

Introduction



Cite this article: Palmer T, Düben P, McNamara H. 2014 Stochastic modelling and energy-efficient computing for weather and climate prediction. *Phil. Trans. R. Soc. A* **372**: 20140118.
<http://dx.doi.org/10.1098/rsta.2014.0118>

One contribution of 14 to a Theme Issue 'Stochastic modelling and energy-efficient computing for weather and climate prediction'.

Subject Areas:

climatology, meteorology, computer modelling and simulation

Author for correspondence:

Tim Palmer
e-mail: t.n.palmer@atm.ox.ac.uk

Stochastic modelling and energy-efficient computing for weather and climate prediction

Tim Palmer¹, Peter Düben¹ and Hugh McNamara²

¹Atmospheric, Oceanic and Planetary Physics, University of Oxford, Clarendon Laboratory, Parks Road, Oxford OX1 3PU, UK

²Mathematical Institute, University of Oxford, Andrew Wiles Building, Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, UK

Being able to predict weather and climate reliably over a range of time scale is crucial if we are to create a society which is resilient to extremes of weather and climate. Although current weather and climate models can show impressive levels of predictive skill, for example in predicting the track of a tropical cyclone over the coming days or the development of El Niño sea surface temperature anomalies over the coming months, the performance of these models is degraded by the development of systematic departures, or biases, between simulated fields and observed fields. The development of these model deficiencies is not associated with a lack of knowledge or understanding of the laws which govern the evolution of weather and climate. Rather, these deficiencies arise because we are not able to solve the associated nonlinear multi-scale evolution equations with sufficient accuracy.

Ever since the first numerical weather prediction models were formulated in the 1940s, these evolution equations have been solved numerically by projecting the equations onto a grid with finite resolution. These so-called dynamical cores would be supplemented by a set of formulae (or parametrizations) derived by a combination of theory and empiricism, which represented the bulk effect of processes unresolved by the dynamical core.

By decreasing the grid spacing sufficiently, the inevitable inaccuracies associated with these parametrizations can be reduced. However, this process is computationally costly. Halving the spacing of the three-dimensional grid can increase the computational

cost of a weather or climate model by up to a factor of 16. In this way, it has been estimated that the development of a climate model which is able to represent deep convective cloud systems (which play a key role in the tropics) will require exascale computational capability. It may be many years before weather and climate centres have dedicated access to such computers, and it is not inconceivable that the power requirements of such computers may make them prohibitively expensive to many such centres.

Recently, however, a new paradigm for solving the equations of motion of weather and climate is beginning to emerge. The basis for this paradigm is the power-law structure observed in many climate variables. This power-law structure indicates that there is no natural way to delineate variables as ‘large’ or ‘small’—in other words, there is no absolute basis for the separation in numerical models between resolved and unresolved variables.

A first step towards making this division less artificial in numerical models has been the generalization of the parametrization process to include inherently stochastic representations of unresolved processes. However, the development of stochastic parametrizations raises a profound question. What is the real information content in the variables of the dynamical core as a function of scale? The information content of variables near the truncation scale of the dynamical core cannot be arbitrary because they are strongly influenced by the stochasticity of the unresolved parametrization packages. As we move to larger scale, the influence of stochasticity will decrease and the real information content will increase—but how quickly? In a three-dimensional system with a shallow power-law spectrum, the increase in information content with scale may be slow.

In principle, a knowledge of scale-dependent information content is extremely relevant for the design of next-generation weather and climate models. Because of the power-hungry nature of current and likely future supercomputers, it is essential that future models are as efficient as possible. A knowledge of scale-dependent information content will help determine the optimal numerical precision with which the variables of a weather or climate model should be represented as a function of scale. A knowledge of scale-dependent information can also help determine whether there is a need for complete determinism in the numerical computations for the dynamical core. By relaxing the traditional assumption that all equations in the dynamical core are computed with bit-reproducible determinism, and that all variables are represented by double-precision floating-point real numbers, a new paradigm for weather and climate models has emerged based on the notion of variable scale-dependent precision and exactness.

This Theme Issue is based on presentations at a workshop at Oriel College Oxford in 2013 that brought together, for the first time, weather and climate modellers on the one hand, and computer scientists on the other hand, to discuss the role of inexact and stochastic computation in weather and climate prediction. As a result of this meeting, a vision was articulated of a new synergism between software and hardware design, based firmly on the laws of physics. The first steps in realizing this new synergism are outlined in the papers in this Theme Issue. The goal, as far as weather and climate prediction is concerned, is the development of truly cloud-resolved simulators: by only computing to levels of precision dictated by the relevant scale-dependent information content, valuable computing (and energy) resources can be used to extend the resolution of the dynamical core, thus taking key cloud-scale processes out of the bulk-formula parametrizations.

While the papers in this Theme Issue are strongly focused on weather and climate prediction, it is likely that the techniques discussed will be relevant to many other problems in computational science. Indeed, it is important to be able to identify these problems—the greater the number of potential customers for new types of super-efficient inexact chips, the more likely it is that the chip manufacturers will be prepared to create such chips in bulk.

We dedicate this Theme Issue to the memory of Dr James Martin, founder of the Oxford Martin School and himself a pioneer in the development of computer software systems. The Oxford Martin School was a generous supporter of Oxford’s contribution to the research presented in this issue, and provided considerable support for the Oriel workshop. James Martin himself was an extremely enthusiastic supporter of the aims of the workshop and gave an inspirational and memorable address to the workshop participants.