



Fast GCM - ice sheet model coupling software OBLIMAP 2.0, including on-line embeddable mapping routines

Thomas Reerink and Roderik van de Wal

Institute for Marine and Atmospheric research Utrecht, Utrecht University, 3508 TA Utrecht, The Netherlands

Correspondence to: Thomas Reerink (tjreerink@gmail.com)

Abstract. This paper accompanies the second OBLIMAP open source release. The package is developed to map climate fields between a general circulation model (GCM) and an ice sheet model (ISM) in both directions by using optimal aligned oblique projections, which minimize distortions. Both grids may be irregularly spaced and the ratio of the grids is allowed to differ largely. The stand-alone version of OBLIMAP is a powerful tool to map various differently gridded datasets on one uniform
5 ISM grid with an optimal centered projection. This is demonstrated by relevant applications concerning the major ice caps. As this applies also for the reverse mapping direction, it can be used as an off-line coupler. Besides, OBLIMAP 2.0 is an embeddable GCM - ISM coupler, suited for high frequent on-line coupled experiments. A new fast scan method is presented as an alternative for the former time consuming grid search strategy, realising a performance gain of several orders of magnitude and enabling the mapping of high resolution datasets with a much larger number of grid nodes. Further a highly flexible masked
10 mapping option is added. The limitations of the fast scan method with respect to unstructured and adaptive grids are discussed together with several proposed parallel implementations in order to achieve another performance gain.

1 Introduction

Ice caps are part of the climate system and interact with the atmosphere and the ocean via various feedback mechanisms. In order to simulate their interaction, ice sheet models (ISM's) need to be coupled with general circulation models (GCM's). An
15 atmospheric GCM (AGCM) and an oceanic GCM (OGCM) can be mutually coupled with e.g. the external OASIS coupler Valcke (2013), in order to study the ocean - atmosphere interaction. In contrast to the geographical coordinates which are used by the GCM's, the ice dynamic equations are often solved on rectangular coordinates, due to the type of its equations. This requires an additional projection step when coupling an ISM with an GCM.

Because ISM's are predominantly influenced by the atmospheric forcing, coupling them with an AGCM is from the ISM
20 perspective self-evident. Nevertheless, if ice shelves are included, the coupling with an OGCM makes sense as well because ice shelves are sensitive to the ocean temperatures, which strongly affect the dynamics of the shelves and the ice sheet behind the ice shelf (e.g. Holland and Jenkins, 1999). On the other hand an ISM model provides output in terms of bedrock, surface height and albedo changes which affect the climate and needs to be provided back to the GCM.



In order to couple GCM's and ISM's, the GCM fields which are defined on a grid representing the curved Earth surface have to be mapped at the ISM grid which coincides with a flat surface, and vice versa. OBLIMAP (Reerink et al., 2010) performs this technical mapping task, which comprises the combination of projection and interpolation in both directions.

The oblique stereographic (SG) and the oblique Lambert azimuthal equal-area (LAEA) and their inverse projections are available in OBLIMAP because they are suited projections given the constraints involved with this type of geographical mappings. The conformal SG projection preserves angles which is a welcome property for direction dependent geometries and velocity fields. The SG projection is nearly area conservative if the projection is optimally aligned. Areas remain conserved under the LAEA projection, while this projection is nearly conformal if optimally aligned. The area-invariant property of this projection is interesting with respect to conservative mapping. However, note that conserved mapping requires in addition a conserved interpolation method as well.

Both projection methods are azimuthal (i.e. perspective), which means that with exception of the point of projection itself, the entire domain can be mapped without any singularities. Either a sphere or an ellipsoid is a natural choice to represent the Earth surface. Therefore in OBLIMAP the SG and LAEA projection can be used in combination with the sphere and the (default WGS84) ellipsoid.

The projected points which are used to obtain interpolated values at the destination grid points, are irregularly distributed with respect to this destination grid points. Furthermore the ratio of the grid resolutions vary orders of magnitude among OBLIMAP's applications. Therefore OBLIMAP uses a quadrant and radius interpolation method (Reerink et al., 2010) in which the weighting is based on Shepard (1968). These interpolation methods are suitable for large resolution differences and data gaps, and robustly deal with the irregular located interpolation source points. Due to this approach OBLIMAP is also capable to cope with departure and destination grid coordinates which are irregularly spaced.

Using an optimized projection is fundamental in OBLIMAP's strategy to provide high mapping quality and to obtain results which are close to conservative after to and fro mapping. To achieve this OBLIMAP uses oblique projections, for which only the geographical coordinates of the center of the area of interest need to be specified in order to obtain an optimal centered projection. With the optimal centered projection the SG projection is close to area preserving and the point of projection, which is the only singular point in the mapping, is by definition at the opposite side of the area of interest. In addition the distortions of the mapped local distances is minimalized. A range of to and fro mapping tests with OBLIMAP using this strategy proved the robustness of this method and showed accurate results (Reerink et al., 2010).

The stand-alone version of OBLIMAP is extremely suitable and convenient for mapping various input datasets on one ISM grid with an optimal centered projection. These input datasets may vary in grid resolution and in grid extent, and can be defined on either a grid representing the curved Earth surface or on an ISM domain coinciding with a flat surface. In the latter case a remapping is required if the selected projection and grid configuration differ.

Examples of typical input datasets which are used to initialize and force an ISM are topographic datasets including surface height, bedrock level and ice thickness fields, atmospheric forcing datasets containing surface massbalance, surface temperature, refreezing and runoff fields, and datasets containing fields like the geothermal heat flux, the ice surface velocities for



validation or for initializing the ice shelf velocity field and for instance the ocean surface temperature which could be used in the ice shelf basal melt parameterisation.

Because the dynamics of the ice caps are supposed to be sensitive to the dynamics in the coastal zones, accurate mapping of the different datasets is essential in order to ensure that their complex mapped patterns, related to the varying topography, fully match.

ISM's usually cover a limited local area compared to GCM's. Therefore only a part of the GCM points will participate in the ISM to GCM mapping. These points are distinguished from the remaining points by a participation mask, and will default be merged with the pre-mapped field values of the points which didn't participate in the mapping.

For example, in order to equally map their initial topographic and atmospheric and other forcing fields on their preferred ISM grids with an optimal centered projection, Helsen et al. (2012, 2013); De Boer et al. (2013, 2014, 2015) used OBLIMAP's first release (Reerink et al., 2010) for preparing their ice sheet modelling input fields. Besides the use of GCM output as atmospheric forcing, the higher resolution fields provided by regional climate models like RACMO (van Meijgaard et al., 2009) might be favoured as forcing of the ISM (Helsen et al., 2013), the more because they deliver direct applicable products like the surface mass balance (SMB) which are of interest for the ice sheet modeller.

With the RACMO2 datasets for Greenland (van Angelen et al., 2014; Noël et al., 2015) and Antarctica (Van Wessem et al., 2014, 2016) higher resolution atmospheric forcing datasets are available. For instance the RACMO2 SMB is a field which contains only relevant valid values at ice covered grid points, distinguished by the assignment of a missing value for the remaining grid points. This requires a method which accurately maps masked fields. Therefore OBLIMAP 2.0 provides the new 'masked mapping' method.

At the time of OBLIMAP's first release OBLIMAP typically had to cope with situations in which GCM fields, which are defined on a rather coarse grid ($\sim 1^\circ$), are mapped on ISM grids with a resolution of about 10 km or coarser. In the meantime new high resolution topographic datasets for Greenland (Bamber et al., 2013) and Antarctica (Fretwell et al., 2013) have become publicly available. The large gain in resolution refinement for these datasets confronted OBLIMAP with a performance challenge which has not been an issue at the time of OBLIMAP's first release. Mapping or remapping for instance the entire Greenland area with a 1x1, 2x2 or 5x5 km resolution is much more demanding with respect to the computational time of OBLIMAP's so called scan phase. The scan phase computes the projection of all grid coordinates and calculates the distances over the surface of the destination grid, and based on them the nearby projected points which contribute to the interpolation are selected. With an alternative 'fast scan method' OBLIMAP 2.0 realizes a large performance gain and therefore enables the mapping of the high resolution datasets.

In order to simulate the interaction of an ice sheet with the ocean and atmosphere in sufficient detail, on-line high frequent GCM - ISM coupling is required in which the mutual feedback processes are implicitly included. A solution which is computational efficient, will use embeddable coupling routines and for instance an embeddable ISM. With on-line coupling we mean that the field exchange takes place during a simultaneous GCM - IM run, which can be achieved with either an external or an embedded coupler, the latter means that the coupling routines are directly called from the GCM code for instance. With OBLIMAP's redesign the mapping routines can be used embedded as well now. The introduction of a so called 'dynamic data



object' solves the main I/O bottleneck. This combination makes the OBLIMAP mapping routines suitable for high frequent on-line coupling, which we consider also as one of the main achievements of OBLIMAP 2.0.

In Sect. 3 a few other new key features are described, whereas a complete description is available in the OBLIMAP manual. The primary objective of this new manual is to explain how the configuration variables can be configured by the user.

5 The OBLIMAP-package, including the stand-alone OBLIMAP 2.0 code, several applications and the OBLIMAP manual, is available at the GMD site (see supplementary material <http://www.geosci-model-dev.net/???-supplement.zip>) and is distributed under the terms of the GNU General Public License.

2 General overview of OBLIMAP 2.0

OBLIMAP is divided into two phases: a prior scan phase and a post scan phase. In the scan phase the departure grid points
10 are projected on the destination surface corresponding with the destination grid. For each destination grid point the projected departure grid points which contribute to the interpolation, depending on the selected interpolation method, have to be selected. As this is time consuming, the resulting grid addresses and relative distances are stored by writing them to the scanned indices and distances file: the SID file.

In the post scan phase the dynamic data object (DDO) is initialized by loading the SID file data. Thereafter the fast mapping
15 of multiple fields, layers and records is that fast that the computational time is more or less negligible for most common applications. In a later stage the same mapping with the same departure grid - destination grid combination can be repeated with the scan phase switched off by reading the earlier produced SID file.

Beside the so-called full scan method a fast scan method is available in OBLIMAP 2.0, which is orders of magnitude faster
20 than the full scan method. The fast scan method can be applied both for regular and gaussian reduced grids. A schematic representation of OBLIMAP 2.0 stand-alone is given in Fig. 1.

In order to couple an ISM on-line with a GCM, one of them should host the other model and the OBLIMAP mapping routines. Due to the complexity of most GCM's, the GCM will be in most cases the obvious host model candidate. The current OBLIMAP 2.0 mapping routines are designed for embedded usage. Of course this approach requires an embeddable ISM as well.

25 Because in the on-line case the OBLIMAP mapping routines are used embedded, it is recommended to conduct the scanning in both mapping directions off-line and prior to the coupled run. In the initialization phase of the coupled run both DDO's are loaded by reading both SID files, and the fast mapping can be repeatedly used as shown in the scheme of Fig. 2, where the time steps of the ISM, the GCM and the coupling interval might differ from each other and change over time.

OBLIMAP 2.0 works with a separate configuration file for each mapping direction. A configuration file is an ascii file containing the configuration variables which enables to configure each mapping. The number and order of the listed configuration
30 variables in the configuration file is not prescribed. Those configuration variables which are not listed keep their predefined OBLIMAP settings. The 67 configuration variables are described in the OBLIMAP manual (see the supplementary OBLIMAP



manual). Just one configuration file for both embedded mapping directions will be sufficient if, as recommended in an on-line coupled experiment, the coupled run itself uses only the fast mapping mode.

The post scan configuration parameters can be changed at any time without the obligation of repeating the scan phase. They could even be changed during an on-line coupled experiment. This contrasts with the scan configuration parameters, as soon as a user decides to alter them, the time consuming scan phase has to be repeated.

3 New OBLIMAP features and achievements

In this section and in Table 1 the most important new or extended OBLIMAP features and achievements are described.

Before we proceed we emphasize that a distinction is made between data gaps and invalid points in OBLIMAP. If for example a forcing field covers the Earth surface up to a latitude of 87° South, it means that there will be no departure grid points available within the polar area for the fields which are mapped on an ISM grid for Antarctica. This is what we call a data gap. In such cases with the quadrant interpolation method OBLIMAP succeeds in establishing an reasonable well interpolation of the forcing fields for the polar area on the ISM grid (Reerink et al., 2010). Data gaps are thus areas where no departure grid points are available. By contrast, invalid points are departure grid points which contain an invalid value. With OBLIMAP 2.0 they can optionally be excluded for the interpolation by masked mapping (see Sect. 3.2) in the post scan phase.

15 3.1 Scanning

The objective of the scan phase is to identify all departure grid points which are projected close to each considered destination grid point and thus contribute to their interpolation. Each projected departure grid point which is indeed used for the interpolation of the considered destination grid point is called a 'contribution' of this destination grid point. The number of detected contributing points per destination grid point varies due to the selected interpolation method, but also due to their availability near grid domain borders for instance. In this stage no distinction is made between masked or non masked invalid value contributions, all of them are involved. The masking is a post-scan feature, see Sect. 3.2.

