X International Conference on Computational Methods for Coupled Problems in Science and Engineering
COUPLED PROBLEMS 2023
M. Papadrakakis, B. Schrefler and E. Oñate (Eds)

EFFICIENT NUMERICAL METHODS FOR SIMULATING CARDIAC ELECTROPHYSIOLOGY WITH CELLULAR RESOLUTION

Fatemeh Chegini[a], Algiane Froehly[b], Ngoc Mai Monica Huynh[c], Luca Pavarino[c], Mark Potse[b], Simone Scacchi[d], and Martin Weiser[a*]

[a] Zuse Institute Berlin, Takustr. 7, D-14195 Berlin, Germany, {chegini,weiser}@zib.de
[b] Inria/IMB/IHU Liryc, U. Bordeaux, Talence, France {mark.potse,algiane.froehly}@inria.fr
[c] Università degli Studi di Pavia, Italy, {ngocmaimonica.huynh,luca.pavarino}@unipv.it
[d] Università degli Studi di Milano, Italy, simone.scacchi@unimi.it

Key words: EMI model, high-order time integration, spectral deferred correction, algebraic adaptivity, multirate integration scheme, ladder methods, BDDC preconditioner

Abstract. The cardiac extracellular-membrane-intracellular (EMI) model enables the precise geometrical representation and resolution of aggregates of individual myocytes. As a result, it not only yields more accurate simulations of cardiac excitation compared to homogenized models but also presents the challenge of solving much larger problems. In this paper, we introduce recent advancements in three key areas: (i) the creation of artificial, yet realistic grids, (ii) efficient higher-order time stepping achieved by combining low-overhead spatial adaptivity on the algebraic level with progressive spectral deferred correction methods, and (iii) substructuring domain decomposition preconditioners tailored to address the complexities of heterogeneous problem structures. The efficiency gains of these proposed methods are demonstrated through numerical results on cardiac meshes of different sizes.

1 INTRODUCTION

Cardiac arrhythmias are significant causes of mortality, and numerical modeling plays a crucial role in understanding their mechanisms and designing effective treatments. While homogenized models like the bidomain and monodomain models represent the action potential at the tissue level, they lack details about arrhythmogenic cellular-level factors such as fibrosis or ion channel distribution. To address this, we consider a recent cell-by-cell model that explicitly represents myocytes. Due to the size of such models, they pose new simulation challenges related to mesh generation, time stepping, and preconditioning.

In this paper, we present a geometry and mesh generation approach for defining large-scale computational meshes and introduce efficient adaptive time stepping and domain decomposition preconditioning techniques. We employ tetrahedral meshes based on the implicit surface meshing of a random network of cells, which mimics real cardiac tissue [11]. This method has been extended to represent 400 cells in 15 million tetrahedra.

Although the cell-by-cell model exhibits localized solution features, the overhead induced by mesh adaptivity during simulation has proven to outweigh efficiency gains. To overcome

this, we combine the spectral deferred correction (SDC) method as a higher-order time stepping approach with computationally inexpensive algebraic adaptivity [2]. This combination, which can be interpreted as multi-rate integration, significantly speeds up simulations in both homogenized and cell-by-cell models [1]. The approach utilizes nested subset selection on the algebraic level and explicit embedded error estimates, incurring only negligible overhead.

Since the dominant diffusion and elliptic constraints in the cell-by-cell models necessitate implicit time-stepping methods, we also propose a novel domain decomposition preconditioner called balancing domain decomposition by constraints (BDDC) [8]. This preconditioner is specifically tailored to the unique structure of cell-by-cell models, and we investigate its convergence properties numerically.

To showcase the overall performance and convergence rate of our proposed approach, we provide demonstrations on 3D cardiac simulations using a combination of SDC, algebraic adaptivity, and BDDC preconditioning.

