

# *A robust algebraic domain decomposition preconditioner for sparse normal equations*

Article

Accepted Version

Al Daas, H., Jolivet, P. and Scott, J. A. ORCID:  
<https://orcid.org/0000-0003-2130-1091> (2022) A robust algebraic domain decomposition preconditioner for sparse normal equations. *SIAM Journal on Scientific Computing*, 44 (3). pp. 1047-1068. ISSN 1095-7197 doi:  
<https://doi.org/10.1137/21M1434891> Available at  
<https://centaur.reading.ac.uk/102548/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1137/21M1434891>

Publisher: Society for Industrial and Applied Mathematics

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

[www.reading.ac.uk/centaur](http://www.reading.ac.uk/centaur)

**CentAUR**

Central Archive at the University of Reading

Reading's research outputs online



40 rates. Currently, there is much less knowledge of preconditioners for LS problems  
 41 than there is for sparse symmetric linear systems and, as observed in Bru et al. [8],  
 42 “the problem of robust and efficient iterative solution of LS problems is much harder  
 43 than the iterative solution of systems of linear equations.” This is, at least in part,  
 44 because  $A$  does not have the properties of differential problems that can make standard  
 45 preconditioners effective for solving many classes of linear systems.

46 Compared with other classes of linear systems, the development of preconditioners  
 47 for sparse LS problems may be regarded as still being in its infancy. Approaches  
 48 include

- 49 • variants of block Jacobi (also known as block Cimmino) and SOR [19];
- 50 • incomplete factorizations such as incomplete Cholesky, QR, and LU  
 51 factorizations, for example, [8, 30, 38, 39];
- 52 • sparse approximate inverses [11].

53 A review and performance comparison is given in [22]. This found that, whilst none  
 54 of the approaches successfully solved all LS problems, limited memory incomplete  
 55 Cholesky factorization preconditioners appear to be the most reliable. The incomplete  
 56 factorization-based preconditioners are designed for moderate size problems because  
 57 current approaches, in general, are not suitable for parallel computers. The block  
 58 Cimmino method can be parallelized easily, however, it lacks robustness as the  
 59 iteration count to reach convergence cannot be controlled and typically increases  
 60 significantly when the number of blocks increases for a fixed problem [17]. Several  
 61 techniques have been proposed to improve the convergence of block Cimmino but  
 62 they still lack robustness [18]. Thus, we are motivated to design a new class of LS  
 63 preconditioners that are not only reliable but can also be implemented in parallel.

64 We restrict our study in this paper to the case where  $C$  is sparse. We observe  
 65 that in some practical applications the matrix  $A$  contains a small number of rows that  
 66 have many more nonzero entries than the other rows, resulting in a dense matrix  $C$ .  
 67 Several techniques, including matrix stretching and using the augmented system, have  
 68 been proposed to handle this type of problem. These result in solving a transformed  
 69 system of sparse normal equations, see for example [40] and the references therein.

70 In [3], Al Daas and Grigori presented a class of robust fully algebraic two-level  
 71 additive Schwarz preconditioners for solving SPD linear systems of equations. They  
 72 introduced the notion of an algebraic local symmetric positive semi-definite (SPSD)  
 73 splitting of an SPD matrix with respect to local subdomains. They used this splitting  
 74 to construct a class of second-level spaces that bound the spectral condition number  
 75 of the preconditioned system by a user-defined value. Unfortunately, Al Daas and  
 76 Grigori reported that for general sparse SPD matrices, constructing the splitting is  
 77 prohibitively expensive. Our interest is in examining whether the particular structure  
 78 of the normal equations matrix allows the approach to be successfully used for  
 79 preconditioning LS problems. In this paper, we show how to compute the splitting  
 80 efficiently. Based on this splitting, we apply the theory presented in [3] to construct  
 81 a two-level Schwarz preconditioner for the normal equations.

82 Note that for most existing preconditioners of the normal equations, there is  
 83 no need to form and store the normal equations matrix  $C$  explicitly. For example,  
 84 the lower triangular part of its columns can be computed one at a time, used  
 85 to perform the corresponding step of an incomplete Cholesky algorithm, and then  
 86 discarded. However, forming the normal equations matrix, even piecemeal, can entail  
 87 a significant overhead and can potentially lead to a severe loss of information in highly  
 88 ill-conditioned cases. Although building our proposed preconditioner does not need  
 89 the explicit computation of  $C$ , our parallel implementation computes it efficiently

90 and uses it to setup the preconditioner. This is mainly motivated by technical  
 91 reasons. As an example, state-of-the-art distributed-memory graph partitioners such  
 92 as ParMETIS [28] or PT-SCOTCH [36] cannot directly partition the columns of  
 93 the *rectangular* matrix  $A$ . Our numerical experiments on highly ill-conditioned LS  
 94 problems showed that forming  $C$  and using a positive diagonal shift to construct the  
 95 preconditioner had no major effect on the robustness of the resulting preconditioner.

96 This paper is organized as follows. The notation used in the manuscript is given  
 97 at the end of the introduction. In [section 2](#), we present an overview of domain  
 98 decomposition (DD) methods for a sparse SPD matrix. We present a framework for  
 99 the DD approach when applied to the sparse LS problem in [section 3](#). Afterwards, we  
 100 show how to compute the local SPSD splitting matrices efficiently and use them in line  
 101 with the theory presented in [3] to construct a robust two-level Schwarz preconditioner  
 102 for the normal equations matrix. We then discuss some technical details that clarify  
 103 how to construct the preconditioner efficiently. In [section 4](#), we briefly discuss how  
 104 the new preconditioner can be implemented on top of the PETSc library [7] and  
 105 we illustrate its effectiveness using large-scale LS problems coming from practical  
 106 applications. Finally, concluding comments are made in [section 5](#).

107 *Notation.* We end our introduction by defining notation that will be used in this  
 108 paper. Let  $1 \leq n \leq m$  and let  $A \in \mathbb{R}^{m \times n}$ . Let  $S_1 \subset \llbracket 1, m \rrbracket$  and  $S_2 \subset \llbracket 1, n \rrbracket$  be  
 109 two sets of integers.  $A(S_1, \cdot)$  is the submatrix of  $A$  formed by the rows whose indices  
 110 belong to  $S_1$  and  $A(\cdot, S_2)$  is the submatrix of  $A$  formed by the columns whose indices  
 111 belong to  $S_2$ . The matrix  $A(S_1, S_2)$  is formed by taking the rows whose indices belong  
 112 to  $S_1$  and only retaining the columns whose indices belong to  $S_2$ . The concatenation  
 113 of any two sets of integers  $S_1$  and  $S_2$  is represented by  $[S_1, S_2]$ . Note that the order  
 114 of the concatenation is important. The set of the first  $p$  positive integers is denoted  
 115 by  $\llbracket 1, p \rrbracket$ . The identity matrix of size  $n$  is denoted by  $I_n$ . We denote by  $\ker(A)$  and  
 116  $\text{range}(A)$  the null space and the range of  $A$ , respectively.

117 **2. Introduction to domain decomposition.** Throughout this section, we  
 118 assume that  $C$  is a general  $n \times n$  sparse SPD matrix. Let the nodes  $V$  in the  
 119 corresponding adjacency graph  $\mathcal{G}(C)$  be numbered from 1 to  $n$ . A graph partitioning  
 120 algorithm can be used to split  $V$  into  $N \ll n$  disjoint subsets  $\Omega_{I_i}$  ( $1 \leq i \leq N$ ) of  
 121 size  $n_{I_i}$ . These sets are called nonoverlapping subdomains. Defining an overlapping  
 122 additive Schwarz preconditioner requires overlapping subdomains. Let  $\Omega_{\Gamma_i}$  be the  
 123 subset of size  $n_{\Gamma_i}$  of nodes that are distance one in  $\mathcal{G}(C)$  from the nodes in  $\Omega_{I_i}$   
 124 ( $1 \leq i \leq N$ ). The overlapping subdomain  $\Omega_i$  is defined to be  $\Omega_i = [\Omega_{I_i}, \Omega_{\Gamma_i}]$ , with  
 125 size  $n_i = n_{\Gamma_i} + n_{I_i}$ .

126 Associated with  $\Omega_i$  is a restriction (or projection) matrix  $R_i \in \mathbb{R}^{n_i \times n}$  given by  
 127  $R_i = I_n(\Omega_i, \cdot)$ .  $R_i$  maps from the global domain to subdomain  $\Omega_i$ . Its transpose  $R_i^\top$   
 128 is a prolongation matrix that maps from subdomain  $\Omega_i$  to the global domain. The  
 129 *one-level additive Schwarz preconditioner* [16] is defined to be

$$130 \quad (2.1) \quad M_{\text{ASM}}^{-1} = \sum_{i=1}^N R_i^\top C_{ii}^{-1} R_i, \quad C_{ii} = R_i C R_i^\top.$$

131 That is,

$$132 \quad M_{\text{ASM}}^{-1} = \mathcal{R}_1 \begin{pmatrix} C_{11}^{-1} & & \\ & \ddots & \\ & & C_{NN}^{-1} \end{pmatrix} \mathcal{R}_1^\top,$$

133 where  $\mathcal{R}_1$  is the one-level interpolation operator defined by

$$134 \quad \mathcal{R}_1 : \prod_{i=1}^N \mathbb{R}^{n_i} \rightarrow \mathbb{R}^n$$

$$135 \quad (u_i)_{1 \leq i \leq N} \mapsto \sum_{i=1}^N R_i^\top u_i.$$

136 Applying this preconditioner to a vector involves solving concurrent local problems  
 137 in the overlapping subdomains. Increasing  $N$  reduces the sizes  $n_i$  of the overlapping  
 138 subdomains, leading to smaller local problems and faster computations. However,  
 139 in practice, the preconditioned system using  $M_{\text{ASM}}^{-1}$  may not be well-conditioned,  
 140 inhibiting convergence of the iterative solver. In fact, the local nature of this  
 141 preconditioner can lead to a deterioration in its effectiveness as the number of  
 142 subdomains increases because of the lack of global information from the matrix  $C$  [16,  
 143 21]. To maintain robustness with respect to  $N$ , an artificial subdomain is added to  
 144 the preconditioner (also known as second-level correction or coarse correction) that  
 145 includes global information.

