
Multilevel Schwarz and Multigrid Preconditioners for the Bidomain System

Simone Scacchi¹ and Luca F. Pavarino²

¹ Dipartimento di Matematica, Università di Pavia, Via Ferrata 1, 27100 Pavia, Italy. simone.scacchi@unipv.it

² Dipartimento di Matematica, Università di Milano, Via Saldini 50, 20133 Milano, Italy. luca.pavarino@mat.unimi.it

Summary. Two parallel and scalable multilevel preconditioners for the Bidomain system in computational electrocardiology are introduced and studied. The Bidomain system, consisting of two degenerate parabolic reaction-diffusion equations coupled with a stiff system of several ordinary differential equations, generates very ill-conditioned discrete systems when discretized with semi-implicit methods in time and finite elements in space. The multilevel preconditioners presented in this paper attain the best performance to date, both in terms of convergence rate and solution time and outperform the simpler one-level preconditioners previously introduced. Parallel numerical results, using the PETSc library and run on Linux Clusters, show the scalability of the proposed preconditioners and their efficiency on large-scale simulations of a complete cardiac cycle.

1 Introduction

We introduce and study two parallel and scalable multilevel preconditioners for the Bidomain system in computational electrocardiology. These preconditioners improve upon the recent studies [2, 7], where one-level block Jacobi preconditioners were found to perform satisfactorily for the simplified Monodomain model but not for the more complex Bidomain system. The latter is a multiscale model of the cardiac bioelectrical activity, consisting of two degenerate parabolic reaction-diffusion equations describing the intra and extracellular potentials of the anisotropic cardiac tissue (macroscale), coupled through the nonlinear reaction term with a stiff system of several ordinary differential equations describing the ionic currents through the cellular membrane (microscale).

The numerical resolution of the Bidomain system is computationally very expensive, because of the interaction of the different scales in space and time, the degenerate nature of the PDEs involved and the very severe ill-conditioning of the discrete systems arising at each time step. Fully implicit methods in time have been considered in few studies, see e.g. [6] and require the solution of nonlinear systems at each time step. Most numerical studies employ semi-implicit (IMEX) methods in

time that only require the solution of linear systems at each time step. Many different preconditioners have been proposed in order to devise efficient iterative solvers for such linear systems: diagonal preconditioners [9], Symmetric Successive Over Relaxation [8, 14], Block Jacobi (BJ) preconditioners with incomplete LU factorization (ILU) for each block [2, 7, 13], multigrid [15].

The multilevel preconditioners presented in this paper attain the best performance to date, both in terms of convergence rate and solution time and outperform the simpler one-level preconditioners previously introduced. Parallel numerical results, using the PETSc library (see [1]) and run on Linux Clusters, show the scalability of the proposed preconditioners and their efficiency on large-scale simulations of a complete cardiac cycle.

2 The Mathematical Model

The macroscopic Bidomain model represents the cardiac tissue as the superposition of two anisotropic continuous media, the intra (i) and extra (e) cellular media, coexisting at every point of the tissue and separated by a distributed continuous cellular membrane. The cardiac tissue is traditionally modeled as an arrangement of fibers rotating clockwise from epicardium to endocardium [11] and, according to [4], presents a laminar organization, which consists of a set of muscle sheets, moving radially from epicardium to endocardium. Therefore, at any point \mathbf{x} , it is possible to identify a triplet of orthonormal principal axes $\mathbf{a}_l(\mathbf{x})$, $\mathbf{a}_t(\mathbf{x})$, $\mathbf{a}_n(\mathbf{x})$, with \mathbf{a}_l parallel to the local fiber direction, \mathbf{a}_t and \mathbf{a}_n tangent and orthogonal to the radial laminae respectively. The anisotropic conductivity properties of the tissue are described by the conductivity coefficients in the intra and extracellular media $\sigma_l^{i,e}$, $\sigma_t^{i,e}$, $\sigma_n^{i,e}$ measured along the corresponding direction \mathbf{a}_l , \mathbf{a}_t , \mathbf{a}_n and by the conductivity tensors $\mathbf{D}_i(\mathbf{x})$ and $\mathbf{D}_e(\mathbf{x})$, given by