2 THE EMI MODEL

The extracellular-membrane-intracellular (EMI) model of electrophysiology [9] describes the myocardium as a collection of pairwise disjoint myocytes $(\Omega_i)_{i=1,\dots,N}$ which, together with the extracellular space Ω_0, cover the whole domain $\Omega \subset \mathbb{R}^d$, $d \in \{2,3\}$, occupied by the myocardium, i.e. $\overline{\Omega} = \bigcup_{i=0}^{N} \overline{\Omega}_i$. Ions can diffuse freely within each myocyte and in the extracellular domain, subject to conductivities σ_i, which leads to electric intra- and extracellular potentials $u_i \in H^1(\Omega_i)$. Gap junctions between myocytes allow ion currents through passive channels, while active channels regulate ion exchange between myocytes and the extracellular space. The transmembrane current $n^T \sigma_i \nabla u_i$ includes the ion current I_{ij}^{ion} and the capacitive current $C_m \dot{v}_{ij}$. The ion current depends on the transmembrane voltage $v_{ij} = u_i - u_j$ defined on the membrane $F_{ij} = \partial\Omega_i \cap \partial\Omega_j$, and w_{ij} comprising ion concentrations and channel states governed by a nonlinear dynamic given by R. This setting results in the partial differential algebraic system

$$\begin{cases} -\operatorname{div}(\sigma_i \nabla u_i) = 0 & \text{in } \Omega_i \\ -n_i^T \sigma_i \nabla u_i = C_m \dot{v}_{ij} + I_{ij}^{\text{ion}}(v_{ij}, w_{ij}) & \text{on } F_{ij},\ i \neq j \\ n^T \sigma_i \nabla u_i + \epsilon u_i = 0 & \text{on } \partial\Omega_i \cap \partial\Omega \\ \dot{w}_{ij} = R(v_{ij}, w_{ij}) & \text{on } F_{ij}, ij = 0 \\ w_{ij} = 0 & ij > 0. \end{cases} \tag{1}$$

Here, n_i denotes the unit outer normal of Ω_i for $i = 0, \dots, N$. We require $I_{ij}^{\text{ion}}(v_{ij}, w_{ij}) = -I_{ji}^{\text{ion}}(v_{ji}, w_{ji})$ and $w_{ij} = w_{ji}$. The small value $\epsilon > 0$ in the Robin boundary condition on $\partial\Omega$ makes the solution unique, which for a pure Neumann problem would be defined only up to a constant. An alternative normalization would be to require $\int_\Omega \phi u \, dx = 0$ for some nonvanishing function ϕ. The parameter values used for the EMI model are given in Tab. 1.

3 MESH GENERATION

Computational meshes were generated as described previously [11], with modifications that allowed us to produce larger samples. First we generated a random network of cell centers and links, which were not allowed to deviate more than 15 degrees from a defined "fiber orientation". Each half-link was then assigned the part of the tissue volume that was nearest to it.

Parameter	Value
$\sigma_{>0}$	$0.4\,\mathrm{S\,m^{-1}}$
σ_0	$2.0\,\mathrm{S\,m^{-1}}$
C_m	$1 \times 10^{-4}\,\mathrm{F\,m^{-2}}$

Table 1: Parameters used in the EMI model, based on [9].

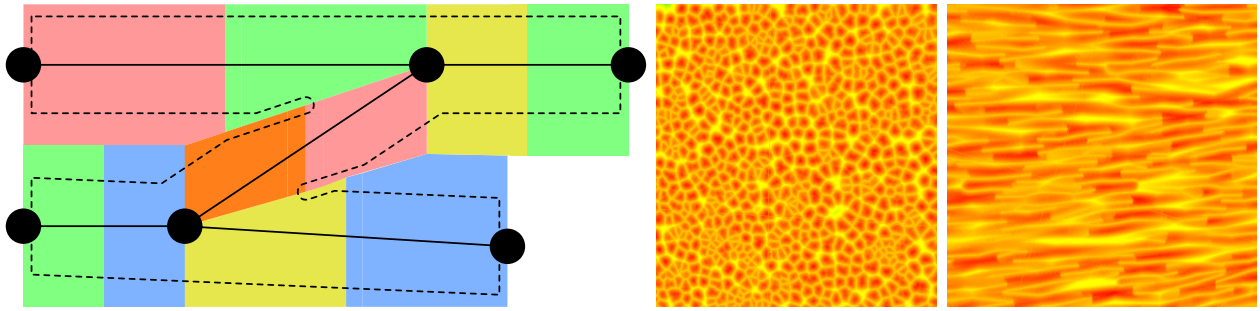

Figure 1: Overview of the mesh generation methods. Left: a network with 6 centers (black dots), 5 links, and halflink domains. The dashed lines represent cell membranes. Middle and right: transverse and longitudinal section of the level-set function, where red colors code the inside of the cells. The narrow spaces between the cells (about one tenth of the cell diameter in reality) made it necessary to discretize the domain boundaries before computation of the level-set function.

