SAND2019-0840PE

Overview of Exascale Computing Project and Next Generation Development in the Energy Exascale Earth System Model

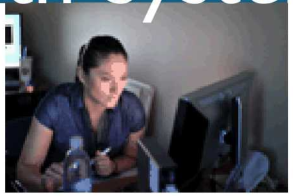

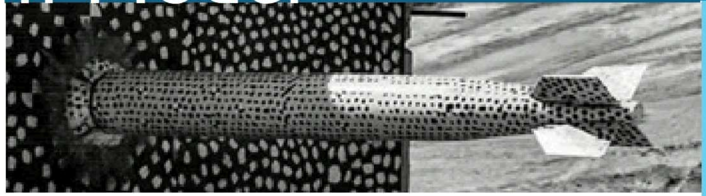

PRESENTED BY

Presented by Benjamin R. Hillman

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & 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.

Earth System Modeling in DOE: E3SM

- Global Earth System Model
- Atmosphere, Land, Ocean and Ice component models
- 8 DOE labs, 12 university subcontracts, 53 FTEs spread over 87 individuals
- Development driven by DOE-SC mission interests: Energy/water issues looking out 40 years
- Key computational goal: Ensure E3SM will run well on upcoming DOE pre-exascale and exascale computers
- E3SM is open source / open development
  - Website: www.e3sm.org
  - Github: https://github.com/E3SM-Project
  - DOE Science youtube channel: https://www.youtube.com/channel/UC_rhpi0lBeD1U-6nD2zvlBA

