Ocean/Biogeochemistry macro-task parallelism in NEMO

E. Maisonnave

WN/CMGC/20/31

Abstract

The offline biogeochemistry (BGC) model TOP-PISCES is coupled to the stand alone NEMO ocean/sea-ice. The individual executables are exchanging coupling variables via OASIS. This configuration produces the same results than the standard online ocean/sea-ice/BGC single executable, even if bit to bit reproducibility is not ensured and changes in some namelist parameters can jeopardize the result. We find no significant bias to the concurrent (instead of sequential) performing of ocean and BGC calculations. The computing performance can be enhanced in coupled mode, but an estimation of the extra cost induced by the exchange between the two components of several 3D variables at each model time step is relatively big (around 20%). The coarsening of the BGC component remains the most efficient solution for a significant performance gain. This OASIS based coupled system can pave the way for a modular and perennial implementation of this coarsening

Table of Contents

Representation of always more phenomena, at always smaller scales in space and time requires to optimally make the most of the available computing power. In this trend, the inclusion of ocean biogeochemistry (BGC), for carbon cycle modelling in ESMs [1], is particularly demanding. As put in evidence in [2], the chemical species reactions and advection can multiply the NEMO ocean model [3] cost by more than a factor 3. On our most powerful supercomputers, this forbids to switch on this BGC module, called TOP-PISCES [4] at global spatial resolution higher than 1°. 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.

We follow the strategy already implemented for NEMO ocean and sea ice submodels [5]: 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 [6]. This implementation, recently updated in NEMO 4.0 at Met Office [7], showed interesting computing performance. The separation of the two submodels in two executables allowed to perform their computations concurrently, increasing what we call the *macro-task parallelism*. The modularity of this configuration is supposed to facilitate the last step of the solution we are proposing here: a multi-grid coupled system.

The splitting of ocean/BGC into distinct modules is described and tested in this document. We first documents (i) an overview of the starting ocean/BGC online configuration (reference), (ii) the study of the existing BGC offline module results vs online ones, (iii) the coupling strategy of this module with an ocean only executable, (iv) the comparison of the new coupled model results with those of the reference configuration and (v) an estimate of the new computing performance. We concludes with assumptions concerning the best way to reduce BGC resolution and how to perform the necessary transformation between the two model grids (coarsening).

Reference configurations

A two steps strategy is proposed to validate the results of the ocean-BGC coupled model:

- As described in Figure 1 (“intermediate configuration”), we force a BGC offline model with ocean variables produced during a previous online ocean/BGC simulation, performed with the same resolution/namelist parameters/input files. If BGC output of the two simulations are identical, it means that the dynamical forcing conditions are the same for the BGC offline (forcing by files) and online models (forcing by arrays).
- In this case, we can replace the forcing routines of the BGC offline model by a coupling interface and receive, from a separated ocean only executable running at the same time, the dynamical forcing conditions needed. Again, the results should be identical to the online reference ones.

Note: To simplify the naming in this document, the SI3 sea-ice module is included in the submodel called "ocean" or "ocean only".

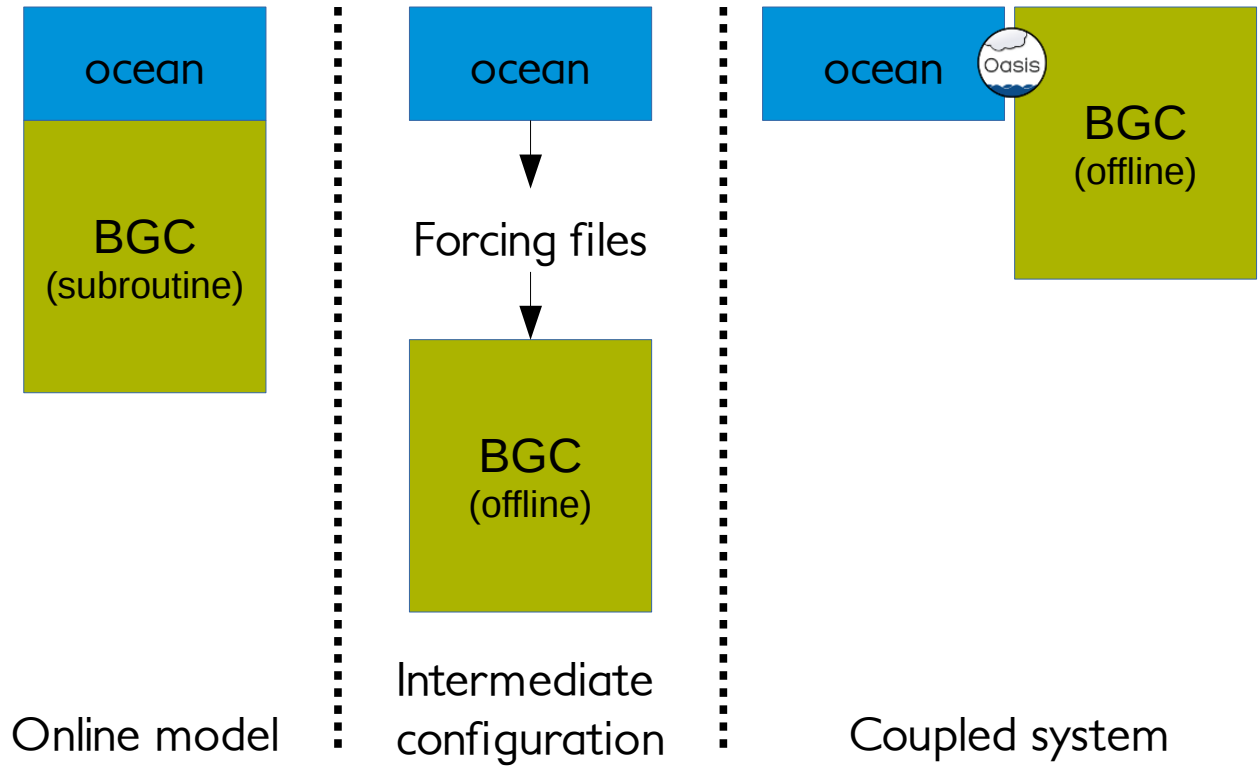

Figure 1: Description of the 3 configurations compared in this study: online reference (BGC called as a subroutine), intermediate (offline BGC forced by the results of a previous ocean only simulation) and coupled (offline BGC coupled with ocean only model)

Online BGC

We start from the NEMO sources released in October 2019 (4.0.1) and the associated input files and namelists. For that reason, and to facilitate the recursive development/test procedure, the global ORCA2 resolution is selected as a test case.

We use this NEMO model, including SI3 and TOP-PISCES modules, to produce a reference one year long simulation. The purpose is double: to save ocean and BGC reference quantities and produce ocean forcing fields that will be used in a following BGC offline simulation.

For reasons explained in the "Offline BGC" paragraph, this reference configuration was slightly changed, but only by namelist, during the implementation procedure. We have to mention here that the time step length is reduced to 3600s (instead of 5400s) for a practical reason. We need to save in files a set of model variables that will be used to force the BGC offline model. Since it is not possible to output variables at non multiple frequency of 1h, the time step length is set to 1h. For the same reasons, it was necessary to call the `sbc` surface module at each model time step, to output every 1h some variables from this `sbc` subroutine.

Table 1 lists these 2D and 3D variables, their name and the routine from where the subroutine

iom_put is called to save them on files (one file per grid type). Notice that the output is not necessarily done immediately before the BGC sub routine call, (sbc: before dynamics, dia_wri: after BGC) which can introduce a small difference between the forcing fields and the fields as they are seen by the BGC online code.

