

Reply to review #3

YAC 1.2.0: New aspects for coupling software in Earth system modelling.  
M. Hanke, R. Redler, T. Holfeld, and M. Yastremsky.

We would like to thank the referee for the concise and very constructive review. We understand the referees concerns. To a large extent we address the comments with the revised version of the manuscript. In black we copied the referees remarks and statements from the pdf file followed by our replies in blue, in *italic* we mark passages that we copied from the new revised version of the manuscript.

---

This version of the paper is significantly improved. The paper's focus is much improved and many problems have been addressed. However, the authors did not adequately address the major shortcoming noted by both reviewers in the original paper, lack of performance data. In addition, there are a few new issues.

Our intention for the paper was to present the aspects of YAC that make it different from other coupling solutions. Advertising YAC or its performance is still not our primary intention with this paper. To clarify this, we changed the sentence regarding the performance measurements in the abstract to:

*Preliminary performance measurements of a set of realistic use cases are presented to demonstrate the potential performance and scalability of our approach.*

Another similar sentence in the introduction was changes to:

*We present preliminary performance results for a set of realistic use cases in Sec. 3.*

In addition, we added the following paragraph to the performance section:

*In this section we present a first set of performance measurements to show the potential of YAC. As mentioned in Sec. 2.1, we do not yet concentrate on writing highly optimised code. Therefore, the measurements in this section are only preliminary. They can be seen as an estimate for the upper bound of the performance of YAC. In addition, scaling characteristics can be derived from them.*

We understand that the reviewer is interested in performance data of YAC. However, currently YAC is not at all optimized for performance. Therefore, the results only give general information on the scalability of our approach and the upper bound for the performance using idealised tests.

Nevertheless, in the latest version (2. revision) of the paper we added additional performance measurements and their analysis. The respective section is now the biggest one of the paper. In order to not further shift the focus of the paper, we are strongly against further extending this section. If the reviewer is still interested in even more performance data, we would like to refer the reviewer to the ESiWACE-project (<https://www.esiwace.eu/>).

- Section 4 has several spelling and gramatical errors, that section should be reviewed carefully

measurment -> measurement

"The performance of the exchange step itself was in the tested setup good for all node counts." needs to be written more clearly

"The measurments 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." is something I think is incorrect.

We have carefully checked section 4 (discussion and summary). Apart from the one spelling error already pointed out by the referee did not detect any further spelling errors. We use British English throughout the paper which is consistent with the Copernicus guidelines. Based on recommendations by a native speaker and as a consequence of the request by the reviewer to revise and add measurements we have revised this section, among them the two sentences pointed out by the referee. For brevity we have not included the revised section here as the details are accessible from the difference section below.

- I think the new section discussing the remapping cost is probably not correct and needs to be updated. That's page 15, lines 12-16. The authors state performance was random and fell over 2 orders of magnitude. I believe the performance should vary with nodes/cores and timing numbers should be reasonably reproducible. My sense, based on the statement in the paper, is that the timing may need to be redone, with appropriate barriers and/or a larger sample size. You cannot measure one exchange in these kinds of tests, the time required is too small and is often shorter than the resolution of the timers. The point is that over the course of a long coupled run, coupling is done thousands or millions of times and it's that aggregate time that's important. At high resolution or high pe counts, the coupling time can be a bottleneck. To get an accurate representation of the time, the ping pong test should couple data back and forth for at least a few wall-seconds in total, maybe 1000s or 100s of thousands of times. I would like to see a scaling curve, similar to Figure 3 for the ping pong test for a couple of resolutions that accurately reflect the communication/remapping cost.

Indeed, in the first set of measurements we did too few exchanges per measurement. New and more accurate measurements have been added to the paper and accordingly the performance section has been completely revised. For brevity we have not included the revised section here as the details are accessible from the difference section below.

In addition, I think the time for remapping and the time for coupling could be separated in the presentation of the data. It would be interesting to know how much of the ping pong time is spent remapping (including halo update and application of mapping weights) and how much is spent in the concurrent transfer of data. There are cases when remapping is not needed.

