A FLAGSHIP EUROPEAN PROGRAMME ON EXTREME COMPUTING AND CLIMATE

About Me

I am a Royal Society Research Professor of Climate Physics at the University of Oxford. This proposal arose out of a 2-day workshop between 21 leading climate modellers (representing Europe's leading climate modelling and climate computing centres – see below) held at the Royal Society's Chicheley Hall in January 2016. As well as these 21 climate scientists, I also represent the European Network for Earth System Modelling whose Board fully supports this proposal.

What is the Challenge and the Vision?

At the 2015 Paris Climate Conference, leaders from 194 countries of the world unanimously acknowledged the serious threat posed by anthropogenic emissions of greenhouse gases. Society must now become resilient to changes in climate over coming decades. Most importantly, this will require quantitative estimates of the changing character of climate extremes. Such extremes include not only exceptional weather events such as violent wind storms and flash floods, but also persistent anomalies in planetary-scale circulation patterns, which lead to pervasive flooding in some regions and seasons, and long-lived drought and extremes of heat in others. However, providing such information will require a step change in the quality of global climate models, which at present are simply not adequate for this purpose[1]. In particular, the resolution of these models must increase to a level which allows both ocean eddies and individual cloud systems to become represented explicitly; this is the best hope for obviating long-standing climate model biases.

Such capability can begin to be realised with the advent of exascale supercomputing (10^{18} floating point operations per second) anticipated around 2023. However, being able to realise exaFLOPS for practical applications will require considerable investment in climate modelling software. Indeed, such is the complexity of current climate system models, such an endeavour will require a pooling of expertise and resources from existing climate institutes. Here we argue for a European Flagship Programme On Extreme Computing and Climate. Drawing on existing climate modelling expertise in Europe and working closely with existing supercomputing centres, EPECC (pronounced "Epic") would oversee the development of cloud- and eddy-resolved global climate system models, and integration of these models into an extreme-scale computing technology platform. Over the term of the programme, EPECC deliverables would be tailor-made for simulating climate extremes accurately on exaFLOP compute infrastructure. The key scientific focus for EPECC will be the drivers of

---

1

http://rspa.royalsocietypublishing.org/content/472/2188/20150772.full?ijkey=DBxz3lBqowVnTMJ&keytype=ref

changing extremes in Europe over the 21st century. EPECC will work in cooperation with national meteorological and oceanographic agencies, to provide information tailored to the needs of their clients in the public and private sector, with the aim of increasing European resilience to a changing climate.

Why Is It Good For Europe?

Whilst there is no doubt that the surface of the planet will continue to warm as a result of past and on going anthropogenic emissions of greenhouse gases, and little doubt that such warming will be linked to an increase in the occurrence of climate extremes around the globe, climate scientists have only rather general knowledge about the way in which such extremes will manifest themselves in any particular region, such as Europe.

For example, the UK suffered severe flooding during the winters 2013/14 and 2015/16 with record-breaking rainfall amounts. A key question repeatedly asked by the public, media and government alike was: Were these floods caused by climate change? Because key infrastructure (such as electricity sub-stations) was threatened from inundation by flood water, it is crucial to know whether such flooding is likely to become the norm in future years. Being better able to anticipate and understand the origins of extremes and how they may change would thus greatly increase the resilience of European society. However, extremes often depend on remote drivers. For instance, the key meteorological reason for these floods was the strength and orientation of the jet stream across the Atlantic. Hence, a crucial question is whether climate change has led to, and more importantly will lead to, persistent and intense jet stream anomalies of the type leading to European flooding. But the current generation of models have not proven adept at anticipating these types of anomalies. Reasons are not hard to find. There is evidence that some of the key drivers of persistent jet stream anomalies lie in the tropics, and current climate models continue to exhibit major errors in tropical cloud and precipitation when compared with observations. This in turn lies in the fact that large-scale tropical circulations are themselves driven by deep convective cloud systems whose kilometre scales are too small to be resolved by current climate models. Secondly, anomalies in both the position and intensity of the jet streams are linked to strong gradients in temperature (and other key variables) between the troposphere and stratosphere. Again because of resolution limitations (vertical and horizontal), these gradients are erroneously weak in current climate models. Finally, there is good evidence that biases and errors in the representation of ocean currents, such as the Gulf Stream, again caused by insufficient ocean model resolution, degrade simulations of climate.