The departure grid indices of each contribution and the distance over the surface associated with the destination grid are written to the SID file. Per destination grid point one line is written to the SID file with the following format (see also the header of each SID file): The indices of the destination grid point are followed by n , the number of contributions for this destination grid point, whereafter the n contributions follow with for each contribution its departure grid indices and its relative distance over the destination surface to the destination grid point. n may differ per destination grid point. This format which copes with the fluctuating amount of data per destination grid point is also the basis of the DDO.

In fact OBLIMAP needs in its scan phase the grid coordinates and has to know the grid shape. The actual values at the grid points are not used and therefore not required. The selection of the projected contributions is always based on their relative distance to the destination grid point measured over the destination grid surface. Also this distance calculation requires nothing more than the coordinates of the destination grid point and the projected departure grid point.



Assuming that the ISM grid area is a rectangle defined by the coordinates of the bottom-left and the upper-right corners, the participation mask for the GCM grid can be constructed with negligible computational costs.

3.1.1 Full scan

In the full scan method for each participating destination grid point a full scan over all projected departure grid points is conducted in order to select all contributions for this destination grid point. Of course this is computational not very efficient, but it is easy to implement and entirely robust with respect to grid configurations around the geographical poles, cyclic grids and data gaps. The full scan method works well for any grid type: regular, gauss-reduced and unstructured grids. OBLIMAP's prior version relied only on the full scan method.

3.1.2 Fast scan

The fast scan algorithm uses the fact that the closest projected departure grid point contributions of the previous neighbour destination grid point are known. The closest contribution is called the pivot. Searching the contributions for the destination grid point under consideration can be conducted within a limited block of grid points surrounding the pivot. If this block is constructed sufficiently wide, no relevant contributions will be missed.

At the very first considered destination grid point a full scan over all departure grid points is conducted in order to ensure that the correct contributions are determined. This yields a pivot for the next destination grid point. So for the second considered destination grid point the contribution scan will be conducted within a limited block of departure grid points, whereafter the pivot will be updated for the next destination grid point. This is repeated until the end of the first row of the 2D-grid. When jumping to the next row of the destination grid, the nearest contribution of the neighbour destination grid point at the previous row is taken as the pivot for the first destination grid point at the next row. And so on.

The use of the participation mask in the reverse ISM to GCM mapping direction, complicates this approach. If there is a certain number of non-participating destination grid points at the start of a row, the 'jump approach' will be applied to the first participating destination grid point at this row. There are situations possible in which there is no participating point available at the previous row either. Or no contribution is detected at all at the previous scanned neighbouring points. In such exceptional cases a full scan has to be conducted for these destination grid points, decreasing the computational time performance.

Evidently in the fast scan method a correct estimate of the block size is essential. In order to do so OBLIMAP estimates a scan block size radius (b), which should be as small as possible to obtain the fastest performance while at the same time it should be large enough to obtain identical results as with the full scan method (which serves as the quality-performance benchmark).

The estimate of b depends on the ratio of the local destination grid resolution over the local departure grid resolution and differs for the quadrant and the radius method, because for the radius method an additional term is included to cover the entire possible area of radius method contributions. Note that the range of the ratio of the grids may be substantial. In the development phase this has been one of the key challenges to make the method general robust. The local grid resolution can differ in each grid direction, in the current method its estimate is based on the distance to the next grid point of the row. In general the



projected grid points will not align with the destination grid points, but show a certain angle. Therefore an ample rounded b is taken which in fact for most cases yields a small under performance in computational time but ensures a robust quality performance.

Fig. 3 shows how b is determined for a regular example of a GCM to ISM mapping using the quadrant method. Fig. 4 shows how b is determined for a regular example of an ISM to GCM mapping using the radius method. Fig. 3 and Fig. 4 describe the main concept of the fast scan method: For each destination grid point a scan over a local block of departure grid points is conducted where the local block has to be well positioned. In fact the fast scan method basically uses the fact that the local block of the previous search will be close to the local block of the next search for common geographic grids.

With the optional and additional 'dynamic block size method' the initial local estimated block size is increased step by step until no additional contributions are found. The method is coded in such way that only added block borders are checked (for efficiency) and it is consistent with the hereafter mentioned cyclic approach.

Beside, additional techniques are implemented: 1. In order to deal with cyclic longitudinal domain boundaries. At the east and west longitudinal border of a global departure grid, points at the other side of the departure grid might contain contributions as well. In these cases the block size is extended cyclic by a special treatment. 2. Some equidistant longitude - latitude geographical grids are constructed in such way that the grid row at highest latitude is so close to the geographical pole that all points in this row are situated nearly at the same location but spread in the full longitudinal width over the grid. This is an awkward case for the fast scan method because this requires a full scan in longitudinal direction in order to obtain identical results with a full scan. The longwinded description of the exact implementation of these additional techniques is omitted here.

Even if the additional dynamic block size method is omitted, the majority of the fast scan mappings yield identical results with the full scan method. However, including it appears to be very effective in obtaining identical results for the exceptional cases.

The fast scan method has one general restriction: real world geographical neighbour points must be neighbour points at the grid. Because by far the majority of geophysical applications satisfy this restriction, this might seem rather trivial. However, an example outside the scope of this work, is a long winded river system which is folded on a more compact rectangular grid. Nearby located grid points may represent distant river locations due to the folding, and therefore erroneously fall inside the scan block. In that kind of cases the fast scan method might fail, and the full scan method should be used. Unstructured grids can be another problematic case for the fast scan method, because their possible deviant way of storing the geographical neighbours on the grid. This neighbour ordering can become less straightforward in case in certain shaped areas the node density increases largely. The irregularity of the unstructured grid itself is not an issue for OBLIMAP.

In combination with the additional dynamic block size method, the fast scan method is guaranteed applicable in case the 'connection criterion' is fulfilled, i.e. geographical neighbours have to be connected via a sequence of direct neighbours at the grid. In that case the fast scan method even works for unstructured grids.



3.2 Masked mapping

Not seldomly field variables which have to be mapped are not everywhere available on the departure grid. In such cases it is important that masked points can be excluded for interpolation in order to guarantee the quality performance of the mapping at the borders of the mask. This applies for instance for field variables which are only available at an ice sheet mask, like the surface mass balance (SMB) and refreezing fields as produced by RACMO. In another example one might argue that the missing data at the exterior of the bottom topography field, might be a minor issue in a single mapping because of the remoteness of the problem to the area of interest. However, after several times of to and fro mapping the error will propagate into the area of interest.

With OBLIMAP's masked mapping, invalid masked points are ignored for the interpolation. If for a certain destination grid point the nearest projected departure grid point has an invalid mask, default this destination grid point will be set to that invalid value. However, this behavior can be changed for each mapped field individually by altering the mask criterion, as with the other option all valid contributions are considered irrespective if the nearest contribution has an invalid mask. For each mapped field a separate invalid value can be specified. The mask of each field for the masked mapping is shaped by the pattern of the invalid value for that field. Each mask is allowed to vary in time, and might vary per vertical layer in case the masking is based on a spatial 3D field. Default a mask of a certain field can be based on the invalid value pattern of that field itself, but the mask can also be based on the pattern of one of the other mapped fields. Because masked mapping is a post-scanning option it can be switched on and off at any time.

Excluding the masked area might allow a significant decrease of α , the angle which controls the standard parallel, which induces a further optimisation of the projection. In case of a masked mapping and a raised message level OBLIMAP will inform the user about an optimal masked α , estimated by:

$$\alpha = \arcsin \left(\frac{1}{R} \sqrt{\frac{1}{2\pi} (\text{COUNT}(\text{mask})) \Delta x \Delta y} \right) \quad (1)$$

with R the Earth radius, Δx and Δy the ISM grid spacing in the x and y direction and the 2D mask is 0 for invalid values and 1 otherwise. Eq. (1) equals Eq. (2.2) in Reerink et al. (2010) except that the total number of ISM nodes $N_x N_y$ is replaced by the number of ISM points which have a valid value, being $\text{COUNT}(\text{mask})$.

3.3 Embedded mapping

Once a certain grid combination has been scanned the scanning phase can be omitted if this mapping is repeated off-line or on-line afterwards. With OBLIMAP 2.0 the data of the SID file (for each destination grid point the indices of each contribution and its relative distance to the destination grid point) is stored in the dynamic data object (DDO). With that the off-line performance becomes faster for the multiple field, multiple layer and multiple record mapping. But more importantly, it solves the I/O bottleneck for high frequent on-line coupling.

The redesign and recoding of the OBLIMAP package enables the embedded calling of the same mapping routines which are used in the off-line stand-alone mode. Embedding the OBLIMAP mapping routines basically requires five code additions



- to the host model: 1. Adding the OBLIMAP initialization routine, which reads the OBLIMAP configuration variables, in the initialization of the host model. 2. Loading the DDO's for both mapping directions by reading both the SID files in the initialization of the host model. 3. Declaring a vector of spatial 3D fields where the vector length equals the number of mapped fields. 4. Calling the OBLIMAP mapping routines in both directions within the host model time loop. 5. Deallocate the DDO
- 5 in the finalizing stage of the host model.

The introduction of the DDO in combination with OBLIMAP's redesign solved the I/O bottleneck for high frequent coupling and enables the embedded calling of the mapping routines, which makes OBLIMAP suitable as GCM - ISM coupling software for on-line coupling projects.

3.4 Nearest point assignment

- 10 The 'nearest point assignment' is a post scan alternative to the quadrant and radius interpolation method. Instead of interpolating the nearby projected source points on the destination nodes, with this option each destination node obtains the field value of the nearest projected source point disregarding any other contributions. This method can be combined with masked mapping, in that case a destination node will be always invalid if the nearest projected point has an invalid mask. Regardless which interpolation method has been used during the scan phase, the 'nearest point assignment' can be used in the post scan phase. This option
- 15 can be considered in case both grids have about the same resolution. The 'nearest point assignment' performs faster than both interpolations methods in the post scan phase.

3.5 Vincenty method for distances on the ellipsoid

- An ISM to GCM mapping projects ISM grid points on the Earth ellipsoid. In that case the interpolation requires the distance over the curved surface for the Shepard distance weighting between each projected ISM point and the GCM point. The
- 20 geodesic, which is the shortest route between two points on the Earth's surface along the great circle, can locally be approximated accurately by the geodesic on the auxiliary Earth sphere. OBLIMAP uses this estimate default for the distances on the Earth ellipsoid, because this saves computational time in the scan phase and because the contributions are located close to each other compared to the Earth radius so the deviations will be small. However, OBLIMAP 2.0 provides the option to calculate the precise geodesics for the ellipsoid by Vincenty's method. Vincenty's inverse numerical approximation (Vincenty, 1975a) is
- 25 implemented in OBLIMAP.

3.6 Mapping multiple layers of spatial 3D fields

- In the most common OBLIMAP applications spatial 2D geographical fields are mapped for none, one or more time records. With OBLIMAP 2.0 it is also possible to map spatial 3D fields for none, one or more time records. The mapping of 3D fields is in fact limited to the mapping of several parallel vertical layers where each layer simply uses the same mapping. The layers are
- 30 assumed to be close to each other in comparison with the Earth radius, because for each layer the same Earth radius is applied.



Furthermore the horizontal grid distribution of all layers is assumed to be identical to the grid distribution of the top layer. And the fields are only interpolated in the horizontal directions for each vertical layer.

The new OBLIMAP 2.0 netcdf I/O routines automatically detect the spatial dimension of each input field and whether it contains the (unlimited) time dimension. The spatial 2D and 3D fields, including or excluding the time dimension, can be mapped simultaneously and in arbitrary order. With this combination a convenient way of mapping dimensionally different fields is achieved.

3.7 Automated selection of scan parameters

The radius method is recommended as soon as the destination grid resolution is four times larger than the departure grid resolution (Reerink et al., 2010). In all other cases the quadrant method is favored. Based on this criterion the interpolation method is default automatically selected in OBLIMAP 2.0 by checking the ratio of the grid resolutions of the grid centers. However it is also possible to select the interpolation method manually.

Similarly the optimal radius is default automatically determined for the radius method, if the radius method is selected. OBLIMAP checks whether the geographical grid has cyclic longitudinal grid borders, if so default the cyclic mode will be switched on automatically. Finally OBLIMAP can also determine automatically the optimal angle α which controls the standard parallel of the projection, but this is not default the case.

3.8 Data architecture, messaging and manual

OBLIMAP 2.0 stores all fields in a vector of fields with each field a spatial 3D field. In case the spatial dimension of a certain field is lower than 3D, its dimension is reduced when written to a netcdf. This applies also to symmetric dimensions, they will be default reduced if they are fully symmetric in one dimension. The field vector is reused and updated for successive records or time steps.

The geographical scientific climate database conventions are followed in the fully recoded I/O-netcdf interface, which has become highly flexible and which has been largely automatized. The package has been professionalized in the sense that for any user warning or error we endeavor to provide a meaningful message (see items 11-13 in Table 1). The embedding of the OBLIMAP mapping routines requires the addition of a minimum of code in the host model, for that reason programmer error messages are added in order to prevent improper software usage. The newly added OBLIMAP manual serves the user to correctly configure the mappings and OBLIMAP's options.



4 Performance and applications

4.1 Computational time performance

Sixty benchmark mapping experiments have been used, representing a diverse set of mappings which differ in number of nodes, grid resolution, mapping direction, interpolation method, location and thus also in projection. This relative arbitrary set of benchmarks has been used to evaluate the computational efficiency of the fast and full scan method.

