



UNIVERSITY OF OKLAHOMA

120 David L. Boren Blvd

Norman, OK 73072

Phone: (405) 325-1819

Fax: (405) 325-1108

[www.weathersphere.org](http://www.weathersphere.org)

## Courtesy of High-Performance Computing—02/25/09

### Weathering well

By: Gemma Church/HPC

The planet's changing climate and freaky weather is hitting the headlines more and more these days. Hurricane Katrina, Antarctica's increasing temperature – and even floods in north Cornwall in the UK – show that, not only can Earth's weather be erratic, it can be devastating.

But how can you predict the unpredictable? Simulations of any weather system must take into account a massive number of variables and contributing factors, making such modelling seemingly impossible. Or so it was until a few years ago.

As supercomputing improved, so did its ability to handle the swarms of data needed to produce meaningful models of the Earth's climate. Thanks to dedicated centres popping up all over the globe, scientists can at last run simulations that will find the most likely answer to how a hurricane hitting the US East coast will affect its landscape, how an ice sheet melts, or even look thousands of years into the past to pinpoint the planet's previous environmental changes.

### Chilling times

In the popular imagination, 'climate change' often conjures up an image of a polar bear clinging onto a melting chunk of ice. Modelling the Earth's ice sheets is a major concern for two facilities in the west of the UK.

Swansea University's Mike Barnsley Centre for Climate Research recently unveiled its aptly-named Blue Ice supercomputer, designed and installed by HPC integrator OCF, which will aid research into understanding the impact of environmental changes, such as rising sea levels and melting glaciers and ice sheets.

The main area of expertise is environmental modelling, as Professor Tavi Murray, scientific director of the centre, explains: 'We are aiming to model a wide variety of environmental systems, such as the ice itself and the coupling between the ice, water and the atmosphere.'

This is where the high-powered HPC resource helps, as Murray adds: 'There are three reasons why we need a lot of computational power. Firstly, we might need to represent features accurately, for example if you are trying to model the outlet glaciers of the Greenland Ice Sheet, that obviously takes a lot of computational power.

'Secondly, the coupling between the ice sheets, water and atmosphere takes up a lot of computational resource and, finally, we are running experiments by running these different models many times.'

This holy trinity of accuracy, computational bulk and repeatability is echoed by climate modellers around the world, with the final point being particularly important. By changing the initial conditions within each model and running simulations multiple times concurrently, scientists can produce a mean result. This means, in the case of the Mike Barnsley Centre, finding the most likely effect of the melting glaciers on the planet. Murray adds: 'This leads to more robust answers, rather than just running one simulation with one set of initial conditions.'


At the University of Bristol, climate modellers are simulating how the planet's climate changed tens of thousands of years ago. Modelling the entire planet over such immense timescales requires an immense supercomputer so the university opened its £7m facility, Blue Crystal, last year. The supercomputer can carry out more than 37 trillion calculations per second.

While most current weather simulations of the UK or Europe may look at a resolution of a few kilometres across and timescales of a few days, for such planet-size, thousand-year timescale simulations the Earth is broken down into areas hundreds of kilometres across. For example, the UK is represented by six or seven boxes. But the simulations can cope with such relatively small timesteps, as Dr Dan Lunt, research fellow at the School of Geographical Science at the University of Bristol, says: 'The new facility has helped our research no end. We used to model timescales of a century or two; now we simulate over thousands of years and can break those simulations down to timesteps of every 30 minutes.'

As Professor Murray mentioned, another big boost to such climate change research is that the scientists can map out many simulations at once to accommodate all the different variables that have to be taken into account. Lunt adds: 'One of the great things about Blue Crystal is that it is so big so we can run up to 100 simulations at a time.'

'Before this HPC facility came along we just had to use best estimates of the boundary conditions. This new HPC facility actually makes the previous work we did look meaningless in comparison.'

The centre is currently trying to work out what happened between a time around 20,000 years ago when the Earth's northern hemisphere was blanketed with ice, to how the planet became the relatively ice-free world we live in today. Blue Crystal is allowing the researchers to look into this time frame in much more detail than ever before, as Lunt explains: 'We used to do one simulation where the ice was at its maximum thickness (around 20,000 years ago), and another at the minimum thickness (the present day).'



*Visual interpretation of climate change data, aided by Blue Ice.*

‘Now we have moved away from such snapshots and have transient models simulating the ice changing over the whole 20,000 years. This simply was not possible five or so years ago.’

But the code needed to be re-jigged to cope with the heavyweight hardware it must run on, as Lunt explains: ‘A lot of the climate model code was written 30 years ago, and it’s beginning to show its age because it simply was not written with hundreds or thousands of processors in mind. At the moment, we can only break the Earth down into up to about 50 chunks – beyond that you have a saturation point where the code does not know how to parallelise the work over multiple processors.’

Lunt adds: ‘We are now developing code that is more parallelisable, which can run over thousands of processors and correspondingly work thousands of times faster. This could potentially bring huge benefits, it may be possible to move away from modelling over 20,000 years to being able to model over one million years.’

## **Oceans 2025**

No, it’s not yet another sequel to a rehashed Hollywood movie, Oceans 2025 is actually concerned with modelling the Earth’s oceans.

Researchers at the National Oceanography Centre at Southampton (NOCS) are part of the much larger Oceans 2025 project, which seeks to increase knowledge of the marine environment. The project, collaboration between seven UK marine research centres, will run from 2007 to 2012 with a £120m budget. Specifically, the NOCS group will aim to deliver the ocean models needed for the next decade of UK marine science.