Due to the asynchronous communication that we use for the data exchange between the source processes and the transfer of the remapping results from the source to the target processes, a clear separation and individual measurements of the steps would

not reflect the actual behaviour of YAC. As an alternative, we now provide additional numbers for our fixed interpolation, which is equivalent to a concurrent transfer of data between source and target – but without halo updates and without an application of mapping weights. (see page 15, line 20 ff)

- There is a set of timing data called "file" on figure 3 but very little information is provided in the text how this is implemented (is it parallel I/O or "read and scatter" or something else) and there is no accompanying discussion. The results are interesting and I feel I am left hanging. I think some discussion needs to be added to discuss that set of data. This is interesting data.

The file interpolation mentioned in figure 3 is introduced earlier, but with a slightly different naming. To avoid confusion, we now use the same name for the interpolation throughout the text (file interpolation). Furthermore we have added some information about how the reading is done:

*To avoid problems with memory consumption and runtime scalability, we use a parallel input scheme to read in the weight file. A subset of the source processes is selected to do the I/O. Each of these processes reads in an individual part of the data from the file. The data is then stored in a distributed directory (Pinar and Hendrickson, 2001). Afterwards, the source processes access the directory to get the data required to do the interpolation of the target points assigned to them.*

- Both reviewers felt the paper fell short on performance data and the paper has not substantially improved in that area. As indicated above, some scaling curves for remapping and data transfer would make the paper significantly better. I also think additional data should be provided on the weights generation, for additional resolutions and more mapping options. The introduction also talked about how high resolution was the reason fast weights generation was important, but there are no high resolution results in the paper. The authors suggested moderate resolution is not interesting and no "significant challenge", but then the authors only present results at moderate resolution. I feel this inconsistency needs to be addressed.

We already provide the numbers for the weight generation for patch recovery, 1<sup>st</sup> order conservative remapping and the file interpolation. In the new revision measurements for the fixed interpolation were added.

We would like to note again that the whole search and weight generation is implemented by making extensive use of asynchronous communication. As already mentioned for the ping pong exchange, artificially breaking the initialisation apart for more detailed timing results would not reflect the real behaviour of our approach. We prefer to only show performance numbers that are relevant for the end-users of YAC. However, the newly added measurements for the fixed interpolation allow a slightly deeper analysis when compared with the other interpolation methods (see revised sections about performance and discussion).

As requested by the reviewer we added measurements for an additional grid configuration. With ~22 km in the atmosphere and ~12 km in the ocean it uses grids

with resolutions two times higher than the original ones.

In the summary (see below) the referee explicitly asks for performance numbers for the nearest neighbour interpolation. YAC supports nearest neighbour interpolation. The nearest neighbour interpolation is listed in the paper as it is part of the functionality of YAC and is working correctly. However, we still see significant performance issues with this interpolation, which are already addressed in the original version of the manuscript. Thus the work on this interpolation is not yet finalised. We strongly prefer not to show performance numbers in a publication for an algorithm, which is in some sense still under construction.

- Several of the authors' responses to the reviewers comments did not seem particularly well thought out. Without focusing on that too much, I think there are a few important points that don't necessarily reflect on the current paper but might on YAC.

First, a round robin or similar decomposition can be an excellent decomposition in cases where the cost per gridcell varies in time and/or space, like in radiation calculations (diurnal), sea ice coverage (seasonal), or land cover models (cost can vary diurnally and seasonally by location). This decomposition probably deserves more consideration than "We do not mind not supporting round robin like decompositions"

We apologise for having been too superficial in our first response. We fully agree with the referee that a round robin or similar decomposition does have advantages. Nevertheless, we do not claim that YAC is already complete. We prioritise our work such that we are able to support those applications for which we would like to use our software first and possibly add more functionality when required at a later stage.

Actually, with some limitations these types of decomposition are supported. The file interpolation makes no use of the halo cells. With this interpolation method YAC could theoretically support this type of decomposition. However, we did not yet test such a use case.