The mapping experiments in Fig. 5 are subdivided by different symbol colours: ISM to GCM mappings are plotted red and purple for the quadrant and radius interpolation method respectively. GCM to ISM mappings are plotted blue and light blue for the quadrant and radius interpolation method respectively. The typical error in the time measurements is twice the size of the plot symbol in the figure, and is obtained by repeating the mapping many times.

The part of the Earth's surface which is covered by ice is limited in comparison with the entire ocean - atmosphere surface. Therefore the number of grid nodes which are involved in the mapping is significantly lower than the total number of GCM nodes. This is relevant with respect to the time performance plotting. Therefore the performance is plotted in Fig. 5 as function of N the number of participating destination grid points multiplied with the number of departure grid points. The data is plotted on a logarithmic scale for both axes in Fig. 5, which improves the visualization of the trends compared to the otherwise rather sparse clustered plotting.

The time spent in the full and fast scan routines are measured by including built-in Fortran time counting routines inside the OBLIMAP code, and are shown for this set of mappings respectively in Fig. 5a and Fig. 5b. Fig. 5c shows the gain factor (equal to the full scan time divided by the fast scan time) for these mappings if the fast scan method is used instead of the full scan method. Fig. 5d shows the total gain achieved per individual 2D field in case the post scan fast mapping is used instead of the full scan method. For the latter the fast mapping time is divided by the number of fields, layers and records because they might differ for each of the mappings.

Fig. 5a shows a strong increase of the computational time with N for the full scan method. Clearly visible are the separate branches for the different mapping directions though their trend is similar. In contrast no significant differences are caused due to the selected interpolation method. As can be seen from Fig. 5b the computational time for the fast scan method is much lower, and though it increases with N , it is important to observe that it levels off for higher N which is also reflected in Fig. 5c where the gain factor strongly increases with N . The scattering in Fig. 5b is larger than in Fig. 5a because the size of the local scan block is sensitive to the individual grid configurations. Because the fast scan times are relative fast, this causes a relative large influence. As a consequence the gain factor and the total gain factor are influenced by this scatter of the fast scan method measurements. The large values of the gain factor in Fig. 5c show that the fast scan method is orders of magnitude faster than the full scan method. Fig. 5d shows a strong increasing trend of the total gain factor up to $\sim 10^6$ at $N = 10^{12}$. The typical fast mapping time of one individual 2D field is about $\sim 10^{-2}$ seconds.



4.2 Masked and non-masked mapping applications

We start with some general remarks concerning the plotting of the pre and post mapped data displayed in Figs. 6-11. The high quality data itself is saved in netcdf files. In order to visualize those fields, python's matplotlib and its basemap extension are used to script the plotting. For plotting fields which are defined on grids which are based on geographical coordinates, a projection has to be specified with basemap.

The plotting interpolates the fields once projected by the plotting. Though the selection of the plotting projection can be independently and arbitrary chosen from the mapping projection, we used for most GCM field plots a plotting projection which is similar to the mapping projection. The fields on the ISM grid are plotted as true grid values, i.e. no plotting interpolation has been applied. In several subfigures one black or a few coloured contours are plotted on top of the data. In several subfigures the ETOPO dataset (Amante and Eakins) is used as background for the masked mapped areas or for a surrounding area with a constant value. Usually this concerns the ocean basin and remote areas. Plotting the data on the ETOPO dataset also serves as a check because all coastline contours should coincide.

The availability of the high resolution topographic and atmospheric forcing datasets for the Greenlandic and Antarctic ice caps opens in combination with the OBLIMAP 2.0 package the opportunity to model the major Earth ice caps or parts of them with an optimal centered projection. These applications are shown here as mapping examples. Additional examples show the local mapping of the Antarctic Peninsula and the mapping of a geothermal heat flux forcing dataset. Table 2 lists the mapping parameters for each mapping on the three ISM grids in the various mapping examples.

In case of remapping a dataset, the coordinates are projected twice, however the fields are only interpolated at the final destination grid in order to minimize the mapping error. The first projection leaves the field data unaffected, only the (x,y)-coordinates are converted to (longitude, latitude)-coordinates. This task can be conducted by the oblimap convert program which is part of the OBLIMAP-package.

4.2.1 Topographic fields for Greenland

The publicly available topographic dataset for Greenland (Bamber et al., 2013) contains the surface topography, the bedrock topography and the ice thickness for the present day situation and is projected by (Bamber et al., 2013) with a polar SG projection on an ISM grid with a 1x1 km resolution. In a post-processing phase some manual corrections are applied to this 1x1 km ISM grid (Bamber et al., 2013).

In order to obtain an optimal centered projection this dataset has been remapped by first applying the inverse polar projection on all coordinates of this data set (see Fig. 6a and Fig. 6c) and thereafter this result is mapped with an optimal centered oblique SG projection on an ISM grid with a 5x5 km resolution (see Fig. 6c and Fig. 6d).

No masked mapping is used for the mapping of the surface topography on the ISM grid in Fig. 6b because it concerns a field which literally levels off to sealevel (the zero level). Because the bedrock topography contains missing values (see the white bottom corner areas in Fig. 6c), a masked mapping is used for this field. Resulting in a properly mapped mask border, as can be seen at the bottom right corner of Fig. 6d.



Note that in fact it would be possible to directly map the irregular spaced measured data points, but then the manual applied corrections are lacking.

4.2.2 Atmospheric forcing fields for Greenland

The RACMO2 atmospheric forcing dataset for Greenland (van Angelen et al., 2014) containing the surface massbalance (SMB), the surface air temperature, the surface refreezing, runoff and other fields, is defined on a gaussian-reduced grid with an approximate horizontal resolution of about 11 km (see Fig. 7a and Fig. 7c). These fields are mapped with the same projection on a grid with the same 5x5 km resolution and extent as used in Sect. 4.2.1. The SMB field in Fig. 7a only contains valid values for ice covered grid points, and is therefore masked mapped on the ISM grid (see Fig. 7b). The same applies for the refreezing in Fig. 7c which is also masked mapped on the ISM grid (see Fig. 7d).

10 4.2.3 Topographic fields for Antarctica

The publicly available Bedmap2 topographic dataset for Antarctica (Fretwell et al., 2013) contains the surface topography, the bedrock topography and the ice thickness for the present day situation and is projected by (Fretwell et al., 2013) with a polar SG projection on an ISM grid with a 1x1 km resolution.

In order to obtain a dataset which can be mapped on any (local) grid with OBLIMAP, the inverse polar projection is applied on all coordinates of this data set (see Fig. 8a and Fig. 8c). Thereafter this result is mapped with the same polar SG projection on an ISM grid with a 20x20 km resolution (see Fig. 8c and Fig. 8d). Though this might seem superfluous, the advantage is that from this longitude - latitude based dataset a grid of any grid extent and resolution can be created with OBLIMAP. In addition any optimum centered local grid can be created as well from this dataset, like for instance the local area of the Antarctic Peninsula (see Sect. 4.2.5). We choose the common polar SG projection for entire Antarctica, however given the position of the Antarctic continent a slightly oblique projection might in fact yield the optimal projection.

No masked mapping is used for the mapping of the surface topography on the ISM grid in Fig. 8b because it concerns a field which literally levels off to sealevel (the zero level). Because the bedrock topography contains missing values (see the white corner areas in Fig. 8c), a masked mapping is used for this field. Resulting in a properly mapped mask border, as can be seen at the corners of Fig. 8d.

25 4.2.4 Atmospheric forcing fields for Antarctica

The RACMO2 atmospheric forcing dataset for Antarctica (Van Wessem et al., 2014) containing the SMB, the surface air temperature, the surface refreezing, runoff and other fields, is defined on a gaussian-reduced grid with an approximate horizontal resolution of about 27 km (see Fig. 9a and Fig. 9c). These fields are mapped with the same projection on a grid with the same 20x20 km resolution and extent as used in Sect. 4.2.3. The SMB field in Fig. 9a only contains valid values for ice covered grid points, and is therefore masked mapped on the ISM grid (see Fig. 9b). The same applies for the refreezing field in Fig. 9c which is also masked mapped on the ISM grid (see Fig. 9d)



The refreezing field in Fig. 9c is an example of a source field with an inadequately defined value for the missing data which equals zero in this case. This zero value is inconvenient here because the field values themselves reach zero in the interior of Antarctica. Coincidentally this implies that a masked mapping based on an invalid value which is taken equal to zero, will affect the mapping of the zero contour inside the interior of Antarctica as well. Despite this error the masked mapping is still to be preferred over the non masked mapping in this case because the latter one yields large errors at grid points in the vicinity of the coastline. Here we circumvent this problem by using the ice-cover mask, which is co-distributed as part of the RACMO2 dataset. In this case the ice-cover mask is used as the mask for the masked mapping of the refreezing, which illustrates the flexibility of the masked mapping options. Preferably the invalid value for the missing values of the source fields have a value outside the range of the actual field values in order to avoid this problem.

10 4.2.5 Local mapping of the Antarctic Peninsula

In contrast to the polar projection used in Sects. (4.2.3-4.2.4), a local mapping of the Antarctic Peninsula demonstrates an oblique projection example with the same datasets. The Bedmap2 surface topography (Fig. 8a) and the Bedmap2 bedrock topography (Fig. 8c) have been mapped on a local 5x5 km ISM grid for the Antarctic Peninsula with an optimal centered oblique projection without using a mask (see Fig. 10a and Fig. 10b). The RACMO2 SMB (Fig. 9a) and the RACMO2 refreezing (Fig. 9c) for the year 2006 have been masked mapped on the same local 5x5 km ISM grid with the same optimal centered oblique projection (see Fig. 10a and Fig. 10b).

4.2.6 Mapping the geothermal heat flux

The spatial variable geothermal heat flux (Shapiro and Ritzwoller, 2004) which is defined on a global regular $1^\circ \times 1^\circ$ longitude - latitude grid (see Fig. 11a) for the present day situation, is another forcing dataset which can be mapped on the same grid. In Fig. 11b this geothermal heat flux has been mapped on the same 5x5 km ISM grid with the same optimal centered projection for Greenland as in Sect. 4.2.1. In Fig. 11c the geothermal heat flux has been mapped on the same 20x20 km ISM grid with the same optimal centered projection for Antarctica as in Sect. 4.2.3. Finally in Fig. 11d the geothermal heat flux has been mapped on the same 5x5 km ISM grid with the same optimal centered projection for the Antarctic Peninsula as in Sect. 4.2.5.

4.3 Coupling and embedding

25 An experiment in which a GCM is on-line coupled with an ISM consists of much more than the technical coupling task. Perfectly mapped fields may require a successive downscaling step. Several decisions have to be made concerning issues like to which degree the on-line coupling will be conducted, which fields are available for a particular GCM - ISM combination and can be effectively used, which model time steps and coupling time step will be used and does that require certain time averaging prior to each coupling step, and for which fields only the perturbations will be used in the coupling. Examining the results should learn whether the coupling is numerically stable, if the feedback mechanisms do properly work and to which



extent the resolution differences limit the coupling of the models. Presenting coupled results requires the evaluation of these issues and the description of the used GCM and ISM, but that is far beyond the scope of this paper.

Instead we only shortly report that we coupled the IMAU-ICE model with CLIMBER-2 (Petoukhov et al., 2000) by using OBLIMAP's mapping routines embedded. We benefit from CLIMBER's Fortran77 implementation in which all relevant variables for the coupling are globally defined. This makes it possible to embed CLIMBER in the ice sheet model without major recoding of CLIMBER. Therefore in this case the ice sheet model is taken as the host model and the OBLIMAP mapping routines are embedded in the ice sheet model as well. The low CLIMBER resolution certainly limits the coupling degree, but appeared to be suitable for practical and technical learning purposes, because the IMAU-ICE - CLIMBER coupled model technically easily operates on a laptop. In general however, due to its complexity, it is preferable to take GCM as the host model.

5 Discussion

OBLIMAP's full scan method is robust and suitable for any GCM - ISM grid combination regardless of the irregular spacing of the grids and their arbitrary ordering including unstructured grids. It is used for benchmarking, but can also be used in special cases like the embedment of a flow line model or the embedment of a very low resolution model like CLIMBER. The withdrawal of the full scan method is its slow performance. This becomes a serious bottleneck for larger grid combinations, i.e. $N \gg 10^{12}$.

The fast scan method is orders of magnitude faster especially for larger grid combinations, but the involved grid ordering has to suffice one of the following rather general conditions (see also Sect. 3.1.2): 1. Geographical neighbour points must be neighbours on the grid (except for cyclic longitudinal borders which are also allowed). 2. The neighbour connectivity criterion. The first condition is the most severe restriction as it will often not hold for unstructured grids, though for many geographical grids it is a rather general and trivial condition. The weaker second condition requires that two geographical neighbour points are connected via a sequence of direct neighbours. In that case the fast scan method works if the dynamic block size is switched on, even for unstructured grids which satisfy the second condition.

For large datasets it is important to note that if suddenly a tremendous slow down is encountered with increasing N , it is likely that the size of the processor memory is the bottleneck. In that case it is recommended to switch to a platform with a larger processor memory.

