https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

Variational Stokes method applied to free surface boundaries in numerical geodynamic models

Timothy Gray[1], Paul Tackley[1], and Taras Gerya[1]

[1]Institute of Geophysics, Department of Earth and Planetary Sciences, ETH Zürich, Sonneggstrasse 5, 8092 Zurich, Switzerland

Correspondence: Timothy Gray (timothy.gray@eaps.ethz.ch)

Abstract.

Accurately and efficiently modelling topographic evolution is a key challenge in geodynamic modelling, which requires the solution of the Stokes equations with free surface boundary conditions. While finite difference methods on staggered grids, as used in geodynamic modelling codes such as StagYY, I3ELVIS and LaMEM, offer strong computational performance and compatibility with multigrid solvers, the use of fixed Eulerian grids complicates the implementation of realistic, deformable free surfaces. Two existing methods are available to model free surface boundary conditions in StagYY: the commonly used sticky-air method, which suffers from limitations relating to high viscosity contrasts, and the "staircase" method, which improves upon the sticky air method by imposing free surface boundary conditions at cell boundaries.

To address the limitations of existing methods of implementing free surface boundary conditions, this study investigates an alternative variational discretisation of the Stokes equations that uses volume fractions to represent a smooth surface within a fixed Eulerian grid, allowing the imposition of accurate free surface boundary conditions while allowing it to bypass the limitations of existing free surface discretisation methods.

The variational Stokes method is demonstrated to be an accurate and computationally efficient alternative to existing methods. It reproduces results comparable to existing methods while reducing computational cost and enabling broader applications, including non-zero surface tractions, complex surface loading, and compatibility with 3D spherical geometries.

1 Introduction

The Stokes equations are the governing equations of mantle convection, and their solution is the key goal of numerical geodynamic models. Numerous discretisation methods exist for solving these equations numerically, including the finite difference method, finite volume method, finite element method, and spectral element method. Among these, the finite difference method on a staggered grid, which is equivalent to the finite volume method, is often used in geodynamic modelling codes (e.g. (Patankar, 1980; Ogawa et al., 1991; Tackley, 1993; Trompert and Hansen, 1996; Gerya and Yuen, 2007; Tackley, 2008; Gerya et al., 2015a; Kaus et al., 2016)) due to its simplicity in discretising the governing equations and compatibility with efficient multigrid solvers in both 2D and 3D.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

A key component of the finite difference method is the transformation of continuous differential equations into a discretised system that can be solved numerically on a fixed Eulerian grid. However, while the finite difference method offers computational efficiency and stability, it introduces challenges when modelling free surface boundary conditions, which are essential for the accurate modelling of topographic evolution on planets. The fixed grids used by the above geodynamical modelling codes are not readily adaptable to modelling surface deformation, necessitating alternative techniques to account for free surface boundary conditions.

Modelling free surface boundary conditions on a fixed Eulerian grid requires two key components: a method for tracking the location of the surface over time, and a method of discretising the Stokes equations in a way that incorporates the free surface boundary condition, with this study focusing on on the latter.

The most straightforward approach for achieving this is the so-called sticky air method. In this method, a relatively low-viscosity layer is added above the material of interest to represent the air (Gerya (2019) and references therein). The presence of this low viscosity layer allows the surface to deform relatively freely and thus approximate a true free surface boundary condition, however computational limitations on viscosity contrast mean that the air layer is unable to be set to the extremely low viscosity of actual air, hence the use of a higher viscosity "sticky" air layer. Although conceptually simple, this method presents several drawbacks. Surface deformation when using a sticky air method is sensitive to the thickness and viscosity of the sticky air layer, with the most accurate results resulting from a thick air layer and a high viscosity contrast (Crameri et al., 2012a). Increasing thickness of the sticky air layer by adding new computational cells to the model adds degrees of freedom to the model, which incurs additional computational cost. Meanwhile, high viscosity contrasts between the air layer and underlying material can to poor numerical performance, particularly when using iterative solvers, significantly increasing computation time (Duretz et al., 2016).

To address the limitations of the sticky air method, Duretz et al. (2016) introduced a novel discretisation method, referred to in this chapter as the "staircase" method. This approach eliminates the need for modelling a sticky air layer by applying zero Dirichlet boundary conditions on pressures and velocities within the air layer, while allowing the surface itself to deform freely via alternative boundary conditions in the interfacial cells. Although this method improves numerical performance by allowing for a relatively thinner air layer and avoiding the viscosity contrast inherent in the sticky air method, it introduces its own set of challenges. The staircase representation can lead to inaccuracies at coarse grid resolutions, and like the sticky air method, can hamper the efficiency of multigrid solvers due to the misalignment of surface at coarse and fine levels. In practice, this often forces solvers to take iterations on a fine grid when using an iterative solver, which can undermine the method's computational efficiency. Furthermore, the existing implementation of this method in StagYY is often susceptible to convergence issues when using multigrid solvers, which further undermines its potential to replace the sticky air method.

Given these challenges, there is a need for a more robust and efficient method to discretise the Stokes equations for the modelling of free surfaces. Such a method should satisfy several key criteria:

- Remove or minimise the viscosity contrast dependence inherent in the sticky air method.
- Remove or minimise the necessity for a thick air layer.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

– Enable efficient and stable solutions in both 2D and 3D, using both direct and iterative (multigrid) solvers, and on both Cartesian and non-Cartesian grids.

A promising direction in this field is the so-called variational Stokes method, introduced by Larionov et al. (2017). This method is conceptually similar to the aforementioned staircase method in that it eliminates the need to explicitly model a low viscosity sticky air layer through the use of alternative boundary conditions near the surface. However, instead of a discrete staircase representation of the surface, this method proposes a smooth representation of free surfaces through the use of volume fractions, which can be readily computed in one of several ways, including a volume of fluid method or a direct Lagrangian representation of the surface. This chapter primarily discusses the implementation of this method into StagYY, the comparison of this method the existing sticky air and staircase methods, and the potential applications of such methods to global scale numerical geodynamic models.

2 Existing methods

Modelling surface deformations accurately in geodynamic models requires the solution of the Stokes equations for fluid flow with a free surface boundary condition (Crameri et al., 2012b). The Stokes equations are:

$$-\nabla p + \nabla \cdot \boldsymbol{\tau} + \rho \mathbf{g} = \mathbf{0}, \tag{1}$$

with pressure p, deviatoric stress tensor $\boldsymbol{\tau}$, density ρ and acceleration due to gravity $\mathbf{g}$.

The free surface boundary condition is that the traction vector acting at the free surface is zero:

$$\boldsymbol{\sigma} \cdot \mathbf{n} = \mathbf{0} \quad \text{on } \partial\Omega_{\text{free surface}}, \tag{2}$$

where $\boldsymbol{\sigma}$ is the Cauchy stress tensor, consisting of pressure p and deviatoric stress tensor $\boldsymbol{\tau}$ where $\boldsymbol{\sigma} = -p\mathbf{I} + \boldsymbol{\tau}$, and $\mathbf{n}$ is the surface normal vector. Conceptually, this can be interpreted as the idea that the modelled fluid, which is typically rock in the case of geodynamic modelling, should not feel any resistance to motion at the surface, either normal or shear, and therefore is able to deform freely.

Two methods for numerically modelling free surface boundary conditions in an Eulerian grid are already implemented in StagYY: the sticky air method (Crameri et al., 2012a), and the staircase method (Duretz et al., 2016).

2.1 Sticky air method

The sticky air method involves the use of a layer of low viscosity material at the top of the model domain, which represents an air layer within which the surface can freely deform. At standard temperature and pressure, real air has a viscosity of approximately 1.8×10^{-5} Pa s — over 20 orders of magnitude lower than typical mantle rock viscosities, which typically range between $10^{19} - 10^{20}$ Pa s. However, the naive solution of applying such a large viscosity contrast to a numerical model

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

would cause the resulting linear system to become extremely poorly conditioned, resulting in low quality results and poor numerical performance (Gerya, 2019).

For this reason, an effective air viscosity is used that is significantly higher than that of actual air, typically on the order of 1,000 times lower then the viscosity of rock in the model, typically resulting in a modelled viscosity on the order of $10^{16}-10^{18}$ Pa s. This enhances numerical performance significantly, while still approximating free deformations of the surface. However, the sticky air approach suffers from limitations in cases when the viscosity contrast or the thickness of the air layer is insufficient. For certain analytical cases, simple criteria can be derived which can be used to determine whether the surface will be traction free as a function of the model geometry and viscosity contrast (Crameri et al., 2012a). These criteria take the form:

$$A\frac{\eta_{\text{air}}}{\eta_{\text{rock}}}\left(\frac{L}{h_{\text{air}}}\right)^3 \ll 1, \tag{3}$$

where η_{air} and η_{rock} are the characteristic viscosities of the sticky air and underlying rock, h_{air} is the sticky air layer thickness, L is a characteristic length (typically the mantle thickness in global scale models). The constant A is a model-dependent parameter which can be derived analytically for a number of simple scenarios (Crameri et al., 2012a). From this relation, it can be concluded that both air layer thickness and viscosity contrast must be sufficiently high in order to approximate a traction-free surface.

The presence of a sticky air layer introduces several issues. The first issue is the requirement to have a thick layer of low viscosity material. This inherently increases the model domain size, and thus decreases numerical performance by having to solve for a greater number of degrees of freedom than are otherwise necessary (Duretz et al., 2016). Furthermore, as the fluid velocities within the low-viscosity air layer are often higher than the high-viscosity rock, timesteps may be limited unnecessarily in order to accommodate deformation of the air layer. As the air layer itself is uninteresting for researchers, serving only to enable a free surface boundary condition, this implies that computational effort is wasted.

The second and most prominent issue relates to the sharp viscosity contrast necessary to obtain accurate solutions when using the sticky air method. The introduction of viscosity contrasts into a model makes the discretised form of the Stokes equations poorly conditioned, and generally more difficult to solve. While direct solvers, which are more commonly used in 2D models, are not as sensitive to sharp viscosity contrasts, iterative solvers including the iterative multigrid approach used for solving 3D models are strongly affected by them. This increase in computational cost has contributed to making use of the sticky air method, and due to the ubiquity of the method, free surface modelling in general less desirable when using 3D models in StagYY.

2.2 Staircase method

In an attempt to overcome the limitations of the sticky air approach, Duretz et al. (2016) proposed an alternative method of modelling surfaces in geodynamic models, referred to in this chapter as the staircase method. This approach eliminates the need for a sticky air layer by avoiding equations in the air domain and instead modelling a free surface boundary condition using

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

modified boundary conditions at interfacial cells. In this way, it is possible to model a true free surface boundary condition without the need for a thick, low-viscosity sticky air layer.

The method relies on there existing a method of tracking the location of the free surface. While Duretz et al. (2016) use a simple marker chain to directly track the interface location over time, in principle any method of tracking free surfaces can be used to do so. Such surface tracking methods implemented in StagYY include the marker-in-cell method, the use of Lagrangian surface tracking methods or the volume of fluid method.

In the staircase method, computational cells with a centre point that is below the location of the surface are considered to be fluid (i.e. rock for geodynamic applications) cells, while cells above are considered to be air cells. These two cell classifications imply the existence of an interface between the cells which follows cell boundaries, and which therefore follows a staircase pattern when using a Cartesian grid (Figure 1). In air cells that are not directly adjacent to the interface, no equations are solved and the pressure and velocities are set to zero through Dirichlet boundary conditions. In rock cells not directly adjacent to the surface, the Stokes equations are discretised as they would be normally, with no additional modifications required. At interfacial cells, the equations are modified to ensure that tractions are zero at the boundaries, thus imposing free surface boundary conditions directly at the interfacial cell boundaries.

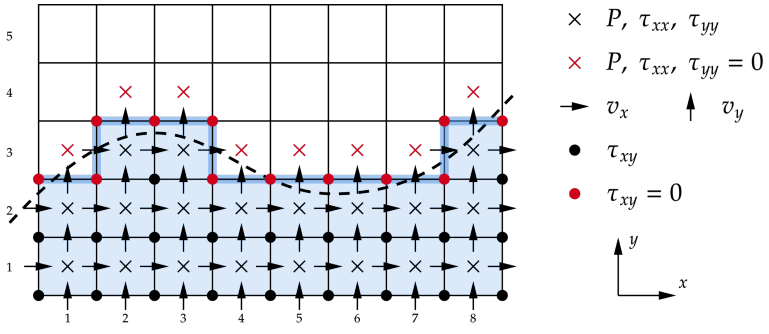

Figure 1. In the staircase method, cells are divided into rock (blue) and air (white). Between them is an interface, at which normal and shear stresses are set to zero in order to create free surface boundary conditions.

The discretised Stokes equations can be formulated in terms of a residual which is to be minimised within each cell. In this form, the $x-$ and $y-$momentum equations in 2D can be stated as (Duretz et al., 2016):

$$F_{v_x} = \frac{1}{h_x}\left(\tau_{xx}^E - \tau_{xx}^W\right) + \frac{1}{h_y}\left(\tau_{xy}^N - \tau_{xy}^S\right) - \frac{1}{h_x}\left(p^E - p^W\right) - F^x_{\text{Dirichlet}} - F^x_{\text{Neumann}} \tag{4}$$

$$F_{v_y} = \frac{1}{h_y}\left(\tau_{yy}^N - \tau_{yy}^S\right) + \frac{1}{h_x}\left(\tau_{yx}^E - \tau_{yx}^W\right) - \frac{1}{h_y}\left(p^N - p^S\right) - \rho g_y - F^y_{\text{Dirichlet}} - F^y_{\text{Neumann}} \tag{5}$$

Where F_{v_x} and F_{v_x} are residuals to be minimised by the solver, τ_{xx}, τ_{yy}, τ_{xy}, and τ_{yx} refer to the components of the deviatoric stress tensor $\boldsymbol{\tau}$, p is the pressure, ρ is the density, g_y is acceleration due to gravity in the y direction, h_x and h_y refer to the grid spacing in the x and y directions, and $F^{x/y}_{\text{Dirichlet}}$ $F^{x/y}_{\text{Neumann}}$ refer to Dirichlet and Neumann boundary conditions

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

respectively. The superscripts N, S, E, and W denote the cardinal directions: north, south, east, and west respectively. With a staggered grid discretisation, these equations are solved at face centred points.

This discretised form of the Stokes equations remains unchanged within the rock layer, and is only modified at interfacial cells, i.e. rock cells that have one or more boundary with an air cell. As an example, for cell (3, 3) in Figure 1, the equations are modified to reflect the fact that the stress at the north and east boundary of the cell are zero. This is reflected in the modified stencil for the $x-$ and $y-$momentum equations at those interfacial points:

$$F_{v_x} = \frac{1}{h_x}\left(-\tau_{xx}^{W}\right) + \frac{1}{h_y}\left(-\tau_{xy}^{S}\right) - \frac{1}{h_x}\left(p^{E} - p^{W}\right) - F^{x}_{\text{Dirichlet}} - F^{x}_{\text{Neumann}} \tag{6}$$

$$F_{v_y} = \frac{1}{h_y}\left(-\tau_{yy}^{S}\right) + \frac{1}{h_x}\left(-\tau_{yx}^{E}\right) - \frac{1}{h_y}\left(p^{N} - p^{S}\right) - \rho g_y - F^{x}_{\text{Dirichlet}} - F^{x}_{\text{Neumann}} \tag{7}$$

Cells above the free surface do not involve the solution of any equations, and therefore a simple implementation of setting all stresses to zero would lead to zero rows and a singular matrix when deriving the linear system of equations. In order to make the equations non-singular, air cells not directly adjacent to rock cells must be identified and Dirichlet conditions applied to explicitly set the pressures and velocities in those cells to zero.

