Geosci. Model Dev., 19, 8877–8894, 2026
https://doi.org/10.5194/gmd-19-8877-2026

Geoscientific Model Development
Open Access
EGU

Lagrangian tracking methods applied to free surface boundaries in numerical geodynamic models

Timothy S. Gray, Paul J. Tackley, and Taras V. Gerya

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

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

Received: 30 December 2025 – Discussion started: 23 January 2026
Revised: 18 May 2026 – Accepted: 6 July 2026 – Published: 21 September 2026

Abstract. A desirable characteristic of mantle convection models is the ability to determine surface topography evolution over time on a global scale. This capability is increasingly important due to growing interest in coupling planetary geodynamics with climate, landscape, habitats, and biological evolution systems. A common way of achieving this in numerical geodynamic models with a fixed Eulerian grid is through the implementation of a free-surface boundary condition using the sticky air method in conjunction with an appropriate method of tracking the free surface. Although existing methods for tracking the interface between the air and rock layers are available, they often struggle to provide high-resolution results on a global scale without incurring significant computational costs. We propose a method for representing surfaces directly using Lagrangian markers that can track the location of a free surface in 2D and 3D models and test it using the finite-volume mantle convection code StagYY. This approach offers a direct, high-resolution representation of the surface without the need for a large number of high-cost volumetric markers throughout the model domain. Benchmarks demonstrate the effectiveness of this method compared to the commonly-used marker-in-cell method. The direct representation of the surface enables additional features such as the direct tracking of sea levels over time, potential for coupling with surface process models on the global scale, and enables the implementation of alternative discretisations of the Stokes equations to improve Stokes solver accuracy near the free surface boundary.

1 Introduction

Modelling the evolution of planetary surfaces over time has attracted considerable attention in geodynamics, since surface topography is a primary observable of a planet's internal dynamics and can also change the dynamics compared to assuming a flat surface (e.g. Crameri et al., 2012b). The growing interest in coupled planetary systems, particularly the coupling between geodynamic processes, climatic processes, and biological evolution, further motivates the study of topographic evolution due to the coupling between these systems occurring primarily at the surface. The study of these couplings, also referred to as *biogeodynamics* (Zerkle, 2018; Stern and Gerya, 2023), requires accurate methods of determining how mantle convection affects the surface, and in turn affects other systems in order to study fundamental questions regarding the habitability of planetary bodies.

Geodynamic models commonly represent a free surface using a fixed, Eulerian grid with "sticky air", a low-viscosity, low-density layer above the surface that approximates a true free surface boundary condition, representing material above the surface (e.g. the various modelling codes in Crameri et al., 2012a and Schmeling et al., 2008). Modelling free surfaces on fixed Eulerian grids requires two key components: an appropriate solver that can model or approximate the solution to the Stokes equation with a free surface boundary condition, and a method of tracking the location of the surface over time.

This study implements a method for tracking free surfaces on a global scale in StagYY (Tackley, 2008), which uses the staggered-grid finite difference discretisation as is common 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., 2015; Kaus et al., 2016). The free surface tracking methods presented are also generally applicable to other Eulerian-grid geodynamic codes.

The standard technique for interface tracking in StagYY and many other codes is the marker-in-cell method, in which Lagrangian markers distributed throughout the model domain carry material type information (rock or air), which is averaged onto the grid to determine density and viscosity fields and infer the surface location (Tackley and King, 2003). Two averaging approaches are common: simple cell (boxcar) averaging, and the more accurate linear shape function averaging used as the default in StagYY, which additionally weights markers by their position within the cell. While effective, both approaches share a fundamental limitation: surface quality depends strongly on the number of markers per cell (here referred to as marker density). For a global-scale 3D model with a typical vertical resolution of $\approx 10\,\mathrm{km}$ (Coltice et al., 2019), resolving topographic variations of order $\pm 10\,\mathrm{km}$ at 100 m sensitivity would require at least 100 markers per cell. Achieving this is computationally prohibitive: marker advection cost scales linearly with marker density, and in practice dominates total runtime at the densities required for accurate surface tracking, accounting for up to 77 % of total runtime at 500 markers per cell in our tests. Doubling vertical marker density requires $4\times$ more markers in 2D and $8\times$ more in 3D, making this approach impractical at global scale.

For clarity, the Lagrangian points already implemented in StagYY are hereafter referred to as "markers", while Lagrangian points introduced in this study specifically for tracking the free surface are referred to as "surface markers".

A successful surface tracking method must provide fine sub-grid resolution of topography on the global scale, remain stable over many timesteps and across diverse geodynamic regimes including subduction, ensure compatibility with the underlying Stokes solver, and reduce computational cost relative to the marker-in-cell method. Interface-tracking approaches meeting these criteria include level-set and volume-of-fluid methods. However, one of the conceptually simplest options is to track interfaces directly using Lagrangian markers, referred to as direct interface tracking in the CFD literature (Tezduyar, 2006), in which the surface is explicitly represented by markers that follow the surface flow. Confining markers to the surface alone avoids the scaling limitations of the full marker-in-cell approach.

Several characteristics of planetary surfaces make this approach particularly attractive. For Earth-like planets, the maximum surface deviation ($\pm 10\,\mathrm{km}$) is small relative to total mantle thickness (2890 km), and the surface height can be treated as a single-valued function of position. Planetary surfaces are sharp interfaces with no diffuse mixed region. Many geodynamic codes already use Lagrangian markers, enabling straightforward implementation by adapting existing routines. The number of surface markers required is small compared to the total marker population, and surface markers are well-suited for parallelisation including on GPUs.

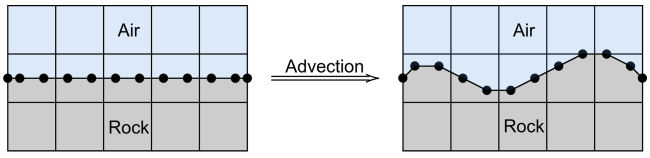

Figure 1. The marker chain divides the domain into regions of rock and air. Lagrangian markers are advected with the velocity field at each timestep to model the evolution of the surface over time.

In this study, we implement Lagrangian interface tracking in StagYY and introduce two methods: the *bilinear method* and the *integral method*, which differ in how they transfer surface information to the underlying grid. We assess their computational performance and accuracy in tracking a free-surface boundary condition, comparing them with the existing marker-in-cell approach.

2 Surface tracking

2.1 Marker chain/mesh

In 2D models, the surface can be tracked by a marker chain which is initialised as an ordered array of Lagrangian markers, each of which stores its position $\boldsymbol{x}_i$ and a type variable to facilitate manipulation of the chain (Fig. 1). The position vector for each marker $\boldsymbol{x}_i$ may be represented in either Cartesian coordinates $\boldsymbol{x}_i = (x_i, z_i)$, or polar coordinates $\boldsymbol{x}_i = (r_i, \phi_i)$, and may be dimensional or dimensionless. The height of each marker (z_i or r_i) is set by an initial surface height function $z_{\mathrm{init}}(x)$ or $r_{\mathrm{init}}(\phi)$.

In 3D models the surface representation extends to a mesh. The initial surface consists of markers arranged in a grid with coordinates $\boldsymbol{x}_i = (x_i, y_i, z_i)$ in Cartesian geometry, or $\boldsymbol{x}_i = (r_i, \theta_i, \phi_i)$ in spherical geometry. As with the 2D case, the initial height of the markers z_i or r_i is given by an initial surface height function $z_{\mathrm{init}}(x, y)$ or $r_{\mathrm{init}}(\theta, \phi)$.

Surface markers are advected with the velocity field, determined from the Stokes solver, using a spatially fourth-order Runge-Kutta scheme, similar to the treatment of the markers used in StagYY. Velocities, which are face-centred in the staggered-grid finite-volume discretisation used by StagYY, are interpolated to surface markers using a second-order quadratic spline interpolation method (Gerya et al., 2021).

2.2 Piecewise linear surface reconstruction method (bilinear method)

It is necessary to be able to determine the location of the surface as a function of position through interpolation, i.e. finding $z(x)$ or $r(\phi)$ in 2D, or $z(x, y)$ or $r(\theta, \phi)$ in 3D as a function of the position of the Lagrangian surface markers.

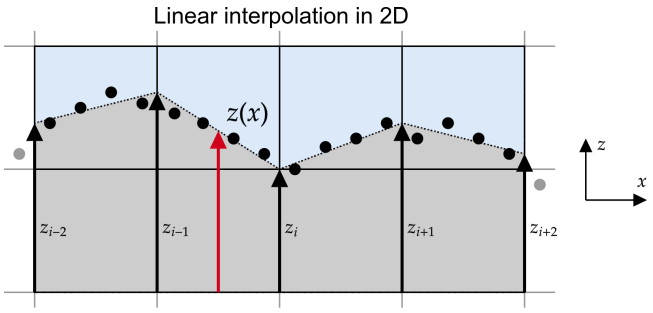

Figure 2. When using the bilinear representation of the surface, the surface markers are interpolated to the nodes of the Eulerian grid through distance-weighted linear interpolation. The surface height at a given point e.g. $z(x)$ (red arrow) may then be computed using linear interpolation.

