

Response to GMD-2016-47-RC2 (Sophie Valcke)

Dear Ms. Valcke,

we thank you for the constructive and detailed comments and questions and hope to give an easy to follow and satisfactory answer. We tried to consider all of your points and some more with the article revision. We also kept some redundancy of the basic aims in order to facilitate following the idea of the article.

Please, keep in mind that all references to the paper given in the following are references to the revised version of the article. Your comments are given in black, our answers in blue.

Best regards, Andreas Will

General comments

This paper presents a detailed analysis of the performance of coupling configurations involving the COSMO-CLM regional model. An extensive literature exists on the performance analysis of individual models or codes but there is much less published on the performance of the coupled system and on the coupling aspects per se. This paper addresses this gap and as the work onto which it is based is sound, it deserves publication. However, I consider it needs major revisions before being published.

Answer: We are very pleased to know that our paper is interesting from your point of view and we did the best to answer your questions.

Specific comments

My first main concern is about the way the results on the optimum configurations (p.25, section 4.4, Fig. 5 and Table 8) are presented; currently, they are difficult to appreciate because there is, on one hand, a lot of information (sometimes superfluous), and on the other hand, some missing details.

Answer: Thank you for that comment. We revised chapter 4, improved the figure description, separated figures 3-4 “time to solution” and “cost” from figure 5 “parallel efficiency”, which belongs to finding of optimum configuration, we introduced a separation of extra cost in 5 components and revised table 8 accordingly considering the reviewer comments. In particular, the last section 3.3 of table 8 is revised and all numbers are presented in a consistent and unique way as % of cost of optimum configuration of CCLM stand-alone. We removed the figures and the discussion of possible improvements of CCLM+MPI-ESM.

First of all, I do not understand why the cost of the CCLM part of CCLM-CLM and CCLM+VEG3D are about doubled compared to the costs of the CCLM stand alone or compared to the CCLM-NEMO-MED12 coupling. On p.25, l.854, it is stated “The corresponding costs are about double the costs of the stand-alone reference: 512.0 and 473.6 CHPSY, respectively”. Can you give an explanation? Is it linked to the fact that CCLM runs in SMT non-alternating mode in stand-alone and in the CCLM-NEMO-MED12 coupling whereas it runs in SMT alternating mode in the CCLM-CLM and CCLM+VEG3D couplings? If so, it should be stated in the text.

Answer: We agree with the reviewer that it is difficult to follow the discussion and improved it (hopefully). You find a new paragraph (section 4.5. line 927 ff) clarifying which the dominating components of extra cost of coupling are. It turned out, it is mainly due to using ST instead of SMT mode and of the double number of cores.

“On p.25, l.854, it is stated “The corresponding costs are about double the costs of the stand-alone reference: 512.0 and 473.6 CHPSY, respectively”. Can you give an explanation? Is it linked to the fact that CCLM runs in SMT non-alternating mode in stand-alone and in the CCLM-NEMO-MED12 coupling where as it runs in SMT alternating mode in the CCLM-CLM and CCLM+VEG3D couplings?”

Answer: You are right to a wide extend. See previous answer.

The fact that the COSMO version used for CCLM+CLM is different from the COSMO version used for CCLM+VEG3D and that the results presented for CCLM+VEG3D are in fact not the optimum ones (128 cores were chosen to be able to compare with CCLM-CLM) is disturbing and the paragraph p.25 l-843-855 is difficult to understand (same thing for p.27, l922-923). I am not sure on how to correct this but this should be simplified maybe simply by removing results for CCLM-CSM?

Answer: Thank you for the comment. We conducted additional measurements comparing cosmo_5.0_clm1 used in CCLM+CLM and cosmo_4.8_clm17 used in CCLM+VEG3D (on another machine since blizzard is not available anymore). This exhibited 45% higher cost of 5.0. We revised the result presentation and in particular this paragraph .

It is not clear on figure 5 if the time to solution includes or not the OASIS interpolations. Can you clarify this? It is written in Table 8 caption that it does not but it should be stated either in Fig 5 captions and in the text, stressing that the interpolation time is relatively small anyway (as quantified in Table 8). Can you also specify that the OASIS interpolation times are provided directly by the lucia tool in table 8 captions (even if this is mentioned in the text on p.21)

Answer: The “computing time” measured by LUCIA and by the “time” function includes interpolation time. We introduced a clear analysis of extra cost, corrected the caption of table 8 and extended the caption of (now) Figure 6. The OASIS interpolation is now given clearly for each coupling.

