

Matrix-free higher-order finite element methods for hyperelasticity

R. Schussnig^{a,*}, N. Fehn^b, P. Munch^{c,d}, M. Kronbichler^a

^a Faculty of Mathematics, Ruhr University Bochum, Germany

^b Institute of Mathematics, University of Augsburg, Germany

^c Department of Information Technology, Uppsala University, Sweden

^d Institute of Mathematics, Technical University of Berlin, Germany

ARTICLE INFO

Keywords:

Finite-strain problem

Matrix-free

Finite-element method

Hyperelasticity

Geometric multigrid

ABSTRACT

This work presents a matrix-free finite element solver for finite-strain elasticity adopting an *hp*-multigrid preconditioner. Compared to classical algorithms relying on a global sparse matrix, matrix-free solution strategies significantly reduce memory traffic by repeated evaluation of the finite element integrals.

Following this approach in the context of finite-strain elasticity, the precise statement of the final weak form is crucial for performance, and it is not clear a priori whether to choose problem formulations in the material or spatial domain. With a focus on hyperelastic solids in biomechanics, the arithmetic costs to evaluate the material law at each quadrature point might favor an evaluation strategy where some quantities are precomputed in each Newton iteration and reused in the Krylov solver for the linearized problem. Hence, we discuss storage strategies to balance the compute load against memory access in compressible and incompressible neo-Hookean models and an anisotropic tissue model. Additionally, numerical stability becomes increasingly important using lower/mixed-precision ingredients and approximate preconditioners to better utilize modern hardware architectures.

Application of the presented method to a patient-specific geometry of an iliac bifurcation shows significant speed-ups, especially for higher polynomial degrees, when compared to alternative approaches with matrix-based geometric or black-box algebraic multigrid preconditioners.

1. Introduction

Implicit numerical solvers for nonlinear PDE problems in structural mechanics typically spend most of the time in the solution of linear systems of equations, arising from the linearization in Newton's method for nonlinear problems. Classical matrix-based finite element solvers assemble a (sparse) system matrix and then, in a separate step, solve the linear system. During the solution phase, matrix-vector products (or related triangular matrix solves) are usually the most important computational kernels. On current hardware, this is rooted in high memory traffic from loading matrix entries into compute units performing the arithmetic work. For higher polynomial degrees $p \geq 2$, the increased number of entries per row leads to reduced throughput per unknown. Matrix-free algorithms, on the contrary, avoid storing the system matrix explicitly to reduce memory traffic, typically yielding improved performance for $p \geq 2$. Whenever the numerical method (or the iterative techniques involved) need to evaluate the nonlinear or linearized operator, the matrix-free approach re-evaluates the spatial integrals of the finite element discretization by numerical

* Corresponding author.

E-mail address: richard.schussnig@rub.de (R. Schussnig).

<https://doi.org/10.1016/j.cma.2024.117600>

Received 22 August 2024; Received in revised form 13 November 2024; Accepted 24 November 2024

Available online 6 December 2024

0045-7825/© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

quadrature rules, which we call on-the-fly evaluation in this text. State-of-the-art matrix-free methods might further incorporate sum-factorization [1], SIMD vectorization (Single Instruction Multiple Data) over multiple elements [2,3], and parallelization via domain decomposition [4]. This can possibly lead to higher performance, as memory traffic and operation counts are significantly reduced for higher-order methods. Comparing matrix-based operator evaluation with its matrix-free counterpart adopting sum factorization with polynomial degree p and dimension d , operation counts and memory traffic reduce from an estimated $\mathcal{O}(p^{2d})$ to $\mathcal{O}(p^{d+1})$ and $\mathcal{O}(p^d)$, respectively.

In structural mechanics, recent developments include the work by Davydov et al. [5], who considered a purely displacement-based approach for finite-strain hyperelasticity in a h -multigrid (geometric multigrid) setup. Brown et al. [6] and Mehraban et al. [7] in comparison adopt a p -multigrid (polynomial) approach with algebraic multigrid (AMG) coarse grid solvers, where the former leverages GPU acceleration. In [8], Fabien presents a hybridizable discontinuous Galerkin solver for linear elasticity (in first order form), combining GPU acceleration, p -multigrid and an AMG coarse grid solver. Kiran et al. [9–11] tackle elastoplasticity with a focus on GPU implementations, whereas the resulting linear systems from lower-order discretizations are solved via Ginkgo [12] or Cusp.

The multigrid strategy used herein to precondition an outer Krylov method builds upon our previous developments [13–16], using hp -multigrid with matrix-free operator evaluation in single precision, and an AMG-preconditioned Krylov solver as coarse-level solver. The implementations of the methods presented in this work are carried out in the software project ExaDG [17] (see [15] for a comprehensive overview and [18] for the exa-scale project as a whole), which implements numerical solvers for many PDE model problems in computational fluid and structural dynamics and is based on the deal.II [19] finite-element library and in particular its matrix-free infrastructure [2,13].

Matrix-free finite element solvers for hyperelastic problems bear great potential to speed up simulations in the context of biomechanics. In fact, soft biological tissue often serves as the prime example for anisotropic hyperelastic continua, and hence, advances in solver design can significantly reduce the simulation times for patient-specific geometries of the aorta or other vessels. This in turn immediately impacts medical device design, surgery planning or enables studies on virtual cohorts to derive statistically sound biomarkers, rendering solver design and advancements highly relevant for such applications.

Constitutive modeling plays a central role in these biomedical applications. The model by Holzapfel et al. [20] is considered herein. It incorporates non-symmetrically dispersed collagen fiber families reinforcing a nearly incompressible neo-Hookean ground material, which is key to capture the load-bearing behavior of arterial tissue and additionally confronts us with new challenges regarding fast integration due to its constituents. We adopt a purely displacement-based formulation similar to [5,6], where we enforce the incompressibility constraint via a penalty term. This standard approach might suffer from locking for high bulk moduli using linear finite elements. Higher-order finite element methods significantly reduce this problem [21–24] and are hence particularly relevant for the target applications in biomechanics, as thin-walled, anisotropic and nearly incompressible structures are especially prone to locking. More involved alternative approaches are available with mixed displacement-pressure formulations [25–28], enhanced strain methods [29,30], local pressure-projection methods [31,32], or non-conforming finite elements such as Crouzeix–Raviart or DG formulations [33–36].

The contributions of this work are three-fold: First, we extend numerically stable formulations from Shakeri et al. [37] for the compressible and nearly incompressible neo-Hookean model to the tissue model and present forward stability test results. This is motivated by current hardware trends shifting towards mixed/low-precision arithmetic and the mixed-precision multigrid strategy used herein. Reformulating the weak form of standard structural mechanics problems in spatial configuration results in different formulations, which might simplify the terms to be integrated or allow for storing less linearization data at integration points. This leads to the second contribution, which is on analyzing these alternative formulations of linear(-ized) operators in terms of precomputing and storing linearization data at quadrature points for the compressible and nearly incompressible constitutive models and a fiber-reinforced tissue model [20]. Third, these ingredients are embedded into an hp -multigrid framework with matrix-free smoothing and level transfer, whose performance is compared to a matrix-based AMG preconditioner in a practically relevant problem. This last aspect thus bridges the gap between theoretical performance improvements and numerical stability considerations to the practically relevant setting, where the applicability of the introduced concepts is demonstrated.

This paper is organized as follows: Section 2 introduces finite-strain elasticity in the classical Lagrangian setting and an alternative formulation integrating over the spatial configuration. Thereafter, the relevant material models are introduced in Section 3. The related weak forms are discussed in Section 4 in terms of their numerical stability, while precomputing strategies are detailed in Section 5. The matrix-free hp -multigrid preconditioner adopting the proposed ingredients is discussed in Section 6. Numerical tests then demonstrate improved stability for the proposed weak forms in Section 7.1, evaluate the linear(-ized) operator in terms of its throughput and memory traffic in Section 7.2, and showcase the framework's applicability to a patient-specific iliac bifurcation with physiological parameters in Section 7.3. A summary and conclusions are then given in Section 8.

2. Continuum mechanics and finite element solver

This section presents standard relations from continuum mechanics to derive the nonlinear boundary value problem related to finite-strain elasticity and follows classical literature (see, e.g., [38,39]). We are interested in finding the map from the material to the spatial configuration $\phi : \Omega_0 \rightarrow \Omega_t$. It connects points $X \in \Omega_0$ to points $x \in \Omega_t$, that is, maps material coordinates in a body's material or undeformed configuration Ω_0 to spatial coordinates in the body's spatial or deformed configuration Ω_t . Introducing a displacement field u , we can express the map as

$$\phi(X) = X + u(X), \quad \text{with} \quad u(X) = x(X) - X.$$

The boundaries of both material and spatial configurations with unit outward normals \mathbf{N} and \mathbf{n} , respectively, are decomposed into non-overlapping Dirichlet and Neumann parts, $\partial\Omega_0 = \Gamma_0^D \cup \Gamma_0^N$ and $\partial\Omega_t = \Gamma_t^D \cup \Gamma_t^N$. We further introduce the deformation gradient \mathbf{F} and its determinant, referred to as the Jacobian J ,

$$\mathbf{F} := \mathbf{I} + \text{Grad } \mathbf{u}, \quad J := \det \mathbf{F}.$$

Within this contribution, we denote with $\text{Grad}(\cdot)$ the gradient with respect to the material coordinates $\mathbf{X} \in \Omega_0$, while its spatial counterpart, $\text{grad}(\cdot)$, denotes the gradient with respect to spatial coordinates $\mathbf{x} \in \Omega_t$.

Following the classical Lagrangian approach, the static linear momentum balance in material configuration Ω_0 reads

$$-\text{Div } \mathbf{P} = \mathbf{B}(\mathbf{X}) \quad \text{in } \Omega_0, \quad (1)$$

with $\text{Div}(\cdot)$ being the divergence with respect to \mathbf{X} , a given body force $[\mathbf{B} \in L^2(\Omega_0)]^d$, the first Piola–Kirchhoff stress tensor $\mathbf{P} := \mathbf{F}\mathbf{S}$, and the second Piola–Kirchhoff stress tensor \mathbf{S} . Eq. (1) is independent of the material model and the underlying constitutive relation. We further introduce strain measures being the right Cauchy–Green tensor \mathbf{C} and the Green–Lagrange strain tensor \mathbf{E} ,

$$\mathbf{C} := \mathbf{F}^T \mathbf{F}, \quad \mathbf{E} := 1/2 (\mathbf{C} - \mathbf{I}),$$

to define $\mathbf{S}(\mathbf{E})$ or $\mathbf{S}(\mathbf{C})$, capturing the material behavior, see Section 3. Finally, Eq. (1) is equipped with suitable Dirichlet and Neumann boundary conditions, $\mathbf{u}|_{\Gamma_0^D} = \mathbf{g}_D \in [H^{1/2}(\Gamma_0^D)]^d$ and $\mathbf{P}\mathbf{N}|_{\Gamma_0^N} = \mathbf{h}_N \in [H^{-1/2}(\Gamma_0^N)]^d$, to close the system. Employing a standard displacement-based finite element formulation, the residual in weak form then reads

$$r_{\Omega_0}(\mathbf{v}, \mathbf{u}) := (\text{Grad } \mathbf{v}, \mathbf{P})_{\Omega_0} - (\mathbf{v}, \mathbf{h}_N)_{\Gamma_0^N} - (\mathbf{v}, \mathbf{B})_{\Omega_0}, \quad (2)$$

where $(\cdot, \cdot)_{\Omega_0}$ denotes the standard inner product of the two arguments integrated over the given domain. To state the weak forms related to the nonlinear finite strain elasticity problem, define the vector-valued Sobolev spaces

$$H_{g_D}^1(\Omega_0) := \left\{ \mathbf{v} \in [H^1(\Omega_0)]^d : \mathbf{v}|_{\Gamma_0^D} = \mathbf{g}_D \right\} \quad \text{and} \quad H_0^1(\Omega) := \left\{ \mathbf{v} \in [H^1(\Omega)]^d : \mathbf{v}|_{\Gamma_0^D} = \mathbf{0} \right\}$$

of square integrable functions with square integrable first derivatives on Ω_0 . Similar definitions for respective counterparts defined on Ω_t are omitted for brevity. In material configuration, this leads to

Problem 2.1. Find $\mathbf{u} \in H_{g_D}^1(\Omega_0)$, such that

$$r_{\Omega_0}(\mathbf{v}, \mathbf{u}) = 0 \quad \forall \mathbf{v} \in H_0^1(\Omega_0). \quad (3)$$

Regarding well-posedness of Problem 2.1 under suitable assumptions, we refer the reader to [40–42], noting that the existence of minimizers of the related energy functional cannot be guaranteed for general hyperelastic materials and load configurations. Herein, we employ standard C^0 -continuous finite-dimensional subspaces and Newton’s method to solve the nonlinear problem (3). In Newton’s method, each iteration updates the initial guess \mathbf{u}_0 via $\mathbf{u}_{k+1} = \mathbf{u}_k + \Delta \mathbf{u}$, $k = 0, \dots, N_{\max}$ starting from \mathbf{u}_0 with $\mathbf{u}_0|_{\Gamma_0^D} = \mathbf{g}_D$. The increment $\Delta \mathbf{u}$ is obtained by solving

Problem 2.2. Find $\Delta \mathbf{u} \in H_0^1(\Omega_0)$ given the previous iterate $\mathbf{u}_k \in H_{g_D}^1(\Omega_0)$ by solving

$$D_{\mathbf{u}} r_{\Omega_0}(\mathbf{v}, \Delta \mathbf{u})|_{\mathbf{u}_k} = -r_{\Omega_0}(\mathbf{v}, \mathbf{u}_k) \quad \forall \mathbf{v} \in H_0^1(\Omega_0). \quad (4)$$

The directional derivative corresponding to Eq. (2) reads

$$D_{\mathbf{u}} r_{\Omega_0}(\mathbf{v}, \Delta \mathbf{u})|_{\mathbf{u}_k} = (\text{Grad } \mathbf{v}, (D_{\mathbf{u}} \mathbf{F}) \mathbf{S} + \mathbf{F} D_{\mathbf{u}} \mathbf{S})_{\Omega_0}, \quad (5)$$

where all terms are evaluated using \mathbf{u}_k . A basic Newton method is provided in Algorithm 1, which for the sake of brevity does not contain a line-search algorithm or load stepping procedure. It highlights the update of stored quadrature point data of central interest within this work, see Section 5.