The boundaries between these domains were discretized inside a tetrahedral mesh. Using these boundaries and the links themselves we determined a level-set function that defined the cell membranes. The level set was discretized by the Mmg remeshing software [4]. The halflink domains were preserved through level-set discretization and then merged for each center, resulting in cell domains that determined the locations of the intercalated disks, the structures that connect the cells electrically and mechanically. The resulting poor-quality mesh was then improved with a second application of Mmg. In contrast to our previous work [11], we merged the extracellular parts of the cell domains before remeshing. This made the remesher's task much easier, because the junctions of multiple domain boundaries, often joining at small angles, were hard to preserve. This modification, together with recent improvements in the Mmg software, allowed us to generate samples that were orders of magnitude larger than the few dozen myocytes reported before.

A domain decomposition for parallel computations was made by combining the intracellular domains of several cells and expanding them until the extracellular volume was completely partitioned as well. Thus, cells were never split across partition boundaries. Necessarily, the partition boundaries could coincide with cell membranes and intercalated disks.

Fig. 1 shows an overview of the mesh generation methods, Fig. 4 the four meshes used in this study, and Fig. 2 the intra- and extracellular domains of a single cell.

4 ALGEBRAIC ADAPTIVITY IN HIGH ORDER TIME STEPPING

The depolarization and repolarization fronts travelling through the domain are highly localized, calling for high spatial resolution at the front position but allowing for much coarser meshes in regions far from the front, reducing the problem size considerably. This can be

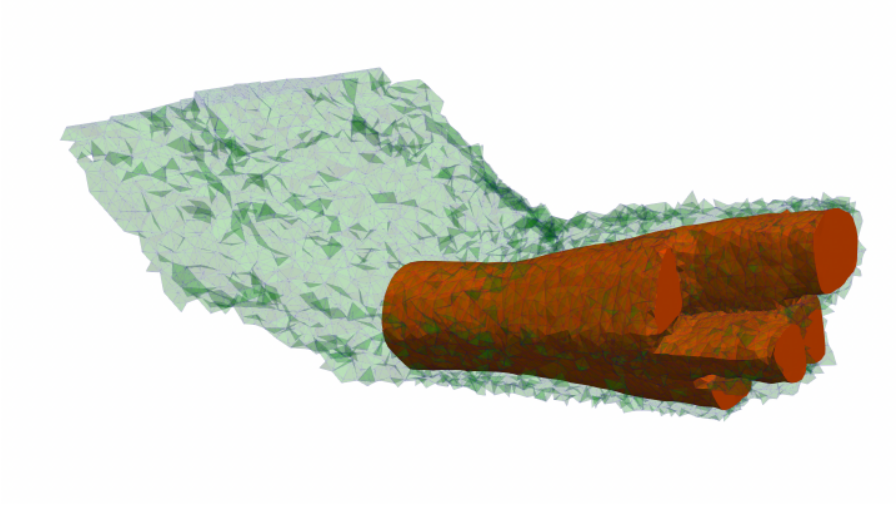

Figure 2: Each myocyte (orange) is surrounded by an extracellular subdomain (green). As illustrated here for a location near a model boundary, the extracellular domains can be considerably larger than the cells when these are sparse.

achieved by adaptive mesh refinement and coarsening [3]. Unfortunately, the overhead of error estimation, mesh modification, and re-assembly often compensates the reduction of computational work due to the smaller problems to be solved. Here we consider a recent approach to spatial adaptivity with spectral deferred correction time stepping [1, 2] that works purely on the algebraic level without incurring significant overhead and thus can achieve a substantial speedup. Combining this algebraic adaptivity with ladder methods [10] improves the efficiency gain further.

