

gmd-2015-267 comment to referee #1

In **red** we repeat the reviewers remarks, in **green** we reply to the remarks, followed by sections and phrases in black (manuscript text in *italic black*) with changes that we applied to the discussion paper to obtain the revised version.

Review #1

We would like to thank the reviewer 1 for the careful reading and the detailed review.

General Comments

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 now provide additional measurements with other interpolation methods (see also our comment to referee 2).

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.

A similar remark was made by referee 2. 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. We now comment on the measurements for the remapping itself without adding a figure.

A new paragraph on the performance of the remapping step was added (see also response to referee 2):

We also did measurements of the remapping step itself. For this we measured the maximum wall clock time required to do a bidirectional data exchange between both toy

gmd-2015-267 comment to referee #1

models. The data exchange includes the exchange of halo data between source processes, the application of the remapping weights, and the transfer of the results to the target processes. All measurements varied in the range of 0.1ms to 10ms. The results have a mostly random pattern and no dependency on the number of nodes per component.

Additionally, it would be useful to show a semi-realistic example where the two models that are coupled together are not collocated 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. Even though in each setup the two components are distributed over the same number of nodes each the processes are not collocated on the same nodes.

We added a sentence at the end of the second paragraph in section "Performance":

Each component is run on its own set of nodes.

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. Navigating the Doxygen site (as recommended in the manuscript), I have been unable to find instructions for downloading / building YAC, so that part was not tested / reviewed, and I cannot comment on the readability of the source code.

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.

A README is shipped with the software which describes the basic steps to build the

gmd-2015-267 comment to referee #1

library and run the tests.

Corrections

Page 10, figure 1: Please expand the caption of this figure. Alone the caption is not enough to determine what each of the subfigures represents, or what the triangles / rectangles are.

The figure caption has been expanded to:

WOA09 January monthly mean sea surface salinity given on white rectangles interpolated to an ICON grid (triangles) using the YAC interpolation stack with 1st order conservative remapping (a) plus 1st-order polynomial fit (b) plus fixed values (c). (for further details see Sec. 2.4)

Page 17, Figure 4 is captioned with...

The values in Fig. 4 (Fig 3 in the revised version) represent the wall clock time required for the call to `yac_csearch`." (see section 4, former section 5) This includes the global search and all communication and computation necessary to compute and distributed the weights required for the interpolation. The figure has been revised (see remarks above), text and captions were adjusted to the changes in the figure.

Other corrections

All corrections suggested by the referee have been applied to the revised version of the manuscript.

Review #2

Again we would like to thank the reviewer 2 for the careful reading and the constructive comments.

General Overview

I am a bit torn ...

Here, R2 provided two alternatives to better set the focus of this paper. Unfortunately, we are not convinced by either of these. The first suggestion would lead us provide insight into the source code via a description of the internal subroutines and interfaces we we think is not doable and would probably go much beyond the scope of this paper. The second suggestion of giving a description about YAC would require major rewriting as R2 already stated. Furthermore it would distract the reader (and us, the authors) from what we would like to communicate. With this paper we wanted to present YAC with a special focus on a set of selected key aspects that we did differently compared to other coupling solutions. The weight generation algorithms we use are more or less the same as in any other coupler and are hence, from our point of view, sufficiently documented. To better set the focus of this paper we modified the title and abstract. Furthermore we added a sentence at the beginning of the last paragraph of the introduction.

New title: *YAC 1.2.0: New aspects for coupling software in Earth system modelling*

Abstract, added: ... *the new aspects* ...

Introduction, last paragraph, added:

With this publication we present YAC with a special focus on a set of selected key aspects that we do differently compared to other coupling solutions.

Specific Comments

Overall, I believe there is a shortfall in the number of cited references.

We thank the referee for pointing us to missing references. Deliberately we do not provide a comprehensive overview on existing coupling software. Instead we already cite an existing review article (Valcke et al., 2012). In addition to that we now cite another review article by Valcke et al. (2012) on "Coupling technologies for Earth System Modelling". We have now added the missing references to coupling software that we mention explicitly, in particular MCT, OASIS3, and OASIS3-MCT. (References to ESMF are already provided.)

gmd-2015-267 comment to referee #2

For the interpolation methods we already provide the reference to the SCRIP library. Appropriate references for the patch recovery are also given in the discussion paper. There is no point to provide references for a nearest-neighbour or average interpolation as this is standard textbook knowledge.

We added references:

Antonov, J., Seidov, D., Boyer, T., Locarnini, R. A., Mishonov, A., Garcia, H., Baranova, O., Zweng, M., and Johnson, D.: World Ocean Atlas 2009, Volume 2: Salinity. S. Levitus, Ed. NOAA Atlas NESDIS 69, U.S. Government Printing Office, Washington, D.C., 184 pages, 2010.

Jacob, R., Larson, J., and Ong, E.: MxN Communication and Parallel Interpolation in CCSM3 Using the Model Coupling Toolkit, Int. J. High Perf. Comp. App., 19, 293–307, 2005.

MPI Forum: MPI: A Message-Passing Interface Standard Version 3.1, Tech. rep., Knoxville, TN, USA, 2015.

Valcke, S., Balaji, V., Craig, A., DeLuca, C., Dunlap, R., Ford, R., Jacob, R., Larson, J., O’Kuinghttons, R., Riley, G., and Vertenstein, M.: Coupling technologies for Earth System Modelling, Geoscientific Model Development, 5, 1589–1596, doi:10.5194/gmd-5-1589-2012, <http://www.geosci-model-dev.net/5/1589/2012/>, 2012

Valcke, S.: The OASIS3 coupler: a European climate modelling community software, Geoscientific Model Development, 6, 373–388, doi:10.5194/gmd-6-373-2013, <http://www.geosci-model-dev.net/6/373/2013/>, 2013.

Valcke, S., Craig, T., and Coquart, L.: OASIS3-MCT User Guide, OASIS3-MCT 3.0, Tech. Rep. 1875, CERFACS/CNRS SUC URA, 2015.

Are the weights generated on the source or target side, on the union of all tasks or is it up to the user? Is remapping done on the source or target side or as part of communicating data between models? Is the communication of data between models separate from remapping or part of remapping?

The search and the calculation of weights is performed on all source processes (owners of the source grid). Here the data are interpolated onto the target before sending. In the discussion paper we already address this question on P11, L18. We have now introduced a separate section to explain this (see also next remark). We did not consider our implementation of the remapping step particularly interesting and therefore a detailed description has been omitted.

First, the source processes communicate with each other for a kind of halo exchange in order to provide the necessary data for the processes that have to calculate the

gmd-2015-267 comment to referee #2

stencil. We follow the referees request and will modify the revised version of the manuscript accordingly by adding an additional paragraph.

The following paragraph was added:

Once the interpolation is finished, each target point that can be interpolated has been assigned to a single source process. The source processes have the weights required to do the remapping of all target points assigned to them. The source processes might require data from other source processes to do the remapping of their target points. Therefore, in the actual remapping step, the respective data exchange between the source processes is performed first. Afterwards, the weights are applied to the source data in order to generate the target data, which is then sent to the target processes.

The authors are extremely familiar with OASIS4, but it's unlikely the readers will be. It is fine to compare to OASIS4 and point readers to an OASIS4 reference, but it is also important to make sure the description does not assume the readers are inherently familiar with OASIS4 and that it does not require a study of OASIS4 as a prerequisite.

R1 made a similar remark. We have now removed the whole section about OASIS4 and kept only those references to OASIS4 where we refer to concepts that we inherited. Understanding YAC does not require studying OASIS4, and we are now convinced that this is not implied by our text anymore.

Deleted old section 2

The mentioning of OASIS4 has been removed from Sec. (YAC – Yet Another Coupler) or weakened.

P4 L 24ff, we removed: ~~*Despite its advantages, the OASIS software suffers from several drawbacks that hamper its use in our case. With YAC, our aim is to inherit successful concepts from OASIS4 (but not software in terms of source code) and at the same time find alternatives were we consider OASIS4 to be less flexible.*~~

P6, L22, we removed: ~~*Compared to OASIS4*~~

P7, L5, we replaced:

~~*based on experience gained from OASIS3 and OASIS4*~~
by
based on our personal experience

P15, L8ff: The whole paragraph has been removed.

We slightly rephrased other parts: P8, L11; P8, L20; P13, L17

How does YAC ensure that the weights generated on the fly are of high quality? The

gmd-2015-267 comment to referee #2

text notes problems at poles in various algorithms, issues with different types of grid cells and different edge options, the requirement that one grid have either convex edges or be rectangular. Are those things checked on the fly? Are the weights somehow checked for conservative or gradient properties independently after they are generated? Are the properties at the pole checked? Is this a concern?

A detailed description of the interpolation algorithms themselves is not supposed to be part of the paper, because from our point of view they do not contain noteworthy differences to other implementations, except for the clipping that has its own paragraph.

We added a reference on first order conservative remapping:

Jones, P. W.: First- and Second-Order Conservative Remapping Schemes for Grids in Spherical Coordinates, Monthly Weather Review, 127, 2204–2210, doi:10.1175/1520-0493(1999)127, [http://dx.doi.org/10.1175/1520-0493\(1999\)](http://dx.doi.org/10.1175/1520-0493(1999)), 1999.

Page 1, Line 20: "not critical" is not a proper description. ESMs have been carrying out computations on moderate to high resolution grids for several years and the (pre) computation of interpolation weights in terms of both performance and quality has been a critical issue. This issue has taken up not insignificant resources in several projects including OASIS, ESMF, and SCRIP.

With this paragraph we refer to past generations of ESMs as stated in the beginning of this sentence. "Past generations" we used here as a synonym for coarse resolution models as it has been mainly used in the past CMIP phases. For those coarse resolution models with a low number of horizontal grid points of the order of less than 100.000, the compute time of the neighbourhood search has not been an issue, not in our models nor – to our knowledge – in other coupled models of these generations. We have revised the whole paragraph to clarify this.