In fact there are various fast scan method variants available in OBLIMAP 2.0, and although the fastest ones work in the most situations, there is a trade off between the best performance and robustness. With robustness we mean that the fast scan method yields in any situation identical results as the full scan method. The differences in performance between the various fast scan methods is limited. In order to avoid that the fast scan method requires expert knowledge to configure it, OBLIMAP 2.0 comes with an option which determines the fast scan parameters automatically. This option is default switched on.

The situation in which a regular longitude - latitude grid is defined up to the pole in such way that for instance near or at 90° North N_{LON} points coincide with almost the same location (while the field values are not identical), requires the addition



of a full longitude search in the vicinity of the poles as discussed in Sect. 3.1.2 in order to obtain identical results with the full scan method. For the sake of guaranteed robustness with the automatic fast scan mode, the full longitude scan is applied for the near polar area of any grid with $|\text{latitude}| \geq 87^\circ$. For this kind of grids themselves this is not a major issue, their scanning is usually fast enough. But this also applies to other high resolution grids, leading to a performance decrease. Automatic detection, whether a situation requires a full longitude scan, would be an advantage in a future OBLIMAP release, avoiding unnecessary performance loss.

In case the quadrant interpolation method is used in combination with the full scan method, sometimes a rather remote contribution is detected in one of the quadrants in comparison with the detected contributions in the other quadrants, because the quadrant interpolation method continues searching for the nearest contribution until it is found (Reerink et al., 2010). This is usually due to missing data or because it concerns a contribution originating from the vicinity of the departure grid border. The remote contribution has a very limited influence on the interpolation due to the distance weighting. However, with the fast scan method the limited block size deselects this remote bycatch and accordingly the results are not identical but deviate insignificantly. In fact the fast scan result is even favourable in this case.

We conclude that the default automatic fast scan method is robust, indispensable for large grid combinations (i.e. for large N) and can be safely used by non-expert users.

OBLIMAP is able to map between models which differ largely in resolution. But in particular if the destination grid is much coarser than the departure grid, the computational time inevitably increases for both the fast scan method and the fast mapping due to the large radius in the radius interpolation method.

Various examples for different resolution combinations show that the masked mapping works well, i.e. artefacts are absent in the vicinity of the mask borders and the fields represent realistic values as shown in Figs. 6-10. The masked mapping method is indispensable for products like the SMB, the refreezing and the runoff, because their values differ strongly along the mask border. In case fields like the surface topography are mapped frequently to and fro, masked mapping prevents the propagation of artefacts into the domain of interest. Besides, these figures show the high quality of the masked and non-masked fields with high resolution.

If the available datasets which might differ in grid extent, grid resolution and grid type are based on longitude - latitude coordinates or as soon they have been converted to that, OBLIMAP is a powerful tool to map them all on an equal arbitrary ISM grid by the same optimal centered projection, and vice versa. Although 'optimal' and 'centered' are the recommended preferences, neither of them is a prerequisite. This is demonstrated for the topographic, atmospheric forcing and geothermal heatflux datasets by the applications in Sect. 4. Similar, the ice surface velocity dataset for Greenland (Rignot and Mouginot, 2012) could also be mapped on the ISM grid for Greenland.

The redesign and recoding of OBLIMAP in combination with the DDO introduction ensure that the embeddable OBLIMAP 2.0 mapping routines are suited to bridge the technical task of on-line coupling in both directions. As indicated in Sect. 3.3, five code items have to be added in the host model in order to embed the OBLIMAP mapping routines. Preferably the ISM is embeddable, because usually it will be beneficial if the GCM is the host model candidate due to its complexity. For this reason we plan to separate the initialization phase and the time loop for IMAU-ICE. At the same time, an embeddable ISM allows



the simultaneous embedment of multiple ISM domains like e.g. Greenland and Antarctica, each with its own projection and configuration file. In addition it enables the simultaneous simulation of several ice cap domains by the ISM in case they are mutual connected by the sealevel evolution, similar with the approach of De Boer et al. (2014).

OBLIMAP is dedicated to the GCM - ISM mapping and coupling task in both mapping directions and developed from that perspective. However, OBLIMAP might be very well applicable to other geophysical mapping problems because the included mapping components like the projection, the regridding and the fast scan method are based on a generic implementation.

Although the subjects which are treated by OBLIMAP for example rather resemble those of a frequently used GCM component coupler like OASIS, it is remarkable how much the implementations differ (cf. Valcke, 2013). One obvious difference is the additional projection step in OBLIMAP. Further, OBLIMAP's interpolation methods are able to cope with the irregular spaced locations where the projected points end up after they have been projected on the destination surface, which make them at the same time suitable for irregular spaced departure and destination grids. Moreover the OBLIMAP interpolation methods have to be robust for relative large resolution differences. OBLIMAP 2.0 comes with the fast scan method in order to reduce the search time. The same objective has been addressed in OASIS4 (Redler et al., 2010) in a parallel approach, though Valcke (2013) reports that its development is not pursued.

Adaptive grids require continuous scanning, each time one of the grids changes the scan has to be repeated. With the current sequential OBLIMAP code only on-line coupling of low resolution adaptive grid models is feasible using the fast scan method. Embedding the mapping routines, including repeated scanning, changes the list of code addition as indicated in Sect. 3.3 as follows: The second item is replaced by an allocation statement inside the time loop in the host model, with a successive call to the embeddable scan routine for each mapping direction. The fifth item, the deallocation of the DDO needs to move inside the time loop at its end. However a parallel implementation will be beneficial for on-line coupling of adaptive grids and will extend the possibilities.

On platforms which have at least the same number of processors available as the number of participating destination grid points, a possible new efficient and straightforward parallel implementation of OBLIMAP's full scan method can be developed. In that case the expected time for the parallel full scan method will be reduced and is expected to be faster than the sequential fast scan method. This yields two advantages: 1. The full scan method becomes fast, c.q. the scanning of any kind of grid including unstructured grids becomes fast. 2. On-line coupling of adaptive grids becomes feasible as long N is not too large but remains limited to ordinary applications.

An alternative hybrid approach combines the fast scan method with a parallel approach. As explained in Sect. 3.1.2 the fast scan method slows down if situations are encountered in which none of the previous scanned neighbours is a participating point at the row jump or if no contributions were detected. An effective hybrid approach therefore could use one processor per row, starting with a full scan at the first participating point of each row. Whereafter each processor continues the scan of its row with the fast scan method. This hybrid approach requires a much smaller number of processors than the first proposed parallel approach, and is expected to be almost as fast.



We expect that both proposed parallel implementations scale well because the slave processors do not need any mutual communication during their task. They only need access to the DDO and their results have to be collected by the master processor afterwards.

The hybrid approach could serve as the starting point for heavier applications. An extra performance gain could be realized by using more processors to carry out each full scan at the start of each row. The full scan over the departure grid points is then split up between several processors, this however is a slightly more complex parallel task. Because after a first scan the pivots off all points are known, in another combinational approach for adaptive grids the pivots of the previous mapping could be used for a local block search if they are remembered and updated. This works under the assumption that the adaptive grid changes are locally smooth. It could potentially realize a large performance gain for both the parallel and the sequential implementation.

6 Conclusions

With a significant extended range of applications than its prior version, OBLIMAP 2.0 has become much faster and easier to configure. A manual has been added to precisely describe OBLIMAP's user options. The increase in performance, the ability to map much larger grid combinations, the extensive and flexible way of masked mapping, and the fact that the OBLIMAP mapping routines can be used embedded in an on-line high frequent coupled application, are among the main achievements of OBLIMAP 2.0 while OBLIMAP's high accuracy and robustness is maintained.

The power of OBLIMAP is its ability to map various kind of datasets which can largely differ in grid resolution and grid extent, by an optimal centered projection on one destination grid of arbitrary configuration. This potential has been demonstrated by a few relevant examples in which topographic, atmospheric forcing and geothermal heat flux fields from various datasets have been mapped on grids for the major ice caps.

The fast and fully sequential OBLIMAP 2.0 package is lightweight and suitable to run on a laptop. Future implementation of one of the discussed parallel approaches offers the possibility of an additional performance gain.

Code availability

The OBLIMAP 2.0 code and the OBLIMAP manual are available as supplementary material (<http://www.geosci-model-dev.net/???-supplement.zip>) and are distributed under the terms of the GNU General Public License.

Acknowledgements. We thank Melchior van Wessem and Jan van Angelen for providing the RACMO2 data of Antarctica and Greenland. We thank Bas de Boer, Michiel Helsen and Sarah Bradley for their user feedback over the years on the OBLIMAP package. This project has been funded by Kennis voor Klimaat (KvK) and by a grant from the Netherlands Earth System Science Centre (NESSC) from the netherlands organization for scientific research (NWO). The RACMO datasets could be provided due to a NESSC-WP3 collaboration.



References

- Amante, C. and Eakins, B.: ETOPO1 1 Arc-Minute Global Relief Model: Procedures, Data Sources and Analysis. NOAA Technical Memorandum NESDIS NGDC-24. National Geophysical Data Center, NOAA, 2009, <http://www.ngdc.noaa.gov/mgg/global/global.html>, doi:10.7289/V5C8276M.
- 5 Bamber, J. L., Griggs, J. A., Hurkmans, R. T. W. L., Dowdeswell, J. A., Gogineni, S. P., Howat, I., Mouginot, J., Paden, J., Palmer, S., Rignot, E., and Steinhage, D.: A new bed elevation dataset for Greenland, *The Cryosphere*, 7, 499–510, doi:10.5194/tc-7-499-2013, <http://www.the-cryosphere.net/7/499/2013/>, 2013.
- De Boer, B., van de Wal, R. S. W., Lourens, L. J., Bintanja, R., and Reerink, T. J.: A continuous simulation of global ice volume over the past 1 million years with 3-D ice-sheet models, *Climate Dynamics*, 41, 1365–1384, 2013.
- 10 De Boer, B., Stocchi, P., and Van De Wal, R. S. W.: A fully coupled 3-D ice-sheet-sea-level model: Algorithm and applications, *Geoscientific Model Development*, 7, 2141–2156, 2014.
- De Boer, B., Dolan, A. M., Bernales, J., Gasson, E., Goelzer, H., Golledge, N. R., Sutter, J., Huybrechts, P., Lohmann, G., Rogozhina, I., Abe-Ouchi, A., Saito, F., and Van De Wal, R. S. W.: Simulating the Antarctic ice sheet in the late-Pliocene warm period: PLISMIP-ANT, an ice-sheet model intercomparison project, *Cryosphere*, 9, 881–903, 2015.
- 15 Fretwell, P., Pritchard, H. D., Vaughan, D. G., Bamber, J. L., Barrand, N. E., Bell, R., Bianchi, C., Bingham, R. G., Blankenship, D. D., Casassa, G., Catania, G., Callens, D., Conway, H., Cook, A. J., Corr, H. F. J., Damaske, D., Damm, V., Ferraccioli, F., Forsberg, R., Fujita, S., Gim, Y., Gogineni, P., Griggs, J. A., Hindmarsh, R. C. A., Holmlund, P., Holt, J. W., Jacobel, R. W., Jenkins, A., Jokat, W., Jordan, T., King, E. C., Kohler, J., Krabill, W., Riger-Kusk, M., Langley, K. A., Leitchenkov, G., Leuschen, C., Luyendyk, B. P., Matsuoka, K., Mouginot, J., Nitsche, F. O., Nogi, Y., Nost, O. A., Popov, S. V., Rignot, E., Rippin, D. M., Rivera, A., Roberts, J., Ross, N., Siegert, M. J.,
- 20 Smith, A. M., Steinhage, D., Studinger, M., Sun, B., Tinto, B. K., Welch, B. C., Wilson, D., Young, D. A., Xiangbin, C., and Zirizzotti, A.: Bedmap2: improved ice bed, surface and thickness datasets for Antarctica, *The Cryosphere*, 7, 375–393, doi:10.5194/tc-7-375-2013, <http://www.the-cryosphere.net/7/375/2013/>, 2013.
- Helsen, M. M., van de Wal, R. S. W., van den Broeke, M. R., van de Berg, W. J., and Oerlemans, J.: Coupling of climate models and ice sheet models by surface mass balance gradients: application to the Greenland Ice Sheet, *The Cryosphere*, 6, 255–272, doi:10.5194/tc-6-255-2012, <http://www.the-cryosphere.net/6/255/2012/>, 2012.
- 25 Helsen, M. M., van de Berg, W. J., van de Wal, R. S. W., van den Broeke, M. R., and Oerlemans, J.: Coupled regional climate–ice-sheet simulation shows limited Greenland ice loss during the Eemian, *Climate of the Past*, 9, 1773–1788, doi:10.5194/cp-9-1773-2013, <http://www.clim-past.net/9/1773/2013/>, 2013.
- Holland, D. and Jenkins, A.: Modeling thermodynamic ice-ocean interactions at the base of an ice shelf, *Journal of Physical Oceanography*, 29, 1787–1800, 1999.
- 30 Noël, B., van de Berg, W. J., van Meijgaard, E., Kuipers Munneke, P., van de Wal, R. S. W., and van den Broeke, M. R.: Evaluation of the updated regional climate model RACMO2.3: summer snowfall impact on the Greenland Ice Sheet, *The Cryosphere*, 9, 1831–1844, doi:10.5194/tc-9-1831-2015, <http://www.the-cryosphere.net/9/1831/2015/>, 2015.
- Petoukhov, V., Ganopolski, A., Brovkin, V., Claussen, M., Eliseev, A., Kubatzki, C., and Rahmstorf, S.: CLIMBER-2: A climate system model of intermediate complexity. Part I: Model description and performance for present climate, *Climate Dynamics*, 16, 1–17, 2000.
- 35 Redler, R., Valcke, S., and Ritzdorf, H.: OASIS4 - a coupling software for next generation earth system modelling, *Geoscientific Model Development*, 3, 87–104, doi:10.5194/gmd-3-87-2010, <http://www.geosci-model-dev.net/3/87/2010/>, 2010.



