

Automatic tuning of iterative pseudo-transient solvers for ~~modelling~~ modeling the deformation of heterogeneous media

Thibault Duret¹, Albert de Montserrat², Rubén Sevilla³, Ludovic Räss⁴, Ivan Utkin⁵, and Arne Spang⁶

¹Institut für Geowissenschaften, Goethe-Universität Frankfurt, Frankfurt, Germany

²Geophysical Fluid Dynamics, Institute of Geophysics, ETH Zurich, Zurich, Switzerland

³Zienkiewicz Centre for Computational Engineering, Faculty of Science and Engineering, Swansea University, Wales, UK

⁴Institute of Earth sciences, University of Lausanne, Lausanne, Switzerland

⁵Laboratory of Hydraulics, Hydrology and Glaciology (VAW), ETH Zurich, Zurich, Switzerland

⁶Bayerisches Geoinstitut, Universität Bayreuth, Bayreuth, Germany

Correspondence: Thibault Duret (duret@uni-frankfurt.de)

Abstract.

Geodynamic modeling has become a crucial tool for investigating the dynamics of Earth deformation across various scales. Such simulations often involve solving mechanical problems with significant material heterogeneities (e.g. ~~strong~~ viscosity contrasts) under nearly incompressible conditions. Recent advancements have enabled the development of iterative solvers based on Dynamic Relaxation or Pseudo-Transient schemes, which require minimal global communication and exhibit quasi-linear scaling on GPU and supercomputing architectures. These solvers incorporate automatic tuning of iterative parameters, including pseudo-time steps and damping coefficients, based on spectral estimates of the discrete operators, ensuring both robust and rapid convergence. We demonstrate the effectiveness of this approach on ~~problems discretized using discretized~~ problems with finite-difference and face-centered finite volume methods, including heterogeneous incompressible Stokes flows. Moreover, the relative algorithmic simplicity of DR-based methods allows ~~for~~ straightforward extensions to compressible flow, multiphase flow, and nonlinear constitutive laws, opening promising avenues for large-scale, high-resolution simulations of geoscientific problems.

1 Introduction

Geodynamic ~~modelling~~ modeling has become an indispensable tool for simulating tectonic processes at various scales, ranging from the rock sample (e.g. ?) and outcrop levels (e.g. ?) to regional (?), global (e.g. ?), and planetary scales (e.g. ?). Geodynamic ~~modelling tools aim at solving conservation equations for~~ modeling tools aim to solve the conservation equations of linear momentum, mass, and energy to produce full-field solutions for velocity, pressure, and temperature fields in space and time (?).

While 2D geodynamic simulations can effectively leverage solvers that ~~utilize~~ employ sparse-direct or direct-iterative schemes (e.g. ???), 3D simulations generally require iterative solvers as direct methods become prohibitively expensive. Most

approaches in this context rely on multigrid solvers (e.g. ~~???????~~), spectral methods (e.g. ~~?~~), or pseudo-transient integration techniques (~~??~~) to achieve efficient solutions.

In geodynamic models, the viscosity field may exhibit smooth (i.e. thermal variations) and sharp (i.e. material interfaces) spatial variations of several orders of magnitude. This peculiarity makes the iterative solution procedure of the near-incompressible and incompressible Stokes equation particularly challenging (e.g. ~~????~~).

Iterative solvers based on pseudo-transient integration or Dynamic Relaxation (DR) have been employed to solve compressible and incompressible geomechanical and geodynamical problems (~~???~~) (e.g. ~~???~~). These methods exhibit interesting parallel scaling properties on hardware accelerators such as ~~graphics processing units~~ Graphics Processing Units (GPU) and have therefore recently been applied to study incompressible Stokes flow and multi-physics coupled problems (~~?????~~) (e.g. ~~?????~~). These solvers implement coupled solving strategies that involve iterative updates of both the velocity and pressure fields within a single iteration loop. Despite their algorithmic simplicity, their application to heterogeneous mechanical problems remains challenging (e.g. ~~?~~). One strategy relies on smoothing sharp contrasts in material properties; however, this approach requires the use of a very fine spatial resolution to match exact flow solutions (e.g. ~~??~~). In practice, Stokes solvers are typically coupled with Marker-In-Cell or Level-Set approaches (e.g. ~~????~~) to handle large viscous deformations and to represent heterogeneous materials. In the latter case, the degree to which discontinuities in material properties are smoothed depends on the spatial extent (or support) of the interpolant ~~'s~~-basis functions (e.g. ~~?~~). The degree of smoothing is therefore not defined ~~in-an-ad hoc manner~~ ~~ad hoc~~ and must be based on a narrow ~~interpolant interpolation~~ basis to best match ~~the~~ exact flow solution (~~?~~). Moreover, current geodynamic pseudo-transient solvers rely on either empirical tuning of iterative parameters (e.g. ~~?~~) or exact derivation of optimal parameters based on simplified model ~~configuration~~ (~~??~~) ~~configurations~~ (e.g. ~~??~~). This situation hinders their applicability in solving "real-life" geodynamic problems, as optimal parameters may adapt to the evolving internal model state and material configuration. Exploring dynamic strategies to select optimal iterative parameters is a motivation to overcome current limitations.

In this study, we aim to solve quasi-steady mechanical problems with large and sharp variations in material properties using ~~DR-type-of-like~~ iterative solvers and automatic tuning of iterative parameters. We present a decoupled solution strategy that combines Schur-complement reduction, Powell–Hestenes iterations, and DR-type of iterative solvers.

2 Mechanical equilibrium

The equations governing quasi-static mechanical equilibrium in a classical continuum medium may be expressed as:

$$\nabla \cdot \boldsymbol{\tau} - \nabla P - \mathbf{fb} = \mathbf{0}, \quad (1)$$

$$\nabla \cdot \mathbf{v} + \frac{1}{K} \frac{dP}{dt} = 0. \quad (2)$$

where ~~$\boldsymbol{\tau}$~~ is the deviatoric stress tensor, P is pressure, and \mathbf{v} is the velocity vector. The term ~~$\mathbf{f} = -\rho \mathbf{g}$~~ ~~$\mathbf{b} = -\rho \mathbf{g}$~~ is a forcing term, where ρ , \mathbf{g} and K correspond to density, gravitational acceleration, and bulk modulus, respectively. The incompressible

limit is reached if $K \rightarrow \infty$. The isotropic viscous constitutive relationship takes the form of:

$$\boldsymbol{\tau} = 2\eta \dot{\boldsymbol{\epsilon}} = 2\eta \left[\mathbf{D} - \frac{1}{3}(\nabla \cdot \mathbf{v})\mathbf{I} \right], \quad (3)$$

where $\dot{\boldsymbol{\epsilon}}$ is the deviatoric strain rate tensor, $\mathbf{D} = 1/2 (\nabla \mathbf{v} + (\nabla \mathbf{v})^T)$, \mathbf{I} is the identity matrix and η is the dynamic viscosity. In the following, the Stokes equations will be solved for velocity and pressure fields over a domain Π , accounting for Neumann boundary conditions ($\boldsymbol{\sigma} \cdot \mathbf{n} = \mathbf{T}_N$) on the boundary $\partial\Pi_N$, and Dirichlet boundary conditions ($\mathbf{v} = \mathbf{v}_D$) on the boundary $\partial\Pi_D$.

3 Discretisations/Discretizations

We consider two different types of numerical ~~discretisation~~-discretization of the Stokes equations, the staggered-grid Finite-Difference (FD) and Face-Centered Finite Volume (FCFV) Method. With either ~~discretisation~~discretization, both the components of the velocity vector and the pressure will constitute the primitives, which we aim to solve for.

3.1 Staggered-Grid Finite-Differences Method

The staggered-grid finite-difference is a popular approach to ~~discretising~~-discretizing the Stokes equation originally proposed by ?. This scheme is constructed using conservative second-order finite difference and a staggered arrangement of both velocity and pressure (Fig 1a, b). The strong form of the Stokes equations (Eq. 2) is ~~discretised~~-discretized using central finite differences:

$$r_{\text{mom}} = \hat{\nabla} \cdot \boldsymbol{\tau} - \hat{\nabla} P - \mathbf{fb}, \quad (4)$$

$$r_{\text{cont}} = -\frac{1}{K} \frac{\hat{\Delta} P}{\hat{\Delta} t} - \hat{\nabla} \cdot \mathbf{v}, \quad (5)$$

$$\boldsymbol{\tau} = 2\eta \left[\hat{\mathbf{D}} - \frac{1}{3}(\hat{\nabla} \cdot \mathbf{v})\mathbf{I} \right], \quad (6)$$

where the hat symbol denotes the discrete representation of the differential operators. The primitive solution fields \mathbf{v} and p are obtained by ~~minimising~~-minimizing the momentum and continuity residuals (r_{mom} and r_{cont}). The staggered FD scheme represents, to our knowledge, the most computationally efficient stencil that satisfies the Inf-Sup stability condition ~~on~~-in regular meshes (?). For flows with smooth viscosity variations, FD achieves second-order accuracy in the L1 norm, while in the presence of sharp material heterogeneities, its accuracy reduces to first-order (?). Geometrically, FD has been successfully implemented in Cartesian and polar coordinates ~~(?)~~(e.g. ?), spherical configurations ~~(??)~~(e.g. ??), and on adaptively refined meshes ~~(??)~~(e.g. ??). However, its application ~~on~~-in non-orthogonal meshes remains ~~uncommon~~-rare (e.g. ?) due to the need for extended stencils and the associated increase in algorithmic complexity. Consequently, the treatment of free-surface boundary conditions remains a notable challenge (e.g. ??).

