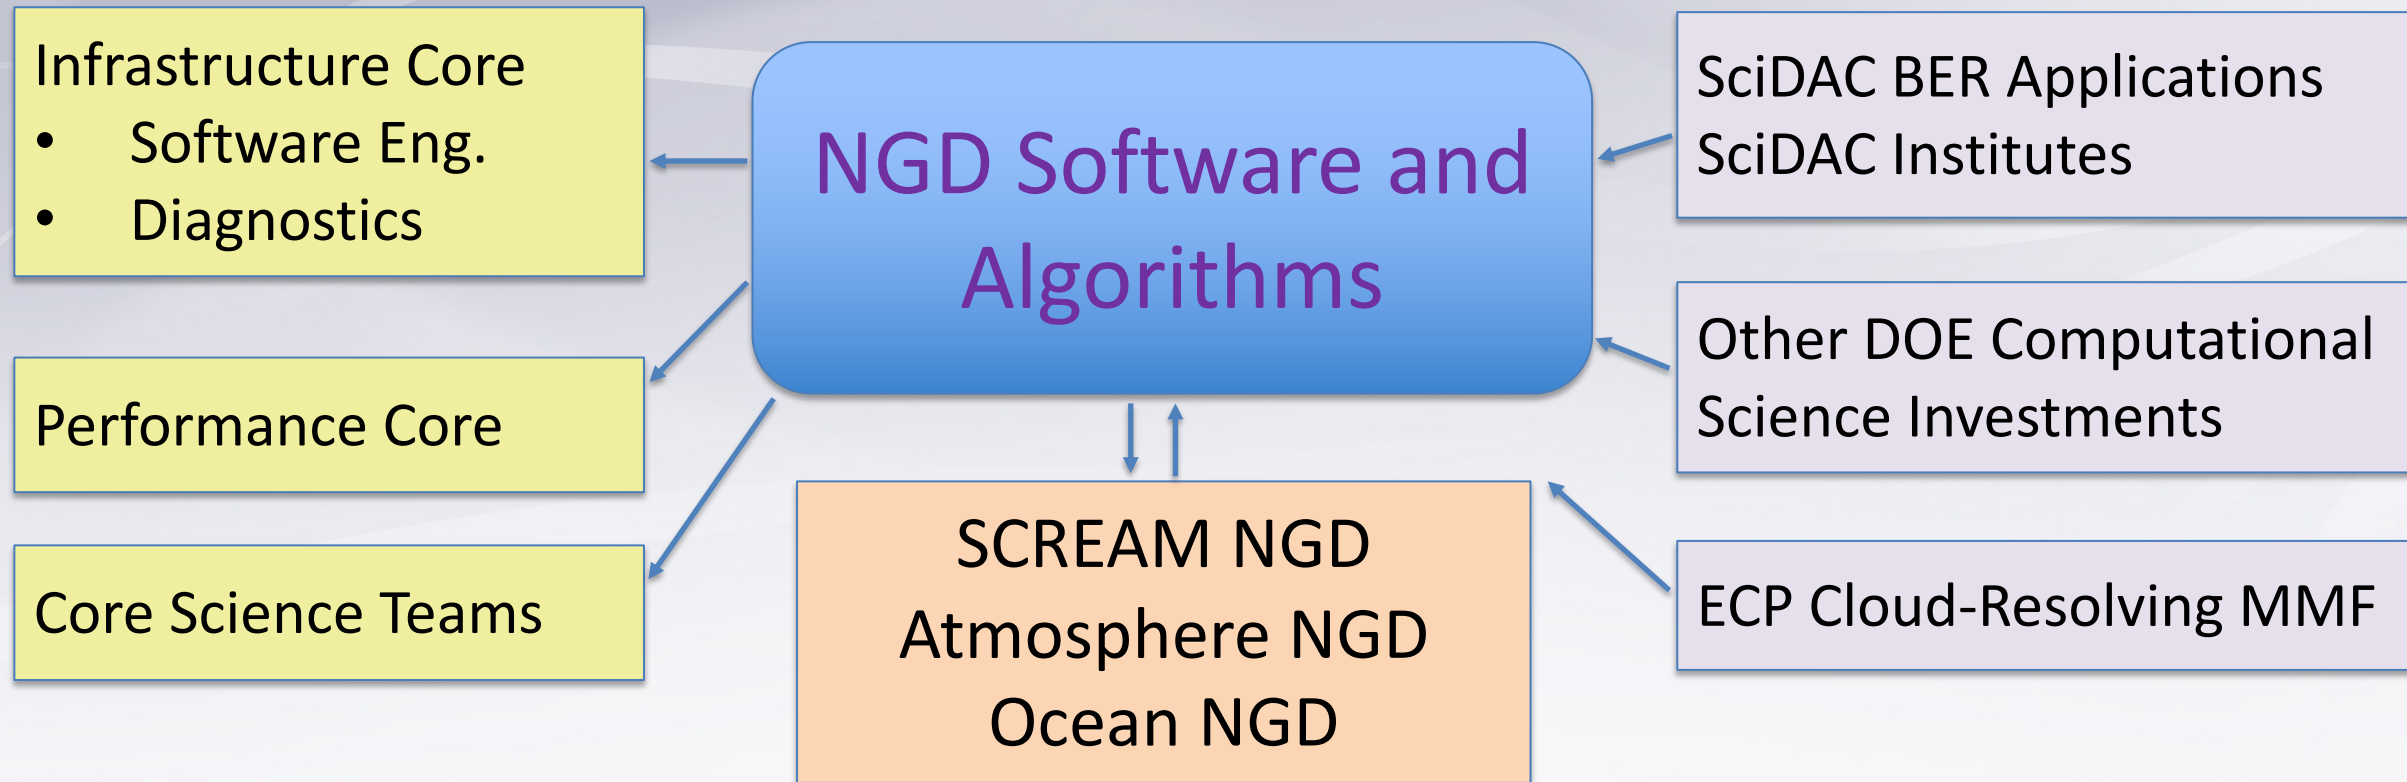


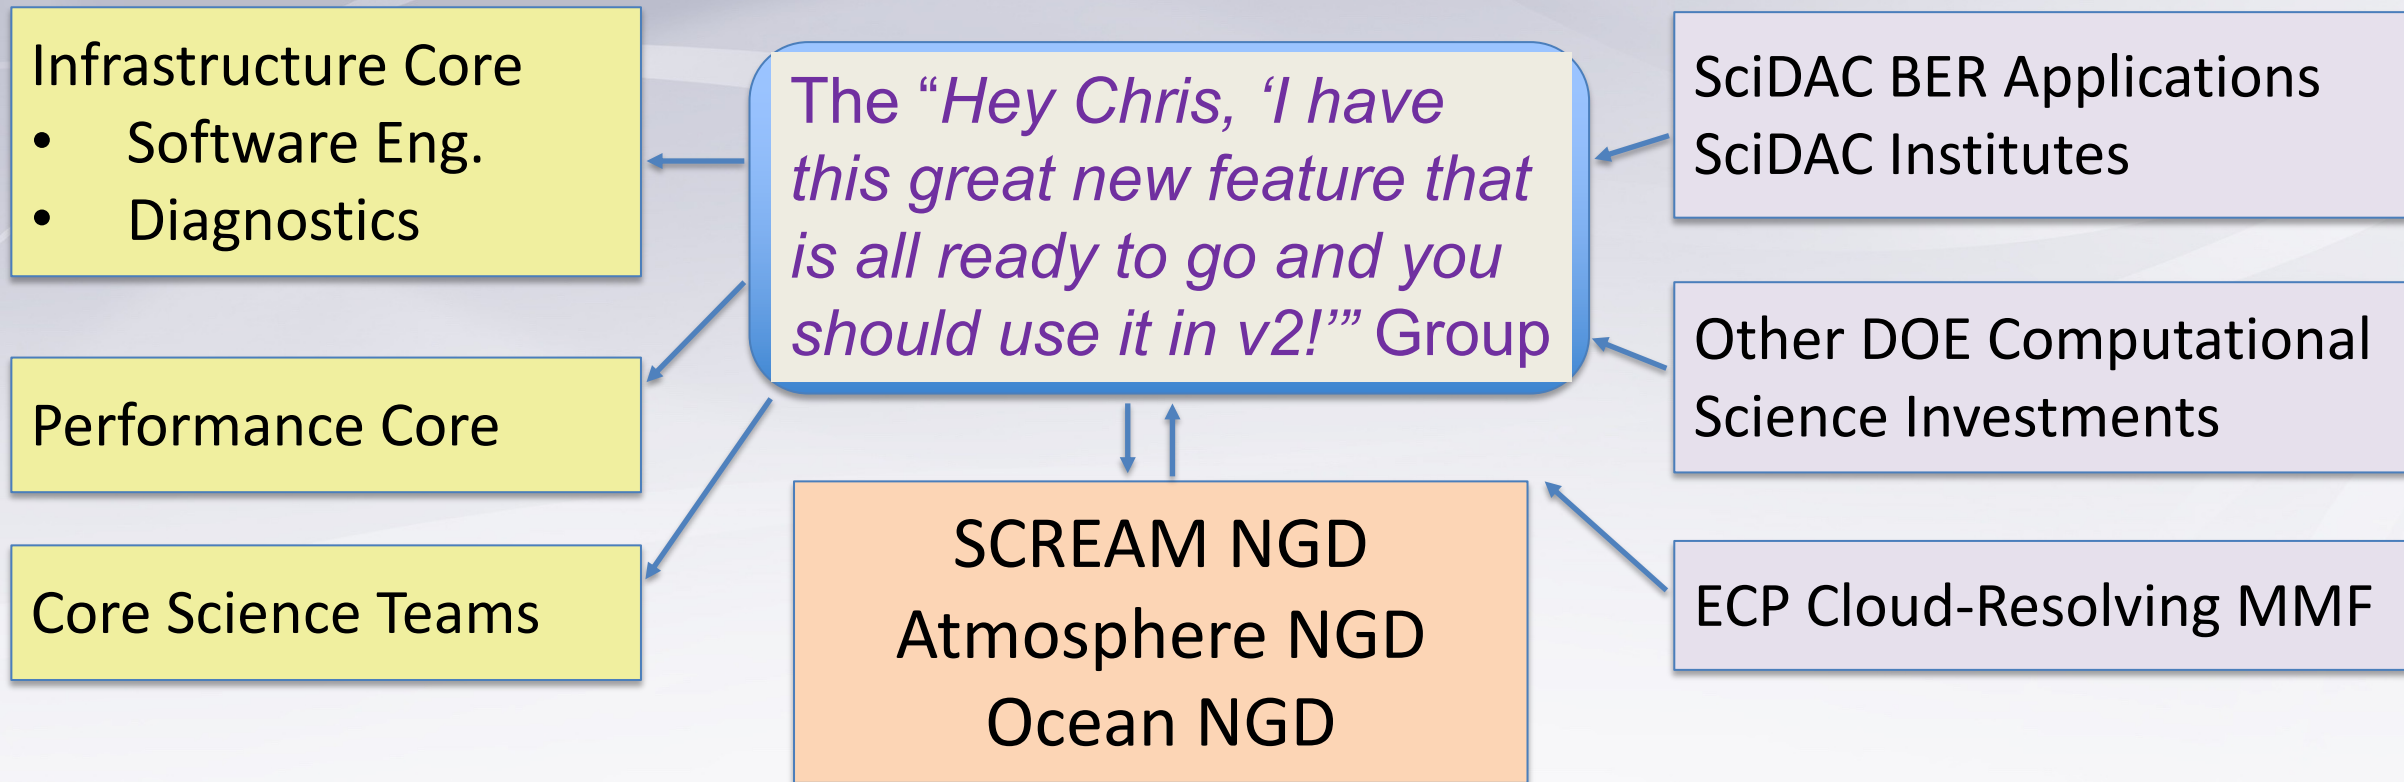
NGD Software and Algorithms



Andy Salinger, E3SM All-Hands Meeting, June 2021

Benj Wagman and Kenny Chowdhary; Andrew Bradley; Hyun Kang; Salil Mahajan & Mike Kelleher;
Vijay Mahadevan, Iulian Grindeanu, & Karen Devine; Andrew Steyer & Cassidy Krause; Josh Fu & Rong-you Chien

NGD Software and Algorithms



Andy Salinger, E3SM All-Hands Meeting, June 2021

Benj Wagman and Kenny Chowdhary; Andrew Bradley; Hyun Kang; Salil Mahajan & Mike Kelleher;
Vijay Mahadevan, Iulian Grindeanu, & Karen Devine; Andrew Steyer & Cassidy Krause; Josh Fu & Rong-you Chien

