High-resolution coupled atmosphere/ocean simulations on today's petascale systems

Mark Taylor
Sandia National Labs
mataylo@sandia.gov
And the members of the
CESM Atmosphere Model Working Group
CESM Software Engineering Working Group

SOS18: 18th SNL / ORNL / ETH Workshop on Computer Science
St. Moritz, March 17, 2014

U.S. Department of Energy
Office of Science

SciDAC
Scientific Discovery through Advanced Computing

Sandia National Laboratories

Outline

- Motivation for Earth system modeling on exascale machines
- Performance of the Community Earth System Model (CESM) on 10PF systems
- The spectral element dynamical core
  - Much of the CESM's scalability is due to the recent adoption of an spectral element dynamical core for the atmosphere component

Exascale Drivers for ESMs

Sandia National Laboratories

- Increased resolution
- Increased resolution (with other improvements) leads to better skill in reproducing observations. For both NWP and climate modeling
- At cloud resolving scales (~1km) we can remove convective parameterizations – one of the biggest sources of uncertainty in climate models
- Regional climate and extreme weather
  - Impact of climate change on tropical cyclones, statistics of precipitation extremes
  - Regional detail of climate change – important for accessing human impacts.

CESM on petascale systems

- Could we run a cloud resolving (1km) Earth system model on an exascale system?
  - 10km resolution on a 10PF system: 10^2 less grid points (ignoring the vertical)
  - 10PF system has 10^2 less computing power.
  - 10km "weather resolving" model should run on today's 10PF systems 10x faster (timestep is 10x larger)
- Examine weather resolving CESM on Mira (10PF BG/Q) and Titan (27PF XK7)
  - 1/8 degree (13.5km resolution) Atmosphere/ Land
  - 1/10 degree (3km – 11km resolution) Ocean / Sea Ice

The Community Earth System Model (CESM)

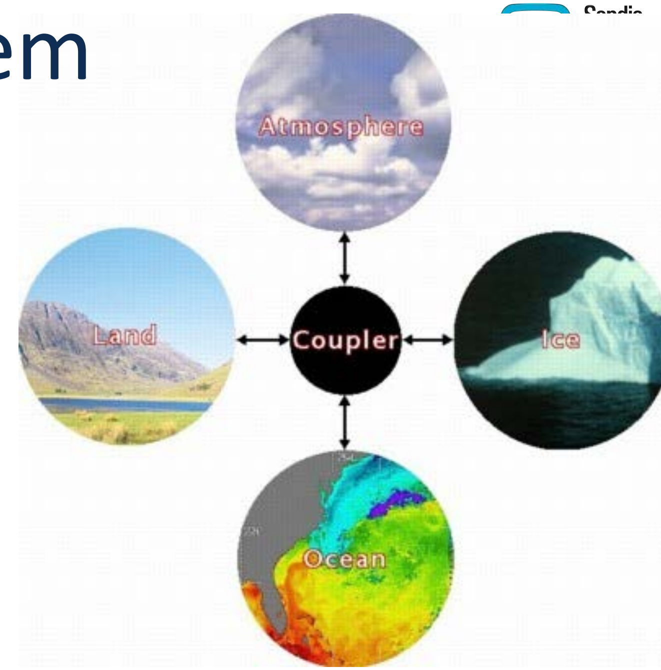

- IPCC-class model developed by NCAR, U.S. National Labs and Universities
- Atmosphere, Land, Ocean and Sea ice component models
- CAM is the atmosphere component model
- Science & policy applications:
  - Seasonal and interannual variability in the climate
  - Explore the history of Earth's climate
  - Estimate future of environment for policy formulation
  - Contribute to assessments

