Using Algebraic Multigrid in Inexact BDDC Domain Decomposition Methods

Axel Klawonn, Martin Lanser, and Oliver Rheinbach

1 Introduction

Traditionally, domain decomposition methods use sparse direct solvers as building blocks, i.e., to solve local subdomain problems and/or the coarse problem. Often, the sparse direct solvers can be replaced by spectrally equivalent preconditioners without loss of convergence speed. In FETI-DP and BDDC domain decomposition methods, such approaches have first been introduced in [9, 8, 4], and have since then successfully been used in large parallel codes [6, 1].

2 An Inexact BDDC Method

2.1 A BDDC Preconditioner for the Assembled System

Let us briefly describe the BDDC preconditioner which can directly be applied to a linear system

$$Au = b \tag{1}$$

arising from a finite element discretization of a partial differential equation on a computational domain $\Omega \subset \mathbb{R}^d$, $d = 2,3$. The variant discussed here was first introduced in [9]. Let Ω_i, $i = 1,\dots,N$, be a nonoverlapping domain decomposition of Ω such that $\overline{\Omega} = \bigcup_{i=1}^{N} \overline{\Omega}_i$. Each subdomain Ω_i is discretized using finite elements,

Axel Klawonn, Martin Lanser
Mathematisches Institut, Universität zu Köln, Weyertal 86-90, 50931 Köln, Germany, e-mail: {axel.klawonn,martin.lanser}@uni-koeln.de

Oliver Rheinbach
Institut für Numerische Mathematik und Optimierung, Fakultät für Mathematik und Informatik, Technische Universität Bergakademie Freiberg, Akademiestr. 6, 09596 Freiberg, e-mail: oliver.rheinbach@math.tu-freiberg.de

the corresponding local finite element spaces are denoted by W_i, $i = 1,\dots,N$, and the product space is defined by $W = W_1 \times \dots \times W_N$. Let us also introduce the global finite element space V^h corresponding to the discretization of Ω and a restriction $R : V^h \to W$. We obtain local problems in the spaces W_i

$$K_i u_i = f_i,\ i = 1,\cdots,N.$$

Introducing the block operators

$$K = \begin{pmatrix} K_1 & & \\ & \ddots & \\ & & K_N \end{pmatrix},\ f = \begin{pmatrix} f_1 \\ \vdots \\ f_N \end{pmatrix},$$

we can write $A := R^T K R$ and $b := R^T f$. Finally, the interface between the subdomains is $\Gamma := \bigcup_{i=1}^{N} \partial\Omega_i \setminus \partial\Omega$. Let us assume that the degrees of freedom (d.o.f.) on the Dirichlet boundary $\partial\Omega_D \subset \partial\Omega$ are eliminated.

We use the index Γ for degrees of freedom on Γ. For degrees of freedom in the interior of the subdomains and on the Neumann boundary $\partial\Omega_N \subset \partial\Omega$, we use the index I. For the construction of a BDDC preconditioner directly applicable to the assembled linear system $Au = b$, we subdivide, as usual in BDDC and FETI-DP methods, the interface Γ into primal (Π) and the remaining dual (Δ) degrees of freedom. As primal variables usually subdomain vertices or averages over edges or faces are chosen.

Let us introduce the space $\widetilde{W} \subset W$ of functions, which are continuous in all primal variables and the restriction operator $\bar{R} : \widetilde{W} \to W$. We can now define a partially assembled system matrix

$$\widetilde{K} := \bar{R}^T K \bar{R} \tag{2}$$

and the corresponding right hand side $\tilde{f} := \bar{R}^T f$. Using a scaled restriction operator $\widetilde{R}_D : V^h \to \widetilde{W}$, we define the BDDC preconditioner by

$$M^{-1}_{BDDC} := \left(\widetilde{R}_D^T - \mathscr{H} P_D\right) \widetilde{K}^{-1} \left(\widetilde{R}_D - P_D^T \mathscr{H}^T\right); \tag{3}$$

