

SCALABLE SOLVERS FOR COMPLEX ELECTROMAGNETICS PROBLEMS

SANTIAGO BADIA^{†,‡}, ALBERTO F. MARTÍN^{†,‡}, AND MARC OLM[†]

ABSTRACT. In this work, we present scalable balancing domain decomposition by constraints methods for linear systems arising from arbitrary order edge finite element discretizations of multi-material and heterogeneous 3D problems. In order to enforce the continuity across subdomains of the method, we use a partition of the interface objects (edges and faces) into sub-objects determined by the variation of the physical coefficients of the problem. For multi-material problems, a constant coefficient condition is enough to define this sub-partition of the objects. For arbitrarily heterogeneous problems, a relaxed version of the method is defined, where we only require that the maximal contrast of the physical coefficient in each object is smaller than a predefined threshold. Besides, the addition of perturbation terms to the preconditioner is empirically shown to be effective in order to deal with the case where the two coefficients of the model problem jump simultaneously across the interface. The new method, in contrast to existing approaches for problems in curl-conforming spaces does not require spectral information whilst providing robustness with regard to coefficient jumps and heterogeneous materials. A detailed set of numerical experiments, which includes the application of the preconditioner to 3D realistic cases, shows excellent weak scalability properties of the implementation of the proposed algorithms.

Keywords: Finite Element Method, Maxwell Equations, Domain Decomposition, Electromagnetics, Solvers.

1. INTRODUCTION

Realistic simulations in electromagnetic problems often involve multiple materials (e.g., dielectric and conducting materials), which may imply high contrasts in the coefficients describing the physical properties of the different materials. Besides, the behaviour of conducting materials may be modelled by highly variable, heterogeneous coefficients. This problem definition inevitably leads to high condition numbers for the resulting linear systems arising from curl-conforming finite element (FE) discretizations of the corresponding partial differential equation (PDE), which pose great challenges for solvers. Furthermore, the design of solvers for $H(\text{curl})$ -conforming approximations poses additional difficulties, since the kernel of the curl operator is non-trivial. Consequently, for realistic electromagnetic simulations in 3D, the use of robust iterative solvers is imperative in terms of complexity and scalability. In this work, we will focus on the development of robust balancing domain decomposition by constraints (BDDC) preconditioners for problems posed in $H(\text{curl})$ involving high variation of the coefficients for the corresponding PDE.

Date: May 7, 2019.

[‡] Centre Internacional de Mètodes Numèrics en Enginyeria, Esteve Terrades 5, E-08860 Castelldefels, Spain. [†] Universitat Politècnica de Catalunya, Jordi Girona 1-3, Edifici C1, E-08034 Barcelona, Spain. SB gratefully acknowledges the support received from the Catalan Government through the ICREA Acadèmia Research Program. MO gratefully acknowledges the support received from the Catalan Government through the FI-AGAUR grant. This work has been partially funded by the project MTM2014-60713-P from the “Ministerio de Economía, industria y Competitividad” of Spain. The authors thankfully acknowledge the computer resources at Marenostrum-IV and the technical support provided by the Barcelona Supercomputing Center (RES-ActivityID: FI-2018-3-0029). E-mails: sbadia@cimne.upc.edu (SB), amartin@cimne.upc.edu (AM), and molm@cimne.upc.edu (MO).

BDDC preconditioners [1] belong to the family of non-overlapping domain decomposition (DD) methods [2]. They can be understood as an evolution of the earlier Balancing DD method [3]. These methods rely on the definition of a FE space with relaxed inter-element continuity, which is defined by choosing some quantities to be continuous across subdomain interfaces, i.e., the *coarse* or *primal* degrees of freedom (DOFs). Then, the continuity of the solution at the interface between subdomains is restored with an averaging operator. The method has two properties that make it an outstanding candidate for extreme scale computing, namely it allows for aggressive coarsening and computations among the different levels can be performed in parallel. Outstanding scalability results have been achieved by an implementation in the scientific computing software FEMPAR [4, 5], which exploits these two properties in up to almost half a million cores and two million subdomains (MPI tasks) [6]. Another work showing excellent scalability properties up to two hundred thousand cores is [7], which is implemented in the software project PETSc [8].

The main purpose of this work is to construct BDDC methods for the linear systems arising from arbitrary order edge (Nédélec) FE discretizations of heterogeneous electromagnetic problems. An analysis for 3D FETI-DP¹ algorithms with the lowest order Nédélec elements of the first kind was given by Toselli in [13], who argued that the difficulty of iterative substructuring methods for edge element approximations mainly lies in the strong coupling between the energy of subdomain faces and edges. In short, no efficient and robust iterative substructuring strategy is possible with the standard basis of shape functions for the edge FE (see [14]). A suitable change of basis was introduced in [13] for lowest order edge elements and box-subdomains. Besides, an extension to arbitrary order edge FEs and subdomain geometrical shapes is presented in [15]. In this work, we will offer some new insights in the definition and construction of the change of basis for the latter general case. As pointed out in [16], the change of variables can be implemented in practice with just a few simple modifications to the standard BDDC algorithm [1].

Modern BDDC methods [17] propose coarse space enrichment techniques that adapt to the variation of coefficients of the problem [18, 19, 20, 21, 22, 23], where coarse DOFs are adaptively selected by solving generalized eigenvalue problems. This approach is backed up by rigorous mathematical theory and has been numerically shown to be robust for general heterogeneous problems. On the other hand, several different scalings have been proposed for the averaging operator in the literature to improve the lack of robustness of the *cardinality* (i.e., arithmetic mean) scaling for coefficient jumps. The *stiffness*² scaling takes more information into account but can lead to poor preconditioner performance with mildly varying coefficients [24]. A robust approach is the *deluxe* scaling, first introduced in [16] for 3D problems in curl-conforming spaces. It is based on the solution of local auxiliary Dirichlet problems to compute efficient averaging operators [25, 17, 26, 7, 27, 28], involving dense matrices per subdomain vertex/edge/face. However, to solve eigenvalue and auxiliary problems is expensive and extra implementation effort is required as coarse spaces in DD methods are not naturally formulated as eigenfunctions.

The main motivation of this paper is to construct robust BDDC preconditioners for problems in curl-conforming spaces that keep the simplicity of the standard BDDC method, i.e., to avoid the spectral solvers of adaptive versions, whereas keeping robustness and low computational cost. In order to do so, we follow the idea of the physics-based BDDC (PB-BDDC) preconditioner, presented in [29] for problems in grad-conforming spaces.

Based on the fact that BDDC methods (and DD methods in general) are robust with regard to jumps in the material coefficients when these jumps are aligned with the partition [2, 30], one can use a physics-based (PB)-partition obtained by aggregating elements of

¹FETI-DP algorithms [9] are closely related to BDDC methods. In fact, it can be shown that the eigenvalues of the preconditioned operators associated with BDDC and FETI-DP are almost identical [10, 11, 12].

²weighted averages with the diagonal entries of the operator for every DOF

the same (or similar) coefficient value. However, using this type of partition can lead to a poor load balancing among subdomains and large interfaces. To overcome this situation, the PB-BDDC respects the original partition (well-balanced) but considers a sub-partition of every subdomain based on the physical coefficients, leading to a partition of the objects into sub-objects defined according to the variation of the coefficients. Consequently, the method is also based on an enrichment of the coarse space but with the great advantage of not requiring to solve eigenvalue or auxiliary problems, i.e., the simplicity of the original BDDC preconditioner is maintained. On the other hand, the PB-BDDC methods involve a richer interface with the application software, e.g., access to the physical properties of the problem, which is in line with the philosophy of the `FEMPAR` library, i.e., a tight interaction of discretization and linear solver steps to fully exploit the mathematical structure of the PDE operator. The PB-BDDC preconditioner turned out to be one order of magnitude faster than the BDDC method with deluxe scaling in [7] for linear elasticity and thermal conductivity problems with high contrast.

Our problem formulation arises from the time-domain quasi-static approximation to the Maxwell's Equations for the magnetic field (see [31]), which involves two different operators, the mass and double curl terms. This fact certainly poses more complexities than the ones faced in [29] for the PB-BDDC solver, since it has to deal with the interplay of both (simultaneous) coefficient jumps. Our solution is to propose a simple technique to recover the scenario where only one coefficient has a jump across interfaces: we will add a perturbation at the preconditioner level so that the perturbed formulation does not involve a jump for the mass-matrix terms across interfaces. The effectiveness of the technique will be empirically shown. In order to extend the PB-BDDC algorithm to heterogeneous materials, a relaxed definition of the PB-partition will be stated where we only require that the maximal contrast of the physical coefficient in each PB-subdomain is smaller than a predefined threshold. The threshold can be chosen so that the condition number is reasonably small while the size of the coarse problem is not *too large*.

The article outline is as follows. The problem is defined in Sect. 2, where basic definitions are introduced. Sect. 3 is devoted to the presentation of the PB-BDDC preconditioner for heterogeneous 3D problems in $H(\text{curl})$. In Sect. 4 we will give some implementation insights, based on our experience through the implementation of the algorithms in the scientific software project `FEMPAR`. In Sect. 5, we present a detailed set of numerical experiments, covering a wide range of cases and applications for the PB-BDDC preconditioner. Finally, some conclusions are drawn in Sect. 6.

2. PROBLEM SETTING

Let us consider the boundary value Maxwell problem on a physical domain $\Omega \subset \mathbb{R}^3$:

$$(1) \quad \begin{aligned} \nabla \times (\alpha \nabla \times \mathbf{u}) + \beta \mathbf{u} &= \mathbf{f} && \text{in } \Omega, \\ \mathbf{n} \times (\mathbf{u} \times \mathbf{n}) &= 0 && \text{on } \partial\Omega, \end{aligned}$$

where $\alpha \geq 0$, $\beta > 0$ are the resistivity and the magnetic permeability of the materials, respectively, \mathbf{n} is a unit normal to the boundary and $\nabla \times$ is the 3D curl operator (see [32]). For the sake of simplicity, we consider homogeneous Dirichlet conditions, i.e., zero tangential traces on $\partial\Omega$. Nevertheless, all the developments in this work can readily be applied to Neumann and/or inhomogeneous conditions, see [14] for proper definitions. In order to pose the weak form of the problem, let us define the functional space

$$H(\nabla \times; \Omega) \doteq \{\mathbf{v} \in L^2(\Omega)^3 : \nabla \times \mathbf{v} \in L^2(\Omega)^3\},$$

and its subspace that satisfies homogeneous Dirichlet boundary conditions,

$$H_0(\nabla \times; \Omega) \doteq \{\mathbf{v} \in H(\nabla \times; \Omega) : \mathbf{n} \times (\mathbf{u} \times \mathbf{n}) = 0 \text{ on } \partial\Omega\}.$$

Besides, we will also make use of the space

$$H^1(\Omega) \doteq \{v \in L^2(\Omega) : \nabla v \in L^2(\Omega)^3\}.$$

Functions in $H(\nabla \times; \Omega)$ are approximated by edge FE methods of arbitrary order, which we represent by $\widehat{X}_h \subset H(\nabla \times; \Omega)$. In addition, functions in $H^1(\Omega)$ are approximated by standard scalar, continuous Lagrangian FE methods, which we represent by $\widehat{V}_h \subset H^1(\Omega)$. The weak form of the boundary value Maxwell problem in Eq. (1) reads: find $\mathbf{u} \in H_0(\nabla \times; \Omega)$ such that

$$\widehat{\mathcal{A}}_h(\mathbf{u}, \mathbf{v}) = (\mathbf{f}, \mathbf{v}), \quad \forall \mathbf{v} \in H_0(\nabla \times; \Omega),$$

where

$$(2) \quad \widehat{\mathcal{A}}_h(\mathbf{u}, \mathbf{v}) = \int_{\Omega} [(\alpha \nabla \times \mathbf{u}) \cdot (\nabla \times \mathbf{v}) + \beta \mathbf{u} \cdot \mathbf{v}] dx, \quad (\mathbf{f}, \mathbf{v}) = \int_{\Omega} \mathbf{f} \cdot \mathbf{v} dx.$$

2.1. Domain partition. Let us consider a bounded polyhedral domain $\Omega \subset \mathbb{R}^3$. Let \mathcal{T}_h be a partition of Ω into a set of tetrahedral or hexahedral cells K . For every cell $K \in \mathcal{T}_h$ consider its set of vertices \mathcal{V}_K , edges \mathcal{S}_K , or faces \mathcal{F}_K . They constitute the set of geometrical entities of the cell (excluding itself) as $\mathcal{G}_K = \mathcal{V}_K \cup \mathcal{S}_K \cup \mathcal{F}_K$. The union of these sets for all cells is represented with $\mathcal{G} \doteq \cup_{K \in \mathcal{T}_h} \mathcal{G}_K$. We consider a partition Θ of the domain Ω into non-overlapping subdomains $\tilde{\Omega}_i$, $i = 1, \dots, \tilde{N}$ obtained by aggregation of elements $K \in \mathcal{T}_h$. These subdomains are assumed to be such that the computational cost of solving the discrete Maxwell problem in the different subdomains leads to a well-balanced distribution of computational loads among processors in memory distributed platforms. We denote by $\Gamma(\Theta)$ the interface of the partition Θ , i.e., $\Gamma(\Theta) = \cup_{\tilde{\Omega}_i \subset \Omega} \partial \tilde{\Omega}_i \setminus \partial \Omega$. Every subdomain $\tilde{\Omega}_i \subset \Omega$ can be also partitioned into the smallest set of subdomains Ω_{ij} , $j = 1, \dots, N_i$, such that the material properties (α, β) in Eq. (1) are constant at every Ω_{ij} . For obvious reasons, we call this sub-partition a PB-partition and will be denoted by Θ_{pb} . Clearly, the resulting global Θ_{pb} is also a partition of Ω , and there is a unique $\mathcal{D} \in \Theta$ for every $\mathcal{D}' \in \Theta_{\text{pb}}$ such that $\mathcal{D}' \subset \mathcal{D}$. We consider a global numbering for the PB-subdomains, i.e., Ω_k , $k = 1, \dots, N$, having a one-to-one mapping between the two indices labels. Analogously, we define the interface of the PB-partition as $\Gamma(\Theta_{\text{pb}}) = \cup_{\Omega_k \subset \Omega} \partial \Omega_k \setminus \partial \Omega$.

