SIAM J. SCI. COMPUT.
Vol. 44, No. 2, pp. B224–B249

PARALLEL NEWTON–KRYLOV BDDC AND FETI-DP DELUXE SOLVERS FOR IMPLICIT TIME DISCRETIZATIONS OF THE CARDIAC BIDOMAIN EQUATIONS*

NGOC MAI MONICA HUYNH†, LUCA F. PAVARINO†, AND SIMONE SCACCHI‡

Abstract. Two novel parallel Newton–Krylov balancing domain decomposition by constraints (BDDC) and dual-primal finite element tearing and interconnecting (FETI-DP) solvers with deluxe scaling are constructed, analyzed, and tested numerically for implicit time discretizations of the three-dimensional bidomain system of equations. This model represents the most advanced mathematical description of the cardiac bioelectrical activity, and it consists of a degenerate system of two nonlinear reaction-diffusion partial differential equations (PDEs), coupled with a stiff system of ordinary differential equations (ODEs). A finite element discretization in space and a segregated implicit discretization in time, based on decoupling the PDEs from the ODEs, yields at each time step the solution of a nonlinear algebraic system. The Jacobian linear system at each Newton iteration is solved by a Krylov method, accelerated by BDDC or FETI-DP preconditioners, both augmented with the recently introduced deluxe scaling of the dual variables. A polylogarithmic convergence rate bound is proven for the resulting parallel Bidomain solvers. Extensive numerical experiments on Linux clusters up to two thousand processors confirm the theoretical estimates, showing that the proposed parallel solvers are scalable and quasi-optimal.

Key words. domain decomposition, FETI-DP and BDDC preconditioners, deluxe scaling, bidomain system

AMS subject classifications. 65N55, 65M55, 65F10, 92C30

DOI. 10.1137/20M1353848

1. Introduction.

The aim of this work is to design, analyze and numerically test balancing domain decomposition by constraints (BDDC) and dual-primal finite element tearing and interconnecting (FETI-DP) algorithms for a preconditioned Newton–Krylov solver for implicit time discretizations of the bidomain system. This model describes the propagation of the electric impulse in cardiac tissue by means of a degenerate parabolic system of two nonlinear reaction-diffusion partial differential equations (PDEs), modeling the evolution of the transmembrane electric potential [10, 28, 29, 30]. These PDEs are coupled through the nonlinear reaction term with a system of ordinary differential equations (ODEs), deriving from a membrane model describing the ionic currents flowing through the cell membrane and the dynamics of the associated gating variables.

Several time discretizations have been proposed for these complex nonlinear cardiac models, employing different semi-implicit, operator splitting, and decoupling techniques; see [10, Chapter 7.2] for a review. The most popular time discretizations have been based on semi-implicit (see, e.g., [5, 8, 11, 38, 39]) and/or operator splitting schemes (see, e.g., [35, 6, 7]). These techniques have been largely preferred to

*Submitted to the journal's Computational Methods in Science and Engineering section July 17, 2020; accepted for publication (in revised form) October 1, 2021; published electronically March 9, 2022.

https://doi.org/10.1137/20M1353848

Funding: This work was supported by INdAM. The work of the first and second authors was supported by MIUR grant 2017AXL54F_002. The work of the third author was supported by MIUR grant 2017AXL54F_003.

†Dipartimento di Matematica, Università degli Studi di Pavia, 27100 Pavia, Italy (ngocmaimonica.huynh@unipv.it, luca.pavarino@unipv.it).

‡Dipartimento di Matematica, Università degli Studi di Milano, 20133 Milano, Italy (simone.scacchi@unimi.it).

fully implicit schemes such as the bidomain monolithic solver proposed in [25]. Fully implicit solvers can be computationally very expensive when the ionic model consists of very stiff and high-dimensional nonlinear systems of ODEs (e.g., [17, 36]). On the other hand, operator splitting and decoupling techniques introduce additional errors which increase the time finite difference errors; see [10, Chapter 7.2] and the references therein. As a possible compromise to balance computational effort and accuracy, we propose here a decoupling solution strategy based on a segregated implicit time discretization of our model, where at each time step we solve the ODEs system of the ionic model first, and then we solve and update the nonlinear system arising from the discretized bidomain equations. This strategy was studied previously in [23, 24, 34], where only overlapping one-level and multilevel Schwarz preconditioners were considered. In this paper, we extend this solution strategy to two of the most efficient preconditioners currently available, the FETI-DP and BDDC dual-primal preconditioners with deluxe scaling. We recall that a practical advantage of these methods with respect to other domain decomposition preconditioners such as multilevel Schwarz is that they can be easily extended to unstructured meshes, because they do not need the implementation of an intergrid operator.

FETI-DP methods were proposed by [14] as an alternative to one-level or two-level finite element tearing and interconnecting (FETI). They have been applied in several contexts, from three-dimensional elliptic problems with heterogeneous coefficients [18] to linear elasticity problems [19, 31]. In the biomechanics field, applications of FETI and FETI-DP have been extensively studied by [1, 3, 20, 38]. BDDC were introduced by [12] as an alternative to FETI-DP for scalar elliptic problems and then analyzed by [21, 22]. Among other applications, BDDC has been applied to the linearized (semi-implicit) bidomain system in [38, 39] and to cardiac mechanics in [8, 27]. FETI-DP and BDDC have been proven to be spectrally equivalent [16, 22]. The other main family of domain decomposition methods, based on the overlapping Schwarz method, has also been applied to the bidomain system; see [26, 33, 5] for semi-implicit time discretizations and [23, 24, 34] for segregated implicit time discretizations.

Our main contribution consists in a novel theoretical estimate for the condition number of the preconditioned operator, using the recent deluxe scaling introduced in [13], for the solution of the nonlinear system arising from a fully implicit discretization of the decoupled cardiac electrical model. We also present the results of extensive numerical tests confirming our optimality bound and the scalability of the proposed solver. Robustness of and computational equivalence between the proposed dual-primal algorithms are shown, thus encouraging further investigations with more complex realistic geometries and the development of monolithic solvers for electromechanical models.

The rest of the paper is structured as follows. In section 2, we briefly introduce the bidomain model describing the electrical activity in the cardiac tissue. In section 3, we formulate the space discretization and the time decoupling strategy. An overview of nonoverlapping domain decomposition spaces and objects follows, and an excursus of FETI-DP and BDDC preconditioners concludes section 4. The novel convergence rate estimate is proved in section 5, while the results of several parallel numerical tests in three dimensions are presented in section 6.

2. The cardiac bidomain model. We consider here the nonlinear parabolic reaction-diffusion system, describing the propagation of electrical signal through the cardiac tissue.

In the bidomain system [10, 28], the cardiac tissue is represented as two interpenetrating domains, the intra- and extracellular domains, coexisting at every point

of the cardiac tissue and connected by a distributed continuous cellular membrane, which (as the intra- and extracellular domains) fills the complete volume. Cardiac cells are arranged in fibers set as laminar sheets running counterclockwise from the epicardium to the endocardium (see [15, 10] for more details). In this way, at each point $\mathbf{x}$ of the cardiac domain $\Omega \subset \mathbb{R}^3$ it is possible to define an orthonormal triplet of vectors $\mathbf{a}_l(\mathbf{x})$, $\mathbf{a}_t(\mathbf{x})$, and $\mathbf{a}_n(\mathbf{x})$ parallel to the local fiber direction and tangent and orthogonal to the laminar sheets, respectively.

We define the conductivity tensors of the two media as

$$D_{i,e}(\mathbf{x}) = \sum_{\bullet=\{l,t,n\}} \sigma_{\bullet}^{i,e} \mathbf{a}_{\bullet}(\mathbf{x}),$$

where $\sigma_{\bullet}^{i,e}$ are conductivity coefficients in the intra- and extracellular domain along the corresponding direction $\mathbf{a}_{\bullet}$, with $\bullet = l, t, n$. In our analysis, we assume here that these coefficients are constant in space. By defining $u_{i,e} : \Omega \times (0,T) \to \mathbb{R}$ as the intra- and extracellular electric potential, $v = u_i - u_e$ as the transmembrane potential, and $w : \Omega \times (0,T) \to \mathbb{R}^{N_w}$ the gating and ionic concentration variables, the parabolic-parabolic formulation of the Bidomain system reads as follows:

$$\begin{cases} \chi C_m \dfrac{\partial v}{\partial t} - \operatorname{div}(D_i \cdot \nabla u_i) + I_{\text{ion}}(v,w) = 0 & \text{in } \Omega \times (0,T), \\ -\chi C_m \dfrac{\partial v}{\partial t} - \operatorname{div}(D_e \cdot \nabla u_e) - I_{\text{ion}}(v,w) = -I_{\text{app}}^e & \text{in } \Omega \times (0,T), \\ \dfrac{\partial w}{\partial t} - R(v,w) = 0, & \text{in } \Omega \times (0,T), \\ v(x,t) = u_i(x,t) - u_e(x,t) & \text{in } \Omega \times (0,T) \end{cases} \tag{2.1}$$

with initial values $v(\mathbf{x},0) = u_i(\mathbf{x},0) - u_e(\mathbf{x},0)$, $w(\mathbf{x},0) = w_0(\mathbf{x})$, where χ is the ratio of membrane area per tissue volume and C_m is the surface capacitance. Assuming that the heart is immersed in a nonconductive medium, we require zero-flux boundary conditions $\mathbf{n}^T D_{i,e} \nabla u_{i,e} = 0$ on $\partial\Omega \times (0,T)$ and compatibility condition $\int_\Omega I_{\text{app}}^e dx = 0$, where $I_{\text{app}}^e : \Omega \times (0,T) \to \mathbb{R}$ is the extracellular applied current. Existence, uniqueness, and regularity results for (2.1) have been extensively studied; see, for example, [9].

The ionic model describing the ionic currents flowing through the cell membrane is defined by the terms I_{ion} and $R(v,w)$ in (2.1). In this work, we consider the very simple (yet macroscopically reliable) Rogers–McCulloch ionic model [32], with only one gating variable and

$$I_{\text{ion}}(v,w) = G\, v \left(1 - \frac{v}{v_{th}}\right)\left(1 - \frac{v}{v_p}\right) + \eta_1 v w, \qquad R(v,w) = \eta_2 \left(\frac{v}{v_p} - w\right),$$

where G, v_{th}, v_p, η_1, and η_2 are constant coefficients.

3. Numerical methods. (a) Space discretization. The cardiac domain Ω is discretized by a structured quasi-uniform grid of hexahedral isoparametric Q_1 elements of maximal diameter h. Let $V_h \subset V$ be the associated finite element space with the same basis functions $\{\varphi_p\}_{p=1}^{N_h}$ for both variables $u_{i,e}$ and w, and let $A_{i,e}$ and M be the stiffness and mass matrices with entries

$$\{A_{i,e}\}_{nm} = \int_\Omega (\nabla \varphi_n)^T D_{i,e} \cdot \nabla \varphi_m dx, \qquad \{M\}_{nm} = \int_\Omega \varphi_n \varphi_m dx.$$

For our purposes, the mass matrix is obtained with the usual mass-lumping technique. We denote by $\mathbf{u}_{i,e}$, $\mathbf{v}$, $\mathbf{w}$, $\mathbf{I}_{\text{ion}}$, and $\mathbf{I}^{\mathbf{i,e}}_{\text{app}}$ the coefficient vectors from the discretization of $u_{i,e}$, v, w, I_{ion}, and $I^{i,e}_{\text{app}}$, respectively. With these choices, we thus need to solve the semidiscrete bidomain system

$$\begin{cases} \chi C_m \mathcal{M} \frac{\partial}{\partial t} \begin{bmatrix} \mathbf{u}_i \\ \mathbf{u}_e \end{bmatrix} + \mathcal{A} \begin{bmatrix} \mathbf{u}_i \\ \mathbf{u}_e \end{bmatrix} + \begin{bmatrix} M\ \mathbf{I}_{\text{ion}}(\mathbf{v},\mathbf{w}) \\ -M\ \mathbf{I}_{\text{ion}}(\mathbf{v},\mathbf{w}) \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ -M\ \mathbf{I}^{\mathbf{e}}_{\text{app}} \end{bmatrix}, \\ \frac{\partial \mathbf{w}}{\partial t} = R(\mathbf{v},\mathbf{w}), \end{cases} \tag{3.1}$$

where $\mathcal{A}$ and $\mathcal{M}$ are the stiffness and mass block-matrices

$$\mathcal{A} = \begin{bmatrix} A_i & 0 \\ 0 & A_e \end{bmatrix}, \quad \mathcal{M} = \begin{bmatrix} M & -M \\ -M & M \end{bmatrix}.$$

(b) Decoupled implicit time discretization. Fully implicit discretizations in time of the bidomain system coupled with ionic models while using physiological coefficients lead to the solution of nonlinear problems at each time step, which can be very expensive from a computational point of view: indeed, realistic ionic models are very complex and can present up to fifty nonlinear ODEs (see, e.g., [17, 36]). Few attempts in this direction have been made by using simple ionic models (e.g., [25]). In the literature, common alternatives consider implicit-explicit schemes and/or operator splitting, where the diffusion terms are treated separately from the reaction (e.g., [6, 7, 8, 38]); see also [10, Chapter 7.2] and the references therein.

Instead, we choose here to decouple the gating variable w from the intra- and extracellular potentials u_i and u_e as in [23]. At each time step, this decoupled backward Euler strategy consists of two more substeps.

