Open access • Journal Article • DOI:10.1137/19M1266964

A Multilevel Schwarz Preconditioner Based on a Hierarchy of Robust Coarse Spaces — Source link

Hussam Al Daas, Laura Grigori, Pierre Jolivet, Pierre-Henri Tournier

Institutions: Centre national de la recherche scientifique

Published on: 25 May 2021 - SIAM Journal on Scientific Computing (Society for Industrial and Applied Mathematics)

Topics: Domain decomposition methods, Preconditioner, Multigrid method, Iterative method and Condition number

A Multilevel Schwarz Preconditioner Based on a Hierarchy of Robust Coarse Spaces

Hussam Al Daas, Laura Grigori, Pierre Jolivet, Pierre-Henri Tournier

A MULTILEVEL SCHWARZ PRECONDITIONER BASED ON A HIERARCHY OF ROBUST COARSE SPACES*

HUSSAM AL DAAS†, LAURA GRIGORI†, PIERRE JOLIVET‡, AND PIERRE-HENRI TOURNIER§

Abstract. In this paper we present a multilevel preconditioner based on overlapping Schwarz methods for symmetric positive definite (SPD) matrices. Robust two-level Schwarz preconditioners exist in the literature to guarantee fast convergence of Krylov methods. As long as the dimension of the coarse space is reasonable, that is, exact solvers can be used efficiently, two-level methods scale well on parallel architectures. However, the factorization of the coarse space matrix may become costly at scale. An alternative is then to use an iterative method on the second level, combined with an algebraic preconditioner, such as a one-level additive Schwarz preconditioner. Nevertheless, the condition number of the resulting preconditioned coarse space matrix may still be large. One of the difficulties of using more advanced methods, like algebraic multigrid or even two-level overlapping Schwarz methods, to solve the coarse problem is that the matrix does not arise from a partial differential equation (PDE) anymore. We introduce in this paper a robust multilevel additive Schwarz preconditioner where at each level the condition number is bounded, ensuring a fast convergence for each nested solver. Furthermore, our construction does not require any additional information than for building a two-level method, and may thus be seen as an algebraic extension.

Key words. domain decomposition, multilevel, elliptic problems, subspace correction

AMS subject classifications. 65F08, 65F10, 65N55

1. Introduction.

We consider the solution of a linear system of equations

$$Ax = b, \tag{1.1}$$

where $A \in \mathbb{R}^{n\times n}$ is a symmetric positive definite (SPD) matrix, $b \in \mathbb{R}^n$ is the right-hand side, and $x \in \mathbb{R}^n$ is the vector of unknowns. To enhance convergence, it is common to solve the preconditioned system

$$M^{-1}Ax = M^{-1}b.$$

Standard domain decomposition preconditioners such as block Jacobi, additive Schwarz, and restricted additive Schwarz methods are widely used [32, 9, 8]. In a parallel framework, such preconditioners have the advantage of relatively low communication costs. However, their role in lowering the condition number of the system typically deteriorates when the number of subdomains increases. Multilevel approaches have shown a large impact on enhancing the convergence of Krylov methods [33, 12, 7, 25, 20, 10, 21, 1, 15, 23, 34, 30]. In multigrid and domain decomposition communities, multilevel methods have proven their capacity of scaling up to large numbers of processors and tackling ill-conditioned systems [37, 4, 19]. While some preconditioners are purely algebraic [7, 20, 10, 26, 29, 16, 1], several multilevel methods are based on hierarchical meshing in both multigrid and domain decomposition communities [35, 9, 25, 15, 23]. Mesh coarsening depends on the geometry of the problem. One has to be careful when choosing a hierarchical structure since it can

*Submitted to the editors June 7, 2019.

†ALPINES, INRIA, Paris, France (aldaas.hussam@gmail.com, laura.grigori@inria.fr).

‡IRIT, CNRS, Toulouse, France (pierre.jolivet@enseeiht.fr).

§LJLL, CNRS, Paris, France (tournier@ljll.math.upmc.fr).

have a significant impact on the iteration count [23, 25]. In [23], the authors propose a multilevel Schwarz domain decomposition solver for the elasticity problem. Based on a heuristic approach and following the maximum independent set method [2], they coarsen the fine mesh while preserving the boundary in order to obtain a two-level method. This strategy is repeated recursively to build several levels. However, they do not provide a bound on the condition number of the preconditioned matrix of the multilevel method. Multilevel domain decomposition methods are mostly based on non-overlapping approaches [35, 9, 25, 23, 37, 4, 30, 34]. Two-level overlapping domain decomposition methods are well studied and provide robust convergence estimates [33, 12, 5]. However, extending such a construction to more than two levels while preserving robustness is not straightforward. In [6], the authors propose an algebraic multilevel additive Schwarz method. Their approach is inspired by algebraic multigrid strategies. One drawback of it is that it is sensitive to the number of subdomains. In [15], the authors suggest applying the two-level Generalized Dryja–Smith–Widlund preconditioner recursively to build a multilevel method. In this case, the condition number bound of the two-level approach depends on the width of the overlap, the diameter of discretization elements, and the diameter of the subdomains. They focus on the preconditioner for the three-level case. One drawback of their approach is that the three-level preconditioner requires more iterations than the two-level variant. In this paper, the only information from the PDE needed for the construction of the preconditioner consists of the local Neumann matrices at the fine level. These matrices correspond to the integration of the bilinear form in the weak formulation of the studied PDE on the subdomain-decomposed input mesh. No further information is necessary: except on the fine level, our method is algebraic and does not depend on any coarsened mesh or auxiliary discretized operator. For problems not arising from PDE discretization, one needs to supply the local SPSD matrices on the finest level. In [3], a subset of the authors propose a fully algebraic approximation for such matrices. However, their approximation strategy is heuristic and may not be effective in some cases.

Our preconditioner is based on a hierarchy of coarse spaces and is defined as following. At the first level, the set of unknowns is partitioned into N_1 subdomains and each subdomain has an associated matrix $A_{1,j} = R_{1,j} A R_{1,j}^\top$ obtained by using appropriate restriction and prolongation operators $R_{1,j}$ and $R_{1,j}^\top$ respectively, defined in the following section. The preconditioner is formed as an additive Schwarz preconditioner coupled with an additive coarse space correction, defined as,

$$M^{-1} = M_1^{-1} = V_1 A_2^{-1} V_1^\top + \sum_{j=1}^{N_1} R_{1,j}^\top A_{1,j}^{-1} R_{1,j},$$

where V_1 is a tall-and-skinny matrix spanning a coarse space obtained by solving for each subdomain $j = 1$ to N_1 a generalized eigenvalue problem involving the matrix $A_{1,j}$ and the Neumann matrix associated with subdomain j. The coarse space matrix is $A_2 = V_1^\top A V_1$. This is equivalent to the GenEO preconditioner, and is described in detail in [33] and recalled briefly in section 2. The dimension of the coarse space is proportional to the number of subdomains N_1. When it increases, factorizing A_2 by using a direct method becomes prohibitive, and hence the application of A_2^{-1} to a vector should also be performed through an iterative method.

Our multilevel approach defines a hierarchy of coarse spaces V_i and coarse space matrices A_i for $i = 2$ to any depth $L+1$, and defines a preconditioner M_i^{-1} such that the condition number of $M_i^{-1} A_i$ is bounded. The depth $L+1$ is chosen such that the coarse space matrix A_{L+1} can be factorized efficiently by using a direct method. At each level i, the graph of the coarse space matrix A_i is partitioned into N_i subdomains, and each subdomain j is associated with a local matrix $A_{i,j} = R_{i,j} A_i R_{i,j}^\top$ obtained by using appropriate restriction and prolongation operators $R_{i,j}$ and $R_{i,j}^\top$, respectively. The preconditioner at level i is defined as,

$$M_i^{-1} = V_i A_{i+1}^{-1} V_i^\top + \sum_{j=1}^{N_i} R_{i,j}^\top A_{i,j}^{-1} R_{i,j},$$

where the coarse space matrix is $A_{i+1} = V_i^\top A_i V_i$.

One of the main contributions of the paper concerns the construction of the hierarchy of coarse spaces V_i for levels i going from 2 to L, that are built algebraically from the coarse space of the previous level V_{i-1}. This construction is based on the definition of local symmetric positive semi-definite (SPSD) matrices associated with each subdomain j at each level i that we introduce in this paper. These matrices are obtained by using the local SPSD matrices of the previous level $i-1$ and the previous coarse space V_{i-1}. They are then involved, with the local matrices $A_{i,j}$, in concurrent generalized eigenvalue problems solved for each subdomain j that allows to compute the local eigenvectors contributing to the coarse space V_i.

We show in Theorem 5.3, section 5, that the condition number of $M_i^{-1} A_i$ is bounded and depends on the maximum number of subdomains at the first level that share an unknown, the number of distinct colors required to color the graph of A_i so that $\left\{span\{R_{i,j}^\top\}\right\}_{1 \leqslant j \leqslant N_i}$ of the same color are mutually A_i-orthogonal, and a user defined tolerance τ. It is thus independent of the number of subdomains N_i.

The main contribution of this paper is based on the combination of two previous works on two-level additive Schwarz methods [3, 33]. The coarse space proposed in [33] guarantees an upper bound on the condition number that can be prescribed by the user. The SPSD splitting in the context of domain decomposition presented in [3] provides an algebraic view for the construction of coarse spaces. The combination of these two works leads to a robust multilevel additive Schwarz method. Here, robustness refers to the fact that at each level, an upper bound on the condition number of the associated matrix can be prescribed by the user a priori. The rest of the paper is organized as follows. In the next section, we present the notations used throughout the paper. In section 2, we present a brief review of the theory of one- and two-level additive Schwarz methods. We extend in section 3 the class of SPSD splitting matrices presented in [3] in order to make it suitable for multilevel methods. Afterwards, we define the coarse space at level i based on the extended class of local SPSD splitting matrices associated with this level. Section 4 describes the partitioning of the domain at level $i+1$ from the partitioning at level i. In Section 5, we explain the computation of the local SPSD matrices associated with each subdomain at level $i+1$. We compute them using those associated with subdomains at level i. Section 6 presents numerical experiments on highly challenging diffusion and linear elasticity problems in two- and three-dimensional problems. We illustrate the theoretical robustness and practical usage of our proposed method by performing strong scalability tests up to 8,192 processes.

Context and notation. By convention, the finest level, on which (1.1) is defined, is the first level. A subscript index is used in order to specify which level an entity is defined on. In the case where additional subscripts are used, the first subscript always denotes the level. For the sake of clarity, we omit the subscript corresponding to level 1 when it is clear from context, e.g., matrix A. Furthermore, the subscripts i and j always refer to a specific level i and its subdomain j, respectively. The number of levels is $L+1$. Let $A_i \in \mathbb{R}^{n_i \times n_i}$ denote symmetric positive definite matrices, each corresponding to level $i = 1, \ldots, L+1$. We suppose that a direct solver can be used at level $L+1$ to compute an exact factorization of A_{L+1}.

