An overview of the new CAM benchmark suite

Patrick Worley

Oak Ridge National Laboratory

CCSM Atmosphere Model Working Group Meeting
February 12, 2008
National Center for Atmospheric Research
Boulder, CO

Acknowledgements

- The work described in this presentation was sponsored by the Atmospheric and Climate Research Division and by the Office of Mathematical, Information, and Computational Sciences, both of the Office of Science, U.S. Department of Energy, under Contract No. DE-AC05-00OR22725 with UT-Battelle, LLC.
- These slides have been authored by a contractor of the U.S. Government under contract No. DE-AC05-00OR22725 with UT-Battelle, LLC. 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.
- 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.

Goals

- Goals of Benchmarking Activity
  - Quantify change in cost of model due to model evolution and to choice of model configuration.
  - Establish new performance baselines upon which to base future performance optimization activities, and with which to evaluate computing platforms.
- Goals of Talk
  - Present draft benchmark suite and initial results.
  - Request feedback on benchmark configurations.
  - Request suggestions for additional/replacement configurations.

Participants

- Draft benchmark specification discussions
  - Phil Rasch (lead)
  - Philip Cameron-Smith
  - Brian Eaton
  - Cecile Hannay
  - Art Mirin
  - Will Sawyer
  - Pat Worley
  - (my apologies to anyone left out)
- Initial benchmarking
  - Pat Worley

Current Benchmark Suite

- cam3.0p1

  default configuration (C0): T42L26, T85L26, FV2x2.5 L26
- cam3.1p2

  C0: T42L26, T85L26, FV2x2.5 L26, FV1.9x2.5 L26
- cam3_5_27

  C0: T42L26, T85L26, FV1.9x2.5 L26

  C1: C0 with 30 levels and FV1.9x2.5 only

  C2: C1 with "cam3.5" aerosols

  C3: C2 with UW physics package

  C4: C3 with Morrison Gettelman cloud parameterization

  C5: C4 with predicted aerosol fields

  C6: C4 with full tropospheric chemistry

Benchmark Suite Comments

- cam3.0p1 => cam3.1p2:
  - new physics interface design (software engineering release)
- cam3.1p2 => cam3_5_27:
  - Neale-Richter convection mods, Vavrus "freeze-dry" mod, surface component coupling every timestep, …
- C0 => C1: 30 levels
- C1 => C2
  - Prescribed aerosol datasets generated by CAM with prognostic aerosol chemistry turned on; reimplementation of existing interpolation scheme

Benchmark Suite Comments

- C2 => C3
  - Replace Holtslag-Boville vertical diffusion with diag_TKE (Grenier-Bretherton); replace Hack shallow convection with UW (McCaa)
- C3 => C4
  - Replace Rasch-Kristjansson microphysics with Morrison Gettelman, including addition of two advected species
- C4 => C5
  - Enable prognostic aerosol code, adding 16 advected species
- C4 => C6
  - Enable full tropospheric chemistry, adding 108 advected species

Benchmark Details

- 30 simulation days (September) with monthly history output
- Platforms:
  - Cray XT3 at Oak Ridge National Laboratory (ORNL): dual-core 2.6 GHz Opteron processors
  - IBM BG/P at ORNL: quad-core 850MHz PowerPC processors (results for cam3_5_27 only)

  both platforms supporting MPI and OpenMP.
- Set optimization as high as possible while still preserving reproducibility. CAM has not been validated on either platform as of yet.
- C6 configuration not available with cam3_5_27. All other configurations were run, using both pure MPI and hybrid MPI/OpenMP.

3.0 (T42) vs. 3.1 (T85) vs. 3.5 (FV1.9)