4.1 Algebraic adaptivity with spectral deferred corrections

Spatial semi-discretization of (1) with finite elements results in a large scale differential-algebraic system

$$B\dot{z} = Az + F(z) \tag{2}$$

of index one for $z = [u, w] \in \mathbb{R}^{N_h+n_h}$. Here, A contains the infinitely stiff diffusion term, and F represents the nonlinear ion channel dynamics. Note that w is restricted to the membranes, but may be vectorial, whereas u is defined on the whole domain, but is scalar. Consequently, $N_h \neq n_h$. Since time derivatives only appear on the membranes, the matrix B is singular, such that implicit L-stable integration methods such as an implicit-explicit (IMEX) Euler scheme

$$(B - \tau A)z_{k+1} = z_k + \tau F(z_k) \tag{3}$$

for a step size τ have to be used [7]. Spectral deferred correction (SDC) methods turn a low order basic scheme such as (3) into a higher order scheme by solving the defect equation $B\dot{\delta z} = A\delta z + F(z + \delta z) - F(z) + [Az + F(z) - \dot{z}]$ for a given approximate trajectory z on a certain time horizon, using the basic scheme [6, 12]. The term in brackets is independent of δz and can be treated explicitly by high-order quadrature on a collocation time grid. Updating the initial guess $z \leftarrow z + \delta z$ and repeating yields a stationary iteration that converges to a high order collocation solution. Choosing a Radau quadrature with n collocation points and taking

m iterations yields an L-stable IMEX Runge-Kutta scheme of order $\min\{2n-1, m\}$ with nm stages.

In electrophysiology simulations, the SDC corrections have a dominantly local support around the fronts. Based on an error estimator based on the linear convergence of the stationary iteration, thresholding corrections for determining the relevant subregion still requiring higher accuracy, and restricting subsequent iterations to that subregion has been shown to reach the desired uniform tolerance [1]. The selection of the subregions and the formulation of the restricted problems can be done purely on the algebraic level of finite element coefficients, as shown in Fig. 5, which is computationally cheap and thus turns the reduction in problem size directly into faster solution. Wall clock time speedup factors over a conventional SDC iteration are given in Tab. 3 for grids of different size. Since the depolarization front width stays fixed, the speedup increases slightly with the domain size.

4.2 Ladder methods

The speedup that can be achieved with the algebraic adaptivity outlined above is limited by the number m of SDC iterations. This holds for the extreme case when after the first iteration, which needs to be performed on the whole domain, the following iterations involve just a negligible number of degrees of freedom. A further improvement is possible if the first iteration does not involve all n collocation points, but just one. So-called ladder methods, refining the collocation grid every second iteration, reduce the computational cost, but also yield a larger discretization error, and have been found not to improve the efficiency of ODE integration because they require more iterations for the same accuracy [10].

This changes if the computational cost for the iterations is not uniform. With algebraic adaptivity, the computational cost of the first iteration is much higher than for the later iterations. Consequently, the upper bound for the speedup is increased from m to nm. As a very coarse model of impact of ladder methods on accuracy and effort, we will assume the following:

1. Standard SDC converges linearly with a contraction factor $\rho < 1$. This is in general satisfied, with $0.15 \leq \rho \leq 0.6$ the usual range for reaction-diffusion equations [12].

2. Algebraic adaptivity does not affect the reached accuracy and hence convergence rate. With a proper choice of drop tolerance, this is in general the case [1].

3. The number of degrees of freedom retained in iteration i scales as Nr^i for some $r < 1$. The actual number of retained degrees of freedom is, of course, highly problem dependent, but a factor $r = 0.5$ appears to be realistic.

4. The larger basic time step in ladder methods increases the contraction factor to ρ^q, with $q < 1$. Asymptotically, this is a rather pessimistic assumption, since after $2n$ iterations, the full collocation grid is reached and the ladder method's contraction rate reaches ρ. Up to then, of course, larger contraction factors are observed.

With these assumptions, and performing $m = 2n-1$ iterations to reach the full time discretization order of Radau collocation, we face a computational effort proportional to