Second, the requirement to provide connectivity information as part of the weights generation is not state-of-the-art. The author's response, "ICON can do this", "models with diffusion and advection operations will have the information readily available", and "it was too hard in OASIS4" are poor excuses. If the tools in YAC are being designed to work generally, then other models will not be ICON, diffusion and advection operators may implicitly carry the connectivity information but accessing that information to pass to YAC might not be straight-forward. In addition, there is no guarantee that the same halo is required for a diffusion or advection operation and for the YAC weights generation. Finally, the fact that OASIS4 was not particularly successful in computing connectivity does not mean it's not possible. Other weights generation tools are able to compute connectivity on the fly.

We fully agree with the referee that the computation of connectivity on the fly is possible. What we tried to say in our first reply was – and obviously we failed in this first attempt – that we find it not straightforward to perform such a computation in a fully parallel, efficient, and robust way. At least this was the lesson that we learned

from OASIS4, and with the first version of YAC we preferred to bypass this complexity and invest our time and effort on other – in our very personal opinion more interesting – functionality like the interpolation stack.

In our first reply we already stated that we see an issue with pure column-based model components like e.g. land models. For the neighbourhood search we currently do require a halo with a width of at least one cell around the compute domain. We use this halo solely to identify neighbouring processes. The stencils used for the interpolation are not limited by the halo and can even go across multiple compute domains. Our point is that most applications that operate on dynamical kernels are able to provide the halo information that YAC requires, and the applications that we have worked with so far are able to provide this halo either directly or with very minor effort.

As the referee states, these remarks do not apply to the manuscript but more to aspects of the software itself. Therefore, we have not changed the manuscript to further discuss these aspects.

In summary, I believe the paper needs to be further revised to

- review section 4 for grammar and spelling.

We have revised section 4.

- show accurate timing data for the ping pong test and if possible, to separate the remapping and coupling times into distinct terms.

We have revised the timing data for the ping pong test. We did not explicitly separate remapping and coupling, but added measurements for the fixed interpolation, which basically is an exchange without remapping.

- add some discussion of the "file" data on figure 3.

We cleaned up the wording to ensure that we use identical phrases throughout the text for this type of interpolation. The technical implementation of the reading of a weight file and the distribution of the weights to the source processes is now described. We compare the file interpolation with the online computation of weights for the conservative remapping.

- extend the weights generation performance data further to include several additional mapping options including especially nearest neighbor.

Above we provide our reasons why we prefer not to show performance data for the nearest neighbour interpolation. We already provide performance data for the most important and quite costly interpolations (patch recovery, 1<sup>st</sup> order conservative remapping, file). In addition, data for the fixed interpolation was added. We are convinced that these numbers provide a firm estimate on the scalability of the approach used in YAC and on the upper bound for the absolute performance that is to be expected from YAC.

Recent changes to our testing system led to changes in the performance data. Especially for high node counts the results are now better compared to the measurements provided in the previous revision of the paper.

- provide timing results for the ping pong, file reading, and weights generation on at least one high resolution configuration as this is a major driver of the current implementation.

We added timing results for a two times higher horizontal resolution (~22 km in the atmosphere and ~12 km in the ocean) and discuss them. As the benchmark runs occupy a significant portion of the whole system at DKRZ, our ability to perform more benchmark tests in near future is limited. We would like to note again that performance is not the main topic of our paper. Nevertheless, we are convinced that the additional performance results now provide a good estimate for the scaling behaviour of YAC and what we can expect from YAC when applied to even higher resolutions.

In summary, as a consequence the performance and discussion sections have been substantially revised.

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

Moritz Hanke<sup>1</sup>, René Redler<sup>2</sup>, Teresa Holfeld<sup>2</sup>, and Maxim Yastremsky<sup>2</sup>

<sup>1</sup>Deutsches Klimarechenzentrum, Hamburg, Germany

<sup>2</sup>Max-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 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~~ Preliminary performance measurements of a set of realistic use cases are presented to demonstrate the potential performance and scalability of our approach. The coupling software is now used for the coupling of the model components in the Icosahedral nonhydrostatic (ICON) general circulation model.

