E3SM
Energy Exascale Earth System Model

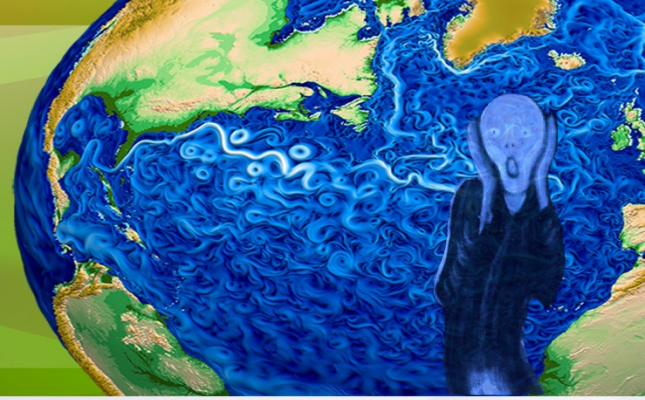

Next-Generation Developments in Software and Algorithms

November 9, 2020

Andy Salinger (Sub-Project Lead)

Hyun Kang, Andrew Bradley, Andrew Steyer, Cassidy Krause

Karen Devine, Iulian Grindeanu, Vijay Mahadevan

Salil Mahajan, Michael Kelleher, Rong-You Chien, Josh Fu

Benj Wagman, Kenny Chowdhury

U.S. DEPARTMENT OF ENERGY

How the NGD Software and Algorithms sub-project Fits in the Larger E3SM Ecosystem

