OASIS based grid coarsening of TOP-PISCES biogeochemistry in the NEMO ocean model: performance

E. Maisonnave*, S. Berthet°, R. Séférian°
* CECI, UMR CERFACS/CNRS No5318, France
° CNRM, Université de Toulouse, Météo-France, CNRS, Toulouse, France

TR/CMGC/21/201

Abstract

Even though the coarsening strategy relevance in production mode could be jeopardised by the disk access cost, and even though, because of the load imbalance, the cost of coarsening operation involving a large number of computing nodes and performed at each time step becomes as expensive as the whole BGC computations on the coarsened grid, the OASIS based coupled configuration is two time cheaper than the full resolution ORCA025 ocean-BGC model. In a context of energy saving and carbon footprint reduction, this result cannot be neglected. But the importance of the coupling cost revealed in this study emphasises the reality of the MPI parallelism limits, that one can not only notice in the exchange of NEMO variable boundaries, but that must be lowered each time than other tasks, usually neglected, have to be operated in addition to those needed for the core model equation resolving: coupling, diagnostics, resolution coarsening or resolution refinement

Table of Contents

Decades of computing power growth of the Von Neumann architectures used in Climate Modelling promoted the development of ever finer resolution configurations. Exponential growth is still there [1] but it is now dependant of chip specialisation [2] that could not necessarily address all the multi-physics computation needs of our community. At the same time, the computing cost of numerical simulations [3] is an emerging concern in low carbon emission laboratories. For these reasons, the reduction of unnecessary model resolution when possible becomes an interesting research topic.

The inclusion of ocean biogeochemistry (BGC) for carbon cycle modelling in ESMs [4] is computationally demanding, particularly if mesoscale eddy representation is targeted. The explicit computation of chemical species reaction and advection can significantly increase the NEMO ocean [5] cost. On our most powerful supercomputers, this makes too expensive to perform, at global spatial resolution higher than 1°, large simulation exercises like CMIP [6] if the BGC module, called TOP-PISCES [7], is enabled .

A finer representation of the ocean circulation would increase ESMs reliability, but the same dependency to resolution is not as clear for BGC processes. This leads to the idea of a coupled system that includes ocean and BGC submodels with different spatial resolution (multi-grid). This strategy was already validated with our model [8]. In this implementation [9], the reduction of the BGC grid resolution (coarsening) was done locally and the ocean dynamics / BGC computations performed sequentially. This choice has two main drawbacks: (i) for parallel computing, the ocean and BGC component grids must be decomposed into the same number of subdomains, which limit the overall scalability to the less scalable component (Amdahl law) and (ii) since the two components are sharing the same physical resources, ocean dynamics and BGC computations must be performed sequentially.

In the new implementation attempt described in this document, we propose to follow the strategy already tested to separate NEMO ocean and sea ice components [10]: the surface module, which includes the sea ice, was separated from the ocean and launched as a separated executable. Communications of surface quantities between the two executables were ensured by the OASIS coupler [11]. The separation of the two submodels in two executables allowed to perform their computations on separate hardware resources, at their optimum parallelism, and to perform them concurrently, increasing what we call the macro-task parallelism [12].

The ocean/BGC splitting into two separate executables and the communication of the coupling fields at their interface is described in a previous document [13]. We reproduced the results of the original mono-executable configuration. The last step towards a multi-grid configuration is described in the present document.

Quickness of the coupling exchanges mainly determines the relevance of our implementation, regarding to performance. In the mono-grid configuration (global 1°, ORCA1), computing performance was enhanced by the possibility of differentiated scalability and the macro-task parallelism. However, the extra cost induced by the exchange between the two components of several 3D variables at each model time step was relatively big (around 20%). The model horizontal resolution targeted in the community (CMIP compatible configurations, global

0.25°, ORCA025) is supposed to increase this coupling cost. This is why a preliminary analysis of the potential gain induced by the BGC resolution coarsening must be estimated beforehand.

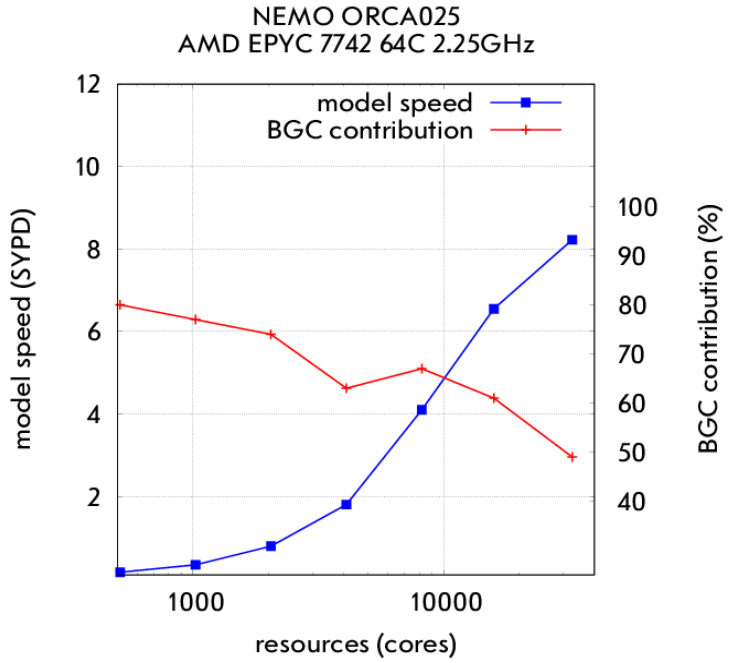

Figure 1: mono-executable configuration scalability (blue) and BGC contribution to the total simulation time (red)

The maximum gain is estimated by measuring the model scalability and the splitting of the computation timings between the two components (BGC and ocean dynamics). The BGC contribution gives the maximum gain limit of BGC horizontal resolution reduction in a BGC grid coarsened configuration. A precise methodology to measure NEMO scalability is definitely required. As described in [15], only a subset of all possible two dimensional domain decompositions must be chosen[1] to plot the model scalability as done in Figure 1. Wrong choices would superpose enough noise to change the analysis. The high number of hardware computing units (cores) per node, and the necessity to make use of all of them (128 in our Météo-France AMD based machine), add another constraint to the possible decompositions number (multiple of 128, or close to it).

It is also important to mention here the strong perturbation in the measurement caused by any disk access (reading or writing). Since it is not possible to use here the appropriate I/O free benchmarking configuration (ibid.), assuming that only numbers related to real production configurations would satisfy the reader curiosity, the control of the time step duration in a large number of attempt was necessary to select non outlier realisations and provide average values for each chosen decomposition.

Figure 1 shows that BGC contribution to the overall model is decreasing with parallelism. The scalability of our model is bounded by the amount of communication necessary to exchange

1 The `ln_listonly` namelist option gives this optimal decomposition for any grid, taking into account land point only subdomains, where no computing resource must be allocated [15]. Wrong numbers necessary bring load imbalance between subdomains.

boundary values at sub-domain limits (ibid.). This amount increases in conjunction with the amount of computations performed in ever small subdomains. A look to the communication analysis provided in timing mode (ibid.) shows that more boundary communications are required in the ocean (particularly ice) related subroutines. The higher BGC scalability induces that its contribution decreases with the resource number. At low parallelism, the potential gain of reducing the BGC contribution by coarsening is 80%, but only 50% at scalability limit. In other words, the coarsened configuration can be 5 time faster on a small number of resources, when the model speed is low. But it cannot be more than 2 time faster on a large number of resources, when the model reaches its maximum speed. We decide that the magnitude of potential gains on the central part of the curb justifies the development of the ocean-BGC coarsened coupled model.

To do so, we choose to coarsen global grids of the NEMO model (ORCA class grids). The specificity of the grid geometry (periodicity but mainly tri-polarity with duplicated fold lines at North Pole) limits the choice of coarsening factors. These limitations are clarified in Sect. 1.

A coarsened BGC model grid is derived from an existing ORCA ocean grid, following rules that guarantee the conservativeness of the coarsening methods. This new grid differs from any existing ORCA configuration, which means that its bathymetry and all the derived input and forcing files has to be build accordingly. The building rules are described in Sect. 2.

The targeted resolution of the ocean grid makes necessary the enabling of the variable volume vertical coordinate system. This increases the number of coupling fields exchanged via OASIS, as listed in Sect. 3. To reduce the coupling cost induced by this increase, new algorithms are implemented in the OASIS library, particularly to handle large 3D fields at every model time step and to provide efficiently the maximum value of source grid neighbours, as proposed in Sect. 4.

We propose at Sect. 5 to review the final measurement of the costs of a coarsened coupled model, an ocean-ORCA025/BGC-ORCA075 in its best synchronised configuration, which makes possible the evaluation of the benefit achieved with our solution.

