

Review of GMD-2018-129

“Interactive ocean bathymetry and coastlines for simulating the last deglaciation with the Max Planck Institute Earth System Model (MPI-ESM-v1.2) by V. L. Meccia and U. Mikolajewicz for Geosci. Model Dev. Discuss. Doi.org/10.5194/gmd-2018-129

Overall Comments:

This paper presents a complicated multi-step algorithm for automatically and successively modifying the MPI ocean model (MPIOM) bathymetry and land/ocean mask and restart input fields in a transient simulation under evolving boundary conditions such as for ice sheet growth and melt on long time scales. The set of time-stepped ICE6-G_C boundary conditions through the deglacial period are used here to demonstrate the utility and feasibility of the method. However, the ultimate goal is to be able to incorporate active solid earth and ice sheet models to drive ocean bathymetry, volume and coastline changes due to isostatic adjustments and added ice sheet meltwater fluxes that are important for simulating climate change over a glacial-interglacial cycle. This paper documents a new procedure for approaching an extremely challenging technical problem. Up to now, when ocean bathymetry or coastlines need to be changed over the course of a long transient simulation, it necessitates much human intervention and hands-on methods that may have been designed to be used once, and thus is usually done infrequently or not attempted at all. The authors have demonstrated the success and feasibility of this new procedure that can automatically be applied at run-time and updated every 10 years for the long durations needed, though it is designed to be highly specific to MPIOM's particular model grid and architecture. I recommend acceptance after some minor revisions that could help clarify the details of the procedure.

We thank Referee #1 for his/her useful comments. We give a detailed response to each issue in what follows.

Specific comments:

1) It should be mentioned in the procedural description that an important feature of the MPIOM is the employment of partial depth bottom cells, which makes their procedure possible. Models without partial bottom cells would be constrained to discrete values of bottom depth relative to the global mean sea surface (i.e., not including the sea surface height).

We include in the manuscript a section in which the model requirements are described:

“ 2 Ocean model requirements

The algorithms presented in this paper are tailored for the coarse resolution setup of MPIOM but should be easily transferable to other model resolutions or other ocean models having similar assumptions and approximations. MPIOM is a free-surface ocean general circulation model with the hydrostatic and Boussinesq approximations and incompressibility is assumed. It solves the primitive equations on an Arakawa-C grid in the horizontal and a z-grid in the vertical (Maier-Reimer 1997). For freshwater, a mass-flux boundary condition is implemented. A detailed description of the model equations and its physical parametrizations is given in Marsland et al. (2003) while its performance as the ocean component of the MPI-ESM is evaluated by Jungclaus et al. (2013). MPIOM includes an embedded dynamic/thermodynamic sea-ice model (Notz et al., 2013) with a viscous-plastic rheology following Hibler (1979). Sea-ice is swimming in the water. Ice shelves are not included. In this paper,

we use the MPIOM coarse resolution configuration with a curvilinear orthogonal grid (GR30) and two poles (Haak et al., 2003), over Greenland and Antarctica. We decide to use the coarse configuration to reduce the computational time, but the algorithms presented in this paper can easily be adapted to higher resolution grids. In the vertical, the model has 40 unevenly spaced levels, ranging from 15 meters near the surface to several hundred meters in the deep ocean. Vertical discretization includes partial vertical grid cells. Therefore, at each horizontal grid point, the deepest wet cell has a thickness that is adjusted to resolve the discretized bathymetry. On the other hand, the surface layer thickness is also adjusted to account for the sea surface elevation and the sea ice/snow where appropriate.”

2) Lines 107-112: This procedure omits any lakes that form other than those connected to the Caspian and Black seas. The existence of large mid-continental post-glacial lakes formed following the melt and retreat at the southern boundary of the massive Laurentide ice sheet may be important for accurately reproducing the deglacial climate state. Drainage from Lake Agassiz, for example, and the routing of this significant source of meltwater to the ocean, is often hypothesized as causing changes in the meridional overturning circulation during the deglacial period. Excluding such lakes may be necessary in this first implementation of the tool, however, I suggest including a short explanation for why this step is required in this first implementation, the potential ramification, and plans for including them in the future.

