Coupling NEMO global ocean with hemispheric Arctic and Antarctic ice models

E. Maisonnave *, R. Bourdallé-Badie °

* CECI, UMR CERFACS/CNRS No5318, France
° Mercator-Ocean International, France

TR/CMGC/22/18

Abstract

Two new NEMO-SAS (stand alone surface module) configurations are derived from a global ORCA grid. A FORTRAN tool is developed to transform the ORCA input files and define the OASIS weights & addresses to map fluxes and surface variables between one global and two hemispheric grids. A tri-component coupled system is set up, including an ocean on the ORCA2 global grid and two hemispheric NEMO-SAS with SI3 sea-ice. Removing North Pole folding communications in ice model helps to improve computing performance, since our ORCA2 coupled configuration is twice faster and two times cheaper than its single executable counterpart. The North Pole folding free hemispheric grid discretisation simplifies the coding of the future neXtSIM model, while relying on the same grid points than the ocean, thus avoiding inaccurate interpolation between ice and ocean grids

Table of Contents

1. Rationale

This work belongs to the set of ongoing developments that aim to change the spatial discretisation of the neXtSIM [1] sea-ice model, from Lagrangian to Eulerian grids. The discretisation choice is guided by several requirements from which :

- the coupling with an ocean model, NEMO [2] as a first step, but not exclusively
- the importance of local conservation at ocean/sea-ice interface
- the coupling with an atmosphere model such as CNRM-CM [3]

The tight coupling between sea-ice and ocean (several fluxes and surface fields exchanged every few time steps) and the modularity required (neXtSIM must be plugged to several ocean models during the SASIP [4] project), in addition to the inclusion of the ocean/sea-ice unit in larger climate systems, argue for the implementation of a coupler interface, for example based on OASIS [5].

Moreover, one of the main objective of the neXtSIM modeling is the enhancement of the representation of near shore phenomena such as leads but also polynyas. This requires to conserve fluxes and surface variables with particular care in coastal areas. Even though the SCRIP interpolation package included in OASIS can perform accurate and locally conservative interpolations, any land-sea mask mismatch between ocean and sea-ice model would necessarily lead to significant errors in these near shore areas. This is why, in conjunction with the fact that interpolation possibly downgrade the overall computing performance, we prefer to keep the same grid in the two models. Even though it means that the neXtSIM spatial discretisation must be done on the NEMO grid (ORCA, for global configurations), the requirement of ocean agnostic development forbids the simple use of NEMO sub-routines in the neXtSIM model development. We think in particular, to the `lbc_lnk` MPI communication routine, needed to provide boundary conditions to subdomain halos. This routine provides the full (and complex) algorithm to fill the halos, in particular, for global grid, in the so-called North Pole Folding (NPF) area. Since the ORCA grid discretisation is only used in the NEMO model, the inclusion of the `lbc_lnk` routine would clearly forbid the new Eulerian-neXtSIM to be ocean model agnostic.

To address this issue, we propose to discretise the neXtSIM model on two grids. In that purpose, the ORCA grid is split at the equator into two hemispheric grids. A new ordering of the ORCA Northern Hemisphere (ONH) grid points allows the NPF removal and to treat the domain as a closed regional grid, since no sea-ice is supposed to be found at their equatorial boundaries. Regarding the ORCA Southern Hemisphere (OSH), a simple cyclic boundary condition update needs to be implemented to perform the subdomain boundary condition exchanges at each eastern and western boundaries of the grid.

This solution has the advantage to ensure the similarity of the grid points in NEMO ocean and sea-ice models. In addition, to ensure the perfect conservation of the coupled quantities exchanged, it increases the coupling computing performance by avoiding to perform any interpolation computation. At the same time, this solution simplifies the neXtSIM MPI parallelisation coding, in an ocean model agnostic way.

As always with OASIS, the coupled configuration is more difficult to handle than a single executable: three executables must be launched and load balanced, and three sets of input files must be provided. In addition, a clean and exhaustive interface must be defined to be able to exchange the coupled quantities in the same way than with a single executable. Hopefully, this work is already partially implemented, tested and maintained in the NEMO model [6,7], and also validated in a climate configuration. The load balancing is highly simplified with the 5.0 version of the coupler [8].