Algorithm 1 Generic Newton’s method to solve $\mathbf{r}(\mathbf{u}) = \mathbf{0}$ with quadrature point data update.

```

1: function NEWTONSOLVER( $\mathbf{u}_0$ ,  $\epsilon_{\text{abs}}$ ,  $\epsilon_{\text{rel}}$ ,  $N_{\max}$ )
2:    $\mathbf{u}_0 \leftarrow \mathbf{g}_D$  on  $\Gamma_0^D$  ▷ initialize iterate and enforce Dirichlet conditions
3:    $k \leftarrow 0$  ▷ initialize counter
4:   while  $\|\mathbf{r}(\mathbf{u}_k)\| > \epsilon_{\text{abs}}$  and  $\|\mathbf{r}(\mathbf{u}_k)\| > \epsilon_{\text{rel}} \|\mathbf{r}(\mathbf{u}_0)\|$  and  $k < N_{\max}$  do
5:     update quadrature point data ▷ see Sec. 5
6:     solve for Newton update:  $\mathbf{K} \Delta \mathbf{u} = -\mathbf{r}(\mathbf{u}_k)$ 
7:      $\mathbf{u}_k \leftarrow \mathbf{u}_k + \Delta \mathbf{u}$  ▷ update iterate
8:      $k \leftarrow k + 1$  ▷ update iteration counter
9:   return  $\mathbf{u}_k$ ,  $k - 1$  ▷ return last iterate and number of iterations

```

The integrals in the weak form of the Newton update step (4) can be equivalently written as integrals over the spatial configuration Ω_t . Specifically, one may rewrite the domain integral involving the stress, $(\text{Grad } \mathbf{v}, \mathbf{P})_{\Omega_0}$, keeping the Neumann data and body force unchanged. Introducing the Kirchhoff stress $\boldsymbol{\tau} := \mathbf{F} \mathbf{S} \mathbf{F}^T$ yields

$$(\text{Grad } \mathbf{v}, \mathbf{P})_{\Omega_0} = (\text{Grad } \mathbf{v}, \boldsymbol{\tau} \mathbf{F}^{-T})_{\Omega_0} = ((\text{Grad } \mathbf{v}) \mathbf{F}^{-1}, \boldsymbol{\tau})_{\Omega_0} = (\text{grad } \mathbf{v}, \boldsymbol{\tau})_{\Omega_0} = (1/J \text{ grad } \mathbf{v}, \boldsymbol{\tau})_{\Omega_t} = (1/J \text{ grad}^S \mathbf{v}, \boldsymbol{\tau})_{\Omega_t},$$

using $\mathbf{P} = \mathbf{F} \mathbf{S}$, $\text{grad } \mathbf{v} = (\text{Grad } \mathbf{v}) \mathbf{F}^{-1}$, the notation $(\cdot)^S := 1/2 [(\cdot) + (\cdot)^T]$ for the symmetric part of a tensor, and symmetry of $\boldsymbol{\tau}$ in the last step. The residual then reads

$$r_{\Omega_t}(\mathbf{v}, \mathbf{u}) := (1/J \text{ grad}^S \mathbf{v}, \boldsymbol{\tau})_{\Omega_t} - (\mathbf{v}(\boldsymbol{\phi}(\mathbf{X})), \mathbf{h}_N)_{\Gamma_0^N} - (\mathbf{v}(\boldsymbol{\phi}(\mathbf{X})), \mathbf{B})_{\Omega_0}, \quad (6)$$

where $\mathbf{v}(\boldsymbol{\phi}(\mathbf{X}))$ is equivalent to the test functions in the reference domain Ω_0 and independent of \mathbf{u} . The spatial counterpart of the weak form of the nonlinear finite strain elasticity problem, Problem 2.1, reads

Problem 2.3. Find $\mathbf{u} \in H_{g_D}^1(\Omega_t)$, such that

$$r_{\Omega_t}(\mathbf{v}, \mathbf{u}) = 0 \quad \forall \mathbf{v} \in H_0^1(\Omega_t).$$

Integrating over the spatial domain Ω_t , the increment $\Delta \mathbf{u}$ is obtained in each iteration of Newton's method by solving

Problem 2.4. Find $\Delta \mathbf{u} \in H_0^1(\Omega_t)$ given the previous iterate $\mathbf{u}_k \in H_{g_D}^1(\Omega_t)$ by solving

$$D_{\mathbf{u}} r_{\Omega_t}(\mathbf{v}, \Delta \mathbf{u}) \Big|_{\mathbf{u}_k} = -r_{\Omega_t}(\mathbf{v}, \mathbf{u}_k) \quad \forall \mathbf{v} \in H_0^1(\Omega_t).$$

The directional derivatives of the stress term are given by

$$D_{\mathbf{u}} (1/J \text{ grad}^S \mathbf{v}, \boldsymbol{\tau})_{\Omega_t} = D_{\mathbf{u}} (\text{grad}^S \mathbf{v}, \boldsymbol{\tau})_{\Omega_0} = (\text{grad}^S \mathbf{v}, D_{\mathbf{u}} \boldsymbol{\tau})_{\Omega_0} + ((\text{Grad } \mathbf{v}) D_{\mathbf{u}} \mathbf{F}^{-1}, \boldsymbol{\tau})_{\Omega_0}.$$

Introducing the contravariant push-forward of the fourth-order material part of the stiffness tensor denoted as $J\mathbb{C}$ [38,43] and transformed onto Ω_t , this can be rewritten as

$$(\text{grad } \mathbf{v}, J\mathbb{C} : \text{grad}^S \Delta \mathbf{u})_{\Omega_0} + (\text{grad } \mathbf{v}, (\text{grad } \Delta \mathbf{u}) \boldsymbol{\tau})_{\Omega_0} = (1/J \text{ grad } \mathbf{v}, J\mathbb{C} : \text{grad}^S \Delta \mathbf{u} + (\text{grad } \Delta \mathbf{u}) \boldsymbol{\tau})_{\Omega_t}. \quad (7)$$

Note that the actual fourth-order tensor \mathbb{C} need not be computed, but rather its action on a symmetric second-order tensor to evaluate $J\mathbb{C} : (\cdot)^S$ in Eq. (7). Again, (6) and (7) are independent of the material model, which enters via $\boldsymbol{\tau}$ and \mathbb{C} to be discussed in Section 3.

Following the strategy of spatial integration comes at the cost of updating the vertex positions of the finite element mesh and related data. When updating the quadrature point data in Algorithm 1, the spatial grid, which approximates the body in its deformed configuration Ω_t , is updated simultaneously.

3. Constitutive modeling

The second Piola–Kirchhoff stress tensor in the residuals (2) or (6) is defined in terms of a strain measure, i.e., $\mathbf{S} = \mathbf{S}(\mathbf{C})$ or $\mathbf{S} = \mathbf{S}(\mathbf{E})$. For hyperelastic continua, the constitutive relation is expressed in terms of the strain–energy density Ψ (per unit reference volume),

$$\mathbf{F}^{-1} \mathbf{P} = \mathbf{S} := \frac{\partial \Psi(\mathbf{E})}{\partial \mathbf{E}} = 2 \frac{\partial \Psi(\mathbf{C})}{\partial \mathbf{C}}.$$

A compressible neo-Hookean model (cNH) is given by [44],

$$\Psi_{\text{cNH}}(\mathbf{C}) = \mu/2 (I_1 - \text{tr } \mathbf{I} - 2 \ln J) + \lambda \ln^2 J, \quad \mathbf{S}_{\text{cNH}} = \mu \mathbf{I} - (\mu - 2\lambda \ln J) \mathbf{C}^{-1}, \quad (8)$$

with the first invariant $I_1 := \text{tr } \mathbf{C}$, and the material parameters being the shear modulus μ and Lamé coefficient λ . For the nearly incompressible neo-Hookean model (iNH) yielding $J \approx 1$, the deformation gradient is split into isochoric and volumetric parts according to Flory [45], yielding

$$\Psi_{\text{iNH}}(\mathbf{C}) = \mu/2 (J^{-2/3} I_1 - \text{tr } \mathbf{I}) + \kappa_b/4 (J^2 - 1 - 2 \ln J), \quad \mathbf{S}_{\text{iNH}} = \mu J^{-2/3} \mathbf{I} + [\kappa_b/2 (J^2 - 1) - \mu/3 J^{-2/3} I_1] \mathbf{C}^{-1}, \quad (9)$$

with bulk modulus κ_b enforcing $J = 1$ as $\kappa_b \rightarrow \infty$, acting as a penalty term.

Now, for the target applications in biomedical engineering and medicine, more involved constitutive relations are required to capture the material behavior. Aortic tissue of prime interest within this work shows an anisotropic stiffening effect under large strains due to collagen fibers reinforcing the ground material. The model by Holzapfel et al. [20], herein simply referred to as the fiber model (fiber), adds exponential terms to the strain–energy density. The main motivation for this choice lies in the fact that we aim to showcase and analyze the potential performance impact in realistic scenarios. Our performance improvements hence directly translate to problems of high practical relevance, which has not been addressed in literature so far. The model includes two fiber families and allows accounting for the non-symmetric fiber dispersion, which is more significant in the tangential plane compared

Table 1

Parameters for the nearly incompressible fiber model [20] fit to material tests of aortic medial tissue, taken from [52,53]. The bulk modulus κ_b , i.e., the penalty term in purely displacement-based formulations corresponds to a Poisson's ratio of 0.49.

μ [kPa]	κ_b [kPa]	a [–]	b [–]	k_1 [kPa]	k_2 [–]	Φ [°]	H_{11} [–]	H_{22} [–]	H_{33} [–]
62.1	3084.3	3.62	34.3	1.4	22.1	27.47	0.9168	0.0759	0.0073

to the out-of-plane direction [46]. The strain-energy density combines the neo-Hookean ground material Ψ_{INH} given in Eq. (9) and collagen fiber contributions via Ψ_c as

$$\Psi_{\text{fiber}}(\mathbf{C}, \mathbf{H}_i) = \Psi_{\text{INH}}(\mathbf{C}) + \sum_{i=4,6} \Psi_c(\mathbf{C}, \mathbf{H}_i) = \Psi_{\text{INH}}(\mathbf{C}) + \sum_{i=4,6} \begin{cases} \frac{k_1}{2k_2} [\exp(k_2 E_i^2) - 1] & \text{if } I_i^* > 1, \\ 0 & \text{else,} \end{cases} \quad (10)$$

where k_1 is a stiffness-like parameter and k_2 a dimensionless shape parameter. The strain energy related to collagen fibers Ψ_c only contributes to the total energy Ψ , when the squared fiber stretches defined as

$$I_i^* = (\mathbf{M}_i \otimes \mathbf{M}_i) : \mathbf{C}, \quad i = 4, 6,$$

signal tension, i.e., when $I_4^* > 1$ or $I_6^* > 1$, while fiber bundles with compressed mean fiber buckle immediately (cf. [47]). The symmetric structure tensor \mathbf{H}_i and strain-like quantity E_i are defined as

$$\mathbf{H}_i = H_{11} \mathbf{M}_1 \otimes \mathbf{M}_1 + H_{22} \mathbf{M}_2 \otimes \mathbf{M}_2 + H_{33} \mathbf{M}_3 \otimes \mathbf{M}_3, \quad E_i = \mathbf{H}_i : (\mathbf{C} - \mathbf{I}) = \text{tr}(\mathbf{H}_i \mathbf{C}) - 1. \quad (11)$$

For the targeted applications in biomechanics and medicine, the orthonormal basis spanned by \mathbf{M}_1 , \mathbf{M}_2 and \mathbf{M}_3 is related to tailored material coordinate systems in the reference configuration [48–50]. Given a suitable material coordinate system, the mean in-plane and out-of-plane angles Φ and $\Theta = 0$ then yield

$$\mathbf{M}_1 = E_1 \cos \Phi - E_1 \sin \Phi, \quad \mathbf{M}_2 = E_2 \sin \Phi + E_2 \cos \Phi, \quad \mathbf{M}_3 = E_3,$$

describing the classical helical patterns of collagen fibers in vascular tissue. Here, E_1 aligns with the circumferential direction, E_2 with the longitudinal direction, and E_3 with the radial direction. This specific representation of \mathbf{H}_i in (11) assumes a bivariate von Mises distribution of the fiber density [20]. Multiplicative decomposition then leads to only three nonzero components of the generalized structure tensor in the reference configuration,

$$H_{11} = \frac{1 - H_{33}}{2} \left(1 + \frac{I_1(a)}{I_0(a)} \right), \quad H_{22} = \frac{1 - H_{33}}{2} \left(1 - \frac{I_1(a)}{I_0(a)} \right), \quad H_{33} = \frac{1}{4b} - \frac{\exp(-2b)}{\sqrt{2\pi b} \text{erf}(\sqrt{2b})},$$

where I_0 and I_1 denote the Bessel functions of the first kind of orders 0 and 1, respectively, and $\text{erf}(\cdot)$ is the error function, see [51] for details. Thus, the parameters a and b together with the mean in-plane angle Φ describe the dispersion of the collagen fibers based on the material coordinate system spanned by E_1 , E_2 and E_3 .

This leads to the second Piola–Kirchhoff stress tensor $\mathbf{S}_{\text{fiber}}$ as a sum of the collagen fiber \mathbf{S}_c and nearly incompressible neo-Hookean \mathbf{S}_{INH} (9) contributions, that is

$$\mathbf{S}_{\text{fiber}} = \mathbf{S}_{\text{INH}}(\mathbf{C}) + \mathbf{S}_c(\mathbf{C}, \mathbf{H}_i) = \mathbf{S}_{\text{INH}}(\mathbf{C}) + \sum_{i=4,6} 2k_1 \exp(k_2 E_i^2) E_i \mathbf{H}_i,$$

where physiological tissue parameters used in this work [52,53] are summarized in Table 1. The second Piola–Kirchhoff stress tensor and the related directional derivative are the only ingredients in the Newton solver specific to the constitutive model and can be found in Appendix.

4. Stable and fast numerics for hyperelasticity

The momentum balance residuals (2) and (6) and their respective directional derivatives feature terms that may suffer from significant numerical instability. Especially when considering ongoing and expected future hardware-driven developments towards mixed- and low-precision strategies [54], the classical formulations have to be reviewed. In the following, we extend numerically stable forms of the strain energy density and the stresses from Shakeri et al. [37] towards the fiber model and furthermore present stable forms of the directional derivatives, which have not been reported in literature yet. Furthermore, we present a fast evaluation scheme for $J^{-2/3}$ based on Newton's method.

Starting with the strain measures and their evaluation, the Green–Lagrange strain tensor defined as

$$\mathbf{E} := 1/2 (\mathbf{C} - \mathbf{I}) = 1/2 (\mathbf{F}^T \mathbf{F} - \mathbf{I}),$$

shows cancellation in the small strain regime since components of \mathbf{F} , which is close to the unit tensor \mathbf{I} , are used in floating-point operations, followed by a subtraction of the unit tensor. To reduce loss of accuracy, we evaluate \mathbf{E} as

$$\mathbf{E} = 1/2 (\text{Grad } \mathbf{u} + \text{Grad}^T \mathbf{u} + \text{Grad}^T \mathbf{u} \text{ Grad } \mathbf{u}). \quad (12)$$

Similar reasoning lies behind computing the Green–Euler strain tensor according to

$$\tilde{\mathbf{b}} := 1/2 (\mathbf{b} - \mathbf{I}) = 1/2 (\mathbf{F} \mathbf{F}^T - \mathbf{I}) = 1/2 (\text{Grad } \mathbf{u} + \text{Grad}^T \mathbf{u} + \text{Grad } \mathbf{u} \text{ Grad}^T \mathbf{u}),$$

where the left Cauchy–Green tensor $\mathbf{b} := \mathbf{F}\mathbf{F}^T$, suffering from similar cancellation for $\mathbf{F} \approx \mathbf{I}$, is avoided. Reformulations of the strain measures also affect the fiber invariants, which are evaluated as

$$E_i := \mathbf{H}_i : (\mathbf{C} - \mathbf{I}) = 2 \mathbf{H}_i : \mathbf{E}, \quad I_i^* := (\mathbf{M}_1 \otimes \mathbf{M}_1) : \mathbf{C} = 2 (\mathbf{M}_1 \otimes \mathbf{M}_1) : \mathbf{E} + \text{tr} (\mathbf{M}_1 \otimes \mathbf{M}_1).$$

Further recurring terms in finite-strain (hyper-)elasticity are, e.g., the inverse right Cauchy–Green tensor \mathbf{C}^{-1} . Cancellation can be reduced by inverting \mathbf{F} , where $\text{cond}(\mathbf{F}) = \sqrt{\text{cond}(\mathbf{C})}$, and computing $\mathbf{C}^{-1} = \mathbf{F}^{-1} \mathbf{F}^{-T}$, rather than inverting \mathbf{C} . Note that \mathbf{F}^{-1} is often needed anyways, such that this approach also requires fewer arithmetic operations.

