

A NUMERICAL STUDY OF FETI ALGORITHMS FOR MORTAR FINITE ELEMENT METHODS*

DAN STEFANICA[†]

Abstract. The finite element tearing and interconnecting (FETI) method is an iterative substructuring method using Lagrange multipliers to enforce the continuity of the finite element solution across the subdomain interface. Mortar finite elements are nonconforming finite elements that allow for a geometrically nonconforming decomposition of the computational domain into subregions and, at the same time, for the optimal coupling of different variational approximations in different subregions.

We present a numerical study of FETI algorithms for elliptic self-adjoint equations discretized by mortar finite elements. Several preconditioners which have been successful for the case of conforming finite elements are considered. We compare the performance of our algorithms when applied to classical mortar elements and to a new family of biorthogonal mortar elements, and we discuss the differences between enforcing mortar conditions instead of continuity conditions for the case of matching nodes across the interface. Our experiments are carried out for both two and three dimensional problems, and include a study of the relative costs of applying different preconditioners for mortar elements.

Key words. FETI algorithms, mortar finite elements, Lagrange multipliers, domain decomposition

AMS subject classifications. 65F10, 65N30, 65N55

PII. S1064827500378829

1. Introduction. The finite element tearing and interconnecting (FETI) method is an iterative substructuring method using Lagrange multipliers which is actively used in industrial-size parallel codes for solving difficult computational mechanics problems. This method was introduced by Farhat and Roux [26]; a detailed presentation is given in [27], a monograph by the same authors. Originally used to solve second order, self-adjoint elliptic equations, it has later been extended to many other problems, e.g., time-dependent problems [18], plate bending problems [19, 24, 41], heterogeneous elasticity problems with composite materials [43, 44], acoustic scattering and Helmholtz problems [22, 23, 28, 29], linear elasticity with inexact solvers [32], and Maxwell's equations [42, 50]. Another Lagrange multiplier based method, the dual-primal FETI method, has recently been introduced by Farhat et al. [20] and Farhat, Lesoinne, and Pierson [21] for two dimensional problems, and a coordinate-free formulation of the FETI method has been analyzed by Brenner [11].

The FETI method is a nonoverlapping domain decomposition method and requires the partitioning of the computational domain Ω into nonoverlapping subdomains. It has been designed for conforming finite elements and makes use of Lagrange multipliers to enforce pointwise continuity across the interface of the partition. After eliminating the subdomain variables, the dual problem, given in terms of Lagrange

*Received by the editors October 2, 2000; accepted for publication (in revised form) May 30, 2001; published electronically November 7, 2001. Part of this work was carried out while the author was at the Courant Institute of Mathematical Sciences and with support in part by the National Science Foundation under grant NSF-CCR-9732208, and in part by the U.S. Department of Energy under contract DE-FG02-92ER25127.

<http://www.siam.org/journals/sisc/23-4/37882.html>

[†]Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139 and Baruch College, City University of New York, 17 Lexington Avenue, New York, NY 10017 (dstefan@math.mit.edu, <http://www-math.mit.edu/~dstefan>).

multipliers, is solved by a projected conjugate gradient (PCG) method. Once an accurate approximation for the Lagrange multipliers has been obtained, the values of the primal variables are obtained by solving a local problem for each subdomain.

It was shown experimentally in [25] that a certain projection operator used in the PCG solver plays a similar role to that of a coarse problem for other domain decomposition algorithms and that certain variants of the FETI algorithm are numerically scalable with respect to both the subproblem size and the number of subdomains. Mandel and Tezaur later showed that for a FETI method which employs a Dirichlet preconditioner, the condition number grows at most in proportion to $(1 + \log(H/h))^2$ if the decomposition of Ω does not have crosspoints, i.e., the points that belong to the closure of more than two subdomains, and as $C(1 + \log(H/h))^3$ in the general case; cf [40, 49]. Here, H is the subdomain diameter and h is the mesh size. Using a different preconditioner, Klawonn and Widlund obtained a FETI method which converges in fewer iterations than the classical FETI method. In [33], they proved an upper bound for the condition number of their method for elliptic problems with heterogeneous coefficients which is on the order of $(1 + \log(H/h))^2$; see section 4.3 for more details.

Rixen and Farhat [43, 44] considered a Dirichlet preconditioner with a maximal number of pointwise continuity conditions at crosspoints, which results in a FETI algorithm with redundant Lagrange multipliers. It was shown in [33] that this algorithm is equivalent to using the preconditioner of Klawonn and Widlund and nonredundant multipliers for the FETI method.

In this paper, we study the numerical convergence properties of a family of FETI algorithms applied to mortar finite elements. Mortar finite elements are nonconforming finite element methods that allow for a geometrically nonconforming decomposition of the computational domain into subregions and, at the same time, for the optimal coupling of different variational approximations in different subregions. Here, optimality means that the global error is bounded by the sum of the local approximation errors on each subregion.

The importance of our study is related to the inherent advantages of mortar methods over the conforming finite elements. For example, the mesh generation is more flexible and can be made quite simple on individual subregions. This also makes it possible to move different parts of the mesh relative to each other, e.g., in a study of time-dependent problems. The same feature is most valuable in optimal design studies, where the relative position of parts of the model is not fixed a priori. The mortar methods also allow for local refinement of finite element models in only certain subregions of the computational domain, and they are also well suited for parallel computing; cf. [30].

We have used geometrically nonconforming mortar finite elements. Three FETI algorithms with different preconditioners for the dual problem have been considered: the Dirichlet preconditioner of Farhat and Roux [26], the block-diagonal preconditioner of Lacour [34], and the new preconditioner of Klawonn and Widlund [33]. These algorithms have been implemented for both the classical mortar finite elements of Bernardi, Maday, and Patera [8], and for the new biorthogonal mortar elements of Wohlmuth [52, 53], in two and three dimensions. Based on the results presented in this paper, we proposed and analyzed a new FETI algorithm with inexact solvers applied to three dimensional mortar finite elements in [47]. We note that a study of a FETI preconditioner for Maxwell's equations on nonmatching grids has been completed by Rapetti and Toselli [42].

Our results show that the Dirichlet preconditioner does not perform well in the

mortar case, since convergence is achieved only after hundreds or thousands of iterations. However, the new preconditioner performs satisfactory; i.e., the number of iterations required to achieve convergence and the condition number of the algorithms depend only weakly on the number of nodes in each subregion and is independent of the number of subregions. For each of the three preconditioners, using the biorthogonal mortars results in algorithms which require less computational effort and fewer iterations than those using the classical mortar finite elements.

We have also studied the extra computational effort, due to the complexity of the mortar conditions, required for the implementation of the FETI algorithm with the new preconditioner. These costs might have been significant, in particular, in the three dimensional case. We conclude that the improvement of the iteration count was enough to offset this extra cost.

In the conforming finite element case, the meshes across the interface match. Therefore, continuity conditions, as well as mortar conditions, may be enforced across the interface. We have studied the differences between the FETI algorithms using both types of constraints in terms of iteration counts and computational costs. We conclude that the new preconditioner for either continuity conditions or for biorthogonal mortars results in the best algorithms.