146 Let  $0 < n_0 \ll n$ . If  $R_0 \in \mathbb{R}^{n_0 \times n}$  is of full row rank, the *two-level additive Schwarz*  
 147 *preconditioner* [16] is defined to be

$$148 \quad (2.2) \quad M_{\text{additive}}^{-1} = \sum_{i=0}^N R_i^\top C_{ii}^{-1} R_i = R_0^\top C_{00}^{-1} R_0 + M_{\text{ASM}}^{-1}, \quad C_{00} = R_0 C R_0^\top.$$

149 That is,

$$150 \quad M_{\text{additive}}^{-1} = \mathcal{R}_2 \begin{pmatrix} C_{00}^{-1} & & & \\ & C_{11}^{-1} & & \\ & & \ddots & \\ & & & C_{NN}^{-1} \end{pmatrix} \mathcal{R}_2^\top,$$

151 where  $\mathcal{R}_2$  is the two-level interpolation operator

$$152 \quad (2.3) \quad \mathcal{R}_2 : \prod_{i=0}^N \mathbb{R}^{n_i} \rightarrow \mathbb{R}^n$$

$$153 \quad (u_i)_{0 \leq i \leq N} \mapsto \sum_{i=0}^N R_i^\top u_i.$$

154 In the rest of this paper, we will make use of the canonical one-to-one correspondence  
 155 between  $\prod_{i=0}^N \mathbb{R}^{n_i}$  and  $\mathbb{R}^{\sum_{i=0}^N n_i}$  so that  $\mathcal{R}_2$  can be applied to vectors in  $\mathbb{R}^{\sum_{i=0}^N n_i}$ .  
 156 Observe that, because  $C$  and  $R_0$  are of full rank,  $C_{00}$  is also of full rank. For any full  
 157 rank  $R_0$ , it is possible to cheaply obtain upper bounds on the largest eigenvalue of the  
 158 preconditioned matrix, independently of  $n$  and  $N$  [3]. However, bounding the smallest  
 159 eigenvalue is highly dependent on  $R_0$ . Thus, the choice of  $R_0$  is key to obtaining a well-  
 160 conditioned system and building efficient two-level Schwarz preconditioners. Two-  
 161 level Schwarz preconditioners have been used to solve a large class of systems arising  
 162 from a range of engineering applications (see, for example, [23, 27, 29, 31, 41, 42, 45]  
 163 and references therein).

164 Following [3], we denote by  $D_i \in \mathbb{R}^{n_i \times n_i}$  ( $1 \leq i \leq N$ ) any non-negative diagonal  
 165 matrices such that

$$166 \quad (2.4) \quad \sum_{i=1}^N R_i^\top D_i R_i = I_n.$$

167 We refer to  $(D_i)_{1 \leq i \leq N}$  as an *algebraic partition of unity*. In [3], Al Daas and Grigori  
 168 show how to select local subspaces  $Z_i \in \mathbb{R}^{n_i \times p_i}$  with  $p_i \ll n_i$  ( $1 \leq i \leq N$ ) such that,  
 169 if  $R_0^\top$  is defined to be  $R_0^\top = [R_1^\top D_1 Z_1, \dots, R_N^\top D_N Z_N]$ , the spectral condition number  
 170 of the preconditioned matrix  $M_{\text{additive}}^{-1} C$  is bounded from above independently of  $N$   
 171 and  $n$ .

172 **2.1. Algebraic local SPSD splitting of an SPD matrix.** We now recall  
 173 the definition of an algebraic local SPSD splitting of an SPD matrix given in [3].  
 174 This requires some additional notation. Denote the complement of  $\Omega_i$  in  $\llbracket 1, n \rrbracket$  by  
 175  $\Omega_{ci}$ . Define restriction matrices  $R_{ci}$ ,  $R_{Ii}$ , and  $R_{\Gamma i}$  that map from the global domain  
 176 to  $\Omega_{ci}$ ,  $\Omega_{Ii}$ , and  $\Omega_{\Gamma i}$ , respectively. Reordering the matrix  $C$  using the permutation  
 177 matrix  $P_i = I_n(\llbracket \Omega_{Ii}, \Omega_{\Gamma i}, \Omega_{ci} \rrbracket, \cdot)$  gives the block tridiagonal matrix

$$178 \quad (2.5) \quad P_i C P_i^\top = \begin{pmatrix} C_{I,i} & C_{I\Gamma,i} & \\ C_{\Gamma I,i} & C_{\Gamma,i} & C_{\Gamma c,i} \\ & C_{c\Gamma,i} & C_{c,i} \end{pmatrix},$$

179 where  $C_{I,i} = R_{Ii} C R_{Ii}^\top$ ,  $C_{\Gamma I,i}^\top = C_{I\Gamma,i} = R_{Ii} C R_{\Gamma i}^\top$ ,  $C_{\Gamma,i} = R_{\Gamma i} C R_{\Gamma i}^\top$ ,  $C_{c\Gamma,i}^\top = C_{\Gamma c,i} =$   
 180  $R_{\Gamma i} C R_{ci}^\top$ , and  $C_{c,i} = R_{ci} C R_{ci}^\top$ . The first block on the diagonal corresponds to the  
 181 nodes in  $\Omega_{Ii}$ , the second block on the diagonal corresponds to the nodes in  $\Omega_{\Gamma i}$ , and  
 182 the third block on the diagonal is associated with the remaining nodes.

183 An *algebraic local SPSD splitting* of the SPD matrix  $C$  with respect to the  $i$ -th  
 184 subdomain is defined to be any SPSD matrix  $\tilde{C}_i \in \mathbb{R}^{n \times n}$  of the form

$$185 \quad P_i \tilde{C}_i P_i^\top = \begin{pmatrix} C_{I,i} & C_{I\Gamma,i} & 0 \\ C_{\Gamma I,i} & \tilde{C}_{\Gamma,i} & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

186 such that the following condition holds:

$$187 \quad 0 \leq u^\top \tilde{C}_i u \leq u^\top C u, \quad \text{for all } u \in \mathbb{R}^n.$$

We denote the  $2 \times 2$  block nonzero matrix of  $P_i \tilde{C}_i P_i^\top$  by  $\tilde{C}_{ii}$  so that

$$\tilde{C}_{ii} = R_i^\top \tilde{C}_i R_i.$$

188 Associated with the local SPSD splitting matrices, we define a multiplicity  
 189 constant  $k_m$  that satisfies the inequality

$$190 \quad (2.6) \quad 0 \leq \sum_{i=1}^N u^\top \tilde{C}_i u \leq k_m u^\top C u, \quad \text{for all } u \in \mathbb{R}^n.$$

191 Note that, for any set of SPSD splitting matrices,  $k_m \leq N$ .

192 The main motivation for defining splitting matrices is to find local seminorms that  
 193 are bounded from above by the  $C$ -norm. These seminorms will be used to determine a  
 194 subspace that contains the eigenvectors of  $C$  associated with its smallest eigenvalues.

195 **2.2. Two-level Schwarz method.** We next review the abstract theory of the  
 196 two-level Schwarz method as presented in [3]. For the sake of completeness, we present  
 197 some elementary lemmas that are widely used in multilevel methods. These will be  
 198 used in proving efficiency of the two-level Schwarz preconditioner and will also help  
 199 in understanding how the preconditioner is constructed.

200 **2.2.1. Useful lemmas.** The following lemma [34] provides a unified framework  
 201 for bounding the spectral condition number of a preconditioned operator. It can be  
 202 found in different forms for finite and infinite dimensional spaces. Here, we follow the  
 203 presentation from [16, Lemma 7.4].

204 **LEMMA 2.1 (Fictitious Subspace Lemma).** *Let  $C \in \mathbb{R}^{n_C \times n_C}$  and  $B \in \mathbb{R}^{n_B \times n_B}$*   
 205 *be SPD. Let the operator  $\mathcal{R}$  be defined as*

$$\begin{aligned} \mathcal{R} : \mathbb{R}^{n_B} &\rightarrow \mathbb{R}^{n_C} \\ v &\mapsto \mathcal{R}v, \end{aligned}$$

208 and let  $\mathcal{R}^\top$  be its transpose. Assume the following conditions hold:

- 209 (i)  $\mathcal{R}$  is surjective;  
 210 (ii) there exists  $c_u > 0$  such that for all  $v \in \mathbb{R}^{n_B}$

$$211 (\mathcal{R}v)^\top C (\mathcal{R}v) \leq c_u v^\top B v;$$

- 212 (iii) there exists  $c_l > 0$  such that for all  $v_C \in \mathbb{R}^{n_C}$  there exists  $v_B \in \mathbb{R}^{n_B}$  such that  
 213  $v_C = \mathcal{R}v_B$  and

$$214 c_l v_B^\top B v_B \leq (\mathcal{R}v_B)^\top C (\mathcal{R}v_B) = v_C^\top C v_C.$$

215 Then, the spectrum of the operator  $\mathcal{R}B^{-1}\mathcal{R}^\top C$  is contained in the interval  $[c_l, c_u]$ .

216 The challenge is to define the second-level projection matrix  $R_0$  such that the two-level  
 217 additive Schwarz preconditioner  $M_{\text{additive}}^{-1}$  and the operator  $\mathcal{R}_2$  (2.3), corresponding  
 218 respectively to  $B$  and  $\mathcal{R}$  in Lemma 2.1, satisfy conditions (i) to (iii) and, in addition,  
 219 ensures the ratio between  $c_l$  and  $c_u$  is small because this determines the quality of the  
 220 preconditioner.

221 As shown in [16, Lemmas 7.10 and 7.11], a two-level additive Schwarz  
 222 preconditioner satisfies (i) and (ii) for any full rank  $R_0$ . Furthermore, the constant  $c_u$   
 223 is bounded from above independently of the number of subdomains  $N$ , as shown in  
 224 the following result [10, Theorem 12].

225 **LEMMA 2.2.** *Let  $k_c$  be the minimum number of distinct colours so that the spaces*  
 226 *spanned by the columns of the matrices  $R_1^\top, \dots, R_N^\top$  that are of the same colour are*  
 227 *mutually  $C$ -orthogonal. Then,*

$$228 (\mathcal{R}_2 u_B)^\top C (\mathcal{R}_2 u_B) \leq (k_c + 1) \sum_{i=0}^N u_i^\top C_{ii} u_i,$$

229 for all  $u_B = (u_i)_{0 \leq i \leq N} \in \prod_{i=0}^N \mathbb{R}^{n_i}$ .

230 Note that  $k_c$  is independent of  $N$ . Indeed, it depends only on the sparsity structure  
 231 of  $C$  and is less than the maximum number of neighbouring subdomains.

232 The following result is the first step in a three-step approach to define a two-level  
 233 additive Schwarz operator  $\mathcal{R}_2$  that satisfies condition (iii) in Lemma 2.1.



234 LEMMA 2.3. [16, Lemma 7.12] Let  $u_{\mathcal{B}} = (u_i)_{0 \leq i \leq N} \in \prod_{i=0}^N \mathbb{R}^{n_i}$  and  $u = \mathcal{R}_2 u_{\mathcal{B}} \in$   
 235  $\mathbb{R}^n$ . Then, provided  $R_0$  is of full rank,

$$236 \quad \sum_{i=0}^N u_i^\top C_{ii} u_i \leq 2 u^\top C u + (2k_c + 1) \sum_{i=1}^N u_i^\top C_{ii} u_i,$$

237 where  $k_c$  is defined in Lemma 2.2.

238 It follows that (iii) is satisfied if the squared localized seminorm  $u_i^\top C_{ii} u_i$  is  
 239 bounded from above by the squared  $C$ -norm of  $u$ .