The stress tensor of nearly incompressible continua and its derivative also contain $(J-1)$ or similar terms. Since we aim to enforce $J = 1$ through a penalty term scaled by the bulk modulus κ_b as in Eq. (9), we inevitably have $J \approx 1$ as the bulk modulus increases. As pointed out by Shakeri et al. [37], numerical stability can be improved significantly by introducing $J_{-1} := J - 1 = \det(\mathbf{F}) - 1$ and avoiding forming $\det(\mathbf{F})$ first and subtracting 1 from the result, but instead using the components of $\text{Grad } \mathbf{u}$ directly. This can further be exploited to rewrite $(J^2 - 1)$ in a more stable manner as $J_{-1}(J_{-1} + 2)$.

Compressible and (nearly) incompressible hyperelastic continua also contain more complex functions taking the Jacobian as an argument. Within the current work, $\ln(J)$ and $J^{-2/3}$ are of interest both in terms of numerical stability and computational efficiency. Herein, SIMD vectorization is enabled straight-forwardly via polynomial or rational approximations to process data on all lanes of a SIMD vector. In addition to that, restricting input argument ranges allows for further optimizations. For $\ln(J)$, we define [37,55]

$$\ln(J) = \ln(J_{-1} + 1) = \ln_{+1}(J_{-1}), \quad \ln_{+1}(x) := 2 \sum_{n=0}^{\infty} \frac{1}{2n+1} \left(\frac{x}{2+x} \right)^{2n+1}.$$

Note that the sum only contains odd powers of $x \in (-1, \infty)$, such that only terms with the same sign are added, guaranteeing numerical stability when summing a fixed number of terms small to large.

For (nearly) incompressible continua, the term $J^{-2/3}$ plays a central role, as its evaluation can be costly. Possible options are: (i) looping over SIMD vectors and resorting to standard techniques for scalar types, (ii) exploiting the floating point representation and approximation via summation (see [56]), or (iii) a Newton solver given a good initial guess $J \approx 1 \Leftrightarrow J^{-2/3} \approx 1$. Depending on the storage strategy used to evaluate the integrals, we consider (iii) for on-the-fly evaluation, but (i) or (ii) in case $J^{-2/3}$ is stored anyways, see Section 5.

A Newton solver for $x = J^{-2/3} \Leftrightarrow f(x) := x^{-3} - J^2 = 0$ uses $f'(x) = -3x^{-4}$, leading to

$$x_{k+1} = x_k - \frac{f(x_k)}{f'(x_k)} = 1/3 (4x_k - J^2 x_k^4) \quad (13)$$

for $k = 0, \dots, N$ starting from $x_0 = 1$ or some previously computed $J^{-2/3}$. Since for the undeformed initial state of the elastic structure we have $\mathbf{u} = \mathbf{0} \Rightarrow \mathbf{F} = \mathbf{I} \Rightarrow J = 1$ and due to incompressibility, $J \approx 1$ holds throughout the entire motion, we employ a fixed number of 3 Newton iterations according to Algorithm 2. This is justified by local quadratic convergence of Newton's method and an excellent initial guess exploiting $J^{-2/3} \approx 1$.

Algorithm 2 Approximation of $J^{-2/3}$ using N Newton steps (13)

```

1: function FASTAPPROXJPOW( $J_{-1}$ ,  $N$ )
2:    $\alpha = 1/3 [(J_{-1} + 2) J_{-1} + 1]$   $\triangleright$  store  $1/3 J^2$ 
3:    $\beta = 4/3 - \alpha$   $\triangleright$  assignment is first iteration w.  $x_0 = 1$ 
4:   for  $n = 1, \dots, N - 1$  do
5:      $\gamma = \beta^2 \beta^2$   $\triangleright \gamma = x_k^4$  by repeated squaring
6:      $\beta \leftarrow 4/3 \beta - \alpha \gamma$ 
7:   return  $\beta$ 

```

Similar to $\ln(J)$ and $J^{-2/3}$, the last ingredient required with regards to the material models considered within this work is a SIMD-compatible $\exp(x)$ appearing in the fiber model (10). Here, we adopt the approach by Proell et al. [57], which is based on [56,58,59], and exploits the floating point representation for fast evaluation.

5. Storage strategies for hyperelasticity

Contrary to matrix-based algorithms, the integrals involved in the residual and its linearization are evaluated repeatedly. For achieving optimal performance, the memory access and arithmetic workload of these operations need to be compared to the capabilities of the underlying hardware. Depending on the material models' complexity, precomputing and storing data on the integration point level might be beneficial if the in-core resources such as arithmetic units or available registers are highly busy. More precisely, we aim to identify certain data that we compute once per nonlinear Newton iteration, store it at the integration points, and load it during operator evaluation, see Algorithm 1, whereas other quantities are computed repeatedly in each matrix-vector product. Based on the observations from Davydov et al. [5], we introduce three stages in this regard: (i) on-the-fly integral evaluation of all terms, (ii) precompute and store scalar quantities where useful, and (iii) precompute and store scalars and second-order tensors where useful. Additionally, the symmetry of stresses and strains and certain intermediate quantities is exploited both in

terms of memory consumption and when performing operations such as addition, multiplication, double contraction, push-forward and others.

In the following, the final weak forms as derived in [Appendix](#) are presented in stable form besides the classical form, while quantities to be stored are highlighted as $\underline{(\cdot)}$ if they are scalar, and as $\underline{\underline{(\cdot)}}$ if they are second-order (possibly symmetric) tensors. Here, the stable form of the fiber model, the linearizations using integration over the spatial configuration, and the storage strategies in particular are novel contributions. Common to all constitutive models is the Newton update step in Algorithm 1. In material configuration, we seek $\Delta \mathbf{u} \in H_0^1(\Omega_0)$, such that there holds

$$\left(\text{Grad } \mathbf{v}, (\text{Grad } \Delta \mathbf{u}) \underline{\underline{\mathbf{S}}} + \underline{\underline{\mathbf{F}}} \text{D}_u \mathbf{S} \right)_{\Omega_0} = (\mathbf{v}, \mathbf{h}_N)_{\Omega_0} - \left(\text{Grad } \mathbf{v}, \underline{\underline{\mathbf{F}}} \underline{\underline{\mathbf{S}}} \right)_{\Omega_0} + (\mathbf{v}, \mathbf{B})_{\Omega_0} \quad \forall \mathbf{v} \in H_0^1(\Omega_0),$$

with the directional derivative $\text{D}_u \mathbf{S}$ defined shortly. As indicated, the tensors \mathbf{S} (symmetric) and \mathbf{F} can be precomputed and stored. Note that the additional tensors and local operations only involve a modest number of simple operations, such as multiplications and additions, to be performed in local variables, typically mapped to registers for execution on hardware. These are cheaper than memory accesses on all temporary hardware architectures. The alternative approach integrating over the spatial domain requires solving for $\Delta \mathbf{u} \in H_0^1(\Omega_t)$, such that there holds

$$\left(\frac{1}{J} \text{grad } \mathbf{v}, J \mathbb{C} : \text{grad}^S \Delta \mathbf{u} + \text{grad } \Delta \mathbf{u} \underline{\underline{\boldsymbol{\tau}}} \right)_{\Omega_t} = (\mathbf{v}(\boldsymbol{\phi}^{-1}), \mathbf{h}_N)_{\Gamma_0^N} - \left(\frac{1}{J} \text{grad}^S \mathbf{v}, \underline{\underline{\boldsymbol{\tau}}} \right)_{\Omega_t} + (\mathbf{v}(\boldsymbol{\phi}^{-1}), \mathbf{B})_{\Omega_0} \quad \forall \mathbf{v} \in H_0^1(\Omega_t),$$

where we store the scalar quantity $1/J$ as well as the symmetric tensor $\boldsymbol{\tau}$. The stress tensors and their derivatives depend on the material model given in the following.

Compressible neo-Hookean model in material configuration.

$$\begin{aligned} \mathbf{S}_{\text{cNH}} &= \mu \mathbf{I} - (\mu - 2\lambda \ln J) \mathbf{C}^{-1} \stackrel{d=3}{=} \mathbf{C}^{-1} \left(2\mu \mathbf{E} + 2\lambda \mathbf{I} \ln_{+1} J_{-1} \right) \\ \text{D}_u \mathbf{S}_{\text{cNH}} &= (\mu - 2\lambda \ln J) 2 \left(\underline{\underline{\mathbf{F}}}^{-1} (\text{Grad } \Delta \mathbf{u}) \underline{\underline{\mathbf{C}}}^{-1} \right)^S + 2\lambda \text{tr} \left(\underline{\underline{\mathbf{F}}}^{-1} \text{Grad } \Delta \mathbf{u} \right) \underline{\underline{\mathbf{C}}}^{-1} \\ &= (\mu - 2\lambda \ln_{+1} J_{-1}) 2 \left(\underline{\underline{\mathbf{F}}}^{-1} (\text{Grad } \Delta \mathbf{u}) \underline{\underline{\mathbf{C}}}^{-1} \right)^S + 2\lambda \text{tr} \left(\underline{\underline{\mathbf{F}}}^{-1} \text{Grad } \Delta \mathbf{u} \right) \underline{\underline{\mathbf{C}}}^{-1} \end{aligned} \quad (14)$$

The notation $\stackrel{d=3}{=}$ indicates that we assume $d = 3$ for the sake of presentation, since a term of the form $\mathbf{I}(1 - d/3)$ cancels here and at similar places in this manuscript.

Spatial integration of the compressible neo-Hookean model.

$$\begin{aligned} \boldsymbol{\tau}_{\text{cNH}} &= \mu \mathbf{b} - (\mu - 2\lambda \ln J) \mathbf{I} \stackrel{d=3}{=} 2\mu \tilde{\mathbf{b}} + 2\lambda \mathbf{I} \ln_{+1} J_{-1} \\ J \mathbb{C}_{\text{cNH}} : (\cdot)^S &= 2(\mu - 2\lambda \ln J) (\cdot)^S + 2\lambda \text{tr}(\cdot) \mathbf{I} = 2 \left((\mu - 2\lambda \ln_{+1} J_{-1}) (\cdot)^S + 2\lambda \text{tr}(\cdot) \mathbf{I} \right) \end{aligned} \quad (15)$$

Nearly incompressible neo-Hookean model in material configuration.

$$\begin{aligned} \mathbf{S}_{\text{iNH}} &= \mu \underline{\underline{J}}^{-2/3} \mathbf{I} + c_1 \mathbf{C}^{-1} \stackrel{d=3}{=} \mathbf{C}^{-1} \left[\kappa_b/2 \underline{\underline{J}}_{-1} (\underline{\underline{J}}_{-1} + 2) \mathbf{I} + 2\mu \underline{\underline{J}}^{-2/3} (\mathbf{E} - 1/3 \mathbf{I} \text{tr } \mathbf{E}) \right] \\ \text{D}_u \mathbf{S}_{\text{iNH}} &= -2\mu/3 \underline{\underline{J}}^{-2/3} (1/J \text{D}_u J) \mathbf{I} + 2c_1 \left[\underline{\underline{\mathbf{F}}}^{-1} (\text{Grad } \Delta \mathbf{u}) \underline{\underline{\mathbf{C}}}^{-1} \right]^S + \left[c_2 (1/J \text{D}_u J) - 2\mu/3 \underline{\underline{J}}^{-2/3} \text{tr} \left(\underline{\underline{\mathbf{F}}}^T \text{Grad } \Delta \mathbf{u} \right) \right] \underline{\underline{\mathbf{C}}}^{-1} \end{aligned}$$

The additional scalars are

$$c_1 := \kappa_b/2(J^2 - 1) - \mu/3 \underline{\underline{J}}^{-2/3} I_1 = \kappa_b/2 \underline{\underline{J}}_{-1} (\underline{\underline{J}}_{-1} + 2) - \mu/3 \underline{\underline{J}}^{-2/3} (d + 2 \text{tr } \mathbf{E}), \quad (16)$$

$$c_2 := 2\mu/9 \underline{\underline{J}}^{-2/3} I_1 + \kappa_b J^2 = 2\mu/9 \underline{\underline{J}}^{-2/3} (d + 2 \text{tr } \mathbf{E}) + \kappa_b J^2. \quad (17)$$

Spatial integration of the nearly incompressible neo-Hookean model.

$$\begin{aligned} \boldsymbol{\tau}_{\text{iNH}} &= \mu \underline{\underline{J}}^{-2/3} \mathbf{b} + c_1 \mathbf{I} \stackrel{d=3}{=} \kappa_b/2 \underline{\underline{J}}_{-1} (\underline{\underline{J}}_{-1} + 2) \mathbf{I} + 2\mu \underline{\underline{J}}^{-2/3} (\tilde{\mathbf{b}} - 1/3 \mathbf{I} \text{tr } \tilde{\mathbf{b}}), \\ J \mathbb{C}_{\text{iNH}} : (\cdot)^S &= -4\mu/3 \underline{\underline{J}}^{-2/3} \text{tr}(\cdot) \underline{\underline{\mathbf{C}}} - 2c_1 (\cdot)^S + c_2 \text{tr}(\cdot) \mathbf{I}, \end{aligned}$$

where we exploit $\text{tr } \mathbf{C} = \text{tr } \mathbf{b} = 2\text{tr } \tilde{\mathbf{b}} + d$ and use c_1 (16) and c_2 (17).

Nearly incompressible fiber model in material configuration.

$$\mathbf{S}_{\text{fiber}} = \mathbf{S}_{\text{iNH}} + \sum_{i=4,6} c_3 \frac{E_i}{E_i} \underline{\underline{\mathbf{H}}}_i, \quad \text{D}_u \mathbf{S}_{\text{fiber}} = \text{D}_u \mathbf{S}_{\text{iNH}} + \sum_{i=4,6} c_3 \left(2k_2 E_i^2 + 1 \right) \left[\underline{\underline{\mathbf{H}}}_i : \left(2\underline{\underline{\mathbf{F}}}^T \text{Grad } \Delta \mathbf{u} \right)^S \right] \underline{\underline{\mathbf{H}}}_i$$

with an additional scalar

$$c_3 := \begin{cases} 2k_1 \exp(k_2 E_i^2) & \text{if } I_i^* > 1, \\ 0 & \text{else.} \end{cases} \quad (18)$$

Table 2

Quantities loaded per integration point for the variants storing scalars or scalars and tensors with related memory using double precision and $d = 3$. The Jacobian matrices of the finite element maps, \mathbf{J}_0 and \mathbf{J}_i , are stored for each domain. \mathbf{u}_k refers to nodal DoF vector entries loaded and interpolated (not stored as vector at integration point level). \mathbf{S} , $\boldsymbol{\tau}$, \mathbf{C} , \mathbf{C}^{-1} , \mathbf{H}_i and $\mathbf{F}\mathbf{H}_i\mathbf{F}^T$, $i = 4, 6$ are symmetric. Numbers in brackets indicate contribution to asymptotic memory traffic in linearized operator application where differing from storage requirement.

	Integration over reference domain Ω_0				Integration over spatial domain Ω_i			
	Quantities		Memory in B		Quantities		Memory in B	
	Scalar	Tensor	Scalar	Tensor	Scalar	Tensor	Scalar	Tensor
cNH	$\mathbf{J}_0, \mathbf{u}_k, \ln J$	$\mathbf{F}, \mathbf{S}, \mathbf{F}^{-1}, \mathbf{C}^{-1}$	104	320	$\mathbf{J}_0, \mathbf{J}_i, \mathbf{u}_k, 1/J, \ln J$	$\boldsymbol{\tau}$	184	208 (136)
iNH	$\mathbf{J}_0, \mathbf{u}_k, J_{-1}, J^{-2/3}, c_1, c_2$	$\mathbf{F}, \mathbf{S}, \mathbf{F}^{-1}, \mathbf{C}^{-1}$	128 (120)	344 (336)	$\mathbf{J}_0, \mathbf{J}_i, \mathbf{u}_k, 1/J, J_{-1}, J^{-2/3}, c_1, c_2$	$\boldsymbol{\tau}, \mathbf{C}$	208	280 (200)
fiber	iNH, $c_3, I_i^*, E_i, \mathbf{H}_i$	iNH	272 (248)	488 (464)	iNH, $c_3, I_i^*, E_i, \mathbf{H}_i$	iNH, $\mathbf{F}\mathbf{H}_i\mathbf{F}^T$	352	520 (328)




Fig. 1. Visualization of a multigrid V-cycle (left) and an exemplary hp -multigrid hierarchy (right), with degrees of freedom indicated by circles.