Before delivering the ONH/OSH grids description to the neXtSIM developers, we propose in this document to set up a coupled model with the desired discretisations, but including the original NEMO SI3 [9] sea-ice model (§2). A tool will be necessary to derived the two ONH/OSH NEMO-SAS input files from the original ORCA grid, in parallel of the OASIS weight and address (W&A) file needed for coupling (§3). The computing performance will be checked (§4), before the final recommendations for an implementation in neXtSIM (§5).

2. Test configuration

The 4.0.6 version of the ocean model, available at development starting date, is chosen to implement and test the three executable system. A branch could be derived from this point if definitive modifications are required in the official NEMO code. Considering their small amount, modifications of a more recent version (e.g. 4.2, when available) would be straightforward.

To simplify the test-and-correct cycles during development and performance measurement, a global ORCA2 resolution is chosen. Another advantage of this light configuration is that its input files can be easily downloaded and perfectly fit to the reference set of parameter files provided via the official portal. And finally, the performance limit can be reached with a small number of computing resources, which speeds up the scaling tests.

The full testing configuration includes:

- a global ORCA2 configuration, ocean only,
- a NEMO-SAS configuration, discretised on an ONH grid derived from ORCA2,
- a similar configuration, but discretised on the complementary OSH grid,
- and a set of XIOS servers for output

The ocean part is rather similar to the one developed in [6] and upgraded in [7]. As explained in these documents, the same released configuration must simply be compiled without the sea ice CPP key but with the OASIS one. At runtime, the `namelist` parameter file must be modified to put the ocean model in SAS coupled mode[1]. None code modification was necessary, except the prefixing of new output files[2], to avoid that the three executable overwrite their output on the same file. Small adjustments were also made to start the model from rest (it was mandatory to start from a restart in the released NEMO-SAS coupled implementation, to ensure reproducibility between coupled and single executable versions). Compared to the released NEMO-SAS coupled model, the ocean must now send its surface variables to (and receive its fluxes from) *two* ice models. Hopefully, the NEMO OASIS interface (`sbccpl`) was designed to perform exchanges with more than one model [10]. A mask variable, which defines the ORCA grid area in contact with each model, must be provided as an input file (see next §) and the number of coupled models can be set in a `namelist` parameter[3].

From the SAS side, the `namelist` parameter files must be modified symmetrically to put the models in SAS coupled mode[4]. The only code modification compared to the release was the addition of a prefix in the SAS coupled fields names[5].

As any OASIS coupled model, a parameter file called `namcouple` has to be provided, describing the exchanged fields (accordingly to their declaration in the three model interfaces). We keep the same 13 coupling fields defined in the previously released NEMO-SAS configuration but we duplicate them in order to serve the two ice models. Unlike in this previous version, and despite the identical position of their grid points, an interpolation[6] have to be defined between the global ORCA and the hemispheric OHN (both directions) grid in one hand, and between the global ORCA and the hemispheric OHN (both directions) grid in the other hand. Predefined W&A files are used in that purpose (see next §). This

The time step is kept unchanged compared to the original ORCA2 configuration. It is the same in NEMO ocean and SAS sea-ice components. The calling frequency of the surface module is kept to the same value in SAS[7], but set to 1 in the ocean. At the same time, the coupling frequency is defined as the surface module calling frequency, that ensures the same update frequency of the incoming fluxes in the ocean model than in single executable mode.

---

[1] `nn_components = 1`

[2] `time.step, run.stat` & `run.stat.nc` in `stpctl.F90`

[3] `nn_cplmodel = 2` and `ln_usecplmask = .true.`

[4] `nn_components = 2`

[5] respectively `south_` and `north_` for OSH and ONH models, in `cpl_oasis.F90`

[6] `MAPPING` keyword transformation

[7] `nn_fsbc = 4`

Finally, the XIOS servers were configured to output variables from the three models. Of course three different contexts had to be defined properly in their respective `xml` configuration files.

3. Transformation tool

The NEMO code modifications are limited and the definition of the OASIS parametrisation rather straightforward. Most of the complexity of the coupled configuration implementation lies in the input file definition. A tool is developed for that purpose. The FORTRAN language is chosen to ensure a quick processing of large grids. Although out of fashion, this widespread language also ensures a quick understanding by any community user, a necessary condition for a true freedom of use.