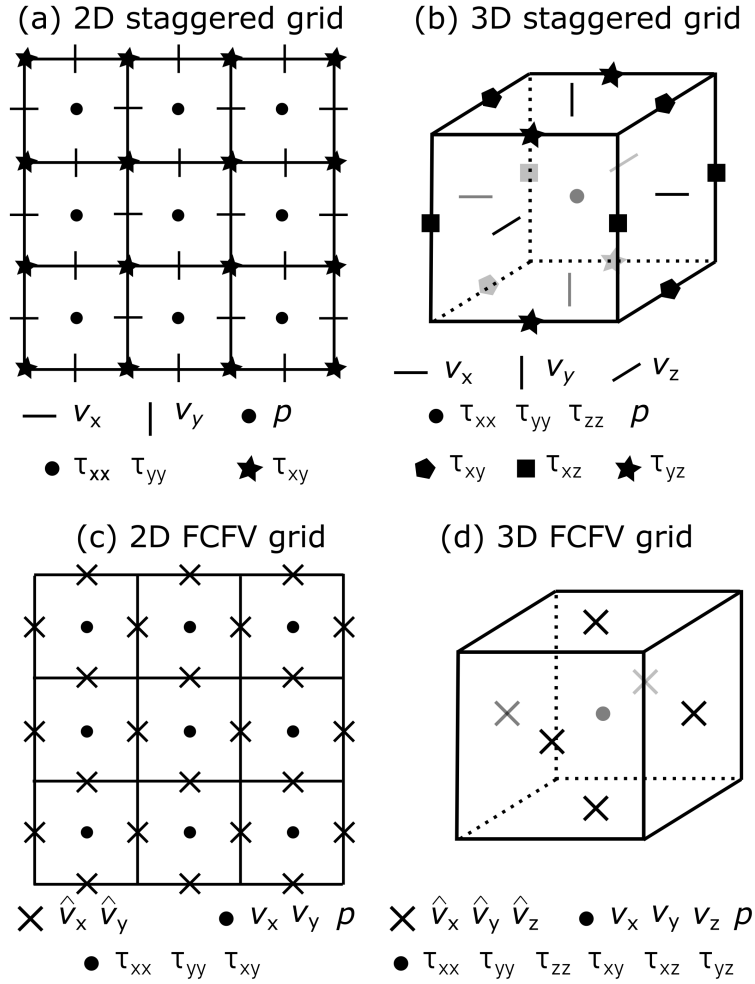


Figure 1. Grids employed for the [discretisation-discretization](#) of mechanical problems. (a) 2D FD staggered grid, (b) 3D staggered grid, (c) 2D FCFV grid and (d) 3D FCFV grid.

3.2 Face-Centered Finite Volume Method (FCFV)

The Face-Centered Finite-Volume (FCFV) method is an emerging [discretisation-discretization](#) scheme rooted in the principles of the [hybridisable-hybridizable](#) discontinuous Galerkin (HDG) method (?). This approach has been successfully employed to achieve full-field solutions for the Stokes equations [across-in](#) various types of viscosity [distributionsdistribution](#), including constant (?), discontinuous (?), and smoothly varying fields (?). The FCFV scheme satisfies the Inf-Sup stability condition, making it particularly suitable for solving the Stokes equations. The discrete system of equations is derived from the weak formulation of the Stokes problem, using a constant degree of approximation for both velocity and pressure fields:

$$85 \quad \mathbf{r}_{\bar{i}}^{\text{mom}} = \sum_{e=1}^{n_{\text{elem}}} \Gamma_{\bar{i}} \left[\left(\left[\boldsymbol{\tau}_e^T \mathbf{n}_{\bar{i}} \right] + \left[P_{\bar{e}} \right] \mathbf{I} \right) \mathbf{n}_{\bar{i}} + 2s \left(\left\{ \left\{ \mathbf{v}_{\bar{e}} \right\} \right\} - \hat{\mathbf{v}}_{\bar{i}} \right) - \chi \mathbf{T}_N \right], \quad (7)$$

$$r_{\text{cont}} = - \sum_{i=1}^{n_{\text{fac}}} \Gamma_i \left(\frac{1}{K} \frac{\hat{\Delta} P}{\hat{\Delta} t} + \mathbf{n}_i \cdot \hat{\mathbf{v}}_i \right), \quad (8)$$

$$\boldsymbol{\tau} = 2\eta \left[\sum_{i=1}^{n_{\text{fac}}} \frac{\Gamma_i}{2\Omega} (\mathbf{n}_i \otimes \hat{\mathbf{v}}_i + \hat{\mathbf{v}}_i \otimes \mathbf{n}_i) - \frac{1}{3} \mathbf{n}_i \cdot \hat{\mathbf{v}}_i \mathbf{I} \right], \quad (9)$$

$$\mathbf{v}_{\bar{e}} = \frac{1}{a} \left(\mathbf{b}\Omega + \sum_{i=1}^{n_{\text{fac}}} \Gamma_i s \hat{\mathbf{v}}_i \right), \quad (10)$$

The primitive variables, namely the hybrid velocity vector ($\hat{\mathbf{v}}$) and pressure (P), are obtained by **minimising** the momentum and continuity residuals ($r_{\text{mom}}, r_{\text{cont}}$). **The FCFV discretisation requires the outward-pointing** The hybrid velocity is an independent variable defined on element faces, different from the element-wise velocity. It serves as a Lagrange multiplier-type trace variable that enforces weak continuity of the velocity field across elements. The symbol n_{fac} corresponds to the number of faces of an element. $[\![\cdot]\!] is the jump operator defined as $[\![\cdot]\!] = \odot_{\text{R}} - \odot_{\text{L}}$, where \odot denotes quantities defined on elements located to the right (R) and left (L) of an interface. $\{\{ \cdot \}\}$ is the average operator defined as $\{\{ \cdot \}\} = \frac{1}{2} (\odot_{\text{R}} + \odot_{\text{L}})$. The$ discretization requires the outward-pointing normal to each face (\mathbf{n}), the face area (Γ), the volume of the element (Ω), the **stabilisation** stabilization parameter (s), the factor a ($a = \sum_{i=1}^{n_{\text{fac}}} \Gamma_i s a = \sum_{i=1}^{n_{\text{fac}}} \Gamma_i s$), the traction vector (\mathbf{T}_N), and the function χ that indicates the presence of **an**-applied traction (0 or 1). The FCFV scheme requires the storage of the hybrid velocity vector (**$\hat{\mathbf{v}}$**) and momentum residual at the midpoint of each element face, while the pressure and stress tensor components, the deviatoric stress tensor and the continuity residual are stored in the centroids of the elements (Fig 1c, d). The **discretisation** discretization considered here yields a first-order approximation of velocity and pressure fields (?), nevertheless second-order solutions can be constructed without altering the sparsity of the **discretisation** discretization (??). In this study, we implement the FCFV approach using structured quadrangular meshes, although the method is versatile and supports arbitrary element geometries, exhibiting low sensitivity to mesh distortions (?). Additionally, FCFV can naturally handle internal boundaries and traction boundary conditions, such as free surfaces. Due to these features, the FCFV method is becoming increasingly popular in the geodynamics community, as highlighted by ?.

4 Dynamic Relaxation with automatic tuning

In this section, we discuss the principles behind the Dynamic Relaxation method and illustrate its basic properties using a simple 1D Poisson example.

4.1 Principle

The dynamic relaxation method is designed to iteratively solve both linear and **non-linear problems** through successive iterations nonlinear problems. The iterative procedure resembles an explicit time integration scheme, which is why it is also often referred to as a pseudo-transient integration method. As in classical iterative procedures, the DR method iteratively refines the unknown \mathbf{u} until the magnitude of the residual of the partial differential operator $\mathbf{r}(\mathbf{u})$ drops below a given tolerance. Once this criterion

is met, \mathbf{u} is considered a solution satisfying $\mathbf{r}(\mathbf{u}) = \mathbf{0}$. Note that the function $\mathbf{r}(\mathbf{u})$ may include physical transient terms; thus,
 115 this method is not limited to solving quasi-static problems (e.g. ?). The algorithm is formulated by introducing the first and
 second-order derivative of the solution \mathbf{u} in pseudo-time (?):

$$\mathbf{M} \frac{d^2 \mathbf{u}}{d\bar{t}^2} + \mathbf{M}c \frac{d\mathbf{u}}{d\bar{t}} = \mathbf{r}(\mathbf{u}) \quad (11)$$

where \bar{t} refers to pseudo-time, c is a damping parameter, and \mathbf{M} is a preconditioner. Following previous work (e.g. ??), the
 second derivative is taken as the central derivative of the first derivative and the first derivative is taken as the average of the
 120 new and old values:

$$\frac{d^2 \mathbf{u}}{d\bar{t}^2} \approx \frac{1}{\Delta\bar{t}} \left(\frac{\Delta \mathbf{u}^{\text{new}}}{\Delta\bar{t}} - \frac{\Delta \mathbf{u}^{\text{old}}}{\Delta\bar{t}} \right) \quad (12)$$

$$\frac{d\mathbf{u}}{d\bar{t}} \approx \frac{1}{2} \left(\frac{\Delta \mathbf{u}^{\text{new}}}{\Delta\bar{t}} + \frac{\Delta \mathbf{u}^{\text{old}}}{\Delta\bar{t}} \right). \quad (13)$$

This leads to update rules for the rate of change of \mathbf{u} , and \mathbf{u} itself:

$$\frac{\Delta \mathbf{u}^{\text{new}}}{\Delta\bar{t}} = b \frac{\Delta \mathbf{u}^{\text{old}}}{\Delta\bar{t}} + a \mathbf{M}^{-1} \mathbf{r}(\mathbf{u}) \quad (14)$$

$$125 \quad \mathbf{u}^{\text{new}} = \mathbf{u}^{\text{old}} + \Delta\bar{t} \frac{\Delta \mathbf{u}^{\text{new}}}{\Delta\bar{t}} \quad (15)$$

where $a = 2\Delta\bar{t}/(2 + c\Delta\bar{t})$ and $b = (2 - c\Delta\bar{t})/(2 + c\Delta\bar{t})$. Alternatively, the DR algorithm can be expressed in analogy to the
 conjugate gradient (CG) method (?), using the variable $\mathbf{p} = \Delta \mathbf{u}/(\beta\Delta\bar{t})$:

$$\mathbf{p}^{\text{new}} = \beta \mathbf{p}^{\text{old}} + \mathbf{M}^{-1} \mathbf{r}(\mathbf{u}) \quad (16)$$

$$\mathbf{u}^{\text{new}} = \mathbf{u}^{\text{old}} + \alpha \mathbf{p}^{\text{new}} \quad (17)$$

130 where \mathbf{p} is analogous to the search direction, $\alpha = 2\Delta\bar{t}^2/(2 + c\Delta\bar{t})$, $\alpha = a\Delta\bar{t}$ and $\beta = b$. The main difference is that the CG
 method involves different expressions for the parameters α and β , which need to be reevaluated at each iteration, requiring
 computationally expensive reduction operations (e.g. \cdot vector norms).

In either case, two iteration parameters need to be determined as a function of the minimum (λ_{\min}) and maximum (λ_{\max})
 eigenvalues of the discrete problem (??):

$$135 \quad \Delta\bar{t} = c_{\text{CFL}} \frac{2}{\sqrt{\lambda_{\max}}} \quad (18)$$

$$c = c_{\text{damp}} 2\sqrt{\lambda_{\min}} \quad (19)$$

where $c_{\text{CFL}} < 1$ and $1/2 < c_{\text{damp}} < 1$. c_{CFL} and c_{damp} are scalar adjustment parameters. Whereas c_{CFL} is strictly bounded
 ($c_{\text{CFL}} < 1$), no analogous upper bound applies to c_{damp} . In our experiments, values in the range $\approx [0.5, 1]$ were found to

provide robust performance. For optimal convergence, at least in the case of linear problems, c_{CFL} should be chosen as large as possible (e.g. $c_{\text{CFL}} = 0.999$). In theory (e.g. ?), the optimal value of c_{damp} is 1.0. Nevertheless, we empirically observed that slightly lower values may further enhance convergence.

