

SciDAC Update

Pat Worley
Oak Ridge National Laboratory

Software Engineering Working Group Meeting
16th Annual CESM Workshop
The Village at Breckenridge
Breckenridge, CO
June 23, 2011



Alternative Title


Selected SciDAC Software Engineering Activities

Pat Worley
Oak Ridge National Laboratory

Software Engineering Working Group Meeting
16th Annual CESM Workshop
The Village at Breckenridge
Breckenridge, CO
June 23, 2011



We are (for 3 more months) ...



**A Scalable and Extensible Earth System Model
for Climate Change Science
SEESM – aka SciDAC CCSM**

and

**Performance Engineering for the
Community Climate System Model
PECCSM – aka SciDAC CCSM SAP**



Some of us also are *or may be* ...

- Development of Frameworks for Robust Regional Climate Modeling
- Ice Sheet Initiative for CLimate ExtremeS (ISICLES), including the following working directly with the ice sheet component in CESM (CISM):
 - High-Performance Adaptive Algorithms for Ice Sheet Modeling (B-ISICLES)
 - A Scalable, Efficient, and Accurate Community Ice Sheet Model (SEA-CISM)
- Investigation of the Magnitudes and Probabilities of Abrupt Climate TransitionS (IMPACTS)
- Ultra High Resolution Global Climate Simulation
- Visual Data Exploration and Analysis of Ultra-Large Climate Data
- ...
- **Climate Science for a Sustainable Energy Future (CSSEF)**
- **Exascale Performance Research for Earth System Simulation (EXPRESS)**
- **SciDAC3 Science Application projects in climate**
- **SciDAC3 Science Application Partnership projects in climate**
- ...


CESM-oriented software engineering activities for these other Department of Energy projects will not be covered in this presentation.

The SciDAC CCSM Consortium:

- **PI:** Phil Jones, Los Alamos National Laboratory
- **Argonne National Laboratory (ANL)** Robert Jacob, Jay Larson, Sheri Mickelson
- **Lawrence Berkeley National Laboratory (LBNL)** William Collins, Michael Wehner, Inez Fung
- **Lawrence Livermore National Laboratory (LLNL)** Phillip Cameron-Smith, Arthur Mirin
- **Los Alamos National Laboratory (LANL)** Scott Elliot, Philip Jones, William Lipscomb, Mat Maltrud
- **National Center for Atmospheric Research (NCAR)** Peter Gent /Jim Hurrell, Andrew Conley, Tony Craig, Jean-Francois Lamarque, Mariana Vertenstein, Trey White, Warren Washington
- **Oak Ridge National Laboratory (ORNL)** Marcia Branstetter, David Bader, David Erickson, Kate Evans, James Hack, Forrest Hoffman, Peter Thornton, Patrick Worley
- **Pacific Northwest National Laboratory (PNNL)** Steven Ghan, Xiaohong Liu, Richard Easter
- **Sandia National Laboratories (SNL)** Mark Taylor

- **Scientific Application Partnerships**
 - **Lawrence Livermore National Laboratory** Arthur Mirin
 - **Oak Ridge National Laboratory** Patrick Worley

- **Centers for Enabling Technology Collaborations**
 - ESG** - Dean Williams **PERI** – Pat Worley
 - VIZ** – Wes Bethel **TOPS** – David Keyes
 - ITAPS** – Lori Diachen



Computational Climate End Station for Climate Change Science
- Warren Washington (PI)



Objectives

The primary objective of the two projects is to develop, test, and exploit a first generation of Earth system models based upon the CCSM that will run efficiently on thousands of processors and include significant model enhancements.

Tasks and Organization

- **Scalable and Extensible Earth System Model (SEESM)**
 - terrestrial carbon cycle and dynamic vegetation
 - atmospheric chemistry and aerosol dynamics
 - ocean ecosystems and biogeochemical coupling
 - feedbacks between atmospheric composition and biogenic emissions
- **Model Integration and Evaluation (SEESM)**
 - integration and unit testing
 - new ice sheet and ocean models
 - new atmospheric dynamics: finite volume (cubed sphere), continuous Galerkin, others
 - frameworks for model evaluation
- **Computational Performance (SEESM and PECCSM)**
 - scalability to thousands of processors, load balance, (fault recovery)
 - performance portability and software engineering

Engineering-related Presentations

AMWG meeting

1. CAM-SE: The HOMME dynamical core

LIWG meeting

2. BISICLES: A high-performance adaptive higher-order thermomechanical ice sheet model
3. Progress towards high-resolution continental-scale ice sheet simulations using a higher-order dynamical core in Glimmer-CISM

OMWG meeting

4. Fully-coupled, Fine-resolution CCSM Simulations: A prototype and planned advancements
5. Status of MPAS ocean model

SEWG meeting

6. MPAS and CESM
7. Introducing Parvis
8. Progress towards accelerating CAM-SE on hybrid multi-core systems
9. CSEG Update

