



PERGAMON

Computers & Geosciences 28 (2002) 723–734

**COMPUTERS &
GEOSCIENCES**

www.elsevier.com/locate/cageo

Parallel numerical modelling of the Antarctic Ice Sheet

Andrea Takeda^a, Simon Cox^a, Antony J. Payne^{b,*}

^a *Department of Electronics and Computer Science, University of Southampton, UK*

^b *Bristol Glaciology Centre, School of Geographical Sciences, University of Bristol, Bristol B58 155, UK*

Received 13 February 2001; received in revised form 23 October 2001; accepted 10 November 2001

Abstract

The Antarctic Ice Sheet comprises the West Antarctic Ice Sheet and the much larger East Antarctic Ice Sheet. Fast flowing ice streams and outlet glaciers are important dynamic components of the ice sheet system, and a grid resolution of at least 20 km is required to identify many of these areas. Previous fine resolution numerical models have focussed on ice flow in West Antarctica or on fine resolution modelling of subsections of the ice sheet, since the size of East Antarctica has generally precluded studies of the whole ice sheet at a resolution adequate to identify complex flow features.

The equations describing ice flow are highly non-linear, making this a computationally intensive problem. We use a staggered grid for calculation of ice diffusivity to overcome numerical instability, and a sparse packing scheme to take account of the irregular boundary of Antarctica. We have developed an efficient parallel temperature-dependent ice flow model of the entire grounded portion of the Antarctic Ice Sheet at a resolution of 20 km. The model was primarily written to run on a commodity cluster of workstations, and performance results for this and other systems are presented. Ice flow patterns at steady state compare well with recently published balance velocity calculations. © 2002 Elsevier Science Ltd. All rights reserved.

Keywords: Antarctica; Parallel computing; Non-linear partial differential equations; Commodity supercomputing; BEDMAP data

1. Introduction

Antarctica is the largest ice mass on Earth, and has a strong influence over global climate and ocean circulation. The ice sheet consists of the West Antarctic Ice Sheet (WAIS) and the much larger eastern part (EAIS).

Fine resolution numerical models have been developed to study ice stream behaviour in West Antarctica (Payne, 1998, 1999). Ice streams are regions of relatively fast flow, and are thought to be a key feature affecting the stability (or instability) of the whole of the WAIS. The fast flowing ice streams are governed by different physics than the surrounding ice sheet. Although a standard ice flow model of the type described here cannot simulate this complex ice stream behaviour, it is

nevertheless able to locate the general patterns of faster flow due to the fine resolution grid employed. Recent remote sensing data and balance velocity calculations identified areas of faster flow stretching for many kilometres inland in East Antarctica (Rémy et al., 1999; Bamber et al., 2000). Fine resolution modelling of the whole of Antarctica, rather than just the West Antarctic Ice Sheet, is therefore important.

The system is governed by a pair of highly non-linear coupled partial differential equations. The model described uses the shallow ice approximation, but the dependence of diffusivity on ice thickness to the power of 5 and on surface slope squared makes this a computationally intensive problem. This complexity, and the vast size of the EAIS, has restricted previous modelling to the study of the WAIS at a high resolution, the whole of Antarctica at a resolution too low to resolve complex flow features, or to fine resolution modelling of subsections of the ice sheet. For example,

*Corresponding author. Tel.: +44-117-954-5972; fax: +44-117-928-7878.

E-mail address: a.j.payne@bristol.ac.uk (A.J. Payne).

Calov et al. (1998) and Savvin et al. (2000) modelled Dronning Maud Land at a resolution of 10.9 km, nested within a 109 km resolution Antarctic model. Here we take a different approach, modelling the whole of Antarctica at 20 km resolution using an efficient parallel thermally coupled ice flow model.

The numerical model includes the assumption that ice deformation is driven solely by shear stresses in the $x-z$ and $y-z$ planes. This assumption is valid for models with a grid resolution of 20 km or more, but at finer resolutions stresses in the $x-y$ plane must also be included. The present model can therefore not be used with a grid resolution finer than 20 km. The main motivation of this work was to increase the resolution of a model for the whole of Antarctica, and the focus of this paper is the parallel implementation of the ice flow model. It should be noted that the model was simplified to a certain extent to facilitate the parallelization process. The results presented reflect the simple nature of the model rather than a sophisticated study of ice mechanics, but nevertheless show some improvement between the 40 and 20 km grid resolution results and some interesting patterns of basal melt. The strategy used for parallel processing revolves around the solution of second order parabolic equations using a preconditioned conjugate gradient method. The same techniques could be used in the parallel implementation of higher order models.

2. Numerical model

The model is time (t) dependent and has horizontal dimensions x and y , and a vertical dimension z . We now briefly summarise the main equations; further details can be found in Payne and Dongelmans (1997). The time-dependent change in ice thickness H is calculated as

$$\frac{\partial H}{\partial t} = b + \frac{\partial}{\partial x} \left(D \frac{\partial s}{\partial x} \right) + \frac{\partial}{\partial y} \left(D \frac{\partial s}{\partial y} \right), \quad (1)$$

where s is surface elevation ($s = H + h$, h being bedrock elevation) and D is the diffusion term:

$$D = 2\tilde{A}(\rho g)^n H^{n+2} \left(\left(\frac{\partial s}{\partial x} \right)^2 + \left(\frac{\partial s}{\partial y} \right)^2 \right)^{n-1/2}. \quad (2)$$

\tilde{A} is the depth averaged flow parameter A . Basal sliding is not included in this model, as a simplified model was chosen to aid initial parallelization. Parameter A introduces the temperature dependence of ice viscosity, and is calculated using an Arrhenius relation with constants based on work by Paterson and Budd (1982) and an enhancement factor of 5.

