



The OASIS3 coupler: a European climate modelling community software

S. Valcke

CERFACS, Sciences de l'Univers au CERFACS, URA1875, 42 Av. G. Coriolis, 31057 Toulouse Cedex 01, France

Correspondence to: S. Valcke (valcke@cerfacs.fr)

Received: 13 July 2012 – Published in Geosci. Model Dev. Discuss.: 31 July 2012
Revised: 11 January 2013 – Accepted: 11 February 2013 – Published: 14 March 2013

Abstract. This paper presents the OASIS3 coupling software used in five of the seven European Earth System Models (ESMs) participating to the Fifth Coupled Model Intercomparison Project (CMIP5). A short history of the coupler development is followed by a technical description of OASIS3. The performances of a few relatively high resolution models coupled with OASIS3 are then described. It is shown that, although its limited field-per-field parallelism will eventually become a bottleneck in the simulation, OASIS3 can still be considered an appropriate tool for most of these relatively demanding coupled configurations. Its successful use in different CMIP5 ESMs is then detailed. A discussion of the benefits and drawbacks of the OASIS3's approach and a presentation of planned developments conclude the paper.

1 Introduction

Coupled General Circulation Models (CGCMs) offer numerical representations of the different components of the climate system and their interactions. CGCMs include numerical codes simulating, for example, the atmosphere, the ocean, the land surface and the sea ice. Since the late 1960s, CGCMs have been used to simulate the climate of the Earth System (Manabe and Bryan, 1969) and, with the constant increase of computing power, climate models have continuously grown in resolution and complexity. Today, CGCMs are used for a variety of purposes from studies of the climate dynamics to decadal and centennial projections of the future climate.

Classically, an Atmospheric General Circulation Model (AGCM) including a land surface model coupled to an Oceanic General Circulation Model (OGCM) incorporating a sea-ice model form the basis of a climate CGCM. More

recently, the trend is to include additional components representing e.g., the atmospheric chemistry, the marine biology or the carbon cycle. In doing so, climate CGCMs are constantly becoming more complete and complex Earth system models (ESMs).

The development of climate CGCMs capable of simulating and assessing the climate system over a wide range of space and time scales is one of the main objectives of the World Climate Research Programme (WCRP) of the World Meteorological Organisation (WMO), with the mission to “facilitate analysis and prediction of Earth system variability and change for use in an increasing range of practical applications of direct relevance, benefit and value to society” (see <http://www.wcrp-climate.org/mission.shtml>). Most importantly, climate CGCMs are the backbone of the climate simulations that form the basis of the periodic Assessment Reports (AR) published by the Intergovernmental Panel of Climate Change (IPCC). For example, 26 CGCMs assembled by 19 groups around the world were used for the 4th Assessment Report released in 2007; currently, 28 groups running a total of 59 CGCMs are participating in the on-going Fifth Coupled Model Intercomparison Project (CMIP5) of which the fifth Assessment Report due in 2013 will be based.

Typically, the different components of a CGCM are developed independently by different research groups that also use these components in stand-alone mode (i.e., uncoupled) to investigate processes or test new physical parameterisations in controlled experiments. This naturally leads to the conclusion that CGCMs should be made up of separate interoperable components. Most commonly in Europe, the OASIS (which stands for Ocean Atmosphere Sea Ice Soil) approach is used, where the component models remain separate applications and an external coupling software performs the

communications, but requires only minimal changes to the component codes.

In this paper, we present in detail the OASIS3 version of the coupler and its use in CMIP5. We start with a short history of the coupler development in Sect. 2. There follows a technical description of OASIS3 in Sect. 3, detailing the Application Programming Interface (API) of the OASIS3 coupling library, the coupling configuration and the coupling field transformations and regriddings offered by the coupler. Some numbers illustrating the performance of OASIS3 for few coupled configurations are then provided in Sect. 4 and its use, in particular in different CMIP5 CGCMs, is detailed in Sect. 5. We finally discuss the benefits and drawbacks of the OASIS3's approach in Sect. 6 and conclude with some perspectives on future developments.

2 History

When research in climate modelling started at CERFACS in 1991, the first objective was to couple the ocean General Circulation Model OPA developed by the Laboratoire d'Océanographie Dynamique et de Climatologie (LODYC) to two different atmospheric General Circulation Models, ARPEGE and LMDz developed by Météo-France and the Laboratoire de Météorologie Dynamique (LMD), respectively. The initial period of investigation led to the conclusion that the technical coupling between the ocean and atmosphere codes should take the form of an external coupler, i.e., a coupling library linked to the components, exchanging the coupling data with a separate application performing the regridding of the coupling fields. This choice ensured a minimal level of changes to the existing codes while focussing on modularity and portability. As the coupling was, at this stage, involving only a relatively small number of 2-D coupling fields at the air-sea interface, efficiency was not considered a major criterion.

Two years later, a first version of the OASIS coupler was used in a 10-yr coupled integration of the tropical Pacific (Terry et al., 1995). At the time, the communication, i.e., the exchange of coupling fields between the atmospheric and oceanic applications, was ensured through CRAY named pipes and ASCII files. In 1995, a major rewriting and the introduction of a new communication library based on Parallel Virtual Machine (PVM, see <http://www.csm.ornl.gov/pvm/>) lead to the OASIS2 version. OASIS2 has been used by different groups in France, at the European Centre for Medium Range Weather Forecast, and at CERFACS, in particular in a heterogeneous computing experiment into which the ocean model and the atmosphere model of a coupled system were respectively run on Météo-France CRAY II in Toulouse and on the Electricité de France CRAY C98 in Paris (Cassou et al., 1998). Between 1996 and 2000, alternative communication techniques based on UNIX System V Interprocess Communication (SVIPC) and on the Message Passing Inter-

face (MPI) were introduced while the community of users was constantly growing in Europe, but also in Australia and in the USA. From 2001 until 2004, the development of OASIS benefited from support from the European Commission in the framework of the PRISM project (Valcke et al., 2006) during which the OASIS3 version, including in particular a new API, was released. As detailed below, the OASIS3 coupler is used today by more than 35 different climate modelling groups around the world and forms the coupling infrastructure of different versions of five of the seven European CGCMs participating in CMIP5.

3 Technical overview

We provide here a technical overview of the use and functionality of the OASIS3 coupler. Further details on OASIS3 technical aspects can be found in the OASIS3 User Guide provided as an electronic supplement to this article.

3.1 Component model interfacing

In a system coupled with OASIS3, the coupler itself forms a separate binary that performs driving and regridding tasks. The original component models remain individual binaries in the UNIX sense with their main characteristics, such as internal parallelisation or I/O, untouched with respect to their uncoupled (stand alone) mode. As the coupler does not control the definition of global parameters (e.g., the total run duration or the calendar), the user has to take care that the component models define these parameters consistently. To interact with the other components through the coupler binary, the component models need to use the OASIS3 coupling interface library to perform the different coupling steps described in more details in the next paragraphs.

– Coupling initialisation

By calling the `prism_init_comp`¹ routine, all component model processes initialise the coupling. The MPI communicators that will be used for the coupling exchanges are established and the OASIS3 coupler transfers to the component models most of the coupling configuration information defined by the user in the “name-couple” configuration file (see Sect. 3.2).

If needed, each component then retrieves a local communicator for their internal parallelisation with a call to the `prism_get_localcomm` routine. In fact, OASIS3 supports two ways of starting the binaries of the coupled application. If a complete implementation of the MPI2 (Gropp et al., 1998) is available, only the OASIS3 binary has to be started by the user; all component binaries are then launched by the OASIS3 coupler at the beginning of the run using the MPI2

¹All OASIS3 API routine names end with `_proto`; for simplicity, we omit this suffix everywhere in this article.

MPI_Comm_Spawn functionality. If only MPI1 (Snir et al., 1998) is available, the OASIS3 binary and the component model binaries must all be started at once in the job script in a “multiple programme multiple data” (MPMD) mode. The advantage of the MPI2 approach is that each component keeps its own internal communication context unchanged with respect to the standalone mode, whereas in the MPI1 approach, OASIS3 needs to recreate a specific communicator for each component model that must be retrieved by the component model (with the `prism_get_localcomm` routine) and then used for its own internal parallelisation.

– Grid data file definition

To perform the regridding of the coupling fields, the OASIS3 coupler needs the definition of all source and target grids. At run time, OASIS3 reads this information from NetCDF auxiliary grid data files `grids.nc`, `areas.nc`, and `masks.nc`. The file `grids.nc` must contain the longitude and the latitude of the grid points onto which the coupling fields are defined and also the longitude and the latitude of the corners of the grid mesh associated to the grid point (used for conservative remapping). The file `areas.nc` contains the surface of all grid meshes, while `masks.nc` contains the mask of each grid point defining if the coupling field value at that point is valid (not masked) or not (masked). The user can construct the auxiliary grid data files before the run, but these files can also be directly created by the component models at the beginning of the run by calling routines `prism_start_grids_writing`, `prism_write_grid`, `prism_write_corner`, `prism_write_mask`, `prism_write_area`, and `prism_terminate_grids_writing` with appropriate arguments containing the grid information.

– Partition definition

The coupling fields of a parallel component model are usually distributed among its different processes. Using the OASIS3 coupling library, each process can send and receive only its part of the coupling field (see paragraph “Coupling field send and receive” below). To do so, each process exchanging coupling data has to describe its local partition in a global index space and has to send this information to the OASIS3 coupler with the `prism_def_partition` routine. This routine has, as input argument, a vector of integers describing the offset and the extent of the local partition in a global index space covering the global grid. Different types of partitions (APPLE, ORANGE, BOX) are supported. Figure 1 illustrates the value of the elements of this vector for very simple APPLE and BOX partitionings.

