Title: Tim Palmer: Climate Change, Chaos and Inexact Computing

Date: May 04, 2016 07:00 PM

URL: http://pirsa.org/16050018

Abstract: <p>How well can we predict our future climate? If the flap of a butterflyâ€™s wings can change the course of weather a week or so from now, what hope trying to predict anything about our climate a hundred years hence? In this talk I will discuss the science of climate change from a perspective which emphasises the chaotic (and hence uncertain) nature of our climate system. In so doing I will outline the fundamentals of climate modelling, and discuss the emerging concept of inexact supercomputing, needed - paradoxically perhaps - if we are to increase the accuracy of predictions from these models. Indeed, revising the notion of a supercomputer from its traditional role as a fast but precise deterministic calculating machine, may be important not only for climate prediction, but also for other areas of science such as astrophysics, cosmology and neuroscience.</p>

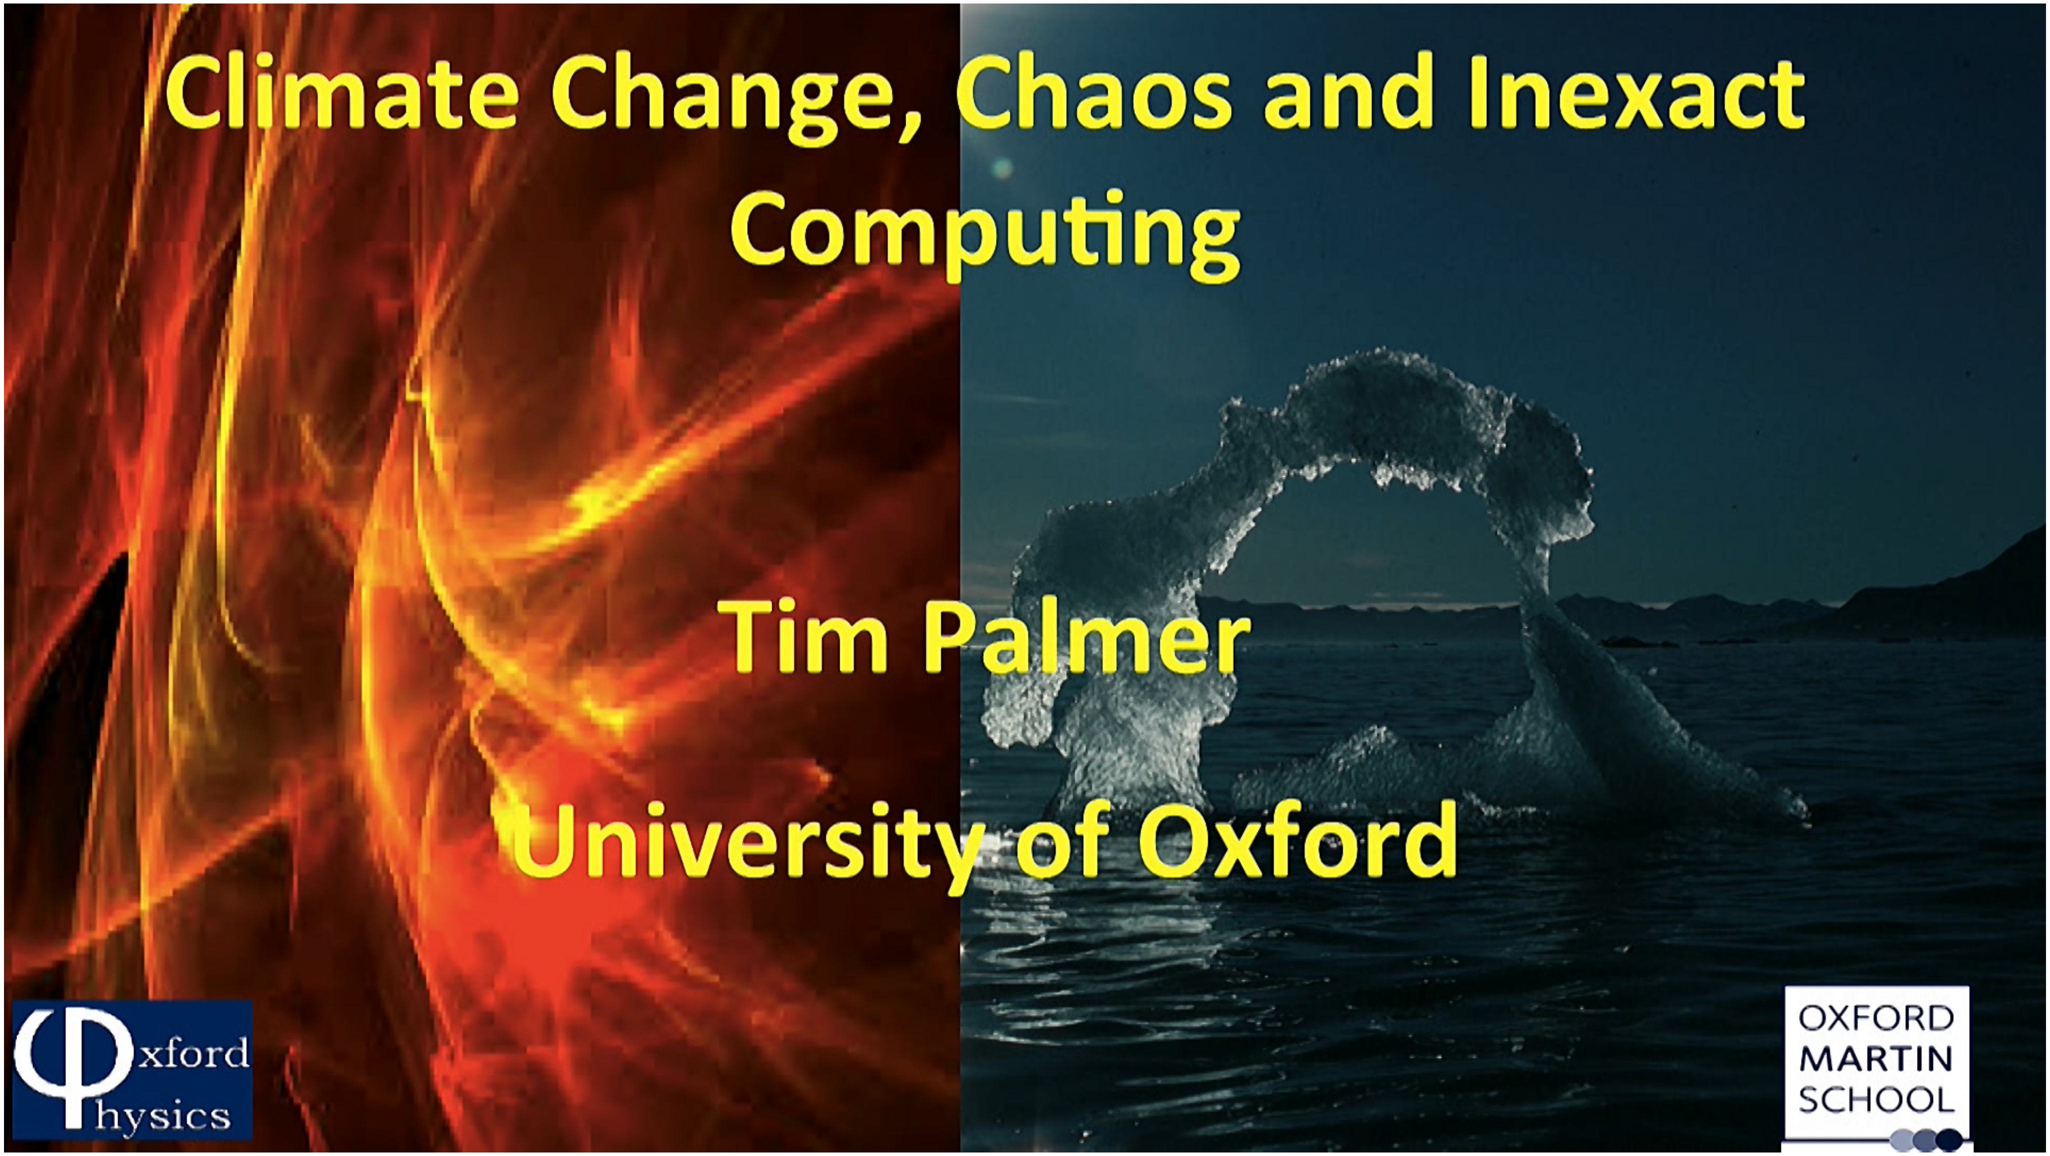
Climate Change, Chaos and Inexact Computing
Tim Palmer
University of Oxford
Oxford Physics
OXFORD MARTIN SCHOOL