We find that past generations of Earth system models with low spatial resolution and relatively simple block-structured grids, the neighbourhood search between any pair of grids, source and target, did not consume a significant amount of CPU time consumption. In the Coupled Model Intercomparison Project Phase 5 (CMIP5) the majority of coupled models were operated at horizontal resolution of 1 degree and more (see Appendix 9.A in Flato, 2013). Furthermore, the connectivity between neighbouring cells could be deduced by the coupler directly. In addition with the very low degree of parallelism, providing efficient algorithms to perform the neighbourhood search did not pose a significant challenge. With the advancement of new numerical models formulated on irregular grids, a trend towards very high resolution of more than 0.1 degree in the numerical grid and the high degree of parallelism with thousands of processes, we now require flexible and efficient algorithms also in the context of coupling.

gmd-2015-267 comment to referee #2

Section 3.2 describes the communication implementation and is far more important than much of the other material before it. I would like the description here to be expanded a bit, especially as related to lines 10-22. You need to describe "the non-blocking buffered send from Oasis4" for those that are not familiar. Please explain how the callback and data pointer work in a bit more detail. This is not entirely clear. Remember that your audience didn't work on either the Oasis4 or YAC implementation. Also, maybe there needs to be some further clarification on what aspect is being described. There is communication associated with weights generation, interpolation, and coupling data. Is this communication approach used for all of these?

We have now explained the non-blocking buffered send

This routine works similar to a MPI_Bsend but does not require to attach and detach message buffers prior to and after the call. Instead, the buffer management is handled internally.

We improved the description of the communication scheme and now make clear that the communication scheme is used throughout the YAC internal workflow.

In the third paragraph of section 2.2 of the revised manuscript we replaced

~~*For the receiving counterpart we provide a modified version of the standard asynchronous receive operation. Instead of the request argument, our implementation requires a callback function and user-defined data pointer. This callback function is called once the respective message is received by the communication layer. The received data and the provided data pointer are passed as arguments to this function.*~~

with

For the receiving counterpart we provide an asynchronous receive operation. In contrast to typical asynchronous MPI receive operations, our implementation has no request argument. Instead the user needs to provide a function pointer to a callback routine. A request for the receive is set up internally by the communication layer. Once the receive request is fulfilled, the communication layer will call the callback function associated to the request. The data that was received is passed as an argument to the callback function.

We added a sentence at the end of section 2.3

It is used for nearly all communication in the initialisation phase of YAC.

Section 3.4. The interpolation stack feature is well thought out and something that other weights generation methods are not able to do easily yet but is needed. Well done.

gmd-2015-267 comment to referee #2

Thanks.

For us this part is more important than the interpolation methods themselves, because it might give others new ideas for future developments.

Section 3.7. Does the current calculation of intersections guarantee conservation? The overlapping areas have to be computed in a way that the partial areas add up to the areas of each grid cell in total. I think page 14, line 3 suggests the areas are handled properly, but maybe a sentence stating this would clarify.

The sum of the partial areas will always add up to the area of the respective grid cell. Everything else is a bug or numerical inaccuracy. We might follow your suggestion and enhance the paragraph in this regard. We have added a sentence on page 14, line 4 to clarify this.

The sum of the partial source cell areas will always add up to the area of the respective target grid cell up to numerical precision.

Section 4. I find the description of the user interface a bit out of place. The article is really focused on the interpolation weights generation. Nothing has been presented about how interpolation is carried out nor how models are coupled. This section, describing how to setup YAC via XML does not seem to fit into the paper.

Without a user interface the coupler library would be hard to use. Therefore we still think it is worthwhile to mention it. We agree with the referee that the section is misplaced in the discussion paper.

Section 4 has been moved into the Appendix.

Section 5. The total time and scaling for weights generation is very good. I think it's also important to show similar timing information for patch (1st, 2nd, and 3rd) and nearest neighbor calculations as well as bilinear if it's available. These are very different algorithms and I think providing scaling information for all three is important for this paper. In our experience, the 1st order conservative weights generation is sometimes the easiest and fastest to carry out. (This is also noted in Section 6, so data to back it up would be very useful)

Additional measurements for patch 3rd and file base interpolation were added. The different interpolation methods show similar scaling behaviour. Therefore, we refrained from adding more measurements, as suggested by the reviewer, because it would not add more information about the scaling of YAC. We currently have some issues with the nearest neighbour interpolation and no time to investigate it that is why it is not included.

Added new diagram with additional measurements.

gmd-2015-267 comment to referee #2

Added description of measured interpolation methods:

Three different interpolation methods were measured: first order conservative interpolation (conserv), patch recovery with a 3rd order polynomial fit (patch), and a file based interpolation using the weights from the first order conservative interpolation (file).

The description of what the measurements include was improved:

The values in Fig. 3 represent the maximum wall clock time required for the call to `yac_csearch`, which is responsible for nearly all of the runtime consumed by YAC in the initialisation phase of the model run. This routine does the complete search, which is comprised of the interpolation-independent global search (see Sec. 2.3) and the interpolation-specific weight computation.

The following sentence was removed, because it was no longer true for the latest measurements:

~~*Our measurements did not show a significant variance. Therefore, each value in Fig. 3 only consists of five individual measurements.*~~

The description of the measurement graphs was simplified to avoid misunderstandings:

For low number of nodes per component (up to 16) the computation required by the interpolation-independent global search and the interpolation-specific weight computation are the dominant factor for the runtime. This work scales well with the number of nodes, resulting in a minimum runtime of around half a second for the `yac_csearch` at 16 nodes or 384 processes per component. For higher node counts the computation time is negligible and the MPI communication contributes the preponderant part to the runtime.

Interpretation of the differences between the three measured interpolation methods was added:

The three measured interpolation methods generate different workloads, with patch recovery being the most time consuming one. With an increasing number of nodes per component the computation becomes less important for the total runtime of `yac_csearch`, the measurements of all three methods converge for high number of nodes as the communication dominates.

Figure 4 title seems incorrect. According to the text, it is the time to generate 1st order conservative mapping weights, not to carry out remapping.

The figure caption has been adjusted:

gmd-2015-267 comment to referee #2

The new text reads as follows:

Time required for global search and calculation of weights for different interpolation methods in both directions between an ICONR2B06 grid and a 405x405-cubed-sphere grid

Section 6. Line 25. The only thing that has been shown in the text is the cost of the weights generation for 1st order conservative. It is a stretch to say "YAC scales reasonably well". Much of the YAC performance is not shown including remapping performance and cost to communicate data back and forth between models, the ability to support concurrency and other issues.

Except for the declaration of the local data on each process, which contains no communication, the measurements show the whole initialisation cost. This has now been clarified in the text and the title of the diagram (see changes listed above). In addition, measurements on the actual remapping step have been added. From our perspective this should cover the performance characteristics of YAC. Figure 4 shows measurements with up to 12.288 MPI processes. In our opinion this shows the ability of YAC to support concurrency.

A new paragraph on the performance of the remapping step was added:
We also did measurements of the remapping step itself. For this we measured the maximum wall clock time required to do a bidirectional data exchange between both toy models. The data exchange includes the exchange of halo data between source processes, the application of the remapping weights, and the transfer of the results to the target processes. All measurements varied in the range of 0.1ms to 10ms. The results have a mostly random pattern and no dependency on the number of nodes per component.

Section 6. Line 30. It would be nice to document the actual time needed to write and read mapping files as this is already available in YAC to have a quantitative comparison against the cost of generating the mapping weights. It would not surprise me if the read operation was more expensive than online weights generation for the test case used, and the actual numbers would add to the paper.

Because weight files need to be written only once, we did not include any measurements on writing of them. The reading of weights files is part of the file based interpolation method and hence is now shown in the additional measurements in Figure 4.

Because the following paragraph was partially wrong and is covered by the additional measurements, it has been removed:

~~*As an alternative the user can precompute the interpolation weights in a preprocessing step and keep them in files such that the weights are read and*~~

gmd-2015-267 comment to referee #2

~~distributed by YAC at run-time. However, the parallel reading of the weight files and a proper distribution of the weights is not trivial and involves communication among the processes as well.~~

Comments that are mainly on the coupler and not the paper:

We do not see any necessity to modify the text of the manuscript for these remarks.

Section 3.3. The requirement on the user to decompose the grid in a way that also includes a halo region and the rank owner of each halo gridpoint seems to be rather inflexible. That datatype and information will almost certainly have to be computed specifically to support the YAC weights generation interface and is unlikely to be a natural part of any model decomposition.

It is available in ICON. Furthermore, each model that uses advection and diffusion operators in a domain-decomposed world must have the knowledge about the halo. This includes ocean and atmosphere models. We agree that purely column based models like land components or chemistry (not chemistry transport) may have a problem here. In our case land and biogeochemistry are part of one component which does know about the composition. We do not see any need to modify the manuscript in this respect.

The halos are probably only needed in a subset of interpolation methods (like bilinear) and maybe YAC should be computing that connectivity, not the model. In addition, in a decomposition like round robin where each process has a random(ish) set of points, the halo description is going to require that "n" halo points be specified for each grid cell, increasing the memory and complexity. YAC would be much more usable if the connectivity were computed within the coupling layer when needed.

The halos are used to identify communication partners in the 1st-order conservative, patch recovery, nearest-neighbour and average interpolation. In addition it is used by the global search. The design is fundamentally based on having the halos. We could identify the owners of the halo points internally, but since ICON already provides us with that information, we did not yet see a need to do that. OASIS4 made an attempt to compute the connectivity within the coupling library. Our personal experience showed us that these attempts failed in the sense that it was not possible to provide a stable and performant algorithm. With every new grid configuration that was introduced to OASIS4 the algorithm needed to be revised. We do not see any need to modify the manuscript in this respect.

We do not mind not supporting round robin like decompositions.

Section 3.5. Do you have a bilinear interpolation option? This option is heavily used in ESMs.