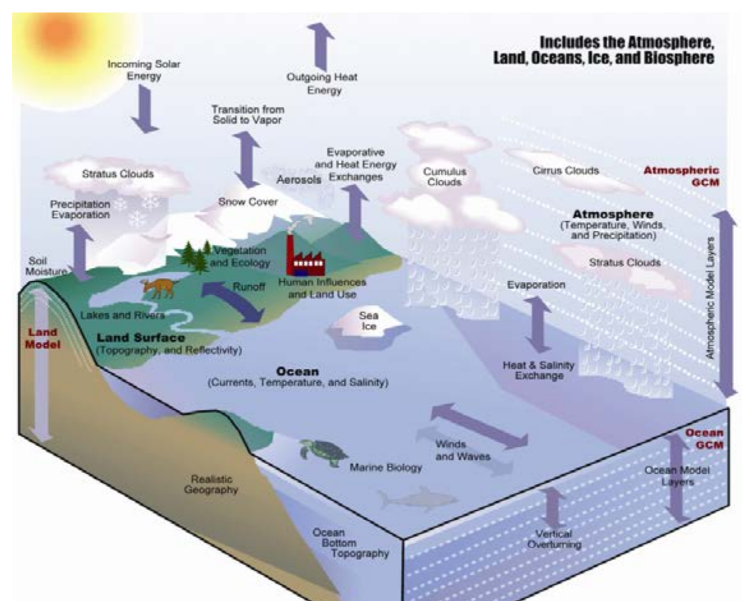

Horizontal Grid Resolution

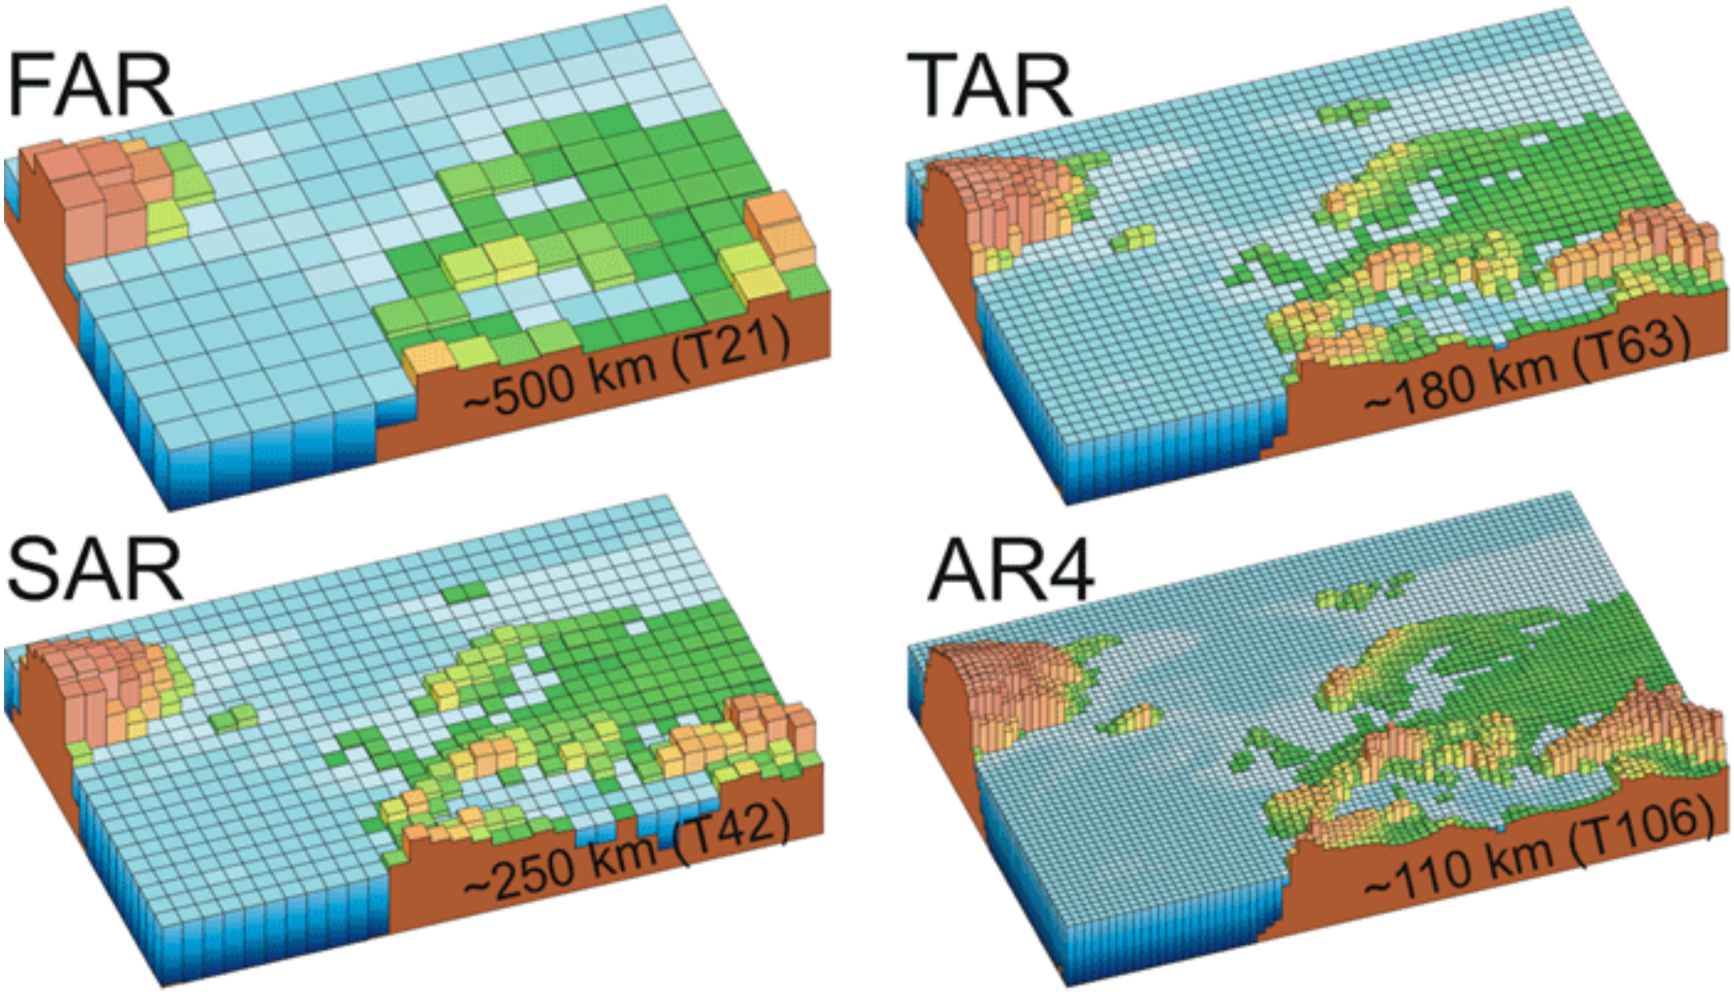

Source: IPCC 4th Assessment Report

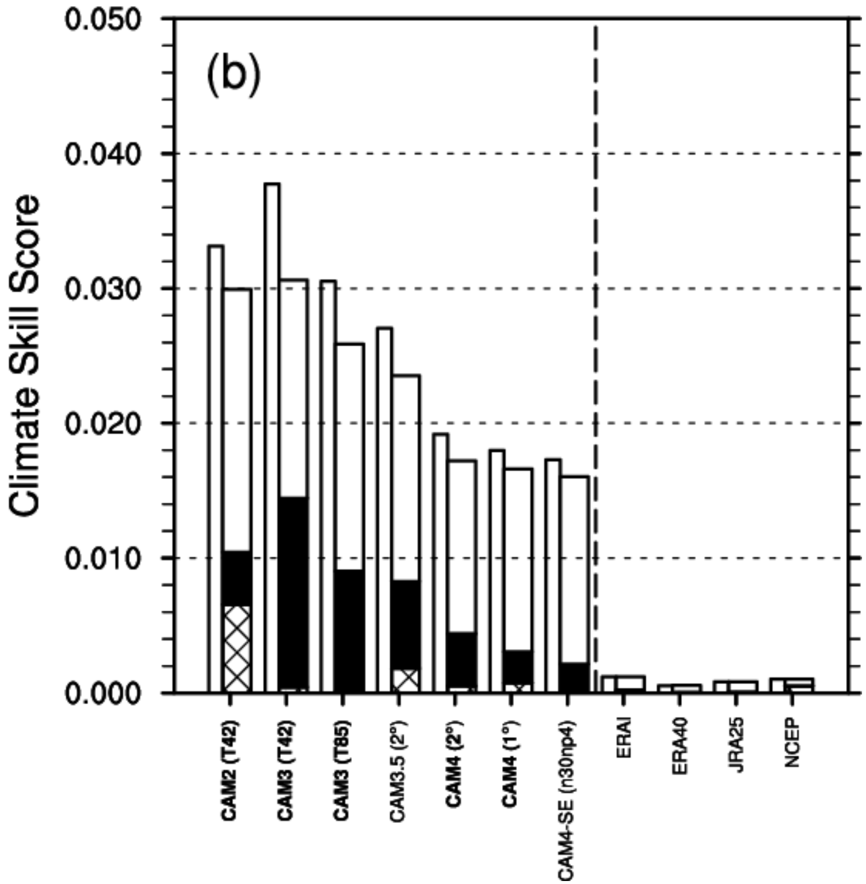

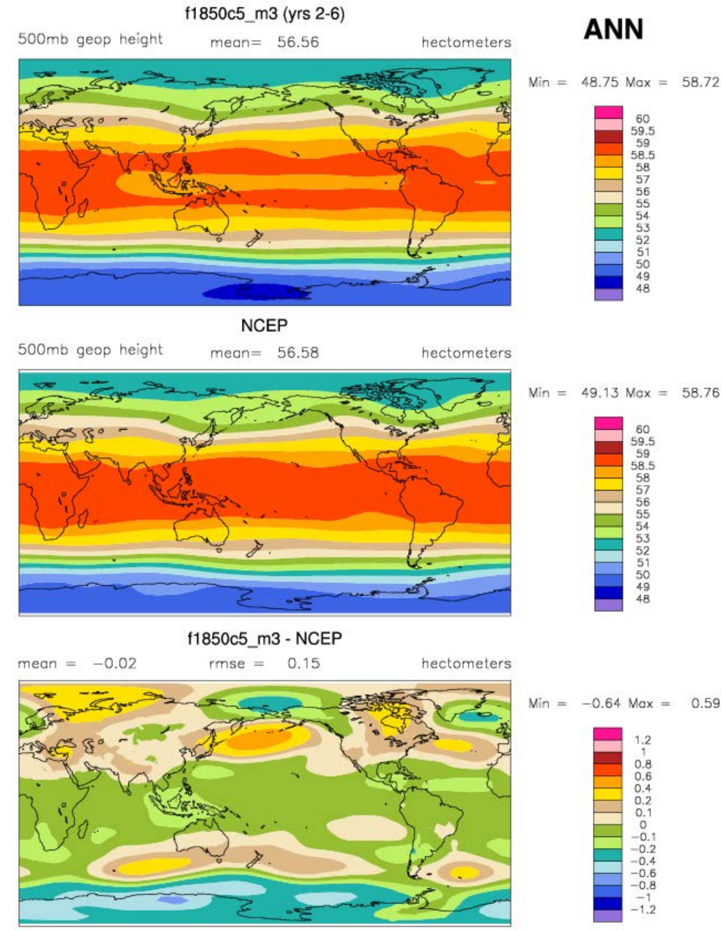

- CAM shows steady improvement as resolution is increased and physics improved
- 500mb geopotential height skill score (30-90N) DJF
- Mean square error from uncond. bias, cond. bias and phase error
- Source: Rich Neale (NCAR)

Taylor diagram

- CCSM3.5
- CESM1.1 (FV)
- CESM1.1 (SE)

	RMSE	Bias
CCSM3.5 (2008)	1.000	1.00
CCSM4 (2010)	0.883	0.78
CESM1.1 (FV) (2012)	0.791	1.58
CESM1.1 (SE) (2012)	0.735	1.34