My second concern is about the definition of the criteria to identify the optimum configuration, which are not clear:

In section 4.2, please specify what you mean by “each component’s gain in speed, Compared to its speed on one node, outweighs the increase in costs.” The units of speed are not the same as the units of cost so they cannot be compared directly. Are you considering the relative gain in speed (in %) is compared to the relative increase of cost (in %)? To help the understanding, one practical example with numbers should be given (maybe at the beginning of current section 4.4?), for example the steps that lead to the identification of the optimum configuration for the CCLM

Answer: Thank you for this comment. We agree that the description of the strategy was not sufficient to understand and reproduce the results presented. We revised (now) section 4.3 describing the strategy and give the numbers in the new section 4.4 describing the application of the strategy.

The optimum configuration is always a compromise between efficiency (depending on models scalability) and availability of resources or time to solution and cost. It is maybe not possible to give an objective definition of what this compromise should be. Thus we introduced a parameter for that compromise, the parallel efficiency: “The optimum configuration is found by starting the measuring of the computing time on one node for all components, doubling the resources and measuring the computing time again and again as long as all component parallel efficiencies remain above 50%. The threshold of 50% is subjective and can be defined by the user, i.e. one could decide to stop at a higher parallel efficiency if costs are a limiting factor.” This definition is the same for both concurrent and sequential

configurations. An additional criterion is introduced, if the increase of cost has no impact on time to solution, in other words, if there is no scalability. In this case the parallel efficiency down to 50% is not used.

I think section 4.2 would be better to place just before 4.4 (i.e switching current 4.2 and 4.3 sections)

Answer: we follow the advice and switched the two §

p.22, l.749: This constraint is effective only for sequential coupling so it should be moved to the paragraph currently starting in line 754.

Answer: We have rewritten the sentence considering the reviewers suggestion (line 826 ff).

p.22 l.763 & l.779: It is not clear who or what decides if the costs are a limiting factor or not. Did you consider the costs was a limiting factor in your identification of the optimum? If so, what was the limit? This should be clarified.

Answer: Thank you for the comment. We rewrote the paragraph (line 836-849). We didn't introduce any other criterion but 50% parallel efficiency and, lowest cost, if no scalability is found. The application of the criteria is described in section 4.4 for each coupling in detail.

My third major concern is about section 4.5.2. I find this section not really relevant in the context of this paper. Of course, one can always get better fictive results by neglecting costly or badly written parts of the code!

Answer: We thank the reviewer for this suggestion. Section 4.5.2 is removed. Figure 6 and 7 are removed. Instead the extra cost of coupling for CCLM-MPIESM are discussed in section 4.5. line 984 ff and in section 4.6.

Then I have the following major remarks:

In general, I think the text is quite heavy with many repetitions. I suggest to make it lighter and more "right-to-the-point". In particular, section 2 describing the components could be reduced and the appendix A, that is not essential to the understanding of the paper could be given as supplementary material.

Answer: We would like to thank the reviewer for highlighting this important aspect of readability. The authors discussed this aspect again. Interestingly, reviewers 1 and 2 have dissenting opinions. You suggest reduction of the content and focusing on finding of an optimum configuration. The second reviewer suggests adding more details on configuration for asserting reproducibility and adding a discussion of the impact of configuration of model physics and dynamics on cost and time to solution. Considering the online publication form, we would like to keep section 2 in the paper. We revised the introductions of the sections in chapter 2 indicating that it is not essential for readers interested in the strategy of finding an "optimum configuration" only. The text is kept as it was with minor corrections. The introduction of the Appendix 1 is revised as well. It is essential for understanding of the coupling physics and dynamics and it is kept as appendix of the article since it does not increase the size of the PDF significantly and allows to keep everything in one document.

With which version of COSMO were the CCLM stand-alone tests done? Is it cosmo_4.8_clm19 like for all coupling but CCLM-CLM? This should be clarified in the text.

Answer: Yes, see our answer below. We modified § 2.1, line156, and § 2.6 accordingly.

p.3, l.69: Please give some details on why the MESSy approach was not considered.

Answer: The CCLM couplings available with MESSy and OASIS are different. A comparison between MESSy and OASIS is planned for CCLM+MPI-ESM couplings. This requires additional developments for a fair comparison which are not finished yet.

p.5, l.597: You mention a “coupling weight” increasing to 1 with time but this coupling weight is not described. Can you explain with more details how it works?