1a. Automated Tuning for E3SM

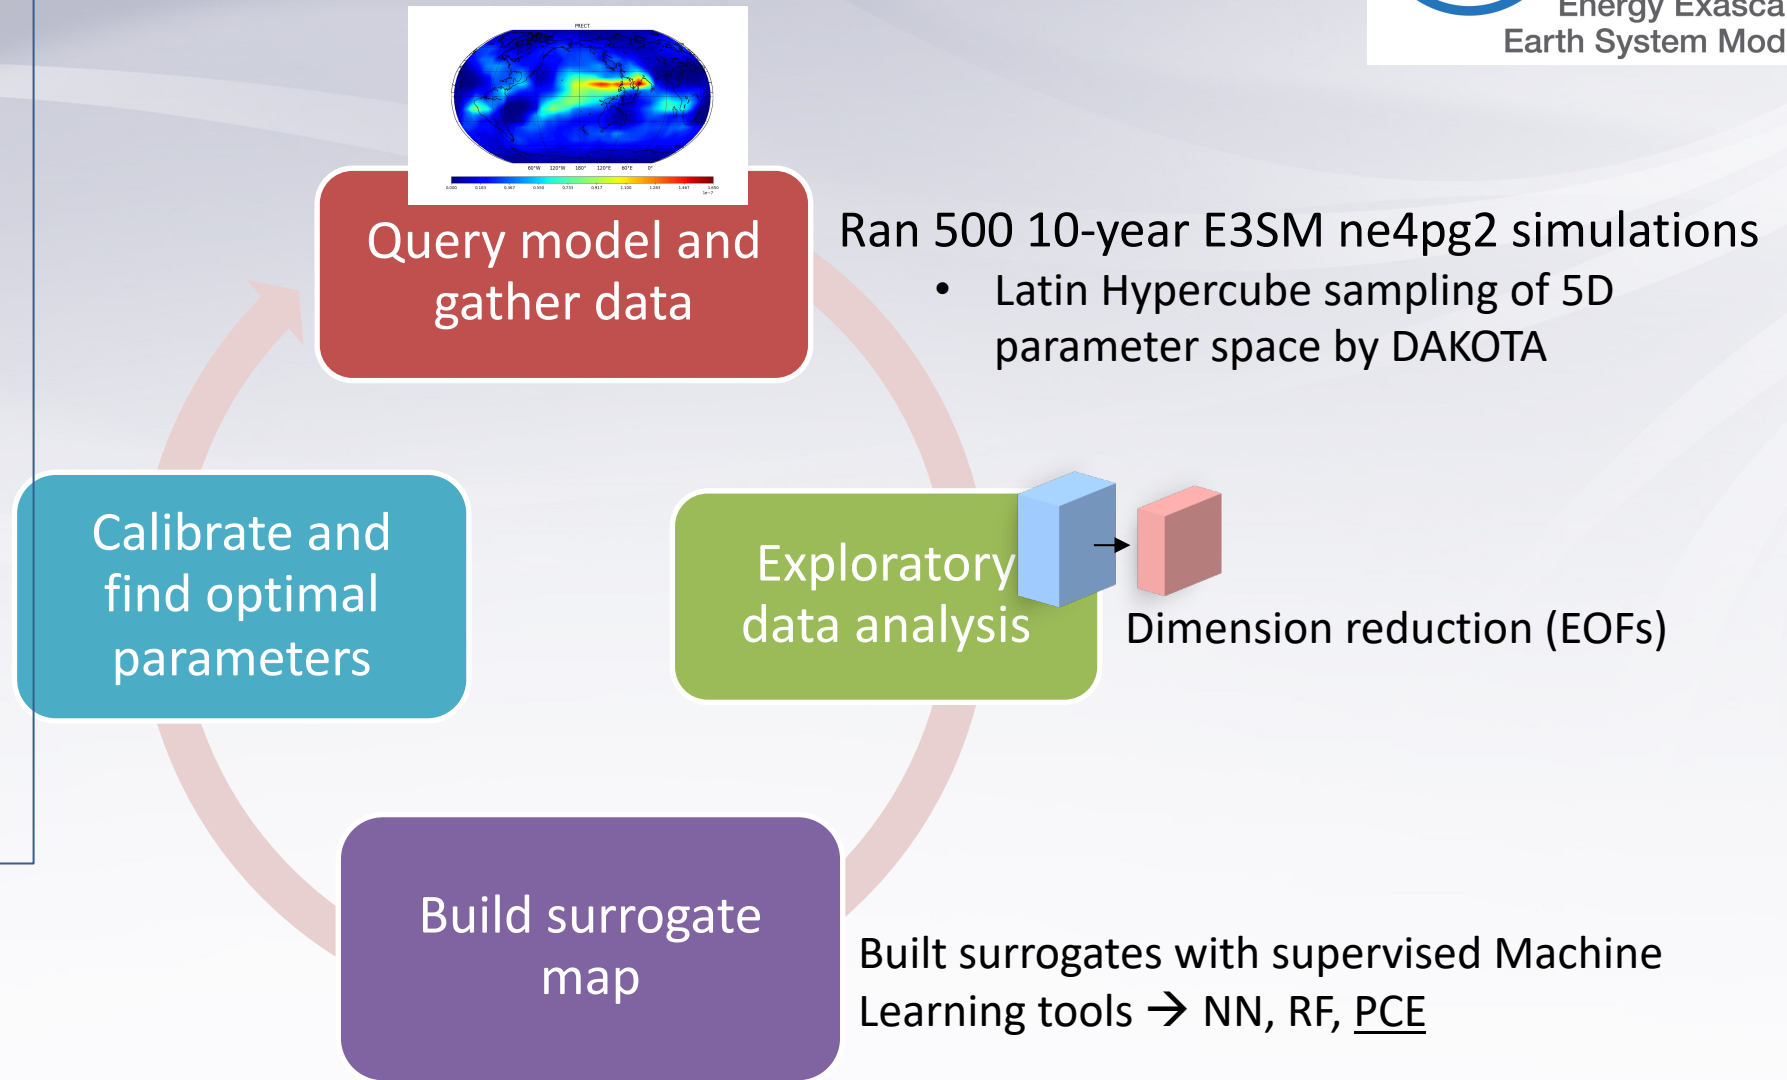
(Started 11/20)

1. **Goal:** Develop optimization algorithms and software infrastructure for simple, flexible, and robust automated model tuning for E3SM.
2. **Motivation:**
 - a. Tuning “by hand” is time-consuming and inherently non-reproducible.
 - b. Formalizing tuning enables Bayesian UQ, e.g. by generating calibrated ensembles.
3. **Challenges:**
 - a. Automated tuning requires surrogate construction or direct model sampling.
 - i. Computationally expensive and time-consuming.
 - b. A single objective function encapsulating multiple quantities of interest (metrics) must be formulated.

1b. E3SM Surrogate Construction

Prototype problem:
Develop methods and algorithms for constructing a surrogate for EAM spatial fields:

- Ultra-low ne4 resolution
- F-case with **5** uncertain parameters.



1c. Recent Progress and Next Steps

1. What have we learned?
 - a. 200-300 5-year E3SM simulations are sufficient to build a surrogate for the ultra-low resolution 5-parameter F-Case spatial fields with less than a 10% relative error.
 - b. Tools: Scikit-learn.
2. Next steps:
 - a. Surrogate construction for spatial fields for ne30pg2.
 - b. Calibration using simple objective functions.

Updates and results can be found under [our confluence page](#):

