



Modified augmented Lagrangian preconditioning for mixed-dimensional beam-solid coupling

Max Firmbach ^a, Ivo Steinbrecher ^a, Alexander Popp ^{a,b}, Matthias Mayr ^{a,b}

^a Institute for Mathematics and Computer-Based Simulation, Universität der Bundeswehr München, Werner-Heisenberg-Weg 39, D-85577 Neubiberg, Germany

^b Data Science & Computing Lab, Universität der Bundeswehr München, Werner-Heisenberg-Weg 39, D-85577 Neubiberg, Germany

ARTICLE INFO

Dataset link: <https://github.com/trilinos/Trilinos>, <https://github.com/4C-multiphysics/4C>

Keywords:

Modified augmented Lagrangian preconditioning
Block triangular preconditioner
Algebraic multigrid
Mixed-dimensional modeling
Beam-solid interaction

ABSTRACT

This paper presents modified augmented Lagrangian block preconditioners for the mixed-dimensional coupling of three-dimensional solid bodies with embedded one-dimensional torsion-free Kirchhoff–Love beams using Lagrange multipliers for constraint enforcement. The finite element discretization of this mixed formulation leads to an indefinite saddle-point system. An augmented Lagrangian formulation is employed to regularize the linear system while maintaining exact enforcement of the coupling constraints. Starting from the corresponding ideal augmented Lagrangian block preconditioner, more practical block-triangular variants are derived in which the solid, beam, and Schur complement blocks can be treated independently. In addition, different variants of Schur complement approximations are introduced. Numerical experiments investigate the dependence on model parameters and mesh refinement, as well as weak and strong scalability for short-fiber composite representative volume elements. For the tested parameter regime, a suitable choice of the Schur complement approximation yields favorable results, indicating the suitability of the proposed approach for large-scale simulations of mixed-dimensional models in solid and structural mechanics, as illustrated by an engineering example involving a composite sandwich plate.

1. Introduction

In both nature and technical systems, fibers embedded into solids improve the functional properties of the coupled system, for example by serving as reinforcements to improve load-bearing capacities under tensile loading. Examples can be found in many fields: In human biology, collagen fibers serve as the ubiquitous load-bearing and reinforcing element on the nanometer scale and, thus, form an important structural basis in both healthy and diseased tissue, for example in arterial walls [1] or for the design of scaffolds in regenerative medicine [2]. In civil engineering, fiber-reinforced concrete, i.e., concrete with randomly dispersed short steel fibers, comes with many advantages such as improved structural strength (especially in the tensile regime), crack width reduction, improved abrasion- and impact-resistance, and even post-cracking tensile strength, cf. [3–5] among others. To analyze and design such systems, numerical modeling and simulation nowadays plays a key role. Naturally, accurate but efficient models are required to obtain useful and predictive simulation capabilities.

Different modeling approaches are available for the embedding of slender fibers into solid continua (see Fig. 1): On the one hand, homogenization techniques incorporate all fiber information into the bulk constitutive law, usually leading to anisotropic

* Corresponding author.

E-mail addresses: max.firmbach@unibw.de (M. Firmbach), ivo.steinbrecher@unibw.de (I. Steinbrecher), alexander.popp@unibw.de (A. Popp), matthias.mayr@unibw.de (M. Mayr).

<https://doi.org/10.1016/j.cma.2026.119425>

Received 4 June 2026; Received in revised form 20 August 2026; Accepted 10 September 2026

Available online 30 September 2026

0045-7825/© 2026 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/>).
Open access funding enabled and organized by Projekt DEAL.




Fig. 1. Spectrum of modeling techniques for fibers embedded into three-dimensional solids [6].

formulations with preferential directions along the fiber orientation [7]. On the other hand, both the embedding solid as well as the fibers can be modeled as distinct 3D solid bodies with coupling conditions along the surfaces of the fibers. While this fully resolved approach delivers the highest modeling accuracy, its computational effort renders it infeasible for practical application scenarios with many fibers. As a remedy, mixed-dimensional modeling techniques exploit the slenderness of the embedded fibers to represent them by reduced-dimensional structural models such as trusses or beams, which are then embedded into three-dimensional solid bodies. A first Gauss-point-to-segment approach using spring-like “junction elements” has been outlined in [8]. A hybrid approach, where fibers are modeled as fully resolved three-dimensional (3D) continua embedded into the solid via the extended finite element method in the zone of interest and are represented by simple truss elements elsewhere, has been proposed in [9]. A variationally consistent overlapping domain decomposition approach for mixed-dimensional beam-solid coupling has been described in [10]. Recently, we have proposed a mixed-dimensional approach to fiber-solid coupling using a mortar-type method for the constraint discretization combined with a penalty regularization for the solution process [6,11,12]. In addition, we have developed an approximate block factorization preconditioner to tackle the specific challenges of the penalty regularization [13]. Furthermore, mixed-dimensional coupling is successfully used for fibers immersed in fluid domains [14–16].

Constraint enforcement based on penalty regularization, as used in our prior work [6,12,13], results in reduced system sizes, a straightforward and robust nonlinear solution procedure, and ease of implementation as well as are readily applicable for many engineering systems. However, their resulting formulation comes with several well-known drawbacks including variational inconsistency, ill-conditioning of the linear system due to the penalty parameter, and inexact constraint enforcement. In contrast, the direct use of a Lagrange multiplier field gives rise to a linear system of equations exhibiting saddle-point structure [6, Equation 38]. Yet, it enables exact constraint enforcement as well as alleviates the user from choosing a penalty parameter, that potentially harms the conditioning and solvability of the arising linear system of equations. Still, the saddle-point structure of the discretized linear system poses a challenge, as preconditioners need to be tailored to the system properties in order to yield an effective iterative solution process. For indefinite linear systems arising in computational physics, a variety of factorization-based preconditioning strategies is available [17]. An extensive summary and classification of preconditioning strategies for saddle-point systems can be found in [18].

So far, most of the research is concerned with the efficient preconditioning of problems related to fluid mechanics, where the discretization of the Navier–Stokes equation with finite elements results in a saddle-point system [19–22]. In solid mechanics, the need for saddle-point block preconditioners arises when additional constraints in the form of Lagrange multipliers are added to the system as in contact mechanics [23–25] or fracture mechanics [26] for example. Furthermore, the coupling of domains with different dimensionality has been realized with Lagrange multipliers. The resulting mixed-dimensional formulation is also of saddle-point structure and, thus, needs special treatment [27–29]. Another popular class of preconditioners for saddle-point systems is based on augmented Lagrangian methods [18]. Recently, such methods have been successfully applied to fictitious domain formulations and the associated interface problems [30,31]. Except for a diagonal preconditioner based on the augmented Lagrangian being outlined in [32,33], their application to mixed-dimensional formulations remains mostly unexplored so far. Furthermore, existing work on specialized preconditioners for mixed-dimensional modeling is limited to transport phenomena in the one-dimensional (1D) domain. To this end, scalable block preconditioners for mixed-dimensional beam-solid coupling problems governed by vector-valued elasticity formulations and using Lagrange multipliers for constraint enforcement remain largely unexplored.

In this contribution, we develop modified augmented Lagrangian block preconditioners for mixed-dimensional beam-solid coupling problems with Lagrange multiplier constraint enforcement. The resulting finite element discretization gives rise to an indefinite saddle-point system whose efficient solution requires suitable Schur complement approximations and is further complicated by the presence of a singular beam sub-problem. The singular beam sub-problem arises because the considered fully embedded beams are a pure Neumann problem, only the coupling to the surrounding solid constrains the beam rigid body modes. Our main idea is to reformulate the original Lagrange multiplier system in an augmented Lagrangian framework. This reformulation is motivated not only by the preconditioner construction, but also by the pure Neumann character of the beam sub-problem, which gets regularized by the augmented term. Starting from the corresponding ideal augmented Lagrangian block preconditioner, we derive practical block-triangular preconditioners that permit independent treatment of the solid, beam, and Schur complement blocks. In this context, we investigate several Schur complement approximations, including variants based on sparse approximate inverses, spectral equivalence properties of the augmentation term, and separate augmentation parameters for the solid and beam contributions. Further, we analyze the influence of physical, geometric, discretization, and solver parameters on the resulting




Fig. 2. Geometrical setup and the corresponding notation showing the three-dimensional solid domain Ω , its boundary $\partial\Omega$ and the one-dimensional beam centerline Λ with its boundary $\partial\Lambda$.

methods and investigate the convergence behavior under mesh-refinement as well as weak and strong scalability on parallel computing clusters. Finally, we use the proposed approach for large-scale mixed-dimensional simulations in solid and structural mechanics.

The remainder of this manuscript is structured as follows. Section 2 presents the relevant equations describing the coupling approach and introduces the finite element discretization used for the mixed-dimensional beam-solid problem, which leads to a block system with saddle-point structure. Section 3 discusses important features of the arising block system, the relevant modeling parameters, and positions the modeling scheme compared to existing approaches and the current state of literature. Section 4 introduces block preconditioners tailored to the linear system, discusses different variants and Schur complement approximations for practical computations and compares the preconditioning approach to existing methods in literature. Section 5 provides numerical experiments to validate and show the findings related to the proposed preconditioner. The results include insights into the choice of governing hyper-parameters, modeling parameter influence, strong and weak scaling behavior as well as the applicability to real world problems. Section 6 summarizes our findings and hints at possible future research directions.

2. Modeling approach of the mixed-dimensional beam-solid problem

We now introduce the governing equations and discretization used to model the coupling of a three-dimensional solid body with one-dimensional fibers. First, we introduce each physical field individually before combining them to form the overall coupled problem. An illustration of the geometrical setting is given in Fig. 2 showing the solid domain Ω , its boundary $\partial\Omega$ and the beam centerline Λ with its boundary $\partial\Lambda$. As the main focus of this publication is the construction of robust and efficient preconditioning techniques, we mainly give a brief overview of the mechanical theory. The interested reader is referred to [6] for more details.

Before dealing with equations, we introduce some notation that is used throughout the remainder of this manuscript: All quantities denoted with $(\cdot)^S$ are associated with the solid continuum, while quantities with superscript $(\cdot)^B$ are related to beam contributions. In addition, let D be a bounded domain in $\mathbb{R}^d$ with $d = 1, 2, 3$. Then, $H^s(D)$ for $s \in \mathbb{N}_0$ are the classical Sobolev spaces. As we solely treat vector-valued unknowns with d components, we also introduce the corresponding product Sobolev spaces $[H^s(D)]^d$. To distinguish continuous quantities from their discrete counterparts, we use the subscript $(\cdot)_h$ to denote the latter. In addition, δ indicates virtual, but kinematically admissible quantities in line with the concept of virtual work. Tensors of second order and higher are denoted in bold font.

2.1. Pure solid problem

The solid domain $\Omega \subset \mathbb{R}^3$ is governed by the equations of quasi-static nonlinear elasticity, stated in the material configuration as

$$\begin{aligned} \text{Div } \mathbf{P} + \hat{\mathbf{b}} &= 0 & \text{in } \Omega, \\ u^S &= u_D^S & \text{on } \Gamma_D, \\ \mathbf{P} \cdot \mathbf{n} &= \hat{\mathbf{t}} & \text{on } \Gamma_N, \end{aligned} \quad (1)$$

where Γ_D and Γ_N denote the Dirichlet and Neumann parts of the boundary $\partial\Omega = \Gamma_D \cup \Gamma_N$ with $\Gamma_D \cap \Gamma_N = \emptyset$, respectively. Hereby, $\mathbf{P}$ represents the first Piola–Kirchhoff stress tensor, while u^S defines the displacement, $\hat{\mathbf{b}}$ the body forces acting on the solid continuum, $\hat{\mathbf{t}}$ the tractions at the Neumann boundary and $\mathbf{n}$ the outward-pointing normal vector in the material configuration. We choose $\mathcal{V}^S(\Omega) := [H^1(\Omega)]^3$ as the underlying Sobolev space for test and basis functions. Further, we transform (1) into the variational formulation:

Find the displacement vector $u^S = (u_x^S, u_y^S, u_z^S) \in \mathcal{V}_D^S(\Omega) := \{u^S \in \mathcal{V}^S(\Omega) \mid \text{trace}_{\Gamma_D} u^S = u_D^S\}$ such that

$$\delta \mathcal{W}^S = 0 \quad \forall \delta u^S \in \mathcal{V}_0^S(\Omega) := \{\delta u^S \in \mathcal{V}^S(\Omega) \mid \text{trace}_{\Gamma_D} \delta u^S = 0\}. \quad (2)$$

Therein, we define the overall virtual work $\delta\mathcal{W}^S$ as difference between the variation of the potential energy Π^S of the solid and the respective virtual work of the external load reading

$$\delta\mathcal{W}^S = \delta\Pi_{\text{int}}^S - \delta\mathcal{W}_{\text{ext}}^S = \int_{\Omega} \mathbf{S} : \delta\mathbf{E} \, dV - \int_{\Omega} \hat{\mathbf{b}} \cdot \delta\mathbf{u}^S \, dV - \int_{\Gamma_N} \hat{\mathbf{t}} \cdot \delta\mathbf{u}^S \, dA,$$

with $\mathbf{S} := \mathbf{F}^{-1} \cdot \mathbf{P}$ and $\mathbf{E}$ representing the second Piola–Kirchhoff stress tensor and the Green–Lagrange strain tensor, respectively. Here, $\mathbf{F}$ describes the deformation gradient. In addition, we define strain tensor $\mathbf{E}$ and stress tensor $\mathbf{S}$ as

$$\mathbf{E} := \frac{1}{2}(\mathbf{F}^T \cdot \mathbf{F} - \mathbf{I}) \quad \text{and} \quad \mathbf{S} = \frac{\partial\mathcal{P}}{\partial\mathbf{E}},$$

with strain energy density $\mathcal{P}$. For a St.-Venant–Kirchhoff material, $\mathcal{P}$ follows as

$$\mathcal{P} := \frac{E^S \nu}{2(1+\nu)(1-2\nu)} (\text{trace } \mathbf{E})^2 + \frac{E^S}{2(1+\nu)} \mathbf{E} : \mathbf{E},$$

where E^S represents the solid Young’s modulus and ν the Poisson’s ratio of the solid body. Without loss of generality, we restrict ourselves to a St.-Venant–Kirchhoff material model in the scope of this publication. By discretizing (2) with continuous linear Lagrangean finite elements, we obtain a linear system to be solved stated as

$$\mathbf{K}^S \mathbf{u}_h^S = \mathbf{f}^S,$$

with stiffness matrix $\mathbf{K}^S \in \mathbb{R}^{n_{\text{DOF}}^S \times n_{\text{DOF}}^S}$, load vector $\mathbf{f}^S \in \mathbb{R}^{n_{\text{DOF}}^S}$ and $\mathbf{u}_h^S \in \mathbb{R}^{n_{\text{DOF}}^S}$ being the vector of the discrete displacement degrees of freedom.

2.2. Pure beam problem

We consider torsion-free beams based on the Kirchhoff–Love formulation [34,35], assuming vanishing shear deformations and no torsion contribution to the internal elastic energy, which results in a purely displacement-based formulation. These assumptions are valid for fibers with high slenderness ratios, a double symmetric cross-section and a straight centerline in the reference configuration. The resulting formulation represents a simplification, as rotational degrees of freedom are not considered explicitly. This modeling choice is made in view of the primary focus of the present work, namely the investigation of the proposed block preconditioning strategy rather than the treatment of advanced beam kinematics. At the same time, the concepts developed in this work are not restricted to the present setting and could, in principle, also be extended to more general geometrically exact beam formulations that account for torsion and shear effects, like the Simo–Reissner beam theory [36,37]. We employ beams with a constant circular cross-section along the beam centerline Λ . For a more detailed introduction on geometrically exact beam formulations, we refer to [38,39]. We define $\mathcal{V}^B(\Lambda) := [H^2(\Lambda)]^3$ and consider the following variational formulation:

Find $\mathbf{u}^B = (u_x^B, u_y^B, u_z^B) \in \mathcal{V}_D^B(\Lambda) := \{\mathbf{u}^B \in \mathcal{V}^B(\Lambda) \mid \text{trace}_{\partial\Lambda_D} \mathbf{u}^B = \mathbf{u}_D^B\}$ such that

$$\delta\mathcal{W}^B = 0 \quad \forall \delta\mathbf{u}^B \in \mathcal{V}_0^B(\Lambda) := \{\delta\mathbf{u}^B \in \mathcal{V}^B(\Lambda) \mid \text{trace}_{\partial\Lambda_D} \delta\mathbf{u}^B = 0\}.$$

The global contribution of the beams to the overall virtual work is again given as the difference of the variation of the potential energy and the variation of the external work

$$\delta\mathcal{W}^B = \delta\Pi_{\text{int}}^B - \delta\mathcal{W}_{\text{ext}}^B \quad \text{with} \quad \Pi_{\text{int}}^B = \frac{1}{2} \int_{\Lambda} E^B A \epsilon^2 + E^B I \kappa^2 \, ds,$$

where E^B denotes the Young’s modulus, A the cross-section area, I the second moment of area, ϵ the axial tension and κ the scalar curvature of the beams, respectively. Due to the representation of the curvature κ by the Frenet–Serret vector, which solely depends on the beam centerline, the finite element discretization of the beam requires a C^1 continuous interpolation of the centerline along Λ . Therefore, we consider continuous finite elements based on cubic Hermite polynomials for the discretization [38], resulting in

$$\mathbf{K}^B \mathbf{u}_h^B = \mathbf{f}^B.$$

Hereby, $\mathbf{K}^B \in \mathbb{R}^{n_{\text{DOF}}^B \times n_{\text{DOF}}^B}$ represents the stiffness matrix, $\mathbf{f}^B \in \mathbb{R}^{n_{\text{DOF}}^B}$ the load vector and $\mathbf{u}_h^B \in \mathbb{R}^{n_{\text{DOF}}^B}$ the discretized beam displacement vector.