Average with inverse distance weighting basically is linear interpolation. (page 12 line

gmd-2015-267 comment to referee #2

8). The patch-recovery with a linear polynomial fit or a 4-nearest-neighbour interpolation would be another alternative. For triangular grids a bilinear interpolation is not defined.

Section 8. The web info is pretty useful. Just FYI, I think it is missing a "getting started" type of documentation for new users to know what calls are needed and how to organize them in their models.

The source code contains trivial toy models that show how to use YAC. We will take up this suggestion and refer to them on the Doxygen page as well.

Technical corrections

The grammar is pretty good in this paper but could be improved ...

For further improvement we prefer to rely on the professional service for the final version of the manuscript offered by the Copernicus team.

Please make sure to have references for both L'Huilier's Theorem and Girard's Theorem in the paper.

We consider this textbook material like the Pythagorean Theorem for which one usually does not provide citations. For both theorems we are not able to locate the original sources where these were published first by Simon Antoine Jean L'Huilier (1750 – 1840) and Albert Girard (1595 – 1632).

Please add a reference to the YAC Doxygen web site the first time it is mentioned in the paper. I know it is listed in section 8, but a reference would also be good.

The YAC Doxygen pages are first mentioned in Sec. 3.1 – Documentation on page 6, line 13 in the Discussion paper. There we already provide the URL as a footnote (bottom of page 6).

Other technical corrections

We have considered all other technical corrections and revised the text where requested by the reviewer.

YAC 1.2.0: ~~An extendable~~ New aspects for coupling software ~~for in~~ Earth system modelling

Moritz Hanke¹, René Redler², Teresa Holfeld², and Maxim Yastremsky²

¹Deutsches Klimarechenzentrum, Hamburg, Germany

²Max-Planck-Institut für Meteorologie, Hamburg, Germany

Correspondence to: René Redler (rene.redler@mpimet.mpg.de)

Abstract.

A light-weight software framework has been developed as a library to realise the coupling of Earth system model components. The software provides a parallelised 2-dimensional neighbourhood search, interpolation, and communication for the coupling between any two model components. The software offers flexible coupling of physical fields defined on regular and
5 irregular grids on the sphere without a priori assumptions about the particular grid structure or grid element types. All supported grids can be combined with any of the supported interpolations. We describe the new aspects of our approach and provide an overview about some of the algorithms we are using and the implemented functionality. The parallel performance is examined with a set of realistic use cases. The coupling software is now used for the coupling of the model components in the Icosahedral nonhydrostatic (ICON) general circulation model.

10 1 Introduction

With coupling, as we understand it in this paper, we focus on the exchange of physical fields between model components formulated on different numerical grids. Here, we concentrate on model components which can be described as ocean, atmosphere, sea ice and alike, and do not address the coupling of processes within these types of components.

In general, the tasks that have to be covered by any coupling software –named coupler hereafter– are the interpolation
15 between source and target elements and the handling of the data exchange between them. The interpolation is usually based on a matrix-vector multiplication. The required coefficients for this multiplication can be provided by an external program or by the coupler itself. A significant part of the computation of the coefficients is the so-called neighbourhood search that determines a mapping between source and target elements. For a more detailed discussion on these aspects see Redler et al. (2010).

~~With~~ We find that past generations of Earth system models ~~and the numerical grids involved with low spatial resolution~~
20 and relatively simple block-structured grids, the neighbourhood search between any pair of grids, source and target, ~~was not critical in terms~~ did not consume a significant amount of CPU time ~~consumption. The implicitly given~~. In the Coupled Model Intercomparison Project Phase 5 (CMIP5) the majority of coupled models were operated at horizontal resolutions coarser than 0.5 degree (see Appendix 9.A in Flato et al., 2013). Furthermore, the connectivity between neighbouring cells, ~~low resolution and thus a small problem size, and the~~ could be deduced by the coupler directly. In addition with the very low degree of

parallelism ~~did not pose any challenge to provide~~, providing efficient algorithms to perform ~~this task~~ the neighbourhood search ~~did not pose a significant challenge~~. With the advancement of new numerical models formulated on irregular grids, a trend towards ~~higher resolution~~ very high resolution of more than 0.1 degree in the numerical grid and the ~~higher high~~ degree of parallelism with thousands of processes, we now require flexible and efficient algorithms also in the context of coupling.

5 A variety of software exists in Earth system modelling for this type of coupling. A very prominent software package taking up this challenge is the Earth System ~~Modelling Framework (ESMF, Hill et al., 2004), mainly written and supported by the~~ National Center of Atmospheric Research (NCAR) Modeling Framework (ESMF, Hill et al., 2004). The software is written in a mixture of C++ and Fortran90. In particular, ESMF allows for a parallel exchange of data. In addition, it offers a rich functionality far beyond the pure coupling as we address in this paper. For example, with ESMF individual physics routines
10 can be encapsulated which allows further splitting of individual model components. ~~According to the ESMF Strategic Plan for 2012-2015 the~~ The software package encompasses about ~~600000 lines of code~~ 800,000 lines of Fortran90 and C++ code¹.

Many coupled climate models use the Ocean Atmosphere Sea-Ice Surface (OASIS) coupler maintained by CERFACS in its version ~~3-3~~ (Valcke, 2013). One major advantage, which we think made OASIS a success story, is the lightweight user interface whose implementation requires only very minor code changes in the user code. OASIS3 supports the most common
15 grid types, among them unstructured grids. For the data transfer, complete fields are collected by individual OASIS3 processes. With a quite recent development OASIS3 uses the Model Coupling Toolkit (~~MCT~~) (MCT, Jacob et al., 2005). This development step now allows for a direct and parallel data exchange between the participating model component processes. In its current version, OASIS3-MCT ~~still requires to perform~~ (Valcke et al., 2015) still requires the neighbourhood search to be performed and calculation of interpolation weights on a single process. Alternatively, external tools like the Earth System ~~Modelling~~
20 Modeling Framework (ESMF) can be used to provide the required interpolation weights a priori.

Further details about those and other important approaches available so far have been assembled in a collection of articles by ~~Valeke et al. (2012b)~~ Valcke et al. (2012c) and in a review article by Valcke et al. (2012a). So why ~~we are are we~~ not going to use an existing software solution and adapt ~~this it~~ to our needs? Our target is to create a software framework which allows us to easily add or replace algorithms in order to test new ideas and concepts not only in theory but also in the real life of numerical
25 modelling. Our primary focus for our software development is on the Icosahedral nonhydrostatic (ICON) general circulation model (Wan et al., 2013) and to a lesser degree on the model components available with the first version of the Max Planck Institute Earth System Model (MPI-ESM1) (Giorgetta et al., 2013). Last but not least we contribute to the development of the Climate Data Operator (CDO²) software.

ICON is designed to run on massively parallel systems. As a consequence we require the coupling software to take full
30 advantage of the hardware and to exchange coupling fields efficiently and in parallel. Furthermore, the handling of large grids on a single CPU may exceed the available memory; parallel algorithms working on local and thus smaller domains seem to be the only way out. The demand to allow the model components in the future to change land-sea masks during the forward integration add the requirement for an efficient and therefore parallel on-line neighbourhood search and the calculation of interpolation

¹https://www.earthsystemcog.org/projects/esmf/sloc_annual

²<https://code.zmaw.de/projects/cdo>

weights rather than to perform these steps off-line in a pre-process [step](#). Last but not least the software has to be sufficiently modular. This is required to be able for example to add new user-defined interpolation schemes to the existing code. These new schemes can be tailored to the specific elements in use or physical properties required. In addition, a modular design would allow an easy replacement or the addition of alternatives to already implemented algorithms for the neighbourhood search, interpolation, communication, or other tasks. As CDO supports polygons beyond just triangles and quadrilaterals provided that the grid definition follows the CF conventions, we require algorithms to work with a variety of polygons as well.

None of the existing software packages offers a combination of parallelised and efficient algorithms in particular for the neighbourhood search and interpolation, an user interface which minimises the interference with the existing model component code, all packed in a concise software package. Several of these features mentioned above have been addressed by OASIS4 (Redler et al., 2010) which is similar in functionality to OASIS3 but with an improved support for the parallelism of climate model components. OASIS4 is able to perform a quite efficient parallel on-line search and to handle the parallel data exchange. However, in its current status, this software does not provide any support for unstructured grids and is restricted to quadrilateral elements. Furthermore, the development and support of OASIS4 has been stopped. Therefore, we started a new approach and created a software framework that fits our specific purpose and allows us to easily use or integrate existing solutions where available. More important, our new framework allows the reuse of successfully tested algorithms and routines within other software projects like the CDO.

~~We start in Sec. ?? with an evaluation of OASIS4 and discuss pros and cons that drive our design decisions for the new software~~ With this publication we present YAC with a special focus on a set of selected key aspects that we do differently compared to other coupling solutions. Our new approach is introduced in Sec. 2. We ~~dedicate Sec. A to describe some key features of the user interface before we~~ investigate the performance of the software in Sec. 3 and further discuss and compare our approach in Sec. 4. We finish this paper with some major conclusions in Sec. 5 and provide an outlook on future work. In addition, we describe some key features of the user interface in Appendix A.

2 YAC - Yet Another Coupler

“A complete rewrite of legacy science software, with a specific focus on clarity, can prove useful. It uncovers bugs, provides a definitive statement of the underlying algorithm, lets scientists easily pose and answer new research questions, makes it easier to develop new visualisations and other interfaces, and makes the code considerably more accessible to interested third parties and the public.” (Barnes and Jones, 2011)

YAC started as a small software project which we used to gain practical experience with a variety of methods, programming paradigms, and software tools which are already well known but have not found ~~its~~ their way into the climate modelling community on a broader scale. Rather than doing this in the context of an abstract software project we chose the coupling for Earth system models as a real use case. We have taken this opportunity to generate a framework which allows us to test new concepts and algorithms. Our goal here is to be able to easily replace algorithms or add them for performance comparison purposes and to allow for an easy extension later towards a richer set of functionality.

