SAND2019-0841C

The Super-Parameterized E3SM

Ben Hillman, Walter Hannah, Chris Jones, Matt Norman, Sarat Sreepathi, Mark Taylor, Dave Bader, Ruby Leung

EXASCALE COMPUTING PROJECT

U.S. DEPARTMENT OF ENERGY

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.

The problem

- Goal to develop capability to assess regional impacts of climate change on the water cycle that directly affect the US economy (e.g., agriculture and energy production)
- Cloud-resolving scales desired to reduce systematic errors in climate simulations due to structural uncertainty in numerical treatments of convection
- Achieving cloud resolving model resolutions for climate-length timescales requires Zettascale computational resources; the desire for grid refinement is outpacing computational resources!

Our solution (project goals)

- Develop a “super-parameterized” earth system model with computational efficiency enabling multi-decadal climate experiments
- Efficient use of next-generation computational resources: leveraging GPU acceleration to maintain target of 5 SYPD while improving physical realism by capturing *some* aspects of a cloud resolving model
- The multi-scale modeling framework/super-parameterization approach is ideal for this task

The Multi-scale Modeling Framework (Super-Parameterization)

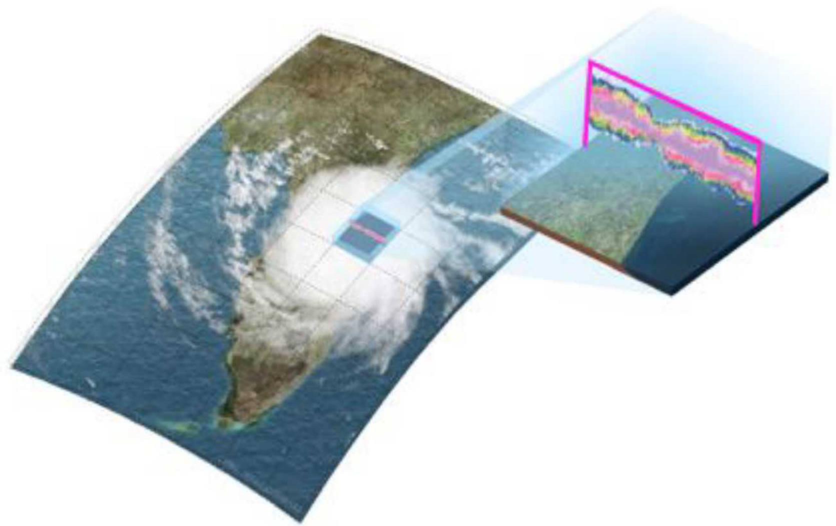

- Addresses structural uncertainty in cloud processes by replacing traditional parameterizations with explicit cloud resolving modeling within each grid cell of global climate model
- Dramatically increases arithmetic intensity, making the MMF approach one of the few ways to achieve exascale performance on upcoming architectures
- Exascale + MMF approach will make it possible for the first time to perform climate simulation campaigns with some aspects of cloud resolving resolutions

Model formulation/coupling schematic

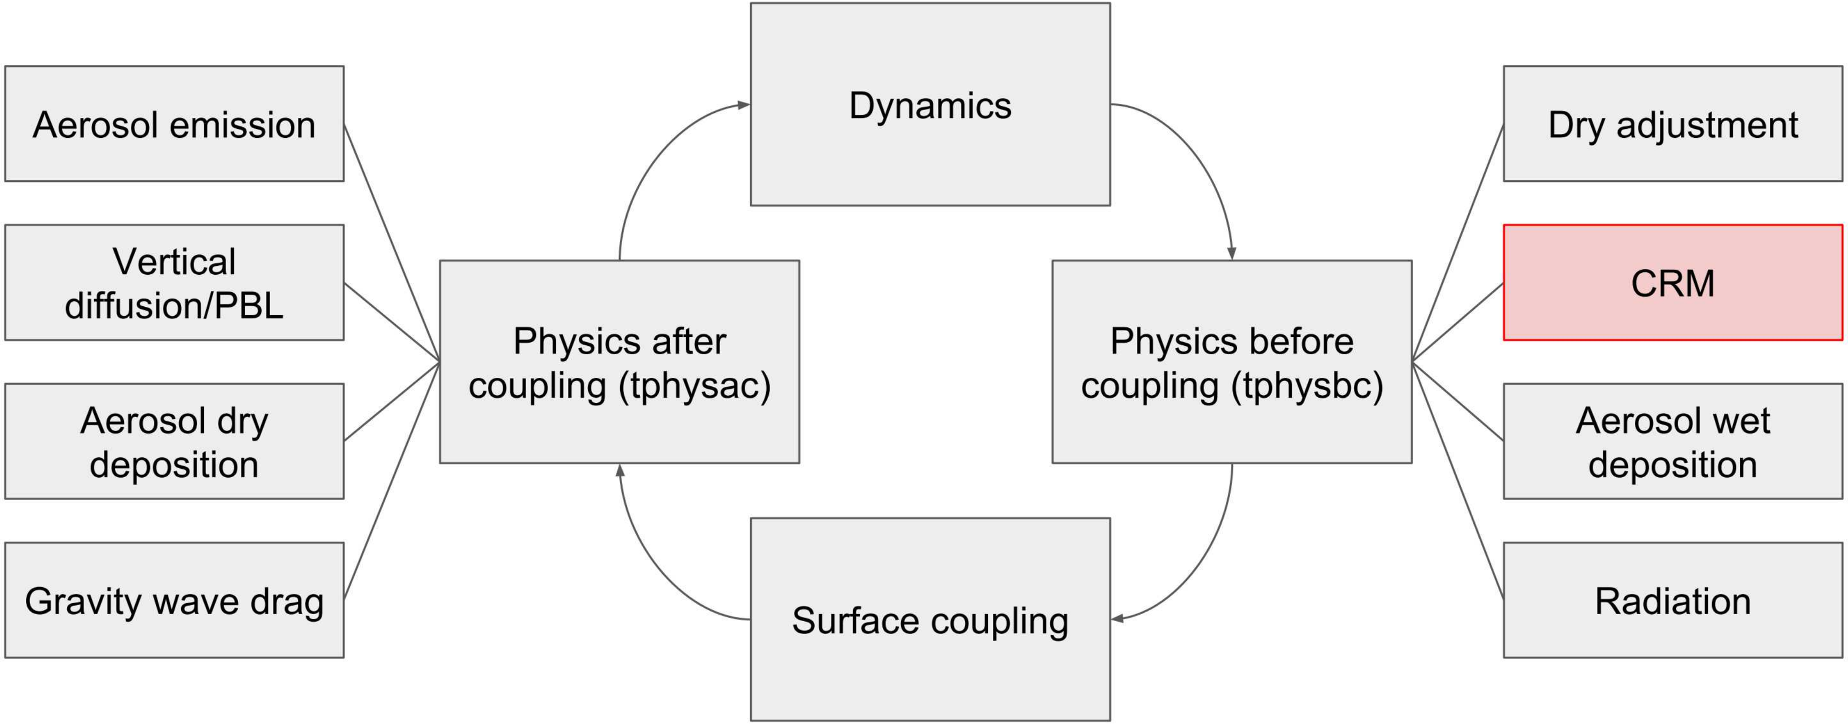

Model overview

Based on E3SM v1 (forked from master, with periodic merges from E3SM master to keep pace with development)

