BDDC Deluxe Domain Decomposition Algorithms

Olof B. Widlund, Courant Institute, New York University

Clark R. Dohrmann, Sandia National Laboratories
Albuquerque, New Mexico

DD22, Lugano

Problems considered

- BDDC domain decomposition algorithms for finite element approximations for a variety of elliptic problems with very many degrees of freedom.
- Among applications, problems formulated in $H(\mathbf{curl})$, $H(\mathrm{div})$, and for Reissner-Mindlin plates.
- Mostly lowest order finite element methods for self-adjoint elliptic problems but we have also helped develop solvers for isogeometric analysis.
- All this aims at developing preconditioners for the stiffness matrices. These approximate inverses are then combined with conjugate gradients or other Krylov space methods.
- Primarily interested in hard problems with very many subdomains and to obtain convergence rates independent of that number and with rates that decrease slowly with the size of the subdomain problems. Most bounds independent of jumps in coefficients between subdomains.

BDDC, finite element meshes, and equivalence classes

- BDDC algorithms work on decompositions of the domain Ω of the elliptic problem into non-overlapping subdomains Ω_i, each often with tens of thousands of degrees of freedom. In between the subdomains the interface Γ. The local interface of Ω_i: $\Gamma_i := \partial\Omega_i \setminus \partial\Omega$. Γ does not cut any elements.
- Many of the finite element nodes are interior to individual subdomains while others belong to several subdomain interfaces.
- The nodes on Γ are partitioned into equivalence classes of sets of indices of the local interfaces Γ_i to which they belong. For 3D and nodal finite elements, we have classes of face nodes, associated with two local interfaces, and classes of edge nodes and subdomain vertex nodes.
- For $H(\mathbf{curl})$ and Nédélec (edge) elements, only equivalence classes of element edges on subdomain faces and on subdomain edges. For $H(\text{div})$ and Raviart-Thomas elements, only degrees of freedom for element faces.

Partial assembly

- These equivalence classes play a central role in the design, analysis, and programming of domain decomposition methods.
- The BDDC (Balancing Domain Decomposition by Constraints) algorithms introduced by Dohrmann in 2003, following the introduction of FETI–DP by Farhat et al in 2000. These two families are related algorithmically and have a common theoretical foundation.
- These preconditioners are based on using partially subassembled stiffness matrices assembled from the subdomain stiffness matrices $A^{(i)}$. We will first look at nodal finite element problems.
- The nodes of $\Omega_i \cup \Gamma_i$ are divided into those in the interior (I) and those on the interface (Γ). The interface set is further divided into a primal set (Π) and a dual set (Δ).

Torn 2D scalar elliptic problem

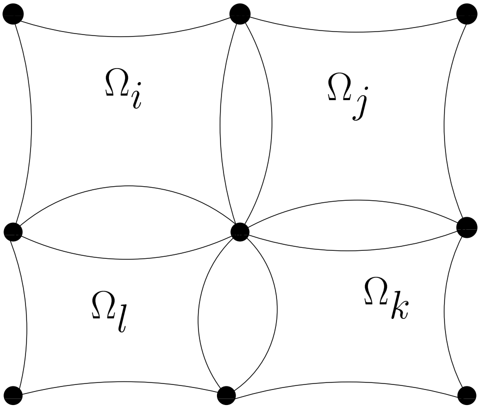

Subdomain matrices

- Represent the subdomain stiffness matrix $A^{(i)}$ as

$$\begin{pmatrix} A_{II}^{(i)} & A_{I\Delta}^{(i)} & A_{I\Pi}^{(i)} \\ A_{\Delta I}^{(i)} & A_{\Delta\Delta}^{(i)} & A_{\Delta\Pi}^{(i)} \\ A_{\Pi I}^{(i)} & A_{\Pi\Delta}^{(i)} & A_{\Pi\Pi}^{(i)} \end{pmatrix}.$$

- This matrix represents the energy contributed by Ω_i.
- We enforce continuity of the primal variables, as in the given finite element model, but allow multiple values of the dual variables when working with the partially subassembled model.