2.3. Coupled beam-solid system

In a final step, we consider the coupled system between the solid and the beam field. For the mixed-dimensional coupling of the 3D solid domain Ω and the 1D beam domain Λ , we introduce a Lagrange multiplier field $\lambda = (\lambda_x, \lambda_y, \lambda_z)$ to enforce

$$g := \mathbf{u}^S - \mathbf{u}^B = 0 \quad \text{on } \Lambda, \tag{3}$$

with the vector-valued gap function. We note that (3) is of Dirichlet type. The constraints given in (3) are referred to as positional coupling, since they enforce the position of the beam cross-section centroid to be coupled to the underlying solid. We remark that due to the torsion-free Kirchhoff–Love beam formulation a pure coupling of positions is sufficient. Considering more complex geometrically exact beam theories makes the coupling of rotational degrees of freedom necessary, as e.g. presented in [40], yet we

do not consider this approach in the scope of this publication. Due to the introduction of the Lagrange multiplier field, the extended overall virtual work of the constraint mixed-dimensional problem is defined as

$$\delta\mathcal{W} = \delta\mathcal{W}^S + \delta\mathcal{W}^B + \delta\Pi^\lambda = 0,$$

with the addition of the variation of the Lagrange multiplier potential $\delta\Pi^\lambda$. The Lagrange multiplier potential Π^λ and its variation $\delta\Pi^\lambda$ are given as

$$\Pi^\lambda = \int_A \lambda g \, ds \quad \text{and} \quad \delta\Pi^\lambda = \int_A \delta\lambda g \, ds + \int_A \lambda \delta g \, ds.$$

We define the product space for our displacement and Lagrange multiplier solution as $\mathcal{V}^S(\Omega) \times \mathcal{V}^B(\Lambda) \times \mathcal{Q}(\Lambda)$ where $\mathcal{Q}(\Lambda)$ denotes the Lagrange multiplier space. Thus, the final variational formulation of the coupled problem reads: Find $u = (u^S, u^B, \lambda) \in \mathcal{V}_D^S(\Omega) \times \mathcal{V}_D^B(\Lambda) \times \mathcal{Q}(\Lambda)$ such that

$$\delta\mathcal{W} = 0 \quad \forall \delta u^S \in \mathcal{V}_0^S(\Omega), \quad \forall \delta u^B \in \mathcal{V}_0^B(\Lambda), \quad \forall \delta\lambda \in \mathcal{Q}(\Lambda). \quad (4)$$

For the spatial discretization of (4), we consider a mortar-type approach for the discretization of the Lagrange multiplier field. As the continuous Lagrange multiplier field is defined along the beam centerline, the interpolation is defined along the one-dimensional beam elements. In the nomenclature of classical contact mechanics, the beam would be considered the source side, and the solid the target side [6,41]. After discretization with continuous finite elements, we obtain the linear system of equations of the coupled problem

$$\mathcal{A}u_h = f, \quad \text{with } \mathcal{A} := \begin{pmatrix} K^S & -M^T \\ -M & D \end{pmatrix} \quad \text{and} \quad f := \begin{pmatrix} f^S \\ f^B \\ g \end{pmatrix}, \quad (5)$$

with the discrete solution vector u_h sought in the corresponding product of solid, beam, and Lagrange multiplier finite element spaces and solved for in every iteration of the nonlinear Newton-type solver. By construction, we consider a non-conforming discretization arising from non-matching (and even mixed-dimensional) meshes across subdomain interfaces $\Omega \cap \Lambda$. This gives rise to the off-diagonal coupling operators $M \in \mathbb{R}^{n_{\text{DOF}}^S \times n_{\text{DOF}}^B}$ and $D \in \mathbb{R}^{n_{\text{DOF}}^B \times n_{\text{DOF}}^B}$. Again, the interested reader is referred to our previous work [6] for detailed derivations. In the absence of explicit Dirichlet boundary conditions on the embedded beams, the respective displacement unknowns are constrained only through the Lagrange multiplier field. Consequently, the discrete beam operator itself is singular, as it inherits the rigid-body nullspace of the corresponding pure Neumann beam problem. Nevertheless, the positional coupling eliminates this nullspace within the coupled system, so the singularity of the isolated beam operator does not compromise the well-posedness of the coupled formulation.

3. Analysis of the coupled problem

Finding a solution of the block system given in (5) in a fast and efficient manner for increasing problem sizes is only feasible with an iterative linear solver. These methods usually converge slowly without a properly constructed preconditioner $\mathcal{P}$. The aim of this publication is to find an appropriate block preconditioner specifically tailored to the underlying physical problem. For the sake of completeness, the preconditioned system reads

$$\mathcal{P}^{-1} \mathcal{A}u_h = \mathcal{P}^{-1} f.$$

In the scope of this publication, we consider block triangular preconditioners of the generic form

$$\mathcal{P} := \begin{pmatrix} K^S & & \\ & K^B & \\ -M & D & S \end{pmatrix} \quad (6)$$

with the Schur complement $S := -M(K^S)^{-1}M^T - D(K^B)^{-1}D^T$. In the following, we see that this naive preconditioning approach is not feasible for the range of applications we consider. Nevertheless we propose to use (6) as a starting point for more sophisticated versions of $\mathcal{P}$ later on. Before constructing these block preconditioners though, we analyze the linear system given in (5) first. We consider the influence of the relevant modeling parameters and highlight important features of the block matrix structure. In a next step, we compare the introduced Lagrange multiplier coupling approach to methods relying on a penalty-based regularization to enforce the coupling constraints. In a last step, we discuss the solvability of the beam sub-problem.

Before starting the analysis of the linear system, we introduce the scaling matrix $W \in \mathbb{R}^{n_{\text{DOF}}^S \times n_{\text{DOF}}^S}$ as it will be used throughout the remainder of this manuscript. We postpone the discussion on how to choose W to later. For now, we state that W is derived from the findings given in [6] being represented by a diagonal matrix of mass-matrix-type at the coupling interface Λ .

3.1. Influence of modeling parameters and modeling assumptions

The mixed-dimensional coupling between a solid continuum and fibers is dominated by a few modeling parameters, which highly influence the conditioning of the linear system given in (5). Those parameters can be related to the material models, the geometry, and the discretization of the coupled problem. We shortly want to introduce these parameters and discuss their importance for

Table 1

Overview of the relevant modeling parameters with their notations and units.

Notation	Parameter	Unit
E^S	Young's modulus of the solid domain	N/m ²
h^S	Solid element size	m
V^S	Volume of the solid domain	m ³
E^B	Young's modulus of the beam domain	N/m ²
h^B	Beam element length	m
R^B	Beam cross-section radius	m
V^B	Volume of the beam domain	m ³
$\mathcal{V} = V^B/V^S$	Beam-to-solid volume ratio	–
$\mathcal{E} = E^B/E^S$	Beam-to-solid stiffness ratio	–
$\mathcal{H} = h^B/h^S$	Beam-to-solid element size ratio	–

our range of applications. We note that our choice regarding the discretization and geometric relations is based on the original mixed-dimensional model [6], in particular the assumption of $h^S \geq 2R^B$, i.e., that the solid mesh size h^S needs to be larger than the beam cross-section dimensions. Up until that point, the mixed-dimensional model is valid and produces similar results to a fully resolved coupling (i.e., coupling on the beam surface). If h^S becomes smaller than the beam radius R^B and violates the modeling assumptions of [6], the mixed-dimensional approach reproduces a singular solution artificially introduced by the 1D-3D modeling assumptions, which is unphysical. Thus, such cases are not considered in this publication.

We further note that the mathematical analysis of the well-posedness of the continuous 1D-3D coupling entails additional difficulties that are not addressed in the present work. In particular, for the standard solid space $[H^1(\Omega)]^3$, the trace of the solid displacement on the 1D beam centerline Λ is not well defined. Conversely, a Lagrange multiplier distributed along Λ acts as a line load on the three-dimensional continuum and does not belong to the dual of the natural solid energy space. A rigorous treatment of the continuous 1D-3D problem therefore requires, for instance, weighted Sobolev spaces or suitably defined fractional/operator-based norms, see [42,43] and related analyses in [27,44,45]. For the cases considered in this work, the mortar-type discretization and mass-type scaling introduced in Section 2.3 are discrete devices that mask the underlying continuous ill-posedness problem. This mathematical limitation is consistent with the modeling restriction $h^S \geq 2R^B$ introduced in [6] based on mechanical considerations. The present formulation is intended for a regime in which the beam cross-section remains unresolved by the solid discretization, rather than as an approximation of the singular continuous 1D-3D problem under arbitrary mesh refinement.

We use [46–48] as guidelines to match the modeling parameters for practical application cases. An overview of all relevant parameters is given in Table 1. We first discuss the parameter related to the material, which is given by the stiffness ratio $\mathcal{E} := E^B/E^S$ of the Young's modulus of the beam and the solid, respectively. We assume $E^B > E^S$, which can be interpreted as the beams acting as a reinforcement of the solid continuum. Usually, the ratio varies between $\mathcal{E} \approx 2$ for metal-based composites and can go up to $\mathcal{E} > 1000$ for niche applications. We mostly consider $\mathcal{E} \in [10, 1000]$, as one would find for example in steel-reinforced concrete, short fiber polymer matrix applications or in carbon reinforced composites. Increasing values of the stiffness ratio result in a scaling mismatch and therefore result in a worse conditioning of (5), thus complicating the solvability with an iterative linear solver.

The second group of parameters is related to geometric and discretization-based quantities given by the beam cross-section radius R^B , the characteristic beam element length h^B and the beam-to-solid volume ratio $\mathcal{V} := V^B/V^S$. In order not to violate the geometrically exact beam finite element formulation, we assume beam slenderness ratios $\xi^B := h^B/(2R^B) \gg 1$. A decreasing beam cross-section radius results in an increasing condition number of the beam sub-matrix in (5). In addition, we assume the solid mesh size h^S to be greater or equal the diameter of the beam element $h^S \geq (2R^B)$ and expect the solid mesh size to be smaller than the characteristic beam element length $h^S \leq h^B$ [6]. We denote the ratio between characteristic beam element length and solid mesh size by $\mathcal{H} := h^B/h^S$. In engineering applications, the amount of beams inside a solid is usually defined by the beam-to-solid volume ratio $\mathcal{V}$, which is indirectly related to the geometric quantities introduced before. We consider the beam-to-solid volume ratio to be at maximum $\mathcal{V} \approx 10\%$. We argue that higher values can be modeled more efficiently by a homogenization of the material properties, without explicitly modeling the beams. We consider application cases like they would appear for steel-reinforced concrete, which usually feature $\mathcal{V} \approx 1\% - 3\%$. The beam-to-solid volume ratio for short fiber polymer matrix composites can be $\mathcal{V} \approx 10\% - 30\%$. High ratios of $\mathcal{V}$ potentially increase the bandwidth of the coupling matrices D and M and therefore also of the Schur complement, leaving preconditioning and solving of (5) an even more difficult task.

3.2. Influence of block matrix structure and sub-solves

By enforcing the positional Dirichlet coupling (3) by a Lagrange multiplier field, the resulting linear operator has saddle-point structure with a zero (3,3)-block on the main diagonal. Due to the saddle-point structure, the overall block system is indefinite. Both the solid and beam stiffness matrices are stemming from finite element discretizations of elliptic operators and feature positive definiteness and are in addition also symmetric. As we use a non-symmetric preconditioner of block triangular structure (as given in (6)), we consider a generalized minimal residual method (GMRES) [49] as linear solver for the remainder of this manuscript. Further, the saddle-point structure of the block matrix restricts the type of preconditioners to be applied. Due to the zero block on the main block diagonal, black-box methods like block Jacobi or block Gauss–Seidel are not applicable, further emphasizing the use of Schur

complement based approximate block factorization preconditioners. As the solid problem is of elliptic nature, standard algebraic multigrid (AMG) [50,51] methods present themselves as favorable option for preconditioning the respective matrix sub-block. For our application case of several short and independent fibers, we consider only interactions between the fibers and the surrounding solid continuum, but not between the fibers. This results in the matrix K^B featuring a block-diagonal sparsity pattern. The size of each diagonal block depends on the amount of beam elements used to discretize a fiber. Therefore, direct methods and the corresponding factorization are still efficient to a certain degree even for systems with many fibers. While multigrid methods would in general work too, their application comes with challenges that are still topic of current research, as highlighted in [13, Remark 4.1]. As we employ different discretizations on the solid, beam and Lagrange multiplier field, and therefore have non-matching interfaces, the coupling matrices M and D have block structure, but feature no other special sparsity pattern, for example such as being diagonal like in [27]. As the coupling matrices M and D potentially have a large bandwidth, the Schur complement itself might also feature many non-zero entries. Still, the system size is usually small, which again makes direct methods feasible.

3.3. Comparison to a penalty-based formulation

Another popular approach in computational mechanics to ease the solvability of (5) is the application of a penalty regularization. Hereby, the positional coupling (3) is enforced using the relation

$$\lambda = \epsilon W^{-1} g \quad (7)$$

with penalty parameter $\epsilon \in \mathbb{R}^+$, resulting in a purely displacement-based formulation

$$\mathcal{A}u_h = f, \quad \text{with } \mathcal{A} = \begin{pmatrix} K^S + \epsilon M^T W^{-1} M & -\epsilon M^T W^{-1} D \\ -\epsilon D^T W^{-1} M & K^B + \epsilon D^T W^{-1} D \end{pmatrix} \quad \text{and } f = \begin{pmatrix} f^S - \epsilon M^T W^{-1} g \\ f^B + \epsilon D^T W^{-1} g \end{pmatrix}. \quad (8)$$

Undeniably, one advantage of the penalty-based formulation is the reduced system size due to eliminating the Lagrange multiplier field as an independent unknown through the penalty regularization (7). Yet, we argue that there are also several disadvantages in using (8), which are mainly related to preconditioning and efficient solvability. The regularization introduces a severe ill-conditioning arising from the penalty parameter ϵ as well as destroys block-diagonal dominance of $\mathcal{A}$ in (8) as discussed in our prior work [13]. To enforce the constraints (3) properly, the penalty parameter should be chosen $\epsilon \approx E^B$ [6], which in return results in a high condition number of the system matrix. Therefore, preconditioners specifically tailored to the problem formulation are necessary. One common approach is the use of block factorization based preconditioners. The appearing approximated Schur complement in such methods usually results in a sparse matrix with comparably high bandwidth due to the introduction of the penalty terms. Therefore, iteratively solving and preconditioning (8) is difficult and computationally demanding, especially leaving traditional smoothing-based methods as they usually appear in multigrid methods inefficient. While progress has been made recently related to preconditioning of the penalty-regularized system based on multilevel algorithms [52], its application to (8) is not trivial as the coupling is enforced on two fundamentally different partial differential equations (PDE). On the other hand, block preconditioning approaches as given in [13] come with a high computational cost.

3.4. Pure Neumann problem related to K^B

As we consider fibers completely embedded into the solid body, usually no explicit Dirichlet conditions are imposed onto the one-dimensional beam domain. When looking at the entire block system, the rigid body modes of the fibers are fully constrained by the Lagrange multiplier field that enforces the positional coupling to the solid domain. Hence, the overall system is uniquely solvable. Things change however, when considering the individual sub-blocks as done inside the block preconditioner and thus the linear solver (given in (6)). Therein, the sub-block related to the beam contribution K^B is singular after discretization due to the missing Dirichlet conditions on the beam centerline Λ and its boundary $\partial\Lambda$. Therefore, the sub-block is not uniquely solvable as solutions are only determined up to the rigid body modes.

Different approaches to solve pure Neumann problems are presented in literature with the so-called null space playing a crucial role. For three-dimensional elasticity problems, the null space is defined by the space of rigid body modes consisting of three translations and three rotations. In our case, the Neumann problem is defined on the domain governed by the equations of a torsion-free Kirchhoff-Love beam formulation. Due to the fact, that the considered beams cannot represent torsion deformation, the null space is defined by only five components, which fully describe the space of rigid body modes.

The null space can be used to enforce additional constraints by Lagrange multipliers, which make the system uniquely solvable [53,54]. However, this requires to solve yet another saddle-point system. Another approach considers a projection operator to remove the null space from the solution of the linear system [53,54]. In addition, one can apply a rank- m correction with m being the number of independent null space components to shift the critical eigenvalues to make the linear operator invertible. However, the resulting projection matrix is usually dense and therefore difficult to treat [55]. While deflation techniques related to iterative linear solvers can avoid the explicit construction of the projected matrix, their implementation is invasive, making the use of already available software libraries inconvenient [56].