As already outlined before, we favour the OASIS-like approach due to its user-friendliness concerning the specification of the Fortran API and later use in a coupled Earth system model. In contrast to other existing software solutions this allows us to meet other boundary conditions posed by the ICON project. ~~Despite its advantages the OASIS software suffers from several drawbacks that hamper its use in our case. With YAC our aim is to inherit successful concepts from OASIS4 (but not software in terms of source code) and at the same time find alternatives were we consider OASIS4 to be less flexible.~~ A modular design allows for other software tools like the CDOs to benefit from subsets of the provided functionality. A subset of internal YAC routines has already found its way into recent CDO versions. With the current version of YAC we provide different algorithms to calculate areas on the sphere which are enclosed by polygons. Tests are available to check the quality of the results for different cell types which allows us to identify the most stable algorithm or replace the ~~selected~~ current one by a new alternative. In the same way, different search algorithms can be implemented and tested against each other with respect to performance and scalability.

2.1 Best practices

Best practices of code development found its way only recently into the literature of climate science (Clune and Rood, 2011; Easterbrook, 2012). Being a small group without formal management we did not set up a formal document about a coding style, but we nevertheless defined a set of rules and principles for our development work in order to streamline and focus our effort.

Programming language

One of the most fundamental choices is the selection of an appropriate programming language. We decided to use C for the following reasons: In our particular case we are convinced that most parts of the software are much easier to program in C rather than Fortran. Specific examples are the allocation and reallocation of memory; in general C is more flexible in the handling of dynamically allocated memory. Non-commercial debugging tools for C are in general more mature compared to those available for Fortran. Debugging of Fortran programs with the GNU debugger (GDB) can be more difficult with some Fortran compilers. For example module subroutines get automatically renamed and the inspection of data structures can sometimes cause problems. In our opinion the same argument holds for Valgrind³ which helps during the development to detect memory leaks in a quite early stage. Writing portable code is far easier compared to Fortran as there is less diversity among C ~~compiler~~ compilers than it is the case with Fortran. Especially the different levels of Fortran standards give the programmer a tough job to write portable code. As already indicated by quoting Barnes and Jones (2011) at the beginning of this section using C makes it much more attractive for a larger class of programmers outside the climate modelling community to join this effort and thus helps to bring in different expertise.

The CDO software in many aspects has requirements quite similar to a coupler with the major difference being that the transfer of data happens between files rather than model components. ~~As in its current implementation~~ Currently, the CDOs use in parts the same ~~algorithms~~ algorithm for the neighbourhood search and calculation of the interpolation weights as OASIS3.

³<http://www.valgrind.org>

Hence, they also suffer more or less from the same lack of performance (when data have to be processed on global high resolution grids) and interpolation errors close to the poles. The CDOs are programmed in C, and thus the CDO software can directly benefit by using parts of YAC.

- 5 Since Earth system model codes such as ICON (and the majority of other climate model code) are written in Fortran we also provide a respective interface.

Test suite

We provide short test programs that serve two purposes. The test programs demonstrate how to use particular subsets of the functions and immediately explain the general usage (interface) and the context in which a particular function is intended to be used. Thus, the test programs itself already serve as a documentation of our software. With the help of these test programs it is easier to debug the code and to detect programming bugs already at early stages during the development process as these short test programs can be run at a high frequency. Furthermore, the tests allow to systematically check for most special cases. Most of the test programs have the character of a unit test which forces us to keep the different parts of the code independent with well defined interfaces. Overall, the tests cover a large portion of the code and quickly highlight unintentional bugs due to unknown dependencies. In addition to these short test programs, we provide a set of examples which focus on the usage of the Fortran and C user API. These simple toy models demonstrate the use of the API, and the code sections can be transferred into real model code.

Documentation

Even though a proper documentation of software is key to any software project it is often neglected because programmers consider it an unattractive and time consuming task. It is very challenging to keep external documentation up to date with the software development unless you have a large development team with sufficient resources to dedicate some of these to documentation. In our case we rely on Doxygen⁴. Having the source code documentation and the source code itself in the same place eases the task of keeping both parts in synchronisation with each other. Our main repository is automatically checked for new commits in regular intervals. If new commits are found the YAC Doxygen web site⁵ is rebuilt, which guarantees the availability of an up-to-date documentation at any time. As the Doxygen configuration file is part of the software, users are able to generate a local copy of the html tree.

Style guide

As indicated above we have not provided a written style guide. Nevertheless we stick to certain coding rules which we established while coding the first routines. The principle is then not to deviate from the once established style anymore thereafter. We use long names for routines, data types and variables which makes the code longer but more readable and easier to understand. Wherever possible we avoid global variables. We restrict access to contents of internal data structures by using information

⁴<http://www.doxygen.org>

⁵<https://doc.redmine.dkrz.de/YAC/html/index.html>

hiding schemes. Access to these data structures is only available through a set of interface routines. We elaborate further on this in Sec. 2.1. ~~Compared to OASIS4 we~~ We have kept the individual functions ~~shorter, which again increases~~ short to increase the readability of the code.

5 Version control

We use the version control system git. The main advantage for us is that git allows for smaller (local) commits during the development of a special feature which are independent of the availability of any network access. Only when a particular development cycle is finished the changes are pushed to the remote directory that is accessible to all, which then also generates a single notification by email. ~~The way git is handling the branches and especially the ease of switching between branches we~~

10 ~~feel very comfortable with.~~

Optimisation

We first concentrate on the development of functionality without a particular focus on optimisation. Highly optimised code often comes at the cost of reduced readability, and the resulting code is often less flexible to changes of the algorithm. For some parts of the workflow, we deviated from this paradigm, because based on ~~experience gained from OASIS3 and OASIS4~~

15 our personal experience we know that high performance is key for certain aspects. Therefore, we implement efficient algorithms for the search (see Sec. 2.3) right from the start. The modular structure of the code facilitates this work as algorithms and even data structures can be replaced without any impact on other functions. We have already minimised the number of messages by packing data. Apart from some special support for completely regular grids in longitude and latitude we do not make use of any implicitly given connectivity as it is the case for block-structured grids. Rather they are treated internally like any other
20 unstructured grid, and we see some potential for optimisation in order to speed up the search on this type of source ~~grids~~grid.

Object-oriented programming

Even though the implementation is done in C, we have tried to use an object-oriented programming paradigm. Most data is stored in objects, which are C-structures that contain a number of related data fields. The data within an object is normally not accessed directly. Instead, it is read or modified using a set of interface routines.

25 The major advantage of this approach is that different parts of the code only depend on the interfaces of each other rather than directly on the actual data. By using the interfaces all access to the data is controlled by the code that is directly associated to the data. This avoids misuse of its content and unwanted side effects that might happen when changing seemingly independent parts of the code. In addition, this allows changing of the data structures while keeping the original interface. Such a change might be required in case an algorithm is modified or replaced.

30 One example of such an object within YAC is the grid. YAC supports multiple grid types. Each grid type has its own method of storing its data, such that memory consumption and access is optimal (e.g. regular grids do not store explicit connectivity information). However, there is one set of interface routines which supports all grid types. All algorithms in YAC access grid

data through this interface. As a result, parts of the code that use this interface work with all grid types. For example, a new grid type would automatically work with all existing interpolation methods.

5 The interpolation method object is another example. In YAC an interpolation consists of a number of arbitrary interpolation methods. An interpolation starts by applying the first method to the data that is to be interpolated. Data points that cannot be handled by an interpolation method are passed to the next one. This is only possible because interpolation methods have a common interface, which is used to pass the data between them without knowledge of the other interpolation method type. The interpolation and a practical example of applying a sequence of interpolation methods is presented in Sec. 2.4

10 2.2 Communication

YAC can be considered as a kind of abstraction which hides the complexity of the Message Passing Interface (MPI) (MPI, MPI Forum, 2015) the communication between model processes. Like MPI, YAC is programmed as a pure library which has to be linked to the model code. YAC enables the model component processes to communicate directly with each other depending on the coupling configuration.

15 Internally, YAC does all communication through its own communication layer. This allows us to do the communication independently of underlying low-level communication routines (e.g. MPI). In addition, this layer is used to enhance the capabilities provided by the MPI library. Currently, there are two implementations of this layer. The standard one is based on MPI. For debugging purposes there is also a second implementation of the communication layer that simulates a parallel environment within a single process. This allows us to simulate a parallel application that uses asynchronous communication without the
20 effects of non-deterministic behaviour.

In the MPI variant we avoid blocking operations in favour of an asynchronous communication scheme. For the sending we inherit ~~the a~~ non-blocking buffered send from OASIS4~~with its internal buffer management~~. ~~This routine works similar to a MPI Bsend but does not require to attach and detach message buffers prior to and after the call. Instead, the buffer management is handled internally~~. For the receiving counterpart we provide ~~a modified version of the standard an~~ asynchronous receive
25 operation. ~~Instead of the request argument~~~~In contrast to typical asynchronous MPI receive operations~~, our implementation ~~requires a callback function and user-defined data pointer~~. ~~This callback function is called once the respective message is received~~~~has no request argument. Instead the user needs to provide a function pointer to a callback routine. A request for the receive is set up internally~~ by the communication layer. ~~The received data and the provided data pointer are passed as arguments to this~~ ~~Once the receive request is fulfilled, the communication layer will call the callback function associated to the request.~~
30 ~~The data that was received is passed as an argument to the callback~~ function.

We use these communication routines to split the complete workload into tasks. Each task has its own callback routine. By calling the asynchronous receive routine and by passing the respective function pointer a task is "activated". Once the respective data has been received by the communication layer, the task is processed. Each task can trigger other tasks by sending the respective messages. This generates dependencies between tasks. All tasks and their dependencies can be seen as a directed acyclic graph (DAG). The main advantage of this communication scheme is that independent tasks can be processed in any order, which should result in a good load-balancing. ~~This concept was already introduced in OASIS4~~~~has a similar~~

functionality. However, the YAC implementation is by far more general, which allows usage of this communication scheme by independent parts of the code without interfering with each other. It is used for nearly all communication in the initialisation phase of YAC.