- search on “Autotuning” in Confluence

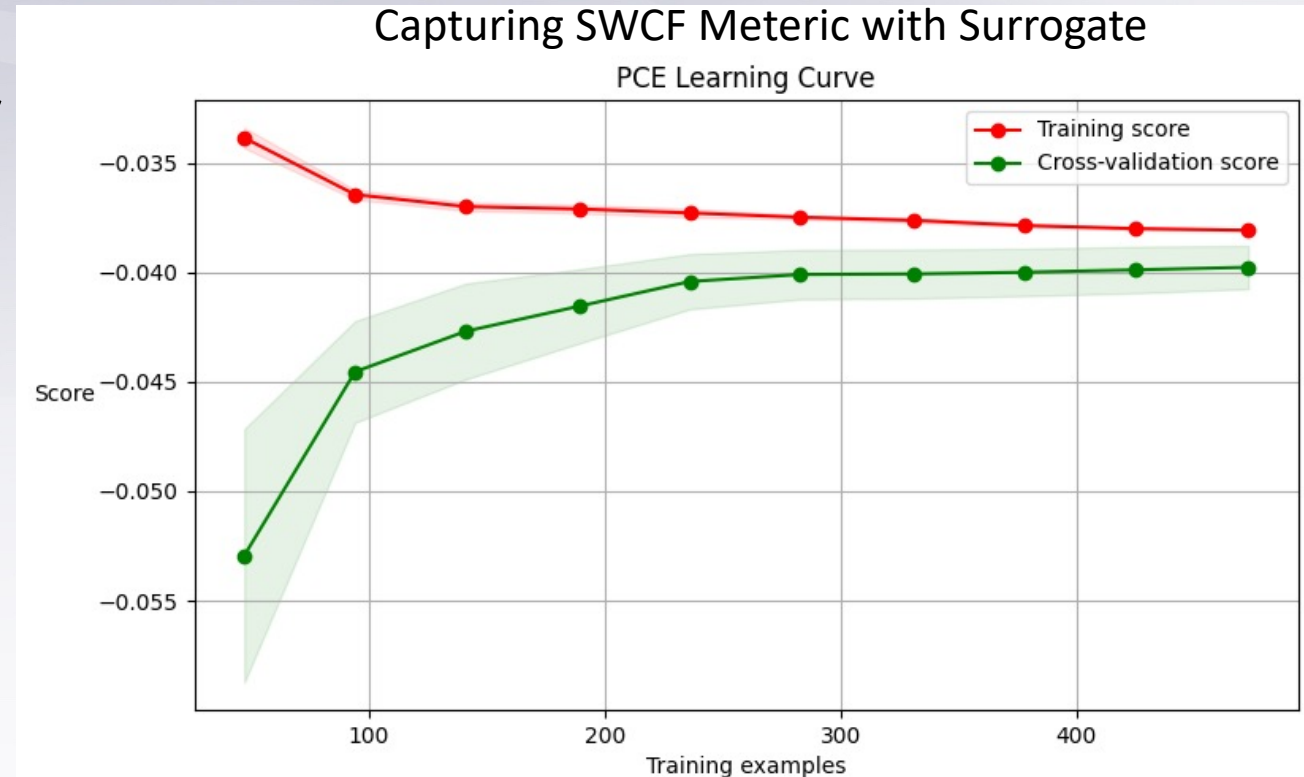


Figure 1 (Surrogate learning curve). Here we show the training and test scores (higher is better) for the surrogate model as a function of the number of training samples. The gap between the two is the overfit error. While it is initially large, the gap lessens and the scores converge. After 250 points or so, additional training data doesn't have a significant impact on the accuracy of the surrogate, and so more data may not be needed.

2a. Semi-Lagrangian Transport on GPUs at Scale

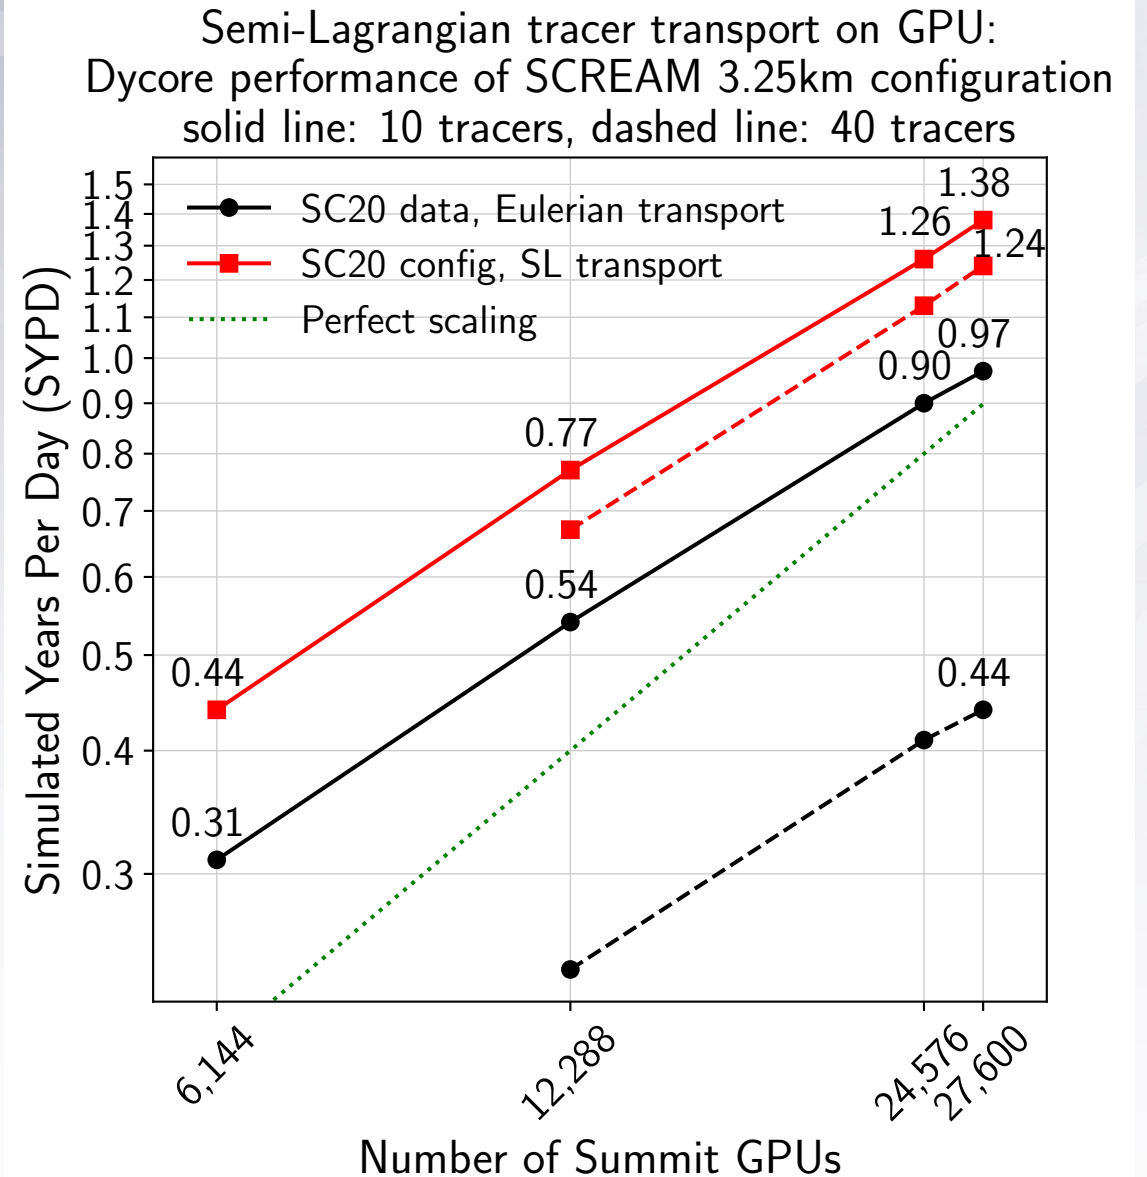
COMPOSE SciDAC: SL Transport

NGD S&A: Upwind SL-MPI, and **now GPUs**

CMDV-SM and SCREAM (HOMMEXX-NH Results)

Dycore Only: Summit 27600 GPUs

Algorithm	SYPD, 10 Tracers	SYPD, 40 Tracers	Cost of 10→40
Eulerian	0.97	0.44	-55%
Semi-Lagrangian	1.38	1.24	-10%

