 to 4,096 procs
- More efficient random walk procedure
- Wrote parallel HDF layer
- Used AVS/Express to visualize molecules and electron trajectories
- Animations of the trajectories showed 3D behavior of walkers for the first time
- “Visualization has provided us with modes of presenting our work beyond our wildest imagination”
- “We have benefited enormously from the support of NERSC staff”

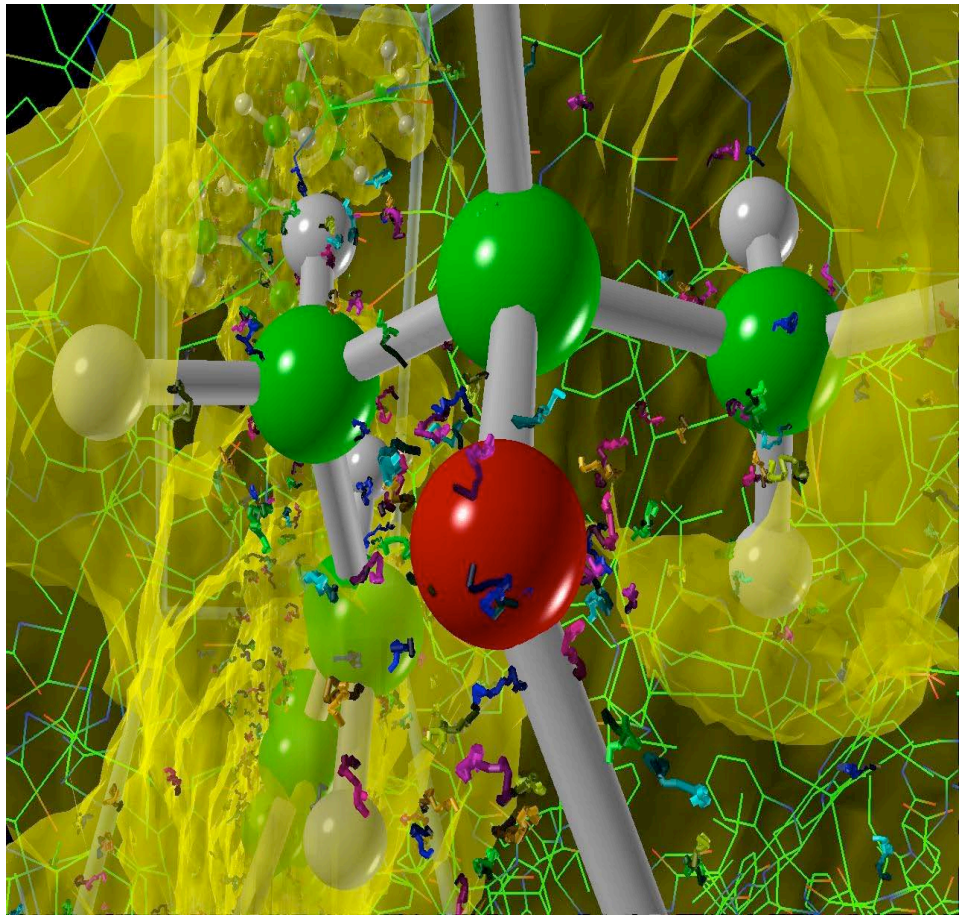

Office of Science
U.S. DEPARTMENT OF ENERGY

INCITE 2004: Turbulence

- Parallel performance profiling
- Implemented advanced math and MPI libraries (ESSL and early version of PE4.2)
- Face to face meetings on performance and visualization at ACTS workshop.
- “We have benefited from consultants' comments on code performance, innovative ideas for improvement, and diagnostic assistance when user or system problems occur. The visualization staff have also been very good.”