2.3 Search

During the initialisation phase a so-called global search is performed. This is done once for each pair of source and target grids for which any interpolation is required. The grids can be distributed among multiple processes, and the global search does a basic intersection computation between all cells of both of these distributed grids. The result is a prerequisite for all interpo-
10 lations. It greatly simplifies the interpolation specific search described in Sec. 2.5. Furthermore, it reduces the communication between source and target processes in the interpolation search step.

In order to emphasise that we support a large variety of cells we call these *polygons* from now on. For the global search we require that within a grid all vertices of the polygons and the polygons itself are unambiguously identifiable by ids. For grids that are distributed among multiple processes the parts of the grid on each process need to have a halo with a width of at least
15 one polygon. Each vertex, edge or polygon must be owned by only one process. In case of halos the user needs to provide the rank of the owner process of the respective data points.

Since computations of intersections between polygons can be computationally intensive, we use bounding circles as a basic estimate for possible intersections. On the sphere bounding circles have better properties than bounding boxes. A particular advantage of circles is the computation of the intersection between two of them on the sphere, because it is simple to program
20 and fast to calculate. Furthermore, our experience with OASIS4 has shown that bounding boxes in the longitude - latitude space can have problems at the pole regions of the sphere.

If the computation of the intersection of the bounding circles yield a possible intersection between the polygons, the exact result needs to be computed. A basic operation required to do this is the computation of the intersection between two polygon edges. We currently differentiate between three different edge types: (1) edges that are located on great circles, (2) circles of
25 longitude, and (3) circles of latitude. Intersection computation involving great circles is done in 3D-Cartesian space. In the pole regions of the sphere this is much more accurate than using trigonometric functions in the longitude - latitude space.

To reduce the number of checks that need to be done in order to determine all intersecting polygons of both grids, we have introduced an intermediate step. For this intermediate step we implemented two different algorithms: a bucket- and a tree-based search algorithm.

30 **Bucket-based search:**

Our first approach was a bucket search. It starts by mapping cells of the source and target grid to an intermediate grid that is relatively coarse compared to the source grid. We use a Gaussian grid as the intermediate grid. Due to the properties of this grid, the mapping operations have a low computational cost. Afterwards, it is sufficient to compute the intersection between source and target cells that are mapped to the same cell of the intermediate grid.

The bucket search suffers from the fact that the "cell densities" of the intermediate grid and the source and/or target grid can deviate significantly. This can result in a high number of source and/or target cells being mapped to a single cell of the intermediate grid. As a result more source and target cells might have to be compared to each other than necessary.

Tree-based search:

As an alternative to the bucket search a search tree was implemented. Based on the bounding circles of all cells of the source grid a search tree is generated. Using this search tree, we can then find all source cells whose bounding circles overlap with bounding circles of the target cells. Since the algorithm is based on the bounding circles of the grid cells, all computation is rather cheap. In our tests we have seen that the tree-based search algorithm is more than twice as fast as the bucket algorithm. Therefore, this is now the preferred algorithm in YAC.

2.4 Interpolation stack

In a typical coupling setup where a data field is sent from a source component to a target component, the grids of both components are not identical. This is particularly true for land-sea masks. Therefore, a target may not be completely covered by valid source polygons such that a particular interpolation may not be well suited to provide data for this target. In order to provide a solution to these problems in YAC all interpolation methods are independent of each other but share a common interface. Multiple interpolation methods as listed in Sec. 2.5 can be combined in the so-called interpolation stack. This allows the user to freely select any sequence of interpolation methods individually for each coupling field.

In Fig. 1 we apply an example of such an interpolation stack for interpolating the World Ocean Atlas 2009 (WOA09) sea surface salinity (Antonov et al., 2010), visible as rectangles, onto the ICON R2B4 grid, visible as grey triangles. As the primary interpolation we choose the ~~1st-order~~ 1st-order conservative remapping (Fig. 1a). Here only those target cells get a value which are completely covered by valid (non-masked) target cells. In a second step we try to interpolate remaining cells using the patch recovery method with a 1st-order polynomial (Fig. 1b). In the third step all remaining cells which still have ~~not got any~~ no value are set to a fixed value, 999 in this case and appear as dark grey areas over continents and narrow bays in Fig. 1c. We note that salinities with values of less ~~that than~~ than 34 are colored in light blue.

The interpolation stack can easily be modified via the XML interface (see below) to become e.g. a 1st-order conservative remapping followed by a 2nd-order polynomial patch-recovery, and then by a 4 nearest-neighbour interpolation.

We provide a small set of parameters to configure the individual interpolations. The complete list of parameters is described in the GUI user manual which is provided as a supplement to this publication.

2.5 Interpolation

~~An interpolation~~ The results of the global search, which are independent of the actual interpolation method being used, act as the starting point for the interpolation. On the target processes, it provides for each local polygon a list of all source processes having overlapping source polygons. Based on this information, the interpolation starts by distributing the target points, which need to be interpolated, among the source processes. For this the To process the received target points, the global search provides on the ~~target-source~~ target-source processes for each ~~source-process the list of local polygons that have an overlap with the~~

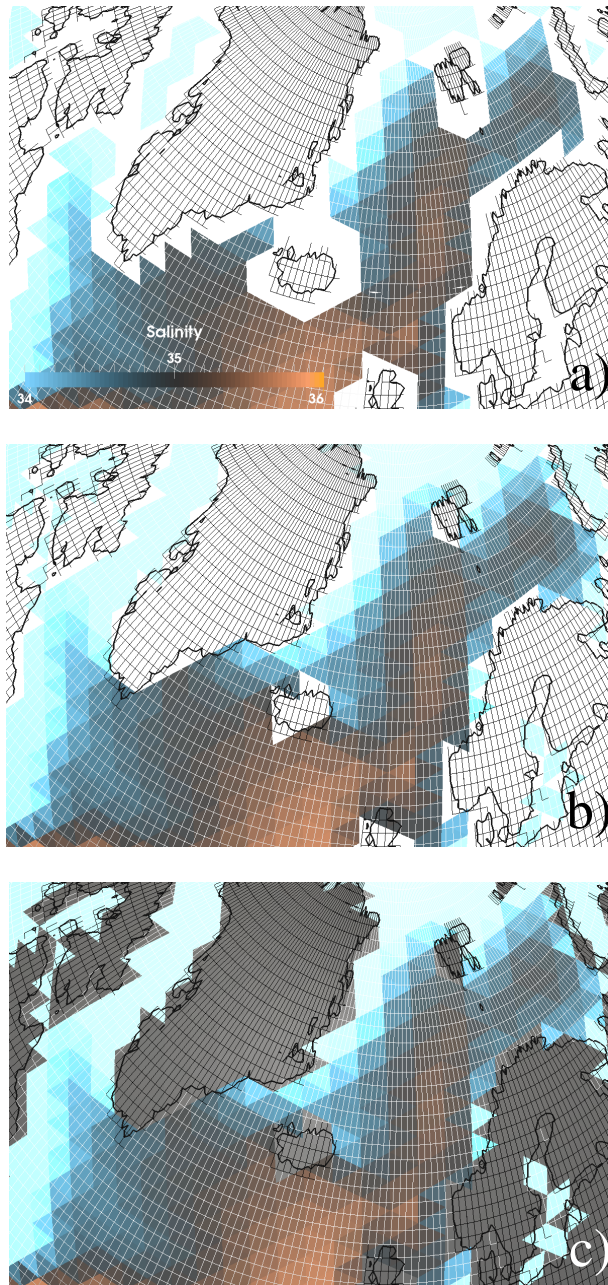


Figure 1. Interpolation of WOA09 January monthly mean sea surface salinity given on white rectangles interpolated to an ICON grid (triangles) using the YAC interpolation stack with 1st-order conservative remapping (a) plus 1st-order polynomial fit (b) plus fixed values (c). (for further details see Sec. 2.4)

polygons of the respective source process. On the source side, for each target polygons that overlaps with the local domain, a list of all source polygons matching the received target polygon a list of all overlapping local polygons.

Once the interpolation is finished, each target polygon is available. This information serves as the starting point for almost all interpolation-specific search steps. A point that can be interpolated has been assigned to a single source process. The source processes have the weights required to do the remapping of all target points assigned to them. The source processes might require data from other source processes to do the remapping of their target points. Therefore, in the actual remapping step, the respective data exchange between the source processes is performed first. Afterwards, the weights are applied to the source data in order to generate the target data, which is then sent to the target processes.

1st-order conservative remapping: Using the information provided by the global search, most of the work required for the 1st-order conservative remapping (Jones, 1999) is done. The remaining task here is to determine for each target point a responsible source process (e.g. the process with the lowest ranks working contributing to the interpolation result of the respective target point) and to collect information for those target polygons that overlap with more than one source domain (see First Interpolate Then Transfer (FITT) pattern, ?) and main remaining task is to calculate the partial overlaps between source and target cells, which are used to generate the interpolation weights. Different algorithms for the calculation of a polygon area on the sphere are implemented (see area.c and test_area.c for a comparison of these algorithms). Like in ESMF as default we are following L'Huilier's theorem.

Patch recovery: Inspired by ESMF we provide the so-called patch recovery interpolation (Zienkiewicz and Zhu, 1992). In our implementation this method is built on the very same set of preliminary information that is available for the conservative remapping. For one particular target polygon we consider all overlapping source polygons for the patch, very similar to the 1st-order conservative remapping. In addition we allow the user to extend this patch by one more row of source polygons surrounding the primary patch. For each source element in the patch we determine sampling points in analogy to Gauss points for a unit triangle and then apply a polynomial fit over these sampling points. The user can choose between 1st, 2nd and 3rd-order polynomials, furthermore the density or number of sampling points can be chosen. The system of linear equations which we have to solve for the polynomial fit can be rewritten in such a way that we are able to calculate a single value (weight) per source polygon in the patch which is only dependent on the geometry. We provide the details about the math on our Doxygen web site. In analogy to the 1st-order conservative remapping these weights are constant in time and can be applied to the physical fields later during the exchange.