In the scope of this publication, we do not further explore null space related approaches, but try to reformulate the linear system (5) itself to avoid the pure Neumann problem in the first place, as is shown in Section 4.

4. Block preconditioning of the mixed-dimensional problem

We now discuss the novelty and main idea of this publication: the construction of modified augmented Lagrangian block preconditioners for the mixed-dimensional beam-solid coupled problem. As our naive preconditioning approach (6) is not directly applicable to the coupled problem and based on the general findings discussed in Section 3, we propose to reformulate (5) into its augmented formulation [57–59] reading

$$\mathcal{A}u_h = f, \text{ with } \mathcal{A} = \begin{pmatrix} K^S + \epsilon M^T W^{-1} M & -\epsilon M^T W^{-1} D & -M^T \\ -\epsilon D^T W^{-1} M & K^B + \epsilon D^T W^{-1} D & D^T \\ -M & D & \end{pmatrix} \text{ and } f = \begin{pmatrix} f^S - \epsilon M^T W^{-1} g \\ f^B + \epsilon D^T W^{-1} g \\ g \end{pmatrix}. \quad (9)$$

The linear system given in (9) can be seen as a combination of the pure Lagrange multiplier approach (5) and the method based on a penalty regularization (8). The augmented Lagrangian formulation combines the advantages of both approaches by keeping the saddle-point structure of the original system and therefore solving exactly for the constraints, but also adding a penalty contribution to stabilize the solution process [60]. The system given in (9) is therefore less sensitive to changes in parameters and better conditioned than the original one. In addition, the Schur complement is modified, which is discussed further in Section 4.2. We opt to use the augmented Lagrangian formulation not only in the preconditioner itself, but already assemble the full system during the nonlinear solution process, as the relevant terms are already available there. This also avoids the complexity of handling two system matrices, namely (5) and (9), at the same time for the nonlinear iteration, the preconditioner and the respective linear solver.

Considering the model parameters described in Section 3.1, their influence on the augmented system (9) does not change in comparison to the pure Lagrange multiplier formulation. Similar to the penalty regularization, we introduce one additional parameter to steer the constraint enforcement, given as the penalty parameter $\epsilon \in \mathbb{R}^+$. However, one important difference is the choice of ϵ : While for the penalty-regularized system $\epsilon \approx E^B$ should hold true, with the augmented formulation we have greater freedom in choosing the penalty parameter to our needs, without having an explicit restriction. As we end up with the original system for $\epsilon = 0$ and still solve for the original solution, ϵ can usually be chosen much smaller than E^B . This in turn allows us to control the ill-conditioning of the related matrix sub-blocks. We later see in Section 4.1, that ϵ plays an important role in the construction of the block preconditioner.

Following the discussion of Section 3.4, the addition of the penalty contribution can be seen as indirect enforcement of the coupling conditions. Matrix D encodes the coupling constraints between the beam and the Lagrange multiplier field and is injective on the rigid-body subspace of the embedded Kirchhoff–Love beams. In particular, it eliminates all five rigid-body modes of the beam, and thus $\text{null}(K^B) \cap \text{null}(D) = \{0\}$. This heals the problematic nature of the pure Neumann problem on the beam sub-block even for small values of the penalty parameter. This, therefore, results in an invertible augmented beam matrix sub-block, thus making direct methods applicable as outlined in Section 3.2. Further, considering the augmentation of the solid matrix still yields a symmetric and positive definite operator, which underlines the applicability of AMG.

The introduction of additional terms increases the density of the stiffness matrix, but the bandwidth can be controlled quite well due to the choice of W being a diagonal matrix [31,44]. In addition, coupling matrices directly acting on the solid and beam domain are added to the off-diagonal blocks, where before plain zero blocks appeared, changing the overall matrix structure and increasing its density. If the beam sub-matrix is block diagonal and the respective sub-blocks are comparably small, as it is the case for many, but short embedded fibers, the addition of penalty contributions is negligible, as the matrix bandwidth will not grow outside of the block diagonal sparsity pattern. All contributions related to the Lagrange multiplier field are not altered. The most influential part is on the solid stiffness K^S increasing the bandwidth of the respective matrix block. The overall system matrix becomes denser, yet this increase in entries can be controlled quite well. An illustration of the sparsity patterns of the saddle-point system (5) and the augmented version (9) for an exemplary case is shown in Fig. 3, with the overall amount of non-zero entries increased by $\approx 50\%$.

For the sake of simplified notation, we consider terms with subscript $(\cdot)_\epsilon$ to be related to the augmentation. Thus we can rewrite the augmented system (9) resulting in the compact notation


$$\mathcal{A}_\epsilon u_h = f_\epsilon, \text{ with } \mathcal{A}_\epsilon = \begin{pmatrix} K_\epsilon^S & -B_\epsilon^T & -M^T \\ -B_\epsilon & K_\epsilon^B & D^T \\ -M & D & \end{pmatrix} \text{ and } f_\epsilon = \begin{pmatrix} f_\epsilon^S \\ f_\epsilon^B \\ g \end{pmatrix}. \quad (10)$$

with the coupling matrix $B_\epsilon := \epsilon D^T W^{-1} M$. We also denote that ϵ has the unit N/m^2 , yet with a little abuse of notation, we omit the explicit mentioning of this property in most parts of the discussion.

4.1. Modified augmented Lagrangian block preconditioning

Based on the augmented linear system given in (10), we now construct a block preconditioner of the form given by $\mathcal{P}$ initially introduced in Section 3. An almost obvious preliminary choice is the use of an augmented Lagrangian preconditioner. To the authors' knowledge, this type of preconditioner has mostly been applied to the Stokes and Oseen equations [19,22] as well as for fictitious domain problems [31], yet its application to problems from computational solid mechanics is mainly limited to contact mechanics [24].

In the following analysis, we assume that $K^S \in \mathbb{R}^{n_{\text{DOF}}^S \times n_{\text{DOF}}^S}$ and $W \in \mathbb{R}^{n_{\text{DOF}}^A \times n_{\text{DOF}}^A}$ are symmetric positive definite. Furthermore, for $M \in \mathbb{R}^{n_{\text{DOF}}^A \times n_{\text{DOF}}^S}$ and $D \in \mathbb{R}^{n_{\text{DOF}}^A \times n_{\text{DOF}}^B}$, we assume $\text{rank}(M) = \text{rank}(D) = n_{\text{DOF}}^A$. Thus, M and D have full row rank, whereas full column rank is not required. To obtain a transparent comparison between the ideal and modified augmented Lagrangian



(a) Sparsity pattern of the coupled problem (5) showing the saddle-point structure (nnz = 302135). (b) Sparsity pattern of the augmented coupled problem (9) (nnz = 447651).

Fig. 3. Comparison of the sparsity patterns of the coupled problem and the augmented version, with the latter one showing an increase in the number of non-zeros by roughly 50%.

preconditioners, the following auxiliary analysis assumes that the rigid-body nullspace of $K^B \in \mathbb{R}^{n_{\text{DOF}}^B \times n_{\text{DOF}}^B}$, has been removed, so that K^B is also symmetric positive definite. This makes the respective inverse expressions and their asymptotic limits well defined and allows the changes in the preconditioned operator to be attributed specifically to the dropping of specific matrix sub-blocks, rather than to the singularity of the isolated beam sub-problem. In practice, the inverse of the singular matrix K^B is never applied, as all beam sub-solves involve the augmented block K_ϵ^B .

Ideal augmented Lagrangian preconditioner. The ideal augmented Lagrangian preconditioner based on the augmented mixed-dimensional beam-solid problem (10) reads

$$\mathcal{P}_\epsilon := \begin{pmatrix} K_\epsilon^S & -B_\epsilon^T & \\ -B_\epsilon & K_\epsilon^B & \\ -M & D & S_\epsilon \end{pmatrix}. \quad (11)$$

Similar to [20], the term “ideal” refers to using all the augmented sub-blocks inside the preconditioner. The respective action of $\mathcal{P}_\epsilon^{-1}$ is given as

$$\mathcal{P}_\epsilon^{-1} := \begin{pmatrix} I & & \\ & I & \\ & & (S_\epsilon)^{-1} \end{pmatrix} \begin{pmatrix} I & & \\ & I & \\ M & -D & I \end{pmatrix} \begin{pmatrix} (K_\epsilon^S & -B_\epsilon^T)^{-1} \\ -B_\epsilon & K_\epsilon^B \\ I \end{pmatrix}. \quad (12)$$

Hereby, S_ϵ is the augmented Schur complement

$$S_\epsilon := -(-M \quad D) \begin{pmatrix} K_\epsilon^S & -B_\epsilon^T \\ -B_\epsilon & K_\epsilon^B \end{pmatrix}^{-1} \begin{pmatrix} -M^T \\ D^T \end{pmatrix} \quad (13)$$

$$= -(-M \quad D) \left(\begin{pmatrix} K^S & \\ & K^B \end{pmatrix} + \epsilon \begin{pmatrix} -M^T \\ D^T \end{pmatrix} W^{-1} \begin{pmatrix} -M & D \end{pmatrix} \right)^{-1} \begin{pmatrix} -M^T \\ D^T \end{pmatrix}. \quad (14)$$

By applying the Sherman–Morrison–Woodbury formula (e.g., [61, Theorem 3.2]), we can rewrite (14) into

$$S_\epsilon^{-1} = - \left((M \quad -D) \begin{pmatrix} K^S & \\ & K^B \end{pmatrix}^{-1} \begin{pmatrix} -M^T \\ D^T \end{pmatrix} \right)^{-1} - \epsilon W^{-1},$$

which in the end simplifies to the statement

$$S_\epsilon^{-1} = S^{-1} - \epsilon W^{-1} = (-M (K^S)^{-1} M^T - D (K^B)^{-1} D^T)^{-1} - \epsilon W^{-1}. \quad (15)$$

The resulting inverse of the augmented Schur complement is a simple combination of S from the original system (5) and the contribution from the augmentation. Based on (15), we can argue that if $\epsilon \rightarrow \infty$ for a fixed mesh size h , the augmentation term ϵW^{-1} dominates S_ϵ^{-1} and leads to a spectrally equivalent representation if the scaling matrix W is properly chosen (e.g. [19, Remark

4.4] and [62]). This asymptotic observation concerns the exact augmented Schur complement only. Applying the ideal augmented Lagrangian preconditioner to our system matrix results in

$$\mathcal{P}_\epsilon^{-1} \mathcal{A}_\epsilon = \begin{pmatrix} I & & \\ & I & \\ & & S_\epsilon^{-1} \end{pmatrix} \begin{pmatrix} I & & \\ & I & \\ M & -D & I \end{pmatrix} \begin{pmatrix} K_\epsilon^S & -B_\epsilon^T \\ -B_\epsilon & K_\epsilon^B \end{pmatrix}^{-1} \begin{pmatrix} I & & \\ & I & \\ -M & D & I \end{pmatrix} \begin{pmatrix} K_\epsilon^S & -B_\epsilon^T & -M^T \\ -B_\epsilon & K_\epsilon^B & D^T \end{pmatrix} = \begin{pmatrix} I & & Q_{13} \\ & I & Q_{23} \\ & & Q_{33} \end{pmatrix},$$

with the individual matrix blocks Q_{13} , Q_{23} , and Q_{33} given as

$$\begin{aligned} Q_{13} &= (K_\epsilon^S)^{-1} (-M^T + B_\epsilon^T Q_{23}), \\ Q_{23} &= (K_\epsilon^B - B_\epsilon (K_\epsilon^S)^{-1} B_\epsilon^T)^{-1} (D^T - B_\epsilon (K_\epsilon^S)^{-1} M^T), \\ Q_{33} &= I. \end{aligned}$$

With these statements, we analyze the asymptotic behavior of the matrix blocks Q_{13} , Q_{23} , and Q_{33} given by the following limits

$$\begin{aligned} \lim_{\epsilon \rightarrow 0} \|Q_{13}\| &= \|(K^S)^{-1} M^T\|, & \lim_{\epsilon \rightarrow \infty} \|Q_{13}\| &= 0, \\ \lim_{\epsilon \rightarrow 0} \|Q_{23}\| &= \|(K^B)^{-1} D^T\|, & \lim_{\epsilon \rightarrow \infty} \|Q_{23}\| &= 0, \\ \lim_{\epsilon \rightarrow 0} \|Q_{33}\| &= 1, & \lim_{\epsilon \rightarrow \infty} \|Q_{33}\| &= 1. \end{aligned}$$

For the exact ideal augmented Lagrangian preconditioner and every penalty parameter $\epsilon > 0$, the preconditioned operator is unit upper triangular, and therefore all its eigenvalues are equal to one. For $\epsilon \rightarrow \infty$, matrix blocks Q_{13} and Q_{23} vanish and, ultimately, the preconditioned operator tends towards identity, $\mathcal{P}_\epsilon^{-1} \mathcal{A}_\epsilon \rightarrow I$. Thus, while its spectrum is independent of the penalty parameter, the preconditioned operator itself approaches the identity as the penalty parameter increases.

Modified augmented Lagrangian preconditioner. As the augmented sub-blocks represent the penalty-based system discussed in Section 3.3 and its preconditioner is non-trivial, as outlined in [13], the application of (12) is not optimal in the context of computationally efficient augmented Lagrangian preconditioning for mixed-dimensional beam-solid problems. The core component of the respective preconditioner would be the same as the one proposed in [13], yet the goal of this publication is to construct a more efficient method. Therefore, we aim to remove block B_ϵ^T from $\mathcal{P}_\epsilon$ to obtain a block triangular preconditioner resulting in a modified version of the ideal augmented Lagrangian preconditioner proposed in [20,63] and expanded in [64] stated as

$$\mathcal{P}_{\epsilon, \text{mod}} := \begin{pmatrix} K_\epsilon^S & & \\ -B_\epsilon & K_\epsilon^B & \\ -M & D & S_\epsilon \end{pmatrix}, \quad (16)$$

with its action defined by

$$\mathcal{P}_{\epsilon, \text{mod}}^{-1} := \begin{pmatrix} I & & \\ & I & \\ & & S_\epsilon^{-1} \end{pmatrix} \begin{pmatrix} I & & \\ & I & \\ M & -D & I \end{pmatrix} \begin{pmatrix} I & & \\ & (K_\epsilon^B)^{-1} B_\epsilon & I \\ & & I \end{pmatrix} \begin{pmatrix} (K_\epsilon^S)^{-1} & & \\ & (K_\epsilon^B)^{-1} & \\ & & I \end{pmatrix}. \quad (17)$$

The action of the preconditioner defined by $\mathcal{P}_{\epsilon, \text{mod}}^{-1}$ enables us to approximate the inverse of each main diagonal block individually, yet it violates the Schur complement assumption given above [20], due to the removal of block B_ϵ^T . Again, we consider the preconditioned operator, this time applying the modified augmented Lagrangian approach resulting in

$$\begin{aligned} \mathcal{P}_{\epsilon, \text{mod}}^{-1} \mathcal{A}_\epsilon &= \begin{pmatrix} I & & \\ & I & \\ & & S_\epsilon^{-1} \end{pmatrix} \begin{pmatrix} I & & \\ & I & \\ M & -D & I \end{pmatrix} \begin{pmatrix} I & & \\ & (K_\epsilon^B)^{-1} B_\epsilon & I \\ & & I \end{pmatrix} \begin{pmatrix} (K_\epsilon^S)^{-1} & & \\ & (K_\epsilon^B)^{-1} & \\ & & I \end{pmatrix} \begin{pmatrix} K_\epsilon^S & -B_\epsilon^T & -M^T \\ -B_\epsilon & K_\epsilon^B & D^T \\ -M & D & I \end{pmatrix} \\ &= \begin{pmatrix} I & Q_{12} & Q_{13} \\ & Q_{22} & Q_{23} \\ & Q_{32} & Q_{33} \end{pmatrix}. \end{aligned}$$

The individual matrix blocks $Q_{ij}, i \in \{1, 2, 3\}, j \in \{2, 3\}$ of $\mathcal{P}_{\epsilon, \text{mod}}^{-1} \mathcal{A}_\epsilon$ are therefore given as

$$\begin{aligned} Q_{12} &= -(K_\epsilon^S)^{-1} B_\epsilon^T, \\ Q_{13} &= -(K_\epsilon^S)^{-1} M^T, \\ Q_{22} &= I - (K_\epsilon^B)^{-1} B_\epsilon (K_\epsilon^S)^{-1} B_\epsilon^T, \\ Q_{23} &= (K_\epsilon^B)^{-1} (D^T - B_\epsilon (K_\epsilon^S)^{-1} M^T), \\ Q_{32} &= S_\epsilon^{-1} (-M (K_\epsilon^S)^{-1} B_\epsilon^T + D (K_\epsilon^B)^{-1} B_\epsilon (K_\epsilon^S)^{-1} B_\epsilon^T), \\ Q_{33} &= S_\epsilon^{-1} (-M (K_\epsilon^S)^{-1} M^T - D (K_\epsilon^B)^{-1} D^T + D (K_\epsilon^B)^{-1} B_\epsilon (K_\epsilon^S)^{-1} M^T). \end{aligned}$$

The resulting asymptotic limits for changing ϵ based on the given factorization are stated as follows:

$$\lim_{\epsilon \rightarrow 0} \|Q_{12}\| = 0, \quad \lim_{\epsilon \rightarrow \infty} \|Q_{12}\| = \|(K^S)^{-1} M^T (M (K^S)^{-1} M^T)^{-1} D\|,$$

$$\begin{aligned}
\lim_{\epsilon \rightarrow 0} \|Q_{13}\| &= \|(K^S)^{-1} M^T\|, & \lim_{\epsilon \rightarrow \infty} \|Q_{13}\| &= 0, \\
\lim_{\epsilon \rightarrow 0} \|Q_{22}\| &= 1, & \lim_{\epsilon \rightarrow \infty} \|Q_{22}\| &= \|I - (K^B)^{-1} D^T (D(K^B)^{-1} D^T)^{-1} D\|, \\
\lim_{\epsilon \rightarrow 0} \|Q_{23}\| &= \|(K^B)^{-1} D^T\|, & \lim_{\epsilon \rightarrow \infty} \|Q_{23}\| &= 0, \\
\lim_{\epsilon \rightarrow 0} \|Q_{32}\| &= 0, & \lim_{\epsilon \rightarrow \infty} \|Q_{32}\| &= \|(D(K^B)^{-1} D^T)^{-1} D\|, \\
\lim_{\epsilon \rightarrow 0} \|Q_{33}\| &= 1, & \lim_{\epsilon \rightarrow \infty} \|Q_{33}\| &= 1.
\end{aligned}$$

In contrast to the ideal augmented Lagrangian preconditioner, choosing $\epsilon \rightarrow 0$ or $\epsilon \rightarrow \infty$ does not result in $\mathcal{P}_\epsilon^{-1} \mathcal{A}_\epsilon \rightarrow I$ as individual terms grow or shrink. This is the direct consequence of dropping B_ϵ^T . Especially for $\epsilon \rightarrow \infty$, Q_{32} does not vanish and $Q_{22} \neq I$, thus destroying the unit upper triangular structure and modifying the eigenvalue spectrum. The competing limits suggest that an effective value of ϵ should lie between the two asymptotic regimes to obtain a favorable eigenvalue spectrum and $\mathcal{P}_\epsilon^{-1} \mathcal{A}_\epsilon \approx I$, but we remark, that the given analysis alone does not establish existence or uniqueness of an optimal value.