Spatial integration of the nearly incompressible fiber model.

$$\boldsymbol{\tau}_{\text{fiber}} = \boldsymbol{\tau}_{\text{iNH}} + \sum_{i=4,6} c_3 \underline{\underline{E_i \mathbf{F} \mathbf{H}_i \mathbf{F}^T}},$$

$$J \mathbb{C}_{\text{fiber}} : (\cdot)^S = J \mathbb{C}_{\text{iNH}} : (\cdot)^S + \sum_{i=4,6} 2c_3 \left(2k_2 \underline{\underline{E_i^2}} + 1 \right) \left[\left(\underline{\underline{\mathbf{F} \mathbf{H}_i \mathbf{F}^T}} \right) : (\cdot) \right] \underline{\underline{\mathbf{F} \mathbf{H}_i \mathbf{F}^T}},$$

with c_3 defined in (18).

An overview of all the variants and the respective precomputed and stored variables is given in Table 2, which also lists the required memory for storing in double precision (1 double-precision variable = 8 bytes, 8 B). For the fiber model, storing scalar quantities requires 248 B (material configuration) or 328 B (spatial configuration), while storing tensorial quantities requires 488 B (material configuration) or 520 B (spatial configuration) per integration point. Depending on the exact form of the integrals and arithmetic operations to evaluate them, the increased memory traffic storing second-order tensors might pay off. Furthermore, this example also shows that it is not clear per se, if integrating over the material or spatial configuration is favorable in terms of memory traffic, since this depends also on whether or not tensors or scalars are precomputed and stored. For the linearized operator, not all tensors need to be loaded when integrating over the spatial domain.

6. Matrix-free preconditioning

The linear systems corresponding to the Newton update step are solved with the preconditioned flexible generalized minimal residual method (FGMRES) [60]. We use a flexible formulation since we use a Krylov solver at the coarse level in the multigrid preconditioner, which renders the operation non-stationary. The Krylov solver only requires the action of the operator on a vector, not the explicit entries of the matrix. The matrix-free evaluation of the action of the matrix on a vector is realized via numerical quadrature, exploiting the tensor product structure of the shape functions and quadrature rule via sum factorization techniques (see [2,3,61]) and employing SIMD vectorization over batches of elements [2].

For constructing a preconditioner that is compatible with the matrix-free evaluation, an hp -multigrid preconditioning strategy with matrix-free smoothers is adopted from Fehn et al. [14], where first the polynomial degree is lowered recursively from p to $\lfloor p/2 \rfloor$, that is halved and rounded down to the nearest integer, going from level l to $l-1$, after which h -coarsening is performed, see Fig. 1. This strategy is denoted as ph -multigrid in [14] and ExaDG. The outer Krylov solver operates in double precision, while the preconditioner operates in single precision, see [2,62]. The individual h -levels are created by uniformly refining an initial coarse grid, where the finest grid is equipped with a potentially higher-order accurate mapping. This higher-order mapping is then interpolated to the coarser levels keeping the order of the mapping constant, i.e., equal to the finest grids' polynomial degree. The mappings are assumed invertible on all levels, which is checked for the grid hierarchies used in the numerical examples.

In a standard multigrid V-cycle (see, e.g., [63–65]), smoothers reduce the high-frequency part of the error associated to each level, and restriction and prolongation operators transfer residuals and corrections between the levels, respectively. The smoother chosen

is a Chebyshev-accelerated Jacobi scheme [13,14,66], which uses the inverse of the matrix diagonal on each level (precomputed before solving) and the level operator application for computing residuals in the Jacobi-type iteration. Note that also the additional fine-scale errors resulting from round-off due to the mixed-precision strategy are well-captured by the multigrid smoothers [62]. For the coarse solver, we use a conjugate gradient method (CG) preconditioned via AMG from the Trilinos ML package [67,68], which runs in double precision only. Employing a constant preconditioner would enable using CG as outer solver. Herein, however, we present a more general approach, noting that it might not be the optimal choice.

Updating the quadrature point data of the multigrid hierarchy means updating data of all operators on all levels using the multigrid transfer operator. Invertibility of the displacement map on all levels is checked, is not violated in the numerical examples discussed in the following, but is not enforced with the present approach.

7. Numerical results

We first conduct forward stability tests comparing double and single precision evaluations of the stress tensor and linearization, whereafter the single-node performance is showcased in a simplified setting. Lastly, we solve a finite-strain elasticity problem using a patient-specific geometry of an iliac bifurcation adopting the fiber model.

7.1. Forward stability test

Forward numerical stability of the first and second Piola–Kirchhoff stress tensors and the directional derivatives is investigated by evaluating them with a pseudo-random sampled second-order tensor \mathbf{G} , whose components fulfill

$$-\epsilon \leq \mathbf{G}_{ij} \leq -\epsilon/10 \vee \epsilon/10 \leq \mathbf{G}_{ij} \leq \epsilon \quad i, j = 1, \dots, d,$$

with a pseudo-random sign and an additional scale ϵ . This scale is used to emulate strain rates ranging from $\mathcal{O}(10^{-8})$ to $\mathcal{O}(10^2)$ in 200 steps, while the interval $[-\epsilon/10, \epsilon/10]$ is not considered to ensure samples of the desired order of magnitude only. To illustrate, assume $\mathbf{G}_{11} = 1.0$, while the remaining entries of \mathbf{G} are zero; this means that $\partial \mathbf{u}_1 / \partial \mathbf{x}_1 = 1.0$, that is, a stretch of 100% in \mathbf{x}_1 -direction, being already well beyond reasonable design limits in most engineering applications. However, note that the randomly generated tensor \mathbf{G} does in general not fulfill $\det(\mathbf{I} + \mathbf{G}) = 1$. For each gradient scale, we generate 10^3 independent samples, set $\text{Grad } \mathbf{u} = \mathbf{G}$ and evaluate the stress tensors and the directional derivatives using double and single precision arithmetic. The relative error ϵ_{rel} between the double precision and single precision representations is then computed as the maximum over all samples and over all corresponding tensor entries.

From this experiment we can infer forward stability (up to the observed limit) also for double precision arithmetic. It has to be mentioned, however, that this is not intended to be a rigorous analysis of numerical stability, but rather serves to showcase the improvements in the small strain limits adopting the stable formulations. Specifically, no analysis has been performed to quantify the effect of the uniform scaling of the tensor on the obtained results.

The investigations here extend the work by Shakeri et al. [37], compared to which we also present results for the directional derivatives, which enter the linear system directly. Fig. 2 illustrates the relative error in stresses (stress) and directional derivatives (D/Du stress) adopting the material or spatial integration strategies (Ω_0 or Ω_i) using stable and standard formulations. The stable formulations yield small relative errors in the small strain limit, while the standard formulations show significant numerical instability. The individual material models also show different behavior in the medium strain range of 10^{-3} to 10^0 . Interestingly, the second Piola–Kirchhoff stress tensor of the compressible neo-Hookean model in stable formulation shows numerical instability for (excessively) large strains. This might be related to the stable form in Eq. (14) using $\mathbf{C}^{-1}\mathbf{E}$ instead of \mathbf{C}^{-1} only in the standard formulation. While the standard formulation is stable for large strains, it is not numerically stable in the small strain limit, such that the stable form is considered regardless. The St.Venant–Kirchhoff model, that is, $\mathbf{S}_{\text{VK}} := \lambda \text{tr}(\mathbf{E})\mathbf{I} + 2\mu\mathbf{E}$ with \mathbf{E} according to Eq. (12), for comparison yields relative errors $\epsilon_{\text{rel}} \in [10^{-5}, 10^{-3}]$ for all strains considered here.

As a next experiment, Fig. 3 presents the accuracy of the fast evaluation of $J^{-2/3}$ and $\exp(x)$ as discussed in Section 4, which are relevant for the nearly incompressible neo-Hookean and fiber models. These results show that the effect of the fast evaluation of $\exp(\cdot)$ exploiting the floating point representation only mildly affects the numerical stability. For large strain scales, we observe large relative errors or even values out of the admissible range for the approximation of $J^{-2/3}$. This is due to the fact that the sampling strategy does not enforce $J \approx 1$, and hence the initial guess $J^{-2/3} = 1$ used in the Newton solver (see Algorithm 2) is inadequate. Note however, that this only occurs for large strain scales, and did not lead to any problems in the results presented within this work. Hence, we employ the Newton solver for $J^{-2/3}$ by default. Fast evaluation of $J^{-2/3}$ has similar effects as shown in Fig. 3(a) when using the iNH model, such that we omit these results.

Summing up, we do not observe excessive round-off errors in the relevant strain range (0 to 100%) employing the stable formulations. The fast evaluation strategies as described in Section 4 only mildly affect numerical stability in that reasonable range.

7.2. Node-level performance analysis

The single-node performance of operator evaluation is tested on an Intel Xeon Platinum 8360Y “Ice Lake” with 36 physical cores with 2.0 GHz base frequency per socket, and two sockets per node with 256 GB of DDR4-3200 memory and 100 Gbit/s full-duplex

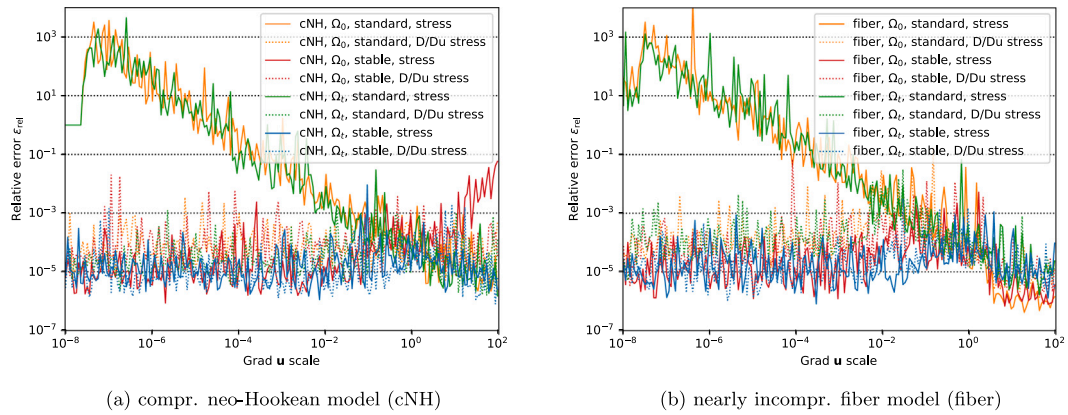


Fig. 2. Relative errors in the stress and directional derivative (D/Du stress) for sampled Grad u comparing standard and stable formulations in the material (Q_0) and spatial (Q_s) configurations. The nearly incompressible neo-Hookean and fiber models yield similar results, as the iNH model is one part of the fiber model.

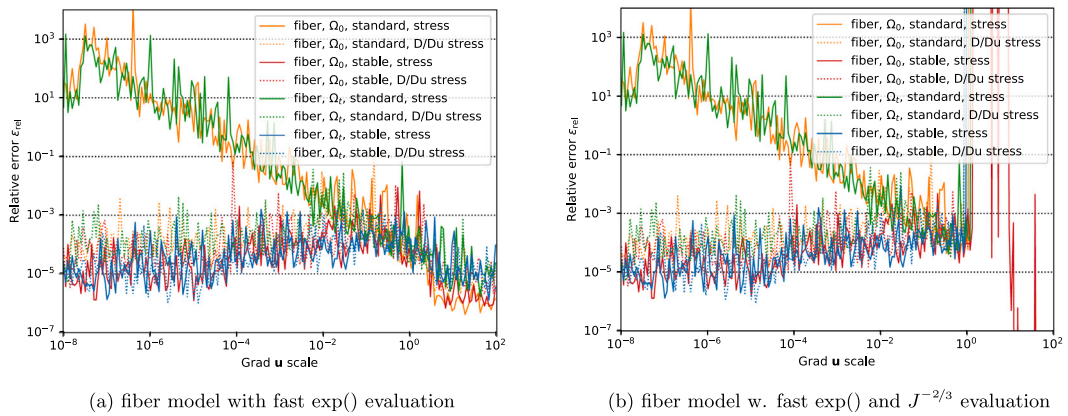


Fig. 3. Relative errors in the stress and directional derivative (D/Du stress) for sampled Grad u comparing standard and stable formulations in the material (Q_0) and spatial (Q_s) configurations with fast evaluation strategies.

Infiniband interconnect. The sum of L2 and L3 caches amounts to 200 MB, which corresponds to 25×10^6 floating point double-precision numbers. The full width of the AVX-512 instruction set (8 double-precision or 16 single-precision floating point numbers) is used in a vectorization-across-cells strategy [3]. The GNU compiler version 12.1.0 with flags “-O3 -march=icelake-server” and OpenMPI version 4.1.3 are used.

The grids are constructed by uniformly refining an initial coarse cube of $N_0 \times N_0 \times N_0$ elements of polynomial degree $p = 1, \dots, 8$ until $2.5 - 5 \times 10^6$ DoFs are reached. With this DoF count, caches are saturated according to our tests (omitted here for brevity). In the tests, we analyze the performance for different polynomial degrees p as a parameter. However, the objective of this work is to identify regimes of high throughput, not exploiting high asymptotic convergence rates of higher-order methods for sufficiently smooth solutions. We thus assume similar accuracy per unknown in all examples within this work, noting that general geometries of practical interest and resulting complexities related to mesh construction limit the practically feasible maximum polynomial degree. Table 3 lists the resulting grid parameters, showing discretizations with similar numbers of DoFs, where the discretizations with degrees $p = 5$ and $p = 7$ differ by 13% and 58%, respectively. To account for non-trivial geometries in a practical setting, the generated cube is deformed, resulting in a non-constant Jacobian of the isoparametric mapping.

