Journal of Computational Acoustics,

FETI-DPH: A DUAL-PRIMAL DOMAIN DECOMPOSITION METHOD FOR ACOUSTIC SCATTERING

Charbel Farhat, Philip Avery and Radek Tezaur
Department of Mechanical Engineering
and Institute for Computational and Mathematical Engineering
Stanford University, Mail Code 3035, Stanford, CA 94305, U.S.A.

Jing Li
Department of Mathematical Sciences
Kent State University, Kent, OH 44242, U. S. A.

Received (to be inserted
Revised by Publisher)

A dual-primal variant of the FETI-H domain decomposition method is designed for the fast, parallel, iterative solution of large-scale systems of complex equations arising from the discretization of acoustic scattering problems formulated in bounded computational domains. The convergence of this iterative solution method, named here FETI-DPH, is shown to scale with the problem size, the number of subdomains, and the wave number. Its solution time is also shown to scale with the problem size. CPU performance results obtained for the acoustic signature analysis in the mid-frequency regime of mockup submarines reveal that the proposed FETI-DPH solver is significantly faster than the previous generation FETI-H solution algorithm.

1. Introduction

Let Θ and Υ denote respectively a three-dimensional impenetrable obstacle, and a three-dimensional bounded domain containing this obstacle (Fig. 1). The scattering of a wave by the obstacle Θ can be modeled by the following boundary value problem (BVP)

Find $u \in H^1(\Omega)$ *such that*

$$\begin{array}{rcll} -\Delta u - k^2 u &=& 0 & \text{in} \quad \Omega \\ \alpha u + \beta \dfrac{\partial u}{\partial \nu} &=& -(\alpha + \beta \dfrac{\partial}{\partial \nu}) e^{ik\mathbf{t}\cdot\mathbf{x}} & \text{on} \quad \partial\Theta \\ \dfrac{\partial u}{\partial \nu} &=& A(u) + g & \text{on} \quad \Sigma \end{array} \tag{1.1}$$

where $k > 0$ is the wave number, $\Omega = \Upsilon \backslash \overline{\Theta}$, $\Sigma = \partial\Upsilon$ is an artificial boundary, $H^1(\Omega)$ is the usual Sobolev space, $\boldsymbol{\nu}$ is the outward normal, $\dfrac{\partial}{\partial \nu}$ denotes the normal derivative, α and β are two constants with either $(\alpha, \beta) = (1, 0)$ or $(\alpha, \beta) = (0, 1)$, $\mathbf{t} \in \mathbb{R}^3$ is the direction of

the incident wave, $\mathbf{x} \in \mathbb{R}^3$, A is a differential operator, and g is a source function which may depend on k. The second of Eqs. (1.1) is a concise representation of Dirichlet and Neumann boundary conditions. The third of Eqs. (1.1) encompasses Neumann, Robin, and many absorbing boundary conditions.

The boundary value problem (1.1) is known as the formulation of the *exterior* Helmholtz problem in a bounded domain. If Θ is replaced by void — in which case $\Upsilon \equiv \Omega$ and $\Sigma \equiv \partial\Omega$ — and the second of Eqs. (1.1) is eliminated (or $\alpha = \beta = 0$), the resulting instance of the BVP (1.1) defines an *interior* Helmholtz problem.

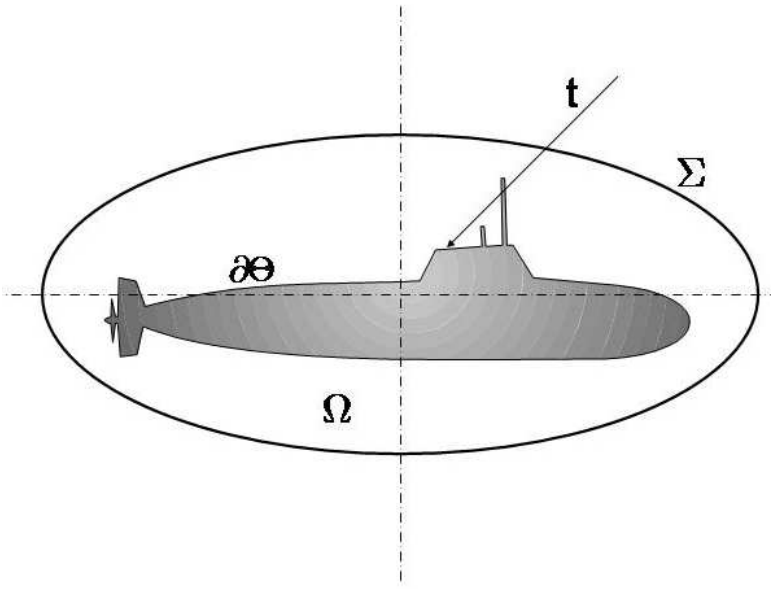

Fig. 1. Formulation of an acoustic scattering problem in a bounded domain

The discretization of the BVP (1.1) by finite elements (FE) leads to a linear algebraic problem of the form

$$(\mathbf{K} - k^2\mathbf{M} + i\mathbf{A})\mathbf{u} = \mathbf{f}, \tag{1.2}$$

where $\mathbf{K}$ and $\mathbf{M}$ denote the usual FE stiffness and mass symmetric matrices, respectively, i is the imaginary complex number satisfying $i^2 = -1$, and $\mathbf{A}$ is a FE symmetric matrix arising from the discretization of the absorbing boundary condition represented by the differential operator A. Hence, $\mathbf{A} = 0$ for interior Helmholtz problems. Otherwise, $\mathbf{A}$ is a sparse matrix when A is a local differential operator.

Let

$$\mathbf{Z} = \mathbf{K} - k^2\mathbf{M} + i\mathbf{A} \tag{1.3}$$

denote the impedance matrix arising from the FE discretization of problem (1.1). Using this notation, problem (1.2) can be re-written as

$$\mathbf{Zu} = \mathbf{f}, \tag{1.4}$$

where $\mathbf{Z}$ is a sparse, symmetric (but not Hermitian) matrix.

In the mid-frequency regime — for example, for $ka \geq 5$, where a denotes a characteristic length of the obstacle Θ — the discretization of the BVP (1.1) by the standard FE method calls for fine meshes and therefore leads to large-scale systems of equations of the form given in (1.4). In this case, solving Eq. (1.4) by a direct method entails memory and CPU requirements that often overwhelm the available computational resources. For this reason, a significant amount of research effort continues to be invested in the development of iterative algorithms for the solution of such problems.

Domain decomposition-based preconditioned Krylov methods have emerged as powerful equation solvers on both sequential and parallel computing platforms [1]. While most successful domain decomposition methods (DDMs) have been designed for the solution of symmetric positive (semi-) definite systems, some of them have targeted indefinite ($\mathbf{A} = 0$) or complex ($\mathbf{A} \neq 0$) problems of the form given in Eq. (1.2) [2,3,4]. Among these, the FETI-H method [4] stands out as a DDM that has been routinely applied to the solution of large-scale, three-dimensional, realistic acoustic scattering problems [5,6,7]. The FETI-H solver is based on a DDM with Lagrange multipliers known as the dual FETI method [8,9,10,11,12]. FETI, which is provably numerically scalable for second- [12] and fourth-order [13] elliptic problems — that is, its iteration count increases only weakly with the size of the subdomain and global problems — was recently transformed into a leaner, faster, and yet equally numerically scalable dual-primal DDM known as FETI-DP [14,15]. Hence, the objective of this paper is to develop an an alternative to FETI-H that is based on FETI-DP rather than on FETI, and assess its performance for the solution of interior and exterior Helmholtz problems. For this purpose, the reminder of this paper is organized as follows.

In Section 2, the FETI-DP method is overviewed and its scalability properties are recalled. In Section 3, this DDM is extended to the solution of interior and exterior Helmholtz problems that give rise to algebraic systems of the form given in Eq. (1.2); this leads to the design of the FETI-DPH method ("H" for Helmholtz problems). In Section 4, the parallel implementation of FETI-DPH adopted in this work is briefly overviewed. In Section 5, the scalability properties of FETI-DPH are assessed numerically, first for three-dimensional academic Helmholtz problems, then for large-scale acoustic scattering problems associated with the signature analysis on an Origin 3200 parallel processor of a mockup submarine. Finally, conclusions are offered in Section 6.

2. The FETI-DP Method

The dual-primal finite element tearing and interconnecting (FETI-DP) method [14,15] is a third-generation FETI method (for example, see [8,9,10]) developed for the fast,

iterative, and parallel solution of systems of equations arising from the FE discretization of static, dynamic, second-order, and fourth-order elliptic partial differential equations (PDEs). When equipped with the Dirichlet preconditioner [11] and applied to plane stress/strain or shell problems, the condition number κ of its interface problem grows asymptotically as [12]

$$\kappa = \mathcal{O}\ (1 + \log^m \frac{H}{h}\ ), \quad m \leq 2, \tag{2.5}$$

where H and h denote the subdomain and mesh sizes, respectively. When equipped with the same Dirichlet preconditioner and an auxiliary coarse problem constructed by enforcing some set of optional constraints at the subdomain interfaces [15], the condition number estimate (2.5) also holds for second-order scalar elliptic problems that model three-dimensional solid mechanics problems [16]. This estimate proves the scalability of the FETI-DP method with respect to all of the problem size, subdomain size, and number of subdomains. More specifically, it suggests that one can expect FETI-DP to solve small-scale and large-scale problems in similar iteration counts (this property is usually referred to as *numerical* scalability). This in turn suggests that if the FETI-DP method is well-implemented on a parallel processor so that for a fixed problem size it can deliver a speed-up that increases almost linearly with the number of processors (this property is usually referred to as *parallel* scalability) , it should be capable of solving an m-times larger problem using an m-times larger number of processors in almost a constant CPU time (this property is usually referred to simply as *scalability*). The scalability of the FETI-DP method was demonstrated in practice for many complex structural mechanics and structural dynamics problems (for example, see [14,15,17] and the references cited therein).

In order to keep this paper as self-contained as possible, the FETI-DP method is next overviewed in the context of the generic symmetric positive semi-definite static problem

$$\mathbf{K}\mathbf{u} = \mathbf{f}, \tag{2.6}$$

where $\mathbf{K}$ has the same meaning as in problem (1.2) and $\mathbf{f}$ is an arbitrary load vector.

