Challenges for weather and climate prediction - a UK perspective

Nigel Wood
Met Office, Exeter, UK

26/01/12

This decade is set to be an interesting one for operational weather and climate modelling. The accuracy of weather forecasts has reached unprecedented and probably unexpected levels: large-scale measures of accuracy continue to improve at the rate of 1 day every 10 years so that today's 3 day forecast is as accurate as the 1 day forecast was 20 years ago.

In order to maintain this level of improvement operational centres need to continue to increase the resolutions of their models. Increasingly this means running models at resolutions of the order of a kilometre. This leads to many challenges. One is how to handle processes that are only barely resolved at those scales. Another is how to present, and also verify, forecasts that are inherently uncertain due to the chaotic nature of the atmosphere.

A more practical issue though is simply how to run the models at these increased resolutions! To do so requires harnessing the power of some of the world's largest supercomputers. However, we are entering a period of radical change in the architecture of these machines: no longer is increased power coming from increased clock speeds; rather the increase in transistors predicted by Moore's law is being achieved by vast increases in the number of processors. The fastest machine currently in the world is the K-machine. This reaches a peak of about 10 petaflops ($O\left(10 \times 1000^5\right)$ flops) which is achieved using around 700,000 cores. It consumes 10Mw of power. The exascale machines (promised for the end of this decade) will be two orders of magnitude faster ($O\left(1000^6\right)$) than the K-machine.

Current operational models scale well to $O\left(10^3\right)-O\left(10^4\right)$ cores with some (non-operational) models achieving good scaling to $O\left(10^5\right)$ cores. It is clear that to continue to capitalise on the ever increasing computer power available, our algorithms need to be fundamentally redesigned to dramatically increase their parallelism. However, that challenge is made more difficult by the fact that the UK Met Office's model (the MetUM) is unified in that the same dynamical core (and increasingly also the same physics packages and settings) is used for all our operational weather and climate predictions. The model therefore has to perform well across a wide range of both spatial scales [$O\left(10^0\right) - O\left(10^4\right)$km] and temporal scales [$O\left(10^0\right) - O\left(10^4\right)$ days] as well as a wide range of platforms: only the largest, flagship climate simulations will have access to the very largest computers; the size of computer used for our daily operational forecasts will be limited by the public purse! and many experimental climate simulations (e.g. paleoclimate studies) will want to run on much smaller clusters.

This talk will start by outlining the current status of the MetUM, then discuss planned developments (focussing on numerical aspects) before going on to highlight recent progress within GungHo! - the project that is redesigning the dynamical core of the model.