forecast 13 October 1987, 12UTC +0 h

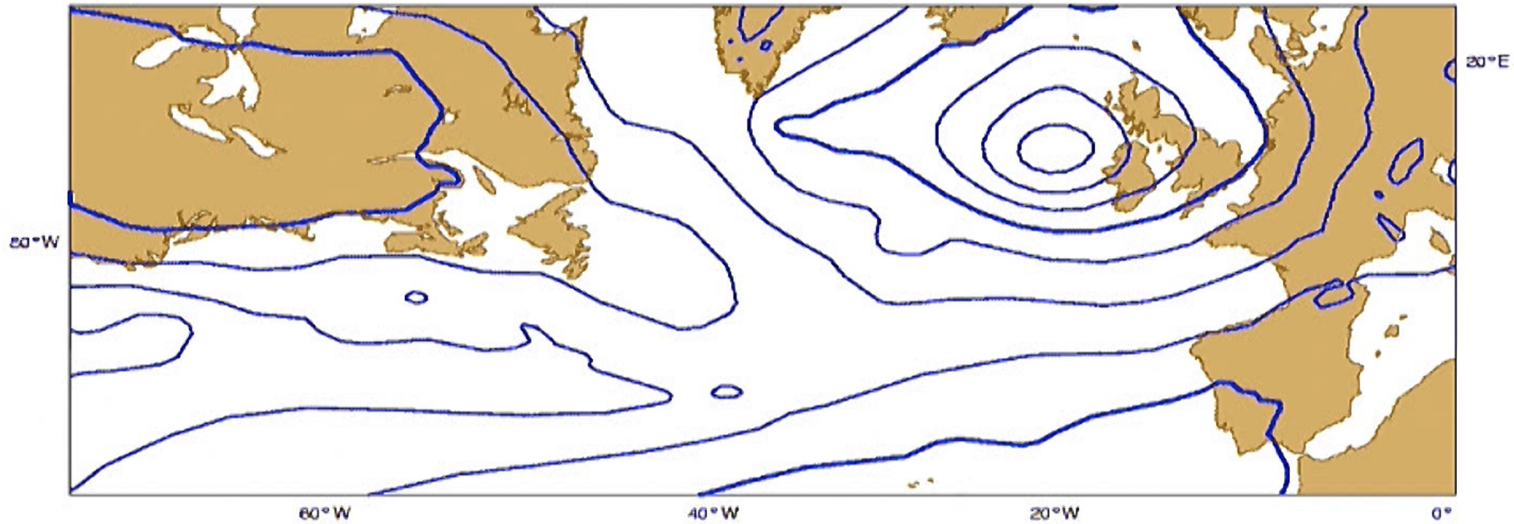

forecast 13 October 1987, 12UTC +0 h

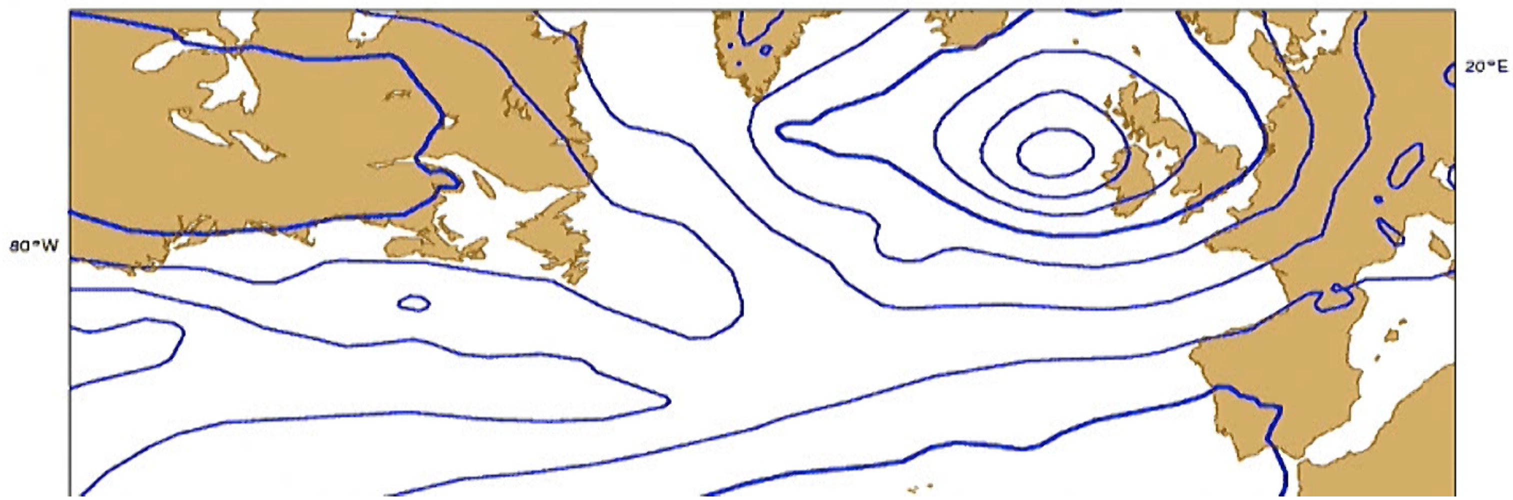

forecast 13 October 1987, 12UTC +29 h

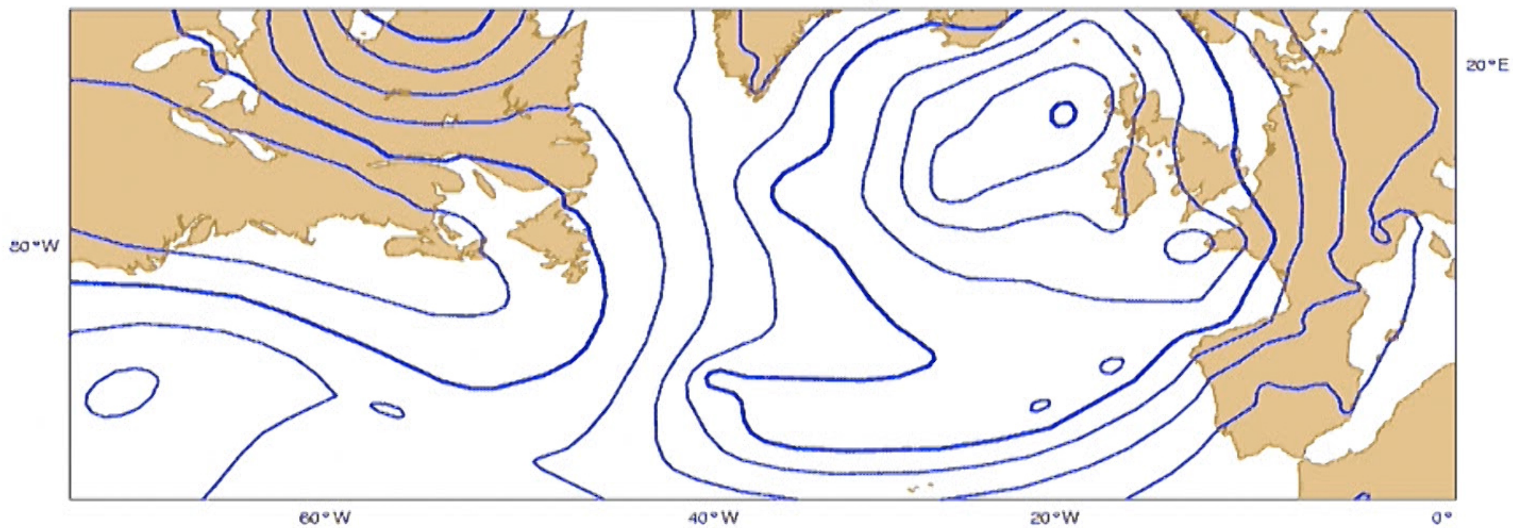

forecast 13 October 1987, 12UTC +29 h

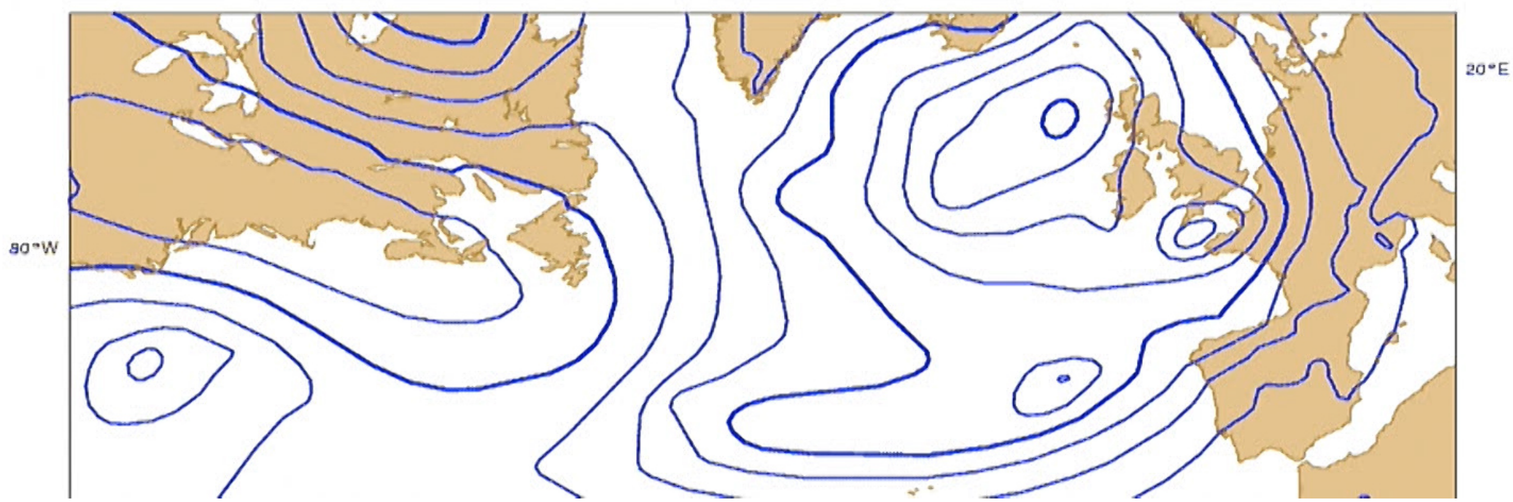

forecast 13 October 1987, 12UTC +66 h

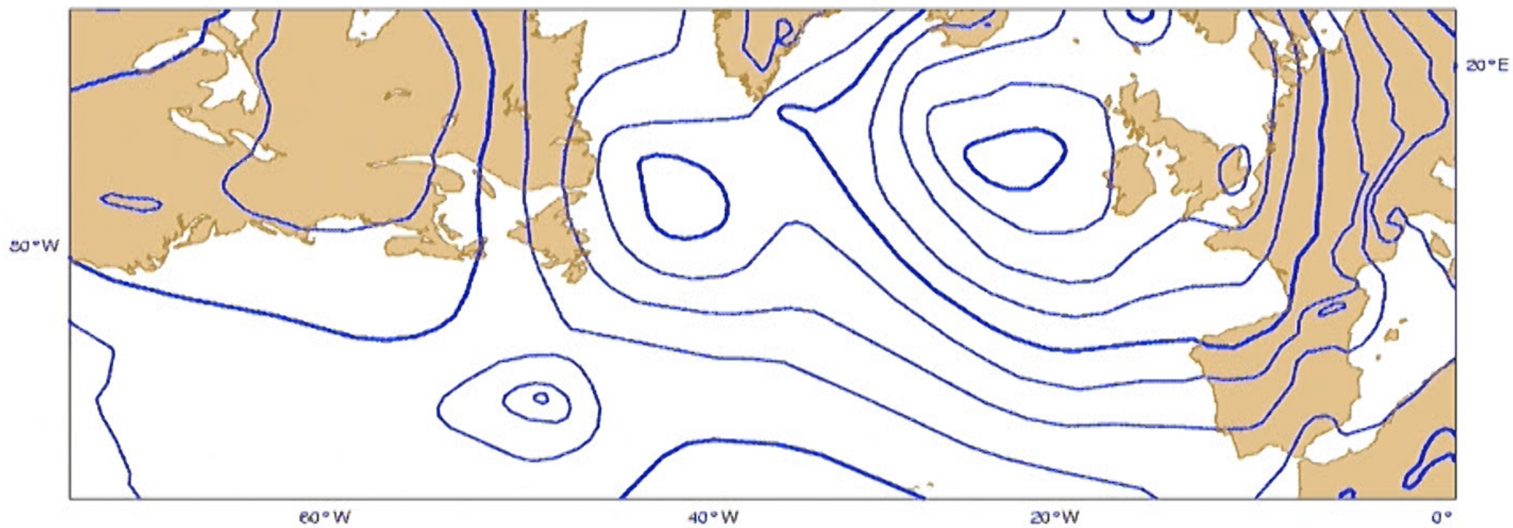

forecast 13 October 1987, 12UTC +66 h

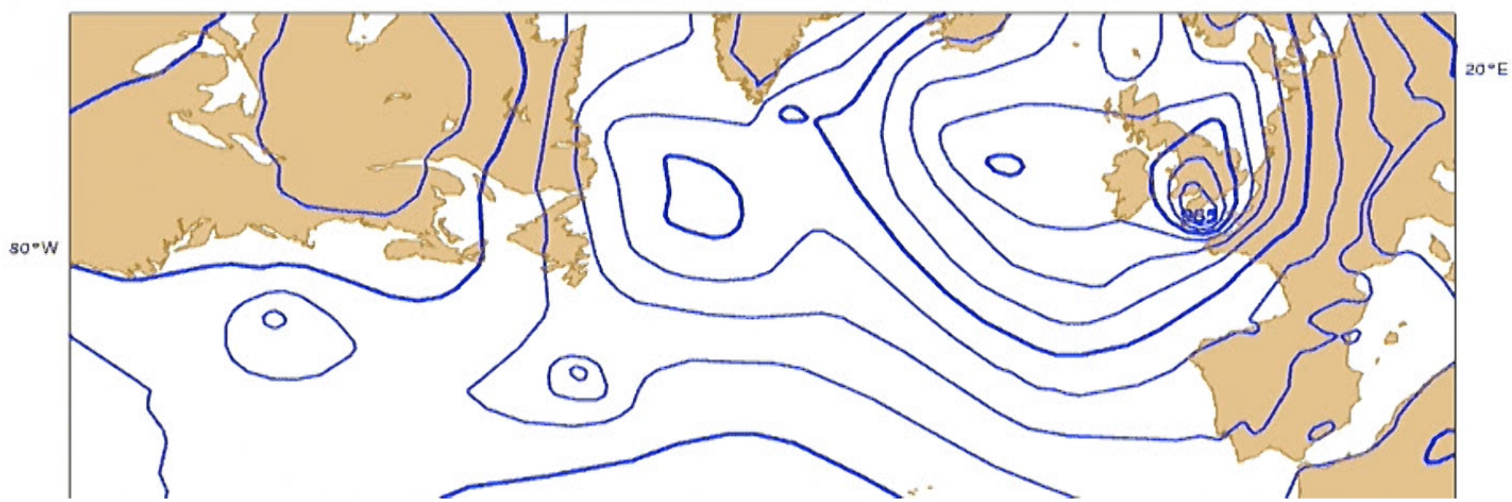

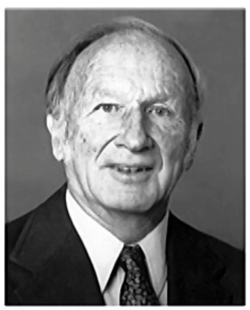

$$\frac{dX}{dt} = -\sigma X + \sigma Y$$

$$\frac{dY}{dt} = -XZ + rX - Y$$

$$\frac{dZ}{dt} = XY - bZ$$

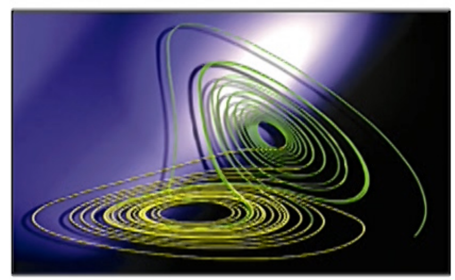

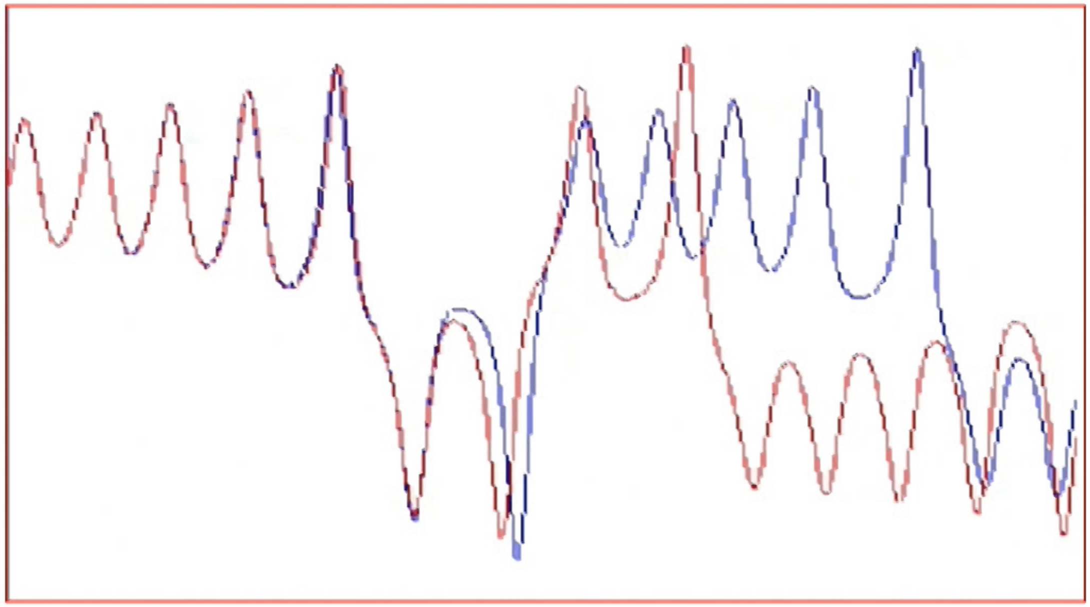

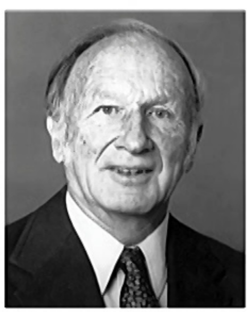

$$\frac{dX}{dt} = -\sigma X + \sigma Y$$

$$\frac{dY}{dt} = -XZ + rX - Y$$

$$\frac{dZ}{dt} = XY - bZ$$

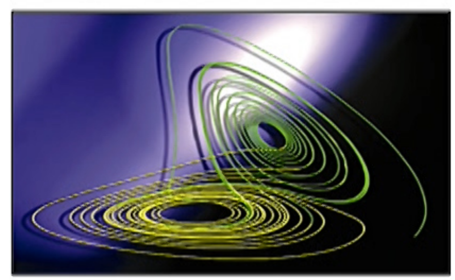

The Telegraph

Home News World Sport Finance Comment Culture Travel Life Women Fashion Luxur
Film Oscars Music Art Books TV and Radio Theatre Comedy Dance Opera Photogr
Book Reviews | Fiction | Non-Fiction | Children's Books | Short Story Club | Book Shop | Ways Wit

HOME » CULTURE » BOOKS » NON-FICTION REVIEWS

The game is up for climate change believers

Charles Moore reviews The Age of Global Warming by Rupert Darwall (Quartet)

Power station emitting steam and smoke Photo: Reuters

By Charles Moore
9:42PM BST 06 Apr 2014
Follow 5,404 followers

Print this article
Share 5K

By Charles Moore
9:42PM BST 06 Apr 2014
Follow 5,404 followers

4870 Comments

Most of us pay some attention to the weather forecast. If it says it will rain in your area tomorrow, it probably will. But if it says the same for a month, let alone a year, later, it is much less likely to be right. There are too many imponderables.

The theory of global warming is a gigantic weather forecast for a century or more. However interesting the scientific inquiries involved, therefore, it can have almost no value as a prediction. Yet it is as a prediction that global warming (or, as we are now ordered to call it in the face of a stubbornly parky 21st century, "global weirding") has captured the political and bureaucratic elites. All the action plans, taxes, green levies, protocols and carbon-emitting flights to massive summit meetings, after all, are not because of what its supporters call "The Science". Proper science studies what is – which is, in principle, knowable – and is consequently very cautious about the future – which isn't. No, they are the result of a belief that something big and bad is going to hit us one of these days.

"The theory of global warming is a gigantic weather forecast for a century or more....therefore it can have almost no value as a prediction "