Individually scaled modified augmented Lagrangian preconditioner. In a final step, we introduce a block preconditioner based on the augmented Lagrangian formulation (9) applying the modified preconditioner (16), but with different scaling parameters for the solid and beam contributions. As the coupled system (5) is composed of two different PDEs, we follow the recently proposed approach from [30] and rewrite the augmented system (9) by introducing two independent penalty parameters ϵ^S and ϵ^B such that

$$\mathcal{A}u_h = f, \text{ with } \mathcal{A} = \begin{pmatrix} K^S + \epsilon^S M^T W^{-1} M & -\epsilon^S M^T W^{-1} D & -M^T \\ -\epsilon^B D^T W^{-1} M & K^B + \epsilon^B D^T W^{-1} D & D^T \\ -M & D & \end{pmatrix} \text{ and } f = \begin{pmatrix} f^S - \epsilon^S M^T W^{-1} g \\ f^B + \epsilon^B D^T W^{-1} g \\ g \end{pmatrix} \quad (18)$$

and therefore

$$\mathcal{A}_{\epsilon^S, \epsilon^B} u_h = f_{\epsilon^S, \epsilon^B}, \text{ with } \mathcal{A}_{\epsilon^S, \epsilon^B} = \begin{pmatrix} K_{\epsilon^S}^S & -B_{\epsilon^S}^T & -M^T \\ -B_{\epsilon^B}^B & K_{\epsilon^B}^B & D^T \\ -M & D & \end{pmatrix} \text{ and } f_{\epsilon^S, \epsilon^B} = \begin{pmatrix} f_{\epsilon^S}^S \\ f_{\epsilon^B}^B \\ g \end{pmatrix} \quad (19)$$

describes the independently row-scaled augmented system. Using the same modified augmented Lagrangian preconditioner as before applied to (19) yields

$$\mathcal{P}_{\epsilon^S, \epsilon^B, \text{mod}} := \begin{pmatrix} K_{\epsilon^S}^S & & \\ -B_{\epsilon^B}^B & K_{\epsilon^B}^B & \\ -M & D & S_{\epsilon^S, \epsilon^B} \end{pmatrix}. \quad (20)$$

We next consider the application of the individually scaled modified augmented Lagrangian preconditioner to the linear operator given in (19) stated by

$$\begin{aligned}
\mathcal{P}_{\epsilon^S, \epsilon^B, \text{mod}}^{-1} \mathcal{A}_{\epsilon^S, \epsilon^B} &= \begin{pmatrix} I & & \\ & I & \\ & & S_{\epsilon^S, \epsilon^B}^{-1} \end{pmatrix} \begin{pmatrix} I & & \\ & I & \\ M & -D & I \end{pmatrix} \begin{pmatrix} I & & \\ & (K_{\epsilon^B}^B)^{-1} B_{\epsilon^B} & \\ & I & \end{pmatrix} \begin{pmatrix} (K_{\epsilon^S}^S)^{-1} & & \\ & (K_{\epsilon^B}^B)^{-1} & \\ I & & I \end{pmatrix} \begin{pmatrix} K_{\epsilon^S}^S & -B_{\epsilon^S}^T & -M^T \\ -B_{\epsilon^B}^B & K_{\epsilon^B}^B & D^T \\ -M & D & \end{pmatrix} \\
&= \begin{pmatrix} I & Q_{12} & Q_{13} \\ Q_{22} & Q_{23} & \\ Q_{32} & Q_{33} & \end{pmatrix},
\end{aligned}$$

with the individual blocks defined as

$$\begin{aligned}
Q_{12} &= -(K_{\epsilon^S}^S)^{-1} B_{\epsilon^S}^T, \\
Q_{13} &= -(K_{\epsilon^S}^S)^{-1} M^T, \\
Q_{22} &= I - (K_{\epsilon^B}^B)^{-1} B_{\epsilon^B} (K_{\epsilon^S}^S)^{-1} B_{\epsilon^S}^T, \\
Q_{23} &= (K_{\epsilon^B}^B)^{-1} (D^T - B_{\epsilon^B} (K_{\epsilon^S}^S)^{-1} M^T), \\
Q_{32} &= S_{\epsilon^S, \epsilon^B}^{-1} (-M (K_{\epsilon^S}^S)^{-1} B_{\epsilon^S}^T + D (K_{\epsilon^B}^B)^{-1} B_{\epsilon^B} (K_{\epsilon^S}^S)^{-1} B_{\epsilon^S}^T), \\
Q_{33} &= S_{\epsilon^S, \epsilon^B}^{-1} (-M (K_{\epsilon^S}^S)^{-1} M^T - D (K_{\epsilon^B}^B)^{-1} D^T + D (K_{\epsilon^B}^B)^{-1} B_{\epsilon^B} (K_{\epsilon^S}^S)^{-1} M^T).
\end{aligned}$$

The respective exact Schur complement is denoted by

$$S_{\epsilon^S, \epsilon^B} := -W(W + \epsilon^S M (K^S)^{-1} M^T + \epsilon^B D (K^B)^{-1} D^T)^{-1} (M (K^S)^{-1} M^T + D (K^B)^{-1} D^T)$$

with its inverse

$$S_{\epsilon^S, \epsilon^B}^{-1} := -(M (K^S)^{-1} M^T + D (K^B)^{-1} D^T)^{-1} (W + \epsilon^S M (K^S)^{-1} M^T + \epsilon^B D (K^B)^{-1} D^T) W^{-1}.$$

We now consider the asymptotic behavior of the matrix blocks under varying ϵ^S , but ϵ^B being fixed:

$$\lim_{\epsilon^S \rightarrow 0} \|Q_{12}\| = 0, \quad \lim_{\epsilon^S \rightarrow \infty} \|Q_{12}\| = \|(K^S)^{-1} M^T (M (K^S)^{-1} M^T)^{-1} D\|,$$




Fig. 4. Model problem consisting of a cantilever beam with one fully embedded fiber.

$$\begin{aligned}
 \lim_{\epsilon^S \rightarrow 0} \|Q_{13}\| &= \|(K^S)^{-1} M^T\|, & \lim_{\epsilon^S \rightarrow \infty} \|Q_{13}\| &= 0, \\
 \lim_{\epsilon^S \rightarrow 0} \|Q_{22}\| &= 1, & \lim_{\epsilon^S \rightarrow \infty} \|Q_{22}\| &= \|(K_{\epsilon^B}^B)^{-1} K^B\|, \\
 \lim_{\epsilon^S \rightarrow 0} \|Q_{23}\| &= \|(K_{\epsilon^B}^B)^{-1} (D^T - B_{\epsilon^B} (K^S)^{-1} M^T)\|, & \lim_{\epsilon^S \rightarrow \infty} \|Q_{23}\| &= \|(K_{\epsilon^B}^B)^{-1} D^T\|, \\
 \lim_{\epsilon^S \rightarrow 0} \|Q_{32}\| &= 0, & \lim_{\epsilon^S \rightarrow \infty} \|Q_{32}\| &= \infty, \\
 \lim_{\epsilon^S \rightarrow 0} \|Q_{33}\| &= 1, & \lim_{\epsilon^S \rightarrow \infty} \|Q_{33}\| &= \infty.
 \end{aligned}$$

For fixed ϵ^B , choosing ϵ^S sufficiently small is favorable because the omitted block $B_{\epsilon^S}^T$ vanishes as $\epsilon^S \rightarrow 0$. In contrast, the preconditioned operator is generically unbounded as $\epsilon^S \rightarrow \infty$, due to Q_{32} and Q_{33} growing. In a final step we investigate the cases for $\epsilon^S \rightarrow 0$ and variable ϵ^B , which only affects Q_{23} :

$$\lim_{\substack{\epsilon^S \rightarrow 0 \\ \epsilon^B \rightarrow 0}} \|Q_{23}\| = \|(K^B)^{-1} D^T\|, \quad \lim_{\substack{\epsilon^S \rightarrow 0 \\ \epsilon^B \rightarrow \infty}} \|Q_{23}\| = \|-(K^B)^{-1} D^T (D(K^B)^{-1} D^T)^{-1} M (K^S)^{-1} M^T\|.$$

As the exact application of $\mathcal{P}_{\epsilon, \text{mod}}^{-1}$ and $\mathcal{P}_{\epsilon^S, \epsilon^B, \text{mod}}^{-1}$ is not feasible in an actual computation, we use approximated versions. Following the discussion in Sections 3.2 and 4, we opt to use AMG on the solid domain and an LU-factorization on the beam and Lagrange multiplier field, respectively. The last missing component of the block preconditioners therefore consists of finding a suitable approximation $\hat{S}_\epsilon$ to the exact augmented Schur complement S_ϵ and respectively $\hat{S}_{\epsilon^S, \epsilon^B}$ for $S_{\epsilon^S, \epsilon^B}$, as its exact computation is not feasible and would result in a dense matrix.

4.2. Choice of the augmented schur complement

In the practical block preconditioners considered in the scope of this publication, the Schur complement is no longer applied exactly, but replaced by an inexpensive matrix whose approximation quality must be assessed separately. Therefore, we introduce a small model problem to be used for numerical studies in the scope of this section, which is shown in Fig. 4. The numerical setup consists of a cantilever beam with dimensions $5 \text{ m} \times 1 \text{ m} \times 1 \text{ m}$ with a single fiber embedded into the center of the solid continuum. Both fiber end-cross-sections are equally far away from the respective perpendicular solid surfaces. The single fiber is discretized with four beam elements with a cross-section radius of $R^B = 0.125 \text{ m}$, beam element length $h^B = 1.0 \text{ m}$ and a Young's modulus of $E^B = 50 \text{ N/m}^2$. The solid continuum uses $E^S = 10 \text{ N/m}^2$ and $\nu = 0.3$ as material parameters. The solid discretization consists of $15 \times 3 \times 3$ elements, respectively (see Fig. 4(b)). The displacements of the solid at the left side of the cantilever beam is fully constrained, while on the right side a surface load p is applied to the structure (see Fig. 4(a)) resulting in a bending deformation. The beam is constrained by the Lagrange multiplier field, attaching its displacement to the one from the surrounding solid.

In a first step, we briefly discuss a practical choice for W , using the ideal augmented Lagrangian preconditioner, and present a numerical study that shows the influence of ϵ on the eigenvalues of the system matrix $\mathcal{A}_\epsilon$ and the preconditioned operator $\mathcal{P}_\epsilon^{-1} \mathcal{A}_\epsilon$. Based on the asymptotic limit of (15), we consider $\hat{S}_\epsilon = -\frac{1}{\epsilon} W$ as Schur complement approximation. Hereby we specifically choose $W = \kappa$, which is defined for each Lagrange multiplier node j as given in [6], yielding

$$\kappa_{jj} := \int_{\Lambda} \phi_j \, ds \, I, \tag{21}$$

with an identity matrix $I \in \mathbb{R}^{3 \times 3}$. Following up on the analysis presented in Section 4.1 and replacing S_ϵ by $\hat{S}_\epsilon$ leaves Q_{13} and Q_{23} unchanged, whereas the multiplier block becomes

$$Q_{33} = \hat{S}_\epsilon^{-1} S_\epsilon = -\epsilon W^{-1} (S^{-1} - \epsilon W^{-1})^{-1} = (I - \frac{1}{\epsilon} S^{-1} W)^{-1}$$




Fig. 5. Comparison of the eigenvalue spectrum of the system matrix $\mathcal{A}_\epsilon$ (top) and the preconditioned operator $\mathcal{P}_\epsilon^{-1} \mathcal{A}_\epsilon$ (bottom) using the ideal augmented Lagrangian preconditioner for the mixed-dimensional beam-solid model problem for different values of ϵ . We note that the horizontal axes are scaled differently for visualization purposes.

resulting in

$$\lim_{\epsilon \rightarrow 0} \|Q_{33}\| = 0, \quad \lim_{\epsilon \rightarrow \infty} \|Q_{33}\| = 1.$$

Despite using a simple approximation for the Schur complement, we are able to recover the properties of the exact Schur complement in the asymptotic limit $\epsilon \rightarrow \infty$ such that $\mathcal{P}_\epsilon^{-1} \mathcal{A}_\epsilon \rightarrow I$. Fig. 5 shows the eigenvalue spectrum for different values of the penalty parameter for the system matrix $\mathcal{A}_\epsilon$ in the top row and $\mathcal{P}_\epsilon^{-1} \mathcal{A}_\epsilon$ in the bottom row. As expected, the maximum eigenvalues of $\mathcal{A}_\epsilon$ increase and are not bounded with increasing values of ϵ . In contrast, for $\mathcal{P}_\epsilon^{-1} \mathcal{A}_\epsilon$ all eigenvalues cluster around one with $\epsilon \rightarrow \infty$, as $\hat{S}_\epsilon$ is becoming spectrally closer to the true S_ϵ . With the observations made in [6,31] and our brief numerical investigation, we consider $W = \kappa$ to be a proper choice [65].

The choice of the penalty parameter ϵ (respectively ϵ^S and ϵ^B) and the respective Schur complement approximation play a central role in the overall performance of the preconditioners $\mathcal{P}_{\epsilon, \text{mod}}$ and $\mathcal{P}_{\epsilon^S, \epsilon^B, \text{mod}}$. As stated in [63] and further investigated in [20], ϵ influences the modified augmented Lagrangian preconditioner differently as (11). The penalty parameter now has to be chosen carefully, usually not being located at the asymptotic limits $\epsilon \rightarrow 0$ and $\epsilon \rightarrow \infty$. Additionally, for $\mathcal{P}_{\epsilon^S, \epsilon^B, \text{mod}}$ an appropriate choice of ϵ^S and ϵ^B is essential to obtain an efficient method [30]. In the following, we consider three types of Schur complement approximations, with variants I and II being related to $\mathcal{P}_{\epsilon, \text{mod}}$ and variant III to $\mathcal{P}_{\epsilon^S, \epsilon^B, \text{mod}}$.

