Chapter 6

Iterative Solvers for EMI Models

Miroslav Kuchta¹ and Kent-André Mardal^{1,2}

Abstract This chapter deals with iterative solution algorithms for the four EMI formulations derived in (17, Chapter 5). Order optimal monolithic solvers robust with respect to material parameters, the number of degrees of freedom of discretization as well as the time-stepping parameter are presented and compared in terms of computational cost. Domain decomposition solver for the single-dimensional primal formulation is discussed.

6.1 Introduction

Spatial discretization of EMI models describing a few cells with a complex/realistic geometry or a large collection of cells leads to linear systems with considerable number of unknowns. In our largest simulations, we may have linear systems involving billions of unknowns at millions of time steps. It is the purpose of this chapter to discuss how such systems can be solved efficiently with available algorithms.

Let us denote the system size as N and let h be a typical grid/mesh size. The complexity of a solution algorithm can then be analyzed in terms of how the computational time grows with N . For instance, a naive Gaussian elimination would perhaps scale as $O(N^3)$ and for linear system involving billions of unknowns such an approach is therefore not feasible during a life-time. Here, we shall aim for algorithms that are *order optimal*, i.e. their solution time scales linearly, $O(N)$, with respect to the number of unknowns. Clearly, $O(N)$ is order optimal in the sense that this is also the complexity of writing the results to file. In general, direct solvers, like for instance LU, do not scale linearly in N . Here, we will use Krylov solvers like the Conjugate

¹Simula Research Laboratory, Norway

²Department of Mathematics, University of Oslo, Oslo, Norway

Gradient (CG) or the Minimal Residual (MinRes) methods. Both methods are known to provide efficient computations when combined with proper preconditioning techniques. Preconditioners based on multigrid and/or domain decomposition methods have been shown to be order-optimal for a variety of problems, but the theory for the EMI problem is currently limited. Efficient solvers for the EMI-problem are discussed in the following.

We say that a solver for the *transient* EMI model (1.30) is order optimal if the solution time grows linearly in Δt^{-1} , where Δt is the time step. Considering the *stationary* problem (5.1), which is to be solved at every step of the temporal loop, it is clear that order optimality of the transient solver requires that the solver of the linear system does not degenerate for small (or large) Δt , i.e. that it is robust with respect to the time stepping.

In this chapter we discuss solution algorithms for the linear systems due to the finite element discretization of (5.1) which are robust in h and well as Δt . Two types of approaches will be considered. Monolithic approaches, where all the unknowns are solved for at once, are the subject of Section 6.2. Section 6.3 then concerns a domain decomposition approach where one iterates between the intra/extra-cellular subproblems. The solvers will be compared in terms of robustness and cost, however, only serial performance will be addressed. We remark that parallel scalable solvers suitable for the mixed formulations of the EMI models (5.12) (5.14) are the balancing domain decomposition methods, see e.g. (24; 25). For the elliptic single-dimensional primal formulation (5.6), the FETI domain decomposition methods, e.g. (13; 18) could be used. Moreover, to simplify the exposition the focus shall be on a single cell model. We remark that all the algorithms presented further can be generalized in a rather straightforward manner to models with multiple *disconnected* cells. However, construction of robust monolithic solvers for *multi-dimensional* formulations for collections of *connected* cells is out of the scope of this manuscript.

6.2 Monolithic Solvers

Let $\mathcal{A}_h x_h = L_h$ be a linear system due to discretization of the continuous problem $\mathcal{A}x = L$ in W' , where W' is a dual space to some Hilbert space W and $\mathcal{A} : W \rightarrow W'$. Note that in case of (5.1) the continuous operator $\mathcal{A}$ depends on material parameters σ_i , σ_e , C_m as well as the time step size Δt . Here the focus is placed on the latter dependence.

The monolithic solvers considered here are preconditioned Krylov methods where the preconditioner $\mathcal{B} : W' \rightarrow W$ shall be constructed such that the number of iterations required for solving problems $\mathcal{B}_h \mathcal{A}_h x_h = \mathcal{B}_h L_h$ is bounded in h and Δt . A constructive framework for establishing $\mathcal{B}$ is operator preconditioning (19), where the structure of the preconditioner reflects mapping properties of $\mathcal{A}$ as an isomorphism

between suitably chosen Hilbert space and its dual space. In particular, if $\mathcal{A} : W \rightarrow W'$ is an isomorphism then a Riesz map with respect to the inner product of W is a suitable preconditioner. As explained in (19), multilevel or domain decomposition algorithms are equivalent to Riesz maps for standard elliptic problems. Furthermore, these schemes also work in a fractional setting (1); a property which will be exploited in this chapter. We shall illustrate the framework briefly below using the single-dimensional formulations (5.6), (5.12) as an example.

In the following we follow the notation of Chapter 5. In particular, $\|\cdot\|_{0,\Omega}$ denotes the L^2 norm of a scalar or vector field over Ω , while $\|\cdot\|_{1,\Omega}$, $\|\cdot\|_{\text{div},\Omega}$ are the $H^1(\Omega)$ and $\mathbf{H}(\text{div},\Omega)$ norms respectively. For a Hilbert space W other than L^2 , H^1 , $\mathbf{H}(\text{div})$ the norm is denoted as $\|\cdot\|_W$ while $(\cdot, \cdot)_W$ denotes the inner product on W . Moreover, we let $(\cdot, \cdot)_\Omega$ be the L^2 inner product. If the domain is clear the subscript will be omitted. By $(\cdot, \cdot)$ we shall also denote a duality pairing between a Hilbert space and its dual. The dual of an operator A with respect to the L^2 inner product is then denoted as A' . Finally, let us introduce scaled, sum and intersection spaces, which will be required for well-posedness of some of the formulations. If W, Q are Hilbert spaces and a an arbitrary positive real number, the scaled space aW , the sum space $W + Q$ and the intersection space $W \cap Q$ are Hilbert spaces with norms $\|x\|_{aW} = a\|\cdot\|_W$,

$$\|x\|_{W+Q} = \inf_{\substack{w+q=x \\ w \in W, q \in Q}} \sqrt{\|w\|_W^2 + \|q\|_Q^2} \quad \text{and} \quad \|x\|_{W \cap Q} = \sqrt{\|x\|_W^2 + \|x\|_Q^2}.$$

