Introduction

We present numerical results that exploit recent algorithmic advances in Balancing Domain Decomposition by Constraints (BDDC) (Dohrmann 2003) preconditioners for Symmetric Positive Definite (SPD) linear systems arising in the contexts of modelling reservoirs and electromagnetics by means of finite element discretizations of elliptic PDEs with highly heterogeneous coefficients.

BDDC methods have proven to be powerful preconditioners for SPD systems, having typical poly-logarithmic condition number bounds of the type (Mandel *et. al.* 2005)

$$C\left(1+\log\left(H/h\right)\right)^2$$

with h the size of the elements, H the maximum diameter of the subdomains, and C a constant independent of the number of subdomains in which the domain has been decomposed, and possibly independent of smooth variations of the coefficients of the PDE (Klawonn *et. al.* 2008, Pechstein and Scheichl 2008). Besides the interesting theoretical aspects, the BDDC algorithm has also proven to be very efficient on distributed memory architectures with tens or even hundreds of thousands of cores, as shown in different studies (Klawonn and Rheinbach 2008, Badia *et. al.* 2014, Zampini 2014).

BDDC methods have been extensively studied and adapted to many different discretizations and PDEs. In the current context it is worth mentioning the works on div-conforming Raviart-Thomas vector fields (Oh *et. al.* 2013, Tu 2005), curl-conforming Nédélec elements (Dohrmann and Widlund 2015), and the mixed-hybrid formulation (Boffi *et. al.* 2013) of porous media flow (Tu 2007). A recent extension of the BDDC methods to the mixed-hybrid formulation of porous media flows in the presence of fractures can be found in (Sistek *et. al.* 2015).

Recently, BDDC methods have been extended with algebraic and analytical techniques in order to tackle SPD systems with high-contrast distributions in the coefficients of the underlying PDE (Mandel *et. al.* 2012, Kim and Chung 2015, Klawonn *et. al.* 2015, Pechstein and Dohrmann 2013). The high arithmetic intensity and the ability to strongly control the number of iterations of outer Krylov iterative accelerators makes them very attractive for future exascale machines, where local flops are cheaper than data movement, and where the number of global synchronization steps needed by the solver should be minimized (Dongarra *et. al.* 2011).

Methods

The BDDC methods belong to the family of non-overlapping Domain Decomposition (DD) preconditioners (Toselli and Widlund, 2005) and they can be viewed as an evolution of the Balancing Neumann-Neumann methods (Mandel 1993), where the local and coarse problems defined by the method are additively, instead of multiplicatively, combined. As usual in DD methods, local problems are solved independently on each subdomain whereas the coarse solver is solved in parallel among the subdomains and it provides the global exchange of information that is mandatory to obtain scalable solvers in terms of number of iterations.

The core of the BDDC algorithm resides in the construction of scalable and cheap preconditioners for the interface Schur complement, which is defined by statically condensing the degrees of freedom belonging to the interior of the subdomains (Toselli and Widlund, 2005) The recipe for the characterization of the methods relies on the selection of the so-called *primal continuity constraints* and on the choice of an averaging/scaling procedure. The interface preconditioner is defined by the algorithm on a partially assembled space, where the continuity of primal degrees of freedom is enforced by means of the solution of the partially assembled Schur complement, which naturally defines the local corrections and the coarse problem of the algorithm (Li and Widlund, 2006). The averaging procedure is thereafter needed to restore the continuity of the remaining *dual* degrees of freedom during Krylov iterations.

In the current work, BDDC methods for SPD systems with heterogeneous coefficients are characterized by the algebraic strategy recently proposed by Pechstein and Dohrmann (2013), which consists in the combination of the so-called *deluxe scaling* (Widlund and Dohrmann 2014) with an adaptive selection of primal constrains achieved by means of the solutions of small and dense generalized eigenvalue problems defined for each open subset of the interface (either an edge or a face); only those eigenvectors having the corresponding eigenvalue larger than a given user-specified threshold are selected and included into the primal space. Lightweight nearest neighbor communications are needed in order to assemble the eigenvalue problems, which can then be efficiently and independently solved on each subdomain. For additional details on the implementation of the method considered in the current work, see ref. (Zampini 2015). In what follows, we will refer to the adopted strategy as *adaptive BDDC*.