Similar to older SP-CAM formulations, but with some key differences:

- Coupling with spectral element dynamical core
- More vertical levels and higher model top
- New radiation scheme (RRTMGP)
- Explicit scalar momentum transport
- Pseudo-random CRM orientation

	SP-E3SM	SP-CAM (2-mom)	SP-CAM (1-mom)
Dynamical core	SE	FV	FV
Micro	sam1mom or m2005	m2005	sam1mom
CRM/GCM levels	58/72	28/30	24/26
Model top	0.1 hPa	3.5 hPa	3.5 hPa
Radiation	RRTMGP	RRTMG	CAMRT
Aerosol	MAM4	MAM3	BAM

GPU Acceleration

- Super-Parameterization is attractive for GPU acceleration because CRMs do not communicate directly through MPI; no communication slowdown at scale
- CRM code ported to offload work to GPUs using directives-based approach: flexible, maintain single codebase for both GPU and CPU-only configurations, use either PGI or XL compilers (OpenACC or OpenMP)
- Up to ~17x speedup possible using GPU, but acceleration is limited by work per node; smaller workload to achieve throughput constraint limits benefit of using accelerators (~4-10x speedup estimated on Summit)

Performance

- SP-E3SM vs E3SM: ~60x slower (0.1 SYPD vs 6 SYPD on Titan)
- SP-E3SM vs estimated “conventional” E3SM with cloud resolving atmosphere (3 km) and eddy resolving ocean: ~20x algorithmic speedup (0.1 SYPD vs 0.005 SYPD)
  - E3SM 3 km figure estimated based on current high-resolution (25 km atmosphere ,18-6 km ocean) throughput
  - We expect ~4-10x speedup on Summit with GPU acceleration (~1 SYPD)

