A Distributed-Memory Framework for the EMI Model

F.Chegini, A. Neic, T. Steinke, M. Weiser

ZIB ZUSE INSTITUTE BERLIN

MICROCARD

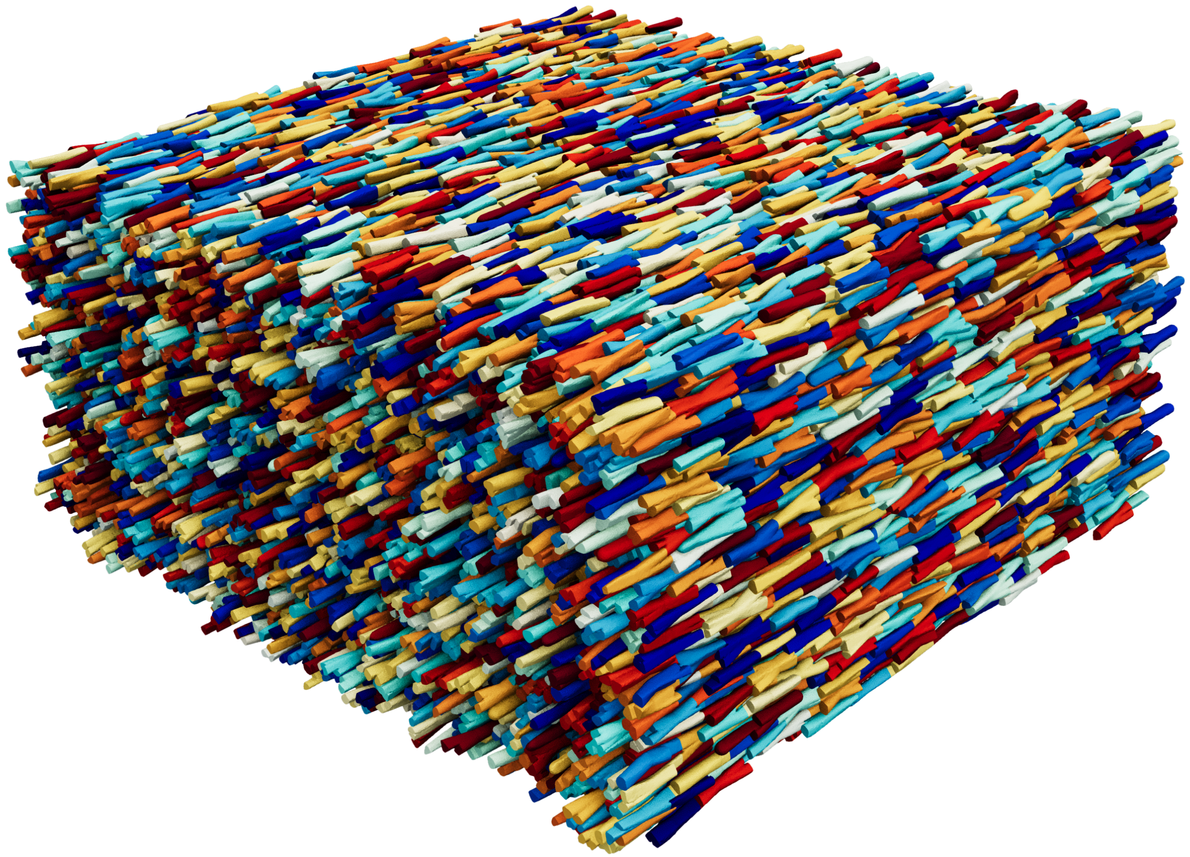

With funding from the:

Federal Ministry of Research, Technology and Space

Co-funded by the European Union

EuroHPC Joint Undertaking

This work was supported by the MICROCARD-2 project (project ID 101172576). The project is supported by the EuroHPC Joint Undertaking and its members (including top-up funding by ANR, BMBF, and Ministero dello sviluppo economico). Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or EuroHPC. Neither the European Union nor EuroHPC can be held responsible for them.

MICROCARD Workshop, June 10th , 2026, Strasbourg

Motivation

disorganization of cardiac tissue at the cellular scale affects propagation of the action potential (AP)
at tissue scale (e.g., in arrhythmia, aging, infarction scars, microscopic conduction block, local heterogeneity)
→ realistic microscopy-based geometries

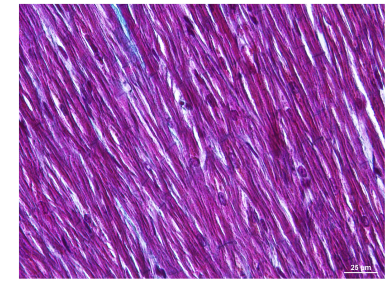

Microscopic view of the tissue

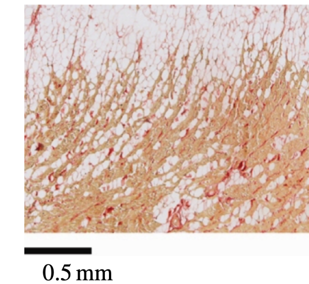

Cardiac muscle infiltrated by fatty tissue

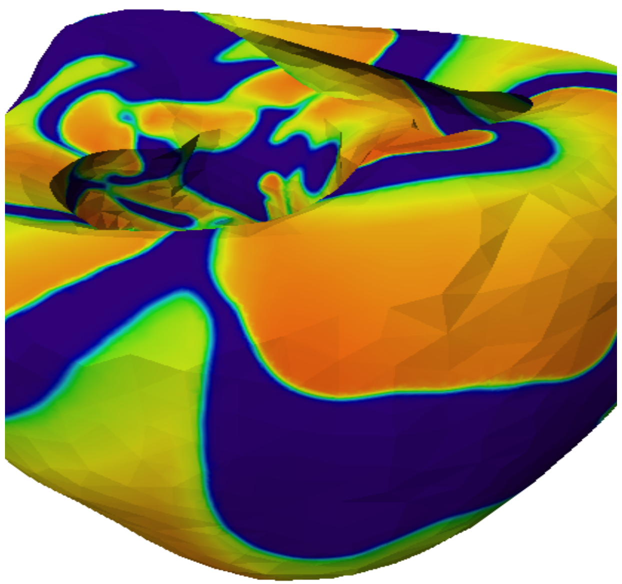

Organ scale
conduction & excitation patterns
mono/bidomain models
100 µm

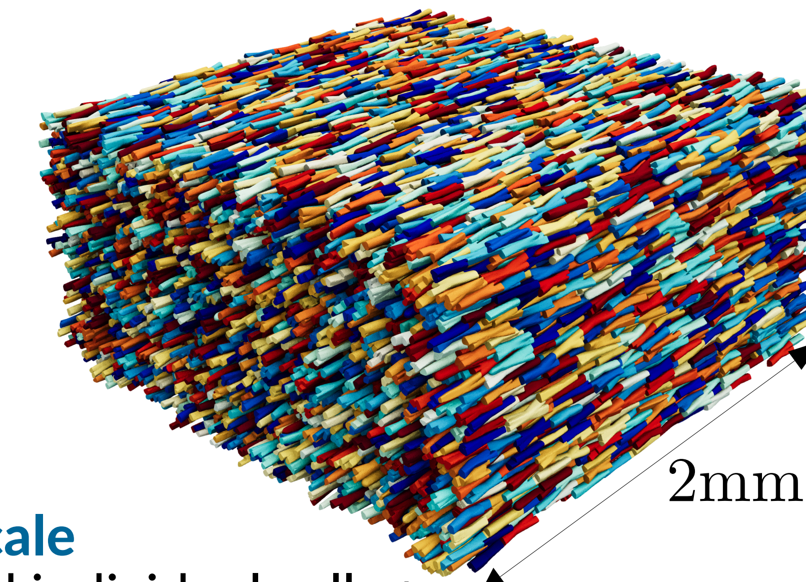

Cellular scale
Discretized individual cells + their interconnections
EMI model
4 µm

