

Towards a fully unstructured ocean model for ice shelf cavity environments: Model development and verification using the Firedrake finite element framework

William I. Scott^{a,b,*}, Stephan C. Kramer^a, Paul R. Holland^c, Keith W. Nicholls^c,
Martin J. Siebert^{a,b,d}, Matthew D. Piggott^{a,b}

^a Science and Solutions for a Changing Planet DTP, Department of Earth Science and Engineering, Imperial College London, Prince Consort Road, South Kensington, London, SW7 2BP, UK

^b Science and Solutions for a Changing Planet DTP, Grantham Institute for Climate Change, Imperial College London, Exhibition Road, South Kensington, London, SW7 2AZ, UK

^c British Antarctic Survey, High Cross, Madingley Road, Cambridge, CB3 0ET, UK

^d Tremough House, University of Exeter, Penryn, Cornwall, TR10 9FE, UK

ARTICLE INFO

Keywords:

FEM
Ice shelf cavities
Firedrake
Adjoint
Verification
MMS

ABSTRACT

Numerical studies of ice flow have consistently identified the grounding zone of outlet glaciers and ice streams (the region where ice starts to float) as crucial for predicting the rate of grounded ice loss to the ocean. Owing to the extreme environments and difficulty of access to ocean cavities beneath ice shelves, field observations are rare. Estimates of melt rates derived from satellites are also difficult to make near grounding zones with confidence. Therefore, numerical ocean models are important tools to investigate these critical and remote regions. The relative inflexibility of structured grid models means, however, that they can struggle to resolve these processes in irregular cavity geometries near grounding zones. To help solve this issue, we present a new nonhydrostatic unstructured mesh model for flow under ice shelves built using the Firedrake finite element framework. We demonstrate our ability to simulate full ice shelf cavity domains using the community standard ISOMIP+ *Ocean0* test case and compare our results against those obtained with the popular MITgcm model. Good agreement is found between the two models, despite their use of different discretisation schemes and the sensitivity of the melt rate parameterisation to grid resolution. Verification tests based on the Method of Manufactured Solutions (MMS) show that the new model discretisation is sound and second-order accurate. A main driver behind using Firedrake is the availability of an automatically generated adjoint model. Our first adjoint calculations, of sensitivities of melt rate with respect to different inputs in an idealised grounding zone domain, are promising and point to the ability to address a number of important questions on ocean influence on ice shelf vulnerability in the future.

1. Introduction

Recent observations have shown that outlet glaciers and ice streams in West Antarctica are retreating at an alarming rate (Mouginot et al., 2014; Scambos et al., 2017). In total there is enough ice in West Antarctica to raise sea level by approximately 3.5 m (Fretwell et al., 2013). Evidently, this has serious implications for communities that live in low lying coastal regions. There are still major uncertainties, however, in predictions of how much ice could be lost and, importantly, on what time scales this may occur. Increased basal melting beneath ice shelves, due to increased ocean heat, is thought to be the main cause for ice loss in West Antarctica (Rignot et al., 2013). Wherever ice

shelves are topographically confined, they provide buttressing forces to the grounded glaciers (Shepherd et al., 2018). When ice shelves thin, a decrease in buttressing can lead to an increase in ice flow speeds and hence the rate at which grounded ice is lost increases. Predicting sea level rise is a pressing issue, but there are also questions about how changes in melt water flux to the ocean can affect global ocean dynamics (Dinniman et al., 2016). Improving understanding of ocean processes beneath ice shelves and their impact on ice shelf vulnerability is therefore an important problem in glaciology and climate science.

Numerical models of ocean flow beneath ice shelves play a key role in our understanding of ocean cavities, since direct observations are so limited. The model needs to resolve the interaction between

* Corresponding author at: Science and Solutions for a Changing Planet DTP, Department of Earth Science and Engineering, Imperial College London, Prince Consort Road, South Kensington, London, SW7 2BP, UK.

E-mail address: wis15@ic.ac.uk (W.I. Scott).

<https://doi.org/10.1016/j.ocemod.2023.102178>

Received 30 September 2022; Received in revised form 30 January 2023; Accepted 8 February 2023

Available online 12 February 2023

1463-5003/© 2023 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license

(<http://creativecommons.org/licenses/by/4.0/>).

salinity, temperature and the flow, along with the thermodynamics at the complex, sloping ice-ocean interface that characterise ice shelf cavity environments. Melting ice injects buoyant, fresh water into the domain, which in turn drives overturning ocean circulation within the cavity. The freezing point of water decreases with depth so, in some cases, ice that melts at depth in the cavity can refreeze closer to the ocean surface, thus transporting ice up the underside of the ice shelf. This is known as the ‘ice pump’ mechanism (Lewis and Perkin, 1986; Jenkins and Bombosch, 1995).

In coastal regions where changes in coastlines and bathymetry can have important effects on dynamics, horizontally unstructured grid models have proven to be extremely useful tools (Kärnä et al., 2018). In an ice shelf cavity setting, this is doubly important since the base of the ice shelf can vary as much as the seafloor, with the formation of basal crevasses and channels that evolve with the ocean, and indeed ice, flow. The lack of flexibility in traditional structured grid ocean models in resolving these features means that there is still large uncertainty in flow dynamics beneath ice shelves (Dinniman et al., 2016). This is compounded by the fact that results using the commonly employed ‘three-equation melt parameterisation’ depend significantly on vertical resolution and implicitly on the model’s choice of vertical discretisation (Gwyther et al., 2020) — an issue that will be explored further in Section 3.2.

Satellite measurements indicate that spatial melt rate patterns near grounding zones (the region separating the grounded ice sheet from the floating ice shelf) are highly variable (Milillo et al., 2019). This implies that there must be a complex combination of processes in the grounding zone that need to be resolved to accurately model melt rates. Due to the limitations on grid resolution imposed by structured grids used by traditional ocean models, it has not been possible to investigate in detail melting at grounding zones. Ensemble simulations from glacial flow models have shown that the largest uncertainties in projected ice loss come from estimating melt rates at the grounding zone (Arthern and Williams, 2017). Similarly, Goldberg et al. (2019) used an ice flow model with an associated adjoint model to investigate how sensitive the volume of ice loss was to spatially varying melt for Dotson and Crosson ice shelves. Again, they found that ice loss was most sensitive to melting at the grounding zone. An extension of this study to the wider Amundsen Sea Embayment confirmed the high sensitivity of ice loss to melting at the grounding zone (Morlighem et al., 2021). Therefore, accurately simulating ocean dynamics near the grounding zone seems to be one of the most important tasks required to help reduce uncertainty in ice loss estimates, and consequently reduce uncertainty in sea level projections.

The finite element method is particularly well suited to solving problems in complicated domains (Elman et al., 2014). The power of the finite element method comes from being able to define the solution using piecewise functions which can be defined on completely arbitrary meshes. This is ideal for modelling flow in complicated domains, such as ice shelf cavities. A fully unstructured finite element ocean model, Fluidity (Piggott et al., 2008), was previously adapted to enable simulation of ice shelf cavities (Kimura et al., 2013). Due to its ability to run on fully unstructured grids, Fluidity was able to resolve all the key features of an ice shelf cavity (pinching grounding lines, sloping ice base and the steep vertical ice front; Kimura et al. 2013), as well as investigate flow within basal crevasses (Jordan et al., 2014).

The model presented here is similarly motivated by the applicability of finite elements on unstructured meshes and the need to retain a full representation of physics, valid at order one aspect ratios, for buoyancy driven flows in domains with complicated geometries, such as ice shelf cavities. A significant departure from previous work is use here of the Firedrake toolkit to implement the underlying discretisation of the finite element method (Rathgeber et al., 2016). The philosophy behind Firedrake has several advantages. By separating the underlying software implementation from the physics of the system, it enables a Firedrake user to quickly develop and test new discretisation methods

on the problem in question. The syntax of the Unified Form Language (UFL) Firedrake employs is designed to closely mimic the mathematical description of the weak form prescribing the finite element discretisation (Alnæs et al., 2014). This, combined with the readability of the Python programming language, helps to improve the longevity of the code by reducing the learning curve required for new users before they can start using and implementing features in the model. ‘Icepack’, a glacier flow modelling package built on top of Firedrake, is a good example of this (Shapiro et al., 2021). The Firedrake framework automatically translates this high-level mathematical description of the numerics into highly optimised low-level C code which leads to an efficient implementation of the finite element method. In addition, Firedrake uses the Portable, Extensible Toolkit for Scientific Computation (PETSc) a state-of-the-art library to efficiently solve the systems that arise from the finite element discretisation (Balay et al., 1997).

There are already ocean models that use horizontally unstructured grids to investigate ice shelf cavities, such as FESOM (Timmermann et al., 2012), FVCOM (Zhou and Hattermann, 2020) and MPAS-O (Ringler et al., 2013; Gwyther et al., 2020). These models typically retain a distinct difference in how they deal with horizontal and vertical motion in the model. In part this means making the hydrostatic approximation: provided the horizontal length scales are much larger than the vertical, the vertical acceleration terms can be neglected, which reduces the computational cost. This approximation is evidently true for the open ocean when typically the meshes used for simulations have horizontally stretched, anisotropic grids. To investigate interesting basal features under ice shelves though, such as crevasses (Jordan et al., 2014) or in the proximity of topographically complex grounding zone regions, one may want to have more flexibility by using meshes where the aspect ratio of the grid becomes close to order one, i.e., the horizontal grid sizes are comparable to the vertical grid cell sizes. In this case the hydrostatic approximation may no longer be valid. Therefore we include the nonhydrostatic vertical acceleration terms as part of the solve. Moreover, by treating the horizontal and vertical equations on the same footing, our model has been designed so that it can easily incorporate isotropic turbulence closure schemes from the CFD literature, similar to work carried out by Yeager (2018), whilst at the same time able to transition to more traditional GFD turbulence closures in high aspect ratio domains, where vertical and horizontal mixing is parameterised separately.

Probably the main advantage of using a framework such as Firedrake, or similar projects such as Fenics (Alnæs et al., 2015), is the fact that it readily facilitates the use of an automatically generated adjoint model, thanks to the work of the dolfin-adjoint project (Mitusch et al., 2019). Adjoint modelling represents an efficient means to calculate gradients of model output functionals with respect to model inputs (Errico, 1997). By solving the adjoint system the gradient of an output functional can be calculated with respect to any number of input parameters, independent of the number of these input parameters. Adjoint modelling is often the only tractable method for calculating gradient information for numerical models that rely on grid-based discretisations, like ocean models, because the number of input parameters scales with the grid resolution (Errico, 1997). Early implementations of adjoint models were developed by hand separately from forward models (Kalnay, 2003). As the complexity of the forward code develops this approach becomes time consuming and error prone. A more sophisticated method uses what is sometimes referred to as automatic differentiation (AD) of the forward code, by using the chain rule on each line of forward code repeatedly to generate an adjoint model (Errico, 1997). Usually, careful intervention by the user is required during development and maintenance, especially with regards to ensuring the computational efficiency of the adjoint model. Hence it is common to refer to this process as algorithmic, not automatic, differentiation. The approach implemented by the dolfin-adjoint project makes use of the high-level symbolic representation of the discrete mathematical problem written in UFL (Mitusch et al., 2019). Since the

mathematical problem is separate from its software implementation, provided the mathematical form of the equations is differentiable then the process is automatic. The adjoint model is consistent with the forward discretisation and automatically inherits the solver strategies and the parallelisability of the forward model (Mitusch et al., 2019).

Sensitivities of a functional with respect to different inputs are very useful in their own right (Errico, 1997). In a time-dependent forward model, the equivalent adjoint model propagates sensitivities backwards through time, from the final time to the initial time. This complementary information can be very useful when trying to understand modelled processes that may not be obvious from considering the forward model alone. Gradient information is also very useful for optimisation problems, in particular parameter estimation and data assimilation. In many ocean modelling studies there are many parameters that are unknown, but can be constrained by observations. Gradient information can be used to update these parameters in an efficient manner by minimising a functional that measures the error between observed and modelled values.

There have been a limited number of studies of an ocean adjoint in an ice shelf cavity context. The first used the MITgcm ocean model to simulate the Pine Island Ice Shelf cavity (Heimbach and Losch, 2012). Goldberg et al. (2020) extended MITgcm's adjoint model, such that it was able to calculate the sensitivity of melt rate to bathymetry. Notably, this was using the open source AD tool, openAD (Utke et al., 2008). Most recently the ECCO2 assimilation framework (which also relies on MITgcm's adjoint capability) was applied to the Amundsen and Bellingshausen seas and ice shelf cavities (Nakayama et al., 2021b). Since all the previous adjoint modelling work for ice shelf ocean cavities have used the MITgcm framework, they inherit the relative inflexibility of MITgcm's structured grid. Although MITgcm's resolution can vary in space (dx can vary horizontally) dz is limited to varying with depth. This has limited practical application for ice shelf cavities since the ocean cavity thickness varies significantly from the grounding zone to the open ocean. We see our model, which can combine flexible unstructured meshes with the adjoint capability, as a tool with significant potential for investigating ocean conditions in the complex domains that are ice shelf cavities.

The outline for the rest of this paper is as follows: Section 2 covers the model discretisation choices. In Section 3.1 the accuracy of the second-order discretisation is verified using a Method of Manufactured Solution-based test which includes melting. We demonstrate the resulting model's capability to run simulations for the 3D ISOMIP+ test case in Section 3.2 and compare results with the MITgcm ocean model. This includes an investigation into the sensitivity of the melt parameterisation to vertical mixing and resolution choice. Finally, we present preliminary adjoint sensitivity results in an idealised domain in Section 3.3. These steps are necessary to prove the capabilities and accuracy of the model, so that future work can tackle the challenging problem of simulating ocean dynamics and melt rates near grounding zones.

2. Model description

2.1. Model equations

As emphasised in Section 1, one of the key motivations is to develop a model with the ability to treat vertical dynamics in the same way as horizontal dynamics, when the mesh cells become close to isotropic, on fully unstructured meshes. This means we solve the full incompressible Navier–Stokes equations for velocity and pressure, and we do not make the hydrostatic approximation.

Conservation of momentum (in strong form) for a rotating fluid under the Boussinesq approximation can be written as

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} + 2\boldsymbol{\Omega} \times \mathbf{u} = -\frac{1}{\rho_0} \nabla p + \nabla \cdot \boldsymbol{\tau} + \mathbf{F}, \quad (1)$$

where

$$\boldsymbol{\tau} = \nu \cdot [\nabla \mathbf{u} + (\nabla \mathbf{u})^T]. \quad (2)$$

The rotation vector of the Earth is given by $\boldsymbol{\Omega}$ and in the applications presented in this work Coriolis is implemented using an f -plane assumption, with $f = -1.409 \times 10^{-4} \text{ s}^{-1}$, corresponding to a latitude of 75° S (Vallis, 2017). ρ_0 is the reference density of sea water, taken as $1027.51 \text{ kg m}^{-3}$ following Asay-Davis et al. (2016). Note Table 1 summarises values for general model parameters. Kinematic eddy viscosity, ν , is a rank two tensor and can be spatially variant. \mathbf{F} is a generic vector body force. The velocity, $\mathbf{u} = (u, v, w)^T$, and pressure, p , can be solved for using the continuity equation as a constraint to enforce the incompressibility condition

$$\nabla \cdot \mathbf{u} = 0. \quad (3)$$

In this work we have not modified the continuity equation to allow for a free surface implementation. The top surface of the model is treated as a rigid lid. Flows with more complicated tidal forcing may require a free surface to accurately capture the dynamics though and we anticipate implementing this feature in future work.