2.1. *Domain decomposition and notation*

Let Ω denote the computational support of a solid mechanics or shell problem whose discretization leads to problem (2.6), $\{\Omega^{(s)}\}_{s=1}^{N_s}$ denote its decomposition into N_s subdomains with matching interfaces $\Gamma^{(s,q)} = \partial\Omega^{(s)} \bigcap \partial\Omega^{(q)}$, and let $\Gamma = \bigcup_{s=1,q>s}^{s=N_s} \Gamma^{(s,q)}$ denote the global interface of this decomposition. *In the remainder of this paper, each interface* $\Gamma^{(s,q)}$ *is referred to as an "edge", whether* Ω *is a two- or three-dimensional domain.* Let also $\mathbf{K}^{(s)}$ and $\mathbf{f}^{(s)}$ denote the contributions of subdomain $\Omega^{(s)}$ to $\mathbf{K}$ and $\mathbf{f}$, respectively, and let $\mathbf{u}^{(s)}$ denote the vector of dof associated with it.

Let N_c of the N_I nodes lying on the global interface Γ be labeled "corner" nodes (see Fig. 2), Γ_c denote the set of these corner nodes, and let $\Gamma' = \Gamma \backslash \Gamma_c$. The selection of the

corner nodes is briefly discussed in Section 2.3. If in each subdomain $\Omega^{(s)}$ the unknowns are partitioned into global corner dof designated by the subscript c, and "remaining" dof designated by the subscript r, $\mathbf{K}^{(s)}$, $\mathbf{u}^{(s)}$ and $\mathbf{f}^{(s)}$ can be partitioned as follows

$$\mathbf{K}^{(s)} = \begin{bmatrix} \mathbf{K}_{rr}^{(s)} & \mathbf{K}_{rc}^{(s)} \\ \mathbf{K}_{rc}^{(s)^T} & \mathbf{K}_{cc}^{(s)} \end{bmatrix}, \qquad \mathbf{u}^{(s)} = \begin{bmatrix} \mathbf{u}_r^{(s)} \\ \mathbf{u}_c^{(s)} \end{bmatrix} \qquad \text{and} \qquad \mathbf{f}^{(s)} = \begin{bmatrix} \mathbf{f}_r^{(s)} \\ \mathbf{f}_c^{(s)} \end{bmatrix}. \tag{2.7}$$

The r-type dof can be further partitioned into "interior" dof designated by the subscript i, and subdomain interface "boundary" dof designated by the subscript b. Hence, $\mathbf{u}_r^{(s)}$ and $\mathbf{f}_r^{(s)}$ can be further partitioned as follows

$$\mathbf{u}_r^{(s)} = \begin{bmatrix} \mathbf{u}_i^{(s)} & \mathbf{u}_b^{(s)} \end{bmatrix}^T \qquad \text{and} \qquad \mathbf{f}_r^{(s)} = \begin{bmatrix} \mathbf{f}_i^{(s)} & \mathbf{f}_b^{(s)} \end{bmatrix}^T, \tag{2.8}$$

where the superscript T designates the transpose.

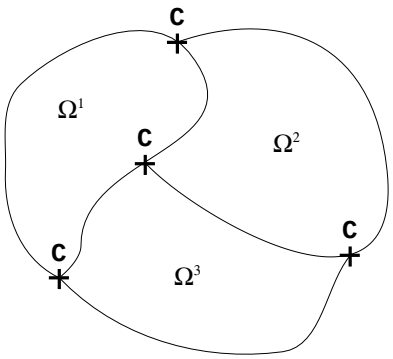

Fig. 2. Sample definition of corner points

Let $\mathbf{u}_c$ denote the global vector of corner dof, and $\mathbf{u}_c^{(s)}$ denote its restriction to $\Omega^{(s)}$. Let also $\mathbf{B}_r^{(s)}$ and $\mathbf{B}_c^{(s)}$ be the two subdomain Boolean matrices defined by

$$\mathbf{B}_r^{(s)}\mathbf{u}_r^{(s)} = \pm\mathbf{u}_b^{(s)} \qquad \text{and} \qquad \mathbf{B}_c^{(s)}\mathbf{u}_c = \mathbf{u}_c^{(s)}, \tag{2.9}$$

where the $\pm$ sign is set by any convention that implies that $\sum_{s=1}^{N_s} \mathbf{B}_r^{(s)}\mathbf{u}_r^{(s)}$ represents the *jump* of the displacement solution $\mathbf{u}$ across the subdomain interfaces. Finally, let

$$\mathbf{f}_c = \sum_{s=1}^{N_s} \mathbf{B}_c^{(s)^T}\mathbf{f}_c^{(s)}. \tag{2.10}$$

In [14] and [15], it was shown that solving problem (2.6) is equivalent to solving the following

subdomain-based problem

$$\mathbf{K}_{rr}^{(s)}\ \mathbf{u}_r^{(s)} + \mathbf{K}_{rc}^{(s)}\mathbf{B}_c^{(s)}\ \mathbf{u}_c + \mathbf{B}_r^{(s)^T}\lambda + \mathbf{B}_r^{(s)^T}\mathbf{Q}_b\ \mu = \mathbf{f}_r^{(s)}, \quad s = 1, ..., N_s \tag{2.11}$$

$$\sum_{s=1}^{N_s} \mathbf{B}_c^{(s)^T}\mathbf{K}_{rc}^{(s)^T}\ \mathbf{u}_r^{(s)} + \sum_{s=1}^{N_s} \mathbf{B}_c^{(s)^T}\mathbf{K}_{cc}^{(s)}\mathbf{B}_c^{(s)}\ \mathbf{u}_c = \mathbf{f}_c, \tag{2.12}$$

$$\sum_{s=1}^{N_s} \mathbf{B}_r^{(s)}\mathbf{u}_r^{(s)} = 0, \tag{2.13}$$

$$\mathbf{Q}_b^T \sum_{s=1}^{N_s} \mathbf{B}_r^{(s)}\mathbf{u}_r^{(s)} = 0, \tag{2.14}$$

where λ is an N_λ-long vector of Lagrange multipliers introduced on Γ' to enforce the continuity (2.13) of the displacement vector $\mathbf{u}$, and μ is another vector of Lagrange multipliers introduced to enforce the optional linear constraints (2.14). These optional constraints, a concept first developed in [18], generate a matrix $\mathbf{Q}_b$ with $N_Q < N_\lambda$ columns defined on Γ'. The word "optional" refers to the fact that Eq. (2.14) and the vector of Lagrange multipliers μ are not necessarily needed for formulating the above subdomain-based problem. Indeed, since the solution of problem (2.6) is continuous across the subdomain interfaces, it satisfies Eq. (2.13) and therefore satisfies Eq. (2.14) for any matrix $\mathbf{Q}_b$. Note also that each of Eqs. (2.11) is a local subdomain equation, whereas Eq. (2.12) is an assembled global equation.

The subdomain-based problem (2.11–2.14) was labeled "dual-primal" in [14] and [15] because it is formulated in terms of two different types of global unknowns: the dual Lagrange multipliers represented by the vector λ, and the primal corner dof represented by the vector $\mathbf{u}_c$.

In the remainder of this paper, the j-th column of $\mathbf{Q}_b$ is denoted by $\mathbf{q}_j$ so that

$$\mathbf{Q}_b = [\mathbf{q}_1 \quad \cdots \quad \mathbf{q}_j \quad \cdots \quad \mathbf{q}_{N_Q}]. \tag{2.15}$$

2.2. *Interface and coarse problems*

Let

$$\widetilde{\mathbf{K}}_{cc} = \begin{bmatrix} \mathbf{K}_{cc} & 0 \\ 0 & 0 \end{bmatrix}, \qquad \mathbf{K}_{cc} = \sum_{s=1}^{N_s} \mathbf{B}_c^{(s)^T}\mathbf{K}_{cc}^{(s)}\mathbf{B}_c^{(s)}, \qquad \mathbf{d}_r = \sum_{s=1}^{N_s} \mathbf{B}_r^{(s)}\mathbf{K}_{rr}^{(s)^{-1}}\mathbf{f}_r^{(s)},$$

$$\text{and} \qquad \mathbf{f}_c^* = \mathbf{f}_c - \sum_{s=1}^{N_s} (\mathbf{K}_{rc}^{(s)}\mathbf{B}_c^{(s)})^T\mathbf{K}_{rr}^{(s)^{-1}}\mathbf{f}_r^{(s)}. \tag{2.16}$$

After some algebraic manipulations aimed at eliminating $\mathbf{u}_r^{(s)}$, $s = 1, ..., N_s$, $\mathbf{u}_c$, and μ, the subdomain-based problem (2.11–2.14) can be transformed into the following symmetric

positive semi-definite interface problem

$$(\mathbf{F}_{I_{rr}} + \widetilde{\mathbf{F}}_{I_{rc}} \widetilde{\mathbf{K}}_{cc}^{*^{-1}} \widetilde{\mathbf{F}}_{I_{rc}}^T)\lambda = \mathbf{d}_r - \widetilde{\mathbf{F}}_{I_{rc}} \widetilde{\mathbf{K}}_{cc}^{*^{-1}} \tilde{\mathbf{f}}_c^*, \tag{2.17}$$

where

$$\begin{aligned}
\mathbf{F}_{I_{rr}} &= \sum_{s=1}^{N_s} \mathbf{B}_r^{(s)} \mathbf{K}_{rr}^{(s)^{-1}} \mathbf{B}_r^{(s)^T}, \qquad \widetilde{\mathbf{F}}_{I_{rc}} = \sum_{s=1}^{N_s} \mathbf{B}_r^{(s)} \mathbf{K}_{rr}^{(s)^{-1}} \widetilde{\mathbf{K}}_{rc}^{(s)}, \\
\widetilde{\mathbf{K}}_{rc}^{(s)} &= \left[\mathbf{K}_{rc}^{(s)} \mathbf{B}_c^{(s)} \quad \mathbf{B}_r^{(s)^T} \mathbf{Q}_b \right], \qquad \tilde{\mathbf{f}}_c^* = \begin{bmatrix} \mathbf{f}_c^* \\ -\mathbf{Q}_b^T \mathbf{d}_r \end{bmatrix}, \\
\widetilde{\mathbf{K}}_{cc}^* &= \widetilde{\mathbf{K}}_{cc} - \begin{bmatrix} \sum_{s=1}^{N_s} (\mathbf{K}_{rc}^{(s)} \mathbf{B}_c^{(s)})^T \mathbf{K}_{rr}^{(s)^{-1}} (\mathbf{K}_{rc}^{(s)} \mathbf{B}_c^{(s)}) & \sum_{s=1}^{N_s} (\mathbf{K}_{rc}^{(s)} \mathbf{B}_c^{(s)})^T \mathbf{K}_{rr}^{(s)^{-1}} (\mathbf{B}_r^{(s)^T} \mathbf{Q}_b) \\ \sum_{s=1}^{N_s} (\mathbf{B}_r^{(s)^T} \mathbf{Q}_b)^T \mathbf{K}_{rr}^{(s)^{-1}} (\mathbf{K}_{rc}^{(s)} \mathbf{B}_c^{(s)}) & \sum_{s=1}^{N_s} (\mathbf{B}_r^{(s)^T} \mathbf{Q}_b)^T \mathbf{K}_{rr}^{(s)^{-1}} (\mathbf{B}_r^{(s)^T} \mathbf{Q}_b) \end{bmatrix}.
\end{aligned} \tag{2.18}$$