$$\mathbf{D}_{i,e}(\mathbf{x}) = \sigma_l^{i,e} \mathbf{a}_l(\mathbf{x})\mathbf{a}_l^T(\mathbf{x}) + \sigma_t^{i,e} \mathbf{a}_t(\mathbf{x})\mathbf{a}_t^T(\mathbf{x}) + \sigma_n^{i,e} \mathbf{a}_n(\mathbf{x})\mathbf{a}_n^T(\mathbf{x}).$$

The intra and extracellular electric potentials u_i , u_e in the cardiac domain Ω are described in the Bidomain model by the following parabolic reaction-diffusion system coupled with a system of ODEs for the ionic variables w :

$$\begin{cases} c_m \frac{\partial v}{\partial t} - \operatorname{div}(\mathbf{D}_i \nabla u_i) + I_{ion}(v, w) = 0 & \text{in } \Omega \times (0, T) \\ -c_m \frac{\partial v}{\partial t} - \operatorname{div}(\mathbf{D}_e \nabla u_e) - I_{ion}(v, w) = -I_{app}^e & \text{in } \Omega \times (0, T) \\ \frac{\partial w}{\partial t} - R(v, w) = 0, & \text{in } \Omega \times (0, T), \end{cases} \quad (1)$$

with boundary conditions $\mathbf{n}^T \mathbf{D}_{i,e} \nabla u_{i,e} = 0$ in $\partial\Omega \times (0, T)$ and initial conditions $v(\mathbf{x}, 0) = v_0(\mathbf{x})$, $w(\mathbf{x}, 0) = w_0(\mathbf{x})$ in Ω . Here c_m is the capacitance per unit area times the surface to volume ratio; $v = u_i - u_e$ is the transmembrane potential; I_{app}^e is the applied current; I_{ion} and R model the ionic currents and depend on the choice of the membrane model. In this work we consider the LR1 model, see [5]. Existence and regularity results for this degenerate system are proved in [3] and [12]. The system uniquely determines v , while the potentials u_i and u_e are defined only up to a same additive time-dependent constant.

3 Discretization and Numerical Methods

System (1) is discretized by the finite element method in space and a semi-implicit method in time. The space discretization is obtained meshing the cardiac domain Ω with a structured grid of hexahedral \mathbb{Q}_1 elements and introducing the associated finite element space V_h . A semidiscrete problem is obtained by applying a standard Galerkin procedure. We denote by M the symmetric mass matrix, by $A_{i,e}$ the symmetric stiffness matrices associated to the intra and extra-cellular anisotropic conductivity tensors, respectively, and by $I_{ion}^h, I_{app}^{e,h}$ the finite element interpolants of I_{ion} and I_{app}^e , respectively. The time discretization is performed by a semi-implicit method using for the diffusion term the implicit Euler method, while the nonlinear reaction term I_{ion} is treated explicitly.

As a consequence, the full evolution system is decoupled by first solving the ODEs system (given the potential \mathbf{v}^n at the previous time-step)

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

and then solving for $\mathbf{u}_i^{n+1}, \mathbf{u}_e^{n+1}$ the linear system

$$\left(\frac{c_m}{\Delta t} \begin{bmatrix} M & -M \\ -M & M \end{bmatrix} + \begin{bmatrix} A_i & 0 \\ 0 & A_e \end{bmatrix} \right) \begin{pmatrix} \mathbf{u}_i^{n+1} \\ \mathbf{u}_e^{n+1} \end{pmatrix} = \frac{c_m}{\Delta t} \begin{pmatrix} M(\mathbf{u}_i^n - \mathbf{u}_e^n) \\ M[-\mathbf{u}_i^n + \mathbf{u}_e^n] \end{pmatrix} + \begin{pmatrix} M[-\mathbf{I}_{ion}^h(\mathbf{v}^n, \mathbf{w}^{n+1})] \\ M[\mathbf{I}_{ion}^h(\mathbf{v}^n, \mathbf{w}^{n+1}) - \mathbf{I}_{app}^{e,h}] \end{pmatrix}, \quad (2)$$