Let $B \in \mathbb{R}^{p \times q}$ be a matrix. Let $P \subset [\![1;p]\!]$ and $Q \subset [\![1;q]\!]$ be two sets of indices. The concatenation of P and Q is represented by $[P,Q]$. We note that the order of the concatenation is important. $B(P,:)$ is the submatrix of B formed by the rows whose indices belong to P. $B(:,Q)$ is the submatrix of B formed by the columns whose indices belong to Q. $B(P,Q) = (B(P,:))(:,Q)$. The identity matrix of size p is denoted I_p. We suppose that the graph of A_i is partitioned into N_i non-overlapping subdomains, where $N_i \ll n_i$ and $N_{i+1} \leqslant N_i$ for $i = 1, \ldots, L$. We note that partitioning at level 1 can be performed by using a graph partitioning library such as ParMETIS [22] or PT-SCOTCH [11]. Partitioning at greater levels will be described later in section 4. In the following, we define for each level $i = 1, \ldots, L$ notations for subsets and restriction operators that are associated with the partitioning. Let $\Omega_i = [\![1;n_i]\!]$ be the set of unknowns at level i and let $\Omega_{i,j,I}$ for $j = 1, \ldots, N_i$ be the subset of Ω_i that represents the unknowns in subdomain j. We refer to $\Omega_{i,j,I}$ as the *interior unknowns* of subdomain j. Let $\Gamma_{i,j}$ for $j = 1, \ldots, N_i$ be the subset of Ω_i that represents the neighbor unknowns of subdomain j, i.e., the unknowns at distance 1 from subdomain j through the graph of A_i. We refer to $\Gamma_{i,j}$ as the *overlapping unknowns* of subdomain j. We denote $\Omega_{i,j} = [\Omega_{i,j,I},\ \Gamma_{i,j}]$, for $j = 1, \ldots, N_i$, the concatenation of interior and overlapping unknowns of subdomain j. We denote $\Delta_{i,j}$, for $j = 1, \ldots, N_i$, the complementary of $\Omega_{i,j}$ in Ω_i, i.e., $\Delta_{i,j} = \Omega_i \setminus \Omega_{i,j}$. In Figure 1.1, a triangular mesh is used to discretize a square domain. The set of nodes of the mesh is partitioned into 16 disjoint subsets $\Omega_{1,j,I}$, which represent a non-overlapping decomposition, for $j = 1, \ldots, 16$ (left). On the left, a matrix A_1 whose connectivity graph corresponds to the mesh is illustrated. The submatrix $A_1(\Omega_{1,j,I}, \Omega_{1,j,I})$ is associated with the non-overlapping subdomain j. Each submatrix $A_1(\Omega_{1,j,I}, \Omega_{1,j,I})$ is colored with a distinct color. The same color is used to color the region that contains the nodes in the non-overlapping subdomain $\Omega_{1,j,I}$. Note that if two subdomains j_1, j_2 are neighbors, the submatrix $A_1(\Omega_{1,j_1,I}, \Omega_{1,j_2,I})$ has nonzero elements. For $j = 1, \ldots, N_i$, we denote by $n_{i,j,I}$, $\gamma_{i,j}$ and $n_{i,j}$ the cardinality of $\Omega_{i,j,I}$, $\Gamma_{i,j}$ and $\Omega_{i,j}$ respectively.

Let $R_{i,j,I} \in \mathbb{R}^{n_{i,j,I} \times n_i}$ be defined as $R_{i,j,I} = I_{n_i}(\Omega_{i,j,I},:)$.
Let $R_{i,j,\Gamma} \in \mathbb{R}^{\gamma_{i,j} \times n_i}$ be defined as $R_{i,j,\Gamma} = I_{n_i}(\Gamma_{i,j},:)$.
Let $R_{i,j} \in \mathbb{R}^{n_{i,j} \times n_i}$ be defined as $R_{i,j} = I_{n_i}(\Omega_{i,j},:)$.
Let $R_{i,j,\Delta} \in \mathbb{R}^{(n_i - n_{i,j}) \times n_i}$ be defined as $R_{i,j,\Delta} = I_{n_i}(\Delta_{i,j},:)$.
Let $\mathcal{P}_{i,j} = I_{n_i}([\Omega_{i,j,I}, \Gamma_{i,j}, \Delta_{i,j}],:) \in \mathbb{R}^{n_i \times n_i}$, be a permutation matrix associated with the subdomain j, for $j = 1, \ldots, N_i$. The matrix of the overlapping subdomain j, $R_{i,j} A_i R_{i,j}^\top$, is denoted $A_{i,j}$. We denote $D_{i,j} \in \mathbb{R}^{n_{i,j}, \times n_{i,j}}$, $j = 1, \ldots, N_i$, any set of non-negative diagonal matrices such that

$$I_{n_i} = \sum_{j=1}^{N_i} R_{i,j}^\top D_{i,j} R_{i,j}.$$

We refer to $\{D_{i,j}\}_{1 \leqslant j \leqslant N_i}$ as the algebraic partition of unity. Let $V_i \in \mathbb{R}^{n_i \times n_{i+1}}$ be a tall-and-skinny matrix of full rank. We denote $\mathcal{S}_i$ the subspace spanned by the columns of V_i. This subspace will stand for the coarse space associated with level i. By convention, we refer to $\mathcal{S}_i$ as subdomain 0 at level i. Thus, we have $n_{i,0} = n_{i+1}$.

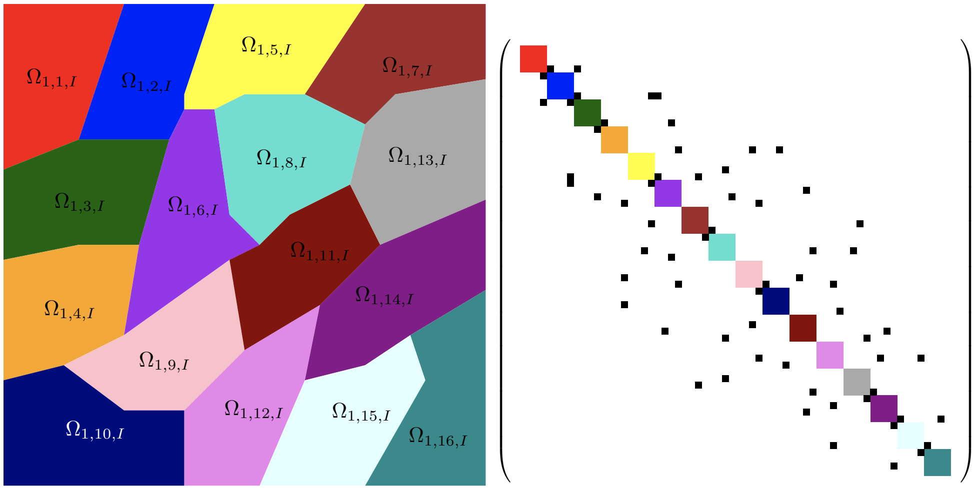

FIG. 1.1. *Left: a triangular mesh is used to discretize the unit square. The set of nodes of the mesh is partitioned into 16 disjoint subsets, non-overlapping subdomains, $\Omega_{1,j,I}$ for $j = 1, \ldots, 16$. Right: Illustration of the matrix A_1 whose connectivity graph corresponds to the mesh on the left. The diagonal block j of A_1 corresponds to the non-overlapping subdomain $\Omega_{1,j,I}$. Each submatrix $A_1(\Omega_{1,j,I}, \Omega_{1,j,I})$ is colored with a distinct color. The same color is used to color the region of the square that contains nodes in $\Omega_{1,j,I}$.*

The interpolation operator at level i is defined as:

$$\begin{aligned} \mathscr{R}_{i,2} \colon \prod_{j=0}^{N_i} \mathbb{R}^{n_{i,j}} &\to \mathbb{R}^{n_i} \\ (u_j)_{0 \leqslant j \leqslant N_i} &\mapsto \sum_{j=0}^{N_i} R_{i,j}^\top u_j. \end{aligned} \tag{1.2}$$

Finally, we denote $\mathcal{V}_{i,j}$ the set of neighboring subdomains of each subdomain j at level i for $(i, j) \in [\![1; L]\!] \times [\![1; N_i]\!]$.

$$\mathcal{V}_{i,j} = \{k \in [\![1; N_i]\!] : \Omega_{i,j} \cap \Omega_{i,k} \neq \emptyset\}.$$

As previously mentioned, partitioning at level 1 can be performed by graph partitioning libraries such as ParMETIS [22] or PT-SCOTCH [11]. Partitioning at further levels will be defined later: the sets $\Omega_{i,j,I}$, $\Omega_{i,j,\Gamma}$, $\Omega_{i,j}$, and $\Delta_{i,j}$ for $i > 1$ are defined in subsection 4.2. The coarse spaces $\mathcal{S}_i$ as well as the projection and prolongation operators $V_i^\top$ and V_i are defined in subsection 3.2. We suppose that the connectivity graph between the subdomains on each level is sparse. This assumption is not true in general, however, it is valid in structures based on locally constructed coarse spaces in domain decomposition as we show in this paper, see [18, Section 4.1 p.81] for the case of two levels.

2. Background. In this section, we review briefly several theoretical results related to additive Schwarz preconditioners. We introduce them for the sake of completeness.

LEMMA 2.1 (fictitious subspace lemma). *Let $A \in \mathbb{R}^{n_A \times n_A}$, $B \in \mathbb{R}^{n_B \times n_B}$ be two* *symmetric positive definite matrices. Let $\mathscr{R}$ be an operator defined as*

$$\begin{aligned}\mathscr{R}\colon \mathbb{R}^{n_B} &\to \mathbb{R}^{n_A}\\ v &\mapsto \mathscr{R}v,\end{aligned}$$

and let $\mathscr{R}^\top$ be its transpose. Suppose that the following conditions hold:

1. *The operator $\mathscr{R}$ is surjective.*
2. *There exists $c_u > 0$ such that*
$$(\mathscr{R}v)^\top A\,(\mathscr{R}v) \leqslant c_u v^\top B v, \quad \forall v \in \mathbb{R}^{n_B}.$$
3. *There exists $c_l > 0$ such that for all $v_{n_A} \in \mathbb{R}^{n_A}, \exists v_{n_B} \in \mathbb{R}^{n_B} | v_{n_A} = \mathscr{R}v_{n_B}$ and*
$$c_l v_{n_B}^\top B v_{n_B} \leqslant (\mathscr{R}v_{n_B})^\top A\,(\mathscr{R}v_{n_B}) = v_{n_A}^\top A v_{n_A}.$$

Then, the spectrum of the operator $\mathscr{R}B^{-1}\mathscr{R}^\top A$ is contained in the segment $[c_l, c_u]$.

