PetaApps: Ultra-High Resolution Climate

John M. Dennis
dennis@ucar.edu
July 1, 2010

PetaApps Project

- NSF PetaApps project: Interactive Ensemble
  - Kinter, Stan (COLA)
  - Kirtman (U of Miami)
  - Collins, Yelick (Berkley)
  - Bryan, Dennis, Loft, Vertenstein (NCAR)
  - Bitz (U of Washington)
- Ultra-High resolution Climate
- Computing
  - ~99,000 core Cray XT5 system at NICS [Kraken]
  - Large TG allocation: 35M CPU hours

Other Ultra-High resolution CCSM efforts

- LLNL Grand Challenge Project
  - 0.25° ATM, LND + 0.1° OCN, ICE
  - Bader, McClean, Bryan, Jones, Dennis, Ivanova, Vertenstein, Craig, Norton, Worley, Boyle, Norton, Jones, Mirin, Maltrud, Jacob
  - 20 year run
- Upcoming DOE runs
  - 0.25° ATM, LND + 0.1° OCN, ICE (FV,HOMME)
  - T341 ATM, LND + 0.1°deg OCN,ICE (spectral)

Funding Sources

- Department of Energy: CCPP Program Grants
  - DE-FC03-97ER62402 [SciDAC]
  - DE-PS02-07ER07-06 [SciDAC]
- National Science Foundation:
  - OCI-0749206 [PetaApps]
  - CCF-0937939 [HECURA]
  - OCE-0825754
  - Cooperative Grant NSF01

Large scale PetaApps run

- 155 year control run
  - 0.1° Ocean model [3600 x 2400 x 42]
  - 0.1° Sea-ice model [3600 x 2400 x 20]
  - 0.5° Atmosphere [576 x 384 x 26]
  - 0.5° Land [576 x 384]
- Statistics
  - ~18M CPU hours
  - 5844 cores for 4-5 months
  - ~100 TB of data generated
  - 0.5 to 1 TB per wall clock day generated

