

gmd-2015-267 comment to referee #1

We would like to thank the reviewer 1 for the careful reading and the detailed review. Please find below our comments on the review.

Overall, this manuscript is well written and very readable. It spends a lot of time discussing the capabilities of OASIS3 and OASIS4 outside of the introduction, and using them to justify decisions made for YAC. A lot of this seems unnecessary, and could easily be removed without negatively impacting the manuscript. Including this makes the manuscript seem almost like it is intended to describe those couplers as much as YAC.

In the "coupling" community, the question arose: why we had to invest into a new coupling software. Therefore, we felt the need to clarify the differences. In addition, OASIS3/4 is our main source of experience regarding coupling. Many of our decisions are based on these couplers, so use them to justify our design and give credit to them. As referee 2 had similar concerns we will shorten the manuscript and concentrate on YAC.

The performance results are overly simplistic. With only a single result provided for readers to gauge the performance of the coupler. While this may be the most expensive part of this coupler, it is also only performed once in a given simulation.

As stated in the paper, we did not really care about the performance of the coupler as long as it was within reasonable bounds. For the paper we put our focus on the design of the coupler itself. With the given performance results we wanted to show that it is working with acceptable performance. To enhance performance results we will do additional measurements with other grid resolutions and other interpolation methods. We will also explicitly measure the time required by the global search and the interpolation computation.

It would be useful if the authors provided performance results for coupling steps including the interpolation from one grid to another and communicating the results.

Our measurements showed that the coupling step itself (including interpolation and exchange) is really fast and hence boring. Furthermore, we do not consider the data exchange as time critical for real coupled applications. Wait time due to load imbalance is typically much more time consuming than the exchange of data. But we may add a respective diagram.

Additionally, it would be useful to show a semi-realistic example where the two models that are coupled together are not colocated on the same nodes / processors.

It seems like the paper was not clear on this part. For example, in the last measurement we used 256 nodes (6.144 process/individual MPI ranks per component (see x-axis label of fig. 4) which adds up to a total of 512 nodes or 12.288 processes. We will further clarify this in the text.

gmd-2015-267 comment to referee #1

It would also be useful to provide performance for the global search for additional interpolation methods, similar to the one-off discussion found in Section 6.

The global search is done only "once for each pair of source and target grids for which any interpolation is required." (see section 3.3) The global search results are available to all interpolation methods, which in turn do additional communication to fulfil their special requirements.

One major issue with the manuscript is that it is not clearly documented where to gain access to the source code or find instructions for building / running the tests that are discussed in the manuscript.

As mentioned in section 8, on the main page of our Doxygen page (<https://doc.redmine.dkrz.de/YAC/html/index.html>) we have a paragraph about code availability:

Code Availability

Tagged versions of the software are available upon request. Please contact Moritz Hanke (hanke at dkrz.de) or Rene Redler (rene.redler at mpimpet.mpg.de) and provide your name, institution and a few lines describing your intention with yac.

In order to guarantee anonymity for the reviewers we provided a link to an archive file to the Topical Editor. This tar file contains the status of our source code at the time when the paper was accepted for publication in GMDD.

We are very sorry that this was not communicated clearly enough. We will add those details about code availability in the revised version of the manuscript.

Page 16, Figure 4 is captioned with...

"The values in Fig. 4 represent the wall clock time required for the call to yac_csearch." (see section 5) This includes the global search and all communication and computation necessary to compute and distributed the weights required for the interpolation.

Corrections

We would like to thank the reviewer once again for the careful reading of our manuscript and the suggestions for corrections. We will of course include all of these in our revised version of the manuscript.