see [9]. Here, $\mathscr{H} : \widetilde{W} \to V^h$ is a discrete harmonic extension operator defined by

$$\mathscr{H} := \begin{pmatrix} 0 & -(K_{II})^{-1} \widetilde{K}_{\Gamma I}^T \\ 0 & 0 \end{pmatrix}, \tag{4}$$

where K_{II} and $\widetilde{K}_{\Gamma I}$ are blocks of the partially assembled stiffness matrix

$$\widetilde{K} = \begin{pmatrix} K_{II} & \widetilde{K}_{\Gamma I}^T \\ \widetilde{K}_{\Gamma I} & \widetilde{K}_{\Gamma\Gamma} \end{pmatrix}, \tag{5}$$

which are common to both, BDDC and FETI-DP methods. The matrix K_{II} is block-diagonal and applications of K_{II}^{-1} only require local solves on the interior parts of the subdomains and are thus easily parallelizable.

Finally, let $P_D : \widetilde{W} \to \widetilde{W}$ be a scaled jump operator defined by

$$P_D = I - E_D := I - \widetilde{R}\widetilde{R}_D^T. \tag{6}$$

In the FETI-DP literature this operator is often defined as $P_D = B_D^T B$; see [12, Chapter 6] and [9] for more details. There, B is the standard jump matrix used in FETI-type methods. Let us remark that the preconditioned system $M_{BDDC}^{-1}A$ has, except for some eigenvalues equal to 0 and 1, the same spectrum as the standard BDDC preconditioner formulated on the Schur complement; see [9, Theorem 1]. Therefore, under sufficient assumptions (see [9, Assumption 1]), the condition number of the preconditioned system is bounded by

$$\kappa(M_{BDDC}^{-1}A) \leq \Phi(H,h). \tag{7}$$

For a homogeneous linear elasticity problem, if appropriate primal constraints are chosen, we obtain the well known BDDC (and FETI-DP) condition number bound with $\Phi(H,h) = C(1+\log(H/h))^2$. Here, H always denotes the maximal diameter of all subdomains and h the minimal diameter of all finite elements.

2.2 Using Inexact Solvers and Implementation Remarks

In this paragraph, we describe the use of inexact solvers in the preconditioner M_{BDDC}^{-1} as suggested in [9] and also provide some remarks on our implementation. We assume that $\widehat{K}^{-1}$ and $\widehat{K}_{II}^{-1}$ are spectrally equivalent preconditioners for $\widetilde{K}$ and K_{II}, respectively. In this paper, we always choose a fixed number of V-cycles of an AMG method for solving problems including $\widetilde{K}^{-1}$ and K_{II}^{-1} for those preconditioners. While $\widehat{K}^{-1}$ requires an MPI parallel implementation of an AMG method, an application of $\widehat{K}_{II}^{-1}$ requires only a sequential AMG, due to the block diagonal structure of K_{II}. Using $\widehat{K}_{II}^{-1}$, we define an approximate discrete harmonic extension $\widehat{\mathscr{H}}$ by

$$\widehat{\mathscr{H}} := \begin{pmatrix} 0 & -\widehat{K}_{II}^{-1}\widetilde{K}_{\Gamma I}^T \\ 0 & 0 \end{pmatrix}. \tag{8}$$

We investigate two different variants of the inexact BDDC preconditioner in this paper, namely

$$\widehat{M}_{BDDC,1}^{-1} := \left(\widetilde{R}_D^T - \mathscr{H}P_D\right)\widehat{K}^{-1}\left(\widetilde{R}_D - P_D^T\mathscr{H}^T\right) \tag{9}$$

and

$$\widehat{M}_{BDDC,2}^{-1} := \left(\widetilde{R}_D^T - \widehat{\mathscr{H}}P_D\right)\widehat{K}^{-1}\left(\widetilde{R}_D - P_D^T\widehat{\mathscr{H}}^T\right). \tag{10}$$