1- Coarsening factor

In this study, we propose to define and implement a three dimensional conservative interpolation between two ORCA grids with different horizontal resolutions and the same vertical resolution. The ratio between the target and source grid horizontal resolutions is called coarsening factor (CF). The ratio is isotropic, to simplify the lateral flux conservation issue. The ratio is bigger than 1 and defines, for each exchanged quantities, the number of source grid point values necessary to calculate each target grid point value (CF^2). This necessarily imposes a relation between the i and j horizontal dimensions of the two grids, which is not simply :

```
jpit = jpis / CF and
jpjt = jpjs / CF
```

On the X (longitudes) axis, computations across the edge of the grid periodicity limit requires the definition of two duplicated columns at each grid j sides. The relation between source and target i dimension is then:

```
(a) jpit = 2 + (jpis-2) / CF
```

Even though we cannot find any specification in the model documentation, we also assume that, to ensure the exact folding of the North pole line, this jpit value must be an even number.

On the Y (latitudes) axis, the tri-polar geometry of the ORCA grid strongly complicates the formula. We call latitude periodicity (`jperio`) the kind of folding strategy applied to allow computation across the North Pole line limits. The last line (`jperio=6`) or the last and half of the penultimate line (`jperio=4`) can either be duplicated. The relation between the j dimension of source and target grid depends on this periodicity, but also changes depending whether CF is an odd or even number.

In Fig 2 and 3, we summarise how the different geometries define the 3 possible relations between the Y grid dimensions.

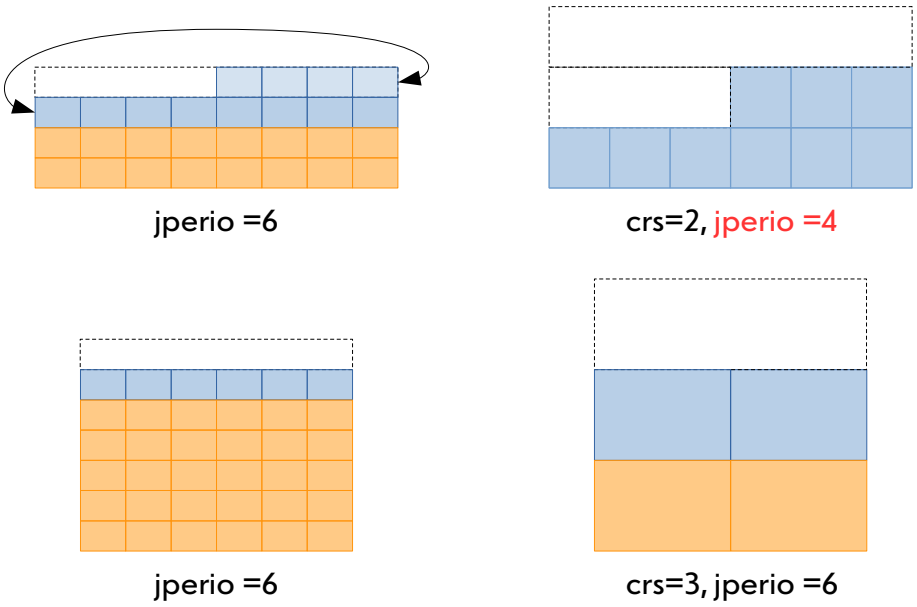

Figure 2: Correspondence between source (upper and lower left) and target (upper and lower right) grid points a coarsening operation applied to a `jperio=6`-periodic source grid, for an even (e.g. =2, upper graphics) or odd (e.g. =3, lower graphics) CF. The duplicated North Pole folding lines are represented in blue (effective grid points) and white (duplicated grid points). The arrowed line represents the operation needed to build the target penultimate half line.

For a source grid with `jperio=6`, odd or even CF are permitted, but for an even number, the latitude periodicity must change since only CF/2 source lines are available at the top of the array to build only half of the penultimate target line. In this case, the number of target lines must be :

```
(b) jpjt = ( jpjs - 1 + CF / 2 ) / CF + 1
```

For even CF numbers, it must be :

```
(c) jpjt = ( jpjs - 1 ) / CF + 1
```

For a source grid with `jperio=4`, not all CF numbers are permitted. As illustrated in Fig 3, half of the penultimate but one line remains unused in case of even CF numbers, which would contradict the conservation constraint.

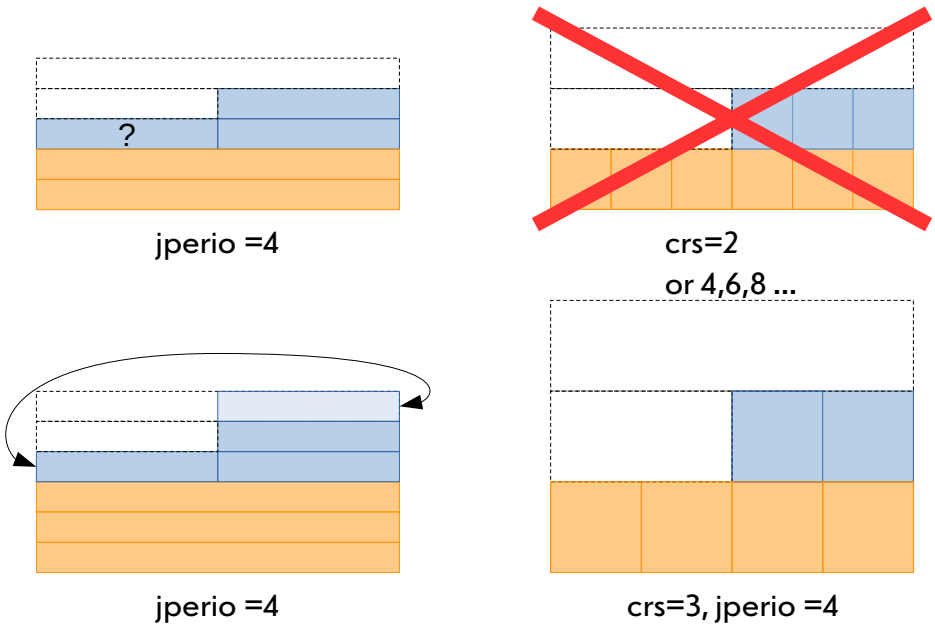

Figure 3: same than Fig 2 for a `jperio=4`-periodic source grid. The arrowed line represents the operation needed to build the target penultimate half line.

With a `jperio=4` latitude periodic source grid, only odd numbers are permitted with the following relation:

```
(d) jpjt = ( jpjs - 1 + ( CF - 1 ) / 2 ) / CF + 1
```

2- Coarsened grid

In this study, we focused on the ORCA025 (¼ degree of horizontal resolution) source grid (`jperio=4`), which resolution would match future requirements for a high resolution of CMIP carbon cycle simulations. The ocean model variables are coarsened to a biogeochemistry grid ORCA075 (necessarily defined with a `jperio=4`), with a CF=3. According to the rules defined in Sect. 1, since dimensions of the source grid are 1442 x 1050, the horizontal dimensions of the BGC grid are 482 x 351.

To build the ORCA075 grids (all T,U,V,F and W), we rely on the definition given at Sect. 3 in Bricaud et al. [9] :

- the grid point positions are defined by a subsampling of the ORCA025 ones,
- the mesh scale factor is given by the sum of the ORCA025 mesh scale factors in both zonal and meridional directions,
- a mask is applied if all corresponding ORCA025 grid points are masked (see Table 3 for a grid type related definition).

The ORCA075 grid properties derive from the ORCA025. Since the BGC model is an independent model in the coupled system we are building, its grid must be defined beforehand, which actually means that a configuration file must be defined and provided as an input of the BGC executable.

To do so, a specific NEMO tool is added in the appropriate subdirectory. As a first step, this tool is able to build, from the ORCA025 quantities, the latitude, longitude and scale factor variables and save them in the corresponding `coordinates.nc` and `bathy_meter.nc` files. In a second step, these files are used to produce the ORCA075 configuration file, launching the existing `DOMAINcfg` tool.

Land-sea masks are the only missing variables. They will be calculated in a third step, at the same time than the coarsening operators (see Sect. 4) and added to the configuration file. Notice that these masks are also derived from the ORCA025 ones. It means that, in the BGC model, the standard building of the U,V,F and W masks from the original T mask must be disabled[2] and these quantities read in the configuration file. Test simulation are necessary to better understand the impact of the new kind of relation between the different grids that this new definition implies.

In addition, in order to keep the coarsened mask as realist as possible (conservation of Panama isthmus or Gibraltar straight), additional and, unfortunately, manual modification can be required.

