

Interactive comment on “The atmosphere-ocean general circulation model EMAC-MPIOM” by A. Pozzer et al.

S. Valcke (Referee)

valcke@cerfacs.fr

Received and published: 28 April 2011

General comment

This paper presents an evolution of the EMAC Atmospheric Chemistry model now coupled to the ocean general circulation model MPIOM. The ECHAM atmospheric component model has already been coupled with the MPIOM ocean model using the OASIS3 coupler resulting the COSMOS coupled system but a new approach, based on the Modular Earth Submodel System (MESSY) is presented here. In the former approach, the atmospheric and oceanic component models must be run as concurrent executables while in the MESSY integrated approach, the ocean model is run as a subroutine of the atmospheric model. A comparative analysis of the performances of the two approaches in terms of computing time is presented in the first part of the paper. The

C170

second part of the paper presents a scientific evaluation of the EMAC-MPIOM coupled system, which shows that its results, when compared to observations and other ocean-atmosphere coupled models, are reasonable for a Coupled General Circulation Model at that resolution.

This paper, comparing an integrated coupling approach with a more traditional coupling approach based on an external coupler, addresses an issue of great importance in the climate modelling community. Therefore it could represent a significant contribution to the current climate science. However, while I agree that the MESSY integrated approach should in principle be more efficient than the OASIS3 approach, the test case presented here and its analysis are not really convincing. Furthermore, I think a better presentation of the advantages and drawbacks of the two approaches should be done. I strongly recommend that the paper be reviewed to address the following comments so to be published and to present a fair view of this important coupling issue.

Specific comments

1. p. 460, L 7-12: The external method is the most widely used method in Europe but it is not the case in the US where most groups follow an integrated approach (e.g. NCAR CESM1-CPL7, GFDL FMS, ESMF). The OASIS coupler is not used in the Unified Model described in Dando 2004; OASIS3 is used in the most recent version of the Hadley Centre ESM, HADGEM3-AO.

2. p.460, L 23 to p.461, L 1-9: The modifications required to fit the MPIOM model in the MESSY framework as a subroutine should be described more explicitly. What do you mean by “use the same high level API to the MPI”? Do you mean they naturally did so or that they had to do so? Also, I suppose that the original code had to be split in an initialisation, run and finalization part; was this easy to do? Were there any conflicts (in term of I/O units, name space, MPI communicators) that needed to be solved? Was the memory required by the integrated system still reasonable? If none of these problems occurred, it probably means that MPIOM follows good coding rules and has a clean

code structure and this should be clearly stated too! In fact, I suspect this would not be the case for all state-of-the-art component models in Europe and many people consider this step as the major difficulty of the integrated approach. Therefore I think it should be discussed in more detail in this paper, which advocates for the integrated approach.

3. P.461, L 25: It should be stated that the ocean and atmosphere subdomain do not match geographically (if I understand well figure 2).

4. P. 463, L 18: With a more sophisticated parallel search algorithm, such as the OASIS4 one, even the gathering would not be required. Maybe it could be mentioned here that the gathering is still required because the SCRIP library search is sequential?

5. P.463, L 27: it is not clear what "point-to-point communication" means. I would propose something like "... OASIS4 coupler supports fully parallel interpolation which means the interpolation is done in parallel for each intersection of source and target subdomains". The reference should be changed for: R. Redler and S. Valcke, 2010. OASIS4, A Coupling Software for Next Generation Earth System Modelling. Geosci. Model Dev., 3, 87-104.

6. P. 464, L 21: Fig 3; I think "top" should be "left" and "bottom" should be "right". If the right panel in fact shows the field after interpolation, the quality of the interpolation should be discussed: why do we see the MPIOM grid lines on the plot (alternation of dark and light blue stripes converging to the grid poles over Greenland and Asia)?