Field name	Variable name	from routine	Grid
Temperature	tsn(jp_tem) (3D)	dia_wri	T
Salinity	tsn(jp_sal) (3D)	dia_wri	T
Mixed layer depth	hmld	zdf_mxl	T
Water balance	emp-rnf	sbc	T
Salt flux	fmmflx	sbc	T
Sea ice cover	fr_i	sbc	T
Solar flux	qsr	sbc	T
Wind speed module	wndm	dia_wri	T
Effective zonal transport	zun(3D)	tra_adv	U
bbl diffusive flux - i	ahu_bbl	tra_bbl	U
Effective meridional transport	zvn (3D)	tra_adv	V
bbl diffusive flux - j	ahv_bbl	tra_bbl	V
Effective vertical transport	zwn (3D)	tra_adv	W
Vert. eddy diff. coef. for T	avt (3D)	dia_wri	W
Runoff	rnf	sbc_rnf	T
Water balance (ts before)	emp_b-rnf	sbc	T
Horizontal divergence	hdivn (3D)	dia_wri	T
Vert. eddy diff. coef. for S	avs (3D)	dia_wri	W

Table 1: List of fields exchanged from ocean to biogeochemistry components, name and grid of corresponding NEMO variable and output subroutine (blue : reference, orange : additional)

Offline BGC

A one year long simulation of the BGC offline model is performed. The model is built with the same routines than the online ocean-sea-ice-BGC, but also includes the routines stored in the src/OFF directory. The offline nemogcm program avoids calling ocean dynamics, physics and surface processes. This part is replaced by the reading (dtadyn) of ocean variables produced by the previously described simulation.

To ensure the comparativeness of the results between BGC online and offline modules, we use the very same subroutines (except src/OFF nemogcm and dtadyn), the same starting conditions and the same namelists in both experiments.

The comparison of BGC model integrated quantities, e.g. the total Chlorophyll concentration, shows significant differences. For that reason, we decide to change some namelist parameters,

so that result reproducibility between online and offline models can be found.

The main modifications are:

- The enabling of the linear free surface (`ln_linssh = .true.`)
- No runoffs, nor nutrients inputs in BGC
- No Sea Surface restoring on T and/or S
- No representation of icebergs
- No double diffusion
- Same time step for biology than for the ocean (`nrdttrc = 1`)

These new conditions lead to acceptable differences. We show values for Chlorophyll concentration in Fig 2. The monthly anomalies after a 12 month long simulation can locally reach 2%. More variables are provided in appendix and exhibits the same magnitude of errors.

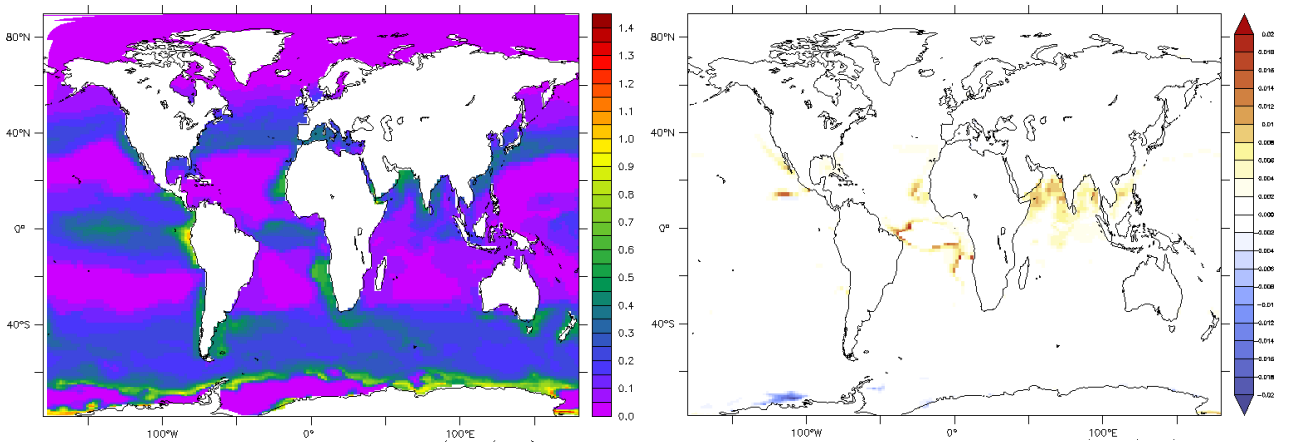

Figure 2: Raw concentration (left) and anomaly of offline model vs reference (right) of Chlorophyll concentration (mg/m3) vertically integrated and averaged during the last 12th month of the simulations

We perform a complementary set of simulations without the PISCES biogeochemistry module. In this configuration, the TOP model only advects the water age tracer. The error order of magnitude is similar to the previously presented Chlorophyll concentration anomaly.

As already mentioned, the time step shift affecting some of the BGC forcing quantities (vs the quantities directly transmitted to the online BGC module) could be the origin of the small mismatch between the two simulations. An additional source of error can be attributed to the special treatment of slopes (`l_ldfslp = .true.`), in link with the reference lateral diffusion scheme we are using. Since our configuration includes the linear free surface computations, it is not possible to switch off these slope computations, that are probably done in a different way by the offline BGC model.

Of course, a much more comprehensive validation of the similarity of the online/offline simulations is needed (and this validation should include the investigation of the `l_ldfslp` option mismatch) to ensure that a coupling of the offline model with the ocean only executable can reproduce the same results than the online configuration ones *for any values* of the ocean/BGC namelists. This important work (necessarily a team work) is postponed until we can prove the validity and the computational efficiency of the coupled system including

coarsened BGC.

Coupling

In a second step, we propose to interface the BGC offline model presented above with the OASIS library and couple it with another executable: an online-like NEMO configuration, excluding the TOP-PISCES routines (`key_top` CPP key disabled). The simulation is supposed to again reproduce exactly the same results than the online reference presented above, assuming that the OASIS library exactly reproduces the online (resp. offline) write (resp. read) mechanisms of coupling quantities listed in Table 1. But this time, OASIS communicates directly at runtime these quantities to the BGC model via MPI messages and without any intermediate files. On the other way round, the BGC model transmits the light absorption coefficient to the ocean (`etot3`), but this quantity is not used (`ln_qsr_bio = .false.`, bio-model light penetration disabled) in order to simplify our problem, while keeping the two-way coupling communication pattern supposed to be used in ESM production simulations.

Implementation

The existing OASIS interface, designed for atmosphere-ocean coupling at surface, has to be adapted to the BGC coupling, which is three-dimensional. In that purpose, the `cpl_oasis` subroutine, in charge of the OASIS library calls, has to be enriched and a new routine `dyncpl`, symmetric to the `sbccpl` surface interface, is developed to ensure the filling of the model variables with the incoming quantities and the filling of the outgoing quantities. Two different `dyncpl` routines are necessary: one for the ocean model and one for the BGC model.

In `cpl_oasis`, we proceed to a re-organisation of the OASIS coupling definition phase (former `cpl_define`). We split this subroutine into three pieces:

- the definition of domain decomposition, called only once for atmosphere (in future ESM configuration) and BGC coupling, since the domain decomposition is the same in both cases (`cpl_define_partition`)
- the field definition is split into two subroutines, `cpl_define_sbc_fields` and `cpl_define_bgc_fields`, called separately in `sbccpl` and `dyncpl`
- a single call the the termination phase of the OASIS coupling definition (`cpl_define_end`)

It is also mandatory to create two separate set of OASIS exchange commands (`cpl_sbc_snd` and `cpl_sbc_rcv` to send/receive coupling field from atmosphere and `cpl_bgc_snd` and `cpl_bgc_rcv` to send/receive coupling field from BGC). Notice that the exchange of 3D quantities with BGC model is performed by taking benefit of the new "bundle 2D" functionality: a 3D NEMO variable is directly provided as argument to the

`oasis_put/oasis_get` API. OASIS is able to process and communicate every horizontal slices (31 levels in our ORCA2 configuration) in the same way.

In the BGC code, input/output ASCII files are renamed to avoid confusion with the files needed/produced by the ocean model:

- input: `namelist_cfg` and `namelist_ref`
- output: `communication_report.txt`, `time.step`, `ocean.output`, `output.namelist.dyn`

This renaming is active only if the BGC offline model is coupled. The NetCDF/XIOS input/output files are unchanged. However, the declaration of a new context must be activated in `iodef.xml`:

```
<context id="bgc" src="./context_bgc.xml"/>
```

and two coupled components must also be declared:

```
<variable id="oasis_codes_id" type="string" >oceanx,bgc</variable>
```

The calling sequence of BGC coupling related subroutines can be summarised in the following table:

Ocean	Biogeochemistry
nemogcm	nemogcm
nemo_init	nemo_init
cpl_init	cpl_init
cpl_define_partition	**cpl_define_partition**
dyn_cpl_init	**dyn_cpl_init**
cpl_define_bgc_fields	**cpl_define_bgc_fields**
cpl_define_end	**cpl_define_end**
DO WHILE	DO WHILE
…	**dyn_cpl_rcv**
DYNAMIC	**cpl_bgc_rcv**
dyn_cpl_snd	PASSIVE TRACERS
cpl_bgc_snd	**dyn_cpl_snd**
dyn_cpl_rcv	**cpl_bgc_snd**
cpl_bgc_rcv	ENDDO
ACTIVE TRACERS	
ENDDO	
cpl_finalize	cpl_finalize

Table 2: Calling sequence of BGC coupling related subroutine, in both ocean and BGC models. In bold, newly created subroutines

The coupled configuration needs to be parametrised through the OASIS `namcouple` ASCII file. No interpolation is performed between the two models but an instantaneous (no time average) transformation must be prescribed to define the 14 ocean to BGC coupling fields and the BGC to ocean one. A simple addition in this namcouple of the LAG parameter, set to the model time step (3600s), is enough to switch from sequential to concurrent mode:

- in sequential mode, BGC and ocean models are waiting each other the coupling fields to start their computations
- in concurrent mode, the BGC input coupling fields are coming from the previous ocean time step and are already available in a file at restart, and the ocean incoming field is needed only at the end of the ocean time step.

