

This is the author's final, peer-reviewed manuscript as accepted for publication (AAM). The version presented here may differ from the published version, or version of record, available through the publisher's website. This version does not track changes, errata, or withdrawals on the publisher's site.

Efficient Algebraic Two-Level Schwarz Preconditioner for Sparse Matrices

HUSSAM AL DAAS, PIERRE JOLIVET, AND TYRONE REES

Published version information

Citation: H Al Daas et al. "EFFICIENT ALGEBRAIC TWO-LEVEL SCHWARZ PRECONDITIONER FOR SPARSE MATRICES ", SIAM Journal on Scientific Computing 2023 45:3, A1199-A1213,

DOI: <https://doi.org/10.1137/22M1469833>

This version is made available in accordance with publisher policies. Please cite only the published version using the reference above. This is the citation assigned by the publisher at the time of issuing the AAM/APV. Please check the publisher's website for any updates.

This item was retrieved from **ePubs**, the Open Access archive of the Science and Technology Facilities Council, UK. Please contact epublications@stfc.ac.uk or go to <http://epubs.stfc.ac.uk/> for further information and policies.

EFFICIENT ALGEBRAIC TWO-LEVEL SCHWARZ PRECONDITIONER FOR SPARSE MATRICES

HUSSAM AL DAAS*, PIERRE JOLIVET†, AND TYRONE REES†

Abstract. Domain decomposition methods are among the most efficient for solving sparse linear systems of equations. Their effectiveness relies on a judiciously chosen coarse space. Originally introduced and theoretically proved to be efficient for positive definite self-adjoint operators, spectral coarse spaces have been proposed in the past few years for indefinite and non-self-adjoint operators. This paper presents a new spectral coarse space that can be constructed in a fully-algebraic way, unlike most existing spectral coarse spaces. We present a theoretical convergence result for Hermitian positive definite diagonally dominant matrices. Numerical experiments and comparisons against state-of-the-art preconditioners in the multigrid community show that the resulting two-level Schwarz preconditioner is efficient, especially for non-self-adjoint operators. Furthermore, in this case, our proposed preconditioner outperforms state-of-the-art preconditioners.

Key words. Algebraic domain decomposition, Schwarz preconditioner, sparse linear systems, diagonally dominant matrices.

1. Introduction. In this paper, we develop an algebraic overlapping Schwarz preconditioner for the linear system of equations

$$Ax = b,$$

for a sparse matrix $A \in \mathbb{C}^{n \times n}$ and a given vector $b \in \mathbb{C}^n$. Solving sparse linear systems of equations is omnipresent in scientific computing. Direct approaches, based on Gaussian elimination, have proved to be robust and efficient for a wide range of problems [23]. However, the memory required to apply sparse direct methods often scales poorly with the problem size, particularly for three-dimensional discretizations of partial differential equations (PDEs). Furthermore, the algorithms underpinning sparse direct software are poorly suited to parallel computation, which makes them difficult to adapt to emerging computing architectures.

Iterative methods for solving linear systems [45] have been an active research topic since early computing days. Their simple structure, at their most basic level requiring only matrix-vector multiplication and vector-vector operations, makes them attractive for tackling large-scale problems. However, since the convergence rate depends on the properties of the linear system, iterative methods are not, in general, robust. For the class of iterative methods known as Krylov subspace methods, we may alleviate this by applying a preconditioner, which transforms the problem into one with more favourable numerical properties. The choice of the preconditioner is usually problem-dependent, and a wide variety of preconditioning techniques have been proposed to improve the convergence rate of iterative methods, see for example the recent survey [44] and the references therein.

Multilevel domain decomposition (DD) and multigrid methods are widely used preconditioners [21, 47, 52, 53, 54]. They have proved to be effective on a wide variety of matrices, but they are especially well suited to sparse Hermitian positive definite (HPD) matrices arising from the discretization of PDEs. Their efficiency stems from a judicious combination of a cheap fine-level approximate solver with a

*STFC Rutherford Appleton Laboratory, Harwell Campus, Didcot, Oxfordshire, OX11 0QX, UK (hussam.al-daas@stfc.ac.uk, tyrone.rees@stfc.ac.uk).

†Sorbonne Université, CNRS, LIP6, 4 place Jussieu, 75252 Paris Cedex 05, France (pierre@joliv.etc).

coarse-space correction. In the last two decades, there has been a great advance in the development of spectral coarse spaces that yield efficient preconditioners. Spectral coarse spaces were initially proposed in the multigrid community for elliptic PDEs with positive definite self-adjoint operators [15, 19, 24, 34], and similar ideas were later picked up by the DD community for the same kind of problems [2, 3, 4, 7, 32, 33, 40, 48, 49]. The spectral coarse space based on Generalized Eigenvalues in the Overlap (GenEO) was introduced in [48] to construct a scalable DD preconditioner for symmetric positive definite matrices arising from the discretization of elliptic PDEs using the finite element method. The past three years have seen several approaches to tackle symmetric indefinite systems and non-self-adjoint problems. For example, spectral coarse spaces for least squares problems and symmetric indefinite saddle-point systems were proposed in [5, 39], where coarse spaces are proposed for related Hermitian positive definite matrices. An exciting new development is that a number of multigrid methods and spectral coarse spaces have been suggested for problems with indefinite or non-self-adjoint operators [9, 10, 11, 12, 13, 14, 22, 36, 37]. These coarse spaces are mainly based on heuristics and show efficiency on several challenging model problems arising from discretized PDEs.

A variety of mathematical tools such as the fictitious subspace lemma [41] and local Hermitian positive semi-definite (HPSD) splitting matrices [2] are now available to analyze and propose effective coarse spaces for positive definite self-adjoint operators. For example, the authors in [2] show that given local HPSD splitting matrices of an HPD matrix, one can construct a spectral coarse space such that the condition number of the preconditioned matrix is bounded by a user-defined number. However, these tools may not be directly used for indefinite or non-self-adjoint operators. An alternative approach to studying the convergence of DD methods in the indefinite or non-self-adjoint case is to use Elman’s theory [25] of GMRES convergence, see for example [8, 10, 18].

In this work, we propose a fully-algebraic spectral coarse space for the two-level Schwarz preconditioner for general sparse matrices. We review the overlapping DD framework in section 2, including a summary of the main features of local HPSD splitting matrices. For each subdomain, we introduce the *local block splitting using lumping in the overlap* of a sparse matrix in section 3. The coarse space is then constructed by solving locally and concurrently a generalized eigenvalue problem involving the local block splitting matrix and the local subdomain matrix. In the case where the matrix is HPD diagonally dominant, we prove that the local block splitting matrices are local HPSD splitting matrices, and in that case, we show that one can bound the condition number of the preconditioned matrix from above by a user-defined number. In comparison to the previously introduced HPSD splitting matrices in [4], the proposed splitting in this paper is significantly cheaper to construct. If n_i is the number of unknowns per subdomain, the proposed splitting can be constructed in $O(n_i)$ operations. The splitting from [4] has a cubic cost $O(n_i^3)$ coming from a preprocessing step requiring the solution of a subdomain-wise local singular value decomposition. We heuristically generalize the proposed method for general sparse matrices.