To complete the description of the problem boundary conditions need to be imposed, which are given as follows

$$\mathbf{u} = \mathbf{u}_D = u_n \mathbf{n} + u_t \mathbf{t} \quad \text{on } \Gamma_D, \quad (4)$$

$$\mathbf{u} \cdot \mathbf{n} = u_n, \quad \tau_{nt} = f_{\text{tang}} \quad \text{on } \Gamma_{D_n}, \quad (5)$$

$$-\rho_0 \tau_{nn} + p = p_{\text{ext}}, \quad \tau_{nt} = 0 \quad \text{on } \Gamma_{p_{\text{ext}}}, \quad (6)$$

Unit normal and tangential vectors are given by \mathbf{n} and \mathbf{t} respectively. On Γ_D Dirichlet boundary conditions for both the normal and tangential components of velocity, u_n and u_t , are specified. This type of boundary can be used for rigid walls by setting $u_n = u_t = 0$ for no-normal, no-slip conditions. Although these boundaries are not used in this work they are included for completeness to ensure the correctness of the weak forms in Section 2.2. In this work all walls are represented using Γ_{D_n} boundaries. On these boundaries only the normal component of velocity is specified and is combined with a condition for the shear stress in the tangential direction

$$\tau_{nt} = \mathbf{n} \cdot \boldsymbol{\tau} \cdot \mathbf{t}. \quad (7)$$

In this work all side walls are free slip boundaries. These are imposed by setting f_{tang} and u_n to zero. On the top and bottom boundaries, namely the floating ice shelf and the seabed, we apply quadratic wall drag

$$f_{\text{tang}} = C_D |\mathbf{u}| \mathbf{u} \cdot \mathbf{t} \quad (8)$$

where C_D is a drag coefficient.

For boundaries, $\Gamma_{p_{\text{ext}}}$, only an external pressure is specified in (6), with τ_{nn} the normal component of the stress vector on the boundary given by

$$\tau_{nn} = \mathbf{n} \cdot \boldsymbol{\tau} \cdot \mathbf{n}. \quad (9)$$

Although we do not consider these types of boundaries in this paper we include their formulation for completeness to ensure that the surface integrals in the weak forms of Section 2.2 are consistent. These types of boundaries are necessary in open domains, where an external hydrostatic pressure field is applied to balance the stress on either side of the open boundary, thus preventing fluid from 'falling' out of the domain.

To incorporate buoyancy effects we make the Boussinesq approximation, where perturbations in density are neglected in all terms except the buoyancy term, which is implemented as a source \mathbf{F} where

$$\mathbf{F} = -\frac{\rho'}{\rho_0} g \mathbf{k}, \quad (10)$$

where $g = 9.81 \text{ m s}^{-2}$ is the gravitational acceleration and \mathbf{k} is the unit vector pointing in the vertical direction. ρ' is a density perturbation

related to temperature, T , and salinity, S , by a linear equation of state, given by

$$\rho' = \rho_0(\alpha_T(T - T_0) + \beta_S(S - S_0)), \quad (11)$$

where $\alpha_T = -3.733 \times 10^{-5} \text{ }^\circ\text{C}^{-1}$ and $\beta_S = 7.843 \times 10^{-4}$ are respectively the expansion and contraction coefficients for temperature and salinity. $T_0 = -1 \text{ }^\circ\text{C}$ and $S_0 = 34.2$ are reference temperature and salinity values used in the linear equation of state again with values from [Asay-Davis et al. \(2016\)](#). Note that a constant hydrostatic pressure term has been subtracted from the momentum equations and incorporated into the definition of pressure, so that p is a perturbation pressure

$$p = P - P_{\text{hyd}}, \quad (12)$$

where P is the full pressure and P_{hyd} is given by

$$P_{\text{hyd}} = -\rho_0 g z, \quad (13)$$

and z is the vertical coordinate.

Temperature and salinity are governed by scalar advection diffusion equations, which can be written in strong form as

$$\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T = \nabla \cdot (\boldsymbol{\kappa}_T \cdot \nabla T) + r_T(T_{\text{res}} - T), \quad (14)$$

and

$$\frac{\partial S}{\partial t} + \mathbf{u} \cdot \nabla S = \nabla \cdot (\boldsymbol{\kappa}_S \cdot \nabla S) + r_S(S_{\text{res}} - S), \quad (15)$$

where $\boldsymbol{\kappa}_T$ and $\boldsymbol{\kappa}_S$ are the (eddy) diffusivity of temperature and salinity respectively. These are spatially variant rank two tensors similar to viscosity. r_T and r_S represent restoring frequencies when relaxing the solution in a sponge region to specified values, T_{res} and S_{res} .

Boundary conditions for temperature and salinity are given as

$$T = T_D \quad \text{on } \Gamma_{D_T}, \quad (16)$$

$$(\boldsymbol{\kappa}_T \cdot \nabla T) \cdot \mathbf{n} = \Phi_T \quad \text{on } \Gamma_{\Phi_T}, \quad (17)$$

$$S = S_D \quad \text{on } \Gamma_{D_S}, \quad (18)$$

$$(\boldsymbol{\kappa}_S \cdot \nabla S) \cdot \mathbf{n} = \Phi_S \quad \text{on } \Gamma_{\Phi_S}, \quad (19)$$

where T_D and S_D are the Dirichlet boundary values for temperature and salinity respectively on walls Γ_{D_T} and Γ_{D_S} given in (16) and (18). Neumann boundary fluxes for temperature (17) and salinity (19) are given by Φ_T and Φ_S respectively. Note the definition of Φ includes diffusivity, so if the tracer is temperature, the units of Φ_T would be m K/s.

The next section gives the weak forms of (1), (3), (14) and (15) necessary for the finite element discretisation.

2.2. Weak form of model equations

2.2.1. Tracers

Choosing a suitable scalar function space, \mathcal{Q} , for a generic tracer, q , and a test function, ϕ , the weak form of the advection diffusion equation ((14) and (15)) can be written as

$$M_q + A_q + K = F_q - S, \quad \forall \phi \in \mathcal{Q} \quad (20)$$

where

$$M_q = \int_{\Omega} \frac{\partial q}{\partial t} \phi \, dx, \quad (21)$$

$$A_q = \int_{\Omega} -q \nabla \cdot (\phi \mathbf{u}) \, dx + \int_{\Gamma} \phi \bar{q} \mathbf{u} \cdot \mathbf{n} \, ds, \quad (22)$$

$$K = \int_{\Omega} \nabla \phi \cdot (\boldsymbol{\kappa} \cdot \nabla q) \, dx - \int_{\Gamma_{D_q}} \phi \mathbf{n} \cdot \boldsymbol{\kappa} \cdot \nabla q \, ds - \int_{\Gamma_{\Phi_q}} \phi \Phi_q \, ds, \quad (23)$$

$$F_q = \int_{\Omega} F \phi \, dx, \quad (24)$$

$$S = \int_{\Omega} \alpha_{\text{res}} q \phi \, dx, \quad (25)$$

where the individual terms are denoted: M_q — time derivative, A_q — advection, K — (eddy) diffusivity, F_q — source and S — sink. Dirichlet boundary values for the tracer are denoted q_D and Neumann boundary conditions as Φ_q . Note \bar{q} is the ‘advected’ scalar value defined in Section 2.4 in the context of surface integrals and Discontinuous Galerkin (DG) finite elements.

2.2.2. Velocity–pressure

To derive the weak forms of (1) and (3) we choose separate, suitable function spaces \mathcal{V} and \mathcal{W} for velocity, $\mathbf{u} = (u, v, w)^T$, and pressure, p respectively. Multiplying (1) by a vector test function $\boldsymbol{\phi} = (\phi_u, \phi_v, \phi_w)^T \in \mathcal{V}$ and integrating over the domain, Ω , gives

$$M + A + C + P + V = F \quad \forall \boldsymbol{\phi} \in \mathcal{V}, \quad (26)$$

where

$$M = \int_{\Omega} \frac{\partial \mathbf{u}}{\partial t} \cdot \boldsymbol{\phi} \, dx, \quad (27)$$

$$A = \int_{\Omega} -\mathbf{u} \cdot [\nabla \cdot (\mathbf{u} \otimes \boldsymbol{\phi})] \, dx + \int_{\Gamma} (\bar{\mathbf{u}} \cdot \boldsymbol{\phi})(\mathbf{u} \cdot \mathbf{n}) \, ds \quad (28)$$

$$C = \int_{\Omega} -f v \phi_u + f u \phi_v \, dx, \quad (29)$$

$$P = \int_{\Omega} \frac{1}{\rho_0} \boldsymbol{\phi} \cdot \nabla p \, dx - \int_{\Gamma_{p_{\text{ext}}}} \frac{1}{\rho_0} (p - p_{\text{ext}})(\boldsymbol{\phi} \cdot \mathbf{n}) \, ds, \quad (30)$$

$$V = \int_{\Omega} \nabla \boldsymbol{\phi} \cdot \boldsymbol{\tau} \, dx - \int_{\Gamma_{D \cup \Gamma_{D_n}}} \boldsymbol{\phi} \cdot (\mathbf{n} \cdot \boldsymbol{\tau}) \, ds + \int_{\Gamma_{D_n}} (\mathbf{n} \cdot \boldsymbol{\tau} \cdot \mathbf{t})(\boldsymbol{\phi} \cdot \mathbf{t}) \, ds - \int_{\Gamma_{D_n}} f_{\text{tang}} \boldsymbol{\phi} \cdot \mathbf{t} \, ds, \quad (31)$$

$$F = \int_{\Omega} \mathbf{F} \cdot \boldsymbol{\phi} \, dx, \quad (32)$$

the individual terms being: M — time derivative, A — advection, C — rotation, P — pressure gradient, V — (eddy) viscosity and F — source. Note $\bar{\mathbf{u}}$, similar to \bar{q} , is an ‘advected’ velocity also defined in Section 2.4. The symbol \otimes represents an outer product between two vector quantities.

Multiplying (3) by a test function $\psi \in \mathcal{W}$ and integrating by parts gives the weak form of the continuity equation

$$\int_{\Omega} -\nabla \psi \cdot \mathbf{u} \, dx + \int_{\Gamma_{p_{\text{ext}}}} \psi \mathbf{u} \cdot \mathbf{n} \, ds + \int_{\Gamma_{D \cup \Gamma_{D_n}}} \psi u_n \, ds = 0, \quad \forall \psi \in \mathcal{W}. \quad (33)$$

Defining ψ in the pressure space ensures that the discretised divergence and gradient operators are the negative transpose of each other and it means there are the same number of discrete continuity equations as pressure degrees of freedom. This is natural, since the pressure can be thought of as a Lagrange multiplier that enforces the incompressibility constraint ([Strang, 2007](#)).

2.3. Finite element discretisation: Fully unstructured and ‘extruded’ meshes

2.3.1. Tracers

On fully unstructured meshes (2D: triangle elements, 3D: tetrahedral elements) we use linear discontinuous basis functions (P1DG) for the tracer test and trial space, \mathcal{Q} . This is because P1DG elements are well suited to advection dominated problems ([Kärnä et al., 2018](#)).

For high aspect ratio domains, where the horizontal extent of the domain is much greater than the depth, for instance when simulating an entire ice shelf cavity, an alternative option to using a fully unstructured mesh is to make use of Firedrake’s in-built ‘extruded’ mesh feature. Extruded meshes offer increased performance and robustness when elements are highly anisotropic, yet retain the flexibility of a horizontally unstructured mesh ([Bercea et al., 2016](#); [McRae et al., 2016](#)).




Fig. 1. Schematic of 3D extruded elements used in Section 3.2. Degrees of Freedom (DOF) are represented by black dots (with arrows representing vector quantities). DOF situated on the edges of the element indicate continuity between elements, whereas DOF displaced into the element indicate discontinuity between elements. The extruded triangular prism elements are created by taking the product of a triangular element in the horizontal with an interval element in the vertical. The element on the left, used to discretise pressure, is a tensor product of P2 in the horizontal with P2 in the vertical. The element in the middle, representing the horizontal velocity component, is formed by taking a product of N2 in the horizontal with P2 in the vertical. This ensures the horizontal component of velocity is discontinuous in the normal direction but continuous in the tangential direction. The element on the right is the vertical velocity component formed from taking a product of P2 in the horizontal with P1DG in the vertical. This ensures that the vertical velocity component is discontinuous in the vertical direction but continuous in the horizontal direction.

Although we do not carry out 2D vertical slice modelling on extruded meshes in this paper we describe the 2D case to help build intuition for the 3D case, which is used in Section 3.2. An extruded 2D vertical slice mesh is made up of quadrilateral elements, which can be thought of as the tensor product of an element interval in the horizontal and an element interval in the vertical. For 3D domains the elements are triangular prisms, which are formed by the tensor product of a triangular element in the horizontal with an interval element in the vertical. This means the horizontal part of the mesh can be fully unstructured. In the vertical direction the mesh is structured. It is possible to choose between a fixed number of columns in the vertical, or a fixed cell height, as well as ‘shaving’ cells in the top and bottom columns to accurately fit the domain geometry.

Finite elements on extruded meshes are defined in terms of a tensor product of a finite element on the horizontal mesh and finite element on the vertical mesh. This approach controls how the element behaves in the horizontal direction and how the element behaves in the vertical. This is particularly useful for more ‘exotic’ element combinations that have different continuity requirements in the horizontal and vertical directions.

On extruded meshes we choose the scalar finite element which mimics the P1DG element. This is formed by taking the tensor product of P1DG on the horizontal element with P1DG on the vertical interval.

2.3.2. Velocity–pressure

The choice of a suitable finite element velocity–pressure element pair is crucial to ensuring an accurate and stable discretisation of the continuous equations (Elman et al., 2014). On fully unstructured meshes we use (vector) P1DG basis functions for the velocity space, \mathcal{V} , and quadratic continuous basis functions (P2) for the pressure space, \mathcal{W} . The resulting P1DG-P2 velocity–pressure element pair has been used with the finite element ocean model Fluidity (Kramer et al., 2010; Kimura et al., 2013; Jordan et al., 2014) and more recently implemented within the Thetis (coastally-focused) ocean model developed using Firedrake (Kärnä et al., 2018; Wallwork et al., 2020). The element pair has been shown to maintain geostrophic balance, because the skew gradient of the pressure maps to the velocity space (Cotter et al., 2009b,a; Cotter and Ham, 2011).

On extruded meshes we have chosen an element pair that has similar properties to P1DG-P2, so that as the horizontal mesh is refined the discretisation is able to deal with order one aspect ratios. Defining the equivalent of P2 on extruded meshes is relatively straightforward: it is the tensor product of a P2 horizontal element with the equivalent P2 element on the vertical interval, as shown on the left side of Fig. 1. Forming the required velocity element on an extruded mesh is more complicated though. Our solver strategy, outlined in Section 2.5, relies on the gradient of functions in the pressure function space mapping

to functions in the velocity trial space. Given our discretisation choice for pressure this places a constraint on the continuity of the velocity element. Along the element boundary the gradient of the pressure is continuous across cells in the tangential direction and varies quadratically. In the normal direction the pressure gradient is discontinuous between cells and varies linearly.