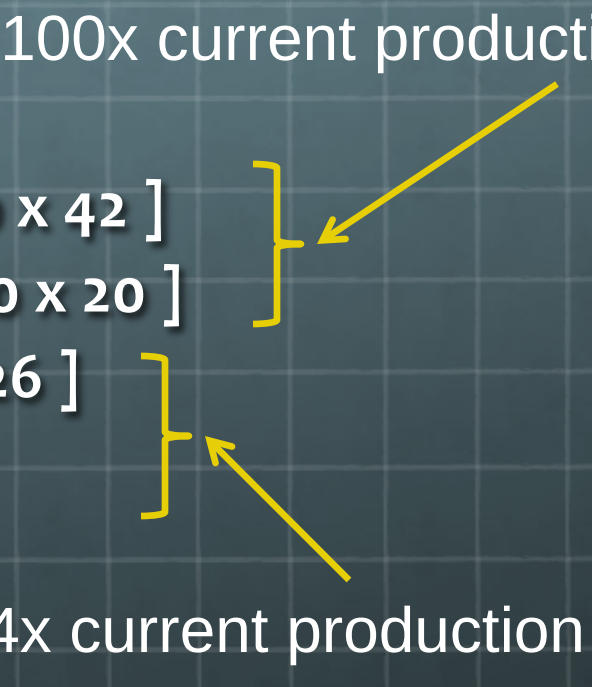

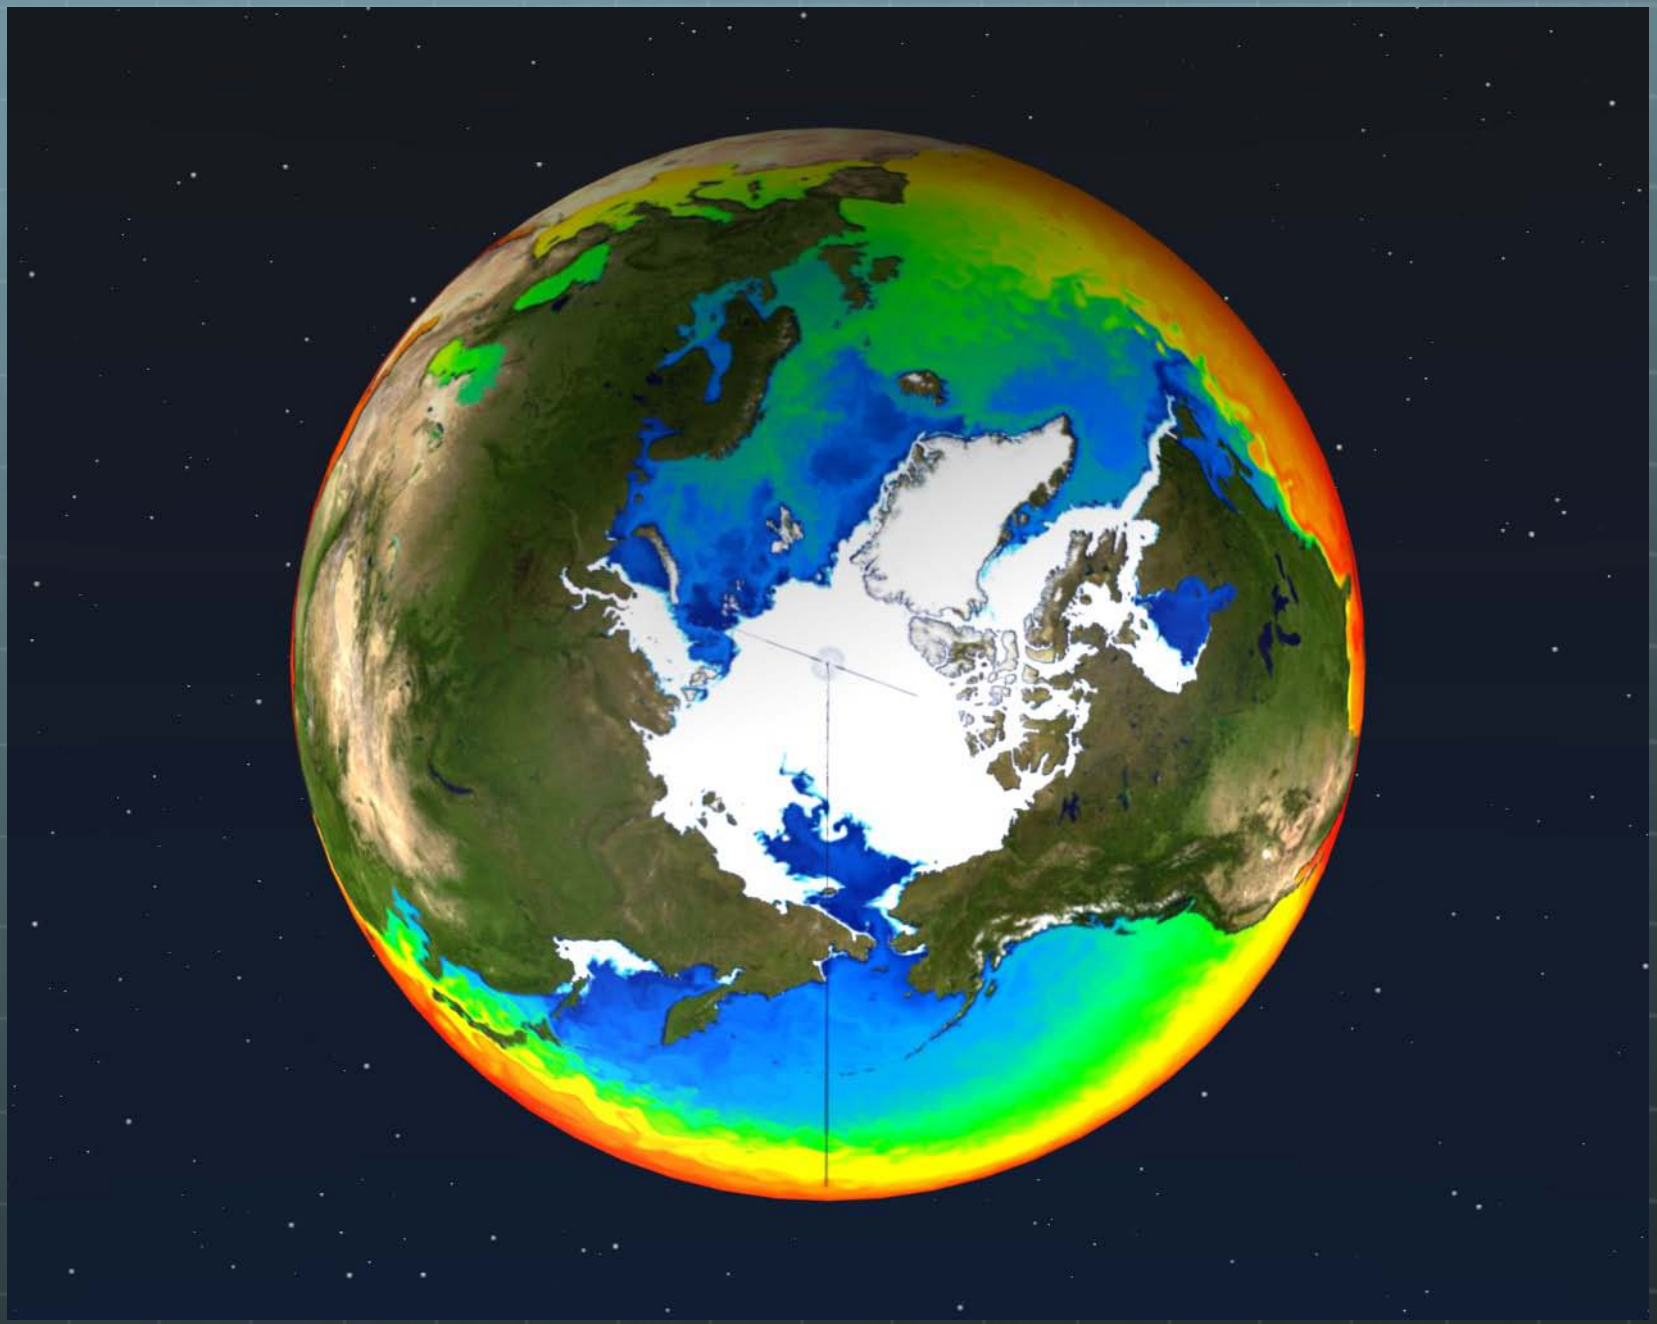

Large scale PetaApps run (con't)

- Work flow
  - Run on Kraken (NICS)
  - Transfer output from NICS to NCAR (100 – 180 MB/sec sustained)
    - Caused noticeable spikes in TG network traffic
  - Archive on HPSS
  - Data analysis using 55 TB project space at NCAR

Issues/challenges with runs

- Reduced cost of simulation by 20%
  - code changes
  - system upgrades
- Very large variability with job run times (MPI message passing)
  - Interference with other jobs
  - 25% of jobs terminated abnormally