Workshop posters

1. Computer Perf. of the CESM
2. Parvis
3. MPAS-Ocean
4. Adaptive Mesh Modeling of Pine Island Glacier
5. A number of MPAS and SE model development AMWG posters

(Again - not a comprehensive overview or status report of the SciDAC project)

1. Recent software engineering activities
2. Model development in support of high resolution studies
3. CAM-SE development
4. Performance evaluation
 - a) Low resolution fully active
 - b) High resolution dycore evaluation
 - c) High resolution fully active
5. New performance analysis tools

Recent Software Engr. Activities

(All in collaboration with CSEG)

- DOE platform support (Atlas, Hopper, Intrepid, JaguarPF)
- CESM maintenance and development, including
 - Communication algorithms (all components)
 - Reproducible global sum (finally done!)
 - Monitoring functionality and performance, and debugging, for all components
- CESM configuration debugging and benchmarking, including
 - POP space-filling-curve decomposition evaluation tool (Dennis/Mickelson/Worley)
 - Extensive low and high resolution studies (see later slides) and advising CESM users on efficient configurations
 - First successful fully active (B1805CN) eighth degree atmosphere/tenth degree ocean run (ne240f23_t12)
 - First CESM run to use (productively) nearly full JaguarPF system (204,408 cores).

(A. Mirin)

Very-high-resolution atmospheric simulations with interactive chemistry and sectional aerosol microphysics

- Explore how regional variations of aerosol and precursor emissions affect global and regional climate as well as key societal impacts, including air quality and hydrology
 - more comprehensive chemistry mechanism, including secondary organic aerosols
 - sectional aerosol microphysics – more precise than modal approach, in that aerosols are partitioned into size bins (currently 8)
- Computational challenges
 - 335 tracers
 - I/O – high resolution requires *pnetcdf*
 - use of module variables within parallel region requires thread-private
- Results to date
 - proof of principle - have successfully executed *trop_mozart* calculations with conceptual 399 tracers at 0.25-deg, including history output and restarts
 - have executed full sectional model with *cam5* physics at 1-deg (and at 0.5-deg, but without restarts – merge with trunk should enable restarts)
- Work supported by LLNL Institutional Computing Grand Challenge program

Improvements to CAM-SE

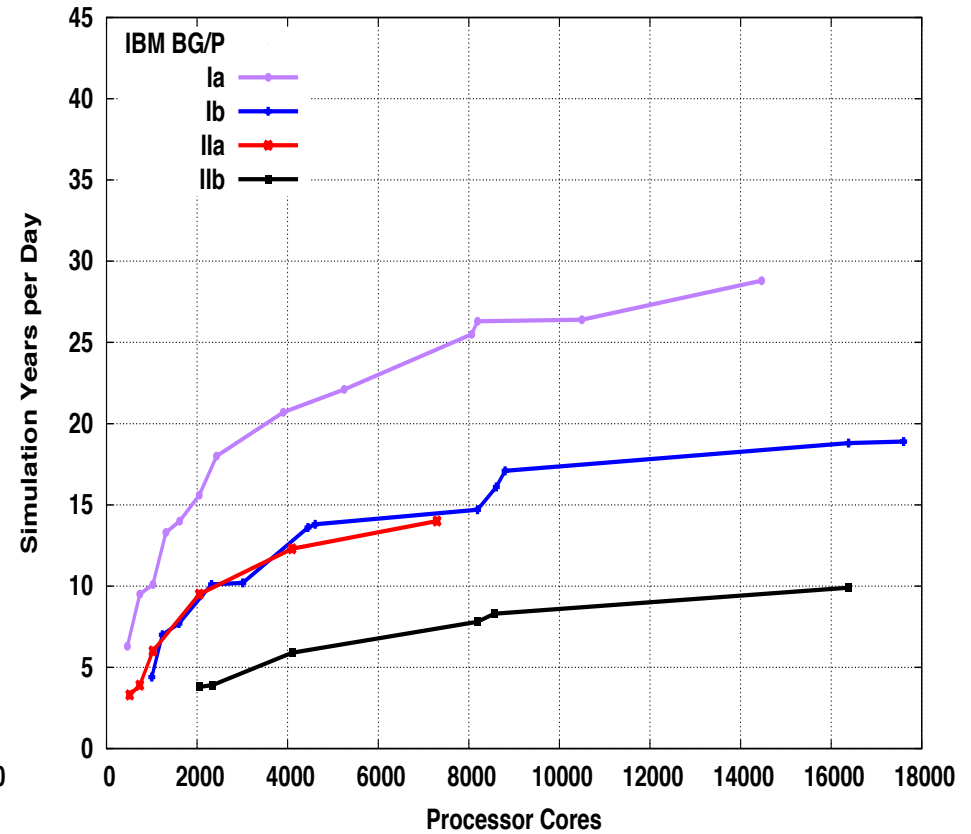

(A. Mirin)

