Dual-Primal Preconditioners for Newton-Krylov Solvers for the Cardiac Bidomain Model

Ngoc Mai Monica Huynh, Luca F. Pavarino, and Simone Scacchi

1 Introduction

We present here an overview of Newton-Krylov solvers for implicit time discretizations of the cardiac Bidomain equations, preconditioned by Balancing Domain Decomposition with Constraints (BDDC) [5] or Dual-Primal Finite Element Tearing and Interconnecting (FETI-DP) [7] algorithms.

The Bidomain model describes the propagation of the electric signal in the cardiac tissue by means of two parabolic partial differential equations (PDEs) [3, 13]; it is coupled through the non-linear reaction term to a system of ordinary differential equations (ODEs), modeling the ionic currents through the cell membrane and the associated opening and closing process of ionic channel gates.

One of the main issues to face when computing these systems is the choice of an appropriate solver, which can combine computational efficiency and accuracy in representing the solution. As a matter of fact, the need of accurately representing phenomena both at macroscopic and at microscopic level leads to time and space discretizations with millions of degrees of freedom (dofs) or more. The solution of the associated large discrete systems for increasing dimensions represent a challenging computation, requiring efficient parallel solvers [4, 14].

In this work we show some parallel numerical results obtained with two non-linear solvers, each of whom derives from a different solution strategy: a monolithic (or *coupled*) solution approach and a staggered (or *decoupled*) approach. Both these approaches arise from an implicit time discretization of the Bidomain model, which

N. M. M. Huynh and L. F. Pavarino
Dipartimento di Matematica, Università degli Studi di Pavia, via Ferrata 5 - 27100 Pavia, Italy,
e-mail: ngocmaimonica.huynh01@universitadipavia.it, luca.pavarino@unipv.it

S. Scacchi
Dipartimento di Matematica, Università degli Studi di Milano, via Saldini 50 - 20133 Milano, Italy,
e-mail: simone.scacchi@unimi.it

is solved coupled to or decoupled from the ionic equations, respectively, as in Refs. [11, 12].

In Sec. 2 a brief description of the model is provided, while in Sec. 3 we present our solution strategies. Parallel numerical experiments in Sec. 4 using the PETSc library [1] end this work.

2 The Bidomain cardiac electrical model

The propagation of the electrical impulse in the cardiac tissue is modeled by a system of two parabolic reaction-diffusion PDEs coupled through the non-linear reaction term to a system of ODEs describing the flow of ionic currents inward and outward the cell membrane:

$$\begin{cases} \chi C_m \dfrac{\partial v}{\partial t} - \operatorname{div}\left(D_i \cdot \nabla u_i\right) + I_{\text{ion}}(v,w) = 0 \\ -\chi C_m \dfrac{\partial v}{\partial t} - \operatorname{div}\left(D_e \cdot \nabla u_e\right) - I_{\text{ion}}(v,w) = -I^e_{\text{app}} \quad \text{in } \Omega \times (0,T), \\ \dfrac{\partial w}{\partial t} - R(v,w) = 0 \end{cases} \tag{1}$$

where u_i and u_e are the intra- and extracellular potentials, $v(x,t) = u_i(x,t) - u_e(x,t)$ is the transmembrane potential and w represents the opening and closing process of the ionic channel gates in the cell membrane. Here, C_m is the membrane capacitance, I_{ion} the ionic membrane current (both for unit area of the membrane surface), χ is the membrane surface to volume ratio and I_{app} is the applied external current. This system is known in the literature as Bidomain model [3, 13].

In this work, we consider a phenomenological ionic model, named the Rogers-McCulloch ionic model [15]. More realistic and complex ionic models have been integrated in different numerical studies, see e.g. Refs. [4, 14].