Validation

Both sequential and concurrent coupling are studied in this paragraph. The sequential coupling is supposed to perfectly reproduce the online model behaviour (besides open questions addressed in the "Offline BGC" paragraph).

In sequential mode, a comparison of BGC variables, similar to the offline/online comparison, shows the same order of magnitude differences. As an example, Figure 3 again details the final chlorophyll concentration diagnostic (monthly average). The difference pattern is similar to the offline-online difference, which suggests that biases introduced by the coupling are mainly introduced by different behaviour of the BGC model in online and offline modes. Compared evolution of other BGC integrated quantities is available in appendix.

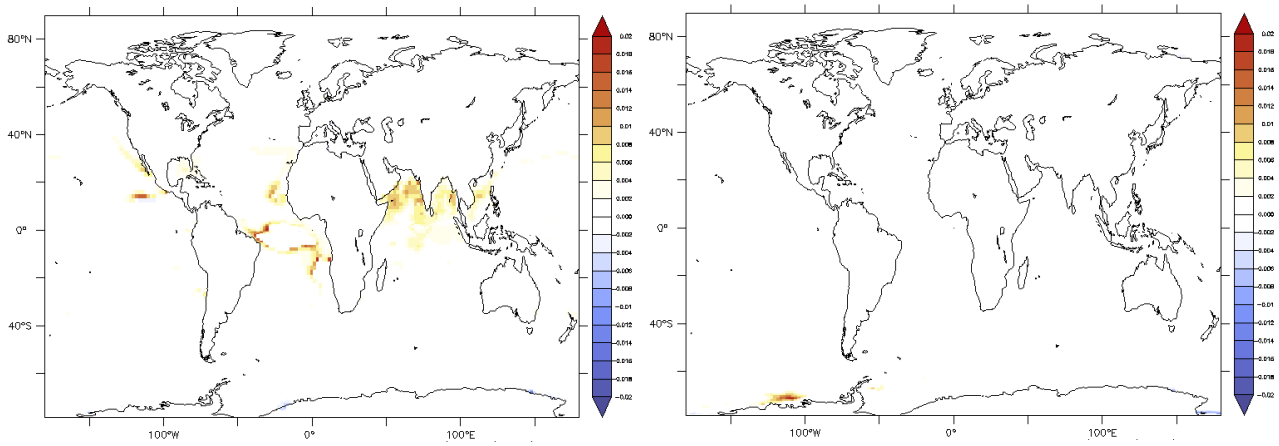

Figure 3: Anomaly of sequentially coupled model vs online (left) and offline (right) models, of Chlorophyll concentration (mg/m3) vertically integrated and averaged during the last 12th month of the simulations

An evaluation of the concurrent coupling option impact is proposed in Figure 4. Differences are several orders of magnitude below online configuration differences (confirmed on other quantities as shown in appendix). This suggests that the concurrent mode, supposed to be more computationally efficient, can be substituted to the sequential mode in the next steps of our study.

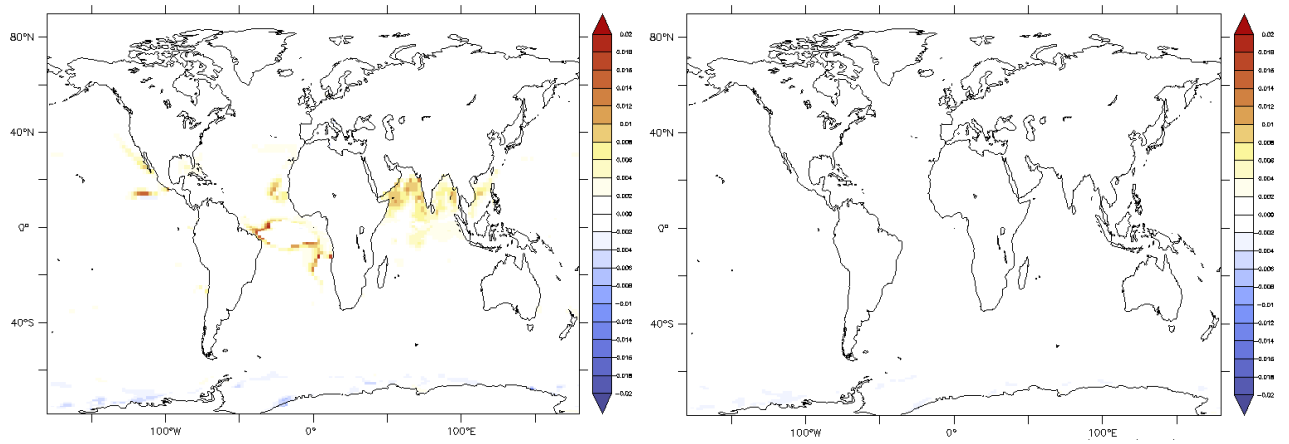

Figure 4: Anomaly of concurrent coupled model vs online (left) and sequential coupled model (right) of Chlorophyll concentration (mg/m3) vertically integrated and averaged during the last 12th month of the simulations

There is no difference in ocean variables (bit to bit reproducibility), which confirms that the BGC model variables do not affect, in any cases, the ocean computations.

Computing performance

The coupled model used to take the measurements presented in this paragraph is set in concurrent mode, to maximise its performance.

An initial attempt to measure the computational performance of the coupled model leads to non reproducible results, in that sense that the total restitution time of two similar simulations can significantly differ.

This result motivates the removal of any model output. As already seen in [2], disk access in the Météo-France production supercomputer (`beaufix`[1]) strongly perturbs timing measurements. The best solution would be to set up an idealised configuration like BENCH (square basin, no input, no output) which reproduces a realistic coupled ocean-BGC pattern of computations. Due to limitations in time and man power, we prefer (i) to remove all unnecessary output (`key_iom` removal), including ASCII files output, (ii) perform an ensemble of simulation to provide an averaged value and (iii) remove outliers from the ensemble, by controlling the total time needed to read input files (mainly located in the `sbc` routine).

In the same perspective,

- measurements are only performed between time steps `nit000+3` et `nitend-3`, to remove any perturbation in link with initialisation or restart writing
- the printing of every time step duration, in the `ocean.output` file, at every time step, is stored in array and its writing postponed to the last time step
- for concurrent mode, the OASIS restart writing, a routine well known for its time

1 https://www.top500.org/system/178075

consumption, must be reduced to one time step slice : the INSTANT transformation must be preferred to the (equivalent in this case) AVERAGE option.