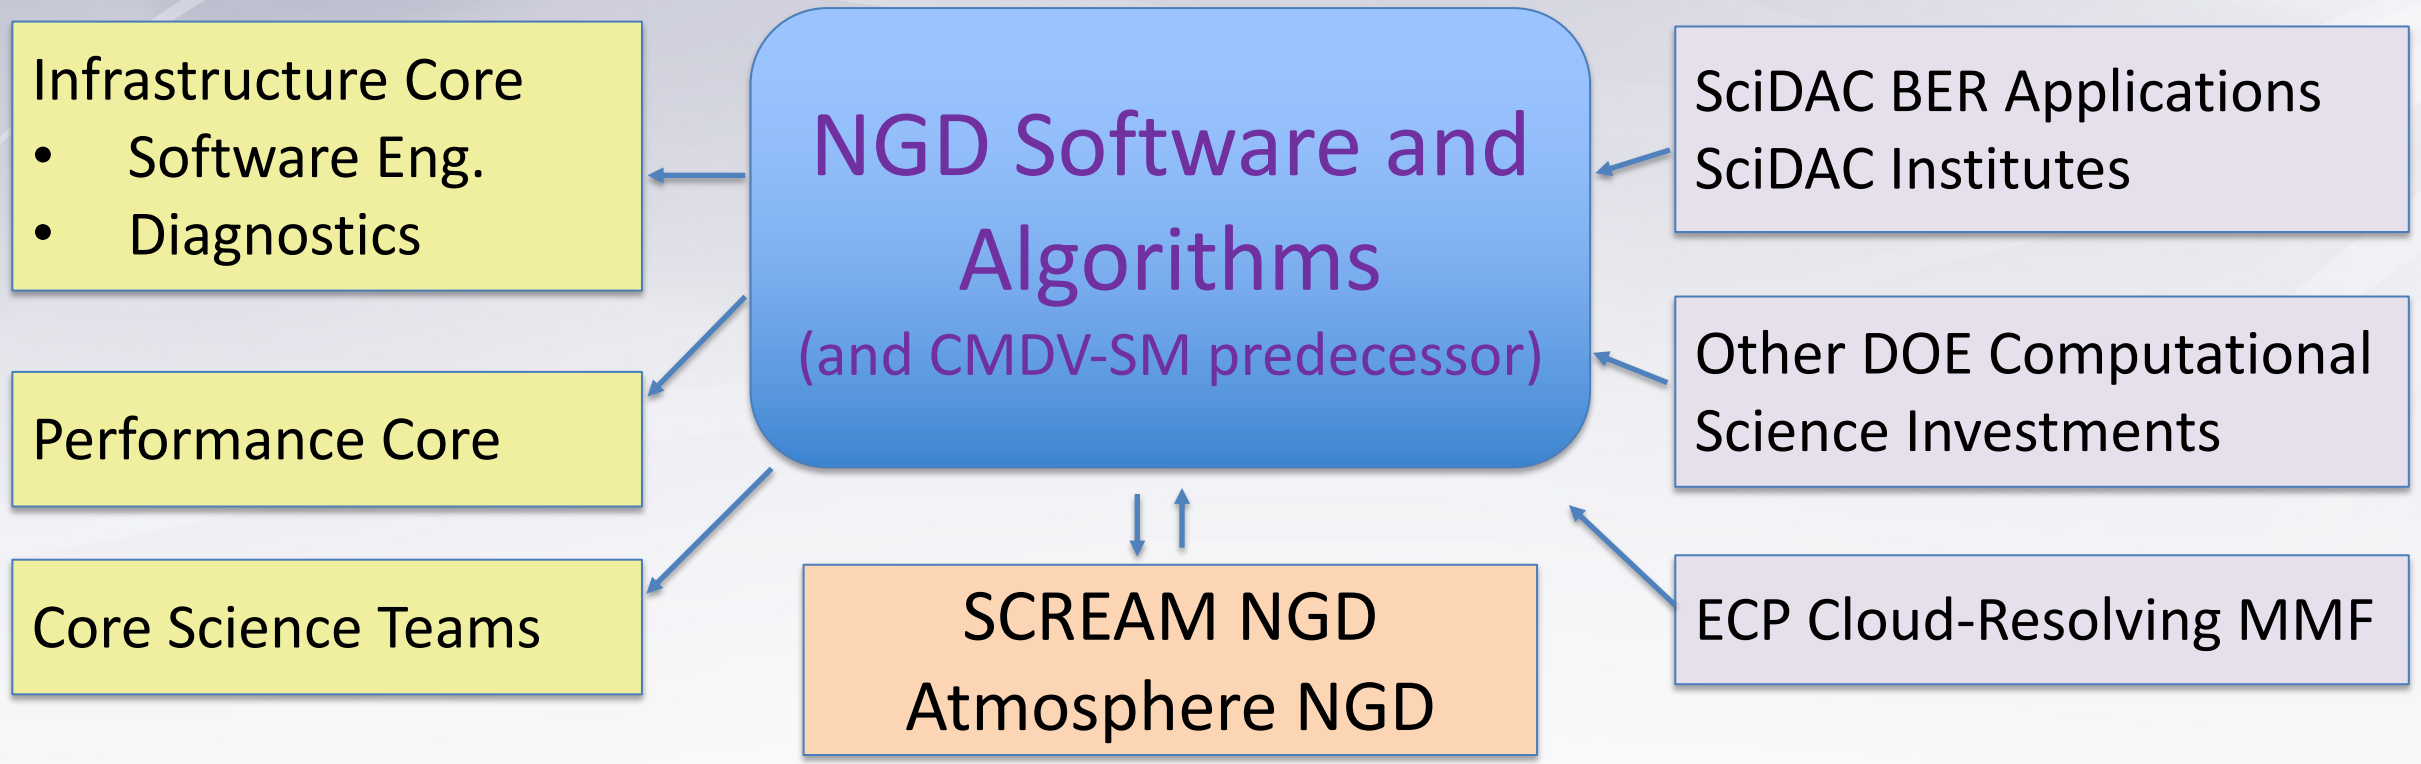

Incorporation into E3SM:

~E3SM v2 | ~E3SM v3 | ~E3SM v4

NGD Software and Algorithms Vision

Make targeted investments for E3SM v3-v4 to:

1. Improve model throughput on Exascale machines
   - Algorithms, Performance
   - *(Note: GPUs covered in Performance and Nonhydrostatic Atm NGD talks)*
2. Create a culture of Verification as we do use-inspired science and move towards decision support missions.
3. Leverage DOE Computational Science investments
   - Entrain talented computational scientists into E3SM

Algorithmic Improvements for Exascale

Many Algorithmic improvements focus on decreasing the communication between processors by increasing the time step size.

- Semi-Implicit Barotropic Solve in the Ocean
- Semi-Lagrangian Transport with upwind-MPI in the Atmosphere
- Implicit-Explicit (IMEX) time integrators for nonhydrostatic atmosphere

Decreasing communication becomes MORE important on exascale architectures

Communication Halo every time step, Δt

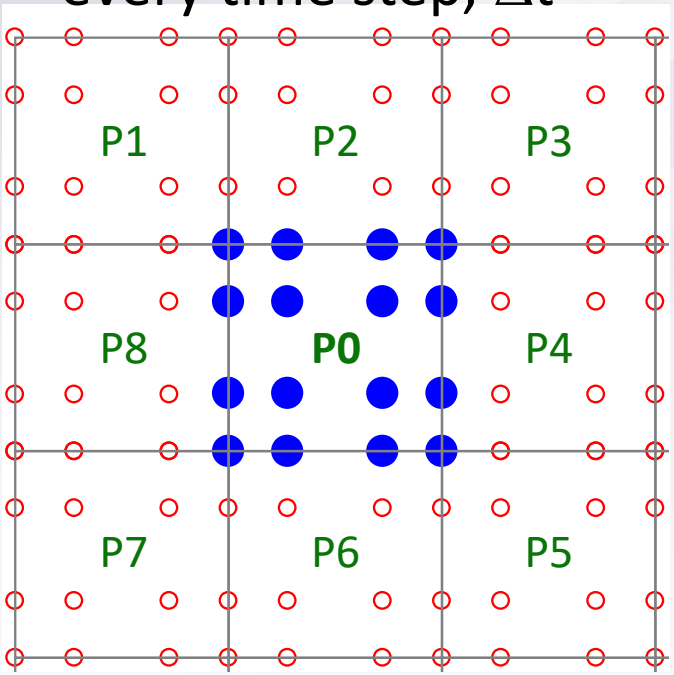

Fig: A "Halo exchange" requires communicating all red points to the processor that owns the blue points.

MPAS Barotropic Mode Solver

Idea: Take Implicit time step over the Ocean barotropic mode (surface gravity wave)

- Currently, Explicit Subcycling (ES) takes 60 explicit time step Halo Exchanges of a 2D model for every 3D time step
- The Key is effective preconditioning of the matrix solve in the Implicit method

Semi-implicit (SI) solver is working!

a. 1.74x speedup of full MPAS-Ocean dynamics

Also developing a semi-implicit solver using AMG preconditioner from ForTrilinos (ECP Software)

- Not yet as scalable
- Comparisons led to improvements in Hand-Rolled solver
- Provides quicker path to GPUs