Proof. We refer the reader to [12, Lemma 7.4 p.164] or [28, 27, 13] for a detailed proof. ∎

LEMMA 2.2. *The operator $\mathscr{R}_{i,2}$ as defined in (1.2) is surjective.*

Proof. The proof follows from the definition of $\mathscr{R}_{i,2}$ (1.2). ∎

LEMMA 2.3. *Let $k_{i,c}$ for $i = 1, \ldots, L$ be the minimum number of distinct colors so that $\left\{span\{R_{i,j}^\top\}\right\}_{1\leqslant j\leqslant N_i}$ of the same color are mutually A_i-orthogonal. Then, we have*

$$\begin{aligned}&(\mathscr{R}_{i,2}u_{\mathcal{B}_i})^\top A_i\,(\mathscr{R}_{i,2}u_{\mathcal{B}_i})\\ &\qquad\leqslant (k_{i,c}+1)\sum_{j=0}^{N_i} u_j^\top \left(R_{i,j}A_iR_{i,j}^\top\right)u_j, \quad \forall u_{\mathcal{B}_i} = (u_j)_{0\leqslant j\leqslant N_i} \in \prod_{j=0}^{N} \mathbb{R}^{n_{i,j}}.\end{aligned}$$

Proof. We refer the reader to [9, Theorem 12 p.93] for a detailed proof. ∎

We note that at level i, the number $k_{i,c}$ is smaller than the maximum number of neighbors over the set of subdomains $[\![1; N_i]\!]$

$$k_{i,c} \leqslant \max_{1\leqslant j\leqslant N_i} \#\mathcal{V}_{i,j}.$$

Due to the sparse structure of the connectivity graph between the subdomains at level i, the maximum number of neighbors over the set of subdomains $[\![1; N_i]\!]$ is independent of the number of subdomains N_i. Then, so is $k_{i,c}$.

LEMMA 2.4. *Let $u_{A_i} \in \mathbb{R}^{n_{A_i}}$ and $u_{\mathcal{B}_i} = \{u_j\}_{0\leqslant j\leqslant N_i} \in \prod_{j=0}^{N_i} \mathbb{R}^{n_{i,j}}$ such that $u_{A_i} = \mathscr{R}_{i,2}u_{\mathcal{B}_i}$. The additive Schwarz operator without any other restriction on the coarse space $\mathcal{S}_i$ verifies the following inequality*

$$\sum_{j=0}^{N_i} u_j^\top \left(R_{i,j}A_iR_{i,j}^\top\right)u_j \leqslant 2u_{A_i}^\top A_i u_{A_i} + (2k_{i,c}+1)\sum_{j=1}^{N_i} u_j^\top R_{i,j}A_iR_{i,j}^\top u_j,$$

where $k_{i,c}$ is defined in Lemma 2.3.

Proof. We refer the reader to [12, Lemma 7.12, p. 175] to view the proof in detail. ∎

LEMMA 2.5. *Let $A, B \in \mathbb{R}^{m\times m}$ be two symmetric positive semi-definite matrices. Let $ker(A)$, $range(A)$ denote the null space and the range of A respectively. Let P_0 be an orthogonal projection on $range(A)$. Let τ be a positive real number. Consider the generalized eigenvalue problem,*

$$P_0 B P_0 u_k = \lambda_k A u_k,$$
$$(u_k, \lambda_k) \in range(A) \times \mathbb{R}.$$

Let P_τ be an orthogonal projection on the subspace

$$Z = ker(A) \oplus span\{u_k | \lambda_k > \tau\},$$

then, the following inequality holds:

$$(u - P_\tau u)^\top B (u - P_\tau u) \leqslant \tau u^\top A u, \quad \forall u \in \mathbb{R}^m. \tag{2.1}$$

Proof. We refer the reader to [3, Lemma 2.4] and [12, Lemma 7.7] for a detailed proof. ∎

2.1. GenEO coarse space. In [33, 12] the authors present the GenEO coarse space which relies on defining appropriate symmetric positive semi-definite (SPSD) matrices $\tilde{A}_j \in \mathbb{R}^{n\times n}$ for $j = 1, \dots, N$. These are the unassembled Neumann matrices, corresponding to the integration on each subdomain of the operator defined in the variational form of the PDE. These matrices are local, i.e., $R_{j,\Delta}\tilde{A}_j = 0$. Furthermore, they verify the relations

$$u^\top \tilde{A}_j u \leqslant u^\top A u, \quad \forall u \in \mathbb{R}^n,$$
$$u^\top \sum_{j=1}^{N} \tilde{A}_j u \leqslant k_{\text{GenEO}} u^\top A u, \quad \forall u \in \mathbb{R}^n,$$

where $k_{\text{GenEO}} \leqslant N$ is the maximum number of subdomains that share an unknown.

2.2. Local SPSD splitting of an SPD matrix. In [3], the authors present the local SPSD splitting of an SPD matrix. Given the permutation matrix $\mathcal{P}_j$, a local SPSD splitting matrix $\tilde{A}_j$ of A associated with subdomain j is defined as

$$\mathcal{P}_j \tilde{A}_j \mathcal{P}_j^\top = \begin{pmatrix} R_{j,I} A R_{j,I}^\top & R_{j,I} A R_{j,\Gamma}^\top & \\ R_{j,\Gamma} A R_{j,I}^\top & \tilde{A}_\Gamma^j & \\ & & 0 \end{pmatrix}, \tag{2.2}$$

where $\tilde{A}_\Gamma^j \in \mathbb{R}^{\gamma_j \times \gamma_j}$ satisfies the two following conditions: For all $u \in \mathbb{R}^{\gamma_j}$,

- $u^\top \left(R_{j,\Gamma} A R_{j,I}^\top\right) \left(R_{j,I} A R_{j,I}^\top\right)^{-1} \left(R_{j,I} A R_{j,\Gamma}^\top\right) u \leqslant u^\top \tilde{A}_\gamma^j u$
- $u^\top \tilde{A}_\Gamma^j u \leqslant u^\top \left(\left(R_{j,\Gamma} A R_{j,\Gamma}^\top\right) - \left(R_{j,\Gamma} A R_{j,\Delta}^\top\right) \left(R_{j,\Delta} A R_{j,\Delta}^\top\right)^{-1} \left(R_{j,\Delta} A R_{j,\Gamma}^\top\right)\right) u.$

The authors prove that the matrices $\tilde{A}_j$ defined in such a way verify the following relations:

$$R_{j,\Delta}\tilde{A}_j = 0, \tag{2.3}$$
$$u^\top \tilde{A}_j u \leq u^\top A u, \quad \forall u \in \mathbb{R}^n, \tag{2.4}$$
$$u^\top \sum_{j=1}^{N} \tilde{A}_j u \leqslant k u^\top A u, \quad \forall u \in \mathbb{R}^n, \tag{2.5}$$

where k is a number that depends on the local SPSD splitting matrices and can be at most equal to the number of subdomains $k \leqslant N$. The authors also show that the local matrices defined in GenEO [33, 12] can be seen as a local SPSD splitting.

In [3], the authors highlight that the key idea to construct a coarse space relies on the ability to identify the so-called local SPSD splitting matrices. They present a class of algebraically constructed coarse spaces based on the local SPSD splitting matrices. Moreover, this class can be extended to a larger variety of local SPSD matrices. This extension has the advantage of allowing to construct efficient coarse spaces for a multilevel structure in a practical way. This is discussed in the following section.

3. Extension of the class of coarse spaces.

In this section we extend the class of coarse spaces presented in [3]. To do so, we present a class of matrices, that is larger than the class of local SPSD splitting matrices. This will be our main building block in the construction of efficient coarse spaces. Furthermore, this extension can lead to a straightforward construction of hierarchical coarse spaces in a multilevel Schwarz preconditioner setting.

3.1. Extension of the class of local SPSD splitting matrices.

Regarding the two-level additive Schwarz method, the authors of [3] introduced the local SPSD splitting related to a subdomain as defined in (2.2). As it can be seen from the theory presented in that paper, it is not necessary to have the exact matrices $R_{j,I}AR_{j,I}^\top$, $R_{j,I}AR_{j,\Gamma}^\top$, and $R_{j,\Gamma}AR_{j,I}^\top$ in the definition of the local SPSD splitting in order to build an efficient coarse space. Indeed, the one and only necessary condition is to define for each subdomain j an SPSD matrix $\tilde{A}_j$ for $j = 1, \ldots, N$ such that:

$$
\begin{aligned}
&R_{j,\Delta}\tilde{A}_j = 0,\\
&u^\top \sum_{j=1}^{N} \tilde{A}_j u \leqslant k u^\top A u, \forall u \in \mathbb{R}^n,
\end{aligned} \tag{3.1}
$$

where k is a number that depends on the local SPSD matrices $\tilde{A}_j$ for $j = 1, \ldots, N$. The first condition means that $\tilde{A}_j$ has the local SPSD structure associated with subdomain j, i.e., it has the following form:

$$
\mathcal{P}_j \tilde{A}_j \mathcal{P}_j^\top = \begin{pmatrix} \tilde{A}_{I,\Gamma}^j & 0 \\ 0 & 0 \end{pmatrix},
$$

where $\tilde{A}_{I,\Gamma}^j \in \mathbb{R}^{n_j \times n_j}$. The second condition is associated with the stable decomposition property [36, 12]. Note that with regard to the local SPSD matrices, the authors in [33] only use these two conditions. That is to say, with matrices that verify conditions (3.1) the construction of the coarse space is straightforward through the theory presented in either [33] or [3]. To this end, we define in the following the local SPSD (LSPSD) matrix associated with subdomain j as well as the associated local filtering subspace that contributes to the coarse space.

Definition 3.1 (local SPSD matrices). *An SPSD matrix $\tilde{A}_{i,j} \in \mathbb{R}^{n_i \times n_i}$ is called local SPSD (LSPSD) with respect to subdomain j if*
- $R_{i,j,\Delta}\tilde{A}_{i,j} = 0$,
- $u^\top \sum_{j=1}^{N_i} \tilde{A}_{i,j} u \leqslant k_i u^\top A_i u$,

where $k_i > 0$.

We note that the local SPSD splitting matrices form a subset of the local SPSD matrices.

3.2. Multilevel coarse spaces. This section summarizes the steps to be performed in order to construct the coarse space at level i once we have the LSPSD matrices associated with each subdomain at that level.