Partially subassembled matrix

- Maintain continuity of the primal variables at the vertices. Partially subassemble and mark with tilde:

$$\begin{bmatrix} A_{II}^{(1)} & A_{I\Delta}^{(1)} & & & & \tilde{A}_{\Pi I}^{(1)T} \\ A_{\Delta I}^{(1)} & A_{\Delta\Delta}^{(1)} & & & & \tilde{A}_{\Pi\Delta}^{(1)T} \\ & & \ddots & & & \vdots \\ & & & A_{II}^{(N)} & A_{I\Delta}^{(N)} & \tilde{A}_{\Pi I}^{(N)T} \\ & & & A_{\Delta I}^{(N)} & A_{\Delta\Delta}^{(N)} & \tilde{A}_{\Pi\Delta}^{(N)T} \\ \tilde{A}_{\Pi I}^{(1)} & \tilde{A}_{\Pi\Delta}^{(1)} & \cdots & \tilde{A}_{\Pi I}^{(N)} & \tilde{A}_{\Pi\Delta}^{(N)} & \tilde{A}_{\Pi\Pi} \end{bmatrix}$$

- BDDC: After solving, enforce the continuity constraints by averaging across the interface.
- FETI–DP: Use Lagrange multipliers instead. They can be interpreted as fluxes.

More on BDDC and FETI–DP

- The partially subassembled stiffness matrix of this alternative finite element model is used to define preconditioners; the resulting linear system is much cheaper to solve than the fully assembled system. Primal variables provide a global component of these preconditioners. Makes matrices invertible.
- In a FETI–DP algorithm, the continuity at the edge nodes enforced by using Lagrange multipliers and the rate of convergence enhanced by also solving a Dirichlet problem on each subdomain in each iteration. The conjugate gradient algorithm is used to find accurate enough values of Lagrange multipliers. There are subtle scaling issues.
- In a BDDC algorithm, continuity is instead restored in each step by computing a weighted average across the interface. This leads to non-zero residuals at nodes next to Γ. In each iteration use a subdomain Dirichlet solve to eliminate them.

Alternative sets of primal constraints

- For scalar second order elliptic equations, approach outlined yields condition number of $C(1+\log(H/h))^2$. Results can be made independent of jumps in the coefficients, if the interface average chosen carefully.
- Good numerical results in 2D but for competitive algorithms in 3D, certain average values (and moments) of the displacement over individual edges (and faces) should also take common values across interface Γ. Same matrix structure as before after a change of variables.
- Reliable recipes exist for selecting small sets of primal constraints for elasticity in 3D, which primarily use edge averages and first order moments as primal constraints. High quality PETSc-based codes have been developed and successfully tested on very large systems. Public domain software in PETSc contributed by Stefano Zampini.

Product spaces

- The BDDC and FETI–DP algorithms can be described in terms of 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: no constraints; $\widehat{W}_\Gamma$: continuity at every point on Γ; $\widetilde{W}_\Gamma$: common values of the primal variables.
- Change variables, explicitly introducing primal variables and complementary sets of dual displacement variables. Simplifies presentation and also makes methods more robust.
- After eliminating the interior variables, write the subdomain Schur complements as

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

- Partially subassemble the $S^{(i)}$, obtaining $\tilde{S}$.

More details on BDDC

- Work with $\widetilde{W}_\Gamma$ and a set of primal constraints. At the end of each iterative step, the approximate solution will be made continuous at all nodal points of the interface; continuity is restored by applying a weighted average operator E_D, which maps $\widetilde{W}_\Gamma$ into $\widehat{W}_\Gamma$.
- In each iteration, first compute the residual of the fully assembled Schur complement. Then apply E_D^T to obtain right-hand side of the partially subassembled linear system. Solve this system and then apply E_D.
- This last step changes the values on Γ, unless the iteration has converged, and results in non-zero residuals at nodes next to Γ.
- In final step of iteration step, eliminate these residuals by solving a Dirichlet problem on each of the subdomains. Accelerate with preconditioned conjugate gradients.
- For any application, only an estimate of $\|E_D\|_{\tilde{S}}$ is needed.

New algorithmic idea

- When designing a BDDC algorithm, we have to choose an effective set of primal constraints and also a recipe for the averaging across interface.
- Traditional averaging recipes found not to work uniformly well for 3D problems in $H(\mathbf{curl})$: DD20 paper by CRD and OBW.
- Alternative found, also very robust for 3D $H(\mathrm{div})$ problems: Duk-Soon Oh, OBW, and CRD; CIMS TR2013-951.
- Both the $H(\mathbf{curl})$ and $H(\mathrm{div})$ problems have two material parameters; complicates the design of the average operator.
- A paper on isogeometric elements, joint with Beirão da Veiga, Pavarino, Scacchi, and Zampini: CIMS TR2013-955 has been been submitted.
- My former student Jong Ho Lee has completed a paper on Reissner-Mindlin plates: CIMS TR2013-958.
- Also work on DG by Dryja and Sarkis and by Chung and Kim.

