Joint Usage/Research Center for Interdisciplinary Large-scale Information Infrastructures
Final Report for JHPCN Joint Research of FY 2025

jh250028

High resolution simulation of cardiac electrophysiology on realistic whole-heart geometries

Kengo Nakajima（The University of Tokyo, Japan）

Abstract
A recent trend in high-fidelity modeling of cardiac electrophysiology is the adoption of the so-called EMI model, which resolves the intra- and extra-cellular spaces of the cells, with the cell membranes being the interface between the two. In comparison with the conventional "homogenized" models, this new mathematical model thus provides a more accurate description of the cardiac electrical activities that are of vital importance for the human body. This project aims to implement and optimize a new simulator based on the EMI model, with the purpose of effectively using modern supercomputers. By combining efficient numerical algorithms with hardware-compatible software implementations, we want to help computational scientists carry out novel "in-silico" experiments. The research is done as a collaboration between The University of Tokyo and Simula Research Laboratory (Norway) to combine and further strengthen the expertise areas of the two partners, namely, high-performance computing and cardiac modeling. During FY2025, this international JHPCN project carried out research on the following topics: (1) Improving the communication scalability in connection with parallelizing the new cell-resolved simulator of cardiac electrophysiology; (2) Studying the numerical stability of the new simulator; (3) GPU-porting of the important components of the new simulator; (4) Performance modeling of Krylov-subspace iterative solvers; (5) Carrying out large-scale cell-resolved simulations.

1. Basic Information

(1) Collaborating JHPCN Centers

Tohoku University

The University of Tokyo

Nagoya University

(2) Theme Area

Large scale computational science area

(3) Project Members and Their Roles

Kengo Nakajima (U Tokyo): Administration, numerical algorithms (alg.)

Xing Cai (Simula/Norway): Administration, numerical alg.

Akihiro Ida (JAMSTEC): Numerical alg.

Toshihiro Hanawa , **Yohei Miki** , **Kazuya Yamazaki** (U Tokyo): Optimization

Akira Naruse (NVIDIA): Optimization

Masatoshi Kawai (Tohoku U): Numerical alg.

Tetsuya Hoshino (Nagoya U): Optimization

Masaharu Matsumoto (Fukushima Univ): Numerical alg.

Akihiro Fujii, Kei Teshigawara, Taketeru Kondo (Kogakuin Univ): Numerical alg.

Johannes Langguth (Simula/Norway): Optimization

Hermenegild Arevalo (Simula/Norway): Cardiac electrophysiology, mathematical modeling, simulations

James D. Trotter (Simula/Norway): Mathematical modeling, optimization

2. Purpose and Significance of the Research

The overall purpose of this joint research project is to enable cell-resolved mathematical modeling of cardiac electrophysiology, with the help of high-performance computing (HPC). Compared with the conventional monodomain & bidomain models, which are based on "homogenization", the novelty of cell-resolved modeling lies in the possibility of studying subcellular heterogeneity and non-uniform details along the cell membranes, due in particular to the newly adopted EMI model that resolves the interior

and exterior of the cardiac cells, as well as the membranes. Such cell-resolved simulations require much higher resolutions than the traditional simulation approach. The presence of different subdomains (interior, exterior & membrane) also requires special work/mesh partitioning before parallel computing can be carried out. Other performance-critical topics such as overlapping computation with communication and proper use of GPUs also require more investigations in this new context. This project thus continued the research efforts of jh230018 & jh240030. A successful outcome of this JHPCN joint project will be unprecedented, cell-resolved simulations, which can lead to new knowledge in the domains of both computational physiology and HPC.

3. Significance as JHPCN Joint Research Project

The proposed research spans computational physiology, numerical analysis and HPC. The necessity of implementing this JHPCN joint research project is due to two aspects. First, The University of Tokyo has world-leading expertise in implementing and optimizing advanced numerical code for running on cutting-edge supercomputers. Such hands-on experience on supercomputing is lacking for the Norwegian partner, Simula Research Laboratory, whose main expertise is on cardiac modeling and simulation. Second, the Wisteria/BDEC-01 and Miyabi systems are two future-oriented hardware platforms. They represent large-scale clusters of non-x86 processors and tightly coupled CPU-GPU superchips. Efficient use of such processors is thus important for preparing the cell-resolved simulator for the latest hardware architectures.