The rest of the paper is structured as follows. In the next section, we describe the mortar finite element method. In section 3, we present the classical FETI method and the Dirichlet preconditioner, and in section 4, we discuss the FETI algorithm for mortars with two different preconditioners. In sections 5 and 6, we present numerical comparisons of the performances of three different FETI algorithms for mortar finite elements, and for two and three dimensional problems, respectively. In the last section, we discuss the differences between enforcing mortar conditions instead of continuity conditions for the case of matching nodes across the interface.

2. Mortar finite elements. The mortar finite element methods were first introduced by Bernardi, Maday, and Patera in [8] for low order and spectral finite elements. A three dimensional version was developed by Ben Belgacem and Maday in [7] and was further analyzed for three dimensional spectral elements in [6]. Another family of biorthogonal mortar elements has recently been introduced by Wohlmuth [52, 53]. See also [45] for mortar hp finite elements and [4, 12, 31] for mortar $H(\text{curl})$ elements. Cai, Dryja, and Sarkis [13] have extended the mortar methods to overlapping decompositions.

Several domain decomposition methods for mortar finite elements have been shown to perform similarly to the case of conforming finite elements; cf. [3, 15] for iterative substructuring methods, [16, 37, 38] for Neumann–Neumann algorithms, and [36, 48] for the FETI method. For other studies of preconditioners for the mortar method, see [14] for a hierarchical basis preconditioner and [1, 2] for iterative substructuring preconditioners. Multigrid methods have also been used to solve mortar problems; cf. [9, 10, 51].

2.1. Two dimensional low order mortar finite elements. For simplicity, we restrict our presentation to mortar elements used in our numerical experiments, i.e., first order mortar finite elements on polygonal domains. The computational domain Ω is decomposed using a nonoverlapping polygonal partition $\{\Omega_i\}_{i=1:N}$. Let $\partial\Omega_D$ be the part of $\partial\Omega$ where Dirichlet conditions are imposed. If an edge of a polygon intersects $\partial\Omega_D$, we require that the entire edge belongs to $\partial\Omega_D$. The partition is said to be geometrically conforming if the intersection between the closure of any two subregions is either empty, a vertex, or an entire edge, and it is geometrically




FIG. 1. *Test functions. Left: classical mortars; right: new mortars.*

nonconforming otherwise.

The interface between the subregions $\{\Omega_i\}_{i=1:N}$, denoted by Γ , is defined as the set of points that belong to the boundaries of at least two subregions.

The mortar finite element space V^h is defined as follows: Any mortar function $v \in V^h$ vanishes at all the nodes on $\partial\Omega_D$. The restriction of v to any Ω_i is a P_1 or a Q_1 finite element function. We do not require pointwise continuity across Γ . Instead, we choose a set of open edges of the subregions $\{\Omega_i\}_{i=1:N}$, called nonmortars, which form a disjoint partition of the interface. The edges of $\{\Omega_i\}_{i=1:N}$, which are part of Γ and were not chosen to be nonmortars, are called mortars. Across each nonmortar side γ , we impose weak continuity conditions for v . Let $\Gamma(\gamma)$ be the union of the parts of the mortars that coincides geometrically with $\bar{\gamma}$. Let v_γ and $v_{\Gamma(\gamma)}$ be the restriction of v to γ and $\Gamma(\gamma)$, respectively. The values of v on the nonmortar γ are then given by the mortar conditions

$$(1) \quad \int_{\gamma} (v_\gamma - v_{\Gamma(\gamma)}) \psi \, ds = 0 \quad \forall \psi \in \Psi^h(\gamma),$$

where $\Psi^h(\gamma)$ is the space of test functions. Here, $\Psi^h(\gamma)$ is a subspace of codimension two of $V^h(\gamma)$, the restriction of V^h to γ . It consists of continuous, piecewise linear functions on γ that are constant in the first and last mesh intervals of $\bar{\gamma}$; cf. Figure 1.

We note that a nonmortar partition of the interface is always possible; cf. Stefanica [46]. The partition is not unique, but any choice can be treated the same from a theoretical point of view.

From (1), it results that the interior nodes of the nonmortar sides are not associated with genuine degrees of freedom in the finite element space V^h , while the values of v at the end points of $\bar{\gamma}$ are genuine degrees of freedom. To emphasize this aspect, we present here the matrix formulation of the mortar conditions, which will be further used in section 4.

Let \bar{v}_γ be the vector of the interior nodal values of v on γ . For simplicity, we

assume that the mesh is uniform on γ of mesh size h . Let $\bar{v}_{\Gamma(\gamma)}$ be the vector of the values of v at the end points of γ and at all the nodes on the edges opposite γ such that the intersection of γ and the supports of the corresponding nodal basis functions are not empty. Then \bar{v}_γ is uniquely determined by $\bar{v}_{\Gamma(\gamma)}$; the matrix formulation of the mortar conditions (1) is

$$(2) \quad M_\gamma \bar{v}_\gamma - N_\gamma \bar{v}_{\Gamma(\gamma)} = 0,$$

or, solving for \bar{v}_γ , $\bar{v}_\gamma = P_\gamma \bar{v}_{\Gamma(\gamma)}$, with $P_\gamma = M_\gamma^{-1} N_\gamma$.

We note that N_γ is a banded matrix with a bandwidth of similar size for both the classical and the new mortars. For the classical mortar method, M_γ is a tridiagonal matrix and the mortar projection matrix P_γ is a full matrix. The projection of a nodal basis function from the mortar side results in a function with support equal to γ . The nodal values of this function decay exponentially to 0 at the end points of γ , away from the nodes on γ opposite the support of the nodal basis function from the mortar side.

Since $V^h(\Omega_i) \subset H^1(\Omega_i)$, we know that $v_\gamma \in H^{1/2}(\gamma)$. Thus, the test functions space $\Psi^h(\gamma)$ may be embedded in the dual space of $H^{1/2}(\gamma)$ with respect to the L^2 inner product, and therefore $\Psi^h(\gamma) \subset H^{-1/2}(\gamma)$. Based on this observation, a space of discontinuous piecewise linear test functions $\Psi_{new}^h(\gamma)$ for low order mortars has been developed by Wohlmuth [52]. The test function associated with the first interior node on γ is the constant 1 on the first mesh interval, decreases linearly from 2 to -1 on the second mesh interval, and vanishes everywhere else. A similar test function is introduced for the last interior node on γ . The test function for any other node on γ has the support on the two mesh intervals having the node as an end point; it increases linearly from -1 to 2 on the first interval and decreases from 2 to -1 on the second; cf. Figure 1.

The new mortar space has similar approximation properties to the classical mortar space; cf. [52]. A major advantage of the new mortar finite element space is that the mortar projection can be represented by a banded matrix, as opposed to the classical mortar finite element method, where the mortar projection matrix is, in general, a full matrix. More precisely, for the new mortar method, $M_\gamma = hI$ is a diagonal matrix and $P_\gamma = N_\gamma/h$ is banded. Therefore, the mortar projection of a nodal basis function on the mortar side vanishes outside the mesh intervals on the nonmortar which intersect its support.

