

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 #1

Received and published: 12 April 2016

The article reports about one part of the DEEP project examining the approach of heterogeneous cluster-computing for Earth System models. In particular, this article is about the GCCM EMAC. A performance analysis shows, that the chemistry submodel MECCA is the bottleneck of usual EMAC simulations. As MECCA is solving the chemistry within each grid box independently, it is embarrassing parallel and thus the ideal candidate for application on a Booster architecture.

The article addresses a very important issue of Earth system modelling. The fact that these codes have a long history brings about that they are not at all optimised for modern computer architectures. Thus the refactoring of the code, to use the possibilities provided by current and future computing architectures is a very important issue.

C1

Nevertheless, there are some issues that need to be improved upon revision. First of all the specific model configuration used for the scaling tests is not provided. But the setup heavily influences the performance of the model. Additionally, the explanations provided in this article are only correct for certain model setups. Therefore the setup needs to be provided. Secondly, this is a GMD article and the authors classified it “Development and technical paper”. Thus I expect the authors to provide much more details about the developments themselves. How exactly have they been implemented? In general, the performance analysis is as long as the description of the developments. The latter should be the main part of the article from my point of view.

Below I give the details about the points raised above and name some additional issues. In summary, I am very much in favour of publishing this article. Nevertheless, major revisions are required.

Specific Comments

1. The authors often use the term “meteorological model” when they mean the dynamical core of the model. Meteorology comprises the dynamical processes as well as the physical processes such as cloud and precipitation formation or radiation. EMAC uses the dynamical core of the ECHAM model, but all physical processes are – by now – modularised as MESSy submodels. In ECHAM the dynamical core operates in spectral space, the physical processes are implemented in grid point space. All MESSy submodels (so far) are implemented as additions in the grid point space. Therefore the authors should be consistent in the terms they use throughout the article. Please change accordingly:

- page 2 line 12
- p.3, l. 2 + l. 6
- p.5 l.24
- p. 10 l.10

C2

- Figure 2, left box

This list is not complete. So please check throughout the article.

2. To really understand the performance analysis, the full EMAC setup should be listed somewhere. If an ESCiMo-Setup or another published setup was used, this could simply be cited. If not, I would prefer to have a description, containing all details relevant for the publication, in the appendix. A zip-file containing the full namelist setup should be provided in the supplement.
3. Connected to the previous point: You characterize MECCA as the submodel computing “the chemical kinetics of the homogeneous gas-phase chemistry of the atmosphere,” ... (e.g., p.3 l.13; p.4 l.29)
In all MESSy setups which are not only focussed on the lower troposphere, heterogeneous processes on ice cloud are included via reaction rates provided by MSBM. As I do not know the specific setup of this study, I can not judge, if this statement is correct with respect to it. As a general statement it is definitely not correct.
4. If MECCA is employed with a pure gas phase mechanism in this test setup, what will be the result, if heterogeneous reactions are included? This would establish a second source of imbalance caused by the appearance of PSCs.
5. SCAV uses KPP as well. In maximum it is called four times during one time step (for grid/subgrid scale liquid/ice clouds). Here the load imbalance is caused by the distribution of clouds over the model domain and additionally, SCAV is column bound. At least in the conclusion or outlook I'd like to see a statement, how easily your developments could be applied to SCAV and (maybe) if you expect performance gains for SCAV as well.
6. Chapter 3: Model developments
What I am really missing in this chapter is the “development” i.e., an explicit

C3

mentioning of the code changes required. You describe them superficially with some words, what - at least for a MESSy developer / user - is really interesting is what the code changes look like. For which of the described changes did you change which code parts? Do they require changes in the MESSy submodel interface layer (SMIL) or in the core layer (SMCL). Here it is most interesting, if these developments require a change in the automatically generated code. If yes, MECCA knows two stages of “automation”. First KPP produces the code automatically from the equation file. Secondly, KP4 can be applied in order to remove the indirect indexing and expand the original KPP code by an additional dimension to enable a better performance due to better cache usage. If there are changes in the automatically generated code parts, did you change the scripts performing the automation or did you just change one MECCA setup (and running xmecca once would destroy all you efforts)?

I assume that this information is only of interest for MESSy or MECCA developers / users and not to the general readership. Therefore I recommend to add the most important information to the article itself and to add a supplement describing the changes in more detail and providing information about how to use it (something like a user manual), which also is in accordance to the GMD guidelines.

Additionally, a description, how to optimise the setup as described in p.8 ll.10-12 should be provided.