Variant I. We consider the first case for the special case of short, independent beams inside the solid continuum, such that we choose $\epsilon \ll E^B$, which is rather untypical for augmented Lagrangian preconditioners, yet not uncommon for the augmented Lagrangian formulation itself. One can see the augmentation as a pure stabilization in this case. Due to ϵ being small, we assume

$$K_\epsilon^S \approx K^S, \quad K_\epsilon^B \approx K^B, \quad B_\epsilon = B_\epsilon^T \approx 0 \quad \text{and thus} \quad S_\epsilon \approx S$$

to hold true. Therefore we assume the Schur complement to be given by

$$S_\epsilon = -M (K_\epsilon^S)^{-1} M^T - D (K_\epsilon^B)^{-1} D^T,$$

being very close to the original Schur complement with a small perturbation, while neglecting the additional coupling contributions given by B_ϵ and B_ϵ^T . With the augmented beam sub-matrix being described by a block-diagonal sparsity pattern, its inverse can be cheaply approximated by a sparse approximate inverse (SPAI) [66]. A common choice for the solid inverse is a SIMPLE-like approximation [25] or a block-diagonal variant [26], resulting in the respective approximated Schur complement given by

$$\hat{S}_\epsilon = -M \text{diag}(K_\epsilon^S)^{-1} M^T - D \text{SPAI}(K_\epsilon^B) D^T. \quad (22)$$

Variant II. For the second case we seek to find a value for ϵ such that

$$\hat{S}_\epsilon = -\frac{1}{\epsilon} W$$




Fig. 6. Comparison of the eigenvalue spectrum of preconditioned operator $\mathcal{P}_{\epsilon, \text{mod}}^{-1} \mathcal{A}_\epsilon$ with Schur complement variant I (top) and the preconditioned operator $\mathcal{P}_{\epsilon, \text{mod}}^{-1} \mathcal{A}_\epsilon$ with Schur complement variant II (bottom) for the mixed-dimensional beam-solid model problem for different values of ϵ .

holds true. By setting $W = \kappa$ as highlighted before, the augmented Schur complement follows the approximation given by

$$\hat{S}_\epsilon = -\frac{1}{\epsilon} \kappa. \quad (23)$$

For the modified augmented Lagrangian preconditioner with variant II, the multiplier block also satisfies

$$\lim_{\epsilon \rightarrow 0} \|Q_{33}\| = 0, \quad \lim_{\epsilon \rightarrow \infty} \|Q_{33}\| = 1.$$

Thus, as $\epsilon \rightarrow 0$, the limiting preconditioned operator becomes singular because $Q_{33} \rightarrow 0$. Although $Q_{33} \rightarrow I$ as $\epsilon \rightarrow \infty$, the full preconditioned operator does not approach the identity, as Q_{12} and Q_{32} approach nonzero limits. In fact, the limiting Q_{22} has eigenvalues, which are not equal to one, hence, the overall eigenvalue spectrum does not cluster around one. To underline our statements, we conduct a small eigenvalue analysis similar to the one presented in Section 4.1 with the resulting eigenvalue spectra given in Fig. 6 for varying penalty parameter values. Considering Schur complement variant I in the top row, the eigenvalues of the preconditioned linear operator cluster around one and move towards zero for increasing penalty parameter values. As the influence of B_ϵ gets stronger with bigger ϵ , the Schur complement approximation $\hat{S}_\epsilon$ given in (22) moves away from the actual S_ϵ , which supports our assumption. For variant II, the eigenvalues are clustered around zero for small penalty parameter values. With an increasing ϵ the eigenvalue spectrum is bounded away from zero, yet does not cluster around one. Increasing the penalty parameter even further worsens the eigenvalue spectrum again, with eigenvalues moving back to zero. This underscores that the removal of B_ϵ^T violates the properties shown in the second row of Fig. 5. We expect that there exists an optimal value for ϵ , such that the eigenvalues of $\mathcal{P}_{\epsilon, \text{mod}}^{-1} \mathcal{A}_\epsilon$ cluster around one. We numerically show this behavior in Section 5.1.

We now turn to the Schur complement approximation for $\mathcal{P}_{\epsilon^S, \epsilon^B, \text{mod}}$. Due to the introduction of two independent penalty parameters and similar to [30], we are able to scale the system (19), such that we can obtain a favorable solution. Since $B_{\epsilon^S}^T$ is omitted, the asymptotic analysis motivates choosing ϵ^S sufficiently small so that the neglected coupling contribution remains weak. This also has the nice side effect, that $K_\epsilon^S \approx K^S$ for small ϵ and thus standard AMG works very effectively. In contrast, we need to choose an appropriate ϵ^B , such that the influence of $K_{\epsilon^B}^B$ is highlighted and the eigenvalues are shifted towards one. While this leads to a severe ill-conditioning of the beam sub-problem, its influence on the actual preconditioner implementation is, in our case, negligible as we solve the beam sub-problem with a direct method.

Variant III. For the third case, we set

$$\hat{S}_{\epsilon^S, \epsilon^B} = -\frac{1}{\epsilon^B} W$$

and choose to use ϵ^B as the overall problem is dominated by the beam penalty contribution. By setting $W = \kappa$ as highlighted before, the augmented Schur complement follows the approximation given by

$$\hat{S}_{\epsilon^S, \epsilon^B} = -\frac{1}{\epsilon^B} \kappa. \quad (24)$$




Fig. 7. Comparison of the eigenvalue spectrum of preconditioned operator $\mathcal{P}_{\epsilon^S, \epsilon^B, \text{mod}}^{-1} \mathcal{A}_{\epsilon^S, \epsilon^B}$ with Schur complement variant III for the mixed-dimensional beam-solid model problem for different values of ϵ^S and ϵ^B .

Considering the respective Schur complement approximation for the individually scaled case we obtain the following asymptotic limits for changing ϵ^S and fixed ϵ^B :

$$\begin{aligned} \lim_{\epsilon^S \rightarrow 0} \|Q_{32}\| &= 0, & \lim_{\epsilon^S \rightarrow \infty} \|Q_{32}\| &= \left\| \epsilon^B W^{-1} D(K_{\epsilon^B}^B)^{-1} K^B \right\|, \\ \lim_{\epsilon^S \rightarrow 0} \|Q_{33}\| &= \left\| \epsilon^B (W + \epsilon^B D(K^B)^{-1} D^T)^{-1} (M(K^S)^{-1} M^T + D(K^B)^{-1} D^T) \right\|, & \lim_{\epsilon^S \rightarrow \infty} \|Q_{33}\| &= \left\| \epsilon^B W^{-1} D(K_{\epsilon^B}^B)^{-1} D^T \right\|. \end{aligned}$$

Using $\epsilon^S \rightarrow 0$ is favorable as Q_{32} vanishes and we are able to recover a block upper-triangular operator. In contrast to the exact Schur complement, Q_{33} does not approach the identity, but a more complicated limit depending on ϵ^B . Thus, we finally consider the special case of $\epsilon^S \rightarrow 0$ under varying ϵ^B resulting in

$$\lim_{\substack{\epsilon^S \rightarrow 0 \\ \epsilon^B \rightarrow 0}} \|Q_{33}\| = 0, \quad \lim_{\substack{\epsilon^S \rightarrow 0 \\ \epsilon^B \rightarrow \infty}} \|Q_{33}\| = \left\| I + (D(K^B)^{-1} D^T)^{-1} M(K^S)^{-1} M^T \right\|.$$

Similar to the cases before, there exists an optimal value for ϵ^B away from the asymptotic limits.

As shown in Fig. 7, the eigenvalues nicely cluster around one for $\epsilon^S \rightarrow 0$ and growing ϵ^B . Yet we remark that there are a few outliers, which we do not show explicitly. For an increasing beam penalty parameter, we expect an over-penalization of the system, thus shifting eigenvalues above one. This assumption is backed by the asymptotic limit given by $\epsilon^B \rightarrow \infty$. We study this effect for different combinations of ϵ^S and ϵ^B in more detail in Section 5.1.

While there exist occasions, where it makes sense to use an approximation of the full augmented Schur complement (15) [67], we omit this variant in the scope of this publication, but mention it here for the sake of completeness.

4.3. Comparison to existing methods in literature

Having discussed all details of the proposed preconditioner, we now position it relative to representative block preconditioners from the literature. In particular, we compare the presented modified augmented Lagrangian methods with operator preconditioning methods for saddle-point systems [27–29,68] and with penalty-based block factorization strategies [13,44,52].

In contrast to our approach, the preconditioners from [27–29,68] are tailored to saddle-point problems stemming from the mixed-dimensional coupling of scalar transport equations and are based on the concept of operator preconditioning, resulting in block diagonal preconditioners. The arising Schur complement is approximated by a spectrally matching fractional Laplacian, which constitutes the main difference. This idea is not directly transferable to the present vector-valued coupling of 3D elasticity and 1D torsion-free Kirchhoff–Love beams without a new operator preconditioning analysis for the beam block and the mixed-dimensional coupling operators. We therefore regard these methods as relevant reference points, but not as out-of-the-box competitors for the present problem class. Still, the application of operator preconditioning to the present problem type is an attractive topic for future research.

For the penalty-based case, the most relevant method is the approximate block factorization preconditioner developed in [13], which addresses exactly the same beam-solid coupling setting. In contrast to the present augmented Lagrangian formulation, that approach solves the penalty-regularized system (8) and relies on an approximate factorization of the resulting 2×2 block structure. It therefore provides the closest predecessor and is considered as the primary baseline for assessing the benefit of the modified augmented Lagrangian preconditioners later in Section 5.4. The multigrid method proposed in [52] is also designed for penalty-regularized coupled problems, however, its application to domains with different degrees of freedom and physics has not yet been explored. Likewise, the exact block preconditioning approach presented in [44] has not been investigated in the given setting too.

5. Numerical examples

To evaluate the proposed block preconditioners with the corresponding Schur complement variants, we carry out a set of experiments designed to quantify the effect of the modeling parameters, assess parallel scalability and demonstrate the relevance




Fig. 8. Visualization of the numerical test setup for different beam-to-solid volume ratios for a constant beam cross-section radius.

of the proposed methods for engineering applications. In a first step we examine the influence of the penalty parameter on the overall convergence behavior of the preconditioned linear solver based on the discussion presented in Section 4. The second numerical experiment is designed to investigate the dependency of the preconditioners on varying model parameters as introduced in Section 3.1. Afterwards, we assess weak and strong scalability of the proposed methods. In a last step we show the applicability of the block preconditioners to an engineering application consisting of a fiber-reinforced hybrid composite and compare the newly proposed approaches to the commonly used incomplete LU (ILU) factorization preconditioner and the penalty-based approximate block factorization preconditioner from [13].

The general numerical setup for the first three experiments in Sections 5.1–5.3 is inspired by the test cases given in [46,47] and the ones introduced in [13]. It mimics a generic short fiber-reinforced material modeled by a representative volume element (RVE). The domain of the solid continuum is given by a unit cube with fixed volume $V^S = 1 \text{ m}^3$, which is filled with randomly oriented fibers of equal length $l = 0.25 \text{ m}$. An illustration of the geometry for different beam-to-solid volume ratios is given in Fig. 8. The solid body is clamped at the bottom face setting the displacement in all directions to $u^S = 0$, while the top surface is displaced in vertical direction by 0.01 m , which is equal to 1% of the cube's edge length. In the following, we consider the solid mesh size h^S , the beam mesh size h^B , the beam Young's modulus E^B and the beam cross-section radius R^B to be varying parameters. Therefore, also the stiffness contrast $\mathcal{E}$ and volume ratio $\mathcal{V}$ vary depending on the chosen values. Based on the findings in [6], we consider beam-to-solid mesh size ratios of $\mathcal{H} \geq 2.5$ to avoid an overconstraining by the discrete Lagrange multipliers. As outlined in Section 2, the solid is modeled as St.-Venant–Kirchhoff material with the Young's modulus being fixed to $E^S = 1 \text{ N/m}^2$ and the Poisson's ratio set to $\nu = 0.3$. For this numerical study, the fibers are represented by a single torsion-free Kirchhoff–Love beam finite element each. This is justified since we consider short, straight fibers, which are well modeled by one element only [6,13] due to the high approximation quality of the torsion-free Kirchhoff–Love beam formulation making use of cubic Hermite polynomials for the interpolation of the centerline displacements [39]. In practice, the formulation therefore allows even relatively long fibers to be approximated using only a few beam elements. We employ the discretization scheme considered above, using continuous linear Lagrangean finite elements for the solid and the Lagrange multiplier field, and cubic Hermite polynomials for the beams. All geometries and meshes are generated using BEAMMe [69].

We use TRILINOS [70,71] and 4C [72] for the simulation of the numerical experiments. We run both codes in a distributed-memory fashion building upon the Message Passing Interface (MPI) communication model. GMRES from BELOS [73] is used to solve the arising linear systems. The iterative solver is assumed to be converged, when the residual is reduced by a factor of 10^8 . The modified augmented Lagrangian block preconditioners are implemented using the block preconditioning package TEKO [74]. We apply the algebraic multigrid implementation from MUELU [75], specifically smoothed aggregation (SA-AMG) and plain aggregation (PA-AMG). We solve the coarsest level of the multigrid hierarchy with the distributed version of SUPERLU [76] as direct solver. The coarse level is reached if 6500 or fewer unknowns remain. On each level, a pre- and post-smoothing is applied with Chebychev polynomials of 3rd order. Coarse levels of the multigrid hierarchy are repartitioned by a coordinate-based approach [77] implemented in ZOLTAN2 to ensure a proper workload per MPI process. Direct sub-block solves are performed by AMESOS2 [73], again using the distributed version of SUPERLU. All simulations are run on CPU nodes of our in-house cluster (one CPU node features 2x Intel Skylake CPUs with 24 cores each).

5.1. Choice of the schur complement and the penalty parameter

The choice of the penalty parameter ϵ (respectively ϵ^S and ϵ^B for variant III of the Schur complement) has a major influence on the overall performance of the proposed block preconditioners, as it steers the approximation quality of the Schur complement variants. In a first study, we examine how the number of iterations of the linear solver changes due to a varying $\epsilon \in [10^{-4} \text{ N/m}^2, 10 \text{ N/m}^2]$ (additionally $\epsilon^S \in \{10^{-1} \text{ N/m}^2, 10^{-3} \text{ N/m}^2, 0.0 \text{ N/m}^2\}$ and $\epsilon^B \in [10^{-2} \text{ N/m}^2, 10^2 \text{ N/m}^2]$), whereas the penalty-regularized formulation from [6] proposes a fixed value $\epsilon \approx E^B$. We utilize the three Schur complement approximations from Section 4.2 given by variant I




Fig. 9. Iteration count of the linear solver based on Schur complement variant I (top) and variant II (bottom) for a changing penalty parameter. In addition the results are shown for different stiffness ratios and a varying beam-to-solid volume ratio of $\nu = 1\%$ (left), $\nu = 3\%$ (middle) and $\nu = 10\%$ (right).

(22), variant II (23) and variant III (24). The dependency of the number of linear iterations on the value of ϵ is shown in Fig. 9 for variants I and II, with the first row corresponding to variant I, while the second row represents variant II. The respective behavior for variant III is visualized in Fig. 10 for varying ϵ^S and ϵ^B . Experiments with a constant stiffness ratio $\mathcal{E} \in \{10, 100, 1000\}$ are distinguished by different colors. Each column corresponds to a different beam-to-solid volume ratio $\nu \in \{1\%, 3\%, 10\%\}$. We assume a constant mesh size ratio of $H = 7.5$, which turns out to be a good choice considering the results given in Section 5.2 and a constant beam cross-section radius of $R^B = 0.01$ m.

For Schur complement variant I, we introduced the assumption, that the respective approximation is valid for $\epsilon \ll E^B$ with its quality degrading for $\epsilon \rightarrow E^B$. This behavior can exactly be observed in the first row of Fig. 9 throughout all columns. For small values of ϵ , the iteration count stays low for different stiffness ratios and volume ratios, only increasing slightly for higher values of $\mathcal{E}$ and ν . Shifting the penalty parameter closer to E^B results in a degradation of the approximation quality of the Schur complement and therefore results in unbounded values for the number of iterations. We conclude that for sufficiently small values of ϵ far away from E^B , variant I is a valid and robust approximation for the Schur complement.

The observed behavior changes considerably for variant II. Similar to the observations made in [20], there exists a value for ϵ that minimizes the iteration count, which can be seen in the second row of Fig. 9. In addition, the Schur complement approximation appears to depend on the stiffness ratio, resulting in a potentially unbounded iteration count. For $\nu = 10\%$, only the characteristic behavior of $\mathcal{E} = 10$ is shown, as higher material contrasts exceed the allowed iteration count and are therefore omitted. The value of ϵ that minimizes the iteration count is also not identical for different stiffness ratios. The optimal penalty parameter for variant II should be chosen higher for larger values of $\mathcal{E}$. Increasing the stiffness contrast reduces the parameter range that leads to fast convergence, thereby complicating the identification of the optimal penalty parameter value. For application cases in our range of consideration, which feature a medium material contrast, a valid choice is $\epsilon \in [0.001E^B, 0.01E^B]$.

Regarding variant III, the behavior is again slightly different as illustrated in Fig. 10. Each diagram plots the number of linear iterations over a changing value of the beam penalty parameter for a unique combination of a fixed beam-to-solid volume ratio and solid penalty parameter. Again, there exists a parameter combination of ϵ^S and ϵ^B , which minimizes the iteration count and in contrast to variant II results in a limited dependency to the stiffness and volume ratio. For a low stiffness ratio, an increasing volume ratio affects the number of linear iterations taken only slightly if ϵ^S is chosen sufficiently small. This can be observed when comparing the iteration results of $\epsilon^S = 10^{-1}$ N/m² and $\epsilon^S = 0.0$ N/m² for a constant value of $\mathcal{E} = 10$. For the first parameter combination, an increasing volume ratio also leads to an increase in iterations, while for the second one the iteration count remains constant. For larger material contrasts, this behavior marginally degrades, with the choice of ϵ^B becoming more important, as the curves of linear iterations for a constant stiffness ratio now feature a distinct minimum. A general observation is, that for $\epsilon^S \rightarrow 0$ the number of linear iterations decreases steadily, while an appropriate choice of the beam penalty parameter reduces the influence on the modeling parameters for $\nu < 10\%$ for the given model problem.