6.2.1 Single-Dimensional Primal Solvers

In order to simplify the analysis let σ_i , σ_e and C_m be positive constants and let us consider a slightly modified (cf. the underlined term) single-dimensional primal formulation (5.6): Find $u_i \in V_i = H^1(\Omega_i)$, $u_e \in V_e = H_{\Gamma_e}^1(\Omega_e)$ such that for all $v_e \in V_e$, $v_i \in V_i$

$$\begin{aligned} & (\sigma_e \nabla u_e, \nabla v_e) + \frac{C_m}{\Delta t} (T_e u_e, T_e v_e) - \frac{C_m}{\Delta t} (T_i u_i, T_e v_e) = \frac{C_m}{\Delta t} (f, v_e), \\ & -\frac{C_m}{\Delta t} (T_e u_e, T_i v_i) + (\sigma_i \nabla u_i, \nabla v_i) + \underline{(u_i, v_i)} + \frac{C_m}{\Delta t} (T_i u_i, T_i v_i) = -\frac{C_m}{\Delta t} (f, v_i). \end{aligned} \quad (6.1)$$

Here, T_e, T_i are the trace operators $T_i u_i = u_i|_\Gamma$ and $T_e u_e = u_e|_\Gamma$. Letting $W = V_e \times V_i$, $u = (u_e, u_i)$ the problem (6.1) can be stated as $\mathcal{A}u = L$ in W' with the operator $\mathcal{A}$ and functional L defined as

$$\mathcal{A} = \begin{pmatrix} -\nabla \cdot (\sigma_e \nabla) + \frac{C_m}{\Delta t} T_e' T_e & -\frac{C_m}{\Delta t} T_e' T_i \\ -\frac{C_m}{\Delta t} T_i' T_e & -\nabla \cdot (\sigma_i \nabla) + I + \frac{C_m}{\Delta t} T_i' T_i \end{pmatrix} \quad \text{and} \quad L(v) = \frac{C_m}{\Delta t} \begin{pmatrix} (f, v_e) \\ -(f, v_i) \end{pmatrix}. \quad (6.2)$$

Note that the underlined lower order term in (6.1) corresponds to I in $\mathcal{A}$.

To precondition (6.1) using the operator preconditioning framework we proceed to show that $\mathcal{A} : W \rightarrow W'$ is an isomorphism. This statement follows from the Lax-Milgram theorem, see e.g. (6, Ch 2.7). Let $\|\cdot\|_W$ be *some* norm of the space W (we shall see shortly that different norms can be considered leading to different preconditioners). Hence coercivity of $\mathcal{A}$, i.e.

$$\text{There exists } \alpha_* > 0 \text{ such that } \alpha_* \|u\|_W^2 \leq (\mathcal{A}u, u), \quad \forall u \in W \quad (6.3)$$

and boundedness of $\mathcal{A}$, i.e.

$$\text{There exists } \alpha^* > 0 \text{ such that } (\mathcal{A}u, v) \leq \alpha^* \|u\|_W \|v\|_W, \quad \forall u, v \in W \quad (6.4)$$

need to be shown. While the details of the proof are beyond the scope of the current text, we remark that using the space $W = H_{1-D}^1(\Omega_e) \times H^1(\Omega_i)$, we obtained the bound

$$(\mathcal{A}u, u) \leq \max \left(\sigma_e + 2C_e \frac{C_m}{\Delta t}, \sigma_i, 1, 2C_i \frac{C_m}{\Delta t} \right) \|u\|_W^2 \quad (6.5)$$

by trace, Cauchy-Schwarz, and Young inequalities. Hence, the operator $\mathcal{A}$ is bounded. Note, however, that the constant α^* depends on Δt and, in particular, it blows up for small time steps. Then, the lower bound is

$$\begin{aligned} (\mathcal{A}u, u) &= \sigma_e \|\nabla u_e\|_0^2 + \sigma_i \|\nabla u_i\|_0^2 + \|u_i\|_0^2 + \frac{C_m}{\Delta t} \|T_e u_e - T_i u_i\|_0^2 \\ &= \min(1, \sigma_i, \sigma_e) \|u\|_W^2 \end{aligned} \quad (6.6)$$

and the coercivity thus holds with constant $\alpha_* = \min(1, \sigma_i, \sigma_e)$ which is independent of the time step. The condition number of the preconditioned system, using a preconditioner deduced from W , is then $\frac{\alpha^*}{\alpha_*}$ and from this theoretical consideration we may then expect the number of iterations to grow linearly as Δt decreases.

The exact preconditioner based on W is block diagonal operator

$$\mathcal{B} = \begin{pmatrix} -\nabla \cdot (\sigma_i \nabla) & 0 \\ 0 & -\nabla \cdot (\sigma_e \nabla) + I \end{pmatrix}^{-1}. \quad (6.7)$$

Observe that the preconditioner is efficient in the sense that its evaluation amounts to solving two smaller subproblems posed on Ω_i and Ω_e respectively. However, from our analysis we expect the performance of $\mathcal{B}$ to deteriorate for small time steps. Indeed, this behavior is confirmed by results¹ in Table 6.1. Therein it can be seen

¹ All the experiments in this chapter are conducted using the manufactured problem (5.17) considered in the convergence study in Chapter 5. In particular, we use uniform meshes obtained by dividing the unit square $(\Omega_e \cup \Omega_i)$ first into n_h^2 squares each of which is then subdivided to 2 isosceles triangles with diameter h . The finite element discretization is as described in Chapter 5. We set $\sigma_i = 1$, $\sigma_e = 2.2$ and $C_m = 1$.

The discrete linear systems $\mathcal{A}_h x_h = L_h$ are solved iteratively using the CG or MinRes methods. Implementation of the methods is provided by (2). The solvers are always started from a random initial vector. As a convergence criterion the relative preconditioned residual norm is required to

that the iterations stabilize with h for $\Delta t > 10^{-4}$ while for smaller time steps finer meshes are needed to attain the bounds.

In an attempt to find a more robust preconditioner let us observe that estimates (6.5), (6.6) show that $u \mapsto (\mathcal{A}u, u)^{1/2}$ defines a different norm on W . Let now $\|u\|_W = (\mathcal{A}u, u)^{1/2}$. Then $(\mathcal{A}u, u) = \|u\|_W^2$ and $(\mathcal{A}u, v) \leq (\mathcal{A}u, u)^{1/2}(\mathcal{A}v, v)^{1/2}$ so that the conditions (6.3) and (6.4) hold with $\alpha^* = 1$, $\alpha_* = 1$. In another words, a Riesz map preconditioner with respect to the inner product $(u, v)_W = (\mathcal{A}u, v)$ is independent of Δt . As the Riesz map in this case is in fact $\mathcal{A}^{-1}$ its exact evaluation (by LU) is not feasible. However, for the purpose of preconditioning, it suffices to replace the mapping by a spectrally equivalent operator. Then, provided that the equivalence is robust in Δt the approximate operator will lead to iterations bounded in time step and the discretization parameter. We remark that since the preconditioner takes the entire $\mathcal{A}$ into account, including the off-diagonal terms cf. block diagonal operator (6.7), we shall refer to it as monolithic.

