



A Numerical Study of Scalable Cardiac Electro-Mechanical Solvers on HPC Architectures

Piero Colli Franzone¹, Luca F. Pavarino^{1*} and Simone Scacchi²

¹ Department of Mathematics, University of Pavia, Pavia, Italy, ² Department of Mathematics, University of Milano, Milan, Italy

We introduce and study some scalable domain decomposition preconditioners for cardiac electro-mechanical 3D simulations on parallel HPC (High Performance Computing) architectures. The electro-mechanical model of the cardiac tissue is composed of four coupled sub-models: (1) the static finite elasticity equations for the transversely isotropic deformation of the cardiac tissue; (2) the active tension model describing the dynamics of the intracellular calcium, cross-bridge binding and myofibril tension; (3) the anisotropic Bidomain model describing the evolution of the intra- and extra-cellular potentials in the deforming cardiac tissue; and (4) the ionic membrane model describing the dynamics of ionic currents, gating variables, ionic concentrations and stretch-activated channels. This strongly coupled electro-mechanical model is discretized in time with a splitting semi-implicit technique and in space with isoparametric finite elements. The resulting scalable parallel solver is based on Multilevel Additive Schwarz preconditioners for the solution of the Bidomain system and on BDDC preconditioned Newton-Krylov solvers for the non-linear finite elasticity system. The results of several 3D parallel simulations show the scalability of both linear and non-linear solvers and their application to the study of both physiological excitation-contraction cardiac dynamics and re-entrant waves in the presence of different mechano-electrical feedbacks.

Keywords: domain decomposition preconditioners, cardiac electro-mechanics, bidomain model, scalable parallel solvers, re-entrant waves, mechano-electric feedback

OPEN ACCESS

Edited by:

Mariano Vázquez,
Barcelona Supercomputing Center,
Spain

Reviewed by:

Alfio Maria Quarteroni,
Politecnico di Milano, Italy
Daniel Hurtado,
Pontificia Universidad Católica de
Chile, Chile

*Correspondence:

Luca F. Pavarino
luca.pavarino@unipv.it

Specialty section:

This article was submitted to
Computational Physiology and
Medicine,

a section of the journal
Frontiers in Physiology

Received: 15 December 2017

Accepted: 08 March 2018

Published: 05 April 2018

Citation:

Colli Franzone P, Pavarino LF and
Scacchi S (2018) A Numerical Study
of Scalable Cardiac
Electro-Mechanical Solvers on HPC
Architectures. *Front. Physiol.* 9:268.
doi: 10.3389/fphys.2018.00268

1. INTRODUCTION

In recent years, several areas of medicine, and in particular cardiology, have undergone a cultural revolution generated by new findings that have emerged from molecular biology. This new knowledge has helped to identify, for each disease and for each patient, the specific mechanisms of the disease and the resulting medical treatments, leading to the so-called personalized medicine. For example, the use of mathematical models with parameters for the individual patient-specific characteristics could allow cardiologists to predict the effectiveness of anti-arrhythmic drug treatments or the proper installation of implantable defibrillators (see e.g., Nordsletten et al., 2011; Constantino et al., 2012; Lamata et al., 2015; Trayanova and Chang, 2016).

The spatio-temporal evolution of the electrical impulse in the cardiac tissue and the subsequent process of cardiac contraction-relaxation are quantitatively described by the cardiac electro-mechanical coupling model, which consists of the following four sub-models:

- the **static finite elasticity model** describing the deformation of cardiac tissue, derived from an anisotropic strain energy function which characterizes the passive mechanical properties of the myocardium;
- the **active tension system** of non-linear ordinary differential equations (ODEs), describing the dynamics of the intracellular calcium, cross-bridge binding and myofibril tension;
- the **anisotropic Bidomain model** of the cardiac tissue, which is a non-linear system of two partial differential equations (PDEs) of reaction-diffusion type, describing the spatio-temporal evolution of the intra- and extracellular electric potentials in the cardiac tissue;
- the **ionic membrane model** of the cardiac myocyte, a stiff system of ODEs describing the dynamics of ionic currents, gating variables, ionic concentrations and stretch-activated channels.

The theoretical and numerical challenges posed by this complex non-linear electro-mechanical model are very interesting. Indeed, the theoretical analysis of the well-posedness of the cardiac electro-mechanical coupling model is still an open problem, as well as the convergence analysis of its finite element approximation. On the numerical level, the very different space and time scales associated with the electrical and mechanical sub-models, as well as their non-linear and multiphysics interactions, make the approximation and simulation of the cardiac electro-mechanical coupling model a very demanding and expensive computational task.

In the last decade, several groups have performed cardiac computational studies based on three-dimensional electrical and electro-mechanical simulations (see Pathmanathan and Whiteley, 2009; Göktepe and Kuhl, 2010; Keldermann et al., 2010; Gurev et al., 2011; Trayanova et al., 2011; Land et al., 2012b; Nobile et al., 2012; Rossi et al., 2012; Dal et al., 2013; Sundnes et al., 2014; Favino et al., 2016). However, the computational costs required by the solution of the mathematical models describing the cardiac bioelectrical and mechanical activity are still too high to allow their use in a clinical setting. Therefore, there is a strong effort in the research community to develop effective computational tools and to speedup the simulation of the cardiac electro-mechanical activity (see e.g., Vázquez et al., 2011; Lafortune et al., 2012; Washio et al., 2013; Aguado-Sierra et al., 2015; Gurev et al., 2015; Land et al., 2015; Augustin et al., 2016).

Among the most efficient high-performance solvers for these complex cardiac models are parallel iterative methods, such as the Preconditioned Conjugate Gradient method (PCG) and Generalized Minimal Residual Method (GMRES), accelerated by proper scalable preconditioners. For the bioelectrical component modeled by the Bidomain system, several types of preconditioners have been proposed, such as Block Jacobi (BJ) preconditioners employing an incomplete LU factorization (ILU) for each block (Colli Franzone and Pavarino, 2004), other kinds of block preconditioners (Gerardo-Giorda et al., 2009; Chen et al., 2017), geometric multigrid (Sundnes et al., 2002; Weber dos Santos et al., 2004), algebraic multigrid (Plank et al., 2007; Pennacchio and Simoncini, 2009, 2011), and domain decomposition preconditioners such as Multilevel Schwarz

(Pavarino and Scacchi, 2008; Scacchi, 2008, 2011; Munteanu et al., 2009; Pavarino and Scacchi, 2011; Charawi, 2017), Neumann-Neumann and BDDC (Zampini, 2013, 2014). For a general introduction to Domain Decomposition methods we refer the interested reader to the monograph (Toselli and Widlund, 2005). More recently, the study of efficient parallel solvers and preconditioners has been extended also to cardiac electro-mechanical models (see e.g., Colli Franzone et al., 2015; Gurev et al., 2015; Pavarino et al., 2015; Augustin et al., 2016; Colli Franzone et al., 2016a,b, 2017) and to cardiac and cardiovascular flow (see e.g., Quarteroni et al., 2017a,b).

The goal of this work is to study the performance of our parallel electro-mechanical solver in three-dimensional left-ventricular simulations on two different HPC (High Performance Computing) architectures. The finite element parallel solver we have developed is based on Multilevel Additive Schwarz preconditioners accelerated by PCG for solving the discretized Bidomain system and on Newton-Krylov methods with Balancing Domain Decomposition by Constraints (BDDC) preconditioners for solving the discretized non-linear finite elasticity system. Extensive numerical simulations have shown the scalability of both linear and non-linear solvers and their effectiveness in the study of the physiological excitation-contraction cardiac dynamics and of re-entrant waves in the presence of different mechano-electrical feedbacks.

The paper is organized as follows. The main four electro-mechanical cardiac sub-models are briefly introduced in section 2 and discretized in time and space in section 3, where the main computational kernels, parallel solvers and preconditioners are also described. Section 4 contains the main results of the paper obtained in large-scale 3D simulations using high-performance parallel architectures.

2. ELECTRO-MECHANICAL CARDIAC MODELS

We consider a cardiac electro-mechanical coupling model consisting of the following four coupled sub-models; see also **Figure 1**.

2.1. Cardiac Tissue Mechanical Model

We assume a quasi-steady state regime and model the cardiac tissue as a non-linear hyperelastic material satisfying the equilibrium equation

$$\text{Div}(\mathbf{FS}) = \mathbf{0}, \quad \mathbf{X} \in \widehat{\Omega}, \quad (1)$$

with appropriate boundary conditions, where we denote by $\mathbf{x} = \mathbf{x}(\mathbf{X}, t)$ the spatial coordinates of the deformed cardiac domain $\Omega(t)$ at time t , by $\mathbf{X} = (X_1, X_2, X_3)^T$ the material coordinates of the undeformed cardiac domain $\widehat{\Omega}$, by $\mathbf{F}(\mathbf{X}, t) = \frac{\partial \mathbf{x}}{\partial \mathbf{X}}$ the deformation gradient and by $\mathbf{u}(\mathbf{X}, t) = \mathbf{x} - \mathbf{X}$ the displacement field. Following the active stress approach, the second Piola-Kirchhoff stress tensor \mathbf{S} is written as the sum of passive (*pas*), volumetric (*vol*) and active (*act*) components, i.e.,

$$\mathbf{S} = \mathbf{S}^{pas} + \mathbf{S}^{vol} + \mathbf{S}^{act}. \quad (2)$$

The passive and volumetric terms of \mathbf{S} are defined as

$$S_{ij}^{pas,vol} = \frac{1}{2} \left(\frac{\partial W^{pas,vol}}{\partial E_{ij}} + \frac{\partial W^{pas,vol}}{\partial E_{ji}} \right) \quad i, j = 1, 2, 3,$$

where $\mathbf{E} = \frac{1}{2}(\mathbf{C} - \mathbf{I})$ is the Green-Lagrange strain tensor and W^{pas} is an exponential strain energy function describing the myocardium as an hyperelastic material transversely isotropic (derived from the orthotropic law proposed in Holzapfel and Ogden, 2009; Eriksson et al., 2013)

$$W^{pas} = \frac{a}{2b} \left(e^{b(I_1-3)} - 1 \right) + \sum_{i=l,n} \frac{a_i}{2b_i} \left(e^{b_i(I_{4i}-1)^2} - 1 \right) + \frac{a_{ln}}{2b_{ln}} \left(e^{b_{ln}I_{8ln}^2} - 1 \right), \quad (3)$$

where a , b , $a_{(l,n)}$, $b_{(l,n)}$ are positive material parameters and

$$I_{4l} = \hat{\mathbf{a}}_l^T \mathbf{C} \hat{\mathbf{a}}_l, \quad I_{4n} = \hat{\mathbf{a}}_n^T \mathbf{C} \hat{\mathbf{a}}_n, \quad I_{8ln} = \hat{\mathbf{a}}_l^T \mathbf{C} \hat{\mathbf{a}}_n.$$

We did not employ an isochoric-deviatoric decomposition of the deformation gradient tensor. The volumetric term $W^{vol} = K(J-1)^2$ is a penalization term added to enforce the nearly incompressibility of the myocardium, where K is a positive bulk modulus and $J = \det \mathbf{F}$. The model is closed by imposing boundary conditions of mixed Dirichlet and traction type.

2.2. Mechanical Active Tension Model

The active tension generation model is based on calcium kinetic and myofilament dynamics. Here we consider the model proposed in Land et al. (2012a), where the active tension T_a depends on the intracellular calcium concentration Ca_i , the fiber stretch $\lambda = \sqrt{\hat{\mathbf{a}}_l^T \mathbf{C} \hat{\mathbf{a}}_l}$, the fiber stretch-rate $\frac{d\lambda}{dt}$ and auxiliary variables included in vector \mathbf{z} , i.e.,

$$\begin{cases} \frac{d\mathbf{z}}{dt} = \mathbf{R}_z \left(\mathbf{z}, Ca_i, \lambda, \frac{d\lambda}{dt} \right) \\ T_a = f_{Ta} \left(\mathbf{z}, \lambda, \frac{d\lambda}{dt} \right). \end{cases}$$

The generated active force is assumed to act only along the fiber direction, so the active Cauchy stress is

$$\sigma^{act}(\mathbf{x}, t) = T_a \mathbf{a}_l(\mathbf{x}) \otimes \mathbf{a}_l(\mathbf{x}),$$

where \mathbf{a}_l is a unit vector parallel to the local fiber direction and T_a is the active fiber stress associated to the deformed cardiac tissue. In the deformed configuration, the unit vector parallel to the local fiber direction can be written as

$$\mathbf{a}_l = \frac{\mathbf{F} \hat{\mathbf{a}}_l}{\|\mathbf{F} \hat{\mathbf{a}}_l\|} = \frac{\mathbf{F} \hat{\mathbf{a}}_l}{\sqrt{\hat{\mathbf{a}}_l^T \mathbf{C} \hat{\mathbf{a}}_l}}, \quad (4)$$

where $\hat{\mathbf{a}}_l$ is the fiber direction in the reference configuration. Then the active stress component \mathbf{S}^{act} of the second Piola-Kirchhoff tensor is given by

$$\mathbf{S}^{act} = J \mathbf{F}^{-1} \sigma^{act} \mathbf{F}^{-T} = J T_a \frac{\hat{\mathbf{a}}_l \otimes \hat{\mathbf{a}}_l}{\hat{\mathbf{a}}_l^T \mathbf{C} \hat{\mathbf{a}}_l}.$$

