

Adaptive BDDC Preconditioning for Unstructured Bidomain Discretizations

Talaat Abdelhamid^[0000-0002-7282-5437],
Ngoc Mai Monica Huynh^[0000-0003-3639-5844],
Stefano Zampini^[0000-0002-0435-0433],
Rongliang Chen^[0000-0002-5253-0412],
Luca F. Pavarino^[0000-0002-3014-4668],
Simone Scacchi^[0000-0001-6011-784X]

1 Introduction

In recent years, there has been increasing interest in cardiac digital twins – personalized computational representations developed to simulate an individual’s heart function [5]. While large-scale modeling and simulation of cardiac function aim to develop clinically applicable tools, many underlying processes—such as electrical excitation, muscle contraction, blood flow, and circulation—remain subjects of ongoing investigation and focus of mathematical and numerical studies. For instance, the mathematical formulations of the electrophysiology of the heart tissue have been extensively studied [3], and many computational approaches for its numerical solution have been introduced over the past decade (see for e.g. [7, 9] and references therein). One of the key challenges in this context is reducing the overall computational time required by the iterative solvers used for the Bidomain equations. The effectiveness of Balancing Domain Decomposition by Constraints (BDDC) preconditioners for finite element formulations of such parabolic partial differential equations (PDEs)

Talaat Abdelhamid
Department of Mathematics, University of Milan, Italy & Physics and Mathematical Engineering
Department, Faculty of Electronic Engineering, Menoufiya University, Egypt, e-mail:
talaat.abdelhamid@unimi.it

Ngoc Mai Monica Huynh · Simone Scacchi
Department of Mathematics, University of Milan, Italy, e-mail: ngoc.huynh@unimi.it,
simone.scacchi@unimi.it

Stefano Zampini
King Abdullah University of Science and Technology, Saudi Arabia, e-mail:
stefano.zampini@kaust.edu.sa

Rongliang Chen
Shenzhen Institutes of Advanced Technology, China, e-mail: rl.chen@siat.ac.cn

Luca F. Pavarino
Department of Mathematics, University of Pavia, Italy, e-mail: luca.pavarino@unipv.it

has been well established [11, 12, 13]. Typically, the condition number κ_2 of the preconditioned system satisfies

$$\kappa_2 \leq C(1 + \log(H/h))^2,$$

where h denotes the mesh element size, H represents the characteristic subdomain diameter, and C is a constant independent of h , H , and the total number of subdomains. The primary contribution of this study is the formulation and computational evaluation of adaptive BDDC (aBDDC) preconditioners for unstructured finite element discretizations of the Bidomain equations on patient-specific left ventricular geometries. Adaptive BDDC have been shown to be more computationally demanding in the setup phase, but the resulting number of linear iterations can be lower than in standard BDDC if a proper primal space can be adaptively constructed. This means that the number of collective communications and data transfers are limited, yielding computationally efficient simulations involving unstructured and large computational meshes [14]. Numerical experiments carried out on a Linux-based high-performance computing cluster indicate that the proposed aBDDC strategy shows improved scalability and efficiency compared to standard BDDC methods.

2 Mathematical Model

This study focuses on the macroscopic Bidomain formulation for cardiac electrophysiology, expressed in a parabolic–elliptic (PE) form [3]. The computational domain $\Omega \subset \mathbb{R}^3$ denotes the open, connected, and bounded region corresponding to the myocardium. The cardiac tissue is represented as the superposition of two overlapping anisotropic conductive media—the intracellular and extracellular spaces—which coexist at each spatial location and are separated by a distributed cellular membrane. For a given spatial domain Ω and temporal interval $(0, T)$, the objective is to determine the intracellular and extracellular electric potentials, $u_i(x, t)$ and $u_e(x, t)$, respectively, as well as the transmembrane potential $v = u_i - u_e$. The Bidomain equations are complemented by a set of Ordinary Differential Equations (ODEs), commonly referred to as the membrane model, which governs the ionic exchanges between the intra- and extracellular domains. The variables of this subsystem include the gating variables $\mathbf{w} : \Omega \times (0, T) \rightarrow \mathbb{R}^{N_w}$ and the ionic concentrations $\mathbf{c} : \Omega \times (0, T) \rightarrow \mathbb{R}^{N_c}$. Together, these components form the complete coupled system, expressed as follows:

