The Simple Cloud-Resolving E3SM Atmosphere Model Running on the Frontier Exascale System

Luca Bertagna (lbertag@sandia.gov)
Sandia National Laboratories
Nov. 15th 2023

Mark Taylor, Peter Caldwell, Luca Bertagna, Conrad Clevenger, Aaron Donahue, James Foucar, Oksana Guba, Benjamin Hillman, Noel Keen, Jayesh Krishna, Matthew Norman, Sarat Sreepathi, Christopher Terai, James White, Danqing Wu, Andrew Salinger, Renata McCoy, L. Ruby Leung, David Bader

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Supercomputing and climate research

- Climate change is one of the biggest threats facing humanity
- Climate models need to capture scales from global down to rain drop size
- Climate models have always been big users of HPC

- 1922: Lewis F. Richardson "human supercomputer" for global weather forecast
- 1977: NCAR first official customer of Cray Research, with Cray-1 supercomputer

Energy Exascale Earth System Model (E3SM)

- DOE's state-of-the-art Earth system model, forked from CESM in 2014.
- Several components: atmosphere, land, ocean, land-ice, sea-ice, etc.
- Broad variety of time/space scales.
- Mostly written in Fortran.
- Developed by hundreds of people, contains snippets of codes written across several decades.

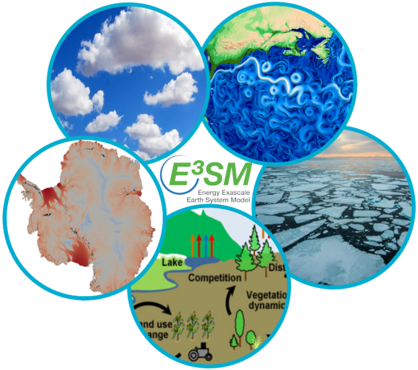

E3SM Atmosphere Model (EAM)

- Broadly speaking, divided in two parts: dynamics and physics
- Dynamics: solves Navier-Stokes equations in rotating spherical framework. It also solves for the transport of tracers in the atmosphere. E3SM uses High Order Methods Modeling Environment (HOMME, M.Taylor) package, which
  - decouples horizontal and vertical differential operators;
  - uses Spectral Element method in horizontal direction;
  - uses Finite Difference methods in vertical direction;
- Physics: approximates processes not resolved by dynamics. Examples include:
  - microphysics: water (vapor, liquid, ice) phase changes and precipitation;
  - macrophysics: subgrid cloud and turbulent processes;
  - radiation: radiative effects on atm temperature;
  - aerosols: cloud and radiative effects of transported particles.
  - deep convection: thermally driven turbulent mixing of air

EAM discretization

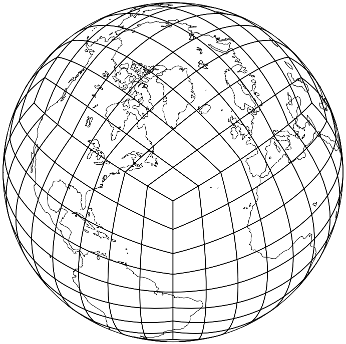

(a) Example of cubed sphere quadrilateral mesh

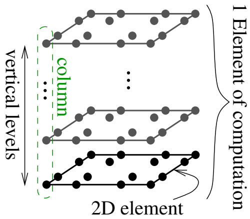

(b) Structure of a vertically extruded 2d Spectral Element.

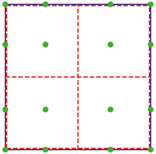

(c) overlap of dynamics (blue line, green dots) and physics (red lines) 2d computational element.

Why cloud-resolving resolutions?

- Global Cloud-Resolving Models (GCRM) avoid the need for convection parameterizations, which are the main source of climate change uncertainty (Sherwood et al., Nature 2014)
- Resolved convection will substantially reduce major systematic errors in precipitation because of its more realistic and explicit treatment of convective storms.
- Improve our ability to assess regional impacts of climate change on the water cycle that directly affect multiple sectors of the US and global economies, especially agriculture and energy production.

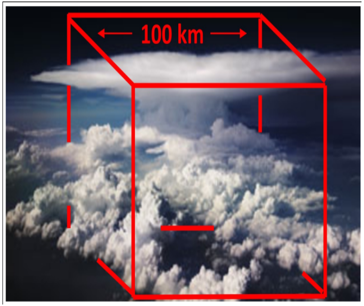

Why cloud-resolving resolutions?