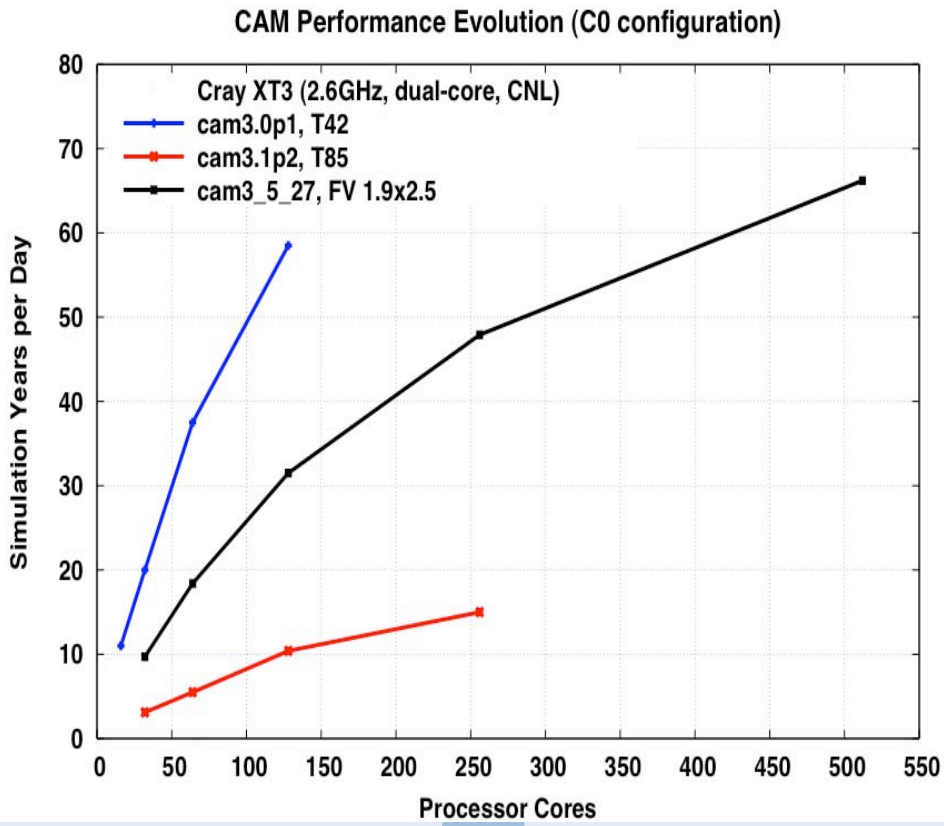

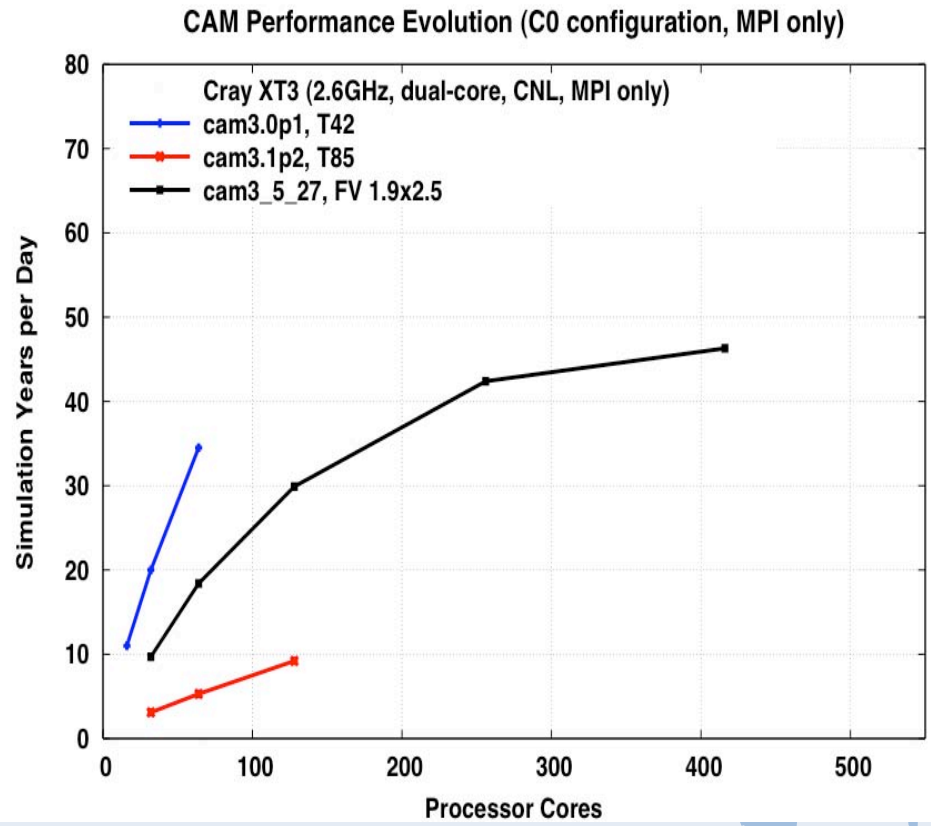

This analysis does not take into account computer performance evolution. Could also add plots for cam3.0 and cam3.1 on older platforms, if this would provide additional insight.