For a 2D vertical slice the horizontal component of this velocity element, u , is formed by taking the tensor product of P1DG on a horizontal interval and P2 on a vertical interval. This ensures that the horizontal component of velocity is discontinuous in the horizontal direction, but continuous between elements on vertical facets. The vertical component of the velocity, w , is defined the other way round, as a tensor product of P2 in the horizontal and P1DG in the vertical. Again, this ensures that the vertical velocity is continuous in the horizontal direction, but discontinuous in the vertical direction.

The 3D equivalent of the 2D velocity element is built up as follows: the horizontal velocity components, u and v , are formed by taking the tensor product of a Nedelec element of the second kind (N2) defined on a triangular element with a P2 element in the vertical. N2 elements are inherently (2D) vector elements, that are tangentially continuous between cells and discontinuous in the normal direction (Kirby et al., 2012). The middle panel of Fig. 1 shows this element discretising the horizontal velocity components. The vertical component of velocity is defined in the same way as on the 2D vertical slice mesh with a tensor product of P2 in the horizontal and P1DG in the vertical and is shown on the right hand side of Fig. 1. Similarly Fig. 2 from Section 2.6 detailing the melt parameterisation shows a 2D vertical schematic representation of the P1DG discretisation on the fully unstructured mesh versus the tensor product discretisation on the vertically structured extruded mesh.

2.4. Discontinuous Galerkin surface integrals

Various terms in the weak form have been integrated by parts, resulting in additional surface integrals, which can be exploited to enforce boundary conditions weakly. With continuous finite element basis functions, contributions from either side of the same internal boundary between elements cancel. This means only surface integrals defined on exterior boundaries need to be considered. In the Discontinuous Galerkin (DG) method however, the surface contributions from either side of the internal boundaries do not cancel, because of the ‘jump’ in the solution value across the boundary.

For the advection term in (28) we choose the advected velocity, $\bar{\mathbf{u}}$, to be the upwind value. On an exterior outflow boundary this simply means $\bar{\mathbf{u}}$ is the solution inside the domain and on inflow boundaries, $\bar{\mathbf{u}}$ is set using the Dirichlet values at the boundary. In the interior, the solution fields on either side of a facet between two adjacent elements

are arbitrarily labelled with '+' and '-'. Averages across interfaces are denoted by $\{\{\cdot\}\}$, while $[[\cdot]]$ denotes a jump across an interface with unit normal \mathbf{n} . Integrals over interior facets are given by $\int_{\Gamma_{\text{int}}} dS$. Using this notation, the additional term for DG advection is

$$\int_{\Gamma_{\text{int}}} \bar{\mathbf{u}} \cdot ([\phi(\mathbf{u} \cdot \mathbf{n})]) dS, \quad (34)$$

where the upwind velocity, $\bar{\mathbf{u}}$ is given by

$$\bar{\mathbf{u}} = \mathbf{u}^+ \quad \text{if } \mathbf{u} \cdot \mathbf{n}^- \leq 0, \quad (35)$$

$$\bar{\mathbf{u}} = \mathbf{u}^- \quad \text{if } \mathbf{u} \cdot \mathbf{n}^- > 0. \quad (36)$$

The viscosity term is discretised using the Symmetric Interior Penalty Galerkin (SIPG) method (Epshteyn and Rivière, 2007). After integrating by parts the discontinuous velocity field leads to an interior facet term

$$- \int_{\Gamma_{\text{int}}} [[\phi \otimes \mathbf{n}]] \cdot \{\{\mathbf{v} \cdot (\nabla \mathbf{u} + (\nabla \mathbf{u})^T)\}\} dS. \quad (37)$$

Adding only this term to (31) would not give a stable scheme. Instead, two additional terms are included which ensure that the form of the diffusivity term is symmetric positive definite given by

$$- \int_{\Gamma_{\text{int}}} ([[\mathbf{u} \otimes \mathbf{n}] + [[\mathbf{u} \otimes \mathbf{n}]^T]) \cdot \{\{\mathbf{v} \cdot \nabla \phi\}\} dS \quad (38)$$

$$+ \int_{\Gamma_{\text{int}}} \sigma [[\phi \otimes \mathbf{n}]] \cdot (\{\{\mathbf{v}\}\} \cdot ([[\mathbf{u} \otimes \mathbf{n}] + [[\mathbf{u} \otimes \mathbf{n}]^T])) dS.$$

The first term in (38) makes the scheme symmetric. The second term is a penalty term, penalising jumps in the solution to ensure that the operator remains coercive which is required for stability. Eqs. (3.20) and (3.23) from Hillewaert (2013) give a minimal value for the penalty parameter, σ as

$$\sigma = \alpha C(p)n_e \{A_f/V_e\}, \quad (39)$$

where n_e is the number of facets for the element type, α is a constant (the theoretical minimum is 1 for stability, but after experimentation we found using a value of 2 improves robustness), A_f is the area of the facet, V_e is the volume of the element. Eq. (3.7) from Hillewaert (2013) defines C as a constant that depends on polynomial degree, p , and element type such that

$$\int_f u^2 dS \leq C A_f/V_e \int_e u^2 dx, \quad (40)$$

is true for all polynomials u of degree p . Values for C are given in Table 3.1 from Hillewaert (2013). The additional SIPG terms are also added on exterior surfaces

$$- \int_{\Gamma_D} ((\mathbf{u} - \mathbf{u}_D) \otimes \mathbf{n} + ((\mathbf{u} - \mathbf{u}_D) \otimes \mathbf{n})^T) \cdot (\mathbf{v} \cdot \nabla \phi) ds \quad (41)$$

$$+ \int_{\Gamma_D} 2\sigma \phi \otimes \mathbf{n} \cdot (\mathbf{v} \cdot ((\mathbf{u} - \mathbf{u}_D) \otimes \mathbf{n} + ((\mathbf{u} - \mathbf{u}_D) \otimes \mathbf{n})^T)) ds.$$

Scalar advection is treated similarly to momentum, with the additional interior facet term, due to the discontinuous finite element space, given by

$$\int_{\Gamma_{\text{int}}} [[\phi \mathbf{u} \cdot \mathbf{n}]] \bar{q} dS, \quad (42)$$

where the upwind value of the tracer, \bar{q} , is defined as

$$\bar{q} = q^+ \quad \text{if } \mathbf{u} \cdot \mathbf{n}^- \leq 0, \quad (43)$$

$$\bar{q} = q^- \quad \text{if } \mathbf{u} \cdot \mathbf{n}^- > 0. \quad (44)$$

On Dirichlet boundaries \bar{q} is again only replaced if the flow is into the domain. Scalar diffusion in (23) is also implemented using the SIPG formulation, where the additional terms are given by

$$- \int_{\Gamma_{\text{int}}} [[\phi \mathbf{n}]] \cdot \{\{\boldsymbol{\kappa} \cdot \nabla q\}\} dS - \int_{\Gamma_{\text{int}}} [[q \mathbf{n}]] \cdot \{\{\boldsymbol{\kappa} \cdot \nabla \phi\}\} dS \quad (45)$$

$$+ \int_{\Gamma_{\text{int}}} \sigma [[\phi \mathbf{n}]] \cdot (\{\{\boldsymbol{\kappa}\}\} \cdot [[q \mathbf{n}]])) dS$$

$$- \int_{\Gamma_{D_q}} (q - q_D) \mathbf{n} \cdot \boldsymbol{\kappa} \cdot \nabla \phi ds + \int_{\Gamma_{D_q}} 2\sigma \phi \mathbf{n} \cdot \boldsymbol{\kappa} \cdot (q - q_D) \mathbf{n} ds,$$

and similar interpretations of each term can be made as for momentum viscosity in (31).

2.5. Timestepping and solver strategy

As outlined in the previous section, (26), (20) and (33), are the conservation of momentum, the incompressibility constraint and advection diffusion equations, and together they form a coupled system that needs to be solved for velocity, pressure, temperature and salinity. As stated before we are solving for the full nonhydrostatic dynamics and we do not carry out any mode splitting to separate the barotropic and baroclinic components of the flow.

First, we solve for the velocity and pressure using (26) and (33) by the established pressure projection approach (Chorin, 1967; Kramer et al., 2010), which splits the solve into two parts. There is an initial prediction step to find an intermediate velocity, where the momentum equation is solved using an implicit Backward Euler timestep for all the terms in the momentum equation, except the pressure, p^n , which uses the value from the previous timestep. The nonlinear advection and drag terms are linearised by setting the advecting velocity to the velocity value from the previous time step. This intermediate velocity, \mathbf{u}_{int} , is not divergence-free, however, so a correction step enforces the incompressibility constraint. This finds the updated pressure, p^{n+1} , needed to obtain a divergence-free velocity field \mathbf{u}^{n+1} . This correction step takes the form of a time-splitting step

$$\mathbf{u}^{n+1} = \mathbf{u}_{\text{int}} + \Delta t \nabla (p^{n+1} - p^n), \quad (46)$$

which is discretised in the same way as the M and P terms in (26), coupled with an incompressibility constraint Eq. (33) for \mathbf{u}^{n+1} .

Firedrake provides an interface to choose solver options to pass to PETSc, the package that implements the solver routines (Balay et al., 1997). Typically, we choose the GMRES Krylov subspace method to solve the momentum equation during the prediction step, because the matrix is not symmetric. It is preconditioned with an algebraic multigrid preconditioner, such as PETSc's default Geometric Algebraic Multigrid (GAMG) or BoomerAMG from the HYPRE suite (Henson and Yang, 2002).

The block structure of the coupled correction solve is simpler than the full Navier–Stokes system (26) and (33) because the matrix that arises in the top left-hand block is only a mass matrix (resulting from the time derivative term in (27)), where in the full system it would contain contributions from all the other terms in the momentum equation. Because of the careful choice of velocity–pressure finite element function spaces (see Section 2.3.2), where the gradient of pressure maps pointwise exactly into the velocity function space, it can be shown that the Schur complement of this system is exactly equivalent to a standard Continuous Galerkin discretisation of the pressure Poisson equation (Cotter and Ham, 2011).

This Poisson equation is solved using the Conjugate Gradient (CG) algorithm, preconditioned with an algebraic multigrid. The efficiency of this solution strategy is highly dependent on the condition number. For an anisotropic domain, this is proportional to the square of the aspect ratio of the grid cells (Kramer et al., 2010). Thus, as the aspect ratio goes to infinity the condition number can become unbounded, with negative implications for the convergence of the solve. However, by preconditioning the CG algorithm with a preconditioner that approximates the solution to the depth averaged equations, convergence can be made independent of the aspect ratio. This so called vertical lumping approach (Kramer et al., 2010) is similar to the idea first used in MITgcm and ensures that when the dynamics are close to a hydrostatic state, the solve is as fast as when only solving the hydrostatic equations (Marshall et al., 1997).

For parallel efficiency, the method does require cells to be aligned in columns in the vertical direction. This can be achieved using Firedrake's in-built extruded mesh feature in combination with the tensor product discretisation discussed in Section 2.3.1. We have found that when running 3D simulations on structured tetrahedral meshes, on aspect ratios larger than 10^2 (e.g., $dx = dy = 2$ km, $dz < 20$ m), using the generic algebraic multigrid preconditioner GAMG, the wall clock time needed to solve one timestep of the model was dominated by the pressure correction solve. However, on extruded meshes with the vertical lumping preconditioner, the pressure correction solve is faster than the momentum solve for the intermediate velocity.

After solving for velocity and pressure we solve for temperature and salinity individually. We use a 3-stage Diagonally Implicit Runge–Kutta (DIRK) time-stepping method (Ascher et al., 1995) for the temperature and salinity advection diffusion equations (20). The method is third order accurate. It is also an L-stable scheme, so it is suitable for stiff problems. We apply a vertex-based slope limiter after each time-step to prevent the solution from becoming unbounded (Kuzmin, 2010). We have found that using a slope limiter is only necessary in the region of the ice front transition, which tends to have the fastest flow as the buoyant melt water plume accelerates up the ice front. Without the limiter a spurious freezing signal can occur, confined to the final grid cell before the ice front transition. With the limiter, however, freezing does not take place.

Although, computational performance and parallel scalability are not directly considered in this paper we expect our P1DG finite element based method with implicit time stepping to be about an order of magnitude slower than a conventional low order (possibly explicit) finite volume model for the same model domain and grid resolution. The ability to take longer timesteps can offset this. More importantly, by using a fully flexible meshing strategy we are able to represent small scale features in a large domain that would be impossible to represent with structured grids, typical of other ocean models. We have run Firedrake simulations of ocean circulation beneath ice shelves in parallel with up to 128 cores. The reader is referred to the main Firedrake paper for a more in depth study of performance (Rathgeber et al., 2016). Future work will investigate the performance of this ocean model when it is applied to more demanding, geometrically complex 3D domains that require a fully flexible unstructured mesh.

2.6. Melt parameterisation

Melt rates are calculated based on conservation of heat and salt at the ice-ocean boundary and under the assumption that the boundary is at the freezing point, commonly referred to as the ‘three-equation’ melt parameterisation (Hellmer and Olbers, 1989; Holland and Jenkins, 1999; Dinniman et al., 2016). Conservation of heat is expressed as

$$Q_i^T - Q_w^T = Q_{\text{latent}}^T, \quad (47)$$

where the latent heat term, Q_{latent}^T is given by

$$Q_{\text{latent}}^T = -\rho_0 w_b L_f, \quad (48)$$

with w_b the unknown melt rate (melting: $w_b > 0$, freezing: $w_b < 0$) and $L_f = 3.34 \times 10^5$ J kg⁻¹ is the latent heat of fusion. The conductive heat flux into the ice, Q_i^T is

$$Q_i^T = -\rho_i c_{pi} w_b (T_i - T_b), \quad (49)$$

where $\rho_i = 920$ kg m⁻³ is the density of ice, $c_{pi} = 2000$ J kg⁻¹ K⁻¹ is the specific heat capacity of ice, T_i is the far field ice temperature, taken to be -20 °C and T_b is the unknown temperature at the ice-ocean boundary. The turbulent heat flux from the ocean to the ice, Q_w^T is given by

$$Q_w^T = -\rho_0 c_p \gamma_T (T_b - T_w), \quad (50)$$

where $c_p = 3974$ J kg⁻¹ K⁻¹ is the specific heat capacity of water and γ_T the turbulent thermal exchange velocity. Commonly γ_T is expressed

as $\Gamma_T u^*$, where Γ_T is a dimensionless constant and $u^* = \sqrt{C_D} |u|$ is the friction velocity (Asay-Davis et al., 2016). We use the velocity at the top surface of the domain to calculate u^* , in effect making the assumption that our computational boundary is in the surface layer (where rotation effects are unimportant) which may not be valid in the presence of strong stratification (McPhee, 2008). T_w is defined by Holland and Jenkins (1999) as the temperature in the mixed layer at the edge of the turbulent ice-ocean boundary layer. In most ocean models this is either taken as the value of the temperature in the top model cell adjacent to the boundary or temperature value representative of some distance from the ice-ocean interface (Gwyther et al., 2020). We take T_w as the model temperature value at the computational boundary. This is perhaps an uncommon choice compared with other ocean models (Dinniman et al., 2016). Unless doing Direct Numerical Simulation (DNS), which is prohibitively expensive for ice-ocean simulations in more than metre scale domains (Couston et al., 2021), the scalar molecular and turbulence boundary layer will not be modelled explicitly; it must be accounted for with a parameterisation. Since we are unable to model these effects explicitly we choose to imagine that the actual ice-ocean interface is offset slightly beyond the computational domain, which is similar to the approach taken for the parameterisation of wall regions in Reynolds Averaged Navier–Stokes (RANS) turbulence models (Yeager, 2018). Fig. 2 shows a schematic of the melt parameterisation. While the finite element solution is valid at any point within each element, and this means it is possible to offset where the tracer is sampled towards the interior of the domain as shown by Kimura et al. (2013), this could cause performance issues when running the model in parallel because the sampled tracer value may not be stored on the same processor.