Unlike most existing spectral coarse spaces, especially those suggested for non-symmetric matrices, we obtain the matrices involved in the local generalized eigenvalue problem efficiently from the coefficient matrix; the preconditioner is therefore fully-algebraic. In order to assess the proposed preconditioner, we provide in section 4 a set of numerical experiments on problems arising from a wide range of applications including the convection-diffusion equation and other linear systems

from the SuiteSparse Collection [20]. Furthermore, we compare our proposed preconditioner against state-of-the-art preconditioners in the multigrid community. Finally, we give concluding remarks and future lines of research in section 5.

The contributions we bring in this paper are:

1. we introduce the local block splitting using lumping in the overlap;
2. we show that the resulting splitting corresponds to local HPSD splitting when the matrix is HPD diagonally dominant, which leads to bounding the condition number of the preconditioned matrix in this case;
3. we demonstrate empirically the efficiency of the resulting preconditioner for general sparse matrices.

Notation. Let $1 \leq n \leq m$ and let $M \in \mathbb{C}^{m \times n}$ be a complex sparse matrix. Let $\llbracket 1, p \rrbracket$ denote the set of the first p positive integers, and let $S_1 \subset \llbracket 1, m \rrbracket$ and $S_2 \subset \llbracket 1, n \rrbracket$. $M(S_1, :)$ is the submatrix of M formed by the rows whose indices belong to S_1 and $M(:, S_2)$ is the submatrix of M formed by the columns whose indices belong to S_2 . $M(S_1, S_2)$ denotes the submatrix formed by taking the rows whose indices belong to S_1 and only retaining the columns whose indices belong to S_2 . $\llbracket S_1 \ S_2 \rrbracket$ means the concatenation of any two sets of integers S_1 and S_2 , where the order of the concatenation is important. I_n is the identity matrix of size n , the transpose matrix of M is denoted M^\top , and the adjoint of M , denoted M^H , is the conjugate transpose of M , i.e., $M^H = \bar{M}^\top$. $\ker(M)$ and $\text{range}(M)$ denote the null space and the range of M , respectively.

2. Domain decomposition. Consider $\mathcal{G}(A + A^\top)$, the undirected adjacency graph of the coefficient matrix in (1), and number its nodes, V , from 1 to n . Using a graph partitioning algorithm, we split V into $N \ll n$ *nonoverlapping subdomains*, i.e., disjoint subsets Ω_{I_i} , $i \in \llbracket 1, N \rrbracket$, of size n_{I_i} . Let Ω_{Γ_i} be the subset, of size n_{Γ_i} , of nodes that are distance one in $\mathcal{G}(A)$ from the nodes in Ω_{I_i} , $i \in \llbracket 1, N \rrbracket$. We define the *overlapping subdomain*, Ω_i , as $\Omega_i = \llbracket \Omega_{I_i}, \Omega_{\Gamma_i} \rrbracket$, with size $n_i = n_{\Gamma_i} + n_{I_i}$. The complement of Ω_i in $\llbracket 1, n \rrbracket$ is denoted by Ω_{c_i} .

Associated with Ω_i is a restriction (or projection) matrix $R_i \in \mathbb{R}^{n_i \times n}$ given by $R_i = I_n(\Omega_i, :)$. R_i maps from the global domain to the overlapping subdomain Ω_i . Its transpose R_i^\top is a prolongation matrix that maps from subdomain Ω_i to the global domain.

We define the *one-level Schwarz preconditioner* as

$$(2.1) \quad M_{\text{ASM}}^{-1} = \sum_{i=1}^N R_i^\top A_{ii}^{-1} R_i,$$

where we assume $A_{ii} = R_i A R_i^\top$ is nonsingular for $i \in \llbracket 1, N \rrbracket$.

Applying this preconditioner to a vector involves solving concurrent local problems in each subdomain. Increasing N reduces the size of the subdomains, leading to smaller local problems and, correspondingly, faster computations. However, in practice, preconditioning by M_{ASM}^{-1} alone is often not enough for the convergence of the iterative solver to be sufficiently rapid. We can improve convergence, while still maintaining robustness with respect to N , by applying a suitably chosen *coarse space*, or *second-level* [2, 8, 21, 27].

Let $\mathcal{R} \subset \mathbb{C}^n$ be a subspace of dimension $0 < n_0 \ll n$ and let $R_0 \in \mathbb{C}^{n_0 \times n}$ be a matrix such that the columns of R_0^H span the subspace \mathcal{R} . Assuming that $A_{00} = R_0 A R_0^H$ is nonsingular, we define the *two-level Schwarz preconditioner* as

$$(2.2) \quad M_{\text{additive}}^{-1} = R_0^H A_{00}^{-1} R_0 + M_{\text{ASM}}^{-1}.$$

Note that R_i is real valued hence its adjoint is its transpose while R_0 may be complex valued thus its adjoint is its Hermitian transpose. Such preconditioners have been used to solve a large class of systems arising from a range of engineering applications (see, for example, [3, 5, 29, 35, 38, 47, 51] and references therein).

We denote by $D_i \in \mathbb{R}^{n_i \times n_i}$, $i \in \llbracket 1, N \rrbracket$, any non-negative diagonal matrices such that

$$\sum_{i=1}^N R_i^\top D_i R_i = I_n.$$

We refer to $(D_i)_{1 \leq i \leq N}$ as an *algebraic partition of unity*.

Variants of one- and two-level preconditioners. The so-far presented Schwarz preconditioners are the additive one-level (2.1) and the additive two-level based on additive coarse space correction (2.2). It was noticed in [17] that scaling the one-level Schwarz preconditioner by using the partition of unity usually yields faster convergence. The resulting one-level preconditioner is referred to as the restricted additive Schwarz preconditioner and is defined as:

$$(2.3) \quad M_{\text{RAS}}^{-1} = \sum_{i=1}^N R_i^\top D_i A_{ii}^{-1} R_i.$$

Furthermore, there are a number of ways of how to combine the coarse space with a one-level preconditioner such as the additive, deflated, and balanced combinations, see for example [50]. Given a one-level preconditioner M_\star^{-1} , where the subscript \star stands for either ASM or RAS, the two-level preconditioner with additive coarse space correction is defined as

$$M_{\star, \text{additive}}^{-1} = R_0^H A_{00}^{-1} R_0 + M_\star^{-1}.$$

The two-level preconditioner based on a deflated coarse space correction is

$$M_{\star, \text{deflated}}^{-1} = R_0^H A_{00}^{-1} R_0 + M_\star^{-1} (I - A R_0^H A_{00}^{-1} R_0).$$

Due to its simple form, the additive two-level Schwarz based on the additive coarse space correction is the easiest to analyze. However, we observe that the deflated variant combined with the restricted additive Schwarz preconditioner has better performance in practice. The theory and presentation in this work employ the additive two-level Schwarz preconditioner using an additive coarse space correction, however, all numerical experiments involving the proposed preconditioner employ the restricted additive two-level Schwarz with deflated coarse space correction so that the two-level preconditioner used in section 4 reads as

$$(2.4) \quad M_{\text{deflated}}^{-1} = R_0^H A_{00}^{-1} R_0 + M_{\text{RAS}}^{-1} (I - A R_0^H A_{00}^{-1} R_0).$$

Note that the aforementioned variants are agnostic to the choice of partitioning and the coarse space. That is, once the restriction operators to the subdomains and the coarse space are set, all these variants are available.

