OASIS Dedicated Support
4th annual summary

E. Maisonnave

TR/CMGC/19/149

Abstract

Following the original IS-ENES-1 program, a dedicated support is provided again to the climate modelling community in Europe, in order to set up, upgrade or enhance coupled systems based on OASIS. After a selection process, ETH Zürich, MetOffice and GEOMAR Kiel laboratories were granted with a total of 3 person-months. During this support, we could upgrade OASIS3-MCT to the current version 4 and make available new coupler functionalities in the coupled systems, such as the parallel computing of interpolation weights. Interfaces are modified to allow single precision computations (ETHZ), concurrent coupling of ocean and ice (MetOffice) or full ocean zoom coupling (GEOMAR). A significant performance improvement is always obtained. Set up configurations are already used for studies (GEOMAR, ETHZ) and our modifications saved in community repositories (FOCI, NEMO). This should facilitate the diffusion of our work to a larger number of laboratories (e.g. AWI, and COSMO & NEMO users) and contribute to the spreading of coupled modelling in our community. In addition, one can find in Appendix the detailed Carbon footprint of this work, with the hope that it could contribute to understand how to make the best benefit of our infrastructure in a sustainable way

Table of Contents

Introduction

The Horizon 2020 European infrastructure project IS-ENES3 (2019-2022) is organising the extension of the existing OASIS support (hotline, training) to a dedicated support, at user site. The 6^{th} work package, named "Services on European ESMs and Software Tool", proposes to provide a technical help to design, upgrade or enhance the implementation of OASIS3-MCT interfaces in models and/or set up a tailored and computationally efficient coupled system. In 2019, a total of 3 person-months of Dedicated User Support was offered to 3 different groups. This service excludes the scientific development tuning, analysis and evaluation of the coupled model components themselves and of the coupled model as a whole, i.e. any comprehensive geophysical study in link with the implemented coupling.

Applicants had to briefly describe their project and their needs filling a questionnaire, available on-line. The collection and a preliminary technical analysis were done at CERFACS and transmitted to a panel of the OASIS Advisory Board members, that had to make the selection, taking into account:

- The originality of the problem: e.g. new physics (ice sheets, hydrology, atmosphere/ocean boundary layer, regional modelling, ...), increased task parallelism (extraction and concurrent running of sub-components e.g. sea-ice), etc.
- The quality of the methodology proposed
- The expected scientific impact of the target coupled system and its long-term support by the applicant group
- The opportunity of development of cooperation with communities outside ENES
- Potential training aspects for new or young users
- The synergy with the OASIS3-MCT development plan

The panel selected the 3 proposals appearing to bring clear benefits to a wide community, and for OASIS testing and demonstration: ETH Zürich, Met Office and GEOMAR Kiel. The panel also suggested that the support should work mainly on maintainable aspects useful for the community. Discussion followed on possible clarification to bring to the call, to help the panel making a more informed choice in the future. Improvements in the call could include:

- more details about the task complexity, to decide whether help is needed or not
- what could be beneficial for the coupler itself in the requested support
- more information about the community that is supposed to use the newly built coupled system: what goes beyond a single application for a given institution
- what is support actually helping

Conclusions of the panel meeting were transmitted to IS-ENES head, which validated the selection.

In the following report, a detailed description of the three technical collaborations is provided, together with their main results. In appendix, we also tried to briefly document the practical issues that have to be addressed to allow such kind of collaborative work. Last but not least, a summary of its costs, and particularly its carbon footprint, is proposed.

Mission #13[1] : ETH Zürich (Switzerland), Land-Climate Dynamics group

June 24 – July 19, 2019

Main Goal: Upgrading and performance enhancement of atmosphere-land coupled model on GPU based supercomputer

---

Summary

To take a maximum benefit of the GPU compliant COSMO model, the last version of OASIS and ClandM were included in the coupled system, the existing OASIS interface in COSMO was modified to allow single precision compiling, and a coupling between heterogeneously (PGI/Intel) compiled components was set up successfully. All together, this upgrade leads practically to the multiplication by a factor 1.5 of the COSMO-CLandM coupled system speed

---

COSMO-Climate Limited-area Model	Community Land Model (CLM, here CLandM)
v5 of the regional atmosphere model, https://www.clm-community.eu	v5, as part of the CESM v2.0.0 coupled framework, http://www.cesm.ucar.edu/
From 6 to 2Km, 383x328 (LR) to 801x801 (HR), 60 vertical levels	From 25 to 2Km, 288x128 (LR) to 1332x811 (HR)
For reference description of the GPU version see [1]	For reference description of the land model component only see [2]

```
Supercomputer:

"piz daint", CSCS, Manno, Switzerland

5704 nodes of 12 cores - Xeon E5-2690v3 12C 2.6GHz, and 1 NVIDIA
Tesla P100 - Memory per node: 64 Gb

https://www.top500.org/system/177824
```

1 Like a popular Sci-Fi movie trilogy, the dedicated support numbering does not start with one. For a better explanation, see Maisonnave, E., Valcke, S. and Foujols, M.-A., 2013: OASIS Dedicated User Support 2009-2012, Synthesis,Technical Report,TR/CMGC/13/19, SUC au CERFACS, URA CERFACS/CNRS No1875, France

Model description

The reference coupled configuration is derived from the first implementation of an OASIS based coupled system developed at ETH Zürich [3]. The atmosphere is represented by the climate compliant version of the DWD-MeteoSwiss COSMO limited area operational model. This model was recently fully re-written to be efficiently handled on GPGPU architectures [1]. A Domain-Specific Language (STELLA [4]) was used to hide the complexity of C++/CUDA coding to the end-users and efficiently perform the dynamical core computations on GPU. In the 5th version of COSMO we used, only I/O and OASIS coupling are still performed on one single core of the CPU host.

The stand alone version of the model leaves unused 11 or the 12 CPU of the node. The OASIS coupled configuration takes benefit of these idled resources to perform the calculations of an alternative land surface model (CLandM), before disabling the existing TERRA land surface subroutines of the COSMO model. CLandM is the land component of the NCAR Earth System (usually named CLM). A previous modification of the CESM structure (see [5]) ensured the two-way communication of surface fields between the two models across the internal CESM coupler and atmosphere forcing modules. CLandM MPI processes are mapped to the 11 available CPU cores of the node and coupled via OASIS to the COSMO MPI process also located on one CPU core.

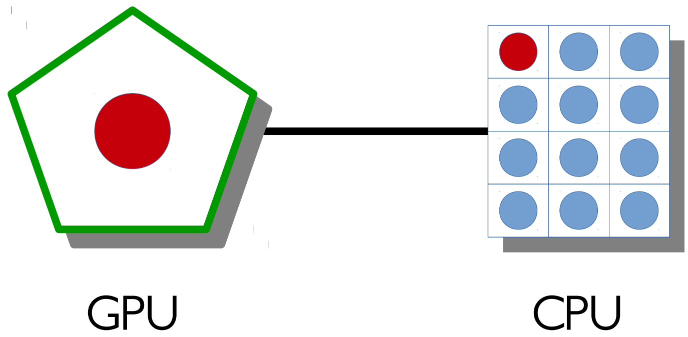

Figure 1: COSMO (red) and CLandM (blue) task binding on GPU (CUDA threads) and CPU (MPI processes)

To take the maximum of performance of GPU resources, the PGI compiler v18.10.0 was used, along with the MPI cray-mpich/v7.7.2 library.

The replacement of TERRA original soil model by CLandM proved its relevance [6] and this additional value is obtained with the same amount of energy, considering that CPUs are still consuming energy while standing. However, the sequentiality of COSMO/CLandM calculations necessarily increases the restitution time compared to COSMO stand alone (even including TERRA subroutine cost).

Upgrade

Changes in both so-called OASIS interfaces on COSMO and CLandM (the set of model subroutines we modified to include OASIS API calls and pre/post processing of the variables needed for coupling) are made necessary by respectively the compliance with single precision computing and a version upgrade. Additional modifications were required (i) to test the Intel compiled version of CLandM, (ii) to allow a fitted definition of the land point mask of this model, (iii) to make possible a concurrent instead of sequential computing of both model and (iv) to provide additional coupling fields in CLandM for future improvement of the momentum flux calculations.

In parallel of these necessary improvements, we decided to upgrade the OASIS library (OASIS-MCT v4 release [7]), keeping in mind the possible benefits of the release (e.g. parallel computations of interpolation weight and addresses in case of further resolution increase). The official release still cannot be used without being adapted for compliance with MCT library included in CESM (see [8]).

CLandM version update

Since a new version of NCAR land model was released for CMIP6 [2], an update of the ETH coupled model is planed to take benefit of the component enhancements. The CLandM OASIS interface implementation, while non intrusive, was designed to exchange information not with the land model directly but through the whole CESM structure, since there is no more stand alone CLandM version available as a CESM independent model. The OASIS coupling fields are transiting from/to the CESM module, which is supposed to read forcing fields (`datm`), and also from coupling routines (`cpl`), which are supposed to transfer variables from/to this module to/from the land surface module. This is why several routines of these libraries, and not those belonging to the land surface module, have to be modified in the new CESM release, to keep ensuring the OASIS coupling.

A particular care must be taken to avoid an overwriting of the incoming coupling fields by forcing values during the restart procedure. In the forcing routine `seq_io_read_avscomp`, the 19 variables previously modified by the incoming coupled values are not taken from the forcing module restart but kept unchanged after their calculation in the coupling interface routines.

Despite modifications on several CESM routines involved in the OASIS coupling, e.g. changes of FORTRAN structure or variable names, it only took a couple of hours to update the CLandM coupling interface without increasing the implementation complexity or intrusiveness. One would imagine how different would be a complete rewriting of new CLandM subroutines (extracted to the CESM structure) directly in COSMO.

COSMO single precision compliance

The full renovation of the COSMO model to make it running efficiently on new architecture such as GPU also included a validation of the model calculations exclusively using single precision floating variables. This remarkable and probably painful effort is particularly fruitful, since it gives an approximative 20-30% speed up at various resolutions. Such performance improvement makes mandatory to include this COSMO single precision version in our coupled system.

To do so, it was not necessary to modify or even recompile our OASIS library: the `oasis_put` and `oasis_get` FORTRAN modules, called to couple model variables, are able to handle both single and double precision real arrays. Considering the low number of exchanged coupling fields in the COSMO-CLandM coupled system, and the weak contribution of coupling procedures to the simulation restitution time, we did not feel necessary to try modifying the OASIS library itself to make possible the handling of single precision arrays during interpolation (MCT matrix multiplication) and communication (MPI) phases. This work is postponed, at least until that a more demanding coupling (e.g. 3D coupling) would be required or when an OASIS based coupled system will be operated on single precision only architectures.