3- Coupling interface

The routines added to the initial NEMO release, similar to the existing atmosphere coupling ones (`sbccpl`), were described in the document related to the first OPA-TOP_PISCES iso-resolution coupling [12]. The algorithm is substantially modified, according to Bricaud et al. and also taking into account the initial implementation of the internal coarsening related subroutines[3]. The light absorption coefficient, a 3D variable produced by the BGC model, is still included in both model interfaces (as a sent field in the BGC model, and as a received field in the ocean one).

2 In the `zgr_read` subroutine

3 See the `crs_` prefixed routines in NEMO 3.6

The new list of ocean to BGC coupling field is provided in Table 1. Runoff or Water balance at 'before' time step can optionally be provided depending on the parametrisation set in the ocean namelist.

Coupling field	Grid	Dimension #
Temperature	T	3
Salinity	T	3
Effective zonal transport	U	3
Effective meridional transport	V	3
Vert. eddy diff. coef. for T (and possibly S)	W	3
Mixed layer depth	T	2
Water balance	T	2
Salt flux	T	2
Sea ice cover	T	2
Solar flux	T	2
Wind speed module	T	2
Sea surface height	T	2
bbl diffusive flux -i	U	2
bbl diffusive flux - j	V	2

Table 1: List of ocean variables provided to the BGC model, the kind of grid where they are discretised and their dimension number

In comparison to the previous interface, the horizontal divergence and the vertical transport were removed, since it seems more accurate to calculate these quantities on the BGC grid, starting from the zonal and vertical effective transports. Also notice that, following Bricaud et al, the logarithm of the vertical diffusion coefficient is coupled instead of the raw value.

In addition to these ocean variables, a set of coarsening factors are also transferred by OASIS, to cope with the new variable volume parametrisation required at such resolution. These quantities are listed in Table 2.

Coupled scale factors	Grid	Dimension
e1t.e2t.e3t (now)	T	3
e1t.e2t.e3t (after)	T	3
e2u .e3u (now)	U	3
e1v .e3v (now)	V	3
e1t.e2t.e3w (now)	W	3
e1t.e2t	T	2
e3t (now)	T	3

Table 2: List of ocean scale factors provided to the BGC model, the kind of grid where they are discretised and their dimension number

They mainly contribute to the conservation of the transferred quantities as follows:

- `log(avt) * e1e2t * e3w_n` in ocean, then `interpolated (avt) / interpolated (e1t.e2t * e3w_n)` in BGC and elevated to the power of 10
- `tsn * e1e2t * e3t_n` in ocean, then `interpolated(tsn) / interpolated(e1e2t * e3t_n)` in BGC
- `un * e2u * e3u_n` in ocean, then `interpolated(un) / interpolated(e2u * e3u_n)` in BGC
- `vn * e1v * e3v_n` in ocean, then `interpolated(vn) / interpolated(e1v * e3v_n)` in BGC

Scale factors in BGC are also calculated from interpolated values :

- `e3t after` in BGC is computed with `interpolated (e3t) / interpolated (e1t.e2t)`

In this study, we limit our implementation to the coupled fields and a basic use in the BGC model. In a future step, a careful work of identification of the BGC variables modified by the coupling field will be led according to Bricaud et al.

4- OASIS adaptation

The OASIS coupling library is the chosen solution to perform the coarsening operations, amongst different others: internal coupling, as previously implemented in [8] or AGRIF based transformations. The latter solution requires to allow the definition of a mother global grid where usually zoomed grids are required and to extend the existing global/regional operations to global/global grids. In this study, we focus on the solution already implemented in order to couple ice [10] or BGC modules [12]. But the solution implemented in these examples was not requiring conservative interpolation between different grids. In this chapter, we describe how the coarsening operations needed were defined within the OASIS formalism (Sect. 4.1) , and how the coupler itself was modified to be able to handle large 3D coupling fields (Sect 4.2) and new operations (geographical maximum, Sect. 4.3).

4.1- Interpolation weight & address

As any other OASIS interpolation, all coarsening operators, including geographical maxima, can be seen as a matrix multiplication of weights by the coupling field values decomposed on a source grid (neighbours), to produce a single value on each target grid point. The definition of these weights depends on the geographical position of the grid points and, for conservativeness, of their areas. In our case, the "boundaries" of the target meshes exactly match the source ones (see Sect. 2). This highly simplifies the computation of the interpolation weights and the identification of the associated source and target addresses. We first

describe the ocean to BGC operations (coarsening), for 2D and 3D quantities, then the reverse operation.

4.1.1- Ocean to BGC 2D fields

A set of interpolations can be produced on the fly thanks to the SCRIP library included in OASIS. We chose not to add the definition of the new coarsening operation at this place, but to implement a new independent FORTRAN program that defines and saves the coarsening operation in a standard interpolation weight & address (W&A) file. This file can be read once at runtime by the OASIS library (MAPPING option of the `namcouple`), that applies the coarsening operators to the coupling fields at each coupling time step.

The definition of the coarsening operators are only valid for grids which definition strictly follows the definition rules of ORCA-like NEMO global grids. In that sense, it would be meaningless to provide the corresponding W&A files in the OASIS framework, which is by definition a community tool, supposed to be able to handle a large set of geophysical grids. For that reason, we chose to develop this program in the NEMO/tools subdirectory. This would also possibly facilitates the upgrade of the comprehensive OASIS based coarsening functions in a future NEMO release.

The proposed implementation is sequential, considering that the production of the full set of coarsening interpolation W&A (for T,U,V,F and W NEMO grids, 2D and 3D, + grid T 3D geographical maximum) and save it on file, for the largest targeted grid (ORCA025), takes about 10 seconds. However, there would be no major issue to produce a parallel version, as the one developed for SCRIP [16].

The algorithm needed to find and associate mesh addresses from source and target grids, as well as the interpolation weights, is straightforward, even though a special care must be given to avoid using twice the source grid points included in the cyclic lateral (resp. polar) fold columns (resp. lines). Because this is not permitted by the association algorithm coded in the MCT library, these bounding lines and colums are not filled by the OASIS interpolation, which means that an additional operation (`lbc_lnk`) is required after receiving the fields in the BGC model.

The program starts by checking whether the grid dimensions and periodicities matches or not the coarsening conditions required by the chosen coarsening factor, as described in Sect. 1. As a consequence of the coarsened grid definition, it is impossible that a non masked target grid point remains without any associated source grid point. To be sure that it is the case, the T,U,V,F and W masks of the target grids are defined at the same time than the source/target grid point association. The two informations are saved in separated `netCDF` files : for each grid type, one for the mask, that will be merged with the NEMO config file, and one for the OASIS W&A file. The two kinds of file will be read by the model (resp. OASIS library) at runtime during the initialisation phase.

If the checking phase is successful, the test continues with the reading of the source grid scale factors that are needed to calculate the weights, and the reading of the source grid masks.

Following the definition of the target grid scale factors, the OASIS weights for 2D field coarsening are defined as :

$$e1x.e2x \;/\; \sum e1x.e2x$$

where `e1x` and `e2x` are the scale factors at each non masked source grid point. The sum is done with every non masked source grid point associated to the target grid point. According to the definition proposed in Sect. 1, the number of associated grid points changes following the grid type (see Table 3). W grid only differs from the T grid by its mask.

Grid type	i-index of associated source	j-index of associated source
T	1 to CF	1 to CF
U	CF	1 to CF
V	1 to CF	CF
F	CF	CF

Table 3: indexes in I and j dimension of the grid points involved in the coarsening operations, among all the CF times CF source grid points possibly associated, depending on the NEMO grid type

A special weighting is required to apply the geographical maximum (grid T only) and a specific W&A file is created for that purpose. In this case, a uniform weight equal to 1 is prescribed for each unmasked source grid point. A specific transformation described in Sect 4.3 allows to provide to the target grid point the geographical maximum amongst the unmasked source grid points.

4.1.2- Ocean to BGC 3D fields

Since the variable volume is activated[4] in our target NEMO configuration, the vertical scale factor (`e3x`) is time dependant. As a consequence, the conservativeness of coarsening operation on 3D quantities must involve a variable that changes during the simulation. As our OASIS weights are calculated once at the simulation start, the standard strategy applied for 2D field coarsening cannot be the same for 3D exchanges.

The 3D coarsening operations are then decomposed in 3 phases:

1. At every time step, in the ocean model, the coupling fields are multiplied by the scale factors, possibly including the time dependant `e3x`. The ocean model sends these modified coupling fields, and it also sends the scale factors themselves,
2. The coupler perform a standard matrix multiplication with the coupling fields as defined in phase 1 and weights equal to 1. Said differently, the OASIS operation is a simple addition of the non masked source grid points. OASIS performs the same operation on the associated scale factors,
3. In the BGC model, both fields and scale factors are received. The final operation consists in dividing the field by the scale factor. Depending on the kind of field, the

4 `ln_linssh = .FALSE.`

conservativeness must be insured by also involving the target grid factors in the conservation equation.

