1. The Role of Climate System Noise in Climate Simulations

- Participating Organization
  - Center for Ocean-Land-Atmosphere Studies (COLA)

    J. Kinter (PI), C. Stan
  - University of Miami

    B. Kirtman
  - University of Washington

    C. Bitz
  - National Center for Atmospheric Research (NCAR)

    J. Dennis(dennis@ucar.edu), R. Loft, M. Vertenstein,
- Scientific Objectives for the next 3 years
  - Examine the impact of noise in the climate system
  - Test hypothesis that noise is “reddened” and influences low-frequency components of climate system
  - Is high-resolution necessary to model climate variability correctly?

2. Current HPC Requirements

- Facilities used:

  NICS: 35M CPU hours [2009]

  NERSC: 2.3M CPU hours [2008 Directors grant] (collaboration)

  LLNL: 7.7M CPU hours [2008 LLNL grant] (collaboration)

- Arch: Cray XT [primary], {BlueGene, Linux Custer} [secondary]
  - 5,800 - 12,000 cores per run
  - 200 jobs x (24 hours per jobs)
  - 2 GB per core
  - File I/O
    - Read from disk: .09 to 30 GB per file, {30 files, 80GB read} per job
    - Written to disk: .09 to 30 GB per file, {180 files, 1044GB written} per job
    - Archived: {60 files, 414 GB} per job
    - Import: minimal
    - Exported: {60 files, 414 GB} per job [Current approach]
- Necessary software, services or infrastructure
  - gridFTP

2. Current HPC Requirements (con't)

- Current primary codes and their methods or algorithms
  - Community Climate System Model consists of components:

    Parallel Ocean Program: method: Finite difference & CG solver, grid size: 3600x2400x42

    Community Atmospheric Model: method: Finite volume, grid size: 576x384x30

    Community Land Model: grid size: 576x384x17

    Community Ice CodE method: Finite difference, grid size: 3600x2400x20
  - All components models support MPI, OpenMP and MPI/OpenMP hybrid
- Known limitations/obstacles/bottlenecks
  - Disk I/O is performed through a single MPI task.
    - 2 of 5 component models now support parallel I/O
  - Disk I/O system variability [*]
- Anything else
  - POP is particularly sensitive to OS jitter. There appears to be scalability issues with POP > 4000 cores on the most recent Cray OS [*]