Nevertheless, despite this OASIS API compliance, the first test with COSMO single precision model in coupled mode leads to an immediate crash of the atmosphere model. A quick check of the incoming coupling fields validated the OASIS exchange procedure, but the examination of the COSMO variables modified by the incoming coupling fields revealed non representable values. We concluded to the non compliance of the OASIS interface of COSMO with single precision calculations. In particular, we identified the possibility of insufficient precision in the calculation of potential temperatures, which has bad consequences within the equation of the short wave flux transfer coefficient:

$$\text{zp\_ref} = (1.e^{-5}\,.(\text{P0+Pp})\,)^{\text{rdocp}}$$
$$\text{zp\_surface} = (1.e^{-5}\,.\,\text{Ps})^{\text{rdocp}}$$

$$\text{transfer\_coeff} = -\ \left(\frac{\text{short wave}}{(\text{C.Ps}).(\text{T\_surf} - \text{zp\_surface}\ .\ \text{T\_air} / \text{zp\_ref}\,)}\right)$$

Low values of `rdocp` and `P0+Pp` could lead to zero value of the temporary `zp_ref` variable, when using single precision floats. In this case, the division by `zp_ref` in the transfer coefficient equation stops the simulation. The solution here simply consists in evaluating the ratio of pressure before calculating potential temperature coefficient:

$$\text{transfer\_coeff} = -\ \left(\frac{\text{short wave}}{(\text{C.Ps}).(\text{T\_surf} - (1.e^{-5}\ .\ \text{Ps}/(\text{P0+Pp}))^{\text{rdocp}}\ .\ \text{T\_air}\,)}\right)$$

Only the operation order is modified. The solution has the advantage of not modifying the model results.

A second set of crashes occurs after several months of simulation at low resolution (but after a few time steps at HR). Despite our modification, the transfer coefficient used by the model to rebuild fluxes with soil/air temperature differences could have strong values. This occurs, randomly, when the temperature differences are below a threshold of 1e-7 K. We proposed to limit the temperature differences between soil and first atmosphere level, adding or removing the 1e-7 value to the difference, which is equivalent to a limitation of the transfer coefficient. A second test of this coefficient is also made and the value set to zero if a limit of 2,000,000 is reached.

This solution, while non conservative, allows to perform a 1 year long LR-simulation and a 1 month long HR-simulation. At LR, temperature difference is corrected less than 100 times per month and practically no call to the second limitation is observed. At HR, the 1st correction is done approx. 2000 times, the second about 10 times. A more careful diagnostic of the flux non conservativeness should be carried, but we assume than it is less important than, for example, any interpolation error of the surface temperature between CLandM and COSMO grids.

Heterogeneous compilation of coupled components

PGI compiling of the COSMO model is necessary to get performances. Actually, comprehensive efforts were made with PGI on `daint` to make efficient the use of GPU accelerators. However, better restitution time are observed if CLandM is compiled using Intel. The obvious question of an heterogeneous coupled simulation, with respect to compiler, rose immediately.

To gradually address the question, we compiled two versions of the OASIS library with the two compilers and tried running the toy coupled system (`test_interpolation`) included in the OASIS release, each one linked with a different OASIS library. The non intuitive success of this test probably relies in the fact that the same MPI library version, while also compiled with a different compiler, is statically linked to both executable.

On a second step, the two COSMO and CLandM model executables are prepared similarly, i.e. linked respectively with PGI and Intel compiled OASIS libraries. Again, simulations were performed successfully, preserving performance of the stand alone models. Performance gain of the coupled model is presented in section "Performance".

CLandM grid point definition according to COSMO grid boundaries

An important issue related to the coupling of two regional models differently discretised on the sphere is the mismatch of their latitude and longitude limits. The direct consequence is that some grid point mesh of one at least of the two models will not intersect any neighbour mesh of the other grid in the calculation of interpolation. The implemented solution for COSMO-CLandM consists in :

1. defining the CLandM domain limits in such a way that COSMO domain is included in the CLandM area.
2. performing a first coupled simulation to interpolate the COSMO land/sea mask on the CLandM grid.
3. redefining accordingly the active grid points of the land model.

This solution not only ensures that interpolations are mostly done between source/target meshes intersections (without any extrapolation from remote positions) but also reduces the number of CLandM active grid points and, consequently, speeds up its restitution time.

The second phase is made possible by a special procedure implemented in both model interfaces and activated via the change of a integer variable of the codes (`IOASISDEBUGLVL=2`). A similar change (and a similar recompiling) is required at OASIS level, to neutralise the extrapolation procedure in case of lack of source/target mesh intersection. The result of this second phase (a new land/sea mask in the CLandM grid) can be not fully satisfactory and may require an additional manual choice of masking. The CERFACS' graphical tool made for this purpose was modified during the support period [1] to avoid using X11 tools not always available on supercomputers.

COSMO/CLandM concurrent computations

The computation sequence of coupled model component depends on the organisation of the coupling fields exchange. For example, if, in model-1, the reception of field-A immediately follows the sending of field-B and if, in model-2 the calculation of field-B depends on field-A reception, model-1 will necessarily wait model-2 before resuming its computations.

To exactly reproduce the calling sequence of TERRA original land model in COSMO, the coupling field exchange is made in such a way that COSMO waits the result of CLandM computations and vice versa. The two components are then running sequentially and half of the computing resources are wasted in waiting the results of each other. A simple check of the `LUCIA` load balancing tool [1], enabled via the OASIS parameter file, validates this hypothesis.

A new organisation of the coupling field exchange is coded in the OASIS interface of the COSMO model. To allow simultaneous computations in COSMO and CLandM, the first call of coupling exchange is changed from a send to a receive. This makes COSMO waiting the first set of incoming coupling field. An appropriate modification of the OASIS parameter file is setting a reading of these fields in a restart file and their delivery to the COSMO model. Symmetrically, the COSMO sent fields are saved in a restart file at the end of the simulation, via an additional call of the OASIS send API in this model.

This operation obviously changes the simulation results, since COSMO coupled variables computed at `n-1`, and no more `n`, are received by CLandM at time step `n`. The robustness of this change has to be fully tested with a long term simulation and the impact on model equilibrium must also be estimated. But the expected effect of this modification on performance is the levelling of the coupled system speed to the slowest model one.

Complementary work

As previously [1], a comprehensive User Guide was delivered at the end of the support period to help ETH scientists to reproduce the different steps necessary to set up a new coupled configuration and launch simulation at optimum speed. This document can be provided on demand.

In addition to the modelling modifications provided by the CLandM new release, ETH scientists could decide to better address the question of the numerical stability of fluxes computation at the land surface. For this purpose, two new quantities (Surface friction velocity and Aerodynamical resistance) are carried from the land surface model to the CESM coupling interface and will be easily available when the COSMO interface will be redesign on purpose.

Performance

The speedup resulting from the 4 improvements described above is measured during a 2 days long simulation, using the LR version of the model on 9 nodes, in a close to the optimum parallel decomposition. Initialisation and termination phases of the execution are excluded. The results are summarised in Table 1. Variability of restitution time, not fully estimated, is below 5%.

	COSMO timing (s)	CLandM timing (s)	Coupled model speedup (%)
reference	530	32.	
gpts bounding	530	8.2	4 %
single precision	305	8.2	41 %
Intel compiling	305	3.9	1 %
Overall improvement (sequential) :			45 %
concurrent cpl	305	0.1	46 %

Table 1: elapsed time of 2 simulated day long run of COSMO-CLandM-LR coupled system, after modification of the reference configuration by (i) bounding the CLandM active grid points to COSMO domain limits, (ii) changing for single precision real variable in COSMO, (iii) compiling CLandM with a different compiler (Intel) and (iv) performing both COSMO and ClandM calculations concurrently

The 2D CLandM computations, also performed at lowest resolution than COSMO (25Km vs 6Km), run much quicker (1 order of magnitude). The consequence is that the improvements made on the CLandM side have a minor impact to the overall coupled system speed. The COSMO calculation precision effect is as expected close to a 40% increase of the speed. The same, even if smaller, effect is obtained with the HR configuration (20%).

However, a load balanced configuration will better benefit from the CLandM related improvements. The Intel compiling can potentially divide by two the CLandM restitution time. The removal of off limit grid points on reasonably well fitted latitude-longitude domains can also significantly speed up the land model. Finally, the concurrent coupling mode, which physical results have to be preliminarily validated, gives the possibility to use the idled CPU of a COSMO-GPU model without any additional cost, nor from respect to energy or restitution time. In its HR configuration, the coupled system speed exhibits a 10% increase.

Mission #14 : MetOffice Exeter (UK), Climate Science IT Applications

August 26 – September 20, 2019
Main Goal: Setup an OASIS coupling between NEMO ocean and SI3 sea ice components

---

Summary

Coming with the new NEMO 4.0 version, the recent upgrade of the sea-ice component from LIM to SI3 makes necessary a check up of the ocean/surface_module coupled interface. Few code modifications, included in a development branch for later trunk update, were necessary to perform test simulations at ORCA1 and ORCA12 resolution and roughly check its validity. Improvement of NEMO speed and cost is real but limited to 10 to 20% and observed with sufficiently high decomposition only. At its best, our coupled configuration is faster (x2) and cheaper (-25%), but, since it is spread on a larger number of resources, it could reduce the actual speed (simulation + scheduling time) of production runs

---

NEMO, ocean (OPA)	NEMO, sea ice (SI3)
v4 of the global ocean model, https://www.nemo-ocean.eu	Included in the surface module (SAS) together with flux computations and icebergs
From ORCA1 to ORCA12, 31 to 75 vertical levels	Same resolutions, same grid, but possibly different decompositions
For reference description of the whole model see [12]	

```
Supercomputer:
CRAY XC40, “xce”, MetOffice/Science Park, Exeter, UK
2496(+6636) nodes of 36 cores - Xeon E5-2695V4, 18C, 2.1GHz
https://www.top500.org/system/178925
```

Model description

The NEMO ocean model is currently included in several configurations of coupled models, for climate modelling but also operational purposes. A large range of resolutions from ORCA1 to ORCA12, and coupled components (from atmosphere-ocean only to ESM components) makes necessary the modularity of the OASIS coupler and the choice of interpolation offered by the SCRIP and ESMF libraries.