240 In the second step, we bound  $u_i^\top C_{ii} u_i$  by the squared localized seminorm defined  
 241 by the SPSD splitting matrix  $\tilde{C}_i$ , which can be bounded by the squared  $C$ -norm (2.6).  
 242 The decomposition of  $u = \sum_{i=0}^N R_i^\top u_i \in \mathbb{R}^n$  is termed *stable* if, for some  $\tau > 0$ ,

$$243 \quad \tau u_i^\top C_{ii} u_i \leq u^\top C u, \quad 1 \leq i \leq N.$$

244 The two-level approach in [3] aims to decompose each  $\mathbb{R}^{n_i}$  ( $1 \leq i \leq N$ ) into two  
 245 subspaces, one that makes the decomposition of  $u$  stable and the other is part of the  
 246 artificial subdomain associated with the second level of the preconditioner. Given the  
 247 partition of unity (2.4),  $u = \sum_{i=1}^N R_i^\top D_i R_i u$  and, if  $\Pi_i = \Pi_i^\top \in \mathbb{R}^{n_i \times n_i}$ , we can write

$$248 \quad u = \sum_{i=1}^N R_i^\top D_i (I_{n_i} - \Pi_i) R_i u + \sum_{i=1}^N R_i^\top D_i \Pi_i R_i u$$

$$249 \quad = \sum_{i=1}^N R_i^\top u_i + \sum_{i=1}^N R_i^\top D_i \Pi_i R_i u, \quad \text{with } u_i = D_i (I_{n_i} - \Pi_i) R_i u.$$

251 Therefore, we need to construct  $\Pi_i$  such that

$$252 \quad \tau u^\top R_i^\top (I_{n_i} - \Pi_i) D_i C_{ii} D_i (I_{n_i} - \Pi_i) R_i u \leq u^\top C u.$$

253 The following lemma shows how this can be done.

254 LEMMA 2.4. [3, Lemma 4.2] Let  $\tilde{C}_i = R_i^\top \tilde{C}_{ii} R_i$  be a local SPSD splitting of  
 255  $C$  related to the  $i$ -th subdomain ( $1 \leq i \leq N$ ). Let  $D_i$  be the partition of unity  
 256 (2.4). Let  $P_{0,i}$  be the projection on  $\text{range}(\tilde{C}_{ii})$  parallel to  $\ker(\tilde{C}_{ii})$ . Define  $L_i =$   
 257  $\ker(D_i C_{ii} D_i) \cap \ker(\tilde{C}_{ii})$ , and let  $L_i^\perp$  denote the orthogonal complementary of  $L_i$  in  
 258  $\ker(\tilde{C}_{ii})$ . Consider the following generalized eigenvalue problem:

$$259 \quad \text{find } (v_{i,k}, \lambda_{i,k}) \in \mathbb{R}^{n_i} \times \mathbb{R}$$

$$260 \quad \text{such that } P_{0,i} D_i C_{ii} D_i P_{0,i} v_{i,k} = \lambda_{i,k} \tilde{C}_{ii} v_{i,k}.$$

262 Given  $\tau > 0$ , define

$$263 \quad (2.7) \quad \mathcal{Z}_i = L_i^\perp \oplus \text{span} \left\{ v_{i,k} \mid \lambda_{i,k} > \frac{1}{\tau} \right\}$$

264 and let  $\Pi_i$  be the orthogonal projection on  $\mathcal{Z}_i$ . Then,  $\mathcal{Z}_i$  is the subspace of smallest  
 265 dimension such that for all  $u \in \mathbb{R}^n$ ,

$$266 \quad \tau u_i^\top C_{ii} u_i \leq u^\top \tilde{C}_i u \leq u^\top C u,$$

267 where  $u_i = D_i (I_{n_i} - \Pi_i) R_i u$ .

268 **Lemma 2.5** provides the last step that we need for condition (iii) in **Lemma 2.1**.  
 269 It defines  $u_0$  and checks whether  $(u_i)_{0 \leq i \leq N}$  is a stable decomposition.

270 **LEMMA 2.5.** *Let  $\tilde{C}_i$ ,  $Z_i$ , and  $\Pi_i$  be as in **Lemma 2.4** and let  $Z_i$  be a matrix whose  
 271 columns span  $Z_i$  ( $1 \leq i \leq N$ ). Let the columns of the matrix  $R_0^\top$  span the space*

$$272 \quad (2.8) \quad \mathcal{Z} = \bigoplus_{i=1}^N R_i^\top D_i Z_i.$$

273 Let  $u \in \mathbb{R}^n$  and  $u_i = D_i (I_{n_i} - \Pi_i) R_i u$  ( $1 \leq i \leq N$ ). Define

$$274 \quad u_0 = (R_0 R_0^\top)^{-1} R_0 \left( \sum_{i=1}^N R_i^\top D_i \Pi_i R_i u \right).$$

275 Then,

$$276 \quad u = \sum_{i=0}^N R_i^\top u_i,$$

277 and

$$278 \quad \sum_{i=0}^N u_i^\top C_{ii} u_i \leq \left( 2 + (2k_c + 1) \frac{k_m}{\tau} \right) u^\top C u.$$

279 Finally, using the preceding results, **Theorem 2.6** presents a theoretical upper  
 280 bound on the spectral condition number of the preconditioned system.

281 **THEOREM 2.6.** *If the two-level additive Schwarz preconditioner  $M_{\text{additive}}^{-1}$  (2.2) is  
 282 constructed using  $R_0$  as defined in **Lemma 2.5**, then the following inequality is  
 283 satisfied:*

$$284 \quad \kappa(M_{\text{additive}}^{-1} C) \leq (k_c + 1) \left( 2 + (2k_c + 1) \frac{k_m}{\tau} \right).$$

285 **2.3. Variants of the Schwarz preconditioner.** So far, we have  
 286 presented  $M_{\text{ASM}}^{-1}$ , the symmetric additive Schwarz method (ASM) and  $M_{\text{additive}}^{-1}$ ,  
 287 the additive correction for the second level. It was noted in [9] that using the  
 288 partition of unity to weight the preconditioner can improve its quality. The  
 289 resulting preconditioner is referred to as  $M_{\text{RAS}}^{-1}$ , the *restricted additive Schwarz* (RAS)  
 290 preconditioner, and is defined to be

$$291 \quad (2.9) \quad M_{\text{RAS}}^{-1} = \sum_{i=1}^N R_i^\top D_i C_{ii}^{-1} R_i.$$

292 This preconditioner is nonsymmetric and thus can only be used with iterative  
 293 methods such as GMRES [37] that are for solving nonsymmetric problems. With  
 294 regards to the second level, different strategies yield either a symmetric or a  
 295 nonsymmetric preconditioner [44]. Given a first-level preconditioner  $M_*^{-1}$  and setting  
 296  $Q = R_0^\top C_{00}^{-1} R_0$ , the balanced and deflated two-level preconditioners are as follows

$$297 \quad (2.10) \quad M_{\text{balanced}}^{-1} = Q + (I - CQ)^\top M_*^{-1} (I - CQ),$$

298 and

$$299 \quad (2.11) \quad M_{\text{deflated}}^{-1} = Q + M_*^{-1} (I - CQ),$$

300 respectively. It is well-known in the literature that  $M_{\text{balanced}}^{-1}$  and  $M_{\text{deflated}}^{-1}$  yield better  
 301 convergence behavior than  $M_{\text{additive}}^{-1}$  (see [44] for a thorough comparison). Although the  
 302 theory we present relies on  $M_{\text{additive}}^{-1}$ , in practice we will use  $M_{\text{balanced}}^{-1}$  and  $M_{\text{deflated}}^{-1}$ . If  
 303 the one-level preconditioner  $M_{\text{+}}^{-1}$  is symmetric, then so is  $M_{\text{balanced}}^{-1}$ , while  $M_{\text{deflated}}^{-1}$   
 304 is typically nonsymmetric. For this reason, in the rest of the paper, we always  
 305 couple  $M_{\text{ASM}}^{-1}$  with  $M_{\text{balanced}}^{-1}$ , and  $M_{\text{RAS}}^{-1}$  with  $M_{\text{deflated}}^{-1}$ . All three variants have the same  
 306 setup cost, and only differ in how the second level is applied.  $M_{\text{balanced}}^{-1}$  is slightly more  
 307 expensive because two second-level corrections (multiplications by  $Q$ ) are required  
 308 instead of a single one for  $M_{\text{additive}}^{-1}$  and  $M_{\text{deflated}}^{-1}$ .

309 **3. The normal equations.** The theory explained thus far is fully algebraic but  
 310 somewhat disconnected from our initial LS problem (1.1). We now show how it can  
 311 be readily applied to the normal equations matrix  $C = A^T A$ , with  $A \in \mathbb{R}^{m \times n}$  sparse,  
 312 first defining a one-level Schwarz preconditioner, and then a robust algebraic second-  
 313 level correction. We start by partitioning the  $n$  columns of  $A$  into disjoint subsets  
 314  $\Omega_{I_i}$ . Let  $\Xi_i$  be the set of indices of the nonzero rows in  $A(:, \Omega_{I_i})$  and let  $\Xi_{c_i}$  be the  
 315 complement of  $\Xi_i$  in the set  $\llbracket 1, m \rrbracket$ . Now define  $\Omega_{\Gamma_i}$  to be the complement of  $\Omega_{I_i}$  in  
 316 the set of indices of nonzero columns of  $A(\Xi_i, :)$ . The set  $\Omega_i = [\Omega_{I_i}, \Omega_{\Gamma_i}]$  defines the  
 317  $i$ -th overlapping subdomain and we have the permuted matrix

$$318 \quad (3.1) \quad A([\Xi_i, \Xi_{c_i}], [\Omega_{I_i}, \Omega_{\Gamma_i}, \Omega_{c_i}]) = \begin{pmatrix} A_{I,i} & A_{\Gamma,i} & \\ & A_{\Gamma,i} & A_{c,i} \end{pmatrix}.$$

319 To illustrate the concepts and notation, consider the  $5 \times 4$  matrix

$$320 \quad A = \begin{pmatrix} 1 & 0 & 6 & 0 \\ 2 & 4 & 0 & 0 \\ 3 & 0 & 0 & 0 \\ 0 & 5 & 0 & 7 \\ 0 & 0 & 0 & 8 \end{pmatrix}$$

321 and set  $N = 2$ ,  $\Omega_{I_1} = \{1, 3\}$ ,  $\Omega_{I_2} = \{2, 4\}$ . Consider the first subdomain. We have

$$322 \quad A(:, \Omega_{I_1}) = \begin{pmatrix} 1 & 6 \\ 2 & 0 \\ 3 & 0 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

323 The set of indices of the nonzero rows is  $\Xi_1 = \{1, 2, 3\}$ , and its complement is  $\Xi_{c_1} =$   
 324  $\{4, 5\}$ . To define  $\Omega_{\Gamma,1}$ , select the nonzero columns in the submatrix  $A(\Xi_1, :)$  and  
 325 remove those already in  $\Omega_{I_1}$ , that is,

$$326 \quad (3.2) \quad A(\Xi_1, :) = \begin{pmatrix} 1 & 0 & 6 & 0 \\ 2 & 4 & 0 & 0 \\ 3 & 0 & 0 & 0 \end{pmatrix},$$

327 so that  $\Omega_{\Gamma_1} = \{2\}$  and  $\Omega_{c_1} = \{4\}$ . Permuting  $A$  to the form (3.1) gives

$$328 \quad A([\Xi_1, \Xi_{c_1}], [\Omega_{I_1}, \Omega_{\Gamma_1}, \Omega_{c_1}]) = \begin{pmatrix} 1 & 6 & 0 & 0 \\ 2 & 0 & 4 & 0 \\ 3 & 0 & 0 & 0 \\ 0 & 0 & 5 & 7 \\ 0 & 0 & 0 & 8 \end{pmatrix}.$$

329 In the same way, consider the second subdomain.  $\Omega_{I2} = \{2, 4\}$  and

$$330 \quad A(:, \Omega_{I2}) = \begin{pmatrix} 0 & 0 \\ 4 & 0 \\ 0 & 0 \\ 5 & 7 \\ 0 & 8 \end{pmatrix},$$