2.3. The Bioelectrical Bidomain Model

We denote by v , u_e , \mathbf{w} , \mathbf{c} the transmembrane potential, the extracellular potential, the gating and ionic concentrations variables on the deformed configuration and by \hat{v} , \hat{u}_e , $\hat{\mathbf{w}}$, $\hat{\mathbf{c}}$ the same quantities on reference configuration. The Bidomain model, written on the deformed configuration $\Omega(t)$ is given in its parabolic-elliptic formulation by

$$\begin{cases} c_m \frac{\partial v}{\partial t} - \text{div}(D_i \nabla(v + u_e)) + i_{ion}(v, \mathbf{w}, \mathbf{c}, \lambda) = i_{app}^i \\ -\text{div}(D_i \nabla v) - \text{div}((D_i + D_e) \nabla u_e) = i_{app}^i + i_{app}^e, \end{cases} \quad (5)$$

where c_m and i_{ion} are the membrane capacitance and ionic current per unit volume, respectively. We apply insulating boundary conditions on $\partial\Omega(t)$, i.e.,

$$\mathbf{n}^T D_i \nabla(v + u_e) = 0 \quad \text{and} \quad \mathbf{n}^T D_e \nabla u_e = 0,$$

with \mathbf{n} being the normal to $\partial\Omega(t)$. In order to satisfy the compatibility condition $\int_{\Omega(t)} (i_{app}^i + i_{app}^e) d\mathbf{x} = 0$, we choose $i_{app}^i = -i_{app}^e = i_{app}$; see e.g., Colli Franzone et al. (2014). In the Lagrangian framework, after the pull-back on the reference configuration $\hat{\Omega} \times (0, T)$, this system becomes

$$\begin{cases} c_m J \left(\frac{\partial \hat{v}}{\partial t} - \mathbf{F}^{-T} \text{Grad} \hat{v} \cdot \mathbf{V} \right) - \text{Div}(J \mathbf{F}^{-1} \hat{D}_i \mathbf{F}^{-T} \text{Grad}(\hat{v} + \hat{u}_e)) \\ + J i_{ion}(\hat{v}, \hat{\mathbf{w}}, \hat{\mathbf{c}}, \lambda) = \hat{J} i_{app}, \\ -\text{Div}(J \mathbf{F}^{-1} \hat{D}_i \mathbf{F}^{-T} \text{Grad} \hat{v}) - \text{Div}(J \mathbf{F}^{-1} (\hat{D}_i + \hat{D}_e) \mathbf{F}^{-T} \text{Grad} \hat{u}_e) = 0, \end{cases} \quad (6)$$

where $\mathbf{V} = \frac{\partial \mathbf{u}}{\partial t}$ is the rate of deformation; see Colli Franzone et al. (2016a) for the detailed derivation. These two partial differential equations (PDEs) are coupled through the reaction term i_{ion} with the ODE system of the membrane model, given in $\Omega(t) \times (0, T)$ by

$$\frac{\partial \mathbf{w}}{\partial t} - \mathbf{R}_w(v, \mathbf{w}) = 0, \quad \frac{\partial \mathbf{c}}{\partial t} - \mathbf{R}_c(v, \mathbf{w}, \mathbf{c}) = 0. \quad (7)$$

The bioelectrical system (Equations 6, 7) is completed by prescribing initial conditions on \hat{v} , \mathbf{w} , \mathbf{c} , insulating boundary conditions on \hat{u}_e , $\hat{u}_i = \hat{v} + \hat{u}_e$, and the intra- and extracellular applied current $\hat{i}_{app} = \hat{i}_{app}^i = -\hat{i}_{app}^e$. We recall that the extracellular potential \hat{u}_e is defined only up to a time dependent constant in space $R(t)$, which can be determined by choosing a reference potential. Here we select as a reference potential the average of the extracellular potential over the cardiac volume, i.e., we require $\int_{\hat{\Omega}} \hat{u}_e(\mathbf{X}, t) J(\mathbf{X}, t) d\mathbf{X} = 0$. Assuming transversely isotropic properties of the intra- and extracellular media, the conductivity tensors on the deformed configuration are given by

$$D_{i,e} = \sigma_t^{i,e} I + (\sigma_l^{i,e} - \sigma_t^{i,e}) \mathbf{a}_l \otimes \mathbf{a}_l,$$

where $\sigma_l^{i,e}$, $\sigma_t^{i,e}$ are the the intra- and extracellular conductivity coefficients measured along the fiber direction \mathbf{a}_l and any cross fiber direction, respectively. From Equation (4), it follows that the tensors $D_{i,e}(\mathbf{x}, t)$ written on the reference configuration are

$$\hat{D}_{i,e}(\mathbf{X}, t) = D_{i,e}(\mathbf{x}(\mathbf{X}, t), t) = \sigma_t^{i,e} I + (\sigma_l^{i,e} - \sigma_t^{i,e}) \frac{\mathbf{F} \hat{\mathbf{a}}_l \hat{\mathbf{a}}_l^T \mathbf{F}^T}{\hat{\mathbf{a}}_l^T \mathbf{C} \hat{\mathbf{a}}_l}. \quad (8)$$

Therefore, the equivalent conductivity tensors appearing into the bidomain model written in the reference configuration are given by

$$J\mathbf{F}^{-1}\widehat{D}_{i,e}(\mathbf{X},t)\mathbf{F}^{-T} = \sigma_t^{i,e} \mathbf{C}^{-1} + (\sigma_l^{i,e} - \sigma_t^{i,e}) \frac{\widehat{\mathbf{a}}_l \widehat{\mathbf{a}}_l^T}{\widehat{\mathbf{a}}_l^T \widehat{\mathbf{C}}_l \widehat{\mathbf{a}}_l}. \quad (9)$$

For the values of the conductivity coefficients of the Bidoman model (see Colli Franzone et al., 2016a).

2.4. The Ionic Membrane Model and Stretch-Activated Channel Currents

The ionic current in the Bidomain model (Equation 6) is given by $i_{ion} = \chi I_{ion}$, where χ is the membrane surface to volume ratio and the ionic current per unit area of the membrane surface I_{ion} is given by the sum $I_{ion}(v, \mathbf{w}, \mathbf{c}, \lambda) = I_{ion}^m(v, \mathbf{w}, \mathbf{c}) + I_{sac}$ of two terms: the ionic term $I_{ion}^m(v, \mathbf{w}, \mathbf{c})$ given by the ten Tusscher model (TP06) (ten Tusscher et al., 2004; ten Tusscher and Panfilov, 2006), available from the cellML depository (models.cellml.org/cellml), and a stretch-activated current term I_{sac} . The TP06 ionic model also specifies the functions $R_w(v, \mathbf{w})$ and $R_c(v, \mathbf{w}, \mathbf{c})$ in the ODE system Equation (Equation 7), consisting of 17 ordinary differential equations modeling the main ionic currents dynamics.

The stretch-activated current (SAC) is modeled as the sum of a non-selective and a potassium selective currents

$$I_{sac} = I_{ns} + I_{Ko},$$

as in Niederer and Smith (2007). The non-selective SAC current is defined by

$$I_{ns} = I_{ns,Na} + I_{ns,K} = g_{ns} \gamma_{sl}(\lambda) [r(v - v_{Na}) + (v - v_K)],$$

with $\gamma_{sl}(\lambda) = 10 \max(\lambda - 1, 0)$, $g_{ns} = 4.13 \cdot 10^{-3} \text{ mS/cm}^2$ and the value of r measures the relative conductance of the ions Na^+ and K^+ and determines the reversal potential v_{ns} of I_{ns} , varying the degree of expression of the ions Na^+ and K^+ . We have chosen $r = 0.2$.

The K^+ selective SAC current is defined by

$$I_{Ko} = g_{Ko} \frac{\gamma_{SL,Ko}}{1 + \exp(-(10 + v)/45)} (v - v_K),$$

where $g_{Ko} = 1.2 \cdot 10^{-2} \text{ mS/cm}^2$ and $\gamma_{SL,Ko} = 3 \max(\lambda - 1, 0) + 0.7$.

3. NUMERICAL METHODS

3.1. Space and Time Discretization

3.1.1. Domain Geometry

We consider an idealized left ventricular geometry $\widehat{\Omega} = \Omega(0)$ modeled as a truncated ellipsoid described in ellipsoidal coordinates by the parametric equations

$$\begin{cases} x = a(r) \cos \theta \cos \phi & \phi_{min} \leq \phi \leq \phi_{max}, \\ y = b(r) \cos \theta \sin \phi & \theta_{min} \leq \theta \leq \theta_{max}, \\ z = c(r) \sin \theta & 0 \leq r \leq 1. \end{cases}$$

Here $a(r) = a_1 + r(a_2 - a_1)$, $b(r) = b_1 + r(b_2 - b_1)$, $c(r) = c_1 + r(c_2 - c_1)$, and $a_1 = b_1 = 1.5$, $a_2 = b_2 = 2.7$, $c_1 = 4.4$, $c_2 = 5$ (all in cm) and $\phi_{min} = -\pi/2$, $\phi_{max} = 3\pi/2$, $\theta_{min} = -3\pi/8$, $\theta_{max} = \pi/8$. We will refer to the inner surface of the truncated ellipsoid ($r = 0$) as endocardium and to the outer surface ($r = 1$) as epicardium. Proceeding counterclockwise from epicardium to endocardium, the cardiac fibers rotate intramurally linearly with the depth, for a total amount of 120° . Considering a local ellipsoidal reference system ($\mathbf{e}_\phi, \mathbf{e}_\theta, \mathbf{e}_r$), the fiber direction $\mathbf{a}_l(\mathbf{x})$ at a point \mathbf{x} is given by $\mathbf{a}_l(\mathbf{x}) = \mathbf{b}_l(\mathbf{x}) \cos(\beta) + \mathbf{n}(\mathbf{x}) \sin(\beta)$, where

$$\begin{aligned} \mathbf{b}_l(\mathbf{x}) &= \mathbf{e}_\phi \cos \alpha(r) + \mathbf{e}_\theta \sin \alpha(r), \text{ with} \\ \alpha(r) &= \frac{2}{3}\pi(1-r) - \frac{\pi}{4}, \quad 0 \leq r \leq 1, \end{aligned}$$

$\mathbf{n}(\mathbf{x})$ is the unit outward normal to the ellipsoidal surface at \mathbf{x} and β is the imbrication angle given by $\beta = \arctan(\cos \alpha \tan \gamma)$, with $\gamma = \theta(1-r)60/\pi$.

3.1.2. Time Discretization

The time discretization of the electromechanical model is performed by the following semi-implicit splitting method, where different electrical and mechanical time steps could be used.

(a) given $v^n, \mathbf{w}^n, \mathbf{c}^n$ at time step t_n , we compute the new variables $\mathbf{w}^{n+1}, \mathbf{c}^{n+1}$ by solving the ODE system of the ionic membrane model (Equation 7) with a first order implicit-explicit (IMEX) method, i.e.,

$$\begin{cases} \frac{\mathbf{w}^{n+1} - \mathbf{w}^n}{\Delta t} - \mathbf{R}_w(v^n, \mathbf{w}^{n+1}) = 0, \\ \frac{\mathbf{c}^{n+1} - \mathbf{c}^n}{\Delta t} - \mathbf{R}_c(v^n, \mathbf{w}^{n+1}, \mathbf{c}^n) = 0; \end{cases}$$

(b) given the calcium concentration Ca_i^{n+1} , which is part of the vector of concentration variables \mathbf{c}^{n+1} , we compute the new deformed coordinates \mathbf{x}^{n+1} , providing the new deformation gradient tensor \mathbf{F}_{n+1} , by solving the variational formulation of the mechanical problem (Equation 1) and the active tension system, i.e.,

$$\begin{cases} \mathbf{z}^{n+1} = \mathbf{z}^n + \Delta t R_z \left(\mathbf{z}^{n+1}, Ca_i^{n+1}, \lambda^{n+1}, \frac{\lambda^{n+1} - \lambda^n}{\Delta t_n} \right) \\ T_a^{n+1} = fTa \left(\mathbf{z}^{n+1}, \lambda^{n+1}, \frac{\lambda^{n+1} - \lambda^n}{\Delta t_n} \right) \\ \text{Div}(\mathbf{F}_{n+1} \mathbf{S}_{n+1}) = 0; \end{cases}$$

(c) given $\mathbf{w}^{n+1}, \mathbf{c}^{n+1}, \mathbf{F}_{n+1}$ and $J_{n+1} = \det(\mathbf{F}_{n+1})$, we compute the new electric potentials v^{n+1}, u_e^{n+1} by solving the variational formulation of the Bidomain system (Equation 6) with a first order IMEX and operator splitting method, consisting of decoupling the parabolic from the elliptic equation, i.e.,

$$\begin{cases} -\text{Div}(J_{n+1} \mathbf{F}_{n+1}^{-1} \widehat{D}_i \mathbf{F}_{n+1}^{-T} \text{Grad} \widehat{v}^n) - \text{Div}(J_{n+1} \mathbf{F}_{n+1}^{-1} (\widehat{D}_i + \widehat{D}_e) \mathbf{F}_{n+1}^{-T} \text{Grad} \widehat{u}_e^n) = 0, \\ c_m J_{n+1} \left(\frac{\widehat{v}^{n+1} - \widehat{v}^n}{\Delta t} - \mathbf{F}_{n+1}^{-T} \text{Grad} \widehat{v}^n \cdot \mathbf{V}^{n+1} \right) - \text{Div}(J_{n+1} \mathbf{F}_{n+1}^{-1} \widehat{D}_i \mathbf{F}_{n+1}^{-T} \text{Grad} (\widehat{v}^{n+1} + \widehat{u}_e^{n+1})) + \\ J_{n+1} i_{ion}(\widehat{v}^n, \widehat{\mathbf{w}}^{n+1}, \widehat{\mathbf{c}}^{n+1}, \lambda^{n+1}) = J_{n+1} \widehat{i}_{app}^{n+1}. \end{cases}$$

In our simulations, we use the electrical time step size $\Delta_{et} = 0.05$ ms, and a mechanical times step five times larger, $\Delta_{mt} = 0.25$ ms. In order to approximate the convective term in the variational formulation of Equation (6), an *upwind* discretization strategy is employed. We refer to Colli Franzone et al. (2015) and Colli Franzone et al. (2016a) for more details about the numerical scheme.