The maximum eigenvalue is found using the Gershgorin circle theorem:

$$\lambda_{\max} = \max_i \left(\left| \frac{\partial \mathbf{r}_i}{\partial \mathbf{u}_i} \right| + \sum_{j \neq i} \left| \frac{\partial \mathbf{r}_i}{\partial \mathbf{u}_j} \right| \right). \quad (20)$$

The minimum eigenvalue can only be approximated, using ~~for example~~, for example, a Rayleigh quotient approach (e.g. ???):

$$\lambda_{\min} = \left| \frac{\Delta \mathbf{u}^T \Delta \mathbf{r}}{\Delta \mathbf{u}^T \mathbf{M} \Delta \mathbf{u}} \right|, \quad (21)$$

where $\Delta \mathbf{u} = \mathbf{u}^{\text{new}} - \mathbf{u}^{\text{old}}$ and $\Delta \mathbf{r} = \mathbf{r}(\mathbf{u}^{\text{new}}) - \mathbf{r}(\mathbf{u}^{\text{old}})$. The evaluation of iterative parameters requires two reductions. Reduction ~~operation~~ operations should be avoided when possible given that reduction algorithms in shared memory architectures are not embarrassingly parallel, and global communication between different nodes is needed ~~when using a distributed parallel approach~~ in distributed memory architectures. Nevertheless, this procedure can be performed every n iterations, together with the evaluation of the stopping criteria. This approach provides successful convergence while ~~minimising~~ minimizing the cost of global communication. It should be noted that the estimate of λ_{\min} relies on differences in the residual. Before the iteration is initiated, this quantity is set to zero by default.

Lastly, the DR method requires the definition of a preconditioner, typically defined as the Jacobi or diagonal preconditioner:

$$\mathbf{M} = \text{diag} \left(\frac{\partial \mathbf{r}}{\partial \mathbf{u}} \right). \quad (22)$$

While this choice allows for trivial inversion and parallelisation, it restricts the convergence property of the iteration scheme. In the context of matrix-free schemes, both the evaluation of the preconditioner and maximum eigenvalue estimates can be achieved by using automatic differentiation of the residual function.

4.2 Example: 1D Poisson problem

An example of the application of the DR method to a variable coefficient Poisson problem is depicted ~~on~~ in Fig. 2. The problem is defined as:

$$\mathbf{r}(\mathbf{u}) = -\nabla \cdot \mathbf{q} - \mathbf{b}, \quad (23)$$

$$\mathbf{q} = -\mathbf{k} \nabla \mathbf{u} \quad (24)$$

where ~~$\mathbf{k} = k(x)$~~ $\mathbf{k} = f(x)$. The differential operators are ~~discretised~~ discretized with both the FD and the FCFV over a 1D domain $\Omega = [-\frac{1}{2}, \frac{1}{2}]$ and Dirichlet boundary conditions $u(-\frac{1}{2}) = 1$, $u(\frac{1}{2}) = 2$ (Fig. 2a). The coefficient k varies with x (Fig. 2b) and is expressed as ~~$k(x) = 1 + 100 \left([1 + \exp(-200(x+0.1))]^{-1} - [1 + \exp(-200(x-0.1))]^{-1} \right)$~~ $f(x) = 1 + 100 \left([1 + \exp(-200(x+0.1))]^{-1} - [1 + \exp(-200(x-0.1))]^{-1} \right)$. The forcing term is set to zero ~~such so~~ that $b(x) = 0$. The problem ~~is converged~~ converges for different resolutions to a relative

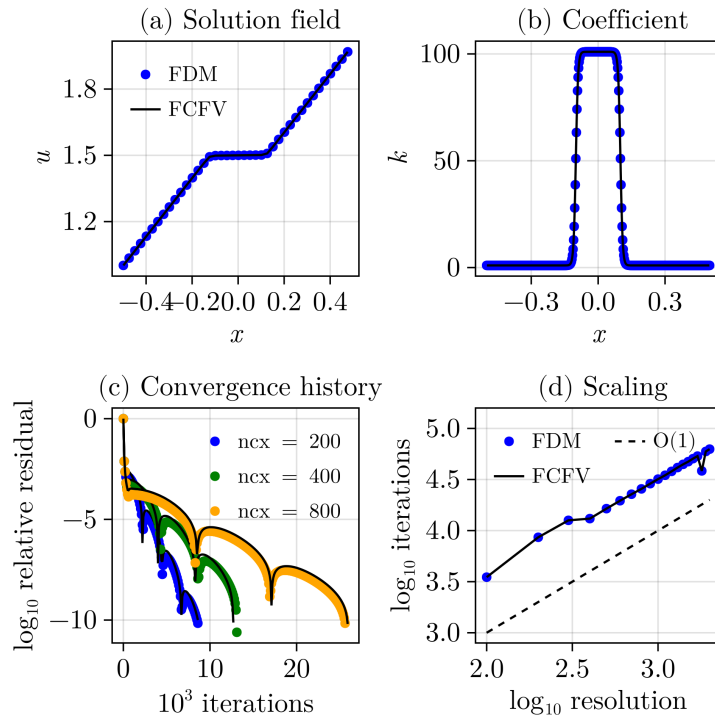


Figure 2. Properties of the DR method for the solution of a 1D Poisson problem with variable coefficient using both FD and FCFV. **(a)** Solution field in space. **(b)** Spatial distribution of the conductivity coefficient k . **(c)** Convergence history for three different resolutions. **(d)** Linear scaling of the iteration count versus the resolution. The dashed line corresponds to a linear scaling trend. In all panels, the colored scatter points correspond to the FD results, while the solid black line represents the FCFV result.

residual of 10^{-10} , with evaluation of iterative parameters and convergence check every 100 iterations. The **typical convergence behaviour**—frequency with which the parameters are reevaluated can influence the convergence behavior. In the present study, a value of 100 was empirically determined to provide an appropriate compromise. The typical convergence behavior of the DR method is shown **on** in Fig. 2c. As a consequence of diagonal preconditioning, the number of iterations needed to reach convergence depends linearly on the size of the problem (Fig. 2d). It is worth noting that for the 1D Poisson case, the FD and the FCFV deliver very similar results, both in terms of the numerical solution and convergence behavior.

The DR method is particularly well-suited to solving symmetric positive-definite problems. The solution of indefinite problems, such as those arising in saddle-point systems such as the Stokes equations, requires additional modifications.

Geodynamic problems often rely on mixed velocity–pressure formulations, and solvers must be able to perform robustly under both weakly compressible and incompressible conditions. These problems are generally *indefinite* or even *singular*, making them well-suited for Schur complement-based approaches (e.g. ???), such as the Powell–Hestenes iteration method (??)(??).

Powell–Hestenes iterations provide a solution strategy based on the augmented Lagrangian formulation, $\mathcal{L}(\mathbf{v}, P)$:

$$180 \quad \mathcal{L}(\mathbf{v}, P) = \frac{1}{2} \int_{\Omega} 2\eta |\dot{\epsilon}|^2 d\Omega - \int_{\Omega} p P \nabla \cdot \mathbf{v} d\Omega + \frac{\gamma}{2} \int_{\Omega} (\nabla \cdot \mathbf{v})^2 d\Omega - \int_{\Omega} \mathbf{b} \cdot \mathbf{v} d\Omega - \int_{\Gamma} \sigma_{BC} \cdot \mathbf{n} \cdot \mathbf{v} d\Gamma. \quad (25)$$

This approach ~~regularises~~ regularizes the original indefinite problem by adding a positive ~~penalisation~~ penalization term, γ , which can be interpreted as a numerical bulk viscosity. Powell–Hestenes iterations involve successive decoupled solutions of the velocity and pressure fields (see Alg. 1). Each velocity update requires the solution of a velocity Schur complement. In matrix-free context, this the residual of the velocity Schur complement can be formulated as a modified momentum equation:

$$185 \quad \mathbf{r}_{SCR}(\mathbf{v}^{new}) = \nabla \cdot \boldsymbol{\tau}(\mathbf{v}^{new}) - \nabla P^{old} - \gamma \nabla r_{cont}^{old} - \mathbf{fb} = \mathbf{0} \quad (26)$$

where the superscript *old* denotes values from the previous Powell–Hestenes iteration. This problem is typically positive semi-definite, and can therefore be directly tackled using the previously described DR method. Moreover, this solve does not need to be exact, which is particularly desirable in the context of iterative methods. After each inexact velocity solve, the pressure is updated using the following:

$$190 \quad P^{new} = P^{old} + \gamma r_{cont}(\mathbf{v}^{new}) \quad (27)$$

Powell–Hestenes iterations are performed until the residuals of the momentum and continuity equations are satisfied to a given tolerance. A known drawback of Powell–Hestenes iterations is their sensitivity to the choice of γ , which is further explored in this study. Typically, this parameter is defined globally, such that a single value of γ is used for the entire computational domain.

195 In the following, we will use the DR method to solve for velocity updates within Powell–Hestenes iterations. We refer to this as the PH/DR scheme. For compressible models, the numerical bulk viscosity used in the Powell–Hestenes method is set to equal the local value of the physical (or effective) bulk viscosity which may thus vary in space. In this case, the iterative procedure converges in a single step.