[*] See supplementary slide

CCSM Sustained Output Bandwidth on Kraken

High = 92 MB/sec

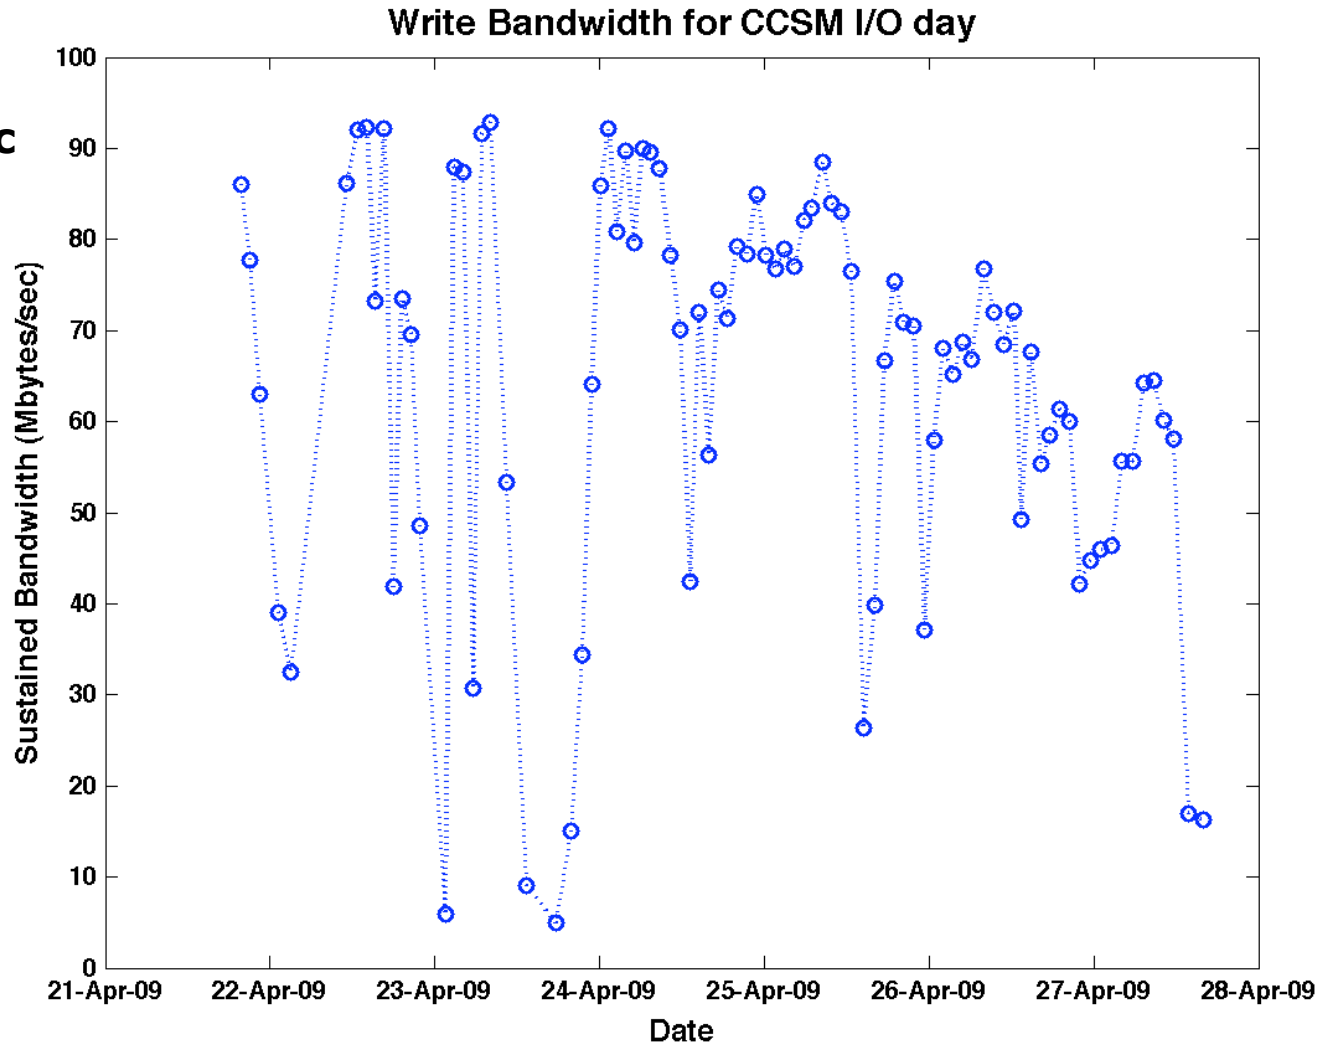

Low = 5 MB/sec

Lustre variability doubled cost of disk I/O from 11% to 23% of total

OS jitter kills POP scalability

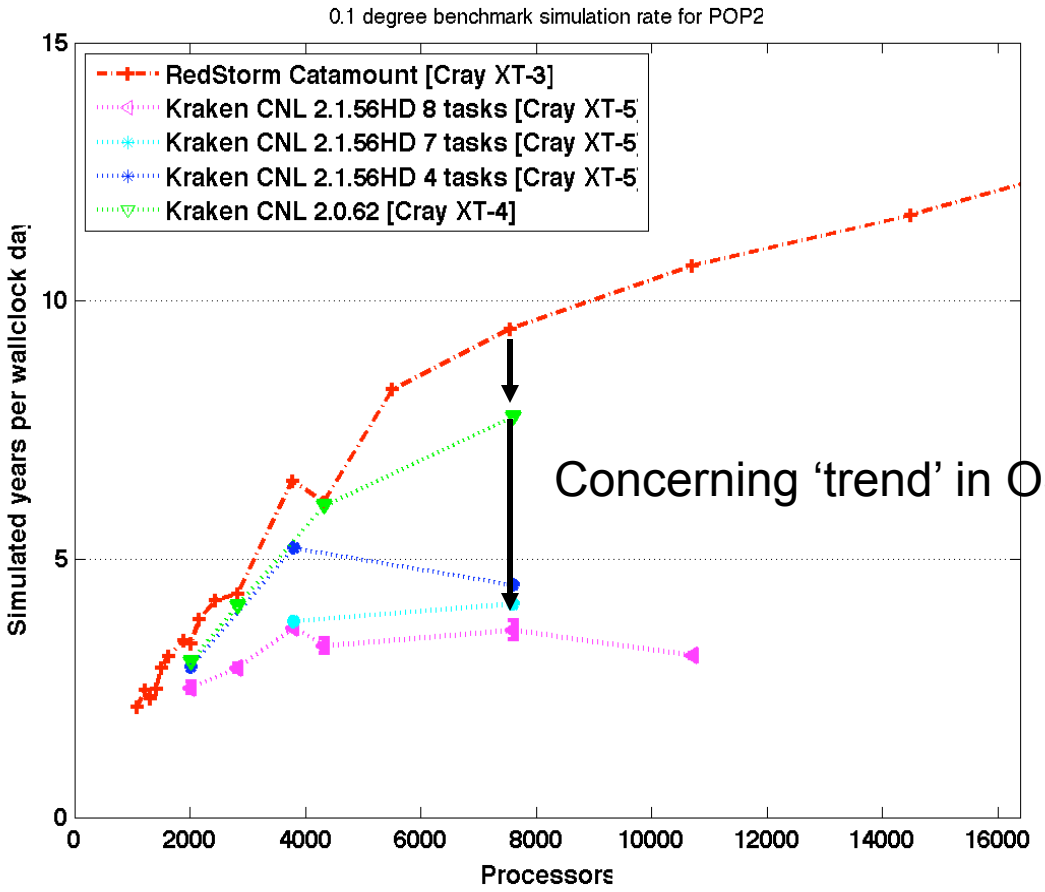

Concerning ‘trend’ in OS/HW evolution?

3. HPC Usage and Methods for the Next 3-5 Years

- Anticipated requirements:
  - 50-100M CPU hours/year
  - 12K-30K cores/job
  - 24 hours/job, 2 GB/core
- Upcoming changes to codes/methods/approaches
  - Improvements to disk I/O subsystem
  - Addition of more scalable atmospheric dynamical core
  - Increase in resolution of existing components
- Changes to Data read/written
  - 5-10x increase in # of files generated [existing size]
  - 5x increase in size of CAM files [.9GB to 4 GB]
- Changes to necessary software, services or infrastructure
  - maintain state-of-art WAN connection
- Anticipated limitations/obstacles/bottlenecks on 10K-1000K PE system.
  - OS jitter a significant obstacle to scalability!

Execution time of Barotropic section of POP on 7600 cores