Let us remark that in $M_{BDDC,1}^{-1}$ the discrete harmonic extension is applied exactly using a direct solver, while in $M_{BDDC,2}^{-1}$ the approximate discrete harmonic extension

$\widehat{\mathscr{H}}$ is used. Assuming that we have chosen an appropriate $\widehat{K}$, i.e., satisfying

$$\tilde{c}u^T\widetilde{K}u \leq u^T\widehat{K}u \leq \widetilde{C}u^T\widetilde{K}u,\ \forall u \in \widetilde{W}, \tag{11}$$

a condition number bound of the same quality as (7) is valid,

$$\kappa(M^{-1}_{BDDC,2}A) \leq \frac{\widetilde{C}C}{\tilde{c}}(1+\Phi(H,h));$$

see [9, Theorem 4]).

Our parallel implementation uses C/C++ and PETSc version 3.6.4 [3]. While the matrix $\widetilde{K}$ is an MPI parallel matrix, all other matrices are completely local to the computational cores. All restrictions and prolongations are performed using PETSc *VecScatter* and *VecGather* operations. More details on the implementation of the linear BDDC preconditioner can be found in [7], where a parallel implementation of an nonlinear inexact BDDC method is applied to hyperelasticity and elasto-plasticity problems.

2.3 The GM (Global Matrix) Interpolation

Good constants $\tilde{c},\widetilde{C}$ in equation (11) are important for fast convergence. It is well known, that for scalability of multigrid methods the preconditioner should preserve nullspace or near-nullspace vectors of the operator. This is especially important for $\widetilde{K}$. It is a bit less important for the blocks $K_{II}^{(i)}$ in K_{II}, where a large portion of the boundary has Dirichlet data. In this latter case, standard methods can also work well.

Since the AMG method should preserve the nullspace of the operator on all levels, these nullspace vectors have to be in the range of the AMG interpolation. While classical AMG guarantees this property only for constant vectors, the global matrix approach (GM), introduced in [2], allows the user to specify certain near-nullspace vectors, which are interpolated exactly from the coarsest to the finest level; details on the method and its scalability can be found in [10, 2]. Since we are interested in linear elasticity problems, we choose the rotations of the body in $\widetilde{W}$ for the exact interpolation. All translations of the body are already interpolated exactly in classical AMG approaches for systems of PDEs since they use classical interpolation applied component-by-component. We partially assemble the rotations of the subdomains Ω_i in the primal variables. In our implementation, we always use BoomerAMG from the hypre package [5], where a highly scalable implementation of the GM2 approach is integrated; see [2]. We will compare the use of the GM2 approach with a hybrid AMG approach for systems of PDEs. By hybrid AMG approaches, we refer to methods, where the coarsening is based on the physical nodes (nodal coarsening) but the interpolation is based on the unknowns. In general, a nodal coarsening approach is beneficial for the solution of systems of PDEs, and all degrees of freedom belonging to the same physical node are either all coarse or fine on a certain level.

The latter fact is also mandatory for the GM approach. Therefore, GM is based on the same nodal coarsening and can also be considered as a hybrid approach.

3 Numerical Results

As model problems, we choose linear elasticity problems in two and in three dimensions. In two dimensions, we consider a beam $\Omega = [0,8] \times [0,1]$ with a homogeneous Dirichlet boundary condition on the left; see also Fig. 1. A constant volume force is applied in y-direction and the material is chosen to be homogeneous with $E = 210$ and $\nu = 0.3$.

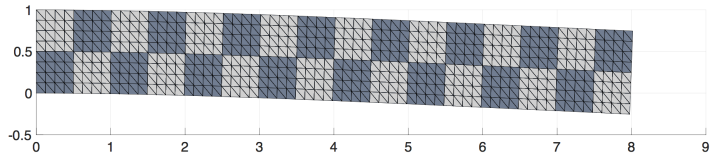

Fig. 1 Beam problem in two dimensions; exemplary decomposition in 32 subdomains depicted.

