Geoscientific
Model Development
Discussions
EGU
Open Access

Interactive comment on "Earth System Modelling on System-level Heterogeneous Architectures: EMAC (version 2.42) on the Dynamical Exascale Entry Platform (DEEP)" *by* M. Christou et al.

Anonymous Referee #2

Received and published: 3 May 2016

The article adresses the important issue of the performance of Earth System Models on future heterogeneous cluster architecture. Within the DEEP project the GCCM EMAC has been ported to a cluster booster system.

Initially, the authors provide a performance analysis to identify the dominant factors in computational workload and to pin down the bottlenecks that prevent scalability. They conclude that with respect to performance issues the choosen EMAC setup could be split into two parts. The KPP based chemistry integrator MECCA as the major contributor to the computational workload is well suited for massive parallel treatment and the candidate for the booster architecture. Whereas the rest has limited scalability due to strong coupling and high communication demands and remains on the cluster.

The authors briefly present the refactoring that is needed to remedy the issues preventing the code to run successfully on the cluster booster system. With this changes the model shows impressively enhanced performance and scalability. In a nutshell the article provides recipes to successfully port ESMs on the future heterogeneous architecture. This is an important result that definitely deserves to be published.

Nonetheless, I require major revisions. There is a considerable lack in detail, important information is missing. The reader simply can't really reconstruct what the authors have done, the results are not reproducable. The isssues of the code refactoring are just presented as keywords. The reader does not really know what has been done. Every code refactoring topic has to be exemplified. The authors should present the old structure (not the whole code, but the essential programm structure that needs to be modified) in comparison with the new logic. This demands a considerable rewriting of the article.

In summary, I strongly support to publish this article. But major revisions are required to assure reproducibility and to render it really beneficial for the modelling community.

Specific Comments

2.1 Phases

-page 3, lines 9-15 : Redundant, you have said that already in the introduction.

-page 3, lines 26-27 : "Furthermore, even a coarser..." What does this sentence mean?

2.2 Dominant factors

Fig. 2 is not mentioned.

-page 4, line 8 : "data size scales with the the square of model resolution" This is only true, if the vertical resolution is unchanged.

-page 4, lines 11-18: The discription of Fig. 4 is very confusing and maybe wrong. Doesn't it start with the transformation from spectral space into Eulerian space, fol-

lowed by grid point calculations and ending with transformation back into spectral space?

-page 4, lines 23-26: "Model performance depends largely on a virtual ..." What do these sentences mean? If the model performance largely depends on that, this should be explained better.

-page 5, line 3 : Fig. 5 is not really needed. The numbers could be mentioned in the text.

2.3 Scalability considerations

In general, chapters 2.2 and 2.3 deal with the same topic. Maybe just one chapter is needed.

-page 6, line 9 : It isn't clear how Fig.7 is made. Does it show the integrated computation time of a column?

-page 6, lines 10-15 : Has been said before.

-page 6, lines 16-20 : Fig. 8, what is the configuration of MareNostrum 3? What is the number of MPI-processes on a node?

3.1 Intranode taskification

Doesn't Fig. 10 contradict with Fig. 9 ? ECHAM in EMAC is not affected by OmpSs and with 32 Nodes 40% of CPU time is with ECHAM. This amount strongly increases with the number of nodes. Why does Fig. 10 not reflect this?

-page 7, lines 5-6 : Refactored, but how? Please exemplify.

3.2 Internode taskification

Please exemplify. What has been changed in the code? Please provide examples from the code. The reader can't really reproduce or reconstruct what you have done.

3.3 Vectorisation

C3

-page 8, lines 21-27 : Redundant, has been said before.

-page 8 line 28 - page 9, line 5 : What have you done? Have you changed the original code? Please exemplify.

Fig. 11 is not really necessary, the numbers are already in the text.

-page 9, lines 11-13 : What does this mean? How has the code been changed? Please exemplify.

4 Attainable performance

-page 9, lines 18-29 : This is very confusing. How is this done? Please provide equations.

5 Conclusions

-page 11, lines 4-7 : How is this done? This has to be exemplified in the article.

Code Availability Is the recoded EMAC from the DEEP project available to the public?