EMI Model

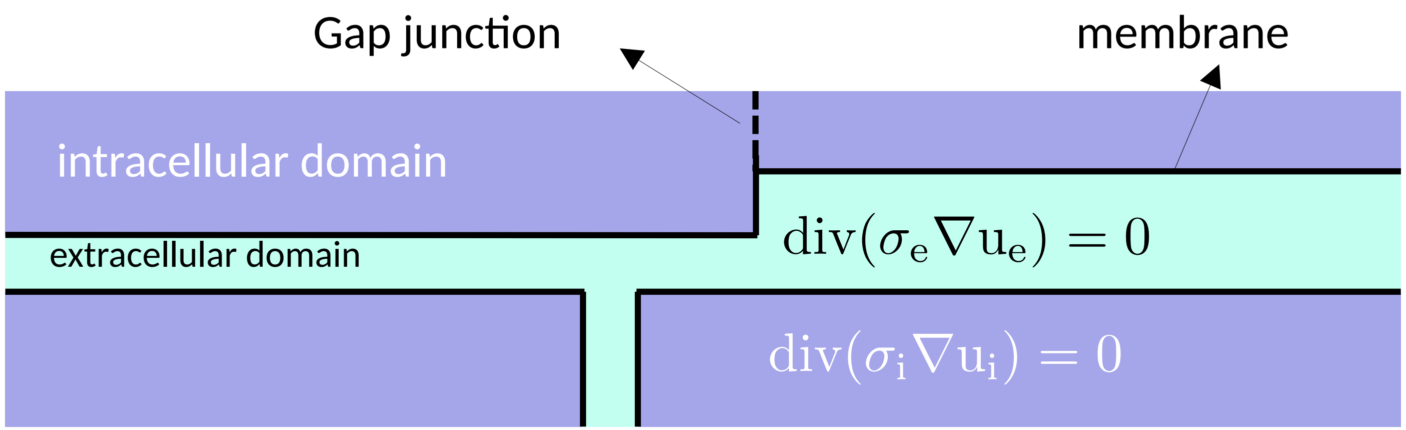

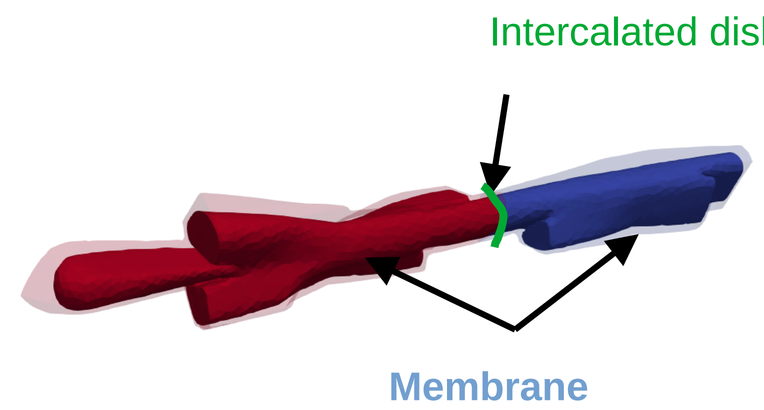

- EMI resolves intracellular space, extracellular space, membranes, and intercellular interfaces explicitly.

$$\mathrm{div}(\sigma_\mathrm{i} \nabla \mathrm{u}_\mathrm{i}) = 0 \quad \text{in } \Omega_i$$

- The transmembrane voltage is a jump.

$$v = u_i - u_j \quad \text{on } \Gamma_{ij} = \overline{\Omega_i} \cap \overline{\Omega_j},\ i > j$$

- The membrane couples elliptic volume equations to ionic dynamics:

$$I_m = C_\mathrm{m} \frac{\partial v}{\partial t} + I_\mathrm{ion}(v, w) - I_\mathrm{stim}.$$

$$\dot{w} = F(v, w) \quad \text{on } \Gamma_{ij}$$

$$\int_0^T \left(\int_{\Omega_e} \sigma_\mathrm{e} \nabla u_e \cdot \nabla \phi_e dx + \int_{\Omega_i} \sigma_\mathrm{i} \nabla u_i \cdot \nabla \phi_i dx + \int_{\Gamma_{ij(i \neq j)}} (C \dot{v}_{ij} + I_{ij}^{\mathrm{ion}}(v_{ij}, w_{ij})) \varphi dx \right) dt = 0$$

Challenges on Current computation

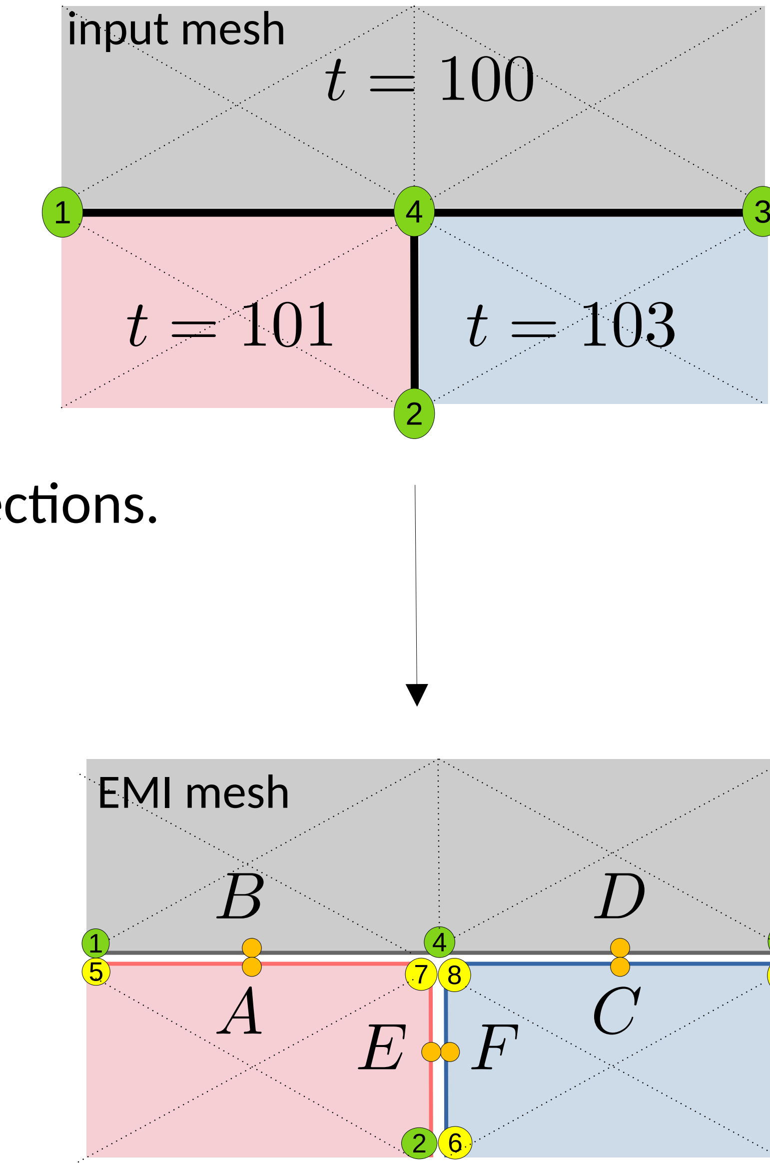

- one physical location on interface may carry several different intersections.
  $\rightarrow$ current computation on the face to avoid ambiguity
- Interface handling, numbering, assembling
- Mapping between volume mesh to the interface mesh

Mesh Processing and Matrix Assembling

Tagged volume mesh → Distributed region → Surface extraction → Interface decoupling → Matrices/operator

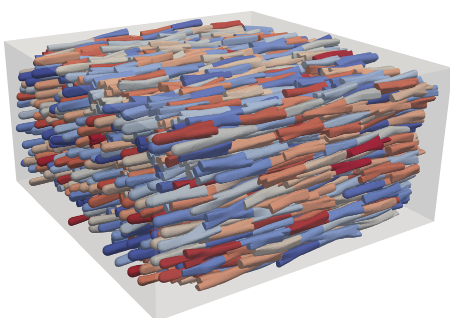

Two interface representations

Γ two-sided interface for volume–surface assembly.