$\Delta t \backslash n_h$	Diagonal (6.7)						Monolithic BoomerAMG($\mathcal{A}$)					
	2^3	2^4	2^5	2^6	2^7	2^8	2^3	2^4	2^5	2^6	2^7	2^8
10^8	8	5	5	5	5	5	8	8	8	8	8	9
10^6	7	7	7	7	7	7	8	8	8	8	8	9
10^4	7	7	7	7	7	7	8	8	8	8	8	9
10^2	6	6	6	6	6	6	7	8	8	8	8	9
1	11	11	11	11	10	10	8	8	8	9	9	9
10^{-2}	26	30	33	32	33	32	8	8	9	9	10	10
10^{-4}	39	70	90	114	139	155	28	26	19	14	11	10
10^{-6}	42	79	134	181	234	300	56	81	101	104	87	62
10^{-8}	42	78	130	179	244	327	81	140	226	320	359	378

Table 6.1: Primal single-dimensional formulation. Number of CG iterations using (left) the diagonal operator (6.7) and (right) the Riesz map with respect to the $\mathcal{A}$ induced norm. Neither preconditioner is robust for $\Delta t \ll 1$.

Since (5.6) leads to symmetric positive-definite matrices, cf. the coercivity condition (6.3), we shall here use algebraic multigrid to construct the approximation of $\mathcal{A}^{-1}$. More precisely, the action of the operator is realized by single V-cycle of BoomerAMG (23). In Table 6.1 it can be seen that this choice leads to h bounded CG iterations for $\Delta t > 10^{-8}$. However, there is a clear sensitivity of the bound to Δt . Thus BoomerAMG approximations of $\mathcal{A}^{-1}$ are not Δt -robust. In fact, estimates of the condition number for the preconditioned problems with $n_h = 2^8$ are 1.2, 1.5, 2.0, 9.2 and 53 for $\Delta t = 1, 10^{-2}, 10^{-4}, 10^{-6}$ and 10^{-8} respectively.

be less than 10^{-12} in magnitude. Unless specified otherwise the preconditioners $\mathcal{B}_h$ are evaluated exactly by LU. Finally, condition number estimates of the preconditioned linear systems are obtained by using iterative Krylov-Schur solver from (10) applied to the generalized eigenvalue problem $\mathcal{A}_h x_h = \lambda_h \mathcal{B}_h^{-1} x_h$. The reported condition number is then $\max |\lambda_h| / \min |\lambda_h|$.

The source code used for the experiments can be found on <https://github.com/MiroK/emi-book-solvers> and is archived at (15).

6.2.2 Single-Dimensional Mixed Solvers

We next consider preconditioners for the single-dimensional mixed formulation (5.12): Find $\mathbf{E} \in S$, $u \in Q$ such that

$$\begin{aligned} (\sigma^{-1} \mathbf{J}, \boldsymbol{\tau}) + \frac{\Delta t}{C_m} (T_n \mathbf{J}, T_n \boldsymbol{\tau}) - (u, \nabla \cdot \boldsymbol{\tau}) &= -(f, T_n \boldsymbol{\tau}), \quad \forall \boldsymbol{\tau} \in S, \\ (q, \nabla \cdot \mathbf{J}) &= 0, \quad \forall q \in Q, \end{aligned}$$

where $S = \mathbf{H}_{\Gamma^N}(\text{div}, \Omega)$, $Q = L^2(\Omega)$ and T_n is the normal trace operator $T_n \boldsymbol{\tau} = \boldsymbol{\tau} \cdot \mathbf{n}_i$. Letting $W = S \times Q$, $x = (\mathbf{J}, u)$ the single-dimensional mixed formulation is equivalent to the problem $\mathcal{A}x = L$ in W' with

$$\mathcal{A} = \begin{pmatrix} \sigma^{-1} I + \frac{\Delta t}{C_m} T_n' T_n & -\nabla \\ \nabla \cdot & 0 \end{pmatrix} \quad \text{and} \quad L(x) = -(f, T_n \mathbf{J}). \quad (6.8)$$

Note that in (6.8) the membrane term $(T_n \mathbf{J}, T_n \boldsymbol{\tau})$ is weighted by $\Delta t/C_m$, cf. (6.2). Thus, unlike in the single-dimensional primal formulation, the term does not dominate for small time steps.

In order to apply the operator preconditioning framework to (6.8) the operator $\mathcal{A}$ needs to be shown to be an isomorphism. To this end we consider $\mathcal{A}$ as an abstract operator over $W = S \times Q$ with the form

$$\mathcal{A} = \begin{pmatrix} A & B' \\ B & 0 \end{pmatrix} \quad \text{where} \quad \begin{aligned} A : S &\rightarrow V', \\ B : S &\rightarrow Q', \end{aligned} \quad (6.9)$$

and apply the Brezzi theory (7). This leads us to the potential preconditioners for the multi-dimensional mixed problem

$$\mathcal{B}_1 = \begin{pmatrix} \sigma^{-1} I - \nabla \nabla \cdot & 0 \\ 0 & I \end{pmatrix}^{-1} \quad \text{and} \quad \mathcal{B}_2 = \begin{pmatrix} \sigma^{-1} I - \nabla \nabla \cdot + \frac{\Delta t}{C_m} T_n' T_n & 0 \\ 0 & I \end{pmatrix}^{-1}, \quad (6.10)$$

which are the Riesz mappings with respect to the inner products of the spaces

$$W_1 = S \times Q \quad \text{and} \quad W_2 = S \cap \sqrt{\frac{\Delta t}{C_m}} N \times Q, \quad (6.11)$$

where $N = \{\mathbf{v} \in S; \|\mathbf{v} \cdot \mathbf{n}\|_{0,\Gamma} < \infty\}$. Note that in W_2 we enforce additional regularity on the normal trace since T_n maps S to $H^{-1/2}(\Gamma)$, e.g. (9, Ch. 2). Moreover the traces are considered in a weighted space.