331 so that  $\Xi_2 = \{2, 4, 5\}$  and  $\Xi_{c2} = \{1, 3\}$ . To define  $\Omega_{\Gamma2}$ , select the nonzero columns in  
332 the submatrix  $A(\Xi_2, :)$  and remove those already in  $\Omega_{I2}$ , that is,

$$333 \quad (3.3) \quad A(\Xi_2, :) = \begin{pmatrix} 2 & 4 & 0 & 0 \\ 0 & 5 & 0 & 7 \\ 0 & 0 & 0 & 8 \end{pmatrix},$$

334 which gives  $\Omega_{\Gamma2} = \{1\}$  and  $\Omega_{c2} = \{3\}$ . Permuting  $A$  to the form (3.1) gives

$$335 \quad A([\Xi_2, \Xi_{c2}], [\Omega_{I2}, \Omega_{\Gamma2}, \Omega_{c2}]) = \begin{pmatrix} 4 & 0 & 2 & 0 \\ 5 & 7 & 0 & 0 \\ 0 & 8 & 0 & 0 \\ 0 & 0 & 1 & 6 \\ 0 & 0 & 3 & 0 \end{pmatrix}.$$

336 Now that we have  $\Omega_{Ii}$  and  $\Omega_{\Gamma i}$ , we can define the restriction operators

$$337 \quad R_1 = I_4(\Omega_1, :) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}, \quad R_2 = I_4(\Omega_2, :) = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}.$$

338 For our example,  $n_{I1} = n_{I2} = 2$  and  $n_{\Gamma1} = n_{\Gamma2} = 1$ . The partition of unity  
339 matrices  $D_i$  are of dimension  $(n_{Ii} + n_{\Gamma i}) \times (n_{Ii} + n_{\Gamma i})$  ( $i = 1, 2$ ) and have ones  
340 on the  $n_{Ii}$  leading diagonal entries and zeros elsewhere, so that

$$341 \quad (3.4) \quad D_1 = D_2 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

342 Observe that  $D_i(k, k)$  scales the columns  $A(:, \Omega_i(k))$ .

343 Note that it is possible to obtain the partitioning sets and the sets of indices  
344 using the normal equations matrix  $C$ . Most graph partitioners, especially those that  
345 are implemented in parallel, require an undirected graph (corresponding to a square  
346 matrix with a symmetric sparsity pattern). Therefore, in practice, we use the graph  
347 of  $C$  to setup the first-level preconditioner for LS problems.

348 **3.1. One-level DD for the normal equations.** This section presents the  
349 one-level additive Schwarz preconditioner for the normal equations matrix  $C =$   
350  $A^\top A$ . Following (2.1) and given the sets  $\Omega_{Ii}$ ,  $\Omega_{\Gamma i}$ , and  $\Xi_i$ , the one-level Schwarz  
351 preconditioner of  $C = A^\top A$  is

$$352 \quad M_{\text{ASM}}^{-1} = \sum_{i=1}^N R_i^\top (R_i A^\top A R_i^\top)^{-1} R_i,$$

$$353 \quad = \sum_{i=1}^N R_i^\top (A(:, \Omega_i)^\top A(:, \Omega_i))^{-1} R_i,$$

354

355 *Remark 3.1.* Note that the local matrix  $C_{ii} = A(:, \Omega_i)^\top A(:, \Omega_i)$  need not be  
 356 computed explicitly to be factored. Instead, the Cholesky factor of  $C_{ii}$  can be  
 357 computed by using a “thin” QR factorization of  $A(:, \Omega_i)$ .

358 **3.2. Algebraic local SPSD splitting of the normal equations matrix.** In  
 359 this section, we show how to cheaply construct algebraic local SPSD splittings for  
 360 sparse matrices of the form  $C = A^\top A$ . Combining (2.5) and (3.1), we can write

$$361 \quad P_i A^\top A P_i^\top = \begin{pmatrix} A_{I,i}^\top A_{I,i} & A_{I,i}^\top A_{I\Gamma,i} & \\ A_{I\Gamma,i}^\top A_{I,i} & A_{I\Gamma,i}^\top A_{I\Gamma,i} + A_{\Gamma,i}^\top A_{\Gamma,i} & A_{\Gamma,i}^\top A_{c,i} \\ & A_{c,i}^\top A_{\Gamma,i} & A_{c,i}^\top A_{c,i} \end{pmatrix},$$

362 where  $P_i = I_n([\Omega_{Ii}, \Omega_{\Gamma i}, \Omega_{ci}], :)$  is a permutation matrix. A straightforward splitting  
 363 of  $P_i A^\top A P_i^\top$  is given by

$$364 \quad P_i A^\top A P_i^\top = \begin{pmatrix} A_{I,i}^\top A_{I,i} & A_{I,i}^\top A_{I\Gamma,i} & 0 \\ A_{I\Gamma,i}^\top A_{I,i} & A_{I\Gamma,i}^\top A_{I\Gamma,i} & 0 \\ 0 & 0 & 0 \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 \\ 0 & A_{\Gamma,i}^\top A_{\Gamma,i} & A_{\Gamma,i}^\top A_{c,i} \\ 0 & A_{c,i}^\top A_{\Gamma,i} & A_{c,i}^\top A_{c,i} \end{pmatrix}.$$

366 It is clear that both summands are SPSD. Indeed, they both have the form  $X^\top X$ ,  
 367 where  $X$  is  $(A_{I,i} \ A_{I\Gamma,i} \ 0)$  and  $(0 \ A_{\Gamma,i} \ A_{c,i})$ , respectively. The local SPSD  
 368 splitting matrix related to the  $i$ -th subdomain is then defined as:

$$369 \quad (3.5) \quad \tilde{C}_{ii} = A(\Xi_i, \Omega_i)^\top A(\Xi_i, \Omega_i) = (A_{I,i} \ A_{I\Gamma,i})^\top (A_{I,i} \ A_{I\Gamma,i}),$$

and

$$\tilde{C}_i = R_i^\top \tilde{C}_{ii} R_i = A(\Xi_i, :)^\top A(\Xi_i, :).$$

Hence, the theory presented in [3] and summarised in subsection 2.2 is applicable. In  
 particular, the two-level Schwarz preconditioner  $M_{\text{additive}}^{-1}$  (2.2) satisfies

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

371 where  $k_c$  is the minimal number of colours required to colour the partitions of  $C$   
 372 such that each two neighbouring subdomains have different colours, and  $k_m$  is the  
 373 multiplicity constant that satisfies the following inequality

$$374 \quad \sum_{i=1}^N R_i^\top \tilde{C}_{ii} R_i \leq k_m C.$$

375 The constant  $k_c$  is independent of  $N$  and depends only on the graph  $\mathcal{G}(C)$ , which is  
 376 determined by the sparsity pattern of  $A$ . The multiplicity constant  $k_m$  depends on  
 377 the local SPSD splitting matrices. For the normal equations matrix, the following  
 378 lemma provides an upper bound on  $k_m$ .

379 **LEMMA 3.2.** *Let  $C = A^\top A$ . Let  $m_j$  be the number of subdomains such that*  
 380  *$A(j, \Omega_{Ii}) \neq 0$  ( $1 \leq i \leq N$ ), that is,*

$$381 \quad m_j = \#\{i \mid j \in \Xi_i\}.$$

*Then,  $k_m$  can be chosen to be  $k_m = \max_{1 \leq j \leq m} m_j$ . Furthermore, if  $k_{\Omega_i}$  is the number  
 of neighbouring subdomains of the  $i$ -th subdomain, that is,*

$$k_{\Omega_i} = \#\{j \mid \Omega_i \cap \Omega_j \neq \emptyset\},$$

382 then

$$383 \quad k_m = \max_{1 \leq j \leq m} m_j \leq \max_{1 \leq i \leq N} k_{\Omega_i}.$$

384 *Proof.* Since  $C = A^\top A$  and  $\tilde{C}_i = A(\Xi_i, \cdot)^\top A(\Xi_i, \cdot)$ , we have

$$385 \quad u^\top C u = \sum_{j=1}^m u^\top A(j, \cdot)^\top A(j, \cdot) u,$$

$$386 \quad u^\top \tilde{C}_i u = \sum_{j \in \Xi_i} u^\top A(j, \cdot)^\top A(j, \cdot) u,$$

$$387 \quad \sum_{i=1}^N u^\top \tilde{C}_i u = \sum_{i=1}^N \sum_{j \in \Xi_i} u^\top A(j, \cdot)^\top A(j, \cdot) u.$$

389 From the definition of  $m_j$ , the term  $u^\top A(j, \cdot)^\top A(j, \cdot) u$  appears  $m_j$  times in the last  
390 equation. Thus,

$$391 \quad \sum_{i=1}^N u^\top \tilde{C}_i u = \sum_{j=1}^m m_j u^\top A(j, \cdot)^\top A(j, \cdot) u,$$

$$392 \quad \leq \max_{1 \leq j \leq m} m_j \sum_{j=1}^m u^\top A(j, \cdot)^\top A(j, \cdot) u,$$

$$393 \quad = \max_{1 \leq j \leq m} m_j (u^\top C u),$$

395 from which it follows that we can choose  $k_m = \max_{1 \leq j \leq m} m_j$ . Now, if  $1 \leq l \leq m$ ,  
396 there exist  $i_1, \dots, i_{m_l}$  such that  $l \in \Xi_{i_1} \cap \dots \cap \Xi_{i_{m_l}}$ . Furthermore,  $m_l \leq \max_{1 \leq p \leq l} k_{\Omega_{i_p}}$ .  
397 Taking the maximum over  $l$  on both sides, we obtain

$$398 \quad k_m \leq \max_{1 \leq i \leq N} k_{\Omega_i}. \quad \square$$

399 Note that because  $A$  is sparse,  $k_m$  is independent of the number of subdomains.

400 **3.3. Algorithms and technical details.** In this section, we discuss the  
401 technical details involved in constructing a two-level preconditioner for the normal  
402 equations matrix.

403 **3.3.1. Partition of unity.** Because the matrix  $A_{\Gamma, i}$  may be of low rank, the  
404 null space of  $\tilde{C}_{ii}$  (3.5) can be large. Recall that the diagonal matrices  $D_i$  have  
405 dimension  $n_i = n_{I_i} + n_{\Gamma_i}$ . Choosing the entries in positions  $n_{I_i} + 1, \dots, n_i$  of the  
406 diagonal of  $D_i$  to be zero, as in (3.4), results in the subspace of  $\ker(\tilde{C}_{ii})$  caused  
407 by the rank deficiency of  $A_{\Gamma, i}$  to lie within  $\ker(D_i C_{ii} D_i)$ , reducing the size of the  
408 space  $\mathcal{Z}$  given by (2.8). In other words, if  $A_{\Gamma, i} u = 0$ , we have  $\tilde{C}_{ii} v = 0$ , where  
409  $v^\top = (0, u^\top)$ , i.e.,  $v \in \ker(\tilde{C}_{ii})$  and because by construction  $D_i v = 0$ , we have  
410  $v \in \ker(\tilde{C}_{ii}) \cap \ker(D_i C_{ii} D_i)$ , therefore,  $v$  need not be included in  $\mathcal{Z}_i$ .

