SCHRIFTENREIHE DES FACHBEREICHS MATHEMATIK

Inexact FETI-DP Methods

by

Axel Klawonn and Oliver Rheinbach

SM-E-609 2005

Universität Duisburg-Essen

Eingegangen am 28.07.2005

INEXACT FETI-DP METHODS

AXEL KLAWONN* AND OLIVER RHEINBACH*

July 2005

Abstract. Inexact FETI-DP domain decomposition methods are considered. Preconditioners based on formulations of FETI-DP as a saddle point problem are used which allow for an inexact solution of the coarse problem. A positive definite reformulation of the preconditioned saddle point problem, which also allows for approximate solvers, is considered as well. In the formulation that iterates on the original FETI-DP saddle point system, it is also possible to solve the local Neumann subdomain problems inexactly. Given good approximate solvers for the local and coarse problems, convergence bounds of the same quality as for the standard FETI-DP methods are obtained. Numerical experiments which compare the performance of the method to FETI-DP are shown for 2D and 3D elasticity using GMRES and CG as Krylov space methods.

Key words. domain decomposition, Lagrange multipliers, FETI, preconditioners, elliptic systems, elasticity, finite elements, parallel computing, inexact.

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

1. Introduction. Domain decomposition methods are preconditioned iterative algorithms for the solution of the large systems obtained from the discretization of partial differential equations. In domain decomposition methods, the domain associated with the partial differential equation is decomposed into a, possibly large, number of subdomains. On these subdomains, local problems are defined which are solved in each iteration step in order to define an approximate inverse of the system matrix. In order to obtain a numerical and parallel scalable algorithm, also a small coarse problem has to be introduced and solved in each iteration step.

In this article, we consider nonoverlapping domain decomposition methods belonging to the family of Dual-Primal Finite Element Tearing and Interconnecting (FETI-DP) methods; see [10, 20, 16]. In FETI-DP methods the continuity of the solution across the subdomain boundaries is enforced by Lagrange multipliers. This results in a mixed linear system with primal variables and Lagrange multipliers as unknowns. The basic idea of FETI-DP domain decomposition methods is to eliminate the primal variables and iterate, usually in combination with a preconditioner, on the resulting Schur complement and the Lagrange multiplier variables. Special attention has to be given to the elimination process of the primal variables since the associated matrix usually is only semidefinite although the overall mixed linear system is uniquely solvable; this is due to local stiffness matrices belonging to subdomains lacking sufficient essential boundary conditions. In FETI-DP methods, a sufficient number of constraints, e.g., continuity across the interface at selected nodes on the subdomain boundaries, is chosen such that the local stiffness matrices become invertible; we note that nodal constraints are only sufficient in two dimensions, more elaborate choices, e.g., averages over edges, are used in three dimensions in order to obtain a good convergence estimate. These additional primal constraints introduce a certain coupling between the otherwise completely decoupled local subdomain problems but this coupling also builds our coarse problem needed for scalability of the algorithm. In standard FETI-DP methods, an inexact solution of the coarse problem is not straightforward, since the coarse problem is, by means of the elimination

*Fachbereich Mathematik, Universität Duisburg-Essen, Universitätsstr. 3, 45117 Essen, Germany. E-mail: {axel.klawonn, oliver.rheinbach}@uni-essen.de, URL: http://www.uni-essen.de/numa.

process, built into the FETI-DP system matrix. Thus, an inexact solution in the elimination process of the primal variables, would lead to a different linear system to be solved and thus to a perturbed solution, different to that of the original problem. FETI-DP methods obtain their numerical and parallel scalability from the use of a coarse problem which is very small compared to the overall problem. This coarse problem traditionally is solved exactly by the use of a direct solver. Nevertheless, if a very large number of subdomains is used or if the problem requires the use of a larger coarse space the cost of solving the coarse problem directly may become high.

Here, we present a family of methods, based on different versions of the FETI-DP saddle point formulation, which allow for the use of inexact solvers for the FETI-DP coarse problem. In one variant, which iterates on the complete saddle point system, it is also possible to solve the local Neumann problems inexactly; see also Klawonn and Widlund [18] for a related approach for the classical, onelevel FETI method; we note that in the approach in [18], the coarse problem has to be solved exactly. The resulting preconditioned systems in our article are based on appropriate block triangular preconditioners and can either be solved by GMRES or by any other Krylov space method suitable for nonsymmetric linear systems, e.g., BiCGSTAB or QMR. We also discuss a positive definite reformulation which can be solved by the method of conjugate gradients. This approach dates back to work on preconditioners for saddle point problems by Bramble and Pasciak [4]; see also Klawonn [15] and Dohrmann and Lehoucq [8]. The inexact solution of the local Dirichlet problems is always possible with FETI methods, e.g., by the use of the non optimal, lumped preconditioner. Let us note that the algorithms presented here also allow for optimal, inexact Dirichlet subdomain solvers.

Another class of nonoverlapping domain decomposition methods which is closely related to the FETI-DP algorithms, are the Balancing Domain Decomposition methods by Constraints (BDDC); see Cros [6], Dohrmann [9], Mandel and Dohrmann [24], Mandel, Dohrmann, and Tezaur [25], or Li and Widlund [22]. An approach to solve the BDDC coarse problem inexactly has successfully been suggested and analyzed by Tu [35, 36]. Such an approach is more straightforward for BDDC methods, since the coarse problem is built into the preconditioner and not into the system matrix. For very recent work on approximate subdomain solvers for BDDC methods, see Li and Widlund [23]. In Gosselet [12] a hybrid domain decomposition method is considered and applied to multi-field problems that iterates on a linear system consisting of primal and dual variables at the same time.

The remainder of this paper is organized as follows. In Section 2, we introduce as an elliptic model problem the system of linear elasticity. In Section 3, we introduce the FETI-DP method and derive it from a saddle point formulation which is the basis for our inexact FETI-DP methods. In Section 4, we present the different, exact and inexact, FETI-DP preconditioners and in Section 5, we review a general convergence theory for block triangular preconditioners applied to symmetric saddle point problems. In Section 6, we apply that theory to our inexact preconditioners, which are of block triangular form. In Section 7, we discuss some performance considerations and in Section 8, we present numerical results for our inexact FETI-DP methods. In Section 9, we conclude this paper with some annotations on solvers and hardware used in our work.

2. Model problem. In this section, we consider the system of linear elasticity as a model problem. We note that other elliptic partial differential equations could be treated as well using the methods provided in this paper.

The equations of linear elasticity model the displacement of a linear elastic material under the action of external and internal forces. The elastic body occupies a domain $\Omega \subset \mathbb{R}^d, d = 2, 3$, which is assumed to be polygonal or polyhedral, respectively. We denote its boundary by $\partial\Omega$ and assume that one part of it, $\partial\Omega_D$, is clamped, i.e., with homogeneous Dirichlet boundary conditions, and that the rest, $\partial\Omega_N := \partial\Omega \backslash \partial\Omega_D$, is subject to a surface force g, i.e., a natural boundary condition. We can also introduce a body force f, e.g., gravity. The appropriate space for a variational formulation is the Sobolev space $H^1_0(\Omega, \partial\Omega_D) := \{v \in (H^1(\Omega))^d : v = 0 \text{ on } \partial\Omega_D\}$. The linear elasticity problem consists in finding the displacement $u \in H^1_0(\Omega, \partial\Omega_D)$ of the elastic body Ω, such that

$$\int_\Omega G(x)\varepsilon(u) : \varepsilon(v)dx + \int_\Omega G(x)\,\beta(x)\,\mathrm{div}u\,\mathrm{div}v\,dx = \langle F, v\rangle \;\; \forall v \in H^1_0(\Omega, \partial\Omega_D). \quad (2.1)$$

Here G and β are material parameters which depend on the Young modulus $E > 0$ and the Poisson ratio $\nu \in (0, 1/2]$; we have $G = E/(1 + \nu)$ and $\beta = \nu/(1 - 2\nu)$. In this article, we only consider the case of compressible elasticity, which means that the Poisson ratio ν is bounded away from $1/2$. Furthermore, $\varepsilon_{ij}(u) := \frac{1}{2}(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i})$ is the linearized strain tensor, and

$$\varepsilon(u) : \varepsilon(v) = \sum_{i,j=1}^{3} \varepsilon_{ij}(u)\varepsilon_{ij}(v), \quad \langle F, v\rangle := \int_\Omega f^T v\,dx \;+\; \int_{\partial\Omega_N} g^T v\,d\sigma.$$

For convenience, we also introduce the notation

$$(\varepsilon(u), \varepsilon(v))_{L_2(\Omega)} := \int_\Omega \varepsilon(u) : \varepsilon(v)dx.$$

The bilinear form associated with linear elasticity is then

$$a(u, v) = (G\,\varepsilon(u), \varepsilon(v))_{L_2(\Omega)} \;+\; (G\,\beta\,\mathrm{div}u, \mathrm{div}v)_{L_2(\Omega)}.$$

The wellposedness of the linear system (2.1) follows immediately from the continuity and ellipticity of the bilinear form $a(\cdot, \cdot)$, where the first follows from elementary inequalities and the latter from Korn's first inequality; see, e.g., [5].