- **Step 1**: Update gating and ionic variables. Given $\mathbf{u}^n_{i,e}$ (hence $\mathbf{v}^n$) at the previous time step t_n, compute $\mathbf{w}^{n+1}$ by solving the membrane model

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

  where $\tau = t_{n+1} - t_n$ is the current time step.
- **Step 2**: Solve the bidomain system. Given $\mathbf{u}^n_{i,e}$ at the previous time step and given $\mathbf{w}^{n+1}$, calculate $\mathbf{u}^{n+1} = (\mathbf{u}^{n+1}_i, \mathbf{u}^{n+1}_e)$ by solving the nonlinear equation $\mathbf{F}(\mathbf{u}^{n+1}) = \mathbf{G}$ derived from the backward Euler scheme applied to the bidomain equations, where

$$\mathbf{F}(\mathbf{u}^{n+1}) = (\chi C_m \mathcal{M} + \tau \mathcal{A}) \begin{bmatrix} \mathbf{u}^{n+1}_i \\ \mathbf{u}^{n+1}_e \end{bmatrix} + \tau \begin{bmatrix} M\mathbf{I}_{\text{ion}}(\mathbf{v}^{n+1}, \mathbf{w}^{n+1}) \\ -M\mathbf{I}_{\text{ion}}(\mathbf{v}^{n+1}, \mathbf{w}^{n+1}) \end{bmatrix},$$
$$\mathbf{G} = \chi C_m \mathcal{M} \begin{bmatrix} \mathbf{u}^n_i \\ \mathbf{u}^n_e \end{bmatrix} + \tau \begin{bmatrix} \mathbf{0} \\ -M\mathbf{I}^{\mathbf{e}}_{\text{app}} \end{bmatrix}.$$

The strategy proposed here for the solution of the nonlinear system in Step 2 consists in a Newton–Krylov approach, as in [23, 24, 34], where the decomposition of the problem is made after the linearization: a Newton scheme is applied as an outer iteration, and the Jacobian linear system arising at each Newton step is solved by a Krylov method with a dual-primal preconditioner.

(c) Newton scheme and properties of the Jacobian bilinear form. The outer Newton iteration can be summarized as

- choose a starting value $\mathbf{u}^0 = (\mathbf{u}_i^0, \mathbf{u}_e^0)$;
- for $k \geq 0$ solve the Jacobian linear system

$$\mathbf{J}\mathbf{F}^k \mathbf{s}^k = -\mathbf{F}(\mathbf{u}^k) \tag{3.2}$$

  until a Newton stopping criterion is satisfied, where $\mathbf{s}^k = (\mathbf{s}_i^k, \mathbf{s}_e^k)$ is the Newton correction at step k and $\mathbf{J}\mathbf{F}^k$ is the Jacobian of $\mathbf{F}$ computed in $\mathbf{u}^k$;
- update: $\mathbf{u}^{k+1} = \mathbf{u}^k + \mathbf{s}^k$.

For our theoretical purposes, we need to associate to problem (3.2) the bilinear form

$$a(s^k, \phi) := \chi C_m \left(s_i^k - s_e^k,\ \varphi_i - \varphi_e\right) + \tau a_i \left(s_i^k,\ \varphi_i\right) + \tau a_e \left(s_e^k,\ \varphi_e\right) + \tau \left(\sum_{l=1}^{N_h} \frac{\partial I_{\text{ion}}}{\partial v_l}(v^{k-1}) \left(s_{i,l}^k - s_{e,l}^k\right) \psi_l,\ \varphi_i - \varphi_e \right)$$

for all $s^k = \left(s_i^k, s_e^k\right) \in V_h \times V_h$ and $\phi = (\varphi_i, \varphi_e) \in V_h \times V_h$, being ψ_l the lth nodal basis function, where $a_{i,e}(\cdot, \cdot)$ are the bilinear forms associated with the diffusion terms and $(\cdot,\ \cdot)$ denotes the usual L^2-inner product. As in [23], it is possible to show that this bilinear form associated to the Jacobian linear system is continuous and coercive with respect to the following norm defined for all $u = (u_i, u_e) \in V_h \times V_h$:

$$|||u|||_\tau^2 := (1+\tau)||u_i - u_e||_{L^2(\Omega)}^2 + \tau a_i \left(u_i,\ u_i\right) + \tau a_e \left(u_e,\ u_e\right).$$

LEMMA 3.1. *Assume that*

$$\chi C_m + \tau \frac{\partial I_{ion}}{\partial v_l}(v^k) \geq c > 0, \qquad c \in \mathbb{R}^+,$$

holds for all $l = 1, \ldots, N_h$ and for all k. Then the bilinear form $a(\cdot, \cdot)$ is continuous and coercive with respect to the norm $|||\cdot|||_\tau$.

Remark 3.2. *We do observe that the above hypothesis of nonnegativity is always satisfied for any time step $\tau \leq 0.37$ ms, using the Rogers–McCulloch ionic model. Indeed, numerical computations of $\chi C_m + \tau \frac{\partial I_{ion}}{\partial v}$ validate this assumption (see Figure* 3.1*).*

As an immediate consequence of the continuity and coercivity of the bilinear form $a(\cdot, \cdot)$, the following lemma holds. We drop the index k from now on, unless an explicit ambiguity occurs.

LEMMA 3.3. *Assuming that the conductivity coefficients are constant in space, the bilinear form $a(\cdot, \cdot)$ satisfies the bounds*

$$a(s,s) \leq (\chi C_m + \tau K_M)\, ||s_i - s_e||_{L^2(\Omega)}^2 + \tau \sigma_M^i |s_i|_{H^1(\Omega)}^2 + \tau \sigma_M^e |s_e|_{H^1(\Omega)}^2,$$
$$a(s,s) \geq (\chi C_m + \tau K_m)\, ||s_i - s_e||_{L^2(\Omega)}^2 + \tau \sigma_m^i |s_i|_{H^1(\Omega)}^2 + \tau \sigma_m^e |s_e|_{H^1(\Omega)}^2,$$

where

$$\sigma_M^{i,e} = \max_{\bullet = \{l,t,n\}} \sigma_\bullet^{i,e}, \qquad \sigma_m^{i,e} = \min_{\bullet = \{l,t,n\}} \sigma_\bullet^{i,e},$$

and K_M, K_m independent from the subdomain diameter H and the mesh size h.

Remark 3.4. *This result is extensible to the case of conductivity coefficients almost constant over each subdomain.*

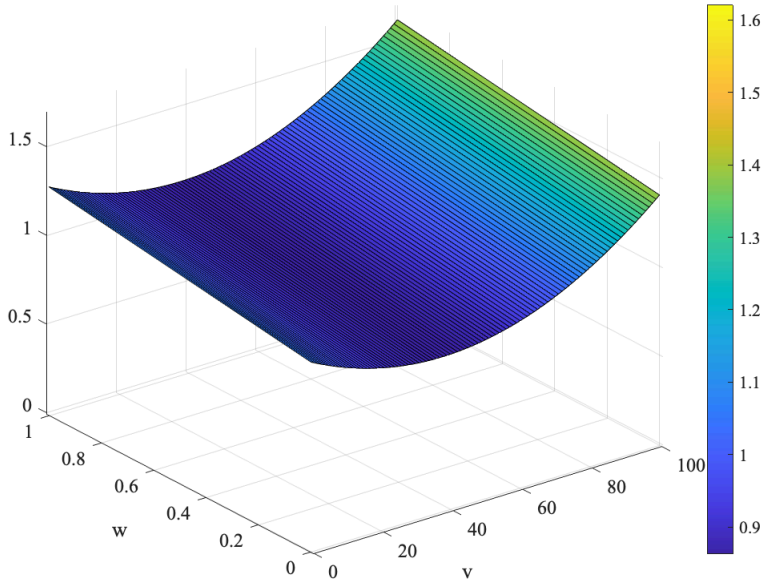

FIG. 3.1. *Surface plot of* $\chi C_m + \tau \frac{\partial I_{ion}}{\partial v}$ *with* $C_m = 1 \frac{mF}{cm^3}$, $\chi = 1$, *and* $\tau = 0.05$ *ms, which are values usually employed in numerical experiments.*

4. Dual-primal preconditioners for Newton–Krylov solvers.

4.1. Nonoverlapping dual-primal spaces. Let Ω_j, $j = 1, \ldots, N$, be a decomposition of the cardiac domain Ω, into nonoverlapping subdomains (or substructures). This decomposition forms a partition of Ω such that $\overline{\Omega} = \cup_{j=1}^{N} \overline{\Omega}_j$, $\Omega_j \cap \Omega_k = \emptyset$ if $j \neq k$ and the intersection of the boundaries $\partial\Omega_j \cap \partial\Omega_k$ is either empty, a vertex, an edge, or a face. Each subdomain is a union of shape-regular conforming finite elements. The interface Γ is the set of points that belong to at least two subdomains,

$$\Gamma := \bigcup_{j \neq k} \partial\Omega_j \cap \partial\Omega_k, \qquad \Gamma^{(j)} = \partial\Omega_j \cap \Gamma.$$

We assume that subdomains are shape regular and have a typical diameter of size H_j, whereas the finite elements are of diameter h; we denote by $H = \max_j H_j$. Let $W^{(j)} = V_h(\Omega_j) \times V_h(\Omega_j)$ be the associated local finite element spaces. We partition $W^{(j)}$ into the interior part $W_I^{(j)}$ and the finite element trace space $W_\Gamma^{(j)}$. Note that we consider variables on the Neumann boundaries $\partial\Omega_N$ as interior to a subdomain. We introduce the product spaces

$$W = \prod_{j=1}^{N} W^{(j)}, \qquad W_\Gamma := \prod_{j=1}^{N} W_\Gamma^{(j)}.$$

Therefore, we define $\widehat{W} \subset W$ as the subspace of functions of W, which are continuous in all interface variables between subdomains, and similarly we denote by $\widehat{W}_\Gamma \subset W_\Gamma$ the subspace formed by the continuous elements of W_Γ. In dual-primal methods, we iterate in the space W while requiring continuity constraints (also called *primal constraints*) to hold throughout the iterations. Primal constraints guarantee that each subdomain problem is invertible and that a good convergence bound can be obtained.

We denote by $\widetilde{W}$ the space of finite element functions in W, which are continuous in all primal variables; clearly we have $\widehat{W} \subset \widetilde{W} \subset W$ and likewise $\widehat{W}_\Gamma \subset \widetilde{W}_\Gamma \subset W_\Gamma$. Let $W_\Pi^{(j)} \subset W_\Gamma^{(j)}$ be the primal subspace of functions which are continuous across the interface and that will be subassembled between the subdomains sharing $\Gamma^{(j)}$. Additionally, let $W_\Delta^{(j)} \subset W_\Gamma^{(j)}$ contain the finite element functions (called dual) which can be discontinuous across the interface and vanish at the primal degrees of freedom.

We then introduce the product subspaces $W_\Pi = \prod_{j=1}^N W_\Pi^{(j)}$ and $W_\Delta = \prod_{j=1}^N W_\Delta^{(j)}$, from which $W_\Gamma = W_\Pi \oplus W_\Delta$. Using this notation, we can decompose $\widetilde{W}_\Gamma$ into a primal subspace $\widehat{W}_\Pi$ which has continuous elements only and a dual subspace W_Δ which contains finite element functions which are not continuous; i.e., we have $\widetilde{W}_\Gamma = \widehat{W}_\Pi \oplus W_\Delta$. In this work, we will denote with subscripts I, Δ, and Π the interior, dual, and primal variables, respectively.

In the nonoverlapping framework, the global system matrix (see in this application (3.2)) is never formed explicitly, but a local version with the same structure is assembled on each subdomain, by restricting the integration set from Ω to Ω_j by defining the local bilinear forms

$$a^{(j)}(s,\phi) = \chi C_m \left(s_i - s_e,\ \varphi_i - \varphi_e\right)_{|\Omega_j} + \tau a_i^{(j)}\left(s_i,\ \varphi_i\right) + \tau a_e^{(j)}\left(s_e,\ \varphi_e\right) + \tau \left(\sum_{l=1}^{N_h} \frac{\partial I_{\text{ion}}}{\partial v_l}(v) \left(s_{i,l} - s_{e,l}\right) \varphi_l,\ \varphi_i - \varphi_e \right)_{|\Omega_j},$$

where $(\cdot\ ,\ \cdot)_{|\Omega_j}$ denotes the restriction of the L^2-inner product on the jth subdomain. These definitions are admissible here, as the proposed theory allows constant nonnegative distribution of the conductivity coefficients among all subdomains, with large jumps aligned to the interfaces.

The reordering of the degrees of freedom leads to a reordered system matrix: assuming that the system (3.2) can be written as $\mathcal{K}u = f$, then we will have

$$\mathcal{K}^{(j)} = \begin{bmatrix} K_{II}^{(j)} & K_{I\Gamma}^{(j)} \\ K_{I\Gamma}^{(j)T} & K_{\Gamma\Gamma}^{(j)} \end{bmatrix}, \qquad \mathcal{K} = \begin{bmatrix} K_{II} & K_{I\Gamma} \\ K_{I\Gamma}^T & K_{\Gamma\Gamma} \end{bmatrix}.$$

As in classical iterative substructuring, we apply the so-called *static condensation*, which consists in eliminating the interior degrees of freedom, thus reducing the problem to one on the interface $\Gamma = \cup_{j=1}^N \partial\Omega_j \backslash \partial\Omega$. The local Schur complement systems are

$$S_\Gamma^{(j)} = K_{\Gamma\Gamma}^{(j)} - K_{I\Gamma}^{(j)T} K_{II}^{(j)-1} K_{I\Gamma}^{(j)}.$$

By defining the unassembled Schur complement matrix $S_\Gamma = \text{diag}\,[S_\Gamma^{(1)}, \ldots, S_\Gamma^{(N)}]$, we obtain the global Schur complement matrix $\widehat{S}_\Gamma = R_\Gamma^T S_\Gamma R_\Gamma^T$, where R_Γ is the direct sum of local restriction operators $R_\Gamma^{(j)}$ returning the local interface components. Thus, instead of solving system (3.2), we solve the Schur complement system

$$\widehat{S}_\Gamma u_\Gamma = \widehat{f}_\Gamma, \tag{4.1}$$

where $\widehat{f}_\Gamma$ is retrieved from the right-hand side of (3.2). Once this problem is solved, the solution u_Γ on the interface is used to recover the solution on the internal degrees of freedom $u_I = K_{II}^{-1}\left(f_I - K_{I\Gamma} u_\Gamma\right)$. The Schur complement matrix $\widehat{S}_\Gamma$ of the Jacobian

bidomain system (3.2) is symmetric, positive semidefinite; hence it is possible to apply the preconditioned conjugate gradient (PCG) method.