The conservation of salt is similar to heat and is given by

$$Q_i^S - Q_w^S = Q_{\text{brine}}^S, \quad (51)$$

where the fresh water melt flux Q_{brine}^S is given by

$$Q_{\text{brine}}^S = -\rho_0 w_b (S_i - S_b), \quad (52)$$

and where S_i is the salinity of the ice, which is taken to be zero, and S_b is the unknown salinity at the ice base. The turbulent salt flux across the ocean boundary layer Q_w^S can be written as

$$Q_w^S = -\rho_0 \gamma_S (S_b - S_w), \quad (53)$$

where γ_S is the turbulent salinity exchange velocity, analogous to γ_T and as before we take S_w as the value of salinity at the edge of the computational domain offset from the ice-ocean boundary. The diffusive flux of salt into the ice, Q_i^S is assumed to be zero. The ice-ocean boundary temperature, T_b is assumed to be at the freezing point

$$T_b = a S_b + b + c P_b, \quad (54)$$

where $a = -5.73 \times 10^{-2}$ °C, $b = 8.32 \times 10^{-2}$ °C and $c = -7.53 \times 10^{-8}$ °C Pa⁻¹ are the coefficients obtained from linearisation of the nonlinear freezing point equation (Millero, 1978) using the values from Asay-Davis et al. (2016) and P_b is the pressure at the boundary. This means there are now three equations ((47) (51) and (54)) to solve for the three unknowns, w_b , S_b , and T_b ; the melt rate and the salinity and temperature values at the boundary, respectively. Rearranging the equations gives a quadratic equation for S_b , and the positive value is taken as the solution. By using the form of the heat and salt flux from (50) and (53), Robin boundary conditions for temperature and salinity can be imposed as shown in Fig. 2 according to

$$\Phi_T = \kappa \mathbf{n} \cdot \nabla T = (\gamma_T + w_b)(T_b - T_w), \quad (55)$$

and

$$\Phi_S = \kappa \mathbf{n} \cdot \nabla S = (\gamma_S + w_b)(S_b - S_w), \quad (56)$$

where Φ_T and Φ_S are substituted into the third term of (23). Since the ice-ocean boundary is not a material interface, we include the melt water correction term accounting for the advection of melt water (with water mass properties corresponding to values at the ice-ocean interface) into the ocean domain as part of the temperature and salt flux boundary conditions following Jenkins et al. (2001).

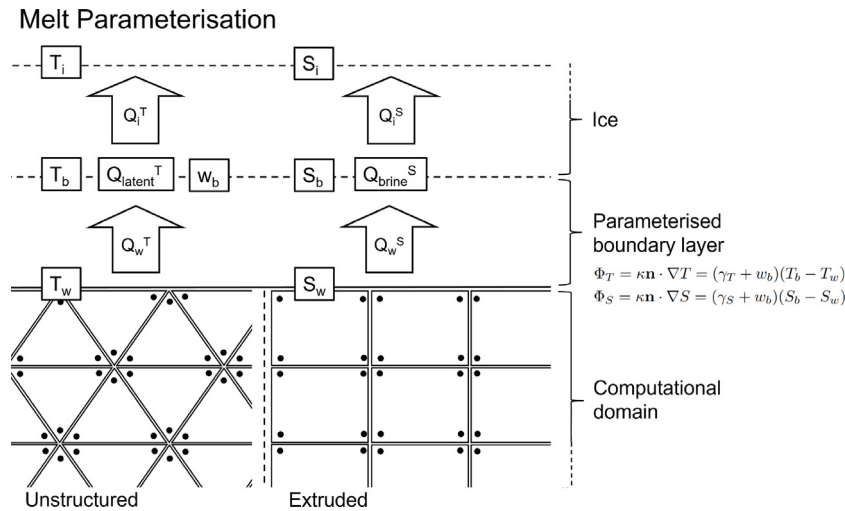


Fig. 2. 2D vertical schematic of melt parameterisation and implementation of tracer boundary conditions. Black dots on the left hand side of the computational domain represent locations of Degrees of Freedom (DOFs) for PIDG finite element, which is used for the tracers and velocity components on vertically unstructured meshes. The right hand side of the figure represents the computational domain for the vertically structured (but horizontally unstructured) ‘extruded’ mesh. Black dots are DOFs for the element used for scalar quantities formed by taking the tensor product of PIDG on a horizontal interval with PIDG defined on a vertical interval. The velocity element (not shown) is also defined using a tensor product (as described in Section 2.3.2) and is discontinuous in the normal direction, but continuous in the tangential direction. Note the offset computational domain so in our modelling we assume that the ice-ocean interface is slightly beyond the computational domain following Yeager (2018). T_w and S_w are the finite element solutions for temperature and salinity at the edge of the computational domain. Other terms are as defined in Section 2.6. After solving for T_b , S_b and w_b , Robin boundary conditions for temperature and salinity are imposed weakly (third term of (23)) using the definition of the heat, Q_w^T , and salt Q_w^S , fluxes from the melt parameterisation.

Table 1
General constants for model equations and melt parameterisation. Note in this work Salinity is treated as unitless, except for thermodynamic relations or equations of state, where it is assumed to be g kg^{-1} .

Parameter	Value	Description
f	$-1.409 \times 10^{-4} \text{ s}^{-1}$	Coriolis parameter
ρ_0	$1027.51 \text{ kg m}^{-3}$	Reference density of sea water
g	9.81 m s^{-2}	Gravitational acceleration
α_T	$-3.733 \times 10^{-5} \text{ }^\circ\text{C}^{-1}$	Temperature expansion coefficient
β_S	7.843×10^{-4}	Salinity contraction coefficient
T_0	$1.0 \text{ }^\circ\text{C}$	Reference temperature in EOS
S_0	34.2	Reference salinity in EOS
L_f	$3.34 \times 10^5 \text{ J kg}^{-1}$	Latent heat of fusion of ice
ρ_i	920 kg m^{-3}	Density of ice
c_{pi}	$2000 \text{ J kg}^{-1} \text{ K}^{-1}$	Specific heat capacity of ice
c_p	$3974 \text{ J kg}^{-1} \text{ K}^{-1}$	Specific heat capacity of water
T_i	$-20 \text{ }^\circ\text{C}$	Far field ice temperature
S_i	0	Far field ice salinity
a	$-5.73 \times 10^{-2} \text{ }^\circ\text{C}$	Coefficient freezing point equation
b	$8.32 \times 10^{-2} \text{ }^\circ\text{C}$	Coefficient freezing point equation
c	$-7.53 \times 10^{-8} \text{ }^\circ\text{C Pa}^{-1}$	Coefficient freezing point equation
Γ_T	0.011	Turbulent heat exchange coefficient
Γ_S	$\Gamma_T/35$	Turbulent salt exchange coefficient
C_D	2.5×10^{-3}	Drag coefficient

2.7. Model implementation

We have built a Python module on top of Firedrake to facilitate simulations of ice shelf cavities. The core part of our Python package is made up of 11 modules. We make use of Python classes to build up the complexity of the code in blocks. Each equation inherits useful attributes from a BaseEquation, i.e. a method for summing the residuals from each term. And each equation is made up of terms, effectively each term in (20) and (26), themselves based on a BaseTerm, which contains the required features of a term, i.e. a method to evaluate the residual form for the term. This approach leads to a compact codebase, aids with debugging, as individual pieces of the code can be isolated and tested, but retains the flexible nature of Firedrake, which is one of the main reasons for using the software. The code is available at <https://github.com/thwaitesproject/thwaites.git>. It has also been stored in this Zenodo repository: <https://doi.org/10.5281/>

[zenodo.7584700](https://zenodo.org/record/7584700), along with instructions for running the examples from Section 3.

The next section presents a simplified verification test applicable to ice-ocean models, which will hopefully prove of value for other models, to ensure that the code is solving the model equations correctly and consistently.

3. Results

3.1. Method of manufactured solution (MMS) tests

The Method of Manufactured Solutions (MMS) is a rigorous method for verifying that the numerical scheme as implemented is solving the model equations accurately and consistently (Farrell et al., 2011; Roache, 1998). This is a separate problem to that of validating whether the model, including the choice of underlying equations, yields a good approximation to the real system. Such a validation requires comparison with observations, which is challenging for simulations of ice shelf cavity circulation. The most reliable way to test whether a numerical approximation is consistent is to ensure that the solution error relative to a known solution decreases at the correct order as the grid is refined (Salari and Knupp, 2000). However, finding analytical solutions to systems of coupled Partial Differential Equation (PDE) that arise in fluid dynamics is notoriously difficult, and only possible for simple geometrical configurations and often simplified Boundary Conditions (BCs). MMS removes much of this difficulty. Instead of solving the PDE system to find the solution, the analytical form of the solution is chosen to begin with. Solution fields are constructed from smoothly varying differentiable functions, often trigonometric and polynomial functions. These functions are substituted into the strong form of the equations and the resulting residuals (which are expected to be non-zero as these functions could not be expected to represent an exact solution) provide additional source terms that can be used to force a version of the simulation for which the chosen functions are exact solutions. The chosen functions can be manufactured to meet specific constraints (for example incompressibility of the flow field) or specific BCs. Alternatively, consistent BCs can be found directly from the chosen solutions simply by evaluating the solution at the boundary for the type of BC required for testing.




Fig. 3. Example MMS solution and error fields. Left panel shows example solution field for horizontal velocity components on the two coarsest grids (10×10 cells and 20×20 cells). Red indicates flow to the right and blue indicates flow to the left. The right panel shows the squared error of the finite element solution relative to the chosen manufactured solution on the two meshes. Note the log scale on the right hand plot. Although the solution fields in the left hand panels at the two resolutions are visually similar, the error is clearly lower on the finer grid in the right hand panel. This plot is intended as a qualitative example of MMS, Fig. 4 demonstrates that the solution fields converge at the correct order.

Following the notation from Farrell et al. (2011), the error on a given mesh is

$$E_{h_1} \approx C h_1^{c_p}, \quad (57)$$

where C is a constant, h_1 is the characteristic mesh size, and c_p is the order of accuracy of the numerical approximation. The error on a finer mesh is

$$E_{h_2} \approx C \left(\frac{h_1}{r} \right)^{c_p}, \quad (58)$$

where the new characteristic mesh size $h_2 = h_1/r$. By taking the ratio of the two errors the convergence rate can be found using

$$c_p \approx \log_r \left(\frac{E_{h_1}}{E_{h_2}} \right). \quad (59)$$

The MMS solutions should be constructed to be as general as possible to ensure good coverage of the terms in the equations (Salari and Knupp, 2000). Although this does not require the chosen solution fields to be realistic, the solution we have chosen for this MMS test is an overturning flow that is a crude representation of an ice-pump mechanism.

For simplicity we have limited ourselves here to flow within a 2D vertical slice. The domain is a square box with length $L = 100$ m, height $H = 100$ m, and with the bottom boundary at a depth $D = 1000$ m. The chosen solution for the horizontal velocity is

$$u = u_0 \frac{x}{L} \cos \left(-\frac{\pi(z + D - H)}{H} \right), \quad (60)$$

where u_0 is a characteristic velocity of 1 m/s, x and z are the horizontal and vertical coordinates. As an example of the MMS procedure, Fig. 3 shows the finite element solution fields for the horizontal velocity component on the two coarsest grids alongside the corresponding error relative to the chosen solution. The error on the finer mesh is noticeably reduced compared to the error on the coarser mesh despite the two solution fields being qualitatively similar.

The chosen solution for vertical velocity is

$$w = u_0 \frac{H}{\pi L} \sin \left(-\frac{\pi(z + D - H)}{H} \right). \quad (61)$$

This has been chosen to ensure that the flow is incompressible so that the continuity equation, $\nabla \cdot \mathbf{u} = 0$, is satisfied. For this simple flow configuration, the vertical velocity does not depend on x . The right hand side is specified as an open boundary and all other boundaries employ free-slip conditions, with no normal flow imposed. The pressure field has been chosen to have homogeneous Neumann boundaries ($\nabla p \cdot \mathbf{n} = 0$)

$$p = p_0 \cos \left(\frac{\pi x}{L} \right) \cos \left(-\frac{\pi(z + D - H)}{H} \right). \quad (62)$$

p_0 is a constant set to 1 Pa to ensure the dimensions remain consistent and the reference density, ρ_0 , is set to 1 kg m^{-3} for this test. The temperature field is

$$T = T_c \sin \left(\frac{4\pi x}{L} \right) + a_T z^2 + b_T z + c_T, \quad (63)$$

where T_c is $0.1 \text{ }^\circ\text{C}$, a_T is $-3.89 \times 10^{-4} \text{ }^\circ\text{C m}^{-2}$, b_T is $-7.54 \times 10^{-1} \text{ }^\circ\text{C m}^{-1}$, and c_T is $-3.64 \times 10^2 \text{ }^\circ\text{C}$. The salinity field is

$$S = S_c \cos \left(\frac{4\pi x}{L} \right) + a_S z^2 + b_S z + c_S, \quad (64)$$

where S_c is 0.345, a_S is $-1.44 \times 10^{-4} \text{ m}^{-2}$, b_S is $-2.81 \times 10^{-1} \text{ m}^{-1}$, and c_S is -1.03×10^2 . The viscosity and diffusivity are $1 \text{ m}^2/\text{s}$.

