Locally conservative OASIS interpolation using target grid nearest neighbours

E. Maisonnave

TR/CMGC/20/166

Abstract

A conservative interpolation between two spherical grids that share no intersection between their unmasked grid point meshes can be achieved by selecting, for each source grid point, a variable number of target grid point neighbours and normalise the interpolation weights by the source/target mesh area ratio. Additional normalisations by source/target distances or Gaussian functions are possible. This interpolation is added to the OASIS library, its results validated and its parallel performance checked.

Table of Contents

Rationale....4
Implementation....5
Validation....7
Performance....8
References....10

Rationale

A recurrent question asked from outside the climate modelling community by computing scientists is: why don't you re-use algorithms or software freely available online ? The following description of the particular issue raised by the river runoff field interpolation should convince anyone that physics sometimes impose to develop specific algorithms.

At the interface of the models which build Earth Systems [1], quantities are exchanged. Rainfall water routed from land surface grid points to river mouths needs to be transferred from land surface or river routing specific models to the debouch grid points of the ocean model included in the system (Figure 1). Considering that Earth Systems are usually discretised on global grids, and integrated during long periods, conservation of the total water quantity matters. The conservative interpolation of runoff required at this stage slightly differs from standard conservative interpolation in that sense that there is possibly no spatial intersection between the source and target grid meshes. It is then impossible to use algorithms relying on those intersections to associate source to target grid points and calculate the appropriate weights. Neither to choose classical nearest neighbour based algorithms, since in the runoff case, the association condition is that all source grid point must find a correspondent in the target grid, and not the opposite.

Several approximate solutions were proposed to address this issue (e.g. [2]). They all apply a global conservation after a non conservative interpolation. The main drawback of this solution is that the cumulated quantity resulting from the non conservation is spread evenly on the target grid points, adding a local error to the initial non conservative interpolation. These answers lead to potentially catastrophic biases that have been emphasised in [3]. To avoid the global conservation step, Aurore Voldoire proposed and tested a locally conservative interpolation of the river outflow. We describe in this work a parallel implementation of this algorithm in the OASIS library.

We would like to distribute the value of each non masked source grid point to its nearest neighbours, which number can be parametrised in the `namcouple` parameter file. The distribution can be :

- uniform : every neighbour receive an even part of the source grid point value
- distance weighted : the inverse of the source/target grid point is used to weight the distribution
- Gaussian weighted : the inverse of the source/target grid point and a Gaussian function are used to weight the distribution

These nearest-neighbour interpolations are not conservative. To impose conservation, to any of these 3 initial distributions, we need to perform an additional weighting by the ratio of the areas between the associated source and target grid points.

Notice that the chosen number of neighbours defines the number of *target* grid neighbours. Consequently, as illustrated in Figure 1, the number of source grid neighbours of each target grid is variable and for that same reason, the weight sum for each target grid point can be different from one. The second weighting, related to conservation, also modulates the weight sum value.

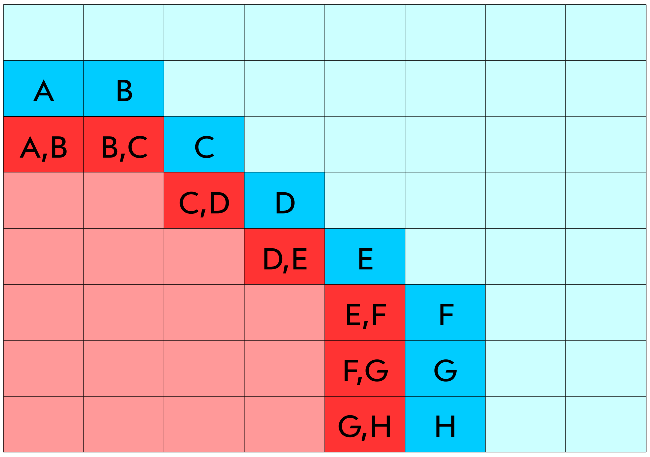

Figure 1: Definition of land points (red), ocean points (blue), land river mouth points (high red) and ocean debouch points (high blue). In this particular case, land and ocean grids are the same and masks are complementary. River mouth (resp. debouch) points belong to land (resp. ocean) masks. The debouch neighbours of each river mouth grid point are named in the river mouth grid mesh

Even though this interpolation can be applied to any kind of grid, it is recommended to reduce the number of involved grid points by setting an appropriate mask, keeping only the river mouth grid points of the land/river routing model grid. It is possible, but not mandatory, to also select a subset of ocean debouch grid points.