Early success

Mesoscale convective systems in SP-E3SM

- 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

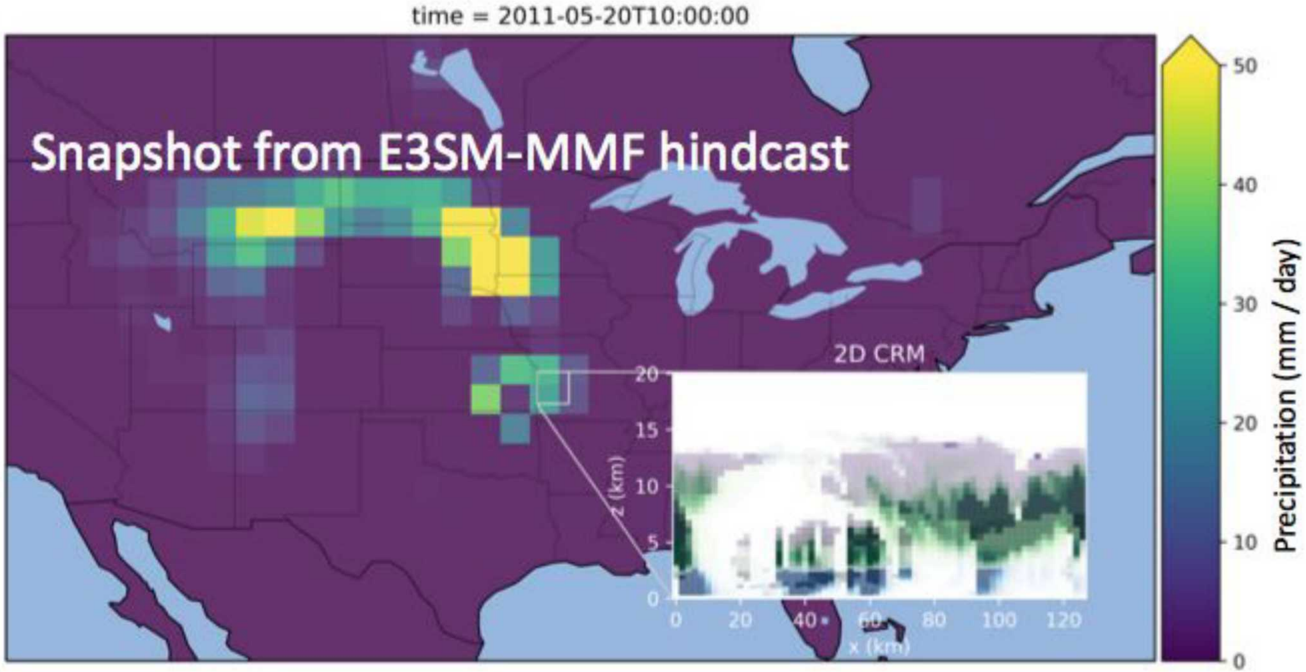

Mesoscale convective systems in SP-E3SM

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

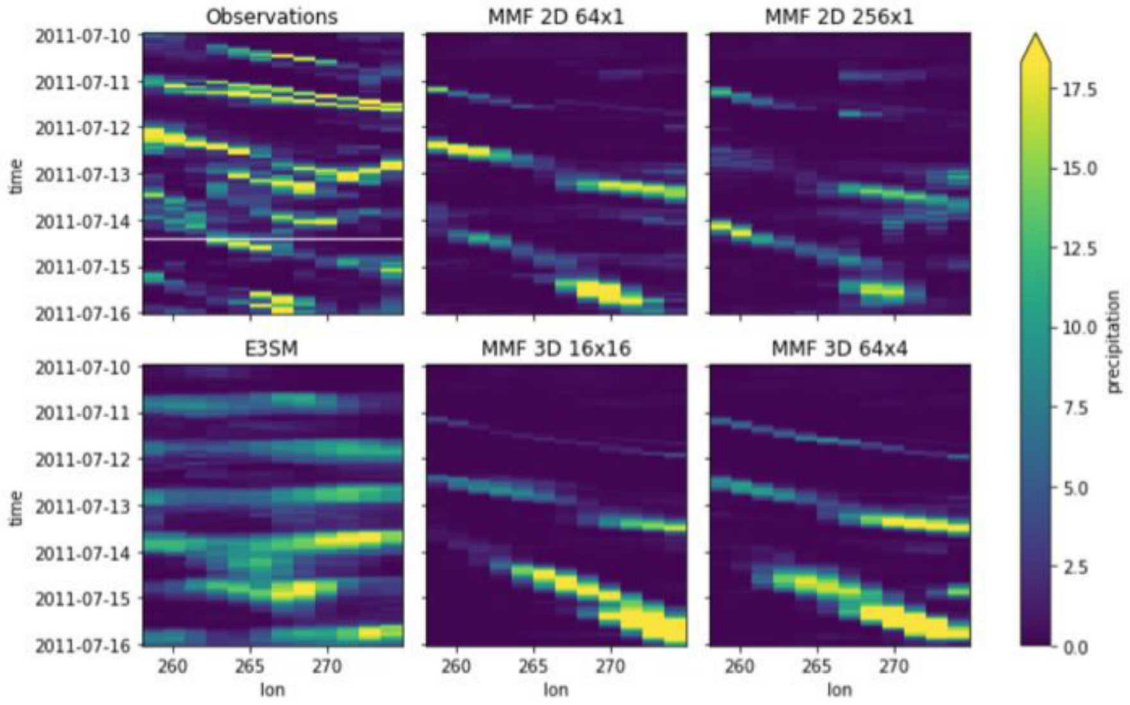

Hovmoller diagrams of precipitation for hindcast initialized on July 10, 2011

Challenges

Grid imprinting

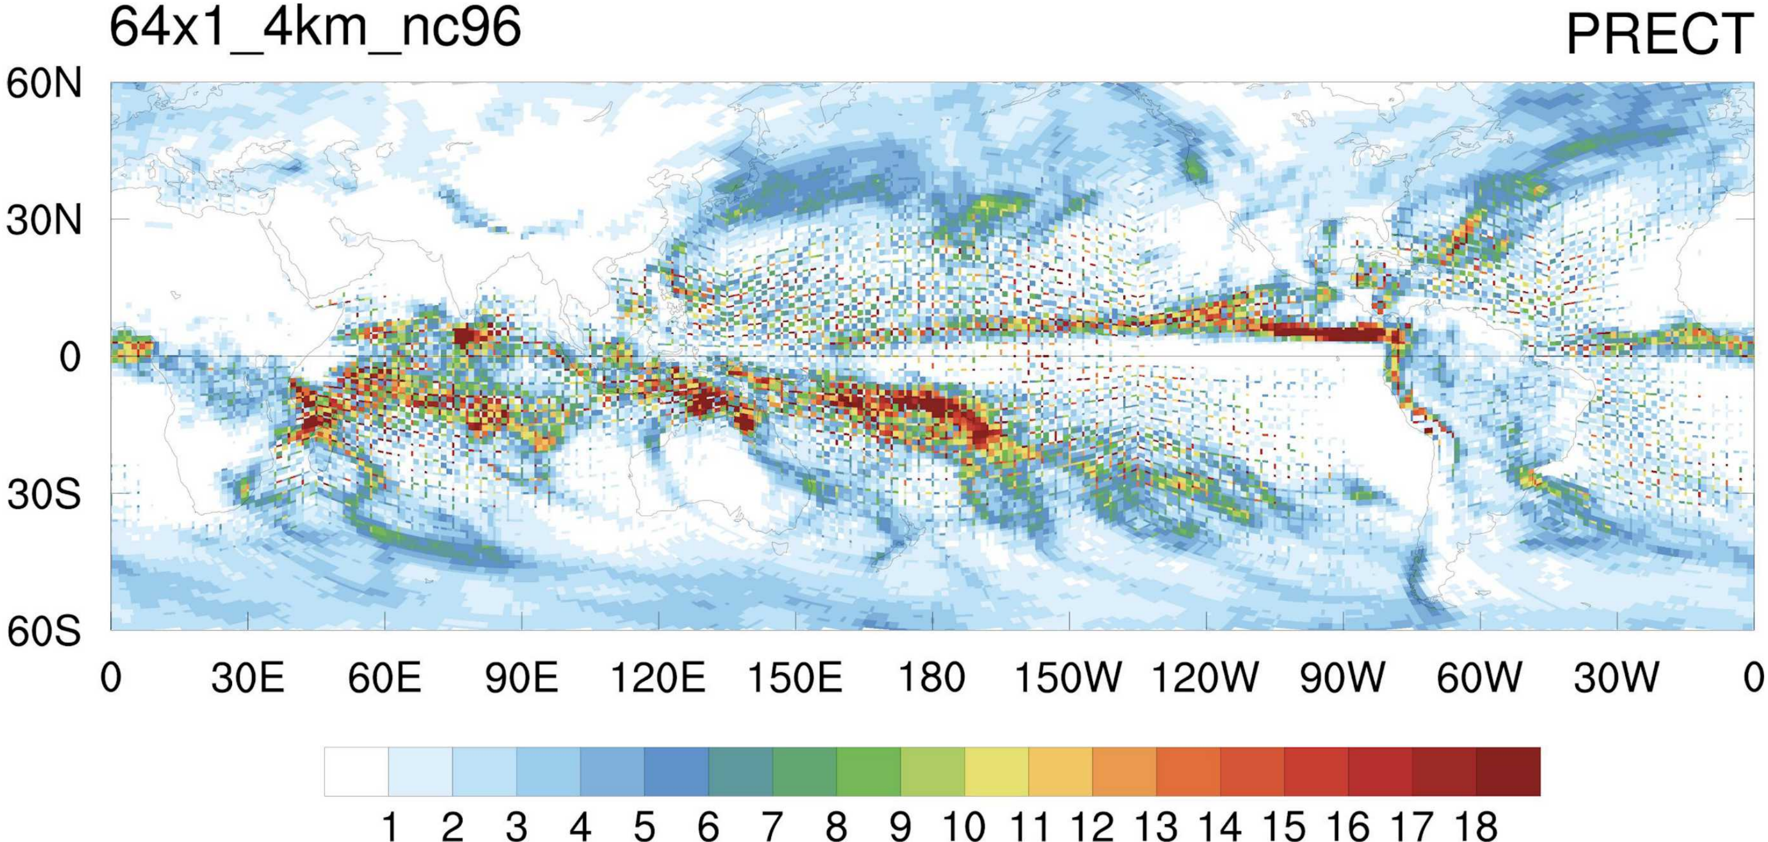

Grid imprinting

