LOCAL FOURIER ANALYSIS OF BDDC-LIKE ALGORITHMS*

JED BROWN†, YUNHUI HE‡, AND SCOTT MACLACHLAN‡

Abstract. Local Fourier analysis is a commonly used tool for the analysis of multigrid and other multilevel algorithms, providing both insight into observed convergence rates and predictive analysis of the performance of many algorithms. In this paper, we adapt local Fourier analysis to examine variants of two- and three-level BDDC algorithms, to better understand the eigenvalue distributions and condition number bounds on these preconditioned operators. This adaptation is based on a new choice of basis for the space of Fourier harmonics that greatly simplifies the application of local Fourier analysis in this setting. The local Fourier analysis is validated by considering the two dimensional Laplacian and predicting the condition numbers of the preconditioned operators with different sizes of subdomains. Several variants are analyzed, showing the two- and three-level performance of the "lumped" variant can be greatly improved when used in multiplicative combination with a weighted diagonal scaling preconditioner, with weight optimized through the use of LFA.

Key words. BDDC, Domain decomposition, Local Fourier analysis, Multiplicative methods

AMS subject classifications. 65N22, 65N55, 65F08

1. Introduction.

Domain decomposition methods are well-studied approaches for the numerical solution of partial differential equations both experimentally and theoretically [1, 10, 12, 27], due to their efficiency and robustness for many large-scale problems, and the need for parallel algorithms. Among the main families of domain decomposition algorithms are Neumann-Neumann [27], FETI [13], Schwarz [12, 27], and Optimized Schwarz [10, 15]. Balancing domain decomposition by constraints (BDDC) is one family of non-overlapping domain decomposition method. While BDDC was first introduced by Dohrmann in [6], several variants have recently been proposed. BDDC-like methods have been successfully applied to many PDEs, including elliptic problems [18, 22], the incompressible Stokes equations [17, 19], H(curl) problems [9], flow in porous media [29], and the incompressible elasticity problem [7, 8]. Theoretical analysis of BDDC has primarily been based on finite-element approximation theory [4, 7, 11, 23, 24]. It has been shown that the condition number of the preconditioned BDDC operator can be bounded by a function of $\frac{H}{h}$(where h is the meshsize, and H is the subdomain size), independent of the number of subdomains [29]. A nonoverlapping domain decomposition method for discontinuous Galerkin based on the BDDC algorithm is presented in [3], and the condition number of the preconditioned system is shown to be bounded by similar estimates as those for conforming finite element methods. BDDC methods in three- or multilevel forms have also been developed [25, 30, 31].

Since BDDC algorithms are widely used to solve many problems with high efficiency and parallelism, better understanding of how this methodology works is useful in the design of new algorithms. Local Fourier analysis (LFA), first introduced by Brandt [2] and well-studied for multigrid methods [5, 26, 28, 32, 33], is an analysis framework that provides predictive performance estimates for many multilevel itera-

*Submitted to the editors DATE.

Funding: The work of J.B. was partially funded by U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research under Award Number DE-SC0016140. The work of S.M. was partially funded by an NSERC Discovery Grant.

†Department of Computer Science, University of Colorado Boulder, 430 UCB, Boulder, CO 80309 (jed@jedbrown.org).

‡Department of Mathematics and Statistics, Memorial University of Newfoundland, St. John's, NL A1C 5S7, Canada (yunhui.he@mun.ca, smaclachlan@mun.ca).

tions and preconditioners. However, to our knowledge, there has been no research applying local Fourier analysis to BDDC-like algorithms. The same is true of the closely related finite element tearing and interconnect (FETI) methodology [13, 14, 16]. Because LFA can reflect both the distribution of eigenvalues and associated eigenvectors of a preconditioned operator, here, we adopt LFA to analyze variants of the common "lumped" and "Dirichlet" BDDC algorithms, based on [20], to guide construction of these methods. To do this, we introduce a novel basis for the Fourier analysis that is well-suited for application to domain decomposition preconditioners.

Applying the two-level BDDC algorithm requires the solution of a Schur complement equation (coarse problem), which usually poses some difficulty with increasing problem size. Two- and three-level variants are, thus, considered in this paper. However, as is well-known in the literature, the performance of BDDC degrades sharply from two-level to three-level methods, particularly for large values of H/h. Since our analysis shows that the largest eigenvalues of the preconditioned operator for the lumped BDDC algorithm are associated with oscillatory modes, we propose variants of BDDC based on multiplicative preconditioning and multigrid ideas. From the condition numbers offered by LFA, we can easily compare the efficiency of these variants. Furthermore, LFA can provide optimal parameters for these multiplicative methods, helping tune and understand sensitivity to the parameter choice.

This paper is organized as follows. In Section 2, we introduce the finite element discretization of the Laplace problem in two dimensions and the lumped and Dirichlet preconditioners. Two- and three-level preconditioned operators are developed in Section 3. In Section 4, we discuss the Fourier representation of the preconditioned operators. Section 5 reports LFA-predicted condition numbers of the BDDC variants considered here. Conclusions are presented in Section 6.

2. Discretization. We consider the two-dimensional Laplace problem in weak form: Find $u \in H_0^1(\Omega) := V$ such that

$$a(u,v) = \int_\Omega \nabla u \cdot \nabla v d\Omega = \langle f, v\rangle, \forall v \in V, \tag{2.1}$$

where $\Omega \subset \mathbb{R}^2$ is a bounded domain with Lipschitz boundary $\partial\Omega$. Here, we consider the Ritz-Galerkin approximation over V_h, the space of piecewise bilinear functions on a uniform rectangular mesh of $\Omega = [0,1]^2$. The corresponding linear system of equations is given as

$$Ax = b. \tag{2.2}$$

We partition the domain, Ω, into N nonoverlapping subdomains, $\Omega_i, i = 1, 2, \cdots, N$, where each subdomain is a union of shape regular elements and the nodes on the boundaries of neighboring subdomains match across the interface $\Gamma = \bigcup \partial\Omega_i \backslash \partial\Omega$. The interface of subdomain Ω_i is defined by $\Gamma_i = \partial\Omega_i \bigcap \Gamma$. Here, we consider $\Omega = [0,1]^2$, with both a discretization mesh (with meshsize h) and subdomain mesh (with meshsize $H = ph$) given by uniform grids with square elements or subdomains.

The finite-element space V_h can be rewritten as $V_h = V_{I,h} \bigoplus V_{\Gamma,h}$, where $V_{I,h}$ is the product of the subdomain interior variable spaces $V_{I,h}^{(i)}$. Functions in $V_{I,h}^{(i)}$ are supported in the subdomain Ω_i and vanish on the subdomain interface Γ_i. $V_{\Gamma,h}$ is the space of traces on Γ of functions in V_h. Then, we can write the subdomain problem with Neumann boundary conditions on Γ_i as

$$A^{(i)}x^{(i)} = \begin{pmatrix} A_{II}^{(i)} & A_{\Gamma I}^{(i)^T} \\ A_{\Gamma I}^{(i)} & A_{\Gamma\Gamma}^{(i)} \end{pmatrix} \begin{pmatrix} x_I^{(i)} \\ x_\Gamma^{(i)} \end{pmatrix} = \begin{pmatrix} b_I^{(i)} \\ b_\Gamma^{(i)} \end{pmatrix}, \tag{2.3}$$

where $x^{(i)} = (x_I^{(i)}, x_\Gamma^{(i)}) \in V_h^{(i)} = (V_{I,h}^{(i)}, V_{\Gamma,h}^{(i)})$, and T is the (conjugate) transpose. Then, the global problem (2.2) can be assembled from the subdomain problems (2.3) as

$$A = \sum_{i=1}^{N} {R^{(i)}}^T A^{(i)} R^{(i)}, \text{ and } b = \sum_{i=1}^{N} {R^{(i)}}^T b^{(i)},$$

where $R^{(i)}$ is the restriction operator from a global vector to a subdomain vector on Ω_i.

2.1. A Partially Subassembled Problem. In order to describe variants of the BDDC methods, we first introduce a partially subassembled problem, following [20], and the corresponding space of partially subassembled variables,

$$\hat{V}_h = V_{\Pi,h} \bigoplus V_{r,h}, \tag{2.4}$$

where $V_{\Pi,h}$ is spanned by the subdomain vertex nodal basis functions (the coarse degrees of freedom). The complementary space, $V_{r,h}$, is the product of the subdomain spaces $V_{r,h}^{(i)}$, which correspond to the subdomain interior and interface degrees of freedom and are spanned by the basis functions which vanish at the coarse-grid degrees of freedom. For a 4×4 mesh, the degrees of freedom in $V_{\Pi,h}$ are those corresponding to the circled nodes at the left of Figure 1, while the degrees of freedom in $V_{r,h}$ correspond to all interior nodes, plus duplicated (broken) degrees freedom along subdomain boundaries.

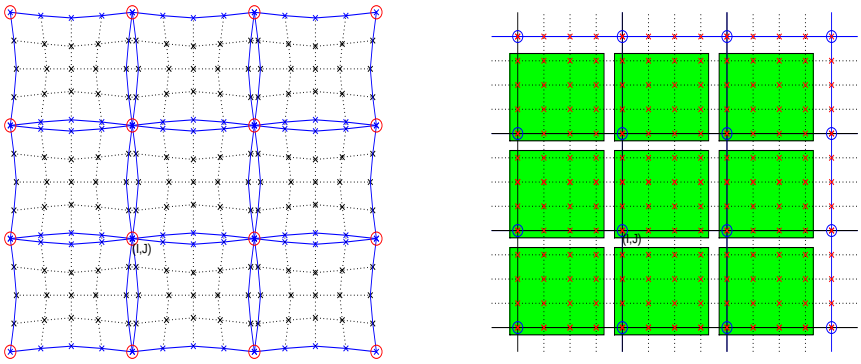

FIG. 1. *At left, the partially broken decomposition given in Equation (2.4), with circled degrees of freedom corresponding to $V_{\Pi,h}$ and all others corresponding to $V_{r,h}$. This matches the periodic array of subdomains induced by the subsets $\mathfrak{S}^*_{I,J}$ introduced in Equation (4.11) for $p = 4$. At right, a non-overlapping decomposition into subdomains of size $p \times p$ for $p = 4$, corresponding to the subsets $\mathfrak{S}_{I,J}$ introduced in Equation (4.7), where LFA works on an infinite grid and characterizes operators by their action in terms of the non-overlapping partition denoted in green.*

The partially subassembled problem matrix, corresponding to the variables in the space $\hat{V}_h$, is obtained by assembling the subdomain matrices (2.3) only with respect to the coarse-level variables; that is,

$$\hat{A} = \sum_{i=1}^{N} (\bar{R}^{(i)})^T A^{(i)} \bar{R}^{(i)}, \tag{2.5}$$

where $\bar{R}^{(i)}$ is a restriction from space $\hat{V}_h$ to $V_h^{(i)}$.

