EarthWorks Progress

David Randall
Colorado State University

EarthWorks

EarthWorks is a five-year project, in which CSU and 3 NCAR laboratories are working together as partners.

The goal is to develop a *global convection-permitting and ocean-eddy permitting coupled model* based on the CESM.

Earthworks is supported by a grant from NSF/CISE. The CSU/NCAR split is 60/40.

EarthWorks is also partnering informally with DOE, through our use of MPAS-Ocean and interactions with Andrew Gettelman and scientists at the Los Alamos National Lab.

This is a team effort.

David Randall[1], James Hurrell[1], Donald Dazlich[1], Lantao Sun[1], Andrew Feder[1],
William Skamarock[2], Brian Medieros[2], Xingying Huang[2],
Sheri Mickelson[2], Supreeth Suresh[2], Thomas Hauser[2], Ming Chen[2],
Dylan Dickerson[2], Brian Dobbins[2], Michael Duda[2], Jim Edwards[2], Chris Fisher[2],
Jihyeon Jang[2], Mariana Vertenstein[2], John Clyne[2], Richard Loft[3], Andrew Gettelman[4],
Phil Jones[5], Luke Van Roeckel[5], John Cazes[6], Gunther Huebler[7],
Raghu Raj Prasanna Kumar[8], Pranay Reddy Kommera[8], Carl Ponder[8]

[1]Colorado State University, [2]National Center for Atmospheric Research, [3]AreandDee, LLC, [4]Pacific Northwest National Laboratory,
[5]Los Alamos National Laboratory, [6]University of Texas at Austin, [7]University of Wisconsin, Milwaukee, [8]NVIDIA

NCAR software engineers have been key to our progress.

Rich Loft and Sheri Voelz created some of the slides in this talk.

The EarthWorks Vision

- Using CESM as a base, create a coupled model that uses the same 3.75-km geodesic grid for atmosphere, ocean, and land surface. Use of the same grid for all components is central to the EarthWorks vision.
- Develop improved parameterizations that work well with both 3.75-km and 120-km grids.
- Use the model to study both weather and climate, on time scales ranging from days to years.

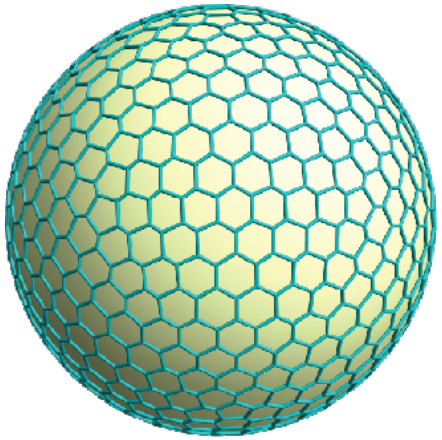

EarthWorks consists of:

- The MPAS non-hydrostatic dynamical core, with a resolved stratosphere and CAM-ish physics
- The MPAS ocean model, developed at Los Alamos
- The MPAS sea ice model, based on CICE
- The Community Land Model (CLM)
- The Community Mediator for Earth Prediction Systems (CMEPS)
- The Common Community Physics Package (CCPP), when ready