We define the Jacobian Bidomain local discrete harmonic extension operators as follows:

$$\mathcal{H}_j : W_\Gamma^{(j)} \longrightarrow W^{(j)}, \qquad \mathcal{H}_j w_\Gamma^{(j)} = \begin{cases} -K_{II}^{(j)-1} K_{I\Gamma}^{(j)} w_\Gamma^{(j)} & \text{on } W_I^{(j)}, \\ w_\Gamma^{(j)} & \text{on } W_\Gamma^{(j)}. \end{cases}$$

We note that the local discrete harmonic extension of a constant vector is the vector itself. From now on, we will use the componentwise notation

$$\mathcal{H}_j w_\Gamma^{(j)} = \left(\mathcal{H}_j^i w_\Gamma^{(j)}, \mathcal{H}_j^e w_\Gamma^{(j)} \right),$$

where the superscripts i, e denote the usual intra- and extracellular components. As for standard elliptic problems (see [37]), the Schur bilinear form can be defined through the action of the Schur matrix and the Jacobian bilinear form

$$a^{(j)}(\mathcal{H}_j u_\Gamma^{(j)}, \mathcal{H}_j v_\Gamma^{(j)}) = v_\Gamma^{(j)} S_\Gamma^{(j)} u_\Gamma^{(j)} = s^{(j)}(u_\Gamma^{(j)}, v_\Gamma^{(j)}).$$

From the definition of $S_\Gamma^{(j)}$, it follows immediately that the bilinear form $s^{(j)}(\cdot,\cdot)$ is symmetric and coercive. Thanks to Lemma 3.3, it is possible to bound the energies related to the local Schur complements,

$$s^{(j)}(u_\Gamma^{(j)}, u_\Gamma^{(j)}) = \min_{v^{(j)}_{|\partial\Omega_j \cap \Gamma} = u_\Gamma^{(j)}} a^{(j)}(\mathcal{H}_j u_\Gamma^{(j)}, \mathcal{H}_j u_\Gamma^{(j)}), \tag{4.2}$$

which allows us to work with discrete harmonic extensions instead of functions defined only on Γ.

4.2. Restriction operators and scaling. We define the restriction operators

$$R_\Delta^{(j)} : W_\Delta \to W_\Delta^{(j)}, \qquad R_{\Gamma\Delta} : W_\Gamma \to W_\Delta,$$
$$R_\Pi^{(j)} : \widehat{W}_\Pi \to W_\Pi^{(j)}, \qquad R_{\Gamma\Pi} : W_\Gamma \to \widehat{W}_\Pi$$

and the direct sums $R_\Delta = \oplus R_\Delta^{(j)}$, $R_\Pi = \oplus R_\Pi^{(j)}$, and $\widetilde{R}_\Gamma = R_{\Gamma\Pi} \oplus R_{\Gamma\Delta}$, which maps W_Γ into $\widetilde{W}_\Gamma$. We also need a proper scaling of the dual variables.

***ρ*-scaling.** Originally proposed for Neumann–Neumann methods, the ρ-scaling is defined for the bidomain model at each node $x \in \Gamma^{(j)}$ as

$$\delta_j^{i,e\,\dagger}(x) = \frac{\sigma_M^{i,e^{(j)}}}{\sum_{k \in \mathcal{N}_x} \sigma_M^{i,e^{(k)}}}, \qquad \sigma_M^{i,e^{(j)}} = \max_{\bullet = \{l,t,n\}} \sigma_\bullet^{i,e^{(j)}}, \tag{4.3}$$

where $\mathcal{N}_x$ is the set of indices of all subdomains with x in the closure of the subdomain. If x is in the interior of a subdomain, then $\mathcal{N}_x$ contains only the index of that subdomain. Moreover, $\mathcal{N}_x$ induces the definition of an equivalence relation that allows the classification of interface degrees of freedom into faces, edges, and vertices equivalence classes.

Deluxe scaling. Recently introduced in [13] and studied in [4], the deluxe scaling computes the average $\bar{w} = E_D w$ for each face $\mathcal{F}$ or edge $\mathcal{E}$ equivalence class. Suppose that $\mathcal{F}$ is shared by subdomains Ω_j and Ω_k. Let $S_\mathcal{F}^{(j)}$ and $S_\mathcal{F}^{(k)}$ be the principal minors

obtained from $S_\Gamma^{(j)}$ and $S_\Gamma^{(k)}$ by removing all the contributions that are not related to the degrees of freedom of the face $\mathcal{F}$. Let $u_{j,\mathcal{F}} = R_{\mathcal{F}} u_j$ be the restriction of u_j to the face $\mathcal{F}$ through the restriction operator $R_{\mathcal{F}}$. The deluxe average across $\mathcal{F}$ is then defined as

$$\bar{u}_{\mathcal{F}} = \left(S_{\mathcal{F}}^{(j)} + S_{\mathcal{F}}^{(k)}\right)^{-1} \left(S_{\mathcal{F}}^{(j)} u_{j,\mathcal{F}} + S_{\mathcal{F}}^{(k)} u_{k,\mathcal{F}}\right).$$

The action of $(S_{\mathcal{F}}^{(j)} + S_{\mathcal{F}}^{(k)})^{-1}$ can be computed by solving a Dirichlet problem over the two subdomains involved by extending to zero the right-hand-side entries that correspond with the interior nodes.

If we consider an edge $\mathcal{E}$ instead, the deluxe average across $\mathcal{E}$ is defined in a similar manner. Suppose for simplicity that $\mathcal{E}$ is shared by only three subdomains with indices j_1, j_2, and j_3; the extension to more than three subdomains is immediate. Let $u_{j,\mathcal{E}} = R_{\mathcal{E}} u_j$ be the restriction of u_j to the edge $\mathcal{E}$ through the restriction operator $R_{\mathcal{E}}$, and define $S_{\mathcal{E}}^{(j_{123})} = S_{\mathcal{E}}^{(j_1)} + S_{\mathcal{E}}^{(j_2)} + S_{\mathcal{E}}^{(j_3)}$; the deluxe average across an edge $\mathcal{E}$ is given by

$$\bar{u}_{\mathcal{E}} = \left(S_{\mathcal{E}}^{(j_{123})}\right)^{-1} \left(S_{\mathcal{E}}^{(j_1)} u_{j_1,\mathcal{E}} + S_{\mathcal{E}}^{(j_2)} u_{j_2,\mathcal{E}} + S_{\mathcal{E}}^{(j_3)} u_{j_3,\mathcal{E}}\right).$$

The relevant equivalence classes, involving the substructure Ω_j, will contribute to the values of $\bar{u}$. These contributions will belong to $\widehat{W}_\Gamma$ after being extended by zero to $\Gamma\backslash\mathcal{F}$ or $\Gamma\backslash\mathcal{E}$; the sum of all contributions will result in $R_*^T \bar{u}_*$. We then add the contributions from the different equivalence classes to obtain

$$\bar{u} = E_D u = u_\Pi + \sum_{*=\{\mathcal{F},\mathcal{E}\}} R_*^T \bar{u}_*,$$

where E_D is a projection. Its complementary projection is given by

$$P_D u := (I - E_D)u = u_\Delta - \sum_{*=\{\mathcal{F},\mathcal{E}\}} R_*^T \bar{u}_*. \tag{4.4}$$

For each subdomain Ω_j we define the scaling matrix

$$D^{(j)} = \begin{bmatrix} D_{*_{k_1}}^{(j)} & & \\ & \ddots & \\ & & D_{*_{k_j}}^{(j)} \end{bmatrix}, \qquad * = \{\mathcal{F}, \mathcal{E}\} \tag{4.5}$$

with $k_1, \ldots, k_j \in \Xi_j^*$ set containing the indices of the subdomains that share the face $\mathcal{F}$ or the edge $\mathcal{E}$ and where the diagonal blocks are given by $D_{\mathcal{F}}^{(j)} = (S_{\mathcal{F}}^{(j)} + S_{\mathcal{F}}^{(k)})^{-1} S_{\mathcal{F}}^{(j)}$ or $D_{\mathcal{E}}^{(j)} = (S_{\mathcal{E}}^{(j_1)} + S_{\mathcal{E}}^{(j_2)} + S_{\mathcal{E}}^{(j_3)})^{-1} S_{\mathcal{E}}^{(j_1)}$.

We can now define the scaled local restriction operators

$$R_{D,\Gamma}^{(j)} = D^{(j)} R_\Gamma^{(j)}, \qquad R_{D,\Delta}^{(j)} = R_{\Gamma\Delta}^{(j)} R_{D,\Gamma}^{(j)},$$

$R_{D,\Delta}$ as direct sum of $R_{D,\Delta}^{(j)}$ and the global scaled operator $\widetilde{R}_{D,\Gamma} = R_{\Gamma\Pi} \oplus R_{D,\Delta} R_{\Gamma\Delta}$.

4.3. FETI-DP preconditioner. The FETI-DP preconditioner was first proposed in [14] as an alternative to one-level and two-level FETI. This class of methods is based on the transposition from the Schur problem (4.1) on $\widehat{W}_\Gamma$ to a minimization problem on $\widetilde{W}_\Gamma$, with continuity constraints on the dual degrees of freedom: find $w_\Gamma \in \widetilde{W}_\Gamma$ which minimizes

$$\begin{cases} \frac{1}{2} w_\Gamma^T \widetilde{S}_\Gamma w_\Gamma - w_\Gamma^T \tilde{f}_\Gamma, \\ B w_\Gamma = 0, \end{cases}$$

where $\tilde{f}_\Gamma = \widetilde{R}_\Gamma^T f_\Gamma$ is given by partially subassembling the Schur complement right-hand side on primal nodes and B is the jump operator with entries $0, \pm 1$. The second part of the system $Bw_\Gamma = 0$ holds if and only if $w_\Gamma \in \widehat{W}$, which means that the columns of B related to primal degrees of freedom are null. By introducing a set of Lagrange multipliers $\lambda \in \Lambda = range(B)$, it is possible to formulate the minimization problem as a saddle point system:

$$\begin{bmatrix} \widetilde{S}_\Gamma & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} w_\Gamma \\ \lambda \end{bmatrix} = \begin{bmatrix} \tilde{f}_\Gamma \\ 0 \end{bmatrix}.$$

As $\widetilde{S}_\Gamma$ is invertible on $range(B)$, the degrees of freedom in $\widetilde{W}_\Gamma$ can be eliminated by a block-Cholesky factorization, reducing the above system to a problem only in the Lagrange multiplier unknowns:

$$F\lambda = d, \qquad \text{where} \quad F = B\widetilde{S}_\Gamma^{-1} B^T, \qquad d = B\widetilde{S}_\Gamma^{-1} \tilde{f}_\Gamma. \tag{4.6}$$

After the solution λ is found, we can retrieve the solution on $\widetilde{W}_\Gamma$ as $w_\Gamma = \widetilde{S}_\Gamma^{-1}(\tilde{f}_\Gamma - B^T\lambda)$. The FETI-DP system (4.6) in our application is symmetric; thus the PCG method works well. In order to ensure fast convergence, a quasi-optimal preconditioner is given by

$$M^{-1}_{\text{FETI-DP}} = B_D \widetilde{S}_\Gamma B_D^T,$$

where B_D is the scaled jump operator, obtained by applying $D^{(j)} : \Lambda \to \Lambda$ scaling matrices that act on the space of the Lagrange multipliers, which are given by (4.5) if the deluxe scaling is used or by the pseudoinverses (4.3) if the standard ρ-scaling is used.

4.4. BDDC preconditioner. BDDC is a two-level preconditioner for the Schur complement system $\widehat{S}_\Gamma u_\Gamma = \widehat{f}_\Gamma$. If we partition the degrees of freedom of the interface Γ into those internal (I) and those dual (Δ), the matrix $\mathcal{K}^{(j)}$ from the problem $\mathcal{K}u = f$ can be written as

$$\mathcal{K}^{(j)} = \begin{bmatrix} K_{II}^{(j)} & K_{I\Gamma}^{(j)} \\ K_{I\Gamma}^{(j)T} & K_{\Gamma\Gamma}^{(j)} \end{bmatrix} = \begin{bmatrix} K_{II}^{(j)} & K_{I\Delta}^{(j)} & K_{I\Pi}^{(j)} \\ K_{I\Delta}^{(j)T} & K_{\Delta\Delta}^{(j)} & K_{\Delta\Pi}^{(j)} \\ K_{I\Pi}^{(j)T} & K_{\Delta\Pi}^{(j)T} & K_{\Pi\Pi}^{(j)} \end{bmatrix}.$$

It is possible to define the BDDC preconditioner using the restriction operators as

$$M^{-1}_{\text{BDDC}} = \widetilde{R}_{D,\Gamma}^T \widetilde{S}_\Gamma^{-1} \widetilde{R}_{D,\Gamma}, \qquad \widetilde{S}_\Gamma = \widetilde{R}_\Gamma S_\Gamma \widetilde{R}_\Gamma^T,$$

where the action of the inverse of $\widetilde{S}_\Gamma$ can be evaluated with a block-Cholesky elimination procedure

$$\widetilde{S}_\Gamma^{-1} = \widetilde{R}_{\Gamma\Delta}^T \left(\sum_{j=1}^N \begin{bmatrix} 0 & R_\Delta^{(j)T} \end{bmatrix} \begin{bmatrix} K_{II}^{(j)} & K_{I\Delta}^{(j)} \\ K_{I\Delta}^{(j)T} & K_{\Delta\Delta}^{(j)} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ R_\Delta^{(j)} \end{bmatrix} \right) \widetilde{R}_{\Gamma\Delta} + \Phi S_{\Pi\Pi}^{-1} \Phi.$$

In this way, the first term is the sum of local solvers on each substructure Ω_j, while the latter is a coarse solver for the primal variables where

$$\Phi = R_{\Gamma\Pi}^T - R_{\Gamma\Delta}^T \sum_{j=1}^N \begin{bmatrix} 0 & R_\Delta^{(j)T} \end{bmatrix} \begin{bmatrix} K_{II}^{(j)} & K_{I\Delta}^{(j)} \\ K_{I\Delta}^{(j)T} & K_{\Delta\Delta}^{(j)} \end{bmatrix}^{-1} \begin{bmatrix} K_{I\Pi}^{(j)} \\ R_{\Delta\Pi}^{(j)} \end{bmatrix} R_\Pi^{(j)},$$

$$S_{\Pi\Pi} = \sum_{j=1}^N R_\Pi^{(j)T} \left(K_{\Pi\Pi}^{(j)} - \begin{bmatrix} K_{I\Pi}^{(j)T} & K_{\Delta\Pi}^{(j)T} \end{bmatrix} \begin{bmatrix} K_{II}^{(j)} & K_{I\Delta}^{(j)} \\ K_{I\Delta}^{(j)T} & K_{\Delta\Delta}^{(j)} \end{bmatrix}^{-1} \begin{bmatrix} K_{I\Pi}^{(j)} \\ R_{\Delta\Pi}^{(j)} \end{bmatrix} \right) R_\Pi^{(j)}$$

are the matrix which maps the primal degrees of freedom to the interface variables and the primal problem, respectively.