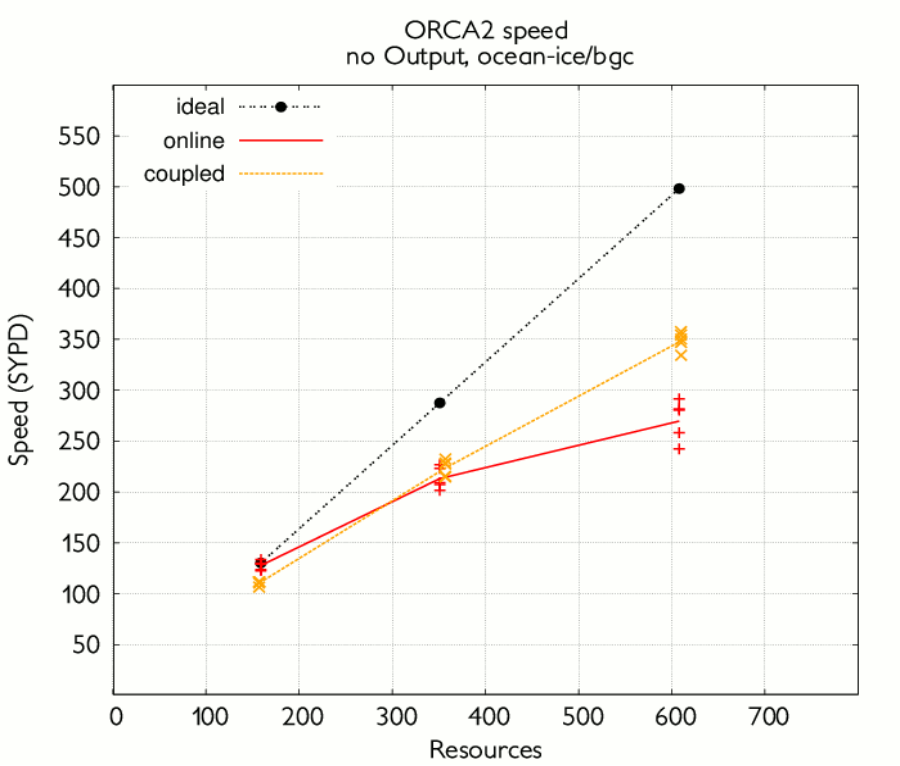

Figure 5: Speed of ORCA2 ocean-ice-BGC reference and coupled configurations, on Intel Broadwell, Météo-France. Comparison of online reference on 159,351 & 608 cores, and ocean+bio coupled configuration on 38+119, 119+238 & 238+372 cores. Member measurements (dots) and ensemble average (line). Perfect scaling in black, referenced from 159 cores

A comprehensive comparison of computing performance requires several measurements, with varying number of resources, since performance is a function of the MPI decomposition. The decomposition choice is constrained. In particular, it is necessary,

- to minimise the subdomain perimeter (to limit communications and halo extra memory requirements) and equalise the subdomain areas (to avoid load imbalance between MPI processes). This leads to magic numbers, available in the `ocean.output` file, since version 4, during the initialisation phase
- to remove the land only processes (information also provided by NEMO at runtime)
- to fully occupy all resources of the computing node (40 on `beaufix`). This constraint could be softened on the last allocated node

and in addition, in coupled mode,

- to prevent mixing of ocean and BGC processes on the same node. This requirement is checked thanks to the `HIPPO_Query_print_affinity` function [8], introduced in our OASIS library. It is fulfilled by association of every process to a given node (`-machinefile` option of `mpirun` command) and by the setting, for our Intel library,

of the environment variable:

```
I_MPI_PIN_PROCESSOR_LIST = allcores,grain=core,shift=1:map=bunch
```

- to choose decompositions close to the online model decomposition choice
- to minimise the load imbalance between ocean and BGC components

This last constraint, combined with the full occupation of node resources, is difficult to satisfy when the allocated nodes are few and leads to higher load imbalance (first point of Figure 6 dashed line) and lower performance (first point of the orange line in Figure 5).

With higher number of resources, the coupled model goes faster than the reference online counterpart. In particular, the measurement made with the largest number of resources shows about 25% of improvement. This can be explained by our ORCA2 model scalability limit, probably lower than the value of 608, allocated to the online model, but bigger than the respective values 238 and 372, allocated to each component in coupled mode.

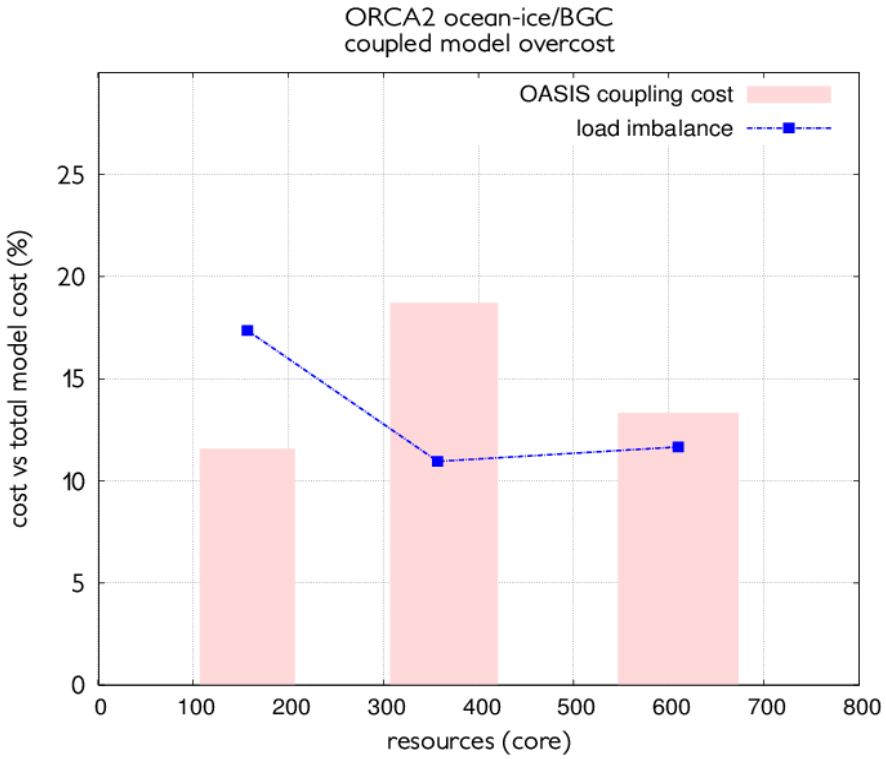

Figure 6: OASIS coupling cost (%) defined as the ratio of (i) the difference of the restitution time of slowest coupled component between a 1 hour coupling frequency simulation (ensemble average) and a 50 days coupling frequency coupled simulation (ensemble average) with (ii) the restitution time of the 1 hour coupling frequency coupled simulation. In blue line, the load imbalance between components, defined as the restitution time difference between the fastest and the slowest components, in the 50 days coupling frequency coupled simulation. Total simulated days : 50.

This good result must be mitigated. With higher resolution configuration, the scalability limit will be harder to reach, due to the limited amount of resources usually available on production

computers. Consequently, with these models, the favourable zone of coupled model performance supremacy will be out of reach or unfavourable with respect to actual speed (including wait for scheduling).

The high volume of MPI exchanges required by the coupling (five to seven 3D variables, 31 vertical levels) at every model time step, justifies a more accurate analysis of the time spent in coupling routines in general, and in OASIS library in particular.

The measurement of the coupling cost for a coupled system, also called load imbalance, was defined in [9] as the normalized difference between the time-processor integral for the whole model vs. the sum of individual concurrent components. This measurement requires the evaluation of the computing time of each individual components. The OASIS library provides a specific tool to do so, LUCIA [10]. Unfortunately, the huge number of coupling related communications and the printing of the corresponding timing on files affect the performance. In a future work, the current LUCIA behaviour would have to be enhanced with a silent measurement strategy: all timings will be stored in arrays, and processed/printed at runtime end only.

For the moment, we prefer to adopt a slightly different strategy to evaluate the coupled component load imbalance. We increase the coupling period from one time step to the total simulation duration, and keep measuring individual component timings between `nit000+3` et `nitend-3`. This removes any cost in link with coupling exchanges and only measures the difference of times needed by the two components to perform their calculations, i.e. load imbalance. A similar strategy, formerly lead in [11], also allows, by comparison of the two coupled simulations (with coupling period of one time step in one hand and the whole simulation duration in the other hand) to evaluate what we could call the "OASIS coupling cost"[2]. During this former attempt, we estimated that the OASIS coupling cost could represent 6% of the total COSMO regional ESM model cost.

Figure 6 shows the two quantities (load imbalance and OASIS coupling cost) for the three previously chosen MPI decompositions. The load imbalance values are high, particularly for the first point, as explained previously by the "full node" constraint, but also for the two other measurements. This probably derives from the other result: the high cost of OASIS coupling. This cost practically prevents to find ideally balanced configurations. The cost of OASIS coupling cannot be lowered below 10%. An additional study was necessary to understand from where this important cost was coming from.