3.1.3. Space Discretization

The cardiac domain is discretized with a structured hexahedral grid T_{h_m} for the mechanical model (Equation 1) and T_{h_e} for the Bidomain model (Equation 6), where T_{h_e} is a refinement of T_{h_m} , i.e., the mechanical mesh size h_m is an integer multiple of the electrical mesh size h_e . We consider the variational formulations of both mechanical and bioelectrical models and then approximate all scalar and vector fields by isoparametric Q_1 finite elements in space. In all our simulations, we employ an electrical mesh size $h_e = 0.01$ cm in order to properly resolve the sharp excitation front, while the smoother mechanical deformation allow us to use a coarse mechanical mesh of size $h_m = 0.08$ cm. The resulting electrical mesh consists of $N_\phi \times N_\theta \times N_k$ elements, whose values will be specified in each numerical test reported in the Results section.

3.2. Computational Kernels and Parallel Solvers

At each time step of the space–time discretization described above, the two main computational kernels are:

(a) the solution of a non-linear system arising from the discretization of the mechanical problem (1); to this end, we use a parallel Newton-Krylov-BDDC (NK-BDDC) solver, where the Krylov method chosen is GMRES and the BDDC preconditioner will be described in the next sections;

(b) the solution of two linear systems deriving from the discretization of the elliptic and parabolic equations in the Bidomain model (Equation 6); to this end, we use a parallel Preconditioned Conjugate Gradient (PCG) method, with Multilevel Additive Schwarz preconditioner for the very ill-conditioned elliptic system and with Block-Jacobi preconditioner for the easier parabolic system.

The parallelization of these two main computational kernels of our electro-mechanical solver is based on the parallel library PETSc (Balay et al., 2012) from the Argonne National Laboratory. All the parallel simulations have been performed on high-performance supercomputers and Linux clusters described in the Result section. For the parallel implementation of the BDDC preconditioner, see Zampini (2016).

3.3. Multilevel Additive Schwarz Preconditioners

We now describe the Multilevel Additive Schwarz preconditioner employed in the PCG solution of the elliptic kernel (b) associated with the Bidomain system. Let Ω^k , $k = 0, \dots, \ell - 1$ be a family of ℓ nested triangulations of Ω , with finer mesh sizes from level 0 to $\ell - 1$, and let A^k be the matrix obtained by discretizing the second equation of Equation (6) on Ω^k ; we have $A^{\ell-1} = A_{bid}$, where A_{bid} is the stiffness matrix related to the elliptic equation of Equation (6) discretized on the fine mesh. Denote by R^k the restriction operators from $\Omega^{\ell-1}$ to Ω^k . We decompose each grid Ω^k , for $k = 1, \dots, \ell - 1$, into N_k overlapping subgrids Ω_i^k for $i = 1, \dots, N_k$, such that the overlap size δ^k at level $k = 1, \dots, \ell - 1$ equals the mesh size h^k of the grid Ω^k . We denote by R_i^k the restriction operator from $\Omega^{\ell-1}$ to Ω_i^k and define $A_i^k := R_i^k A^k R_i^{kT}$. The Multilevel Additive Schwarz (MAS(ℓ)) preconditioner is given by

$$B_{MAS}^{-1} := R^{0T} A^{0-1} R^0 + \sum_{k=1}^{\ell-1} \sum_{i=1}^{N_k} R_i^{kT} A_i^{k-1} R_i^k.$$

The resulting PCG algorithm has a convergence rate independent of the number of subdomains N_k (scalability), the number of levels ℓ (multilevel optimality), while it depends linearly on the ratio H_k/h_k of subdomain to element size on level k (optimality); see Pavarino and Scacchi (2008), Scacchi (2008), and Pavarino and Scacchi (2011) for the theoretical details.

3.4. Iterative Substructuring, Schur Complement System and BDDC Preconditioners

We then turn to the BDDC preconditioner used in the mechanical computational kernel (a) above, i.e., the Jacobian system arising at each iteration of the Newton method applied to the non-linear elasticity system (Equation 1). For sake of simplicity, in the following sections we will denote the reference domain by Ω instead of $\widehat{\Omega}$. We consider a decomposition of Ω into N non-overlapping subdomains Ω_i of diameter H_i

$$\Omega = \bigcup_{i=1}^N \Omega_i,$$

and set $H = \max H_i$. We first reduce the Jacobian system

$$Kx = f, \quad (10)$$

arising at each Newton step of the mechanical solver, to the interface

$$\Gamma := \left(\bigcup_{i=1}^N \partial \Omega_i \right) \setminus \partial \Omega,$$

by eliminating the interior degrees of freedom (dofs) associated with the basis functions having support in each subdomain's interior and obtaining the Schur complement system

$$S_\Gamma x_\Gamma = g_\Gamma. \tag{11}$$

Here $S_\Gamma = K_{\Gamma\Gamma} - K_{\Gamma I} K_{II}^{-1} K_{I\Gamma}$ and $g_\Gamma = f_\Gamma - K_{\Gamma I} K_{II}^{-1} f_I$ are obtained from the global system (Equation 10) by reordering the finite element basis functions into interior (denoted by the subscript I) and interface (denoted by the subscript Γ) basis functions

$$\begin{pmatrix} K_{II} & K_{I\Gamma} \\ K_{\Gamma I} & K_{\Gamma\Gamma} \end{pmatrix} \begin{pmatrix} x_I \\ x_\Gamma \end{pmatrix} = \begin{pmatrix} f_I \\ f_\Gamma \end{pmatrix}. \tag{12}$$

The Schur complement system (Equation 11) is solved iteratively by the GMRES method, where only the action of S_Γ on a given vector is required and S_Γ is never explicitly formed; instead, a block diagonal problem on the interior dofs is solved while computing the matrix vector product. Once the interface solution x_Γ has been determined, the interior dofs x_I can be found by solving local problems on each subdomain Ω_i . We then solve by the GMRES method the preconditioned Schur complement system

$$M_{\text{BDDC}}^{-1} S_\Gamma x_\Gamma = M_{\text{BDDC}}^{-1} g_\Gamma, \tag{13}$$

where M_{BDDC}^{-1} is the BDDC preconditioner, defined in Equation (17) below.

Balanced Domain Decomposition by Constraints (BDDC) preconditioners were introduced by Dohrmann (2003) and first analyzed by Mandel and Dohrmann (2003) and Mandel et al. (2005). In these methods all local and coarse problems are treated additively and the user selects the so-called primal continuity constraints across the subdomains' interface. Usual choices of primal constraints are e.g., point constraints at subdomain vertices and/or averages or moments over subdomains edges or faces. Closely related to BDDC methods are FETI and FETI-DP algorithms, as well as the previous balancing Neumann-Neumann methods; for more details, we refer the interested reader to the domain decomposition monograph (Toselli and Widlund, 2005, Ch. 6). See also Brands et al. (2008) and Klawonn and Rheinbach (2010) for FETI-DP algorithms applied in other fields of computational biomechanics.

3.4.1. Subspace Decompositions

Let V be the Q_1 finite element space for displacements and $V^{(i)}$ be the local finite element space defined on subdomain Ω_i that vanish on $\partial\Omega_i \cap \partial\Omega_D$. This local space can be split into a direct sum of its interior (I) and interface (Γ) subspaces $V^{(i)} = V_I^{(i)} \oplus V_\Gamma^{(i)}$ and we can define the associated product spaces as

$$V_I := \prod_{i=1}^N V_I^{(i)}, \quad V_\Gamma := \prod_{i=1}^N V_\Gamma^{(i)}.$$

While our finite element approximations are continuous across the interface Γ , the functions of V_Γ are generally discontinuous across Γ . We then define the subspace

$$\widehat{V}_\Gamma := \{\text{functions of } V_\Gamma \text{ that are continuous across } \Gamma\},$$

and the intermediate subspace

$$\widetilde{V}_\Gamma := V_\Delta \oplus \widehat{V}_\Pi,$$

defined by further splitting the interface dofs (denoted by the subscript Γ) into primal (subscript Π) and dual (subscript Δ) dofs. Here:

(a) the subspace \widehat{V}_Π consists of functions which are continuous at selected *primal* variables. These can be e.g., the subdomain basis functions associated with subdomains' vertices and/or edge/face basis functions with constant values at the nodes of the associated edge/face. A change of basis can be performed so that each primal variable correspond to an explicit dof.

(b) the subspace $V_\Delta = \prod_{i=1}^N V_\Delta^{(i)}$ is the product space of the local subspaces $V_\Delta^{(i)}$ of *dual* interface functions that vanish at the primal dofs.

3.4.2. Restriction and Scaling Operators

The definition of our dual-primal preconditioners require also the following restriction and interpolation operators, associated with boolean matrices (with $\{0, 1\}$ elements):

$$\begin{aligned} R_{\Gamma\Delta} : \widetilde{V}_\Gamma &\longrightarrow V_\Delta, & R_{\Gamma\Pi} : \widetilde{V}_\Gamma &\longrightarrow \widehat{V}_\Pi, \\ R_\Delta^{(i)} : V_\Delta &\longrightarrow V_\Delta^{(i)}, & R_\Pi^{(i)} : \widehat{V}_\Pi &\longrightarrow \widehat{V}_\Pi^{(i)}, \end{aligned} \tag{14}$$

where $\widehat{V}_\Pi^{(i)}$ is the local *primal* subspace. Moreover, we define the pseudo-inverse counting functions $\delta_i^\dagger(x)$, which are defined at each dof x on the interface of subdomain Ω_i by

$$\delta_i^\dagger(x) := \frac{1}{\mathcal{N}_x}, \tag{15}$$

with \mathcal{N}_x the number of subdomains sharing x . We finally define scaled local restriction operators $R_{D,\Delta}^{(i)}$ by scaling by δ_i^\dagger the only nonzero element of each row of $R_\Delta^{(i)}$. We then define the scaling matrix

$$R_{D,\Gamma} := \text{the direct sum } R_{\Gamma\Pi} \oplus R_{D,\Delta}^{(i)} R_{\Gamma\Delta}. \tag{16}$$

3.4.3. Choice of Primal Constraints

The efficiency of BDDC (and more in general dual-primal) preconditioners is strongly dependent of the choice of primal constraints. The simplest choice of selecting the subdomains vertices as primal dofs is not always sufficient to obtain scalable and fast preconditioners. Therefore, richer (and computationally more expensive) primal sets have been developed in order to obtain faster preconditioners. These stronger preconditioners are based on larger coarse problems employing also edge and/or face based primal dofs, see e.g., Toselli and Widlund (2005).

3.4.4. Matrix Form of the BDDC Preconditioner

Analogously to the dual-primal splitting introduced before, we partition the local dofs into interior (I), dual (Δ), and primal (Π) dofs, so that the local stiffness matrix $K^{(i)}$ associated to subdomain Ω_i can be written as

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

The BDDC preconditioner is then defined as

$$M_{\text{BDDC}}^{-1} = R_{D,\Gamma}^T \tilde{S}_{\Gamma}^{-1} R_{D,\Gamma}, \quad (17)$$

where the scaled restriction matrix $R_{D,\Gamma}$ has been defined in Equations (14, 16), and

$$\tilde{S}_{\Gamma}^{-1} = R_{\Gamma\Delta}^T \left(\sum_{i=1}^N \begin{bmatrix} 0 & R_{\Delta}^{(i)T} \end{bmatrix} \begin{bmatrix} K_{II}^{(i)} & K_{\Delta I}^{(i)T} \\ K_{\Delta I}^{(i)} & K_{\Delta\Delta}^{(i)} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ R_{\Delta}^{(i)} \end{bmatrix} \right) R_{\Gamma\Delta} + \Phi S_{\Pi\Pi}^{-1} \Phi^T. \quad (18)$$

The first term in Equation (18) represent the sum of local problems on each subdomain Ω_i , with Neumann data on the local dual dofs and with zero Dirichlet data on the local primal dofs. The second term in Equation (18) represents a coarse problem for the primal variables involving the coarse matrix

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

and a matrix Φ mapping primal to interface dofs

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

The columns of Φ are associated with coarse basis functions defined as the minimum energy extension into the subdomains with respect to the original bilinear form and subject to the chosen set of primal constraints.

For compressible linear elasticity problems it can be shown that the BDDC algorithm is scalable and quasi-optimal, satisfying a condition number bound (see e.g., Toselli and Widlund, 2005, Ch. 6.4) as

$$\text{cond}(M_{\text{BDDC}}^{-1} S_{\Gamma}) \leq C \left(\frac{H}{h} \right) \left(1 + \log \frac{H}{h} \right)^2,$$

with $C(\frac{H}{h_m}) = \alpha$ constant if the primal space is sufficiently rich, while $C(\frac{H}{h}) = \alpha \frac{H}{h}$ if the primal space is the minimal one spanned by the dofs associated with the subdomain vertices. We recall that H is the characteristic subdomain size and $h = h_m$ is the characteristic mechanical mesh size defined in section 3.1. We could not prove a similar bound for the convergence rate of our non-symmetric NK-BDDC preconditioned operator, since our complex non-linear elasticity problem (Equation 1) involves an exponential strain energy function. Nevertheless, the numerical results presented in the next section suggests that such a bound holds also for our operator and demonstrate the effectiveness and scalability of the NK-BDDC method.