Stiffness scaling in frequency space

- New average operator E_D across a face $F \subset \Gamma$, common to two subdomains Ω_i and Ω_j, defined in terms of two Schur complements, sub-blocks of Dirichlet-Neumann matrices,

$$S_F^{(k)} := A_{FF}^{(k)} - A_{FI}^{(k)} A_{II}^{(k)^{-1}} A_{IF}^{(k)}, \; k = i, j.$$

- The deluxe averaging operator is then defined by

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

- This part of the action of E_D can be implemented by solving a Dirichlet problem on $\Omega_i \cup \Gamma_{ij} \cup \Omega_j$. Here, Γ_{ij} interface between the two subdomains. This adds to the costs.
- Similar formulas for subdomain edges and other equivalence classes of interface variables. The operator E_D is assembled from these components.

BDDC deluxe

- Core of any estimate for a BDDC algorithm is the norm of the average operator E_D. By an algebraic argument known, for FETI–DP, since 2002,

$$\kappa(M^{-1}A) \leq \|E_D\|_{\tilde{S}}.$$

- We will show that the analysis of BDDC deluxe essentially can be reduced to bounds for individual subdomains.
- Arbitrary jumps in two coefficients can often be accommodated.
- Analysis of traditional BDDC requires an extension theorem; deluxe version does not.
- These methods often competitive with the overlapping Schwarz methods.

Analysis of BDDC deluxe

- Instead of estimating $(R_F^T \bar{w}_F)^T S^{(i)} R_F^T \bar{w}_F$, estimate the norm of $R_F^T (w^{(i)} - \bar{w}_F)$. By simple algebra, we find that

$$w_F^{(i)} - \bar{w}_F = (S_F^{(i)} + S_F^{(j)})^{-1} S_F^{(j)} (w_F^{(i)} - w_F^{(j)}).$$

- By more algebra, noting that $R_F S^{(i)} R_F^T = S_F^{(i)}$:

$$(R_F^T (w_F^{(i)} - \bar{w}_F))^T S^{(i)} (R_F^T (w_F^{(i)} - \bar{w}_F)) =$$

$$(w_F^{(i)} - w_F^{(j)})^T S_F^{(j)} (S_F^{(i)} + S_F^{(j)})^{-1} S_F^{(i)} (S_F^{(i)} + S_F^{(j)})^{-1} S_F^{(j)} (w_F^{(i)} - w_F^{(ji)}))$$

$$\leq 2 (w_F^{(i)})^T S_F^{(j)} (S_F^{(i)} + S_F^{(j)})^{-1} S_F^{(i)} (S_F^{(i)} + S_F^{(j)})^{-1} S_F^{(j)} w_F^{(i)} +$$

$$2 (w_F^{(j)})^T S_F^{(i)} (S_F^{(i)} + S_F^{(j)})^{-1} S_F^{(i)} (S_F^{(i)} + S_F^{(j)})^{-1} S_F^{(j)} w_F^{(j)}.$$

Proof continued

- We can now prove that

$$S_F^{(j)}(S_F^{(i)} + S_F^{(j)})^{-1} S_F^{(i)} (S_F^{(i)} + S_F^{(j)})^{-1} S_F^{(j)} \leq S_F^{(i)}$$

  and that

$$S_F^{(j)}(S_F^{(i)} + S_F^{(j)})^{-1} S_F^{(i)} (S_F^{(i)} + S_F^{(j)})^{-1} S_F^{(j)} \leq S_F^{(j)}.$$

- Follows easily by considering action of these operators on any eigenvector of the generalized eigenvalue problem

$$S_F^{(i)} \phi = \lambda S_F^{(j)} \phi$$

  and using only that all eigenvalues are strictly positive.
- Thus, $(R_F^T(w_F^{(i)} - \bar{w}_F))^T S^{(i)} R_F^T(w_F^{(i)} - \bar{w}_F)$
  $\leq 2 {w_F^{(i)}}^T S_F^{(i)} w_F^{(i)} + 2 {w_F^{(j)}}^T S_F^{(j)} w_F^{(j)}.$
- Each of the two terms local to only one subdomain.

Proof continued

- Now remains to estimate ${w_F^{(i)}}^T S_F^{(i)} w_F^{(i)}$ by ${w^{(i)}}^T S^{(i)} w^{(i)}$, after possible shift with an element of the primal space.
- Routine for H^1. Completely standard estimates in the domain decomposition literature: Face lemma; compare energy of the extension of face values by zero with that of the minimal energy extension. Factor of $C(1+\log(H/h))^2$ results.
- Also estimate contributions from subdomain edges, etc. For many edges, more than two Schur complements enter in calculation of the average across the interface; a cheaper approach exists. Same bound results. Again, for H^1, estimates could have been done 20 years ago.
- For $H(\text{div})$ and $H(\textbf{curl})$, considerable new efforts required.