Algorithm 1 Powell–Hestenes with inner Dynamic Relaxation velocity solve (PH/DR)

Require: Initial guess for velocity \mathbf{v}^0 , pressure P^0 , penalty parameter $\gamma > 0$, tolerance ϵ_{PH}

1: $k \leftarrow 0$

2: **repeat**

3: Compute continuity residual: $r_{\text{cont}}^k = -\beta \frac{\hat{\Delta} P^k}{\Delta t} - \nabla \cdot \mathbf{v}^k$

4: Inner DR solve for velocity \mathbf{v}^{k+1} :

Initialize $\mathbf{v}^{(0)} = \mathbf{v}^k$, $m = 0$, a , b

repeat

(a) Compute Schur complement residual: $\mathbf{r}_{\text{SCR}}^m = \nabla \cdot \boldsymbol{\tau}(\mathbf{v}^m) - \nabla P^k - \gamma \nabla r_{\text{cont}}^k - \mathbf{b}$

(b) Update pseudo-rate of velocity: $\frac{d\mathbf{v}^{m+1}}{d\tau} = b \frac{d\mathbf{v}^m}{d\tau} + a \mathbf{M}^{-1} \mathbf{r}_{\text{mom}}^m \frac{d\mathbf{v}^{m+1}}{d\tau} = b \frac{d\mathbf{v}^m}{d\tau} + a \mathbf{M}^{-1} \mathbf{r}_{\text{SCR}}^m$

(c) Update velocity: $\mathbf{v}^{m+1} = \mathbf{v}^m + \Delta t \cdot \frac{d\mathbf{v}^{m+1}}{d\tau} = \mathbf{v}^m + \Delta t \cdot \frac{d\mathbf{v}^{m+1}}{d\tau}$

(d) Every n iterations, update DR parameters a , b

(e) $m \leftarrow m + 1$

(f) **until** $\|\mathbf{r}_{\text{SCR}}\| < \epsilon_{\text{DR}}$

Set $\mathbf{v}^{k+1} = \mathbf{v}^m$

5: Update continuity residual: $r_{\text{cont}}^{k+1} = -\beta \frac{\hat{\Delta} P^k}{\Delta t} - \nabla \cdot \mathbf{v}^{k+1}$

6: Update pressure: $P^{k+1} = P^k + \gamma r_{\text{cont}}^{k+1}$

7: $k \leftarrow k + 1$

8: **until** $\|\mathbf{r}_{\text{mom}}\| < \epsilon_{\text{PH}}$ and $\|r_{\text{cont}}\| < \epsilon_{\text{PH}}$

9: **return** (\mathbf{v}^k, P^k)

The model implementation is performed in the Julia language (?) using the `ParallelStencil.jl` package (?), which allows to effectively target different CPU and GPU hardware backends. The computations were performed on a single NVIDIA V100 SXM2 GPU with 32 GB of DRAM.

6 Solution of mechanical problems: a pathological case

In this section, we present the model configuration used to test the PH/DR solver in the context of the FD. We then provide several analyses that offer analyzes that provide quantitative insight into the scaling behaviour behavior of the PH/DR scheme.

6.1 Model configuration: multiple inclusions

In order to test the robustness of the proposed solution procedure and to evaluate its scaling behaviour behavior, we have designed a model configuration that combines known challenging elements for geodynamic flow solvers. These elements

include the presence of large and sharp viscosity variations, as well as the enforcement of the incompressibility constraint. This type of flow problem is known to be pathological for geodynamic solvers that rely on iterative solution procedures (???)

210 Both viscosity variations and incompressibility contribute to a deterioration in the conditioning of the discrete system. While this does not pose a problem for direct solvers, it is problematic for iterative solvers, whose performance strongly depends on the spectral properties of the discrete system.

The selected model configuration consists of a ~~cubic domain of volume equal to 1.0~~ unit volume cubic domain, centered at the origin. The ~~reference viscosity is set to $\eta_{\text{ref}}=1$, and the viscosity~~ viscosity within the inclusions varies according to the viscosity contrast ($\Delta\eta = \frac{\eta_{\text{max}}}{\eta_{\text{min}}}$). The background viscosity η_{ref} is defined as $\eta_{\text{ref}} = 10^{\log_{10}(\eta_{\text{min}}) + \log_{10}(\eta_{\text{max}})}$ and is set to 1. A total of 50 spherical inclusions, with variable radii and positions, are distributed throughout the domain (see Appendix A1). Weak inclusions have their viscosity set to the minimum viscosity value ($\eta_{\text{weak}} = \eta_{\text{min}}$), while strong inclusions have their viscosity set to the maximum viscosity value ($\eta_{\text{strong}} = \eta_{\text{max}}$). The flow is driven by boundary conditions that prescribe pure shear deformation in the x - y plane, with extension along the x -axis. No flow occurs across the z boundaries, and free-slip
220 conditions are applied to all faces of the domain.

An example of the model results is presented in Fig. 3, showcasing the model configuration and the resulting flow field structure for a viscosity contrast spanning four orders of magnitude. The computations were performed with a numerical resolution of 320^3 cells (66 M dof) for the FD (Fig. 3a) and 256^3 cells (168 M dof) for the FCFV (Fig. 3b). These resolutions correspond to the largest resolution that can be achieved on a single GPU with 32 GB of DRAM (NVIDIA Tesla V100-SXM2-
225 32GB).

Overall, the two methods exhibit a consistent and satisfactory qualitative agreement. Small differences between the two methods can be observed, particularly in regions where inclusions are in close proximity. These discrepancies stem from the different spatial resolution used for the FD and the FCFV ~~method~~ methods, their different order of accuracy ~~as well as~~, and their different treatments of interfaces. A more detailed comparison of the FD and FCFV results is provided in the Appendix A2.
230 The stopping criteria of the PH/DR scheme was set to 10^{-6} , the inner DR stopping criteria was set to 10^{-3} , the penalty parameter was set to 15 times the mean of the viscosity field ($\gamma = 15\langle\eta\rangle$).

6.1.1 Scaling: problem size

The first test investigates the effect of problem size. To this end, we vary the model resolution by increasing the number of cells (N_x) uniformly along each dimension. The lowest resolution is 32^3 cells, while the highest resolution is 320^3 cells (131
235 M dof) for the FD and 256^3 cells (168 M dof) for the FCFV.

The results show that for both the FD and the FCFV methods, 4 to 7 Powell–Hestenes iterations are required to reach convergence, regardless of the value of N_x . This number of Powell–Hestenes iterations appears to be largely insensitive to the problem size (Fig. 4a). In contrast, the total number of DR iterations accumulated throughout the solution process increases linearly with the problem size (Fig. 4b). This scaling behavior is consistent with that observed for the 1D Poisson problem discussed earlier (Sec. 4), and reflects the influence of the diagonal preconditioner. When normalized by the number of cells along
240 one axis, the number of DR iterations required for convergence falls between 70 and 30 per N_x . The iteration count per N_x

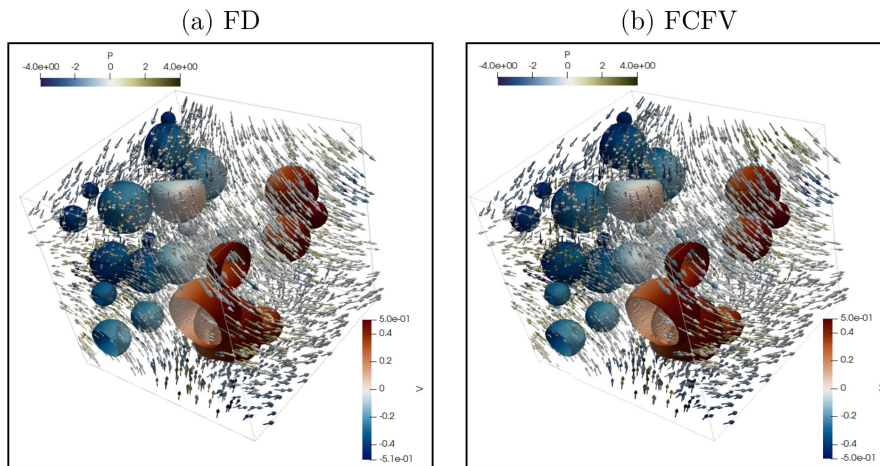


Figure 3. Example of multiple inclusion simulation result. Panels (a) and (b) shows the resulting flow field for the FD (320^3 cells) and the FCFV (256^3 cells). The inclusions are coloured by the magnitude of the velocity vector (range: $[0, 0.75]$). The black arrows indicate the velocity vectors.

tends to decrease at higher resolutions, where the flow field is better resolved (Fig. 4c). The wall-time is reported as a function of the total number of degrees of freedom, N_{dof} . For both ~~the FD and the FD~~ and FCFV, the observed scaling is approximately linear with respect to the problem size, which can be attributed to the combined effects of diagonal preconditioning and GPU parallelism (Fig. 4d). ~~Interestingly the FCFV and the~~ Interestingly, both FCFV and FD achieve comparable wall times for similar number of degrees of freedom and iterative tolerance ($1e-6$).

6.1.2 Scaling: penalty factor

We further explore the influence of the penalty factor on the scaling of the PH/DR solver. In contrast to the previous tests, we vary the value of γ independently of the values of the viscosity field. We perform the computations for three different values of $\Delta\eta$.

The number of Powell–Hestenes iterations strongly depends on the magnitude of the penalty parameter γ . In particular, very small values of γ lead to a large number of Powell–Hestenes iterations that exceed 50 (Fig. 5a). This number decreases to fewer than 10 when $\gamma > 100$, across all ~~considered~~ values of $\Delta\eta$. The total number of DR iterations is also higher for small values of γ , reflecting the increased number of Powell–Hestenes iterations required for convergence. However, excessively large values of γ also degrade convergence. Specifically, higher values of γ increase the spectral radius of the discrete velocity Schur complement, which in turn increases the number of DR iterations needed to reach convergence (Fig. 5b). An optimal range for γ is found between 10 and 1000, depending on the value of $\Delta\eta$. The wall-time exhibits a similar dependence on γ , closely following the behavior of the total number of DR iterations, with minimum values observed within the range 10–1000 (Fig. 5c).

Problem size

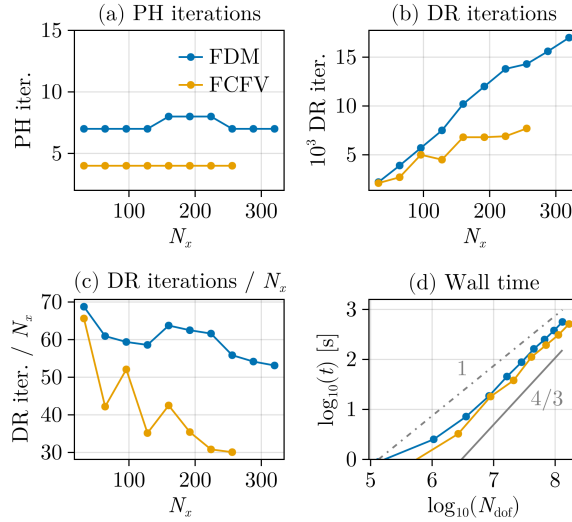


Figure 4. Influence of problem size on the convergence of the PH/DR flow solver for a viscosity contrast of $1e4$ using the FD and the FCFV methods. (a) reports the total number of Powell–Hestenes iterations, (b) the total number of DR iterations for the entire solution procedure, (c) the number of DR iterations per grid point along one dimension, and (d) the wall-time.