- Reerink, T. J., Kliphuis, M. A., and van de Wal, R. S. W.: Mapping technique of climate fields between GCM's and ice models, *Geoscientific Model Development*, 3, 13–41, doi:10.5194/gmd-3-13-2010, <http://www.geosci-model-dev.net/3/13/2010/>, 2010.
- Rignot, E. and Mouginot, J.: Ice flow in Greenland for the International Polar Year 2008-2009, *Geophysical Research Letters*, 39, doi:10.1029/2012GL051634, 2012.
- 5 Shapiro, N. and Ritzwoller, M.: Inferring surface heat flux distributions guided by a global seismic model: Particular application to Antarctica, *Earth and Planetary Science Letters*, 223, 213–224, doi:10.1016/j.epsl.2004.04.011, 2004.
- Shepard, D.: A two-dimensional interpolation function for irregularly-spaced data, *Proceedings-1968 ACM National Conference*, pp. 517–524, 1968.
- Valcke, S.: The OASIS3 coupler: a European climate modelling community software, *Geoscientific Model Development*, 6, 373–388, doi:10.5194/gmd-6-373-2013, <http://www.geosci-model-dev.net/6/373/2013/>, 2013.
- 10 van Angelen, J., van den Broeke, M., Wouters, B., and Lenaerts, J.: Contemporary (1960-2012) Evolution of the Climate and Surface Mass Balance of the Greenland Ice Sheet, *Surveys in Geophysics*, 35, 1155–1174, doi:10.1007/s10712-013-9261-z, 2014.
- van Meijgaard, E., van Uft, L. H., Van de Berg, W. J., Bosveld, F. C., Van den Hurk, B. J. J. M., Lenderink, G., and Siebesma, A. P.: The KNMI regional atmospheric climate model RACMO version 2.1, *Tech. Rep. 302*, KNMI, P.O. box 201, 3730 AE, De Bilt, the Netherlands,
- 15 2009.
- Van Wessem, J., Reijmer, C., Morlighem, M., Mouginot, J., Rignot, E., Medley, B., Joughin, I., Wouters, B., Depoorter, M., Bamber, J., Lenaerts, J., Van De Berg, W., Van Den Broeke, M., and Van Meijgaard, E.: Improved representation of East Antarctic surface mass balance in a regional atmospheric climate model, *Journal of Glaciology*, 60, 761–770, doi:10.3189/2014JoG14J051, 2014.
- Van Wessem, J., Ligtenberg, S., Reijmer, C., Van De Berg, W., Van Den Broeke, M., Barrand, N., Thomas, E., Turner, J., Wuite, J., Scambos, T., and Van Meijgaard, E.: The modelled surface mass balance of the Antarctic Peninsula at 5.5 km horizontal resolution, *The Cryosphere*, 20 10, 271–285, doi:10.5194/tc-10-271-2016, <http://www.the-cryosphere.net/10/271/2016/>, 2016.
- Vincenty, T.: Direct and inverse solutions of geodesics on the ellipsoid with application of nested equations., *Surv Rev*, pp. 88–93 [addendum: *Surv Rev* 23(180):294 (1976)], http://www.ngs.noaa.gov/PUBS_LIB/inverse.pdf, 1975a.



Table 1. Shortlist of the new OBLIMAP 2.0 features and changes. The stars indicate the categories of the new features. For the full list and more details see the supplementary OBLIMAP manual. The terms in bold are used for short reference to the described feature.

Feature number	Most important	Extended functionality	User comfort	Description of the new OBLIMAP 2.0 features
1	★	★		The mapping of large grid combinations has become feasible with the fast scanning option, as a result of the dramatic performance improvement in the scanning phase.
2	★	★		With masked mapping points which are part of an invalid mask will not contribute to the interpolation. The invalid mask covers the area for which the field values of a certain user specified field matches with a certain user specified invalid value.
3	★	★		With multiple masked mapping each field can use its own mask. This mask is allowed to vary in time, and might vary per vertical layer in case the masking is based on a spatial 3D field.
4		★	★	The redesign of OBLIMAP allows embedded calling of the OBLIMAP mapping routines.
5	★	★		The introduction of a dynamic data object avoids superfluous reading of the scanned file. This improved the post scanning phase performance, in particular in combination with multiple record, multiple field and multiple layer (3D) mapping.
6	★	★		On-line coupling of an ISM with a GCM is possible by using the OBLIMAP routines embedded. Prior off-line scanning is recommended for both mapping directions, so the embedded OBLIMAP mapping routines can use the superfast post scan mapping relying on the dynamic data object.
7		★		Nearest point assignment , a post scan alternative to the quadrant and radius interpolation method. I.e. each destination node obtains the field value of the nearest projected point, which implies that no interpolation is required. This option can be considered in case both grids have about the same resolution.
8	★	★		Multiple field , multiple record and multiple layer mapping . The 2D and 3D fields can be mapped simultaneously, while each field might also contain the unlimited time dimension. In the embedded case the fields will be mapped at each coupling time step.
9		★		The precise calculation with the Vincenty method of the great distance over the ellipsoidal arc is added as an option.
10		★	★	With the automatic scan option the scan parameters are determined by OBLIMAP itself.
11			★	Automatic OBLIMAP advice concerning the correct and optimal settings of the scan parameters.
12			★	Extended OBLIMAP messaging including four levels of message intensity.
13			★	Automatic dimension shape determination while reading the netcdf input files.
14			★	Separate configuration files have to be used for each mapping direction.
15		★		The ISM grid is allowed to be irregularly spaced.
16			★	A manual accompanies OBLIMAP 2.0, which has been added to guide on user level.

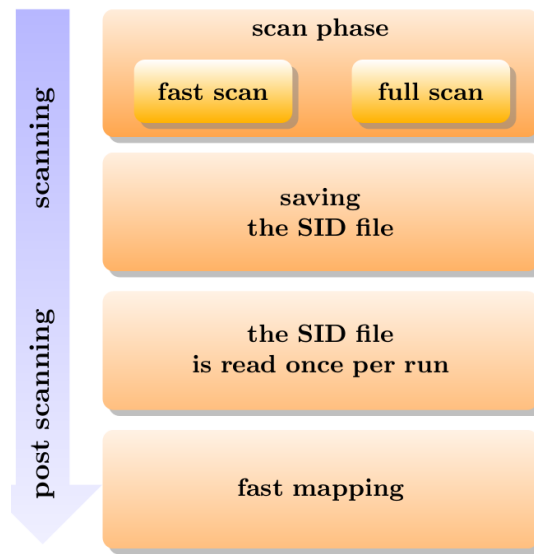


Figure 1. Schematic representation of the OBLIMAP 2.0 stand-alone version. The scanning consists of the scan phase and results in a SID file. The post scanning consists of reading the SID file and loading its content in the DDO, whereafter the fast mapping of multiple fields, layers and records can be repeated as often as required.

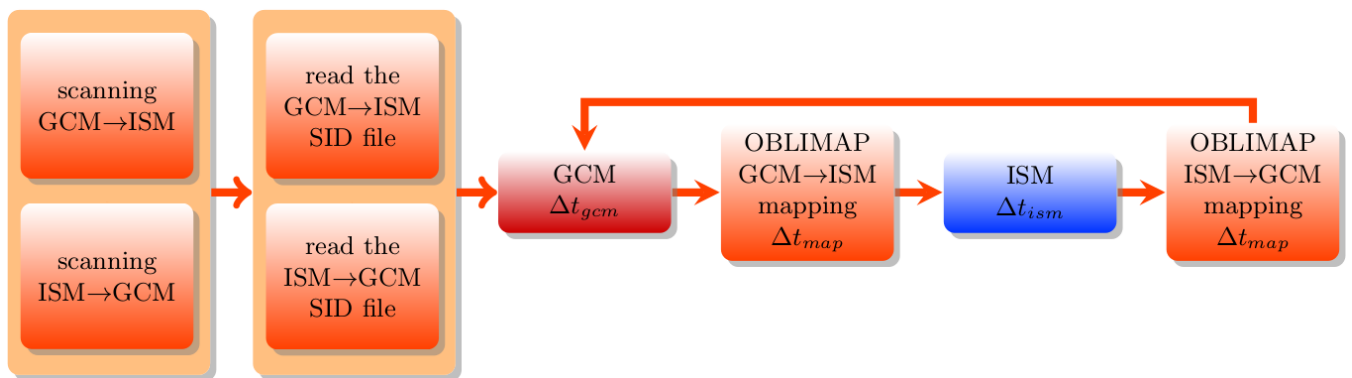


Figure 2. Schematic representation of an embedded ISM within a GCM run. At the initialization of the on-line coupled run the SID files are read, which have been created by off-line scanning prior to this run. The GCM and ISM are coupled with the embedded OBLIMAP routines at each coupling interval Δt_{map} . The GCM and the ISM evolve with their own time step Δt_{gcm} and Δt_{im} depending on their specific numerical stability criteria.



Table 2. The grid sizes N_x , N_y and the grid resolution Δ of the three ISM grids which result from the mapping applications in Sect. 4.2, using the oblique SG projection for the WGS84 ellipsoid with the projection parameters α and the coordinates of the projection center (λ_M, ϕ_M) .

ISM grid	N_x	N_y	Δ (km)	α ($^\circ$)	λ_M ($^\circ$)	ϕ_M ($^\circ$)
Greenland	301	551	5	7.1	319.0	72.0
Antarctica	281	281	20	19.0	0.0	-90.0
Peninsula	271	351	5	5.54	293.5	-70.2

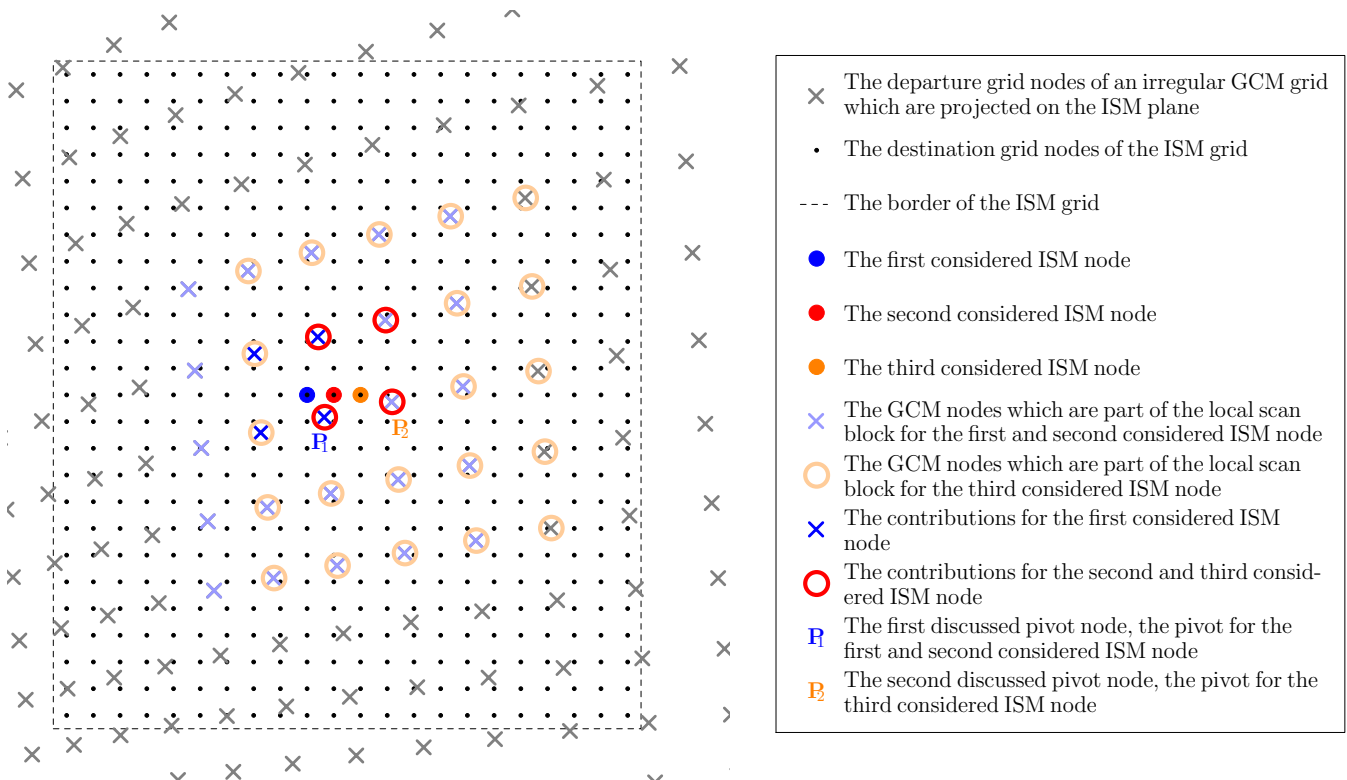


Figure 3. Situation sketch of the procedure to find a next scan block for a GCM to ISM mapping, supposed that the quadrant interpolation method is used. Consider the situation for which the four contributions (X) are found for the first considered ISM node (●). In order to find the contributions for the second considered ISM node (●), a scan over a local block is conducted. This local block is constructed by taking the nearest contribution (the pivot) and extending the block in all directions by b . In this case P_1 is the pivot and the block is extended by two X nodes in each direction. For the second considered ISM node this results in a different set of contributions (○). Because P_1 is the nearest contribution for the second considered ISM node as well, P_1 stays the pivot. Therefore the scan block remains unchanged in the next scan. Though for the third considered ISM node (●) the contributions are the same as for the second considered ISM node, now P_2 becomes the pivot of the third considered ISM node. So the next scan is conducted over the ○-marked GCM nodes.