5. Convergence rate estimate. It has been proven that FETI-DP and BDDC methods are spectrally equivalent [16]. In this perspective, we are able to prove a convergence rate estimate for the preconditioned operator, which holds for both methods when the same coarse space is chosen. We observe that, as in most of the convergence bounds for FETI-DP and BDDC operators, also in this application the condition number is independent of the number of subdomains. We first recall some useful technical results that will be employed in the proof of the convergence rate estimate. These results can be found in Appendix A of [37].

THEOREM 5.1 (trace theorem). *Let Ω_j be a polyhedral domain, and define the discrete harmonic extension of the Laplacian operator $\mathcal{H}_j^\Delta u_\Gamma$ on Ω_j as*

$$u = \mathcal{H}_j^\Delta u_\Gamma \Leftrightarrow \begin{cases} -\Delta u = 0 & \text{in } \Omega_j, \\ u = u_\Gamma & \text{on } \Gamma^{(j)}, \\ u = 0 & \text{on } \partial\Omega_j \backslash \Gamma^{(j)}. \end{cases}$$

Then,

$$|u|^2_{H^{1/2}(\Gamma^{(j)})} \sim |\mathcal{H}_j^\Delta u_\Gamma|^2_{H^1(\Omega_j)}.$$

PROPOSITION 5.2 (Poincaré–Friedrichs inequality). *Let Ω be Lipschitz continuous with diameter H. Then, there exists a constant C_f, which depends only on the shape of Ω but not on its size, such that*

$$||u||^2_{L^2(\Omega)} \leq C_f H^2 |u|^2_{H^1(\Omega)}$$

for all $u \in H^1(\Omega)$ with vanishing mean value on Ω.

We will write with $A \lesssim B$ whenever $A \leq cB$ with c constant independent from the diameter H, the mesh size h, the time step τ, and the conductivity coefficients; similarly, we will write $A \sim B$ whenever $A \lesssim B$ and $B \lesssim A$. The main result of the paper is the following optimality bound.

THEOREM 5.3. *If the deluxe scaling is used, the condition numbers of the FETI-DP and BDDC preconditioned operators satisfy*

$$cond\ (P^{-1}Q) \leq \max_{\substack{k=1,\dots,N \\ \star=i,e}} \frac{\tau\sigma_M^{\star(k)} + H^2\left(\chi C_m + \tau K_M\right)}{\sigma_m^{\star(k)}} \left(1 + \log\left(\frac{H}{h}\right)\right)^3, \tag{5.1}$$

where $\sigma_M^{i,e} = \max_{\bullet=\{l,t,n\}} \sigma_\bullet^{i,e}$, $\sigma_m^{i,e} = \min_{\bullet=\{l,t,n\}} \sigma_\bullet^{i,e}$, P^{-1} *denotes the FETI-DP or BDDC preconditioner,* Q *is the system matrix, and* K_M *is a constant independent of the subdomain diameter* H *and mesh size* h.

The core of the proof relies on the following lemma.

LEMMA 5.4. *Let* $\chi C_m + \tau \frac{\partial I_{ion}}{\partial v_l}(v^k) \geq c > 0$ *with* $c \in \mathbb{R}^+$. *Let the primal set be spanned by the vertex nodal finite element functions and the subdomain edges averages. If the projection operator is scaled by the deluxe scaling, then, for all* $u \in \widetilde{W}_\Gamma$,

$$|E_D u|^2_{S_\Gamma} \lesssim \max_{\substack{k=1,\dots,N \\ \star=i,e}} \frac{\tau\sigma_M^{\star(k)} + H^2\left(\chi C_m + \tau K_M\right)}{\tau\sigma_m^{\star(k)}} \left(1 + \log\frac{H}{h}\right)^3 |u|^2_{S_\Gamma},$$

where $\sigma_M^{i,e} = \max_{\bullet=\{l,t,n\}} \sigma_\bullet^{i,e}$, $\sigma_m^{i,e} = \min_{\bullet=\{l,t,n\}} \sigma_\bullet^{i,e}$, *and* K_M *is a constant independent of the subdomain diameter* H *and mesh size* h.

Proof. Instead of proving the bound for the projection operator E_D, we prove it for the complementary projection P_D (4.4). Moreover, it is sufficient to compute only the local bounds, as it holds that $|P_D u|^2_{S_\Gamma} = \sum_{j=1}^N |R_{\partial\Omega_j} P_D u|^2_{S_\Gamma^{(j)}}$. Thus, for all $u \in \widetilde{W}_\Gamma$,

$$|R_{\partial\Omega_j} P_D u|^2_{S_\Gamma^{(j)}} \leq |\Xi_j^*| \sum_{\substack{*=\{\mathcal{F},\mathcal{E}\} \\ *\in\Xi_j^*}} |R_*^T \left(u_{j,*}^{i,e} - \bar{u}_*^{i,e}\right)|^2_{S_\Gamma^{(j)}},$$

where Ξ_j^* is the index set containing the indices of the subdomains that share the face $\mathcal{F}$ or the edge $\mathcal{E}$. Let us distinguish between face and edge contributions.

Face contributions. Suppose that the face $\mathcal{F}$ is shared by subdomains Ω_j and Ω_k. Then, by simple algebra, it follows that

$$u_{j,\mathcal{F}}^{i,e} - \bar{u}_{\mathcal{F}}^{i,e} = \left(S_{\mathcal{F}}^{(j)} + S_{\mathcal{F}}^{(k)}\right)^{-1} S_{\mathcal{F}}^{(k)} \left[\left(u_{j,\mathcal{F}}^{i,e} - \bar{u}_{j,\mathcal{F}}^{i,e}\right) - \left(u_{k,\mathcal{F}}^{i,e} - \bar{u}_{k,\mathcal{F}}^{i,e}\right) + \left(\bar{u}_{j,\mathcal{F}}^{i,e} - \bar{u}_{k,\mathcal{F}}^{i,e}\right)\right],$$

where we add and subtract the mean value $\bar{u}_{\cdot,\mathcal{F}}^{i,e}$ of u over $\mathcal{F}$ on the subdomain $\cdot$. Therefore, by noticing that $R_{\mathcal{F}} S_\Gamma^{(j)} R_{\mathcal{F}}^T = S_{\mathcal{F}}^{(j)}$, it follows that

$$\begin{aligned}\left|R_{\mathcal{F}}^T \left(u_{j,\mathcal{F}}^{i,e} - \bar{u}_{\mathcal{F}}^{i,e}\right)\right|^2_{S_\Gamma^{(j)}} \leq\ & 2|u_{j,\mathcal{F}}^{i,e} - \bar{u}_{j,\mathcal{F}}^{i,e}|^2_{S_{\mathcal{F}}^{(j)}} + 2|u_{k,\mathcal{F}}^{i,e} - \bar{u}_{k,\mathcal{F}}^{i,e}|^2_{S_{\mathcal{F}}^{(k)}} \\ & + \left|\left(S_{\mathcal{F}}^{(j)} + S_{\mathcal{F}}^{(k)}\right)^{-1} S_{\mathcal{F}}^{(k)} \left(\bar{u}_{j,\mathcal{F}}^{i,e} - \bar{u}_{k,\mathcal{F}}^{i,e}\right)\right|^2_{S_{\mathcal{F}}^{(j)}}\end{aligned}$$

For all $u_{j,\mathcal{F}}^{i,e} \in \widetilde{W}_\Gamma$, where we take advantage of two inequalities arising from the generalized eigenvalue problem $S_{\mathcal{F}}^{(j)}\phi = \lambda S_{\mathcal{F}}^{(k)}\phi$ and by observing that all eigenvalues are strictly positive.

It is sufficient now to estimate $|u_{j,\mathcal{F}}^{i,e} - \bar{u}_{j,\mathcal{F}}^{i,e}|^2_{S_{\mathcal{F}}^{(j)}}$ and

$$\left|\left(S_{\mathcal{F}}^{(j)} + S_{\mathcal{F}}^{(k)}\right)^{-1} S_{\mathcal{F}}^{(k)} \left(\bar{u}_{j,\mathcal{F}}^{i,e} - \bar{u}_{k,\mathcal{F}}^{i,e}\right)\right|^2_{S_{\mathcal{F}}^{(j)}};$$

we highlight that in case the subdomain face averages are also included in the primal space, the latter is zero. Starting from the first term, we make use of the bilinear form associated to the Schur complement

$$
\begin{aligned}
|u_{j,\mathcal{F}}^{i,e} - \bar{u}_{j,\mathcal{F}}^{i,e}|_{S_{\mathcal{F}}^{(j)}}^2 &\leq (\chi C_m + \tau K_M)\, ||\mathcal{H}_j^{\Delta}\left(u_j - \bar{u}_{j,\mathcal{F}}^{i,e}\right)||_{L^2(\Omega_j)}^2 \\
&\quad + \tau \sum_{\star=i,e} \sigma_M^{\star} |\mathcal{H}_j^{\Delta}\left(u_{j,\mathcal{F}} - \bar{u}_{j,\mathcal{F}}^{\star}\right)|_{H^1(\Omega_j)}^2 \\
&\lesssim \left[\tau\sigma_M^{i,e} + H^2\left(\chi C_m + \tau K_M\right)\right] |u_{j,\mathcal{F}} - \bar{u}_{j,\mathcal{F}}^{i,e}|_{H^{1/2}(\Gamma^{(j)})}^2,
\end{aligned}
$$

applying the ellipticity Lemma 3.3 and the Poincaré–Friedrichs inequality (Proposition 5.2) combined with the trace theorem, Theorem 5.1. As we are already taking in consideration the discrete restriction of a function $u_j \in \widetilde{W}_{\Gamma}$ on the face $\mathcal{F}$, the notations $u_j - \bar{u}_{j,\mathcal{F}}^{i,e}$ and $I^h(\Theta_{\mathcal{F}}(u_j - \bar{u}_{j,\mathcal{F}}^{i,e}))$ are essentially the same. Therefore, it is possible to apply [37, Lemma 4.26] and the trace theorem to get

$$
|u_{j,\mathcal{F}}^{i,e} - \bar{u}_{j,\mathcal{F}}^{i,e}|_{S_{\Gamma}^{(j)}}^2 \lesssim \left[\tau\sigma_M^{i,e} + H^2\left(\chi C_m + \tau K_M\right)\right] \left(1 + \log\frac{H}{h}\right)^2 |\mathcal{H}_j^{\Delta} u_j|_{H^1(\Omega_j)}^2.
$$

Regarding the second term, let $\mathcal{E} \subset \partial\mathcal{F}$ be a primal edge such that we can add and subtract $\bar{u}_{j,\mathcal{E}}^{i,e} = \bar{u}_{k,\mathcal{E}}^{i,e}$. Then, using the same inequalities from the generalized eigenvalue problem,

$$
|\left(S_{\mathcal{F}}^{(j)} + S_{\mathcal{F}}^{(k)}\right)^{-1} S_{\mathcal{F}}^{(k)} \left(\bar{u}_{j,\mathcal{F}}^{i,e} - \bar{u}_{k,\mathcal{F}}^{i,e}\right)|_{S_{\mathcal{F}}^{(j)}}^2 \leq 2\, |\bar{u}_{j,\mathcal{E}}^{i,e} - \bar{u}_{j,\mathcal{F}}^{i,e}|_{S_{\mathcal{F}}^{(j)}}^2 + 2\, |\bar{u}_{k,\mathcal{E}}^{i,e} - \bar{u}_{k,\mathcal{F}}^{i,e}|_{S_{\mathcal{F}}^{(k)}}.
$$

It is sufficient now to esteem the first term on the right-hand side, as we can deal with the other in the same fashion. Combining the result of ellipticity (Lemma 3.3), the Poincaré–Friedrichs inequality, and the trace theorem, we get

$$
\begin{aligned}
|\bar{u}_{j,\mathcal{E}}^{i,e} - \bar{u}_{j,\mathcal{F}}^{i,e}|_{S_{\mathcal{F}}^{(j)}}^2 &= a^{(j)}\left(\mathcal{H}_j^{i,e}\left(\bar{u}_{j,\mathcal{E}}^{i,e} - \bar{u}_{j,\mathcal{F}}^{i,e}\right),\ \mathcal{H}_j^{i,e}\left(\bar{u}_{j,\mathcal{E}}^{i,e} - \bar{u}_{j,\mathcal{F}}^{i,e}\right)\right) \\
&= a^{(j)}\left(\mathcal{H}_j^{i,e}\overline{\left(u_j - \bar{u}_{j,\mathcal{F}}^{i,e}\right)}_{j,\mathcal{E}},\ \mathcal{H}_j^{i,e}\overline{\left(u_j - \bar{u}_{j,\mathcal{F}}^{i,e}\right)}_{j,\mathcal{E}}\right) \\
&\overset{\substack{Lemma \\ 3.3}}{\leq} (\chi C_m + \tau K_M) \left\|\mathcal{H}_j^{\Delta}\overline{\left(u_j - \bar{u}_{j,\mathcal{F}}^{i,e}\right)}_{j,\mathcal{E}}\right\|_{L^2(\Omega_j)}^2 \\
&\quad + \tau \sum_{\star=i,e} \sigma_M^{\star} \left|\mathcal{H}_j^{\Delta}\overline{\left(u_j - \bar{u}_{j,\mathcal{F}}^{i,e}\right)}_{j,\mathcal{E}}\right|_{H^1(\Omega_j)}^2 \\
&\lesssim \sum_{\star=i,e} \left[\tau\sigma_M^{\star} + H^2\left(\chi C_m + \tau K_M\right)\right] \left|\overline{\left(u_j - \bar{u}_{j,\mathcal{F}}^{i,e}\right)}_{j,\mathcal{E}}\right|_{H^1(\Omega_j)}.
\end{aligned}
$$

