Submission No: 18.

Submission to the House of Representatives Industry, Science, and Innovation Committee's Inquiry Into Meteorological Forecasting

Prepared by J. Walter Larson, PhD

Background/Biographical Information

I am a permanent resident of Australia but am employed as a computational scientist at Argonne National Laboratory in Chicago, and am a senior fellow in the computation institute at the University of Chicago. I am also currently an adjunct senior lecturer in the Department of Computer Science at the ANU, and a visiting fellow in the Research School of Physical Sciences and Engineering at the ANU. From 1996-1999 I worked in the field of data assimilation at NASA's Data Assimilation office. From 2000-2005 I was the Department of Energy (DOE) lead on the design and development of the flux coupling infrastructure for the US Community Climate System Model. From 2006-2008 I was employed by the predecessor to the NCI National Facility and during this time worked with Australian researchers on computational issues regarding the Australian Community Climate and Earth System Simulator (ACCESS). In sum, I have just under fifteen years' experience working in the area of high-performance computing in climate and weather applications. I mention these affiliations to give the committee some understanding of my qualifications and experience, and their relevance to some of the terms of reference in this inquiry. That said, ***the opinions and ideas expressed in my submission are mine alone, were prepared on my own time with my own resources, and should not be taken in any way to reflect the opinions or attitudes of my employers past or present, or any other organisations with which I am presently or was previously affiliated.***

Response to Inquiry's Terms of Reference

Innovation in long-term meteorological forecasting methods and technology

Long-standing inadequate financial support for research in the fields of numerical weather prediction and computational science have had a damaging effect on the Australian climate/weather/ocean (CWO) modelling community. Traditionally, Australia has enjoyed a reputation of "punching above its weight" in the CWO arena. I believe this enviable position is in peril. A notable consequence of poor long-term support of the

CWO community is the adoption of the UK Hadley Centre Unified Model (UM) as the atmosphere model in the Predictive Ocean and Atmosphere Model for Australia (POAMA) and the Australian Community Climate and Earth System Simulator (ACCESS) systems developed by Bureau of Meteorology (BoM) and CSIRO with some input from the Australian university community. It is my understanding that one of the chief reasons for adoption of the UM was its support for four-dimensional data assimilation (4DDA), the current "best-practice" in data assimilation for initialisation of forecast models, and something absent from POAMA.

Some things no doubt will also be said about how the lack of computational resources in Australia is at the root of this problem. Poor support for computational platforms for research is part of the problem. Among the semi-annual "top 500" listing of the world's fastest supercomputers (http://www.top500.org), Australia has only one machine on the list, which is owned by a computer animation company, not a government research body. New Zealand, by comparison has three machines on the Top 500. More hardware alone will not solve Australia's problems in the CWO prediction/modelling enterprise. A larger problem in my view is the lack of support for the emerging field of computational science, an interdisciplinary area that combines computer science, high-performance computing, software engineering, and numerical analysis. Computational scientists seek to solve algorithmic problems relevant to computer modelling in many fields of science and engineering. Because researchers are not well-supported in this area, we are missing opportunities for collaborative innovation in CWO modelling. We are not training future generations of people who have the necessary skills to develop superior, performance-portable algorithms in support of the types of short-to-medium-range weather forecasting, seasonal-to-interannual prediction, climate, and other environmental modelling the Australian taxpayer expects from our CWO forecasting and research bodies. And, needless to say, we are diminishing our future national competitiveness in this field.

I suspect some other submissions may point out the poor career path on offer to many CWO researchers. This is due in large part to the feast-or-famine funding model CWO research bodies are forced to accept, restricting their ability to offer researchers job security or globally competitive salaries. Prospects for computational scientists within the Australian CWO community are even worse. Almost universally, computational scientists and software engineers are misclassified as "IT officers," rather than the seasoned, highly-trained professionals (many with postgraduate degrees) that they are.

Strategies, systems and research overseas that could contribute to Australia's innovation in this area

Australia could learn a lot from looking at aspects of the American model for computational science.

Both the US DOE and the National Science Foundation recognise computational science as a legitimate research discipline, and regularly offer funding opportunities. The ARC should be urged to follow suit.

The major academic-sector climate system model in the US is the Community Climate System Model (CCSM). The model is developed under interagency support between NSF and DOE. DOE supports much of the computational science aspects of the model such as algorithm development and performance engineering, while NCAR handles the bulk of the basic science input. The partnership has been in place since 2000, and to date DOE has invested over US$40M in CCSM. CCSM's governance structure engages the academic community as well as DOE and NSF scientists, and is worth considering for future CWO systems development.

The best example of combined software and performance engineering in the CWO arena is the US Weather Research and Forecasting Model (WRF; http://wrf-model.org). WRF has a governance structure similar to CCSM and thus engages the research community well. The committee may be interested to know that portions of WRF have been ported to graphical processing units (GPUs), and effort to implement more GPU kernels for WRF is underway.

It is worth mentioning both CCSM and WRF are open-source; both models are freely available for download at no cost. This means the models are widely used, and bugs are found and fixed. I believe the more "closed" approach in place here marginalises Australian researchers in the CWO field.

I urge the committee to recommend increased funding to put in place closer collaboration between the CWO and computational science communities. Sadly, the first step in this direction will involve increased funding to ensure the health of these communities separately before they are collaboration-ready.

Closing Statement

I have identified what I believe to be present and growing threats to Australia's national prestige and competitiveness in the field of meteorological and environmental prediction. I have pointed out that the major research bodies in this area (BoM and CSIRO) need a return to adequate support to fund operational forecasting, near-term research and development, and long-term capability-building. Furthermore, the problem of poor support for high-performance computing and computational science must also be addressed so that CWO researchers have collaborators for future model development. I would be eager to discuss at length this submission and happy to answer any questions it raises with the committee, up to and including appearing before a committee hearing.

Respectfully submitted 24 April 2009,

J. Walter Larson, PhD.