We first provide a comparison of the preconditioners $M^{-1}_{BDDC,1}$ and $M^{-1}_{BDDC,2}$ using a hybrid AMG approach or the GM2 approach for $\widehat{K}$, respectively; see Fig. 2 for the results. Let us remark that we always use the standard hybrid approach for the approximation of the discrete harmonic extension $\widehat{\mathscr{H}}$ in the case of $M^{-1}_{BDDC,2}$, since this appears to be sufficient so far; also see the remark above on the large Dirichlet boundary. We always use an *HMIS* coarsening, $extended+i$ interpolation, and a threshold of 0.375 for the detection of strong coupling. The interpolation operators of the AMG method are truncated to a maximum of P_{max} entries per row, to keep the operator complexity low and to obtain sufficient weak scalability. We always choose P_{max} such that the operator complexity of the hybrid approach and GM2 approach are similar, to provide a fair comparison. We always use preconditioned GMRES with a relative stopping criteria of 10^{-8}.

In Fig. 2, we present results for the two dimensional beam which is decomposed into 512 subdomains. We increase the problem size by increasing the subdomain size. As primal constraints, we only consider subdomain vertices. We use piecewise quadratic finite elements and thus, the smallest problem carries 882 and the largest problem 136K degrees of freedom per subdomain. We always use one MPI rank per subdomain but use two MPI ranks for each core of the JUQUEEN BlueGene/Q at Forschungszentrum Jülich, Germany, to make use of the hardware threads. Therefore, we have 500 MB of memory available for each subdomain. Using direct solvers for the discrete harmonic extension (i.e., $M^{-1}_{BDDC,1}$), we always have slightly lower GMRES iteration counts and faster runtimes compared to $M^{-1}_{BDDC,2}$, but $M^{-1}_{BDDC,2}$ is more memory efficient. The largest problem, which can be solved with $M^{-1}_{BDDC,1}$

carries 81K d.o.f. per subdomain ($H/h = 100$), while $M^{-1}_{BDDC,2}$ can handle problems twice as large, with 136K d.o.f. per subdomain ($H/h = 130$).

As expected, BDDC using the GM2 approach clearly outperforms the hybrid approach. While the iteration count grows with H/h for the hybrid approach, it stays nearly constant for the GM2 approach. For the problem with $H/h = 120$, $M^{-1}_{BDDC,2}$ with GM2 is six times faster than $M^{-1}_{BDDC,2}$ combined with the hybrid approach, and for $H/h = 130$, $M^{-1}_{BDDC,2}$ with the hybrid approach does not fit in the memory. Choosing $P_{max} = 2$ solves this problem, but the number of iterations is even higher.

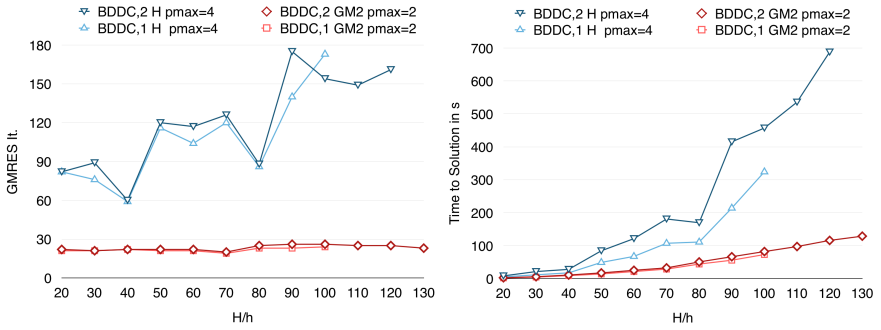

Fig. 2 Comparison for growing H/h and 512 subdomains of the different preconditioners $M^{-1}_{BDDC,1}$ using direct solvers (UMFPACK) for the discrete harmonic extension and $M^{-1}_{BDDC,2}$ using an inexact discrete harmonic extension. Both variants are equipped with hybrid AMG (marked with an H) or GM2, respectively. *pmax* denotes the truncation of the interpolation matrices. **Left:** GMRES iterations. **Right:** Time to solution. Computation performed on JUQUEEN BlueGene/Q at FZ Jülich, Gemany.