3 Dual-Primal Newton-Krylov methods

Space and time discretizations. The cardiac domain Ω is discretized in space with a structured quasi-uniform grid of hexahedral finite elements, leading to the semi-discrete system

$$\begin{cases} \chi C_m \mathcal{M} \dfrac{\partial}{\partial t} \begin{bmatrix} \mathbf{u}_i \\ \mathbf{u}_e \end{bmatrix} + \mathcal{A} \begin{bmatrix} \mathbf{u}_i \\ \mathbf{u}_e \end{bmatrix} + \begin{bmatrix} M\, \mathbf{I}_{\text{ion}}(\mathbf{v},\mathbf{w}) \\ -M\, \mathbf{I}_{\text{ion}}(\mathbf{v},\mathbf{w}) \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ -M\, \mathbf{I}^{\mathbf{e}}_{\text{app}} \end{bmatrix}, \\ \dfrac{\partial \mathbf{w}}{\partial t} = R\left(\mathbf{v},\mathbf{w}\right), \end{cases} \tag{2}$$

with the stiffness and mass block-matrices $\mathcal{A} = \begin{bmatrix} A_i & 0 \\ 0 & A_e \end{bmatrix}$, $\mathcal{M} = \begin{bmatrix} M & -M \\ -M & M \end{bmatrix}$.

Regarding the time discretization, in the literature it is very common to adopt operator splitting strategies [2, 16] or implicit-explicit (IMEX) schemes [4, 17], in order to avoid the elevated computational costs related to the solution of the non-linear discrete problem. Here we propose two ways for the solution of the discretized system using the Backward Euler method: a monolithic, or *coupled*, solution strategy where at each time step we solve the non-linear system with the discrete Bidomain coupled with the ionic model, as in Refs. [8, 12], and a staggered, or *decoupled*, solution approach (as in Refs. [9, 11]). Both approaches rely on a preconditioned Krylov method nested within a Newton loop.

***Coupled* solution approach.** The monolithic strategy can be summarized in algorithmic steps as follows. At the n-th time step, solve the non-linear system $\mathbf{F}_{\text{coupled}}(\mathbf{s}^{n+1}) = \mathbf{0}$, with $\mathbf{s}^{n+1} = (\mathbf{u}_i^{n+1}, \mathbf{u}_e^{n+1}, \mathbf{w}^{n+1})$:

$$\mathbf{F}_{\text{coupled}}\left(\mathbf{s}^{n+1}\right) = \begin{cases} (\chi C_m \mathcal{M} + \tau \mathcal{A}) \begin{bmatrix} \mathbf{u}_i^{n+1} \\ \mathbf{u}_e^{n+1} \end{bmatrix} + \tau \begin{bmatrix} M\mathbf{I}_{\text{ion}}(\mathbf{v}^{n+1}, \mathbf{w}^{n+1}) \\ -M\mathbf{I}_{\text{ion}}(\mathbf{v}^{n+1}, \mathbf{w}^{n+1}) \end{bmatrix} - \mathbf{G} \\ \mathbf{w}^{n+1} - \tau R(\mathbf{v}^{n+1}, \mathbf{w}^{n+1}) - \mathbf{w}^n \end{cases}$$

with $\mathbf{G} = \chi C_m \mathcal{M} \begin{bmatrix} \mathbf{u}_i^n \\ \mathbf{u}_e^n \end{bmatrix} + \tau \begin{bmatrix} \mathbf{0} \\ -M\mathbf{I}_{\text{app}}^{\mathbf{e}} \end{bmatrix}$ and being $\tau = t_{n+1} - t_n$. This non-linear system is solved with a Newton method:

1. compute and solve the Jacobian linear system $\mathbf{DF}(\mathbf{s}^n)\ \mathbf{ds}^{n+1} = -\mathbf{F}(\mathbf{s}^n)$, where $\mathbf{ds}^{n+1} := (\mathbf{du}_i^{n+1}, \mathbf{du}_e^{n+1}, \mathbf{dw}^{n+1})$ is the increment at step $n+1$;
2. update $\mathbf{u}_i^{n+1} = \mathbf{u}_i^n + \mathbf{du}_i^{n+1}$, $\mathbf{u}_e^{n+1} = \mathbf{u}_e^n + \mathbf{du}_e^{n+1}$ and $\mathbf{w}^{n+1} = \mathbf{w}^n + \mathbf{dw}^{n+1}$.

Since the linear system in Step 1 is non-symmetric (due to the presence of the gating term), it is necessary to use the Generalized Minimal Residual method (GMRES) for its solution.

