

Researcher: Climate models can run for months on supercomputers—but my new algorithm can make them ten times faster

May 4 2024, by Samar Khatiwala



Credit: Pixabay/CC0 Public Domain

Climate models are some of the most complex pieces of software ever written, able to simulate a vast number of different parts of the overall system, such as the atmosphere or ocean. Many have been developed by hundreds of scientists over decades and are constantly being added to

and refined. They can run to over a million lines of computer code—tens of thousands of printed pages.

Not surprisingly, these models are expensive. The simulations take time, frequently several months, and the supercomputers on which the models are run consume a lot of energy. But a new algorithm [I have developed](#) promises to make many of these climate model simulations ten times faster, and could ultimately be an important tool in the fight against climate change.

One reason climate modeling takes so long is that some of the processes being simulated are intrinsically slow. The ocean is a good example. It takes a few thousand years for water to circulate from the surface to the [deep ocean](#) and back (by contrast, the atmosphere has a "mixing time" of weeks).

Ever since the first climate models were developed in the 1970s, scientists realized this was going to be a problem. To use a model to simulate [climate change](#), it has to be started from conditions representative of before industrialization led to the release of greenhouse gases into the atmosphere.

To produce such a stable equilibrium, scientists "spin-up" their model by essentially letting it run until it stops changing (the system is so complex that, as in the real world, some fluctuations will always be present).

An initial condition with minimal "drift" is essential to accurately simulate the effects of human-made factors on the climate. But thanks to the ocean and other sluggish components this can take several months even on large supercomputers. No wonder climate scientists have called this bottleneck one of the ["grand challenges"](#) of their field.

Can't just throw more computers at the problem

You might ask, "why not use an even bigger machine?" Unfortunately, it wouldn't help. Simplistically, supercomputers are just thousands of individual computer chips, each with dozens of processing units (CPUs or "cores") connected to each other via a high-speed network.

[One of the machines](#) I use has over 300,000 cores and can perform almost 20 quadrillion arithmetic operations per second. (Obviously, it is shared by hundreds of users and any single simulation will only use a tiny fraction of the machine.)

A climate model exploits this by subdividing the surface of the planet into smaller regions—subdomains—with calculations for each region being performed simultaneously on a different CPU. In principle, the more subdomains you have the less time it should take to perform the calculations.

That is true up to a point. The problem is that the different subdomains need to "know" what is happening in adjoining ones, which requires transmitting information between chips. That is much slower than the speed with which modern chips can perform arithmetic calculations, what computer scientists call "bandwidth limitation." (Anyone who has tried to stream a video over a slow internet connection will know what that means.)

There are therefore diminishing returns from throwing more computing power at the problem. Ocean models especially suffer from such poor "scaling."

Ten times faster

This is where the new computer algorithm that I've developed and published in [Science Advances](#) comes in. It promises to dramatically

reduce the spin-up time of the ocean and other components of earth system models. In tests on typical [climate models](#), the algorithm was on average about ten times faster than current approaches, decreasing the time from many months to a week.

The time and energy this could save [climate scientists](#) is valuable in itself. But being able to spin-up models quickly also means that scientists can calibrate them against what we know actually happened in the real world, improving their accuracy, or to better define the uncertainty in their climate projections. Spin-ups are so time consuming that neither are currently feasible.

The new algorithm will also allow us to perform simulations in more spatial detail. Currently, ocean models typically don't tell us anything about features smaller than 1° width in longitude and latitude (about 110km at the equator). But many critical phenomena in the ocean occur at far smaller scales—tens of meters to a few kilometers—and higher spatial resolution will certainly lead to more accurate climate projections, for instance of sea level rise, storm surges and hurricane intensity.

How it works

Like so much "new" research it is based on an old idea, in this case one that goes back centuries to the Swiss mathematician Leonhard Euler. Called "sequence acceleration," you can think of it as using past information to extrapolate to a "better" future.

Among other applications, it is widely used by chemists and material scientists to calculate the structure of atoms and molecules, a problem that happens to take up more than half the world's supercomputing resources.

Sequence acceleration is useful when a problem is iterative in nature, exactly what climate model spin-up is: you feed the output from the model back as an input to the model. Rinse and repeat until the output becomes equal to the input and you've found your equilibrium solution.

In the 1960s Harvard mathematician D.G. Anderson came up with a clever way to combine multiple previous outputs into a single input so that you get to the final solution with far fewer repeats of the procedure. About ten times fewer as I found when I applied his scheme to the spin-up problem.

Developing a [new algorithm](#) is the easy part. Getting others to use it is often the bigger challenge. It is therefore promising that the UK Met Office and other climate modeling centers are trying it out.

The next major IPCC report is due in 2029. That seems like a long way off but given the time it takes to develop models and perform simulations, preparations are already underway. Coordinated by an [international collaboration](#) known as the [Coupled Model Intercomparison Project](#), it is these simulations that will form the basis for the report. It is exciting to think that my algorithm and software might contribute.

More information: Samar Khatiwala, Efficient spin-up of Earth System Models using sequence acceleration, *Science Advances* (2024). [DOI: 10.1126/sciadv.adn2839](https://doi.org/10.1126/sciadv.adn2839)

This article is republished from [The Conversation](#) under a Creative Commons license. Read the [original article](#).

Provided by The Conversation

Citation: Researcher: Climate models can run for months on supercomputers—but my new algorithm can make them ten times faster (2024, May 4) retrieved 19 May 2024 from <https://phys.org/news/2024-05-climate-months-supercomputers-algorithm-ten.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.