## 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 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).

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. 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 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.

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 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 can be encapsulated which allows further splitting of individual model components. The software package encompasses about 800,000 lines of Fortran90  
10 and C++ code<sup>1</sup>.

Many coupled climate models use the Ocean Atmosphere Sea-Ice Surface (OASIS) coupler maintained by CERFACS in its version 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 grid types, among them unstructured grids. For the data transfer, complete fields are collected by individual OASIS3 processes. With a  
15 quite recent development OASIS3 uses the Model Coupling Toolkit (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 (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 Modeling Framework (ESMF) can be used to provide the required interpolation weights a priori.

20 Further details about those and other important approaches available so far have been assembled in a collection of articles by Valcke et al. (2012b) and in a review article by Valcke et al. (2012a). So why are we not going to use an existing software solution and adapt 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 modelling. Our primary focus for our software development is on the Icosahedral nonhydrostatic (ICON) general circulation model (Wan  
25 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 (CDO<sup>2</sup>) software.

ICON is designed to run on massively parallel systems. As a consequence we require the coupling software to take full advantage of the hardware and to exchange coupling fields efficiently and in parallel. Furthermore, the handling of large grids on  
30 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 weights rather than to perform these steps off-line in a pre-process step. Last but not least the software has to be sufficiently

---

<sup>1</sup>[https://www.earthsystemcog.org/projects/esmf/sloc\\_annual](https://www.earthsystemcog.org/projects/esmf/sloc_annual)

<sup>2</sup><https://code.zmaw.de/projects/cdo>



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-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 have 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.

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 ~~investigate the performance of the software present preliminary performance results for a set of realistic use cases~~ in Sec. 3 ~~and further discuss and compare our approach.~~ Further discussion of the results are 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 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. 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 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 Valgrind<sup>3</sup> 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 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. Currently, the CDOs use in parts the same algorithm for the neighbourhood search and calculation of the interpolation weights as OASIS3. 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.

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.

---

<sup>3</sup><http://www.valgrind.org>

## 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 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 Doxygen<sup>4</sup>. 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 site<sup>5</sup> 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 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 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. We have kept the individual functions short to increase the readability of the code.

---

<sup>4</sup><http://www.doxygen.org>

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

## 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.

## 5 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 have deviated from this paradigm, because based on 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  
10 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 unstructured grid, and we see some potential for optimisation in order to speed up the search on this type of source grid.

## 15 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.

The major advantage of this approach is that different parts of the code only depend on the interfaces of each other rather than  
20 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.

One example of such an object within YAC is the grid. YAC supports multiple grid types. Each grid type has its own method  
25 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.

The interpolation method object is another example. In YAC an interpolation consists of a number of arbitrary interpolation  
30 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

## 2.2 Communication

YAC can be considered as a kind of abstraction which hides the complexity of the Message Passing Interface (MPI, MPI Forum, 2015) for 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.

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 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 a non-blocking buffered send from OASIS4. 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 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 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. 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-

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 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 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 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 have implemented two different algorithms: a bucket- and a tree-based search algorithm.

#### 20 **Bucket-based search:**

Our first approach was has been 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.

25 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:**

30 As an alternative to the bucket search a search tree was-is 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.

5 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 conservative remapping (Fig. 1a). Here only those target cells get a value which are completely covered by valid (non-masked) ~~target-source~~ 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 no value  
10 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~~ salinity with value of less than 34 ~~are-colored-is coloured~~ in light blue.

The interpolation stack can easily be modified via the XML-Extensible Markup Language (XML) interface (see ~~below~~Appendix) 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.


15 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

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  
20 overlapping source polygons. Based on this information, the interpolation starts by distributing the target points, which need to be interpolated, among the source processes. To process the received target points, the global search provides on the source processes for each received target polygon a list of all overlapping local polygons.