Answer: Thank you for this suggestion. The function used is given now in the text, line 601-605..

p.21, l.724: can you justify the formula used to approximate the time for 40 levels based on the time for 45 levels; why not simply use : $T_{40} = T_{45} \times 40/45$

Answer: The scaling of 80% of the computing time (and not 100% as suggested by your comment) is already explained in the footnote.4, line 760.

p.24, 806-807: I do not understand how one can conclude that “COSMO-CLM inv ST and SMT mode exhibits a very similar PE for the same number of processes ...” The curves are distinct. Do you mean that we should compare the SMT results for a specific number of cores with the ST results with twice as many cores (to get to the same number of processes)?

Answer: Yes, this is what we wanted to say. We agree that the explanation is weakly understandable and improved it. See line 808 ff.

p.24, l.808: you write “an increased loss of PE between 160 and 80 grid points per process.” but the reader cannot directly infer the number of grid points per process given the number of cores (which is the information provided on the figure), so the corresponding number of cores should be mentioned to help the reader.

Answer: We thank the reviewer for this comment and revised the paragraph for a better explanation. See line 814.

16. p.24, l.813-814: I do not fully understand this sentence. First I am not sure what the “component interface” is. Is it the coupling interface? If so, I do not understand how to reconcile this with the fact that the coupling interface time probably includes the time for interpolations (which are done either on the source side before the sending or on the target side after the receiving) and that the “time to solution” does not.

Answer: We thank the reviewer for raising this point. We agree that this sentence is not correct. We revised this paragraph and moved this explanation to “extra time and costs” in §4.5 , line 1009.

p.24, l.815-819: I do not understand the meaning of the sentence “Hereby, the number of cores and the threading mode (ST or SMT) are kept constant.” I propose to remove this sentence and rewrite the following ones as: “ COSMO-CLM components of concurrent couplings should be compared to stand-alone COSMO-CLM in SMT mode because in both cases two threads per core are used to run COSMO-CLM. Conversely, COSMO-CLM components of sequential couplings should be compared to stand-alone COSMO-CLM in ST mode because in both cases only one thread per core is used to run COSMO-CLM.” if I am right in my interpretation.

Answer: Thank you for this suggestion. The reference for each coupling is described now as suggested in section 4.4

p.24, l.826-827: It is written “However, as mentioned in section 2.6 CLM is coupled to cosmo_5.0_clm1 model version which is a more recent version than cosmo_4.8_clm19 used for all other couplings “ but I don’t see this mentioned in section 2.6.

Answer: We thank the reviewer for this remark. We corrected section 2.6 accordingly.. “The model version cosmo_4.8_clm19 is the recommended version of the CLM-Community (Kotlarski et al., 2014) and it is used as basis of the development of the couplings. CCLM as part of the CCLM+CLM coupled system is used in a slightly different version (cosmo_5.0_clm1). The way this affects the performance results is presented in section 4.5, line 954 ff. In addition, the reviewer can see, in the figure below, a scalability comparison between the 2 versions. This reveals (even though the machine is not the same than the one used in the article) the cost of the 5.0 version are 45% higher than for 4.8.




Figure 1: Time to solution of 5.0 and 4.8 COSMO-CLM versions in dependence on core number on Cray XC30 at CSCS, Lugano.

p.26, l.898-900: I propose to rephrase these two sentences for “It is not surprising that the couplings with soil-vegetation models shows only moderate extra costs as they replace the use of TERRA, the internal soil-vegetation model activated in stand-alone versions of COSMO-CLM.”

Answer: We changed the sentences. They are now in § 4.5, line 953.

p.28, l.949-955: This paragraph is not clear. Going from non-alternating to alternating reduces the time to solution by 35.1 %. Improving the performances of the derivative calculation reduces the time to solution by 9.2%. Going from 16 cores in SMT mode to 32

cores in ST mode results in a reduction of time to solution by 25.5 %. But then why is the “discrepancy” calculated by comparing this 25.5% to the 9.2% linked to the improvement of the derivative calculations? It should be calculated by comparing the 25.5% to the non-alternating to alternating gain of 35.1%, shouldn't it?

Answer: Thank you for asking for clarification of this puzzling result. We explain this complex result now in more detail. See §4.6, line 1017 ff

p.30, l.1040-1043: The 10% variation in the time to solution results should be introduced in the text and not only in the conclusion.

Answer: Thank you for the suggestion. We added this information at the end of §4.1

Minor remarks and technical corrections

I think it would be less confusing to use CCLM everywhere and not sometimes CCLM and sometimes COSMO-CLM

Answer: COSMO-CLM is the official name chosen by the CLM community, with CCLM as the official abbreviation when there is not enough space like in figures. We use now CCLM nearly everywhere. However, to avoid confusion between CLM and CCLM, the full name COSMO-CLM is used more than once.

p.1, l.8: The OASIS3-MCT interface is not really described in the paper. I suggest changing “present” for “use”.

