Efficient and Objective Testing of Solution Correctness for HPC Applications

The Challenge

- The Energy Exascale Earth System Model (E3SM) solves a large set of equations in millions of grid cells. We do not have a single metric to describe solution fidelity (i.e., model's predictive skill)
- A typical comparison of two solutions involves comparing "climate statistics" of tens to hundreds of physical quantities. Multiple years of simulations are needed to produce data to derive such statistics, costing a lot of computer time
- Simpler, cheaper, and more objective methods are needed for the testing of solution correctness, especially for high-resolution simulations

The Test Procedure

1. Perform short simulations (minutes to days) with a benchmark code in a trusted environment using a range of time steps (Δt)
2. Take the benchmark solution computed with smallest Δt as reference, compute solution RMS difference for other simulations with longer Δt. Characterize self-convergence
3. Conduct test simulations. Calculate solution difference relative to reference solution from benchmark code
4. Deviation of convergence behavior from benchmark indicates non-negligible role of "new errors" (see schematic in Figure 1 and application examples)

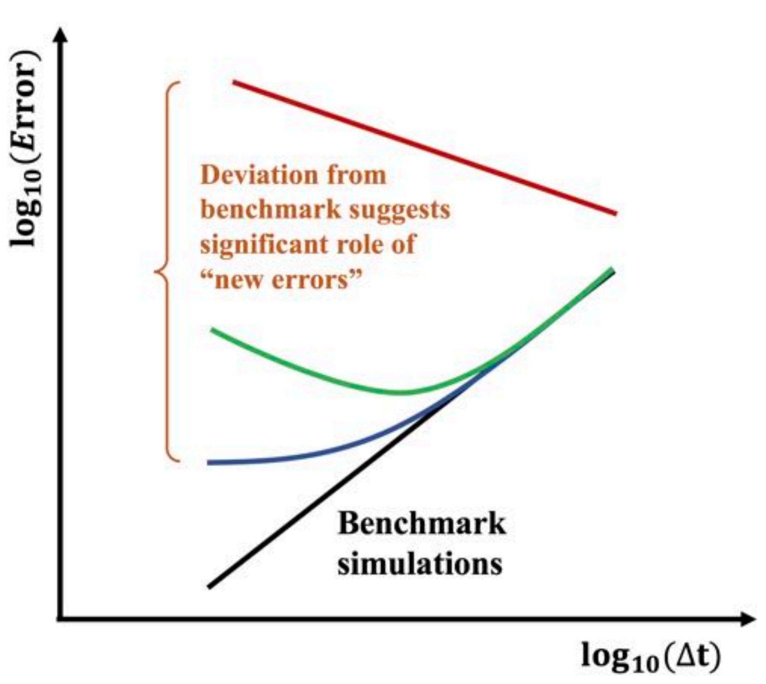

Figure 1. A schematic showing solution self-convergence in the benchmark (black) and possible convergence pathways that indicate the impact of "new errors" (colored lines).

Advantages of the Method

- No code modifications needed
- Easy to perform and automate
- Objective and computationally efficient
- Applicable to a subset of the code, useful for debugging
- Capable of detecting changes in both global and/or regional features of the results

Key Results

- A simple, quantitative and objective method is proposed using the concept of time step convergence
- Results from short convergence tests are found to be good predictors of E3SM's long-term behavior related to fast physics, making the test computationally very efficient
- The method is expected to be applicable to other models that solve time evolution equations