To avoid unrealistic routing, the user will take care of excluding the river mouth located on inland areas (lakes …) from the interpolation process and treat them in a separated (global) interpolation. The total sum of these runoff remains low and a global conservation distributed on every ocean point has a marginal effect on the water exchange at their surface.

Implementation

The OASIS coupling library [4] comes with the SCRIP interpolation package [5], that facilitates the exchange of coupling fields between non identically gridded models. Despite its patina, this package merely fits user expectations. But since its first implementation, performance enhancements [6], adaptation to grid specificities or bug corrections [7] slightly makes the code more complex and error prone. A complete rewriting or replacement of the interpolation package is then unavoidable at short to medium term. Taking into account its foretold obsolescence, the proposed implementation of a new locally conservative interpolation in this package relies on modularity : the recycling of existing functions of the code are favoured instead of a new writing. Since our new interpolation implies a nearest neighbour search, we base our implementation on the existing interpolation calculating a distance-weighted or Gaussian-weighted average of n nearest neighbours (`remap_dist_gaus_wgt`). We choose to duplicate and adapt this routine, instead of adding a new option to the existing one. We prefer to have to maintain (for a limited period) two simple routines instead of a single complex one.

At the opposite of the original nearest neighbour interpolation, the neighbours here are in the target grid and have to be searched at the vicinity of each source grid point. To implement that, a switch between the two set of arrays is simply performed at the beginning of the algorithm and the result stored back in the opposite arrays. In between, the double weight normalisation, described in Section "Rationale", is performed. This normalisation choice is facilitated by defining the type of interpolation via new keywords of the `namcouple` file:

- LOCCUNIF, for defining a locally conservative interpolation with uniform weights for all target neighbours
- LOCCDIST, for defining the same interpolation, but weight given by the inverse of the source/target distance-weighted
- LOCCGAUS, for defining the same interpolation, but weight given by the inverse of the source/target distance-weighted and a Gaussian function

In any case, the neighbours number must be defined in the `namcouple`, jointly with the Gaussian variance, for LOCCGAUS option only.

The area ratio needed for conservation is calculated thanks to the grid mesh area values that must be provided by the user in the `areas.nc` file.

In a final phase, weights and addresses (W&A) are reordered before being saved in a specific (`rmp`) file. The same original SCRIP (sequential) routine is re-used for that purpose.

Validation

We take benefit of the OASIS coupling toy available in the `example/test_interpolation` directory to test our interpolation and measure its performance. The CNRM-CM6 dataset of OASIS auxiliary files (`grids.nc`, `masks.nc`, `areas.nc`), both low and high resolution, is necessary to conservatively interpolate values from river mouth grid points only (specific mask). The original `namcouple` is modified to calculate the weights of the new interpolation and to exchange a slowly spatially varying analytic field from the TRIP grid (global regular .5 degree resolution) to the two ORCA1 (LR) and ORCA025 (HR) NEMO grids.

We first controlled the conservation by adding a global CONSERV post-processing to the list of OASIS computations. The log file of the source model includes information about the global sum of the coupling field on unmasked source and target grid points, before the global conservation, together with the correction added by the global CONSERV operation. In any case, this number is strictly smaller than the double precision minimum floating value, proving the conservativeness of our interpolation.

In a second step, we check the effect of the local conservation, comparing the result of a standard nearest-neighbour (distance weighted) interpolation (Figure 2, left) and of our locally conservative (distance normalisation) interpolation (Figure 2, right). It reveals highly different raw values but also different spreading. The locally conservative interpolation distributes on less grid points (4 only per TRIP grid point) but absolute values are also significantly bigger because of the area ratio (about 4 in this HR case) systematically applied to each contribution. Unsurprisingly, bays or straits receive higher quantities since they can gather several TRIP grid point contributions.

