# SciDAC-2 CCSM Consortium: Software Engineering Update

Patrick Worley
Oak Ridge National Laboratory
(On behalf of all the consorts)

Software Engineering Working Group Meeting
12th Annual CCSM Workshop
June 20, 2007
Village at Breckenridge
Breckenridge, CO

# Acknowledgements

- The work described in this presentation was sponsored by the Atmospheric and Climate Research Division, Office of Science, U.S. Department of Energy. In particular, Worley was supported under Contract No. DE-AC05-00OR22725 with UT-Battelle, LLC.

- This work used resources of the National Center for Computational Sciences at Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under Contract DE-AC05-00OR22725, and of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC03-76SF00098.

# SciDAC-2

Two coordinated SciDAC-2 projects that include CCSM software engineering activities:

- SEESM: A Scalable and Extensible Earth System Model for Climate Change Science (Science Application: Climate Modeling and Simulation**)**
  - Immediate software (and performance) engineering needs, 5 year duration
- PENG: Performance Engineering for the Next Generation Community Climate System Model (Science Application Partnership: Computer Science)
  - More speculative, longer-term performance engineering activities, 3 year duration

# Project Goals

- Software
  - Performance scalability
  - Performance portability
  - Software engineering
- Model Development
  - Better algorithms
  - New physical processes (esp. chemistry, biogeochemistry)

# Talk Overview

- Pointer to March 2007 presentation
- Pointer to Mirin presention
- Collaborations
- Other Recent Activities

# March 2007 Talk Overview

- Recent and Ongoing Activities:
  - CAM
  - POP / Glimmer / CICE
  - CLM
    - Especially CLM Memory Scaling Results!
  - MCT
  - CCSM
  - Recent Performance Data
- Timing Library Update and Discussion

# Mirin Presentation

*Extending Scalability of the Community Atmosphere Model*

- Complete or in progress:
  - Allow auxiliary processes, in preparation for exploiting additional parallelism in atmospheric chemistry and cloud resolving models
  - Allow the latitude/vertical decomposition to have fewer subdomains than the latitude/longitude decomposition in the FV dycore
  - Allow the number of active MPI processes to be smaller in the dynamics than in the physics

# Mirin Presentation

*Extending Scalability of the Community Atmosphere Model*

- Future work
  - Decompose tracer advection with respect to tracer index (implementing latitude/vertical/tracer decomposition)
  - Consider finer vertical decomposition for tracer advection (vs. main dynamics)
  - Consider overlap of tracer advection (*n* tracer subcycle) with main dynamics (corresponding to *(n+1)* tracer subcycle)
  - Performance analysis and optimization of initial implementation

See also poster *Extending Scalability of the Community Atmosphere Model*.

# Collaborations

All activities are collaborative. Significant on-going collaborative efforts include:

1. Parallel I/O (with John Dennis, Jim Edwards, and CSEG)
2. Porting to BG/L (with John Dennis, Jim Edwards, and CSEG)
3. Sequential CCSM (with CSEG)
4. HOMME/CAM integration (with Jim Edwards)
5. Benchmarking cubed sphere FV on Cray XT4 (with Chris Kerr, S-J Lin, Bill Putman)

# Other Recent Activities

- Scalable, reproducible global mean calculation in CAM
  - 0.5x0.625 resolution, 960 processors on XT4:
    original: 1.53 secs/day; new algorithm: .023 secs/day
- High resolution (.25 degree) benchmarking on Atlas system at LLNL (also evaluating compilers and utility of OpenMP):
  - 672 processors: 0.48 simulation years per wallclock day
  - 1344 processors: 0.97 ypd
  - 2688 processors: 1.28 ypd

  (without scalable global mean)
- Optimizing and benchmarking POP performance on the Cray XT4

# Other Recent Activities

- Porting code that allows CAM to use more processes in physics than in dynamics into CSU version of CAM that uses superparameterization (for Marat Khairoutdinov at CSU)
  - Motivation is to be able to run on systems without OpenMP (IBM BG/L and Cray XT4).
  - More efficient than running with OpenMP (16 threads per process) on blueice at NCAR for T42/SLD on 512 processors (using SMT, so 1024 virtual processors).
- Optimizing and benchmarking CAM performance on Intel clusters (with Dmitry Shkurko at Intel)
  - See poster *CAM performance on Intel-based clusters*.