$$W := \sum_{i=0}^{m-1} nr^i = n\frac{1-r^{2n-1}}{1-r}$$

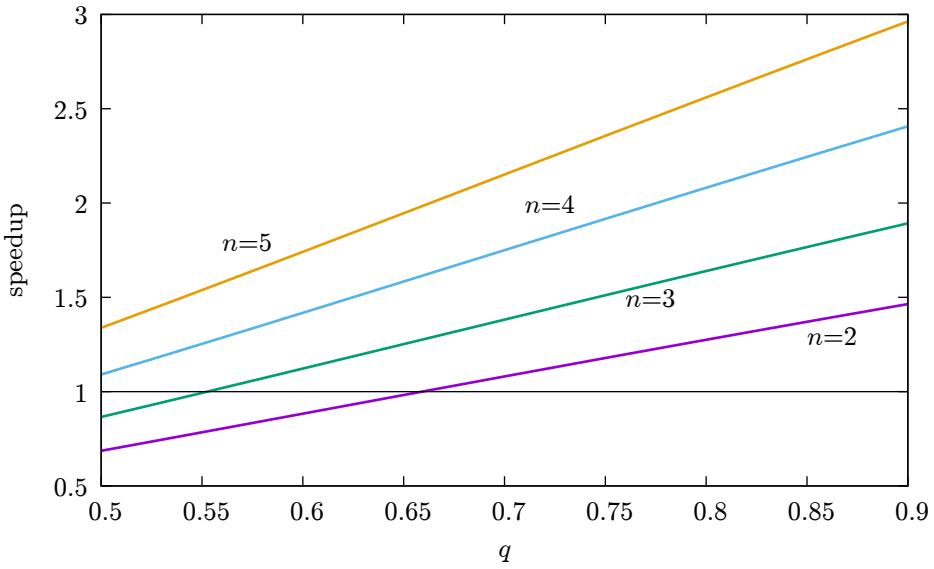

Figure 3: Expected speedup due to ladder methods applied in addition to algebraic adaptivity for different numbers n of collocation points and larger contraction factors ρ^q.

for standard SDC with algebraic adaptivity. For ladder methods with algebraic adaptivity, the number of iterations is about $m = (2n-1)/q$ to reach the same truncation error. With a slower convergence rate ρ^q, the number of retained degrees of freedom scales as Nr^{iq} and thus decreases slower, but there are also just $\lfloor i/2+1 \rfloor$ collocation points in each iteration. Thus, with $a = r^q$ and $m = \lceil (2n-1)/q \rceil$, the computational work is roughly proportional to

$$\sum_{i=0}^{\lceil (2n-1)/q \rceil} \left(\left\lfloor \frac{i}{2} \right\rfloor + 1 \right) r^{qi} \approx \frac{2 - a + a^{m+1}((m+2)a - (m+3))}{2(1-a)^2} =: W_{\text{lad}}. \tag{4}$$

The expected speedup W/W_{lad} due to using ladder methods with algebraic adaptivity is shown in Fig. 3 for $r = 0.5$ and different values of n and q. As can be expected, larger speedups are predicted for more collocation points (larger n) and less impact of the ladder approach on the convergence rate (larger q). Also apparent is that for small n and q, an actual slowdown must be expected.

The empirical speedup observed in the examples considered here is reported in the last column of Tab. 3.

5 DOMAIN DECOMPOSITION PRECONDITIONERS

The basic IMEX Euler scheme (3) requires the solution of globally coupled equation systems of the form $B - \tau A$. The matrix $B - \tau A$ is positive definite if the nullspace of the pure Neumann problem is eliminated by applying Robin boundary conditions as defined in (1). For large grids with many myocytes as shown in Fig. 4, iterative solvers are essential, and the method of choice is preconditioned conjugate gradients (PCG).

Among the numerous available options we consider a BDDC preconditioner [5] specifically tailored to the geometrical structure of the EMI model [8]. BDDC splits the mesh in several non-overlapping subdomains and decomposes the system matrix $M = B - \tau A$ into a corresponding block-diagonal system by replicating degrees of freedom associated to the interface

mesh	vertices	myocytes	extr. regions	degrees of freedom
(a)	332,179	44	44	394,106
(b)	522,489	86	86	628,586
(c)	1,314,675	208	208	1,664,019
(d)	2,521,502	415	415	3,286,868

Table 2: Properties of the meshes shown in Fig. 4.

between subdomains in each block. Approximate global continuity is preserved by adding few constraints such as continuity at subdomain corners or vanishing average jumps over subdomain interfaces. Since there are very few such constraints, the Schur complement with respect to the block-diagonal part is relatively small, and can therefore be solved efficiently using direct solvers. This yields an approximate solution with artificial jumps across subdomain interfaces, which are subsequently removed by averaging the solution at the interfaces and computing local solutions on subdomains given these Dirichlet boundary data. The resulting approximate solution is usually quite good, and thus BDDC used as a preconditioner exhibits very low condition numbers of the preconditioned system.