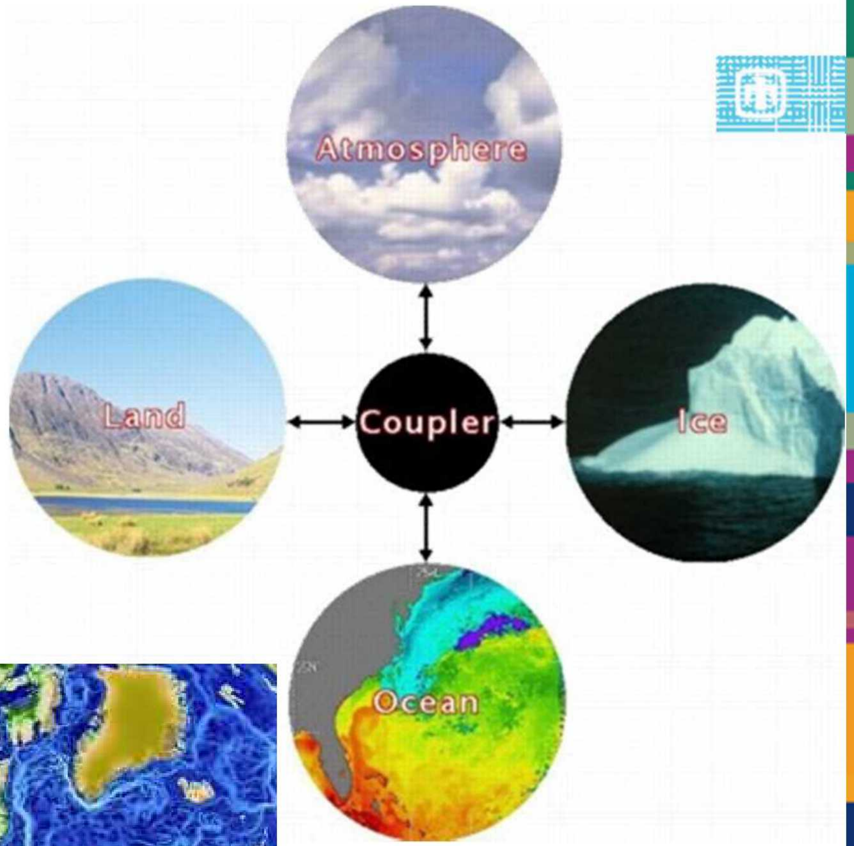

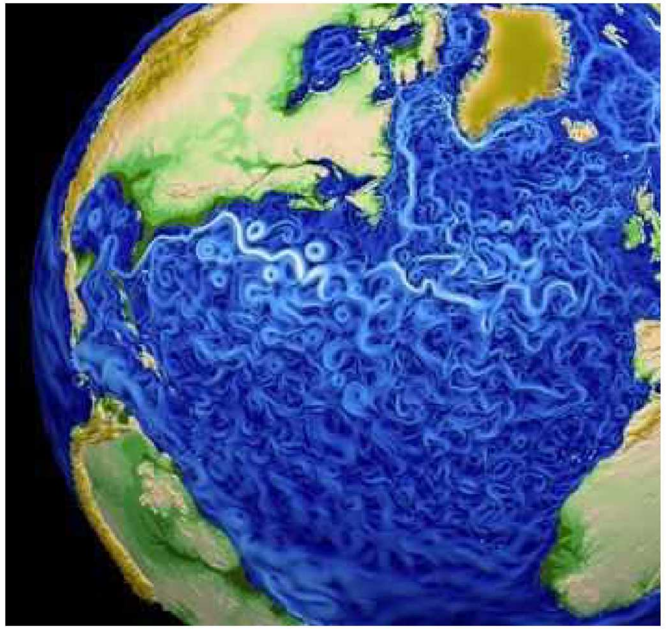

E3SM
Energy Exascale
Earth System Model

Atmosphere Component

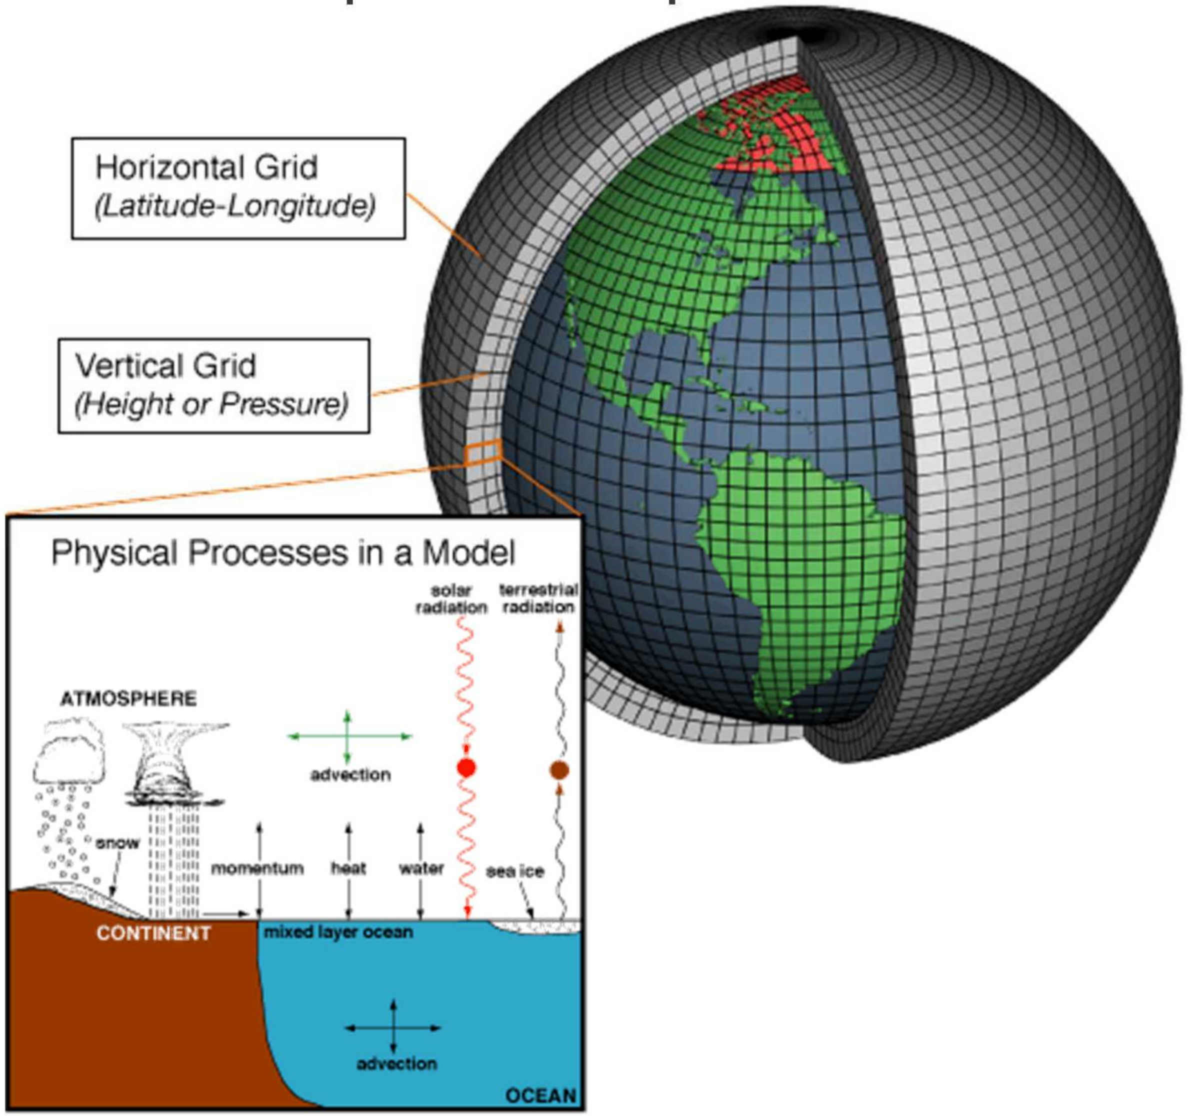

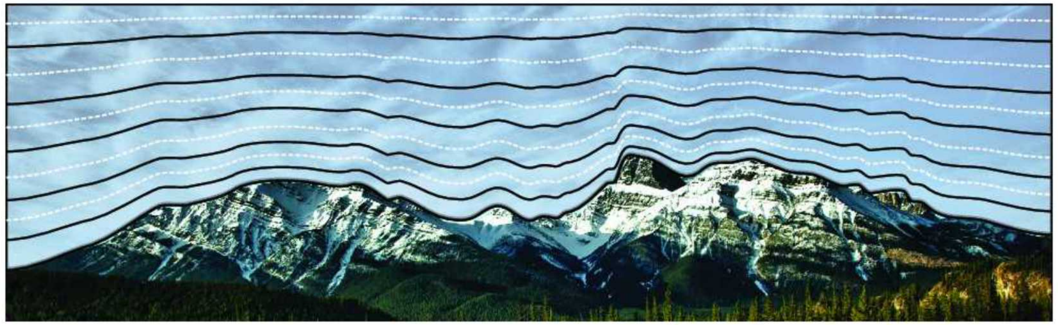

hydrostatic-pressure terrain-following coordinates

Column Physics
- Subgrid parametrizations: precipitation, radiative forcing, etc.
- Embarrassingly parallel with 2D domain decomposition

Dynamical Core
- Solves the Atmospheric Primitive Equations
- Linear transport of ~40 atmospheric species
- Scalability bottleneck

Terrain following figure: D. Hall, CU Boulder
Source: http://celebrating200years.noaa.gov/breakthroughs/climate_model/welcome.html

E3SM Energy Exascale Earth System Model

Two paths toward high resolution, global simulations

- Simplified Cloud Resolving E3SM Atmosphere Model (SCREAM)
  - Produce a *global cloud resolving model* (3 km effective grid spacing)
  - Simplified physics (no deep convection, new microphysics)
  - Goal: to make it possible to run very high resolution, global *atmosphere-only* simulations for short-medium timescales (months to years)

- Exascale Computing Project (ECP) E3SM using the Multi-scale Modeling Framework (E3SM-MMF)
  - Produce a *fully coupled global climate model* with *some aspects of global cloud resolving models*
  - Use the Multi-scale Modeling Framework (MMF) to replace traditional cloud parameterization with *embedded* cloud resolving models
  - Increased computational efficiency over global cloud resolving model, but compromises and challenges arising due to coupling models of different scales
  - Goal: to make it possible to run *climate-scale coupled simulations* (decades to centuries) with some aspects of explicit cloud simulation

Simple Cloud Resolving E3SM Atmosphere Model (SCREAM)

Collaborators: Peter Caldwell (LLNL), Andy Salinger (SNL), Luca Bertagna (SNL), Peter Bogenschutz (LLNL), Andrew Bradley (SNL), Aaron Donahue (LLNL), Jim Foucar (SNL), Chris Golaz (LLNL), Oksana Guba (SNL), Jorge Guerra (UCD), Ben Hillman (SNL), Noel Keen (LBNL), Andrew Steyer (SNL), Mark Taylor (SNL), and Paul Ullrich (UCD)

Context

- Atmospheric General Circulation models (GCMs) are the workhorse of both climate and weather forecasting
- To make century-scale simulations practical, **climate simulations use much coarser resolution than weather models**

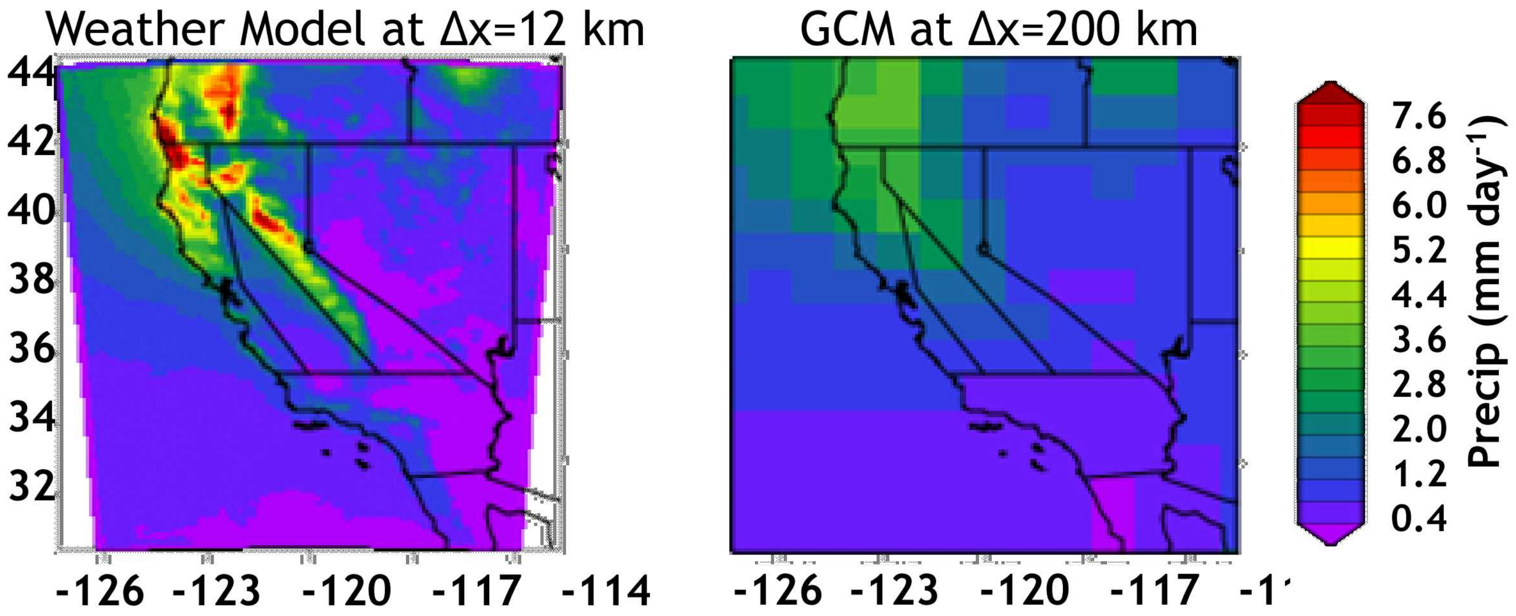

Fig: Precipitation climatology over CA from weather and climate models

(Content courtesy Peter Caldwell at LLNL)

E3SM Energy Exascale Earth System Model

The problem with coarse resolution

- Coarse resolution results in heavy reliance on sophisticated sub-grid parameterizations
  - Uncertain sub-grid parameterizations are a major cause of climate-change uncertainty

Fig: How do we parameterize this sub-grid variability?