Nearest-neighbour: Points that need to be interpolated can be initially assigned to multiple source processes. For the nearest-neighbour search we first start by assigning each target point to only one source process. Based on the global search data, we then determine for each target point the n closest source points in the local data, with n being prescribed by the user through the XML configuration file. Afterwards bounding circles are constructed that have the target points at their center. The diameter of these circles is chosen such that all points found for the respective target points are within the circle. In case a halo polygon is within a bounding circle, a search request is issued to the actual owner of the halo cell containing all relevant information. It can occur that a search request is forwarded multiple times until the final result is found. In case the number of found source points is lower than n, the bounding circle is iteratively enlarged using the connectivity information until enough points have been found. The n geographic inverse distances are calculated one-on the sphere. The inverse distances are normalised and stored as weights. As an alternative, simple weights with values 1/n are stored to provide the arithmetic average.

Average: For the average interpolation we take the source polygon which contains the target point and compute the arithmetic average from all vertices of the polygon. We also provide the option to apply inverse distance weighting to the source vertices
5 to do the interpolation.

Precalculated weights: The file interpolation is able to read in NetCDF files that contain weights. Currently, it only supports a file format that is a simplified version of the weight files generated by the SCRIP library (Jones, 1998). However, it has the potential to be able to handle weight files generated by other programs (e.g. ESMF).

The common use case for applying this interpolation method is to compute the weights once and then reuse them in order to
10 save time in the initialisation of the model run. This approach is only feasible when reading in and distributing the weights is faster than computing them. Measurements show that computation of the weight at the start of the model run is not necessarily a significant performance factor (see Sec. 3). It depends on the number of processes and the number and complexity of the interpolations used.

Due to our concept of the interpolation stack (see Sec. 2.4) another potential use case is conceivable: If there are target
15 points that require a special handling, which is not covered by any interpolation method currently available in YAC, the user can provide a file that contains the weights for only these points. These weights could be tuned by hand. The remaining points could then be interpolated using any standard interpolation or even another weight file.

Fixed value: Last but not least we provide the option to assign a user-defined fixed value to target polygons or -points which is particularly useful when selected as part of the interpolation stack.

20 **2.6 Weight file generation**

YAC is able to write the weights generated by an interpolation stack to a decomposition independent weight file, which is supported by the file interpolation method (see 2.5). To activate this, the user has to specify it in the XML configuration file. Currently, this is supported for any interpolation stack, except for stacks containing the fixed interpolation (see 2.5). Fixed interpolation is not supported because instead of weights it only generates target point lists that have to be assigned a certain
25 value. And this type of "interpolation" is usually not covered by typical weight file formats. However, it could be added if necessary.

2.7 Cells intersection computation

For 1st-order conservative remapping the coupler needs to compute the area of the overlap region between intersecting cells. It basically consists of three subtasks: (1) computation of the intersection point between edges, (2) clipping of two polygons,
30 and (3) computation of the area of the overlap region.

Most couplers support edges that are represented by a section of a great circle on the sphere. If not identical, two great circles intersect twice. Usually, only one of both points is of interest to the coupler. There are three basic methods to compute this intersection point. The simplest method assumes that all edges are straight lines in latitude-longitude space, which makes intersection computation rather simple. For longitude and latitude circle edges (edges of cells of a Gaussian grid) this is accurate. However, for great circle edges this gets more and more inaccurate the closer the edge is to a pole. A second method

uses trigonometric functions to compute the intersection point of great circle edges. Theoretically, this method should be very accurate, but due to numerical inaccuracy we observe problems when computing intersections close to the pole. The

5 SCRIIP library (Jones, 1998) provides the option to apply a coordinate transformation (e.g. Lambert equivalent azimuthal projection) for cells that are close to the pole. This is controlled with threshold latitudes. All cells that are above (for the northern hemisphere) or below (for the southern hemisphere) this latitude are transformed. The SCRIIP User Guide provides an example threshold value of ± 1.5 (in radian). Even though the transformation improves accuracy, it also generates a discontinuity. ~~When using a similar approach~~ A similar approach is used in OASIS4 ~~we observed~~, where this technique causes problems in some

10 very rare cases, because it generated "holes" in the grid; some parts of the sphere seemingly were not covered by the global grid. In YAC we apply vector operations in three-dimensional space to compute all intersections involving great circle edges. This method is much more robust and does not require any special handling of close-to-pole cases. In addition to great circle edges, YAC also explicitly supports latitude and longitude circle edges. Unfortunately, latitude circle edges introduce the possibility that two intersection points exist, which can be both within the bounds of the edges involved. This can occur when a great

15 circle edge intersects with a latitude circle edge. The computation of these points itself is not an issue, but it makes the clipping more complicated.

The clipping of two polygons computes the intersection between them. Typical clipping algorithms assume that all edges of the polygons are straight lines. Due to the type of edges supported by YAC, this is not the case here. We tried to approximate the latitude circle edges using sections of great circles, but this increases computation time and makes the cell concave even though

20 in latitude-longitude space it is convex. Currently, we use a modified version of the Sutherland-Hodgman clipping algorithm (Sutherland and Hodgeman, 1974). This algorithm requires that one cell is either a convex cell that has only great circle edges or a rectangular cell consisting of latitude and longitude circle edges. The second cell can be either convex or concave and can have any combination of edge types.




Figure 2. Clipping of ICON triangle and ECHAM rectangle

To compute the area of a spherical triangle consisting of great circle edges there are two basic formulas: Girard's Theorem and L'Huilier's Theorem. Originally, we used Girard's Theorem. While testing we noticed an increasing error with decreasing cell size. L'Huilier's Theorem is more complex but yields better results.

The output of the clipping is a polygon that is potentially concave and may contain all edge types. In a first step we assume that all edges of the polygon are great circle sections. We split the polygon into triangles and compute their area using the above mentioned theorem. In a second step all latitude circle edges are handled. For each of these edges the error that is made when assuming them to be a great circle edge is computed and either added or subtracted from the overall area depending on the respective case. The sum of the partial source cell areas will always add up to the area of the respective target grid cell up to numerical precision.

Figure 2 depicts a special clipping case near a pole. The triangle consists of great circle edges and is directly on the pole. The rectangle is a typical Gaussian grid cell consisting of latitude and longitude circle edges. The upper part of the figure shows a three-dimensional view directly from above the pole. The lower part depicts the same case in the two-dimensional "Lon-Lat-space". The result of the clipping is a concave polygon with all supported edge types.

To illustrate the difference between great circle and latitude circle edges in the area computation, we have compared the area of rectangular cells that only differ in the type of the edges. The reference cells are comprised of longitude and latitude circles edges (typical Gaussian grid cells). The other cells have great circle edges. For high resolution grids (edge length of around 0.5°) the area difference is up to 1 ‰ for cells close to the pole. The area difference is higher for a low resolution grid (edge length of around 5°) where the error goes up to 71 ‰.

3 Performance

We developed two toy models to do performance measurements: perf_toy_icon and perf_toy_cube. The toy model perf_toy_icon is based on an unstructured grid that consists of triangular cells. The other toy is based on a cubed sphere grid. Both toy models cover the whole sphere. Each toy model has one component, which defines an output field that is coupled to an input field defined on the component of the respective other model. Thus, each component acts as source and target component, and consequently the search is performed within each (source) component.

Figure 3 shows scaling measurements done on mistral at ~~DKRZ~~the German Climate Compute Center (DKRZ), which at the time of testing was equipped with two Intel Xeon E5-2680 v3 12-core processors per node. We used the Intel compiler version 16.0.2 with Intel MPI version 5.1.109.2.150. The measurements ran using 24 MPI processes per node. The toy model perf_toy_icon used an ICONR2B06 grid, which has 245760 cells and around 35km distance between the centers of neighbouring cells. In a typical atmosphere-ocean coupling setup, the ocean model component usually has a higher grid resolution than the atmosphere model component. To mimic this we used a 405x405-cubed-sphere grid (984150 cells; around 25km distance between cell centers) for perf_toy_cube, which represents the ocean in our setup. ~~The coupling was done using a first order conservative interpolation.~~Three different interpolation methods were measured: 1st-order conservative interpolation

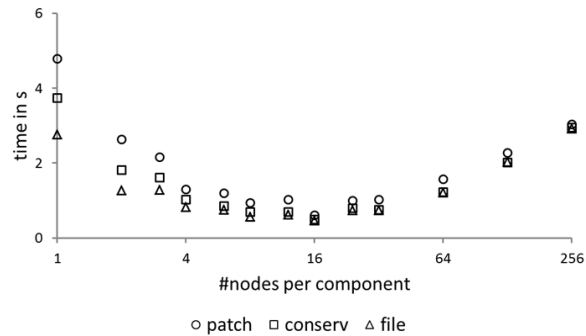


Figure 3. Performance—Time required for global search and calculation of YAC weights for 1st-order conservative remapping different interpolation methods in both directions between an ICONR2B06 grid and a 405x405-cubed-sphere grid

(conserv), patch recovery with a 3rd order polynomial fit (patch), and a file based interpolation using the weights from the 1st-order conservative interpolation (file). Each component is run on its own set of nodes.

5 The values in Fig. 3 represent the maximum wall clock time required for the call to yac_csearch, which is responsible for nearly all of the runtime consumed by YAC in the initialisation phase of the model run. This routine does the complete search, which is comprised of the global-interpolation-independent global search (see Sec. 2.3) and the interpolation-specific search and weight calculation weight computation. Due to the different resolutions, the total workload generated by this routine differs between the two components. By adjusting the number of processes for each component individually it would be possible to minimise the load imbalance between them. To simplify the experiment, we still used the same number of processes for each one. Our measurements did not show a significant variance. Therefore, each value in Fig. 3 only consists of five individual measurements.