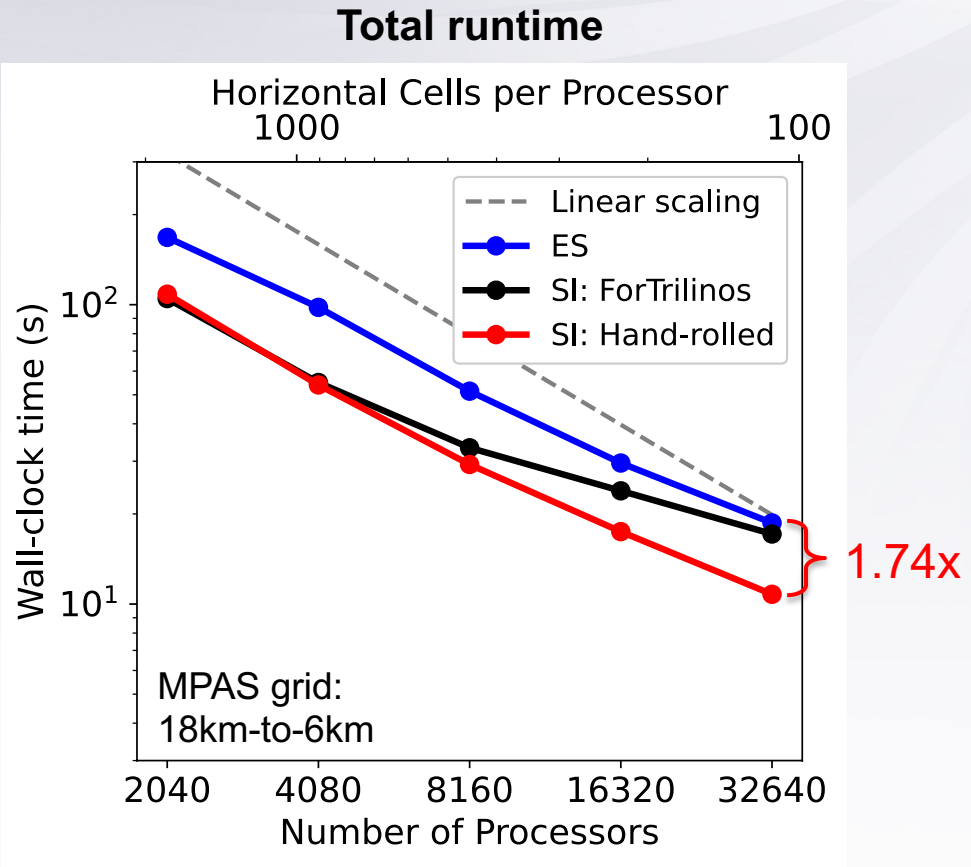

Semi-Lagrangian Tracer Transport Upwind MPI

1. Semi-Lagrangian tracer transport takes long time step (~6Δt) [COMPOSE SciDAC]

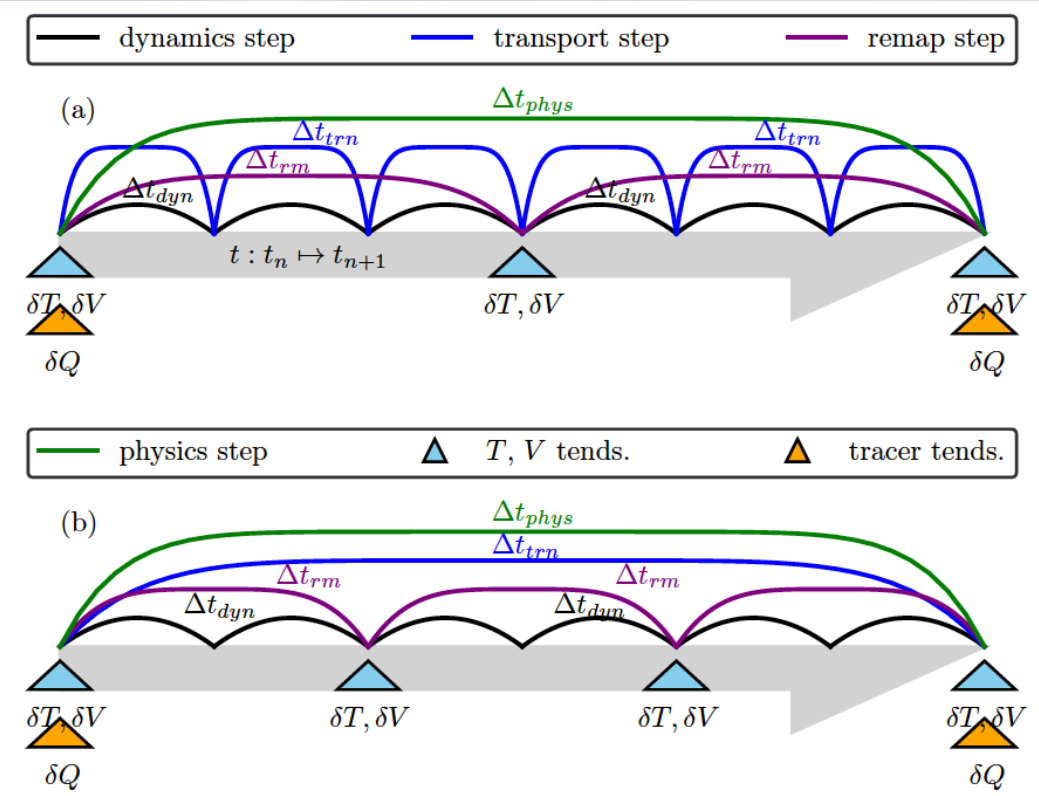

2. Extra speedup: communicate only what is needed, not full halo.

- 2x speed-up

1/6 as many halo exchanges for transport

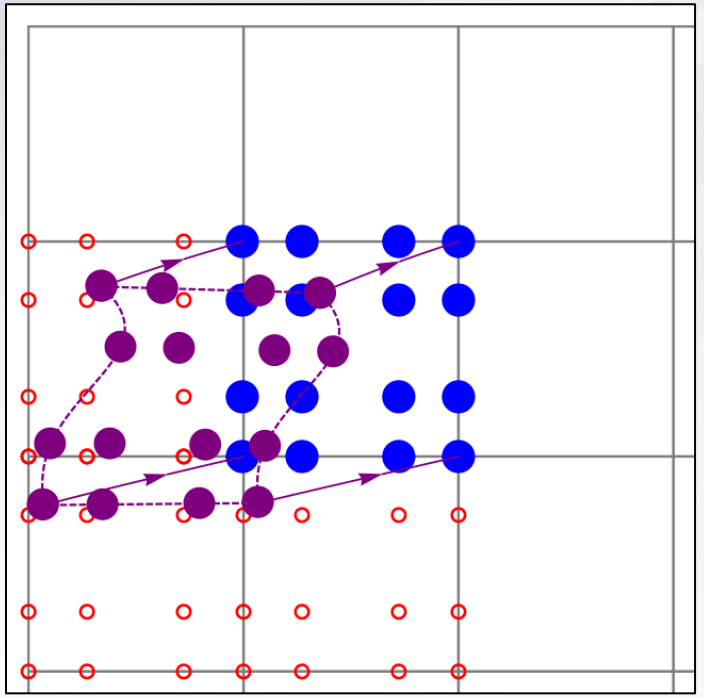

Fig: Request red-owning processes to evaluate the purple points. These then update the blue points.

