Parallel Sums and Adaptive BDDC Deluxe

Olof B. Widlund[1] and Juan G. Calvo[2]

1 Introduction

There has recently been a considerable activity in developing adaptive methods for the selection of primal constraints for BDDC algorithms and, in particular, for BDDC deluxe variants. The primal constraints of a BDDC or FETI–DP algorithm provide the global, coarse part of such a preconditioner and are of crucial importance for obtaining rapid convergence of these preconditioned conjugate gradient methods for the case of many subdomains. When the primal constraints are chosen adaptively, we aim at selecting a primal space, which for a certain dimension of the coarse space, provides the fastest rate of the convergence for the iterative method. In the alternative, we can try to develop criteria which will guarantee that the condition number of the iteration stays below a given tolerance.

A particular inspiration for our own work has been a talk, see Dohrmann and Pechstein [2012], by Clark Dohrmann at DD21, held in Rennes, France, in June 2012. Dohrmann had then started joint work with Clemens Pechstein, see also Pechstein and Dohrmann [2016].

Much of this work for BDDC and FETI-DP iterative substructuring algorithms, which has been supported by theory, has been confined to developing primal constraints for equivalence classes with two elements such as those related to subdomain edges for problems defined on domains in the plane; see a recent survey paper, Klawonn et al. [2016b]. In our context, the equivalence classes are sets of finite element nodes which belong to the boundaries of more than one subdomain with the equivalence relation defined by the sets of subdomain boundaries to which the nodes belong. While it is important to further study the best way of handling all cases, the basic issues appear to be well settled when the equivalence classes all have just two elements.

Courant Institute, 251 Mercer Street, New York, NY 10012, USA widlund@cims.nyu.edu ·
CIMPA, Universidad de Costa Rica, San Jose, Costa Rica, 11501 juan.calvo@ucr.ac.cr

We note that this work is relevant for problems posed in $H(\mathrm{div})$ even in three dimensions (3D) since the degrees of freedom on the interface between subdomains for Raviart-Thomas and Brezzi-Douglas-Marini elements are associated only with faces of the elements, see Oh et al. [2015], Zampini [2016]. (These papers also concern BDDC three–level algorithms choosing two levels of primal constraints adaptively.) But for other elliptic problems in 3D, there is a need to develop algorithms and results for equivalence classes with three or more elements.

There is early work by Mandel, Šístek, and Sousedík, who developed condition number indicators, cf. Mandel and Sousedík [2007], Mandel et al. [2012]. Talks by Clark Dohrmann and Axel Klawonn at DD23, held on Jeju Island, the Republic of Korea in July 2015, see Klawonn et al. [2016a], reported on recent progress to give similar algorithms a firm theoretical basis. A talk by Hyea Hyun Kim in the same mini-symposium also reported considerable progress for a different kind of algorithm. Her main new algorithm for problems in three dimensions is similar but not the same as ours; see further Kim et al. [2015]. Our main result, developed independently, was reported on by the first author in the same mini-symposium; see further Calvo and Widlund [2016] and, for applications to isogeometric analysis, Beirão da Veiga et al. [2015].

This paper will focus on using parallel sums for general equivalence classes. Such an approach for equivalence classes with two elements has proven very successful in simplifying the formulas and arguments; see in particular Pechstein and Dohrmann and section 2. Parallel sums for equivalence classes with more than two elements have also been quite successfully in numerical experiments by Simone Scacchi and Stefano Zampini, reported in Beirão da Veiga et al. [2015], for problems arising in isogeometric analysis and also by Zampini in a study of 3D problems formulated in $H(\mathbf{curl})$, based in part on Dohrmann and Widlund [2016], and reported on in this mini-symposium.

In this paper, we will focus on low order, nodal finite element approximations for scalar elliptic problems in three dimensions,

$$-\nabla \cdot (\rho(x)\nabla u) = f(x), \quad x \in \Omega, \quad \rho(x) > 0, \tag{1}$$

resulting in a linear system of equations to be solved using BDDC domain decomposition algorithms, especially its deluxe variant. We will always assume that the choice of boundary conditions results in a positive definite, symmetric stiffness matrix.