We will only consider compressible elastic materials. It is therefore sufficient to discretize our elliptic problem of linear elasticity (2.1) by low order, conforming finite elements, e.g., linear or trilinear elements. Let us assume that a triangulation τ^h of Ω is given which is shape regular and has a typical diameter of h. We denote by $W^h := W^h(\Omega)$ the corresponding conforming finite element space of finite element functions. The associated discrete problem is then to find $u_h \in W^h$, such that

$$a(u_h, v_h) = \langle F, v_h\rangle \quad \forall v_h \in W^h. \quad (2.2)$$

When there is no risk of confusion, we will drop the subscript h.

3. The FETI-DP saddle point formulation. Let a domain $\Omega \subset \mathbf{R}^d$, $d = 2, 3$ be decomposed into N nonoverlapping subdomains Ω_i of diameter H, each of which is the union of finite elements with matching finite element nodes on the boundaries of neighboring subdomains across the interface $\Gamma := \bigcup_{i\neq j} \partial\Omega_i \cap \partial\Omega_j$, where $\partial\Omega_i, \partial\Omega_j$ are the boundaries of Ω_i, Ω_j, respectively. The interface Γ is the union of edges and vertices (in 2D) and faces, edges, and vertices (in 3D). Here, for simplicity, we regard

edges in 2D and faces in 3D as open sets, that are shared by two subdomains, edges in 3D as open sets that are shared by more than two subdomains, and vertices, in 2D and 3D, as endpoints of edges; see, e.g., Toselli and Widlund [33, Chapter 4.2]. For a more detailed definition of faces, edges, and vertices; see Klawonn and Widlund [20, Section 3] and Klawonn and Rheinbach [16, Section 2].

For each subdomain Ω_i, $i = 1, \ldots, N$ we assemble the local stiffness matrices $K^{(i)}$ and load vectors $f^{(i)}$. We denote the unknowns on each subdomain, e.g., the displacements in the case of elasticity, by $u^{(i)}$.

In order to obtain again the solution of the original finite element problem (2.2), we need to enforce the continuity of the $u^{(i)}$ across the subdomain interfaces. Each nodal vector $u^{(i)}$ can be divided into a set of interior unknowns, $u_I^{(i)}$, associated with nodes in the interior of Ω_i, and interface variables, $u_\Gamma^{(i)}$, associated with nodes on the interface Γ. A subset of the interface variables, denoted by $u_\Pi^{(i)}$ or primal variables, we will enforce the continuity by global subassembly of the subdomain stiffness matrices $K^{(i)}$. For all other interface variables, denoted by $u_\Delta^{(i)}$ or dual displacement variables, we will introduce Lagrange multipliers to enforce continuity. We denote the variables that are not primal by $u_B^{(i)} = [u_I^{(i)T}, u_\Delta^{(i)T}]^T$ and partition the local stiffness matrices accordingly,

$$K^{(i)} = \begin{bmatrix} K_{BB}^{(i)} & K_{\Pi B}^{(i)\,T} \\ K_{\Pi B}^{(i)} & K_{\Pi\Pi}^{(i)} \end{bmatrix},\ K_{BB}^{(i)} = \begin{bmatrix} K_{II}^{(i)} & K_{\Delta I}^{(i)T} \\ K_{\Delta I}^{(i)} & K_{\Delta\Delta}^{(i)} \end{bmatrix},\ K_{\Pi B}^{(i)} = [K_{\Pi I}^{(i)}, K_{\Pi\Delta}^{(i)}].$$

Next, we subassemble the primal variables $u_\Pi^{(i)}, i = 1, \ldots, N$. Denoting by $R_\Pi^{(i)}$ the standard prolongation matrices, which map from the local subdomain variables $u_\Pi^{(i)}$ to the global variables $\tilde{u}_\Pi$, we obtain

$$\widetilde{K}_{\Pi\Pi} = \sum_{i=1}^{N} R_\Pi^{(i)} K_{\Pi\Pi}^{(i)} R_\Pi^{(i)\,T},\ \ \widetilde{K}_{\Pi B}^{(i)} = R_\Pi^{(i)} K_{\Pi B}^{(i)},\ \ \tilde{u}_\Pi = \sum_{i=1}^{N} R_\Pi^{(i)} u_\Pi^{(i)},\ \ \tilde{f}_\Pi = \sum_{i=1}^{N} R_\Pi^{(i)} f_\Pi^{(i)}.$$

Defining the block matrices

$$K_{BB} = \mathrm{diag}_{i=1}^{N}(K_{BB}^{(i)}), \quad \widetilde{K}_{\Pi B} = [\,\widetilde{K}_{\Pi B}^{(1)}, \ldots, \widetilde{K}_{\Pi B}^{(N)}\,],$$

we obtain the partially assembled matrix $\widetilde{K}$ and corresponding right hand side $\tilde{f}$, i.e.,

$$\widetilde{K} = \begin{bmatrix} K_{BB} & \widetilde{K}_{\Pi B}^T \\ \widetilde{K}_{\Pi B} & \widetilde{K}_{\Pi\Pi} \end{bmatrix}, \qquad \tilde{f} = \begin{bmatrix} f_B \\ \tilde{f}_\Pi \end{bmatrix},$$

where $f_B = [\,f_B^{(1)\,T}, \ldots, f_B^{(N)\,T}\,]^T$. Choosing a sufficient number of primal variables $u_\Pi^{(i)}$ to constrain our solution, results in a symmetric positive definite matrix $\widetilde{K}$; see [20].

We introduce a jump operator B with entries $0, -1$ or 1 and Lagrange multipliers λ to enforce continuity on the remaining interface variables $u_\Delta^{(i)}$. For convenience, we always use the full set of (redundant) Lagrange multipliers wherever more than two subdomains share a single node. We introduce the notation

$$u_B = [\,u_B^{(1)\,T}, \ldots, u_B^{(N)\,T}\,]^T, \qquad u = [\,u_B^T,\, \tilde{u}_\Pi^T\,]^T.$$

Now, we can formulate the FETI-DP saddle point problem,

$$\begin{bmatrix} \widetilde{K} & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} u \\ \lambda \end{bmatrix} = \begin{bmatrix} \tilde{f} \\ 0 \end{bmatrix}, \qquad u \in \mathbf{R}^n,\, \lambda \in \mathbf{R}^m, \tag{3.1}$$

from which the solution of the original finite element problem (2.2) can be obtained by subassembling the solution u in the interface variables $u_\Delta^{(i)}$. We will also use the notation

$$\mathcal{A}x = \mathcal{F},$$

where

$$\mathcal{A} := \begin{bmatrix} \widetilde{K} & B^T \\ B & 0 \end{bmatrix}, \quad x := \begin{bmatrix} u \\ \lambda \end{bmatrix}, \quad \mathcal{F} := \begin{bmatrix} \tilde{f} \\ 0 \end{bmatrix}.$$

So far, we have not discussed how to choose the decomposition of the local interface vectors $u_\Gamma^{(i)}$ into primal variables $u_\Pi^{(i)}$ and remaining (dual) interface vectors $u_\Delta^{(i)}$. One immediate possibility is to choose selected vertices (or corners) of Ω_i as nodes where the variables $u_\Gamma^{(i)}$ are assembled; cf. Farhat et al. [10] or Mandel and Tezaur [26]. This yields good convergence bounds of the order of $(1+\log(H/h))^2$ in two dimensions, see [26], but not in three dimensions, see Klawonn, Widlund, and Dryja [21]. Numerical evidence for this deterioration in three dimensions was first given in Farhat, Lesoinne, and Pierson [11]; cf. also Klawonn, Rheinbach, and Widlund [17, Table 5, Figure 2].

To obtain scalable algorithms with condition number bounds of the order of $(1+\log(H/h))^2$ in three dimensions, Klawonn, Widlund, and Dryja [21] for scalar, second order elliptic equations, introduced averages over selected subdomain edges or faces to be continuous across the interface. For linear elasticity in three dimensions, Klawonn and Widlund [20] introduced averages and first order moments over selected edges as primal variables. For some very hard cases with large coefficient jumps, e.g., in the diffusion coefficient or the stiffness of the material, some vertices have to be selected as primal variables as well, in order to obtain a condition number bound of the order of $(1+\log(H/h))^2$ which is robust with respect to the coefficient jumps. Numerical results with vertex and face average constraints for three dimensional elasticity problems were already presented in Farhat, Lesoinne, and Pierson [11]; see also the doctoral dissertation of Pierson [27]. There are two different possibilities to implement the average and moment constraints over faces or edges. The first is the introduction of optional Lagrange multipliers; see Farhat, Lesionne, and Pierson [11], Pierson [27], or Klawonn and Widlund [20]. The second is to apply a transformation of basis,

$$\overline{K}^{(i)} = T^T K^{(i)} T,$$

introducing the averages and moments explicitly as new variables. Then, the dual interface variables $u_\Delta^{(i)}$ have zero average or first order moment on the selected edges or faces; see Klawonn and Widlund [20] and, for numerical results and a parallel implementation, Klawonn and Rheinbach [16]. In this article, we will always use a transformation of basis when edge averages are used as primal constraints; for a detailed algorithmic description of applying the transformation of basis, see [16] or [20]. The choice of a good coarse problem, i.e., the selection of vertex, edge, and face constraints, is of vital importance to the convergence and scalability of FETI-DP methods. A detailed description is beyond the scope of this paper; for details see, e.g., [20] or [16]. In our experiments, in two dimensions, we will always use continuity at all vertices and continuity of all edge averages. This version will be denoted Algorithm B; see [21], where an analogous notation was introduced for three dimensions. In three dimensions, in our experiments, we will always use continuous edge averages at

all edges and no vertex or face constraints. This version is denoted Algorithm D_E; see [16]. For further details on the implementation and other algorithmic choices, see [21], [20], [16], or [11].