2.2. Lumped and Dirichlet Preconditioners. In order to define the preconditioners under consideration for (2.2), we introduce a positive scaling factor, $\delta_i(\boldsymbol{x})$, for each node $\boldsymbol{x}$ on the interface Γ_i of subdomain Ω_i. Let $\mathcal{N}_{\boldsymbol{x}}$ be the set of indices of the subdomains that have $\boldsymbol{x}$ on their boundaries. Define $\delta_i(\boldsymbol{x}) = 1/|\mathcal{N}_{\boldsymbol{x}}|$, where $|\mathcal{N}_{\boldsymbol{x}}|$ is the cardinality of $\mathcal{N}_{\boldsymbol{x}}$. The scaled injection operator, $\mathcal{R}_1$, is defined so that each column of $\mathcal{R}_1$ corresponds to a degree of freedom of the global problem (2.2). For subdomain interior and coarse-level variables, the corresponding column of $\mathcal{R}_1$ has a single entry with value 1. Columns that correspond to an interface degree of freedom $\boldsymbol{x} \in \Gamma_{i,h}$ (the set of nodes in Γ_i) have $|\mathcal{N}_{\boldsymbol{x}}|$ non-zero entries each of $\delta_i(\boldsymbol{x})$.

Based on the partially subassembled problem, the first preconditioner introduced for solving (2.2) is

$$M_1^{-1} = \mathcal{R}_1^T \hat{A}^{-1} \mathcal{R}_1.$$

The preconditioned operator $M_1^{-1}A$ has the same eigenvalues as the preconditioned FETI-DP operator with a lumped preconditioner, except for some eigenvalues equal to 0 and 1 [14, 20]. We refer to M_1 as the *lumped* preconditioner.

A similar preconditioner for A augments this using discrete harmonic extensions in the restriction and interpolation operators [20], giving

$$M_2^{-1} = (\mathcal{R}_1^T - \mathcal{H} J_D) \hat{A}^{-1} \underbrace{(\mathcal{R}_1 - J_D^T \mathcal{H}^T)}_{:=\mathcal{R}_2}, \tag{2.6}$$

where $\mathcal{H}$ is the direct sum of $\mathcal{H}^{(i)} = -(A_{II}^{(i)})^{-1}(A_{\Gamma I}^{(i)})^T$, which maps the jump over a subdomain interface (given by J_D) to the interior of the subdomain by solving a local Dirichlet problem, and gives zero for other values. For any given $v \in \hat{V}_h$, the component of $J_D^T v$ on subdomain Ω_i is given by

$$\left(J_D^T v(\boldsymbol{x})\right)^{(i)} = \sum_{j \in \mathcal{N}_{\boldsymbol{x}}} \left(\delta_j(\boldsymbol{x}) v^{(i)}(\boldsymbol{x}) - \delta_i(\boldsymbol{x}) v^{(j)}(\boldsymbol{x})\right), \ \forall \boldsymbol{x} \in \Gamma_{i,h}. \tag{2.7}$$

Extending the interface values using the discrete harmonic extension minimizes the energy norm of the resulting vector [27], giving a better stability bound. Furthermore, the preconditioned operator $M_2^{-1}A$ has the same eigenvalues as the BDDC operator [18], except for some eigenvalues equal to 1 [20]. We refer to M_2 as the *Dirichlet* preconditioner.

Standard bounds (see, e.g., [20]) on the condition numbers of the preconditioned operators are that, for $M_1^{-1}A$, there exists $\mathfrak{C}_{1,0} \geq 0$ such that $\kappa \leq \mathfrak{C}_{1,0} \frac{H}{h}(1 + \log \frac{H}{h})$ and, for $M_2^{-1}A$, there exists $\mathfrak{C}_{2,0} \geq 0$ such that $\kappa \leq \mathfrak{C}_{2,0}(1 + \log \frac{H}{h})^2$.

3. Two- and Three-level Variants. In both of the above preconditioned operators, we need to solve the following partially subassembled problem, now written in block form

$$\hat{A}\hat{x} = \begin{pmatrix} A_{rr} & \hat{A}_{\Pi r}^T \\ \hat{A}_{\Pi r} & A_{\Pi\Pi} \end{pmatrix} \begin{pmatrix} \hat{x}_r \\ \hat{x}_\Pi \end{pmatrix} = \begin{pmatrix} A_{rr} & 0 \\ \hat{A}_{\Pi r} & \hat{S}_\Pi \end{pmatrix} \begin{pmatrix} I & A_{rr}^{-1}\hat{A}_{\Pi r}^T \\ 0 & I \end{pmatrix} \begin{pmatrix} \hat{x}_r \\ \hat{x}_\Pi \end{pmatrix} = \begin{pmatrix} \hat{d}_r \\ \hat{d}_\Pi \end{pmatrix} = \hat{d}, \tag{3.1}$$

where $\hat{S}_\Pi = A_{\Pi\Pi} - \hat{A}_{\Pi r} A_{rr}^{-1} \hat{A}_{\Pi r}^T$, $\hat{x}_r$ contains the subdomain interior and interface degrees of freedom, and $\hat{x}_\Pi$ corresponds to the coarse-level degrees of freedom, which

are located at the corners of the subdomains. We write $\hat{A}$ in (3.1) in factorization form to easily separate the action on the coarse degrees of freedom, and to find the corresponding symbol of $\hat{A}^{-1}$. If we define

$$P = \begin{pmatrix} -A_{rr}^{-1}\hat{A}_{\Pi r}^T \\ I \end{pmatrix},$$

then the Schur complement is the Galerkin coarse operator, $\hat{S}_\Pi = P^T\hat{A}P$, and block-factorization solve for $\hat{A}$ can be seen to be equivalent to a two-level additive multigrid method with exact F-relaxation using

$$S_F = \begin{pmatrix} A_{rr}^{-1} & 0 \\ 0 & 0 \end{pmatrix}.$$

In the partially subassembled problem (3.1), we need to solve a coarse problem related to $\hat{S}_\Pi$. We can either solve this coarse problem exactly (corresponding to a two-level method, where the Schur complement is inverted exactly) or inexactly (as a three-level method), where the lumped and Dirichlet preconditioners defined above are used recursively to solve this problem.

3.1. Exact and Inexact Solve for the Schur Complement. Let

$$\mathcal{K}_1 = \begin{pmatrix} A_{rr} & 0 \\ \hat{A}_{\Pi r} & \hat{S}_\Pi \end{pmatrix}, \quad \mathcal{K}_2 = \begin{pmatrix} I & A_{rr}^{-1}\hat{A}_{\Pi r}^T \\ 0 & I \end{pmatrix},$$

and note that the product of $\mathcal{K}_1$ and $\mathcal{K}_2$ is $\hat{A}$. For $i = 1, 2, j = 0, 1, 2$, let $G_{i,j}$ denote the preconditioned operators for two- and three-level variants of BDDC, where i and j denote using M_i and $M_{s,j}$ (with $M_{s,0} := \hat{S}_\Pi$) as preconditioners for the fine and coarse problems, respectively, where $M_{s,j}^{-1}$ stands for applying the preconditioner M_j to the Schur complement problem. By standard calculation, we can write

$$G_{i,j} = \mathcal{R}_i^T\mathcal{K}_2^{-1}\mathcal{P}_j\mathcal{K}_1^{-1}\mathcal{R}_i A,$$

with

$$\mathcal{P}_j = \begin{pmatrix} I & 0 \\ 0 & M_{s,j}^{-1}\hat{S}_\Pi \end{pmatrix}.$$

Remark 3.1. When $j = 0$, $G_{i,j}$ is a two-level method, solving the Schur complement problem exactly, as $\mathcal{P}_0 \equiv I$. Note that, for the three-level variants ($j = 1, 2$),

$$\mathcal{P}_j\mathcal{K}_1^{-1} = \begin{pmatrix} I & 0 \\ 0 & M_{s,j}^{-1}\hat{S}_\Pi \end{pmatrix}\begin{pmatrix} A_{rr}^{-1} & 0 \\ -\hat{S}_\Pi^{-1}\hat{A}_{\Pi r}A_{rr}^{-1} & \hat{S}_\Pi^{-1} \end{pmatrix} = \begin{pmatrix} A_{rr}^{-1} & 0 \\ -M_{s,j}^{-1}\hat{A}_{\Pi r}A_{rr}^{-1} & M_{s,j}^{-1} \end{pmatrix}.$$

Thus, $G_{i,j}$ can be applied without directly applying the inverse of $\hat{S}_\Pi$.

Standard bounds (see, e.g., [31]) on the condition numbers of the three-level preconditioned operators are that there exists $\mathfrak{C}_{i,j}$ such that $\kappa(G_{i,j}) \leq \mathfrak{C}_{i,j}\Upsilon_i\Upsilon_j$, where $\Upsilon_1 = \frac{H}{h}(1 + \log\frac{H}{h})$ and $\Upsilon_2 = (1 + \log\frac{H}{h})^2$.

3.2. Multiplicative Preconditioners. As we shall see, the bounds above are relatively sharp and the performance of both preconditioners degrades with subdomain size and number of levels. To attempt to counteract this, we consider multiplicative combinations of these preconditioners with a simple diagonal scaling operator,

mimicking the use of weighted Jacobi relaxation in classical multigrid methods. We use $G^f_{i,j}$ to denote the multiplicative preconditioned operator based on $G_{i,j}$ with diagonal scaling on the fine level. Here,

$$G^f_{i,j} = G_{i,j} + \omega D^{-1} A(I - G_{i,j}),\ i = 1, 2,\ j = 0, 1, 2, \tag{3.2}$$

where D is the diagonal of A and ω is a chosen relaxation parameter.