2.2.1 Implementation in StagYY

While we here focus on the mantle convection code StagYY (Tackley, 2008), the implementation and related considerations are broadly applicable to codes using the same discretisation, many of which were cited earlier.

The staircase method is implemented in StagYY by imposing zero viscosities at cell centres in the cells that have been flagged as air cells. For viscosities centred at cell edges, alternative viscosity interpolation routines are used in order to ensure that viscosities on the boundary are also set to zero. Setting these viscosities to zero has the same effect as setting the stresses at those locations to zero. Furthermore, pressures within the air layer are also set to zero, as the continuity equation no longer applies above the surface.

In StagYY, tracking the free surface location typically requires the use of Lagrangian tracers to determine rock and air distributions within the model. The advection of these tracers necessitates interpolation of velocities from the Eulerian grid to the positions of individual tracers. StagYY provides two interpolation methods: a first-order method using bilinear interpolation of the four neighbouring velocity points, and a divergence-free method that uses modified quadratic splines in the direction parallel to each velocity component and linear interpolation in the directions perpendicular to each velocity component, therefore involving six neighbouring velocity points.

However, when employing the staircase discretisation, interpolating velocities for tracers at or near the surface presents a challenge, as velocities in cells above the interface are not explicitly calculated, which would lead to errors when trying to determine velocities near the surface. Two approaches can be used to address this issue:

Velocity copying In this method, velocities from the highest cell where they are solved for are simply copied upwards into the air layer. This approach is valid because stresses within the air layer are zero, implying that velocity gradients in

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

this region should also be zero. However, while this method is computationally straightforward, it does not inherently preserve the continuity equation. This can result in issues such as tracer bunching or the unintended advection of tracers outside the model domain unless corrected by other methods.

Continuity-preserving velocity field This method ensures that the continuity equation is satisfied within the air layer without solving the Stokes or continuity equations there explicitly. Starting with the horizontal velocity at the highest cell at which it was solved for, a linear horizontal velocity profile is imposed for each column of cells within the air layer. Subsequently, vertical velocities are computed for each cell in the column, beginning from the surface, such that the continuity equation is satisfied in every air cell. This approach generates a velocity field in the air layer that is both physically consistent and free from the bunching of the copying method.

While originally implemented only for 2D geometries (Duretz et al., 2016), the method was later extended to 3D in StagYY (Tackley, 2021). Furthermore, due to the method's generality, it was able to be extended to non-Cartesian geometries, including the yin-yang geometry used by StagYY to model full 3D models on the global scale.

2.2.2 Limitations of the staircase method

The staircase implementation of the surface comes with some limitations. The first limitation is one of accuracy. By imposing free surface boundary conditions at a staircase boundary between cells, it is able to approximate the solution to the free surface problem at high resolutions. Unlike the sticky air method, however, which implicitly accounts for surface position at a sub-grid level through the viscosity field, the staircase method's discrete interface is less accurate, especially at coarser resolutions.

While Duretz et al. (2016) did not demonstrate this issue explicitly, their study compared the staircase method to a similarly discretised sticky air method that also used a discrete viscosity jump between air and rock layers. When comparing the staircase method to a sticky air method with a smooth viscosity transition that accounts for the sub-grid location of the surface, the staircase method performs less favourably, introducing a systematic error that persists across all resolutions. These results can be seen in Section 5.2.

Another limitation of the staircase method is the absence of explicit velocity solutions above the surface. This creates challenges when using Lagrangian tracers near the surface, as velocity interpolation from the Eulerian grid requires contributions from neighbouring grid points. When air layer velocities are set to zero, inaccuracies arise in the advection of tracers near the surface. While the two velocity field extrapolation methods presented above provide a partial solution to this problem, an ideal method would be able to capture near-surface velocities without recourse to extrapolation.

A further limitation of the staircase method is its numerical performance and stability. Models employing the staircase method often fail due to numerical instabilities that are generally absent when using a sticky air layer. These instabilities are particularly pronounced with iterative multigrid solvers and are likely caused by the misalignment of surface at coarse and fine levels, as well as the aforementioned inaccuracies of the discrete representation of the surface, particularly when restricting to coarse grids.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

3 Variational Stokes method

A discretisation of the Stokes equation on staggered grids that implicitly models free surface boundary conditions was proposed by Larionov et al. (2017), henceforth referred to as the variational method. The method is conceptually similar to the staircase method in that it does not involve the solution of any equations within the air layer, instead modelling the free surface boundary condition through the use of alternative boundary conditions at the interface. The methods differ in how they implement these alternative boundary conditions. While the staircase method used a discrete staircase representation of the surface at cell boundaries, the variational method attempts to construct a smooth representation of the free surface boundary condition at the interface by using volume fractions.

3.1 Variational form of the Stokes equations

The variational form of the Stokes equations considers the steady-state, incompressible Stokes equations in the absence of body forces with variables velocity $\mathbf{v}$, pressure p, and the deviatoric stress tensor $\boldsymbol{\tau}$. The continuity and momentum equations respectively are given by:

$$\nabla \cdot \mathbf{v} = 0, \tag{8}$$

$$-\nabla p + \nabla \cdot \boldsymbol{\tau} = \mathbf{0}, \tag{9}$$

where the deviatoric stress tensor $\boldsymbol{\tau}$ is can be expressed in terms of velocities by the constitutive law for an incompressible Newtonian viscous fluid:

$$\boldsymbol{\tau} = \eta \left(\nabla \cdot \mathbf{v} + (\nabla \cdot \mathbf{v})^T \right). \tag{10}$$

The equations as presented can be reformulated in a variational form, where the problem is expressed as a functional to be maximised or minimised in order to recover the original equations (Equation 11). This formulation differs slightly from that derived by Larionov et al. (2017), as the original authors considered the time-dependent Stokes equations.

$$\max_{p,\boldsymbol{\tau}} \min_{\mathbf{v}} \iiint_{\Omega_F} -p \nabla \cdot \mathbf{v} + \boldsymbol{\tau} : \left(\frac{\nabla \mathbf{v} + (\nabla \mathbf{v})^T}{2} \right) - \frac{1}{4\eta} ||\boldsymbol{\tau}||_F^2 \, dV. \tag{11}$$

where the integral is computed over the total volume of the fluid domain Ω_F, and $||\boldsymbol{\tau}||_F$ refers to the Frobenius norm of $\boldsymbol{\tau}$. It can be confirmed by differentiating with respect to the primitive variables that the original steady-state, incompressible Stokes equations are recovered under this formulation. For example, differentiating with respect to pressure p recovers the incompressible continuity equation, differentiating with respect to the deviatoric stress tensor $\boldsymbol{\tau}$ recovers the constitutive law between stress and velocity, and differentiating with respect to velocity recovers the momentum equations.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

The domain of integration is the entire fluid filled domain, denoted by Ω_F, which does not include regions containing air. The domain boundaries $\partial\Omega$ therefore include the solid boundaries surrounding the model, as well as the location of the free surface which can be tracked with an appropriate method. It can be shown using the divergence theorem that the variational formulation of Equation 11 will not only provide the solution to the original Stokes equations, but will automatically do so subject to the free surface boundary conditions (Larionov et al. (2017), supplemental material). Thus, the boundaries satisfy the condition:

$$(-p\mathbf{I}+\boldsymbol{\tau})\cdot\mathbf{n}=0 \quad \text{on } \partial\Omega. \tag{12}$$

The variational formulation (Equation 11) may be discretised, remaining formulated in terms of a variational problem:

$$\max_{p,\boldsymbol{\tau}}\min_{\mathbf{v}} \mathbf{p}^T\mathbf{G}^T\mathbf{v}+\boldsymbol{\tau}^T\mathbf{D}\mathbf{v}-\frac{1}{4}\boldsymbol{\tau}^T\mathbf{M}^{-1}\boldsymbol{\tau} \tag{13}$$

Where the matrix $\mathbf{G}$ refers to the discrete gradient operator, and due to the adjoint relationship between the discrete gradient operator and the discrete divergence operator, $\mathbf{G}^T\mathbf{v}$ refers to discrete divergence of $\mathbf{v}$. The matrix $\mathbf{D}$ is the discrete form of the symmetric gradient operator $\frac{1}{2}(\nabla+\nabla^T)$, while the matrix $\mathbf{M}$ is a diagonal matrix of viscosity coefficients located at the appropriate stress locations corresponding to the structure of $\boldsymbol{\tau}$. Again, evaluating the derivatives with respect to the primary variables recovers the original discretised Stokes equations in the form of a system of equations. By differentiating with respect to velocity $\mathbf{v}$, deviatoric stress $\boldsymbol{\tau}$, and pressure $\mathbf{p}$ respectively produces three sets of equations which can be expressed as a block matrix:

$$\begin{bmatrix} 0 & \mathbf{D}^T & \mathbf{G} \\ \mathbf{D} & -\frac{1}{2}\mathbf{M}^{-1} & 0 \\ \mathbf{G}^T & 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \boldsymbol{\tau} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}. \tag{14}$$

By rearranging the central row as:

$$2\mathbf{M}\mathbf{D}\mathbf{v}=\boldsymbol{\tau}, \tag{15}$$

it is possible to eliminate the central block of equations, resulting in a more familiar form of the discretised Stokes equations in terms of velocities and pressures:

$$\begin{bmatrix} 2\mathbf{D}^T\mathbf{M}\mathbf{D}\mathbf{v} & \mathbf{G} \\ \mathbf{G}^T & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}. \tag{16}$$

With appropriate discretised operators $\mathbf{G}$ and $\mathbf{D}$, this results in an identical set of equations to the velocity-pressure form of the Stokes equations as solved by StagYY and many other numerical geodynamic codes.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

3.2 Volume fractions

Further progress on the free surface problem requires the use of cellwise volume fractions. The volume fraction function, denoted by ϕ, is defined as the fraction of material in a given computational cell, with the material for geodynamic applications being rock. The volume fraction function is therefore 0 if the cell is completely empty, and 1 if it is completely full (Figure 2). In the vicinity of the surface, partially filled cells can be found.

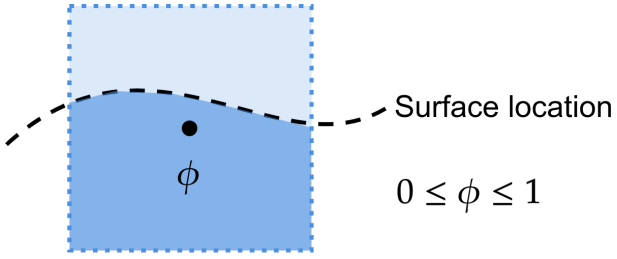

Figure 2. The volume fraction function ϕ for a given control volume is defined as 0 if the cell is completely empty, and 1 if it is completely full.

In a staggered grid discretisation, variables are not necessarily found only at cell centres but staggered at cell centres, faces, edges, and nodes. The volume fractions associated with variables at their appropriate locations is denoted by ϕ^P for pressure points, which are cell centered, and ϕ^τ for stresses, which may be located at cell vertices or cell centres. For example, volume fractions associated with normal stresses $\phi^{\tau_{xx}}$ and $\phi^{\tau_{yy}}$ are located at cell centres and are identical to ϕ^P, while $\phi^{\tau_{xy}}$ is located at cell vertices in 2D, or at edge centres in 3D.

A key result of the work of Larionov et al. (2017) is that the free surface problem can be treated by modifying the discretised form of the variational Stokes formulation using volume fractions multiplied by the original variables. This is motivated by the fact that the continuous variational formulation in 11 includes volume integrals over the fluid filled domain, which when discretised can be approximated using volume fractions. The full discretised form of the free surface problem can be given by:

$$\max_{p,\tau} \min_{\mathbf{v}} \mathbf{p}^T \phi^P \mathbf{G}^T \mathbf{v} + \boldsymbol{\tau}^T \phi^\tau \mathbf{D}\mathbf{v} - \frac{1}{4} \boldsymbol{\tau}^T \mathbf{M}^{-1} \phi^\tau \boldsymbol{\tau} \tag{17}$$

Which when differentiated with respect to velocity $\mathbf{v}$, deviatoric stress $\boldsymbol{\tau}$, and pressure $\mathbf{p}$ respectively produces a block matrix form of the Stokes equations including volume fractions:

$$\begin{bmatrix} 0 & \mathbf{D}^T \phi^\tau & \mathbf{G} \phi^P \\ \phi^\tau \mathbf{D} & -\frac{1}{2} \mathbf{M}^{-1} \phi^\tau & 0 \\ \phi^P \mathbf{G}^T & 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \boldsymbol{\tau} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}. \tag{18}$$

As before, is possible to eliminate the stress equations in order to rearrange the problem into a formulation suitable for a pressure-velocity solver by using the fact that:

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

$$2\mathbf{M}\phi^{\tau}\mathbf{D}\mathbf{v} = \phi^{\tau}\boldsymbol{\tau}. \tag{19}$$

The pressure-velocity formulation is then given by:

$$\begin{bmatrix} 2\mathbf{D}^T\mathbf{M}\phi^{\tau}\mathbf{D} & \mathbf{G}\phi^P \\ \phi^P\mathbf{G}^T & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \tag{20}$$

As the volume fraction function is zero when a cell is completely outside the fluid domain, this formulation indicates that the discretised continuity equation $\mathbf{G}^T\mathbf{v} = \mathbf{0}$ should only be solved when the corresponding cell centred volume fraction is nonzero. For cases in 2D, Figure 3 demonstrates the location of the control volume that must be nonzero for the continuity equation to be solved at that point. Likewise when solving the momentum equations, no equations will be solved when both ϕ^{τ} and ϕ^P are simultaneously zero. In both of these cases, the system of equations becomes singular. To avoid this, in all equations where a zero is encountered, the pressure or velocity is set to zero. The locations of the volume fractions surrounding v_x and v_z points in 2D are shown in Figure 4. Only one of the pictured volume fractions is required to be nonzero in order for the momentum equation to be defined in the cells pictured.

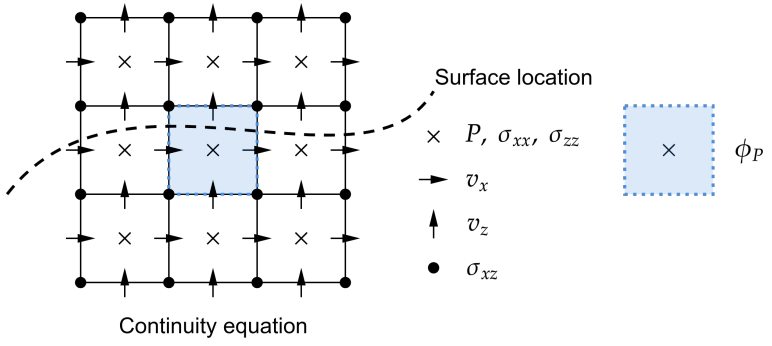

Figure 3. The continuity equation is solved in a given cell if the volume fraction centred on the central pressure point is nonzero, i.e. $\phi_P > 0$.