In various contexts (CMIP6 exercise, operational forecasting …), speed and cost are more than crucial quantities that this support proposes to deals with. Met Office IT group identified the ocean/sea-ice interface as a potential source of computing performance enhancement. As demonstrated e.g. in [13], the separation of the ocean surface module (SAS), including sea ice model, from the ocean related routines of NEMO (OPA) in two executables, and their coupling via OASIS, allows the concurrent performing of the two set of calculations, increasing the model speed and, in some appropriate conditions of load balancing, reducing its cost.

The recent upgrade of the sea-ice component from LIM to SI3 makes necessary the check up of the OPA/SAS coupled interface. Component scalability changes resulting in particular from [14] may have jeopardized the expected advantages of the OASIS coupling. This is what this study proposes to verify.

Ocean-Surface module coupling enabling

NEMO

The separate compiling of the two SAS and OPA executables is still functional in NEMO 4.0 version without any additional modification. It simply consists in compiling one of the ORCA reference version only using `OCE` sources and activating the CPP `key_mpp_mpi key_iomput key_oasis3` on one hand, and the SAS reference version using `OCE`, `SAS` and `ICE` sources and activating the CPP `key_si3 key_iomput key_mpp_mpi key_oasis3` on the other hand. Both executables must be launched simultaneously (as XIOS server, if needed) thanks to the MPMD syntax of the CRAY `aprun` tool. No modification either to the `namcouple` OASIS parameter file available with NEMO 3.6 version, that defined the ocean/sea-ice coupling fields and their exchange frequencies. However, one will think to modify the `iodef.xml` XIOS configuration file to (i) allow the simultaneous use of XIOS and OASIS[2] and (ii) the concurrent service of output by the XIOS servers of both OPA and SAS executables[3]. This latter can also require, for a better readability, the splitting of output fields and files definition into separate XML files, by creating a new `context_sas.xml` file.

As observed but not reported in the previous version, the initialisation phase of the coupled model is the most error prone of any other model sequences, especially if made from rest. In

2 `oasis_server` variable to `true`

3 `oasis_codes_id` variable to "`oceanx,sas`"

this case, since model exchanges did not started yet (coupling fields are exchanged during the first time step), surface model variables that require ocean values to be built could be wrongly initialised, and this leads to instabilities that could be fatal.

A special restart procedure must then be followed to avoid this kind of inconsistency. The OPA ocean component does not need any restart file but it is necessary to provide it to the SAS component. This file can be produced e.g. by a short uncoupled simulation starting also from rest. The restarting procedure of the coupled configuration follows the same rules than the reference uncoupled one.

For the same reasons, even though OASIS (since version 4.0) is able to create its own restart files with zero values, it is necessary to provide coupling fields restart files with realistic values. To do so, one could activate the EXPOUT options in `namcouple` to produce NetCDF output files of all coupling fields during a short OASIS zeroed restart simulation, then transform them to produce OASIS restart files with the right format.

Code modification were necessary to avoid a performance leak and clarify the origin (SAS or OPA model) of several log output file produced by the model. In order to let SAS producing, in addition to the ocean module, communication, timing and possibly `run.stat` (ASCII or NetCDF format) report files, a “`sas_`” suffix is added to the corresponding file names and files produced by SAS executable in addition to the original ones (in this latter configuration, only one set of files is produced for both ocean and surface module routines). We also noticed that the concurrent update by SAS and OPA of the time step value in the time.step file considerably damaged the model performance. A separate writing, instead of the SAS routine disabling, is proposed. All code modifications are included in a development branch for later trunk update. We emphasise that these minor changes, mainly affecting the SAS module (coupled version) should not have any impact on other reference versions.

OASIS

Any regular ocean grid includes grid points located over land. In NEMO, all model computations, though meaningless, are also performed at these positions. The MPI subdomain decomposition procedure can attribute to some tasks a set of grid points exclusively located on land. In this case, it is strongly recommended to stop the simulation and launch it again on a smaller number of resources. During the next initialisation phase, no resources are attributed to these “land-only subdomains” and the decomposition process leaves holes in the total global grid. OASIS is supposed to be able to handle such non contiguous decomposition, since it is now possible to define a larger global domain size thanks to an optional argument (`isize`) of the `oasis_def_partition` primitive. Unfortunately, the OPA/SAS coupling requires the reading/writing of coupled fields in restart files, which stops the MCT library. The definitive problem fixing probably requires an internal modification of the library, but cannot be made immediately available (OASIS ticket #2472). In the meantime, MetOffice users could rely on a modified version of the coupler, which simply avoid to perform the test that stops the simulation. In the present report, there is no removal of “land-only” subdomains.

Our tests put in evidence a second weakness of the coupling library. Using the `LUCIA` load imbalance measurement tool [10] during the performance optimisation phase, and more specifically its compiling with CRAY FORTRAN, put in evidence the non portability of EOF detection in `LUCIA` log files. A bug fix was proposed by Mirek Andrejczuk (OASIS ticket #2473).

Performance measurement

Low resolution test configuration

To be able to conveniently handle the model (input/output file processing and visualisation, short number of resources but quick simulations and favourable scheduling on supercomputer) and to keep our results as close as possible as those that could be achieved with a production run, we based our implementation, tests, first performance measurements and validation on the ORCA1 stand alone configuration in use at MetOffice for post-AMIP6 simulations (`MO_GO8`). We kept this configuration as close as possible to the production version. In particular, we let the model, that includes SI3 sea-ice and icebergs, writing a comprehensive set of output variables via XIOS servers at daily pentad and monthly periods. Notice that the SAS independent executable (the whole surface module set of routines `SBC`) not only includes the sea-ice model but also iceberg related routines and flux computations.

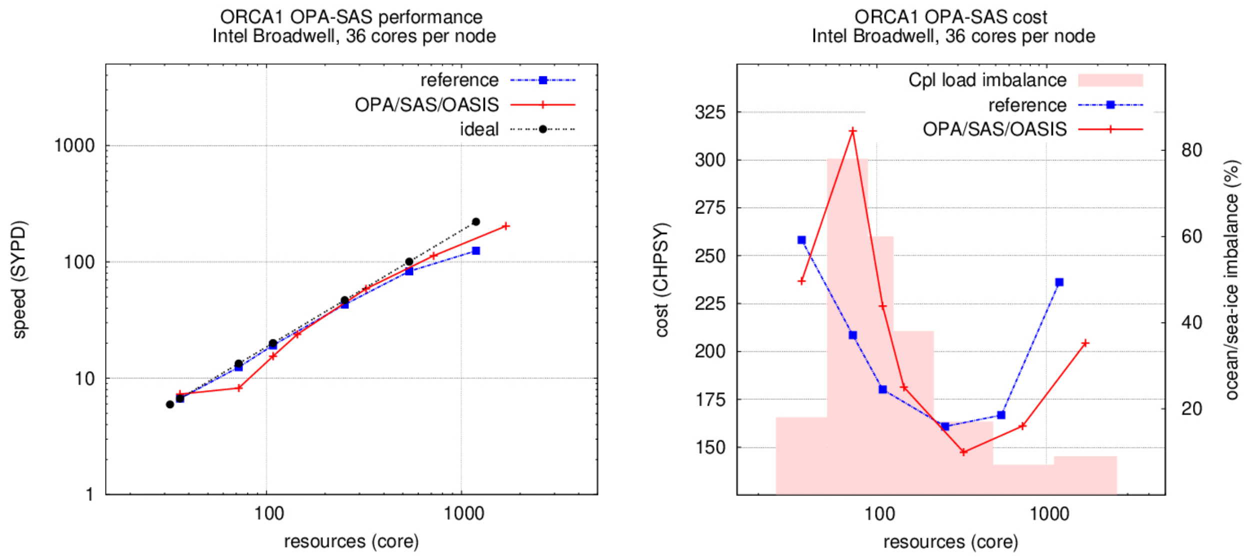

Figure 2: Speed and cost of the NEMO 4.0 model, ORCA1 grid, no land-only subdomain removal, with ocean and surface module (including ice) components coupled via OASIS (red) and in regular uncoupled mode (reference, blue). On right plot, load imbalance of the coupled experiments, i.e. ratio between fastest component waiting time in coupling and total elapsed time (in time loop). 1 node is always allocated to XIOS. In regular mode, 1-2-3-7-15-33 nodes are allocated to NEMO. In coupled mode, the same node number is allocated to the OPA and 1-1-1-2-5-14 additional nodes are allocated to SAS. The first point comes from a single node experiment, which OPA/SAS core allocation is 32/4.

We measured the performance of the ORCA1 configuration with regular and coupled ocean/ice models. On the figure 2, models speed (left) and cost & coupled model load imbalance (right), excluding initialisation/termination phases, can be compared. Resources number is given excluding the extra node allocated to XIOS but including both OPA & SAS executables allocation in the OPA/SAS coupled case. Resources are also allocated to land-only subdomains (see previous chapter) but it does not change the main conclusions concerning performance (a simple quasi linear shift of 1/3 in #resources must be done to extrapolate real speed and cost).

A first problem bounds OPA/SAS advantages where low number of resources are allocated. The narrow list of choices for decomposition, given that one node can only be filled with MPI tasks of the same executable, leads to huge load imbalance, except on a single node[4]. The consequence is that we waste too much resources for SAS and, even if the speed is higher than the reference configuration for the same number of resources allocated to the reference NEMO and the coupled OPA, the speed comparing the total number of resources is smaller, and the cost is higher.

When using a higher number of resources, the load imbalance can be reduced and the coupling slightly improves the speed. Cost also is better, even though the SAS scalability becomes less good and the resource ratio OPA/SAS decreases.

At this point, a closer look to the routine consumption ranking, available in the new `sas_timing.output` file, put the stress on the slowing down of the `step` routine. Based on previous tests with 4.0 reference version, the cost of this routine should be much less prominent. Additional diagnostic would be necessary to find the origin of this limitation, on this machine and for this set of parameter.

A slow down due to the `LUCIA` log file writing was also put in evidence during measurements. This is related, on this machine, to the high dependence of the performance to the disk access speed. Optimal ratio of OPA/SAS allocated resources are calculated with our load balancing tool, but it is disabled to take measurement of speed and cost. We assume that the optimal ratio is not modified by the `LUCIA` enabling/disabling. The simulation slowing down created by the `LUCIA` log file output at runtime put once again in evidence the urgent need for reconsidering the way load imbalance is measured in OASIS. The next rewriting of this tool must take into account the extreme dependency of performance with disk writing and must avoid log file production.

We can conclude that the improvement of ORCA1 speed and cost generated by the splitting and the coupling via OASIS of ocean and sea ice tasks into two executables is real but limited to 10 to 20% and observed with sufficiently high decomposition.