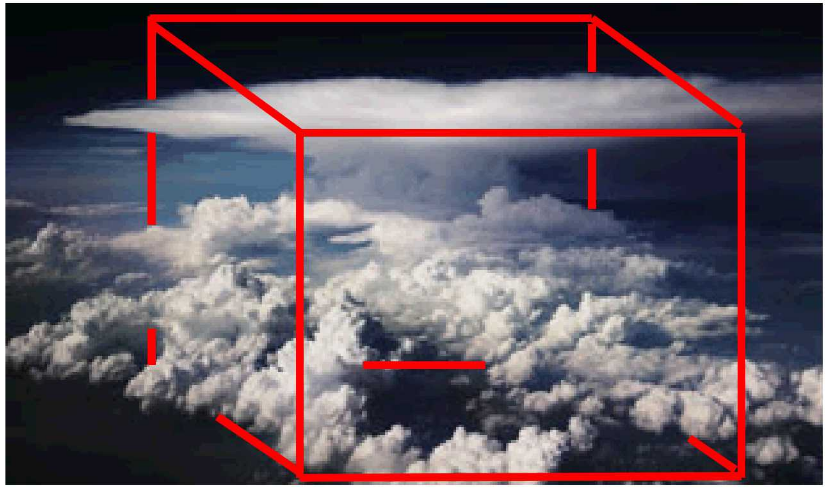

(Content courtesy Peter Caldwell at LLNL)

Computing paradigm shift

- Since chips aren't getting faster, increases in speed are coming primarily from *increased parallelism*
  - Higher resolution means more columns, making it a good fit for exascale machines
  - Caveat: fluid dynamics Δt gets smaller as Δx becomes finer (CFL constraint)
    - this requires us to be clever about making predictions from shorter runs

(Content courtesy Peter Caldwell at LLNL)

E3SM
Energy Exascale
Earth System Model

Plans for SCREAM

- Develop a Global Cloud-Resolving Model (GCRM) for the E3SM coupled system which:
  - is targeted at **3 km resolution**
    - Explicitly resolves large convective events, thought to be the main source of parametric uncertainty (e.g. Webb et al., 2015)
  - is written from scratch in **templated C++**
    - Removes legacy code/enables modern software design
    - Enables use of Kokkos, an ASC-developed library for machine-portable parallelism (e.g. GPUs!)
  - is as **simple** as possible (but not simpler)
    - Makes complete rewrite possible
    - Improves computational efficiency
    - Makes model easier to understand and therefore trust
  - uses modern software best practices (e.g. testing)
    - Tested code is more believable

(Content courtesy Peter Caldwell at LLNL)

The components of a global cloud resolving model

1. Resolved-scale **fluid dynamics** treated by a Spectral Element (**SE**) approach

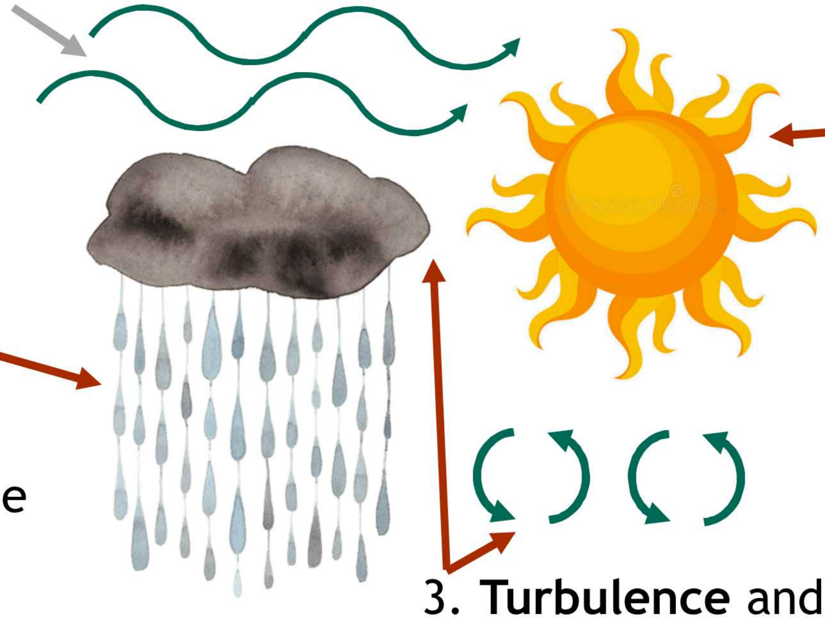

2. **Radiation** handled by externally-developed, GPU-ready **RRTMGP** package

3. **Turbulence** and **cloud formation** handled by Simplified Higher-Order Closure (**SHOC)**

4. **Microphysical processes** handled by Predicted Particle Properties (**P3)** scheme

(Content courtesy Peter Caldwell at LLNL)

E3SM Energy Exascale Earth System Model

Impact for DOE

- Reducing *subgrid assumptions* and *adding rigorous testing* improves our ability to answer questions about:
  - Science (understanding nuclear winter, hurricanes, shifts in weather patterns, etc.)
  - Energy (determining future energy and water demands, siting power plants, etc.)
  - National Security (contaminant dispersion, non-proliferation, and resource availability)

(Content courtesy Peter Caldwell at LLNL)

E3SM Energy Exascale Earth System Model

Super-Parameterized E3SM (SP-E3SM)

Collaborators: Mark Taylor, Ben Hillman (SNL), Jayesh Krishna, Danqing Wu, Nichols Romero (ANL), David Randall, Don Dazlich, Mark Branson (CSU), Philip Jones, Rob Aulwes (LANL), David Bader, Walter Hannah, Jungmin Lee (LLNL), Matthew Norman, Sarat Sreepathi, Marcia Branstetter (ORNL), Ruby Leung, Mikhail Ovchinnikov, Chris Jones, Guangxing Lin (PNNL), Mike Pritchard, Hossein Parishani, Chris Terai (UCI)

The Exascale Computing Project (ECP)

- Created in support of President Obama's National Strategic Computing initiative
- A collaborative effort of Two US Department of Energy (DOE) organizations:
  - Office of Science (DOE-SC)
  - National Nuclear Security Administration (NNSA)
- A 10-year project to accelerate the development of a **capable** exascale ecosystem
  - Led by DOE laboratories
  - Executed in collaboration with academia and industry

A **capable** exascale computing system will have a well-balanced ecosystem (software, hardware, applications)

ECP has formulated a holistic approach that uses co-design and integration to achieve capable exascale

Application Development	Software Technology	Hardware Technology	Exascale Systems
Science and mission applications	Scalable software stack	Hardware technology elements	Integrated exascale supercomputers