The joint effort between The University of Tokyo and Simula for developing an initial, preliminary implementation of a new simulator based on the EMI model was started in FY2023 and continued in FY2024. During FY2025, the current project enhanced the research activities, by improving the parallel performance and scalability of the cell-resolved simulator of the EMI model, as well as working further on a partial GPU porting of the EMI simulator. Therefore, the work done in FY2025 continued to pave the way for the new cell-resolved simulator to be used for realistic cardiac simulations (which will require using very large-scale GPU clusters) in the future.

4. Outline of Research Achievements until FY2024

With respect to a complete simulator that is based on the cell-resolved EMI model, there only existed a very preliminary implementation until FY2024. Specifically, the preliminary simulator only had a sequential version and a shared-memory version based on OpenMP threads for CPU execution.

The main research effort prior to FY2025 focused on developing several components needed for an MPI-based parallel version of the EMI simulator that aims at distributed-memory supercomputers. The developed components prior to FY2025 include (1) a special mesh partitioning strategy that targets the separated intra- and extra-cellular spaces; (2) use of the “sector cache” technique for optimizing sparse matrix-vector multiplications on the A64FX processor; (3) an elaborate communication-computation overlap strategy; (4) the use of lookup tables related to concurrently solving many systems of ordinary differential equations for CPU execution; (5) testing of various reordering schemes for the

degrees of freedom involved in finite element discretization; (6) development of an alternating subdomain strategy to separate the updates of the cell interiors from the exterior; and (7) optimization of the assembly procedure connected with finite element discretization for both CPU and GPU execution.

Based on the preliminary implementations of a sequential and a shared-memory version of the EMI simulator, as well as the various components mentioned above to prepare a full MPI-based parallelization, the research work in FY2025 (see below) resulted in an MPI-based parallel simulator for execution on CPU clusters, as well as a big step toward full GPU porting.

5. Details of FY2025 Research Achievements

(1) Overview

Before presenting the details of the research achievements during FY2025, we will quickly recap the cell-resolved simulation approach based on the EMI model. Unlike the traditional simulation approach that lets the intracellular & extracellular spaces and the cell membranes unphysiologically "coexist" everywhere, the new simulation approach fully resolves the individual cardiac cells (termed the *intracellular subdomains*), the surrounding area in-between the individual cells (termed the *extracellular subdomain*), and the cell membranes that are the interfaces between the intracellular and extracellular subdomains. Inside each intracellular subdomain (i.e., each cardiac cell), an individual Laplace equation applies, whereas in the entire extracellular subdomain a global Laplace equation applies. The coupling between the intracellular and extracellular subdomains is through boundary conditions applicable on the cell membranes, whose dynamics are described by a set of nonlinear ordinary differential equations. The "EMI" mathematical model is as follows:

$$
\begin{aligned}
\nabla \cdot \sigma_i \nabla u_i^k &= 0 && \text{in } \Omega_i^k,\\
\nabla \cdot \sigma_e \nabla u_e &= 0 && \text{in } \Omega_e,\\
C_m \frac{\partial v^k}{\partial t} + I_{\text{ion}}^k(v^k, s^k) &= n_e \cdot \sigma_e \nabla u_e = -n_i^k \cdot \sigma_i \nabla u_i^k && \text{at } \Gamma_k,\\
v^k &= u_i^k - u_e && \text{at } \Gamma_k,\\
\frac{\partial s^k}{\partial t} &= F_k(s^k, v^k) && \text{at } \Gamma_k,\\
C_g \frac{\partial w^k}{\partial t} + I_{\text{gap}}^{k,j}(w^k) &= n_i^j \cdot \sigma_i \nabla u_i^j = -n_i^k \cdot \sigma_i \nabla u_i^k && \text{at } \Gamma_{k,j},\\
w^k &= u_i^k - u_i^j && \text{at } \Gamma_{k,j},
\end{aligned}
$$

The overall numerical strategy that is adopted by us is to "decouple" the intracellular subdomains from the extracellular subdomain (operator-splitting) at each time step. That is, solution of the extracellular Laplace equation alternates with the solution of the individual intracellular Laplace equations. The connection between the two subdomain types is through updated boundary conditions.