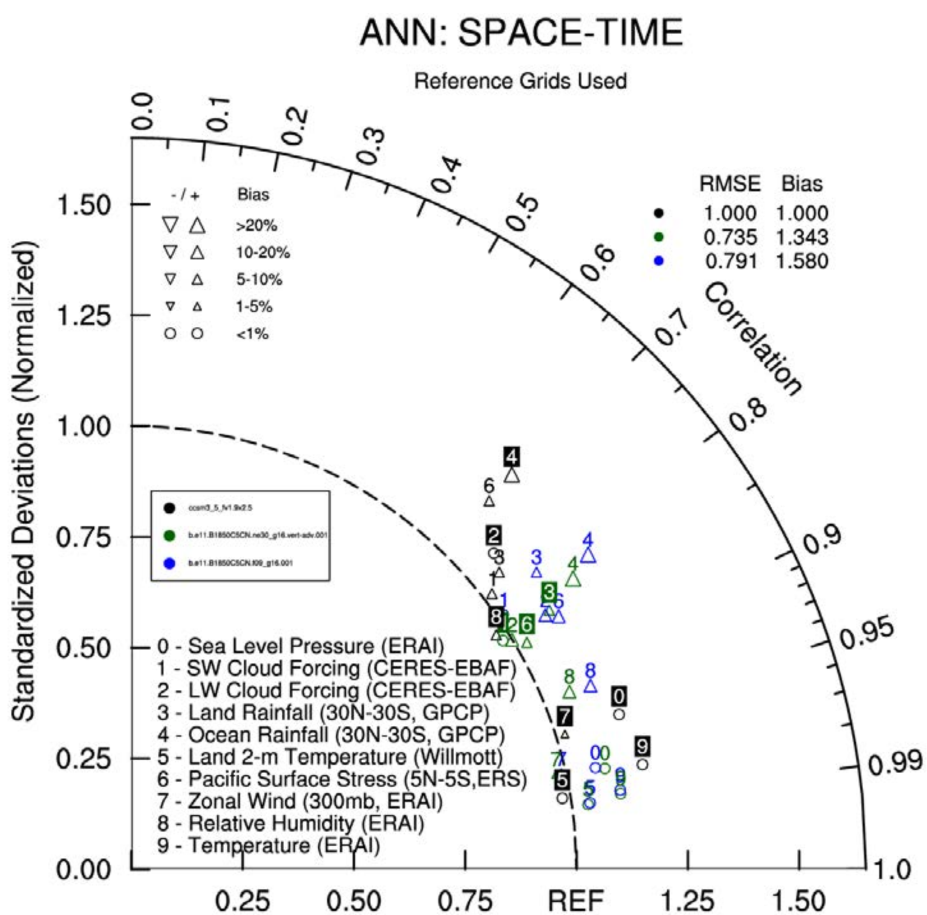

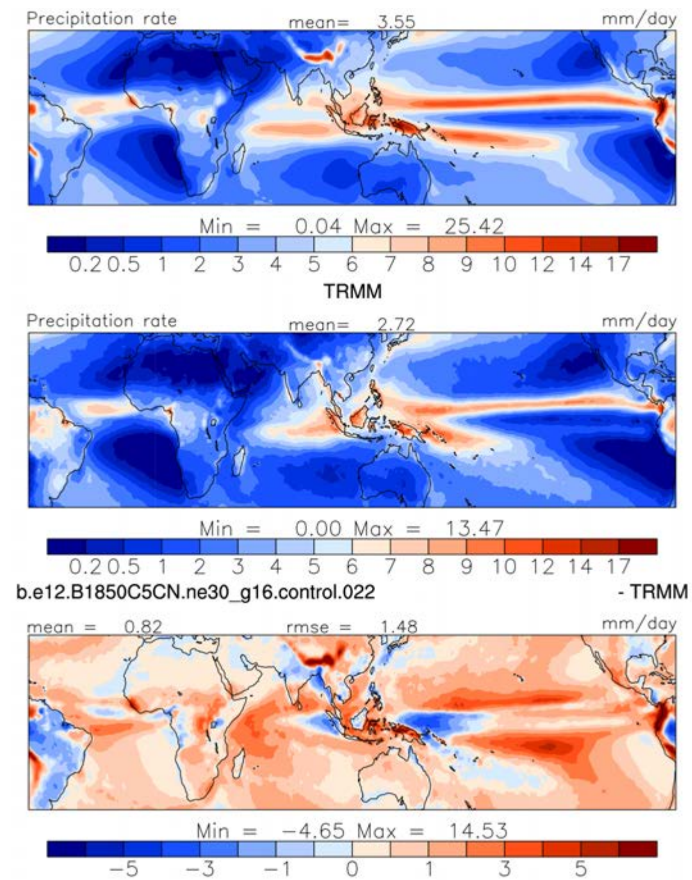

Source: Cécile Hannay, AMWG Meeting, 2012

Kinetic Energy Spectra

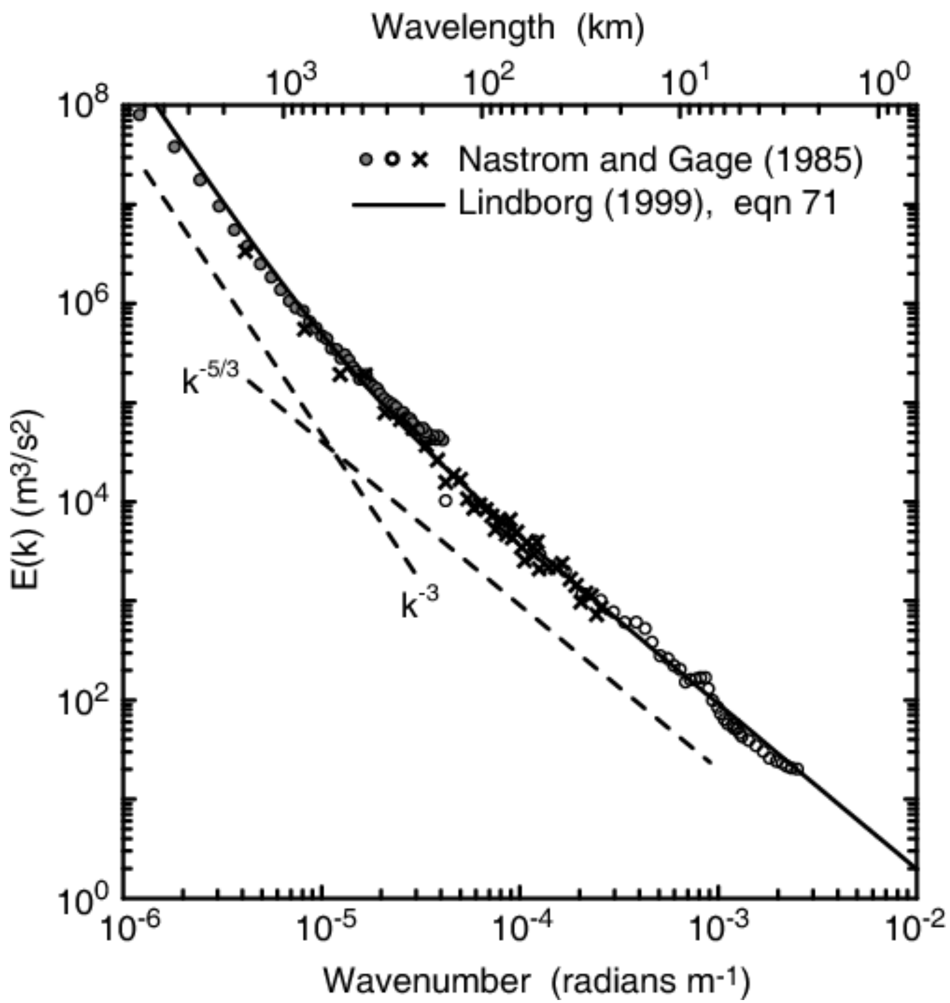

- Nastrom-Gage transition in KE spectra
- Mesoscale shallowing:
- Transition from a -3 regime (representative of quasi-2d large scale flow) to a -5/3 regime (associated with increased variability, increased frequency of extreme events)
- Resolving the -5/3 regime considered necessary if not sufficient to simulate correct mesoscale variability

KE spectra from aircraft observations (symbols, Nastrom and Gage 1985) and functional fit (solid line, Lindborg, 1999). Figure from Skamarock, 2004.