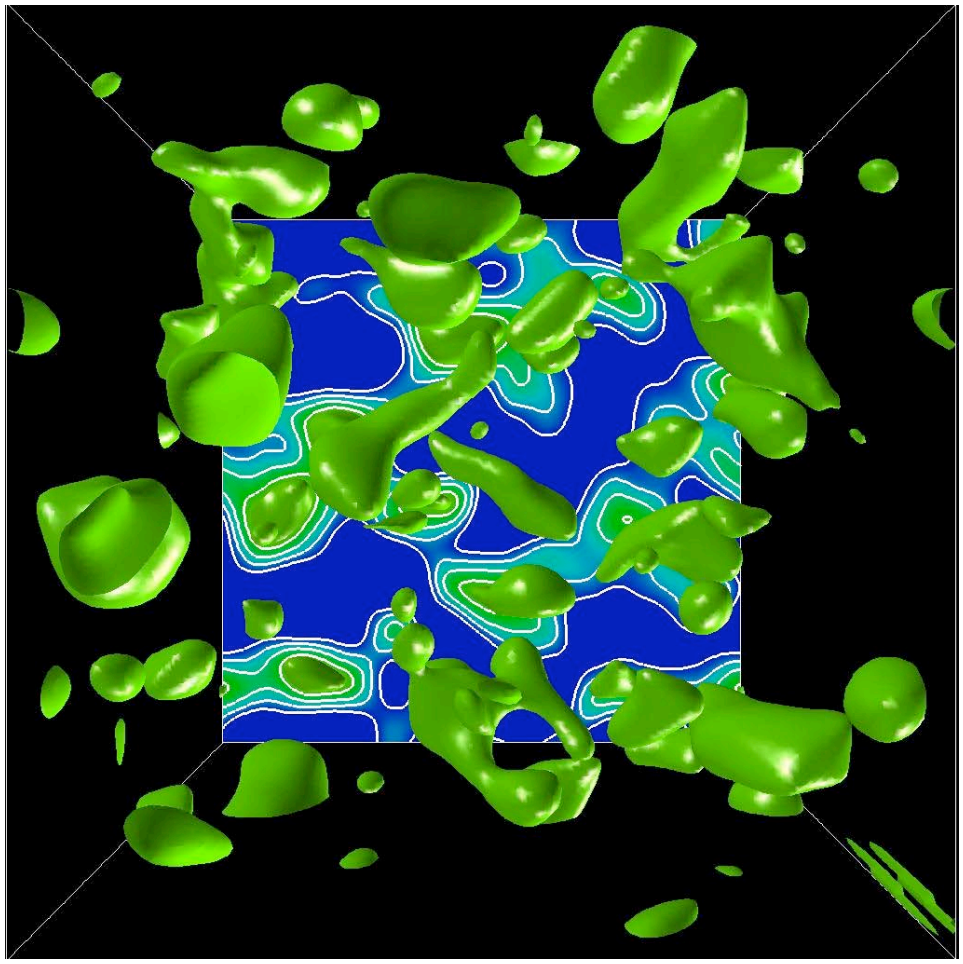

Office of Science
U.S. DEPARTMENT OF ENERGY

INCITE 2004: Thermonuclear Supernovae

- Resolved problems with parallel I/O by switching to a 64-bit environment
- Tuned network connections and replaced scp with hsi : transfer rate went from 0.5 to 70 MB/sec
- Created automatic procedure for code checkpointing
- “We have found NERSC staff extremely helpful in setting up the computational environment, conducting calculations, and also improving our software”