The FETI-DP method is a DDM that solves the original problem (2.6) by applying a preconditioned conjugate gradient (PCG) algorithm to the solution of the corresponding dual interface problem (2.17). At the n-th PCG iteration, the matrix-vector product $(\mathbf{F}_{I_{rr}} + \widetilde{\mathbf{F}}_{I_{rc}} \widetilde{\mathbf{K}}_{cc}^{*^{-1}} \widetilde{\mathbf{F}}_{I_{rc}}^T)\lambda^n$ incurs the solution of an auxiliary problem of the form

$$\widetilde{\mathbf{K}}_{cc}^* \mathbf{z} = \widetilde{\mathbf{F}}_{I_{rc}}^T \lambda^n. \tag{2.19}$$

From the fifth of Eqs. (2.18), it follows that the size of this auxiliary problem is equal to the sum of the number of corner dof, N_c^{dof}, and the number of columns of the matrix $\mathbf{Q}_b$, N_Q.

For $N_Q = 0$ — that is, for $\mathbf{Q}_b = 0$, the auxiliary problem (2.19) is a coarse problem, and $\widetilde{\mathbf{K}}_{cc}^*$ is a sparse matrix whose pattern is that of the stiffness matrix obtained when each subdomain is treated as a "superelement" whose nodes are its corner nodes. This coarse problem ensures that the FETI-DP method equipped with the Dirichlet preconditioner (see Section 2.4) is numerically scalable for plate and shell problems as well as two-dimensional plane stress/strain problems [15,12]. However, for $\mathbf{Q}_b = 0$, the FETI-DP method equipped with the Dirichlet preconditioner is not numerically scalable for three-dimensional solid problems.

For any choice of $\mathbf{Q}_b \neq 0$, $\widetilde{\mathbf{K}}_{cc}^*$ remains a sparse matrix. If $\mathbf{Q}_b$ is constructed edge-wise — that is, if each column of $\mathbf{Q}_b$ is constructed as the restriction of some operator to a specific edge of Γ' — the sparsity pattern of $\widetilde{\mathbf{K}}_{cc}^*$ becomes that of a stiffness matrix obtained by treating each subdomain as a superelement whose nodes are its corner nodes augmented by virtual mid-side nodes. The number of dof attached to each virtual mid-side node is equal to the number of columns of $\mathbf{Q}_b$ associated with the edge on which lies this mid-side node. If N_Q is kept relatively small, the auxiliary problem (2.19) remains a relatively small coarse problem. This coarse problem was labeled "augmented" coarse problem in [14] in order to

distinguish it from the smaller coarse problem obtained with $\mathbf{Q}_b = 0$. Furthermore, each column of $\mathbf{Q}_b$ is referred to as an "augmentation coarse mode". When these augmentation coarse modes are chosen as the translational rigid body modes of each edge of Γ', each three consecutive columns of $\mathbf{Q}_b$ can be written as follows

$$\begin{aligned} \mathbf{q}_x &= [0 \cdots 0\ [1\ 0\ 0\ \cdots\ 1\ 0\ 0\ \cdots\ 1\ 0\ 0]\ 0 \cdots 0\ ]^T, \\ \mathbf{q}_y &= [0 \cdots 0\ [0\ 1\ 0\ \cdots\ 0\ 1\ 0\ \cdots\ 0\ 1\ 0]\ 0 \cdots 0\ ]^T, \\ \mathbf{q}_z &= [0 \cdots 0\ [0\ 0\ 1\ \cdots\ 0\ 0\ 1\ \cdots\ 0\ 0\ 1]\ 0 \cdots 0\ ]^T, \end{aligned} \tag{2.20}$$

and the FETI-DP method equipped with the Dirichlet preconditioner becomes numerically scalable for three-dimensional second-order elasticity problems [16]. In (2.20), each of the vectors $\mathbf{q}_x$, $\mathbf{q}_y$, and $\mathbf{q}_z$ has non-zero entries only between the pair of brackets [ ] delimiting the dof attached to the edge associated with these three vectors. Within these brackets, the sequence 1 0 0 corresponds to a translational rigid body mode of this edge in the x direction, the sequence 0 1 0 corresponds to a translational rigid body mode in the y direction, and the sequence 0 0 1 corresponds to a translational rigid body mode in the z direction.

2.3. *Corner selection*

From the definitions of $\mathbf{F}_{I_{rr}}$, $\widetilde{\mathbf{K}}^*_{cc}$, and $\mathbf{d}_r$ given in (2.18) and (2.16), it follows that the corner nodes must be chosen such that $\mathbf{K}^{(s)}_{rr}$ is non-singular. From Eq. (2.17), it follows that when $\mathbf{Q}_b \neq 0$, the selection of the corner nodes must furthermore guarantee that $\widetilde{\mathbf{K}}^*_{cc}$ is non-singular. From the theory and results exposed in [19] and [20], it follows that the corner nodes must also include crosspoints — that is, points that belong to three or more subdomains — when these are attached to a beam, plate, or shell element. A corner selection algorithm that meets all of these requirements was proposed and discussed in [21].

2.4. *Local preconditioning*

So far, two conventional local preconditioners have been developed for the FETI-DP method:

1. The Dirichlet preconditioner which can be written as

$$\begin{aligned} \overline{\mathbf{F}}^{D^{-1}}_{I_{rr}} &= \sum_{s=1}^{N_s} \mathbf{W}^{(s)} \mathbf{B}^{(s)}_r \begin{bmatrix} 0 & 0 \\ 0 & \mathbf{S}^{(s)}_{bb} \end{bmatrix} \mathbf{B}^{(s)^T}_r \mathbf{W}^{(s)}, \\ \text{where} \qquad \mathbf{S}^{(s)}_{bb} &= \mathbf{K}^{(s)}_{bb} - \mathbf{K}^{(s)^T}_{ib} \mathbf{K}^{(s)^{-1}}_{ii} \mathbf{K}^{(s)}_{ib}, \end{aligned} \tag{2.21}$$

   the subscripts i and b have the same meaning as in Section 2.1, and $\mathbf{W}^{(s)}$ is a subdomain diagonal scaling matrix that accounts for possible subdomain heterogeneities [22]. The roots of this preconditioner and its mechanical interpretation can be found in [11]. It is mathematically optimal in the sense that it leads to the condition number estimate (2.5).

2. The lumped preconditioner which can be written as

$$\overline{\mathbf{F}}_{I_{rr}}^{L^{-1}} = \sum_{s=1}^{N_s} \mathbf{W}^{(s)} \mathbf{B}_r^{(s)} \begin{bmatrix} 0 & 0 \\ 0 & \mathbf{K}_{bb}^{(s)} \end{bmatrix} \mathbf{B}_r^{(s)^T} \mathbf{W}^{(s)}. \tag{2.22}$$

   The roots of this preconditioner and its mechanical interpretation can be found in [9]. This preconditioner is not mathematically optimal in the sense defined above; however, it decreases the cost of each iteration in comparison with the Dirichlet preconditioner, often with a modest increase in the iteration count.

The Dirichlet preconditioner is less economical per iteration than the lumped preconditioner. However, experience has shown that it is more computationally efficient when $\mathbf{K}$ in problem (2.6) arises from the discretization of shell problems. When $\mathbf{K}$ arises from plane stress/strain and solid mechanics problems, the lumped preconditioner becomes more computationally efficient even though it is not mathematically optimal.

3. The FETI-DPH Method for Interior and Exterior Helmholtz Problems

In the context of Eq. (1.4), $\mathbf{Z}_{rr}^{(s)}$ replaces $\mathbf{K}_{rr}^{(s)}$ in the formulation of the FETI-DP method. For interior Helmholtz problems or exterior ones but away from the artificial boundary Σ,

$$\mathbf{Z}_{rr}^{(s)} = \mathbf{K}_{rr}^{(s)} - k^2 \mathbf{M}_{rr}^{(s)}, \tag{3.23}$$

and therefore the direct application of FETI-DP to the solution of problem (1.4) raises the following questions:

1. Away from or in the absence of the artificial boundary Σ, $\mathbf{A}_{rr}^{(s)} = 0$, which begs to ask whether the subdomain matrix $\mathbf{Z}_{rr}^{(s)}$ can become in this case singular (or nearly singular)?

2. For the exterior Helmholtz problem ($\mathbf{A} \neq \mathbf{0}$), the dual interface problem (2.17) is complex-valued. For the interior Helmholtz problem ($\mathbf{A} = \mathbf{0}$), it is real-valued. In the latter case, what is the algebraic type (positive definite or indefinite) of this interface problem?

3. The restriction of the augmentation coarse modes (2.20) to one dof per node does not address neither the physical nor the mathematical nature of problem (1.2). Which augmentation coarse modes do?

The significance of the above three questions can be summarized as follows.