The internal OASIS timing counters are activated (`LOGPRT namcouple` option) but the existing instrumentation is not precise enough to identify the slow down origin. A larger code instrumentation again leads to measurement perturbation (catastrophic total restitution increase) but suggests that the observed OASIS coupling cost finds its origin in several parts of the code (array copies) rather than on a single bottleneck. However, a finer evaluation of the issue is ongoing and it seems that there is room or further improvements.

2 Even though the extra cost caused by the coupling exchanges not only comes from operations performed by OASIS routines (like MPI communications or arrays copies) but also from model MPI process synchronisation during coupling exchanges

We can temporarily conclude that

- the ocean-BGC coupled model can exhibit better computing performance, in particular if the subdomain decomposition leads to computations/communications ratio that put the performance just below the scalability limit
- the coupling cost, caused by OASIS coupling extra cost and load imbalance between components is non negligible (around 20% in our case) but can be reduced

This contrasted result suggests that the only clear performance gain can only be ensured with the radical cost lowering of the most time consuming component, the BGC model.

Coarsening

This cost lowering was previously reached by decreasing the BGC component resolution in a previous version of the NEMO code [12]. In this attempt, the online model was modified to perform the computations of the BGC subroutines in smaller arrays. The transformation of ocean fine resolution arrays to BGC coarsened ones was done by a comprehensive set of additional routines that conserve the exchanged quantities (coarsening).

The advantage of this strategy is the compactness of the solution (single executable, no external coupling library needed, no component load balancing procedure) but it also has drawbacks:

- both ocean and BGC grids must be subject to the same MPI decomposition,
- a non negligible amount of extra routines must be maintained,
- the modularity is limited: the coupling with another BGC model (e. g. MEDUSA) requires additional coding, and the inclusion of AGRIF zoom seems error prone

In addition, the efficient coding of transformations from grid to grid is not trivial and the actual implementation suffers from several limitations (limited choices of decomposition and coarsening factor).

This is why we propose to follow the path drawn by this study and take benefit of the new modularity provided by OASIS between the ocean and BGC components to suggest the implementation of a multi-grid solution.

The definition of a BGC offline model operated at a lower resolution than the ocean component is trivial. One would only take care to set the resolution in accordance with the appropriate coarsening ratio (usually 3, but the 5 value can also be explored). The main difficulty relies in coarsening operation coding and the correct use of the coarsened values by BGC.

The necessary coarsening operations were recently summarised in [13]. We do not see many difficulties to let OASIS performing the coarsening operations and provide the coupled

quantities on BGC model grid. The average of the 3x3 source grid points pad in one target grid point, taking care of masked grid points, can be performed by a standard `MAPPING` operation, after a special but simple interpolation weight & address computation. The difference of land-sea mask by vertical levels can slow down the operation but is not an insuperable issue. Neither the logarithmic average, suggested for vertical mixing coarsening. The main difficulty relies in the construction of coarsened vertical factors. For vertical gradient operators, in order to preserve the ocean grid thickness, we must communicate to the BGC model the maximum of the vertical dimensions of the 3x3 source grid points pad. This operation is not permitted by any current OASIS transformation. A possible but time consuming modification of the MCT operators is required, particularly the `m_MatAttrVectMul.F90` multiplication operation:

```
AV%rAttr(m,i) = yAV%rAttr(m,i) + wgt * xAV%rAttr(m,col)
```

The second difficulty is the development of additional procedures in BGC to be able to combine the coarsened variables. These operations are already implemented in the former suffixed `crs_` NEMO routines but need to be extended to take into account more NEMO physics/dynamics configurations.

Conclusion

The offline biogeochemistry (BGC) model TOP-PISCES is coupled to the stand alone NEMO ocean/sea-ice. The individual executables are exchanging coupling variables via OASIS. This configuration produces the same results than the standard online ocean/sea-ice/BGC single executable, even if bit to bit reproducibility is not ensured and changes in some namelist parameters can jeopardize the result. We find no significant bias to the concurrent (instead of sequential) performing of ocean and BGC calculations. The computing performance can be enhanced in coupled mode, but an estimation of the extra cost induced by the exchange between the two components of several 3D variables at each model time step is relatively big (around 20%). The coarsening of the BGC component remains the most efficient solution for a significant performance gain. This OASIS based coupled system can pave the way for a modular and perennial implementation of this coarsening.

Acknowledgement: Many thanks to all NEMO community members involved in the preparatory tasks in link with this study. At LOCEAN: Olivier Aumont, Julie Deshayes, Christian Ethé, Claire Lévy, Gurvan Madec, Sébastien Masson, Renaud Person & Clément Rousset. At Météo-France, Sarah Berthet. At Mercator-Ocean, Romain Bourdallé-Badie & Clément Bricaud. At LEGOS: Rachid Benshila. At MetOffice: Mike Bell, Richard Hill, Colin Jones & Marc Stringer. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 821926 (IMMERSE)

Reference

[1] 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*

[2] 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

[3] 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)

[4] 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)

[5] 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

[6] 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

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

[8] Piacentini, A., 2020: Rationales for the implementation choices in HIPPO, personal communication, to be published

[9] Balaji, V., Maisonnave, E., Zadeh, N., Lawrence, B. N., Biercamp, J., Fladrich, U., Aloisio, G., Benson, R., Caubel, A., Durachta, J., Foujols, M.-A., Lister, G., Mocavero, S., Underwood, S., and Wright, G., 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

[10] Maisonnave, E. & Caubel, A., 2014: LUCIA, load balancing tool for OASIS coupled systems, Technical Report, TR/CMGC/14/63, SUC/CERFACS, URA CERFACS/CNRS No1875, France

[11] Maisonnave, E., Hill, R., Jamil, O. & Valcke, S., 2013: OASIS Dedicated User Support 2012, Annual report ,Technical Report, **TR/CMGC/13/18**, SUC au CERFACS, URA CERFACS/CNRS No1875, France

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

[13] Bricaud, C., Le Sommer, J., Gurvan, M., Calone, C., Deshayes, J., Ethe, C., Chanut, J., and Levy, M., 2020: Multi-grid algorithm for passive tracer transport in NEMO ocean circulation model: a case study with NEMO OGCM (version 3.6), *Geosci. Model Dev. Discuss.*, https://doi.org/10.5194/gmd-2019-341, in review

Appendix

Daily evolution of globally averaged 3D ocean variables of BGC model, raw value (upper plot) and anomalies (lower plot)

Left column: online (black), offline (red) and sequentially coupled (green). Anomalies vs online

Right column: online (black), concurrently coupled (red) and sequentially coupled (green)

Anomalies vs online, except blue: sequentially coupled vs concurrently coupled

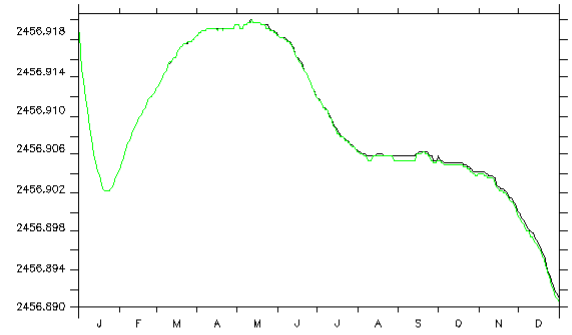

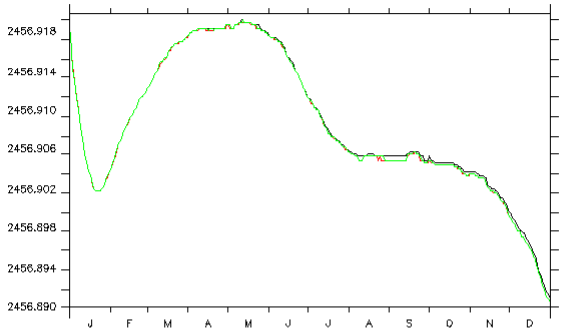

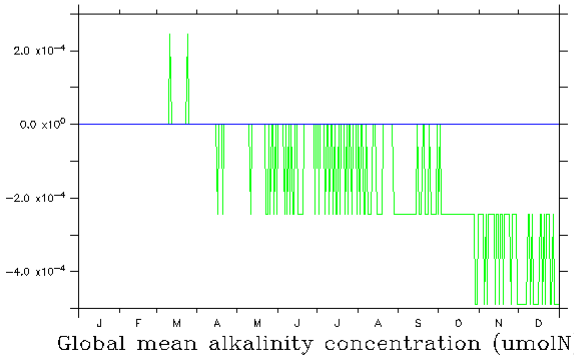

Global mean alkalinity concentration (umolN)

