The Cryosphere Discuss., 2, 75–109, 2008
www.the-cryosphere-discuss.net/2/75/2008/
© Author(s) 2008. This work is distributed under
the Creative Commons Attribution 3.0 License.

The Cryosphere Discussions

The Cryosphere Discussions is the access reviewed discussion forum of *The Cryosphere*

The ISMIP-HOM benchmark experiments performed using the Finite-Element code Elmer

O. Gagliardini[1] and T. Zwinger[2]

[1]LGGE, CNRS, UJF-Grenoble I, BP 96, 38402 Saint-Martin d'Hères Cedex, France
[2]CSC-Scientific Computing Ltd., Keilaranta 14, P.O. Box 405, 02101 Espoo, Finland

Received: 20 December 2007 – Accepted: 3 January 2008 – Published: 19 February 2008

Correspondence to: O. Gagliardini (gagliardini@lgge.obs.ujf-grenoble.fr)

Published by Copernicus Publications on behalf of the European Geosciences Union.

Abstract

The aim of this paper is to describe in detail how the benchmark tests ISMIP-HOM (Ice Sheet Model Intercomparison Project – Higher-Order ice-sheet Model) has been performed using the open-Source finite element code Elmer (http://www.csc.fi/elmer). The ISMIP-HOM setup consists of five diagnostic and one prognostic experiments, for both 2-D and 3-D geometries. For all the tests, the full-Stokes equations are solved. Some FE technical points, such as mesh characteristics, stabilisation methods, numerical methods used to solve the linear system and parallel performance are discussed. For all these setups, the CPU time consumption is analysed in comparison to the accuracy of the solution. Some general rules are then inferred that optimise the computing time versus the accuracy of the results.

1 Introduction

Following the EISMINT (European Ice Sheet Modelling INiTiative, Payne et al., 2000) benchmark experiments, the Ice Sheet Model Intercomparison Project (ISMIP) aims to provide a comparison of the new generation of ice-sheet flow models. It is composed of three different tests. The comparison of the ice-dynamic response of Antarctic and Greenland ice sheet models applied to climatic warming (ISMIP-POLICE), coordinated by Philippe Huybrechts, is dedicated to global ice-sheet models. The simulation setup for Heinrich-type ice-sheet instabilities (ISMIP-HEINO), coordinated by Reinhard Calov and Ralf Greve, aims to compare the ability of 3-D ice-sheet flow models to simulate large scale surges induced by an activation and deactivation wave of the temperate basal area. Since the model domain represents a small synthetic circular ice-sheet, this experiment is not limited to global ice-sheet flow models and can be performed by higher order flow models. The last experiment that compares Higher-Order Models (ISMIP-HOM, http://homepages.ulb.ac.be/~fpattyn/ismip/) is coordinated by Frank Pattyn and is addressed to all 2-D and 3-D ice flow models.

TCD

2, 75–109, 2008

In the present paper, we discuss in detail the results obtained by performing the ISMIP-HOM experiments with the FE open source code Elmer. The complexity of the equations to be solved in the proposed tests is such that most of the standard FE software for fluid dynamic should be able to perform the experiments (at least the five diagnostic ones). Therefore, the interest of such experiments is more on the comparison of the results obtained for a various range of all the parameters that control the simulation, rather than in producing the results themselves. In this paper, we present a detailed study on the influence of the mesh characteristics (e.g. number of nodes, type of elements, ratio of horizontal to vertical number of elements), numerical methods used to solve the linear system (direct or iterative methods), and the influence of the tolerance criteria used to control the convergence of the non-linear iteration due to the non-linear ice behaviour. To compare these results, one should check the accuracy of the solution as well as the CPU time consumption. This last information, even if being machine dependent, should have been part of the model output for the ISMIP-HOM experiments because it is primordial to know for which size of problem each model can be reasonably applied. All the results presented here have been performed on a the HP ProLiant Cluster of the Center for Scientific Computing (CSC, Finland, http://www.csc.fi/metacomputer/sepeli.phtml.en).

The results from all the applicants are presented in Pattyn et al. (2008). The number of participants is surprisingly large (28). Among these, only two full-Stokes model, including Elmer, have been ran on all the tests. This paper presents in details the results obtained with Elmer.

2 Equations and numerical methods

We briefly summarise the equations to be solved and the numerical methods used to solve them by the finite element (FE) method.

TCD
2, 75–109, 2008

2.1 Equations to be solved

For all these experiments, we are interested in solving the gravity-driven flow of ice over a rigid bedrock for various boundary condition at the ice-bedrock boundary. The constitutive law for the ice behaviour is given by a Norton-Hoff type law (Glen's law in Glaciology):

$$\tau_{ij} = 2\eta D_{ij}\,, \tag{1}$$

where $\boldsymbol{\tau}$ and $\boldsymbol{D}$ are the deviatoric stress and strain-rate tensors, respectively. Both linear and non-linear ice rheology are envisaged. In this later case, the effective viscosity η is strain-rate dependent and denotes:

$$\eta = \frac{A^{-1/n}}{2}\dot{\epsilon}_e^{(1-n)/n}\,, \tag{2}$$

where the second invariant of the strain-rate $\dot{\epsilon}_e$ is such that

$$2\dot{\epsilon}_e^2 = \mathrm{tr}\,\boldsymbol{D}^2 = \boldsymbol{D} : \boldsymbol{D}\,. \tag{3}$$

Since the HOM experiments are isothermal, the fluidity parameter A, in Eq. (2), is a constant.

Since ice has a very high viscosity and can be regarded as incompressible, the mass conservation and the conservation of linear momentum reduce to the well-known set of Stokes equations:

$$\mathrm{div}\,\boldsymbol{u} = \mathrm{tr}\,\boldsymbol{D} = 0\,, \tag{4}$$

and

$$\mathrm{div}\,\boldsymbol{\sigma} + \rho\boldsymbol{g} = 2\eta\,\mathrm{div}\,\boldsymbol{D} + \mathrm{grad}\,p + \rho\boldsymbol{g} = 0\,, \tag{5}$$

where ρ is the ice density, $\boldsymbol{g}$ the gravity vector and the stress tensor $\boldsymbol{\sigma}$ has been decomposed into its deviatoric part – replaced by the strain-rate using Eq. (1) – and

CC BY

the isotropic pressure p. For both 2-D and 3-D experiments, the z-axis is opposite to the gravity and the x-axis points in the principal direction of the flow.