With the notation introduced before, we can rewrite (3.1) as

$$\begin{bmatrix} K_{BB} & \widetilde{K}_{\Pi B}^T & B_B^T \\ \widetilde{K}_{\Pi B} & \widetilde{K}_{\Pi\Pi} & 0 \\ B_B & 0 & 0 \end{bmatrix} \begin{bmatrix} u_B \\ \tilde{u}_\Pi \\ \lambda \end{bmatrix} = \begin{bmatrix} f_B \\ \tilde{f}_\Pi \\ 0 \end{bmatrix}. \tag{3.2}$$

Eliminating u_B by one step of block Gaussian elimination, we obtain the reduced system

$$\begin{bmatrix} \widetilde{S}_{\Pi\Pi} & -\widetilde{K}_{\Pi B} K_{BB}^{-1} B_B^T \\ -B_B K_{BB}^{-1} \widetilde{K}_{\Pi B}^T & -B_B K_{BB}^{-1} B_B^T \end{bmatrix} \begin{bmatrix} \tilde{u}_\Pi \\ \lambda \end{bmatrix} = \begin{bmatrix} \tilde{f}_\Pi - \widetilde{K}_{\Pi B} K_{BB}^{-1} f_B \\ -B_B K_{BB}^{-1} f_B \end{bmatrix}, \tag{3.3}$$

where $\widetilde{S}_{\Pi\Pi} = \widetilde{K}_{\Pi\Pi} - \widetilde{K}_{\Pi B} K_{BB}^{-1} \widetilde{K}_{\Pi B}^T$. Here, we will also use the notation

$$\mathcal{A}_r x_r = \mathcal{F}_r,$$

where

$$\mathcal{A}_r = \begin{bmatrix} \widetilde{S}_{\Pi\Pi} & -\widetilde{K}_{\Pi B} K_{BB}^{-1} B_B^T \\ -B_B K_{BB}^{-1} \widetilde{K}_{\Pi B}^T & -B_B K_{BB}^{-1} B_B^T \end{bmatrix}, \; x_r := \begin{bmatrix} \tilde{u}_\Pi \\ \lambda \end{bmatrix},$$

and

$$\mathcal{F}_r := \begin{bmatrix} \tilde{f}_\Pi - \widetilde{K}_{\Pi B} K_{BB}^{-1} f_B \\ -B_B K_{BB}^{-1} f_B \end{bmatrix}.$$

By also eliminating the primal variables $\tilde{u}_\Pi$, we obtain the reduced system

$$F\lambda = d, \tag{3.4}$$

where

$$\begin{aligned} F &:= B_B K_{BB}^{-1} B_B^T + B_B K_{BB}^{-1} \widetilde{K}_{B\Pi} \widetilde{S}_{\Pi\Pi}^{-1} \widetilde{K}_{\Pi B} K_{BB}^{-1} B_B^T = B \widetilde{K}^{-1} B^T, \\ d &:= B_B K_{BB}^{-1} f_B + B_B K_{BB}^{-1} \widetilde{K}_{\Pi B}^T \widetilde{S}_{\Pi\Pi}^{-1} (\tilde{f}_\Pi - \widetilde{K}_{\Pi B} K_{BB}^{-1} f_B) = B \widetilde{K}^{-1} \tilde{f}. \end{aligned}$$

The linear system (3.4) is the standard, exact FETI-DP system which is solved using preconditioned conjugate gradients and an appropriate preconditioner M^{-1}; cf. Section 4.

4. Exact and inexact FETI-DP methods. In the standard, exact FETI-DP methods two different preconditioners are commonly used, the theoretically optimal Dirichlet preconditioner M_D and the lumped preconditioner M_L. To define these preconditioners, we partition the local matrices according to interior and interface variables, $u_I^{(i)}$ and $u_\Gamma^{(i)}$, and obtain

$$K^{(i)} = \begin{bmatrix} K_{II}^{(i)} & K_{\Gamma I}^{(i)\,T} \\ K_{\Gamma I}^{(i)} & K_{\Gamma\Gamma}^{(i)} \end{bmatrix}$$

with

$$K_{II} = \mathrm{diag}_{i=1}^{N}(K_{II}^{(i)}), \qquad K_{\Gamma I} = \mathrm{diag}_{i=1}^{N}(K_{\Gamma I}^{(i)}), \qquad K_{\Gamma\Gamma} = \mathrm{diag}_{i=1}^{N}(K_{\Gamma\Gamma}^{(i)}).$$

The Dirichlet preconditioner M_D is defined by

$$M_D^{-1} = B_D R_\Gamma^T (K_{\Gamma\Gamma} - K_{\Gamma I} K_{II}^{-1} K_{\Gamma I}^T) R_\Gamma B_D^T$$

and the lumped preconditioner M_L by

$$M_L^{-1} = B_D R_\Gamma^T K_{\Gamma\Gamma} R_\Gamma B_D^T$$

where

$$R_\Gamma = \mathrm{diag}_{i=1}^{N}(R_\Gamma^{(i)}).$$

Here, the matrices $R_\Gamma^{(i)}$ are restriction matrices which restrict the degrees of freedom of a subdomain to the interface. The matrices B_D are scaled variants of the jump operator B where the contribution from and to each interface node is scaled by the inverse of the multiplicity of the node. The multiplicity of a node is defined as the number of subdomains it belongs to. It is well known that for heterogeneous problems a more elaborate scaling is necessary, see, e.g., [20].

The original or standard, exact FETI-DP method is the method of conjugate gradients applied to the symmetric positive definite system

$$F\lambda = d$$

with the preconditioners M_D^{-1} or M_L^{-1}. We note that only for the Dirichlet preconditioner, we have the polylogarithmic condition number bounds mentioned before; see [26], [21], [20]. The term "exact" refers here to the exact solution of the coarse problem given by $\widetilde{S}_{\Pi\Pi}$ and the exact solution of the local Neumann subdomain problems $K_{BB}^{(i)}$. When the Dirichlet preconditioner is used, we of course also have to solve the local Dirichlet problems $K_{II}^{(i)}$ exactly.

Now, we are going to present new, inexact FETI-DP methods by solving the saddle point problems (3.1) and (3.3) iteratively, using block triangular preconditioners and a suitable Krylov space method.

For the saddle point problems (3.1) and (3.3), we introduce the block triangular preconditioners $\hat{\mathcal{B}}_L$ and $\hat{\mathcal{B}}_{r,L}$, respectively, as

$$\hat{\mathcal{B}}_L^{-1} = \begin{bmatrix} \hat{K}^{-1} & 0 \\ M^{-1}B\hat{K}^{-1} & -M^{-1} \end{bmatrix}, \qquad \hat{\mathcal{B}}_{r,L}^{-1} = \begin{bmatrix} \hat{S}_{\Pi\Pi}^{-1} & 0 \\ -M^{-1}B_B K_{BB}^{-1} \widetilde{K}_{\Pi B}^T \hat{S}_{\Pi\Pi}^{-1} & -M^{-1} \end{bmatrix}.$$

where $\hat{K}^{-1}$ and $\hat{S}_{\Pi\Pi}^{-1}$ are assumed to be spectrally equivalent preconditioners for $\widetilde{K}$ and $\widetilde{S}_{\Pi\Pi}$, respectively, with bounds independent of the discretization parameters h, H. The matrix block M^{-1} is assumed to be a good preconditioner for the FETI-DP system matrix F and can be chosen as the Dirichlet or the lumped preconditioners M_D^{-1} and M_L^{-1}, respectively. We will denote the corresponding right preconditioners by the subscript R, i.e., we have $\hat{\mathcal{B}}_R = \hat{\mathcal{B}}_L^T$ and $\hat{\mathcal{B}}_{r,R} = \hat{\mathcal{B}}_{r,L}^T$.