A related example is the case of the 2003 European heat wave. Again, even though there may be general reasons to expect the frequency and intensity of heat waves to increase under the impact of climate change, the key meteorological reason for the 2003 heat wave was a persistent deviation of the Atlantic jet stream from its normal west-to-east direction. To estimate reliably the likelihood of another 2003-like heat wave in Europe, it is crucially important to be able to estimate how enhanced greenhouse-gas concentrations in the atmosphere will affect the probability of these (so-called) persistent anticyclonically blocked circulation patterns. Whilst increases in model resolution have improved simulation of blocking in recent years, current climate models still do poorly in simulating such very long lived blocked circulation patterns, let alone their drivers, for many of the same reasons as for the flooding events described above.

Research over the last decade or so – particularly in the field of numerical weather prediction - has shown that systematic errors will reduce as the resolution of our models (in both atmosphere and ocean components) increases. Over these years, the horizontal resolution of the very best climate models has reduced from some hundreds of kilometres to some tens of kilometres. However, it is now widely believed that many of the gravest and most stubborn errors are related to poor representation of processes that occur on scales of 10 km or less. To resolve processes on a 10km scale requires a grid an order of magnitude finer. Indeed, developing the capacity to simulate the global climate system on 1 km scales (or even finer) will start to allow deep convective cloud systems, crucial for determining the structure of the tropics, to become resolved explicitly, rather than represented by highly simplified sub-grid parametrisation formulae, currently the case. Much higher resolution will also allow the sharp troposphere/stratosphere interface to be represented much more accurately than is currently possible, and higher resolution in the oceans will allow ocean eddies to be resolved and many important bathymetric features to be represented at all latitudes – both are believed to be crucial in shaping the large ocean currents, such as the Gulf Stream (again with substantial implications for European climate).

A 1km global resolution is currently too fine to be run in an operational sense on current supercomputers, but becomes feasible with next generation exascale computers, at least for the multi-decadal timescales simulations needed to inform on matters of societal resilience. Bearing in mind the recent US Presidential Initiative designed to accelerate progress in the development of exascale supercomputing, it is anticipated that prototype exascale computers will be available in 2023 or thereabouts. We propose that the complexity of climate on the one hand, and the importance of reliable predictions on the other, demands *dedicated* exascale computing for climate. This can be achieved through a EU Flagship Programme. This in no way reduces the continuing need for national supercomputing capability for weather and climate modelling.

What would it take to do it?

Europe has had world-leading climate modelling and numerical weather forecasting capabilities. However, no single nation, let alone institute has the capability to develop and operationally use a 1km (or finer) global coupled (atmosphere-ocean) model in the coming years, and few have begun to think systematically about what drives future extremes. Experience in simulating climate at 1 km resolution is beginning to be explored in regional climate models and in regional numerical weather prediction. We must draw on this expertise for developing next generation global climate models. In addition, many centres are beginning to develop prototype solutions, as evidenced by a new generation of global model dynamical core developments in France, Germany, and the UK. However, to port these applications to an exaFLOP environment and use them in a more operational manner will require large investments in software engineering. The scale of the requisite effort and the commonalities entailed demands a coordinated approach, one that is designed not around a single solution, but rather one that develops tools and standards to define a framework for porting algorithms already in development. Only such an approach will ensure European competitiveness in the use of exaFLOP computing to study weather and climate. By developing frameworks designed to accommodate existing developments, EPECC will strengthen on-going national activities, be in a better position to assess the contribution of structural uncertainties in estimates of future extremes, and open the door for further coordination in the application of high-performance computing to study climate and provide climate services, particularly as it relates to extremes.