A straightforward approach is to construct a piecewise continuous linear or bilinear representation of the surface using a first-order accurate distance-weighted linear interpolation (Gerya, 2019, Chap. 18). This method was employed with a 2D marker chain in the context of modelling free surfaces in numerical geodynamic models by Duretz et al. (2016), and a similar method is used in 3D models using the regional scale geodynamic code LaMEM (Kaus et al., 2016). The method requires two steps: the interpolation of surface height data from Lagrangian surface markers to nodal points using distance-weighted linear interpolation, and the linear or bilinear interpolation of nodal values to arbitrary locations within the domain. Once surface heights are interpolated to nodes, determining the surface height at any point can be accomplished in constant $O(1)$ time, independently of the number of surface markers.

For consistency, this approach is referred to throughout as the *bilinear method*, reflecting the bilinear interpolation used in 3D models. In 2D, the process is technically linear, but the term *bilinear method* is retained for uniformity.

2.3 Direct interpolation (integral method)

Another method of determining the surface location at arbitrary points is through direct interpolation from the marker chain or mesh, without first mapping to the Eulerian grid. In 2D, this is achieved by locating the markers on the chain immediately to the left and right of the interpolation point and linearly interpolating between them (Fig. 2). Higher order methods such as piecewise quadratic or cubic spline interpolation are possible, but in practice simple linear interpolation produces adequate results.

As advection allows markers to move laterally, the chain must be re-sorted by horizontal coordinate after each timestep. Since the chain is typically nearly sorted after advection, adaptive algorithms such as insertion sort are efficient for this purpose, and an $O(\log n)$ binary search can then be used to locate neighbouring markers for interpolation.

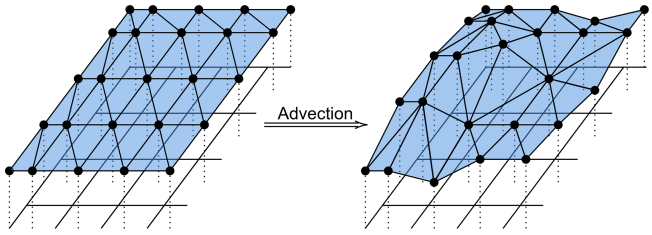

Figure 3. In 3D, Lagrangian markers are advected with the velocity field, and the Delaunay triangulation recomputed at each timestep in order to create a mesh that continuously tracks the surface.

In 3D, direct interpolation is more involved since the surface is represented by planes rather than line segments. As the surface markers are generally unstructured, a Delaunay triangulation of the horizontal marker coordinates is used (Fig. 3). The triangulation is computed using GEOMPACK2 (Joe, 1991) in $O(n \log n)$ time and must be recomputed at each timestep; adaptive methods exist to reduce this cost (Phongthanapanich and Dechaumphai, 2004). Methods for performing Delaunay triangulations directly on spheres also exist (Renka, 1997), though for the yin-yang grid (Tackley, 2008) the errors introduced by projecting the spherical patch into a Cartesian plane are negligible. Each Delaunay triangle defines a plane through three markers $\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{c}$, with normal vector $\boldsymbol{n} = (\boldsymbol{b} - \boldsymbol{a}) \times (\boldsymbol{c} - \boldsymbol{a})$, from which the surface height at any point within the triangle can be recovered. An efficient $O(\log n)$ triangle location method is available in GEOMPACK2; if the point lies outside the triangulation, the nearest triangle is used for extrapolation.

2.4 Interpolating to the Eulerian grid

In order to couple the surface tracking method to the Stokes solver and to output results, the surface heights of the marker chain or mesh must be interpolated to the Eulerian grid. When using the bilinear method, this step is handled automatically by interpolating surface heights to the nodes of the Eulerian grid. For the integral method, a cellwise integration approach is used in both 2D and 3D.

The cellwise integration method computes the area (in 2D) or volume (in 3D) enclosed by the surface within each computational cell, and derives the equivalent topographic height that preserves this area or volume. This is physically motivated by the need to maintain conservation of mass of both rock and air within the model. In 2D the integral is evaluated using the trapezium rule by looping over marker chain intervals, which is accomplished in linear time $O(n)$ as the chain is already sorted. In 3D, Delaunay triangles are intersected with the footprint of each cell column using a polygon intersection algorithm (O'Rourke et al., 1982); the volume beneath each intersection polygon is then computed and summed to obtain the surface height per column (Fig. 4).

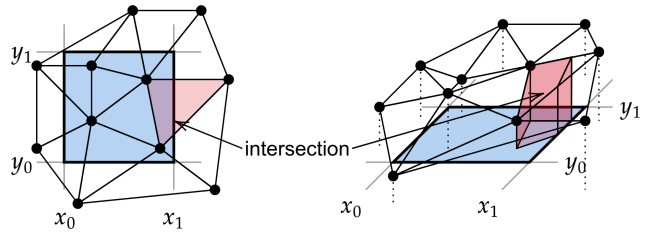

Figure 4. Piecewise integration under the Delaunay mesh using a polygon clipping algorithm is used to compute the volume enclosed by the surface within a given computational cell.

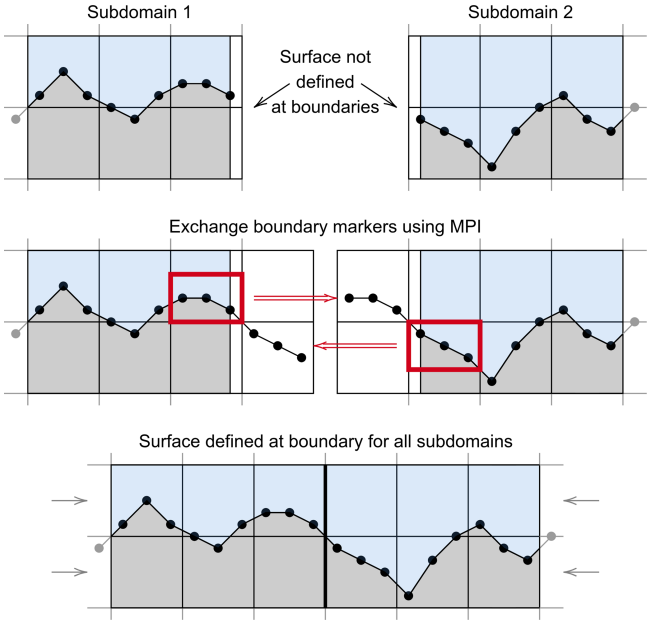

Figure 5. Parallelisation of the marker chain/mesh involves the use of an overlap region communicated between subdomains via MPI. This approach is also applicable to 3D cases.

Both 2D and 3D operations scale as $O(n)$ with respect to the number of surface markers.

When solving on multiple CPU cores, the domain is split into parallel subdomains. Surface markers near subdomain boundaries are exchanged via MPI to form a one-cell-wide overlap region, ensuring the surface is defined at all locations for all subdomains (Fig. 5).

3 Stabilisation

When implemented naively, the surface marker chain/mesh does not satisfy the criterion of being stable over multiple timesteps. There are three main reasons for this. Firstly, while the solution to the Stokes equations ensures conservation of mass, the surface marker chain/mesh itself does not, and may shift up and down over time. Secondly, regions of divergence can cause the surface markers to become widely spaced, compromising accuracy. Thirdly, regions of convergence can cause entrainment of markers into the mantle, a feature which is unphysical and must be corrected (Fig. 6). There are several ways that these stability issues can be addressed.

3.1 Surface height adjustment and reinitialisation

Minor numerical errors in the Stokes solution and/or marker advection can cause the surface to drift upward or downward over time. To correct this, the height of all surface markers is adjusted at each timestep to preserve the total area or volume of rock and air in the model, and therefore mass under the assumption of near-incompressibility. This is achieved by integrating under the marker chain or mesh using the same trapezium rule or Delaunay-based approach described in Sect. 2.4, with corrections for the spherical area element $\mathrm{d}A = r^2 \sin\theta \, \mathrm{d}\theta \, \mathrm{d}\phi$ (Kreyszig, 2010) applied where necessary.

Although reinitialisation is essential for preserving the integrity of the marker chain over multiple timesteps, it does come with trade-offs: the process may degrade information about the surface geometry and introduce artificial numerical diffusion. Therefore, reinitialisation should only be applied when needed.

Two reinitialisation strategies are implemented. The first is *periodic global reinitialisation*, in which the entire surface marker chain/mesh is reinitialised to align with the current surface level at every timestep. This is computationally efficient and maintains a consistent number of surface markers, but introduces numerical diffusion that can smooth surface features over time (Fig. 7). It does not require the chain to be sorted or the Delaunay triangulation to be computed, making it compatible with the bilinear method. Despite the diffusion it introduces, some smoothing is beneficial for stability, and the effect is relatively mild in practice.

The second strategy is *interval reinitialisation*, in which markers are added or removed locally based on the spacing between adjacent markers (in 2D) or the area of Delaunay triangles (in 3D). In 2D, a marker is added at the midpoint of any interval exceeding $2\Delta x_{\mathrm{init}}$, and two markers are merged if their spacing falls below $0.2\Delta x_{\mathrm{init}}$; empirical testing confirmed these thresholds give good results. The 3D analogue uses triangle area ratios A/A_{init} with an upper threshold of 3 and a lower threshold of 0.2. This approach avoids the numerical diffusion inherent in global reinitialisation by only modifying the chain where necessary. However, it requires either sorting or Delaunay triangulation of the surface markers, which precludes its use with the bilinear method.