Figure 4. Situation sketch of the procedure to find a next scan block for an ISM to GCM mapping, supposed that the radius interpolation method is used. Consider the situation for which the contributions (the ISM nodes in the blue shaded circle) are found for the first considered GCM node (×). In order to find the contributions for the second considered GCM node (×), a scan over a local block is conducted. This local block is constructed by taking the nearest contribution (the pivot) and extending the block in all directions by b (indicated by the arrows in the figure). In this case the pivot is the ISM node under the ×-marker and the ●-marked block is extended as indicated by the blue arrow. The contributions for the second considered GCM node lay in the red shaded circle. The pivot shifts to the GCM node under the ×-marker. The scan for the third considered GCM node (×) is conducted over the ○-marked ISM nodes in order to find the orange shaded contributions. Thereafter the pivot shifts again to the ISM node under the ×-marker for the next scan.

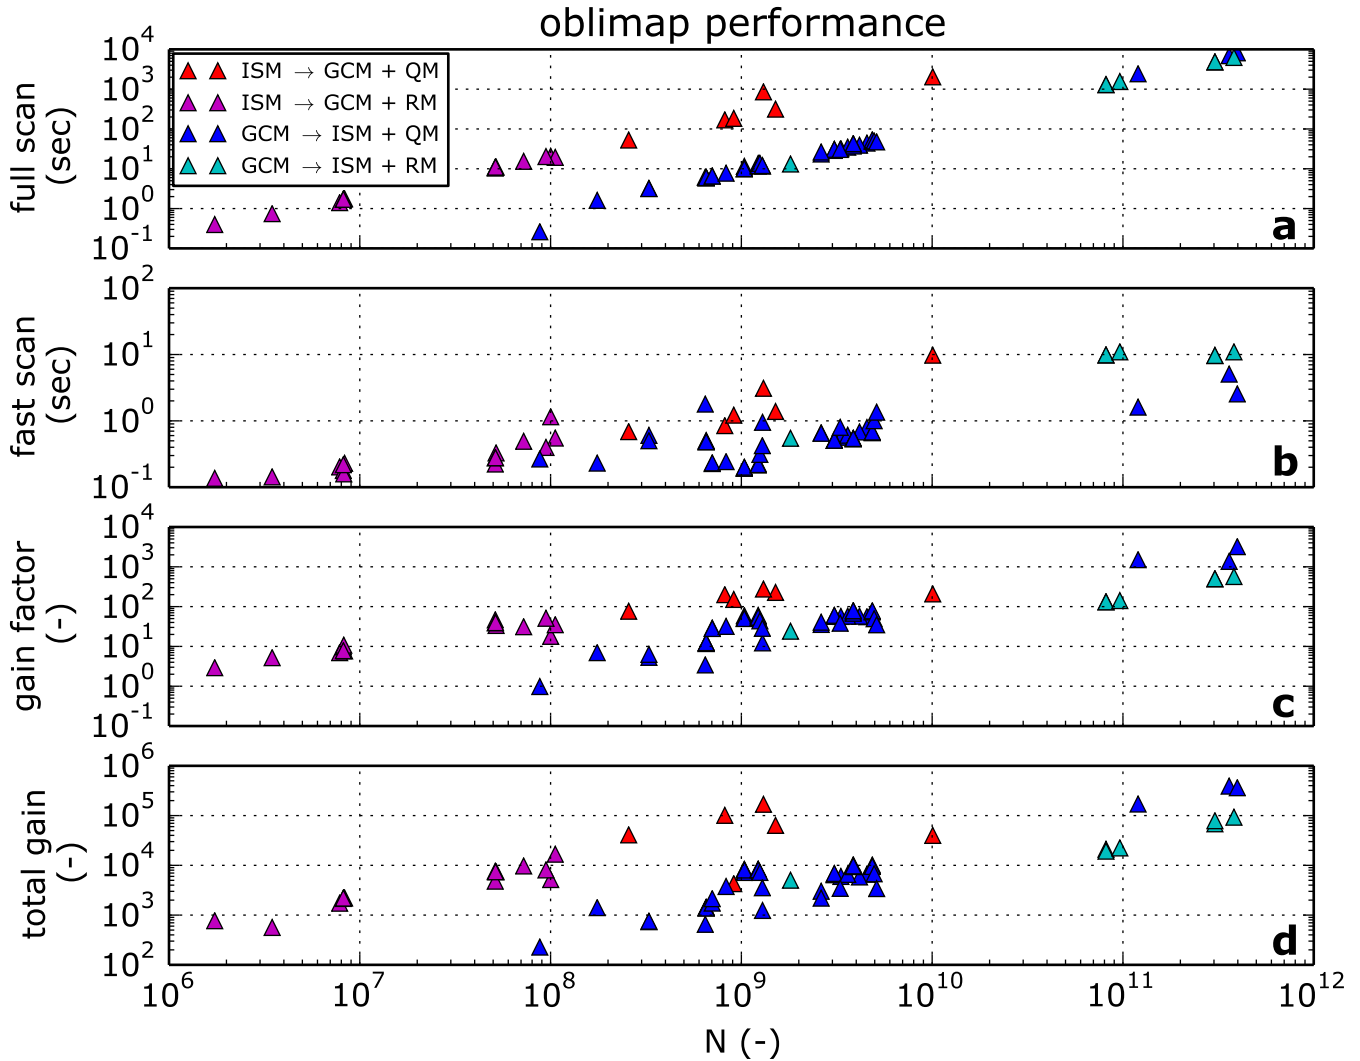


Figure 5. Sixty mapping experiments representing a diverse set of mappings which differ in number of nodes, grid resolution, mapping direction, interpolation method, location and thus also in projection, have been used to compare the time performance of the full scan method (a) and the fast scan method (b) as function of N , the number of participating destination grid points multiplied with the number of departure grid points. Subfigures c and d show the gain factor if respectively the fast scan method or the fast mapping per individual 2D field is used instead of the full scan method. Four colours distinguish between experiments which differ in mapping direction and which use either the quadrant interpolation method (QM) or the radius interpolation method (RM).

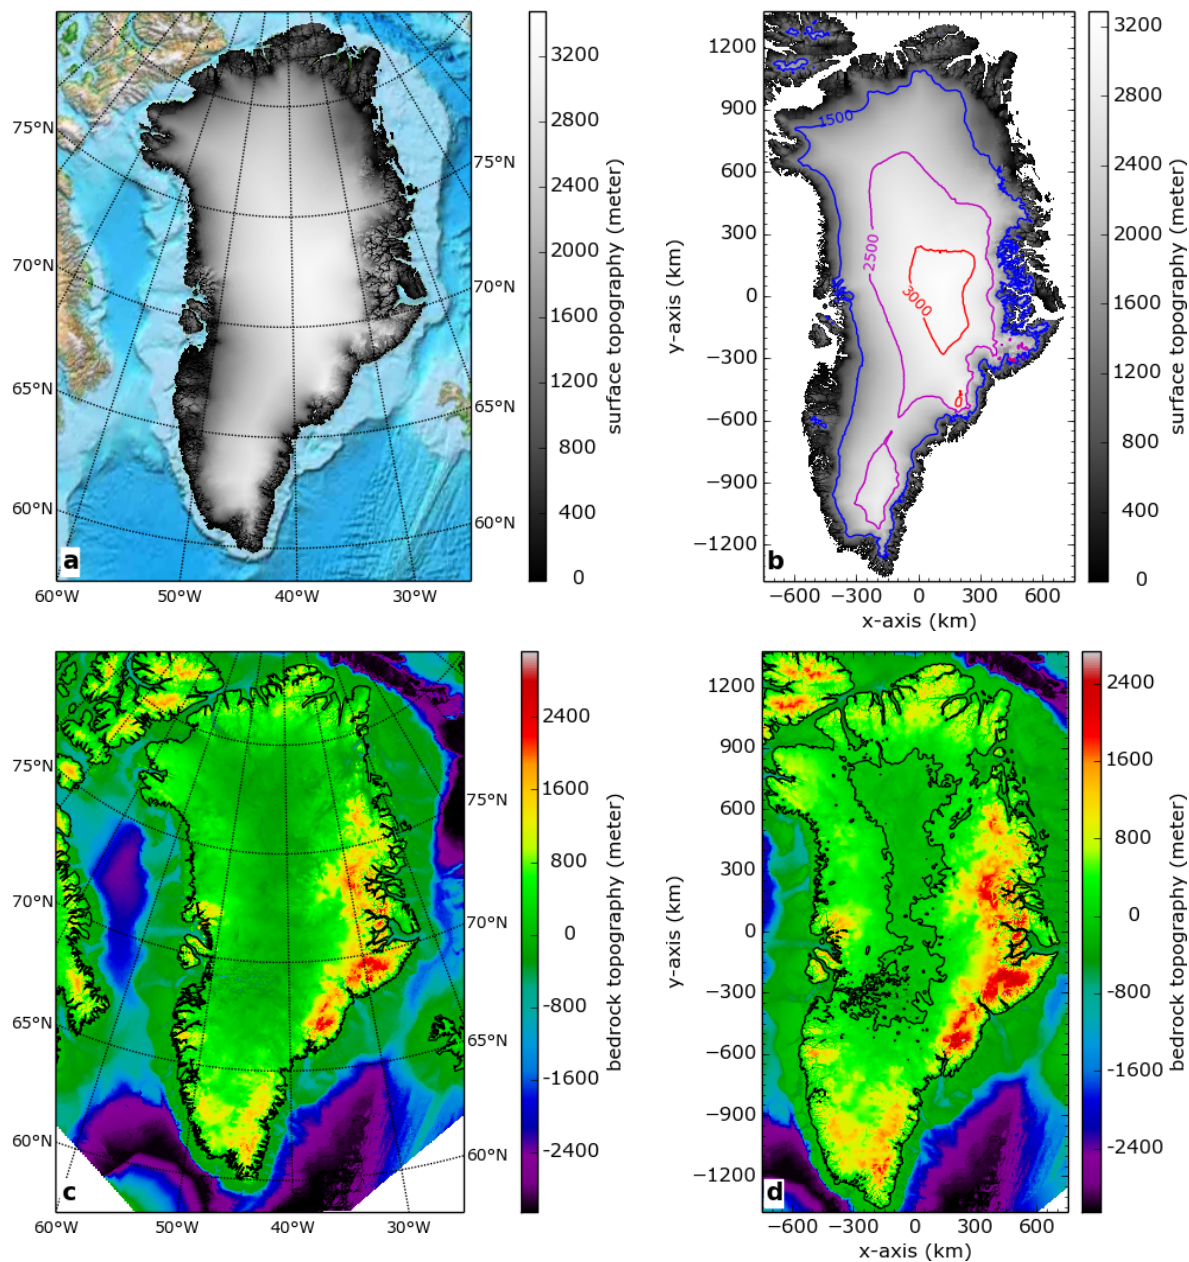


Figure 6. Subfigure **a** and **c** show respectively the surface topography and the bedrock topography of the 1x1 km resolution dataset for Greenland after an inverse projection on the Earth WGS84 ellipsoid. In **a** a coloured ETOPO background replaces the part at sealevel. In **c** the missing values are white coloured and the continental contour is plotted black. In **b** the surface topography has been mapped on a 5x5 km ISM grid with an optimal centered projection without using a mask, but points at sealevel have been plotted white in order to visualize the coastline contours. In **d** the bedrock topography has been masked mapped on a 5x5 km ISM grid with an optimal centered projection, the bottom right corner shows the proper resulting mask border. In **b** a few coloured contours are plotted on top of the data, and in **d** the zero contour is plotted black.