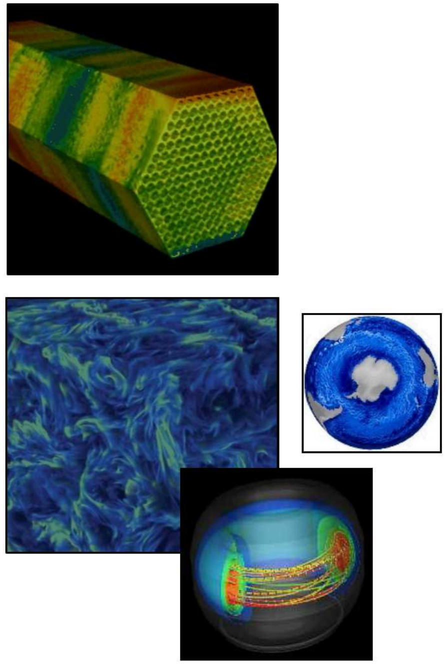

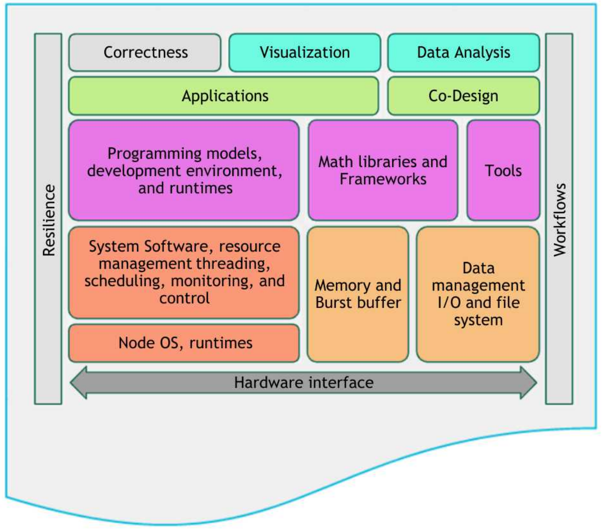

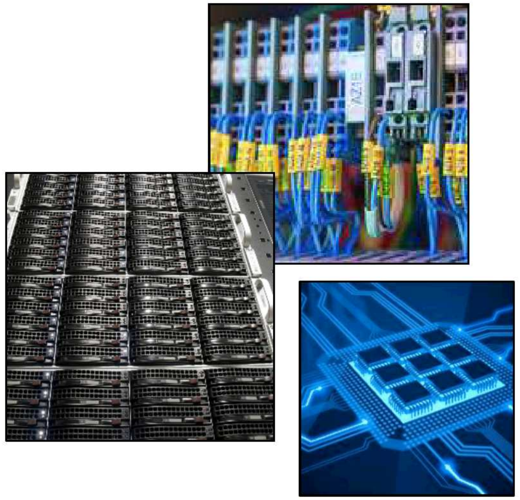

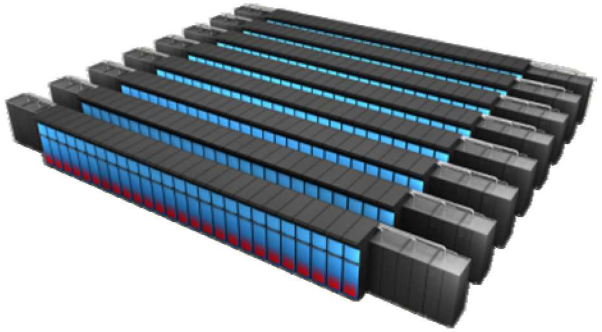

EXASCALE COMPUTING PROJECT

E3SM-MMF Cloud Resolving Climate Model

- Develop capability to assess *regional* impacts of climate change on the water cycle that directly affect the US economy such as agriculture and energy production.
- A cloud resolving climate model is needed to reduce major systematic errors in climate simulations due to structural uncertainty in numerical treatments of convection – such as convective storm systems
- Challenge: Cloud resolving climate model using traditional approaches requires Zettascale resources.
- E3SM-MMF: Use a multiscale approach ideal for new architectures to achieve cloud resolving convection on Exascale resources

Convective storm system nearing the Chicago metropolitan area
http://www.spc.noaa.gov/misc/AbtDerechos/derechofacts.htm

Exascale Challenge Problem

Develop an Earth system model with a fully weather resolving atmosphere and cloud-resolving super-parameterization, an eddy resolving ocean and ice components, all while obtaining the necessary throughput to run 10-100 member ensembles of 100-year simulations.

Definitions:

- Cloud-resolving: 1km grid spacing in both horizontal and vertical directions.
- Weather resolving: 50-25km horizontal resolution, ~1km vertical (the resolution of today's global operational forecast models).
- Eddy resolving ocean/ice: minimum 18 km resolution in equatorial regions, decreasing to 6 km in polar regions to capture the reduction in eddy size with decreasing Rossby radius of deformation, with O(100) levels in the vertical.
- Necessary throughput: 5 simulated-years-per-day.

Exascale Cloud Resolving Modeling

Exascale will make "conventional" cloud resolving simulations routine for shorter simulations

- process studies, weather prediction

For cloud resolving **climate** simulations, we need fundamentally new approaches.

- New algorithms which can take advantage of exascale architectures

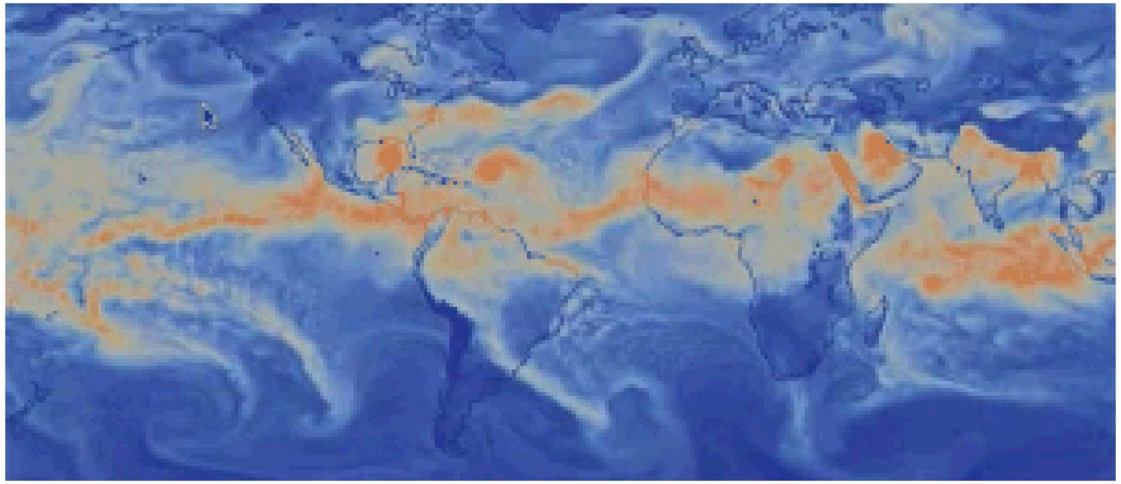

The Multi-scale Modeling Framework (MMF)