7. P.465, L 5-7: I think it is a bit vague to state "differences in the achieved efficiency can be attributed to the different coupling methods." The different load balancing that can be reached in the sequential versus concurrent execution of the components also has a strong impact on the achieved efficiency.

8. P. 465 and 466, section 4 in general: The comparison between EMAC-MPIOM and COSMOS performances is not convincing. First, I do not understand how the bias of 58 sec is evaluated and why it is entirely attributed to the SCRIP weight calculation. There-

C172

fore, I do not understand how one can simply subtract 58 sec to the EMAC-MPIOM results and then compare with COSMOS results. For a more convincing comparison, the EMAC-MPIOM system should be modified so to read pre-calculated values from files (as it is done in the OASIS coupler); this should be quite straightforward to implement. Second, as stated in the paper, the performances of COSMOS are not unambiguous as they depend not only on the coupling exchanges but also on the load balance between its ocean and atmosphere components. Third, the OASIS3 coupler can be used with more than one dedicated task for the grid transformations and this functionality can be used to parallelise the coupling on a field-per-field basis and therefore to reduce the bottleneck caused by the sequential interpolation. This functionality, although implemented and tested in COSMOS for the CMIP5 simulations, has not been used here. So for all these reasons, I think the conclusion of a 10% improvement in run-time performance for EMAC-MPIOM attributed to the coupling method cannot really be justified at this point. The issues mentioned here should be addressed to allow a precise comparison and conclusion. (In all cases, I doubt that the comparison at that low resolution would lead to any firm conclusion; the benefits of the integrated approach on the performance of the coupling will be obvious only at much higher resolution, I think.)

9. P. 468, L 15: It is weird to state that the SST are compared to the AMIP ones. A direct mention of the SST observation data sets used in AMIP would be better.

10. P. 469, L 26: Fig 8a should include the East Pacific ocean up to the South America coast to show the variability near the coast.

11. P. 470, L 4: On Fig 8b, why are the results from the PI simulation compared to the not detrended HadISST data? Either the data should be detrended or it would seem more logical to compare it with the results of the TRANS simulation.

12. P 470, L 9-10: Can you detail on what is based the assertion "in the Indian Ocean the model reproduces the observed patterns reasonably well ..."? I think "reasonably well" is a bit strong here!

C173

13. P 470, L 22: “seems to predict a too high ice coverage” is quite a weak assertion; the predicted ice coverage is much too important as is obvious on Fig. 9.

14. P 471, L 20: On Fig 11, why are only the first 50 years of the simulation TRANS shown?

15. P 472, L 5: I think it is not right to state “convection occurs mainly in the Weddel Sea and Ross Sea, although small convective events occur around the Antarctic coast”. On Fig 11, it is quite evident that the major convection site is outside the Weddel sea off the Antarctic coast between 0 and 45 deg East.

16. P. 474, L 13: this sentence should be reviewed after answering concerns raised above in 8.

17. P 474, L 15 to 18: I agree that this could be the conclusion of a precise comparison of the two approaches (see comment #8 above). However, the drawbacks of the integrated approach should also be mentioned (see my comment 2. Above). Another drawback of the integrated approach that should be mentioned is that it will scale only if the two components have the same scaling behaviour, as they are run sequentially one after the other on the same number of processes.

Technical corrections

1. p.458, L 2: “System” not “Sytem”

2. p.462, L 2: The reference for submodel CHANNEL should be moved to p. 461, the first time CHANNEL is mentioned

3. p.463, L 17: “case” not “chase”

4. p.468, L 21: Add “Although” before “The deviation . . .”

5. p.469, L 18: “CMIP3” not “CMPI3”

6. p. 469, L 26: “see Fig. 8a” would be more precise than “see Fig. 8”

C174

7. p.470, L 4: “Fig. 8b” would be more precise than “Fig. 8”

8. p.474, L 7: Units of climate sensitivity should be in K/Wm-2 not KWm-2

Interactive comment on Geosci. Model Dev. Discuss., 4, 457, 2011.

C175