***Decoupled* solution approach.** As alternative to the previous strategy, the staggered approach requires first the solution of the ionic model, then solve and update the Bidomain equations. For each time step n,

a. given the intra- and extracellular potentials at the previous step, hence $\mathbf{v} := \mathbf{u}_i^n - \mathbf{u}_e^n$, compute the gating

$$\mathbf{w}^{n+1} - \tau R(\mathbf{v}, \mathbf{w}^{n+1}) = \mathbf{w}^n;$$

b. solve and update the Bidomain non-linear system. Given $\mathbf{u}_{i,e}^n$ at the previous time step and given $\mathbf{w}^{n+1}$, compute $\mathbf{u}^{n+1} = (\mathbf{u}_i^{n+1}, \mathbf{u}_e^{n+1})$ by solving the system $\mathbf{F}_{\text{decoupled}}(\mathbf{u}^{n+1}) = \mathbf{G}$

$$\mathbf{F}_{\text{decoupled}}\left(\mathbf{u}^{n+1}\right) = (\chi C_m \mathcal{M} + \tau \mathcal{A}) \begin{bmatrix} \mathbf{u}_i^{n+1} \\ \mathbf{u}_e^{n+1} \end{bmatrix} + \tau \begin{bmatrix} M\mathbf{I}_{\text{ion}}(\mathbf{v}^{n+1}, \mathbf{w}^{n+1}) \\ -M\mathbf{I}_{\text{ion}}(\mathbf{v}^{n+1}, \mathbf{w}^{n+1}) \end{bmatrix},$$

$$\mathbf{G} = \chi C_m \mathcal{M} \begin{bmatrix} \mathbf{u}_i^n \\ \mathbf{u}_e^n \end{bmatrix} + \tau \begin{bmatrix} \mathbf{0} \\ -M\mathbf{I}_{\text{app}}^{\mathbf{e}} \end{bmatrix}.$$

The Jacobian linear system associated to the non-linear problem in step (b) is symmetric, thus allowing us to use the Conjugate Gradient (CG) method within each Newton iteration.

Dual-primal preconditioners. In both approaches, a linear system has to be solved within each Newton loop, either by GMRES (in case of the *coupled* approach) or by CG method (for the *decoupled* case), preconditioned by a dual-primal substructuring algorithm.

In this work, we focus on the most common dual-primal iterative substructuring algorithms, the BDDC and FETI-DP methods.
FETI-DP methods were first proposed in Ref. [7] and are based on the transposition of the linear system to a constrained minimization problem.
Conversely, BDDC methods were introduced in Ref. [5] as an alternative to FETI-DP and provide a preconditioner for the discretized linear problem.

Convergence rate bound. In Ref. [10] these two algorithms are shown to be spectrally equivalent, thus allowing us to derive a convergence rate estimate for the preconditioned operator, which holds for both preconditioners in case the same coarse space is chosen. In the *coupled* approach, the bound is related to the residual at the m-th iteration of GMRES, while in the *decoupled* strategy the bound is for the condition number. Details on the derivation of both bounds can be found in the works of the authors [9, 8].

4 Numerical experiments

The parallel numerical experiments are performed on an idealized left ventricular geometry, modeled as a portion of half truncated ellipsoid, see Fig. 1.
Boundary and initial conditions represent an isolated tissue with resting potential. We simulate the initial excitation process on the time interval $[0, 2]$ ms following an extracellular stimulus.

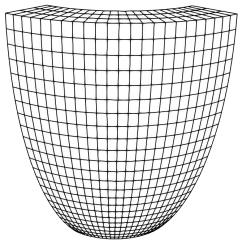

Fig. 1: Computational domain.

Different supercomputers are taken into account: the Galileo cluster from the Cineca centre (a Linux Infiniband cluster equipped with 1084 nodes, each with 36 2.30 GHz Intel Xeon E5-2697 v4 cores and 128 GB/node, for a total of 39024 cores, `www.hpc.cineca.it`) for the tests related to the *coupled* solution approach and the weak scaling of the *decoupled* case; the Linux cluster Indaco at the University of Milan (a Linux Infiniband cluster with 16 nodes, each carrying 2 processors Intel

Xeon E5-2683 v4 2.1 GHz with 16 cores each, `https://www.indaco.unimi.it/`) for the strong scaling of the *decoupled* approach.