Ice temperature evolves according to

$$\frac{\partial T}{\partial t} = \frac{k}{\rho c} \frac{\partial^2 T}{\partial z^2} - U \cdot \nabla T - w \frac{\partial T}{\partial z} + \frac{g(s-z)}{c} \nabla s \cdot \frac{\partial U}{\partial z}, \quad (3)$$

where U is the horizontal velocity vector, k is the conductivity of ice ($6.62 \times 10^7 \text{ J m}^{-1} \text{ K}^{-1} \text{ yr}^{-1}$) and c its specific heat capacity ($2009 \text{ J kg}^{-1} \text{ K}^{-1}$). The terms represent vertical diffusion, horizontal advection, vertical advection and dissipation respectively. Following Paterson (1981), horizontal diffusion is assumed to be negligible.

Surface air temperature (T_a) is used as the upper boundary condition, and is calculated as a function of surface elevation (s), where sea level air temperature (SLAT) is 258 K, T_a is in K and s is in m:

$$T_a = \text{SLAT} - 1.2 \times 10^{-2} s. \quad (4)$$

Given the simplified nature of the model, this calculation of SAT is considered to be sufficient for the work described here.

Geothermal heat flux (G) forms the lower boundary condition at the bedrock (h) and has a value of $-5.46 \times 10^{-2} \text{ W m}^{-2}$. Basal temperature is constrained by the pressure melting point. The melt rate is negligible for the Antarctic Ice Sheet, so in this model the mass balance is assumed to be the accumulation rate alone. The spatially variable accumulation rate is read into the model from the data set compiled by Vaughan et al. (1999).

The model uses the BEDMAP topography data set (Lythe and Vaughan, 2000), smoothed from the original 5 km resolution to a 20 km resolution. The current topography is depressed due to ice loading, and would need to be rebounded artificially to approximate ice-free conditions. The ice evolution model was simplified to aid parallelization at this stage, so does not include isostatic response. Current topography is therefore used as the initial condition without rebounding the bedrock. This results in a topography that is too low for early stages of ice evolution, but is correct for the maturing ice sheet as it reaches steady state. If the bedrock had been rebounded at the start of the calculation and isostasy not incorporated, the surface elevation would be too high throughout the simulation, affecting the surface air temperature.

It was necessary to smooth the topography data set because steep bedrock gradients both around areas of high elevations and deep troughs otherwise cause numerical stability problems in the model. Bedrock is smoothed over the four nearest neighbours on the 20 km grid. Ice shelves are not included in the model because the physical description of floating ice is different from that of grounded ice. The present day grounding line is therefore held fixed within the model and cannot migrate during the model run. Ice thickness is prescribed to be zero at this grounding line. The artificial retention of ice inside the grounding line that would otherwise flow out to ice shelves may lead to a slightly thicker ice surface elevation just inside the grounding line. Bedrock

topography outside of the grounding line is not included in the model computations.

3. Numerical method

Ice flow is governed by a highly non-linear parabolic equation and is solved using a semi-implicit finite difference method. The temperature equation is a diffusion-advection problem, and is solved as a 1D problem with correction terms to incorporate the effects of horizontal advection. The ice sheet model uses a preconditioned conjugate gradient method to solve the ice flow equation, and a column based, semi-implicit method to solve the ice temperature evolution equation. The conjugate gradient method has the advantage of using a sparse matrix to store grid points, which saves computer memory. It is also reasonably straightforward to parallelize, as it requires only matrix-vector products and communication of points at the boundaries of regions on different processors.

Ice diffusivity is calculated on a staggered grid, with D being evaluated at x -midway and y -midway points staggered in the direction of flow. This is comparable with the Method 2 scheme of Hindmarsh and Payne (1996), who provide a more detailed discussion. First order horizontal upwinding is used in the calculation of velocity. This is known to introduce artificial diffusion, but has the advantage of improving numerical stability.

4. Parallel implementation

The ice flow equation (Eq. (1)) is solved on a two dimensional grid using a parallel sparse conjugate gradient method (Takeda et al., 2000). The ice temperature evolution equation is then solved in the vertical dimension, as an individual column of ice for each grid cell using a tridiagonal solver (Press et al., 1996).

The mass balance and topography data sets consist of square grids 281×281 cells in size, and are taken from the Budd data sets. The grids contain the Antarctic Ice Sheet but surrounding ocean is also included, covering approximately one third of the grid area. To overcome the problem of including irrelevant sea areas in expensive but unnecessary computations, land values are indexed, so every grid point that falls on the ice mass rather than the surrounding ocean has its own identity number. A lookup subroutine then allows the original co-ordinates of a grid point to be determined, given only its identity number. The model has a fixed grounding line, so the spatial extent of the ice sheet does not vary through time.

The map of the Antarctic Ice Sheet is split into strips of approximately equal size by dividing the total number of land cells by the number of processors. Each




Fig. 1. Antarctica is divided into 8 strips of approximately equal area to run on 8 processors. Although included in domain decomposition, ice shelves are excluded from ice thickness and temperature calculations.

processor's computational domain is then extended to the next coastal point if necessary to create a complete strip, as shown in Fig. 1.

