Proceedings of the
OASIS User meeting 2009

CERFACS, Toulouse, France, May 25-26, 2009

organised in the framework of the
IS-ENES EU funded project

assembled by Sophie Valcke
CERFACS Technical Report TR/CMGC/09-89

Table of contents

Introduction

Coupled General Circulation Models (CGCMs) of the sea-ice-ocean-atmosphere-biosphere system are the most comprehensive tools used today to study the interactions between the different components of the global climate system, understand its natural variability and predict its evolution. Great care must be given to the computing and technical aspects of the complex and CPU-intensive simulations based on such CGCMs.

In 1991, CERFACS decided to develop a software interface to couple existing numerical General Circulation Models of the ocean and the atmosphere. Today, both the widely used OASIS3 version, which is the result of more than 15 years of evolution, and the newer fully parallel OASIS4 version, which writing started during the PRISM EU project and which is presently developed up to now thanks to an active collaboration between NEC Laboratories Europe IT Research Division in Germany, the Centre National de la Recherche Scientifique (CNRS) and CERFACS, are available.

The objective of the workshop was to gather scientists and engineers and offer them a unique opportunity to communicate on their recent results in coupled climate modelling, to detail their difficulties and successes in setting up their coupled system, and to share their view for further development of the OASIS coupler. Strategy on how to prepare climate models and the OASIS coupler in particular for the next generation of high performance computer architecture were also addressed.

This meeting was organized in the framework of the IS-ENES project but everyone, involved or not in IS-ENES, was welcome to participate.

The workshop was very successful thanks to the active participation of all people present. The presentations and discussions that took place during these 2 days will now help the OASIS team to identify the preferred short-, medium- and long-term developments for OASIS, quantify the work associated, and define precisely the OASIS development strategy, in particular during the IS-ENES project. Discussions on OASIS future developments are summarized at the end of section 1.

Different presentations were done illustrating the wide use of the OASIS coupler in the climate modelling community. A summary of each presentation can be found in the next sections. The slides of the presentations can be downloaded from following web page: http://www.cerfacs.fr/~coquart/pagecerfacs/projet_isenes/OASIS_meeting_2009 .

1. The OASIS coupler: history, community, current status and plans for future developments

S. Valcke (CERFACS)

After a brief overview of the different techniques available to assemble existing model codes, we present here the current status of the OASIS coupler widely used in the climate modeling community.

1.1 Different technical approaches to assemble component model codes

Coupling component model codes means managing synchronised exchange of information ("coupling fields") between those codes, and transforming the coupling information provided by one code to ensure that it can be "ingested" by the other code. Ideally, the technical solution chosen should be easy to implement with existing codes, flexible, efficient and portable. Different technical approaches exist to assemble component model codes.

i. Merging of the codes

The most natural approach is to merge the existing codes into one new application, which means that one code remains a main program and calls the other code as a subroutine. The coupling information can be exchanged by argument passing or by sharing a common module. This approach ensures efficient memory exchanges and portability (in as much as the original codes are themselves portable). However this approach is not flexible, as the coupling algorithm and the coupling exchanges must be hard-coded while merging the codes, and supposes that the user implements and uses in the source or in the target code his own transformations and interpolations. Other disadvantages of this approach are that many conflicts in namespaces, I/O, etc. are likely to appear and that the memory requested by the resulting code may be very large, depending of course on how the original codes are programmed. This first approach is in general not recommended.

ii. Direct use of existing communication protocols

The second approach is to keep the original component models separate but implement the coupling exchanges directly where needed in the codes using an existing protocol such as MPI, CORBA, Unix pipes, or files. Compared to the first one, the advantage of this approach is that no conflict will appear but it is not either flexible as the coupling exchanges, specific to each coupling configuration, will also be hard-coded in the codes. It also requires that the scientist masters the communication protocol and implements his own transformations and interpolations. Finally, its portability depends on the portability of the chosen communication protocol.

iii. Use of a coupling framework

The third solution is to use a coupling framework, such as ESMF (http://www.esmf.ucar.edu) or FMS (http://www.gfdl.noaa.gov/ fms/). This approach supposes that the user splits the original codes into elemental units (at least in initialisation, main, and termination units), adapts the units to the framework standard data structure and calling interface, and finally uses the framework to build and control a hierarchical merged application integrating the different units. This approach is fully flexible (the different units can be easily reused in different applications), allows the user to use the different tools offered by the framework (parallelisation, regridding, time management, etc.) and ensure efficient coupling exchanges within the merged application. But it requires a deeper level of interference in the codes and imposes strict coding rules in order to take full advantage of the framework functionalities. This approach is therefore probably the most recommended one in a controlled top-down development environment.

iv. Use of a coupler

In many cases though, the different component models chosen to form an ESM come from different research groups that also use these components independently in stand-alone mode for other research purposes and that are not likely to follow strict coding rules imposed by external constraints. In this case, a less intrusive approach based on the use of a coupler and associated coupling library, such as MpCCI (http://www.mpcci.de), MCT (http://www.mcs.anl.gov/mct), PALM (http://www.cerfacs.fr/globc/PALM WEB/index.html) or OASIS (https://oasistrac.cerfacs.fr/) is probably the best trade-off that can be chosen. In particular, it ensures that the original codes will run as separate executables with main characteristics (e.g. internal parallelisation) unchanged with respect to the uncoupled mode. The drawback is that the execution of the resulting coupled model may in some cases be less efficient than a more integrated one-executable approach. This approach is flexible as the coupling exchanges generally follow the principle of “end-point” data exchange (see section 1.3 below). The portability of the coupling depends on the portability of the coupler, criteria usually of great importance for the coupler development teams. This approach also allows the user to take advantage of the different transformation and regridding routines offered with the coupler.

1.2 The OASIS coupler: historic and community

In 1991, CERFACS decided to develop a software interface to couple existing numerical General Circulation Models of the ocean and the atmosphere. Today, both the widely used OASIS3 version, which is the result of more than 15 years of evolution, and the newer fully parallel OASIS4 version, which writing started during the PRISM EU project and which is presently developed thanks to an active collaboration between NEC Laboratories Europe IT Research Division in Germany, the Centre National de la Recherche Scientifique (CNRS) and CERFACS, are available.

The OASIS community has steadily grown since its first release. The OASIS3 version is currently used by about 30 modelling groups in Europe, Australia, Asia and North America, on the different computing platforms used by the climate modelling community. A beta release of the newer fully parallel OASIS4 version is currently used in the framework of the GEMS project (lead by ECMWF) by Météo-France, KNMI (Netherlands), and MPI-M (Germany) for 3D coupling between atmospheric dynamic and atmospheric chemistry models, by IPSL in the framework of the ANR-funded CICLE project (ANR-05-CIGC-04), and by SMHI (Sweden) and by BoM (Australia) for 2D ocean-atmosphere coupling.

1.3 Data exchanges with OASIS

To exchange coupling information with other components, a component model needs to call few specific routines of the OASIS coupling library for its initialisation, grid and partition definition, field declaration, field Get and Put actions (to receive or send a field by respectively) and termination. The main difference between OASIS3 and OASIS4 coupling library API (Application Programming Interface) is the grid definition. For OASIS3, this is done for the global grid either in an external file or by the master process only. With OASIS4, each process has to describe its local part of the grid by giving the localisation (longitude and latitude) of the local grid corners and points. Besides this difference, OASIS3 and OASIS4 API were kept as close as possible to facilitate the migration from one version to the other.

In OASIS3 and OASIS4, the communication follows the “end-point” principle, i.e. there is no reference in the component model code to the origin of a Get action or to the destination of a Put action; the source and target component models (coupling exchange) or the input or output file (I/O) are set externally by the user. This ensures an easy transition between different coupling configurations, in particular from the coupled mode (Get/Put actions leading to a coupling exchange performed using MPI) to the forced mode (Get/Put action leading to reading/writing from/to a file using the GFDL mpp io library, Balaji 2001), totally transparent for the component model itself. Furthermore, the Get/Put routines can be called at each time step in the component model code but the receiving/sending actions will effectively be performed only at appropriate times from/to the appropriate source/target following the configuration externally defined by the user.

Internally, the data exchanges, including the regridding, are realized very differently between OASIS3 and OASIS4. With OASIS3, each couplings field is gathered from the source model processes to the OASIS3 central process which performs: 1) the neighbourhood search, i.e. the identification for each target point of the source grid points that will participate to the calculation of its regridded value, 2) the regridding per se and 3) the sending of the resulting field to the target processes. With the last pseudo-parallel version of OASIS3, it is now possible to include more than one regridding central process in a coupled model with each regridding process treating a subset of the coupling field; this results in a pseudo-parallelisation version of OASIS3 on a field-per-field basis.

With OASIS4, the neighbourhood search is done thanks to an efficient multigrid algorithm by the source communication library for each pair of source and target process if an intersection between the source and target local domain is found. For each exchange, only the useful source points (i.e. the one participating to the regridding) are effectively sent by each source process to the central parallel Transformer which calculates the regridded values and send them to the relevant target process. The neighbourhood search and the regridding are therefore performed in a fully parallel and efficient mode for each intersection of source and target process domain.

1.4 Regridding algorithms available in OASIS

The following 2D regridding algorithms based on the SCRIP library (Jones 1999) are available in OASIS3 and OASIS4, with 3D extensions in OASIS4 only (for more details, the reader is referred to the SCRIP User Guide available at http://climate.lanl.gov/Software/SCRIP/):

- N nearest-neighbour: the N closest source neighbours are used. The weight of each neighbour is inversely proportional to d, its Great Circle distance to the target point, or to exp(−1/2 · d2/s2) where s2 is the variance of a Gaussian function.
- Bilinear: the 4 enclosing source neighbour points are used and their respective weight is evaluated using a general bilinear iteration in a continuous local coordinate system.
- Bicubic: the value of the source field, its gradients and cross gradient with respect to the local directions I and j at the 4 enclosing source neighbour points are used. For Reduced Gaussian grid, a standard bicubic algorithm with the 16 enclosing source neighbours is used.
- 1st order conservative remapping: the weight of a source cell is proportional to area of the source cell intersected by target cell. Using the divergence theorem, the SCRIP library evaluates this area with the line integral along the cell borders enclosing the area.

1.5 Plans for future developments

Due to economical reasons, NLE-IT had to withdraw in February 2009 from its collaboration with CERFACS and CNRS on the development of OASIS. Within IS-ENES however, collaborative development of OASIS is going on between CERFACS, CNRS (devoting one engineer full-time to this task) and the DKRZ from Hamburg, Germany. In the coming 4 years, IS-ENES will fund 63 additional persons-months at CERFACS and 35 persons-months at DKRZ for OASIS development.