Implicit-Explicit (IMEX) Time Integrators for Nonhydrostatic Atmosphere Dycore

- Nonhydrostatic atmosphere needed to capture convective instabilities at cloud-resolving resolutions
- New Nonhydrostatic atmosphere would cause 100x decrease in explicit time step size compared to hydrostatic model
  - Capture vertical motion implicitly: IMEX
- HEVI: horizontally-explicit, vertically implicit
- Real Topography was a challenge

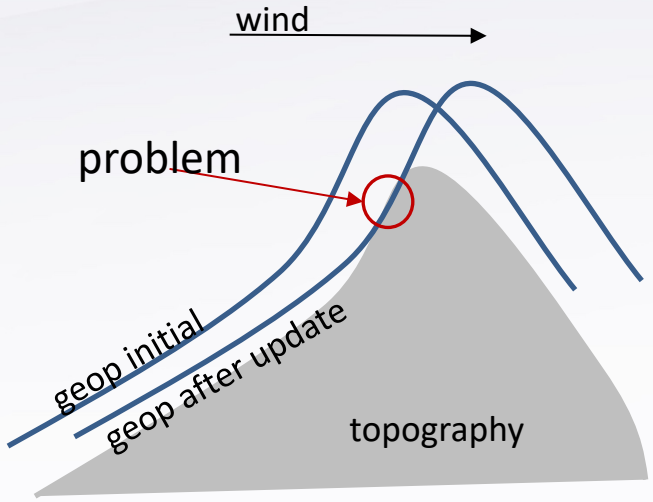

A zoo of methods developed and tested

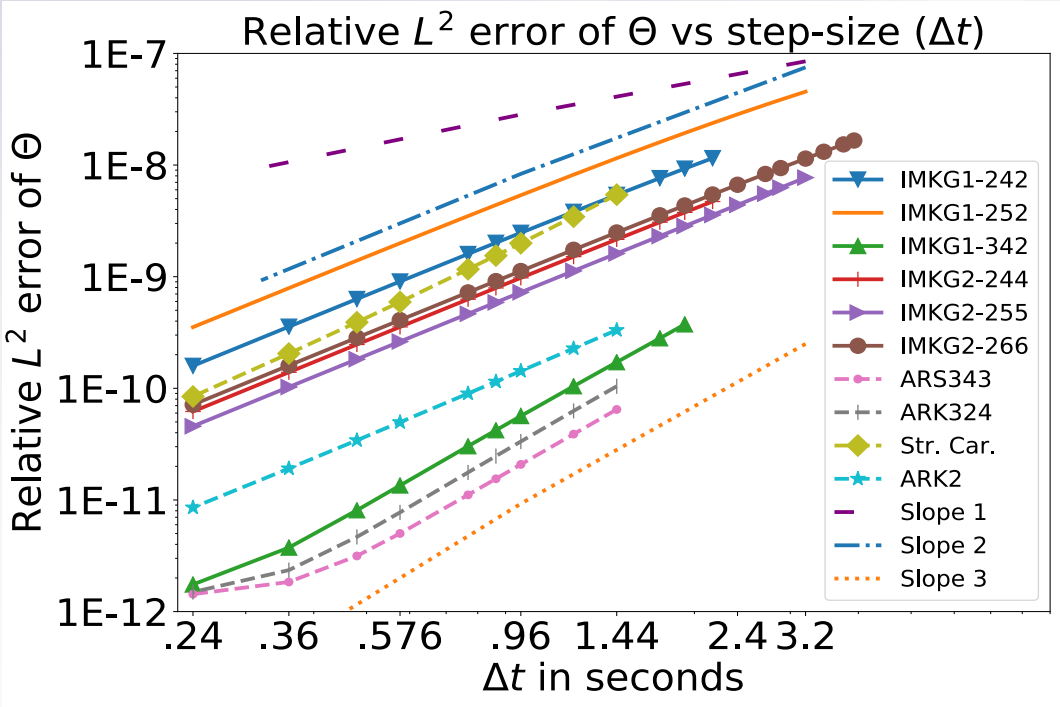

Result: Nonyhdrostatic model can now run at the same Δt as Hydrostatic model

ALGORITHMS

Delivered a 2.2x Faster and Improved E3SM Atmosphere through Algorithms R&D

Algorithms

1. "PhysGrid": Subgrid physics on independent discretization
2. Semi-Lagrangian Tracer Transport with property preservation and optimized MPI communication (6x speed-up)
3. New Implicit-Explicit time integrator for non-hydrostatic model chosen from a dozen methods