The result of this is a set of equations that does not include the solution of any pressures in cells that do not have any part within the fluid domain, and also does not include the solution of velocities that are more than 1.5 cell distances away from the surface. The discretisation of the same problem shown for the staircase method in Figure 1 is repeated for the variational Stokes method in Figure 5.

A benefit of this discretisation over the staircase method is that it does explicitly calculate velocities within 1.5 cell distances of the surface. This means that interpolating velocities directly at the surface, either using a bilinear or mixed quadratic spline approach, is possible without the need to extrapolate surface velocities.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

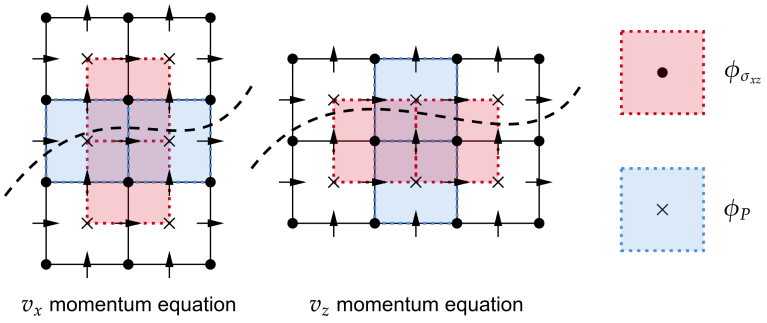

Figure 4. In 2D, The v_x and v_z momentum equations are solved only if at least one of the pictured volume fractions surrounding the velocity point are nonzero.

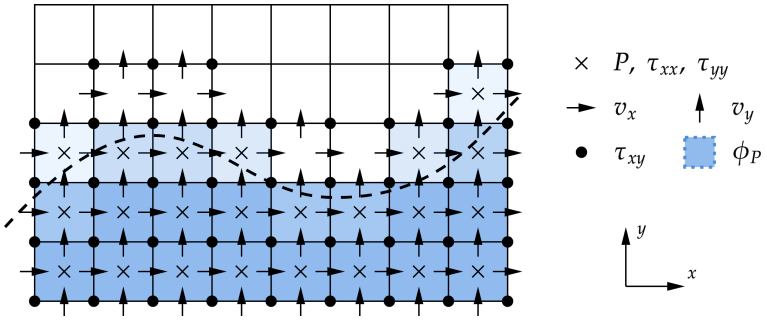

Figure 5. The discretised Stokes equations on a simplified domain when using the variational method. Cell centred (pressure point) volume fractions ϕ_P dictate at which location the continuity equation is solved; cells with zero volume fraction (white) are not solved.

3.3 Right hand side

Adding a right hand side to the Stokes equations in variational form is achieved trivially and requires no special modification to the method. However, for an approach consistent with the method used for discretising the Stokes solver, it is recommended to consider the density as the product of background density and an appropriate volume fraction function.

This results in the following system of equations:

$$\begin{bmatrix} 2\mathbf{D}^T\mathbf{M}\phi^\tau\mathbf{D} & \mathbf{G}\phi^P \\ \phi^P\mathbf{G}^T & 0 \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \phi^{v_z}\rho\mathbf{g} \\ 0 \end{bmatrix}, \tag{21}$$

where ϕ^{v_z} is the volume fraction determined at z (vertical) velocity points, ρ is the density field (also evaluated at z velocity points), and $\mathbf{g}$ is the gravitational acceleration vector.

In order for this method to work correctly it is necessary to set the density of the air layer to some reference value similar to that of the lithosphere in order to prevent the effect of an air layer being counted twice. This approach results in minor errors when considering lateral density variations, discussed further in Section 4.4.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

3.4 Prescribed traction boundary conditions

The variational Stokes discretisation also enables the implementation of prescribed traction boundary conditions at the free surface. Such conditions can occur when there is an external loading on the surface which is not explicitly modelled by the Stokes solver. Prescribed traction boundary conditions are treated by considering the energy per unit time required to "push" against the surface as a result of the Cauchy stress tensor of the prescribed traction $\boldsymbol{\sigma}_{BC}$, which can be split into pressure and deviatoric components $\boldsymbol{\sigma}_{BC} = p_{BC}\mathbf{I} - \boldsymbol{\tau}_{BC}$.

$$\iint_{\partial\Omega_A} \mathbf{v} \cdot (p_{BC}\mathbf{I} - \boldsymbol{\tau}_{BC})\mathbf{n}\, dA, \tag{22}$$

which has units of energy over time, i.e. power, and can be conceptually interpreted as the power required to push against the surface with a velocity v against the surface stress field $\boldsymbol{\sigma}_{BC} = \boldsymbol{\tau}_{BC} - p_{BC}\mathbf{I}$.

By applying the divergence theorem, it is possible to obtain this expression in the form of a volume integral over the air part of the domain Ω_A:

$$\iiint_{\Omega_A} p_{BC}\nabla\cdot\mathbf{v} - \boldsymbol{\tau}_{BC} : \left(\frac{\nabla\mathbf{v} + (\nabla\mathbf{v})^T}{2}\right) + \mathbf{v}\cdot(\nabla p_{BC} - \nabla\cdot\boldsymbol{\tau}_{BC})\, dV \tag{23}$$

As this is the air part of the domain, it is necessary to introduce the air fraction ξ, which is related to the volume fraction for each control volume by $\xi = 1 - \phi$. The discretised form of the above with the appropriate air fractions is given by:

$$-p_{BC}^T \xi^p \mathbf{G}^T\mathbf{v} - \boldsymbol{\tau}_{BC}^T \xi^\tau \mathbf{D}\mathbf{v} + \mathbf{v}^T\xi^v(\mathbf{G}p_{BC} + \mathbf{D}^T\boldsymbol{\tau}_{BC}) \tag{24}$$

This results in a right hand side which can be simply added to the existing right hand side of the Stokes equations:

$$\begin{bmatrix} \mathbf{G}\xi^p p_{BC} - \xi^v\mathbf{G}p_{BC} + \mathbf{D}^T\xi^\tau\boldsymbol{\tau}_{BC} - \xi^v\mathbf{D}^T\boldsymbol{\tau}_{BC} \\ 0 \end{bmatrix}. \tag{25}$$

It should be noted that this right hand side will only affect the Stokes equations at cells where the appropriate volume fractions and air fractions are simultaneously nonzero, i.e. only within interfacial cells.

3.5 Applicability to non-Cartesian geometries

The discrete operators $\mathbf{G}$ and $\mathbf{D}$ are generalised versions of the gradient and divergence operators, which also exist for spherical geometries. As such, it is possible to implement the method also in spherical geometry. For example, the discrete gradient operator $\mathbf{G}$ is a discretised form of the gradient operator $\nabla f(x, y, z)$, which in 3D Cartesian coordinates is given by:

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

$$\nabla f = \frac{\partial f}{\partial x}\hat{\mathbf{i}} + \frac{\partial f}{\partial y}\hat{\mathbf{j}} + \frac{\partial f}{\partial z}\hat{\mathbf{k}}, \tag{26}$$

where $\hat{\mathbf{i}}, \hat{\mathbf{j}}, \hat{\mathbf{k}}$ are the unit vectors in the x, y, and z directions, respectively. In spherical coordinates for a scalar field $f(r, \theta, \phi)$, the gradient operator is likewise given by:

$$\nabla f = \frac{\partial f}{\partial r}\hat{\mathbf{r}} + \frac{1}{r}\frac{\partial f}{\partial \theta}\hat{\boldsymbol{\theta}} + \frac{1}{r\sin\theta}\frac{\partial f}{\partial \phi}\hat{\boldsymbol{\phi}}. \tag{27}$$

where $\hat{\mathbf{r}}, \hat{\boldsymbol{\theta}}, \hat{\boldsymbol{\phi}}$ are the unit vectors in the r, θ, and ϕ directions.

Discretised forms of these operators already exist in the formulation of the Stokes equations used in StagYY, hence no additional work is required to convert the variational Stokes method to spherical coordinates.

4 Implementation in StagYY

The implementation of the variational Stokes method involves several steps: determination of volume fractions, the treatment of boundary conditions, the treatment of near-surface densities and viscosities, and the application of the method to the multigrid solver.

4.1 Obtaining volume fractions

Volume fractions ϕ are a prerequisite for implementing the variational Stokes method. There are several possible methods for obtaining volume fractions in StagYY.

4.1.1 Tracer ratio averaging

The simplest method of obtaining volume fractions is by calculating cell averages of Lagrangian tracers, which are distributed everywhere in the domain and track material type, "rock" or "air". Specifically, the volume fraction of rock is given by the ratio: number of rock tracers divided by total number of tracers (rock and air), similar to in Tackley and King (2003). A cell containing only air tracers will have a volume fraction $\phi = 0$, while a cell containing only rock will have a volume fraction $\phi = 1$. By determining the fraction of rock and air tracers for each of the required volume fraction control volumes, an estimate of volume fractions at all required locations can be determined in a straightforward manner.

While conceptually simple and trivial to implement, this method suffers from the downside of being dependent on tracer density, here defined as number of tracers per cell. Low tracer densities will provide poor resolution for the method to work with, reducing the degree of accuracy. Additionally, due to the strong dependence of surface tracking on tracer density, the surface topography generated by using this method also required a high number of tracers per cell in order to provide accurate results. A demonstration of the tracer density dependence of the tracer averaging method can be seen in Figure 15.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

4.2 Air tracers

With the variational Stokes method in conjunction with an appropriate surface tracking method, the existence of air tracers becomes unnecessary. However, for several reasons it may be desirable to keep the air tracers. These reasons include:

- In StagYY, the tracer density distribution is corrected using the solution of a Poisson equation to redistribute tracers according to tracer density (Tackley, 2025). Removing air tracers would prevent this method from functioning properly.
- There is minimal theoretical performance gain from removing the air tracers when running models in parallel, as tracer advection time is constrained by the subdomain with the largest number of tracers.
- Exchanging of tracers between the air and rock layers as a solution for the entrainment problem becomes impossible when there are no air tracers.
- When using cell averaging to determine volume fractions, air tracers are necessary to compute the correct volume fractions in surface cells.
- Many other features of StagYY rely on the existence of air tracers for backwards compatibility reasons.

This leaves the question of how to treat air tracers in the model. The first option is to do nothing with air tracers. As air velocities are zero, this means that the air tracers remains stationary unless near the surface. However, errors can occur when the surface rises and falls, causing bunching of tracers. This bunching is corrected for through the redistribution of tracers, however this redistribution subsequently causes errors from tracers being moved for reasons not related to the Stokes solver.

For this reason, the second and best option is to copy air velocities into the air layer, ideally using the continuity preserving scheme. This ensures that air tracer density remains consistent due to continuity being satisfied in the air layer, and does not affect the underlying solution through unphysical tracer redistribution.

4.3 Boundary conditions

The variational Stokes method is compatible with the implementation of common boundary conditions, including free-slip or no-slip boundary conditions. As a result, no special modifications to boundary conditions are required.

4.4 Near-surface densities and viscosities

In order to capture near-surface densities and viscosities, a reference value must be chosen for density and viscosity near the surface, to which the appropriate volume fraction is multiplied in order to obtain a fractional value. This approach works for situations in which density and viscosity are constant, as the reference values are known. However, in situations where there are lateral hetereogeneities in near-surface density and viscosity, it can result in minor errors as the reference values may not correspond to the actual density near-surface values at all points.

A potential solution to this problem is to copy reference values from cells immediately below the surface, i.e. the closest cell that does not contain air. In testing, while this approach appeared to provide more accurate results for the density field,

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

testing with the 2D subduction benchmark presented by Schmeling et al. (2008) resulted in differing subduction behaviour as a consequence of subtle variations in the viscosity field near the air-mantle-slab triple junction. In some cases due to the copying of these viscosities, slab break-off was made impossible as it became impossible to create a mantle wedge. As a consequence of this negative result, the results shown consider the use of fixed reference densities and viscosities. Finding the correct treatment for near-surface viscosities is a key future research direction.

4.5 Multigrid solver

A key feature of the variational Stokes method is its ability to be applied to multigrid solvers. The application of volume fraction based methods to multigrid solvers has been demonstrated in the past. For example, Rauwoens et al. (2015) demonstrated the use of a volume fraction based on the solver developed by (Botto, 2013) for solving Poisson equations with a multigrid solver.

The multigrid implementation can be achieved by restricting the volume fraction field to a coarse grid through simple averaging of fine grid volume fractions (Figure 6).

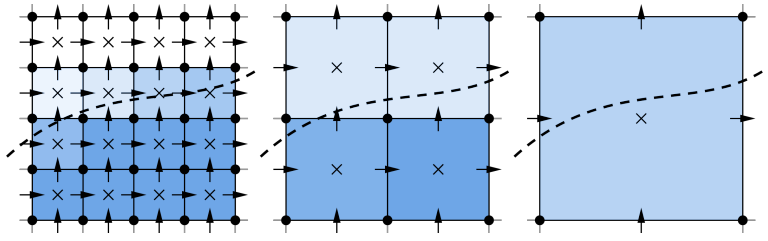

Figure 6. Restriction of the volume fraction function can be achieved by combining neighbouring volume fractions on the fine grid to form the coarse grid.

This coarse grid restriction can easily be achieved in 2D and 3D by averaging 4 or 8 fine grid cells respectively. The extension to 3D is particularly important, as 3D models require the use of iterative multigrid solvers, in contrast to 2D models which also have the option of using direct solvers.

5 Results

Results are demonstrated in both 2D and 3D for the three free surface modelling methods: the sticky air method, staircase method, and the newly implemented variational method. These benchmarks aim to demonstrate the accuracy and efficiency of the new variational method in comparison to the existing free surface modelling methods.

Unless stated otherwise, the surface is tracked for all methods using Lagrangian surface tracking with a bilinear interpolation step, similar to that described in Duretz et al. (2016). This method was found to be able to accurately track free surfaces with minimal additional computational cost.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

5.1 2D benchmark summary

A summary of the 2D benchmarks run can be found in Table 1.

Table 1. Summary of 2D benchmarks for free surface discretisation

Benchmark	Geometry	Resolution	Description
2D inclusion	Cartesian	Various	Benchmark comparison using a test code with the analytical solution of Stokes flow around a viscous inclusion, originally presented by Duretz et al. (2016).
2D relaxation	Cartesian	512×128	Isostatic relaxation of an initially non-flat surface based on Case 1 presented by Crameri et al. (2012a).
2D plume	Cartesian	1024×256	Dynamic topography resulting from a mantle plume based on Case 2 presented by Crameri et al. (2012a).
2D subduction	Cartesian	512×128	Simple subduction zone benchmark setup based on Schmeling et al. (2008).
2D constant viscosity	Spherical annulus	Various	Constant viscosity models with $\mathrm{Ra} = 10^7$ used for investigating performance scaling with resolution.
2D imposed pressure boundary conditions	Cartesian	512×64	Deformation of the surface in response to a pressure boundary condition imposed at the surface.

A summary of the benchmarks run in 2D.

5.2 2D inclusion benchmark

Analytical solutions to simple problems involving the Stokes equation exist for benchmarking our methods. An analytical solution was employed by Duretz et al. (2016) to benchmark the staircase method and examine its performance when scaling with resolution.