$H(\mathbf{curl})$ problems in 3D

- Consider variational problem: Find $\mathbf{u} \in H_0(\text{curl};\Omega)$ such that

$$a(\mathbf{u},\mathbf{v})_\Omega = (\mathbf{f},\mathbf{v})_\Omega \quad \forall \mathbf{v} \in H_0(\text{curl};\Omega),$$

where $\mathbf{u} \times \mathbf{n} = 0$ on $\partial\Omega$ and where

$$a(\mathbf{u},\mathbf{v})_\Omega := \int_\Omega [\alpha \nabla\times\mathbf{u}\cdot\nabla\times\mathbf{v} + \beta\mathbf{u}\cdot\mathbf{v}]\, dx, \quad (\mathbf{f},\mathbf{v})_\Omega := \int_\Omega \mathbf{f}\cdot\mathbf{v}\, dx.$$

- Here, $\alpha(x) \geq 0$ and $\beta(x)$ strictly positive.
- For piecewise constant coefficients,

$$a(\mathbf{u},\mathbf{v})_\Omega = \sum_{i=1}^{N} (\alpha_i(\nabla\times\mathbf{u},\nabla\times\mathbf{v})_{\Omega_i} + \beta_i(\mathbf{u},\mathbf{v})_{\Omega_i}).$$

Pioneering work by Andrea Toselli

- For Nédélec elements, the obvious finite element basis results in a poor result; the coupling between the faces and wire basket is far too strong.
- Following Toselli DD16 talk, IMA J. Numer. Anal. vol. 26, 2006, we change the variables associated with the wire basket of the interface: Use a constant along each subdomain edge and the gradient of the standard $W_{\text{grad}}^{h_i}$ basis functions for all the interior nodes of the subdomain edges.
- After this change of variables, quite stable decompositions can be found. Many basis functions associated with the wire basket will then have non-zero coefficients on the edges of the triangulation, which have an end point on the wire basket.
- To establish the stability of our decompositions, a new auxiliary result was needed; you will not need to know any trace theorems to read our proof.

$H(\mathbf{curl})$ results

- Toselli primarily advocates the use of two primal variables for each subdomain edge: the average and first moment. So do we. Our technical tools differ a lot from his.
- We have improved Toselli's condition number bound from

$$C \max_i (1 + \frac{\beta_i H_i^2}{\alpha_i})(1 + \log(H_i/h_i))^4$$

to

$$C \max_i \min((H_i/h_i)^2, (1 + \frac{\beta_i H_i^2}{\alpha_i}))(1 + \log(H_i/h_i))^2.$$

- We have fewer restrictions on the coefficients than Toselli; our constant C is independent of the α_i and β_i.
- So far, we have not mastered the case where $\frac{\beta_i H_i^2}{\alpha_i}$ is large. For $H(\text{div})$, the same simple primal space always works; not so for $H(\mathbf{curl})$.

Nédeléc spaces

- There are now two relevant finite element spaces, namely $W_{\mathrm{curl}}^{h_i}$ of lowest order triangular Nédeléc elements and $W_{\mathrm{grad}}^{h_i}$ of standard piecewise linear continuous elements, on the same triangulation τ_h.
- The Nédeléc elements are $H(\mathbf{curl})-$conforming with constant tangential components on each edge of the triangulation with these values common across element edges.
- In 3D, the Nédeléc element space can be characterized as the range of an interpolation operator Π^{h_i} :

$$\Pi^{h_i}(\mathbf{u}) := \sum_{edges} u_e \mathbf{N}_e, \quad u_e = (1/|e|) \int_e \mathbf{u} \cdot \mathbf{t}_e ds,$$

where $\mathbf{t}_e$ is a unit vector in the direction of the edge e and $\mathbf{N}_e \in W_{\mathrm{curl}}^{h_i}$ is a piecewise linear, vector-valued function such that $\mathbf{N}_e \cdot \mathbf{t}_e = 1$ and with $\mathbf{N}_e \cdot \mathbf{t}_{e'} = 0$ on the other edges.