DEFINITION 3.2 (coarse space based on LSPSD matrices). *Let $\tilde{A}_{i,j} \in \mathbb{R}^{n_i \times n_i}$ for $j = 1, \ldots, N_i$ be LSPSD matrices. Let $D_{i,j} \in \mathbb{R}^{n_{i,j}}$ for $j = 1, \ldots, N_i$ be the partition of unity. Let $\tau_i > 0$ be a given number. For a subdomain $j \in [\![1; N_i]\!]$, let*

$$G_{i,j} = D_{i,j} \left(R_{i,j} A_i R_{i,j}^\top \right) D_{i,j}.$$

Let $\tilde{P}_{i,j}$ be the projection on range$(R_{i,j} \tilde{A}_j R_{i,j}^\top)$ parallel to ker$(R_{i,j} \tilde{A}_j R_{i,j}^\top)$. Let $K_{i,j} =$ ker$(R_{i,j} \tilde{A}_{i,j} R_{i,j}^\top)$. Consider the generalized eigenvalue problem:

$$\begin{aligned} &\tilde{P}_{i,j} G_{i,j} \tilde{P}_{i,j} u_{i,j,k} = \lambda_{i,j,k} R_{i,j} \tilde{A}_{i,j} R_{i,j}^\top u_{i,j,k}, \\ &(u_{i,j,k}, \lambda_{i,j,k}) \in \mathit{range}(R_{i,j} \tilde{A}_{i,j} R_{i,j}^\top) \times \mathbb{R}. \end{aligned} \tag{3.2}$$

Set

$$Z_{i,j} = K_{i,j} \oplus \mathit{span} \left\{ u_{i,j,k} | \lambda_{i,j,k} > \tau_i \right\}. \tag{3.3}$$

Then, the coarse space associated with LSPSD matrices $\tilde{A}_{i,j}$ for $j = 1, \ldots, N_i$ at level i is defined as:

$$\mathcal{S}_i = \bigoplus_{j=1}^{N_i} R_{i,j}^\top D_{i,j} Z_{i,j}. \tag{3.4}$$

Following notations from section 1, the columns of V_i span the coarse space $\mathcal{S}_i$. The matrix A_{i+1} is defined as:

$$A_{i+1} = V_i^\top A_i V_i. \tag{3.5}$$

The local SPSD splitting matrices at level 1 will play an important role in the construction of the LSPSD matrices at subsequent levels. In the following, we present an efficient approach for computing LSPSD matrices for levels greater than 1.

4. Partitioning for levels strictly greater than 1. In this section, we explain how to obtain the partitioning sets $\Omega_{i,j,I}$ for $(i,j) \in [\![2; L]\!] \times [\![1; N_i]\!]$. Once the sets $\Omega_{i,j,I}$ for $j = 1, \ldots, N_i$ are defined at level i, the following elements are readily available: sets $\Gamma_{i,j}, \Delta_{i,j}$, and $\Omega_{i,j}$; restriction operators $R_{i,j,I}, R_{i,j,\Gamma}, R_{i,j,\Delta}$, and $R_{i,j}$; permutation matrices $\mathcal{P}_{i,j}$ for $j = 1, \ldots, N_i$. The partition of unity is constructed in an algebraic way. The mth diagonal element of $D_{i,j}$ is 1 if $m \leqslant n_{i,j,I}$ and 0 otherwise.

4.1. Superdomains as unions of several subdomains. In this section, we introduce the notion of a *superdomain*. It refers to the union of several neighboring subdomains. Let $\mathcal{G}_{i,1}, \ldots, \mathcal{G}_{i,N_{i+1}}$ be disjoint subsets of $[\![1; N_i]\!]$, where $\bigcup_{j=1}^{N_{i+1}} \mathcal{G}_{i,j} = [\![1; N_i]\!]$. We call the union of the subdomains $\{k \in [\![1; N_i]\!] : k \in \mathcal{G}_{i,j}\}$ superdomain j, for $j = 1, \ldots, N_{i+1}$. Figure 4.1 gives an example of how to set superdomains. Though this definition of superdomains may look somehow related to the fine mesh, it is in practice done at the algebraic level, as explained later on. Note that the indices of columns and rows of A_{i+1} are associated with the vectors contributed by the subdomains at level i in order to build the coarse space $\mathcal{S}_i$, see Figure 4.2. Hence, defining subdomains on the structure of A_{i+1} is natural once we have the subsets $\mathcal{G}_{i,j}$, for $j = 1, \ldots, N_{i+1}$.

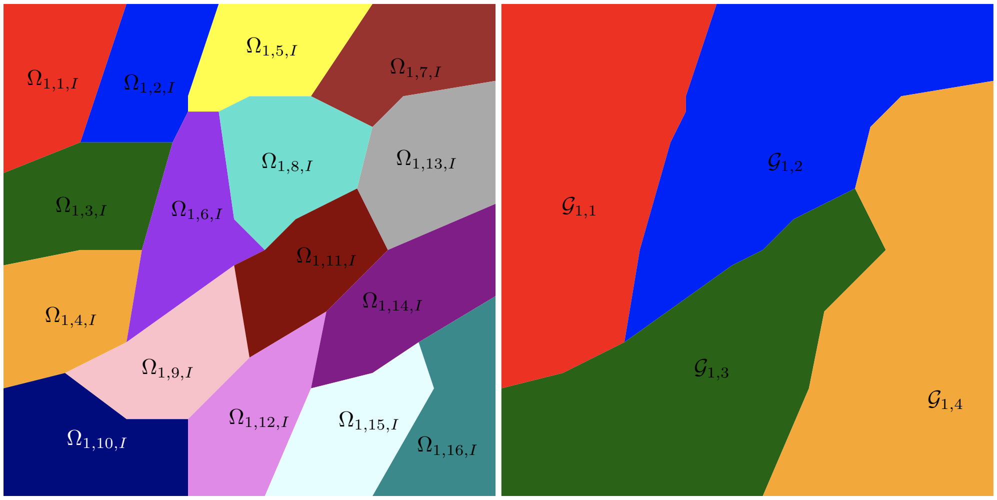

FIG. 4.1. *Left: 16 subdomains at level 1. Right: 4 superdomains at level 1.* $\mathcal{G}_{1,j} = [\![4(j-1)+1; 4(j-1)+4]\!]$.

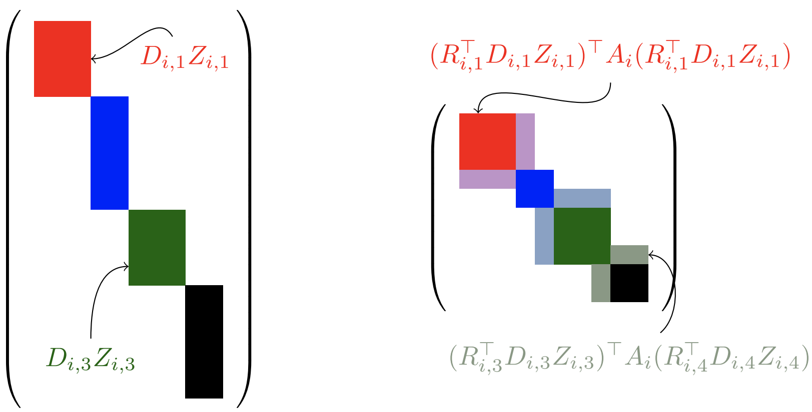

FIG. 4.2. *Illustration of the correspondence of indices between the columns of V_i (left) and the rows and columns of A_{i+1} (right). Having no overlap in V_i is possible through a non-overlapping partition of unity.*

4.2. Heritage from superdomains. Let $e_{i,j}$ be the set of indices of the vectors that span $R_{i,j}^\top D_{i,j} Z_{i,j}$ in the matrix V_i for some $(i,j) \in [\![1; L-1]\!] \times [\![1; N_i]\!]$, see Figure 4.2. We define $\Omega_{i+1,j,I} = \cup_{k\in\mathcal{G}_{i,j}} e_{i,k}$, for $j = 1, \ldots, N_{i+1}$. We denote $\Omega_{i+1,j,\Gamma}$ the subset of $[\![1; n_{i+1}]\!] \setminus \Omega_{i+1,j,I}$ whose elements are at distance 1 from $\Omega_{i+1,j,I}$ through the graph of A_{i+1}. We note that

$$\Omega_{i+1,j,\Gamma} \subset \bigcup_{p\in\mathcal{G}_{i,j}} \bigcup_{k\in\mathcal{V}_{i,p}} e_{i,k},$$

where $\mathcal{V}_{i,j}$ represents the set of subdomains that are neighbors of subdomain j at level i for $j = 1, \ldots, N_i$. The overlapping subdomain j is defined by the set $\Omega_{i+1,j} = [\Omega_{i+1,j,I}, \Omega_{i+1,j,\Gamma}]$. The rest of the sets, restriction, and prolongation operators can be defined as given in section 1.

5. LSPSD matrices for levels strictly greater than 1. In [33, 12, 3], different methods are suggested to obtain local SPSD splitting matrices at level 1. These matrices are used to construct efficient two-level additive Schwarz preconditioners. Here in this section, we do not discuss the construction of these matrices at level 1. We suppose that we have the local SPSD matrices $\tilde{A}_{1,j} \in \mathbb{R}^{n_1 \times n_1}$ for $j = 1, \dots, N_1$. We focus on computing LSPSD matrices $\tilde{A}_{i,j} \in \mathbb{R}^{n_i \times n_i}$ for $(i, j) \in [\![2; L]\!] \times [\![1; N_i]\!]$. We also suppose that the coarse space $\mathcal{S}_1$ is available, i.e., the matrices V_1 and $A_2 = V_1^\top A_1 V_1$ are known explicitly.

PROPOSITION 5.1. *Let i be a fixed level index, and let $\tilde{A}_{i,j}$ be an LSPSD of A_i, (see Definition 3.1), associated with subdomain j, for $j = 1, \dots, N_i$. Let $\mathcal{G}_{i,1}, \dots, \mathcal{G}_{i,N_{i+1}}$ be a set of superdomains at level i associated with the partitioning at level $i+1$, see subsection 4.1. Let $V_i^\top$ be the restriction matrix to the coarse space at level i. Then, the matrix $\tilde{A}_{i+1,j}$ which is defined as:*

$$\tilde{A}_{i+1,j} = \sum_{k \in \mathcal{G}_{i,j}} V_i^\top \tilde{A}_{i,k} V_i,$$

satisfies the conditions in Definition 3.1. That is, $\tilde{A}_{i+1,j}$ is LSPSD of A_{i+1} with respect to subdomain j for $j = 1, \dots, N_{i+1}$.

Proof. To prove that $\tilde{A}_{i+1,j}$ is LSPSD of A_{i+1} with respect to subdomain j, we have to prove the following:

- $R_{i+1,j,\Delta} \tilde{A}_{i+1,j} = 0$
- $u^\top \sum_{j=1}^{N_{i+1}} \tilde{A}_{i+1,j} u \leqslant k_{i+1} u^\top A_{i+1} u$ for all $u \in \mathbb{R}^{n_{i+1}}$.

First, note that $R_{i,k} \tilde{A}_{i,j} = 0$ for all non-neighboring subdomains k of subdomain j. This yields $Z_{i,k}^\top D_{i,k} R_{i,k} \tilde{A}_{i,j} = 0$ for these subdomains k.