High resolution

The same set of performance simulation, led during 100 time steps, is produced with a higher resolution configuration (ORCA12, 1/12 degree horizontal resolution) of the very same code. Only namelists are modified to switch to appropriate dynamical and physical schemes. No additional correction is necessary in the code to fully perform coupled tests or even longer simulations. A version incompatibility between OASIS and XIOS libraries prevent to take a comprehensive measure of the original configuration at high scalability. Consequently, the results given below do not include the 12 nodes usually allocated to XIOS at ORCA12 resolution.

4 In single node mode, it is possible to mix OPA and SAS related MPI tasks on the node. XIOS is switched off to simplify the test

As previously, the uncoupled procedure reduces simulation speed and increase its cost if a relatively few number of resources are allocated. But this time, the reason why a good load balancing cannot be achieved at lower decomposition is the total memory requirement of the model (the SAS module, which includes most of the memory bound 3D NEMO variables, was impossible to launch on less than 8 nodes).

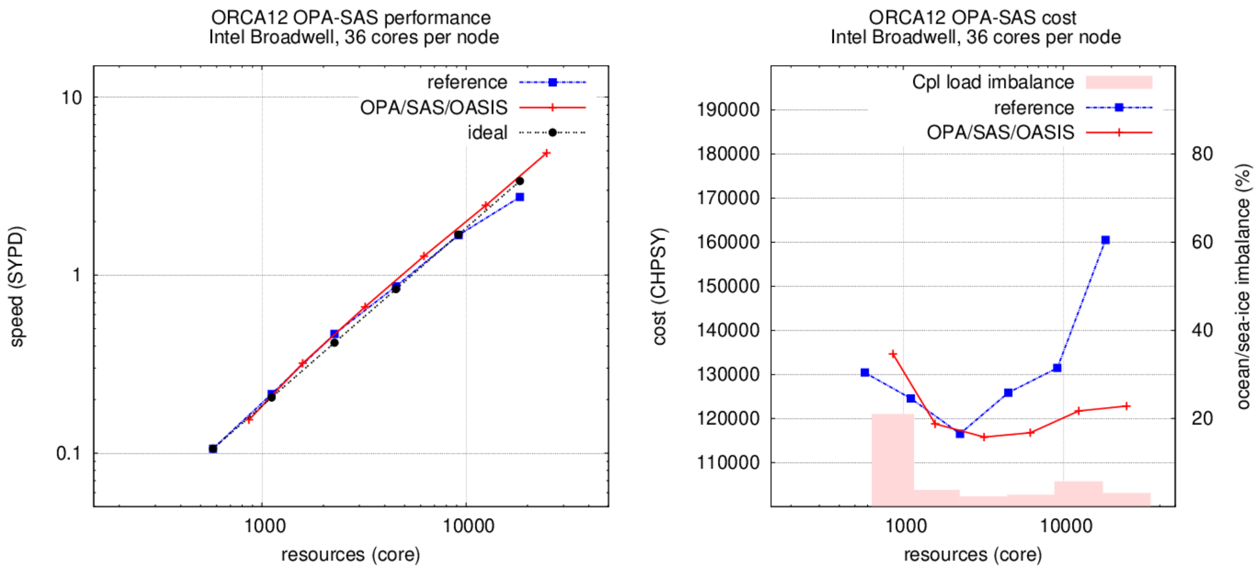

Figure 3: Same than Fig 2. for ORCA12 grid, without XIOS output. In regular mode, 16-31-63-126-255-510 nodes are allocated to NEMO. In coupled mode, the same node number is allocated to the OPA and 8-12-24-44-92-180 additional nodes are allocated to SAS

Using a larger number of resources, the coupled mode slowly become faster and cheaper than the reference configuration. However, this superiority occurs for such a high number of resources that they are usually uneasy to obtain from the batch scheduler. At its best, our coupled configuration is faster (x2) and cheaper (-25%), but its extra resource requirement could reduce the actual speed (simulation + scheduling time) of production runs.

Validation

Three 5 year long simulations are produced with the same experimental protocol (same ORCA1 grid, same namelist parameters, restart from a 5 year long spin up simulation made with the reference configuration).

The regular NEMO configuration (REF), in which surface module is called at each time step, is used for reference. An additional simulation (HLF) is made with a calling period of the surface module (`nn_fsbc`) equal to 2. The third simulation (CPL) is performed with the OPA/SAS coupled model, with models and coupling time steps equal to the reference model time step.

The total ice volumes and the corresponding difference (HFL-REF and CPL-REF) are plotted in Fig. 4. This quantity usually integrates any coupling anomaly in time and helps to put in evidence possible missing link between OPA and SAS (e.g. runoff or calving). There is no clear evidence of such missing coupled quantities in these results. Even though trends can be observed in both Arctic and Antarctic ice volume produced by the OPA/SAS model, their slope is still compatible with the HLF simulation and the global difference compared to the reference remains around 1%. However, a closer look to other surface quantities in polar regions would be necessary to fully validate our coupled configuration.

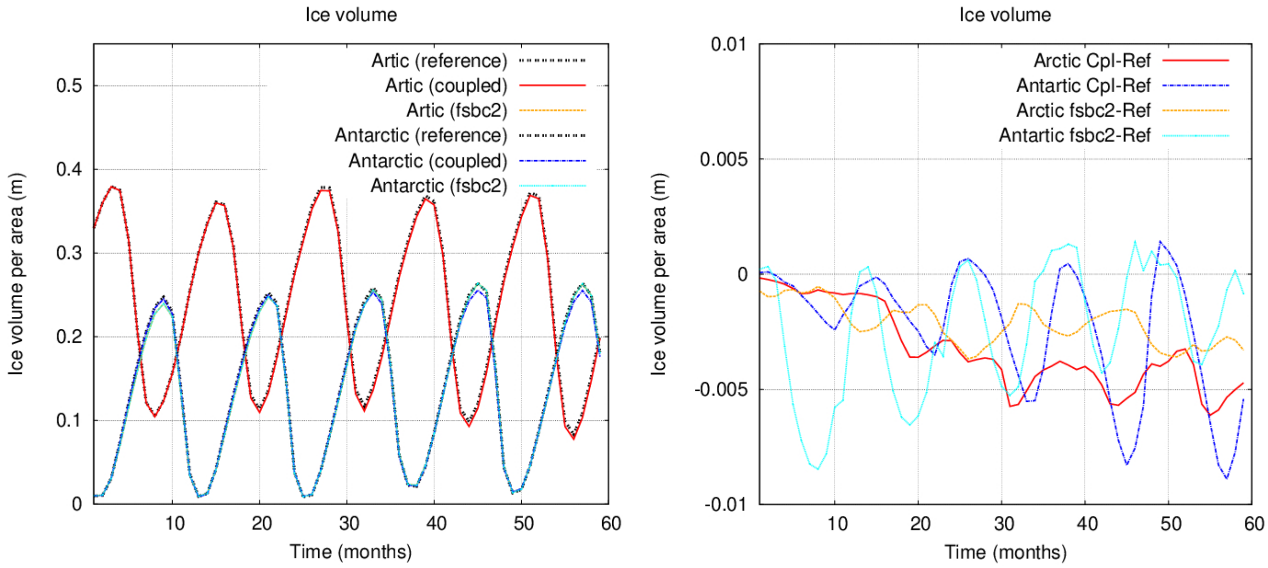

Figure 4: Arctic and Antarctic ice volume for reference version, same model with ice model calling period (`nn_fsbc`) of two time steps and OPA/SAS coupled configuration with coupling at every ocean time step

Further improvements

We found computing performance sensitive to disk access on the `xce` machine, which means that a definitive decision on whether including the coupled configuration in the official laboratory suite must wait a more comprehensive evaluation with various configuration (resolution, namelist parameters ...) Meanwhile, a recording of NEMO code modifications was performed on a MetOffice official repository branch and a short user guide delivered to the first users.

Other improvements can be imagined to further reduce the cost and/or increase the speed of the coupled configuration. First, the SAS subdomain shape could be modified to increase their Y-size in tropical regions, where no sea-ice related computations occur, by modifying their definition in `mppinit` subroutines. Second, a finer mapping of SAS/OPA processes, e.g. a mixing of different model MPI tasks on the same node, could contribute to take a better benefit of the available bandwidth and reduce the cost of the OASIS intra-node communications. Third, to better address suitable scientific questions, resolution of sea-ice (or ocean) could be downgraded. This option was already technically tested in [13] and could easily be activated with our new model version.

We finally mentioned other developments that should follow the present implementation: the addition of the atmosphere coupled model to the OPA/SAS coupled configuration (also tested with a different atmosphere model in [13]) and the same operation of decoupling of the TOP-PISCES biogeochemistry module and its replacement with the MEDUSA model. This interface, widely needed in our community (LOCEAN, Mercator Ocean, Météo-France, GFDL ...) is already implemented [15] without OASIS but would benefit of the modularity that our coupler can offer.

Mission #15 : GEOMAR Kiel (Germany), Marine Meteorology team

September 30 – October 19, 2019

Main Goal: Extending OASIS coupling between OpenIFS and NEMO ocean to an AGRIF zoom

---

Summary

The necessary removal of the on-disk coupling procedure of the ocean zoom surface fields in the GEOMAR OpenIFS-NEMO-AGRIF coupling (FOCI) required the upgrade of both NEMO and OpenIFS interfaces. The inexpensive OASIS coupling that has been set up allowed to increase the OpenIFS horizontal resolution to 25Km. The CPU cost of the coupled system, that includes the North Atlantic zoom AGRIF, is estimated to approximately 30 time less than the CPU cost of the corresponding global ORCA12 based configuration

---

OpenIFS, atmosphere	NEMO, ocean
cy40 of the global model, https://confluence.ecmwf.int/display/OIFS	v3.6 of the global ocean model, https://www.nemo-ocean.eu
From T159 (85Km) to T799 (25km), 91 vertical levels	ORCA05, 46 vertical levels. Includes AGRIF zooms (e. g. North Atlantic)
For reference description of the IFS version from which OpenIFS is derived see [16]	For reference description of the whole model see [12]

```
Supercomputer:

"mistral", DKRZ, Hamburg, Germany

1368 nodes of 24 cores - Xeon E5-2680v3 12C 2.5GHz - Memory per
node: 64 Gb

https://www.top500.org/system/178567
```

Model description

The ocean-atmosphere coupled model currently used at GEOMAR Kiel substantially differs from the standard NEMO based coupled models of the climate community. It is derived from the former Kiel Climate Model [17], including ECHAM atmosphere. The NEMO global ocean horizontal resolution is 0.5 degrees (ORCA05) but it can be increased in selected areas (Southern ocean, North Atlantic …) taking benefit of the AGRIF functionality. In the new FOCI (Flexible Ocean and Climate Infrastructure) configuration [18], ECHAM is replaced by the licensed software of ECMWF IFS, OpenIFS [19]. Our study will focus on two horizontal resolutions: a coarse T159 (~125km) and an accurate T799 (~25km). The global ocean NEMO v3.6 in ORCA05 grid includes a zoom over North Atlantic ocean (VIKING) at 1/10 degree resolution. A runoff remapping tool complements this climate model. NEMO outputs are speed up by the XIOS I/O server (detached mode) but OpenIFS still ensures its own output (serial mode, GRIB format).