- Capability to use more processes for physics than dynamics has been added
 - important for scenarios with intense physics
 - such capability already exists for other dycores
- Ability to invoke OpenMP within element has been added
 - OpenMP is with respect to level, tracer index, or horizontal
 - OpenMP paradigm to date is across elements, which competes with MPI-parallelism
 - OpenMP within element enables greater degree of parallelism
 - decision of which OpenMP paradigm to use must be decided at compile time (due to non-support for nested constructs)
- For ne30np4 (1-deg) configuration with trop_mozart chemistry and cam5 physics, using (the maximum) 5400 tasks, increasing from 1 to 4 OpenMP threads results in a factor of 2.2 speedup of dynamics on Intrepid.

CESM Performance: Low Res.

I: 1.9x2.5_gx1v6 (FV dycore)
 II: 0.9x1.25_gx1v6 (FV dycore)

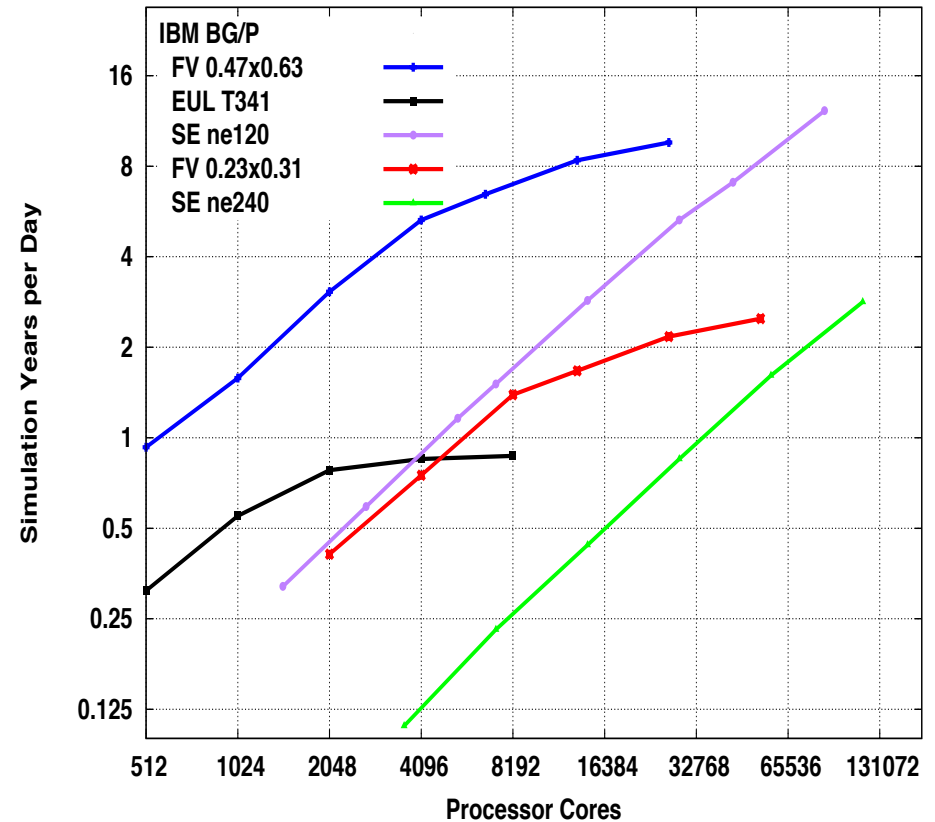

a: B1850CN
 b: B1850_CAM5



- Using more expensive “CAM5” physics decreases throughput by less than a factor of 2 for same processor core count .

CAM Performance: High Res.

