A Modest Proposal for Climate Models

By John Benson

June 2024

1. Introduction

This post started with the article in the May 3 issue of Science, "Climate modelers grapple with their own emissions." Note that I am a member of both the American Association for The Advancement of Science (AAAS) which publishes Science, and the Institute of Electrical and Electronic Engineers (IEEE) which published IEEE Spectrum. These organizations' publications provide much content for my writing.

Over the decades, supercomputer simulations of Earth's climate have yielded unprecedented insights into how the interplay of atmosphere, ocean, and land shapes the planet's response to rising levels of greenhouse gases. But as these climate models have grown in complexity, researchers have started to worry the simulations have a substantial climate footprint of their own. Running them can take weeks or longer on a supercomputer, consuming megawatts of electric power—some of it from fossil fuels.[1]

With plans taking shape for a new round of modeling for the next U.N. climate assessment, scientists are seeking to better measure the pollution their simulations create—and curb modeling that might be excessive or redundant. "We are all aware that in everything we do we make a contribution to climate change," says Helene Hewitt, a climate scientist at the United Kingdom's Met Office and co-chair of the Coupled Model Intercomparison Project (CMIP), which coordinates global climate modeling. "It's only right that we think of our carbon budget across the board..."

2. How Climate Models Work

Climate models divide the atmosphere into thousands of boxes and use the equations of fluid dynamics to calculate how mass and energy move between them. The individual calculations aren't taxing. "Climate models are very sparse in computation," says Sarat Sreepathi, a computational earth scientist at Oak Ridge National Laboratory. But they require tracking reams of data about each of the 100-kilometer-wide boxes, meaning a model's performance depends as much on computer memory as on raw computing power. And given that models can run simulations 500 years into the future, with each year dependent on the one before it, there's a limit to how much code can be run in parallel to speed the process. "What everyone needs is a small supercomputer for a long time," Sreepathi says. The long run times push up power consumption and emissions.

3. How Much Greenhouse Gas Is Emitted?

Nearly 50 modeling centers worldwide contributed to the last round of CMIP, which ended in 2022, simulating hundreds of thousands of years and creating 40 petabytes of data. But fewer than a dozen centers tracked and reported the computing and energy resources used to run their experiments, and only eight provided enough to estimate their carbon footprint. Together those centers accounted for the release of nearly 1700 tons of carbon dioxide (CO_2), a study led by Mario Acosta, a climate modeler at the Barcelona Supercomputing Center, reported last month in Geoscientific Model Development.

---

[1] Paul Voosen, Science, "Climate modelers grapple with their own emissions," May 3, 2024, https://www.science.org/content/article/climate-modelers-grapple-their-own-carbon-emissions

In the grand scheme of global warming, those emissions are minimal, says Bryan Lawrence, a climate scientist at the U.K.'s National Centre for Atmospheric Science and study co-author—no more than what 200 average Americans emit in a year. Cryptocurrency mining, meanwhile, has emitted many millions of tons of CO_2. But that's no excuse not to measure the emissions, Lawrence says. "It is small beer," he says. "On the other hand, we may not need to do as much small beer."

Over the past few years, Acosta, Lawrence, and others have pushed centers to measure their efficiency and share best practices to boost it. During the last round of CMIP, they asked contributors to measure variables such as the energy used to simulate a year, how long each year takes to run, and the memory used. But it was tough to get traction, especially outside Europe. In some cases, groups may have been reluctant share proprietary information; more often they simply had little interest. "They are focused on the scientific problems," Acosta says. "It is difficult to motivate people to collect this stuff."

4. Potential for Conservation

Still, Acosta and his co-authors found that experiments often used half or less of their computer's theoretical daily capacity, squandering computing power. Some of the waste might reflect time needed to switch between different experiments, but errors in the models' code or the supercomputer's workflow routines might also be reducing efficiency. Fixing such errors could save significant energy and human resources, Acosta says.

Making hard choices about just how many model runs are needed could also limit emissions, Lawrence says. The last round of CMIP proposed 190 different modeling experiments for centers to run—not just to provide baseline climate projections, but also to study aspects of climate change from geoengineering to paleoclimate. Many centers scrambled to do as many runs as possible, he says. "Modeling centers feel like they have to do [all of] CMIP to be in the game. It should not mean you have to do every bloody experiment."

CMIP now plans to endorse a much more limited core group of model runs, while also helping coordinate community-led experiments. "We're trying to do this as lean as possible while still maintaining scientific quality and avoiding a huge overload of emissions," says Eleanor O'Rourke, director of the CMIP International Project Office. CMIP will also encourage centers to collect the data needed to estimate their carbon footprint, including emissions from staff travel.

Acosta thinks there is a moral imperative for modeling centers to practice the carbon discipline they preach. "We cannot say [other] people should reduce travel emissions while we are spending hours and hours on our supercomputers to run these scenarios."

5. The Proposal

Final author's comment: The one thing I didn't see in the above model was what percentage of CMIP and other related modeling U.S. supercomputer fleets perform, and what type of facilities are used. However, a bit of research and the site linked below led me to believe that U.S. personnel and institutions are heavily involved.

https://wcrp-cmip.org/cmip-governance/cmip-panel/

The CMIP Panel is a Working Group on Climate Modelling (WGCM) subcommittee charged with overseeing the evolving design of the CMIP experiments and with coordinating the various activities contributing to the science of CMIP. It works with the scientific leaders of specialized model intercomparisons that contribute to CMIP, and it has traditionally been responsible for identifying experts to develop the forcing datasets needed for some of the experiments.

There are 29 CMIP panel members as of March 2023. Of these the following six (20%) are U.S. Panel Members, and their organization.

John Dunne, Co-Chair; U.S. Organization: National Oceanic & Atmospheric Administration (NOAA), Geophysical Fluid Dynamics Laboratory

Robert Pincus, Core Panel, Data Access Co-Lead; U.S. Organization: Lamont-Doherty Earth Observatory, Columbia University

Paul Durack, Working Group on Climate Modeling Infrastructure Panel Co-Chair, Climate Forcings Co-Lead; U.S. Organization: Lawrence Livermore National Lab (LLNL), Program for Climate Model Diagnosis & Intercomparison

Vaishali Naik, Climate Forcings Co-Lead; U.S. Organization: NOAA

Sasha Ames, Data Citation Co-Lead; U.S. Organization: LLNL

Forrest Hoffman, Model Benchmarking Co-Lead; U.S. Organization: Oak Ridge National Lab

Assuming that U.S. government labs do much of the models' heavy lifting, the U.S. Department of Energy and Department of Commerce (for NOAA) controls these and would be ideally suited to coordinate efforts to conserve energy and reduce greenhouse gas emissions. My proposal follows.

The electric utility grid management organizations in the U.S., generally allow each major electricity-user to select the energy-provider(s) from which it is supplied during a given time-frame. It is suggested that any supercomputer fleet that runs climate models should, as a minimum have energy provided from 100% zero-carbon sources (these include renewables, nuclear and large hydro).at the times they are running climate models. Preferably, these fleets should be supplied with zero-carbon energy at all times.

The benefits of doing this include:

1. Assuming this increases demand for zero-carbon generation, this will drive expansion of this resource.

2. This will avoid making the problem that these models are attempting to understand worse, or at least the perception that this is happening.

3. Since most of these fleets are in government-sponsored facilities, this will help drive the federal governments' consumption to zero-carbon sources.