260 The magnitude of the penalty factor must be selected carefully to avoid an excessive number of iterations ~~and degraded~~
, which would degrade the wall-time. To guide this choice, we systematically varied the viscosity contrast and computed
the optimal penalty value (γ^{opt}) for each corresponding $\Delta\eta$. The results reveal a linear relationship between γ^{opt} and $\Delta\eta$ in
 \log_{10} – \log_{10} space (Fig. 5d). This relationship allows for a least-squares fit with a slope of 0.3 and an intercept of 0.575. The
resulting fit can be used to select an appropriate penalty parameter when varying the viscosity contrast. This particular scaling
265 produces the best performance for the given model configuration; however, it should not be regarded as generally applicable
to all cases. In most of the examples presented hereafter, we choose γ to be proportional to the mean of the viscosity field
($\gamma \propto \langle \eta \rangle$).

6.1.3 Scaling: viscosity contrast

We investigate the effect of viscosity contrast on the PH/DR solver by systematically varying it over a range spanning two to
270 six orders of magnitude. As expected, increasing the viscosity contrast impacts the performance of the PH/DR solver.

The number of Powell–Hestenes iterations required for convergence ranges from 6 to 9 (Fig. 6a). The decrease in the number
of Powell–Hestenes iterations with increasing $\Delta\eta$ is attributed to the scaling of the penalty factor with the mean of the viscosity
field. Therefore, higher values of $\Delta\eta$ lead to higher values of γ , which in turn reduces the number of Powell–Hestenes iterations

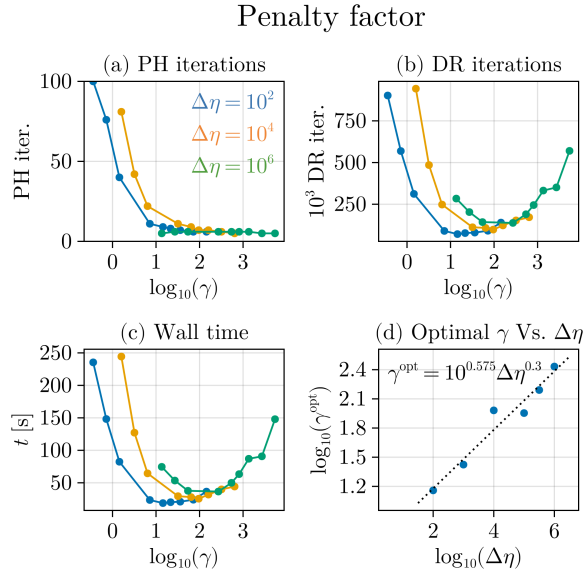


Figure 5. Influence of the penalty factor on the convergence of the PH/DR flow solver for a resolution of 128^3 cells and 3 values of $\Delta\eta$. (a) reports the total number of Powell–Hestenes iterations, (b) the total number of DR iteration for the entire solution procedure, (c) the wall time, and (d) the optimal value of γ as function of $\Delta\eta$

required for convergence. The total number of DR iterations scales approximately linearly with the viscosity contrast (Fig. 6b).
 275 Consequently, the wall-time also exhibits a linear dependency on the viscosity contrast (Fig. 6c).

6.1.4 Scaling: DR solve tolerance

Another important numerical parameter of the PH/DR procedure is the tolerance ϵ_{DR} selected to exit the inner DR iterations (ϵ_{DR}). We have systematically varied the value of ϵ_{DR} from 10^{-5} to $10^{-0.5}$.

The magnitude of ϵ_{DR} has a significant impact on the number of Powell–Hestenes iterations. Small values lead to more
 280 accurate inner DR solves, resulting in fewer Powell–Hestenes iterations. In contrast, increasing ϵ_{DR} causes a quadratic increase
 in the number of Powell–Hestenes iterations (Fig. 7a). The trend differs for the total number of DR iterations. Very small and
 very large values of ϵ_{DR} lead to a high total iteration count, while a minimum is observed for ϵ_{DR} in the range 10^{-3} to 10^{-2} .
 High values of ϵ_{DR} result in inaccurate inner solves, increasing the number of Powell–Hestenes iterations. In contrast, very
 285 low values lead to over-solving of the inner system, unnecessarily increasing the total DR iterations (Fig. 7b). This **behaviour**
behavior directly affects the wall-time, which can vary by up to a factor of 2 depending on the choice of ϵ_{DR} . Therefore, ϵ_{DR}
 must be carefully tuned to avoid these two limiting cases.

Viscosity contrast




Figure 6. Influence of viscosity contrast on the convergence of the PH/DR flow solver for a resolution of 128^3 cells and using γ^{opt} . (a) reports the total number of Powell–Hestenes iterations, (b) the total number of DR iteration for the whole solution procedure, (c) the wall time.

7 Compressibility, multi-physics and ~~non-linear~~ nonlinear models

In this section, we show how the proposed PH/DR solution strategy can be further extended to solve steady mechanical problems that involve enriched physical models. These models are performed using the FD method, as the FCFV method still needs
 290 to be developed to tackle two-phase flow and ~~non-linear~~ nonlinear mechanical problems.

7.1 Compressible mechanical problem

Although the previous examples were performed in the incompressible limit, the PH/DR scheme can adequately handle compressible deformation or flow. Here, we show examples of compressible flow performed in 2D, using ~~the~~ FD. The model configuration consists of a ~~$[0, 1] \times [0, 1]$ domain~~ $\Omega \in [0, 1] \times [0, 1]$ with a circular inclusion of radius 0.2 (Fig. 8a).
 295 The background viscosity is set to 1 and the inclusion viscosity is 100. Pure shear background deformation is applied at the boundaries. Compressibility is accounted for by including a bulk viscosity term:

$$\nabla \cdot \boldsymbol{\tau} - \nabla P = \mathbf{0}, \quad (28)$$

Inner DR solver tolerance

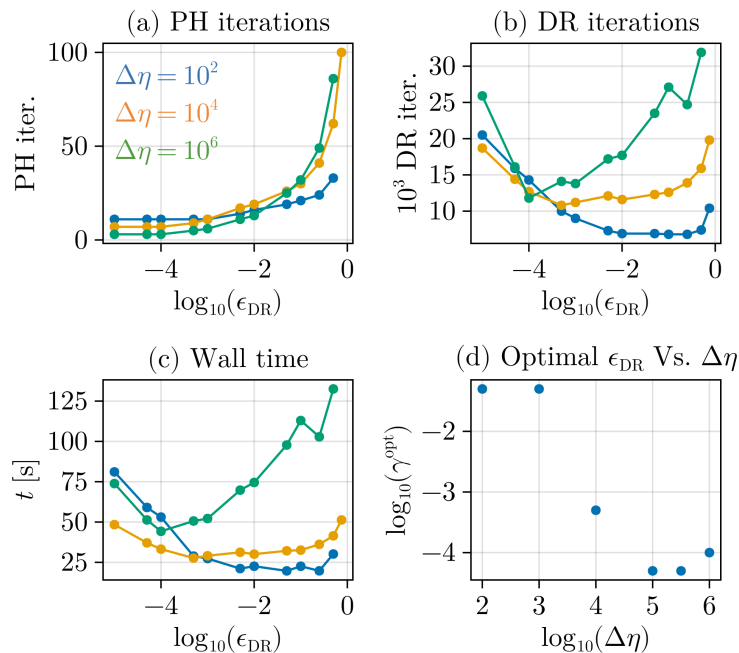


Figure 7. Influence of the inner DR solve tolerance on the convergence of the PH/DR flow solver for a resolution of 128^3 cells and $\Delta\eta = 10^4$. (a) reports the total number of Powell–Hestenes iterations, (b) the total number of DR iteration for the whole solution procedure, (c) the wall time. (d) Influence of the viscosity contrast.

$$\nabla \cdot \mathbf{v}^s + \frac{P}{\eta_b} = 0 \quad (29)$$

The degree of compressibility varies with the value of bulk viscosity (η_b). To solve this problem, one can use PH/DR, setting the penalty factor independently of the bulk viscosity. Here, we express the penalty factor as the quasi-harmonic average of a reference numerical bulk viscosity ($\gamma_{\text{num}} = 60\langle\eta\rangle$) and the physical bulk viscosity, $\gamma_{\text{phys}} = \eta_b$, so that $\gamma = (1/\gamma_{\text{phys}} + 1/\gamma_{\text{num}})^{-1}$. In this case, several Powell–Hestenes iterations will be needed to reach convergence. Alternatively, one can set the penalty factor to the bulk viscosity ($\gamma = \eta_b$). In that case, PH/DR becomes equivalent to a basic DR scheme and only one Powell–Hestenes iteration is needed to reach convergence. We observe that in the case of compressible problems ($\eta_b = 2.0$, equivalent to a Poisson ratio of 0.286), the basic DR can be beneficial, leading to a lower number of iterations [in comparison to PH/DR compared to PH/DR](#) (Fig. 8b). The PH/DR solver exhibits a characteristic sawtooth convergence history that reflects the update of pressure at each Powell–Hestenes iteration (Fig. 8b). In contrast, in weakly compressible cases ($\eta_b = 2e3$, equivalent to a Poisson ratio of 0.499), the PH/DR schemes outperform the standard DR iteration approach. In the context of geodynamic problems involving weakly compressible-to-incompressible materials, PH/DR is preferred.

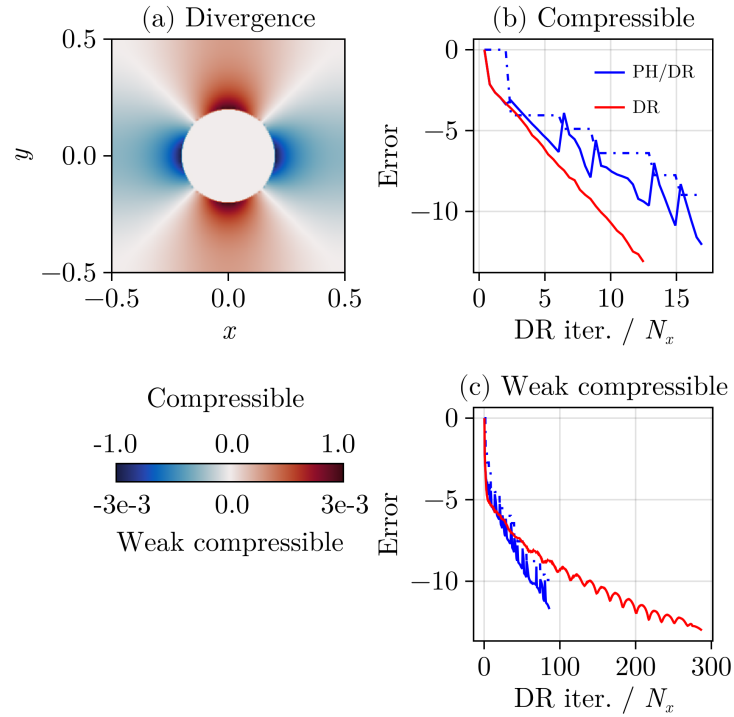


Figure 8. Example of compressible flow: pure shear deformation with viscosity heterogeneity. (a) Divergence field either compressible case ($\eta_b = 2$) or weakly compressible case ($\eta_b = 2e3$). Convergence history of DR and PH/DR for the (b) compressible and (c) weakly compressible case. For the PH/DR case, the solid line represents the velocity error norm and the dotted line corresponds to the pressure error norm. For the PH/DR case, iterations are summed over the Powell–Hestenes needed for converging the problem. In (a) and (b), x and y have arbitrary units.