– Coupling field declaration

In the definition phase, each component process exchanging coupling data has to declare the fields to be sent or received during the simulation by calling `prism_def_var` with, as arguments, the different characteristics of the field (e.g., a symbolic name, field array rank and shape, data type, etc.)

– End of definition phase

Each component process then has to close the definition phase by calling the routine `prism_enddef`. When all components have closed their definition phase, the communication pattern between the component processes and the OASIS3 processes is established. As a result of the information provided by the user in the `namcouple` configuration file (see Sect. 3.2 below), each coupler process knows, for each coupling field it will manage, the symbolic name used in the source component (from which it will receive the field) and the corresponding symbolic name used in the target component (to which it will send the field after regridding). To establish the communication pattern of the coupling exchanges between the component processes and the coupler processes, each process communicates the symbolic name of its coupling fields (declared either with the `prism_def_var` for the component model processes or in the `namcouple` configuration file for the coupler processes) to all other processes. When a match of the symbolic names is found between a component and a coupler process, a communication “channel” is created. This “channel” will later be used to perform an exchange of the corresponding coupling field between the component and the coupler using MPI.

– Coupling field send and receive

The sending and receiving of a coupling field is implemented in the time step loop of the components by respectively calling routines `prism_put` or `prism_get` with, as argument, an array that contains the (part of) the coupling field to be sent or that will store the data received. MPI is used below these routines to perform the communication of the data. The `prism_get` is blocking, i.e., it will return only when the coupling data is effectively received, but the `prism_put` is not blocking so it will return even if the exchange is not completed after possibly MPI buffering of the coupling data.

These routines follow the *end-point exchange* principle in the sense that the target component of a sending call or the source component of a receiving call are not defined in the call. In fact, a source component does not know to which other component the coupling field will go to and a target component does not respectively know where the coupling data comes from. The match between sending and receiving actions is done through the coupler and based on the coupling field source and target symbolic names provided by the

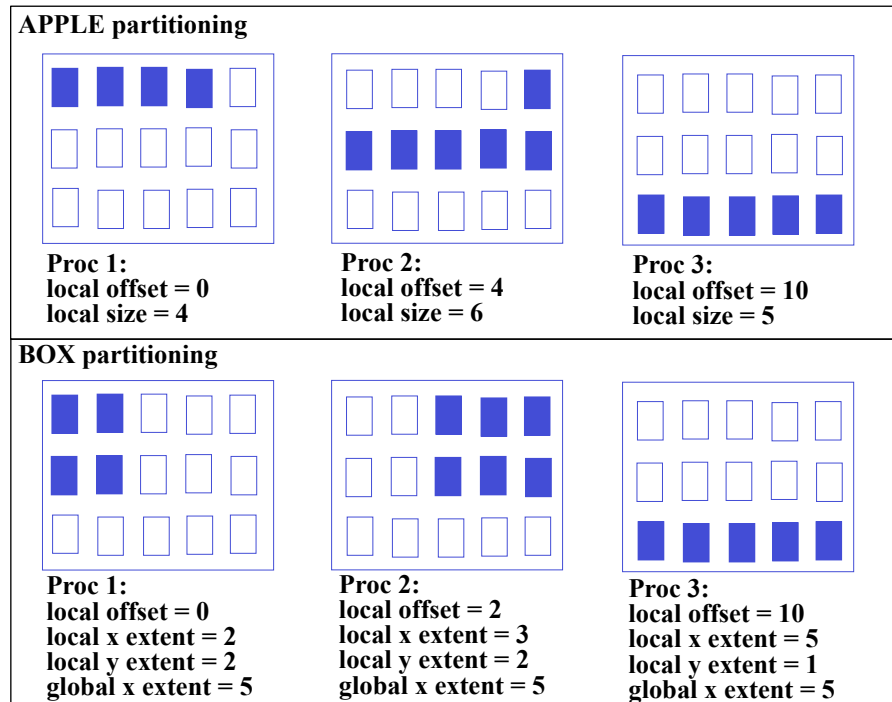


Fig. 1. Very simple APPLE (top) and BOX (bottom) partitionings.

user in the `namcouple` configuration file as explained above.

The user may also decide that the source of a receiving action (`prism_get`) or the target of a sending action (`prism_put`) is a disk file. In this case, the coupling library automatically performs the corresponding I/O action from or to the file indicated by the user in the `namcouple` configuration file, using the GFDL `mpp_io` library (Balaji, 2001), which is interfaced in the coupling library. This functionality implies that the component models can be run in coupled or forced mode transparently.

The sending and receiving routine can be called at each time step anywhere in the component code. The time (in model seconds since the beginning of the current run) at which the call is valid is given as argument and the sending/receiving is actually performed only if the time is equal to an integer number of coupling periods (specified in the `namcouple` by the user, see Sect. 3.2). A change in the coupling period is, therefore, also totally transparent for the component model itself.

If a “lag” is specified for the coupling field by the user in the `namcouple`, a sending action performed by the source model at a particular time will automatically match a receiving action performed by the target model at a later time. In this case, the last sending action of the source model at the end of a run should match the first receiving action of the target model in the next run. OA-

SIS3 automatically supports this transition from one run to the next by writing the last coupling field sent by the source model at the end of a run to a “coupling restart file” that is read in by the coupler at the beginning of the next run to fulfil the first receiving action of the target model.

As the coupler does not provide advanced control on the coupling exchanges, the user has to take few precautions in order to make sure that the coupling exchanges take place as intended. For example, the coupling period defined in the `namcouple` has to be a multiple of the time step duration so that the sending and receiving routines are called at the coupling frequency or more frequently. The user also has to ensure that the coupling algorithm does not lead to a deadlock in the simulation, which would be the case, for example, if two components were both waiting on a `prism_get` for coupling data coming from the other component.

– Coupling termination

All processes of all component models must finalise the coupling exchanges by calling the routine `prism_terminate` after the time step loop. The OASIS3 binary itself will terminate after all component processes have called this routine.

```
#####
# Field 2 #
CONSFTOT SOHEFLDO 6 86400 4 flxat.nc EXPORTED
atmo toce LAG=+14400 SEQ=+2
P 0 P 2
LOCTRANS CHECKIN SCRIPR CHECKOUT
#
ACCUMUL
INT=1
CONSERV LR SCALAR LATLON 10 FRACAREA FIRST
INT=1
#####
```

Fig. 2. Second part of the `namcouple` for one coupling field with source symbolic name `CONSFTOT` and target symbolic name `SOHEFLDO`. On the first line, 6 is an index used to identify corresponding CF standard name and units in auxiliary file, 86400 is the coupling period, 4 is the number of transformations, `flxat.nc` is the name of the coupling restart field, and `EXPORTED` is the type of the coupling field. On the second line, `atmo` and `toce` are the source and target grid acronyms; a lag of 14 400 and a sequence index of 2 are then specified. On the third valid line, the source and target grid characteristics are provided: the source grid is periodic (P) with no overlapping grid point and the target grid is also periodic (P), but with 2 overlapping grid points. On the fourth line, the 4 transformations are listed. Then one additional line is provided for each transformation with some specifications.

3.2 Coupling configuration

The OASIS3 configuration file `namcouple` is a text file that must be written by the user before the run to define, below specific keywords, all information necessary to configure a particular coupled run.

If the OASIS3 coupler binary runs on many processes, the user has to provide one `namcouple` file per coupler process specifying the coupling information relevant to the coupling field(s) that will be treated by each coupler process. This possibility of using many processes for the OASIS3 coupler, each one treating a subset of the coupling fields, is called the OASIS3 “field-per-field” parallelisation.

The first part of `namcouple` is devoted to general coupling parameters such as the number of models involved in the simulation, the number of coupling fields, the MPI mode (MPI1 or MPI2, see paragraph “Coupling initialisation” above). The second part provides specific information on each coupling field. In particular, for each field, it specifies the symbolic name used in the source component and the symbolic name used in the target component; this is how the link between two component models, which do not a priori know of each other, is defined through the coupler. For each coupling field, the user has to specify in the `namcouple` the coupling pe-

riod, the lag, if any, and the list of transformations (including the regridding) to be performed by OASIS3 coupler.

Figure 2 illustrates the second part of the `namcouple` for one coupling field with source symbolic name `CONSFTOT` and target symbolic name `SOHEFLDO`.

3.3 Coupling field transformations and regriddings

For each coupling exchange, the local parts of the source field sent by the different source component processes are gathered on one OASIS3 coupler process that performs the transformations and/or regridding needed to express the coupling field on the grid of the target model. The resulting field is then distributed to the target model processes.

All steps needed for the transformation and regridding are, therefore, done for the whole coupling field by one coupler process. These steps include the neighbourhood search, i.e., the determination of the source points that contribute to the calculation of the regridded value for each target point, the calculation of the corresponding weight and the regridding per se. The first two steps are done only once at the beginning of the run. The weights and addresses are stored in the coupler process memory and then used to perform the regridding of the coupling fields for all exchanges.

3.3.1 Transformations and regriddings offered in OASIS3

In OASIS3, the following transformations on 2-D coupling field in the Earth spherical coordinate system are available.