Fig. 10. Iteration count of the linear solver based on Schur complement variant III for changing solid and beam penalty parameters. In addition the results are shown for different stiffness ratios and a varying beam-to-solid volume ratio of $\nu = 1\%$ (left), $\nu = 3\%$ (middle) and $\nu = 10\%$ (right).

The overall results for variant II are very similar to the findings from [20,24] showing the effect of ϵ related to problems from fluid and contact mechanics. While Schur complement variant I is robust w.r.t. different stiffness and volume ratios, the respective penalty parameter has to be chosen sufficiently small and far away from E^B . In addition, the calculation of the sparse approximate inverse of the beam matrix sub-block as well as two triple matrix products to form the Schur complement present a computational complexity, which should not be underestimated. In contrast, Schur complement variant II is represented by a simple diagonal and, thus, is computationally cheap. Yet, it lacks independence from the stiffness and volume ratio and needs a matching ϵ to be effective. Variant III combines the best of the two other Schur complement approaches, also being represented by a diagonal and able to partly regain independence to the modeling parameters and low iteration counts by steering ϵ^S and ϵ^B correctly. The presented results are mostly in line with the behavior observed in [30]. In the following numerical experiments, we further highlight all three variants discussing the need for parameter robustness or performance.

5.2. Robustness under varying model parameters

Now, we study the robustness of the block preconditioners w.r.t. model parameter changes and how these influence the number of iterations of the linear solver. Based on the initial results obtained in Section 5.1, in particular the challenges in identifying an optimal penalty parameter for Schur complement variant II we only use variant I with a constant penalty parameter $\epsilon = 0.01 \text{ N/m}^2$ and variant III with fixed $\epsilon^S = 0.0 \text{ N/m}^2$ and $\epsilon^B = 5.0 \text{ N/m}^2$. For the parameter study, we consider varying values of the beam-to-solid volume ratio $\nu \in \{1\%, 3\%, 10\%\}$, the stiffness ratio $\mathcal{E} \in \{10, 100, 1000\}$ and the beam cross-section radius $R^B = \{0.003 \text{ m}, 0.006 \text{ m}, 0.010 \text{ m}\}$. We also study changing values of the mesh size ratio $\mathcal{H}$ by maintaining a constant beam mesh size and varying the solid mesh size. Smaller solid element sizes are generated by uniform mesh refinement of the base configuration. The results are visualized in Figs. 11 and 12 for variants I and III, respectively, showing the number of linear iterations plotted over




Fig. 11. Visualization of the iteration count based on different solid mesh sizes for different stiffness ratios, a varying beam-to-solid volume ratio of $\nu = 1\%$ (top), $\nu = 3\%$ (middle) and $\nu = 10\%$ (bottom) and changing beam cross-section radius $R^B = 0.003$ m (left), $R^B = 0.006$ m (middle) and $R^B = 0.010$ m (right) for Schur complement variant I.

an increasing number of degrees of freedom (DOFs). Each row of the plot shows different values of the volume ratio with $\nu = 1\%$ (top), $\nu = 3\%$ (middle) and $\nu = 10\%$ (bottom). The columns are related to the beam cross-section radius with $R^B = 0.003$ m (left), $R^B = 0.006$ m (middle) and $R^B = 0.01$ m (right). As outlined in Section 3.1, all parameter combinations considered in the following remain within the intended validity regime of the mixed-dimensional model. The parameter variations are chosen to represent practically relevant changes in quantities such as the beam cross-section radius, beam volume ratio, and beam-to-solid stiffness ratio. Accordingly, the results are interpreted primarily with respect to the influence of these application-relevant parameters on solver performance. The cases combining the finest solid discretization with a beam cross-section radius of 0.010 m are close to the limit of the underlying modeling assumption but remain within its intended regime.

In the following, we consider the parameter combination $\nu = 1\%$ and $R^B = 0.01$ m in the upper right corner of Figs. 11 and 12 as baseline. For this case, both preconditioners show a limited dependency on the modeling parameters with decreasing solid mesh size, as the iteration count stays perfectly constant for variant I and marginally increases for variant III.

We first discuss the behavior of the block preconditioner using Schur complement variant I when the beam cross-section radius is reduced. Hereby, the block preconditioner loses some of its parameter independence showing a weak relation of the iteration count w.r.t. the stiffness ratio in case of a decreasing solid mesh size. In addition, the number of iterations reacts more sensitive to increasing stiffness ratios with more iterations necessary to find a solution to the linear problem. Reducing the cross-section radius while keeping ν constant effectively results in more fibers and, thus, more beam elements interacting with the solid continuum. It is observed in [27], that the cross-section radius of the one-dimensional domain is a relevant parameter for parameter robustness, specifically related to the discretization approach itself. In our case, R^B has a similar influence at first glance, but due to constant beam-to-solid volume ratios enforced for the robustness study the parameter is tightly connected to the amount of interaction partners present.

In a second step, we increase ν to 3%, which also shows a dependency of the iteration count to the stiffness ratio with decreasing solid mesh size. In addition, for large solid element mesh sizes the iteration count seems to decrease first, while increasing again for




Fig. 12. Visualization of the iteration count based on different solid mesh sizes for different stiffness ratios, a varying beam-to-solid volume ratio of $\nu = 1\%$ (top), $\nu = 3\%$ (middle) and $\nu = 10\%$ (bottom) and changing beam cross-section radius $R^B = 0.003\text{ m}$ (left), $R^B = 0.006\text{ m}$ (middle) and $R^B = 0.010\text{ m}$ (right) for Schur complement variant III.

decreasing values of h^S . This seems to be related to an overconstraining of the coupling conditions by the Lagrange multiplier field and is not discussed further in the scope of this publication, as we mainly deal with the construction of block preconditioners and do not specifically consider an in-depth analysis of the discretization and respective formulation. For the parameter combination $\nu = 10\%$ and $R = 0.003\text{ m}$, the numerical solver breaks down due to extremely populated matrix sub-blocks arising from the multitude of coupling contributions, which marks the limit of the coupling approach and Schur complement variant I itself. While showing near-independence in parameters for low values of ν and $R = 0.01\text{ m}$, the iteration count is showing an increasing dependency on mesh refinement for growing stiffness ratios. In addition, this effect is amplified for examples with small beam radii and large beam-to-solid volume ratios. In general, the iteration count increases modestly for a decreasing R^B and increasing ν .

Now we discuss the results for the block preconditioner using Schur complement variant III. Again, we consider a reduction of the beam cross-section radius first keeping ν fixed and thus looking at each row of Fig. 12 independently. Hereby, the preconditioner shows to be robust with a decreasing mesh size as the number of iterations as well as the overall iteration behavior remains similar. This holds true for all stiffness ratios considered. Only for the highest material contrast, an increase in linear iterations can be observed, which amplifies for increasing values of ν . Especially for $\nu = 10\%$ and $\varepsilon = 1000$, this effect is pronounced as the iteration count reduces drastically with smaller mesh size.

In a next step we consider fixed beam radii and vary the volume ratio, thus considering each column separately at a time. For an increasing volume ratio, the number of linear iterations for the lowest stiffness contrast $\varepsilon = 10$ remains constant for each beam cross-section radius. For $\varepsilon = 100$ the growth in iterations is still small, yet the highest material contrast $\varepsilon = 1000$ again shows to be sensitive to parameter changes. In contrast to Schur complement variant I, we are mostly able to maintain independence from the modeling parameters across different beam cross-section radii and volume ratios for low to medium material contrast values. For high stiffness ratios, the iteration count reduces with decreasing mesh size, which is the opposite behavior as for variant I, where the number of iterations increases.

Table 2

Number of DOFs n_{DOF} of each part of the discretization for different problem sizes used for the scaling study.

ID	n_{DOF}^S	n_{DOF}^B	n_{DOF}^A	n_{DOF}^{total}
1	89373	4584	2292	96249
2	346053	18336	9168	373557
3	1361613	73344	36672	1471629
4	5401533	293364	146682	5841579

The overall iteration count for both block preconditioners stays between 10 and 50 iterations for the given parameter setup, which is still in an acceptable range for our applications. It should also be highlighted that the preconditioners are still parameter robust for a variety of combinations.

5.3. Strong and weak scalability


In the following, we study the parallel efficiency of the proposed block preconditioners by performing a strong and weak scaling study. We conduct the numerical experiment with a fixed parameter configuration based on the findings from Sections 5.1 and 5.2. To offer enough computational work and a robust setup, we consider a beam-to-solid volume ratio of $\mathcal{V} = 3\%$, a stiffness ratio of $\mathcal{E} = 10$ and a mesh size ratio of $H = 7.5$, while the beam cross-section radius is set to $R^B = 0.01$ m. The respective penalty parameter for each Schur complement is given by $\epsilon = 0.01$ N/m² for variant I, $\epsilon = 0.15$ N/m² for variant II and $\epsilon^S = 0.0$ N/m², $\epsilon^B = 5.0$ N/m² for variant III. We adapt the cube length of the solid domain for each problem size such that the overall geometric ratios are preserved during weak scaling. The number of DOFs of each discretization for the different problem sizes is shown in Table 2. Each problem scale starts at an initial processor count P featuring around $n_{DOF}^{total}/P \approx 90000$ DOFs per processor, which is halved several times for the strong scaling setup, until we reach the strong scaling limit or the maximum number of processors $P = 512$ in our setup. Hereby, n_{DOF}^{total} describes the total problem size. All timings are obtained by averaging the respective numbers from 20 individual simulation runs.

As initial verification of our test setup, we check the iteration count of the linear solver for each problem scale. For variant I and II the number of iterations stays in between 30 and 38 for all different processor counts, while for variant III the values range from 15 to 20. The actual scaling results are given as timings and are summarized in Fig. 13(a) for Schur complement variant I, Fig. 13(b) for variant II and Fig. 13(c) for variant III. We additionally compute the standard deviation for each data point and plot it as error bars. Owing to the near-identical timings for the generation of one data point over 20 runs, the deviation is negligible and therefore the actual error bar is hard to visualize, yet we state that the standard deviation is at most 0.1229 for all data points. We differentiate between the setup, the solve and the combined setup + solve phase. Lines of the same color represent strong scaling for a constant problem size, with the bright gray line representing ideal strong scaling. The black dashed lines illustrate weak scaling over a constant problem size per MPI process.


Regarding strong scalability, we are able to achieve good scaling results for the solve and setup phase of the first three problem scales until performance degrades when reaching a load per MPI process of $n_{DOF}^{total}/P \approx 12000$. This holds true for all variants, whereas variant I shows slightly worse strong scaling. For bigger problem sizes and higher processor counts, this behavior is more pronounced. The problem with ID 4 (see Table 2) already hits the strong scaling limit at $n_{DOF}^{total}/P \approx 24000$, with both solve and setup timings rising again for smaller values of n_{DOF}^{total}/P . For variant II, the effect is not as dramatic considering the same values for n_{DOF}^{total}/P . While strong scaling also slows down for higher processor counts, the setup and solve times are still decreasing. This behavior also holds true for the block preconditioner using Schur complement variant III. The difference in the strong scaling efficiency comes from the different Schur complement approaches themselves. For Schur complement variant II and III, a simple inversion of a diagonal matrix is sufficient, while for variant I we employ a direct solver for preconditioning the respective sub-block. As we go up to $n_{DOF}^A = 146682$ for the Lagrange multiplier block (cf. Table 2), the compute time of the direct method dominates the whole setup phase due to limited exploitation of parallelism and higher order of computational complexity in comparison to e.g., AMG, especially for an increasing number of processors.

The results for weak scalability show a similar picture. For variant I, the setup phase is dominated by the factorization of the Lagrange multiplier sub-block, resulting in a reduced weak scalability to large problem sizes and processor counts. We achieve efficiencies up to 89% for the setup phase and acceptable 87% for the solve phase. The overall block preconditioner reaches a weak scaling efficiency of around 88%. For Schur complement variant II, the setup phase shows very good weak scalability up to 97% and the solve phase reaches 92%. The combined setup and solve phase achieves a weak scaling efficiency of around 95%. For variant III, we achieve similar solve efficiencies of approximately 93%, yet the efficiency of the setup phase degrades to 88%. As we set $\epsilon^S = 0.0$ N/m² for variant III, the number of non-zeros of the solid sub-matrix is greatly reduced as the penalty contribution is neglected. This results in less computational work per MPI process compared to variant II and results in a less favorable parallel scaling. The overall solve phase reaches a parallel efficiency of 89%.


We conclude that variant II shows the best strong and weak scalability. Variant III slightly outperforms variant I in terms of the reported parallel efficiencies. The overall timings suggest that the block preconditioner using Schur complement variant III enables the fastest solution process, being considerably faster than the other two for the given problem setup. We remark that the scalability results reported here are representative for the short-fiber regime with block diagonal beam sub-problems. For longer or more continuously connected fibers, where the beam block loses its block diagonal structure and the coupling matrices feature more non-zeros, the cost of the direct sub-solves and the overall setup phase may scale less favorably.



(a) Schur complement variant I scaling results for the short fiber composite RVE for different problem sizes.



(b) Schur complement variant II scaling results for the short fiber composite RVE for different problem sizes.



(c) Schur complement variant III scaling results for the short fiber composite RVE for different problem sizes.

Fig. 13. Weak and strong scaling results for Schur complement variant I (top), variant II (middle) and variant III (bottom). Ideal strong scaling is shown for reference, along with a line corresponding to the estimated weak scaling efficiency over the problem scales. Ideal weak scaling is a horizontal line.




Fig. 14. A hybrid composite plate with a prescribed deformation at its free end.

5.4. Application: Hybrid composite plate

To assess the performance of the presented block preconditioners under an application-oriented configuration, we consider their application to a hybrid composite plate (inspired by the example shown in [47]) with different layers of fiber orientations and distributions. The problem setup is shown in Fig. 14(a) with a visualization of the cross-section given in Fig. 14(c). We consider the problem setup to resemble a sandwich plate with the dimensions $2.0\text{ m} \times 4.0\text{ m} \times 0.25\text{ m}$. The top and bottom sheets of the structure consist of two continuous fiber layers each rotated by 45° and -45° , respectively. The individual fibers are discretized adaptively based on their length with short fibers being represented by only one beam element, while longer ones contain up to 11 elements. The beam cross-section radius is set to $R^B = 0.0075\text{ m}$ and the Young's modulus to $E^B = 210\text{ N/m}^2$. The inner part of the structure consists of a composite core with short fibers of constant length $l = 0.25\text{ m}$. Each fiber is discretized by one torsion-free Kirchhoff–Love beam element with $R^B = 0.005\text{ m}$ and $E^B = 70\text{ N/m}^2$. The solid part of the domain features the same material properties throughout the whole continuum: a St-Venant–Kirchhoff material with Young's modulus $E^S = 30\text{ N/m}^2$ and Poisson's ratio $\nu = 0.3$. It is discretized by first-order hexahedral finite elements. The left side of the structure is fully clamped, the right side is displaced by 0.05 m in negative e_3 -direction forcing the plate to bend (see Fig. 14(b)), while all other sides are subject to natural boundary conditions. We consider a hierarchy of three different meshes for the solid domain, which are created by uniform mesh refinement of the base configuration. The numbers of DOFs for the beam and the Lagrange multiplier field remain constant.

We study GMRES as linear solver with different preconditioners, namely a naive approach with a textbook ILU preconditioner, the approximate block factorization preconditioner from [13] as well as the proposed augmented Lagrangian block preconditioners with Schur complement variants I, II and III from Section 4. The factorization-based preconditioner reads

$$P_{\text{fact}} = \begin{pmatrix} \hat{S} & -B^T \\ K_\epsilon^B & I \end{pmatrix} \begin{pmatrix} I \\ -(K_\epsilon^B)^{-1} B_\epsilon \end{pmatrix}$$

with $\hat{S} = K_\epsilon^S - B_\epsilon \text{SPAI}(K_\epsilon^B) B_\epsilon^T$. While the block factorization preconditioner is not directly depending on a specific choice of the penalty parameter, we use $\epsilon = 210.0\text{ N/m}^2$ to fulfill the constraints of the penalty-regularized system (8) properly. For variant I, the penalty parameter is set to $\epsilon = 0.01\text{ N/m}^2$, while we use $\epsilon = 0.2\text{ N/m}^2$ for variant II. As the interaction of the continuous fiber layers with the surrounding solid is non-local, the resulting coupling matrices D and M as well as the contributions due to the augmented Lagrangian formulation feature proportionally many non-zero entries. Instead of SA-AMG, we therefore employ PA-AMG to reduce the fill-in on coarse levels of the augmented solid sub-problem for both preconditioner variants. A more detailed insight to this effect and how it influences AMG is given in [44,78] for example, while a more general overview on the solution process of these types of systems is provided in [55]. For Schur complement variant III, we consider three different parameter combinations. Option IIIa considers $\epsilon^S = 10^{-3}\text{ N/m}^2$, while variant IIIb and IIIc set the penalty parameter related to the solid contribution to zero. Both make use of PA-AMG, while the last option IIIc uses SA-AMG as setting $\epsilon^S = 0\text{ N/m}^2$ avoids excessive fill-in due from penalty contributions on coarse multigrid levels. All different combinations of variant III use $\epsilon^B = 1.0\text{ N/m}^2$. For the beam and Schur complement sub-blocks, we use an LU factorization. In the case of a diagonal Schur complement approximation, we directly invert the matrix sub-block. An overview of all different parameter combinations for the modified Lagrangian preconditioner cases is provided in Table 3.