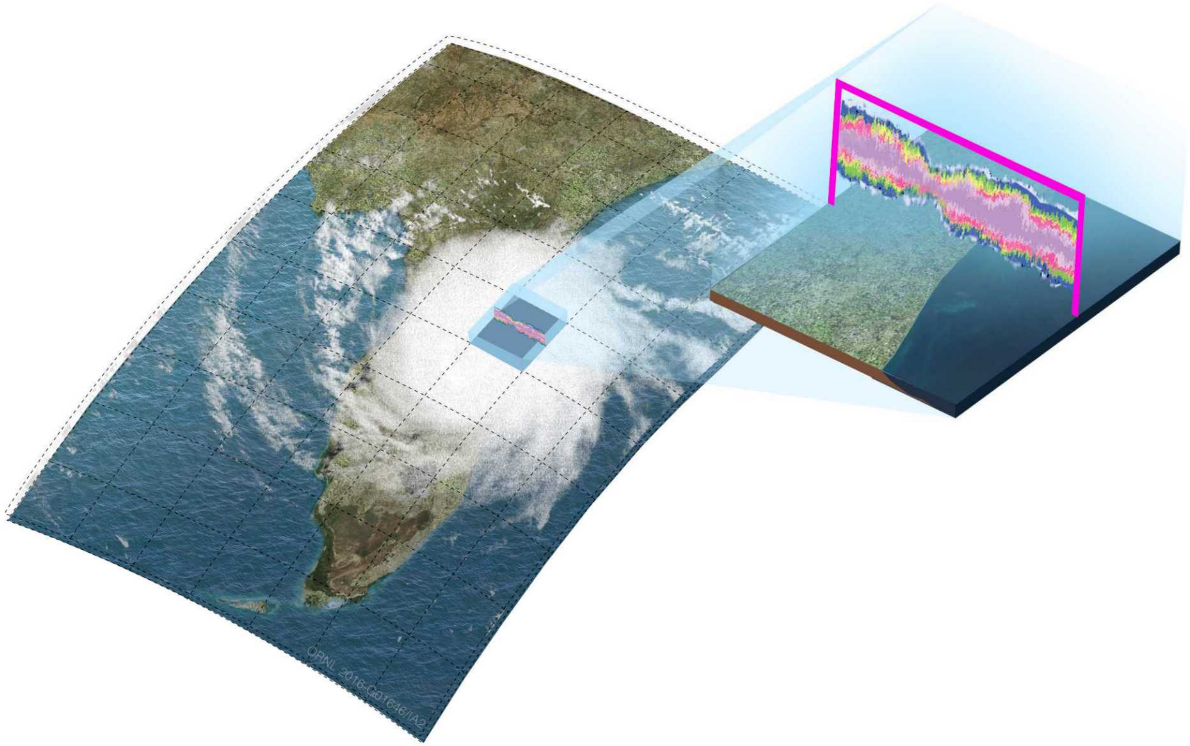

- Address structural uncertainty in cloud processes by replacing traditional parameterizations with cloud resolving model within each grid cell of global climate model
- Ideal approach for exascale
- Exascale + MMF will make it possible for the first time to perform climate simulation campaigns with *some* aspects of cloud resolving resolutions

ECP
EXASCALE COMPUTING PROJECT

Performance

- SP-E3SM is ***~60x slower*** than E3SM (0.1 SYPD vs 6 SYPD on Titan)
- SP-E3SM should be ***~20x faster*** than “conventional” E3SM running with a cloud resolving atmosphere (3 km) and eddy resolving ocean (0.1 SYPD vs 0.005 SYPD)

EXASCALE COMPUTING PROJECT

GPU Acceleration

- Super-parameterization is ideal for GPU acceleration
- Directives-based approach (OpenACC or OpenMP)
- Acceleration limited by work per node and throughput constraint

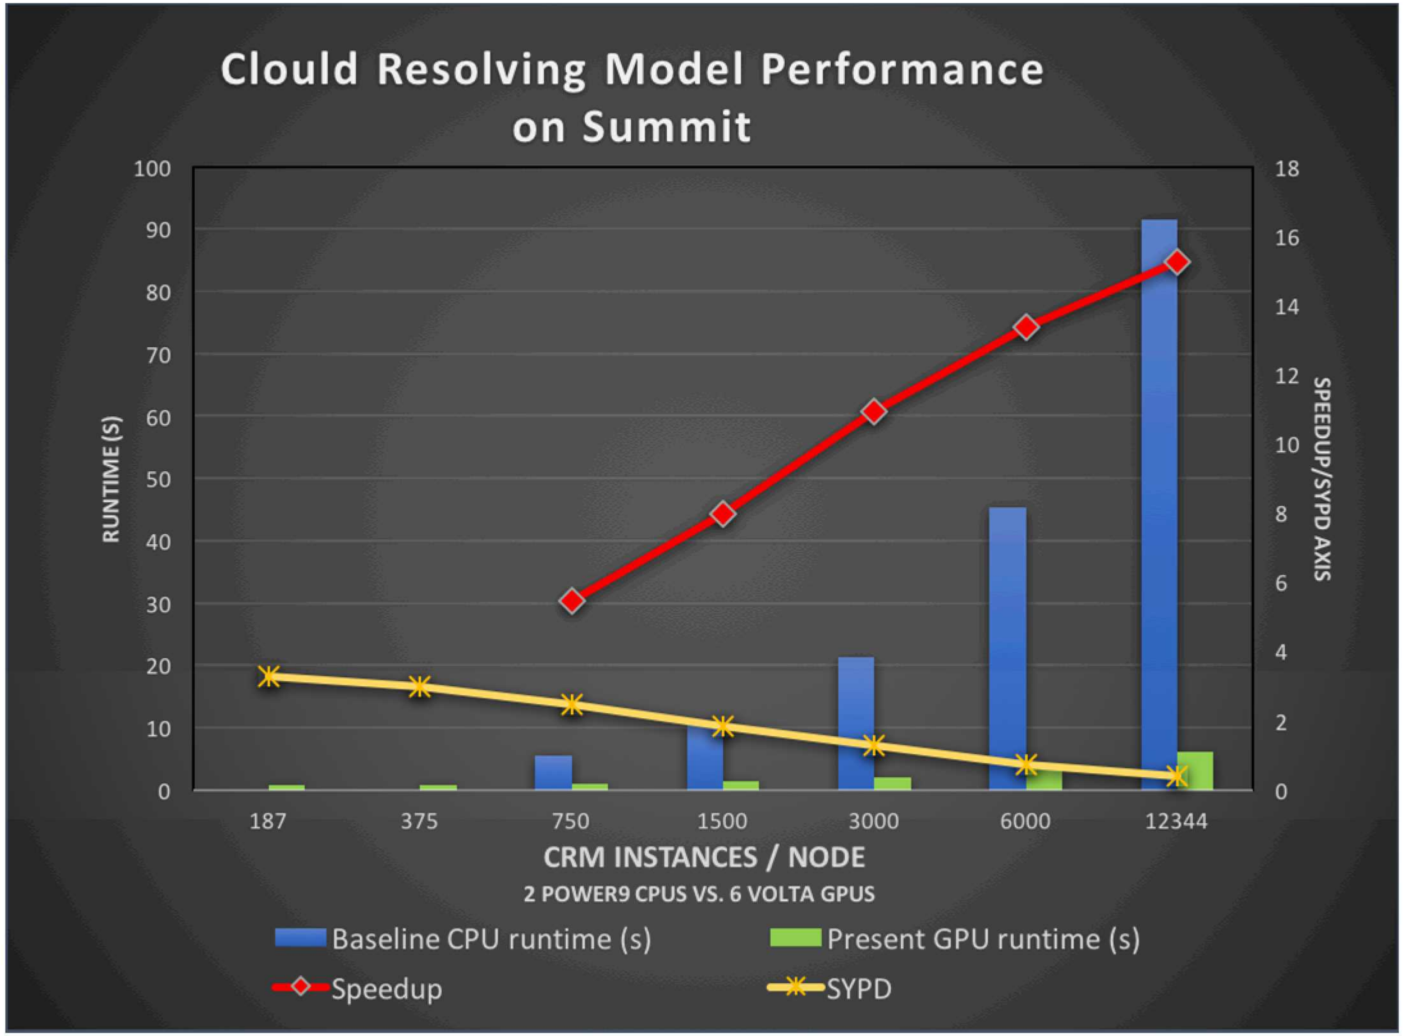

Project breakdown

- 50% MMF science
- 50% Computational work (GPU porting)
  - ECP Expansion: shift to 40/60 as we add A21 support

- LLNL,PNNL: MMF model development
- ORNL: GPU porting (atmosphere)
- LANL: GPU porting (ocean/ice, expanding to A21 support)
- SNL: Software engineering, RRTMGP, baselines
- ANL: I/O (expansion: A21 atmosphere)
- University subcontracts (CSU, UCI): superparameterization pioneers providing critical experience and advice

Early successes and challenges: Mesoscale Convective Systems (MCS)

- 5-10 day hindcasts with E3SM and SP-E3SM at ne30 resolution
- Hindcasts coincide with observed MCSs from May-Aug 2011
- Initial conditions generated from extended E3SM simulation with horizontal winds nudged to ERA-I