310 7.2 Hydro-mechanical coupling

We further applied the PH/DR scheme to the solution of a coupled multi-physics problem. The equations governing deformation and fluid flow in a Darcy poro-viscous medium can be expressed following ?:

$$\nabla \cdot \bar{\tau} - \nabla \bar{P} = 0, \quad (30)$$

$$\nabla \cdot \mathbf{v}^s + \frac{\bar{P} - P^f}{(1 - \phi)\eta_\phi} = 0, \quad (31)$$

$$315 \quad \nabla \cdot \mathbf{q}^D - \frac{\bar{P} - P^f}{(1 - \phi)\eta_\phi} = 0. \quad (32)$$

Here, the Darcy flux is expressed as $\mathbf{q}^D = -k^f/\eta^f \nabla P^f$. The overbar denotes phase-averaged quantities, ϕ is porosity, and the superscripts s and f indicate solid and fluid properties, respectively. k and η refer to permeability and viscosity. An elliptic equation for fluid pressure can be obtained by substituting the Darcy flux into Equation (32).

The PH/DR scheme is extended by including fluid pressure updates within the inner DR iterations. The update parameters
 320 for fluid pressure are obtained from the automatic DR procedure described in Section 4. The pseudo-time steps and damping coefficients for velocity and fluid pressure are determined independently, with each variable assigned its own value. As in previous examples, the total pressure is updated at each Powell–Hestenes iteration.

The test case we consider uses a configuration similar to that in the previous example (Sec. 7.1). The shear viscosity of the inclusion is assigned a value 100 times greater than that of the matrix, while the bulk viscosity is taken as twice the shear
 325 viscosity. In addition, the background porosity is set to 1e-3, and the fluid permeability-viscosity ratio (k^f/η^f) is set to 1e-3. The penalty factor is defined as

$$\gamma = \left(\frac{1}{\gamma_{\text{phys}}} + \frac{1}{\gamma_{\text{num}}} \right)^{-1},$$

where $\gamma_{\text{phys}} = (1 - \phi)\eta_\phi$ and $\gamma_{\text{num}} = 5\langle\eta^s\rangle$.

Figure 9 shows examples of solution fields and the scaling properties of the PH/DR scheme. We compute the two-phase flow
 330 problem with a resolution varying from 31^2 to 992^2 cells and apply a tolerance of 1×10^{-7} . For the chosen parameters, pronounced differences are observed between the total and fluid pressure fields (Fig. 9a,b). The total pressure field is sharply defined and zero inside the stiff inclusion, in contrast to the fluid pressure field, which is smoother and diffuses within the inclusion. We observe the characteristic sawtooth convergence history for both momentum and fluid continuity (Fig. 9c).

We find that low-resolution models require a relatively high number of iterations (Fig. 9d). Nevertheless, the number of
 335 DR iterations per N_x drops to a stable low value (approximately 20) for resolutions greater than 100^2 cells. This indicates a linear dependence of the number of iterations on the problem size, which is consistent with the behavior documented in the previous examples.

7.3 Visco-elasto-viscoplasticity

In this example, we show how ~~non-linear~~ nonlinear material properties can be incorporated into the PH/DR solution strategy,
 340 considering the case of non-associated plasticity. With such ~~non-linear~~ nonlinear problems, parameters such as the effective viscosity vary throughout the ~~non-linear~~ nonlinear solution procedure, which directly affects the eigenvalues of the discrete operator. It is thus important to reevaluate the maximum pseudo-time step (Eq. (18)) during the iterations to ensure stability of the procedure. Here, we reevaluate the pseudo-time step every 100 iterations by applying the Gershgorin circle theorem (Eq. (21)).

345 We adopt a single-phase formulation and apply a visco-elasto-viscoplastic material model (?). An additive decomposition of the deviatoric strain rate and the divergence rate is assumed:

$$\dot{\epsilon} = \dot{\epsilon}^v + \dot{\epsilon}^c + \dot{\epsilon}^p \tag{33}$$

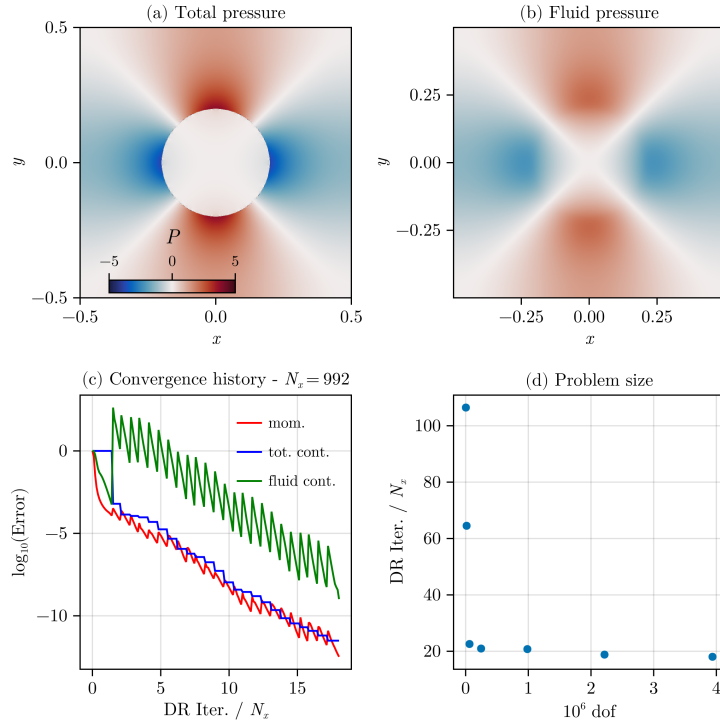


Figure 9. Poro-viscous hydro-mechanical simulation. (a) Total pressure field. (b) Fluid pressure field. (c) Convergence history for a resolution of 992^2 . (d) Scaling of iterations relative to the problem size. In (a) and (b), x and y have arbitrary units.

$$\nabla \cdot \mathbf{v} = (\nabla \cdot \mathbf{v})^e + (\nabla \cdot \mathbf{v})^p \quad (34)$$

where the superscripts v, e, and p correspond to the viscous, elastic, and plastic components, respectively.

350 The viscoplastic Drucker–Prager yield and potential functions are defined as:

$$f = \tau_{II} - P \sin \Phi - C \cos \Phi - \dot{\lambda} \eta^{vp} \quad (35)$$

$$q = \tau_{II} - P \sin \Psi \quad (36)$$

where τ_{II} , C , Φ , Ψ , and η^{vp} denote the second deviatoric stress invariant, the cohesion, the friction angle, the dilatancy angle, and the viscoplastic viscosity, respectively.

355 The model domain spans $[0, 2] \times [0, 1]$ is $\Omega \in [0, 2] \times [0, 1]$ km² and contains a circular inclusion of radius 1e-1 km (Fig. 10a). A pure shear loading rate of $1e-15$ s⁻¹ is applied, inducing horizontal extension. The inclusion has a low viscosity (1e10 Pa.s), while the surrounding matrix is effectively elasto-plastic (1e23 Pa.s). Both shear and bulk moduli are set to 3e10 Pa. The cohesion is set to 50 MPa, the friction and dilatancy angles are set to 35° and 5° , respectively, and the viscoplastic viscosity is 2e20 Pa.s. The simulation proceeds over 50 time steps, each of size $\Delta t = 4e10$ s.

360 The temporal evolution of the second invariant of the deviatoric strain rate is shown in Fig. 10a, b, and c, while the evolution of the mean stress is shown in Fig. 10d. The model undergoes an initial phase of elastic loading (Fig. 10a), after which shear banding initiates, propagates (Fig. 10b), and reflects off the domain boundaries (Fig. 10c). As a result, stress drops after 30 ky witnessing underlying structural softening (Fig. 10d). With a relative tolerance of $1e-6$, the number of iterations per N_x varies between 5 and 25 (Fig. 10e). Low iteration counts correspond to linear-elastic steps, while higher counts occur during
365 ~~non-linear~~nonlinear steps involving significant changes in the solution pattern (e.g. τ -reflections). The convergence history for the most challenging step is shown in Fig. 10f, where a sawtooth pattern is also observed.

Overall, the model converges reliably without the need for manual tuning of numerical parameters throughout the simulation.

7.4 Practical Performance of the PH/DR Method

In this section, we investigate the benefits of the proposed PH/DR method for solving two practical geodynamic problems.
370 The main benefit of the PH/DR method over the APT approach is its ability to automatically adapt to transient changes in solution patterns, caused by variation of the model geometry or by the emergence of internal features. The PH/DR approach can therefore determine parameters that ensure convergence while also significantly reducing the overall wall time. The results presented below are intended to be indicative. In principle, the APT method could deliver results that are at least as good as those of the PH/DR method. However, this would require manually tuning the iteration parameters at every time step, which
375 is not a desirable task. Several simulations were performed using the open-source code `JustRelax.jl` (?), which features both the PH/DR solver and the Accelerated Pseudo-Transient (APT) method (?).

The first model is a simplified subduction setup based on the configuration of ?. Simulations were performed over a model duration of 25 Myr, after which the slab tip reaches a depth of approximately 500 km (Fig. 11). We observe that the use of the PH/DR scheme significantly reduces cumulative wall time compared to the APT method (Fig. 11a). The number of iterations
380 per grid point in one dimension is also well behaved, fluctuating around 20 iterations per grid point. For the same tolerance, the APT scheme requires about 200 iterations per grid point and exhibits larger fluctuations (Fig. 11b). Overall, both approaches converge at each iteration and yield similar results (Fig. 11c, d).