The test involves Stokes flow around a viscous inclusion in a 2D Cartesian domain, for which there exists an analytical solution (Schmid and Podladchikov, 2003). In the case of a vanishing inclusion viscosity, the boundary conditions at the interface are identical to free surface boundary conditions. Dirichlet boundary conditions are imposed at the other edges of the model, which correspond to the analytical solution. This test was implemented not in StagYY but in a test code written in Julia, based on the original MATLAB version presented in Duretz et al. (2016).

An improved version of the sticky-air method was implemented, differing from the results presented in (Duretz et al., 2016). The results of that paper considered viscosity as a step function, where a cell whose midpoint was below the surface had the rock viscosity, and the cells whose midpoint was above had the air viscosity. This resulted in an unrealistic representation of viscosity, especially compared to how viscosity is typically computed in numerical geodynamic models. In these results,

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

viscosity near the surface is multiplied with the appropriate volume fraction in order to obtain a smooth representation of viscosity near the surface, providing more accurate sticky air results than the original paper. This approach also highlights the limitations of the staircase discretisation arising from a discrete staircase representation of the surface.

The three surface modelling methods are compared against the analytical solution in terms of pressure and velocity errors. Figures 7, 8, and 9 show the absolute errors of the sticky air method (with $\eta_{\text{rock}}/\eta_{\text{air}} = 10^2$), the staircase method, and the variational method with respect to the analytical solution when run with a model resolution of 100×100. Only cells entirely within the fluid domain are compared, as cells within the air layer do not correspond to physical values for this case.

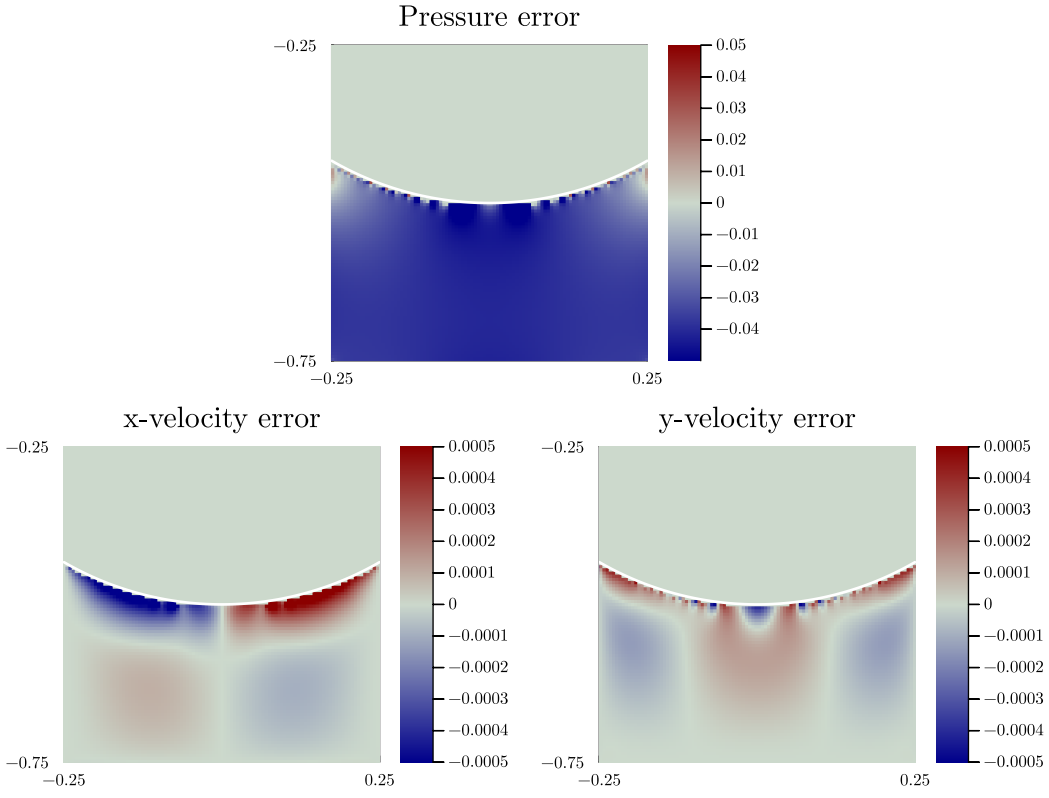

Figure 7. The sticky air method with $\eta_{\text{rock}}/\eta_{\text{air}} = 10^2$ produces significant errors in the pressure and velocity fields when compared to the analytical solution.

The results qualitatively show that the variational method reduces errors in solution in comparison to both the sticky air method and the staircase method. The difference is most apparent when considering the pressure error. The sticky air method, as a result of its insufficient viscosity contrast, produces a large error in pressure throughout the domain. While the staircase method is able to reduce the error far from the surface, this comes at the cost of high errors close to the surface arising from the discrete nature of the staircase approach. The variational method appears to be able to overcome many of these near- and far-field errors, and provide a visibly closer fit to the analytical solution overall. A similar situation exists for the velocity errors.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

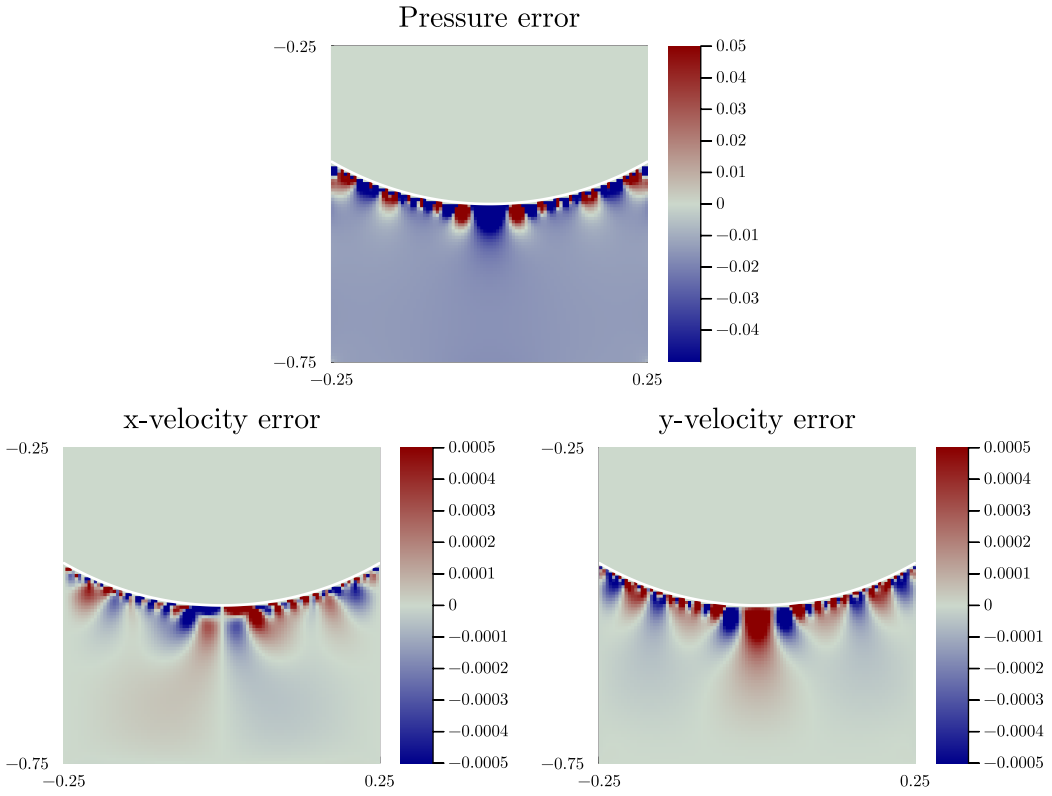

Figure 8. The staircase method has considerably reduced pressure and velocity errors due to viscosity contrast independence. However, significant errors remain close to the surface as a result of the discrete staircase nature of the surface.

In order to quantitatively analyse the performance of each method, models were run at a range of resolutions ranging from 10×10 to 1200×1200 cells. The errors in pressure and velocities using both the L_1 and L_∞ can then be compared and plotted as a function of grid spacing to investigate the scaling performance of each method. As the sticky air method is viscosity dependent, three viscosity contrasts were compared, namely $\eta_{\text{rock}}/\eta_{\text{air}} = 10^2$, 5×10^3, and 10^3.

Figures 10 and 11 consider the pressure errors in the L_1 and L_∞ norm respectively over the i cells in the model. The definition of these errors is given by:

$$||P - P_a||_{L_1} = \frac{1}{V} \sum_{i \in \Omega} P_i - P_{i,a}, \tag{28}$$

$$||P - P_a||_{L_\infty} = \max_{i \in \Omega} (P_i - P_{i,a}), \tag{29}$$

where Ω denotes the fluid domain, and V is the volume of the fluid domain, and P_a is the analytical solution.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

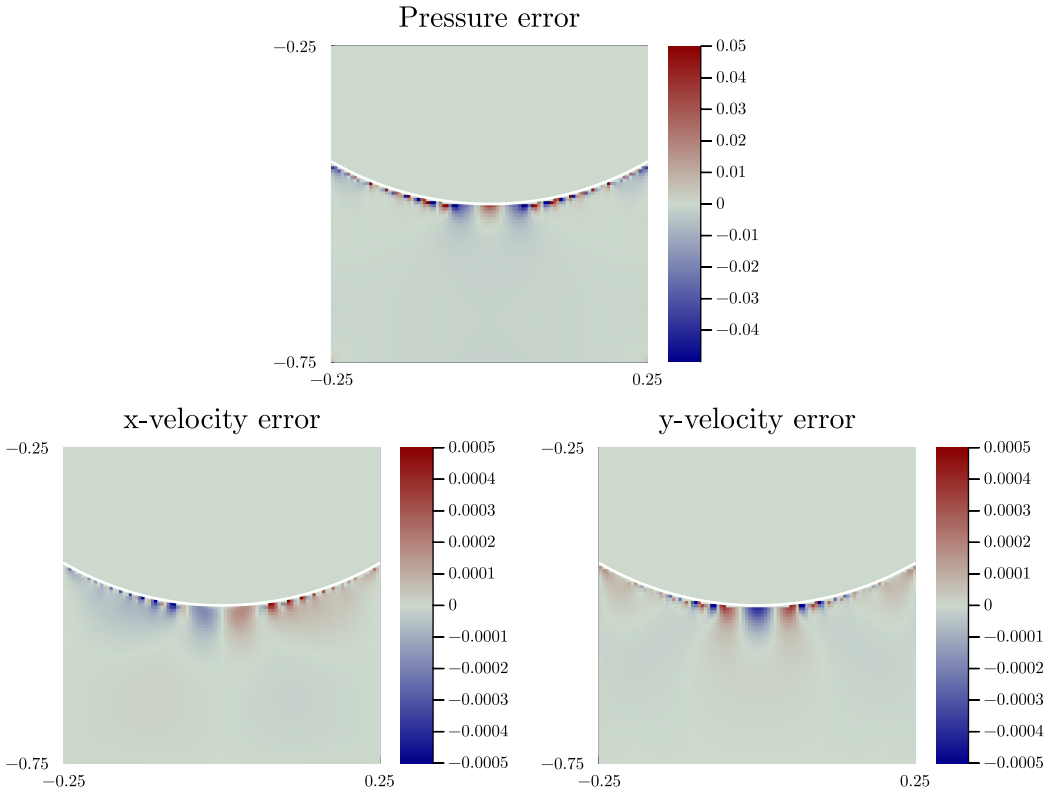

Figure 9. The variational method significantly reduces errors in comparison to the sticky air and staircase method, qualitatively performing the best out of the methods compared for matching the analytical solution.

These pressure error tests reveal that the pressure errors with the variational method have first order scaling with resolution in the L_1 norm, similarly to the staircase method. The L_∞ pressure error does not scale with grid spacing, a feature common to all methods.

In Figures 13 and 12, the errors in velocity (in each dimension) are compared in the L_1 and L_∞ norm respectively, with similar definitions to the pressure errors.

The results from this analytical test reveal that the variational Stokes method is overall more accurate than the staircase method, and does not suffer from the viscosity contrast dependence of the sticky air method which becomes apparent at higher resolution.

5.3 2D relaxation benchmark

The 2D surface relaxation benchmark (Figure 14) is based on Case 1 as presented by Crameri et al. (2012a). It considers the time dependent relaxation of an initially non-flat surface with initial topography $z_{\text{init}}(x)$ given by the function:

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

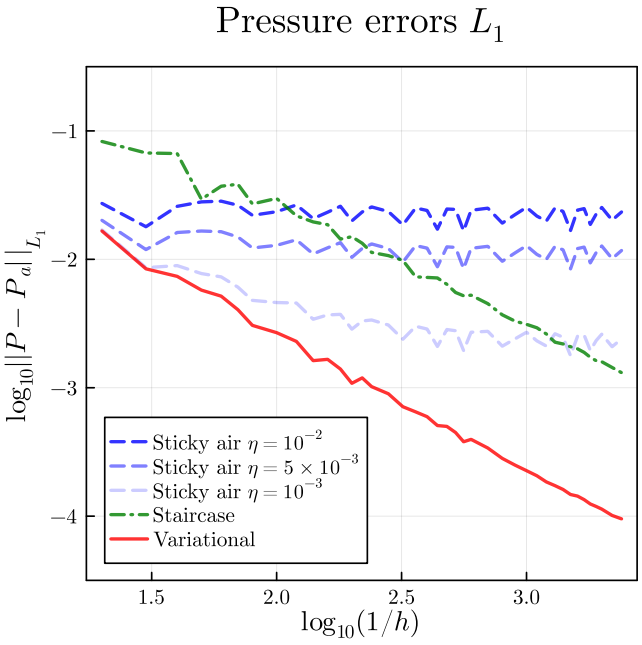

Figure 10. The L_1 error in pressure P_{err} follows a first order scaling with increasing grid resolution $1/h$. The sticky air method quickly fails depending on viscosity, and the staircase method also achieves first order scaling but is overall less accurate.

$$z_{\text{init}}(x) = 7\cos\left(\frac{2\pi x}{2800\text{ km}}\right)\text{km}. \tag{30}$$

This problem is simple enough to have an analytical solution for surface position as a function of time as derived using a three-layer model (Ramberg, 1981). Starting from an initial maximum of 7 km, the maximum topography as a function of time $z_{\text{max}}(t)$ is given by:

$$z_{\text{max}}(t) = 7\exp\left(\frac{t}{14.825\text{ ky}}\right)\text{km}. \tag{31}$$

The model domain is a 2800×800 km Cartesian box consisting of three compositional layers: a 600 km thick mantle layer, a more viscous 100 km thick lithosphere, and a 100 km thick air layer. The resolution is 512×128, with 100 tracers per cell for advecting material properties. When using the sticky air method, a viscosity of 10^{18} Pa s was used. Symmetric boundary conditions are imposed at the horizontal edges of the model, while a no-slip boundary condition is imposed at the bottom boundary, and a free slip condition is imposed at the top. Acceleration due to gravity g is set to 10 m s^{-2}. For consistency with the benchmarks run by Crameri et al. (2012a), models were run until $t = 100$ ky to show the system reaching equilibrium. Timesteps were limited to 500 years in order to give high temporal resolution for plots of model evolution.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

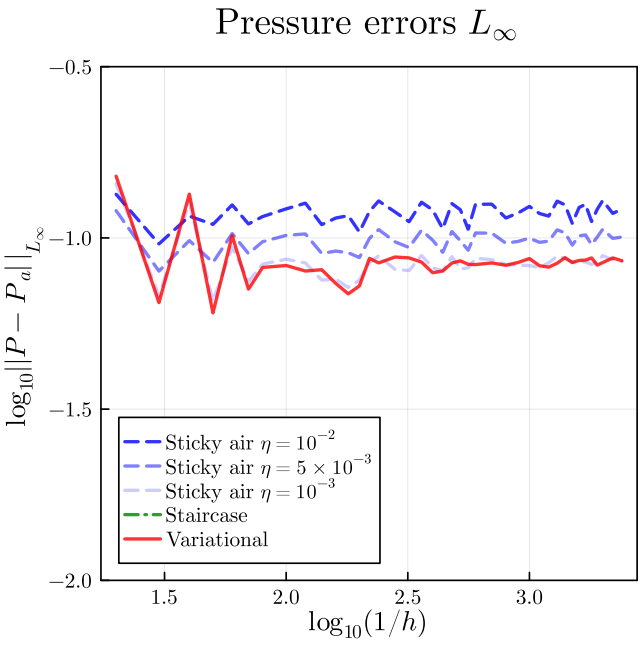

Figure 11. The L_∞ error in pressure does not scale with resolution, with the variational method obtaining similar errors to the staircase method and the sticky air method with high viscosity over all resolutions.