Enhanced precipitation on corner and edge GLL nodes

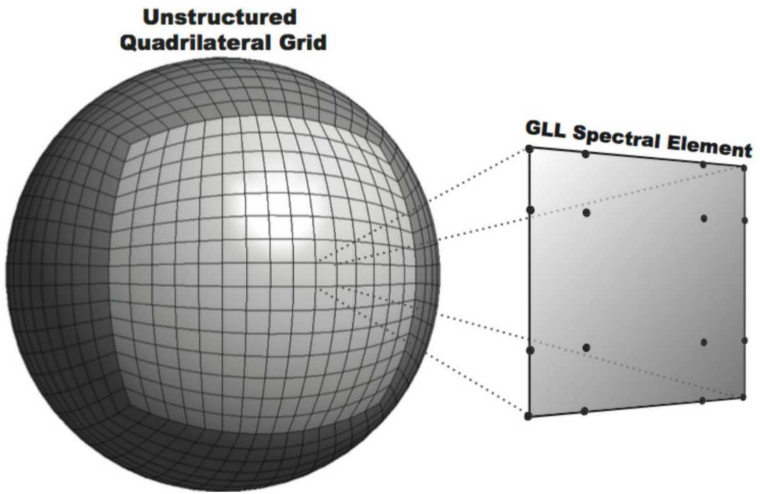

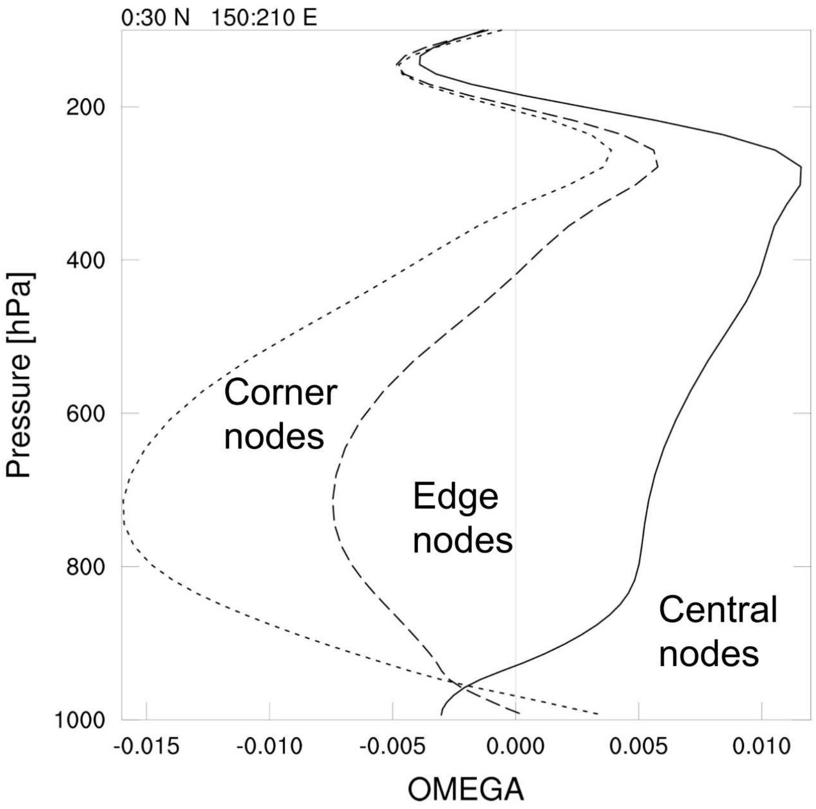

Not evident in fields that are not directly related to convection, but shows up strongly in large-scale vertical velocity

Problematic precipitation

- Distribution of daily rain rates (year 2000, 60S to 60N)
- SP-E3SM has more extremes at both ends: more heavy rain events, and more light drizzle rain events