Two kinds of new input files are necessary to define and use the new grids :
- the NEMO configuration files, required to define the hemispheric grids, on one hand,
- the OASIS W&A files required to perform the remapping between global and hemispheric grids, on the other hand.

Both operations are based approximately on the same transformation, as described in Fig. 1. The ORCA grid has `jpiglo x jpjglo` points, including one or one and a half lines of duplicated points at NPF.

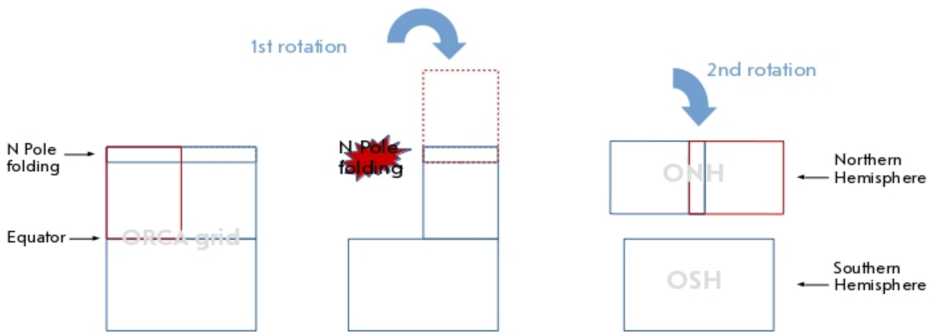

Figure 1: Definition of the ORCA grid transform to OSH and ONH grids. On the left , the ORCA global grid is represented (Northern and Southern Hemisphere and NPF area). On the right, the two hemispheric grids include the same ORCA grid points, but ordered differently. At center, the intermediate step necessary to transform the upper ORCA grid into the ONH grid

The lowest part of the ORCA grid (`j` index from `1` to `jpjglo/2`) is extracted from the global ORCA grid and copied without any change in the new OSH grid. Notice that the 2 duplicated columns at eastern and western boundaries are also defined in the OSH grid to keep an East-West periodicity[8]. The ORCA-OSH transformation is straightforward, as it associates each point of the lowest part of the ORCA grid to one similar grid point of OSH, in the same order.

The transformation that implies the upper part of the ORCA grid is slightly more complex, while already used for Arctic only configuration in [11]. The `jpiglo/2` first columns are rotated and put at the top of the last `jpiglo/2` to `jpiglo` columns. The overall columns are rotated again to make the new ONH `jpiglo` number bigger that `jpjglo`. The useless duplicated grid points of the ORCA grid are not transformed, leaving the new ONH grid without NPF and, since no boundary cyclic conditions are required there, the first and last column of the ORCA grid are no used either [9].

Even though the new neXtSIM model is supposed to be finally independent from the NEMO model, we chose to develop the transformation tool in the NEMO directory, to simplify maintenance & sharing, and centralise the developments. However, no specific NEMO subroutine was used and the program[10] can be compiled on another environment.