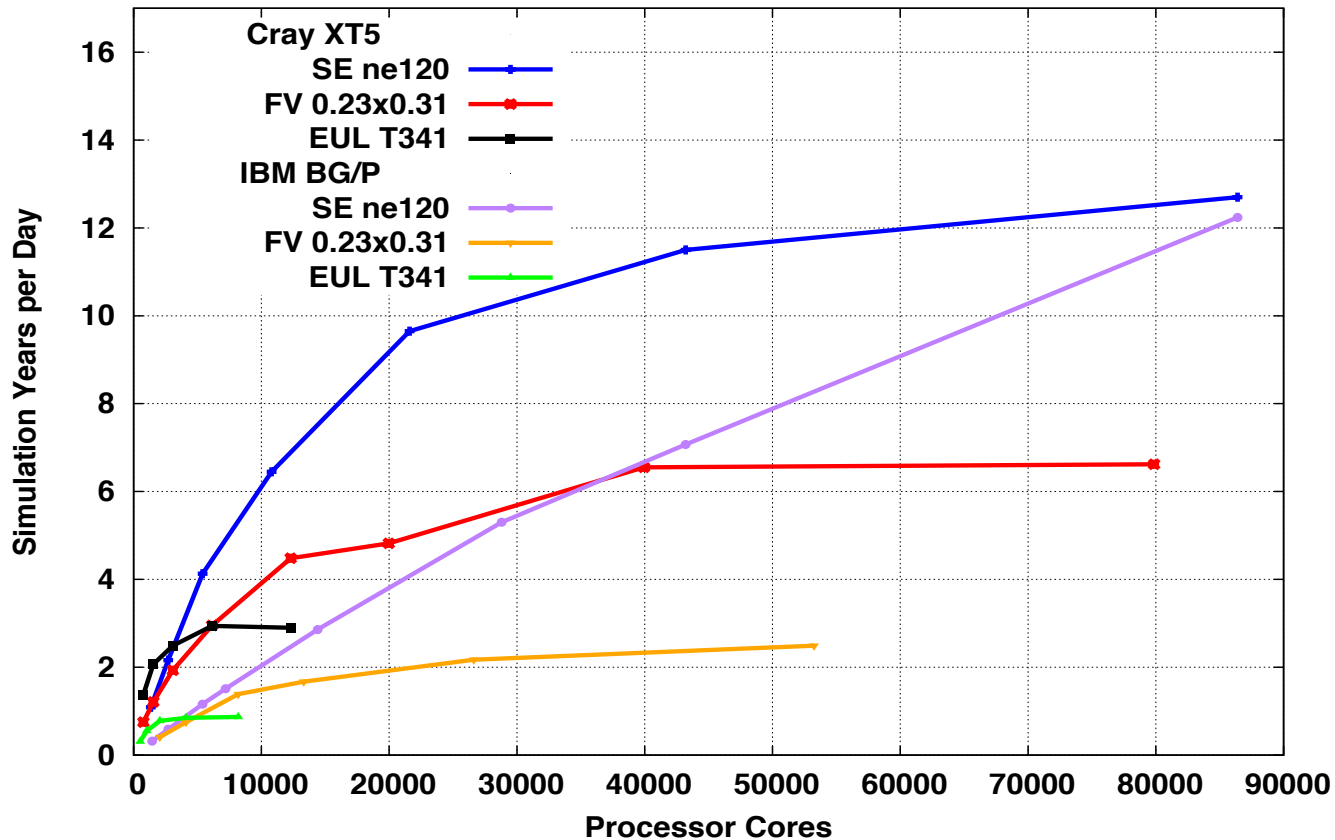
F_1850



- Atmosphere-only timings
- EUL T341, SE ne120, and FV 0.23x0.31 achieve similar solution fidelity; comparing among these SE much more scalable on both systems.

CAM Performance: High Res.

F_1850




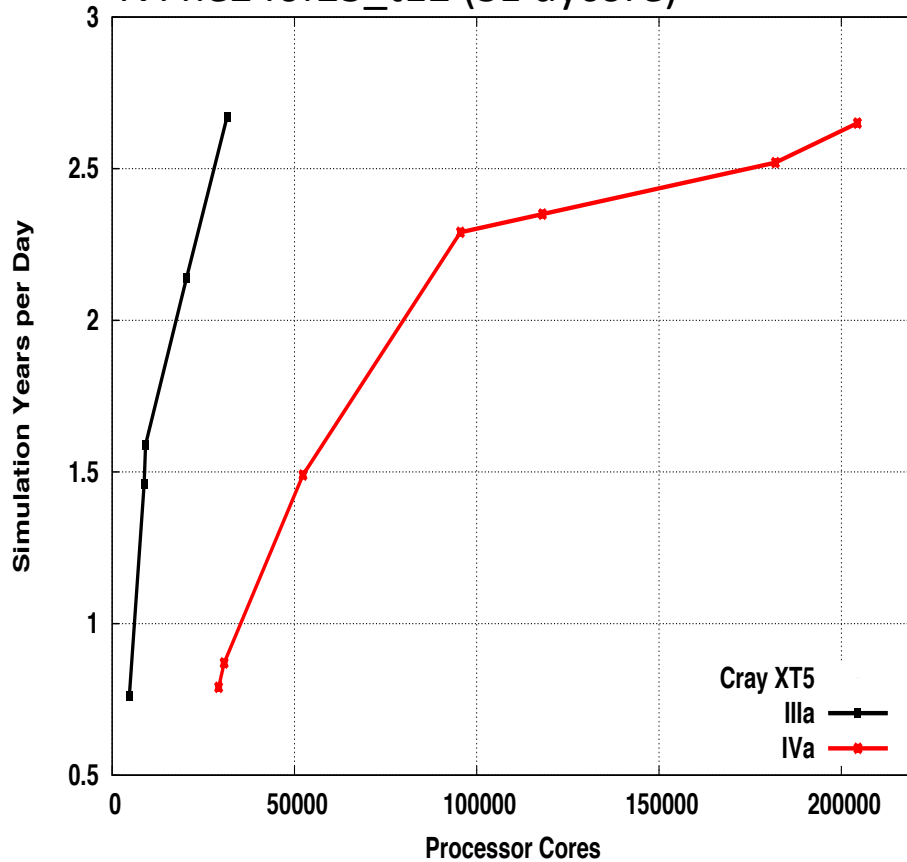
- SE scales very well on BG/P, especially compared to XT5. Similar behavior holds for ne240 (see previous slide), but BG/P ne240 perf. is still half that of XT5 up to max size attempted.
- Other dycores do not show same preference for BG/P over XT5.