A consequence of the mapping properties of the normal trace operator is the fact that the term $(T_n \mathbf{J}, T_n \boldsymbol{\tau})$ cannot be bounded using the $\mathbf{H}(\text{div})$ norm and in turn the Brezzi conditions are not satisfied on W_1 . We remark that if such a bound were possible the boundedness constant would depend on Δt , cf. (6.5). As the Brezzi conditions do not hold on W_1 we expect preconditioner $\mathcal{B}_1$ to perform poorly. Indeed, Table 6.2 shows

that the condition numbers of $\mathcal{B}_{1,h}\mathcal{A}_h$ are not bounded in h . The condition numbers for fixed h can be seen to grow with the time step. In Table 6.3 the ill-posedness of the problem in W_1 manifests as unstable iterations or failure of MinRes to converge. We note that for small Δt the iterations appear bounded in h .

$\Delta t \backslash n_h$	$\mathcal{B}_1$						$\mathcal{B}_2$					
	2^2	2^3	2^4	2^5	2^6	2^7	2^2	2^3	2^4	2^5	2^6	2^7
10^2	4223	7540	14766	29437	4785	9569	13.52	13.52	13.52	13.52	13.52	13.52
10^1	75.0	132	259	515	1028	2057	2.28	2.28	2.28	2.28	2.28	2.28
1	8.86	15	27	52	104	207	2.20	2.20	2.20	2.20	2.20	2.20
10^{-1}	2.27	2.85	4.11	6.66	12	22	2.20	2.20	2.20	2.20	2.20	2.20
10^{-2}	2.20	2.20	2.20	2.20	2.59	3.60	2.20	2.20	2.20	2.20	2.20	2.20

Table 6.2: Conditioning of single-dimensional mixed formulation. (Left) Posing the problem in $H(\text{div}) \times L^2$ violates Brezzi conditions. (Right) Preconditioner based on W_2 in (6.11) is inf-sup stable with the inf-sup constant depending on Δt for $\Delta t/C_m > 1$.

Posing (6.8) in W_2 it can be shown that the Brezzi conditions hold with the constants independent of the time step.

We shall not prove the validity of the Brezzi theory here. Instead, Table 6.2 offers numerical evidence that the discrete condition is satisfied. Therein condition numbers of the $\mathcal{B}_2$ preconditioned problems are reported and boundedness in h can be observed. Moreover, it can be seen that the inf-sup constant depends on Δt , in particular, the bound goes to 0 as Δt grows. However, on the subspace $\{(\mathbf{J}, u) \in W_2; (u, 1)_{0,\Omega_t} = 0\}$ the inf-sup condition holds independent of Δt . The single run-away mode appears to have no effect on the MinRes iterations, cf. Table 6.3, see also (20).

$\Delta t \backslash n_h$	$\mathcal{B}_1$							$\mathcal{B}_2$						
	2^3	2^4	2^5	2^6	2^7	2^8		2^3	2^4	2^5	2^6	2^7	2^8	
10^8	–	–	–	–	–	–		10	10	10	10	10	10	
10^6	–	–	–	–	–	–		15	12	12	12	12	12	
10^4	–	–	–	–	–	–		15	15	15	16	16	16	
10^2	156	–	–	–	–	–		15	15	15	15	16	16	
10^1	26	36	57	86	129	209		14	14	14	14	14	14	
1	35	50	72	107	155	234		16	16	16	16	16	16	
10^{-1}	16	19	23	32	44	60		16	17	17	18	18	18	
10^{-2}	18	19	21	27	34	45		19	19	19	19	20	20	
10^{-4}	19	19	20	20	20	20		19	19	20	20	20	21	
10^{-6}	19	19	20	20	21	21		19	19	20	20	21	21	
10^{-8}	19	19	20	20	21	21		19	19	20	20	21	21	

Table 6.3: MinRes iterations with preconditioned single-dimensional mixed formulation using preconditioners (6.10). Failure to converge within 500 iterations is denoted by –. Convergence of $\mathcal{B}_2$ seems unaffected by the single Δt unbounded inf-sup mode.

Considering order optimality of the $\mathcal{B}_2$ based solver we recall that the iterations in Table 6.3 were run with the exact preconditioner. In particular, the $S \cap N$ Riesz map

was computed by LU. In practical computations such realization is not feasible and scalable approximation is needed. For standard $\mathbf{H}(\text{div})$ inner product, i.e. the one used in $\mathcal{B}_1$, the Riesz map can be efficiently realized by multigrid algorithms (14). However, in the authors' experience this approach does not work equally well with the $S \cap N$ inner product.

We have seen in Sections 6.2.1 and 6.2.2 that construction of order optimal monolithic solvers for the single-dimensional EMI formulations presents a challenge. In primal formulation Δt robustness of the monolithic approach was problematic. For the mixed formulation, order optimality required a specialized solver for the Riesz mapping over the subspace of $\mathbf{H}(\text{div})$. These two problems are addressed by solvers for multi-dimensional formulations.

6.2.3 Multi-Dimensional Solvers

In this section we consider the construction of preconditioners for the operators

$$\mathcal{A}_p = \begin{pmatrix} -\nabla \cdot (\sigma_e \nabla) & 0 & -T'_e \\ 0 & -\nabla \cdot (\sigma_i \nabla) & T'_i \\ -T_e & T_i & -\frac{\Delta t}{C_m} I \end{pmatrix} \quad \text{and} \quad \mathcal{A}_m = \begin{pmatrix} \sigma^{-1} I - \nabla & T'_n \\ \nabla \cdot & 0 & 0 \\ T_n & 0 & -\frac{C_m}{\Delta t} I \end{pmatrix}, \quad (6.12)$$

which induce respectively the multi-dimensional primal weak formulation (5.8) and the multi-dimensional mixed weak formulation (5.14). In order to discuss well-posedness let

$$W_p = H^1_{\Gamma^D_e}(\Omega_e) \times H^1(\Omega_i) \times H^{-1/2}(\Gamma) \cap \sqrt{\frac{\Delta t}{C_m}} L^2(\Gamma) \quad (6.13)$$

and

$$W_m = \mathbf{H}_{\Gamma^N_e}(\text{div}, \Omega) \times L^2(\Omega) \times H^{1/2}(\Gamma) \cap \sqrt{\frac{C_m}{\Delta t}} L^2(\Gamma). \quad (6.14)$$

We remark that the fractional space $H^{-1/2}$, in which the membrane current density I_m in (5.6) is sought, and $H^{1/2}$, which is the space of the transmembrane potential v in (5.14), are dictated by the mapping properties of the trace operators T_e, T_i and T_n . For example, as $T_i : H^1(\Omega_i) \rightarrow H^{1/2}(\Gamma)$, $I_m \in (H^{1/2})'$ so that the term $(T_i u_i, I_m)$ is bounded. Note also that only the spaces involving Γ are now Δt -dependent, cf. e.g. W_2 in (6.11).