The second model is a crustal-scale shear banding setup, which features a visco-elasto-viscoplastic rheology and temperature-dependent creep laws. The reader is referred to ? for the detailed model configuration and material parameters. The simulations are run for
385 100 time steps, in which the solution pattern evolves significantly in response to shear banding and frictional viscoplasticity. The results presented in Fig. 12 show the cumulative wall time and the number of iterations over 100 time steps for various model resolutions. Models performed with the PH/DR scheme consistently result in 2 to 12 times shorter wall times (Fig. 12a). Unlike APT models, simulations using the PH/DR method provide a reliable and well-defined range of iterations per time step. Nevertheless, the overall shear banding patterns are similar (Fig. 12c,d). The minor differences may be attributed to the fact
390 that APT models could not reach convergence at every time step.




Figure 10. Visco-Elasto-Viscoplastic shear banding. (a), (b) and (c) Model evolution (strain rate field). (d) Evolution of the mean stress. (e) Number of iteration per grid point in x over simulation time. (f) Convergence history for the most demanding time step.

8 Discussion

Our study presents a quantitative assessment of the performance of the DR method, extends it to incompressible flow problems through a Powell–Hestenes strategy, and demonstrates its applicability to nonlinear and multiphysics problems.

The efficiency of pseudo-transient (PT) solvers relies on a careful selection of iterative parameters (see Appendix A3). The proposed DR and PHDR solvers are not expected to outperform standard PT solvers with optimally chosen parameters. However, the DR-based methods introduced here are designed to automatically determine these iterative parameters. This automated selection introduces only a minor computational overhead and has a negligible impact on the convergence history compared

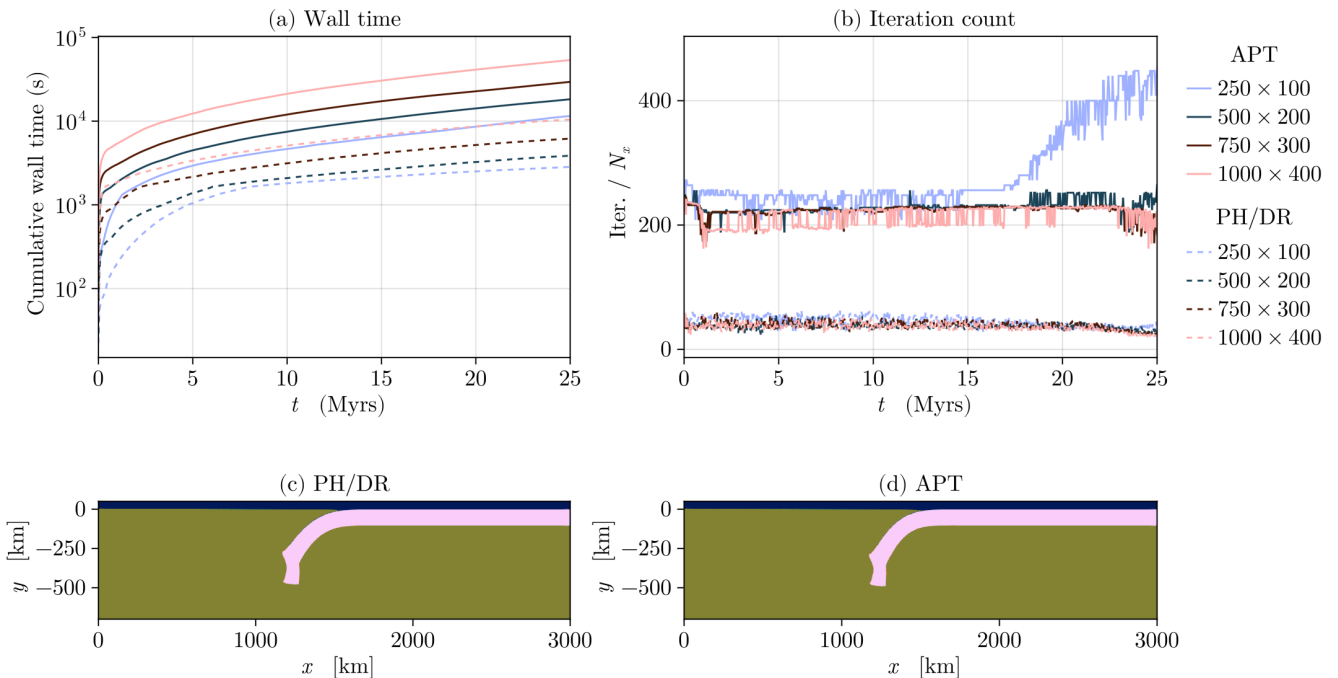


Figure 11. (a) Cumulative wall times over 25 Myr of model simulation steps using both PH/DR and APT methods at various resolutions. (b) Iterations per time step for the two methods. (c) Map of material phases at 25 Myr using the PH/DR method. The mantle is green, the slab is pink and the sticky air is blue. (d) Same as (c) for the APT method.

to an optimally tuned PT solver (see Appendix A3–A4). In practice, DR-based schemes can offer substantial advantages over standard pseudo-transient solvers that do not feature automatic tuning of iteration parameters.

400 Unlike multigrid methods, DR-based approaches exhibit a total iteration count that increases linearly with problem size (see Fig. 2 and Fig. 4b). Although this behavior may appear disadvantageous, the DR methods are straightforward to parallelize and require only a few global communication steps. In all the presented examples, global communication operations (e.g., norms and inner products) and iterative parameter updates were performed every 100 DR iterations. This makes DR-based methods particularly suitable for GPU-type GPU parallelism.

405 In particular, we show that, when using a single GPU, the wall time for three-dimensional incompressible and heterogeneous Stokes problems scales quasi-linearly with problem size (Fig. 4d). This property is attractive for large-scale 3D geodynamic simulations and makes DR-based algorithms competitive with legacy multigrid-based solvers.

Furthermore, DR-based solvers can be naturally extended to multi-GPU parallelism. Computational frameworks such as ImplicitGlobalGrid.jl (?) provide efficient communication scheduling schemes, enabling near-ideal scaling behavior on-in modern computing clusters (?).




Figure 12. (a) Cumulative wall times over 100 time steps for both PH/DR and APT methods at various resolutions. (b) Iterations per time step for the two methods. The number of pseudo-transient iterations per time step was capped at $250e3$; as a result, some APT simulations did not reach the convergence criteria (relative tolerance of $1e-6$). (c) Map of the second deviatoric strain rate invariant at time step 100 (1 Myr) obtained with the PH/DR method. (d) Same as (c) for the APT method.

In the future, several improvements may be envisaged, in particular regarding preconditioning. To our knowledge, Jacobi preconditioning is known to be relatively inefficient in comparison to other preconditioners (e.g. $\bar{\bar{}}$ -incomplete Cholesky). Nevertheless, the application of Jacobi is fully parallel, and thus efficiently implemented on GPUs. Recently, iterative implementations of incomplete LU and Cholesky preconditioners have been introduced (??). Although these methods are fully parallel, achieving effective preconditioning requires accurate application of the LU/Cholesky operators, often necessitating many iterations. Whether such an approach would offer benefits within the framework of DR solvers remains to be established. Multigrid preconditioning represents another possible avenue. In principle, it could maintain a constant iteration count as the resolution increases. However, its implementation requires hierarchical data structures and coarse-grid solvers on GPUs, and it remains uncertain whether this added complexity would yield tangible reductions in wall time, or enhance the already quasi-linear wall time scaling.

DR solvers are traditionally used in conjunction with the Finite Element Method (FEM; $\bar{\bar{}}$) (e.g. ??). Here, we have shown that such solvers can also be applied to finite difference methods, as well as finite volume or low-order discontinuous

Galerkin methods. In practice, efficient implementation is facilitated by the structured nature of these problems, which allows information from neighboring cells to be easily gathered for stencil evaluations. Extending DR solvers to unstructured FEM or FCFV discretizations in the context of GPU computing requires specialized algorithms to improve memory access patterns (e.g., ??)(e.g., ??). The combination of these approaches with DR solvers may provide an alternative to traditionally employed solvers, a possibility that remains to be explored.

The relatively low algorithmic complexity of DR methods makes them particularly attractive for solving coupled-nonlinear and multiphysics coupled problems. While we have focused on basic hydro-mechanical and linear plasticity models, future implementations could naturally extend these solvers to include thermal or chemical couplings.

9 Conclusions

In this contribution, we have employed DR-based techniques to develop robust and high-performance solvers for geodynamic problems. These solvers incorporate automatic tuning of iterative parameters—including pseudo-time steps and damping coefficients—based on estimates of the spectral bounds of the discrete operators.

We applied this approach. The approach was applied to problems discretized on structured grids using both finite-difference (FD) and face-centered finite-volume/discontinuous Galerkin (FCFV) methods. To handle incompressible deformations, DR-based solvers were combined with Powell—Hestenes iterations, resulting in enabling the successful solution of heterogeneous Stokes problems.

When implemented on GPU hardware-accelerators, the solvers exhibit a quasi-linear scaling of solution time with problem size, demonstrating excellent computational efficiency. Moreover, the method can be trivially extended to handle compressible flow, multiphase flow, and nonlinear constitutive laws.

This flexibility stems from the algorithmic simplicity of the DR-based methods, which makes them particularly well-suited for large-scale geoscientific simulations and opens new paths for exploring complex coupled physical processes.

Appendix A: Appendix

A1 Three-dimensional multiple inclusion problem

The position and type of inclusion used in the example in the main text are provided in the following table (Tab. A1). The extent of the model domain and the viscosity contrast are detailed in the main text.