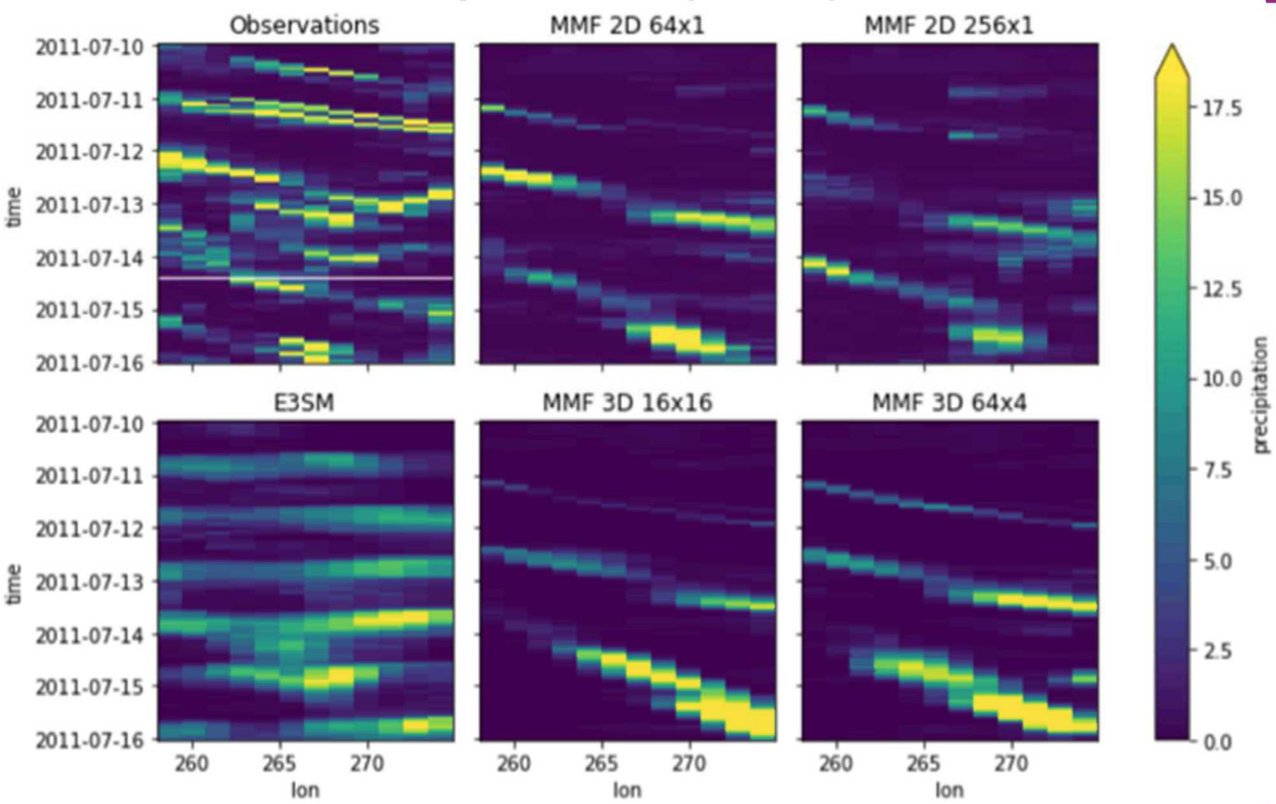

Realistic eastward propagation of MCS in SP-E3SM, regardless of CRM configuration

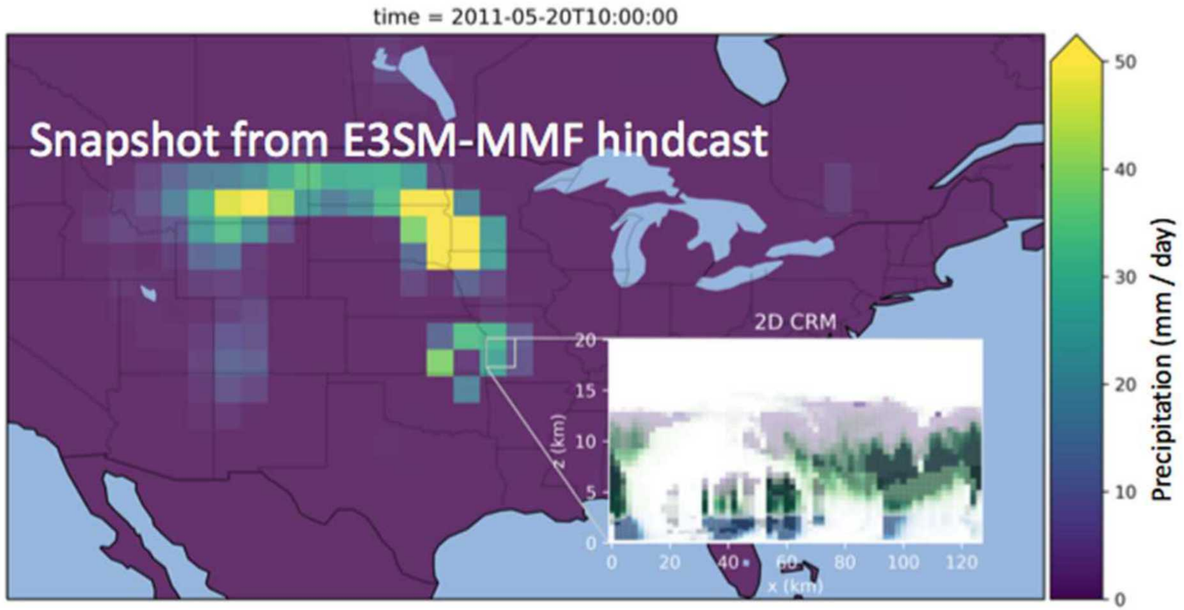

Courtesy Chris Jones (PNNL)