- Time accumulation or averaging: this operation is performed over all coupling field arrays given as argument to `prism_put` calls performed during the coupling period.
- Correction: external data are read from a file and used to modify the coupling field; this transformation can be used, for example, to perform flux correction on the field.
- Combination: a linear combination of the coupling field is performed with other coupling fields.
- Addition or multiplication by a scalar: this operation can be used, for example, to transform the units of the coupling field.
- Extrapolation: the field is extrapolated over its masked points .
- Interpolation: the different algorithms implemented in the Spherical Coordinate Remapping and Interpolation Package (SCRIP) library (Jones, 1999) are available (see also <http://climate.lanl.gov/Software/SCRIP/> and the SCRIP User Guide at this address):

- *N* nearest-neighbour, possibly Gaussian-weighted: the values of the *N* nearest neighbours on the source grid are weighted by the inverse of their great circle distance to the target grid point. A Gaussian function can also be applied to provide even more weight to the closest source neighbours.
 - Bilinear: the interpolation is based on a local bilinear approximation, which uses the value of the coupling field at the 4 enclosing source grid points.
 - Bicubic: the interpolation is based on a local bicubic approximation and uses the value of the coupling field at the 4 enclosing source grid points, its gradients in each horizontal direction and its cross gradient. It is usually used to interpolate coupling field for which it is important to conserve the higher order property, such as the curl of the wind.
 - 2-D conservative remapping: the contribution of a source cell is proportional to the fraction of the target cell it intersects. This remapping is also taken from the SCRIP library. The area of the intersection is calculated by computing line integrals along the borders and using the divergence theorem.
- In some cases, a target cell may intersect non-masked source grid cells for only a part of its surface (e.g., in the case of non-matching sea-land masks in the ocean and atmosphere models). In these cases, the different types of normalisation available in the SCRIP library will give different results. With the `DESTAREA` normalisation, the whole area of the target cell is used. This ensures local conservation, but the value of the target coupling field may become unrealistic. For example, let us consider the case of a target cell that intersects only one non-masked source cell for one tenth of its surface, the rest intersecting masked source cells and suppose that a solar flux (in $W\ m^{-2}$) is exchanged between these cells. The value of the solar flux coming from the non-masked source cell divided by the ratio of the surfaces (i.e., 10) will be assigned to the target cell: the value of the solar flux will be 10 times smaller than a realistic one, but the energy (in W) received by the 10 times bigger target cell will correspond exactly to the energy leaving the non-masked source cell. With the `FRACAREA` normalisation, only the fraction of the target cell intersected by some source cells is used; local conservation is not ensured anymore, but the values of the target coupling field always remain much closer to the source original values. In OASIS3, we have added the `FRACNNEI` option, which acts as the `FRACAREA` normalisation and in addition attributes the value of the source nearest non-masked neighbour value to target cells that intersect only masked source cells.
- User-defined regridding: the coupler performs the mapping of the source field on the target grid using weights and addresses pre-defined by the user in an external file.
 - Forced global conservation: this performs a global modification of the target coupling field so that conservation is globally ensured. Different options are available:
 - With `GLOBAL`, the field is integrated on both source and target grids, without considering values of masked points, and the residual (target – source) is uniformly distributed on the target grid; this option ensures global conservation of the field.
 - With `GLBPOS`, the same operation is performed except that the residual is distributed proportionally to the value of the original field; this option ensures the global conservation of the field and does not change the sign of the field.
 - With `BASBAL`, the operation is analogous to `GLOBAL` except that the residual is multiplied by the ratio of the total non-masked target surface over the total non-masked source surface; this option ensures that the energy received is proportional to the total non-masked surface of the target grid.
 - With `BASPOS`, the operation is analogous to `GLBPOS` except that the residual is multiplied by the ratio of the total non-masked target surface over the total non-masked source surface; this option ensures that the energy received is proportional to the total non-masked surface of the target grid and does not change the sign of the field.
 - Recreation of subgrid scale variability: this operation can be useful when the source grid has a relatively lower resolution than the target grid. Two types of subgrid interpolation can be performed, depending if the field is of type solar or non-solar.

OASIS3 can also be used off-line in the interpolator-only mode to transform and regrid fields contained in files without running any model. This functionality can be really useful to test off-line the quality of the interpolation for a particular set of source and target grids without having to interface and couple the real components.

3.3.2 Regridding of vector fields

For vector coupling fields, such as wind or ocean current, using a spherical or local coordinate system, regridding the vector components separately as scalar fields will lead to wrong results as their coordinate system is not an absolute reference system.

Therefore, OASIS3 offers the possibility to associate two coupling fields as vector components of one same field, perform if needed a rotation from the local coordinate system

to the spherical coordinate system, project the resulting vector components in a Cartesian coordinate system, regrid the resulting 3 Cartesian components on the target grid, project the result back in the spherical coordinate system, perform if needed a local rotation to project the interpolated vector field on the target local coordinate system.

These steps result in a correct regridding of vector coupling fields.

3.3.3 Grids supported

The transformations listed in Sect. 3.3.1 are available for fields provided on any type of 2-D “logically-rectangular” or unstructured grids, except the bilinear and bicubic interpolations that are available only for logically-rectangular and Gaussian Reduced grids.

Logically-rectangular grids are grids for which the longitudes and the latitudes of the grid points can be described by two arrays `longitude(i, j)` and `latitude(i, j)`, where i and j are respectively the first and second index dimensions. Regular latitude-longitude, stretched or/and rotated grids can be expressed as logically-rectangular grids. Unstructured grids do not have any particular structure. Gaussian Reduced grids are composed of a certain number of latitude circles, each one being divided into a varying number of longitudinal segments (see also http://en.wikipedia.org/wiki/Gaussian_grid).

4 OASIS3 performances

To have an exact measure of the CPU time required by the coupler, we performed measurements of OASIS3 coupling exchanges on the Bullx Curie platform at the “Très Grand Centre de Calcul” (TGCC) in Bruyères-le-Châtel near Paris². We used a toy coupled model composed of two components simulating no dynamics and no physics, but performing realistic “ping-pong” coupling exchanges between a 0.25 degree logically rectangular grid (1021 × 1442 grid points) for the first component and a Gaussian Reduced T799 grid (843 000 grid points) for the second component, for different number of cores. The codes were compiled with Intel Fortran 11 and Bullx MPI 1.1.10.1 (based on OpenMPI) was used.

In a ping-pong exchange, the processes of the first component send a coupling field to OASIS3 that gathers, interpolates and sends it to the processes of the second component that receive the coupling field, and send it back to the first code still through OASIS3 process for the interpolation. The elapse time of the ping-pong exchange is measured in the first component as the difference between the time just before the sending and the time just after the reception. The results are shown at Fig. 3 for 1 to 2048 cores for each code.

Each ping-pong exchange takes about 0.2 s at low number of cores and about 0.3 s for 256 cores and more. The fact that




Fig. 3. Elapse time (in second) on Curie Bullx for one ping-pong exchange with OASIS3 between two toy components, one using a 0.25 degree logically rectangular grid (1021 × 1442 grid points) and the other using a Gaussian Reduced T799 grid (843 000 grid points) for 1 to 2048 cores per component.

the time for an exchange does not decrease with higher number of component cores (i.e., the lack of scalability in OASIS3) is expected as each exchange always goes through the non-parallel OASIS3 single process. The significant increase in the ping-pong time at higher number of cores (0.3 versus 0.2 s) can be explained by the fact that, with components running on a higher number of cores, the single OASIS3 process receives many more different messages from the component processes. For real scalable components, it is expected that the elapse run time spent in the components will decrease with increased number of cores. For a real coupled system, it is therefore most likely that the time spent in the coupler will become proportionally more important when the parallelism of the components will increase. This is what we will call the “bottleneck” effect of OASIS3.

OASIS3 has also been used in few relatively high resolution coupled simulations, which also give indications on the coupler performance. The overhead introduced by the coupling in these simulations depends strongly on the coupling configuration. In many coupled system, the component models running concurrently are not perfectly well load balanced; in this case, at each coupling exchange, the “fast” component waits for the “slow” component and the OASIS3 coupler can perform its work during this “waiting” time as illustrated on the panel A of Fig. 4. Considering the elapse time of the simulation, OASIS3 cost can, therefore, be totally or partially “hidden” by the component unbalance. A more exact measure of the coupling overhead can be done when the component models run sequentially as illustrated on panel B of Fig. 4. In that case, the coupling overhead is exactly the time needed for the communication of the coupling fields and their transformation by the coupler.