γ unique physical faces for ionic states and currents.

Challenges in large scale simulation

- The number of communication during the mesh processing
- Efficient memory usage → Memory preallocation for PETSc library
- solver/preconditioner (PETSc/Ginkgo)

Computation at each time step

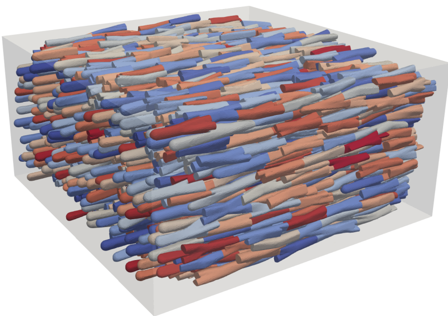

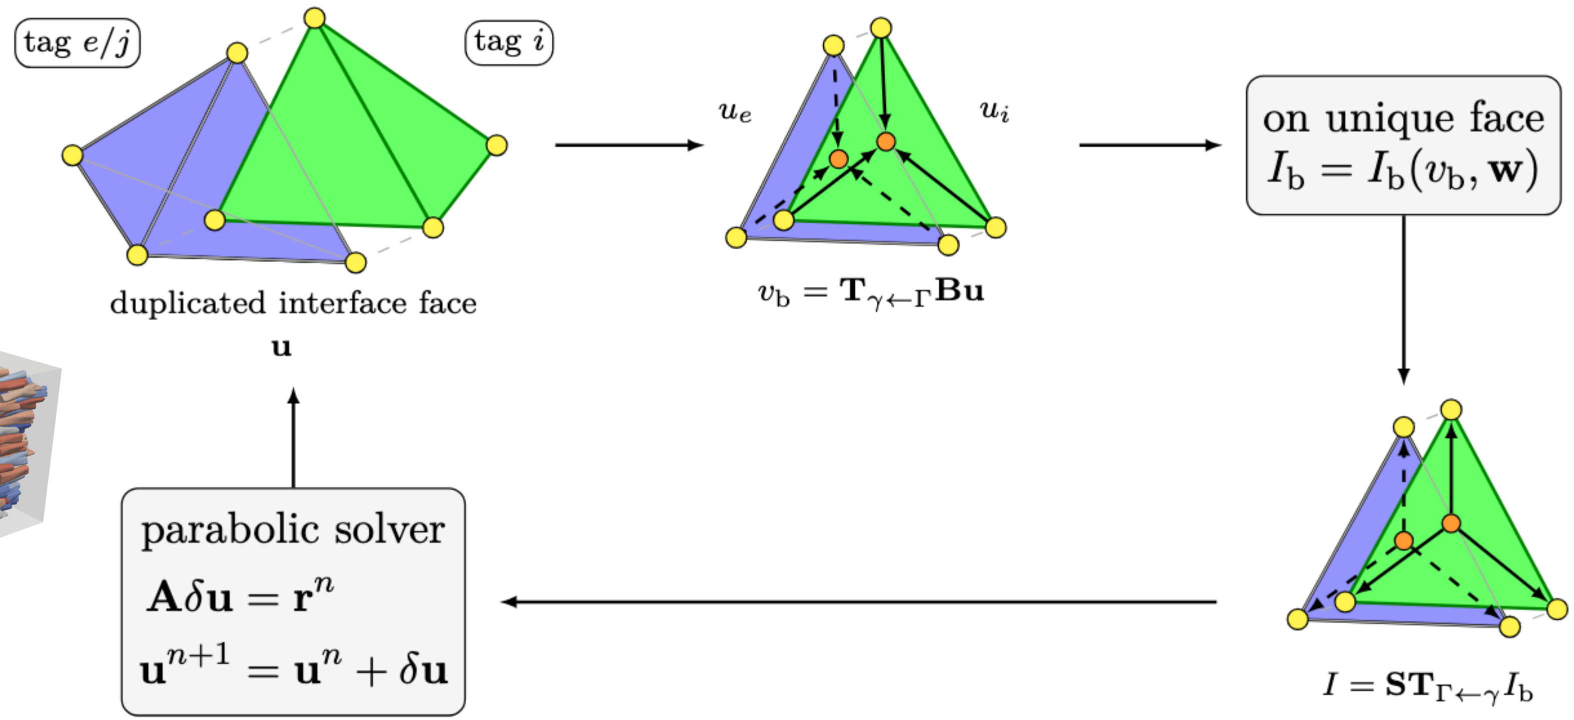

Semi-implicit method

$$\mathbf{A} = \mathbf{K} + \frac{C_\mathrm{m}}{\Delta t}\mathbf{M}_\Gamma$$

$$\mathbf{r}^n = -\left(\mathbf{S}\mathbf{T}_{\Gamma\leftarrow\gamma} I_b^n + \mathbf{K}\mathbf{u}^n\right) + \mathbf{M} I_\mathrm{vol}^n$$