7. p. 8 ll. 1-19: It is not clear to me, what these paragraphs are about: ll.1-6 are a repetition of what was written earlier. ll. 7-12: claim that the code can be optimised tailormade for each architecture but does not tell how. ll.13-19 simply state that these changes increase the performance, but does not show any proof. Somehow I miss the point here...
8. Section 3.3
In this section the authors must be much more precise. It is not clear, what exactly

C4

the authors are discussing here. Principally, the ECHAM (and thus the MESSy) grid point code is decomposed in a way of artificial latitude bands. ECHAM provides a so-called "local loop": here a loop over the second horizontal dimension is established, reducing the size of the fields forwarded to the individual submodels within the local loop by one horizontal dimension. The length of the remaining horizontal dimension is called vector length and can be chosen by namelist for optimisation on different computing architectures.

Thus MECCA, which is called in the local loop is called with this one horizontal dimension only. For a better performance the original KPP output can be expanded by this additional horizontal dimension times the vertical dimension. This code is automatically produced by calling KP4. It is not clear to me, about which of these different aspects the authors talk exactly. So please clarify this issue.

9. Code Availability: You state the general terms for the MESSy code. Nevertheless, you are presenting new code developments, therefore it would be good to know, if it will become part of the official MESSy version soon and if it is possible to get hold of the code prior to this.

Minor issues

- paragraph p.3 ll.22-27, it would be good to have examples here. You provide them on page 5 last two lines, but they are already here useful to understand which kind of processes you are talking about.
- I have mixed feelings about Fig. 1. It looks nice, but is it really required? Additionally, as processes as deposition occur on the right hand side, this figure is in contradiction to Fig. 2.
- p.4, l.1: Shouldn't it be Figure 2?
- p.4, ll.8-10: it scales with the square of the horizontal resolution.

C5

- p.5 l.13/ Table 1: Table 1 does not give a clue about the model setup, it only contains different numbers calculated from the resolution of the model. In my opinion the first column does not contain any information I would expect in a table. That the numbers indicate the possibility for strong scaling should be stated in the text not in the table itself. What is meant by "42 coefficients"? Which coefficients? Do you mean because of T42? Than the reader anyhow understands what T42 refers to, or he/she does not become any wiser by reading "42 coefficients" (my opinion).

The table would be much better readable, if columns and rows would be switched

- Table 2: Here the last statement for table 1 applies even more: the table is much better readable, if switch columns and rows are switched.
- p.6 ll. 3-6: This is not fully correct. From your description I visualise a decomposition where the domain is split up in rectangular grid boxes. But ECHAM is using a decomposition where each task gets two (independent) latitudinal bands (of arbitrary length in the longitudinal range). These bands are usually not even adjacent to each other.
- p.6 l. 8: add location of natural and anthropogenic emissions.
- p.6 ll. 16-19: Why are you only describing the results for ECHAM seen in Fig. 8? Add something about "computing time for MESSy is still decreasing" and "computing time for MECCA decreases stronger than MESSy". Additionally, please point to the logarithmic scale and to the fact, that "MESSy" means without MECCA.
- p.6 ll.21-23: I am missing a conclusion, you just describe the plot.
- p.7 last line: "each time step", "step" is missing.
- p.8 l.31: What is meant by "domain-specific language" ?

C6

- p.9 l.29 ; It would be helpful for the reader if you could mention the colours of the respective lines in the text. Do I assume right that “MPI” is the sum of ECHAM+MESSy and MECCA? First I thought it is the MPI communication time. Please try to clarify your description.
- Fig. 1: For me this figure produces more questions instead of assisting in understanding the distribution on the two different computing architecture parts. The “base model cluster” part contains a picture of a cloud, i.e., physical processes, and the “atmospheric chemistry booster” part does not only contain the chemical mechanism, but also deposition, thus it is not quite clear where the separation between cluster and booster should appear.
- Fig. 3: It is not clear on which ground the colours of the boxes are chosen. Personally I think, the figure overemphasises the dynamical core (including the transformations from grid point to spectral and vice versa). Because the grid point calculations contain much more sub processes which are completely left out by this figure.
- Fig. 5: Please provide a more descriptive caption for this figure. Not every reader is familiar with Scalasca output.
- Fig. 7: Personally I think the ferret labels should be removed from the graphic. You can acknowledge use of the ferret program in the acknowledgements.
- Fig. 8: Not the impact on run time, but the run time itself for different numbers of nodes is shown in the figure. I assume that all the tests are performed without any output. Could you mention this somewhere (e.g. in the setup description to be added to this article?) How exactly do you deduce the time for MESSy? Do you assume GPC is MESSy (including MECCA)-time and the rest is ECHAM time?
- Fig. 11: Its MESSy not MESSY (2x)

C7

Interactive comment on Geosci. Model Dev. Discuss., doi:10.5194/gmd-2015-262, 2016.