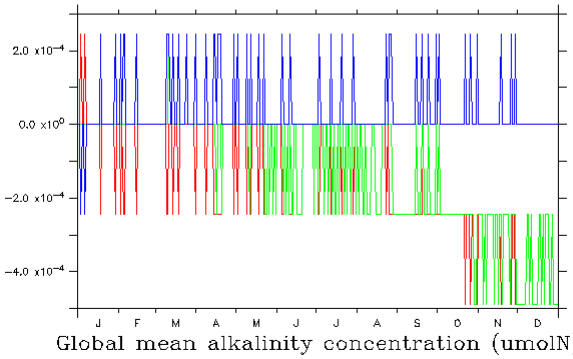

Global mean alkalinity concentration (umolN)

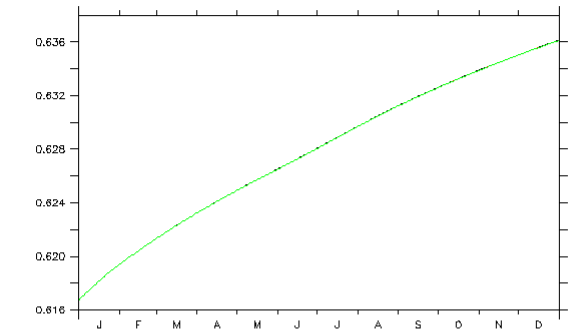

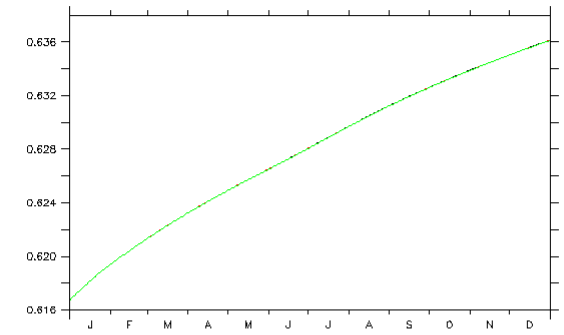

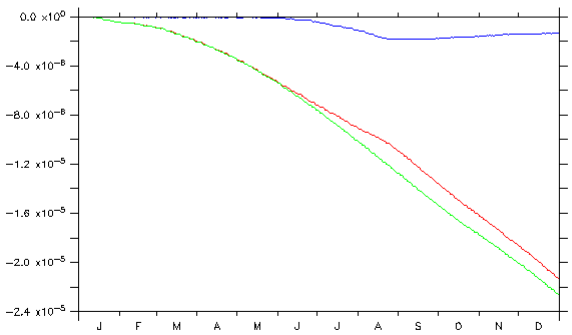

Global mean iron concentration (umolN)

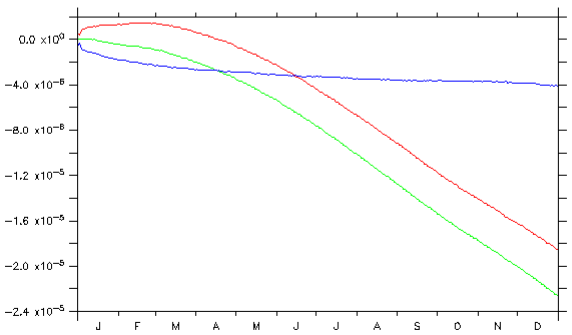

Global mean iron concentration (umolN)

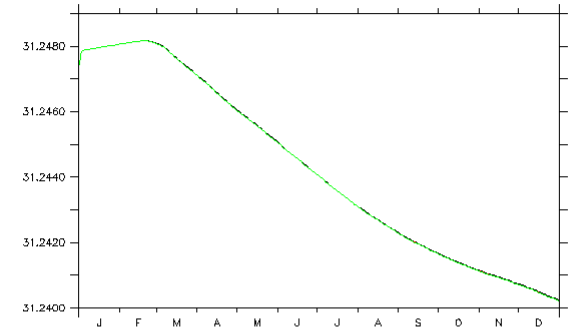

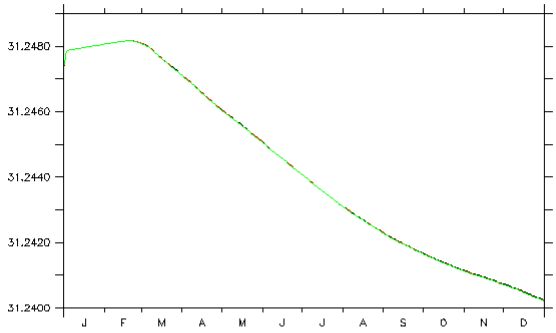

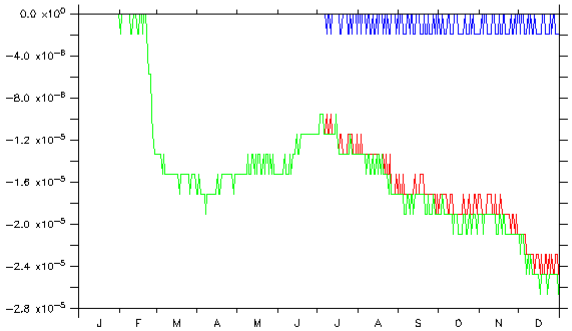

Global mean nitrate concentration (umolN)

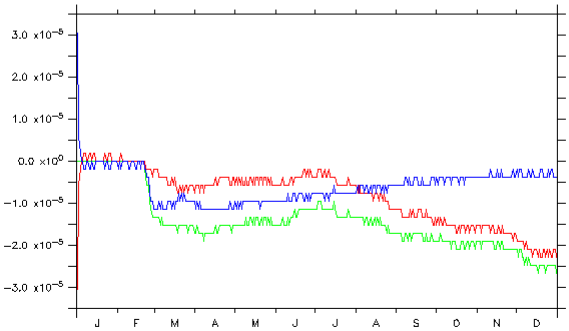

Global mean nitrate concentration (umolN)

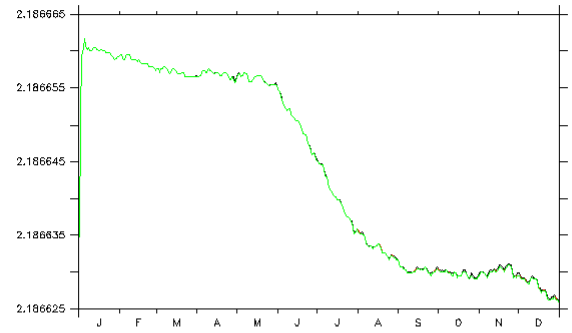

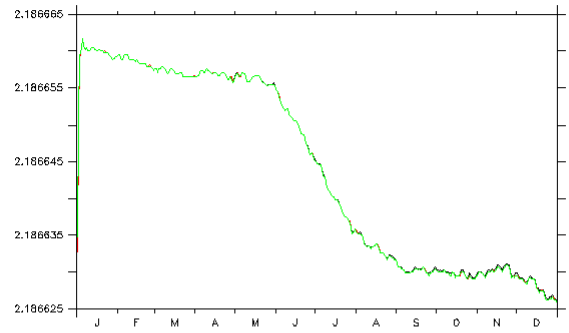

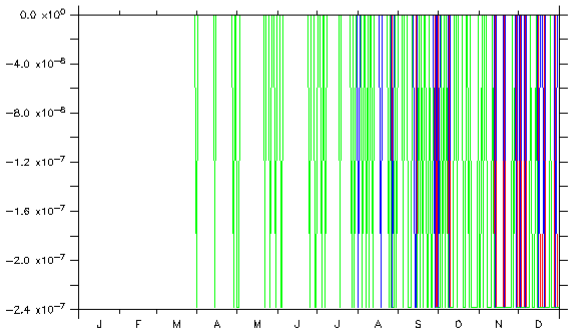

Global mean phosphorus concentration (umolN)

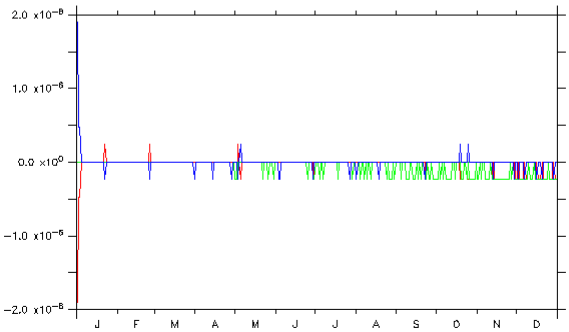

Global mean phosphorus concentration (umolN)