Another variant is the use of multiplicative preconditioning on the coarse level with a similar diagonal scaling. We use $G^c_{i,j}$ to denote the resulting multiplicative preconditioner. Here,

$$G^c_{i,j} = \mathcal{R}_i^T \mathcal{K}_2^{-1} \mathcal{P}_j^c \mathcal{K}_1^{-1} \mathcal{R}_i A,\ i, j = 1, 2, \tag{3.3}$$

where

$$\mathcal{P}_j^c = \begin{pmatrix} I & 0 \\ 0 & G_{c,j} \end{pmatrix},$$

in which

$$G_{c,j} = M_{s,j}^{-1} \hat{S}_\Pi + \omega D_s^{-1} \hat{S}_\Pi (I - M_{s,j}^{-1} \hat{S}_\Pi),$$

where D_s is the diagonal of $\hat{S}_\Pi$.

Instead of using a single sweep of Jacobi in $G_{c,j}$, we can consider a symmetrized Jacobi operator $G^s_{c,j}$, where $I - G^s_{c,j} = (I - \omega_1 D_s^{-1} \hat{S}_\Pi)(I - M_{s,j}^{-1} \hat{S}_\Pi)(I - \omega_2 D_s^{-1} \hat{S}_\Pi)$; that is,

$$G^s_{c,j} = G_{c,j} + \omega_2 (I - G_{c,j}) D_s^{-1} \hat{S}_\Pi,$$

then $G^c_{i,j}$ changes to

$$G^{s,c}_{i,j} = \mathcal{R}_i^T \mathcal{K}_2^{-1} \mathcal{P}_j^{s,c} \mathcal{K}_1^{-1} \mathcal{R}_i A,\ i, j = 1, 2. \tag{3.4}$$

When $\omega_1 = \omega_2$, $G^{s,c}_{i,j}$ is a symmetric preconditioner for A, although we note that our LFA predicts a positive real spectrum for the nonsymmetric forms, $G^f_{i,j}$ and $G^c_{i,j}$, as well.

Finally, we can also apply the multiplicative operators based on diagonal scaling on both the fine and coarse levels. We denote this as

$$G^{f,c}_{i,j} = G^c_{i,j} + \omega_2 D^{-1} A(I - G^c_{i,j}),\ i = 1, 2,\ j = 1, 2, \tag{3.5}$$

where D is the diagonal of A and ω_2 is a chosen relaxation parameter.

In the following, we focus on analyzing the spectral properties of the above preconditioned operators by local Fourier analysis [28]. The main focus of this work is on the operators $\mathcal{K}_1, \mathcal{K}_2$, and $\mathcal{P}_j$, because the Fourier representations of other operators are just combinations of these three and some simple additional terms.

4. Local Fourier Analysis. To apply LFA to the BDDC-like methods proposed here, we first review some terminology of classical LFA. We consider a two-dimensional infinite uniform grid, $\mathbf{G}_h$, with

$$\mathbf{G}_h = \left\{ \boldsymbol{x}_{i,j} := (x_i, x_j) = (ih, jh), (i, j) \in \mathbb{Z}^2 \right\}, \tag{4.1}$$

and Fourier functions $\psi(\boldsymbol{\theta}, \boldsymbol{x}_{i,j}) = e^{\iota \boldsymbol{\theta}\cdot\boldsymbol{x}_{i,j}/h}$ on $\mathbf{G}_h$, where $\iota^2 = -1$ and $\boldsymbol{\theta} = (\theta_1, \theta_2)$. Let L_h be a Toeplitz operator acting on $\mathbf{G}_h$ as

$$L_h \stackrel{\wedge}{=} [s_{\boldsymbol{\kappa}}]_h \ (\boldsymbol{\kappa} = (\kappa_1, \kappa_2) \in \mathbb{Z}^2);\ L_h w_h(\boldsymbol{x}) = \sum_{\boldsymbol{\kappa}\in \boldsymbol{V}} s_{\boldsymbol{\kappa}} w_h(\boldsymbol{x} + \boldsymbol{\kappa} h),$$

with constant coefficients $s_{\boldsymbol{\kappa}} \in \mathbb{R}$ (or $\mathbb{C}$), where $w_h(\boldsymbol{x})$ is a function on $\mathbf{G}_h$. Here, $\boldsymbol{V}$ is taken to be a finite index set. Note that since L_h is Toeplitz, it is diagonalized by the Fourier modes $\psi(\boldsymbol{\theta}, \boldsymbol{x})$.

DEFINITION 4.1. *If for all grid functions, $\psi(\boldsymbol{\theta}, \boldsymbol{x})$,*

$$L_h \psi(\boldsymbol{\theta}, \boldsymbol{x}) = \widetilde{L}_h(\boldsymbol{\theta}) \psi(\boldsymbol{\theta}, \boldsymbol{x}),$$

we call $\widetilde{L}_h(\boldsymbol{\theta}) = \sum_{\boldsymbol{\kappa}\in \boldsymbol{V}} s_{\boldsymbol{\kappa}} e^{\iota \boldsymbol{\theta}\boldsymbol{\kappa}}$ the symbol of L_h.

Remark 4.2. In Definition 4.1, the operator L_h acts on a single function on $\mathbf{G}_h$, so $\widetilde{L}_h$ is a scalar. For an operator mapping vectors on $\mathbf{G}_h$ to vectors on $\mathbf{G}_h$, the symbol will be extended to be a matrix.

4.1. Change of Fourier Basis. Here, we discuss domain decomposition methods. While the classical basis set for LFA, denoted $\mathbf{E}_h$ below, could be used, we find it is substantially more convenient to make use of a transformed "sparse" basis, introduced here as $\mathbf{F}_H$. This basis allows a natural expression of the periodic structures in domain decomposition preconditioners. We treat each subdomain problem as one macroelement patch, and each subdomain block in the global problem is diagonalized by a coupled set of Fourier modes introduced in the following. Because each subdomain has the same size, $p \times p$, we consider the high and low frequencies for coarsening by factor p, given by

$$\boldsymbol{\theta} \in T^{\text{low}} = \left[-\frac{\pi}{p}, \frac{\pi}{p}\right)^2,\ \boldsymbol{\theta} \in T^{\text{high}} = \left[-\frac{\pi}{p}, \frac{(2p-1)\pi}{p}\right)^2 \setminus \left[-\frac{\pi}{p}, \frac{\pi}{p}\right)^2.$$

Let $\boldsymbol{\theta}^{(q,r)} = (\theta_1^{(q)}, \theta_2^{(r)})$, where $\theta_1^{(q)} = \theta_1^{(0)} + \frac{2\pi q}{p}$ and $\theta_2^{(r)} = \theta_2^{(0)} + \frac{2\pi r}{p}$ for $0 \le q, r < p$. For any given $\boldsymbol{\theta}^{(0,0)} \in T^{\text{low}}$, we define the p^2-dimensional space

$$\mathbf{E}_h(\boldsymbol{\theta}^{(0,0)}) := \text{span}\{\psi(\boldsymbol{\theta}^{(q,r)}, \boldsymbol{x}_{s,t}) = e^{\iota \boldsymbol{\theta}^{(q,r)}\cdot\boldsymbol{x}_{s,t}/h} : q, r = 0, 1, \cdots, p-1\}, \tag{4.2}$$

as the classical space of Fourier harmonics for factor p coarsening.

For any $\boldsymbol{x}_{s,t} \in \mathbf{G}_h$, we consider a grid function defined as a linear combination of the p^2 basis functions for $\mathbf{E}_h(\boldsymbol{\theta}^{(0,0)})$ with frequencies $\{\boldsymbol{\theta}^{(q,r)}\}_{q,r=0}^{p-1}$ and coefficients $\{\beta_{q,r}\}_{q,r=0}^{p-1}$ as

$$e_{s,t} := \sum_{q,r=0}^{p-1} \beta_{q,r} \psi(\boldsymbol{\theta}^{(q,r)}, \boldsymbol{x}_{s,t}).$$

We note that any index (s,t) has a unique representation as $(pm + k, pn + \ell)$ where

$(m,n) \in \mathbb{Z}^2$ and $k, \ell \in \{0, 1, \cdots, p-1\}$. From (4.2), we have

$$\begin{aligned} e_{pm+k,pn+\ell} &= \sum_{q,r=0}^{p-1} \beta_{q,r} e^{\iota(\theta_1^{(0)} + \frac{2\pi q}{p})x_s/h} e^{\iota(\theta_2^{(0)} + \frac{2\pi r}{p})x_t/h} \\ &= \sum_{q,r=0}^{p-1} \beta_{q,r} e^{\iota\theta_1^{(0)} x_s/h} e^{\iota\frac{2\pi q(pm+k)}{p}} e^{\iota\theta_2^{(0)} x_t/h} e^{\frac{2\pi r(pn+\ell)}{p}} \\ &= \sum_{q,r=0}^{p-1} \beta_{q,r} e^{\iota\frac{2\pi qk}{p}} e^{\iota\theta_1^{(0)} x_s/h} e^{\frac{2\pi r\ell}{p}} e^{\iota\theta_2^{(0)} x_t/h} \\ &= \left(\sum_{q,r=0}^{p-1} \beta_{q,r} e^{\iota\frac{2\pi qk}{p}} e^{\frac{2\pi r\ell}{p}} \right) \left(e^{\iota\boldsymbol{\theta}^{(0,0)} \cdot \boldsymbol{x}_{s,t}/h} \right). \end{aligned}$$

Thus, we can write

$$e_{pm+k,pn+\ell} = \hat{\beta}_{k,\ell} e^{\iota\boldsymbol{\theta}\cdot\boldsymbol{x}_{s,t}/H}, \tag{4.3}$$

with

$$\boldsymbol{\theta} = p\boldsymbol{\theta}^{(0,0)}, \text{ and } \hat{\beta}_{k,\ell} = \sum_{q,r=0}^{p-1} \beta_{q,r} e^{\iota\frac{2\pi qk}{p}} e^{\frac{2\pi r\ell}{p}}. \tag{4.4}$$

Thus, for any point (s,t) with $\text{mod}(s,p) = k$ and $\text{mod}(t,p) = \ell$, $e_{s,t}$ can be reconstructed from a single Fourier mode with coefficient $\hat{\beta}_{k,\ell}$. Thus, on the mesh $\mathbf{G}_h$ defined in (4.1), the periodicity of the basis functions in $\mathbf{E}_h(\boldsymbol{\theta}^{(0,0)})$ can also be represented by a pointwise basis on each $p \times p$-block.