Kinetic Energy Spectra

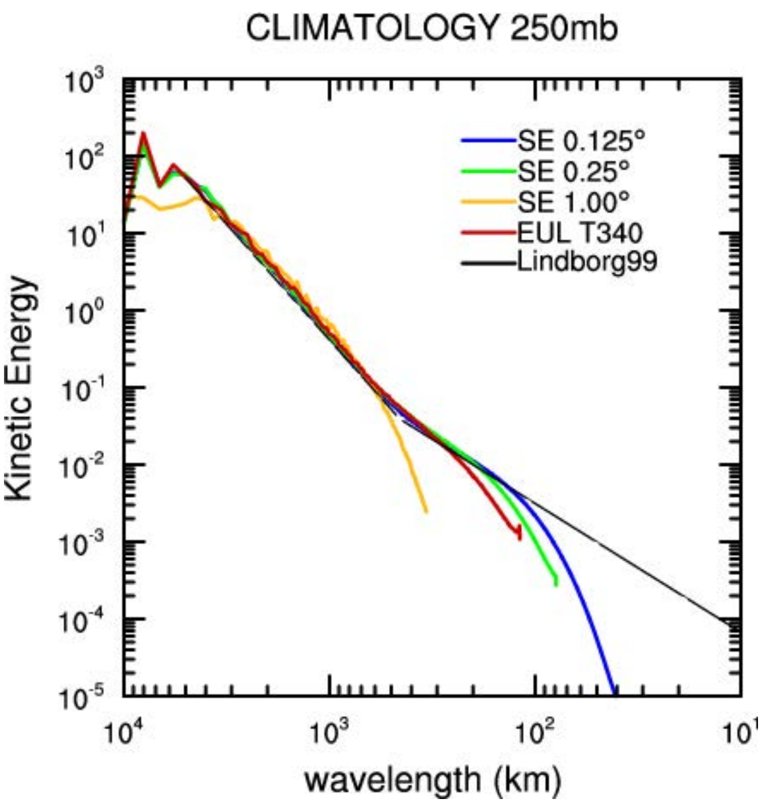

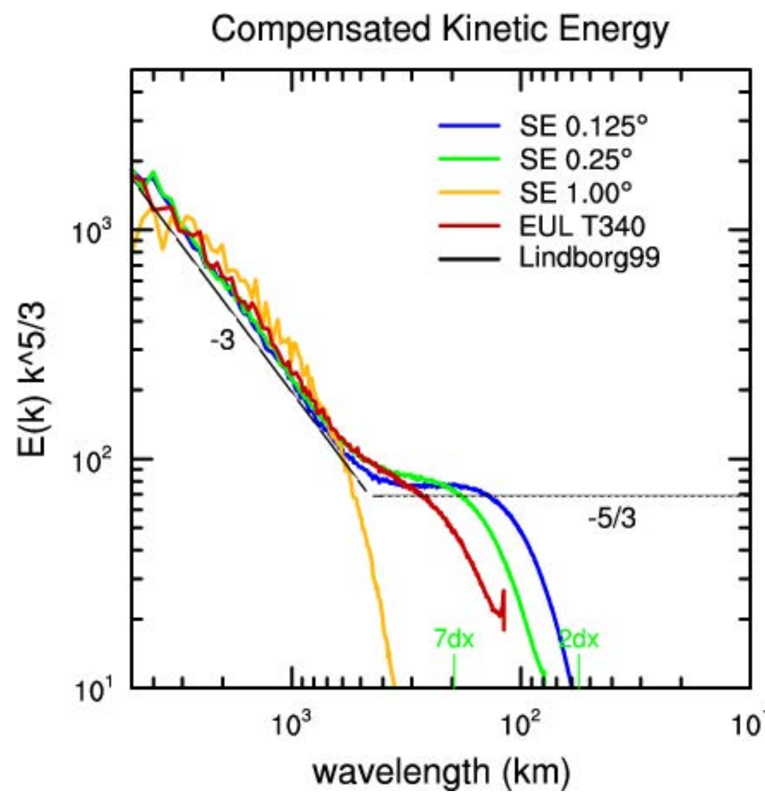

- Observations: KE transitions from a -3 scaling to -5/3, which continues to ~2km and probably even further.
- Transition occurs at ~400km, but requires 13.5km resolution to capture in CAM4

Resolving Tropical Cyclones at 27km Resolution

Sandia National Laboratories

Prescribed SST

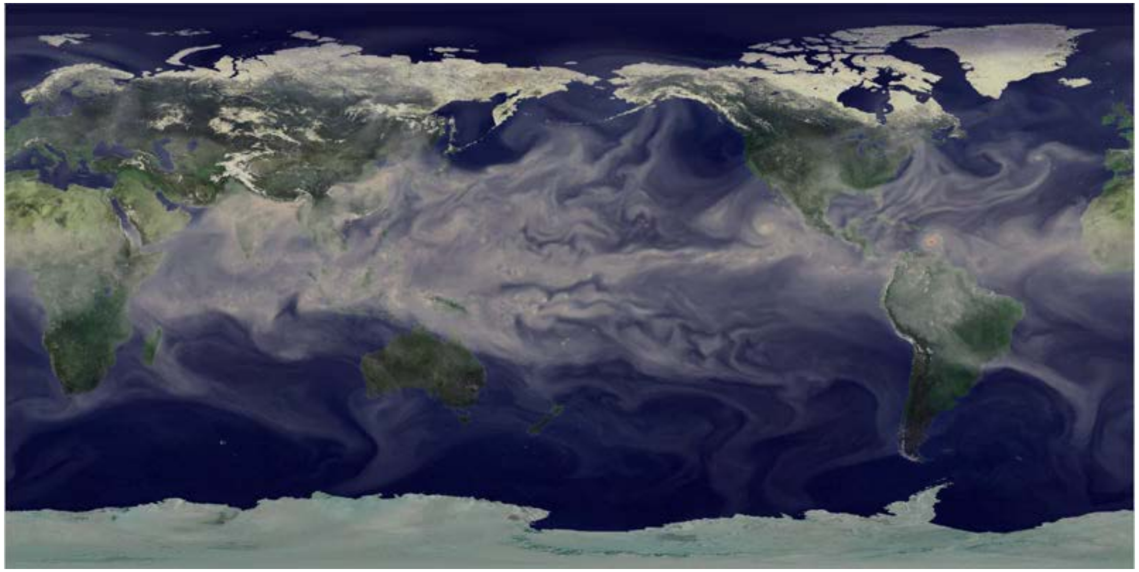

Fully Coupled

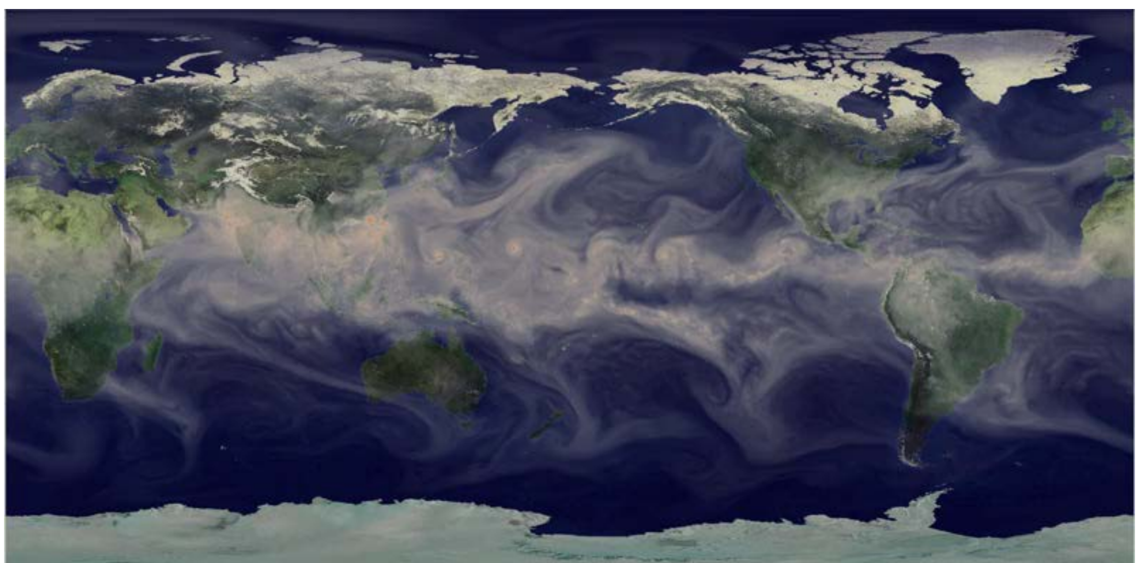

High resolution does a much better job at capturing storms like tropical cyclones

20 day loop, September conditions, atmosphere water content

CESM at 27km does a reasonable job at capturing tropical cyclones, including category 5 storms

At 27km resolution there is sufficient storm activity that the model can be *tuned* (with prescribed SST) to match observations (~80 TC per year, or ~50 category 1-5 per year)

Our first coupled simulations (27km atm/10km ocean) have SST biases resulting in very little Atlantic storms and too many Pacific storms.