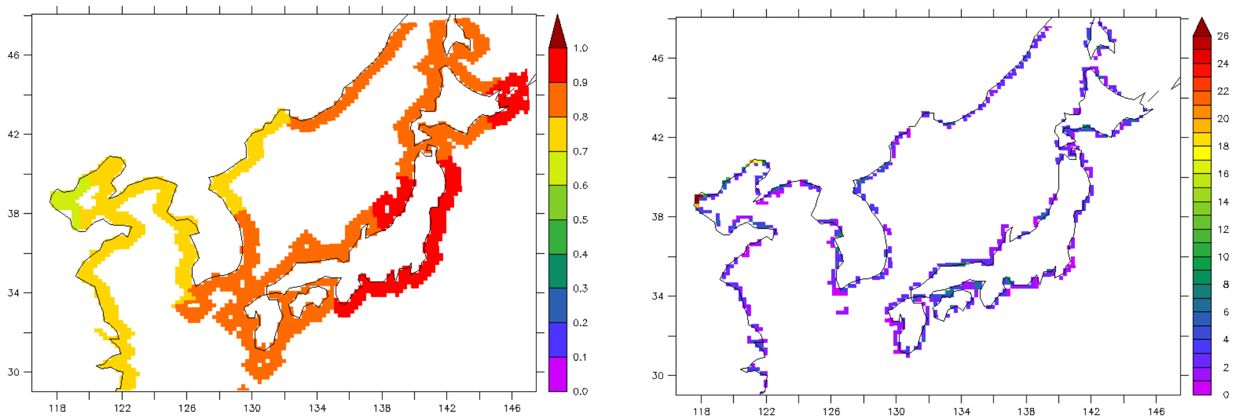

Figure 2: Interpolation of a slowly varying analytical field, with distance weighted nearest-neighbours method (left) and locally conservative normalised by distance (right). The source (target) grid is the TRIP (NEMO) global ½ (one) degree. Both interpolation requires 4 neighbours. No global conservation applied

In a next step, we evaluate the capacity of the (target) neighbour number parameter to spread the incoming values to a larger number of target grid point. The experiment result can be seen in Figure 3. Increasing the number of target neighbours not only increases the number of non zero value of the interpolated field, but also smooths the peak values.

Further analysis (e.g. involving the Gaussian normalisation) are left to future tests that must be carried with realistic runoff input, or real Earth systems. Only partial results can be achieved with analytical fields. However, we propose two hypothesis:

1. the kind of normalisation has one or two order of magnitude less impact on results than the change from nearest-neighbour to locally conservative interpolation.
2. the Gaussian normalisation must be handled with care, since it computations can produce out of order numbers that stop the simulation. In particular, it has been found that distances between inland river mouth (e.g. Lop Nor, Victoria Lake ...) and ORCA ocean debouch grid points are too big to be transformed by a Gaussian function.

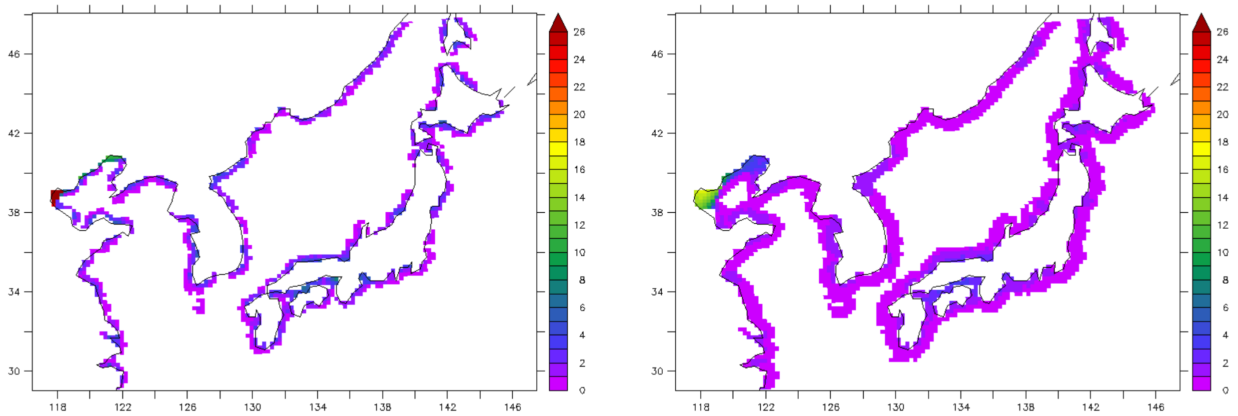

Figure 3: Same than Figure 3, with locally conservative interpolation normalised by distance using 8 (left) and 32 (right) neighbours

The last validation consists in making sure that interpolation W&A are independent from the model parallelism. This was checked for both MPI and OpenMP decompositions.