Based on (4.3), we consider a "sparse" p^2-dimensional space as follows

$$\mathbf{F}_H(\boldsymbol{\theta}) := \text{span}\{\varphi_{k,\ell}(\boldsymbol{\theta}, \boldsymbol{x}_{s,t}) = e^{\iota\boldsymbol{\theta}\cdot\boldsymbol{x}_{s,t}/H} \chi_{k,\ell}(\boldsymbol{x}_{s,t}) : k, \ell = 0, 1, \cdots, p-1\}, \tag{4.5}$$

where $\boldsymbol{\theta} \in [-\pi, \pi)$ and

$$\chi_{k,\ell}(\boldsymbol{x}_{s,t}) = \begin{cases} 1, & \text{if } \text{mod}(s,p) = k, \text{ and } \text{mod}(t,p) = \ell, \\ 0, & \text{otherwise.} \end{cases}$$

Note that, with this notation, (4.3) can be rewritten as

$$e_{pm+k,pn+\ell} = \hat{\beta}_{k,\ell} \varphi_{k,\ell}(\boldsymbol{\theta}, \boldsymbol{x}_{s,t}). \tag{4.6}$$

THEOREM 4.3. *$\mathbf{E}_h(\boldsymbol{\theta}^{(0,0)})$ and $\mathbf{F}_H(p\boldsymbol{\theta}^{(0,0)})$ are equivalent.*

Proof. While the derivation above shows directly that $\mathbf{E}_h(\boldsymbol{\theta}^{(0,0)}) \subset \mathbf{F}_H(p\boldsymbol{\theta}^{(0,0)})$, we revisit this calculation now to show that the mapping $\{\beta_{q,r}\} \to \{\hat{\beta}_{k,\ell}\}$ is invertible and, hence, $\mathbf{F}_H(p\boldsymbol{\theta}^{(0,0)}) \subset \mathbf{E}_h(\boldsymbol{\theta}^{(0,0)})$ as well.

Let $\mathcal{X}$ be an arbitrary vector with size $p^2 \times 1$, denoted as

$$\mathcal{X} = \begin{pmatrix} \mathcal{X}_0 & \mathcal{X}_1 & \cdots & \mathcal{X}_{p-2} & \mathcal{X}_{p-1} \end{pmatrix}^T,$$

where

$$\mathcal{X}_r = \begin{pmatrix} \beta_{0,r} & \beta_{1,r} & \cdots & \beta_{p-2,r} & \beta_{p-1,r} \end{pmatrix}, \quad r = 0, 1, \cdots, p-1.$$

Then, we define a $p^2 \times 1$ vector, $\hat{\mathcal{X}}$, based on (4.4), as follows

$$\hat{\mathcal{X}} = \begin{pmatrix} \hat{\mathcal{X}}_0 & \hat{\mathcal{X}}_1 & \cdots & \hat{\mathcal{X}}_{p-2} & \hat{\mathcal{X}}_{p-1} \end{pmatrix}^T,$$

where

$$\hat{\mathcal{X}}_\ell = \begin{pmatrix} \hat{\beta}_{0,\ell} & \hat{\beta}_{1,\ell} & \cdots & \hat{\beta}_{p-2,\ell} & \hat{\beta}_{p-1,\ell} \end{pmatrix}, \quad \ell = 0, 1, \cdots, p-1,$$

in which

$$\hat{\beta}_{k,\ell} = \sum_{r=0}^{p-1} \Big(\sum_{q=0}^{p-1} \beta_{q,r} e^{\iota \frac{2\pi q k}{p}} \Big) e^{\frac{2\pi r \ell}{p}}, \quad q, r = 0, 1, \cdots, p-1.$$

Let $\mathcal{T}$ be the matrix of this transformation, $\hat{\mathcal{X}} = \mathcal{T}\mathcal{X}$, and

$$\mathcal{T}_1 = \begin{pmatrix} (e^{\iota \frac{2\pi}{p} 0})^0 & (e^{\iota \frac{2\pi}{p} 1})^0 & (e^{\iota \frac{2\pi}{p} 2})^0 & \cdots & (e^{\iota \frac{2\pi}{p} (p-1)})^0 \\ (e^{\iota \frac{2\pi}{p} 0})^1 & (e^{\iota \frac{2\pi}{p} 1})^1 & (e^{\iota \frac{2\pi}{p} 2})^1 & \cdots & (e^{\iota \frac{2\pi}{p} (p-1)})^1 \\ (e^{\iota \frac{2\pi}{p} 0})^2 & (e^{\iota \frac{2\pi}{p} 1})^2 & (e^{\iota \frac{2\pi}{p} 2})^2 & \cdots & (e^{\iota \frac{2\pi}{p} (p-1)})^2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ (e^{\iota \frac{2\pi}{p} 0})^{p-2} & (e^{\iota \frac{2\pi}{p} 1})^{p-2} & (e^{\iota \frac{2\pi}{p} 2})^{p-2} & \cdots & (e^{\iota \frac{2\pi}{p} (p-1)})^{p-2} \\ (e^{\iota \frac{2\pi}{p} 0})^{p-1} & (e^{\iota \frac{2\pi}{p} 1})^{p-1} & (e^{\iota \frac{2\pi}{p} 2})^{p-1} & \cdots & (e^{\iota \frac{2\pi}{p} (p-1)})^{p-1} \end{pmatrix}.$$

Note that $\mathcal{T}_1 \mathcal{X}_r$ defines a vector whose $(k+1)$-th entry is $\sum_{q=0}^{p-1} \beta_{q,r} e^{2\pi q k/p}$ and, thus, we see that $\mathcal{T} = \mathcal{T}_1 \otimes \mathcal{T}_1$.

Note that $\mathcal{T}_1$ is a $p \times p$ Vandermonde matrix based on values $d_k = e^{\iota \frac{2\pi k}{p}}$, where $k = 0, 1, 2, \cdots, p-1$. It is obvious that $d_j \neq d_k$ if $j \neq k$. Consequently, $\det(\mathcal{T}_1) \neq 0$. Thus, $\mathcal{T}_1$ is invertible, and so is $\mathcal{T}$. It follows that $\mathbf{E}_h(\boldsymbol{\theta}^{(0,0)})$ and $\mathbf{F}_H(p\boldsymbol{\theta}^{(0,0)})$ are equivalent. ∎

Remark 4.4. Let $z = e^{\iota 2\pi/p}$, be the primitive p-th root of unity, and note that $(\mathcal{T}_1)_{i,j} = z^{(j-1)(i-1)}$. Thus, $\widetilde{\mathcal{T}}_1 = \frac{1}{\sqrt{p}} \mathcal{T}_1$ is the unitary discrete Fourier transform (DFT) matrix with $\widetilde{\mathcal{T}}_1^{-1} = \widetilde{\mathcal{T}}_1^T$, where T denotes the conjugate transpose. Thus, $\mathcal{T}_1^{-1} = \frac{1}{p} \mathcal{T}_1^T$. Similarly, $\mathcal{T}$ is a scaled version of the two-dimensional unitary Fourier transform matrix, and $\mathcal{T}^{-1} = \frac{1}{p^2} \mathcal{T}^T$.

In the rest of this paper, we use the basis of $\mathbf{F}_H$ as the foundation for local Fourier analysis on the $p \times p$ periodic structures of the BDDC operators. The "sparse" (or "pointwise") nature of the basis in $\mathbf{F}_H$ allows a natural expression of the operators in BDDC and, as such, is more convenient than the equivalent "global" basis in $\mathbf{E}_h$.

Note that the presentation above assumes that the original Fourier space, $\mathbf{E}_h$, is considered with harmonic frequencies in domain $[-\frac{\pi}{p}, \frac{(2p-1)\pi)}{p})^2$, and the sparse basis in $\mathbf{F}_H$ considers a single mode, $\boldsymbol{\theta} \in [-\pi, \pi)^2$. In both cases, it is clear that any frequency set covering an interval of length 2π in both x and y components can be used instead.

4.2. Representation of the Original Problem. On $\mathbf{G}_h$, we call each node, (I, J), where $\text{mod}(I, p) = 0$ and $\text{mod}(J, p) = 0$ a coarse-level point index. We construct a collective grid set associated with (I, J) for each subdomain as

$$\mathfrak{S}_{I,J} = \big\{ \boldsymbol{x}_{(I+k, J+\ell)} : k, \ell = 0, 1, \cdots, p-1 \big\}. \tag{4.7}$$

The degrees of freedom in A can be divided into subsets, $\mathfrak{S}_{I,J}$, whose union provides a disjoint cover for the set of degrees of freedom on the infinite mesh $\mathbf{G}_h$. Throughout the rest of this paper, the index (I, J) corresponds to the coarse point at the lower-left corner of the subdomain under consideration, unless stated otherwise. The left of Figure 2 shows the meshpoints for this decomposition for $p = 4$.

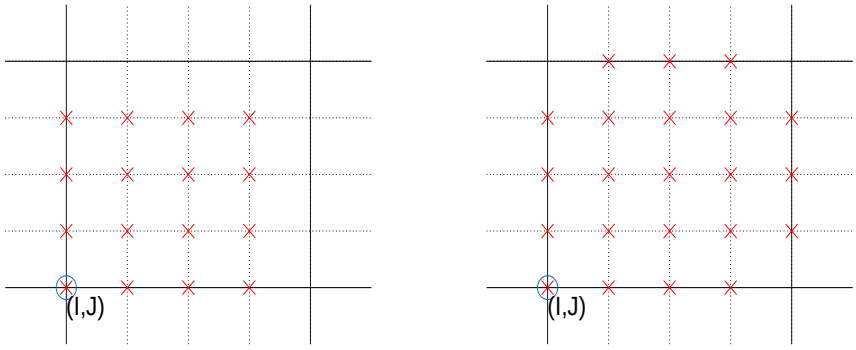

FIG. 2. *At left, the location of degrees of freedom in $\mathfrak{S}_{I,J}$ defined in Equation (4.7) for one subdomain with $p = 4$. At right, the location of degrees of freedom in $\mathfrak{S}^*_{I,J}$ defined in Equation (4.11) for one subdomain with $p = 4$.*

For each $\mathfrak{S}_{I,J}$, we use a row-wise ordering of the grid points (lexicographical ordering). This will fix the ordering of the symbols in the following; for any other ordering, a permutation operator would need to be applied. In the following, we do not show the specific position of each element in a vector or matrix, and they are assumed to be consistent with the ordering of the grid points. Based on the set $\mathfrak{S}_{I,J}$, we define the p^2-dimensional space

$$\mathcal{E}(\boldsymbol{\theta}) = \operatorname{span}\big\{\boldsymbol{\varphi}_{k,\ell}(\boldsymbol{\theta}) : k, \ell = 0, 1, \cdots, p-1\big\}, \tag{4.8}$$

where $\boldsymbol{\varphi}_{k,\ell}(\boldsymbol{\theta}) = \big(\varphi_{k,\ell}(\boldsymbol{\theta}, \boldsymbol{x}_{I+s,J+t})\big)_{s,t=0}^{p-1}$ is a $p^2 \times 1$ vector with only one nonzero element, defined in (4.5), in the position corresponding to $(I + k, J + \ell)$. For both $\mathcal{E}(\boldsymbol{\theta})$ and $\boldsymbol{\varphi}_{k,\ell}(\boldsymbol{\theta})$, we have simply taken the infinite mesh representation of $\mathbf{F}_H$ and truncated it to a single $p \times p$ block of the mesh, which is sufficient to define the symbol of A in this basis. Let Φ_h be a $p^2 \times p^2$ diagonal matrix, whose diagonal elements are functions $\varphi(\boldsymbol{\theta}, \boldsymbol{x}) = e^{\iota \boldsymbol{\theta} \cdot \boldsymbol{x}/H}$, where $\boldsymbol{x} \in \mathfrak{S}_{I,J}$, so $\mathcal{E}(\boldsymbol{\theta}) = \operatorname{Range}(\Phi_h)$.