$$\begin{cases} \chi C_m \frac{\partial v}{\partial t} - \operatorname{div}(D_i \nabla v) - \operatorname{div}(D_i \nabla u_e) + \chi I_{ion}(v, \mathbf{w}, \mathbf{c}) = I_{app}^i & \text{in } \Omega \times (0, T), \\ -\operatorname{div}(D_i \nabla v) - \operatorname{div}((D_i + D_e) \nabla u_e) = I_{app}^i + I_{app}^e & \text{in } \Omega \times (0, T), \\ \frac{\partial \mathbf{w}}{\partial t} - \mathbf{R}(v, \mathbf{w}) = 0 & \frac{\partial \mathbf{c}}{\partial t} - \mathbf{C}(v, \mathbf{w}, \mathbf{c}) = 0 & \text{in } \Omega \times (0, T), \\ \mathbf{n}^T D_i \nabla (v + u_e) = 0 & \mathbf{n}^T D_e \nabla u_e = 0 & \text{on } \partial\Omega \times (0, T), \\ v(x, 0) = v_0(x), \quad \mathbf{w}(x, 0) = \mathbf{w}_0(x), \quad c(x, 0) = c_0(x) & \text{in } \Omega \end{cases}$$

Within this framework, I_{ion} represents the ionic current as defined by the selected electrophysiological model. In the present work, we employ the ten Tusscher-Panfilov 2006 (TP06) model [10]. The functions \mathbf{R} and \mathbf{C} characterize the time-dependent behavior of the gating variables \mathbf{w} and the ionic concentration variables \mathbf{c} , respectively. C_m denotes the membrane capacitance, $C_m = 1 \mu F/cm^3$, while χ is the ratio of membrane area per tissue volume. The applied currents in the intracellular and extracellular spaces, denoted by I_{app}^i and I_{app}^e , must fulfill the compatibility constraint $\int_{\Omega} (I_{app}^i + I_{app}^e) dx = 0$. For convenience, we assume that the externally applied extracellular current is the negative of its intracellular counterpart, i.e., $I_{app}^e = -I_{app}^i$. The directional dependence of cardiac tissue conductivity, arising from the organized orientation of myocardial fibers, is captured by the conductivity tensors $D_{i,e}(x) = \sigma_l^{i,e} \mathbf{a}_l(x) \mathbf{a}_l^T(x) + \sigma_t^{i,e} \mathbf{a}_t(x) \mathbf{a}_t^T(x) + \sigma_n^{i,e} \mathbf{a}_n(x) \mathbf{a}_n^T(x)$. Here $\mathbf{a}_l, \mathbf{a}_t, \mathbf{a}_n \in [L^\infty(\Omega)]^3$ represent a local orthonormal triplet, where \mathbf{a}_l aligns with the cardiac fiber orientation, and \mathbf{a}_t and \mathbf{a}_n are tangent and normal to the myocardial sheets, respectively, both orthogonal to \mathbf{a}_l . Lastly, $\sigma_{l,t,n}^{i,e}$ are the corresponding conductivities in the intra- and extracellular spaces along these principal directions.

Space and time discretization. By considering a quasi-uniform tessellation of Ω , denoted by \mathcal{T}_h , with a maximum diameter h and by denoting $V^h(\Omega) \subset V$ the corresponding conforming finite element space, with the same basis functions $\{\varphi_p\}_{p=1}^{N_h}$ for all variables $u_{i,e}$, w and c and let $A_{i,e}$ and M be the usual stiffness and mass matrices. Therefore, we can express the discrete-in-space PE formulation of the Bidomain problem in the following form:

$$\begin{cases} \chi C_m M \dot{\mathbf{v}} + A_i \mathbf{v} + A_e \mathbf{u}_e + \chi M \mathbf{I}_{ion}(\mathbf{v}, \mathbf{w}, \mathbf{c}) = M \mathbf{I}_{app}^i \\ A_i \mathbf{v} + (A_i + A_e) \mathbf{u}_e = 0 \\ \dot{\mathbf{w}} = \mathbf{R}(\mathbf{v}, \mathbf{w}), \quad \dot{\mathbf{c}} = \mathbf{C}(\mathbf{v}, \mathbf{w}, \mathbf{c}), \quad \mathbf{v} = \mathbf{u}_i - \mathbf{u}_e \end{cases} \quad (1)$$