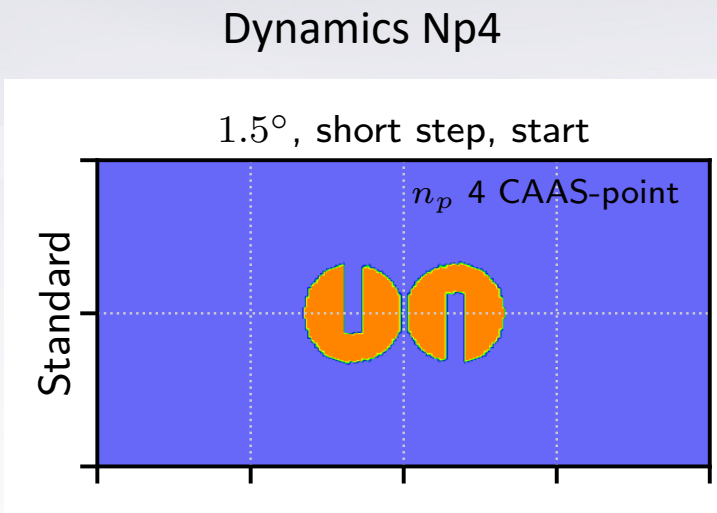
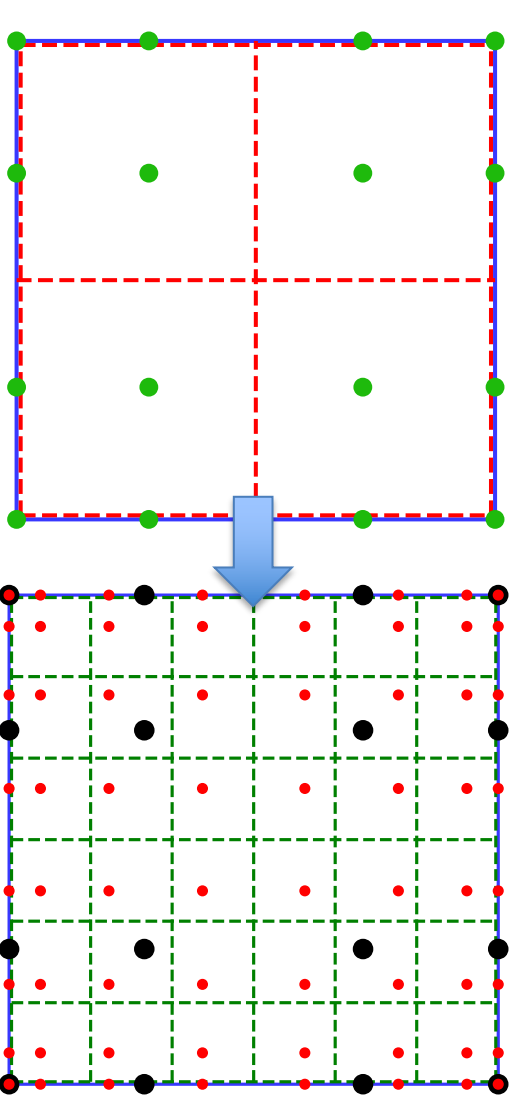


2b. SL Transport with Super-Resolved Tracers

Why: Enable highly resolved tracers and, optionally, denser physics grid, for fixed dynamics cost

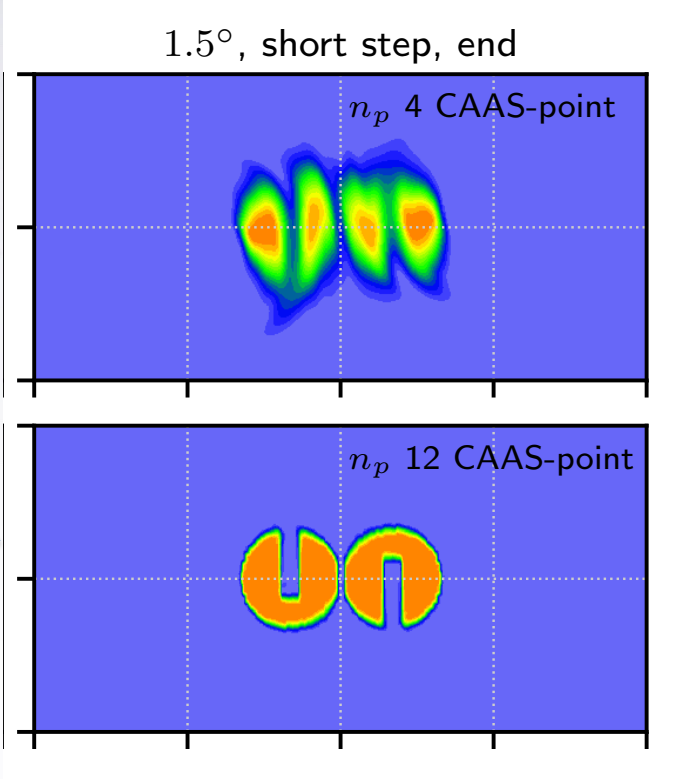
- EAMv1 had 1 Grid: Atmosphere
- EAMv2 has 2 Grids: Physics and dynamics
- EAMv3 with 3 Grids? Physics, dynamics, and tracers

Preserve sub-dynamics-grid filamentary structure of aerosol plumes for applications such as Marine Cloud Brightening. 4-8x greater spatial resolution than with standard method.



Tracer Np4

Tracer Np12



3a. MPAS-O 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!

- **1.74x speedup** of full MPAS-Ocean dynamics

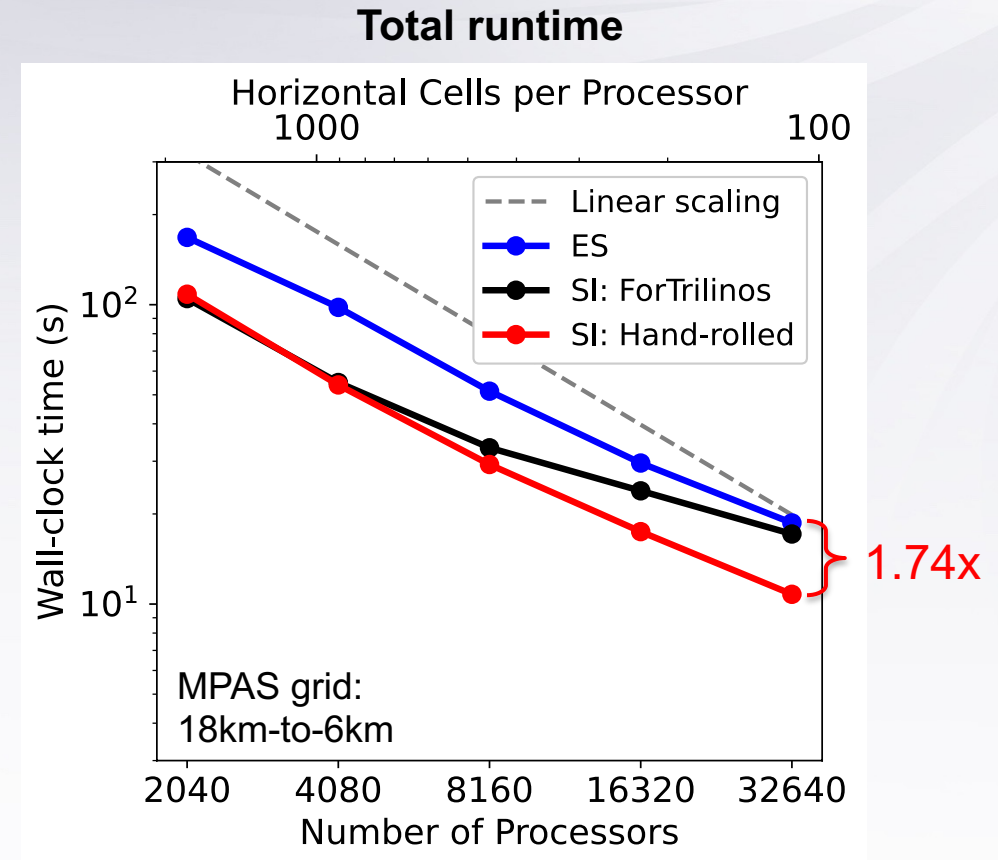
Also developing a semi-implicit solver using RAS preconditioner from **ForTrilinos** (ECP Software)

- Not yet as scalable, but getting optimized
- Provides quicker path to GPUs

+ GPU porting of the semi-implicit solver (just initiated)

- OpenACC directives, MAGMA for linear algebras

Hyun Kang



3b. MPAS-O Baroclinic Time Stepping

Idea: Change of the baroclinic time stepping method to reduce compute time

- A current time stepping method (the predictor-corrector method) repeats a whole time-stepping procedure **twice** to achieve the second-order accuracy.
- A new time stepping method, the second-order Adams-Bashforth method (AB2), computes the time-stepping procedure **once**.