1 Validation of "PhysGrid" discretization

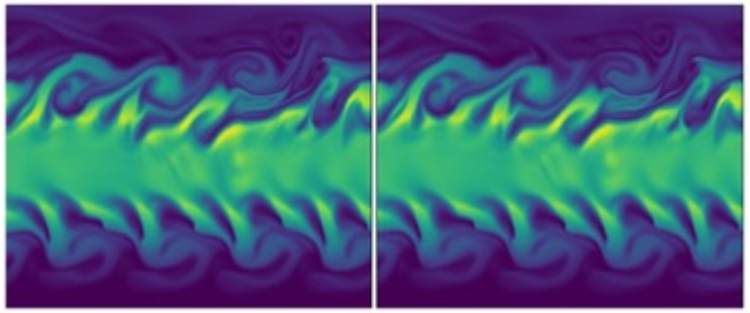

EAM performance on 68 nodes of Compy

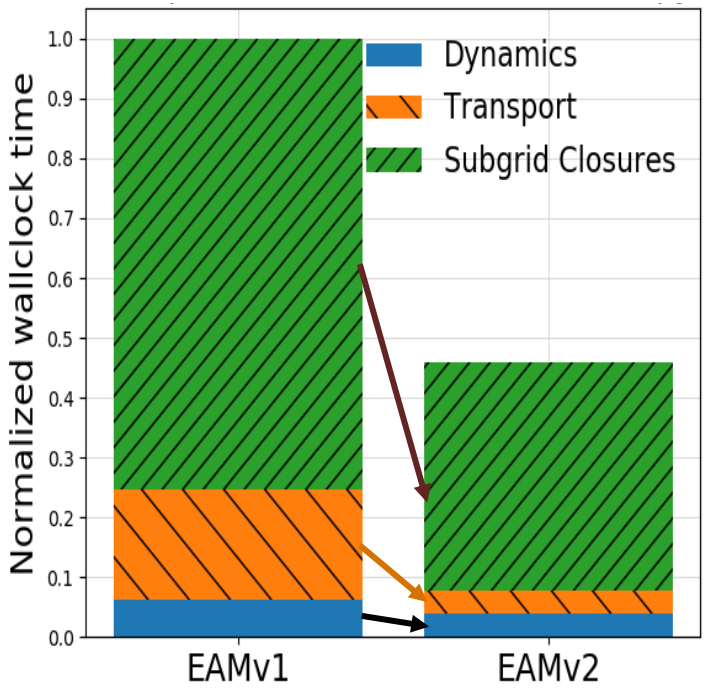

Funding Sources

- SciDAC: COMPOSE
- ECP: E3SM-MMF
- CMDV Software Modernization
- E3SM SCREAM
- E3SM NGD Software and Algorithms
- SciDAC: Nonhydrostatic Atm

2 New effort: SL Transport in Ocean Model

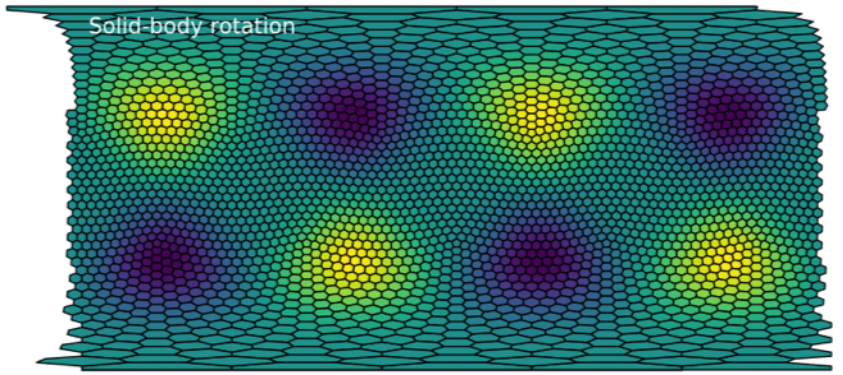

Creating a Culture of Verification

Verification: "Evidence that the equations and solution algorithms are coded correctly"

- Pre-requisite for validation, tuning, or trustable predictions

✓ Funded under CMDV-SM and NGD S&A

✓ *Baked into SCREAM, EAGLES, Ocean NGD*

Metric: Number of Log-Log Plots at ESMD PI Meeting!