Due to the Δt -weighted (penalty) terms in $\mathcal{A}_p$ and $\mathcal{A}_m$ the operators cannot be established as isomorphisms on W_p and W_m by straightforward application of the Brezzi theory. Instead, $\mathcal{A}_p$ for $\Delta t/C_m \leq 1$ and $\mathcal{A}_m$ for $\Delta t/C_m \geq 1$ can be analyzed using the framework for saddle point systems with penalty, see (4, Ch. 3.4). A crucial assumption then is that the system without the penalty term satisfies the Brezzi conditions. While here the well-posedness shall be demonstrated only by numerical

experiments, let us point out an important difference between the operators. To this end let $c > 0$ be a constant and consider a piecewise constant potential u such that $u_i = c$, $u_e = 0$ and set $\mathbf{J} = \mathbf{0}$. Then for any $\boldsymbol{\tau} \in \mathbf{H}(\text{div})$ by divergence theorem

$$(\nabla \cdot \boldsymbol{\tau}, u)_{\Omega} + (T_{\mathbf{n}} \boldsymbol{\tau}, v) = (\nabla \cdot \boldsymbol{\tau}, u)_{\Omega_i} + (T_{\mathbf{n}} \boldsymbol{\tau}, v) = c(T_{\mathbf{n}} \boldsymbol{\tau}, 1) + (T_{\mathbf{n}} \boldsymbol{\tau}, v)$$

and in turn, if $v = -c$ we have that

$$\begin{pmatrix} \sigma^{-1} I & -\nabla T'_{\mathbf{n}} \\ \nabla \cdot & 0 & 0 \\ T_{\mathbf{n}} & 0 & 0 \end{pmatrix} \begin{pmatrix} \mathbf{J} \\ u, \\ v \end{pmatrix} = \begin{pmatrix} c(T_{\mathbf{n}} \boldsymbol{\tau}, 1) + (T_{\mathbf{n}} \boldsymbol{\tau}, v) \\ 0 \\ 0 \end{pmatrix} = \begin{pmatrix} \mathbf{0} \\ 0 \\ 0 \end{pmatrix}.$$

Thus the operator $\mathcal{A}_m$ without the penalty term is singular with a one dimensional kernel spanned by vector $(\mathbf{0}, u, -c)$ and the Brezzi conditions do not hold. We shall see implications of this property on the convergence of the iterative solvers shortly. We remark that $\mathcal{A}_p$ without the penalty term is non-singular.

Assuming that $\mathcal{A}_p : W_p \rightarrow W'_p$ is an isomorphism we consider as a preconditioner for the single-dimensional primal formulation the operator

$$\mathcal{B}_p = \begin{pmatrix} -\nabla \cdot (\sigma_e \nabla) & 0 & 0 \\ 0 & -\nabla \cdot (\sigma_i \nabla) + I & 0 \\ 0 & 0 & (-\Delta + I)^{-1/2} + \frac{\Delta t}{C_m} I \end{pmatrix}^{-1}. \quad (6.15)$$

Observe that the preconditioner consists of (approximate) solvers for three decoupled subproblems posed on Ω_e , Ω_i and Γ . Moreover, the problems on extra and intracellular domains are standard elliptic operators for which efficient (black-box) multigrid techniques exist, e.g. (23). The problem on the interface is then less standard as it concerns a fractional Helmholtz operator. However, efficient multilevel solvers have been established e.g. in (5; 1). We remark that if the discrete spaces for I_m (or v) contain only a few thousands of degrees of freedom an eigenvalue realization of the fractional preconditioner is feasible cf. (16). As the interfacial spaces here are small, cf. Table (5.1), we use further the spectral approach.

$\Delta t \backslash n_h$		n_h						
		2^2	2^3	2^4	2^5	2^6	2^7	2^8
Δt	10^8	10	10	10	10	10	10	10
	10^6	10	10	10	10	10	10	10
	10^4	15	15	15	15	15	15	15
	10^2	18	18	18	17	17	17	17
	1	27	27	27	26	26	26	26
Δt	10^{-2}	43	47	46	44	44	43	43
	10^{-4}	49	60	62	60	57	56	56
	10^{-6}	51	60	62	61	59	59	59
	10^{-8}	49	60	62	61	59	59	59

Table 6.4: Fractional preconditioner (6.15) for multi-dimensional primal formulation. (Left) Condition numbers are bounded in h . Growth in Δt for $\Delta t > 1$ is caused by a single mode. (Right) MinRes iterations counts are bounded in h and time step.

Using (6.15) Table 6.4 reports the condition numbers of the preconditioned problems $\mathcal{B}_{p,h}\mathcal{A}_{p,h}$. It can be seen that the results are bounded in h , thus providing a numerical evidence for $\mathcal{A}_p$ being an isomorphism on W_p in (6.13). Moreover, it can be seen that the conditioning of the problem deteriorates with Δt large. However, in this case it is only a single mode $u_i = 0$, $u_e = \text{const}$, $I_m = 0$ which causes the blowup. Considering the MinRes iteration counts, the presence of this mode seems to have almost no impact on boundedness in h or Δt .

We finally come back to the multi-dimensional mixed formulation. Based on the space W_m in (6.14) let a preconditioner for the multi-dimensional mixed formulation be

$$\mathcal{B}_m = \begin{pmatrix} \sigma^{-1}I - \nabla \nabla \cdot & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & (-\Delta + I)^{1/2} + \frac{C_m}{\Delta t}I \end{pmatrix}^{-1}. \quad (6.16)$$

We remark that the preconditioner uses a standard $\mathbf{H}(\text{div})$ inner product, cf. $\mathcal{B}_2$ in (6.10), which can be efficiently realized by multigrid methods, e.g. (14). Note also that the fractionality of the Laplacian is $1/2$, cf. $-1/2$ of the multi-dimensional primal preconditioner (6.15). In addition to the previously mentioned multilevel methods, problems $(-\Delta u + u)^s x = b$ for $0 < s < 1$ can be efficiently solved by a number of approaches, see (3) and references therein.