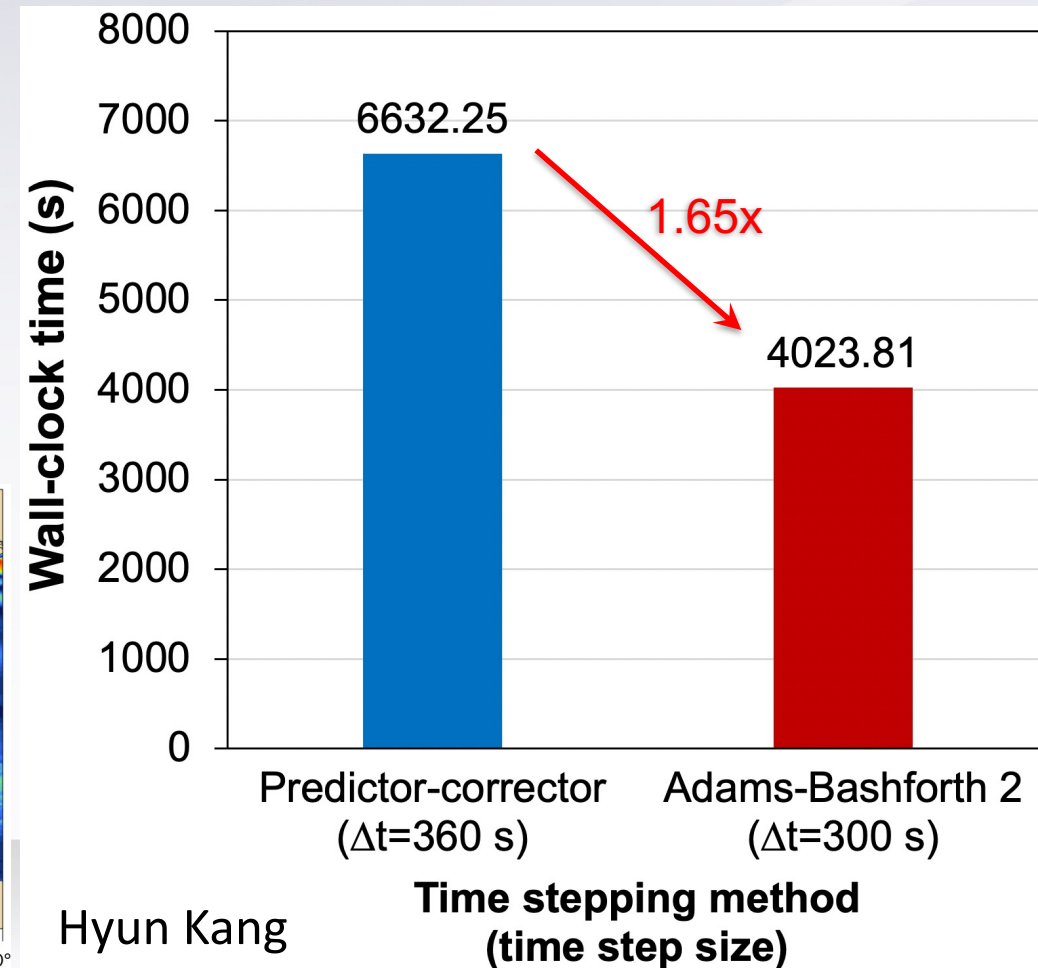

Preliminary results

- 1.65x speedup of full MPAS-Ocean dynamics
- Similar quality of results

Global ocean test case (30to10 km mesh, Sfc. Kenergy at 30 day)

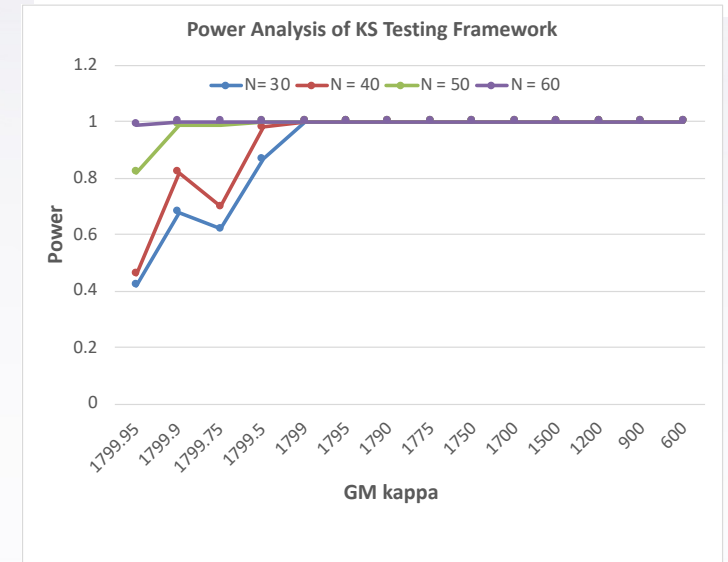
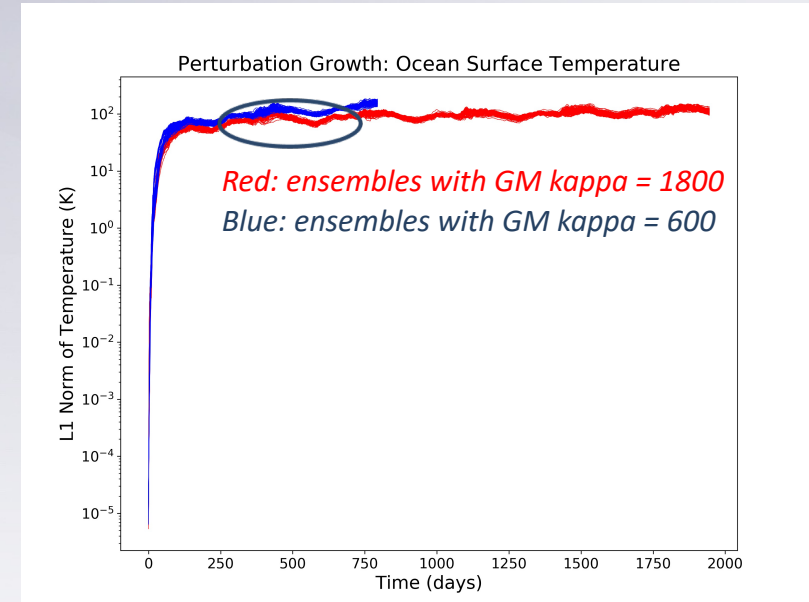
Predictor-corrector (current)

Adams-Bashforth 2 (plan)



4a. Climate Reproducibility: Multivariate Tests for MPAS-O

- a. Approach:
 - i. Test null hypothesis that model climate after a non-b4b modification is unchanged
- b. Strategy:
 - i. Short G-compset ensembles:
 - 1. 2-yr simulations, 30 ensembles each of modified and original model
 - ii. Compare pdfs at each grid point between ensembles:
 - 1. using KS-test and Cucconi test, with corrections for multi-testing
 - iii. Variables evaluated: SSH, U, V, T, Salinity
- c. Results:
 - i. Power analysis - Tests catch small changes to model climate, e.g.:
 - 1. GM Kappa from 1800 to 1799 or lower
 - 2. Redi Kappa from 0.0 to 5.0 or more
 - 3. KPP Richardson number from 0.25 to 0.255 or more
- d. Prototype test is ready:
 - i. Ensemble generation, post processing and tests by one driver script
 - ii. Currently using csh, python, mpi4py and R
 - iii. Some parts ported to EVV, ongoing work for full integration into CIME testing
 - iv. Ongoing work to improve scaling - reduce testing time.