2 Equivalence classes and BDDC algorithms

BDDC algorithms, see, e.g., Li and Widlund [2006], are domain decomposition algorithms based on the decomposition of the domain Ω of an elliptic operator into non-overlapping subdomains Ω_i, each often associated with tens of thousands of degrees of freedom. The subdomain interface Γ_i of Ω_i does not cut through any elements and is defined by $\Gamma_i := \partial\Omega_i \setminus \partial\Omega$. The equivalence classes are associated with the subdomain faces, edges, and vertices of $\Gamma := \cup_i \Gamma_i$, the interface of the entire decomposition. Thus, for a problem in three dimensions, a subdomain face is associated with the degrees of freedom of the nodes belonging to the interior of the intersection of two boundaries of two neighboring subdomains Ω_i and Ω_j. Those of a subdomain edge are typically associated with a set of nodes common to three or more subdomain boundaries, while the endpoints of the subdomain edges are the subdomain vertices which are associated with even more subdomains.

Given the stiffness matrix $A^{(i)}$ of the subdomain Ω_i, we obtain a subdomain Schur complement $S^{(i)}$ by eliminating the interior variables, i.e., all those that do not belong to Γ_i. We will also work with principal minors of these Schur complements associated with faces, F, and edges, E, denoting them by $S^{(i)}_{FF}$ and $S^{(i)}_{EE}$, respectively.

The interface space is divided into a *primal* subspace of functions which are continuous across Γ and a complementary, *dual* subspace for which we will allow multiple values across the interface during part of the iteration. In our study, all the subdomain vertex variables will always belong to the primal set. We have three product spaces of finite element functions/vectors defined by their interface nodal values:

$$\widehat{W}_\Gamma \subset \widetilde{W}_\Gamma \subset W_\Gamma.$$

W_Γ is a product space without any continuity constraints across the interface. Elements of $\widetilde{W}_\Gamma$ have common values of the primal variables but allow multiple values of the dual variables while the elements of $\widehat{W}_\Gamma$ are continuous at all nodes on Γ. We will change variables, explicitly introducing the primal variables and a complementary sets of dual variables in order to simplify the presentations. We note that the change of basis will not in any way change the results of the computation. After eliminating the interior variables, we can then write the subdomain Schur complements as

$$S^{(i)} = \begin{pmatrix} S^{(i)}_{\Delta\Delta} & S^{(i)}_{\Delta\Pi} \\ S^{(i)}_{\Pi\Delta} & S^{(i)}_{\Pi\Pi} \end{pmatrix}.$$

We will partially subassemble the $S^{(i)}$, obtaining $\widetilde{S}$, enforcing the continuity of the primal variables only. Thus, we then work in $\widetilde{W}_\Gamma$. In each step of the iteration, we solve a linear system with the coefficient matrix $\widetilde{S}$. Solving these linear systems will be considerably much faster than if we work with the fully assembled system if the dimension of the primal space is modest. At the end of each iteration, the approximate solution is made continuous at all nodal points of the interface by applying a weighted averaging operator E_D. We always accelerate the iteration with the preconditioned conjugate gradient algorithm.

BDDC deluxe. When designing a BDDC algorithm, we have to choose an effective set of primal constraints and also a good recipe for the averaging across the interface. Our paper concerns the choice of the primal constraints while we will always use the deluxe recipe in the construction of the averaging operator E_D.

We note that in work on three-dimensional problems formulated in $H(\mathbf{curl})$, it was found that traditional averaging recipes did not always work well; cf. Dohrmann and Widlund [2016]. The same is true for problems in $H(\text{div})$; see Oh et al. [2015]. This occasional failure has its roots in the fact that there are two sets of material parameters in these applications. The deluxe scaling that was then introduced has also proven quite successful for a variety of other applications.

A face component of the average operator E_D across a subdomain face $F \subset \Gamma$, common to two subdomains Ω_i and Ω_j, is defined in terms of principal minors $S_{FF}^{(k)}$ of the $S^{(k)}, k = i, j$:

$$\bar{w}_F := (E_D w)_F := (S_{FF}^{(i)} + S_{FF}^{(j)})^{-1}(S_{FF}^{(i)} w_F^{(i)} + S_{FF}^{(j)} w_F^{(j)}).$$

Here $w_F^{(i)}$ is the restriction of $w^{(i)}$ to the face F, etc.

Deluxe averaging operators are also developed for subdomain edges and the operator E_D is assembled from all these components; see further section 3. Our bound for this operator will be obtained from bounds for certain eigenvalues for the individual equivalence sets and will include factors that depend quadratically on the number of equivalence classes associated with the faces and edges of the individual subdomains. We have found that the performance consistently is far better than these bounds.