Using (6.16) Table 6.5 reports the condition numbers of the preconditioned multi-dimensional mixed formulation. The conditioning can be seen to be stable in h , while in agreement with the limit singularity property, there is a growth with Δt . Given that only a single mode is responsible for the lack of Δt -boundedness and recalling results of Table 6.4 or 6.3 we might expect that also here the MinRes solver will not be affected. However, Table 6.6 shows that this is not the case. In fact, as Δt grows the iterations become unstable in h . Figure 6.1 then shows a typical convergence behavior of the solver. We see that the relative preconditioned residual norm is quickly reduced to about 10^{-9} in approximately 30 iterations (regardless of the mesh resolution). Afterwards the convergence stalls. We conclude that for robust preconditioning the nullspace of $\mathcal{A}_m$ must be addressed.

$\Delta t \backslash n_h$	$\mathcal{B}_m \mathcal{A}_m$						$\mathcal{B}_m^0 \mathcal{A}_m^0$					
	2^2	2^3	2^4	2^5	2^6	2^7	2^2	2^3	2^4	2^5	2^6	2^7
10^2	17566	17627	17733	17810	17857	17882	4.40	4.42	4.51	4.58	4.63	4.67
10^1	177	178	179	180	180	180	4.39	4.41	4.50	4.57	4.62	4.66
1	3.31	3.73	3.96	4.09	4.18	4.27	3.31	3.73	3.96	4.09	4.18	4.27
10^{-1}	1.08	1.16	1.30	1.56	1.96	2.48	2.61	2.61	2.61	2.61	2.64	3.14
10^{-2}	1.03	1.03	1.03	1.03	1.03	1.03	2.62	2.62	2.62	2.62	2.62	2.62

Table 6.5: Conditioning of the multi-dimensional mixed formulation. (Left) Problem considered on W_m in (6.14) yields operator $\mathcal{A}_m$ with a singular limit as $\Delta t \rightarrow \infty$. (Right) Formulation including an additional constraint on the transmembrane potential $\int_{\Gamma} v \, dS = 0$ yields h and Δt robust condition numbers.

Motivated by the observation that $\mathcal{A}_m$ becomes singular in the limit $\Delta t = \infty$ with a mode $\mathbf{J} = \mathbf{0}$, $u_e = 0$, $u_i = c$, $v = -c$ in the kernel, let $W_e^0 = W_e \times \mathbb{R}$ and let us define a new solution operator for the multi-dimensional mixed formulation $\mathcal{A}_m^0 : W_e^0 \rightarrow (W_e^0)'$ and its preconditioner $\mathcal{B}_m^0 : (W_e^0)' \rightarrow W_e^0$ as

$$\mathcal{A}_m^0 = \begin{pmatrix} \sigma^{-1}I - \nabla & T'_n & 0 \\ \nabla \cdot & 0 & 0 \\ T_n & 0 & -\frac{C_m}{\Delta t}I \\ 0 & 0 & I \end{pmatrix}, \mathcal{B}_m^0 = \begin{pmatrix} \sigma^{-1}I - \nabla \nabla \cdot & 0 & 0 \\ 0 & I & 0 \\ 0 & 0 & (-\Delta + I)^{1/2} + \frac{C_m}{\Delta t}I \\ 0 & 0 & 0 \end{pmatrix}^{-1} \mu I, \quad (6.17)$$

where $\mu = \min(1, C_m/\Delta t)$. Note that $\mathcal{A}_m^0$ includes an additional unknown, a single scalar, which enforces $(1, v)_\Gamma = 0$. The new constraint thus eliminates constant transmembrane potential and in turn the new operator is non-singular². Considering results reported in Tables 6.5, 6.6 it can be seen that the new preconditioned formulation leads to condition numbers and iteration counts, which are bounded in both the mesh size and the time step. Note that the extra unknown has only a small impact on the number of iterations compared to the $\mathcal{B}_m$ preconditioner.

$\Delta t \backslash n_h$	$\mathcal{B}_m \mathcal{A}_m$						$\mathcal{B}_m^0 \mathcal{A}_m^0$					
	2 ³	2 ⁴	2 ⁵	2 ⁶	2 ⁷	2 ⁸	2 ³	2 ⁴	2 ⁵	2 ⁶	2 ⁷	2 ⁸
10 ⁸	457	–	241	–	–	–	28	32	34	36	37	37
10 ⁶	54	174	82	301	258	227	28	32	34	36	37	37
10 ⁴	37	41	45	45	47	47	28	32	34	36	37	37
10 ²	32	37	39	39	41	41	28	32	34	36	37	37
1	28	30	32	34	34	35	30	33	35	35	37	38
10 ⁻²	19	22	25	29	33	35	20	23	29	33	35	39
10 ⁻⁴	10	10	10	11	11	12	12	12	12	13	13	14
10 ⁻⁶	7	7	9	9	9	9	9	9	10	10	10	10
10 ⁻⁸	7	7	7	7	7	7	9	9	9	9	9	9

Table 6.6: MinRes iterations for the multi-dimensional mixed formulation. No convergence in 500 iterations is indicated by –. (Left) For large Δt the iterations are unstable since operator $\mathcal{A}_m$ becomes singular in the limit as $\Delta t \rightarrow \infty$. (Right) Constraining transmembrane potential the operator $\mathcal{A}_m^0$ does not have the limit singular property.

² A physical motivation for the constraint can be found in considering the EMI interface equation $C_m \frac{\partial v}{\partial t} = I_m + I_{\text{ion}}$ on Γ . By integrating the left-hand side we obtain $C_m \frac{d}{dt}(1, v)_\Gamma$. Recall that $I_m = \sigma_i \nabla u_i \cdot \mathbf{n}_i$ on Γ and $-\nabla \cdot (\sigma_i \nabla u_i) = 0$ in Ω_i . Then, setting $I_{\text{ion}} = 0$ and integrating the right-hand side we have

$$(1, I_m + I_{\text{ion}})_\Gamma = (1, I_m)_\Gamma = (1, -\nabla \cdot (\sigma_i \nabla u_i))_{\Omega_i} = 0.$$

Thus we obtain a conservation relation $C_m \frac{d}{dt}(1, v)_\Gamma = 0$.

Fig. 6.1: (Left) Convergence of the MinRes method for (6.12)-based multi-dimensional mixed formulation with $\Delta t = 10^{-8}$ and problem from Table 6.6. For all mesh resolutions the norm is reduced in cca. 30 iterations. Afterwards the reduction stalls. (Right) Convergence of the Robin/Neumann domain decomposition algorithm. For small Δt the algorithm becomes sensitive to mesh size and can even diverge.