In the case of the prognostic experiments, the surface elevation $z{=}z_s(x,y,t)$ is part of the solution. For this free surface, the following equation serves as the kinematic boundary condition:

$$\frac{\partial z_s}{\partial t} + \boldsymbol{u}_\perp.\boldsymbol{\nabla}_\perp z_s = u_z + a \quad \text{for all } z = z_s(x,y,t), \tag{6}$$

where $\boldsymbol{u}_\perp{=}(u_x,u_y)$, the operator $\nabla_\perp$ stands for the gradient evaluated in the horizontal directions, i.e. $\nabla_\perp{=}(\partial./\partial x, \partial./\partial y)$, and a is the accumulation-ablation function, considered as a vertical flux. For experiment F, the accumulation is assumed to be zero.

2.2 Boundary conditions

For each of these conservation equations, a various number of boundary conditions have to be applied, depending on the test case. Three different kinds of boundaries can be distinguished for all these tests.

2.2.1 Ice-bedrock interface $z{=}z_b(x,y,t)$

For the Stokes equations, either no sliding is assumed ($\boldsymbol{u}(x,y,z_b){=}0$) as in experiments A, B and E, or a friction law linking the sliding velocity to the basal shear stress applies, as for tests C and D. Such friction law is of the form:

$$\boldsymbol{t}_i.\boldsymbol{u} = A_s \boldsymbol{t}_i.(\boldsymbol{\sigma}.\boldsymbol{n}_b) = A_s \tau_{bi} \quad (i = 1,2), \tag{7}$$

where $\boldsymbol{n}_b$ is the unit normal vector pointing into the bedrock and $\boldsymbol{t}_i$ the unit tangent vectors ($i{=}1$ in 2-D and $i{=}1,2$ in 3-D). The sliding parameter A_s is only a function of space ($A_s{=}1/\beta^2$ in C, D and F). Note if using a higher order model, the basal shear stress (basal drag, τ_b) should not be equated with the basal driving stress ($\rho g(z_s{-}z_b)\boldsymbol{\nabla}_\perp z_s$), but is part of the solution.

CC BY

2.2.2 Upper surface $z=z_s(x,y,t)$

For all these tests, the upper surface is a stress free surface, which implies that $\boldsymbol{n}_s \cdot (\boldsymbol{\sigma} \cdot \boldsymbol{n}_s) = p_{\text{atm}} \approx 0$, where $\boldsymbol{n}_s$ is the unit normal vector of the surface. Note that this condition does not imply that the isotropic pressure p vanishes at the surface. As a consequence of the variational formulation used in Elmer, this stress free condition is inherent (natural boundary condition), such that no explicit condition has to be applied for the Stokes problem.

For the kinematic boundary condition at the free surface (Eq. 6), the rate of accumulation, a, has to be prescribed in form of a body force.

2.2.3 Lateral boundaries of the model

For experiments A, B, C, D and F, the periodicity of the boundary is applied by linking the unknowns at the nodes of the two corresponding boundaries together.

2.3 Finite element formulation

We hereafter only discuss some technical aspects of interest for the ISMIP-HOM experiments. For a detailed presentation of the FE formulation in Elmer, the reader can consult the Elmer documentations (http://www.csc.fi/elmer).

2.3.1 Stokes equation

The ice flow is described by the Stokes problem for an incompressible fluid, corresponding to the solution of Eqs. (4) and (5). A numerical solution of this set of equations is obtained by either using the classical P2P1 elements (quadratic interpolation of the velocity and linear interpolation of the pressure), the stabilised method (Franca and Frey, 1992) or the residual free bubbles method (Baiocchi et al., 1993).

The calculation of the stress from the velocity and isotropic pressure fields is a matter of interest because different methods can lead to sensibly different solutions. In

CC BY

this study, the deviatoric stress field is obtained by solving Eq. (1) using the following variational form with the scalar test functions Φ:

$$\int_V \tau_{ij}\Phi \, \mathrm{d}V = 2\int_V \eta D_{ij}\Phi \, \mathrm{d}V \ , \tag{8}$$

where D_{ij} and η are computed from the nodal velocities using the derivative of the basis functions.

2.3.2 Free surface equation

The free surface elevation z_s is discretized as

$$z_s(x,y,t) = \Psi_i(x,y) z_s^i(t) \ , \tag{9}$$

where z_s^i is the (nodal) elevation at the i-th node inside the local element of the discretized ice sheet surface, and Ψ_i stands for the nodal interpolation functions.

The discrete variational form of Eq. (6) is obtained by spatial integration over the ice sheet surface, using the test function Φ. At time step $t+1$ the following set of equations is then solved

$$\frac{\partial z_s^i}{\partial t}\int_{z_s(t)} \Psi_i \Phi \, \mathrm{d}z_s(t) + z_s^i \int_{z_s(t)} \boldsymbol{u}_\perp . \boldsymbol{\nabla}_\perp \Psi_i \Phi \, \mathrm{d}z_s(t) = \int_{z_s(t)} (u_z + a)\Phi \, \mathrm{d}z_s(t) \ , \tag{10}$$

where $z_s(t)$ is the ice sheet surface at time step t, and $\boldsymbol{u}_\perp = (u_x, u_y)$ and u_z are taken from the solution of the Stokes problem at time t.

Due to the hyperbolic nature of Eq. (10) the standard Galerkin method (i.e. $\Phi \equiv \Psi$) does not apply. Stabilisation is obtained by applying the stabilised method introduced by Franca et al. (1992).