The following “high-resolution” coupled simulations correspond to the first case illustrated on Fig. 4a):

² <http://www-hpc.cea.fr/en/complexe/tgcc-curie.htm>

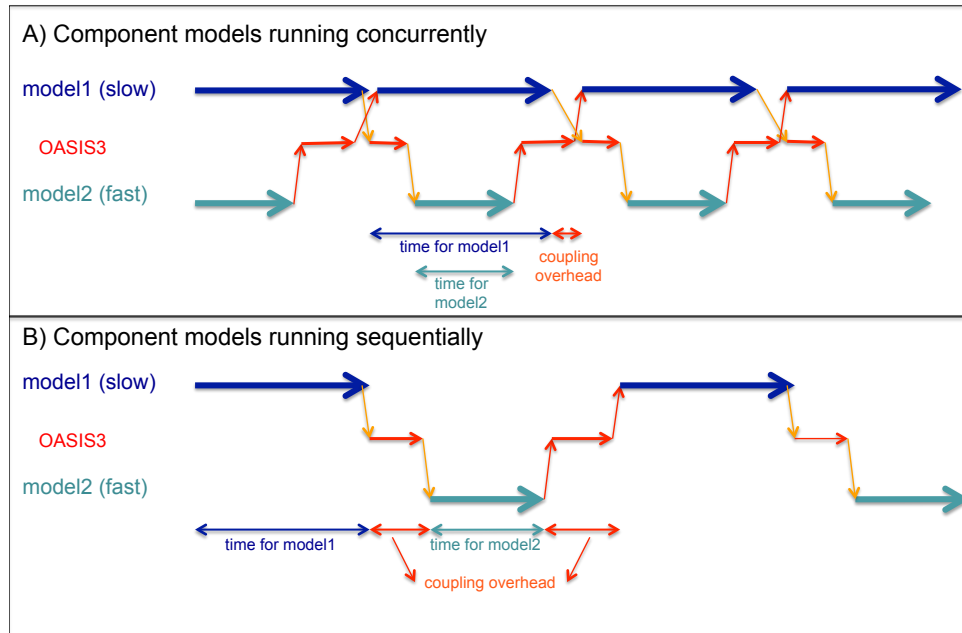


Fig. 4. Different coupling configurations and their impact on coupling overhead. In (A) the component models run concurrently and OASIS3 can perform its work when the fast model is waiting for the slow model; the coupling overhead (in elapse time) is, therefore, much smaller than in (B) where the component models run sequentially and where the coupling overhead is exactly the sum of the communication time and the OASIS3 time for the coupling exchanges in both directions.

- In a high-resolution version of the UK Hadley Centre coupled model³, OASIS3 is used to couple the atmospheric Unified Model (UM), with a horizontal resolution of 432×325 grid points and 85 vertical levels, to the ocean NEMO including the CICE sea ice at a horizontal resolution of 0.25 degree (ORCA025 configuration, $\sim 1\,500\,000$ points) and 75 depth levels. The coupling exchanges are performed every 3 h and the coupled model is run on an IBM Power6 with 192 cores for the UM, 88 cores for NEMO, and 8 cores for OASIS3. In this case, there is less than 2 % overhead in the elapse time of the simulation compared to the UM stand-alone elapse time (R. Hill, personal communication, 2010).
- OASIS3 was also used in a high-resolution version of IPSL Earth System Model to couple the LMDz atmospheric model with 589 000 points horizontally ($\sim 1/3^\circ$) and 39 vertical levels to the NEMO ocean model in the ORCA025 configuration and 75 depth levels on the CINES SGI ALTIX ICE “Jade” (Meurtdesoif et al., 2010). The coupling exchanges were performed every 2 h. The coupled system used up to 2191 cores, with 2048 for LMDz, 120 for NEMO and 23 for OASIS3. Even if the coupling overhead was not measured exactly, no significant slow down was observed during the 10 yr long simulation realised on Jade compared

³<http://www.metoffice.gov.uk/research/modelling-systems/unified-model/climate-models/hadgem3>

to LMDz stand-alone elapse time (A. Caubel, personal communication, 2011).

- In 2010, a high-resolution version of the EC-Earth coupled system (see <http://ecearth.knmi.nl/>) based on the atmospheric model IFS T799 (~ 20 km, 843 000 grid points) with 62 vertical levels and on the NEMO ocean model using the ORCA025 configuration and 45 depth levels was assembled. It was run on the Ekman cluster (1268 nodes of 2 quadripro AMD Opteron2374HE processors, i.e., a total of 10 144 cores⁴) with different numbers of cores for each component and OASIS3. In all configurations, a load unbalance exists between the ocean and atmosphere components, but the cost of the coupler can never be totally hidden. When IFS, NEMO and OASIS3 were run on 800, 256 and 1 core, respectively, the overhead, i.e., the proportional increase of the total elapse time of the simulation with respect to elapse time of the slowest component (IFS), was observed to be 11 %. This overhead decreased to 1.3 % when OASIS3 was run on 10 cores, illustrating the benefits of its field-per-field parallelisation.

In the ARPEGE-NEMIX coupled system developed at CERFACS, the ocean component is a mixed-layer version of the NEMO ocean model and the ocean and atmospheric components run one after the other on different sets of cores;

⁴See also <http://www.pdc.kth.se/resources/computers/ekman>

this configuration corresponds to the second case illustrated in Fig. 4b). It was tested at high resolution on the Bullx Curie machine. NEMIX was run in the ORCA025 configuration with 46 levels vertically and ARPEGE used the T799 grid with 843 000 grid points at each of its 31 vertical levels. 12 coupling fields were exchanged every 3 h through 12 OASIS3 processes. For a relatively low number of processes, the coupling overhead observed was still relatively small. For example, with 64 processes for NEMIX and 48 processes for ARPEGE, the coupling overhead was about 4 %. But when NEMIX and ARPEGE were run on 512 and 496 processes, respectively, the coupling overhead went up to about 20 %. We conclude here that, as suspected, OASIS3 becomes a bottleneck in the simulation for high-resolution models run on a high number of cores because of its limited field-per-field parallelisation. Indeed, as OASIS3 parallelism is exactly limited by the number of coupling fields, the time spent in the coupler becomes relatively more important when the time spent in the components decreases (because of increased parallelism of the components).

5 Use of OASIS3 in CMIP5

Since the first version released in 1993, the OASIS user community has been steadily and regularly increasing from the initial users mainly in France. In Europe, OASIS3 is currently used by CERFACS, Météo-France, the Institut Pierre Simon Laplace (IPSL), the European Centre for Medium range Weather Forecasts (ECMWF) for their operational seasonal prediction suite, the EC-Earth consortium gathering 25 ECMWF member states (that extends ECMWF seasonal forecast system into a real ESM), the Max-Planck Institute for Meteorology (MPI-M) in Germany, the National Centre for Atmospheric Science (NCAS) and the MetOffice in the UK, the Swedish Meteorological and Hydrological Institute (SMHI) in Sweden, the “Koninklijk Nederlands Meteorologisch Instituut” (KNMI) in the Netherlands, the “Centro Euro-Mediterraneo per i Cambiamenti Climatici” (CMCC) in Italy, etc. OASIS3 is also used in the USA (Oregon State University, Hawaii University), in Canada (Environment Canada, Université du Québec à Montréal), in Peru (Instituto Geofísico del Peru), in Japan on the Earth Simulator super computer (the Japan Marine Science and Technology Center), in China (Meteorological National Centre, China Academy of Sciences), in Australia (the Bureau of Meteorology Research Center, the Commonwealth Scientific and Industrial Research Organisation, and the University of Tasmania), etc. The list of all known groups and coupled models that used or are using version 3 of the OASIS coupler, together with the associated computing platform, is given in Appendix A.

In particular, OASIS3 is the coupling software used in different versions of 5 European ESMs participating to CMIP5. The components used in each of these ESMs and the OA-

SIS3 functions activated for each of them are described hereafter in more detail. Two European climate modelling groups are not using OASIS3 in CMIP5, i.e., the Met Office Hadley Centre for their HadCM3 and HadGEM2 models (although OASIS3 is used in their more recent HadGEM3 version, see Sect. 4) and the Norwegian Climate Centre (NCC) in their NorESM1 model.

5.1 CNRM-CM5

CNRM-CM5, assembled by Météo-France and CERFACS, is used in CMIP5 for the decadal and the long-range simulations (Voltaire et al., 2011). CNRM-CM5 is composed of 3 codes: the atmospheric component ARPEGE-Climat 5.1 including the surface module SURFEX, the ocean NEMO V3.2 interfaced with the sea-ice module GELATO and the runoff routing model TRIP. The atmospheric spectral model operates on a T127 triangular truncation with 31 vertical levels and the coupling fields are expressed on a reduced Gaussian grid equivalent to a spatial resolution of about 1.4 degree in both longitude and latitude (with a total of 24 572 grid points). NEMO uses the ORCA1 configuration (362 × 292 grid points horizontally) with 42 levels vertically. TRIP is used for river routing and has a 100 × 100 horizontal resolution.

All coupling fields are exchanged every day between the components. At the beginning of day n , each component receives its input coupling fields sent by the corresponding source component and interpolated by OASIS3 at the end of day $n - 1$. A lag (see Sect. 3.1) corresponding to the time step length of the source component is, therefore, specified for each field.

6 coupling fields are transferred from the ocean to the atmosphere using bilinear regridding for the surface temperature, sea ice extent and albedo fields, and bicubic regridding for the surface current fields, while 16 fields are transferred from the atmosphere to the ocean using bicubic regridding for the wind stress fields and conservative remapping with option `FRACNNEI` for the water, solar and non-solar heat fluxes.

As the coastlines in the ocean and in the atmosphere models do not match, overall global operations are used to force the global conservation with options `GLOBAL`, `GLBPOS`, `BASBAL` or `BASPOS` (see Sect. 3.3.1). For example, the water and solar heat fluxes undergo a global conservation of type `BASPOS` while a global conservation of type `BASBAL` is applied to the non-solar heat flux.

To avoid any drift in the water budget, accumulated snow over Antarctic and Greenland in ARPEGE-SURFEX is artificially discharged in the ocean. The accumulated snow over Antarctic is distributed over ocean grid points South of 60° S whereas the accumulated snow over Greenland is evenly distributed over all ocean grid points. Technically, this is achieved in OASIS3 by defining additional masks for the atmospheric grid (i.e., with only the Antarctic points or only

the Greenland points non-masked) and for the ocean grid (with only the points South of 60° non-masked) and by performing a global conservation of type GLOBAL, therefore resulting in a distribution of the total field on the target non-masked points.