Finally, when a large number of subdomains are considered, the direct solution of the BDDC coarse problem by means of a Cholesky parallel solver is no longer feasible in the presence of large contrasts in the coefficients, with the size of the coarse problem being equal to the number of primal constraints imposed. A remedy consists in considering the multilevel extension of the BDDC algorithm, where the subdomains at the finest level are regarded as elements at a coarser level and aggregated into coarse subdomains; the parallel Cholesky solver can then be replaced by one application of a BDDC preconditioner defined on the coarse discretization (Mandel *et. al.* 2008). Also, using adaptive selection of constraints at the coarser level provides high-quality spectrally equivalent BDDC solvers, which do not compromise the global convergence properties of the method at the finest level. The computation of the local and coarse corrections involved in the coarser levels of the BDDC preconditioner can be overlapped and independently solved on different MPI communicators (Zampini 2015).

Numerical results

The distributed memory results presented in this work have been obtained using the C++ finite element library DOLFIN (Logg and Wells 2012) with the PETSc library as a backend (Balay *et. al.* 2014); the implementation of the BDDC preconditioner has been contributed to PETSc by the corresponding author (Zampini 2015). SPD linear systems are always solved using the Preconditioned Conjugate Gradient (PCG) method using random right-hand sides, null initial guesses, relative residual reduction 10^{-8} as a stopping criterion, and the BDDC as a preconditioner. The domain considered is the unit cube discretized with tetrahedral meshes, and decomposed using mesh partitioners; a one-to-one mapping between MPI processes and subdomains is always considered. Numerical results have been obtained on the Cray XC40 PizDora of the Swiss National Supercomputing Centre CSCS, equipped with two 12-core Intel Haswell CPUs clocked at 2.6 Ghz and with 64 GB of DRAM (http://www.cscs.ch), and on the Cray XC40 Shaheen at KAUST, which features dual 16-core Haswells clocked at 2.3 GHz and 128 GB of DRAM per node (http://hpc.kaust.edu.sa). Executions up to 18432 cores are presented in this extended abstract. Higher core counts will be available in the full paper.

It is well established that Raviart Thomas vector fields can be used for finite element based simulations in reservoir modeling, since they naturally conserve the mass, being div-conforming, as an alternative to finite volume methods (see e.g. Boffi *et. al.* 2013). A commonly used variational problem for assessing the quality of a Raviart-Thomas solver (see e.g. Kolev and Vassilevsky 2012) reads as follows: given $\alpha, \beta \in L^\infty(\Omega), \alpha > 0, \beta \geq 0, \mathbf{f} \in (L^2(\Omega))^3$, find $\mathbf{u} \in H_0(\text{div}, \Omega)$ s.t.

$$\int_\Omega (\alpha \, \text{div}\mathbf{u} \, \text{div}\mathbf{v} + \beta \mathbf{u} \cdot \mathbf{v}) dx = \int_\Omega \mathbf{f} \cdot \mathbf{v} dx, \; \forall \mathbf{v} \in H_0(\text{div}, \Omega)$$

where $H_0(\text{div}, \Omega)$ is the proper subspace of $H(\text{div}, \Omega)$ consisting in functions with vanishing normal component on $\partial\Omega$. Figure 1 contains the results of a weak scaling test for the PCG solver for an element-wise random distribution of the coefficients in the range $[10^{-3}, 10^{3}]$ (thus 6 orders of magnitude in the contrast) for both material parameters. Results compare the setup and solving times for adaptive BDDC using the parallel Cholesky solver (labeled by direct MUMPS in the legend) from the MUMPS library (Amestoy *et. al.* 2001) as a coarse solver, and the adaptive multilevel BDDC considering an extra level with different aggregation strategies (CR labels in the legend). As an example, CR equal to 48 means that 48 subdomains at the finest level are aggregated into one coarse subdomain. Such aggregation is accomplished by partitioning the coarser mesh using the connectivity graph of the finer subdomains. The threshold used for selecting the primal constraints at the finest level is equal to 10 for both the adaptive BDDC and the adaptive multilevel BDDC; in addition, a threshold value equal to 5 has been used to select the primal constraints at the coarser level for the adaptive multilevel BDDC. The size of the subdomains is kept fixed (around 87K degrees of freedom) and the number of subdomains is increased from 1536 to 18432, with the size of the finest problem increasing proportionally from 135M to 1.6B degrees of freedom. The number of iterations required by PCG to reach convergence varies between 20 and 24 for any of the test considered (data not shown), proving the robustness of the methods considered for three-dimensional div-conforming discretizations with high-contrast coefficients; however, setup and solving times differ dramatically between adaptive BDDC and adaptive multilevel BDDC when using more than 3K cores. Indeed, coarse problem factorization and solution phases respectively dominate the setup and solve phases of the adaptive BDDC method. On the other hand, the adaptive multilevel BDDC exhibits a very mild dependence on the number of subdomains in both phases. It must be noted that the parallel Cholesky solver failed to compute the numerical factorization of the coarse problem with 18432 cores by reason of exceeding memory capacity.