We also present a weak scaling study for the best performing combination of $M^{-1}_{BDDC,2}$ and the GM2 approach using $H/h = 80$ and $H/h = 100$; see Fig. 3. While a radical truncation of $P_{max} = 2$ works fine for up to 8 192 subdomains, $P_{max} = 4$ is necessary for the larger configurations. All in all, the parallel efficiency of 91% on 131K MPI ranks and 65K cores and a total problem size of 10 billion degrees of freedom is satisfying.

Finally, we present a weak scaling study in three dimensions. We again consider a linear elastic material and deform a heterogeneous cube. We have a single spherical stiff inclusion ($E = 21\,000$, $nu = 0.3$) in each subdomain. The remaining material is softer with $E = 210$, $nu = 0.3$. This time, we choose piecewise linear finite elements, $H/h = 20$, and, as primal constraints, we enforce continuity in all subdomain vertices and in the midpoints of all edges. We use the same AMG settings as before. In Fig. 4, we again observe a sufficient weak scaling behavior using $M^{-1}_{BDDC,2}$ with the GM2 approach, while the hybrid approach cannot deliver satisfying convergence behavior, since it cannot fulfill (11) with good bounds.

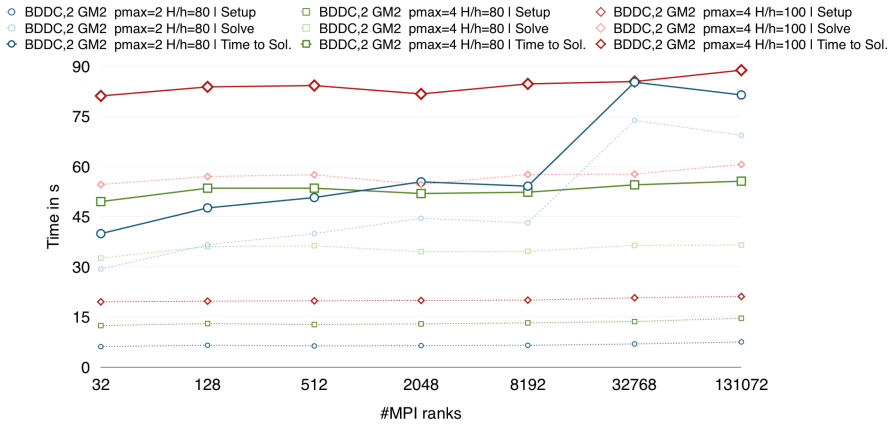

Fig. 3 Weak scalability for $H/h = 80$ and $H/h = 100$ and different truncations *pmax*. *Setup* denotes the BDDC setup time, including all AMG setup times and *Solve* the time spent in the GMRES iteration. Computation performed on JUQUEEN BlueGene/Q at FZ Jülich, Gemany.

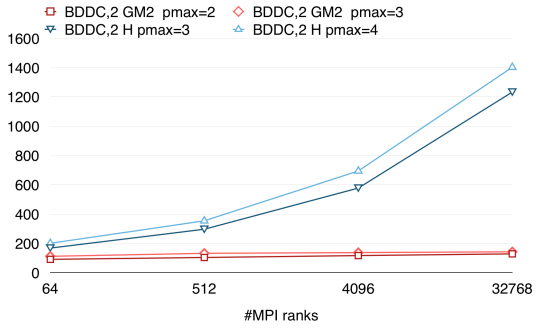

Fig. 4 Heterogeneous and linear elastic material in three dimensions; $H/h = 20$. See Fig. 2 for the remaining notation. Good scalability is achieved using the GM2 interpolation. Computation performed on JUQUEEN BlueGene/Q at FZ Jülich, Gemany.

