On the Robustness and Prospects of Adaptive BDDC Methods for Finite Element Discretizations of Elliptic PDEs with High-Contrast Coefficients

Stefano Zampini

جامعة الملك عبدالله
للعلوم والتقنية
King Abdullah University of
Science and Technology

Computer, Electrical and Mathematical Sciences & Engineering,
Extreme Computing Research Center
Saudi Arabia

SHAXC '16 @ KAUST, May 10th, 2016

- Iterative solution of large and sparse elliptic linear systems arising from FEM

$$Ax = b, \quad \kappa(A) = \frac{\lambda_M(A)}{\lambda_m(A)} \propto h^{-2}$$

- Preconditioned conjugate gradient

$$M^{-1}Ax = M^{-1}b, \quad \kappa(M^{-1}A) = \frac{\lambda_M(M^{-1}A)}{\lambda_m(M^{-1}A)}$$

- Use domain decomposition preconditioners:
  - Couple local Cholesky solvers through a suitable coarse space
  - Obtain convergence rates which are independent of the number of subdomains, of the discretization space, and slowly deteriorates with the size of the subdomain problems.
  - Accomodate heterogeneities in the coefficients of PDE.

[A. Toselli and O. B. Widlund, Domain Decomposition Methods: algorithm and theory, 2005]

The characterization of the BDDC method [Dohrmann, SISC 25, 2003] relies on

- the selection of primal (coarse) degrees of freedom
- the choice of an averaging procedure.

Typical preconditioned spectra for elliptic problems (if a suitable primal space is found from the theory)

$$\kappa(M^{-1}A) \leq C(1+\log(H/h))^2, \quad H/h \propto n_{loc}$$

with H maximum diameter of sudomains, h the mesh size and C independent on the number of subdomains and possibly independent on the coefficients of the PDE [Mandel, Dohrmann, Tezaur, Appl. Numer. Math. 54, 2005], [Li and Widlund, IJNME, 2008].
Dual of the FETI-DP method [C. Farhat *et. al* , IJNME 50, 2001].

- Contact problems [P. Avery, G. Rebel, M. Lesoinne, C. Farhat. CMAME 93, 2004]
- Indefinite problems [J. Li, X. Tu. NLAA 16, 2009], [C. Farhat, J. Li ANM 54, 2005], [C. Farhat, J. Li, P. Avery IJNME 63, 2005]
- Porous media flow [X. Tu. ETNA 20, 2005]
- Electromagnetic problems [Y. J. Li, J. M. Jin IEEE Trans. Antennas Propag. 54, 2006]
- Incompressible Stokes [J. Li, O. B. Widlund. SISC 44, 2006]
- Linear elasticity [A. Klawonn, O. B. Widlund CPAM 59, 2006]
- Stokes problem [H. H. Kim, C. O. Lee, E. H. Park SISC 47, 2010]
- Stokes–Darcy coupling [J. Galvis, M. Sarkis. CAMCS 5, 2010]
- Almost Incompressible Elasticity [L. Pavarino, O. B. Widlund, S. Z. SISC 32, 2010]

Introduction: some references

- Spectral Elements [L. F. Pavarino, CMAME 196, 2007]
- Lowest order Nédélec elements [A. Toselli, IMAJNA 26, 2006], [C. Dohrmann , O. B.Widlund, CPAM 2015], [S. Z. submitted 2016]
- Discontinuous Galerkin [M. Dryja, J. Galvis, M. Sarkis. J. Complexity 23, 2007]
- Mortar discretizations [H. H. Kim. SINUM 46, 2008], [H. H. Kim, M. Dryja, O. B. Widlund. SINUM 47, 2009]
- Reissner-Mindlin plates and Tu-Falk elements [J. H. Lee, SINUM, 2015]
- Naghdi shells and MITC elements [L. Beirão da Veiga, C. Chinosi, C. Lovadina, L. F. Pavarino. Comp. Struct. 102, 2012]
- IsoGeometric Analysis [L. Beirão da Veiga, L. Pavarino, S. Scacchi, O. B. Widlund, S.Z. SISC, 36, 2014]
- Lowest order Raviart-Thomas elements [D.-S. Oh , O. B. Widlund, S. Z., C. Dohrmann , TR, 2015]

Introduction: adaptive coarse space

- Convergence of DD algorithms usually deteriorates when jumps in the coefficients are not aligned with Γ.

- Convergence of DD algorithms usually deteriorates when jumps in the coefficients are not aligned with Γ.
- Coarse space has to be adaptively generated using generalized eigenvalue problems [J. Mandel and B. Sousedík, DD XVI, 2007], [B. Sousedík, J. Šístek, and J. Mandel, Computing, 95, 2013], [A. Klawonn, P. Radtke and O. Rheinbach, SINUM 53, 2015] [A. Klawonn, P. Kühn and O. Rheinbach, TR 2015], [H.H. Kim and E.T. Chung, SIAM J. Multiscale Model. Simul., 13, 2015, H.H. Kim, E.T. Chung and J. Wang CMAME, 2015], [C. Pechstein, C. R. Dohrmann, Seminar talk 2013], [J. Calvo and O. B. Widlund, submitted, 2015],[L. Beirão da Veiga, L. Pavarino, S. Scacchi, O. B. Widlund, S.Z. Submitted, 2015]