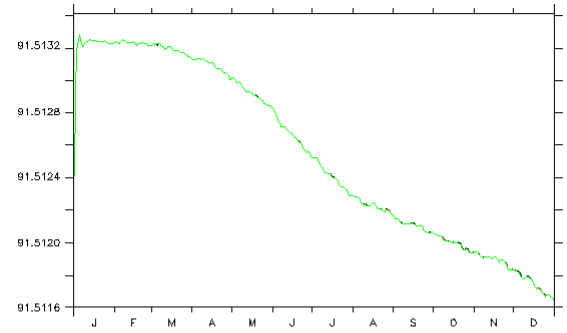

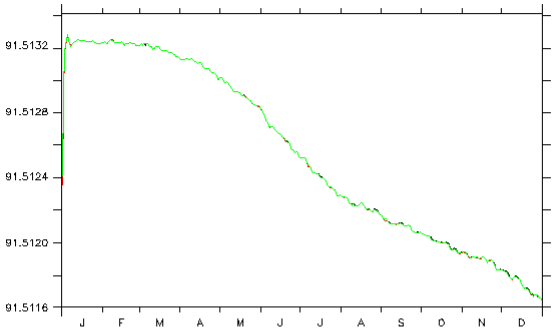

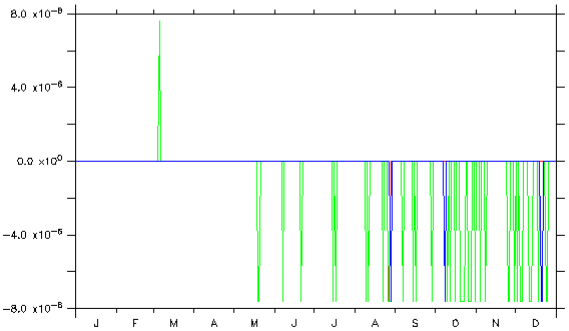

Global mean silicate concentration (umolN)

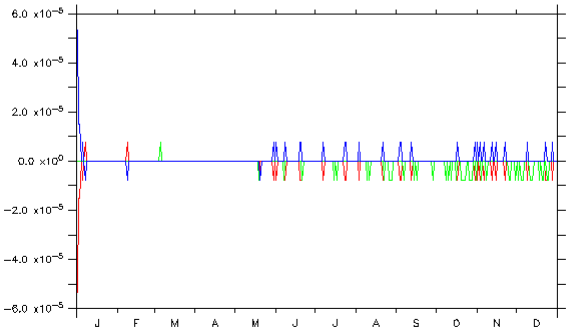

Global mean silicate concentration (umolN)

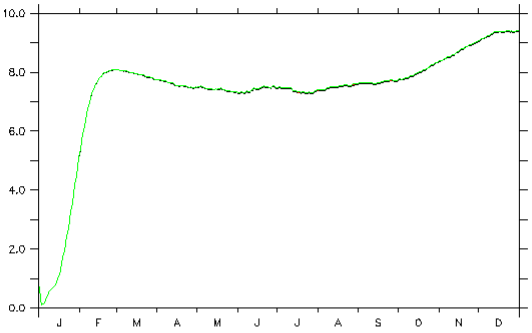

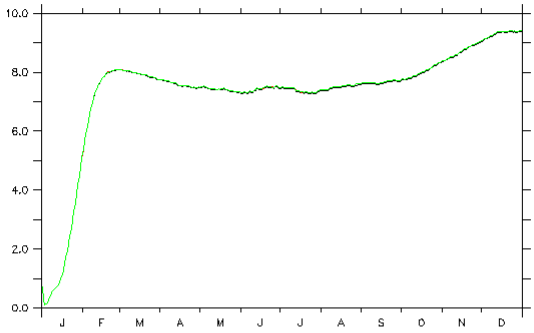

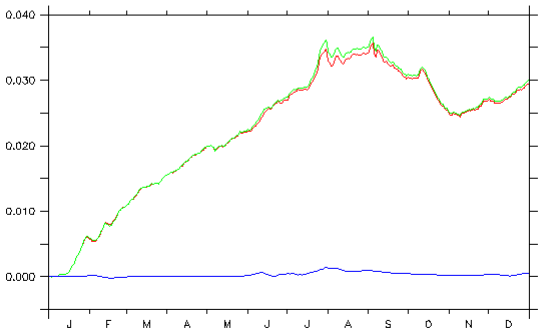

Total Carbon export at 100m (PgC/yr)

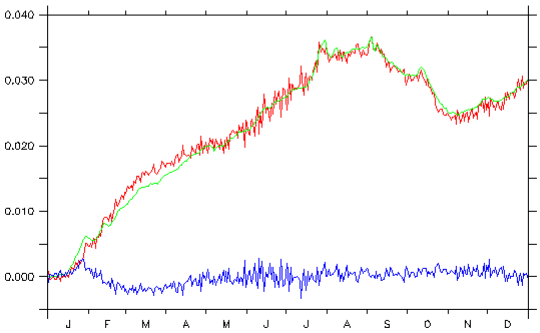

Total Carbon export at 100m (PgC/yr)

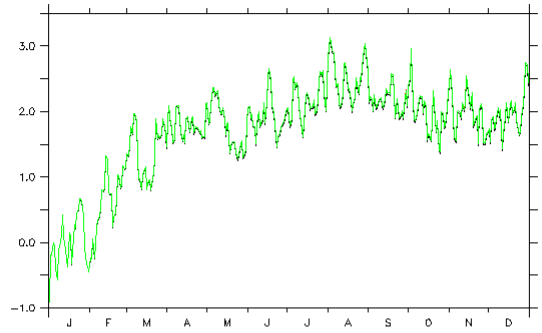

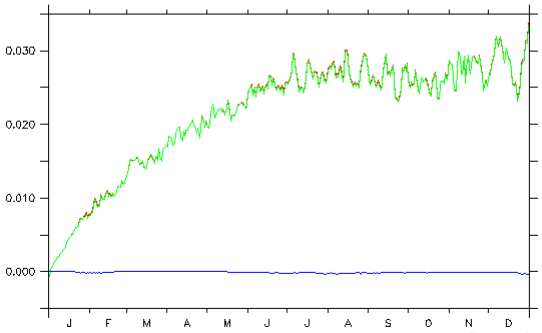

Total Flux of Carbon out of the ocean (PgC/yr)

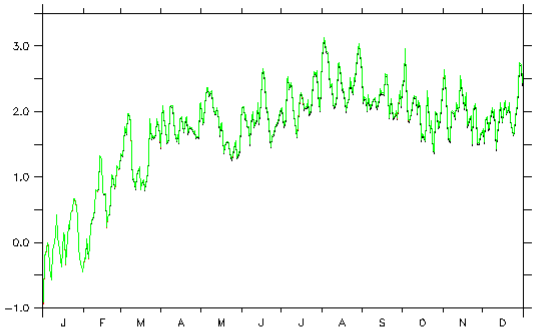

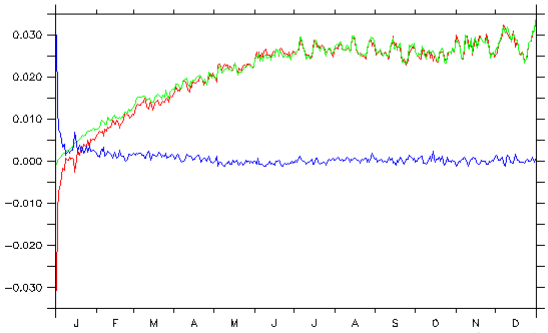

Total Flux of Carbon out of the ocean (PgC/yr)