Answer: Thank you for the comment. We realised that we introduced a confusion by using “interface” for model routines where coupling is performed instead for the OASIS3-MCT API (widely known as “PSMILE library”). We modified the text: “We present a unified interface, based on OASIS3-MCT coupling library”

p.3, l.58 & p.6, l.166: Valcke 2013 refers to a paper describing the “old” OASIS3 version and not the more recent OASIS3-MCT version. The reference Valcke et al., 2013 should be used instead.

Answer: We changed the reference.

p.3, l.67: I propose changing “is based” for “would be based”

Answer: We changed the text, l67.

p. 4, l.94: Please add “depends” after “but” in “but on the coupling method

Answer: Done

p.6, l.168: Please add “which” after “data” in “amount of data is a requirement”

Answer: Done

p.6, l.186-187: The sentence “The coupling of COSMO-CLM with the global ocean model NEMO is realized by means of two different regional versions of the NEMO model ...” sounds weird to me because of the opposition between “global” and “re- gional”. I suggest simply “COSMO-CLM is coupled to two different regional versions of the NEMO model ...”

Answer: Done, line 191.

p.10, l.337: The fact that each component needs to be a separate executable is not a constraint anymore with the last OASIS3-MCT_3.0 version; maybe this could be mentioned.

Answer: We added a remark on that feature in OASIS3, See line 340.

p.11, l.339 & l.366: Please change “whose” for “which”

Answer: Done, line 344 and 373

p.11, l.343: I suggest changing “is directly executed via the Message Passing Interface” for “is directly executed via the Model Coupling Toolkit (MCT, Jacob et al 2005) based on the Message Passing Interface (MPI)” and add the reference “Jacob, R., J. Larson, and E. Ong: MxN Communication and Parallel Interpolation in CCSM3 Using the Model Coupling Toolkit. Int. J. High Perf. Comp. App., 19(3), 293-307 2005 “

Answer: We thank the reviewer for this suggestion making this point more clear. We changed the text accordingly, line 349 ff.

P.11, l.357-358: I suggest changing “This component partitioning does not have to be the same” for “The component partitioning and grid do not have to be the same”

Answer: We thank the reviewer for this suggestion and changed the text accordingly in line 364.

p.11, l.361: I suggest adding “and accumulation” after “time averages” I propose “average or accumulation”

Answer: Done, line 368

p.11, l.373: I suggest changing “OASIS3-MCT includes the MPI library” for “OASIS3-MCT includes the MCT library based on MPI ” (but this is redundant with p.11, l.343 -see also my remark #9 above

Answer: We followed the reviewers suggestion and keep the redundancy for better readability, line 380.

p.13, l.428: Please add a) after 4.1

Answer: Done, line 435.

p.13, l. 442: Please change “interpolation” for “coupling” as it is not only the interpolation that is improved but the interpolation and the communication.

Answer: We added “and communication” for clarity, line 449.

p.20, l.687-691: I am not convinced these are effectively the two main goal of performance analysis. These sentences are unnecessary and contribute to the heaviness of the text (see also my first “Important remark” above.

Answer: Thank you for the comment. We removed the discussion of what is not done and changed the text accordingly. See line 718 ff.

p.21, l.722: Please change “compansated” for “compensated”

Answer: Done, line 761.

p.22, l.737-738, I suggest rephrasing the sentence “In a perfectly scaling parallel application the costs would remain constant if the resources are doubled, the parallel efficiency would be 100 %, the speed would be doubled and the speed-up would be 200 %.” for “If the resources of a perfectly scaling parallel application are doubled, the speed would be doubled and therefore the cost would remain constant, the parallel efficiency would be 100 %, and the speed-up would be 200 %.”

Answer: We thank the reviewer for this suggestion and changed the text accordingly, line 777.

p.23, l.791: Please change “CPUh” for “core hours” to be coherent with the rest of the text.

Answer: We thank the reviewer for this suggestion and changed the text accordingly, line 797.

Table 8 should be placed after Table 7 and not after all the Figures.

Answer: We thank the reviewer for this suggestion. Table 8 is now located after figure 6 showing the optimum configurations.

p.26, l906: Please change “atmosphere” for “coupled model”

Answer: We thank the reviewer for this suggestion and changed the text accordingly, line 991 ff.

p.30, l.1031: Please change “scaling;” by “scaling,”

Answer: We thank the reviewer for this suggestion and changed the text accordingly, line 1077.