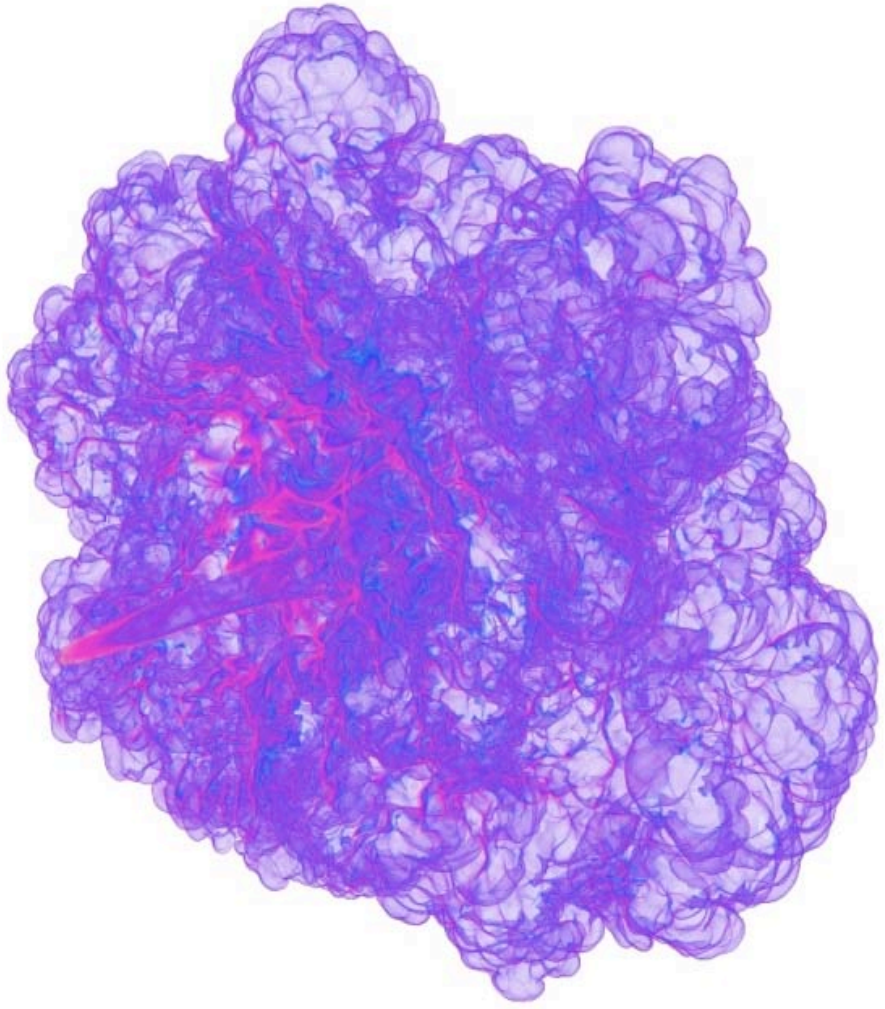

Office of Science
U.S. DEPARTMENT OF ENERGY

Publications in 2005 from cycles in 2004

- P.K. Yeung, D.A. Donzis and K.R. Sreenivasan (2005) “High-Reynolds-number simulation of turbulent mixing” Physics of Fluids Vol. 17, paper 081703 (August 2005).
- P.K. Yeung, S.B. Pope and D.A. Donzis (2005) “Acceleration and dissipation statistics in large numerical simulations of isotropic turbulence“ To be submitted to Physics of Fluids
- Turbulence Data is available in an online database http://www.ae.gatech.edu/people/pyeung/DNSdb.html
- A. Aspuru-Guzik, R. Salomon-Ferrer, B. Austin, R. Perusqua-Flores, M. A. Griffin, R. A. Oliva, D. Skinner, D. Domin, and W. A. Lester, Jr., Zori 1.0 : A parallel quantum Monte Carlo electronic structure package. Journal of Computational Chemistry. Journal of Computational Chemistry 26 p. 856, 2005
- T. Plewa and T. Linde, paper in press.

2005 : A work very much in progress

- INCITE 2005: MHD, Combustion, MD
  - Face to face meetings with Sandia
  - Telecons with U. Chicago
  - Elevated queue priority for INCITE 2005
  - Heavy usage of seaborg
- NERSC goal is to provide INCITE 2005 with Comprehensive Scientific Support
- Usage of allocations so far
  - INCITE 4 : 81%
  - INCITE 5 : 71% (+ non-seaborg usage)
  - INCITE 6: 75%

INCITE 4: Prof. Fausto Catteneo @ U. Chicago

Science Goals:
Understanding accretion in the cosmos through simulation and experiment. This work intends to model an experiment being done in a Princeton lab to understand magneto-rotational instability.

NERSC Work:
Nek5000 is a mature code

Review and analysis of parallel I/O strategy.