The main objectives for the coming year regarding OASIS3 is to officially release the OASIS3_3 pseudo-parallel version (see section 1.3) including also CMCC parallelisation as an option (see section 13). This new version should include some optimisation of the “scriprmp” and “extrap” routines. A need for alternative conservative remapping algorithms, more precise than the SCRIP especially near the pole, was also clearly during the meeting. Finally, the new version should be released with the MetOffice FCM build system. But besides those few and well identified developments, the resources should mainly be devoted to user support which can be time consuming considering the large community of OASIS3 users.

Most of the developments should in fact address the OASIS4 coupler. Considering the User meeting discussions, the following ones should have high priority:

- fully validate the current OASIS4 2D transformations, including the parallel 2D conservative remapping, and ensure OASIS3 reproducibility
- test and optimise the coupler into real coupled systems on HPC platforms
- develop automatic test suites
- develop a Graphical User Interface (GUI) to help the user building the OASIS4 XML configuration files and simplify this configuration
- support vector fields
- support user-defined weights and addresses for a customised regridding
- offer forced global conservation
- allow more than one components to run sequentially within one executable
- support unstructured grids (in collaboration with the Alfred Wegener Institute in the framework of the German ScalES project, see section 14)

The following aspects were also discussed during the User meeting and should be considered in the coming 4 years:

- testing and validation of the 3D parallel interpolations including vertical interpolation
- full support of regional models (i.e. nearest neighbour value for target point falling outside the source domain)
- storage of weights and addresses (or at least the nearest neighbour info) to avoid redoing the full calculation at the beginning of each run
- testing and optimisation of the coupler for the next generation of HPC platforms
- mixed MPI-openMP parallelisation and the thread safety

Within IS-ENES, comprehensive services will also be established around OASIS through a portal offering documentation, user guides, tutorial, FAQs, user forum and tips for best practices. In addition, IS-ENES will fund CERFACS for 3 persons-month/year for the next 4 years to provide personal technical help to implement new coupled models or improve existing configurations based on the OASIS3 or OASIS4 versions, or to help the groups to migrate from OASIS3 to OASIS4 for better performance. As the Application Programming Interface (API) of OASIS4 was designed to be as close as possible to OASIS3 API, the transition should occur smoothly in the community.

2. Ocean-atmosphere coupling with OASIS at Météo-France

F. Sevault and D. Salas-y-Melia (Météo-France/CNRM)

2.1 Introduction

Ocean-atmosphere coupling has been used for many years in the Large Scale Meteorology and Climate Research group, with the OASIS tool, and in collaboration with CERFACS. This presentation will focus on the work done on the Mediterranean Region, and in a second part on the issue of the flux conservation with OASIS in the case of global experiments.

2.2 Regional coupling on the Mediterranean Sea

The interest of regional coupling is to provide a state of equilibrium between the atmosphere and ocean models, so that the atmosphere model receives a high resolution sea surface temperature (SST) coming from an ocean model which has received the appropriate atmospheric fluxes. This system allows for example to study the climate evolution following scenarios of climate change, and also the behaviour of the Mediterranean Sea in past periods of the 20th century.

A first coupled atmosphere-ocean regional climate model on the Mediterranean region was set up a few years ago, using ARPEGE-Climate V3, OASIS2.4 and OPAMED8, regional version of OPA8.2 (LOCEAN ex-LODYC), and it was used to perform to a 21st century climate change scenario for the Mediterranean (1).

Today the ocean model has evolved to NEMOMED8, regional version of NEMO2 (LOCEAN), with the same grid as OPAMED8, which is tilted so that the Gibraltar Strait is represented by two points, and with a resolution of about 10 km on the Mediterranean. The OASIS3 coupler is used. And two atmosphere models can be used:

- ARPEGE-Climate, global atmospheric model used with a stretched grid centred on theTyrrhenian Sea, and a resolution of about 50 km on the Mediterranean; the sequence MASK-EXTRAP- MOZAIC or -INTERP are used in OASIS for the interpolation from atmosphere to ocean, and MOZAIC-FILLING for ocean to atmosphere (the ocean fields on the Mediterranean are filled with global data, with observations or large scale model outputs previously interpolated on the atmosphere grid);
- ALADIN-Climate, regional atmospheric model with resolutions of 50, 25 or 12 km on the Mediterranean; the sequence MASK-EXTRAP-SCRIPR is used in OASIS for the interpolation from atmosphere to ocean, and SCRIPR-FILLING for ocean to atmosphere (a part of the Atlantic ocean on the Gulf of Gascogne is needed to complete the ALADIN grid). This regional atmosphere model is forced by ARPEGE-Climate outside its domain, eventually following reanalysis treated by the spectral nudging technique, method which leads to hindcast runs following the chronology of the past years.

For these two coupled models, with atmosphere and ocean grids not exactly fitting, there is no issue of conservation of the fluxes during the interpolation of the fluxes and SST with OASIS.

For the CIRCE European project (2), a tri-coupled model has been prepared, with the addition of the NEMO global ocean model with its ORCA2 grid (about 1° resolution on the Mediterranean). It is composed of ARPEGE-Climate-V4.6, NEMOMED8, NEMO-ORCA2 and the OASIS3 coupler. The aim is to provide a 100-year climate scenario (A1B type of IPCC) with a small-scale grid on the Mediterranean region, but with global models, so that there is no arbitrary forcing outside the Mediterranean. This scenario among others will be used by regional modellers for their finer grid coupled models. Technically the atmospheric model sends its fluxes to both ocean models, each ocean model sends its SST, which are combined in OASIS before sending one SST field to ARPEGE-Climate. The MASKP and

BLASNEW options of OASIS are used respectively for the two steps of the combination of the SST field. The two ocean models communicate at the Gibraltar Strait, with the exchange of files respectively written and read by each model.

2.3 Ocean-atmosphere coupling: flux conservation issue at the atmosphere-ocean interface

For global coupled atmosphere-ocean models, the issue of the conservation of the fluxes exchanged via the OASIS coupler is crucial for the stability of the system. Indeed an inconsistency between the global mean net heat flux computed by the atmosphere model and the same quantity after interpolation (seen by the ocean) can generate a climate drift.

The CONSERV option is the first solution offered, but it is not fully satisfactory, because the redistribution of the difference is made globally, and not locally. Then the SUBGRID option is proposed for the solar and non-solar fluxes. For the former fluxes, SUBGRID answers the problem and is largely used. For the latter fluxes, there is a problem of coherence in time with the SST fields which are used. Let us try to explain this issue, considering that the reader knows the OASIS tool (3):

The SUBGRID option for non-solar flux Φ computes at the ocean point i, with o designing ocean, a the atmosphere field interpolated to the ocean grid, and $\partial\Phi/\partial SST$ the derivative of the flux sent by the atmosphere model:

$$\Phi_o(i) = \Phi_a + \frac{\partial \Phi_a}{\partial SST_a}(SST_o(i) - SST_a)$$

Thus at the timestep t for the ocean grid point i it can be written as (int_ao for interpolation atmosphere to ocean):

$$\Phi_o^t(i) = \mathrm{int\_}ao\left[\Phi_a^t + \frac{\partial \Phi_a}{\partial SST_a}^t\right] \times \left[SST_o^t(i) - \mathrm{int\_}ao\left(SST_a^t\right)\right]$$

But SST_a is in fact the SST_o of the previous coupling timestep, usually one day in climate coupling systems, already interpolated once from ocean to atmosphere.

So it gives:

$$\Phi_o^t(i) = \mathrm{int\_}ao\left[\Phi_a^t + \frac{\partial \Phi_a}{\partial SST_a}^t\right] \times \left[SST_o^t(i) - \mathrm{int\_}ao\left(\mathrm{int\_}oa(SST_o^{t-dt}(i))\right)\right]$$

Considering that the SST received by the atmosphere is corrected by the Gibbs effect in the ARPEGE-Climate model, then the last term of this equation leads to a loss of about 1% of the fluxes after the interpolation.

The solution proposed is to replace SST_o (t-dt) by SST_o(t) in the last term and to modify the weights of the interpolation to get int_*ao*(int_*oa*) in one operation.

Then no more time shift is introduced in the sequence, and the global loss of the non-solar flux is now of about 0.01W/m² after the interpolation.

References

(1) Somot et al, 2008: 21st century climate change scenario for the Mediterranean using a coupled atmosphere-ocean regional climate model, Global and Planetary Change.

(2) http://www.circeproject.eu

(3) S. Valcke, 2006: OASIS3 User Guide (prism_2-5). PRISM Support Initiative No 3, 68 pp.

3. Heat and water conservations in the IPSL model: interpolation, tiling and time scheme

O. Marti, A. Caubel and J. Bellier (LSCE, IPSL)

3.1 Interpolating the wind stress.

The interpolation of the wind stress, or of any vector field, from the atmosphere to the ocean is a very peculiar problem. The wind stress is defined in a local referential by two components (eastward and northward). Between two grid points, the components are defined in two different local referentials. Near the poles, the change of the local referential becomes very large. When we interpolate by doing a weighted average of 16 components (for a bicubic interpolation), we use 16 different definitions of the local referential. This may yield to very strange wind stresses near the poles !

To overcome this problem, we have adopted in IPSL CM4 the method developed at UCL to couple CLIO and LMDZ. We first compute the wind stress components in an uniform geocentric referential, linked to the Earth. The 3 components are interpolated towards the ocean. The method gives a vertical (normal to the Earth surface) wind stress component. It should be 0 when a horizontal vector is interpolated. Its computation in the validation step of the method, allows us to check that it is negligible. The components are interpolated toward the ocean in the eastward/northward referential. The last step consists to compute the component in the referential of the ORCA model. The wind stress, which in LMDZ is a vector defined at the same location that the scalar variables, is interpolated twice: toward the u and v grids of ORCA.

3.2 Sub-surface tiling

Each atmospheric column has four types of sub-surfaces: land, ocean, sea-ice and glacier. The coupling is the same whatever the sub-surface model is. For instance, the coupling follows the same method if the SST is read or is computed by a full oceanic model or by a very simplified ocean (slab ocean). In our approach, the radiative code sees only one surface, with mean properties, and computes only one net flux in both shortwave and longwave domain. Only the turbulent fluxes (sensible, latent, momentum) are computed separately on each sub-surface, and the tendency of the atmospheric column is the weighted sum of tendencies computed by each sub-surface.