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  
25 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 main remaining task is to compute the partial overlaps between  
30 source and target cells, which are used to generate the interpolation weights. Different algorithms for the calculation of a



**Figure 1.** 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)

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’Huillier’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-  
5 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.

10 **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~~centres. The diameter of these circles is chosen such that all points found for the respective target points are within the circle. In case a  
15 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 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.

20 **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 to do the interpolation.

**~~Precalculated weights~~File:** The file interpolation is able to read in ~~NetCDF files that contain weights~~interpolation weights from Network Common Data Form (NetCDF) files. 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 or CDO).

To avoid problems with memory consumption and runtime scalability, we use a parallel input scheme to read in the weight file. A subset of the source processes is selected to do the I/O. Each of these processes reads in an individual part of the data from the file. The data is then stored in a distributed directory (Pinar and Hendrickson, 2001). Afterwards, the source processes  
30 access the directory to get the data required to do the interpolation of the target points assigned to them.

The common use case for applying this interpolation method is to compute the weights once and then reuse them in order to 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 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 ~~-(see Sec. 2.4 and Fig. 1 c).~~

## 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 value. ~~And this~~ 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, 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 SCRIP 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 SCRIP User Guide provides an example threshold value of  $\pm 1.5$  (in radian). Even though the transformation improves accuracy, it also generates a discontinuity. A similar approach is used in OASIS4, where this technique causes problems in some very rare cases, because it ~~generated~~ generates "holes" in the grid; some parts of the sphere seemingly ~~were~~ are 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 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 [have](#) tried to approximate the latitude circle edges using sections of great circles, but this increases computation time and makes the cell concave even though 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 [have](#) used Girard's Theorem. While testing we [have](#) 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  $0.5^\circ$ ) 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  $5^\circ$ ) where the error goes up to 71‰.

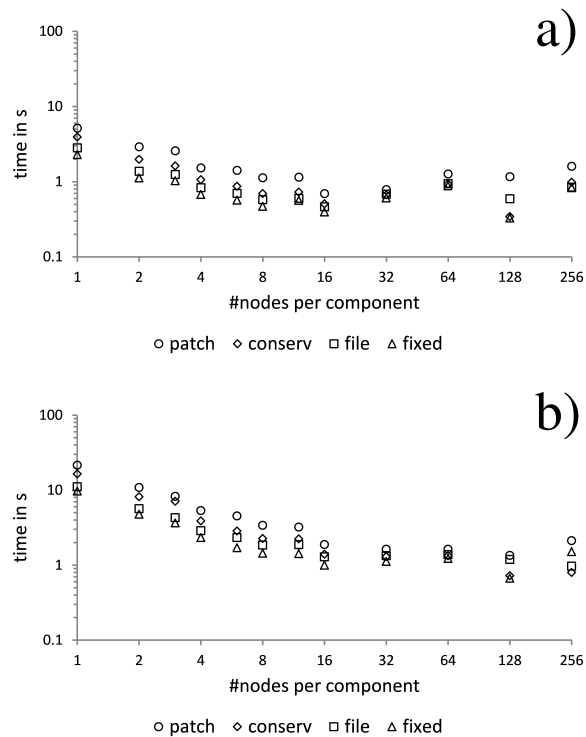
### 3 Performance

In this section we present a first set of performance measurements to show the potential of YAC. As mentioned in Sec. 2.1, we do not yet concentrate on writing highly optimised code. Therefore, the measurements in this section are only preliminary. They can be seen as an estimate for the upper bound of the performance of YAC. In addition, scaling characteristics can be derived from them.