3.2 The need for smoothing

A fundamental problem with the Lagrangian surface marker method is the undesirable entrainment of markers into the mantle or air layer instead of remaining at the surface. The

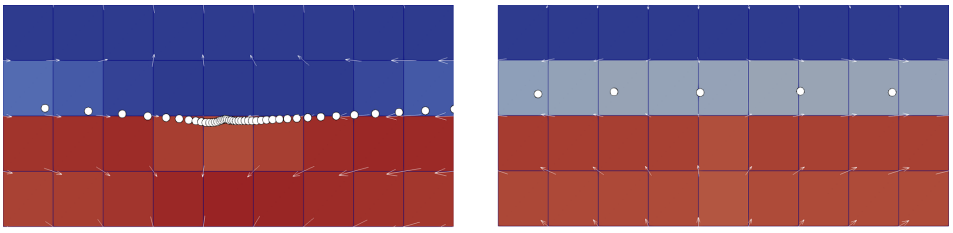

(a) (b)

Figure 6. Without reinitialisation, regions of convergence **(a)** experience bunching of the surface markers, while in regions of divergence **(b)** the marker density becomes low. These effects must be corrected using reinitialisation to maintain surface marker density.

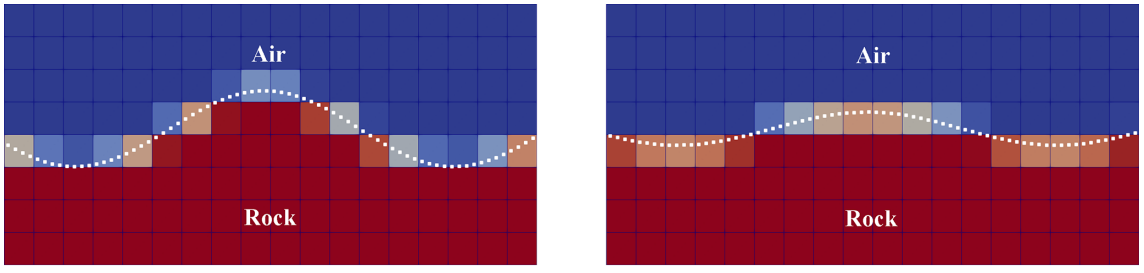

(a) (b)

Figure 7. Illustration of the smoothing effect caused by global periodic reinitialisation over 20 timesteps using bilinear interpolation, shown at timestep 1 **(a)** and timestep 20 **(b)**. There is no coupling between the surface marker chain and the underlying model; all deformation of the surface is caused by numerical diffusion resulting from global reinitialisation.

effect is most pronounced when using the integral method, and results in the formation of unrealistically deep trenches or high mountains that can cause numerical instabilities (Fig. 8). Solving this problem requires a method of smoothing the Lagrangian markers.

A natural choice for smoothing is diffusion, which is conservative and physically motivated. However, it can smooth model features unnecessarily, and periodic global reinitialisation, if used, already introduces some numerical diffusion. The method implemented here for the integral method is instead a novel slope limiting approach.

3.3 Slope limiting

The slope limiting method is physically motivated by the behaviour of loosely packed granular materials: a pile of sand under gravity forms a cone with a maximum angle of repose ϕ, related to the internal coefficient of static friction by $\tan(\phi) = \mu_s$. Slopes below this angle are stable; those above erode until the angle is reached. This behaviour is expressed as a steady-state PDE based on the diffusion equation:

$$\nabla^2 z \, H\left(||\nabla z|| - s_{\text{crit}}\right) = 0, \tag{1}$$

where H is the Heaviside step function and s_{crit} is the critical slope. The Heaviside function ensures that only regions where the slope exceeds s_{crit} are modified, limiting the smoothing effect to regions of high deformation such as subduction or collision zones. The nonlinearity introduced by the Heaviside function is handled using an iterative solver in both 2D and 3D. Applied to sinusoidal test surfaces (Figs. 9 and 10), the method flattens steep wave flanks while leaving shallower regions unaffected.

A key advantage of slope limiting over diffusion is that it is timestep independent, producing consistent results across different temporal and spatial scales. A limitation is its relatively poor computational performance (see Sect. 5.6.1), and the fact that it only modifies the surface location and not the underlying markers, which can become unmixed as a result. Solutions to this are discussed in Sect. 4.1.

4 Coupling between surface tracking and the Stokes solver

Although it is useful to track the location of the free surface by itself, the utility of the method is significantly increased by coupling the surface representation to the mechanical solver such that the height of the surface affects cell-averaged density (hence buoyancy) and viscosity. There are two methods of coupling: indirectly through manipulation of the existing Lagrangian volumetric markers in the model, which then determine the density and viscosity fields, and directly through modification of these fields using volume fraction functions.

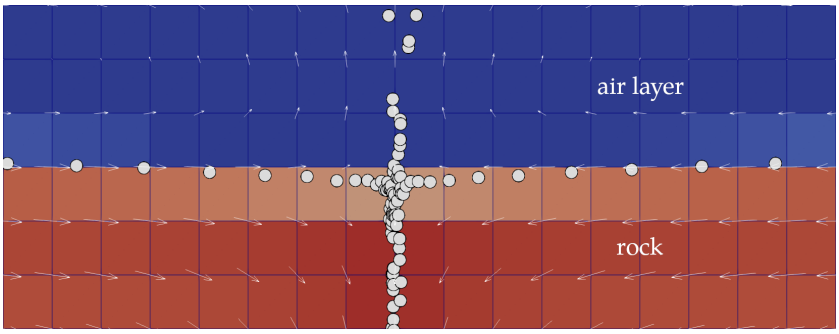

Figure 8. In the absence of an appropriate smoothing algorithm, surface markers can become entrained within rock (red) or the sticky air layer (blue), causing significant numerical instability.

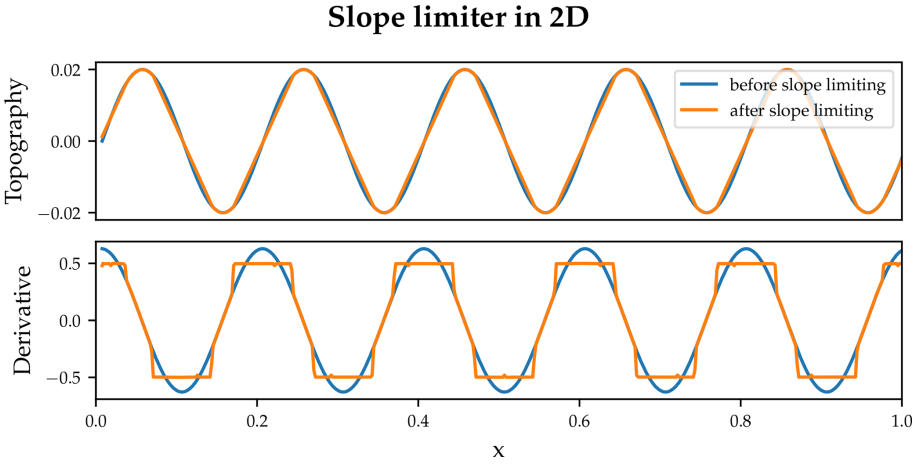

Figure 9. Slope limiter in 2D showing topography (top) before and after the slope limiter has been applied to an initial surface $z = 0.02\sin(5\pi x)$, and its derivative (bottom) showing the limiting of the slope to a maximum of $s_{\text{crit}} = 0.5$.

4.1 Marker unmixing

During advection, rock markers can end up above the surface and air markers below it. This can also occur when the surface location is directly manipulated by a smoothing algorithm. The resulting erroneous marker distribution leads to numerical instability and must be corrected. Two methods were implemented (Fig. 11).

In the *bouncing* method, erroneous markers are reflected back to their correct side of the surface, with the displacement equal to their distance from the surface. No markers are created or destroyed, making this method mass-conserving and applicable to models with many marker types. A downside is that bouncing can create gaps in the marker field in convergent regions where the surface is modified by smoothing but the underlying markers are not, which over time can cause numerical instability.

In the *exchange* method, erroneous markers are directly converted to the appropriate type: air markers below the surface become rock markers, and vice versa. This avoids the gap-forming tendency of the bouncing method. For the conversion of air markers to rock markers, a random template marker is selected from the cell immediately below the surface. This works well for simple compositional setups, but becomes unsuitable when markers carry physical properties such as temperature or viscosity, or when many distinct compositions are present. The method is not strictly volume conserving, though on average the number of conversions in each direction remains roughly equal.

4.2 Density and viscosity coupling

StagYY typically determines the density and viscosity fields from the marker-in-cell marker distribution. With Lagrangian surface tracking, greater accuracy can be achieved by coupling the surface location directly to these fields through volume fractions. The volume fraction ϕ of a control volume is defined as the fraction of that volume occupied by rock, ranging from 0 (completely empty) to 1 (completely full), with partially filled cells near the surface having $0 < \phi < 1$ (Fig. 12).