Early successes and challenges: grid imprinting in convective fields

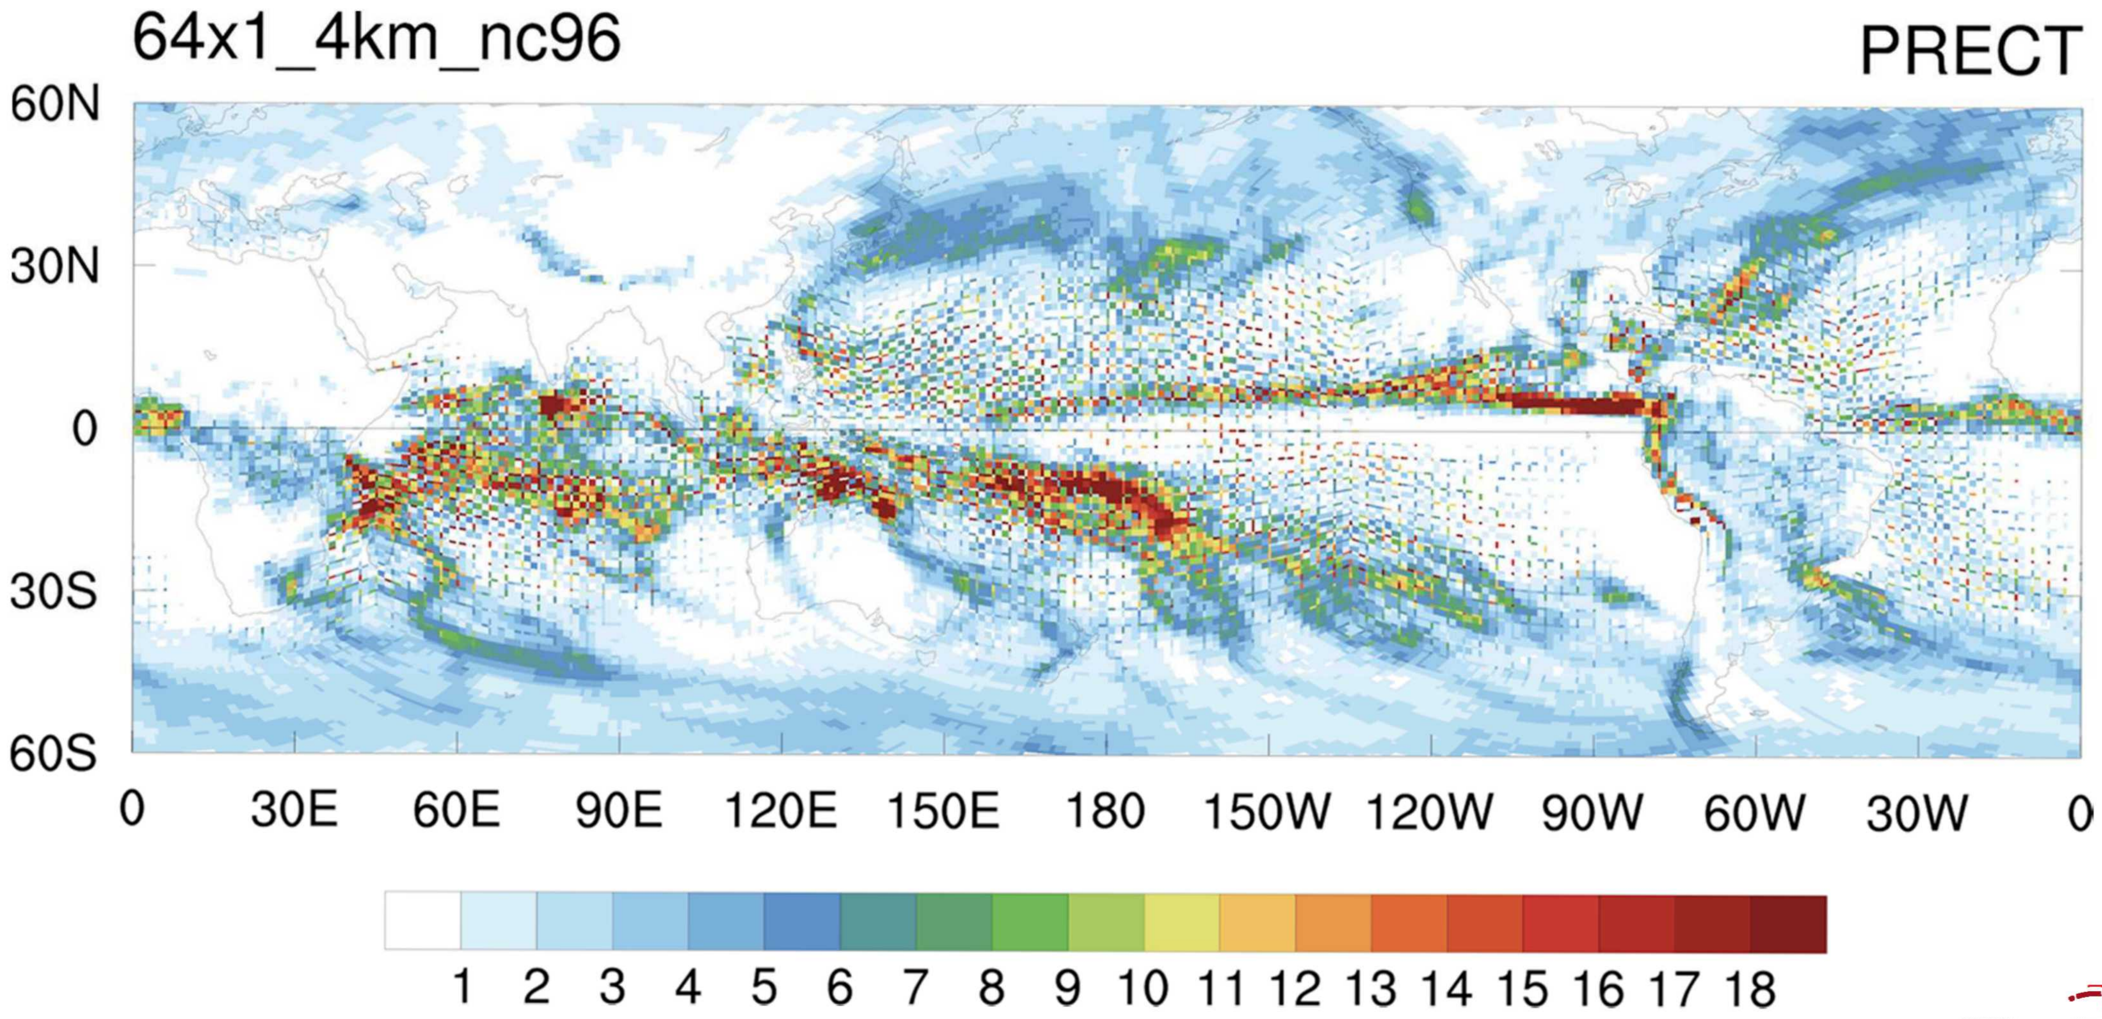

Early successes and challenges: grid imprinting in convective fields

Enhanced precipitation on corner and edge GLL nodes

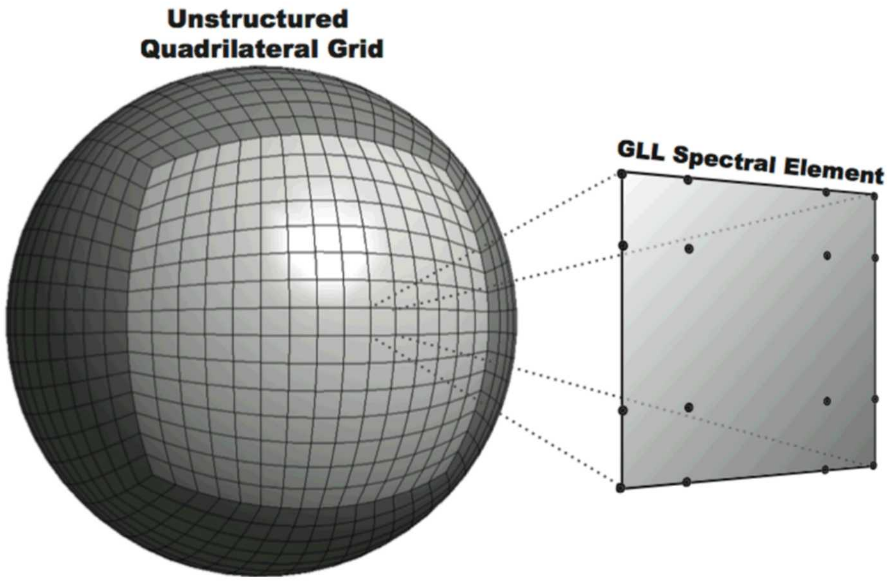

Mitigated with smaller grid spacing, longer GCM timestep, more CRM columns

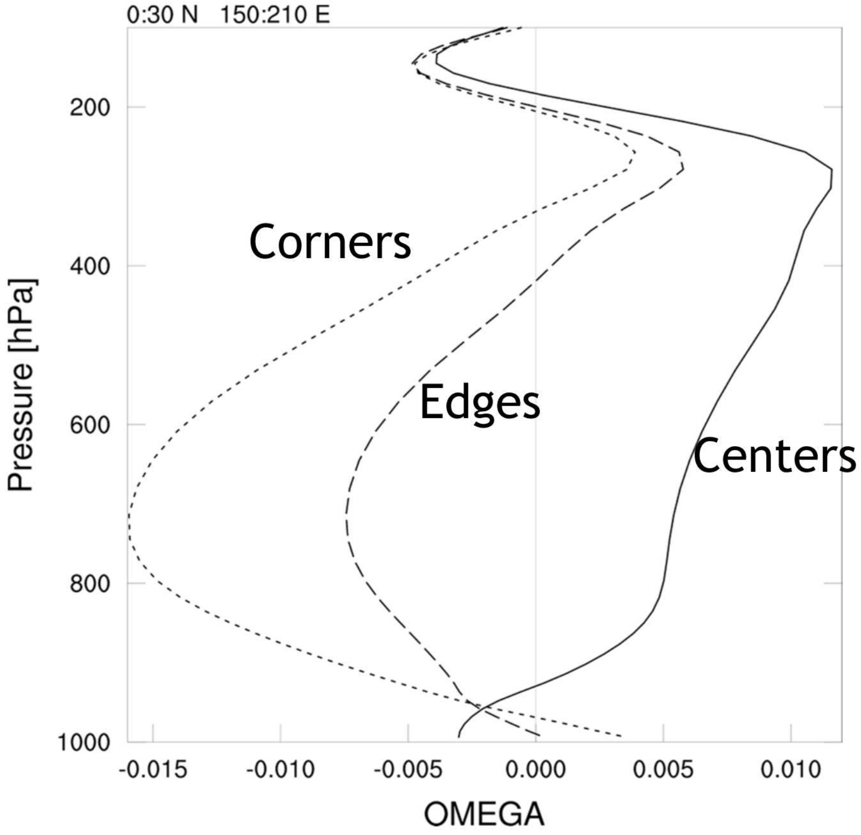

Not evident in fields that are not directly related to convection, but shows up strongly in large-scale vertical velocity