2.2. Finite Element spaces. Let us define the FE spaces $X_h^i \doteq \widehat{X}_h(\tilde{\Omega}_i) \subset H(\nabla \times; \tilde{\Omega}_i)$ for every subdomain $\mathcal{D} \in \Theta$, and the corresponding Cartesian product space $X_h = \prod_{i=1}^{\tilde{N}} X_h^i$. Note that functions belonging to this space are allowed to have discontinuous tangent traces across the interface $\Gamma(\Theta)$. The global space in which the global problem is sought, i.e., \widehat{X}_h , can be understood as the subspace of functions in X_h that have continuous tangent traces across $\Gamma(\Theta)$. We can now define the subdomain FE operator $\mathcal{A}_h^i : X_h^i \rightarrow X_h^{i'}$, $i = 1, \dots, \tilde{N}$, as $\mathcal{A}_h^i(\mathbf{u}_i, \mathbf{v}_i) = \int_{\tilde{\Omega}_i} [(\alpha \nabla \times \mathbf{u}_i) \cdot (\nabla \times \mathbf{v}_i) + \beta \mathbf{u}_i \cdot \mathbf{v}_i] dx$ for all $\mathbf{u}_i, \mathbf{v}_i \in X_h^i$. Then, the sub-assembled operator $\mathcal{A}_h : X_h \rightarrow X_h'$ is defined as $\mathcal{A}_h(\mathbf{u}, \mathbf{v}) = \prod_i^{\tilde{N}} \mathcal{A}_h^i(\mathbf{u}_i, \mathbf{v}_i)$, in which contributions between subdomains have not been assembled. The assembled operator $\widehat{\mathcal{A}}_h : \widehat{X}_h \rightarrow \widehat{X}_h'$ (see Eq. (2)) is the Galerkin projection of the operator \mathcal{A}_h onto \widehat{X}_h .

The space of edge FE functions can be represented as the range of an interpolation operator π^h , which is well-defined for sufficiently smooth functions $\mathbf{u} \in H(\nabla \times; \Omega)$, by

$$(3) \quad \pi^h(\mathbf{u}) \doteq \sum_a \sigma^a(\mathbf{u}) \boldsymbol{\varphi}^a$$

where $\sigma^a(\mathbf{u})$ are the evaluation of the moments, i.e., the DOF values, and $\boldsymbol{\varphi}^a$ are the elements of the unique basis of functions that satisfies $\sigma_a(\boldsymbol{\varphi}^b) = \delta_{ab}$, i.e., the shape functions.

The reader is referred to [14] for a comprehensive definition of edge moments and the construction of polynomial spaces and basis of shape functions for the tetrahedral/hexahedral edge FE of arbitrary order.

2.3. Objects. In this section we introduce the definitions of global objects, or simply *globs*, which are heavily used in DD preconditioners (see, e.g., [2]). Given a geometrical entity $s \subset \Gamma(\Theta)$ and a subdomain partition Θ , we denote by $\text{neigh}_\Theta(s)$ the set of subdomains in Θ that contain s . Then, we define a geometrical object as the maximal set λ of geometrical entities in $\Gamma(\Theta)$ with the same $\text{neigh}_\Theta(s)$ subdomain set. We denote by $\text{neigh}_\Theta(\lambda)$ the set of subdomains in Θ containing λ and by $\text{ndof}(\lambda)$ the total number of DOFs placed on $s \in \lambda$. An object λ such that $\text{ndof}(\lambda) > 0$ is a face F if $|\text{neigh}_\Theta(\lambda)| = 2$ or an edge E if $|\text{neigh}_\Theta(\lambda)| > 2$. In addition, an object such that $\text{ndof}(\lambda) = 0$ is a corner. Grouping together the objects of the same type, we obtain the set of corners Λ_C , edges Λ_E and the set of faces Λ_F . Therefore, the set of *globs* is defined as $\Lambda(\Theta) = \Lambda_C \cup \Lambda_E \cup \Lambda_F$.

Remark 2.1. *This definition differs from the standard one (see, e.g., [33]). It is intentionally done in order to isolate globs that do not contain DOFs, i.e., Λ_C , which can be omitted in the rest of our exposition.*

Once *globs* are defined, let us also introduce the set of PB-*globs*, denoted by $\Lambda_{\text{pb}}(\Theta)$, as classification of all $s \subset \Gamma(\Theta)$ into Λ_C , Λ_E or Λ_F by considering the previous definitions based on $\text{neigh}_{\Theta_{\text{pb}}}(s)$ rather than $\text{neigh}_\Theta(s)$. $\Lambda_{\text{pb}}(\Theta)$ is a sub-partition of $\Lambda(\Theta)$ where coefficients are subdomain-wise constant within each $\lambda \in \Lambda_{\text{pb}}(\Theta)$.

3. PHYSICS-BASED BDDC

3.1. Change of basis. Any BDDC method that employs a standard 3D edge FE basis of shape functions is bound to show a factor dependent on the element size h^{-2} in the condition number [13], which precludes scalability. A key aspect of the curl-conforming edge FE spaces is the fact that $\nabla V_h^i \subset X_h^i$. One of the main ingredients of any BDDC method are the averaging operators $\mathcal{W}_h : X_h \rightarrow \widehat{X}_h$ (see detailed exposition in Sect. 3.2) that restore the continuity of the solution at the interface among subdomains. Since the averaging operators are usually based on some algebraic operations over DOF values, they are, more precisely, scaling matrices that depend on the basis being used to describe X_h (and X_h^i , by restriction to every subdomain). A key property that must hold for such operator to end up with a stable decomposition is the following: Given a function $\mathbf{u} \in X_h$ such that its local component in every processor belongs to ∇V_h^i , the restriction of the resulting function $\mathcal{W}_h \mathbf{u} \in \widehat{X}_h$ to every subdomain must belong to ∇V_h^i too. Otherwise, the energy of such functions is much increased after the averaging operation, and thus, the decomposition is not *scalable*. A key result in this direction is the decomposition proposed in [13] in the framework of FETI-DP methods for problems in $H(\nabla \times; \Omega)$.

Edge FE space moments can be assigned to edges/faces of the mesh (see [14]). Let us denote by $X_h^i(I)$ the subspace of functions of X_h^i such that their DOF values are not *located* on some $E \in \Lambda_E$ or $F \in \Lambda_F$, i.e., the DOFs are interior. Clearly, $X_h^i = \{X_h^i(I)\} \oplus \{X_h^i(F)\}_{F \subset \partial \tilde{\Omega}_i} \oplus \{X_h^i(E)\}_{E \subset \partial \tilde{\Omega}_i}$. On the other hand, a function $\mathbf{v}^i \in X_h^i(E)$ for a coarse edge $E \subset \partial \tilde{\Omega}_i$ admits a unique decomposition as follows (see [13, 25] for more details):

$$(4) \quad \mathbf{v}^i \cdot \mathbf{t}_E = s_{0,E}(\mathbf{v}^i) \Phi_E \cdot \mathbf{t}_E + \sum_{j=1}^{n_E-1} w_{jE}(\mathbf{v}^i) \nabla \phi_E^j \cdot \mathbf{t}_E, \quad \forall E \in \Lambda_E,$$

with $\phi_E^j \in V_h^i(E)$ being the Lagrangian shape functions related to the internal nodes of E (i.e., nodes $\xi \in E$ such that $\xi \notin \partial E$) and n_E their cardinality, whereas $s_{0,E}(\mathbf{v}^i) \doteq \int_E \mathbf{v}^i \cdot \mathbf{t}_E ds$. It is clear from Eq. (4) that two kinds of DOFs arise in the new basis for each subdomain edge $E \in \Lambda_E$: a DOF associated with the basis function Φ_E , which represents

the average tangent value over the coarse edge E , and DOFs associated with gradients of scalar, Lagrangian shape functions placed at its internal nodes. An illustration for the variables in the old (original) and new basis for a given E is presented in Fig. 1. Thus, we have that $X_h^i(E) \doteq \nabla V_h^i(E) \oplus \Phi_E$. As a result, X_h^i admits the unique decomposition:

$$(5) \quad X_h^i = \{X_h^i(I)\} \oplus \{X_h^i(F)\}_{F \subset \partial\tilde{\Omega}_i} \oplus \{\nabla V_h^i(E)\}_{E \subset \partial\tilde{\Omega}_i} \oplus \{\Phi_E\}_{E \subset \partial\tilde{\Omega}_i},$$

where Φ_E is the tangential vector such that $\Phi_E \cdot \mathbf{t}_E = 1$, \mathbf{t}_E being the unit tangent to $E \in \Lambda_E$.

Let us now describe the relation between the original set of DOFs (old basis in the global space) and the one that arises from Eq. (5) (new basis) for \hat{X}_h . A function $\mathbf{u} \in \hat{X}_h$ can be written in the old basis as $\mathbf{u} = \sum_a u^a \boldsymbol{\varphi}^a$, where $\boldsymbol{\varphi} = \{\boldsymbol{\varphi}^1, \dots, \boldsymbol{\varphi}^n\}$ is the set of global edge shape functions. Furthermore, consider the set of new basis functions $\boldsymbol{\psi} = \{\boldsymbol{\psi}^1, \dots, \boldsymbol{\psi}^n\}$, where old basis elements $\boldsymbol{\varphi}^a$ associated to $E \in \Lambda_E$ are replaced by its corresponding functions in Eq. (4) (i.e., interior and face edge functions, Lagrangian shape functions gradients, and the coarse edge functions). The interpolation operator π^h (see Eq. (3)) induces the change of basis matrix, whose entries are computed by evaluating the original edge moments σ^a for the introduced set of new basis functions $\boldsymbol{\psi}$ as (Einstein notation)

$$u_{\text{old}}^a = \sigma^a(\mathbf{u}_{\text{new}}) = \sigma^a(\boldsymbol{\psi}^b) u_{\text{new}}^b = \hat{Q}_{ab} u_{\text{new}}^b,$$

or in compact form, $\mathbf{u}_{\text{old}} = \hat{Q} \mathbf{u}_{\text{new}}$. Furthermore, we can readily define the inverse change of basis as $\mathbf{u}_{\text{new}} = \hat{Q}^{-1} \mathbf{u}_{\text{old}}$. The usual restriction matrix $R_i : \hat{X}_h \rightarrow X_h^i$ is used to obtain local restrictions of the global change of basis as $Q_i = R_i \hat{Q} R_i^T$; we abuse notation, using the same notation for the restriction with respect to the two bases, since it will be clear from the context. Finally, local restrictions lead to the change of basis $Q = \prod_i Q_i$, which will be applied for functions defined on X_h . A detailed exposition of an implementation strategy for the change of basis is found in Sect. 4.2.

3.2. Preconditioner. Similarly to other BDDC methods, we associate coarse DOFs to some of the *globals* in $\Lambda_{\text{pb}}(\Theta)$. In particular, BDDC methods for 3D curl-conforming spaces associate two coarse DOFs to every $E \in \Lambda_E$, defined as

$$(6a) \quad s_{0,E}(\mathbf{v}^i) \doteq \int_E \mathbf{v}^i \cdot \mathbf{t}_E ds,$$

$$(6b) \quad s_{1,E}(\mathbf{v}^i) \doteq \int_E s \mathbf{v}^i \cdot \mathbf{t}_E ds,$$

where s is an arc-length parameter, $s \in [-|E|/2, |E|/2]$. Thus, the expression Eq. (6b) refers to the first order moment of the tangent component of the solution on the edge E , in contrast to the zero-order moment in Eq. (6a). In the new basis (see Eq. (4)), it is easy to check that $s_{0,E}(\mathbf{v}^i) = s_{0,E}(\Phi_E)$ and $s_{1,E}(\mathbf{v}^i) = s_{1,E}(\sum_{j=1}^{n_E-1} w_{jE} \nabla \phi_E^j)$ [13]. Let us define the subspace \tilde{X}_h as

$$\tilde{X}_h \doteq \{\mathbf{w} \in X_h : s_{0,E}^{\mathcal{D}} = s_{0,E}^{\mathcal{D}'}, s_{1,E}^{\mathcal{D}} = s_{1,E}^{\mathcal{D}'} \forall E \in \Lambda_E, \forall \mathcal{D}, \mathcal{D}' \in \text{neigh}_{\Theta_{\text{pb}}}(E)\}$$

i.e., the subspace $\tilde{X}_h \subset X_h$ such that for all $\mathbf{w} \in X_h$, coarse DOFs (6a) and (6b) are continuous across subdomain interfaces $\Gamma(\Theta)$ for all $E \in \Lambda_E$. Clearly, $\hat{X}_h \subset \tilde{X}_h \subset X_h$.