At each time step, to avoid a distortion of the domain mesh caused by the moving free surface, the nodes of the domain mesh are re-distributed by solving a fictive elasticity problem. To this end, the domain is given artificial elastic properties, and it is deformed according to the displacement, $z_s(t)$, of the free surface.

CC BY

2.4 Numerical methods

2.4.1 Linear system

The discretization of partial differential equations by the FE method leads to a linear system to be solved. There are two categories of method to solve this linear system: direct and iterative methods. The first leads to an exact solution up to machine precision, but do not apply to very large system. The latters utilise a converging sequence of approximate solutions.

Two-dimensional problems were solved by the direct Unsymmetric MultiFrontal method (UMFPACK, Davis, 2004), whereas a Biconjugate Gradient Stabilised method (BiCGStab, Kelley, 1995) using an incomplete LU decomposition preconditioning was applied to 3-D problems. The convergence of the iterative method is obtained when the relative change of the residuals is lower than the criterion ϵ_L.

2.4.2 Non-linear system

The non-Newtonian stress-strain relation introduces non-linearities into the system. Linearization of those terms implies the application of an iterative scheme. The variables used for the evaluation of the velocity dependent effective viscosity, η, for the $(n+1)$-th non-linear step are taken from the previous iteration step. Convergence is checked upon the global change of the field variables

$$\frac{1}{N}\sum_{i=1}^{N}|\boldsymbol{U}_i^{n+1}-\boldsymbol{U}_i^{n}|<\epsilon_{NL}\ll 1, \tag{11}$$

where $\boldsymbol{U}_i$ stands for the solution vector at the i-th (out of total N) nodes.

TCD
2, 75–109, 2008

2.4.3 Coupled solvers and time stepping scheme

Applying constant boundary conditions, the prognostic tests are integrated over time as long as a steady state is reached. For time-integration, an implicit scheme is applied. On a single time-level, the two solvers are run iteratively until the convergence criterion ϵ_C is reached. Using this implicit scheme, the time step can be increased with no decrease of the accuracy for the solution. If the time step is small enough, an explicit scheme is equivalent.

For a good convergence of each problem, one should have $\epsilon_L < \epsilon_{NL} < \epsilon_C$.

2.4.4 Mesh generation

Implied by to the domain geometry, experiments A, B, C, D and F were performed using layered meshes consisting of N_x horizontal and N_z vertical layers in 2-D and N_x and N_y horizontal and N_z vertical layers in 3-D. The mesh is then composed of quadrilateral elements with 4 nodes in 2-D and hexahedral elements with 8 nodes in 3-D.

For experiment E, in order to compare the influence of longitudinal as well as vertical element density, four different computational meshes using layers and one using unstructured meshing technique have been created. To this end, the 2×51 vertices (with two pairs coinciding) defining the free surface as well as the bedrock have been read into the commercial pre-processor Gambit. Thereafter the nodes have been connected with straight lines along the bedrock as well as free surface. Also vertical connections between the corresponding points on the free surface and the bedrock have been applied.

Consequently, the initial geometry consisted of 50 sections (48 of quatri-lateral and 2 of triangular topology), which further have been meshed with different vertical and horizontal resolutions (see Fig. 1).

This has the advantage that the 2×49+2 supporting points for which output is required inherently are part of every mesh and in neither run output variables have to be interpolated there. Table 1 shows the parameters for the different meshes, which are

displayed in Fig. 2.

3 General settings using FEM

3.1 Mesh influence

3.1.1 Mesh density

Mesh density is one of the most important parameter that controls the accuracy of the solution, but the drawback is that increasing the mesh resolution increases the CPU time consumption. Due to the numerical methods used, the CPU time is found to be proportional to the number of degrees of freedom to the power 1.27 and 1.11 in 2-D and 3-D applications, respectively (see Fig. 3 and its caption). For the Stokes problem, the degrees of freedom are $3N$ in 2-D and $4N$ in 3-D, where N is the total number of nodes. In order to run one of the diagnostic experiments, the CPU cost is approximately $0.0024\,N^{1.27}$ s for 2-D problems, whereas it is $0.061\,N^{1.11}$ s in 3-D. This difference of regression between 2-D and 3-D cases, certainly results from the different methods used to solve the linear system (direct in 2-D and iterative in 3-D). Note the discrepancy of the CPU time consumption for the different tests, indicating that the CPU time consumption depends on the geometry of the problem to be solved (and not only the mesh density).

3.1.2 Ratio of horizontal to vertical layers

The in total six variations of the experiments B and D were performed using a layered mesh consisting of N_x and N_z horizontal and vertical layers, respectively. The ratio N_x/N_z was varied from 1 up to 10, while keeping an almost constant value for N. Since we imposed that N_x is always a multiple of 4 to have one node exactly at the minimum and maximum value of the sinusoidal functions (bedrock equation for test B and sliding

CC BY

function for test D), this number cannot be exactly the same for all the meshes. The number of nodes was 10 278±267.

The comparison of the tests is done on the minimum and maximum values of the field variables computed at the surface and the bottom. In the whole range of the tested ratios N_x/N_z, the maximum relative difference for the velocity, stress and pressure fields is less than 1% for test B. For test D, this relative difference is also smaller than 1% for all the fields, excepted for the minimum value of shear stress at the bottom, which should vanish in the limit of $A_s \rightarrow \infty$. Relative to the maximum value of the shear stress at the bottom (≈100 kPa for the 5 tests), this error is lower than 5% for all the meshes, but increasing with a growing ratio N_x/N_z.

The aspect ratio H/L of the experiments, where H and L are, respectively, the height and the length of the domain, varies from 0.006 up to 0.20. As commonly admitted for the FE method, elements should not be to much elongated. For the six experiments B and D, one cannot conclude from these results which ratio N_x/N_z between 1 to 10 is more adapted for these tests. From these results, we show that for these particular cases where the difference between horizontal and vertical gradients is large, elongated quadrilateral elements can be used. For L=160 km, this leads to an element aspect ration of 0.0125. Therefore, the same ratio N_x/N_z=240/120=2 was used to produce the submitted results. This result from 2-D tests was also extended to 3-D, so that the 3-D meshes were built with N_x/N_z=N_y/N_z=2 for experiments A and C.