The answer to the first question impacts the solution of the local subdomain problems (2.11). In practice (finite arithmetic), given a wave number k, it is difficult to determine whether a subdomain impedance matrix $\mathbf{Z}_{rr}^{(s)}$ is singular or not. However, $\mathbf{Z}_{rr}^{(s)}$ is typically very ill-conditioned when k^2 coincides with or is very close to an eigenvalue of the pencil

$(\mathbf{K}_{rr}^{(s)}, \mathbf{M}_{rr}^{(s)})$. This can be alarming because round-off errors are known to give rise to erroneous solutions to very ill-conditioned problems. Furthermore, it should be noted that FETI-DP as well as many other DDMs rely on a direct method for solving the local subdomain problems. In the low-frequency regime, $\mathbf{Z}_{rr}^{(s)}$ can be expected to be positive definite. However, in the mid-frequency regime — that is, for higher values of the wave number k — $\mathbf{Z}_{rr}^{(s)}$ can become indefinite in which case pivoting may become required. Since pivoting degrades the performance of a direct skyline or sparse solver, it follows that it may also degrade the overall CPU performance of FETI-DP.

The objective of the second question is to determine the Krylov method that is most suitable for solving the dual interface problem (2.17) for interior Helmholtz problems.

The third question aims at recognizing that for acoustic scattering problems, there are perhaps other important modes besides the subdomain constant modes that can accelerate the iterative solution of the dual interface problem (2.17).

Next, the three questions raised above are addressed in some details, and the Dirichlet and lumped preconditioners summarized in Section 2.4 are tailored to Helmholtz problems. This leads to the development of yet another FETI method for the parallel iterative solution of problem (1.4). This method is named here the FETI-DPH method.

3.1. *On the algebraic type of the real-valued subdomain impedance matrices*

Consider the case where $\mathbf{A}_{rr}^{(s)} = 0$. The corner nodes selected by the FETI-DP solver are such that each subdomain matrix $\mathbf{K}_{rr}^{(s)}$ is non-singular. However, at least in theory, $\mathbf{Z}_{rr}^{(s)} = \mathbf{K}_{rr}^{(s)} - k^2\mathbf{M}_{rr}^{(s)}$ will still be singular when k^2 coincides with an eigenvalue of the pencil $(\mathbf{K}_{rr}^{(s)}, \mathbf{M}_{rr}^{(s)})$. This issue was addressed in [4] by introducing in the FETI method a regularization procedure that prevents all subdomain problems from being singular for any wave number k, without destroying the sparsity of the local matrices $\mathbf{K}_{rr}^{(s)} - k^2\mathbf{M}_{rr}^{(s)}$, and without affecting the solution of the global problem (1.4). The resulting DDM was named the FETI-H method (H for Helmholtz). Unfortunately, the regularization procedure characterizing FETI-H transforms each real-valued subdomain problem into a complex-valued one. For exterior Helmholtz problems, this is not a significant issue because the global problem is complex-valued. However, for interior Helmholtz problems for which Eq. (1.4) is real-valued ($\mathbf{A} = 0$), the regularization procedure of the FETI-H method is unjustifiable from both computational resource and performance viewpoints. For this reason, an alternative approach is adopted here for addressing both the algebraic type and the potential singularity of a local subdomain impedance matrix $\mathbf{Z}_{rr}^{(s)}$. This alternative approach exploits the fact that whereas $\mathbf{K}^{(s)}$ can be singular in the FETI method, $\mathbf{K}_{rr}^{(s)}$ is always non-singular in the FETI-DP method.

In practice, if k^2 approaches an eigenvalue of the pencil $(\mathbf{K}_{rr}^{(s)}, \mathbf{M}_{rr}^{(s)})$ or coincides with it, $\mathbf{Z}_{rr}^{(s)}$ becomes only very ill-conditioned because of the effect of finite arithmetic. One approach for preventing this from happening is to somehow make sure that, given a subdomain $\Omega^{(s)}$, any value of the wave number k in a mid-frequency regime of interest is such that $k^2 < k_1^2$, where k_1^2 denotes the smallest eigenvalue of the pencil $(\mathbf{K}_{rr}^{(s)}, \mathbf{M}_{rr}^{(s)})$. This

approach is enabled by the fact that $k_1^2 > 0$ in FETI-DP (whereas $k_1^2 \geq 0$ in FETI), and motivated by the following example.

Consider an interior Helmholtz problem with $\Omega = [0,\ a] \times [0,\ a]$ and homogeneous Dirichlet boundary conditions at the four points defining this square domain Ω (Fig. 3). In this case, the eigenvalues of the pencil $(\mathbf{K}, \mathbf{M})$ are given by

$$k_n^2 = \frac{2n^2\pi^2}{a^2}, \quad n = 1,\ ...,\ \infty. \tag{3.24}$$

Next, suppose that Ω is decomposed into $\sqrt{N_s} \times \sqrt{N_s}$ subdomains (Fig. 3). In this case,

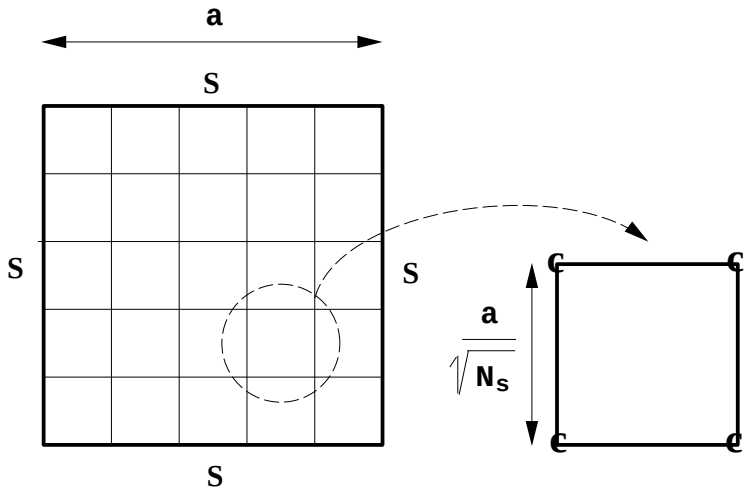

Fig. 3. Interior Helmholtz problem with homogeneous Dirichlet boundary conditions at the four points defining the square domain — Uniform decomposition into N_s subdomains.

each subdomain is a square of size

$$a^{(s)} = \frac{a}{\sqrt{N_s}}. \tag{3.25}$$

In the FETI-DP method, each subdomain away from the boundary conditions is treated as if it has Dirichlet boundary conditions at its four corner nodes (see Fig. 3). Hence, from Eq. (3.24) and Eq. (3.25) it follows that the first eigenvalue of the pencil $(\mathbf{K}_{rr}^{(s)}, \mathbf{M}_{rr}^{(s)})$ is

$$k_1^2 = \frac{2N_s\pi^2}{a^2}. \tag{3.26}$$

It follows that the impedance matrix $\mathbf{Z}_{rr}^{(s)}$ of the typical subdomain shown in Fig. 3 is positive definite for all values of k^2 satisfying

$$(ka)^2 < 2N_s\pi^2. \tag{3.27}$$

The condition (3.27) suggests that given a wave number k, the subdomain matrices $\mathbf{Z}_{rr}^{(s)}$ associated with an interior Helmholtz problem or subdomains that are away from the artificial boundary Σ can be prevented from being singular or very ill-conditioned by increasing the number of subdomains N_s — or equivalently, decreasing the subdomain sizes. For example, for the interior Helmholtz problem considered here and a mid-frequency regime characterized by $ka = 20$, Eq. (3.27) reveals that it suffices to decompose Ω into 20 subdomains to prevent all subdomain impedance matrices from being singular. This illustrates the practical feasibility of criterion (3.27).

In summary, for wave numbers in the low and mid-frequency regimes, the global impedance matrix $\mathbf{Z}$ may be indefinite, but the local subdomains impedances $\mathbf{Z}_{rr}^{(s)}$ can be maintained positive definite by choosing appropriately the number of subdomains.

3.2. *On the algebraic type of the dual interface problem*

Recall that in the context of Eq. (1.4), $\widetilde{\mathbf{K}}_{cc}^{*}$ (2.18) has the meaning of the impedance matrix obtained when each subdomain is treated as a superelement whose nodes are the subdomain corner nodes, and virtual mid-side nodes when $\mathbf{Q}_b \neq 0$. It follows that when the global impedance matrix $\mathbf{Z}$ (1.3) is indefinite, $\widetilde{\mathbf{K}}_{cc}^{*}$ is also indefinite. Hence, for interior Helmholtz problems, even when the subdomain impedance matrices $\mathbf{Z}_{rr}^{(s)}$ can be maintained positive definite by choosing appropriately the number of subdomains N_s, the matrix $(\mathbf{F}_{I_{rr}} + \widetilde{\mathbf{F}}_{I_{rc}} \widetilde{\mathbf{K}}_{cc}^{*^{-1}} \widetilde{\mathbf{F}}_{I_{rc}}^{T})$ (2.17) will be in general indefinite. For exterior Helmholtz problems, the matrices $\mathbf{F}_{I_{rr}}$, $\widetilde{\mathbf{F}}_{I_{rc}}$, and $\widetilde{\mathbf{Z}}_{cc}^{*}$ are complex-valued, and therefore the dual interface problem (2.17) is complex-valued. These observations suggest that for robustness and code unification, a preconditioned generalized minimum residual (PGMRES) algorithm is more appropriate than a PCG algorithm for solving iteratively the dual interface problem (2.17).

3.3. *Wave-based augmentation coarse modes*

In [4], it was advocated to equip the regularized FETI method with a coarse problem based on plane waves when addressing Helmholtz applications. By construction, the FETI-DP method embeds a coarse problem that does not however recognize any specific character of the Helmholtz equation. For this reason, the plane waves introduced in [4] for equipping the regularized FETI method with an auxiliary coarse problem are used again here to augment the coarse problem characterizing the FETI-DP method. Also, this section explains why plane waves can in general accelerate the convergence of a FETI method applied to the solution of a Helmholtz problem.