The linear solver is assumed to be converged, when the residual is reduced by a factor of 10^6 . The corresponding tolerance is lower than in the previous numerical examples, but this is justified for the application setting, where such a residual reduction is usually sufficient to obtain a good approximation of the linear system within the nonlinear solution process and, in addition, allows comparability to the application example results from [13]. The solution methods and their iteration counts and timings

Table 3

Block-LU and modified augmented Lagrangian preconditioner setup for the different Schur complement variants I, II and III based on the penalty parameter and numerical methods applied to approximate the solid, beam and Schur complement contributions, respectively.

Variant	ϵ [N/m ²]	$\hat{p}^S$	$\hat{p}^B$	$\hat{p}^A$
Block-LU	210.0	PA-AMG	LU	–
I	0.01	PA-AMG	LU	LU
II	0.2	PA-AMG	LU	–
	ϵ^S [N/m ²]	ϵ^B [N/m ²]		
IIIa	10 ^{−3}	1.0	PA-AMG	LU
IIIb	0.0	1.0	PA-AMG	LU
IIIc	0.0	1.0	SA-AMG	LU

Table 4

Comparison of averaged linear solver timings per nonlinear iteration for a hybrid composite sandwich plate (Block-LU – a GMRES solver using an approximate block factorization preconditioner, Variant I, II and III — a GMRES solver using the block preconditioners proposed in Section 4 with the respective Schur complement variant). The best value in each category is highlighted in bold font.

n^{proc}	n_{DOF}^S	n_{DOF}^B	n_{DOF}^I	n_{DOF}^{total}	Method	#iter	CPU time [s]					
							T_{setup}^S	T_{setup}^B	T_{setup}^I	T_{solve}	T_{total}	
8	430353	43344	21672	495369	Block-LU	74	3.638	0.112	–	21.153	62.473	83.625
					Variant I	47	0.902	0.104	5.160	7.255	13.827	21.082
					Variant II	81	0.949	0.104	–	1.052	22.477	23.532
					Variant IIIa	79	0.779	0.115	–	0.912	19.735	20.647
					Variant IIIb	23	0.247	0.107	–	0.375	1.5631	1.9383
					Variant IIIc	20	0.859	0.111	–	0.977	1.5017	2.4783
64	3255903	43344	21672	3320919	Block-LU	72	7.962	0.116	–	37.104	131.414	168.518
					Variant I	37	0.745	0.188	0.945	3.690	13.967	17.662
					Variant II	82	0.781	0.198	–	0.967	19.829	20.796
					Variant IIIa	75	0.770	0.189	–	0.954	11.505	12.459
					Variant IIIb	37	0.550	0.184	–	0.745	4.391	5.136
					Variant IIIc	21	3.037	0.182	–	3.2243	3.0494	6.274
512	25308603	43344	21672	25373619	Block-LU	61	8.731	0.149	–	32.224	119.987	152.213
					Variant I	33	1.754	0.434	0.383	4.132	13.427	17.559
					Variant II	82	1.770	0.470	–	2.061	32.953	35.014
					Variant IIIa	55	1.579	0.388	–	1.944	12.750	14.694
					Variant IIIb	54	0.875	0.342	–	1.258	8.926	10.184
					Variant IIIc	23	3.397	0.314	–	3.943	4.201	8.144

are summarized in Table 4. The timings are broken down into the setup time for each individual sub-block, namely the multigrid method on the solid field T_{setup}^S , the direct method used for the beam contribution T_{setup}^B and the different variants for the Schur complement T_{setup}^A . In addition, we consider the overall setup time for the block preconditioner T_{setup} , the solve time of the linear method T_{solve} and finally the total time spent for solving the linear system T_{total} . All timings are obtained as averages of ten individual simulation runs. We also consider distributed sparse direct solvers, yet none of them is able to solve the linear system in feasible time, thus we do not report results for those methods.

We consider the naive preconditioning approach first, which does not result in convergence for any of the three meshes, emphasizing the importance of block preconditioning for this type of problem. In a next step, we consider the approximate block factorization preconditioner from [13] to serve as baseline for the following discussion. The preconditioner achieves iteration counts of ≈ 70 for the first two mesh configurations and 61 for the finest mesh. The proposed block preconditioner with Schur complement variant II performs slightly worse, showing a constant iteration count of ≈ 80 for all different mesh refinements. Still, the number of iterations taken to find a solution for both approaches is rather high compared to variant I, which shows a decrease in the iteration count for an increasing problem size, reducing from 47 to 33. This behavior is most likely connected to the ratio of solid and beam interaction pairs related to the overall number of solid elements, which determines the density of the Lagrange multiplier matrix sub-block. With an increasing number of solid elements and a constant number of beam elements, the sub-matrix is getting sparser, helping to find better approximations of the Schur complement. Variant IIIa shows a similar characteristic with the iteration count decreasing from 79 to 55. In contrast, the parameter combinations related to IIIb lead to an increase in iterations from 23 to 54. Only variant IIIc shows a low and constant number of iterations of ≈ 20 over all three mesh configurations. Overall, variant III outperforms variant II and the block-LU preconditioner, since setting $\epsilon^S \rightarrow 0$ eliminates the approximation error introduced by the dropped off-diagonal block B_{ϵ}^T , which in turn leads to lower iteration counts. Moreover, choosing $\epsilon^S = 0$ N/m² removes the additional augmentation of the solid block, making SA-AMG affordable for the respective sub-problem because excessive fill-in on coarse levels is avoided, thereby enabling proper scalability.

Further, we compare the timings of the individual preconditioners used for each of the sub-blocks T_{setup}^S , T_{setup}^B as well as T_{setup}^A . As we employ PA-AMG for the matrix block representing the solid stiffness, we can expect reasonable scalability to bigger problem sizes. While roughly double the time is necessary between the base problem and the finest refinement for the multigrid setup, the method still stays below 2 s, while the problem size increases by a factor of $\approx 60\times$. The setup cost of the block-LU preconditioner stands out being considerably higher compared to the other preconditioning variants. This relates to the use of the augmented solid Schur complement for the setup, which features substantially more non-zeros than the augmented solid sub-matrix considered in all other methods. The setup time of the SA-AMG method lies in between the benchmark preconditioner and the other variants, increasing about a factor of four from ≈ 0.9 s to ≈ 3.4 s between the smallest and the biggest problem. Due to the beam sub-block being represented by a block-diagonal sparsity pattern, the direct factorization works fast with a slight increase in the setup time mainly due to communication overhead due to the increased number of parallel processes. The setup time of the factorization for the Lagrange multiplier block drastically decreases from the first to the second mesh refinement step. As the amount of DOFs stays constant and the overall density of the matrix block decreases, reduced timings are to be expected. The timings between variant I, II and IIIa are similar as the individual preconditioners act on the same matrix sub-blocks. For IIb, one can expect a faster setup time for the PA-AMG method, as the missing penalty contribution results in a lower bandwidth and, thus, a reduced time. Another difference is the construction of the Schur complement block, which for variant II and III reduces to the inversion of a diagonal matrix and therefore we neglect time measurements for this case. Next, we consider the total setup time T_{setup} , with all variants showing a different behavior over all mesh refinements. Variant I results in an expensive construction procedure for the block preconditioner due to the more expensive Schur complement approximation requiring two triple matrix products and the SPAI calculation. While variant II shows a slightly increasing trend, variant I gives a decrease in timings being dominated by the factorization of the Lagrange multiplier block. The setup time of variant IIIa behaves very similar to variant II as both use the same methods and operate on the same sparsity patterns. The block preconditioner related to variant IIb achieves the overall lowest setup time, as the augmented Lagrangian terms of the solid contribution are omitted, which represent the major bottleneck for the preconditioner setup due to an increased bandwidth of the sparsity pattern. The setup cost of variant IIc is comparable to the one of variant I, as SA-AMG is more expensive to construct, with both being the overall most costly methods in regard to the setup of the modified augmented Lagrangian preconditioners presented. Still, all methods are able to beat the approximate block factorization preconditioner.

For the timings of the solve phase T_{solve} , the behavior switches, with variant I showing faster solve times due to a lower iteration count compared to variant II. While variant II shows a constant iteration count over all problem sizes, the solve time still increases. While the augmentation adds proportionally more non-zero values to the matrix for each refinement step, the matrix–vector products inside the Krylov method therefore also get more expensive and do not scale as efficiently. Variant IIIa performs as good and in some cases even better than variant I. Overall, the solve timings related to variant IIb and variant IIc are the lowest. Lastly, we discuss the total time spent in the linear solver. For the base problem, the total time T_{total} for variants I, II and IIIa is nearly identical, with IIb and IIc both outperforming the other preconditioners by $\approx 10\times$. This changes for the other meshes with variant I being $\approx 2\times$ faster than variant II for the biggest problem. Especially in the largest mesh refinement case, a clear ordering of the methods becomes visible with variant II performing the worst followed by variant I. The variants IIIa, IIb and IIc perform consistently better in exactly that order. The best method is $\approx 4\times$ faster than the worst. Again, all methods are able to beat the benchmark case.

We conclude, that most preconditioning methods show a bounded iteration count over all different problem sizes. While variant I is robust under mesh refinement, its implementation comes with more effort and complexity, which also shows up in the timings. Variant II is simple to construct, yet still robust to a decreasing mesh size with the high iteration count as major drawback. Variant III leaves more room to steer the Schur complement approximation, thus all of its parameter combinations outperform the other variants in terms of timings. We consider the block preconditioner using variant IIc to perform the best for this problem type as it shows nearly constant iteration counts over all three mesh configurations and also provides almost always the best time-to-solution. Compared with the block-LU preconditioner, the advantages of modified augmented Lagrangian preconditioning for the saddle-point problem are clear and highlight its importance for the presented application case. This result, however, was to be expected as the approximate block factorization preconditioner is based on the penalty-regularized system, and operates on matrices with much higher bandwidth and thus considerably more non-zeros.

6. Conclusion and outlook

In this work, we proposed and analyzed modified augmented Lagrangian block preconditioners for the mixed-dimensional beam-solid coupling of a three-dimensional solid with embedded one-dimensional torsion-free Kirchhoff–Love beams and constraint enforcement through Lagrange multipliers. Starting from the underlying variational formulation and its finite element discretization, we studied the structure of the resulting linear system, the presence of a pure Neumann sub-problem on the beam sub-block, and the role of physically relevant modeling parameters.

To overcome the given limitations, we reformulated the discrete coupled problem in an augmented Lagrangian fashion. This formulation preserves the exact enforcement of the coupling constraints while adding a mild regularization that also improves solvability in the case of singular sub-blocks. Based on this augmented system, we derived an ideal augmented Lagrangian block preconditioner and then constructed more practical, block-triangular “modified” variants that allow for separate and efficient treatment of the solid, beam and Schur complement blocks. We analyze the asymptotic limits of the preconditioned operator using the presented block preconditioners and discuss practical approximations of the appearing inverses. The solid sub-problem is handled by AMG, while the beam and Schur complement blocks can be treated by sparse direct methods or are inverted directly in case of diagonal matrices. Central to the effectiveness and performance of the approach is an appropriate Schur complement

approximation and the choice of the penalty parameter. We introduced three Schur complement variants: one that approximates the Schur complement of the original system using simple inverses of the solid and beam operators, a second one that exploits a scaling matrix, and a third one that introduces different penalty parameters for the solid and beam sub-problem to improve the approximation with that scaling matrix. This scaling matrix has been chosen as a mass-like matrix at the coupling interface of the Lagrange multipliers to retain spectral equivalence of the Schur complement approximation.

The numerical experiments suggest that the proposed modified augmented Lagrangian block preconditioner, together with suitable choices of the Schur complement approximation and penalty parameter, exhibits only a mild dependence on a broad range of modeling parameters relevant for short-fiber-reinforced materials. For this tailored approximation and given parameter combinations, the method shows iteration counts with a moderate mesh dependence over the admissible refinement range for moderate stiffness contrasts and beam-to-solid volume ratios. It maintains acceptable iteration numbers even in more challenging regimes characterized by small beam radii, large stiffness ratios, and high beam-to-solid volume ratios. At the same time, the results also indicate that small beam radii combined with high stiffness ratios remain extremely challenging for the proposed method. Moreover, the present study is restricted to short-fiber configurations, the behavior for long or continuous fibers is not yet investigated. The influence of the penalty parameter on convergence has been studied systematically for all three Schur complement variants, revealing the expected trade-off between performance and independence from model parameters. Overall, the numerical results indicate that suitable choices of the Schur complement approximation and penalty parameter can reduce the sensitivity of the iteration counts to simultaneous variations in the considered modeling parameter range.

The parallel scalability studies show that the proposed preconditioners are suitable for large-scale parallel simulations of short-fiber mixed-dimensional beam-solid coupling scenarios. For representative RVE problems with several million degrees of freedom, we observe favorable weak and strong scaling on distributed-memory architectures, with good efficiency of both the AMG-based solid sub-solver and the block preconditioners as a whole, especially for Schur complement approximation variant II and III. For longer or more continuously connected fibers, the beam and coupling blocks feature more non-zeros, so the setup cost and direct sub-solves are expected to scale less favorably, which is left for future work.

Finally, we studied an application-oriented test case where the proposed coupling and preconditioning strategy has been applied to a hybrid composite plate with layered fiber reinforcement and a short fiber-reinforced core. The problem features moderate stiffness contrasts and heterogeneous length scales, posing a challenging mixed-dimensional and multi-material setting. The results demonstrate that not only textbook preconditioners do not yield convergence, but that the newly proposed modified augmented Lagrangian block preconditioners enable convergence and remain robust and scalable in the given configuration for an appropriate choice of the Schur complement approximation and penalty parameter. Compared to a block factorization preconditioner based on the penalty-regularized system, the presented method is competitive regarding iteration counts and is overall considerably faster.

Further work will extend the proposed coupling scheme and block preconditioner to be usable with a Simo-Reissner beam formulation. One difficulty in this regard is the introduction of additional rotational DOFs in the beam domain, which enables the modeling of more complex interaction scenarios with initially curved fibers. Treating the coupling of both the solid and beam domain in a purely positional way without any additional modifications is not sufficient anymore, as the rotations are not constrained properly, thus leaving components of the rigid body modes unconstrained. Application areas of interest are mixed-dimensional beam-solid volume and surface coupling. Addressing these challenges will further broaden the applicability of the proposed methods and enable the efficient simulation of a wider range of fiber-reinforced structures and mixed-dimensional beam-solid interaction problems.

CRedit authorship contribution statement

Max Firmbach: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Ivo Steinbrecher:** Writing – review & editing, Software, Conceptualization. **Alexander Popp:** Writing – review & editing, Supervision, Funding acquisition. **Matthias Mayr:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Funding acquisition, Conceptualization.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used ChatGPT-5.5 Instant to fix writing errors and to improve the readability and language in some parts of the manuscript. After using these tools, the authors reviewed and edited the content and take full responsibility for the content of the published article.

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 work described in this contribution has been funded by the *Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)* within the project “Stable discretization methods and scalable solvers for embedded fiber/solid coupling” (project number 528397555) as well as *dtec.bw - Digitalization and Technology Research Center of the Bundeswehr* under the project “hpc.bw - Competence Platform for High Performance Computing”. dtec.bw is funded by the European Union – NextGenerationEU. All authors read and approved the final manuscript.

Data availability

Building blocks of the modified augmented Lagrangian preconditioner developed and applied in this study are openly available in **Trilinos** at <https://github.com/trilinos/Trilinos> [70,71]. The implementation of the mixed-dimensional beam-solid coupling is available in **4C** at <https://github.com/4C-multiphysics/4C> [72]. All other data that support the findings of this study are available from the corresponding author upon reasonable request.