The EMI geometry suggests a natural decomposition of the computational domain Ω along the cell membranes into subdomains Ω_i representing the myocytes and extracellular domain. Since the potentials u_i are discontinuous across membranes, there are two different degrees of freedom associated to each interface vertex, and both of them are replicated in both subdomains. This BDDC design has been proved to exhibit the expected almost mesh-independent condition bound and demonstrated to yield a highly effective preconditioner in 2D problems [8].

Here, we consider larger 3D problems. This leads to a large extracellular subdomain, which is challenging for direct solvers and causes imbalance with the many smaller myocyte subdomains. We therefore subdivide the extracellular space into several subdomains. Introducing an interface condition of the same mathematical structure as gap junctions between these subdomains allows extending the BDDC design seamlessly to subdivided extracellular space. On these interfaces, we impose an ion current $I_{\text{ion}}(v) = \sigma_{\text{virt}} v$ with a large conductivity σ_{virt}, realizing an interior penalty discontinuous Galerking discretization, such that the solution is almost continuous.

6 NUMERICAL RESULTS

We have applied the algorithmic approaches described in Sec. 4 and 5 to four grids of different size shown in Fig. 4. Their properties are given in Tab. 2. Since the potential values jump across membranes, the number of degrees of freedom in a piecewise linear finite element discretization exceeds the number of mesh vertices by a small fraction.

We simulated the EMI model (1) on these grids with SDC on a Radau IIA time grid with three collocation points. The number of degrees of freedom retained in every SDC iteration is shown in Fig. 5 over the course of the whole simulation. The clearly visible significant reduction of system size translates into a reduction of wall clock simulation time, given as speedup over plain SDC in Tab. 3. Since the domains are larger than the one considered before [1], but the excitation front width remains the same, the observed speedup is slightly larger (4–6 here, 3–4 before). Since the computational cost of the SDC iterations with algebraic adaptivity