The tests consist of (i) applying the linearized operator corresponding to the directional derivative, that is, the left-hand side of the Newton update step (4) on a vector as repeatedly executed within a Krylov solver, and (ii) evaluating the nonlinear residual, i.e., the right-hand side of said equations and storing it in a vector. The full node is used with 72 threads and an MPI-only parallelization, measuring the throughput in DoF/s and total memory transfer in B/DoF via LIKWID [69]. The results are derived by measuring multiples of 100 repetitions that are stopped once at least 1 s of execution time has been elapsed. The quadrature point and mapping data are updated only once in the beginning for the linearized operator application, but for every evaluation of the nonlinear residual. This is motivated by the fact that in a Newton scheme, the linearized operator is repeatedly evaluated in a single Newton iteration, but the residual needs to be evaluated with the current iterate. Tests are repeated 10 times, where the best 3 runs are averaged

Table 3

Number of continuous finite elements N_{el} and DoFs per polynomial degree p on a l times refined initial coarse grid of $N_0 \times N_0 \times N_0$ elements.

p	1	2	3	4	5	6	7	8
N_0	3	3	1	3	5	1	1	3
l	5	4	5	3	2	4	4	2
N_{el}	884 736	110 592	32 768	13 824	8000	4096	4096	1728
MDoFs	2.74	2.74	2.74	2.74	3.09	2.74	4.33	2.74

to reduce effects stemming from other jobs on the compute system. Fig. 4 depicts the obtained results for the application of the linearized operator/single evaluation of the residual for various polynomial degrees and switching between the approaches adopting integration over the spatial or material configuration and precomputing strategies. The following observations are made:

- (i) The residual evaluation involves less terms in the related integrals than the linearized operator application, but requires updating all integration point data. Therefore, integration over the reference or spatial configuration yield significantly different results, where the former competes with the fastest precomputing strategy for the linearized operator application, and the latter requires an update of all geometry-related data structures in the implementation of the deal.II library, including the detection of possible data compression and reuse of metric terms of the underlying finite element grid [3,19], and can thus only achieve a throughput of $5\text{--}8 \times 10^7$ DoF/s.
- (ii) Increasing the polynomial degree for linearized operator application reduces the memory transfer per DoF as the ratio $p^3/(p+1)^3$ between unique DoFs and quadrature points gets more favorable, such that the throughput is significantly (up to 5 times) increased up to a polynomial degree $p = 4, 5, 6$. For higher polynomial degrees, the memory transfer per DoF increases slightly due to cache effects in the element-wise integration [3].
- (iii) Regarding the precomputing and integration strategies in linearized operator application, we observe that when integrating over Ω_0 , storing scalars is faster than recomputing all data, which is faster than storing tensors. Integrating over Ω_t , this trend changes: storing tensors is the fastest option, followed by storing scalars, followed by recomputing all data. This sequence only mildly depends on the material model, but is of course highly dependent on the ratio of the number of operator evaluations to number of data updates. The fastest variant for integration over Ω_0 achieves in general better performance than the fastest variant for integration over Ω_t .
- (iv) The amount of linearization data stored per integration point does not directly translate to memory traffic as already indicated in Table 2. Storing tensors hence might circumvent loading the non-symmetric Jacobians (\mathbf{J}_0 and \mathbf{J}_t) of the finite element maps and the DoF vector of the linearization point needed otherwise — see, e.g., Eq. (15) for τ_{cNH} which is symmetric, but depends on $\text{Grad } \mathbf{u}$ (hence requiring \mathbf{J}_0 and \mathbf{u}_k , see Table 2).
- (v) Comparing the constitutive models, we see that the fiber model yields the lowest throughput and has the highest memory traffic. However, the peak throughput achieved for each of the constitutive models is $\approx 12.5 \times 10^8$ DoF/s (fiber), $\approx 17.5 \times 10^8$ DoF/s (iNH and cNH), which is due to the higher complexity of the fiber model.
- (vi) Using $p = 4$, the variant storing tensors is close to saturating the memory bandwidth of the machine (approx. 260–280 GB/s), whereas the variant storing scalars reaches a lower memory transfer (approx. 220 GB/s), but is limited by the additional computations and unstructured data access. Recomputing all linearization data yields 140–160 GB/s.
- (vii) Jacobian-free Newton–Krylov solvers might be an attractive alternative to the present approach. Based on the current results, integration has to be carried out over the reference configuration, or the mapping and/or other integration point data has to be updated selectively.

Given these observations, it is thus not clear a priori, which of the presented strategies reliably delivers the highest throughput. This decision has to be made based on measurements and hence depends on the problem at hand and the target hardware. Additionally, the fraction of updates performed per linearized operator application or residual evaluation impacts the obtained results as well. The choice of executing 100 repetitions here is therefore to be taken into account when interpreting the results. Nonetheless, the data transfer measurements shown in Fig. 4 and theoretical considerations of Section 5 provide guidelines for other hardware with possibly different arithmetic performance and memory bandwidth. For the practical influence of these low-level measurements, we apply the proposed schemes to a real-life example in the following section.

For a broader perspective, we further compare these results to linear elasticity, where $\mathbf{P} := \lambda \mathbf{I} \text{Div } \mathbf{u} + 2\mu \text{Grad}^s \mathbf{u}$, and an alternative matrix-based implementation based on identical integration routines, which assembles and stores the system matrix. Taking the fastest variant of the fiber model, i.e., integrating over the reference configuration and storing scalars, as baseline, Table 4 lists the relative throughput and memory transfer. The linear elastic model yields 2.6–3.1 times higher throughput having 0.17–0.36 times the memory transfer, while the matrix based implementation delivers merely 0.71 ($p = 1$) to 0.01 ($p = 6$) times the throughput. The memory transfer using the matrix-based variant of the fiber model is higher by a factor of 1.40 for $p = 1$, which further increases up to 112.0 for $p = 6$, taking the matrix-free implementation as baseline.

7.3. Application to biomechanics: iliac bifurcation

The numerical results in this section focus on the application to a patient-specific geometry of an iliac bifurcation, which can be considered a prototypical configuration of practical interest. The spatial approximation via a pure-hex finite element grid follows

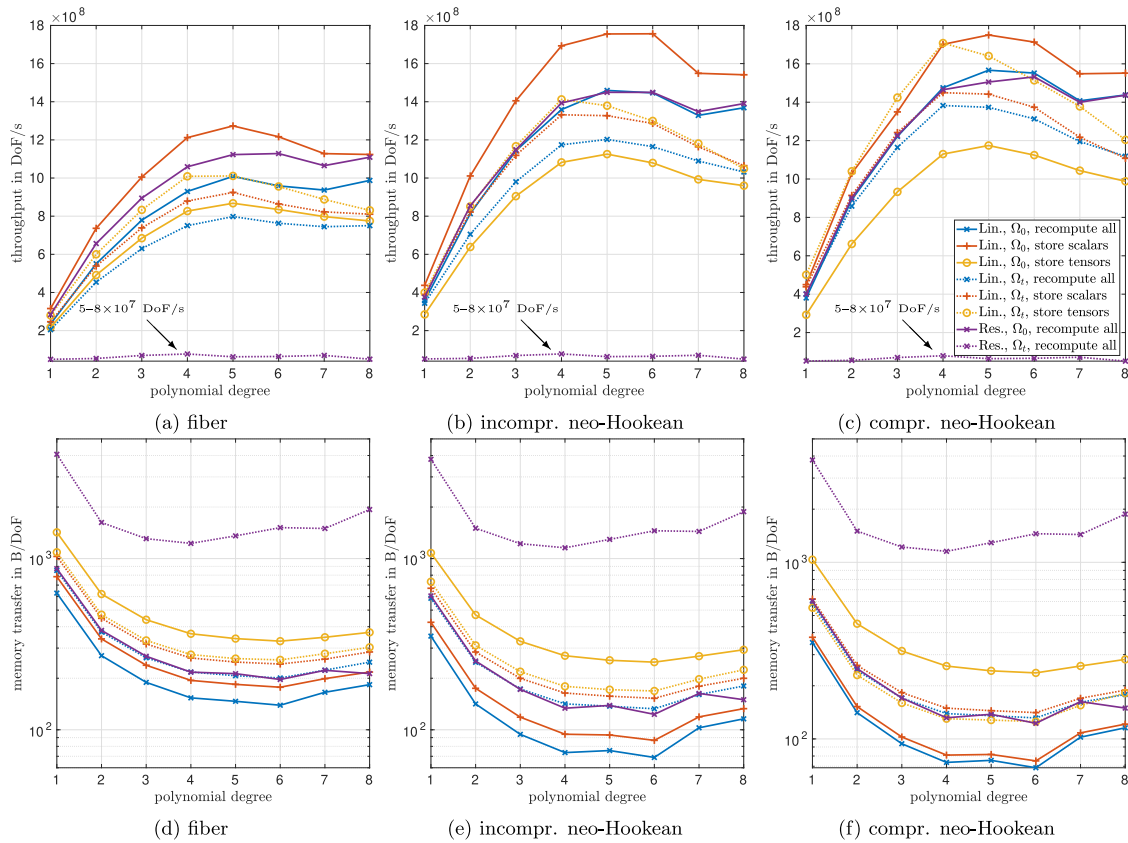


Fig. 4. Throughput (top row) and memory transfer (bottom row) using various material models considering integration over the reference configuration (Ω_0) or the spatial configuration (Ω_t) and precomputing (i) no data, (ii) scalar values, or (iii) tensorial quantities. Throughput for residual evaluation integrating over Ω_t yields values of $5 - 8 \times 10^7$ DoF/s.

Table 4

Relative throughput and memory transfer of linearized operator application compared to a baseline method (value 1.0). The baseline method refers to the fiber model, integration over Ω_0 and a matrix-free evaluation that stores scalars. This baseline method is modified in two ways: (i) by changing the material model to linear elasticity, and (ii) by choosing a matrix-based implementation for the fiber model.

Polynomial degree p		1	2	3	4	5	6
(i) linear elasticity, matrix-free	rel. throughput	2.66	2.73	2.69	2.71	2.63	2.93
	rel. memory transfer	0.36	0.23	0.25	0.22	0.24	0.17
(ii) fiber model, matrix-based	rel. throughput	0.71	0.15	0.06	0.03	0.02	0.01
	rel. memory transfer	1.39	7.32	20.71	43.47	72.84	112.29

ideas from Bošnjak et al. [70], resulting in higher-order discretizations as shown in Fig. 5 for refinement level $l = 0, 1, 2$ and mapping degree $p = 1, 2$. To demonstrate the advantages of the present approach, we aim for a problem size of 10^6 DoFs on the finest level, emulating engineering-size structural mechanics problems. The polynomial degree is varied from $p = 1$ to $p = 5$, executing uniform refinement to generate the nested multigrid hierarchy and reaching the target number of DoFs based on a coarse grid with 78 cells yielding 516 DoFs for $p = 1$. Given the fact that the coarse mesh is non-trivial and already contains a certain number of elements to resolve the topology, the resulting discretizations yield slightly differing DoF numbers, see Table 5.

Contrary to the previous example, which lies in the regime with saturated caches, this practical example of limited size showcases the effects of (partial) caching. Such a setup is relevant for practical application, when the fastest time to solution given sufficient compute resources is of interest. For the discretizations as listed in Table 5 and memory traffic related to integration point data, see Table 2, we estimate that roughly 50%–70% of the overall data can be cached in some cases. Polynomial degrees $p = 2, 3, 4$ are favorable in this regard, but this also depends on the spatial or material integration approach chosen and if linearization data is stored. However, even when precomputing scalar or tensorial quantities, a significant fraction might be cached as well.

This test involves an initial boundary value problem, adding an acceleration term, $\rho \ddot{u}$, to the linear momentum balance equation employing a standard single-step WBZ- α time integration scheme [71] with spectral radius $\rho_\infty = 0.8$ and a time step size $\Delta t = 0.1$ ms. The fiber model is employed with parameters listed in Table 1, a physiological density of $\rho = 1200$ kg/m², and the local coordinate

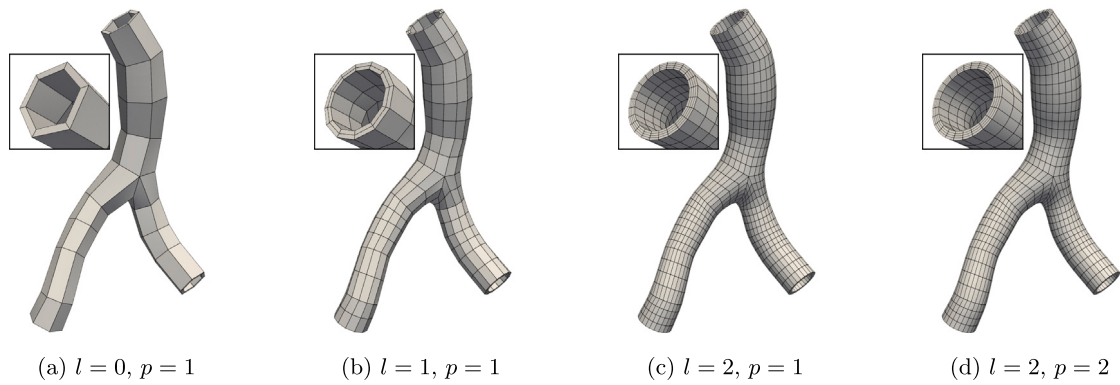


Fig. 5. Discretizations of the iliac bifurcation using $l = 0, 1, 2$ refinement levels and mapping degree $p = 1, 2$.

Table 5

Number of continuous finite elements N_{el} and DoFs per polynomial degree p on an l -times uniformly refined initial coarse grid of the iliac bifurcation with 78 elements. The discretizations left to the vertical double line are intermediate h -refined levels with $p = 1$, whereas entries on the right list fine-level discretizations.

p	1	1	1	2	1	1	2	3	4	5
l	0	1	2	2	3	4	3	2	2	2
N_{el}	78	624	4992	4992	39 936	319 488	39 936	4992	4992	4992
DoFs	516	2961	19 245	136 701	136 701	1.03×10^6	1.03×10^6	442 221	1.03×10^6	1.98×10^6

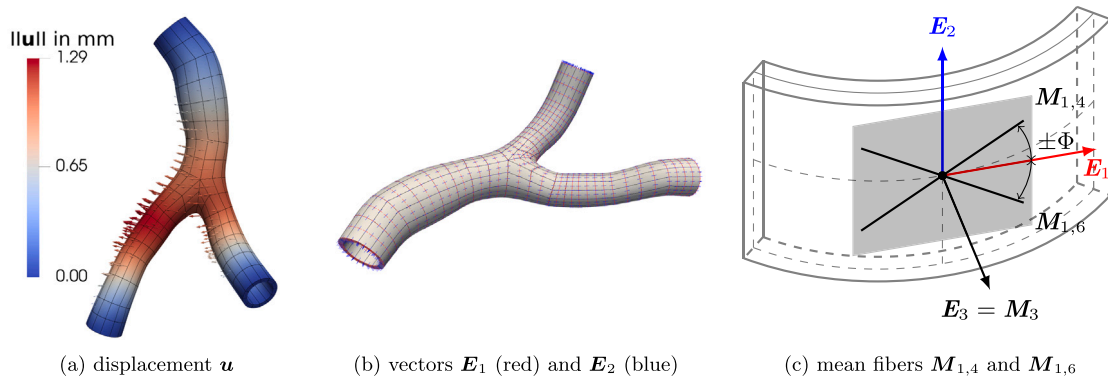


Fig. 6. Iliac bifurcation structural mechanics: tissue displacement due to pressure differential (a), material orientation (b) and sketch for derived mean fiber orientations (c). The circumferential (E_1 , red) and longitudinal direction vectors (E_2 , blue) are used to describe the helical fiber reinforcement on each h -level. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

systems as shown in Fig. 6(b) (see [48,50] for details). Focusing on a simplified structural mechanics problem, the fluid flow's effect on the tissue is roughly approximated by a uniform pressure differential acting on the vessel wall. Naturally, a fluid–structure interaction approach might in fact yield quite different results. In our simplified setup, the in- and outlets of the vessel are fixed; a rough simplification, which should be replaced for a more realistic model. Due to the lack of viscoelastic support on the exterior, the internal pressure “straightens” the curved vessel. For the present purposes, it suffices to choose the pressure differential such that a displacement of ≈ 1.3 mm is obtained as shown in Fig. 6(a), referring the interested reader to [48–50].

Regarding the preconditioner settings, we employ 6 smoothing sweeps of the Jacobi-type smoother, see Section 6. The coarse-level preconditioner uses the Amesos-KLU direct solver on the coarsest algebraic level once the total number of degrees of freedom on the level falls below 2000, which is in this present example reached immediately with 516 DoFs on level $l = 0$. However, AMG is also used for comparison with a matrix-based preconditioner on the finest level, simply referred to as AMG approach. It is important to note that the matrix–vector product of both alternatives are realized in matrix-free fashion in the Krylov solver, but the AMG preconditioner utilizes the system matrix for setup and matrix–vector products internally. If not stated otherwise, the matrix-free preconditioner operates in single precision, while the AMG preconditioner operates in double precision. When used on the fine level, the AMG preconditioner adopts a single V-cycle with Chebyshev smoother (degree 6). All these settings were tuned to yield shortest runtimes in the present scenario with a polynomial degree of $p = 5$, while increasing the number of sweeps or changing the smoother type in the AMG case did not speed-up runtimes significantly for the other polynomial degrees employed. In fact, a time-dependent

Table 6

Overview of the employed variants of the finite strain elasticity solver.