StagYY uses a fully staggered grid (Fig. 13), a standard layout in geodynamic modelling (Patankar, 1980; Ogawa et al., 1991; Gerya, 2019), in which the density field is de-

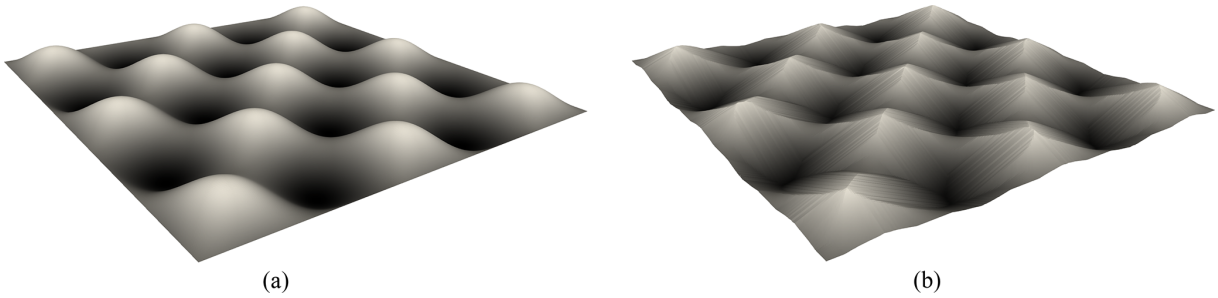

Figure 10. Effect of the slope limiting algorithm with $s_{\text{crit}} = 0.5$ in 3D on an initial surface $z = 0.05\sin(5\pi x)\sin(5\pi y)$, before **(a)** and after **(b)** slope limiting.

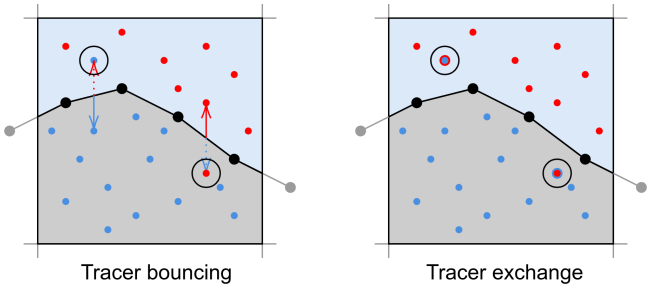

Figure 11. Two methods of dealing with markers which are erroneously on the wrong side of the surface. Air markers (red) should remain above the surface, while rock markers (blue) should remain below. In the bouncing method, erroneous markers (circled) are bounced back to their side, while with the exchange method they are directly converted to the opposite type.

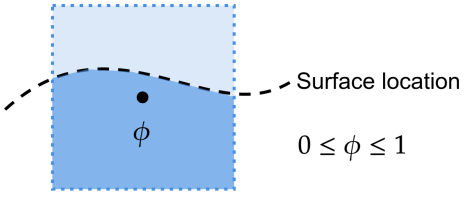

Figure 12. The volume fraction function ϕ for a given control volume is defined as 0 if the control volume is completely empty, and 1 if it is completely full.

fined at cell centres and v_z points, and viscosity at up to four locations (P, σ_{xz}, and in 3D also σ_{xy} and σ_{yz} points). Volume fractions are computed for each of these control volumes offset from the main cells.

As an example, the effective density used in the z-momentum equation is given by:

$$\rho_{\text{eff}} = \rho_{\text{ref}}\,\phi_{v_z}, \tag{2}$$

where ρ_{ref} is a reference near-surface rock density and ϕ_{v_z} is the volume fraction for the v_z control volume (Fig. 14). A similar approach is applied to the viscosity field at each of the relevant staggered grid locations.

This coupling scheme requires that the density and viscosity fields not already account for the sticky air layer; otherwise, its contribution would be counted twice, once through the field values themselves and once through ϕ. We avoid this by setting the air's density and viscosity to the same reference values used in the coupling. One limitation of this is that, because a single reference value is assumed, the approach loses accuracy where surface density or viscosity varies strongly along the surface. In that case, the reference value can instead be taken locally from the cell immediately below the surface at each point.

Volume fractions are computed geometrically in 2D using the marker chain, and in 3D using the same Delaunay-based approach described in Sect. 2.4, both in linear time $O(n)$. In all benchmark tests, computing volume fractions incurred negligible cost compared to other surface marker operations.

5 Tests and results

Three methods were chosen for testing and benchmarking: the existing marker-in-cell method (baseline), the *bilinear method* using surface markers with linear/bilinear interpolation in 2D/3D (Kaus et al., 2016; Duretz et al., 2016), and the *integral method* using cellwise integration to evaluate surface heights on the Eulerian grid. All methods were tested using the sticky air method to isolate the effect of surface tracking. In all benchmarks, 100 volumetric markers per cell are used for the marker-in-cell method, a density chosen to represent a reasonable compromise between accuracy and computational cost as demonstrated in the introduction. The same marker population is retained for the Lagrangian methods to ensure a fair comparison by isolating the effect of the surface tracking approach alone. Default parameter combinations for 2D and 3D benchmarks are summarised in Tables 1 and 2.

5.1 2D relaxation benchmark

The 2D surface relaxation benchmark 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:

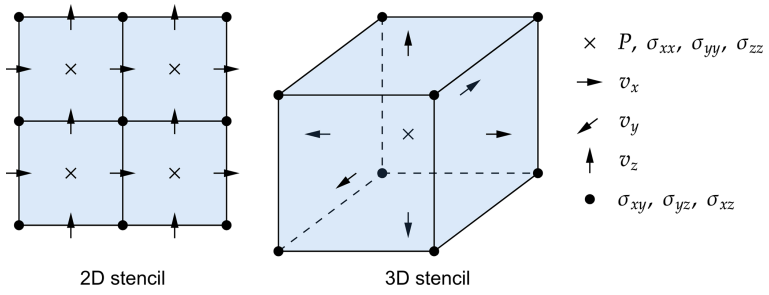

Figure 13. The staggered-grid stencil in 2D and 3D used by StagYY. Note that the vertical axis is denoted by z for both 2D and 3D.

Table 1. Summary of Lagrangian surface tracking options used in 2D benchmarks.

Method	Surface MPC	Reinitialisation	Smoothing method	Unmixing method
Marker-in-cell	–	–	–	bouncing
Bilinear method	8	global reinitialisation after each timestep	–	bouncing
Integral method	8	interval method $\Delta x/\Delta x_{\text{init}} \in [0.2, 2.0]$	slope limiter $s_{\text{crit}} = 0.5$	exchanging

A summary of the combination of options used for the three methods compared in 2D benchmarks. An initial density of 8 surface markers per cell (MPC) and 100 volumetric markers per cell are used in all 2D models.

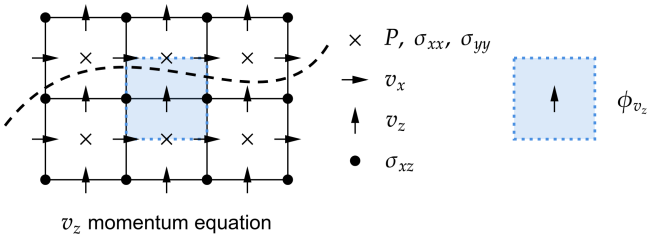

Figure 14. The ϕ_{v_z} volume fraction is used to couple the surface location to the density used in the z-momentum equation near the surface.

$$z_{\text{init}}(x) = 7\cos\left(\frac{2\pi x}{2800\,\text{km}}\right)\text{km}. \tag{3}$$

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:

$$z_{\text{max}}(t) = 7\exp\left(\frac{t}{14.825\,\text{kyr}}\right)\text{km}, \tag{4}$$

where t is time and 14.825 kyr is the analytical relaxation timescale derived from the model viscosity structure.

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 markers per cell for advecting material properties and tracking the surface in the case of the marker-in-cell method. The sticky air layer has a viscosity of 10^{18} Pa s. 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 $t = 100$ kyr to show the system reaching equilibrium. Timesteps were limited to 500 years in order to give high temporal resolution for plots of model evolution.

The three surface tracking methods can be compared against the analytical solution when plotting maximum topography with time, as in Fig. 15.

All methods are qualitatively able to track the analytical solution closely in agreement with the results presented in Crameri et al. (2012a). Minor divergence between the methods becomes more apparent after 30 kyr, with the marker-in-cell method failing to match the analytical solution as closely as Lagrangian tracking methods after this point. At 100 kyr, while the other methods have converged to a maximum topography near zero, the marker-in-cell method resulted in around 200 m of topography, a result that can also be seen in Fig. 2 of Crameri et al. (2012a). This result can be attributed to the maximum topography being affected by the noise inherent in the marker-in-cell method when using a finite number of randomly initialised markers throughout the model. Meanwhile, the Lagrangian surface marker-based tracking methods all converge to near zero at 100 kyr, indicating that Lagrangian surface tracking methods are less susceptible to this noise.

It should be noted that part of the remaining discrepancy between the methods presented here and the analytical solution can also be explained by the limitations of the sticky air method (i.e. finite viscosity of air) used to produce the

Table 2. Summary of Lagrangian surface tracking options used in 3D benchmarks.

Method	Surface MPC	Reinitialisation	Smoothing method	Unmixing method
Marker-in-cell	–	–	–	bouncing
Bilinear method	16	global reinitialisation after each timestep	–	bouncing
Integral method	16	interval method $\Delta x/\Delta x_{\text{init}} \in [0.2, 3.0]$	slope limiter $s_{\text{crit}} = 0.5$	exchanging

A summary of the combination of options used for the three methods compared in 3D benchmarks. An initial density of $16 = 4 \times 4$ surface markers per cell (MPC) is used in all 3D models.