2.2. The three dimensional case. For three dimensional problems, the mortars and nonmortars are open faces of the subregions which form the nonconforming decomposition of the computational domain Ω .

To introduce the mortar finite element space, we follow the outline from the previous section. Let $\{\Omega_i\}_{i=1:N}$ be a nonoverlapping polyhedral partition of Ω . If a face or an edge of a polyhedron intersects $\partial\Omega_D$ at an interior point, then the entire face or edge is assumed to belong to $\partial\Omega_D$. The partition is said to be geometrically conforming if the intersection between the closures of any two subregions is either empty, a vertex, an entire edge, or an entire face, and it is nonconforming otherwise.

The nonmortars $\{\mathcal{F}_l\}_{l=1}^L$ are faces of the subregions which form a disjoint partition of the interface Γ . The faces of $\{\Omega_i\}_{i=1:N}$ that are part of Γ and were not chosen to be nonmortars are called mortars.

We now describe the test functions associated with an arbitrary nonmortar face \mathcal{F} . Let $\Gamma(\mathcal{F})$ be the union of parts of mortar faces opposite \mathcal{F} . The test function space $\Psi^h(\mathcal{F})$ is a subset of $V^h(\mathcal{F})$, the restriction of V^h to \mathcal{F} , such that the value of a test function at a node on $\partial\mathcal{F}_l$ is a convex combination of its values at the neighboring

interior nodes of \mathcal{F}_l . If $V^h(\mathcal{F})$ is a P_1 or a Q_1 space, then the dimension of $\Psi^h(\mathcal{F})$ is equal to the number of interior nodes of \mathcal{F} .

The mortar finite element space V^h consists of functions v which vanish at all the nodal points of $\partial\Omega_D$. Its restriction to any Ω_i is a P_1 or a Q_1 finite element function. The values of a mortar function $v \in V^h$ on any nonmortar face \mathcal{F} are given by the mortar conditions

$$\int_{\mathcal{F}} (v_{\mathcal{F}} - v_{\Gamma(\mathcal{F})}) \psi \, ds = 0 \quad \forall \psi \in \Psi^h(\mathcal{F}).$$

A version of the new mortars for the three dimensional case, based on biorthogonal test functions such as those described in section 2.1, has been developed by Wohlmuth. For details, we refer the reader to [53].

3. The classical FETI algorithm. In this section, we review the original FETI method of Farhat and Roux [27] for elliptic problems discretized by *conforming finite elements*. To simplify our presentation, we discuss only the Poisson equation with mixed Neumann–Dirichlet boundary conditions. The extension of the algorithm to the case of other self-adjoint elliptic equations is straightforward.

Let $f \in L^2(\Omega)$. We look for a solution $u \in H^1(\Omega)$ of the mixed boundary value problem

$$(3) \quad \begin{cases} -\Delta u = f & \text{on } \Omega, \\ u = 0 & \text{on } \partial\Omega_D, \\ \frac{\partial u}{\partial n} = 0 & \text{on } \partial\Omega_N, \end{cases}$$

where $\partial\Omega = \partial\Omega_N \cup \partial\Omega_D$. For unique solvability, we require that $\partial\Omega_D$ has positive Lebesgue measure.

On Ω , we consider P_1 or Q_1 finite elements with mesh size h . The finite element mesh is partitioned along mesh lines into N nonoverlapping subdomains $\Omega_i \subset \Omega$, $i = 1 : N$. Since the finite element mesh is conforming, the boundary nodes of the subdomains match across the interface. A subdomain Ω_i is said to be floating if $\partial\Omega_i \cap \partial\Omega_D = \emptyset$, and nonfloating otherwise.

As in other substructuring methods, the first step of the FETI method consists of eliminating the interior subdomain variables, which results in a Schur complement formulation of our problem. Let $S^{(i)}$ be the Schur complement matrix of Ω_i , and let f_i be the contribution of Ω_i to the load vectors. Let $S = \text{diag}_{i=1}^N S^{(i)}$ be a block-diagonal matrix, and let f be the vector $[f_1, \dots, f_N]$. We denote by u_i the vector of nodal values on $\partial\Omega_i$ and by u the vector $[u_1, \dots, u_N]$.

If Ω_i is a floating subdomain, then $S^{(i)}$ is a singular matrix, and its kernel is generated by a vector Z_i with entries corresponding to the nodes of $\partial\Omega_i$ equal to 1 and entries corresponding to all the other nodes equal to 0. Let Z consist of all the column vectors Z_i . Then $\text{Ker} S = \text{Range} Z$.

Let B be the matrix of constraints which measures the jump of a given vector u across the interface; B will also be referred to as the Lagrange multiplier matrix. Each row of the matrix B is associated with two matching nodes across the interface, and has values 1 and -1 , respectively, at the two nodes, and zero entries everywhere else. A finite element function with corresponding vector values u is continuous if and only if $Bu = 0$.

For a method without redundant constraints and multipliers, the number of pointwise continuity conditions required at crosspoints, i.e., the points that belong to the closure of more than two subdomains, and, therefore, the number of corresponding

rows in the matrix B , is one less than the number of the subdomains meeting at the crosspoint. There exist several different ways of choosing which conditions to enforce at a crosspoint, all of them resulting in algorithms with similar properties.

An alternative suggested in [43, 44] is to connect all the degrees of freedom at the crosspoints by Lagrange multipliers and use a special scaling, resulting in a method with redundant multipliers; see section 4.3 for further details.

Let W_i be the space of the degrees of freedom associated with $\partial\Omega_i \setminus \partial\Omega_D$, and let W be the direct sum of all spaces W_i . If $U = \text{Range}B$ is the space of the Lagrange multipliers, then $S : W \rightarrow W$ and $B : W \rightarrow U$. Let B^t be the transpose of B . By introducing Lagrange multipliers λ for the constraint $Bu = 0$, we obtain a saddle point Schur formulation of (3),

$$(4) \quad \begin{cases} Su + B^t\lambda = f, \\ Bu = 0. \end{cases}$$

3.1. Algebraic formulation. In the FETI method, the primal variable u is eliminated from (4), and the resulting equation for the dual variable λ is solved by a PCG method.

We note that S is singular if there exist at least one floating subdomain among the subdomains Ω_i , $i = 1 : N$. Let S^\dagger be the pseudoinverse of S ; i.e., for any $b \perp \text{Ker}S$, $S^\dagger b$ is the unique solution of $Sx = b$ such that $S^\dagger b \in \text{Range}S$. The first equation in (4) is solvable if and only if

$$(5) \quad f - B^t\lambda \perp \text{Ker}S.$$

If (5) is satisfied, then

$$(6) \quad u = S^\dagger(f - B^t\lambda) + Z\alpha,$$

where $Z\alpha$ is an element of $\text{Ker}S = \text{Range}Z$ to be determined.

