Assessing CESM scalability for hierarchical model ensembles

Principal Investigator
*Ryan L Sriver, Assistant Professor, Department of Atmospheric Sciences, University of Illinois at Urbana-Champaign, Email: rsriver@illinois.edu

Co-Principal Investigator
Hui Li, Graduate Student, Department of Atmospheric Sciences, University of Illinois at Urbana- Champaign, Email: huili7@illinois.edu

*Corresponding author

Executive summary (150 words)

In this project, we test the scalability of the Community Earth System Model (CESM) across a wide range of model resolutions, complexities, and configurations. Results enable us to assess the cost and applicability of using CESM in a hierarchical ensemble framework, in which we run a large number of short simulations using reduced complexity versions of CESM to assess the predictability of complex 3-D atmospheric behavior. We further performed benchmarking and load balancing of the high-resolution versions of the model with various coupling strategies and output frequencies, which provides references for the next phase of high resolution climate simulations focusing on climate change assessments and coupled climate feedbacks.

Descriptions of research activities and results

Climate models are valuable tools for understanding the fundamental physical processes and interactions within the Earth system and how the system is changing. Comprehensive coupled climate models with high complexity are designed to provide the most realistic representation of the Earth system. Such models have been widely used to assess the physical processes and large-scale dynamics influencing climate change impacts, as well as potential future changes under anthropogenic global warming. Examples include those used in the Coupled Model Intercomparison Project (CMIP) Phase 5. A major limitation of these comprehensive climate models is their high computational cost, making it difficult to quantify uncertainties related to internal variability of the relevant physical processes and model structures. Characterizing uncertainties surrounding structural model errors and limitations is critical to developing sound adaptation and risk management strategies, particularly at local to regional scales.

Simple climate models with reduced complexity and/or resolution have a much lower demand on computational resources, and can be used in large ensemble experiments. The tradeoff here is that simple models contain simplified physics and lack key coupled interactions, thus leading to potentially underconfident assessments of climate change projections and variability. Here we tested a hierarchical ensemble framework method using the Community Earth System Model (CESM), in which we perform 1000s of short simulations with varying initial conditions and uncertain parameter values to develop predictive relationships between simple models and complex 3-D atmospheric behavior. This new modeling framework will free us from the typical computational bottlenecks associated with running dynamic coupled models at fine-scale grid resolutions, and it increases the scalability of climate model ensemble experiments to infinitely large numbers of cores. This type of hierarchical ensemble framework can be useful for assessing model uncertainties and provide fundamental insights into probabilistic climate prediction.

In this project, we use Blue Waters to test the scalability of CESM across a wide range of model configurations, resolutions, and complexities. We also examine the reproducibility of recent perturbed physics ensemble efforts (performed at NCAR) utilizing atmosphere-only configurations of CESM featuring Community Atmosphere Model (CAM4/CAM5). Blue Waters provides several distinct advantages over other computer systems (e.g. XSEDE) for conducting this research. Blue Water's petascale computing environment offers unique capabilities to handle the computational demands associated with executing large CESM ensemble experiments, as well as controlling the high-frequency input and output of information.

We perform scalability analysis of a series of model configurations varying atmospheric model versions (CAM4 vs. CAM5), resolution (1° vs. 0.25°), coupling strategy (atmosphere-only vs. fully coupled) and output frequency (monthly vs. daily and 6-hourly). We also test the sensitivity of model performance to the number of OpenMP threads per task for selected configurations. The scaling results and PE costs and are summarized in Figure 1 and Figure 2, respectively.