Using [37, Lemmas 4.16, 4.17, and 4.19] and [37, Lemma 4.30], it follows that

$$
\| \overline{\left(u_j - \bar{u}_{j,\mathcal{F}}^{i,e}\right)}_{j,\mathcal{E}} \|^2 \leq C \left(1 + \log\frac{H}{h}\right)^3 |u_j|_{H^{1/2}(\partial\Omega_j)}^2.
$$

This means that

$$
|\bar{u}_{j,\mathcal{E}}^{i,e} - \bar{u}_{j,\mathcal{F}}^{i,e}|_{S_{\mathcal{F}}^{(j)}}^2 \lesssim C \left(1 + \log\frac{H}{h}\right)^3 \sum_{\star=i,e} \left[\tau\sigma_M^{\star} + H^2\left(\chi C_m + \tau K_M\right)\right] |\mathcal{H}_j^{\Delta} u_j|_{H^1(\Omega_j)}^2.
$$

To conclude, the face contribution gives the bound

$$|P_D u|^2_{S_\Gamma^{(j)}} \lesssim \sum_{\star=i,e} \sum_{\mathcal{F}\in\Xi_j^{\mathcal{F}}} \left[\tau\sigma_M^\star + H^2\left(\chi C_m + \tau K_M\right)\right] \left(1+\log\frac{H}{h}\right)^3 |\mathcal{H}_j^\Delta u_j|^2_{H^1(\Omega_j)}$$
$$\leq \max_{\star=i,e} \sum_{\mathcal{F}\in\Xi_j^{\mathcal{F}}} \frac{\tau\sigma_M^\star + H^2\left(\chi C_m + \tau K_M\right)}{\tau\sigma_m^\star} \left(1+\log\frac{H}{h}\right)^3 |u_j|^2_{S_\Gamma^{(j)}}.$$

Edge contributions. For simplicity, suppose that an edge $\mathcal{E}$ is shared only by three substructures, each with indexes j_1, j_2, and j_3. The extension to the case of more subdomains is then similar. Define $S_\mathcal{E}^{(j_{123})} := S_\mathcal{E}^{(j_1)} + S_\mathcal{E}^{(j_2)} + S_\mathcal{E}^{(j_3)}$. Then, the average operator is given by

$$\bar{u}_\mathcal{E}^{i,e} := \left(S_\mathcal{E}^{(j_{123})}\right)^{-1} \left(S_\mathcal{E}^{(j_1)} u_{j_1,\mathcal{E}} + S_\mathcal{E}^{(j_2)} u_{j_2,\mathcal{E}} + S_\mathcal{E}^{(j_3)} u_{j_3,\mathcal{E}}\right).$$

Proceeding in the same fashion as for the face contribution, it follows that

$$u_{j_1,\mathcal{E}} - \bar{u}_\mathcal{E}^{i,e} = \left(S_\mathcal{E}^{(j_{123})}\right)^{-1} \left[\left(S_\mathcal{E}^{(j_2)} + S_\mathcal{E}^{(j_3)}\right) u_{j_1,\mathcal{E}} - S_\mathcal{E}^{(j_2)} u_{j_2,\mathcal{E}} - S_\mathcal{E}^{(j_3)} u_{j_3,\mathcal{E}}\right],$$

which leads to

$$|R_\mathcal{E}^T \left(u_{j_1,\mathcal{E}} - \bar{u}_\mathcal{E}^{i,e}\right)|^2_{S_\Gamma^{(j_1)}} \leq 3u_{j_1,\mathcal{E}}^T S_\mathcal{E}^{(j_1)} u_{j_1,\mathcal{E}} \ + 3u_{j_2,\mathcal{E}}^T S_\mathcal{E}^{(j_2)} u_{j_2,\mathcal{E}} \ + u_{j_3,\mathcal{E}}^T S_\mathcal{E}^{(j_3)} u_{j_3,\mathcal{E}},$$

where we use analogous inequalities as in the face case.

Since we have included the edge averages into the primal space, we have the same average value for the three subdomains. By adding and subtracting the mean value over the edges $\bar{u}_{j_1,\mathcal{E}}^{i,e}$, we can get the estimate for the edges by using [37, Lemmas 4.16, 4.17, and 4.19]:

$$u_{j_1,\mathcal{E}}^T S_\mathcal{E}^{(j_1)} u_{j_1,\mathcal{E}} \leq \left[\tau\sigma_M^{i,e} + H^2\left(\chi C_m + \tau K_M\right)\right] \left(1+\log\frac{H}{h}\right) |\mathcal{H}_{j_1}^\Delta u_{j_1,\mathcal{E}}|^2_{H^1(\Omega_j)}.$$

In conclusion, the edge estimate gives

$$|P_D u|^2_{S_\Gamma^{(j)}} \leq \max_{\star=i,e} \sum_{\mathcal{E}\in\Xi_j^{\mathcal{E}}} \frac{\tau\sigma_M^\star + H^2\left(\chi C_m + \tau K_M\right)}{\tau\sigma_m^\star} \left(1+\log\frac{H}{h}\right) |u_j|^2_{S_\Gamma^{(j_1)}},$$

where the index j collects all contributions from the subdomains that share the edge $\mathcal{E}$. ∎

6. Parallel numerical results. We report here the results of several parallel numerical tests which confirm our theoretical estimates and study the performance of the proposed preconditioners with respect to the discretization parameters.

The weak scaling tests (with fixed local problem size per processor while the total problem size increases with the processor count) are performed on the supercomputer Galileo from Cineca center, 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. Meanwhile, the strong scaling tests (with fixed total problem size while the local problem size per processor decreases with the inverse of the processor count) are computed on the 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, for a total amount of 512 cores. The optimality tests are carried out on the cluster Eos at University of Pavia, a Linux Infiniband cluster with 21 nodes, each carrying 2 processors Intel Xeon Gold 6130 2.1 GHz with 16 cores each, for a total of 672 cores. Our C code is based on the PETSc library [2] from Argonne National Laboratory.

In Tests 1, 2, and 4, dual-primal preconditioners are applied with the standard rho-scaling, as Test 3 shows an almost computational equivalence between rho- and deluxe scaling for this application model, as concerns for nonlinear and linear iterations numbers per time step.

All our numerical experiments are carried both on a thin slab and on an idealized left ventricular geometry, modeled as a truncated ellipsoid. The latter is described in ellipsoidal coordinates by the parametric equations

$$\begin{cases} \mathbf{x} = a(r)\cos\theta\cos\varphi, & \theta_{\min} \le \theta \le \theta_{\max}, \\ \mathbf{y} = b(r)\cos\theta\sin\varphi, & \varphi_{\min} \le \varphi \le \varphi_{\max}, \\ \mathbf{z} = c(r)\sin\varphi, & 0 \le r \le 1, \end{cases}$$

where $a(r) = a_1 + r(a_2 - a_1)$, $b(r) = b_1 + r(b_2 - b_1)$, and $c(r) = c_1 + r(c_2 - c_1)$ with $a_{1,2}$, $b_{1,2}$, and $c_{1,2}$ given coefficients defining the main axes of the ellipsoid.

The fibers rotate intramurally linearly with the depth for a total amount of 120^o proceeding counterclockwise from epicardium ($r = 1$, outer surface of the truncated ellipsoid) to endocardium ($r = 0$, inner surface). Regarding the physiological coefficients in Table 6.1, we refer to the original paper [32] for the parameters of the ionic membrane model, while we refer to [10] for the bidomain and monodomain parameters.

The external stimulus of $I_{\text{app}} = 100$ mA/cm^3, needed for the potential to start propagating, is applied for 1 ms to the surface of the domain representing the endocardium. Instead, if a slab geometry is considered, the stimulus is applied in one corner of the domain, over a spheric volume of radius 0.1 cm.

Figures 6.1 and 6.2 show the time evolution of the extracellular and transmembrane potentials, respectively, from the epicardial view of a portion of the idealized left ventricle when the external stimulus I_{app} is applied at an epicardial location. We consider insulating boundary conditions, resting initial conditions, and a fixed time step size $\tau = 0.05$ ms.

In order to test the efficiency of our solver on parallel architectures, we also compute the parallel speedup $S_N = \frac{T_1}{T_N}$, the ratio between the runtime T_1 on 1 processor, and the average runtime T_N on N processors.

We use the default nonlinear solver (scalable nonlinear equations solvers (SNES)) in the PETSc library [2], which consists of a Newton method with cubic backtracking linesearch. We adopt the default SNES convergence test as stopping criterion, based

TABLE 6.1

Conductivity coefficients for the bidomain model and physiological parameters for the Rogers–McCulloch ionic model.

Bidomain conductivity coeff.		Ionic parameters	
σ_l^i	$3 \times 10^{-3}\Omega^{-1}$ cm^{-1}	G	1.2 Ω^{-1} cm^{-2}
σ_t^i	$3.1525 \times 10^{-4}\Omega^{-1}$ cm^{-1}	η_1	4.4 Ω^{-1} cm^{-1}
σ_n^i	$3.1525 \times 10^{-5}\Omega^{-1}$ cm^{-1}	η_2	0.012
σ_l^e	$2 \times 10^{-3}\Omega^{-1}$ cm^{-1}	v_{th}	13 mV
σ_t^e	$1.3514 \times 10^{-3}\Omega^{-1}$ cm^{-1}	v_p	100 mV
σ_n^e	$6.757 \times 10^{-4}\Omega^{-1}$ cm^{-1}	C_m	1 mF/cm^2

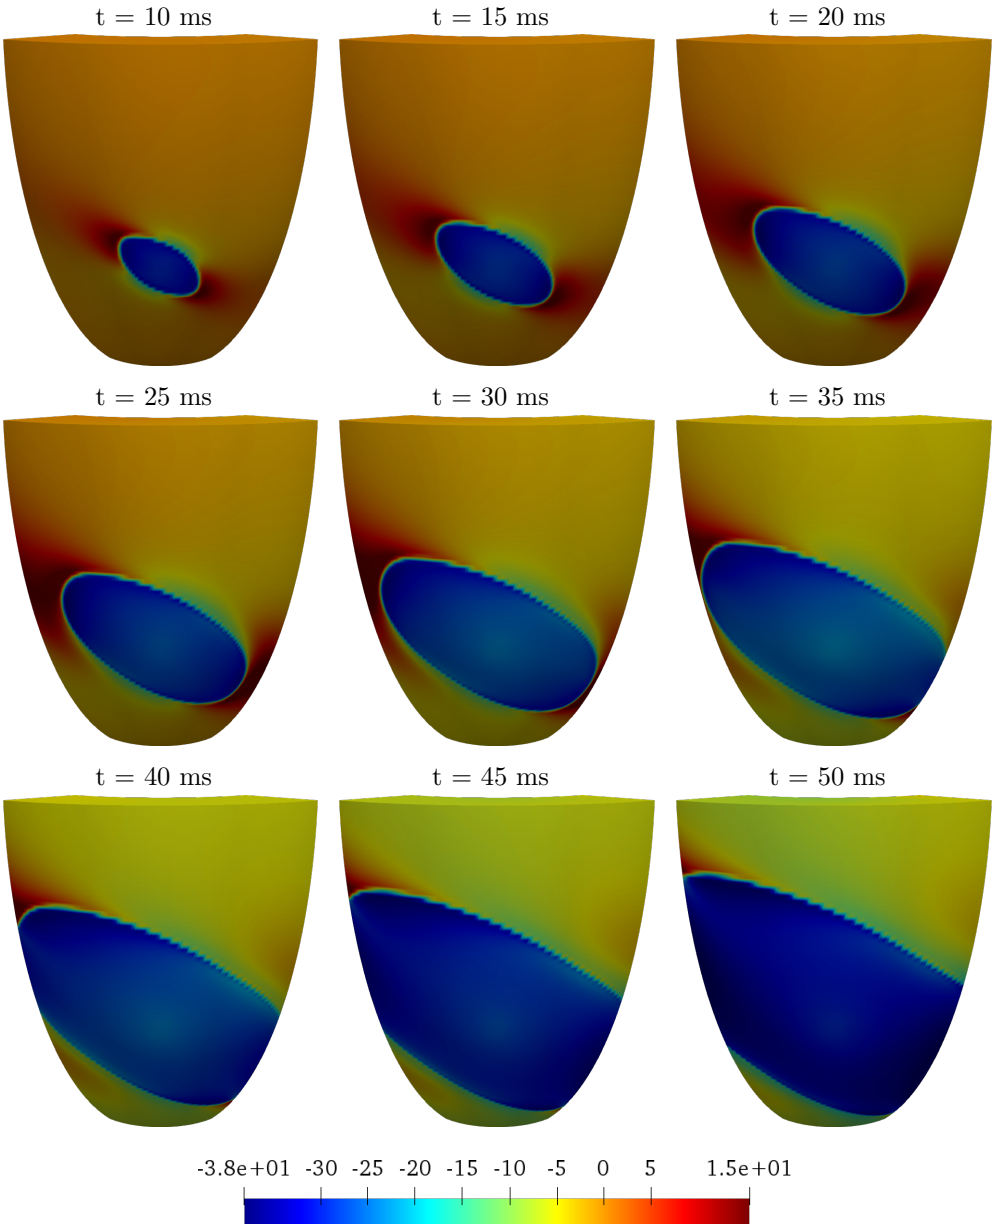

FIG. 6.1. *Snapshots (every 5 ms) of extracellular potential u_e time evolution. For each time frame, we report the epicardial view of a portion of the left ventricle, modeled as a truncated ellipsoid.*

on the comparison of the L^2-norm of the nonlinear function at the current iterate and at the current step with specified tolerances. Since the linear system arising from the discretization of the Jacobian problem at each Newton step is symmetric, we solve it with the PCG method, with the BDDC or FETI-DP preconditioners from the PETSc library, or with the *Boomer* Algebraic MultiGrid (bAMG) preconditioner from the Hypre library. We use the default Hypre parameters strong threshold = 0.25, number of smoothing levels = 25, and number of levels of aggressive coarsening = 0. In the