The model was primarily written to run on a cluster of eight Compaq Alpha PCs using MPI/Pro (version 1.2.3)¹ for Windows NT (Hebert et al., 1998). On the Alpha cluster, where communication latency (overhead of setting up a message) is expensive compared with computation, it is better to send a smaller number of large messages than a larger number of small messages. Whilst not optimal, column decomposition was therefore found to be more efficient than block decomposition on the modest number of processors employed. Quinn (1994) showed that the block-oriented approach is superior to the column oriented-approach when

$$n > \frac{\lambda}{1 - \frac{2}{\sqrt{p}} \beta} \quad (5)$$

where n is the grid size, p the number of processors, λ the message latency and β the time needed to transmit a single value. For a model running on 8 processors on the Alpha cluster, the latency is $196 \mu\text{s}$ and the time required to pass one message $1.7 \mu\text{s}$ (Lancaster and Takeda, 1999). The grid size would therefore have to be over 393×393 for a block decomposition to be preferable to column decomposition.

¹ MPI/Pro for Windows NT, MPI Software Technology Inc., <http://www.mpi-softtech.com>.




Fig. 2. Ice thickness computation at point (ij) uses six nearest “east-west” diffusion neighbours at $(i - \frac{1}{2}, j - 1)$, $(i - \frac{1}{2}, j)$, $(i - \frac{1}{2}, j + 1)$, $(i + \frac{1}{2}, j - 1)$, $(i + \frac{1}{2}, j)$, $(i + \frac{1}{2}, j + 1)$, and six nearest “north-south” diffusion neighbours $(i - 1, j - \frac{1}{2})$, $(i - 1, j + \frac{1}{2})$, $(i, j - \frac{1}{2})$, $(i, j + \frac{1}{2})$, $(i + 1, j - \frac{1}{2})$, $(i + 1, j + \frac{1}{2})$. Calculation of diffusion at point $(i + \frac{1}{2}, j)$ uses ice thickness points $(i, j - 1)$, (i, j) , $(i, j + 1)$, $(i + 1, j - 1)$, $(i + 1, j)$, $(i + 1, j + 1)$. Calculation of diffusion at point $(i, j + \frac{1}{2})$ uses ice thickness from points $(i - 1, j + 1)$, $(i - 1, j)$, $(i, j + 1)$, (i, j) , $(i + 1, j + 1)$, $(i + 1, j)$.

All computations on the sparse matrix representing the discretized version of Eq. (1) proceed in parallel and the array is globally reduced to a single processor at the end for output to disk. If the resolution of the grid was further enhanced we would consider using the parallel I/O facilities of MPI 2.²

For numerical stability, calculation of ice diffusivity takes place on a staggered grid system, shown in Fig. 2. Along with the irregular boundaries of our grid, this prevented us from using the standard parallel template libraries for our solvers, e.g. KELP (Baden, 1996). The staggered diffusion grid means that the split between processors dissects the east-west diffusion term for ice thickness values on the processor border. For this reason there is an overlap, with diffusivity at the border being calculated by both processors. Diffusivity in the north-south direction is calculated on the same rows as the ice thickness points. It is therefore not affected by the processor split, so the calculation does not need to be duplicated.

Information regarding the north, south, east and west neighbours of individual grid points, and of the relevant neighbours on the staggered diffusion grid (Fig. 2) is stored in two arrays. Using these, it is possible to identify the neighbours of the grid points located on the processor boundaries, and to work out the limits of each processor's computational domain. Packing of the diffusivity terms in north-south and east-west directions can then be run in parallel using these limits for each

processor. These arrays are used to pack the information from the sparse irregular grid into arrays to be sent between processors. This helps to increase the latency tolerance of the solver, which is important for running on clusters of PCs.

5. Parallel performance

The code was originally designed to run on a cluster of eight 500 MHz, 21164 Compaq/DEC Alphas running Windows NT and MPI/Pro. It has also been run on a PC cluster running under Red Hat Linux 6.1 using LAM MPI version 6.3.1, which has two dual Pentium III 450 MHz processors and 16 single processors. Earlier work described by Takeda et al. (2000) was extended to include further benchmarking on the UK National CSAR Cray T3-E and Origin 2000.

The model was run for a reduced simulation time (1,500 years) on different numbers of processors to assess parallel performance. The results are plotted in Fig. 3. Efficiency is defined as

$$100 \times \frac{\text{runtime on 1 processor}}{\text{runtime on } n \text{ processors} \times n}$$

The model scales most efficiently on the Linux PC cluster, with an efficiency of 100% on two processors and 83% on 16 processors. On the Alpha cluster, the model did not scale as well as on the Linux cluster. However, the run time on a single Alpha processor is 4926 s compared with 6958 s on the Linux cluster, so communication latency effects will be comparatively more expensive on the Alpha cluster.

²Message Passing Interface Forum, “MPI2: Extensions to the Message Passing Interface”. <http://www.mpi-forum.org/docs>.

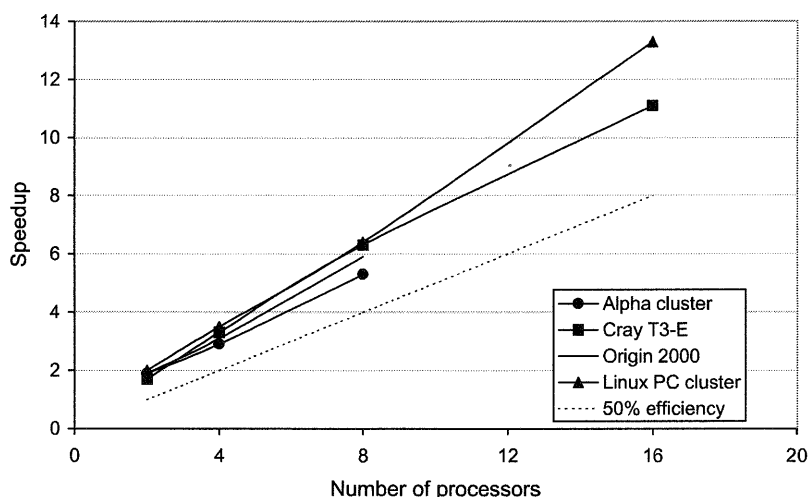


Fig. 3. Speedup of model on two PC clusters, an Origin 2000 and a Cray T3-E.