The runoff field modelled by the land scheme (included in the atmosphere model) is transferred to TRIP with a 1-nearest-neighbour regridding, so to avoid smoothing the extrema of the field, followed by a global conservation of type GLBPOS. TRIP then uses this information to calculate the runoff at its discharge coastal points, which is sent to NEMO. In order to remap the runoff appropriately, additional coupling masks were defined in OASIS3 both for the TRIP grid, with only the land discharge coastal point unmasked, and for the NEMO grid with only the ocean points belonging to a narrow band along the coast left unmasked. The runoff is remapped from the TRIP land discharge coastal points to the ORCA1 ocean coastal band with a 6 nearest-neighbour Gaussian-weighted interpolation ensuring an uneven distribution of the runoff with the target points closer to the coast receiving more runoff. A transformation of type GLBPOS is then applied to ensure global conservation (see Maisonave et al., 2008 for details). The time averaging and the multiplication by a scalar (to transform units) are also activated for some coupling fields.

5.2 IPSL-CM5

IPSL-CM5 (Dufresne et al., 2013) is developed by IPSL and includes 5 component models representing the Earth System climate and its carbon cycle: LMDz (atmosphere), NEMO (ocean, oceanic biogeochemistry and sea-ice), ORCHIDEE (continental surfaces and vegetation), INCA (atmospheric tropospheric chemistry) and Reprobus (atmospheric stratospheric chemistry). INCA, Reprobus and ORCHIDEE are directly included in LMDz that is coupled to NEMO through OASIS3. In CMIP5, three different versions of IPSL-CM5 are used that differ by the atmospheric model with two different sets of parameterisation and two horizontal resolution (LMDZ5A with 96 × 95 L39) and (LMDZ5B with 144 × 144 L39) while NEMO is always used in the ORCA2 configuration (149 × 182 grid points horizontally) with 30 levels vertically. The coupling configuration is the same in the three versions.

As for CNRM-CM5, a lag is specified for each field and each component receives, at the beginning of day n , its input coupling fields sent by the corresponding source component and interpolated by OASIS3 at the end of day $n - 1$.

24 coupling fields are exchanged in the coupled system, 7 from the ocean to the atmosphere and 17 from the atmosphere to the ocean. The fields transferred from the ocean to the atmosphere (surface temperature over water and ice, sea-ice extent, surface currents and albedo) are remapped by OASIS3 using a pre-defined set of weights and addresses. This set of weights and addresses was computed using an in-

house conservative algorithm: the source cells and the target cell are described by 8 points (4 corners and 4 middle point on each edge) and, for each target cell, the resulting polygons are projected on a plane passing through the centre of the target cell with a projection that conserves surfaces; the weight of each source cell is then evaluated using a general algorithm calculating the intersection between polygons⁵). Of course, the resulting sets of weights and addresses are different for the low-resolution and the high-resolution versions of the coupled model. OASIS3 is also used to perform some time averaging on the ocean-to-atmosphere fields.

In the atmosphere-to-ocean direction, the 3 Cartesian components of the wind stress and the 10 m wind speed are first extrapolated over land and then interpolated to the U and V ocean grid using an in-house bicubic interpolation method, directly implemented in OASIS3 sources. The water fluxes (precipitation, snow and evaporation) as well as the solar and non-solar heat fluxes are remapped using a set of weights and addresses also pre-defined using the in-house conservative algorithm described in the previous paragraph. For the evaporation, the solar and non-solar fluxes, values of the field averaged over each cell and specific values over ice are transferred.

5.3 CMCC-ESM

Three different coupled models are used at CMCC for CMIP5. They are all based on the OPA8.2 ocean model, using the ORCA2 configuration (182 × 149 grid points) with 31 vertical levels and including LIM for the sea-ice, coupled by OASIS3 to the atmosphere model ECHAM5.

- CMCC-CM (Scoccimarro et al., 2011) is the standard CMCC climate model used for CMIP5 pre-industrial simulations, decadal simulations and centennial projections. In CMCC-CM, ECHAM5 is run with a horizontal triangular truncation T159 (480 × 240 grid points) and 31 vertical levels.
- CMCC-CMS is very close to CMCC-CM except that the atmosphere model runs at higher resolution to resolve the stratosphere with a horizontal triangular truncation T63 with 95 vertical levels.
- In CMCC-ESM (Vichi et al., 2011), a lower resolution is used for ECHAM5, i.e., a horizontal triangular truncation T31 (96 × 48 grid points) with 19 vertical levels, but the processes related to the biological and geochemical parts of the carbon cycle are represented by SILVA for the land and vegetation (interfaced directly in ECHAM5) and by PELAGOS for the ocean biogeochemistry (interfaced directly in OPA).

⁵See also <http://dods.ipsl.jussieu.fr/omamce/IPSLCM5/DocIPSLCM5/MOZAIC/>

The coupling period of all fields is 160 min in CMCC-CM and one day in CMCC-CMS and CMCC-ESM. Besides this difference, the coupling configurations of the three systems are very similar. In both directions, the ocean values of the fields are first extrapolated over land in order to avoid contamination by land values (except for the continental water flux and the total solar and non-solar heat flux provided over the ocean only). In the ocean-to-atmosphere direction, the sea surface temperature is interpolated with a nearest-neighbour interpolation, the sea ice extent with a conservative remapping while the snow thickness, the sea ice thickness and the two vector components of the sea water velocity are interpolated with a bilinear algorithm. In CMCC-CM, OASIS3 also performs a time averaging of the ocean-to-atmosphere coupling fields. In the atmosphere-to-ocean direction, the eastward and northward vector components of the wind stress over the open sea and over the sea ice are interpolated on the ocean U or V grids using a bicubic interpolation. The solar and non-solar heat flux over the open sea and the sea ice are transformed using a bilinear interpolation. The water and snowfall fluxes are also remapped using a bilinear interpolation and global conservations (of types GLOBAL for the water and GLBPOS for the snow) are then applied. Finally the continental water flux and the integral of the total solar and non-solar heat flux over the ocean are remapped using a bicubic algorithm. In addition, for CMCC-ESM only, the 10 m wind speed and the atmospheric CO_2 partial pressure are transferred from the atmosphere to the ocean using a bicubic interpolation.

5.4 EC-Earth V2.3

The EC-Earth model is a state-of-the-art Earth System Model based on ECMWF Seasonal Forecasting System. Different partners in the consortium further develop the baseline model into an Earth System Model used for different climate studies and for CMIP5 in particular. Currently, the EC-Earth consortium consists of 25 academic institutions and meteorological services from 10 countries in Europe. EC-Earth component models are IFS for the atmosphere, including the land and vegetation HTESSEL component, and NEMO for the ocean including the LIM2 sea-ice model.

In EC-Earth V2.3 used for CMIP5 (Hasselger et al., 2011), the IFS version used is cycle 31r1 and runs with horizontal triangular truncation of T159 (i.e., with 35 718 grid points in the horizontal for the reduced Gaussian grid) and 62 vertical levels; NEMO uses the ORCA1 configuration with 362×292 grid points horizontally and 42 vertical levels. 39 coupling fields are exchanged every 3 h, with 30 from the atmosphere to the ocean and 9 from the ocean to the atmosphere. The atmosphere-to-ocean fields are the fraction of water and ice to T , U and V grids, wind stress vector components over water and ice to U and V grids, precipitation-evaporation over water and ice, snow evaporation over ice, solar and non-solar heat flux over water, sensitivity of non-

solar heat flux (needed by LIM2), evaporation flux derivative over water, reference sea temperature for non-solar flux adjustment, net downward surface solar radiation over ice, sea ice surface albedo, net downward surface non-solar heat flux over sea-ice, non-solar heat flux and evaporation derivative over ice, reference sea ice temperature for non-solar flux adjustment, and land surface and drainage runoff. The ocean-to-atmosphere fields are the fractions of water and ice, the sea and ice surface temperatures, the sea-ice albedo and thickness, the snow thickness over the ice and the two vector components of the ocean currents. All coupling fields in both directions are remapped using the SCRIP first-order conservative remapping and no other transformation is performed by OASIS3.

5.5 MPI-ESM

In MPI-ESM (Giorgetta et al., 2012 and Jungclaus et al., 2012) the atmospheric circulation model ECHAM6, including the dynamical land vegetation JSBACH, is coupled via OASIS3 to the ocean and sea ice model MPIOM that also includes the marine biogeochemistry model HAMOCC.

Different versions of MPI-ESM are used for CMIP5: MPI-ESM-MR, MPI-ESM-LR and MPI-ESM-P. In MPI-ESM-MR and MPI-ESM-LR (but not in MPI-ESM-P used for paleoclimatic simulations), a dynamic feedback of vegetation and land use is fully included, land cover change based on data read from an external file is considered, and orbital parameters are calculated at every radiation time step. All 3 versions use a spherical harmonic truncation T63 in ECHAM6. In MPI-ESM-MR, ECHAM6 is run using with 95 vertical levels and the MPI-OM tripolar ocean grid has an approximate horizontal resolution of 0.4 degree and 40 vertical levels. In MPI-ESM-LR and MPI-ESM-P, ECHAM6 uses 47 vertical levels and MPI-OM uses a bi-polar grid with poles located over Greenland and Antarctica with a horizontal resolution of about 1.5 degree, still with 40 vertical levels.