Now, let $m \in [\![1; n_{i+1}]\!] \setminus \Omega_{i+1,j}$. We will show that the mth row of $\tilde{A}_{i+1,j}$ is zero. Following the partitioning of subdomains at level $i+1$, there exists a subdomain Ω_{p_0} such that the mth column of V_i is part of $R_{i,p_0}^\top D_{i,p_0} Z_{i,p_0}$. We denote this column vector by v_m. Furthermore, the subdomain p_0 is not a neighbor of any subdomain that is a part of the superdomain $\mathcal{G}_{i,j}$. Hence, $v_m^\top \tilde{A}_{i,k} = 0$ for $k \in \mathcal{G}_{i,j}$. The mth row of $\tilde{A}_{i+1,j}$ is given as $v_m^\top \sum_{k \in \mathcal{G}_{i,j}} \tilde{A}_{i,k} V_i$. Then, $v_m^\top \sum_{k \in \mathcal{G}_{i,j}} \tilde{A}_{i,k} = 0$, and the mth row of $\tilde{A}_{i+1,j}$ is zero.

To prove the second condition, we have

$$u^\top \sum_{j=1}^{N_{i+1}} \tilde{A}_{i+1,j} u = u^\top \sum_{j=1}^{N_{i+1}} \sum_{k \in \mathcal{G}_{i,j}} V_i^\top \tilde{A}_{i,k} V_i u.$$

Since $\{\mathcal{G}_{i,j}\}_{1 \leqslant j \leqslant N_{i+1}}$ form a disjoint partitioning of $[\![1; N_i]\!]$, we can write

$$\begin{aligned} u^\top \sum_{j=1}^{N_{i+1}} \tilde{A}_{i+1,j} u &= u^\top \sum_{k=1}^{N_i} V_i^\top \tilde{A}_{i,k} V_i u, \\ &= u^\top V_i^\top \sum_{k=1}^{N_i} \tilde{A}_{i,k} V_i u. \end{aligned}$$

$\tilde{A}_{i,k}$ is an LSPSD matrix of A_i for $k = 1, \ldots, N_i$. Hence, we have

$$u^\top \sum_{j=1}^{N_{i+1}} \tilde{A}_{i+1,j} u \leqslant k_i u^\top V_i^\top A_i V_i u,$$
$$\leqslant k_i u^\top A_{i+1} u.$$

We finish the proof by setting $k_{i+1} = k_i$. ∎

Figure 5.1 gives an illustration of the LSPSD construction provided by Proposition 5.1. Figure 5.1 (top left) represents the matrix A_1. The graph of A_1 is partitioned into 16 subdomains. Each subdomain is represented by a different color. Figure 5.1 (top right) represents the matrix V_1 whose column vectors form a basis of the coarse space $\mathcal{S}_1$. Colors of columns of V_1 correspond to those of subdomains in A_1. Figure 5.1 (bottom left) represents the matrix $A_2 = V_1^\top A_1 V_1$. Note that column and row indices of A_2 are associated with column indices of V_1. Four subdomains are used at level 2. The partitioning at level 2 is related to the superdomain $\mathcal{G}_{1,j} = [\![4(j-1)+1; 4(j-1)+4]\!]$ for $j = 1, \ldots, 4$. Figure 5.1 (bottom right) represents an LSPSD matrix of A_2 with respect to subdomain 1 at level 2.

Theorem 5.2 shows that the third condition of the fictitious subspace lemma Lemma 2.1 holds at level i for $i = 1, \ldots, L$.

Theorem 5.2. *Let $\tilde{A}_{i,j}$ be an LSPSD of A_i associated with subdomain j, for $(i,j) \in [\![1; L]\!] \times [\![1; N_i]\!]$. Let $\tau_i > 0$, $Z_{i,j}$ be the subspace associated with $\tilde{A}_{i,j}$, and $P_{i,j}$ be the projection on $Z_{i,j}$ as defined in Lemma 2.5. Let $u_i \in \mathbb{R}^{n_i}$ and let $u_{i,j} = \left(D_{i,j}\left(I_{n_{i,j}} - P_{i,j}\right) R_{i,j} u_i\right)$ for $(i,j) \in [\![1; L]\!] \times [\![1; N_i]\!]$. Let $u_{i,0}$ be defined as,*

$$u_{i,0} = \left(V_i^\top V_i\right)^{-1} V_i^\top \left(\sum_{j=1}^{N_i} R_{i,j}^\top D_{i,j} P_{i,j} R_{i,j} u_i\right).$$

Let $m_i = (2 + (2k_{i,c} + 1)k_i \tau_i)^{-1}$. Then,

$$u_i = \sum_{j=0}^{N_i} R_{i,j}^\top u_{i,j},$$

and

$$m_i \sum_{j=0}^{N_i} u_{i,j}^\top R_{i,j} A_i R_{i,j}^\top u_{i,j} \leqslant u_i^\top A_i u_i. \tag{5.1}$$

Proof. We have

$$\sum_{j=0}^{N_i} R_{i,j}^\top u_{i,j} = V_i \left(V_i^\top V_i\right)^{-1} V_i^\top \left(\sum_{j=1}^{N_i} R_{i,j}^\top D_{i,j} P_{i,j} R_{i,j} u_i\right) + \sum_{j=1}^{N_i} R_{i,j}^\top u_{i,j}$$

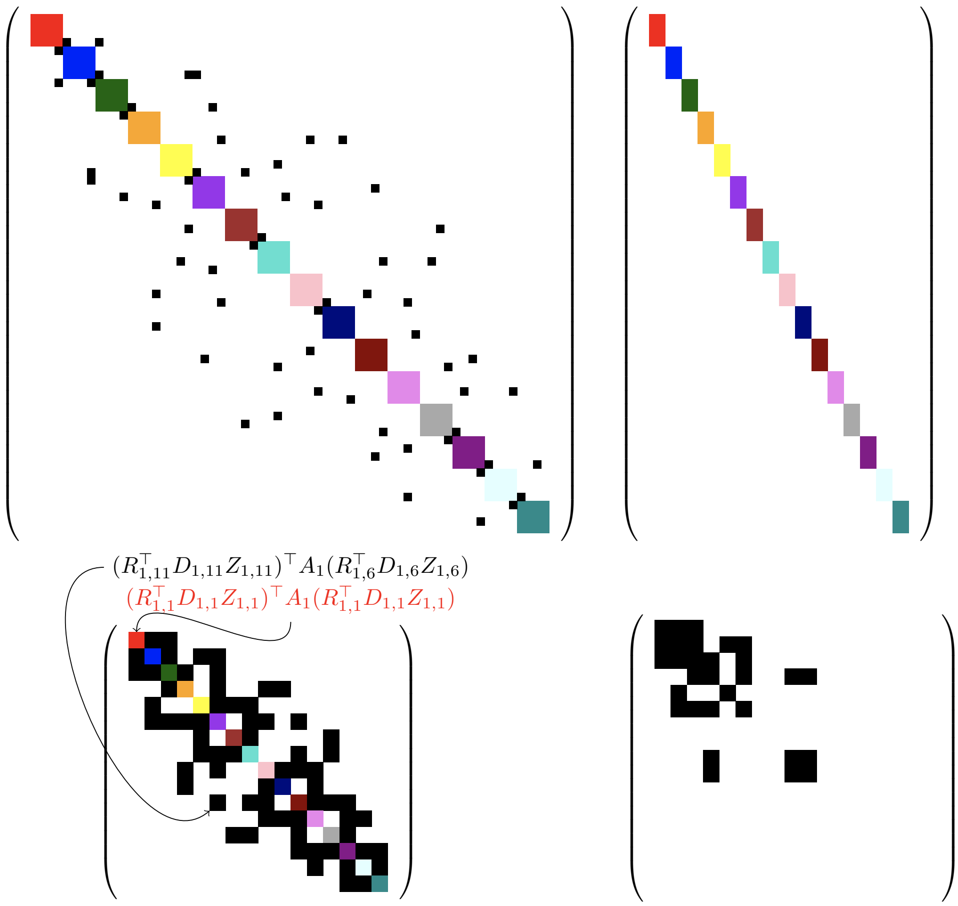

FIG. 5.1. *Illustration of the LSPSD construction provided by Proposition 5.1. Top left: the matrix* A_1*, top right:* V_1*, bottom left: the matrix* $A_2 = V_1^\top A_1 V_1$*, bottom right:* $\tilde{A}_{2,1} = \sum_{j \in \mathcal{G}_{1,1}} V_1^\top \tilde{A}_{1,j} V_1$*, where* $\mathcal{G}_{1,1} = 1, \ldots, 4$

Since for all $y \in \mathcal{S}_i$, $V_i \left(V_i^\top V_i\right)^{-1} V_i^\top y = y$, we have

$$\begin{aligned}
\sum_{j=0}^{N_i} R_{i,j}^\top u_{i,j} &= \sum_{j=1}^{N_i} R_{i,j}^\top D_{i,j} P_{i,j} R_{i,j} u_i + \sum_{j=1}^{N_i} R_{i,j}^\top \left(D_{i,j} \left(I_{n_{i,j}} - P_{i,j}\right) R_{i,j} u_i\right), \\
&= \sum_{j=1}^{N_i} R_{i,j}^\top D_{i,j} R_{i,j} u_i, \\
&= u_i.
\end{aligned}$$

To prove the inequality (5.1), we start with the inequality from Lemma 2.4. We have

$$\sum_{j=0}^{N_i} u_{i,j}^\top R_{i,j} A_i R_{i,j}^\top u_{i,j} \leqslant 2 u_i^\top A_i u_i + (2k_{i,c} + 1) \sum_{j=1}^{N_i} u_{i,j}^\top R_{i,j} A_i R_{i,j}^\top u_{i,j}, \tag{5.2}$$

where we chose $u_{\mathcal{B}_i}$ in Lemma 2.4 to be $(u_{i,j})_{j=0,\ldots,N_i}$ and $u_{A_i} = u_i$. In Definition 3.2, we defined $Z_{i,j}$, such that for all $w \in \mathbb{R}^{n_{i,j}}$ we have

$$\left((I_{n_{i,j}} - P_{i,j})w\right)^\top \left(D_{i,j} R_{i,j} A_i R_{i,j}^\top D_{i,j}\right) \left((I_{n_{i,j}} - P_{i,j})w\right) \leqslant \tau_i w^\top \left(R_{i,j} \tilde{A}_{i,j} R_{i,j}^\top\right) w.$$

Hence, in the special case $w = R_{i,j} u_i$, we can write

$$\begin{aligned}\left((I_{n_{i,j}} - P_{i,j})R_{i,j}u_i\right)^\top &\left(D_{i,j} R_{i,j} A_i R_{i,j}^\top D_{i,j}\right) \left((I_{n_{i,j}} - P_{i,j})R_{i,j}u_i\right)\\ &\leqslant \tau_i (R_{i,j}u_i)^\top \left(R_{i,j} \tilde{A}_{i,j} R_{i,j}^\top\right) (R_{i,j}u_i).\end{aligned}$$