where $\mathbf{v}^n = \mathbf{u}_i^n - \mathbf{u}_e^n$. The iteration matrix is symmetric semidefinite, having the zero eigenvalue associated to the $(\mathbf{1}, \mathbf{1})$ eigenvector, therefore, as in the continuous model, \mathbf{u}_i^n and \mathbf{u}_e^n are determined only up to the same additive time-dependent constant, chosen by imposing the condition $\mathbf{1}^T M \mathbf{u}_e^n = 0$. From [2] we know that the iteration matrix is very ill conditioned and we need an efficient preconditioner.

4 Parallel Implementation and Preconditioners

The parallel strategy consists of partitioning the computational domain into subdomains of the same size and assign them to different processors. The linear system (2) is solved with the parallel PCG method of the PETSc library. We will compare three different preconditioners.

Block Jacobi Preconditioner (BJ), i.e. a block diagonal matrix with blocks built from the local restriction of matrix A to each subdomain; on each block, we use an ILU(0) solver.

V-cycle Multigrid Preconditioner (MG): the linear system at each time step is solved with a five-level V-cycle Multigrid method (MG(5)). The smoother used for all but the coarsest level is a single iteration of CG with BJ-ILU(0) preconditioner. On the coarsest level we solve the system using the PCG preconditioned by BJ-ILU(0).

Symmetrized Multiplicative Multilevel Schwarz Preconditioner (SMMS). Let be $\Omega^{(i)}, i = 0, \dots, M$ a family of nested triangulations of Ω , coarsening from M to 0, and $A^{(i)}$ the matrix obtained by discretizing (1) on $\Omega^{(i)}$; so $A^{(M)} = A$. $R^{(i)}$ are the restriction operators from $\Omega^{(i+1)}$ to $\Omega^{(i)}$. We decompose Ω into N overlapping subdomains, hence each grid $\Omega^{(i)}$ is decomposed into N overlapping subgrids $\Omega_k^{(i)}$

for $k = 1, \dots, N$, such that the overlap $\delta^{(i)}$ at level $i = 1, \dots, M$ is equal to the mesh size $h^{(i)}$ of the grid $\Omega^{(i)}$. Let $R_k^{(i)}$ be the restriction operator from $\Omega^{(i)}$ to $\Omega_k^{(i)}$ and define $A_k^{(i)} := R_k^{(i)} A^{(i)} R_k^{(i)T}$. The action of this preconditioner on a given residual \mathbf{r} is given by:

$$\begin{aligned}
\mathbf{u}^{(M)} &\leftarrow \sum_{k=1}^N R_k^{(M)T} A_k^{(M)-1} R_k^{(M)} \mathbf{r} \\
\mathbf{r}^{(M-1)} &\leftarrow R^{(M-1)} (\mathbf{r} - A^{(M)} \mathbf{u}^{(M)}) \\
\mathbf{u}^{(M-1)} &\leftarrow \sum_{k=1}^N R_k^{(M-1)T} A_k^{(M-1)-1} R_k^{(M-1)} \mathbf{r}^{(M-1)} \\
&\dots \\
\mathbf{u}^{(0)} &\leftarrow A^{(0)-1} \mathbf{r}^{(0)}, \quad \mathbf{u}^{(1)} \leftarrow \mathbf{u}^{(1)} + R^{(0)T} \mathbf{u}^{(0)} \\
\mathbf{u}^{(1)} &\leftarrow \mathbf{u}^{(1)} + \sum_{k=1}^N R_k^{(1)T} A_k^{(1)-1} R_k^{(1)} (\mathbf{r}^{(1)} - A^{(1)} \mathbf{u}^{(1)}) \\
&\dots \\
\mathbf{u}^{(M)} &\leftarrow \mathbf{u}^{(M)} + R^{(M-1)T} \mathbf{u}^{(M-1)} \\
\mathbf{u}^{(M)} &\leftarrow \mathbf{u}^{(M)} + \sum_{k=1}^N R_k^{(M)T} A_k^{(M)-1} R_k^{(M)} (\mathbf{r}^{(M)} - A^{(M)} \mathbf{u}^{(M)}) \\
\mathbf{u} &\leftarrow \mathbf{u}^{(M)}
\end{aligned}$$

We implemented this method with 5 levels, hence in the remainder we denote it by SMMS(5). For details see [10].