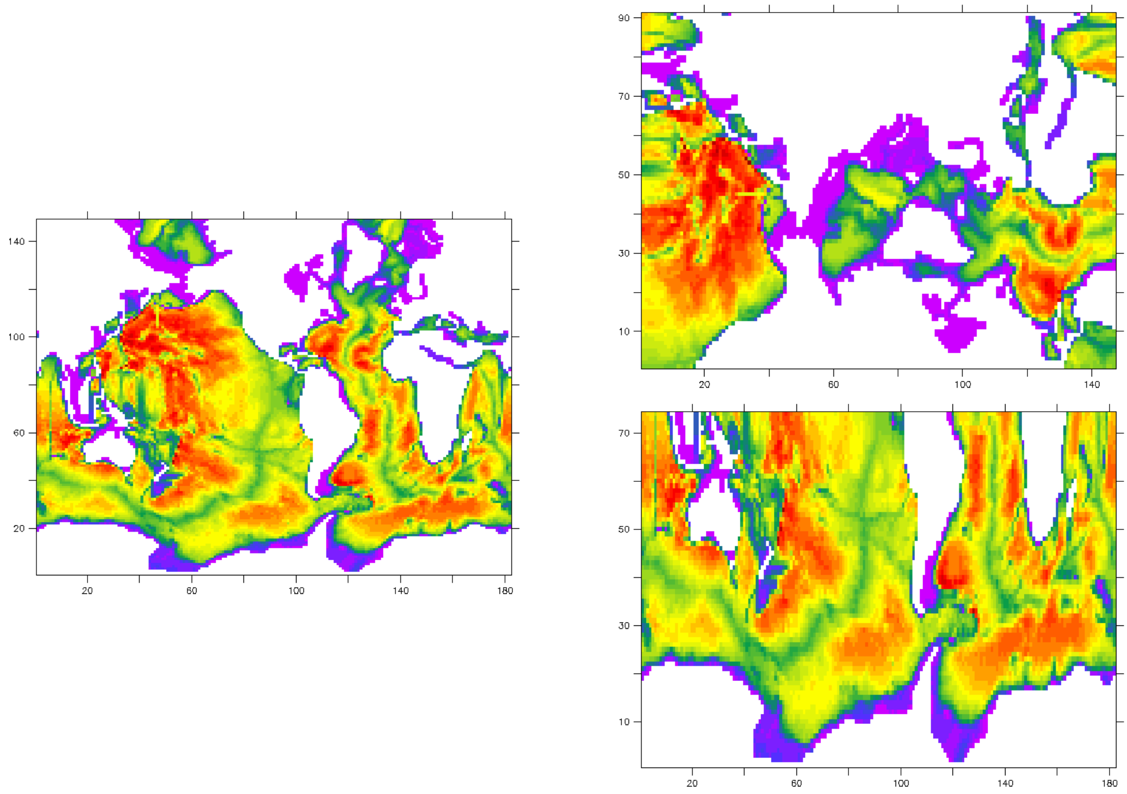

Figure 2: ORCA2 grid bathymetry (left) and the same after splitting (+ rotating) between Northern and Southern hemispheric grids (right)

[8] The `jperio` variable in the `DOMAINcfg.nc` file of OSH will have to be set to 1. In the next NEMO version 4.2, this requirement will disappear with the duplicated column.

[9] The `jperio` variable in the `DOMAINcfg.nc` file of ONH will have to be set to 0

[10] named `tools/SPLIT_ORCA/src/split_orca.F90`

The definition of the initial ORCA grid relies on the `meshmask.nc` file coming from a previous single executable simulation. We prefer to deduce our transformation from this output file rather than the `DOMAINcfg` input file, which describe the same grid geometry but include possible configuration driven modification (like straights opening or closed sea definition). The `meshmask` file provides the dimensions and the periodicity, in addition to the land-sea masks of the T,U,V and F grids. Notice that both `jperio=4` and `jperio=6` geometries are supported.

A `namelist_cfg` file, specific for the tool parametrisation, is required to define the mesh file name, in addition to the name of a second input file. This name is needed optionally, if the user wants to transform a NEMO input file defined on the source ORCA grid. In any cases, the tool creates the OASIS W&A files and the `cplmask` file required to gather, in the ORCA grid of the ocean model, the coupling fields coming from the two hemispheric grids.

The core algorithm that associates one point of the T,U,V,F ORCA grids to the corresponding point of the T,U,V,F OSH grids in one hand, and of the ONH grids on the other hand, is adapted from the program already used to build the OASIS W&A file of the 3D ocean-BGC coupling described in [12].

In a first step, the algorithm defines the addresses of the non masked grid points of the two corresponding grids and save in file the values in the appropriate OASIS compliant format. The same netCDF file includes the weights, always equal to one, since there is one and only one source grid point corresponding to each target grid point.

In a second step, if needed, a netCDF ORCA input file is read. Masked and non masked grid points are transformed and saved into two files corresponding to the two new input files of the two ONH and OSH SAS models. Every 2D, 3D or 4D variables of the file are processed, unless their first two dimensions do not match with the ORCA grid horizontal dimensions. Double or simple precision real types are supported.

The files described in the following list were needed to start a simulation of the official ORCA2 configuration. Obviously, more files would have to be transformed depending on the `namelist` option chosen.