6.3 Domain Decomposition Solvers

Having seen that monolithic solvers for the EMI equations can be sensitive to spatial and temporal resolution we next briefly discuss robustness of the non-overlapping domain decomposition (DD) approach. The focus here shall only be on the single-dimensional primal formulation and the Robin/Neumann DD algorithm in the form presented in (12, Chapter 4). In particular, our implementation shall not include any coarse space, cf. (21), or a preconditioner, see e.g. (8) for interpretation of DD as Steklov-Poincaré operators in fractional Sobolev spaces.

For the sake of self-containedness we review the variational formulation of the Robin/Neumann algorithm. Let $V_e = H^1_{\Gamma_e}(\Omega_e)$, $V_i = H^1(\Omega_i)$ and v^0, u_e^0 be the given initial transmembrane and extracellular potentials. A single iteration of the DD algorithm then produces a new approximation v^1 in three steps. (i) We find $u_i \in V_i$ such that

$$(\sigma_i \nabla u_i, \nabla v_i) + \frac{C_m}{\Delta t} (T_i u_i, T_i v_i) = -\frac{C_m}{\Delta t} (f, v_i) + \frac{C_m}{\Delta t} (T_e u_e^0, T_i v_i), \quad \forall v_i \in V_i.$$

(ii) Using the computed intracellular potential the new extracellular potential $u_e \in V_e$ is found such that

$$(\sigma_e \nabla u_e, \nabla v_e) = (-\sigma_i \nabla u_i \cdot n_i, T_e v_e), \quad \forall v_e \in V_e.$$

(iii) Finally $v^1 = T_i u_i - T_e u_e$. The iterations continue by assigning $v^0 = v^1$ and $u_e^0 = u_e$ until convergence, which in the following is determined as $\|v^k - v^{k-1}\|_0 / \|v^{k-1}\|_0 < \epsilon$, where $\epsilon = 10^{-5}$. We remark that the tolerance was chosen such that in the refinement

studies, where convergence of DD was observed, the approximation error in the $H^1 \times H^1$ norm decreased with the expected order as $O(h)$ (we recall that $P_1 - P_1$ elements were used).

To test robustness of the DD solver we consider the setup of the single cell experiment used with the monolithic approaches in Section 6.2. Here the Robin problem (i) and the Neumann problem (ii) shall be solved by LU to eliminate effects of inexact subdomain solvers. We remark that due to the conforming triangulation, see Figure 5.1, and P_1 discretization of both V_i and V_e the discrete transmembrane potential is computed simply by interpolation. Figure 6.1 plots evolution of the potential difference $\|v_h^k - v_h^{k-1}\|_0 / \|v_h^{k-1}\|_0$ with DD iterations for several mesh resolutions and time step values $\Delta t \leq 1$. It can be seen that for $\Delta t = 1$ the algorithm converges in about 5 iterations irrespective of the spatial discretization. For smaller timesteps convergence is delayed on finer meshes and for $\Delta t = 10^{-4}$ the iterations diverge.

Considering the results in Figure 6.1 we conclude that in the form presented here domain decomposition is not a robust algorithm for the single-dimensional primal formulation of the EMI equations. However, due to its simplicity (relative to e.g. the multi-dimensional formulations) and speed (see Section 6.4), DD might be the method of choice in practical cardiac modeling. It is therefore natural to ask whether the divergence conditions are likely to arise in real applications. To address this question we perform a scaling analysis based on the membrane dynamical condition $C_m \frac{\partial v}{\partial t} \sim \sigma_i \nabla u_i \cdot \mathbf{n}_i$. Letting L be a characteristic length scale and $\nabla' = L \nabla$ the equation becomes $(LC_m / \sigma_i) \frac{\partial v}{\partial t} = \nabla' u_i \cdot \mathbf{n}_i$ and the pre-factor $T = LC_m / \sigma_i$ can be seen to have the unit of seconds. Following (22) let us insert $C_m = 1 \mu\text{F}/\text{cm}^2$, $\sigma_i = 10 \text{mS}/\text{cm}$, $L = 100 \mu\text{m}$, where the length scale is determined by the size of a typical cell used by the authors'. We remark that this a sensible choice given the setup of our experiments. As a result $T = 10^{-6} \text{s}$ and thus $\Delta t = 1$ in the experiments reported here corresponds to a time step of 10^{-6} seconds. Note that this is the finest time step considered in (11). Furthermore, therein the spatial resolution is $2 \mu\text{m}$ while here with $n_h = 2^8$ and $L = 100 \mu\text{m}$ the mesh size is approximately $0.4 \mu\text{m}$. In summary, the conditions for divergence of the DD algorithm discussed here can be encountered not far away from the parameter regime in (11).

6.4 Solver Comparison

To complete our discussion of EMI solvers we finally address the speed of the different algorithms. Recall that until this point results for all solvers, but the monolithic single-dimensional primal one, were obtained using LU in the construction. Such an implementation, however, is not scalable. Here we show that the proposed algorithms are order optimal if LU is replaced by suitable multilevel methods.

Referring to the legend of Figure 6.2 we shall compare 8 different methods. Single-dimensional primal formulation (sp) shall be solved either with the BoomerAMG($\mathcal{A}$)-preconditioner (mono), or diagonal preconditioner (6.7). In the latter case all the blocks of the diagonal preconditioner are approximated by single algebraic multigrid V-cycle. The single-dimensional mixed (sm) formulations shall use the $\mathcal{B}_1$ preconditioner (6.10) with the leading block approximated by $\mathbf{H}(\text{div})$ multigrid HypreAMS of (14). We recall that the solver in general is not independent of the discretization, cf. Table 6.3, however, HypreAMS does not work well for the robust preconditioner $\mathcal{B}_2$. Finally, the multi-dimensional primal (mp) and multi-dimensional mixed (mm) formulations shall use (6.15) and (6.16) with the Ω_i , Ω_e and Ω subproblems of the preconditioners approximated by multigrid (BoomerAMG in case of (6.15) and HypreAMS for (6.16)). The fractional operators will be computed exactly. In addition, two implementations of Robin/Neumann domain decomposition algorithm are considered with the subproblems solved either exactly by LU or by 4 V-cycles of BoomerAMG. Finally, a reference solution time is provided by timings of exact solution of the single-dimensional formulation (sp-LU).