Let $G = BZ$. Substituting (6) into the second equation in (4), it follows that

$$(7) \quad BS^\dagger B^t\lambda = BS^\dagger f + G\alpha.$$

An important role in the FETI algorithm is played by V , a subset of U defined by $V = \text{Ker}G^t$. In other words, $V = \text{Ker}G^t \perp \text{Range}G = B\text{Range}Z = B\text{Ker}S$. Let $P = I - G(G^tG)^{-1}G^t$ be the orthogonal projection onto V . It is easy to see that G^tG is nonsingular, by using the fact that $\text{Ker}B \cap \text{Range}Z = \text{Ker}B \cap \text{Ker}S = \emptyset$. Since $P(G\alpha) = 0$, if P is applied to (7), it results that

$$(8) \quad PBS^\dagger B^t\lambda = PBS^\dagger f.$$

We now return to the necessary condition (5). Since $\text{Ker}S = \text{Range}Z$, we obtain that (5) is equivalent to $f - B^t\lambda \perp \text{Range}Z$, which leads to $Z^t(f - B^t\lambda) = 0$ and therefore to

$$(9) \quad G^t\lambda = Z^t f.$$

Let $F = BS^\dagger B^t$, $d = BS^\dagger f$, and $e = Z^t f$. We concluded that we have to solve the dual problem (8) for λ subject to the constraint (9); with the new notations,

$$(10) \quad PF\lambda = Pd,$$

$$(11) \quad G^t\lambda = e.$$

We note that, from (7), it follows that $\alpha = (G^t G)^{-1} G^t (F\lambda - d)$. Therefore, after an approximate solution for λ is found, the primal variable u is obtained from (6) by solving a Neumann or a mixed boundary problem on each floating and nonfloating subdomain, respectively, corresponding to a vector multiplication by S^\dagger .

The main part of the FETI algorithm consists of solving (10) for the dual variable λ , which is done by a PCG method. Since λ must also satisfy the constraint (11), let $\lambda_0 = G(G^t G)^{-1} e$ be the initial approximation. Then $G^t \lambda_0 = e$ and $\lambda - \lambda_0 \in \text{Ker} G^t = V$. If all the increments $\lambda_k - \lambda_{k-1}$, i.e., the search directions, are in V , then (11) will be satisfied.

One possible preconditioner for (10) is of the form PM , where $M = BSB^t$. When a vector multiplication by M is performed, N independent Dirichlet problems have to be solved in each iteration step. Therefore, M is known as the Dirichlet preconditioner. We note that the Schur complement matrix S is never computed explicitly, since only the action of S on a vector is needed.

Mandel and Tezaur [40] have shown that the condition number of this FETI method is bounded from above by $C(1 + \log(H/h))^3$, where C is a positive constant independent of h, H . If there are no crosspoints in the partition of Ω , then this polylogarithmic bound improves to $C(1 + \log(H/h))^2$.

We conclude this section with a few comments on the PCG algorithm:

Projected Preconditioned Conjugate Gradient Iteration (PCG)

$$\lambda_0 = G(G^t G)^{-1} e, r_0 = Pd - PF\lambda_0, n = 1$$

while $(Mr_{n-1}, r_{n-1}) \geq \text{tol}$

$$w_{n-1} = Pr_{n-1}$$

$$z_{n-1} = Mw_{n-1}$$

$$y_{n-1} = Pz_{n-1}$$

$$\beta_n = (y_{n-1}, r_{n-1}) / (y_{n-2}, r_{n-2}) \quad (\beta_1 = 0)$$

$$p_n = y_{n-1} + \beta_n p_{n-1} \quad (p_1 = y_0)$$

$$\alpha_n = (y_{n-1}, r_{n-1}) / (Fp_n, p_n)$$

$$\lambda_n = \lambda_{n-1} + \alpha_n p_n$$

$$r_n = r_{n-1} - \alpha_n PFp_n$$

$$n = n + 1$$

end

In each iteration step of the PCG algorithm, the residual and the search directions are projected onto the space V , i.e., $w_{n-1} = Pr_{n-1}$ and $y_{n-1} = Pz_{n-1}$. This projection step plays the role of a coarse problem which is solved in each iteration and is the reason why the FETI method is numerically scalable, even though it lacks an explicit coarse space construction. We note that $r_{n-1} \in V$ at every step. Therefore, it follows that $w_{n-1} = r_{n-1}$, and thus only one projection onto V is required per iteration step. This observation is particularly important for some of the algorithms suggested in [33].

4. The FETI algorithm for mortars. As we have seen in section 3, in the classical FETI algorithm the computational domain Ω is partitioned into nonoverlapping subregions, multiple degrees of freedom are introduced for the matching nodes across the interface, and pointwise continuity across the interface is enforced by a Lagrange multiplier matrix B . This methodology is very similar to that used in [5], where a saddle point formulation for the mortar finite element method has been introduced.

In fact, the FETI method can be applied without any algorithmic changes for a mortar finite element discretization of Ω , using the nonoverlapping partition $\{\Omega_i\}_{i=1:N}$ considered in section 2. We recall that this partition may be geometrically nonconforming and the nodes across the interface do not necessarily match. To keep the

presentation clear, we assume that each subregion Ω_i has a diameter of order H and that its triangulation has a mesh size of order h . The matrix S is again a block-diagonal matrix $\text{diag}_{i=1}^N S^{(i)}$, where the local Schur complement matrices $S^{(i)}$ are obtained from the finite element discretizations on individual subregions. We have to solve the problem

$$\begin{cases} Su + B^t \lambda = f, \\ Bu = 0, \end{cases}$$

where the matrix B enforces mortar conditions across the interface. The dual problem is obtained as in section 3.1. It results in solving

$$(12) \quad PF\lambda = Pd,$$

with a PCG method, with initial approximation $\lambda_0 = G(G^tG)^{-1}e$, and with all the search directions in V .

The price we pay for the inherent flexibility of the mortar finite elements is due to the fact that the matrix B is more complicated in the mortar case, compared to that of the classical FETI method with conforming finite elements. The matrix B has one block, B_γ , for each nonmortar side γ . We adopt the matrix formulation of the mortar conditions from section 2.1. Let M_γ and N_γ be the matrices which multiply the nonmortar and mortar nodal values in the mortar conditions across γ , respectively. Then B_γ consists of the columns of M_γ and $-N_\gamma$ for the nodes of γ and those on the mortars opposite γ , and has zero columns corresponding to all the other nodes.

We note that the mortar conditions are all associated with the interior nodes on the nonmortar sides. Therefore, the problem of choosing the crosspoints constraints does not arise in the mortar case.

In our numerical experiments, we have implemented three different preconditioners suggested in the FETI literature for the dual problem (12). In sections 4.1–4.3, we present each of them briefly.

4.1. The Dirichlet preconditioner. In [26], Farhat and Roux introduced the Dirichlet preconditioner for the FETI method

$$(13) \quad PM = P B S B^t.$$

This preconditioner was shown to perform well for conforming finite elements; see, e.g., [26] for numerical results and [40] for condition number estimates.