We note that $\widehat{K}^{-1}$ can also be defined using the following exact factorization of $\widetilde{K}^{-1}$, i.e.,

$$\begin{bmatrix} K_{BB} & \widetilde{K}_{\Pi B}^T \\ \widetilde{K}_{\Pi B} & \widetilde{K}_{\Pi\Pi} \end{bmatrix}^{-1} = \begin{bmatrix} I & -K_{BB}^{-1}\widetilde{K}_{\Pi B}^T \\ 0 & I \end{bmatrix} \begin{bmatrix} K_{BB}^{-1} & 0 \\ 0 & \widetilde{S}_{\Pi\Pi}^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -\widetilde{K}_{\Pi B} K_{BB}^{-1} & I \end{bmatrix}.$$

We note that $\widetilde{K}_{\Pi B} K_{BB}^{-1} =: \overline{K}_{\Pi B}$ is always built explicitly in a preprocessing step, since we need it to form $\widetilde{S}_{\Pi\Pi}$. To obtain a preconditioner $\widehat{K}^{-1}$, we can now replace K_{BB}^{-1} and $\widetilde{S}_{\Pi\Pi}^{-1}$ by good preconditioners $\widehat{K}_{BB}^{-1}$ and $\hat{S}_{\Pi\Pi}^{-1}$. This yields the preconditioner

$$\hat{K}^{-1} = \begin{bmatrix} I & -\overline{K}_{\Pi B}^T \\ 0 & I \end{bmatrix} \begin{bmatrix} \hat{K}_{BB}^{-1} & 0 \\ 0 & \hat{S}_{\Pi\Pi}^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -\overline{K}_{\Pi B} & I \end{bmatrix}. \tag{4.1}$$

We note that the application of $\widehat{K}^{-1}$ to a vector only involves one application of $\widehat{K}_{BB}^{-1}$ and $\hat{S}_{\Pi\Pi}^{-1}$ each. Such a factorization was also the basis for iterative substructuring methods with inexact Dirichlet solvers; see, e.g., Smith, Bjørstad, and Gropp [29, Chapter 4.4] or Toselli and Widlund [33, Chapter 4.3] and the references given therein.

It is also possible to use exact local solvers, i.e., $\widehat{K}_{BB}^{-1} = K_{BB}^{-1}$, and to solve only the coarse problem inexactly. This variant is closely related to preconditioning the reduced system (3.3) by an appropriate block triangular preconditioner.

Our inexact FETI-DP methods are now given by using a Krylov space method for nonsymmetric systems, e.g., GMRES, to solve the preconditioned systems

$$\hat{\mathcal{B}}_L^{-1} \mathcal{A} x = \hat{\mathcal{B}}_L^{-1} \mathcal{F}$$

and

$$\hat{\mathcal{B}}_{r,L}^{-1} \mathcal{A}_r x_r = \hat{\mathcal{B}}_L^{-1} \mathcal{F}_r,$$

respectively.

Let us note that we can also use a positive definite reformulation of the two preconditioned systems, which allows the use of conjugate gradients. For this reformulation, a special inner product and a scaling of the preconditioners $\widehat{K}$ and $\hat{S}_{\Pi\Pi}$ have to be used; see Sections 5 and 6 for further details.

5. A review of block triangular preconditioners for symmetric saddle point problems. In this section, we review some theoretical convergence results for block triangular preconditioners applied to symmetric saddle point problems. This theory will then be used in the next section to derive convergence estimates for the full and the reduced preconditioned system. In general, using block triangular preconditioners leads to nonsymmetric preconditioned systems, even when the original saddle point problem is symmetric. Thus, Krylov space methods which are well suited for nonsymmetric linear systems have to be chosen, e.g., GMRES, BiCGSTAB, QMR or variants of these methods. In some cases, the preconditioned system is symmetric positive definite in a certain inner product; then, a conjugate gradient method can be used; see Bramble and Pasciak [4]. Let us note that the theory for deriving a priori GMRES convergence bounds for block triangular preconditioners is not complete and still an area of research. In contrast to conjugate gradient methods, eigenvalue bounds are in general not sufficient for convergence estimates of GMRES. Our presentation in this section is based on Klawonn [15].

We consider a mixed linear system of the form

$$\mathcal{A} x = \mathcal{F}, \tag{5.1}$$

where we have

$$\mathcal{A} = \begin{bmatrix} A & B^T \\ B & -C \end{bmatrix}, \; x = \begin{bmatrix} u \\ \lambda \end{bmatrix}, \; \mathcal{F} = \begin{bmatrix} f \\ g \end{bmatrix}.$$

We assume that $A \in \mathbb{R}^{n\times n}$ is a symmetric positiv definite matrix, $C \in \mathbb{R}^{m\times m}$ a symmetric positive semi definite matrix, and $B \in \mathbb{R}^{m\times n}$ a matrix with full rank. Furthermore, we define left and right block triangular preconditioners

$$\hat{\mathcal{B}}_L = \begin{bmatrix} \hat{A} & 0 \\ B & -\hat{C} \end{bmatrix}, \quad \hat{\mathcal{B}}_R = \begin{bmatrix} \hat{A} & B^T \\ 0 & -\hat{C} \end{bmatrix}.$$

Here, we assume that there exist constants $\alpha_0, \alpha_1 > 0$, such that

$$\alpha_0 \; u^T \hat{A} u \leq u^T A u \leq \alpha_1 \; u^T \hat{A} u \quad \forall u \in \mathbb{R}^n \tag{5.2}$$

and constants $\gamma_0, \gamma_1 > 0$ such that

$$\gamma_0 \; \lambda^T \hat{C} \lambda \leq \lambda^T S_C \lambda \leq \gamma_1 \; \lambda^T \hat{C} \lambda \quad \forall \lambda \in \mathbb{R}^m, \tag{5.3}$$

where the Schur complement S_C is defined as $S_C := C + BA^{-1}B^T$.

In our analysis, we only consider the case of inexact preconditioners $\hat{A}$, i.e., we exclude the case $\hat{A} = A$. In our application on FETI-DP methods, the exact case relates to the standard, exact FETI-DP method and therefore, we do not have to analyze it here. Nevertheless, an exact solver for A can be applied with GMRES; see Klawonn [15] and Simoncini [28] for numerical results and eigenvalue bounds.

To the best of our knowledge, the first GMRES convergence analysis for block triangular preconditioners applied to symmetric saddle point problems was given in Klawonn [14, 15], where the following assumption was made for the preconditioner $\hat{A}$,

$$1 < \alpha_0 \leq a_1, \tag{5.4}$$

which can be always obtained by an appropriate scaling. We will briefly review those results using our notation. We first introduce the symmetric positive definite matrices $\mathcal{H}$ and $\tilde{\mathcal{H}}$,

$$\mathcal{H} = \begin{bmatrix} A - \hat{A} & 0 \\ 0 & \hat{C} \end{bmatrix}, \quad \tilde{\mathcal{H}} = \begin{bmatrix} A & 0 \\ 0 & S_C \end{bmatrix}.$$

From a direct calculation, we obtain the symmetric matrix

$$\mathcal{H}\hat{\mathcal{B}}_L^{-1}\mathcal{A} = \begin{bmatrix} A\hat{A}^{-1}A - A & (A-\hat{A})\hat{A}^{-1}B^T \\ B\hat{A}^{-1}(A-\hat{A}) & C + B\hat{A}^{-1}B^T \end{bmatrix}.$$

To apply the theory proven in [15], we note that the equality $\mathcal{H}\hat{\mathcal{B}}_L^{-1}\mathcal{A} = \mathcal{A}\hat{\mathcal{B}}_R^{-1}\mathcal{H}$ holds; see also [15, Remark 2]. The next lemma is proven in [15, Lemma 3.3].

LEMMA 5.1. *There exist positive constants $\tilde{C}_0, \tilde{C}_1$, such that*

$$\tilde{C}_0 \; x^T \tilde{\mathcal{H}} x \leq x^T \mathcal{H}\hat{\mathcal{B}}_L^{-1}\mathcal{A} x \leq \tilde{C}_1 \; x^T \tilde{\mathcal{H}} x \quad \forall x \in \mathbb{R}^{n+m},$$

where $\tilde{C}_0 = \min\{(\alpha_0 - 1), 1\}/3$ and $\tilde{C}_1 = 3\max\{(\alpha_1 - 1), 1\}$.

Using (5.2), (5.3), and (5.4), we obviously have the following spectral equivalence:

$$\min\{\frac{\alpha_1}{\alpha_1 - 1}, \gamma_0\} \; x^T \mathcal{H} x \leq x^T \tilde{\mathcal{H}} x \leq \max\{\frac{\alpha_0}{\alpha_0 - 1}, \gamma_1\} \; x^T \mathcal{H} x \quad \forall x \in \mathbb{R}^{n+m}. \tag{5.5}$$

Combining Lemma 5.1 and (5.5), we obtain, see also [15, Lemma 3.4],

LEMMA 5.2. *We have*

$$C_0\ x^T\mathcal{H}x \leq x^T\mathcal{H}\hat{\mathcal{B}}_L^{-1}\mathcal{A}x \leq C_1\ x^T\mathcal{H}x \quad \forall x \in \mathbb{R}^{n+m}$$

with positive constants $C_0 = \left(\frac{1}{3}\min\{(\alpha_0-1),1\}\min\{\frac{\alpha_1}{\alpha_1-1},\gamma_0\}\right)$ *and* $C_1 = \left(3\max\{(\alpha_1-1),1\}\max\{\frac{\alpha_0}{\alpha_0-1},\gamma_1\}\right)$.