Let $\mathbf{r}$ denote the residual associated with the iterative solution of the dual interface problem (2.17). From Eqs. (2.11–2.14) and Eq. (2.17), it follows that

$$\mathbf{r} = \mathbf{d}_r - \widetilde{\mathbf{F}}_{I_{rc}} \widetilde{\mathbf{K}}_{cc}^{*^{-1}} \tilde{\mathbf{f}}_c^{*} - (\mathbf{F}_{I_{rr}} + \widetilde{\mathbf{F}}_{I_{rc}} \widetilde{\mathbf{K}}_{cc}^{*^{-1}} \widetilde{\mathbf{F}}_{I_{rc}}^{T})\lambda = \sum_{s=1}^{N_s} \mathbf{B}_r^{(s)} \mathbf{u}_r^{(s)}, \tag{3.28}$$

which shows that the residual $\mathbf{r}$ represents the jump of the iterate solution across the

subdomain interfaces. Hence, the FETI-DP method converges when the iterate $\mathbf{u}^n$ becomes continuous across all the subdomain interfaces.

From Eq. (2.14), Eq. (2.18), Eq. (2.17) and Eq. (2.13), it follows that at each iteration of a PGMRES algorithm applied to the solution of problem (2.17), FETI-DP forces the jump of the solution across the subdomain interfaces to be orthogonal to the subspace represented by the matrix $\mathbf{Q}_b$. This feature is a strategy for designing an auxiliary coarse problem that, when $\mathbf{Q}_b$ is well chosen, accelerates the convergence of a DDM [18].

Suppose that the space of traces on Γ' of the solution of problem (1.4) is best approximated by a set of orthogonal vectors $\{\mathbf{v}_{j_E}\}_{j=1}^{N_v}$, where the subscript E indicates that $\mathbf{v}_{j_E}$ is non-zero only on edge $E \in \Gamma'$. Then, the residual $\mathbf{r}$ defined in Eq. (3.28) can be approximated as

$$\mathbf{r} \approx \sum_{j=1}^{N_v} \alpha_j \mathbf{v}_{j_E}, \tag{3.29}$$

where $\{\alpha_j\}_{j=1}^{N_v}$ is a set of N_v coefficients. If each augmentation coarse mode is chosen as

$$\mathbf{q}_j = \mathbf{v}_{j_E}, \qquad j = 1, \cdots, N_Q, \tag{3.30}$$

Eq. (2.14) simplifies to

$$\alpha_j = 0, \qquad j = 1, \cdots, \ N_Q. \tag{3.31}$$

In this case, Eq. (3.31) implies that at each iteration of the PGMRES algorithm, the first N_Q components of the residual $\mathbf{r}$ in the independent set $\{\mathbf{v}_{j_E}\}_{j=1}^{N_v}$ are zero. Hence, if a few vectors $\{\mathbf{v}_{j_E}\}_{j=1}^{N_Q}$, $N_Q << N_v$, that dominate the expansion (3.29) can be found, choosing these vectors as coarse augmentation modes can be expected to accelerate the convergence of the iterative solution of the dual interface problem (2.17). It remains to exhibit such a set of orthogonal vectors $\mathbf{v}_{j_E}$ and construct a computationally efficient matrix $\mathbf{Q}_b$.

Recall that the vectors $\mathbf{v}_{j_E}$ are supposed to approximate the traces on Γ' of the solution of problem (1.4), and this algebraic problem arises from the FE discretization of the BVP (1.1). It follows that free-space solutions of the BVP (1.1) are good candidates for constructing $\mathbf{v}_{j_E}$ and therefore $\mathbf{Q}_b$. Plane waves are free-space solutions of the BVP (1.1), and the following real-valued representation of plane waves is also a free-space solution of the BVP (1.1):

$$v = \sum_{j=1}^{\infty} \big(c_{0_j} \sin(k\theta_j \cdot X) + c_{1_j} \cos(k\theta_j \cdot X)\big), \tag{3.32}$$

where $\theta_j \in \mathbb{R}^3$ is an arbitrary vector of unit length defining the direction of propagation of a plane wave, and c_{0_j} and c_{1_j} are two real coefficients. From Eq. (3.32) and Eq. (3.30) it follows that one choice for $\mathbf{Q}_b$ is the matrix composed of blocks of two columns, where the columns of each block are associated with one direction of propagation θ_j and one edge E of the mesh partition. This choice can be expressed as follows:

$$\mathbf{q}_{b_l}[m] = \sin(k\theta_j \cdot X_m), \quad \mathbf{q}_{b_{l+1}}[m] = \cos(k\theta_j \cdot X_m), \quad l = 2(j-1)+1, \quad m = 1, \cdots, N_I - N_c, \tag{3.33}$$

where $\mathbf{q}_b[m]$ designates the entry of $\mathbf{q}_b$ associated with the dof attached to the m-th node on an edge $E \in \Gamma'$, and $X_m \in \mathbb{R}^3$ denotes the coordinates of this m-th node. Hence, if N_E denotes the number of edges of the mesh partition, and N_θ the number of considered directions of wave propagation, the total number of augmentation coarse modes is given by by

$$N_Q = 2N_E N_\theta. \tag{3.34}$$

To these modes one can add the edge-based constant modes as these are also free-space solutions of the BVP (1.1) with $k = 0$.

In this paper, the directions θ_j are generated as follows. A generic cube is discretized into $n_c \times n_c \times n_c$ points. A direction θ_j is defined by connecting the center of the cube to a point lying on a face of the cube (Fig. 4). Since each direction θ_j is used to define both a cosine and a sine mode, only one direction is retained for each pair of opposite directions. This results in a total number of directions

$$N_\theta = \frac{\left(n_c^3 - (n_c - 2)^3\right)}{2}. \tag{3.35}$$

Hence, for a specified N_θ, n_c is chosen so that N_θ is as close as possible to the value given by the above equation.

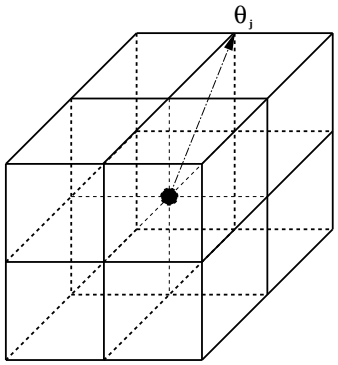

Fig. 4. Generation of the directions of wave propagation ($n_c = 3$).

3.4. *Filtering the coarse space*

There are at least two mechanisms that can cause the matrix $\mathbf{Q}_b$ described in Section 3.3 to be rank deficient:

1. If a direction θ_j turns out to be perpendicular to an edge $E \in \Gamma'$, $\theta_j \cdot X_m$ becomes constant for all $X_m \in E$, the rank of the two-column block of $\mathbf{Q}_b$ associated with the

edge E and the direction θ_j becomes equal to one, and therefore $\mathbf{Q}_b$ becomes rank deficient.

2. In the appendix, it is shown that an interplay between a low mesh resolution and N_θ can also cause $\mathbf{Q}_b$ to become rank deficient.

From Eq. (2.14) or the definition of $\widetilde{\mathbf{K}}^*_{cc}$ given in the fifth of Eqs. (2.18), it follows that whenever $\mathbf{Q}_b$ is rank deficient, $\widetilde{\mathbf{K}}^*_{cc}$ becomes singular. Singular coarse problems can be solved by a number of different techniques among which the drop-tolerance-based direct method described in [23]. However, preventing $\mathbf{Q}_b$ from being rank deficient — and therefore preventing $\widetilde{\mathbf{K}}^*_{cc}$ from being singular — is desirable as this improves the computational efficiency as well as the robustness of the solution of the coarse problem (2.19), and therefore enhances the efficiency and robustness of the overall solution method.

Let ϵ_{QR} denote a small tolerance value. The following describes a simple procedure for filtering the matrix of augmentation coarse modes $\mathbf{Q}_b$ and transforming it into a matrix $\mathbf{Q}^*_b$ that has full column rank:

1. Perform the **QR** factorization [24] of $\mathbf{Q}_b$.

2. Construct $\mathbf{Q}^*_b$ as the union of the columns $\mathbf{q}_j$ of $\mathbf{Q}_b$ for which $\mathbf{R}_{jj} > \epsilon_{QR}$.

The following observations are worthy noting:

1. Since each column of $\mathbf{Q}_b$ has non-zero entries only for the dof associated with a corresponding edge E, the **QR** factorization of $\mathbf{Q}_b$ entails only *local* computations that can be performed on an edge-by-edge basis. Hence, these computations are amenable to an efficient parallelization.

2. A too small value of ϵ_{QR} can result in a matrix $\mathbf{Q}^*_b$ that is still rank deficient, thereby defeating the purpose of the filtering procedure. On the other hand, a too large value of ϵ_{QR} can only cause an excessive filtering. Since the columns of $\mathbf{Q}_b$ serve only the purpose of defining *optional* constraints aimed at accelerating the convergence of the iterative DDM (see Eq. (2.14)), it follows that this DDM is better served by setting ϵ_{QR} to a value that is sufficiently large to prevent $\mathbf{Q}^*_b$ from being rank deficient, even if such a tolerance can cause occasionally an excessive filtering of the matrix $\mathbf{Q}_b$.

Remark. From Section 3.3 (and in particular Eq. (3.34)), it follows that given a mesh partition characterized by a number of subdomains N_s, fixing the number of wave directions N_θ determines the size of the coarse problem (2.19) and therefore fixes in principle the cost of an iteration of the extension of the FETI-DP method developed here, independently of the value of the wave number k. However, when the tolerance ϵ_{QR} is fixed, N_s and N_θ are kept constant, but k is varied, the filtering procedure described above can result in a size of the coarse problem (2.19) that varies with k. Essentially, this is because the columns of $\mathbf{Q}_b$ and the matrix $\widetilde{\mathbf{K}}^*_{cc}$ — depend on k. For a fixed value of ϵ_{QR}, a higher value of k

typically results in filtering out fewer columns of $\mathbf{Q}_b$ and therefore in generating a larger coarse problem (2.19) than a lower value of k. For this reason, in practice, the cost of one iteration of the extension of the FETI-DP solver developed in this paper can be expected to vary with k even when the mesh partition (N_s) and the number of wave directions (N_θ) are fixed.