Local HPSD splitting matrices of sparse HPD matrix. A local HPSD matrix associated with subdomain i is any HPSD matrix of the form

$$P_i \tilde{A}_i P_i^\top = \begin{pmatrix} A_{Ii} & A_{\Gamma, i} \\ A_{\Gamma, i} & \tilde{A}_{\Gamma, i} \end{pmatrix},$$

where $P_i = I_n([\Omega_{Ii} \ \Omega_{\Gamma i} \ \Omega_{ci}], :)$ is a permutation matrix, $A_{Ii} = A(\Omega_{Ii}, \Omega_{Ii})$, $A_{I\Gamma, i} = A_{\Gamma I, i}^\top = A(\Omega_{Ii}, \Omega_{\Gamma i})$, and $\tilde{A}_{\Gamma, i}$ is any HPSD matrix such that the following inequality holds

$$0 \leq u^\top \tilde{A}_i u \leq u^\top A u, \quad u \in \mathbb{C}^n.$$

First presented and analyzed in [2], local HPSD splitting matrices provide a framework to construct robust two-level Schwarz preconditioners for sparse HPD matrices. Recently, this has led to the introduction of robust multilevel Schwarz preconditioners for finite element SPD matrices [3], sparse normal equations matrices [5], and sparse general SPD matrices [4].

3. Two-level Schwarz preconditioner for sparse matrices. We present in this section the construction of a two-level preconditioner for sparse matrices. First, we introduce a new local splitting matrix associated with subdomain i that uses local values of A to construct a preconditioner which is cheap to set up. In the special case where A is HPD diagonally dominant, we prove that these local matrices are local HPSD splitting matrices of A . We demonstrate empirically that these matrices, when used to construct a GenEO-like coarse space, produce a two-level Schwarz preconditioner that outperforms existing two-level preconditioners in many applications, particularly in the difficult case where A is not symmetric.

3.1. Local block splitting matrices of A using lumping in the overlap.

DEFINITION 3.1. *Local block splitting matrix.* Given the overlapping partitioning of A presented in [section 2](#), we have for each $i \in \llbracket 1, N \rrbracket$

$$P_i A P_i^\top = \begin{pmatrix} A_{Ii} & A_{I\Gamma i} & \\ A_{\Gamma I i} & A_{\Gamma i} & A_{\Gamma ci} \\ & A_{c\Gamma i} & A_{ci} \end{pmatrix}$$

Let s_i be the vector whose j th component is the sum of the absolute values of the j th row of $A_{\Gamma ci}$, and let $S_i = \text{diag}(s_i)$. Define $\tilde{A}_{\Gamma i} = A_{\Gamma i} - S_i$. The local block splitting matrix of A associated with subdomain i is defined to be $\tilde{A}_i = R_i^\top \tilde{A}_{ii} R_i$, where

$$\tilde{A}_{ii} = \begin{pmatrix} A_{Ii} & A_{I\Gamma i} \\ A_{\Gamma I i} & \tilde{A}_{\Gamma i} \end{pmatrix}.$$

Note that we only require the sum of the absolute values of each row in the local matrix $A_{\Gamma ci}$ to construct the local block splitting matrix of A associated with subdomain i . Then, each of these values is subtracted from the corresponding diagonal entry of the local matrix $A_{\Gamma i}$. We can therefore construct \tilde{A}_{ii} cheaply and concurrently for each subdomain.

The following lemma shows that if A is HPD diagonally dominant, the local splitting matrices defined in [Definition 3.1](#) are local HPSD splitting matrices.

LEMMA 3.2. *Let A be HPD diagonally dominant. The local block splitting matrix \tilde{A}_i defined in [Definition 3.1](#) is local HPSD splitting matrix of A with respect to subdomain i .*

Proof. First, note that the j th diagonal element of \tilde{A}_i is

$$\tilde{A}_i(j, j) = \begin{cases} A(j, j) & \text{if } j \in \Omega_{Ii}, \\ A(j, j) - s_i(j) & \text{if } j \in \Omega_{\Gamma i}, \\ 0 & \text{if } j \in \Omega_{ci}, \end{cases}$$

where s_i is the vector whose j th component is the sum of the absolute values of the j th row of $A_{\Gamma_{ci}}$. Therefore, by construction, \tilde{A}_i is Hermitian diagonally dominant, hence HPSD. Furthermore, $A - \tilde{A}_i$ is Hermitian and diagonally dominant, hence HPSD. By the local structure of \tilde{A}_i , we conclude it is HPSD splitting of A with respect to subdomain i . \square

3.2. Coarse space. In this section, we present a coarse space for the two-level Schwarz preconditioner. For each $i \in \llbracket 1, N \rrbracket$, given the local nonsingular matrix A_{ii} , the local splitting matrix $\tilde{A}_{ii} = R_i \tilde{A}_i R_i^\top$, and the partition of unity matrix D_i , let $L_i = \ker(D_i A_{ii} D_i)$ and $K_i = \ker(\tilde{A}_{ii})$. Now, define the following generalized eigenvalue problem:

$$(3.1) \quad \begin{aligned} &\text{find } (\lambda, u) \in \mathbb{C} \times \mathbb{C}^{n_i} \text{ such that} \\ &\Pi_i D_i A_{ii} D_i \Pi_i u = \lambda \tilde{A}_{ii} u, \end{aligned}$$

where Π_i is the projection on $\text{range}(\tilde{A}_{ii})$.

Given a number $\tau > 0$, the coarse space we propose is defined to be the space generated by the columns of the matrix

$$R_i^H = [R_1^\top D_1 Z_1 \quad \cdots \quad R_N^\top D_N Z_N],$$

where Z_i is the matrix whose columns form a basis of the subspace

$$(3.2) \quad (L_i \cap K_i)^{\perp \kappa_i} \oplus \text{span} \left\{ u \mid \Pi_i D_i A_{ii} D_i \Pi_i u = \lambda \tilde{A}_{ii} u, |\lambda| \geq \frac{1}{\tau} \right\},$$

where $(L_i \cap K_i)^{\perp \kappa_i}$ is the complementary subspace of $(L_i \cap K_i)$ inside K_i . Note that the subspace $(L_i \cap K_i)^{\perp \kappa_i}$ is spanned by the generalized eigenvectors corresponding to infinite eigenvalues. In inexact arithmetic, those eigenvalues are close to ε^{-1} , where ε is the machine precision. Furthermore, for $\tau > 0$, those infinite eigenvalues still satisfy $|\lambda| \geq \frac{1}{\tau}$. For further analysis of such a generalized eigenvalue problem, we refer the reader to [5, Section 3.3.2].

