

Introduction

NEMO (Nucleus for European Modelling of the Ocean) is a 3-dimensional ocean model that is used for oceanography, climate modelling as well as operational ocean forecasting. It includes submodels that describe sea-ice and biogeochemistry. Many processes are parameterised, e.g. convection and turbulence. It is used by hundreds of institutes all over the world. The open-source code consists of 100k lines of code, is developed in France and the UK by the NEMO development team and is fully written in Fortran 90. The MPI (Message-Passing Interface) paradigm is used to parallelise the code and depending on the configuration of the model, it can run on a single processor or scale up to 1000 or more processors.

Model description

NEMO is a finite-difference model with a regular domain decomposition and a tripolar grid to prevent singularities. It calculates the incompressible Navier-Stokes equations on a rotating sphere. The prognostic variables are the three-dimensional velocity, temperature and salinity and the surface height. To further simplify the equations it uses the boussinesq and hydrostatic approximations, which e.g. remove convection. It can use a linear or non-linear equation of state. The top of the ocean is implemented as a free surface, which requires the solution of an elliptic equation:

$$\left[\nabla \times \left[\frac{1}{H} \mathbf{k} \times \nabla \left(\frac{\partial \psi}{\partial t} \right) \right] \right]_z = [\nabla \times \bar{\mathbf{M}}]_z$$

to find $\partial_t \psi$. For this purpose, it uses either a successive overrelaxation or a preconditioned conjugate-gradient method. Both methods require the calculation of global variables, which incurs a lot of communication (both global and with its nearest neighbours) when multiple processors are used.

Problem

Both methods to calculate the free surface are iteratively solved (see the manual, section 10.7), and are implemented as single subroutines of less than 200 lines. Their convergence rate is very slow and on average they need hundreds or even thousands of iterations to reach a satisfactory tolerance. At this moment, the scalability of NEMO is limited by the communication overhead of the free surface solver. As new computer architectures will consist of many more cores (100k+), it is imperative to improve the scaling of NEMO.

Goal

To find new algorithms or to change the existing algorithms that improve the scalability of the NEMO model. This could consist e.g. of implementing a deflation method or improving the preconditioner. SARA will provide technical assistance with analysing, implementing and running NEMO.

Contact

For further information please contact John Donners (john.donners@sara.nl or 020-5923055).

Links

<http://www.nemo-ocean.eu>
<http://tinyurl.com/y96s8ol>

Website of NEMO
NEMO documentation with detailed description