4.2. A block-diagonal preconditioner. In [34, 35], Lacour suggested another preconditioner designed specifically for a mortar version of the FETI algorithm and without a counterpart in the conforming case. Let $\text{diag} B_\gamma B_\gamma^t$ be the block-diagonal matrix which has a block $B_\gamma B_\gamma^t$ of size equal the number of interior nodes on γ for each nonmortar γ . We note that $\text{diag} B_\gamma B_\gamma^t$ is the block-diagonal part of the matrix $B B^t$. In the three dimensional case, each block corresponds to a nonmortar face \mathcal{F} , and the block-diagonal matrix is $\text{diag} B_\mathcal{F} B_\mathcal{F}^t$. To simplify our presentation, we will use the same notation, $\text{diag} B_\gamma B_\gamma^t$, for the three dimensional block-diagonal matrix.

The preconditioner $P\bar{M}$ is defined as follows:

$$(14) \quad P\bar{M} = P(\text{diag} B_\gamma B_\gamma^t)^{-1} B S B^t (\text{diag} B_\gamma B_\gamma^t)^{-1}.$$

4.3. A new preconditioner. In [33], Klawonn and Widlund studied a FETI method for elliptic problems with heterogeneous coefficients, discretized by conforming finite elements, and designed a new preconditioner for these types of problems. They used this preconditioner to show the connection between FETI methods and Neumann–Neumann methods, in particular the balancing method [39].

In the case of no coefficient jump, as in our Poisson problem, the new preconditioner has the form

$$(15) \quad P\widehat{M} = P(BB^t)^{-1}BSB^t(BB^t)^{-1}.$$

Klawonn and Widlund established the following upper bound for the condition number of their FETI algorithms, which is valid for all cases, including when the partition contains crosspoints

$$\kappa(P\widehat{M}PF) \leq C(1 + \log(H/h))^2.$$

In the same paper, it is proven that the preconditioner \widehat{M} with a minimal number of pointwise continuity conditions at the crosspoints, and therefore of Lagrange multipliers, results in a similar algorithm as the FETI method with redundant Lagrange multipliers of Rixen and Farhat [43, 44]. Since the Lagrange multipliers in the mortar case are not associated with the vertices of the subregions, a FETI algorithm with redundant multipliers cannot be implemented for mortars.

To use the new preconditioner of Klawonn and Widlund for the FETI method with mortars, we must show that the matrix BB^t is nonsingular in the mortar case. The number of columns of B is equal to the number of nodes from W , while the number of rows of B is equal to the number of Lagrange multipliers. Since each Lagrange multiplier is associated with an interior node on a nonmortar side, it results that B has fewer rows than columns. Therefore, if we show that the rank of B is equal to its number of rows, we may conclude that BB^t is nonsingular. We consider the minor of B consisting of the columns corresponding to the interior nodes of the nonmortars. The resulting block-diagonal square matrix $\text{diag}M_\gamma$ is nonsingular, since each block M_γ is a diagonally dominant matrix for the classical mortar elements, and the identity matrix for the new mortar elements.

We conclude this section with comments on the difficulties of extending the convergence analysis of the FETI method from conforming finite elements to mortar elements.

The numerical evidence presented in sections 5.1 and 6.1 strongly suggested that, in the mortar case, the new preconditioner performed as well as it did in the conforming finite element case. Extending the analysis of Klawonn and Widlund to the mortar case would require establishing the bound

$$(16) \quad |B^t(BB^t)^{-1}Bw|_S^2 \leq C(1 + \log(H/h))^2 |w|_S^2 \quad \forall w \in W, \quad Bw \in V.$$

In the conforming finite element case, the matrix BB^t is equal to twice the identity matrix, except for the entries corresponding to crosspoint conditions, where it has a block-diagonal structure. The inverse matrix $(BB^t)^{-1}$ has the same block-diagonal structure, and (16) can be obtained by using Sobolev-type inequalities for finite element functions (also known as cutoff estimates).

In the mortar finite element case, the matrix BB^t is not close to a multiple of the identity and is no longer block diagonal. There are two types of entries outside the diagonal blocks. Some correspond to Lagrange multipliers associated with the first and

last interior points of the nonmortars. Others occur because there exist nodal basis functions associated with points on the mortar sides, the support of which intersects more than one nonmortar; see Figure 4 for the sparsity pattern of BB^t .

An important consequence of such structure of BB^t is that the matrix $(BB^t)^{-1}$ is a full matrix in the mortar case. Therefore, cutoff estimates can no longer be used to prove (16). An alternative solution would require new tools in order to reduce the inequality (16) to the two subdomain case. This case would then be solved by deriving some specific stability estimates for the mortar projection.

5. Numerical results for two dimensional problems. In this section, we present numerical results for the FETI method for a mortar finite element discretization of a two dimensional problem. We have tested each of the three preconditioners of sections 4.1–4.3 on nonconforming discretizations of the computational domain.

Our interests were three fold:

- to compare the convergence performances of the different FETI preconditioners for mortar methods, based on iteration counts and estimates for the condition numbers;
- to apply the FETI algorithms for the new mortar finite elements and compare the iteration counts and the flop counts to those obtained for the classical mortar finite elements;
- to analyze the extra computational effort, due to the complexity of the mortar conditions, required for the implementation of the FETI algorithm with the new preconditioner.

As the model problem in two dimensions, we chose the Poisson equation on the unit square $\Omega = [0, 1]^2$ with zero Dirichlet boundary conditions. The right-hand side function f in (3) was selected such that the exact solution of the problem is known.

The computational domain Ω was partitioned into 16, 32, 64, and 128 geometrically nonconforming rectangular subregions, respectively; see Figure 2. On each subregion, we considered Q_1 elements of mesh size h , and, to make the comparisons easier, all the subregions had diameters of the same order, H . For each partition, the number of nodes on each edge, H/h , was taken to be, *on average*, 4, 8, 16, and 32, respectively, for different sets of experiments. Across the partition interface Γ the meshes did not necessarily match. A saddle point formulation of the problem was used, and mortar conditions were enforced across Γ .

We report the iteration counts and the flop counts of the algorithms. We also report condition number estimates for the algorithms, obtained directly from the conjugate gradient iteration. The PCG iteration was stopped when the residual norm had decreased by a factor of 10^{-6} . All the experiments were carried out in MATLAB.

