This is a pre-copyedited version of a contribution published in Haynes, R. et al. (eds) Domain Decomposition Methods in Science and Engineering XXV. Lecture Notes in Computational Science and Engineering, vol 138 published by Springer Nature Switzerland AG. The definitive authenticated version is available online via https://link.springer.com/chapter/10.1007/978-3-030-56750-7_26.

A closer look at local eigenvalue solvers for adaptive FETI-DP and BDDC

Axel Klawonn[1,2], Martin J. Kühn[3], and Oliver Rheinbach[4]

1 Introduction

In order to obtain a scalable domain decomposition method (DDM) for elliptic problems, a coarse space is necessary and an associated coarse problem has to be solved in each iteration. In the presence of arbitrary, large coefficient jumps or in case of almost incompressible elastic materials, the convergence rate of standard DDM deteriorates. In recent years, many authors have proposed the use of different (local, generalized) eigenvalue problems to develop problem dependent, adaptive coarse spaces in order to ensure or accelerate the convergence of the method; see, e.g., [2, 17, 6, 18, 5, 4, 22, 23, 14, 9, 10, 1, 19, 3, 20]. These methods are very robust and in many cases, a condition number estimate of the form

$$\text{cond} \leq C\,\text{TOL} \tag{1}$$

exists. Here, TOL is an a priori, user defined tolerance for the solution of the eigenvalue problems and $C > 0$ a constant that only depends on geometric parameters, e.g., maximum number of edges of a subdomain; cf., (3). However, in order to make their use feasible, many issues have to be considered in the parallel implementation. In [16, 13], we have seen that the computational overhead of the solution process of the local eigenvalue problems in adaptive FETI-DP is not negligible. Consequently, in this paper, we focus on some aspects of the local eigenvalue solution process that

[1]University of Cologne, Faculty of Mathematics and Natural Sciences, Mathematical Institute, Weyertal 86-90, 50931 Köln, Germany, e-mail: axel.klawonn@uni-koeln.de, url: http://www.numerik.uni-koeln.de
[2]University of Cologne, Center for Data and Simulation Science, url: http://www.cds.uni-koeln.de
[3]CERFACS (Centre Européen de Recherche et de Formation Avancée en Calcul Scientifique), 42 Avenue Gaspard Coriolis, 31057 Toulouse Cedex 01, France; martin.kuehn@cerfacs.fr.
[4]Technische Universität Bergakademie Freiberg, Fakultät für Mathematik und Informatik, Institut für Numerische Mathematik und Optimierung, 09596 Freiberg, Germany; oliver.rheinbach@math.tu-freiberg.de

have not been studied or documented elsewhere. Certainly, load balancing of the eigenvalue problems is a very important task but this issue is out the scope of this paper and will be discussed in detail in [13].

2 Model problem, domain decomposition, and notation

As a model problem, we consider three-dimensional linear elasticity, discretized with piecewise quadratic conforming finite elements. The domain is decomposed into nonoverlapping subdomains. Due to page restrictions, for further details, we refer to [10, Section 2] or [12, Section 2].

As it is standard in FETI-DP, we assemble the local stiffness matrices $K^{(i)}$ and compute the local Schur complements $S^{(i)}$ on the interface, $i = 1, \ldots, N$. Starting with the block-diagonal matrix S built from the the local Schur complements, we get the global matrix $\widetilde{S}$ by finite element subassembly in only a few a priori chosen primal variables (i.e., all vertices). In order to enforce continuity on the remaining a priori dual degrees of freedom on the interface, we introduce a jump operator B as well as a scaled variant B_D. Different scaling choices are available in the literature. We then obtain the FETI-DP system, which is reduced to the Lagrange multipliers enforcing continuity in the a priori dual variables,

$$M_D^{-1} F := B_D S B_D^T B \widetilde{S}^{-1} B \lambda = B_D S B_D^T d = M_D^{-1} d$$

with corresponding right hand side d; see, e.g., [12, Section 3] for further details.