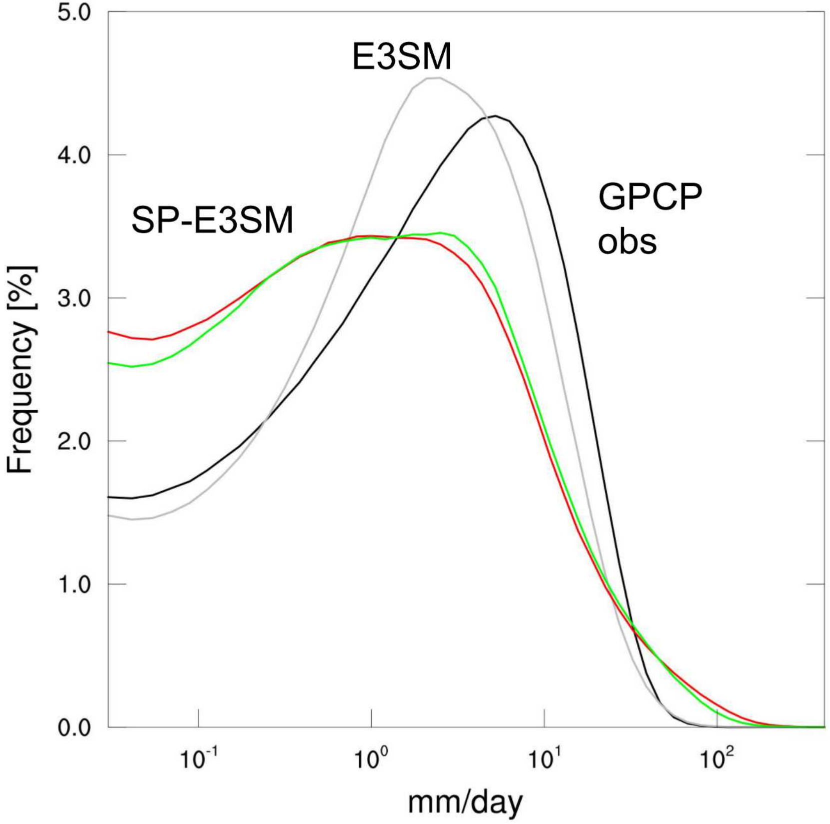

Dry bias

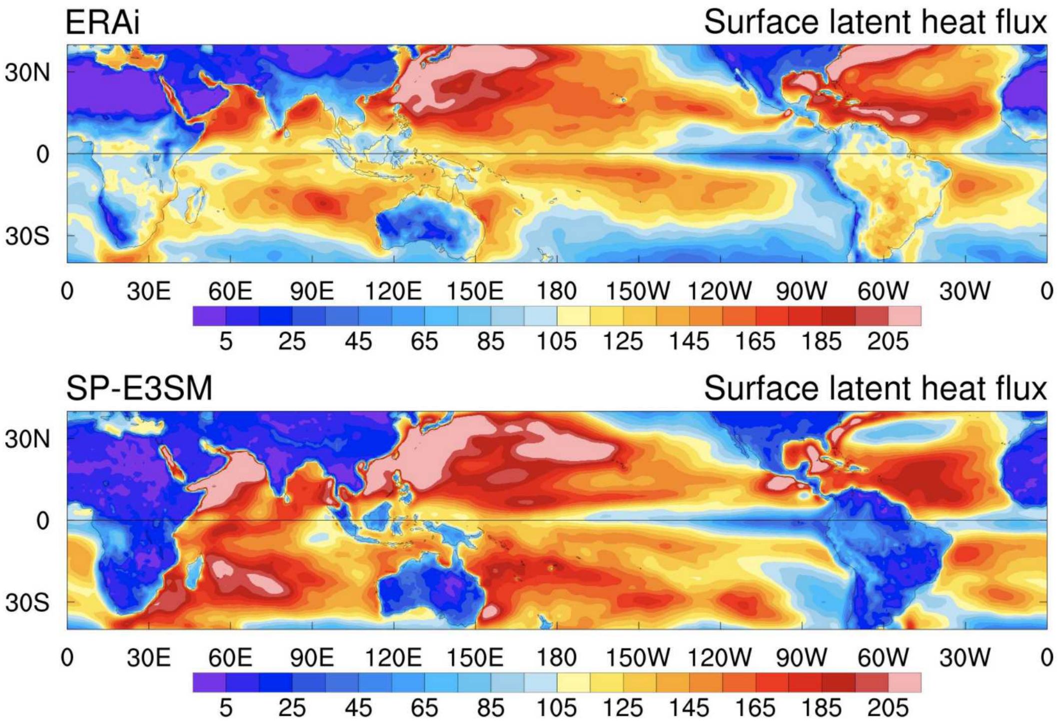
ERAi
Surface latent heat flux
30N
0
30S
0 30E 60E 90E 120E 150E 180 150W 120W 90W 60W 30W 0
5 25 45 65 85 105 125 145 165 185 205
SP-E3SM
Surface latent heat flux
30N
0
30S
0 30E 60E 90E 120E 150E 180 150W 120W 90W 60W 30W 0
5 25 45 65 85 105 125 145 165 185 205

Dry bias

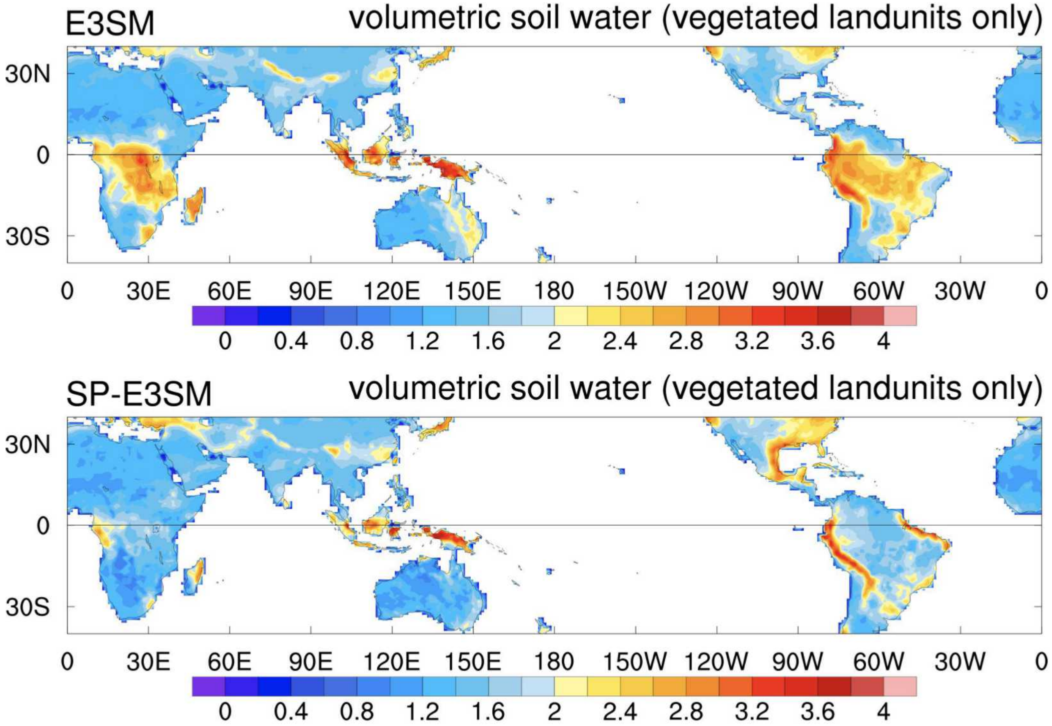
E3SM
volumetric soil water (vegetated landunits only)
30N
0
30S
0 30E 60E 90E 120E 150E 180 150W 120W 90W 60W 30W 0
0 0.4 0.8 1.2 1.6 2 2.4 2.8 3.2 3.6 4
SP-E3SM
volumetric soil water (vegetated landunits only)
30N
0
30S
0 30E 60E 90E 120E 150E 180 150W 120W 90W 60W 30W 0
0 0.4 0.8 1.2 1.6 2 2.4 2.8 3.2 3.6 4

Take-aways

Successes:

- Ideal approach to leverage next generation HPC resources while improving physical realism
- Propagating MCS features are well-captured
- Ripe opportunity to explore coupling between small and large-scale processes

Remaining questions and biases:

- Coupling with SE dycore leads to grid imprinting in precipitation field
- Precipitation statistics and evaporation bias