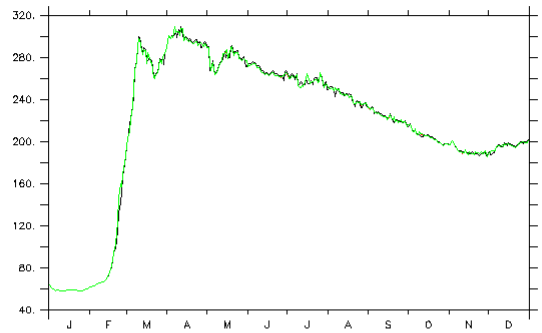

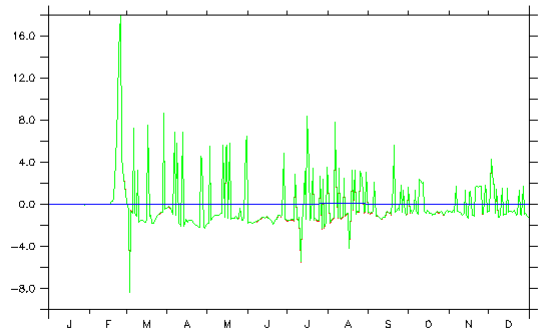

Total denitrification (TgN/yr)

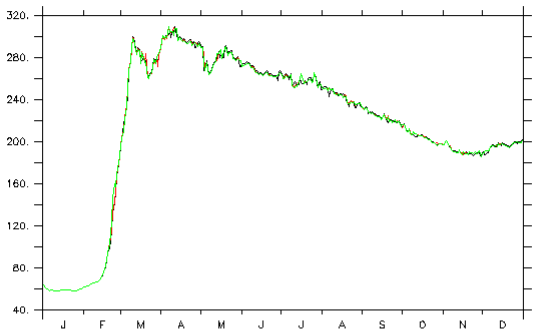

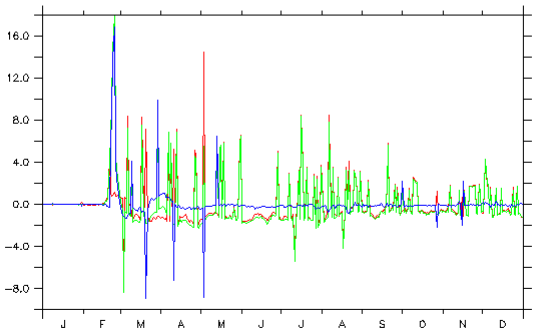

Total denitrification (TgN/yr)

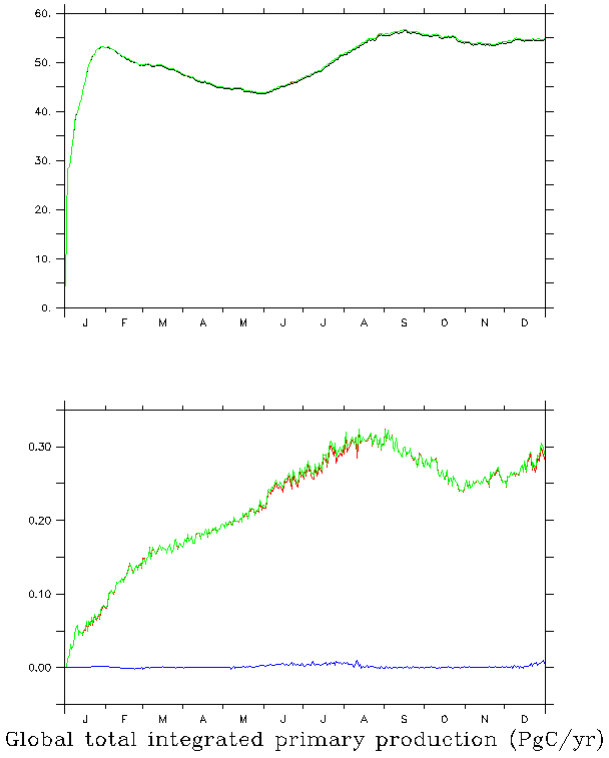

Global total integrated primary production (PgC/yr)

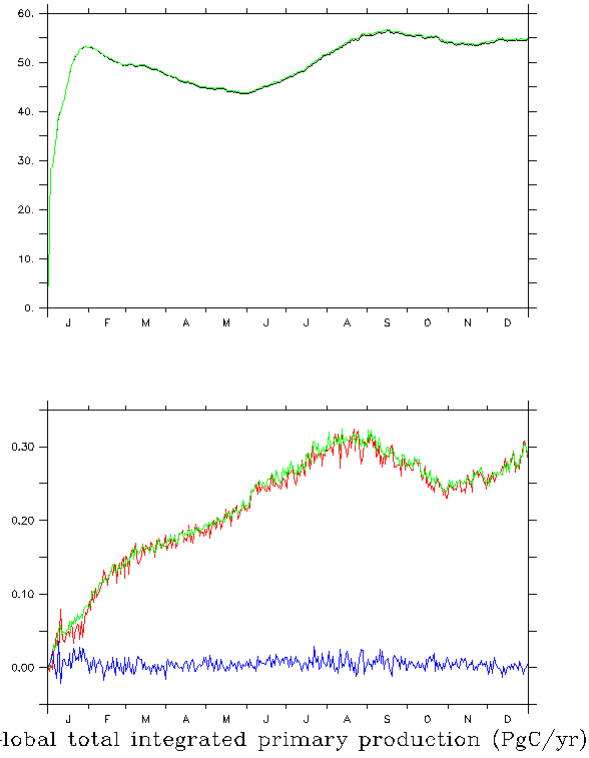

Global total integrated primary production (PgC/yr)

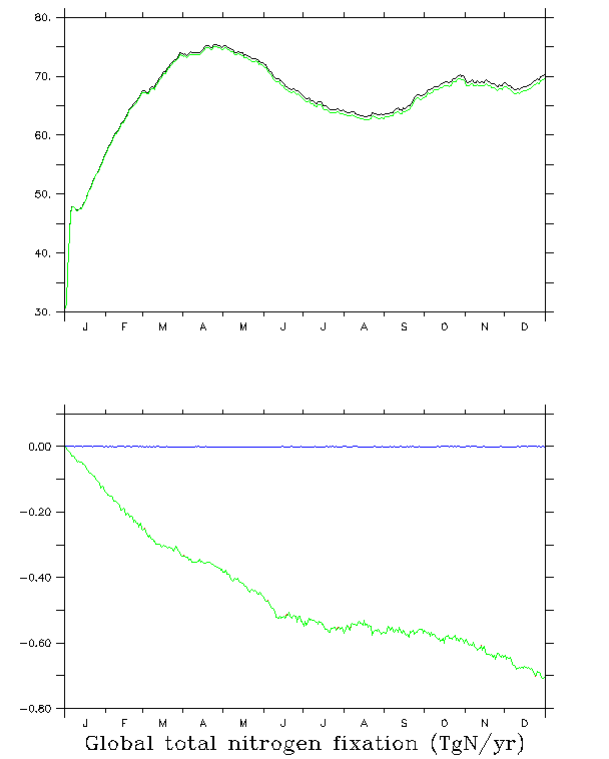

Global total nitrogen fixation (TgN/yr)

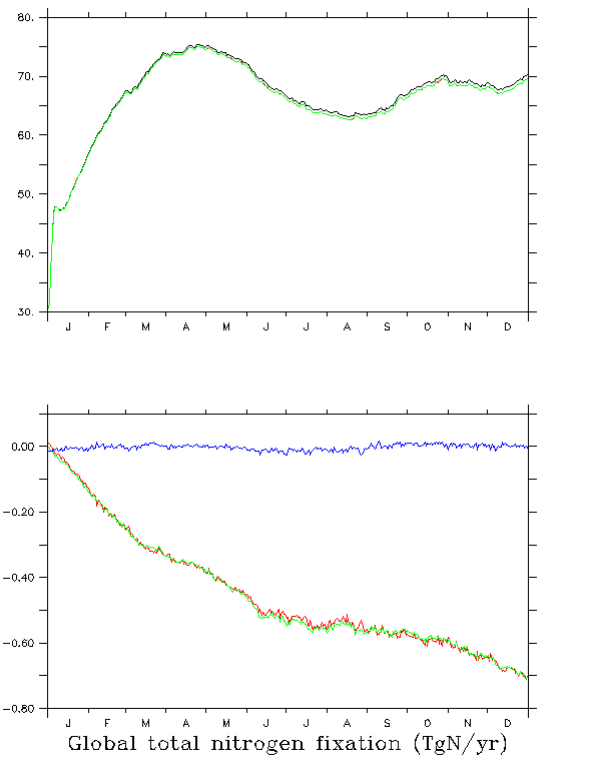

Global total nitrogen fixation (TgN/yr)