5 Numerical Results

The numerical experiments were performed on two distributed memory parallel architectures, the IBM CLX/1024 Linux cluster of the Cineca Consortium (www.cineca.it), with 1024 processors Intel Xeon Pentium IV (3 GHz, 512 KB cache) grouped into 512 nodes of 2 processors (total RAM = 1 TB), and the Ulisse Linux cluster of the Department of Mathematics of the University of Milan (cluster.mat.unimi.it), with 72 processors Xeon (2.4 GHz) grouped into 36 nodes of 2 processors. Our FORTRAN code is based on the parallel library PETSc from the Argonne National Laboratory [1].

Test 1: standard speedup. We simulate the initial depolarization of a thin slab of cardiac tissue, having dimensions of $2.56 \times 2.56 \times 0.01 \text{ cm}^3$, applying a stimulus of 200 mA/cm^3 for 1 ms on a small volume of $2 \times 2 \times 2$ elements at a vertex of the domain. The global mesh is fixed to be of $257 \times 257 \times 2$ nodes (264196 unknowns) and the number of subdomains varies from 1 to 16. The model is run for 40 time steps of 0.05 ms , i.e. for a time interval of 2 ms on the Linux cluster of the University of Milan. In table 1, we report the average number of PCG iterations per time step, needed to reduce the l^2 norm of the residual smaller than 10^{-4} , the average condition number per time step and the average time needed to solve the linear system. Both the multilevel methods are scalable, in fact the iterations remain almost constant increasing the number of subdomains. The BJ speedup is low, because the number of iterations increases with the processors. The multilevel preconditioners behave well up to 8 processors, but with 16 the local problems are too small and the communication costs deteriorate the parallel performance.

Test 2: scaled speedup. In this test, we vary the number of subdomains from 8 to 128, keeping fixed the local mesh in each subdomain to $48 \times 48 \times 48$ nodes

Table 1. Test 1, standard speedup. IT:= average PCG iterations per time step; COND:= average condition number per time step; TIME:= average execution time per time step in seconds.

# SUB	BJ			MG(5)			SMMS(5)		
	IT.	COND.	TIME	IT.	COND.	TIME	IT.	COND.	TIME
1	95	1817	23.00	3	1.04	9.11	-	-	-
2	108	2209	22.27	3	1.04	4.63	3	1.04	4.95
4	109	2229	10.40	4	1.08	2.92	3	1.04	2.49
8	111	2367	5.31	4	1.11	1.58	3	1.04	1.28
16	114	2745	3.47	4	1.13	0.78	3	1.04	0.71

(221184 unknowns), hence varying the global number of degrees of freedom (d.o.f.) from 1.7×10^6 in the smallest case with 8 subdomains to 2.8×10^7 in the largest with 128 subdomains. As in test 1, we simulate the initial depolarization of a cardiac slab, running the model for 40 time steps on the CLX cluster of CINECA. Table 2 reports the average number of PCG iterations, the average condition number and the average solving time per time step. These results show the parallel scalability of the proposed multilevel methods, that have constant iteration counts, while the one-level BJ preconditioner has increasing iteration counts as expected. The solving time is also scalable, increasing of only 15~20 % going from 8 to 128 processors; for SMMS(5) this increase is due only to communications, because the iterations remain constant.

Table 2. Test 2, scaled speedup. Same format as in Table 1.