where the bold font refers to a vector representation of the highlighted quantity and $\dot{\mathbf{x}}$ its time derivative. We assume interpolated ionic currents, since I_{ion} is nonlinear in v . As for time discretization, we consider a standard discrete-in-time approximation of the time derivative $\dot{v} \approx (1/\Delta t)(v^k - v^{k-1})$, where the time interval $(0, T)$ has been divided into K sub-intervals with length $\Delta t = t^k - t^{k-1}$. The equations in system (1) are discretized in time with an IMEX (IMplicit-EXplicit) scheme, where the diffusion terms are treated implicitly and the reaction explicitly. We employ here a decoupling strategy between the PDEs and the ODEs systems: this can be motivated by the different spatial and time scales involved, as well as the strength of the coupling between the models. This proposed decoupling is shown in Alg. 1.

Adaptive BDDC preconditioner. The adaptive BDDC (aBDDC) strategy enhances the standard BDDC framework [4] by enriching the primal space with additional constraints derived from localized generalized eigenvalue problems. In standard BDDC methods, the coarse space is usually spanned by vertex constraints and edge/face averages. However, such choice may not be sufficient to ensure good convergence for complex problems, while being computational demanding for problems with a high number of degrees of freedom. On the contrary, the choice of spaces spanned by specific local eigenvectors for aBDDC methods can lead to a decrease

Algorithm 1 Decoupling IMEX scheme for the PE Bidomain model.

```

1: Notation: given any variable  $x$ , we use the short notation  $x^k = x(t^k)$ .
2: for  $k = 1, \dots, K$  (time loop) do
3:   Given  $\tilde{\mathbf{v}} := \mathbf{v}^{k-1}$ , solve the ODE system:
4:      $\mathbf{w}^k = \mathbf{w}^{k-1} + \Delta t \mathbf{R}(\tilde{\mathbf{v}}, \mathbf{w}^{k-1}) \rightarrow$  obtain  $\mathbf{w}^k$ 
5:      $\mathbf{c}^k = \mathbf{c}^{k-1} + \Delta t \mathbf{C}(\tilde{\mathbf{v}}, \mathbf{w}^{k-1}, \mathbf{c}^{k-1}) \rightarrow$  obtain  $\mathbf{c}^k$ 
6:   Given  $\tilde{\mathbf{v}}$ , solve the elliptic PDE:
7:      $(A_i + A_e)\mathbf{u}_e^k = -A_i\tilde{\mathbf{v}} \rightarrow$  obtain  $\mathbf{u}_e^k$ 
8:   Given  $\tilde{\mathbf{u}}_e = \mathbf{u}_e^k$ , solve the parabolic PDE:
9:      $(C_m M + \Delta t A_i)\mathbf{v}^k = \mathbf{f}(\mathbf{v}^{k-1}, \tilde{\mathbf{u}}_e, \tilde{\mathbf{w}}, \tilde{\mathbf{c}}) \rightarrow$  obtain  $\mathbf{v}^k$ 
10:   where  $\mathbf{f}(\tilde{\mathbf{v}}, \tilde{\mathbf{u}}_e, \tilde{\mathbf{w}}, \tilde{\mathbf{c}}) = C_m M \tilde{\mathbf{v}} + \Delta t (M \mathbf{I}_{app}^i - A_e \tilde{\mathbf{u}}_e - M \mathbf{I}_{ion}(\tilde{\mathbf{v}}, \tilde{\mathbf{w}}, \tilde{\mathbf{c}}))$ .
11: end for

```

in the linear iteration counts, which, in turn, reduces both the number of collective communications and the volume of data movement required by the linear solver [14]. In the following, we briefly present a simplified formulation involving equivalence classes containing only two subdomains; see [8, 14] for more extensive explanations. We first need to define the parallel sum of two symmetric positive definite matrices \mathbf{A} and \mathbf{B} , as $\mathbf{A} : \mathbf{B} := (\mathbf{A}^{-1} + \mathbf{B}^{-1})^{-1}$. The generalized eigenvalue problem formulation for the three-dimensional case, considering a generic interface face \mathcal{F} shared by two neighboring subdomains Ω_i and Ω_j , can be expressed as:

$$\tilde{\mathbf{S}}_{\mathcal{F}\mathcal{F}}^{(i)} : \tilde{\mathbf{S}}_{\mathcal{F}\mathcal{F}}^{(j)} \phi = \lambda \mathbf{S}_{\mathcal{F}\mathcal{F}}^{(i)} : \mathbf{S}_{\mathcal{F}\mathcal{F}}^{(j)} \phi. \quad (2)$$