Our algorithms are applied within the ocean model and therefore they work on the ocean domain. You are right that mid-continental post-glacial lakes are important for reproducing the deglacial climate state. But, actually, this is a problem that should be treated in the land-model instead of the ocean one. As a matter of fact, including such lakes when considering changes in the routing of the meltwater to the ocean is an ongoing work (as a follow-up of Riddick et al., 2018). For the ocean model, the freshwater fluxes into the ocean is a forcing and the algorithms presented in this paper do not treat the problem of how that forcing is derived.

Because we are solving only the ocean domain, we are interested only in lakes that are connected to the ocean, that is the Black Sea. The Caspian Sea is, indeed, an exemption because it is not connected to the oceans. However, the Caspian Sea is much larger than the other minor lakes. We decided to include it to solve the SST there that might impact on the climate of Central Asia. Therefore, solving the SST of the Caspian Sea might be important for coupled climate models.

We clarify this issue at the beginning of section 2 (*now 3*) *Methodology*:

“Finally, we check for the presence of lakes in the GR30 bathymetry; the Caspian Sea and the Black Sea (under LGM condition, for example) are the only cases that are permitted. Because we are dealing with an ocean model, we are interested in lakes that are connected to the ocean, that is the Black Sea. However, we include the Caspian Sea in our calculations because of its potential impact on the climate of Central Asia. Solving the SST of the Caspian Sea, which is much larger than other minor lakes, might be important for coupled climate simulations. All other lakes need to be removed from the ocean domain either by connecting them to the open ocean or by considering them as land. The atmospheric model component allows accounting for lakes on land (only the thermal component). In the framework of our model system, the adequate place to calculate water storage in lakes is the hydrological discharge model.”

3) As discussed again below, I found section 2.4, which describes the method for

redistributing mass and tracers vertically and horizontally in the process of adjusting the restart files, difficult to follow. For example, what is meant by “vertical re-location” in line 259. A schematic diagram depicting the procedure following changes in depth would help to clarify this procedure.

Thanks for this comment. We realize that we were not clear enough and we reformulate part of section 2.4 (now 3.4) *Adaptation of the restart file in order to conserve mass and tracers*:

“Our approach consists of the following steps:

(a) Vertical redistribution of water and tracers. In this first step, we keep the land-sea mask fixed and we only deal with changes in depth. 2D fields of SSH and 3D fields of tracers are vertically adjusted to the new depth. The strategy here is to conserve the volume and amount of tracers within the water column in each grid point. Considering an individual wet point, the SSH is modified according to changes in depth in order to preserve the ocean volume locally. For example, consider a wet grid point in which the depth is 120.44 meters and the SSH from the restart file is -0.71 meters. The height of the water column results in 120.44 – 0.71 meters and the vertical levels for this configuration are shown in Fig. 5a. After changing the bathymetry, the depth at the same grid point is 122.16 meters. Because the grid area is unchanged, the SSH is lowered to -2.43 meters to conserve the volume of the water column and the vertical levels are adjusted as shown in Fig. 5b. As pointed out before, in MPIOM, the thickness of the uppermost or first layer depends on SSH, whereas the thickness of the deepest or last wet cell is adjusted to the bathymetry. The vertical distribution of tracers is consistently moved along the vertical, taking into account the new layers thickness, in order to preserve the total amount of them within the water column. The behaviour of the algorithms is displayed in Fig. 5 which shows an example of vertical profiles of temperature (Fig. 5c) and salinity (Fig. 5d). This way, the vertical profiles displayed in blue (Fig. 5c and d) are the ones from the original restart. The orange lines (Fig 5c and d) represent the original profiles shifted downward according to the change in depth. The resulting profiles after redistributing vertically the tracers to the new layer's thicknesses are displayed in green (Fig. 5d and d). Values of tracers are constant within each vertical layer of the model (stepped profile). As a result of deepening the bathymetry, the thickness of the bottom (surface) layer increase (decrease), whereas the middle layers remain unchanged (Figs. 5a and b). Therefore, to conserve tracers along the water column, vertical profiles are modified.