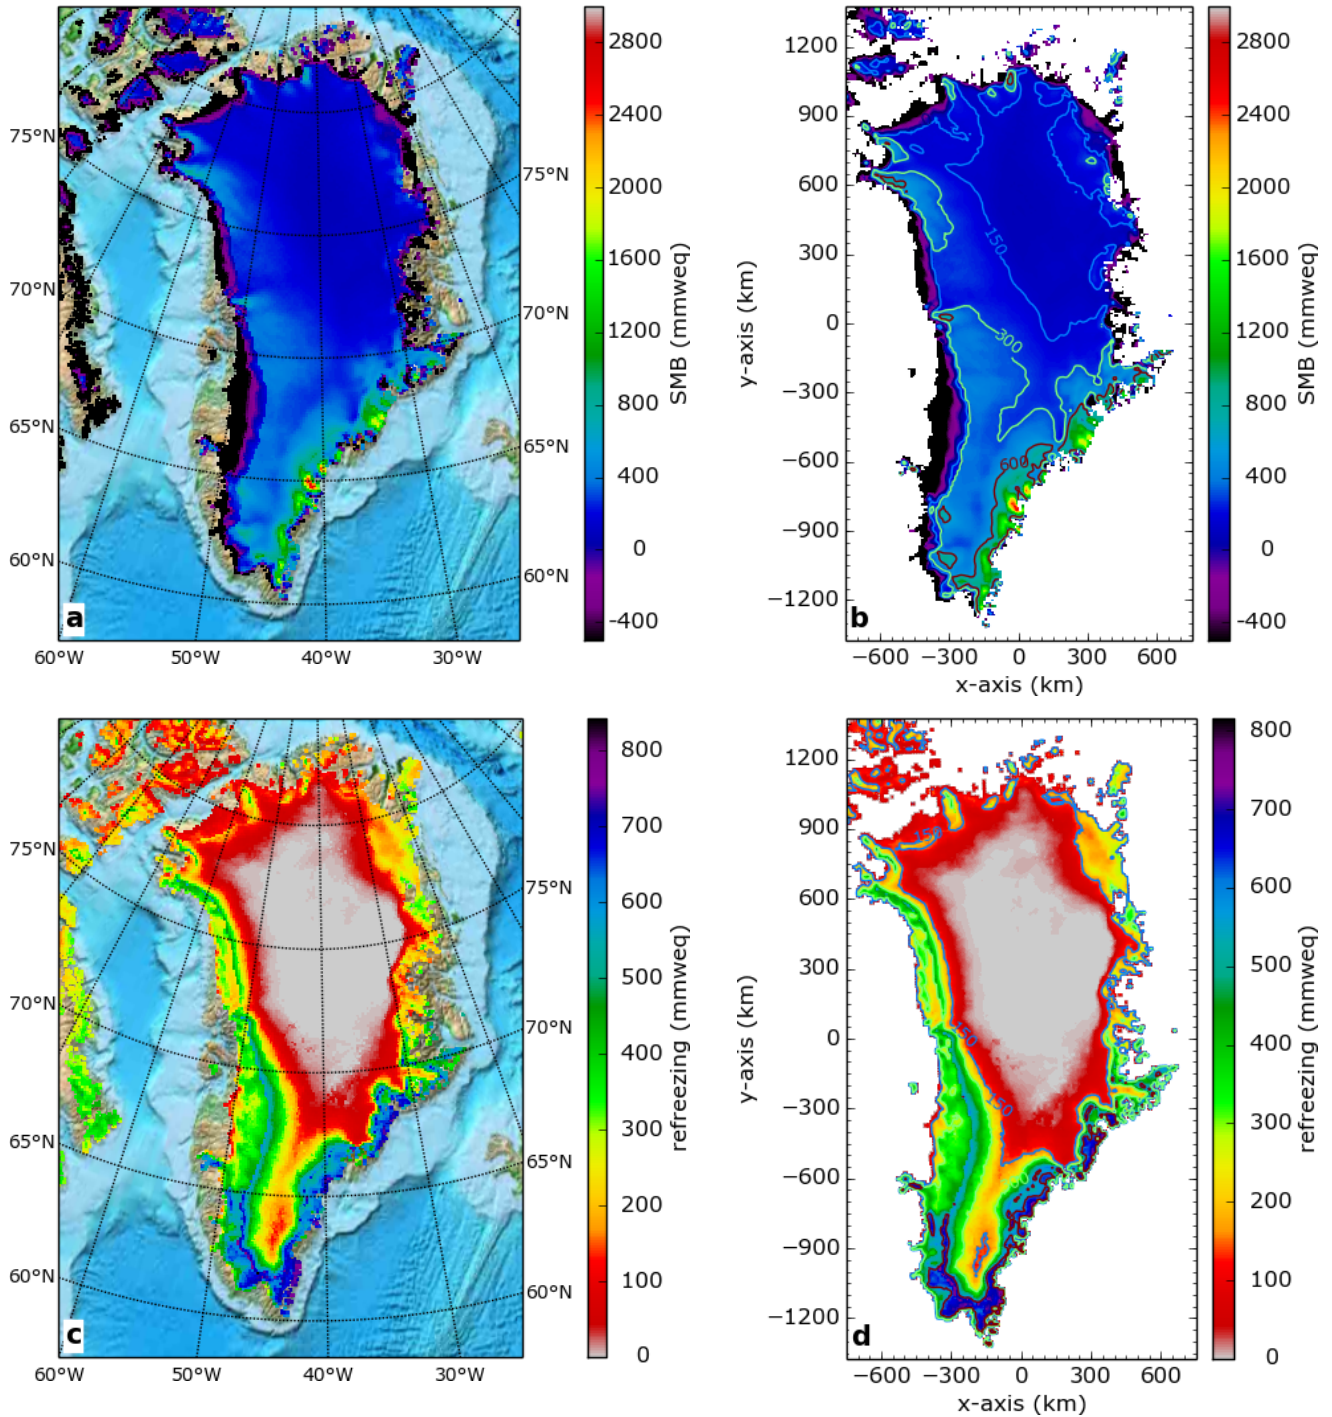


Figure 7. Subfigure **a** and **c** show respectively the surface massbalance (SMB) and the refreezing for the year 2006 of the RACMO2 dataset for Greenland on a gaussian-reduced grid with an horizontal resolution of about 11 km, the coloured ETOPO background replaces the masked area. In **b** and **d** the SMB and the refreezing have been masked mapped on a 5x5 km ISM grid with an optimal centered projection. Both fields are expressed in millimeter water equivalent (mmweq).

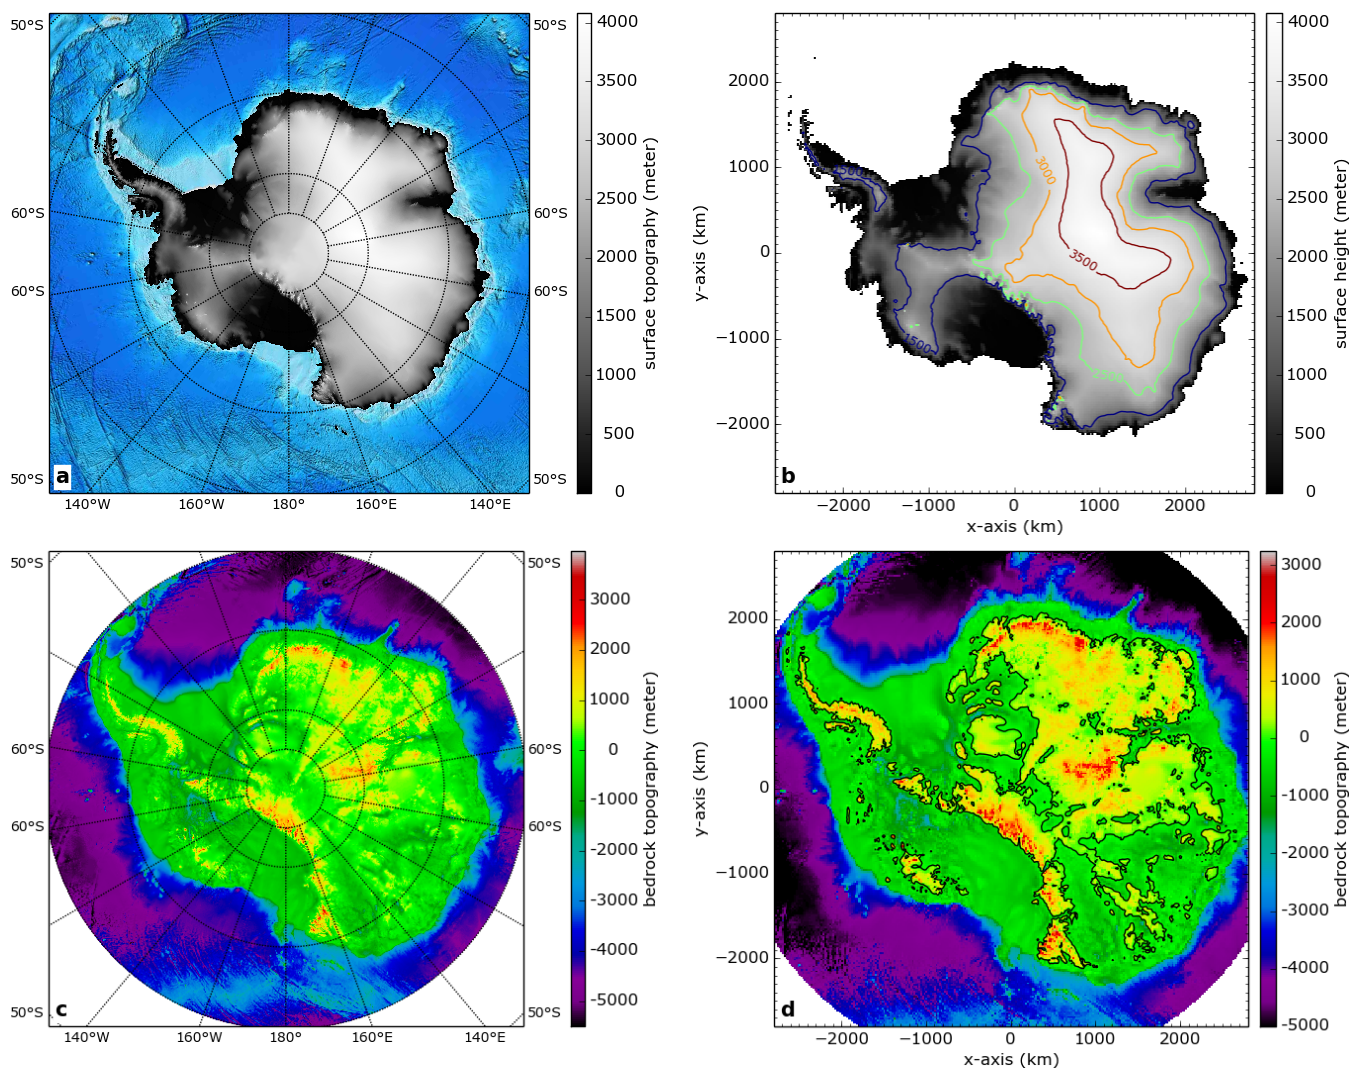


Figure 8. As Fig. 6 but for Antarctica, the Antarctic ISM grid resolution in **b** and **d** is 20x20 km.