References

- [1] G. Holzapfel, Collagen in arterial walls: Biomechanical aspects, in: P. Fratzl (Ed.), *Collagen: Structure and Mechanics*, Springer, Boston, MA, 2008, pp. 285–324.
- [2] G.M. Cunniffe, F.J. O'Brien, Collagen scaffolds for orthopedic regenerative medicine, *Biomater. Regen. Med.* 63 (2011) 66–73.
- [3] P.-C. Aitcin, *High-Performance Concrete*, E & FN Spon, London, 1998.
- [4] E. Fehling, M. Schmidt, J.C. Walraven, T. Leutbecher, S. Fröhlich, *Ultra-High Performance Concrete UHPC*, Ernst & Sohn, Berlin, 2014.
- [5] A.E. Naaman, Half a century of progress leading to ultra-high performance fiber reinforced concrete: Part 1 – overall review, in: R.D. Toledo Filho, F.A. Silva, E.A.B. Koenders, E.M.R. Fairbairn (Eds.), *2nd International RILEM Conference on Strain Hardening Cementitious Composites (SHCC2-Rio)*, 2011, pp. 17–26.
- [6] I. Steinbrecher, M. Mayr, M.J. Grill, J. Kremheller, C. Meier, A. Popp, A mortar-type finite element approach for embedding 1D beams into 3D solid volumes, *Comput. Mech.* 66 (6) (2020) 1377–1398.
- [7] B.D. Agarwal, L.J. Broutman, K. Chandrashekhara, *Analysis and Performance of Fiber Composites*, Wiley, 2017.
- [8] D. Durville, Finite element simulation of textile materials at mesoscopic scale, in: *Finite Element Modelling of Textiles and Textile Composites*, Saint Petersburg, 2007.
- [9] B. Lé, G. Legrain, N. Moës, Mixed dimensional modeling of reinforced structures, *Finite Elem. Anal. Des.* 128 (2017) 1–18.
- [10] U. Khristenko, S. Schuß, M. Krüger, F. Schmidt, B. Wohlmuth, C. Hesch, Multidimensional coupling: A variationally consistent approach to fiber-reinforced materials, *Comput. Methods Appl. Mech. Engrg.* 382 (2021) 113869.
- [11] I. Steinbrecher, N. Hagmeyer, C. Meier, A. Popp, A consistent mixed-dimensional coupling approach for 1D Cosserat beams and 2D surfaces in 3D space, *Comput. Mech.* 76 (5) (2025) 1233–1260.
- [12] I. Steinbrecher, A. Popp, C. Meier, Consistent coupling of positions and rotations for embedding 1d Cosserat beams into 3D solid volumes, *Comput. Mech.* 69 (2022) 701–732.
- [13] M. Firmbach, I. Steinbrecher, A. Popp, M. Mayr, An approximate block factorization preconditioner for mixed-dimensional beam-solid interaction, *Comput. Methods Appl. Mech. Engrg.* 431 (2024) 117256.
- [14] N. Hagmeyer, M. Mayr, A. Popp, A fully coupled regularized mortar-type finite element approach for embedding one-dimensional fibers into three-dimensional fluid flow, *Internat. J. Numer. Methods Engrg.* 125 (8) (2024) e7435.
- [15] N. Hagmeyer, M. Mayr, I. Steinbrecher, A. Popp, One-way coupled fluid-beam interaction: Capturing the effect of embedded slender bodies on global fluid flow and vice versa, *Adv. Model. Simul. Eng. Sci.* 9 (2022) 9.
- [16] F. Lespagnol, C. Grandmont, P. Zunino, M.A. Fernández, A mixed-dimensional formulation for the simulation of slender structures immersed in an incompressible flow, *Comput. Methods Appl. Mech. Engrg.* 432 (2024) 117316.
- [17] M.F. Murphy, G.H. Golub, A.J. Wathen, A note on preconditioning for indefinite linear systems, *SIAM J. Sci. Comput.* 21 (6) (2000) 1969–1972.
- [18] M. Benzi, G.H. Golub, J. Liesen, Numerical solution of saddle point problems, *Acta Numer.* 14 (2005) 1–137.
- [19] M. Benzi, M.A. Olshanskii, An augmented Lagrangian-based approach to the Oseen problem, *SIAM J. Sci. Comput.* 28 (6) (2006) 2095–2113.
- [20] M. Benzi, Z. Wang, Analysis of augmented Lagrangian-based preconditioners for the steady incompressible Navier–Stokes equations, *SIAM J. Sci. Comput.* 33 (5) (2011) 2761–2784.
- [21] H.C. Elman, Preconditioning for the steady-state Navier–Stokes equations with low viscosity, *SIAM J. Sci. Comput.* 20 (4) (1999) 1299–1316.
- [22] P.E. Farrell, L. Mitchell, F. Wechsung, An augmented Lagrangian preconditioner for the 3D stationary incompressible Navier–Stokes equations at high Reynolds number, *SIAM J. Sci. Comput.* 41 (5) (2019) A3073–A3096.
- [23] M.F. Adams, Algebraic multigrid methods for constrained linear systems with applications to contact problems in solid mechanics, *Numer. Linear Algebra Appl.* 11 (2–3) (2004) 141–153, URL <https://onlinelibrary.wiley.com/doi/10.1002/nla.374>.
- [24] A. Franceschini, M. Ferronato, M. Frigo, C. Janna, A reverse augmented constraint preconditioner for Lagrange multiplier methods in contact mechanics, *Comput. Methods Appl. Mech. Engrg.* 392 (2022) 114632.
- [25] T.A. Wiesner, M. Mayr, A. Popp, M.W. Gee, W.A. Wall, Algebraic multigrid methods for saddle point systems arising from mortar contact formulations, *Internat. J. Numer. Methods Engrg.* 122 (15) (2021) 3749–3779.
- [26] A. Franceschini, N. Castelletto, M. Ferronato, Block preconditioning for fault/fracture mechanics saddle-point problems, *Comput. Methods Appl. Mech. Engrg.* 344 (2019) 376–401.
- [27] N. Dimola, M. Kuchta, K.-A. Mardal, P. Zunino, Robust preconditioning of mixed-dimensional PDEs on 3d-1d domains coupled with Lagrange multipliers, in: A. Linninger, K.-A. Mardal, P. Zunino (Eds.), *Quantitative Approaches To Microcirculation: Mathematical Models, Computational Methods and Data Analysis*, Springer Nature Switzerland, Cham, 2024, pp. 137–171.
- [28] M. Kuchta, K.-A. Mardal, M. Mortensen, Preconditioning trace coupled 3d-1d systems using fractional Laplacian, *Numer. Methods Partial Differential Equations* 35 (1) (2019) 375–393.
- [29] M. Kuchta, M. Nordaas, J.C.G. Verschaeve, M. Mortensen, K.-A. Mardal, Preconditioners for saddle point systems with trace constraints coupling 2D and 1D domains, *SIAM J. Sci. Comput.* 38 (6) (2016) B962–B987.
- [30] M. Benzi, M. Feder, L. Heltai, F. Mugnai, Augmented Lagrangian preconditioners for fictitious domain formulations of elliptic interface problems, 2026, <http://dx.doi.org/10.48550/arXiv.2603.12993>, Preprint.
- [31] M. Benzi, M. Feder, L. Heltai, F. Mugnai, Scalable augmented Lagrangian preconditioners for fictitious domain problems, *Comput. Methods Appl. Mech. Engrg.* 450 (2026) 118522.
- [32] A. Budiša, W.M. Boon, X. Hu, Mixed-dimensional auxiliary space preconditioners, *SIAM J. Sci. Comput.* 42 (5) (2020) A3367–A3396.
- [33] A. Budiša, X. Hu, Block preconditioners for mixed-dimensional discretization of flow in fractured porous media, *Comput. Geosci.* 25 (2) (2021) 671–686.
- [34] G. Kirchhoff, Ueber das Gleichgewicht und die Bewegung eines unendlich dünnen elastischen stabes, *J. Die Reine Angew. Math.* 56 (1859) 285–313.
- [35] A.E.H. Love, *A Treatise on the Mathematical Theory of Elasticity*, Dover, 1944.
- [36] E. Reissner, On one-dimensional finite-strain beam theory: The plane problem, *Z. Angew. Math. Phys. (ZAMP)* 23 (5) (1972) 795–804.
- [37] J.C. Simo, A finite strain beam formulation. The three-dimensional dynamic problem. Part I, *Comput. Methods Appl. Mech. Engrg.* 49 (1) (1985) 55–70.
- [38] C. Meier, A. Popp, W.A. Wall, A locking-free finite element formulation and reduced models for geometrically exact Kirchhoff rods, *Comput. Methods Appl. Mech. Engrg.* 290 (2015) 314–341.

- [39] C. Meier, A. Popp, W.A. Wall, Geometrically exact finite element formulations for slender beams: Kirchhoff–Love theory versus SimoReissner theory, *Arch. Comput. Methods Eng.* 26 (1) (2019) 163–243.
- [40] I.S. Steinbrecher, Mixed-Dimensional Finite Element Formulations for Beam-To-Solid Interaction (Ph.D. thesis), Universität der Bundeswehr München, 2022.
- [41] A. Popp, M.W. Gee, W.A. Wall, A finite deformation mortar contact formulation using a primal–dual active set strategy, *Internat. J. Numer. Methods Engrg.* 79 (11) (2009) 1354–1391.
- [42] L. Cattaneo, P. Zunino, Numerical investigation of convergence rates for the FEM approximation of 3D–1D coupled problems, in: A. Abdulle, S. Deparis, D. Kressner, F. Nobile, M. Picasso (Eds.), *Numerical Mathematics and Advanced Applications – ENUMATH 2013*, Springer International Publishing, Cham, 2015, pp. 727–734.
- [43] C. D’Angelo, Finite element approximation of elliptic problems with Dirac measure terms in weighted spaces: Applications to one- and three-dimensional coupled problems, *SIAM J. Numer. Anal.* 50 (1) (2012) 194–215.
- [44] D. Cerroni, F. Laurino, P. Zunino, Mathematical analysis, finite element approximation and numerical solvers for the interaction of 3D reservoirs with 1D wells, *Int. J. Geomath.* 10 (1) (2019) Article Number: 4.
- [45] M. Kuchta, F. Laurino, K.-A. Mardal, P. Zunino, Analysis and approximation of mixed-dimensional PDEs on 3D–1D domains coupled with Lagrange multipliers, *SIAM J. Numer. Anal.* 59 (1) (2021) 558–582.
- [46] C. Lauff, M. Krause, M. Schneider, T. Böhlke, On the influence of the fiber curvature on the stiffness of long fiber reinforced composites, *Internat. J. Numer. Methods Engrg.* 126 (15) (2025) e70094.
- [47] C. Lauff, M. Schneider, T. Böhlke, Microstructure generation of long fiber reinforced hybrid composites using the fused sequential addition and migration method, *J. Thermoplast. Compos. Mater.* 38 (8) (2025) 2855–2893.
- [48] M. Schneider, An algorithm for generating microstructures of fiber-reinforced composites with long fibers, *Internat. J. Numer. Methods Engrg.* 123 (24) (2022) 6197–6219.
- [49] Y. Saad, M.H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* 7 (3) (1986) 856–869.
- [50] M. Griebel, D. Oeltz, M.A. Schweitzer, An algebraic multigrid method for linear elasticity, *SIAM J. Sci. Comput.* 25 (2) (2003) 385–407.
- [51] P. Vaněk, J. Mandel, M. Brezina, Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems, *Computing* 56 (1996) 179–196.
- [52] A. Budiša, X. Hu, M. Kuchta, K.-A. Mardal, L. Zikatanov, Algebraic multigrid methods for metric-perturbed coupled problems, *SIAM J. Sci. Comput.* 46 (3) (2024) A1461–A1486.
- [53] P. Bochev, R.B. Lehoucq, On the finite element solution of the pure Neumann problem, *SIAM Rev.* 47 (1) (2005) 50–66.
- [54] M. Kuchta, K.-A. Mardal, M. Mortensen, On the singular Neumann problem in linear elasticity, *Numer. Linear Algebra Appl.* 26 (1) (2018) e2212.
- [55] M. Benzi, C. Faccio, Solving linear systems of the form $(A + \gamma U U^T)x = b$ by preconditioned iterative methods, *SIAM J. Sci. Comput.* 46 (2) (2024) S51–S70.
- [56] R. Baggio, A. Franceschini, N. Spiezia, C. Janna, Rigid body modes deflation of the preconditioned conjugate gradient in the solution of discretized structural problems, *Comput. Struct.* 185 (2017) 15–26.
- [57] M. Fortin, R. Glowinski, *Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-Value Problems*, North-Holland, Amsterdam, 1983.
- [58] M.R. Hestenes, Multiplier and gradient methods, *J. Optim. Theory Appl.* 4 (5) (1969) 303–320.
- [59] M.J.D. Powell, A method for nonlinear constraints in minimization problems, 1969.
- [60] E. Burman, P. Hansbo, M.G. Larson, The augmented Lagrangian method as a framework for stabilised methods in computational mechanics, *Arch. Comput. Methods Eng.* 30 (4) (2023) 2579–2604.
- [61] C. Bacuta, A unified approach for Uzawa algorithms, *SIAM J. Numer. Anal.* 44 (6) (2006) 2633–2649.
- [62] G.H. Golub, C. Greif, On solving block-structured indefinite linear systems, *SIAM J. Sci. Comput.* 24 (6) (2003) 2076–2092.
- [63] M. Benzi, M.A. Olshanskii, Z. Wang, Modified augmented Lagrangian preconditioners for the incompressible Navier–Stokes equations, *Internat. J. Numer. Methods Fluids* 66 (4) (2011) 486–508.
- [64] M. Benzi, Z. Wang, A parallel implementation of the modified augmented Lagrangian preconditioner for the incompressible Navier–Stokes equations, *Numer. Algorithms* 64 (1) (2013) 73–84.
- [65] A.J. Wathen, Realistic eigenvalue bounds for the Galerkin mass matrix, *IMA J. Numer. Anal.* 7 (4) (1987) 449–457.
- [66] M.J. Grote, T. Huckle, Parallel preconditioning with sparse approximate inverses, *SIAM J. Sci. Comput.* 18 (3) (1997) 838–853.
- [67] X. He, C. Vuik, C.M. Klaij, Combining the augmented Lagrangian preconditioner with the simple Schur complement approximation, *SIAM J. Sci. Comput.* 40 (3) (2018) A1362–A1385.
- [68] T. Baerland, M. Kuchta, K.-A. Mardal, Multigrid methods for discrete fractional Sobolev spaces, *SIAM J. Sci. Comput.* 41 (2) (2019) A948–A972.
- [69] BeamMe Authors, BeamMe – A general purpose 3D beam finite element input generator, 2025, <https://beamme-py.github.io/beamme>. (Accessed 13 November 2025).
- [70] M.A. Heroux, R.A. Bartlett, V.E. Howle, R.J. Hoekstra, J.J. Hu, T.G. Kolda, Lehoucq, K.R. Long, R.P. Pawłowski, E.T. Phipps, A.G. Salinger, H.K. Thornquist, R.S. Tuminaro, J.M. Willenbring, A. Williams, K.S. Stanley, An overview of the Trilinos project, *ACM Trans. Math. Software* 31 (3) (2005) 397–423.
- [71] M. Mayr, A. Heinlein, C.A. Glusa, S. Rajamanickam, M. Arnst, R.A. Bartlett, L. Berger-Vergiat, E.G. Boman, K.D. Devine, G. Harper, M.A. Heroux, M. Hoemmen, J.J. Hu, B. Kelley, D.P. Kouri, P. Kuberly, K. Kim, K. Liegois, C.C. Ober, R.P. Pawłowski, C. Pearson, M. Perego, E.T. Phipps, D. Ridzal, N.V. Roberts, C.M. Siefert, H.K. Thornquist, R. Tomasetti, C.R. Trott, R.S. Tuminaro, J.M. Willenbring, M. Wolf, I. Yamazaki, Trilinos: Enabling scientific computing across diverse hardware architectures at scale, *ACM Trans. Math. Software* 52 (2) (2026) 12.
- [72] 4C, 4C: A comprehensive multiphysics simulation framework, 2025, <https://www.4c-multiphysics.org>. (Accessed 13 November 2025).
- [73] E. Bavier, M. Hoemmen, S. Rajamanickam, H. Thornquist, Amesos2 and belos: Direct and iterative solvers for large sparse linear systems, *Sci. Program.* 20 (3) (2012) 243875.
- [74] E.C. Cyr, J.N. Shadid, R.S. Tuminaro, Teko: A block preconditioning capability with concrete example applications in Navier–Stokes and MHD, *SIAM J. Sci. Comput.* 38 (5) (2016) S307–S331.
- [75] L. Berger-Vergiat, C.A. Glusa, G. Harper, J.J. Hu, M. Mayr, A. Prokopenko, C.M. Siefert, R.S. Tuminaro, T.A. Wiesner, MueLu User’s Guide, Technical Report SAND2023-12265, Sandia National Laboratories, Albuquerque, NM (USA) 87185, 2023.
- [76] X.S. Li, J.W. Demmel, SuperLU-DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems, *ACM Trans. Math. Softw.* 29 (2) (2003) 110–140.
- [77] M. Deveci, S. Rajamanickam, K. Devine, U. Catalyurek, Multi-jagged: A scalable parallel spatial partitioning algorithm, *IEEE Trans. Parallel Distrib. Syst.* 27 (2015) 1.
- [78] S.J. Thomas, S. Ananthan, S. Yellapantula, J.J. Hu, M. Lawson, M.A. Sprague, A comparison of classical and aggregation-based algebraic multigrid preconditioners for high-fidelity simulation of wind turbine incompressible flows, *SIAM J. Sci. Comput.* 41 (5) (2019) S196–S219.