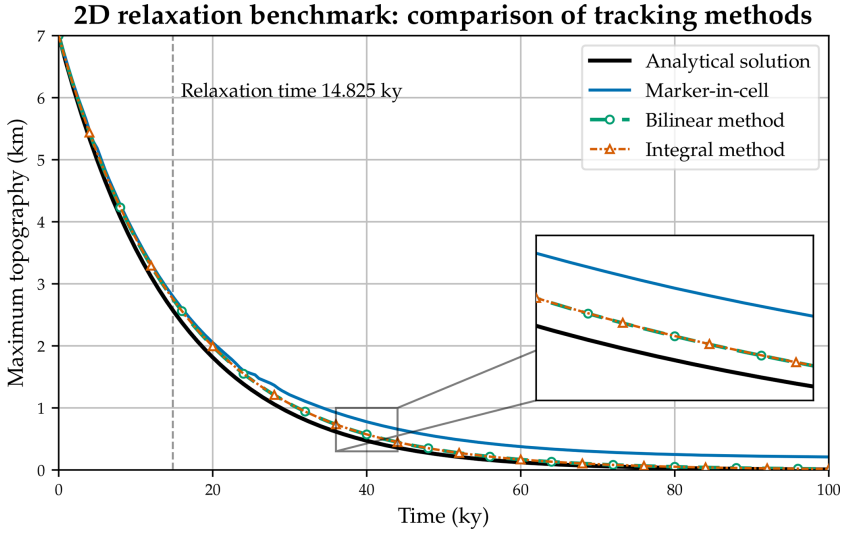

Figure 15. Maximum topography with time for the 2D relaxation benchmark. Both Lagrangian marker based methods closely match the analytical solution. The existing marker-in-cell method, while also tracking the analytical solution closely initially, fails to converge with the analytical solution near the end of the model run.

solution to the Stokes equations rather than the surface tracking method itself. The primary value of this benchmark is therefore in the relative comparison between methods: the Lagrangian tracking methods converge more closely to zero at 100 kyr than the marker-in-cell method, demonstrating a clear reduction in marker-density-dependent noise that is independent of the Stokes solver used.

5.2 2D plume benchmark

The 2D plume benchmark is based on Case 2 presented in Crameri et al. (2012a), considering a buoyant plume (radius 100 km, density 3200 kg m^{-3}, i.e. 100 kg m^{-3} less than the surrounding mantle and lithosphere) rising through a 2800×850 km domain at 1024×256 resolution over 20 Myr. The sticky air layer has a viscosity of 10^{19} Pa s. Free slip boundary conditions are imposed at the horizontal edges and top of the model, while a no-slip boundary condition is imposed at the bottom. All other parameters follow Crameri et al. (2012a).

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 of 800–850 m after 20 Myr.

The results of the plume benchmark further demonstrate the strengths of Lagrangian tracking methods in isolating the topographic response of the model without the noise associated with the marker-in-cell method. The marker-in-cell method introduces random noise, which results in a $\approx$ 100 m discrepancy in topography at the centre of the plume. Lagrangian tracking methods closely matched the results demonstrated by Crameri et al. (2012a) for both the maximum topography at 20 Myr and the evolution of topography in the intervening time.

5.3 2D scaling tests

Scaling tests used a non-dimensional model in spherical annulus geometry with constant viscosity ($\text{Ra} = 10^7$), run for 5000 timesteps to ensure statistically steady-state convection (Fig. A1 in the Appendix). Figure 17 shows performance as a function of surface marker density for both methods. The bilinear method consistently accounts for approximately 3 % of total runtime across all tested MPC values, while the integral method peaks at 11 % at 64 MPC. In both cases the dominant cost is marker unmixing, which requires looping through all

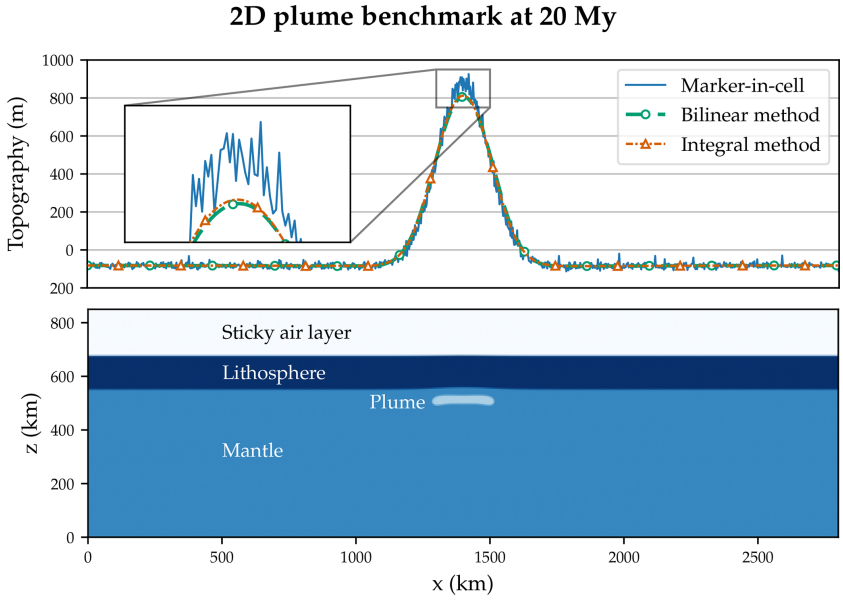

Figure 16. After 20 Myr, all methods produce topography profiles broadly consistent with those reported in Crameri et al. (2012a) (their Fig. 7), both in peak amplitude and overall shape. The marker-in-cell method demonstrates significantly more noise compared to the Lagrangian methods and a systematically raised maximum topography.

markers in the model to determine their position relative to the surface: a procedure taking $O(1)$ time per marker for the bilinear method but $O(\log n)$ for the integral method. Tests across four resolutions (128×16 to 1024×128) at a fixed density of 8 MPC confirm this pattern: unmixing scales with marker count rather than surface marker count, leaving the bilinear method below 3 % of total runtime at all resolutions tested.

5.4 3D plume benchmark

The 2D plume benchmark was extended to 3D on a $1400 \times 1400 \times 850$ km Cartesian domain at $256 \times 256 \times 64$ resolution, with the plume geometry modified to a sphere of initial diameter 100 km. A no-slip boundary condition is imposed at the lower boundary, while free slip conditions are imposed on all other boundaries. All other physical parameters are identical to the 2D case. As no prior 3D results exist for comparison, results are assessed qualitatively.

The performance of different tracking methods, as illustrated in Fig. 18, highlights the inherent noise challenges of the marker-in-cell method in 3D. This method struggles with noise generation due to the inherent dependence on marker density and positioning. Nonetheless, it is still able to capture the main features of the mantle plume, particularly the central topographic uplift created by the buoyant upwelling. Both the bilinear and integral methods eliminate this noise, providing visually similar results that isolate the plume's influence on topography, with minor differences between the two methods visible in the cross-sectional profile (Fig. 19), possibly a side effect of the differing smoothing methods used.

5.5 3D constant-viscosity Cartesian convection

A simple Cartesian box setup (Fig. 20) with a relatively coarse $64 \times 64 \times 32$ resolution was used to compare the three methods under constant-viscosity convection with non-dimensional parameters ($\mathrm{Ra} = 10^7$). Visual inspection of Fig. 21 confirms that both the bilinear and integral methods eliminate the marker-density-dependent noise of the marker-in-cell method, producing visually smooth topography that isolates the underlying plume and downwelling structure.

5.6 3D scaling tests

The same Cartesian box model was used to test scaling of surface tracking methods with respect to model resolution and marker density.

5.6.1 Resolution dependence

The resolution dependence test evaluated the model at four resolutions ranging from $32 \times 32 \times 16$ to $256 \times 256 \times 128$. Figure 22 shows that smooth topography is achieved consistently across all length scales using the integral method, in contrast to the grainy topography produced by the marker-in-cell method; similar results are obtained with the bilinear method. Higher resolutions resolve greater detail, particularly around downwellings, while lower resolutions produce more diffuse but still smooth features. The median topography increases with resolution while the mean remains at 0.0, consistent with deeper trenches being better resolved at higher resolutions. This resolution dependence can

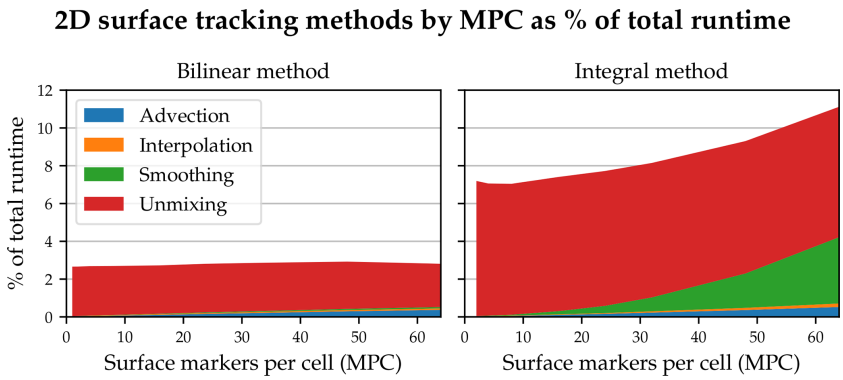

Figure 17. 2D scaling test evaluating the impact of surface marker density (MPC) on tracking performance. The integral method demands more computational effort than the bilinear method, reaching a maximum of 11 % of total computational effort at 64 MPC compared to 3 % for the bilinear method. Marker unmixing accounts for the majority of computational cost in both methods.