CESM Performance: High Res.

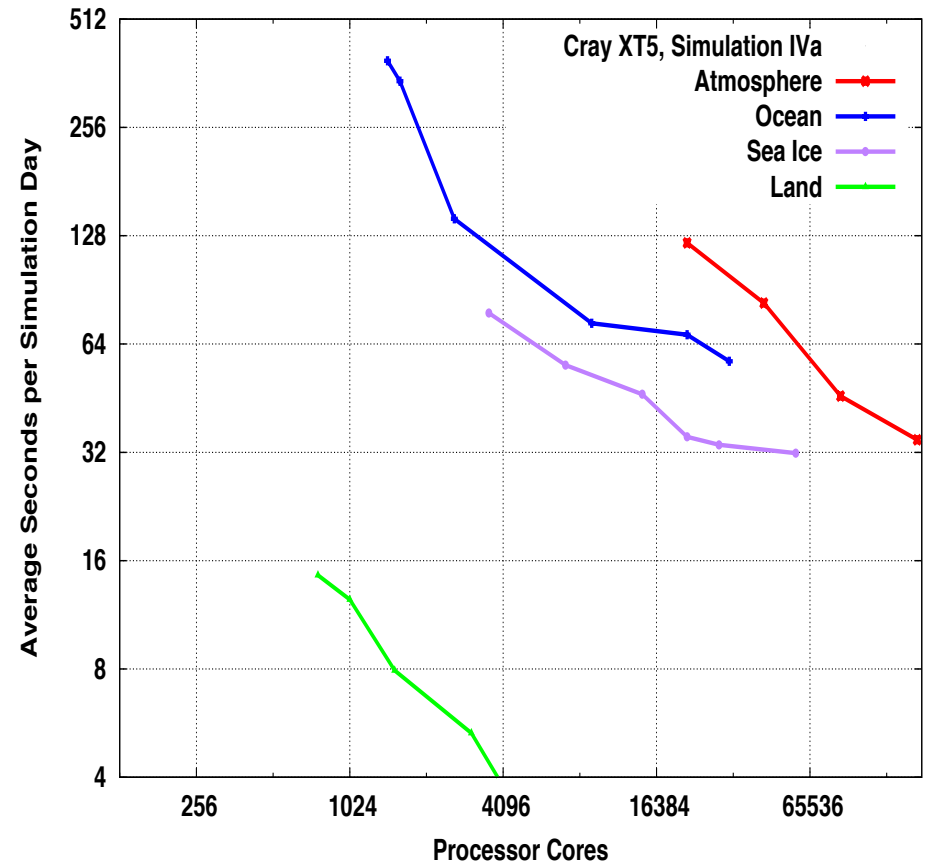

III: f02_t12 (FV dycore)

a: B1850CN

IV: ne240f23_t12 (SE dycore)