4. RESULTS

In this section, we report the results of several 3D parallel simulations with our electro-mechanical Bidomain solver, using two HPC architectures:

- the Marconi-A2 supercomputer of the Cineca Lab (<http://www.hpc.cineca.it/hardware/marconi>), an Intel OmniPath cluster with 3,600 nodes, each with 68 1.40 GHz Intel Xeon Phi 7250 Knights Landing (KNL) cores and 16 GB/node, for a total 244,800 cores;
- the Mira BG/Q supercomputer of the Argonne National Lab (<https://www.alcf.anl.gov/mira>), an IBM BG/Q machine with 49,152 nodes, each with 16 1.60 GHz PowerPC A2 cores and 16 GB/node, for a total 786,432 cores.

4.1. Test 1: Double Reentry Simulation With the Electro-Mechanical Bidomain Model (Figures 2, 3)

We start by studying the performance of our electro-mechanical Bidomain solver on a closed ellipsoidal ventricular geometry during a double reentry dynamics initiated by an S1–S2 protocol.

Figure 2 shows the snapshots of the transmembrane potential and mechanical deformation time evolution every 50 ms, computed on 256 KNL processors of Marconi-A2. At each time instant, we report the epicardial lateral view (top panel) and selected horizontal and vertical transmural sections (bottom panel). After three S1 stimulations applied at the apex every 500 ms (not shown), an S2 cross-gradient stimulation (visible as a vertical strip in the $t = 0$ panel) is applied 280 ms. after the last S1 stimulus, and this instant is taken as the reference time $t = 0$ ms for this simulation. Two counter-rotating scroll waves are generated by the S2 stimulus, with transmural filaments located near the apex and rotation period of about 250 ms (see the panels $t = 0, 250, 500$ ms). The lateral epicardial view of the upper panels shows mostly one of the two scroll waves, but the second almost-symmetric one is visible in the transmural sections of the lower panels.

This reentry dynamics is visible also in **Figure 3** that reports the waveforms at epicardial sites P1, P2, P3 (shown in **Figure 3A**) of the transmembrane potential V (**Figure 3B**), extracellular potential u_e (**Figure 3C**), fiber stretch λ (**Figure 3D**), active tension T_a (**Figure 3E**), intracellular calcium concentration Ca_i (**Figure 3F**).

4.2. Test 2: Weak Scalability of the Elliptic Bidomain - TP06 Solver (Figures 4, 5)

Figures 4, 5 (left columns) report the results of weak scalability tests on MIRA BG/Q for the elliptic solver (PCG-MAS(4)) required by the bioelectrical Bidomain - TP06 model on a half ellipsoidal domain representing an idealized half left ventricle. The number of processors is increased from 1K to 163K cores of the Mira BG/Q supercomputer of the Argonne National Lab. **Figure 4A1** reports the condition number (blue), iteration counts (red), solution times (yellow) of the PCG - MAS(4) solver. Both a fixed half ellipsoidal domain (**Figure 4A1**, top plot) and an increasing ellipsoidal domain (**Figure 4A1**, bottom plot) are considered, where in both cases the local meshsize (hence the local problem size on each processor) is kept fixed at $H/h = 16$. The results clearly show the very good scalability of the PCG - MAS(4) solver, since all quantities are bounded from above as the processor count is increased from 1K to 163K cores (a factor

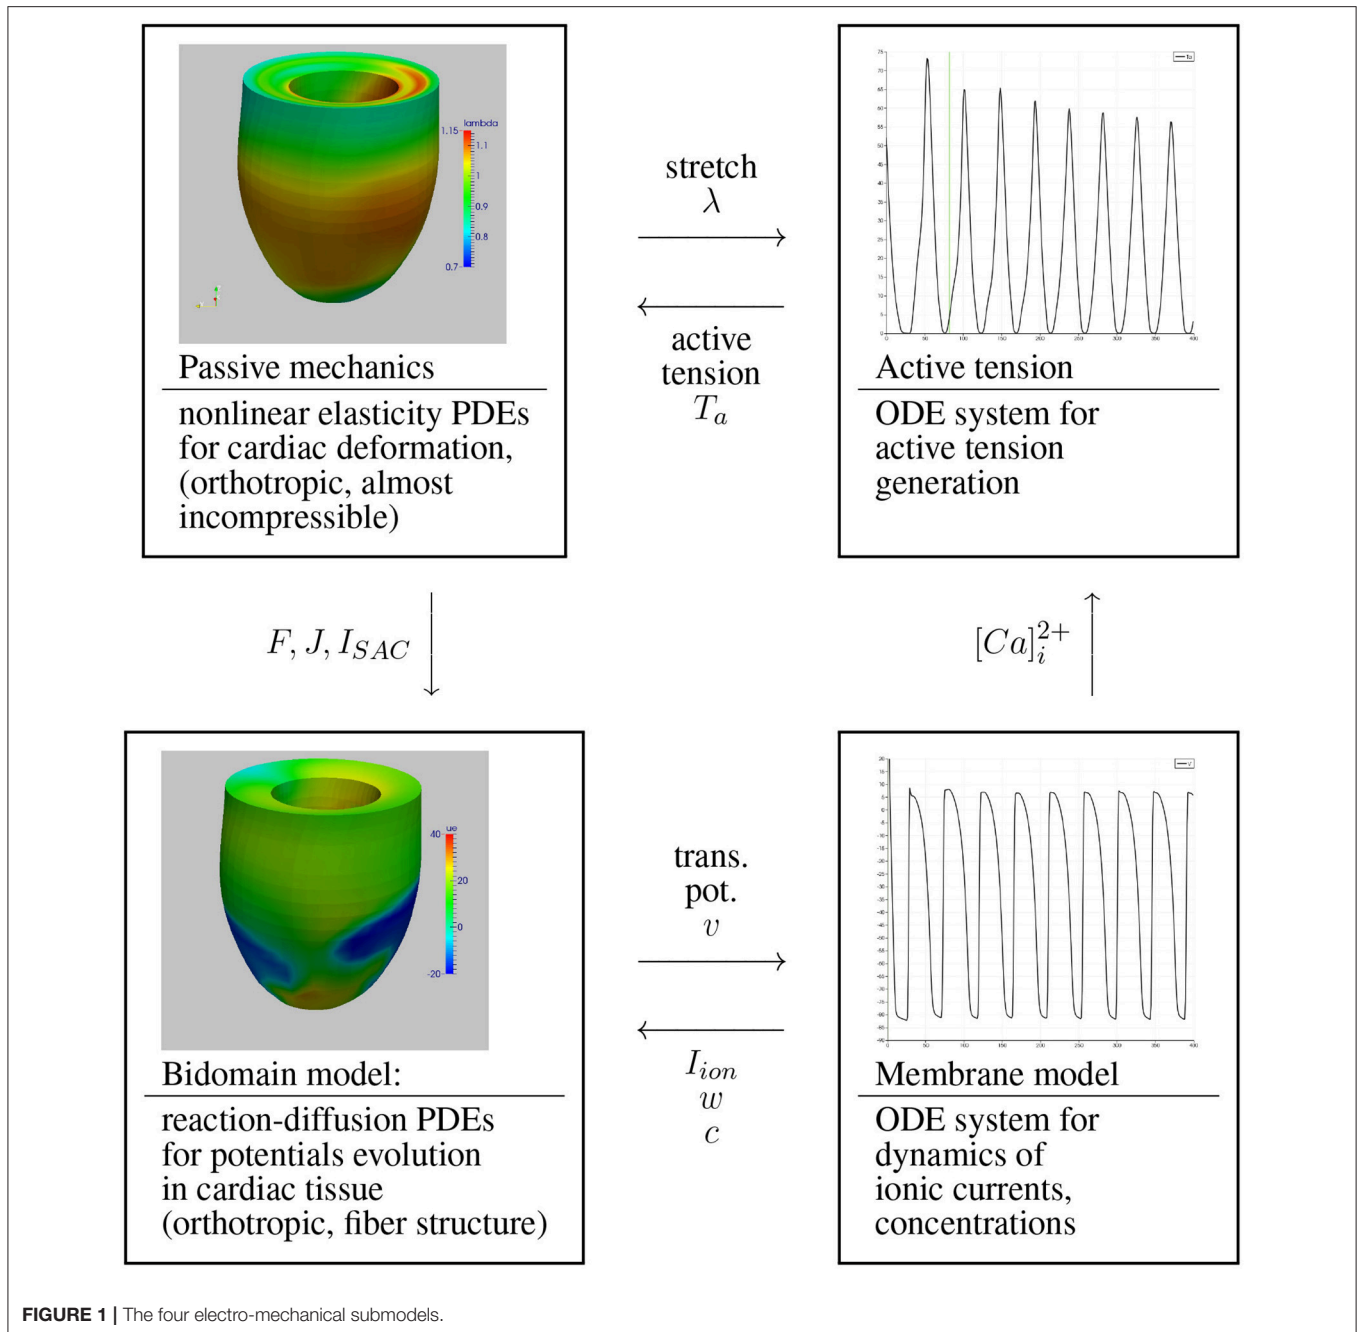


FIGURE 1 | The four electro-mechanical submodels.

163) and therefore the global problem size increases from about $O(10^6)$ to $O(10^8)$ degrees of freedom. In particular, we remark that in spite of this problem size increase of a factor 163, the CPU times are almost constant in the case of an increasing half ellipsoid (**Figure 4A1**, bottom plot) or increase by only a factor 2–3 in the case of a fixed half ellipsoid (**Figure 4A1**, top plot), while being almost constant between 16K and 128K cores.

Analogously, **Figures 4, 5** (right columns) report the results of weak scalability tests on Marconi - A2 for the elliptic solver (PCG-MAS(4)) and also the non-linear mechanical solver (NK-BDDC), described in section 4.3 below. As before, the results

clearly show the very good scalability of the PCG-MAS(4) solver, since all quantities associated with the elliptic solver are bounded from above.

In order to study more in detail the weak scalability test on a fixed half ellipsoid (**Figure 4A1**, top plot), we report in **Figures 5A1, B1, C1** the percent summary (given by the LogView PETSc subroutine) of the main PETSc functions called by the PCG - MAS(4) elliptic solver. These PETSc functions, shown in the legend of each plot, range from inner products (VecTDoc) and vector norms (VecNorm) to the whole PCG solver (KSPSolve) and application of

