

Interactive comment on “Coupling of the regional climate model COSMO-CLM using OASIS3-MCT with regional ocean, land surface or global atmosphere model: description and performance” by Stefan Weiher et al.

S. Valcke (Referee)

sophie.valcke@cerfacs.fr

Received and published: 8 June 2016

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.

C1

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.

1. 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.

2. 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, l.922-923). I am not sure on how to correct this but this should be simplified maybe simply by removing results for CCLM-CSM?

3. 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 captions 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)

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

C2

4. 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 than 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

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

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

7. 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.

My third major concern is about section 4.5.2

8. 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!

Then I have the following major remarks:

9. 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.

10. 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.

C3

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

12. 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?

13. 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 : $T_{40} = T_{45} \times 40/45$

14. p.24, 806-807: I do not understand how one can conclude that “COSMO-CLM in 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)

15. 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.

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.

17. 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

C4

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.

18. 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.

19. 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. “

20. 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?

21. 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.

Minor remarks and technical corrections

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

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

24. 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.,

C5

2013 should be used instead.

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

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

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

28. 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 “regional”. I suggest simply “COSMO-CLM is coupled to two different regional versions of the NEMO model ...”

29. 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.

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

31. 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 “

32. 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”

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

34. 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

C6

35. p.13, l.428: Please add a) after 4.1
36. 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.
37. 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.
38. p.21, l.722: Please change “compansated” for “compensated”
39. 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 %.”
40. p.23, l.791: Please change “CPUh” for “core hours” to be coherent with the rest of the text.
41. Table 8 should be placed after Table 7 and not after all the Figures.
42. p.26, l906: Please change “atmosphere” for “coupled model”
43. p.30, l.1031: Please change “scaling;” by “scaling,”

Interactive comment on Geosci. Model Dev. Discuss., doi:10.5194/gmd-2016-47, 2016.