4.1.3- BGC to ocean 3D field

On the other way round, the only quantity to communicate is the light absorption coefficient, that does not require conservative interpolation. We take benefit of the previously described program (see 4.1.1), that already establishes the source-target grid association. By changing the exchange order, any ocean grid point receives a coupled value from one single BGC grid point, with an interpolation weight coefficient equal to one. In the future, this interpolation could be replace by a smoother one, if the resulting coupling field on the ocean grid is found too patchy.

4.2- Volume partitioning

The coupling of 3D variables is already possible in the current version of the OASIS coupler but suffers restrictions that are not compatible with our needs. The 3D interpolation of ocean quantities between different grids necessarily implies to deal with a changing bathymetry between vertical layers, thus different mask at each levels. The standard so called "bundle management" provided in OASIS only propose to group 2D coupling fields processed identically. In our case, due to a different mask at each level, a different interpolation is required at every vertical level of our 3D fields. For computing performance reasons, the high number of 3D fields exchanged in our ocean-BGC interface forbids to define a specific interpolation at each level. For the same reasons, we also prefer to avoid an interpolation operation on unmasked arrays, larger thus more costly to manipulate. It is then necessary to define a new method to efficiently interpolate our 3D fields.

We proposed to consider every 3D field as a single 1D vector. The interpolation matrix defined in the Sect. 4.1.2 is a simple identity matrix. The source/target grid correspondence takes into account the specific mask at each layer. The grid point index varies from 1 to the total grid point number in the 3 dimensions. This interpolation is externally built, but could possibly be included in the SCRIP interpolation package for a runtime computation, if necessary.

In order to apply this new interpolation to our 3D coupling fields, it must be possible to partition the global domain, now a 3D domain, into also 3D subdomains. A new partitioning option (`CUBE`) is added and made available in the OASIS API (`oasis_part_setup`) and the get-put interface modified to be able to treat 3D input arrays as unbundled data (`mod_oasis_getput_interface.F90`).

The purpose of these modifications, to test the possibility to perform coarsened ocean-BGC simulation with OASIS, is too specific to justify an immediate inclusion in the next OASIS release. This is why we delay the additional coding of input/output in 3D netCDF variables that must be mandatory for diagnostics in production mode.

The impact of this new partitioning option on computing performance is evaluated with the original ORCA1 ocean-BGC coupled system used in [12]. In this configuration, both models are discretised on the same grid, no interpolation are needed. A simple field mapping from different MPI decompositions in horizontal sub-domains is the only OASIS transform needed. We compare the time spent to map the source information onto the 256 target horizontal sub-domains, using 128 cores of a single node of our machine.

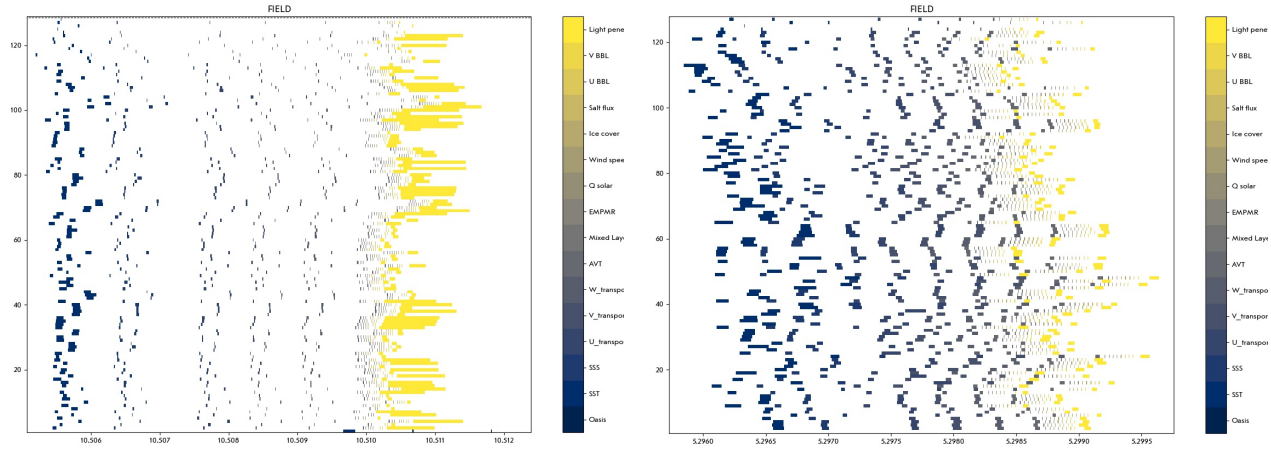

Figure 4: Comparison of two timelines (x-time axis limits differ), for one coupling exchange, showing mapping operation timing (color boxes) on the 128 resources of the ocean model, when vertical fields are horizontally decomposed level by level following a BOX decomposition (left) or are decomposed together following the new CUBE decomposition (right)

We take benefit of the new load balancing analysis available in OASIS [17] and its visualisation tool [18] to produce the two plots shown in Fig 4. We extract from the complete timeline of OASIS related coupling events only one sequence of coupling field exchanges, i.e. 14 coupling fields from ocean to BGC, from which 6 3D arrays (dark color), and 1 BGC to ocean coupling field (yellow color). The mapping operations are shown on the ocean part only (128 resources, y-axis), where most of the time is spent. The BOX decomposition tends to slower the overall mapping operation, not only per se, but also because it seems to need more synchronisation at each coupling field exchange, since the white areas that separates each of them are larger. In addition, it appears that the decomposition declaration phase in OASIS[5] (not shown) is much faster in CUBE configuration. This validated, from performance point of view, the use of our new CUBE decomposition to efficiently exchange the ocean-BGC 3D fields.

4.3- Geographical maximum

At Sect. 3.1.4 of their paper, Bricaud et al. define an extra variable in the CRS grid, necessary to conserve the source grid thickness in the vertical gradient operator computations. This variable, named `e3tmax`, is calculated as the maximum value of the `n` source grid points `e3t` vertical scale factors.

5 During the `OASIS_enddef` call, at initialisation phase

The OASIS coupler is designed to apply weight factors to n source grid point values, with

```
n = number of source grid point neighbours of each target grid point
```

and to add up the results of these multiplications to build the target grid point value. To substitute a maximum operator to this multiplication+sum operations, a modification is necessary in the OASIS library dedicated to the mapping: MCT [19]. An argument is added to the sMatAvMult_DataLocal subroutine to simply switch from the standard matrix multiplication and accumulation to a maximum selection of the source contributions. This argument is a character string that can be changed, in the namcouple OASIS parameter file, from the default avg (average) option to max, opening the possibility to define other operators in the future. This namcouple new parametrisation choice also implies modifications in the PSMILE library[6], to communicate the information from the namcouple interface to the MCT subroutine. The computing performance of this new operation is evaluated using again the OASIS internal load balancing tool: it shows no significant difference compared to the performance of a standard matrix multiplication on the same ORCA1 gridded array.

5- Performance

The main justification of the reduction of BGC related computations is an anticipation of a limitation of our hardware resources. This limit is already reached during large simulation campaign such as CMIP. During the previous CMIP6 exercise led with ESMs, the ORCA1 horizontal resolution was targeted. ORCA025 is the next identified step towards representation of mesoscale dynamics. For that reason, we proposed to lead our computing performance measurement with an ocean model of this resolution and a BGC module resolution reduced with a coarsening factor of 3 (ORCA075). All the measurements presented below are made on the ATOS Sequana XH2000 machine recently provided at Météo-France[7].

5.1- Synchronisation

As any other OASIS coupled system, computations performed by the two ocean and BGC models during one coupling time step (here equal to one model time step or 900s) can occur sequentially or concurrently, taking benefit of the LAG option, where initial coupling conditions are provided on a restart file. Considering that ocean and BGC processes are mapped on independent hardware resources (cores), it becomes always more valuable to keep both ocean and BGC resources busy and compute both ocean and BGC equations concurrently.

6 mod_oasis_namcouple, mod_oasis_map, mod_oasis_coupler and mod_oasis_advance
7 https://www.top500.org/system/179853/

The ocean-BGC exchanges are strongly unbalanced. The volume of exchanged data is clearly bigger in the ocean-BGC way than in the opposite BGC-ocean one (only one 3D field). In addition, this latter has no effect on model conservativeness and can easily be delayed of one time step for computational efficiency reasons. Concurrent coupling, that implies that both model were receiving coupling fields from the previous time step, is not as critical as in ocean-ice coupling, for example. For these reasons, the concurrent two way coupling is the configuration chosen to validate the physical results.