In [2], the authors presented a theoretical framework to study spectral coarse spaces combined with a Schwarz preconditioner. The highlight of the paper is that finding a local HPSD splitting matrix for each subdomain provides a systematic way to obtain a spectral coarse space that guarantees a user-provided upper bound on the condition number of the preconditioned matrix. However, constructing those local splitting matrices is not a simple task for a general HPD matrix. The proposed local matrices \tilde{A}_i may not be local HPSD splitting matrices of A in general (except if A is HPD diagonally dominant). Nonetheless, they seem to lead to efficient coarse spaces. Note that in the case where A is sparse HPD diagonally dominant, \tilde{A}_i is a local HPSD splitting matrix of A , and the definition of the coarse space matches the one defined in [2]. Therefore, the two-level Schwarz preconditioner using the coarse space defined guarantees an upper bound on the condition number of the preconditioned matrix

$$\kappa_2(M_{\text{ASM,additive}}^{-1} A) \leq (k_c + 1) \left(2 + (2k_c + 1) \frac{k_m}{\tau} \right),$$

where k_c is the number of colours required to colour the graph of A such that every two neighbouring subdomains have different colours and k_m is the maximum number of overlapping subdomains sharing a row of A . Therefore, when A is sparse HPD diagonally dominant, the upper bound on $\kappa_2(M_{\text{ASM,additive}}^{-1} A)$ is independent of N and can be controlled by using the value τ .

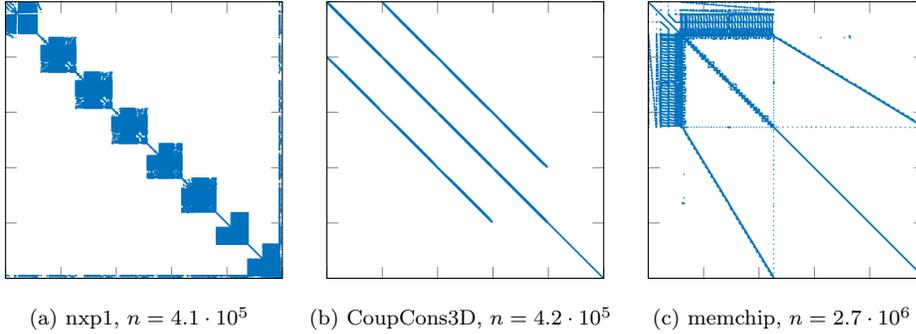


FIG. 1. Nonzero sparsity pattern of some of the test matrices from Table 1.

4. Numerical experiments. In this section, we validate the effectiveness of the two-level method when compared to other preconditioners. Table 1 presents a comparison between four preconditioners: M_{deflated}^{-1} (2.4), BoomerAMG [26], GAMG [1], and AGMG [42]. Throughout the paper, the default PETSc 3.17 parameters for GAMG and BoomerAMG are used. The results are for the right-preconditioned GMRES [46] with a restart parameter of 30 and a relative tolerance set to 10^{-8} . We highlight the fact that our proposed preconditioner can handle unstructured systems, not necessarily stemming from standard PDE discretization schemes, by displaying some nonzero patterns in Figure 1. For preconditioners used within PETSc [6] (all except AGMG), the systems are solved using 256 MPI processes. After loading them from the disk, their symmetric part $A^T + A$ is first renumbered by ParMETIS [31]. The resulting permutation is then applied to A and the corresponding linear systems are solved using a random right-hand side. The initial guess is always zero. The code that implements these steps is given in Appendix A. For our DD method, we leverage the PCHPDDM framework [30] which is used to assemble spectral coarse spaces using (3.1). The new option `-pc_hpddm_block_splitting`, introduced in PETSc 3.17, is used to compute the local splitting matrices of A from subsection 3.1. In all experiments, if a threshold τ for the eigenvalues in (3.2) is specified along with a maximum number of eigenvectors `nev` to be selected, we compute the eigenvectors associated with the `nev` largest eigenvalues (in absolute value). Then, at most `nev` eigenmodes (λ, u) are selected and they verify $|\lambda| \geq \frac{1}{\tau}$. At most 300 eigenpairs are computed on each subdomain and the threshold parameter τ from (3.2) is set to 0.6. These parameters provided good numerical performance after a quick trial-and-error approach on a single problem. We did not want to adjust them for each problem individually, but it will be shown next that they are fine overall without additional tuning.

Furthermore, a single subdomain is mapped to each process, i.e., $N = 256$ in (2.3). Eventually, the exact subdomain and second-level operator LU factorizations are computed.

4.1. Adaptive preconditioning. Whilst Table 1 highlights that tuning τ for individual problems is not necessary to successfully solve a range of problems, it does not validate the ability of our preconditioner to concurrently select the most appropriate local eigenpairs to define an adaptive preconditioner. To that end, for problem G3_circuit, we consider the effect of varying the threshold τ on

Identifier	n	$\text{nnz}(A)$	AGMG	BoomerAMG	GAMG	M_{deflated}^{-1}	n_0
light_in_tissue	29,282	406,084	15	‡	53	6	7,230
finan512	74,752	596,992	9	7	8	6	2,591
consp ϕ	83,334	6,010,480				93	31,136
Dubcova3	146,689	3,636,643		72	71	7	21,047
CO	221,119	7,666,057		25		26	56,135
nxp1	414,604	2,655,880	†	†	†	20	19,707
CoupCons3D	416,800	17,277,420		†	26	20	28,925
parabolic_fem	525,825	3,674,625	12	8	16	5	24,741
Chevron4	711,450	6,376,412		‡	†	5	22,785
apache2	715,176	4,817,870	14	11	35	8	45,966
tmt_sym	726,713	5,080,961	14	10	17	5	28,253
tmt_unsym	917,825	4,584,801	23	13	18	6	32,947
ecology2	999,999	4,995,991	18	12	18	6	34,080
thermal2	1,228,045	8,580,313	18	14	20	26	40,098
atmosmodj	1,270,432	8,814,880	†	8	17	7	76,368
G3_circuit	1,585,478	7,660,826	25	12	35	8	71,385
Transport	1,602,111	23,487,281	18	10	98	9	76,800
memchip	2,707,524	13,343,948	†	15	†	36	57,942
circuit5M_dc	3,523,317	14,865,409	†	5		7	8,629

TABLE 1

Preconditioner comparison. Iteration counts are reported. M_{deflated}^{-1} is the restricted two-level overlapping Schwarz preconditioner as in (2.4). No value denotes iteration count exceeds 100. † denotes either a failure in constructing the preconditioner or a breakdown in GMRES. ‡ denotes the problem is complex-valued and the preconditioner is unavailable. Matrix identifiers that are emphasized correspond to symmetric positive definite matrices, otherwise, matrices are non-self-adjoint.