The efficiency of the test case on the Cray T3-E is 84% on 4 processors, 79% on 8 processors, but only 69% on 16 processors. The faster interconnects on the Cray T3-E should allow lower communication costs, so the comparatively low efficiency is unexpected. It is likely that this lower efficiency is due to the size of the test case—the machine simply may not have enough work to do between communication exchanges. It was difficult to assess performance on the Origin 2000 because the system is configured to swap the running model between different nodes, even pausing it at times. Despite the unreliable times recorded, the model still appears to run better on the Origin 2000 than on the Alpha cluster. Single processor run time at 4749 s was also faster than on the Cray T3-E and the Linux PC cluster.

The Antarctic ice flow model was successfully parallelized, allowing good performance on clusters of PCs (over 80% on the Linux cluster). However, the current computational model is too small to maximise parallel efficiency on the Cray T3E. A more detailed model at a finer grid resolution would have greater computational requirements, so would run more efficiently on a parallel computer.

6. Model validation

Results from the uncoupled parallel ice flow model, where ice temperature is calculated but the flow parameter A is held constant with a value of $A = 10^{-16} \text{ Pa}^{-3} \text{ a}^{-1}$, were compared against the EIS-MINT benchmarks for an idealized geometry (Huybrechts et al., 1996). The results compared closely, with the ice divide of 3420.51 m being well within the published error limit of 3419.9 ± 1.7 m. The results were the same to within machine error (i.e. to 15 decimal

places) when the model was run on more than one processor. This comparison was a necessary exercise in proving that parallel processing does not significantly affect the model's results.

The Antarctic model is run for 150,000 years from the initial bedrock start, by which time steady state has been reached by the temperature and ice thickness parts of the model. The basal temperature of the ice sheet takes longer to reach steady state than ice thickness does, with slight oscillatory behaviour observed in the early stages of evolution for areas of thinner ice cover. Two experiments were run to check for hysteresis in the model, i.e. to find out whether the initial conditions affect the model's final results. The standard test case (with $\text{GHF} = 54.6 \text{ mW m}^{-2}$) was run for 150,000 years, at which point the geothermal heat flux was increased to 100 mW m^{-2} and the model run for a further 150,000 years. A second run was carried out with a geothermal heat flux of 100 mW m^{-2} for 150,000 years, which was then decreased to the standard value of 54.6 mW m^{-2} for a further 150,000 years. The results of these experiments showed that the model's initial conditions did not affect the steady state ice sheet.

The parallel Antarctic model uses a horizontal grid of 281×281 points, with 20 km grid spacing. The effects of using a coarser grid were studied by restricting the map of Antarctica onto increasingly coarse resolution grids using bilinear interpolation (Press et al., 1996). A numerical solution will deviate increasingly from the "true" solution (taken to be the 281^2 result) as grid resolution becomes coarser.

The areas of the modelled ice sheet's base at pressure melting point are shown more clearly on the 281^2 (20 km) grid than on the 40 km resolution 141^2 grid (Fig. 4). In particular, coastal patterns of basal melt are more clearly defined and stretch further inland using the




Fig. 4. Areas of basal melt (shown in white) for (A) 141^2 grid and (B) 281^2 grid.

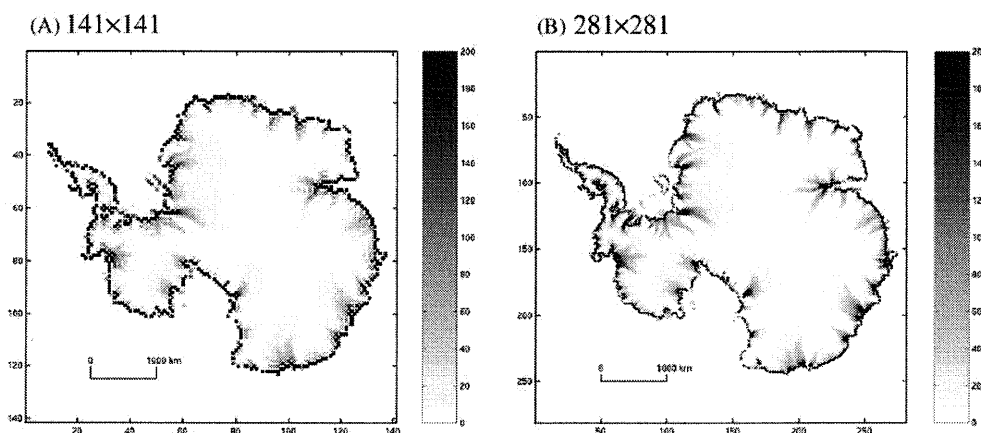


Fig. 5. Absolute flux (m^2a^{-1}) for (A) 141^2 grid and (B) 281^2 grid.