The research results obtained in FY2025 can be summarized for the following topics:

(2) Communication scalability

We further improved the special mesh partitioner that fits with the adopted operator-splitting numerical strategy, where the intracellular subdomains and the extracellular subdomain are updated alternatingly during each time step. Specifically, we now first partition the extracellular subdomain, and then incorporate the partitioned extracellular pieces as "super-nodes" into another graph that is used to partition the intracellular subdomains.

Communication volume; Old partitioning strategy (Round robin)

	1	2	3	4	5	6	7	8
1		177.0	180.6	75.1				
2	119.8		130.3	157.3	80.9			
3	70.4			52.0	94.3	88.7	143.3	182.8

4					137.2	222.3	168.4	59.0
5				69.6		119.0	152.2	168.3
6					121.8		160.5	148.0
7	173.8	165.9	146.0	70.4				69.0
8	66.9	147.2	166.9	171.5	69.4			
SUM	**430.9**	**490.1**	**623.8**	**595.9**	**503.6**	**430.0**	**624.5**	**627.1**

New partitioning strategy (multi-constraint graph partitioning)

	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
1							141.4	
2			263.0	335.7	30.0			
3							186.6	443.7
4		219.0	87.9			31.6	6.9	
5			39.3			28.2		
6		9.4					207.7	202.9
7			214.5	4.4	39.3	358.0		
8	143.8	395.9					82.3	
SUM	**143.8**	**624.3**	**604.8**	**340.1**	**69.2**	**417.8**	**625.0**	**646.6**

As an example, we consider a setup with 256 cardiac cells and 8 MPI processes, such that the extracellular subdomain is partitioned into 8 pieces. The old partitioning strategy (from FY2024) assigns each extracellular piece to an MPI process and then distributes the 256 cells to the MPI processes in a *round robin* fashion. The improved strategy of FY2025 performs a multi-constraint, weighted graph partitioning based on a graph of the subdomains and interfaces, ensuring that each MPI process receives a balanced number of extracellular and intracellular elements. The two tables above show the volumes of point-to-point communication among 8 MPI processes resulting from the two strategies. This type of communication is related to updating electric potentials on the cell membranes, which is needed after every extracellular subdomain update, before updating the individual intracellular subdomains. The total communication volume is reduced from 4325.9 KiB to 3471.5 KiB and the total number of messages is also reduced since each process has fewer neighbors.

(3) Numerical scalability

We looked at two important aspects that both impact the numerical scalability. The first aspect is the number of iterations needed to achieve convergence between the intra- and extracellular subdomains per time step. This, together with the total number of linear solver iterations and time usage, determines if the operator-splitting numerical strategy is better than a monolithic strategy that simultaneously updates the extra- and intracellular subdomains.

The operator-splitting and monolithic strategies are compared against each other for a benchmark problem having a single cardiac cell. For the monolithic strategy, we've used an EMI simulator from OpenCARP that is in development. A direct solver (CHOLMOD) is found to provide the best performance in this case. For the splitting strategy, we've used 2 iterations between the intra- and extracellular subdomains per time step. Linear systems are solved by the conjugate gradient (CG) method. The two methods produce visually identical results. The serial performance of both solvers is shown in the table below:

	Splitting (Our solver)	Monolithic (OpenCARP)
ODE part	15 s	160 s
PDE part	174 s	211 s
Total	189 s	371 s

The second aspect particularly investigates the CG method when preconditioned by algebraic multigrid (AMG). This is very important for the extracellular subdomain in large-scale simulations. Despite NVIDIA AmgX's performance-oriented design, it failed to converge. During FY2026, we

aim to integrate non-smoothed aggregation into our solver and apply OpenACC for enhanced GPU acceleration. We can mention that, from the middle of FY.2025, Professor Fujii's team from Kogakuin University joined the project. The AMGS library (https://hpcl.info.kogakuin.ac.jp/docs/amgs/), which Prof. Fujii has been developing since around 2000, has a strong track record in large-scale parallel simulations using both CPUs and GPUs (doi: 10.1109/ACCESS.2024.3386226). The use of AMGS will be an important topic in FY.2026.