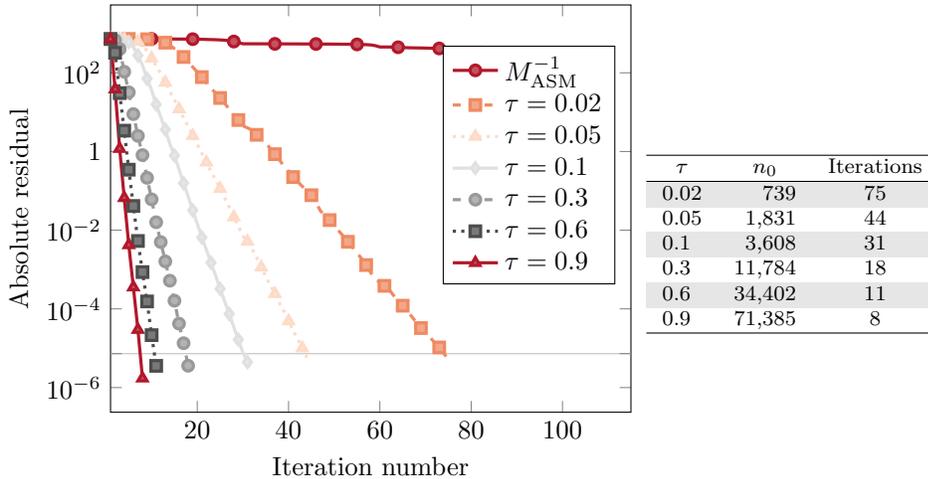


FIG. 2. Influence of the threshold parameter τ on the convergence of preconditioned GMRES for problem *G3_circuit* ($m = 1,585,478$).

the performance of our two-level preconditioner. Figure 2 depicts the history of the residual norm for different values of τ demonstrating the adaptiveness of the constructed coarse space. The table on the right of the figure prints the size of the coarse space corresponding to each selected threshold value τ along with the iteration count required to reach an unpreconditioned relative residual norm lower than 10^{-8} .

Dimension	k	N	n	ν				
				1	10^{-1}	10^{-2}	10^{-3}	10^{-4}
2	1	1,024	$6.3 \cdot 10^6$	23 <small>(52,875)</small>	20 <small>(52,872)</small>	19 <small>(52,759)</small>	20 <small>(47,497)</small>	21 <small>(28,235)</small>
3	2	4,096	$8.1 \cdot 10^6$	18 <small>(1.8 · 10⁵)</small>	14 <small>(1.8 · 10⁵)</small>	11 <small>(1.6 · 10⁵)</small>	16 <small>(97,657)</small>	29 <small>(76,853)</small>

TABLE 2

Iteration counts of the proposed preconditioner for solving the two- and three-dimensional convection-diffusion problem from (4.1) with order k Lagrange finite element space. The number of subdomains is N and the size of the discrete system is n . After each iteration count for each ν , the size of the coarse space problem is typeset between parentheses.

4.2. Linear systems arising from a PDE. A convection-diffusion problem will now be investigated. It reads:

$$(4.1) \quad \begin{aligned} \nabla \cdot (Vu) - \nu \nabla \cdot (\kappa \nabla u) &= 0 \text{ in } \Omega \\ u &= 0 \text{ in } \Gamma_0 \\ u &= 1 \text{ in } \Gamma_1. \end{aligned}$$

The problem is SUPG-stabilized [16] and discretized by FreeFEM [28]. It is important to keep in mind that the proposed preconditioner is algebraic, thus there is no specific transfer of information from the discretization kernel to the solver backend. Still, as the yielded linear systems are obtained from a standard finite element discretization, we can relax the selection parameters from those used in Table 1 and set $\tau = 0.3$ and at most 60 eigenmodes computed per subdomain (instead of the previous values $\tau = 0.6$ and 300 eigenmodes). The domain Ω is either the unit square or the unit cube meshed semi-structurally to account for boundary layers, see an example of such a mesh in Figure 3a. The value of ν is constant in Ω . The value of κ is given in Figure 3b. The value of the velocity field V is either:

$$V(x, y) = \begin{pmatrix} x(1-x)(2y-1) \\ -y(1-y)(2x-1) \end{pmatrix} \quad \text{or} \quad V(x, y, z) = \begin{pmatrix} 2x(1-x)(2y-1)z \\ -y(1-y)(2x-1) \\ -z(1-z)(2x-1)(2y-1) \end{pmatrix},$$

in 2D and 3D, respectively. These are standard values taken from the literature [43]. The definition of Γ_0 and Γ_1 may be inferred by looking at the two- and three-dimensional solutions in Figures 3c to 3e and Figures 3f to 3h, respectively. The iteration counts reported in Table 2 show that the proposed preconditioner handles this problem, even as ν tends to zero. In 2D, the operator, resp. grid, complexity is of at most 1.008, resp. 1.43. In 3D, these figures are 1.02 and 1.7, respectively.

For comparison, GAMG and BoomerAMG iteration counts are also reported in Table 3 and Table 4, respectively. For these test cases, when it comes to the algorithmic cost of the proposed methodology, we need to emphasize that when both algebraic multigrid and domain decomposition methods converge, our preconditioner is about an order of magnitude slower than multigrid (order of 10 seconds vs. order of a second). Indeed, for these configurations, subdomains are small, with around 2k (resp. 6k) unknowns per subdomain in 3D (resp. 2D). Setting up the one-level preconditioner and solving (3.1) has a negligible cost, but the assembly and exact factorization of the coarse operator become non-negligible. To alleviate this, one could think of solving inexactly the second-level system, or to reduce the threshold criterion τ , which would yield a slightly less robust but cheaper-to-assemble preconditioner. Such an in-depth investigation of the algorithmic performance of the proposed preconditioner goes beyond the scope of this paper.