EMI Simulation

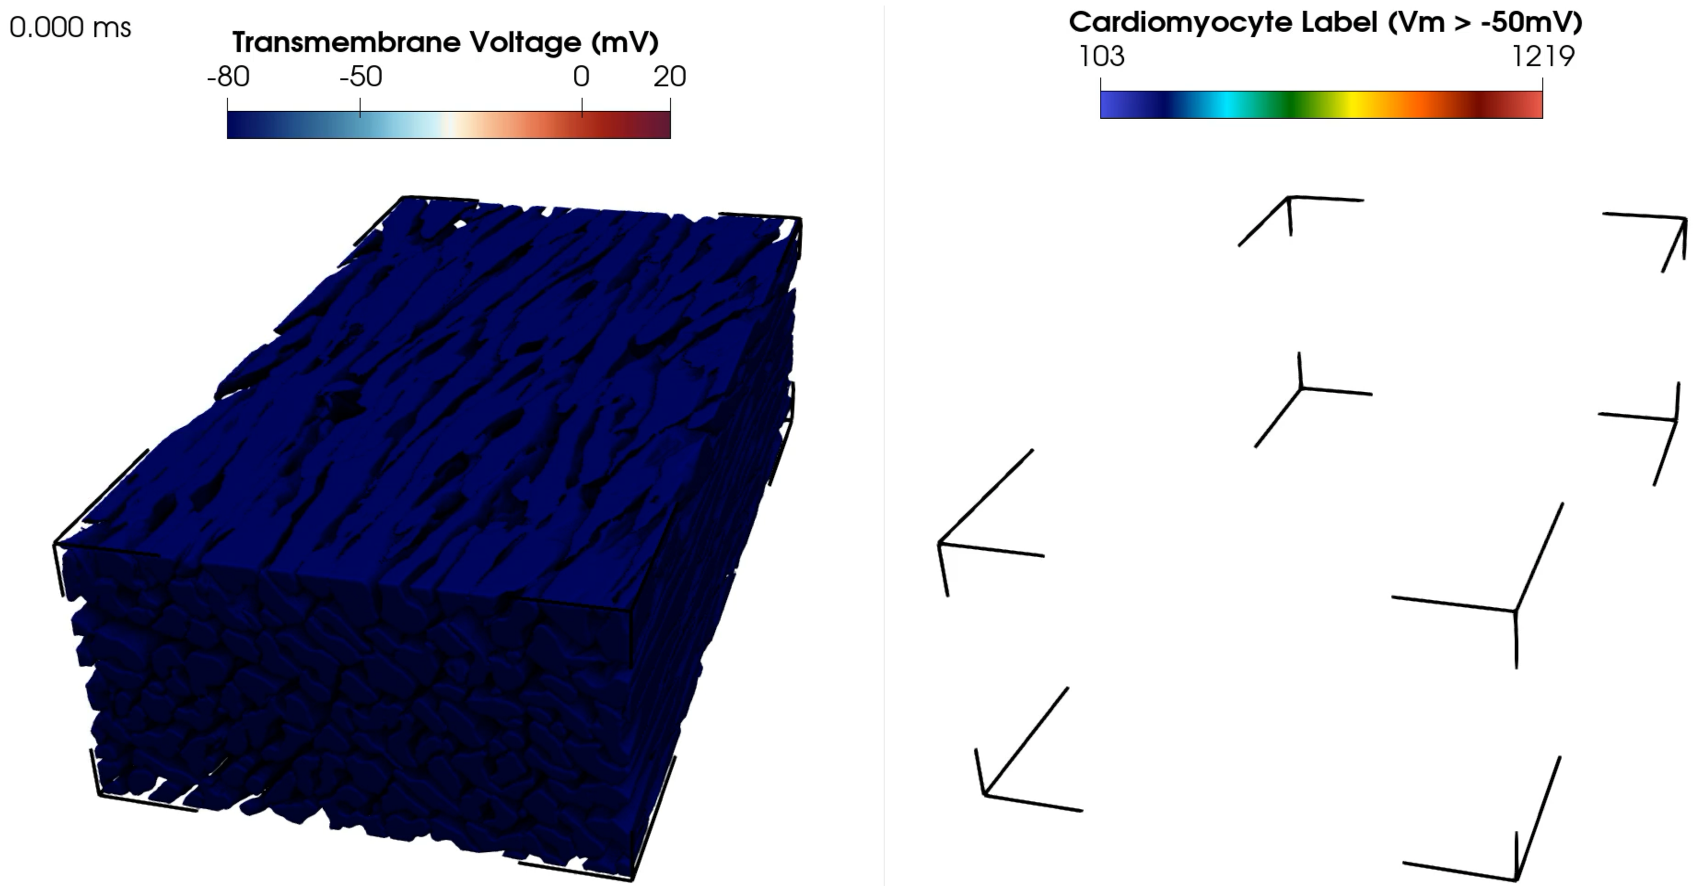

A deep learning-enabled toolkit for the 3D segmentation of ventricular cardiomyocytes

https://pubmed.ncbi.nlm.nih.gov/41031902/

202 x 280 x 113 µm3 With 556 myocytes

EMI Simulation

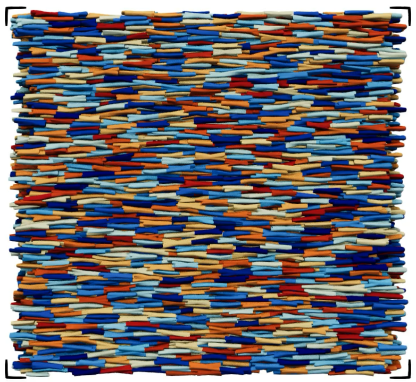

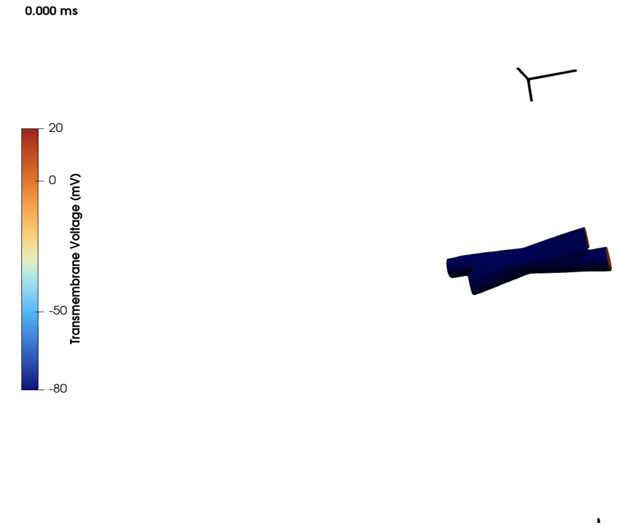

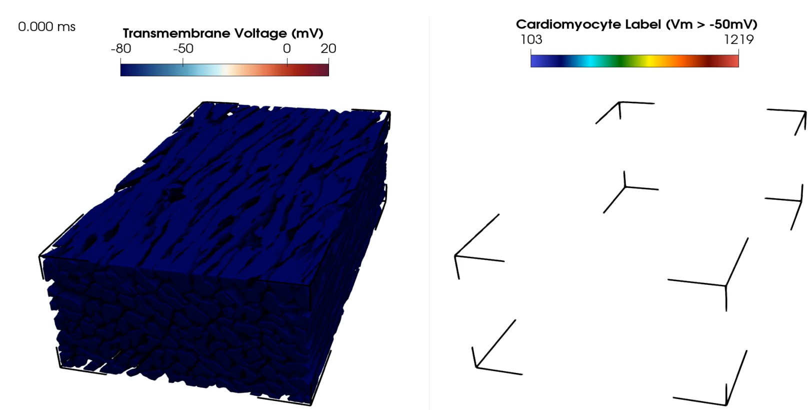

Weak-Scaling Benchmark with Ginkgo Backend

Lego Meshes

grid	#cells	#elem	#Γ	#Γ±	#DOFs
$4 \times 4 \times 4$	64	529 511	88 076	176 152	139 177
$8 \times 4 \times 4$	128	1 061 804	176 798	353 596	277 706
$8 \times 8 \times 4$	256	2 123 938	353 484	706 968	548 814
$8 \times 8 \times 8$	512	4 247 815	707 926	1 415 852	1 085 449

Weak Scaling

grid	#GPUs	avg. itr	EMI init (s)	EMI compute (s)	Ionics compute (s)	total (s)	η_p
$4 \times 4 \times 4$	1	203.9	4.41	19.96	0.13	29.58	1.00
$8 \times 4 \times 4$	2	299.9	5.12	31.58	0.14	42.19	0.70
$8 \times 8 \times 4$	4	414.3	9.75	50.99	0.22	66.70	0.44
$8 \times 8 \times 8$	8	439.6	11.71	60.77	0.23	80.30	0.37

communication overhead **additive-Schwarz** solve
→ stronger preconditioners on the larger runs

$$\eta_p = T_1/T_p$$

the weak-scaling efficiency
relative to the single-GPU baseline

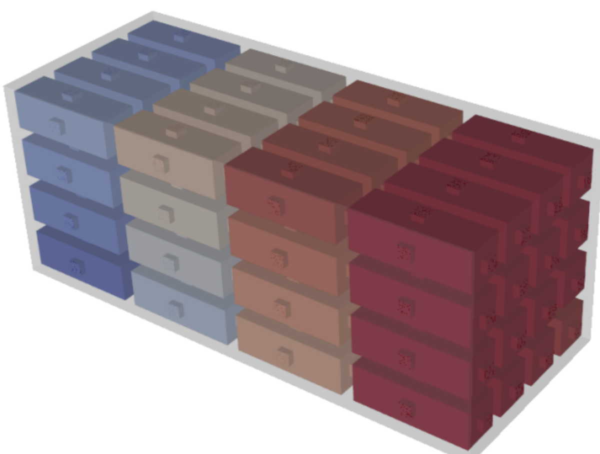

Alps GH200 GPUs,
one NVIDIA Grace Hopper GH200 GPU.