- Convergence of DD algorithms usually deteriorates when jumps in the coefficients are not aligned with Γ.
- Coarse space has to be adaptively generated using generalized eigenvalue problems [J. Mandel and B. Sousedík, DD XVI, 2007], [B. Sousedík, J. Šístek, and J. Mandel, Computing, 95, 2013], [A. Klawonn, P. Radtke and O. Rheinbach, SINUM 53, 2015] [A. Klawonn, P. Kühn and O. Rheinbach, TR 2015], [H.H. Kim and E.T. Chung, SIAM J. Multiscale Model. Simul., 13, 2015, H.H. Kim, E.T. Chung and J. Wang CMAME, 2015], [C. Pechstein, C. R. Dohrmann, Seminar talk 2013], [J. Calvo and O. B. Widlund, submitted, 2015],[L. Beirão da Veiga, L. Pavarino, S. Scacchi, O. B. Widlund, S.Z. Submitted, 2015]
- Given a user-defined threshold λ, we can guarantee that

$$\kappa \leq C\lambda$$

- Convergence of DD algorithms usually deteriorates when jumps in the coefficients are not aligned with Γ.
- Coarse space has to be adaptively generated using generalized eigenvalue problems [J. Mandel and B. Sousedík, DD XVI, 2007], [B. Sousedík, J. Šístek, and J. Mandel, Computing, 95, 2013], [A. Klawonn, P. Radtke and O. Rheinbach, SINUM 53, 2015] [A. Klawonn, P. Kühn and O. Rheinbach, TR 2015], [H.H. Kim and E.T. Chung, SIAM J. Multiscale Model. Simul., 13, 2015, H.H. Kim, E.T. Chung and J. Wang CMAME, 2015], [C. Pechstein, C. R. Dohrmann, Seminar talk 2013], [J. Calvo and O. B. Widlund, submitted, 2015],[L. Beirão da Veiga, L. Pavarino, S. Scacchi, O. B. Widlund, S.Z. Submitted, 2015]
- Given a user-defined threshold λ, we can guarantee that
$$\kappa \leq C\lambda$$
- Similar techniques have been also studied for enriching the coarse space of Schwarz algorithms and FETI
[J. Galvis and Y. Efendiev, Multiscale Models Simul. 8, 2010], [N. Spillane and D. J. Rixen, IJNME, 2013], [N. Spillane *et. al.*,C.R. Math. Acad. Sci. Paris 2013], [N. Spillane *et. al.*, Num. Math, 126, 2014]
- For multigrid, BootstrapAMG [A. Brandt *et. al.*, SISC, 33, 2011].

Adaptive BDDC satisfies 3 of the pillars for exascale algorithms [J. Dongarra, et al, Int. J. High Perf. Comp. Appl. 6, 2011]

- Reduces the global synchronization steps and MatVecs in Krylov methods
- Increases arithmetic intensity of the preconditioning step
- Increases concurrency of the preconditioning step

Adaptive BDDC satisfies 3 of the pillars for exascale algorithms [J. Dongarra, et al, Int. J. High Perf. Comp. Appl. 6, 2011]

- Reduces the global synchronization steps and MatVecs in Krylov methods
- Increases arithmetic intensity of the preconditioning step
- Increases concurrency of the preconditioning step

Other key features of the algorithm

- Cholesky based
- Local and coarse problem additively combined (overlap)
- Multilevel extensions with high F/C coarsening ratios $O(10^2) - O(10^4)$ (however, convergence is exponentially dependent on the number of levels)

Adaptive BDDC satisfies 3 of the pillars for exascale algorithms [J. Dongarra, et al, Int. J. High Perf. Comp. Appl. 6, 2011]

- Reduces the global synchronization steps and MatVecs in Krylov methods
- Increases arithmetic intensity of the preconditioning step
- Increases concurrency of the preconditioning step

Other key features of the algorithm

- Cholesky based
- Local and coarse problem additively combined (overlap)
- Multilevel extensions with high F/C coarsening ratios $O(10^2) - O(10^4)$ (however, convergence is exponentially dependent on the number of levels)

Can it compete with optimally designed MG?

Adaptive BDDC satisfies 3 of the pillars for exascale algorithms [J. Dongarra, et al, Int. J. High Perf. Comp. Appl. 6, 2011]

- Reduces the global synchronization steps and MatVecs in Krylov methods
- Increases arithmetic intensity of the preconditioning step
- Increases concurrency of the preconditioning step

Other key features of the algorithm

- Cholesky based
- Local and coarse problem additively combined (overlap)
- Multilevel extensions with high F/C coarsening ratios $O(10^2) - O(10^4)$ (however, convergence is exponentially dependent on the number of levels)

Can it compete with optimally designed MG?

Usually not, but there are problems where MG fails.

- Ω subdivided in N non-overlapping open subdomains

$$\overline{\Omega} = \bigcup_{i=1}^{N} \overline{\Omega}_i, \quad \Omega_j \cap \Omega_i = \emptyset, \quad \Gamma = \bigcup_{i \neq j} \partial\Omega_j \cap \partial\Omega_i.$$

- Ω subdivided in N non-overlapping open subdomains

$$\overline{\Omega} = \bigcup_{i=1}^{N} \overline{\Omega}_i, \quad \Omega_j \cap \Omega_i = \emptyset, \quad \Gamma = \bigcup_{i \neq j} \partial\Omega_j \cap \partial\Omega_i.$$

- A is never assembled explicitly; matrix vector product as

$$\mathrm{A} = R^T \mathrm{A}^* R, \quad \mathrm{A}^* = \mathrm{diag}(\mathrm{A}^{(i)}),$$

with $\mathrm{A}^{(i)}$ the matrix of the FEM problem on Ω_i.

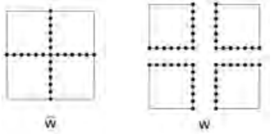

- Block factorization for A based on the split $\widehat{\mathbf{W}} = \widehat{\mathbf{W}}_\Gamma \oplus \mathbf{W}_I$

$$A^{-1} = \begin{bmatrix} I_{II} & -A_{II}^{-1}A_{I\Gamma} \\ & I_{\Gamma\Gamma} \end{bmatrix} \begin{bmatrix} A_{II}^{-1} & \\ & S_\Gamma^{-1} \end{bmatrix} \begin{bmatrix} I_{II} & \\ -A_{I\Gamma}^T A_{II}^{-1} & I_{\Gamma\Gamma} \end{bmatrix}$$

with $S_\Gamma = A_{\Gamma\Gamma} - A_{I\Gamma}^T A_{II}^{-1} A_{I\Gamma}$.