The core of any estimate for a BDDC algorithm is the norm of the averaging operator E_D. By an algebraic argument known, for FETI–DP since 2002, we know that the condition number of the iteration satisfies

$$\kappa(M_{BDDC}^{-1}\widehat{S}) \leq \|E_D\|_{\widetilde{S}}; \tag{2}$$

see Klawonn et al. [2002]. Here M_{BDDC}^{-1} denotes the BDDC preconditioner and $\widehat{S}$ the fully assembled Schur complement of the problem. Instead of developing an estimate for E_D, we will work with $P_D := I - E_D$ and estimate $(R_F^T(w_F^{(i)} - \bar{w}_F))^T S^{(i)} R_F^T(w_F^{(i)} - \bar{w}_F)$. Here R_F denotes the restriction to the face F. We find, following Pechstein, that the sum of this quadratic form and a similar contribution from the neighboring subdomain Ω_j equals

$$(w_F^{(i)} - w_F^{(j)})^T (S_{FF}^{(i)} : S_{FF}^{(j)})(w_F^{(i)} - w_F^{(j)})$$

where

$$A : B := A(A+B)^{-1}B$$

is the parallel sum of A and B; cf. Anderson Jr. and Duffin [1969]. We note that if A and B are positive definite, then $A : B = (A^{-1} + B^{-1})^{-1}$. If $A+B$ is only positive semi-definite, we can replace $(A+B)^{-1}$ by $(A+B)^{\dagger}$, any generalized inverse. The quadratic form can be estimated from above by

$$2(w_F^{(i)} - w_{F\Pi})^T (S_{FF}^{(i)} : S_{FF}^{(j)})(w_F^{(i)} - w_{F\Pi}) + 2(w_F^{(j)} - w_{F\Pi})^T (S_{FF}^{(i)} : S_{FF}^{(j)})(w_F^{(j)} - w_{F\Pi})$$

where $w_{F\Pi}$ is the restriction of an arbitrary element of the primal space to the face. We note that each of these terms can be estimated by an expression which is local to only one subdomain.

With $w_{F\Delta}^{(i)} := w_F^{(i)} - w_{F\Pi}$, we now estimate $w_{F\Delta}^{(i)T}(S_{FF}^{(i)} : S_{FF}^{(j)})w_{F\Delta}^{(i)}$ by the energy of $w^{(i)}$. We then need the minimum norm extension of any finite element function defined on F, which will provide a uniform bound for any extension of the values on F to the rest of Γ_i. We find that the relevant matrix is

$$\widetilde{S}_{FF}^{(i)} := S_{FF}^{(i)} - S_{F'F}^{(i)T} S_{F'F'}^{(i)-1} S_{F'F}^{(i)}.$$

Here $S_{F'F'}^{(i)}$ is the principal minor of $S^{(i)}$ with respect to $\Gamma_i \setminus F$ and $S_{F'F}^{(i)}$ an off-diagonal block of $S^{(i)}$. By appropriate choices of the primal space and of $w_{F\Pi}$, we are able to show that

$$w_{F\Delta}^{(i)T}(\widetilde{S}_{FF}^{(i)} : \widetilde{S}_{FF}^{(j)})w_{F\Delta}^{(i)} \leq w^{(i)T} S^{(i)} w^{(i)},$$

where $w^{(i)}$ is an arbitrary extension of the values of $w_F^{(i)}$.

For an adaptive algorithm, we can complete the estimate by using a generalized eigenvalue problem:

$$\widetilde{S}_{FF}^{(i)} : \widetilde{S}_{FF}^{(j)} \phi = \lambda S_{FF}^{(i)} : S_{FF}^{(j)} \phi. \tag{3}$$

Primal constraints are then generated by using the eigenvectors of a few of the smallest eigenvalues of (3) and making $(\widetilde{S}_{FF}^{(i)} : \widetilde{S}_{FF}^{(j)})(w_F^{(i)} - w_F^{(j)})$ orthogonal to these eigenvectors.

A bound can now be obtained in terms of the smallest eigenvalue associated with the eigenvectors not used in deriving the primal constraints. Numerical studies show a very rapid decay of the eigenvalues of $S_{FF}^{(i)-1}(S_{FF}^{(i)} - \widetilde{S}_{FF}^{(i)})$; this property can also be proven assuming that Ω_i is Lipschitz and the coefficient $\rho(x)$ a constant. Therefore only a few primal constraints will greatly improve the bound on the norm of $(E_D w)_F$.