- `DOMAINcfg`, for grid configuration. The scalar variable `jperio`, `jipglo` and `jpjglo` have to be modified separately. `ORCA` and `ORCA_index` variables must be removed to avoid the hemispheric grid to be treated as global ORCA grids[11]
- Salinity for surface damping
- SAS and ice model restart files

[11] Hopefully, these parameters have few or even no influence on the SAS part of the code (including ice) which means that the new hemispheric SAS computations must be the same than the ORCA global SAS ones

- The W&A file needed to interpolate atmosphere surface variables. For that, a special treatment is performed by our tool, not only to split the grid points but also to correct the addresses included in the `dst*` variables

The Fig. 2 shows the transformation of the bathymetry variable in our ORCA2 case. Notice that the OHN grid is now described following a quasi-polar projection, with array dimensions of 147x91 (182x74 for OSH).

4. Performance

The measurement of the NEMO computing performance must be done following a precise methodology [14]. In particular, a BENCH configuration is mandatory to avoid any disk access perturbations. In our case, the measurement is mainly needed to check that the coupling does not contribute significantly to downgrade the performance (evaluation of the *coupling cost*, see [15]). In that perspective, we use the developed ORCA2 single executable (SE) and coupled (CPL) configurations and simply reproduce the measurement to exclude outliers. Any XIOS output is switched off to limit the perturbations.

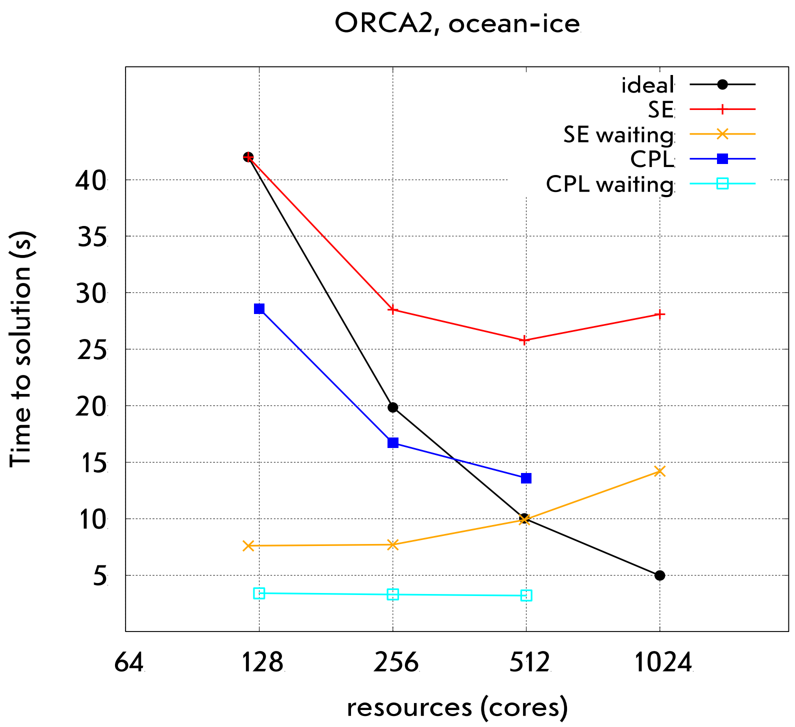

Figure 3: Compared scalability (time to solution) for single executable (SE) and coupled (CPL) configurations. In addition to the total time, the time spent to wait the boundary conditions from the neighbouring subdomains is provided (for the ocean component only in CPL configuration)

In CPL mode, a strict load balancing adjustment between models is needed. 128 cores per node are available in the Météo-France machine. Assuming that no depletion is needed by the executable (e.g. for memory bounding reasons), it means that the total resources allocated to our model must be a multiple of this number. With the additional requirement to select the "magic numbers" that minimise inner load imbalance in the three NEMO models. Land only subdomains are switched off. SE and CPL measurement are compared for the same number of fully allocated nodes (1,2,4 and possibly 8 nodes with 128 cores and MPI processes).