The main goals of the new developments are the following:

- to redistribute the radiative fluxes, computed in the atmospheric column, on each sub-surface taking into account the local properties of each sub-surface;
- to establish a clear interface between the atmospheric boundary layer code and the surface model, whatever it is.

An absolute requirement is energy and water conservation. In the following paragraphs, subscript i stands for a subsurface i of relative fraction w_i . For each atmospheric column, one has $\Sigma_i w_i = 1$.

i. Short wave flux

The net shortwave flux at surface F^{SW} has been computed by the radiative code for the whole atmospheric columns with an albedo $r = \Sigma_i \, w_i \, r_i$ where r_i is the albedo of sub-surface i. Assuming that the downward shortwave flux is the same above all the sub-surfaces, the net shortwave flux F_i^{sw} for each sub-surface i may be written as Dufresne and Grandpeix (1996): $F_i^{SW} = (1- r_i)/(1 - r)\, F^{SW}$. One may verify that energy conservation is ensured.

ii. Longwave flux

The net longwave flux at surface Flw has been computed by the radiative code for the whole atmospheric columns with an emissivity ε and a temperature T_r, with $\varepsilon = \Sigma_i w_i \varepsilon_i$ and $T_r = \Sigma_i w_i T_i$, where ε_i is the emissivity of sub-surface i and T_i its temperature. Assuming that the downard longwave flux is the same above all the sub-surfaces, the net longwave flux F_i^{lw} for each sub-surface i reads (Dufresne and Grandpeix 1996):

$$F_i^{lw} = \varepsilon_i/\varepsilon \, \{ F^{lw} + \partial F^{lw}/\partial T_r \, (T_i - T_r) \} \text{ with } \partial F^{lw}/\partial T_r = 4\varepsilon\sigma T_r^3$$

iii. Interface for coupling the turbulent fluxes

At the beginning of a time step, the atmospheric column has a profile of temperature and humidity. The column is divided in sub-columns, corresponding to the sub-surfaces. For each sub-column i, the vertical diffusion is computed independently throughout the whole column. The result is a set of turbulent fluxes (latent and sensible) for each sub-surface, and a vertical profile of tendencies for each sub-column. The fluxes are sent to their respective sub-surfaces, and the tendencies are averaged over the sub-columns in a conservative way. In addition an interface model was also introduced to disconnect more easily surface processes from the atmosphere. The diffusion scheme was rewritten to systematically force the boundary layer by surface fluxes. The computation of surface fluxes is done in an independent model that requires providing this model with the sensitivity of the turbulent flux to temperature, in order to preserve the properties of the semi-implicit scheme. With this formulation the flux model can be either a routine in the atmospheric model, an ocean model or a land surface scheme.

3.3 Interpolation by polygon intersection (even at the pole…).

In IPSL CM4 model, the heat and water fluxes, the sea surface temperature, the sea ice temperature, and the sea ice fraction are interpolated between ocean and atmosphere by OASIS, using the interpolation scheme called 'MOZAIC'. OASIS does not compute any weight for this scheme. The user should compute the weights outside OASIS, and write the file in the format specified by OASIS. This weight computation is the purpose of the software package MOSAIC.

The basic of the weight generator is to compute the common surface between any atmosphere grid box with any ocean grid box. With correct normalisation, the ratio between the total surface and the common surface became an interpolation weight. The algorithm used to compute the common intersection between the polygons has been designed and programmed by Jacques Bellier. For MOSAIC, we have to use the algorithm on the sphere. To do that, we project the coordinates of the polygons on a plane, using a projection that conserves surfaces. The pole of the projection is the centre of one of the two polygons. A standalone program is available to compute weights. Some IPSL specificities are hard-wired in it, and it is not very versatile, re-usable, or adaptable. The POLYGON library is re-usable.

3.4 Time scheme and conservation.

At the beginning of each coupling time step, the coupler exchanges the fields between models. The fields are averaged over a coupling period, say one day. The time scheme is the following:

- ORCA computes sea surface temperature and sea ice properties (surface temperature, albedo, fraction) during day n.
- The surface properties are sent to LMDz at the end of day n. LMDZ receives them at the beginning of the day $n+1$ and uses them to run over day $n+1$.
- LMDz sends the fluxes computed during day $n+1$, averaged over the day.
- ORCA receives these fluxes at the beginning of day $n+2$, and uses them as surface conditions during day $n+2$.

This means that the fluxes used by ORCA during day $n+2$ are computed by LMDz using the sea-ice fraction of day n. In IPSL CM4, LMDz sends separately the flux over sea-ice and over free ocean. The flux seen by atmosphere is

$$Q_{total} = Q_{oce}(\text{day}=n+1).f_{oce}(\text{day}=n) + Q_{ice}(\text{day}=n+1).f_{ice}(\text{day}=n)$$

when the flux applied to ocean (the following day) is

$$Q_{total} = Q_{oce}(\text{day}=n+1).f_{oce}(\text{day}=n+2) + Q_{ice}(\text{day}=n+1).f_{ice}(\text{day}=n+2).$$

$f_{ice}(\text{day}=n+2)$ and $f_{oce}(\text{day}=n+2)$ evolves during the day. So, $f_{ice}(\text{day}=n+2)$ and $f_{oce}(\text{day}=n+2)$ are different than $f_{ice}(\text{day}=n)$ and $f_{oce}(\text{day}=n)$. The flux conservation is challenged!

In IPSL CM5, LMDz sends the flux over sea-ice Q_{ice} and the total flux $Q_{total} = Q_{oce}.f_{oce} + Q_{ice}.f_{ice}$. The flux over free ocean is computed in ORCA as $Q_{oce} = (Q_{total} - Q_{ice}.f_{ice})/f_{oce}$ (ocean fraction can not reach zero in our sea-ice model). The total flux is strictly conserved. This method has been tested in NEMO-ECHAM and NEMO-LMDZ. It will be available in the next NEMO release.

4. CMCC_MED: a 3 components CGCM

E. Scoccimarro, S. Gualdi, A. Bellucci, A. Sanna, P.G. Fogli, E. Manzini, M. Vichi, P. Oddo, A. Navarra (CMCC/INGV)

The model developed at the Euro-Mediterranean Center for Climate Change (CMCC/INGV) is a global coupled ocean-atmosphere general circulation model, specifically designed to investigate the interactions between the Mediterranean region and the global climate system.

A peculiar aspect of the present model is the use of two distinct ocean models: a coarse-resolution global ocean model and a high-resolution eddy-resolving regional model for the Mediterranean Sea.

In particular, the global ocean component adopted is version 8.2 of the Océan Parallélisé (OPA) model, in the global ORCA2 configuration with 31 vertical levels. The horizontal resolution is $2^o x 2^o \cos \varphi$ with a meridional refinement near the equator, approaching a minimum 0.5o grid spacing. The Mediterranean Sea model is a regional configuration of NEMO (Nucleus for European Modeling of the Ocean) model, with a $1/16^o$ horizontal resolution and 71 levels along the vertical, including an "Atlantic box", at the western edge of the model domain. This model version is currently being used to perform forecasting activities in the Mediterranean Sea by the Italian Operational Oceanography National Group (GNOO) (http://gnoo.bo.ingv.it).

The atmospheric model component is ECHAM5, with a T159 horizontal resolution, and 31 hybrid sigma-pressure levels along the vertical.

The global ocean model does also include a rather coarse representation of the Mediterranean Sea. However, only the high-resolution SST from the eddy-resolving Mediterranean model is transferred to the overlying atmosphere, which is overwritten on the coarser SST signal from the global model.

The Mediterranean regional model does not include the Black Sea. However, the water exchange between the Black Sea and the Mediterranean Sea is included by diagnosing the hydrological budget (precipitation + river discharge - evaporation) over the Black Sea region in the atmospheric model. The inferred water mass transport is then communicated through the coupler to the Mediterranen Sea model, as river runoff in the north-eastern Aegean Sea.

The communication between the atmospheric model and the two ocean models is performed through the OASIS3 coupler, responsible for synchronization (2.5 hours as coupling frequency) and field interpolation (Figure 1). The OASIS3 version used is the one modified by the CMCC-SCO Division designed to run per-field parallel coupling: the coupler optimization and the new implementation is described by I.Epicoco in section 13. The global-Mediterranean connection occurs through the exchange of dynamical and tracer fields via simple input/output operations. In particular, horizontal velocities, tracers and sea-level are transferred from the global ocean to the Mediterranean model through the open boundaries in the Atlantic box. Similarly, vertical profiles of temperature, salinity and horizontal velocities at Gibraltar Strait are transferred from the regional Mediterranean model to the global ocean. The ocean-to-ocean exchange occurs with a 8h frequency, with the exchanged variables being averaged over the daily time-window. Coupling between oceans needs improvement, coherently with steps done in OASIS4 version, particularly in terms of not gridded data handling.