We have 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 ?? shows scaling measurements.~~ All scaling measurements presented here were done on mistral at 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 5.1.2.150. The measurements ~~ran~~ have been run using 24 MPI processes per node. ~~The~~ and performance results are provided for two different grid configurations. For the first configuration the toy model `perf_toy_icon` ~~used~~ uses an ICONR2B06 grid, which has 245760 cells and around 35km distance between the ~~centers~~ centres of neighbouring cells. In a typical atmosphere-ocean coupling setup, the ocean model component ~~usually~~ usually has a higher grid resolution than the atmosphere model component. To mimic this we ~~used~~ use a 405x405-cubed-sphere grid (984150 cells; around 25km distance between cell ~~centers~~ centres) for `perf_toy_cube` in this configuration, which represents the ocean in our setup. ~~Three~~ The second configuration uses grids with around four times more cells: ICONR2B07 (983040 cells) in perf\_toy\_icon and 810x810-cubed-sphere grid (3936600 cells) in perf\_toy\_cube. Four different interpolation methods ~~were~~ are measured: 1st-order conservative interpolation (`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), and an interpolation that assigns a fixed value to all target points (fixed). Each component is run on its own set of nodes.

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



**Figure 3.** Time required for global search and calculation of weights for different interpolation methods in both directions between a) an ICONR2B06 and 405x405-cubed-sphere grid and b) an ICONR2B07 and 810x810-cubed-sphere grid

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~~ use the same number of processes for each one.

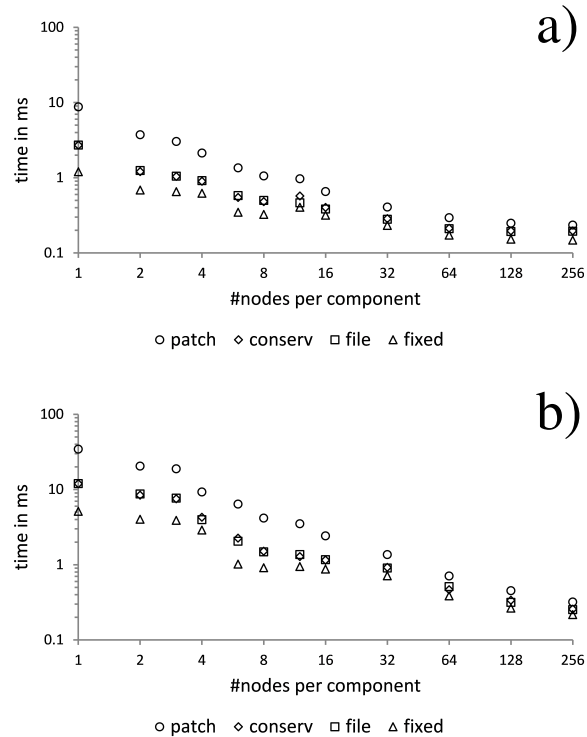
~~For low number of nodes per component (A special case is the fixed interpolation, it does not compute any weights. The measurements for yac\_csearch using the fixed interpolation mainly represent the time consumed by the interpolation-independent global search. Here, the measured time represents the lower bound for the initialisation of the search in the respective grid configuration.~~

25 ~~For up to 16~~ nodes per component the computation required by the interpolation-independent global search and the interpolation-specific weight computation are the two dominant factors 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~~ for up to 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.

30 ~~The three measured interpolation methods generate different workloads, with patch recovery being~~ From the measurements it is evident, that the patch recovery interpolation is in our case 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

35 ~~three methods converge for high number of nodes as the communication dominates.~~ interpolation method. The file interpolation does not compute weights, instead it reads them from file and distributes them among the source processes. The time required for the reading and distribution of the weights is the difference between the file and the fixed measurements.

~~Since, the wall clock time for yac\_earch in our test case is still 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. Even with 1024 nodes per component, which could not be tested due to the limitations of the available hardware,~~



**Figure 4.** Time required for ping-pong exchange between a) an ICONR2B06 and 405x405-cubed-sphere grid and b) an ICONR2B07 and 810x810-cubed-sphere grid

5 The measurements for the ~~results should still be acceptable~~ ping-pong exchange depicted in Fig. 4 represent the time required by perf\_toy\_icon to execute MPI\_Barrier followed by a put and get call. Perf\_toy\_cube executes the required matching calls. To improve accuracy we have measured the total time it took to execute a significant number of the ping-pong exchanges and divided the result by the number of exchanges.

