Geoscientific
Model Development
Discussions

GMDD

4, C346–C352, 2011

Interactive
Comment

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

A. Pozzer et al.

pozzer@cyi.ac.cy

Received and published: 9 June 2011

We thank referee #2 for her/his comments.

p. 458, l. 9 As mentioned in the paper (p. 465, l. 13) "The two climate models were run with no output for one month at T31L19 resolution for the atmosphere and at GR30L40 resolution for the ocean. The radiation in the atmosphere was calculated every 2 simulation hours." In both codes (COSMOS and EMAC/MPIOM) the same convective and large scale cloud parameterisations have been used for the atmosphere and the same algorithms for advection and diffusion in the ocean, respectively. Yet, we will extend the description of the model set-ups on p.465, because we think that the introduction should not contain too many details.

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper

p. 458, l. 11 Following the replies to S. Valcke, we will reformulate this sentence, mentioning that the two models show a comparable run-time performance.

p. 461, l 10-12 We will expand the section. The CHANNEL submodel allows an easy management of the memory, internal data exchange and output. In our case additional definitions of 3D and 2D objects became necessary, which are consistent with the dimensioning of the original MPIOM arrays.

p. 461, l. 27-28 As explained on p. 461, l. 27: "An example is shown in Fig. 2, where possible parallel domain decompositions of EMAC and MPIOM are presented". Thus, Fig. 2 shows just a **possible** decomposition, which, as explained in the code, depends on NPX and NPY (see also p. 461, l. 20). Hence, the figure is just an example to show the spatial coverage and the possible decomposition configuration, being by no means the only one which is possible.

p. 462, l. 13-14 We will refer to the table in this paragraph. The amount of exchanged data depends on the model set-up and the chosen resolutions of the components. In the case of a pure physically coupled setup without any chemical coupling, the exchanged fields are listed in Tab. 1. The fields are exchanged at regular time intervals, defined via namelist (as described in the electronic supplement), which can almost arbitrarily be chosen by the user (e.g., 2 hours, 12 hours, 24 hours, etc.). We will add more information in the revised manuscript.

p. 464, l. 6-13 As indeed written in the electronic supplement, we used a *bilinear* remapping method for interpolating scalar fields, and a *conservative* remapping method for interpolating flux fields. The conservative remapping has been specifically designed for flux exchange, being conservative on the global but also on the local scale. While the conservative remapping requires more time during the initialisation phase for calculating the relative weights, during the integration phase all the transformations require the same amount of time, because the same al-

Interactive
Comment

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper

gorithm for all methods is used, only with different weights. We will add this information to Tab. 1, to make the manuscript more complete.

p. 465, l. 5-7 We agree that the comparison can be rather difficult. However, the same machine, compiler and compiler options, libraries etc. have been used in the compilation of both codes, in order to avoid the problem mentioned by the referee. Although a direct comparison of the time consumed for the coupling procedure would be technically interesting, it would require an appropriate instrumentation of both codes, which is beyond the scope of our study. Here, our intention was to assess the overall performance to ensure, that our coupling approach, in comparison to an established approach, does neither deteriorate the overall run-time performance (in which the user is ultimately interested in), nor the results. Moreover, as mentioned in the reply to S. Valcke, this overall performance is not simply a result of the "time spent in the information exchange", but rather a complicated combination of scalability, load-balance and potential limitations in the parallel decomposition.

For the same reason, i.e., because we were only interested in the overall performance, we are very hesitating to perform numerous additional simulations without and with only one-directional coupling. Given that in both approaches the component models are almost identical, we do not expect different conclusions.

p. 465, l. 8-13 We completely agree on this point, and we will add this information (machine type, compiler version and options etc.) to the revised manuscript. As mentioned in our reply to S. Valcke, we will underline the fact that the main goal of the comparison is to show that the new coupling method did not degrade the overall run-time performance. It is certainly not possible to extrapolate from our example general statements valid for all cases.

p. 465, l. 27 This is indeed a difficult question which we cannot answer at the moment. In fact, MPIOM has been implemented in the MESSy system as submodel,

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper

hence it does currently not run separately from the entire EMAC system. To run MPIOM/MESSy "stand-alone", i.e., MPIOM as MESSy basemodel, is an ongoing project. Thus it is currently not possible to compare the original MPIOM model with the "modified MPIOM" to assess the exact performance difference. We speculate, however that the removal of the output routines and the on-line diagnostic slightly improves the MPIOM performance, although we expect this gain to be marginal.

p. 466, l. 2 This information is located here `http://www.messy-interface.org`, under the link "ECHAM5/MESSy Performance".

p. 466, l. 3-12 We see the misunderstanding. Yes, we subtracted the 58 seconds, knowing that this is a constant difference between EMAC-MPIOM and COSMOS (see also our reply to S. Valcke). The calculation of the interpolation weights in the initialisation phase was indeed meant with "most expensive procedure". Once these weights have been calculated, the interpolation during the time integration phase of the model is extremely fast (as explained at p. 463,l. 9-25). We will clarify this in the revised manuscript.