The following key ingredient in the BDDC method is the averaging operator $\mathcal{W}_h : X_h \rightarrow \hat{X}_h$, defined as some weighted average of the DOF values at the interface. This operator is in practice defined as a matrix for a particular choice of the basis functions for X_h^i . Let us consider the new basis functions in $\boldsymbol{\psi}$. Given a fine edge/face $f \subset \Gamma(\Theta)$, we define a

weight for each $\mathcal{D} \in \text{neigh}_\Theta(f)$ as

$$(7) \quad \delta_{\mathcal{D}}^\dagger(f) = \frac{\sum_{\mathcal{D}' \in \text{neigh}_{\Theta_{\text{pb}}}(f) \cap \mathcal{D}} \chi_{\mathcal{D}'}}{\sum_{\mathcal{D}' \in \text{neigh}_{\Theta_{\text{pb}}}(f)} \chi_{\mathcal{D}'}}$$

where the choice of χ defines the scaling: the *cardinality* scaling with $\chi = 1.0$ or the α -based, β -based scaling with $\chi = \alpha$ or $\chi = \beta$, respectively. Besides, one can consider a weighted coefficient for χ , $\omega = \alpha + \beta h^2$, which is also subdomain-wise constant within all $E \in \Lambda_E$ in our definitions if regular structured meshes are considered. We note that all the expressions for the scalings are constant on *glob*s by construction, due to objects generation based on Θ_{pb} with constant coefficients. Then, we define the weighted function $\mathcal{W}_h \mathbf{v} \in \widehat{X}_h$ as follows. First, we compute for every subdomain Ω_i the weighted local functions as

$$\mathbf{w}^i = \mathbf{v}_I^i + \sum_{F \subset \Lambda_F} \delta_{\mathcal{D}}^\dagger(F) \mathbf{v}_F^i + \sum_{E \subset \Lambda_E} \delta_{\mathcal{D}}^\dagger(E) \mathbf{v}_E^i,$$

where \mathbf{v}^I , \mathbf{v}_F^i , and \mathbf{v}_E^i include the components related to interior, face, and edge DOFs in Eq. (5), respectively, and $\delta_{\mathcal{D}}^\dagger$ the corresponding weight for every interface edge/face f in $\mathcal{D} \in \text{neigh}_\Theta(f) \cap \Omega_i$. Next, we sum the values of DOFs on different subdomains that represent the same DOFs in X_h , i.e., assemble the DOFs as

$$\mathbf{v} = \sum_i R_i^T \mathbf{w}^i.$$

Next, we recover the sub-assembled bilinear form \mathcal{A}_h , whereas $\widehat{\mathcal{A}}_h$ and $\widetilde{\mathcal{A}}_h$ are the Galerkin projection of \mathcal{A}_h onto \widehat{X}_h and \widetilde{X}_h , respectively. We additionally define the harmonic extension operator \mathcal{E} , that, given $\mathbf{u} \in \widehat{X}_h$, provides $\mathbf{u} + \delta \mathbf{u}_I$, where $\delta \mathbf{u}_I \in X_h^i(I)$ is a bubble function that vanishes on the interface $\Gamma(\Theta)$ and holds:

$$\langle \mathcal{A}_h^i \delta \mathbf{u}_I^i, \mathbf{v}_I^i \rangle = -\langle \mathcal{A}_h^i \mathbf{u}^i, \mathbf{v}_I^i \rangle, \quad \forall \mathbf{v}_I^i \in X_h^i(I).$$

Let us denote the Galerkin projection of $\widehat{\mathcal{A}}_h$ onto the global bubble space $\widehat{X}_h(I) \doteq \{\mathbf{v} \in \widehat{X}_h \text{ such that } \mathbf{v} = \mathbf{0} \text{ on } \Gamma(\Theta)\}$ by $\widehat{\mathcal{A}}_{h,0}$. Thus, the action of the harmonic extension operator can be written as

$$\mathcal{E} \mathbf{u} \doteq (1 - \widehat{\mathcal{A}}_{h,0}^{-1} \widehat{\mathcal{A}}_h) \mathbf{u}.$$

We finally define the operator $\mathcal{H} = \mathcal{E} \mathcal{W}_h$. We can now state the BDDC preconditioner as

$$(8) \quad \mathcal{P} \doteq \widehat{\mathcal{A}}_{h,0}^{-1} + \mathcal{H} (\widetilde{\mathcal{A}}_h)^{-1} \mathcal{H}^T.$$

Note that having the expression of the operators associated with the new basis is essential in order to apply the averaging operator \mathcal{W}_h . Nevertheless, it is possible to employ the original operators in the standard basis and work with the change of basis matrix $\widehat{\mathcal{Q}}$ [16]. In this case, the only difference with regard to Eq. (8) is the application of the averaging operator as

$$(9) \quad \mathcal{H} = \mathcal{E} \widehat{\mathcal{Q}} \mathcal{W}_h \mathcal{Q}^{-1} \quad \text{or} \quad \mathcal{H}^T = \mathcal{Q}^{-T} \mathcal{W}_h^T \widehat{\mathcal{Q}}^T \mathcal{E}^T.$$

Application details for the change of basis are exposed in Sect. 4.2. Therefore, the definition of the preconditioner is the one of the standard BDDC [1] with a set of *glob*s generated by a partition based on coefficients and a modification of the averaging operator to take into account this fact. Besides, one can work with the standard basis of edge FEs and use strategically the change of basis required to attain a scalable algorithm in the application of the weighting operator.

3.3. Perturbed PB-BDDC preconditioner. The presented PB-BDDC preconditioner has been shown to be robust with the jump of coefficients in the steady Poisson equation [29]. However, the problem in Eq. (1) adds the complexity of the interplay between the two different parameters α and β across the interface. Following the robust approach in [29], our idea is to get rid of the jump of one coefficient across the interface so the preconditioner has not to deal with the interplay between the two of them and the scenario where the method is successful is recovered. In order to decide which coefficient is affected, we consider the locality of the mass matrix operator in front of the double curl terms. The main idea is to add a perturbation in the original formulation of the preconditioner so we end up with common information for the mass matrix operator for DOFs that are replicated among different subdomains, i.e., located on the interface $\Gamma(\Theta)$. Therefore, the problem posed in \tilde{X}_h will only contain a jump in the double curl term across the interface.

Given a function $\mathbf{u}^i \in X_h^i$, we can define its extension as a global function $\tilde{\mathbf{u}}^i \in \hat{X}_h$ such that all DOFs belonging to $\tilde{\Omega}_i$ are identical to the ones of \mathbf{u}^i and the rest are zero. The extended function has support on $\tilde{\Omega}_i$ and its neighbours, denoted by $\tilde{\Omega}_i$. The perturbed preconditioner for a local subdomain $\tilde{\Omega}_i$ is expressed as:

$$\tilde{\mathcal{A}}_h^i(\mathbf{u}^i, \mathbf{v}^i) = \int_{\tilde{\Omega}_i} (\alpha \nabla \times \mathbf{u}^i) \cdot (\nabla \times \mathbf{v}^i) dx + \int_{\tilde{\Omega}_i} \beta \tilde{\mathbf{u}}^i \cdot \tilde{\mathbf{v}}^i dx.$$

Therefore, entries for interface DOFs in the local mass matrix will be *fully-assembled* instead of *partially* assembled, leading to common information at the interface across all subdomains. In the situation where no jump occurs for the mass matrix coefficients at the interface among subdomains, we consider the original preconditioner presented in Sect. 3.2, avoiding the perturbed formulation for obvious reasons.

Remark 3.1. *The definition of the original problem is not modified. We only consider the perturbed local operator $\tilde{\mathcal{A}}_h^i$ in the formulation for the preconditioner.*

3.4. Relaxed PB-BDDC. In previous sections, the definition of Θ_{pb} (and consequently the definition of PB-*globs*) is based on the requirement that coefficients are constant in each PB-subdomain, i.e., different subparts with constant coefficients can be identified in a subdomain, e.g. a problem composed by different homogeneous materials. However, physical coefficients may vary across a wide spectrum of values, even in a small spatial scale. Besides, the requirement that coefficients have to be constant in each PB-subdomain may result in an over-partitioned domain where coefficient jumps are not significant among different PB-subdomains. In order to address these situations and to deal with a more general applicability of the preconditioner, we introduce the relaxed PB-BDDC (rPB-BDDC) extension of the preconditioner. In short, relaxed PB (rPB)-subdomains are not determined by constant coefficients within the original partition but we only require that the maximal contrast in each PB-subdomain is less than some predefined tolerance r . We define the maximal contrast independently for each coefficient present in the problem Eq. (1), thus defining two (different) thresholds. Then, one can find a rPB-partition, which we denote by Θ_{pb}^r , such that

$$\frac{\alpha_{\max}(\mathcal{D})}{\alpha_{\min}(\mathcal{D})} < r_\alpha \quad \text{and} \quad \frac{\beta_{\max}(\mathcal{D})}{\beta_{\min}(\mathcal{D})} < r_\beta \quad \forall \mathcal{D} \in \Theta_{\text{pb}}^r$$

where $\{r_\alpha, r_\beta\} \geq 1$. Hence, the choice of both thresholds will determine the partition Θ_{pb}^r as a sub-partition of the original partition Θ . Note that if we consider $r_\alpha = r_\beta = \infty$, we recover the original partition Θ , while lower values for the thresholds lead to an increasing number of subparts, consequently *globs*, and thus richer coarse spaces. The rPB-BDDC preconditioner can be defined for any value of the threshold $r > 1$. By tuning r one can obtain the right balance between computational time and robustness.

As coefficients α, β are no longer constant in each rPB-subdomain, we propose to use averaged coefficients in Eq. (7) in order to define the averaging operator. The averaged coefficients, denoted as $\bar{\alpha}$ and $\bar{\beta}$, are computed in the rPB-subdomain Ω_k simply as

$$\bar{\alpha} = \frac{1}{|\Omega_k|} \int_{\Omega_k} \alpha dx, \quad \bar{\beta} = \frac{1}{|\Omega_k|} \int_{\Omega_k} \beta dx.$$

Hence, the definition of the averaging operator Eq. (7) is not modified and all DOFs on the same coarse geometrical entity are weighted by the same constant value. Thus, under the Θ_{pb}^r , the preconditioner expression is written exactly in the same form as in Sect. 3.2.

4. IMPLEMENTATION ASPECTS

In this section, we expose implementation strategies for some key points that the authors find of interest for potential users/developers of similar methods, namely an edge partition algorithm to avoid problematic cases in (unstructured) PB-partitions, the construction of the change of basis, the implementation of the original BDDC constraints and the aggregation of cells into rPB-subdomains based on heterogeneous coefficients α, β and the thresholds r_α, r_β . For a comprehensive implementation strategy of arbitrary order curl-conforming tetrahedral/hexahedral FEs, the reader is referred to [14].

4.1. Coarse edge partition. Special care has to be taken with the general definition for subdomain edges presented in Sect. 2.3. In particular, when *globs* are generated based on Θ_{pb} or a partition obtained with graph partitioners, e.g. METIS, the presented definition of E in Sect. 2.3 may not be sufficient for expressing the function and the coarse DOFs in the new basis. We detail the pathological cases identified in [25] (cases [1] and [2] below), and extension of a case in [25] (case [3]) plus an additional case (case [4]), for which we provide examples. We propose a unique cure, based on the partition of problematic coarse edges E into coarse sub-edges E_j such that the problematic cases are solved.

- [1] *Disconnected components.* We say that fine edges $e \in E$ are connected if they have an endpoint in common. Consequently, if a coarse edge E has m disconnected components, it has $2m$ endpoints. Note that, while this fact does not preclude the invertibility of the change of basis, if each of the components is treated as a coarse edge we recover original meaningful definitions for continuity constraints across subdomains.
- [2] *Interior node in touch with another subdomain.* This case occurs when an internal node v to E does not have the same set of subdomains as $\text{neigh}_{\Theta_{\text{pb}}}(E)$, i.e., is shared by $\text{neigh}_{\Theta_{\text{pb}}}(E)$ plus additional subdomains. In fact, v is then an element of Λ_C in the classification provided in Sect. 2.3. We recall that the change of variables is made for gradients of scalar, Lagrangian functions $\nabla\phi_E$ defined on all internal nodes of E . However, if we consider a nodal shape function associated to v , it will be coupled with other internal nodal DOFs for E , thus introducing a coupling between an external subdomain to $\text{neigh}_{\Theta_{\text{pb}}}(E)$ and itself, which is clearly not present in the original basis. A remedy for it consists of simply splitting the coarse edge E into two sub-edges at the problematic node v . Let us denote by V_p the subset of this kind of nodes for all $e \in E$.
- [3] *Edge n -furcation.* This situation occurs when a coarse edge E that does not have disconnected components has more than two endpoints. At some internal node the coarse edge is n -furcated into n edges, so the definition of the shape function Φ_E in the new basis loses its original meaning. Furthermore, this fact precludes the locality of the change of basis for every edge E . In this case, a simple remedy is again to split the edge into sub-edges at any node shared by more than two edges.
- [4] *Closed loop.* In this case we cannot identify endpoints for a coarse edge and therefore define a unique orientation for it. Furthermore, the new set of basis functions

is not well defined since the definition in Sect. 3.1 relies on the fact that every edge has 2 end points, thus not being applicable in this case. In this situation, an internal node for the coarse edge E must be chosen as start/end point (common in all subdomains) to assign an orientation to the edge and be treated as an edge endpoint in the change of basis definition.

In order to address all the presented problematic cases we propose a simple algorithm based on a classification $\forall e \in E$ into sub-edges. Our goal is to find a partition of $e \in E \in \Lambda_E$ into E_j such that every E_j is constructed connecting fine edges that share (only) one vertex with the following edge. Therefore, every coarse sub-edge $E_j \subset E$ has a unique starting point, a chain of connected *fine* edges sharing only one node and a unique end-point, which defines its unique orientation across all subdomains. Let us consider the set of nodes $V = \cup_{e \in E} (v \in \partial e)$, where the number of occurrences for each node $v \in V$ is denoted by $\text{count}(v)$. First, we can identify the set of nodes where E is n -furcated as

$$n\text{-furcation nodes } V_N \doteq \{v \in V \setminus V_p \mid \text{count}(v) > 2\}$$

We note that V_p is already identified in the *glob* generation algorithm. Then, we can find a partition of the set of nodes into the two following subsets:

$$\begin{aligned} \text{Edge boundary nodes } V_B &\doteq \{v \in V \mid \text{count}(v) = 1\} \cup V_p \cup V_N \\ \text{Interior nodes } V_I &\doteq \{v \in V \setminus V_B\}. \end{aligned}$$

Note that by definition of interior nodes, they are such that $\text{count}(v) = 2$. Such classification is performed by simply counting the number of appearances of nodes plus setting problematic nodes belonging to other objects as edge boundary nodes. Then, the coarse edge partitioning Alg. 1 finds paths from one edge boundary node (with a global criteria to select it) until the following edge boundary node. Furthermore, in this procedure we identify the direction of every fine edge with regard to its container coarse edge.

```

Data:  $V = V_B \sqcup V_I$ ,  $e \in E$ 
Result:  $E_j$  s.t.  $E = \sqcup E_j$ 
 $j \leftarrow 0$ 
while  $\text{card}(V) > 0$  do
  if  $\text{card}(V_B) > 0$  then
    | Find  $v^s \in V_B$  with minimum global id
  else
    | Find  $v^s \in V_I$  with minimum global id
    |  $V_I \leftarrow V_I \setminus v^s$  and  $V_B \leftarrow V_B \cup v^s$ 
  end
   $j \leftarrow j + 1$ 
  Find  $v^e$  s.t.  $\{v^s, v^e\} \in \partial e$  with minimum global id     $E_j \leftarrow \{e\}$ 
  update counters and subsets (Alg. 2)
  while  $v^e \in V_I$  do
    |  $v^s \leftarrow v^e$ 
    | Find  $v^e$  s.t.  $\{v^s, v^e\} \in \partial e$      $E_j \leftarrow E_j \cup e$ 
    | update counters and subsets (Alg. 2)
  end
end

```

Algorithm 1: Edge partition algorithm

From this point onwards, we consider that each $E \in \Lambda_E$ is a (sub-)edge of the original coarse edges such that they do not present problematic cases.