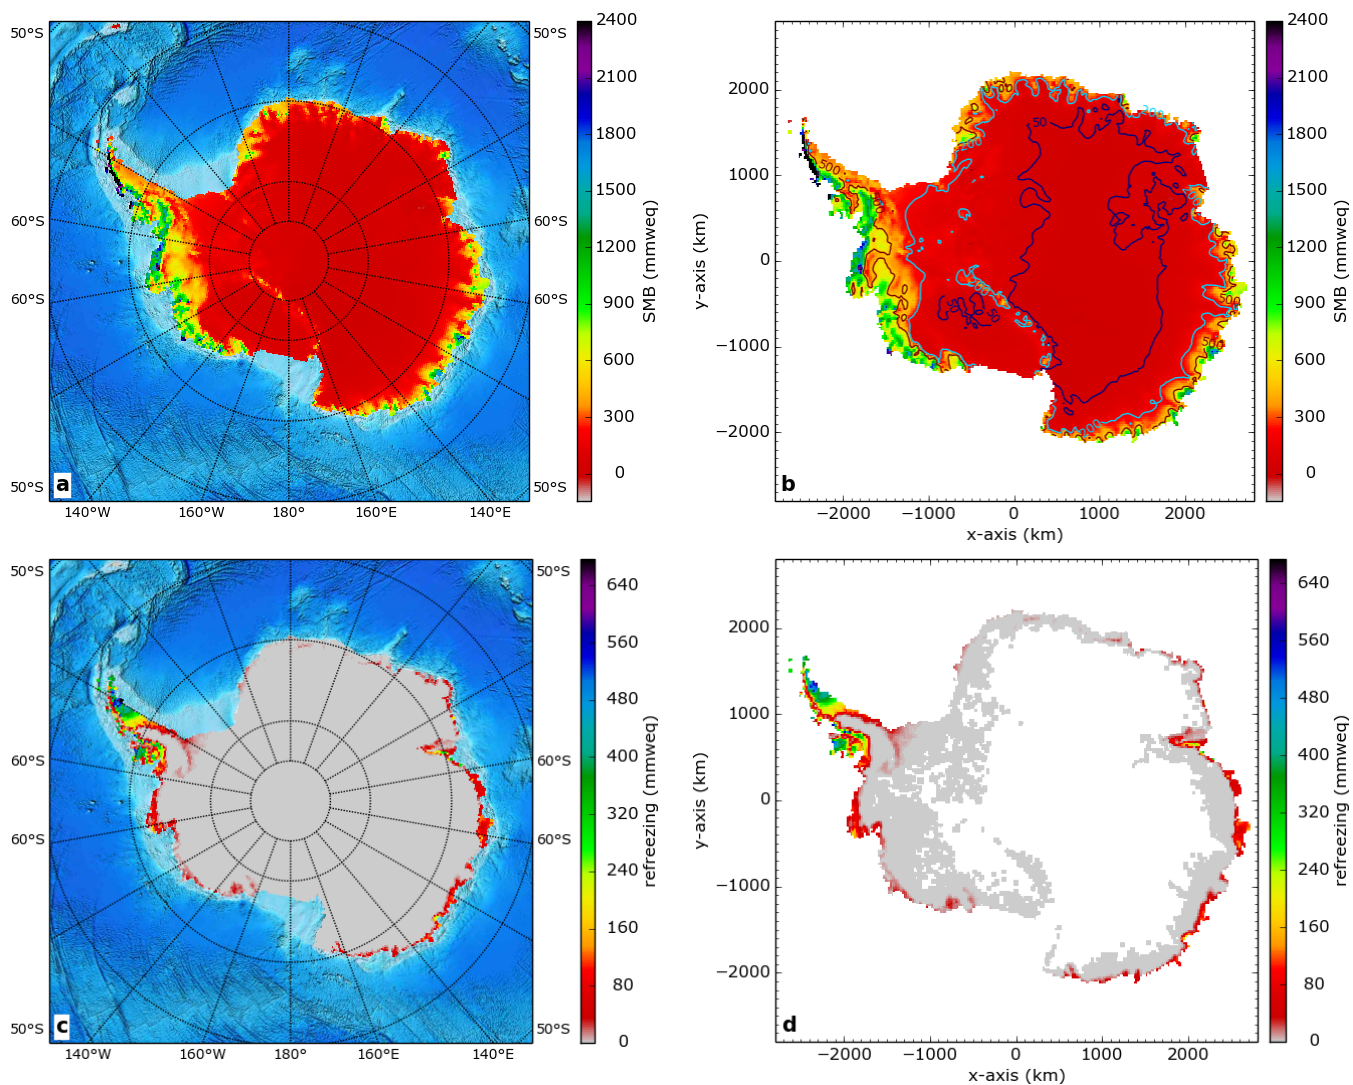


Figure 9. As Fig. 7 but for Antarctica, the Antarctic RACMO2 grid resolution in **a** and **c** is about 27 km and the Antarctic ISM grid resolution in **b** and **d** is 20x20 km.

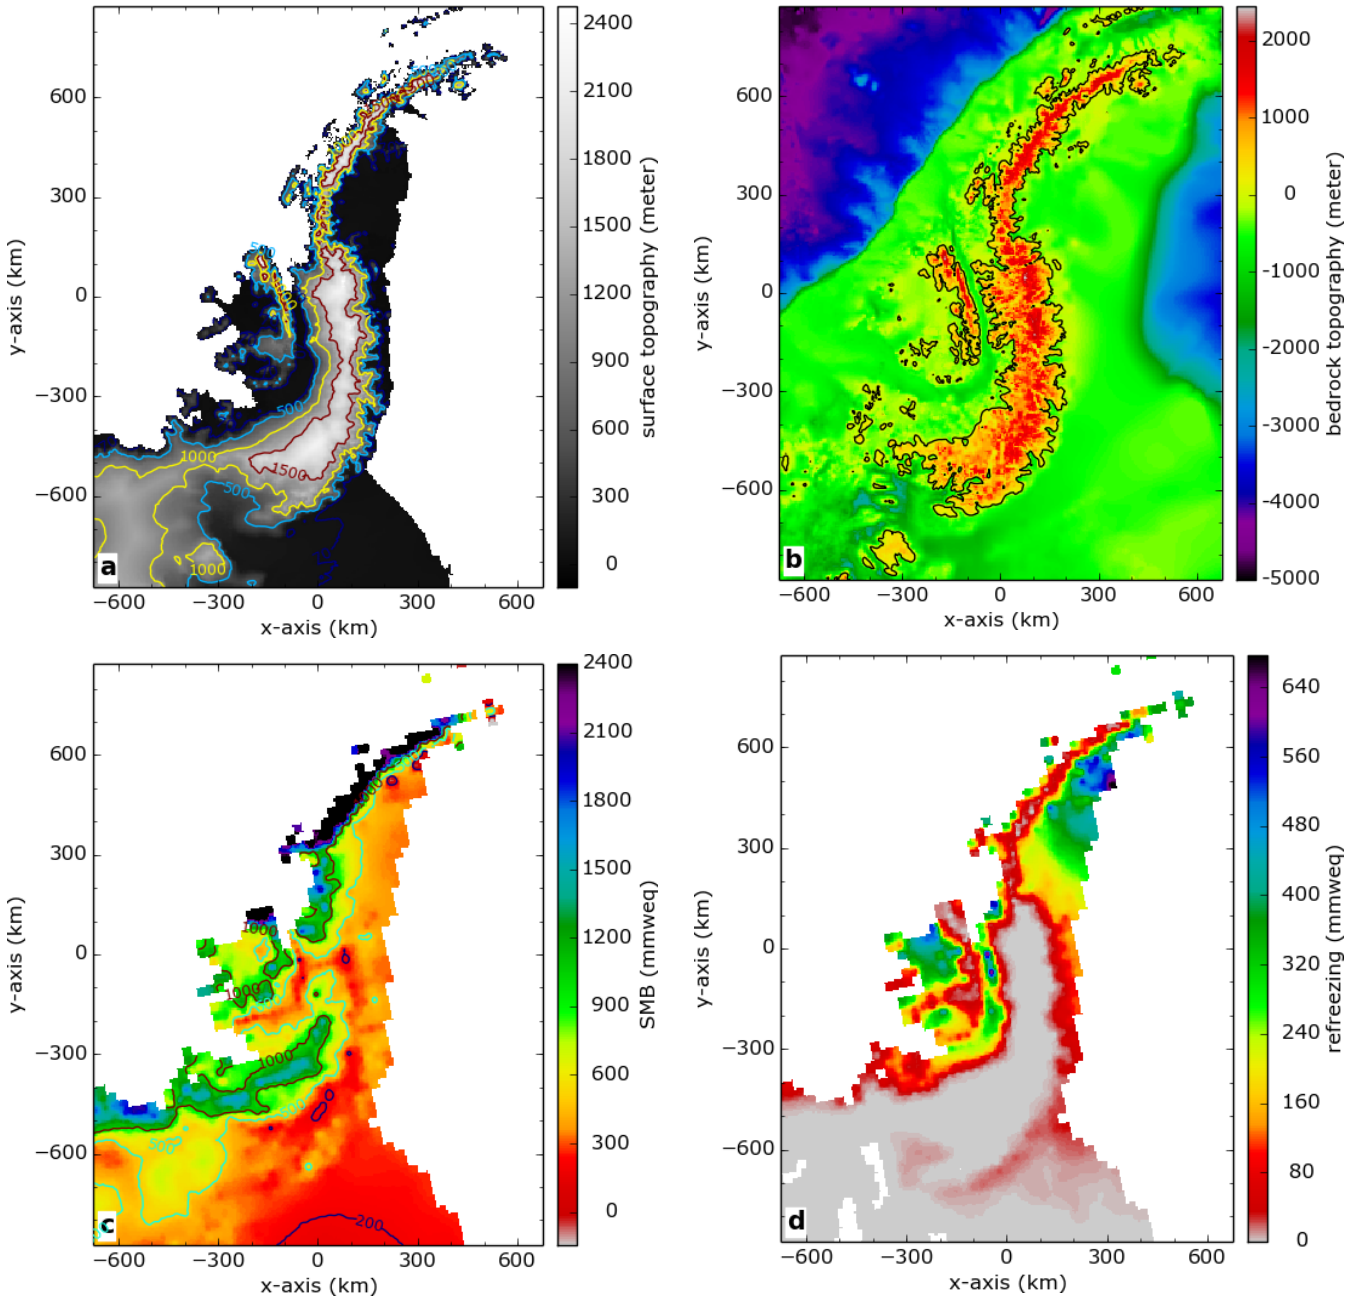


Figure 10. The Bedmap2 surface topography (Fig. 8a) and the bedrock topography (Fig. 8c) have been mapped on a local 5x5 km ISM grid for the Antarctic Peninsula with an optimal centered oblique projection without using a mask (see **a** and **b**). In **a** points at sealevel have been plotted white in order to visualize the coastline contours. The RACMO2 SMB (Fig. 9a) and the refreezing (Fig. 9c) for the year 2006 have been masked mapped on the same local 5x5 km ISM grid with the same optimal centered projection (see **c** and **d**).

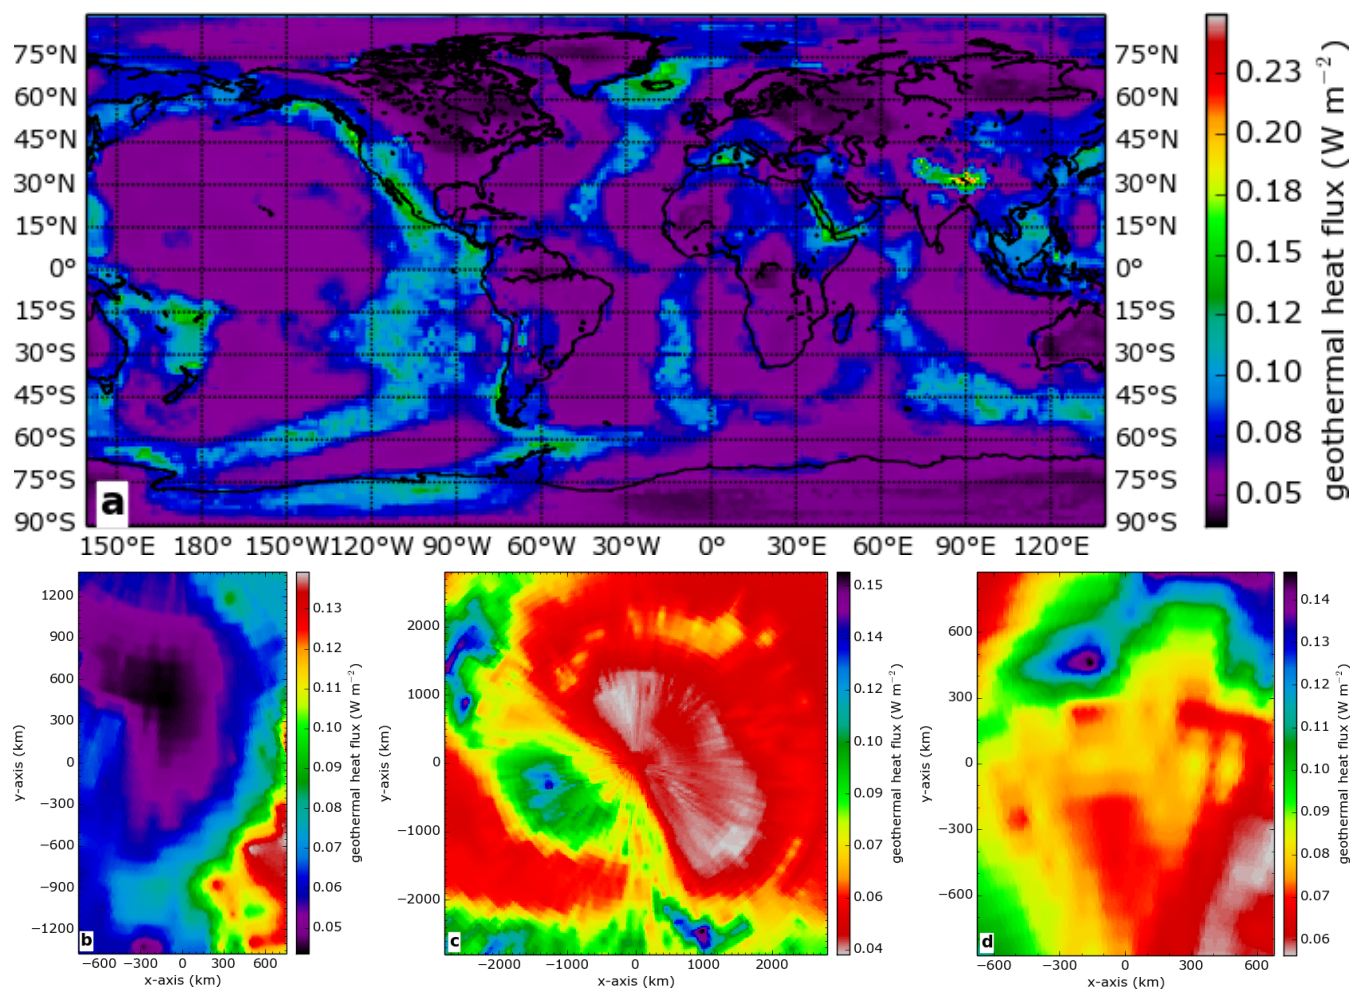


Figure 11. Subfigure **a** shows the geothermal heat flux (Shapiro and Ritzwoller, 2004) on a global regular $1^\circ \times 1^\circ$ longitude - latitude grid, the continental contours are plotted black. This geothermal heat flux has been mapped with optimal centered projections on a 5×5 km ISM grid for Greenland (**b**), a 20×20 km ISM grid for Antarctica (**c**) and a 5×5 km ISM grid for the Antarctic Peninsula (**d**).