We now present some implementation details. We did not compute the Schur complements explicitly, nor their pseudoinverses, but only the stiffness matrices for each subdomain. To multiply a vector by a Schur complement matrix, we solved, in each subregion, a Poisson problem with Dirichlet boundary conditions. To multiply a vector by S^\dagger , we solved a Poisson problem with mixed boundary conditions in each nonfloating subregion and with Neumann boundary conditions in each floating subregion; see, e.g., [17]. We stored only the interior-boundary and boundary-boundary blocks of the local stiffness matrix and the Cholesky factor of the interior-interior block, which is symmetric and positive definite. To have a uniquely solvable problem on the floating subregions, we required that the solution of the local Neumann problem be orthogonal to $\text{Ker}S$, i.e., to the constant functions on the subregion. A simple way of enforcing this orthogonality condition was by adding a Lagrange multiplier

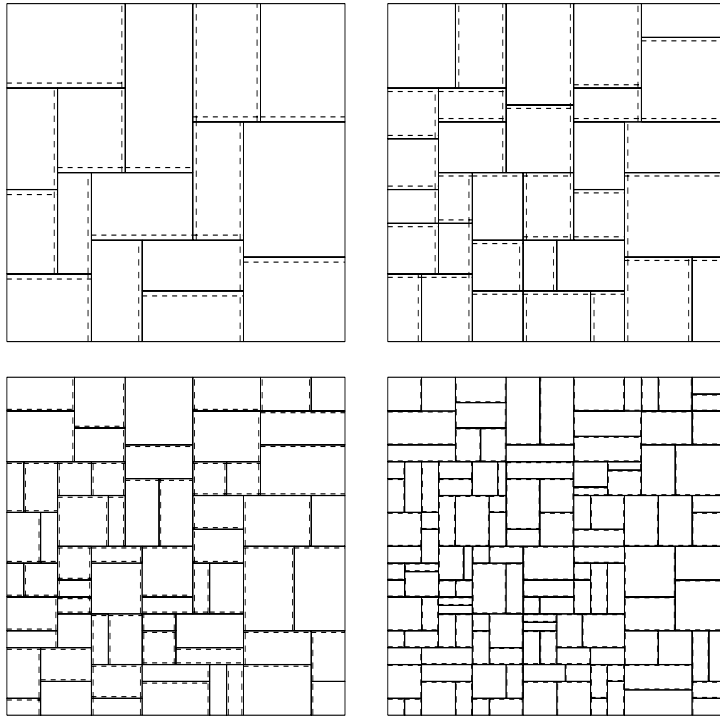


FIG. 2. Geometrically nonconforming partitions of Ω . Upper left: 16 subdomains; upper right: 32 subdomains; lower left: 64 subdomains; lower right: 128 subdomains.

and storing the LU components of the extended stiffness matrix.

5.1. Convergence properties of the FETI algorithms. We now turn to the main part of this section, a discussion of the performance of the FETI algorithms for mortar finite elements with the new preconditioner \widehat{M} , (15), the preconditioner \overline{M} , (14), and the Dirichlet preconditioner M , (13). In Table 1 we report the iteration count, the condition number approximation, and the flop count for the aforementioned preconditioners.

The FETI algorithm with the Dirichlet preconditioner M required hundreds of iterations to converge, and the computational costs were one to two orders of magnitude bigger than for the other preconditioners. The iteration count grew faster than polylogarithmically as a function of the number of nodes on each subdomain edge, H/h , and appeared to grow linearly with the number of subdomains. The Dirichlet preconditioner is therefore noncompetitive, since many domain decomposition methods have convergence rates independent of the number of subdomains. The condition numbers estimates were on the order of 10^4 – 10^6 , unusually large for these types of algorithms. Moreover, our estimates are likely to be smaller than the actual condition numbers, since there was no convincing convergence pattern for the condition number approximation obtained in the iteration; see section 6.1 and Figure 3 therein for more details. Thus, unlike in the case of FETI algorithms with conforming finite elements, the Dirichlet preconditioner M did not yield a numerically scalable method for mortar finite element methods.

The new preconditioner \widehat{M} scaled similarly to M in the conforming case. When

TABLE 1

Convergence results; two dimensional geometrically nonconforming partition; classical mortar elements.

N	H/h	New precondition.			Block-diag. precondition.			Dirichlet precondition.		
		Iter	Cond	Mflops	Iter	Cond	Mflops	Iter	Cond	Mflops
16	4	10	4.14	6.9e-1	21	26.95	1.4e+0	111	7.3e+3	7.2e+0
16	8	12	5.14	8.2e+0	21	29.86	1.4e+1	240	4.2e+4	1.6e+2
16	16	13	6.44	1.5e+2	23	36.53	2.6e+2	320	6.5e+4	3.7e+3
16	32	14	7.35	3.4e+3	23	38.03	5.6e+3	348	6.8e+4	7.4e+4
32	4	11	6.53	1.9e+0	23	34.51	3.7e+0	223	1.2e+4	3.5e+1
32	8	13	7.58	1.9e+1	24	45.96	3.5e+1	455	7.1e+4	6.6e+2
32	16	14	8.86	3.5e+2	26	61.97	6.5e+2	528	1.0e+5	1.3e+4
32	32	16	9.79	9.1e+3	27	65.39	1.5e+4	569	1.2e+5	3.2e+5
64	4	14	7.23	6.1e+0	32	47.99	1.3e+1	578	9.1e+4	2.2e+2
64	8	16	8.76	5.7e+1	35	72.62	1.2e+2	1012	7.5e+5	3.5e+3
64	16	18	10.68	1.0e+3	36	91.43	2.0e+3	1266	1.2e+6	7.1e+4
64	32	20	12.40	2.5e+4	39	94.47	4.8e+4	1324	1.4e+6	1.5e+6
128	4	14	7.60	1.3e+1	36	64.53	3.0e+1	1144	9.2e+5	9.2e+2
128	8	17	9.56	1.3e+2	40	82.09	2.9e+2	1350	6.8e+5	1.0e+4
128	16	19	11.73	2.3e+3	41	96.60	4.8e+3	1436	9.9e+6	1.7e+5
128	32	21	13.03	5.7e+4	41	99.82	1.1e+5	-	-	-

the number of nodes on each subdomain edge, H/h , was fixed and the number of subdomains, N , was increased, the iteration count showed only a slight growth. When H/h was increased, while the partition was kept unchanged, the increase in the number of iterations was quite satisfactory and very similar to that of the conforming case. The condition number estimates exhibited a similar dependence on the number of subdomains and on the number of nodes on each subdomain edge.

The block-diagonal preconditioner \bar{M} had good convergence properties as well. The iteration counts showed just a small increase when the number of nodes on each subdomain edge was increased, while the partition was kept unchanged. There seemed to be a stronger than desired dependence of the iteration counts on the number of subdomains, which was less than optimal. The condition number estimates followed a similar pattern but were significantly larger than the condition number estimates for the new preconditioner \tilde{M} .

Overall, the block-diagonal preconditioner \bar{M} required about twice as many iterations to convergence and twice as much computational effort as \tilde{M} ; cf. Table 1. Therefore, even though multiplying \bar{M} by a vector required less computational effort than when \tilde{M} was used, the increase in the iteration count resulted in flop counts which were twice as large. This suggests that dropping the nonzero diagonal terms of BB^t relaxed the weak continuity conditions for mortar finite elements more than is optimal.

We conclude that, among the three preconditioners for FETI algorithms for mortar finite element methods analyzed here, the new preconditioner \tilde{M} is the best.

5.2. New mortars vs. classical mortars. Another objective of our study was to compare the performance of the FETI algorithms for new mortar element methods with the performance of the FETI algorithms for classical mortar element methods.

We used the same nonconforming partitions of our computational domain (see Figure 2) and considered new mortar finite elements on the subdomains. As explained in section 2.1, the only difference between the two mortars methods is due to different mortar conditions across the interface Γ . This results in different Lagrange multiplier

TABLE 2

Convergence results; two dimensional geometrically nonconforming partition; new mortar elements.