A2 Two-dimensional heterogeneous Stokes flow with FD and FCFV

To verify our implementation of the PH/DR iterative solvers, the results were compared against solutions obtained using direct solvers. This verification could only be conducted in 2D due to the prohibitive computational cost of direct solvers in 3D. Both FD and FCFV models were executed using configurations analogous to the 3D setups. The inclusion positions are given in the table below (Tab. A2): weak inclusions have a viscosity of $1e-3$, while strong inclusions have a viscosity of 1000. No

x_0	y_0	z_0	r	type	x_0	y_0	z_0	r	type
0.333	0.123	0.436	0.134	weak	0.337	-0.037	0.339	0.133	weak
0.153	-0.427	0.431	0.149	strong	0.326	0.296	-0.209	0.104	weak
0.079	-0.367	0.318	0.053	strong	0.148	0.138	0.047	0.086	strong
-0.217	-0.020	0.319	0.030	strong	-0.144	-0.041	0.090	0.094	weak
0.328	-0.128	0.302	0.103	weak	-0.394	-0.197	0.060	0.109	weak
-0.337	0.244	0.123	0.114	weak	0.293	-0.456	-0.114	0.026	strong
-0.152	-0.347	-0.493	0.045	strong	-0.479	-0.123	-0.069	0.033	weak
-0.039	0.359	0.077	0.112	weak	0.478	-0.103	-0.098	0.046	strong
-0.463	0.289	-0.317	0.110	weak	-0.495	0.189	-0.309	0.086	strong
0.280	0.260	-0.373	0.131	strong	-0.064	0.386	0.034	0.098	strong
-0.209	0.470	0.245	0.110	strong	-0.474	0.403	-0.387	0.030	weak
0.302	-0.241	0.099	0.080	weak	-0.227	0.025	0.456	0.048	weak
-0.159	-0.379	0.434	0.009	strong	0.323	-0.293	0.446	0.008	strong
-0.356	-0.048	0.309	0.129	strong	-0.293	-0.331	0.225	0.072	weak
-0.047	-0.326	-0.434	0.113	strong	-0.364	0.014	0.270	0.096	weak
-0.289	-0.273	0.458	0.083	weak	0.298	-0.073	0.390	0.079	weak
-0.482	0.327	0.168	0.033	weak	0.462	-0.020	-0.047	0.036	strong
-0.145	-0.190	0.389	0.111	strong	-0.084	0.479	-0.000	0.133	strong
0.440	0.226	-0.178	0.054	weak	-0.381	0.448	0.218	0.008	strong
-0.471	0.265	0.311	0.056	weak	-0.123	0.103	-0.014	0.039	weak
0.091	-0.407	0.242	0.013	strong	0.477	0.380	0.386	0.082	weak
-0.402	0.069	-0.425	0.137	weak	-0.430	0.403	-0.313	0.027	strong
-0.198	-0.122	-0.371	0.035	strong	-0.134	-0.081	-0.222	0.018	weak
-0.114	0.403	0.414	0.091	strong	0.353	0.131	-0.090	0.095	weak
-0.332	-0.051	0.374	0.014	strong	0.218	0.146	-0.143	0.002	strong

Table A1. Tabulated values of x_0 , y_0 , z_0 , r , and type (weak/strong). Each block of 5 rows is shown side by side.

smoothing of the viscosity contrast was applied. A background pure shear rate and free-slip boundary conditions were imposed on the box boundaries.

455 ~~The~~ FD models solved with a direct solver follow the approach of ?, whereas the FCFV direct-solver implementation is based on ?. All simulations were performed on meshes of 128^2 cells. The resulting pressure fields are shown in Fig. A1.

For both discretizations, good agreement is obtained between models solved iteratively and those solved directly, demonstrating that the FD and FCFV implementations are consistent with previous studies. In addition, this comparison highlights differences between the two schemes: while the FD approach exhibits spurious oscillations at sharp viscosity contrasts, the
 460 FCFV method yields a smoother pressure field. It should be noted that the mesh used in the FCFV simulations is similar to that of the FD models and does not explicitly capture the material interface through traction continuity enforcement.

x_0	y_0	r	type
0.0	0.0	0.2	weak
0.2	0.4	0.09	strong
-0.3	0.4	0.05	strong
-0.4	-0.3	0.08	weak
0.0	-0.2	0.08	strong
-0.3	0.2	0.1	weak
0.4	-0.2	0.07	strong
0.3	-0.4	0.08	weak
0.35	0.2	0.07	weak
-0.1	-0.4	0.07	strong

Table A2. Tabulated values of x_0 , y_0 , r , and type (weak/strong). Each block of 5 rows is shown side by side.

A3 Comparison with ~~results of ?~~ the pseudo-transient formulation

The DR and PH/DR schemes are formulated in a manner similar to pseudo-transient (PT) solvers. The solution fields are updated iteratively, with update steps determined through stability analysis to ensure stable integration in pseudo-time. Most
 465 importantly, damping is applied, which effectively modifies the form of the equation to be integrated in pseudo-time, resulting in a second-order integration scheme following ?. The pseudo-time step and damping parameters constitute the two numerical parameters that control the convergence behavior.

Optimal values for these parameters can be derived analytically (e.g. ?). However, they depend on the spectral properties of the discrete operator ??. Consequently, determining iteration parameters a priori that remain optimal throughout simulations
 470 involving evolving geometry or time-dependent material properties is highly challenging. This may result in degraded convergence, stagnation, or even divergence.

The DR and PH/DR schemes address this difficulty by estimating spectral bounds and adjusting the iteration parameters accordingly. In practice, the main additional components are: (1) a routine to estimate the maximum eigenvalue (Gershgorin

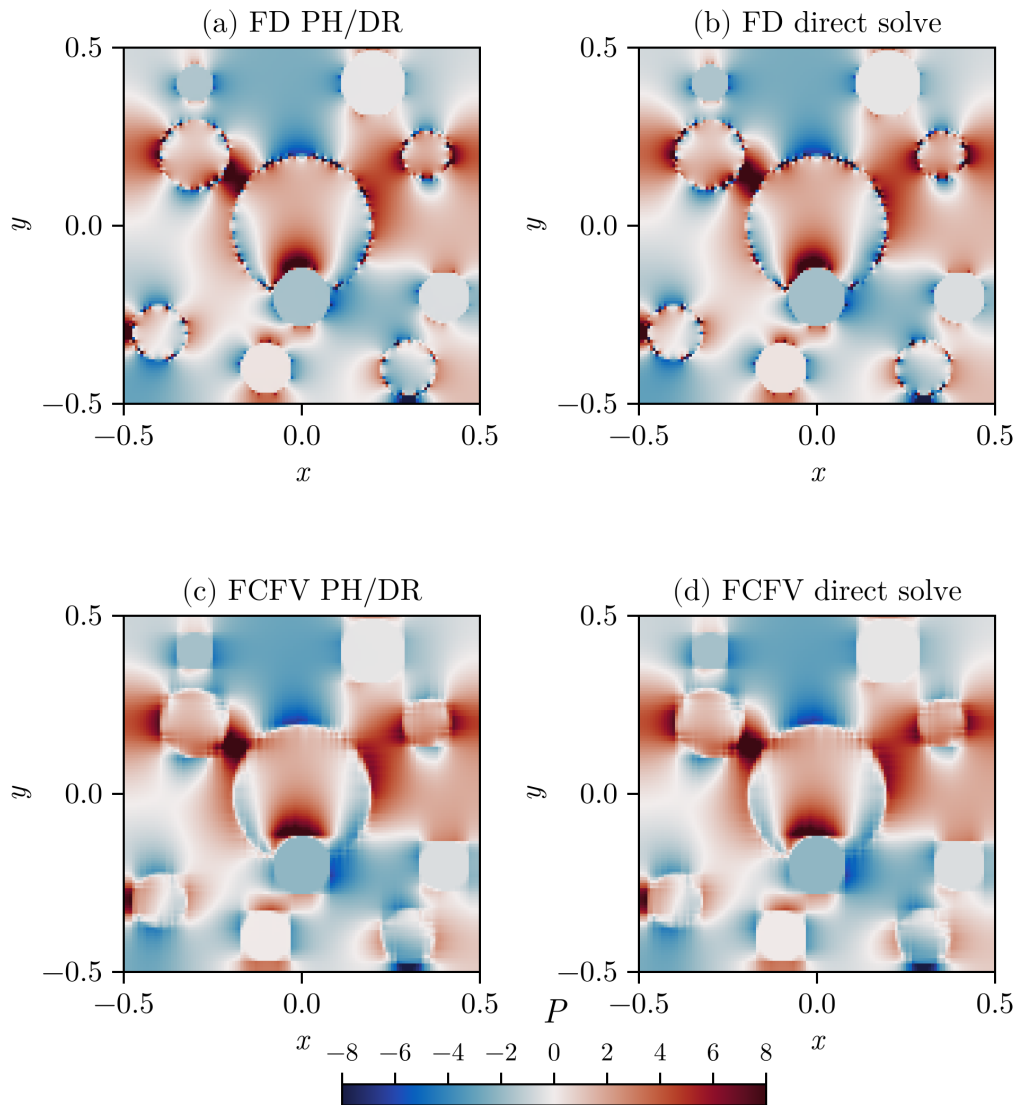


Figure A1. (a) Pressure field obtained with FD using a PH/DR scheme. (b) Pressure field obtained with a direct solution scheme. (c) Pressure field obtained with FCFV using a PH/DR scheme. (d) Pressure field obtained with FCFV using a direct solution scheme. In all panels, x and y have arbitrary units.

475 analysis), (2) a routine to approximate the minimum eigenvalue (Rayleigh quotient), and (3) a recalculation of the iteration parameters based on these estimates.

In addition, the PH/DR scheme features an outer loop for pressure updates. During the velocity update cycles, the pressure is held fixed. At the end of each such cycle, the pressure is updated using a numerical bulk viscosity and the continuity residual.

following the classical Powell—Hestenes approach. The pressure is thus updated less frequently than the velocity, which constitutes a key difference from the pseudo-transient approach presented in ?.

480 **A4** Comparison with the pseudo-transient formulation and the results of ?

We employed the PT solvers developed by ? to assess the relative performance of the PH/DR solver, using the single-inclusion benchmark configuration. The inclusion has a viscosity of $1e-3$, while the surrounding matrix has a viscosity of 1. A smoothing of the viscosity contrast similar to that used in ? was applied. Both the PT and PH/DR models were run to the same relative tolerance of $1e-6$. The resulting model fields are shown in Fig. A2.

485 The PT models of ? were computed using optimally tuned iterative parameters and thus represent the best possible performance bound. The present results demonstrate that the additional computational cost associated with the automatic parameter selection inherent to the PH/DR scheme is very small. Both methods converge in fewer than 20 iterations per grid point for this model configuration. Because pressure updates are less frequent than velocity updates, the convergence of the pressure residual does not closely track that of the momentum residuals.

490 *Code availability.* The various software developed and used in this study is licensed under MIT License. The latest versions of the code is available in Zenodo (?).

Author contributions. **TD** Original study design, code and algorithm development, simulations, figures, writing, review. **AdM** Development, study design, writing, review **RS** Development of discretization, verification, writing, and review **LR** Writing, figures, concept, and review **IU** Writing, concept, review **AS** Code testing, writing, review

495 *Competing interests.* **LR** is a member of the editorial board of GMD.

Acknowledgements. We thank the Swiss Geocomputing Centre at the University of Lausanne for providing access to the Octopus super-computer. **LR** and **IU** gratefully acknowledge funding from the Swiss University Conference and the Swiss Council of Federal Institutes of Technology through the Platform for Advanced Scientific Computing (PASC) program. This research was also supported by the Swiss National Supercomputing Centre (CSCS) under project ID c23 and c44, awarded via the PASC projects GPU4GEO and dGPU4GEO. We

500 thank Pascal Aellig for providing access to a GH200 node.




Figure A2. (a) Pressure field obtained by ?. (b) Pressure field obtained in this study. (c) Convergence of momentum and continuity residuals.

In (a) and (b), x and y have arbitrary units.