281^2 grid. Ice flux patterns also reflect this improvement (Fig. 5), and a log scale shows the East Antarctic coastal features much more clearly for the 281^2 grid than the 141^2 (Fig. 6). Huybrechts (1990) used a 40 km grid resolution, which was not sufficiently detailed to allow identification of individual flow features in the East Antarctic Ice Sheet. The extra computational cost of the 20 km resolution model is therefore justified, as the finer grid appears to be necessary to resolve these flow features using the present model. The model's assumption of ice flow being driven by shear stresses in the $x-z$ and $y-z$ planes is only appropriate for a grid resolution over 10 times the average ice thickness, in this case approximately 2 km. A more detailed model including longitudinal stresses would be required for a resolution finer than 20 km.

7. Ice thickness

The modelled ice thickness (Fig. 7) was compared against the BEDMAP ice thickness data-set (Lythe et al., 2000). The major ice domes, ice divides and general shape of the ice sheet is well modelled because this stems principally from the grounding line, which is taken from current conditions and held constant in the model. An experimental run with flat bedrock topography correctly reproduced the extent of the ice sheet, demonstrating the constraining effect of the grounding line.

The difference between the modelled ice thickness and the BEDMAP data set is plotted in Fig. 8. The positive (black) areas are where the modelled ice sheet is too thin, and the negative (white) areas are where the modelled ice sheet is too thick. There is generally a good comparison

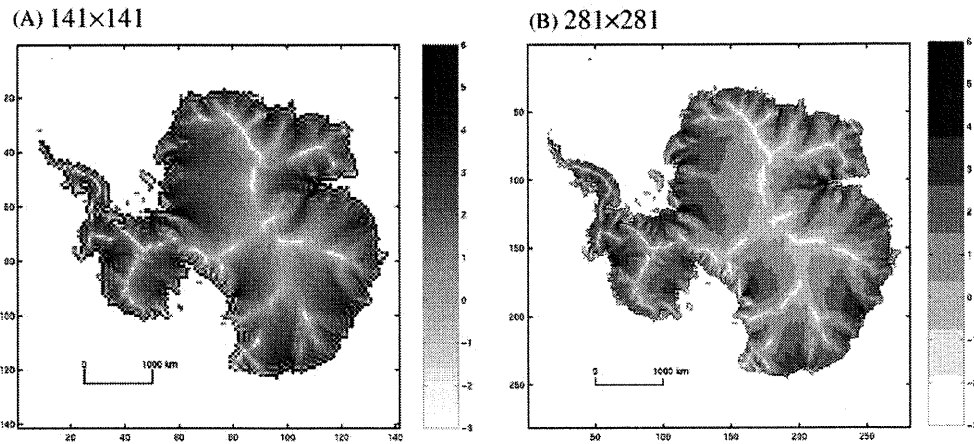


Fig. 6. Log absolute flux for (A) 141^2 grid and (B) 281^2 grid.




Fig. 7. Modelled ice thickness (m) at steady state.

between the data set and the modelled ice sheet, with most areas being within ± 200 m of the data. The East Antarctic Ice Sheet is slightly too thin in the east, notably to the south and west of the Lambert Glacier. The feedback between elevation and surface air temperature causes areas where surface elevations are lower than in reality to experience warmer surface air temperatures than they should. The warmer temperatures will lead to faster ice flow, preventing ice from reaching its correct surface elevation.

Other areas where the modelled ice sheet is significantly different from the BEDMAP data tend to be coastal areas near ice shelves. In reality, ice is drained into floating ice shelves, and then calved into the sea as icebergs. The model has an artificial grounding line, so there is an abrupt boundary at the edge of the ice sheet. This prevents ice from leaving the grounded ice sheet, so it builds up to a greater thickness than that observed in

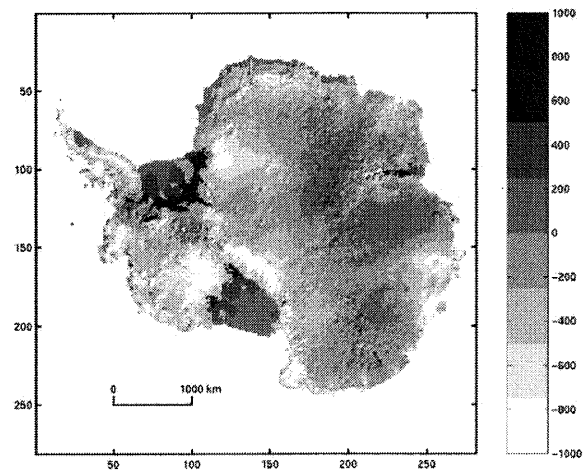


Fig. 8. BEDMAP ice thickness data set minus modelled ice thickness. Coastal differences in ice thickness are due to ice being trapped behind model's artificial grounding line. In reality, ice is thinner because it flows out into ice shelves.

reality along the coastline. The region of the WAIS where ice streams feed into the Ross Ice Shelf is too thick in the modelled ice sheet. In reality, the ice drains away from this region at a faster rate than the simplified model suggests. Ice shelves are not included in the current model because the shallow ice approximation on which the model is based is not applicable to ice shelf flow. Overall, the ice thickness results are within the tolerances expected from the input data and the model assumptions.

8. Balance velocities

If an ice sheet is at steady state the depth-averaged horizontal velocity should balance the ice accumulated

upstream of a particular point. Bamber et al. (2000) have calculated balance velocities for the Antarctic Ice Sheet using recent data sets collected by remote sensing. The ice sheet is assumed to be at steady state, so balance velocities can be calculated as the depth averaged velocities required to maintain the present surface profile.