N	H/h	New precondition.			Block-diag. precondition.			Dirichlet precondition.		
		Iter	Cond	Mflops	Iter	Cond	Mflops	Iter	Cond	Mflops
16	4	10	4.20	6.8e-1	18	18.40	1.2e+0	55	897	3.5e+0
16	8	12	5.15	8.1e+0	19	20.85	1.3e+1	70	908	4.7e+1
16	16	13	6.50	1.5e+2	20	25.40	2.3e+2	70	1018	8.0e+2
16	32	14	7.43	3.4e+3	21	30.11	5.1e+3	80	1090	2.0e+5
32	4	11	6.55	1.9e+0	19	23.58	3.0e+0	91	913	1.4e+1
32	8	13	7.61	1.9e+1	20	34.45	2.9e+1	101	1565	1.4e+2
32	16	14	8.87	3.5e+2	22	48.14	5.5e+2	114	1873	2.9e+3
32	32	16	9.80	9.1e+3	23	50.63	1.3e+4	117	2108	6.6e+4
64	4	14	7.29	6.0e+0	27	32.49	1.0e+1	278	1.0e+4	1.1e+2
64	8	16	8.69	5.6e+1	28	43.37	9.6e+1	292	1.6e+4	1.0e+3
64	16	18	10.83	1.0e+3	29	56.19	1.6e+3	297	1.9e+4	1.7e+4
64	32	20	12.57	2.5e+4	31	70.74	3.9e+4	312	2.2e+4	3.1e+5
128	4	14	7.60	1.3e+1	30	41.78	2.4e+1	614	1.0e+5	4.9e+2
128	8	17	9.57	1.3e+2	33	52.36	2.4e+2	677	1.3e+5	4.8e+3
128	16	19	11.75	2.2e+3	35	62.54	4.1e+3	755	1.6e+5	8.9e+4
128	32	21	13.07	5.6e+4	36	66.66	9.7e+4	–	–	–

matrices B .

We ran the same set of experiments for the new mortars discretization, for preconditioners \widehat{M} , \overline{M} , and M . The PCG iteration was stopped when the residual norm had decreased by a factor of 10^{-6} . We report the iteration count, the condition number approximation, and the flop counts in Table 2.

Due to the inherent simplification of the Lagrange multiplier matrix B for the new mortar constraints, we expected the results for the new mortar method to be similar but somewhat better than those for the classical mortar method. Indeed, this was confirmed by our numerical results.

The iteration counts for the new preconditioner \widehat{M} were identical to those for \widehat{M} for classical mortars. The condition number estimates and the flop counts were slightly smaller than in the classical mortar case but essentially the same. Therefore, \widehat{M} scaled just as well as in the classical mortar case.

For the new mortar conditions, the matrix BB^t had fewer nonzero entries outside its block-diagonal structure and fewer terms to be dropped in order to obtain $\text{diag} B_\gamma B_\gamma^t$. Therefore, the block-diagonal preconditioner \overline{M} was closer to \widehat{M} than in the classical mortar case. This resulted in lower iteration counts and condition numbers for \overline{M} than in the classical mortar case, and, consequently, in lower flop counts.

Once again, the iteration counts increased moderately with H/h and seemed to have a stronger than desired dependence on the number of subdomains. Overall, the block-diagonal preconditioner still performed worse than \widehat{M} for the new mortars with iteration counts one and a half times higher and with significantly bigger condition numbers.

An even greater improvement over the classical mortar case was obtained for the Dirichlet preconditioner M . The number of iterations required for M in the new mortar case was less than half the number of iterations required in the classical mortar case, and a similar improvement can be observed for the flop counts. The condition number estimates were about an order of magnitude less than for the classical mortar case.

Despite these improvements, the FETI algorithm with the Dirichlet preconditioner still required hundreds of iterations to converge and the iteration count appeared to grow linearly with the number of subdomains. The condition number estimates were on the order of 10^3 – 10^5 , much higher than desired. Therefore, as in the classical mortar case, applying the preconditioner M for the FETI method did not result in a scalable algorithm.

Once again, among the three preconditioners for FETI algorithms methods, the new preconditioner \widehat{M} was the best. There was no significant improvement when using the new mortar elements instead of the classical mortar elements for the FETI algorithm with optimal preconditioner \widehat{M} .

5.3. Complexity study of the preconditioners. The last topic of this section is an analysis of how expensive it is to apply the preconditioners \widehat{M} and \overline{M} , compared to applying the Dirichlet preconditioner M .

In each iteration step, we compute one vector multiplication by the preconditioner, which requires solving two systems with the matrix BB^t , and $\text{diag}B_\gamma B_\gamma^t$, respectively. In section 4.2, we mentioned that $\text{diag}B_\gamma B_\gamma^t$ is obtained from BB^t by eliminating the nonzero entries outside the diagonal blocks. There are two such types of entries; see section 4.3 for more details. While in the three dimensional case the sparsity pattern of BB^t is complex, in the two dimensional case there are relatively few nonzero entries; see Figure 4.

It is easy to see that $\text{diag}B_\gamma B_\gamma^t$ has a bandwidth of order H/h , the number of interior nodes on an arbitrary nonmortar. The matrix BB^t is also banded, but in this case the band depends on the ordering of the nodes on the interface, and it is possible to have bandwidth of order $1/h$. Therefore, multiplying a vector by $(BB^t)^{-1}$ is potentially an expensive operation. To minimize the effect of these vector multiplications, we computed the Cholesky factorizations of BB^t and $\text{diag}B_\gamma B_\gamma^t$ just once and stored the factors. Then, solving systems with BB^t or $\text{diag}B_\gamma B_\gamma^t$ amounted only to one back and one forward solve.

Our results showed that the costs of a vector multiplication by $(BB^t)^{-1}$ were between two and ten times larger than those associated with $(\text{diag}B_\gamma B_\gamma^t)^{-1}$. However, due to the sparsity pattern of BB^t , even the costs associated with $(BB^t)^{-1}$ were relatively small compared to those for other operations performed during one iteration, e.g., multiplying a vector by the Schur complement or by its pseudoinverse; cf. Table 3. These low relative costs result in very similar flop counts per iteration step for the two preconditioners, almost identical for the case of many nodes per subdomain edge, i.e., $H/h = 16$ and $H/h = 32$.

As expected, the costs associated with $(BB^t)^{-1}$ in each iteration step decreased significantly, from six percent to less than .05 percent, when the partition was fixed and H/h increased, since the costs of multiplying S and S^\dagger by a vector rose much faster than those corresponding to $(BB^t)^{-1}$.

From the flop counts reported in Table 1, it is clear that the improvement of the iteration count easily offsets the small extra costs due to the complexity of \widehat{M} and \overline{M} .

6. Numerical results for three dimensional problems. In this section, we report numerical results for the FETI method for mortar finite element discretizations of a three dimensional problem. As before, we compare the performance of different FETI preconditioners and discuss the effects of using the new mortar finite elements instead of the classical ones. We include a study of the costs of applying the new preconditioner for the for classical mortar elements and more details on the convergence rate of the condition number approximation for some of our algorithms.