Performance

The hybrid MPI/OpenMP parallelism of weight computations is inherited from the original nearest-neighbours routine [6]. In addition, our new interpolation is computed in less grid points than for other coupling fields, since runoff involves coastal grid points only. This helps to keep performance at the same level, or higher, than for the other interpolations of the SCRIP package. The original instrumentation is used to measure the performance of our algorithm on the `belenos` machine[1].

1 `https://www.top500.org/system/179853/`

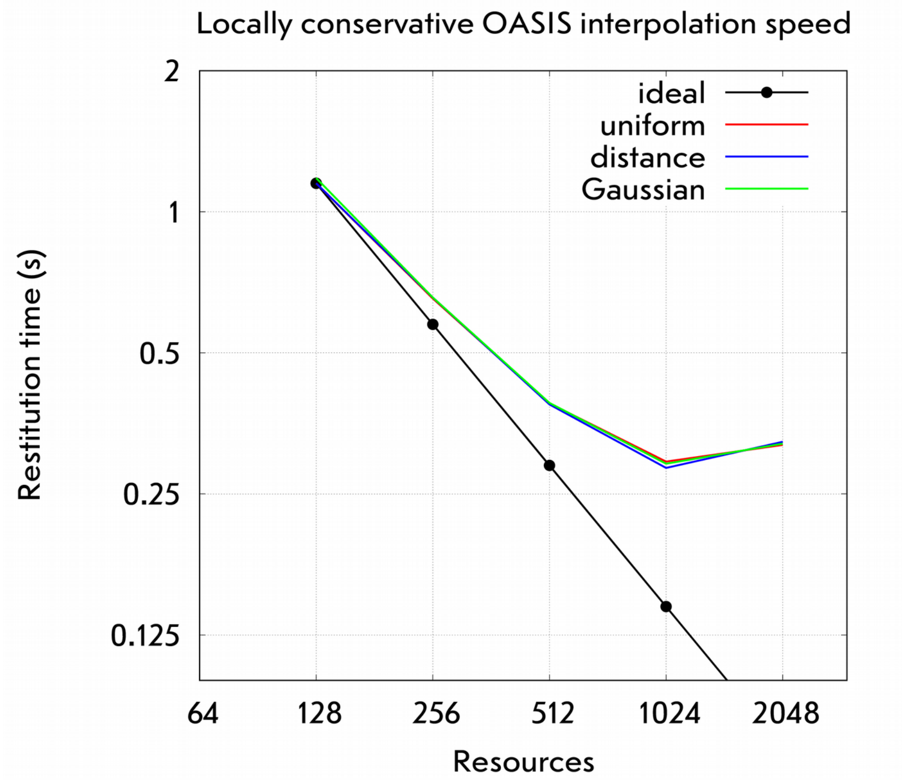

Figure 4: Performance of the locally conservative interpolation algorithm (W&A computations only), from TRIP ½ degree to ORCA025 grids, using 1 to 16 nodes of the `belenos` Météo-France supercomputer

As it can be seen in Figure 4, the W&A computations needed to define the interpolation from TRIP to ORCA025 grid takes less than one second. The W&A re-ordering phase, two order of magnitude faster, is excluded from the result. The MPI scalability quick decrease is not a consequence of slower MPI communications. Our interpolation only involves unmasked grid points, which reduce the number to 6,000 at HR. Rapidly, the MPI decomposition leads to a reduction of computation/communication ratio and a strong computing load unbalance between poorly fed CPU. Because of this particularity, it is recommended to perform the W&A computations in a preparation phase, involving a reasonable number of resources, and re-use the W&A file during the production phase.

Acknowledgements

The author is in debt to Sophie Valcke and Aurore Voldoire for their review. 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. The project IS-ENES3 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 824084.

References

[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, ***11*** **(12)**, 4182-4227

[2] Maisonnave, E. and Terray, L., 2008: How to couple runoff conserving anomaly (in french), Technical Report, **TR/CMGC/08/71**, Cerfacs, France

[3] Voldoire, A., 2020: River to ocean models interpolation. [Research Report] CNRM, Université de Toulouse, Météo-France, CNRS ⟨meteo-02986574⟩

[4] Craig, A., Valcke, S. and 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, https://doi.org/10.5194/gmd-10-3297-2017

[5] Jones, P., 1999: Conservative remapping: First-and second-order conservative remapping, Mon. Weather Rev., **127**, 2204–2210

[6] 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

[7] Maisonnave, E., 2019: OASIS Gaussian interpolation with locally relative source-target distances, Technical Report, **TR/CMGC/19/83**, CECI, UMR CERFACS/CNRS No5318, France