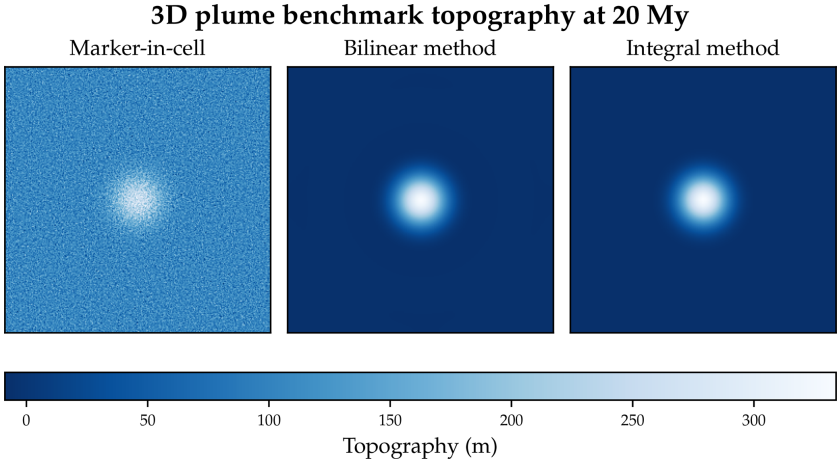

Figure 18. 3D plume benchmark topography at 20 Myr (top-down view). The marker-in-cell method with 100 markers per cell captures the topography generated by the mantle plume but with visually apparent noise. Both the bilinear and integral methods eliminate this noise.

only be observed using the Lagrangian tracking methods, as the marker-in-cell results are dominated by marker-density-dependent noise at equivalent marker counts.

The performance of each method at different resolutions is shown in Fig. 23. As in 2D, marker unmixing dominates runtime for both methods. For the bilinear method, surface location lookups are $O(1)$, keeping unmixing cost roughly constant across resolutions. For the integral method, the $O(\log n)$ lookup cost means unmixing generally increases with resolution. The slope limiting algorithm adds further cost to the integral method, scaling with resolution due to the increased iterations required to relax the marker mesh.

5.6.2 Marker density dependence

A key question is whether the implemented methods eliminate the dependence on volumetric marker density discussed in the introduction. Figure 24 shows surface topography produced by the integral method across a range of volumetric marker densities: no qualitative differences are visible, in stark contrast to marker-in-cell results. The same result is obtained with the bilinear method. This confirms that model performance can be improved by reducing marker density without compromising topographic quality, since the density and viscosity fields near the surface are determined directly from the tracked surface position rather than from the marker distribution.

6 Summary and discussion

6.1 Summary of key results

This study implemented and benchmarked two Lagrangian surface marker tracking methods, the bilinear method and the integral method, in the mantle convection code StagYY, comparing them against the existing marker-in-cell approach. Both new methods produce high-resolution, low-noise surface topography that is independent of volumetric marker density, unlike the marker-in-cell method, allowing signifi-

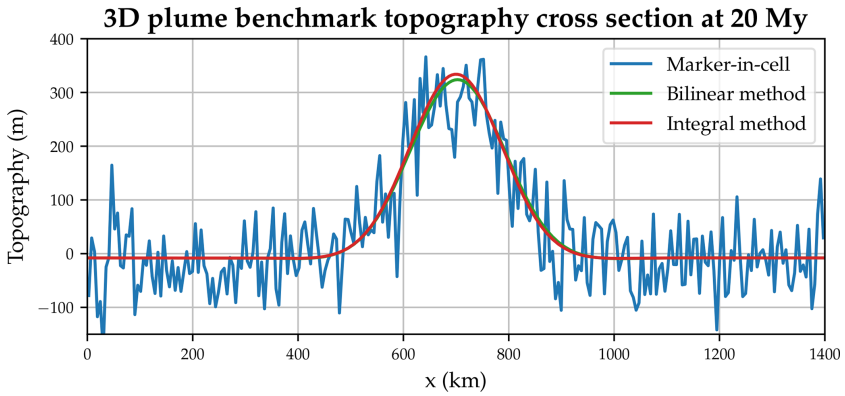

Figure 19. 3D plume benchmark topography at 20 Myr (cross section). Both the bilinear and integral methods effectively capture topography without the noise inherent in the marker-in-cell method. Minor differences between the two Lagrangian methods are visible, likely a consequence of the differing smoothing approaches.

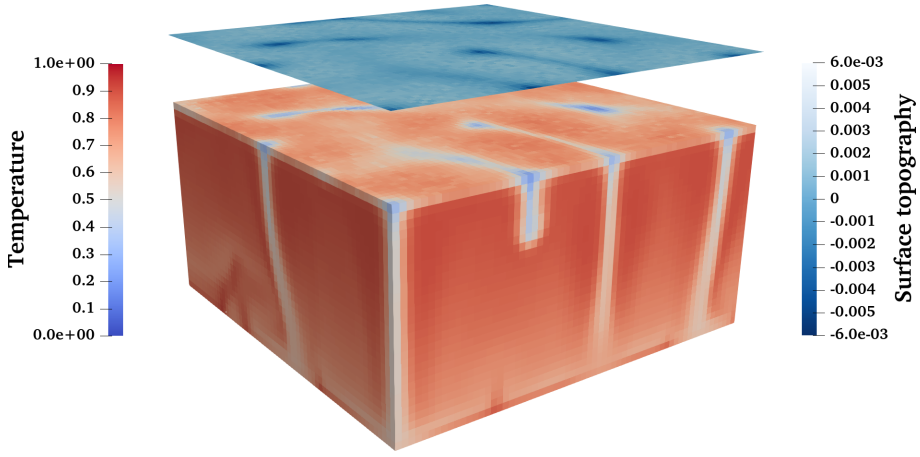

Figure 20. A nondimensional constant-viscosity Cartesian box test model ($\mathrm{Ra} = 10^7$). A sticky air discretisation with a viscosity contrast of $\eta_{\mathrm{air}}/\eta_{\mathrm{rock}} = 10^{-3}$ is used. A relatively coarse resolution of $64 \times 64 \times 32$ is used to better highlight the effect of low volumetric marker densities on topography.

cant reductions in volumetric marker count and associated computational cost. The bilinear method offers the best computational performance, consistently accounting for less than 3 % of total runtime, and is straightforward to implement, though it introduces some artificial numerical diffusion due to periodic global reinitialisation. The integral method produces topography of comparable quality but at greater computational cost, due primarily to the $O(\log n)$ surface lookups required during marker unmixing and the overhead of slope limiting. Both methods can be used to obtain volume fractions for coupling the surface location to the density and viscosity fields used by the Stokes solver (ρ, η coupling), and this coupling is found to be essential for smooth and accurate surface tracking results.

6.2 Discussion of key results

6.2.1 Novelty and relation to prior work

The bilinear method presented here is not entirely new: a closely related approach is already used in the regional-scale code LaMEM (Kaus et al., 2016), and a simplified 2D version was employed by Duretz et al. (2016) to test free surface discretisation methods. A related approach using a Delaunay-triangulated surface embedded in an Eulerian mesh was previously implemented in the code DOUAR (Braun et al., 2008). The novelty of this study lies in the extension and thorough benchmarking of both methods to 3D spherical geometry within StagYY, the development of the slope-limiting stabilisation algorithm for the integral method, the introduction of ρ, η coupling to the Stokes solver, and the demonstration that volumetric marker density can be substantially reduced without loss of topographic accuracy.

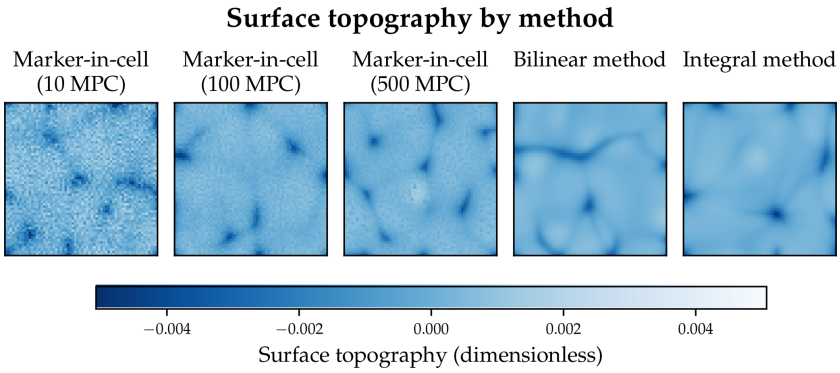

Figure 21. Surface topography generated by the marker-in-cell method at various marker densities, compared against the bilinear and integral Lagrangian tracking methods at a density of 100 markers per cell. Both Lagrangian approaches yield smoother topography, achieving a quality comparable or exceeding the high-density marker-in-cell results.

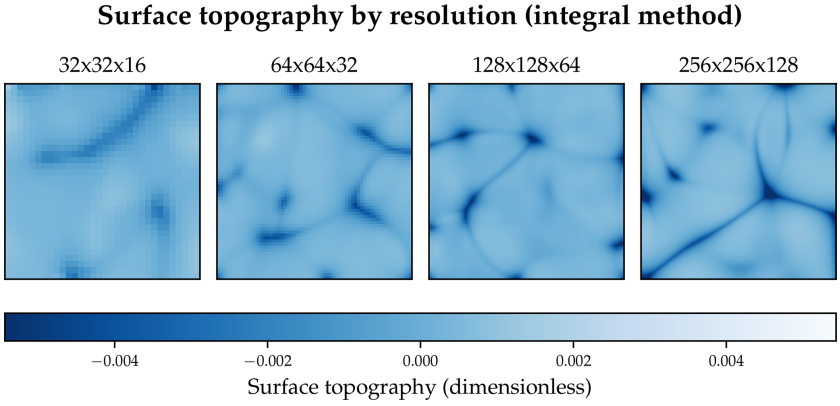

Figure 22. 3D surface topography tracked using the integral method at different resolutions. Smooth topography is consistent across length scales, with higher resolutions resolving greater detail around downwelling regions. Similar results are obtained with the bilinear method.