3.5. *Tailoring the Dirichlet and lumped preconditioners*

$\mathbf{M}_{rr}^{(s)}$ is a mass matrix; hence, in three dimensions and at the element level, this matrix is proportional to h^3. On the other hand, $\mathbf{K}_{rr}^{(s)}$ is a stiffness matrix; in three dimensions and at the element level, it is proportional to h. It follows that for a sufficiently fine mesh, $\mathbf{Z}_{rr}^{(s)}$ is dominated by $\mathbf{K}_{rr}^{(s)}$. These observations suggest that the local matrices $\mathbf{Z}_{rr}^{(s)}$ can be preconditioned by Dirichlet and lumped constructs (see Section 2.4) that are based on the $\mathbf{K}_{rr}^{(s)}$ stiffness matrices rather than the $\mathbf{Z}_{rr}^{(s)}$ impedance matrices. This results in real-valued $\overline{\mathbf{F}}_{I_{rr}}^{D^{-1}}$ and $\overline{\mathbf{F}}_{I_{rr}}^{L^{-1}}$ local preconditioners and therefore reduces storage cost. This intuitive reasoning can be mathematically justified (for example, see [25,26] and the references cited therein).

3.6. *A subdomain-based preconditioned GMRES solver with an auxiliary coarse problem*

In summary, the FETI-DP method is extended here to address the solution of interior and exterior Hemlmholtz problems by

1. Replacing the CG interface solver by the GMRES interface solver.
2. Adapting the Dirichlet and lumped preconditioners as explained in Section 3.5.
3. Choosing as augmentation coarse modes the plane wave modes originally introduced in [4], justified in Section 3.3, and expressed in real-valued form.

4. Preliminary Parallel Implementation

All computational steps of the FETI-DPH method described above, except the solution of the coarse problem (2.19), can be carried out on a subdomain-by-subdomain basis. Hence, these computations are trivially parallelized by assigning one or several subdomains to each processor. In this work, all subdomain problems are solved by a sparse direct method. If the given parallel processor consists of a network of shared memory parallel boxes (or boards), the coarse problem is duplicated in each box. Otherwise, it is duplicated in each processor. In all cases, this coarse problem is solved by a *sequential* sparse direct method. As long as the size of the coarse problem is small compared to the size of the global problem, and the target number of processors N_p is also small, this serial step should not affect the overall performance of the parallel implementation of the FETI-DPH solver adopted in this work. Otherise, it can degrade the overall parallel performance (see next section for details).

5. Numerical Scalability and Performance Assessment

In this section, attention is focused on the numerical scalability and CPU performance of the FETI-DPH method presented in this paper. For this purpose, two different problems are considered. The first one is an academic Helmholtz problem that has the merit of being simple to reproduce by the interested reader. The second problem involves a mockup submarine and highlights the potential of the FETI-DPH method for realistic acoustic scattering applications. All computations are performed in double-precision arithmetic on a Silicon Graphics Origin 3200 computer equipped with R12000 400 MHz processors. This parallel system consists of 5 "boxes" of 8 processors each. On each box, the 8 processors share 8 Gigabytes of real memory. The 5 boxes are interconnected by a Brocade Silkwork 2400 switch.

As coded, the FETI-DPH solver relies on MPI for interprocessor communication within a box and across boxes, and allows the number of subdomains, N_s, to be different from the number of processors, N_p. In this work, the filtering tolerance is set to

$$\epsilon_{QR} = 10^{-2} \tag{5.36}$$

and the convergence of this iterative solver is declared when the relative residual satisfies

$$RE^n = \frac{\|\mathbf{Z}\mathbf{u}^n - \mathbf{f}\|_2}{\|\mathbf{f}\|_2} \leq 10^{-6}. \tag{5.37}$$

5.1. *A three-dimensional waveguide problem*

Here, Θ is replaced by void (and therefore $\alpha = \beta = 0$), $\Upsilon \equiv \Omega$, $\Sigma \equiv \partial\Omega$, and the computational domain is the cube defined by $\Omega = [0,\ 1] \times [0,\ 1] \times [0,\ 1]$. The following boundary conditions are applied on two parallel faces of this computational domain:

$$\begin{aligned} u &= 1 \quad \text{on} \quad y = 0 \\ \frac{\partial u}{\partial \nu} + iku &= 0 \quad \text{on} \quad y = 1. \end{aligned}$$

In all cases, the computational domain Ω is uniformly discretized by eight-noded brick elements and the resulting mesh is uniformly partitioned into N_s subdomains.

First, the number of subdomains is fixed to $N_s = 5 \times 5 \times 5 = 125$ and the mesh size is varied as $\frac{1}{50} \leq h \leq \frac{1}{20}$. Two wave numbers are considered, $k = 4$ and $k = 20$. The performance results of the FETI-DPH solver equipped with the Dirichlet preconditioner are reported in Table 1 for $k = 4$ and Table 2 for $k = 20$, where N_{itr} denotes the number of iterations for convergence. For $k = 4$, the generated meshes appear to be well-resolved as measured by the number of elements per wave length, $G = \frac{2\pi}{kh}$. In this case, the FETI-DPH solver exhibits a perfect numerical scalability with repect to the problem size when $N_\theta = 3$. The contrast between the performance results obtained with $N_\theta = 0$ and $N_\theta = 3$ highlights

the significant effect on convergence of the augmentation of the coarse problem. For $k = 20$, most of the considered meshes appear to be under-resolved. In this case too, the FETI-DPH solver exhibits a perfect numerical scalability with repect to the problem size; however, it requires twice as many iterations as for the case $k = 4$. To determine whether this increase of the iteration count is due to an effect of the wave number k, or an effect of the mesh resolution G, or to both effects, Table 3 reports the performance results obtained for three fixed values of G that are representative of under-resolved, well-resolved, and over-resolved meshes, and a wave number $4 \leq k \leq 24$. These performance results indicate that for a fixed value of N_θ, the FETI-DPH solver converges in a rather small number of iterations for all values $4 \leq k \leq 24$; nevertheless, for any of the considered mesh resolutions, the convergence of the FETI-DPH solver does not exhibit a perfect numerical scalability with respect to the wave number when N_s and N_θ — and therefore the size of the augmented coarse problem — are fixed. However, the performance results summarized in Table 3 also show that when k is increased, the iteration count of FETI-DPH can be maintained constant if N_θ is increased appropriately. In other words, FETI-DPH can achieve numerical scalability with respect to the wave number if the size of the augmented coarse problem is increased appropriately.

Table 1. Waveguide problem: numerical scalability with respect to the mesh size

$k = 4$, $N_s = 125$, Dirichlet preconditioner

h	G	N_{itr} ($N_\theta = 0$)	N_{itr} ($N_\theta = 3$)
1/20	31.4	72	4
1/30	47.1	56	4
1/40	62.8	39	4
1/50	78.5	82	4

Table 2. Waveguide problem: numerical scalability with respect to the mesh size

$k = 20$, $N_s = 125$, Dirichlet preconditioner

h	G	N_{itr} ($N_\theta = 0$)	N_{itr} ($N_\theta = 3$)
1/20	6.28	264	7
1/30	9.42	392	9
1/40	12.6	331	9
1/50	15.7	363	10

Next, the mesh size is fixed to $h = \dfrac{1}{40}$ and the number of subdomains is varied as $8 \leq N_s \leq 125$. The performance results reported in Table 4 and Table 5 show that in this case, the augmented FETI-DPH solver exhibits numerical scalability with respect to the number of subdomains in both of the low-frequency ($k = 4$) and mid-frequency ($k = 20$) regimes.

Table 3. Waveguide problem: numerical scalability with respect to the mesh size

$N_s = 125$, Dirichlet preconditioner

h	k	G	N_{itr} ($N_\theta = 0$)	N_{itr} ($N_\theta = 3$)
1/5	4	6.75	5	1
1/10	8	6.75	35	15
1/15	12	6.75	61	5
1/20	16	6.75	99	6
1/25	20	6.75	292	8
1/30	24	6.75	594	17 (3 for $N_\theta = 13$)
1/10	4	15.7	26	7
1/20	8	15.7	54	4
1/30	12	15.7	77	5
1/40	16	15.7	250	5
1/50	20	15.7	363	10
1/60	24	15.7	617	14 (4 for $N_\theta = 13$)
1/20	4	31.4	72	4
1/40	8	31.4	67	6
1/60	12	31.4	93	5
1/80	16	31.4	350	6
1/100	20	31.4	453	11
1/120	24	31.4	652	20 (5 for $N_\theta = 13$)

Table 4. Waveguide problem: numerical scalability with respect to the number of subdomains

$k = 4$, $h = 1/40$, $G = 62.8$, Dirichlet preconditioner

N_s	N_{itr} ($N_\theta = 0$)	N_{itr} ($N_\theta = 3$)
8	7	5
64	21	4
125	39	4

Table 5. Waveguide problem: numerical scalability with respect to the number of subdomains

$k = 20$, $h = 1/40$, $G = 12.6$, Dirichlet preconditioner

N_s	N_{itr} ($N_\theta = 0$)	N_{itr} ($N_\theta = 3$)
8	35	28
64	237	15
125	331	9

Next, the mesh size is set to $h = \frac{1}{40}$ and the number of subdomains is fixed to $N_s = 125$. The performance results reported in Table 6 and Table 7 show that in both of the low- and mid-frequency regimes, the effect of coarse problem augmentation on the performance of the FETI-DPH solver dominates that of conventional preconditioning.

Table 6. Waveguide problem: performance of the preconditioner

$k = 4$, $h = 1/40$, $G = 62.8$, $N_s = 125$

Preconditoner	N_{itr} ($N_\theta = 0$)	N_{itr} ($N_\theta = 3$)
None	144	18
Lumped	51	5
Dirichlet	39	4

Table 7. Waveguide problem: performance of the preconditioner

$k = 20$, $h = 1/40$, $G = 12.6$, $N_s = 125$

Preconditoner	N_{itr} ($N_\theta = 0$)	N_{itr} ($N_\theta = 3$)
None	532	26
Lumped	288	10
Dirichlet	331	9

In summary, the performance results presented above suggest that the FETI-DPH solver presented in this paper is numerically scalable with respect to both the problem size and the number of subdomains, and that the augmentation of its coarse problem by plane waves improves significantly its iteration count. However, does augmentation necessarily improve its CPU performance? To answer this question, the mesh size is finally fixed to $h = \frac{1}{60}$, the wave number is varied between $k = 4$ and $k = 20$, the number of subdomains is varied between $N_s = 64$ and $N_s = 216$, and the FETI-DPH solver is applied with and without coarse problem augmentation to the solution on an 8-processor Origin 3200 of the resulting systems of equations. The performance results reported in Table 8 reveal that:

- As in the case of many DDMs, the CPU performance of FETI-DPH depends on the number of subdomains. Here, this is particulary noticeable not only because the size of the (augmented) coarse problem depends on N_s, but also because the coarse problem is solved by a sequential algorithm.

