Geosci. Model Dev. Discuss., 8, C4016–C4017, 2016
www.geosci-model-dev-discuss.net/8/C4016/2016/

Geoscientific Model Development Discussions

Interactive comment on "Performance and results of the high-resolution biogeochemical model PELAGOS025 within NEMO" *by* I. Epicoco et al.

G. Williams (Referee)

gethin.williams@bristol.ac.uk

Received and published: 15 February 2016

Summary ——-

The scalability of an on-line coupled model–PELAGOS025 and NEMO–is investigated on two different compute architectures: an IBM BlueGene/Q and an IBM iDataPlex using Intel SandyBridge processors.

Comments ——

This paper constitutes a very welcome investigation of the impact upon parellel scaling of coupling two component models together in an on-line fashion. Earth System Models often represent different climatic factors by coupling together different sub-models into a coherent whole. However these components sub-models may well make rad-

ically different demands of the computation, memory and storage capacities of the computer used to run them. The key finding of this paper is that the best domain decomposition for the overall model will vary, depending upon which sub-models have been coupled together.

Also welcome in this paper is an aspect of performance analysis using a code profiler. Profiling tools tell us which subroutines account for the majority of the run-time of the model and so help focus our efforts in code optimisation, when trying to reduce the time to solution.

While this section contained some interesting comparisons of the most costly routines, as experienced on the two compute architectures, and the relative scaling of this routines, I would welcome a higher level review of the profile information. For example, I think readers would benefit a great deal from a short analysis of just a couple of the most costly routines, explaining why each of them consumed a lot of the run-time. Were the routines compute or memory-bandwidth bound? Is there a better implementation which could address the observed bottle-neck?

It was observed that the coupled model utilised only 2.7 Gflops per node on the BlueGene/Q, which is only 0.25% of the peak capability of 204.8 Gflops/node. This is a startlingly low utilisation of the processor by any standards and could serve as a salutary example for programmers. However, I have to disagree with the analysis of the limiting factor for scaling on the BlueGene. While it is true that using all 4 rather than just one of the hardare threads will surely help, a naiive prediction is that this would yield 1% of peak as a utilisation. That is 99% of the CPU cycles are still essentially wasted. Surely there must be some other large bottleneck in the code?

Interactive comment on Geosci. Model Dev. Discuss., 8, 10585, 2015.