6.2.2 Implementation

The implementation of Lagrangian surface markers, especially the bilinear method, is relatively straightforward. Since Lagrangian markers already exist in StagYY, as they do for many other numerical geodynamic modelling codes, existing routines for manipulating these markers were easily adapted for the new surface markers. Reconstructing the surface using the bilinear representation was similarly straightforward and could likely be implemented in other numerical geodynamic codes with minimal effort.

In contrast, the integral method is significantly more complex to implement, particularly due to the need for surface marker reinitialisation and smoothing. The 3D implementation requires a significant amount of code for performing complex operations, such as computing the Delaunay triangulation and determining cellwise volume fractions. Developers considering these methods should carefully weigh the cost-benefit trade-offs, especially given the integral method's performance limitations discussed below.

6.2.3 Bilinear method

The bilinear method was able to produce high-resolution topography and demonstrated the best computational performance among the methods tested. Key surface marker operations, such as marker advection and surface reconstruction, are performed in linear time $O(n)$ with respect to the number of surface markers. Once surface locations are interpolated to nodes, determining the surface height at a specific point can be achieved in constant time $O(1)$, which significantly benefits operations requiring frequent surface height lookups, such as marker unmixing.

Global reinitialisation introduces implicit numerical diffusion that stabilises the marker chain and avoids the need for explicit smoothing, improving performance. However, this diffusion can gradually reduce the accuracy of the surface representation, potentially at the cost of precision when considering sub-grid scale features.

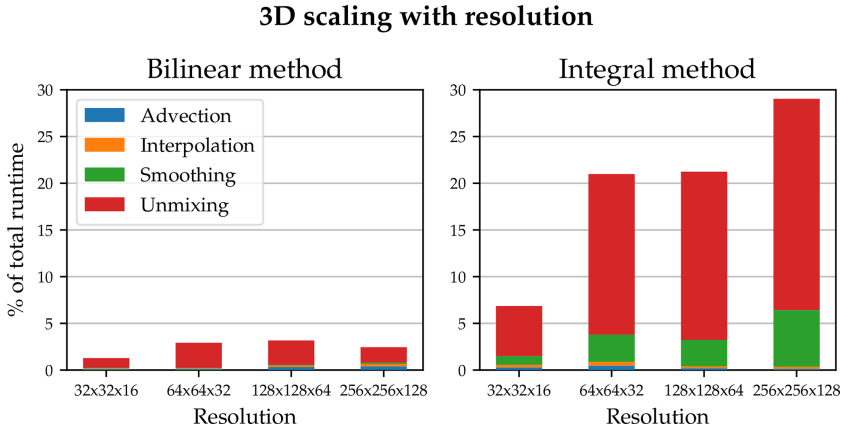

Figure 23. Performance of each method with resolution as a percentage of total runtime. The bilinear method performs significantly better than the integral method in 3D, with marker unmixing the dominant cost in both cases.

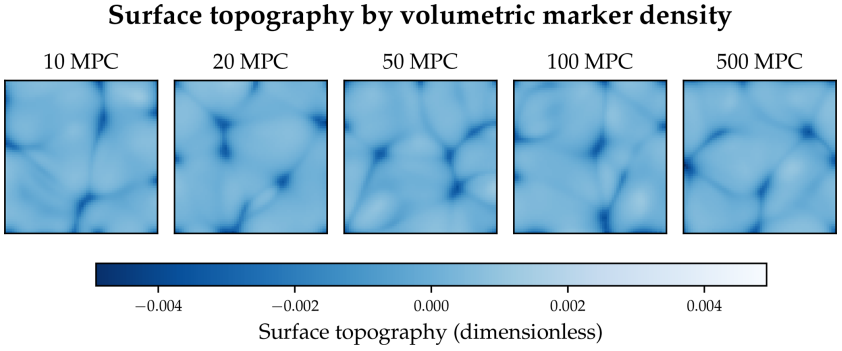

Figure 24. Surface topography as determined using the integral method with varying markers per cell. No qualitative differences are seen in surface topography reconstruction. A similar result is obtained for the bilinear method.

6.2.4 Integral method

The integral method can produce similarly high-quality topography to the bilinear method but is less performant. Its main advantage is the absence of artificial numerical diffusion, as surface markers move independently of the Eulerian grid and are only adjusted where needed. In 3D, Delaunay triangulation offers a natural surface representation suitable for visualising fine-scale features, tracking sea level, and potentially integrating surface process models. However, the method's reliance on several $O(n \log n)$ operations such as marker sorting and reinitialisation limits its efficiency, and slope limiting adds further cost that cannot be solved in linear time. The largest performance bottleneck in both 2D and 3D is marker unmixing, which requires $O(\log n)$ per marker, leading to $O(m \log n)$ overall with m markers. This cost remains significant even after optimisation of the unmixing algorithm.

6.2.5 Comparison with the marker-in-cell method

Both implemented methods improve upon the existing marker-in-cell method in two key respects. Firstly, they bypass the need for high volumetric marker densities to obtain high-quality surface topography, allowing for significantly decreased computational effort and memory consumption. Secondly, the implemented methods are more accurate, reducing topographic noise on all scales. The trade-off is a modest increase in computational effort for the surface tracking operations themselves, which in all tests remained well below the cost of volumetric marker advection at the densities required by the marker-in-cell method.

6.2.6 Necessity of density and viscosity coupling

A key result of this study is the necessity of coupling the location of the surface to the underlying physical model through the density and viscosity fields, referred to as ρ, η coupling in Sect. 4.2. This coupling is essential to avoid topographic noise that arises from determining density and viscosity based on a finite number of discrete markers with a random initial distribution. The requirement for density coupling was also noted by Duretz et al. (2016); the methods presented here extend this further with application to the viscosity field as well.

6.3 Applications and future directions

6.3.1 Surface processes

A potential future research direction would be the direct implementation of surface processes into StagYY. While the diffusion and slope limiting methods presented here can be considered very simple representations of erosion and sedimentation, a direct representation of the surface coupled to the underlying model theoretically allows for fully coupled modelling of surface processes on the global scale for the first time. Prior work coupling geodynamic codes to surface process models has been carried out at regional scales (e.g. Simpson, 2004; Burov and Toussaint, 2007; Braun, 2010; Bahadori et al., 2022; Wolf et al., 2025), but global-scale coupling remains an open problem. A natural direction would be the incorporation of a surface process model such as Badlands (Salles, 2016) directly into StagYY, modifying the surface at each timestep in a way that feeds back into the underlying model through the existing marker unmixing and ρ, η coupling machinery.

6.3.2 Alternative Stokes solvers

While the results presented here have all employed the sticky air method to approximate the free surface boundary condition, the direct surface representation is compatible with alternative solvers. The staircase discretisation of Duretz et al. (2016) eliminates the need for a sticky air layer entirely by applying alternative boundary conditions to cells near the surface, requiring only a method to determine whether a cell centre lies above or below the surface, which is readily provided by a marker chain or mesh. The variational Stokes discretisation of Larionov et al. (2017) is also compatible, requiring volume fraction functions that are already computed as part of the ρ, η coupling described in Sect. 4.2.

6.3.3 Sea level modelling

Accurately tracking the evolution of sea levels over time is a key capability of the direct surface tracking methods presented here. Sea level tracking is a necessary step toward coupling geodynamic models with climate and biological evolution models: climate models such as FOAM (Tobis et al., 1997) and biological models like Gen3sis (Hagen et al., 2021) require information about continental configurations that can be more accurately provided by dynamically computing sea level changes from an initial ocean volume. As StagYY already includes functionality to estimate volatile outgassing rates including water, it would theoretically be possible to incorporate a changing sea volume into the model, enabling more integrated models linking deep Earth dynamics with climate, landscape, and biosphere evolution.