As already described, the source terms are derived by substituting the solutions into the strong form of the equations. The source term for velocity is thus given by

$$\mathbf{S}_u = \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} + \frac{1}{\rho_0} \nabla p - \nabla \cdot \boldsymbol{\tau} + \frac{\rho'}{\rho_0} \mathbf{g} \mathbf{k}. \quad (65)$$

Note the coriolis term is not included because the domain is a 2D vertical slice. The implementation of rotation has been verified against analytical solutions of Ekman spirals (not shown). The source term for temperature is given by

$$\mathbf{S}_T = \frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T - \nabla \cdot \kappa_T \nabla T. \quad (66)$$

The source term for salinity is given by

$$\mathbf{S}_S = \frac{\partial S}{\partial t} + \mathbf{u} \cdot \nabla S - \nabla \cdot \kappa_S \nabla S. \quad (67)$$

Use of symbolic computation tools, such as SymPy (Meurer et al., 2017), make this process relatively straightforward to implement. For example, the source term for the horizontal velocity component in this case is given by

$$S_u = \frac{-p_0 \pi \sin(\pi x/L) \cos(\pi(-H+D+z)/H)}{L \rho_0} + \frac{u_0^2 x \sin(\pi(-H+D+z)/H)^2}{L^2} + \frac{u_0^2 x \cos(\pi(-H+D+z)/H)^2}{L^2} + \frac{\pi^2 \nu u_0 x \cos(\pi(-H+D+z)/H)}{H^2 L} \quad (68)$$

Note we chose this term for its compactness because it does not include buoyancy terms.

Dirichlet BCs for temperature and salinity are applied on the left and right boundaries, with Neumann conditions on the bottom boundary. Values for these boundary conditions are found in a similar manner to defining source terms, the solution itself is substituted to give the exact form of the (generally inhomogeneous) boundary conditions (Salari and Knupp, 2000). Particular attention is given to the boundary conditions at the ice-ocean interface situated at the top of the domain, as verifying the accuracy of the model implementation at this location is a key reason for carrying out the MMS test. The boundary conditions at the ice-ocean interface are represented by Robin boundary conditions, which define the gradient of the temperature and salinity fields at the boundary as a function of the temperature and salinity at the edge of the computational domain (see Section 2.6 for more information). The default boundary condition for the melt boundary is given by (55) and (56), which are imposed weakly in the third term of (23). For an MMS test the temperature boundary condition applied at the top is given by

$$\Phi_T = \kappa \mathbf{n} \cdot \nabla T^a + \Phi_T^h - \Phi_T^a. \quad (69)$$

The first term in (69) represents the temperature gradient of the analytical temperature field, T^a , (found from (63)), at the boundary with the eddy diffusivity as a multiplication factor as discussed previously. This is the only term needed to implement Neumann boundary conditions for conventional MMS tests. The second term Φ_T^h is the flux boundary condition calculated by the melt parameterisation using the modelled temperature (and salinity) values, equivalent to (55) (and (56)). The final term, Φ_T^a , is the flux calculated by the melt parameterisation using the analytical temperature field. Provided the temperature, salinity and velocity (through u^*) fields converge at the correct order, then the melt rate error (calculated using the exact temperature, salinity and velocity fields) should also converge at the correct order because the contribution from Φ_T^h and Φ_T^a cancel out.

This is a spatial convergence test so the solution fields have been chosen to be stationary (Farrell et al., 2011). The simulation is spun up from rest. Backward Euler is used for timestepping the momentum equations and the scalar equations. The timestep is initially set to $4 \text{ m}/(u_0 n_x)$, chosen for robustness in the initial spin-up, where n_x is the number of cells in the x direction. The timestep is increased to $L/(u_0 n_x)$ after 100 time steps. For convenience Table 2 summarises parameters and constants used for the MMS convergence test.

Velocity and tracer fields are discretised with P1DG elements and so second order convergence is expected, which means that halving the grid size should reduce the error by a factor of four. Pressure is discretised with P2 elements so third order might be expected; however, coupling errors associated with the velocity mean that in practice the convergence is about second-order (Cotter and Ham, 2011).

Fig. 4 plots the errors for each solution field, with logarithmic axes. Velocity, temperature and melt rate are all very close to the expected second order convergence or higher. Salinity and pressure are slightly lower, though never below 1.8 and the convergence rate is increasing towards second order as the grid is refined. The integrated melt is also initially above second order but drops to 1.82 convergence. Overall these convergence results give confidence in the accuracy and consistency of the numerical discretisation and hopefully will prove to

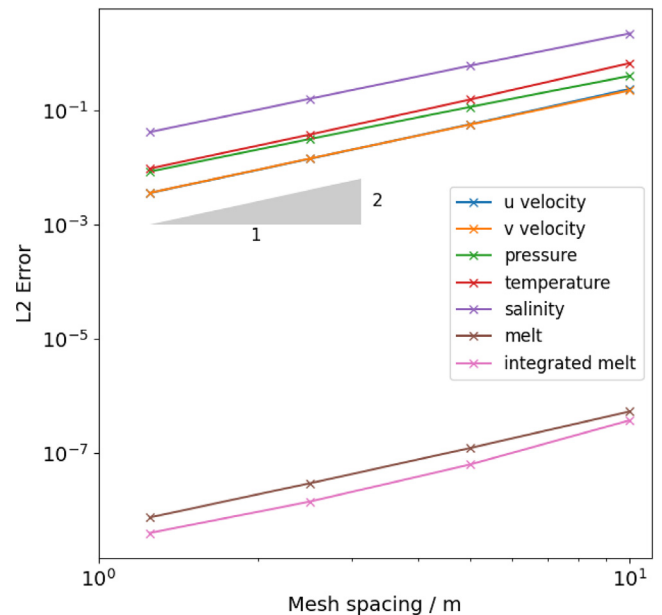


Fig. 4. MMS convergence tests for temperature, salinity, melt, integrated melt over the top boundary, velocity and pressure. The test was carried out on four grids with grid resolution of 10 m (10×10 cells), 5 m (20×20 cells), 2.5 m (40×40 cells), 1.25 m (80×80). The error is given in the L_2 norm. Both axes use a logarithmic scale. A triangle with slope of two has been added to the plot to help show that all fields achieve the expected second order convergence.

Table 2

Parameters and constants used for MMS convergence test in Section 3.1.

Parameter	Value	Description
L	100 m	Domain length
H	100 m	Domain height
D	1000 m	Domain depth
u_0	1 m s ⁻¹	Characteristic velocity
p_0	1 Pa	Characteristic pressure
ρ_0	1 kg m ⁻³	Reference density
T_c	0.1 °C	Coefficient in chosen temperature solution
a_T	-3.89×10^{-4} °C m ⁻²	Coefficient in chosen temperature solution
b_T	-7.54×10^{-1} °C m ⁻¹	Coefficient in chosen temperature solution
c_T	-3.64×10^2 °C	Coefficient in chosen temperature solution
S_c	0.345	Coefficient in chosen salinity solution
a_S	-1.44×10^{-4} m ⁻²	Coefficient in chosen salinity solution
b_S	-2.81×10^{-1} m ⁻¹	Coefficient in chosen salinity solution
c_S	-1.03×10^2	Coefficient in chosen salinity solution
μ	1 m ² s ⁻¹	Viscosity
κ	1 m ² s ⁻¹	Diffusivity
n_x	Variable	Number of grid points
Δt	4 m/($u_0 n_x$)	Timestep (first 100 steps)
Δt	$L/(u_0 n_x)$	Timestep (after 100 steps)

be a useful tool as the model continues to be developed. In the following section we demonstrate our ability to simulate more complex ice shelf cavity geometries in 3D using the ISOMIP+ Ocean0 test case.

3.2. 3D simulations of an idealised ice shelf cavity: ISOMIP+ Ocean0

The aim of this section is to demonstrate our ability to simulate ocean flow in a fully three-dimensional ice shelf cavity. The Ice Shelf Ocean Model Intercomparison Project+ (ISOMIP+) consists of a set of idealised ice shelf cavity geometries and parameters for a set of common experiments (Asay-Davis et al., 2016). Here we use the Ocean0 experiment, which has been used for a number of studies to introduce and compare different ocean models for ice shelf cavity applications (Zhou and Hattermann, 2020; Gwyther et al., 2020; Favier et al., 2019). This experiment reaches a quasi-steady state in the cavity within a few months of simulation time, so it is a convenient test

Table 3
Parameters and constants used for the ISOMIP+ *Ocean0* test case in Section 3.2.

Parameter	Value	Description
L	~400 km	Domain length
W	~80 km	Domain width
D	720 m	Maximum depth
dx, dy	2 km	Default horizontal grid size
nz	36 (30) layers	Default (under ice) vertical grid for Firedrake
dz	20 m	Default vertical grid size for MITgcm
h_{zmin}	10 m	Minimum ocean thickness for Firedrake
h_{zmin}	40 m	Minimum ocean thickness for MITgcm
Δt	900 s	Timestep
μ_h	$6 \text{ m}^2 \text{ s}^{-1}$	Horizontal viscosity
μ_v	$1 \times 10^{-3} \text{ m}^2 \text{ s}^{-1}$	Vertical viscosity (ISOMIP+ switch)
κ_h	$1 \text{ m}^2 \text{ s}^{-1}$	Horizontal diffusivity
κ_v	$5 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$	Default vertical diffusivity (ISOMIP+ switch)

Table 4
Parameters and constants used for adjoint sensitivity calculations in the idealised grounding zone domain from Section 3.3. Note the values used for the expansion and contraction coefficients for temperature and salinity are changed from Table 1.

Parameter	Value	Definition
L	10 km	Domain length
H	2–100 m	Domain height
D	600 m	Domain depth
dx	500 m	Horizontal grid size
dz	2 m	Vertical grid size
Δt	300 s	Timestep
T_{init}	$1 \text{ }^\circ\text{C}$	Initial and restoring temperature
S_{init}	$34.4 \text{ }^\circ\text{C}$	Initial and restoring salinity
$t_{restore}$	1 day	Restoring period
μ_h	$0.25 \text{ m}^2 \text{ s}^{-1}$	Horizontal viscosity
μ_v	$1 \times 10^{-3} \text{ m}^2 \text{ s}^{-1}$	Vertical viscosity
κ_h	$0.25 \text{ m}^2 \text{ s}^{-1}$	Horizontal diffusivity
κ_v	$1 \times 10^{-3} \text{ m}^2 \text{ s}^{-1}$	Vertical diffusivity
α_T	$-2 \times 10^{-4} \text{ }^\circ\text{C}^{-1}$	Temperature expansion coefficient
β_S	7×10^{-4}	Salinity contraction coefficient

case for carrying out model comparisons and testing model parameter choices (Asay-Davis et al., 2016). There is not an ‘exact’ solution that all models are aiming to replicate. For this reason, we have found it helpful to run simulations of MITgcm alongside our model. MITgcm is a finite volume, z -layer (fixed vertical resolution), ocean model widely used in a number of applications including investigations of ice shelf-ocean cavities (Losch, 2008; Dansereau et al., 2014; Seroussi et al., 2017; Nakayama et al., 2017; Kimura et al., 2017; Nakayama et al., 2021a; Naughten et al., 2021). It also has a nonhydrostatic option, selected for this comparison, so that both models solve the same underlying set of equations, albeit with different discretisations.

The *Ocean0* domain is approximately 400 km long, 80 km wide and has a maximum depth of 720 m. The bathymetry is specified by an analytical profile representing an idealised fjord, and the ice draft is from an ice flow simulation from the corresponding MISMIP+ experiment (Asay-Davis et al., 2016). The target ISOMIP+ grid resolution is specified as 2 km in the horizontal and 36 vertical layers, with the caveat that different vertical meshing strategies will impose constraints on the resolution (Asay-Davis et al., 2016). The ISOMIP+ protocol does not specify exactly how the cavity thickness should be defined as the ‘real’ cavity pinches to zero thickness at the grounding line. For MITgcm simulations we have followed the suggested 40 m minimum thickness for z -layer models, so that there are always at least two cells with the specified 20 m vertical resolution. To create the 3D Firedrake mesh we used the inbuilt ‘extruded’ mesh capability of Firedrake. First, we used QGIS (QGIS Development Team, 2022) to generate the horizontal extent of the mesh. We chose the grounding line by tracing the 10 m ocean thickness contour (ice draft–bathymetry). We used the *qmesh* package (Avidis et al., 2018) to generate the horizontal surface mesh that forms the base of the columns in the full 3D extruded mesh. *qmesh* is effectively a wrapper for Gmsh (Geuzaine and Remacle, 2009) to

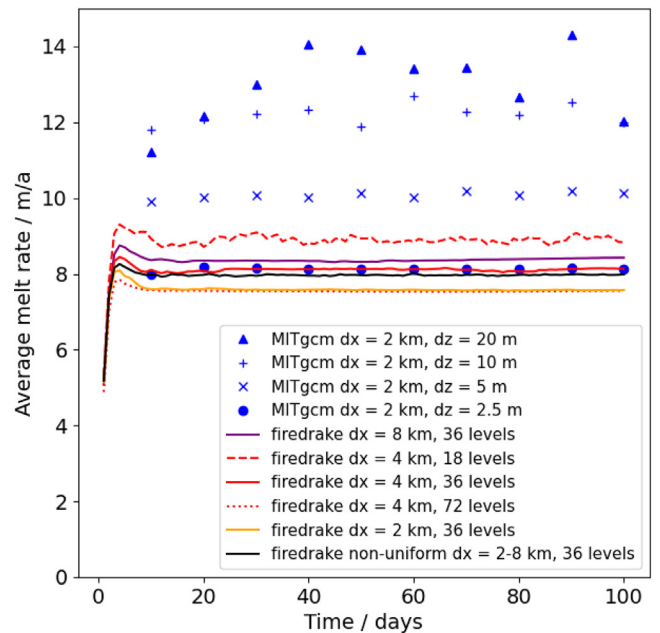


Fig. 5. Integrated melt rate over deepest region, where ice draft is deeper than 300 m following Asay-Davis et al. (2016). Melt rates are instantaneous values, not time averaged. MITgcm values are plotted every 10 days with blue symbols dependent on vertical resolution. Triangles — 20 m, Pluses — 10 m, Crosses — 5 m, Circles — 2.5 m. Firedrake values are plotted as lines, with colour dependent on horizontal resolution (and line style dependent on vertical resolution). Orange — 2 km, Red — 4 km (Dashed — 18 vertical levels outside the cavity, Continuous — 36 vertical levels outside the cavity, Dotted — 72 levels outside the cavity), Purple — 8 km, Black — 2 km in vicinity of grounding zone and western boundary coarsening to 8 km outside the cavity.

allow meshing of domains generated from GIS data. The surface mesh is then extended in the vertical direction to create columns of extruded triangular prism elements (Bercea et al., 2016; McRae et al., 2016). We use 30 layers under the ice, with a transition to 36 layers in the open ocean region, with this transition occurring over a single horizontal grid cell. Squashing the domain under the ice gives a terrain following, quasi-sigma style mesh within the cavity, with higher vertical resolution towards the grounding line. We use the tensor product elements described in Section 2.3.2 for the discretisation of the velocity–pressure finite element pair, and for tracers the tensor product element described in 2.3.1.

The temperature and salinity profiles for initialisation and restoring on the northern boundary are the ‘warm’ profiles given in Table 6 from Asay-Davis et al. (2016). We use a timestep of 900 s, with Backward Euler time stepping for the momentum equations and a Diagonally Implicit Runge–Kutta method, DIRK33, for the tracers. Horizontal and vertical kinematic viscosity are specified at $6 \text{ m}^2/\text{s}$ and $1 \times 10^{-3} \text{ m}^2/\text{s}$ respectively, and horizontal and vertical diffusivity are specified at $1 \text{ m}^2/\text{s}$ and $5 \times 10^{-5} \text{ m}^2/\text{s}$ as per the ISOMIP+ protocol (Asay-Davis et al., 2016). We include the ‘switch’ to higher vertical mixing values when the stratification is unstable, however, the scheme has negligible effect on calculated melt rates due to the build up of stable stratification at the ice-ocean boundary during melting.