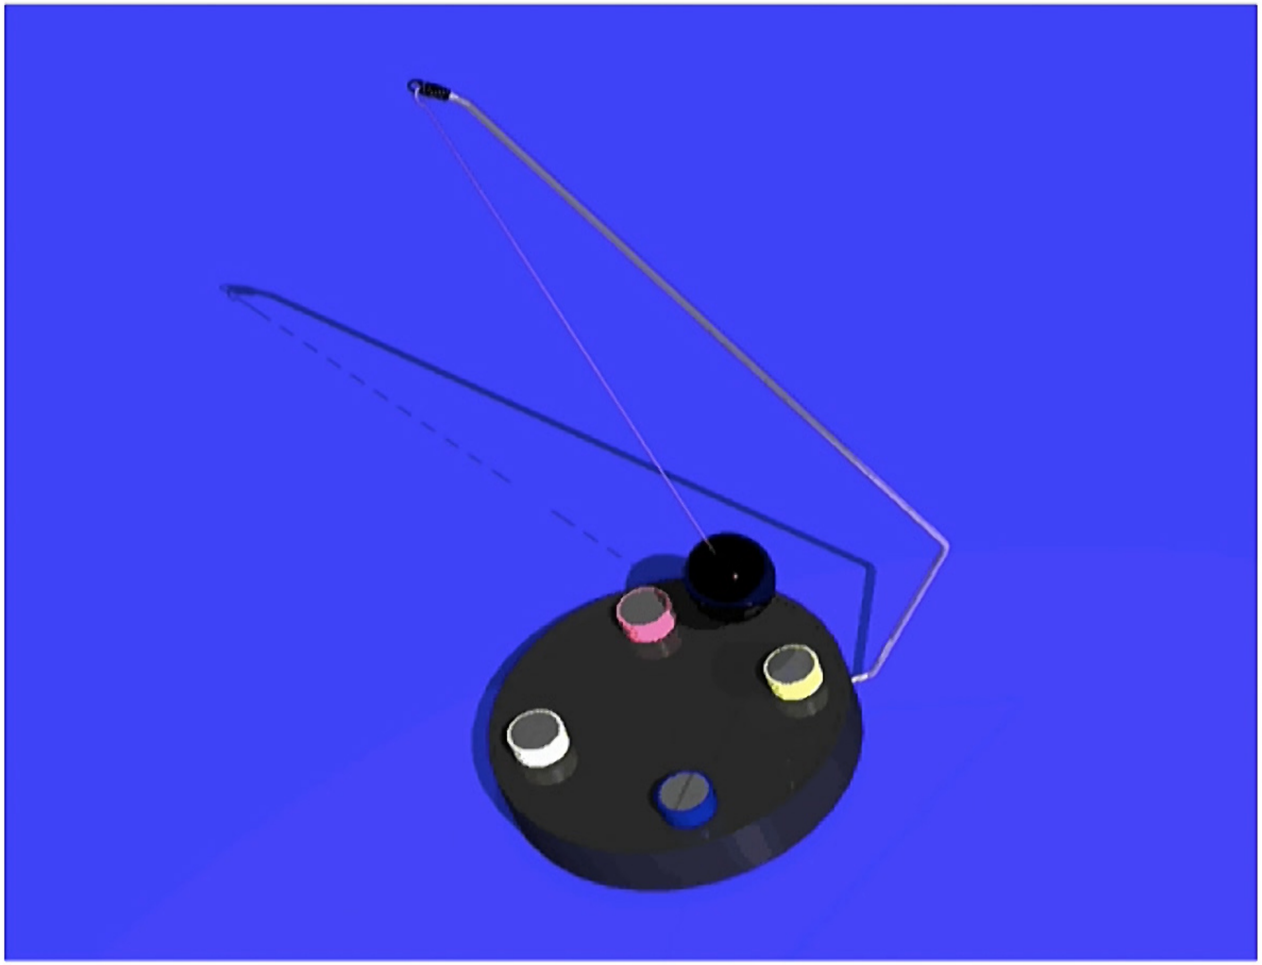

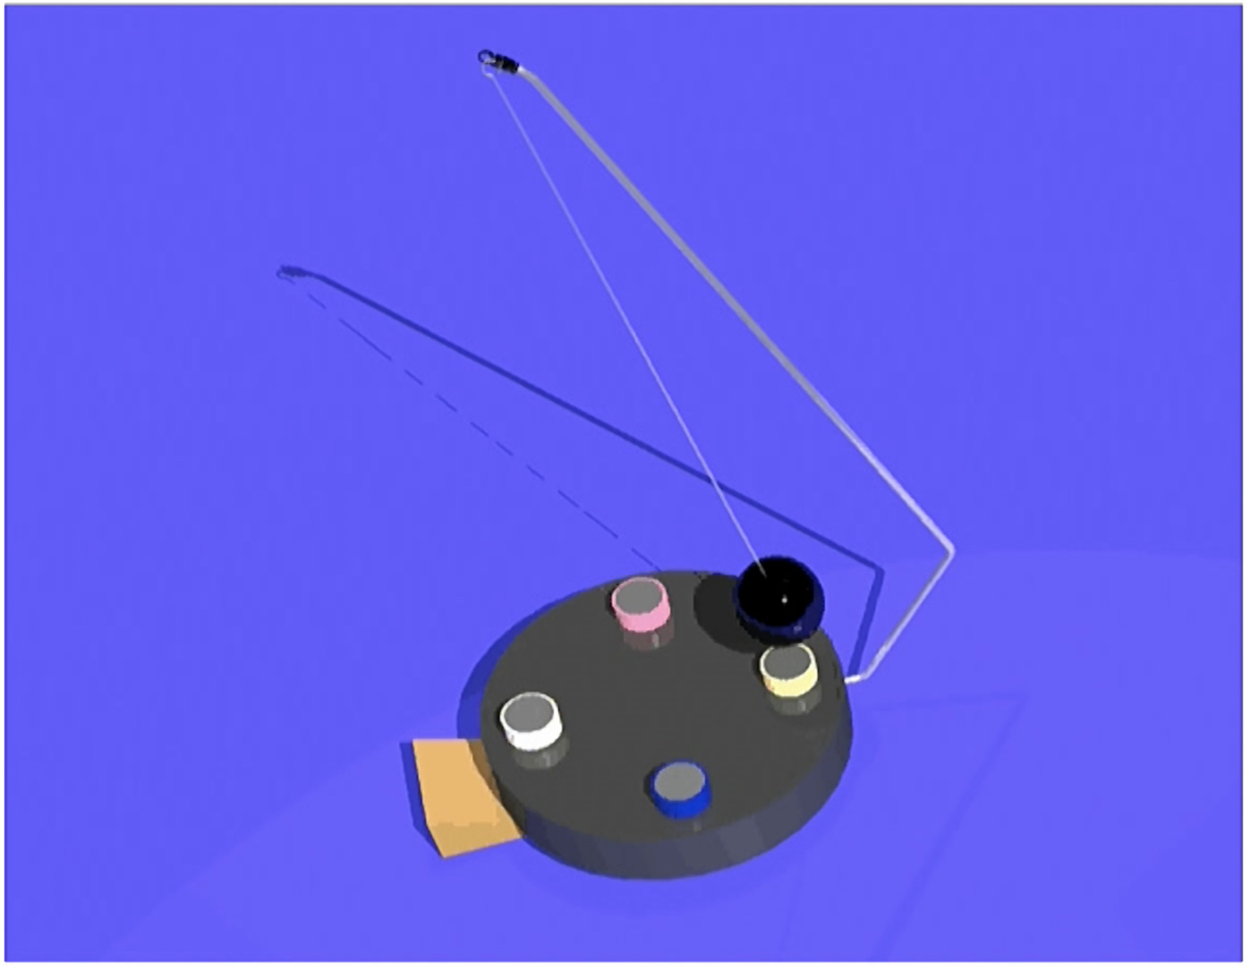

Global Land and Ocean Temperature Anomalies, January-December

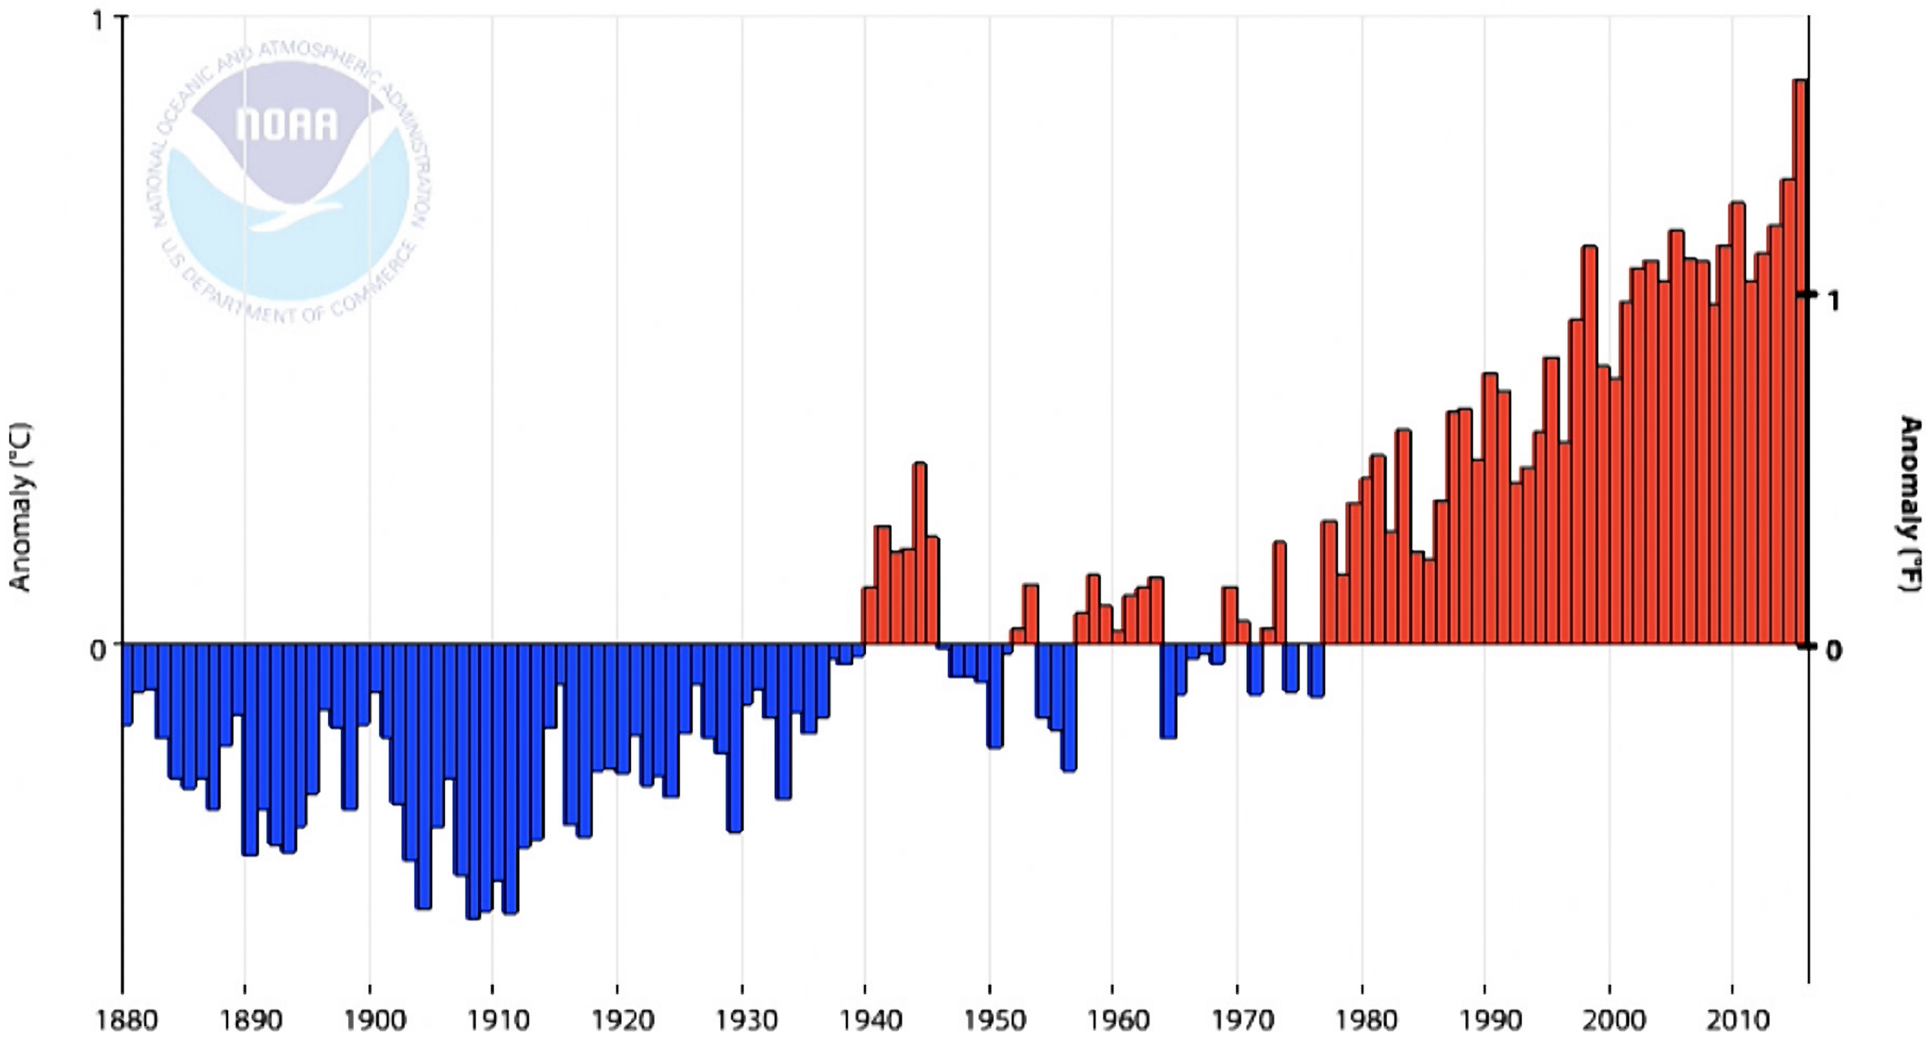

The Key Question that we can't easily answer:

How Big is the Wedge?