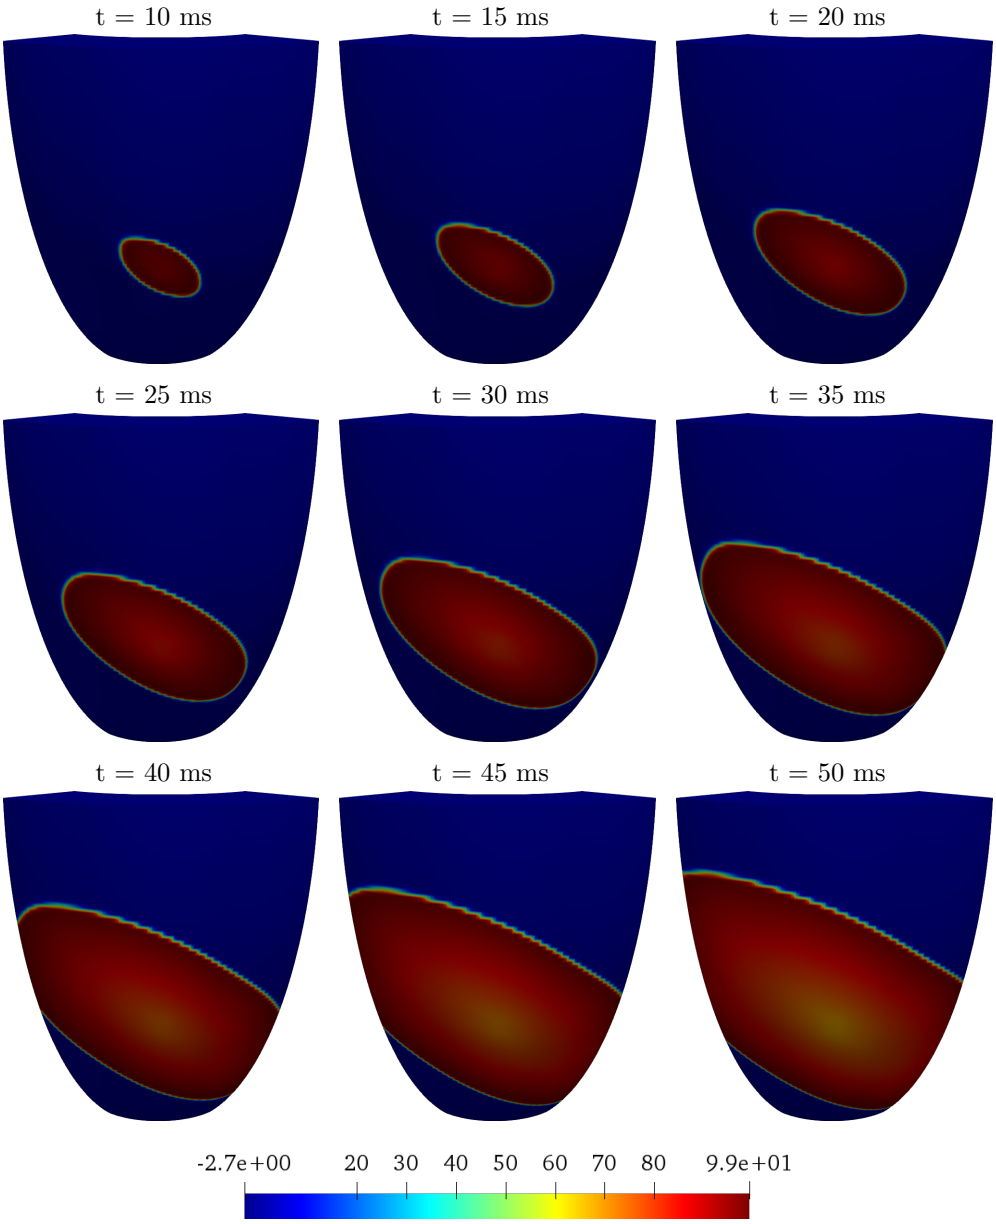

FIG. 6.2. *Snapshots (every* 5 *ms) of transmembrane potential* v *time evolution. For each time frame, we report the epicardial view of a portion of the left ventricle, modeled as a truncated ellipsoid.*

strong scaling tests, when testing the performance of the proposed solver against two different ionic models (Rogers–McCulloch and Luo–Rudy phase 1), the generalized minimal residual (GMRES) method is applied. The convergence criterion of the linear solver is based on the decreasing of the residual norm (default from PETSc).

All parameters can be found in Table 6.2.

Test 1: Weak scaling. The first set of tests we report here is a weak scaling test on both slab and ellipsoidal domains, performed on the Galileo cluster. For both cases,

TABLE 6.2

PETSc SNES and Krylov subSPace solvers (KSP) tolerances. r_{tol} is the relative convergence tolerance, a_{tol} is the absolute convergence tolerance, d_{tol} is the KSP divergence convergence tolerance, and s_{tol} is the convergence tolerance related to the solution change between Newton steps.

KSP	$r_{tol} = 1e^{-08}$	$a_{tol} = 1e^{-10}$	$d_{tol} = 1e^{+04}$
SNES	$r_{tol} = 1e^{-04}$	$a_{tol} = 1e^{-08}$	$s_{tol} = 1e^{-08}$

TABLE 6.3

Weak scaling test for the bidomain decoupled solver on the cluster Galileo. Slab domain, local mesh $16 \cdot 16 \cdot 16$ elements. Simulations of 2 ms of cardiac activation with $dt = 0.05$ ms (40 time steps). Comparison of Newton–Krylov solvers with bAMG, BDDC, and FETI-DP preconditioners. Average Newton iterations per time step (nit); average conjugate gradient iterations per Newton iteration (lit); average CPU solution time per time step (time) in seconds.

procs	mesh	dofs	bAMG			BDDC			FETI-DP		
			nit	lit	time	nit	lit	time	nit	lit	time
32	$64 \cdot 64 \cdot 32$	278,850	1.25	106	4.9	1.0	22	6.1	1.25	10	6.0
64	$128 \cdot 64 \cdot 32$	553,410	1.25	132	6.8	1.0	27	6.2	1.25	11	6.0
128	$128 \cdot 128 \cdot 32$	1,098,306	1.25	180	9.4	1.0	32	7.6	1.25	10	7.4
256	$256 \cdot 128 \cdot 32$	2,188,098	1.25	237	15.2	1.0	39	7.2	1.25	10	7.9
512	$256 \cdot 256 \cdot 32$	4,359,234	1.25	318	20.0	1.0	48	10.1	1.25	10	11.1
1024	$512 \cdot 256 \cdot 32$	8,701,506	1.25	405	29.6	1.0	63	13.8	1.25	10	18.7
2048	$512 \cdot 512 \cdot 32$	17,369,154	1.25	536	40.3	1.0	78	33.5	1.25	10	63.2

TABLE 6.4

Weak scaling test for the bidomain decoupled solver on the cluster Galileo. Ellipsoidal domain, local mesh $16 \cdot 16 \cdot 16$ elements. Simulations of 2 ms of cardiac activation with $dt = 0.05$ ms (40 time steps). Comparison of Newton–Krylov solvers with bAMG, BDDC, and FETI-DP preconditioners. Average Newton iterations per time step (nit); average conjugate gradient iterations per Newton iteration (lit); average CPU solution time per time step (time) in seconds.

procs	mesh	dofs	bAMG			BDDC			FETI-DP		
			nit	lit	time	nit	lit	time	nit	lit	time
32	$64 \cdot 32 \cdot 64$	278,850	1.0	86	3.3	1.0	30	5.4	1.0	20	4.7
64	$64 \cdot 64 \cdot 64$	549,250	1.07	124	6.0	1.07	37	6.2	1.07	20	6.5
128	$64 \cdot 128 \cdot 64$	1,090,050	1.20	207	11.3	1.20	26	7.5	1.2	19	6.6
256	$64 \cdot 256 \cdot 64$	2,171,650	1.42	348	22.2	1.42	25	8.7	1.42	17	10.7
512	$128 \cdot 256 \cdot 64$	4,309,890	1.42	335	21.3	1.42	27	10.5	1.42	18	11.4
1024	$256 \cdot 256 \cdot 64$	8,586,370	out of memory			1.42	28	12.5	1.42	19	11.0
2048	$512 \cdot 256 \cdot 64$	17,139,330	out of memory			1.42	28	26.6	1.42	19	21.4

we fix the local mesh size to $16 \cdot 16 \cdot 16$, and we increase the number of subdomains from 32 to 2048, thus resulting in an increasing slab geometry and in an increasing portion of ellipsoid. From Tables 6.3 and 6.4, it is evident that the dual-primal algorithms have a better performance than bAMG: the average number of linear iterations per Newton iteration is clearly lower and does not increase with the number of subdomains, except for BDDC on the slab domain, where the linear iterations increase unexpectedly. Moreover, the reported average CPU times (in seconds) per Newton step (see also Figure 6.3) are slightly better for the BDDC and FETI-DP preconditioners, except for FETI-DP on the slab domain and 2048 processors. In the harder ellipsoidal tests, both BDDC and FET-DP are scalable and outperform bAMG when the number of processors increases past 128, indicating lower computational complexity and interprocessor communications.

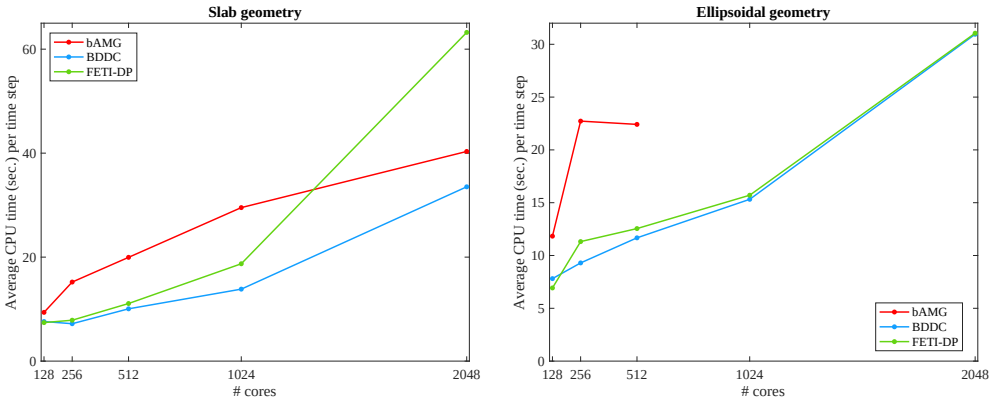

FIG. 6.3. *Weak scaling test on the cluster Galileo. Slab (left) and ellipsoidal (right) domains, local mesh* $16 \cdot 16 \cdot 16$ *elements. Simulations of* 2 *ms of cardiac activation with* $dt = 0.05$ *ms (*40 *time steps). Comparison of average CPU time per time step, in seconds.*

TABLE 6.5

Strong scaling test for the bidomain decoupled solver on the cluster Indaco. Slab domain, global mesh $192 \cdot 192 \cdot 32$ *elements,* $2,458,434$ *dofs. Simulations of* 2 *ms of cardiac activation with* $dt = 0.05$ *ms (*40 *time steps). Comparison of Newton–Krylov solvers with bAMG, BDDC, and FETI-DP preconditioners. Average Newton iterations per time step (nit); average conjugate gradient iterations per Newton iteration (lit); average CPU solution time per time step (time) in seconds; parallel speedup with respect to* 32 *(*S_{32}*) and* 64 *(*S_{64}*) processors.*

procs	bAMG					BDDC					FETI-DP				
	nit	lit	time	S_{32}	S_{64}	nit	lit	time	S_{32}	S_{64}	nit	lit	time	S_{32}	S_{64}
32	1.25	250	116.0	-	-	1.0	27	348.2	-	-	1.25	11	352.2	-	-
64	1.25	252	62.7	1.8	-	1.22	32	59.5	5.8	-	1.25	17	53.0	6.6	-
128	1.25	252	33.6	3.5	1.8	1.22	37	21.9	15.8	2.7	1.25	21	19.9	17.6	2.6
256	1.25	252	18.6	6.2	3.4	1.22	22	10.4	33.5	5.7	1.25	13	8.9	39.7	5.9

TABLE 6.6

Strong scaling test for the bidomain decoupled solver on the cluster Indaco. Ellipsoidal domain, global mesh $128 \cdot 128 \cdot 64$ *elements,* $2,163,330$ *dofs. Simulations of* 2 *ms of cardiac activation with* $dt = 0.05$ *ms (*40 *time steps). Comparison of Newton–Krylov solvers with bAMG, BDDC and FETI-DP preconditioners. Average Newton iterations per time step (nit); average conjugate gradient iterations per Newton iteration (lit); average CPU solution time per time step (time) in seconds; parallel speedup with respect to* 32 *(*S_{32}*) and* 64 *(*S_{64}*) processors.*

procs	bAMG					BDDC					FETI-DP				
	nit	lit	time	S_{32}	S_{64}	nit	lit	time	S_{32}	S_{64}	nit	lit	time	S_{32}	S_{64}
32	1.92	311	188.4	-	-	1.92	36	571.8	-	-	1.92	14	558.2	-	-
64	1.92	310	113.4	1.7	-	1.92	30	129.1	4.4	-	1.92	19	129.7	4.3	-
128	1.92	310	60.5	3.1	1.9	1.92	40	40.2	14.2	3.2	1.92	24	42.4	13.2	3.1
256	1.92	311	32.2	5.8	3.1	1.92	23	15.1	37.9	8.5	1.92	14	19.0	29.4	6.8

Test 2: Strong scaling. We now perform a strong scaling test for the two geometries on the Indaco cluster. For the thin slab geometry, we fix the global mesh to $192 \cdot 192 \cdot 32$ elements, and we increase the number of subdomains. We fix the global mesh to $128 \cdot 128 \cdot 64$ elements for the portion of ellipsoid instead. We observe from Tables 6.5 and 6.6 that, as the local number of degrees of freedom decrease, the preconditioner with the better balance in terms of average linear iterations and CPU time per time step is FETI-DP. In both cases, BDDC and FETI-DP preconditioners