Note that each subdomain contains p^2 degrees of freedom, and that the corresponding symbol is not a scalar due to the definition of the Fourier basis in (4.8). We treat the block symbol as a system, presented as a $p^2 \times p^2$ matrix. Let $A_{I,J}$ be the periodic Laplace operator on $\mathfrak{S}_{I,J}$. Then, its symbol $\widetilde{A}$ satisfies

$$A_{I,J}\boldsymbol{\phi}(\boldsymbol{\theta}, \boldsymbol{x}) = \widetilde{A}(\boldsymbol{\theta})\boldsymbol{\phi}(\boldsymbol{\theta}, \boldsymbol{x}), \ \forall \boldsymbol{\phi}(\boldsymbol{\theta}, \boldsymbol{x}) \in \mathcal{E}(\boldsymbol{\theta}), \tag{4.9}$$

where $\widetilde{A}$ is a $p^2 \times p^2$ matrix. Equation (4.9) is equivalent to

$$A_{I,J} \sum_{0 \le k,\ell \le p-1} \alpha_{k,\ell} \boldsymbol{\varphi}_{k,\ell} = \Phi_h \widetilde{A} \boldsymbol{\alpha}, \tag{4.10}$$

for any vector $\boldsymbol{\alpha}$, whose elements are denoted as $\alpha_{k,\ell}$. Since (4.10) holds for any $\alpha_{k,\ell}$, we have $\widetilde{A} = \Phi_h^T A_{I,J} \Phi_h$, where T is the (conjugate) transpose. Note that $\Phi_h^{-1} = \Phi_h^T$ and the entries in these matrices have the same form, $e^{\pm \iota \boldsymbol{\theta} \cdot \boldsymbol{x}_{I,J}/H}$.

We consider the action of $\widetilde{A}(\boldsymbol{\theta})$ on a vector in terms of the coefficients of the Fourier basis functions. Considering a point in $\mathfrak{S}_{I,J}$, if the values of a function at neighbouring points are expressed by $\alpha_{k,\ell}\boldsymbol{\varphi}_{k,\ell}$, the entries in $\widetilde{A}(\boldsymbol{\theta})\boldsymbol{\alpha}$ give the coefficients of the Fourier expansion of the original operator A on G_h acting on the function in $\mathcal{E}(\boldsymbol{\theta})$ with coefficient $\boldsymbol{\alpha}$. We note that a similar approach was employed for LFA for vector finite-element discretizations in [21].

4.3. Representation of Preconditioned Operators. Now we turn to calculating the Fourier representations of M_1^{-1} and M_2^{-1}. First, we define a collective grid set associated with (I, J) for the partially subassembled problem for each subdomain as

$$\mathfrak{S}^*_{I,J} = \{\boldsymbol{x}_{(I+k,J+\ell)} : k, \ell = 0, 1, \cdots, p\} \setminus \{\boldsymbol{x}_{(I+p,J)}, \boldsymbol{x}_{(I,J+p)}, \boldsymbol{x}_{(I+p,J+p)}\}, \tag{4.11}$$

see the right of Figure 2. We first consider the stencil of M_1^{-1} acting on one subdomain, $\mathfrak{S}^*_{I,J}$.

Recall the scaling operator, $\mathcal{R}_1$, where each column of $\mathcal{R}_1$ corresponding to a degree of freedom of the global problem in the interiors and at the coarse-grid points has a single nonzero entry with value 1, and each column of $\mathcal{R}_1$ corresponding to an interface degree of freedom has two nonzero entries, each with value $\frac{1}{2}$. Since we consider periodic Fourier modes on each subdomain, the interface degrees of freedom share the same values scaled by an exponential shift. For example, at the left of Figure 2, the degrees of freedom located at the left boundary and the right boundary have the same coefficient of the (shifted) exponential, as do the degrees of freedom located at the bottom and top. Thus, $\mathcal{R}_1$ is its own Fourier representation, since the neighborhoods do not contribute to each other. Note that $\mathcal{R}_1$ maps the p^2-dimensional Fourier basis from $\mathcal{E}(\boldsymbol{\theta})$, used to express $\widetilde{A}(\boldsymbol{\theta})$ onto a $(p+1)^2 - 3$ dimensional space with similar sparse basis on $\mathfrak{S}^*_{I,J}$ that is suitable for expressing the symbol of $\hat{A}$ and its inverse.

We now focus on $\hat{A}$ presented on one subdomain. Let $\hat{A}^{(I,J)}$ be a $(p+1)^2 \times (p+1)^2$ matrix, which is the partially subassembled problem on one subdomain including its four neighbouring coarse-grid degrees of freedom, as

$$\hat{A}^{(I,J)} = \begin{pmatrix} A_{rr}^{(I,J)} & \left(\hat{A}_{\Pi r}^{(I,J)}\right)^T \\ \hat{A}_{\Pi r}^{(I,J)} & A_{\Pi\Pi}^{(I,J)} \end{pmatrix} = \begin{pmatrix} A_{rr}^{(I,J)} & 0 \\ \hat{A}_{\Pi r}^{(I,J)} & \hat{S}_{\Pi}^{(I,J)} \end{pmatrix} \begin{pmatrix} I & \left(A_{rr}^{(I,J)}\right)^{-1} \left(\hat{A}_{\Pi r}^{(I,J)}\right)^T \\ 0 & I \end{pmatrix}, \tag{4.12}$$

where $A_{rr}^{(I,J)}$ is a $\left((p+1)^2 - 4\right) \times \left((p+1)^2 - 4\right)$ matrix corresponding to the interior and interface degrees of freedom on the subdomain and $A_{\Pi\Pi}^{(I,J)}$ corresponds to the four coarse-level variables on one subdomain. Note that $A_{\Pi\Pi}^{(I,J)} = \frac{2}{3} I$ and $\hat{S}_{\Pi}^{(I,J)} = A_{\Pi\Pi}^{(I,J)} - \hat{A}_{\Pi r}^{(I,J)} (A_{rr}^{(I,J)})^{-1} (\hat{A}_{\Pi r}^{(I,J)})^T$. We use index (I, J) as a superscript in order to distinguish this from the matrix in (3.1), but note that it is independent of the particular subdomain, (I, J), under consideration. Let $\widetilde{\hat{A}}$ be the Fourier representation of the partially subassembled problem with the corresponding symbol being a $\big((p+$

$1)^2-3\big)\times\big((p+1)^2-3\big)$ matrix,

$$\widetilde{\hat{A}}=\begin{pmatrix}\widetilde{A}_{rr} & 0\\ \widetilde{A}_{\Pi r} & \widetilde{S}_\Pi\end{pmatrix}\begin{pmatrix}\widetilde{I} & (\widetilde{A}_{rr})^{-1}\widetilde{A}_{\Pi r}^T\\ 0 & \widetilde{I}\end{pmatrix}=\widetilde{\mathcal{K}}_1\widetilde{\mathcal{K}}_2,$$

where $\widetilde{A}_{rr}$ is a $\big((p+1)^2-4\big)\times\big((p+1)^2-4\big)$ Fourier representation of $A_{rr}^{(I,J)}$ computed as was done for $\widetilde{A}$ above and $\widetilde{S}_\Pi$ is the representation of the global Schur complement, $\hat{S}_\Pi$. Let $S_0=\hat{A}_{\Pi r}^{(I,J)}(A_{rr}^{(I,J)})^{-1}(\hat{A}_{\Pi r}^{(I,J)})^T$ be a 4×4 matrix corresponding to the vertices adjacent to one subdomain, representing one macroelement of the coarse-level variables. Direct calculation shows this matrix has the same nonzero structure as the element stiffness matrix for a symmetric second-order differential operator on a uniform square mesh, with equal values for the connections from each node to itself (denoted s_1), its adjacent vertices (s_2), and its opposite corner (s_3). Since $\hat{S}_\Pi^{(I,J)}=\frac{2}{3}I-S_0$ gives the macroelement stiffness contribution, assembling the coarse-level stiffness matrix over 2×2 macroelement patches yields $\widetilde{S}_\Pi$ as the symbol of the 9-point stencil given by

$$\begin{bmatrix}-s_3 & -2s_2 & -s_3\\ -2s_2 & \frac{8}{3}-4s_1 & -2s_2\\ -s_3 & -2s_2 & -s_3\end{bmatrix},$$

acting on the coarse points.