TABLE 3

Complexity study of one iteration step for the new preconditioner and the block-diagonal preconditioner; two dimensional geometrically nonconforming partition.

		New preconditioner			Block-diagonal preconditioner		
N	H/h	Mflops for $(BB^t)^{-1}$	Mflops per iteration	Ratio	Mflops for $(diag B_\gamma B_\gamma^t)^{-1}$	Mflops per iteration	Ratio
16	4	3.9e-3	6.9e-2	.06	1.4e-3	6.7e-2	.02
16	8	9.7e-3	6.8e-1	.02	4.6e-3	6.7e-1	.007
16	16	2.2e-2	1.1e+1	.002	1.1e-2	1.1e+1	.001
16	32	4.6e-2	2.4e+2	.0002	2.4e-2	2.4e+2	.0001
32	4	1.1e-2	1.7e-1	.07	3.2e-3	1.6e-1	.02
32	8	3.1e-2	1.5e+0	.02	9.9e-3	1.5e+0	.007
32	16	6.5e-2	2.5e+1	.003	2.4e-2	2.5e+1	.001
32	32	1.4e-1	5.7e+2	.0003	5.2e-2	5.7e+2	.00009
64	4	4.1e-2	4.3e-1	.10	6.9e-3	3.9e-1	.02
64	8	1.1e-1	3.5e+0	.03	2.2e-2	3.4e+0	.007
64	16	2.3e-1	5.6e+1	.004	5.4e-2	5.6e+1	.001
64	32	4.7e-1	1.2e+3	.0004	1.2e-1	1.2e+3	.00009
128	4	1.2e-1	9.3e-1	.13	1.4e-2	8.2e-1	.02
128	8	3.1e-1	7.4e+0	.04	4.4e-2	7.2e+0	.006
128	16	6.8e-1	1.2e+2	.006	1.1e-1	1.2e+2	.0009
128	32	1.4e+0	2.7e+3	.0005	2.3e-1	2.7e+3	.00009

As the model problem in three dimensions, we chose the Poisson equation on the unit cube $\Omega = [0, 1]^3$ with zero Dirichlet boundary conditions. The right-hand side was selected such that the exact solution is known. The computational domain Ω was partitioned into 8, 16, and 32 nonconforming parallelepipeds, respectively. We chose these partitions such that in each case there exist floating subdomains, i.e., interior subdomains.

The subdomains of the partition had diameter of order H , and Q_1 elements of mesh size h were used in each subdomain. The number of nodes on each edge was, *on average*, 4, 8, and 16. Across the partition interface Γ the meshes did not match, and mortar conditions for three dimensional elements were enforced; cf. section 2.2. This results in a Lagrange multipliers matrix B , which, as explained for the two dimensional case, plays a very important role in all FETI algorithms.

We report the iteration counts, the condition number approximations, and the flop counts of the algorithms. The PCG iteration was stopped when the residual norm had decreased by a factor of 10^{-6} . All the experiments were carried out in MATLAB.

6.1. Convergence properties of the FETI algorithms. We did not compute the Schur complements explicitly but only stored those components of the stiffness matrices which were relevant for the multiplication of a vector by the Schur complement matrix and by the pseudoinverse of the Schur complement. We tested the performance of the same preconditioners as in the two dimensional case, i.e., the new preconditioner \widehat{M} , cf. (15); the preconditioner \overline{M} , cf. (14); and the Dirichlet preconditioner M , cf. (13). We report the iteration count, the condition number estimate, and the flop count of the algorithms in Table 4.

As in the two dimensional case, the FETI algorithm with the Dirichlet preconditioner M did not scale well and required thousands of iterations to converge. Since it soon became clear that M was not an optimal preconditioner, and due to significant computational costs, we performed only tests for every partition of Ω in the case of 4 nodes on each edge, and for the 8 and the 16 subdomains partitions for the case of 8

TABLE 4

Convergence results; three dimensional geometrically nonconforming partition; classical mortar elements.

N	H/h	New precondition.			Block-diag. precondition.			Dirichlet precondition.		
		Iter	Cond	Mflops	Iter	Cond	Mflops	Iter	Cond	Mflops
8	4	11	4.31	1.8e+0	33	55	3.9e+0	866	2.2e+6	9.9e+1
8	8	14	6.54	4.9e+1	33	70	7.3e+1	7984	4.4e+7	1.7e+4
8	16	16	7.90	1.4e+3	38	81	2.4e+3	–	–	–
16	4	13	6.77	6.6e+0	36	75	9.7e+0	2985	1.2e+7	7.7e+2
16	8	15	8.20	1.6e+2	37	85	1.7e+2	14169	2.7e+8	6.3e+4
16	16	17	9.28	4.1e+3	50	176	6.5e+3	–	–	–
32	4	14	8.29	3.1e+1	44	141	2.7e+1	4156	3.5e+7	2.4e+3
32	8	16	9.77	1.1e+3	55	350	5.5e+2	–	–	–
32	16	17	10.69	2.7e+4	69	523	2.9e+4	–	–	–

nodes on each edge of the subdomains.

The iteration count seemed to be a linear function of the number of subdomains, and it grew by an order of magnitude when H/h was doubled while keeping the 8 and the 16 subdomain partitions fixed. The computational costs were also at least two orders of magnitude bigger than those for the other preconditioners and deteriorated as the number of nodes per subdomain edge increased.

The condition number estimates followed a similar dependence pattern on H/h and the number of subdomains. They were on the order of 10^6 – 10^8 , much worse than even in the two dimensional case. In Figure 3, we present the convergence pattern of the condition number estimates for the eight subdomains partition with $H/h = 4$. For M , the PCG iteration was stopped when the residual norm had decreased by a tolerance factor of 10^{-6} , while for \widehat{M} and \overline{M} the tolerance was set at 10^{-10} .

For the Dirichlet preconditioner, there was no clear convergence pattern for the condition number estimates. This suggests that the (extremely large) condition number approximations reported in Table 4 are just lower bounds for the actual condition number. For the other two preconditioners, convergence was achieved early in the iteration count. The estimates reported in Table 4 are within one percent of the condition number corresponding to a tolerance of 10^{-10} .

The new preconditioner \widehat{M} scaled similarly to the two dimensional case and to the Dirichlet preconditioner in the conforming case. The number of iterations grew very slowly when the number of nodes on each subdomain edge (i.e., H/h) was fixed and the number of subdomains was increased. When the partition was kept unchanged and H/h was increased, the iteration count increased slightly, and it seemed to have a polylogarithmic dependence on H/h .

The convergence analysis for the block-diagonal preconditioner \overline{M} is particularly interesting in the three dimensional case; \overline{M} was a possible alternative to \widehat{M} since it required significantly less computational effort per iteration step. However, our results showed a much stronger than desired dependence of the iteration count for \overline{M} on the number of nodes on each subdomain edge. This dependence grew stronger as the number of subdomains in the partition increased. The number of iterations increased with the number of subdomains, another undesirable property. The condition number estimates followed a similar pattern and were significantly larger than those corresponding to \widehat{M} . Their relatively large values, on the order of 10^2 , and their dependence on the number of subdomains were unsatisfactory.