The coupling configurations of the different MPI-ESM versions are very similar. The fields sent from the ocean to the atmosphere are the sea surface temperature, sea ice thickness and fractional area, snow thickness over ice, eastward and northward ocean velocity vector components, the CO_2 transfer coefficient and partial pressure in the ocean upper layer. In this direction, all coupling fields are transferred using first order conservative remapping with FRACAREA normalization. In the other direction, the coupling fields (i.e., the eastward and northward wind stress vector components over water and ice, the snow flux over ice, the water flux in the ocean, the heat flux over water and ice, the residual heat flux, the surface shortwave flux, the wind speed at 10 m, the atmospheric CO_2 concentration, and the ocean-atmosphere CO_2 flux) are first extrapolated over land and then interpolated using first order conservative remapping (again with FRACAREA normalisation) besides the wind stress components in MPI-ESM-MR that use a bicubic interpolation. A

global conservation of type GLOBAL is also imposed on the ocean-atmosphere CO₂ flux and on the water and snow fluxes.

The coupling period is one day for all fields. All fields are exchanged with a lag which means that coupling restart files are used for the first reception of each run and that the fields sent at the last time step of a particular day are received at the first time step of the next day.

6 Discussion and next developments

With the OASIS “multiple binary” approach, the component models remain separate applications and communicate with the coupler through a coupling library offering a relatively simple and flexible API; this ensures a very low degree of intrusiveness in the original codes. Coupling using OASIS will, therefore, not cause any conflict between the original codes, for example, in terms of namespace or I/O. Also, interfacing the component codes with the OASIS3 coupling library can be done in a very generic way (i.e., without specifying, for example, the target of a sending action, the source of a receiving action, the coupling frequency or the coupling transformations) and the configuration of each particular coupled simulation is done by the user in an external text file. Given the size of the current OASIS user community, one can conclude that these original design choices of low-intrusiveness and flexibility were appropriate, especially in the European context of very heterogeneous development environments. It should also be noted that the wide use of the OASIS coupler in climate models naturally emerged as a bottom-up process: OASIS was first developed and used for 2 different French climate models at a time where ocean-atmosphere coupled simulations were just starting and became progressively used by more and more groups. Another important aspect of OASIS success is also the constant user support offered by its developers and the great care taken to constantly integrate community developments in the official version.

However, the OASIS approach suffers from some drawbacks directly linked to the fact that the components remain separate binaries running concurrently on different sets of cores. Indeed, this implies that the coupling exchanges are realised through less efficient MPI communication rather than through the memory within the computing cores. Another disadvantage of this approach is that it is not well suited if the component models are “naturally sequential”, i.e., if one component necessarily waits for some input coupling data when the other is running and vice-versa. This type of coupling would be more efficient if the components were run one after the other on the same set of cores. Finally, a multiple-binary coupled system is in general more difficult to debug for the user and to manage for the operating system.

Because of all these reasons, another approach, the “integrated framework” approach, is followed by some climate groups in Europe, for example, in the EMAC-MPIOM cou-

pled system (Pozzer et al., 2011), but more commonly in the USA, as in the Earth System Modelling Framework (ESMF) (Hill et al., 2004), in the Community Climate System Model (CCSM) (Craig et al., 2012) and in the GFDL Flexible Modelling System (FMS) (Donner et al., 2011). In this approach, the original codes need to be split into initialisation, running and finalisation units, which interfaces need to be standardised, and the resulting subcomponents are integrated into one single application following a coupling algorithm chosen by the user. This approach is, therefore, much more intrusive than the OASIS approach, but offers more opportunities for optimisation as the components can be run concurrently or sequentially on the same set of cores. A review of the different coupling technologies currently used in Earth System Modelling is given in Valcke et al. (2012a). It remains to be seen if, in the longer term for petascale and exascale applications, the more easy-to-use but less efficient multiple-binary OASIS approach will still offer an acceptable coupling solution or if the more intrusive, but more efficient integrated approach will necessarily have to be adopted. In any case, as it is likely that to address efficiently the exascale the model codes will need significant rewriting, the adoption of the initialisation-running-finalisation code structure, recognised as best-practice and favouring the more efficient integrated approach, should be encouraged (Valcke and Dunlap, 2011).

Nevertheless, recent work was done to increase the parallelism of OASIS3. Within the EU FP5 PRISM project (see <http://prism.enes.org/>) and the EU FP7 IS-ENES project (see <https://is.enes.org>), CERFACS, CNRS (Centre National de la Recherche Scientifique, France) and DKRZ (Deutsches Klimarechenzentrum GmbH, Hamburg, Germany) have completely rewritten a new parallel coupler, OASIS4 (Redler et al., 2010). In particular, OASIS4 includes a neighbourhood search library, originally developed by NEC Laboratories Europe – IT Research Division (NLE-IT), performing a fully parallel calculation of the source neighbourhood weights and addresses needed for the regridding of the coupling fields. First versions of OASIS4 have been used by Météo-France, ECMWF, KNMI and MPI-M in the framework of the EU GEMS project (Hollingsworth et al., 2008) for 3-D coupling between atmospheric dynamic and atmospheric chemistry models, and by SMHI, the Alfred Wegener Institute for Polar and Marine Research (AWI in Germany), the BoM in Australia for ocean-atmosphere 2-D regional or global coupling. In the framework of the METAFOR project (Lawrence et al., 2012), OASIS4 was adapted to allow the use of the Common Information Model standard to configure the coupling exchanges. However, performance analyses done during IS-ENES lead to the conclusion that OASIS4 parallel neighbourhood search library presents some fundamental weaknesses in its design. In particular, the support of unstructured grids was not originally included and it would be very difficult to add it in the current code. Also, it is now very clear that the library was developed with efficiency as the prime criteria, leaving aside readability and ease of development.

It was, therefore, decided in July 2011 not to pursue further the development of OASIS4, but to devote the development efforts to the OASIS3-MCT solution.

OASIS3-MCT is an evolution of OASIS3 interfaced with the Model Coupling Toolkit (MCT, see www.mcs.anl.gov/research/projects/mct/) developed by the Argonne National Laboratory in the USA. MCT does not perform the calculation of the regridding source neighbour weights and addresses, but implements fully parallel regridding (as a parallel matrix vector multiplication) and parallel distributed exchanges of the coupling fields, based on pre-computed regridding weights and addresses. Its design philosophy, based on flexibility and minimal invasiveness, is close to the OASIS3 approach. MCT has proven parallel performance and is, most notably, the underlying coupling software used in National Centre for Atmospheric Research Community Earth System Model 1 (NCAR CESM1). Another advantage of the OASIS3-MCT solution is that it will be totally transparent for the OASIS3 user, as the current OASIS3 communication library API provides all the information needed for MCT and will, therefore, not need to evolve. First tests done with up to 4000 cores on the Bullx Curie machine at the TGCC are very encouraging and it is, therefore, very likely that OASIS3-MCT will provide an efficient and easy-to-use solution to remove the foreseen OASIS3 bottleneck. A first official version of OASIS3-MCT was released in August 2012 (Valcke et al., 2012b).

However, in the longer term, we have to prepare for the time when real-time fully parallel calculation of the regridding weights and addresses will become a clear requirement. This functionality will be needed when the component models will run on adaptive grids (which grid point locations change during the run) or when the sequential calculation of the weights and addresses will not be possible anymore because the memory of one core will not be sufficient to hold the definition of the entire source grid. Even if this is some years off for most climate modelling groups, we are currently evaluating on how Open-PALM, another coupler developed at CERFACS (Piacentini et al., 2011) could answer future coupled climate modelling needs. Open-PALM was originally designed to perform the communication and synchronisation of the software components of a data assimilation suite; it, therefore, addresses the particular issue of “dynamic” coupling in the sense that the software components to be coupled can be started and stopped “dynamically” during the run. Open-PALM has proven to be a flexible and powerful dynamic coupler and it is now used by about 40 different groups in France for different multi-disciplinary coupling in different domains, such as aeronautics and space, computational fluid dynamics, combustion but also atmospheric chemistry, hydrology and oceanography. Since January 2011, Open-PALM has been developed in collaboration with ONERA, the French Aerospace Laboratory. In particular, the geometrical interpolation library CWIPI developed at ONERA and based on previous work done at EDF

(Electricité de France) is interfaced in Open-PALM since April 2011. The CWIPI library is designed for finite elements (unstructured) grids in the 3-D space and offers online parallel computation of the weights and addresses for linear interpolations. Open-PALM and its CWIPI library have already shown good performance for up to 12 000 cores (Duchaine et al., 2011), but it is obvious that they do not cover all needs of the climate modelling community at this point. In particular, no conservative remapping or 2nd order interpolation are currently available in CWIPI. Also, it is currently not possible to use a set of weights and addresses pre-calculated off line, which is in some cases essential (for example, when manual input is required to better model the discharge of water runoffs into specific regions of the ocean as a coupling exchange between a river routing model and an ocean model). Evaluation of the work required to adapt Open-PALM to climate modelling requirements and/or possibly merge some OASIS3-MCT functionalities in Open-PALM is, therefore, on-going.

7 Conclusions

The OASIS3 coupler software is widely used in the climate modelling community and, in particular, in five of the seven European Earth System Models participating to CMIP5, as detailed in Sect. 5. To exchange coupling data with the other components of the Earth System, the component models simply need to call few coupling library routines (see Sect. 3.1). The configuration of the coupling exchanges is done externally by the user in a text file (see Sect. 3.2). The coupling exchanges go through the coupler processes that gather the coupling fields and perform their regridding and other transformations (see Sect. 3.3). The coupling overhead introduced by OASIS3 in most of current coupled systems is very reasonable. However, it becomes non-negligible in high-resolution models run on a high number of cores (see Sect. 4). Therefore, within the framework of funded projects such as the EU FP7 IS-ENES project and its follow-on IS-ENES2, recently accepted for the 2013–2016 period, work continues to increase the parallelism of the coupler with the new OASIS3-MCT version, which first version was officially released in August 2012 (see Sect. 6). In the coming years, active user support will also continue through the ENES portal offering source code, documentation, user guides, tutorial, FAQs, and a user forum (see <http://oasis.enes.org>). In the longer term, to prepare for the time when real-time fully parallel calculation of the regridding weights and addresses will be required, we plan to assess the coupler currently developed by CERFACS and ONERA, Open-PALM and its new parallel real-time interpolation library CWIPI and to explore alternative existing coupling technologies.