4.2. Change of basis. In this section we provide some implementation details of the change of basis described in Sect. 3.1. In the application of the averaging operator in

Data: v^s, v^e, V_B, V_I
Result: V_B, V_I
for k *in* s, e **do**
 $\text{count}(v^k) \leftarrow \text{count}(v^k) - 1$
 if $\text{count}(v^k) = 0$.and. $v^k \in V_B$ **then**
 $V_B \leftarrow V_B \setminus v^k$
 else if $\text{count}(v^k) = 0$.and. $v^k \in V_I$ **then**
 $V_I \leftarrow V_I \setminus v^k$
 end
end

Algorithm 2: Update counters and subsets

Eq. (9), we note that one must apply the global change of basis and its restriction to subdomains. A practical implementation of their application in both cases can be performed with the local restriction of the change of basis to the subdomains, i.e., $\mathcal{Q}_i = R_i \widehat{\mathcal{Q}} R_i^T$, thus it can be performed in parallel in distributed memory environments. The application of the inverse of the change of basis in the sub-assembled space \mathcal{Q}^{-1} or \mathcal{Q}^{-T} (Eq. (9)) can be performed in parallel, i.e., relying on the restriction of the operators to the subdomains, given its definition (see Sect. 3.1). On the other hand, the sparsity pattern of the global change of basis $\widehat{\mathcal{Q}}$ can be exploited in order to achieve a parallel implementation of the application of $\widehat{\mathcal{Q}}$ and $\widehat{\mathcal{Q}}^T$ to a function $\mathbf{u} \in \widehat{X}_h$ that only relies on restricted (to the subdomains) information.

Proposition 4.1. *The expression $R_i \widehat{\mathcal{Q}} \mathbf{u} = \mathcal{Q}_i \mathbf{u}^i$ holds, where $\mathbf{u} \in \widehat{X}_h$ and $\mathbf{u}^i \in X_h^i$.*

Proof. By definition of the change of basis matrix, it is easy to check that $R_i \widehat{\mathcal{Q}} \mathbf{u}$ only depends on the DOF values of \mathbf{u}^i (in the new base), i.e., $R_i \widehat{\mathcal{Q}} \mathbf{u} = R_i \widehat{\mathcal{Q}} R_i^T R_i \mathbf{u}$. Thus we can write

$$R_i \widehat{\mathcal{Q}} \mathbf{u} = R_i \widehat{\mathcal{Q}} R_i^T R_i \mathbf{u} = \mathcal{Q}_i \mathbf{u}^i.$$

□

Unfortunately, this reasoning cannot be applied to the transpose of the change of basis.

Proposition 4.2. *Consider arbitrary local weighting diagonal matrices W_j for every subdomain such that $\mathbf{u} = \sum_j R_j^T W_j R_j \mathbf{u}$, i.e., they form a partition of the unity. Then, the expression $R_i \widehat{\mathcal{Q}}^T \mathbf{u} = R_i \sum_j R_j^T \mathcal{Q}_j^T W_j \mathbf{u}^j$ holds, where $\mathbf{u} \in \widehat{X}_h$ and $\mathbf{u}^j \in X_h^j$.*

Proof. Using the fact that $\sum_j R_j^T W_j R_j$ is the identity matrix and the fact that $R_j \widehat{\mathcal{Q}} = R_j \widehat{\mathcal{Q}} R_j^T R_j$ (as above), it holds:

$$R_i \widehat{\mathcal{Q}}^T \mathbf{u} = R_i \sum_j \widehat{\mathcal{Q}}^T R_j^T W_j R_j \mathbf{u} = R_i \sum_j R_j^T R_j \widehat{\mathcal{Q}}^T R_j^T W_j R_j \mathbf{u} = R_i \sum_j R_j^T \mathcal{Q}_j^T W_j \mathbf{u}^j.$$

□

Therefore, the application of the change of basis can rely only on restrictions of the same to subdomains whereas the application of the transpose change of basis can be performed in parallel with subdomain restrictions plus nearest neighbour communications. With this purpose in mind, we detail here how to implement the restriction of the change of basis local to subdomains. Let us define a partition of the DOFs in $X_h^{i'}$ into three subsets of DOFs, namely: the DOFs placed on $e \in E$, the DOFs placed on interface edges/faces $f \notin E$ such that $\partial f \cap E \neq \emptyset$, and the remaining DOFs in $X_h^{i'}$, denoted by \mathbf{u}_E , \mathbf{u}_f and \mathbf{u}_r , respectively. Furthermore, let us consider that DOFs in \mathbf{u}_E are sorted such that DOFs belonging to the same coarse edge E are found in consecutive positions. Note that shape

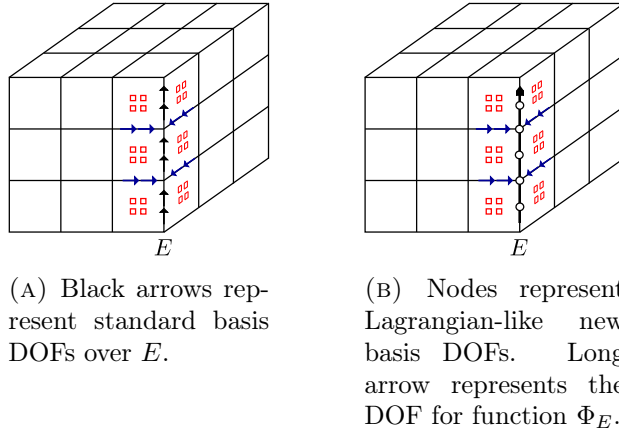


FIGURE 1. Standard (old) and new basis DOFs for 3D hexahedra second order edge FE over E . Additional depicted DOFs values are affected by the change of basis for E , while the remaining DOFs are invariant under the change of basis.

functions associated to \mathbf{u}_f , \mathbf{u}_r are common in both (i.e., old and new) bases. For edge FEs of order k , k moments (i.e., DOFs) are defined on each $e \in \mathcal{T}_h$. Let us denote by n_e the number of fine edges $e \in E$. Then, the total number of DOFs on a coarse edge E is kn_e . On the other hand, the number of Lagrangian-like DOFs interior to E (i.e., excluding ∂E) is $(kn_e - 1)$, i.e., the number of shape functions of the type $\nabla\phi_E$. The change of basis is completed with the addition of the function Φ_E to the new basis so that the dimension of both bases coincides. For the sake of illustration, both sets of basis functions restricted to E are depicted in Fig. 1.

Let us denote by n_E the number of coarse edges $E \in \Lambda_E$ for a given subdomain. Then, we define the change of basis \mathcal{Q}^{E_j} , $j = \{1, \dots, n_E\}$ local to every $E_j \in \Lambda_E$ as

$$(11) \quad \mathcal{Q}_{ab}^{E_j} = \sigma_a(\nabla\phi_E^b), \quad \text{for } b = 1, \dots, kn_e^{E_j} - 1 \quad \mathcal{Q}_{a, kn_e^{E_j}}^{E_j} = \sigma_a(\Phi_{E_j}),$$

where σ_a , $a = 1, \dots, kn_e^{E_j}$, are the (original basis) edge moments defined on E_j (the superscript in $n_e^{E_j}$ has been introduced to show that it depends on the coarse edge). We can now define the change of basis $\mathcal{Q}^E = \text{diag}(\mathcal{Q}^{E_1}, \dots, \mathcal{Q}^{E_{n_E}})$ local to coarse edges. We remark that the same orientation for every coarse edge $E \in \Lambda_E$ must be defined on the set of subdomains $\mathcal{D} \in \text{neigh}_\Theta(E)$. Otherwise, the definition of the new basis function Φ_E is not consistent across subdomains. In addition, the change of basis must take into account the effect of the new DOF values \mathbf{u}_E associated to $\nabla\phi_E^j$ and Φ_E for E in the old values \mathbf{u}_f . Thus, we evaluate for all indices b of shape functions associated to \mathbf{u}_E

$$(12) \quad \mathcal{Q}_{cb}^f = \sigma_c(\nabla\phi_E^b),$$

where c corresponds to the index of all moments associated to \mathbf{u}_f . We recall that, by definition, $\sigma_c(\Phi_E) = 0$. The application of the moments σ_c to the (original) shape functions associated to \mathbf{u}_f results in $\sigma_c(\varphi^b) = \delta_{cb}$. Finally, DOFs in \mathbf{u}_r are invariant under the change of basis. Note that the definition of Eqs. (11) and (12) related to the gradients of the scalar shape functions coincides with the so-called discrete gradient operator related to these functions, as used in [15]. However, we prefer to motivate the change of basis with the usage of the Nédélec interpolator, since it naturally provides the definition of the entries related to the unit tangent function, while suitable eigenvectors to complete the change of basis are computed in [15]. The structure of the change of basis restricted to a subdomain is

$$\mathbf{u}_{\text{old}} = \begin{bmatrix} \mathbf{u}_E \\ \mathbf{u}_f \\ \mathbf{u}_r \end{bmatrix}_{\text{old}} = \begin{bmatrix} \mathcal{Q}^E & 0 & 0 \\ \mathcal{Q}^f & I & 0 \\ 0 & 0 & I \end{bmatrix} \begin{bmatrix} \mathbf{u}_E \\ \mathbf{u}_f \\ \mathbf{u}_r \end{bmatrix}_{\text{new}} = \mathcal{Q}_i \mathbf{u}_{\text{new}},$$

where it becomes clear the fact that the inverse of the change of basis is well defined if and only if \mathcal{Q}^E is invertible. In turn, \mathcal{Q}^E will be invertible if and only if every change of basis local to $E \in \Lambda_E$ (Eq. (11)) is invertible.

Remark 4.3. *Although it is used in this exposition for the sake of clarity, we do not require any particular ordering of DOFs in a practical implementation of \mathcal{Q}_i .*

4.3. BDDC constraints. In this subsection we propose a practical manner of computing the BDDC constraints Eqs. (6a) and (6b) for local problems. In our implementation, constraints over local problems are strongly imposed through the usage of Lagrange multipliers on the original basis. Therefore, the local matrix \mathcal{A}_h^i is extended with the discrete version of the constraints C in order to obtain constrained (Neumann) local problems.

The computation of constraints requires to integrate zero and first order moments for the solution over all coarse edges E . We note that the first constraint Eq. (6a) can be easily implemented for k -order edge FEs as

$$(13) \quad \begin{aligned} s_{0,E}(\mathbf{u}^i) &\doteq \int_E \mathbf{u}^i \cdot \mathbf{t}_E ds = \int_E \left(\sum_{a=1}^{kn_e} u^a \boldsymbol{\varphi}^a \right) \cdot \mathbf{t}_E ds = \sum_{a=1}^{kn_e} u^a \int_E \boldsymbol{\varphi}^a \cdot \mathbf{t}_E ds \\ &= \sum_{a=1}^{kn_e} u^a \int_E \boldsymbol{\varphi}^a \cdot \mathbf{t}_E \left(\sum_{b=1}^k p_b \right) ds = \sum_{a=1}^{kn_e} (\mathbf{t}_e \cdot \mathbf{t}_E) u^a = \sum_{a=1}^{kn_e} C^a u^a, \end{aligned}$$

where we used the partition of the scalar, unit function into the set of Lagrangian test functions p_b belonging to the polynomial space $\widehat{V}_h(E)$ of order $k-1$. These functions are used for defining the k local moments on every e as $\sigma^b(\mathbf{u}) = \int_e \mathbf{u} \cdot \mathbf{t}_e p_b$, $b = \{1, \dots, k\}$ (see [14] for details). Its duality with basis shape functions, i.e., $\sigma^b(\boldsymbol{\varphi}^a) = \delta_{ba}$, has been used in Eq. (13). Thus, to compute the first constraint one only needs to add ± 1 at the corresponding entry in C for each DOF, where the sign is determined by the agreement between fine and coarse edge orientations, i.e., $\mathbf{t}_e \cdot \mathbf{t}_E$. On the other hand, the computation of the second constraint Eq. (6b) requires to define an arc-length parameter over E . A practical implementation of the constraint Eq. (6b) can avoid it by considering the constraint in the new basis. Since ϕ_E^j vanish at ∂E , integration by parts yields [13]

$$\begin{aligned} s_{1,E}(\mathbf{u}^i) &\doteq \int_E s \mathbf{u}^i \cdot \mathbf{t}_E ds = \int_E s (u_E \Phi_E + \sum_{a=1}^{kn_e-1} u^a \nabla \phi_E^a) \cdot \mathbf{t}_E ds = \\ &= \sum_{a=1}^{kn_e-1} u^a \int_E \phi_E^a ds = - \sum_{a=1}^{kn_e-1} C_{\text{new}}^a u^a, \end{aligned}$$

where the contribution of Φ_E is null due to the antisymmetry of the product $s\Phi_E$ (we recall that $s \in [-|E|/2, |E|/2]$) over E . Then, we can apply the change of basis to obtain the expression in the original basis, i.e., $C = C_{\text{new}} \mathcal{Q}_i^{-1}$.

4.4. Building Θ_{pb}^r . In the rPB-BDDC method, a Θ_{pb}^r partition is used such that the maximal contrast, for each one of the coefficients, is lower than a predefined tolerance r in each subdomain. Our goal is to identify a partition of every subdomain into $\mathcal{D}' \in \Theta_{\text{pb}}^r$ subdomains where the thresholds $\frac{\alpha_{\text{max}}(\mathcal{D}')}{\alpha_{\text{min}}(\mathcal{D}')} < r_\alpha$ and $\frac{\beta_{\text{max}}(\mathcal{D}')}{\beta_{\text{min}}(\mathcal{D}')} < r_\beta$ are respected. It can be accomplished using different algorithms. One approach is to consider a seed cell and aggregate the surrounding cells such that the contrast(s) are below the given threshold(s),

proceeding recursively until no neighbouring cells can be aggregated. We take another seed among the non-aggregated cells and proceed again until all cells have been processed.