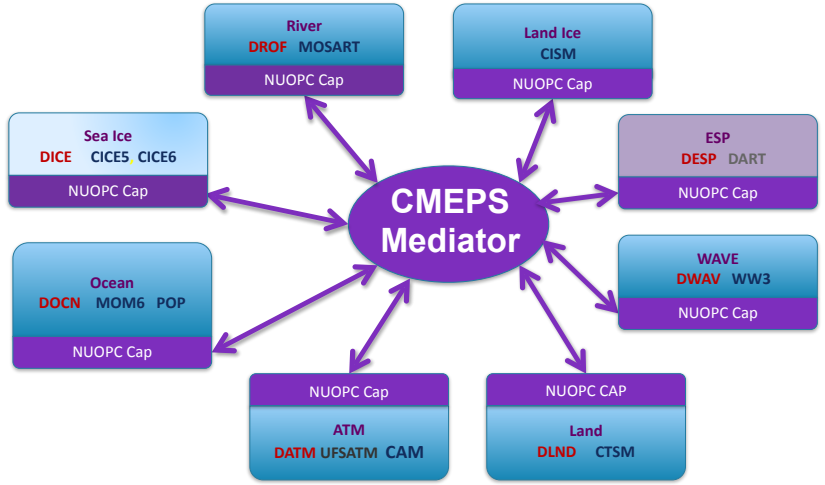

Scientific Goals

- Replace the deep convection and gravity-wave drag parameterizations with direct simulation.
- Demonstrate the resolution-dependent advantages and disadvantages of using the same grid for all ESM components.
- Use the model to study the interactions of mesoscale weather systems with larger spatial scales, on time scales ranging from days to years.
- Use lower-resolution versions of the model to study century-scale climate change.
- Perform deterministic forecasts (days to S2S) with a higher-resolution version of the same model that is used for climate-change experiments.
- Use the model to understand deficiencies of lower-resolution versions of CESM.
- Use the model to create training data sets for machine learning.

Computational Goals

- Find and fix issues with the CESM infrastructure that prevent it from being used with very high resolution.
- Provide CPU compatibility for low resolution experiments and for short tests with ultra-high resolution.
- Provide CPU compatibility (atmosphere and ocean) for ultra-high resolution simulations.
- Achieve about one SYPD with 3.75 km global grid spacing on a DOE "leadership-class" machine in 2025.

Issues encountered (and fixed!)

- **Initialization (abnormally long times)**
  - Traced to an issue in the ESMF framework (needed to turn off subcomponents), resulted in a patch release.
  - *Impact: 5.7x speedup, 2x reduction in memory use during initialization.*
- **Initialization (abnormally long times, MPI communication)**
  - Traced to an issue in the ESMF framework, Alltoall MPI call (~164MB on 164K cores)
  - Workaround was created to eliminate the need for this call, long term fix still needed
  - *Impacts: 41K cores: 45 mins to ~2 mins*
- **Large slowdown in history I/O bandwidth**
  - Traced to the ROMIO MPI-IO implementation in PnetCDF, resulted in a problem report and workaround.
  - *Impact: expected history I/O performance restored*
- **Run after restart errors**
  - Traced to an issue with the PIO2 (parallel I/O) infrastructure in CESM, resulted in a patch release.
  - *Impact: correct model restarts restored*
- **I/O Performance**
  - Sort variables by size and write all similar variables together
  - *Impact: I/O made several times faster*