$\widetilde{A}_{\Pi r}$ is the representation of the contribution from interior and interface degrees of freedom to the coarse degrees of freedom, and has only 12-nonzero elements per subdomain, with 3 contributing to each corner of the subdomain. We take the coarse-level point $\boldsymbol{x}_{I,J}$ as an example. At the right of Figure 2, $\boldsymbol{x}_{I,J}$ obtains contributions from the points $\boldsymbol{x}_{I+1,J},\boldsymbol{x}_{I+1,J+1},\boldsymbol{x}_{I,J+1}$ and the corresponding stencils are

$$\begin{bmatrix}* & -\frac{1}{6}\end{bmatrix},\ \begin{bmatrix} & -\frac{1}{3}\\ * & \end{bmatrix},\ \begin{bmatrix}-\frac{1}{6}\\ *\end{bmatrix},$$

where $*$ denotes the position on the grid at which the discrete operator is applied, namely $\boldsymbol{x}_{I,J}$. The symbols of these three stencils are given by $-\frac{1}{6}e^{\iota\theta_1/p}$, $-\frac{1}{3}e^{\iota(\theta_1+\theta_2)/p}$, $-\frac{1}{6}e^{\iota\theta_2/p}$, respectively. Since $\boldsymbol{x}_{I,J}$ is adjacent to three other subdomains, the coarse degree of freedom at $\boldsymbol{x}_{I,J}$ also obtains contributions from those subdomains, and the other 9 contributing stencils are computed similarly. Finally, the representation of $M_1^{-1}A$ is given by

$$\widetilde{G}_{1,0}(\boldsymbol{\theta})=\widetilde{\mathcal{R}}_1^T(\widetilde{\hat{A}})^{-1}\widetilde{\mathcal{R}}_1\widetilde{A}=\widetilde{\mathcal{R}}_1^T\widetilde{\mathcal{K}}_2^{-1}\widetilde{\mathcal{K}}_1^{-1}\widetilde{\mathcal{R}}_1\widetilde{A}.$$

For the Dirichlet preconditioner in (2.6), we also need to know the LFA representation of the operators J_D and $\mathcal{H}$. Since J_D is a pointwise scaling operator, its symbol in the pointwise basis of $\mathbf{F}_H$ is itself. According to the definition of $\mathcal{H}$, the symbol of $\mathcal{H}$ is given by $\widetilde{\mathcal{H}}=\widetilde{A}_{rr,I}^{-1}\widetilde{A}_{\Gamma,I}^T$, where $\widetilde{A}_{rr,I}$ is the submatrix of $\widetilde{A}_{rr}$ corresponding to the interior degrees of freedom, and $\widetilde{A}_{\Gamma,I}^T$ is the submatrix of $\widetilde{\hat{A}}$ corresponding to the contribution of the interface degrees of freedom to the interior degrees of freedom. Both of these are computed in a similar manner to $\widetilde{A}$ and $\widetilde{\hat{A}}$ as described above. Thus, the LFA representation of $M_2^{-1}A$ can be written as

$$\widetilde{G}_{2,0}(\boldsymbol{\theta})=(\widetilde{\mathcal{R}}_1^T-\widetilde{\mathcal{H}}\widetilde{J}_D)\widetilde{\mathcal{K}}_2^{-1}\widetilde{\mathcal{K}}_1^{-1}(\widetilde{\mathcal{R}}_1-\widetilde{J}_D^T\widetilde{\mathcal{H}}^T)\widetilde{A}.$$

The details of the 3-level variants of LFA are similar to those given above. We now consider a segment of the infinite mesh given, on the fine level, by a $p \times p$ array of subdomains, with each subdomain of size $p \times p$ elements. On the first coarse level (corresponding to the Schur complement $\hat{S}_\Pi$ in (3.1)), we then consider a single $p \times p$ subdomain of the infinite coarse mesh, and apply the same technique recursively. To accommodate this, we adapt the fine-level Fourier modes to be $\varphi^*(\boldsymbol{\theta}, \boldsymbol{x}) := e^{\iota \boldsymbol{\theta} \cdot \boldsymbol{x}/H'}$, where $H' = p^2 h$. The coarse-level Fourier modes are then the same as (4.8). Thus, $\widetilde{G_{i,j}}(\boldsymbol{\theta})$ is a $p^4 \times p^4$ matrix for the three-level variants.

5. Numerical Results.

5.1. Condition Numbers of Two-level Variants. In the LFA setting, $\boldsymbol{\theta} = (\theta_1, \theta_2) \in [-\pi, \pi)^2$. Here we take $d\theta = \pi/n$ as the discrete stepsize and sample the Fourier space at $2n$ evenly distributed frequencies in θ_1 and θ_2 with offset $\pm d\theta/2$ from $\theta_1 = \theta_2 = 0$ to avoid the singularity at zero frequency. For each frequency on the mesh, we compute the eigenvalues of the two-level operators, and define $\kappa := \frac{e_{\max}}{e_{\min}}$, where $e_{\min}$ and $e_{\max}$ are the smallest and biggest eigenvalues over all frequencies.

Table 1 shows the condition numbers for the two-level preconditioners with variation in both subdomain size, p, and sampling frequency, n. When $n = 2$, the condition number prediction is notably inaccurate, but we obtain a consistent prediction for $n \geq 4$ (and very consistent for $n \geq 8$). For $\widetilde{G}_{1,0}$, the condition number increases quickly with p as expected. Compared with $\widetilde{G}_{1,0}$, $\widetilde{G}_{2,0}$ has a much smaller condition number that grows more slowly with p. For $\widetilde{G}_{1,0}$, we know there exists $\mathfrak{C}_{1,0}$ such that the true condition number of the preconditioned system (on a finite grid) is bounded by $\mathfrak{C}_{1,0}\frac{H}{h}(1+\log\frac{H}{h})$ [20]; from this data, we see that our LFA prediction is consistent with this, with constant $\mathfrak{C}_{1,0} \approx 0.6$. For $\widetilde{G}_{2,0}$, we know there exists $\mathfrak{C}_{2,0}$ such that the true condition number of the preconditioned system (on a finite grid) is bounded by $\mathfrak{C}_{2,0}(1+\log\frac{H}{h})^2$ [20]; from this data, again we see that our LFA prediction is consistent with this, with constant $\mathfrak{C}_{2,0} \approx 0.4$.

TABLE 1
LFA-predicted condition numbers of two-level preconditions as a function of subdomain size, p, and sampling frequency, n.

	$\widetilde{G}_{1,0}$				$\widetilde{G}_{2,0}$			
n \ p	4	8	16	32	4	8	16	32
2	4.14	11.11	27.95	67.55	2.23	3.02	3.94	5.01
4	4.36	11.94	30.27	73.44	2.32	3.15	4.13	5.26
8	4.42	12.18	30.94	75.16	2.34	3.19	4.17	5.32
16	4.44	12.25	31.12	75.61	2.35	3.19	4.19	5.33
32	4.44	12.26	31.16	75.72	2.35	3.20	4.19	5.34
64	4.44	12.27	31.17	75.75	2.35	3.20	4.19	5.34
128	4.44	12.27	31.18	75.76	2.35	3.20	4.19	5.34
$\mathfrak{C}_{i,0}(n = 32)$	0.47	0.50	0.52	0.53	0.41	0.34	0.29	0.27

Optimizing the weight parameters for $\widetilde{G}^f_{1,0}$ and $\widetilde{G}^f_{2,0}$ by systematic search with different n and p, we see that the optimal parameter ω is dependent on p, but largely independent of n. Table 2 shows that significant improvement can be had for the M_1 preconditioner, but not for M_2, see Table 3. We again see small n (e.g., $n = 4$ or 8) is enough to obtain a consistent prediction for these condition numbers.

TABLE 2
Condition numbers for two-level lumped preconditioner with fine-grid multiplicative combination with diagonal scaling, $\widetilde{G}^f_{1,0}$. In brackets, value of weight parameter, ω, that minimizes condition number.

n \ p	4	8	16	32
2	2.06(2.1)	3.18(2.3)	5.43(2.5)	9.71(2.6)
4	2.17(1.5)	3.29(2.3)	5.64(2.5)	9.99(2.6)
8	2.18(1.4)	3.32(2.3)	5.70(2.5)	10.08(2.6)
16	2.18(1.4)	3.32(2.3)	5.72(2.5)	10.10(2.6)
32	2.18(1.4)	3.33(2.3)	5.72(2.5)	10.10(2.6)
64	2.18(1.4)	3.33(2.3)	5.72(2.5)	10.10(2.6)

TABLE 3
Condition numbers for two-level Dirichlet preconditioner with fine-grid multiplicative combination with diagonal scaling, $\widetilde{G}^f_{2,0}$. In brackets, value of weight parameter, ω, that minimizes condition number.

n \ p	4	8	16	32
2	1.82(2.2)	2.36(1.7)	3.12(2.0)	4.20(1.8)
4	2.03(1.1)	2.54(1.6)	3.33(2.0)	4.44(1.8)
8	2.07(1.1)	2.59(1.6)	3.39(2.0)	4.50(1.8)
16	2.08(1.1)	2.60(1.6)	3.40(2.0)	4.52(1.8)
32	2.08(1.1)	2.60(1.6)	3.40(2.0)	4.52(1.8)
64	2.08(1.1)	2.61(1.6)	3.40(2.0)	4.52(1.8)

In order to see the sensitivity of performance to parameter choice, we consider the condition numbers for the two-level lumped and Dirichlet preconditioners in multiplicative combination with diagonal scaling on the fine grid with $p = 8$, as a function of ω, in Figure 3. We see that the condition number of $\widetilde{G}^f_{1,0}$ shows strong sensitivity to small values of ω. For $\widetilde{G}^f_{2,0}$, however, many allowable parameters obtain a good condition number.

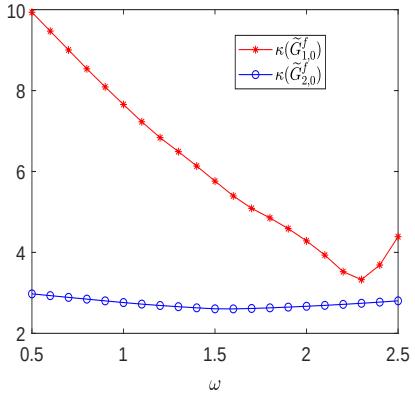

FIG. 3. *Condition numbers for two-level lumped and Dirichlet preconditioners in multiplicative combination with diagonal scaling on the fine grid with $p = 8$, as a function of relaxation parameter, ω.*

5.2. Eigenvalue Distribution of Two-level Variants. In this section, we take $n = 32$, yielding $2n$ points in each dimension and $(2n)^2 = 4096$ values of $\boldsymbol{\theta}$, although similar results are seen for smaller values of n. We also consider only $p = 8$, although similar results are seen for other values of p. For $\widetilde{G}^f_{1,0}$ and $\widetilde{G}^f_{2,0}$, we use the optimal values of ω, shown in the tables above. The histograms in Figure 4 show the density of eigenvalues for the two-level preconditioned operators. For these values of n and p, our LFA computes a total of 262144 eigenvalues, giving 64 eigenvalues for each of 4096 sampling points. For all cases, the eigenvalues around 1 (represented in two bins in the histogram, covering the interval from 0.9 to 1.1) appear with dominating multiplicity, accounting for about 200,000 of the computed eigenvalues.

