Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 108C (2017) 735–744

Procedia Computer Science

International Conference on Computational Science, ICCS 2017, 12-14 June 2017, Zurich, Switzerland

Exploring an Ensemble-Based Approach to Atmospheric Climate Modeling and Testing at Scale

Salil Mahajan[1], Abigail L. Gaddis[1], Katherine J. Evans[1], and Matthew R. Norman[2]

[1] Computational Earth Sciences and Climate Change Science Institute,
Oak Ridge National Laboratory, Oak Ridge, TN, USA
mahajans@ornl.gov; gaddisal@ornl.gov; evanskj@ornl.gov;
[2] Oak Ridge Leadership Computing Facility and Climate Change Science Institute,
Oak Ridge National Laboratory, Oak Ridge, TN, USA
normanmr@ornl.gov

Abstract

A strict throughput requirement has placed a cap on the degree to which we can depend on the execution of single, long, fine spatial grid simulations to explore global atmospheric climate behavior. Alternatively, running an ensemble of short simulations is computationally more efficient. We test the null hypothesis that the climate statistics of a full-complexity atmospheric model derived from an ensemble of independent short simulation is equivalent to that from an equilibrated long simulation. The climate of short simulation ensembles is statistically distinguishable from that of a long simulation in terms of the distribution of global annual means, largely due to the presence of low-frequency atmospheric intrinsic variability in the long simulation. We also find that model climate statistics of the simulation ensemble are sensitive to the choice of compiler optimizations. While some answer-changing optimization choices do not effect the climate state in terms of mean, variability and extremes, aggressive optimizations can result in significantly different climate states.

Peer-review under responsibility of the scientific committee of the International Conference on Computational Science

Keywords: reproducibility, climate simulation, ensemble testing

1 Introduction

Traditional modeling studies addressing climate science questions rely largely on several long, temporally dependent simulations. Relying on single or small long ensembles of simulations is becoming a burden for state-of-the-art high resolution global climate models because the approach does not scale with model complexity. To gain more realism in the model, whether through grid refinement or new model features, the model will require more compute cycles to complete because no realistic atmospheric model achieves perfect weak or strong scaling at the levels of parallelism needed for decent throughput.

1877-0509 © 2017 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of the scientific committee of the International Conference on Computational Science
10.1016/j.procs.2017.05.259

As a result, the workload per compute node decreases "at scale" and the model spends relatively more time exchanging data between nodes. Also, the MPI messages sent between compute nodes become smaller such that nearly all of the time is spent in latency costs. It appears that latency will not improve significantly with future computing architectures, so this problem will only become worse. The Accelerated Climate Model for Energy (ACME) V1 prototype simulations spend more than 90% of their time waiting for MPI data within the atmospheric dynamical core, and that waiting is dominated by latency costs. Clearly, there is a barrier that cannot be removed without rethinking our simulation strategy. If one can cast the methodology to address key climate science questions in a manner amenable to Short Independent Simulation Ensembles (hereafter, SISE) as opposed to (or in cooperation with) Single, Long Runs (hereafter, SLRs), the negative effects of the throughput constraint can be ameliorated. For example, running 100 independent one-year-long ensembles instead of a single 100-year run produces a 100x greater workload per node and, therefore, significantly reduced relative MPI and PCI-e overheads (i.e., better parallel scaling).

Given the dissolution of Moore's Law for the throughput of climate simulations with the next generation of computing, we investigate the viability of SISE for testing and scientific analysis in an effort to help climate community achieve its science goals. It is encouraging that there are already scientific studies in climate that have used SISE for scientific benefit. For instance, Verma et al. [12] have conducted an ensemble of fully-coupled sulfate aerosol forced short runs (1 year) starting with initial conditions from various points of a pre-industrial control run trajectory. Also, the Large Ensemble Community Project [4] has provided wealth of data for analysis. Using ensembles for tuning of model parameters is also of interest, for example [13], used simulation ensembles of a few days with perturbed parameters.

In harnessing modern hybrid computer architectures, climate simulations generate machine round-off level answer changes from identically configured runs via refactoring source code, compiler changes, compiler optimization choices and processor layout changes. The non-linear chaotic nature of the climate system causes these minute perturbations to grow quickly. It is critical to ensure that any of these changes are not systematic, leading to a climate state that is distinguishable from the validated model baseline. Recognizing this, SISE have also been used for verification testing of a new model simulation for these development scenarios [1, 10].

First, we verify the utility of SISE for detecting model changes that lead to statistically disparate climates as demonstrated by [1] with a different strategy. This entails testing the null hypothesis that two SISE belong to the same population. The computational benefits of using SISE in lieu of SLRs in practice are explained. We investigate whether SISE can in fact replace SLRs at equilibrium to represent the natural variability of the atmospheric model, given that the non-linear growth of small state differences limit deterministic predictability to time scales of 10-20 days. [7] show that non-linear atmospheric dynamics can also generate variability of planetary-scale features on decadal time scales, which will not be captured by SISE. Atmospheric internal variability is also understood to generate low-frequency modes of variability like the North Atlantic Oscillation (NAO). NAO variability ranges from sub-seasonal to multidecadal timescales and atmospheric models forced with prescribed SST climatology alone can generate NAO variability across these timescales with little role for the ocean [11, 6].

2 Experiments and Computational Benefits of SISE

We use a version of ACME that predates the first release of the model, used as a baseline for verification. It matches the configuration used in [8] in terms of the model tag and specifies the same active atmosphere and land surface components with a data-only ocean model. However,