From a performance point of view, the BGC feedback to ocean adds a synchronisation at each time step and slows down the model that need to compute the reverse coarse to fine interpolation. Actually, the extra cost added is of the same order of magnitude than the cost of one 3D ocean to BGC coarsening operation. It is not negligible, but for commodity reasons, we decide to lead the scalability test in a one way coupling mode.

5.2- Scalability

The scalability of the coupled system is measured following the same requirements than those detailed in introduction chapter. Only a discrete subset of 2D decompositions is allowed to avoid wasting cores, but in this case, for both ocean and BGC model grids. In addition, as explained in Sect. 5.1, the concurrent mode of coupling is preferred, which means that a balance between the two component speed must be found, thus ends to reduce the possible values to a few set of ordered pairs. It is necessary to precise that the load imbalance between the two coupled components, by minimisation of the waiting time for the fastest model, would not have been achieved without the help of the coupling events measurement and visualisation tool available in the future OASIS-MCT 5.0 version. The difficulty comes from the irregular frequency of the coupling exchanges, due to the heterogeneity of computation, communication and disk access per coupling time step, but also per model subdomain.

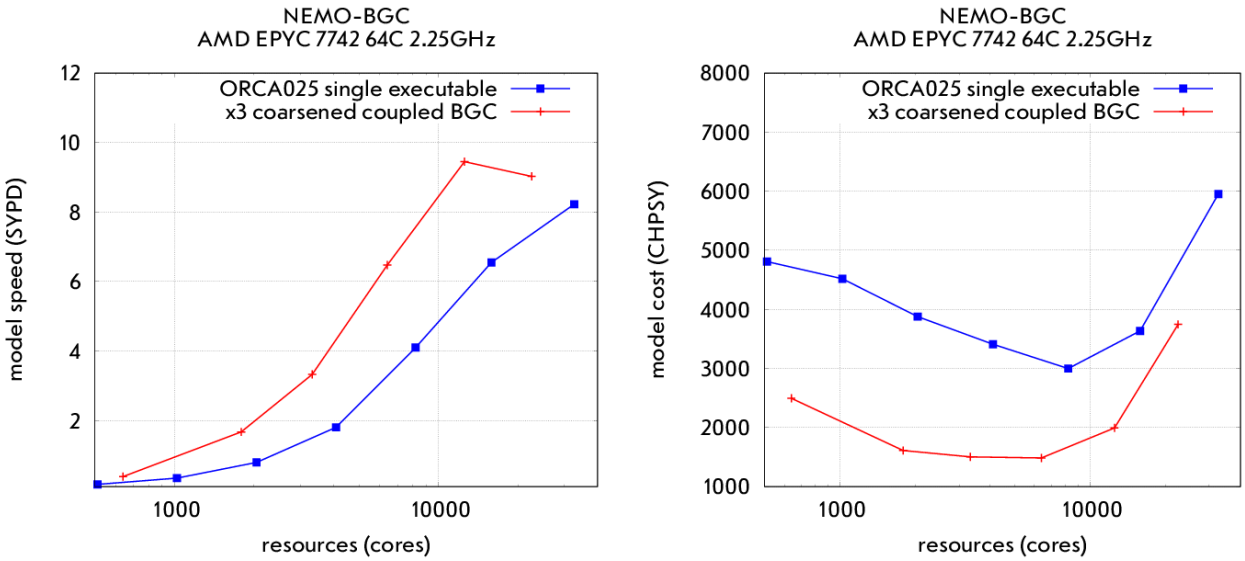

Figure 5: Speed (left) and cost (right) of mono-executable (blue) and coarsened coupled (red) configurations. X-axis shows the total number of hardware resources used (i.e. the sum of the resources needed for ocean and for BGC models in coupled case)

On Fig. 5 (left), the total models speed, excluding initialisation and termination phases, are compared. Below 20,000 cores, the coarsened coupled system is about two time faster than the ORCA025 mono-executable. This ratio decreases at above 20,000 cores, where parallel efficiency is already low. On the right Fig 5, cost is plotted instead of speed. A computation cost is defined in core.hour per simulated year (CHPSY) as the product of the time to solution by the number of cores used to achieve it [20]. Here, the cost ratio between the two configurations falls below 2 when using more than 6,000 computing cores.

The two fold increase of the model speed was predicted in introduction as the upper limit of the possible gain induced by the coarsening. But this was valid for the highly parallel decompositions, i.e. the very right side of the plot. On the left side (low decomposition), the ratio has the value of five.

5.3 - Coupling cost

The reasons of this loss of performance are many. Regarding cost, the upper limit of possible gain is calculated considering that the coarsened BGC computations have a zero cost. In reality, this cost is not zero and can be estimated as the cost of the ORCA025 BGC model divided by a factor CF^2 (= 9 in our case).

Figure 6b represents the kind of activities taking place during one time step on every resources allocated to the coupled model. This picture generalises the observations made with the new OASIS load imbalance measurement tool [18]. The white area symbolises the activity of the resources allocated to the ocean model. At the end of the period, their activity switches to OASIS interpolation and remapping operations (red box) necessary to coarsened the coupling fields and send them to the BGC resources. The BGC computations are pictured in green and take place on a distinct set of resources. The non reducible period spent to wait the reception of coupling fields is represented in orange.

On Figure 6a), the ratio of each coloured box areas by the white box area is represented. In other words, the cost of each kind of non ocean computation is compared to the main ORCA025 ocean cost. These costs are estimated with two methods, that approximately give the same results:

1. our new OASIS load balancing analysis is able to measure, vith `MPI_Wtime` commands, the total time spent during receiving or sending operations. This roughly gives the OASIS mapping (for ocean resources) and coupling load imbalance (for BGC resources)

2. in a special experiment, the exchanges are reduced to one at the beginning and one at the end of the simulation and the same `MPI_Wtime`, called in the model, gives the time spent during the time loop. By subtracting the two time loop timings (coupled at time step frequency and coupled twice per simulation) for the two models, we get an approximation of the two coupling costs

Focusing on the central part of the plot, the green curb shows that the ORCA075 BGC computations require 20% more resources than the ORCA025 ocean computations, which roughly matches with the ratio between the ORCA025 BGC extra cost (x3, see Figure 1) and the estimated coarsening effect (/9). We see that the additional cost required by the OASIS coarsening operations has an equivalent magnitude (30%). A small additional losses, due to the impossibility to fully balance BGC and ocean model speeds, must be added to the OASIS coupling cost. This cost is almost scalable, even if we observe an increase at model scalability limit. As described in introduction, the scalability of BGC routines is higher than the ocean part, which explains the decreasing contribution of BGC computations to the overall cost: with more than 10,000 resources, ocean parallel efficiency starts to decrease while the BGC one is still good.

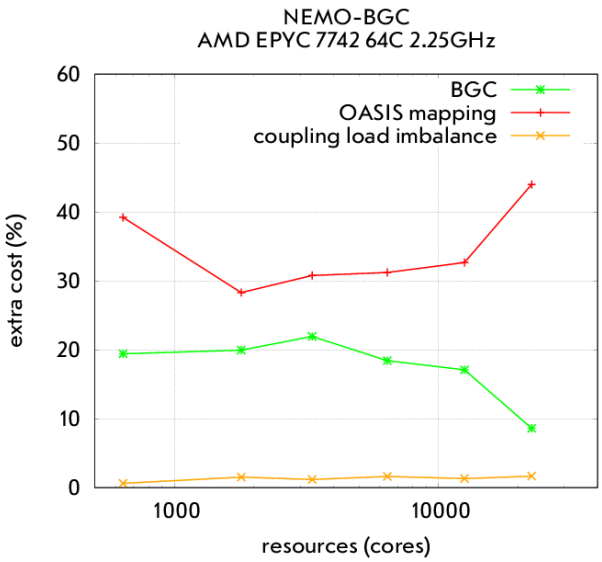

Figure 6a) Ratio of additional cost to the ORCA025 ocean model

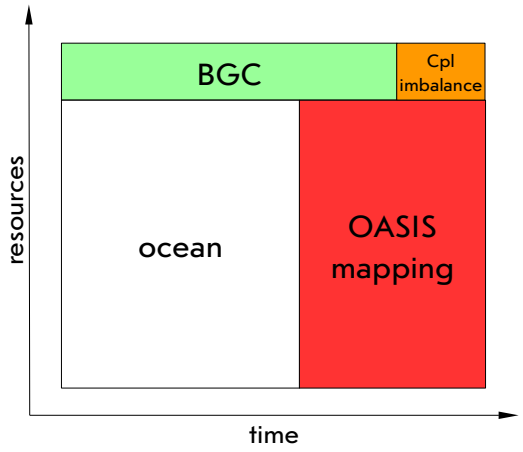

Figure 6b) Schematic picture of processes activity during one coupling time step

Regarding speed, one could hope that the concurrently coupled configuration would go at the same pace than the ORCA025 mono-executable one. This is not the case, because of the additional time needed to perform the OASIS remapping computations, that cannot be deported on the BGC resources. This explains that the coarsened coupled model speed plotted in Fig. 5 (left) is two time smaller than expected.