CESM at 13.5km Resolution

Sandia National Laboratories

Prescribed SST

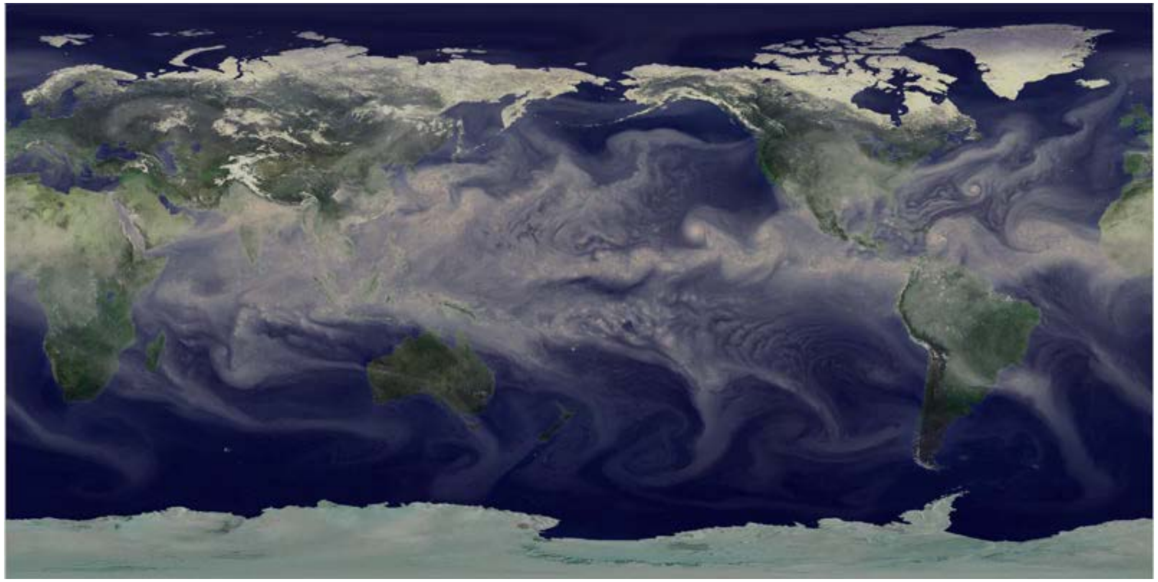

Fully Coupled

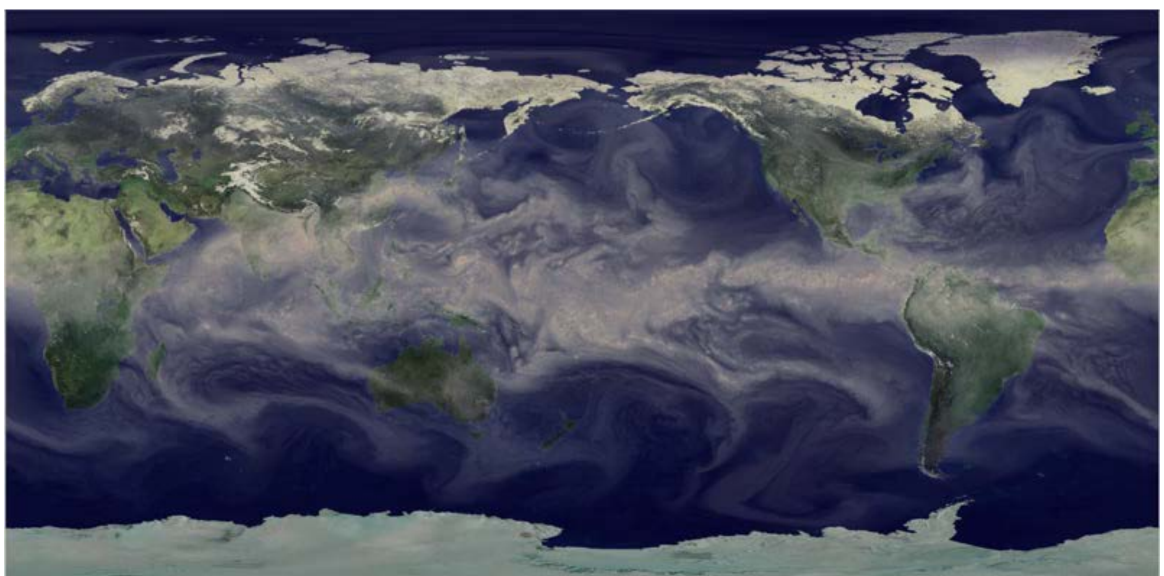

Exploratory results from CESM simulations at 13.5km.

With prescribed SSTs untuned model is too energetic. (top panel, September conditions)

Short coupled simulations (2 years) starts with same SST bias as 27km atmosphere

Bottom panel shows 20 day loop from June conditions

27km

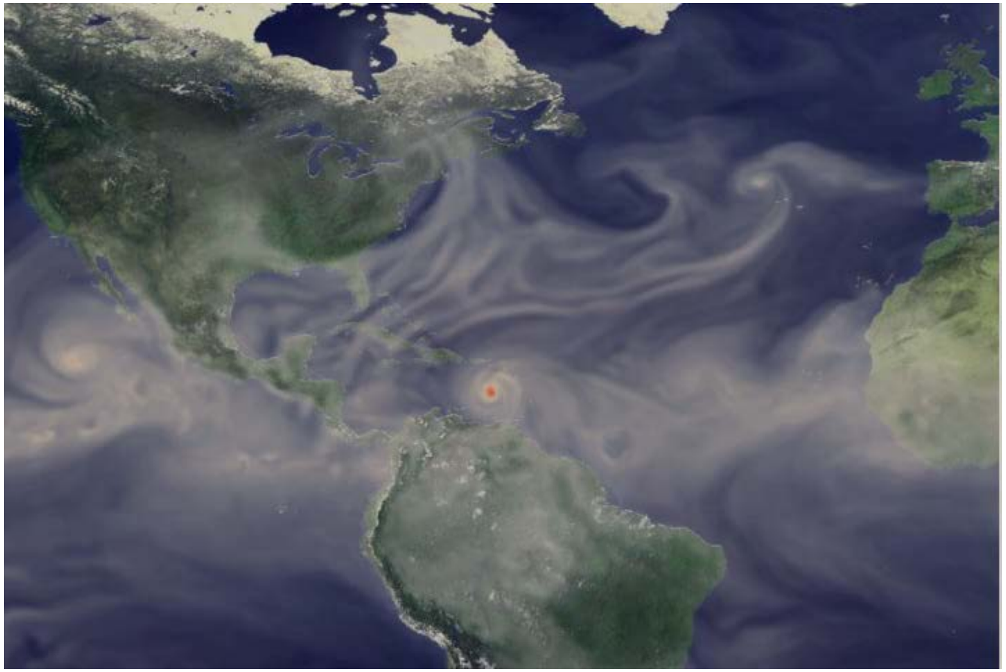

13.5km

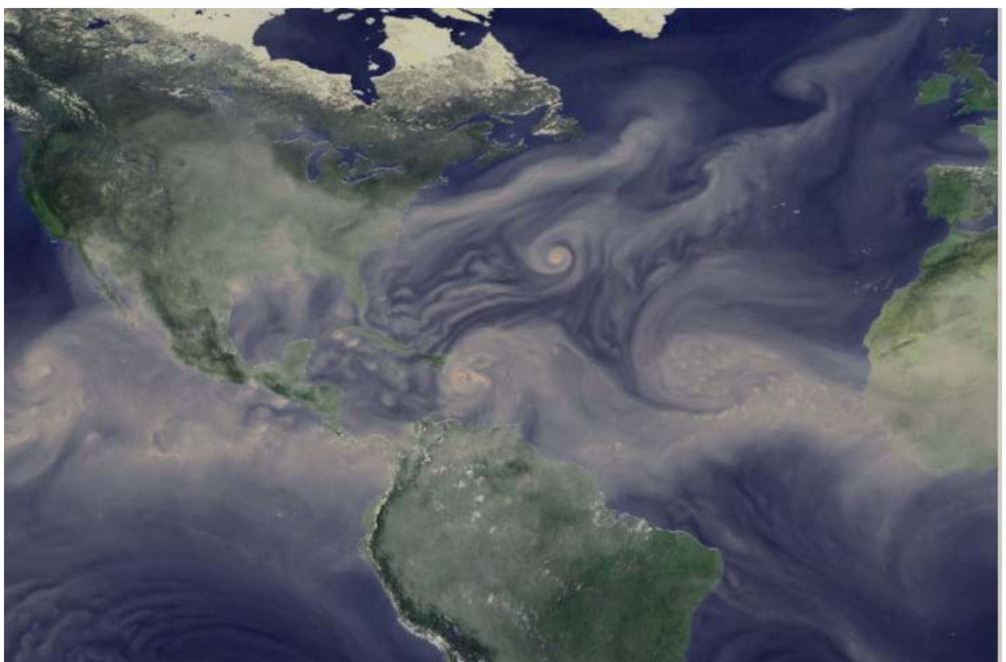

CESM Category 5 Gulf of Mexico hurricane example comparing 27km and 13.5km resolutions