Variant	Preconditioner	Preconditioner evaluation	Operator evaluation	Floating point precision	Domain of integration	Precomputing strategy
$hpMG-MF-MP-\Omega_0$ -none	hp -multigrid	Matrix-free	Matrix-free	Mixed	Material config. Ω_0	Recompute all
$hpMG-MF-MP-\Omega_0$ -scalar	hp -multigrid	Matrix-free	Matrix-free	Mixed	Material config. Ω_0	Store scalars
$hpMG-MF-MP-\Omega_0$ -tensor	hp -multigrid	Matrix-free	Matrix-free	Mixed	Material config. Ω_0	Store tensors
$hpMG-MF-MP-\Omega_r$ -none	hp -multigrid	Matrix-free	Matrix-free	Mixed	Spatial config. Ω_r	Recompute all
$hpMG-MF-MP-\Omega_r$ -scalar	hp -multigrid	Matrix-free	Matrix-free	Mixed	Spatial config. Ω_r	Store scalars
$hpMG-MF-MP-\Omega_r$ -tensor	hp -multigrid	Matrix-free	Matrix-free	Mixed	Spatial config. Ω_r	Store tensors
AMG- Ω_0 -none	AMG	Matrix-based	Matrix-free	Double	Material config. Ω_0	Recompute all
AMG- Ω_0 -scalar	AMG	Matrix-based	Matrix-free	Double	Material config. Ω_0	Store scalars
AMG- Ω_0 -tensor	AMG	Matrix-based	Matrix-free	Double	Material config. Ω_0	Store tensors
AMG- Ω_r -none	AMG	Matrix-based	Matrix-free	Double	Spatial config. Ω_r	Recompute all
AMG- Ω_r -scalar	AMG	Matrix-based	Matrix-free	Double	Spatial config. Ω_r	Store scalars
AMG- Ω_r -tensor	AMG	Matrix-based	Matrix-free	Double	Spatial config. Ω_r	Store tensors
$hpMG-MF-DP-\Omega_0$ -tensor	hp -multigrid	Matrix-free	Matrix-free	Double	Material config. Ω_0	Store tensors
$hpMG-MB-DP-\Omega_0$ -tensor	hp -multigrid	Matrix-based	Matrix-based	Double	Material config. Ω_0	Store tensors

problem was chosen for this test due to the AMG preconditioner encountering convergence issues for higher-order finite element discretizations. With this numerical setup, we investigate the wall time per Newton solve averaged over 5 time steps and record the number of average FGMRES iterations (absolute tolerance of 10^{-12} , relative tolerance of 10^{-3} , maximal Krylov space dimension of 30 before a restart) to solve the arising Newton update steps in the nonlinear solver, which considers an absolute tolerance of 10^{-8} and relative tolerance of 10^{-3} . Due to the small time step size, convergence is reached in one to two Newton steps per time step due to the rather loose convergence criteria.

Results comparing integration over the material or spatial configuration and the different storage strategies are depicted in Fig. 7. The related speed-up of the matrix-free over the matrix-based preconditioner is reported in Tables 7–9. The considered variants summarized in Table 6 encompass matrix-free and matrix-based preconditioners, single-precision and double-precision floating point arithmetic, integration over the spatial configuration (Ω_r) or the material configuration (Ω_0), and precomputing strategies storing scalars or tensors, or recomputing all quantities.

Based on this data, we make the following observations:

- (i) The linearized operators formed by integration of the respective weak forms over the spatial and material configurations yield different approximations of the same operator. Once the nonlinear solver converges, both operators are the same up to round-off. The performance of the AMG preconditioner for higher polynomial degrees and integrating over Ω_r suffers remarkably. This is due to the AMG smoother settings aimed at fastest time to solution, which is not the most robust choice. That is, more expensive smoothers and/or multiple V-cycles in the AMG case yield nearly identical iteration counts irrespective of the domain of integration, similar to the matrix-free approach. The remaining combinations result in low and almost p -independent iteration counts. Here, the matrix-free hp -multigrid has been found to be more robust than the AMG variant used in a black-box fashion.
- (ii) The overall throughput of the proposed matrix-free approach is 1.34–52 times higher than the throughput of the AMG approach (Table 7, top). Factoring out bad preconditioner performance, we note improvements of 1.32–19 times, respectively (Table 7, bottom).
- (iii) Storing tensors is the fastest option in most cases, see Table 8. For the matrix-free approach and integrating over the spatial configuration, speed-ups of 1.18–2.09 are observed, while integrating over Ω_0 , we note a relative speed of 0.86–1.64 (storing scalars is faster for $p = 1, 2$). For the AMG approach, however, integration of the operators is no longer the dominating part of the algorithm, such that only mild improvements stemming from faster integration are observed.
- (iv) Integrating over the material configuration is in almost all cases faster than integrating over the spatial configuration with the only exception being the tensor-storing matrix-free variant, see Table 9. For the matrix-free variant, speed-up factors are 0.95–1.33, while with the AMG preconditioner leads to factors of 1.00–1.99, with higher values related to increased iteration counts.
- (v) In summary, storing tensorial quantities barely pays off compared to storing scalars for the present constitutive models, since the most complex operations involve the stored scalars and many tensorial quantities cannot be precomputed as they depend on the solution vector. Updating the mapping data for the approach integrating over Ω_r is costly and this variant is thus not preferable over the alternative integrating over Ω_0 despite the more involved integrals. This depends on the constitutive law and ratio of linear iterations required per Newton step. However, both formulations perform similar here when precomputing tensors.

Note that especially point (iii) in the above list is opposed to the throughput example from Section 7.2, where storing scalars was faster than recomputing all quantities, which was faster than storing tensors. In the present example, storing tensors is the fastest option, followed or tied with the variant storing scalars, while the variant recomputing all terms was consistently the slowest. This is due to the small problem size compared to the available 200 MB of combined L2 and L3 cache, which fits potentially large portions of the (integration point) data depending on the approach chosen. Recomputing all data is the most compute intense approach,

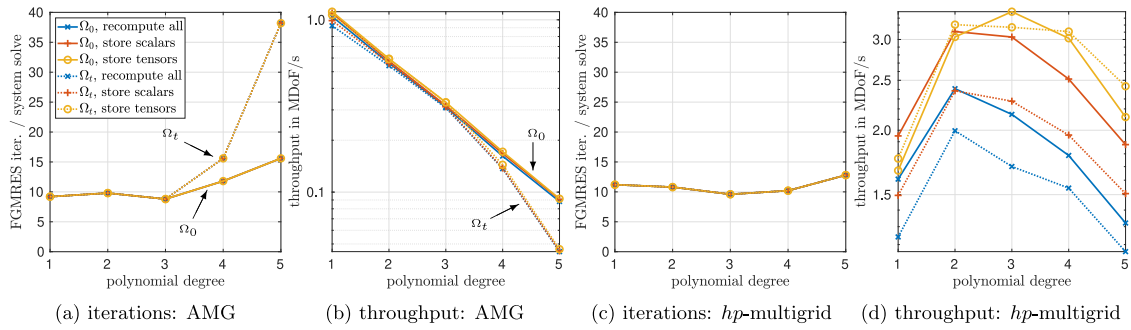


Fig. 7. FGMRES iterations until convergence adopting the AMG preconditioner on the fine grid directly in variants AMG- Ω_0/Ω_t -none/scalar/tensor (a) or matrix-free geometric multigrid in variants hp MG-MF-MP- Ω_0/Ω_t -none/scalar/tensor (c) and respective average throughput for a single system solve (b, d).

Table 7

Overall (top) and per iteration (bottom) speed-up of the matrix-free hp -multigrid (hp -MG) preconditioner (hp MG-MF-MP- Ω_0/Ω_t -none/scalar/tensor) over double-precision matrix-based AMG preconditioner (AMG- Ω_0/Ω_t -none/scalar/tensor). Both variants use matrix-free operator evaluation in the Krylov solver. Computed from data displayed in Fig. 7.

Polynomial degree p			1	2	3	4	5
Speed-up	hp -MG vs. AMG	Ω_0 , recompute all	1.54	4.29	6.88	11.05	14.97
		Ω_0 , store scalars	1.78	5.40	9.53	15.00	20.68
		Ω_0 , store tensors	1.50	5.10	10.23	17.59	23.18
		Ω_t , recompute all	1.34	3.68	5.52	11.30	25.76
		Ω_t , store scalars	1.51	4.27	7.28	14.20	33.08
		Ω_t , store tensors	1.65	5.42	9.56	21.58	52.22
Speed-up/iter.	hp -MG vs. AMG	Ω_0 , recompute all	1.87	4.73	7.51	9.55	12.29
		Ω_0 , store scalars	2.17	5.95	10.40	12.97	16.97
		Ω_0 , store tensors	1.82	5.62	11.16	15.20	19.02
		Ω_t , recompute all	1.64	4.06	6.02	7.39	8.63
		Ω_t , store scalars	1.84	4.71	7.94	9.28	11.08
		Ω_t , store tensors	2.01	5.97	10.42	14.11	17.50

Table 8

Speed-up of the tensor-storing variants (hp MG-MF-MP- Ω_0/Ω_t -tensor or AMG- Ω_0/Ω_t -tensor) over the scalar-storing and recompute all variants (hp MG-MF-MP- Ω_0/Ω_t -none/scalar or AMG- Ω_0/Ω_t -none/scalar). Computed from data displayed in Fig. 7.

Polynomial degree p			1	2	3	4	5
Speed-up	AMG	Ω_0 , recompute all	1.07	1.06	1.06	1.06	1.04
		Ω_0 , store scalars	1.02	1.03	1.04	1.02	1.01
		Ω_0 , store tensors	1.00	1.00	1.00	1.00	1.00
		Ω_t , recompute all	1.15	1.09	1.08	1.05	1.03
		Ω_t , store scalars	1.08	1.06	1.06	1.04	1.02
		Ω_t , store tensors	1.00	1.00	1.00	1.00	1.00
Speed-up	hp -MG	Ω_0 , recompute all	1.04	1.26	1.58	1.69	1.61
		Ω_0 , store scalars	0.86	0.98	1.12	1.20	1.13
		Ω_0 , store tensors	1.00	1.00	1.00	1.00	1.00
		Ω_t , recompute all	1.42	1.61	1.86	2.01	2.09
		Ω_t , store scalars	1.18	1.34	1.39	1.59	1.62
		Ω_t , store tensors	1.00	1.00	1.00	1.00	1.00

while storing scalar or tensorial data trades arithmetic operations for loading data from main memory. If a potentially large fraction of the integration point data resides in cache, as is the case for some of the variants in this example, loading precomputed data comes with a smaller penalty with regards to achievable throughput. In the present example, trends are hence partially different from the previously presented throughput results. The results in Section 7.2 show trends for saturated caches, while the results here show trends for engineering-size problems with plenty compute resources used to reduce solver turnaround times.

For the present scenario, using the standard formulations with impaired numerical stability in the small strain limit does not impact the results significantly in terms of iteration counts (results omitted for brevity), hinting at the effectiveness of the smoothers at reducing the fine-scale errors introduced by numerical round-off. This observation, however, does not extend to other scenarios straight-forwardly and cannot be generalized. Since performance is not affected tremendously, stable reformulations are thus to be preferred, while further benefits in the light of low-precision arithmetic are to be investigated in the future.

In a final comparison, we investigate the relative throughput of double-precision matrix-free and matrix-based hp -multigrid variants, the latter of which also consider matrices for the matrix–vector products on all levels for this comparison, taking the matrix-free mixed-precision hp -multigrid solver as baseline. Integration over the reference configuration and storing tensorial quantities are

Table 9

Speed-up integrating over the material configuration (variants indicated with Ω_0) rather than over the spatial configuration (variants indicated with Ω_s). Computed from data displayed in Fig. 7.

Polynomial degree p			1	2	3	4	5
Speed-up Ω_0/Ω_s	hp -MG	Recompute all	1.29	1.21	1.26	1.16	1.14
		Store scalars	1.30	1.30	1.33	1.28	1.24
		Store tensors	0.95	0.95	1.07	0.97	0.87
	AMG	Recompute all	1.13	1.04	1.01	1.18	1.95
		Store scalars	1.10	1.03	1.02	1.22	1.99
		Store tensors	1.05	1.01	1.00	1.19	1.96

Table 10

Relative throughput of hp -multigrid solvers using double precision in matrix-free and matrix-based settings, taking the matrix-free mixed-precision version as baseline (see Table 6). Integration over the reference configuration and storing tensorial quantities are considered. Iteration counts are identical for all variants.

Polynomial degree p	1	2	3	4	5
rel. throughput, mixed precision, matrix-free (hp MG-MF-MP- Ω_0 -tensor)	1.00	1.00	1.00	1.00	1.00
rel. throughput, double precision, matrix-free (hp MG-MF-DP- Ω_0 -tensor)	0.57	0.57	0.62	0.58	0.58
rel. throughput, double precision, matrix-based (hp MG-MB-DP- Ω_0 -tensor)	0.40	0.14	0.07	0.05	0.04
avg. FGMRES iterations per system solve	11.2	10.8	9.6	10.2	12.8

considered. As indicated in Table 10, iterations are independent of the choice of mixed/double precision (and of course for matrix-free/matrix-based variants). The throughput using a matrix-free double precision variant is roughly 0.60 of the mixed-precision equivalent, while a matrix-based implementation achieves 0.40 of the throughput for $p = 1$, which drastically decreases for higher polynomial degrees in the present example. These factors match the ones presented in Section 7.2. For a mixed-precision matrix-based variant, similar trends as shown here for double precision are to be expected. Again, this comparison underlines that for lower order finite elements, matrix-free and matrix-based solution strategies are competitive. These results indicate that for higher polynomial degrees, assembled sparse matrices are a poor format for achieving high performance in this application. Combining Table 10 and Fig. 7, we deduce that for a problem with the same number of DoFs using $p = 1$ and a matrix-based approach (1.1 MDoF/s \times 0.4) vs. $p = 2, 3, 4$ and a matrix-free approach (3.3 MDoF/s \times 0.6), the latter is roughly 4.5 times faster, and using a single-precision preconditioner for the latter (3.3 MDoF/s \times 1.0), this factor increases to 7.5.

Closing the discussion of this practical application, it should be noted that the numerical results presented here strongly depend on the tuning of the preconditioner and problem setup, where difficulties were encountered reaching convergence for the AMG-preconditioned solver and higher polynomial degrees, as this preconditioner is not particularly well suited for this scenario. The AMG-preconditioned variant also adopts the matrix-free operator evaluation in the FGMRES solver, such that only the impact of the matrix-free vs. matrix-based preconditioners can be evaluated. Results for matrix-free vs. matrix-based approaches are given in Table 10. Hence, all other results do not show the full speed-ups achievable via matrix-free solution strategies. Furthermore, the AMG preconditioner operates in double precision, while the matrix-free hp -multigrid alternative operates in single precision. Assuming that an AMG preconditioner delivering constant iteration counts exists, factoring out the increasing iteration counts reveals that especially for higher polynomial degrees, the matrix-free variants are significantly faster. Due to the limitations of this comparison, we refrain from extending these observations towards $p = 1$, as more involved or better tuned AMG preconditioners might outperform the matrix-free variant for this case, while factors of larger than 4 and up to 19 (for $p = 2, 3, 4, 5$) might be much harder to compensate. Overall, this example represents practically relevant scenarios, and the results shown are hence very promising beyond academic setups. Furthermore, we want to emphasize at this point that choosing a Poisson's ratio $\nu \rightarrow 0.5$ (fully incompressible case) and hence $\kappa_b \rightarrow \infty$ impacts performance significantly through ill-conditioning of the linear system, which is a well-studied limitation of the purely displacement-based formulations in general. For such scenarios, the problem formulation itself, the restriction and prolongation operators, and the smoothers might need to be adapted, while the presented combination of methods is an excellent starting point as demonstrated.