Alternatively, one can first determine the maximum and minimum values for α and β in a given subdomain $\mathcal{D} \in \Theta$. With this information and the thresholds r_α , r_β , we can determine the number of sub-intervals for every subdomain and coefficient as follows. First, we compute $\ell_\alpha(\mathcal{D})$ and $\ell_\beta(\mathcal{D})$ as the smallest positive integers for which

$$\frac{\alpha_{\max}(\mathcal{D})}{\alpha_{\min}(\mathcal{D})} < r_\alpha^{\ell_\alpha(\mathcal{D})}, \quad \frac{\beta_{\max}(\mathcal{D})}{\beta_{\min}(\mathcal{D})} < r_\beta^{\ell_\beta(\mathcal{D})},$$

respectively. One can now define the intervals

$$(14) \quad I_{i,j} \doteq [r_\alpha^{i-1} \alpha_{\min}(\mathcal{D}), r_\alpha^i \alpha_{\min}(\mathcal{D})] \times [r_\beta^{j-1} \beta_{\min}(\mathcal{D}), r_\beta^j \beta_{\min}(\mathcal{D})],$$

for $i \in [1, \ell_\alpha(\mathcal{D})]$, $j \in [1, \ell_\beta(\mathcal{D})]$. Cells with their coefficients on the same interval $I_{i,j}$ are aggregated in a PB-subdomain. Those cells that have coefficients across multiple sub-intervals are treated as additional PB-subdomains. This definition allows one to isolate cells that contain abrupt jumps in the value of the coefficients, while the user has the freedom to select the thresholds such that unnecessary extra PB-subdomains are avoided.

For the sake of illustration, we include an example where all the different subset indices are presented for a unit cube domain: the original partition into $P = 3 \times 3 \times 3$ subdomains in Fig. 2a, the aggregation of cells into subsets based on $\log(\beta) = 3 \sin(3\pi y)$ (see Fig. 3a) for $r = 10^3$ in Fig. 3b, combined with an analogous partition for $\log(\alpha) = 3 \sin(3\pi x)$ leading to a coefficient-based partition in Fig. 2b and the final rPB-partition Θ_{pb}^r in Fig. 2c.

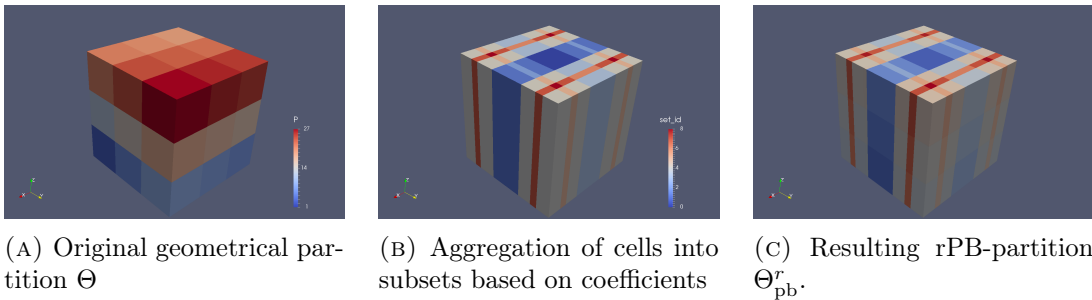


FIGURE 2. Partitions for scalar coefficients described by $\log(\alpha) = 3 \sin(3\pi x)$ and $\log(\beta) = 3 \sin(3\pi y)$ with an initial $3 \times 3 \times 3$ partition of the unit cube.

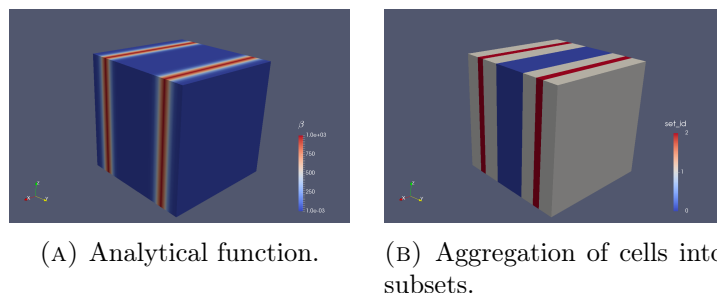


FIGURE 3. Aggregation of cells into subsets based on $\log(\beta) = 3 \sin(3\pi y)$ with threshold $r = 10^3$.

We note that the preconditioner is very robust despite the presence of subdomains with large aspect ratios. Support to this observation can be found in [29] (see Remark 3.12 and the corresponding numerical experiments).

5. NUMERICAL RESULTS

In this section we evaluate the weak scalability of the proposed preconditioner for the problem in Eq. (1), within the preconditioned conjugate gradient (CG) Krylov iterative solver. The robustness of the rPB-BDDC-CG solver is tested in 3D simple domains, which are discretized either with structured or unstructured meshes. As performance metrics, we focus on the number of rPB-BDDC preconditioned CG iterations required to attain the convergence criteria, which is defined as the reduction of the initial residual algebraic ℓ_2 -norm by a factor 10^{-6} . On the other hand, the total computation time will be presented, which will include both preconditioner set-up and the preconditioned iterative solution of the linear system in all the experiments reported. The particular definition of coefficients α and β and its distribution will be specified throughout the section for each case.

We have also applied a plain CG solver for the different tests, but the results of the latter are not reported here due to its poor performance. As an example, for the same problem set-up as in Fig. 8, only 48 processors, and $H/h = 30$, it did not satisfy the convergence criteria within a prescribed maximum of 20,000 iterations, for which it consumed 144 secs, far beyond the proposed rPB-BDDC preconditioner.

5.1. Experimental framework. The rPB-BDDC methods have been implemented in the scientific software project **FEMPAR** [4]. **FEMPAR**, developed by the Large Scale Scientific Computing (LSSC) team at CIMNE-UPC, is a parallel hybrid OpenMP/MPI, software package for the massively parallel FE simulation of multiphysics problems governed by PDEs. **FEMPAR** offers a set of flexible data structures and algorithms for each step in the simulation pipeline, which can be customized in order to meet particular application problem needs. See [5] for a thorough coverage of the software architecture of **FEMPAR**. Among other features, it provides the basic tools for the efficient parallel distributed-memory implementation of substructuring DD solvers [33, 6], based on a *fully-distributed* implementation of data structures involved in the parallel simulation. The parallel codes in **FEMPAR** heavily use standard computational kernels provided by BLAS and LAPACK. Besides, through proper interfaces to several third party libraries, the local constrained Neumann problems and the global coarse-grid problem can be solved via sparse direct solvers. **FEMPAR** is released under the GNU GPL v3 license, and is more than 300K lines of Fortran200X code long following object-oriented design principles. In this work, we use the overlapped BDDC implementation proposed in [34], with excellent scalability properties. It is based on the overlapped computation of *coarse* and *fine* duties. As long as coarse duties can be fully overlapped with fine duties, perfect weak scalability can be attained. We refer to [6] for more details.

The experiments in this section have been performed on the MareNostrum-IV [35] (MN-IV) supercomputer, hosted by the Barcelona Supercomputing Center (BSC). In all cases, we consider a one-to-one mapping among subdomains, cores and MPI tasks. Provided that the algorithm allows for a high degree of overlapping between fine and coarse duties, an additional MPI task is spawn into a full node (i.e., 48 cores) in order to perform the coarse problem related tasks. The multi-threaded PARDISO solver in Intel MKL is used to solve the coarse-grid problem within its computing node.

Unless otherwise stated, the problem Eq. (1) will be solved in the unit cubic domain $\Omega = [0, 1]^3$ with Dirichlet homogeneous boundary conditions on the whole boundary and the forcing term $\mathbf{f} = 1$. Let us denote by h the usual mesh element size, and by H the size of the subdomain. Then, local problem sizes can be characterized in a structured mesh and partition by $\frac{H}{h}$. In order to perform a weak scalability analysis, we build a set of structured meshes consisting on $(4\frac{H}{h}k \times 4\frac{H}{h}k \times 3\frac{H}{h}k)$ hexahedra. A uniform partition of the meshes into $P = (4k \times 4k \times 3k) = 48k^3$ subdomains is considered, where local problem sizes are $(\frac{H}{h})^3$.

5.2. Homogeneous problem. Let us first consider homogeneous coefficients $\alpha = \beta = 1.0$ for the whole domain Ω . We test the problem with different local problem sizes $\frac{H}{h}$ and FE orders. In this case, the PB-BDDC preconditioner reduces to the standard BDDC preconditioner since $\Theta = \Theta_{\text{pb}}$. In Fig. 4, we present weak scalability results for the homogeneous problem up to 16464 subdomains with different local problem sizes $\frac{H}{h} = \{10, 20, 30\}$, where the largest case has more than 10^9 DOFs. We present the number of solver iterations until convergence in Fig. 4a, and employed wall clock times in Fig. 4b, which are composed by the preconditioner set-up time and the solution time with the BDDC preconditioned solver. The plots indicate that both the algorithm and its implementation in **FEMPAR** have excellent weak scalability properties. Provided that the algorithm overlaps fine and coarse tasks, coarse tasks computing times are masked as long as they not exceed computing times for local problems, which allow us to observe excellent weak scalability times for the largest case $\frac{H}{h} = 30$ in Fig. 4b. The size of both local and coarse problems is presented in Fig. 4c. Finally, we add a plot (Fig. 4d) of the time needed to set-up the change of basis, which includes the edge partition algorithm (Alg. 1) to detect problematic cases, to show that consumed time is only dependent on the local problem size and not significant compared to the one spent in the solver run.

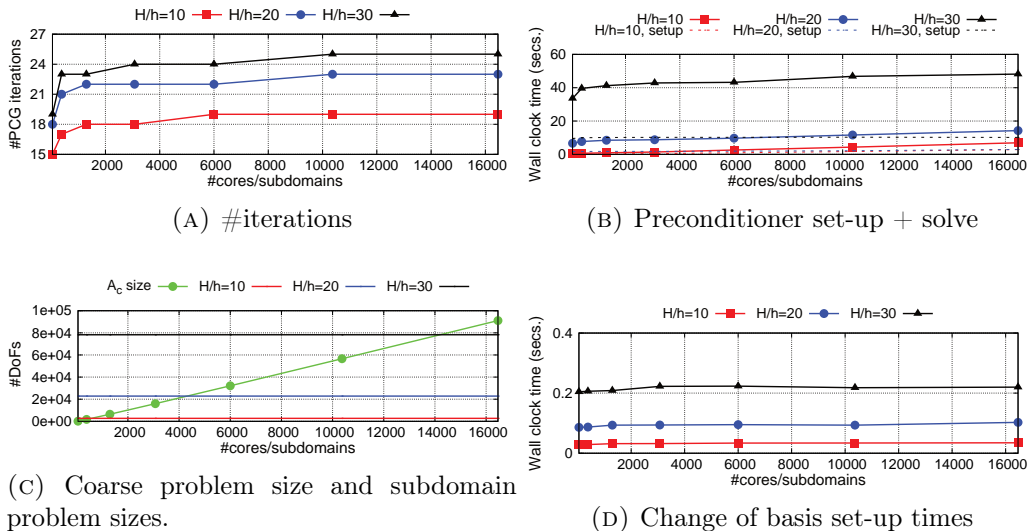


FIGURE 4. Weak scalability results for first order edge FE with a constant distribution of materials for different local problem $\frac{H}{h}$ sizes. Subdomain problem sizes in 4c are given for comparison purposes against the coarse problem size.

Fig. 5 shows the weak scalability results for an homogeneous problem with constant coefficients and different FE orders up to 4. First, the number of iterations is (asymptotically) constant, thus the method is scalable. In this case, local problem sizes are such that the coarse problem is larger from a small number of subdomains (see Fig. 5c), thus it is reflected in the solver times plot in Fig. 5b. In Fig. 5d, the time spent in the set-up of the change of basis is presented. Out of the presented results for the homogeneous case, a clear conclusion can be drawn for the standard BDDC: the algorithm and its implementation have excellent weak scalability properties.

5.3. Multi-material problem.

5.3.1. Checkerboard distribution. The checkerboard arrangement of coefficients is a widely used distribution of materials to test the robustness of the BDDC algorithms for problems in $H(\text{curl})$ against the jump of coefficients across the interface [13, 25]. In short, it is a

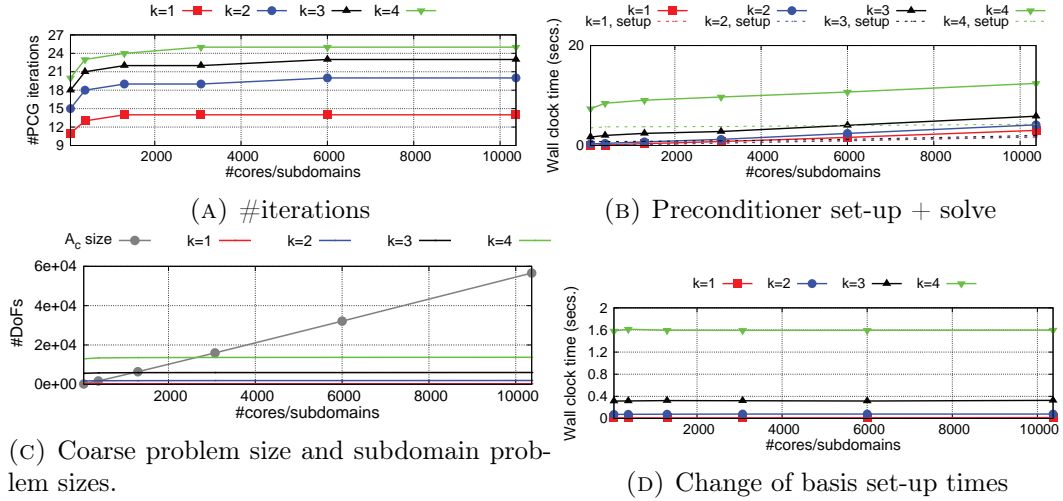


FIGURE 5. Weak scalability results for different order edge FE with a constant distribution of materials and a local problem $\frac{H}{h} = 10$. Subdomain problem sizes in 5c are given for comparison purposes against the coarse problem size.

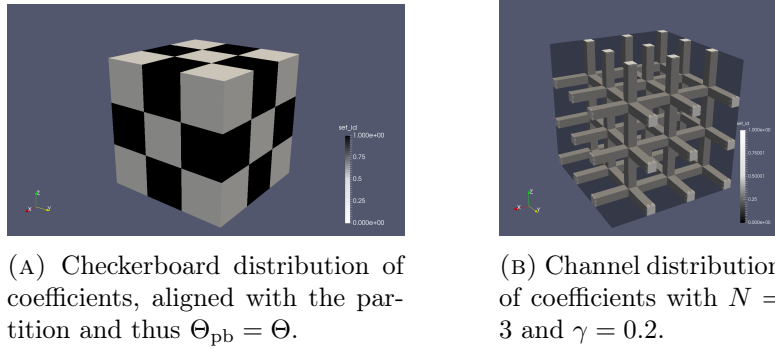


FIGURE 6. Bi-material distribution cases in a structured $3 \times 3 \times 3$ partition of the unit cube $[0, 1]^3$.

bi-material distribution of subdomain-wise constant coefficients such that every subdomain presents a jump of coefficients through the faces to all its neighbours. For the sake of ease, let us distinguish between black and white subdomain materials, see Fig. 6a. Note that in the checkerboard distribution case, the jumps of coefficients are aligned with the partition, thus $\Theta_{pb} = \Theta$.