Bonus

Radiative transfer in SP-E3SM

- RRTMG replaced with RTE/RRTMGP: optimized for parallel applications
- (show cost of radiation in relation to rest of the model)
- Primary goal: utilize GPU port of radiation kernels
- Bonuses: radiation code is more flexible, portable, and testable than RRTMG
- Available for use in standard E3SM already (add “-rad rrtmgp” to CAM_CONFIG_OPTS)

Cloud-radiative coupling

- Radiative heating is updated every GCM (physics) timestep
- Calculated through the GCM interface to take advantage of updated aerosol properties
- Radiative fluxes and heating rates calculated using *cloud-scale*, but *time-averaged* cloud properties
- Cloud scale heating rate applied at *next* CRM call (next physics timestep)

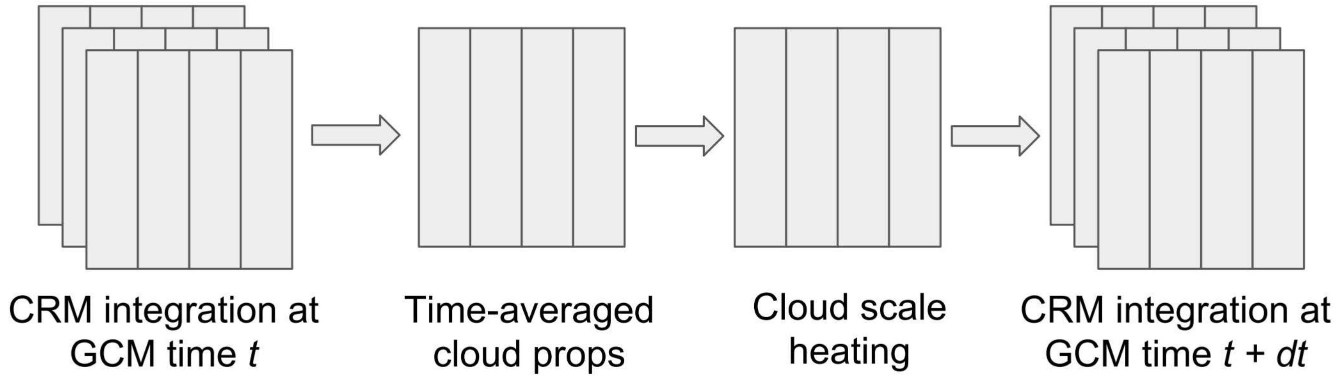

Reduced radiation

- Reduce cost of radiation by calculating heating on *groups* of columns
- Average cloud properties over group of columns, calculate heating from average, apply back to same group of CRM columns
- Maintain *some* heterogeneity of cloud properties, but at reduced cost (versus doing radiation on full CRM resolution)

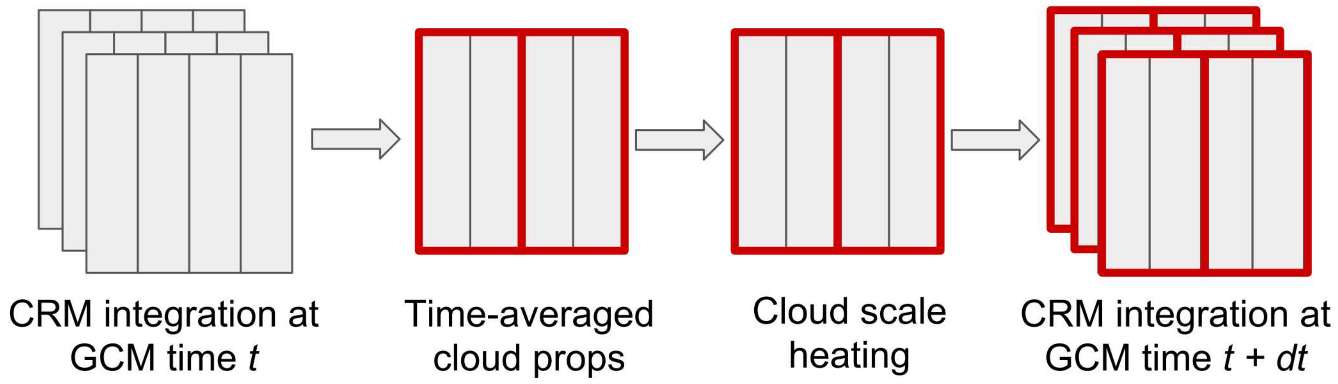

Reduced radiation

- Shortwave and longwave differences compensate
- Little effect on net TOA flux
- Can we use some version of this to reduce cost, while maintaining accuracy?

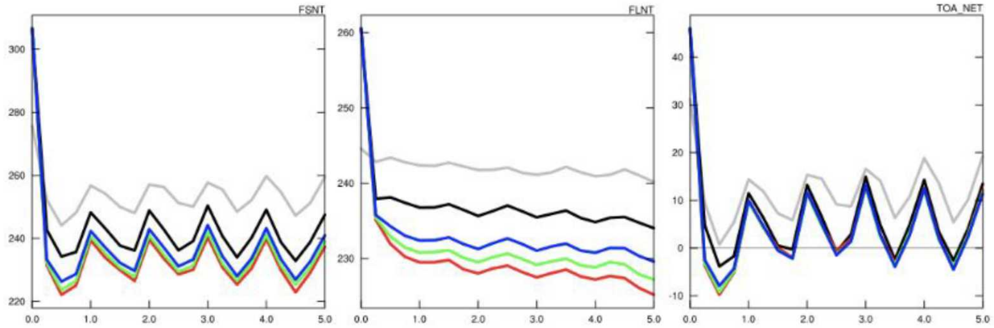

- gray = non-SP (ZM)
- black = SP1 crm_nx=32, crm_nx_rad=32
- blue = SP1 crm_nx=32, crm_nx_rad=4
- green = SP1 crm_nx=32, crm_nx_rad=2
- red = SP1 crm_nx=32, crm_nx_rad=1

No McICA, nxrad = 64 (-50.677)

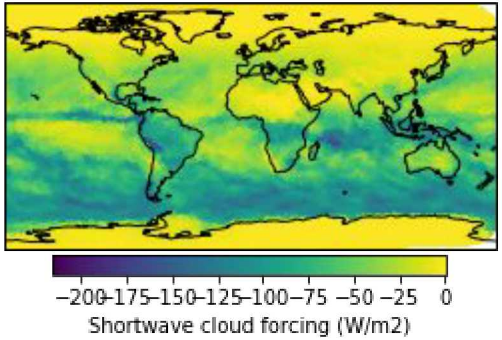

No McICA, nxrad = 1 (-59.288)