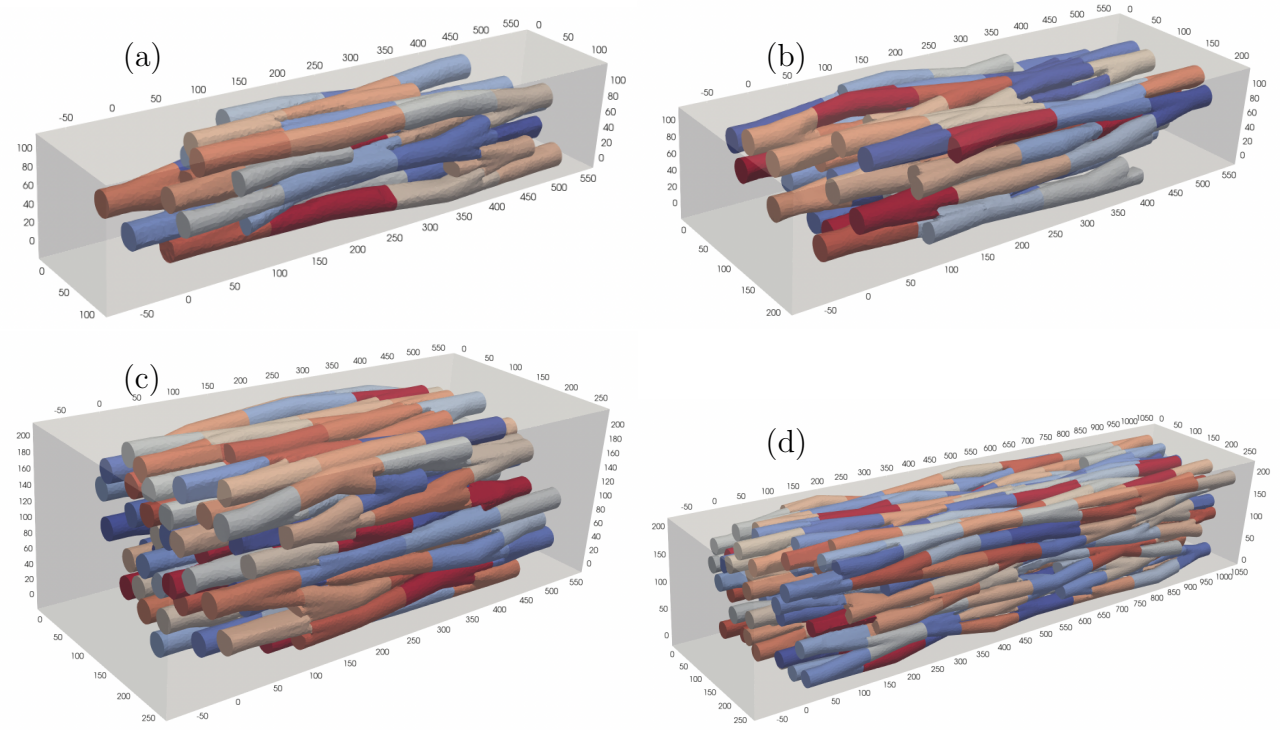

Figure 4: Different 3D meshes, detail in Table 2. Dimensions are indicated in micrometers; the largest model (d) is 1 mm long. Individual myocytes are differentiated by color.

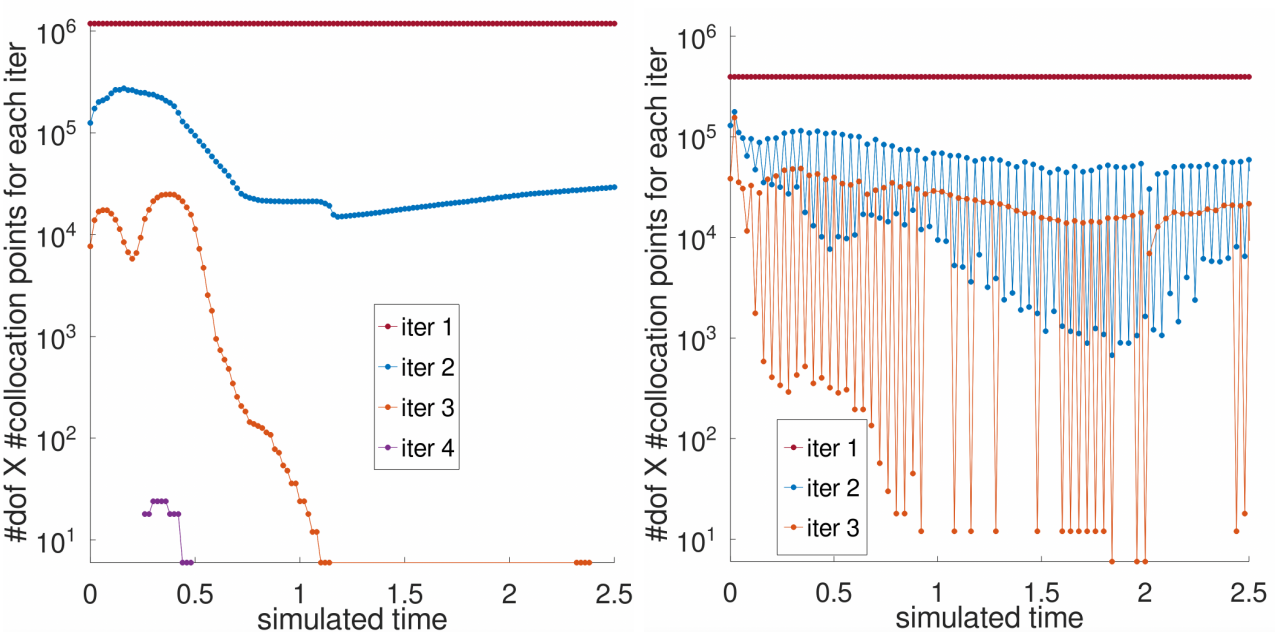

Figure 5: Number of degrees of freedom of the EMI model on grid (a) from Tab. 2 used in the SDC iterations versus simulated time. *Left:* pure algebraic adaptivity. *Right:* algebraic adaptivity combined with ladder methods.

is no longer uniform, ladder methods yield a further moderate but consistent increase of the speedup by 20–50%, which is well within the range predicted by the quite coarse efficiency model worked out in Sec. 4.2. The number of retained dofs, however, exhibits a strong oscillation, which suggests that the combination of algebraic adaptivity and ladder methods deserves more investigation.

mesh	AA	AA + Ladder	ratio
(a)	4.526	5.342	1.180
(b)	4.753	5.973	1.257
(c)	5.432	8.252	1.519
(c)	5.863	8.847	1.509

Table 3: Computational speedup due to algebraic adaptivity (AA) applied to standard SDC and to ladder SDC methods, over non-adaptive SDC. With 3 collocation points used, the additional speedup due to ladder methods is within the range expected from the coarse efficiency model (4).