8. Summary and concluding remarks

This work presents a matrix-free finite element solver for finite-strain elasticity, where several variants for fast and numerically stable integration are discussed. We devise stable reformulations of the classical continuum mechanical relations for the anisotropic hyperelastic model for vascular tissue by Holzapfel et al. [20], which are demonstrated in a simplified forward stability test to not suffer from excessive round-off errors in the small strain limit. The results also encompass stable reformulations of the directional derivatives for this fiber model as well as for compressible and nearly incompressible neo-Hookean models extending the work by [37] in this regard.

We further discuss variants integrating over the material configuration or, alternatively, over the spatial configuration and precompute and store data in integration points extending work by [5] towards more complex constitutive models. Changing the domain of integration alters the integrals to be evaluated, where, depending on the material model and precomputing strategy considered, any of the two integration approaches might turn out to be favorable.

In the presented tests, integrating over the reference configuration and storing scalars turned out to be faster in most cases, as the constitutive models used here feature complex terms for scalar quantities and many tensor-valued operations involve the current iterate and thus cannot be precomputed. Comparing matrix-free and matrix-based preconditioners, we observe increased robustness with respect to the polynomial degree p and significantly reduced time to solution for the matrix-free approach using higher polynomial degrees $p > 1$ due to increased iteration counts using an AMG preconditioner in a black-box fashion.

Even for linear elements, where the matrix-based AMG preconditioner is found to perform well, the matrix-free approach was found to yield speed-ups of 1.34–1.78 in the present setup. For higher polynomial degrees, speed-up factors of 3.68–19.02 have been recorded. These results are primarily due to faster operator evaluation, i.e., matrix–vector products, in the matrix-free case. Here, cache effects might play a central role in the comparison as well, as the problem is purposely chosen to be of engineering size. However, the presented results showcase one particular problem size, while optimized caching strategies go beyond the scope of this contribution. The presented results have not exhaustively analyzed AMG settings tailored to finite-strain hyperelasticity. For example, more sophisticated smoothers or coarsening strategies, not available in the matrix-free hp -multigrid solver, have not been considered. In consequence, the present results have to be interpreted with caution, and further research would be necessary to quantify these options. Nonetheless, factors of 3.68–19.02 for $p = 2, 3, 4, 5$ might be difficult to overcome.

Altogether, the matrix-free finite element solvers for finite-strain problems presented in this work show excellent properties also when tackling challenging real-world applications. The additional implementation effort of matrix-free methods and the different variants discussed in the present work appears justified in light of the significant speed-up that can be achieved over matrix-based methods, in particular when considering higher-order finite element methods, which additionally counteracts locking effects [21–24].

While the present results were recorded on a particular hardware, the computational models quantifying the memory transfer and arithmetic work allow for predictions also on evolving hardware: In general, moderate polynomial degrees $p = 2, \dots, 6$ give the most favorable memory access per DoF. In consequence, the best performance can also be expected in this regime for memory-limited cases. As the proposed models are relatively heavy on operations at quadrature points with many register spills to local memory, architectures with limited cache sizes such as GPUs can be expected to benefit even more from the storage option of tensors than the results presented here.

Ongoing developments are centered around further improving the routines to update quadrature point data to experiment with Jacobian-free Newton–Krylov methods, which avoid the formulation and optimization of the linearized operator, but require fast residual evaluation (see, e.g. [72]). The multigrid solver can be improved following many ideas: First, resolving the incompressibility constraint via mixed formulations instead of the penalty-based approach prevents locking for coarse discretizations and may hence better approximate the solution on coarser multigrid levels. Second, structure-preserving strategies regarding the finite element mapping and restriction/prolongation operators might significantly improve robustness of the solver when facing large deformations. Similar concepts might also be used to derive tailored smoothers for the fully incompressible case. Third, the outer FGMRES solver might be replaced by a CG method once the preconditioner is constant, which is the case for the specific iliac bifurcation example shown here, but the principal solver design allows for rather large, general coarse-level problems, without relying too much on problem-dependent tuning or a fixed number of AMG V-cycles, a fixed number of AMG-preconditioned Krylov solver iterations or low-tolerance preconditioned iterative solvers.

CRedit authorship contribution statement

R. Schussnig: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Conceptualization. **N. Fehn:** Writing – review & editing, Software, Methodology, Investigation. **P. Munch:** Writing – review & editing, Software, Methodology, Investigation. **M. Kronbichler:** Writing – review & editing, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition.

Declaration of competing interest

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

Acknowledgments

The authors gratefully acknowledge the scientific support and HPC resources provided by the Erlangen National High Performance Computing Center (NHR@FAU) of the Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU). NHR funding is provided by federal and Bavarian state authorities. NHR@FAU hardware is partially funded by the German Research Foundation (DFG) – 440719683. Moreover, this work was partially supported by the German Ministry of Education and Research through project “PDExa: Optimized software methods for solving partial differential equations on exascale supercomputers”, grant agreement no. 16ME0637K.

Appendix. Material models: Directional derivatives

Providing the strain energy density suffices to define a hyperelastic material model. Therefore, we present further derivations here to not clutter the main document. For the compressible neo-Hookean model (8), the directional derivative of the second Piola–Kirchhoff stress reads

$$\mathbf{D}_u \mathbf{S}_{\text{cNH}} = -(\mu - 2\lambda \ln J) \mathbf{D}_u \mathbf{C}^{-1} + 2\lambda (1/J \mathbf{D}_u J) \mathbf{C}^{-1},$$

with the following terms not being specific to the material model:

$$\begin{aligned} \mathbf{D}_u \mathbf{C}^{-1} &= (\mathbf{D}_u \mathbf{F}^{-1}) \mathbf{F}^{-\text{T}} + \mathbf{F}^{-1} \mathbf{D}_u \mathbf{F}^{-\text{T}} = 2 [(\mathbf{D}_u \mathbf{F}^{-1}) \mathbf{F}^{-\text{T}}]^{\text{S}}, \\ \mathbf{D}_u \mathbf{F} &= \text{Grad } \Delta \mathbf{u}, \quad \mathbf{D}_u \mathbf{F}^{-1} = -\mathbf{F}^{-1} (\text{Grad } \Delta \mathbf{u}) \mathbf{F}^{-1}, \quad 1/J \mathbf{D}_u J = \text{tr} (\mathbf{F}^{-1} \text{Grad } \Delta \mathbf{u}). \end{aligned}$$

Likewise, for nearly incompressible neo-Hookean materials (9), we obtain

$$\begin{aligned} \mathbf{D}_u \mathbf{S}_{\text{iNH}} &= -2\mu/3 J^{-2/3} (1/J \mathbf{D}_u J) \mathbf{I} + [\kappa_b/2 (J^2 - 1) - \mu/3 J^{-2/3} I_1] \mathbf{D}_u \mathbf{C}^{-1} \\ &\quad + [(1/J \mathbf{D}_u J) (2\mu/9 J^{-2/3} I_1 + \kappa_b J^2) - \mu/3 J^{-2/3} \mathbf{D}_u I_1] \mathbf{C}^{-1}, \end{aligned} \quad (19)$$

with $\mathbf{D}_u I_1 = \text{tr} (\mathbf{D}_u \mathbf{F}^{\text{T}} \mathbf{F} + \mathbf{F}^{\text{T}} \mathbf{D}_u \mathbf{F}) = 2 \text{tr} (\mathbf{F}^{\text{T}} \text{Grad } \Delta \mathbf{u})$. For the fiber model [20], we sum contributions from the nearly incompressible neo-Hookean ground material (19) and collagen fibers to obtain

$$\mathbf{D}_u \mathbf{S}_{\text{fiber}} = \mathbf{D}_u \mathbf{S}_{\text{iNH}} + \mathbf{D}_u \mathbf{S}_{\text{c}} = \mathbf{D}_u \mathbf{S}_{\text{iNH}} + \sum_{i=4,6} 2k_1 \exp(k_2 E_i^2) (2k_2 E_i^2 + 1) (\mathbf{H}_i : \mathbf{D}_u \mathbf{C}) \mathbf{H}_i, \quad (20)$$

where we use

$$\mathbf{D}_u E_i = \mathbf{D}_u (\mathbf{H}_i : \mathbf{C}) = \mathbf{H}_i : \mathbf{D}_u \mathbf{C}, \quad \text{and} \quad \mathbf{D}_u \mathbf{C} = 2 (\mathbf{F}^{\text{T}} \text{Grad } \Delta \mathbf{u})^{\text{S}}.$$

Turning our attention to the approach integrating over the spatial configuration, we have

$$J_{\mathbb{C}} := \chi \left(4 \frac{\partial^2 \Psi(\mathbf{C})}{\partial \mathbf{C} \otimes \partial \mathbf{C}} \right),$$

where the contravariant push-forward of a fourth-order tensor is defined as $\chi(\cdot)_{ijkl} := \mathbf{F}_{iA} \mathbf{F}_{jB} (\cdot)_{ABCD} \mathbf{F}_{kC} \mathbf{F}_{lD}$ using Einstein's summation convention. For the compressible neo-Hookean model (8), we have

$$\boldsymbol{\tau}_{\text{cNH}} = \mu \mathbf{b} - (\mu - 2\lambda \ln J) \mathbf{I}, \quad \text{and} \quad \frac{\partial^2 \Psi_{\text{cNH}}(\mathbf{C})}{\partial \mathbf{C} \otimes \partial \mathbf{C}} = 1/2 (\mu - 2\lambda \ln J) \mathbf{C}^{-1} \odot \mathbf{C}^{-1} + 1/2 \mathbf{C}^{-1} \otimes \mathbf{C}^{-1},$$

with the left Cauchy–Green tensor $\mathbf{b} := \mathbf{F} \mathbf{F}^{\text{T}}$ and using $\frac{\partial \mathbf{C}^{-1}}{\partial \mathbf{C}} = -\mathbf{C}^{-1} \odot \mathbf{C}^{-1}$ (see [38]),

$$(\mathbf{C}^{-1} \odot \mathbf{C}^{-1})_{ABCD} := -\frac{\partial \mathbf{C}_{AB}^{-1}}{\partial \mathbf{C}_{CD}} = 1/2 (\mathbf{C}_{AC}^{-1} \mathbf{C}_{BD}^{-1} + \mathbf{C}_{AD}^{-1} \mathbf{C}_{BC}^{-1}),$$

such that we further have

$$\begin{aligned} J_{\mathbb{C}\text{NH}} : (\cdot)^{\text{S}} &= 2(\mu - 2\lambda \ln J) \chi (\mathbf{C}^{-1} \odot \mathbf{C}^{-1}) : (\cdot)^{\text{S}} + 2\lambda \chi (\mathbf{C}^{-1} \otimes \mathbf{C}^{-1}) : (\cdot)^{\text{S}} \\ &= 2(\mu - 2\lambda \ln J) \mathbb{S} : (\cdot)^{\text{S}} + 2\lambda (\mathbf{I} \otimes \mathbf{I}) : (\cdot)^{\text{S}} \\ &= 2(\mu - 2\lambda \ln J) (\cdot)^{\text{S}} + 2\lambda \text{tr} (\cdot) \mathbf{I}, \end{aligned} \quad (21)$$

using $\chi (\mathbf{C}^{-1} \odot \mathbf{C}^{-1})_{ijkl} = \mathbb{S}_{ijkl} := 1/2 (\delta_{ik} \delta_{jl} + \delta_{il} \delta_{jk})$, and $\chi (\mathbf{C}^{-1} \otimes \mathbf{C}^{-1}) = \mathbf{I} \otimes \mathbf{I}$, and where the symmetry of the argument is exploited in the inner product with the fourth-order tensor \mathbb{S} . For the nearly incompressible neo-Hookean model (9), similar steps lead to

$$\begin{aligned} \boldsymbol{\tau}_{\text{iNH}} &= \mu J^{-2/3} \mathbf{b} + [\kappa_b/2 (J^2 - 1) - \mu/3 J^{-2/3} I_1] \mathbf{I}, \\ \text{and} \quad \frac{\partial^2 \Psi_{\text{iNH}}(\mathbf{C})}{\partial \mathbf{C} \otimes \partial \mathbf{C}} &= -\mu/3 J^{-2/3} \mathbf{I} \otimes \mathbf{C}^{-1} + [\mu/6 J^{-2/3} I_1 - \kappa_b/4 (J^2 - 1)] \mathbf{C}^{-1} \odot \mathbf{C}^{-1} + (\mu/18 J^{-2/3} I_1 + \kappa_b/4 J^2) \mathbf{C}^{-1} \otimes \mathbf{C}^{-1}, \end{aligned}$$

which together with $\chi (\mathbf{I} \otimes \mathbf{C}^{-1}) : (\cdot) = (\mathbf{C} \otimes \mathbf{I}) : (\cdot) = \mathbf{C} [\mathbf{I} : (\cdot)] = \text{tr} (\cdot) \mathbf{C}$ then finally gives a compact expression similar to Eq. (21) only involving second-order tensors,

$$J_{\mathbb{C}\text{iNH}} : (\cdot)^{\text{S}} = -4\mu/3 J^{-2/3} \mathbf{C} \text{tr} (\cdot) + [2\mu/3 J^{-2/3} I_1 - \kappa_b (J^2 - 1)] (\cdot)^{\text{S}} + (2\mu/9 J^{-2/3} I_1 + \kappa_b J^2) \text{tr} (\cdot) \mathbf{I},$$

again exploiting symmetry of the argument in the inner product with the fourth-order symmetric identity tensor \mathbb{S} . For the model including non-symmetrically dispersed fibers [20], we have

$$\begin{aligned} \boldsymbol{\tau}_{\text{fiber}} &= \boldsymbol{\tau}_{\text{iNH}} + \boldsymbol{\tau}_{\text{c}} = \boldsymbol{\tau}_{\text{iNH}} + \sum_{i=4,6} 2k_1 \exp(k_2 E_i^2) E_i \mathbf{F} \mathbf{H}_i \mathbf{F}^{\text{T}}, \\ \text{and} \quad \frac{\partial^2 \Psi_{\text{fiber}}(\mathbf{C})}{\partial \mathbf{C} \otimes \partial \mathbf{C}} &= \frac{\partial^2 \Psi_{\text{iNH}}(\mathbf{C})}{\partial \mathbf{C} \otimes \partial \mathbf{C}} + \sum_{i=4,6} k_1 \exp(k_2 E_i^2) (2k_2 E_i^2 + 1) \mathbf{H}_i \otimes \mathbf{H}_i. \end{aligned}$$

The directional derivative of the fiber contribution is easiest found by pushing forward the fiber contribution in (20) and using $\text{Grad } \Delta \mathbf{u} = (\text{grad } \Delta \mathbf{u}) \mathbf{F}$, yielding

$$\begin{aligned} & \sum_{i=4,6} 2k_1 \exp(k_2 E_i^2) (2k_2 E_i^2 + 1) [\mathbf{H}_i : (\mathbf{F}^T (\text{grad } \Delta \mathbf{u}) \mathbf{F} + \mathbf{F}^T (\text{grad }^T \Delta \mathbf{u}) \mathbf{F})] \mathbf{F} \mathbf{H}_i \mathbf{F}^T \\ &= \sum_{i=4,6} 2k_1 \exp(k_2 E_i^2) (2k_2 E_i^2 + 1) 2 [(\mathbf{F} \mathbf{H}_i \mathbf{F}^T) : \text{grad } \Delta \mathbf{u}] \mathbf{F} \mathbf{H}_i \mathbf{F}^T \end{aligned}$$

For symmetric arguments, this then finally leads to

$$J_{\mathbb{C}_{\text{fiber}}} : (\cdot)^S = J_{\mathbb{C}_{\text{INH}}} : (\cdot)^S + \sum_{i=4,6} 4k_1 \exp(k_2 E_i^2) (2k_2 E_i^2 + 1) [(\mathbf{F} \mathbf{H}_i \mathbf{F}^T) : (\cdot)] \mathbf{F} \mathbf{H}_i \mathbf{F}^T.$$

Data availability

Data will be made available on request.