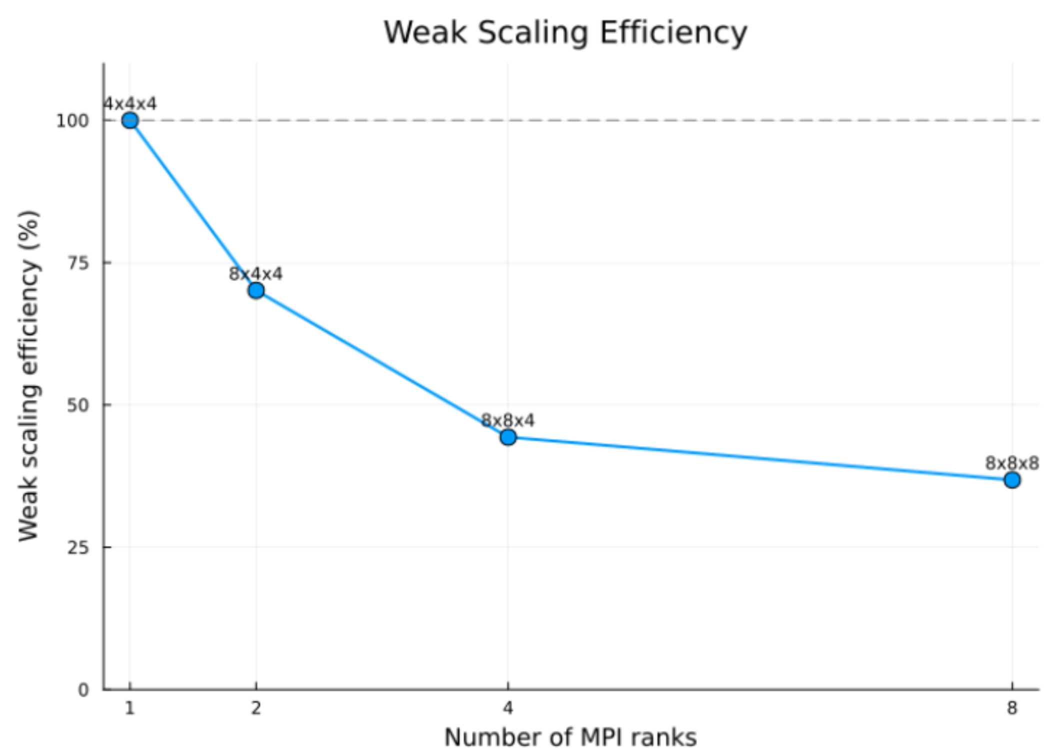

Strong Scaling using PETSc solver

mesh	#vertices	#dofs	#elems_volume	#elems_surface	#intra	#extra	domain size	dimension
$\mathcal{T}_2$	29m	37m	174m	30m	5641	8693	1 mm^3	10x10x10
$\mathcal{T}_3$	59m	76m	350m	66m	12411	16721	2 mm^3	20x10x10
$\mathcal{T}_4$	88m	114m	524m	101m	19063	24668	3 mm^3	30x10x10
$\mathcal{T}_5$	113m	146m	669m	132m	25023	32325	3 mm^4	20x20x10
$\mathcal{T}_6$	146m	190m	868m	169m	32143	40344	5 mm^3	50x10x10
$\mathcal{T}_7$	232m	302m	1.3b	272m	51578	62627	8 mm^3	40x20x10

mesh	#ranks	itr	EMI compute (s)	speedup	ionic compute(s)
$\mathcal{T}_2$	144	54.47	1669.74	1.00	0.33
$\mathcal{T}_2$	288	53.92	1481.32	1.13	0.10
$\mathcal{T}_2$	576	54.35	394.40	4.23	0.05
$\mathcal{T}_2$	1152	55.58	317.67	5.26	0.02
$\mathcal{T}_5$	576	55.02	2900.49	1.00	0.27
$\mathcal{T}_5$	1152	56.31	1882.73	1.54	0.76
$\mathcal{T}_5$	2304	56.20	792.06	3.66	0.34
$\mathcal{T}_5$	4608	56.72	578.45	5.01	0.03
$\mathcal{T}_7$	1152	64.99	4201.74	1.00	0.33
$\mathcal{T}_7$	2304	63.58	2000.55	2.10	0.12
$\mathcal{T}_7$	4608	61.55	1005.29	4.18	0.05
$\mathcal{T}_7$	9216	64.28	923.53	4.55	0.02

negligible

Impact of ionic models on the computation time

mesh	#vertices	#dofs	#elems_volume	#elems_surface	#intra	#extra	domain size	dimension
$\mathcal{T}_2$	29m	37m	174m	30m	5641	8693	$1\,mm^3$	10x10x10
$\mathcal{T}_3$	59m	76m	350m	66m	12411	16721	$2\,mm^3$	20x10x10
$\mathcal{T}_4$	88m	114m	524m	101m	19063	24668	$3\,mm^3$	30x10x10
$\mathcal{T}_5$	113m	146m	669m	132m	25023	32325	$3\,mm^4$	20x20x10
$\mathcal{T}_6$	146m	190m	868m	169m	32143	40344	$5\,mm^3$	50x10x10
$\mathcal{T}_7$	232m	302m	1.3b	272m	51578	62627	$8\,mm^3$	40x20x10

Ionic model	itr	EMI init (s)	EMI compute (s)	ionic init (s)	ionic compute (s)
AlievPanfilov	1430	38.47	587.45	0.11	0.06
Grandi	1086	39.06	459.99	0.08	19.40
LuoRudy91	923	39.31	384.08	0.09	0.28
OHara	1007	38.61	417.98	0.30	2.47
Tomek	876	39.27	368.34	0.24	3.96

Memory preallocation in PETSc

Mesh	Master branch	Preallocation branch	Reduction
$\mathcal{T}_1$	91 GB	18.53 GB	4.91×
$\mathcal{T}_2$	981 GB	179 GB	5.4×

Matrix	Used [GB]	Allocated [GB]	Overalloc [GB]
B	0.30	0.30	0.00
B_i	0.30	0.30	0.00
B_{sM}	0.15	0.15	0.00
unique_to_both	0.05	0.05	0.00
both_to_unique	0.05	0.05	0.00
stiffness_emi	0.68	0.68	0.00
mass_emi	0.68	0.68	0.00
mass_surf_emi	0.36	0.36	0.00
lhs_emi	0.86	0.86	0.00
Total	3.45	3.45	0.00