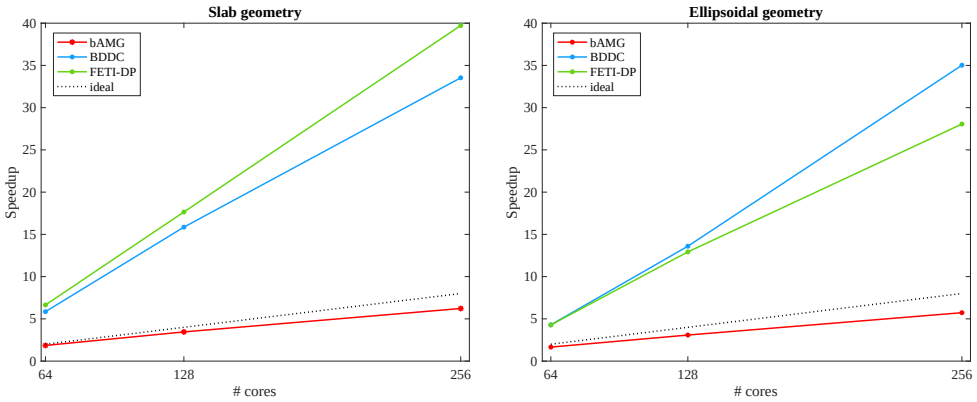

FIG. 6.4. *Strong scaling test on the cluster Indaco. Left: slab domain with global mesh* $192 \cdot 192 \cdot 32$. *Right: ellipsoidal domain with global mesh* $128 \cdot 128 \cdot 64$. *Simulations of* 2 *ms of cardiac activation with* $dt = 0.05$ *ms (*40 *time steps). Comparison of actual parallel speedup computed with respect to* 32 *cores (ideal speedup dotted). Performed on the Indaco cluster.*

TABLE 6.7

Strong scaling test for the bidomain decoupled solver on the cluster Indaco. Slab domain, global mesh $192 \cdot 192 \cdot 32$ *elements,* $2,458,434$ *dofs. Simulations of* 2 *ms of cardiac activation with* $dt = 0.05$ *ms (*40 *time steps). Comparison of Newton–Krylov solvers with BDDC preconditioner using Rogers–McCulloch (RMC) and Luo–Rudy phase* 1 *(LR*1*) ionic models. Average Newton iterations per time step (nit); average conjugate gradient iterations per Newton iteration (lit); average CPU solution time per time step (time) in seconds; parallel speedup* (S_N) *computed with respect to* $N = 32$ *and* $N = 64$ *processors.*

procs	RMC					LR1				
	nit	lit	time	S_{32}	S_{64}	nit	lit	time	S_{32}	S_{64}
32	1.25	16.97	220.25	-	-	2.85	16.97	502.25	-	-
64	1.25	19.92	62.07	3.55	-	2.85	19.57	140.92	3.56	-
128	1.25	15.3	19.2	11.47	3.23	2.85	15.0	43.9	11.44	3.21
256	1.25	17.45	5.8	37.97	10.7	2.85	29.5	17.0	38.08	10.68

outperform the ideal speedup, while bAMG is suboptimal (see Fig. 6.4). Moreover, we compare the performance of the Newton–Krylov solver with BDDC preconditioner using the Rogers–McCulloch and Luo–Rudy phase 1 ionic models in Tables 6.7 and 6.8. In this case the Jacobian linear system is solved with the GMRES method. By increasing the complexity of the ionic current, we observe an increase in the average number of Newton iterations from 1–2 per time step using the Rogers–McCulloch model to 2–3 per time step with the Luo–Rudy phase 1 model. On the other hand, the average numbers of linear iterations per time step for the two ionic models are comparable, indicating that our dual-primal solver retains its good convergence properties even for more complex ionic models. As a consequence, the CPU times for the Luo–Rudy phase 1 model increase due to the increase of nonlinear iterations, but the associated parallel speedups of the two models are comparable.

Test 3: Optimality tests. Tables 6.9 and 6.10 report the results of optimality tests, for both slab and ellipsoid geometries, carried out on the Eos cluster. We fix the number of processors (subdomains) to $4 \cdot 4 \cdot 4$, and we increase the local size H/h from 8 to 24, thus reducing the finite element size h. We focus only on the behavior of the BDDC preconditioner, as FETI-DP has been proven to be spectrally equivalent. We consider both scalings (ρ-scaling on top, deluxe scaling at the bottom of each table), and we test the solver for increasing primal spaces: V includes only

TABLE 6.8

Strong scaling test for the bidomain decoupled solver on the cluster Indaco. Ellipsoidal domain, global mesh $128{\cdot}128{\cdot}64$ *elements,* $2,163,330$ *dofs. Simulations of* 2 *ms of cardiac activation with* $dt = 0.05$ *ms (*40 *time steps). Comparison of Newton–Krylov solvers with BDDC preconditioner using Rogers–McCulloch (RMC) and Luo–Rudy phase* 1 *(LR*1*) ionic models. Average Newton iterations per time step (nit); average conjugate gradient iterations per Newton iteration (lit); average CPU solution time per time step (time) in seconds; parallel speedup (*S_N*) computed with respect to* $N = 32$ *and* $N = 64$ *processors.*

procs	RMC					LR1				
	nit	lit	time	S_{32}	S_{64}	nit	lit	time	S_{32}	S_{64}
32	2	21.1	436.5	-	-	3.95	20.5	862.25	-	-
64	2	26.9	99.3	4.39	-	3.95	25.9	194.87	4.43	-
128	2	21.4	27.27	16.0	3.64	3.95	20.8	53.47	16.12	3.64
256	2	30.0	8.17	53.42	12.15	3.95	29.5	16.08	53.62	12.11

TABLE 6.9

Optimality test on the Eos cluster for PCG-BDDC. Slab domain, $4{\cdot}4{\cdot}4$ *subdomains, increasing local size from* $4 \cdot 4 \cdot 4$ *to* $24 \cdot 24 \cdot 24$*. Comparison between different scaling and different primal sets (*V *= vertices,* E *= edges,* F *= faces). Average Newton iterations (nit), average number of linear iterations (lit), average CPU time in seconds, and average condition number (cond) per time step.*

ρ-scaling												
H/h	V				V+E				V+E+F			
	nit	lit	time	cond	nit	lit	time	cond	nit	lit	time	cond
4	1.24	26	1.7	8.4	1.24	11	0.9	1.9	1.24	9	0.9	1.7
8	1.21	47	3.4	24.1	1.24	14	1.4	2.6	1.24	12	1.4	2.5
12	1.04	66	10.3	42.9	1.17	18	6.4	3.2	1.21	15	4.5	3.2
16	out of memory				1.0	20	11.9	3.7	1.0	20	11.6	3.7
20	out of memory				1.0	22	34.2	4.2	1.0	20	32.5	4.2
24	out of memory				1.0	23	83.1	4.6	1.0	21	80.5	4.5
deluxe scaling												
H/h	V				V+E				V+E+F			
	nit	lit	time	cond	nit	lit	time	cond	nit	lit	time	cond
4	1.24	26	1.9	8.4	1.24	11	0.9	1.9	1.24	9	1.0	1.7
8	1.24	47	4.1	24.0	1.24	14	1.7	2.6	1.24	12	1.8	2.5
12	1.07	65	14.7	42.7	1.17	18	5.5	3.2	1.21	15	7.8	3.2
16	1.0	80	30.0	63.7	1.0	20	19.1	3.7	1.0	20	21.4	3.7
20	1.0	90	93.8	86.3	1.0	22	73.9	4.2	1.0	20	70.0	4.2
24	1.0	99	211.9	110.1	1.0	24	205.8	4.5	1.0	21	247.3	4.5

vertex constraints, V+E includes vertex and edge constraints, and V+E+F includes vertex, edge, and face constraints. We consider a time interval of 2 ms during the cardiac activation phase. The time step is $dt = 0.05$ ms for a total amount of 40 time steps. Similar results hold for both geometries. Despite a higher average CPU time when using the deluxe scaling, all the other parameters are quite similar between the two scalings. We observe almost linear dependence of the condition number if the coarsest primal space (i.e., V) is chosen (see also Figures 6.5 and 6.6 bottom), while we obtain quasi-optimality if we enrich the primal space by adding edges (V+E) and faces (V+E+F).

Test 4: Whole beat (activation-recovery) simulations. In this last set of tests (performed on the Indaco cluster), we compare the performance of our dual-primal and multigrid preconditioners during a whole beat, i.e., during a complete activation-recovery interval over the computational domain. We fix the number of subdomains to $128 = 8 \cdot 8 \cdot 2$ and the global mesh size to $192 \cdot 96 \cdot 24$, obtaining local problems with 8,450 degrees of freedom. We consider either a portion of ellipsoid,

TABLE 6.10

Optimality test on the Eos cluster for PCG-BDDC. Ellipsoidal domain, $4 \cdot 4 \cdot 4$ subdomains, increasing local size from $8 \cdot 8 \cdot 8$ to $24 \cdot 24 \cdot 24$. Comparison between different scaling and different primal sets (V = vertices, E = edges, F = faces). Average Newton iterations (nit), average number of linear iterations (lit), average CPU time in seconds, and average condition number (cond) per time step.

ρ-scaling												
H/h	V				V+E				V+E+F			
	nit	lit	time	cond	nit	lit	time	cond	nit	lit	time	cond
8	2.0	50	5.5	30.1	2.0	17	2.4	4.3	2.0	16	2.4	4.0
12	2.0	66	11.8	54.6	2.0	19	5.6	5.4	2.0	18	5.6	4.9
16	2.0	80	33.9	80.6	2.0	21	16.9	6.3	2.0	21	16.9	5.7
20	2.0	91	90.3	108.2	2.0	22	44.7	6.9	2.0	21	44.9	6.3
24	1.46	100	206.3	137.2	1.46	24	109.3	7.5	1.46	23	84.0	6.8
deluxe scaling												
H/h	V				V+E				V+E+F			
	nit	lit	time	cond	nit	lit	time	cond	nit	lit	time	cond
8	2.0	49	6.7	31.0	2.0	17	3.1	4.3	2.0	16	3.0	4.0
12	2.0	54	27.2	54.6	2.0	19	9.6	5.4	2.0	18	9.6	4.9
16	2.0	79	54.1	80.6	2.0	21	32.1	6.2	2.0	21	32.7	5.7
20	2.0	90	142.4	108.2	2.0	22	125.8	7.0	2.0	21	111.1	6.3
24	1.46	99	329.3	137.1	1.46	24	236.7	7.5	1.46	22	247.1	6.8

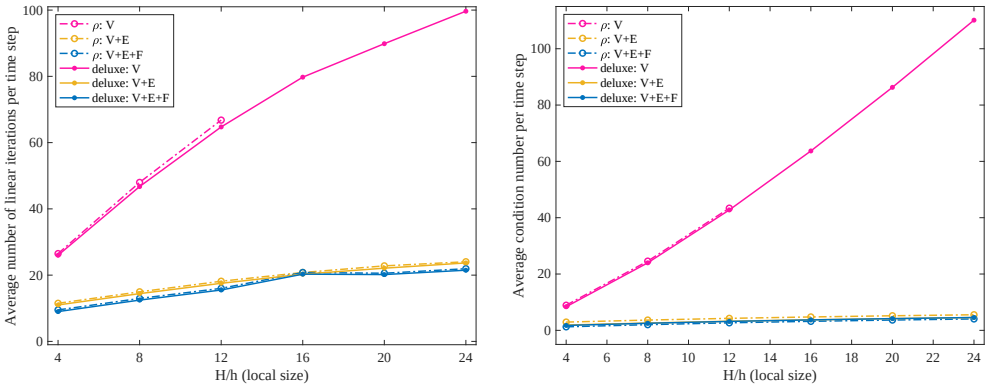

FIG. 6.5. *Optimality test on the Eos cluster for PCG-BDDC. Slab domain, $4 \cdot 4 \cdot 4$ subdomains, increasing local size from $4 \cdot 4 \cdot 4$ to $24 \cdot 24 \cdot 24$. Comparison between different scaling (dash-dotted ρ-scaling, continuous deluxe scaling) and different primal sets (V = vertices, E = edges, F = faces). Average number of linear iterations (left) and average condition numbers (right) per time step.*

defined by $\varphi_{\min} = -\pi/2$, $\varphi_{\max} = 0$, $\theta_{\min} = -3/8\pi$, and $\theta_{\max} = \pi/8$, or a slab of dimensions $1.92 \times 1.92 \times 0.48$ cm^3. The ellipsoid test is on a time interval of $[0, 170]$ ms for a total of 3400 time steps, while the slab test is on the time interval $[0, 120]$ ms for a total of 2400 time steps. Both time intervals are enough to complete the activation and recovery phases over the whole domains, given the short action potential duration of the Rogers–McCulloch ionic model considered. In Figures 6.7 and 6.8 we report the trend of the average number of linear iteration per time step during the simulation. The number of iterations remains bounded and almost constant during the test. Moreover, we notice a huge difference between the multigrid and the dual-primal preconditioners, with a reduction of more than 90% for the latter. If we focus on the trend of the dual-primal preconditioners' average number of linear iterations (Figures 6.7 and 6.8, on the right), we see that on both domains FETI-DP is

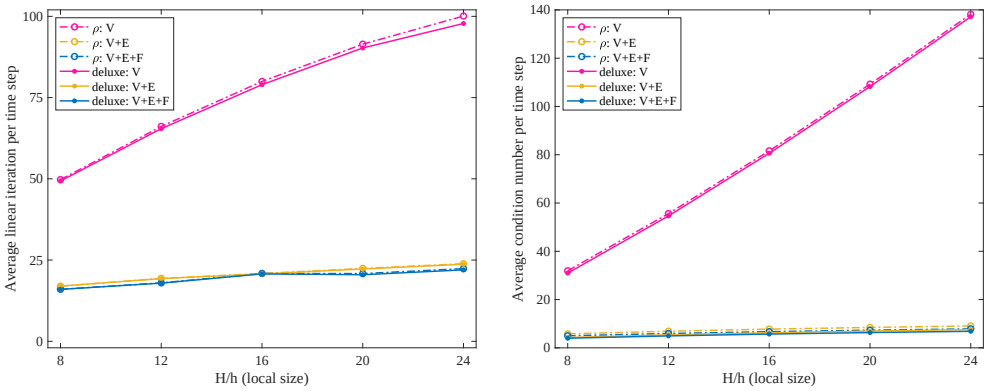

FIG. 6.6. *Optimality test on the Eos cluster for PCG-BDDC with different scaling (dash-dotted ρ-scaling, continuous deluxe scaling) and primal sets (V = vertices, E = edges, F = faces). Ellipsoidal domain, $4 \cdot 4 \cdot 4$ subdomains, increasing local size from $8 \cdot 8 \cdot 8$ to $24 \cdot 24 \cdot 24$. Average number of linear iterations (left) and average condition numbers (right) per time step.*