Note that there is a gap in the spectrum of $\widetilde{G}_{1,0}$ that increases in size with p (not shown here). A notable difference between $\widetilde{G}_{1,0}$ and $\widetilde{G}_{2,0}$ is that, while there is still a small gap in the spectrum of $\widetilde{G}_{2,0}$, it is not very prominent. Note also that the spectra are real-valued, with only roundoff-level errors in the imaginary component. Comparing the eigenvalues for $\widetilde{G}^f_{1,0}$ and $\widetilde{G}^f_{2,0}$ with those for $\widetilde{G}_{1,0}$ and $\widetilde{G}_{2,0}$, we see that the eigenvalues are much more tightly clustered for $\widetilde{G}^f_{1,0}$, but still exhibit a gap in the spectrum. The eigenvalues of $\widetilde{G}^f_{2,0}$, in contrast, appear to lie in a continuous interval. We note that little improvement is seen in the spectrum of $\tilde{G}^f_{2,0}$, in comparison with $\widetilde{G}_{2,0}$. Also interesting to note is that, in contrast to all other cases, the smallest eigenvalue of $\widetilde{G}^f_{1,0}$ is less than 1.

Remark 5.1. As the LFA predicts both eigenvectors and eigenvalues, we can examine the frequency composition of the eigenvectors associated with these eigenvalues. The largest eigenvalue of $\widetilde{G}_{1,0}$ is found to be dominated by oscillatory modes, but this is not true for $\widetilde{G}_{2,0}$. This motivates the proposed multiplicative method based on simple diagonal scaling, which is well known to effectively damp oscillatory errors in the classical multigrid setting.

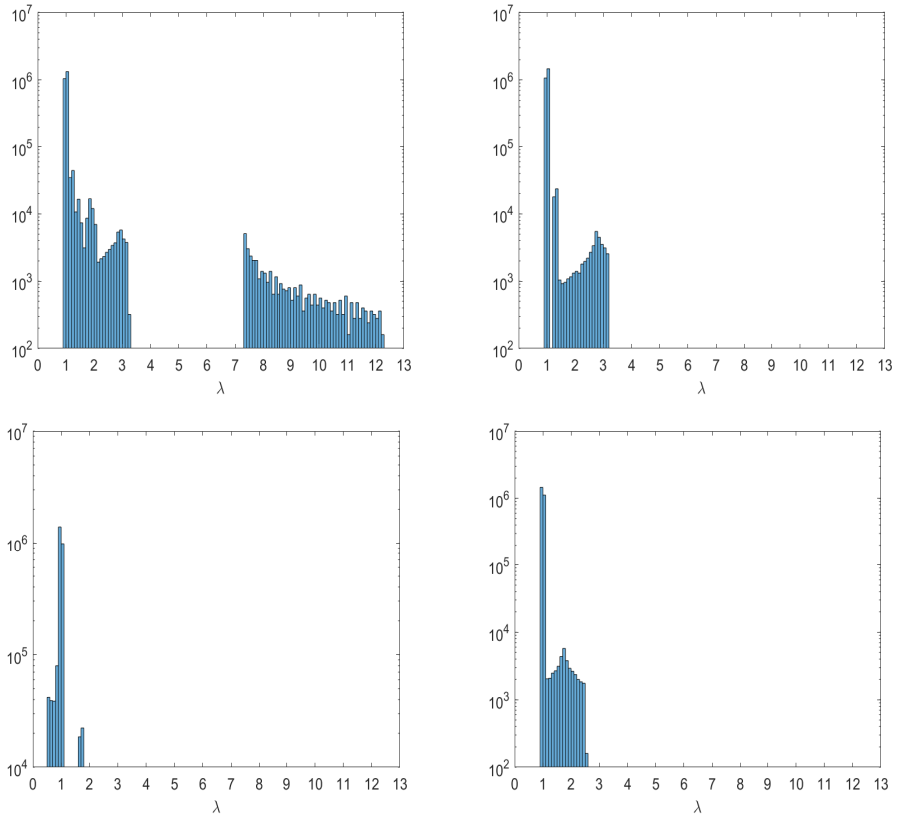

FIG. 4. *Histograms showing density of eigenvalues for two-level preconditioned operators with* $p = 8$. *Top left:* $\widetilde{G}_{1,0}$, *Top right:* $\widetilde{G}_{2,0}$, *Bottom left:* $\widetilde{G}^f_{1,0}$, *Bottom right:* $\widetilde{G}^f_{2,0}$.

5.3. Condition Numbers of Three-level Variants. For the three-level preconditioned operators, we need to find all the eigenvalues of a $p^4 \times p^4$ matrix for each sampled value of $\boldsymbol{\theta}$. For the two-level variants, we saw that sampling with $n = 4$ is sufficient to give useful accuracy of the LFA predictions. Here, we also see similar behavior in Table 4, which shows the condition numbers of $\widetilde{G}_{i,j}(i, j = 1, 2)$ for varying p an n. We see that, as expected from the theory, these condition numbers show degradation from the two-level case. It is not surprising that $\widetilde{G}_{2,2}$ has the smallest condition number of these variants, since M_2 is applied to both fine and coarse levels.

Table 5 presents the condition number of variants $\widetilde{G}^f_{i,j}$ and $\widetilde{G}^c_{i,j}$, based on the multiplicative combination with diagonal scaling on the fine level and coarse level, respectively, and some improvement is offered. For fixed p, the optimal ω is found to be robust to n (not shown here). In general, we see better performance for $\widetilde{G}^f_{i,j}$ in comparison to $\widetilde{G}^c_{i,j}$, and $\widetilde{G}^f_{1,1}$ offers significant improvement over $\widetilde{G}_{1,1}$. For other values of i, j, however, only small improvements are seen.

TABLE 4
Condition numbers of three-level preconditioners with no multiplicative relaxation.

p	$\widetilde{G}_{1,1}$	$\widetilde{G}_{1,2}$	$\widetilde{G}_{2,1}$	$\widetilde{G}_{2,2}$
$4(n=2)$	9.18	5.43	7.27	4.24
$4(n=4)$	9.65	5.68	7.63	4.47
$4(n=8)$	9.79	5.74	7.73	4.53
$4(n=16)$	9.82	5.76	7.76	4.54
$4(n=32)$	9.83	5.76	7.77	4.55
$8(n=2)$	46.66	15.46	24.73	7.55
$8(n=4)$	50.00	16.15	26.53	7.94
$8(n=8)$	50.96	16.33	27.05	8.04

TABLE 5
Condition numbers of three-level preconditioners with fine-scale or coarse-scale multiplicative preconditioning. All results were computed with $n=4$, and the experimentally optimized weight, ω, is shown in brackets.

p	$\widetilde{G}^f_{1,1}$	$\widetilde{G}^f_{1,2}$	$\widetilde{G}^f_{2,1}$	$\widetilde{G}^f_{2,2}$
4	6.80(1.4)	4.28(1.4)	6.14(1.6)	4.04(1.1)
8	28.75(1.7)	9.16(1.7)	20.94(1.6)	6.73(1.5)

p	$\widetilde{G}^c_{1,1}$	$\widetilde{G}^c_{1,2}$	$\widetilde{G}^c_{2,1}$	$\widetilde{G}^c_{2,2}$
4	6.04(1.6)	5.47(1.1)	4.67(1.6)	4.30(1.0)
8	31.91(2.0)	15.17(1.4)	15.57(2.1)	7.46(1.2)

In order to see the sensitivity of performance to parameter choice, we consider three-level preconditioners with weighted multiplicative preconditioning on both fine and coarse scales, $\widetilde{G}^{f,c}_{1,1}$ and $\widetilde{G}^{f,c}_{2,2}$, with $p=4$ and $n=4$. At the left of Figure 5, we present the LFA-predicted condition number for $\widetilde{G}^{f,c}_{1,1}$ with variation in ω_1 and ω_2. Here, we see strong sensitivity to "small" values of ω_1, for example $\omega_1 < 1.5$, and also to large values of ω_1 with small values of ω_2. We note general improvement, though, in the optimal performance for large ω_1 with suitably chosen ω_2, albeit with diminishing returns as ω_1 continues to increase. Fixing $\omega_1 = 4$, we find $\omega_2 = 1.7$ offers best performance, with optimal condition number of 2.66. At the right of Figure 5, we consider $\widetilde{G}^{f,c}_{2,2}$ as a function of ω_1 and ω_2. Here, we see stronger sensitivity to large values of ω_2, and to large values of ω_1 and small values of ω_2, but a large range of parameters that give generally similar performance. Fixing $\omega_1 = 4$, we find that $\omega_2 = 1.2$ achieves the optimal condition number of 3.72. Similar performance was seen for $\widetilde{G}^{f,c}_{1,2}, \widetilde{G}^{f,c}_{2,1}$, and $\widetilde{G}^{s,c}_{i,j}$. Slight improvements can be seen by allowing even larger values of ω_1, giving an LFA-predicted condition number for $\widetilde{G}^{f,c}_{1,1}$ of 2.25 with $\omega_1 = 5.0$ and $\omega_2 = 2.0$, but a much smaller band of values of ω_2 leads to near-optimal performance as ω_1 increases. For $\widetilde{G}^{f,c}_{2,2}$, this sensitivity does not arise, but the improvements are even more marginal, achieving an LFA-predicted condition number of 3.63 for $\omega_1 = 5.8$ and $\omega_2 = 1.3$.

Motivated by Figure 5, we fix $\omega_1 = 4$ with $n = 4$, and optimize the condition numbers for the three-level preconditioners with two multiplicative preconditioning steps per iteration, either both on the coarse level, $\widetilde{G}^{s,c}_{i,j}$, or one on each level, $\widetilde{G}^{f,c}_{i,j}$,

with respect to ω_2. From Table 6, notable improvement is seen for all i, j with $\widetilde{G}^{f,c}_{i,j}$, particularly for $\widetilde{G}^{f,c}_{1,1}$ and $\widetilde{G}^{f,c}_{2,1}$. We also note that there is little variation in the optimal parameter for each preconditioner between the $p = 4$ and $p = 8$ cases. It is notable that we are able to achieve similar performance for the multiplicative preconditioner based on M_1 as seen for M_2, and that both show significant improvement from the classical three-level results shown in Table 4, when used in combination with multiplicative preconditioning on both fine and coarse levels.