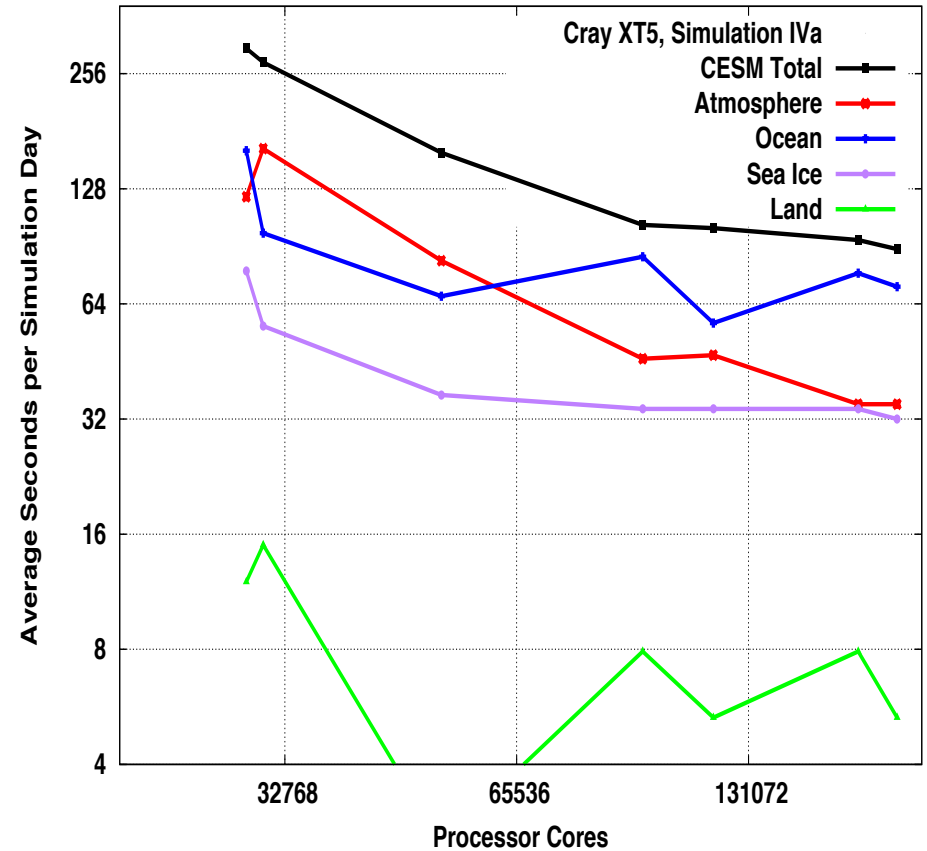
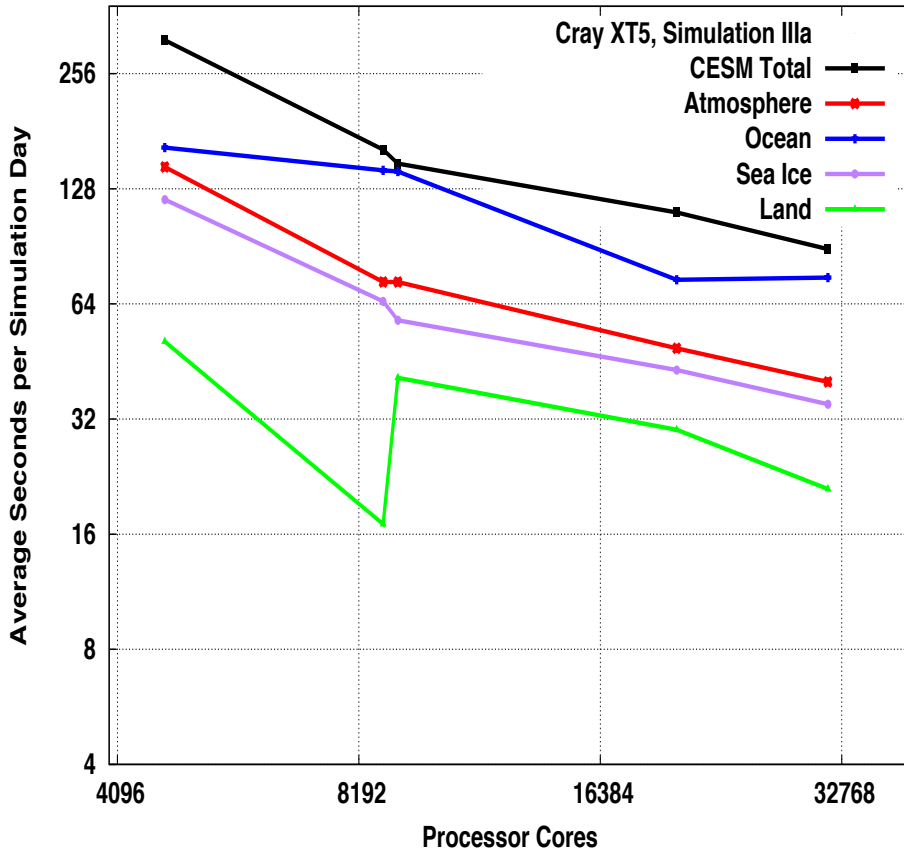


- Two views of same data. Both IIIa and IVa curves include practical maximum processor core count for each configuration on this XT5 system.



- As part of optimization process, varied processor core count for each component. Plotting component-specific performance within a coupled run as a function of processor core count assigned to that component. Other than atmosphere, component performance the same for IIIa and IVa (as it should be).