T85: 3.0 vs. 3.1 vs. 3.5

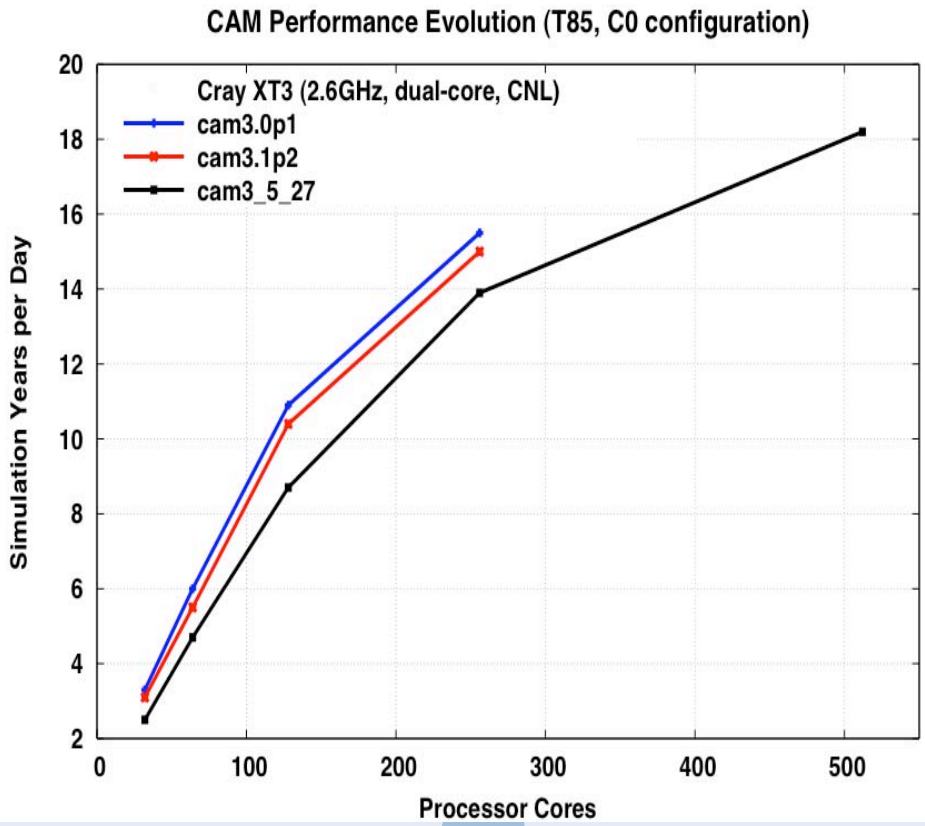

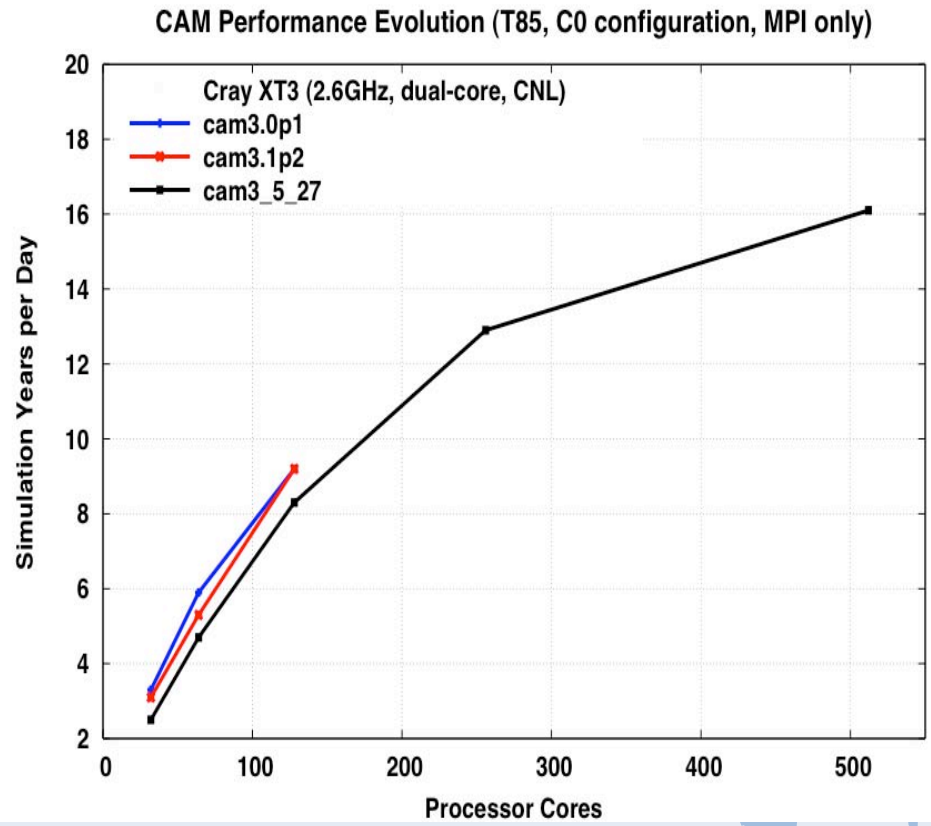

For small processor counts, cam3.5 is approx. 30% slower than cam3.0 . However, cam3.5 is also more scalable than the released version of cam3.0 or cam3.1 .