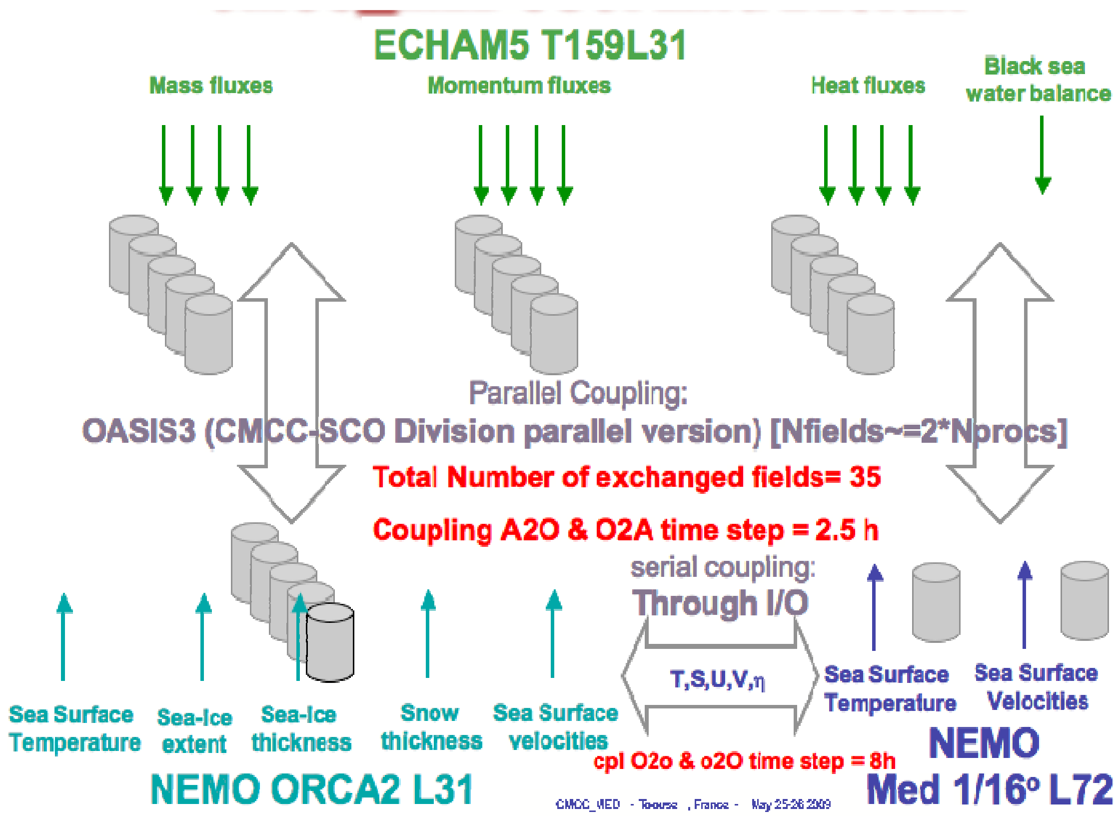
ECHAM5 T159L31
Mass fluxes
Momentum fluxes
Heat fluxes
Black sea water balance
Parallel Coupling:
OASIS3 (CMCC-SCO Division parallel version) [Nfields~=2*Nprocs]
Total Number of exchanged fields= 35
Coupling A2O & O2A time step = 2.5 h
serial coupling:
Through I/O
T,S,U,V,η
Sea Surface Temperature
Sea-Ice extent
Sea-Ice thickness
Snow thickness
Sea Surface velocities
Sea Surface Temperature
Sea Surface Velocities
NEMO ORCA2 L31
cpl O2o & o2O time step = 8h
NEMO
Med 1/16° L72

5. Coupled modeling at ECMWF: Waves, Ocean and Chemistry

K. S. Mogensen, J. Bidlot and J. Flemming (ECMWF)

5.1 Introduction

ECMWF's activities in coupled modelling consist in coupling our atmospheric model (Integrated Forecasting System or IFS in the following) to different component models (in no particular order): several different chemical transport models (CTM's), two different oceans models (HOPE and NEMO) and a wave model (WAM). Each of the systems uses different coupling frameworks.

5.2 Coupling of IFS to chemical transport models

The aim of the coupled IFS-CTM system was to build an OASIS4 coupled system that relied on existing CTM's to exploit the existing IFS data assimilation capability without direct integration of chemistry, deposition and emission injection into the IFS model. Three CTM's models were chosen as candidates for the future operational system: MOZART, TM5 and MOCAGE. During the coupling, up to 8 3D-grid-point meteorological fields (u,v,T,q, convective mass fluxes) and up to 20 2D-grid-point surface fields are sent from the IFS to the CTM's. From the CTM's to the IFS 18 3D-grid-point fields of concentrations, production and loss rate of NO_x, O_3, SO_2, and HCHO are passed. The coupled IFS-OASIS4-CTM system is scientifically sound but it is less efficient than a fully integrated CTM-IFS (with the chemistry within the IFS code) system would be, so the long term plan is to develop such a fully integrated system. However experimentation will be based on the OASIS4 coupled system in the next few years.

5.3 Coupling of IFS to ocean models

The present operational coupled ocean-atmosphere system uses the IFS model coupled to the HOPE ocean model. This system is used for both seasonal forecasting (system 3) and as part of the ensemble prediction system (EPS). Within the EPS system the coupled IFS-HOPE system is running once per day (0Z) for the model integration from day 10 to 15 and once per week (Thursday 0Z) for the model integration for monthly forecasting from day 10 to 32. For all 0Z EPS runs from day 0 to 10 are excecuted without any coupling to HOPE. The HOPE model (as implemented at ECMWF) is shared memory (OpenMP) parallel only and the coupling is done via files using OASIS2. The IFS-HOPE system is not actively developed anymore but only maintained for operations.

In the future ECMWF is going to replace the HOPE model with the NEMO model for seasonal forecasting (system 4) and EPS/Monthly forecasting. For this new system, the coupling will be done with OASIS3 via MPI1. Implementing coupling via OASIS4 has been tried early in the development process, but at the time, there were serious problems with the interpolation between the tri-polar ORCA grids and the reduced Gaussian grid of the IFS. A prototype of the coupled IFS-NEMO system exists and is currently being developed further to be able to meet all requirements of the operational atmosphere-ocean coupled systems.

Initially the IFS-NEMO system will not use the LIM ice model component of the NEMO framework, but the use of the LIM ice model will be explored later on. The long term plan on how to couple the atmosphere and the ocean are also being considered, since it is clearly not ideal to use OASIS4 for the CTM coupling and OASIS3 for the NEMO coupling.

5.4 ECMWF Wave Model coupling

The atmospheric coupling to the wave model (WAM) is active for all operational forecasts done by ECMWF (medium range/monthly/seasonal). The WAM model runs on an irregular lat/lon grid with the land points removed. For the coupled IFS-WAM model, the wave model is just a subroutine call in the IFS time stepping and no external coupler is used. The regridding from the atmospheric grid to the WAM grid is done as part of this subroutine call. This is a very different approach than the approach used in the various OASIS couplers, where the regridding is done within the coupler. Potentially, the single executable approach is more efficient if all components (the atmosphere model, the regridding and the wave model) scale well, since all components have all processors/threads available.

5.5 Conclusions

The OASIS couplers have been and continue to be valuable tools for ECMWF. We are currently using 3 different versions of the OASIS coupler:

- OASIS2 has been used for operational coupled ocean-atmosphere seasonal and monthly forecasting for many years. Soon this system will be retired, but it has served us well so far.
- OASIS4 has been used to get the IFS-CTM system up and running and will be used for pre-operational applications in the MACC project. Long term plans to fully integrate chemistry in the IFS are being pursued.
- OASIS3 has been used to get the IFS-NEMO system up and running and will be used for the operational seasonal and EPS forecast in the years to come. Long term plans are being considered.

6. OASIS at MPI-M

L. Kornblueh, J. Jungclaus, R. Budich and many more ... (MPI-M)

6.1 OASIS usage

OASIS4 from our perspective offers a very usable tool for our climate simulations. Beside the large user community, OASIS is as well a model independent quasi standard. It is offered as an Open Source package keeping the freedom to change the code in any aspect. Another very important aspect is the ongoing support.

Only minor drawbacks are the lot of legacy heritage. The site centric, non standard build system which does not allow to build the system as an independent library and binary (that could be done easily) is the most annoying feature. The steep initial learning curve is an inherent problem one cannot really bypass.

Our application environment has recently changed from the NEC SX-6 at DKRZ offering a maximum of 192 tasks. Currently two big production systems are available as a Sun Linux Cluster (Opteron Shanghai/IB) run by DKRZ and MPIM and the new production system at DKRZ an IBM AIX Cluster (Power6/IB) offering around 16000 tasks.

Our coupled model is supposed to be supported and run as well on a lot of other machines. The list of the required platforms spans from SGI Altix over Cray XT systems to the latest NEC SX-9 at many sites.

A change to OASIS4 is expected during 2010 reducing the serial part in simulations, using an improved more user-friendly interface, improving debugging due to cleanly implemented new code base with no legacy inheritance, and last allowing for rewriting our model interfaces in a clean way.

Our coupled model setups are currently including ECHAM5/MPI-OM, ECHAM5/MPI-OM/HAMOCC, and ECHAM5/JSBACH/MPI-OM/HAMOCC. We decided to couple only the atmosphere-ocean interface, because otherwise memory requirements are exploding and communications cost is dramatically increasing. This is not caused by OASIS capabilities but by hardware restrictions.

Research questions in focus are defined by our last big experiments: the simulation of the last Millenium and the simulation of super volcano eruptions. Those will allow us to address questions including (but not limited to):

- To what extent are the observed pre-industrial climate variations driven by natural forcings (orbital, solar, volcanic)?
- How did the climate system respond to human activities (land use changes, industrialization)?
- How did the carbon cycle respond to natural and anthropogenic disturbances and how important are carbon-climate feedbacks?
- What are the relations between forcing, climatic states, and variability patterns?
- And can ensemble simulations support the interpretation or proxy-based reconstructions?

6.2 OASIS3 lessons learned

OASIS3 is in use for more than 10 years at MPIM. It has been used for IPCC simulations and Millennium (tens of thousands of simulated years). We learned the stable MPI based software needs to be checked on all available MPI software stacks, because standards are sometimes not sufficiently understood. Clean standard conforming code allows for memory checking and debugging and is therefore important to speed up developments. As models are hardly following this requirement, it helps, if all libraries do. A good development team helps in getting your model running.

6.3 OASIS4, the future

OASIS4 is in our view, the coupler for the future. It has a large potential user community and will hopefully be a standard soon. It is open source, offers a lot of the advantage learned from OASIS3 lessons and has continued support. However, some questions arise with respect to future hardware developments: is it thread-safe? , can it be used in hybrid mode?, what is about the adaptation to new hardware/software models (accelerators, multicore)?

Important features for MPI-M new model development are the support of fully unstructured arbitrary triangular or hexagonal/pentagonal grids. A build system which is more in line with the Open Source community would as well be really great.

7. NEMO + LIM + IFS + OASIS = EC-Earth

A. Sterl, C. Severijns (KNMI)

EC-Earth is a new Earth System Model (ESM) that is being developed by a consortium of at present 23 European institutes. The consortium is open for institutes from ECMWF member states. The core of EC-Earth is formed by the IFS, the operational weather forecast model of ECMWF, for the atmosphere, and the NEMO/LIM ocean/sea-ice model from IPSL and the university of Louvain-La-Neuve for the sea-ice. The ocean and atmosphere components are coupled through OASIS. In the future this core will be developed into a true ESM by incorporating other components, e.g., for atmosphere and ocean chemistry.

The prime advantage of using a weather forecasting model for climate studies is the operational infrastructure; an enormous amount of observations can be assimilated into the model and its behaviour can be verified against observations from the daily and seasonal to decadal time scales. The model is constantly improved by dedicated experts, and the expertise of the Seasonal Prediction group at ECMWF can be called upon.