The balance velocities in the ice streams calculated by Bamber et al. (2000) are generally 150 ma^{-1} but increase to 250 ma^{-1} in a few places. The depth averaged horizontal velocities calculated in the parallel model are slightly lower, at a little over 150 ma^{-1} in the fast flow regions. Ice stream flow is governed by different physics to ice sheet flow, so the present model cannot be used to simulate these fast flow environments accurately. However, many of the ice streams identified by Bamber et al. (2000) do appear as patches of faster flow in the modelled ice sheet shown in Fig. 9. A log scale is used in Fig. 9 as it highlights areas of faster flow more clearly than a map of flux (e.g. Fig. 5A). This suggests that the model is at least predicting the pattern of fast ice flow correctly, implying that basal melt is an important precondition for ice-stream flow.

There is a region of enhanced flow in the Siple Coast area of the modelled WAIS that is in the same place as the well-studied ice streams A-E. The Lambert Glacier region is particularly well-resolved, with a similar pattern of flow observed in Fig. 9 as the balance velocity map produced by Bamber et al. (2000). Many other glaciers such as the Byrd Glacier appear on the map of modelled velocities, in addition to numerous fast flow features stretching inland from the coast of the East Antarctic Ice Sheet. These fast flow patterns have not been observed in previous models of East Antarctica

because of their generally poor grid resolution (Bamber et al., 2000).

The depth-averaged modelled horizontal velocities shown in Fig. 9 can be compared with the basal temperature field. Faster flowing ice is warm and relatively soft, so is related to areas of basal melting. The faster flow will help to maintain basal melting conditions because of the feedback between velocity and temperature, and the basal melting conditions themselves help to maintain faster flow via a basal slip mechanism. Basal meltwater helps to lubricate the ice directly, and may also saturate basal till, possibly allowing subglacial deformation. The areas of the ice sheet that have a non-zero basal melt rate are plotted in Fig. 4B.

The patterns of faster flow in Fig. 9 correspond to areas of Fig. 4B where the ice sheet base is at pressure melting point. The West Antarctic ice streams draining into the Ross Ice Shelf appear as an area of faster flow, and are represented as two patches where basal melt rate is greater than zero. Generally, the patterns of non-zero basal melt rate are larger and have less well defined boundaries than the fast flow features shown in Fig. 9.

9. Subglacial lakes

A comprehensive inventory of over 70 lakes has been compiled by Siegert et al. (1996). The presence of subglacial lakes suggests that these parts of the ice sheet base are at pressure melting point, so their locations can be compared with the modelled basal thermal regime.

Approximately half of the 77 mapped subglacial lakes fall within the zones of basal melting calculated by the model (Fig. 10). Lake Vostok is the largest lake in the region, with an area of $14,000 \text{ km}^2$, and this lies just outside the area of basal melting in the modelled ice sheet. Most of the Dome C lakes fall within the model's zone of basal melting. The majority of lakes in the South Pole area and all of those in the Hercules Dome region are excluded from the basal melting zones. The smoothing of topography that was necessary to maintain model stability may also affect the model's ability to locate subglacial lakes. The smaller topographic locations required for development of some lakes may have been smoothed out. It is also possible that inaccuracies in the geothermal heat flux may affect the number of mapped subglacial lakes coinciding with areas of basal melting in the model.

10. Spatially variable GHF

The geothermal heat flux (GHF) beneath ice sheets is difficult to measure in the field because of the huge depth of ice overlying the bedrock (frequently over 3 km in




Fig. 9. Modelled depth-average horizontal ice velocities. Scale is in $\log_{10} \text{ ma}^{-1}$.

Antarctica), so an approximate value is used in ice sheet modelling studies. The value normally used for ice sheet modelling where Precambrian shield forms the bedrock is 42 mW m^{-2} . Hansen and Greve (1996) suggested from measurements beneath the Antarctic Ice Sheet that this value should be much higher. They found that a GHF value of $50\text{--}60 \text{ mW m}^{-2}$ reproduces the present ice sheet topography fairly well. Following Huybrechts (1990), we have used a GHF value of 54.6 mW m^{-2} .




Fig. 10. Areas of modelled ice sheet at pressure melting point are white. Lakes mapped by Siegert et al. (1996) are shown as triangles.

Although a spatially uniform value for GHF across the whole of the continent is used in ice sheet models, it is likely that geological variations in the bedrock actually produce different GHF values in different parts of the Antarctic Ice Sheet. The younger rocks of West Antarctica are likely to have a higher geothermal heat flux than the geologically older rocks of East Antarctica. However, few measurements exist to allow realistic inclusion of a range of values for different areas of the ice sheet. Siegert (2000) calculated from the heat budget on subglacial lakes that the GHF of the Antarctic Plate varies between 37 and 65 mW m^{-2} .

To allow for geological variation the model was run with the GHF in West Antarctica set to 65 mW m^{-2} , the upper limit of Siegert (2000), and the GHF in East Antarctica set to the standard value of 54.6 mW m^{-2} . Although this resulted in a thinner WAIS, there was little impact on velocity patterns. An increase in areas of the WAIS where basal ice is at pressure melting point was observed in the model results (Fig. 11), although the lack of field observations for basal melting prevents a direct comparison. The areas of basal melting in the WAIS are extended when a higher GHF is used, and a small patch of basal melting appears beneath the main WAIS ice cap where the ice is thickest.