A similar approach as in COSMOS, namely to read in the weights from pre-calculated tables, would indeed be possible, but we did not implement it for 3 reasons:

- The "on-line" calculation of the weights renders the model system more flexible and easier to be handled by the user.
- The additional time spent for the weights calculation during the initialisation phase of the model is negligible compared to the overall time spent for a complete simulation.
- We expect that the additional time spent for the weights calculation during the initialisation phase of the model will be small compared to the initiali-

sation time required for other processes, once chemical processes will be included.

p. 466, l. 13-15 We answered to a similar question in the reply to S. Valcke (see answer no. 17). It is rather complicated to reliably estimate, which coupling approach will give the best performance/scalability, because these are influenced by a combination of core availability, load balance, model resolution (and complexity) and single component scalability. We will add to the manuscript:

1. The scaling and load imbalance issues cannot be regarded separately. The external approach has advantages for the scalability, the internal approach for the load balance.
2. The scalability issue is an intrinsic problem of the component models rather than a coupling method issue, and this must be solved inside the single components.
3. A general statement about the performance and scaling features of the internal versus the external coupling method is simply impossible, because it depends on the coupled components and the computer/network architecture.

p. 468, l. 23 The largest bias is a 4-6 K difference. We will add this to the revised manuscript.

p. 469, l. 23 - p. 470, l. 10 We refer to Fig. 8a at p. 460, l. 23 and l. 26, while we refer to Fig. 8b at p. 470, l. 4. It will be clarified in the revision.

p. 470, l. 17 September and March have been chosen as representative for the ice/ice free season. The maximum (minimum) ice coverage in the northern (southern) hemisphere is reached in March, while the minimum (maximum) ice coverage in the northern (southern) hemisphere is reached in September. These figures are standard and used in other publications as well (Jungclaus et al., 2007).

Full Screen / Esc

Printer-friendly Version

Interactive Discussion

Discussion Paper

p. 471, l. 4-15 It is indeed rather difficult to find the real cause of the ice overestimation. The influence of the resolution is at this stage indeed speculative, additional high resolution simulations are beyond the scope of this study, thus we will remove this argument from the revised text.

p. 472, l. 17 We used PI to have a non-transient, statistically robust estimate of the jet stream. For the revision, we will substitute it with data from the TRANS simulation between 1968 and 1996, as in the NCEP/NCAR reanalysis data and change the text accordingly.

p. 473, l. 25 - p. 474, l. 4 In this comparison, model results from different setups are used, and this indeed requires clarification. The COSMOS results are from a simulation with T61L31/GR15L40 with homogeneously distributed GHGs in the atmosphere, whereas the EMAC/MPIOM simulation was in T31L19/GR30L40 resolution with GHGs nudged towards realistic distributions from a previous EMAC simulation. We hence do not claim that the 0.55 K difference is due to the different coupling methods, but rather to the different simulation setups.

p. 475, l. 11-16 We will remove the outlook from this section.

p. 481, Fig. 1: We will modify the figures as suggested. We will use different colours for different executables.

p. 483, Fig. 3 The errors of the SCRIP transformations have been quantified in the paper of Jones (1999, Tab. 2), who also present a detailed quantification of the transformation errors for a variety of grid combinations. We will add this information to the revised manuscript.

p. 481, Fig. 4 We will correct the label to "task number".

p. 493, Fig. 12 We will add the Y-axis label (hPa).

We will include all the suggested technical corrections, and we thank the referee for the detailed and constructive comments.

References

P.W. Jones. First- and second-order conservative remapping schemes for grids in spherical coordinates. *Monthly Weather Review*, 127:2204–2210, 1999.

J. H. Jungclaus, N.Keenlyside, M. Botzet, H. Haak, J.-J. Luo, M. Latif, J. Marotzke, U. Mikolajewicz, and E. Roeckner. Ocean circulation and tropical variability in the coupled model echam5/mpi-om. *J. Clim.*, 19:3952–3972, 2007.

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