(4) Partial GPU porting [3,5]

In addition to experimenting with existing GPU implementations of the AMG preconditioner (see above), we also studied the effect of GPU porting of two other parts of our simulator. The first part is the ODE solver, which is needed to simulate the membrane dynamics. We've used lookup tables to store pre-evaluated outputs of expensive mathematical functions, for the purpose of saving time. The second part is the calculation of right-hand side vectors, which need to be updated for each intra- or extracellular iteration. Instead of a standard element-by-element assembly procedure, we've adopted an equivalent matrix-vector multiplication. The first table below shows the runtime (in seconds) of the ODE solver on CPU (using 128 OpenMP threads on a dual-socket AMD Epyc 7763) and GPU (NVIDIA A100) with and without lookup tables on a problem with 741,353 systems of ODEs.

ODE solver	CPU [s]	GPU [s]
No lookup tables	66.9	7.1
Lookup tables	21.1	4.1

The next table below shows the runtime of right-hand side calculation on a domain with 4,434,789 tetrahedral elements. Replacing the elementwise strategy, which was originally used on the CPU, with a method based on matrix-vector multiplication results in a speedup of about 14x. This strategy also greatly eases the porting to the GPU, where it provides a further speedup of 4.5x.

Right-hand side assembly	CPU [s]	GPU [s]
Elementwise	121.0	N/A
Matrix-vector multiplication	8.5	1.9

As another subject that is relevant for porting the simulator to GPUs, we also studied the effect of using alternative multi-GPU communication libraries other than the standard MPI library. Specifically, we examined the NCCL, RCCL and NVSHMEM communication libraries, in the context of implementing the CG iterative solver for execution on multi-GPU platforms. The following figure reports the speedup results of new multi-GPU implementations of CG (termed *aCG*) using new communication libraries, over single-GPU performance. Six sparse matrices from SuiteSparse were used, and the two tested hardware platforms were Wisteria-Aquarius (NVIDIA A100 GPUs) and LUMI (AMD MI250x GPUs). It is observed that the NCCL library, in particular, has the potential of outperforming the standard MPI counterpart.

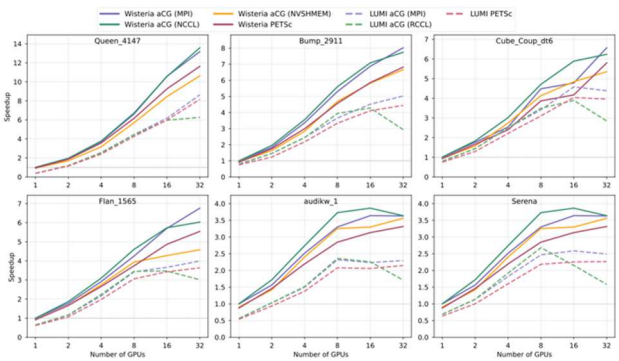

(5) Efficient utilization of CPU/GPU

It is also important to take advantage of the strong connectivity with the GPU via NVLink-C2C and effectively utilize the CPU portion of the GH200 Super Chip installed in Miyabi-G. In multigrid

methods, the problem size becomes small at coarse grid levels, making them unsuitable for GPUs. Therefore, we are investigating approaches that effectively utilize CPUs for solving at the coarse grid levels. Traditionally, the data transfer overhead between CPUs and GPUs has been a serious issue, but NVLink C2C now allows CPUs to directly access data stored in GPU memory [16]. Preliminary studies on geometric multigrid solvers in [1,4] indicate that, for small-scale problems, a performance improvement of approximately 20% can be expected. This will be one of the main topics in FY.2026.