The current version 2 of EC-Earth is based on cycle 31r1 of IFS with some improvements (gravity wave drag, dry mass conservation) from later cycles. A resolution of TL159L62 is used. NEMO/LIM version 2 is used in the ORCA1 configuration with 42 vertical levels. The time step of both models is one hour with a full coupling through OASIS at every three hours.

A throughput of about 3.4 years/day is reached on ECMWF's IBM Power5+ system. This version of EC-Earth will be used to perform runs for the forthcoming CMIP5 climate model intercomparison. At the same time, version 3 of EC-Earth is under development. It's major progress will be the use of NEMO/LIM version 3.

At each time step the current ocean surface values (SST, sea ice coverage, albedo, ...) are passed form NEMO to IFS, while energy and momentum fluxes are passed the opposite direction. A complication arises here as IFS does not have a fixed land-sea mask. Instead, each grid box is subdivided in eight tiles representing different surface types (open water, sea ice, and 6 different forms of land). To conserve fluxes per tile-type, not only the fluxes, but also the tile-fractions have to be passed from IFS to NEMO (see Appendix). The tri-polar grid employed in NEMO creates some extra problems for the coupling. One relates to the large difference in grid-size between the two models (IFS and NEMO/LIM), causing the local SST to deviate much from the average SST that is used to calculate fluxes. A linear correction is applied which unfortunately leads to non-conservation of fluxes. A second problem is caused by the fact that mesh boundaries near the poles of the NEMO grid are curved lines, while the implementation of the 1st order conservative mapping scheme in OASIS assumes straight lines. This also leads to a non-conservation of fluxes between IFS and NEMO.

Despite these problems the coupled model works satisfactorily and is able to reproduce essential aspects of the observed climate. The Reichler & Kim (2008) performance index shows that it is better than the average climate model used for IPCC's Fourth Assessment Report. Its ENSO activity is very realistic with an amplitude of the SST anomalies in the eastern tropical Pacific of 1-2 K and a timescale of 3-5 years. Also other aspects of inter- and intra-annual variability, like NAO and the seasonal cycle, are very realistic. The greatest model bias encountered so far is a much too high SST (locally exceeding observations by more than 5 K) in the Southern Ocean (south of 40◦S). This problem is apparently caused by too much vertical exchange and is currently under investigation.

Reference

Reichler, T., and J. Kim (2008) How well do coupled models simulate today's climate? Bull. Am. Meteorol. Soc., 89, 303-311, doi: 10.1175/BAMS-89-3-303.

Flux coupling in EC-Earth

Global conservative regridding of a flux F requires that

$$(1)\ F = F^*$$

where the asterisk indicates the flux after regridding.

The flux can be written as the sum over all grid cells in both the source and target grids

$$(2a)\ F = \sum_i f_i A_i$$
$$(2b)\ F^* = \sum_j f_j^* A_j^*$$

where f is the grid cell mean flux and A is the grid cell area.

Each grid cell consist of a number of tiles thus

$$(3a)\ F = \sum_j \sum_t \alpha_{i,t} f_{i,t} A_i$$
$$(3b)\ F^* = \sum_j \sum_t \alpha_{j,t}^* f_{j,t}^* A_j^*$$

Since both summations have a finite range, they can be swapped

$$(4a)\ F = \sum_t \sum_i \alpha_{i,t} f_{i,t} A_i$$
$$(4b)\ F^* = \sum_t \sum_j \alpha_{j,t}^* f_{j,t}^* A_j^*$$

In the coupling flux conservation should be guaranteed for each tile

$$(5) \sum_i \alpha_{i,t} f_{i,t} A_j = \sum_j \alpha_{j,t}^* f_{j,t}^* A_j^*$$

The atmosphere model computes the fluxes $f_{i,t}$ for each tile and each grid cell of the atmospheric grid. The tile fractions for land are prescribed and those of open ocean and sea ice are received from NEMO. The latter are given for the binary ocean mask used by NEMO. The fractions are adjusted to match the land distribution so that the total of land, ocean and sea ice is everywhere equal to one in IFS.

IFS sends the tile fractions $\alpha_{i,t}$ and the tile fraction weighted fluxes $\alpha_{i,t} f_{i,t}$ to NEMO using the first order conservative regridding method of OASIS. This regridding method takes care of the area weights and is local conservative. This latter property implies that equation (5) also is valid for a single grid cell in the target grid and can be simplified to

$$(6)\quad \sum_i \alpha_{i,t} f_{i,t} A_i w_{i,j} = \alpha^*_{j,t} f^*_{j,t} A^*_j$$

where $w_{i,j}$ is the overlap of the grid cell i in the source grid and the grid cell j in the target grid.

The coupling now works as follows: IFS computes the fluxes for each grid cell and tile type and sends the fields $\alpha_{i,t} f_{i,t}$ and $\alpha_{i,t}$ to NEMO

$$(7a)\quad \alpha_{i,t} f_{i,t} \Rightarrow (\alpha_{j,t} f_{j,t})^* = \alpha^*_{j,t} f^*_{j,t}$$

$$(7b)\quad \alpha_{i,t} \Rightarrow \alpha^*_{j,t}$$

and NEMO computes the flux for each grid cell and tile type

$$(8)\quad f^*_{j,t} = {(\alpha_{j,t} f_{j,t})^*} \big/ {\alpha^*_{j,t}}$$

8. Introduction to the Australian Climate Ocean-Sea ice Model

D. Bi (CAWCR - CSIRO)

The Australian Community Ocean Model (AusCOM) is an IPCC class coupled ocean-sea ice model developed by the Australian climate sciences community (including the Bureau of Meteorology, CSIRO and the Australian universities) for climate research and applications. AusCOM comprises the Geophysical Fluid Dynamics Laboratory Modular Ocean Model (MOM4p1), the Los Alamos National Laboratory sea ice model (CICE4.0), and a data atmospheric model. Numerical coupling is via the OASIS3.2-5 coupler. AusCOM is one of the core parts of the Australian Community Climate and Earth System Simulator (ACCESS). Specifically, the ACCESS model is built by coupling the UK Met Office atmospheric model UM (Unified Model), and other sub-models as required, to AusCOM, under the same coupling framework.

The horizontal grid of the AusCOM ocean model MOM4p1 is a curvilinear, globally orthogonal, tripolar grid, which is designed for the purpose of avoiding the North Pole singularity, providing reasonably fine resolution in the Arctic Ocean, and thus enhancing computational efficiency and accuracy of the model. The AusCOM sea ice component CICE is also configured on this tripolar grid and has the same horizontal resolution (360 x 300).

In AusCOM simulations, at the beginning of a coupling interval, the atmospheric 'raw' data, on the NT62 Gaussian grid, is read in by the data model and passed into CICE via oasis3 coupler, using the SCRIPR 1st order conservative remapping algorithm. A boundary layer in CICE processes these received 'raw' data, using standard bulk formula, to obtain the atmospheric forcing for the ocean and sea ice. Forcing fields required by MOM4 are then passed into ocean after being adjusted by the ice existence where appropriate. At the same time, MOM4 sends the oceanic fields needed by sea ice into CICE, completing the coupling process, and every component model moves on.

It is worth being mentioned that

- Coupling between ocean and sea ice in AusCOM is conveniently performed via oasis3, i.e., data is exchanged directly without being transformed by the coupler. Furthermore, since both MOM4p1 and CICE are on an Arakawa B-grid, the exchanged coupling fields need no additional handling such as grid shifting after being received.
- Sea ice model works as a 'coupling buffer' in the AusCOM system. Namely, atmospheric data is received and processed here, and then passed into ocean. On the other hand, when an active atmospheric model such as UM is in place, the required oceanic forcing is also firstly passed into ice model and then, together with the required ice variables, sent to UM by the coupler. This is the coupling approach for the ACCESS fully coupled model.
- While the frequency for atmosphere-sea ice coupling in AusCOM is set to 4 time daily, i.e., the temporal resolution of most of the atmospheric forcing fields, the ocean-sea ice coupling interval can be 1, 2, 3 or 6 hours, depending on the experiment design and the model time steps.

AusCOM has originally developed on the NEC SX6 platform at the Australian Bureau of Meteorology/CSIRO High Performance Computing and Communications Centre, and has recently been ported onto the Australian National Computational Infrastructure (NCI) SGI AC/XE Clusters. Simulations with integration of decades and centuries under various atmospheric forcing have been performed.

This presentation focuses on the recent experiments with atmospheric forcing from the CLIVAR Group for Ocean Model Development (WGOMD) Coordinated Ocean Reference Experiments (CORE) datasets, including both the Normal Year (NY) and the Inter-Annual (IA) Forcing. We examine model performance in terms of key global parameters that are of importance to climate studies. These include the meridional overturning circulations, mass transport through key straits, water mass properties, and the seasonality of sea ice thickness and areal concentration. A realistic simulation of AusCOM under climatological and interannually varying atmospheric forcing is a key indicator of expected performance of the fully coupled ACCESS model.

Generally speaking, AusCOM produces an oceanic climatology comparable to results from other ocean-sea ice models of similar class in the world. However, like other models, direct comparison of the model simulation against observations reveals remarkable biases in most of the shown climatological indicators. For example, the annual cycle of sea ice distribution is not very realistically simulated in both hemispheres, particularly off Antarctic where the sea ice retreats too rapidly and extensively in summer, mostly resulting from the warm biases in the ocean surface temperature, which is possibly associated with the unrealistic deep mixing in the ocean under ice that brings up the warm water from depth.

For the CORE IA forcing run, we compared observed and modelled SST and upper ocean heat content in the equatorial Pacific. The correspondence is quite remarkable, showing reasonable simulation of the annual cycle and inter-annual variability, suggestive of good reproduction of the equatorial wave dynamics. There is little or no net drift of the ocean temperature field overall for this region. This indicates good simulation of at least some of the tropical processes important for ENSO.

The computational performance of AusCOM depends on the sub-models, especially the ocean model which takes most of the computing resource. The one-processor oasis3 coupler uses a fairly small fraction of the computing time and therefore the AusCOM coupled system achieves remarkable scalability (as that for the standalone ocean model) on the NCI SGI platforms. This is also proven to be the case for the ACCESS fully coupled system in our recent ACCESS test runs. Therefore we are quite confident that the mono-cpu oasis3 prism 2-5 coupler will not be the ‘bottleneck’ for our current IPCC class coupled models in terms of the computing efficiency. However, should the need arise in the near future, (e.g., for new version of models with considerably higher horizontal resolution, or/and significantly higher atmosphere-sea ice coupling frequency), we are prepared to upgrade the coupler to oasis4 or the multi-executable oasis3.3 (Pseudo-parallélisation coupler).