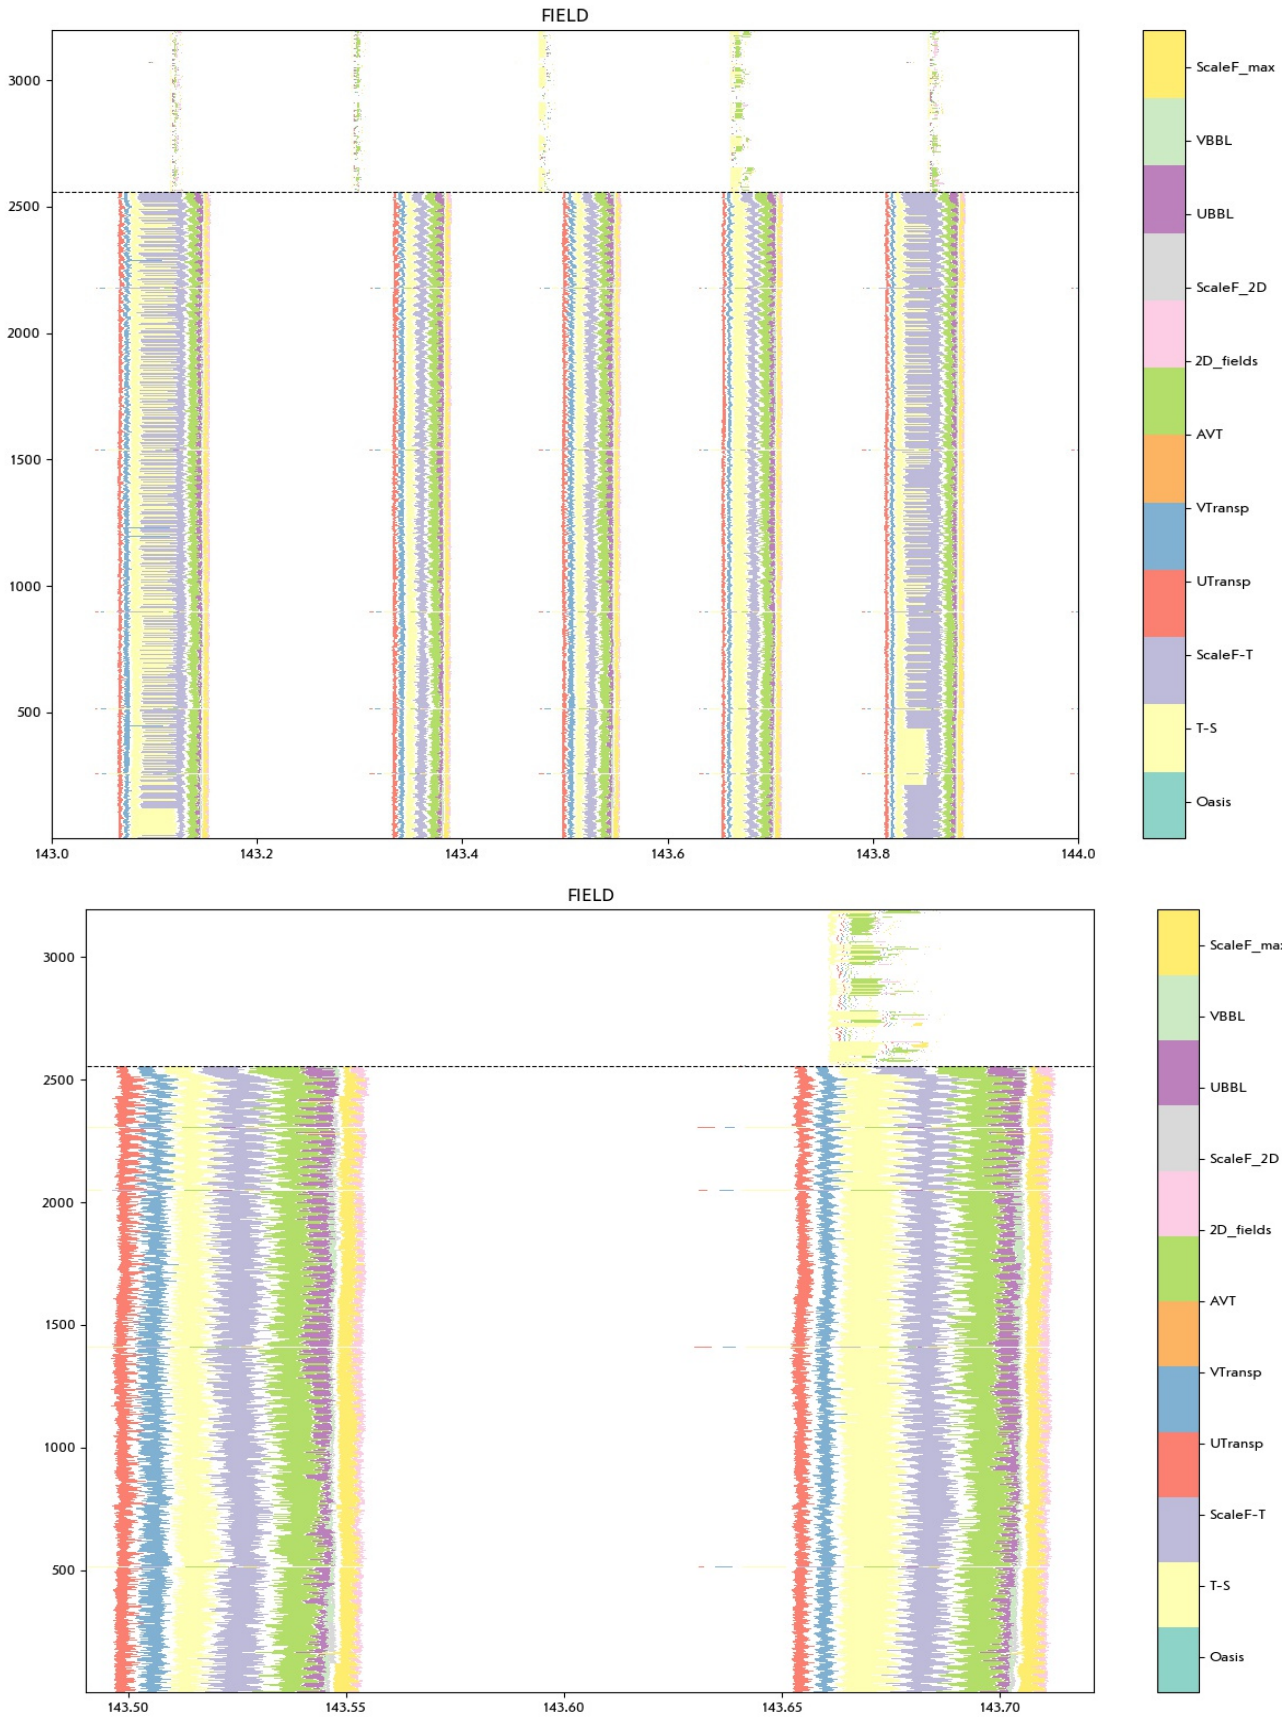

Figure 7: Coupling event timeline, displayed for 1 second (upper) and zoomed to 0.2 second (lower) elapsed time (x-axis) of a sequential one way coupling ocean_ORCA025-BGC_ORCA075 coarsened simulation performed on 2557 (ocean, lower part of the plot) and 640 (BGC, upper part of the plot) cores (y-axis). Each colour stands for one multiple coupling field

An finer analysis of the coupling exchange timings is done with the targeted resolution model (ocean_ORCA025-BGC_ORCA075). The pyLucia graphical tool is used to visualise the time spent during each exchange (Fig. 7). Thanks to this tool, the optimum balancing between ocean and BGC computation can be find. It clearly appears that the time spent in mapping and interpolations, i.e. the coarsening operations performed in the ocean model (below the black vertical separation line) and plotted in colour, is a linear function of the 3D coupling fields number, since the 2D field exchanges can be neglected.