Up to now, CERFACS, devoting one person full time to OASIS, has been able to provide the services needed to maintain a strong community network around the coupler:

Table A1. Use of OASIS3: centres, coupled models and computing platforms.

Centre	Coupled model	Platform
CERFACS (FR)	ARPEGE4.T63/NEMO-ORCA2-LIM/TRIP	Linux Cluster, CRAY XD1, VPP5000, NEC SX6-SX8
CERFACS (FR)	ARPEGE.T63/NEMIX-ORCA2	CRAY XD1
CERFACS (FR)	ARPEGE.T359/NEMO-ORCA012	NEC SX9
CERFACS (FR)	ARPEGE.T799/NEMIX-ORCA025	Bullx, Altix ICE
Météo-France (FR)	ARPEGE.V4.6/NEMO-ORCA2/NEMOmed8	NEC-SX8
Météo-France (FR)	ALADIN-Climat/NEMOmed8/TRIP	NEC-SX9
Météo-France (FR)	ARPEGE.V5/NEMOV3-ORCA1/TRIP	NEC-SX9
Météo-France (FR)	ARPEGE.V5.1/NEMO1	IBM Power 6
IPSL (FR)	LMDZ/NEMO-ORCA2/ – LMDZmed/NEMOmed8	
IPSL (FR)	LMDz(144 × 142)/NEMO-ORCA2	NEC SX8
OMP (FR)	MESO-NH/SYMPHONIE	Linux Opteron cluster
LGGE (FR)	MAR/NEMO-LIM	
ECMWF	IFS.T399/NEMO-ORCA1	IBM Power 6
MPI-M (DE)	ECHAM5/MPIOM	IBM Power 6, SUN Linux
MPI-M (DE)	REMO/MPIOM	IBM Power 575
Met Office (UK)	UM Atm(192 × 145)/NEMO-ORCA1	IBM Power6, NEC SX6
Met Office (UK)	UM Atm(432 × 325)/NEMO-ORCA025	IBM Power6
NCAS/Reading (UK)	ECHAM4/NEMO-ORCA2	NEC SX6-SX8
NCAS/Reading (UK)	HadAM3/NEMO-ORCA2	NEC SX8
IFM-GEOMAR (DE)	ECHAM5.T63/NEMO-ORCA2	NEC SX6, SX8, SX9
CMCC (IT)	ECHAM5.T31L31/OPA8.2-ORCA2	NEC SX9
CMCC (IT)	ECHAM5.T159L31/OPA8.2-ORCA2	IBM Power6
CMCC (IT)	ECHAM5.T63L95/OPA8.2-ORCA2	
CMCC (IT)	ECHAM5.T159/OPA8.2-ORCA2/ – NEMOMed.1/16	
ENEA (IT)	RegCM/MITgcm	IBM-SP5
LMNCP (IT)	WRF/ROMS	
SMHI (SE)	EC-Earth: IFS.T159/NEMO-ORCA1	Linux Cluster
SMHI (SE)	EC-Earth: IFS.T799/NEMO-ORCA025	Ekman AMD Opteron
KNMI (NL)	ECHAM5/MPIOM	NEC SX-8
KNMI (NL)	EC-Earth: IFS.T159/NEMO-ORCA2	SGI Altix, IBM Power
DMI (DK)	ECHAM(global)/HIRLAM (reg)	NEC SX6
U.Bergen (NO)	MM5/ROMS	
ICHEC (IE)	EC-Earth: IFS.T159/NEMO-ORCA1	SGI Altix ICE, Bull clust.
ICHEC (IE)	ROMS/WRF	
NUI Galway (IE)	ECHAM5/REMO/MPI-OM	
ETH (CH)	COSMO-CLM/CLM	
ETH (CH)	COSMO-CLM/ROMS	
U. Castille (ES)	PROMES/U. Madrid ocean	IBM Power 6 – Intel Itanium
NHM Service (RS)	ECHAM5/MPIOM	Linux Fedora core 14
CMC (CA)	GEM/NEMO	IBM Power5
UQAM (CA)	GEMDM 3.3.2/NEMOv.2.3.0	Linux PC
MM (MA)	ARPEGE-Climat.V5.1/NEMO-ORCA1	IBM
INMT (TN)	ARPEGE-Climat.V5.1/NEMO-ORCA2	IBM Regata Series
Oregon St U (USA)	PUMA/UVic	Linux cluster
Hawaii U (USA)	ECHAM4/POP	Linux cluster
JAMSTEC (JP)	ECHAM/OPA8.2	NEC SX8
JAMSTEC (JP)	ECHAM5.T106L31/NEMO.0.5	NEC SX8
Met.Nat.Center (CN)	GRAPES (201 × 161)/ECOM-si	IBM cluster
IAP (CN)	CREM(reg.)/POM2000	SGI
IAP (CN)	ECHAM/MPIOM	Linux_x64
CSIRO (AU)	ACCESS : UMv7.3/MOM4p1/CICE	SUN + SGI clusters
BoM (AU)	BAM3/ACOM2	NEC SX6, SUN
BoM (AU)	TCLAPS/MOM	NEC SX-6
U Tasmania (AU)	Data atm. model/MOM4	SGI O3400 – Compaq5

development, maintenance, integration, user support, etc. This has been possible thanks to numerous collaborations on specific developments and to temporary, but important funding streams (such as the METAFOR and IS-ENES EU projects). The investment of the French CNRS, also devoting one engineer full time in OASIS, is in that respect particularly valuable. However, a concrete involvement of the whole community in terms of funding and governance is now needed given the foreseen jump in complexity of the coupling problem on massively parallel platforms. In the framework of the recently funded IS-ENES2 project, it is planned to set the basis of an efficient community development process including planning, prototyping, implementation, testing, and quality assurance. These are essential aspects of this important shared piece of software that the coupler represents for the European climate modelling community.

OASIS represents about 35 person-year of work and is used by about 35 modelling groups around the world. The average of 1.0 person-year/group is certainly much less than the time it would have taken for each group to develop its own coupler. Therefore, OASIS is and will certainly remain a great example of successful community software for years to come.

Supplementary material related to this article is available online at: <http://www.geosci-model-dev.net/6/373/2013/gmd-6-373-2013-supplement.pdf>.

Acknowledgements. OASIS3 development, maintenance and support has been in the last 10 years partly funded by PRISM (EU 6th Framework Programme project no. EVR1-CT2001-40012), CICLE (French Agence National de la Recherche 2006–2009 project no. ANR-05-CIGC-04), METAFOR (EU 7th Framework Programme e-infrastructure project no. 211753) and IS-ENES (EU 7th Framework Programme e-infrastructure project no. 228203). We would also like to warmly thank the numerous past, present, local, national and international collaborators, especially Laure Coquart and Eric Maisonnave, colleagues at CERFACS, who made possible the successful development and use of the OASIS coupler. We also would like to thank the referees of this article, especially Mick Carter, who really helped improve its quality. The publication of this article is financed by CNRS-INSU.

Edited by: S. Easterbrook



The publication of this article is financed by CNRS-INSU.