10 A put consists of three main parts. At first, the source processes exchange halo data among themselves, if required by the interpolation. Afterwards, each source process computes the target points assigned to it using the interpolation weights. The results are then sent directly to the targets.

~~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. As for~~

the `yac_csearch`, the fixed interpolation is a special case. It does not require halo exchanges between source processes, the application of the remapping weights, and the transfer of the results or the computation of target points, but it does send target points containing the fixed values from the source 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. Again, the time shown here for the fixed interpolation can be considered as the lower bound for an exchange.

The patch recovery interpolation, with the parameters used in this setup, has the most weights per target point, which is why it is the slowest interpolation method for the ping-pong exchange. The conservative and file interpolation use identical weights and a nearly identical communication patterns for the source-to-source and source-to-target communication, resulting in identical interpolation results on the target processes. Hence, it is comprehensible why also the time measurements for these two methods are very similar.

Both tested grid configurations exhibit similar behaviour for the measurements of `yac_csearch` and the ping-pong. Obviously, the second configuration takes longer due to the higher resolution. A higher resolution enables better scaling, which is most visible for high node counts in Fig. 4 b).

#### 4 Discussion and Outlook

In the context of describing YAC Here, 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 For the first development step we have focused on solving the, we focus on solving existing problems and on providing the basic functionality. At the same time it is important to have future problems in mind, we take potential future problems into consideration and design the software such that a later integration of new programming modules is not hampered or blocked. We already went have gone through this exercise when we have added the first-order conservative remapping or and the patch-recovery interpolation (Sec. 2.5). In a second step we have put our focus on the optimisation, for example by focus on optimisation, such as providing an alternative to the bucket search (Sec. 2.3) or by introducing the and introducing 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 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 prove to be very useful during the development and also in the exploration phase. It However, it remains a challenge to achieve full test coverage and to test all possible use cases and all internal functionality. Thinking about Designing reasonable tests and the implementation of those test takes tests take some time. At first Initially, this tends to slow down the generation of source code for the library part library source code. 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 very early stage and improve the robustness of existing code; 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 in terms of the number of lines of code, the compile time and the development cycle. This efficiency allows us to test various methods and to experiment within a concise software

framework. ~~We are able to try out and compare algorithms and methods. Since most parts~~ In addition, the code development for YAC also provides direct contributions to the development of CDO, since most parts of YAC are programmed in C ~~we directly contribute with our methods to the CDO development.~~ As the CDO software is used worldwide in climate research, the parts of YAC, which are already used ~~included~~ in the CDO, will probably be used more extensively than YAC itself.

In Sec. 3 we provide some preliminary performance measurements for ~~a~~ typical model or grid configuration ~~configurations~~. The initialisation phase of YAC scales reasonably well with the number of MPI processes provided that the local problem size remains sufficiently large. With an increasing number of MPI 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. ~~The performance of the exchange step itself was in~~ As already mentioned in Sec. 3 the time required for the tested setup good for all node counts. ~~The measurments suggest that the time for interpolation-independent global search corresponds with the measurements for the fixed interpolation. From this we can derive that the previous statements on the scaling behaviour applies to the global search as well as the interpolation-specific weight computation, which is the difference between fixed and the other interpolation methods.~~

In our test case the file interpolation uses precomputed conservative interpolation weights and can thus be compared with the (online) conservative interpolation. For small node counts of up to 16 nodes the online calculation of the weights is slower than the reading and distributing of the precomputed weights, but not by a huge margin. For 32 to 64 nodes there is hardly any difference between the two interpolation methods. Depending on the chosen interpolation, the grid configuration, and the number of MPI processes involved, there is the potential for gaining a few seconds for lower node numbers, but the advantage diminishes to fractions of a second for higher node numbers (greater than 16 in our case).