# SUB	D.O.F.	BJ			MG(5)			SMMS(5)		
		IT.	COND.	TIME	IT.	COND.	TIME	IT.	COND.	TIME
8	1769472	84	1461	29.9	4	1.26	16.1	4	1.14	16.2
16	3538944	93	2119	35.7	4	1.34	16.9	4	1.14	18.1
32	7077888	106	3543	44.8	5	1.43	16.8	4	1.14	16.3
64	14155776	115	4984	42.3	5	1.61	17.8	4	1.15	18.9
128	28311552	121	5165	51.4	5	1.58	18.5	4	1.14	19.6

Table 3. Test 3, complete cardiac cycle. IT:= average PCG iterations per time step; TIME:= average execution time per time step in seconds; TOTAL TIME:= total simulation time

PREC	aver. IT.	aver. TIME	TOTAL TIME
BJ	205	46.02 sec	29 h 49 m
MG(5)	8	11.11 sec	7 h 21 m
SMMS(5)	6	9.67 sec	6 h 26 m

Test 3: complete cardiac cycle. In this last test, we simulate a complete heartbeat (400 ms) in a portion of ventricle having dimension $2 \times 2 \times 0.5 \text{ cm}^3$, discretized by a cartesian grid of $200 \times 200 \times 50$ nodes (4×10^6 d.o.f.). We run the simulation on 36 processors of the Linux cluster of Milan. Table 3 reports the average

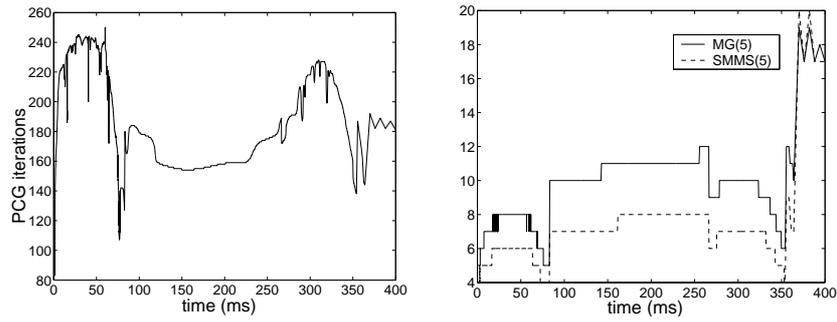


Fig. 1. Test 3. Time evolution of the PCG iterations with BJ preconditioners (left) and multilevel preconditioners MG(5), SMMS(5) (right).

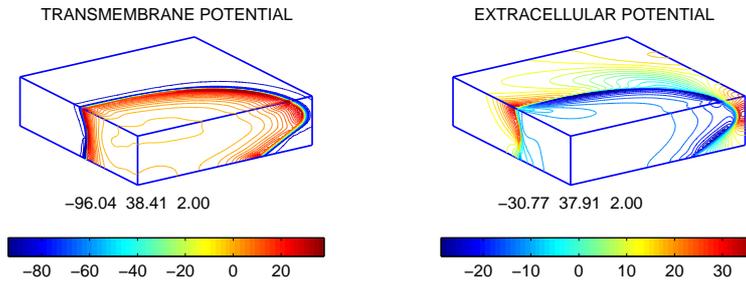


Fig. 2. Test 3. Patterns of level lines of the transmembrane and extracellular potentials during the excitation phase ($t=40$ ms). Reported below each panel are the minimum, maximum and step in mV of the displayed map.

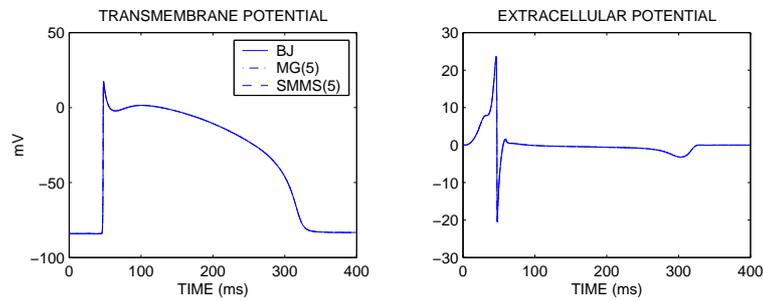


Fig. 3. Test 3. Time evolution at a fixed point of the transmembrane and extracellular potentials, computed with the three methods.