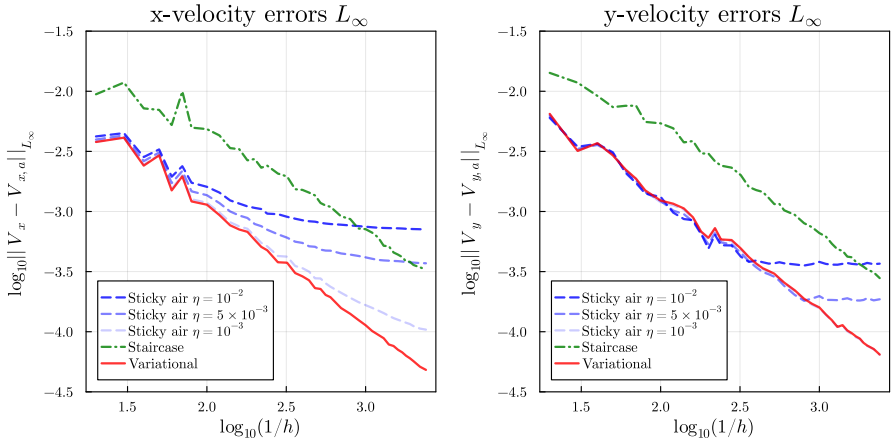

Figure 12. The L_1 error in velocity for the variational method achieves first order scaling with grid spacing, as does the staircase method. The scaling of the sticky air method is viscosity contrast dependent. The variational method is overall more accurate than the staircase method.

Two tests were performed using this model setup. In the first, discretisation methods are compared with the analytical solution to confirm that each method obtains the correct solution. In the second, three methods of acquiring volume fractions when using the variational Stokes method are compared.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

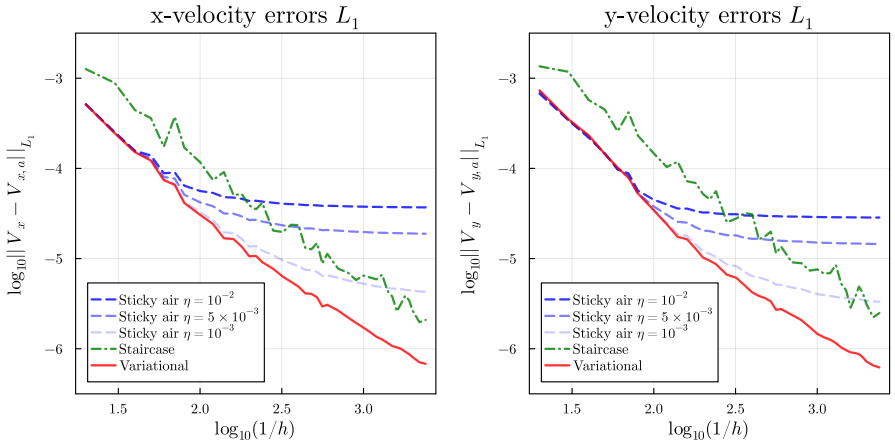

Figure 13. The L_∞ error in velocity for the variational method also achieves first order scaling with grid spacing, as does the staircase method. The scaling of the sticky air method is viscosity contrast dependent. The variational method is overall more accurate than the staircase method.

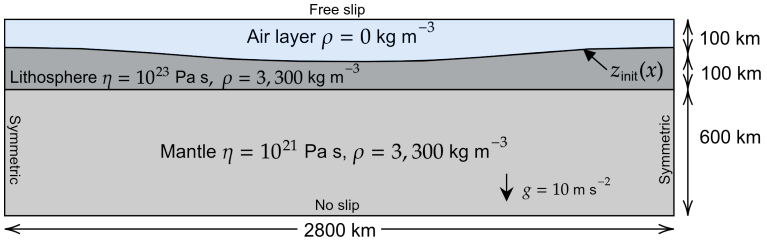

Figure 14. 2D surface relaxation benchmark based on Case 1 presented by Crameri et al. (2012a). An initially non-flat surface in a three-layer model relaxes until equilibrium is reached after approximately 100 ky.

5.3.1 Comparison of discretisation method

In Figure 15, the three discretisation methods are compared against the analytical solution. The result of this benchmark shows that all methods are able to very closely track the analytical solution. The staircase and variational methods in particular are close to the analytical solution, perhaps as a result of minor surface stresses created by the interface to the sticky air layer preventing the sticky air method from relaxing as quickly.

5.3.2 Comparison of volume fraction acquisition methods

A second test was performed to assess the performance of different methods of determining volume fractions. Several different methods are considered: tracer ratio averaging with 10 and 100 tracers per cell (TPC), Lagrangian surface tracking, and the volume of fluid method.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

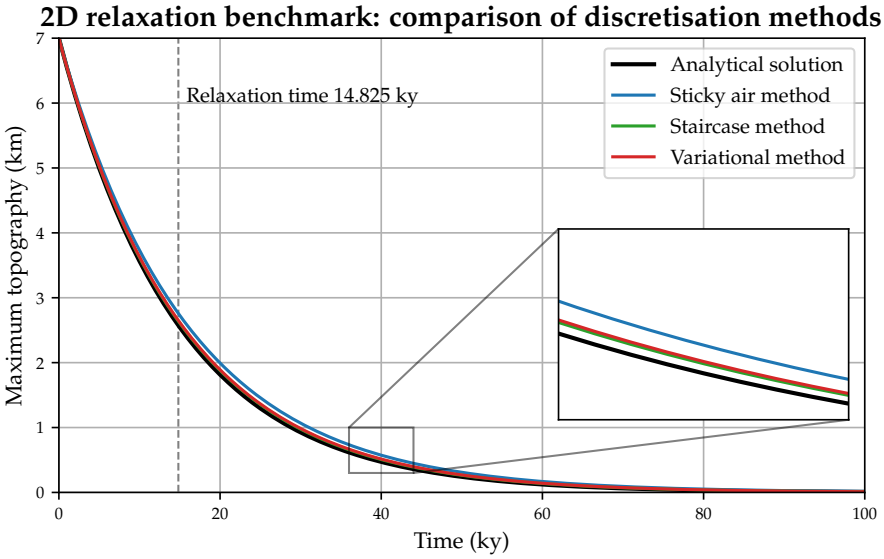

Figure 15. 2D relaxation benchmark comparing performance of different discretisation methods with an analytical solution. All methods tested are able to closely track the analytical solution.

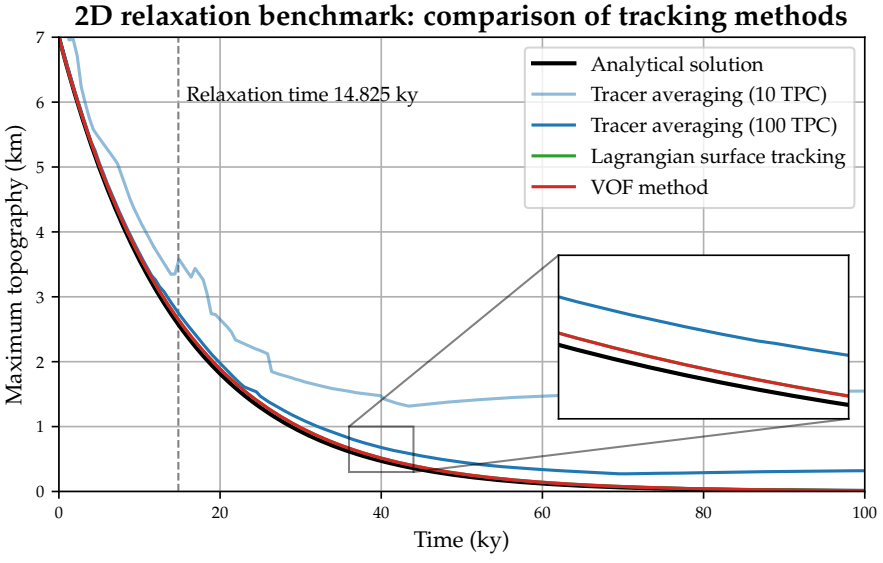

Figure 16. A comparison of tracking methods in conjunction with the variational method. Tracking volume fractions using cell averaging provides poor results, and for this reason tracking using another surface tracking method is recommended.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

The results of this test demonstrates that the simple tracer averaging method is dependent on tracer density, as demonstrated by the 10 tracer per cell solution tracking the analytical solution poorly in comparison to other methods. Increasing to 100 TPC improves the tracking performance, however in this case noise associated with the tracer distribution prevents the surface from fully relaxing, as can be seen by the residual topography at 100 ky. All other methods, which rely on a direct tracking of the surface location, provide very similar, high quality results that closely track the analytical solution

5.4 2D plume benchmark

The 2D plume benchmark (Figure 17) is based on Case 2 presented in Crameri et al. (2012a). It considers a buoyant plume rising through the mantle over a period of 20 My, and the dynamic topography resulting from this plume's impingement on the lithosphere.

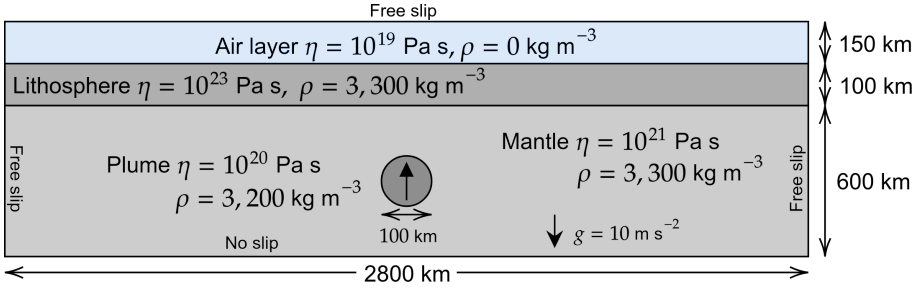

Figure 17. 2D plume benchmark based on Case 2 in Crameri et al. (2012a). A buoyant plume rises through the mantle and impinges upon the lithosphere, producing dynamic topography over a time period of 20 My.

The domain is a 2800×850 km Cartesian box domain run at a resolution of 1024×256, consisting of three compositional layers: a 600 km thick mantle, a 100 km thick lithosphere, and a 150 km thick air layer (Figure 17). When using a sticky air layer, a viscosity of 10^{19} Pa s is used. The buoyant mantle plume, with radius 100 km, has a density of 3,200 kg m^{-3}; 100 kg m^{-3} less than that of the surrounding mantle and lithosphere. 100 tracers per cell are used for advecting material properties. Acceleration due to gravity g is set to 10 m s^{-2}. Free slip boundary conditions are imposed at the horizontal edges of the model, while a no-slip boundary condition is imposed at the bottom boundary, and a free slip condition is imposed at the top. The model is run for 20 My, which allows for the plume to rise through the mantle and impinge upon the lithosphere causing deformation of the surface.

Analytical solutions for this benchmark are not available. However, results from multiple numerical geodynamic models, including finite element models with true deformable free surfaces, were run as part of the study of Crameri et al. (2012a) and consistently show a maximum surface deformation 800-850 km after 20 My.

The results of the 2D topography benchmark demonstrate that there are only very minor deviations between results obtained using the sticky air method, the staircase method, and the new variational method. All results are also in agreement with the range of results presented in Crameri et al. (2012a). This result validates the new variational method's ability to capture 2D dynamic topography.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

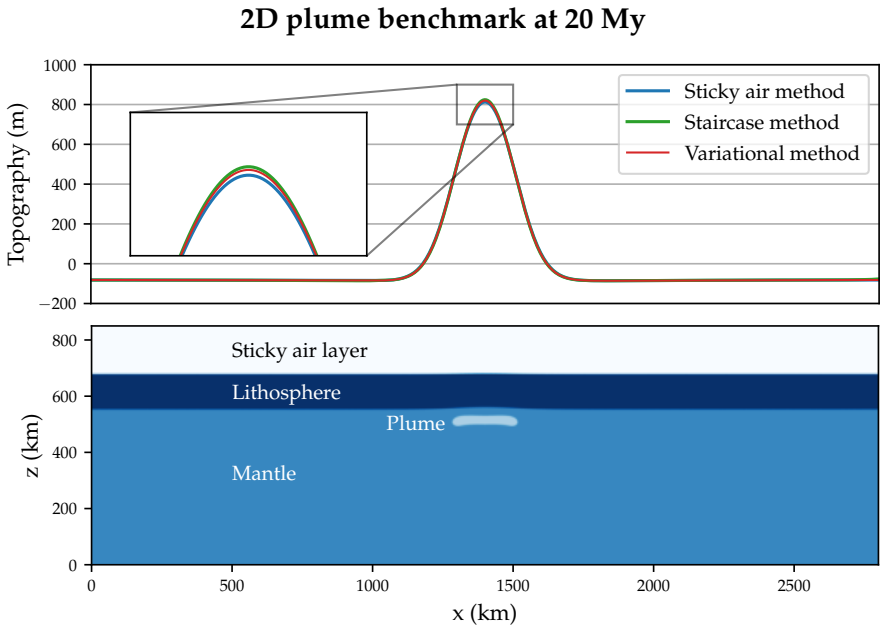

Figure 18. All three discretisation methods provide very similar results, with deviations in maximum topography on the order of metres. These results are also in agreement with the benchmark results presented by Crameri et al. (2012b).

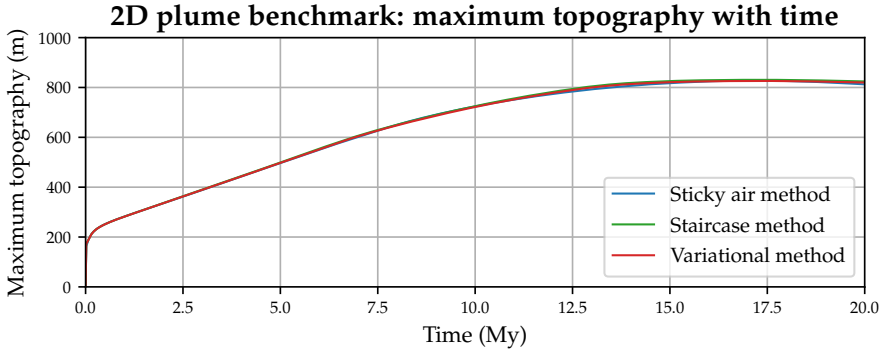

Figure 19. The development of 2D plume topography over time is similarly very similar between all methods considered, with only very minor deviation between them.