The right-hand side involves the parallel sum of the principal submatrices of the local Schur complements $\mathbf{S}^{(i)}$ and $\mathbf{S}^{(j)}$, both of which are symmetric positive definite. Conversely, the left-hand side corresponds to the parallel sum associated with the minimal-energy extensions of the degrees of freedom defined on the interface \mathcal{F} to the remaining portion of the subdomain boundary Γ_i . Formally, the Schur complement for a subdomain can be expressed as:


$$\mathbf{S}^{(i)} = \begin{pmatrix} \mathbf{S}_{\mathcal{F}^c \mathcal{F}^c}^{(i)} & \mathbf{S}_{\mathcal{F}^c \mathcal{F}}^{(i)} \\ \mathbf{S}_{\mathcal{F}^c \mathcal{F}}^{(i)\top} & \mathbf{S}_{\mathcal{F}\mathcal{F}}^{(i)} \end{pmatrix},$$

where \mathcal{F}^c represents the remaining of the subdomain boundary, excluding \mathcal{F} . This allows us to express: $\tilde{\mathbf{S}}_{\mathcal{F}\mathcal{F}}^{(i)} := \mathbf{S}_{\mathcal{F}\mathcal{F}}^{(i)} - \mathbf{S}_{\mathcal{F}^c \mathcal{F}}^{(i)\top} \mathbf{S}_{\mathcal{F}^c \mathcal{F}^c}^{(i)-1} \mathbf{S}_{\mathcal{F}^c \mathcal{F}}^{(i)}$. The primal space is then built by collecting the k smallest eigenvectors that satisfy Equation (2) and such that $\lambda_k \leq \frac{1}{\text{tol}}$, being λ_k the k eigenvalues smaller than a fixed tolerance (see [15] for other details). This criterion ensures that only the most significant constraints contributing to the global coupling are considered.

3 Numerical Results

The three-dimensional geometry of a left ventricle is reconstructed from a computed tomography (CT) scan of a 67-year-old male volunteer. Five unstructured tetrahedral finite element meshes, with different levels of refinement, are then generated using the VMTK meshing toolkit [6], as summarized in Table (a), which lists the number of elements and vertices for each refinement level. A stimulation current of 350 mA/cm^3 was applied for 1 ms within a sub-epicardial spheroidal region of radius 0.05 cm. Cardiac fibers are generating using the `lifex-fiber` tool [2]. The linear systems arising at each time step are solved iteratively with a conjugate gradient (CG) method preconditioned either with standard BDDC or adaptive BDDC (aBDDC) – more extensive tests and details in [1].

Mesh	# vertices	# elements
Mesh1	104 395	496 905
Mesh2	511 026	2 987 019
Mesh3	1 098 025	5 361 953
Mesh4	2 004 108	12 520 975
Mesh5	4 394 120	21 484 023



(a) Data of the refined meshes of the left ventricular geometry.

(b) Geometry of left ventricle epicardium and endocardium.

Fig. 1 Meshes details (Table (a)) and visual representation of the left ventricle geometry (Fig. (b)) employed in the numerical tests.

Test 1 – varying the adaptive tolerance. In this test, we employ Mesh5, the finest discretization, consisting of 4 394 120 degrees of freedom and 21 484 023 tetrahedral elements. The mesh is partitioned into 8, 16, 32, and 64 subdomains using the METIS graph partitioning library (<https://github.com/KarypisLab/METIS>). Our investigation focuses on how variations in the adaptive tolerance parameter tol ($= 2, 5, 10, 15, 20$) influence several performance indicators: the size of the coarse problem (NP), the spectral condition number of the preconditioned system (κ_2), the number of Conjugate Gradient (CG) iterations (It), and the CPU time (t_{cpu}) required both to construct the preconditioner and to solve the linear system. As illustrated in Figure 2, decreasing the tolerance value leads to a larger coarse problem size, consistent with theoretical expectations. Correspondingly, the spectral condition number and the number of CG iterations are reduced as tol decreases, confirming the improved efficiency of the adaptive strategy. This reduction in iteration count also results in lower CPU times for the smallest tolerance values.