411 **3.3.2. The eigenvalue problem.** The generalized eigenvalue problem  
412 presented in Lemma 2.4 is critical in the construction of the two-level preconditioner.  
413 Although the definition of  $\mathcal{Z}_i$  from (2.7) suggests it is necessary to compute the null

414 space of  $\tilde{C}_{ii}$  and that of  $D_i C_{ii} D_i$  and their intersection, in practice, this can be  
 415 avoided. Consider the generalized eigenvalue problem

$$416 \quad (3.6) \quad D_i C_{ii} D_i v = \lambda \tilde{C}_{ii} v,$$

417 where, by convention, we set  $\lambda = 0$  if  $v \in \ker(\tilde{C}_{ii}) \cap \ker(D_i C_{ii} D_i)$  and  $\lambda = \infty$  if  
 418  $v \in \ker(\tilde{C}_{ii}) \setminus \ker(D_i C_{ii} D_i)$ . The subspace  $\mathcal{Z}_i$  defined in (2.7) can then be written as

$$419 \quad \text{span} \left\{ v \mid D_i C_{ii} D_i v = \lambda \tilde{C}_{ii} v \text{ and } \lambda > \frac{1}{\tau} \right\}.$$

420 Consider also the shifted generalized eigenvalue problem

$$421 \quad (3.7) \quad D_i C_{ii} D_i v = \lambda(\tilde{C}_{ii} + sI_{n_i})v,$$

422 where  $0 < s \ll 1$ . Note that if  $s$  is such that  $\tilde{C}_{ii} + sI_{n_i}$  is numerically of full rank,  
 423 (3.7) can be solved using any off-the-shelf generalized eigenproblem solver. Let  $(v, \lambda)$   
 424 be an eigenpair of (3.7). Then, we can only have one of the following situations:

- 425 •  $v \in \text{range}(\tilde{C}_{ii}) \cap \ker(D_i C_{ii} D_i)$  or  $v \in \ker(\tilde{C}_{ii}) \cap \ker(D_i C_{ii} D_i)$ . In which  
 426 case,  $(v, 0)$  is an eigenpair of (3.6).
- 427 •  $v \in \text{range}(\tilde{C}_{ii}) \cap \text{range}(D_i C_{ii} D_i)$ . Then,

$$428 \quad \frac{\|D_i C_{ii} D_i v - \lambda \tilde{C}_{ii} v\|_2}{\lambda \|v\|_2} = s,$$

429 and, as  $s$  is small,  $(v, \lambda)$  is a good approximation of an eigenpair of (3.6)  
 430 corresponding to a finite eigenvalue.

- 431 •  $v \in \ker(\tilde{C}_{ii}) \cap \text{range}(D_i C_{ii} D_i)$ . Then,  $D_i C_{ii} D_i v = \lambda s v$ , i.e.,  $\lambda s$  is a nonzero  
 432 eigenvalue of  $D_i C_{ii} D_i$ . Because  $D_i$  is defined such that the diagonal values  
 433 corresponding to the boundary nodes are zero, the nonzero eigenvalues of  
 434  $D_i C_{ii} D_i$  correspond to the squared singular values of  $A(:, \Omega_{I_i})$ . Hence, all  
 435 the eigenpairs of (3.6) corresponding to an infinite eigenvalue are included in  
 436 the set of eigenpairs  $(v, \lambda)$  of (3.7) such that

$$437 \quad (3.8) \quad \sigma_{\min}^2(A(:, \Omega_{I_i})) \leq \lambda s \leq \sigma_{\max}^2(A(:, \Omega_{I_i})),$$

438 where  $\sigma_{\min}(A(:, \Omega_{I_i}))$  and  $\sigma_{\max}(A(:, \Omega_{I_i}))$  are the smallest and largest  
 439 singular values of  $A(:, \Omega_{I_i})$ , respectively.

Therefore, choosing

$$s = O(\|\tilde{C}_{ii}\|_2 \varepsilon),$$

440 where  $\varepsilon$  is the machine precision, ensures  $\tilde{C}_{ii} + sI_{n_i}$  is numerically invertible and  
 441  $s \ll 1$ . Setting  $s = \|\tilde{C}_{ii}\|_2 \varepsilon$  in (3.8), we obtain

$$442 \quad \sigma_{\min}^2(A(:, \Omega_{I_i})) \leq \lambda \|\tilde{C}_{ii}\|_2 \varepsilon \leq \sigma_{\max}^2(A(:, \Omega_{I_i})).$$

By (3.5), we have

$$\|\tilde{C}_{ii}\|_2 \leq \|C_{ii}\|_2,$$

and because  $\Omega_{I_i} \subset \Omega_i$ , it follows that

$$\|C_{ii}^{-1}\|_2 = \|(A(:, \Omega_i)^\top A(:, \Omega_i))^{-1}\|_2 \leq \sigma_{\min}^2(A(:, \Omega_{I_i})).$$

443 Hence, if  $(v, \lambda)$  is an eigenpair of (3.7) with  $v \in \ker(\tilde{C}_{ii}) \cap \text{range}(D_i C_{ii} D_i)$ , then

444 
$$(\kappa(C_{ii})\varepsilon)^{-1} \leq \lambda,$$

445 where  $\kappa(C_{ii})$  is the condition number of  $C_{ii}$  and  $Z_i$  can be defined to be

446 (3.9) 
$$\text{span} \left\{ v \mid D_i C_{ii} D_i v = \lambda(\tilde{C}_{ii} + \varepsilon \|\tilde{C}_{ii}\|_2 I_{n_i}) v \text{ and } \lambda \geq \min \left( \frac{1}{\tau}, (\kappa(C_{ii})\varepsilon)^{-1} \right) \right\}.$$

447  $Z_i$  is then taken to be the matrix whose columns are the vertical concatenation of  
448 corresponding eigenvectors.

449 *Remark 3.3.* Note that solving the generalized eigenvalue problem (3.7) by an  
450 iterative method such as Krylov–Schur [43] does not require the explicit form of  $C_{ii}$   
451 and  $\tilde{C}_{ii}$ . Rather, it requires solving linear systems of the form  $(\tilde{C}_{ii} + sI_{n_i})u = v$ ,  
452 together with matrix–vector products of the form  $(\tilde{C}_{ii} + sI_{n_i})v$  and  $C_{ii}v$ . It is clear  
453 that these products do not require the matrices  $\tilde{C}_{ii}$  and  $C_{ii}$  to be formed. Regarding  
454 the solution of the linear system  $(\tilde{C}_{ii} + sI_{n_i})u = v$ , Remark 3.1 also applies to the  
455 Cholesky factorization of  $\tilde{C}_{ii} + sI_{n_i} = X^T X$ , where  $X^T = (A(\Xi_i, \Omega_i)^T \sqrt{s}I_{n_i})$ , that  
456 can be computed by using a “thin” QR factorization of  $X$ .

457 From Remarks 3.1 and 3.3, and applying the same technique therein to factor  
458  $C_{00} = R_0 C R_0^T = (AR_0^T)^T (AR_0^T)$ , we observe that given the overlapping partitions  
459 of  $A$ , the proposed two-level preconditioner can be constructed without forming the  
460 normal equations matrix. Algorithm 3.1 gives an overview of the steps for constructing  
461 our two-level Schwarz preconditioner for the normal equations matrix. The actual  
462 implementation of our proposed preconditioner will be discussed in greater detail  
in subsection 4.1.

---

**Algorithm 3.1** Two-level Schwarz preconditioner for the normal equations matrix.

---

**Input:** matrix  $A$ , number of subdomains  $N$ , threshold  $\tau$  to bound the condition number.

**Output:** two-level preconditioner  $M^{-1}$  for  $C = A^T A$ .

- 1:  $(\Omega_{I_1}, \dots, \Omega_{I_N}) = \text{Partition}(A, N)$
  - 2: **for**  $i = 1$  to  $N$  in parallel **do**
  - 3:    $\Xi_i = \text{FindNonzeroRows}(A(:, \Omega_{I_i}))$
  - 4:    $\Omega_i = [\Omega_{I_i}, \Omega_{\Gamma_i}] = \text{FindNonzeroColumns}(A(\Xi_i, :))$
  - 5:   Define  $D_i$  as in subsection 3.3.1 and  $R_i$  as in section 2
  - 6:   Perform Cholesky factorization of  $C_{ii} = A(:, \Omega_i)^T A(:, \Omega_i)$ , see Remark 3.1
  - 7:   Perform Cholesky factorization of  $\tilde{C}_{ii} = A(\Xi_i, \Omega_i)^T A(\Xi_i, \Omega_i)$ , possibly using a small shift  $s$ , see Remark 3.3
  - 8:   Compute  $Z_i$  as defined in (3.9)
  - 9: **end for**
  - 10: Set  $R_0^T = [R_1^T D_1 Z_1, \dots, R_N^T D_N Z_N]$
  - 11: Perform Cholesky factorization of  $C_{00} = (AR_0^T)^T (AR_0^T)$
  - 12: Set  $M^{-1} = M_{\text{additive}}^{-1} = \sum_{i=0}^N R_i^T C_{ii}^{-1} R_i$  or  $M_{\text{balanced}}^{-1}$  (2.10) or  $M_{\text{deflated}}^{-1}$  (2.11)
- 

463

464 **4. Numerical experiments.** In this section, we illustrate the effectiveness of  
465 the new two-level LS preconditioners  $M_{\text{balanced}}^{-1}$  and  $M_{\text{deflated}}^{-1}$ , their robustness with  
466 respect to the number of subdomains, and their efficiency in tackling large-scale sparse



TABLE 1  
*Test matrices taken from the SuiteSparse Matrix Collection.*

Identifier	$m$	$n$	$\text{nnz}(A)$	$\text{nnz}(C)$	$\text{condest}(C)$
mesh_deform	234,023	9,393	853,829	117,117	$2.7 \cdot 10^6$
EternityIIE	262,144	11,077	1,503,732	1,109,181	$5.1 \cdot 10^{19}$
lp_stocfor3	23,541	16,675	72,721	223,395	$4.0 \cdot 10^{10}$
deltaX	68,600	21,961	247,424	2,623,073	$3.7 \cdot 10^{20}$
sc205-2r	62,423	35,213	123,239	12,984,043	$1.7 \cdot 10^7$
stormg2-125	172,431	65,935	433,256	1,953,519	$\infty$
Rucci1	1,977,885	109,900	7,791,168	9,747,744	$2.0 \cdot 10^8$
image_interp	232,485	120,000	711,683	1,555,994	$4.7 \cdot 10^7$
mk13-b5	270,270	135,135	810,810	1,756,755	$\infty$
pds-100	514,577	156,016	1,096,002	1,470,688	$\infty$
fome21	267,596	216,350	465,294	640,240	$\infty$
sgpf5y6	312,540	246,077	831,976	2,761,021	$6.0 \cdot 10^6$
Hardesty2	929,901	303,645	4,020,731	3,936,209	$1.2 \cdot 10^{10}$
Delor338K	450,807	343,236	4,211,599	44,723,076	$1.5 \cdot 10^7$
watson_2	677,224	352,013	1,846,391	3,390,279	$1.0 \cdot 10^7$
LargeRegFile	2,111,154	801,374	4,944,201	6,378,592	$3.0 \cdot 10^8$
cont11.1	1,961,394	1,468,599	5,382,999	18,064,261	$2.0 \cdot 10^{10}$

467 and ill-conditioned LS problems selected from the SuiteSparse Matrix Collection [13].  
 468 The test matrices are listed in Table 1. For each matrix, we report its dimensions,  
 469 the number of entries in  $A$  and in the normal equations matrix  $C$ , and the condition  
 470 number of  $C$  (estimated using the MATLAB function `condest`).