Potential Amplifiers of the effect of anthropogenic carbon emissions

Water Vapour

Clouds

Carbon cycle

Ice albedo

Methane

Potential Amplifiers of the effect of anthropogenic carbon emissions

Water Vapour

Clouds

Carbon cycle

Ice albedo

Methane

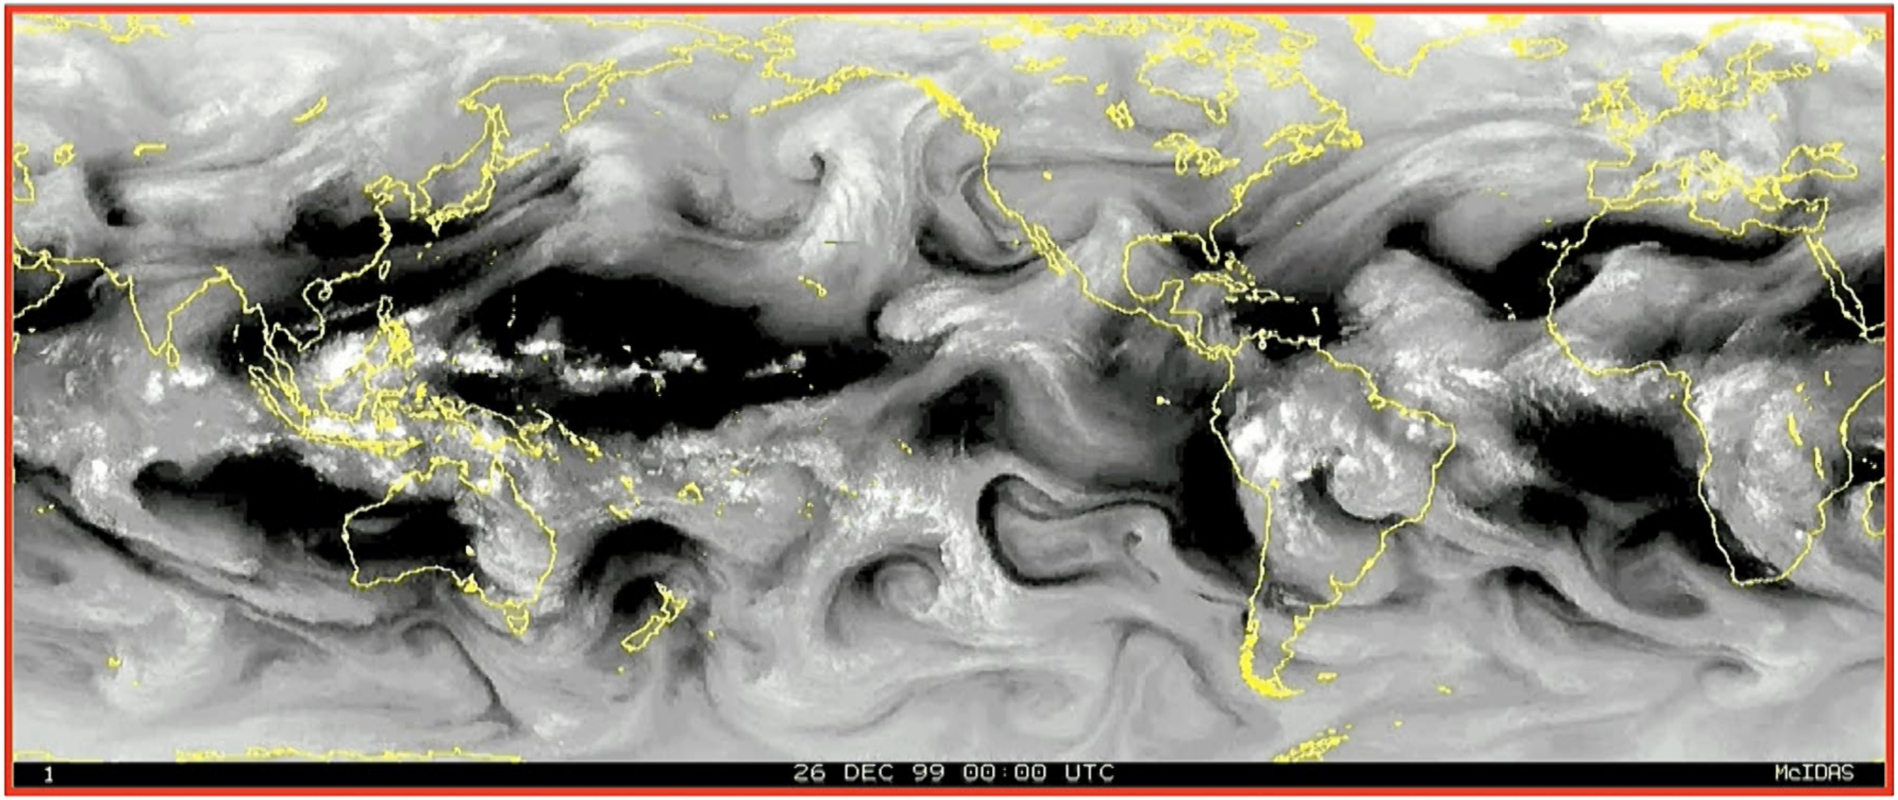
1
26 DEC 99 00:00 UTC
McIDAS

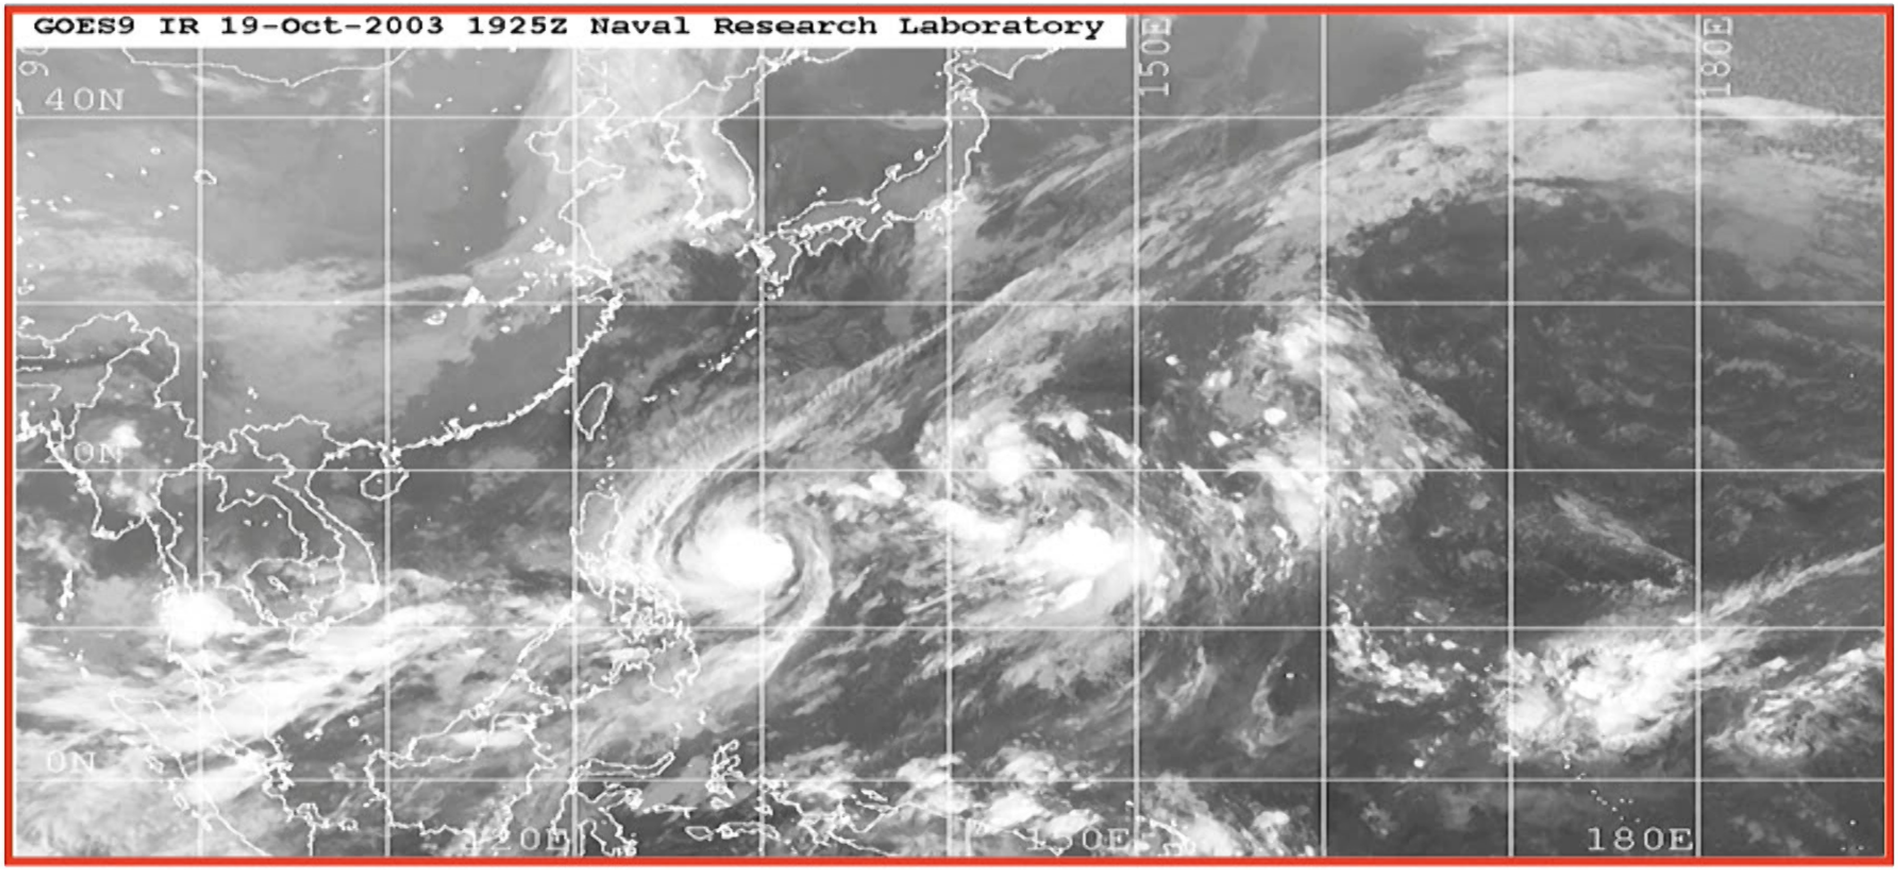
GOES9 IR 19-Oct-2003 1925Z Naval Research Laboratory
40N
150E
180E
120E
180E