2nd order Hierarchical FEM

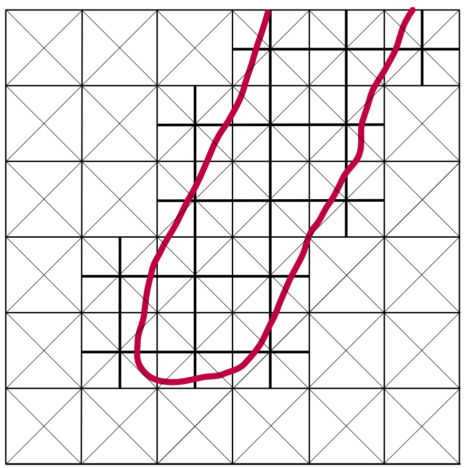

Mesh refinement

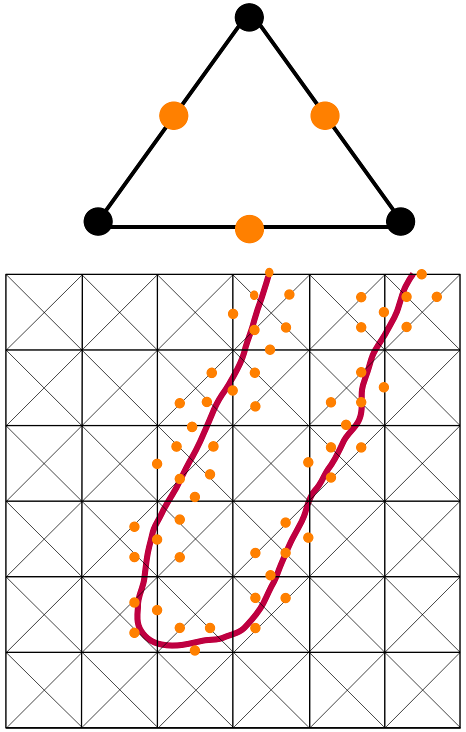

Adaptive 2nd order FE

Adaptive resolution: Hierarchical FEs allow selective addition of nodes or basis functions in specific regions needing higher accuracy.
Multi-resolution approach: Complexity can be increased locally without refining the entire mesh.
Efficiency: Model complexity is enhanced without significantly increasing the number of elements or memory usage.

2nd order Hierarchical FEM

$$\mathbf{A}^{(2)} = \mathbf{K}^{(2)} + \frac{C_m}{\Delta t}\mathbf{M}_\Gamma^{(2)}.$$

$$\mathbf{r}^{n,(2)} = -\left(\mathbf{S}^{(2)}\mathbf{T}_{\Gamma\leftarrow\gamma}\mathbf{I}_b^{n,(2)} + \mathbf{K}^{(2)}\mathbf{u}^{n,(2)}\right) + \mathbf{M}^{(2)}\mathbf{I}_{\mathrm{vol}}^{n,(2)}.$$

$$\downarrow$$

$$\mathbf{I}_b^{n,(2)} = \mathbf{I}_{\mathrm{ion}}(\mathbf{v}_\gamma^{n,(2)}, \mathbf{w}^n) - \frac{1}{2}\mathbf{T}_{\gamma\leftarrow\Gamma}\mathbf{B}_{\mathrm{i}}^{(2)}\mathbf{I}_{\mathrm{nodal}}^{n,(2)}.$$

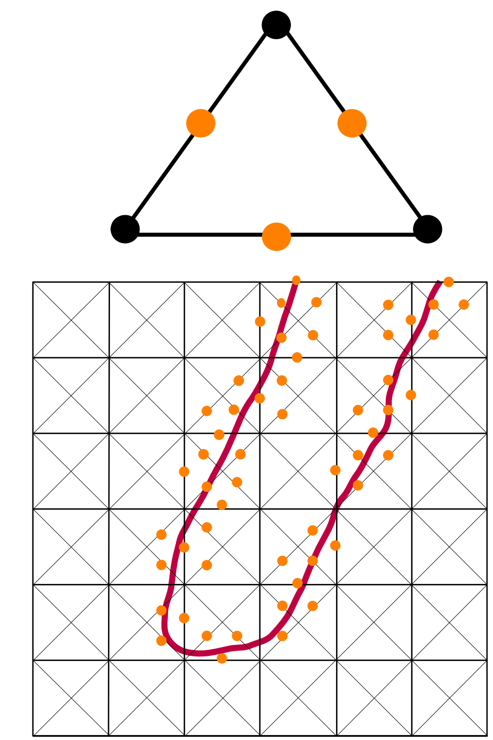

Adaptive 2nd order FE

extract the P1-compatible part of the hierarchical solution

$$\mathbf{u}_{\mathrm{out}}^{(1)} = \mathbf{R}_{1\leftarrow 2}\mathbf{u}^{(2)}.$$

