Geophysical Research Abstracts
Vol. 20, EGU2018-13493, 2018
EGU General Assembly 2018

Towards an optimal use of numerical precision in Earth Science models: the case of NEMO

Oriol Tintó (1,2), Mario Acosta (1), Miguel Castrillo (1), Kim Serradell (1), Alícia Sanchez (1), Ana Cortés (2), Francisco J. Doblas-Reyes (1,3)

(1) Earth Sciences Department, Barcelona Supercomputing Center, Spain, (2) Universitat Autònoma de Barcelona, Spain, (3) Catalan Institution for Research and Advanced Studies, ICREA, Spain

Computing power and our ability to make use of it constrain the quality of weather and climate predictions. The lifespan of models, typically much longer than systems, evince that optimizations should not focus on specific platforms but on long-term system trends. Exploiting vectorization and reducing data movement is important today but in next-generation platforms will be crucial. One field of research that could help scientific codes to better exploit these characteristics is the use of mixed precision algorithms.
It has been a widely extended practice in scientific computing to use 64-bit to represent data without considering which level of precision is really needed. In many applications, 32-bit should provide enough accuracy, and in other cases, 64-bit is not enough. In climate science, the inherent difficulties collecting data imply a considerable level of uncertainty, which suggest that the general use of 64-bit to represent the data may be a waste of resources, while on the other hand, some specific algorithms could not afford a reduction on the precision used and instead they could benefit from an increment of it. These factors suggest that in the future more attention has to be paid to the floating point precision in scientific software, in order to use resources wisely and also avoid losing accuracy.
There is a direct relation between the precision used and the cost of a computation. Current processors have vectorial instructions for 32-bit and 64-bit data, performing the former the double of operations per cycle than the latter. Moreover, next generations of processors will also include 16-bit. Furthermore, the gap between processor and memory speeds implies that the time to bring the data to the CPU can be as long as the time used by the CPU to compute the operations itself and in some cases, data movement becomes the computational performance bottleneck. In conclusion, in both CPU and memory bandwidth bound computations, using more precision to represent the data requires more resources, resources that potentially can be saved if the required precision is properly analyzed.
Our work aims to find out which precision is required in each part of an ES model, allowing to save resources, speed up the models and allow new kinds of experiments. To do so, we use an emulator that allows us to easily explore the variable space to identify the precision needed in each part of the code, making possible to keep the quality of the results reducing the precision only where doing so does not impact the model outputs. We have demonstrated the potential of this approach using this methodology with the NEMO ocean model.