Probability of detecting a change in model climate when GM Kappa is changed from the default value (1800) for different ensemble sizes (N) at the QU240 resolution.

Mahajan et al. (PASC'21, in print)

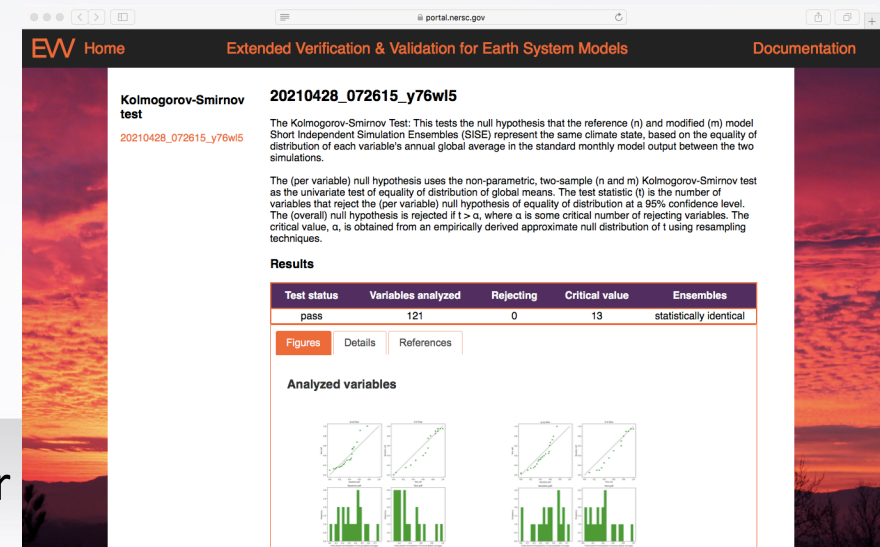
4b. Climate Reproducibility: Multivariate Tests for SCM

1. Quick Multivariate Statistical Tests for non-b4b EAM modifications to physical parameterizations:

- a. Uses Single Column Model
 - i. SCM sensitive to initial conditions
- b. Generate ensembles:
 - i. Perturb initial conditions at machine-precision
- c. Evaluate pdfs of output variables:
 - i. Multivariate testing framework
- d. Detects known changes to:
 - i. Clubb scheme parameters
 - ii. ZM scheme parameters
 - iii. Evaluating other parameterizations
- e. Formal power analysis ongoing

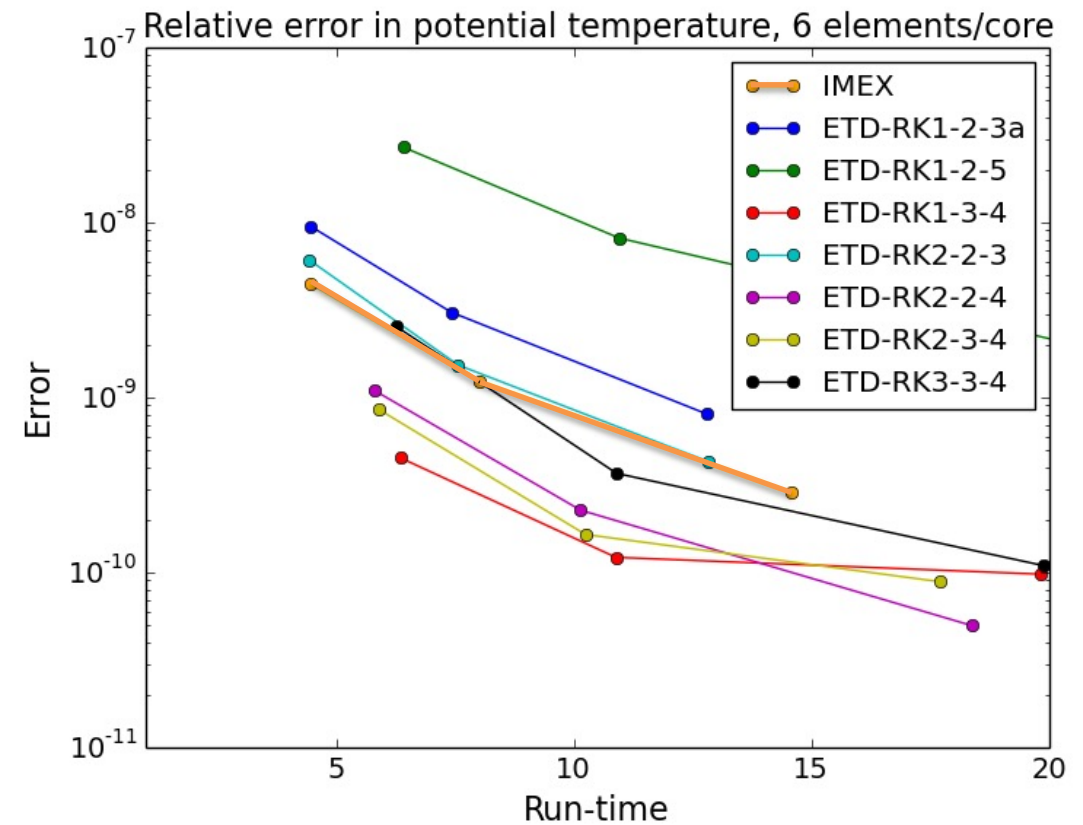

2. Revived EAM reproducibility tests:

- More flexible, adapts to newly re-named EAM / ELM or CAM / CLM
- Fixed time-stepping issues with **MVK** and TSC
 - Use default coupling frequency for **MVK**
 - Specify `se_tstep` as 1/12th of `dtime` for **TSC**
- Create unique ensembles using clock dependent seed for **MVK** test



5. ETD time-integration algorithms for Nonhydrostatic atmosphere

Exponential Time Differencing is alternative to IMEX




Efficiency study: Several ETD methods compared against optimized IMEX method in HOMME-NH.