9. OASIS: A useful coupler for SINTEX-F model on the Earth Simulator

J.-J. Luo, S. Masson, A. Caubel, S. Shingu, C. B. Montegut, W. Sasaki, and T. Yamagata
(FRCGC JAMSTEC)

The OASIS coupler has been used for the development of two SINTEX-F coupled models on the Japan Earth Simulator (ES) since 2001. We found that the OASIS coupler is a quite convenient and useful standardized interface for developing climate models. We started from OASIS2.4.0 for our early model development, which is just applicable for one-node (8 CPUs) computations on the ES. We then turned to OASIS2.4.1 with the availability of MPI1 interface which is necessary for multi-nodes computation. Since the atmospheric model is much heavier than oceanic component, we have also modified the OASIS coupling strategy in order to shorten the communication time. This turned out to save about 23% of CPU time for one-month integration of the SINTEX-F model. In addition, we have improved the model coupling physics which is important to reduce some common biases of the tropical Pacific climate and ENSO simulation (Luo et al. 2005, J. Climate). The SINTEX-F model has been widely used for various climate variability and prediction studies and shown excellent performance (see http://www.jamstec.go.jp/frcgc/research/d1/iod/index.html, click “seasonal prediction”).

Since 2005, we have started to develop a new coupled model (SINTEX-F2) with high-resolution ocean and atmosphere components by using OASIS3 coupler. Because of the dramatic increase of coupling fields after adding a sea ice model, we found that the OASIS3 coupler became less efficient relative to the CPU time consumed by ocean and atmosphere models. To solve this problem, we have developed a pseudo-parallelized version of OASIS3 which allows the coupler to use multi-CPUs; individual processor is able to deal with a subset of coupling fields independently. We found that this approach is much more efficient than simply increasing the CPU numbers for the ocean or atmosphere model. The pseudo-parallelized version of OASIS3 works well for both high and ultra-high resolution coupled models of SINETX-F2.

After the renewing of the Earth Simulator (ES2) in April 2009, however, we have been facing serious problems. The two coupled models become slower despite that the ES2 is about 12 times faster than the ES. The major problem appears to come from a slow MPI communication and shorter vector lengths in all components (OASIS, OGCM, and AGCM) as well as a slow NETCDF I/O. This requires to optimize the physical models and to have a more efficient OASIS coupler; this is particularly true in the future under the stress of increasing CPU speed of super-computers.

10.OASIS High End Computing on NEC SX9

E. Maisonnave (CERFACS)

Taking advantage of recent Météo-France NEC SX9 Operation Health Check (OHC), several configurations of ARPEGE-NEMO climate model have been set up in order to check OASIS3 capability at high end resolutions.

Tested resolutions are t359 (approximately 50Km square) with 31 vertical levels for ARPEGE and ORCA ¼ degree with 50 vertical levels for NEMO. Those resolutions are similar to the higher operational configurations in use at the moment in Europe (ECMWF, MetOffice ...)

The new Météo-France supercomputer is composed by 6+7 vector nodes of 16 processors each of 102 Gflops peak performance and 1 Tb memory per node. Due to OHC time constraints, no fine optimization has been implemented. In particular, during simulation, data processing has not been done on local processor disks but through GPFS global file system.

Our coupling method implies that ocean and atmosphere perform simultaneously a coupled time step (namcouple LAG mode), each model using the coupled field averaged at the previous coupling time step.

Moreover, we used the MPI communication library with bufferized messages (namcouple NOBSEND option disabled), reducing communication time spent at “send” step.

ARPEGE and NEMO are parallel codes: allocating appropriate processor number to each model, response time difference between each model could be minimized. However, due to vector processor efficiency (a fast simulation could be processed with less than 10 processors), possible combinations in processor repartition are limited and a model still remains significantly slower than the other. If OASIS duration (communication and interpolation processing) stays lower than this difference, there is no coupling additional time compared to slower stand alone simulation duration: OASIS calculations and communications are completed during the time interval between the end of fastest model and the end of slowest model calculations.

In order to measure those quantities, clock times are collected (thanks to light modifications under CPP key within OASIS “psmile” library) before and after each OASIS exchanges. Measures are done on elapsed time and not on CPU time (OASIS CPU time are supposed to be independent of model parallelism ratio). The balance we proposed to tune here only influences elapsed time performances. But measures on elapsed time are machine load dependent, particularly if we share node with other users: an ensemble of 5 to 9 simulations (of 4 days of climate) will be processed and uncertainty evaluated.

Two kind of measures are done to evaluate (a) the total coupled simulation time and (b) the calculation duration of each component of the coupled system.

(a): Within OASIS, the difference between the very first “prism_get” of the whole simulation and the very last “prism_put” represents the whole climate simulation duration, excluding model restart read/write operations.

(b): Within models, we measure the interval (at each coupling time step) between the instant after the “prism_get” of last coupling field received at coupling time step N and the instant before the “prism_get” of the first coupling field exchanged at time step N+1: in this way, we evaluate duration of calculation processed by each model between two calls to OASIS (supposing that the time for the non-blocking “prism_put” calls is negligible in this measure).

This counter allows us to finely determine the respective number of processors which minimize the difference between oceanic and atmospheric coupling step durations. If 4 processors are allocated to NEMO (due to speedup optimum considerations), 6 to 7 processors for ARPEGE are most suitable. An important extra cost due to OASIS (more than 50%) incites us to optimize OASIS coupling technique.

In a first step, the OASIS sequential exchange (namcouple SEQ mode) is preferred to the standard one: instead of exchanging all the fields before doing any interpolation calculation, the sequential exchange realize field by field the sequence “get field – process oasis computations – put field”. So, OASIS do not need to wait the slowest model to begin to processes the first interpolation. And interpolated fields are ready to be used by the slowest model as soon as it needs them. This optimization reduces the total elapsed time of our coupled simulation a 20% in the best case.

The second optimization consists on using the OASIS parallelism by field. This configuration called “Oasis pseudo parallelism” is implemented in OASIS3 version since Arnaud Caubel – Sébastien Masson – Jing-Jia Luo IPSL-JAMSTEC joint experiment on Earth Simulator supercomputer. Within this configuration, OASIS is launched several times with different namcouples, each namcouple describing a subset of one or several of the initial coupled fields needed. Several OASIS executables process a subset of the initial coupling fields: communications and interpolation calculations are done in parallel. This optimization also reduces the total elapsed time of our coupled simulation by 18% in the best case.

Combining those two optimizations (Sequential mode + Pseudo parallelism), the cumulated gain compared to the non-optimized run varies between 15 and 25 %. Compared to the slowest model on a stand alone mode, the extra cost of coupling oscillates now between 0 and 25 %. The worst figure is obtained in the case of a load balanced configuration (4 processors for NEMO, 7 for ARPEGE): the difference between model durations is lower than the extra cost induced by the coupling.

To be totally sure that OASIS3 could handle model resolution higher than the present European most demanding configuration ones, we attempt to increase oceanic resolution, using the MERCATOR state-of-the-art operational model (1/12th degree, 50 vertical levels).

At such resolution, even during an OHC period, the machine load is so important that a limited number of tests is possible. A 9 member ensemble test of 2 simulated days has been processed, using 44 processors for NEMO and 4 for ARPEGE (3 nodes, half of the total amount of machine processors). Even if load balancing could not be totally reached, both Sequential mode + Pseudo parallelism optimizations help us to reduce the total simulation time from 8 to 4 hours, with only a 10 % additional time compared to the oceanic stand alone simulation.

Those experiments prove the OASIS3 capability to drive high end resolution coupled simulations on vector machines (even with experimental configurations using a 1/12th degree ocean) with reasonable additional time. Most of the time, when balance between model component duration cannot be reached, this additional time even could be nullified.

Author is grateful to R. Bourdallé-Badie, O. Le Galloudec (MERCATOR) and M. Déqué (CNRM) for providing their NEMO and ARPEGE configurations. We also would like to thank S. Valcke (CERFACS) for fruitful discussions, M. Pithon (Météo-France), N. Monnier and I. d'Ast (CERFACS) for their help in code porting.

11.High resolution coupled simulations on the Earth Simulator

S. Masson, G. Madec, C. Talandier, R. Benshila, P. Terray, A. Caubel, E. Maisonnave, M.-A. Foujols (IPSL); J.-J. Luo, T. Izumo, C. de Boyer Montegut, T. Yamagata, K. Takahashi (JAMSTEC)

This presentation is a quick overview of the recent high-resolution coupled simulation performed on the Earth Simulator in the frame of an EU-Japan collaboration. OASIS 3 in its pseudo-parallel version was used to couple NEMO 2.3 and ECHAM 5.3.

The main part of our results focuses on the comparison of 2 coupled simulations sharing exactly the same atmospheric component and coupling interface but with different resolution of the oceanic component: ORCA05 and ORCA025. We must underline that each resolution needs a specific set of physical parameters. Therefore these two experiments differ not only from their resolution but also from some part of their physics. They underline the impact of the oceanic component on the climate and its variability, but results involving only the impact of ocean resolution remain hazardous at this stage of our research.

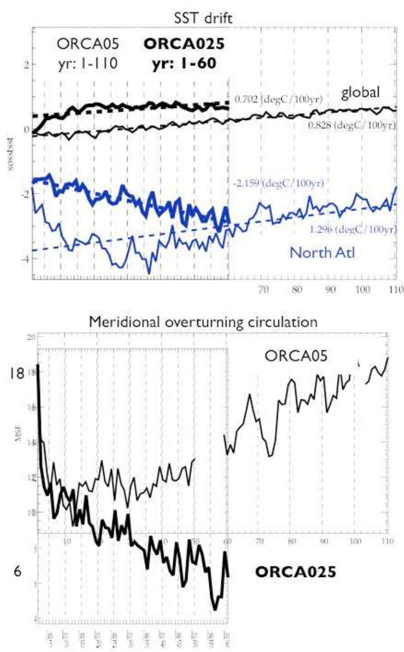

Major climatological SST biases are very similar in both experiments (and close to all NEMO-ECHAM coupled models) with a zonal distribution: warm bias in the ACC and the tropics (mostly offshore of the coastal upwellings) and cold bias at mid-latitudes (especially in the north Atlantic). The global SST drift is moderate (0.7° to 0.8° over 100 years, see black curve in the upper panel figure 1) but once again, the North Atlantic differentiates from the rest of the world (blue curve). With ORCA05, SST strongly drops to reach a bias of less than -4℃ over a large region after 30 years. This cooling is associated with a collapse of the thermohaline circulation within the first 10 years of the simulation using ORCA05 (from 18Sv to 10Sv, see bottom panel figure 1). We must note that this decrease is particularly strong during the first years of the simulation (more than 4Sv the first year). After about 50 years, tendencies reverse in the experiment using ORCA05: amplitude of the cold SST bias decreases (-2℃ after 100 years) and thermohaline circulation reaccelerates to reach its realistic initial value of 18Sv after 100 years of simulation. At the beginning of the experiment, ORCA025 simulation shows a much weaker decrease of the north Atlantic SST bias. But after 60 years, the tendency is not reversed and SST bias is still increasing whereas it is not the case in the ORCA05 test. The meridional circulation displays the same evolution, following closely ORCA05 drop during the first 10 years. Afterwards it doesn't show any reversal of the tendency and it is still decreasing down to 6Sv. Unfortunately, Our experiment using ORCA025 is only 60-year long making impossible to conclude about longer-term evolution of the north Atlantic SST of meridional circulation.