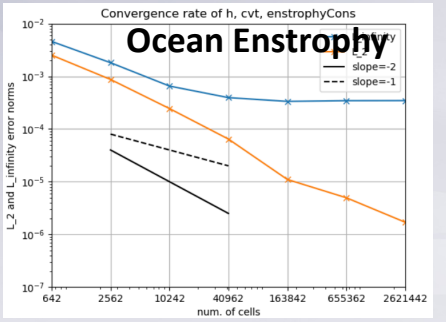

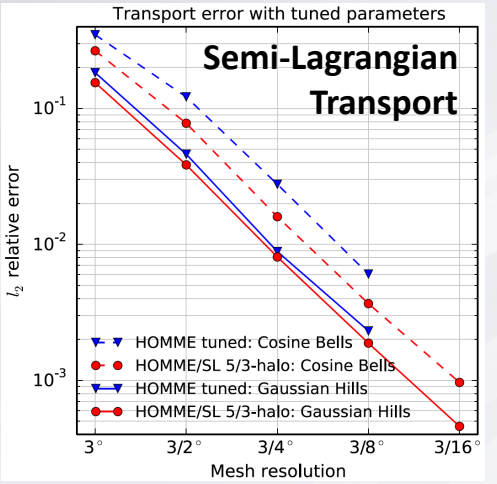

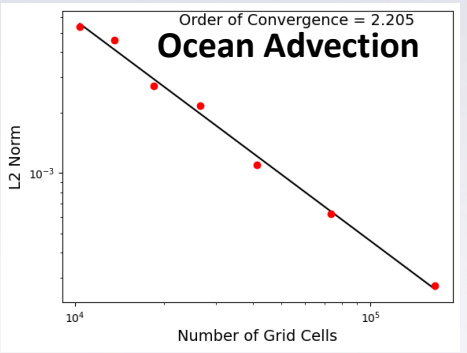

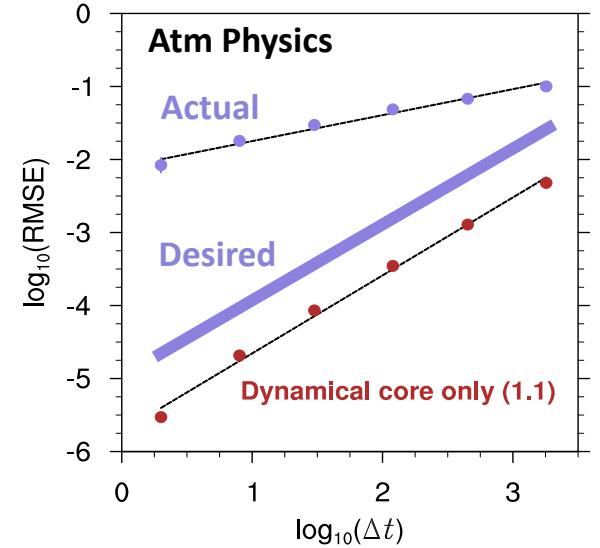

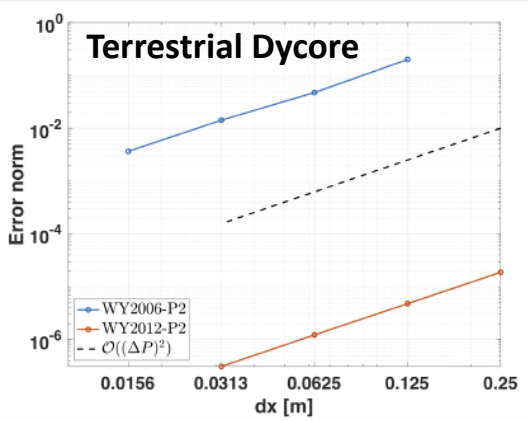

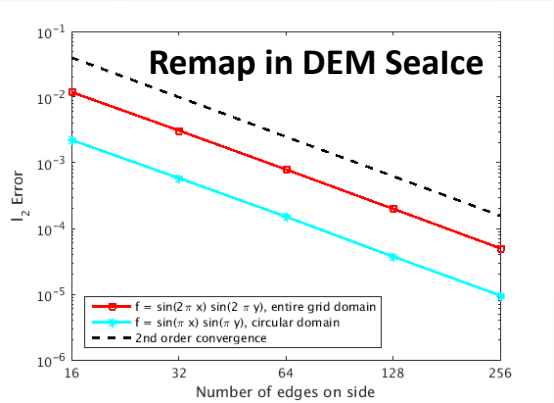

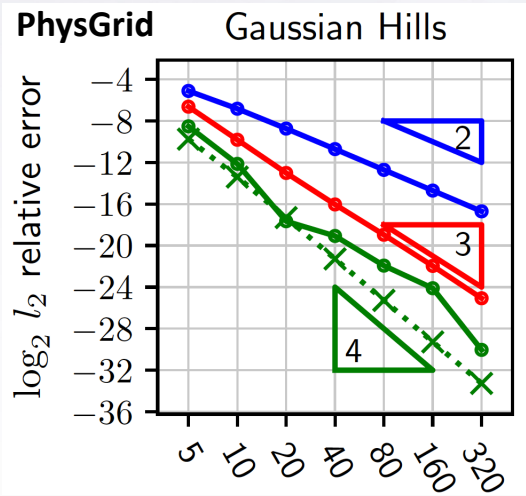

Multivariate Tests for Climate Reproducibility

1. Developed a ML based multivariate testing framework for climate reproducibility for MPAS-O using short ensembles:
2. Confirmed rigorously that we can switch computers in the middle of a science campaign

SST Difference for GM Kappa changing from 1800 to 600, shown:

Ensemble testing framework can detect change from 1800 to 1799.

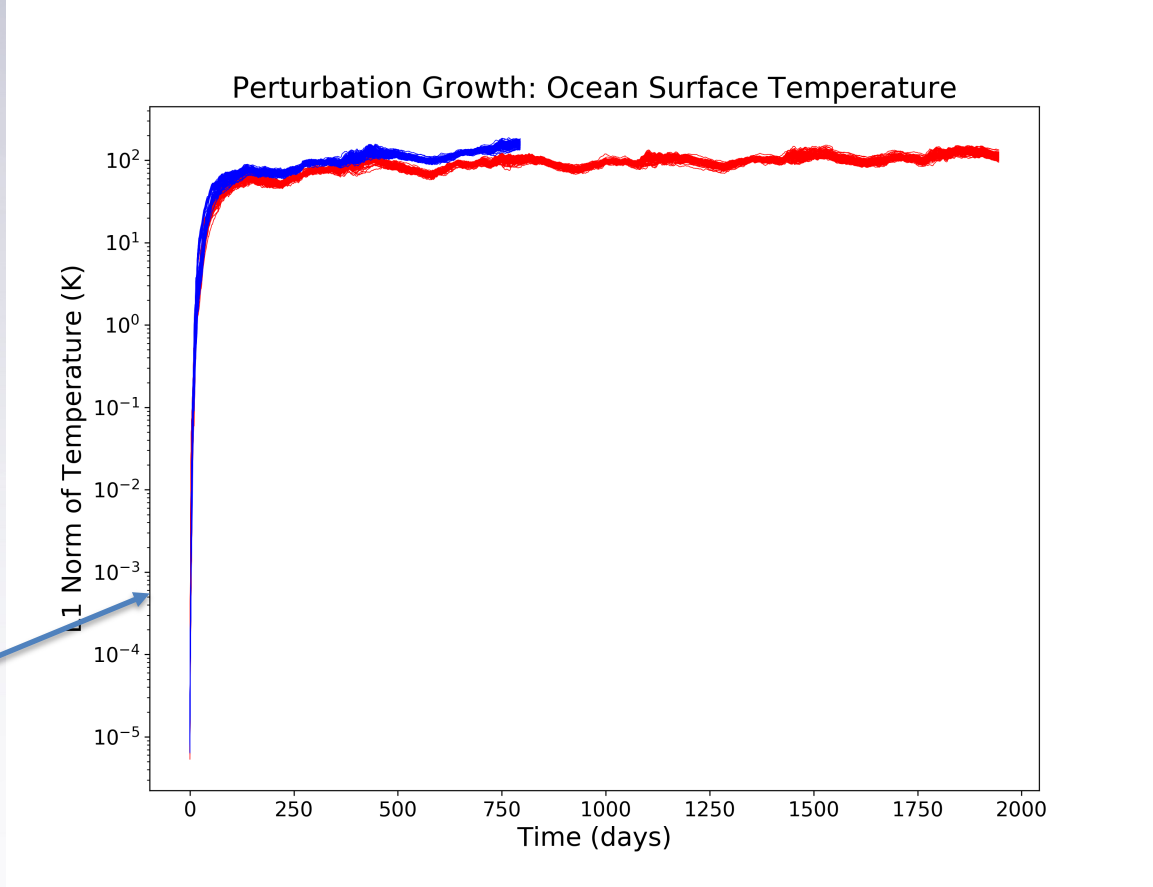

New “coupled” partitioning strategy decreases time for parallel mesh intersection workflow

Fast computation of component-mesh intersections and projection weights for solution transfer for high-fidelity cases