This model can be handled, in various configurations (different resolution, different zoom location) thanks to the community working environment ESM-Tools developed at AWI, Bremerhaven by Dirk Barbi [20]. In this environment, several other OASIS based coupled models are available, from which the former KCM model and the FESOM based system AWI-CM [21,22].

The existing FOCI configuration relies on the version 2 of the OASIS3-MCT coupler. Only coupled fields discretised on the NEMO global grid are exchanged through the coupling library. The existing OpenIFS and NEMO coupling interfaces are not prepared to allow the exchange of AGRIF child grid fields. Consequently, AGRIF needs to read its forcing flux condition in a file (NEMO forced mode). This file is updated online by fluxes provided by OASIS, thanks to the EXPOUT option. The main drawback of this solution is the prohibitive cost of disk writing at high resolution. The purpose of this dedicated support is to substitute to this on-disk coupling solution a full OASIS coupling, through standard MPI communications.

Coupling interface of regional grid: implementation

Upgrading OASIS

The newly available version 4 of OASIS3-MCT coupler is substituted to the existing version 2, mainly for 2 reasons:
- the GAUSWGT interpolation (nearest neighbours with Gaussian weighted distances) was chosen by GEOMAR to perform the transformation between every grids of the FOCI coupled system. The bug correction available in the 4th version seemed mandatory to allow an optimum setting of the coupling,
- the hybrid MPI/OpenMP generation of interpolation weights will significantly facilitate the evaluation of the different parameters of the GAUSWGT.

This new OASIS version was included in the ESM-Tools suite. In addition, the `examples/test_interp` toy model was compiled and configured on `mistral` to allow an offline and efficient generation of the interpolation weights. Using more than 32 nodes in MPI/OpenMP hybrid mode, all RMP files were generated in less than 2 minutes, the maximum time reached for 49 neighbours searched in AGRIF grid (869x884 grid points) to fill the T799 OpenIFS target grid points (843,490).

AGRIF-NEMO update

The necessary removal of on-disk coupling procedure requires the upgrade of both NEMO and OpenIFS interface, and the modification of OASIS parameters. As shown in Fig 5, the coupling exchanges must take place not only in NEMO parent grid related subroutines, but also on the AGRIF part of the code. In consequence, coupling fields must be duplicated in the `namcouple` file and the corresponding OASIS restart files must be created. Notice that it is not necessary to call twice the OASIS sending routine in OpenIFS, but simply modify the `namcouple` in order to associate outgoing coupling field names to both NEMO parent and child incoming coupling field names. In addition, longitudes, latitudes and land/sea masks of the AGRIF grid must be described in the OASIS auxiliary files.

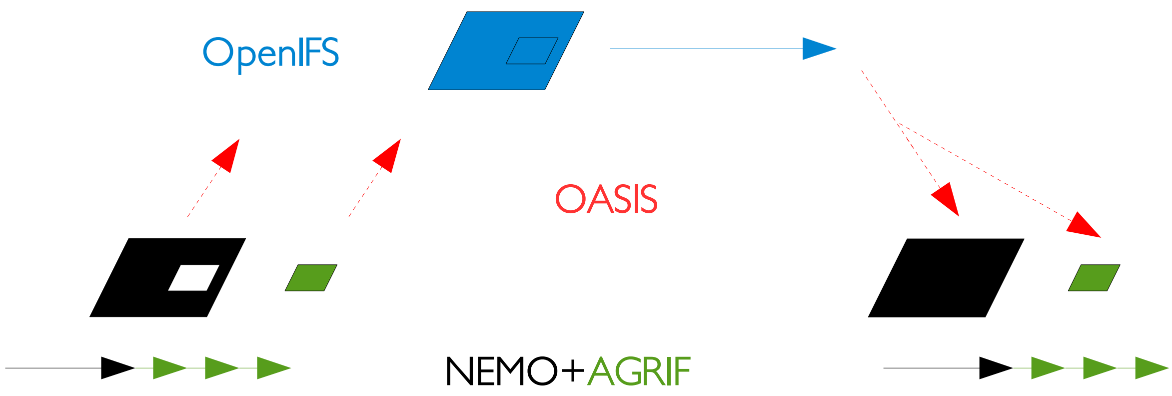

Figure 5: Sequence of coupling exchanges between OpenIFS and NEMO, including an AGRIF zoom

The direct providing of fluxes to the AGRIF subroutines enhances the computing performances of the coupled system. Without putting at risk this enhancement, we can also allow AGRIF subroutines to communicate their surface fields to OpenIFS.

In this case, a special procedure is required to be able to combine, on the OpenIFS grid, two fields coming from two different ocean grids. When OpenIFS receives the two kind of fields, they are already interpolated by OASIS to its own grid. But a combining procedure is required to form a single array that will be used in the model:

1. In AGRIF routines, a variable (`AgrSpg`) is defined to zero with a 4 grid point large zone at the grid boundary, and a linear transition to the inner part of the grid, which is set to one (see Fig 6, left)
2. A new OASIS mask (`agr2`) of the AGRIF grid in defined in auxiliary files, in conjunction with a separate grid set of variables (latitude, longitude, area). This mask is equal to zero (non masked values following the OASIS rules) everywhere, except in a 2 grid point large boundary. It must not include the original land/sea mask of the AGRIF grid
3. The `AgrSpg` variable is interpolated by OASIS, taking into account the `agr2` mask. Doing so, any OpenIFS grid point located outside the zero value boundary of the AGRIF zoom grid will be filled with a zero value, though grid points located inside the AGRIF zoom limit will be filled with the `AgrSpg` zero-to-one values (see Fig 6, right)

4. The interpolated `AgrSpg` variable is received by OpenIFS, which uses it as a weight function to merge surface variables coming from the parent and the child grid :

```
Combined_field = AgrSpg.AGRIF_field + (1-AgrSpg).Global_field
```

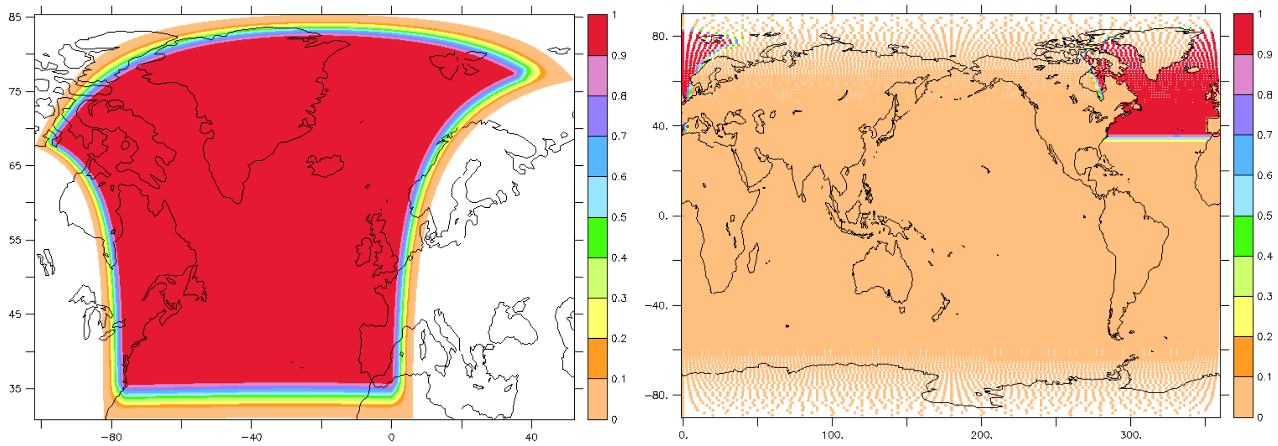

Fig 6: 2D function defining the weight of coupled fields sent by AGRIF zoom (relatively to the weight of coupled fields coming from the global NEMO grid), as defined in the AGRIF NEMO grid (left) and interpolated on the OpenIFS grid (right)

Several modifications were also necessary in both NEMO and OpenIFS interfaces to ensure all new coupling fields exchanges.

In NEMO,

- to avoid calling the final `oasis_enddef` coupling field definition procedure but for the last AGRIF zoom coupling field definition
- to remove the coupling frequency checking, since this frequency is now different for parent and child coupling fields

In addition, we found necessary to define the `namcouple` LAG parameter in link with the outgoing NEMO parent and child coupling fields with the same value equal, to the parent time step. A different value leads to the OASIS "model seems to be running backwards" message.

As previously explained, a new variable (`AgrSpg`) has to be defined and sent within the `sbccpl` routine. The length of the transition zone is arbitrarily set to 48 + the length of the AGRIF sponge zone, which is

```
2 + agr_coeff.2 + 2
```

with `agr_coeff` equal to the zonal/meridional AGRIF zoom factor between parent and child grids.

Code modifications will be reported to the `SVN` NEMO repository, thanks to the NEMO system team, to facilitate the definition of future coupled configuration including AGRIF zooms.

In OpenIFS,

- to allow the reception of the NEMO parent/child weight function coupling field, a new coupled field must be declared to OASIS (`cplng_data_mod.F90`)
- to use this field by multiplying and combining all parent/child incoming fields (`foci_updclie.F90` and `foci_get_ice_state.F90`)

A new OpenIFS namelist parameter is created (`NAMECECFG/FOCI_CPL_NB_OCE_ZOOM`). When it is defined and strictly positive, this parameter enable the AGRIF coupling fields reading. Notice that this parameter can be larger than 1, because fields coming from several AGRIF zooms can be received. However, the geographical zones of every zoom must not intersect the others (coupling of recursive zooms is not implemented).

Modifications were committed to the ESM-tools community repository. Further effort would be needed to include them in the official OpenIFS version.