ENSO variability is also affected by the changes in the oceanic component between our two simulations. Changes in periodicity (around 3-4 years) are difficult to see, however month-to-month interannual variability and dissymmetry between Niño and Niña events are strongly different. In the observations, a weak signal appears at the east of the Pacific in May and seams to propagate westward until winter where the major pattern is visible in the central and eastern Pacific. Decomposition between positive and negative events shows that Niño

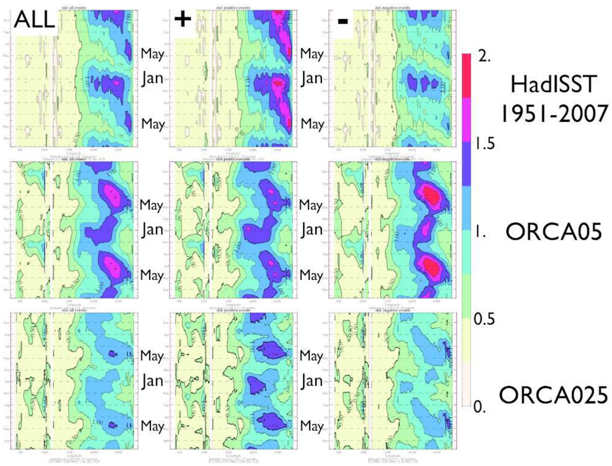

signal is stronger with la Niña mostly visible in the central Pacific in winter with a slight eastward propagation at the beginning of the year (first line of figure 2). Biases of ORCA05 interannual variability (second line of figure 2) are characterized by a much too strong signal in spring-summer time and the too weak amplitude of the variability in the eastern Pacific in winter. The spring-summer bias is mostly associated with a totally unrealistic signal of la Niña which amplitude exceeds 2℃ in June-July. Winter bias appears to be shared between Niño and Niña. Interannual variability of ORCA025 (third line of figure 2) is completely different. Opposite to ORCA05, its amplitude is always too small especially in winter. The strong Niña signal almost disappears which is an excellent point in spring-summer but winter variability is underestimated. Regarding El Niño, spring signal is slightly stronger than the winter one which is not realistic. In winter, the signal is stronger in the central Pacific than in the eastern Pacific, which disagrees with the observations.

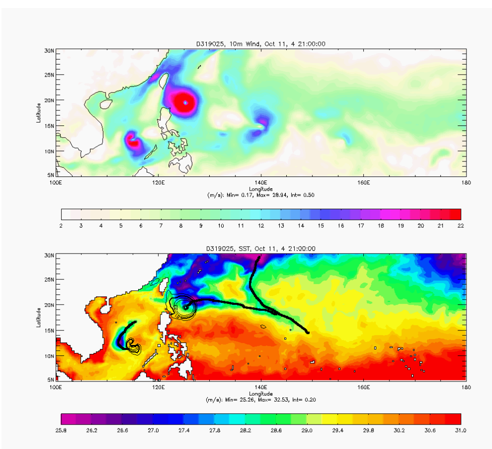

Before concluding, we would like to add a few lines about one of our most spectacular result of our experiment with a very high resolution coupled model: ORCA025-T319. Thanks to is atmospheric resolution, this coupled model configuration is able to reproduce typhoons an cyclones at the proper place and time with realistic amplitude of atmospheric and oceanic signals. Figure 3 is showing the wind speed (upper panel) and the SST (lower panel) in the north-eastern tropical Pacific. Two typhoons are clearly visible with wind speed exceeding 25m/s. The impact on SST is also remarkably decreased (over 5℃) along the typhoon trace (black curve).

In conclusion, these first results are showing that modifying the oceanic component of a coupled model can have deep impact on the modelled climate and its variability. However several unexplained questions are rising from this presentation:

- Why and how ORCA05 Atlantic meridional circulation is able to recover after its strong collapse at the beginning of the experiment?
- What would be its evolution on a longer time period? Is the meridional circulation of ORCA025 definitely collapsed?
- How could we prevent the strong decrease of the meridional circulation in the first years of the simulation?
- Will it help to maintain a realistic amplitude during the whole simulation?
- Which physical processes are involved in the modification of the ENSO variability between ORCA05 and ORCA025 experiments?

12. Using OASIS in HadGEM3

R. Hill (Met Office Hadley Centre)

12.1 Introduction

The Met Office HadGEM3 coupled model is the first Met Office coupled model to employ separate model components for climate modelling, coupled via the OASIS coupler. The Met Office previously used a single executable containing atmosphere, ocean and sea-ice components. HadGEM3 also been employed as the basis for seasonal forecasting and is expected to be used in shorter-range forecasting.

12.2 Background and History

The Met office has been testing and applying OASIS3 and OASIS4 in "serious" climate modelling since about 2006. Prior to that, there had been some investigation into the use of OASIS1 and OASIS2; however no infrastructure had evolved to support the use of OASIS couplers within the Met Office model control system. The adoption of the IPSL NEMO and LANL CICE models to replace the existing Unified Model ocean and sea-ice models was the catalyst for serious development using OASIS.

Initially, development and testing work was carried out using a prototype version of OASIS4, with a 2 degree resolution global model; however project time-scales for an operational model to be available meant that the switch was made to OASIS3.

12.3 Current Status

The Hadgem3 model employs OASI3 on a single executable, with a UM atmosphere on a regular C-grid of 192 columns East-West (1.875°) x 144 rows North-South (1.25°) x 38 vertical levels and a combined NEMO-CICE executable on a tripolar 1° C-grid (NEMO) and B-grid (CICE).

Remapping weights files are generated off-line prior to a run and adjusted by hand, where necessary, to cater for grid points which the SCRIP algorithms have difficulty with (e.g. awkward land-sea masks and quirks of the tripolar grid).

OASIS restart files are not used, since all the necessary data for reproducible restarts and continuation runs is stored in the respective component model restart and dump files. This also cuts down on the already high number of files which have to be managed by the Met Office archiving systems.

Rotation of vector fields between the atmosphere and NEMO-CICE grids is carried out explicitly in the NEMO component on receipt of incoming fields or just prior to sending outgoing fields. OASIS3 is not used to perform the rotation, partly due to the complexities introduced by the atmosphere having one row fewer of V grid data compared with the U grid data.

Coupled climate runs of up to 100 years have been completed on NEC SX8 machines. The coupled model has already been adopted operationally by seasonal forecasters. The coupled model has recently been ported to the Met Office's new IBM Power6 machine.

12.4 Performance

Due to limited availability of resources on the NEC, the coupled model on the SX8 was not well load balanced. The atmosphere is considerably slower than the ocean and even with a processor configuration of 1x6-1-1x1 (Atm-OASIS-NEMOCICE) the ocean spends a considerable time waiting for the atmosphere to catch up.

The IBM affords better opportunities for load balancing, with a ratio of 5 or 6 atmosphere processors to 1 NEMOCICE currently being used e.g. 8x10-1-1x15(15x1). Scalability of the coupled model is no more than adequate, but this does not appear to be due to any OASIS3 bottleneck, rather due to atmosphere scalability being quite bad beyond about 128 CPUs. Coupling appears to be almost a fixed cost regardless of CPU numbers.

12.5 Future Work

Although the use of a single OASIS3 instance appears adequate for current models, the Met Office aims to develop a high resolution model which may additionally involve more frequent coupling. There is a suspicion that OASIS3 may become a limiting factor in this, hence there is an interest in employing multiple instances of OASIS3 in "pseudo-parallel" mode and ultimately upgrading to OASIS4.

The continual need for upgrading to new releases of the UM atmosphere, NEMO, CICE, NetCDF etc presents an ongoing and complex configuration management task.

13. Oasis3: an MPI1/2 per-field parallel approach

I. Epicoco, S. Mocavero, G. Aloisio (CMCC)

13.1 Coupled set-up

The coupled model analyzed in this work is a 3 component model designed at CMCC (S. Gualdi, E. Scoccimarro et al.). It is made of Echam5 T159L31, OPA 8.2 2◦ global ocean and Nemo 1/16◦ for Mediterranean sea, coupled with the sequential version of Oasis3 2.5. The main goal of the activity is to reduce the elapsed time of the whole coupled model currently in production on NEC SX9 nodes. The NEC SX9 cluster is made of 7 nodes with 16 processors each. In order to achieve the best load balancing among the models, we assigned 20 CPUs for Echam5, 14 CPUs for Nemo, 1 CPU for OPA Global and 1 CPU for Oasis3. With this configuration, the time spent for the coupling takes almost 33% of the entire simulation. Considering that the models are hiding during the coupling transformations, a deep analysis and optimization of the coupler source code could drastically reduce the elapsed time of the whole coupled model. The CMCC_MED model uses Oasis for exchanging a total of 35 fields with a coupling period of 2h 40_ for a total of 279 coupling steps in one month of 31 days.

13.2 Optimisation of the extrap routine

With this configuration, the most time consuming transformations are the extrap and scriprmp. The extrap function performs the extrapolation of the fields over its masked points using the source grid. Since the weights used for extrapolation depend only on the source grid, it is convenient to group the fields into different datasets characterized by the same source grid. Within a given dataset, a field is then tagged through the NIO parameter that defines if the weights must be computed and written to file (NIO=1) or read from file (NIO=0). It is worth noting here that the NIO parameter is taken into account only for the first field of a given dataset, then for all of the other fields belonging to the same dataset its value is ignored; for these fields, the weights are always read from memory. With an assertion analysis it is evident that the weights written into the nweigth auxiliary file are never read; thus we can safely delete the file writing. However improvements were very limited since the writing into the file is performed only for the first coupling step.