Equivalently,

$$u_{i,j}^\top R_{i,j} A_i R_{i,j}^\top u_{i,j} \leqslant \tau_i (R_{i,j}u_i)^\top R_{i,j} \tilde{A}_{i,j} R_{i,j}^\top (R_{i,j}u_i).$$

Plugging this inequality in (5.2) gives

$$\sum_{j=0}^{N_i} u_{i,j}^\top R_{i,j} A_i R_{i,j}^\top u_{i,j} \leqslant 2u_i^\top A_i u_i + (2k_{i,c}+1)\,\tau_i \sum_{j=1}^{N_i} (R_{i,j}u_i)^\top R_{i,j} \tilde{A}_{i,j} R_{i,j}^\top (R_{i,j}u_i).$$

Since $\tilde{A}_{i,j}$ is local, we have

$$(R_{i,j}u_i)^\top R_{i,j} \tilde{A}_{i,j} R_{i,j}^\top (R_{i,j}u_i) = u_i^\top \tilde{A}_{i,j} u_i, \text{ for } j = 1, \ldots, N_i.$$

By using the fact that $\tilde{A}_{i,j}$ is LSPSD of A_i for $j = 1, \ldots, N_i$, we obtain the following:

$$\sum_{j=0}^{N_i} u_{i,j}^\top R_{i,j} A_i R_{i,j}^\top u_{i,j} \leqslant 2u_i^\top A_i u_i + (2k_{i,c}+1)\, k_i \tau_i u_i^\top A_i u_i.$$

Multiplying both sides with m_i ends the proof, i.e.,

$$m_i \sum_{j=0}^{N_i} u_{i,j}^\top R_{i,j} A_i R_{i,j}^\top u_{i,j} \leqslant u_i^\top A_i u_i. \qquad \square$$

In [3], the authors presented the minimal subspace that replaces $Z_{i,j}$ (defined in (3.3) and used in Theorem 5.2) that is required to prove Theorem 5.2. The main difference with respect to the subspace that we define in (3.3) is that it is not necessary to include the entire kernel of the LSPSD matrix, $K_{i,j}$, in $Z_{i,j}$, see Definition 3.2. Nevertheless, in this work, we include the entire kernel of the LSPSD matrix in the definition of $Z_{i,j}$. This allows us to ensure that the kernels of Neumann matrices are transferred across the levels, see Theorem 5.4. And in addition, this corresponds to the definition used in GenEO [12, Lemma 7.7] and to its implementation in the HPDDM library [19].

Theorem 5.3 provides an upper bound on the condition number of the preconditioned matrix $M_i^{-1} A_i$ for $i = 1, \ldots, L$.

THEOREM 5.3. *Let M_i be the additive Schwarz preconditioner at level i combined with the coarse space correction induced by $\mathcal{S}_i$ defined in (3.4). The following inequality holds,*

$$\kappa\left(M_i^{-1} A_i\right) \leqslant (k_{i,c}+1)\left(2 + (2k_{i,c}+1) k_i \tau_i\right).$$

Proof. Lemma 2.2, Lemma 2.3, and Theorem 5.2 prove that the multilevel preconditioner verifies the conditions in Lemma 2.1 at each level i. Hence, the spectrum of the preconditioned matrix $M_i^{-1}A_i$ is contained in the interval $[(2+(2k_{i,c}+1)k_i\tau_i)^{-1}, k_{i,c}+1]$. Equivalently, the condition number of the preconditioned matrix at level i verifies the following inequality

$$\kappa\left(M_i^{-1}A_i\right) \leqslant (k_{i,c}+1)\left(2+(2k_{i,c}+1)k_i\tau_i\right). \qquad \square$$

Proposition 5.1 shows that the constant k_i associated with the LSPSD matrices at level i is independent of the number of levels and bounded by the number of subdomains at level 1. Indeed,

$$k_1 \geq k_i \text{ for } i = 2, \dots, L.$$

Furthermore, in the case where the LSPSD matrices at the first level are the Neumann matrices, k_i is bounded by the maximum number of subdomains at level 1 that share an unknown.

The constant $k_{i,c}$ for $i = 1, \dots, L$ is the minimum number of distinct colors so that $\left\{span\{R_{i,j}^\top\}\right\}_{1\leqslant j\leqslant N_i}$ of the same color are mutually A_i-orthogonal. Both constants k_i and $k_{i,c}$ are independent of the number of subdomains for each level i.

The constant τ_i can be chosen such that the condition number of the preconditioned system at level i is upper bounded by a prescribed value. Hence, this allows to have a robust convergence of the preconditioned Krylov solver at each level.

Algorithm 5.1 presents the construction of the multilevel additive Schwarz method by using GenEO. The algorithm iterates over the levels. At each level, three main operations are performed. First, the construction of the LSPSD matrices. At level 1, the LSPSD matrices are the Neumann matrices, otherwise, Proposition 5.1 is used to compute them. Once the LSPSD matrix is available, the generalized eigenvalue problem in (3.2) has to be solved concurrently. Given the prescribed upper bound on the condition number, $Z_{i,j}$ can be set. Finally, the coarse space is available and the coarse matrix is assembled.

The following Theorem 5.4, describes how the kernel of Neumann matrices are transferred across the levels.

THEOREM 5.4. *Suppose that $\tilde{A}_{1,j}$ is the Neumann matrix associated with the subdomain $\Omega_{1,j}$ for $j \in [\![1; N_1]\!]$. For $(i,j) \in [\![2; L]\!] \times [\![1; N_i]\!]$, let*

- *$\tilde{A}_{i,j}$ be the LSPSD matrices associated with $A_{i,j}$ defined in Proposition 5.1,*
- *$\mathcal{G}_{i-1,j}$ be the corresponding superdomains,*
- *$\mathcal{G}_{i-1,j}^1$ be the union of subdomains at level 1 which contribute hierarchically to obtain $\mathcal{G}_{i-1,j}$,*
- *$\tilde{A}_{\mathcal{G}_{i-1,j}}$ be the Neumann matrix associated with $\mathcal{G}_{i-1,j}^1$ (seeing $\mathcal{G}_{i-1,j}^1$ as a subdomain),*
- *$A_{\mathcal{G}_{i-1,j}}$ be the restriction of A to the subdomain $\mathcal{G}_{i-1,j}^1$.*

Then, the kernel of $\tilde{A}_{\mathcal{G}_{i-1,j}}$ is included in the kernel of $\left(\prod_{l=1}^{i-1} V_l\right)\tilde{A}_{i,j}\left(\prod_{l=1}^{i-1} V_l\right)^\top$.

Proof. First, note that for any LSPSD matrix computed as in Proposition 5.1, we have

$$\left(\prod_{l=1}^{i-1} V_l\right)\tilde{A}_{i,j}\left(\prod_{l=1}^{i-1} V_l\right)^\top = \left(\prod_{l=1}^{i-1} V_l\right)\left(\prod_{l=1}^{i-1} V_l\right)^\top \sum_{k\in\mathcal{G}_{i,j}^1}\tilde{A}_{1,k}\left(\prod_{l=1}^{i-1} V_l\right)\left(\prod_{l=1}^{i-1} V_l\right)^\top.$$

Algorithm 5.1 Multilevel GenEO

Require: $A_1 = A \in \mathbb{R}^{n\times n}$ SPD, $L+1$ number of levels, N_i number of subdomains at each level, $\mathcal{G}_{i,j}$ sets of superdomains

Ensure: preconditioner at each level i, M_i^{-1} with bounded condition number of $M_i^{-1}A_i$

```
for i = 1, ..., L do
  for each subdomain j = 1, ..., N_i do
    A_{i,j} = R_{i,j} A_i R_{i,j}^⊤ (local matrix associated with subdomain j)
    if i = 1 then
      local SPSD Ã_{i,j} is Neumann matrix of subdomain j
    else
      compute local SPSD matrix as
```

$$\tilde{A}_{i,j} = \sum_{k\in\mathcal{G}_{i,j}} V_{i-1}^\top \tilde{A}_{i-1,k} V_{i-1}$$

```
    end if
    solve the generalized eigenvalue problem (3.2), set Z_{i,j} as in (3.3)
  end for
```

$\mathcal{S}_i = \bigoplus_{j=1}^{N_i} D_{i,j} R_{i,j}^\top Z_{i,j}$, V_i basis of $\mathcal{S}_i$

coarse matrix $A_{i+1} = V_i^\top A_i V_i$, $A_{i+1} \in \mathbb{R}^{n_{i+1}\times n_{i+1}}$

end for

$M_i^{-1} = V_i A_{i+1}^{-1} V_i^\top + \sum_{j=1}^{N_i} R_{i,j}^\top A_{i,j}^{-1} R_{i,j}$

Moreover, due to the fact that $\tilde{A}_{\mathcal{G}_{i-1,j}}$ and $\tilde{A}_{1,k}$ are Neumann matrices, we have

$$u^\top \tilde{A}_{\mathcal{G}_{i-1,j}} u \leqslant u^\top \sum_{k\in\mathcal{G}_{i,j}^1} \tilde{A}_{1,k} u \leqslant k_1 u^\top \tilde{A}_{\mathcal{G}_{i-1,j}} u.$$

On one hand, the kernels of $\tilde{A}_{1,k}$ for $k \in \mathcal{G}_{i,j}^1$ are included, by construction, in the image of V_1, see Definition 3.2. So is their intersection which is the kernel of $\sum_{k\in\mathcal{G}_{i,j}^1} \tilde{A}_{1,k}$. On the other hand, the previous two-sided inequality implies that the kernels of $\tilde{A}_{\mathcal{G}_{i-1,j}}$ and $\sum_{k\in\mathcal{G}_{i,j}^1} \tilde{A}_{1,k}$ are identical. Hence, the kernel of $\tilde{A}_{\mathcal{G}_{i-1,j}}$ is included in the image of $QQ^\top$, where $Q = \left(\prod_{l=1}^{i-1} V_l\right)$. ∎

Theorem 5.4 proves that the kernel of the Neumann matrix of a union of subdomains at level 1 that hierarchically contribute to form a subdomain at level i is conserved by the construction of the hierarchical coarse spaces. For example in the case of linear elasticity, it is essential to include the rigid body motions in the coarse space in order to have a fast convergence. As these are included in the kernel of the Neumann matrix of the subdomain, the hierarchical coarse space includes them, consequently.