3 Adaptive FETI-DP

In adaptive FETI-DP, as proposed in two dimensions in [17], and in three dimensions in [10], local generalized eigenvalue problems are solved on each pair of subdomains Ω_i and Ω_j sharing either a face $\mathcal{Z} = \mathcal{F}$ or an edge $\mathcal{Z} = \mathcal{E}$. By extracting all the rows of B and B_D corresponding to dual degrees of freedom of Ω_i and Ω_j and belonging to the closure of $\mathcal{Z}$, we can define the localized variants $B_\mathcal{Z}$, $B_{D,\mathcal{Z}}$ and $P_{D,\mathcal{Z}} := B_{D,\mathcal{Z}}^T B_\mathcal{Z}$. The localized Schur complement is defined as $S_{ij} := \operatorname{blockdiag}\big(S^{(i)}, S^{(j)}\big)$. The local generalized eigenvalue problem then writes: find $w_{ij} \in (\ker S_{ij})^\perp$ with $\mu_{ij} > \mathrm{TOL}$, such that

$$(P_{D,\mathcal{Z}^{ij}} v_{ij}, S_{ij} P_{D,\mathcal{Z}^{ij}} w_{ij}) = \mu_{ij} (v_{ij}, S_{ij} w_{ij}) \quad \forall v_{ij} \in (\ker S_{ij})^\perp, \tag{2}$$

cf. [17, Sections 3 and 4] and [10, Section 5] or [16, Section 5] for a more detailed description.

The constraints obtained on a local basis can then be enforced by different techniques. Here, we use the generalized transformation-of-basis approach proposed

in [11] and obtain the preconditioner $\widehat{M}_T^{-1}$, the modified system matrix $\widehat{F}$, and the condition number bound

$$\kappa(\widehat{M}_T^{-1}\widehat{F}) \leq 4 \max\{N_{\mathcal{F}}, N_{\mathcal{E}} M_{\mathcal{E}}\}^2 \text{TOL}, \tag{3}$$

where $N_{\mathcal{F}}$ denotes the maximum number of faces of a subdomain, $N_{\mathcal{E}}$ the maximum number of edges of a subdomain, $M_{\mathcal{E}}$ the maximum multiplicity of an edge; see [12].

4 Numerical results

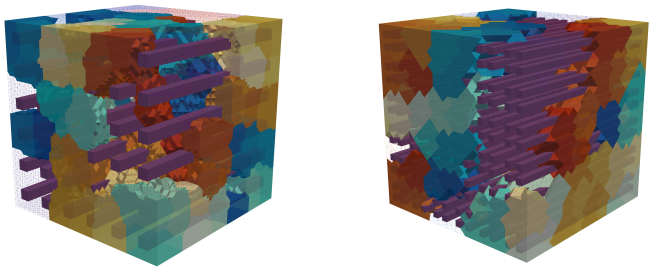

Fig. 1 A composite material on the unit cube for 216 subdomains: 36 beams (left) and 64 beams (right) of a stiff material with $E_2 = 1e+6$, shown in dark purple, are surrounded by a soft matrix material with $E_1 = 1$. A part of the mesh with $1/h = 54$ (left), $1/h = 30$ (right) and the irregular decomposition using METIS is shown in different (half-transparent) colors.

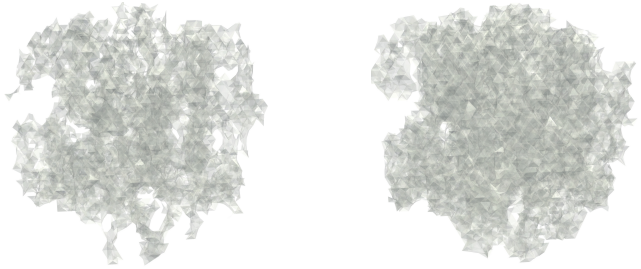

Fig. 2 Stiff material in a foam-like structure with ∼15% (left) and ∼26% (right) high coefficients with $E_2 = 1e+6$. The structure is surrounded by a soft matrix material with $E_1 = 1$. The stiff material is shown smoothed and half-transparent, the surrounding matrix material is not shown.

In the following, we consider the unit cube with zero Dirichlet boundary conditions on the face with $x_1 = 0$ and zero Neumann boundary conditions elsewhere. The domain decomposition is obtained from the METIS partitioner using the options `-ncommon=3` and `-contig`. We apply our adaptive method to four different materials.

The first material is considered for $1/h \in \{30, 54\}$ and 36 beams with a Young modulus of 1e+6 that run from the face with $x_1 = 0$ to the face with $x_1 = 1$; see fig. 1 (left). The remaining part of the material has a Young modulus of one. In the second material, we have a larger number of 64 thinner beams; see fig. 1 (right).