The performance analysis of the extrap function highlighted also some numerical problems due to the replication of the same source code on two different branches. In particular, the extrapolation of the first fields of the dataset is performed during the evaluation of the weights; all of the other fields are extrapolated using a different branch. Unfortunately the compiler optimizes the two branches in different way, introducing some optimizing transformations for the floating-point operations. The experiments showed that if we change the order position of a field into the namcouple file, its values, after the extrapolation, differs with an order of magnitude of 10E−14% that is generally negligible. However, if we change the order of more then one field, this displacement produces a difference on the netcdf output files at the end of one month simulation of 0,25%. It is relevant to underline that this discrepancy is only due to a different order of the fields into the namcouple file. The adopted solution consists on splitting the evaluation of the weights from the extrapolation itself. In this case all of the fields will be extrapolated using the same piece of code avoiding differences between the first field of a dataset with the others.

13.3 Optimisation of the scriprmp routine

The scriprmp routine implements the interpolation techniques offered by Los Alamos National Laboratory SCRIP1.4 library. In particular, it performs a remapping of the fields using weights and addresses values that are evaluated taking into account the source grid, the target grid, the type of interpolation to be used and the normalization option. For each field, the scriprmp function checks if the file containing the remapping weights exists. If not, they are evaluated first and then written into file for the further coupling steps. At each coupling step an access to file is performed. The main optimization introduced within the scriprmp is the management of the remapping weights into the main memory in order to reduce the time spent for I/O operations. The optimization reduced the elapsed time for the scriprmp function of 40%. All of the optimizations introduced on both the scriprmp and the extrap functions reduced the whole coupling time of 27%.

13.4 Parallelisation on a field-per-filed basis

In order to further reduce the elapsed time of the coupling transformations, a parallel approach to the algorithm has been developed. The adopted solution consists on the distribution of the fields among the available processes using MPI library. Each Oasis process is then in charge to compute all of the foreseen transformations for a given field. The implementation is driven by two main factors: to balance the load among the Oasis processes; to reduce the communications at minimum.

For the first point it is necessary to consider that a given field could require different transformations from another one; moreover the fields are also defined on different grids at different resolutions; this imply that the computing time for a field can not be considered equal to the computing time for the others. Even if the best scheduling approach is a dynamic allocation of the fields to the available processes, this choice introduces an overhead of the same order of magnitude of the computing time; for this reason a static scheduling algorithm has been implemented. The fields are allocated to the processes taking into account the sequencing index (SEQ), and the correlation among fields that can happen when a field is a linear combination of other fields (BLASNEW and BLASOLD transformations). The proposed approach does not foreseen MPI communications among Oasis processes unless the transformation of a field depends from others. In order to avoid communications, the scheduling algorithm aggregates those fields that depend on each other and assigns them to the same process. At each coupling step, the master process of Oasis gets the fields from the models and scatters them to the slaves according to the distribution policy established by the scheduling algorithm.

Each slave performs the coupling transformation on the assigned fields and sends them to the master that exports them to the models. The parallel algorithm has been evaluated from a performance point of view. The analysis of scalability has been made with the configuration used for the CMCC_MED coupled model on NEC SX9 nodes. The analysis showed that the algorithm reach a 50% of efficiency with 13 processors. The performance model demonstrated that the scalability is heavily limited only by the coarse grained parallelization; the communication overhead takes almost the 2% of the computing time. The parallel implementation has been verified with a bit-to-bit comparison against the output got from the original OASIS3 version after a 2 month simulation using restart files. The current version has been tested only on a subset of the whole available transformations, namely with those ones used for the CMCC_MED model:

- Time transformations: LOCTRANS, AVERAGE
- Pre-processing transformations: MASK, EXTRAP, NINENN, INVERT
- Interpolation transformations: SCRIPR, DISTWGT, CONSERV, BILINEAR, BICUBIC
- Cooking stage: CONSERV, GLOBAL, BLASNEW (only CONSTANT)
- Post-processing transformations: REVERSE

A qualitative comparison between the proposed approaches with respect to the pseudo-parallel implementation of Oasis described in 1.3 has been analyzed. In the pseudo-parallel approach, each Oasis process must have its own namcouple file carefully created by the modeler. Each process is then independent and unaware of the existence of others and it communicates directly to the models exchanging the fields included into its namcouple file. Such approach implements a distributed communication with the models avoiding the bottleneck represented by a single master process in charge to exchange the fields with the models and to coordinate the slaves and the need of a huge quantity of memory. The main disadvantage of the pseudo-parallel approach regards the configuration; indeed, the user is charged with the burden of creating namcouple files each time the number of Oasis processes changes.

14. First steps to include unstructured models to OASIS4

K. Fieg, W. Hiller (AWI)

The demand to include models defined on unstructured grids into an existing system of conventional model components using structured grids is increasing. Since January 2009, the German governmental funded project “ScalES” focuses on the technical part of that coupling and gives the opportunity, to study scalability of earth system models (ESM), exemplified by COSMOS, on many- and multiprocessor systems.

We aim to couple the finite element ocean model FEOM using the fully parallel coupler OASIS4, because for highly parallel applications (up to 2000 CPUs), using a serial coupler would become the bottleneck.

Coupling unstructured grids will cause divers technical and physical problems and following questions arise:

1. The administrative overhead needed to handle the geometry of an unstructured grid model makes a big difference between structured an unstructured grids. The grid used in FEOM consists of triangles, defined by nodes. A so called “Connectivity Matrix” defines the neighborhood of triangles by addressing the nodes.
2. Metadata standards for unstructured grids have to be defined.
3. The ratio between matching source and target grid points can vary significantly. Therefore, new performing search algorithms and interpolation schemes have to be implemented into OASIS4.
4. The ratio of source to target process of the different climate components can vary significantly. The mapping of partitioned grid geometry can become extremely complicated and lead to a strong load imbalance.
5. Coupling unstructured grids to structured grids on a large scale base was never tried. So strange and unpredictable things can happen!
6. The scientist should be supported to find an adequate resolution in time and space for the coupling, so that the relevant physical details don‘t get lost.
7. Mass conservation has to be secured and model drift avoided, even if the element shape and the land / ocean interface is extremely different.

At the end, the coupling efforts for unstructured grids will have to deal with mainly the same class of problems as coupling structured grids, but in a more “extreme” way. As a first step, we plan to couple our unstructured grid to an intermediate structured two dimensional lon/lat grid. This will be finished until end of the year. Later, the coupler interface as well as the xml communication structure will be extended to unstructured grids by amending unstructured grid nomenclatures. Furthermore, we plan to integrate an additional mass conserving interpolation routine (CISL) into OASIS4.

List of participants

André, Jean-Claude	Centre Européen de Recherche et Formation Avancée en Calcul Scientifique (CERFACS), France
Bergman, Tommi	Center for Science Ltd. (CSC – IT), Finland
Bi, Daohua (Dave)	The Centre for Australian Weather and Climate Research / Commonwealth Scientific and Industrial Research Organisation (CAWCR/CSIRO), Australia
Budich, Reinhard	Max-Planck-Institut für Meteorologie (MPI-M), Germany
Caubel, Arnaud	Institut Pierre Simon Laplace (IPSL), France
Coquart, Laure	Centre Européen de Recherche et Formation Avancée en Calcul Scientifique (CERFACS), France
Deque, Michel	Météo-France/ Centre National de la Recherche en Météorologie (CNRM), France
Epicoco, Italo	Centro Euro-Mediterraneo per i Cambiamenti Climatici (CMCC), Italy
Fieg,Kerstin	Alfred Wegener Institute for Polar and Marine Research (AWI), Germany
Fladrich, Uwe	Sveriges Meteorologiska och Hydrologiska Institut (SMHI), Sweden
Fogli, Pier Giuseppe	Centro Euro-Mediterraneo per i Cambiamenti Climatici, (CMCC), Italy
Gayler, Veronika	Max-Planck-Institut für Meteorologie (MPI-M), Germany
Giraud, Luc	Institut national de recherche en informatique et automatique (INRIA), France
Hill, Richard	Met Office, UK
Kornblueh, Luis	Max-Planck-Institut für Meteorologie (MPI-M), Germany
Latour, Jean	Centre Européen de Recherche et Formation Avancée en Calcul Scientifique (CERFACS), France
Luo, Jing-Jia	Frontier Research Center for Global Change /Japan Agency for Marine-Earth Science and Technology (FRCGC/JAMSTEC), Japan
Maisonnave, Eric	Centre Européen de Recherche et Formation Avancée en Calcul Scientifique (CERFACS), France
Marti,Olivier	Institut Pierre Simon Laplace (IPSL), France
Masson, Sébastien	Laboratoire d'Océanographie et du Climat: Expérimentation et Approches Numériques / Institut Pierre Simon Laplace (LOCEAN/IPSL), France
Mc Kinstry, Alastair	Irish Centre for High-End Computing (ICHEC), Ireland
Mocavero, Silvia	Centro Euro-Mediterraneo per i Cambiamenti Climatici, (CMCC), Italy
Mogensen, Kristian	European Centre for Medium range Weather Forecast, ECMWF
Pel, Vincent	Cray, France
Puri, Kamal	The Centre for Australian Weather and Climate Research (CAWCR), Australia
Redler, Rene	NEC Laboratories Europe – IT Research Division (NLE-IT), Germany
Ritzdorf, Hubert	NEC Europe Ltd, Germany
Rosenhauer, Mathis	Max-Planck-Institut für Meteorologie (MPI-M), Germany
Sanna, Antonella	Centro Euro-Mediterraneo per i Cambiamenti Climatici, (CMCC), Italy
Sasaki, Wataru	Japan Agency for Marine-Earth Science and Technology (JAMSTEC), Japan
Scoccimarro, Enrico	Istituto Nazionale di Geofisica e Vulcanologia / Centro Euro-Mediterraneo per i Cambiamenti Climatici (INGV/CMC), Italy
Sevault, Florence	Météo-France, France
Sterl, Andreas	Koninklijk Nederlands Meteorologisch Instituut (KNMI), The Netherlands
Valcke, Sophie	Centre Européen de Recherche et Formation Avancée en Calcul Scientifique (CERFACS), France
van Noije, Twan	Koninklijk Nederlands Meteorologisch Instituut (KNMI), The Netherlands
Voldoire, Aurore	Météo-France, France
Wang, Shiyu	Irish National Meteorological Service (Met Eireann), Ireland
Yang, Shuting	Danish Meteorological Institute (DMI), Denmark