Validation

The validity of our implementation is checked with a low resolution version of the FOCI model. In addition, a simplified procedure is set up to be able to quickly compile and launch the tests version, without using the ESM-Tools command. However, the modified code is kept in the original directory, to make possible the `git` repository update at the end of the support period.

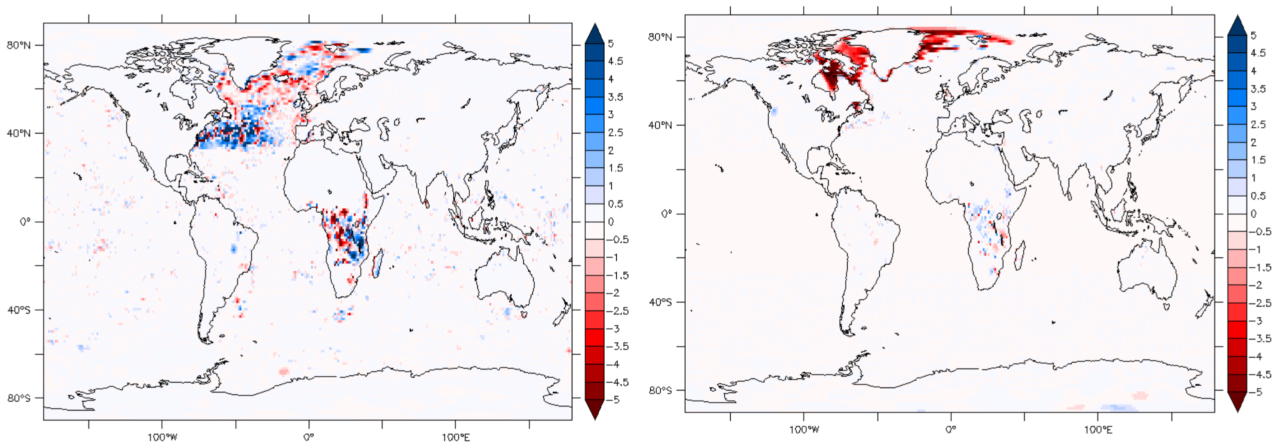

Figure 7: Impact of AGRIF zoom two-way coupling (zoom surface variables used by OpenIFS instead of global ocean surface variables) on latent heat flux in W/m2 (left) and skin temperature in K (right), averaged during the first 12 hours

Exchanges effect

The new communication to OpenIFS of AGRIF surface fields in the zoom area (North Atlantic) is checked by comparison of the results of two short simulations, with (two-way) and without (one-way) OASIS exchanges of these variables. By chance, the parent and child models prognostic variables are initialised with different values, which makes easier the identification of the child coupled fields effect in OpenIFS after a short period (12h). After such a short timing, variables like skin temperatures and latent fluxes cannot be significantly perturbed by model variability. Fig 7 shows the differences two-way – one-way, that clearly sign the effect of different surface values coming from the AGRIF zoom.

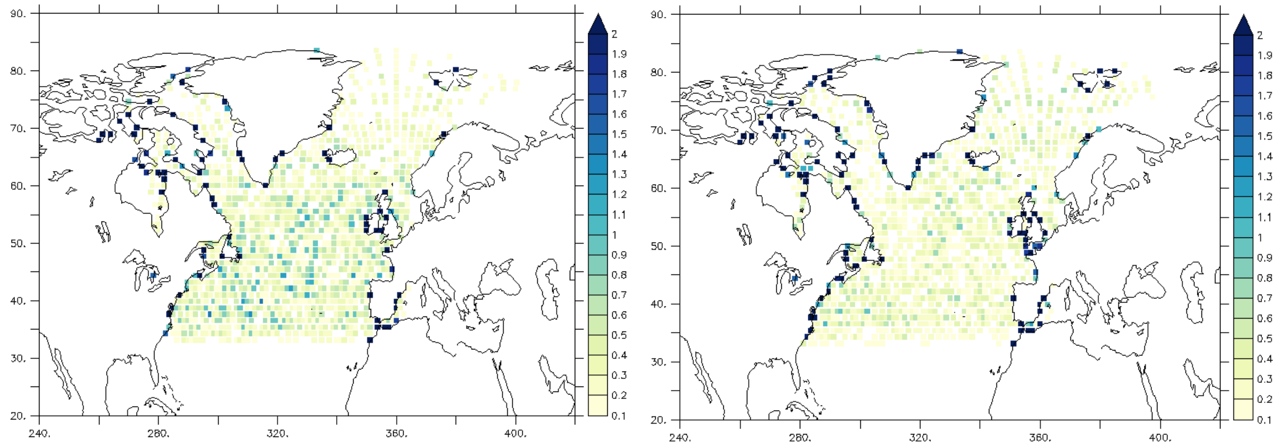

Figure 8: Interpolation error (%) measured on an analytical field, from AGRIF zoom grid (10km) to OpenIFS reduced Gaussian grid (125km), GAUSWGT SCRIP interpolation, variance (VAR) set to 2.0, and number of neighbours equal to 40 (left) and 90 (right)

Interpolation parameter tuning

The GAUSWGT SCRIP/OASIS interpolation is used for all FOCI exchanges. The SCRIP routine finds, in the source grid, a variable number of nearest neighbours for each unmasked target grid point and calculates the Gaussian distributed weights for all its neighbours. Variance of the Gaussian distribution function (GV) and number of neighbours (NN) are two parameters that can be set in the `namcouple` file. We propose to investigate the importance of these two parameters, particularly in the case of the AGRIF to OpenIFS interpolation. In that perspective, we use the OASIS toy models `test_interpolation` to generate (in parallel) and estimate the interpolation error of an analytic field [23].

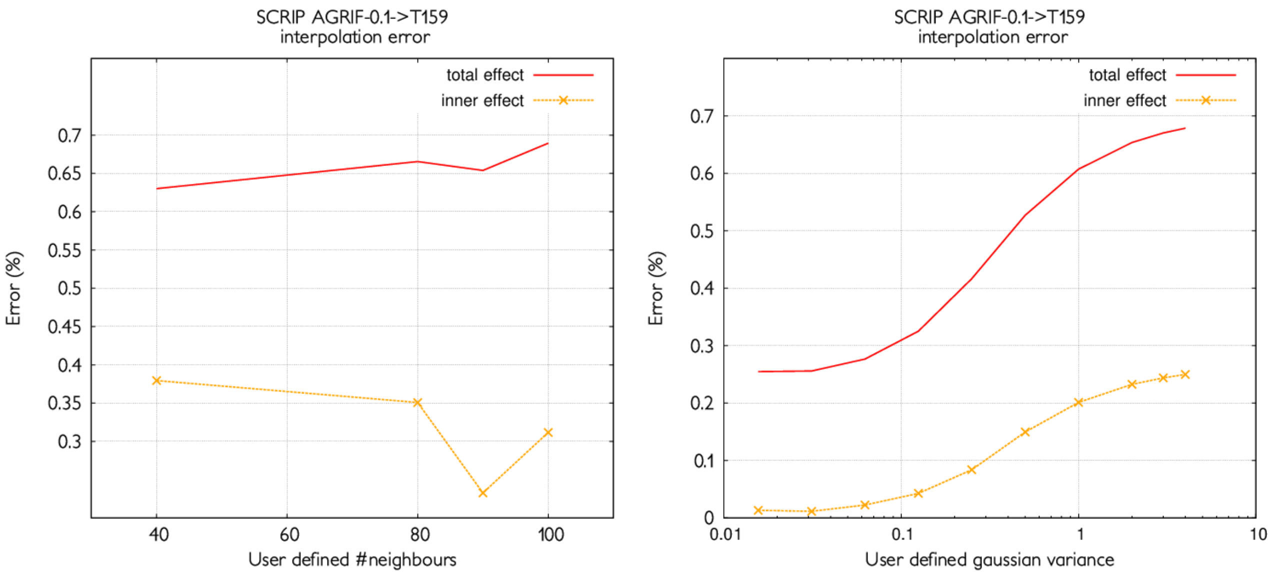

Figure 9: Average interpolation error in the whole AGRIF domain (red) and in a 55E-20E-35N-45N box (orange) as a function of the number of neighbours with constant Gaussian variance = 2.0 (left) and as a function of Gaussian variance with constant number of neighbours equal to 90 (right)

The percentage of this error is represented in Fig 8, for two different set of parameters. Error can be split in two different areas, near the coast line (most important values) and in the open seas (negligible).

The average of this error is plotted in Fig 9, in the whole grid (red) and in a box located in an open sea region (orange), when NN (right) or GV (left) is changed. Stronger beneficial effect on accuracy is found when GV is changed, compared to NN. The same minimum is found for both coastal or open sea grid points (GV = ~ 1/32). However, this result must be mitigated taking into account the small spatial variability of the analytic field chosen (COS(lon)*COS(lat)).

Then we tried to evaluate the interpolation parametrisation effect to a real field like the model SST. Since it is difficult to evaluate local conservation (and the corresponding interpolation error) of spatial values, we propose to compare SST time variability of two simulations produced with a different set of GV/NN parameters. This time variability seems to be better preserved with small values of the GV parameter (see Figure 10) but most of the information is lost due to the strong difference in resolution between AGRIF zoom and atmosphere.