3.2 Convergence criteria

The convergence of the solution is controlled by three criteria:

- ϵ_L for the convergence of the linear solution if an iterative method is used to solve the linear system (3-D experiments),
- ϵ_{NL} for the convergence of the non-linear iterations due to the non linear ice-rheology,

CC BY

– ϵ_C for the convergence of the coupled problem for the prognostic experiments.

The two criteria ϵ_L and ϵ_{NL} can be increased in order to decrease the number of iterations and decrease the CPU time consumption. To quantify the influence of the convergence criteria, experiment A040 has been performed for $\epsilon_L = 10^{-7}$ up to $\epsilon_L = 10^{-2}$ and $\epsilon_{NL} = 10\epsilon_L$. As shown in Fig. 4b, the error of the solution is directly proportional to the two convergence criteria. Increasing the convergence criteria by one order of magnitude will decrease the accuracy of the solution by one order of magnitude. Most of the win in CPU is raised by the decrease of the number of the non-linear iterations as indicated in Fig. 4a.

In all experiments submitted $\epsilon_L = 10^{-6}$ was set for the 3-D experiments (a direct method was used for the 2-D experiments) and $\epsilon_{NL} = 10^{-5}$ was taken to be the non-linear convergence criterion for 2-D and 3-D non-linear problems.

3.3 Numerical methods

Numerical stabilisation of the Stokes equation was achieved by different methods:

– quadratic *P2P1* elements using quadratic velocity basis function and linear pressure basis function,

– linear elements stabilised using the residual free bubbles method (Baiocchi et al., 1993),

– linear elements using the stabilisation method by Franca and Frey (1992).

For a given number of elements, the first two methods conduct to the same size of the linear system, whereas for the third the linear system is just half the size. Nevertheless, we showed that in order to maintain the accuracy of the solution, the number of elements should be increased for the last method. The two first methods conduct to the same results in the case of linear (or plane) boundaries (as for experiments A, B, C and D), but the number of mesh nodes for the bubbles is half of that one for the *P2P1*

CC BY

method, so that tests A, B, C, D and F were performed using the bubbles method. Since the domain boundaries in experiment E are not straight lines, linear elements associated to the bubbles method have occurred to be less adapted to this problem than non-linear elements (see Sect. 4.3).

3.4 Parallel runs

Parallel runs were performed for experiment A with $L=5$ km. The mesh partitioning was created such that the associated nodes of the periodic boundary conditions belong to the same partition. A good partitioning should minimise the number of sheared nodes in order to decrease the information exchange between the processors. It is found that the efficiency of the parallelization decreases as the number of sheared nodes, relative to the total number of nodes, increases. A vertical partitioning was found not to converge, certainly because in that case the number of sheared nodes is too large. As shown in Fig. 3, when using an horizontal partitioning, the efficiency is close to 1, i.e. the decrease of the CPU time consumption is approximately proportional to the number of partitions.

4 Specific characteristics of experiments

The aim of this part is to present some specific results obtained for each of the experiments. A complete description of all the ISMIP-HOM experiments can be found in Pattyn and Payne (2006) (http://homepages.ulb.ac.be/~fpattyn/ismip/) and in Pattyn et al. (2008).

4.1 Experiments A and B

Results of experiments A are shown in Figs. 1 to 7 of the complementary material file (http://www.the-cryosphere-discuss.net/2/75/2008/tcd-2-75-2008-supplement.zip) and the settings for the submitted experiments are given in Table 3.

CC BY

Results of experiments B are shown in Figs. 8 to 12 of the complementary material file (http://www.the-cryosphere-discuss.net/2/75/2008/tcd-2-75-2008-supplement.zip), as well as the results for Experiment A in the plane $y=L/4$. Settings for the submitted experiments B are given in Table 2.

As shown in Fig. 2 of the complementary material file, the shape of the horizontal surface velocity for $L=5$ km is totally different of that for the five other cases. The surface velocity is larger over the bump, anti-correlated with the ice thickness. This can be explained by the mass conservation: variation of horizontal flux cannot be balanced anymore by vertical flux at the free surface since the vertical velocity would be too large for the given flow depth. Therefore, the horizontal flux is more or less constant inducing larger velocity for smaller depth, and vice-versa. This feature is not observed in 3-D certainly because one part of the ice can flow around the obstacle by its lateral sides.

The maximum of the absolute value of $\Delta p(z_b)=p(z_b)-p_h(z_b)=p(z_b)-\rho g H$ decreases with increasing L. For $L=160$ km, the isotropic pressure at the base is very close to the hydrostatic pressure ($|\Delta p(z_b)|<4$ kPa, to be compared to the mean basal pressure $p\approx 9000$ kPa).

For the 3-D tests A, the output variables plotted in the plane $y=L/4$ are very close to that of the 2-D tests B. This indicates that the third direction y, perpendicular to the mean flow direction, does not play a significant role for the flow. Even if the bedrock gradient in the y- and x-directions are of the same order of magnitude, the 3-D flow conditions are not far from the plane-strain assumption made for the 2-D flow. This difference between the curves decreases with increasing length of the domain L.

4.2 Experiments C and D

Results of experiments C are shown in Figs. 13 to 21 of the complementary material file and the settings for the submitted experiments are given in Table 3.

Results of experiments D and C in the plane $y=L/4$ are shown in Figs. 22 to 27 of the complementary material file. Settings for the submitted experiments D are given in Table 2. As for experiment B, the shape of the horizontal surface velocity for $L=5$ km

CC BY

differs from the five others. The same explanation as in experiment B still applies to this test.