- ETD is competitive with IMEX
- ETD shows better strong scalability than IMEX

6. Geometric-inference based partitioners for MOAB-E3SM: optimizing intersection vs runtime projection

Goal: Faster coupler weight generation workflows:

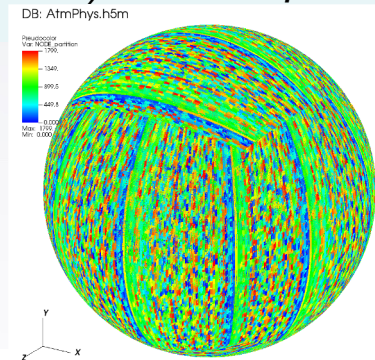
- **Offline:** mbtempest (*MOAB + TempestRemap*)
- **Online:** E3SM-MOAB
- Offline strategy maximizes geometric overlap in parallel partition of component meshes for faster intersection mesh computation
 - 6x speedup for high-res cases
- Inferred online partitioner implementation for E3SM coupler with MOAB-Zoltan interfaces is currently being verified
 - Runtime communication minimization experiments during Ax projection operation by using ATM physics decomposition directly
 - Latest Zoltan release (v3.9.0) and Upcoming MOAB release (v5.3.0)



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




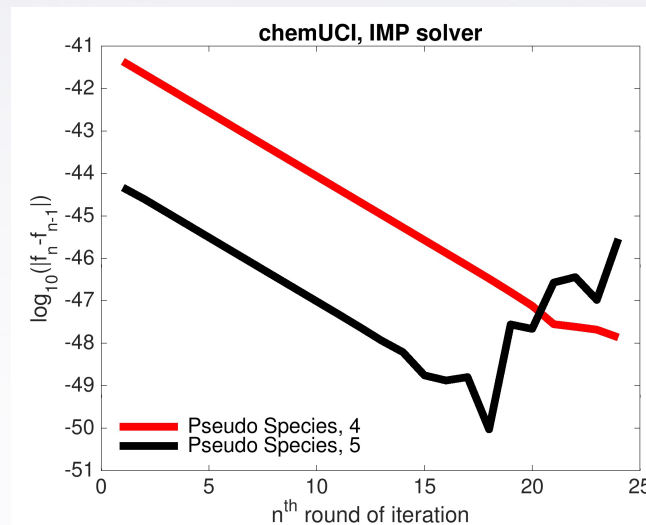
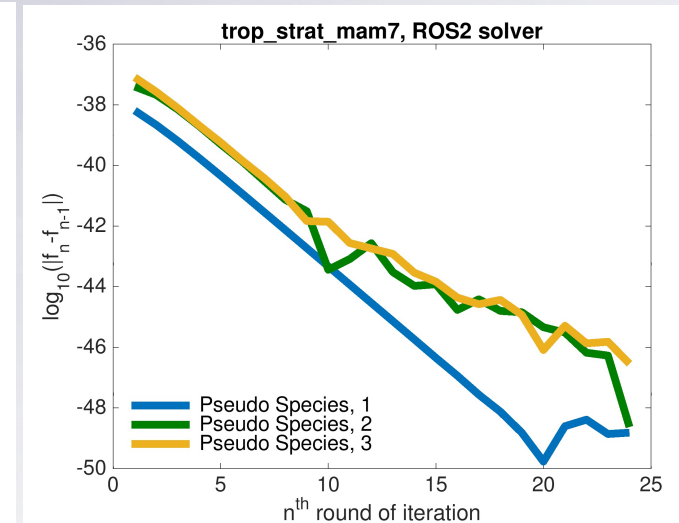
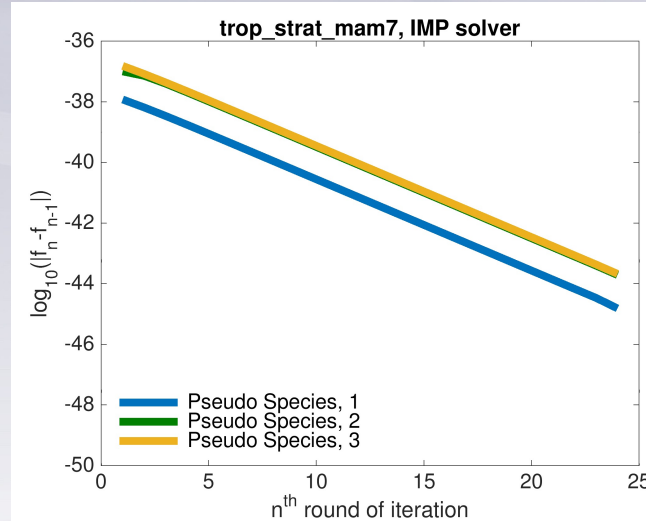
Atmosphere partition inferred from ocean's RCB partitioning tree



Atmosphere (NE120) physics partition layout determines communication effort for online remapping (1800 tasks)

7a. Atmosphere Chemistry Solver

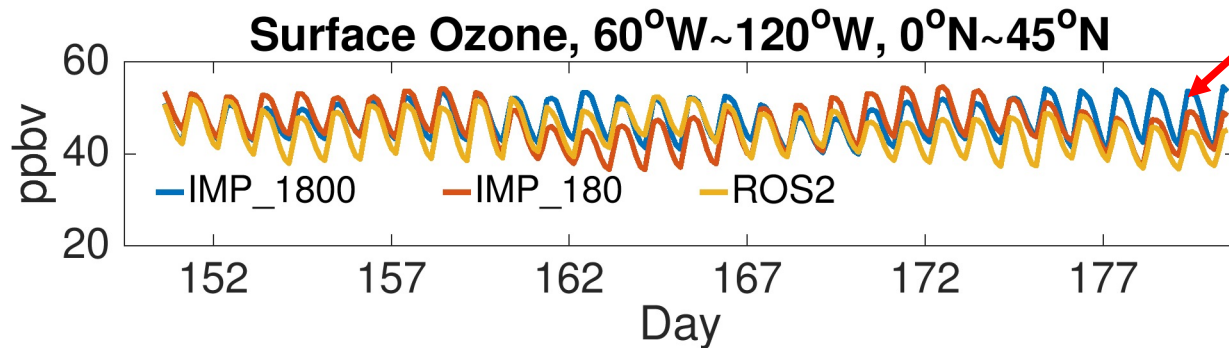
- Two mechanisms: the **two stage Rosenbrock solver (ROS2)** has been tested within UCI chemistry and TROP_STRAT_MOZART_MAM7
- In the box model, the ROS2 solver compares well with the current implicit solver IMP
 - ROS2 solver can have second-order $o(\Delta t^2)$ convergent rate, compare to first-order IMP
 - UCI chemistry can have a 20% improvement
 - TROP_STRAT_MOZART_MAM7 can have a 25% improvement
- Implementation in E3SM is ongoing




7b. Atm Chem

Model Evaluation

1. 6 months simulation of surface ozone with trop_strat_mozart_mam7 is shown
2. In order to fulfill the regional model's request, a finer time resolution (180s) within the chemistry mechanism is used when applying the ROS2 solver
3. Mean difference shown that using 180s within the chemistry solver may enlarge the amplitude of ozone, when compared to the IMP solver (1800s)



Josh Fu and Rong-you Chien



Ideas for Future Software and Algorithms Developments (or SCREAM, Ocean, Perf, Infrastructure)

1. Automatic tuning and Bayesian calibration
2. Verification/Correctness
 - Convergence in Δz and Δt
 - Energy Conservation
3. Optimal vertical coordinate r-adaptivity
4. Time Stepping: adaptive, RRM- Δt
5. High-Order MPAS-Ocean Discretization
6. MPAS Sealce Load Balance / Perf
7. ML: low-hanging fruit from AI4ESP
 - SCREAM data to learn ZM-fixer
8. Infrastructure for large ensembles (1000 simultaneous ne30 runs)
9. Spin-up
 - Ocean, Land, Land-Ice, GCM