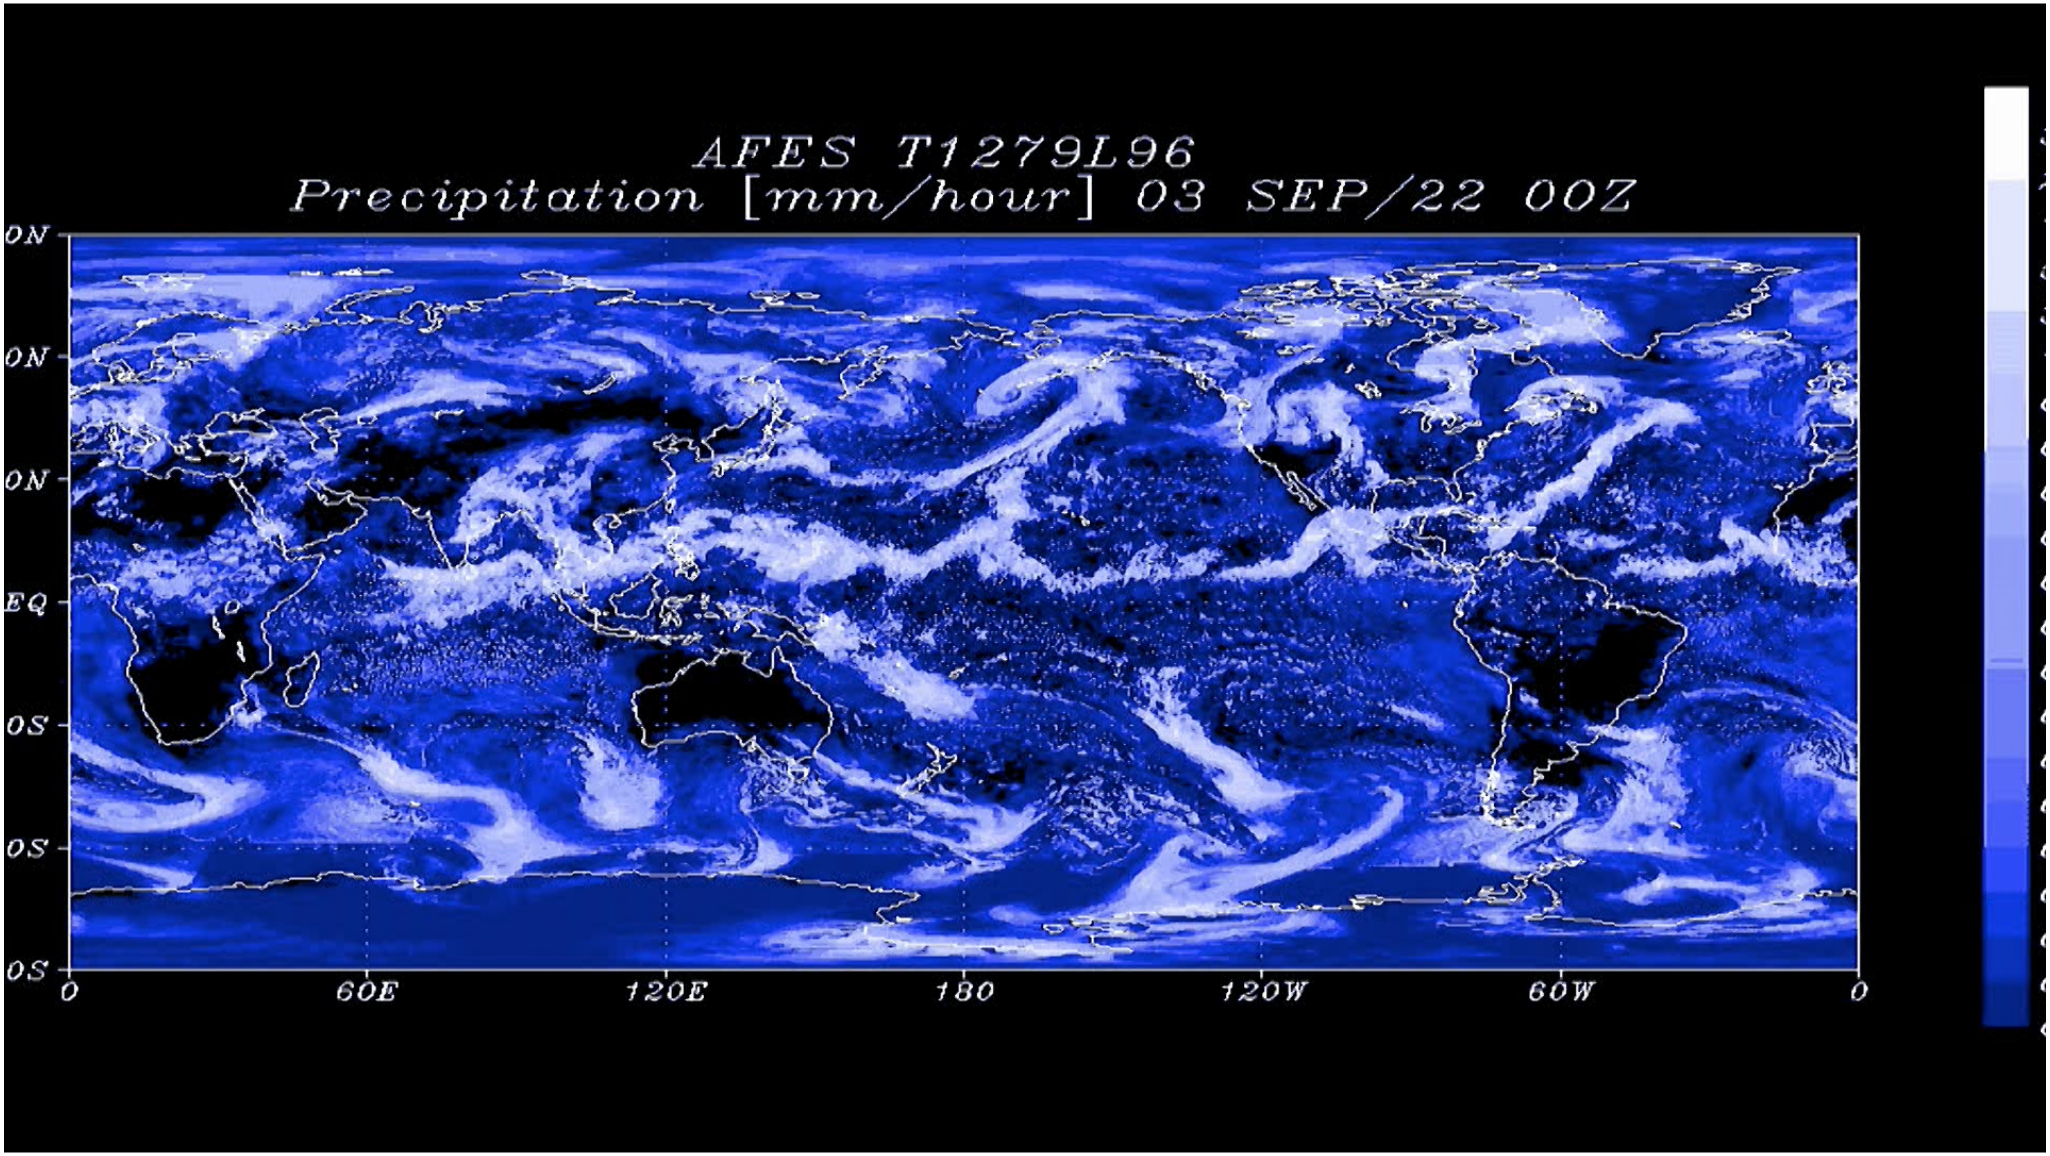
AFES T1279L96
Precipitation [mm/hour] 03 SEP/22 00Z
EQ
0
60E
120E
180
120W
60W
0

Comprehensive climate models are based on the primitive laws of physics eg

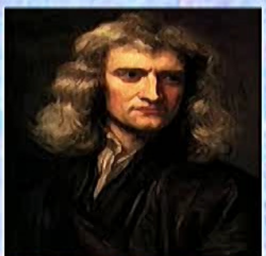

$$\mathbf{F} = m\mathbf{a}$$

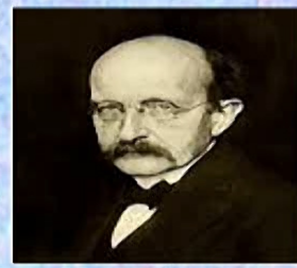

$$E = \hbar\omega$$

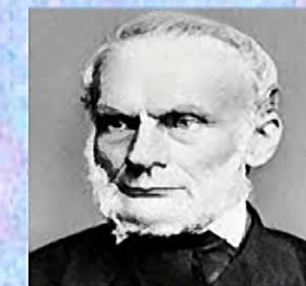

$$\delta Q = TdS$$

$$\rho\left(\frac{\partial}{\partial t}+\mathbf{u}.\nabla\right)\mathbf{u}=\rho\mathbf{g}-\nabla p+\mu\ \nabla^2\mathbf{u}$$

$$\rho\left(\frac{\partial}{\partial t}+\mathbf{u}.\nabla\right)\mathbf{u}=\rho\mathbf{g}-\nabla p+\mu\ \nabla^2\mathbf{u}$$

Unpacks into billions of individual equations, describing scales of motion from planetary scales to microscopic scales.

Even the world's biggest computers aren't big enough to represent all scales of motion

Simplified approximate formulae to describe the effect of atmospheric processes (eg clouds) that the simulator can't resolve