Our C code is based on the parallel library PETSc [1] from the Argonne National Laboratory. BDDC and FETI-DP preconditioners are built-in in PETSc library, both applied with default parameters (coarse space made up of vertices and edge averages, direct subdomain solver with a LU factorization, etc), while *Boomer* Algebraic MultiGrid (bAMG) is from the Hypre library [6]. In our tests, we always assign one subdomain to each processor. In the strong scaling tests, part of the speedup comes from the superlinear computational complexity of the sparse subdomain solvers based on LU factorization.

We manually implement the Newton method for the coupled case, with an absolute residual stopping criterion with tolerance 10^{-4}, while for the decoupled case we use the default non-linear solver (SNES) from PETSc library and we adopt the default SNES convergence test as stopping criterion, based on the comparison of the L^2-norm of the non-linear function at the current iterate and at the current step (see PETSc manual [1] for tolerance values and further details). The linear systems arising in Steps 1 and (b) of the two approaches are solved with GMRES and CG methods respectively, both using PETSc default stopping criteria and default tolerances. We compare the following quantities: the average Newton iterations per time step *nit*, the average linear iterations per Newton iteration *lit* and the average CPU solution time per time step *time* in seconds.

Coupled tests. The linear system arising from the discretization of the Jacobian problem at each Newton step is solved with GMRES method, preconditioned by BDDC preconditioners and bAMG.

Coupled weak scaling. We report here a weak scaling test. We fix the local mesh to $12 \cdot 12 \cdot 12$ elements and we increase the number of subdomains (and therefore the number of processors) from 32 to 256, yielding an ellipsoidal portion of increasing dimensions. It is clear from Table 1 that BDDC performs better than bAMG in terms of average number of linear iterations per non-linear step, as this parameter is lower for BDDC and does not increase with the number of processors. As a matter of fact, there is an increasing reduction rate up to 90% for the average linear iterations. In contrast, BDDC's average CPU time is higher than bAMG CPU time (we do not have a clear explanation of this fact), but we remark that BDDC timings do not increase significantly when the number of processors is increased from 32 to 256, while bAMG timings more than double.

Coupled strong scaling. We fix the global mesh to $128 \cdot 128 \cdot 24$ elements (resulting in more than 1 million of global dofs) and we increase the number of processors from 32 to 256. As the number of processor increases, the local number of dofs decreases and BDDC's average number of linear iterations and CPU times decrease (see Table 2), while bAMG iterations increase and the CPU timings decrease less than expected, even if they are lower than BDDC timings. Moreover, in order to test the efficiency of the proposed solver on the parallel architecture, we compute the parallel speedup $\frac{T_1}{T_N}$, which is the ratio between the runtime T_1 needed by 1 (or N_1)

Table 1: *Coupled* weak scaling test. Local mesh of $12 \cdot 12 \cdot 12$ elements. Comparison of Newton-Krylov solvers preconditioned by BDDC and bAMG. Cluster: Galileo.

procs.	global n. dofs	BDDC			bAMG		
		nit	lit	time	nit	lit	time
32	180,075	2	45	6.8	2	142	1.5
64	356,475	2	32	6.9	2	145	1.9
128	705,675	2	23	7.0	2	158	2.1
256	1,404,075	2	23	8.5	2	212	3.2

processor and the average runtime T_N needed by N processors to solve the problem. Here, we set $N_1 = 32$. While bAMG is sub-optimal, BDDC outperforms the ideal linear speedup.

Table 2: *Coupled* strong scaling test. Global mesh of $128 \cdot 128 \cdot 24$ elements (1,248,075 dofs). Comparison of Newton-Krylov solvers preconditioned by BDDC and bAMG. Parallel speedup (S_p), with ideal speedup in brackets. Cluster: Galileo.

procs.	BDDC				bAMG			
	nit	lit	time	S_p	nit	lit	time	S_p
32	2	37	189.3	-	2	187	15.1	-
64	2	44	59.1	3.2 (2)	2	222	9.2	1.6 (2)
128	2	29	20.1	9.4 (4)	2	240	5.3	2.8 (4)
256	2	46	10.2	18.5 (8)	2	280	3.2	4.7 (8)

Decoupled tests. The outer Newton loop is solved with the non-linear solver SNES of the PETSc library, which implements a Newton method with cubic backtracking linesearch. The linear system arising from the discretization of the Jacobian problem at each Newton step is solved with the CG method, preconditioned by BDDC or FETI-DP preconditioners.