Execution time for non-I/O CCSM day

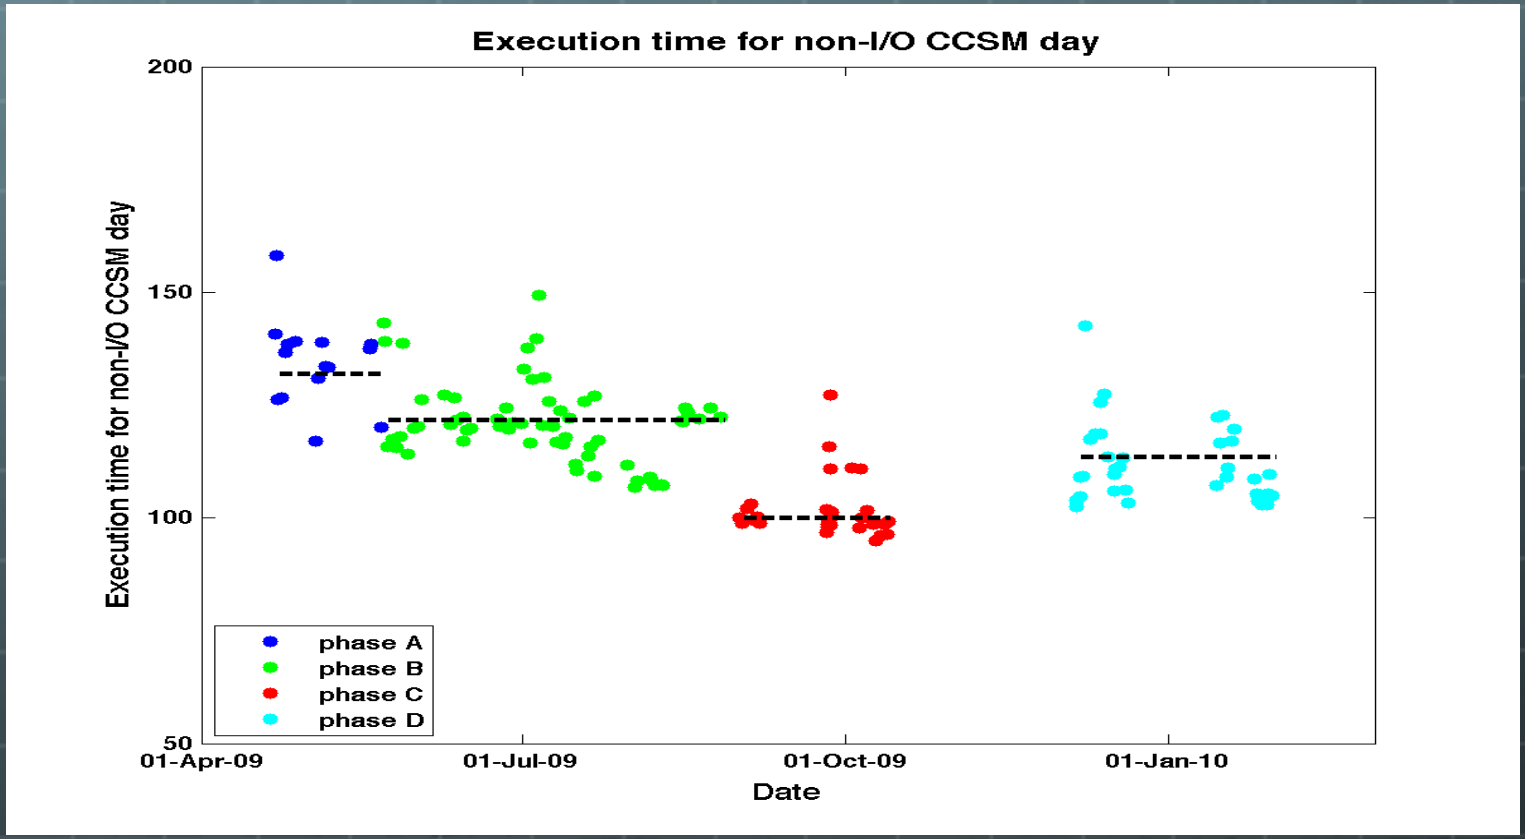

Issues/challenges with runs (con't)

- Very large variability with I/O performance
  - 2-10x slowdown common
  - 300x slowdown was observed
  - Interference with other jobs?