Visualization

Road Ahead:
Heavy interaction with LBL Viz

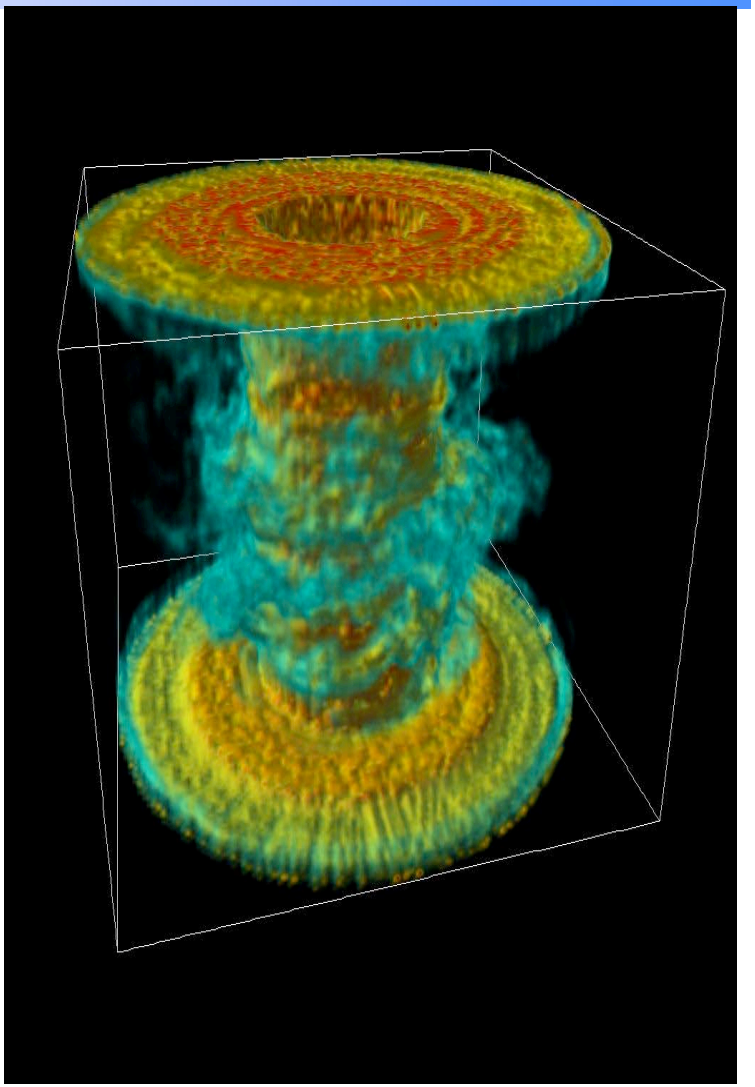

Office of Science
U.S. DEPARTMENT OF ENERGY

INCITE 5 : J. Chen @ Sandia

Science Goals:

DNS of turbulent non-premixed flame that will serve as a benchmark for future theory and experiment.

http://www.ca.sandia.gov/TNF

NERSC Work:

Improved S3D to match Seaborg's SMP structure. Improved memory re-use and vectorization of math functions.

Unexpected Set back:

Usage shifted by 3 months due to surprise lack of re-ignition.

Road Ahead:

Possible shortage of compute time.

Heavy 2048-4096 way seaborg usage.

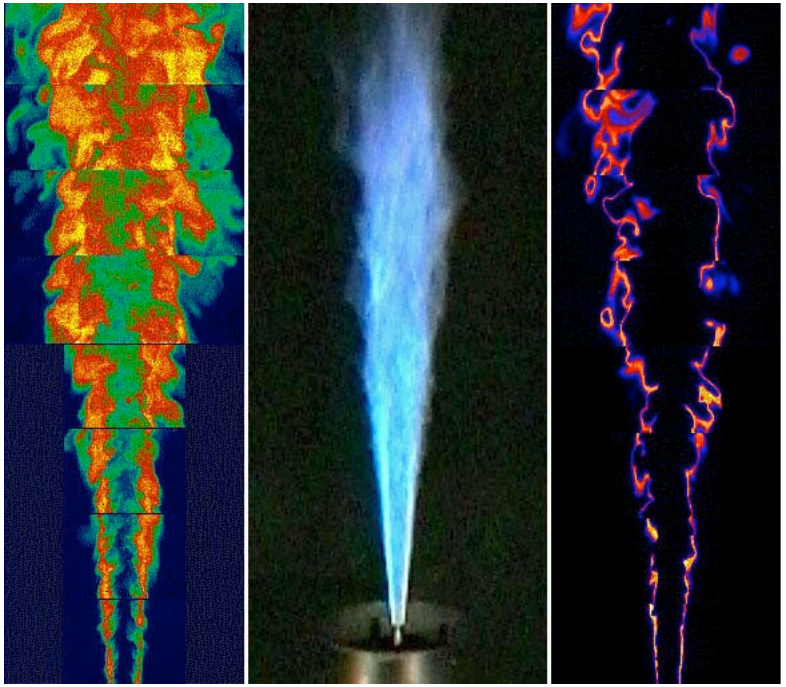

S3D : OMB Joule Metric

Office of Science
U.S. DEPARTMENT OF ENERGY

INCITE 6 : Prof. V. Dagget @ U. Washington

Science Goals:

Catalog dynamical shapes of proteins by systematically unfolding them.