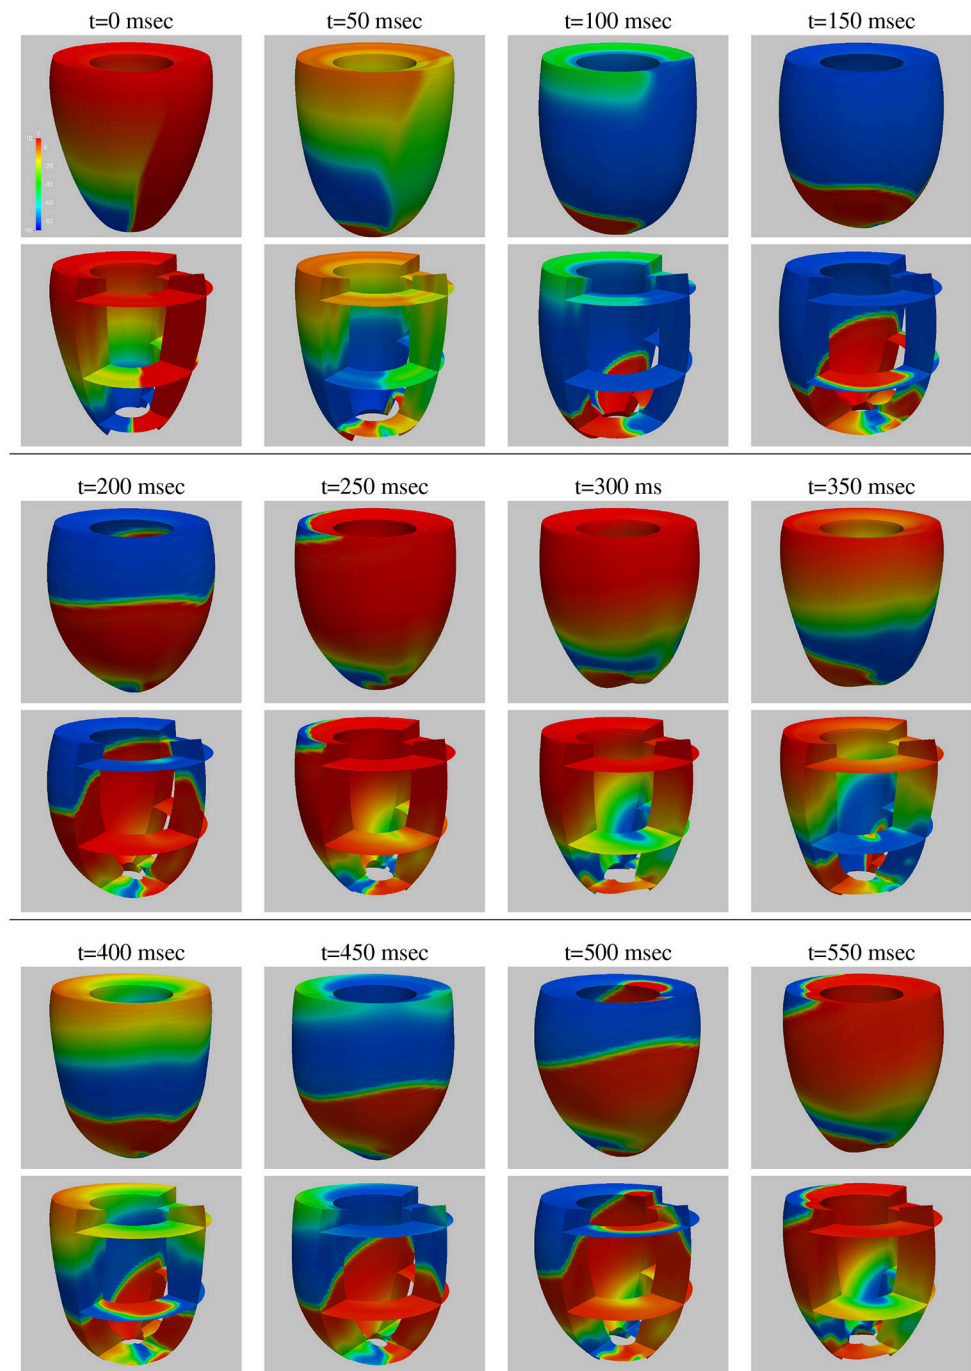
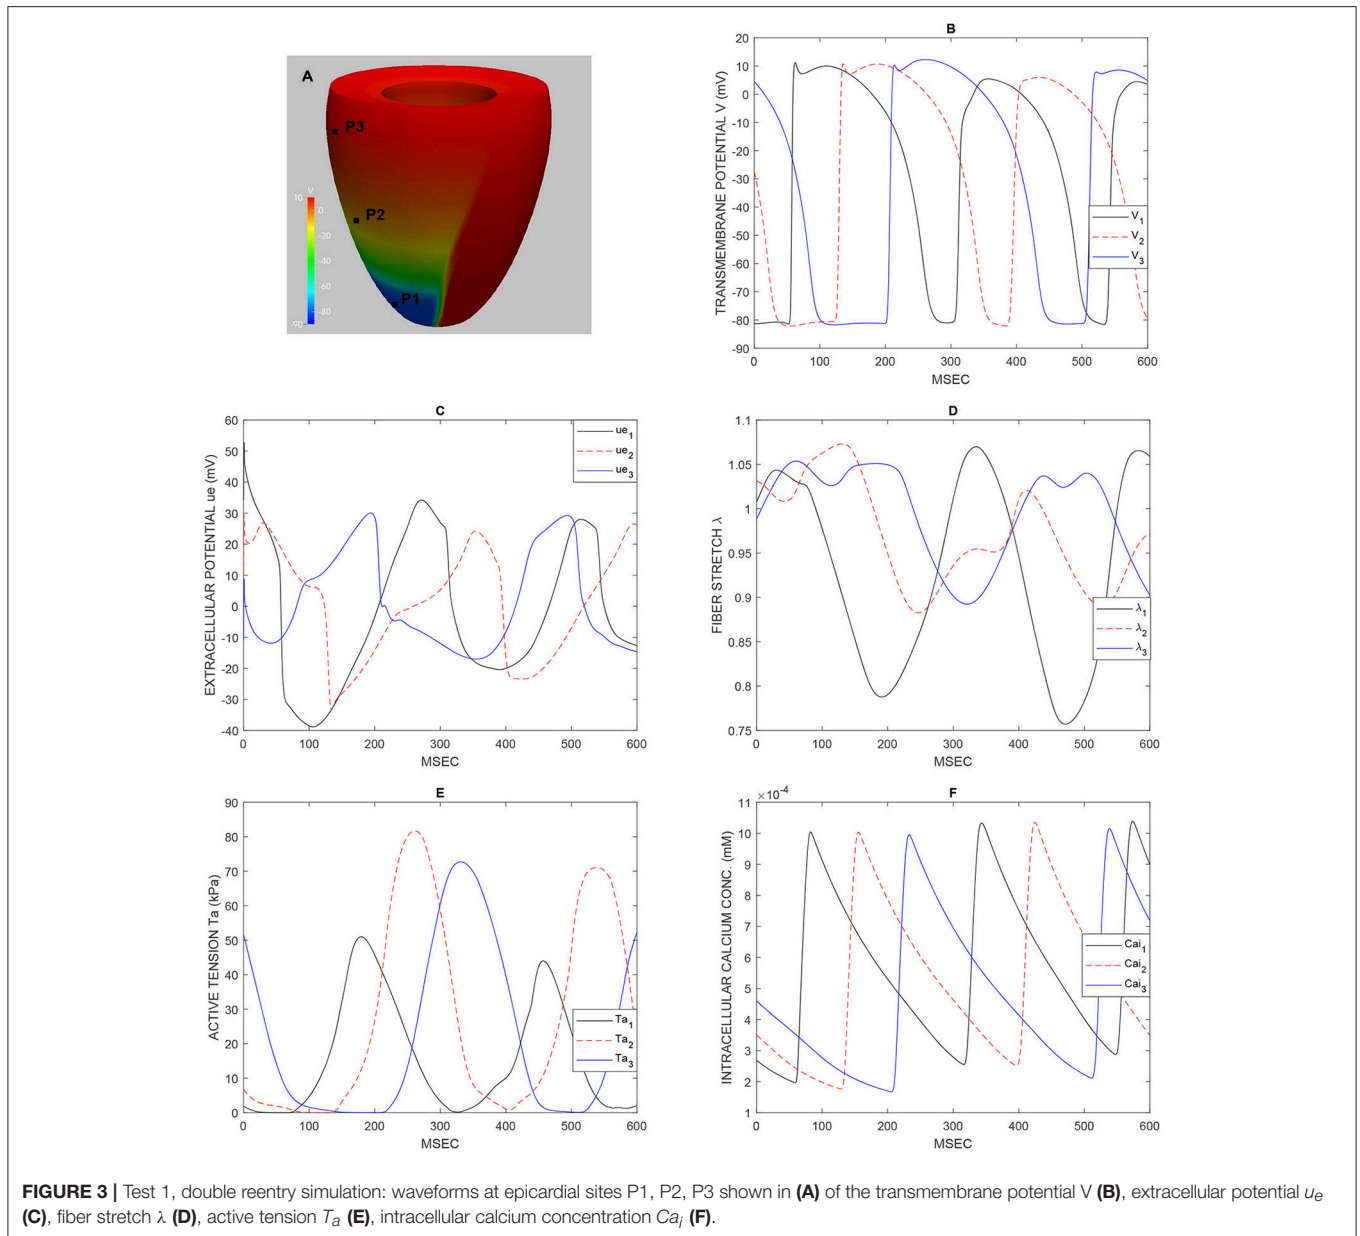


FIGURE 2 | Test 1, double reentry simulation: snapshots (every 50 ms) of the transmembrane potential and mechanical deformation time evolution. At each time instant, we report the epicardial view (**Top**) and selected horizontal and vertical transverse sections (**Bottom**).

the MAS(4) preconditioner (PCApply). In particular, we report the percent of: CPU time as a fraction of the KSPSolve time (**Figure 5A1**), flops (**Figure 5B1**), messages (**Figure 5C1**). When one of these PETSc functions has a negligible percentage, the corresponding legend shows it equal to 0. After an initial increase in some cases, all

reported quantities are very scalable up to 64K cores, and most up to 163K cores, except the VecTDot percent of flops (in **Figure 5B1**). As expected, the percentage of time (**Figure 5A1**) and flops (**Figure 5B1**) are dominated by the PCG solver (KSPSolve), followed by matrix multiplications (MatMult) and inner products (VecTDot). The percentage

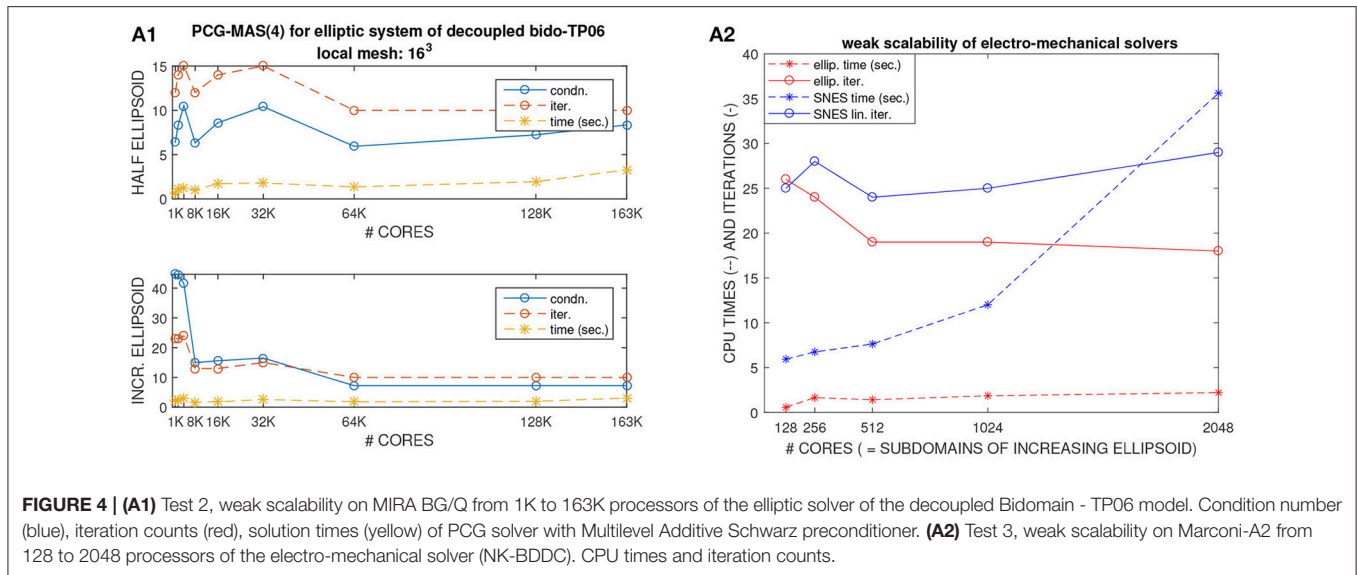


of messages (Figure 5C1) is dominated by vector scattering (VecScatterBegin), matrix multiplications (MatMult) and PCG (KSPSolve).

4.3. Test 3: Weak Scalability of the Electro-Mechanical Solver (Figures 4, 5)

We then study the weak scalability of our electro-mechanical solver from 128 to 2048 KNL processors of Marconi-A2, in particular of the two main computational kernels: the non-linear mechanical solver (NK-BDDC) and the linear elliptic Bidomain solver (PCG - MAS(4)). Figure 4A2 reports the CPU times and iteration counts for both solvers, while Figures 5A2,B2,C2 reports the percent summary of the main PETSc functions called by the electro-mechanical solver.

In this weak scaling test, the local meshsize (hence the local problem size on each processor) is kept fixed at $H/h = 16$, while the global problem size grows proportionally to the processor count by assigning one subdomain to each processor. Hence, the computational domain consists of increasing portions or an ellipsoidal domain. The results in Figure 4A2 clearly show the very good scalability of the PCG - MAS(4) elliptic linear solver, since both its CPU times and iteration counts are bounded from above as the processor count is increased to 2,048 cores. On the other hand, the timings of the non-linear SNES solver are not scalable beyond 512 processors, even if the iteration counts are. This is due to the non-scalability of the coarse solver (Mumps) employed in the BDDC preconditioner.



In order to study more in detail this scalability test, we report in **Figures 5A2,B2,C2** the percent summary (given by the LogView PETSc subroutine) of the main PETSc functions called by the electro-mechanical solver. These PETSc functions, shown in the legend of each plot, range from inner products (VecTDoc) and vector norms (VecNorm) to the linear solvers (KSPSolve) and preconditioner applications (PCApply) required by both the linear (PCG-MAS(4)) and non-linear (NK-BDDC) solvers. In particular, we report the percent of: CPU time (**Figures 5A2**), flops (**Figure 5B2**) and messages (**Figure 5C2**). When one of these PETSc functions has a negligible percentage, the corresponding legend shows it equal to 0). All reported percentages are very scalable, showing quite flat plots, except the time percentages (**Figure 5A2**), where the KSPSolve and PCApply percentages grow considerably beyond 512 cores, due mostly to the growth of MatSolve and PCSetUp, which we know already from **Figures 5A1,A2** are due to the nonscalable direct coarse solve (Mumps) of the BDDC preconditioner called by the non-linear SNES solver. As expected, the percentage of time (**Figure 5A2**) and flops (**Figure 5B2**) are dominated by the PCG solver (KSPSolve), followed by PCApply and MatSolve. The percentage of messages (**Figure 5C2**) is dominated by vector scattering (VecScatterBegin), matrix multiplications (MatMult) and linear solves (KSPSolve).