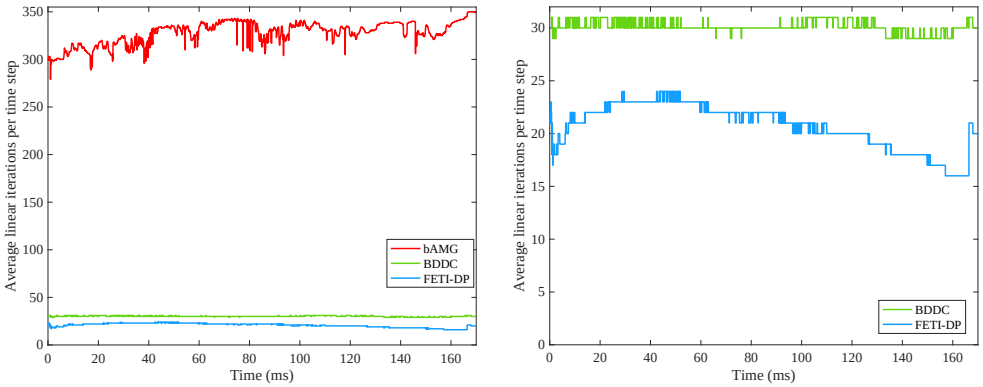

FIG. 6.7. *Full activation-recovery simulation on ellipsoidal domain, time interval* $[0, 170]$ *ms,* $3,400$ *time steps. Fixed number of subdomains* $8 \cdot 8 \cdot 2$ *and fixed global mesh* $192 \cdot 96 \cdot 24$. *Average number of linear iterations per time step (left), zoom over dual-primal preconditioner (right).*

affected by the different phases of the action potential: there is an initial peak during the activation phase, followed by an increase in the number of linear iterations as the potential propagates in the tissue and by a slow decrease as wider portions of tissue return to resting. Similar behavior can be observed for the BDDC preconditioner on the slab domain: there is an initial peak corresponding to the activation phase, followed by a constant period, as the tissue turn to resting. This trend is not visible for BDDC on the ellipsoidal domain, due to the complexity of the geometry. We also observe a better performance of the dual-primal preconditioners in terms of average CPU time per time step (see Table 6.11).

7. Conclusions. We have constructed dual-primal preconditioners for fully implicit discretizations of the Bidomain system, which are solved through a decoupling strategy. We have proved a convergence bound of the preconditioned FETI-DP and BDDC bidomain operators with deluxe scaling. Parallel numerical tests validate the bound and show the efficiency and robustness of the solver, thus enlarging the class of methods available for the efficient and accurate numerical solution of this complex

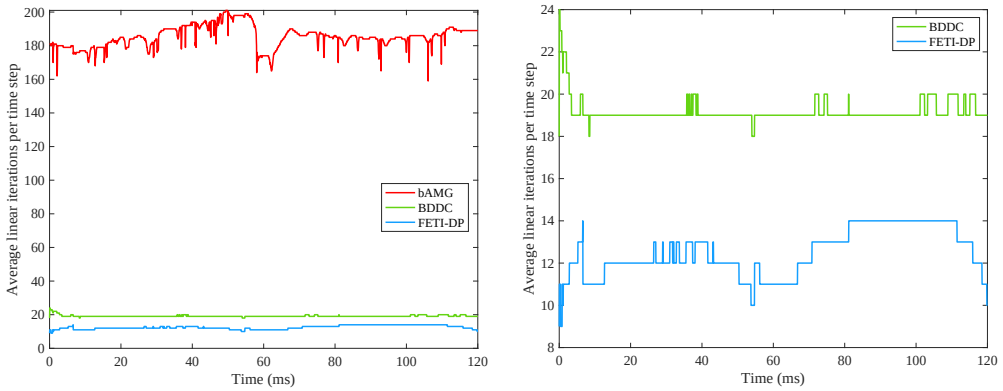

FIG. 6.8. *Full activation-recovery simulation on slab domain, time interval* $[0, 120]$ *ms,* $2,400$ *time steps. Fixed number of subdomains* $8 \cdot 8 \cdot 2$ *and fixed global mesh* $192 \cdot 96 \cdot 24$. *Average number of linear iterations per time step (left), zoom over dual-primal preconditioner (right).*

TABLE 6.11
Whole beat simulation on time interval $[0, 170]$ *ms,* $3,400$ *time steps for the ellipsoidal domain and time interval* $[0, 120]$ *ms,* $2,400$ *time steps for the slab. Fixed number of subdomains* $8 \cdot 8 \cdot 2$ *and fixed global mesh* $192 \cdot 96 \cdot 24$. *Comparison of average Newton steps, average linear iterations (lit), and average CPU time (in seconds) per time step.*

	procs	dofs	bAMG			BDDC			FETI-DP		
			nit	lit	time	nit	lit	time	nit	lit	time
Slab	128	8,450	1.4	185	11.28	1.4	19	8.02	1.4	12	7.62
Ellipsoid	128	8,450	1.97	328	13.24	1.97	30	8.85	1.97	21	8.05

biophysical reaction-diffusion model. Additional research is needed in order to assess the performance of the proposed dual-primal solvers for more realistic ionic models and with respect to optimized multigrid solvers.

REFERENCES

[1] C.M. AUGUSTIN, G.A. HOLZAPFEL, AND O. STEINBACH, *Classical and all-floating FETI methods for the simulation of arterial tissues*, Internat. J. Numer. Methods Engrg., 99 (2014), pp. 290–312.
[2] S. BALAY ET AL., *PETSc Web Page*, 2019, https://petsc.org/release/.
[3] D. BRANDS, A. KLAWONN, O. RHEINBACH, AND J. SCHRÖDER, *Modelling and convergence in arterial wall simulations using a parallel FETI solution strategy*, Comput. Methods Biomech. Biomed. Engrg., 11 (2008), pp. 569–583.
[4] L. BEIRÃO DA VEIGA, L.F. PAVARINO, S. SCACCHI, O. WIDLUND, AND S. ZAMPINI, *Isogeometric BDDC preconditioners with deluxe scaling*, SIAM J. Sci. Comput., 36 (2014), pp. A1118–A1139.
[5] L. A. CHARAWI, *Isogeometric overlapping Schwarz preconditioners for the Bidomain reaction–diffusion system*, Comput. Methods Appl. Mech. Engrg, 319 (2017), pp. 472–490.
[6] H. CHEN, X. LI, AND Y. WANG, *A splitting preconditioner for a block two-by-two linear system with applications to the Bidomain equations*, J. Comput. Appl. Math., 321 (2017), pp. 487–498.
[7] H. CHEN, X. LI, AND Y. WANG, *A two-parameter modified splitting preconditioner for the bidomain equations*, Calcolo, 56 (2019), 21.
[8] P. COLLI FRANZONE, L.F. PAVARINO, AND S. SCACCHI, *A numerical study of scalable cardiac electro-mechanical solvers on HPC architectures*, Front. Physiol., 9 (2018), 268.
[9] P. COLLI FRANZONE AND G. SAVARÉ, *Degenerate evolution systems modeling the cardiac electric field at micro-and macroscopic level*, in Evolution Equations, Semigroups, and Functional Analysis, Springer, Cham, 2002, pp. 49–78.

[10] P. Colli Franzone, L.F. Pavarino, and S. Scacchi, *Mathematical Cardiac Electrophysiology*, Springer, Cham, 2014.
[11] P. Colli Franzone, L.F. Pavarino, and S. Scacchi, *Parallel multilevel solvers for the cardiac electro-mechanical coupling*, Appl. Numer. Math., 95 (2015), pp. 140–153.
[12] C.R. Dohrmann, *A preconditioner for substructuring based on constrained energy minimization*, SIAM J. Sci. Comput., 25 (2003), pp. 246–258.
[13] C.R. Dohrmann and O.B. Widlund, *A BDDC algorithm with deluxe scaling for three-dimensional $H(curl)$ problems*, Commun. Pure Appl. Math., 69 (2016), pp. 745–770.
[14] C. Farhat, M. Lesoinne, P. LeTallec, K. Pierson, and D. Rixen, *FETI-DP: A dual–primal unified FETI method–part* I: *A faster alternative to the two-level FETI method*, Internat. J. Numer. Methods Engrg., 50 (2001), pp. 1523–1544.
[15] I.J. LeGrice, B.H. Smaill, L.Z. Chai, S.G. Edgar, J.B. Gavin, and P.J. Hunter, *Laminar structure of the heart: Ventricular myocyte arrangement and connective tissue architecture in the dog*, Amer. J. Physiol. Heart Circ. Physiol., 269 (1995), pp. H571–H582.
[16] J. Li and O.B. Widlund, *FETI-DP, BDDC, and block Cholesky methods*, Internat. J. Numer. Methods Engrg., 66 (2006), pp. 250–271.
[17] C. Luo and Y. Rudy, *A model of the ventricular cardiac action potential. Depolarization, repolarization, and their interaction*, Circ. Res., 68 (1991), pp. 1501–1526.
[18] A. Klawonn, O.B. Widlund, and M. Dryja, *Dual-primal FETI methods for three-dimensional elliptic problems with heterogeneous coefficients*, SIAM J. Numer. Anal., 40 (2002), pp. 159–179.
[19] A. Klawonn and O.B. Widlund, *Dual-primal FETI methods for linear elasticity*, Comm. Pure Appl. Math., 59 (2006), pp. 1523–1572.
[20] A. Klawonn and O. Rheinbach, *Highly scalable parallel domain decomposition methods with an application to biomechanics*, ZAMM Z. Angew. Math. Mech., 90 (2010), pp. 5–32.
[21] J. Mandel and C.R. Dohrmann, *Convergence of a balancing domain decomposition by constraints and energy minimization*, Numer. Linear Algebra Appl., 10 (2003), pp. 639–659.
[22] J. Mandel, C.R. Dorhmann, and R. Tezaur, *An algebraic theory for primal and dual substructuring methods by constraints*, Appl. Numer. Math, 54 (2005), pp. 167–193.
[23] M. Munteanu and L.F. Pavarino, *Decoupled Schwarz algorithms for implicit discretizations of nonlinear monodomain and bidomain systems*, Math. Models Methods Appl. Sci., 19 (2009), pp. 1065–1097.
[24] M. Munteanu, L.F. Pavarino, and S. Scacchi, *A scalable Newton–Krylov–Schwarz method for the bidomain reaction-diffusion system*, SIAM J.Sci. Comput., 31 (2009), pp. 3861–3883.
[25] M. Murillo and X-C. Cai, *A fully implicit parallel algorithm for simulating the non-linear electrical activity of the heart*, Numer. Linear Algebra Appl., 11 (2004), pp. 261–277.
[26] L. F. Pavarino and S. Scacchi, *Multilevel additive Schwarz preconditioners for the bidomain reaction-diffusion system*, SIAM J. Sci. Comput., 31 (2008), pp. 420–443.
[27] L.F. Pavarino, S. Scacchi, and S. Zampini, *Newton–Krylov-BDDC solvers for nonlinear cardiac mechanics*, Comput. Methods Appl. Mech. Engrg., 295 (2015), pp. 562–580.
[28] M. Pennacchio, G. Savaré, and P. Colli Franzone, *Multiscale modeling for the bioelectric activity of the heart*, SIAM J. Math. Anal., 37 (2005), pp. 1333–1370.
[29] A. Quarteroni, A. Manzoni, and C. Vergara, *The cardiovascular system: Mathematical modelling, numerical algorithms and clinical applications*, Acta Numer., 26 (2017), pp. 365–590.
[30] A. Quarteroni, T. Lassila, S. Rossi, and R. Ruiz-Baier, *Integrated heart: Coupling multiscale and multiphysics models for the simulation of the cardiac function*, Comput. Methods Appl. Mech. Engrg., 314 (2017), pp. 345–407.
[31] O. Rheinbach, *Parallel Scalable Iterative Substructuring: Robust Exact and Inexact FETI-DP Methods with Applications to Elasticity*, Ph.D. thesis, University of Duisburg-Essen, 2006.
[32] J.M. Rogers and A.D. McCulloch, *A collocation-Galerkin finite element model of cardiac action potential propagation*, IEEE Trans. Biomed. Engrg., 41 (1994), pp. 743–757.
[33] S. Scacchi, *A hybrid multilevel Schwarz method for the bidomain model*, Comput. Methods Appl. Mech. Engrg., 197 (2008), pp. 4051–4061.
[34] S. Scacchi, *A multilevel hybrid Newton–Krylov–Schwarz method for the bidomain model of electrocardiology*, Comput. Methods Appl. Mech. Engrg., 200 (2011), pp. 717–725.
[35] J. Sundnes, G. T. Lines, X. Cai, B. F. Nielsen, K.-A. Mardal, and A. Tveito, *Computing the Electrical Activity of the Heart*, Springer, Cham, 2006.
[36] K.H.W.J. Ten Tusscher, D. Noble, P-J. Noble, and A.V. Panfilov, *A model for human ventricular tissue*, Amer. J. Physiol. Heart Circ. Physiol., 286 (2004), pp. H1573–H1589.

[37] A. Toselli and O. Widlund, *Domain Decomposition Methods: Algorithms and Theory*, Springer, Cham, 2006.
[38] S. Zampini, *Dual-primal methods for the cardiac bidomain model*, Math. Models Methods Appl. Sci., 24 (2014), pp. 667–696.
[39] S. Zampini, *Inexact BDDC methods for the cardiac bidomain model*, in Domain Decomposition Methods in Science and Engineering XXI, Springer, Cham, 2014, pp. 247–255.