TC tracks

10-year simulation on 30-km grid using 1990-1999 SSTs

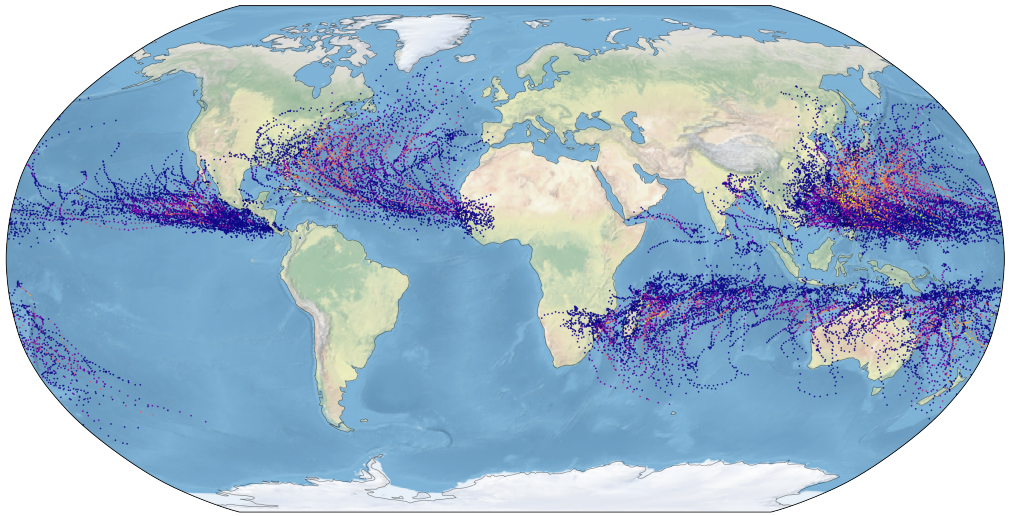

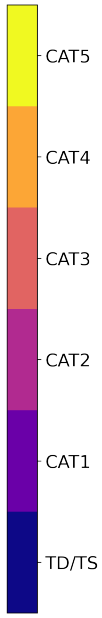

Observations for 1990-1999

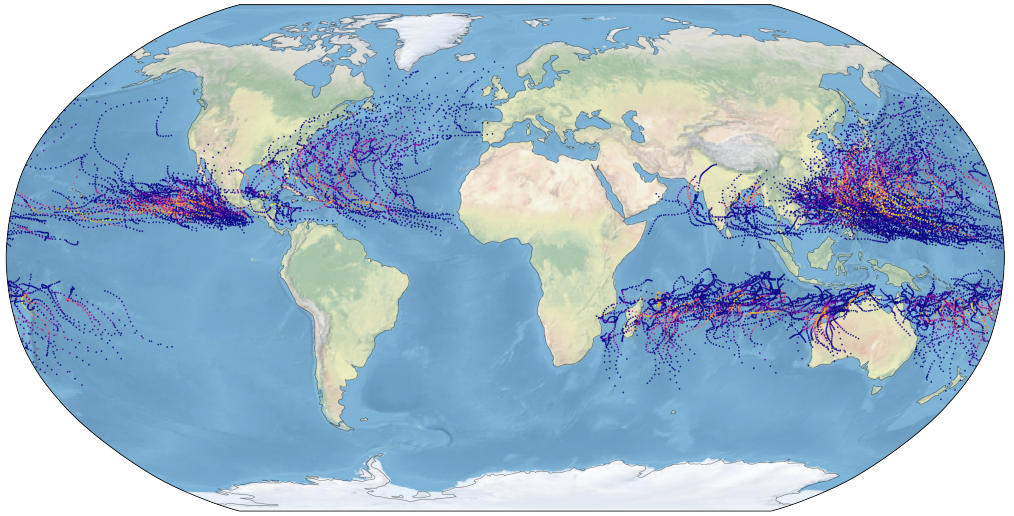

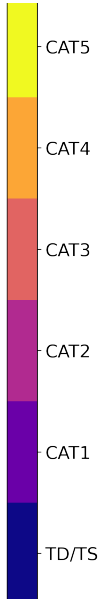

Results from CSU graduate student Andrew Feder

EarthWorks Version 2.2 Software Release

Release Date: June 3, 2024

GitHub Location: github.com/EarthWorks.org/EarthWorks

- Multi-Platform Support
  - We have added support for the GH1, a Grace-Hopper system at the Texas Advanced Computing Center
  - GH1 Grace (CPU) testing has been performed for the FHS94, FKESSLER, and QPC6 (Aquaplanet) compsets only.
  - GH1 Hopper (GPU) offload testing has been performed for the FHS94 (Held-Suarez) test case only.
- Multi-Component GPU Offload
  - Includes MPAS dynamical core, PUMAS microphysics, and RRTMG-P radiative transfer physics code (CLUBB/ GPU will be in next release)
  - Note: performance of the GPU offload has not been fully optimized
- Easier to run "out-of-the-box" model configurations
- More integrated testing
- Developers guide

Current tasks

Running short ultra-high resolution simulations to identify issues

Testing the atmospheric model with GPUs

Maintaining the repository and making sure it's in sync with CESM

Developing workflows for data visualization and analysis on the native MPAS grid and at scale

Next steps

- Optimize RRTMGP and CLUBB.
- Continue to improve performance at the ultra-high resolution on CPUs/GPUs.
- Continue to identify and fix I/O performance issues.
- Study the performance of fully GPU ported atmosphere model on Grace-Hopper.

Looking ahead

The CESM SSC have agreed to designate the model as CESM-E. Support of the model is divided between NCAR and CSU.

CSU and NCAR are currently working on a follow-on proposal (start date 2025).

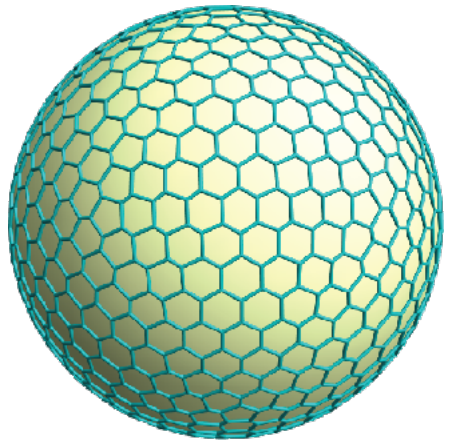

Thanks