- In all frequency regimes, the FETI-DPH solver is faster with a suitable coarse problem augmentation than without, from both iteration count and CPU time viewpoints.

- For this three-dimensional waveguide problem, the FETI-DPH solver delivers a scal-

able CPU performance with respect to the wavenumber when the number of subdomains is appropriately chosen and N_θ is appropriately increased with k.

Table 8. Waveguide problem: CPU performance on an 8-processor Origin 3200

$h = 1/60$, Dirichlet preconditioner

k	G	N_θ	N_s	size of $\tilde{\mathbf{K}}^*_{cc}$	N_{itr}	CPU time
4	94.2	0	27	28	24	105.2 s.
4	**94.2**	**0**	**64**	**72**	**37**	**65.5 s.**
4	94.2	0	125	144	67	67.0 s.
4	94.2	1	125	744	6	25.9 s.
4	**94.2**	**1**	**216**	**1,330**	**5**	**21.3 s.**
4	94.2	1	1,000	6,534	4	28.7 s.
12	31.4	0	27	28	82	191.8 s.
12	**31.4**	**0**	**64**	**72**	**56**	**82.3 s.**
12	31.4	0	125	144	93	85.6 s.
12	31.4	2	125	1,144	7	29.0 s.
12	**31.4**	**2**	**216**	**2,050**	**6**	**24.5 s.**
12	31.4	2	1,000	1,0134	5	37.6 s.
20	18.8	0	27	28	188	358.0 s.
20	**18.8**	**0**	**64**	**72**	**268**	**323.0 s.**
20	18.8	0	125	144	478	547.5 s.
20	18.8	3	125	1,344	10	32.8 s.
20	**18.8**	**3**	**216**	**2,410**	**6**	**26.9 s.**
20	18.8	3	1,000	11,934	9	58.2 s.

The bold font is used to highlight for each wave number
the parameter configuration leading to optimal CPU performance

5.2. *Acoustic signature of a mockup submarine*

Next, the FETI-DPH solver equipped with the Dirichlet preconditioner is applied to the acoustic signature analysis at $k = 4$ of the mockup submarine shown in Fig. 5, where $L = 10$, $D = 1$, $l = 0.5$, $h_e = 0.5$, and $w = 0.25$. For this application of the BVP (1.1), the artificial boundary Σ is chosen as an ellipsoid (Fig. 6) positioned so that the distance between each end of the main tube and the neighboring extreme point of Σ located on the main axis of the ellipsoid is $m_1 = 0.5$, and the distance between the top of the submarine tower and the neighboring extreme point of Σ located on the minor axis of the ellipsoid is $m_2 = 1.52$. The direction of the incident wave is set to $\mathbf{t} = [\sqrt{2}/2,\ \sqrt{2}/2,\ 0]^T$. The computational domain delimited by the ellipsoid Σ and the surface of the mockup submarine is discretized by four-noded tetrahedral elements. A first mesh M1 with 264,436 grid points ($h \approx \dfrac{1}{16}, G \approx 25$) is generated for computations on eight processors of the Origin 3200 system.

Using the mesh M1, N_s and N_θ are varied until the optimal CPU performance of the FETI-DPH solver is reached. This process is documented in Table 9. Again, augmentation

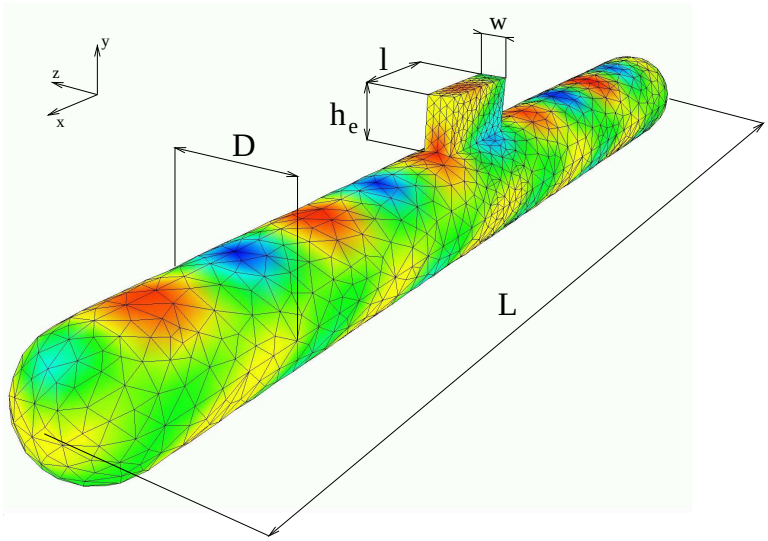

Fig. 5. Mockup submarine: geometric features

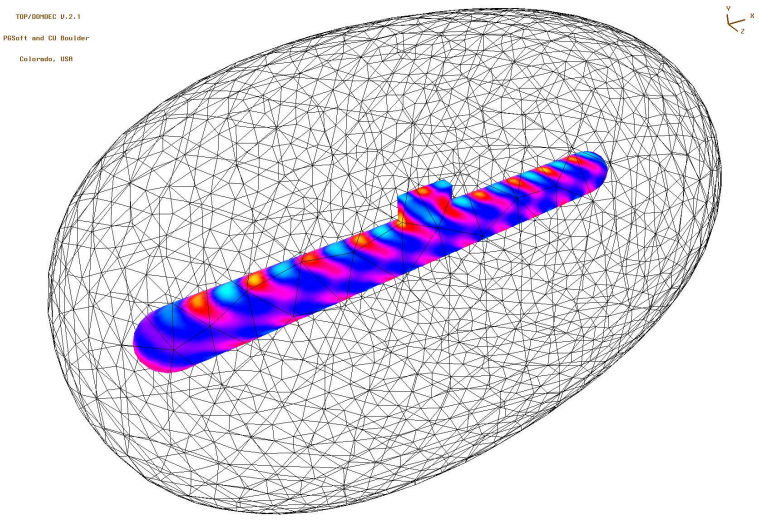

Fig. 6. Mockup submarine: computational domain delimited by an ellipsoidal artificial boundary Σ, and real part of the trace of the scattered field on the surface of the scatterer

of the coarse problem is reported to be effective at reducing both the iteration count and the solution time, and the performance of the augmented FETI-DPH solver is shown to be scalable with respect to the number of subdomains.

Table 9. Mockup submarine problem: CPU performance on an 8-processor Origin 3200

$k = 4$, $h \approx 1/16$, $G \approx 25$, $N^{dof} = 264{,}436$, Dirichlet preconditioner

N_θ	N_s	size of $\widetilde{\mathbf{K}}^*_{cc}$	N_{itr}	CPU time
0	8	7	121	441.5 s.
0	25	54	247	294.9 s.
0	50	81	327	305.5 s.
2	102	2,212	37	64.7 s.
2	201	4,398	21	44.0 s.
2	301	6,650	16	42.3 s.

Next, the number of subdomains and the number of plane wave directions are fixed to $N_s = 201$ and $N_\theta = 2$, respectively, the number of processors is varied between $N_p = 1$ and $N_p = 32$, and the obtained performance results are reported in Table 10. These results show that excellent to reasonable parallel speed-ups and efficiencies are achieved for $2 \leq N_p \leq 16$. In this range of number of processors, the amount of CPU time consumed by the sequential solution of the coarse problem varies between 2.6 % and 29.2 % of the total CPU time consumed by the parallel FETI-DPH solver. Once the ratio of sequential to parallel work reaches 37.4 %, increasing the number of processors from 16 to 32 decreases the parallel efficiency of the FETI-DPH solver from 57 % to 36 %.

Table 10. Mockup submarine problem: parallel scalability on a 32-processor Origin 3200

$k = 4$, $h \approx 1/16$, $G \approx 25$, $N^{dof} = 264{,}436$, $N_\theta = 2$, $N_s = 201$, Dirichlet preconditioner

N_p	CPU (solution coarse problem)	CPU (total)	Parallel speed-up (overall)	Parallel efficiency (overall)
$1 \times 1 = 1$	5.7 s.	217.8 s.	1	100 %
$1 \times 2 = 2$	5.7 s.	118.6 s.	1.8	92 %
$4 \times 1 = 4$	7.3 s.	60.6 s.	3.6	90 %
$4 \times 2 = 8$	7.3 s.	36.8 s.	5.9	74 %
$4 \times 4 = 16$	7.0 s.	24.0 s.	9.1	57 %
$4 \times 8 = 32$	7.0 s.	18.7 s.	11.6	36 %

The pollution effect of discretization by linear elements dictates decreasing the mesh size with increasing the wave number as to maintain constant the product k^3h^2. For this reason, the acoustic signature analysis at $k = 5.4$ of the mockup submarine described above is performed using a mesh M2 with 1,038,974 grid points. In this case, the best CPU performance of the FETI-DPH solver is obtained using $N_\theta = 2$ and $N_s = 501$ (see Table 11). The performance results gathered in Table 11 also show that when the wave number is increased from $k = 4$ to $k = 5.4$ and the mesh is refined to maintain constant k^3h^2, the

FETI-DPH solver achieves numerical scalability with respect to k. Mesh M2 is 3.9 times as large as mesh M1, but the FETI-DPH CPU time associated with mesh M2 is reported to be 6.8 times larger than that associated with mesh M1. This is mainly due to the fact that to maintain numerical scalability when increasing the wave number from $k = 4$ to $k = 5.4$, the coarse problem has to be increased by a factor equal to 2.9. In such as case, solving the coarse problem in serial mode while performing the rest of the computations in parallel mode prevents the FETI-DPH solver from achieving CPU time scalability with respect to the wave number.