Hiptmair–Xu lemma

- A key to our work is: For any $\mathbf{u}_h \in W^{h_i}_{\text{curl}}(\Omega_i)$, there exist $\boldsymbol{\Psi}_h \in (W^{h_i}_{\text{grad}}(\Omega_i))^3$, $p_h \in W^{h_i}_{\text{grad}}(\Omega_i)$, and $\mathbf{q}_h \in W^{h_i}_{\text{curl}}(\Omega_i)$, such that

$$
\begin{aligned}
\mathbf{u}_h \quad & = \mathbf{q}_h + \Pi^{h_i}(\boldsymbol{\Psi}_h) + \nabla p_h, \\
\|\nabla p_h\|^2_{L^2(\Omega_i)} \quad & \le C\left(\|\mathbf{u}_h\|^2_{L^2(\Omega_i)} + H^2\|\nabla \times \mathbf{u}_h\|^2_{L^2(\Omega_i)}\right), \\
\|h^{-1}\mathbf{q}_h\|^2_{L^2(\Omega_i)} \quad & +\|\boldsymbol{\Psi}_h\|^2_{H(\text{grad},\Omega_i)} \le C\|\nabla \times \mathbf{u}_h\|^2_{L^2(\Omega_i)}.
\end{aligned}
$$

- Note that these bounds are local. For convex subdomains, we can also estimate the $L^2(\Omega_i)-$norms of ∇p_h and $\boldsymbol{\Psi}_h$ in terms of $\|\mathbf{u}_h\|_{L^2 i(\Omega_i)}$. We assume that our subdomains convex polytops.
- This result essentially borrowed from Hiptmair, Xu, and Zou. Also central in work on algebraic multigrid algorithms for $H(\mathbf{curl})$; see Hiptmair and Xu, SINUM 45(6), 2007. (Effectively, Poisson solvers can be used.)

$H(\mathrm{div})$ and Reissner–Mindlin plates

- Work has been completed by Duk-Soon Oh, CRD, and OBW on BDDC deluxe algorithms for lowest order Raviart-Thomas elements.
- In variational form: Find $\mathbf{u} \in H_0(\mathrm{div}, \Omega)$, such that

$$\int_\Omega [\alpha \operatorname{div} \mathbf{u} \operatorname{div} \mathbf{v} + \beta \mathbf{u} \cdot \mathbf{v}] dx = \int_\Omega \mathbf{f} \cdot \mathbf{v} dx$$

  where $\alpha(x) \geq 0$ and $\beta(x)$ strictly positive. We assume that these functions are constant in each subdomain Ω_i.
- By developing two sets of inequalities and using the fact that BDDC deluxe allows the use of subdomain bounds, we have a result which is uniform for all choices of α_i and β_i.
- Again we use a result from Hiptmair and Xu, in this case, on the decomposition of $H(\mathrm{div})$ functions.

Results

- We have:
$$\kappa(M^{-1}A) \leq C(1 + \log(H/h))^2.$$
- In our H(div) experiments, condition numbers in the range of 2 to 5. Examples have been constructed where standard BDDC performs much worse with condition numbers in the thousands. Good numerical results with variable coefficients inside subdomains too.
- Recent work by Jong Ho Lee gives similar bound for BDDC deluxe and numerically thin Reissner–Mindlin plates approximated by low order mixed finite elements due to Falk and Tu. In his experiments, condition numbers rarely exceed 4. The thickness can vary between $t = 1$ and $t = 10^{-5}$. Without preconditioning the condition number exceeds 10^{11} for $t = 10^{-5}$.

Isogeometric analysis

- In joint work by Beirão da Veiga, Pavarino, Scacchi, OBW and Zampini, BDDC deluxe has been used to substantially improve the iteration count in comparison with the standard BDDC.
- In this application finite elements are replaced by nonuniform rational B-splines (NURBS), borrowed from computer aided design.
- We often have fat interfaces since the NURBS do not have a nodal basis.
- We have bounds for the condition numbers and observe in our experiments that the iteration count does not grow with the degree of the underlying B-splines. Work is continuing to decrease the dimension of the primal space and to extend the theory.

Perspectives

- I think, we are quite early in this game. One thing is clear: many problems can equally well be solved using traditional BDDC.
- What about irregular subdomains? This was the topic of my Jerusalem DD talk. There is one table in the paper based on CRD's DD20 talk. The primal set is different and very large.
- Would the deluxe version help when we have large variations of coefficients inside subdomains?
- Can we use the present tool set for 3D $H(\mathbf{curl})$ to develop and analyze other DD methods?
- CRD gave a very interesting talk at DD21 indicating that eigenvectors of the generalized eigenvalue problems with a pair of Schur complements can be used to improve the set of primal constraints. Work along these lines has begun with some of our Italian friends. (There is also a lot of activity elsewhere.)