Speed and I/O Performance of Oceanography Code (NEMO) Enhanced by HECToR dCSE Team

Andrew Coward, The National Oceanography Centre, Southampton (NOCS)
Fiona Reid, Edinburgh Parallel Computing Centre (EPCC)
HECToR CSE Team, Numerical Algorithms Group Ltd (NAG)

HPC experts from EPCC, working under NAG's Computational Science and Engineering (CSE) support service for HECToR, the UK's national academic supercomputing facility, have improved the performance of the NEMO (Nucleus for European Modelling of the Ocean) code for scalar MPP architectures and reduced the amount of storage resource required.

Dr Andrew Coward, who is the manager of the Global Ocean Modelling Consortium on HECToR, estimated that their group used around 6M Allocation Units (AUs) running NEMO last year. Reducing the wall clock time of NEMO by up to 25%, as enabled by this dCSE project, could result in a saving in AU cost of as much as £95,000 per year (up to £400,000 for the remainder of the service), for only six months of person effort. Other consortia using NEMO on HECToR have used around 40M AUs over the same period, potentially meaning multi-million pound savings overall.

Commenting on the dCSE success, Andrew Coward said "*The NEMO code is constantly evolving with two major releases a year. Fiona's work will certainly inform our decisions on how best to run the code on HECToR but changes for particular architectures aren't guaranteed to make it into the base code unless a general benefit is evident. More groups are moving away from the vector parallel machines (mainly NECs) to MPP scalars so the code structure should begin to evolve in our favour.*"

HECToR

HECToR is managed by EPSRC on behalf of the participating Research Councils with a mission to support capability science and engineering in UK academia. The Cray XT supercomputers, located at the University of Edinburgh, are managed by UoE HPCx Ltd. The CSE Support Service is provided by NAG Ltd and ensures users have access to appropriate HPC expertise to effectively exploit advanced supercomputers for their science. A critical feature of the CSE Support Service is the distributed CSE (dCSE) programme which, through lightweight peer review, delivers dedicated performance and scalability projects on specific codes in response to proposals from users. The dCSE programme now consists of nearly 40 focused projects complementing the traditional HPC user applications support and training also provided by NAG.

The dCSE projects completed so far have delivered outstanding examples of the cost savings and new science that can be enabled through dedicated CSE effort. The NEMO project reported here adds to these success stories with a successful performance improvement.

Project Background

The objectives of this dCSE project were to improve the performance of the NEMO code for scalar MPP architectures and reduce the amount of storage resource required. NEMO is a modelling framework for oceanographic research, operational oceanography seasonal forecast and climate studies. The key tasks of the project were (a) to address the potential I/O bottleneck arising from the more complex models, higher spatial resolutions and larger numbers of processors needed by the researchers; and (b) to investigate the performance of nested models in NEMO (which enable different parts of the ocean to be modelled with different resolution within the same global model).

Dr Andrew Coward from NOCS (National Oceanography Centre, Southampton - the national focus for oceanography in the UK) was the Principal Investigator on the project. Dr Steven Alderson and Mrs Beverly de Cuevas, also from NOCS, were the Co-Investigators. Dr Fiona Reid of EPCC carried out the work in collaboration with the NEMO developers and the NAG CSE team.

NEMO

NEMO is a state-of-the-art modelling framework for oceanographic research, operational oceanography seasonal forecast and climate studies. The framework allows several ocean related components e.g. sea-ice, biochemistry, ocean dynamics, tracers etc to work either together or separately. Further information on NEMO can be found at http://www.nemo-ocean.eu.

Project Results

The investigations into gridding and the removal of land-only cells under this dCSE project resulted in significant reductions to the resource usage for a given simulation, by as much as 25% at larger processor counts, and a reduced runtime. NEMO has been converted to use netCDF 4.0 for its main output files resulting in a reduction in output file size of up to 3.55 times relative to the original netCDF 3.x code. It is expected that a production length research simulation should benefit in reduced run time also, due to fewer I/O bottlenecks resulting from the reduced I/O data sizes.

NCAS (National Centre for Atmospheric Science), another significant consortium of users on HECToR, only uses code that has been filtered through the main NEMO developers and the UK Met Office, which will take some time. Dr Coward has switched permanently to using the new I/O option as disk space is a valuable and limited resource. It also allows either more or longer time series to be kept on-line for analysis by the wider Ocean modelling community.

A full technical report can be found at http://www.hector.ac.uk/cse/distributedcse/reports/

For more information contact: HECToR CSE Team
The Numerical Algorithms Group Ltd, Wilkinson House, Jordan Hill Road, Oxford, OX2 8DR, United Kingdom
Telephone: 01865 511 245 Email: hector-cse@nag.co.uk Web: http://www.hector.ac.uk/cse/