- Block factorization for A based on the split $\widehat{\mathbf{W}} = \widehat{\mathbf{W}}_\Gamma \oplus \mathbf{W}_I$

$$A^{-1} = \begin{bmatrix} I_{II} & -A_{II}^{-1}A_{I\Gamma} \\ & I_{\Gamma\Gamma} \end{bmatrix} \begin{bmatrix} A_{II}^{-1} & \\ & S_\Gamma^{-1} \end{bmatrix} \begin{bmatrix} I_{II} & \\ -A_{I\Gamma}^T A_{II}^{-1} & I_{\Gamma\Gamma} \end{bmatrix}$$

with $S_\Gamma = A_{\Gamma\Gamma} - A_{I\Gamma}^T A_{II}^{-1} A_{I\Gamma}$.

- Block preconditioner $\kappa(M^{-1}A) = \kappa(M_\Gamma^{-1}S_\Gamma)$

$$M^{-1} = \begin{bmatrix} I_{II} & -A_{II}^{-1}A_{I\Gamma} \\ & I_{\Gamma\Gamma} \end{bmatrix} \begin{bmatrix} A_{II}^{-1} & \\ & M_\Gamma^{-1} \end{bmatrix} \begin{bmatrix} I_{II} & \\ -A_{I\Gamma}^T A_{II}^{-1} & I_{\Gamma\Gamma} \end{bmatrix}$$

- Block factorization for A based on the split $\widehat{\mathbf{W}} = \widehat{\mathbf{W}}_\Gamma \oplus \mathbf{W}_I$

$$A^{-1} = \begin{bmatrix} I_{II} & -A_{II}^{-1}A_{I\Gamma} \\ & I_{\Gamma\Gamma} \end{bmatrix} \begin{bmatrix} A_{II}^{-1} & \\ & S_\Gamma^{-1} \end{bmatrix} \begin{bmatrix} I_{II} & \\ -A_{I\Gamma}^T A_{II}^{-1} & I_{\Gamma\Gamma} \end{bmatrix}$$

with $S_\Gamma = A_{\Gamma\Gamma} - A_{I\Gamma}^T A_{II}^{-1} A_{I\Gamma}$.

- Block preconditioner $\kappa(M^{-1}A) = \kappa(M_\Gamma^{-1} S_\Gamma)$

$$M^{-1} = \begin{bmatrix} I_{II} & -A_{II}^{-1}A_{I\Gamma} \\ & I_{\Gamma\Gamma} \end{bmatrix} \begin{bmatrix} A_{II}^{-1} & \\ & M_\Gamma^{-1} \end{bmatrix} \begin{bmatrix} I_{II} & \\ -A_{I\Gamma}^T A_{II}^{-1} & I_{\Gamma\Gamma} \end{bmatrix}$$

- DD solvers: BDD [J. Mandel, 1993]
- Algebraic solvers: ShyLU [S. Rajamanickam et al. IPDPS, 2012] , STRUMPACK [P. Ghysels et. al. 2015] and many others (search the web for "Schur hybrid solvers")

- Basic idea of BDDC: instead of S_Γ, invert $\widetilde{S}_\Gamma$, defined on a partially assembled space

$$\widehat{\mathbf{W}} \subset \widetilde{\mathbf{W}} = \mathbf{W}_I \oplus \widetilde{\mathbf{W}}_\Gamma \subset \mathbf{W}, \quad \widetilde{\mathbf{W}}_\Gamma = \widehat{\mathbf{W}}_\Pi \oplus \mathbf{W}_\Delta$$

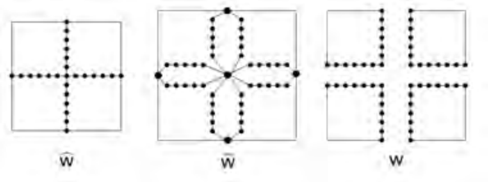

- Discontinuous on dual dofs Δ, continuous on primals dofs Π.
- Primal vertices to prevent subdomains from floating.
- Additional primal dofs (functionals) can be needed for edges and faces of Γ to obtain an algorithmically scalable and robust method.

$$M_\Gamma^{-1} = \widetilde{R}_{D,\Gamma}^T \widetilde{S}_\Gamma^{-1} \widetilde{R}_{D,\Gamma},$$

$$M_\Gamma^{-1} = \widetilde{R}_{D,\Gamma}^T \widetilde{S}_\Gamma^{-1} \widetilde{R}_{D,\Gamma},$$

Block Cholesky

$$\widetilde{S}_\Gamma^{-1} = R_{\Gamma\Delta}^T \left(\sum_{i=1}^N \begin{bmatrix} 0 & R_\Delta^{(i)T} \end{bmatrix} \begin{bmatrix} A_{II}^{(i)} & A_{I\Delta}^{(i)} \\ A_{I\Delta}^{(i)T} & A_{\Delta\Delta}^{(i)} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ R_\Delta^{(i)} \end{bmatrix} \right) R_{\Gamma\Delta} + \Phi S_{\Pi\Pi}^{-1} \Phi^T.$$

$$M_\Gamma^{-1} = \widetilde{R}_{D,\Gamma}^T \widetilde{S}_\Gamma^{-1} \widetilde{R}_{D,\Gamma},$$