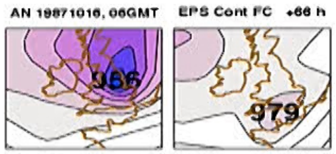

Oct 87 Storm Ensemble Reforecast

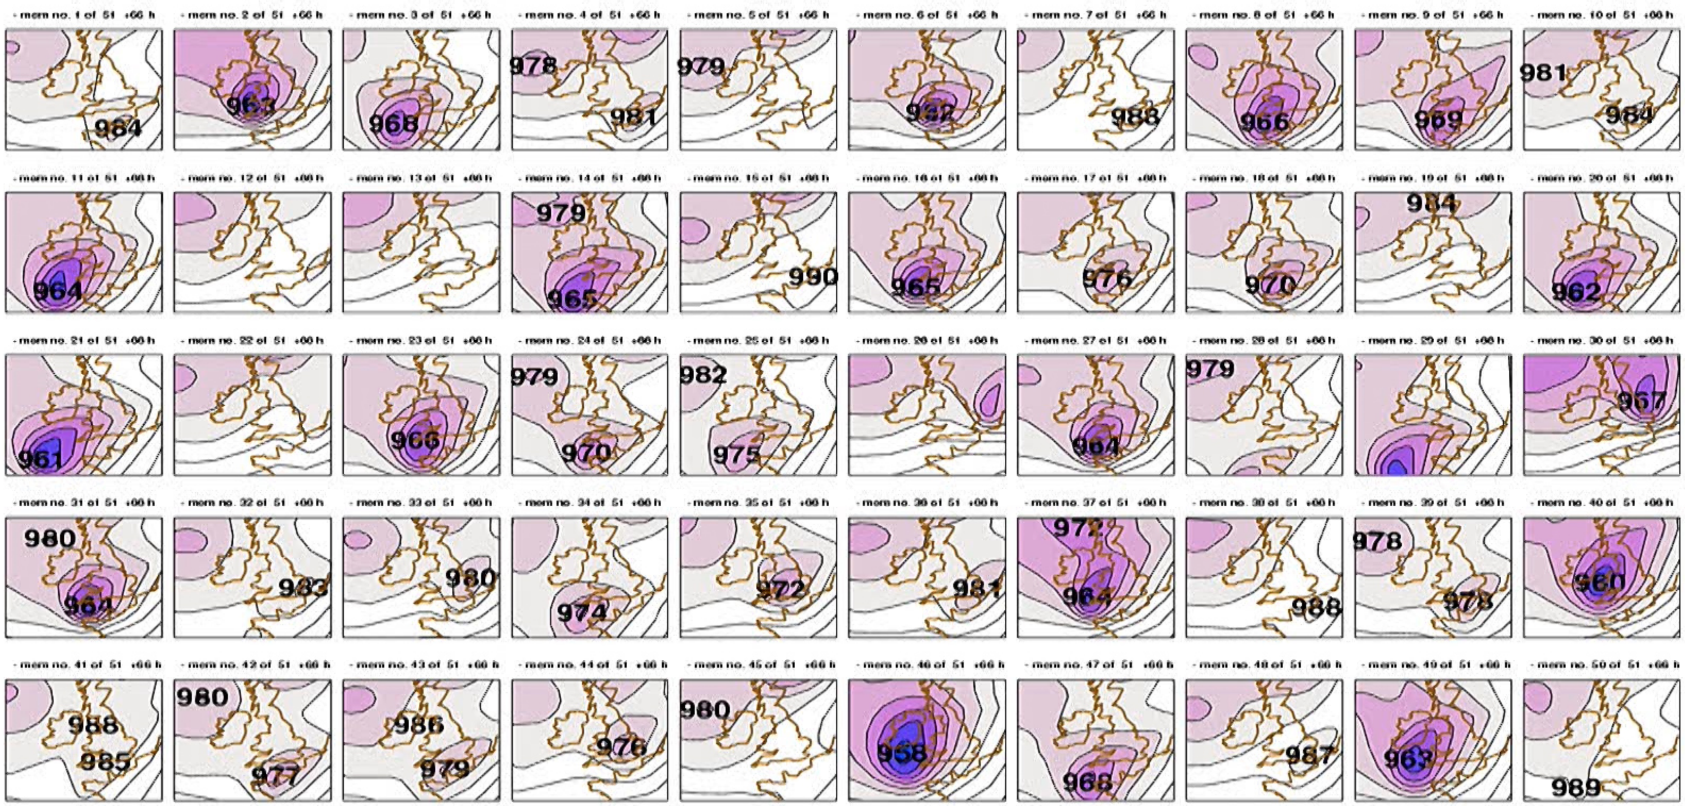

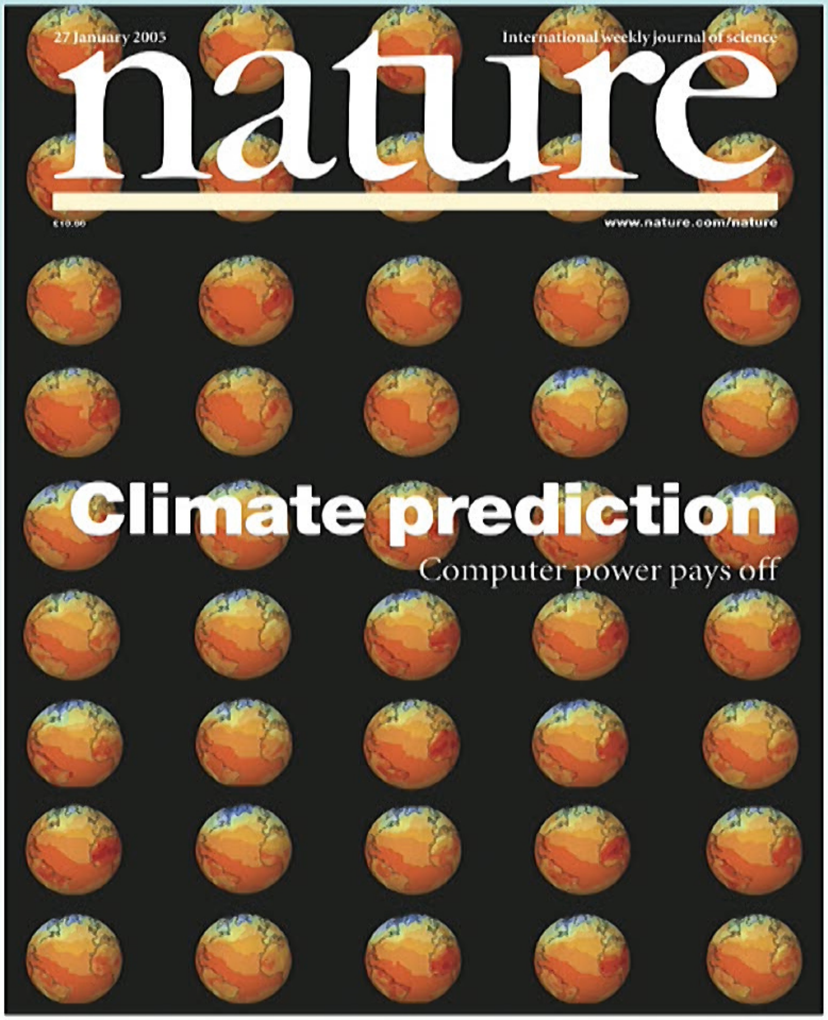

Like modern weather forecasting, climate predictions are based on ensembles of forecasts using slightly different computer model equations ie different

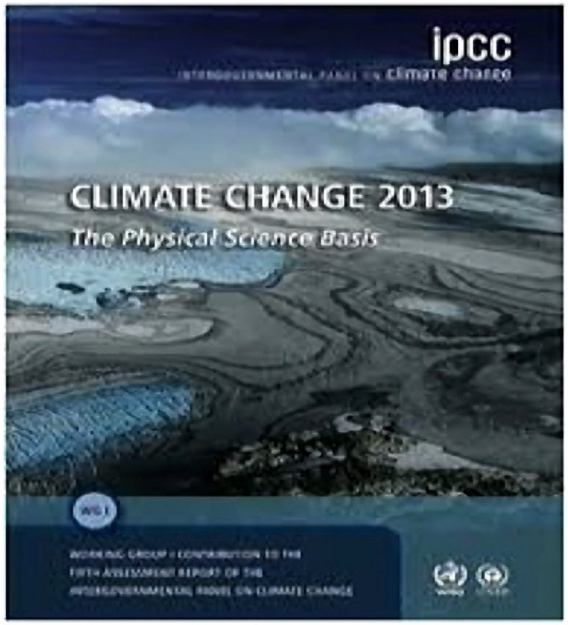
ipcc
climate change
CLIMATE CHANGE 2013
The Physical Science Basis

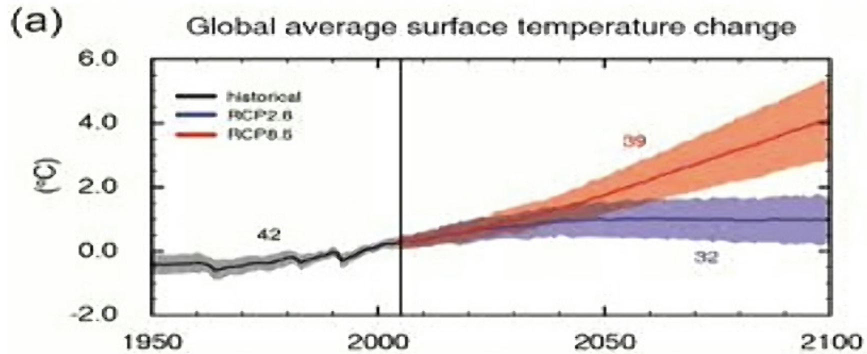
(a)
Global average surface temperature change
6.0
4.0
2.0
0.0
-2.0
(°C)
Historical
RCP2.6
RCP8.5
39
42
32
1950
2000
2050
2100

The fate of humanity depends on how clouds will respond to our emissions of CO_2

This is therefore an issue we have to understand better!

How?

SCIENCEINSIDER

Breaking news and analysis from the world of science policy

OAK RIDGE NATIONAL LABORATORY/FLICKR (CC BY-NC-ND 2.0)

The Titan supercomputer at Oak Ridge National Laboratory in Tennessee is currently the world's second fastest computer.

Obama orders effort to build first exascale computer

Tweet | Share | G+1 | 9

Staff Writer

Email Robert

By Robert F. Service | 30 July 2015 5:15 pm | 7 Comments

The United States is now committed to building an exascale computer, some 30 times more powerful than today's top machine. Yesterday, President Barack Obama signed an executive order creating a national strategic computing initiative, which aims to coordinate high-performance computing research and development between federal agencies. The order should make it easier for agencies to justify increasing their budget requests to Congress for supercomputing R&D.

"This is an extremely important step for high performance computing in the U.S.," says Horst Simon, deputy director of the Lawrence Berkeley National

---

the WHITE HOUSE PRESIDENT BARACK OBAMA | Contact Us › | Get I

BRIEFING ROOM | ISSUES | THE ADMINISTRATION | PARTICIPATE | 1600 PENN

HOME · BRIEFING ROOM · PRESIDENTIAL ACTIONS · EXECUTIVE ORDERS

Briefing Room

- Your Weekly Address
- Speeches & Remarks
- Press Briefings
- Statements & Releases
- White House Schedule
- Presidential Actions
  - **Executive Orders**
  - Presidential Memoranda
  - Proclamations
- Legislation
- Nominations & Appointments
- Disclosures

The White House
Office of the Press Secretary

For Immediate Release July 29, 2015

Executive Order -- Creating a National Strategic Computing Initiative

EXECUTIVE ORDER

\- - - - - - -

CREATING A NATIONAL STRATEGIC COMPUTING INITIATIVE

By the authority vested in me as President by the Constitution and the laws of the United States of America, and to maximize benefits of high-performance computing (HPC) research, development, and deployment, it is hereby ordered as follows:

Section 1. Policy. In order to maximize the benefits of HPC for

---

Exaflop = 10^{18} floating point operations per second

$$\text{Eg} \ -1.2345 = -12345 \times 10^{-4}$$

The (scientific default) 64-bit floating point number

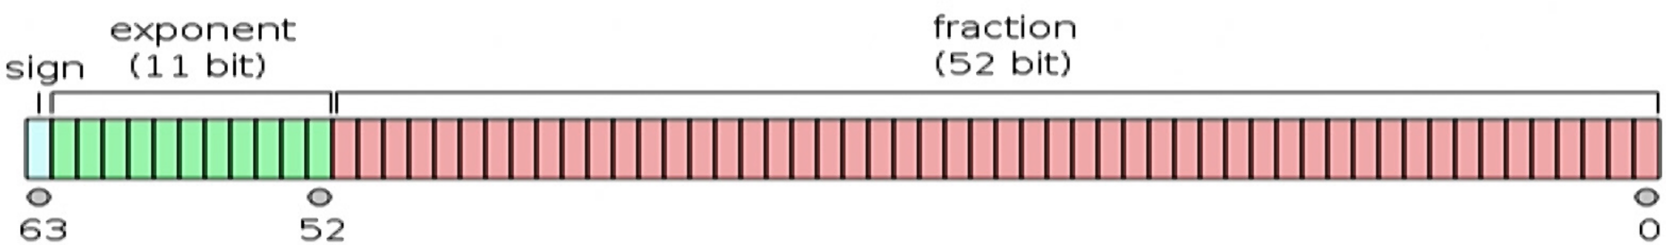

EXTREMETECH

Search Extremetech

Computing Mobile Internet Gaming Electronics Extreme

HOME › COMPUTING › SUPERCOMPUTING DIRECTOR BETS $2,000 THAT WE WON'T HAVE EXASCALE CO

Supercomputing director bets $2,000 that we won't have exascale computing by 2020

By Joel Hruska on May 17, 2013 at 10:11 am | 35 Comments

Share This article

Over the past year, we've covered a number of the challenges facing the supercomputing industry in its efforts to hit exascale compute levels by the end of the decade. The problem has been widely discussed at supercomputing conferences, so we're not surprised that Horst Simon, the Deputy Director at the Lawrence Berkeley National Laboratory's NERSC (National Energy Research Scientific Computing Center), has spent a significant amount of time talking about the problems with reaching exascale speeds.

But putting up $2000 of his own money in a bet that we *won't* hit exascale by 2020? That caught us off guard.

The exascale rethink

Simon lays out, in a 60+ page slideshow, why he doesn't think we'll hit the exascale threshold within seven years. The bottom line is this: Hitting exascale compute levels requires a fundamental rethink of virtually the entire computation process. One of the biggest problems standing in our way is power — not just the power required to run a task on a CPU, but the power required to share that data across the chip, node, and cluster. Data has to be written back to RAM, then shared across multiple systems. Caches must be kept coherent, calculation results written to storage, and new information loaded into RAM.

The Cost of Data Movement

"Hitting exascale compute levels requires a fundamental rethink of virtually the entire computation process. One of the biggest problems standing in our way is power – not just the power required to run a task on a CPU, but the power required to share that data across the chip, node and cluster.

It's theoretically possible to hit exascale computing with current levels of technology if you can afford to dedicate 100MW of power to the task, but the challenge is to bring exascale into the 10-20MW range"

1MW=$$10^6$ per year

$$\rho\left(\frac{\partial}{\partial t}+\mathbf{u}.\nabla\right)\mathbf{u}=\rho\mathbf{g}-\nabla p+\mu\ \nabla^2\mathbf{u}$$

Do we need 64 bits/variable?

What is the real information stored in the bits as a function of scale?