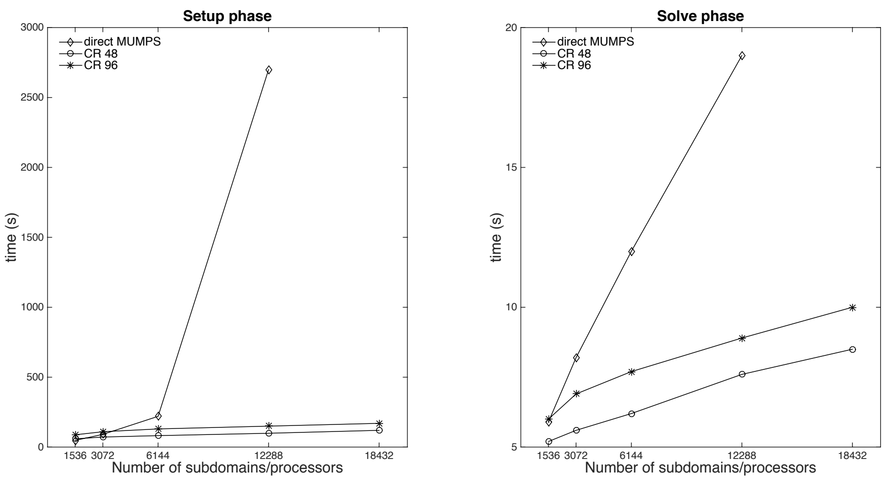

Figure 1 *Setup and solving times for the adaptive multilevel BDDC (CR48 and CR96) are compared with the adaptive BDDC using Cholesky parallel coarse solver (direct MUMPS); three-dimensional div-conforming finite elements discretizations with element-wise random material parameters, with 6 orders of magnitude in the contrast.*

The Eddy's formulation of Maxwell's equations governs the behavior of electro-magnetic geophysical survey data in the low frequency regime (Carcione 2010, Haber 2015). The variational formulation of the problem reads as follows: given $\alpha, \beta \in L^{\infty}(\Omega), \alpha \geq 0, \beta > 0, \mathbf{f} \in (L^2(\Omega))^3$, find $\mathbf{u} \in H_0(\mathrm{curl}, \Omega)$

$$\int_{\Omega} (\alpha \nabla \times \mathbf{u} \cdot \nabla \times \mathbf{v} + \beta \mathbf{u} \cdot \mathbf{v}) dx = \int_{\Omega} \mathbf{f} \cdot \mathbf{v} dx, \forall \mathbf{v} \in H_0(\mathrm{curl}, \Omega)$$