This result suggests that some improvements can be done by grouping coupling fields exchange. 3D fields are already produced by gathering vertical levels with the `CUBE` decomposition but an additional grouping can be applied to these fields (or the 2D fields) with the already available "multiple field" function. It is possible to group all the fields that share the same OASIS operation chain and source/target grids. This last constraint unfortunately limits the grouped fields to 10 but allows to significantly reduce the coupling operation duration. This is this final `namcouple` configuration that was used to produce the Fig. 7.

We assume that OASIS interpolations such as these coarsening operations, that requires more communications than computations, could have a scalability limit. But as shown in Fig. 6a, this limit is higher than the components scalability limit. For that reason, it is not valuable to try to perform the interpolations on model decomposed on less resources (the BGC model).

Despite the grouping improvement, a closer look to the exchange timing (time zoom, lower Fig. 7) reveals that an important jitter (asynchronism) between processes delays each coupling exchange. This can be evaluated by measuring the white areas included in sequence of coupling exchanges (coloured areas). At this stage of our study, there is no obvious solution to reduce this probably machine dependant jitter, which is found to be once again a major concern for the OASIS performance. However, the same effect can be observed on the BGC side, during the receiving operation. Here, an improvement was observed when the receive sequence as coded in the BGC interface matches the ocean interface sending sequence. This last modification was used to produce Fig. 7 and kept during the performance measurement shown in the next section.

It is also interesting to notice on upper Fig. 7 how irregular the coupling period can be, mainly because the irregularity of the ocean model computation or memory/disk access (white areas) or because of the OASIS coarsening operations themselves. As already observed by [21], this variability of the coupling time step duration strongly complicates the load balancing of the two concurrent component computations. Hopefully, the timeline (as the one shown in Fig 7) automatically produced by OASIS and available in the next coupler release, even if it does not allow to reach the perfect balance, strongly helps the user.

5.4- Maximum improvement

In order to minimise computing resource waste, large simulation campaign such as a CMIP exercise should only start when an optimum parallel decomposition of the used model is

found. Maximum speed, minimum cost or a trade off between the two are the usual criteria that guide the choice. In the case of OASIS based coupled models, the problem is more complex, since as many as the component number included in the coupled system must be tuned. In addition, the balance between concurrent components must be found to minimise the load imbalance.

In the case of our coarsened coupled model (ocean-ORCA025/BGC-ORCA075), we choose a decomposition that minimises the total cost (according to right Fig. 5) but also the OASIS mapping cost (according to Fig. 6a). We split the resources of 50 nodes (6400 cores) between 5097 (ocean) and 1256 (BGC) cores. The ocean (resp. BGC) grid decomposition cuts 120 (resp. 48) subdomains in length and 62 (resp. 35) in width for an average subdomain size equal to 14x19 (resp. 12x12).

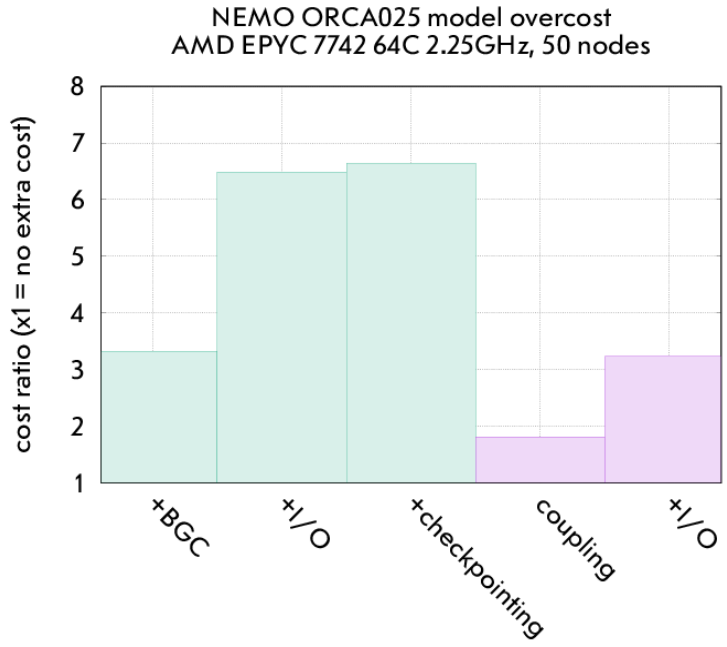

Figure 8: cost (or slowing down) ratio regarding to an ocean only ORCA025 simulations. Additional modules (BGC, Outputs, Checkpointing operation) in single executable are plotted in green. Additional modules (coupling of coarsened BGC, Output) in coarsening coupled system are plotted in pink

At the same time, we measure the performance of an ORCA025 ocean-BGC mono-executable configuration used on the same node number. An appropriate timing instrumentation allows to discriminate ocean (with ice) dynamics and physics from BGC related computations. From this information, we can deduce the number reported in the first box of Figure 8 ("+BGC" label): the extra cost (or slow down ratio) of the ORCA025 BGC computations (~x3.4) regarding to the ORCA025 ocean ones.

From the comparison between the time to solution of the mono-executable and the coarsened coupled configurations, we deduce that the coarsened BGC model slows down the ORCA025 simulation by approximately a factor of 1.8 ("coupling" label in Fig. 8). In other

words, for the same and optimum number of resources, the ORCA025 NEMO ocean model is 3.4 times more expensive (or 3.4 times slower) when BGC model is activated, but only 1.8 times more expensive if the same BGC routines are activated on a coarsened ORCA075 grid.

In production mode, a subset of the NEMO variables are output at runtime, usually via XIOS servers. We try to estimate the cost of our configurations when a comprehensive volume of output are performed by a set of 4 XIOS servers. In coupled mode, the servers can produce output for both ocean and BGC components, on their respective grids, with the appropriate set of `XML` parametrisation files. The production of such output amount further slows down the pace of our simulations. For the reasons already stated and related to the high variability of disk access response time on a production supercomputer, it is difficult to quantify this slow down. In Fig. 8, one can see that the extra cost of output increases the ORCA025 ocean computations slowing down from a factor 6.5 (mono-executable case) and 3.2 (coupled case). Differently said, the output causes the same slowing down to the ocean-BGC computations (x2), whatever the resolution of the BGC grid is. This is a good property of the coupled configuration, because one could have thought that the same extra time would have been added to the simulation time in both mono-executable or coarsened cases. Two reasons can explain this good property:

- the disk writing cost of ORCA075 data is probably smaller,
- the additional diagnostics performed on client side of the BGC model requires global communications, that are done faster on the 10 nodes required by the coarsened BGC processes than on the whole 50 nodes of the mono-executable.

For that same reasons, the relative cost of output is supposed to increase with parallelism. This could so dramatically affect the computing performance in production mode that an optimisation such as the BGC coupled coarsening configuration described here could have much less importance than a good output strategy.

We also mention in Fig. 8 the additional cost of the checkpoint/restart operations ("+checkpointing" label) for a simulation chunk length of 1 simulated year. If the cost remains marginal on 50 nodes of our supercomputer, its importance grows on higher decompositions. Again, an efficient checkpointing strategy is probably of more importance at scale than any coarsening implementation.

Even though the coarsening strategy relevance in production mode could be jeopardised by the disk access cost, and even though, because of the load imbalance, the cost of coarsening operation involving a large number of computing nodes and performed at each time step becomes as expensive as the whole BGC computations on the coarsened grid, the OASIS based coupled configuration is two time cheaper than the ORCA025 ocean-BGC model. In a context of energy saving and carbon footprint reduction, this result cannot be neglected. But the importance of the coupling cost revealed in this study emphasises the reality of the MPI parallelism limits, that one can not only notice in the exchange of NEMO variable boundaries, but that must be lowered each time than other tasks, usually neglected, have to be operated in addition to those needed for the core model equation resolving: coupling, diagnostics, resolution coarsening or resolution refinement.

6- Perspectives

The coarsening implementation should now focus on the validation of the coupling field exchanged and of their use by the BGC model. This before any attempt to quantify the impact of the BGC coarsening in the whole ocean-BGC coupled model results.

We propose a two step validation:

- a coupled simulation without any coarsening but including the ocean/BGC coupling fields in a variable volume configuration must be performed and compared to the original mono-executable simulation. Results should be approximately the same,
- this validated interface must be used to perform an ocean_ORCA025/BGC_ORCA075 simulation and to compare the original and coarsened coupling fields.

The computing cost reduction of an ESM that includes a BGC component is not a new problem that the community try to address. In a larger perspective, the energy cost or the carbon footprint of any activity is now questionable, particularly in a community that risks its credibility [22]. The original solution proposed here, based on the widely used and efficient OASIS coupler, is interesting from several points of view.