The French/European NEMO ocean model will be used for the majority of the ocean modelling carried out by NOCS. NEMO (Nucleus for European Modelling of the Ocean) allows ocean-related components (for example sea-ice, biochemistry, ocean dynamics and so on) to work either together or separately. Unlike previous ocean models used by UK researchers, the NEMO model has not been specifically optimised for use on UK supercomputers.

Dr Andrew Coward, researcher at NOCS, believes the ability to add complexity to the model is important and says: ‘Many uncertainties exist about the role of ocean biogeochemistry in

the climate system. Models of ocean biogeochemistry are themselves complex but life in the ocean is also critically dependent on details of the physical environment that it finds itself in.

‘Thus complex biogeochemistry models need to be coupled with high resolution ocean circulation models in order to provide investigative tools. Such tools are needed to elucidate the biological feedbacks in the climate system. NEMO provides the ideal framework for studies of this type but the sheer size and complexity of the coupled ocean-biogeochemistry system will continue to provide a supercomputing challenge.’

### **Eye of the storm**

Over the pond, a US centre called the CHAMPS lab (Coastal Hydroscience Analysis, Modelling and Predictive Simulations laboratory) at the University of Central Florida (UCF) is concerned with climate modelling. It is investigating coastal hydroscience and the related issue of how systems such as hurricanes bring water onto land in and around the state of Florida.

The CHAMPS Lab is currently working on a model to define the extreme flooding scenarios that could occur once every 100 to 500 years. The supercomputing system deals with around 850,000 computation points per simulation and maps the landscape every second over a three- or five-day period, as Scott Hagen, associate professor of civil engineering at UCF, explains: ‘Because we have such a small time step our matrices need to be solved so many times. Being able to break that model down into interconnecting subdomains means we can do the simulation in terms of hours instead of days, thanks to our HPC facilities.’

The team uses high-resolution LIDAR images and puts that data into a finite element model, which is made up of interconnected triangles to map the water and land’s surface and elevation. The next step for these models is to improve the resolution further, as Hagen adds: ‘Our present models use 850,000 computational points and we are limited by the smallest size of around 40 metres for each triangle’s side. It would be great to get down to the metre range to bring down the approximation but a scale of less than five metres would increase the number of points to make this kind of work infeasible at present.’



‘One area that does require improvement is working out the extent of the inland flooding in localised areas. For example, you might want to zoom into a small area and see how it will be affected by such flooding – this is where the five-metre mark would become important.’

But, as the need for more HPC hardware increases, John Michalakes, lead software developer for Weather Research and Forecast (WRF) model at National Center for Atmospheric Research (NCAR), warns that researchers must box clever: ‘Today, numerical weather

prediction (NWP) is a terascale (10<sup>12</sup> floating point operations per second) application. The current trend in HPC is to reach petascale (10<sup>15</sup> floating point operations per second) and beyond by essentially building larger and larger clusters, which is okay for applications where you can scale the problem by making it larger. Many weather applications, especially on the research side, can do this.'

As part of a multi-institutional collaboration, the NCAR developed the WRF model, which is in wide use for operational weather forecasting and for atmospheric simulation. WRF-generated forecasts are, amongst other applications, repackaged as commercial forecast products and then seen on nightly television in more than 200 markets around the US. NCAR scientists are currently running simulations of processes within the hurricane eyewall at resolutions of 60 metres.

Michalakes adds: 'But it's important to realise that, while the amount of computations possible in a given amount of time on the system increases, the speed of the run does not necessarily increase. So weather and climate applications that have a time to solution requirement – real-time forecasting or very long climate simulations – may not benefit from conventional petascale systems. For these applications, faster processors, not just more of them, are needed.'

### **Oklahoma!**

The Center for Analysis and Prediction of Storms (CAPS) at the University of Oklahoma is developing thunderstorm-scale numerical weather predictions. That is, models using resolutions of 3km or less.

Keith Brewster, senior research scientist and associate director at CAPS, explains: 'In 2008 we did some real-time forecasts at 2km, and some 1km grid scale forecasts locally in near-real-time. The high resolution allows us to explicitly model the flow within the thunderstorm; this allows us to better gauge the threat of severe weather (high wind, hail or tornadoes). The development and movement of storms, however, is often steered by the larger-scale flows, so we need to get that correct as well. Forecasts of more than one or two hours require a large domain, as individual storms and weather systems in the active spring season can travel quite fast (up to 30 m/s).'

But the team is modelling to even higher resolutions, as Brewster explains: 'For research, we are modelling the creation and development of tornadoes within thunderstorms. For that we are using the model run at grid resolutions as small as 20 metres. These cannot be done in real-time but they are very computationally demanding – so, to get turnaround in a reasonable length of time, days not months, they must run on HPC resources and run efficiently.'

### **Tip of the iceberg**

While such simulations are, of course, helping us understand the planet's climate, there does seem to be something a little counterintuitive about using power-hungry servers to try to

calculate how draining the Earth's resources has affected its climate. Last year, the UK's Met Office spent £33m on a new supercomputer to calculate how climate change will affect Britain – only to find the new machine has a giant carbon footprint of its own, emitting around 14,400 tonnes of CO<sub>2</sub> a year, which is equivalent to the CO<sub>2</sub> emitted by 2,400 homes, according to reports.

So, as climatologists demand increasingly more power from their supercomputers, could turning the machines off ultimately be the way to curb the very effects they are trying to understand?