11. Comparison between data sets

The model described in this paper was run using the BEDMAP topography and the accumulation data set of Vaughan et al. (1999). We also ran the model with the earlier data sets compiled by Budd and Smith (1984) for




Fig. 11. Areas of basal melt are shown in white for (A) model with $\text{GHF} = 65 \text{ mW m}^{-2}$ in WAIS and $\text{GHF} = 54.6 \text{ mW m}^{-2}$ in EAIS, and (B) $\text{GHF} = 54.6 \text{ mW m}^{-2}$ in both WAIS and EAIS. Grid points on x and y axes have 20 km spacing.

comparison with the new data. The original Budd ice thickness data set is known to contain errors, as demonstrated by Bamber and Huybrechts (1996) in a comparison with a new digital elevation model of Antarctica. Fig. 12 shows the difference between the Budd ice thickness data set and the ice sheet modelled using Budd accumulation and topography data sets. The areas shown in white in Fig. 12 are where the modelled ice sheet is too thick, and the black areas are where the model under-predicts ice cover. Comparison with Fig. 8 shows the great improvement in the model when the new BEDMAP topography and Vaughan et al. (1999) accumulation data are used.

The Budd data sets show large areas where the modelled ice sheet is 1000 km thicker than the data set, whereas the ice sheet modelled using the BEDMAP topography rarely reaches such extremes (Fig. 8). The modelled Antarctic Peninsula in particular is much thicker in the model run with the Budd data sets, and it is likely that this is due to inaccuracies in the accumulation data set that have been eliminated in the new Vaughan data set. Coastal areas of East Antarctica have modelled ice cover that is far too thick when the Budd data sets are used, whereas the model run with the new data sets presents far more accurate results for these areas. It is possible that this is related to patterns of basal melt (Fig. 13), which differ slightly from those seen in the BEDMAP model run (Fig. 4B).

Differences in basal melt patterns are particularly noticeable in the Lambert Glacier region, where long stretches of basal ice at pressure melting point stretch inland. The pattern of basal melting in Fig. 4B reflects the balance velocities of Bamber et al. (2000) more accurately than the results using the original Budd data

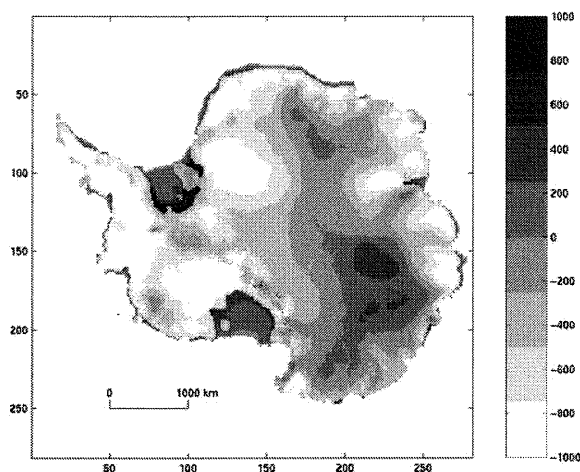


Fig. 12. Difference in metres between Budd ice thickness data set and ice thickness modelled using Budd topography and accumulation data sets.




Fig. 13. Areas of basal melt are shown in white for model run with Budd data sets.

sets. The new BEDMAP data set contains more topographic detail than the original Budd data set, and may therefore allow more distinct channelling of ice flow and associated basal melt. The surface topography is derived internally, so this shows that the model is able to reproduce the surface elevation correctly. The Siple Coast has a single patch of basal melting in Fig. 13 whereas the region is split into two areas of melt in Fig. 4B, giving a more accurate representation of the melt associated with the five ice streams in the area. The Byrd Glacier appears as the only area of basal melt in the Transantarctic Mountains using the BEDMAP data, whereas it appears as one of several glaciers in the Budd data run.

12. Conclusions

The main focus of this paper is the design a parallel ice sheet model enabling faster runtime of a 20 km resolution model for the whole of Antarctica. The model was simplified to ease parallelization at this stage, but even without more sophisticated features the results show some interesting patterns of basal melt.

Our model requires the coupled solution of highly non-linear partial differential equations on a staggered grid over a domain with irregular boundary. The scheme is semi-implicit, and thermal coupling requires a reduction in timestep to ensure numerical stability. The coupled model therefore takes much longer to run than the uncoupled version, due to the necessary

timestep reduction in addition to the extra computation required. Parallelization of the code has therefore been essential in making a high-resolution thermally coupled model of the entire Antarctic ice sheet feasible. A simplified model was used in this initial work, but the same techniques could be used in the parallel implementation of higher order models.

The scaling tests show that the model runs most efficiently on the Linux based PC cluster. However, in terms of run-time, the model runs fastest on the Alpha cluster and the Origin 2000. Running on multiple processors considerably reduces run time, so the goal of efficient parallelization has been achieved. The model was validated against the EISMINT benchmarks, and compared against the ice thickness data set. Parallelization has made it possible to run a finer resolution model of both East and West Antarctica within a reasonable time scale, allowing features of the EAIS to be seen at a higher resolution than previously possible, and results compare well with balance velocity patterns calculated by Bamber et al. (2000).

The original Budd data set is known to contain errors, as demonstrated by Bamber and Huybrechts (1996) in a comparison with a new digital elevation model of Antarctica. The BEDMAP topography and Vaughan et al. (1999) accumulation data used here clearly improve the results from the ice sheet model. The ice sheet modelled using the BEDMAP topography data has more sharply defined areas of basal melting along the coastal regions than the model using the original Budd data. Use of the BEDMAP data sets allows a closer comparison to be made between flow features in the modelled ice sheet and the patterns of balance velocities calculated by Bamber et al. (2000).