We compare the algorithms using the single cell setup of the previous experiments with $\Delta t = 10^{-3}$ where this value was chosen with the intention to not favor any of the methods. Using tolerance $\epsilon = 10^{-3}$ for DD and 10^{-12} for CG/MinRes, Figure 6.2 reports solution times³ of the algorithms for $2^3 \leq n_h \leq 2^{10}$. It can be seen that with the exception of LU-based solvers all the methods are indeed order optimal. The monolithic approach (sp-mono) and the multi-dimensional primal (mp) solver are then the fastest solvers while solvers for the mixed formulations (sm- $\mathcal{B}_1$, mm) are the slowest. For $n_h = 2^{10}$ the solution times of the solvers were cca. 26s, 140s, 640s, 570s respectively. We remark that the mixed formulations have approximately 5 times more unknowns compared to the single-dimensional ones. Thus the cost per degree of freedom is comparable between the two approaches. Considering the scalable DD implementation (DD-AMG) the timing for $n_h = 2^{10}$ is cca. 160s. The method thus has a similar cost to multi-dimensional primal formulation.

³ The timings include setup times of the preconditioners for monolithic methods. For DD the subproblems were assembled and their solvers constructed only once.

Fig. 6.2: Solution times (including preconditioner setup) of the solvers for the EMI system (5.1) in terms of number of unknowns N or mesh size h . Single cell model with $\Delta t = 10^{-3}$ is used. All but multi-dimensional primal solver with LU (sp-LU) and LU-based domain decomposition solver (DD-LU) are order optimal. The former scales as $N^{3/2}$. Legend is shared between the figures.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

References

1. Bærland T, Kuchta M, Mardal KA (2019) Multigrid methods for discrete fractional Sobolev spaces. *SIAM Journal on Scientific Computing* 41(2):A948–A972
2. Balay S, Abhyankar S, Adams MF, Brown J, Brune P, Buschelman K, Dalcin L, Dener A, Eijkhout V, Gropp WD, Karpeyev D, Kaushik D, Knepley MG, May DA, McInnes LC, Mills RT, Munson T, Rupp K, Sanan P, Smith BF, Zampini S, Zhang H, Zhang H (2019) PETSc Web page. <https://www.mcs.anl.gov/petsc>, URL <https://www.mcs.anl.gov/petsc>
3. Bonito A, Pasciak J (2015) Numerical approximation of fractional powers of elliptic operators. *Mathematics of Computation* 84(295):2083–2110
4. Braess D (2007) *Finite Elements: Theory, Fast Solvers, and Applications in Solid Mechanics*. Cambridge University Press, Cambridge
5. Bramble J, Pasciak J, Vassilevski P (2000) Computational scales of Sobolev norms with application to preconditioning. *Mathematics of Computation* 69(230):463–480
6. Brenner S, Scott R (2007) *The mathematical theory of finite element methods*, vol 15. Springer Science & Business Media
7. Brezzi F (1974) On the existence, uniqueness and approximation of saddle-point problems arising from Lagrangian multipliers. *Revue française d’automatique, informatique, recherche opérationnelle Analyse numérique* 8(2):129–151
8. Discacciati M, Quarteroni A, Valli A (2007) Robin–Robin domain decomposition methods for the Stokes–Darcy coupling. *SIAM Journal on Numerical Analysis* 45(3):1246–1268
9. Girault V, Raviart PA (2012) *Finite Element Methods for Navier-Stokes Equations: Theory and Algorithms*, vol 5. Springer Berlin Heidelberg, Berlin, Heidelberg
10. Hernandez V, Roman JE, Vidal V (2005) SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Trans Math Software* 31(3):351–362
11. Jæger KH, Tveito A (2020) Efficient numerical solution of the EMI model representing extracellular space (E), the cell membrane (M) and the intracellular space (I) of a collection of cardiac cells. Preprint
12. Jæger KH, Hustad KG, Cai X, Tveito A (2020) Operator splitting and finite difference schemes for solving the EMI model. In: Tveito A, Mardal KA, Rognes ME (eds) *Modeling excitable tissue - The EMI framework*. Simula Springer Notes in Computing, SpringerNature
13. Klawonn A, Widlund OB, Dryja M (2002) Dual-primal FETI methods for three-dimensional elliptic problems with heterogeneous coefficients. *SIAM Journal on Numerical Analysis* 40(1):159–179
14. Kolev TV, Vassilevski PS (2012) Parallel auxiliary space AMG solver for $H(\text{div})$ problems. *SIAM Journal on Scientific Computing* 34(6):A3079–A3098
15. Kuchta M, Mardal KA (2020) Software for EMI - Iterative solvers for EMI models. URL <https://doi.org/10.5281/zenodo.3771212>

16. Kuchta M, Nordaas M, Verschaeve JCG, Mortensen M, Mardal KA (2016) Preconditioners for saddle point systems with trace constraints coupling 2D and 1D domains. *SIAM Journal on Scientific Computing* 38(6):B962–B987
17. Kuchta M, Mardal KA, Rognes ME (2020) Solving the EMI equations using finite element methods. In: Tveito A, Mardal KA, Rognes ME (eds) *Modeling excitable tissue - The EMI framework*, Simula Springer Notes in Computing, SpringerNature
18. Mandel J, Dohrmann CR, Tezaur R (2005) An algebraic theory for primal and dual substructuring methods by constraints. *Applied Numerical Mathematics* 54(2):167–193
19. Mardal KA, Winther R (2011) Preconditioning discretizations of systems of partial differential equations. *Numerical Linear Algebra with Applications* 18(1):1–40
20. Nielsen BF, Mardal KA (2013) Analysis of the minimal residual method applied to ill posed optimality systems. *SIAM Journal on Scientific Computing* 35(2):A785–A814
21. Smith B, Bjorstad P, Gropp W (2004) *Domain Decomposition: Parallel Multilevel Methods for Elliptic Partial Differential Equations*. Cambridge University Press, Cambridge
22. Tveito A, Jæger KH, Kuchta M, Mardal KA, Rognes ME (2017) A cell-based framework for numerical modeling of electrical conduction in cardiac tissue. *Frontiers in Physics* 5:48
23. Yang UM, et al. (2002) Boomerang: a parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics* 41(1):155–177
24. Zampini S (2016) PCBDDC: a class of robust dual-primal methods in PETSc. *SIAM Journal on Scientific Computing* 38(5):S282–S306
25. Zampini S, Tu X (2017) Multilevel balancing domain decomposition by constraints deluxe algorithms with adaptive coarse spaces for flow in porous media. *SIAM Journal on Scientific Computing* 39(4):A1389–A1415