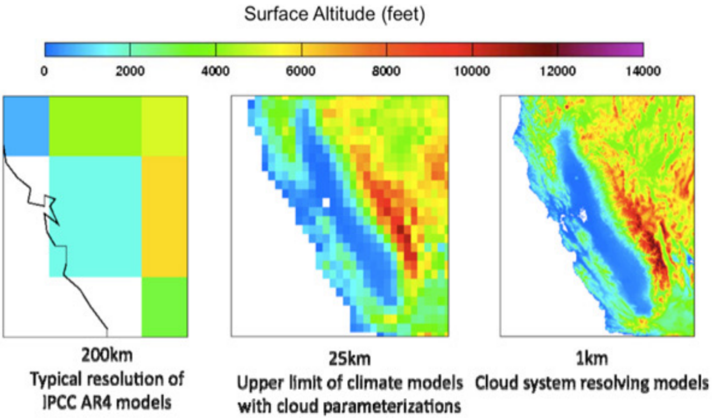

E3SM Simulations

Typical resolution: 100 km, 64 SYPD on 85 CPU nodes

Cloud-resolving: 3.25 km, 1.2 SYPD on 32,000 GPUs

GCRMs are incredibly expensive and typically run on the world's largest computers

https://cs.lbl.gov/news-media/news/news-archive/2009/green-flash-project-runs-first-prototype-successfully/

Simple Cloud-Resolving E3SM Atmosphere Model

The Simple Cloud-Resolving E3SM Atmosphere Model (SCREAM) is a particular configuration of EAM which uses the following

- resolution: 3.25km horizontal, 128 vertical levels
- dynamics: HOMME non-hydrostatic dycore (Taylor, JAMES, 2020)
- microphysics: Predicted Particle Properties (Morrison, Milbrandt, J.Atm.Sci. 2015)
- macrophysics: Simple High Order Closure (Bogenschutz, Kruger, JAMES 2015)
- radiation: RTE+RRTMGP package (Pincus et al, JAMES 2019)
- aerosol: prescribed
- ~~deep convection~~

Note: at 3.25km horizontal resolution, with 128 vertical levels, we have ~7.2B degrees-of-freedom *per variable* on the dynamics grid.

Performance portability strategies

While GPUs are "king" in the exascale era, CPU performance is still crucial in E3SM.

Performance Portability: capability of a code base to run “efficiently” on a variety of computer architectures. Three main approaches:

- **compiler directives**: hint/force compiler on how to optimize (e.g., OpenMP, OpenACC).
- **general purpose lib**: delegate architecture-specific choices to a library (e.g., Kokkos, Raja, etc.)
- **domain-specific lang/lib**: add intermediate compilation step, to generate optimal source for a given architecture (e.g., psyclone, gridtools, etc.).

Note: maintaining multiple versions of E3SM (one for each HPC architecture) is not viable approach. Performance portability **is a must** in the path to exascale.

The Kokkos programming model

- C++ library for on-node parallelism
- Provides constructs for expressing parallelism: execution space, execution policy, parallel operation.
- Provides constructs for multi-dimensional arrays: data type, memory space, layout, memory access/handling.
- Supports several back ends: OpenMP, Pthreads, Cuda, HIP, SYCL, etc.
- Very Reliable: large pool of (world-class) developers, heavily tested, countless downstream apps, closely follows new HPC architectures.

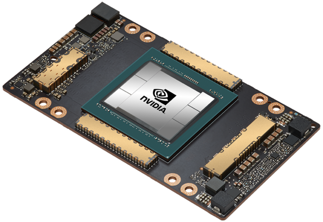

Innovations: software engineering

- Rewrite all atm model in C++, using Kokkos for on-node parallelism.
- Use hierarchical parallelism
  - Implemented via Kokkos TeamPolicy
  - Helps maintain cache friendly code structure for CPUs
  - Threads share common intermediate work, while minimizing index book keeping.
  - Allows for large kernels, minimize kernel launch overheads
- Implement "Pack" data structure to enhance C++ vectorization
  - Necessary to be competitive with Fortran excellent auto-vectorization
  - Support for packed conditionals
- Use workspace manager to use and share minimally sized scratch memory within parallel loops.

Innovations: algorithmic

- Discrete Hamiltonian formulation of the Nonhydrostatic Equations
  - Combined with mimetic horizontal discretization (4th order spectral elements) and mimetic vertical discretization (finite differences) ensures mass and energy conservation
- Timestepping: HEVI-IMEX
  - Horizontally Explicit Vertically Implicit coupled with Runge-Kutta IMEX
  - 5 stage high-CFL explicit Butcher table (3rd order) coupled to 2nd order implicit Butcher table
- Multi-tracer efficient Semi-Lagrangian transport, with mass conservation and monotonicity preservation
- Physical parameterizations run a lower resolution grid as compared to the dynamics (roughly 2.25 times less horiz dofs)

Strong scaling at 110km

- On CPU, C++ code as fast as original Fortran (or faster)
- On CPU, scaling limit reached at 1 elem/core
- On GPU, performance deteriorates at $\lesssim$200 elem/GPU
- Performance portability
  - IBM P9, AMD Epyc
  - NVIDIA V100, A100
  - AMD MI250