# POP Benchmarking

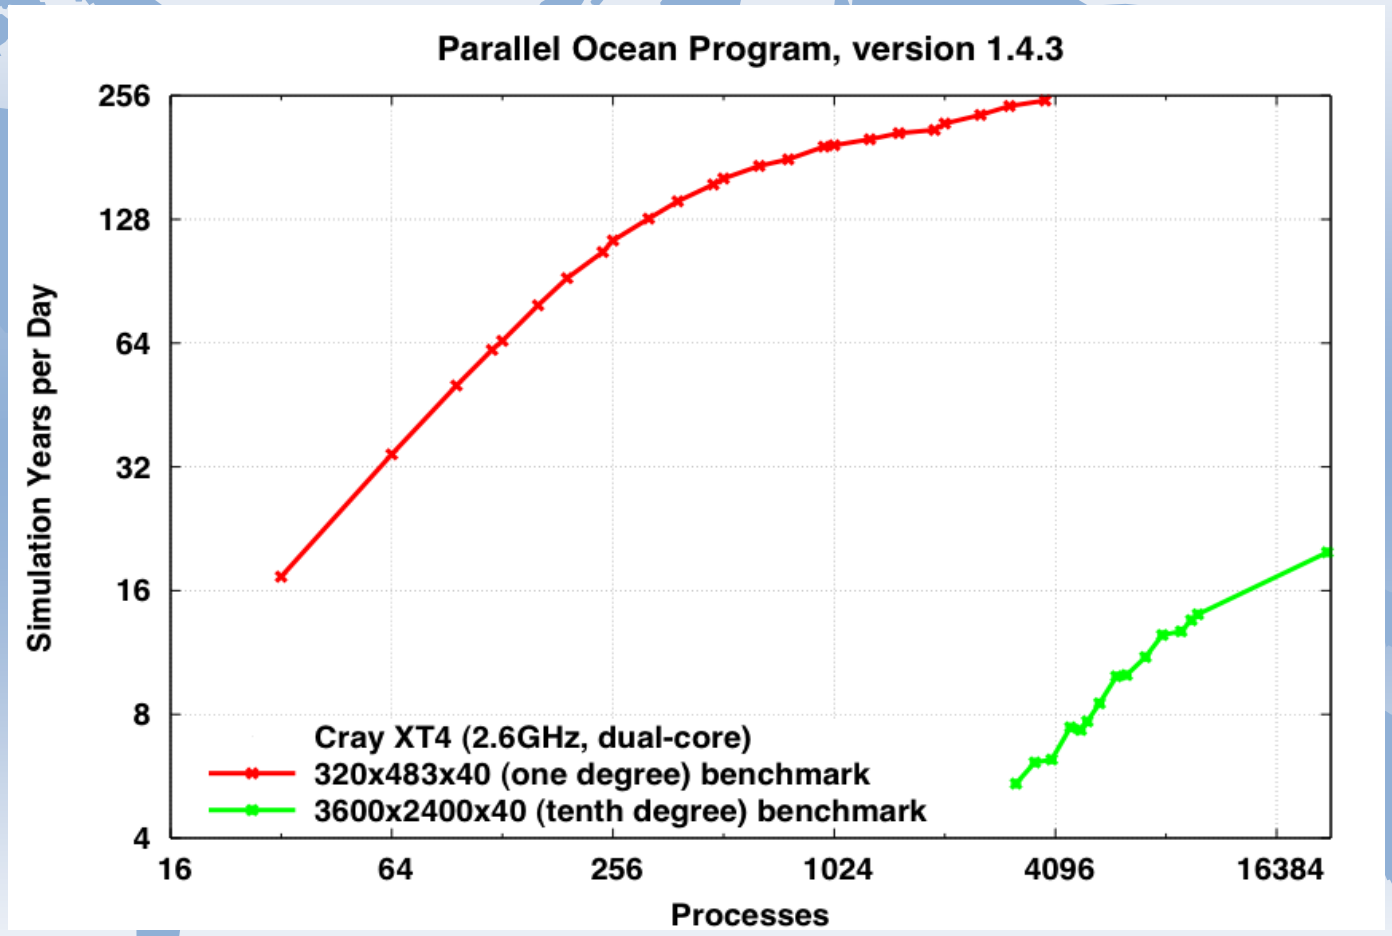


For both one degree and tenth degree benchmark problems, POP can use thousands of MPI processes effectively on the Cray XT4. Using Chronopoulos-Gear variant of conjugate gradient solver and (Cray XT-specific) modification to MPI allreduce algorithm.