Block Cholesky

$$\widetilde{S}_\Gamma^{-1} = R_{\Gamma\Delta}^T \left(\sum_{i=1}^{N} \begin{bmatrix} 0 & R_\Delta^{(i)T} \end{bmatrix} \begin{bmatrix} A_{II}^{(i)} & A_{I\Delta}^{(i)} \\ A_{I\Delta}^{(i)T} & A_{\Delta\Delta}^{(i)} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ R_\Delta^{(i)} \end{bmatrix} \right) R_{\Gamma\Delta} + \Phi S_{\Pi\Pi}^{-1} \Phi^T.$$

Primal basis

$$\Phi = R_{\Gamma\Pi}^T - R_{\Gamma\Delta}^T \left(\sum_{i=1}^{N} \begin{bmatrix} 0 & R_\Delta^{(i)T} \end{bmatrix} \begin{bmatrix} A_{II}^{(i)} & A_{I\Delta}^{(i)} \\ A_{I\Delta}^{(i)T} & A_{\Delta\Delta}^{(i)} \end{bmatrix}^{-1} \begin{bmatrix} A_{I\Pi}^{(i)} \\ A_{\Delta\Pi}^{(i)} \end{bmatrix} R_\Pi^{(i)} \right).$$

$$M_\Gamma^{-1} = \widetilde{R}_{D,\Gamma}^T \widetilde{S}_\Gamma^{-1} \widetilde{R}_{D,\Gamma},$$

Block Cholesky

$$\widetilde{S}_\Gamma^{-1} = R_{\Gamma\Delta}^T \left(\sum_{i=1}^{N} \begin{bmatrix} 0 & R_\Delta^{(i)T} \end{bmatrix} \begin{bmatrix} A_{II}^{(i)} & A_{I\Delta}^{(i)} \\ A_{I\Delta}^{(i)T} & A_{\Delta\Delta}^{(i)} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ R_\Delta^{(i)} \end{bmatrix} \right) R_{\Gamma\Delta} + \Phi S_{\Pi\Pi}^{-1} \Phi^T.$$

Primal basis

$$\Phi = R_{\Gamma\Pi}^T - R_{\Gamma\Delta}^T \left(\sum_{i=1}^{N} \begin{bmatrix} 0 & R_\Delta^{(i)T} \end{bmatrix} \begin{bmatrix} A_{II}^{(i)} & A_{I\Delta}^{(i)} \\ A_{I\Delta}^{(i)T} & A_{\Delta\Delta}^{(i)} \end{bmatrix}^{-1} \begin{bmatrix} A_{I\Pi}^{(i)} \\ A_{\Delta\Pi}^{(i)} \end{bmatrix} R_\Pi^{(i)} \right).$$

Primal coarse problem

$$S_{\Pi\Pi} = R_\Pi^T S_{\Pi\Pi}^* R_\Pi, \quad S_{\Pi\Pi}^* = \mathrm{diag}(S_{\Pi\Pi}^{(i)})$$

$$S_{\Pi\Pi}^{(i)} = A_{\Pi\Pi}^{(i)} - \begin{bmatrix} A_{I\Pi}^{(i)T} & A_{\Delta\Pi}^{(i)T} \end{bmatrix} \begin{bmatrix} A_{II}^{(i)} & A_{I\Delta}^{(i)} \\ A_{I\Delta}^{(i)T} & A_{\Delta\Delta}^{(i)} \end{bmatrix}^{-1} \begin{bmatrix} A_{I\Pi}^{(i)} \\ A_{\Delta\Pi}^{(i)} \end{bmatrix}$$

[C. Dohrmann, SISC, 2003], [C. Dohrmann, NLAA 2007], [J. Li and O. B. Widlund 2008]

$$M_{\Gamma}^{-1} = \widetilde{R}_{D,\Gamma}^{T} \widetilde{S}_{\Gamma}^{-1} \widetilde{R}_{D,\Gamma},$$

Block Cholesky

$$\widetilde{S}_{\Gamma}^{-1} = R_{\Gamma\Delta}^{T} \left(\sum_{i=1}^{N} \begin{bmatrix} 0 & R_{\Delta}^{(i)T} \end{bmatrix} \begin{bmatrix} A_{II}^{(i)} & A_{I\Delta}^{(i)} \\ A_{I\Delta}^{(i)T} & A_{\Delta\Delta}^{(i)} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ R_{\Delta}^{(i)} \end{bmatrix} \right) R_{\Gamma\Delta} + \Phi S_{\Pi\Pi}^{-1} \Phi^{T}.$$

Primal basis

$$\Phi = R_{\Gamma\Pi}^{T} - R_{\Gamma\Delta}^{T} \left(\sum_{i=1}^{N} \begin{bmatrix} 0 & R_{\Delta}^{(i)T} \end{bmatrix} \begin{bmatrix} A_{II}^{(i)} & A_{I\Delta}^{(i)} \\ A_{I\Delta}^{(i)T} & A_{\Delta\Delta}^{(i)} \end{bmatrix}^{-1} \begin{bmatrix} A_{I\Pi}^{(i)} \\ A_{\Delta\Pi}^{(i)} \end{bmatrix} R_{\Pi}^{(i)} \right).$$

Primal coarse problem

$$S_{\Pi\Pi} = R_{\Pi}^{T} S_{\Pi\Pi}^{*} R_{\Pi}, \quad S_{\Pi\Pi}^{*} = \mathrm{diag}(S_{\Pi\Pi}^{(i)})$$

$$S_{\Pi\Pi}^{(i)} = A_{\Pi\Pi}^{(i)} - \begin{bmatrix} A_{I\Pi}^{(i)T} & A_{\Delta\Pi}^{(i)T} \end{bmatrix} \begin{bmatrix} A_{II}^{(i)} & A_{I\Delta}^{(i)} \\ A_{I\Delta}^{(i)T} & A_{\Delta\Delta}^{(i)} \end{bmatrix}^{-1} \begin{bmatrix} A_{I\Pi}^{(i)} \\ A_{\Delta\Pi}^{(i)} \end{bmatrix}$$