FV: 3.0 vs. 3.1 vs. 3.5

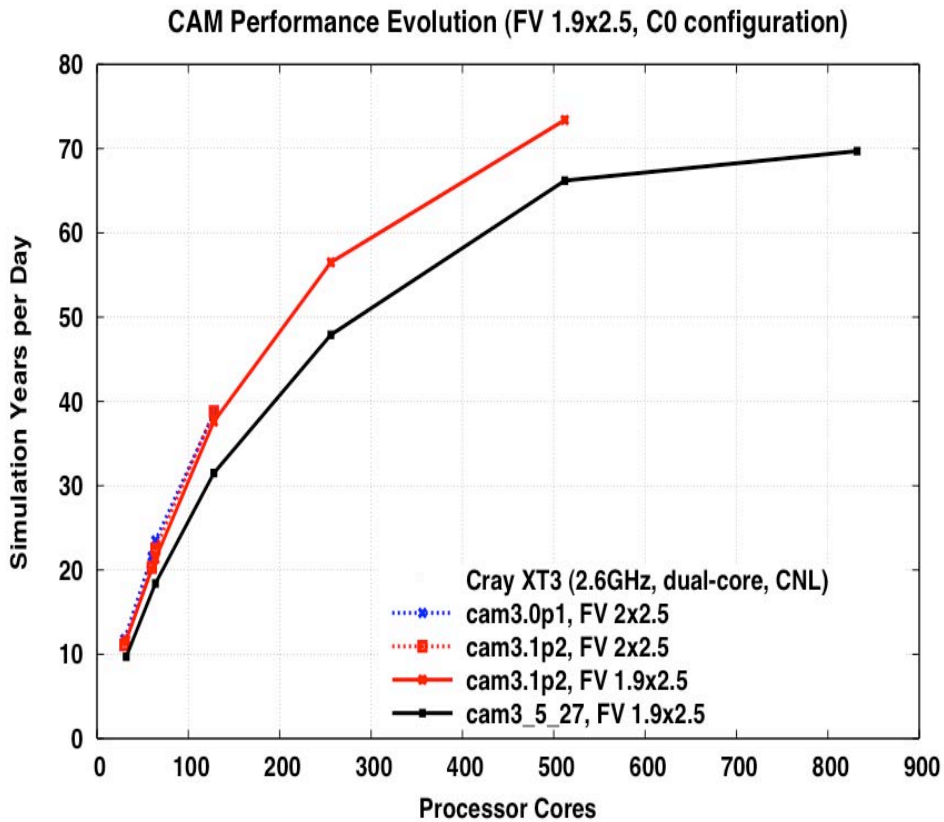

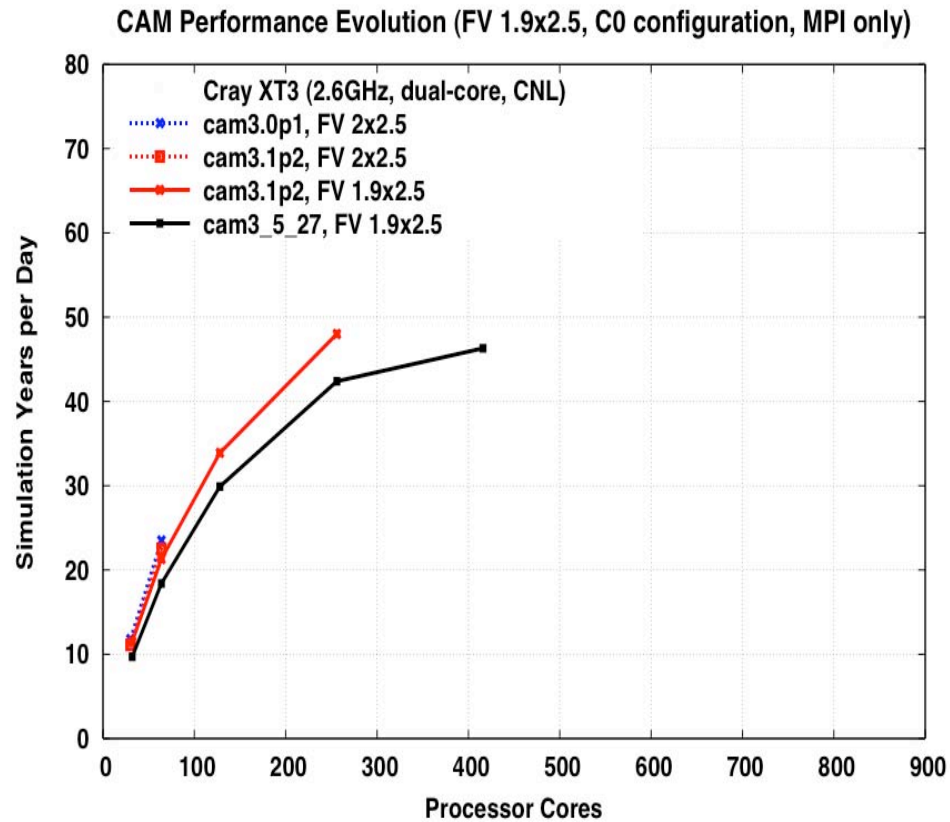

A number of data are missing. FV 2x2.5 scalability is similar to that of FV 1.9x2.5 . For small processor counts, cam3_5_27 is 20% slower than cam3.1 . Not all scalability options in cam3_5_27 are being exercised here.