Test 2 – optimality. We now assess the optimality of both standard and adaptive BDDC preconditioners, focusing on their performance under mesh refinement. The analysis considers the five mesh levels (Mesh1—Mesh5), whose properties are sum-

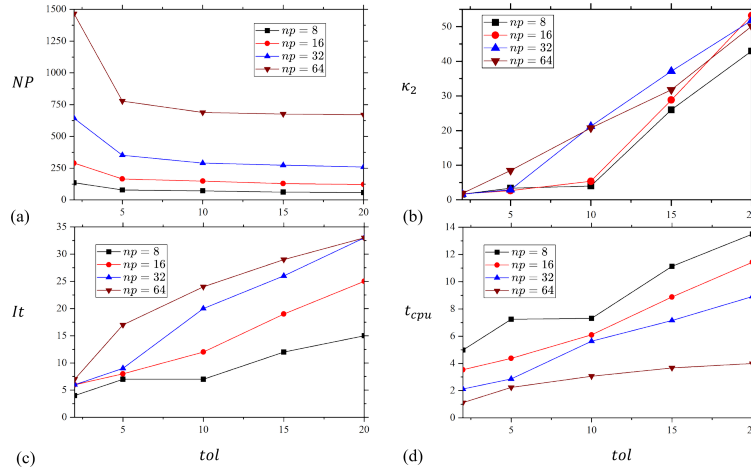


Fig. 2 Test 1 – varying the adaptive tolerance. (a) Coarse problem size (NP), (b) spectral condition number (κ_2), (c) average CG iterations (It) per time step, and (d) average CPU time (t_{cpu}) in seconds per time step as a function of tol and for different number of cores (np in labels).

marized in Table (a). The number of subdomains, corresponding to the number of processing cores, is fixed at 16. Table 1 presents a comparison between the standard BDDC method (with a coarse space composed of vertex constraints and edge averages (V+E)) and the adaptive BDDC (aBDDC) method with a tolerance parameter of $tol = 2$. In both approaches, the spectral condition number and the number of Conjugate Gradient (CG) iterations remain bounded as the mesh is refined, confirming the solvers' optimal scalability. However, the aBDDC method features a substantially larger coarse space than the standard BDDC, which results in fewer CG iterations and up to a 65% reduction in CPU time for Mesh5.

Mesh	BDDC (V+E)				aBDDC ($tol = 2$)			
	NP	κ_2	It	t_{cpu}	NP	κ_2	It	t_{cpu}
Mesh1	133	17.30	21	0.11	170	1.86	5	0.04
Mesh2	128	28.72	32	1.51	213	2.00	7	0.49
Mesh3	131	37.39	30	2.02	229	1.69	5	0.65
Mesh4	118	35.86	35	9.59	273	1.82	7	3.21
Mesh5	141	33.58	32	10.91	290	1.73	6	3.58

Table 1 Test 2 – optimality test. Standard BDDC (with coarse space V+E) and adaptive BDDC (aBDDC, with $tol = 2$). We report the spectral condition number (κ_2), the average CG iterations (It) per time step and the average CPU time (t_{cpu}) in seconds per time step – with respect to progressive mesh refinement, with 16 processors (and corresponding partitions).

Test 3 – strong scaling. In this section, we present the strong scalability analysis conducted to assess the efficiency and parallel performance of the BDDC preconditioners. The test is performed using Mesh5. As shown in Table 2, increasing the

number of cores (or subdomain partitions) from 8 to 64 leads to a consistent decrease in the CPU time required to solve the elliptic problem (t_{cpu}) across all configurations. The parallel speedup for BDDC seems to stagnate for the 64 processors case: we believe it is related to the parallel architecture considered in these tests, where we had to consider 2 computational nodes for 64 processors, thus requiring additional internodes communication with respect to the other cases. We finally observe that the setup cost of aBDDC is generally about 10 times larger than the one of standard BDDC. However, the setup is done once for all. Since in a heartbeat Bidomain simulation we have to solve thousands of linear systems with the same matrix, despite the larger setup costs, aBDDC results anyway cheaper than standard BDDC thanks to the significant reduction of CG iterations.

cores	BDDC (V+E)					aBDDC ($tol = 2$)				
	NP	κ_2	It	t_{cpu}	S_p	NP	κ_2	It	t_{cpu}	S_p
8	67	32.23	20	15.48	1.0	159	1.61	4	5.91	1.0
16	131	33.72	30	10.26	1.5	342	1.75	5	3.99	1.5
32	317	34.33	28	4.38	3.6	789	1.92	5	2.04	2.9
64	889	45.95	33	4.27	3.6	1909	2.33	7	1.40	4.2

Table 2 Test 3 - strong scaling tests. Standard BDDC (with coarse space V+E) and adaptive BDDC (aBDDC, with $tol = 2$). We compare the robustness of the coarse problem size (NP), the spectral condition number (κ_2), the average CG iterations (It) per time step, the average CPU time (t_{cpu}) in seconds per time step, and the parallel speedup (S_p), with respect to an increasing number of cores. The parallel speedup S_p is computed as the ratio between the CPU time with minimum number of cores (in this case 8) and the CPU time with np cores.