471 In subsection 4.1, we discuss our implementation based on the parallel backend [7].  
 472 In particular, we show that very little coding effort is needed to construct all the  
 473 necessary algebraic tools, and that it is possible to take advantage of an existing  
 474 package, such as HPDDM [27], to setup the new preconditioners efficiently. We  
 475 then show in subsection 4.2 how  $M_{\text{balanced}}^{-1}$  and  $M_{\text{deflated}}^{-1}$  perform compared to other  
 476 preconditioners when solving challenging LS problems. The preconditioners we  
 477 consider are:

- 478 • limited memory incomplete Cholesky (IC) factorization specialized for  
 479 the normal equations matrix as implemented in HSL\_MI35 from the HSL  
 480 library [25] (note that this package is written in Fortran and we run it using  
 481 the supplied MATLAB interface with default parameter settings);
- 482 • one-level overlapping Schwarz methods  $M_{\text{ASM}}^{-1}$  and  $M_{\text{RAS}}^{-1}$  as implemented in  
 483 PETSc;
- 484 • algebraic multigrid methods as implemented both in BoomerAMG from the  
 485 HYPRE library [20] and in GAMG [1] from PETSc.

486 Finally, in subsection 4.3, we study the strong scalability of  $M_{\text{balanced}}^{-1}$  and its robustness  
 487 with respect to the number of subdomains by using a fixed problem and increasing  
 488 the number of subdomains.

489 With the exception of the serial IC code HSL\_MI35, all the numerical experiments  
 490 are performed on Irène, a system composed of 2,292 nodes with two 64-core AMD  
 491 Rome processors clocked at 2.6 GHz and, unless stated otherwise, 256 MPI processes  
 492 are used. For the domain decomposition methods, one subdomain is assigned per  
 493 process. All computations are performed in double-precision arithmetic.

494 In all our experiments, the vector  $b$  in (1.1) is generated randomly and the  
 495 initial guess for the iterative solver is zero. When constructing our new two-level  
 496 preconditioners, with the exception of the results presented in Figure 1, at most 300  
 497 eigenpairs are computed on each subdomain and the threshold parameter  $\tau$  from (3.9)  
 498 is set to 0.6. These parameters were found to provide good numerical performance  
 499 after a very quick trial-and-error approach on a single problem. We did not want to  
 500 adjust them for each problem from Table 1, but it will be shown next that they are  
 501 fine overall without additional tuning.

502 **4.1. Implementation aspects.** The new two-level preconditioners are  
 503 implemented on top of the well-known distributed memory library PETSc. This  
 504 section is not aimed at PETSc specialists. Rather, we want to briefly explain what  
 505 was needed to provide an efficient yet concise implementation. Our new code is open-  
 506 source, available at <https://github.com/prj-/aldaas2021robust>. It comprises fewer  
 507 than 150 lines of code (including the initialization and error analysis). The main  
 508 source files, written in Fortran, C, and Python, have three major phases, which we  
 509 now outline.

510 **4.1.1. Loading and partitioning phase.** First, PETSc is used to load the  
 511 matrix  $A$  in parallel, following a contiguous one-dimensional row partitioning among  
 512 MPI processes. We explicitly assemble the normal equations matrix using the routine  
 513 `MatTransposeMatMult` [32]. The initial PETSc-enforced parallel decomposition of  $A$   
 514 among processes may not be appropriate for the normal equations, so ParMETIS is  
 515 used by PETSc to repartition  $C$ . This also induces a permutation of the columns  
 516 of  $A$ .

517 **4.1.2. Setup phase.** To ensure that the normal equations matrix  $C$  is definite  
 518 and its Cholesky factorization is breakdown free,  $C$  is shifted by  $10^{-10}\|C\|_F I_n$  (here  
 519 and elsewhere,  $\|\cdot\|_F$  denotes the Frobenius norm). Note that this is only needed  
 520 for the construction of the preconditioner; the preconditioner is used to solve the  
 521 original LS problem. Given the indices of the columns owned by a MPI process,  
 522 we call the routine `MatIncreaseOverlap` on the normal equations matrix to build an  
 523 extended set of column indices of  $A$  that will be used to define overlapping subdomains.  
 524 These are the  $\Omega_i$  as defined in (3.1). Using the routine `MatFindNonzeroRows`, this  
 525 extended set of indices is used to concurrently find on each subdomain the set of  
 526 nonzero rows. These are the sets  $\Xi_i$  as illustrated in (3.2) and (3.3). The subdomain  
 527 matrices  $C_{ii}$  from (2.1) as well as the partition of unity  $D_i$  as illustrated in (3.4) are  
 528 automatically assembled by PETSc when using domain decomposition preconditioners  
 529 such as PCASM or PCHPDDM. The right-hand side matrices of the generalized  
 530 eigenvalue problems (3.6) are assembled using `MatTransposeMatMult`, but note that  
 531 this product is this time performed concurrently on each subdomain. The small  
 532 shift  $s$  from (3.7) is set to  $10^{-8}\|\tilde{C}_{ii}\|_F$ . These matrices and the sets of overlapping  
 533 column indices are passed to PCHPDDM using routine `PCHPDDMSetAuxiliaryMat`.  
 534 The rest of the setup is hidden from the user. It includes solving the generalized  
 535 eigenvalue problems using SLEPc [24], followed by the assembly and redistribution of  
 536 the second-level operator using a Galerkin product (2.2) (see [26] for more details on  
 537 how this is performed efficiently in PCHPDDM).

538 **4.1.3. Solution phase.** For the solution phase, users can choose between  
 539 multiple Krylov methods, including LSQR [35] and GMRES. We use left-  
 540 preconditioned LSQR (see, for example, [6, Algorithm 2]) and right-preconditioned  
 541 GMRES. Each iteration of LSQR requires matrix–vector products with  $A$  and  $A^\top$ . For

TABLE 2

Preconditioner comparison when running LSQR. Iteration counts are reported.  $M_{ASM}^{-1}$  and  $M_{balanced}^{-1}$  are the one- and two-level overlapping Schwarz preconditioners, respectively. † denotes iteration count exceeds 1,000. ‡ denotes either a failure in computing the preconditioner because of memory issues or a breakdown of LSQR.

Identifier	$M_{balanced}^{-1}$	$M_{ASM}^{-1}$	BoomerAMG	GAMG	HSL_MI35
mesh_deform	13	27	‡	35	<b>5</b>
EternityILE	<b>43</b>	91	‡	63	199
lp_stocfor3	<b>34</b>	136	‡	513	211
deltaX	<b>23</b>	98	‡	784	640
sc205-2r	<b>54</b>	61	‡	195	97
stormg2-125	<b>42</b>	174	‡	†	†
Rucci1	<b>21</b>	484	118	364	†
image_interp	<b>11</b>	409	40	203	†
mk13-b5	19	21	<b>11</b>	‡	<b>11</b>
pds-100	18	202	<b>16</b>	35	110
fome21	20	104	<b>16</b>	20	41
sgpf5y6	224	264	‡	163	<b>110</b>
Hardesty2	<b>30</b>	913	88	404	†
Delor338K	<b>10</b>	11	‡	†	829
watson_2	<b>15</b>	109	‡	64	<b>73</b>
LargeRegFile	41	109	19	‡	<b>12</b>
cont11.1	<b>30</b>	490	53	723	‡

542 GMRES, instead of using the previously explicitly assembled normal equations matrix,  
 543 we use an implicit representation of the operator that computes the matrix–vector  
 544 product with  $A$  followed by the product with  $A^\top$ . The type of overlapping Schwarz  
 545 method (additive or restricted additive) as well as the type of second-level correction  
 546 (balanced or deflated) may be selected at runtime by the user. This flexibility is  
 547 important because LSQR requires a symmetric preconditioner.

**4.2. Numerical validation.** In this section, we validate the effectiveness of the two-level method when compared to other preconditioners. Table 2 presents a comparison between five preconditioners: two-level additive Schwarz with balanced coarse correction  $M_{balanced}^{-1}$ , one-level additive Schwarz  $M_{ASM}^{-1}$ , BoomerAMG, GAMG, and HSL\_MI35. The first level of the one- and two-level methods both use the additive Schwarz formulation; the second level uses the balanced deflation formulation (2.10). The results are for the iterative solver LSQR. If  $M$  denotes the preconditioner, LSQR terminates when the LS residual satisfies

$$\frac{\|(AM^{-1})^\top (Ax - b)\|_2}{\|A\|_{M,F} \|Ax - b\|_2} < 10^{-8},$$

548 where  $\|A\|_{M,F} = \sum_{i=1}^n \lambda_i(M^{-1}A^\top A)$  is the sum of the positive eigenvalues of  
 549  $M^{-1}A^\top A$  that is approximated by LSQR itself. Note that if  $M^{-1} = W^{-1}W^{-\top}$ ,  
 550 then  $\|A\|_{M,F} = \|AW^{-1}\|_F$ .

551 It is clear that both the one- and two-level Schwarz methods are more robust  
 552 than the other preconditioners as they encounter no breakdowns and solve all the LS  
 553 problems using fewer than 1,000 iterations. Because HSL\_MI35 is a sequential code that  
 554 runs on a single core, there was not enough memory to compute the preconditioner

TABLE 3

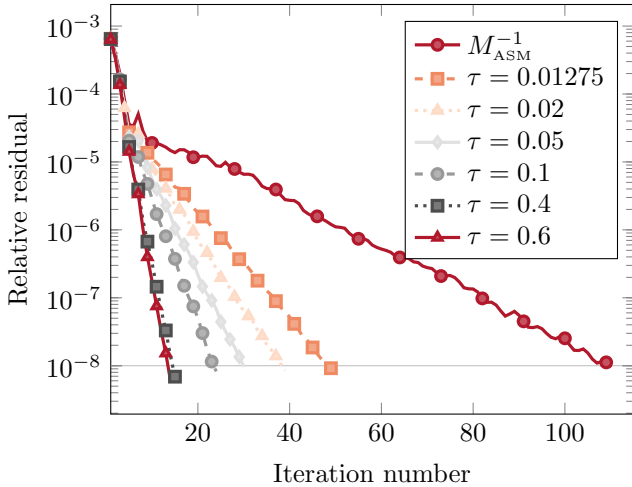
Preconditioner comparison when running GMRES. Iteration counts are reported.  $M_{RAS}^{-1}$  and  $M_{deflated}^{-1}$  are the one- and two-level overlapping Schwarz preconditioners, respectively. † denotes iteration count exceeds 1,000. ‡ denotes either a failure in computing the preconditioner because of memory issues or a breakdown of GMRES.