15 The scaling measurements can be grouped into two parts, the first group being composed of runs with 1 to 16 For low number of nodes per component and the second group consisting of runs with more than (up to 16 nodes. In the first group, the computation time) the computation required by the interpolation-independent global search and the interpolation-specific weight calculation are the dominant factor computation are the two dominant factors for the runtime. This work scales well with the number of nodes. At, resulting in a minimum runtime of around half a second for the yac_csearch at 16 nodes per component (or 384 processes per component) it has an efficiency of around 44% for our test setup. At that point, For higher node counts the computation time is negligible and the MPI communication contributes the preponderant part to the runtime. In this case, it seems to increase linearly with the

The three measured interpolation methods generate different workloads, with patch recovery being the most time consuming one. With an increasing number of nodes per component the computation becomes less important for the total runtime of yac_csearch, the measurements of all three methods converge for high number of nodes as the communication dominates.

Since, the wall clock time for yac_csearch in our test case is still below four around three seconds for 256 nodes per component, which is a total number of 12288 processes for the whole setup, the relatively bad scaling for high numbers of

nodes does not seem to be a problem. ~~An extrapolation of the measurements leads to around 13 seconds for a potential setup~~
Even with 1024 nodes per component, which ~~should be acceptable~~ could not be tested due to the limitations of the available
5 ~~hardware, the results should still be acceptable.~~

~~We also did measurements of the remapping step itself. For this we measured the maximum wall clock time required to do
a bidirectional data exchange between both toy models. The data exchange includes the exchange of halo data between source
processes, the application of the remapping weights, and the transfer of the results to the target processes. All measurements
varied in the range of 0.1ms to 10ms. The results have a mostly random pattern and no dependency on the number of nodes
10 per component.~~

4 Discussion and Outlook

In the context of describing YAC we provide some insight into the search algorithm we are using, the supported interpolations,
the parallelism, and elaborate on some design aspects and best practices. With a first development step we have focused on
solving the existing problems and on providing the basic functionality. At the same time it is important to have future problems
15 in mind and design the software such that a later integration of new programming modules is not hampered or blocked. We
already went through this exercise when we added the first-order conservative remapping or the patch-recovery interpolation
(Sec. 2.5). In a second step we have put our focus on the optimisation, for example by providing an alternative to the bucket
search (Sec. 2.3) or by introducing the asynchronous collective communication routines.

YAC comprises of roughly 38,000 lines of source code for the library plus another 32,000 lines of code for unit tests. On
20 our machine, the YAC source code compiles in less than a minute and the unit tests take a few seconds to run. These tests have
proven to be very useful during the development and also in the exploration phase. It remains a challenge to achieve full test
coverage and to test all possible use cases and all internal functionality. Thinking about reasonable tests and the implementation
of those test takes some time. At first this tends to slow down the generation of source code for the library part. But by strictly
following the rule of writing tests along with the development of library code we are able to detect programming errors at a
25 very early stage; new development is much less likely to break existing code. We are thus convinced that the unit tests help to
significantly speed up the development process.

Overall, we consider YAC as being rather efficient when speaking about the number of lines of code, the compile time and
the development cycle. This efficiency allows us to test methods and to experiment within a concise software framework. We
are able to try out and compare algorithms and methods. Since most parts are programmed in C we directly contribute with our
30 methods to the CDO development. As the CDO software is used worldwide in climate research, the parts of YAC, which are
already ~~transferred into the CDO. These routines are thus used in a much wider context compared to what we can do ourself
with YAC alone~~ used in the CDO, will probably be used more extensively than YAC itself.

In Sec. 3 we provide some performance measurements for a typical model or grid configuration. The initialisation phase
of YAC scales reasonably well with the number of processes provided that the local problem size remains sufficiently large.
With an increasing number of processes and a thus reduced local problem size the communication starts to dominate the time

required to perform the neighbourhood search in the initialisation phase but remains within a few seconds. ~~As an alternative the user can precompute the interpolation weights in a preprocessing step and keep them in files such that the weights are read and distributed by YAC at run-time. However, the parallel reading of the weight files and a proper distribution of the weights is not trivial and involves communication among the processes as well. Nevertheless~~The performance of the exchange step itself was in the tested setup good for all node counts. The measurements suggest that the time for the exchange step is mainly influenced by external factors like network utilisation of other applications running at the same time. Even though the most time consuming parts of YAC are not yet fully optimised, we are already satisfied with the achieved performance.

We shall note here that the time required to perform the search depends on the actual grid configuration, its masks, and the selected interpolation stack. When the distance between a target point and the nearest non-masked source points is large, the nearest-neighbourhood search can become very time-consuming. In this case the required search radius might be very large and can contain the regions of multiple processes. This increases communication and computation time. Likewise, the calculation for the patch recovery is very costly when the 3rd order polynomial fit is selected together with the maximum number of fix points per source triangle. Here, reducing the number of fix points, selecting a lower order for the fit, or reducing the local problem size by employing more processes can significantly reduce the CPU time. We cannot provide a simple recipe to select the best interpolation configuration. Like it is the case with the tuning of model physics for new model configurations the interpolation results need to be checked, and interpolation parameters need to be adjusted when necessary.

The communication scheme described in 2.2 works much better than in OASIS4. Due to the general interface, different parts of the code can successfully use the same mechanism without any knowledge of the other parts. For example, implementing the communication for new interpolation methods can be done without any impact on existing communication related code. As a disadvantage this communication scheme makes the code in some parts hard to read. For some sections of the code we have improved the understandability by including communication diagrams in the Doxygen documentation.

Right from the beginning we have implemented algorithms from which we expect a reasonable weak and strong scaling behaviour. More sophisticated interpolation schemes like a second-order conservative remapping or a bicubic-spline interpolation can help to improve the quality of interpolation results. ~~A~~The build-in support for the coupling of vector fields greatly simplifies handling of these fields. Further down the road is the support for changing land-sea masks at run time. In the future, the abstraction of the communication layer might enable a thread-level parallel functional decomposition of the work performed by YAC in the initialisation phase.

5 Conclusions

In support of constructing coupled Earth system models, we ~~redesign and develop~~redesigned and developed a coupler from scratch. Our focus lies on the neighbourhood search, the calculation of interpolation weights and the data exchange between model components. The efficient and fully parallelised algorithms directly support unstructured and block-structured numerical grids. This efficiency allows for the online calculation of interpolation weights during the initialisation phase of each model run. Furthermore, we offer the alternative to read in interpolation weights which are generated offline by the climate data operator

(CDO) software. This significantly widens the scope of numerical model configurations for which our coupling software can be applied. The software is available to the climate modelling community and is used in the next generation Max Plank Institute for Meteorology Earth System Model. As an added value, the intentional choice of the programming language — in this case C — allows us to directly transfer parts of our software into the CDO and thus contribute to its improvement.

6 Code Availability

Information about access to our software is provided on our YAC Doxygen web site <https://doc.redmine.dkrz.de/YAC/html/index.html> under section code availability.

10 Appendix A: User Interface

~~Like it As~~ is the case with OASIS, our user interface is split into two parts. First we provide a set of library function calls that are to be called from the user code. Using these interfaces, the user makes static information, which is “known” by the model component at run time, available to YAC. This is required to perform the search and data exchange. To allow for some flexibility to explore the functionality of the coupler, a set of free parameters has to be provided at run time like the activation and deactivation of coupling fields or a particular choice of an interpolation or the definition of an interpolation stack. In OASIS3 this information is provided via the namcouple file, an ASCII formatted file, while in OASIS4 this kind of information is provided via XML formatted files.

XML

In the recent literature (Ford et al., 2012, and articles therein) it has been pointed out that XML is the preferred language to describe metadata. Similar to OASIS4 we use XML files for the description of the coupling. [Compared to OASIS4 the XML structure in YAC is simpler, more readable and smaller in size.](#) In the coupling XML file the user can specify the fields that have to be coupled. Other metadata are provided in the same file as well, like the interpolation sequence and the coupling frequency. In order to facilitate the generation of the coupling XML file a Java based Graphical user interface is available.

~~A drawback of the design of the configuration XML of OASIS4 is the high complexity of nested elements. Each element representing one couple contains a lot of sub-elements which includes information that in our opinion exceeds the XML purpose of describing the mere coupling, but instead describes some general properties of components and transients (physical coupling fields). This makes the XML instances hard to read and maintain. The aim of redesigning the coupling configuration XML structure is to simplify and reduce it to the elements necessary for coupling. The resulting new XML structure is much simpler, more readable, and smaller in size.~~

For defining the structure of XML elements, we provide an XML Schema Definition (XSD). It describes the elements, attributes, types, keys, and key references (Fallside and Walmsley, 2004). The XSD allows to validate the coupling XML

instances. Furthermore, it helps to adjust programs that access the XML files, since it ensures that the XML elements satisfy a predefined structure.

We provide a component XSD and a coupling XSD. An XML instance of the component XSD defines the element structure for basic information about components: model name, transients, grids, and timing information. It is used as the input for our Graphical User Interface (GUI). The GUI simplifies the generation of an XML instance of the coupling XSD, which provides detailed information on the actual coupling like which transients are to be coupled, the configuration of the interpolations, and the role of the transients, whether they act as source or target.

Graphical User Interface – the XML GUI

With the design of a minimised XML Schema, the complexity of XML instance files is reduced in a way that it ~~gets well~~ becomes easily readable for humans. Still, it is a tedious task to generate XML files by hand and to get all references correct. To ease the editing process of the coupling configuration, we provide a Java GUI that allows the user to create and manipulate the coupling XML configurations with a few clicks. The user can load component XML instances, couple them together, manipulate all coupling configuration settings, and produce a valid coupling XML file. For the implementation of the GUI we chose Java Swing because it provides a simple means for the creation of lightweight platform-independent, portable GUIs (Eckstein et al., 1998), and because Java provides libraries for easy XML parsing and file I/O.