The discontinuity at $x=3L/4$ observed on most of the $\Delta p(z_b)$ curves is difficult to explain (see Fig. 27 on the complementary material file). It was found whatever the mesh discretization. A mesh with no node exactly at the position $x=3L/4$ revealed the same discontinuity. This discontinuity also appears for the 3-D experiments C (see Fig. 27 on the complementary material file). More detailed study of the deviatoric stress fields over the whole domain indicates that the horizontal normal component of the deviatoric stress changes its sign at $x=L/4$ and $x=3L/4$ from bed to surface, and that its horizontal gradient at this point is very high, as shown on Fig. 5. Contrary to experiment B, the absolute value of $\Delta p(z_b)$ is not a decreasing function of L. For all the tests, the maximum absolute value of $\Delta p(z_b)$ is approximately 30 kPa, which is about 0.3% of the basal hydrostatic pressure.

Note that, since the basal slope of the bedrock is very very small ($\boldsymbol{n}_b \approx [0.0017\ 0\ 0.999998]$), the shear stress $S_{xz}(z_b)$ is very close to the basal drag $\tau_b = \sigma_{nt}|_{z_b} = \boldsymbol{t}_b . \boldsymbol{\sigma} \boldsymbol{n}_b$. This explains why the minimal value of $\sigma_{xz}|_{z_b}$ is very close to zero for all the tests.

A good method to verify the model results is to compare the component of the load caused by the gravity force parallel to the bed $\rho g 1000L \sin(0.1^\circ)$ to the integral of the basal drag over the whole bedrock $\int_0^L \sigma_{nt} dx$. Obviously, these two quantities should cancel each other. For all the six experiments, this balance was found to be fulfilled with a relative error of less than 0.6%.

4.3 Experiment E

Experiment E000 (Arolla flow line without slip section) was run on all five grids presented in Sect. 2.4. Comparison of the CPU time spent for read-in of data and solution of the flow problem are depicted in Fig. 6.

The corresponding results for the Cartesian components of the surface velocities as well as the Cartesian component $\tau_{xy}(z_b)$ of the deviatoric stress tensor and the

CC BY

pressure difference $\Delta p(z_b)$ at the bedrock are shown in Fig. 7.

The results show a quantitative equal behaviour of the surface velocity field between the five different meshes. Nevertheless, a higher frequency as well as amplitude of the oscillation in the stress and pressure curves obtained at the bedrock can be seen on the grids using more than just one mesh interval between the supporting points of the geometry. Comparison to a run with second order elements on the mesh `E_20_20_2nd` revealed that this is a numerical artifact caused by the use of first order elements. Figure 8 clearly shows, that the curve for $\Delta p(z_b)$ for the second order elements is missing these oscillations.

The linear interpolation in between the data points used to construct the mesh is certainly responsible for these oscillations. These findings lead to the conclusion that not only the numerical method applied to the equation system, but also rather the method used to obtain a computational grid from a digital elevation model (DEM) can influence the characteristics of the solution. In our case we stick to the instructions of the test coordinator to use linear connections along bedrock and free surface between the given DEM points. We agree that this linear interpolation was necessary for an objective comparison of all the models with different node numbers on the boundaries. For other applications, we would recommend to non-linear interpolation between the given DEM points. In Zwinger et al. (2007), we chose Nonuniform Rational B-Splines (Les and Wayne, 1995) interpolation in order to get a smooth bedrock geometry, which enabled us to do the runs using linear elements.

As one can deduce from the horizontal position of the inserted markers from the result obtained with the mesh `E_1_20` (where the basal mesh points coincide with the supporting points of the geometry), the peaks in the solution obtained with the linear element mesh `E_20_20` occur exactly at these supporting points of the geometry. The most evident explanation hence is that the kinks in the bedrock topography occurring at mesh points (due to the interconnection using straight lines) are responsible for these oscillations. This phenomenon is omitted using the second order element functions, since the first derivative of the test functions does not show unsteadiness and the

kinks are being smoothed out.

Experiment E001 (Arolla flow line with slip section) was run on three of the grids. The results are displayed in Fig. 9. Here the situation concerning the oscillations of the stress and pressure curves is even more pronounced, as singular points, where a sudden change from no-slip to slip is imposed, are imposed on the bedrock boundary condition. Again, an additional run on the mesh `E_20_20` revealed that – apart from the peaks at the singular points – oscillations can be avoided by applying second order elements. The comparison between different meshes and element types is depicted in Fig. 10.

The submitted experiments have been run using the setting summarised in Table 4.

4.4 Experiment F

Results of experiments F are shown in Figs. 30 and 31 of the complementary material file and the settings for the submitted experiments are summarised in Table 5.

Since the central area over the bump is the zone of interest, the horizontal element size is increased from the centre to the lateral boundaries with a factor 2. To ensure the mass conservation, periodic boundary conditions are applied between the upstream and downstream boundaries. For the two lateral boundaries, a null flux is prescribed and the periodicity for the velocity within the plane of the boundary.

The steady state solution is reached by running a time dependent simulation with constant boundary conditions using a one year time step. As expected, the surface elevation velocity is a decreasing function of time. As shown in Fig. 11b (thick lines), the CPU consumption needed for one time step is a decreasing function of time, indicating that the number of needed coupled iterations during one time step is also decreasing with time (from 6 to 1). The two tests were stopped when the norm of the relative difference between two time step of all the variables (velocity and surface elevation) was lower than 10^{-4}. As shown in Fig. 11b (thin lines), it does correspond to a maximal vertical velocity of the surface of few centimetres per year. It took a time-integration equivalent to 197 and 228 years to reach this convergence criteria for experiments

F000 and F001, respectively.

Since there is no accumulation on the surface, the initial volume V_0=9.9685×10^{12} m^3 should be conserved. As shown in Fig. 11a, for both experiments, the relative change of volume is of the order of the precision (<10^{-6}).

5 Conclusions

All the ISMIP-HOM benchmark tests have been performed using the same finite element code Elmer. Comparison of the results with the other full-Stokes codes indicates a very good agreement and, consequently, a relatively high level of confidence can be attributed to our results. Some FE technical points, such as mesh discretization, numerical methods and parallel issue have been discussed. The CPU time consumption for all the tests have been studied as a function of the model inputs.