We first test the robustness of the algorithm against the contrast of the coefficients. Consider a $3 \times 3 \times 3$ partition of a unit cube domain with a checkerboard distribution of coefficients such that $\alpha_{\text{white}} = \beta_{\text{white}} = 1.0$ and $\alpha_{\text{black}} = 10^i$, $\beta_{\text{black}} = 10^{-i}$. The contrast is defined here as $\frac{\alpha_{\text{black}}}{\beta_{\text{black}}}$. With the variation of the value for i in the range $[-5, 5]$ we test all the possible scenarios, namely the mass dominated problem ($i < 0$) and the curl dominated problem ($i > 0$). The number of iterations with the contrast of the coefficients for different configurations of the preconditioner is presented in Fig. 7. Out of the plot, the most salient property is the robustness of the perturbed preconditioner (see Sect. 3.3) with the contrast of the coefficients. In fact, the original formulation of the preconditioner suffers from a large number of iterations when the contrast between the two coefficients is large, specially in the curl-dominated case. Therefore, the proposed perturbation of the preconditioner is essential to achieve a robust preconditioner, in the case where both

$\frac{H}{h} / \mathbf{P}$	2^3	3^3	4^3	5^3	6^3	7^3	8^3	$\frac{H}{h} / \mathbf{P}$	2^3	3^3	4^3	5^3	6^3	7^3	8^3
4	14	24	35	38	40	40	41	4	8	9	10	10	11	12	12
8	26	37	61	65	70	69	70	8	12	14	16	16	17	17	17
12	31	52	72	78	82	82	84	12	15	22	21	21	21	21	21

(A) Standard BDDC preconditioner

(B) Perturbed BDDC preconditioner

TABLE 1. Weak scalability in terms of number of iterations for both preconditioners. Checkerboard distribution of materials with $\alpha_{\text{white}} = 10^2$, $\beta_{\text{white}} = 1.0$ and $\alpha_{\text{black}} = 10^4$ and $\beta_{\text{black}} = 10^{-2}$.

coefficients α and β jump across the interface. Clearly, the perturbed formulation only has a (negligible) negative impact in the case $i = 0$, since actually no jump occurs across the interface. In the curl-dominated limit, α and ω -based scalings show the same behaviour, as it is suggested by the definition of ω when $\alpha \gg \beta$. On the other hand, when the coefficient β becomes dominant, the choice of *cardinal* and ω -based scalings also leads to good scalability results in the limit. In summary, the combination of the perturbed formulation and α -based scaling is the most robust approach. Unless otherwise stated, this combination will be used throughout the section.

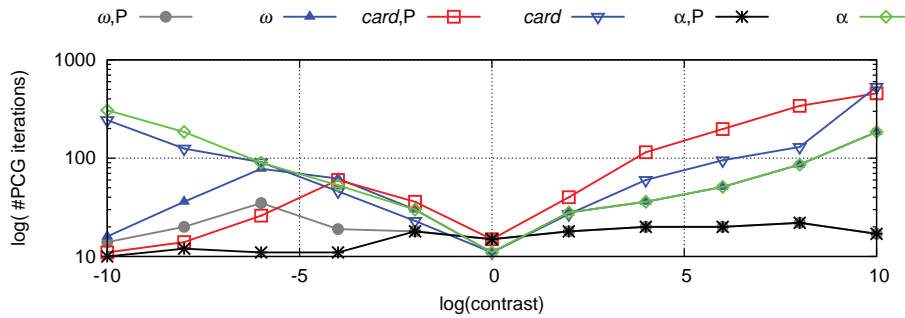


FIGURE 7. Number of iterations for first order edge FE with a $3 \times 3 \times 3$ partition of the unit cube and $H/h = 8$. A checkerboard arrangement of materials is defined: $\alpha_{\text{white}} = 1.0$, $\beta_{\text{white}} = 1.0$ and $\alpha_{\text{black}} = 10^i$ and $\beta_{\text{black}} = 10^{-i}$, leading to a contrast $= \frac{\alpha_{\text{black}}}{\beta_{\text{black}}} = \frac{10^i}{10^{-i}} = 10^{2i}$. Labels include scaling information, where P denotes perturbation of the preconditioner.

Let us now consider a checkerboard arrangement of coefficients such that $\alpha_{\text{white}} = 10^2$, $\beta_{\text{white}} = 1.0$ and $\alpha_{\text{black}} = 10^4$ and $\beta_{\text{black}} = 10^{-2}$. In order to show the importance of the perturbed formulation of the preconditioner for jumps of both coefficients across interfaces, we collect the number of iterations for the original and perturbed preconditioner in Tabs. 1a and 1b, respectively. The problem is solved with a $P = N \times N \times N$ partition of the unit cube and the ω -based scaling is employed in both cases. Iteration counts for the perturbed preconditioner are noticeably lower in all cases without exception.

Once we have shown the importance of the perturbed formulation, we present a weak scalability analysis up to 16,464 subdomains and the checkerboard arrangement of materials with the perturbed preconditioner. Problem sizes in this experiment coincide to the ones presented for the homogeneous problem in Fig. 4c. As expected, plots in Fig. 8 show excellent scalability properties of the preconditioner in this case, i.e., the preconditioner is robust with jumps of coefficients across the interface. Although higher values of $\frac{H}{h}$ lead to a significantly higher number of iterations, these ones are (asymptotically) constant and remain in a reasonable range, see Fig. 8a. On the other hand, Fig. 9a presents the number of iterations for different order FEs and problem size $H/h = 4$, which also is

shown to be scalable. Out of the contrast and scalability results, we would like to remark the following issues. First, the perturbed formulation of the preconditioner is essential to achieve a robust preconditioner. Second, the method is weakly scalable for problems with high coefficient jumps across interfaces for different local sizes and FE orders.

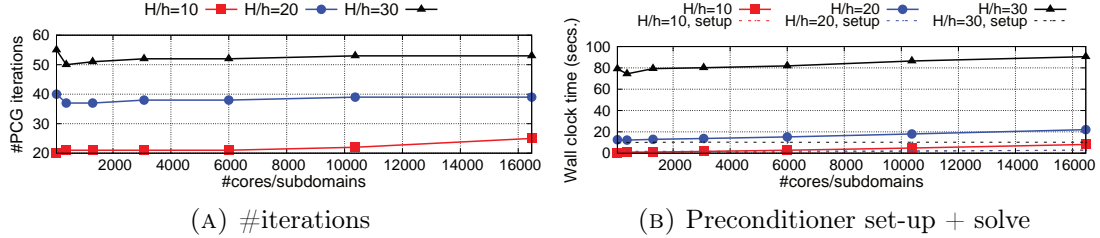


FIGURE 8. Weak scalability results for first order edge FE with a checkerboard distribution of materials: $\alpha_{\text{white}} = 10^2$, $\beta_{\text{white}} = 1.0$ and $\alpha_{\text{black}} = 10^4$ and $\beta_{\text{black}} = 10^{-2}$.

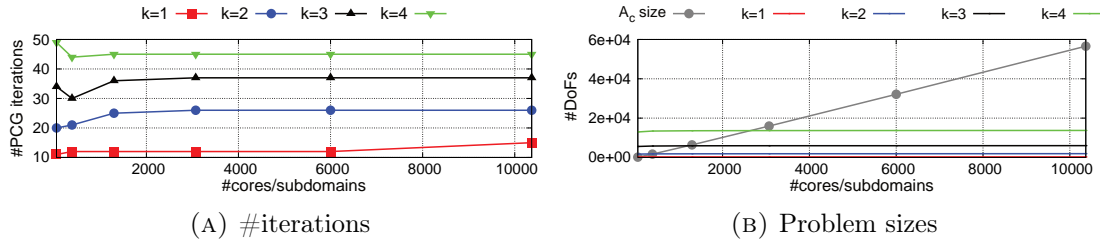


FIGURE 9. Weak scalability results for different order edge FEs and $H/h = 4$ with a checkerboard distribution of materials: $\alpha_{\text{white}} = 10^2$, $\beta_{\text{white}} = 1.0$ and $\alpha_{\text{black}} = 10^4$ and $\beta_{\text{black}} = 10^{-2}$.

In order to show the robustness of the method not only with structured, regular hexahedral meshes we solve the problem for a spherical domain and partition with a graph partitioner METIS. Let us consider a spherical domain with $R = 0.5$, discretized with an unstructured tetrahedral mesh containing around 50,000 cells. In order to achieve high contrast of coefficients across interfaces, a bi-material distribution of coefficients is assigned such white or black subdomain-wise constant materials are randomly assigned, see Fig. 10b for an illustration. The definition of the sets of coefficients is $\alpha_{\text{white}} = 10^2$, $\beta_{\text{white}} = 1.0$ and $\alpha_{\text{black}} = 10^4$ and $\beta_{\text{black}} = 10^{-2}$. Fig. 11 shows the number of iterations with the original and the perturbed formulation of the preconditioner. The perturbed preconditioner, combined with a α -based scaling, is the unique method shown to be robust with regard to the coefficients contrast, reproducing the behaviour observed in the structured case.

5.3.2. Multiple channels. In this arrangement of materials we have a background domain, denoted by black region, and a set of inclusions, denoted by white regions, that cross the domain from one boundary to the opposite one in parallel to the axes directions. We include one channel per direction per subdomain so that with an increasing number of P subdomains we are solving a harder problem with P channels. Channels are parallel to the axes and are positioned in the lowest (i.e., minimum x, y, z coordinates) corners within every subdomain. They have a squared cross-section of size γH , thus occupying a $\gamma^2 |\tilde{\Omega}_i|$ volume in every subdomain, see Fig. 6b for an illustration of a $3 \times 3 \times 3$ partition of the unit cube with the described channel inclusions. The distribution of coefficients is such that contains coefficient jumps within each subdomain and also across all interfaces, thus the perturbed formulation of the PB-BDDC preconditioner will be employed.

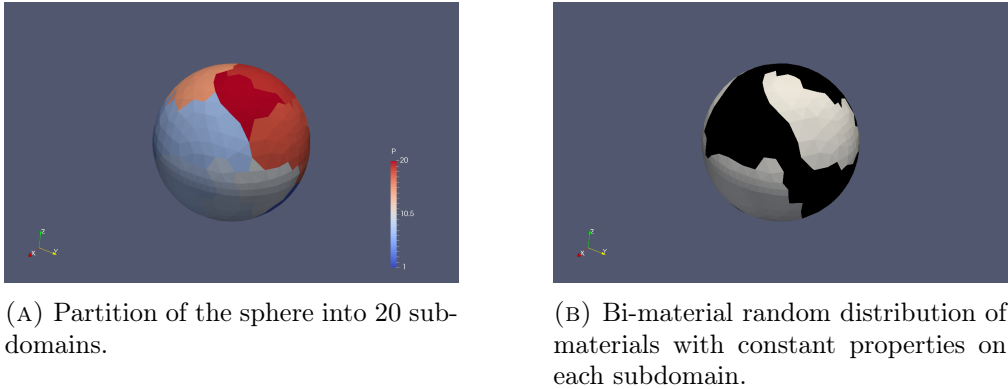


FIGURE 10. Sphere partition and distribution of materials.

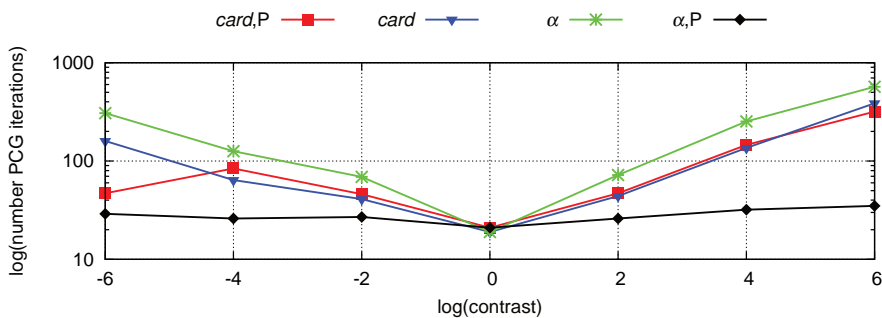


FIGURE 11. Robustness for a tetrahedral mesh for different scalings. Material parameters defined as $\alpha_{\text{white}} = 10^i$, $\alpha_{\text{black}} = 1.0$, $\beta_{\text{white}} = 10^{-i}$ and $\beta_{\text{black}} = 1.0$, contrast defined $\frac{\alpha_{\text{white}}}{\beta_{\text{white}}}$. Mesh partitioned into 20 subdomains and random assignment of materials. Labels include scaling information, where P denotes perturbation of the preconditioner.

Let us first compare the number of iterations for the PB-BDDC preconditioner against the ones that one would have with the standard BDDC, where the definition of *globs* is generated with the original partition Θ . In Sect. 5.3.1, the effectiveness of the perturbation formulation has been empirically shown. Consequently, in order to provide a fair comparison among them, the perturbed formulation is considered for both preconditioners. On the other hand, while α -based scaling is shown to be the most robust approach for PB-BDDC preconditioner, it miserably fails when considered with the standard *globs*, i.e., given by Θ . In this case, a better result is obtained with *cardinality* scaling. Tabs. 2a and 2b show iteration counts for the described BDDC and PB-BDDC preconditioners, respectively, for the solution of the channels problem with $\gamma = 0.5$ and a partition of the unit cube into $P = 3 \times 3 \times 3$ subdomains with local size $H/h = 8$. We define the coefficients $\alpha_{\text{black}} = \beta_{\text{black}} = 1.0$, while we distinguish between $\alpha_{\text{white}} = 10^i$ and $\beta_{\text{white}} = 10^{-i}$, which allow us to define the contrast as $\frac{\alpha_{\text{white}}}{\beta_{\text{white}}}$. As expected, the PB-BDDC preconditioner is robust with the contrast of coefficients. On the other hand, the number of iterations increases for the BDDC preconditioner in the curl-dominated case.

