

Interactive comment on “Coupling technologies for Earth System Modelling” by S. Valcke et al.

S. Valcke et al.

valcke@cerfacs.fr

Received and published: 25 October 2012

We would like to thank Referee #2 for his positive review of our paper. You will find hereafter a reply to his requests and comments.

1) Referee #2 first comment concerns the traceability and reproducibility of each approach: "It would be useful to mention issues relating to traceability and reproducibility of each approach. i.e. it is often important to be able to demonstrate the provenance and integrity of results from ESMs. This can be hampered, for example, by an inability to reproduce results due to parallel computation operations, or by the absence of details about regridding weights (e.g. if those weights are calculated and held in memory at run time rather than being available for inspection outside of the model.)"

To answer this comments, we propose to add few lines on result reproducibility for each coupling approach. Some additions are longer than other as reproducibility is an issue

C818

for some, but not an issue for others.

-for ESMF: "For the most part, ESMF methods do not modify user data numerically, and thus have no effect on the bit-for-bit characteristics of the model code. However, in the case of the sparse matrix multiplication used to apply interpolation weights, user data is directly manipulated by ESMF. In order to help users with the implementation of their bit-for-bit requirements, while also considering the associated performance impact, the ESMF sparse matrix implementation provides three levels of bit-for-bit support. The strictest level ensures that the numerical results are bit-for-bit identical, even when executing across different numbers of processors. In the relaxed level, bit-for-bit reproducibility is guaranteed when running across an unchanged number of processors, while the lowest level makes no guarantees about bit-for-bit reproducibility. This lowest level provides the greatest performance potential for those cases where numerical round-off differences are acceptable."

-for CPL7, in the 2nd paragraph on p.1993: "Mapping weights are generated offline to ensure quality and reproducibility. The CESM coupling operations produce roundoff level differences when processor counts or decompositions are varied, but an optional switch enforces bit-for-bit reproducibility when desired at some performance cost."

-for FCM, before the last paragraph on p.1995: "Since coupling involves parallel floating-point (FP) calculations, one key concern is that of reproducibility across different parallel decompositions, since FP parallel arithmetic operations can sometimes be non-associative. Ensuring bitwise reproducibility across the FMS coupler involves a cost which can be as high as 10\% on some systems. Users have the option of requiring reproducibility but this feature can be turned off if the entire run will be done without changing parallel decomposition."

-for OASIS, before the last paragraph on p. 1996: "But even in this case, reproducibility with respect to parallelism is not an issue as all transformations are done, for each particular coupling field, on only one coupler process."

C819

-for MCT, before the last sentence on p.1997: "This operation can optionally be constrained to give identical results on different numbers of processes."

-for BFG, at the end of the first BFG paragraph on p.1999: "Reproducibility is an issue for coupling systems and, as BFG targets existing coupling technologies, BFG relies on the support for reproducibility that they provide. The wrapper code produced by BFG does not itself cause any reproducibility issues."

2) Referee #2 second comment is about a comparative performance analysis of the various approaches: "Having mentioned the "3%" cost relating to ESMF, it would be useful to be able to compare performance and overheads of the various other approaches, if such figures exist, though I appreciate that it is unlikely that directly comparable figures currently exist for this purpose."

Indeed, even if the overhead of the different approaches were measured for few coupled configurations, these figures are of course extremely sensitive to the particular coupling parameters. And as these overheads were not measured for the same coupling configurations for the different approaches, we think it could be more misleading than informative to add them in the text. Qualitative comparison of the performances of the different approaches is an under-researched problem and there are plans in the IS-ENES2 EU project, which is a follow-on of the IS-ENES project funded over the 2013-2017 period, to develop a benchmark suite to address this issue. We propose to add this statement in the second paragraph of the conclusion section.

3) Referee #2 then has few specific comments.

3.1) "ESMF, page 1990, line 23. Could the term "fully parallel" be expanded or explained? i.e. is the word "fully" necessary? If not, it may lead the reader to surmise that there might also exist a solution which is "partially parallel" as well as one which is "non-parallel"."

The interpolation is composed of two main steps, i.e. the interpolation weight gener-

C820

ation and the application of these weights to the source grid values to calculate the target grid values. "Fully parallel" here meant that both steps are parallel but we agree that this is not clear and may be confusing. To be more specific, we propose to change the sentence "Generation of interpolation weights and their application is fully parallel." by "Generation of interpolation weights and their application via sparse matrix multiplication are both implemented as parallel operations."

3.2) "ESMF, page 1991, line 15. 3% of runtime is not necessarily "negligible". It may be considered "acceptably small" for such a potentially complex task compared with other model component costs, but, given the potential vast amount of computer resources involved in ESMs, extensive optimisation effort is often given to speeding up models by fractions of a percent."

We agree with this remark. As explained above, we now propose not to give any precise number for any approach. The figures we could have for the different approaches were measured for different coupling configurations and could be misleading. We therefore propose to remove the 3% mentioned for ESMF and to refer, in the paper conclusions, to the development of a comparative benchmark suite planned in IS-ENES2, as explained above in 2).

3.3) "OASIS3, page 1996, line 16. Worth pointing out that the official distribution of SCRIP has remained largely undeveloped since 1999 and that it has various known problems which have partly been addressed within a local copy of SCRIP which forms part of the OASIS3 source (and other versions outside of OASIS3.)"

The local copy of SCRIP in OASIS3 does not address known SCRIP problems, in particular near the poles, so we are not sure what the referee refers to. It just offers the options to assign the value of the nearest non-masked source neighbour to target grid points for which the original SCRIP algorithm would not assign any value at all (for example, for the conservative remapping for a non masked target cell that does not intersect any non masked source cells). We propose to complement the sentence mention-

C821

ing SCRIP with: "For regridding, OASIS3 includes the SCRIP library \citep{Jones99}, adding few specific options such that the possibility to assign the value of the nearest non-masked source neighbour to target grid points for which the original SCRIP algorithm would not assign any value at all."

3.4) "MCT, page 1997, line 12. Unsure what is meant by "parallel data movement". Does this mean "parallel coupling field exchanges"?"

Yes it does. We propose to change "parallel data movement" with "parallel coupling field exchanges".

3.5) "MCT, page 1997, line 19. Unsure what is meant by "virtual linearization"."

It just means that multidimensional index spaces are unrolled into a 1D index space. We propose to change "A domain decomposition descriptor (DDD) object uses virtual linearization to represent multidimensional index spaces" with "A domain decomposition descriptor (DDD) object uses a 1D global index space (i.e. a linearization) to represent multidimensional index spaces."

3.6) "BFG, page 1999, line 29 and page 2000, line 2000. References to the now discontinued OASIS4 may lead the reader to conclude that OASIS4 is still of potential interest. It would be better to remove references to OASIS4 and/or replace with "OASIS"."

We agree. We propose to change these sentences with "The work was carried out on a development version of the now withdrawn OASIS4 \citep{Redler10} and efforts are ongoing to extend BFG to support the most recent version of OASIS, OASIS3-MCT \citep{Valcke12b}." where {Valcke12b} is a reference to a recent technical report on the new OASIS3-MCT version.

3.7) BFG, page 2000, line 22-24, the final sentence needs rewriting or removing (the "Finally" and "we" are out of context for this document.)

We agree. We propose to change this sentence with: "In addition, the feasibility of us-
C822

ing BFG2 to couple models that conform to other frameworks by generating appropriate adaptor code will be investigated."

We hope these additions will answer Referee #2 's review.

Interactive comment on Geosci. Model Dev. Discuss., 5, 1987, 2012.