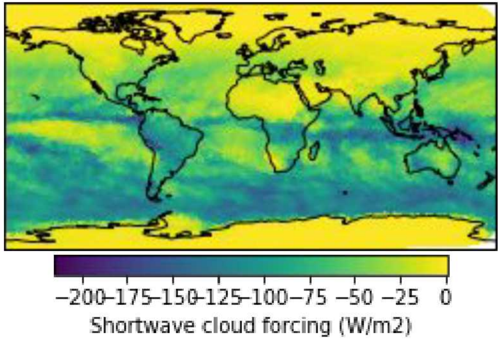

Difference (-8.611)

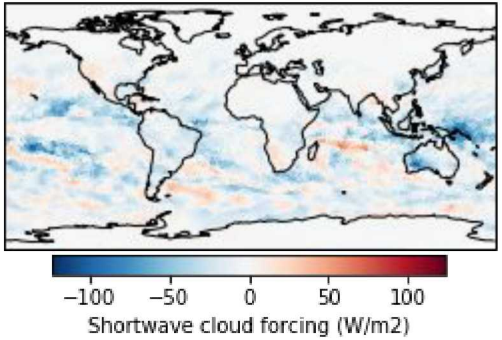

Early climatology

- Climatology of precipitation looks reasonable
- Familiar biases
- Large bias over Amazon...

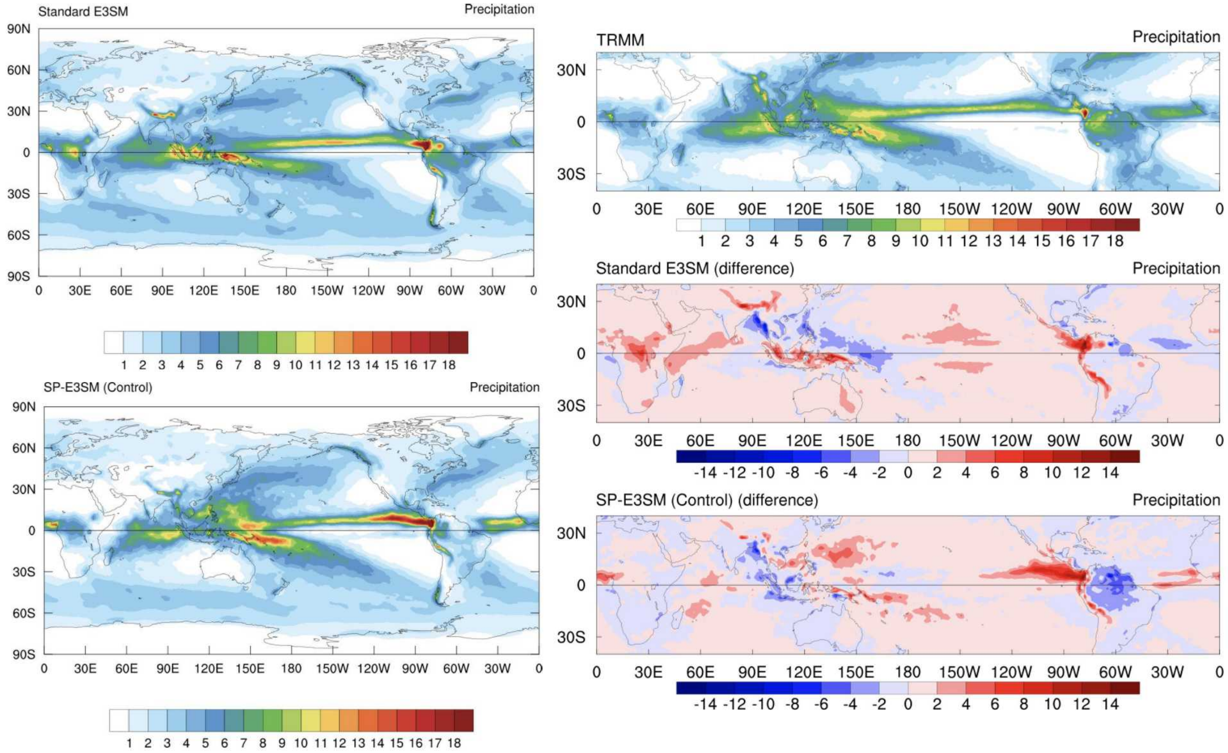

Precip distribution improved with larger CRM

- Distribution of daily rain rates (year 2000, 60S to 60N)
- SP-E3SM has more extremes at both ends: more heavy rain events, and more light drizzle rain events