Write bandwidth for CCSM I/O day

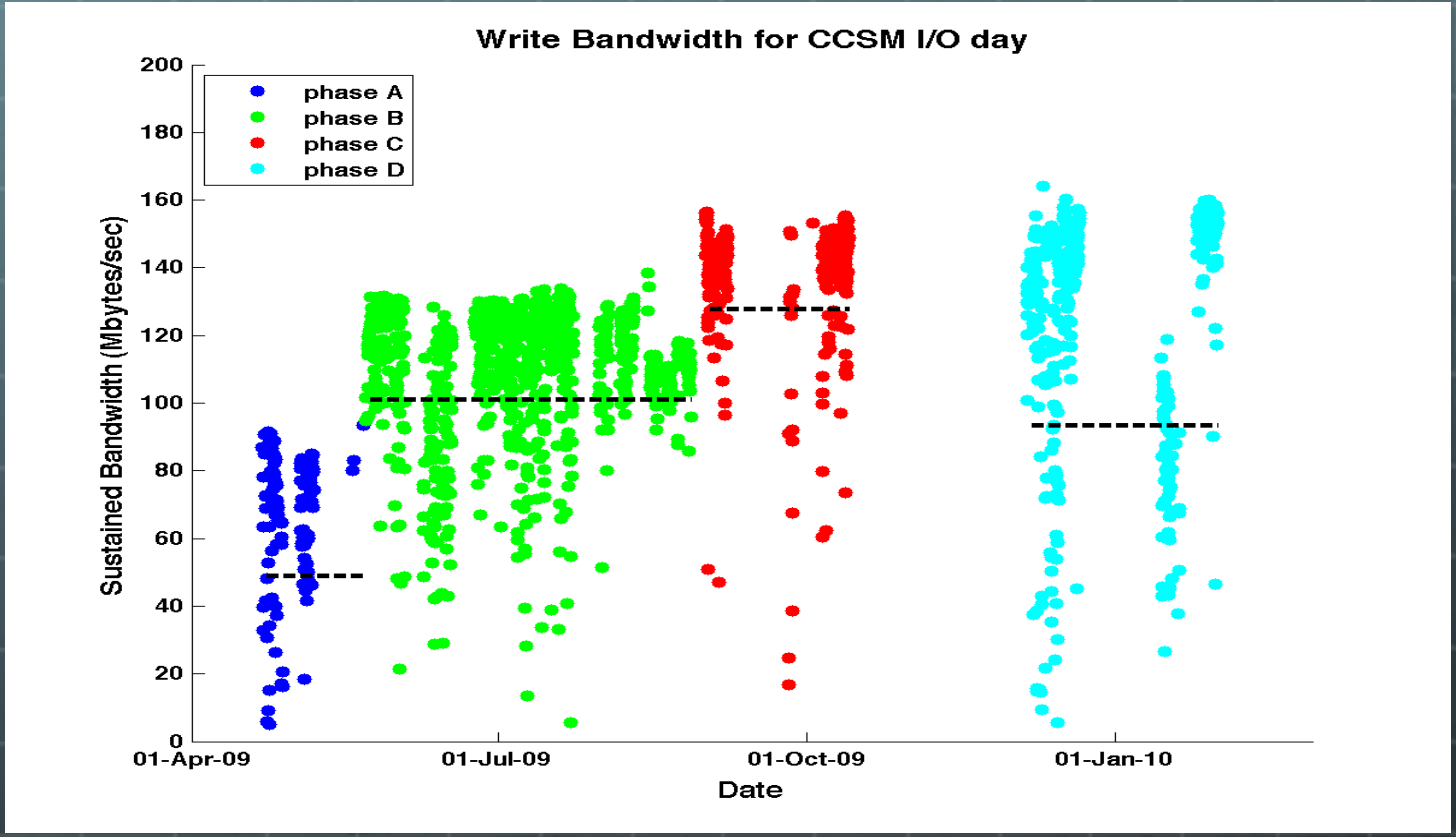

Acknowledgements

- NCAR:
  - D. Bailey
  - F. Bryan
  - T. Craig
  - B. Eaton
  - J. Edwards [IBM]
  - N. Hearn
  - K. Lindsay
  - N. Norton
  - M. Vertenstein
- COLA:
  - J. Kinter
  - C. Stan
- U. Miami
  - B. Kirtman
- U.C. Berkeley
  - W. Collins
  - K. Yelick (NERSC)
- U. Washington
  - C. Bitz
- NICS:
  - M. Fahey
  - P. Kovatch
- ANL:
  - R. Jacob
  - R. Loy
- LANL:
  - E. Hunke
  - P. Jones
  - M. Maltrud
- LLNL
  - D. Bader
  - D. Ivanova
  - J. McClean (Scripps)
  - A. Mirin
- ORNL:
  - P. Worley

and many more…

- Grant Support:
  - DOE
    - DE-FC03-97ER62402 [SciDAC]
    - DE-PS02-07ER07-06 [SciDAC]
  - NSF
    - Cooperative Grant NSF01
    - OCI-0749206 [PetaApps]
    - CNS-0421498
    - CNS-0420873
    - CNS-0420985
- Computer Allocations:
  - TeraGrid TRAC @ NICS
  - DOE INCITE @ NERSC
  - LLNL Grand Challenge
- Thanks for Assistance:
  - Cray, NICS, and NERSC

Questions

dennis@ucar.edu