As expected, Fig. 3 shows that the SE time to solution (time loop only, excluding restart reading and writing phases) decreases with parallelism, but stops scaling when decomposition reaches 500 subdomains, with local dimensions 7x7 (red line). The increasing waiting time observed for boundary exchanges (`lbc_lnk`) among all MPI processes is identified as the scalability bottleneck (orange line). This is clearly not the case for CPL (blue lines). It can be explained by the respective decompositions of the three executables. Most of the resources used are allocated for the ocean, and only a small number for the ice models (between 4 and 12). This leads to practically keep the same performance in the ocean, and perform the sea-ice computations in parallel (at the same speed) instead of sequentially.

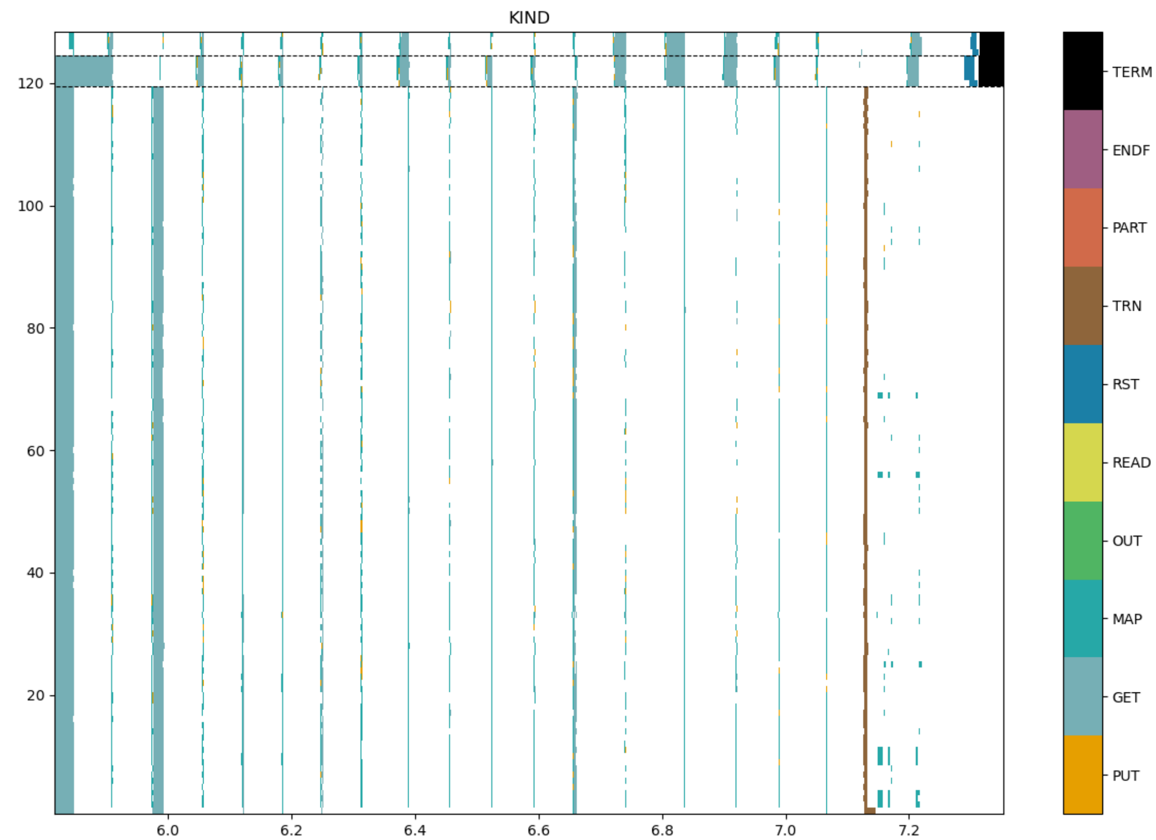

Figure 4: Timeline of OASIS related events during the time loop phase of an ocean-only/Arctic-SAS/Antarctic-SAS coupled simulation. Events occurring in the ocean are display in the first 119 lines (MPI decomposition), Arctic-SAS ones in the next 5 and Antarctic-SAS ones in the last 4

It also means that scaling is not yet limited in the ice models, which enhances the overall performance on 2 or 4 nodes. The small level of decomposition of these SAS models also favorably influences the speed of forcing reading in files. To summarise, these measurements seem to ensure that, in any cases, the CPL configuration would give comparable or better results than SE (+30% of speed in some cases).