The melt parameterisation specified by ISOMIP+ does not include heat flux into the ice and uses constant turbulent exchange coefficients based on Jenkins et al. (2010). For this comparison both models use the same turbulent exchange coefficients in the melt parameterisation, $\Gamma_T: 0.011$ and $\Gamma_S = \Gamma_T/35$. The vertical averaging scheme for melting is turned off in MITgcm (Losch, 2008). Drag boundary conditions are applied on the top and bottom interfaces with a drag coefficient, $C_D = 2.5 \times 10^{-3}$. Model parameters are summarised in Table 3.

Fig. 5 shows melt rates for the first 100 days of simulation integrated over the region where the ice draft is deeper than 300 m, following Asay-Davis et al. (2016). We have found that melt rates are very

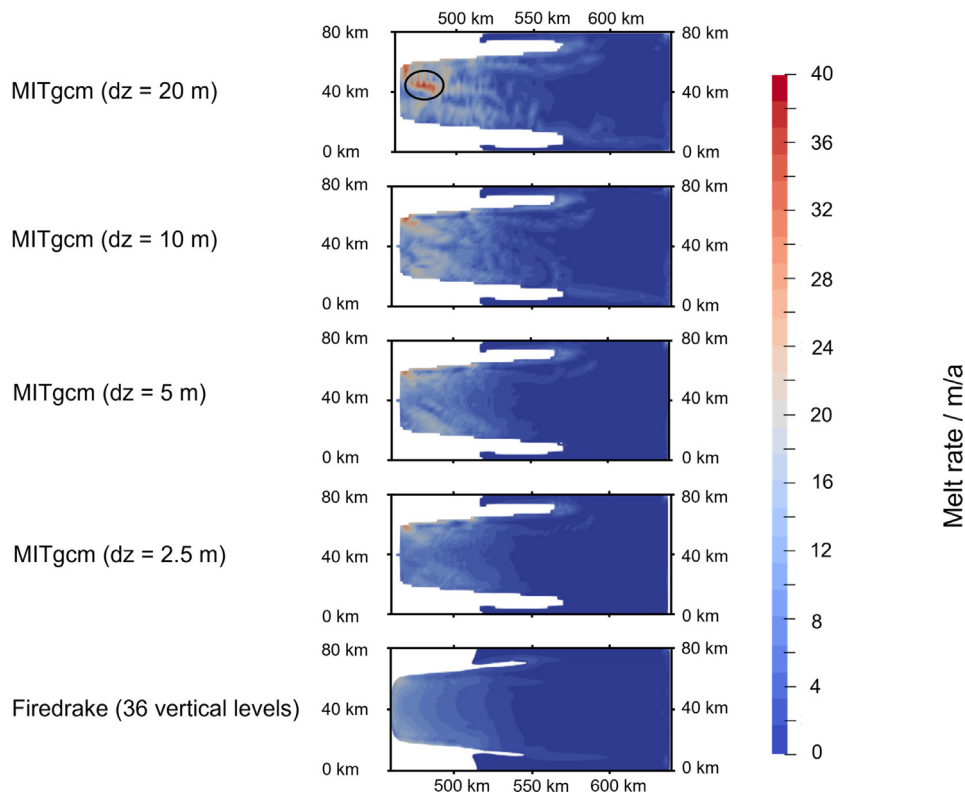


Fig. 6. Plan view of basal melt at 50 days for MITgcm at target ISOMIP+ resolution for z -layer models (dx : 2 km, dz : 20 m), along with finer vertical resolution (dz : 10 m, 5 m, 2.5 m) and Firedrake at target ISOMIP+ resolution for terrain following model (dx : 2 km, 36 levels with 30 levels inside the ice shelf cavity). The circle in the top panel highlights a time varying, grid resolution dependent feature that is referred to in the text.

sensitive to the grid resolution. MITgcm melt rates are plotted in blue, with vertical resolution given by the symbol (triangles: 20 m (target ISOMIP+ resolution), pluses: 10 m, crosses: 5 m, circles: 2.5 m). Clearly, integrated melt rates calculated by MITgcm are highly dependent on vertical resolution, as refining the mesh three times leads to nearly a factor of two difference and there is no sign of convergence. Our model also shows vertical dependence of melt rate as evidenced by the red lines in Fig. 5 (dashed: 18 levels, continuous: 36 levels, dotted: 72 levels), however, the spread is lower than with the z -layer model. These results are consistent with a recent study by Gwyther et al. (2020) using a range of models applied to the same ISOMIP+ *Ocean0* experiment and we will refer to their work later in the discussion.

Fig. 6 shows plan view snapshots of melt rate at 50 days for MITgcm at different vertical resolutions, compared with Firedrake at the target resolution specified by ISOMIP+ for layered models (36 layers outside the cavity, 30 layers inside the cavity). As the vertical resolution of MITgcm increases the spatial melt rate pattern can be seen to approach that from Firedrake. The streaky, time evolving artefacts present in MITgcm at coarser vertical resolutions (circled in Fig. 6) disappear as the grid is refined. Note that while time-mean fields of melt rate act to smooth out the ‘streaks’ there are still noisy grid scale features after time averaging (not shown).

Reducing the vertical diffusivity and viscosity linearly at the same rate as the grid refinement does not bring the ‘streaks’ back (not shown). This suggests it is not a numerical instability of the type dependent on grid Reynolds number (the Reynolds number evaluated at the grid scale, $Re_{\Delta x} = U\Delta x/\nu$ where U is a characteristic velocity, Δx is the grid resolution and ν is the eddy viscosity) or grid Péclet number (similarly $Pe_{\Delta x} = U\Delta x/\kappa$ where κ is the eddy diffusivity). The streaks may therefore be a grid discretisation error caused by the step representation of the boundary that is inherent to z -layer models, even with a partial cell representation (Losch, 2008). Both models show highest melting at the grounding line, which is consistent

with the depth dependent melting point, warm initial and restoring conditions at the bottom of the water column, and steeper basal slopes close to the grounding line that enhance buoyancy driven flow. This melt pattern is also consistent with other published ISOMIP+ *Ocean0* results (Asay-Davis et al., 2016; Zhou and Hattermann, 2020; Gwyther et al., 2020).

Within the cavity, both models show a clearly defined western (defined as positive y) boundary current, due to the effects of rotation. This can be seen in Fig. 7, which is a cross section of the cavity at $x = 480$ km (close to the grounding line). There is strong northward flow along the western boundary together with colder (and fresher) conditions, which results from melt water accumulating and preferentially flowing through this part of the domain. The middle panel of Fig. 7 shows MITgcm at the higher vertical resolution. Refining MITgcm leads to better qualitative agreement with Firedrake. The return boundary flow and the stratification at the top of the domain is qualitatively much better resolved by MITgcm at the higher resolution, particularly in the eastern half of the domain. At this location the cavity thickness is close to 100 m, so with 30 vertical layers, Firedrake’s vertical grid resolution is similar to MITgcm’s vertical grid resolution of 2.5 m, so it is not unexpected that the higher resolution MITgcm simulation yields results much closer to those from Firedrake.

In reality, melting leads to the injection of cold, and importantly, buoyant fresh water at the top of the domain. If the ice shelf base has a sufficient slope the buoyant melt water drives a shear flow and thus enhances mixing. However, melting can also lead to a negative feedback mechanism since melt water forms a stratified layer, especially for low basal slope angles that effectively shuts down vertical mixing (McPhee, 2008; Vreugdenhil and Taylor, 2019). This leads to reduced melt rates because the ice base is insulated from the warmer, saltier water below. It is the competition between the buoyancy driven shear and build up of stratification (as well as external sources, such as tidal currents) that controls mixing. Modelling vertical mixing accurately is therefore

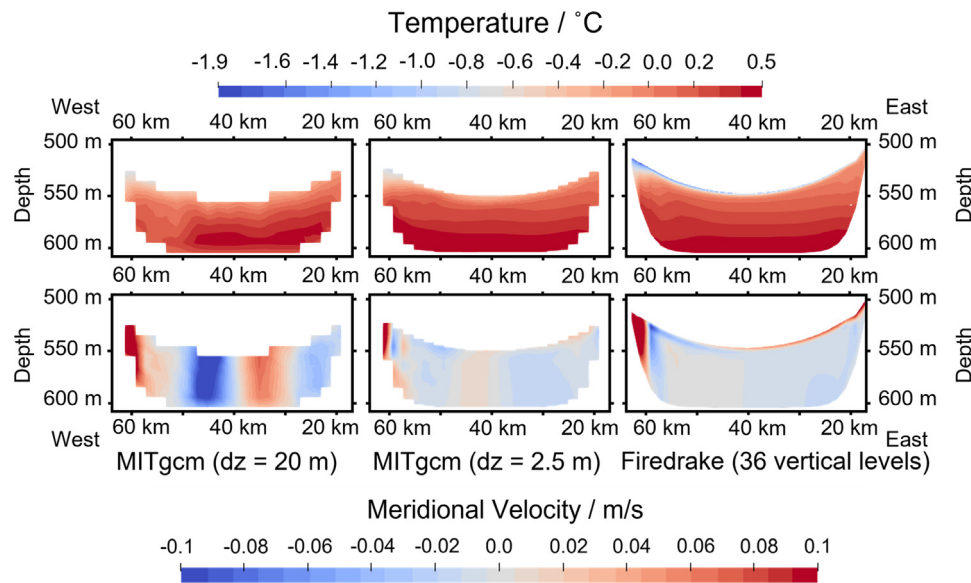


Fig. 7. Cross section across ice shelf cavity at $x = 480$ km, close to the grounding line. Top panels show a snapshot of temperature at 50 days, bottom panels show a snapshot of meridional velocity at 50 days. Plots are stretched one hundred times in the vertical. From left to right the panels correspond to MITgcm at the target ISOMIP+ resolution for z -layer models ($dx = 2$ km, $dz = 20$ m), MITgcm fine ($dx = 2$ km, $dz = 2.5$ m), Firedrake ($dx = 2$ km, 36 levels with 30 levels under the ice shelf). Note the cavity is about 100 m thick at this location, so the vertical size of the Firedrake layers is similar to MITgcm at the fine resolution shown in the middle panel. The alternating barotropic flow pattern, evident in the MITgcm meridional velocity at the ISOMIP+ target resolution ($dx: 2$ km, $dz: 20$ m), is reduced by plotting time averaged fields and is closer in magnitude to the simulation with finer vertical resolution. Time-mean fields of temperature are qualitatively very similar to snapshots after 50 days.

of fundamental importance to represent flow in ice shelf cavities and to calculate the associated melt rate. The total amount of mixing simulated by a numerical ocean model is a combination of resolved turbulence mixing, explicitly specified mixing that aims to account for the unresolved turbulent mixing (here we use a simple eddy diffusivity closure with constant diffusivity) and a grid size dependent implicit contribution, often referred to as numerical mixing. This spurious, numerical mixing is often dictated by the details of the discretisation of the advection and diffusion operators (Griffies et al., 2000). Refining the grid resolution decreases the amount of spurious numerical mixing added by the model. In the context of ice shelf cavity modelling, reduced mixing can lead to increased stratification and in turn lower melt rates if the basal slopes are not too steep. Terrain following models, with typically higher resolution at the grounding zone resulting from the squashed vertical grids will have lower numerical mixing and thus resolve more of the stratification effects than a z -layer model run at typical coarse vertical resolutions, as seen in Fig. 7. This means terrain following models in typical configurations tend to calculate lower melt rates. z -layer models, like MITgcm, can produce lower melt rates at higher resolution as shown in Fig. 5. The specified vertical mixing has to be higher than the spurious mixing value for grid convergence of the melt rate, otherwise as the grid is refined the total amount of mixing generated by the model decreases, leading to a reduction in melt rates.

The dependence of melt rate on vertical resolution has been investigated previously in a separate study also using the *Ocean0* experiment (Gwyther et al., 2020). Gwyther et al. (2020) used three ocean models with different vertical coordinate systems (sigma: ROMS, z : COCO, hybrid: MPAS-O) and showed that the two models with constant, 20 m vertical resolution at the boundary specified by ISOMIP+ (COCO and MPAS-O) systematically calculate higher melt rates than the sigma-coordinate based model, which is consistent with the results presented here. Gwyther et al. (2020), following the ISOMIP+ protocol, framed the problem as a matter of tuning the turbulent exchange coefficients from the melt parameterisation to obtain the same melt rates between models. The ISOMIP+ protocol suggests a target melt rate averaging 30 m/a in the grounding zone region (where the ice draft is below 300 m depth). Gwyther et al. (2020) showed that the sigma model ROMS was incapable of reaching this target melt rate purely by

adjusting the exchange coefficients. Here we take a different approach and frame the problem in terms of grid convergence, with the aim of achieving a melt rate that is unaffected by grid resolution and keeps the exchange coefficients fixed. Given the current implementation of the melt parameterisation and turbulence closure (constant eddy diffusivity and viscosity) this means we need the temperature, salinity and velocity fields to remain constant at the boundary as the grid is refined to achieve a fixed melt rate.

Gwyther et al. (2020) show that sigma-coordinate based models can produce higher melt rates with elevated vertical diffusivity at the boundary. Fig. 8 shows integrated melt rates for Firedrake run at different vertical resolutions and vertical diffusivities. The lines in red are the same as from Fig. 5, which is the integrated melt rate at the target ISOMIP+ vertical diffusivity of 5×10^{-5} m^2/s . With this diffusivity the melt rate still depends on grid resolution as discussed. The blue and orange lines in Fig. 8 show the melt rate with diffusivities of 5×10^{-6} m^2/s and 5×10^{-7} m^2/s respectively. Despite the order of magnitude change, there is no significant difference in melt rate. This suggests that numerical mixing dominates when the specified mixing is specified at these low levels. As a result, the melt rate would be expected to be highly dependent on grid resolution and we expect refining the vertical resolution further would reduce melt rates. In contrast, by increasing the vertical diffusivity by ten times the amount of the ISOMIP+ specification to 5×10^{-4} m^2/s , shown by the black lines in 8, the integrated melt rate calculated with 72 layers is very close to the melt rate calculated with 36 layers (the specified number of layers for the ISOMIP+ profile). Since the melt rate appears to be converging towards a consistent value as the grid is refined, the specified diffusivity must be dominating over the spurious, resolution dependent, numerical mixing. This comparison suggests that a rough estimate for the numerical mixing for a Firedrake simulation with the ISOMIP+ target of 36 layers is likely between 5×10^{-6} m^2/s and 5×10^{-5} m^2/s . Since the upper bound has the same value as specified in the ISOMIP+ protocol, this serves to highlight the importance of running ice-ocean simulations at multiple resolutions to determine how sensitive the melt rate is to resolution for a particular model and setup configuration.