Decoupled weak scaling. We fix here the local mesh size to $16 \cdot 16 \cdot 16$ and we increase the number of processors from 32 to 2048. Also in this case, the good performance of the dual-primal algorithms is confirmed by the average number of linear iterations per Newton step, which is low and remains stable as the number of subdomains increases (see Table 3).

Decoupled strong scaling. We now compare the performance of the dual-primal preconditioners while varying the number of processors from 64 to 256 over a time interval of $[0, 100]$ ms, for a total of 2000 time steps. The global mesh is fixed to $192 \cdot 96 \cdot 24$ elements (936,050 dofs). We can observe an overall reduction of the CPU time while increasing the number of subdomains from 64 to 128. As concerns FETI-DP behavior, the increase of average CPU time and average number of linear

Table 3: *Decoupled* weak scaling test. Local mesh of $16 \cdot 16 \cdot 16$ elements. Comparison of Newton-Krylov solvers preconditioned by BDDC and FETI-DP. Cluster: Galileo.

procs	dofs	BDDC			FETI-DP		
		nit	lit	time	nit	lit	time
32	278,850	1	30	5.4	1	20	4.7
64	549,250	1	37	6.2	1	20	6.5
128	1,090,050	1	26	7.5	1	19	6.6
256	2,171,650	1	25	8.7	1	17	10.7
512	4,309,890	1	27	10.5	1	18	11.4
1024	8,586,370	1	28	12.5	1	19	11.0
2048	17,139,330	1	28	26.6	1	19	21.4

iterations between 128 and 256 processors is unexpected and further investigations should be devoted to explain this result (Figure 2).

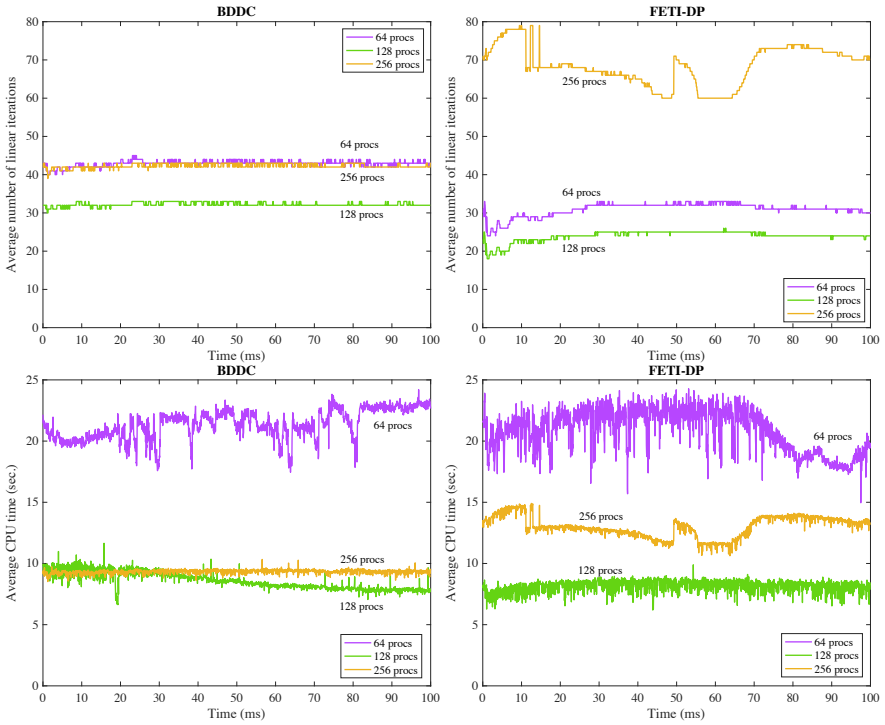

Fig. 2: *Decoupled* strong scaling. Global mesh of $192 \cdot 96 \cdot 24$ elements (936,050 dofs). Comparison between BDDC (left column) and FETI-DP (right column) preconditioners. Top: average number of linear iterations per time step; bottom: average CPU time in seconds of each SNES solver call. Cluster: Indaco.