Figure 5: Example of vertical redistribution of water and tracers for a single wet grid point. Resulting vertical level configuration (a) before and (b) after changing the bathymetry. Blue and green areas represent the first and last vertical layer thicknesses from the original restart file and after the vertical redistribution, respectively. Hatched areas in (b) represent common thickness layer for both configurations. Vertical profiles of (c) temperature and (d) salinity for the original restart fields (blue), the original profiles shifted downward according to the deepening in bathymetry (orange), and after applying the vertical redistribution in which the profiles are adapted to the model layers (green). Because values of tracers are constant within a model layer, the resulting profiles are stepped. Dots in (c) and (d) represent the upper limit of the first and the lower limit of the last vertical layer.

(b) Horizontal smoothing. The previous step is applied to each wet grid point independently, considering only changes in depth. Therefore, the resulting SSH field might present large gradients between adjacent grid points. To fix this, the SSH field is smoothed by taking into consideration the conservation of mass and tracers. That is, when necessary, values of SSH are modified by moving a volume of water with its tracer properties between adjacent ocean grid points. The maximum permitted horizontal SSH gradient between neighbouring points is set to 0.2 meters, which seems to ensure numerical stability in the ocean model.

(c) Horizontal re-location of water, tracers, sea ice and snow on sea ice when the land-sea mask changes. In step (a) we describe the procedure for dealing with changes in depth only. In this step, the new wet (dry) points resulting from changes in the land-sea mask are filled (emptied). We avoid performing any kind of interpolation in this stage because it would not account for conservation of mass and tracers. Instead, in order to conserve properties, the necessary amount of water and tracers to fill new wet points is taken from other boxes. The simplest approach would be to take water from all ocean boxes. However, this would involve the artificial long-distance transfer of water mass properties. Therefore, we decide to use only adjacent ocean boxes. That is, small volumes of water with its

properties coming from adjacent points is placed into the new wet point until completely filling it. Similarly, the amount of water and tracers from a point which is dried is re-located among the neighbouring wet grid points. This operation is repeated for sea ice and snow on sea ice. There needs to be a compromise between involving only a few neighbouring grid points and the risk of obtaining large horizontal gradients of SSH. Sensitivity tests were performed to achieve the optimal balance for both, filling and emptying procedures.

(d) Horizontal smoothing. Again, we apply step (b) to obtain a sufficiently smooth SSH field to ensure numerical stability when running the model.”

Thus, a figure was added to the revised manuscript and figure numbering has changed accordingly.

4) There is no mention of what is done to adjust velocity components and other related fields that restart the flow fields following changes in land/ocean mask and bathymetry.

The aim of adapting the restart file when bathymetry and land-sea mask change is to account for the conservation of mass and tracers. Therefore, the modification of the fields is done for sea surface height, sea-ice, snow on sea ice (because they are key variables for the total ocean mass), temperature, salinity and passive tracers (because we want to conserve tracers). We do not perform any computation for the other variables (including velocity components) and, therefore, their values remain unmodified. Thus, when wetting a new grid point, the values of the restart file for velocity, for example, will be zero because it is the value for a dry/land point. During the restart procedure MPIOM anyway guarantees that velocity on land points is set to zero. Considering the horizontal resolution that is currently being applied in long simulations with climate models, the advection of momentum is of minor importance. Far from the equator, velocity can be approximated pretty well by frictional geostrophy, as done in the LSG model (Maier-Reimer et al., 1993). Even though velocity is formally a prognostic variable of the ocean model, it is de facto a diagnostic variable whereas the main prognostic ones are temperature and salinity and sea ice. This fact is exploited in typical set-up procedures, where the ocean is initialized with fields of temperature and salinity (from climatology or other model runs) and at rest. However, after one month the velocity field is adapted to the hydrographic fields.

We mention that in section 2.4 (now 3.4) *Adaptation of the restart file in order to conserve mass and tracers*:

“The last modelled state of the ocean with its ocean configuration (restart file) will be used as the initial state for the later setup. Hence, the 2D and 3D fields should be adapted to the new bathymetry and land-sea mask. When carrying out this task, our aim is to account for the conservation of mass and tracers not only at global but also at regional scale. Therefore, the variables that are adapted in this step are SSH, sea-ice, snow on sea ice (for conserving mass) and tracers (for conserving them). From here on, when referring to tracers, we mean temperature, salinity and any passive tracer that MPIOM prognostically resolves (age tracer, radioactive tracer, CFC, etc.). The other model variables (like for example velocities) are not being modified. During the restart process, MPIOM multiplies the velocities with the land-sea mask, thus non-zero velocities are not a problem. However, on the coarse horizontal resolution applied in these very long climate model simulations, the velocities in the ocean are determined essentially by geostrophy and friction and after one month of simulation, the velocity field has adapted to the hydrographic fields. Our approach consists of the following steps:”

5) Section 3 describes how freshwater fluxes are added to the ocean from the melt of grounded ice sheets by the river discharge model, effectively increasing ocean volume.