3 Equivalence classes with more than two elements

We begin this section by considering parallel sums of more than two operators. We will work with symmetric matrices which all are at least positive semi-definite. For three positive definite matrices, we can define their parallel sum by

$$A : B : C := (A^{-1} + B^{-1} + C^{-1})^{-1},$$

with similar formulas for four or more matrices. A quite complicated formula for $A : B : C$ is given in Tian [2002] for the general case when some or all of the matrices might be only positive semi-definite. It is also shown, in [Tian, 2002, Theorem 3], that $A : B : C = (A^\dagger + B^\dagger + C^\dagger)^\dagger$ if and only if the three operators A, B, and C have the same range. In our context, this is not always the case since the matrix $\widetilde{S}^{(i)}_{EE}$, defined below, will be singular if Ω_i is an interior subdomain while it will be non-singular if $\partial\Omega_i$ intersects a part of $\partial\Omega$ where a Dirichlet condition is imposed. This issue can be avoided by making all operators non-singular by adding a small positive multiple of the identity to the singular operators.

We will first focus on a case of an equivalence class common to three subdomains as arising for most subdomain edges in a three-dimensional finite element context if the subdomains are generated using a mesh partitioner. We will use the notation $S^{(i)}_{EE}$, $S^{(j)}_{EE}$, and $S^{(k)}_{EE}$ for the principal minors, of the degrees of freedom of an edge E, of the subdomain Schur complements of the three subdomains that have this subdomain edge in common. The Schur complements of the Schur complements representing the minimal energy extensions to individual subdomains, of given values on the subdomain edge E, will be denoted by $\widetilde{S}^{(i)}_{EE}$, $\widetilde{S}^{(j)}_{EE}$, etc., and are defined by

$$\widetilde{S}^{(i)}_{EE} := S^{(i)}_{EE} - S^{(i)T}_{E'E} S^{(i)-1}_{E'E'} S^{(i)}_{E'E}. \tag{4}$$

Here $S^{(i)}_{E'E'}$ is the principal minor of $S^{(i)}$ of $\Gamma_i \setminus E$ and $S^{(i)}_{E'E}$ an off-diagonal block.

We can now introduce the deluxe average over the edge E by

$$\bar{w}_E := (S^{(i)}_{EE} + S^{(j)}_{EE} + S^{(k)}_{EE})^{-1} (S^{(i)}_{EE} w^{(i)}_E + S^{(j)}_{EE} w^{(j)}_E + S^{(k)}_{EE} w^{(k)}_E).$$

By using elementary inequalities, we can now obtain a bound of the square of the norm of an edge component of $P_D w$ by

$$3 w^{(i)T}_{E\Delta} S^{(i)}_{EE} : (S^{(j)}_{EE} + S^{(k)}_{EE}) w^{(i)}_{E\Delta}$$

and two similar terms obtained by changing the superscripts appropriately.

Returning to the search for adaptive primal spaces, we note that ideally, we would now like to prove that the three operators $T^{(i)}_E := S^{(i)}_{EE} : (S^{(j)}_{EE} + S^{(k)}_{EE})$, $T^{(j)}_E := S^{(j)}_{EE} : (S^{(i)}_{EE} + S^{(k)}_{EE})$, and $T^{(k)}_E := S^{(k)}_{EE} : (S^{(i)}_{EE} + S^{(j)}_{EE})$ all can be bounded uniformly from above by

$$S^{(i)}_{EE} : S^{(j)}_{EE} : S^{(k)}_{EE} := (S^{(i)-1}_{EE} + S^{(j)-1}_{EE} + S^{(k)-1}_{EE})^{-1}. \tag{5}$$

If this were possible, we could use that same matrix for estimates for $w_{E\Delta}^{(i)}, w_{E\Delta}^{(j)}$, and $w_{E\Delta}^{(k)}$; we could use arguments very similar to those of the previous section. But we are not that lucky; good bounds are only possible if $S_{EE}^{(i)}, S_{EE}^{(j)}$, and $S_{EE}^{(k)}$ are spectrally equivalent with good bounds. However, it is easy to find interesting examples where this does not hold. We therefore have to find a different common upper bound for $T_E^{(i)}$, $T_E^{(j)}$, and $T_E^{(k)}$ and accomplish this by using the trivial inequality

$$T_E^{(i)} \leq T_E^{(i)} + T_E^{(j)} + T_E^{(k)},$$

and define our generalized eigenvalue problem as

$$(\widetilde{S}_{EE}^{(i)} : \widetilde{S}_{EE}^{(j)} : \widetilde{S}_{EE}^{(k)})\phi = \lambda(T_E^{(i)} + T_E^{(j)} + T_E^{(k)})\phi. \quad (6)$$

We note that these arguments extend directly to equivalence classes with more than three elements.