Coupling cost is an important issue to fix before being able to deliver a good performance with an OASIS coupled system. In addition to the component load balancing work previously described, we must mention that we tried to minimise the coupled field communication cost by grouping the exchanges by grid kind. The load balancing information provided by OASIS can be visualised on Fig. 4. The colored areas mainly represent the moments when processes are stopped to wait coupling fields (blue boxes, GET legend). This "waiting time" or load imbalance between component stays small compared to the period when models are doing their calculations (white areas).

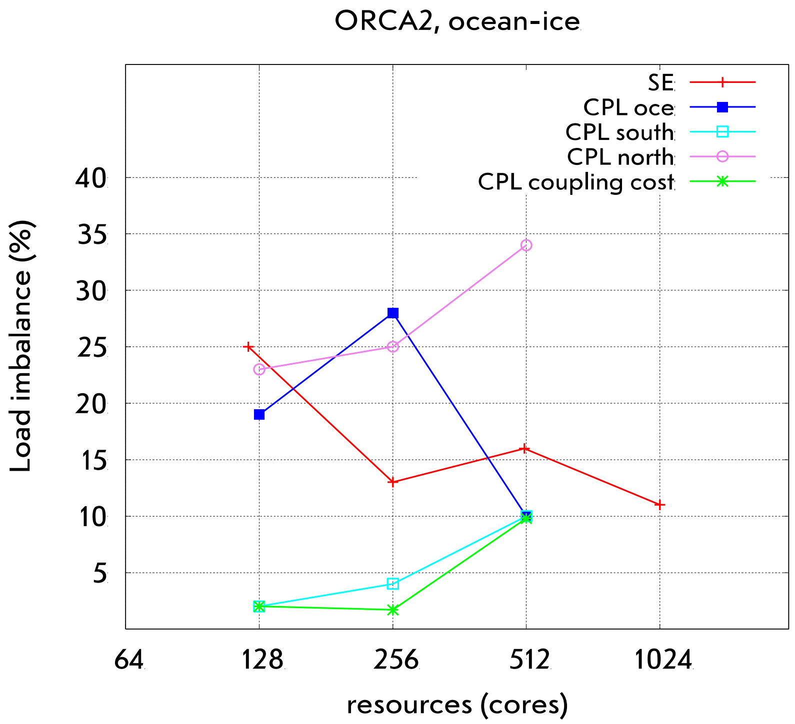

Figure 5: Compared load imbalance for single executable (SE) and coupled (CPL) configurations. In SE, the load imbalance is calculated as the ratio between the difference between slowest and fastest subdomains and the time to solution. In CPL, this quantity is given for the three components. In addition, the coupling cost is provided for this configuration

The Fig. 5 shows the average value of this CPL coupling cost. A good load balancing can be reached on 1 or 2 node. At 4 nodes, the NEMO best decomposition and core number per node requirements makes the coupling increasing by 10% the total simulation cost. One would also notice on this same graphic the surprisingly high load imbalance of the SAS ONH model, probably due to the small number of subdomains available at that resolution and the high number of masked grid points, unfavourably spread on the grid. This hypothesis would have to be confirmed with higher horizontal resolution and larger decompositions.

5. Perspective

A coupled configuration is now able to be adapted to include the neXtSIM sea-ice model instead of SI3, as soon as this model will be discretised on the Arakawa grids. Two hemispheric grids are defined to simplify their MPI spatial domain decomposition. In a second step, the existing NEMO interface is also ready to technically receive fluxes from an atmosphere model. In this case, the developers will take care of the interpolation used. Our implementation avoids any interpolation downgrading between ocean and sea-ice and it would be unfortunate that this downgrading took place between the sea-ice and the atmosphere. One would take care of the errors produced by a too big resolution difference between atmosphere and ocean, particularly near the ocean-sea-ice transition zone, as emphasised in [15].

A secondary result of this development seems to prove the good effect on performance brought by this extra level of macro-task parallelism (two coupled sea-ice model instead of one). Our configuration also offers a test bed for the developers convinced of the weakness of the NPF implementation, by proposing an ORCA like grid, free of extra communication in that region.