For low to medium number of nodes the exchange step scales well. Even for higher node counts the runtime still decreases with increasing number of nodes per component. As explained in Sec. 3, the results for the fixed interpolation mainly consist of 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.~~ time required to send the interpolation results from the source to the target processes. In addition to this, the patch recovery interpolation also requires a data exchange between the source processes and the remapping. The time required to perform these two steps account for the difference between the measurements of the two methods. The remapping should scale nearly perfectly with the number of nodes. This is probably the reason why the patch recovery interpolation scales better than the fixed interpolation for low to medium number of nodes, depending on the problem size. For higher number of nodes the time required for the remapping should be negligible. The source-to-source and the source-to-target communication should have similar scaling behaviour. This can be seen especially well in Fig. 4 a), there both methods show nearly identical scaling behaviour for node counts higher than 16.

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~~ encompass the compute domains of several MPI 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 as it is done here.~~ Reducing the number of fix points, selecting a lower order for the fit, or reducing the local problem size by employing more MPI 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~~ Similar to tuning 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, ~~which~~ allows different parts of the code ~~can to~~ 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~~ However, a disadvantage of this communication scheme is that it makes the code in some parts hard to read. For some sections of the code we ~~have improved the understandability by including~~ include communication diagrams in the Doxygen documentation ~~for clearer illustration of the interdependency between functions.~~

~~Right from the beginning~~ In its current status YAC provides all functionality required to couple the Earth system model components we have in our focus and at the same time leaving ample room for extensions and improvements. From the very 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~~ bicubic spline interpolation can help to improve the quality of interpolation results. ~~The build-in~~ A built-in support for the coupling of vector fields greatly ~~simplifies~~ will greatly simplify the 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 ~~have~~ 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 Planck 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.

### Appendix A: User Interface

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 American Standard Code for Information Interchange \(ASCII\)](#) 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.

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 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 5 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.

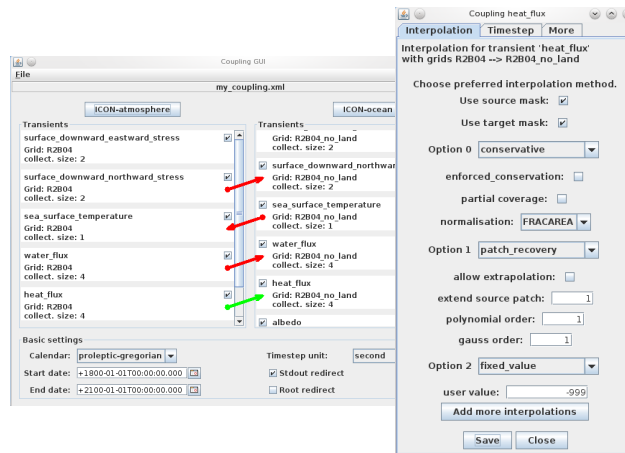


Figure 5. Coupling GUI with detail window.

## Application Programming Interface - API

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 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

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  
5 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

- 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, 10 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.
- 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, 15 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.
- 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– 20 866, 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 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, 25 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 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.
- 30 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.
- 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. 35
- 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.
- MPI Forum: *MPI: A Message-Passing Interface Standard Version 3.1*, Tech. rep., Knoxville, TN, USA, 2015.
- Pinar, A. and Hendrickson, B.: Communication Support for Adaptive Computation, <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.32.7896>, proceedings of the Tenth SIAM Conference on Parallel Processing for Scientific Computing, Portsmouth, Virginia, USA, 5 March 12-14, 2001. SIAM, 2001, 2001.

- 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.
- 10 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., 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.
- 15 Valcke, S., Redler, R., and Budich, R., eds.: *Earth System Modelling - Volume 3: Coupling Software and Strategies*, Springer, 96 pages, 2012b.
- 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.
- 640 Wan, H., Giorgetta, M. A., Zängl, G., Restelli, M., Majewski, D., Bonaventura, L., Fröhlich, K., Reinert, D., Rípodas, P., and Kornblueh, 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.
- 645