Performance of CESM on petascale systems:

Mira (IBM BG/Q at Argonne) 10PF

Titan (Cray XK7 at Oak Ridge) 27PF

Weather resolving ESM at 10PF: should be about 10x faster than a cloud resolving ESM at 1EF

Weather Resolving CESM on Mira

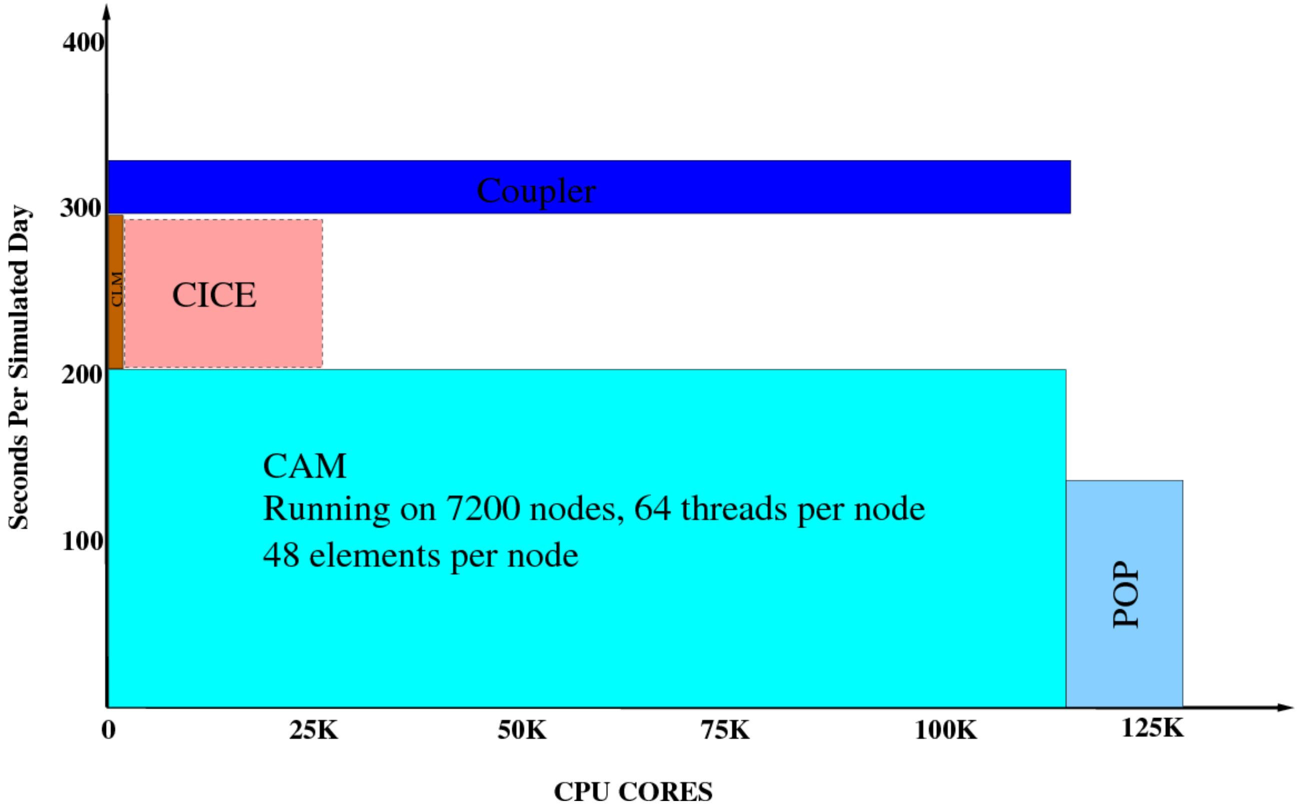

CESM:
13km atm/land
10km ocean/ice
8K BG/Q nodes.
0.7 SYPD

- Processor layout and performance (drawn to scale).
- Atm,ocean and ice components running near the limit of there current scalability (20% of Mira)
- Performance rate not practical for climate (goal: 5 SYPD)
- Atmosphere dominates the cost of the system. But ocean and ice components also not running fast enough for climate

Atmosphere dynamical cores in the CESM

- CAM-SE: CAM with the Spectral Element (SE) dynamical core from HOMME
- CESM1.2 Release: June 2013
  - Switch to CAM-SE on a cubed-sphere grid for high-res simulations
  - Motivation: parallel scalability for performance at high-resolution
- Previous CAM Dynamical Cores:
  - CAM-FV Lin-Rood FV lat/lon
  - CAM-EUL Global Spectral

Atmosphere dynamical core scalability

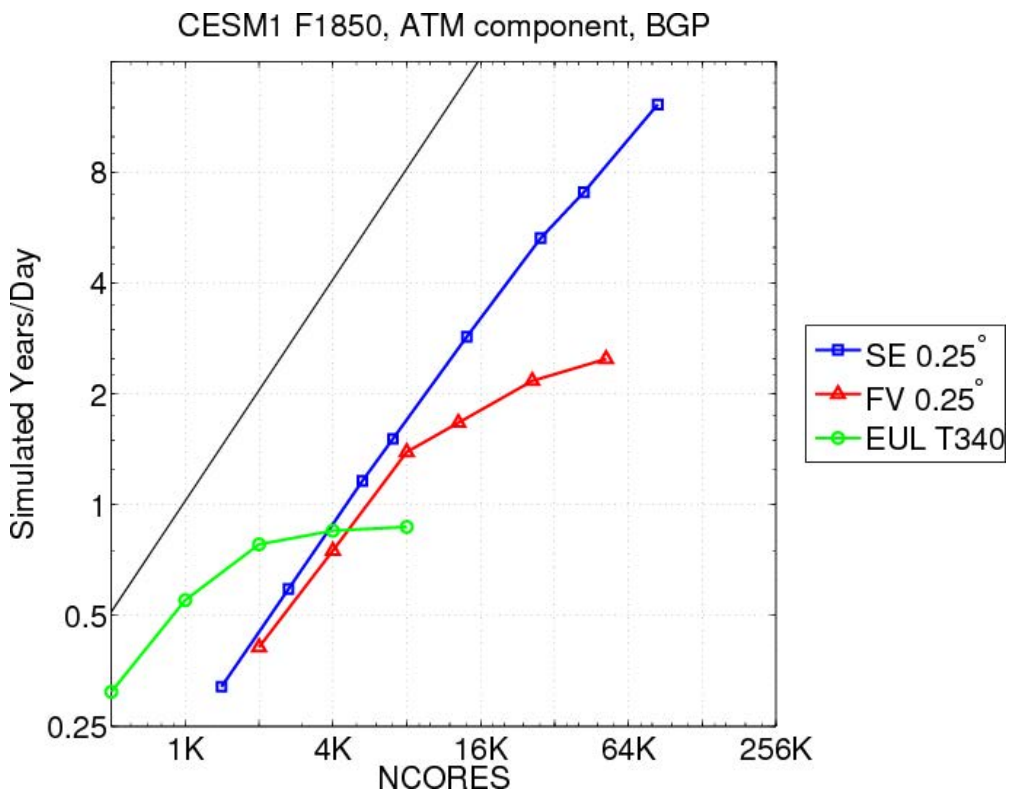

- CAM (CESM Atmosphere component) running on IBM BG/P Intrepid
- CAM4 physics at 0.25° (27km) resolution
- Compare CAM with SE, FV and EUL (global spectral) dycores
- CAM-SE achieves near perfect scalability to 1 element per core (86,000 cores).