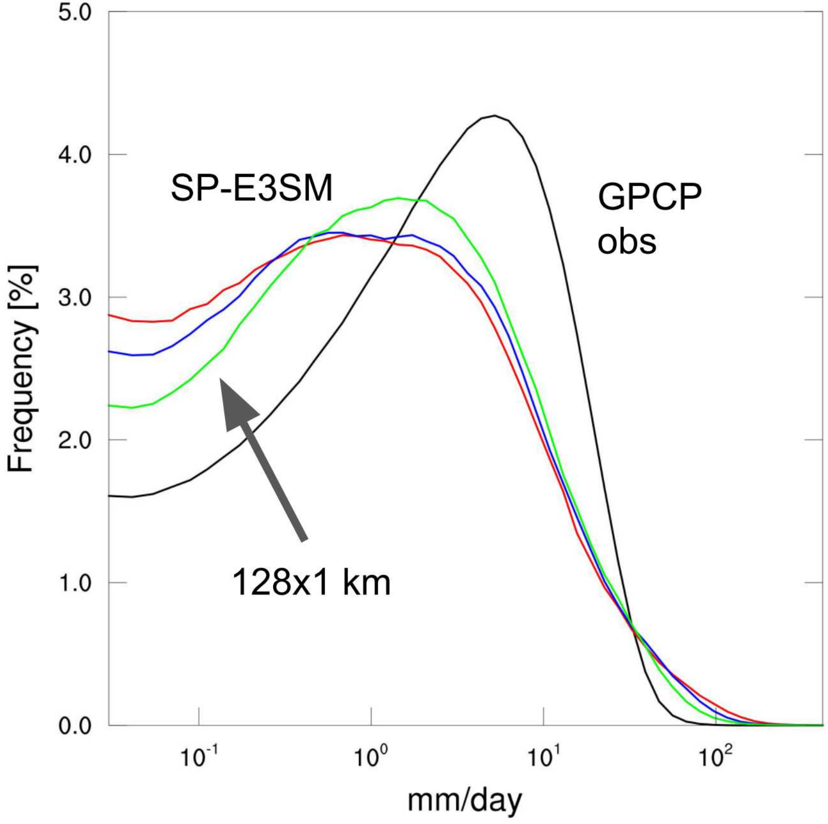

Grid imprinting reduced with smaller dx

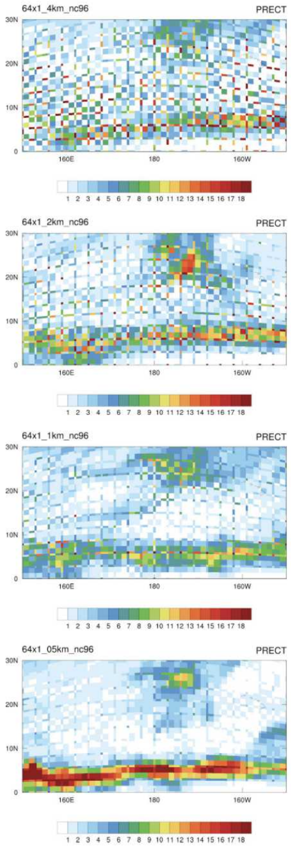

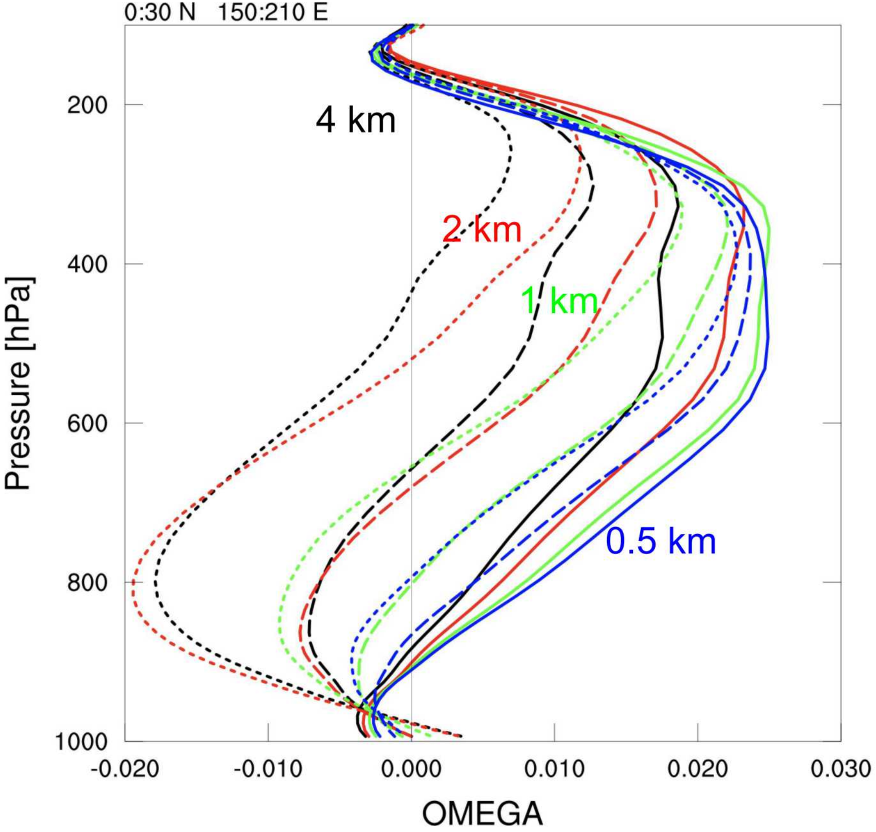

Dry bias

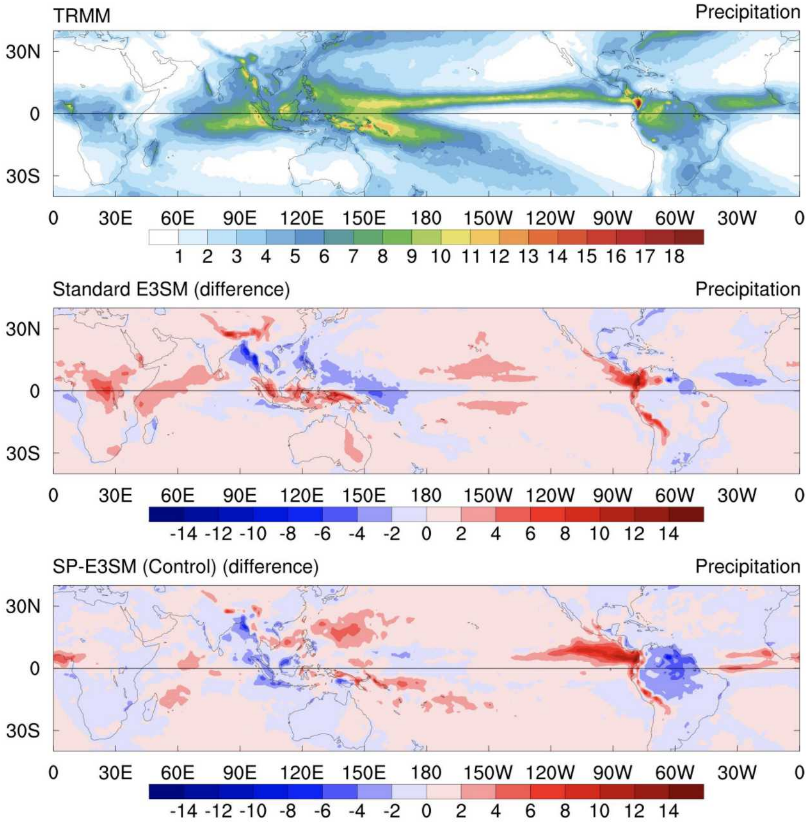
TRMM
Precipitation
30N
0
30S
0 30E 60E 90E 120E 150E 180 150W 120W 90W 60W 30W 0
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18
Standard E3SM (difference)
Precipitation
30N
0
30S
0 30E 60E 90E 120E 150E 180 150W 120W 90W 60W 30W 0
-14 -12 -10 -8 -6 -4 -2 0 2 4 6 8 10 12 14
SP-E3SM (Control) (difference)
Precipitation
30N
0
30S
0 30E 60E 90E 120E 150E 180 150W 120W 90W 60W 30W 0
-14 -12 -10 -8 -6 -4 -2 0 2 4 6 8 10 12 14