The implicit assumption made here is that the diffusivity and viscosity do not vary close to the boundary. This is not necessarily physically




Fig. 8. Integrated melt rate with Firedrake run at different vertical resolutions and vertical diffusivity values. The melt rate is integrated over the deepest region, where ice draft is deeper than 300 m following Asay-Davis et al. (2016). As in Fig. 5 the line style indicates vertical resolution: Dashed — 18 vertical levels outside the cavity, Continuous — 36 vertical levels outside the cavity (target ISOMIP+ resolution), Dotted — 72 levels outside the cavity. The colour of the lines indicates the vertical diffusivity: Black — $K_v = 5 \times 10^{-4} \text{ m}^2/\text{s}$, Red — $K_v = 5 \times 10^{-5} \text{ m}^2/\text{s}$ (target ISOMIP+ diffusivity), Blue — $K_v = 5 \times 10^{-6} \text{ m}^2/\text{s}$, Orange — $K_v = 5 \times 10^{-7} \text{ m}^2/\text{s}$.

justifiable, as Monin–Obukhov theory implies that with strong melting and insufficient shear turbulence to break down the stratification, the mixing rate (and hence the values of diffusivity and viscosity) will be much lower, leading to overestimated melt rates (Vreugdenhil and Taylor, 2019; McPhee, 2008). We have assumed fixed diffusivity and viscosity primarily for reasons of simplicity: it is the ISOMIP+ specification and with a new model we wanted to have confidence that we could understand the general behaviour before implementing more complicated turbulence closure schemes. Even if the model is responsible for explicitly calculating diffusivity and viscosity, the pertinent length scales close to the boundary are very small, sub metre scale (McPhee, 2008), so it may be impractical to account for these effects in large scale ocean simulations. Potentially the melt parameterisation could be modified to incorporate the enhanced stratification effects that arise during melting (similar to the transition of scalar properties across the viscous boundary later, which is accounted for; Holland and Jenkins, 1999), but this leaves the question of where to sample T_w and S_w open. Unless the sampling distance is chosen for physical reasons it seems like it will inherently depend on the grid resolution. It is not obvious how applicable an ice shelf ocean cavity simulation will be for different forcing scenarios or geometrical configurations if the exchange coefficients have been tuned specifically to compensate biases arising from a somewhat arbitrary choice of grid resolution.

Overall, our integrated melt rate of 7.6 m/a for the target ISOMIP+ resolution of 2 km horizontally and 36 vertical layers (orange line Fig. 5) is consistent with, and within the spread of, other published results when using the same turbulent exchange coefficients in the melt parameterisation ($\Gamma_T: 0.011$ and $\Gamma_S = \Gamma_T/35$). Published results for terrain following models with 36 layers include FVCOM: 10 m/a (Figure 9 from Zhou and Hattermann, 2020) and ROMS: 4 m/a (Figure 4 from Gwyther et al., 2020). Similarly, examples of z-level models at high vertical resolution ($dz = 2 \text{ m}$) include COCO: 5 m/a and MPAS-O: 9 m/a (Figure 5 from Gwyther et al., 2020), as well as POP2x: 7 m/a (Figure 9 from Asay-Davis et al., 2016). Also note the MITgcm

integrated melt found here for this comparison is consistent with COCO and MPAS-O, which both gave around 15 m/a with 20 m vertical resolution (Figure 5 from Gwyther et al., 2020).

We also investigated melt rate sensitivity to horizontal resolution within the cavity. We ran Firedrake at two coarser horizontal resolutions, additionally shown in Fig. 5 (purple: 8 km, red: 4 km), to compare with the target ISOMIP+ resolution (orange: 2 km). Although the spatial patterns of melt rate are still broadly equivalent, with highest melting at the grounding line seen in all simulations (not shown), the integrated melt rates have not converged to a consistent result. To investigate the potential of a horizontally unstructured mesh to focus resolution we ran a preliminary simulation with the target ISOMIP+ 2 km resolution only along the western boundary and the grounding zone region, relaxing to 8 km outside the cavity. These areas were chosen to best capture the observed spatial melt pattern, with highest melting at the grounding line, as well as the fast-flowing western boundary current present due to Coriolis. The integrated melt rate for this non-uniform mesh is shown by the black line in Fig. 5. A positive result is that the melt rate predicted with the non-uniform mesh is closer to the value from the uniform 2 km mesh than the 8 km, and indeed the 4 km, uniform mesh. However, considering the Degrees of Freedom (DOFs) associated with each mesh this may not be surprising. The number of combined velocity, pressure, temperature and salinity DOFs is 964,059 for the 8 km mesh, 3,448,681 for the 4 km mesh and 13,693,896 for the 2 km mesh. The non-uniform mesh has 5,011,517 DOFs, which is actually 1.45 times more DOFs than the 4 km mesh. In more complicated cavity geometries, the ability to vary the mesh resolution flexibly, especially in the presence of large channels and crevasses, may prove crucial (Zhou and Hattermann, 2020).

As emphasised in Section 1 one of the main motivations for using the Firedrake framework to simulate ice-ocean interactions is the availability of an automatically generated adjoint model. In the final section we show our first steps at using this capability to investigate sensitivities in an idealised ice shelf cavity.

3.3. Preliminary adjoint sensitivity calculations in an ice shelf cavity

We present preliminary results of sensitivity information obtained with the adjoint to our model. A simplified domain has been chosen to make interpretation of the adjoint sensitivity patterns easier. The domain is a vertical 2D slice within 10 km of an idealised grounding zone domain, with a 2 m wall at the grounding line and a 100 m wall at the open ocean. The domain has a uniform seabed depth of 600 m. The temperature and salinity are initialised with constant temperature and salinity of 1 °C and 34.4, respectively. The temperature and salinity fields are relaxed to these values in a sponge region, which linearly ramps up over the final four grid cells to the right hand side of the domain with a restoring period of one day at the boundary. The grid is made up of triangles arranged in columns, similar to a z-layer discretisation. For an example the reader is referred to Kimura et al. (2013). The grid resolution is 500 m in the horizontal and 2 m in the vertical. Horizontal viscosity is 0.25 m^2/s and vertical viscosity is $1 \times 10^{-3} \text{ m}^2/\text{s}$, matching the grid aspect ratio. Diffusivity for both tracers is set equal to viscosity. We use a P1DG-P2 discretisation for velocity and pressure, and P1DG is used for the tracers. The timestep is 300 s. Table 4 summarises parameters and constants used for the adjoint sensitivity experiment. The forward simulation is run for 50 days to reach a spun-up state. Fig. 9 shows the spun-up state at 50 days. The flow is characteristic of the ice pump mechanism with a clockwise overturning flow, concentrated near the ice ocean boundary, driven by buoyant melt water. At the grounding line the water column is well mixed in the vertical, though the temperature field becomes progressively more stratified away from the grounding line. This type of flow is similar to that described in Holland (2008) where background tidal mixing dominates at the grounding zone if the overturning circulation weakens. In this instance it is likely caused by the coarse grid resolution

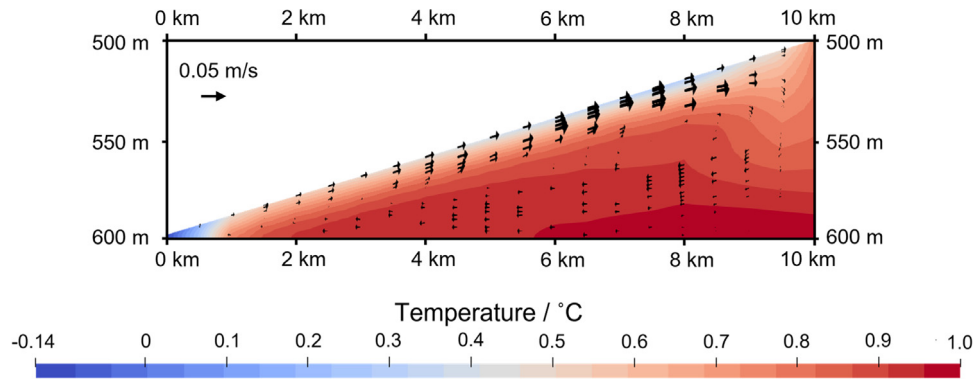


Fig. 9. Temperature and velocity fields at 50 days after the forward model has spun up. After 30 days of simulation time there is negligible change in the fields or melt rate (not shown).

and constant eddy diffusivity values, because at the 2 m high grounding zone wall on the left hand side of the domain there is only one cell to resolve the overturning circulation. It is likely that refining the grid resolution (effectively reducing implicit spurious mixing as described in Section 3.2) or using a turbulence model would change the rate of mixing in this region and the water column would become more (vertically) stratified.

The model is run for a further 20 days with the adjoint model. An objective functional is chosen to be the total basal melt beneath the ice shelf at the final time step. Although using the average basal melt beneath the ice shelf may have been a more conventional quantity, this can be obtained by dividing the results by the length of the ice shelf. We calculate sensitivities of this functional to perturbations of the spatially varying temperature and salinity fields at each time level, as well as to spatially varying perturbations of the viscosity and diffusivity fields which are kept fixed in time. These fields are known as ‘controls’ (Errico, 1997).

The accuracy of the adjoint sensitivity can be verified by a Taylor remainder convergence test (Farrell et al., 2013). If the gradient is correct then with $J(m)$ as the functional, m as the control, and δm as a small perturbation in an arbitrary direction, $|J(m + \delta m) - J(m) - \nabla J \cdot \delta m|$ should converge to zero with second-order accuracy as the magnitude of the perturbation is reduced, i.e. halving the size of the perturbation should cause the result to decrease by a factor of four. We have carried out such Taylor remainder convergence tests for each control field, with the melt functional evaluated after ten timesteps (3000 s) from the 50 day checkpoint. The errors do indeed display convergence at second-order and verify that the gradient information is being calculated accurately by the adjoint model.

Fig. 10 shows adjoint sensitivity of the final total melt flux (note defined here as a volume flux, not a mass flux) to temperature and salinity at three times in the simulation. The adjoint model is effectively a transpose of a linearised version of the full, non-linear forward model. This linear forward model is referred to as the Tangent Linear Model (TLM) and it relates perturbations of the control fields to changes in the objective function. The adjoint model reverses the flow of information (because it is the transpose of the TLM), instead calculating the sensitivity of the objective function to the controls. For time dependent models the adjoint sensitivities thus ‘flow backwards’ in time, from the time when the objective function is evaluated towards the earlier times, provided that the controls had an impact on the objective function (Errico, 1997). The top panel of Fig. 10 shows six hours before the final time, where the sensitivity fields are concentrated at the ice-ocean boundary, where the objective function is evaluated. Since we have evaluated the objective function at the final time, this is equivalent to adding an ‘initial’ condition to the adjoint model. As such even though the forward model is steady, the adjoint sensitivities evolve backwards with time in the opposite sense to the forward model. At two days before the final time (middle panel) the sensitivity patterns

are concentrated closer to the grounding line and at the bottom of the cavity. Finally, five days before the final time (bottom panel) most of the sensitivity of the objective function to the temperature and salinity fields has been lost. The counterclockwise motion of the sensitivity fields is effectively the ice-pump mechanism, but viewed in reverse. The time for the sensitivity fields to be lost is consistent with the forward velocity. Flow speeds are on the order of 0.05 m/s, and an advection path length of 20 km (twice the domain length) gives a circulation time of 4.6 days.

An interesting feature of Fig. 10 is that the signs of adjoint temperature and salinity are opposite. In the same region a small positive perturbation to temperature is required to increase the final total melt flux, whereas an increase in salinity would decrease the final total melt flux. Initially this seems counter intuitive, since increasing salinity decreases the freezing point, which should lead to more melting and thus have the same effect as increasing temperature. This dependency has been found before in an adjoint simulation of Pine Island Glacier ice shelf cavity (Heimbach and Losch, 2012). Heimbach and Losch (2012) suggest that buoyancy effects may account for the sensitivity difference, since increasing temperature and decreasing salinity both increase buoyancy. This in turn should strengthen the overturning ice pump and lead to higher melt rates. Scaling the sensitivity fields by the expansion coefficients in the linear equation of state ($\alpha_T = -2 \times 10^{-4} \text{ }^\circ\text{C}^{-1}$, $\beta_S = 7 \times 10^{-4}$) the two sensitivity fields are almost identical, in sign and magnitude strongly suggesting that buoyancy is controlling the sensitivity pattern. Even though this is a highly idealised experiment, the adjoint model shows potential to investigate the importance of subglacial hydrology and the formation of buoyant plumes at the grounding line and hence the overall effect on melting, as a complimentary tool to the forward model.

Fig. 11 shows patterns of adjoint sensitivity for vertical viscosity, temperature diffusivity and salinity diffusivity. Sensitivity with respect to the vertical components were an order of magnitude stronger than the horizontal components, which is not surprising considering how sensitive the melt rate and hence the flow is to vertical stratification as explored in Section 3.2. Sensitivity for all three fields is concentrated towards the grounding line and the ice-ocean interface. This is intuitive because the overturning flow brings water through the grounding zone region. Results from Yeager (2018) using a $k - \epsilon$ Reynolds Averaged Navier–Stokes (RANS) model to investigate ocean conditions near grounding lines, suggest that changes in diffusivity of several orders of magnitude may occur within 1 km of the grounding line. Even though these large changes in diffusivity would not linearly relate to changes in melt rate, the fact that total melt rates are sensitive to grounding line eddy diffusivity values, which probably vary significantly, highlights the importance of modelling these values accurately.

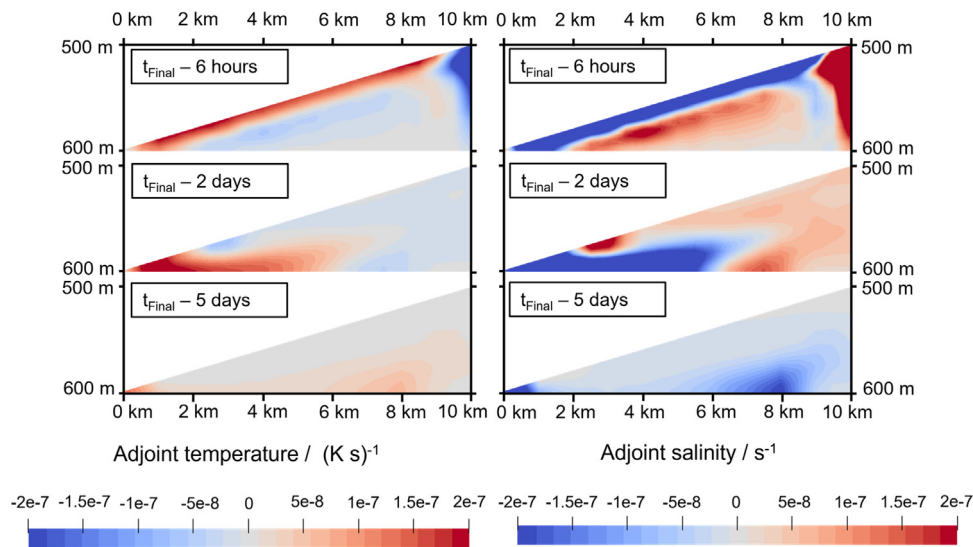


Fig. 10. Adjoint sensitivity fields of final total melt flux with respect to temperature (left) and salinity (right) at three time intervals: 6 h before the final time, two days before the final time and five days before the final time. Note the opposite sign of adjoint temperature and salinity. The sensitivities are calculated in the L_2 norm, $\langle \partial J / \partial m, \delta m \rangle = \delta J$, where $\langle \cdot, \cdot \rangle$ implies taking the inner product and integrating over the domain. By using this integral definition of sensitivity, the sensitivities are not dependent on the mesh resolution (an important consideration if the grid size varies within the mesh). The units for adjoint temperature in this case become $1/(K\ s)$ to ensure that the units of the objective function, J are correct (volume flux: m^2/s) given a temperature perturbation field, δm , with units of K .