The third and fourth materials are stiff foam-like materials surrounded by a soft matrix material. They are obtained by using a pseudo-random number generator and adjacency structures of the tetrehedra and they merely differ by the amount of stiff material inside the unit cube; see fig. 2.

We now focus on the solution process for the local eigenvalue problems. In recent works, [10, 16, 13], we have developed heuristic strategies to discard eigenvalue problems based on the coefficients, or in a more realistic setting, based on scaling information or the entries of the stiffness matrix which are more likely to be available. In following, we will show that our most recent heuristic strategy (see [13]) is successful in discarding unnecessary eigenvalue problems without, on the other hand, discarding necessary ones.

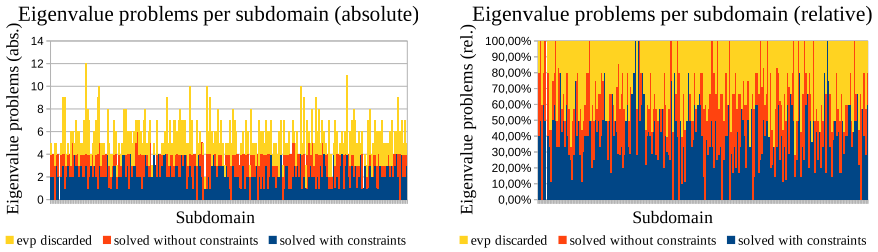

Fig. 3 Number of eigenvalue problems for each subdomain: discarded by our heuristic strategy [13] in yellow; solved and yielding constraints in blue; solved but not resulting in constraints in red; for the composite material with 36 beams; in absolute (left) and relative (right) numbers.

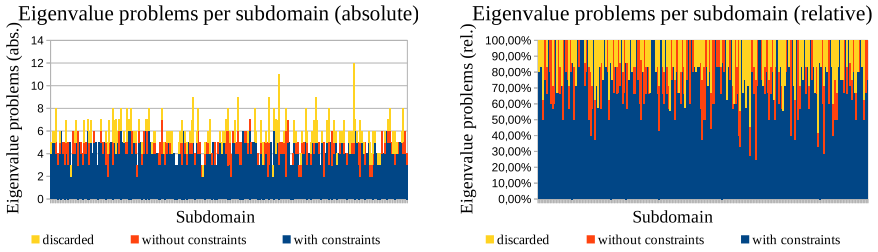

Fig. 4 Number of eigenvalue problems for each subdomain: discarded by our heuristic strategy [13] in yellow; solved and yielding constraints in blue; solved but not resulting in constraints in red; for the composite material with 64 beams; in absolute (left) and relative (right) numbers.

In the present implementation, for every face or edge, we consider the diagonal entries of the local subdomain stiffness matrices corresponding to interior nodes of these faces or edges. We have two criteria. First, if the ratio of the smallest and the largest entry is larger than a certain threshold or possibly second, if all entries are large, then the corresponding eigenvalue problem is solved. Otherwise it is discarded. For a more detailed description, see [13].

There are different situations in which eigenvalue problems become superfluous for the reduction of the condition number. One obvious reason is the nonexistence of jumps in the neighborhood of the face or edge. One could then apply slab techniques; see, e.g. [21, 7]. In our heuristics, we focus on these eigenvalue problems which we classify as *unnecessary*.

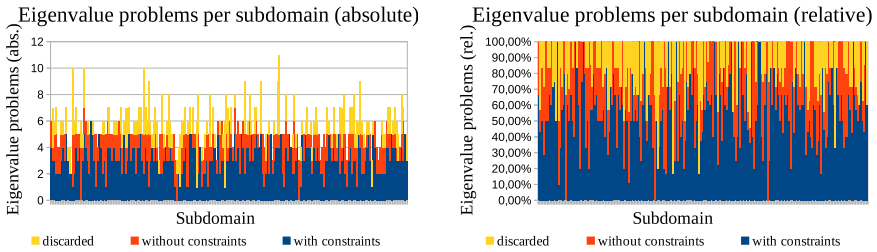

Fig. 5 Number of eigenvalue problems for each subdomain: discarded by our heuristic strategy [13] in yellow; solved and yielding constraints in blue; solved but not resulting in constraints in red; for the foam-like composite with ∼15% high coefficients; in abs. (left) and rel. (right) numbers.