A shown in Pattyn et al. (2008), the Elmer results are very close to that one of other full-Stokes model for each tests. In Pattyn et al. (2008), it is also shown that the observed difference between the models are the result of the model hypothesis (from SIA to full-Stokes), and not caused by the choice of the numerical methods and their implementation. Models built on the same hypothesis give similar results. For the full-Stokes models, this clearly indicates that the results are very close to the exact solution of the problem. As complementary material, the output figures for all the tests are attached to this paper, as well as the original output files in the same format asked by the ISMIP-HOM conveners. One aspect of this article is then to provide detailed results of all ISMIP-HOM runs obtained with Elmer. Such material should be used in the future in helping the developments of new higher-order models, by allowing an easy comparison with some existing well documented benchmark tests.

Acknowledgements. We would like to express our gratitude to J. Ruokolainen for his support in numerical matters.

References

Baiocchi, C., Brezzi, F., and Franca, L. P.: Virtual bubbles and the Galerkin least squares method, Comp. Meths. Appl. Mech. Engrg., 105, 125–141, 1993. 80, 86

Davis, T. A.: A column pre-ordering strategy for the unsymmetric-pattern multifrontal method, ACM Trans. Math. Software, 30, 165–195, 2004. 82

Franca, L. P. and Frey, S. L.: Stabilized finite element methods: II. the incompressible Navier-Stokes equations, Comput. Methods Appl. Mech. Eng., 99, 209–233, 1992. 80, 86

Franca, L. P., Frey, S. L., and Hughes, T. J. R.: Stabilized finite element methods: I. application to the advective-diffusive model, Comput. Methods Appl. Mech. Eng., 95, 253–276, 1992. 81

Kelley, C.: Iterative methods for linear and nonlinear equations, Frontiers in Applied Mathematics, 16, SIAM, 1995. 82

Les, P. and Wayne, T.: The NURBS Book, 1995. 90

Pattyn, F., Perichon, L., Aschwanden, A., Breuer, B., de Smedt, B., Gagliardini, O., Gudmundsson, G. H., Hindmarsh, R., Hubbard, A., Johnson, J. V., Kleiner, T., Konovalov, Y., Martin, C., Payne, A. J., Pollard, D., Price, S., Rückamp, M., Saito, F., Souček, O., Sugiyama, S., and Zwinger, T.: Benchmark experiments for higher-order and full Stokes ice sheet models (ISMIP-HOM), The Crysosphere Discuss., 2, 111–151, 2008. 77, 87, 92

Pattyn, F. and Payne, T.: ICE SHEET MODEL INTERCOMPARISON PROJECT: Benchmark experiments for numerical Higher-Order ice-sheet Models, http://homepages.ulb.ac.be/~fpattyn/ismip/, 2006. 87

Payne, A. J., Huybrechts, P., Abe-Ouchi, A., Calov, R., Fastook, J. L., Marshall, R. G. S. J., Marsiat, I., Ritz, C., Tarasov, L., and Thomassen, M. P. A.: Results from the EISMINT model intercomparison: the effects of thermomechanical coupling, J. Glaciol., 46, 227–238, 2000. 76

Zwinger, T., Greve, R., Gagliardini, O., Shiraiwa, T., and Lyly, M.: A full Stokes-flow thermomechanical model for firn and ice applied to the Gorshkov crater glacier, Kamchatka, Ann. Glaciol., 45, 29–37, 2007. 90

Table 1. Mesh parameters for experiment E.

Identifier	ratio hori over vert.	number of nodes/cells	element type
E_1_10	1/10	541/500	linear quad/tri
E_1_20	1/20	1031/1000	linear quad/tri
E_20_20	20/20	20771/19800	linear quad/tri
E_20_20_2nd	20/20	41542/19800	2nd order quad/tri
E_unstruct	≈20 m	2405/4290	linear tri

TCD
2, 75–109, 2008

Table 2. Settings for submitted experiments A and C.

$N_x \times N_y \times N_z$	60×60×30
Element	Brick (8 nodes)
Linear Method	Iterative (BiCGStab)
Preconditioning	ILU0 (incomplete LU)
ϵ_L	10^{-6}
ϵ_{NL}	10^{-5}
Stabilisation	Bubbles

Table 3. Settings for submitted experiments B and D.

$N_x \times N_z$	240×120
Element	Quadrilateral (4 nodes)
Linear Method	Direct
ϵ_{NL}	10^{-5}
Stabilisation	Bubbles

Table 4. Settings for submitted experiment E.

Mesh	`E_20_20_2nd`
Element	2nd order quad/tri
Linear Method	Direct
ϵ_{NL}	10^{-5}
Stabilisation	Stabilised method

Table 5. Settings for submitted experiment F.

$N_x \times N_y \times N_z$	60×60×40
Element	Brick (8 nodes)
Linear Method	Indirect ILU0
ϵ_L	10^{-5}
ϵ_C	10^{-4}
Stabilisation	Bubbles
Time step	dt=1 a

TCD
2, 75–109, 2008

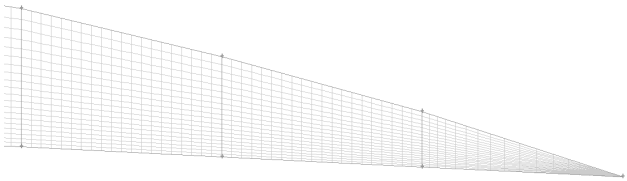

Fig. 1. Three sections (horizontal length 100 meter) of the lower end of Arolla glacier geometry flow-line (experiment E) model with mesh (E_20_20) applied.