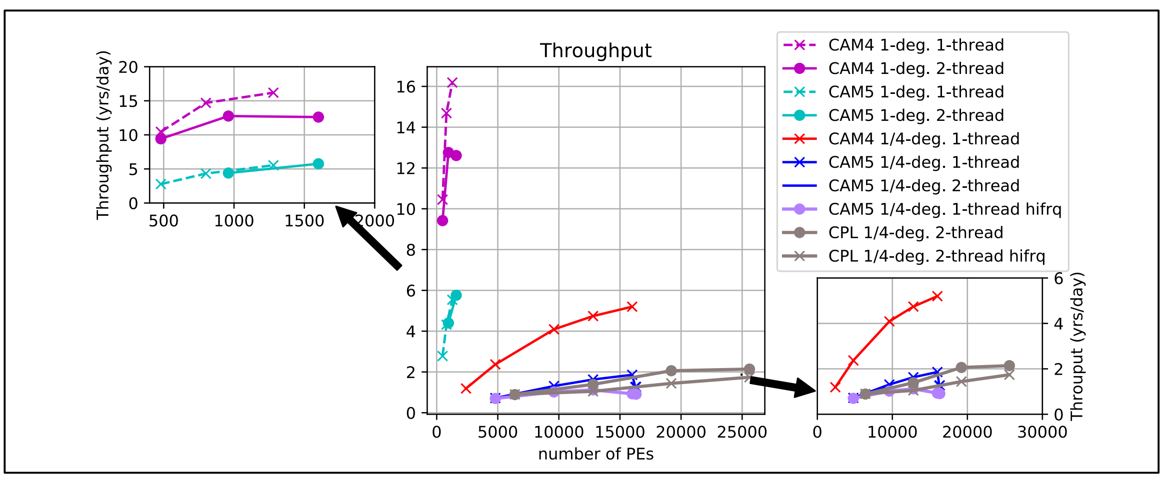

Figure 1. Summary of model throughput (simulation years per day). Experiments vary in coupling strategy (CAM indicates atmosphere-only, CPL* indicates fully coupled),*

resolution (1-deg. and 1/4 deg.), atmosphere model versions (CAM4 and CAM5), number of OpenMP threads per task (1-thread and 2-thread), and output frequency ("hifrq" indicates daily and 6-hourly output).

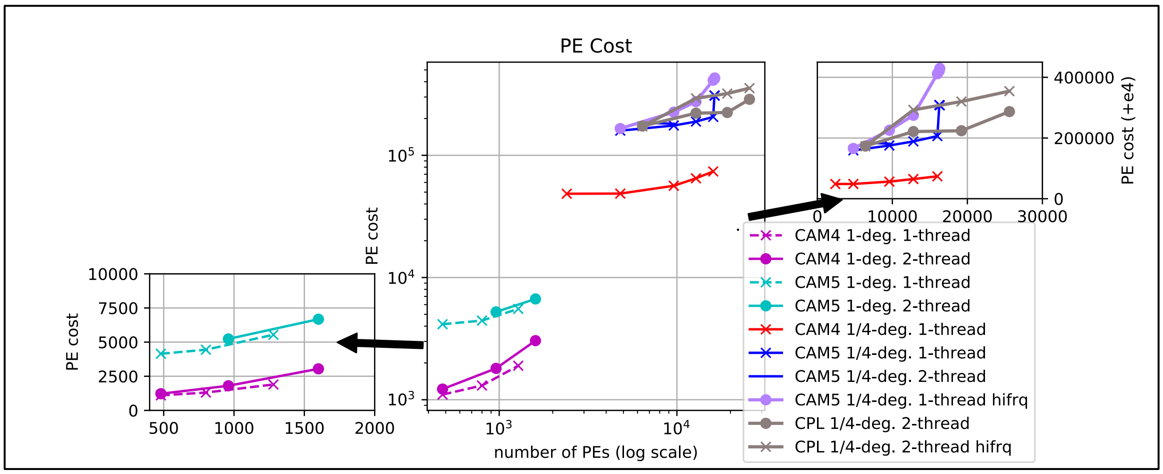

Figure 2. Summary of PE cost (PE hours per simulation year). Legends are consistent with Figure 1.

We find that CAM5 in general is more computationally expensive and less scalable than CAM4. The relatively low resolution (1˚) CAM models achieve optimal efficiency on roughly ~1000 cores, and the higher resolution (0.25˚) models are most efficient on ~16000 cores. PE cost is relatively insensitive to ocean coupling when considering the 0.25˚ atmosphere model. Increasing file output frequency can increase the model cost by 20%~30%.

These results enable us to assess the applicability and cost of using CESM in the hierarchical ensemble framework, and provide important references for optimum PE configuration to achieve the most efficient use of computational resources. These results will serve as the basis for new ensembles that we are designing in collaboration with the National Center for Atmospheric Research (NCAR) and Los Alamos National Laboratory. The scaling and benchmarking analysis of the high resolution configurations will also be used for the next phase modeling research on investigating the relationship between tropical cyclones and climate.