5.5 2D subduction benchmark

In order to test the ability of the implemented methods to track surfaces near convergent margins, particularly subduction zones, the simple isothermal viscous subduction benchmark proposed by Schmeling et al. (2008) was used with the model setup given in Figure 20.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

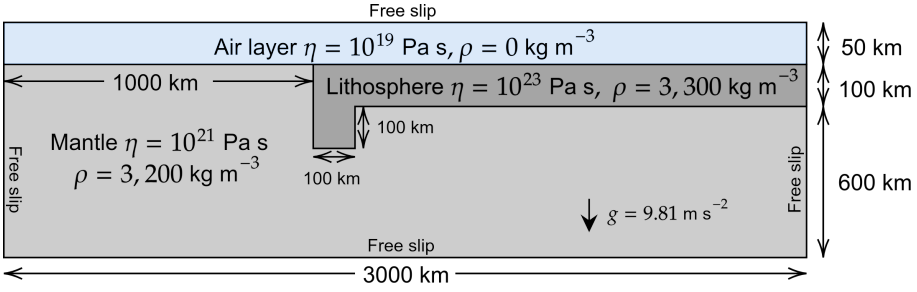

Figure 20. The 2D subduction benchmark is based on Case 1 presented in Schmeling et al. (2008), and considers subduction initiation in this simplified viscous setup. Diagram not to scale.

The model setup consists of a gravitationally unstable lithospheric slab sinking and initiating a new subduction zone. The models are run on a 3000×750 km Cartesian box domain at a resolution of 512×64. The lithosphere has a thickness of 100 km, a density 100 kg m^{-3} greater than the surrounding mantle, and a viscosity 100 times greater than the surrounding mantle. At the top of the model is a 50 km thick air layer. When using the sticky air method, a viscosity of 10^{19} Pa s, 100 times lower than the viscosity of the mantle and 10,000 times lower than the viscosity of the slab, is used. Free slip boundary conditions are imposed on model boundaries. Acceleration due to gravity g is set to 9.81 m s^{-2}. The models were run for a total of 100 My.

The results for this benchmark presented in Schmeling et al. (2008) are notable for the large differences in results produced as a result of relatively minor changes in the viscosity field, such as the use of different viscosity interpolation schemes. This fundamental instability is a result of minute differences in viscosity near the mantle-lithosphere-air triple junction causing radically different subduction initiation behaviour, which then causes different behaviour throughout the model run. In the models run for this study, a geometric average is used for viscosity interpolation, however, radically different results can be obtained by using arithmetic or harmonic averaging.

The topography generated by the three methods, plotted at 40 My in Figure 21, reveals that the variational method and staircase method produce extremely similar results, in contrast to the sticky air method, which exhibits slightly faster subduction. Figure 22 demonstrates the subtle differences between the profiles at 40 My for the sticky air method and the variational method.

Figure 23 further demonstrates the differing rates of subduction between the three methods by considering the depth of the tip of the subducting slab over time.

These results are not consistent with results obtained from the sticky air method with a higher viscosity contrast, in which case subduction was shown to be faster. While the two true free surface methods would be expected to show the fastest subduction due a lack of resistance from the air layer, they instead produce slightly slower, delayed subduction. The cause of this behaviour is unknown, but likely related to subtle differences in the treatment of viscosities near air-mantle-lithosphere triple junction affecting the initiation of subduction.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

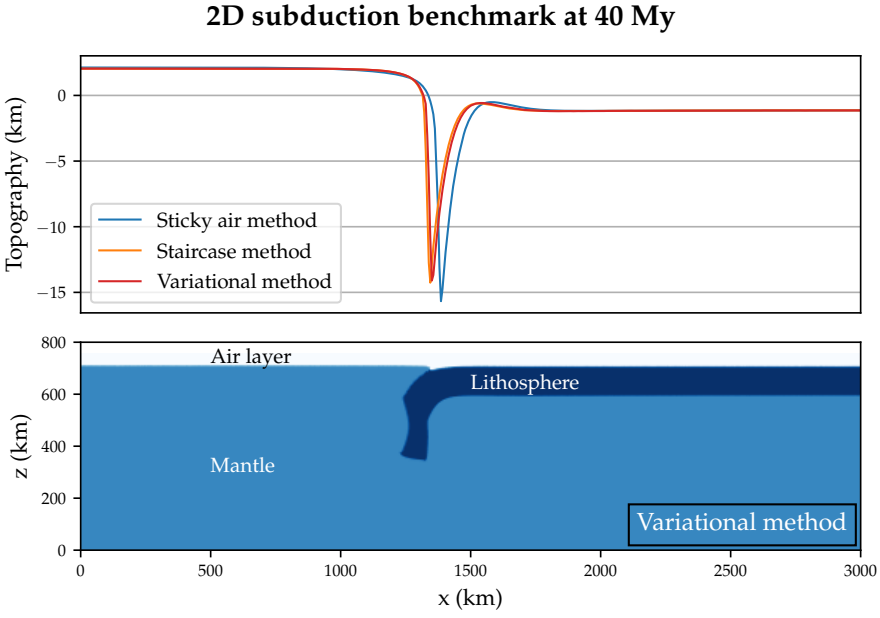

Figure 21. 2D subduction benchmark topography after 40 My for the three discretisation methods. There are slight differences in subduction performance when using a sticky air method as opposed to either the staircase or variational methods, which produce very similar results.

5.6 2D scaling tests

Scaling tests in 2D were performed using a non-dimensional model in spherical annulus geometry (Hernlund and Tackley, 2008). The model setup involves a constant viscosity mantle with Rayleigh number $\mathrm{Ra} = 10^7$, an air layer with a thickness 10% of the mantle's thickness, and, in the case of the sticky air method, a viscosity contrast of $\eta_{\mathrm{rock}}/\eta_{\mathrm{air}} = 10^3$. Each scaling model was run for 5,000 timesteps to ensure steady-state convection was achieved. The computational performance of the Stokes solver—typically the most resource-intensive component of numerical geodynamic models—was assessed across three discretisation methods. Tests were conducted at various resolutions, ranging from 128×16 to 1024×128 to evaluate solver performance at different resolutions. Simulations were run in parallel on 4 cores for lower resolutions and 8 cores for higher resolutions, using an AMD Ryzen 9 3900XT processor.

The performance results, illustrated in Figure 24, show that all methods perform comparably when using a direct solver—the most common approach for 2D models. For these tests the MUMPS solver (Amestoy et al., 2001, 2019) is used via the PETSc toolkit (Balay et al., 2025). This consistency aligns with the fact that the size of the linear system remains identical across discretisation methods. However, when employing an iterative multigrid solver, significant performance improvements are achieved with the staircase and variational methods, the latter offering the most notable gains over the sticky air method.

Interestingly, at the highest resolution of 1024×128, the staircase method exhibited worse performance than the sticky air method. This outcome, verified through repeated simulations, may indicate that this specific resolution generates a linear

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

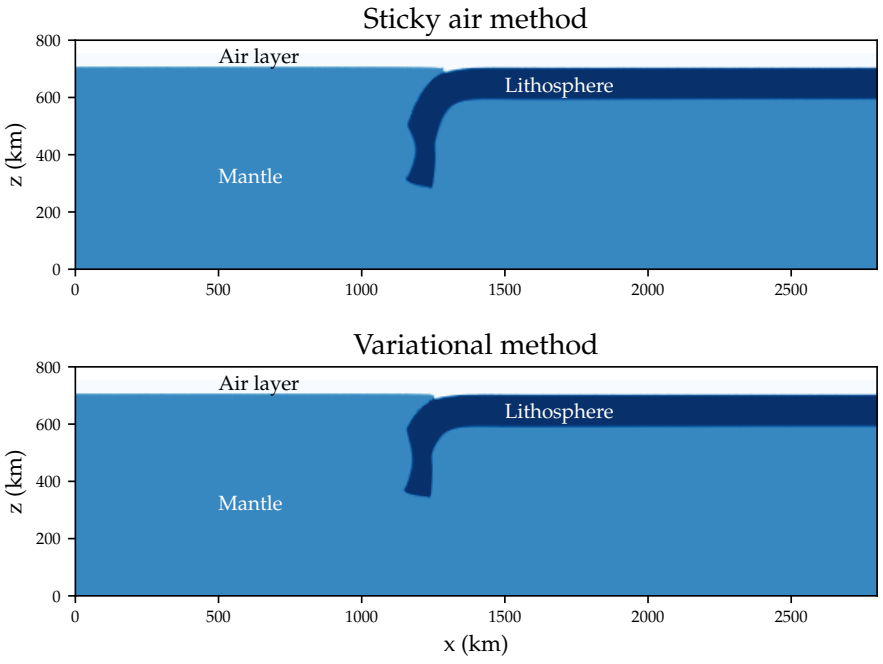

Figure 22. The differences in 2D subduction profiles at 40 My are visible when comparing slab profiles. Compared to the sticky air method, the variational method (as well as the similarly performing staircase method) exhibit slightly delayed subduction initiation, resulting in a slightly less developed slab at this timestep.

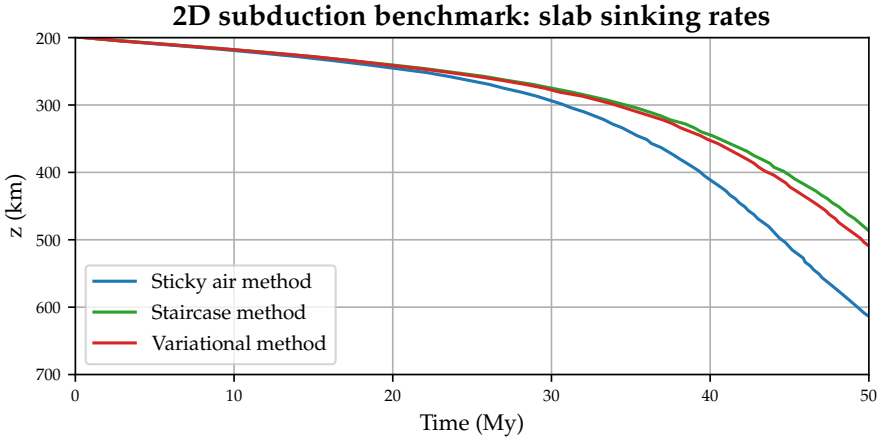

Figure 23. Slab sinking rates can be assessed by plotting the maximum depth of the subducting slab over time. This demonstrates the differences in sinking rates and slightly delayed subduction initiation seen when using the staircase and variational methods.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

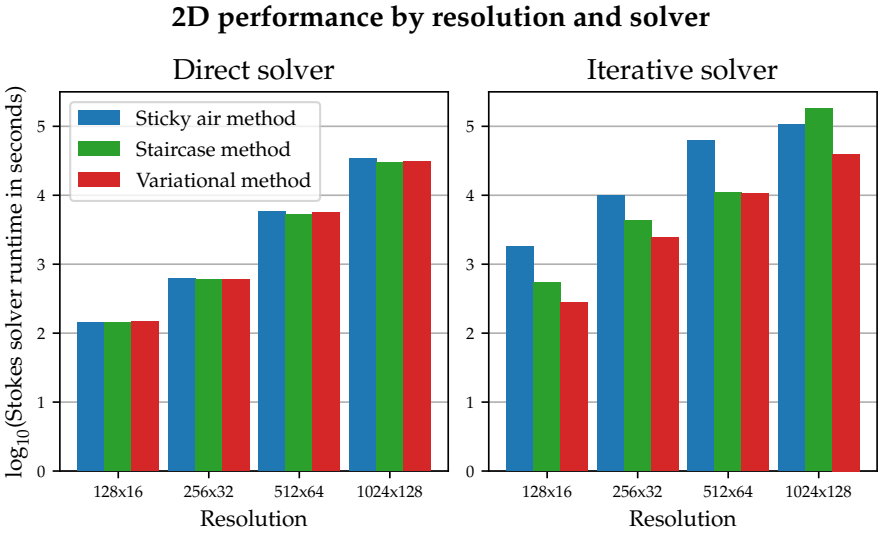

Figure 24. Performance across all resolutions is consistent between discretisation methods when using a direct solver (typical for 2D models). However, with an iterative multigrid solver, significant performance gains are observed when using the variational method.

Table 2. Summary of 3D benchmarks for free surface discretisation

Benchmark	Geometry	Resolution	Description
3D plume	Cartesian	$256 \times 256 \times 128$	Dynamic topography resulting from a rising mantle plume.
3D constant viscosity	Cartesian	Various	Constant viscosity models with $\mathrm{Ra} = 10^7$ used for investigation of scaling with resolution.

A summary of the benchmarks run in 3D.

system that is particularly challenging to solve. Such behaviour can conceivably arise when multigrid cell boundaries are poorly aligned with the surface location at certain resolutions, reducing the effectiveness of coarse-grid iterations and necessitating more fine-grid iterations. This problem does not appear to affect the variational method, which is able to approximate the free surface boundary condition with sub-grid resolution on coarse grids as well as fine grids.

5.7 3D benchmark summary

Similarly to 2D, it is possible to run a set of benchmarks to test the efficiency and accuracy of surface modelling methods in 3D. A summary of the models run in 3D can be found in Table 2

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

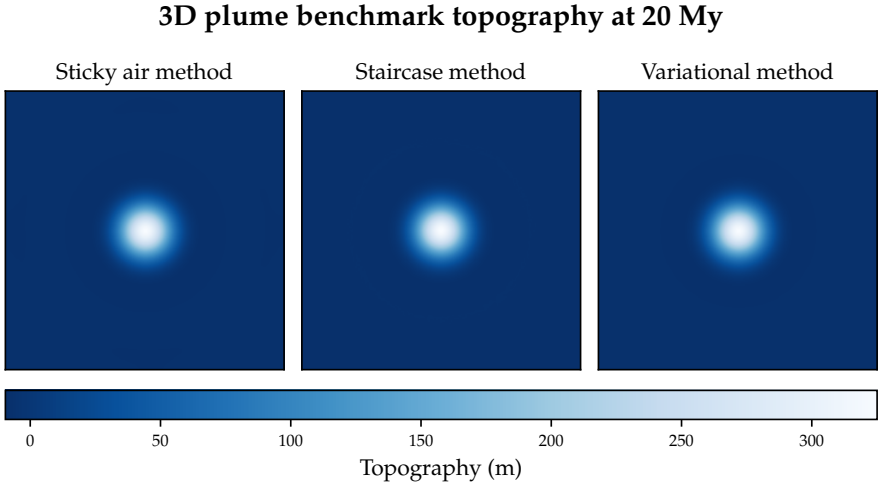

Figure 25. The top down view of 3D plume benchmark topography reveals no significant differences between the three methods.

