ACME
Accelerated Climate Modeling for Energy

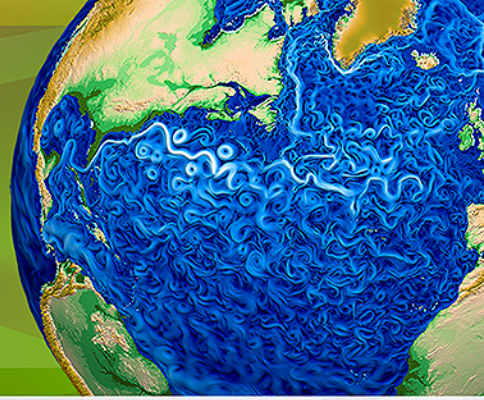

ACME Group Overview and Roadmaps: Performance

Patrick Worley – Performance Co-Lead

May 5, 2015
ACME All-Hands Meeting
Tysons Corner, VA

U.S. DEPARTMENT OF ENERGY

Performance “Swim Lanes”

1. Preparing ACME (project and code) for next generation architectures
2. Evaluating and optimizing ACME v1.0 for target systems
3. Advising on and optimizing performance for ACME development runs

1. “to boldly go where **no code has gone before**”

2. Performance “Sherpas”, hauling the code to the Summit

3. Performance mechanics, tweaking to get the last bit of performance

Performance Activities

1. Performance data capture and analysis
2. Inter-node performance
   - scalable algorithms, load balancing, interconnect, I/O
3. On-node performance
   - multi/many-core CPU, GPU and other, e.g. SIMD/vector, accelerators, hierarchical memory, …
4. NERSC/LCF application readiness and early science programs

for each component and for the full coupled model

Performance Interactions

- Coupled Model
  - performance evaluation and prediction for science simulations
- Science Teams
  - evaluation and optimization of new features; advice and monitoring for development runs
- SE
  - infrastructure enhancements, code enhancements, platform and compiler support, debugging
- Workflow
  - provenance and performance tracking
- Computing Centers (NERSC, ALCF, OLCF)
  - readiness and early science programs, INCITE/ALCC allocations

Progress: Highlights

- Enhanced performance data and (some) provenance being collected for all runs on Hopper, Edison, Mira, and Titan
- Performance of water cycle science case benchmark measured and analyzed on Edison, Mira, and Titan
  - CAM_SE (ne120, 64 levels, 50 tracers), CLM4, POP2 (tx0.1), CICE)
  - Significant progress has been made on including CLM4.5 and COSP (for ATM), but not yet working on all systems.
  - MPAS-Ocean and MPAS-LI are just now available for initial examination.