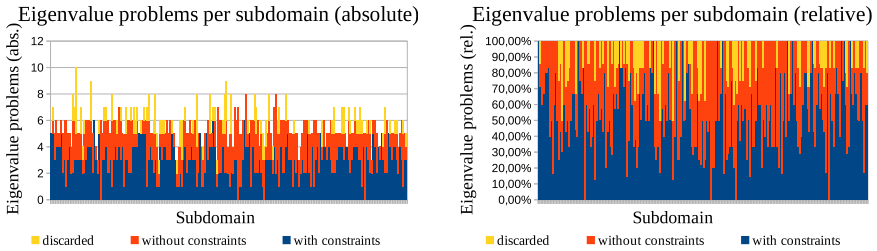

Fig. 6 Number of eigenvalue problems for each subdomain: discarded by our heuristic strategy [13] in yellow; solved and yielding constraints in blue; solved but not resulting in constraints in red; for the foam-like composite with ∼26% high coefficients; in abs. (left) and rel. (right) numbers.

However, there might be eigenvalue problems on coefficient distributions which satisfy all assumptions for weighted Poincare inequalities. For arbitrary situations, the numerical necessity of certain eigenvalue problems becomes even more complex and is not yet fully understood. In [8], we have presented a short numerical study which gives a little more insight for typical situations but which might also raise new questions since not all configurations introduce as many bad modes as expected.

In the case of 36 beams, we see that we can discard a large number of eigenvalue problems (i.e., 37%) while 38% of eigenvalue problems yield large eigenvalues and thus adaptive constraints; see fig. 3. In the case of 64 beams, we see that we only discard 16% of eigenvalue problems but, here, more than 70% of eigenvalue problems yield large eigenvalues and thus adaptive constraints; see fig. 4. In both cases, the percentage of eigenvalue problems that were solved without yielding constraints is small, i.e., 25% and 14%. For both foam-like composites a little more

than 50% of the eigenvalue problems have to be solved to reduce the condition number to about TOL (here TOL $= 50 \log\left(N/n_i\right)^{1/3}$, where n_i is the number of local nodes). Between 26% and 39% of the eigenvalue problems were solved there without yielding constraints; see fig. 5 and fig. 6. Still 10% to 20% of the eigenvalue problems are detected as discardable, and thus the total algorithm is accelerated. We can summarize that our strategy can successfully identify and discard many unnecessary eigenvalue problems while keeping all necessary ones.

In order to give more insight into the eigenvalue problems that are solved, we present the number of constraints yielded for each eigenvalue problem for each material in four different pie charts; see fig. 5 and fig. 6. We see that the number of constraints for each eigenvalue problem range from 1 to 28. However, a large majority always gives between 2 and 12 constraints.

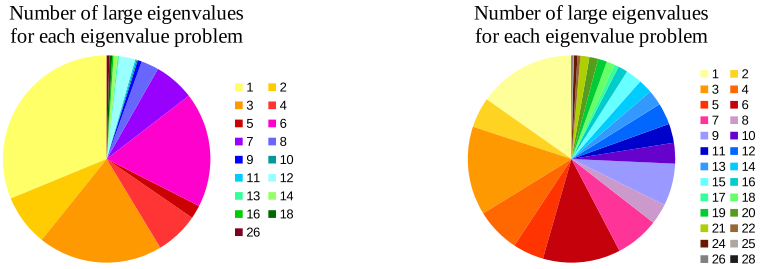

Fig. 7 Number of eigenvalue problems with given number of large eigenvalues for the composite material with 36 (left) and 64 (right) beams. In these presentations, only the blue marked eigenvalue problems of fig. 3 and fig. 4 are considered to give more details.

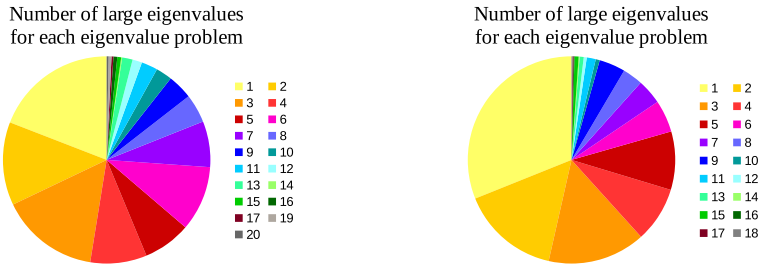

Fig. 8 Number of eigenvalue problems with given number of large eigenvalues for the foam-like composite material with 15% (left) and 26% (right) high coefficients. In these presentations, only the blue marked eigenvalue problems of fig. 5 and fig. 6 are considered to give more details.