6. Numerical experiments. In this section, the developed theory is validated numerically with FreeFEM [14] for finite element discretizations and HPDDM [19] for domain decomposition methods. We present numerical experiments on two highly challenging problems illustrating the efficiency and practical usage of the proposed method. For both problems, we use $N_1 = 2{,}048$ MPI processes (equal to the number of subdomains at level 1), and the domain partitioning is performed using ParMETIS

[22], with no control on the alignments of subdomain interfaces. We compare the two-level GenEO preconditioner and its multilevel extension by varying N_2 between 4 and 256. For the two-level method, N_2 corresponds to the number of MPI processes that solve the coarse problem in a distributed fashion using MKL CPARDISO [17]. For the multilevel method, N_3 is set to 1, i.e., a three-level method is used. The goal of these numerical experiments is to show that when one switches from a two-level method with an exact coarse solver, to our proposed multilevel method, the number of outer iterations is not impacted. Thus, three levels are sufficient. As an outer solver, since all levels but the coarsest are solved approximately, the flexible GMRES [31] is used. It is stopped when relative unpreconditioned residuals are lower than 10^{-6}. Subdomain matrices $\{A_{i,j}\}_{1\leqslant i\leqslant 2, 1\leqslant j\leqslant N_i}$ are factorized concurrently using MKL PARDISO, and eigenvalue problems are solved using ARPACK [24]. In both, two- and three-level GenEO, we factorize the local matrices $A_{1,j}$ for $j \in [\![1; N_1]\!]$ and solve the generalized eigenvalue problems concurrently at the first level. For this reason, we do not take into account the time needed for these two steps which are performed without any communication between MPI processes. We compare the time needed to assemble and factorize A_2 in the two-level approach against the time needed to assemble A_2 and local SPSD matrices $\tilde{A}_{2,j}$ for $j \in [\![1; N_2]\!]$, solve the generalized eigenvalue problems concurrently on the second level, assemble, and factorize the matrix A_3 in the three-level approach. We also compare the time spent in the outer Krylov solver during the solution phase. Readers interested by a comparison of the efficiency of GenEO and multigrid methods such as GAMG [1] are referred to [18]. FreeFEM scripts used to produce the following results are available at the following URL: https://github.com/prj-/aldaas2019multi[1].

6.1. Diffusion test cases. The scalar diffusion equation with highly heterogeneous coefficient κ is solved in $[0, 1]^d$ ($d = 2$ or 3). The strong formulation of the equation is:

$$\begin{aligned} -\nabla \cdot (\kappa \nabla u) &= 1 \quad \text{in } \Omega, \\ u &= 0 \quad \text{on } \Gamma_D, \\ \frac{\partial u}{\partial n} &= 0 \quad \text{on } \Gamma_N. \end{aligned}$$

The exterior normal vector to the boundary of Ω is denoted n. Γ_D is the subset of the boundary of Ω corresponding to $x = 0$ in 2D and 3D. Γ_N is defined as the complementary of Γ_D with respect to the boundary of Ω. We discretize the equation using $\mathbb{P}_2$ and $\mathbb{P}_4$ finite elements in the 3D and 2D test cases, respectively. The number of unknowns is 441×10^6 and 784×10^6, with approximately 28 and 24 nonzero elements per row in the 3D and 2D cases, respectively. The heterogeneity is due to the jumps in the diffusion coefficient κ, see Figure 6.1, which is modeled using a combination of jumps and channels, cf. the file `coefficients.idp` from https://github.com/prj-/aldaas2019multi.

The results in two dimensions are reported in Table 6.1. The number of outer iterations for both two- and three-level GenEO is 32. The size of the level 2 operator is $n_2 = 25 \times 2{,}048 = 51{,}200$. In all numerical results, the number of eigenvectors per subdomain, here 25, is fixed. This is because ARPACK cannot a priori compute all eigenpairs below a certain threshold, and an upper bound has to be provided instead.

[1] note to reviewers: the repository is now public

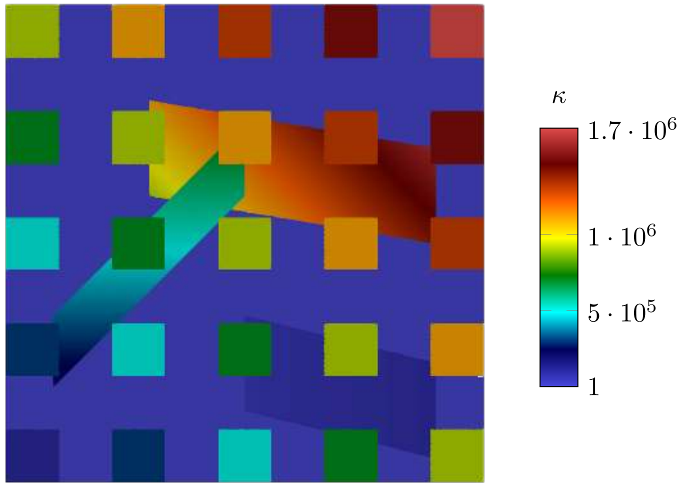

FIG. 6.1. *Variation of the coefficient κ used for the diffusion test case*

HPDDM is capable of filtering the eigenpairs for which eigenvalues are above the user-specified GenEO threshold from Lemma 2.5. However, this means that the coarse operator may be unevenly distributed. With a fixed number of eigenvectors per subdomain, it is possible to use highly optimized uniform MPI routines and block matrix formats. Hence, for performance reasons, all eigenvectors computed by ARPACK are kept when building coarse operators. It is striking that the multilevel method does not deteriorate the numerical performance of the outer solver. For the two-level method, the first column corresponds to the time needed to assemble the Galerkin operator A_2 from (3.5) (assuming V_1 has already been computed by ARPACK), and to factorize it using N_2 MPI processes. For the three-level method, the first column corresponds to the time needed to assemble level 2 local subdomain matrices $\{A_{2,j}\}_{1\leqslant j\leqslant N_2}$, level 2 local SPSD matrices, solve the generalized eigenvalue problem (3.2) concurrently, assemble the Galerkin operator A_3 and factorize it on a single process. The size of the level 3 operator is $n_3 = 20 \times N_2$. For both two- and three-level methods, the second column is the time spent in the outer Krylov solver once the preconditioner has been set up. In the last column of the three-level method, the number of inner iterations for solving systems involving A_2, which is not inverted exactly anymore, is reported. For all tables, this column is an average over all successive outer iterations. Another important numerical property of our method is that, thanks to fully controlled bounds at each level, the number of inner iterations is low, independently of the number of superdomains N_2. Because this problem is not large enough, it is still tractable by a two-level method, for which HPDDM was highly optimized for. Thus, there is no performance gain to be expected at this scale. However, one can notice that the construction of the coarse operator(s) scales nicely with N_2 for the three-level method, whereas the performance of the direct solver MKL CPARDISO quickly stagnates because of the finer and finer parallel workload granularity.

The results in three dimensions are reported in Table 6.2. The number of outer iterations for both the two- and three-level GenEO is 19. The observations made in two dimensions still hold, and the dimensions of A_2 and A_3 are the same. Once again, it is important to note that the number of outer iterations is the same for both methods.

	two-level GenEO			three-level GenEO			
N_2	CS	solve	% of nnz A_2	CS	solve	inner it.	% of nnz A_3
4	2.4	11.9	0.19	6.5	27.4	14	56.0
16	1.8	11.3		3.6	15.4	15	19.0
64	1.9	12.1		3.0	16.7	14	5.5
256	2.4	18.4		2.8	13.9	13	1.4

TABLE 6.1
Diffusion 2D test case, comparison between two- and three-level GenEO. The percentage of nonzero entries in A_1 is 0.3%.

	two-level GenEO			three-level GenEO			
N_2	CS	solve	% of nnz A_2	CS	solve	inner it.	% of nnz A_3
4	7.0	20.9	0.36	16.9	43.6	17	62.0
16	5.0	19.8		7.7	26.7	17	28.0
64	5.1	20.1		5.8	32.7	15	8.9
256	5.2	24.1		5.3	22.6	14	2.6

TABLE 6.2
Diffusion 3D test case, comparison between two- and three-level GenEO. The percentage of nonzero entries in A_1 is 0.5%.

6.2. Linear elasticity test cases. The system of linear elasticity with highly heterogeneous elastic moduli is solved in 2D and 3D. The strong formulation of the equation is given as:

$$\begin{aligned} \operatorname{div} \sigma(u) + f &= 0 \quad \text{in } \Omega, \\ u &= 0 \quad \text{on } \Gamma_D, \\ \sigma(u) \cdot n &= 0 \quad \text{on } \Gamma_N. \end{aligned} \tag{6.1}$$

The physical domain Ω is a beam of dimensions $[0, 10] \times [0, 1]$, extruded for $z \in [0, 1]$ in 3D. The Cauchy stress tensor $\sigma(\cdot)$ is given by Hooke's law: it can be expressed in terms of Young's modulus E and Poisson's ratio ν.

$$\sigma_{ij}(u) = \begin{cases} 2\mu\varepsilon_{ij}(u) & i \neq j, \\ 2\mu\varepsilon_{ii}(u) + \lambda \operatorname{div}(u) & i = j, \end{cases}$$

where

$$\varepsilon_{ij}(u) = \frac{1}{2}\left(\frac{\partial u_i}{\partial x_i} + \frac{\partial u_j}{\partial x_j}\right), \mu = \frac{E}{2(1+\nu)}, \text{and } \lambda = \frac{E\nu}{1-2\nu}.$$

The exterior normal vector to the boundary of Ω is denoted n. Γ_D is the subset of the boundary of Ω corresponding to $x = 0$ in 2D and 3D. Γ_N is defined as the complementary of Γ_D with respect to the boundary of Ω. We discretize (6.1) using the following vectorial finite elements: $(\mathbb{P}_2, \mathbb{P}_2, \mathbb{P}_2)$ in 3D and $(\mathbb{P}_3, \mathbb{P}_3)$ in 2D. The number of unknowns is 146×10^6 and 847×10^6, with approximately 82 and 34 nonzero elements per row in the 3D and 2D cases, respectively. The heterogeneity is due to the jumps in E and ν. We consider discontinuous piecewise constant values for E and ν: $(E_1, \nu_1) = (2 \times 10^{11}, 0.25)$, $(E_2, \nu_2) = (10^7, 0.45)$, see Figure 6.2.

Results in two (resp. three) dimensions are reported in Table 6.3 (resp. Table 6.4). The number of outer iterations are 73 and 45 respectively. For these test cases, we

FIG. 6.2. *Variation of the structure coefficients used for the elasticity test case*

	two-level GenEO			three-level GenEO			
N_2	CS	solve	% of nnz A_2	CS	solve	inner it.	% of nnz A_3
4	4.8	52.7		22.5	179.3	31	43.0
16	3.9	50.3	0.18	9.3	124.9	57	17.0
64	4.0	53.1		7.2	71.5	34	4.9
256	4.8	63.2		6.8	71.2	44	1.4

TABLE 6.3
Elasticity 2D test case, comparison between two- and three-level GenEO. The percentage of nonzero entries in A_1 is 0.4%.

slightly relaxed the criterion for selecting eigenvectors in coarse spaces, which explains why the iteration counts increase. However, the same observations as for the diffusion test cases still hold. The dimension of the level 2 matrix is $n_2 = 50 \times 2{,}048 = 1.02 \cdot 10^5$, while for the level 3 matrix it is $n_3 = 20 \times N_2$. This means that 50 (resp. 20) eigenvectors are kept per level 1 (resp. level 2) subdomains. We observe that the number of iterations of the inner solver increases slowly when increasing the number of subdomains from 4 to 256 in the 2D case and remains almost constant in the 3D case. In terms of runtime, the two-level GenEO is faster than three-level GenEO for these matrices of medium dimensions.