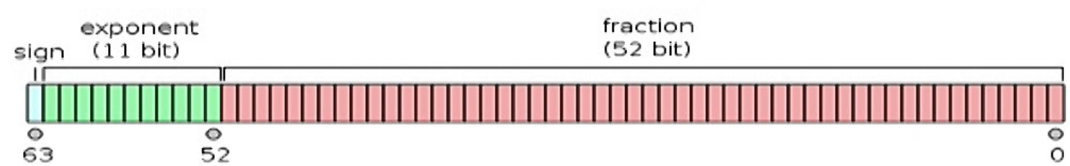

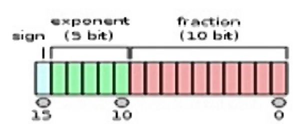

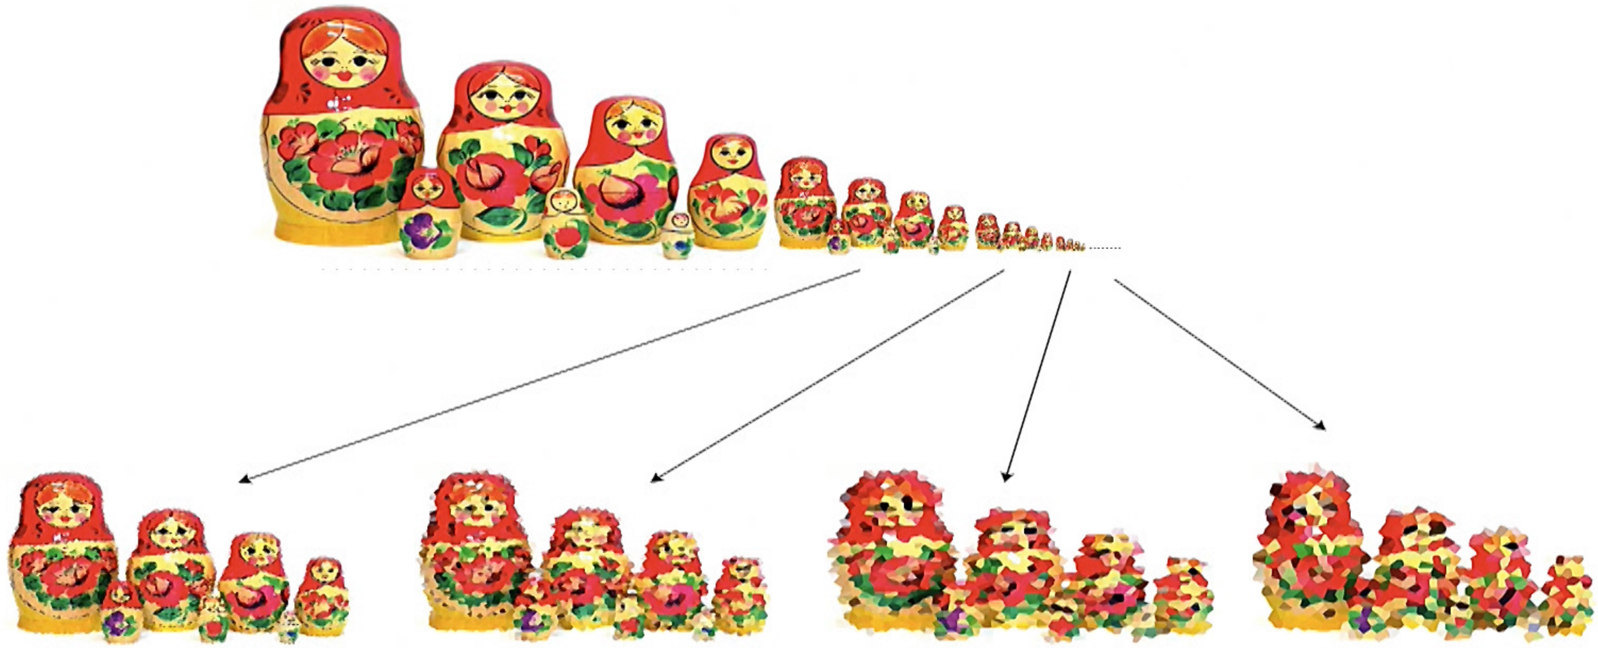

If we represent the smaller dolls (which dominate the computational load) imperfectly, this will allow us to represent more dolls!

More accuracy with less precision!

A future route for supercomputing?

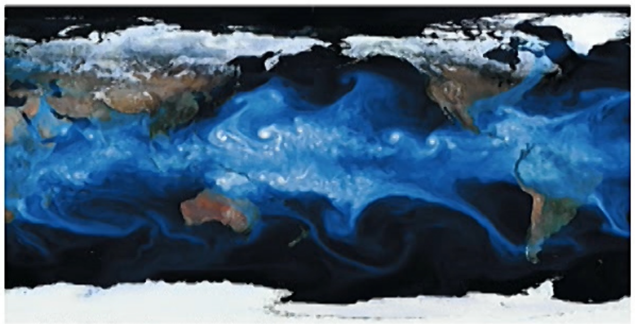

COMMENT

A simulation of Earth's atmosphere generated by the Community Atmosphere Model.

Build imprecise supercomputers

Energy-optimized hybrid computers with a range of processor accuracies will advance modelling in fields from climate change to neuroscience, says **Tim Palmer**.

Today's supercomputers lack the power to model accurately many aspects of the real world, from the impact of cloud systems on Earth's climate to the processing ability of the human brain. Rather than wait decades for sufficiently powerful supercomputers — with their potentially unsustainable energy demands — it is time for researchers to reconsider the basic concept of the computer. We must move beyond the idea of a computer as a fast but otherwise traditional "Turing machine", churning through calculations bit by bit in a sequential, precise and reproducible manner.

In particular, we should question whether all scientific computations need to be performed deterministically — that is, always producing the same output given the same input — and with the same high level of precision. I argue that for many applications they do not.

Energy-efficient hybrid supercomputers with a range of processor accuracies need to be developed. These would combine conventional energy-intensive processors with low-energy, non-deterministic processors, able to analyse data at variable levels of precision. The demand for such machines could be substantial, across diverse sectors of the scientific community.

MORE WITH LESS

Take climate change, for example. Estimates of Earth's future climate are based on solving nonlinear (partial differential) equations for fluid flow in the atmosphere and oceans. Current climate simulators — typically with grid cells of 100 kilometres in width — can resolve the large, low-pressure weather systems typical of mid-latitudes, but not individual clouds. Yet modelling cloud systems accurately is crucial for reliable estimates of the impact of anthropogenic emissions on global temperature[1].

The resolution of the computational grid is determined by the available computing power. Current petaflop computers can perform up to 10^{15} additions or multiplications — floating-point operations — per second (flops). By the early 2020s, next-generation exaflop supercomputers, capable of 10^{18} operations per second, will be able to resolve the largest and most vigorous types of thunderstorm[2]. But cloud physics on scales smaller than a grid cell will still have to be approximated, or

32 | NATURE | VOL 526 | 1 OCTOBER 2015

© 2015 Macmillan Publishers Limited. All rights reserved

- Energy efficient heterogeneous architecture
  - Small numbers of energy-intensive high-precision deterministic processors
  - Large numbers of low-energy low- precision potentially non-deterministic processors

Superefficient inexact chips

http://news.rice.edu/2012/05/17/computing-experts-unveil-superefficient-inexact-chip/

Prototype
Probabilistic
CMOS
Chip

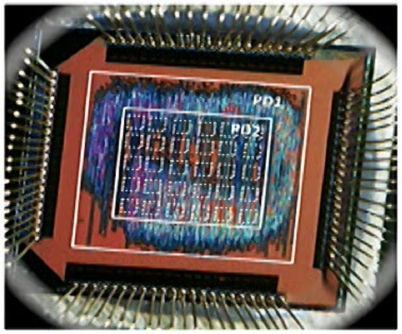

Krishna Palem.
Rice University

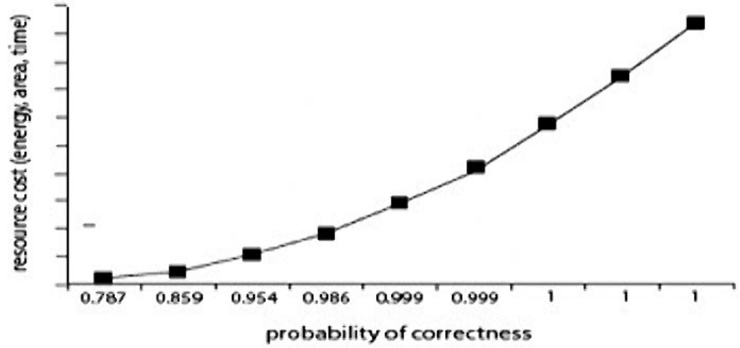

The chip that produced the frame with the most errors (right) is about 15 times more efficient in terms of speed, space and energy than the chip that produced the pristine image (left).

Are there precedents for energy-efficient imprecise supercomputing?

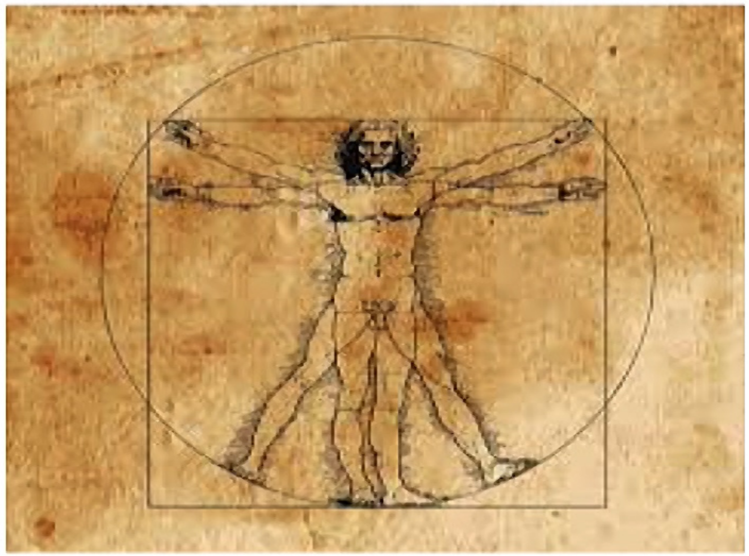

Yes, us!

Are there precedents for energy-efficient imprecise supercomputing?

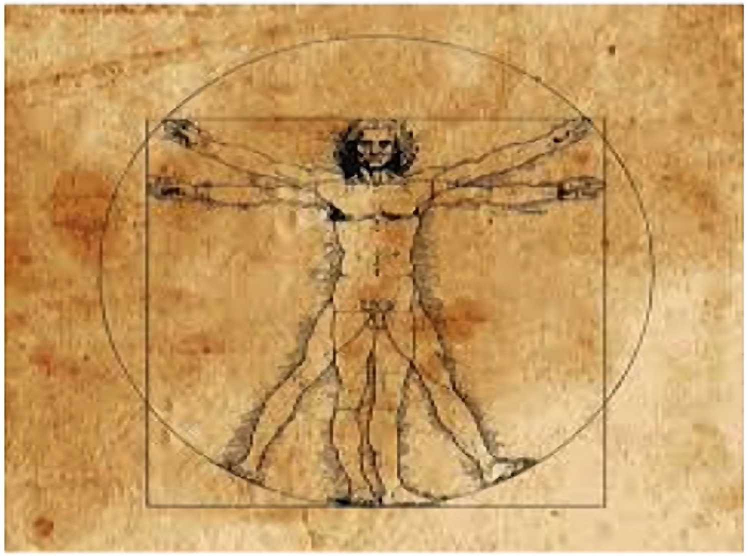

- 100 billion neurons
- Signal processing capability of 1-10 exaflops
- All from 20W!

Yes, us!

Other areas that could benefit from imprecise computing

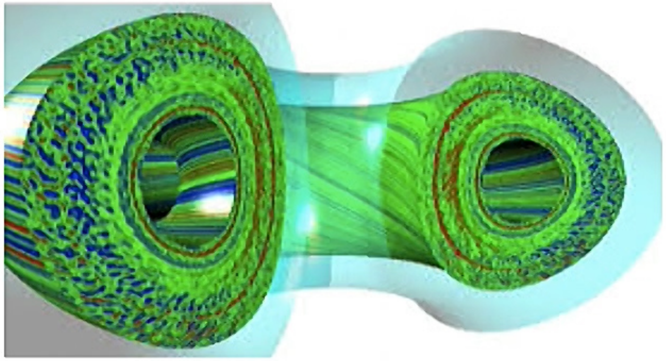

So – what side are you on?

Believer

or

Denier

No!