5.8 3D plume benchmark

The 2D plume benchmark presented in Section 5.4 which was originally presented by Crameri et al. (2012a) can be extended into 3D. In comparison to the 2D case, the width of the domain has been halved in order to reduce the considerable additional computational effort required for the 3D case. The domain then becomes a $1400 \times 1400 \times 850$ km Cartesian box run at a resolution of $256 \times 256 \times 64$, a resolution which was chosen to balance the high computational cost of 3D models with accuracy. The model consists of a 600 km thick mantle, a 100 km thick lithosphere, and a 150 km thick air layer. The geometry of the mantle plume is modified to a sphere with an initial diameter of 100 km. A rigid (no-slip) boundary condition is imposed at the lower boundary, while free slip boundary condition is imposed on all other boundaries. All other physical parameters remain identical between the 2D and 3D models, including the density and viscosity parameters. As this 3D version is a novel benchmark, there are no previous results available for comparison.

The results at 20 My, shown in both top down view and cross section view in Figures 25 and 26, reveal no significant differences between the three methods when using this benchmark, in accordance with the results found for the corresponding 2D results. There is a minor offset in cross sectional topography using the staircase method, however the cause of this is unknown and the magnitude of the discrepancy is very small. Importantly, the variational method provides virtually identical results to the sticky air method for this benchmark.

5.9 3D scaling tests

3D scaling tests are used to assess the performance of the three methods. A constant viscosity Cartesian box setup featuring convection with a Rayleigh number of $\mathrm{Ra} = 10^7$ was run at a variety of resolutions ranging from $32 \times 32 \times 16$ to $256 \times 256 \times 128$. When using the sticky air method, an air layer with viscosity 1000 times less than the mantle and 10% the thickness of the

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

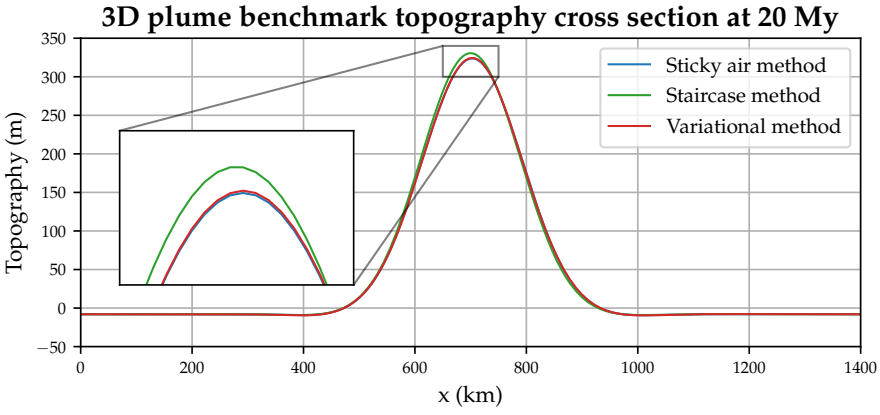

Figure 26. In cross section, all three methods provide virtually identical results with the exception of the staircase method, which appears to feature a minor offset when compared to the other two methods.

mantle was used. Models were run for a total of 5000 timesteps to ensure a steady state convection was reached. As direct solvers are unavailable for 3D cases, only iterative multigrid solvers can be compared, in contrast with the 2D performance benchmarks shown in Section 5.6.

The results of this performance benchmark, which consider the time required to solve the Stokes equations (8 cores, Ryzen 9 3900XT), reveal that the variational method can provide consistently improved Stokes solver performance in comparison to the sticky air method, with the greatest improvement seen at lower resolutions. At lower resolutions, the staircase method also provided superior performance to the sticky air method, with the exception of the highest resolution used, $256 \times 256 \times 128$. The variational method performed better than the staircase method for all except the $128 \times 128 \times 64$ model, where the performance was similar.

These results demonstrate that the variational method offers a more efficient solution to the Stokes equations with a free surface boundary condition than the sticky air method across all tested cases. This efficiency, coupled with consistent surface tracking performance across resolutions, highlights the potential of the variational Stokes method for 3D geodynamic simulations.

5.10 Prescribed traction boundary conditions

One of the key features enabled by the variational Stokes method is the ability to apply tractions directly at the free surface. This capability holds significant potential for applications in global-scale geodynamic models. Here, its implementation is tested under simple imposed pressure boundary conditions in both 2D and 3D configurations.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

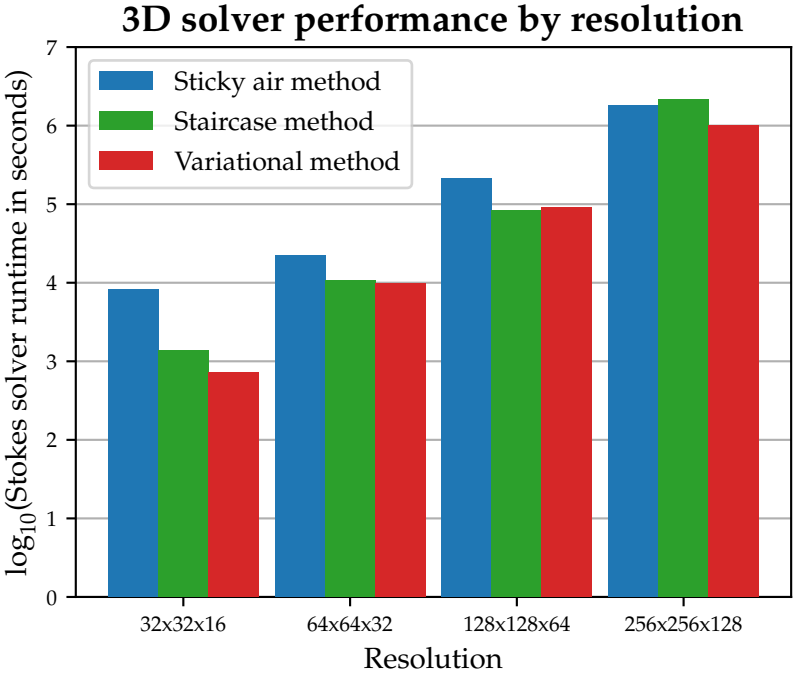

Figure 27. The 3D scaling tests, conducted using an iterative solver on a constant-viscosity Cartesian box domain at various resolutions, demonstrate that the variational method consistently outperforms the sticky air method. The staircase method shows variable performance results.

5.10.1 2D imposed pressure boundary conditions

To evaluate the application of imposed traction boundary conditions in 2D, a simple setup illustrated in Figure 28 is used. The model consists of a 2800×800 km domain run at a resolution of 512×64, consisting of an initially flat 600 km thick mantle, a 100 km thick lithosphere, and a 100 km thick air layer. A sinusoidal pressure boundary condition is applied to the free surface, given by:

The imposed pressure boundary condition tested is a sinusoidal profile given by:

$$P_{BC} = 10^9 \left(1 + \cos(8\pi x / 2800 \text{ km})\right) \text{ Pa}. \tag{32}$$

The analytical solution for surface height can be derived using the relationship for lithostatic pressure, $P = \rho g h$, resulting in:

$$h_{\text{analytical}} = \frac{10^9 \cos(8\pi x / 2800 \text{ km}) \text{ Pa}}{3300 \text{ km } 9.81 \text{ms}^{-2}}. \tag{33}$$

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

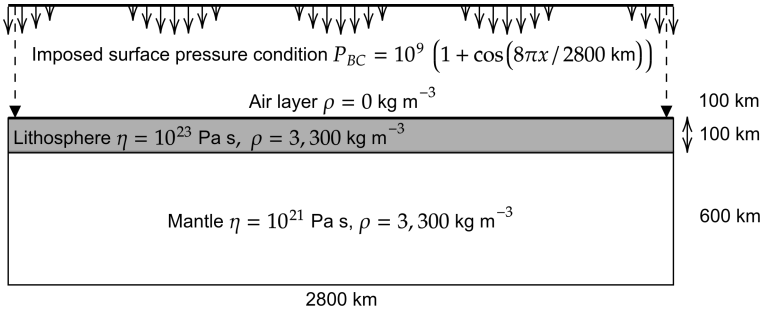

Figure 28. Imposed pressure boundary condition setup in 2D. The model is run at a resolution of 512×64 for 4 My to investigate the effect of the imposed surface pressure on topography.

The model is run for 4 My to allow surface deformation to reach equilibrium. The resulting topography is compared to the analytical solution in Figure 29.

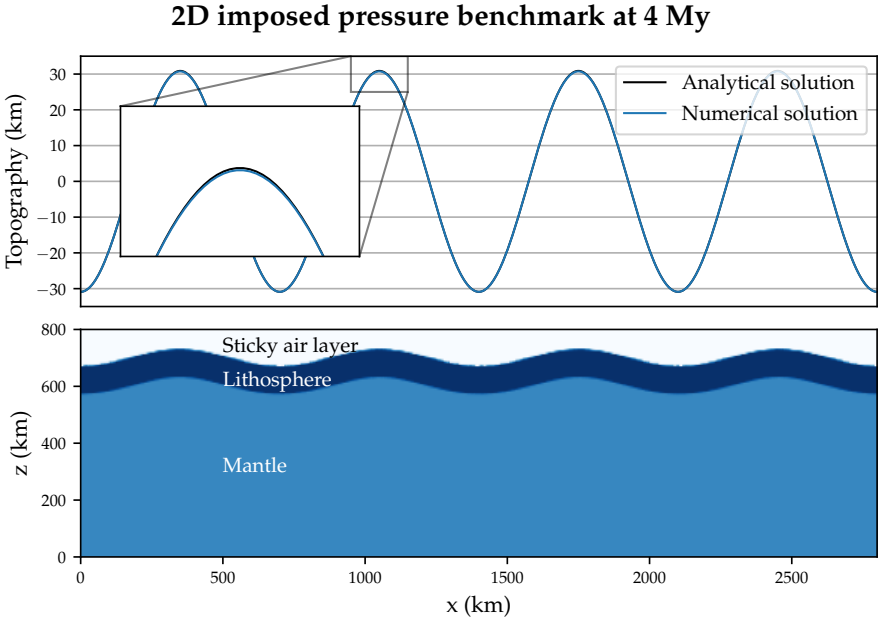

Figure 29. The topography generated by the imposed pressure boundary condition in 2D follows a sinusoidal pattern corresponding to the applied pressure. After 4 My, the numerical solution has almost exactly converged to this analytical solution.

The results show that the surface topography relaxes to the predicted analytical solution, exhibiting sinusoidal deformation due to the imposed pressure boundary condition. The time-dependent relaxation process, influenced by the model's viscosity and setup, is shown in Figure 31.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

5.10.2 3D imposed pressure boundary conditions

The same test can be repeated in 3D to assess the applicability of applied traction boundary conditions to 3D models. In 3D, the following imposed pressure boundary condition was applied:

$$P_{BC} = 10^9 \left(1 + \cos(4\pi x/1400 \text{ km})\right)\left(\cos(4\pi y/1400 \text{ km})\right) \text{ Pa}. \tag{34}$$

From this an analytical solution for maximum topographic height can be derived:

$$h_{\text{analytical}} = \frac{10^9 \cos(4\pi x/1400 \text{ km}) \cos(4\pi y/1400 \text{ km}) \text{ Pa}}{3300 \text{ km } 9.81 \text{ms}^{-2}}. \tag{35}$$

An extension of the 2D model consisting of a $1400 \times 1400 \times 800 km$ Cartesian domain on a grid resolution of $256 \times 256 \times 128$ was tested. The results after 4 My of deformation can be seen in Figure 30.

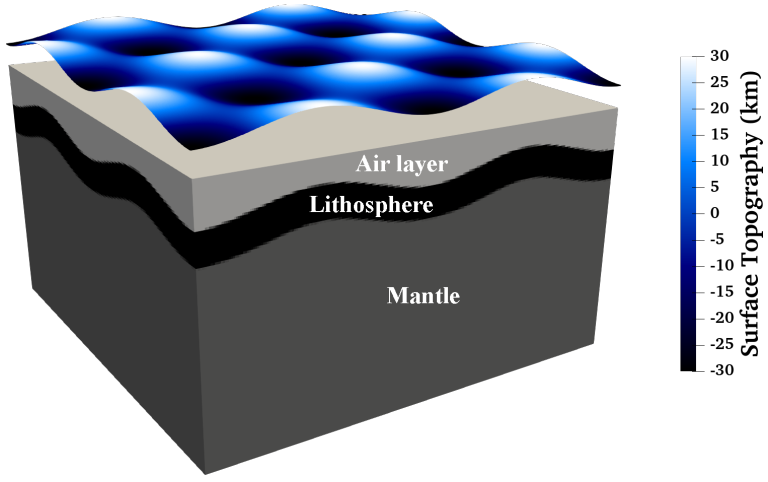

Figure 30. In 3D, similar results to the 2D results can be obtained from an imposed sinusoidal surface condition, shown here at 4 My as run on a $256 \times 256 \times 128$ grid.

As in the 2D case, the deformation evolves over time rather than occurring instantaneously. The relaxation curves, which reflect the impact of model viscosity and geometry, are shown in Figure 31. Despite differences in the rate of convergence between 2D and 3D cases, which are a result of the different imposed pressure boundary conditions, the theoretical maximum topography is the same for both cases.

For very large imposed pressures, larger than those tested here, the efficiency of iterative Stokes solvers may be reduced. In 3D models, more fine grid iterations are sometimes required instead of the coarse multigrid iterations due to poor conditioning of the linear system resulting from the large imposed pressure conditions. However, for realistic pressure magnitudes (e.g.,

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

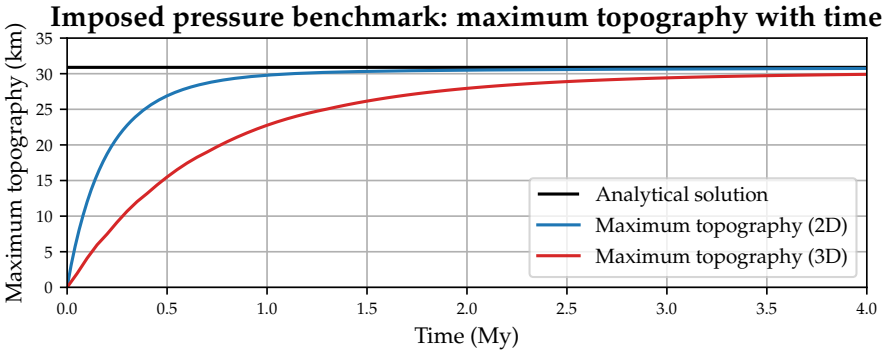

Figure 31. Starting from an initially flat state, the maximum topography converges over time to the theoretical maximum. The rate of convergence is related to the viscosity of the mantle and lithosphere, and the model geometry, hence the difference in convergence rates in 2D and 3D.

ocean or glacial loading), the effect on solver performance is minimal, making this approach suitable for most practical global-scale geodynamic applications. In 2D when using a direct solver, performance is unaffected by the imposition of these boundary conditions.

6 Summary and discussion

6.1 Summary of key results

The key results of this chapter can be summarised as follows:

- The existing sticky air method suffers from limitations in terms of performance, and the requirement to implement an air layer of sufficiently high thickness and low viscosity.
- The staircase method, while bypassing many of the limitations of the sticky air method, does not perfectly replicate the results of high-viscosity contrast sticky air models in all cases and numerical instability can lead to inconsistent performance.
- The implementation of the variational Stokes method is possible in StagYY through modification of the stencil by multiplication with appropriate volume fractions.
- Volume fractions can be generated in a number of ways including tracer ratio averaging, Lagrangian marker chains, or the volume of fluid method. Tracer averaging is not recommended, as it requires high tracer densities in order to provide accurate results.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

- In most scenarios, the variational Stokes method provided virtually identical results to the sticky air method or staircase method, while significantly decreasing the computation cost of solving the Stokes equations when using iterative or multigrid solvers.
- The variational Stokes method is applicable to all 2D and 3D geometries implemented in StagYY, including spherical and yin-yang geometries.
- The variational Stokes method enables the implementation of non-zero boundary conditions at the surface such as the imposition of pressures or tractions with numerous potential geodynamic applications.