References

- [1] S.A. Orszag, Spectral methods for problems in complex geometries, *J. Comput. Phys.* 37 (1) (1980) 70–92.
- [2] M. Kronbichler, K. Kormann, A generic interface for parallel cell-based finite element operator application, *Comput. & Fluids* 63 (2012) 135–147.
- [3] M. Kronbichler, K. Kormann, Fast matrix-free evaluation of discontinuous Galerkin finite element operators, *ACM Trans. Math. Software* 45 (3) (2019) 29/1–40.
- [4] C. Burstedde, L.C. Wilcox, O. Ghattas, P4est: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees, *SIAM J. Sci. Comput.* 33 (3) (2011) 1103–1133.
- [5] D. Davydov, J.-P. Pelteret, D. Arndt, M. Kronbichler, P. Steinmann, A matrix-free approach for finite-strain hyperelastic problems using geometric multigrid, *Internat. J. Numer. Methods Engrg.* 121 (2020) 2874–2895.
- [6] J. Brown, V. Barra, N. Beams, L. Ghaffari, M. Knepley, W. Moses, R. Shakeri, K. Stengel, J.L. Thompson, J. Zhang, Performance portable solid mechanics via matrix-free p -multigrid, 2022, <http://dx.doi.org/10.48550/arXiv.2204.01722>, arXiv.
- [7] A. Mehraban, H. Tufo, S. Sture, R. Regueiro, Matrix-free higher-order finite element method for parallel simulation of compressible and nearly-incompressible linear elasticity on unstructured meshes, *Comput. Model. Eng. Sci.* 129 (3) (2021) 1283–1303.
- [8] M.S. Fabien, A GPU-accelerated hybridizable discontinuous Galerkin method for linear elasticity, *Commun. Comput. Phys.* 27 (2) (2020) 513–545.
- [9] U. Kiran, D. Sharma, S.S. Gautam, A GPU-based framework for finite element analysis of elastoplastic problems, *Computing* 105 (8) (2023) 1673–1696.
- [10] U. Kiran, D. Sharma, S.S. Gautam, Development of GPU-based matrix-free strategies for large-scale elastoplasticity analysis using conjugate gradient solver, *Internat. J. Numer. Methods Engrg.* 125 (7) (2024) e7421.
- [11] U. Kiran, D. Sharma, S.S. Gautam, An efficient framework for matrix-free SpMV computation on GPU for elastoplastic problems, *Math. Comput. Simulation* 216 (2024) 318–346.
- [12] H. Anzt, T. Cojean, G. Flegar, F. Göbel, T. Grützmacher, P. Nayak, T. Ribizel, Y.M. Tsai, E.S. Quintana-Ortí, Ginkgo: A modern linear operator algebra framework for high performance computing, *ACM Trans. Math. Software* 48 (1) (2022) 1–33.
- [13] M. Kronbichler, W.A. Wall, A performance comparison of continuous and discontinuous Galerkin methods with fast multigrid solvers, *SIAM J. Sci. Comput.* 40 (5) (2018) A3423–A3448.
- [14] N. Fehn, P. Munch, W.A. Wall, M. Kronbichler, Hybrid multigrid methods for high-order discontinuous Galerkin discretizations, *J. Comput. Phys.* 415 (2020) 109538.
- [15] N. Fehn, Robust and efficient discontinuous Galerkin methods for incompressible flows, (Ph.D. thesis), Technical University Munich, 2021.
- [16] P. Munch, T. Heister, L. Prieto Saavedra, M. Kronbichler, Efficient distributed matrix-free multigrid methods on locally refined meshes for FEM computations, *ACM Trans. Parallel Comput.* 10 (1) (2023) 3/1–38.
- [17] N. Fehn, M. Kronbichler, P. Munch, ExaDG, 2024, URL <https://github.com/exadg/exadg>.
- [18] D. Arndt, N. Fehn, G. Kanschat, K. Kormann, M. Kronbichler, P. Munch, W.A. Wall, J. Witte, ExaDG: High-order discontinuous Galerkin for the exa-scale, in: H.-J. Bungartz, S. Reiz, B. Uekermann, P. Neumann, W.E. Nagel (Eds.), *Software for Exascale Computing-SPPEXA 2016-2019*, Springer International Publishing, Cham, 2020, pp. 189–224.
- [19] D. Arndt, W. Bangerth, M. Bergbauer, M. Feder, M. Fehling, J. Heinz, T. Heister, L. Heltai, M. Kronbichler, M. Maier, P. Munch, J.-P. Pelteret, B. Turcksin, D. Wells, S. Zampini, The deal. II library, version 9.5, *J. Numer. Math.* 31 (3) (2023) 231–246.
- [20] G.A. Holzapfel, J.A. Niestrawska, R.W. Ogden, A.J. Reinisch, A.J. Schriefl, Modelling non-symmetric collagen fibre dispersion in arterial walls, *J. R. Soc. Interface* 12 (106) (2015) 20150188.
- [21] U. Heisserer, S. Hartmann, A. Düster, Z. Yosibash, On volumetric locking-free behaviour of p -version finite elements under finite deformations, *Commun. Numer. Methods Eng.* 24 (11) (2008) 1019–1032.
- [22] M. Suri, I. Babuška, C. Schwab, Locking effects in the finite element approximation of plate models, *Math. Comp.* 64 (210) (1995) 461–482.
- [23] M. Suri, Analytical and computational assessment of locking in the hp finite element method, *Comput. Methods Appl. Mech. Engrg.* 133 (3) (1996) 347–371.
- [24] L. Radtke, M. König, A. Düster, The influence of geometric imperfections in cardiovascular FSI simulations, *Comput. Math. Appl.* 74 (7) (2017) 1675–1689.
- [25] R. Stenberg, M. Suri, Mixed hp finite element methods for problems in elasticity and Stokes flow, *Numer. Math.* 72 (1996) 367–389.
- [26] M.E. Canga, E.B. Becker, An iterative technique for the finite element analysis of near-incompressible materials, *Comput. Methods Appl. Mech. Engrg.* 170 (1999) 79–101.
- [27] A.M. Maniatty, Y. Liu, O. Klaas, M.S. Shephard, Higher order stabilized finite element method for hyperelastic finite deformation, *Comput. Methods Appl. Mech. Engrg.* 191 (2002) 1491–1503.
- [28] P.E. Farrell, L.F. Gatica, B.P. Lamichhane, R. Oyarzúa, R. Ruiz-Baier, Mixed Kirchhoff stress–displacement–pressure formulations for incompressible hyperelasticity, *Comput. Methods Appl. Mech. Engrg.* 374 (2021) 113562.
- [29] J.C. Simo, M.S. Rifai, A class of mixed assumed strain methods and the method of incompatible modes, *Internat. J. Numer. Methods Engrg.* 29 (1990) 1595–1638.
- [30] J.C. Simo, F. Armero, Geometrically non-linear enhanced strain mixed methods and the method of incompatible modes, *Internat. J. Numer. Methods Engrg.* 33 (1992) 1413–1449.
- [31] J.-S. Chen, C. Pan, A pressure projection method for nearly incompressible rubber hyperelasticity, part I: Theory, *J. Appl. Mech.* 63 (1996) 862–868.
- [32] Y. Yu, H. Baek, M.L. Bittencourt, G.E. Karniadakis, Mixed spectral/ hp element formulation for nonlinear elasticity, *Comput. Methods Appl. Mech. Engrg.* 213–216 (2012) 42–57.

- [33] R.S. Falk, Nonconforming finite element methods for the equations of linear elasticity, *Math. Comp.* 57 (1991) 529–550.
- [34] A. Ten Eyck, A. Lew, Discontinuous Galerkin methods for non-linear elasticity, *Internat. J. Numer. Methods Engrg.* 67 (2006) 1204–1243.
- [35] H. Kabaria, A.J. Lew, B. Cockburn, A hybridizable discontinuous Galerkin formulation for non-linear elasticity, *Comput. Methods Appl. Mech. Engrg.* 283 (2015) 303–329.
- [36] D. Di Pietro, S. Lemaire, An extension of the Crouzeix–Raviart space to general meshes with application to quasi-incompressible linear elasticity and Stokes flow, *Math. Comp.* 84 (2014) 1–31.
- [37] R. Shakeri, L. Ghaffari, J.L. Thompson, J. Brown, Stable numerics for finite-strain elasticity, *Internat. J. Numer. Methods Engrg.* 125 (21) (2024) e7563.
- [38] G.A. Holzapfel, *Nonlinear solid mechanics - A continuum approach for engineering*, John Wiley & Sons, Chichester, 2000.
- [39] J. Bonet, R.D. Wood, *Nonlinear continuum mechanics for finite element analysis*, second ed., Cambridge University Press, Cambridge, 2008.
- [40] J.M. Ball, Convexity conditions and existence theorems in nonlinear elasticity, *Arch. Ration. Mech. Anal.* 63 (4) (1976) 337–403.
- [41] P.G. Ciarlet, Finite element method for elliptic problems, *Society for Industrial and Applied Mathematics*, USA, 2002.
- [42] V. Ruas, Existence and stability of asymmetric finite-element approximations in nonlinear incompressible analysis, *SIAM J. Numer. Anal.* 26 (5) (1989) 1031–1059.
- [43] P. Wriggers, *Nonlinear finite element methods*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [44] L.R.G. Treloar, The mechanics of rubber elasticity, *Proc. R. Soc. A* 351 (1666) (1976) 301–330.
- [45] P.J. Flory, Thermodynamic relations for high elastic materials, *Trans. Faraday Soc.* 57 (1961) 829–838.
- [46] A.J. Schriefl, G. Zeindlinger, D. Pierce, P. Regitnig, G.A. Holzapfel, Determination of the layer-specific distributed collagen fibre orientations in human thoracic and abdominal aortas and common iliac arteries, *J. R. Soc. Interface* 9 (71) (2012) 1275–1286.
- [47] G.A. Holzapfel, R.W. Ogden, On the tension–compression switch in soft fibrous solids, *Eur. J. Mech. A Solids* 49 (2015) 561–569.
- [48] R. Schussnig, D.R.Q. Pacheco, M. Kaltenbacher, T.-P. Fries, Semi-implicit fluid–structure interaction in biomedical applications, *Comput. Methods Appl. Mech. Engrg.* 400 (2022) 115489.
- [49] R. Schussnig, M. Rolf-Pissarczyk, G. Holzapfel, T.-P. Fries, Fluid–structure interaction simulations of aortic dissection, *Proc. Appl. Math. Mech.* 20 (1) (2021) e202000125.
- [50] R. Schussnig, K. Bäuml, T.-P. Fries, Multi-layered tissue models in patient-specific simulations of aortic dissection, *Proc. Appl. Math. Mech.* 21 (1) (2021) e202100090.
- [51] M.P. Wollner, Semi-analytical investigation of the transmural alignment of vascular smooth muscle cells, (Master's Thesis), Technical University Dresden, 2020.
- [52] H. Weisbecker, D.M. Pierce, P. Regitnig, G.A. Holzapfel, Layer-specific damage experiments and modeling of human thoracic and abdominal aortas with non-atherosclerotic intimal thickening, *J. Mech. Behav. Biomed. Mater.* 12 (2012) 93–106.
- [53] M. Rolf-Pissarczyk, M.P. Wollner, D.R.Q. Pacheco, G.A. Holzapfel, Efficient computational modelling of smooth muscle orientation and function in the aorta, *Proc. R. Soc. A* 477 (2021) 20210592.
- [54] A. Abdelfattah, H. Anzt, E. Boman, E. Carson, T. Cojean, J. Dongarra, A. Fox, M. Gates, N.J. Higham, X.S. Li, J. Loe, P. Luszczyk, S. Pranesh, S. Rajamanickam, T. Ribizel, B.F. Smith, K. Świrydowicz, S. Thomas, S. Tomov, Y.M. Tsai, U.M. Yang, A survey of numerical linear algebra methods utilizing mixed-precision arithmetic, *Int. J. Supercomput. Appl. High Perform. Comput.* 35 (4) (2021) 344–369.
- [55] N.H.F. Beebe, *The mathematical-function computation handbook – programming using the MathCW portable software library*, Springer International Publishing, Cham, 2017.
- [56] N.N. Schraudolph, A fast, compact approximation of the exponential function, *Neural Comput.* 11 (4) (1999) 853–862.
- [57] S.D. Proell, J. Brotz, M. Kronbichler, W.A. Wall, C. Meier, A highly efficient computational approach for fast scan-resolved microstructure predictions in metal additive manufacturing on the scale of real parts, *Addit. Manuf.* 92 (2024) 104380.
- [58] A.C.I. Malossi, Y. Ineichen, C. Bekas, A. Curioni, Fast exponential computation on SIMD architectures, in: *Proc. of HIPEAC-WAPCO*, Amsterdam NL, Vol. 56, 2015, p. 224.
- [59] F. Perini, R.D. Reitz, Fast approximations of exponential and logarithm functions combined with efficient storage/retrieval for combustion kinetics calculations, *Combust. Flame* 194 (2018) 37–51.
- [60] Y. Saad, *Iterative methods for sparse linear systems*, Society for Industrial and Applied Mathematics, Panama, 2003.
- [61] P.E.J. Vos, S.J. Sherwin, R.M. Kirby, From h to p efficiently: Implementing finite and spectral/ hp element methods to achieve optimal performance for low- and high-order discretisations, *J. Comput. Phys.* 229 (13) (2010) 5161–5181.
- [62] M. Kronbichler, K. Ljungkvist, Multigrid for matrix-free high-order finite element computations on graphics processors, *ACM Trans. Parallel Comput.* 6 (1) (2019).
- [63] A. Brandt, Multi-level adaptive solutions to boundary-value problems, *Math. Comp.* 31 (138) (1977) 333–390.
- [64] W. Hackbusch, *Multi-grid methods and applications*, Springer, Berlin Heidelberg, 1985.
- [65] U. Trottenberg, C. Oosterlee, A. Schüller, *Multigrid*, Elsevier Academic Press, London, 2001.
- [66] M. Adams, M. Brezina, J. Hu, R. Tuminaro, Parallel multigrid smoothing: Polynomial versus Gauss–seidel, *J. Comput. Phys.* 188 (2) (2003) 593–610.
- [67] M.A. Heroux, J.M. Willenbring, A new overview of the Trilinos project, *Sci. Program.* 20 (2) (2012) 83–88.
- [68] M.W. Gee, C.M. Siefert, J.J. Hu, R.S. Tuminaro, M.G. Sala, *ML 5.0 smoothed aggregation user's guide*, Technical Report SAND2006-2649, Sandia National Laboratories, 2006.
- [69] J. Treibig, G. Hager, G. Wellein, Likwid: A lightweight performance-oriented tool suite for x86 multicore environments, in: *2010 39th International Conference on Parallel Processing Workshops*, IEEE, 2010, pp. 207–216.
- [70] D. Bošnjak, A. Pepe, R. Schussnig, D. Schmalstieg, T.-P. Fries, Higher-order block-structured hex meshing of tubular structures, *Eng. Comput.* 40 (2) (2024) 931–951.
- [71] W.L. Wood, M. Bossak, O.C. Zienkiewicz, An alpha modification of Newmark's method, *Internat. J. Numer. Methods Engrg.* 15 (10) (1980) 1562–1566.
- [72] P. Munch, V. Ivannikov, C. Cyron, M. Kronbichler, On the construction of an efficient finite-element solver for phase-field simulations of many-particle solid-state-sintering processes, *Comput. Mater. Sci.* 231 (2024) 112589.