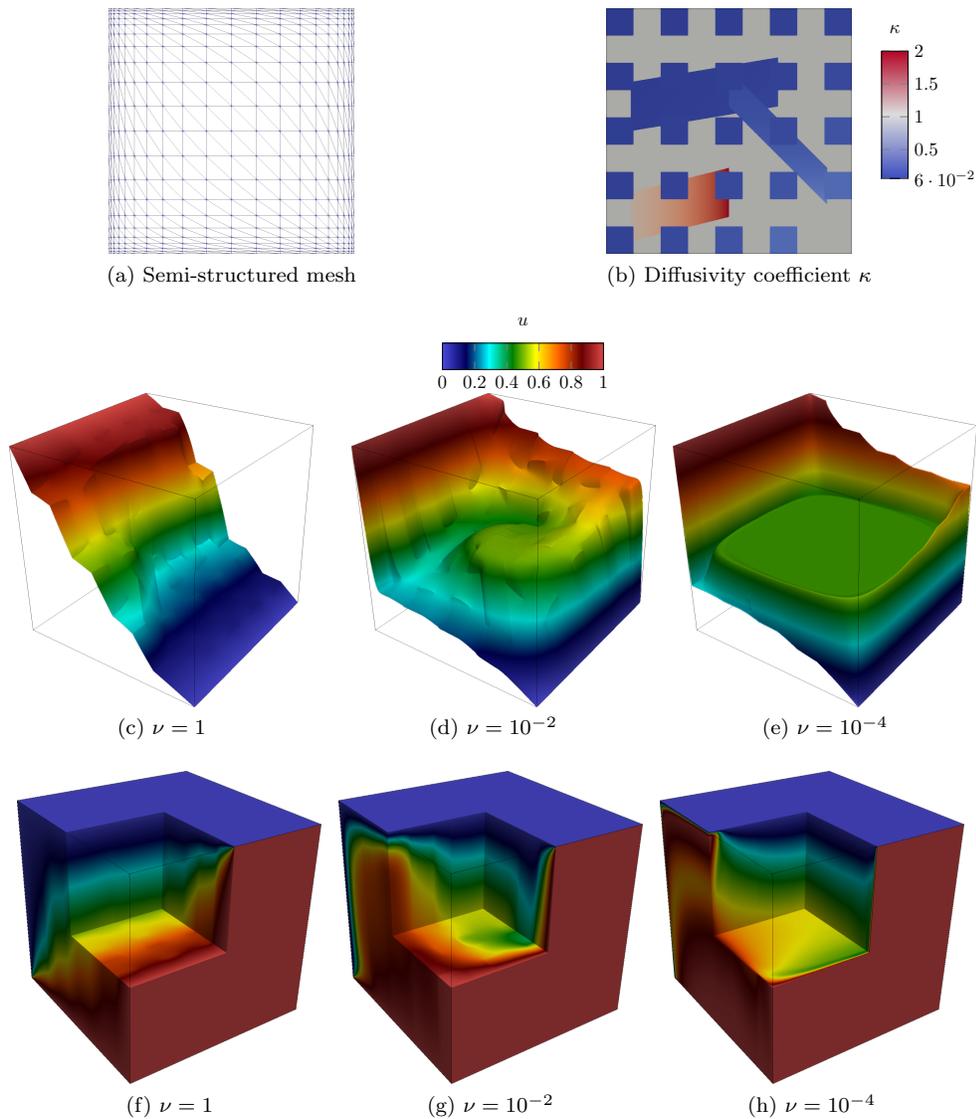


FIG. 3. (a) Mesh, (b) diffusivity coefficient, and solutions of some of the (c)–(e) two- and (f)–(h) three-dimensional test cases from Table 2.

Dimension	n	ν				
		1	10^{-1}	10^{-2}	10^{-3}	10^{-4}
2	$6.3 \cdot 10^6$	42	48	88	†	†
3	$8.1 \cdot 10^6$	40	38	65	†	†

TABLE 3

Iteration counts of GAMG for solving the two- and three-dimensional convection-diffusion problem from (4.1). † denotes either a failure to converge or a breakdown in GMRES.

Dimension	n	ν				
		1	10^{-1}	10^{-2}	10^{-3}	10^{-4}
2	$6.3 \cdot 10^6$	50	49	19	7	†
3	$8.1 \cdot 10^6$	12	9	7	†	†

TABLE 4

Iteration counts of BoomerAMG for solving the two- and three-dimensional convection-diffusion problem from (4.1). † denotes either a failure to converge or a breakdown in GMRES.

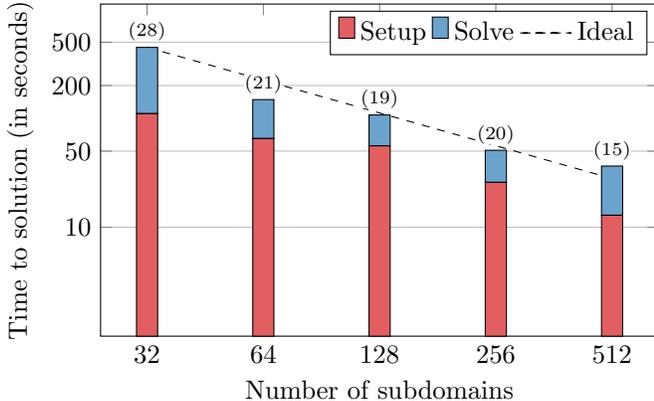


FIG. 4. Strong scaling analysis of the proposed preconditioner for solving (4.1) with $\nu = 10^{-2}$ in 3D and an increasing number of MPI processes. Numbers in parentheses are iteration counts.

Finally, we present a strong scalability experiment of the proposed preconditioner to assess its computational complexity. We consider the three-dimensional problem from Table 2 with $\nu = 10^{-2}$, see Figure 3g. Starting with $N = 32$ subdomains, we increase N by a factor of 2 up to 16 and show in Figure 4 the runtime on the first process, as reported by PETSc -log_view option. We also plot the ideal linear decrease with slope -1 . The runtime is comprised of both setup and solve times. Setup time includes generalized eigenvalue problem solving, computing and factoring the coarse space matrix, and factoring the local one-level subdomain matrix. Solve time accounts for solving the linear system by using GMRES preconditioned by the two-level method M_{deflated}^{-1} . We notice that setup times dominate. However, we observe a linear decrease in the setup and the overall runtime which behaves similarly to the ideal scaling. For a larger number of subdomains, a multilevel strategy should be employed so that setup times remain tractable, but this goes beyond the scope of this paper.

5. Conclusion. We presented in this work a fully-algebraic two-level Schwarz preconditioner for large-scale sparse matrices. The proposed preconditioner combines a classic one-level Schwarz preconditioner with a spectral coarse space. The latter is constructed efficiently by solving concurrently in each subdomain a local generalized eigenvalue problem whose pencil matrices are obtained algebraically and cheaply from the local coefficient matrix. Convergence results were obtained for diagonally dominant HPD matrices. The proposed preconditioner was compared to state-of-the-art multigrid preconditioners on a set of challenging matrices arising from a wide range of applications including a convection-dominant convection-diffusion equation. The numerical results demonstrated the effectiveness and robustness of the proposed preconditioner, especially for highly non-symmetric matrices.

Acknowledgments. We are grateful to two anonymous referees for their constructive feedback that helped us improve the clarity of the manuscript. This work was granted access to the GENCI-sponsored HPC resources of TGCC@CEA under allocation A0110607519.

Appendix A. Code reproducibility.

```

1 #include <petsc.h>

static char help[] = "Solves a linear system after having repartitioned its
symmetric part.\n\n";

int main(int argc, char **args)
6 {
    Vec          b;
    Mat          A,T,perm;
    KSP          ksp;
    IS          is,rows;
11 PetscBool    flg;
    PetscViewer  vwr;
    char        name[PETSC_MAX_PATH_LEN];
    MatPartitioning mpart;

16 PetscFunctionBeginUser;
    PetscCall(PetscInitialize(&argc,&args,NULL,help));
    PetscCall(MatCreate(PETSC_COMM_WORLD,&A));
    PetscCall(PetscOptionsGetString(NULL,NULL,"-mat_name",name,sizeof(name),&flg));
    PetscCheck(flg,PETSC_COMM_WORLD,PETSC_ERR_USER,"Missing -mat_name");
21 PetscCall(PetscViewerBinaryOpen(PETSC_COMM_WORLD,name,FILE_MODE_READ,&vwr));
    PetscCall(MatLoad(A,vwr));
    PetscCall(PetscViewerDestroy(&vwr));
    PetscCall(MatPartitioningCreate(PETSC_COMM_WORLD,&mpart));
    PetscCall(MatTranspose(A,MAT_INITIAL_MATRIX,&T));
26 PetscCall(MatDuplicate(A,MAT_COPY_VALUES,&perm));
    PetscCall(MatAXPY(perm,1.0,T,DIFFERENT_NONZERO_PATTERN));
    PetscCall(MatPartitioningSetAdjacency(mpart,perm)); // partition A^T+A
    PetscCall(MatPartitioningSetFromOptions(mpart));
    PetscCall(MatPartitioningApply(mpart,&is));
31 PetscCall(MatDestroy(&perm));
    PetscCall(MatDestroy(&T));
    PetscCall(MatPartitioningDestroy(&mpart));
    PetscCall(ISBuildTwoSided(is,NULL,&rows));
    PetscCall(ISDestroy(&is));
36 PetscCall(MatCreateSubMatrix(A,rows,rows,MAT_INITIAL_MATRIX,&perm));
    PetscCall(ISDestroy(&rows));
    PetscCall(MatHeaderReplace(A,&perm)); // only keep the permuted matrix
    PetscCall(KSPCreate(PETSC_COMM_WORLD,&ksp));
    PetscCall(KSPSetOperators(ksp,A,A));
41 PetscCall(KSPSetFromOptions(ksp)); // parse command-line options
    PetscCall(MatCreateVecs(A,NULL,&b)); // vector with a compatible dimension
    PetscCall(PetscOptionsGetString(NULL,NULL,"-rhs_name",name,sizeof(name),&flg));
    if (!flg) PetscCall(VecSetRandom(b,NULL)); // random right-hand side
46 else {
    PetscCall(PetscViewerBinaryOpen(PETSC_COMM_WORLD,name,FILE_MODE_READ,&vwr));
    PetscCall(VecLoad(b,vwr)); // right-hand side from file
    PetscCall(PetscViewerDestroy(&vwr));
}
    PetscCall(KSPSolve(ksp,b,b)); // in-place solve
51 PetscCall(KSPDestroy(&ksp));
    PetscCall(VecDestroy(&b));
    PetscCall(MatDestroy(&A));
    PetscCall(PetscFinalize());
    return 0;
56 }