The initial GUI is showing a pane that is split into 2 parts, left and right. The user can load a component XML to each side, and then couple the transients together. An arrow indicates the coupling direction. With a click on a couple, the GUI presents a detail window where the user can configure the interpolation settings, the time intervals, time lags, and some debug settings. Figure 4 shows a screenshot of the GUI in action. We provide a more detailed description of the GUI in our manual which is available as a supplement to this publication.

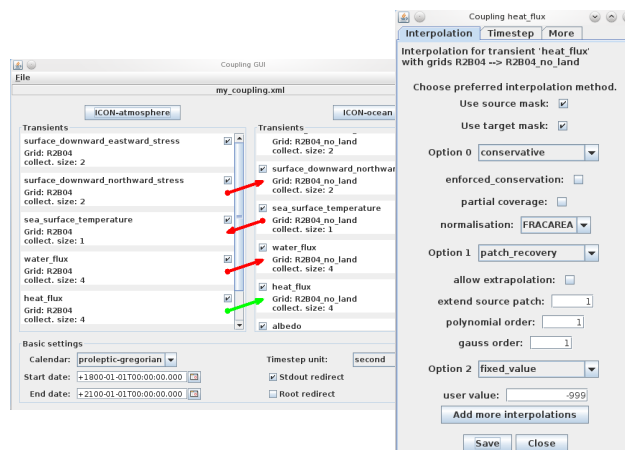


Figure 4. Coupling GUI with detail window.

Application Programming Interface - API

5 With our API we closely follow the philosophy of OASIS4. For a more detailed discussion about the main principle the reader is referred to Redler et al. (2010), for a detailed description of the YAC API the reader is referred to our Doxygen web site for YAC. As it is the case for the internal library routines the API is programmed in C. As ICON and most other climate codes are programmed in Fortran we provide a Fortran wrapper for the API. In order to get rid of the problem with underscores in symbol names which are handled different by individual Fortran compilers we rely on the ISO_C_BINDING module which
10 was introduced with the advent of the Fortran 2003 standard. The Fortran interface uses overloading with respect to data types and certain data structures. The Fortran interface accepts geographical grid information (longitudes and latitudes coordinates) as well as coupling fields in REAL or DOUBLE PRECISION. As the C API only accepts double Fortran REAL is internally converted to DOUBLE PRECISION before being passed to the C API. In Fortran grid coordinates for regular grids in longitude and latitude are passed through the same interface routines (due to overloading) as the grid data for unstructured grids while
15 in C different interfaces are provided for the different grid types. We do not provide any direct access to our internal data structure. Instead, we hand back opaque handles (integer values) for components, grids, masks, and fields.

Acknowledgements. We are grateful to our colleagues Jörg Behrens, Hendrik Bockelmann, and Thomas Jahns at DKRZ for very fruitful discussions during the development and adding the configure mechanism (TJ). Uwe Schulzweida took the burden and transferred several of our internal routines into the CDO package and was thus able to do far more testing of our algorithms than anticipated and thus helped us to make our software more robust for production use. We appreciate the luxury of having the freedom by our institutes, DKRZ and MPI-M to do some development outside of any project requirements and without knowing beforehand where this will lead us.

References

- 5 Antonov, J., Seidov, D., Boyer, T., Locarnini, R. A., Mishonov, A., Garcia, H., Baranova, O., Zweng, M., and Johnson, D.: World Ocean Atlas 2009, Volume 2: Salinity. S. Levitus, Ed. NOAA Atlas NESDIS 69, U.S. Government Printing Office, Washington, D.C., 184 pages, 2010.
- Barnes, N. and Jones, D.: Clear Climate Code: Rewriting Legacy Science Software for Clarity, *IEEE Software*, 28, 36–42, 2011.
- Clune, T. L. and Rood, R. B.: Software Testing and Verification in Climate Model Development, *IEEE Software*, 28, 49 – 55, 2011.
- 10 Easterbrook, S.: Code Design and Quality Control, in: Earth System Modelling - Volume 2: Algorithms, Code Infrastructure and Optimisation, edited by Bonaventura, L., Redler, R., and Budich, R., Springer Briefs in Earth System Sciences, pp. 51 – 65, Springer Heidelberg, 2012.
- Eckstein, R., Loy, M., and Wood, D.: Java Swing, O'Reilly & Associates, Inc., 1998.
- Fallside, D. C. and Walmsley, P.: XML Schema Part 0: Primer Second Edition, W3C recommendation, 2004.
- 15 Flato, G., Marotzke, J., Abiodun, B., Braconnot, P., Chou, S., Collins, W., Cox, P., Driouech, F., Emori, S., Eyring, V., Forest, C., Gleckler, P., Guilyardi, E., Jakob, C., Kattsov, V., Reason, C., and Rummukainen, M.: Evaluation of Climate Models, book section 9, pp. 741–866, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, doi:10.1017/CBO9781107415324.020, www.climatechange2013.org, 2013.
- Ford, R., Riley, G., and Redler, R., eds.: Earth System Modelling - Volume 5: Tools for Configuring, Building and Running Mode, Springer
20 Briefs in Earth System Sciences, Springer Heidelberg, 2012.
- Giorgetta, M. A., Jungclaus, J. H., Reick, C. H., Legutke, S., Bader, J., Böttinger, M., Brovkin, V., Crueger, T., Esch, M., Fieg, K., Glushak, K., Gayler, V., Haak, H., Hollweg, H.-D., Ilyina, T., Kinne, S., Kornblueh, L., Matei, D., Mauritsen, T., Mikolajewicz, U., Mueller, W. A., Notz, D., Pithan, F., Raddatz, T., Rast, S., Redler, R., Roeckner, E., Schmidt, H., Schnur, R., Segschneider, J., Six, K., Stockhause, M., Timmreck, C., Wegner, J., Widmann, H., Wieners, K.-H., Claussen, M., Marotzke, J., and Stevens, B.: Climate and carbon cycle changes
25 from 1850 to 2100 in MPI-ESM simulations for the coupled model intercomparison project phase 5, *Journal of Advances in Modeling Earth Systems*, 5, 572–597, doi:10.1002/jame.20038, 2013.
- Hill, C., DeLuca, C., Balaji, V., Suarez, M., and da Silva, A.: Architecture of the Earth System Modeling Framework, *Computing in Science and Engineering*, 6, 2004.
- Jacob, R., Larson, J., and Ong, E.: MxN Communication and Parallel Interpolation in CCSM3 Using the Model Coupling Toolkit, *Int. J. High Perf. Comp. App.*, 19, 293–307, 2005.
- 30 Jones, P. W.: A User's Guide for SCRIP: A Spherical Coordinate Remapping and Interpolation Package, <http://oceans11.lanl.gov/svn/SCRIP/trunk/SCRIP/doc/SCRIPusers.pdf>, 1998.
- Jones, P. W.: First- and Second-Order Conservative Remapping Schemes for Grids in Spherical Coordinates, *Monthly Weather Review*, 127, 2204–2210, doi:10.1175/1520-0493(1999)127<2204:FASOCR>2.0.CO;2, [http://dx.doi.org/10.1175/1520-0493\(1999\)127<2204:FASOCR>2.0.CO;2](http://dx.doi.org/10.1175/1520-0493(1999)127<2204:FASOCR>2.0.CO;2), 1999.
- 35 FASOCR>2.0.CO;2, 1999.
- MPI Forum: MPI: A Message-Passing Interface Standard Version 3.1, Tech. rep., Knoxville, TN, USA, 2015.
- Redler, R., Valcke, S., and Ritzdorf, H.: OASIS4 – a coupling software for next generation earth system modelling, *Geoscientific Model Development*, 3, 87–104, doi:10.5194/gmd-3-87-2010, <http://www.geosci-model-dev.net/3/87/2010>, 2010.
- Sutherland, I. E. and Hodgeman, G. W.: Reentrant Polygon Clipping, *Communications of the ACM*, 17, 32–42, doi:10.1145/360767.360802, <http://dl.acm.org/citation.cfm?doid=360767.360802>, 1974.

- Valcke, S.: The OASIS3 coupler: a European climate modelling community software, *Geoscientific Model Development*, 6, 373–388, doi:10.5194/gmd-6-373-2013, <http://www.geosci-model-dev.net/6/373/2013/>, 2013.
- 5 Valcke, S., Balaji, V., Craig, A., DeLuca, C., Dunlap, R., Ford, R., Jacob, R., Larson, J., O’Kuinghttons, R., Riley, G., and Vertenstein, M.: Coupling technologies for Earth System Modelling, *Geoscientific Model Development*, 5, 1589–1596, doi:10.5194/gmd-5-1589-2012, <http://www.geosci-model-dev.net/5/1589/2012/>, 2012a.
- Valcke, S., Redler, R., and Budich, R., eds.: *Earth System Modelling - Volume 3: Coupling Software and Strategies*, Springer, 96 pages, 10 2012b.
- Valcke, S., Redler, R., and Budich, R., eds.: *Earth System Modelling - Volume 3: Coupling Software and Strategies*, Springer, 96 pages, 2012c.
- Valcke, S., Craig, T., and Coquart, L.: *OASIS3-MCT User Guide*, OASIS3-MCT 3.0, Tech. Rep. 1875, CERFACS/CNRS SUC URA, Toulouse, France, 2015.
- 15 Wan, H., Giorgetta, M. A., Zängl, G., Restelli, M., Majewski, D., Bonaventura, L., Fröhlich, K., Reinert, D., Rípodas, P., and Kornbluh, L.: The ICON-1.2 hydrostatic atmospheric dynamical core on triangular grids - Part 1: Formulation and performance of the baseline version, *Geosci. Model Dev. Discuss.*, 6, 59–119, 2013.
- Zienkiewicz, O. and Zhu, J.: *The Superconvergent Patch Recovery and a Posteriori Error Estimates. Part 1: The Recovery Technique*, *International Journal for Numerical Methods in Engineering*, 33, 1331–1364, 1992.