Applications

Example 1. Mixed precision in the E3SM atmosphere model

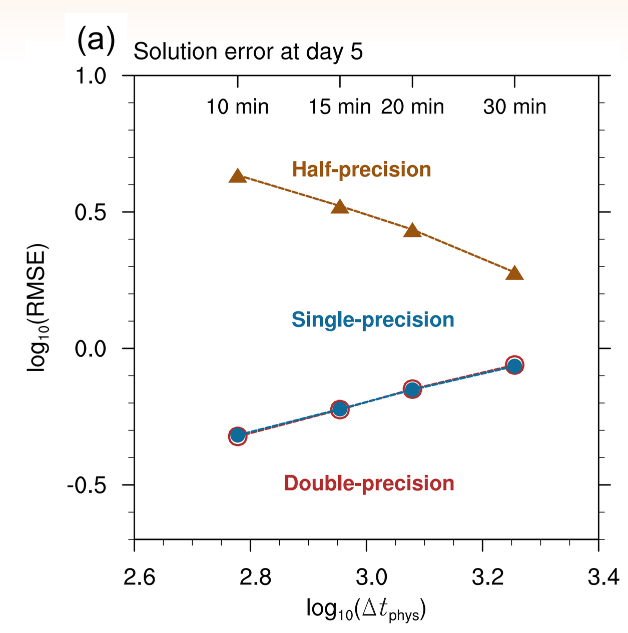

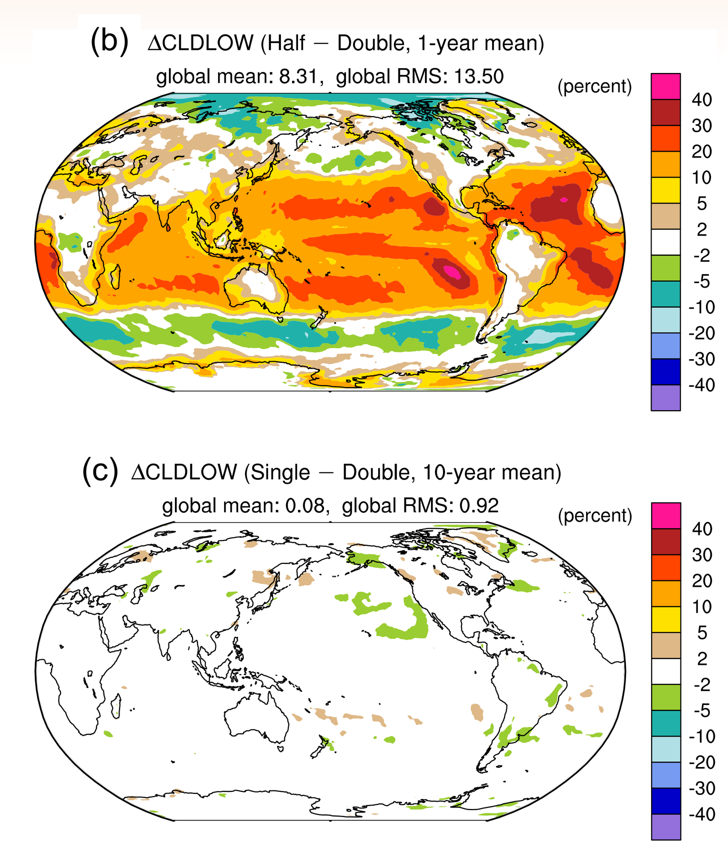

Figure 2. The E3SM atmosphere model is configured to use reduced (single or half) precision in its turbulence parameterization with double precision for the rest of the model. (a) Convergence test based on 5-day simulations suggests that single precision gives essentially the same results while the solutions computed using half precision are dominated by precision error. (b)-(c) This conclusion is confirmed by climate simulations performed with the default Δt. Shown is the annual mean low cloud fraction difference between mixed-precision and double-precision simulations. *(Zhang et at., 2019, submitted)*

Example 2. Software and Hardware Changes

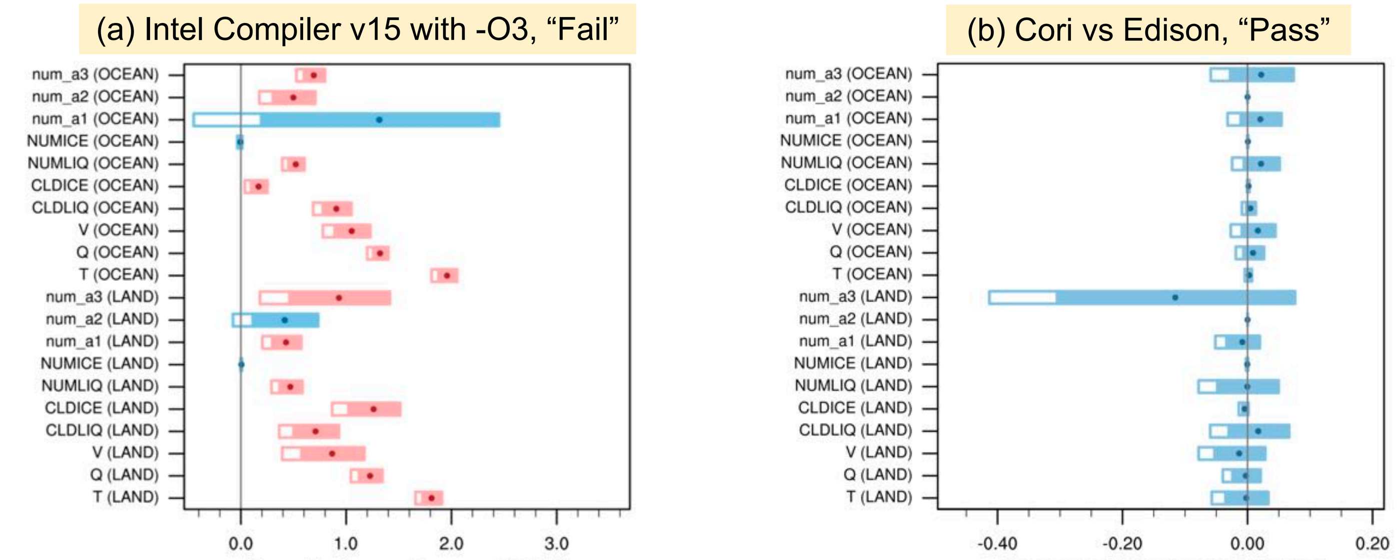

Figure 3. (a) Intel compiler version 15 with –O3 optimization is found to give incorrect solutions. (b) Solutions from from Cori and Edison at NERSC are found to be practically equivalent. Ensemble simulations using 12 different initial conditions were used to add confidence. Shown here are results of statistical testing of the ensemble differences after 5 minutes of model time. Panel (a) is reproduced from *Wan et al. (2015), doi:10.5194/gmd-10-537-2017.*

The team: Shixuan Zhang[1], Hui Wan[1], Phil Rasch[1], Balwinder Singh[1], Vince Larson[2] and Carol S. Woodward[3] ([1]PNNL, [2]U. Wisconsin – Milwaukee, [3]LLNL)

This work is part of the SciDAC project "*Assessing and Improving the Numerical Solution of Atmospheric Physics in E3SM*".