CAM5 27km & 13.5km on BG/Q

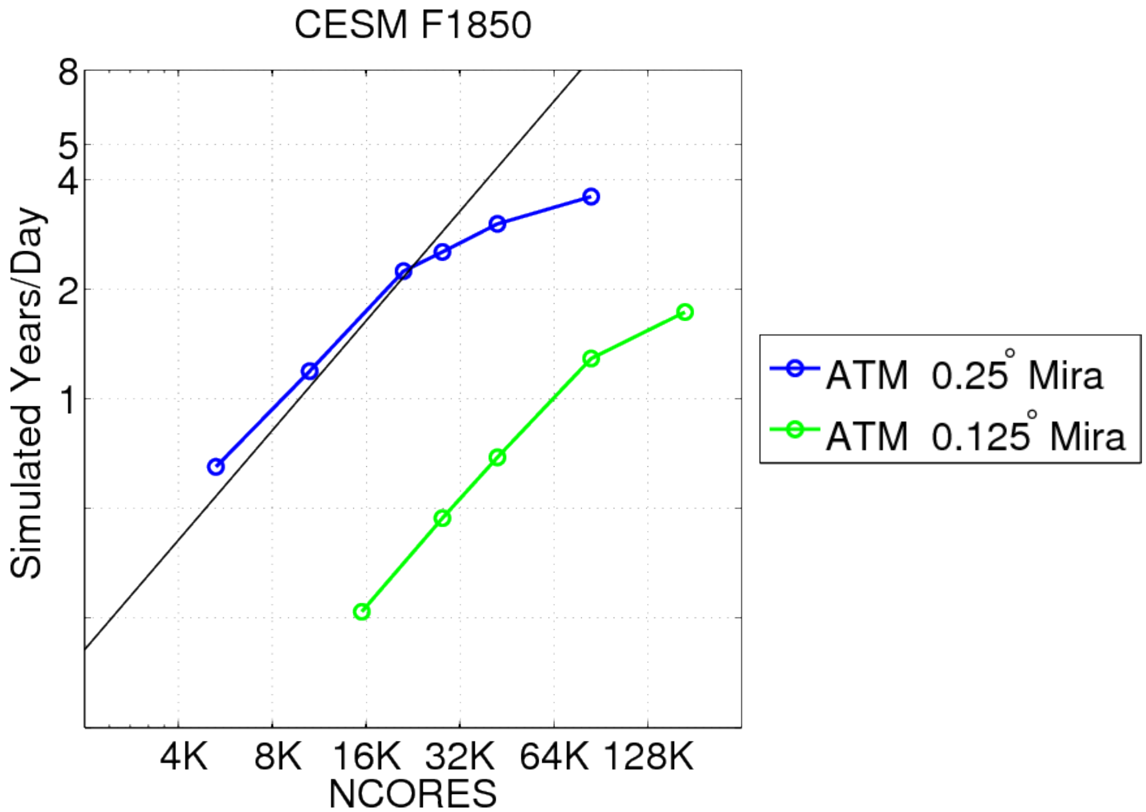

- BG/Q system with 16 cores per node, each with 4 FPU
- CESM running with 4 MPI tasks per node, each with 16 openMP threads
- Best performance: 0.25° 3.6 SYPD
- Best performance: 0.125° 1.7 SYPD

Weather Resolving CESM on Mira

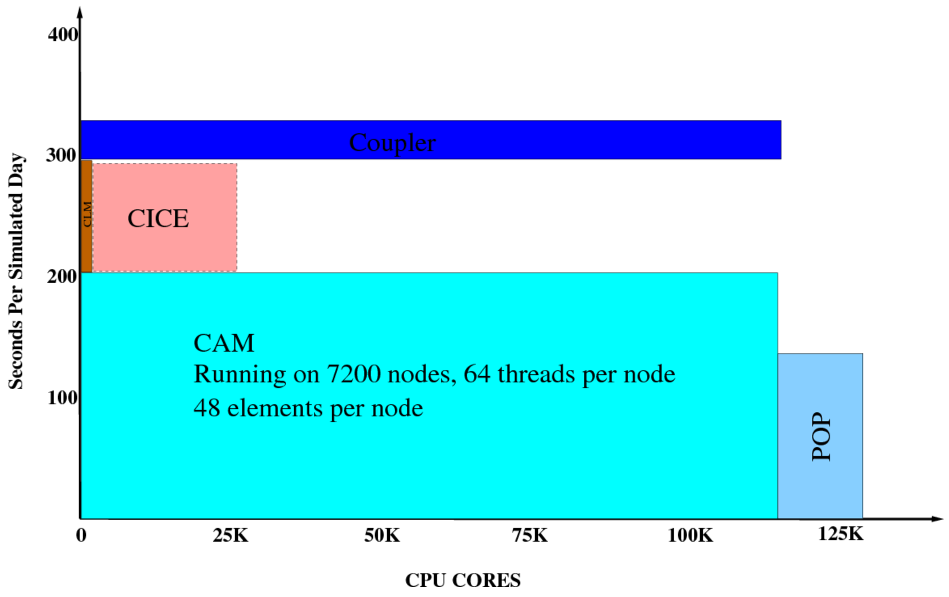

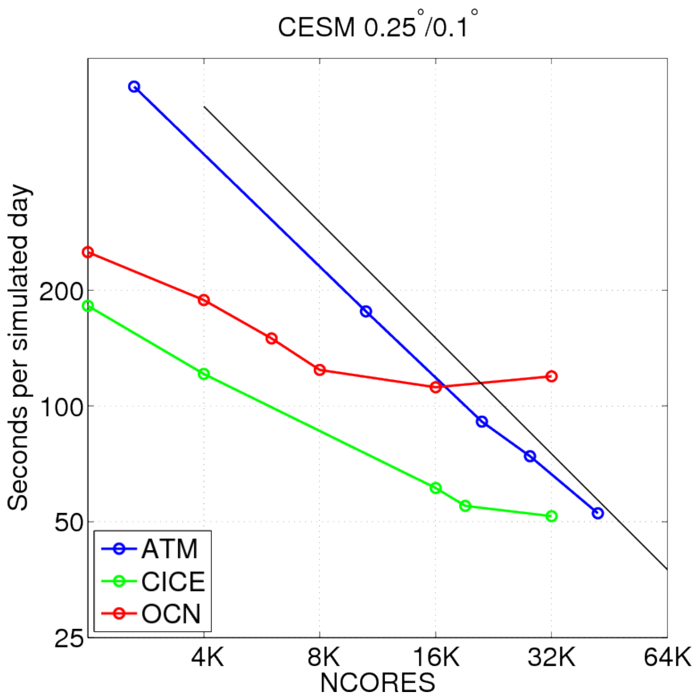

POP and CICE ocean models: new dycore could dramatically improve scaling, making their cost negligible as compared to the atmosphere

Strong scaling of CAM (27km), POP (10km) and CICE (10km) Within the CESM, running on Titan.

Heroic efforts needed in atmosphere to improve performance on Mira

Summary

Sandia National Laboratories

- Weather resolving CESM runs on 20% of Mira at 0.7 SYPD.
- Weak scaling: Cloud resolving CESM could run on 20% of an exascale systems at 0.07 SYPD
- Current model still needs much work on petascale systems to achieve 5 SYPD
  - improved scaling to use more nodes, and better node performance (accelerators and better threading performance)

Publications

- Taylor, Fournier, *A compatible and conservative spectral finite element method on unstructured grids*, J. Comput. Phys. 2010
- Taylor*, Conservation of mass and energy for the moist atmospheric primitive equations on unstructured grids*, Springer, Lecture Notes in Computational Science and Engineering, 2011
- Dennis, Edwards, Evans, Guba, Lauritzen, Mirin, St-Cyr, Taylor, Worley, *CAM-SE: A scalable spectral element dynamical core for the Community Atmosphere Model*, Int. J. High Perf. Comput. Appl., 2012
- Reed, Jablonowski, Taylor, *Tropical cyclones in the spectral element configuration of the Community Atmosphere Model*, Atmos. Sci. Let. 2012
- Levy, Overfelt, Taylor, *A Variable Resolution Spectral Element Dynamical Core in the Community Atmosphere Model*, SNL Tech note 2013-0697J, 2013.
- Zarzycki, Levy, Overfelt, Taylor, Jablonowski, *Using Variable Resolution Meshes to Model Tropical Cyclones in The Community Atmosphere Model,* MWR 2013
- Evans, Lauritzen, Mishra, Neale, Taylor, Tribbia, *AMIP Simulation with the CAM4 Spectral Element Dynamical Core*, J. Climate, 2013
- Lauritzen, Bacmeister, Dubos, Lebonnois, Taylor, *Held-Suarez simulations with the Community Atmosphere Model Spectral Element (CAM-SE) dynamical core: a detailed global axial angular momentum analysis using Eulerian and floating Lagrangian vertical coordinates*, J. Adv. Model. Earth Syst., 2014
- Guba, Taylor, St.Cyr, *Optimization based limiters for the spectral element method,* J. Comput. Phys.. 2014