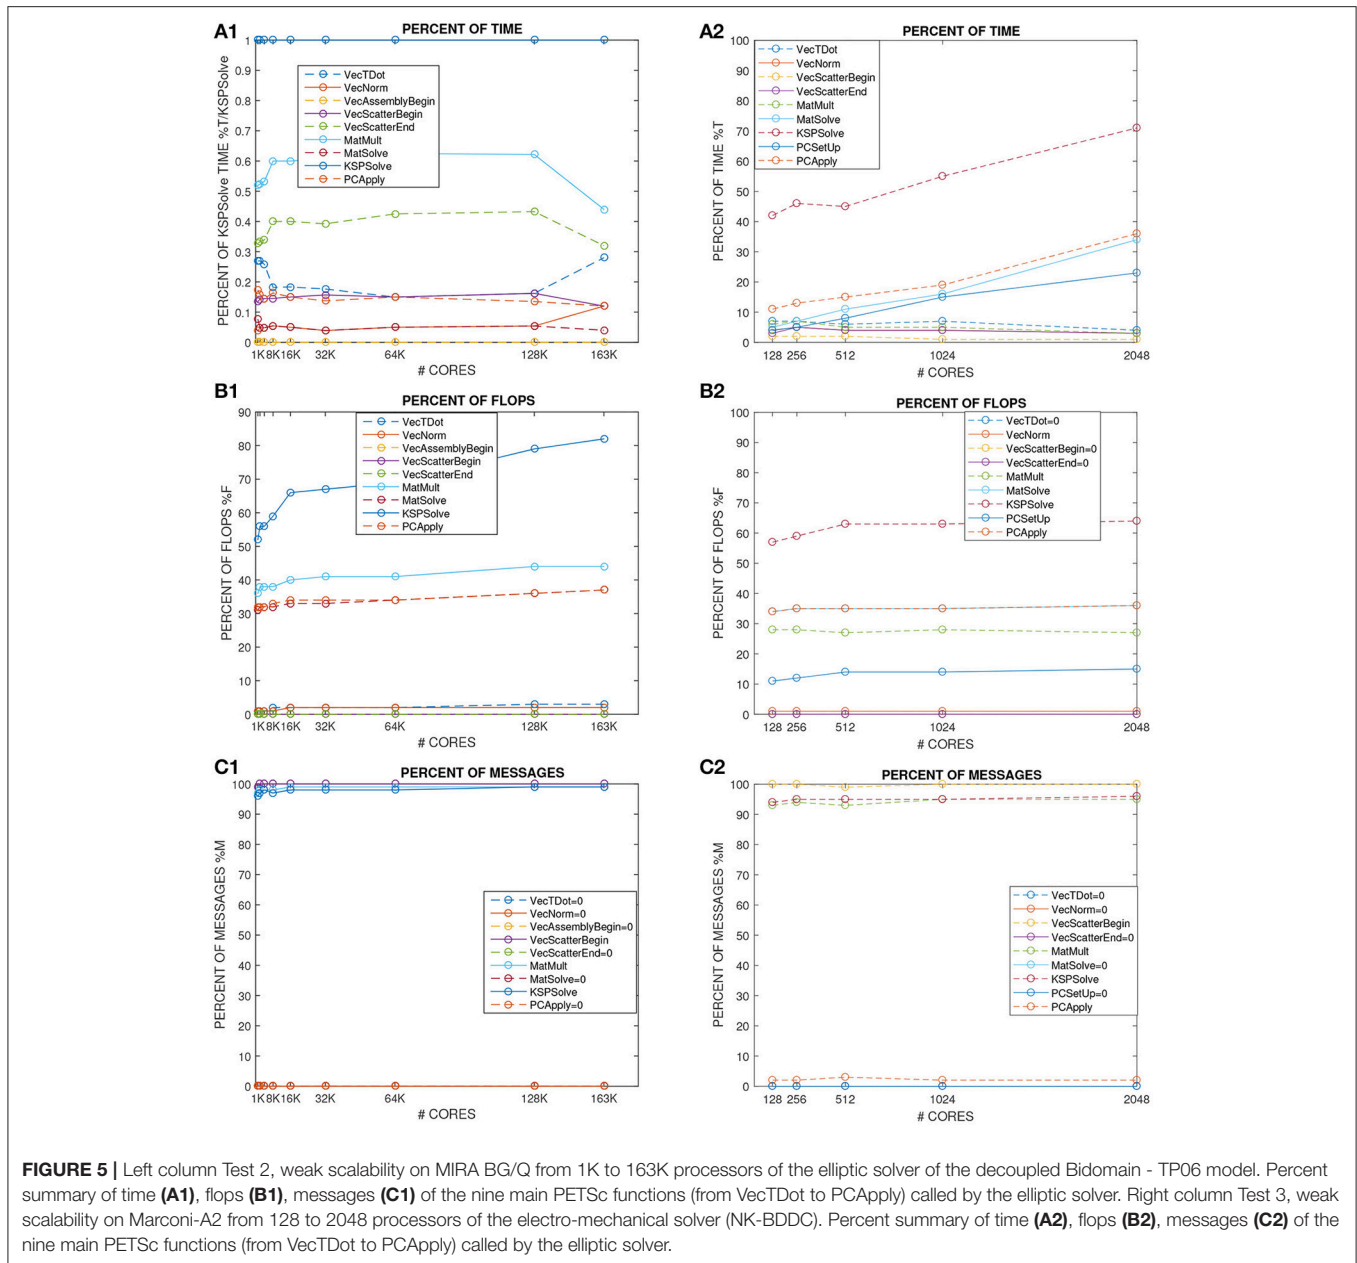
4.4. Test 4: Strong Scalability of the Non-linear Electro-Mechanical Bidomain Solver (Figures 6, 7)

Figure 6 reports the results of strong scalability tests on Marconi-A2 for the non-linear electro-mechanical Bidomain model on an ellipsoidal domain during the time interval [0 100] ms. We study the time evolution of CPU times and iterations of the two main computational kernels of our electro-mechanical model: the non-linear mechanical solver (NK-BDDC) and linear Bidomain solver

(PCG - MAS(3) for the elliptic solve and PCG-BJ for the parabolic solve).

The global mesh size is fixed to $384 \times 192 \times 48$ finite elements while the number of processors is increased from $32 = 8 \times 4 \times 1$ (with local mesh $48 \times 48 \times 48$) to $256 = 16 \times 8 \times 2$ (with local mesh $24 \times 24 \times 24$). **Figure 6A** shows the timings of the NK-BDDC solver: after an initial superlinear speedup from 32 to 64 cores, the timings still reduce when going to 128 and 256 cores but with worse speedups (see also **Figure 7A**) and start to increase at 512 cores or more (not shown). **Figure 7B** shows the number of Newton iterations for each NK-BDDC solve, which remain constant at 4 iterations independently of the number of processors. **Figure 7C** reports the cumulative GMRES iterations for each NK-BDDC mechanical solve, which increase in time since the Jacobian mechanical system becomes increasingly ill-conditioned due to the spreading of the electrical activation front and subsequent mechanical contraction. The number of iterations is reduced when going from 64 to 128 and to 256 cores, but unexpectedly in the 32 core test we got the lowest iteration counts after 20 ms. **Figure 7D** shows that the number of PCG iterations for each Bidomain elliptic solve are almost constant independently of the number of processors used. The timings of each Bidomain elliptic (**Figure 7E**) and parabolic (**Figure 7F**) solve show a reduction when the number of processors is increased, but with reduced speedup when using 256 cores or more.

As before, we now study in **Figure 7** the percent summary (given by the LogView PETSc subroutine) of the main PETSc functions in this strong scaling test for the electro-mechanical solver. We report the percent of: flops (**Figure 7C**), CPU time (**Figure 7D**), messages (**Figure 7E**), reductions (**Figure 7F**). Again we find quite flat plots, except the time percentages (**Figure 7D**), where the the KSPSolve percentage grows considerably due mostly to the growth of PCApply and MatSolve, which again we attribute mostly to the nonscalable direct coarse solve (Mumps) of the BDDC preconditioner called by the



non-linear SNES solver. The percentage of time (Figure 7D), flops (Figure 7C) and reductions (Figure 7F) are dominated by the PCG solver (KSPSolve), but in Figure 7C the percent of flops of KSPSolve and PCSetUp decreases when the processor count increases, while the percentages of MatSolve, PCApply and MatMult increase. The percentage of messages (Figure 7E) are dominated by vector scattering (VecScatterBegin), linear solves (KSPSolve) and matrix multiplications (MatMult).

5. DISCUSSION

We have developed a high-performance parallel solver for cardiac electro-mechanical 3D simulations. After numerical

discretization in space with Q_1 finite elements and IMEX operator splitting finite differences in time, the main computational kernels at each time step require: (a) the solution of a non-linear system deriving from the discretization of the cardiac mechanical problem (1) by a parallel Newton-Krylov-BDDC (NK-BDDC) solver, where the Krylov method chosen is GMRES; (b) the solution of the two linear systems deriving from the discretization of the elliptic and parabolic equations in the Bidomain model (Equation 6) by a parallel PCG method with Multilevel Additive Schwarz and Block-Jacobi preconditioners, respectively. The parallelization of our solver has been based on simulations have been performed on the parallel library PETSc (Balay et al., 2012) from the Argonne

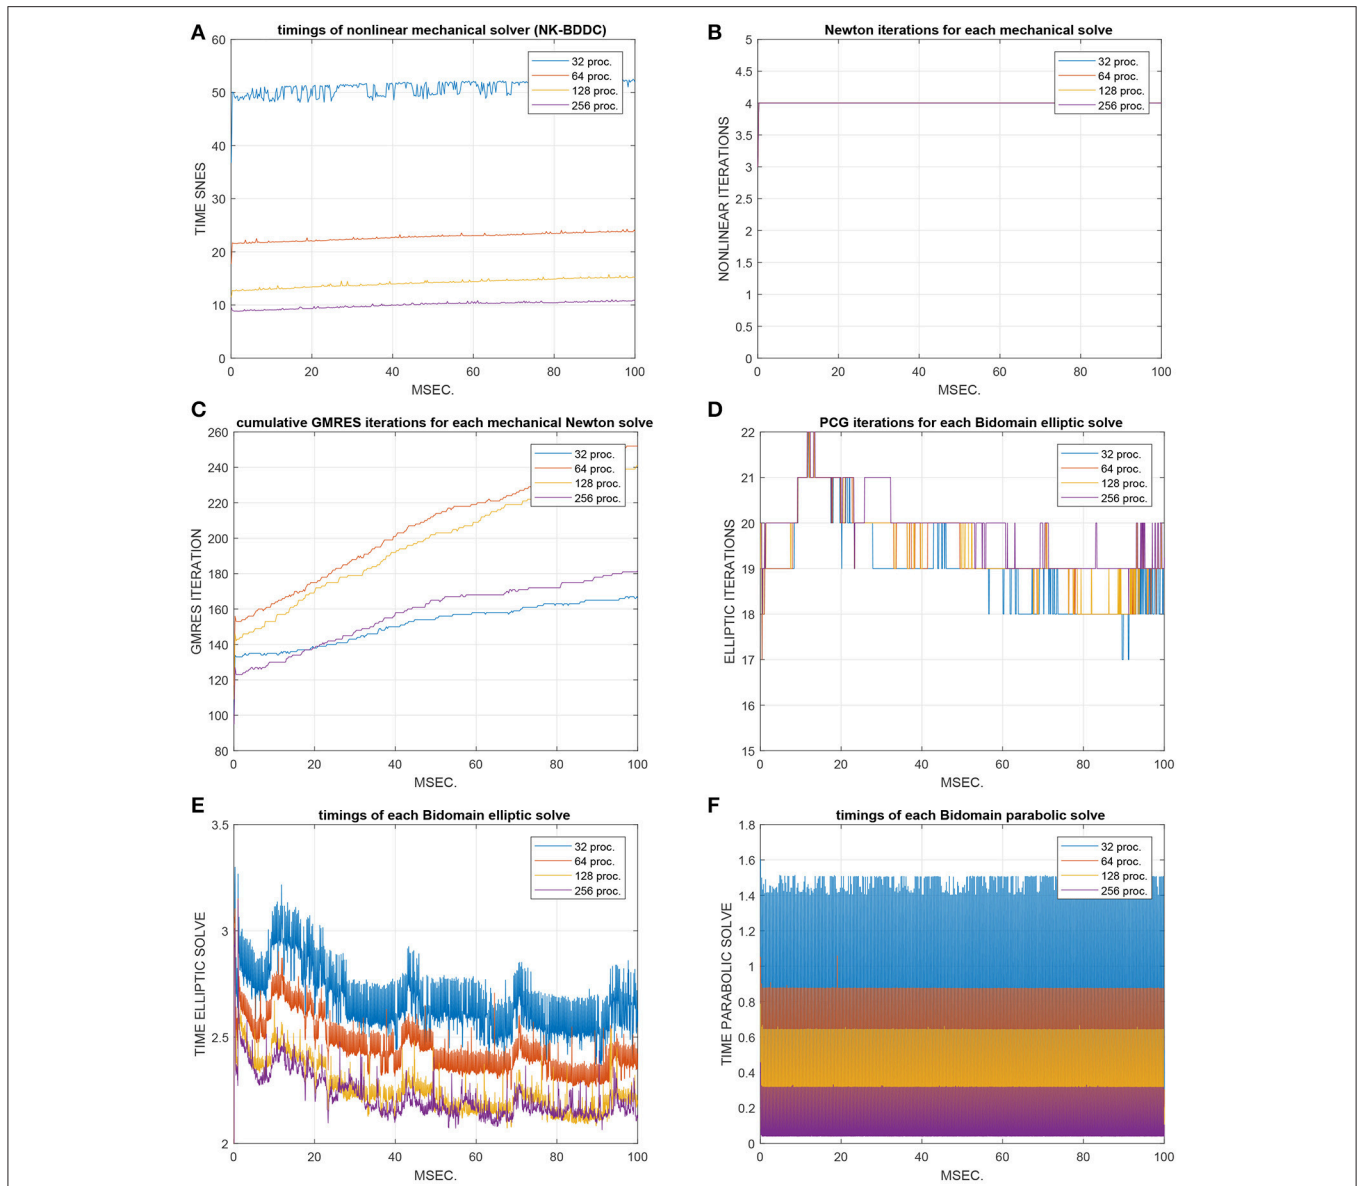


FIGURE 6 | Test 4: time evolution over the [0 100] ms interval of CPU times and iterations of the nonlinear mechanical solver (NK-BDDC) in strong scalability tests from 32 to 256 processors of Marconi-A2. **(A)** timings of NK-BDDC solver. **(B)** Newton iterations for each NK-BDDC solve. **(C)** cumulative GMRES iterations for each NK-BDDC solve. **(D)** PCG iterations for each Bidomain elliptic solve. **(E)** timings of each Bidomain elliptic solve. **(F)** timings of each Bidomain parabolic solve.