TCD
2, 75–109, 2008

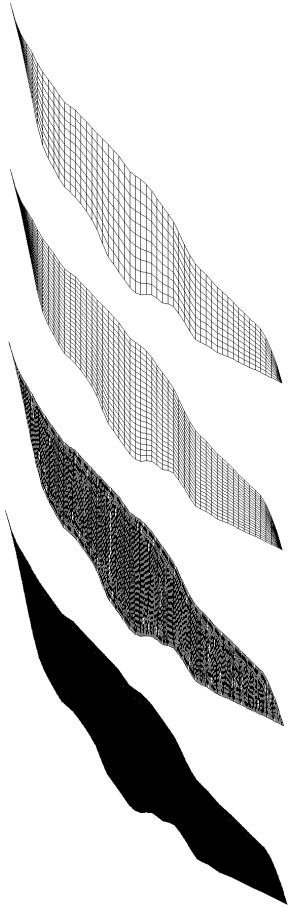

Fig. 2. The meshes `E_1_10` (up,left), `E_1_20` (up,right), `E_unstruct` (bottom,left) and `E_20_20` (bottom,right). Vertical scales are 10 times exaggerated.

TCD
2, 75–109, 2008

CC BY

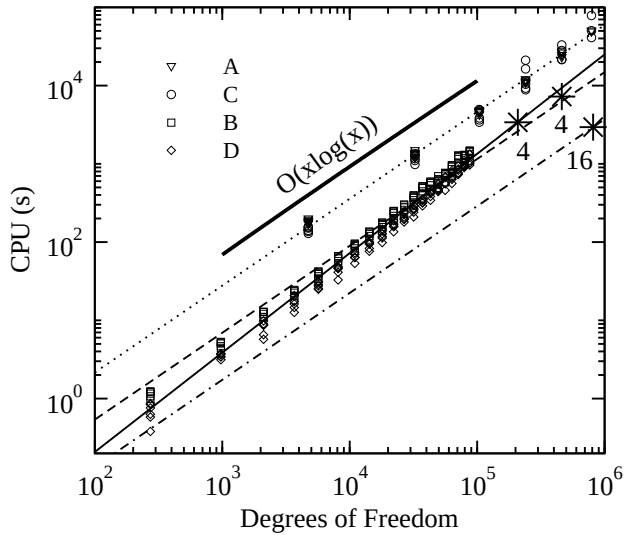

Fig. 3. CPU time consumption (s) as a function of the total degrees of freedom, for 2-D and 3-D experiments A, B, C and D. Power regressions $y=0.0006x^{1.27}$ and $y=0.013x^{1.11}$ for 2-D and 3-D experiments are plotted using straight and dotted lines, respectively. Dashed and dash-dot lines are parallel to the 3-D regression line with a power law factor divided by 4 and 16, respectively. The stars associated with the number indicate results for parallel runs for experiment A with $L=5$ km and the corresponding number of processors.

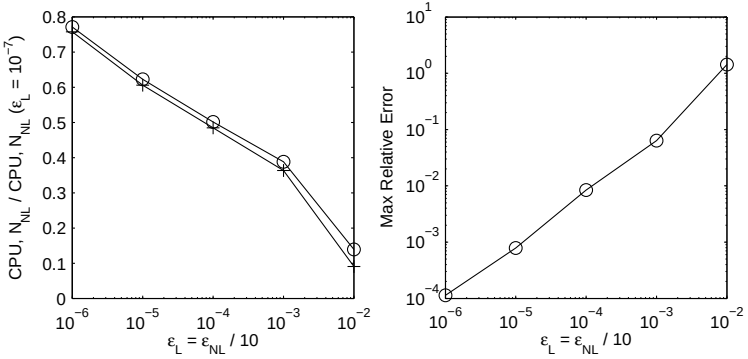

Fig. 4. **(a)** CPU time consumption (circle) and number of non-linear iterations N_{NL} (cross) relative to the reference run ($\epsilon_L = 10^{-7}$) and **(b)** maximum of the relative difference with the reference run ($\epsilon_L = 10^{-7}$) on the velocities as a function of $\epsilon_L = \epsilon_{NL}/10$, for experiment A040.

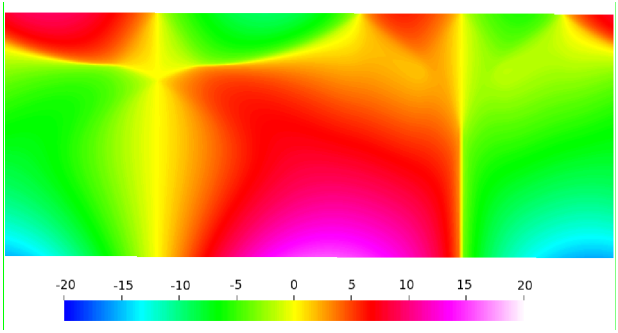

Fig. 5. S_{xx} [kPa] for experiment D005.

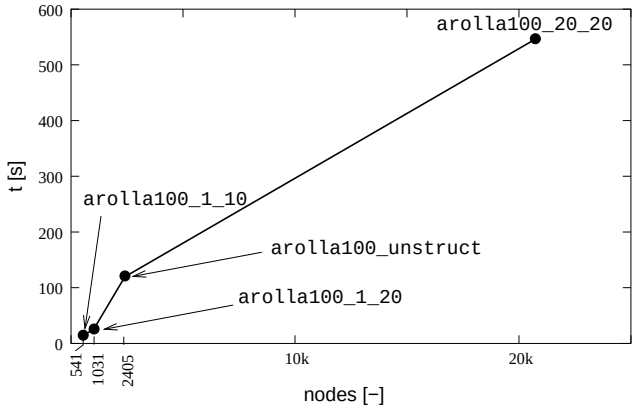

Fig. 6. Experiment E000: CPU time consumption.