where $H_0(\text{curl},\Omega)$ is the proper subspace of $H(\text{curl},\Omega)$ consisting in functions with vanishing tangential trace on $\partial\Omega$ (see e.g. Toselli and Widlund 2005); **u** is the intensity of the electric field, $1/\alpha$ is the magnetic permeability, and β the electrical conductivity. Curl-conforming finite elements such as the lowest order Nédélec elements are generally used to discretize the problem. As a test case for the adaptive BDDC, we consider the decomposition in 40 subdomains of a 32x32x32 hexahedral grid, with each hexahedron further decomposed into 6 tetrahedra, for a total of 240K degrees of freedom; a special change of basis is needed by the BDDC to accommodate the Nédélec elements on the subdomain edges (Dohrmann and Widlund 2015). The coefficients of the PDE are element-wise randomly chosen in the interval $[10^{-p}, 10^{p}]$ and the threshold value for selecting the primal constraints is equal to 10. Figure 2 shows the iteration counts and the total number of constraints (in parenthesis) selected by the adaptive procedure for different orders of magnitude in the contrast of the coefficients (i.e., 2p). The contrast in α is increased from left to right, whereas the contrast in β from top to bottom. Iteration counts are independent of the contrast in each test case, providing experimental evidence of the robustness of the adaptive BDDC method for three-dimensional curl-conforming finite element discretizations. Moreover, it is worth noticing that the increase in the number of primal constraints is less pronounced when α is constant (first column). This is an important consideration in the context of electromagnetic geophysical inversion, where the magnetic permeability may be assumed to be constant in the absence of magnetic sources.

$\alpha \rightarrow$ $\beta \downarrow$	0	2	4	6	8
0	15 (627)	23 (627)	30 (984)	31 (1459)	32 (1718)
2	15 (627)	23 (627)	30 (983)	31 (1457)	32 (1715)
4	15 (627)	23 (627)	30 (976)	31 (1446)	31 (1705)
6	16 (627)	23 (629)	28 (985)	32 (1400)	35 (1669)
8	22 (644)	27 (697)	29 (1066)	32 (1447)	35 (1708)

Figure 2 *Number of iterations and size of the coarse problem for adaptive BDDC for three-dimensional curl-conforming finite elements discretizations using element-wise random coefficients. Orders of magnitude in the contrast for α are provided in the first row, for β in the first column.*

Conclusions

Recent advances in BDDC methods have contributed to the robustness of the methods with respect to highly heterogeneous coefficients distributions of elliptic PDEs, providing algebraic black-box techniques that can improve upon request the quality of the preconditioners.

The experimental results provided in the present study proved the robustness of the adaptive BDDC method for SPD systems arising in three-dimensional div- and curl- conforming finite element discretizations with highly heterogeneous distributions of the coefficients of the underlying PDEs. Large parallel numerical results until 1.6B degrees of freedom on 18K cores have shown the efficacy of the adaptive multilevel BDDC method for Raviart-Thomas vector fields. Future work may explore the usage of hardware accelerators for speeding up the setup phase of the adaptive BDDC method.

The application of the adaptive BDDC algorithm on a mixed formulation of porous media flows with heterogeneous distributions of coefficients is currently under study, together with the extension to the mixed-hybrid formulation. Future investigations will also consider the comparison of the adaptive BDDC method with the existing high performance implementations of the auxiliary space multigrid solvers for div- (Kolev and Vassilevsky 2012) and curl- (Kolev and Vassilevsky 2009) conforming elements using real-world distributions of the coefficients in real applications.

Acknowledgements

The authors wish to thanks the Swiss National Supercomputing Centre CSCS and in particular its director Thomas Schulthess for having given them the opportunity to run the experimental tests provided in the current work.

References

Amestoy P. R., Duff I. S., Koster J. and L'Excellent J. Y. [2001] A fully asynchronous multi- frontal solver using distributed dynamic scheduling, *SIAM Journal of Matrix Analysis with Applications* **23**, 15-41.

Badia S., Martin A. F. and Principe J. [2014] A highly scalable parallel implementation of balancing domain decomposition by constraints, *SIAM Journal of Scientific Computing* **36**, C190-C218.

Balay S., Abhyankar S., Adams M. F., Brown J., Brune P., Buschelman K., Eijkhout V., Gropp W. D., Kaushik D., Knepley M. G., Curfmann McInnes L., Rupp K., Smith B. F., Zhang H. [2014] PETSc Web page, http://www.mcs.anl.gov/petsc.

Boffi D., Brezzi F. and Fortin M. [2013] Mixed finite element methods and applications, *Springer Series in Computational Mathematics* **44**.

Carcione J. M. [2010] Simulation of electromagnetic diffusion in anisotropic media, *Progress in Electromagnetics Research B.* **26**, 425-450.

Dohrmann, C. R. [2003] A preconditioner for substructuring based on constrained energy minimization. *SIAM Journal of Scientific Computing* **25,** 246-258.