Acknowledgements N.M.M.H., L.F.P., and S.S. have been supported by grants of INDAM-GNCS. L.F.P. has been supported by grants of MIUR (PRIN_202232A8AN_002, PRIN_P2022B38NR_001), funded by European Union - Next Generation EU. S.S., N.M.M.H. and T.A. have been supported by grants of MIUR (PRIN_202232A8AN_003 and PRIN_P2022B38NR_002), funded by European Union - Next Generation EU. R.C. has been supported by grants of NSF and MOST of China (62161160312 and 2023YFC2507501). Computational resources are provided by INDACO Platform, which is a project of High Performance Computing at the University of MILAN <http://www.unimi.it>.

References

1. Abdelhamid, T., Huynh, N.M.M., Zampini, S., Chen, R., Pavarino, L.F., Scacchi, S.: Adaptive BDDC preconditioners for the bidomain model on unstructured ventricular finite element meshes. *Comput. Methods Appl. Mech. Eng.* **447**, 118366 (2025)
2. Africa, P.C., Piersanti, R., Fedele, M., Dedè, L., Quarteroni, A.M.: *lifex-fiber*: an open tool for myofibers generation in cardiac computational models. *BMC Bioinformatics* **24**, 143 (2023)
3. Colli Franzone, P., Pavarino, L.F., Scacchi, S.: *Mathematical Cardiac Electrophysiology*. Springer (2014)

4. Dohrmann, C.R.: A preconditioner for substructuring based on constrained energy minimization. *SIAM J. Sci. Comput.* **25**(1), 246–258 (2003)
5. Drummond, D., Gonsard, A.: Definitions and characteristics of patient digital twins being developed for clinical use: Scoping review. *J. Med. Internet Res.* **26**, e58504 (2024)
6. Fedele, M., Quarteroni, A.M.: Polygonal surface processing and mesh generation tools for the numerical simulation of the cardiac function. *Int. J. Num. Meth. Biomed. Eng.* **37**(4), e3435 (2021)
7. Huynh, N.M.M.: Newton–Krylov–BDDC deluxe solvers for non-symmetric fully implicit time discretizations of the bidomain model. *Numer. Math.* **152**(4), 841–879 (2022)
8. Mandel, J., Sousedik, B., Šístek, J.: Adaptive BDDC in three dimensions. *Math. Comput. Simul.* **82**(10), 1812–1831 (2012)
9. Pavarino, L.F., Scacchi, S.: Multilevel additive Schwarz preconditioners for the bidomain reaction–diffusion system. *SIAM J. Sci. Comput.* **31**, 420–443 (2008)
10. ten Tusscher, K.H.W.J., Panfilov, A.V.: Alternans and spiral breakup in a human ventricular tissue model. *Am. J. Physiol. Heart Circ. Physiol.* **291**, H1088–H1100 (2006)
11. Zampini, S.: Balancing Neumann–Neumann methods for the cardiac bidomain model. *Numer. Math.* **123**(2), 363–393 (2013)
12. Zampini, S.: Dual-primal methods for the cardiac bidomain model. *Math. Mod. Meth. Appl. Sci.* **24**(4), 667–696 (2014)
13. Zampini, S.: Inexact BDDC methods for the cardiac bidomain model. In: J. Erhel et al. (ed.) *Domain Decomposition Methods in Science and Engineering XXI*, pp. 247–255. Springer International Publishing, Cham (2014)
14. Zampini, S.: PCBDDC: A class of robust dual-primal methods in PETSc. *SIAM J. Sci. Comput.* **38**, S282–S306 (2016)
15. Zampini, S., Tu, X.: Multilevel balancing domain decomposition by constraints deluxe algorithms with adaptive coarse spaces for flow in porous media. *SIAM J. Sci. Comput.* **39**, A1389–A1415 (2017)