```

REFERENCES

- [1] M. F. ADAMS, H. H. BAYRAKTAR, T. M. KEAVENY, AND P. PAPADOPOULOS, *Ultrascaleable 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, SC04, IEEE Computer Society, 2004, pp. 34:1–34:15.
- [2] 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.
- [3] H. AL DAAS, L. GRIGORI, P. JOLIVET, AND P.-H. TOURNIER, *A multilevel Schwarz preconditioner based on a hierarchy of robust coarse spaces*, SIAM Journal on Scientific

- Computing, 43 (2021), pp. A1907–A1928.
- [4] H. AL DAAS AND P. JOLIVET, *A robust algebraic multilevel domain decomposition preconditioner for sparse symmetric positive definite matrices*, SIAM Journal on Scientific Computing, 44 (2022), pp. A2582–A2598.
 - [5] H. AL DAAS, P. JOLIVET, AND J. A. SCOTT, *A robust algebraic domain decomposition preconditioner for sparse normal equations*, SIAM Journal on Scientific Computing, 44 (2022), pp. A1047–A1068.
 - [6] S. BALAY, S. ABHYANKAR, M. F. ADAMS, S. BENSON, J. BROWN, P. BRUNE, K. BUSCHELMAN, E. M. CONSTANTINESCU, L. DALCIN, A. DENER, V. EIJKHOUT, W. D. GROPP, V. HAPLA, T. ISAAC, P. JOLIVET, D. KARPEEV, D. KAUSHIK, M. G. KNEPLEY, F. KONG, S. KRUGER, D. A. MAY, L. C. MCINNES, R. T. MILLS, L. MITCHELL, T. MUNSON, J. E. ROMAN, K. RUPP, P. SANAN, J. SARICH, B. F. SMITH, S. ZAMPINI, H. ZHANG, H. ZHANG, AND J. ZHANG, *PETSc web page*, 2021, <https://petsc.org>.
 - [7] P. BASTIAN, R. SCHEICHL, L. SEELINGER, AND A. STREHLOW, *Multilevel spectral domain decomposition*, 2021, <https://arxiv.org/abs/2106.06404>.
 - [8] M. BONAZZOLI, X. CLAEYS, F. NATAF, AND P.-H. TOURNIER, *Analysis of the SORAS domain decomposition preconditioner for non-self-adjoint or indefinite problems*, Journal of Scientific Computing, 89 (2021), p. 19.
 - [9] N. BOOTLAND AND V. DOLEAN, *On the Dirichlet-to-Neumann coarse space for solving the Helmholtz problem using domain decomposition*, in Numerical Mathematics and Advanced Applications ENUMATH 2019, F. J. Vermolen and C. Vuik, eds., Cham, 2021, Springer International Publishing, pp. 175–184.
 - [10] N. BOOTLAND, V. DOLEAN, I. G. GRAHAM, C. MA, AND R. SCHEICHL, *GenEO coarse spaces for heterogeneous indefinite elliptic problems*, 2021, <https://arxiv.org/abs/2103.16703>.
 - [11] N. BOOTLAND, V. DOLEAN, I. G. GRAHAM, C. MA, AND R. SCHEICHL, *Overlapping Schwarz methods with GenEO coarse spaces for indefinite and non-self-adjoint problems*, IMA Journal of Numerical Analysis, (2022), <https://doi.org/10.1093/imanum/drac036>.
 - [12] N. BOOTLAND, V. DOLEAN, P. JOLIVET, F. NATAF, S. OPERTO, AND T. P.-H., *Several ways to achieve robustness when solving wave propagation problems*, 2021, <https://arxiv.org/abs/2103.06025>.
 - [13] N. BOOTLAND, V. DOLEAN, P. JOLIVET, AND P.-H. TOURNIER, *A comparison of coarse spaces for Helmholtz problems in the high frequency regime*, Computers & Mathematics with Applications, 98 (2021), pp. 239–253.
 - [14] N. BOOTLAND, V. DOLEAN, F. NATAF, AND P.-H. TOURNIER, *Two-level DDM preconditioners for positive Maxwell equations*, 2020, <https://arxiv.org/abs/2012.02388>.
 - [15] M. BREZINA, C. HEBERTON, J. MANDEL, AND P. VAN, *An iterative method with convergence rate chosen a priori*, Tech. Report 140, University of Colorado, Denver, 1999.
 - [16] A. N. BROOKS AND T. J. HUGHES, *Streamline upwind/Petrov-Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier–Stokes equations*, Computer methods in applied mechanics and engineering, 32 (1982), pp. 199–259.
 - [17] 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.
 - [18] X.-C. CAI AND O. B. WIDLUND, *Domain decomposition algorithms for indefinite elliptic problems*, SIAM Journal on Scientific and Statistical Computing, 13 (1992), pp. 243–258.
 - [19] T. CHARTIER, R. D. FALGOUT, V. E. HENSON, J. JONES, T. MANTEUFFEL, S. MCCORMICK, J. RUGE, AND P. S. VASSILEVSKI, *Spectral AMG ϵ (ρ AMG ϵ)*, SIAM Journal on Scientific Computing, 25 (2003), pp. 1–26.
 - [20] T. A. DAVIS AND Y. HU, *The University of Florida sparse matrix collection*, ACM Transactions on Mathematical Software, 38 (2011), pp. 1–28.
 - [21] V. DOLEAN, P. JOLIVET, AND F. NATAF, *An introduction to domain decomposition methods. Algorithms, theory, and parallel implementation*, Society for Industrial and Applied Mathematics, 2015.
 - [22] V. DOLEAN, P. JOLIVET, P.-H. TOURNIER, L. COMBE, S. OPERTO, AND S. RIFFO, *Large-scale finite-difference and finite-element frequency-domain seismic wave modelling with multi-level domain-decomposition preconditioner*, 2021, <https://arxiv.org/abs/2103.14921>.
 - [23] I. S. DUFF, A. M. ERISMAN, AND J. K. REID, *Direct Methods for Sparse Matrices, Second Edition*, Oxford University Press, London, 2017.
 - [24] Y. EFENDIEV, J. GALVIS, AND P. S. VASSILEVSKI, *Spectral element agglomerate algebraic multigrid methods for elliptic problems with high-contrast coefficients*, in Domain Decomposition Methods in Science and Engineering XIX, Y. Huang, R. Kornhuber, O. Widlund, and J. Xu, eds., Berlin, Heidelberg, 2011, Springer Berlin Heidelberg, pp. 407–