Two level preconditioner

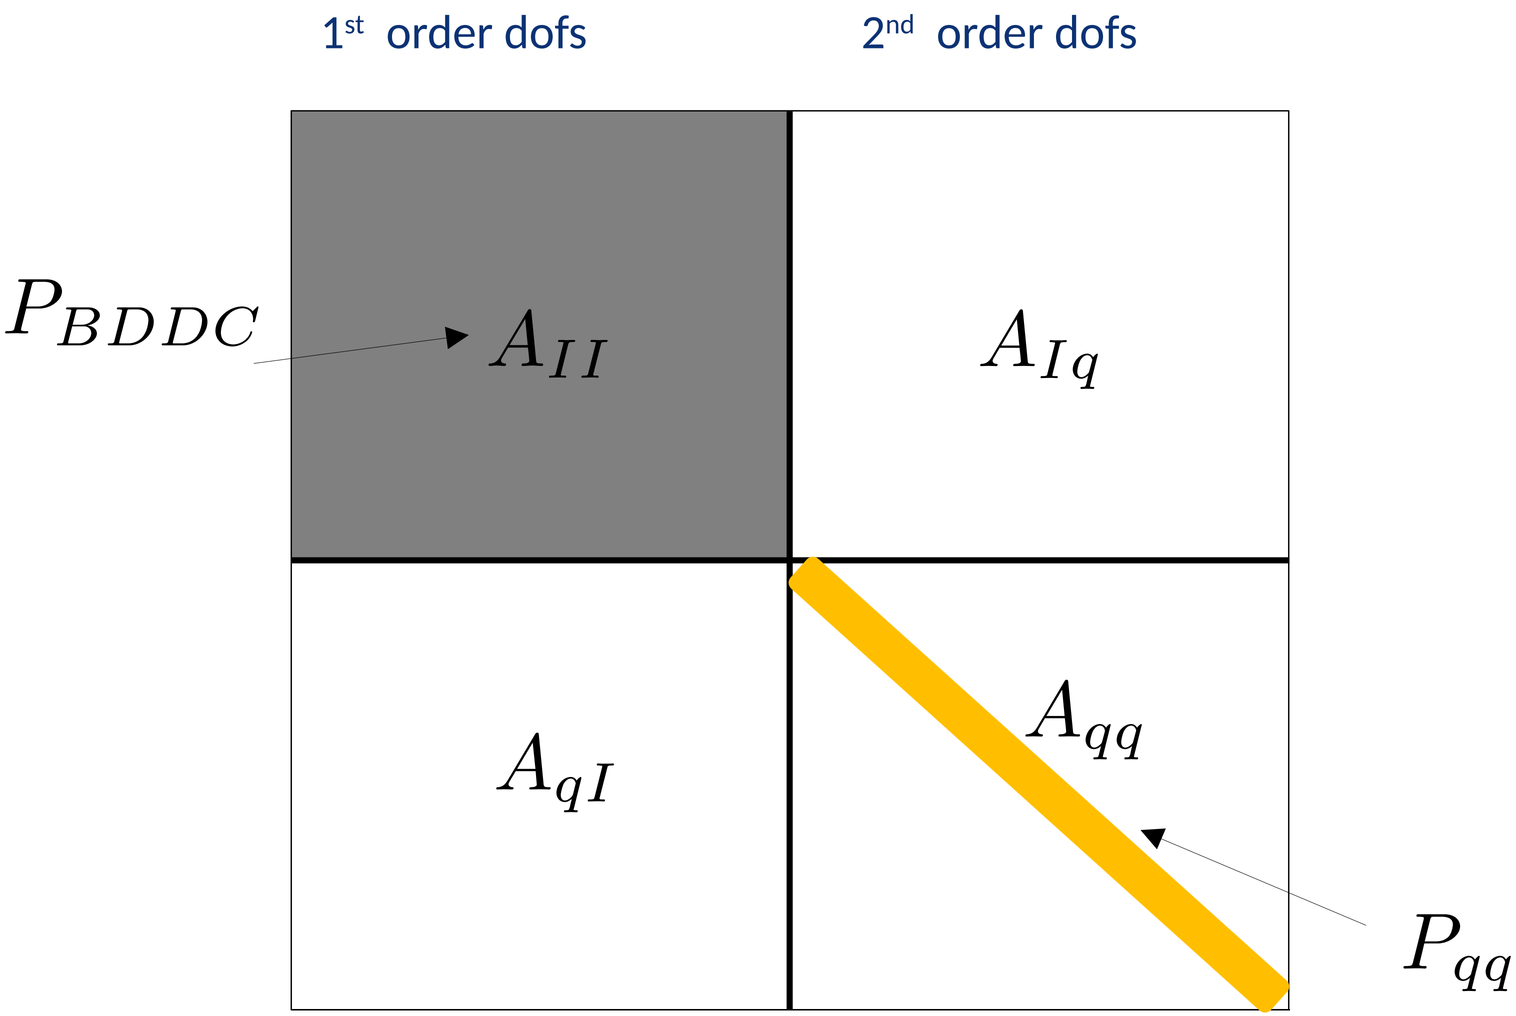

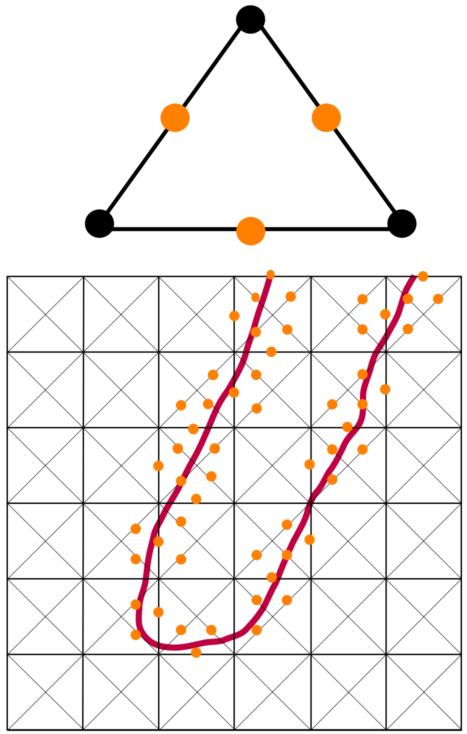

Adaptive 2nd order FE

EMI in hierarchical FEM

accuracy

level	volume elements	membrane faces	P1 DOFs	P2 DOFs	P1 error	P2 error	P1/P2 error ratio
base	2 906	784	609	5 983	$1.5548 \cdot 10^{-2}$	$5.4106 \cdot 10^{-3}$	2.87
refine1	23 248	3 136	4 347	39 257	$9.6461 \cdot 10^{-3}$	$1.9314 \cdot 10^{-3}$	4.99
refine2	185 984	12 544	32 837	280 637	$4.9143 \cdot 10^{-3}$	$4.8963 \cdot 10^{-4}$	10.04
refine3	1 487 872	50 176	255 241	2 113 669	$2.0678 \cdot 10^{-3}$	reference	–

64x less elements

EMI in hierarchical FEM

accuracy

level	volume elements	membrane faces	P1 DOFs	P2 DOFs	P1 error	P2 error	P1/P2 error ratio
base	2 906	784	609	5 983	$1.5548 \cdot 10^{-2}$	$5.4106 \cdot 10^{-3}$	2.87
refine1	23 248	3 136	4 347	39 257	$9.6461 \cdot 10^{-3}$	$1.9314 \cdot 10^{-3}$	4.99
refine2	185 984	12 544	32 837	280 637	$4.9143 \cdot 10^{-3}$	$4.8963 \cdot 10^{-4}$	10.04
refine3	1 487 872	50 176	255 241	2 113 669	$2.0678 \cdot 10^{-3}$	reference	–

64x less elements

Efficiency with two-level preconditioner

level	method	preconditioner	DOFs	avg. itr	total itr	solver time (s)	time/itr (s)
base	P1	GAMG on P1 matrix	609	39.40	401	0.06	0.000155
base	P2	single-block GAMG	5 983	81.60	817	0.15	0.000184
base	P2	split: GAMG + Jacobi	5 983	117.00	1 165	0.15	0.000127
refine1	P1	GAMG on P1 matrix	4 347	62.00	631	0.16	0.000248
refine1	P2	single-block GAMG	39 257	156.00	1 558	1.14	0.000735
refine1	P2	split: GAMG + Jacobi	39 257	190.00	1 897	0.59	0.000313
refine2	P1	GAMG on P1 matrix	32 837	99.60	1 016	0.54	0.000533
refine2	P2	single-block GAMG	280 637	309.00	3 106	14.62	0.004706
refine2	P2	split: GAMG + Jacobi	280 637	368.00	3 671	6.05	0.001649
refine3	P1	GAMG on P1 matrix	255 241	153.40	1 579	4.70	0.002979
refine3	P2	single-block GAMG	2 113 669	619.60	6 213	245.14	0.039456
refine3	P2	split: GAMG + Jacobi	2 113 669	751.80	7 491	95.33	0.012726

4.70/0.59 = 8x faster

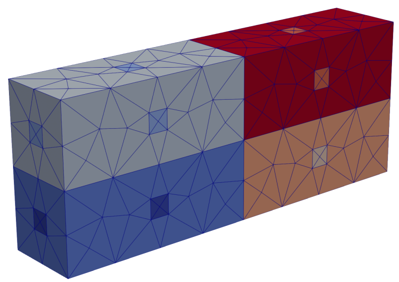

EMI in hierarchical FEM

accuracy

level	volume elements	membrane faces	P1 DOFs	P2 DOFs	P1 error	P2 error	P1/P2 error ratio
base	2 906	784	609	5 983	$1.5548 \cdot 10^{-2}$	$5.4106 \cdot 10^{-3}$	2.87
refine1	23 248	3 136	4 347	39 257	$9.6461 \cdot 10^{-3}$	$1.9314 \cdot 10^{-3}$	4.99
refine2	185 984	12 544	32 837	280 637	$4.9143 \cdot 10^{-3}$	$4.8963 \cdot 10^{-4}$	10.04
refine3	1 487 872	50 176	255 241	2 113 669	$2.0678 \cdot 10^{-3}$	reference	–