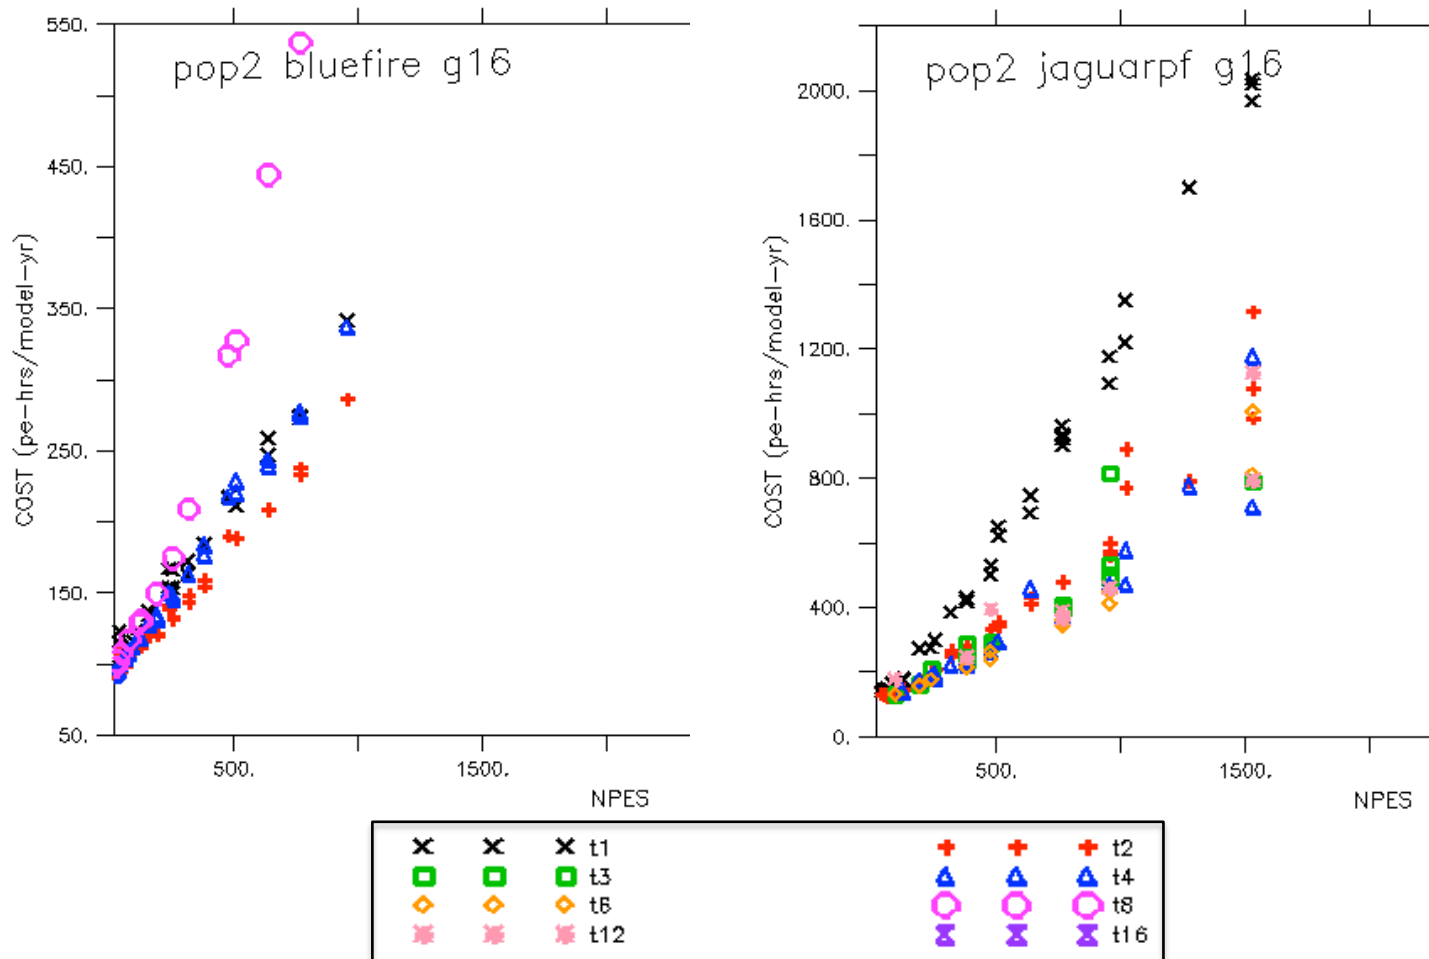
CESM Performance Diagnosis



- Plotting time spent in each component for best observed CESM performance as function of total CESM processor core count. Ocean runs concurrently with other components; land and sea ice run sequentially with atmosphere. Load balancing the components can lead to nonmonotonic behavior in component timing curves.