Dohrmann C. R. and Widlund O. B. [2015] A BDDC algorithm with deluxe scaling for three-dimensional H(curl) problems, *appeared electronically in April 2015 in Communications in Pure and Applied Mathematics*.

Dongarra J. *et. al.* [2011] The international exascale software project roadmap, *International Journal of High Performance Computing and Applications* **6**, 1-58.

Haber E. [2015] Computational methods in geophysical electromagnetics, *SIAM*.

Kim H. H. and Chung E. T. [2015] A BDDC algorithm with optimally enriched coarse spaces for two- dimensional elliptic problems with oscillatory and high contrast coefficients, *Submitted.*

Klawonn A., Rheinbach O. and Widlund O. B. [2008] An analysis of a FETI-DP algorithm on irregular subdomains in the plane, *SIAM Journal of Numerical Analysis* **46**, 2484-2504.

Klawonn A. and Rheinbach O. [2010]. Highly scalable parallel domain decomposition methods with an application to biomechanics, *ZAMM Journal of Applied Mathematics and Mechanics* **90**, 5-32.

Klawonn A., Radtke P. and Rheinbach O. [2015] FETI-DP methods with an adaptive coarse space. *SIAM Journal of Numerical Analysis* **53**, 297-320.

Kolev T. V., Vassilevsky P. S. [2009] Parallel auxiliary space AMG solver for H(curl) problems, *Journal of Computational Mathematics* **27**, 604-632.

Kolev T. V., Vassilevsky P. S. [2012] Parallel auxiliary space AMG solver for H(div) problems, *SIAM Journal of Scientific Computing* **34**.

Logg A. and Wells G. N. [2012] DOLFIN: automated finite element computing, *ACM Transactions on Mathematical Software* **37**.

Li J. and Widlund O. B [2006] FETI-DP, BDDC, and block Cholesky methods, *International Journal of Numerical Methods in Engineering*, **66**, 250-271.

Mandel J. [1993] Balancing domain decomposition, *Communications in Applied Numerical Methods* **9**, 233-241.

Mandel J., Dohrmann C. R. and Tezaur R. [2005] An algebraic theory for primal and dual substructuring methods by constraints, *Applied Numerical Mathematics* **54**, 167-193.

Mandel J., Sousedik B. and Dohrmann C. R. [2008] Multispace and multilevel BDDC, *Computing* **3**, 55-85

Mandel J., Sousedik B. and Sistek J. [2012] Adaptive BDDC in three dimensions, *Mathematics and Computers in Simulation* **82**, 1812-1831.

Oh D. -S., Widlund O. B. and Dohrmann C. R. [2013] A BDDC algorithm for Raviart–Thomas vector fields, *Courant Institute Technical Report* TR2013-951.

Pechstein C. and Dohrmann C. R. [2013] Modern domain decomposition methods BDDC, deluxe scaling, and an algebraic approach. Seminar talk, Linz, http://people.ricam.oeaw.ac.at/c.pechstein/pechstein-bddc2013.pdf

Pechstein C. and Scheichl R. [2008] Analysis of FETI methods for multiscale PDEs, *Numerische Mathematik* **111**, 293-333.

Sistek J, Brezina J. and Sousedik B. [2015] BDDC for mixed-hybrid formulation of flow in porous media with combined mesh dimensions, http://arxiv.org/abs/1504.07085.

Toselli A. and Widlund O. B. [2005] Domain decomposition methods - Algorithms and theory, *Springer Verlag.*

Tu X. [2005] A BDDC algorithm for a mixed formulation of flow in porous media, *Electronic Transactions on Numerical Analysis* **20**, 164-179

Tu X. [2007] A BDDC algorithm for flow in porous media with a hybrid finite element discretization, *Electronic Transactions on Numerical Analysis* **26**, 146-160.

Widlund O.B. and Dohrmann C. R. [2014] BDDC deluxe Domain Decomposition, *proceedings of the XXII International conference on Domain Decomposition Methods in Science and Engineering*.

Zampini S. [2014] Dual-primal methods for the cardiac bidomain model. *Mathematical Models and Methods in Applied Sciences* **24**, 667-696.

Zampini S. [2015] PCBDDC: a class of robust dual-primal methods in PETSc, *Submitted*.