Section 2 describes the procedure for changing bathymetry and ocean volume using the ICE6-G_C data, implicitly changing volume due to melting ice sheets. Figure 8 shows the procedure works out as the volume change from these two processes match, but it reads like the ocean volume is being changed twice here. Is it because the bathymetry changes are made as a result of the meltwater added slowly over previous interval of time (10 years) since last bathymetric changes? Thus, new volume, added through bathymetry changes, lags or catches up to the volume change due to freshwater added through meltwater over the preceding interval? A schematic showing all of these complicated steps would help clarify.

Ocean volume is not being changed twice. Section 3 describes the transient simulation we performed in order to test the algorithms. The ICE6-G_C reconstructions were used to derive the HR topography and to compute the time-dependent freshwater fluxes into the ocean as in eq. (4). This step is necessary here because the HR topography is prescribed in this experiment and will not be needed when coupling the climate model with the ice-sheet and solid earth models. The ice-sheet growth or decay and the resulting net freshwater flux into the ocean is the only responsible process for the changes in ocean volume and ocean surface area (fig. 6). Therefore, the changes in ocean volume should match the net freshwater fluxes into the ocean (fig. 8), except for the delay caused by the time needed within the hydrological discharge model to transport the water to the ocean.

When running the model for 10 years with a fixed bathymetry, the imbalance of net freshwater fluxes into the ocean affects the mean SSH, which is simply a consequence that during the 10 year simulation the ocean volume has changed and is not matching exactly the bathymetry any more. After 10 years, the bathymetry and land-sea mask change and the mass of water is distributed to the new configuration. This way, the mean SSH is being preserved within the simulation. This would work perfectly if a) the model ocean bathymetry had the same horizontal resolution than the topography used to compute the freshwater fluxes into the ocean; b) the data used for computing the freshwater fluxes and HR topography was consistent because it accounts for conservation of water. However, a) reducing the resolution from HR to GR30 results in a smoother bathymetry that might result in differences in ocean volume (between the one that would be for HR and the one that results for GR30); b) we aim at writing an algorithm independently of the data used as forcing and so we consider the potential inconsistencies in the reconstructions. Hence, the aim of step 2.3 is to correct these two possible sources of inconsistencies. The strategy is to match the last ocean volume state (GR30, which already accounts for the accumulated freshwater fluxes during the previous 10 years) with the ocean volume of the new configuration (GR30 that might contain artificial changes in ocean volume due to the loss of bathymetric details when reducing resolution and to potential inconsistencies in the HR topography). The resulting ocean GR30 bathymetry accounts for changes in the ocean volume only due to the imbalanced net freshwater fluxes.

We understand that this issue is not clear enough in the original manuscript and we modify section 2.3 (now 3.3) *Matching changes in ocean volume and freshwater fluxes into the ocean*:

“The growth or decay of ice sheets and the resulting net freshwater flux into the ocean is the only responsible mechanism to change the volume of the ocean in MPIOM, as incompressibility is assumed. Otherwise, effects like thermal expansion could be important as well. When running the model with a fixed bathymetry, the net freshwater fluxes into the ocean affect the mean SSH and consequently the thickness of the uppermost ocean layer. When a new ocean bathymetry is derived in a formally independent process, the mass of water is distributed to the new configuration. Then, both estimates of the ocean volume should be consistent, and therefore, the mean SSH and mean thickness of the surface

layer should be preserved within the simulation for all restart points. However, de facto, this is not always the case mainly for two reasons. On the one hand, the HR reconstructions might show inconsistencies if they do not account for water conservation. On the other hand, reducing the resolution from HR to GR30 can cause disagreements in the ocean volume due to the loss of details in the bathymetry field. The aim of this step is to remove these two possible sources of inconsistencies. The procedure is to match the last GR30 ocean volume, which already accounts for the freshwater fluxes into the ocean, with the ocean volume of the new GR30 configuration, by performing the following steps:

and later in the same section:

“In this way, the resulting ocean GR30 bathymetry accounts for changes in the ocean volume due only to the freshwater fluxes into the ocean. There might exist slight discrepancies produced by the last step. However, by removing possible artificial changes in ocean volume, the procedure ensures that the mean SSH is reasonably well preserved, independently of the freshwater fluxes and the prescribed HR dataset.”