CC BY

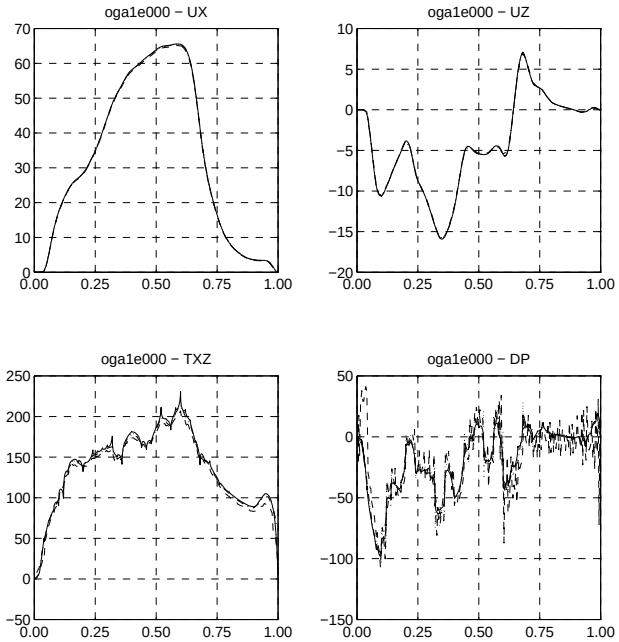

Fig. 7. Experiment E000 on mesh `E_1_20` (dash-dot line), `E_unstruct` (dotted line), `E_20_20` (dashed line) and `E_20_20_2nd` (solid line).

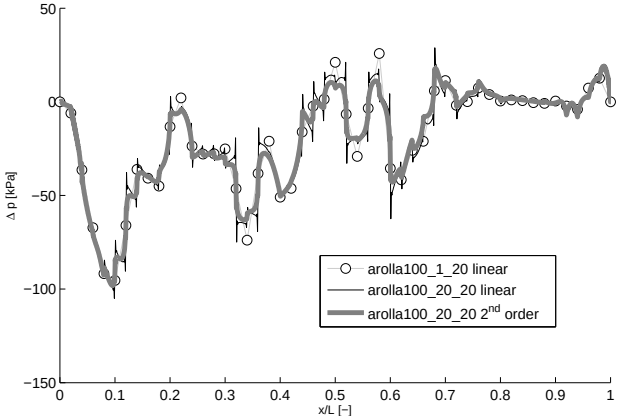

Fig. 8. Experiment E000: comparison of $\Delta p(z_b)$ along the bedrock between linear and second-order element runs.

TCD
2, 75–109, 2008

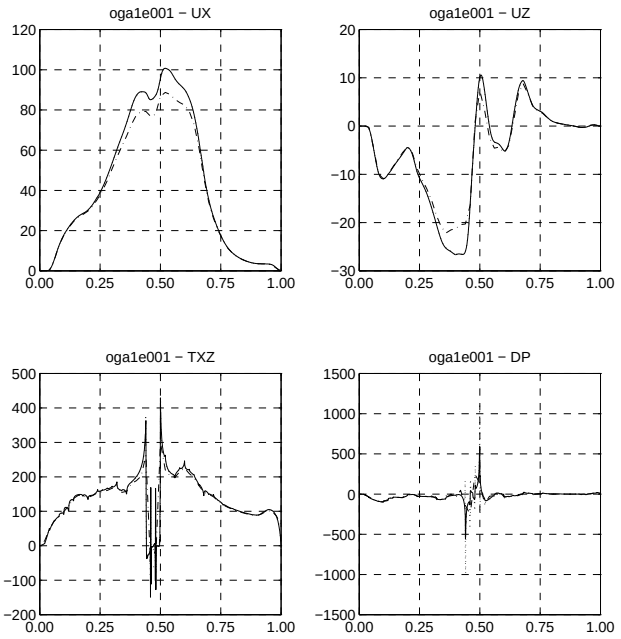

Fig. 9. Experiment E000 on mesh `E_1_20` (dash-dot line), `E_20_20` (dashed line) and `E_20_20_2nd` (solid line).

TCD
2, 75–109, 2008

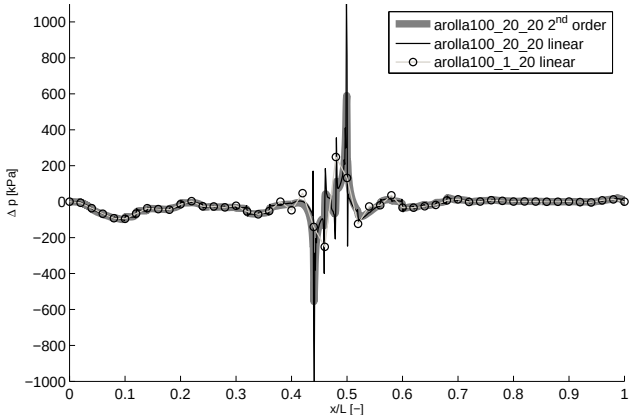

Fig. 10. Experiment E001: comparison of $\Delta p(z_b)$ along the bedrock between linear and second-order element runs.

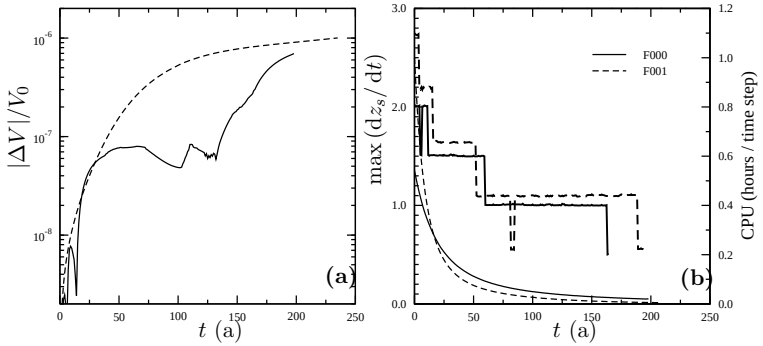

Fig. 11. **(a)** Change of total ice volume relative to the initial volume and **(b)** maximum value of the surface change $\mathrm{d}z_s/\mathrm{d}t$ (thin lines) and CPU time consumption per one time step (thick lines) as a function of the time t for experiments F000 (solid lines) and F001 (dashed lines).

CC BY