Appendix A: Spherical annulus benchmark

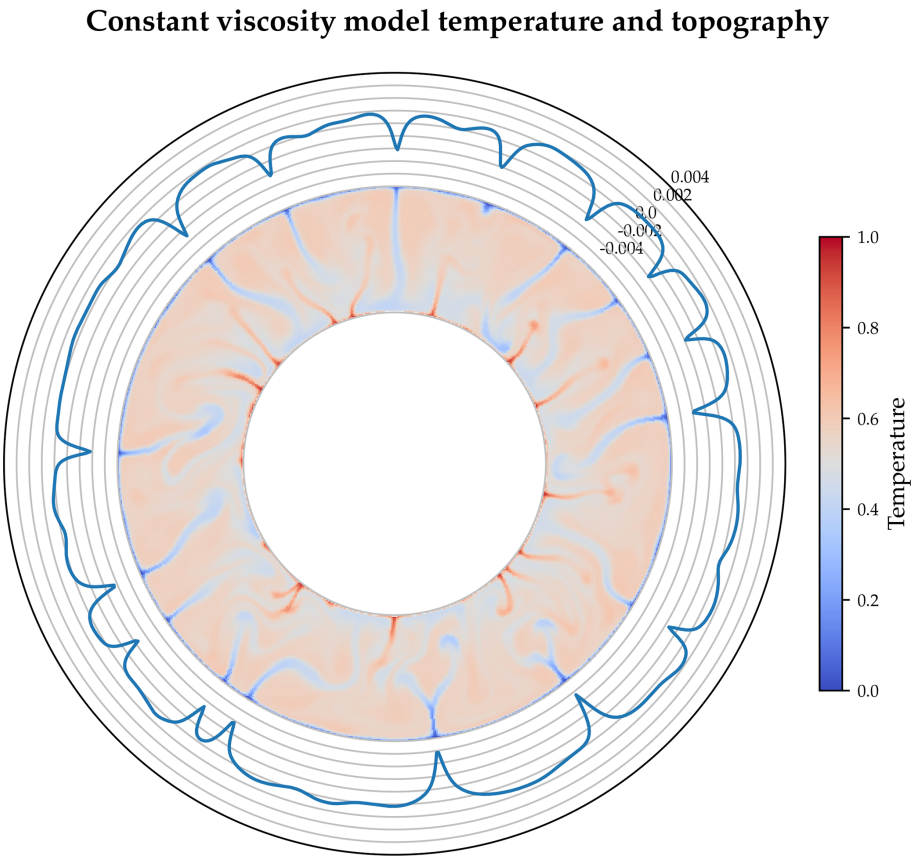

Figure A1. Example of topography generated using a simple constant viscosity model in 2D spherical geometry. The underlying temperature field shows the location of various plumes and downwellings, which are captured in the surface topography generated using the bilinear method. This model forms the basis of the 2D scaling tests performed.

Code and data availability. The code StagYYFS, a testbed code based on StagYY, may be used to produce the benchmark results and figures used in this paper. It is archived on Zenodo under the GPLv3 license under https://doi.org/10.5281/zenodo.18281860 (Tackley and Gray, 2025).

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 and editing. Taras Gerya: conceptual input; manuscript review.

Competing interests. The contact author has declared that none of the authors has any competing interests.

Disclaimer. Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims made in the text, published maps, institutional affiliations, or any other geographical representation in this paper. The authors bear the ultimate responsibility for providing appropriate place names. Views expressed in the text are those of the authors and do not necessarily reflect the views of the publisher.

Acknowledgements. We thank the two anonymous reviewers for constructive reviews that improved this manuscript.

Financial support. This work was supported by Swiss National Science Foundation (grant no. 192296).

Review statement. This paper was edited by Boris Kaus and reviewed by two anonymous referees.

References

Bahadori, A., Holt, W. E., Feng, R., Austermann, J., Loughney, K. M., Salles, T., Moresi, L., Beucher, R., Lu, N., Flesch, L. M., Calvelage, C. M., Rasbury, E. T., Davis, D. M., Potochnik, A. R., Ward, W. B., Hatton, K., Haq, S. S. B., Smiley, T. M., Wooton, K. M., and Badgley, C.: Coupled influence of tectonics, climate, and surface processes on landscape evolution in southwestern North America, Nat. Commun., 13, 4437, https://doi.org/10.1038/s41467-022-31903-2, 2022.

Braun, J.: The many surface expressions of mantle dynamics, Nat. Geosci., 3, 825–833, https://doi.org/10.1038/ngeo1020, 2010.

Braun, J., Thieulot, C., Fullsack, P., DeKool, M., Beaumont, C., and Huismans, R.: DOUAR: A new three-dimensional creeping flow numerical model for the solution of geological problems, Phys. Earth Planet. In., 171, 76–91, https://doi.org/10.1016/j.pepi.2008.05.003, 2008.

Burov, E. and Toussaint, G.: Surface processes and tectonics: Forcing of continental subduction and deep processes, Global Planet. Change, 58, 141–164, 2007.

Coltice, N., Husson, L., Faccenna, C., and Arnould, M.: What drives tectonic plates?, Sci. Adv., 5, eaax4295, https://doi.org/10.1126/sciadv.aax4295, 2019.

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, Geophys. J. Int., 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, Geophys. Res. Lett., 39, https://doi.org/10.1029/2011GL050046, 2012b.

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

Gerya, T., Duretz, T., and Ludovic, R.: New continuity-based velocity interpolation scheme for staggered grids, EGU General Assembly 2021, online, 19–30 Apr 2021, EGU21-15308, https://doi.org/10.5194/egusphere-egu21-15308, 2021.

Gerya, T. V.: Introduction to Numerical Geodynamic Modelling, Cambridge University Press, ISBN 978-1-107-14314-2, 2019.

Gerya, T. V. and Yuen, D. A.: Robust characteristics method for modelling multiphase visco-elasto-plastic thermomechanical problems, Phys. Earth Planet. In., 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, 2015.

Hagen, O., Flück, B., Fopp, F., Cabral, J. S., Hartig, F., Pontarp, M., Rangel, T. F., and Pellissier, L.: gen3sis: A general engine for eco-evolutionary simulations of the processes that shape Earth's biodiversity, PLOS Biol., 19, e3001340, https://doi.org/10.1371/journal.pbio.3001340, 2021.

Joe, B.: GEOMPACK – a software package for the generation of meshes using geometric algorithms, Adv. Eng. Softw. Workst., 13, 325–331, https://doi.org/10.1016/0961-3552(91)90036-4, 1991.

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, 978–983, John von Neumann Institute for Computing (NIC), NIC Series, https://juser.fz-juelich.de/record/507751/files/nic_2016_kaus.pdf (last access: 26 August 2026), 2016.

Kreyszig, E.: Advanced Engineering Mathematics, John Wiley & Sons, ISBN 978-0-470-45836-5, 2010.

Larionov, E., Batty, C., and Bridson, R.: Variational stokes: a unified pressure-viscosity solver for accurate viscous liquids, ACM T. Graphic., 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, J. Fluid Mech., 233, 299–328, https://doi.org/10.1017/S0022112091000496, 1991.

O'Rourke, J., Chien, C.-B., Olson, T., and Naddor, D.: A new linear algorithm for intersecting convex polygons, Comput. Vision Graph., 19, 384–391, https://doi.org/10.1016/0146-664X(82)90023-5, 1982.

Patankar, S. V.: Numerical Heat Transfer and Fluid Flow, Hemisphere Publishing Corporation, New York, ISBN 9780891165224, 1980.

Phongthanapanich, S. and Dechaumphai, P.: Adaptive Delaunay triangulation with object-oriented programming for crack propagation analysis, Finite Elem. Anal. Des., 40, 1753–1771, https://doi.org/10.1016/j.finel.2004.01.002, 2004.

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

Renka, R. J.: Algorithm 772: STRIPACK: Delaunay triangulation and Voronoi diagram on the surface of a sphere, ACM T. Math. Software, 23, 416–434, https://doi.org/10.1145/275323.275329, 1997.

Salles, T.: Badlands: A parallel basin and landscape dynamics model, SoftwareX, 5, 195–202, https://doi.org/10.1016/j.softx.2016.08.005, 2016.

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, Phys. Earth Planet. In., 171, 198–223, https://doi.org/10.1016/j.pepi.2008.06.028, 2008.

Simpson, G.: Dynamic interactions between erosion, deposition, and three-dimensional deformation in compressional fold belt settings, J. Geophys. Res., 109, https://doi.org/10.1029/2003JF000111, 2004.

Stern, R. J. and Gerya, T. V.: Chapter 13 – Co-Evolution of Life and Plate Tectonics: The Biogeodynamic Perspective on the Mesoproterozoic-Neoproterozoic Transitions, in: Dynamics of Plate Tectonics and Mantle Convection, edited by: Duarte, J. C., 295–319, Elsevier, ISBN 978-0-323-85733-8, https://doi.org/10.1016/B978-0-323-85733-8.00013-5, 2023.

Tackley, P. J.: Effects of strongly temperature-dependent viscosity on time-dependent, three-dimensional models of mantle convection, Geophys. Res. Lett., 20, 2187–2190, https://doi.org/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, Phys. Earth Planet. In., 171, 7–18, https://doi.org/10.1016/j.pepi.2008.08.005, 2008.

Tackley, P. and Gray, T.: A testbed for free-surface methods in geodynamical simulations using the finite volume (staggered-grid finite difference) discretization, Zenodo [code, data set], https://doi.org/10.5281/zenodo.18281860, 2025.

Tackley, P. J. and King, S. D.: Testing the tracer ratio method for modeling active compositional fields in mantle convection simulations, Geochem. Geophy. Geosy., 4, https://doi.org/10.1029/2001GC000214, 2003.

Tezduyar, T. E.: Interface-tracking and interface-capturing techniques for finite element computation of moving boundaries and interfaces, Comput. Method. Appl. M., 195, 2983–3000, https://doi.org/10.1016/j.cma.2004.09.018, 2006.

Tobis, M., Schafer, C., Foster, I., Jacob, R., and Anderson, J.: FOAM: Expanding the Horizons of Climate Modeling, in: SC '97: Proceedings of the 1997 ACM/IEEE Conference on Supercomputing, 27–27, https://doi.org/10.1145/509593.509620, 1997.

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.

Wolf, S. G., Huismans, R. S., and Braun, J.: Tectonics and Surface Processes During Collisional Orogenesis – Exploring the Parameter Space of the Beaumont Number, J. Geophys. Res.-Sol. Ea., 130, https://doi.org/10.1029/2024jb030439, 2025.

Zerkle, A. L.: Biogeodynamics: bridging the gap between surface and deep Earth processes, Philos. T. R. Soc. A, 376, 20170401, https://doi.org/10.1098/rsta.2017.0401, 2018.