[C. Dohrmann, SISC, 2003], [C. Dohrmann, NLAA 2007], [J. Li and O. B. Widlund 2008]

Coarse problem obtained by subassembling, we can rercurse → multilevel BDDC

More details on BDDC method and its implementation in PETSc: [S. Z. PCBDDC : a class of robust dual-primal methods in PETSc, SIAM CS&E special issue on software, 2015]

More details on BDDC method and its implementation in PETSc: [S. Z. PCBDDC : a class of robust dual-primal methods in PETSc, SIAM CS&E special issue on software, 2015]

- 92% efficiency on 195K cores and 12.4B dofs (hexahedra, boxes, Poisson)

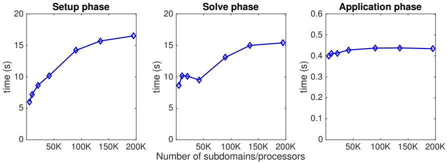

More details on BDDC method and its implementation in PETSc: [S. Z. PCBDDC : a class of robust dual-primal methods in PETSc, SIAM CS&E special issue on software, 2015]

- 92% efficiency on 195K cores and 12.4B dofs (hexahedra, boxes, Poisson)

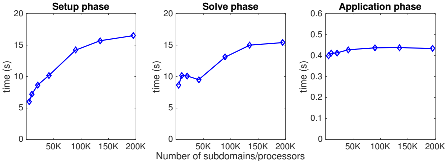

- 99% efficiency on 32K cores and 1B dofs (tetrahedra, ParMetis, H(div))

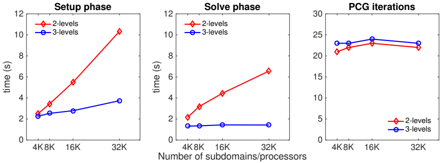

More details on BDDC method and its implementation in PETSc: [S. Z. PCBDDC : a class of robust dual-primal methods in PETSc, SIAM CS&E special issue on software, 2015]

- 92% efficiency on 195K cores and 12.4B dofs (hexahedra, boxes, Poisson)

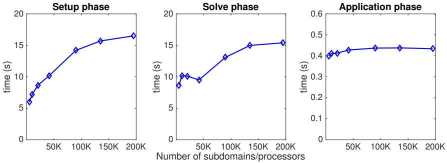

- 99% efficiency on 32K cores and 1B dofs (tetrahedra, ParMetis, H(div))

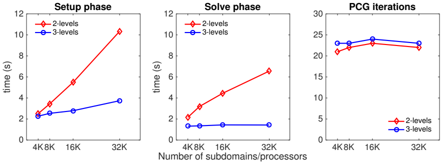

Other implementations:

- Klawonn and Rheinbach [A. Klawonn and O. Rheinbach, ZAMM-Z, 90, 2010], [A. Klawonn, M. Lanser, and O. Rheinbach, SISC, 2014].
- Badia's group in Barcelona [S. Badia, A. F. Martin and J. Principe, SISC, 2014].
- BDDCML (J. Sístek, http://users.math.cas.cz/sistek/software/bddcml.html).

With elliptic problems: $\kappa \leq C||\widetilde{R}_{D,\Gamma}\widetilde{R}_{\Gamma}^{T}||^{2}_{\widetilde{S}_{\Gamma}}$ [J. Mandel, C. Dohrmann 2003]

With elliptic problems: $\kappa \leq C||\widetilde{R}_{D,\Gamma}\widetilde{R}_{\Gamma}^T||^2_{\widetilde{S}_{\Gamma}}$ [J. Mandel, C. Dohrmann 2003]

Key for efficiency, explicit knowledge of $S^{(i)} = A^{(i)}_{\Gamma\Gamma} - A^{(i)T}_{I\Gamma} A^{(i)-1}_{II} A^{(i)}_{I\Gamma}$

With elliptic problems: $\kappa \leq C||\widetilde{R}_{D,\Gamma}\widetilde{R}_{\Gamma}^T||^2_{\widetilde{S}_{\Gamma}}$ [J. Mandel, C. Dohrmann 2003]

Key for efficiency, explicit knowledge of $S^{(i)} = A^{(i)}_{\Gamma\Gamma} - A^{(i)T}_{I\Gamma} A^{(i)-1}_{II} A^{(i)}_{I\Gamma}$

- Solve for each subdomain face F (dense blocks) [C. Pechstein, C. R. Dohrmann, 2013]

$$(\widetilde{S}^{(i)}_F : \widetilde{S}^{(j)}_F)\phi = \mu(S^{(i)}_F : S^{(j)}_F)\phi$$

where

$$A : B = (A^{-1} + B^{-1})^{-1}, \quad \widetilde{S}^{(k)}_F = S^{(k)}_{FF} - S^{(k)}_{FF'} S^{(k)-1}_{F'F'} S^{(k)}_{F'F}$$

With elliptic problems: $\kappa \leq C||\widetilde{R}_{D,\Gamma}\widetilde{R}_{\Gamma}^T||^2_{\widetilde{S}_\Gamma}$ [J. Mandel, C. Dohrmann 2003]

Key for efficiency, explicit knowledge of $S^{(i)} = A^{(i)}_{\Gamma\Gamma} - A^{(i)T}_{I\Gamma} A^{(i)-1}_{II} A^{(i)}_{I\Gamma}$

- Solve for each subdomain face F (dense blocks) [C. Pechstein, C. R. Dohrmann, 2013]

$$(\widetilde{S}_F^{(i)} : \widetilde{S}_F^{(j)})\phi = \mu(S_F^{(i)} : S_F^{(j)})\phi$$

  where

$$A : B = (A^{-1} + B^{-1})^{-1}, \quad \widetilde{S}_F^{(k)} = S_{FF}^{(k)} - S_{FF'}^{(k)} S_{F'F'}^{(k)-1} S_{F'F}^{(k)}$$

