Optimization of an ocean model using performance tools

Oriol Tintó Prims
Barcelona Supercomputing Center
C/ Jordi Girona, 29
08034 Barcelona (Spain)
+34 93 413 40 51
oriol.tinto@bsc.es

Miguel Castrillo
Barcelona Supercomputing Center
C/ Jordi Girona, 29
08034 Barcelona (Spain)
+34 93 413 40 51
miguel.castrillo@bsc.es

Harald Servat
Barcelona Supercomputing Center
C/ Jordi Girona, 31
08034 Barcelona (Spain)
+34 93 413 77 16
harald.servat@bsc.es

German Llort
Barcelona Supercomputing Center
C/ Jordi Girona, 31
08034 Barcelona (Spain)
+34 93 413 77 16
german.llort@bsc.es

Kim Serradell
Barcelona Supercomputing Center
C/ Jordi Girona, 29
08034 Barcelona (Spain)
+34 93 413 40 51
kim.serradell@bsc.es

Oriol Mula-Valls
Barcelona Supercomputing Center
C/ Jordi Girona, 29
08034 Barcelona (Spain)
+34 93 413 40 51
oriol.mula@bsc.es

Francisco Doblas-Reyes
Barcelona Supercomputing Center
C/ Jordi Girona, 29
08034 Barcelona (Spain)
+34 93 567 99 76
francisco.doblas-reyes@bsc.es

ABSTRACT

Do the model users know about the computational performance of the models they are using? Do the model developers know about the computational performance of the models they are developing? The landscape of the computational science has a lot of examples of scientists using huge amounts of resources to run models without worrying too much about computational performance. The proper use of these resources usually is not the main priority and is not strange at all if someone that finally decides to study the computational performance of his model finds out that it is far from being good. In this work, we would like to show that a performance analysis of scientific models is useful and, without any doubt, is worth.

We analyze and optimize the NEMO ocean model, widely used in dozens of projects (in the CMIP5 project, 5 of the 28 earth system models participating were using NEMO as the ocean engine) consuming billions of computing hours each year. The NEMO model is an state-of-the-art ocean Global Circulation Model used for oceanographic research, climate studies, seasonal forecast and also for operational oceanography where the time-to-solution is crucial. The total amount of computing resources used in NEMO simulations can easily exceed a billion of computing hours and for this reason a proper analysis and optimization of the computational performance of the model is mandatory. The development of NEMO began at the eighties and many scientists have been involved in its development.

In order to increase the model's performance and reduce the time-to-solution, we start with a performance analysis on the model. First of all we must take into account that the NEMO model includes many different modules and options that have an impact in the computational behavior of the application, so we need to set-up correctly which simulation we want to perform. In this work, we studied the standard configuration, using the ocean and sea-ice modules on a low resolution grid. The advantage of performing a strong scaling analysis using a low resolution grid is that the problems related with small subdomains that limit the scalability of the model are already visible using a small number of CPUs, so the issues that will be relevant at higher resolutions using hundreds of thousands of CPUs can be faced much more easily.

In a preliminary analysis we measured the model's simulation speed for different number of processors between 16 and 128 cores. After that and with a first guess of how is the model's performance, we carried out our analysis using the open-source performance tools developed at the Barcelona Supercomputing Center (BSC). These tools demonstrated to be really powerful and useful for scientific models. By using them not only is possible to analyze both computation and communication, but also to dissect the information with a very high precision. Therefore, the degree of detail can vary from computing general metrics like parallel efficiency or load balance, to computational performance metrics for a specific line of code. In addition to these features, it is also possible to simulate how the performance will evolve in different scenarios, so it is possible to know how a model would run in a different architecture or the sensitivity to different parameters, like the network latency, the processor speed and many others.

To proceed with this analysis it was necessary to collect information in traces about different NEMO executions and study aspects regarding communication and computation with high precision, in order to find where and why the model is not scaling.

The poster illustrates a real case study for this methodology, which we use to identify bottlenecks, develop optimizations and evaluate their impact on the performance. We put an emphasis on the results, which we can summarize in one sentence: the optimized version of NEMO can simulate 37% faster than the previous version.

This methodology can be extended to any earth system model running in an HPC platform.

All the simulations were done in the Marenostrum 3, hosted by the BSC. Marenostrum 3 is a 1.1 Petaflops supercomputer composed by 3056 nodes with 2 16-CPU SandyBridge-EP processors.

Categories and Subject Descriptors

D.2.8 [**Metrics**] Performance measures

D.4.8 [**Performance**] Measurements, Modeling and prediction, Simulation

J.2 [**Physical Sciences and Engineering**]: Ocean Global Circulation Model

General Terms

Performance

Keywords

Performance Analysis, Performance Optimization, High Performance Computing, Ocean, Climate, Earth System