6.2 Implementation

6.2.1 Variational Stokes implementation

The variational Stokes method can be implemented with relative ease by modifying the viscosity fields within the Stokes solver. Since this method requires adjustments only to local stencils, the implementation is straightforward. It could also be adapted to other geodynamic codes that employ staggered-grid finite differences, such as I2ELVIS/I3ELVIS (Gerya and Yuen, 2007; Gerya et al., 2015b) or LaMEM (Kaus et al., 2016), without significant difficulty.

6.2.2 Volume fraction implementation

Determining volume fractions poses greater challenges, as the simplest method of determining them, cellwise tracer averaging, was determined to be insufficient for accurate results. Alternative surface tracking methods are required, such as the Lagrangian surface tracking, volume of fluid methods, or the level-set method, to compute volume fractions accurately.

6.3 Applications

6.3.1 More efficient free surface modelling

Free surface modelling is often neglected in favour of free-slip boundary conditions, primarily due to the limitations imposed by the commonly used sticky air method. However, free-slip boundary conditions cannot capture all physical phenomena arising at the surface, meaning that physical accuracy is compromised (Crameri et al., 2012b). The variational Stokes method offers improved computational efficiency compared to the both sticky air method, and the staircase method, reducing the need to rely on free-slip boundary conditions. This development makes accurate free surface boundary modelling a more practical and attractive choice for researchers.

6.3.2 Prescribed traction boundary conditions

Non-zero prescribed traction boundary conditions at the free surface present a valuable opportunity for modelling coupled planetary systems, particularly in scenarios involving glaciers and oceans. As an example, the effects of glacial loading on the

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

lithosphere, exemplified by post-glacial rebound, can now conceivably be integrated into global scale geodynamic models with a free surface. With coupling to a method of tracking sea levels over time, the loading imposed by oceans on the lithosphere can also be considered for the first time. This development opens up many opportunities for future study.

6.4 Limitations and potential improvements

Only one critical limitation of the variational Stokes method was noted. The treatment of near-surface viscosities introduces errors in the viscosity field near the surface when lateral surface viscosity variations are present, either using a representative air viscosity value or the copying method discussed in Section 4.4. Furthermore, it was found that the treatment of near-surface viscosities had a very strong effect on certain dynamics, as demonstrated in the subduction benchmarks of Section 5.5. As a result, further work needs to be done in order to determine the correct way of determining near-surface viscosities. Due to the fact that the subduction benchmarks presented both here and in the original paper (Schmeling et al., 2008) demonstrate very strong differences in subduction dynamics as a result of very subtle changes in viscosity field, it is likely that new, more stable benchmarks will have to be introduced in order to test the validity any future methods.

6.5 Future work

6.5.1 Applicability to incompressible flows

The variational Stokes method as presented was derived for the case of incompressible Stokes flow, while StagYY and many other mantle convection codes are able to solve the more general problem of compressible Stokes flow. To implement a similar method for compressible Stokes flow would likely require a re-derivation of the variational formulation as presented in Equation 11, and subsequent rederivation of the discretised form. The result would likely look similar to the current form, but with additional terms and additional volume fractions to consider compressibility. The main changes in the equations are the addition of density to the continuity equation and the subtraction of net expansion/contraction from the normal stress terms.

However, it should be noted that for Earth-like planets, the compressibility of rocks near the surface is small, with compressibility mainly causing a gradual increase of density with depth (pressure), therefore becoming important only in regions that are unaffected by the variational Stokes method. It is very likely that any future correction of the variational Stokes method for compressible fluids would lead to only very minor differences in solution compared to the incompressible version.

6.5.2 Improved volume fraction multigrid restriction

Currently when using multigrid methods, restriction of volume fraction fields to coarse grids is achieved by simple averaging of appropriate cells. This is appropriate for equally spaced Cartesian grids, but less accurate for variable spaced grids or spherical geometries. Although this does not impact fine-grid volume fractions, it could introduce minor errors during coarse-grid iterations, potentially hindering convergence of multigrid methods. A weighted averaging method, accounting for cell areas/volumes, would provide a more accurate solution. The implementation of such an improved method is conceptually simple, but difficult to implement in StagYY due to the way the code is structured.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

6.5.3 Viscosity interpolation

The subduction benchmarks by Schmeling et al. (2008) highlight the sensitivity of results to viscosity interpolation schemes. While geometric averaging was used in this study, future work should explore alternative interpolation schemes (e.g. arithmetic or harmonic) to identify the most accurate approach for use in conjunction with the variational Stokes method.

6.5.4 Solid-fluid coupling

While not implemented in this study, the variational Stokes method as proposed in Larionov et al. (2017) also considers the case of solid-fluid coupling. This can be implemented similarly to the free surface discretisation, requiring only modification of the Stokes equations through the incorporation of volume fractions. This approach retains compatibility with the variational free surface discretisation.

The implementation of fluid-solid coupling could have potential applications for the modelling of undeformable regions within the mantle. For example, it could be used to approximate a rigid continental structure without the need to impose a large viscosity contrast with the mantle. Modelling continents as rigid blocks in this way could potentially resolve two issues with existing continental modelling using the methodology of Rolf and Tackley (2011). Firstly, by imposing a rigid continent through boundary conditions, it could avoid the entrainment of continental material into the mantle by imposing a solid-fluid interface boundary condition directly at the continental margins. Secondly, by avoiding the large viscosity contrast currently required to model continents without them deforming unrealistically, there may be potential for performance gains similar to the performance gain resulting from eliminating the sticky air layer.

Code availability. The Julia script used to produce the benchmark results and figures used in this paper is archived on Zenodo under the MIT license under https://doi.org/10.5281/zenodo.17956246 (Gray, 2025). No input data or additional scripts are required. The data plotted in Figures 7-13 can be generated by running the script.

Author contributions. **Timothy Gray**: original study design; code and algorithm development; scaling and performance benchmarks; figure creation. **Paul Tackley**: development of StagYY; conceptual input; manuscript review. **Taras Gerya**: conceptual input; manuscript review.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. This work was supported by Swiss National Science Foundation #192296.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

References

Amestoy, P., Duff, I. S., Koster, J., and L'Excellent, J.-Y.: A Fully Asynchronous Multifrontal Solver Using Distributed Dynamic Scheduling, SIAM Journal on Matrix Analysis and Applications, 23, 15–41, 2001.

Amestoy, P., Buttari, A., L'Excellent, J.-Y., and Mary, T.: Performance and Scalability of the Block Low-Rank Multifrontal Factorization on Multicore Architectures, ACM Transactions on Mathematical Software, 45, 2:1–2:26, 2019.

Balay, S., Abhyankar, S., Adams, M. F., Benson, S., Brown, J., Brune, P., Buschelman, K., Constantinescu, E., Dalcin, L., Dener, A., Eijkhout, V., Faibussowitsch, J., Gropp, W. D., Hapla, V., Isaac, T., Jolivet, P., Karpeev, D., Kaushik, D., Knepley, M. G., Kong, F., Kruger, S., May, D. A., McInnes, L. C., Mills, R. T., Mitchell, L., Munson, T., Roman, J. E., Rupp, K., Sanan, P., Sarich, J., Smith, B. F., Suh, H., Zampini, S., Zhang, H., Zhang, H., and Zhang, J.: PETSc/TAO Users Manual, Tech. Rep. ANL-21/39 - Revision 3.23, Argonne National Laboratory, https://doi.org/10.2172/2476320, 2025.

Botto, L.: A geometric multigrid Poisson solver for domains containing solid inclusions, Computer Physics Communications, 184, 1033–1044, https://doi.org/10.1016/j.cpc.2012.11.008, 2013.

Crameri, F., Schmeling, H., Golabek, G. J., Duretz, T., Orendt, R., Buiter, S. J. H., May, D. A., Kaus, B. J. P., Gerya, T. V., and Tackley, P. J.: A comparison of numerical surface topography calculations in geodynamic modelling: an evaluation of the 'sticky air' method, Geophysical Journal International, 189, 38–54, https://doi.org/10.1111/j.1365-246X.2012.05388.x, 2012a.

Crameri, F., Tackley, P. J., Meilick, I., Gerya, T. V., and Kaus, B. J. P.: A free plate surface and weak oceanic crust produce single-sided subduction on Earth, Geophysical Research Letters, 39, https://doi.org/10.1029/2011GL050046, _eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2011GL050046, 2012b.

Duretz, T., May, D., and Yamato, P.: A free surface capturing discretization for the staggered grid finite difference scheme, Geophysical Journal International, 204, 1518–1530, https://doi.org/10.1093/gji/ggv526, 2016.

Gerya: Introduction to Numerical Geodynamic Modelling, Cambridge University Press, ISBN 978-1-107-14314-2, https://sfx.ethz.ch/sfx_lib4ri?url_ver=Z39.88-2004&ctx_ver=Z39.88-2004&ctx_enc=info:ofi/enc:UTF-8&rfr_id=info:sid/sfxit.com:opac_856&url_ctx_fmt=info:ofi/fmt:kev:mtx:ctx&sfx.ignore_date_threshold=1&rft.object_id=4930000000052441&svc_val_fmt=info:ofi/fmt:kev:mtx:sch_svc&, 2019.

Gerya, T. V. and Yuen, D. A.: Robust characteristics method for modelling multiphase visco-elasto-plastic thermo-mechanical problems, Physics of the Earth and Planetary Interiors, 163, 83–105, https://doi.org/10.1016/j.pepi.2007.04.015, 2007.

Gerya, T. V., Stern, R. J., Baes, M., Sobolev, S. V., and Whattam, S. A.: Plate tectonics on the Earth triggered by plume-induced subduction initiation, Nature, 527, 221–225, https://doi.org/10.1038/nature15752, bandiera_abtest: a Cg_type: Nature Research Journals Number: 7577 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Geodynamics;Tectonics Subject_term_id: geodynamics;tectonics, 2015a.

Gerya, T. V., Stern, R. J., Baes, M., Sobolev, S. V., and Whattam, S. A.: Plate tectonics on the Earth triggered by plume-induced subduction initiation, Nature, 527, 221–225, https://doi.org/10.1038/nature15752, bandiera_abtest: a Cg_type: Nature Research Journals Number: 7577 Primary_atype: Research Publisher: Nature Publishing Group Subject_term: Geodynamics;Tectonics Subject_term_id: geodynamics;tectonics, 2015b.

Hernlund, J. W. and Tackley, P. J.: Modeling mantle convection in the spherical annulus, Physics of the Earth and Planetary Interiors, 171, 48–54, https://doi.org/10.1016/j.pepi.2008.07.037, accepted: 2017-06-08T21:22:19Z Publisher: Elsevier, 2008.

https://doi.org/10.5194/egusphere-2025-6354
Preprint. Discussion started: 28 December 2025

Kaus, B. J., Popov, A. A., Baumann, T., Pusok, A., Bauville, A., Fernandez, N., and Collignon, M.: Forward and Inverse Modelling of Lithospheric Deformation on Geological Timescales Forward and Inverse Modelling of Lithospheric Deformation on Geological Timescales, in: Proceedings of nic symposium, vol. 48, pp. 978–983, John von Neumann Institute for Computing (NIC), NIC Series, https://juser.fz-juelich.de/record/507751/files/nic_2016_kaus.pdf, 2016.

Larionov, E., Batty, C., and Bridson, R.: Variational stokes: a unified pressure-viscosity solver for accurate viscous liquids, ACM Transactions on Graphics, 36, 101:1–101:11, https://doi.org/10.1145/3072959.3073628, 2017.

Ogawa, M., Schubert, G., and Zebib, A.: Numerical simulations of three-dimensional thermal convection in a fluid with strongly temperature-dependent viscosity, Journal of Fluid Mechanics, 233, 299–328, https://doi.org/10.1017/S0022112091000496, 1991.

Patankar, S. V.: Numerical Heat Transfer and Fluid Flow, Hemisphere Publishing Corporation, New York, 1980.

Ramberg, H.: Gravity, Deformation, and the Earth's Crust: In Theory, Experiments, and Geological Application, Academic Press, ISBN 978-0-12-576860-3, google-Books-ID: NFESAQAAIAAJ, 1981.

Rauwoens, P., Troch, P., and Vierendeels, J.: A geometric multigrid solver for the free-surface equation in environmental models featuring irregular coastlines, Journal of Computational and Applied Mathematics, 289, 22–36, https://doi.org/10.1016/j.cam.2015.03.029, 2015.

Rolf, T. and Tackley, P. J.: Focussing of stress by continents in 3D spherical mantle convection with self-consistent plate tectonics, Geophysical Research Letters, 38, https://doi.org/10.1029/2011GL048677, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2011GL048677, 2011.

Schmeling, H., Babeyko, A. Y., Enns, A., Faccenna, C., Funiciello, F., Gerya, T., Golabek, G. J., Grigull, S., Kaus, B. J. P., Morra, G., Schmalholz, S. M., and van Hunen, J.: A benchmark comparison of spontaneous subduction models—Towards a free surface, Physics of the Earth and Planetary Interiors, 171, 198–223, https://doi.org/10.1016/j.pepi.2008.06.028, 2008.

Schmid, D. W. and Podladchikov, Y. Y.: Analytical solutions for deformable elliptical inclusions in general shear, Geophysical Journal International, 155, 269–288, https://doi.org/10.1046/j.1365-246X.2003.02042.x, 2003.

Tackley, P.: A Three-Dimensional Version of the Free Surface Capturing Discretization for Staggered Grid Finite Difference Schemes: Implementation into StagYY, Tech. Rep. EGU21-16228, Copernicus Meetings, https://doi.org/10.5194/egusphere-egu21-16228, conference Name: EGU21, 2021.

Tackley, P. J.: Effects of strongly temperature-dependent viscosity on time-dependent, three-dimensional models of mantle convection, Geophysical Research Letters, 20, 2187–2190, https://doi.org/10.1029/93GL02317, _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/93GL02317, 1993.

Tackley, P. J.: Modelling compressible mantle convection with large viscosity contrasts in a three-dimensional spherical shell using the yin-yang grid, Physics of the Earth and Planetary Interiors, 171, 7–18, https://doi.org/10.1016/j.pepi.2008.08.005, 2008.

Tackley, P. J.: The tracer nudging method for correcting and preventing uneven tracer distributions in geodynamical models, EGUsphere, 2025, 1–12, https://doi.org/10.5194/egusphere-2025-1354, 2025.

Tackley, P. J. and King, S. D.: Testing the tracer ratio method for modeling active compositional fields in mantle convection simulations, Geochem. Geophys. Geosyst., 4, doi:10.1029/2001GC000 214, 2003.

Trompert, R. A. and Hansen, U.: The application of a finite-volume multigrid method to 3-dimensional flow problems in a highly viscous fluid with a variable viscosity, Geophys. Astrophys. Fluid Dyn., 83, 261–291, 1996.