Extreme weather and the role of deep convection globally in influencing the circulation patterns leading to extremes are at the nexus of two of the World Climate Research Programmes Grand Science Challenges. In addition, the goal of developing a 1km global climate model is part of the European Network for Earth System Modelling's grand challenge for the future[2]. In this document it is noted that getting to this scale will help solve other key problems such as seasonal to decadal weather/climate prediction, and understanding of drivers of past climates. Indeed, the scientific and computational challenges are very similar for weather and climate prediction, and the 1km global model target will benefit both applications alike.

By linking these to advances in extreme computing, EPECC will sustain and advance European Leadership at the very forefront of science, with associated benefits for European climate research institutions, meteorological and oceanographic agencies, and leadership in the application of high-performance computing. The elements of this proposal for a flagship programme on extreme computing and climate can be summarised as follows.

1. In partnership with existing European Climate Centres (including the European Centre of Excellence in Weather and Climate Forecasting, ESIWACE), a European flagship programme is proposed to focus on drivers of 21st Century climate extremes, with a focus on European extremes but not excluding extremes in other parts of the globe, providing reliable estimates of the occurrence of climate extremes in the coming decades. A global modelling approach is essential as many of the drivers of European extremes may be remote.

2. By leveraging existing supercomputer centres, the programme will provide resources for *dedicated* Exascale computing for simulating climate extremes.

[2] https://verc.enes.org/community/about-enes/the-future-of-enes/ENES%20foresight.pdf

3. The key scientific challenge will be the development of the next generation of global climate models, better able to represent small-scale processes and perform efficiently on Exascale computing facilities. By 2025, this would enable ensembles of 1 km global simulations allowing a delivery of around one year's simulation per one day elapse/real time, with an expected tenfold increase in resolution a decade later. Work to develop common dynamical cores would be supplemented by the development of sets of sub-grid parametrisation schemes in collaboration with existing centres around Europe.

4. Strong collaborative links will be developed with existing climate modelling institutes, climate service centres, computing centres and with national meteorological and oceanographic agencies. Partnering centres and institutions will play an important role in EPECC's governance. There will be strong engagement with both climate and weather prediction communities as the underlying scientific and computing challenges are similar.

5. Tools will be developed to allow the analysis of ensembles of high-resolution integrations (which inevitably involve large datasets) to be performed efficiently.

6. A project plan will be developed with a work breakdown consistent with the goals of and developed with, emerging global exascale computing projects such as ESIWACE and the European extreme-scale demonstrator programme. An appropriate project management structure would complement this.

7. An adequate team of computational climate scientists will be built up supported by a larger pool of software engineers, distributed in a coarse-grained way at the major participating modelling/computing institutions. Human resources will be distributed in order to leverage contributions from existing climate centres and meteorological/oceanographic agencies, with perhaps a physical location for coordination and workshops/meetings.

8. Advanced training will be provided to a new generation of scientists and software engineers from the different European Climate Centres. It is these scientists and engineers that will develop, maintain or use the climate and numerical weather forecasting models that will run on future extreme scale computing technology platforms.

The scientists participating in the Chicheley Hall Workshop in January 2016 ("Advancing the Climate Modelling Enterprise in Europe") were, in alphabetical order: Peter Bauer (ECMWF, UK), Stephen Belcher (Met Office, UK), Sandrine Bony (LMD, France), Francisco Doblas-Reyes (Barcelona Computer Centre, Spain), Ralf Doescher (SMHI, Sweden), Jean-Louis Dufresne (LMD, France), Brian Hoskins (Imperial College, UK), Sylvie Joussaume (IPSL, France + ENES Chair), Bryan Lawrence (NCAS, UK), Jochem Marotzke (MPI, Germany), David Salas Y Melia (Météo-France, France), Antonio Navarra (CMCC, Italy), Adrian New (NOC, UK), Tim Palmer (U. Oxford, UK), Thomas Schulthess (ETH Zurich, CSCS, Switzerland), Pier Siebesma (KNMI, Netherlands), Julia Slingo (Met Office, UK), Piotr Smolarkiewicz (ECMWF, UK), Bjorn Stevens (MPI, Germany), Rowan Sutton (U. Reading, UK), Pier Luigi Vidale (U. Reading, UK).