64x less elements

Efficiency with two-level preconditioner

level	method	preconditioner	DOFs	avg. itr	total itr	solver time (s)	time/itr (s)
base	P1	GAMG on P1 matrix	609	39.40	401	0.06	0.000155
base	P2	single-block GAMG	5 983	81.60	817	0.15	0.000184
base	P2	split: GAMG + Jacobi	5 983	117.00	1 165	0.15	0.000127
refine1	P1	GAMG on P1 matrix	4 347	62.00	631	0.16	0.000248
refine1	P2	single-block GAMG	39 257	156.00	1 558	1.14	0.000735
refine1	P2	split: GAMG + Jacobi	39 257	190.00	1 897	0.59	0.000313
refine2	P1	GAMG on P1 matrix	32 837	99.60	1 016	0.54	0.000533
refine2	P2	single-block GAMG	280 637	309.00	3 106	14.62	0.004706
refine2	P2	split: GAMG + Jacobi	280 637	368.00	3 671	6.05	0.001649
refine3	P1	GAMG on P1 matrix	255 241	153.40	1 579	4.70	0.002979
refine3	P2	single-block GAMG	2 113 669	619.60	6 213	245.14	0.039456
refine3	P2	split: GAMG + Jacobi	2 113 669	751.80	7 491	95.33	0.012726

4.70/0.59 = 8x faster

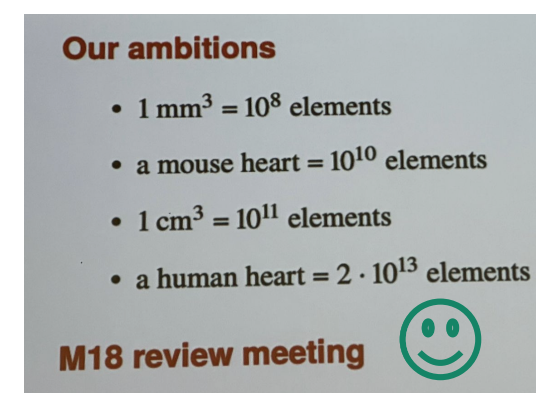

Improve Accuracy & Efficiency

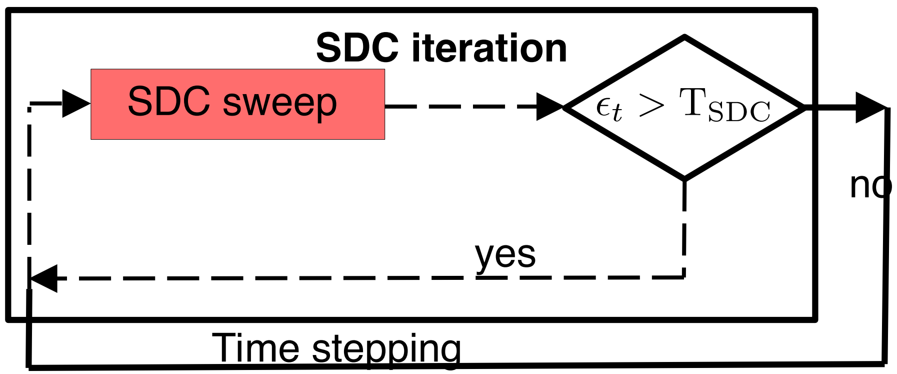

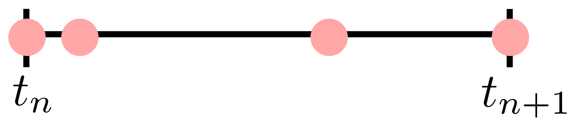

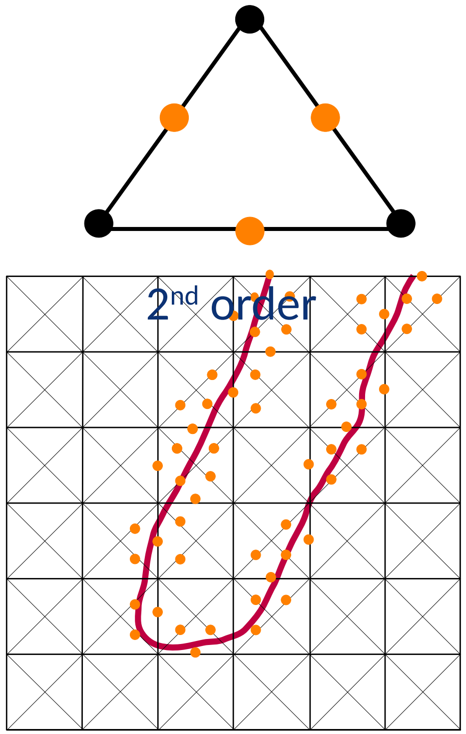

Adaptive 2nd order FE

Accuracy and Efficiency achievement

- Hierarchical quadratic FE → coarser mesh **8x faster**
- Larger time step in SDC
- Algebraic adaptivity
- Two level block preconditioner

Coming soon ☺

Conclusion

(3 years ago in Strasbourg)

Tagged volume mesh → Distributed region → **Surface extraction** → Interface decoupling → Matrices/operator

- Consistent distributed-memory EMI pipeline:
  - interface reconstruction, mesh decoupling, operator assembly, membrane-model evaluation, and large-scale linear solves
- Face-centered piecewise-constant membrane evaluation avoids ambiguities at complex cell junctions.
- Exact PETSc preallocation substantially reduces memory usage for large EMI meshes.
- Demonstrated scalability in weak- and strong-scaling studies.
- Hierarchical P2 FEM improves accuracy compared with P1 on the same mesh levels.
  - Higher accuracy on the coarser mesh + more efficient
  - Later SDC + AA + two level preconditioner → coming soon

Thanks for your attention!

design discussion

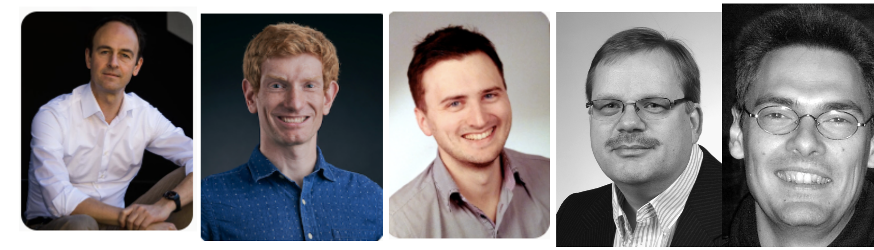

electrophysiological perspective

solver/preconditioner & scalability

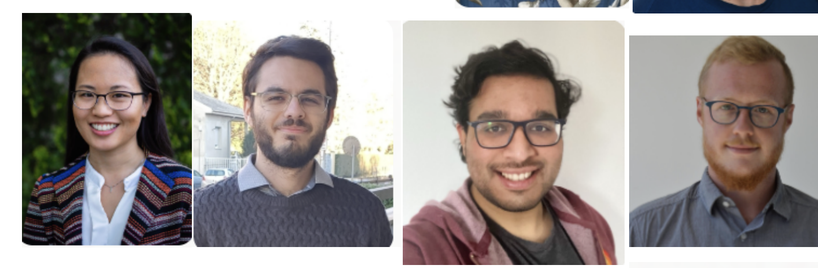

benchmark

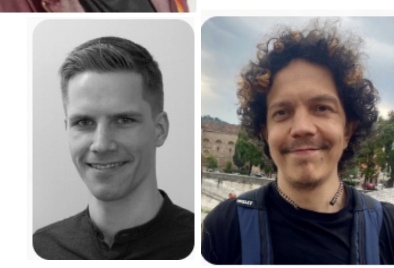

Microscopic simulation

partitioning

With funding from the:

Federal Ministry of Research, Technology and Space

EuroHPC Joint Undertaking

Co-funded by the European Union

This work was supported by the MICROCARD-2 project (project ID 101172576). The project is supported by the EuroHPC Joint Undertaking and its members (including top-up funding by ANR, BMBF, and Ministero dello sviluppo economico). Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or EuroHPC. Neither the European Union nor EuroHPC can be held responsible for them.