- Add to the primal space

$$\widehat{\mathbf{W}}_\Pi \leftarrow \{(S_F^{(i)} : S_F^{(j)})\phi_k \mid 1/\mu_k > \lambda\}$$

With elliptic problems: $\kappa \leq C||\widetilde{R}_{D,\Gamma}\widetilde{R}_{\Gamma}^T||^2_{\widetilde{S}_\Gamma}$ [J. Mandel, C. Dohrmann 2003]

Key for efficiency, explicit knowledge of $S^{(i)} = A^{(i)}_{\Gamma\Gamma} - A^{(i)T}_{I\Gamma} A^{(i)-1}_{II} A^{(i)}_{I\Gamma}$

- Solve for each subdomain face F (dense blocks) [C. Pechstein, C. R. Dohrmann, 2013]

$$(\widetilde{S}_F^{(i)} : \widetilde{S}_F^{(j)})\phi = \mu(S_F^{(i)} : S_F^{(j)})\phi$$

  where

$$A : B = (A^{-1} + B^{-1})^{-1}, \quad \widetilde{S}_F^{(k)} = S^{(k)}_{FF} - S^{(k)}_{FF'} S^{(k)-1}_{F'F'} S^{(k)}_{F'F}$$

- Add to the primal space

$$\widehat{\mathbf{W}}_\Pi \leftarrow \{(S_F^{(i)} : S_F^{(j)})\phi_k \mid 1/\mu_k > \lambda\}$$

- More complicated formulas for edges in 3D lead to fully provable condition number bounds [A. Klawonn, M. Kühn and O. Rheinbach, 2015], [H.H. Kim, E.T. Chung and J. Wang, 2015], [J. Calvo, O. B. Widlund, 2015]

Current BDDC code in PETSc 3.6 considers

$$(S_F^{(i)} : S_F^{(j)})\phi = \mu(\widetilde{S}_F^{(i)} : \widetilde{S}_F^{(j)})\phi$$
$$(S_E^{(i)} : S_E^{(j)} : S_E^{(k)})\phi = \mu(\widetilde{S}_E^{(i)} : \widetilde{S}_E^{(j)} : \widetilde{S}_E^{(k)})\phi$$

- MUMPS or MKL_PARDISO: factorize $\mathrm{A}^{(i)}$ and compute $S^{(i)}$ at once
- Factorization of $A_{II}^{(i)}$ could be reused.
- $S_F^{(i)}$ explicitly inverted.
- $\widetilde{S}_F^{(i)-1}$ obtained by explicitly inverting $S^{(i)}$
- $S^{(i)-1}$ is reused to solve the substructure correction
- Nearest neighbor communication to assemble the "sum of Schurs"
- LAPACK used to solve each local GEP

Computational challenges in 3D (2D is easy)

- Costs of the factorizations: flops $O(n_i^2)$, memory $O(n_i^{4/3})$.

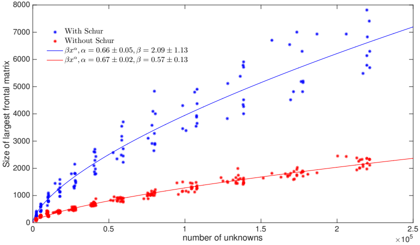

- Costs of the factorizations: flops $O(n_i^2)$, memory $O(n_i^{4/3})$.

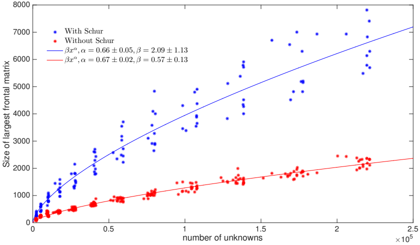

- Explicit inversion of $S^{(i)}$: $O(n_{\Gamma^i}^3)$, n_{Γ^i} could be poorly load balanced

procs	S		S^{-1}		GEP	
	min	max	min	max	min	max
8192	1.02	2.23	0.25	4.82	0.05	0.52
16384	0.98	2.32	0.23	5.28	0.08	0.49
32768	0.94	2.30	0.24	5.57	0.06	0.68

- Precompute basis functions $O(n_i^{4/3})$

Future directions

Future directions

- Exploit many-/multi- cores acceleration (does not solve the complexity problem)

Future directions

- Exploit many-/multi- cores acceleration (does not solve the complexity problem)
- Multiple subdomains per core
  flops $O(n_i^2)$, mem $O(n_i^{4/3}) \rightarrow$ flops $O((n_i/p)^2)$, mem $O((n_i/p)^{4/3})$

- Exploit many-/multi- cores acceleration (does not solve the complexity problem)
- Multiple subdomains per core
  flops $O(n_i^2)$, mem $O(n_i^{4/3}) \rightarrow$ flops $O((n_i/p)^2)$, mem $O((n_i/p)^{4/3})$
- H-solvers for $S^{(i)}$

- Exploit many-/multi- cores acceleration (does not solve the complexity problem)
- Multiple subdomains per core
  flops $O(n_i^2)$, mem $O(n_i^{4/3}) \rightarrow$ flops $O((n_i/p)^2)$, mem $O((n_i/p)^{4/3})$
- H-solvers for $S^{(i)}$
- Exact factorizations of $A_{II}^{(i)}$ are not strictly required. However, the solver $\overline{A}_{II}^{(i)-1}$ has to be spectrally equivalent.