XT3 FV: C0 vs. C1 vs. C2 vs. C3 vs. C4 vs. C5

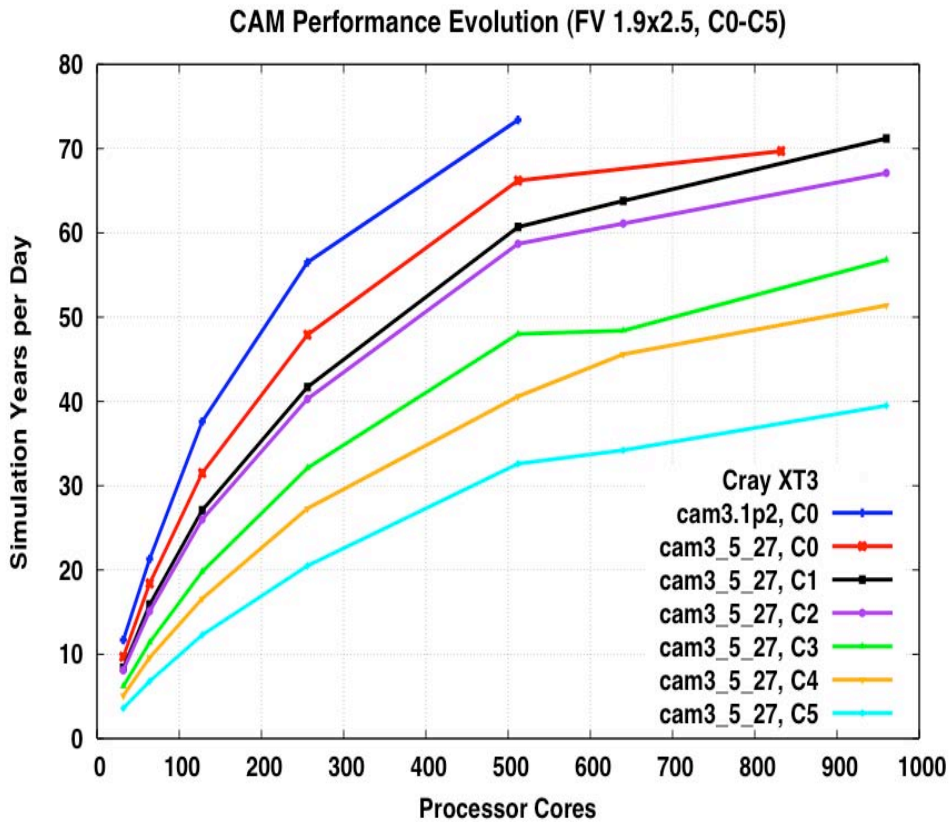

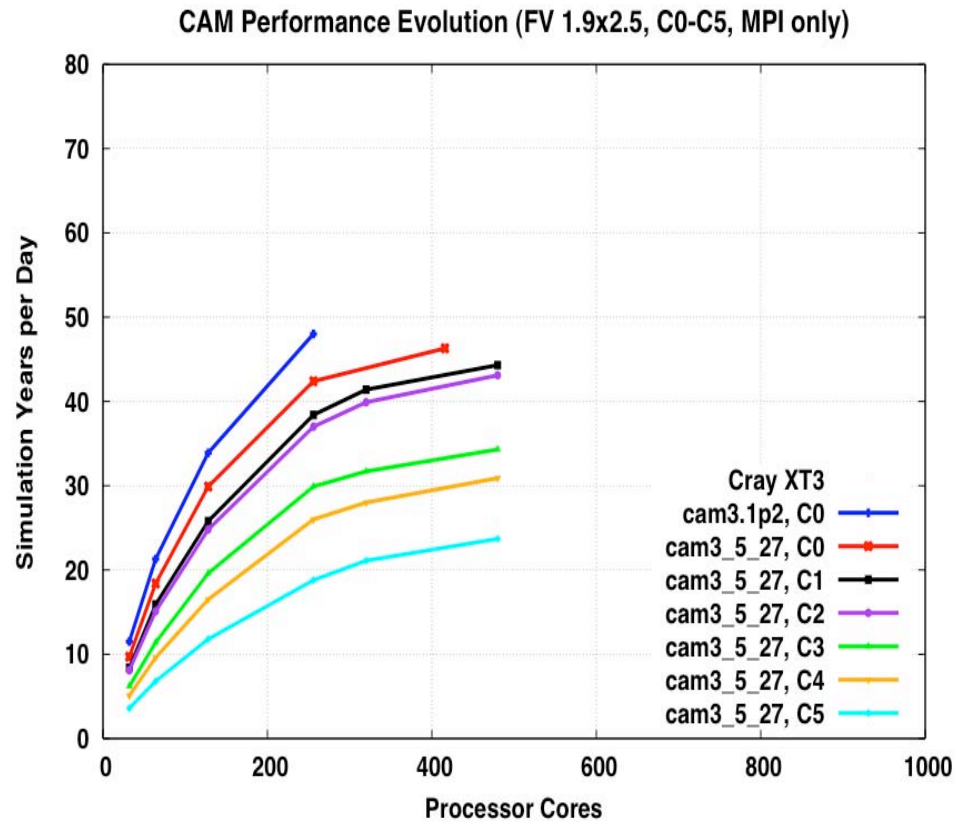

For 32 and 512 processor cores, the normalized cost progression is
.83 => 1.00 => 1.15 => 1.20 => 1.56 => 1.90 => 2.69 and
.90 => 1.00 => 1.09 => 1.13 => 1.38 => 1.63 => 2.03, respectively.