To show the potential of our method at larger scales, a three-dimensional linear elasticity problem of size 593×10^6 is now solved on $N_1 = 16{,}384$ processes and $N_2 = 256$ superdomains. With the two-level method, A_2 is assembled and factorized in 40.8 seconds. With the three-level method, this step now takes 35.1 seconds, see Table 6.5. There is a two iterations difference in the iteration count. Not taking into account the preconditioner setup, the problem is solved in 222.5 seconds in the two-level case and 90.1 seconds in the multilevel case. In this test case the cost of applying the two-level preconditioner on a given vector is approximately twice the cost of applying the multilevel variant. At this regime, it is clear that there are important gains for the solution phase. At even greater scales, gains for the setup phase are also expected. Moreover, another interesting fact to note regarding computation time is that the generalized eigenvalue problems solved concurrently at the first level to obtain V_1 actually represents a significant part of the total time of 377.6 seconds (resp. 244.8 seconds) with the two- (resp. three-)level method: 78.2 seconds. This cost can be reduced by taking a larger number of (smaller) subdomains, with the drawback of increasing the size of V_1 and thus A_2. This drawback represents a clear bottleneck for the two-level method but is alleviated by using the three-level method, making it a good candidate for problems at greater scales.

7. Conclusion.

In this paper, we reviewed general properties of overlapping Schwarz preconditioners and presented a framework for its multilevel extension. We

	two-level GenEO			three-level GenEO			
N_2	CS	solve	% of nnz A_2	CS	solve	inner it.	% of nnz A_3
4	28.5	46.9	0.38	78.9	296.7	23	43.0
16	17.3	35.4		24.5	124.5	23	19.0
64	15.0	33.2		15.4	62.2	21	7.9
256	13.6	40.7		10.6	50.7	23	2.5

TABLE 6.4
Elasticity 3D test case, comparison between two- and three-level GenEO. The percentage of nonzero entries in A_1 is 3.3%.

	two-level GenEO		three-level GenEO		
N_2	CS	solve	CS	solve	inner it.
256	40.8	222.5	35.1	90.1	11

TABLE 6.5
Elasticity 3D test case, comparison between two- and three-level GenEO

generalized the local SPSD splitting presented in [3] to cover a larger set of matrices leading to more flexibility for building robust coarse spaces. Based on local SPSD matrices on the first level, we presented how to compute local SPSD matrices for coarser levels. The multilevel solver based on hierarchical local SPSD matrices is robust and guarantees a bound on the condition number of the preconditioned matrix at each level depending on predefined values. Numerical experiments illustrate the theory and prove the efficiency of the method on challenging problems of large size arising from heterogeneous linear elasticity and diffusion problems with jumps in the coefficients of multiple orders of magnitude.

8. Acknowledgments. We would like to thank the anonymous referees for their comments and remarks that helped us improve the clarity of this manuscript. This work was granted access to the HPC resources of TGCC@CEA under the allocation A0050607519 made by GENCI. The work of the second author was supported by the NLAFET project as part of European Union's Horizon 2020 research and innovation program under grant 671633.

REFERENCES

[1] M. F. Adams, H. H. Bayraktar, T. M. Keaveny, and P. Papadopoulos, *Ultrascalable Implicit Finite Element Analyses in Solid Mechanics with over a Half a Billion Degrees of Freedom*, in Proceedings of the 2004 ACM/IEEE Conference on Supercomputing, SC '04, IEEE Computer Society, 2004.

[2] M. F. Adams and J. W. Demmel, *Parallel Multigrid Solver for 3D Unstructured Finite Element Problems*, in Proceedings of the 1999 ACM/IEEE Conference on Supercomputing, SC '99, ACM, 1999.

[3] H. Al Daas and L. Grigori, *A class of efficient locally constructed preconditioners based on coarse spaces*, SIAM Journal on Matrix Analysis and Applications, 40 (2019), pp. 66–91.

[4] S. Badia, A. Martín, and J. Principe, *Multilevel balancing domain decomposition at extreme scales*, SIAM Journal on Scientific Computing, 38 (2016), pp. C22–C52.

[5] P. E. Bjørstad, M. J. Gander, A. Loneland, and T. Rahman, *Does SHEM for Additive Schwarz Work Better than Predicted by Its Condition Number Estimate?*, in International Conference on Domain Decomposition Methods, Springer, 2017, pp. 129–137.

[6] A. Borzì, V. De Simone, and D. di Serafino, *Parallel algebraic multilevel Schwarz preconditioners for a class of elliptic PDE systems*, Computing and Visualization in Science, 16 (2013), pp. 1–14.

[7] M. Brezina, A. Cleary, R. Falgout, V. Henson, J. Jones, T. Manteuffel, S. McCormick, and J. Ruge, *Algebraic Multigrid Based on Element Interpolation (AMGe)*, SIAM Journal on Scientific Computing, 22 (2001), pp. 1570–1592.

[8] X.-C. Cai and M. Sarkis, *A restricted additive Schwarz preconditioner for general sparse linear systems*, SIAM Journal on Scientific Computing, 21 (1999), pp. 792–797.

[9] T. F. Chan and T. P. Mathew, *Domain decomposition algorithms*, Acta Numerica, 3 (1994), pp. 61—143.

[10] T. Chartier, R. D. Falgout, V. E. Henson, J. Jones, T. Manteuffel, S. McCormick, J. Ruge, and P. S. Vassilevski, *Spectral AMGe (ρAMGe)*, SIAM Journal on Scientific Computing, 25 (2003), pp. 1–26.

[11] C. Chevalier and F. Pellegrini, *PT-SCOTCH: A tool for efficient parallel graph ordering*, Parallel Computing, 34 (2008), pp. 318–331. Parallel Matrix Algorithms and Applications.

[12] V. Dolean, P. Jolivet, and F. Nataf, *An introduction to domain decomposition methods*, Society for Industrial and Applied Mathematics, 2015. Algorithms, theory, and parallel implementation.

[13] M. Griebel and P. Oswald, *On the abstract theory of additive and multiplicative Schwarz algorithms*, Numerische Mathematik, 70 (1995), pp. 163–180.

[14] F. Hecht, *New development in FreeFem++*, Journal of Numerical Mathematics, 20 (2012), pp. 251–266.

[15] A. Heinlein, A. Klawonn, O. Rheinbach, and F. Röver, *A Three-Level Extension of the GDSW Overlapping Schwarz Preconditioner in Three Dimensions*, technical report, Universität zu Köln, November 2018.

[16] V. E. Henson and U. M. Yang, *BoomerAMG: A parallel algebraic multigrid solver and preconditioner*, Applied Numerical Mathematics, 41 (2002), pp. 155–177. Developments and Trends in Iterative Methods for Large Systems of Equations.

[17] Intel, *MKL web page*. https://software.intel.com/en-us/intel-mkl, 2019.

[18] P. Jolivet, *Domain decomposition methods. Application to high-performance computing*, theses, Université de Grenoble, Oct. 2014.

[19] P. Jolivet, F. Hecht, F. Nataf, and C. Prud'homme, *Scalable domain decomposition preconditioners for heterogeneous elliptic problems*, in Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC13, ACM, 2013.

[20] J. Jones and P. Vassilevski, *AMGe Based on Element Agglomeration*, SIAM Journal on Scientific Computing, 23 (2001), pp. 109–133.

[21] D. Kalchev, C. Lee, U. Villa, Y. Efendiev, and P. Vassilevski, *Upscaling of mixed finite element discretization problems by the spectral AMGe method*, SIAM Journal on Scientific Computing, 38 (2016), pp. A2912–A2933.

[22] G. Karypis and V. Kumar, *Multilevel k-way partitioning scheme for irregular graphs*, Journal of Parallel and Distributed Computing, 48 (1998), pp. 96–129.

[23] F. Kong and X.-C. Cai, *A highly scalable multilevel Schwarz method with boundary geometry preserving coarse spaces for 3D elasticity problems on domains with complex geometry*, SIAM Journal on Scientific Computing, 38 (2016), pp. C73–C95.

[24] R. Lehoucq, D. Sorensen, and C. Yang, *ARPACK users' guide: solution of large-scale eigenvalue problems with implicitly restarted Arnoldi methods*, vol. 6, Society for Industrial and Applied Mathematics, 1998.

[25] J. Mandel, B. Sousedík, and C. R. Dohrmann, *Multispace and multilevel BDDC*, Computing, 83 (2008), pp. 55–85.

[26] O. Marques, A. Druinsky, X. S. Li, A. T. Barker, P. Vassilevski, and D. Kalchev, *Tuning the coarse space construction in a spectral AMG solver*, Procedia Computer Science, 80 (2016), pp. 212–221. International Conference on Computational Science 2016, ICCS 2016, 6-8 June 2016, San Diego, California, USA.

[27] S. V. Nepomnyaschikh, *Mesh theorems of traces, normalizations of function traces and their inversions*, Russian Journal of Numerical Analysis and Mathematical Modelling, 6 (1991), pp. 1–25.

[28] ———, *Decomposition and fictitious domains methods for elliptic boundary value problems*, 1992.

[29] Y. Notay, *An aggregation-based algebraic multigrid method*, Electronic Transactions on Numerical Analysis, 37 (2010), pp. 123–146.

[30] C. Pechstein, *Finite and boundary element tearing and interconnecting solvers for multiscale problems*, vol. 90, Springer Science & Business Media, 2012.

[31] Y. Saad., *A Flexible Inner–Outer Preconditioned GMRES Algorithm*, SIAM Journal on Scientific Computing, 14 (1993), pp. 461–469.

[32] Y. Saad, *Iterative Methods for Sparse Linear Systems*, Society for Industrial and Applied Mathematics, 2nd ed., 2003.

[33] N. Spillane, V. Dolean, P. Hauret, F. Nataf, C. Pechstein, and R. Scheichl, *Abstract robust coarse spaces for systems of PDEs via generalized eigenproblems in the overlaps*, Numerische Mathematik, 126 (2014), pp. 741–770.

[34] J. Toivanen, P. Avery, and C. Farhat, *A multilevel feti-dp method and its performance for problems with billions of degrees of freedom*, International Journal for Numerical Methods in Engineering, 116 (2018), pp. 661–682.

[35] A. Toselli and O. Widlund, *Domain Decomposition Methods - Algorithms and Theory*, Springer Series in Computational Mathematics, Springer Berlin Heidelberg, 2005.

[36] J. Xu, *Theory of Multilevel Methods*, PhD thesis, Cornell University, 1989.

[37] S. Zampini, *PCBDDC: A class of robust dual-primal methods in PETSc*, SIAM Journal on Scientific Computing, 38 (2016), pp. S282–S306.