- Exploit many-/multi- cores acceleration (does not solve the complexity problem)
- Multiple subdomains per core
  flops $O(n_i^2)$, mem $O(n_i^{4/3}) \rightarrow$ flops $O((n_i/p)^2)$, mem $O((n_i/p)^{4/3})$
- H-solvers for $S^{(i)}$
- Exact factorizations of $A_{II}^{(i)}$ are not strictly required. However, the solver $\overline{A}_{II}^{(i)-1}$ has to be spectrally equivalent.

Can we partially factorize (with a given accuracy) and obtain $\overline{A}_{II}^{(i)-1}$ and a (possibly hierarchical) $\overline{S}^{(i)}$?

Numerical results: experimental setting

- FENICS library for finite elements: http://fenicsproject.org.
- ParMETIS for mesh partitioning.
- PETSc-dev (branch stefano_zampini/feature-pcbddc-saddlepoint)
- One MPI process/subdomain/core.
- Intel MKL 11.2.2 as BLAS/LAPACK backend.
- MUMPS for local problems and Schur complements.
- PCG with random rhs, zero initial guesses, rtol 1.e-8.

Let $\mathbf{U} = \{\mathbf{u} \in H(\mathrm{div}, \Omega) \;:\; \mathbf{u} \cdot \mathbf{n} = 0 \text{ on } \partial\Omega\}$, $\mathrm{P} = L^2(\Omega)$. Find $(\mathbf{u}, p) \in \mathbf{U} \times \mathrm{P}$ s.t. $\forall\ (\mathbf{v}, q) \in \mathbf{U} \times \mathrm{P}$

$$\begin{cases} \int_\Omega \mathbf{u}\ K(x)^{-1}\ \mathbf{v}\ dx + \int_\Omega p \ \mathrm{div}\ \mathbf{v}\ dx & = 0,\ \forall\ \mathbf{v} \in \mathbf{U}, \\ \int_\Omega \mathrm{div}\ \mathbf{u}\ q\ dx & = \int_\Omega gq,\ \forall\ q \in \mathrm{P}, \end{cases}$$

where $K(x)$ uniformly positive definite tensor (mixed BC can be handled as well).
Discretized with RT_0 or BDM_1 elements for velocities and C^0 polynomials elements (discontinuous) for pressures (LBB stable)

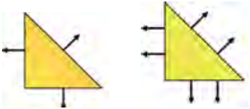

[Brezzi and Fortin, Mixed and Hybrid Finite elements methods, 1993]

Joint work with X. Tu.

[S. Z. and X. Tu, 2016, submitted]

SPE10 benchmark http://www.spe.org/web/csp/datasets/set02.htm:

- $\Omega = 1200\text{ft} \times 2200\text{ft} \times 170\text{ft}$
- mesh $60 \times 220 \times 85$, each hexahedron subdivided in 6 tetrahedra
- 6.7M cells
- dofs: 20.2M with RT_0, 45M with BDM_1.
- Diagonal permeability tensor's coefficients from SPE10

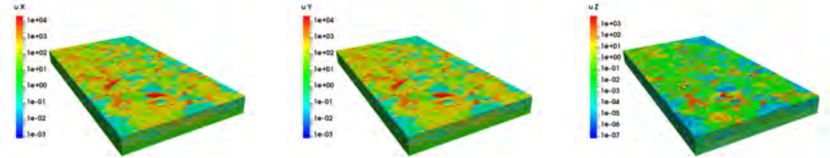

Condition number and number of iterations as a function of eigenvalue threshold λ and number of subdomains N. Left RT_0, right BDM_1

N	$\lambda = 10$	$\lambda = 5$	$\lambda = 2.5$	$\lambda = 1.5$
1024	15.9/25	7.77/17	3.57/11	1.73/6
2048	15.0/25	7.76/17	3.51/11	1.65/6
4096	15.4/25	8.19/18	3.42/11	1.68/6
8192	16.5/26	7.69/17	3.51/11	1.67/6

N	$\lambda = 10$	$\lambda = 5$	$\lambda = 2.5$	$\lambda = 1.5$
1024	16.1/24	7.51/16	3.49/10	1.60/6
2048	16.7/25	7.45/16	3.53/10	1.61/6
4096	15.5/24	7.53/16	3.57/10	1.58/6
8192	15.9/24	7.77/17	3.53/10	1.59/6

Coarsening ratios (F/C,Γ/C) as a function of eigenvalue threshold λ and number of subdomains N. Left RT_0, right BDM_1

N	$\lambda = 10$	$\lambda = 5$	$\lambda = 2.5$	$\lambda = 1.5$
1024	745/25	481/16	281/9	172/5
2048	402/17	272/11	165/7	105/4
4096	218/11	153/8	98/5	64/3
8192	121/8	89/5	60/4	41/2

N	$\lambda = 10$	$\lambda = 5$	$\lambda = 2.5$	$\lambda = 1.5$
1024	1378/61	795/35	408/18	230/10
2048	761/42	458/25	243/13	140/7
4096	420/29	263/18	146/10	86/6
8192	236/20	155/13	90/7	54/4

Setup/Solve times for 2-levels BDDC as a function of eigenvalue threshold λ and number of subdomains N. Left RT_0, right BDM_1

N	$\lambda = 10$	$\lambda = 5$	$\lambda = 2.5$	$\lambda = 1.5$
1024	3.4/1.1	3.9/0.9	5.3/**0.8**	8.7/0.9
2048	**3.3**/1.1	4.4/1.1	6.0/1.2	11.1/1.2
4096	5.6/1.7	7.4/1.7	11.0/1.5	20.8/2.1
8192	9.6/3.0	15.8/4.4	21.0/3.6	48.1/3.9

N	$\lambda = 10$	$\lambda = 5$	$\lambda = 2.5$	$\lambda = 1.5$
1024	25.6/4.7	26.1/3.4	31.0/2.9	47.7/3.1
2048	10.6/2.4	12.0/**1.9**	18.2/2.1	61.5/3.6
4096	**7.9**/2.4	9.8/2.2	27.4/3.0	130.9/4.3
8192	15.7/3.5	18.2/5.8	50.2/5.9	**/**