From this lemma immediately follows that the eigenvalues of $\hat{\mathcal{B}}_L^{-1}\mathcal{A}$ are real, positive, and contained in the intervall $[C_0, C_1]$; cf. also [15, Theorem 3.5]. We can now use the bounds given in Lemma 5.2 to provide a convergence bound for GMRES minimizing the residual in an arbitrary norm equivalent to the $\mathcal{H}$-norm; see [15, Theorem 3.7], where this result is given for right preconditioning with $\hat{\mathcal{B}}_U^{-1}$ and the $\mathcal{H}^{-1}$-inner product. The result is based on the fact that the method of conjugate residuals and GMRES both minimize the same residual in the norm used and that the for symmetric positive definite matrices, a convergence bound for the method of conjugate residuals can be given in terms of the condition number of the preconditioned system; see [15, Theorem 3.7] for further details.

THEOREM 5.1. *Let* $\hat{\mathcal{H}}$ *be a symmetric positive definite matrix, such that* $\bar{C}_0\mathcal{H} \leq \hat{\mathcal{H}} \leq \bar{C}_1\mathcal{H}$ *with positive constants* $\bar{C}_0, \bar{C}_1$. *Then, we have*

$$\frac{\|r^{(k)}\|_{\hat{\mathcal{H}}}}{\|r^{(0)}\|_{\hat{\mathcal{H}}}} \leq \frac{\bar{C}_1}{\bar{C}_0} 2 \left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^k,$$

where $r^{(0)}$ *and* $r^{(k)}$ *are the initial and* k*-th residual of GMRES, respectively, and* $\kappa := \kappa(\hat{\mathcal{B}}_L^{-1}\mathcal{A}) \leq \frac{C_1}{C_0}$ *is the condition number of* $\hat{\mathcal{B}}_L^{-1}\mathcal{A}$ *in the* $\mathcal{H}$*-inner product.*

Let us note that for the block triangular preconditioner, to the best of our knowledge, no complete theory exists for a priori GMRES convergence bounds in the Euclidean norm. Nevertheless, the Euclidean inner product is usually used to implement this preconditioning approach with GMRES. Let us note that recently Simoncini [28] has given an eigenvalue analysis of block triangular preconditioners with right preconditioning without the scaling assumption (5.4). The bounds given in [28] also depend on $\alpha_0, \alpha_1, \gamma_0$, and γ_1. It is shown that the eigenvalues of $\mathcal{A}\hat{\mathcal{B}}_R^{-1}$ become complex when the smallest eigenvalue $\lambda_{\min}$ of $A\hat{A}^{-1}$ is smaller than one. In that case, the eigenvalues lie in a disk centered at unity with radius $\sqrt{1-\lambda_{\min}}$.

Since $\hat{\mathcal{B}}_L^{-1}\mathcal{A}$ is symmetric positive definite in the $\mathcal{H}$-inner product, we can also apply the method of conjugate gradients using this special inner product; see Bramble and Pasciak [4] or Dohrmann and Lehoucq [8]. Since this is a non standard implementation of the cg-method, we provide a version of this algorithm to solve $\hat{\mathcal{B}}_L^{-1}\mathcal{A}x = \hat{\mathcal{B}}_L^{-1}\mathcal{F}$ in Figure 5.1. Here, `xstart` is our initial guess. We note that due to this special implementation, no application of $\hat{C}$ or $\hat{A}$ are needed. This is important since in our applications, we are usually only able to apply $\hat{C}^{-1}$ and $\hat{A}^{-1}$ to a vector.

It is well-known that a convergence bound for conjugate gradients can be given in terms of the square root of the spectral condition number of the preconditioned system. From Lemma 5.2, we immediately obtain an upper bound for the spectral condition number of $\hat{\mathcal{B}}_L^{-1}\mathcal{A}$.

6. Analysis of the preconditioners. In this section, we will apply the general theory for block triangular preconditioners presented in Section 5 to our inexact FETI-DP methods given in Section 4. We only have to identify the matrix blocks in our

$$
\begin{aligned}
\mathtt{z} &= \hat{\mathcal{B}}_L^{-1}(\mathbf{f} - \mathcal{A}\,\mathtt{xstart}) \\
\mathtt{Hz} &= \mathcal{H}\hat{\mathcal{B}}_L^{-1}(\mathbf{f} - \mathcal{A}\,\mathtt{xstart}) \\
\mathtt{p} &= \mathtt{z} \\
\mathtt{zHz} &= \langle \mathtt{z}, \mathtt{Hz} \rangle \\
\mathtt{lzl} &= \|\mathtt{z}\| \\
\text{Until } \|\mathtt{z}\| \ / \ \mathtt{lzl} &< 10^{-7} \\
\mathtt{Ap} &= \mathcal{A}\,\mathtt{p} \\
\mathtt{HBAp} &= \mathcal{H}\hat{\mathcal{B}}_L^{-1}\,\mathtt{Ap} \\
\alpha &= \mathtt{zHz}/\,\langle \mathtt{HBAp}, \mathtt{p} \rangle \\
\mathtt{x} &= \mathtt{x} + \alpha\,\mathtt{p} \\
\mathtt{z} &= \mathtt{z} - \alpha\hat{\mathcal{B}}_L^{-1}\,\mathtt{Ap} \\
\mathtt{Hz} &= \mathtt{Hz} - \alpha\,\mathtt{HBAp} \\
\mathtt{zHzo} &= \mathtt{zHz} \\
\mathtt{zHz} &= \langle \mathtt{z}, \mathtt{Hz} \rangle \\
\beta &= \mathtt{zHz}/\mathtt{zHzo} \\
\mathtt{p} &= \mathtt{z} + \beta\,\mathtt{p}
\end{aligned}
$$

FIG. 5.1. *Conjugate gradient algorithm in the $\mathcal{H}$-inner product.*

inexact FETI-DP methods with those in the general presentation and provide concrete estimates for the constants α_0, α_1 in (5.2) and γ_0, γ_1 in (5.3) in order to obtain our convergence estimates.

6.1. Preconditioning the original system. We first consider the original FETI-DP system (3.1). Here, we have

$$
\mathcal{A} = \begin{bmatrix} \widetilde{K} & B^T \\ B & 0 \end{bmatrix}, \qquad \hat{\mathcal{B}}_L^{-1} = \begin{bmatrix} \widehat{K}^{-1} & 0 \\ M^{-1}B\widehat{K}^{-1} & -M^{-1} \end{bmatrix}.
$$

Hence, we also have

$$
A := \widetilde{K}, \ \ \hat{A} := \widehat{K}, \ \ C := 0, \ \ \hat{C} := M, \ \ S_C := F,
$$

and B is the same matrix as in the original FETI-DP method.

We assume that $\widehat{K}$ is a good preconditioner for $\widetilde{K}$ with optimal spectral bounds α_0 and α_1 which are independent of the discretization parameters h, H. Good examples for such preconditioners are based on geometric and algebraic multigrid methods. Let us note that in some of our experiments incomplete Cholesky decompositions are used although the bounds then will not be optimal.

The spectral bounds γ_0 and γ_1 in (5.3) are given by the eigenvalue bounds of the standard, exact FETI-DP method. There exists a constant $C > 0$, independent of h, H, such that

$$
\lambda^T M \lambda \leq \lambda^T F \lambda \leq C\,(1 + \log(H/h))^2 \lambda^T M \lambda \quad \forall \lambda \in \text{range}\,(B);
$$

see Mandel and Tezaur [26] for second and fourth order, scalar elliptic equations in two dimensions, Klawonn, Widlund, and Dryja [21] for second order, scalar elliptic equations with large coefficient jumps in three dimensions, and Klawonn and Widlund [20] for linear elasticity problems with large coefficient jumps in three dimensions. Thus, we have

$$\gamma_0 := 1, \;\; \gamma_1 := C\,(1+\log(H/h))^2.$$

From these estimates we see that, asymptotically, for our inexact FETI-DP method operating on the original system (3.1), we obtain convergence bounds of the same quality as for the standard, exact FETI-DP methods. This holds for GMRES as well as for conjugate gradients.

6.2. Preconditioning the reduced system. We now consider the reduced FETI-DP system (3.3). Therefore, we have

$$\mathcal{A}_r = \begin{bmatrix} \widetilde{S}_{\Pi\Pi} & -\widetilde{K}_{\Pi B} K_{BB}^{-1} B_B^T \\ -B_B K_{BB}^{-1} \widetilde{K}_{\Pi B}^T & -B_B K_{BB}^{-1} B_B^T \end{bmatrix},$$

$$\hat{\mathcal{B}}_{r,L}^{-1} = \begin{bmatrix} \hat{S}_{\Pi\Pi}^{-1} & 0 \\ -M^{-1} B_B K_{BB}^{-1} \widetilde{K}_{\Pi B}^T \hat{S}_{\Pi\Pi}^{-1} & -M^{-1} \end{bmatrix}$$

and identify $\mathcal{A}$ and $\hat{\mathcal{B}}_L$ from Section 5 with $\mathcal{A}_r$ and $\hat{\mathcal{B}}_{r,L}$, respectively. Hence, we also have

$$A := \widetilde{S}_{\Pi\Pi},\; \hat{A} := \hat{S}_{\Pi\Pi},\; C := B_B K_{BB}^{-1} B_B^T,\; \hat{C} := M,\; B := -B_B K_{BB}^{-1} \widetilde{K}_{\Pi B}^T\,.$$

As before, we also assume here that $\hat{S}_{\Pi\Pi}$ is a good preconditioner for $\widetilde{S}_{\Pi\Pi}$ with optimal spectral bounds α_0 and α_1 which are independent of the discretization parameters h, H.