(6) Performance modeling of the CG solver [2,3]

We developed performance models for two types of inter-GPU communication overhead that arise in any multi-GPU implementation of CG. The first type of overhead is associated with the peer-to-peer communication needed before each GPU can execute its designated part of a parallel sparse matrix-vector multiplication (SpMV). The resulting performance model can handle the general situation where the GPUs may have varying numbers of neighbors for communication (and with different message sizes) and that there may be contention for the communication bandwidth available on a heterogeneous interconnect. The second type of overhead is associated with MPI_Allreduce (or similar) needed for summing up each GPU's contribution to a parallel dot-product. The resulting performance model is in the form of tabulated costs of MPI_Allreduce (or similar) corresponding to different numbers of GPUs used. These costs are measured per target GPU cluster, by one of the microbenchmarks from the OSU suite. Besides the two cost models of communication overhead, we also developed performance models for the three computational kernels involved in each CG iteration: vector addition, dot-product and SpMV. These models were used to validate the single-GPU performance with respect to the theoretical expectation.

Following up earlier work on performance modeling and optimization of SpMV on Wisteria's A64FX CPUs, we confirmed that CPU cache partitioning improves performance by 10-20% in the SpMV kernel, a critical part of the CG solver.

(7) Large-scale tests

A main result of FY2025 was the completion of an MPI-based full parallelization of the EMI simulator for execution on a cluster of multicore CPUs. This parallel simulator represents a major step forward, compared with the sequential and shared-memory versions of the EMI simulator prior to FY2025.

To ensure the correctness of the new MPI parallelization of the complete EMI simulator (termed *DEMI* below), we compared the simulation results from DEMI against an independent shared-memory implementation for a benchmark of a 13x65 configuration of cells. The following figure confirms that the two parallel versions of the EMI simulator produced qualitatively same numerical solutions. (The actual differences are due to slightly different discretization techniques adopted in space and time.)

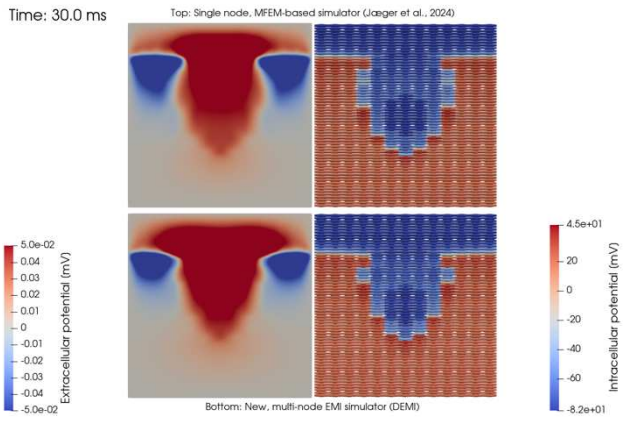

To study the parallel scalability of DEMI, we also carried out another benchmark of a 52x130

configuration of cells. This problem size is the largest according to the current literature on EMI simulations. The following table reports the total time usage of DEMI as different numbers of nodes on Wisteria-Odyssey were used. For each Odyssey node, four MPI processes were always used, each spawning 12 OpenMP threads.

# nodes	# MPI procs	# CPU cores	Total time
4	16	192	109687 s
8	32	384	57526 s
16	64	768	33654 s
32	128	1536	20342 s

6. Self-review of Current Progress and Future Prospects

The research activities in FY2025 have been carried out according to the original plan. However, we see the need of scalable parallel implementations of the algebraic multigrid (AMG) preconditioner, which is desirable for both multi-CPU and multi-GPU executions of the CG iterative solver. Although the AmgX library from NVIDIA failed to provide convergence for the CG tests that are relevant for our EMI simulations, we envision that the adoption of non-smoothed aggregation in the building-up phase of an AMG preconditioner can be helpful. This optimism is based on the positive experience with such techniques that have been reported in the literature. Contributions from the AMGS library by Prof. Fujii are expected in FY.2026 on this issue.

Finally, the achievement rate for the FY.2025 is 85%.

#Section 7, achievements of this FY, should be input on the JHPCN website.

https://jhpcn-kyoten.itc.u-tokyo.ac.jp/abstract/jh250028