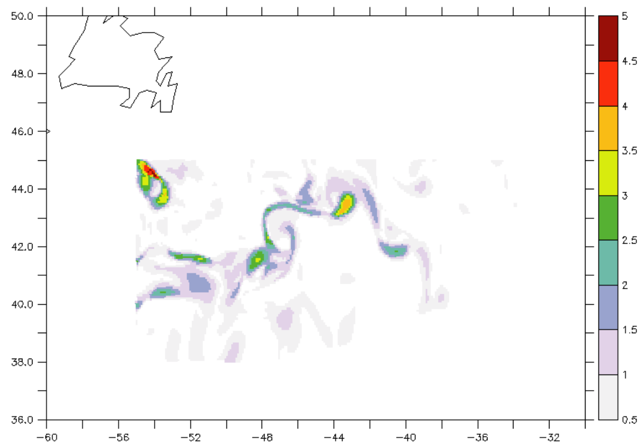

Figure 10: SST time variability (standard deviation, K), at coupling frequency (3h) during 15 days, on source grid AGRIF (top) and target grid OpenIFS, GAUSWGT SCRIP interpolation with 90 neighbours and Gaussian variance = 2 (bottom left) and with 100 neighbours and optimised Gaussian variance = 1/32 (bottom left)

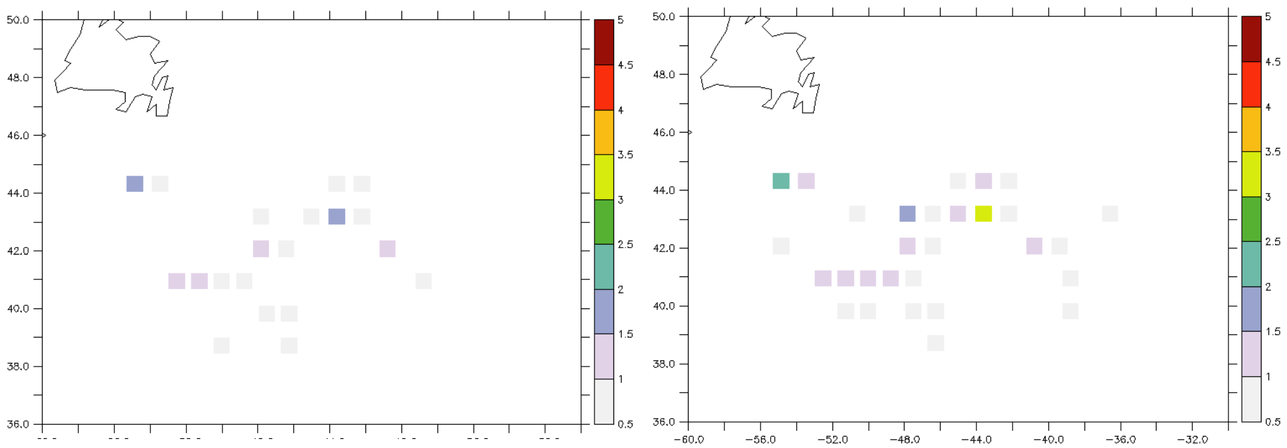

Increasing resolution

The final target of this implementation is to take the most of a compute efficient coupling in order to increase the spatial resolution of the atmosphere model. A T799 (25km) configuration of OpenIFS is prepared (input files), and coupled to the same NEMO ORCA05 + AGRIF zoom 10km. The atmosphere resolution in our new configuration is finer than global ocean one, but two time coarser than the zoom. Coupling time step also decreased from 3 to 1h. This resolution increase can be done with a minimum modification of OASIS input files. Parameters in `namcouple` must be adjusted and new auxiliary/restart file produced. Interpolation weight files are automatically generated. However, a new bug is identified in the recent update of the GAUSWGT related routine and its solution proposed (ticket #2500).

As in low resolution case, a preliminary study is led to find the ideal NN/GV parameter set. But the error is now practically negligible (Fig 11).

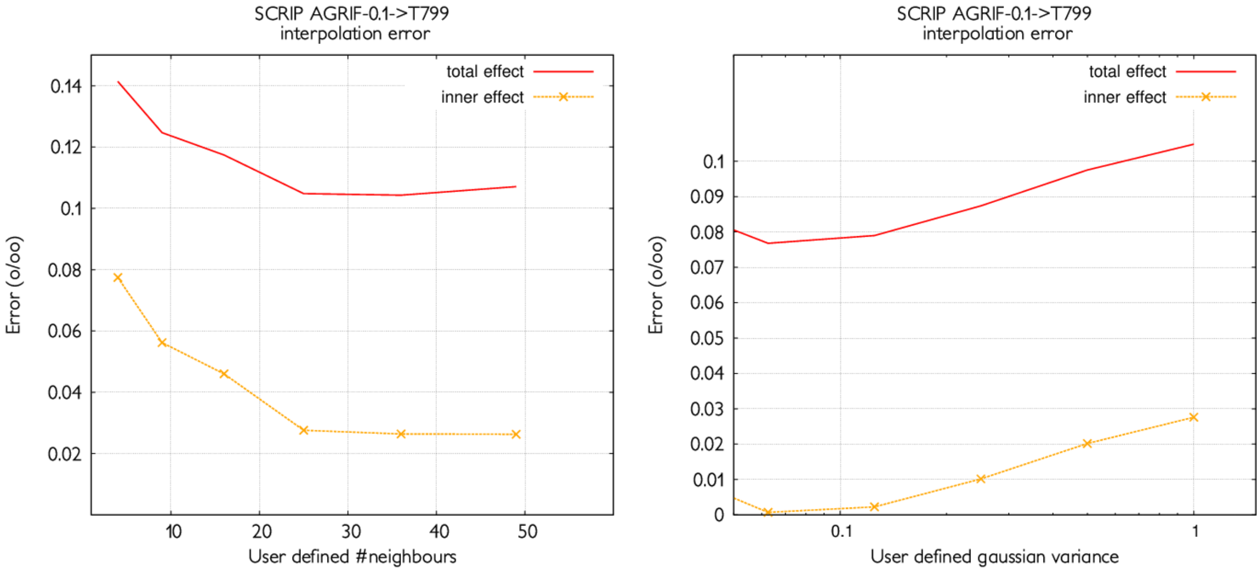

Figure 11: Average interpolation error in the whole AGRIF domain (red) and in a 55E-20E-35N-45N box (orange) as a function of the number of neighbours with Gaussian constant variance = 1. (left) and as a function of Gaussian variance with constant number of neighbours equal to 25 (right)

Moreover, a comparison of the SST time variability, of two one month long simulations with different NN and VAR, shows small differences, bigger with NN=25/VAR=1/32, but not significant (maximum standard deviation lower than .001 K).

These two one month long simulations, produced after a two month long spinup, give a sufficient guarantee of model stability and validates the technical implementation of the FOCI-AGRIF coupled model. After a computing resource tuning of the 2 models (load balancing), a speed of 1 SYPD is measured, for a cost of 17,000 CHPSY. The AGRIF zoom multiplies by a factor 10 the cost of the global ORCA05 model but this cost has to be compared with the cost of the ORCA10 (global 1/10 degree) that would be necessary to set AGRIF spatial resolution of the North Atlantic ocean globally. Following [23], this cost could be roughly estimated to 300 times the ORCA05 one. In conclusion, the set up of this configuration will help to lead selected studies relying on high resolution modelling that costs approximately 30 times less than the same studies relying on the standard ORCA12 configuration.

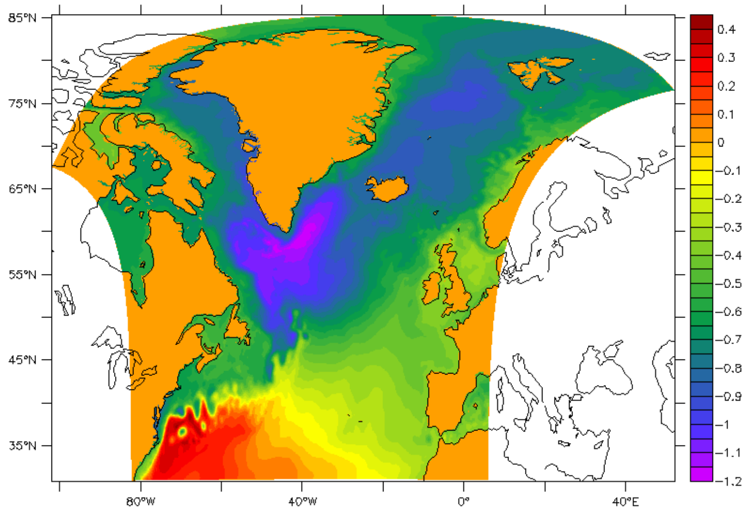

Fig 12: Sea Surface Height five day average at the end of a FOCI-AGRIF 3 month long simulation

Our new configuration already put in evidence the capacity of 2 way coupling to increase both average and variance of surface exchanges, e.g. latent flux, in eddy-rich regions. In addition to a porting on HLRN supercomputers, more studies are planned in the months to come, like upgrading similarly the FOCI-ECHAM configuration, or changing the zoom location (Southern ocean).

Community impact

During the selection procedure, the panel emphasised the importance of not restraining the support to a one to one collaboration but rather prefer actions that could have a broader impact on communities. We tried to quantify this community impact, in a table that summarises (i) the oral communications organised and the origin of the participant/audience, (ii) code updates in official centralised repositories, from which OASIS `gitlab` and (iii) written communications (emails) to laboratories making part of the hosting laboratory working network. This counting necessarily neglects any action in link with our work, organised by the hosting laboratory, that could take place after the dedicated support period.

	ETH	MetOffice	GEOMAR
Talks/meetings	* Starting meeting (3 people)	* Starting meeting (8 people, internal, across teams) * Closing presentation (10 people internal, across teams)	* Monday's informal modelling meeting (3 to 4 people, internal, across teams) * Closing presentation (15 people, internal, across teams) * ESM-Tools Workshop meeting (by GEOMAR staff, 10 people, external)
Repository updates	none	NEMO, MetOffice branch OASIS tickets	NEMO, v4 and trunk FOCI git repository (DKRZ) OASIS3-MCT, v4 and master
Networks	MeteoSwiss (1 email)	Exchanges with ESM, HPC and biogeochemistry community of MetOffice (~10 emails)	Exchanges with DKRZ (4 emails)

Table 2: Quantification of community level communications during support

Of course, community impact cannot be fully evaluated on such short period and using such dangerously formal criteria, but it gives an idea on how practicality results of the support can be used outside the hosting laboratory.