# CAM Benchmarking on new Intel Clusters

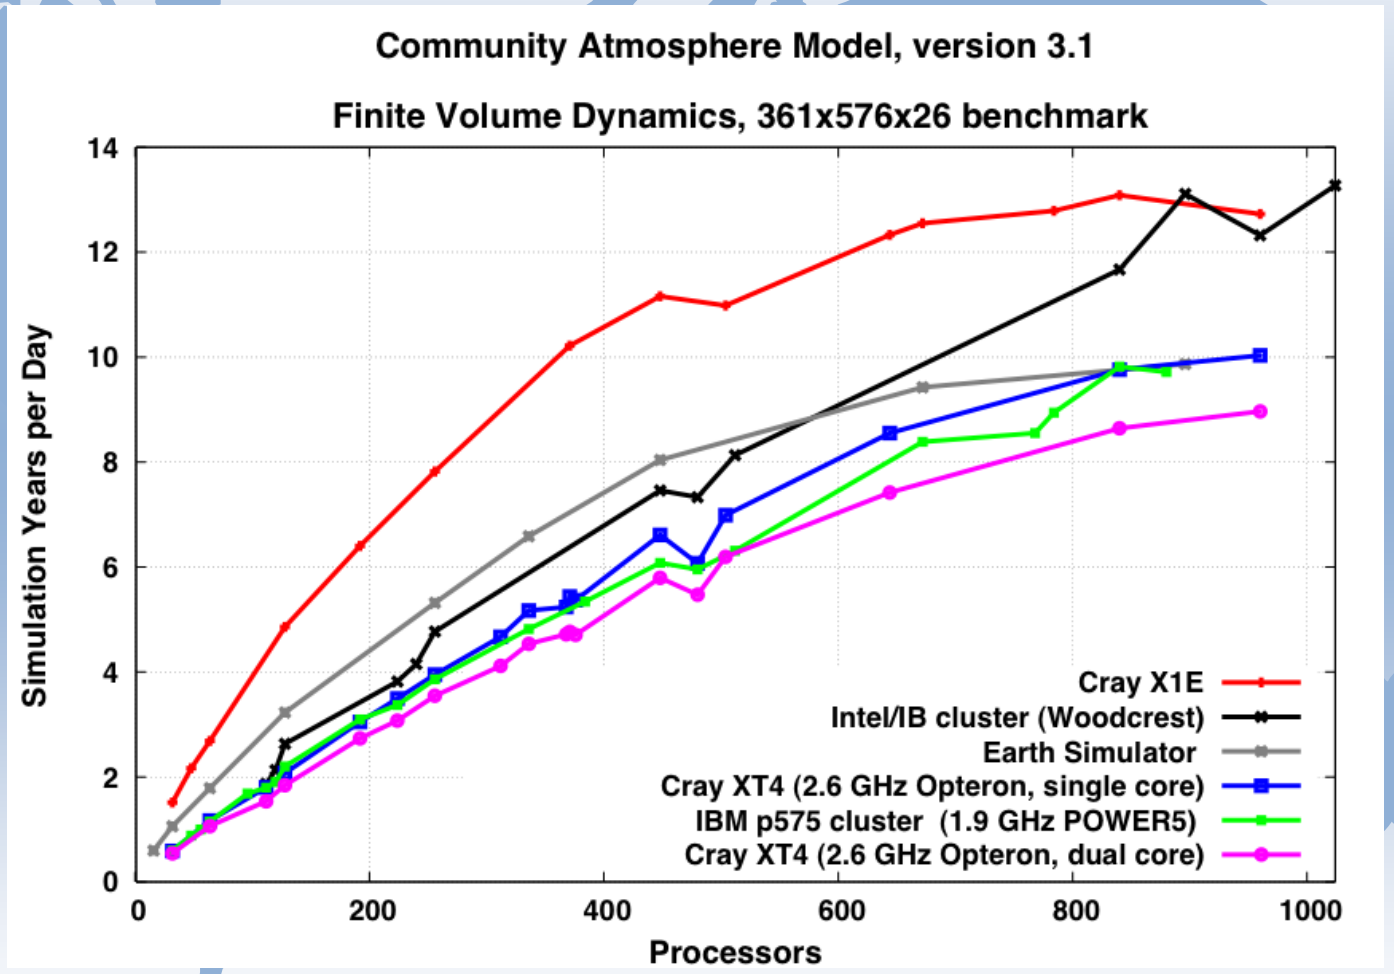

# Cubed Sphere FV Dycore Performance

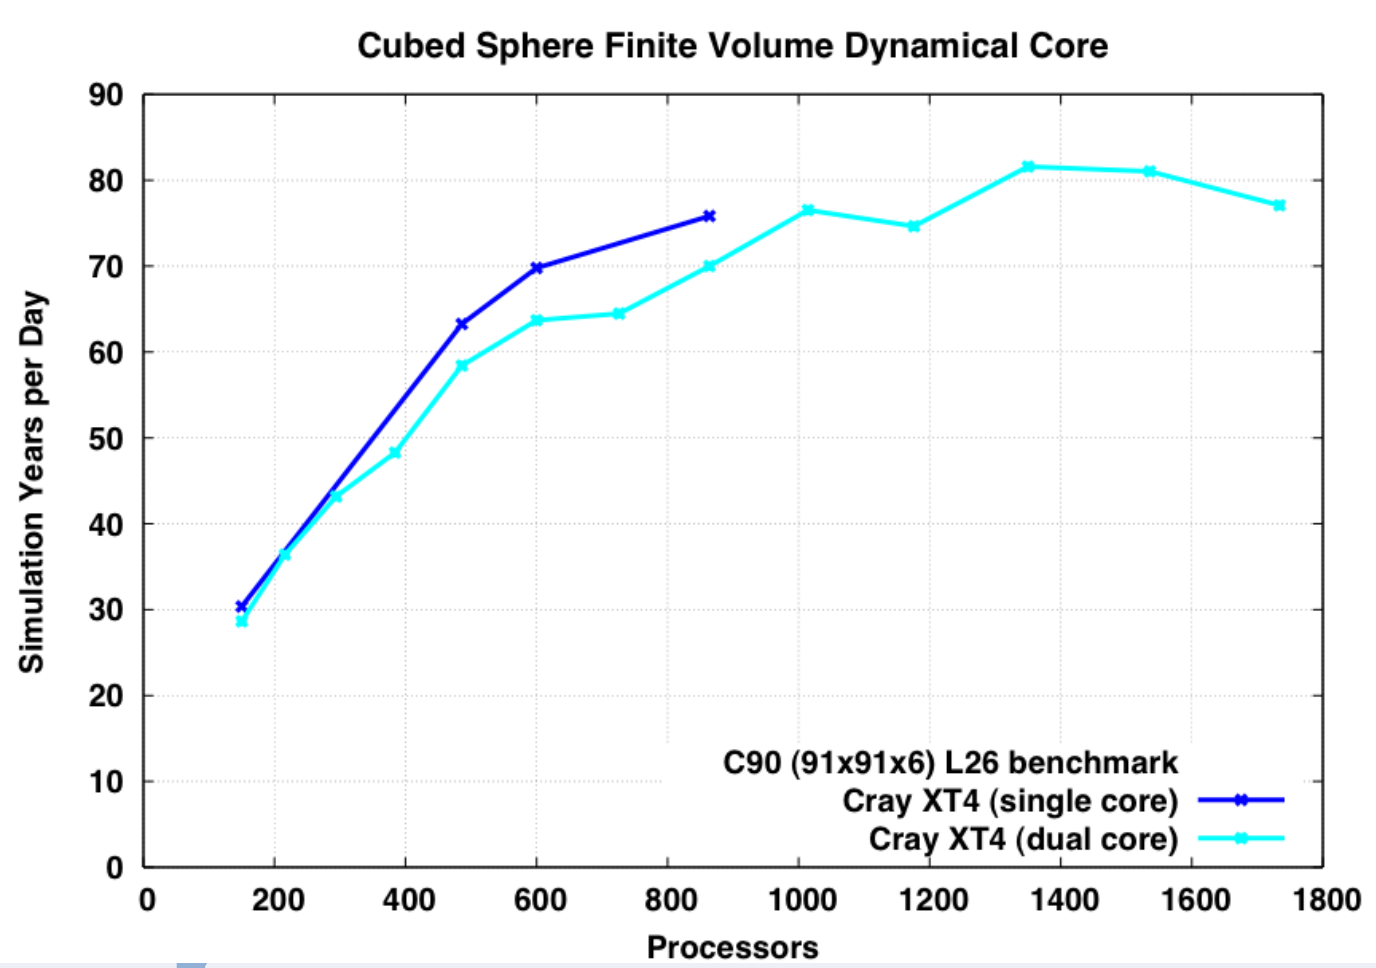


Initial results. Working with C. Kerr, S-J Lin and W. Putman to evaluate performance on Cray XT4. Repeat from March presentation.

# Questions? Comments?