NERSC work:

Load balance in the *ilmm* code
Scalar optimizations
Batch work flow planning

Road ahead:

Finishing up simulations

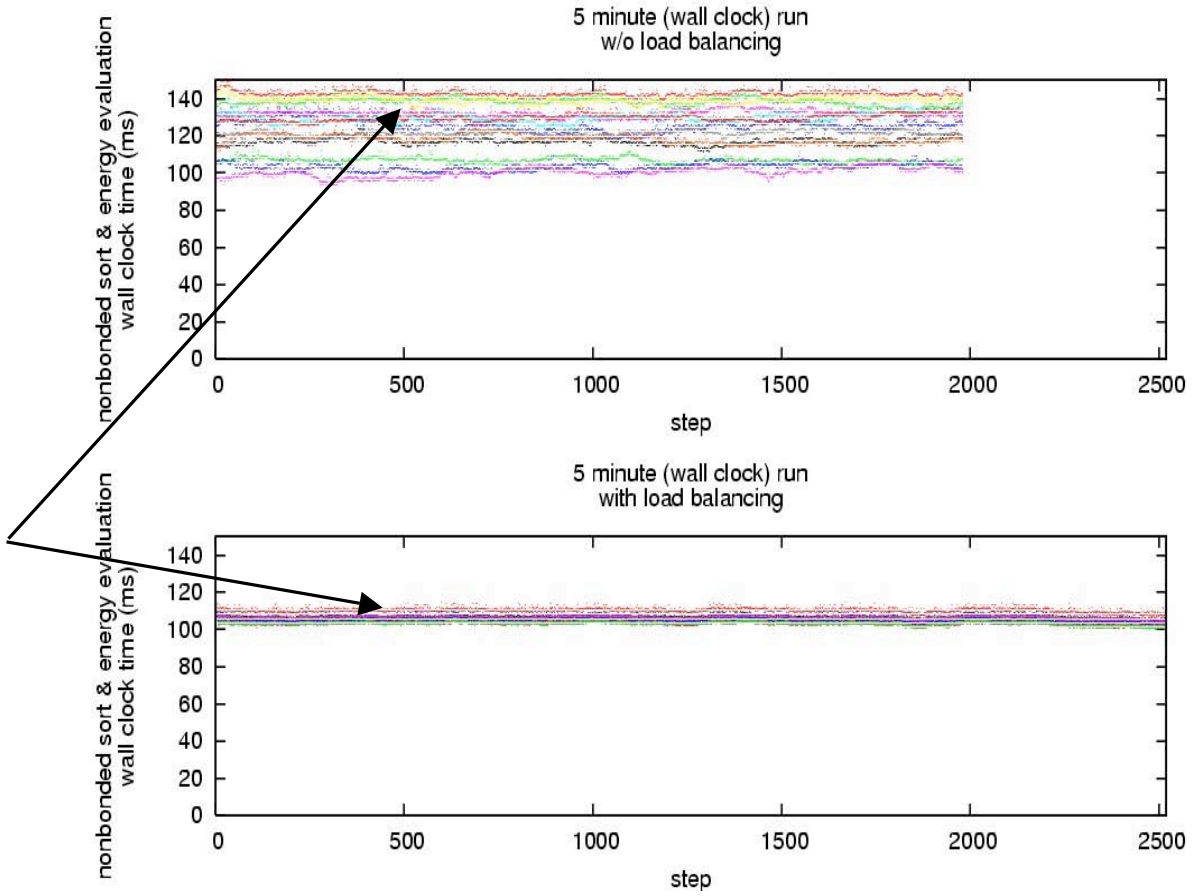

Office of Science
U.S. DEPARTMENT OF ENERGY