Acknowledgements
This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084

Bibliography

[1] Bianco, M., Diamanti, T., Fuhrer, O., Gysi, T., Lapillonne, X., Osuna, C. and Schulthess, T., 2013: A GPU Capable Version of the COSMO Weather Prediction Model, International Supercomputing Conference, Leipzig, Germany
[2] Lawrence, D.M. R.A. Fisher, C.D. Koven, K.W. Oleson, S.C. Swenson, G. Bonan, N. Collier, B. Ghimire, L. van Kampenhout, D. Kennedy, E. Kluzek, P.J. Lawrence, F. Li, H. Li, D. Lombardozzi, W.J. Riley, W.J. Sacks, M. Shi, M. Vertenstein, W.R. Wieder,, C. Xu, A.A. Ali, A.M. Badger, G. Bisht, M.A. Brunke, S.P. Burns,, J. Buzan, M. Clark, A. Craig, K. Dahlin, B. Drewniak, J.B. Fisher, M. Flanner, A.M. Fox, P. Gentine, F.Hoffman, G. Keppel-Aleks, R., Knox, S. Kumar, J. Lenaerts, L.R. Leung, W.H. Lipscomb, Y. Lu, A., Pandey, J.D. Pelletier, J. Perket,, J.T. Randerson, D.M. Ricciuto, B.M., Sanderson, A. Slater, Z.M. Subin, J. Tang, R.Q. Thomas, M. Val Martin, and X. Zeng, 2018: The Community Land Model version 5: Description of new features, benchmarking, and impact of forcing uncertainty, Submitted to J. Adv. Model. Earth Syst.,
[3] Maisonnave, E., 2011: OASIS4 coupling interface implementation on ETHZ land-atmosphere coupled model, Techn. documentation, WN/CMGC/11/15, pp13, SUC au CERFACS, URA CERFACS/CNRS No1875, France
[4] T. Gysi, C. Osuna, O. Fuhrer, M. Bianco and T. C. Schulthess, 2015: STELLA: A domain-specific tool for structure gridmethods in weather and climate models, , New York, NY, USA
[5] Maisonnave, E., 2011: A simple OASIS interface for CESM, Techn. documentation, TR/CMGC/11/63, pp12, SUC au CERFACS, URA CERFACS/CNRS No1875, France
[6] Davin, E. L., Maisonnave E. and Seneviratne, S. I., 2016: Is land surface processes representation a possible weak link in current Regional Climate Models ?, Environ. Res. Lett., 11
[7] Valcke, S., Craig, A. and Coquart, L., 2018: OASIS3-MCT User Guide, OASIS3-MCT4.0, Techn. documentation, TR-CMGC-18-77, pp64, CECI, Université de Toulouse, CNRS, CERFACS, France
[8] Maisonnave, E. and Valcke, S., 2012: OASIS Dedicated User Support 2011, Annual report, Techn. documentation, TR/CMGC/12/18, pp28, SUC au CERFACS, URA CERFACS/CNRS No1875, France
[9] Maisonnave, E., 2019: persona version 2.1, how to graphically patch a mask variable, Working note, WN/CMGC/19/107, CECI, UMR CERFACS/CNRS No5318, France
[10] Maisonnave, E. and Caubel, A., 2014: LUCIA, load balancing tool for OASIS coupled systems, Techn. documentation, TR/CMGC/14/63, pp16, SUC au CERFACS, URA CERFACS/CNRS No1875, France
[11] Maisonnave, E., 2013: A regional coupling with OASIS3-MCT, Techn. documentation, WN/CMGC/13/34, pp6, SUC au CERFACS, URA CERFACS/CNRS No1875, France
[12] Madec, G., Bourdallé-Badie, R, Bouttier, P.-A., Bricaud, C., Bruciaferri, D., Calvert, D., … Vancoppenolle. M., 2017: NEMO ocean engine (Version v3.6.1), Notes du pôle de modélisation de l'Institut Pierre-Simon Laplace (IPSL), http://doi.org/10.5281/zenodo.3248739
[13] Maisonnave, E. and Masson, S., 2015: Ocean/sea-ice macro task parallelism in NEMO Technical Report, TR/CMGC/15/54, SUC au CERFACS, URA CERFACS/CNRS No1875, France
[14] Maisonnave, E. and Masson, S., 2019: NEMO 4.0 performance: how to identify and reduce unnecessary communications , Technical Report, TR/CMGC/19/19, CECI, UMR CERFACS/CNRS No5318, France
[15] Berthet, S., Séférian, R., Bricaud, C., Chevallier, M., Voldoire, A., & Ethe, C., 2019: Evaluation of an Online Grid-Coarsening Algorithm in a Global Eddy-Admitting Ocean Biogeochemical

Model, Journal of Advances in Modeling Earth Systems; Washington Vol. 11, Iss. 6, 1759-1783. DOI:10.1029/2019MS001644
[16] Haiden, T., Janousek, M., Bauer, P., Bidlot, J., Ferranti, L., Hewson, T., Prates, F., Richardson, D. S., and Vitart, F., 2014: Evaluation of ECMWF forecasts, including 2013–2014 upgrades, Technical Memorandum no. 742, European Centre for Medium Range Weather Forecasts
[17] Park, W., N. Keenlyside, M. Latif, A. Stroh, R. Redler, E. Roeckner, and G. Madec, 2009: Tropical Pacific Climate and Its Response to Global Warming in the Kiel Climate Model, Journal of Climate, 22(1), 71-92, doi:10.1175/2008jcli2261.1
[18] Matthes, K. et al., The Flexible Ocean and Climate Infrastructure (FOCI): Mean State and Variability, GMD Discussion submitted.
[19] Open Integrated Forecasting System, https://confluence.ecmwf.int//display/OIFS/Single+column+model+40r1+release+notes
[20] https://www.esm-tools.net/
[21] Sidorenko, D., Rackow, T., Jung, T., Semmler, T., Barbi ,D., Danilov, S., Dethloff ,K., Dorn, W., Fieg, K., Goessling, H. F., Handorf, D., Harig, S., Hiller, W., Juricke, S., Losch, M., Schröter, J., Sein, D. V., Wang, Q., 2015: Towards multi-resolution global climate modeling with ECHAM6–FESOM. Part I: model formulation and mean climate., Climate Dynamics, 44(3-4), pp.757-780.
[22] Rackow, T., Goessling, H.F., Jung, T., Sidorenko, D., Semmler, T., Barbi, D. and Handorf, D., 2016:Towards multi-resolution global climate modeling with ECHAM6-FESOM. Part II: climate variability. Climate Dynamics, pp.1-26.
[23] Maisonnave, E., 2019: OASIS Gaussian interpolation with locally relative source-target distances, Technical Report, TR/CMGC/19/83, CECI, UMR CERFACS/CNRS No5318, France
[24] Maisonnave, E., Lévy, C. and Masson, S., 2019: Implementation and computing performance of NEMO ORCA-Km configuration, Technical Report, TR/CMGC/19/18, CECI, UMR CERFACS/CNRS No5318, France

Appendix

Practical issues

We propose here to summarise management problems that were solved or postponed during our support. This highlights how unexpected and then how difficult it is, to provide, at European scale, a cross laboratory support by organising the on-site collaboration with a specialist. Hopefully, hosting laboratories largely contributed to facilitate the organisation, which ensures good working conditions everywhere. The two most expensive and constrainted budget items are discussed : travel and accommodation.

Although now mainly accepted and even recommended in our community[5] [24], train travel still requires important effort to replace the standard flight solution. Its price necessarily increases the total budget or reduces other expenses such as accommodation. If time spent on train can be converted in working periods, depending on level of comfort and coach equipment, long distance journeys necessarily require to renounce week-end resting stages, without any compensation. In addition, side cost such as meals or accommodation during stops can be hardly refunded. International train journeys can be difficult to book, if not impossible when using institutional booking software/procedures. The alternative free booking, if allowed, is time consuming. Due to the lack/destruction of railways infrastructure in our continent, lost connections are frequent and can lead to costly (and partially covered) replacement solutions and delays. Long term policy of night train reduction worsens the problem and has made practically impossible any trip to Scandinavia/North Eastern Europe. The European Train Timetable was a valuable help to find our path. Commuting possibilities in the 3 cities of Zurich, Exeter and Kiel are mostly reduced to the bus network, with reasonably good frequencies and price, although the accessibility to monthly ticket was diverse. Due to this same drawback, bicycle renting was not an option.

Increasing price of accommodation in big but also small cities makes difficult not to overcome the initial funding possibilities and requires very early reservations. Better than using on-line hosting platforms, that are part of the issue, we preferred to seek institution help (Kiel University guest-house) or book University student rooms that are available during summertime (University of Exeter).

Costs/Sustainability

Budget, energy consumption and carbon footprint are provided in the following table. Computations and train transport are the only two items considered in this summary. Everyday consumption, from which electricity supply for workstation (Intel Atom N270 or Arm Cortex A53) and supercomputer login nodes, is neglected. Energy/CO2 emission conversion (Carbon Intensity) for transport and supercomputing is country and machine dependant.

5 Janisch, T., and Hilty, L., 2017: Changing university culture towards reduced air travel – Background Report for the 2017 Virtual Conference on University Air Miles Reduction. Zurich, Switzerland: ETH Sustainability. doi:10.5167/uzh-14212

A comparison of transportation and computation Carbon footprints clearly put in evidence that a comprehensive effort made to avoid air-plane journeys is not sufficient to lower the total amount of greenhouse gases emissions by more than 50%. The reason is that even short set up tests led with high resolution models, like NEMO-SI3 at ORCA12 resolution, were sufficient to strongly increase the ecological impact of our work.

	Cost	Travel		Computing			Total Carbon footprint
	(€)	(Km)	(KgCO2e)[6][7]	(Core.h)	(kWh)	(KgCO2e)[8]	(KgCO2e)
ETHZ	3670	2,700	14	1,200	7[9]	0	14
MetOffice	2790	1,200[10]	20	40,000	750[11]	445	465
GEOMAR	1140	1,950[12]	30	8,500	114[13]	86	144
Total	7600	5,850	64	49,700	871	531	623

This quantitative cost analysis makes no sense without the corresponding qualitative analysis of the work produced. In particular, one would paid attention to the capacity of some coupled models, set up during this program, to reduce the computing cost compared to other configurations, while giving the same quality of results, from a geophysical point of view.

6 SNCF carbon intensity high speed train : 2,4 gCO2equ/Km, intercity : 8.1 gCO2equ/Km, from https://www.oui.sncf/aide/calcul-des-emissions-de-co2-sur-votre-trajet-en-train and https://ressources.data.sncf.com/explore/dataset/emission-co2-tgv/

7 Paris/London Eurostart journey carbon intensity 4.1 KgCO2equ https://www.eurostar-treadlightly.com/en/environment.php and 41,15 gCO2equ/Km for National Railways https://www.gov.uk/government/publications/greenhouse-gas-reporting-conversion-factors-2019

8 Carbon intensity of High voltage in Switzerland (~29), UK (593) and Germany (599), according to Moro A., Lonza L., 2018: Electricity carbon intensity in European Member States: Impacts on GHG emissions of electric vehicles, Transportation Research Part D: Transport and Environment, 64 , pp. 5-14.

9 CSCS CRAY XC50 consumption for 388,000 cores: 2,384kW, PUE < 1.25, see TOP500

10 Return ticket is not taken into account since it was used for another purpose than this dedicated support (IS-ENES Sea-Ice Workshop)

11 MetOffice CRAY XC40 consumption for 90,000 cores: 1,348kW, PUE = 1.25, see TOP500

12 Single ticket is not taken into account since it was used for another purpose than this dedicated support (IS-ENES Sea-Ice Workshop)

13 DKRZ mistral consumption for 99,000 cores : 1,116kW, PUE = 1.19. See TOP500 and https://www.dkrz.de/communication/news-archive/en-energie-effizienz-des-rechnersystems-mistral