This is the recipe that we have used in most of our numerical experiments, which have proven quite successful; cf. Calvo and Widlund [2016] for many more details. However, it deserves to be noted that the distribution of the eigenvalues associated with the subdomain edges, in our experience, is less favorable than those of the subdomain faces but that we can benefit from the fact that the number of degrees of freedom of an edge typically is much smaller than that of a face.

Given the success, by others, with using parallel sums of each of the two sets of three Schur complements, we have also carried out experiments with that alternative generalized eigenvalue problem. The performance is very similar to that of our algorithm.

In our experiments, we have compared the performance of our adaptive algorithms with standard choices of the primal spaces. In choosing our primal constraints, we have, in some of our experiments, used tolerances introduced in Kim et al. [2015]. We have found that our adaptive algorithm also works quite well for irregular subdomains generated by the METIS mesh partitioner.

References

William N. Anderson Jr. and Richard J. Duffin. Series and parallel addition of matrices. *J. Math. Anal. Appl.*, 26:576–594, 1969.

Lourenço Beirão da Veiga, Luca F. Pavarino, Simone Scacchi, Olof B. Widlund, and Stefano Zampini. Adaptive selection of primal constraints for isogeometric BDDC deluxe preconditioners. Technical Report TR2015-977, Courant Institute, New York University, 2015.

Juan G. Calvo and Olof B. Widlund. An adaptive choice of primal constraints for BDDC domain decomposition algorithms. TR2015-979, Courant Institute, New York University, 2016.

Clark R. Dohrmann and Clemens Pechstein. Constraint and weight selection algorithms for BDDC. Slides for a talk by Dohrmann at DD21 in Rennes, France, June 2012. URL=http://www.numa.uni-linz.ac.at/ clemens/dohrmann-pechstein-dd21-talk.pdf, 2012.

Clark R. Dohrmann and Olof B. Widlund. A BDDC algorithm with deluxe scaling for three-dimensional H(curl) problems. *Comm. Pure Appl. Math.*, 69(4):745–770, 2016.

Hyea Hyun Kim, Eric T. Chung, and Junxian Wang. BDDC and FETI-DP algorithms with adaptive coarse spaces for three-dimensional elliptic problems with oscillatory and high contrast coefficients. http://arxiv.org/abs/1606.07560, 2015.

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(1):159–179, April 2002.

Axel Klawonn, Martin Kühn, and Oliver Rheinbach. Adaptive coarse spaces for FETI–DP in three dimensions. *SIAM J. Sci. Comput.*, 2016a. To appear.

Axel Klawonn, Patrick Radtke, and Oliver Rheinbach. A comparison of adaptive coarse spaces for iterative substructuring in two dimensions. *ETNA*, 46:75–106, 2016b.

Jing Li and Olof B. Widlund. FETI–DP, BDDC, and Block Cholesky Methods. *Internat. J. Numer. Methods Engrg.*, 66(2):250–271, 2006.

Jan Mandel and Bedřich Sousedík. Adaptive selection of face coarse degrees of freedom in the BDDC and the FETI-DP iterative substructuring methods. *Comput. Methods Appl. Mech. Engrg.*, 196(8):1389–1399, 2007.

Jan Mandel, Bedřich Sousedík, and Jakub Šístek. Adaptive BDDC in three dimensions. *Math. Comput. Simulation*, 82(10):1812–1831, 2012.

Duk-Soon Oh, Olof B. Widlund, Stefano Zampini, and Clark R. Dohrmann. BDDC algorithms with deluxe scaling and adaptive selection of primal constraints for Raviart-Thomas vector fields. Technical report, Courant Institute, New York University, 2015. TR2015-978.

Clemens Pechstein and Clark R. Dohrmann. Modern domain decomposition methods, BDDC, deluxe scaling, and an algebraic approach. Talk by Pechstein in Linz, Austria, December 2013, URL: http://people.ricam.oeaw.ac.at/c.pechstein/pechstein-bddc2013.pdf.

Clemens Pechstein and Clark R. Dohrmann. A Unified Framework for Adaptive BDDC. Technical Report 2016-20, Johann Radon Institute for Computational and Applied Mathematics (RICAM), 2016. URL:http://www.ricam.oeaw.ac.at/files/reports/16/rep16-20.pdf.

Yongge Tian. How to express a parallel sum of k matrices. *J. Math. Anal. Appl.*, 266(2):333–341, 2002.

Stefano Zampini. PCBDDC: a class of robust dual-primal preconditioners in PETSc. *SIAM J. Sci. Comput.*, 2016. To appear.