National Laboratory and large-scale 3D simulations have been run on high-performance supercomputers.

We have investigated the performance of the parallel electro-mechanical solver in both physiological excitation-contraction cardiac dynamics and pathological situations characterized by re-entrant waves.

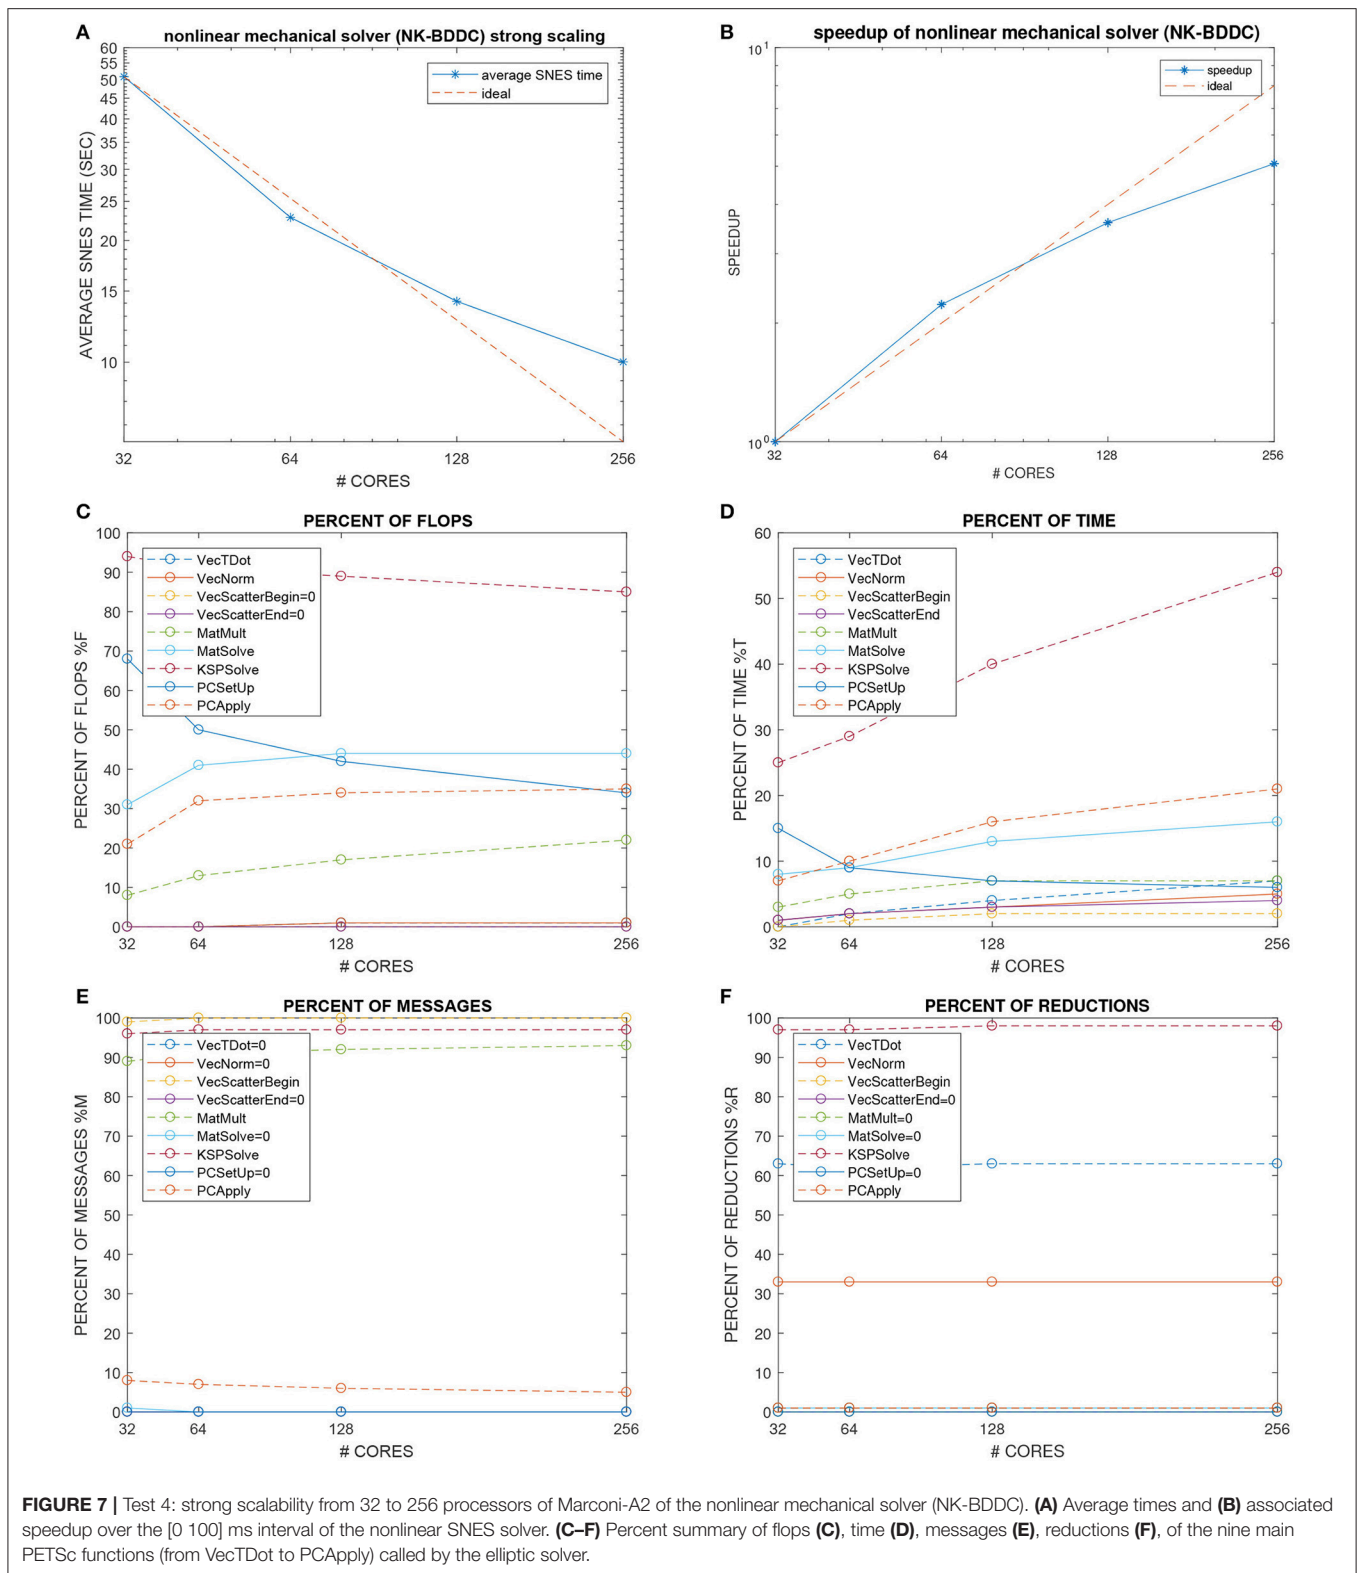
5.1. Bidomain Solver

The results have shown that the electrical Bidomain solver is scalable, in terms of both weak and strong scaling, and is robust with respect to the deformation induced by the mechanical contraction. Bidomain weak scaling tests have been

performed both on the Mira BG/Q and Marconi-A2 clusters. The two architectures and the number of cores used are different, although the load per core is the same. Thus, we can not compare fairly the performances obtained on the two architectures. However, the CPU times reported in **Figure 4A**, bottom and **Figure 5A** have the same order of magnitude, showing that the solution of the Bidomain linear systems on the two architectures exhibit comparable costs.

5.2. Mechanical Solver

The results have shown that also the mechanical NK-BDDC solver is scalable in terms of non-linear and linear iterations



counts, but the CPU timings, especially in the weak scaling test, do not present a scalable behavior. Our results seem to indicate that this increase of CPU timings can be attributed to the increase of computational costs required by the BDDC

coarse solver. A possible remedy would be to employ a multilevel BDDC solver, where the coarse problem is solved recursively by a BDDC method with additional local and coarse problems, or to employ an adaptive selection of BDDC primal constraints. The

nonscalability and ill-conditioning of the nonlinear mechanical system could also be associated with: (a) the penalty formulation employed to enforce the almost incompressibility of the cardiac tissue; (b) the presence of the stress induced by the active tension contraction model; (c) the particular mechanical boundary condition enforcing zero displacements on a fixed endocardial basal ring and fixed intracavitary endocardial pressure.

5.3. Comparison With Previous Studies

So far, only few studies have developed and investigated parallel numerical solvers for cardiac electro-mechanics. Lafortune et al. (2012) have proposed a fully explicit Monodomain-mechanical solver, obtaining good strong scalability results up to 500 cores. The advantage of our approach with respect to that presented in Lafortune et al. (2012) is that our solver, resulting from a semi-implicit time discretization of the electro-mechanical model, allows larger time step sizes and time adaptivity. Augustin et al. (2016) have developed a very effective electro-mechanical solver, tested on highly accurate patient-specific geometric models and based on Algebraic Multigrid (AMG) preconditioners for both the Bidomain and mechanical systems. The strong scalability results they have reported show a very good performance of AMG applied to the non-linear mechanical system, whereas the AMG preconditioner is less effective for the Bidomain linear system. The advantage of our solver compared to that introduced in Augustin et al. (2016) is that both Multilevel Additive Schwarz and BDDC preconditioners should be more robust

than AMG when high order finite elements or isogeometric analysis (see e.g., Charawi, 2017) discretizations are employed. On the other hand, while BDDC preconditioners can be easily constructed for unstructured meshes, Multilevel Additive Schwarz methods are more difficult to implement in case of such grids.

5.4. Future Work

In order to improve our mechanical solver, further studies could consider the following issues: (a) mixed formulations of the mechanical system based on inf-sup stable displacement-pressure discrete spaces; (b) alternative active tension contraction models; (c) alternative mechanical boundary conditions and pressure-volume relationships involving multielement Windkessel models.

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct and intellectual contribution to the work, and approved it for publication.

ACKNOWLEDGMENTS

This work was partially supported by grants of Istituto di Matematica Applicata e Tecnologie Informatiche IMATI - C.N.R., Pavia, Italy, and of Istituto Nazionale di Alta Matematica (INdAM), Italy.