Identifier	$M_{deflated}^{-1}$	$M_{RAS}^{-1}$	BoomerAMG	GAMG	HSL_MI35
mesh_deform	6	27	21	50	<b>5</b>
EternityII.E	<b>5</b>	93	†	97	186
lp_stocfor3	<b>21</b>	†	†	†	198
deltaX	<b>6</b>	93	†	†	†
sc205-2r	<b>12</b>	125	†	490	69
stormg2-125	<b>23</b>	‡	‡	‡	†
Rucci1	<b>10</b>	958	213	882	†
image_interp	<b>10</b>	971	67	476	†
mk13-b5	14	18	21	‡	<b>12</b>
pds-100	<b>10</b>	84	23	51	115
fome21	<b>10</b>	55	22	29	41
sgpf5y6	116	†	†	249	<b>100</b>
Hardesty2	<b>26</b>	†	155	†	†
Delor338K	<b>5</b>	9	†	†	†
watson_2	<b>7</b>	134	252	96	<b>73</b>
LargeRegFile	<b>6</b>	21	23	‡	11
cont11.l	<b>45</b>	†	172	†	‡

555 for problem cont11.l. For many of the problems, the iteration count for HSL\_MI35  
556 can be reduced by increasing the parameters that determine the number of entries in  
557 the IC factor (the default values are rather small for the large test examples). LSQR  
558 preconditioned with BoomerAMG breaks down for several problems, as reported by  
559 PETSc error code KSP\_DIVERGED\_BREAKDOWN. GAMG is more robust but  
560 requires more iterations for problems where both algebraic multigrid solvers are  
561 successful. Note that even with more advanced options than the default ones set  
562 by PETSc, such as PMIS coarsening [14] with extended classical interpolation [15]  
563 for BoomerAMG or Schwarz smoothing for GAMG, these solvers do not perform  
564 considerably better numerically. We can also see that the two-level preconditioner  
565 outperforms the one-level preconditioner consistently.

566 Table 3 presents a similar comparison, but using right-preconditioned GMRES  
567 applied directly to the normal equations (1.2). A restart parameter of 100 is used. The  
568 relative tolerance is again set to  $10^{-8}$ , but this now applies to the unpreconditioned  
569 residual. We switch from  $M_{ASM}^{-1}$  to  $M_{RAS}^{-1}$  (2.9), which is known to perform better  
570 numerically. For the two-level method, we switch from  $M_{balanced}^{-1}$  to  $M_{deflated}^{-1}$  (2.11).  
571 Switching from LSQR to GMRES can be beneficial for some preconditioners, e.g.,  
572 BoomerAMG now converges in 21 iterations instead of breaking down for problem  
573 mesh\_deform. But this is not always the case, e.g., HSL\_MI35 applied to problem  
574 deltaX does not converge within the 1,000 iteration limit. The two-level method  
575 is the most robust approach, while the restricted additive Schwarz preconditioner  
576 struggles to solve some problems, either because of a breakdown (problem stormg2-  
577 125) or because of slow convergence (problems lp\_stocfor3, sgpf5y6, Hardesty2, and  
578 cont11.l).

579 Recall that for the results in Tables 2 and 3, the two-level preconditioner was



$\tau$	$n_0$	Iterations
0.01275	2,400	49
0.02	2,683	39
0.05	3,049	30
0.1	3,337	24
0.4	8,979	15
0.6	30,246	14

FIG. 1. Influence of the threshold parameter  $\tau$  on the convergence of preconditioned LSQR for problem *watson\_2* ( $m = 677,224$  and  $n = 352,013$ ).

580 constructed using at most 300 eigenpairs and the threshold parameter  $\tau$  was set  
 581 to 0.6. Whilst this highlights that tuning  $\tau$  for individual problems is not necessary  
 582 to successfully solve a range of problems, it does not validate the ability of our  
 583 preconditioner to concurrently select the most appropriate local eigenpairs to define  
 584 an adaptive preconditioner. To that end, for problem *watson\_2*, we consider the effect  
 585 on the performance of our two-level preconditioner of varying  $\tau$ . Results for LSQR  
 586 with  $M_{ASM}^{-1}$  and  $M_{balanced}^{-1}$  are presented in Figure 1. Here, 512 MPI processes are  
 587 used and the convergence tolerance is again  $10^{-8}$ . We observe that the two-level  
 588 method consistently outperforms the one-level method. Furthermore, as we increase  
 589  $\tau$ , the iteration count reduces and the size  $n_0$  of the second level increases. It is also  
 590 interesting to highlight that the convergence is smooth even with a very small value  
 591  $\tau = 0.01275$ ,  $n_0 = 2,400$  compared to the dimension  $3.52 \cdot 10^5$  of the normal equations  
 592 matrix.

593 **4.3. Performance study.** We next investigate the algorithmic cost of the two-  
 594 level method. To do so, we perform a strong scaling analysis using a large problem  
 595 not presented in Table 1 but still from the SuiteSparse Matrix Collection, *Hardesty3*.  
 596 The matrix is of dimension  $8,217,820 \times 7,591,564$ , and the number of nonzero entries  
 597 in  $C$  is 98,634,426. In Table 4, we report the number of iterations as well as the  
 598 eigensolve, setup, and solve times as the number  $N$  of subdomains ranges from 16 to  
 599 4,096. The times are obtained using the PETSc `-log_view` command line option. For  
 600 different  $N$ , the reported times on each row of the table are the maximum among  
 601 all processes. The setup time includes the numerical factorization of the first-level  
 602 subdomain matrices, the assembly of the second-level operator and its factorization.  
 603 Note that the symbolic factorization of the first-level subdomain is shared between  
 604 the domain decomposition preconditioner and the eigensolver because we use the  
 605 Krylov-Schur method as implemented in SLEPc, which requires the factorization of  
 606 the right-hand side matrices from (3.7). The Cholesky factorizations of the subdomain  
 607 matrices and of the second-level operator are performed using the sparse direct solver  
 608 MUMPS [5]. For small numbers of subdomains ( $N < 128$ ), the cost of the eigensolves  
 609 are clearly prohibitive. By increasing the number of subdomains, thus reducing their

TABLE 4

*Strong scaling for problem Hardesty3 ( $m = 8,217,820$  and  $n = 7,591,564$ ) for  $N$  ranging from 16 to 4,096 subdomains. All times are in seconds. Column 2 reports the LSQR iteration count. Column 4 reports the setup time minus the concurrent solution time of the generalized eigenproblems, which is given in column 3.*

$N$	Iterations	Eigensolve	Setup	Solve	$n_0$	Total	Speedup
16	113	2,417.4	24.5	301.3	4,800	2,743.2	–
32	117	1,032.7	14.1	154.2	9,600	1,201.0	2.3
64	129	887.2	11.4	112.3	19,200	1,010.9	2.7
128	144	224.1	6.9	55.4	38,400	286.3	9.6
256	97	128.0	6.7	32.2	76,800	166.9	16.4
512	87	45.5	13.0	26.9	153,391	85.3	32.2
1,024	85	23.8	20.2	35.3	303,929	79.3	34.6
2,048	55	14.6	31.4	43.2	497,704	89.1	30.8
4,096	59	11.7	30.8	44.9	695,774	87.3	31.4

size, the time to construct the preconditioner becomes much more tractable and overall, our implementation yields good speedups on a wide range of process counts. Note that the threshold parameter  $\tau = 0.6$  is not attained on any of the subdomains for  $N$  ranging from 16 up to 256, so that  $n_0 = 300 \times N$ . For larger  $N$ ,  $\tau = 0.6$  is attained, the preconditioner automatically selects the appropriate eigenmodes, and convergence improves (see column 2 of Table 4). When  $N$  is large ( $N \geq 1,024$ ), the setup and solve times are impacted by the high cost of factorizing and solving the second-level problems, which, as highlighted by the values of  $n_0$ , become large. Multilevel variants [4] could be used to overcome this but goes beyond the scope of the current study.

**5. Concluding comments.** Solving large-scale sparse linear least-squares problems is known to be challenging. Previously proposed preconditioners have generally been serial and have involved incomplete factorizations of  $A$  or  $C = A^\top A$ . In this paper, we have employed ideas that have been developed in the area of domain decomposition, which (as far as we are aware) have not previously been applied to least-squares problems. In particular, we have exploited recent work by Al Daas and Grigori [3] on algebraic domain decomposition preconditioners for SPD systems to propose a new two-level algebraic domain preconditioner for the normal equations matrix  $C$ . We have used the concept of an algebraic local SPSD splitting of an SPD matrix and we have shown that the structure of  $C$  as the product of  $A^\top$  and  $A$  can be used to efficiently perform the splitting. Furthermore, we have proved that using the two-level preconditioner, the spectral condition number of the preconditioned normal equations matrix is bounded from above independently of the number of the subdomains and the size of the problem. Moreover, this upper bound depends on a parameter  $\tau$  that can be chosen by the user to decrease (resp. increase) the upper bound with the costs of setting up the preconditioner being larger (resp. smaller).

The new two-level preconditioner has been implemented in parallel within PETSc. Numerical experiments on a range of problems from real applications have shown that whilst both one-level and two-level domain decomposition preconditioners are effective when used with LSQR to solve the normal equations, the latter consistently results in significantly faster convergence. It also outperforms other possible preconditioners, both in terms of robustness and iteration counts. Furthermore, our numerical experiments on a set of challenging least-squares problems show that the two-level

643 preconditioner is robust with respect to the parameter  $\tau$ . Moreover, a strong  
 644 scalability test of the two-level preconditioner assessed its robustness with respect  
 645 to the number of subdomains.

646 Future work includes extending the approach to develop preconditioners for  
 647 solving large sparse–dense least-squares problems in which  $A$  contains a small number  
 648 of rows that have many more entries than the other rows. These cause the normal  
 649 equations matrix to be dense and so they need to be handled separately (see, for  
 650 example, the recent work of Scott and Tůma [40] and references therein). As already  
 651 observed, we also plan to consider multilevel variants to allow the use of a larger  
 652 number of subdomains and processes.

653 **Acknowledgments.** This work was granted access to the GENCI-sponsored  
 654 HPC resources of TGCC@CEA under allocation A0090607519. The authors would  
 655 like to thank L. Dalcin, V. Hapla, and T. Isaac for their recent contributions to PETSc  
 656 that made the implementation of our preconditioner more flexible. Finally, we are  
 657 grateful to two anonymous reviewers for their constructive feedback.