Setup/Solve times for 3-levels BDDC (coarsening ratio 16, coarse threshold 10, 2 Chebyshev its, eigs computed) as a function of eigenvalue threshold λ and number of subdomains N. Left RT_0, right BDM_1

N	$\lambda = 10$	$\lambda = 5$	$\lambda = 2.5$	$\lambda = 1.5$
1024	3.7/1.4	6.4/1.2	6.5/1.2	6.6/1.6
2048	**2.9**/**1.1**	3.9/1.2	4.4/1.1	5.5/1.5
4096	4.9/1.2	5.3/1.4	5.8/1.3	6.0/1.2
8192	7.8/2.0	7.8/1.6	7.9/2.1	9.1/1.8

N	$\lambda = 10$	$\lambda = 5$	$\lambda = 2.5$	$\lambda = 1.5$
1024	24.2/5.7	25.2/3.9	28.3/4.4	41.7/5.3
2048	10.2/2.8	10.7/2.2	14.7/2.9	20.2/3.6
4096	**6.2**/2.1	7.2/**1.8**	12.3/2.9	11.5/2.6
8192	9.2/2.4	11.7/2.5	11.9/2.1	13.9/2.3

Let $\mathbf{U} = \{\mathbf{u} \in H(\text{curl}, \Omega) \ : \ \mathbf{u} \times \mathbf{n} = 0 \text{ on } \partial\Omega\}$.
Consider the following problem: find $\mathbf{u} \in \mathbf{U}$ s.t. $\forall\ \mathbf{v} \in \mathbf{U}$

$$\int_\Omega \alpha \nabla \times \mathbf{u} \cdot \nabla \times \mathbf{v}\ dx + \int_\Omega \beta\ \mathbf{u}\ \cdot \mathbf{v}\ dx = \int_\Omega \mathbf{f} \cdot \mathbf{v}$$

Arises in time-domain quasi-static approximation of Maxwell's equations:

- $\alpha = \mu_0^{-1}$ (μ_0 the magnetic permeability)
- $\beta = \sigma/\delta_t$, with $\sigma > 0$ conductivity (anisotropic case could be handled as well)
- δ_t the time step.

or in block preconditioning of frequency domain problems.
Discretized with Nédélec elements on tetrahedra.

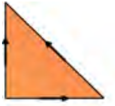

[Brezzi and Fortin, Mixed and Hybrid Finite elements methods, 1993]

[S. Z., submitted, 2015].

Threshold test, fixed mesh. Dofs from 800K to 11M (depending on the order), 64 subdomains.

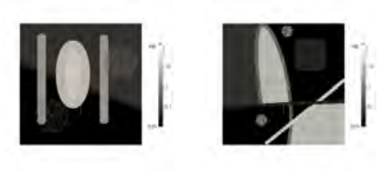

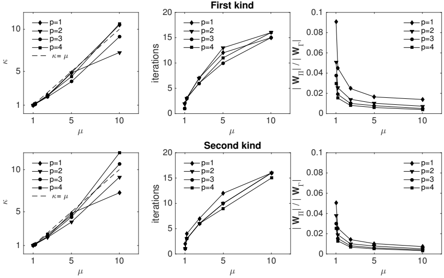

Threshold 1.01, fixed mesh. Dofs from 50K to 11M, 64 subdomains.

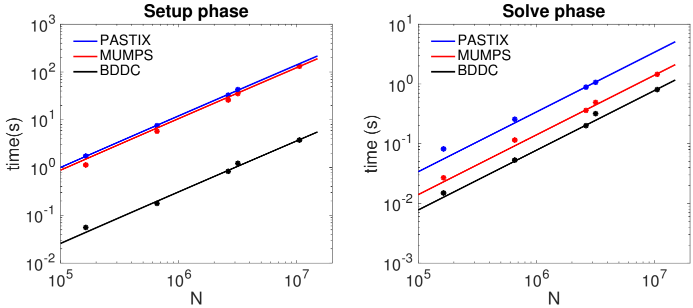

Comparision against MUMPS and PASTIX parallel Cholesky

BDDC for CSEM: reservoir test

- Domain 10km $\times$ 8km $\times$ 4km
- Mesh $100 \times 80 \times 40$ + unif. refin.
- 18M dofs, $\lambda = 10$
- $\delta_t = 0.01$

$$\sigma = \begin{cases} 3.6 & \text{Ocean} \\ 0.5 & \text{Background} \\ 0.02 & \text{Reservoir} \end{cases}$$

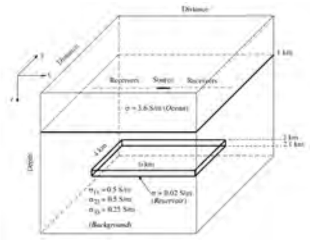

Figure taken from [Carcione, Progr. Electromagnetics Res. B, 26, 2010]

N	t_{set}	t_{sol}	it	n_I	n_Γ	mem_I	mem_c	rel_c	eff
256	110.8	16.1	25	75.1	8.5	136	0.09	0.01	4.3
512	31.4	6.3	25	38.1	5.5	110	0.32	0.02	5.6
1024	11.1	2.5	23	19.5	3.4	89	0.92	0.03	7.0
2048	8.2	1.5	21	9.9	2.1	69	2.70	0.06	5.9
4096	11.6	1.3	19	5.1	1.3	57	7.95	0.10	3.3
8192	28.1	2.9	19	2.6	0.8	49	23.7	0.17	0.7

n_I, n_Γ average number (K) of subdomain and local interface dofs, mem_I and mem_c in GB, eff = eqs (K) /core/second