Acknowledgements

We are grateful to R. Greve and an anonymous reviewer for their useful comments on an earlier draft of this paper. Time on the Cray T3-E and the Origin 2000 was provided by CSAR at Manchester for NERC standard grant GR3/12917, and we would like to thank MPI Pro for supporting software on the Alpha cluster. ALT is grateful to EPSRC for a studentship.

References

- Baden, S.B., 1996. Software infrastructure for non-uniform scientific computations on parallel processors. *Applied Computing Review ACM* 4 (1), 7–10.
- Bamber, J.L., Huybrechts, P., 1996. Geometric boundary conditions for modelling the velocity field of the Antarctic Ice Sheet. *Annals of Glaciology* 23, 365–373.
- Bamber, J.L., Vaughan, D.G., Joughin, I., 2000. Widespread complex flow in the interior of the Antarctic Ice Sheet. *Science* 287 (5456), 1248–1250.
- Budd, W.F., Jenssen, D., Smith, I.N., 1984. A three-dimensional time-dependent model of the west Antarctic Ice Sheet. *Annals of Glaciology* 5, 29–36.
- Calov, R., Savvin, a., Greve, R., Hansen, I., Hutter, K., 1998. Simulation of the Antarctic Ice Sheet with a three-dimensional polythermal ice sheet model, in support of the epica project. *Annals of Glaciology* 27, 201–206.
- Hansen, I., Greve, R., 1996. Polythermal modelling of steady states of the Antarctic Ice Sheet in comparison with the real world. *Annals of Glaciology* 23, 382–387.
- Hebert, S., Seefeld, W., Skjellum, A., Taylor, C.D., Dimitrov, R., 1998. MPI for NT: Two generations of implementations and experience with the Message Passing Interface for clusters and SMP environments. In: *Proceedings 1998 International Conference on Parallel and Distributed Processing Techniques and Applications*, Las Vegas, pp. 309–316.
- Hindmarsh, R.C.A., Payne, A.J., 1996. Time-step limits for stable solutions of the ice-sheet equation. *Annals of Glaciology* 23, 74–85.
- Huybrechts, P., 1990. A 3D model for the Antarctic Ice Sheet: a sensitivity study on the glacial-interglacial contrast. *Climate Dynamics* 5, 79–92.
- Huybrechts, P., Payne, A.J., Abe-Ouchi, A., Calov, R., Fabre, A., Fastook, J.L., Greve, R., Hindmarsh, R.C.A., Hoydal, O., Jóhannesson, T., et al., 1996. The EISMINT benchmarks for testing ice-sheet models. *Annals of Glaciology* 23, 1–12.
- Lancaster, D.J., Takeda, K., 1999. Comparative performance of a commodity Alpha cluster running Linux and Windows NT. In: *Proceedings 1st Institute of Electrical and Electronics Engineers (IEEE) Computer Society International Workshop on Cluster Computing (IWCC-99)*, Melbourne, pp. 19–25.
- Lythe, M.B., Vaughan, D.G., BEDMAP Consortium, 2000. BEDMAP—bed topography of the Antarctic, 1:10 000 000 scale map. BAS (Misc) 9. Cambridge, British Antarctic Survey, 1pp.
- Paterson, W.S.B., 1981. *The Physics of Glaciers*, 2nd Edn. Pergamon, New York, 380pp.
- Paterson, W.S.B., Budd, W.F., 1982. Flow parameters for ice sheet modelling. *Cold Regions Science and Technology* 6, 175–177.
- Payne, A.J., 1998. Dynamics of the siple coast ice streams, west antarctica: results from a thermomechanical ice sheet model. *Geophysical Research Letters* 25 (16), 3173–3176.
- Payne, A.J., 1999. A thermomechanical model of ice flow in west Antarctica. *Climate Dynamics* 15, 115–125.
- Payne, A.J., Dongelmans, P.W., 1997. Self-organization in the thermomechanical flow of ice sheets. *Journal of Geophysical Research* 102 (B6), 12219–12234.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 1996. *Numerical Recipes in Fortran 77. The Art of Scientific Computing*, 2nd Edn. McGraw-Hill Inc., New York, 933pp.
- Quinn, M.J., 1994. *Parallel Computing: Theory and Practice*, 2nd Edn. McGraw-Hill Inc., New York, 446pp.

- Rémy, F., Shaeffer, P., Legrésy, B., 1999. Ice flow physical processes derived from the ERS-1 high-resolution map of the Antarctic and Greenland ice sheets. *Geophysical Journal International* 139, 645–656.
- Savvin, A., Greve, R., Calov, R., Mücke, B., Hutter, K., 2000. Simulation of the Antarctic Ice Sheet with a three-dimensional polythermal ice sheet model, in support of the EPICA project. II. Nested high-resolution treatment of Dronning Maud, Land Antarctica. *Annals of Glaciology* 30, 69–75.
- Siegert, M.J., 2000. Antarctic subglacial lakes. *Earth Science Reviews* 50, 29–50.
- Siegert, M.J., Dowdeswell, J.A., Gorman, M.R., McIntyre, N.F., 1996. An inventory of Antarctic sub-glacial lakes. *Antarctic Science* 8 (3), 281–286.
- Takeda, A.L., Cox, S.J., Payne, A.J., 2000. Parallel numerical modelling of ice flow in Antarctica. In: *Proceedings of International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 2000)*, Las Vegas, June 2000, pp. 335–341.
- Vaughan, D.G., Bamber, J.L., Giovinetto, M., Russell, J., Cooper, A.P.R., 1999. Reassessment of net surface mass balance in Antarctica. *Journal of Climate* 12, 933–946.