658 **Code reproducibility.** Interested readers are referred to <https://github.com/prj-/aldaas2021robust/blob/main/README.md>  
 659 for setting up the appropriate requirements, compiling, and running our proposed preconditioner. Fortran, C, and  
 660 Python source codes are provided.  
 661

662

## REFERENCES

- 663 [1] M. F. ADAMS, H. H. BAYRAKTAR, T. M. KEAVENY, AND P. PAPADOPOULOS, *Ultrascaleable*  
 664 *implicit finite element analyses in solid mechanics with over a half a billion degrees of*  
 665 *freedom*, in Proceedings of the 2004 ACM/IEEE Conference on Supercomputing, SC04,  
 666 IEEE Computer Society, 2004, pp. 34:1–34:15.
- 667 [2] E. AGULLO, A. BUTTARI, A. GUERMOUCHE, AND F. LOPEZ, *Implementing multifrontal*  
 668 *sparse solvers for multicore architectures with sequential task flow runtime systems*,  
 669 ACM Transactions on Mathematical Software, 43 (2016), [http://buttari.perso.enseeiht.](http://buttari.perso.enseeiht.fr/qr_mumps)  
 670 [fr/qr\\_mumps](http://buttari.perso.enseeiht.fr/qr_mumps).
- 671 [3] H. AL DAAS AND L. GRIGORI, *A class of efficient locally constructed preconditioners based on*  
 672 *coarse spaces*, SIAM Journal on Matrix Analysis and Applications, 40 (2019), pp. 66–91.
- 673 [4] H. AL DAAS, L. GRIGORI, P. JOLIVET, AND P.-H. TOURNIER, *A multilevel Schwarz*  
 674 *preconditioner based on a hierarchy of robust coarse spaces*, SIAM Journal on Scientific  
 675 Computing, 43 (2021), pp. A1907–A1928.
- 676 [5] P. R. AMESTOY, I. S. DUFF, J.-Y. L’EXCELLENT, AND J. KOSTER, *A fully asynchronous*  
 677 *multifrontal solver using distributed dynamic scheduling*, SIAM Journal on Matrix Analysis  
 678 and Applications, 23 (2001), pp. 15–41, <http://mumps.enseeiht.fr>.
- 679 [6] S. R. ARRIDGE, M. M. BETCKE, AND L. HARHANEN, *Iterated preconditioned LSQR method for*  
 680 *inverse problems on unstructured grids*, Inverse Problems, 30 (2014), p. 075009.
- 681 [7] S. BALAY, S. ABHYANKAR, M. F. ADAMS, J. BROWN, P. BRUNE, K. BUSCHELMAN, L. DALCIN,  
 682 A. DENER, V. ELJKHOUT, W. D. GROPP, D. KARPEYEV, D. KAUSHIK, M. G. KNEPLEY,  
 683 D. A. MAY, L. C. MCINNES, R. T. MILLS, T. MUNSON, K. RUPP, P. SANAN, B. F. SMITH,  
 684 S. ZAMPINI, H. ZHANG, AND H. ZHANG, *PETSc web page*, 2021, <https://petsc.org>.
- 685 [8] R. BRU, J. MARÍN, J. MAS, AND M. TŮMA, *Preconditioned iterative methods for solving linear*  
 686 *least-squares problems*, SIAM Journal on Scientific Computing, 36 (2014), pp. A2002–  
 687 A2022.
- 688 [9] X.-C. CAI AND M. SARKIS, *A restricted additive Schwarz preconditioner for general sparse*  
 689 *linear systems*, SIAM Journal on Scientific Computing, 21 (1999), pp. 792–797.
- 690 [10] T. F. CHAN AND T. P. MATHEW, *Domain decomposition algorithms*, Acta Numerica, 3 (1994),  
 691 pp. 61–143.
- 692 [11] X. CUI AND K. HAYAMI, *Generalized approximate inverse preconditioners for least-squares*  
 693 *problems*, Japan Journal of Industrial and Applied Mathematics, 26 (2009).
- 694 [12] T. A. DAVIS, *Algorithm 915, SuiteSparseQR: multifrontal multithreaded rank-revealing sparse*  
 695 *QR factorization*, ACM Transactions on Mathematical Software, 38 (2011).

- 696 [13] T. A. DAVIS AND Y. HU, *The University of Florida sparse matrix collection*, ACM Transactions  
697 on Mathematical Software, 38 (2011), pp. 1–28.
- 698 [14] H. DE STERCK, R. D. FALGOUT, J. W. NOLTING, AND U. M. YANG, *Distance-two interpolation*  
699 *for parallel algebraic multigrid*, Numerical Linear Algebra with Applications, 15 (2008),  
700 pp. 115–139.
- 701 [15] H. DE STERCK, U. M. YANG, AND J. J. HEYS, *Reducing complexity in parallel algebraic*  
702 *multigrid preconditioners*, SIAM Journal on Matrix Analysis and Applications, 27 (2006),  
703 pp. 1019–1039.
- 704 [16] V. DOLEAN, P. JOLIVET, AND F. NATAF, *An introduction to domain decomposition methods.*  
705 *Algorithms, theory, and parallel implementation*, Society for Industrial and Applied  
706 Mathematics, 2015.
- 707 [17] I. S. DUFF, R. GUIVARCH, D. RUIZ, AND M. ZENADI, *The augmented block Cimmino distributed*  
708 *method*, SIAM Journal on Scientific Computing, 37 (2015), pp. A1248–A1269.
- 709 [18] A. DUMITRAȘC, P. LELEUX, C. POPA, D. RUIZ, AND S. TORUN, *The augmented block Cimmino*  
710 *algorithm revisited*, 2018, <https://arxiv.org/abs/1805.11487>.
- 711 [19] T. ELFVING, *Block-iterative methods for consistent and inconsistent linear equations*,  
712 Numerische Mathematik, 35 (1980), pp. 1–12.
- 713 [20] R. D. FALGOUT AND U. M. YANG, *hypr: a library of high performance preconditioners*,  
714 Computational Science—ICCS 2002, (2002), pp. 632–641.
- 715 [21] M. J. GANDER AND A. LONELAND, *SHEM: an optimal coarse space for RAS and its multiscale*  
716 *approximation*, in Domain Decomposition Methods in Science and Engineering XXIII, C.-  
717 O. Lee, X.-C. Cai, D. E. Keyes, H. H. Kim, A. Klawonn, E.-J. Park, and O. B. Widlund,  
718 eds., Cham, 2017, Springer International Publishing, pp. 313–321.
- 719 [22] N. I. M. GOULD AND J. A. SCOTT, *The state-of-the-art of preconditioners for sparse linear*  
720 *least-squares problems*, ACM Transactions on Mathematical Software, 43 (2017), pp. 36:1–  
721 35.
- 722 [23] A. HEINLEIN, C. HOCHMUTH, AND A. KLAWONN, *Reduced dimension GDSW coarse spaces for*  
723 *monolithic Schwarz domain decomposition methods for incompressible fluid flow problems*,  
724 International Journal for Numerical Methods in Engineering, 121 (2020), pp. 1101–1119.
- 725 [24] V. HERNANDEZ, J. E. ROMAN, AND V. VIDAL, *SLEPc: a scalable and flexible toolkit for the*  
726 *solution of eigenvalue problems*, ACM Transactions on Mathematical Software, 31 (2005),  
727 pp. 351–362, <https://slepc.upv.es>.
- 728 [25] *HSL. A collection of Fortran codes for large-scale scientific computation*, 2018. [http://www.](http://www.hsl.rl.ac.uk)  
729 [hsl.rl.ac.uk](http://www.hsl.rl.ac.uk).
- 730 [26] P. JOLIVET, F. HECHT, F. NATAF, AND C. PRUD'HOMME, *Scalable domain decomposition*  
731 *preconditioners for heterogeneous elliptic problems*, in Proceedings of the International  
732 Conference on High Performance Computing, Networking, Storage and Analysis, SC '13,  
733 New York, NY, USA, 2013, ACM, pp. 80:1–80:11.
- 734 [27] P. JOLIVET, J. E. ROMAN, AND S. ZAMPINI, *KSPHPDDM and PCHPDDM: extending PETSc*  
735 *with advanced Krylov methods and robust multilevel overlapping Schwarz preconditioners*,  
736 Computers & Mathematics with Applications, 84 (2021), pp. 277–295.
- 737 [28] G. KARYPIS AND V. KUMAR, *Multilevel k-way partitioning scheme for irregular graphs*, Journal  
738 of Parallel and Distributed computing, 48 (1998), pp. 96–129.
- 739 [29] F. KONG AND X.-C. CAI, *A scalable nonlinear fluid–structure interaction solver based on a*  
740 *Schwarz preconditioner with isogeometric unstructured coarse spaces in 3D*, Journal of  
741 Computational Physics, 340 (2017), pp. 498–518.
- 742 [30] N. LI AND Y. SAAD, *MIQR: a multilevel incomplete QR preconditioner for large sparse least-*  
743 *squares problems*, SIAM Journal on Matrix Analysis and Applications, 28 (2006), pp. 524–  
744 550.
- 745 [31] P. MARCHAND, X. CLAEYS, P. JOLIVET, F. NATAF, AND P.-H. TOURNIER, *Two-level*  
746 *preconditioning for h-version boundary element approximation of hypersingular operator*  
747 *with GenEO*, Numerische Mathematik, 146 (2020), pp. 597–628.
- 748 [32] M. MCCOURT, B. F. SMITH, AND H. ZHANG, *Sparse matrix–matrix products executed through*  
749 *coloring*, SIAM Journal on Matrix Analysis and Applications, 36 (2015), pp. 90–109.
- 750 [33] *Intel MKL Sparse QR*, 2018.
- 751 [34] S. V. NEPOMNYASCHIKH, *Mesh theorems of traces, normalizations of function traces and their*  
752 *inversions*, Russian Journal of Numerical Analysis and Mathematical Modelling, 6 (1991),  
753 pp. 1–25.
- 754 [35] C. C. PAIGE AND M. A. SAUNDERS, *LSQR: an algorithm for sparse linear equations and sparse*  
755 *least squares*, ACM Transactions on Mathematical Software, 8 (1982), p. 43–71.
- 756 [36] F. PELLEGRINI AND J. ROMAN, *SCOTCH: a software package for static mapping by*  
757 *dual recursive bipartitioning of process and architecture graphs*, in High-Performance



- 758 Computing and Networking, Springer, 1996, pp. 493–498.
- 759 [37] Y. SAAD AND M. H. SCHULTZ, *GMRES: a generalized minimal residual algorithm for solving*  
760 *nonsymmetric linear systems*, SIAM Journal on Scientific and Statistical Computing, 7  
761 (1986), pp. 856–869.
- 762 [38] J. A. SCOTT AND M. TŪMA, *Preconditioning of linear least squares by robust incomplete*  
763 *factorization for implicitly held normal equations*, SIAM Journal on Scientific Computing,  
764 38 (2016), pp. C603–C623.
- 765 [39] J. A. SCOTT AND M. TŪMA, *Solving mixed sparse–dense linear least-squares problems by*  
766 *preconditioned iterative methods*, SIAM Journal on Scientific Computing, 39 (2017),  
767 pp. A2422–A2437.
- 768 [40] J. A. SCOTT AND M. TŪMA, *Strengths and limitations of stretching for least-squares problems*  
769 *with some dense rows*, ACM Transactions on Mathematical Software, 41 (2021), pp. 1:1–  
770 1:25.
- 771 [41] B. F. SMITH, P. E. BJØRSTAD, AND W. D. GROPP, *Domain decomposition: parallel multilevel*  
772 *methods for elliptic partial differential equations*, Cambridge University Press, 1996.
- 773 [42] N. SPILLANE, V. DOLEAN, P. HAURET, F. NATAF, C. PECHSTEIN, AND R. SCHEICHL, *Abstract*  
774 *robust coarse spaces for systems of PDEs via generalized eigenproblems in the overlaps*,  
775 Numerische Mathematik, 126 (2014), pp. 741–770.
- 776 [43] G. W. STEWART, *A Krylov–Schur algorithm for large eigenproblems*, SIAM Journal on Matrix  
777 Analysis and Applications, 23 (2002), pp. 601–614.
- 778 [44] J. M. TANG, R. NABBEN, C. VUIK, AND Y. A. ERLANGGA, *Comparison of two-level*  
779 *preconditioners derived from deflation, domain decomposition and multigrid methods*,  
780 Journal of Scientific Computing, 39 (2009), pp. 340–370.
- 781 [45] J. VAN LENT, R. SCHEICHL, AND I. G. GRAHAM, *Energy-minimizing coarse spaces for two-level*  
782 *Schwarz methods for multiscale PDEs*, Numerical Linear Algebra with Applications, 16  
783 (2009), pp. 775–799.