We tested both unpreconditioned CG and BDDC-preconditioned CG for solving the full systems. As expected, BDDC reduces the iteration count dramatically, and leads to resolution-independent convergence due to the bounded condition number of the preconditioned system, see Tab. 4. Despite the increased computational effort per iteration introduced by the preconditioner, the computational time is also reduced by a large factor.

dofs	CG-noPrec			BDDC		
	itr	CPU time(s)	k_2	itr	CPU time(s)	k_2
33,698	1111	6.595	24,390	11	0.105	52
75,122	2388	58.711	83,610	12	0.307	134
132,930	2500	69.238	32,560	12	0.630	157
207,122	5209	608.205	138,080	10	0.771	229
297,698	4606	400.290	305,810	10	1.245	196

Table 4: Iteration count and condition number(k_2) estimate for unpreconditioned and BDDC-preconditioned CG on problems of different size.

7 Conclusion

We have made significant progress towards achieving efficient large-scale EMI model simulations, focusing on three key aspects: (i) robust creation of artificial but structurally realistic myocardium geometries and grids, (ii) efficient spatial adaptivity for higher-order time stepping, achieved through a combination of algebraic adaptivity and ladder SDC methods, and (iii) development of BDDC preconditioners tailored for complex 3D piecewise continuous elliptic problems.

The numerical results demonstrate that the proposed methods are effective, efficient, and allow scaling EMI computations to larger domain sizes than previously attainable.

Acknowledgement. This work was supported by the European High-Performance Computing Joint Undertaking EuroHPC under grant agreement No 955495 (MICROCARD) co-funded by the Horizon 2020 programme of the European Union (EU), the French National Research Agency ANR, the German Federal Ministry of Education and Research, and the Italian ministry of economic development.

REFERENCES

[1] F. Chegini, T. Steinke, and M. Weiser. Efficient adaptivity for simulating cardiac electrophysiology with spectral deferred correction methods. In *Submitted*, 2022.

[2] F. Chegini and M. Weiser. Adaptive multirate integration of cardiac electrophysiology with spectral deferred correction methods. In *7th Int. Conf. Comp. Math. Biomed. Eng.*, pages 528–531, 2022.

[3] P. Colli Franzone, P. Deuflhard, B. Erdmann, J. Lang, and L.F. Pavarino. Adaptivity in space and time for reaction-diffusion systems in electrocardiology. *SIAM J. Sci. Comput.*, 28(3):942–962, 2006.

[4] C. Dapogny, C. Dobrzynski, and P. Frey. Three-dimensional adaptive domain remeshing, implicit domain meshing, and applications to free and moving boundary problems. *J. Comp. Phys.*, 262:358–378, 2014.

[5] C.R. Dohrmann. A preconditioner for substructuring based on constrained energy minimization. *SIAM J. Sci. Comput.*, 25:246–258, 2003.

[6] A. Dutt, L. Greengard, and V. Rokhlin. Spectral deferred correction methods for ordinary differential equations. *BIT*, 40(2):241–266, 2000.

[7] E. Hairer and G. Wanner. *Solving ordinary differential equations. II: Stiff and differential-algebraic problems. 2nd rev. ed.* Springer, Berlin, 1996.

[8] N.M.M. Huynh, F. Chegini, L. F. Pavarino, M. Weiser, and S. Scacchi. Convergence analysis of BDDC preconditioners for hybrid DG discretizations of the cardiac cell-by-cell model, arXiv:2212.12295 2022.

[9] K.H. Jæger, A.G. Edwards, A. McCulloch, and A. Tveito. Properties of cardiac conduction in a cell-based computational model. *PLoS Comput. Biol.*, 15(5):e1007042, 2019.

[10] M.L. Minion. Semi-implicit spectral deferred correction methods for ordinary differential equations. *Comm. Math. Sci.*, 1(3):471–500, 2003.

[11] M. Potse, L. Cirrottola, and A. Froehly. A practical algorithm to build geometric models of cardiac muscle structure. In *8th European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS)*, Oslo, Norway, 2022.

[12] M. Weiser. Faster SDC convergence on non-equidistant grids by DIRK sweeps. *BIT Numerical Mathematics*, 55(4):1219–1241, 2015.