

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

I. Epicoco et al.

italo.epicoco@unisalento.it

Received and published: 15 March 2016

Dear Referee

Thanks for your comments. Our reply is online

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 im-
C4230

plementation which could address the observed bottle-neck?

Actually there is a previous work on the NEMO code explaining why the routines consumed most of their run-time. In that paper the roofline line model is used to explain the code behaviour and to identify some actions for improving the performance. In our work we have just referenced that paper and reported the main results.

It was observed that the coupled model utilised only 2.7 Gflops per node on the Blue- Gene/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 hardware threads will surely help, a naive 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?

We agree with you that the lack of performance are not related only to the low exploitation of the hardware threads. Other critical factors are low level of vectorization introduced by the compiler and low level of operational intensity which limits the performance to the memory bandwidth.

We have modified the paper according to your comments. Here attached you can find the updated version (latexdiff has been used to highlight our updates)

thanks again

Please also note the supplement to this comment:

<http://www.geosci-model-dev-discuss.net/8/C4230/2016/gmdd-8-C4230-2016-supplement.pdf>

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

C4232