BGP FV: C0 vs. C1 vs. C2 vs. C3 vs. C4 vs. C5

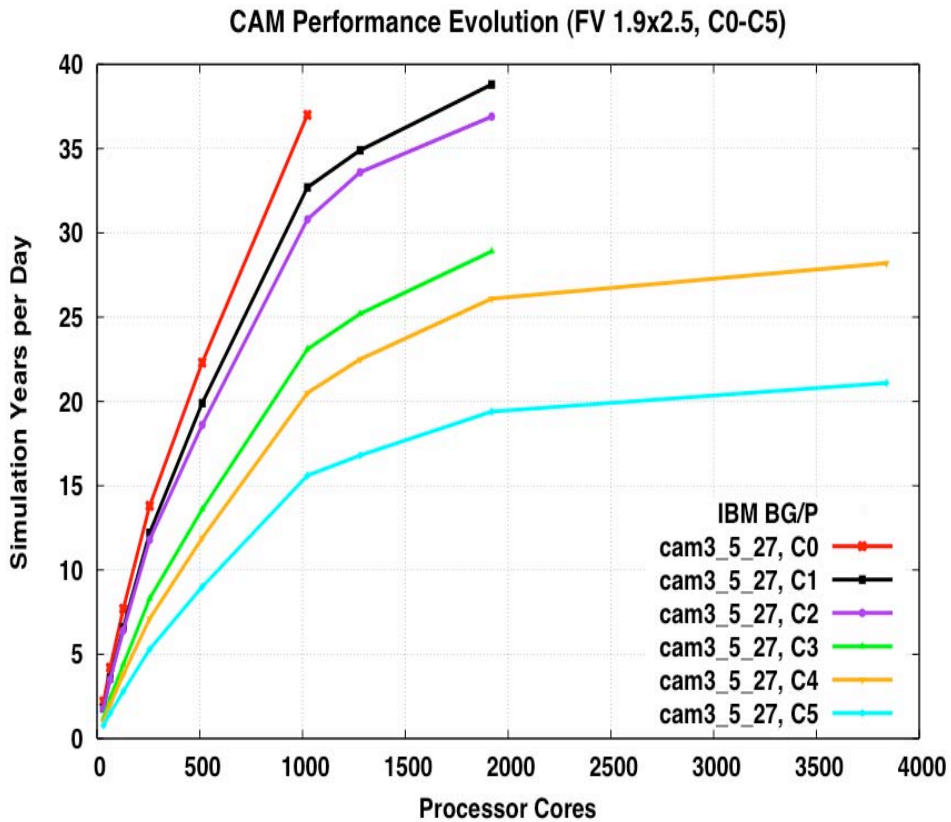

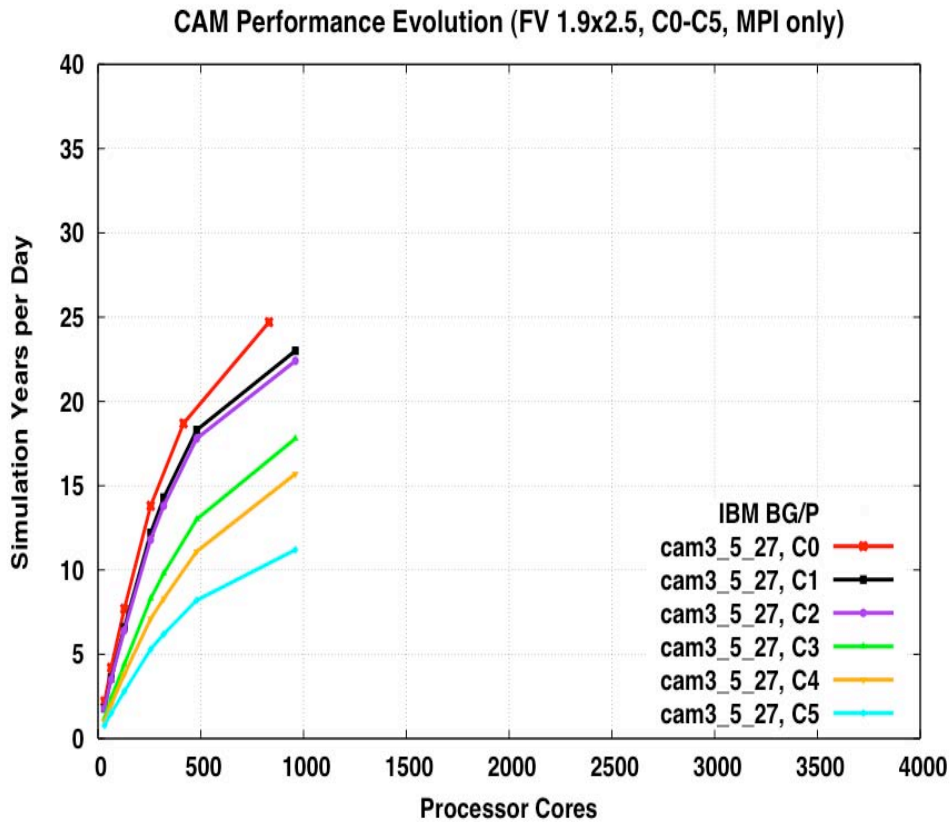

For 32 and 1024 processor cores, the normalized cost progression is
1.00 => 1.22 => 1.22 => 1.83 => 2.20 => 2.82 and
1.00 => 1.13 => 1.20 => 1.60 => 1.80 => 2.37, respectively.

Current Status

- Have lots of data, but will need to repeat experiments once the benchmarks are refined. Not worth doing in-depth analysis until this occurs. Also need to define run and reporting rules precisely, so that others can contribute to this effort.
- Cost increase from C0 => C5 is less than a factor of 3. C6 benchmark will be interesting, however, as it may be appreciably more expensive.
- Performance scaling is reasonable, but there are so many optimization options now (see Mirin talk) that it will take awhile to determine the optimal performance.
- CAM performance scalability is improved by running in hybrid MPI/OpenMP mode, and we hope that this option will be available in the full CCSM on these platforms.