Its interface was built for a specific BGC model (TOP-PISCES) in a specific configuration (global ORCA025 with variable volume meshes) but its modularity facilitates the adaptation for other configurations or even other BGC models such as MEDUSA [23], previously used in NEMO, or even external frameworks such as the GFDL's COBALT [24].

The elevated cost of the coupling functions (the coarsened model is two time more expensive than its cost in the previous mono-executable implementation), mainly due to the huge number of MPI communications needed, can be reduced if more assumptions are made regarding to the quality of the coupling fields and their conservation. First, a significant gain in both computation and communication speed could be achieve in the coupler itself by decreasing the precision of the exchanged variables (following [25]). Second, it does not seems unthinkable to reduce the BGC model time step, considering that its horizontal resolution was also reduced. The OASIS library easily performs time averages and it would not be technically difficult to evaluate the effect of such modification on the result quality. Third, since the origin of the coupling cost is the high number of 3D scale factors exchanged, it would be interesting to carefully check their necessity and make a trade off between cost and quality of the quantities conservation. Finally, a further increase of the coarsening coefficient can be considered, even though this would mainly reduce the BGC computation cost but would let the coupling cost unchanged.

Despite its advantages, the implemented solution has also drawbacks. Despite its light intrusiveness, even though the interface is well separated from the rest of the code and does not change the maintenance cost, the coupled ocean-BGC model, as any other OASIS based

coupled system, is more difficult to handle than the single executable NEMO. Additional input files are needed and the building procedure is not fully automatic. As usual, a question must be addressed before any release : how many people will chose the coupled coarsened ocean-BGC model because of its computing cost and despite the complexity of its setting ?

Even more fundamentally, the proposed modifications increase the code complexity, thus the human cost of its maintenance. In a context of ever growing parametrisation number, and ever languishing human resources, it would be good to prioritise the modifications to be done on NEMO. At the same moment, the pole folding communication improvement, the inclusion of double halos to reduce boundary communication frequency, introduction of OpenMP parallelism or the change of all the real variables to single precision, are all planed in the same effort to optimise the NEMO behaviour on our supercomputers. Who are the people really able to maintain this effort and who will not suffer of the code growing complexity ? The computing performance is necessarily led by the research program of our community. In our case, the BGC resolution reduction is a simple idea to facilitate the resolution increase of the ocean component in the ESM models. It would be important to really evaluate the necessity of such increase, or the necessity to increase it uniformly on the whole globe, before deciding to modify the code as proposed in this document. Before any use in scientific studies, a comprehensive comparison of our results with the full resolution ocean-BGC simulations are mandatory.

Acknowledgement: Computing resources were provided by Météo-France. The authors wish to acknowledge use of the Ferret program for analysis and graphics in this report (Ferret is a product of NOAA's Pacific Marine Environmental Laboratory) and 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 programme under grant agreement No 821926 (IMMERSE)

References

[1] Etiemble, D, 2018: 45-year CPU evolution: one law and two equations. Second Workshop on Pioneering Processor Paradigms, Feb 2018, Vienne, Austria. Hal-01719766

[2] Thompson, N., & Spanuth, S. , 2018: The decline of computers as a general purpose technology: why deep learning and the end of Moore's Law are fragmenting computing. Available at SSRN 3287769

[3] Berthoud, F., Bzeznik, B., Gibelin, N., Laurens, M., Bonamy, C. et al., 2020: Estimation de l'empreinte carbone d'une heure.coeur de calcul (in french) UGA - Université Grenoble Alpes; CNRS, INP Grenoble, INRIA, hal-02549565v4f

[4] Séférian, R., Nabat, P., Michou, M., Saint Martin, D., Voldoire, A., Colin, J., ... & Sénési, S., 2019: - Evaluation of CNRM Earth System Model, CNRM ESM2 1: Role of Earth System Processes in - - Present Day and Future Climate. - Journal of Advances in Modeling Earth Systems, **11** (12), 4182-4227

[5] Madec, G. and NEMO System Team, 2019: "NEMO ocean engine", Scientific Notes of Climate Modelling Center (27) – ISSN 1288-1619, Institut Pierre-Simon Laplace (IPSL)

[6] Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., and Taylor, K. E., 2016: Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization, Geosci. Model Dev., **9**, 1937–1958, https://doi.org/10.5194/gmd-9-1937-2016

[7] NEMO TOP Working Group, 2019: "Tracer in Ocean Paradigm (TOP) — The NEMO passive tracer engine",Scientific Notes of Climate Modelling Center (28) – ISSN 1288-1619, Institut Pierre-Simon Laplace (IPSL)

[8] 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, **11**, Iss. 6, 1759-1783. DOI:10.1029/2019MS001644

[9] Bricaud, C., Le Sommer, J., Madec, G., Calone, C., Deshayes, J., Ethe, C., Chanut, J., and Levy, M., 2020: Multi-grid algorithm for passive tracer transport in the NEMO ocean circulation model: a case study with the NEMO OGCM (version 3.6), Geosci. Model Dev., **13**, 5465–5483, https://doi.org/10.5194/gmd-13-5465-2020

[10] 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

[11] 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

[12] Donahue, A. S., & Caldwell, P. M., 2020: Performance and accuracy implications of parallel split physics-dynamics coupling in the energy exascale Earth system atmosphere model. Journal of Advances in Modeling Earth Systems, 12, e2020MS002080. https://doi.org/10.1029/2020MS002080

[13] Maisonnave, E., 2020: Ocean/Biogeochemistry macro-task parallelism in NEMO, Working Note, **WN/CMGC/20/31**, CECI, UMR CERFACS/CNRS No5318, 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] Irrmann, G., Masson, S., Maisonnave, E., Guibert, D., & Raffin, E., 2021: Improve Ocean Modelling Software NEMO 4.0 benchmarking and communication efficiency , *Geosci. Model Dev. Discuss. [preprint]*, doi:10.5194/gmd-2021-320, in review
[16] Piacentini, A., Maisonnave, E., Jonville, G., Coquart, L. and Valcke, S., 2018: A parallel SCRIP interpolation library for OASIS , Working Note, **WN/CMGC/18/34**, CECI, UMR CERFACS/CNRS No5318, France
[17] 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
[18] Piacentini, A., & Maisonnave, E., 2020: Interactive visualisation of OASIS coupled models load imbalance , Technical Report, **TR/CMGC/20/177**, CECI, UMR CERFACS/CNRS No5318, France
[19] Larson, J., Jacob, R. & Ong, E., 2005: The Model Coupling Toolkit: A New Fortran90 Toolkit for Building Multiphysics Parallel Coupled Models, Int. J. High Perf. Comp. App., **19(3)**, 277-292
[20] Balaji, V., Maisonnave, E., Zadeh, N., Lawrence, B. N., Biercamp, J., Fladrich, U., Aloisio, G., Benson, R., Cau9el, A., Durachta, J., Foujols, M.-A., Lister, G., Mocavero, S., Underwood, S., & Wright, G., 2017: CPMIP: measurements of real computational performance of Earth system models in CMIP6, Geosci. Model Dev., 10, 19–34, https://doi.org/10.5194/gmd-10-19-2017
[21] Tintó Prims, O., Castrillo, M., Acosta, M., Mula-Valls, O., Sanchez Lorente, A., Serradell, K., Cortés, A. & DoblasReyes, F., 2018: Finding, analysing and solving MPI communication bottlenecks in Earth System models, J. of Comp. Science, https://doi.org/10.1016/j.jocs.2018.04.015
[22] Attari, S.Z., Krantz, D.H. & Weber, E.U., 2016: Statements about climate researchers' carbon footprints affect their credibility and the impact of their advice. *Climatic Change* **138,** 325–338 https://doi.org/10.1007/s10584-016-1713-2
[23] Yool, A., Popova, E. E., and Anderson, T. R.: MEDUSA-2.0: an intermediate complexity biogeochemical model of the marine carbon cycle for climate change and ocean acidification studies, Geosci. Model Dev., 6, 1767-1811, doi:10.5194/gmd-6-1767-2013, 2013.
[24] Stock, C. A., Dunne, J. P., Fan, S., Ginoux, P., John, J., Krasting, J. P., et al., 2020: Ocean biogeochemistry in GFDL's Earth System Model 4.1 and its response to increasing atmospheric CO2. Journal of Advances in Modeling Earth Systems, 12, e2019MS002043. https://doi.org/10.1029/2019MS002043
[25] Váňa, F., Düben, P., Lang, S., Palmer, T., Leutbecher, M., Salmond, D. and Carver, G., 2017: Single precision in weather forecasting models: An evaluation with the IFS. Monthly Weather Review, 145(2), pp.495-502.