5 Conclusion

We designed and numerically tested two different solution strategies for the solution of implicit time discretizations of the Bidomain model. Each of these solvers is preconditioned by a dual-primal substructuring algorithm, which perform better than the algebraic multigrid method in terms of number of iterations, scalability, and speedup, even if the computational times of algebraic multigrid are still better for these parameter settings. Future works should extend these solver to the solution of coupled cardiac electro-mechanical models and to more complex ionic models.

References

1. S. Balay et al. PETSc users manual. Technical Report ANL-95/11 - Revision 3.15, Argonne National Laboratory, 2021.
2. H. Chen, X. Li, and Y. Wang. A two-parameter modified splitting preconditioner for the Bidomain equations. *Calcolo*, 56(2), 2019.
3. P. Colli Franzone, L. F. Pavarino, and S. Scacchi. *Mathematical cardiac electrophysiology*, volume 13 of *MS&A. Modeling, Simulation and Applications*. Springer, Cham, 2014.
4. P. Colli Franzone, L. F. Pavarino, and S. Scacchi. A numerical study of scalable cardiac electro-mechanical solvers on HPC architectures. *Frontiers in physiology*, 9:268, 2018.
5. C. R. Dohrmann. A preconditioner for substructuring based on constrained energy minimization. *SIAM J. Sci. Comput.*, 25(1):246–258, 2003.
6. R. D. Falgout and U. M. Yang. hypre: A library of high performance preconditioners. In *International Conference on Computational Science*, pages 632–641. Springer, 2002.
7. C. Farhat et al. FETI-DP: a dual-primal unified FETI method. I. A faster alternative to the two-level FETI method. *Internat. J. Numer. Methods Engrg.*, 50(7):1523–1544, 2001.
8. N. M. M. Huynh. Newton-Krylov-BDDC deluxe solvers for non-symmetric fully implicit time discretizations of the Bidomain model. *arXiv preprint arXiv:2102.08736*, 2021.
9. N. M. M. Huynh, L. F. Pavarino, and S. Scacchi. Parallel Newton-Krylov-BDDC and FETI-DP deluxe solvers for implicit time discretizations of the cardiac Bidomain equations. *arXiv preprint arXiv:2101.02959*, 2021.
10. J. Li and O. B. Widlund. FETI-DP, BDDC, and block Cholesky methods. *International journal for numerical methods in engineering*, 66(2):250–271, 2006.
11. M. Munteanu and L. F. Pavarino. Decoupled Schwarz algorithms for implicit discretizations of nonlinear Monodomain and Bidomain systems. *Mathematical Models and Methods in Applied Sciences*, 19(07):1065–1097, 2009.
12. M. Murillo and X.-C. Cai. A fully implicit parallel algorithm for simulating the non-linear electrical activity of the heart. *Numer. linear algebra with appl.*, 11(2-3):261–277, 2004.
13. M. Pennacchio, G. Savaré, and P. Colli Franzone. Multiscale modeling for the bioelectric activity of the heart. *SIAM Journal on Mathematical Analysis*, 37(4):1333–1370, 2005.
14. A. Quarteroni, T. Lassila, S. Rossi, and R. Ruiz-Baier. Integrated Heart—Coupling multiscale and multiphysics models for the simulation of the cardiac function. *Computer Methods in Applied Mechanics and Engineering*, 314:345–407, 2017.
15. J. M. Rogers and A. D. McCulloch. A collocation-Galerkin finite element model of cardiac action potential propagation. *IEEE Trans. Biomed. Engrg*, 41(8):743–757, 1994.
16. J. Sundnes, G. T. Lines, and A. Tveito. An operator splitting method for solving the bidomain equations coupled to a volume conductor model for the torso. *Mathematical biosciences*, 194(2):233–248, 2005.
17. S. Zampini. Dual-primal methods for the cardiac bidomain model. *Mathematical Models and Methods in Applied Sciences*, 24(04):667–696, 2014.