Example Analysis

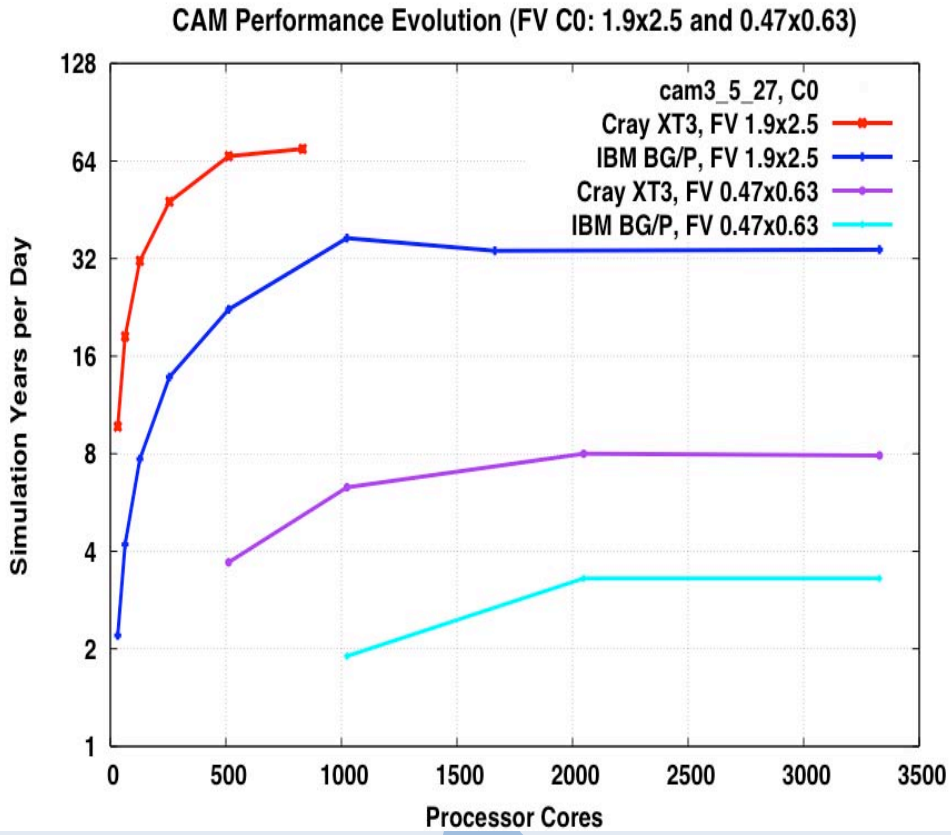

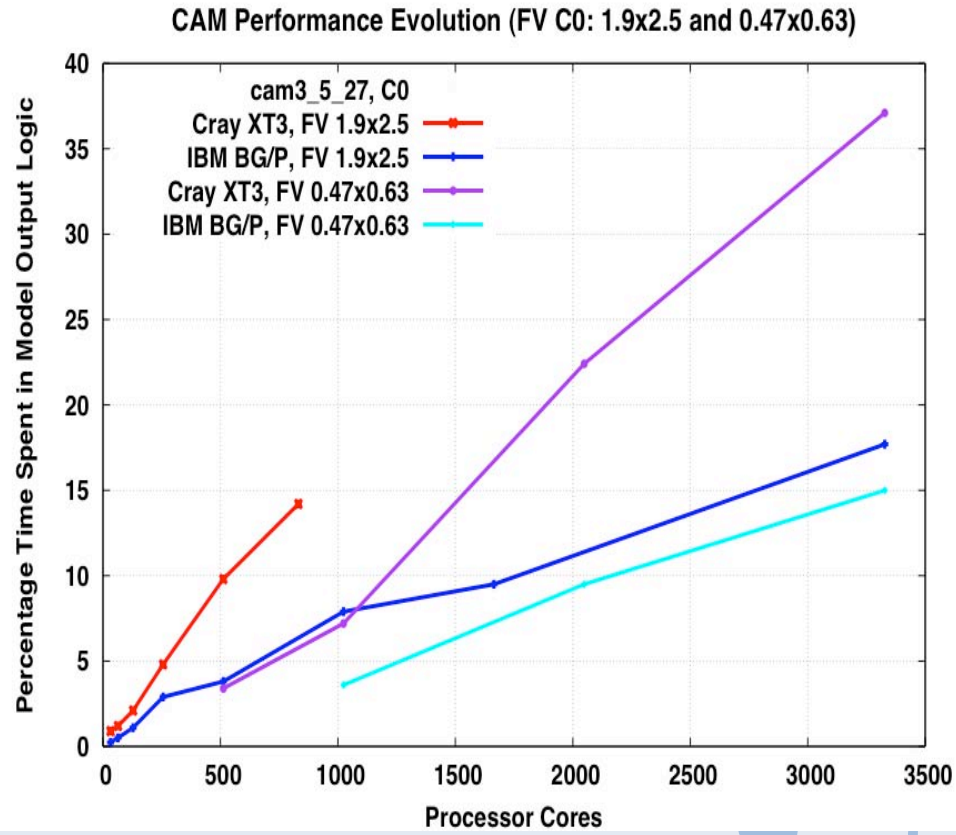

Comparing performance of 1.9x2.5 and 0.47x0.63 C0 benchmarks, and looking at time spent in subset of timers associated with history and restart output (so lower bound on I/O). A performance scaling problem is occurring in clm_driver_io for the 0.47x0.63 benchmark.

What's Next?

- "Freeze" version of model representing current CAM, something more recent than cam3_5_27, and collect data for configuration C6.
  - Should we wait for the new radiative transfer scheme RRTM?
  - Should we wait for the new aerosol package from Ghan and Liu?
- Fine tune existing benchmarks (suggestions?) Any interest in C1-C6 for spectral dycores?
- Add high resolution configuration(s):
  - FV 0.47x0.63? Others?
  - How many vertical levels? Should we add vertical resolution scaling studies?

What's Next?

- Anything to be learned from lower resolution benchmarks (e.g., FV 3x4)?
- Add a WACCM benchmark?
- Add an aquaplanet benchmark (something suitable for initial comparison with alternative dycores)?
- Collect data on NCAR systems. On other systems?
- Modify timers to better identify performance limiters (such as I/O).
- Note: benchmarking can be expensive. It is important to prioritize benchmarks, and to clearly state WHY we are collecting the data.

Performance Benchmarking Costs

- "Fair" benchmarking requires some level of optimization for each configuration on each system.
- Ensembles (collecting data at different dates and times) and/or real-time monitoring (babysitting) are required to eliminate misleading performance perturbations.
- Analysis of detailed timing data is needed to understand results and to identify performance perturbations.
- Scaling studies are important, at least, for baseline studies and after major upgrades (system and model), and are expensive in both computer and people time.

Questions? Comments?