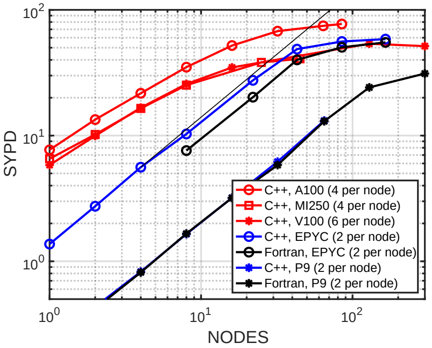

Strong scaling at 3.25km

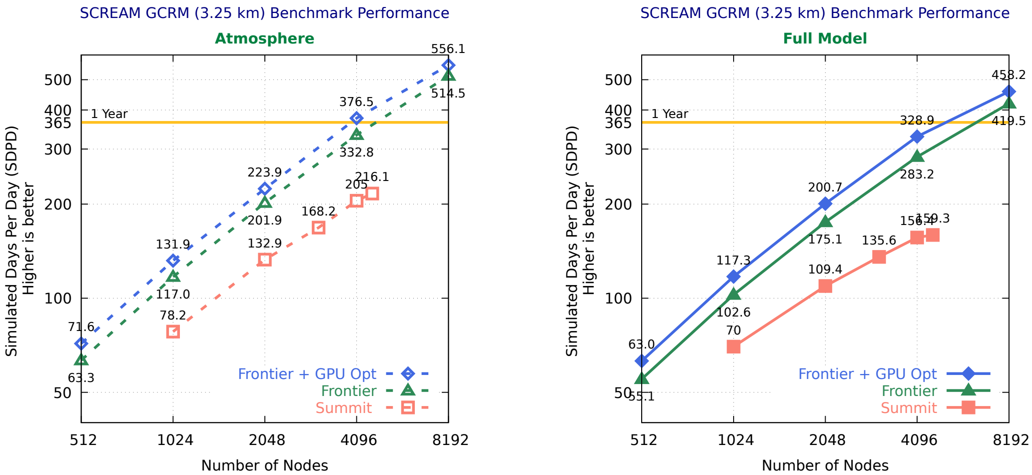

First Global Cloud-Resolving Model to run on an Exascale computer, break the 1 SYPD barrier at cloud resolving resolutions, and run on both AMD and NVIDIA GPUs (as well as conventional CPUs).

SCREAM climate science simulations

- **DYnamics of the Atmospheric general circulation Modeled On Nonhydrostatic Domains (DYAMOND)** [COMPLETED CY2022]
  - Used for model intercomparison and initial model fidelity assessments
  - 40 day simulations (for each season), initialized with real world conditions
  - I/O: 2.6TB per simulated day
- **Cess-Potter: Climate sensitivity** [ONGOING]
  - Used to estimate climate sensitivity w.r.t an increased Sea Surface Temperature
  - A pair of 1 year simulation in progress on Frontier
  - I/O: 0.3 TB per simulated day
- **Atmosphere Model Intercomparison Protocol (AMIP) Simulations** [PLANNED CY2024]
  - A pair of 30 year simulations following the to be used to study atmosphere internal variability and extreme events

All cloud resolving (3.25km) benchmarks in our Gordon Bell submission use the model configured exactly as in the DYAMOND, Cess-Potter and upcoming AMIP simulations

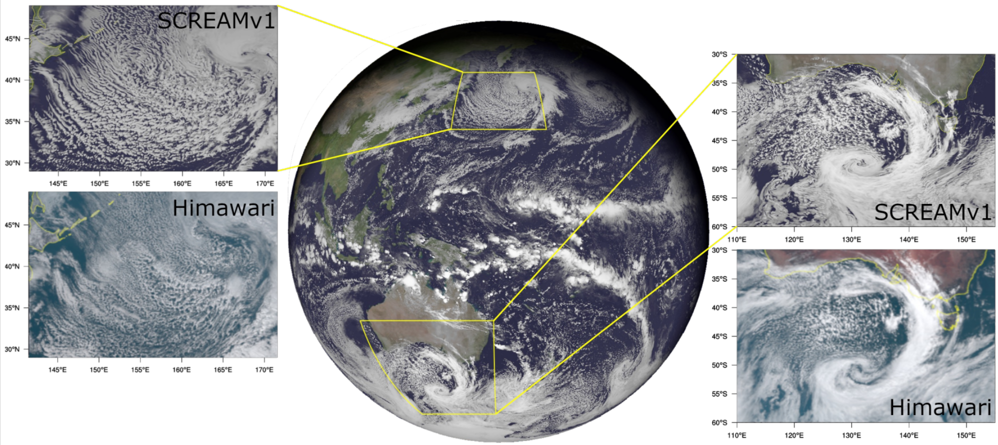

Image from a 40 day SCREAM DYAMOND simulation. Shortwave cloud radiative flux compared to Himawari visible satellite image two days into the simulation (January 22, 2020 at 2:00:00 UTC).

Conclusions

- SCREAM is the first GCRM to run on an Exascale computer, break the 1 SYPD barrier at <5km horizontal resolution, and run on both AMD and NVIDIA GPUs (as well as conventional CPUs).
- Our Gordon Bell configuration is being used for important simulation campaigns, and will help answer important modeling questions.
- E3SM established itself as a reference point for other atmosphere code bases when it comes to exascale performance.
- Kokkos efficiently maps to current architectures, as well as reduces disruption from future ones.
- C++ syntax is *much* heavier than Fortran, which can be a barrier for domain scientists. Good software design is crucial to separate concerns, increase productivity, and make code maintenance manageable.