References

- Balaji, V.: Parallel Numerical Kernels for Climate Models, in: ECMWF TeraComputing Workshop 2001, ECMWF, World Scientific Press, 2001.
- Cassou, C., Noyret, P., Sevault, E., Thuau, O., Terray, L., Beaucourt, D., and Imbard, M.: Distributed Ocean-Atmosphere Modelling and Sensitivity to the Coupling Flux Precision: the CATHODE Project, *Mon. Weather Rev.*, 126, 1035–1053, 1998.
- Craig, A. P., Vertenstein, M., and Jacob, R.: A New Flexible Coupler for Earth System Modeling developed for CCSM4 and CESM1, *Int. J. High Perform. C.*, 26–1, 31–42, doi:10.1177/1094342011428141, 2012.
- Donner, L. J., Wyman, B. L., Hemler, R. S., Horowitz, L. W., Ming, Y., Zhao, M., Golaz, J. C., Ginoux, P., Lin, S. J., Schwarzkopf, M. D., Austin, J., Alaka, G., Cooke, W. F., Delworth, T. L., Freidenreich, S. M., Gordon, C. T., Griffies, S. M., Held, I. M., Hurlin, W. J., Klein, S. A., Knutson, T. R., Langenhorst, A. R., Lee, H.-C., Lin, Y., Magi, B. I., Malyshev, S. L., Milly, P. C. D., Naik, V. and Nath, M. J., Pincus, R., Ploshay, J. J., Ramaswamy, V., Seman, C. J., Shevliakova, E., Sirutis, J. J., Stern, W. F., Stouffer, R. J., Wilson, R. J., Winton, M., Wittenberg, A. T., and Zeng, F.: The dynamical core, physical parameterizations, and basic simulation characteristics of the atmospheric component AM3 of the GFDL Global Coupled Model CM3, *J. Climate*, 24, 3484–3519, 2011.
- Duchaine, F., Morel, T., and Piacentini, A.: On a first use of CWIPI at CERFACS, CERFACS Technical Report TR/CMGC/11/3, Toulouse, France, 15 pp., 2011.
- Dufresne, J.-L., Foujols, M.-A., Denvil, S., Caubel, A., Marti, O., Aumont, O., Balkanski, Y., Bekki, S., Bellenger, H., Benshila, R., Bony, S., Bopp, L., Braconnot, P., Brockmann, P., Cadule, P., Cheruy, F., Codron, F., Cozic, A., Cugnet, D., de Noblet, N., Duvel, J.-P., Ethé, C., Fairhead, L., Fichet, T., Flavoni, S., Friedlingstein, P., Grandpeix, J.-Y., Guez, L., Guilyardi, E., Hauglustaine, D., Hourdin, F., Idelkadi, A., Ghattas, J., Jous-saume, S., Kageyama, M., Krinner, G., Labetoulle, S., Lahellec, A., Lefebvre, M.-P., Lefevre, F., Levy, C., Li, Z. X., Lloyd, J., Lott, F., Madec, G., Mancip, M., Marchand, M., Masson, S., Meurdesoif, Y., Mignot, J., Musat, I., Parouty, S., Polcher, J., Rio, C., Schulz, M., Swingedouw, D., Szopa, S., Talandier, C., Terray, P., and Viovy, N.: Climate change projections using the IPSL-CM5 Earth System Model: from CMIP3 to CMIP5, *Clim. Dynam.*, accepted, 2013.
- Giorgetta, M. A., Jungclaus, J. H., Reick, C. H., Legutke, S., Brovkin, V., Crueger, T., Esch, M., Fieg, K., Glushak, K., Gayler, V., Haak, H., Hollweg, H.-D., Kinne, S., Kornblueh, L., Matei, D., Mauritsen, T., Mikolajewicz, U., Mueller, W., Notz, D., Raddatz, T., Rast, S., Roeckner, E., Schmidt, H., Schnur, R., Segschneider, J., Six, K. D., Martina Stockhause, M., Wegner, J., Widmann, H., Wieners, K.-H., Claussen, M., Marotzke, J., and Stevens, B.: Climate change from 1850 to 2100 in MPI-ESM simulations for the Coupled Model Intercomparison Project 5, *J. Adv. Model. Earth Syst.*, in review, 2012.
- Gropp, W., Huss-Lederman, S., Lumsdaine, A., Lusk, E., Nitzberg, B., Saphir, W., and Snir, M.: MPI – The Complete Reference, 2, The MPI Extensions, MIT Press, 1998.
- Hazeleger, W., Wang, X., Severijns, C., Stefanescu, S., Bintanja, R., Sterl, A., Wyser, K., Semmler, T., Yang, S., van den Hurk, B., van Noije, T., van der Linden, E., and van der Wiel, K.:

- EC-Earth V2.2: description and validation of a new seamless earth system prediction model, *Clim. Dynam.*, 39, 2611–2629, doi:10.1007/s00382-011-1228-5, 2011.
- Hill, C., DeLuca, C., Balaji, V., Suarez, M., and da Silva, A.: Architecture of the Earth System Modeling Framework, *Comput. Sci. Eng.*, 6, 18–28, 2004.
- Hollingsworth, A., Engelen, R. J., Textor, C., Benedetti, A., Boucher, O., Chevallier, F., Dethof, A., Elbern, H., Eskes, H., Flemming, J., Granier, C., Kaiser, J. W., Morcrette, J.-J., Rayner, P., Peuch, V. H., Rouil, L., Schultz, M. G., Simmons, A. J., and The GEMS Consortium: Toward a Monitoring and Forecasting System For Atmospheric Composition: The GEMS Project, *B. Am. Meteorol. Soc.*, 89, 1147–1164, 2008.
- Jones, P.: Conservative remapping: First- and second-order conservative remapping, *Mon. Weather Rev.*, 127, 2204–2210, 1999.
- Jungclaus, J. H., Fischer, N., Haak, H., Lohmann, K., Marotzke, J., Matei, D., Mikolajewicz, U., Notz, D. and von Storch, J. S.: Characteristics of the ocean simulations in MPIOM, the ocean component of the MPI-Earth System Model, *J. Adv. Model. Earth Syst.*, accepted, 2012.
- Lawrence, B. N., Balaji, V., Bentley, P., Callaghan, S., DeLuca, C., Denvil, S., Devine, G., Elkington, M., Ford, R. W., Guilyardi, E., Lautenschlager, M., Morgan, M., Moine, M.-P., Murphy, S., Pascoe, C., Ramthun, H., Slavin, P., Steenman-Clark, L., Tossaint, F., Treshansky, A., and Valcke, S.: Describing Earth system simulations with the Metafor CIM, *Geosci. Model Dev.*, 5, 1493–1500, doi:10.5194/gmd-5-1493-2012, 2012.
- Maisonnave, E. and Terray, L.: Runoff: couplage conservatif par anomalie, CERFACS Technical Report TR/CMGC/08/71, Toulouse, France, 29 pp., available at: http://www.cerfacs.fr/globc/publication/technicalreport/2008/TR_runoffs.pdf, 2008.
- Manabe, S. and Bryan, K.: Climate Calculations with a Combined Ocean-Atmosphere Model, *J. Atmos. Sci.*, 26, 786–789, 1969.
- Meurtdesoif, Y., Caubel, A., and Masson, S.: Modélisation couplée océan-atmosphère à très haute résolution, *La Gazette du CINES, Grands Challenges GENCI-CINES 2010*, 22–25, available at: http://www.cines.fr/IMG/pdf/Recueil_definitif_des_articles_V10.pdf, 2012.
- Piacentini, A., Morel, T., Thévenin, A., and Duchaine, F.: O-PALM : An Open Source Dynamic Parallel Coupler, in: *Proceedings of the IV International Conference on Computational Methods for Coupled Problems in Science and Engineering – Coupled Problems 2011*, Kos Island, Greece, 2011.
- Pozzer, A., Jöckel, P., Kern, B., and Haak, H.: The Atmosphere-Ocean General Circulation Model EMAC-MPIOM, *Geosci. Model Dev.*, 4, 771–784, doi:10.5194/gmd-4-771-2011, 2011.
- Redler, R., Valcke, S., and Ritzdorf, H.: OASIS4 – a coupling software for next generation earth system modelling, *Geosci. Model Dev.*, 3, 87–104, doi:10.5194/gmd-3-87-2010, 2010.
- Scoccimarro, E., Gualdi, S., Bellucci, A., Sanna, A., Fogli, P. G., Manzini, E., Vichi, M., Oddo, P., and Navarra, A.: Effects of Tropical Cyclones on Ocean Heat Transport in a High Resolution Coupled General Circulation Model, *J. Climate*, 24, 4368–4384, 2011.
- Snir, M., Otto, S., Huss-Lederman, S., Walker, D., and Dongarra, J.: *MPI – The Complete Reference*, vol. 1, The MPI Core, MIT Press, 1998.
- Terray, L., Thual, O., Belamari, S., Déqué, M., Dandin, P., Lévy, C. and Delecluse, P.: Climatology and interannual variability simulated by the arpege-opa model, *Clim. Dynam.*, 11, 487–505, 1995.
- Valcke, S. and Dunlap, R.: Workshop Proceedings: Coupling Technologies for Earth System Modelling: Today and Tomorrow, CERFACS Technical Report TR/CMGC/11/39, Toulouse, France, 68 pp., available at: http://www.cerfacs.fr/globc/publication/technicalreport/2011/Proceedings_of_the_workshop_final.pdf, 2011.
- Valcke, S., Guilyardi, E., and Larsson, C.: PRISM and ENES: A European approach to Earth system modelling, *Concurr. Comp.-Pract. E*, 18, 231–245, 2006.
- Valcke, S., Balaji, V., Craig, A., DeLuca, C., Dunlap, R., Ford, R. W., Jacob, R., Larson, J., O’Kuinghtons, R., Riley, G. D., and Vertenstein, M.: Coupling technologies for Earth System Modelling, *Geosci. Model Dev.*, 5, 1589–1596, doi:10.5194/gmd-5-1589-2012, 2012.
- Valcke, S., Craig, T., and Coquart, L.: OASIS3-MCT User Guide, OASIS3-MCT 1.0. CERFACS Technical Report, TR/CMGC/12/49, Toulouse, France, 46 pp., available at: http://pantar.cerfacs.fr/globc/publication/technicalreport/2012/oasis3mct_UserGuide.pdf, 2012.
- Vichi, M., Manzini, E., Fogli, P. G., Alessandri, A., Patara, L., Scoccimarro, E., Masina S., and Navarra, A.: Global and regional ocean carbon uptake and climate change: sensitivity to a substantial mitigation scenario, *Clim. Dynam.*, 37, 1929–1947, doi:10.1007/s00382-011-1079-0, 2011.
- Voltaire, A., Sanchez-Gomez, E., Salas y Mélia, D., Decharme, B., Cassou, C., Sénési, S., Valcke, S., Beau, I., Alias, A., Chevallier, M., Déqué, M., Deshayes, J., Douville, H., Fernandez, E., Madec, G., Maisonnave, E., Moine, M.-P., Planton, S., Saint-Martin, D., Szopa, S., Tyteca, S., Alkama, R., Belamari, S., Braun, A., Coquart, L., and Chauvin, F.: The CNRM-CM5.1 global climate model: Description and basic evaluation, *Clim. Dynam.*, Special Issue, doi:10.1007/s00382-011-1259-y, 2011.