Minor comments by line:

63: “In the frame of the project...” -- awkward phrase to start the sentence.

Reformulated to:

“Our long-term goal in the context of the project “From the Last Interglacial to the Anthropocene: Modeling a Complete Glacial Cycle – (PalMod)” is to simulate the last termination with a coupled ice sheet-solid earth-climate model with interactive coastlines and topography forced only with solar insolation and greenhouse gases concentration.”

150: OK--omitting Arctic and Southern Oceans in this list because they are contiguous with the other major ocean basins?

Yes. Atlantic-Pacific-Indian Oceans was changed to World Oceans:

“The strategy is to keep only the wet points that are directly connected to one of the following basins: World Oceans, Mediterranean Sea, Red Sea, Black Sea and Caspian Sea.”

Accordingly, it also was changed in line 163 of the original manuscript.

171: “Specific regions are examined in detail and modified if necessary.” Also, “...look at the HR land-sea mask...” Suggests human oversight here, but I suspect this is not the case. This step must use some rather specialized coding because many specific regions in the HR mask are checked against the new GR30 mask. What methods are used to make this more automatic? Are multiple solutions possible to obtain using the fraction ocean in the new GR30 to identify pathways that connect new regions?

Yes, this step is done automatically by the code. We explain it in the revised manuscript:

“Specific regions are considered in detail for further checking and the GR30 land-sea mask is, therefore, modified if necessary. First, we check if North and South America are connected by land or artificially separated by the remapping. Then, we check some straits or channels (Strait of Gibraltar,

Bab-el-Mandeb, Bosphorus, Denmark Strait, Faroe-Shetland Channel, Northwest Passage, Nares Strait and the Strait of Sicily), islands (Indonesia and Japan) and peninsulas (Florida, Thailand-Malaysia, Kamchatka, Italy and the Scandinavian Peninsula). The strategy here is to automatically control if the straits/channels are open or closed and if the islands/peninsulas are isolated from or connected to the mainland in the HR land-sea mask. To automatically perform this task, the algorithm finds the path of connection between two points apart. This is done in a restricted domain around the region of interest. For example, when checking the opening or closure of a strait, the points to be connected are wet points located in each side of the strait. If the algorithm finds that the path of connection between both points is always within the ocean, that means that the strait is open. Instead, if the path of connection is blocked by land, that means that the strait is closed. The location of each pair of points was manually and carefully decided for each region and is fixed in the code. It was tested that those points do not change from wet to dry or vice versa during the last deglaciation. The approach is applied to each specific region mentioned before and both resolutions, HR and GR30. When necessary, the GR30 land-sea mask is regionally modified to be consistent with the HR data. The information of the fraction ocean is used to decide about the path of the opening or closure. Being the fraction ocean a float number it is highly unlikely to obtain multiple solutions. In that case, the algorithm would choose the first solution found.”

Section 2.3 and section 3, lines 300-308: Globally adjusting ocean depth to keep global mean SSH constant, and volume changes through adding freshwater from melting ice sheets. Are the steps employed in Section 2.3, done every timestep after freshwater from melting grounded ice sheets is added, thus increasing global ocean volume through increases in SSH?

In the transient simulation we performed, the freshwater fluxes into the ocean are being incorporated every time step, whereas the procedure described in section 2.3 is being applied only for constructing a new ocean bathymetry, in this case, every 10 years. We clarify this aspect in the revised manuscript. Please, see the answer to your point 5) of “*Specific comments*” for more details.

248: Do the final changes made to depth as described in step 2.3(d) require iterations back to (3)?