The authors want to acknowledge Sophie Valcke, Rym Msadek, Einar Ólason & Pierre Rampal for the introductory discussion that led to this implementation. They also acknowledge the use of the Ferret program for analysis and graphics in this report (Ferret is a product of NOAA's Pacific Marine Environmental Laboratory). Thanks to Thomas Williams & Colin Kelley for the development of the Gnuplot program, which analysis and graphics are displayed in this report, in addition to graphics from Matplotlib, a Sponsored Project of NumFOCUS, a 501(c)(3) non profit charity in the United States. This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement No 823988 (ESIWACE-2). Computing resources were provided by Météo-France.

References

[1] Rampal, P., Bouillon, S., Ólason, E., & Morlighem, M., 2016: neXtSIM: a new Lagrangian sea ice model, *The Cryosphere*, **10**, 1055– 1073, https://doi.org/10.5194/tc-10-1055-2016

[2] Gurvan Madec and NEMO System Team, "NEMO ocean engine", *Scientific Notes of Climate Modelling Center (27)* - ISSN 1288-1619, Institut Pierre-Simon Laplace (IPSL), http://doi.org/10.5281/zenodo.1464816

[3] Voldoire et al., 2019: Evaluation of CMIP6 DECK experiments with CNRM-CM6-1, *Journal of Advances in Modeling Earth Systems*, https://doi.org/10.1029/2019MS001683

[4] https://sasip-climate.github.io/

[5] Craig A., Valcke S., & Coquart L., 2017: Development and performance of a new version of the OASIS coupler, OASIS3-MCT_3.0, *Geosci. Model Dev.* , **10**, pp. 3297-3308, doi:10.5194/gmd-10- 3297-2017

[6] Maisonnave, E., & 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

[7] Maisonnave, E., Coquart, L., & Piacentini, A., 2020: A better diagnostic of the load imbalance in OASIS based coupled systems, Technical Report, **TR/CMGC/20/176**, CECI, UMR CERFACS/CNRS No5318, France

[8] Maisonnave, E., 2019: OASIS Dedicated Support, 4th annual summary, Technical Report, **TR/CMGC/19/149**, CECI, UMR CERFACS/CNRS No5318, France

[9] NEMO Sea Ice Working Group, Sea Ice modelling Integrated Initiative (SI3) — The NEMO sea ice engine", *Scientific Notes of Climate Modelling Center (31)* - ISSN 1288-1619, Institut Pierre-Simon Laplace (IPSL), http://doi.org/10.5281/zenodo.1471700

[10] Masson, S. et al., 2012: Tropical Channel NEMO-OASIS-WRF Coupled simulations at very high resolution. In *13th WRF Users' Workshop* (pp. 25-29)

[11] Dupont, F., Higginson, S., Bourdallé-Badie, R., Lu, Y., Roy, F., Smith, G. C., Lemieux, J.-F., Garric, G., & Davidson, F., 2015: A high-resolution ocean and sea-ice modelling system for the Arctic and North Atlantic oceans, *Geosci. Model Dev.*, **8**, 1577–1594, https://doi.org/10.5194/gmd-8-1577-2015

[12] Maisonnave, E., Berthet, S., & Séférian, R., 2021: OASIS based grid coarsening of TOP-PISCES biogeochemistry in the NEMO ocean model: performance, Technical Report, **TR/CMGC/21/201**, CECI, UMR CERFACS/CNRS No5318, France

[13] Irrmann, G., Masson, S., Maisonnave, E., Guibert, D., & Raffin, E., 2022: Improving ocean modeling software NEMO 4.0 benchmarking and communication efficiency , *Geosci. Model Dev.* , **15**, 1567–1582, doi:10.5194/gmd-15-1567-2022

[14] Balaji, V. et al., 2017: CPMIP: Measurements of Real Computational Performance of Earth System Models in CMIP6 , *Geosci. Model Dev.*, **46**, 19-34, doi:10.5194/gmd-10-19-2017

[15] Maisonnave, E., & Voldoire, A., 2018: Exchange grid: implementation of a test configuration to evaluate fluxes modification , Working Note, **WN/CMGC/18/47**, CECI, UMR CERFACS/CNRS No5318, France