- 414.
- [25] S. C. EISENSTAT, H. C. ELMAN, AND M. H. SCHULTZ, *Variational iterative methods for nonsymmetric systems of linear equations*, SIAM Journal on Numerical Analysis, 20 (1983), pp. 345–357.
 - [26] R. D. FALGOUT AND U. M. YANG, *hypr: a library of high performance preconditioners*, Computational Science—ICCS 2002, (2002), pp. 632–641.
 - [27] M. J. GANDER AND A. LONELAND, *SHEM: an optimal coarse space for RAS and its multiscale approximation*, in Domain Decomposition Methods in Science and Engineering XXIII, C.-O. Lee, X.-C. Cai, D. E. Keyes, H. H. Kim, A. Klawonn, E.-J. Park, and O. B. Widlund, eds., Cham, 2017, Springer International Publishing, pp. 313–321.
 - [28] F. HECHT, *New development in FreeFem++*, Journal of Numerical Mathematics, 20 (2012), pp. 251–265.
 - [29] A. HEINLEIN, C. HOCHMUTH, AND A. KLAWONN, *Reduced dimension GDSW coarse spaces for monolithic Schwarz domain decomposition methods for incompressible fluid flow problems*, International Journal for Numerical Methods in Engineering, 121 (2020), pp. 1101–1119.
 - [30] P. JOLIVET, J. E. ROMAN, AND S. ZAMPINI, *KSPHPDDM and PCHPDDM: extending PETSc with advanced Krylov methods and robust multilevel overlapping Schwarz preconditioners*, Computers & Mathematics with Applications, 84 (2021), pp. 277–295.
 - [31] G. KARYPIS AND V. KUMAR, *Multilevel k -way partitioning scheme for irregular graphs*, Journal of Parallel and Distributed computing, 48 (1998), pp. 96–129.
 - [32] A. KLAWONN, M. KÜHN, AND O. RHEINBACH, *Adaptive coarse spaces for FETI-DP in three dimensions*, SIAM Journal on Scientific Computing, 38 (2016), pp. A2880–A2911.
 - [33] A. KLAWONN, P. RADTKE, AND O. RHEINBACH, *FETI-DP methods with an adaptive coarse space*, SIAM Journal on Numerical Analysis, 53 (2015), pp. 297–320.
 - [34] T. V. KOLEV AND P. S. VASSILEVSKI, *AMG by element agglomeration and constrained energy minimization interpolation*, Numerical Linear Algebra with Applications, 13 (2006), pp. 771–788.
 - [35] F. KONG AND X.-C. CAI, *A scalable nonlinear fluid–structure interaction solver based on a Schwarz preconditioner with isogeometric unstructured coarse spaces in 3D*, Journal of Computational Physics, 340 (2017), pp. 498–518.
 - [36] T. A. MANTEUFFEL, S. MÜNZENMAIER, J. RUGE, AND B. S. SOUTHWORTH, *Nonsymmetric reduction-based algebraic multigrid*, SIAM Journal on Scientific Computing, 41 (2019), pp. S242–S268.
 - [37] T. A. MANTEUFFEL, J. RUGE, AND B. S. SOUTHWORTH, *Nonsymmetric algebraic multigrid based on local approximate ideal restriction (ℓ AIR)*, SIAM Journal on Scientific Computing, 40 (2018), pp. A4105–A4130.
 - [38] P. MARCHAND, X. CLAEYS, P. JOLIVET, F. NATAF, AND P.-H. TOURNIER, *Two-level preconditioning for h -version boundary element approximation of hypersingular operator with GenEO*, Numerische Mathematik, 146 (2020), pp. 597–628.
 - [39] F. NATAF AND P.-H. TOURNIER, *A GenEO domain decomposition method for saddle point problems*, 2021, <https://arxiv.org/abs/1911.01858>.
 - [40] F. NATAF, H. XIANG, V. DOLEAN, AND N. SPILLANE, *A coarse space construction based on local Dirichlet-to-Neumann maps*, SIAM Journal on Scientific Computing, 33 (2011), pp. 1623–1642.
 - [41] 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.
 - [42] Y. NOTAY, *An aggregation-based algebraic multigrid method*, Electronic Transactions on Numerical Analysis, 37 (2010), pp. 123–146.
 - [43] Y. NOTAY, *Aggregation-based algebraic multigrid for convection–diffusion equations*, SIAM Journal on Scientific Computing, 34 (2012), pp. A2288–A2316.
 - [44] J. W. PEARSON AND J. PESTANA, *Preconditioners for Krylov subspace methods: An overview*, GAMM-Mitteilungen, 43 (2020), p. e202000015.
 - [45] Y. SAAD, *Iterative methods for linear systems of equations: A brief historical journey*, (2019), pp. 197–216.
 - [46] Y. SAAD AND M. H. SCHULTZ, *GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM Journal on Scientific and Statistical Computing, 7 (1986), pp. 856–869.
 - [47] B. F. SMITH, P. E. BJØRSTAD, AND W. D. GROPP, *Domain decomposition: parallel multilevel methods for elliptic partial differential equations*, Cambridge University Press, 1996.
 - [48] 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.
- [49] N. SPILLANE AND D. RIXEN, *Automatic spectral coarse spaces for robust finite element tearing and interconnecting and balanced domain decomposition algorithms*, International Journal for Numerical Methods in Engineering, 95 (2013), pp. 953–990.
 - [50] J. M. TANG, R. NABBEN, C. VUIK, AND Y. A. ERLANGGA, *Comparison of two-level preconditioners derived from deflation, domain decomposition and multigrid methods*, Journal of Scientific Computing, 39 (2009), pp. 340–370.
 - [51] J. VAN LENT, R. SCHEICHL, AND I. G. GRAHAM, *Energy-minimizing coarse spaces for two-level Schwarz methods for multiscale PDEs*, Numerical Linear Algebra with Applications, 16 (2009), pp. 775–799.
 - [52] O. WIDLUND AND A. TOSELLI, *Domain decomposition methods-algorithms and theory*, Springer, 2005.
 - [53] J. XU AND L. ZIKATANOV, *Algebraic multigrid methods*, Acta Numerica, 26 (2017), pp. 591–721.
 - [54] H. ZHANG, *Algebraic multigrid methods and their applications*, PhD thesis, Pennsylvania State University, 2017.