Finally, we focus on the important topic of block sizes in the SLEPc Krylov-Schur solver. As motivated by [17] and our tests in [10, 12], we have already opted for an approximate solution of the eigenvalue problems by carrying out only a few steps of the iterative block scheme. Justified by the idea that the LOBPCG block solver of [15] could accelerate the convergence on extreme eigenvalues we have mostly

used a block size of 10. Here, we study the timings of the global algorithms by varying the block size of the local Krylov-Schur algorithm for our four materials.

In table 1, we have presented iteration counts and estimated condition numbers in order to show that the chosen block size does not effect the convergence of the global PCG scheme. That means that the constraints obtained with different block sizes do not differ in quality. In fig. 9, we see that the use of smaller block sizes or even a single vector iteration might be favorable with respect to time to solution.

block size Krylov-Schur	36 beams ($1/h = 36$)		64 beams ($1/h = 30$)		~15% foam-like ($1/h = 30$)		~26% foam-like ($1/h = 30$)	
	κ	*its*	κ	*its*	κ	*its*	κ	*its*
1	5.48e+01	60	6.29e+1	62	7.21e+01	62	5.99e+01	63
3	5.48e+01	60	6.29e+1	62	7.21e+01	61	5.99e+01	62
6	5.48e+01	60	6.29e+1	62	7.21e+01	61	5.99e+01	61
9	5.48e+01	60	6.30e+1	63	7.21e+01	62	5.99e+01	64
12	5.48e+01	60	6.29e+1	62	7.21e+01	61	5.99e+01	61
15	5.48e+01	61	6.29e+1	62	7.21e+01	62	5.99e+01	61
18	5.49e+01	61	6.30e+1	63	7.21e+01	61	5.99e+01	61

Table 1 Condition number and iteration count of the global FETI-DP solver for different composite materials for 216 subdomains with different block sizes for the iterative Krylov-Schur eigenvalue solver and TOL = $50 \log\left(N/n_i\right)^{1/3}$, where n_i is the number of local nodes.

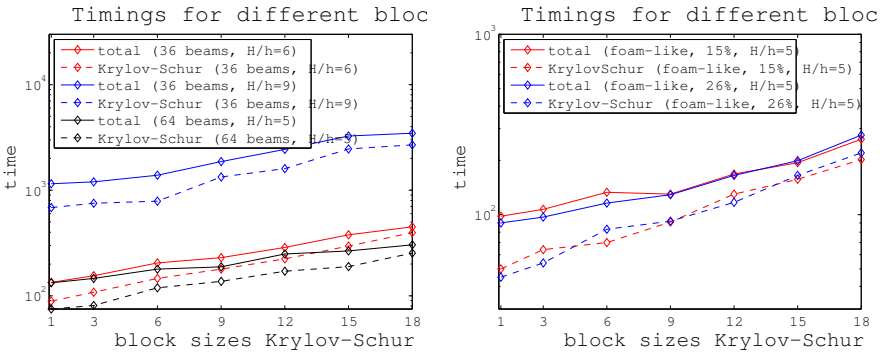

Fig. 9 Total global time and total local time needed by the Rayleigh-Ritz procedures in the Krylov-Schur scheme to approximately compute the largest eigenvectors of the generalized eigenvalue problems. Composite with 36 beams and 64 beams (left) and foam-like composite with ~15% and ~26% high coefficients (right).

References

1. 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. SIAM J. Sci. Comput., 39(1):A281–A302, 2017.
2. Petter Bjørstad, Jacko Koster, and Piotr Krzyżanowski. Domain decomposition solvers for large scale industrial finite element problems. In Applied Parallel Computing. New Paradigms for HPC in Industry and Academia, volume 1947 of Lecture Notes in Comput. Sci., pages 373–383. Springer, Berlin, 2001.
3. Juan G. Calvo and Olof B. Widlund. An adaptive choice of primal constraints for BDDC domain decomposition algorithms. Electron. Trans. Numer. Anal., 45:524–544, 2016.