For the Schur complement S_C we have again

$$S_C = C + BA^{-1}B^T = B_B K_{BB}^{-1} B_B^T + B_B K_{BB}^{-1} \widetilde{K}_{\Pi B}^T \widetilde{S}_{\Pi\Pi}^{-1} \widetilde{K}_{\Pi B} K_{BB}^{-1} B_B^T = F.$$

Since $\hat{C} = M$, we have

$$\gamma_0 := 1, \;\; \gamma_1 := C\,(1+\log(H/h))^2.$$

From these estimates, we see that asymptotically, we again obtain convergence bounds of the same quality as for the standard, exact FETI-DP methods and the inexact FETI-DP methods operating on the original system (3.1). As for the latter method, these bounds hold for GMRES as well as for conjugate gradients.

7. Performance considerations. We will use GMRES or CG to solve the systems (3.1) and (3.2) iteratively using the preconditioners $\hat{B}_L^{-1}$ and $\hat{B}_{r,L}^{-1}$, respectively.

We restrict our rough cost estimate to the use of GMRES. The method for the reduced system (3.2) iterates simultaneously on $\tilde{u}_\Pi$ and λ. Since the dimension of $\tilde{u}_\Pi$ is small, the computational cost spent in the inner products of the Krylov method is comparable to that of the original, exact FETI-DP method. In fact, the dimension of $[\,\tilde{u}_\Pi^T, \lambda^T\,]^T$ is smaller than or equal the number of Lagrange multipliers in the original (one-level) FETI method; see, e.g., Klawonn and Widlund [19] or Toselli and Widlund [33]. This is a first indication that the communication cost is also comparable to the original FETI method.

A more careful analysis, also considering the reuse of the results of gather operations in the different blocks, shows that our method operating on the reduced system, needs one more scatter operation with $\tilde{u}_\Pi$ in the application of the system matrix compared to FETI-DP, although this may vary in concrete implementations.

Note that the significant computational cost in applying the system matrix $\mathcal{A}_r$ to a vector is the same as in applying F except for the matrix-vector product with $\widetilde{S}_{\Pi\Pi}^{-1}$ which is now shifted to the preconditioner and replaced by $\hat{S}_{\Pi\Pi}^{-1}$. Although $\hat{S}_{\Pi\Pi}^{-1}$ appears in two blocks of the preconditioner, in the implementation, the product with a vector has to be carried out only once in each iteration.

We repeat that in all methods considered in this paper the matrix-matrix product $K_{BB}^{-1}\widetilde{K}_{\Pi B}^T$ is built explicitly in a preprocessing step, as is done generally in standard, exact FETI-DP.

The computational effort for the methods that iterate on displacement variables u and Lagrange multipliers λ simultaneously seems to be considerably higher, at least at first glance, because of the higher computation cost in the inner products of the Krylov subspace method. But, fortunately, the inner products $u_B^T v_B$ are perfectly parallel and can be calculated by each processor seperately. Only the scalar results of $u_B^{(i)\,T} v_B^{(i)}$ have to be communicated among the processors.

8. Numerical results. In this section we present numerical results for the preconditioners analyzed in the previous sections. We apply the preconditioners to 2D and 3D linear elasticity problems. In the tables of this section we denote the iterative substructuring method using the preconditioner $\hat{\mathcal{B}}_L^{-1}$ for the system (3.1) by inexact FETI-DP or iFETI-DP. The method using the preconditioner $\hat{\mathcal{B}}_{r,L}^{-1}$ for the system (3.2) iterating on the variables $[\,\tilde{u}_\Pi^T, \lambda^T\,]^T$, is denoted as inexact reduced FETI-DP or irFETI-DP. We always state the Krylov subspace method which is used as accelerator, either GMRES or CG, and which part of the preconditioner is solved inexactly and by which method. We generally use left preconditioning with GMRES so that we can use the same implementation for the preconditioner as for CG.

8.1. Direct solvers. We use the modified FETI-DP formulations for structured benchmark problems in 2D and 3D using exact solvers (Cholesky or LU decomposition) for the coarse grid problem and the local subdomain problems to verify that the methods perform well in the best case. We use GMRES with left preconditioning and CG to solve the preconditioned problem. In order to fulfill assumption (5.4), we scale the results obtained from the direct solver forward backward substitutions for the decomposition of $\widetilde{S}_{\Pi\Pi}$ by a factor of 0.9999^2 when using CG with irFETI-DP. When using CG with iFETI-DP we scale with 0.99999^2, i.e, we have $\hat{S}_{\Pi\Pi}^{-1} = 0.99999^2 \cdot \widetilde{S}_{\Pi\Pi}^{-1}$ and $\hat{K}_{BB}^{-1} = 0.99999^2 \cdot K_{BB}^{-1}$.

For the 2D problems in Table 8.1 we have chosen a larger coarse grid problem than would be necessary for a compressible elasticity problem in two dimensions. It is well known that vertex constraints are sufficient for this case to ensure scalability.

We see from Table 8.1 that the algorithms perform as expected and converge independently of the number of subdomains. The iteration count and estimated eigenvalues are almost identical to the numbers that we get from the original FETI-DP method.

The eigenvalue estimates are obtained from the Lanczos process in the conjugate gradient method. Note that the accuracy of these estimates is much lower than the number of given digits suggests. Nevertheless, we can see that the considered FETI-DP algorithms lead to Lanczos eigenvalue estimates which are very close to each

other.

8.2. Inexact solvers for the coarse problem. We first investigate the effect of an approximate solver for the coarse grid problem on the iteration count and the estimated eigenvalues of the preconditioned operators.

We use cycles of the Fraunhofer SCAI parallel algebraic multigrid package SAMG [32, 31, 30, 34] by Stüben and Clees to precondition $\widetilde{S}_{\Pi\Pi}$. In the experiments in Table 8.2 we use V-cycles and ILU(0)-smoothing within SAMG. We see that the method performs remarkably well in 2D and 3D for the structured benchmark problems. The iteration count is comparable to that of the original FETI-DP method, cf. Table 8.1, and the estimated condition number is only slightly higher. A higher condition number is, of course, expected. If we increase the number of SAMG cycles we get closer to the eigenvalues of the original FETI-DP method.

Whenever scaling is used ("sc.") it is calculated using three digits of an eigenvalue estimate obtained from a few, typically less than ten, CG iterations with $\tilde{S}_{\Pi\Pi}$ preconditioned by SAMG. From the experiments we do not see the necessity of scaling when using GMRES.

In Table 8.3 we present a larger test problem with 13 824 subdomains which we are unable to solve with our current FETI-DP implementation since the direct factorization of $\widetilde{S}_{\Pi\Pi}$ requires too much memory ($>$ 2 GB); see Section 9 for further details.

Next, we consider a larger and unstructured mesh. The mechanical part shown in Figure 8.1, is discretized using 1 291 933 linear, tetrahedral elements. The resulting problem has 841 836 degrees of freedom; see [16]. We have partitioned the mechanical part into 1 024 and 2 048 very small subdomains in order to obtain a coarse problem of reasonable size. For this we use the graph partitioning software ParMetis [13]. Typically, it is more efficient to partition this mechanical part into a considerably smaller number of subdomains; see [16]. From Table 8.6 we see that even for this industrial benchmark problem the GMRES iteration count remains acceptable and compares well with standard FETI-DP using CG.

8.3. Inexact solvers for the local problems. We only present preliminary results for the inexact solution of the local Neumann problems. For a first set of experiments, presented in Table 8.4, we use incomplete Cholesky decompositions with a threshold of 10^{-4} for a structured 3D elasticity problem. We use renumbering [7] before the incomplete factorization. For these calculations the incomplete Cholesky factorization (ICC) uses on average about 50% of the memory required for an exact factorization. Here, we only present results for GMRES. From these results we see that the GMRES iteration count is still comparable to standard FETI-DP using CG. Here, we do not use any scaling to satisfy (5.4) and we also note that ICC is not an optimal preconditioner.

In another set of experiments, presented in Table 8.5, we consider a larger number of smaller subdomains in 2D and use incomplete Cholesky decompositions for the local Neumann problems with a threshold of 10^{-2}. The local subdomain sizes remain fixed. In these experiments the incomplete Cholesky factorizations use less than 65% of the memory required for the exact factorizations. We see that the number of iterations remains bounded as the number of subdomains increases from 16 to 1 024.