The following experiment evaluates the weak scalability properties for a channel-type distribution of materials. In Figs. 12 and 13 we present weak scalability results for different problem sizes and FE orders. The most salient property out of these plots is that the number of iterations is asymptotically constant for all cases. However, coarse problem sizes become larger as the partition into PB-subdomains generates a higher number of coarse DOFs, see Fig. 12c. In this context, the coarse problem is larger than local problem

$\frac{H}{h}/c$	10^{-4}	10^{-2}	1.0	10^2	10^4
4	36	29	13	31	74
8	67	36	16	38	104

(A) BDDC preconditioner

$\frac{H}{h}/c$	10^{-4}	10^{-2}	1.0	10^2	10^4
4	14	14	11	13	14
8	18	19	16	17	20

(B) PB-BDDC preconditioner

TABLE 2. Comparison in number of iterations for both preconditioners in a $3 \times 3 \times 3$ partition. Channel distribution of materials with $\gamma = 0.5$ and $\alpha_{\text{black}} = \beta_{\text{black}} = 1.0$, $\alpha_{\text{white}} = 10^i$ and $\beta_{\text{white}} = 10^{-i}$. Contrast defined as $c = \frac{\alpha_{\text{white}}}{\beta_{\text{white}}}$.

sizes from an small number of subdomains, thus coarse tasks will predominate computing times precluding wall clock time scalability, as it is shown in Fig. 12b. In Figs. 12d and 13b we present scalable wall clock times for the change of basis set-up, for different local problem sizes and FE orders.

We would like to remark that the proposed PB-BDDC preconditioner is weakly scalable for the number of iterations until convergence not only with regard to the jump of coefficients across interfaces but also for distributions of different materials within each subdomain. A multilevel version of the preconditioner for curl-conforming spaces [15], not addressed in this work, is expected to push forward the limits of the computing times scalability results.

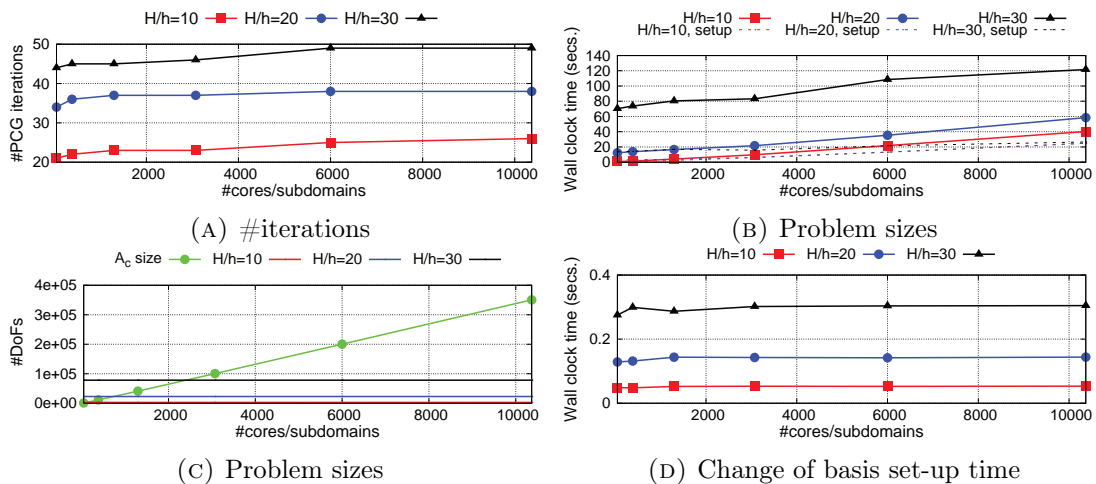


FIGURE 12. Weak scalability results for first order edge FE with a channel distribution of materials: $\alpha_{\text{white}} = 10^2$, $\beta_{\text{white}} = 1.0$ and $\alpha_{\text{black}} = 10^4$ and $\beta_{\text{black}} = 10^{-2}$.

5.4. Heterogeneous problems. In this section we study the scalability of the rPB-BDDC method for problems where the coefficients α, β are described by continuous (at least element-wise) functions, which contain high contrasts for their maximum and minimum values. In order to build the rPB-partition, the approach based on the aggregation of cells with their coefficients on the same interval Eq. (14) is used, see Sect. 4.4.

5.4.1. Periodic analytical functions. In this case, α and β are defined as exponential functions with a sinusoidal exponent such that the function is periodic on the domain and the number of peaks scales with the number of original subdomains in Θ , thus solving a harder problem as we increase the number of processors. In particular, let us consider $\log(\alpha) = \frac{c_{\text{max}}}{2} \sin(N_x \pi x)$ and $\log(\beta) = \frac{c_{\text{max}}}{2} \sin(N_y \pi y)$, where N_x, N_y denotes the number

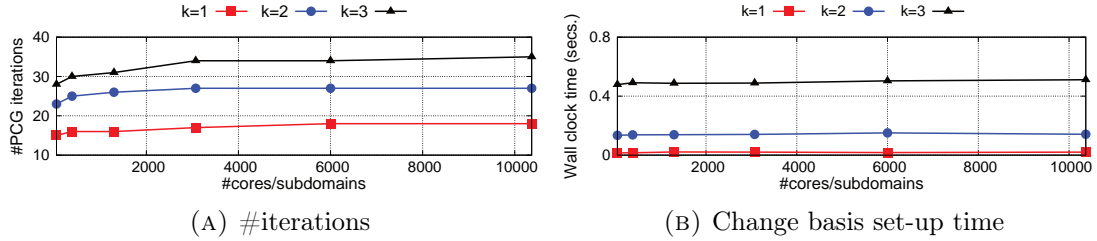


FIGURE 13. Weak scalability results for different order edge FE and $H/h = 4$ with a channel distribution of materials: $\alpha_{\text{white}} = 10^2$, $\beta_{\text{white}} = 1.0$ and $\alpha_{\text{black}} = 10^4$ and $\beta_{\text{black}} = 10^{-2}$.

of subdomains per (x, y) direction in a $P = N_x \times N_y \times N_z$ structured partition, (see β depicted in Fig. 3a for the case $c_{\max} = 6$, $N_x = N_y = 3$). Clearly, the maximum contrast within each coefficient is given by $r_{\max} = 10^{c_{\max}}$. We present weak scalability results up to 3072 subdomains for two different thresholds $r = \{r_{\max}, 10^3\}$, and for a local problem size of $\frac{H}{h} = 20$ in three different scenarios: coefficients α (Fig. 14), β (Fig. 15) or both are heterogeneous (Fig. 16), being set to $\alpha = 1.0$, $\beta = 1.0$ otherwise. Out of these plots, we can draw some conclusions. First, the case where only β is heterogeneous converges in a lower number of iterations compared to problems with heterogeneous α . Secondly, the consideration of lower values for r consequently results in larger coarse problems, but its size is only (approximately) doubled when only one coefficient is heterogeneous or (approximately) quadrupled when both are defined heterogeneous. In fact, in the range of subdomains considered in this experiment, the coarse problem computational times in all cases can be masked by local problem ones (Figs. 14b, 15b and 16b). Finally, and most salient, the rPB-BDDC method with $r = 10^3$ is weakly scalable in all cases with an excellent reduction in the number of iterations and computing times compared to the case where $r = r_{\max}$.

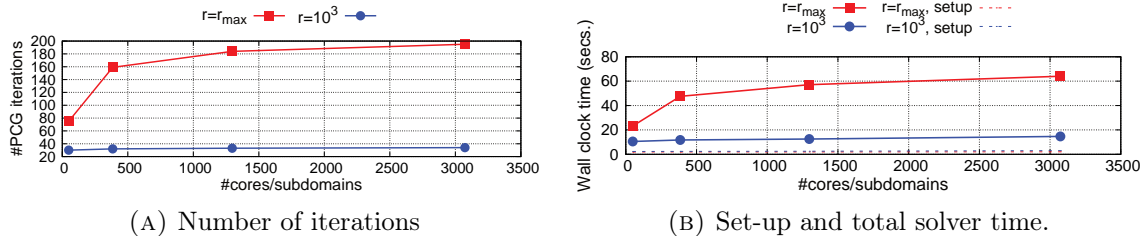


FIGURE 14. Weak scalability for the rPB-BDDC when only an heterogeneous α is considered, $\beta = 1.0$.

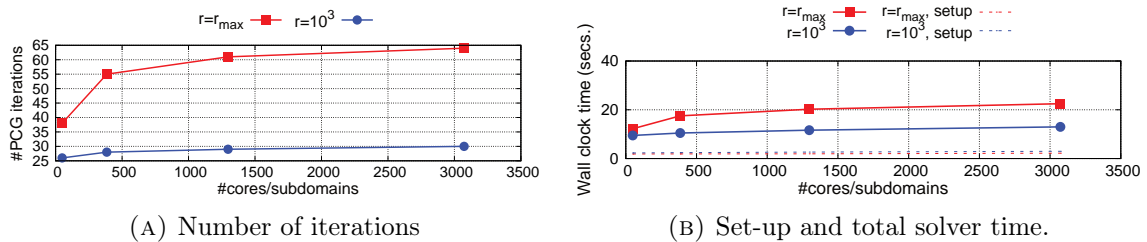


FIGURE 15. Weak scalability for the rPB-BDDC when only an heterogeneous β is considered, $\alpha = 1.0$.

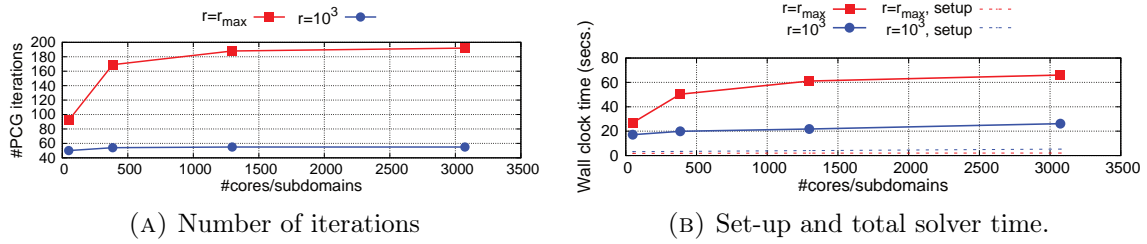


FIGURE 16. Weak scalability for the rPB-BDDC when both coefficients are heterogeneous.

5.4.2. *High Temperature Superconductors.* Next, we study the scalability of the algorithm with a practical application, the modelling of High Temperature Superconductors (HTS). The problem consists in the magnetization of a superconducting cube completely surrounded by a dielectric material (see Fig. 18a), subjected to an external AC magnetic field. The formulation Eq. (1) arises in the time-domain quasi-static approximation of the Maxwell's equations for solving the magnetic field, see [31] for details. Furthermore, the standard Backward Euler method is used to perform time integration over a time interval $[0, T]$, so let us define a time partition $\{0 = t^0, t^1, \dots, t^N = T\}$ into N time elements. Then, the form Eq. (1) can be used to compute the magnetic field for a particular time t^n , provided the solution on the previous time \mathbf{u}^{n-1} . The coefficient β is affected by the current time step size $\Delta t = (t^n - t^{n-1})$ as $\beta = \frac{\mu_0}{\Delta t}$, where $\mu_0 = 4\pi \cdot 10^{-7}$ is the magnetic permeability of the vacuum. While the dielectric material is modelled with a constant value for $\alpha = 10^{-3}$, the superconducting material behaviour is modelled with the stiff nonlinear dependence of the resistivity α with the solution as $\alpha = \alpha_0 \left(\frac{\|\nabla \times \mathbf{u}\|}{J_c} \right)^m$, with $m = 100$, $J_c = 10^{-8}$ and $\alpha_0 = 10^{-12}$. The equivalence with Eq. (1) is completed by considering the source term $\mathbf{f} = \beta \mathbf{u}^{n-1}$ and the strong imposition of an external magnetic field $\mathbf{u}^n \times \mathbf{n} = \mathbf{u}_0^n$ over the whole boundary. For the time step t^n , the weak form of the nonlinear problem reads: find $\mathbf{u}^n \in \widehat{X}_h$ such that

$$(15) \quad (\alpha(\mathbf{u}^n) \nabla \times \mathbf{u}^n, \nabla \times \mathbf{v}) + \beta(\mathbf{u}^n, \mathbf{v}) = \beta(\mathbf{u}^{n-1}, \mathbf{v}) \quad \forall \mathbf{v} \in \widehat{X}_h.$$

In order to derive the linearized form with Newton's method we consider the current approximation $\mathbf{u}^{n,k}$ and a (small) correction $\delta \mathbf{u}^{n,k}$ for the iterate k such that $\mathbf{u}^{n,k+1} = \mathbf{u}^{n,k} + \delta \mathbf{u}^{n,k}$. We plug the expression in Eq. (15), consider a first order Taylor expansion of $\alpha(\mathbf{u}^{n,k+1})$ around $\mathbf{u}^{n,k}$ and neglect the quadratic terms with respect to $\delta \mathbf{u}^{n,k}$, which yields the linearized problem: find $\delta \mathbf{u}^{n,k} \in \widehat{X}_h$ such that

$$(16) \quad \mathcal{J}(\mathbf{u}^{n,k}, \delta \mathbf{u}^{n,k}, \mathbf{v}) = -\mathcal{R}(\mathbf{u}^{n-1,k}, \mathbf{u}^{n,k}, \mathbf{v}) \quad \forall \mathbf{v} \in \widehat{X}_h,$$

where

$$(17a) \quad \begin{aligned} \mathcal{J}(\mathbf{u}^{n,k}, \delta \mathbf{u}^{n,k}, \mathbf{v}) &= (\alpha(\mathbf{u}^{n,k}) \nabla \times \delta \mathbf{u}^{n,k}, \nabla \times \mathbf{v}) + \beta(\delta \mathbf{u}^{n,k}, \mathbf{v}) + \\ &\quad + (\alpha'(\mathbf{u}^{n,k}) \delta \mathbf{u}^{n,k} \nabla \times \mathbf{u}^{n,k}, \nabla \times \mathbf{v}), \\ \mathcal{R}(\mathbf{u}^{n-1,k}, \mathbf{u}^{n,k}, \mathbf{v}) &= -\beta(\mathbf{u}^{n-1}, \mathbf{v}) + \beta(\mathbf{u}^{n,k}, \mathbf{v}) + (\alpha(\mathbf{u}^{n,k}) \nabla \times \mathbf{u}^{n,k}, \nabla \times \mathbf{v}). \end{aligned}$$

Therefore, the rPB-BDDC preconditioner is applied to the linearized problem Eq. (16) at every nonlinear iteration of every time step. We will focus on the performance of the linear solver, and the reader is directed to [31] for a detailed exposition of the composition of the used transient nonlinear solver.