Performance of ACME water cycle benchmark

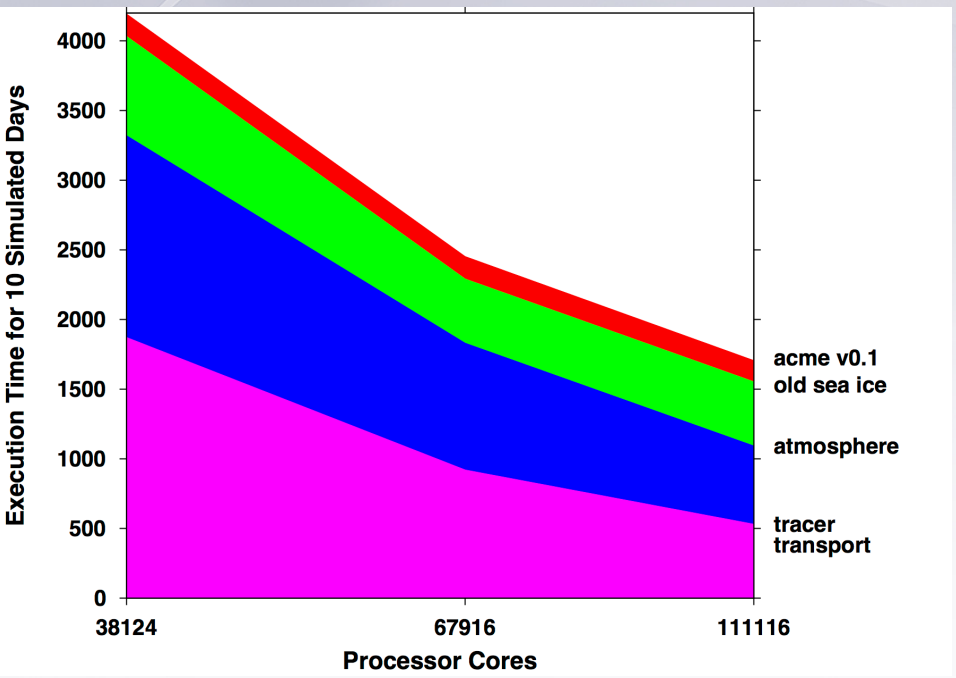

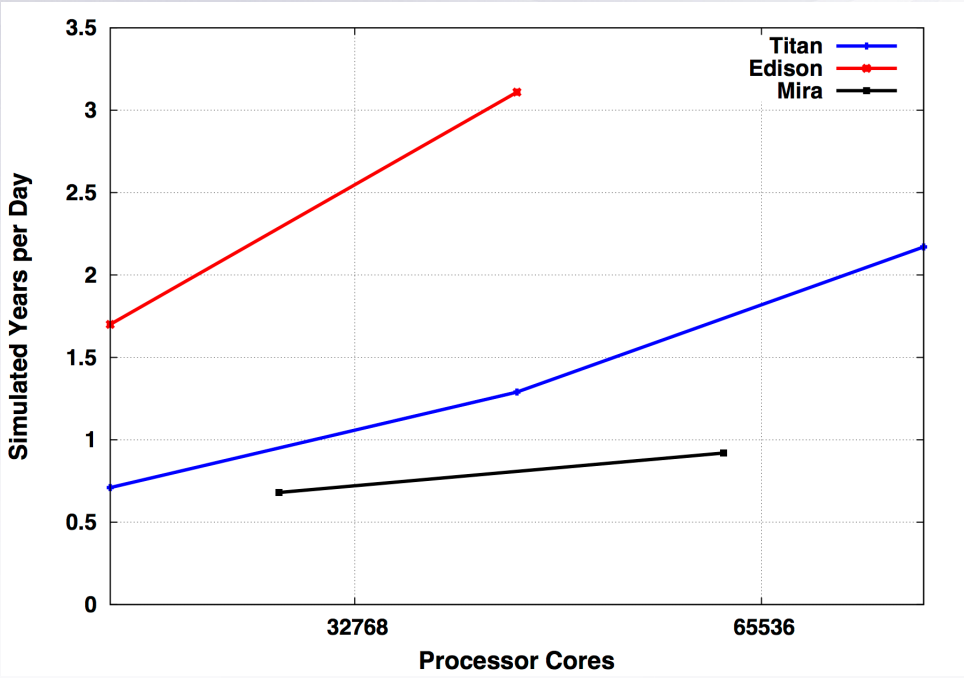

- Performance of ACME water cycle benchmark on Titan.
- Performance of ATM in ACME water cycle benchmark

Progress: Highlights

- NERSC Exascale Science Application Program (NESAP) award
  - provides resources and early access to prepare for NERSC-8 / Cori Intel MIC system being deployed in 2016.
- OLCF Center for Accelerated Application Readiness (CAAR) program award
  - provides resources and access to early hardware to prepare for the Summit system that will replace Titan in 2018.
- Tracer transport kernels developed for evaluation, optimization and comparison between OpenMP, OpenACC, and CUDA Fortran
  - full mini-app, with MPI, under development

Progress: Highlights

- GPU implementation of CAM-SE tracer transport was ported into ACME and performance further enhanced
  - support for limiter8, required for production runs, is in a feature branch and will merged to master shortly
- Compiler flags were identified that halved compile time of the model without diminishing performance on Mira and on Titan.
- Support was added for Cray and Intel compilers on Titan
  - without GPU, both produce faster executables than PGI
  - Cray compiler has better OpenACC and Co-Array Fortran support

Progress: Highlights

- PIO settings were identified that improved I/O performance significantly on Mira (6x for water cycle benchmark) and Titan. Settings were exposed to allow runtime control (via env_run.xml).
- CAM settings were identified that decreased the overhead of load balancing by up to 18X on Titan, and were modestly better on Mira.
- Much other work is “in flight”, and good progress is being made. The performance group has also made significant contributions in platform support, debugging, benchmarking for and advising on development runs, etc.

Performance Team

- ANL
  - Azamat Mametjanov
  - Jayesh Krishna
- LANL
  - *Douglas Jacobsen*
  - Phillip Jones (L)
- LBNL
  - Hans Johansen
  - Noel Keen
- ORNL
  - Benjamin Mayer
  - Matthew Norman
  - Patrick Worley (L)
- SNL
  - Irina Demeshko
  - Mark Taylor

Performance: Roadmaps

- Continue tracking model development
  - add CLM4.5 with river basin grid, MPAS-Ocean, MPAS-LI, MPAS-ICE, CAM with elevation classes, ... to water cycle science case benchmark
  - add cryosphere science case benchmark
  - add biogochemistry science case benchmark
- Integrate provenance and performance data capture into workflow
  - develop dashboard and associated analysis tools for tracking performance data project-wide

Performance: Roadmaps

- Continue improving parallel implementations or developing new parallel algorithms, e.g.
  - overlapping computation, data movement, and/or communication
  - one-sided (MPI-3, Co-Array Fortran), two-sided, and collective communication protocol optimizations
  - new decompositions (for both improved load-balancing and minimizing communication overhead)
  - additional functional parallelism (e.g., between ICE and ATM, within a component, …)
  - level-dependent subcycling
  - PIO2

  for each component and for the coupled model

Performance: Roadmaps

- Continue improving implementations or developing new on-node algorithms (serial and parallel), e.g.
  - memory layout and access pattern
  - SIMD/vectorization
  - OpenMP (nested, scheduling, memory/core affinities, …)
  - overlapping computation and data movement
  - OpenACC/OpenMP4 as a replacement for CUDA for running on GPU accelerators
  - complexity minimization for new code or new configurations or at new scale

  for each component

Performance: Roadmaps

- Continue preparing for next generation architectures:
  - continue development of existing and identify new kernels and mini-apps
  - evaluate and implement strategies for avoiding/maintaining platform-specific code (OpenACC, (nested) OpenMP, CUDA, AVX, MKL, …)
  - investigate porting, testing, evaluation, and optimization on early hardware

  and other NESAP and CAAR program commitments. Also,
  - prepare proposal for ALCF Early Science Program (Theta and/or Aurora, to be decided)

Performance: Roadmaps

- Continue to be responsive to performance needs of science team developments and of coupled model simulations.