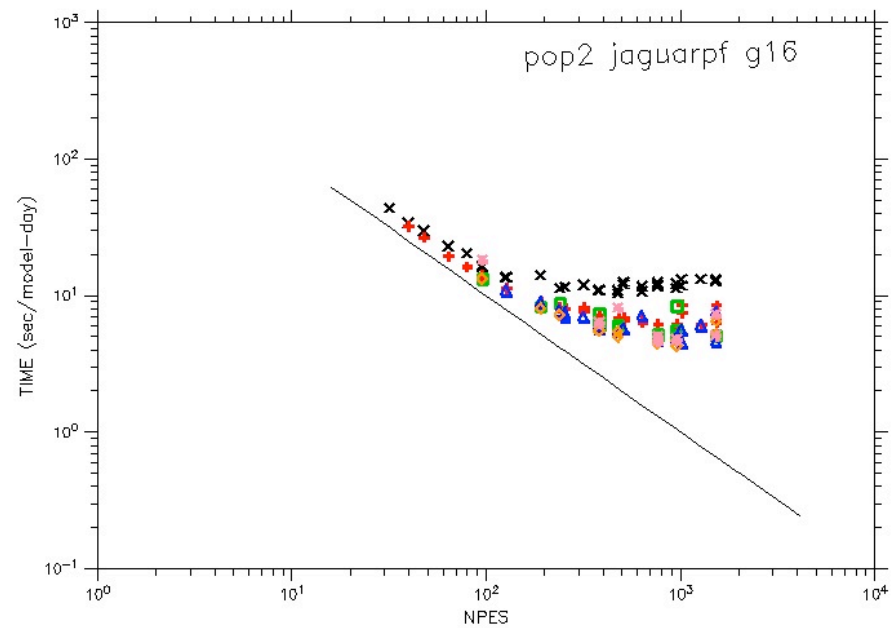
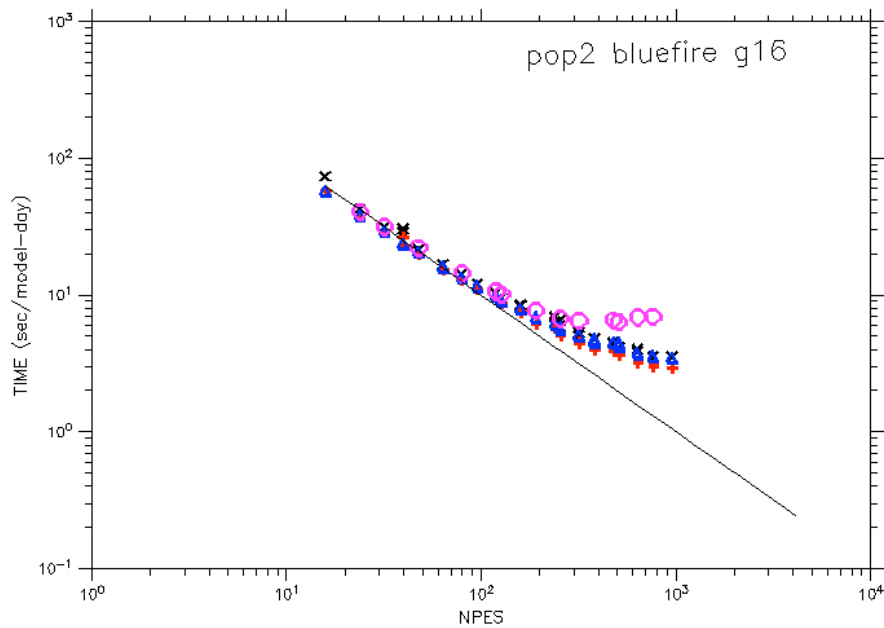
Performance Evaluation Tools

(A. Craig)



Performance Evaluation Tools

(A. Craig)



Ongoing Activities

In addition to (never ending) maintenance and development activities:

1. Spectral element dycore performance diagnosis
2. Ocean performance diagnosis and optimization, including
 - Re-examining reproducibility implementation and overhead
3. Sea ice performance diagnosis and optimization, including
 - CICE space-filling-curve decomposition evaluation tool
 - Adaptive MPI framework evaluation
4. CPL performance diagnosis and optimization
 - Just now showing up as a performance issue in very high resolution configurations
5. New architecture evaluations
 - Hopper (especially Gemini network and support for one-sided communications)
 - GPU accelerators
6. Improved default configuration and performance option settings based on empirical performance data

Acknowledgements

(P. Worley)

- The work of P. Worley was sponsored by the Climate and Environmental Sciences Division of the Office of Biological and Environmental Research and by the Office of Advanced Scientific Computing Research, both in the Office of Science, U.S. Department of Energy, under Contract No. DE-AC05-00OR22725 with UT-Battelle, LLC.
- This work used resources of the Oak Ridge Leadership Computing Facility, located in 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, of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-06CH11357, 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-AC02-05CH11231.
- These slides have been authored by a contractor of the U.S. Government under contracts No. DE-AC05-00OR22725 and No. DE-AC52-07NA27344. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.