REFERENCES

- Aguado-Sierra, J., Santiago, A., Rivero, M. I., López-Yunta, M., Soto-Iglesias, D., Dux-Santoy, L., et al. (2015). "Fully-coupled electromechanical simulations of the LV dog anatomy using HPC: model testing and verification," in *Statistical Atlases and Computational Models of the Heart - Imaging and Modelling Challenges. STACOM 2014. Lecture Notes in Computer Science*, Vol. 8896, eds O. Camara, T. Mansi, M. Pop, K. Rhode, M. Sermesant, and A. Young (Cham: Springer), 114–122. doi: 10.1007/978-3-319-14678-2_12
- Augustin, C. M., Neic, A., Liebmann, M., Prassl, A. J., Niederer, S. A., Haase, G., et al. (2016). Anatomically accurate high resolution modeling of human whole heart electromechanics: a strongly scalable algebraic multigrid solver method for nonlinear deformation. *J. Comput. Phys.* 305, 622–646. doi: 10.1016/j.jcp.2015.10.045
- Balay, S., Buschelman, K., Gropp, W. D., Kaushik, D., Knepley, M., McInnes, L. C., et al. (2012). *Petsc Users Manual*. Tech. Rep. ANL-95/11 - Revision 3.3, Argonne National Laboratory.
- Brands, D., Klawonn, A., Rheinbach, O., and Schröder, J. (2008). Modelling and convergence in arterial wall simulations using a parallel FETI solution strategy. *Comput. Methods Biomech. Biomed. Eng.* 11, 569–583. doi: 10.1080/10255840801949801
- Charawi, L. A. (2017). Isogeometric overlapping schwarz preconditioners for the bidomain reaction–diffusion system. *Comput. Methods Appl. Mech. Eng.* 319, 472–490. doi: 10.1016/j.cma.2017.03.012
- Chen, H., Li, X., and Wang, Y. (2017). A splitting preconditioner for a block two-by-two linear system with applications to the bidomain equations. *J. Comput. Appl. Math.* 321, 487–498. doi: 10.1016/j.cam.2017.03.017
- Colli Franzone, P., and Pavarino, L. F. (2004). A parallel solver for reaction - diffusion systems in computational electrocardiology. *Math. Mod. Meth. Appl. Sci.* 14, 883–911. doi: 10.1142/S0218202504003489
- Colli Franzone, P., Pavarino, L. F., and Scacchi, S. (2014). *Mathematical Cardiac Electrophysiology*. New York, NY: Springer.
- Colli Franzone, P., Pavarino, L. F., and Scacchi, S. (2015). Parallel multilevel solvers for the cardiac electro-mechanical coupling. *Appl. Numer. Math.* 95, 140–153. doi: 10.1016/j.apnum.2014.11.002
- Colli Franzone, P., Pavarino, L. F., and Scacchi, S. (2016a). Bioelectrical effects of mechanical feedbacks in a strongly coupled cardiac electro-mechanical model. *Math. Mod. Meth. Appl. Sci.* 26, 27–57. doi: 10.1142/S0218202516500020
- Colli Franzone, P., Pavarino, L. F., and Scacchi, S. (2016b). Joint influence of transmural heterogeneities and wall deformation on cardiac bioelectrical activity: a simulation study. *Math. Biosci.* 280, 71–86. doi: 10.1016/j.mbs.2016.08.003
- Colli Franzone, P., Pavarino, L. F., and Scacchi, S. (2017). Effects of mechanical feedback on the stability of cardiac scroll waves: a bidomain electro-mechanical simulation study. *Chaos* 27:093905. doi: 10.1063/1.4999465
- Constantino, J., Hu, Y., and Trayanova, N. A. (2012). A computational approach to understanding the cardiac electromechanical activation sequence in the normal and failing heart, with translation to the clinical practice of CRT. *Prog. Biophys. Mol. Biol.* 110, 372–379. doi: 10.1016/j.pbiomolbio.2012.07.009
- Dal, H., Göktepe, S., Kaliske, M., and Kuhl, E. (2013). A fully implicit finite element method for bidomain models of cardiac electromechanics. *Comput. Methods Appl. Mech. Eng.* 253, 323–336. doi: 10.1016/j.cma.2012.07.004
- Dohrmann, C. R. (2003). A preconditioner for substructuring based on constrained energy minimization. *SIAM J. Sci. Comp.* 25, 246–258. doi: 10.1137/S1064827502412887
- Eriksson, T. S. E., Prassl, A. J., Plank, G., and Holzapfel, G. A. (2013). Influence of myocardial fiber/sheet orientations on left ventricular mechanical contraction. *Math. Mech. Solids* 18, 592–606. doi: 10.1177/1081286513485779
- Favino, M., Pozzi, S., Pezzuto, S., Prinzen, F. W., Auricchio, A., and Krause, R. (2016). Impact of mechanical deformation on pseudo-ECG: a simulation study. *Europace* 18, iv77–iv84. doi: 10.1093/europace/euw353
- Gerardo-Giorda, L., Mirabella, L., Nobile, F., Perego, M., and Veneziani, A. (2009). A model-based block-triangular preconditioner for the bidomain system in electrocardiology. *J. Comput. Phys.* 228, 3625–3639. doi: 10.1016/j.jcp.2009.01.034
- Göktepe, S., and Kuhl, E. (2010). Electromechanics of the heart: a unified approach to the strongly coupled excitation–contraction problem. *Comput. Mech.* 45, 227–243. doi: 10.1007/s00466-009-0434-z
- Gurev, V., Lee, T., Constantino, J., Arevalo, H., and Trayanova, N. A. (2011). Models of cardiac electromechanics based on individual hearts imaging

- data. *Biomech. Model. Mechanobiol.* 10, 295–306. doi: 10.1007/s10237-010-0235-5
- Gurev, V., Pathmanathan, P., Fettebert, J.-L., Wen, H.-F., Magerlein, J., Gray, R. A., et al. (2015). A high-resolution computational model of the deforming human heart. *Biomech. Model. Mechanobiol.* 14, 829–849. doi: 10.1007/s10237-014-0639-8
- Holzzapfel, G. A., and Ogden, R. W. (2009). Constitutive modelling of passive myocardium. A structurally-based framework for material characterization. *Philos. Trans. R. Soc. Lond. A* 367, 3445–3475. doi: 10.1098/rsta.2009.0091
- Keldermann, R. H., Nash, M. P., Gelderblom, H., Wang, V. Y., and Panfilov, A. V. (2010). Electromechanical wavebreak in a model of the human left ventricle. *AJP Heart Circul. Physiol.* 299, H134–H143. doi: 10.1152/ajpheart.00862.2009
- Klawonn, A., and Rheinbach, O. (2010). Highly scalable parallel domain decomposition methods with an application to biomechanics. *ZAMM* 90, 5–32. doi: 10.1002/zamm.200900329
- Lafortune, P., Aris, R., Vázquez, M., and Houzeaux, G. (2012). Coupled electromechanical model of the heart: parallel finite element formulation. *Int. J. Numer. Methods Biomed. Eng.* 28, 72–86. doi: 10.1002/cnm.1494
- Lamata, P., Cookson, A., and Smith, N. (2015). Clinical diagnostic biomarkers from the personalization of computational models of cardiac physiology. *Ann. Biomed. Eng.* 44, 46–57. doi: 10.1007/s10439-015-1439-8
- Land, S., Gurev, V., Arens, S., Augustin, C. M., Baron, L., Blake, R., et al. (2015). Verification of cardiac mechanics software: benchmark problems and solutions for testing active and passive material behaviour. *Proc. R. Soc. A Math. Phys. Eng. Sci.* 471:20150641. doi: 10.1098/rspa.2015.0641
- Land, S., Niederer, S. A., Aronsen, J. M., Espe, E. K. S., Zhang, L. L., Louch, W. E., et al. (2012a). An analysis of deformation-dependent electromechanical coupling in the mouse heart. *J. Physiol.* 590, 4553–4569. doi: 10.1113/jphysiol.2012.231928
- Land, S., Niederer, S. A., and Smith, N. P. (2012b). Efficient computational methods for strongly coupled cardiac electromechanics. *IEEE Trans. Biomed. Eng.* 59, 1219–1228. doi: 10.1109/TBME.2011.2112359
- Mandel, J., and Dohrmann, C. R. (2003). Convergence of a balancing domain decomposition by constraints and energy minimization. *Numer. Lin. Alg. Appl.* 10, 639–659. doi: 10.1002/nla.341
- Mandel, J., Dohrmann, C. R., and Tezaur, R. (2005). An algebraic theory for primal and dual substructuring methods by constraints. *Appl. Numer. Math.* 54, 167–193. doi: 10.1016/j.apnum.2004.09.022
- Munteanu, M., Pavarino, L. F., and Scacchi, S. (2009). A scalable Newton–Krylov–Schwarz method for the bidomain reaction-diffusion system. *SIAM J. Sci. Comp.* 31, 3861–3883. doi: 10.1137/08074355X
- Niederer, S. A., and Smith, N. P. (2007). A mathematical model of the slow force response to stretch in rat ventricular myocytes. *Biophys. J.* 92, 4030–4044. doi: 10.1529/biophysj.106.095463
- Noble, F., Quarteroni, A., and Ruiz-Baier, R. (2012). An active strain electromechanical model for cardiac tissue. *Int. J. Numer. Methods Biomed. Eng.* 28, 52–71. doi: 10.1002/cnm.1468
- Nordsletten, D., Niederer, S., Nash, M., Hunter, P., and Smith, N. (2011). Coupling multi-physics models to cardiac mechanics. *Prog. Biophys. Mol. Biol.* 104, 77–88. doi: 10.1016/j.pbiomolbio.2009.11.001
- Pathmanathan, P., and Whiteley, J. P. (2009). A numerical method for cardiac mechanoelectric simulations. *Ann. Biomed. Eng.* 37, 860–873. doi: 10.1007/s10439-009-9663-8
- Pavarino, L. F., and Scacchi, S. (2008). Multilevel additive Schwarz preconditioners for the bidomain reaction-diffusion system. *SIAM J. Sci. Comp.* 31, 420–443. doi: 10.1137/070706148
- Pavarino, L. F., and Scacchi, S. (2011). Parallel multilevel Schwarz and block preconditioners for the bidomain parabolic-parabolic and parabolic-elliptic formulations. *SIAM J. Sci. Comp.* 33, 1897–1919. doi: 10.1137/100808721
- Pavarino, L. F., Scacchi, S., and Zampini, S. (2015). Newton-Krylov-BDDC solvers for nonlinear cardiac mechanics. *Comput. Meth. Appl. Mech. Eng.* 295, 562–580. doi: 10.1016/j.cma.2015.07.009
- Pennacchio, M., and Simoncini, V. (2009). Algebraic multigrid preconditioners for the bidomain reaction–diffusion system. *Appl. Numer. Math.* 59, 3033–3050. doi: 10.1016/j.apnum.2009.08.001
- Pennacchio, M., and Simoncini, V. (2011). Fast structured AMG preconditioning for the bidomain model in electrocardiology. *SIAM J. Sci. Comp.* 33, 721–745. doi: 10.1137/100796364
- Plank, G., Liebmann, M., Weber dos Santos, R., Vigmond, E. J., and Haase, G. (2007). Algebraic multigrid preconditioner for the cardiac bidomain model. *IEEE Trans. Biomed. Eng.* 54, 585–596. doi: 10.1109/TBME.2006.889181
- Quarteroni, A., Lassila, T., Rossi, S., and Ruiz-Baier, R. (2017a). Integrated heart—coupling multiscale and multiphysics models for the simulation of the cardiac function. *Comput. Methods Appl. Mech. Eng.* 314, 345–407. doi: 10.1016/j.cma.2016.05.031
- Quarteroni, A., Manzoni, A., and Vergara, C. (2017b). The cardiovascular system: mathematical modelling, numerical algorithms and clinical applications. *Acta Numer.* 26, 365–590. doi: 10.1017/S0962492917000046
- Rossi, S., Ruiz-Baier, R., Pavarino, L. F., and Quarteroni, A. (2012). Orthotropic active strain models for the numerical simulation of cardiac biomechanics. *Int. J. Numer. Methods Biomed. Eng.* 28, 761–788. doi: 10.1002/cnm.2473
- Scacchi, S. (2008). A hybrid multilevel Schwarz method for the bidomain model. *Comput. Meth. Appl. Mech. Eng.* 197, 4051–4061. doi: 10.1016/j.cma.2008.04.008
- Scacchi, S. (2011). A multilevel hybrid Newton–Krylov–Schwarz method for the bidomain model of electrocardiology. *Comput. Meth. Appl. Mech. Eng.* 200, 717–725. doi: 10.1016/j.cma.2010.09.016
- Sundnes, J., Lines, G., Mardal, K., and Tveito, A. (2002). Multigrid block preconditioning for a coupled system of partial differential equations modeling the electrical activity in the heart. *Comput. Meth. Biomech. Biomed. Eng.* 5, 397–409. doi: 10.1080/1025584021000025023
- Sundnes, J., Wall, S., Osnes, H., Thorvaldsen, T., and McCulloch, A. (2014). Improved discretisation and linearisation of active tension in strongly coupled cardiac electro-mechanics simulations. *Comput. Meth. Biomech. Biomed. Eng.* 17, 604–615. doi: 10.1080/10255842.2012.704368
- ten Tusscher, K. H. W. J., Noble, D., Noble, P. J., and Panfilov, A. V. (2004). A model for human ventricular tissue. *Am. J. Phys. Heart Circ. Physiol.* 286, H1573–H1589. doi: 10.1152/ajpheart.00794.2003
- ten Tusscher, K. H. W. J., and Panfilov, A. V. (2006). Alternans and spiral breakup in a human ventricular tissue model. *Am. J. Phys. Heart Circ. Physiol.* 291, H1088–H1100. doi: 10.1152/ajpheart.00109.2006
- Toselli, A., and Widlund, O. B. (2005). *Domain Decomposition Methods: Algorithms and Theory*. Berlin: Springer.
- Trayanova, N. A., and Chang, K. C. (2016). How computer simulations of the human heart can improve anti-arrhythmia therapy. *J. Physiol.* 594, 2483–2502. doi: 10.1113/JP270532
- Trayanova, N. A., Constantino, J., and Gurev, V. (2011). Electromechanical models of the ventricles. *AJP Heart Circul. Physiol.* 301, H279–H286. doi: 10.1152/ajpheart.00324.2011
- Vázquez, M., Aris, R., Houzeaux, G., Aubry, R., Villar, P., Garcia-Barnés, J., et al. (2011). A massively parallel computational electrophysiology model of the heart. *Int. J. Numer. Meth. Biomed. Eng.* 27, 1911–1929. doi: 10.1002/cnm.1443
- Washio, T., Okada, J., Takahashi, A., Yoneda, K., Kadooka, Y., Sugiura, S., et al. (2013). Multiscale heart simulation with cooperative stochastic cross-bridge dynamics and cellular structures. *Multiscale Model. Simul.* 11, 965–999. doi: 10.1137/120892866
- Weber dos Santos, R., Plank, G., Bauer, S., and Vigmond, E. (2004). Parallel multigrid preconditioner for the cardiac bidomain model. *IEEE Trans. Biomed. Eng.* 51, 1960–1968. doi: 10.1109/TBME.2004.834275
- Zampini, S. (2013). Balancing neumann–neumann methods for the cardiac bidomain model. *Numer. Math.* 123, 363–393. doi: 10.1007/s00211-012-0488-2
- Zampini, S. (2014). Dual-primal methods for the cardiac bidomain model. *Math. Mod. Meth. Appl. Sci.* 24, 667–696. doi: 10.1142/S0218202513500632
- Zampini, S. (2016). PCBDDC: a class of robust dual-primal methods in PETSc. *SIAM J. Sci. Comput.* 38, S282–S306. doi: 10.1137/15M1025785

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2018 Colli Franzone, Pavarino and Scacchi. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.