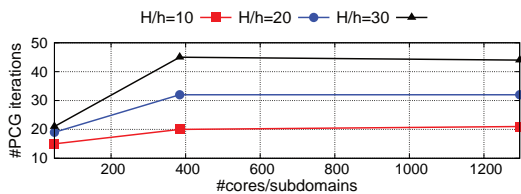
The problem is solved in $\Omega = [0, 40]^3 \text{ mm}^3$, composed by an outer dielectric Ω_{air} material which includes a concentric superconducting cube Ω_{hts} of size 10 mm such that

\mathbf{P}	# Average iters.	Average size(A_c)	size(A_c) ratio
49	16.8	196.7	$1.31c_0$
385	18.9	2424.5	$2.02c_0$
1297	21.7	7728.7	$1.91c_0$

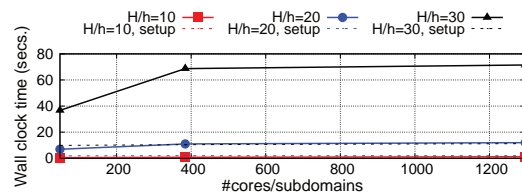
TABLE 3. Average metrics for the simulation of the time interval $T = [0, 5]$ ms. c_0 denotes the number of coarse DOFs of the original, geometrical partition.

$\Omega = \Omega_{\text{hts}} \cup \Omega_{\text{air}}$, see Fig. 18a. There is no source term and Dirichlet-type boundary conditions are imposed over the entire boundary as the time-dependent magnetic field $\mathbf{u}_{\text{ext}} = \frac{B_0}{\mu_0}[0, 0, \sin(2\pi\omega t)]$, where $B_0 = 200$ mT and $\omega = 50$ Hz. We solve the problem in the time interval $[0, 5]$ ms, which corresponds to a quarter of a full cycle in the applied \mathbf{u}_{ext} . Initial conditions are simply $\mathbf{u}^0 = \mathbf{0}$. The partition Θ_{pb}^r is obtained in all simulations for $r = 10^2$. The nonlinear scheme is stopped when the L_2 -norm of the nonlinear residual (Eq. (17a)) is below 10^{-4} , while the convergence criteria for the rPB-BDDC preconditioned linear solver is the reduction of the initial L_2 -norm of the residual of the linearized system by 10^{-8} .

We first present weak scalability results for the first set-up and solve with the rPB-BDDC preconditioner in Fig. 17, i.e., the first linearized problem (Eq. (16)) for the first time step. We include results for $\frac{H}{h} = \{10, 20, 30\}$. As expected, the method shows good weak scalability properties in number of iterations (see Fig. 17a) and computing times (see Fig. 17b).



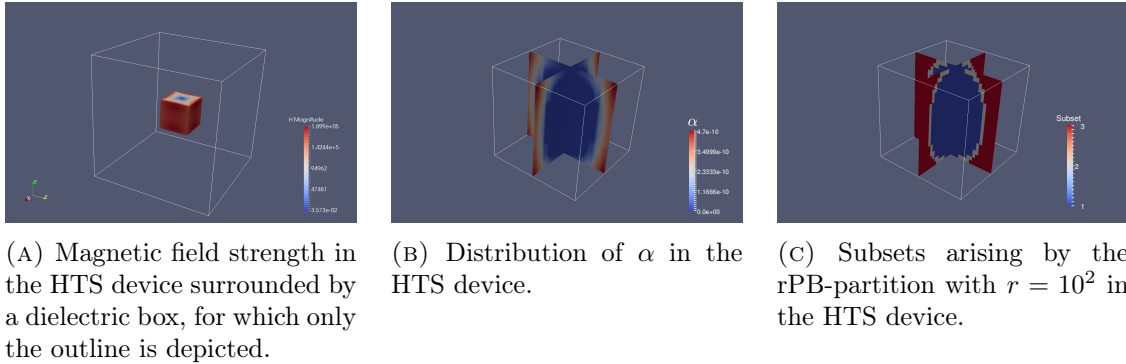
(A) Number of iterations.



(B) Wall clock time for the total solver and only the set-up phase.

FIGURE 17. Weak scalability for the first linear solver in the HTS problem with $r = 10^2$.

Next, we present average counters for the total number of linear solver applications for the simulation of the whole time interval $[0, 5]$ ms in Table 3, for a local problem size of $H/h = 10$ and different partitions. The resulting aggregation of cells into subsets based on their physical coefficient α (see Sect. 4.4) for $t = 4$ ms is depicted in Fig. 18c. We can identify two main regions in the distribution of α (see Fig. 18b): an inner region that is still not magnetized (i.e., with null resistivity) and a surrounding region, separated by a thin layer. Therefore, the selected value for r allows us to capture the behaviour of the different regions in Ω_{hts} . Out of the results in Tab. 3, the most salient property is the (asymptotic) scalability in the average number of iterations. Besides, we show how the coarse problem size for the presented cases is only (approximately) doubled regarding to the size that would be obtained with the partition Θ instead of the Θ_{pb} .

FIGURE 18. Domain and HTS device for $t = 4$ ms.

6. CONCLUSIONS

In this work, we have proposed an extension of the BDDC preconditioners for arbitrary order curl-conforming spaces that are robust for heterogeneous problems with high contrast of coefficients. The main idea is to enrich the continuity constraints enforced among subdomains (i.e., coarse DOFs) for those which contain high contrast of coefficients. The approach, which is shown to be robust for the grad-conforming case in [29], makes use of the knowledge about the physical coefficients to define a sub-partition of the original edge FE-based definition of coarse objects (edges and faces). The motivation for that is the well-known robustness of DD methods when there are only jumps of physical coefficients across the interface between subdomains. However, our case is more complex than the one in [29] for Poisson and elasticity problems, since two different coefficients are involved in the time-domain quasi-static approximation to the Maxwell's equations. Our solution is to add a perturbation term to the preconditioner so as we recover a scenario similar to the one in which only one coefficient jumps across the interface. A relaxed definition of the PB subdomains, where we only require that the maximal contrast of the two physical coefficients is smaller than a predefined thresholds, allows one to extend the range of applicability of the preconditioner to truly heterogeneous materials. Our preconditioners, which use the crucial change of variables in [13] to obtain weakly scalable algorithms for problems in $H(\text{curl})$ with few modifications to the standard BDDC algorithm in [1], are empirically shown to be robust with the contrast of coefficients. We would like to remark that our preconditioners maintain the simplicity of the standard BDDC and do not require to solve any eigenvalue or auxiliary problem.

We devoted a section to describe all the non-trivial implementation issues behind the method based on our experience through the implementation of the preconditioners in FEMPAR. Its task-overlapping implementation of the PB-BDDC preconditioner allows one to mask the computing times for the coarse problem, as long as they do not exceed local solvers time. With such implementation, we have been able to provide notable weak scalability results in the application of our new preconditioners to a wide range of multi-material and heterogeneous electromagnetics problems, including realistic 3D problems where coefficients can be defined by arbitrary functions, even dependent on the solution itself. In the future, the multilevel extension of the algorithm is expected to push forward the limits of its scalability properties.

REFERENCES

- [1] C. R. Dohrmann. A preconditioner for substructuring based on constrained energy minimization. *SIAM Journal on Scientific Computing*, 25(1):246–258, 2003.
- [2] A. Toselli and O. Widlund. *Domain Decomposition Methods: Algorithms and Theory*. Springer-Verlag, Berlin, 2005.

- [3] J. Mandel. Balancing domain decomposition. *Communications in Numerical Methods in Engineering*, 9(3):233–241, 1993.
- [4] S. Badia, A. Martín, and J. Principe. FEMPAR Web page. <http://www.fempar.org>, 2018.
- [5] S. Badia, A. F. Martín, and J. Principe. FEMPAR: An object-oriented parallel finite element framework. *Archives of Computational Methods in Engineering*, 25(2):195–271, 2018.
- [6] S. Badia, A. F. Martín, and J. Principe. Multilevel balancing domain decomposition at extreme scales. *SIAM Journal on Scientific Computing*, pages C22–C52, 2016.
- [7] S. Zampini. PCBDDC: A class of robust dual-primal methods in PETSc. *SIAM J. Sci. Comput.*, 38:S282–S306, 2016.
- [8] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, K. Rupp, B. F. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc Web page. <http://www.mcs.anl.gov/petsc>, 2018.
- [9] C. Farhat, M. Lesoinne, and K. Pierson. A scalable dual-primal domain decomposition method. *Numerical Linear Algebra with Applications*, 7(7-8):687–714, 2000.
- [10] J. Mandel, C. R. Dohrmann, and R. Tezaur. An algebraic theory for primal and dual substructuring methods by constraints. *Applied Numerical Mathematics*, 54(2):167–193, July 2005.
- [11] J. Li and O. B. Widlund. FETI-DP, BDDC, and block Cholesky methods. *International Journal for Numerical Methods in Engineering*, 66(2):250–271, 2006.
- [12] S. C. Brenner and L.-Y. Sung. BDDC and FETI-DP without matrices or vectors. *Computer Methods in Applied Mechanics and Engineering*, 196(8):1429–1435, January 2007.
- [13] A. Toselli. Dual-primal FETI algorithms for edge finite-element approximations in 3D. *IMA Journal of Numerical Analysis*, 26(1):96–130, 2006.
- [14] M. Olm, S. Badia, and A. F. Martín. On a general implementation of h- and p-adaptive curl-conforming finite elements. *Advances in Engineering Software*, 2019.
- [15] S. Zampini, P. Vassilevski, V. Dobrev, and T. Kolev. Balancing domain decomposition by constraints algorithms for curl-conforming spaces of arbitrary order. In P. E. Bjørstad, S. C. Brenner, L. Halpern, H. H. Kim, R. Kornhuber, T. Rahman, and O. B. Widlund, editors, *Domain Decomposition Methods in Science and Engineering XXIV*, pages 103–116. Springer International Publishing, 2017.
- [16] C. R. Dohrmann and O. B. Widlund. Some recent tools and a BDDC algorithm for 3D problems in $H(\text{curl})$. In R. Bank, M. Holst, O. Widlund, and J. Xu, editors, *Domain Decomposition Methods in Science and Engineering XX*, pages 15–25, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [17] C. Pechstein and C. R. Dohrmann. A unified framework for adaptive BDDC. *Electronic Transactions on Numerical Analysis*, 46:273–336, 2017.
- [18] J. Mandel and B. Sousedík. Adaptive selection of face coarse degrees of freedom in the BDDC and the FETI-DP iterative substructuring methods. *Computer Methods in Applied Mechanics and Engineering*, 196(8):1389–1399, January 2007.
- [19] B. Sousedík, J. Šístek, and J. Mandel. Adaptive-Multilevel BDDC and its parallel implementation. *Computing*, 95(12):1087–1119, December 2013.
- [20] J. Calvo and O. Widlund. An adaptive choice of primal constraints for BDDC domain decomposition algorithms. *Electronic Transactions on Numerical Analysis*, 45:524–544, 2016.
- [21] N. Spillane, V. Dolean, P. Hauret, F. Nataf, C. Pechstein, and R. Scheichl. A robust two-level domain decomposition preconditioner for systems of PDEs. *Comptes Rendus Mathématique*, 349(23):1255 – 1259, 2011.
- [22] N. Spillane, V. Dolean, P. Hauret, F. Nataf, and D. J. Rixen. Solving generalized eigenvalue problems on the interfaces to build a robust two-level FETI method. *Comptes Rendus Mathématique*, 351(5):197 – 201, 2013.
- [23] A. Klawonn, P. Radtke, and O. Rheinbach. Adaptive coarse spaces for BDDC with a transformation of basis. In T. Dickopf, M. J. Gander, L. Halpern, R. Krause, and L. F. Pavarino, editors, *Domain Decomposition Methods in Science and Engineering XXII*, pages 301–309. Springer International Publishing, 2016.
- [24] C. Pechstein and R. Scheichl. Analysis of FETI methods for multiscale PDEs. Part ii: interface variation. *Numerische Mathematik*, 118(3):485–529, Jul 2011.
- [25] C. R. Dohrmann and O. B. Widlund. A BDDC algorithm with deluxe scaling for three-dimensional $H(\text{curl})$ problems. *Communications on Pure and Applied Mathematics*, 69(4):745–770, 2016.
- [26] D. Oh, O. Widlund, Z. Zampini, and C. Dohrmann. BDDC algorithms with deluxe scaling and adaptive selection of primal constraints for Raviart-Thomas vector fields. *Math. Comp.*, 87:659–692, 2018.
- [27] S. Zampini. Adaptive BDDC deluxe methods for $H(\text{curl})$. In C.-O. Lee, X.-C. Cai, D. E. Keyes, H. H. Kim, A. Klawonn, E.-J. Park, and O. B. Widlund, editors, *Domain Decomposition Methods in Science and Engineering XXIII*, pages 285–292. Springer International Publishing, 2017.

- [28] S. Zampini and X. Tu. Multilevel balancing domain decomposition by constraints deluxe algorithms with adaptive coarse spaces for flow in porous media. *SIAM Journal on Scientific Computing*, 39(4):A1389–A1415, 2017.
- [29] S. Badia, A. F. Martín, and H. Nguyen. Physics-Based Balancing Domain Decomposition by Constraints for Multi-Material Problems. *Journal of Scientific Computing*, pages 1–30, 2018.
- [30] A. Klawonn, O. Widlund, and M. Dryja. Dual-primal FETI methods for three-dimensional elliptic problems with heterogeneous coefficients. *SIAM Journal on Numerical Analysis*, 40(1):159–179, 2002.
- [31] M. Olm, S. Badia, and A. F. Martín. Simulation of high temperature superconductors and experimental validation. *Computer Physics Communications*, 237:154–167, 2019.
- [32] P. Monk. *Finite Element Methods for Maxwell's Equations*. Oxford Science Publications, 2003.
- [33] S. Badia, A. F. Martín, and J. Principe. Implementation and scalability analysis of balancing domain decomposition methods. *Archives of Computational Methods in Engineering*, 20(3):239–262, September 2013.
- [34] S. Badia, A. F. Martín, and J. Principe. A highly scalable parallel implementation of balancing domain decomposition by constraints. *SIAM Journal on Scientific Computing*, 36(2):C190–C218, January 2014.
- [35] Marenostrium IV website. <https://www.bsc.es/marenostrium/marenostrium>, 2018.