- Zoltan partitioner supported by SciDAC
- MOAB mesh database was developed under SciDAC, NEAMS, ASCR

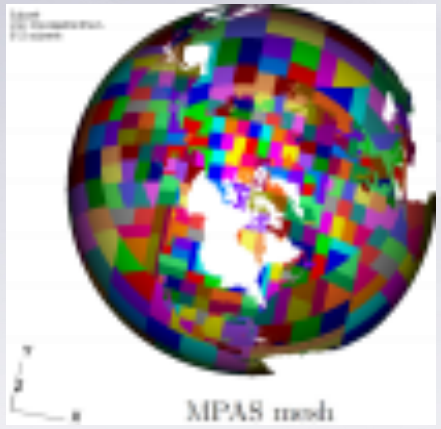

Ocean partitioned via Zoltan’s RCB (recursive coordinate bisection)

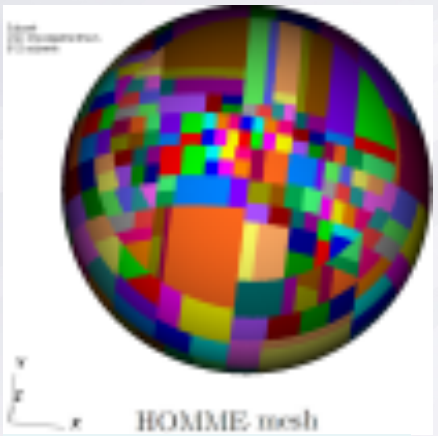

Atmosphere partition inferred from ocean’s RCB partitioning tree

NE256-PG2: OCN (3.7M) -> ATM (1.57M); 1024 cores (bebop KNL)		
Strategy	Coverage time (s)	Intersection time (s)
RCB OCN + RCB ATM	1.81	20.49
Gnomonic RCB OCN + Inferred ATM	**0.43**	**3.08**

New Task: "Optimization/UQ/ML-based Calibration"

1. Looking to automate parts of the Calibration process
   - ➢ Clear definition of Objective Function, quantifying the trade-offs
     - Increase reproducibility of the calibration process
   - ➢ Opportunities for accelerating tuning:
     - Create Reduced Order Models from previous tunings
     - Multi-Fidelity algorithms
   - ➢ Bayesian calibration – draw ensembles from different tunings
2. Staffed for FY21!
   - + E3SM-trained Climate Scientist
   - + SciDAC-trained UQ/Opt/ML expert

Questions?