Table 11. Mockup submarine problem: CPU performance on a 16-processor (4×4) Origin 3200

Dirichlet preconditioner

k	h	k^3h^2	N^{dof}	N_θ	N_s	size of $\widetilde{\mathbf{K}}^*_{cc}$	N_{itr}	CPU time
4	1/16	0.25	264,436	2	201	4,398	21	24.0 s.
5.4	1/25	0.25	1,038,974	2	501	12,836	22	162.5 s.

Finally, the FETI-DPH solver is benchmarked against the previously developed FETI-H counterpart. For this purpose, similar parallel implementations are used for both solvers. The wave numbers $k = 4$ and $k = 8$ are considered, mesh M1 is employed, and N_θ and N_s are varied until the optimal CPU performance on eight processors of the Origin 3200 system is reached separately for each of the two solvers. These optimal performance results are summarized in Table 12. Again, the reader can observe that because of the serialization of the solution of the coarse problem, the FETI-DPH solver cannot achieve simultaneously numerical scalability with respect to k and best parallel CPU performance. However, in both cases, the FETI-DPH solver is reported to outperform its FETI-H counterpart by more than a factor two.

Table 12. Mockup submarine problem: FETI-DPH vs. FETI-H on an 8-processor Origin 3200

$h = 1/16$, $N^{dof} = 264{,}436$, Dirichlet preconditioner, augmented coarse problem

k	G	N_s DPH (H)	N_θ DPH (H)	size of $\widetilde{\mathbf{K}}^*_{cc}$ DPH (H)	N_{itr}	CPU time
4	25.00	201 (301)	2 (2)	4,398 (2,709)	21 (48)	44.0 s. (85.2 s.)
8	12.50	201 (301)	4 (3)	8,076 (8,127)	59 (76)	124.9 s. (322.1 s.)
8	12.50	301 (NA)	6 (NA)	16,792 (NA)	25 (NA)	217.5 s. (NA)

6. Conclusions

The FETI-DPH iterative solver presented in this paper is an extension of the FETI-DP method to Helmholtz problems in which the interface problem is solved by a preconditioned generalized minimum residual algorithm instead of a conjugate gradient algorithm, and the basic coarse problem is augmented with plane waves. Numerical studies performed

for waveguide and acoustic scattering problems show that this solver achieves numerical scalability with respect to both the problem size and the number of subdomains. They also reveal that when the coarse problem size is increased with the wave number, the FETI-DPH solver also achieves numerical scalability with respect to the wave number. A parallel implementation of FETI-DPH in which the solution of the coarse problem is simply serialized is found to deliver good parallel efficiency for up to 16 processors. It also delivers CPU time scalability with respect to the problem size. However, if this implementation is to be executed efficiently on a larger number of processors, and/or if it is to achieve CPU time scalability with respect to the wave number, it requires first parallelizing the coarse problem solver.

7. Acknowledgments

Charbel Farhat acknowledges partial support by the Sandia National Laboratories under Contract No. 29341, and partial support by HPTi, Inc. Philip Avery and Radek Tezaur acknowledge partial support by HPTi, Inc., and partial support by the Office of Naval Research under Grant N00014-04-1-0241. Jing Li acknowledges partial support by the National Science Foundation under Grant No. DMS-0209297. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Sandia National Laboratories, or HPTI, Inc., or the Office of Naval Research, or the National Science Foundation.

Appendix A

The objective of this appendix is to show that an interplay between a low mesh resolution and the number of wave directions N_θ can cause the matrix $\mathbf{Q}_b$ to become rank deficient. Hence, this interplay is one among several motivations for filtering the matrix of augmentation coarse modes $\mathbf{Q}_b$ as described in Section 3.4.

For simplicity, the following proposition considers the case of a two-dimensional problem, a uniform discretization characterized by a mesh size h, and an edge E of the mesh partition that is aligned with the x-axis and begins at its origin. In this case,

$$\theta_j = \begin{bmatrix} \cos\vartheta_j \\ \sin\vartheta_j \end{bmatrix}, \qquad X_m = \begin{bmatrix} x_m \\ 0 \end{bmatrix} = \begin{bmatrix} mh \\ 0 \end{bmatrix}, \qquad \theta_j \cdot X_m = x_m \cos\vartheta_j, \tag{A.1}$$

and Eqs. (3.33) become

$$\mathbf{q}_{b_l}[m] = \sin(kx_m \cos\vartheta_j), \quad \mathbf{q}_{b_{l+1}}[m] = \cos(kx_m \cos\vartheta_j), \quad l = 2(j-1)+1, \quad m = 1, \cdots, N_I - N_c. \tag{A.2}$$

Proposition. If the augmentation coarse problem contains two directions θ_i and θ_j whose angles with the x-axis ϑ_i and ϑ_j satisfy

$$\cos\vartheta_i \pm \cos\vartheta_j = \frac{2n\pi}{kh}, \tag{A.3}$$

where n is an integer, then

1. ϑ_i and ϑ_j satisfy

$$\cos(kx_m \cos\vartheta_i) = \cos(kx_m \cos\vartheta_j), \quad \forall X_m \in E. \tag{A.4}$$

2. The matrix $\mathbf{Q}_b$ defined by Eqs. (A.2) is rank deficient.

Proof. Since E is aligned with the x-axis and begins at its origin, and the mesh is assumed to be uniform, then

$$\forall\ X_m \in E,\ x_m = mh. \tag{A.5}$$

From Eq. (A.3) and Eq. (A.5), it follows that

$$\begin{aligned}
\cos(kx_m \cos\vartheta_i) &= \cos(kx_m(\tfrac{2n\pi}{kh} \mp \cos\vartheta_j)) \\
&= \cos(\mp kx_m \cos\vartheta_j + x_m \frac{2n\pi}{h}) \\
&= \cos(\mp kx_m \cos\vartheta_j + (mh)\frac{2n\pi}{h}) \\
&= \cos(\mp kx_m \cos\vartheta_j + 2nm\pi).
\end{aligned}$$

Hence, $\forall\ X_m \in E$,

$$\cos(kx_m \cos\vartheta_i) = \cos(kx_m \cos\vartheta_j).$$

From Eq. (A.2) and the above result, it follows that the matrix $\mathbf{Q}_b$ has at least two identical columns and therefore is rank deficient. □

Since $|\cos\vartheta_i \pm \cos\vartheta_j| \leq 2$, the result stated in the above proposition is significant when $\dfrac{2\pi}{kh} \leq \dfrac{2}{n}$ — and therefore, when the mesh resolution is less or equal to 2 elements per wave length ($n = 1$) — and when the augmentation coarse problems contains two directions θ_i and θ_j for which $\cos\vartheta_i = -\cos\vartheta_j$ ($n = 0$).

References

1. L. F. Pavarino and A. Toselli, *Recent Developments in Domain Decomposition Methods* (Springer, 2002).
2. X. C. Cai and O. Widlund, *SIAM J. Sci. Statist. Comput.* **13**, 243 (1992).
3. J. Xu and X. C. Cai, *Math. Comp.* **59**, 311 (1992).
4. C. Farhat, A. Macedo, and M. Lesoinne, *Numer. Math.* **85**, 283 (2000).
5. R. Tezaur, A. Macedo, and C. Farhat, *Internat. J. Numer. Meths. Engrg.* **51**, 1175 (2001).
6. R. Tezaur, A. Macedo, C. Farhat, and R. Djellouli, *Internat. J. Numer. Meths. Engrg.* **53**, 1461 (2002).
7. C. Farhat, R. Tezaur, and R. Djellouli, *Inverse Problems* **18**, 1229 (2002).
8. C. Farhat, *J. Comput. Sys. Engrg.* **2**, 149 (1991).
9. C. Farhat and F. X. Roux, *Internat. J. Numer. Meths. Engrg.* **32**, 1205, (1991).
10. C. Farhat and F. X. Roux, *SIAM J. Sc. Stat. Comp.* **13**, 379 (1992).
11. C. Farhat, J. Mandel, and F. X. Roux, *Comput. Meths. Appl. Mech. Engrg.* **115**, 367 (1994).
12. J. Mandel and R. Tezaur, *Numer. Math.* **88**, 543 (2001).

13. J. Mandel, R. Tezaur, and C. Farhat, *SIAM J. Numer. Anal.* **36**, 1370 (1999).
14. C. Farhat, M. Lesoinne, and K. Pierson, *Numer. Lin. Alg. Appl.* **7**, 687 (2000).
15. C. Farhat, M. Lesoinne, P. LeTallec, K. Pierson, and D. Rixen, *Internat. J. Numer. Meths. Engrg.* **50**, 1523 (2001).
16. A. Klawonn, O. B. Widlund, and M. Dryja, *SIAM J. Numer. Anal.* **40**, 159 (2002).
17. M. Bhardwaj, K. Pierson, G. Reese, T. Walsh, D. Day, K. Alvin, J. Peery, C. Farhat, and M. Lesoinne, "Salinas: a scalable software for high-performance structural and solid mechanics simulations," in *Proc. IEEE/ACM SC2002 Conference*.
18. C. Farhat, P. S. Chen, F. Risler, and F. X. Roux, *Internat. J. Numer. Meths. Engrg.* **42**, 257 (1998).
19. C. Farhat and J. Mandel, *Comput. Meths. Appl. Mech. Engrg.* **155**, 129 (1998).
20. C. Farhat, P. S. Chen, J. Mandel, and F. X. Roux, *Comput. Meths. Appl. Mech. Engrg.* **155**, 153 (1998).
21. M. Lesoinne, "A FETI-DP corner selection algorithm for three-dimensional problems," in *Proc. 2002 Domain Decomposition Conference*.
22. D. Rixen and C. Farhat, *Internat. J. Numer. Meths. Engrg.* **44**, 489 (1999).
23. C. Farhat and M. Géradin, *Internat. J. Numer. Meths. Engrg.* **41**, 675 (1998).
24. G. H. Golub and C. F. Van Loan, *Matrix Computations* (Johns Hopkins, 1983).
25. H. Yserentant, *Numer. Math.* **54**, 719 (1988).
26. A. Klawonn, "Preconditioners for Indefinite Problems," Ph.D. thesis, Wilhelms-Universitat, 1995.