4. Clark Dohrmann and Clemens Pechstein. In C. Pechstein, Modern domain decomposition solvers - BDDC, deluxe scaling, and an algebraic approach. NuMa Seminar, http://people.ricam.oeaw.ac.at/c.pechstein/pechstein-bddc2013.pdf, 2013.
5. Victorita Dolean, Frédéric Nataf, Robert Scheichl, and Nicole Spillane. Analysis of a two-level Schwarz method with coarse spaces based on local Dirichlet-to-Neumann maps. Comput. Methods Appl. Math., 12(4):391–414, 2012.
6. Juan Galvis and Yalchin Efendiev. Domain decomposition preconditioners for multiscale flows in high-contrast media. Multiscale Model. Simul., 8(4):1461–1483, 2010.
7. Sabrina Gippert, Axel Klawonn, and Oliver Rheinbach. Analysis of FETI-DP and BDDC for linear elasticity in 3D with almost incompressible components and varying coefficients inside subdomains. SIAM J. Numer. Anal., 50(5):2208–2236, 2012.
8. Alexander Heinlein, Axel Klawonn, and Martin Joachim Kühn. Local spectra of adaptive domain decomposition methods. Submitted to the Proceedings of the 25th International Conference on Domain Decomposition Methods, DD XXV, St. John's NL, July, 23-27, 2018. Also Preprint at https://kups.ub.uni-koeln.de/9019/.
9. Hyea Hyun Kim and Eric T. Chung. A BDDC algorithm with enriched coarse spaces for two-dimensional elliptic problems with oscillatory and high contrast coefficients. Multiscale Model. Simul., 13(2):571–593, 2015.
10. Axel Klawonn, Martin Kühn, and Oliver Rheinbach. Adaptive coarse spaces for FETI-DP in three dimensions. SIAM J. Sci. Comput., 38(5):A2880–A2911, 2016.
11. Axel Klawonn, Martin Kühn, and Oliver Rheinbach. Coarse spaces for FETI-DP and BDDC methods for heterogeneous problems: Connections of deflation and a generalized transformation-of-basis approach. Technical report, 2017. Preprint 2017-01, http://tu-freiberg.de/fakult1/forschung/preprints; submitted for publication.
12. Axel Klawonn, Martin Kühn, and Oliver Rheinbach. Adaptive FETI-DP and BDDC methods with a generalized transformation of basis for heterogeneous problems. Electron. Trans. Numer. Anal., 49:1–27, 2018.
13. Axel Klawonn, Martin Kühn, and Oliver Rheinbach. Parallel adaptive FETI-DP using lightweight asynchronous dynamic load balancing. Technical report, 2018. Technical Report CDS-2019-7, https://kups.ub.uni-koeln.de/9368/; submitted for publication.
14. Axel Klawonn, Patrick Radtke, and Oliver Rheinbach. FETI-DP methods with an adaptive coarse space. SIAM J. Numer. Anal., 53(1):297–320, 2015.
15. Andrew V. Knyazev. Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. SIAM J. Sci. Comput., 23(2):517–541, 2001.
16. Martin Joachim Kühn. Adaptive FETI-DP and BDDC methods for highly heterogeneous elliptic finite element problems in three dimensions. PhD thesis, Universität zu Köln, 2018.
17. 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.
18. Frédéric Nataf, Hua Xiang, and Victorita Dolean. A two level domain decomposition preconditioner based on local Dirichlet-to-Neumann maps. C. R. Math. Acad. Sci. Paris, 348(21-22):1163–1167, 2010.
19. 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. Math. Comp., 87(310):659–692, 2018.
20. Clemens Pechstein and Clark R. Dohrmann. A unified framework for adaptive BDDC. Electron. Trans. Numer. Anal., 46:273–336, 2017.
21. Clemens Pechstein and Robert Scheichl. Analysis of FETI methods for multiscale PDEs. Numer. Math., 111(2):293–333, 2008.
22. Nicole Spillane, Victorita Dolean, Patrice Hauret, Frédéric Nataf, Clemens Pechstein, and Robert Scheichl. Abstract robust coarse spaces for systems of PDEs via generalized eigenproblems in the overlaps. Numer. Math., 126(4):741–770, 2014.
23. Nicole Spillane and Daniel J. Rixen. Automatic spectral coarse spaces for robust finite element tearing and interconnecting and balanced domain decomposition algorithms. Internat. J. Numer. Methods Engrg., 95(11):953–990, 2013.