No. It is important for the model stability that the final depth satisfies eq. (2). For example, a new wet point must have a depth smaller than the thickness of the surface layer in the model. This is to avoid involving more than one layer when adapting the restart file. Going back to this criteria might destroy some corrections done in step 2.3 (now 3.3) as stated in the manuscript:

“In this way, the resulting ocean GR30 bathymetry accounts for changes in the ocean volume due only to the freshwater fluxes into the ocean. There might exist slight discrepancies produced by the last step. However, by removing possible artificial changes in ocean volume, the procedure ensures that the mean SSH is reasonably well preserved, independently of the freshwater fluxes and the prescribed HR dataset.”

Yet, we demonstrated that water is being conserved in a long-term transient simulation (fig. 8) indicating that, if the discrepancies still exist, they are not large enough to affect the mass conservation.

259: Section 2.4(a) What is meant by “vertical re-location”? I find this section difficult to understand the actual details of the method even after looking at Figure 5. A schematic illustrating the method would be helpful, especially for locations that are already “wet” point that become deeper. How are tracers at mid-depth changed? Also does the process of vertical

re-location” result in lateral gradients at depth in the ocean, even after horizontal smoothing?
263: “new layer’s thickness”?

We better explained the methods and we added a new figure in the revised manuscript. Please, see the answer to your point 3) of “*Specific comments*” for more details.

300: The “instantaneous time derivative of the gridded ice thickness” is computed for the meltwater fluxes added by the hydrological discharge model. Does this gridded ice sheet thickness come from the ICE6-G_C data interpolated in time to every 10 years? Does “instantaneously” mean the meltwater flux calculation is done every time step, or for every 10-year interval?

Yes, the gridded ice sheet thickness comes from the interpolation in time to every 10 years. The derivative is done once for every 10 years interval. We clarify this in the revised manuscript:

“The interpolated ICE6-G_C reconstructions were also used to compute the time-dependent freshwater fluxes into the ocean. First, the 10-year interval time derivative of the gridded ice thickness is calculated. Only the ice-sheet thicknesses at grounded points are considered. The time rate of change of this quantity is then divided by the density ratio between ice and freshwater to obtain the extra freshwater flux into the ocean:

eq. (4)

where *Ice* is the ice thickness of the grounded-ice sheets and *R* the density ratio between ice and freshwater. The resulting value is considered constant for a period of 10 years, although it is introduced to the model every time step. The extra freshwater is transported into the ocean through a hydrological discharge (HD) model which considers the changes in river routing (Riddick et al., 2018).”

366: “...called with a maximum of three input files.” The description of the tool software and scripts is short. A bit more information about how it is used in practice and integrated into the model run-time would be helpful. For example, how does it interface with the model during run-time? Which input files are needed? Is the tool launched in the main run script, at the start of a restart submission using files from the previous submission? Because restart files are generated, does this mean that the 10 year time interval between bathymetry changes fixes the maximum number of years between resubmission?

A more detailed description is included in the revised manuscript:

“The principal tool consists of shell scripts that are called with a maximum of three input files. All the calculations are performed with CDO commands and programs written in FORTRAN. The tool can easily be included at the end of the main run script without the necessity of interrupting the simulation. There are two shell scripts that need to be called after the restart file is written by the model. The first one generates the new bathymetry file for running MPIOM. Two input files are required to run this script. The first one corresponds to a NetCDF file containing the new HR bathymetry. The second input is an ASCII file which corresponds to the previous GR30 bathymetry as it was read by MPIOM. The output of this shell script is an ASCII file containing the new GR30 bathymetry to be read by the model. As a result, this script replaces the old bathymetry file to run MPIOM with the new one. The second shell script adapts the restart file generated by the model to the new ocean configuration. This script needs three input files. The first and second ones correspond to the old and new bathymetry files as read by MPIOM, respectively. The last input is the restart file generated by the model in NetCDF format. The output is the modified restart file in NetCDF format to replace the original one. The

execution of this tool needs the restart file generated by the model as input. Therefore, it can be called only after a restart file is generated. Contrary, it is possible to resubmit the job without applying the tool, that is with fixed bathymetry, land-sea mask and, therefore, unmodified restart file. This allows for a shorter number of years between resubmissions than the ones required for changing the bathymetry. Consequently, the tool is easy to apply and it is fast, taking less than a minute to run on a workstation.”