9. Annotations. In general, we use a sequential direct solver in our FETI-DP production code for the factorization of $\widetilde{S}_{\Pi\Pi}$. In this production code and on the specific hardware used in this paper our application calls UMFPACK 4.3 [7] in the

Direct Solvers							
	2D				3D		
N	64	256	1 024	4 096	64	512	4 096
1/H	8	64	32	64	4	8	16
H/h	8	8	8	8	4	4	4
$\dim(\widetilde{S}_{\Pi\Pi})$	322	1 410	5 890	24 066	324	3 528	32 400
FETI-DP							
CG							
It.	**10**	**11**	**11**	**10**	**14**	**15**	**15**
$\lambda_{\min}$	1.008	1.011	1.007	1.007	1.029	1.026	1.022
$\lambda_{\max}$	2.219	2.344	2.348	2.342	4.107	4.064	4.062
iFETI-DP							
GMRES							
It.	**8**	**8**	**7**	**6**	**14**	**13**	**12**
CG (sc.)							
It.	**10**	**10**	**9**	**9**	**15**	**14**	**13**
$\lambda_{\min}$	1.001	1.001	1.004	1.002	1.020	1.025	1.026
$\lambda_{\max}$	2.219	2.341	2.342	2.341	4.107	4.065	4.063
irFETI-DP							
GMRES							
It.	**8**	**8**	**7**	**6**	**14**	**13**	**12**
CG (sc.)							
It.	**10**	**10**	**9**	**9**	**14**	**13**	**12**
$\lambda_{\min}$	1.001	1.004	1.001	1.001	1.029	1.029	1.030
$\lambda_{\max}$	2.220	2.341	2.344	2.333	4.107	4.066	4.064

TABLE 8.1

Exact solvers: Comparison of standard FETI-DP with the inexact variants, denoted inexact FETI-DP (iFETI-DP) and inexact FETI-DP on the reduced system (irFETI-DP). Here we use direct solvers for the local subdomain problems and the coarse grid problem. The GMRES iteration count is given for left preconditioning. For the CG accelerated method we use scaling (irFETI-DP: $\hat{S}^{-1}_{\Pi\Pi} = 0.9999^2 \cdot \widetilde{S}^{-1}_{\Pi\Pi}$; iFETI-DP: $\hat{S}^{-1}_{\Pi\Pi} = 0.99999^2 \cdot \widetilde{S}^{-1}_{\Pi\Pi}$, $\hat{K}^{-1}_{BB} = 0.99999^2 \cdot K^{-1}_{BB}$). The dual Schur complement F is always preconditioned by the Dirichlet preconditioner M_D^{-1}.
— 2D linear elasticity on the unit square for $N = 64$ to $N = 4\,096$ subdomains, Q1-elements, $E = 1, \nu = 0.4$, GMRES restart: 50 iterations, Algorithm B (vertex and edge averages), relative residual reduction of 10^{-7}.
— 3D linear elasticity on the unit cube for $N = 64$ to $N = 4\,096$ subdomains, P1-elements, $E = 210, \nu = 0.29$, GMRES restart: 50 iterations, Algorithm D_E (only edge averages, no vertices), relative residual reduction of 10^{-7}.

	Inexact solver for the coarse problem						
	SAMG						
	2D				3D		
N	64	256	1 024	4 096	64	512	4 096
1/H	8	64	32	64	4	8	16
H/h	8	8	8	8	4	4	4
$\dim(\widetilde{S}_{\Pi\Pi})$	322	1 410	5 890	24 066	324	3 528	32 400
iFETI-DP							
GMRES							
It.	**8**	**8**	**8**	**7**	**14**	**13**	**12**
irFETI-DP							
GMRES							
It.	**8**	**8**	**7**	**7**	**14**	**13**	**12**
GMRES (sc.)							
It.	**8**	**8**	**7**	**7**	**14**	**14**	**13**	
CG (sc.)							
It.	**9**	**9**	**9**	**10**	**14**	**15**	**15**
$\lambda_{\min}$	1.013	1.021	1.035	1.037	1.032	1.029	1.028
$\lambda_{\max}$	2.258	2.419	2.472	2.545	4.276	4.523	4.864

Table 8.2

Inexact solver for the coarse grid problem: Performance of the inexact variants of FETI-DP, denoted inexact FETI-DP (iFETI-DP) and inexact FETI-DP on the reduced system (irFETI-DP). We use direct solvers for the local problems and SAMG [32, 31, 30, 34] to precondition $\widetilde{S}_{\Pi\Pi}$ using two V-cycles per outer Krylov subspace iteration. The V-cycles use one sweep of ILU(0) as pre- and post-smoother. For some of the calculations we use scaling ("sc.").
— 2D linear elasticity on the unit square for $N = 64$ to $N = 4\,096$ subdomains, Q1-elements, $E = 1, \nu = 0.4$, GMRES restart: 50 iterations, Algorithm B (vertex and edge averages), relative residual reduction of 10^{-7}.
— 3D linear elasticity on the unit cube for $N = 64$ to $N = 4\,096$ subdomains, P1-elements, $E = 210, \nu = 0.29$, GMRES restart: 50 iterations, Algorithm D_E (only edge averages, no vertices), relative residual reduction of 10^{-7}.

Fig. 8.1. *Mechanical part courtesy of GETRAG FORD Transmissions GmbH, Cologne, Germany.*

Inexact solver for the coarse problem (3D)								
				irFETI-DP / SAMG				
				GMRES	GMRES (sc.)	CG (sc.)		
N	1/H	H/h	$\dim(\widetilde{S}_{\Pi\Pi})$	It.	It.	It.	$\lambda_{\min}$	$\lambda_{\max}$
13 824	24	4	114 264	**12**	**12**	**14**	1.0302	4.731

Table 8.3

A problem with a large number of subdomains: Inexact FETI-DP on the reduced system (irFETI-DP). Exact solvers for the local problem and SAMG [32, 31, 30, 34] using two V-cycles per outer Krylov subspace iteration to precondition $\widetilde{S}_{\Pi\Pi}$. The V-cycles use one sweep of ILU(0) as pre- and post-smoother; scaling was used for two of the calculations ("sc.").
— 3D linear elasticity on the unit cube, $N = 13\,824$ ($24\times24\times24$) subdomains, 192 ($4^3\times3$) d.o.f. per subdomain, 1 167 051 total d.o.f.; P1-elements, $E = 210, \nu = 0.29$, GMRES restart: 50 iterations, Algorithm D_E (edge averages), relative residual reduction of 10^{-7}.

Inexact solvers for the local problems 3D, icc(1e-4)			
	N	8	64
	1/H	2	4
	H/h	14	14
FETI-DP CG			
	It.	**18**	**23**
	$\lambda_{\min}$	1.03	1.03
	$\lambda_{\max}$	12.94	10.45
iFETI-DP GMRES			
	It.	**22**	**22**

Table 8.4

Inexact solver for the local Neumann problems: Comparison of standard, exact FETI-DP with inexact FETI-DP (iFETI-DP) using incomplete Cholesky with a threshold of 10^{-4} for the local Neumann problems and a direct solver for the coarse grid problem. In these experiments the incomplete Cholesky factorization uses about half of the memory that is required for the total factorization.
— 3D linear elasticity on the unit cube for $N = 8$ and $N = 64$ subdomains, 8 232 ($=14^3\times3$) d.o.f. per subdomain, 59 049 and 446 631 total d.o.f., respectively; P1-elements, $E = 210, \nu = 0.29$, GMRES restart: 50 iterations, Algorithm D_E (edge averages), relative residual reduction of 10^{-7}.

	Inexact solvers for the local problems 2D, icc(1e-2)								
	N	16	36	64	100	256	400	768	1 024
iFETI-DP GMRES									
	It.	14	14	13	13	12	12	11	11

Table 8.5

Inexact solver for the local Neumann problems: Comparison of standard FETI-DP with inexact FETI-DP (iFETI-DP) using incomplete Cholesky with a threshold of 10^{-2} *for the local Neumann problems and a direct solver for the coarse grid problem. In these experiments the incomplete Cholesky factorization uses less than two thirds of the memory that is required for the total factorization.*

— 2D linear elasticity on the unit cube for $N = 4$ *to* $N = 1\,024$ *subdomains,* 128 *(*$= 8^2 \times 2$*) d.o.f. per subdomain. Q1-elements,* $E = 1, \nu = 0.4$*, GMRES restart:* 50 *iterations, Algorithm B (vertices and edge averages), relative residual reduction of* 10^{-7}*.*

	Mechanical Part		
	N	1 024	2 048
FETI-DP CG			
	It.	**41**	**48**
	$\lambda_{\min}$	1.04	1.04
	$\lambda_{\max}$	33.31	44.63
irFETI-DP / SAMG GMRES			
	It.	**53**	**64**

Table 8.6

Mechanical part, see Figure 8.1, 841 836 d.o.f., partitioned into $1\,024$ *and* $2\,048$ *subdomains. The coarse problem has a size of* $10\,380$ *and* $19\,515$ *d.o.f., respectively. We use* 4 *V cycles of SAMG [32, 31, 30, 34] in each outer Krylov subspace iteration. Two Gauss-Seidel sweeps as pre- and two as post-smoother are used. The Gauss-Seidel sweeps use CF-ordering for pre-smoothing and the reverse for post-smoothing.*

— 3D linear elasticity on the unit cube for $N = 1\,024$ *and* $N = 2\,048$ *subdomains, P1-elements,* $E = 210, \nu = 0.29$*, GMRES restart:* 100 *iterations, Algorithm* D_E *(edge averages), relative residual reduction of* 10^{-7}*.*

(faster) 32 bit integer and 32 bit pointer mode. So even though on our 64 bit hardware the application is not subject to the 2 GB limit the factors of every single matrix cannot grow larger than 2 GB. This is typically not a limitation except for very large subdomains or very large coarse matrices as one of coarse matrices considered for this paper. The calculations were carried out on our 16 processer (8 dual Opteron 248, 2.2 Ghz) Opteron cluster in Essen. Some of the calculations were carried out in parallel using PETSc; see [2], [1], [3]. The algebraic multigrid solver SAMG always ran in parallel, using two threads.

