Geosci. Model Dev. Discuss., 2, C441–C445, 2009
www.geosci-model-dev-discuss.net/2/C441/2009/

**Geoscientific
Model Development
Discussions**

Interactive comment on "OASIS4 – a coupling software for next generation earth system modelling" *by* R. Redler et al.

R. Redler et al.

rene.redler@zmaw.de

Received and published: 17 November 2009

We thank referee 2 for the detailed review and suggestions to improve the text. We seriously considered the remarks made and changed the text accordingly.

General comments

1) data transfer

Both the neighborhood search and the data transfers are parallel. The parallel neighborhood search (over the intersections of source and target process local domain) is described in detail in section 5.1 and the the M x N data transfer via the Transformer processes is detailed in the first three paragraphs of section 6.

2) decompositions

This is now addressed in section 5.

3) memory use

In the whole search and regridding process, we never require a global grid assembled on the single process. We have added a sentence at the end of the first paragraph in sec 5.1. and comment on the memory use.

4) processor layout

Currently, we provide support for a fully concurrent component processor layout. Work is in progress to handle a sequential layout. We now mention this in the final section 9 (Conclusions and future perspectives).

5) local-only

Referee 1 made a similar comment. We do not describe the option of a local search anymore. We still describe the problem that arises when working on partitioned grid. To make users, not yet familiar with parallelism, aware of the problem we keep the figure to highlight the interpolation error.

6) use at very high resolution

We do not see any specific concerns with respect to use at high resolution. Nevertheless, we prefer not to include any such speculations in the text.

7) grammatical errors

We have corrected those errors that we are aware of and we followed the advice given by the referee under specific comments. For further corrections we have to rely on the technical editor.

Specific comments

pg 800, line 15

Instead of listing the technical problems that can appear when merging two independently developed codes into one (conflicts with I/O, namespace, memory requirements, etc.) we think it is more interesting to contrast the two current approaches followed to achieve component interoperability and explain why we choose the one implemented with OASIS. This is now detailed in the 5th paragraph of the introduction.

pg 806, line 9-11

The sentence has been revised. We now explain our concept for the general design is more detail.

pg 808, line 19

Contrary to what is stated by the referee we have chosen to communicate the target grid points to the source component and do the neighborhood search on the source side to communicate only useful source points. This is explained in detail in section 5.1. We have included a reference to section 5.1.

pg 818, line 25

We now comment on the accuracy of the different interpolation techniques and explain the larger error we observe at the coast line. The error near the coastline is larger in this regridding involving the BT42 grid than in the regridding involving the LMDz grid described above; this is linked to the fact that the BT42 grid has a much wider masked domain compared to the LMDZ domain, and not to the interpolation algorithm.

pg 822, line 5

Information about the decomposition is now added.

In order not to increase the length of the paper too much, we prefer to provide tables rather than graphs. With graphs we would not be able to display the numbers with the same accuracy. We included a detailed comparison in the 1-1-1 case between OASIS3 and OASIS4 in a new subsection in section 8. We show that even for relatively low

resolution, OASIS4 is more efficient than OASIS3 and the difference increases with increasing resolution as can be expected.

pg 824, line 14

We repeated a few test cases again, but on a different system. Measured time differs by a few milliseconds, for both the search and the data exchanges. We now mention this is the text.

pg 826, , lines 3-11

We improved the explanation of the general communication implementation in particular between the the transformer and the component processes.

As now stated in the text (see section 6), the different Transformer processes work in parallel over the regridding of the different source and target process intersections ("intersection regridding lists"). When the two components performing a ping-pong exchange are not parallel, there are only two such intersections, one when considering one component as the source and another one when considering the other component as the source; in this case, there are in fact only two intersection regridding lists and only two Transformer processes can therefore get some work to do. We hope this is now stated more clearly in the text.

pg 826, line 24 (line 27)

We agree with the referee that we cannot rule out other effects besides cache effects.

We now mention cache effects as one possibility. Another reason could be a better scheduling of the processes and the workload. A finer partitioning of the target grid may fit better to the remote partitioning thus reducing the overhead during the initial search.

Technical corrections

We appreciate the detailed list of comments concerning grammatical corrections. We

have modified the text accordingly without listing each modification in detail at this place.

pg 816, line 25-, sentence makes no sense

We think this remark refers to

"resulting in an effective parallelisation of the Transformer over these lists"

So we changed it for

"this ensures that during the run the Transformer processes work in parallel over the regridding of the different source and target process intersections."

We also clarified the Transformer loop structure.

Interactive comment on Geosci. Model Dev. Discuss., 2, 797, 2009.