PCG iterations per time step, the average execution time per time step and the total simulation time. MG(5) and SMMS(5) are respectively 4 and 4.7 times faster than BJ. The detailed iteration counts as a function of time during the complete heartbeat are shown in Fig. 1 (left panel for BJ and right panel for MG(5) and SMMS(5)). Figure 2 shows the spatial maps of the transmembrane and extracellular potentials computed 40 ms after the stimulus was given at a vertex of the domain, i.e. during the excitation phase. Figure 3 shows the transmembrane and extracellular potentials computed in a fixed point of the domain by the three methods (the graphics are perfectly superimposed).

References

- [1] S. Balay et al. PETSc users manual. Technical Report ANL-95/11 - Revision 2.1.5, Argonne National Laboratory, 2004.
- [2] P. Colli Franzone and L. F. Pavarino. A parallel solver for reaction-diffusion systems in computational electrocardiology. *Math. Models Methods Appl. Sci.*, 14(6):883–911, 2004.
- [3] P. Colli Franzone and G. Savaré. Degenerate evolution systems modeling the cardiac electric field at micro- and macroscopic level. In A. Lorenzi and B. Ruf, editors, *Evolution Equations, Semigroups and Functional Analysis*, pages 49–78. Birkhäuser, 2002.
- [4] I. J. Le Grice et al. Laminar structure of the heart: ventricular myocyte arrangement and connective tissue architecture in the dog. *Am. J. Physiol. (Heart Circ. Physiol.)*, 269(38):H571–H582, 1995.
- [5] C. Luo and Y. Rudy. A model of the ventricular cardiac action potential: depolarization, repolarization, and their interaction. *Circ. Res.*, 68(6):1501–1526, 1991.
- [6] M. Murillo and X.C. Cai. A fully implicit parallel algorithm for simulating the nonlinear electrical activity of the heart. *Numer. Lin. Alg. Appl.*, 11(2-3):261–277, 2004.
- [7] L. F. Pavarino and P. Colli Franzone. Parallel solution of cardiac reaction-diffusion models. In R. Kornhuber et al., editor, *Domain Decomposition Methods in Science and Engineering*, volume 40 of *Springer LNCSE*, pages 669–676. Springer-Verlag, 2004.
- [8] M. Pennacchio and V. Simoncini. Efficient algebraic solution of reaction-diffusion systems for the cardiac excitation process. *J. Comput. Appl. Math.*, 145(1):49–70, 2002.
- [9] K. Skouibine and W. Krassowska. Increasing the computational efficiency of a bidomain model of defibrillation using a time-dependent activating function. *Ann. Biomed. Eng.*, 28:772–780, 2000.
- [10] B.F. Smith, P.E. Bjørstad, and W.D. Gropp. *Domain Decomposition*. Cambridge University Press, Cambridge, 1996.
- [11] D. Streeter. Gross morphology and fiber geometry in the heart. In R. Berne, editor, *Handbook of Physiology*, volume 1 of *The Heart*, pages 61–112. Williams & Wilkins, Baltimore, 1979.
- [12] M. Veneroni. Reaction-diffusion systems for the macroscopic bidomain model of the cardiac electric field. Technical report, I.M.A.T.I.-C.N.R., 2006.

- [13] E. J. Vigmond, F. Aguel, and N. A. Trayanova. Computational techniques for solving the bidomain equations in three dimensions. *IEEE Trans. Biomed. Eng.*, 49(11):1260–1269, 2002.
- [14] R. Weber dos Santos. *Modelling Cardiac Electrophysiology*. PhD thesis, Univ. of Rio de Janeiro, Dept. of Mathematics, 2002.
- [15] R. Weber dos Santos, G. Plank, S. Bauer, and E. J. Vigmond. Parallel multigrid preconditioner for the cardiac bidomain model. *IEEE Trans. Biomed. Eng.*, 51(11):1960–1968, 2004.