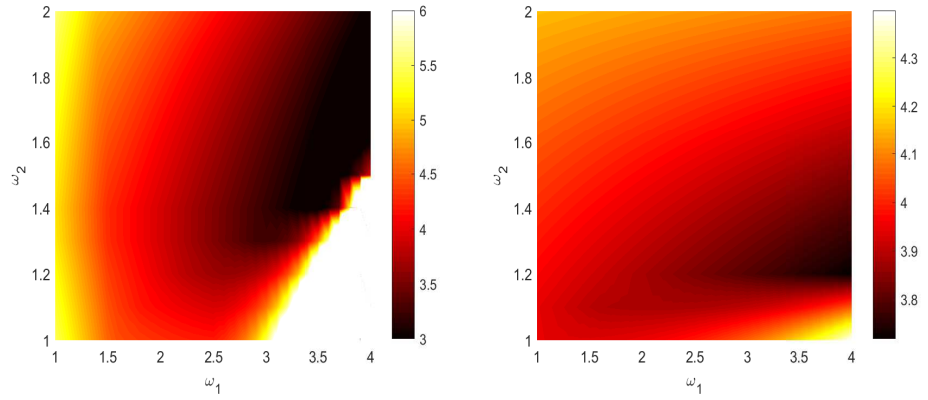

FIG. 5. *Condition number of three-level preconditioners with multiplicative preconditioning on both the fine and coarse scales as a function of ω_1 and ω_2, with $p = 4$ and $n = 4$. At left, condition number for $\widetilde{G}^{f,c}_{1,1}$; at right, condition number for $\widetilde{G}^{f,c}_{2,2}$.*

TABLE 6

Condition numbers of three-level preconditioners with symmetric weighting of multiplicative preconditioning on the coarse scale, $\widetilde{G}^{s,c}_{i,j}$, and weighting of multiplicative preconditioning on both fine and coarse scales, $\widetilde{G}^{f,c}_{i,j}$. All results were computed with $n = 4$, and the experimentally optimized weight, ω_2, is shown in brackets.

p	$\widetilde{G}^{s,c}_{1,1}$	$\widetilde{G}^{s,c}_{1,2}$	$\widetilde{G}^{s,c}_{2,1}$	$\widetilde{G}^{s,c}_{2,2}$
4	5.43(1.4)	5.34(0.9)	4.22(1.3)	4.18(0.9)
8	17.45(1.2)	14.13(1.0)	8.31(1.1)	6.88(0.9)

p	$\widetilde{G}^{f,c}_{1,1}$	$\widetilde{G}^{f,c}_{1,2}$	$\widetilde{G}^{f,c}_{2,1}$	$\widetilde{G}^{f,c}_{2,2}$
4	2.66(1.7)	3.85(1.3)	3.24(1.8)	3.72(1.2)
8	5.16(1.8)	7.59(1.7)	4.88(1.8)	5.70(1.5)

6. Conclusions. In this paper, we quantitatively estimate the condition numbers of variants of BDDC algorithms, using local Fourier analysis. A new choice of basis is proposed to simplify the LFA, and we believe this choice will prove useful in analysing many domain decomposition algorithms in the style used here. Multiplicative preconditioners with these two domain decomposition methods are discussed briefly, and both lumped and Dirichlet variants can be improved in this way. The coarse problem involved in these domain decomposition methods can be solved by similar methods. LFA analysis of three-level variants is also considered. Degradation in convergence is well known when moving from two-level to three-level variants

of these algorithms. We show that the LFA presented above, in combination with the use of multiplicative preconditioners on the coarse and fine levels provide ways to mitigate this performance loss. Future work includes extending these variants of the preconditioned operators, using LFA to optimize the resulting algorithms, and considering other types of problems with similar preconditioners.

REFERENCES

[1] J. H. Bramble, J. E. Pasciak, J. P. Wang, and J. Xu, *Convergence estimates for product iterative methods with applications to domain decomposition*, Math. Comp., 57 (1991), pp. 1–21.

[2] A. Brandt, *Multi-level adaptive solutions to boundary-value problems*, Math. Comp., 31 (1977), pp. 333–390.

[3] S. C. Brenner, E.-H. Park, and L.-Y. Sung, *A BDDC preconditioner for a symmetric interior penalty Galerkin method*, Electron. Trans. Numer. Anal., 46 (2017), pp. 190–214.

[4] S. C. Brenner and L.-Y. Sung, *BDDC and FETI-DP without matrices or vectors*, Computer Methods in Applied Mechanics and Engineering, 196 (2007), pp. 1429–1435.

[5] W. L. Briggs, V. E. Henson, and S. F. McCormick, *A multigrid tutorial*, SIAM, 2000.

[6] C. R. Dohrmann, *A preconditioner for substructuring based on constrained energy minimization*, SIAM Journal on Scientific Computing, 25 (2003), pp. 246–258.

[7] C. R. Dohrmann, *An approximate BDDC preconditioner*, Numerical Linear Algebra with Applications, 14 (2007), pp. 149–168.

[8] C. R. Dohrmann, *Preconditioning of saddle point systems by substructuring and a penalty approach*, in Domain decomposition methods in science and engineering XVI, vol. 55 of Lect. Notes Comput. Sci. Eng., Springer, Berlin, 2007, pp. 53–64.

[9] C. R. Dohrmann and O. B. Widlund, *A BDDC algorithm with deluxe scaling for three-dimensional H(curl) problems*, Communications on Pure and Applied Mathematics, 69 (2016), pp. 745–770.

[10] V. Dolean, P. Jolivet, and F. Nataf, *An introduction to domain decomposition methods: Algorithms, theory, and parallel implementation*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2015.

[11] M. Dryja, J. Galvis, and M. Sarkis, *BDDC methods for discontinuous Galerkin discretization of elliptic problems*, J. Complexity, 23 (2007), pp. 715–739.

[12] M. Dryja and O. B. Widlund, *Towards a unified theory of domain decomposition algorithms for elliptic problems*, in Third International Symposium on Domain Decomposition Methods for Partial Differential Equations (Houston, TX, 1989), SIAM, Philadelphia, PA, 1990, pp. 3–21.

[13] M. Dryja and O. B. Widlund, *A FETI-DP method for a mortar discretization of elliptic problems*, Lecture Notes in Computational Science and Engineering, 23 (2002), pp. 41–52.

[14] C. Farhat, M. Lesoinne, P. LeTallec, K. Pierson, and D. Rixen, *FETI-DP: a dual–primal unified FETI method Part I: A faster alternative to the two-level FETI method*, International Journal for Numerical Methods in Engineering, 50 (2001), pp. 1523–1544.

[15] M. J. Gander, F. Magoulès, and F. Nataf, *Optimized Schwarz methods without overlap for the Helmholtz equation*, SIAM J. Sci. Comput., 24 (2002), pp. 38–60.

[16] J. Li, *A dual-primal FETI method for incompressible Stokes equations*, Numerische Mathematik, 102 (2005), pp. 257–275.

[17] J. Li and O. Widlund, *BDDC algorithms for incompressible Stokes equations*, SIAM Journal on Numerical Analysis, 44 (2006), pp. 2432–2455.

[18] J. Li and O. Widlund, *FETI-DP, BDDC, and block Cholesky methods*, International Journal for Numerical Methods in Engineering, 66 (2006), pp. 250–271.

[19] J. Li and O. Widlund, *A BDDC preconditioner for saddle point problems*, in Domain decomposition methods in science and engineering XVI, vol. 55 of Lect. Notes Comput. Sci. Eng., Springer, Berlin, 2007, pp. 413–420.

[20] J. Li and O. Widlund, *On the use of inexact subdomain solvers for BDDC algorithms*, Computer Methods in Applied Mechanics and Engineering, 196 (2007), pp. 1415–1428.

[21] S. P. MacLachlan and C. W. Oosterlee, *Local Fourier analysis for multigrid with overlapping smoothers applied to systems of PDEs*, Numer. Linear Algebra Appl., 18 (2011), pp. 751–774.

[22] J. Mandel and M. Brezina, *Balancing domain decomposition for problems with large jumps in coefficients*, Math. Comp., 65 (1996), pp. 1387–1401.

[23] J. Mandel and C. R. Dohrmann, *Convergence of a balancing domain decomposition by constraints and energy minimization*, Numerical Linear Algebra with Applications, 10 (2003), pp. 639–659.
[24] J. Mandel, C. R. Dohrmann, and R. Tezaur, *An algebraic theory for primal and dual substructuring methods by constraints*, Applied Numerical Mathematics, 54 (2005), pp. 167–193.
[25] J. Mandel, B. r. Sousedík, and C. R. Dohrmann, *Multispace and multilevel BDDC*, Computing, 83 (2008), pp. 55–85.
[26] K. Stüben and U. Trottenberg, *Multigrid methods: Fundamental algorithms, model problem analysis and applications*, Multigrid Methods, (1982), pp. 1–176.
[27] A. Toselli and O. Widlund, *Domain decomposition methods: algorithms and theory*, vol. 34 of Springer Series in Computational Mathematics, Springer-Verlag, Berlin, 2005.
[28] U. Trottenberg, C. W. Oosterlee, and A. Schüller, *Multigrid*, Academic Press, Inc., San Diego, CA, 2001. With contributions by A. Brandt, P. Oswald and K. Stüben.
[29] X. Tu, *Domain decomposition algorithms: methods with three levels and for flow in porous media*, PhD thesis, New York University, Graduate School of Arts and Science, 2006.
[30] X. Tu, *Three-level BDDC in three dimensions*, SIAM J. Sci. Comput., 29 (2007), pp. 1759–1780.
[31] X. Tu, *Three-level BDDC in two dimensions*, Internat. J. Numer. Methods Engrg., 69 (2007), pp. 33–59.
[32] P. Wesseling, *An introduction to multigrid methods*, Pure and Applied Mathematics (New York), John Wiley & Sons, Ltd., Chichester, 1992.
[33] R. Wienands and W. Joppich, *Practical Fourier analysis for multigrid methods*, CRC press, 2004.