Sandia contributions and opportunities: updating radiative transfer

- The RRTMGP radiation package has been imported and interfaced into both the E3SM-MMF and the standard E3SM used outside of ECP (which will also be used by the SCREAM project)
- Implementing RRTMGP is the first step towards accelerating the radiative transfer, which is responsible for 1/3 the cost of the atmospheric physics
- The new RRTMGP code and interface will serve as the groundwork for our efforts to port the radiation calculations to the GPU and further reduce the computational cost of the atmospheric physics

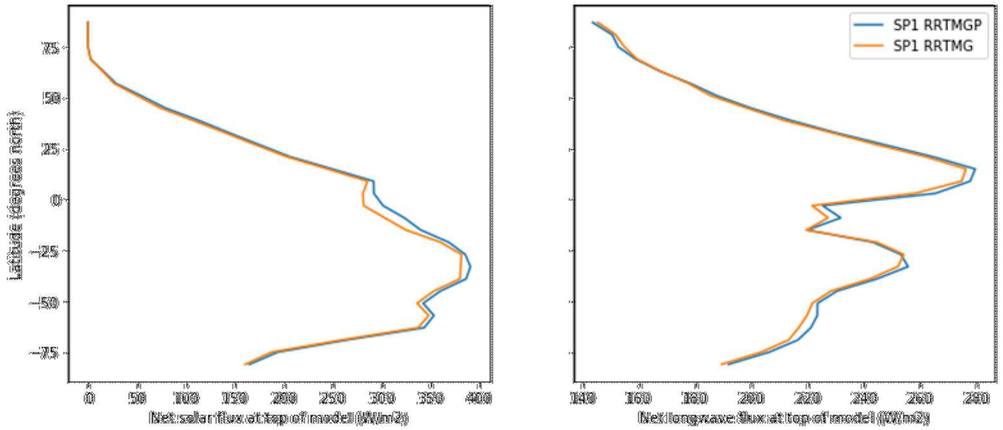

Answers are comparable between RRTMGP and RRTMG, but RRTMGP represents a significant improvement in terms of flexibility and maintainability over RRTMG. Shown are zonal mean shortwave and longwave fluxes, which are comparable between RRTMG and RRTMGP.

Sandia contributions and opportunities: accelerating radiative transfer

- Calculating radiation on *groups* of CRM columns can significantly reduce the cost of the radiation, but with a penalty in accuracy and a loss of cloud-radiative feedbacks – highlights importance of small-scale variability in radiative heating!
- Currently exploring other ways of speeding up radiation
- E3SM-MMF: ideal framework to explore cloud-radiative interactions on multiple scales (how small-scale interactions affect the large-scale)

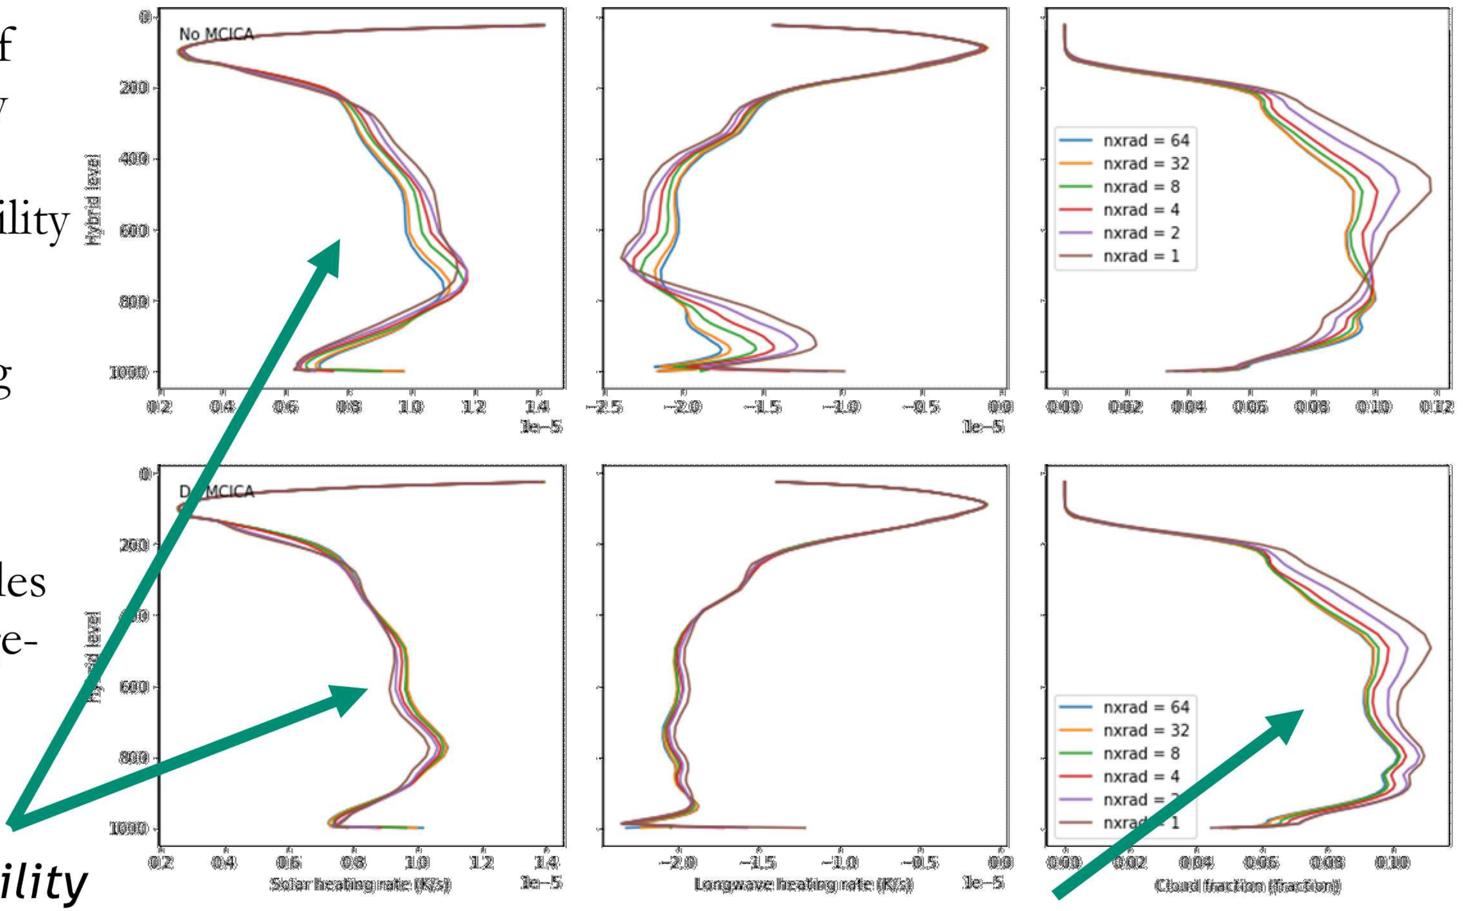

Importance of *cloud variability* on the *radiative heating*

Importance of *heating variability* on the *clouds*

Sandia contributions and opportunities: evaluating Arctic clouds

- Can the Multi-scale Modeling Framework improve simulation of Arctic clouds (LDRD work led by Erika Roesler)?
- Simulations using CAM5 and pre-ECP version of the MMF based on CAM5 (SP-CAM)
- Comparison between models and CALIPSO satellite retrievals reveals persistent underestimate of liquid-containing clouds
- SP-CAM produces *more total cloud*, but does so by producing *excessive ice cloud*; this points to deficiencies with ice microphysics, rather than model resolution