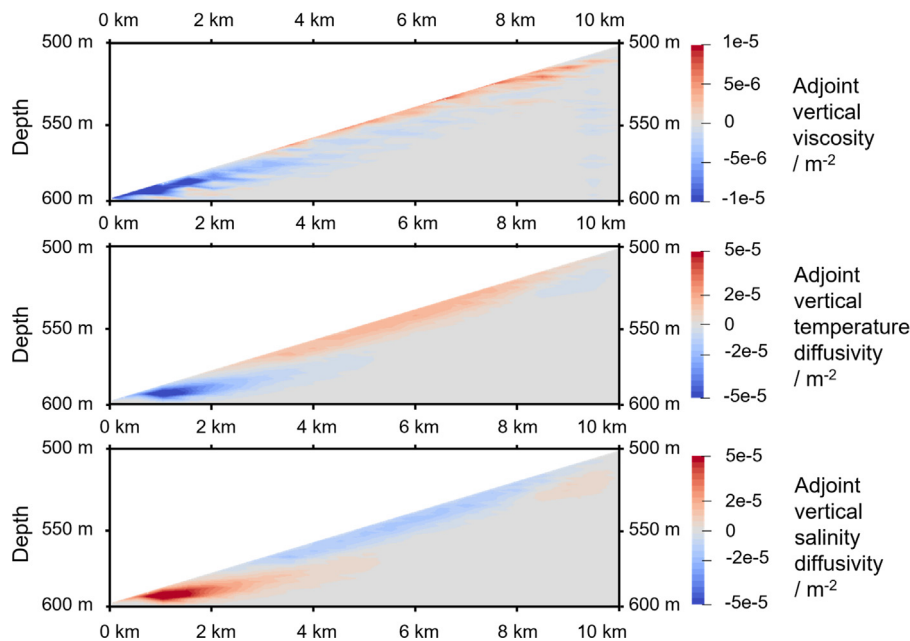


Fig. 11. Sensitivity of total basal melt at the final time step to vertical viscosity and eddy diffusivity for temperature and salinity calculated in the L_2 norm (see the caption of Fig. 10 for an explication of units). Note the change in colour scale for the adjoint vertical viscosity compared with adjoint temperature and salinity vertical diffusivity.

4. Conclusion and future work

We have presented a new model for ocean flow in complex ice shelf cavity environments. By carrying out an MMS test we confirmed that the numerical discretisation of velocity, pressure and tracer fields, and consequently the melt rate, are second order accurate, as expected from our model discretisation choice. We believe that it is important to carry out rigorous code verification checks, for which the MMS technique is ideally suited, because comparing solely with other models may mask model errors. This is especially true considering the strong dependence of the melt parameterisation on grid resolution, as seen in the simulations of the ISOMIP+ test case.

The model is capable of running robustly in 3D domains of a full ice shelf cavity, demonstrated by simulations of the ISOMIP+ test

case. Our simulation results compare favourably when run alongside MITgcm, and the melt rates are consistent with the spread found in the literature (Asay-Davis et al., 2016; Zhou and Hattermann, 2020; Gwyther et al., 2020). We find melt rates are highly sensitive to grid resolution as has been observed before. Ultimately the balance between specified mixing and spurious numerical mixing dictates how sensitive the melt rate is to changes in the grid resolution (and discretisation choice). Our study emphasises the importance of running ice-ocean cavity simulations at multiple resolutions to gauge this sensitivity as it will always be problem specific.

As discussed in Section 3.2, given the assumption of fixed diffusivity and viscosity close to the boundary, finding a melt rate that is insensitive to grid resolution using the existing melt rate parameterisation may not be directly applicable to real ice shelf cavity environments.

A detailed investigation into the melt parameterisation is beyond the scope of this paper, and developing a new, more robust melt parameterisation for large scale applications is not our immediate aim. Studies of the ice-ocean boundary layer are underway that use a hierarchy of models, from Direct Numerical Simulations (Middleton et al., 2021; Couston et al., 2021), to Large Eddy Simulations (Vreugdenhil and Taylor, 2019; Begeman et al., 2022; Vreugdenhil et al., 2022), as well as theoretical models (Jenkins, 2016, 2021). It is likely that these combined approaches will be needed to develop a parameterisation that is consistent with the physics explicitly resolved by the numerical model (and the model's own mixing schemes) and is applicable to the unique environment of the ice shelf-ocean boundary layer where the effects of stratification and basal slope are important. Potentially a model like the one implemented here using Firedrake with flexible grids in the horizontal and vertical will be a useful tool to bridge the gap between high resolution models of boundary flow and the necessarily coarser regional and global scale ocean models.

Ultimately model validation will remain a challenge until more direct observations of sub-ice shelf ocean conditions are made in tandem with measurements of melt rate. Adjoint models are powerful tools to assess model uncertainties and incorporate observations into numerical ocean models (Errico, 1997; Heimbach and Losch, 2012; Goldberg et al., 2020; Nakayama et al., 2021b). We showed preliminary adjoint sensitivity results for an idealised grounding zone domain, which clearly showed a 'reverse' ice-pump mechanism. The link between buoyancy and melt rate also suggests the importance of subglacial outflow into the cavity. The sensitivity of melt rate to vertical mixing (especially at the grounding zone) was also identified by the adjoint model. This again suggests the importance of accurately modelling, or accounting for, mixing within ice shelf cavities. We plan to implement additional turbulence closure schemes such as RANS and LES to investigate these effects, with the caveat that the turbulence closure problem is inherently uncertain and thus challenging to model. In future work, we aim to apply the model to more complicated geometries, in particular ocean flow in realistic grounding zone regions and the impact that basal crevasses have on the ocean flow. We anticipate that the use of fully unstructured meshes will be necessary to resolve these features accurately. We also intend to extend our adjoint capability to solve optimisation problems, to help constrain unknown parameters, such as turbulent mixing coefficients under the ice shelf, based on observations.

CRediT authorship contribution statement

William I. Scott: Conceptualization, Methodology, Software, Validation, Writing – original draft, Visualization. **Stephan C. Kramer:** Conceptualization, Methodology, Software, Writing – review & editing, Supervision. **Paul R. Holland:** Conceptualization, Writing – review & editing Supervision. **Keith W. Nicholls:** Conceptualization, Writing – review & editing, Supervision. **Martin J. Siegert:** Conceptualization, Writing – review & editing, Supervision. **Matthew D. Piggott:** Conceptualization, Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

Acknowledgements

We would like to thank Benjamin Yeager for useful discussions and initial help setting up MITgcm. Simulations were carried out using the high performance cluster of the Research Computing Service at Imperial College London. We would like to thank Xylar Asay-Davies and one other anonymous reviewer for their constructive comments that helped to improve this paper.

This work was supported by the Natural Environment Research Council (NE/S007415/1 and NE/G018391/1). This work is also from the MELT project, a component of the International Thwaites Glacier Collaboration (ITGC). Support from National Science Foundation (NSF: Grant NSF PLR 1739003) and Natural Environment Research Council (NERC: Grant NE/S006656/1 and NE/S006427/1). Logistics provided by NSF-U.S. Antarctic Program and NERC-British Antarctic Survey. ITGC Contribution No. ITGC-097. This work was also supported by the Engineering and Physical Sciences Research Council (EP/R029423/1).

Code availability

The code is available at <https://github.com/thwaitesproject/thwaites.git>. It has also been stored in this Zenodo repository: <https://doi.org/10.5281/zenodo.7584700>, along with instructions for running the examples from Section 3.

References

- Alnæs, M., Blechta, J., Hake, J., Johansson, A., Kehlet, B., Logg, A., Richardson, C., Ring, J., Rognes, M.E., Wells, G.N., 2015. The FEniCS project version 1.5. Arch. Numer. Softw. 3 (100), <http://dx.doi.org/10.11588/ans.2015.100.20553>, URL: <https://journals.ub.uni-heidelberg.de/index.php/ans/article/view/20553>, Number: 100.
- Alnæs, M.S., Logg, A., Ølgaard, K.B., Rognes, M.E., Wells, G.N., 2014. Unified form language: A domain-specific language for weak formulations of partial differential equations. ACM Trans. Math. Software 40 (2), 9:1–9:37. <http://dx.doi.org/10.1145/2566630>.
- Arthern, R.J., Williams, C.R., 2017. The sensitivity of West Antarctica to the submarine melting feedback. Geophys. Res. Lett. 44 (5), 2352–2359. <http://dx.doi.org/10.1002/2017GL072514>, URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/2017GL072514>, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/2017GL072514>.
- Asay-Davis, X.S., Cornford, S.L., Durand, G., Galton-Fenzi, B.K., Gladstone, R.M., Gudmundsson, G.H., Hattermann, T., Holland, D.M., Holland, D., Holland, P.R., Martin, D.F., Mathiot, P., Pattyn, F., Seroussi, H., 2016. Experimental design for three interrelated marine ice sheet and ocean model intercomparison projects: MISIMP v. 3 (MISMIP +), ISOMIP v. 2 (ISOMIP +) and MISOMIP v. 1 (MIS-OMIP1). Geosci. Model Dev. 9 (7), 2471–2497. <http://dx.doi.org/10.5194/gmd-9-2471-2016>, URL: <https://gmd.copernicus.org/articles/9/2471/2016/>, Publisher: Copernicus GmbH.
- Ascher, U.M., Ruuth, S.J., Wetton, B.T.R., 1995. Implicit-explicit methods for time-dependent partial differential equations. SIAM J. Numer. Anal. 32 (3), 797–823. <http://dx.doi.org/10.1137/0732037>, URL: <https://epubs.siam.org/doi/10.1137/0732037>, Publisher: Society for Industrial and Applied Mathematics.
- Avdis, A., Candy, A.S., Hill, J., Kramer, S.C., Piggott, M.D., 2018. Efficient unstructured mesh generation for marine renewable energy applications. Renew. Energy 116, 842–856. <http://dx.doi.org/10.1016/j.renene.2017.09.058>, URL: <https://www.sciencedirect.com/science/article/pii/S0960148117309205>.
- Balay, S., Gropp, W.D., McInnes, L.C., Smith, B.F., 1997. Efficient management of parallelism in object oriented numerical software libraries. In: Arge, E., Bruaset, A.M., Langtangen, H.P. (Eds.), *Modern Software Tools in Scientific Computing*. Birkhäuser Press, pp. 163–202.
- Begeman, C.B., Asay-Davis, X., Van Roekel, L., 2022. Ice-shelf ocean boundary layer dynamics from large-eddy simulations. Cryosphere 16 (1), 277–295. <http://dx.doi.org/10.5194/tc-16-277-2022>, URL: <https://tc.copernicus.org/articles/16/277/2022/>, Publisher: Copernicus GmbH.
- Bercea, G.T., McRae, A.T.T., Ham, D.A., Mitchell, L., Rathgeber, F., Nardi, L., Luporini, F., Kelly, P.H.J., 2016. A structure-exploiting numbering algorithm for finite elements on extruded meshes, and its performance evaluation in Firedrake. Geosci. Model Dev. 9 (10), 3803–3815. <http://dx.doi.org/10.5194/gmd-9-3803-2016>, URL: <https://gmd.copernicus.org/articles/9/3803/2016/>, Publisher: Copernicus GmbH.
- Chorin, A.J., 1967. The numerical solution of the Navier-Stokes equations for an incompressible fluid. Bull. Amer. Math. Soc. 73 (6), 928–931. <http://dx.doi.org/10.1090/S0002-9904-1967-11853-6>, URL: <https://www.ams.org/bull/1967-73-06/S0002-9904-1967-11853-6/>.

- Timmermann, R., Wang, Q., Hellmer, H.H., 2012. Ice-shelf basal melting in a global finite-element sea-ice/ice-shelf/ocean model. *Ann. Glaciol.* 53 (60), 303–314. <http://dx.doi.org/10.3189/2012AoG60A156>, URL: <https://www.cambridge.org/core/journals/annals-of-glaciology/article/iceshelf-basal-melting-in-a-global-finiteelement-seaiceiceshelfocean-model/B8724B766A10BD37FBB6B3F6C70F9089>, Publisher: Cambridge University Press.
- Utke, J., Naumann, U., Fagan, M., Tallent, N., Strout, M., Heimbach, P., Hill, C., Wunsch, C., 2008. OpenAD/F: A modular open-source tool for automatic differentiation of fortran codes. *ACM Trans. Math. Software* 34 (4), 18:1–18:36. <http://dx.doi.org/10.1145/1377596.1377598>, URL: <https://doi.org/10.1145/1377596.1377598>.
- Vallis, G.K., 2017. *Atmospheric and Oceanic Fluid Dynamics*. Cambridge University Press, Google-Books-ID: BTgoDwAAQBAJ.
- Vreugdenhil, C.A., Taylor, J.R., 2019. Stratification effects in the turbulent boundary layer beneath a melting ice shelf: Insights from resolved large-eddy simulations. *J. Phys. Oceanogr.* 49 (7), 1905–1925. <http://dx.doi.org/10.1175/JPO-D-18-0252.1>, URL: <https://journals.ametsoc.org/view/journals/phoc/49/7/jpo-d-18-0252.1.xml>, Publisher: American Meteorological Society Section: Journal of Physical Oceanography.
- Vreugdenhil, C.A., Taylor, J.R., Davis, P.E.D., Nicholls, K.W., Holland, P.R., Jenkins, A., 2022. The ocean boundary layer beneath larsen c ice shelf: Insights from large-eddy simulations with a near-wall model. *J. Phys. Oceanogr.* 52 (8), 1903–1926. <http://dx.doi.org/10.1175/JPO-D-21-0166.1>, URL: <https://journals.ametsoc.org/view/journals/phoc/52/8/JPO-D-21-0166.1.xml>, Publisher: American Meteorological Society Section: Journal of Physical Oceanography.
- Wallwork, J.G., Barral, N., Kramer, S.C., Ham, D.A., Piggott, M.D., 2020. Goal-oriented error estimation and mesh adaptation for shallow water modelling. *SN Appl. Sci.* 2 (6), 1053. <http://dx.doi.org/10.1007/s42452-020-2745-9>, URL: <https://doi.org/10.1007/s42452-020-2745-9>.
- Yeager, B., 2018. Moving mesh finite element modeling of ocean circulation beneath ice shelves. Imperial College London, URL: <http://spiral.imperial.ac.uk/handle/10044/1/73697>.
- Zhou, Q., Hattermann, T., 2020. Modeling ice shelf cavities in the unstructured-grid, finite volume community ocean model: Implementation and effects of resolving small-scale topography. *Ocean Model.* 146, 101536. <http://dx.doi.org/10.1016/j.ocemod.2019.101536>, URL: <https://www.sciencedirect.com/science/article/pii/S1463500319301738>.