4 Conclusion

We have shown that a classical AMG approach based on nodal coarsening for systems of PDEs is not sufficient as a preconditioner of the partially coupled matrix in the inexact BDDC approach introduced in Li and Widlund [9], since, for elasticity, it does not fulfill (11) with good bounds. This can be resolved using the GM2 approach, which preserves the nullspace of the partially assembled stiffness matrix in the inexact BDDC method [9]. Our results show that the inexact BDDC approach from [9] using a classical AMG preconditioner with GM2 interpolation is highly parallel scalable and memory efficient.

Acknowledgements This work was supported in part by the German Research Foundation (DFG) through the Priority Programme 1648 "Software for Exascale Computing" (**SPPEXA**) under grants KL 2094/4-1, KL 2094/4-2, RH 122/2-1, and RH 122/3-2. The authors also gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for providing computing time on the GCS Supercomputer SuperMUC at Leibniz Supercomputing Centre (LRZ, www.lrz.de) and **JUQUEEN** [11] at Jülich Supercomputing Centre (JSC, www.fz-juelich.de/ias/jsc). GCS is the alliance of the three national supercomputing centres HLRS (Universität Stuttgart), JSC (Forschungszentrum Jülich), and LRZ (Bayerische Akademie der Wissenschaften), funded by the German Federal Ministry of Education and Research (BMBF) and the German State Ministries for Research of Baden-Württemberg (MWK), Bayern (StMWFK) and Nordrhein-Westfalen (MIWF).

References

1. Santiago Badia, Alberto F. Martn, and Javier Principe. On the scalability of inexact balancing domain decomposition by constraints with overlapped coarse/fine corrections. Parallel Computing, 50:1 – 24, 2015.
2. Allison H. Baker, Axel Klawonn, Tzanio Kolev, Martin Lanser, Oliver Rheinbach, and Ulrike M. Yang. Scalability of classical algebraic multigrid for elasticity to half a million parallel tasks. In Hans-Joachim Bungartz, Philipp Neumann, and Wolfgang E. Nagel, editors, Software for Exascale Computing - SPPEXA 2013-2015, pages 113–140, Cham, 2016. Springer International Publishing.
3. Satish Balay, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, Modern Software Tools in Scientific Computing, pages 163–202. Birkhauser Press, 1997.
4. Clark R. Dohrmann. An approximate BDDC preconditioner. Numer. Linear Algebra Appl., 14(2):149–168, 2007.
5. Van E. Henson and Ulrike M. Yang. Boomeramg: A parallel algebraic multigrid solver and preconditioner. Applied Numerical Mathematics, 41:155–177, 2002.
6. Axel Klawonn, Martin Lanser, and Oliver Rheinbach. Toward extremely scalable nonlinear domain decomposition methods for elliptic partial differential equations. SIAM Journal on Scientific Computing, 37(6):C667–C696, 2015.
7. Axel Klawonn, Martin Lanser, and Oliver Rheinbach. Nonlinear BDDC methods with inexact solvers. Technical report, 2017. In preparation.
8. Axel Klawonn and Oliver Rheinbach. Inexact FETI-DP methods. Internat. J. Numer. Methods Engrg., 69(2):284–307, 2007.
9. Jing Li and Olof B. Widlund. On the use of inexact subdomain solvers for BDDC algorithms. Comput. Meth. Appl. Mech. Engrg., 196:1415–1428, 2007.
10. John Ruge and Klaus Stüben. Efficient solution of finite difference and finite element equations by algebraic multigrid (AMG). In J.D. Paddon and H. Holstein, editors, The Institute of Mathematics and its Applications Conference Series, volume 3, pages 169–212. Clarenden Press, Oxford, 1985.
11. Michael Stephan and Jutta Docter. JUQUEEN: IBM Blue Gene/Q Supercomputer System at the Jülich Supercomputing Centre. Journal of large-scale research facilities, 1:A1, 2015.
12. Andrea Toselli and Olof Widlund. Domain Decomposition Methods - Algorithms and Theory, volume 34 of Springer Series in Computational Mathematics. Springer, 2004.