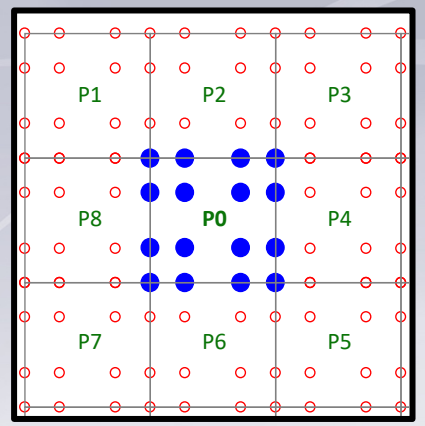

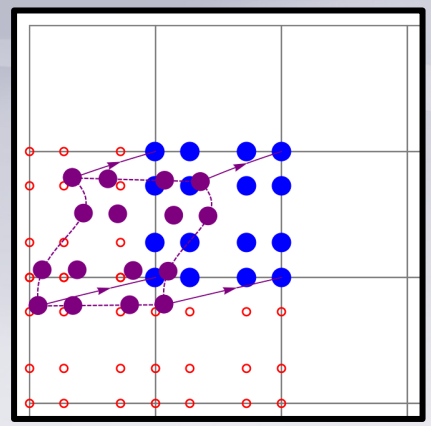

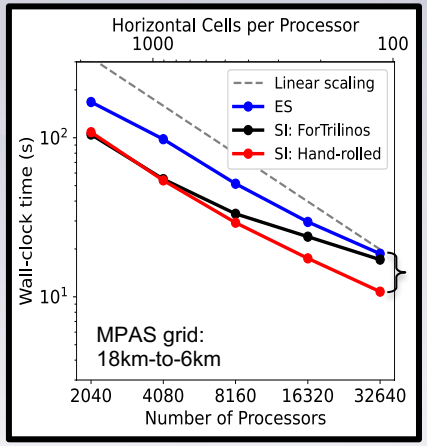

1.74x

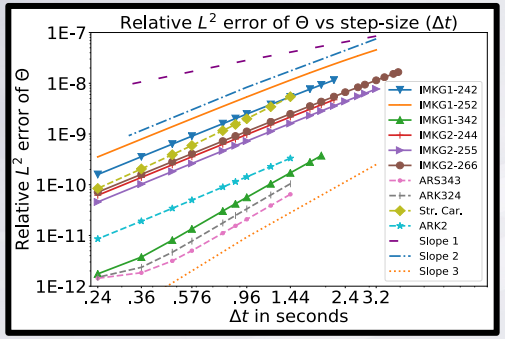

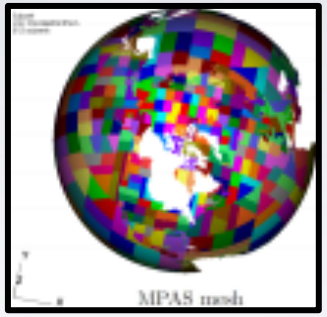

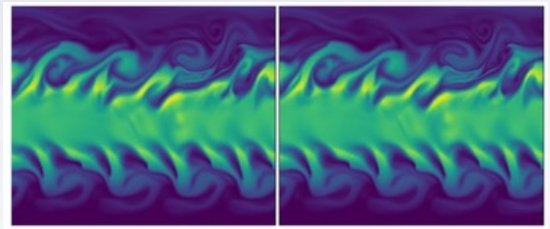

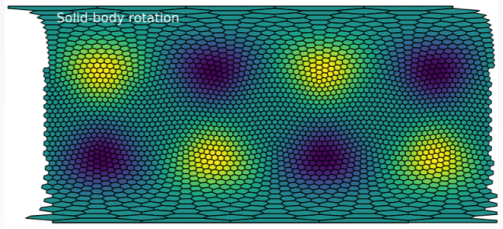

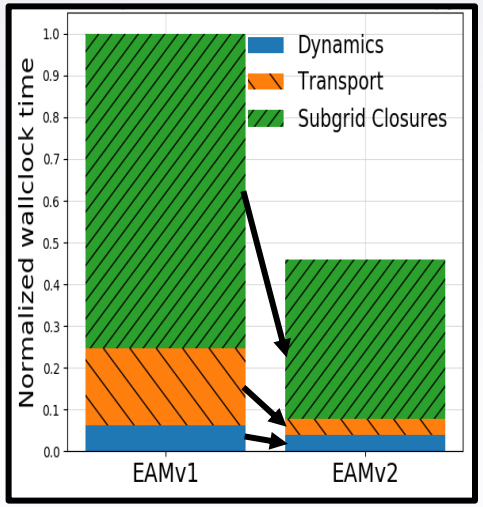

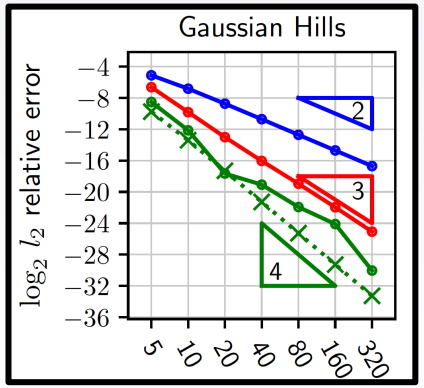

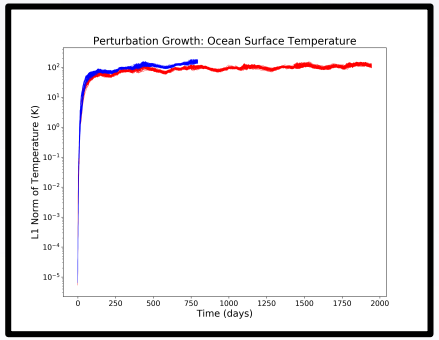

Entraining Computational Scientists into E3SM

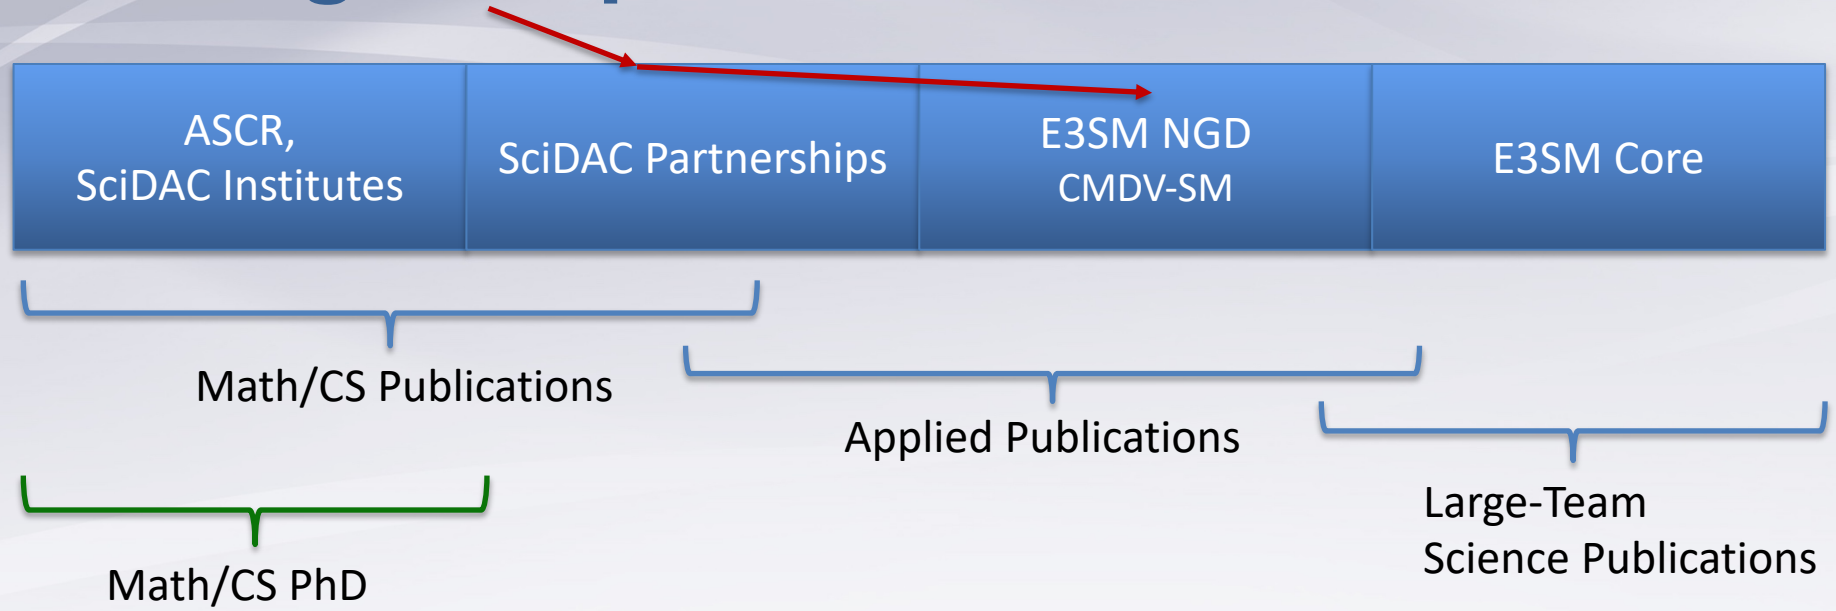

Plan:

1. Bring in Math/CS researchers into SciDAC, and encourage them to move to the right (more applied) over career:
   - ➢ Climate is appealing
2. Find people at labs where you can have a career being more applied.
   - Work to make these staff visible (E3SM Award; Advocate for them at annual performance reviews)