Acknowledgments. The authors would like to thank Tanja Clees, Klaus Stüben, and the Fraunhofer Institute for Algorithms and Scientific Computing (SCAI), Sankt Augustin, Germany, for providing SAMG and also Holger Hein, Hans-Jürgen Philipsenburg, and GETRAG FORD Transmissions GmbH, Cologne, Germany.

REFERENCES

[1] Satish Balay, Kris Buschelman, Victor Eijkhout, William D. Gropp, Dinesh Kaushik, Matthew G. Knepley, Lois Curfman McInnes, Barry F. Smith, and Hong Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 2.1.5, Argonne National Laboratory, 2004.

[2] Satish Balay, Kris Buschelman, William D. Gropp, Dinesh Kaushik, Matt Knepley, Lois Curfman McInnes, Barry F. Smith, and Hong Zhang. PETSc home page. 2001.

[3] Satish Balay, Victor Eijkhout, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press, 1997.

[4] James H. Bramble and Joseph E. Pasciak. A preconditioning technique for indefinite systems resulting from mixed approximations of elliptic problems. *Mathematics of Computation*, 50(181):1–17, 1988.

[5] Philippe G. Ciarlet. *Mathematical Elasticity Volume I: Three–Dimensional Elasticity*. North-Holland, 1988.

[6] Jean-Michel Cros. A preconditioner for the Schur complement domain decomposition method. In O. Widlund I. Herrera, D. Keyes and R. Yates, editors, *Domain Decomposition Methods in Science and Engineering*, pages 373–380. National Autonomous University of Mexico (UNAM), Mexico City, Mexico, ISBN 970-32-0859-2, 2003. Proceedings of the 14th International Conference on Domain Decomposition Methods in Science and Engineering; http://www.ddm.org/DD14.

[7] Timothy A. Davis. A column pre-ordering strategy for the unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software*, 30(2):165–195, June 2004.

[8] Clark Dohrmann and Richard B. Lehoucq. A primal based penalty preconditioner for elliptic saddle point systems. Technical Report SAND 2004-5964J, Sandia National Laboratories, 2004.

[9] Clark R. Dohrmann. A preconditioner for substructuring based on constrained energy minimization. *SIAM J. Sci.Comput.*, 25(1):246–258, 2003.

[10] Charbel Farhat, Michel Lesoinne, Patrick LeTallec, Kendall Pierson, and Daniel Rixen. FETI-DP: A dual-primal unified FETI method - part i: A faster alternative to the two-level FETI method. *Int. J. Numer. Meth. Engrg.*, 50:1523–1544, 2001.

[11] Charbel Farhat, Michel Lesoinne, and Kendall Pierson. A scalable dual-primal domain decomposition method. *Numer. Lin. Alg. Appl.*, 7:687–714, 2000.

[12] Pierre Gosselet. *Méthodes de décomposition de domaine et méthodes d'accéleration pour les problèmes multichamps en non-linéaire*. PhD thesis, Université Paris 6, 2003.

[13] George Karypis, Kirk Schloegel, and Vipin Kumar. ParMetis - parallel graph partitioning and sparse matrix ordering, version 3.1. Technical report, University of Minnesota, Department of Computer Science and Engineering, August 2003.

[14] Axel Klawonn. *Preconditioners for Indefinite Problems*. PhD thesis, Westfälische Wilhelms-Universität Münster, 1996.

[15] Axel Klawonn. Block-Triangular Preconditioners for Saddle Point Problems with a Penalty Term. *SIAM J. Sci. Comp.*, 19(1):172–184, January 1998.

[16] Axel Klawonn and Oliver Rheinbach. A parallel implementation of Dual-Primal

FETI methods for three dimensional linear elasticity using a transformation of basis. Technical Report SM-E-601, Department of Mathematics, Universität Duisburg-Essen, Germany, February 2005. Schriftenreihe des Fachbereichs Mathematik; http://www.uni-essen.de/ingmath/Axel.Klawonn/TR-SM-E-601.pdf.

[17] Axel Klawonn, Oliver Rheinbach, and Olof B. Widlund. Some computational results for Dual-Primal FETI methods for three dimensional elliptic problems. In R. Kornhuber, R.H.W. Hoppe, D.E. Keyes, J. Périaux, O. Pironneau, and J. Xu, editors, *Domain Decomposition Methods*. Springer-Verlag, Lecture Notes in Computational Science and Engineering, 2003. Proceedings of the 15th International Conference on Domain Decomposition Methods, Berlin, July 21-25, 2003.

[18] Axel Klawonn and Olof B. Widlund. A domain decomposition method with Lagrange multipliers and inexact solvers for linear elasticity. *SIAM J. Sci. Comput.*, 22(4):1199–1219, 2000.

[19] Axel Klawonn and Olof B. Widlund. FETI and Neumann–Neumann iterative substructuring methods: Connections and new results. *Comm. Pure Appl. Math.*, 54:57–90, January 2001.

[20] Axel Klawonn and Olof B. Widlund. Dual-Primal FETI methods for linear elasticity. Technical Report 2004-855, Courant Institute of Mathematical Sciences, New York University, New York, USA, September 2004.

[21] Axel Klawonn, Olof B. Widlund, and Maksymilian Dryja. Dual-Primal FETI methods for three-dimensional elliptic problems with heterogeneous coefficients. *SIAM J.Numer.Anal.*, 40, 159-179 2002.

[22] Jing Li and Olof B. Widlund. FETI-DP, BDDC, and Block Cholesky Methods. Technical Report TR2004-857, Department of Computer Science, Courant Institute of Mathematical Sciences, New York University, December 2004. http://cs.nyu.edu/csweb/Research/TechReports/TR2004-857/TR2004-857.pdf.

[23] Jing Li and Olof B. Widlund. On the use of inexact subdomain solvers for BDDC algorithms. Technical Report TR2005-871, Department of Computer Science, Courant Institute of Mathematical Sciences, New York University, July 2005.

[24] Jan Mandel and Clark R. Dohrmann. Convergence of a balancing domain decomposition by constraints and energy minimization. *Numer. Lin. Alg. Appl.*, 10:639–659, 2003.

[25] Jan Mandel, Clark R. Dohrmann, and Radek Tezaur. An algebraic theory for primal and dual substructuring methods by constraints. *Appl. Numer. Math.*, 2004. Sixth IMACS International Symposium on Iterative Methods in Scientific Computing, Denver, Colorado, USA, 2003. To appear.

[26] Jan Mandel and Radek Tezaur. On the convergence of a dual-primal substructuring method. *Numer. Math.*, 88:543–558, 2001.

[27] Kendall H. Pierson. *A family of domain decomposition methods for the massively parallel solution of computational mechanics problems*. PhD thesis, University of Colorado at Boulder, Aerospace Engineering, 2000.

[28] Valeria Simoncini. Block triangular preconditioners for symmetric saddle-point problems. *Applied Numerical Mathematics*, 49:63–80, 2004.

[29] Barry F. Smith, Petter Bjørstad, and William Gropp. *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*. Cambridge University Press, New York, 1996.

[30] Klaus Stüben. An Introduction to Algebraic Multigrid. In *Multigrid* [34], pages 413–532. Also available as GMD Report 70, November 1999.

[31] Klaus Stüben. A review of algebraic multigrid. *J. Comp. Appl. Math.*, 128:281–309, 2001. Also available as GMD Report 69, November 1999.

[32] Klaus Stüben and Tanja Clees. *SAMG User's Manual, Release 22c (May 2005)*. Fraunhofer SCAI, Sankt Augustin, Germany, May 2005. http://www.scai.fraunhofer.de/samg.htm.

[33] Andrea Toselli and Olof B. Widlund. *Domain Decomposition Methods - Algorithms and Theory*, volume 34 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin Heidelberg New York, 2005.

[34] Ulrich Trottenberg, Cornelis W. Oosterlee, and Anton Schüller. *Multigrid*. Academic Press, London, San Diego, 2001.

[35] Xuemin Tu. Three-level BDDC in two dimensions. Technical Report TR2004-856, Department of Computer Science, Courant Institute, New York University, 2004.

[36] Xuemin Tu. Three-level BDDC in three dimensions. Technical Report TR2005-862, Department of Computer Science, Courant Institute, New York University, 2005.