Extra Slides

CAM-SE

- Hydrostatic equations
- Vertical Discretization
  - hybrid pressure/terrain following coordinate (Simmons & Burridge, 1981)
  - Vertically Lagrangian (S.J. Lin 2004) with monotone remap
- Horizontal Discretization:
  - Conservative spectra elements (Taylor & Fournier, JCP 2010)
  - SE monotone limiter for tracers (Guba et al., JCP, under review)
  - Hyperviscosity used for KE dissipation and numerical stabilization (Dennis et al., IJHPCA 20120)
- Runge-Kutta time-stepping
  - Tracers: 3 stage 2^{nd} order SSP
  - Dynamics: 5 stage 3^{rd} order high-CFL

Spectral Element Method

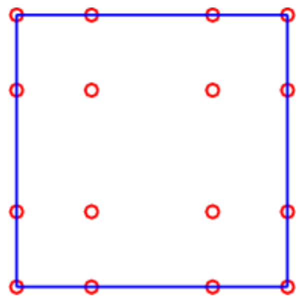

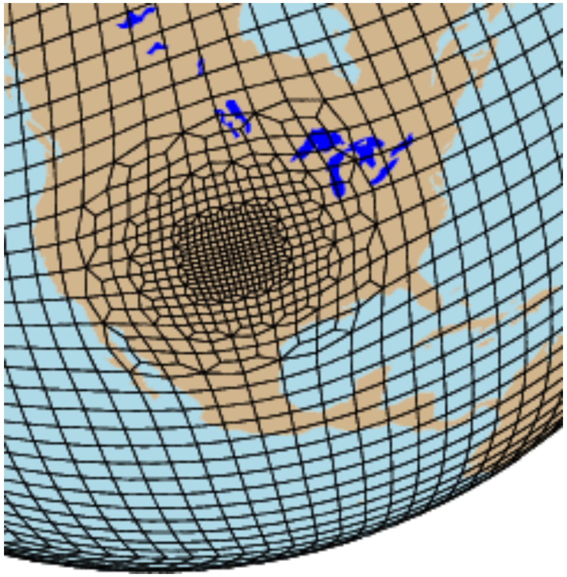

- Q^p-Q^p continuous Galerkin finite element (typically p=3)
- Maps well to modern computers: Arithmetically dense computations, high data locality
- Unstructured, conforming quad meshes (cubed-sphere for uniform grids)
- Mimetic: discrete operators preserve adjoint and annihilator properties of div, grad and curl. Leads to local conservation of mass, tracer mass, energy, 2D PV and linear balance preservation (Taylor & Fournier, 2010)

CESM High-Res Configuration Titan ¼° atmosphere/land - 1/10° ocean/ice

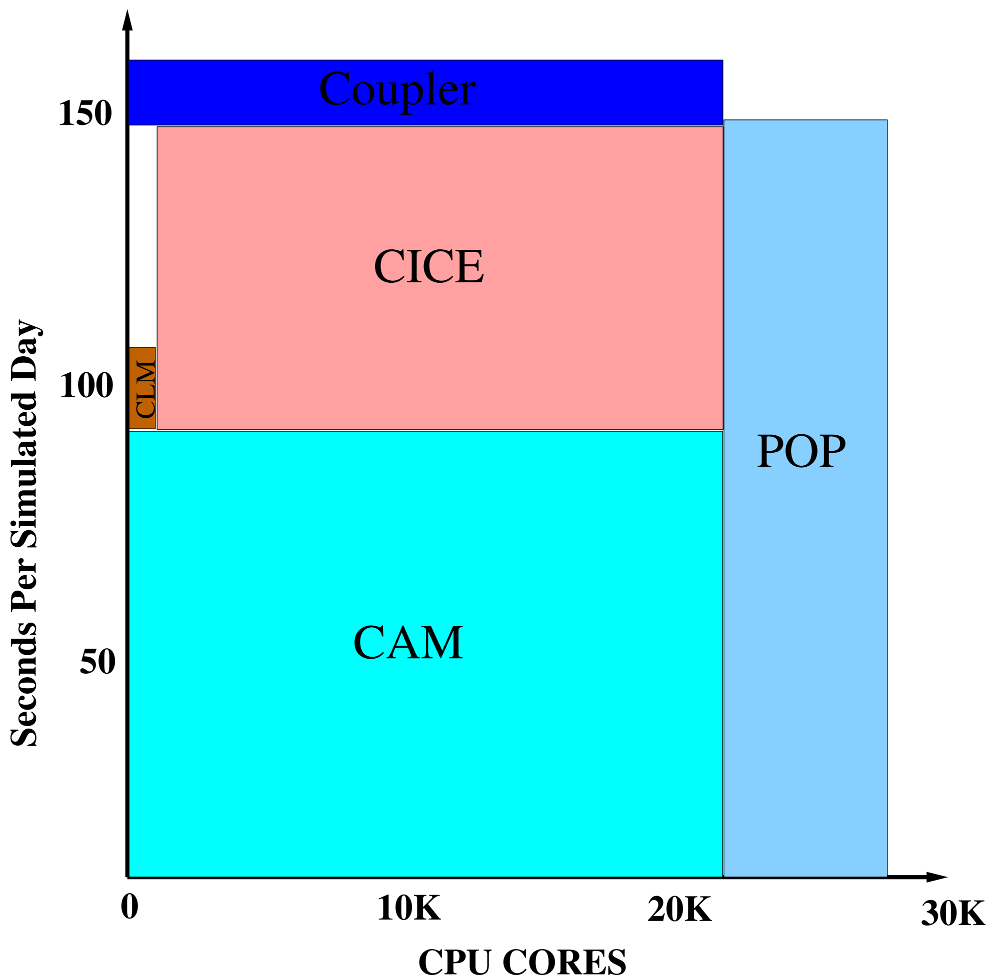

CESM Processor layout and performance (drawn to scale) for an efficient fully coupled CESM configuration running on Titan at a cost of 0.5M core hours per simulated year.

For simplicity the coupler is depicted as running entirely in serial with respect to ATM, CICE an CLM but it does in fact have subcomponents running concurrently.

Conservation

- Mass and Tracer Mass: machine precision and consistent
- Total Energy:
  - KE, PE, IE term-by-term balance preserved to time truncation error
  - KE dissipation added back as heating
  - TE dissipation (S&B81 vertical discretization) < 0.01 W/m^2
  - TE dissipation with vertically lagrangian code ~ 0.2 W/m^2
  - Taylor, Springer LNCSE 2011

Conservation

- Axial Angular Momentum
  - CAM-FV has poor total AAM conservation (spurious sources ~ physical sources)
  - CAM-SE: Spurious sources of total AAM about 1e-3 smaller than physical sources.
  - Not understood why CAM-SE does so well: Lauritzen et al, under review, J. Adv. Model. Earth Syst., 2013

SE/CG and DG

- Spectral elements: Continuous Galerkin (CG) method, closely related to Discontinuous Galerkin (DG) methods.
- Dynamical cores for these methods often have both CG and DG options: HOMME, NUMA, KIAPS?
- SE/CG and DG are quite similar, with a few key distinctions

SE/CG and DG

Sandia National Laboratories

- Stabilization: computational modes present in 2dx-4dx scales are stable but need to be damped to prevent accumulation of noise:
  - CG: hyperviscosity
  - DG: controlled by upwinding
- KE/Enstrophy dissipation: At 1 degree, all models seem to need about 1 W/m^2 of horizontal KE dissipation:
  - CG: hyperviscosity
  - DG: ? (lack of good hyperviscosity prevents CAM-DG)
- CG has better conservation due to its mimetic properties (DG only conserve quantities solved for in conservation form)
- DG has less communication (no corner coupling)
- DG can handle non-conforming grid refinement