

Benchmarking Open Cardiac Electrophysiology Simulators: MonoAlg3D and OpenCARP

Lucas M R de Lima¹, Rafael S Oliveira², Fernando O Campos³, Lucas A Berg⁴, Joventino O Campos¹, Rodrigo W dos Santos¹

¹ Universidade Federal de Juiz de Fora, Juiz de Fora, Brazil

² Universidade Federal de São João Del Rei, São João Del Rei, Brazil

³ King's College London, London, United Kingdom

⁴ University of Oxford, Oxford, United Kingdom

Abstract

This study benchmarks two open-source cardiac simulators, MonoAlg3D (GPU) and openCARP (MPI), across scenarios of increasing complexity. Despite producing numerically congruent results, their performance profiles revealed distinct architectural strengths. MonoAlg3D on a single GPU demonstrated a computational efficiency comparable to openCARP when parallelized across many CPU cores. These findings highlight two distinct parallelization paradigms: MonoAlg3D is optimized for high-throughput simulations on single GPU nodes, while openCARP is designed for large-scale distributed computing on HPC clusters.

1. Introduction

Computational modeling of cardiac electrophysiology is a critical tool for understanding the mechanisms of arrhythmias and for the development of new therapeutic strategies. A variety of simulators, each with unique numerical methods and parallelization strategies, are available to the research community. However, the choice of a simulator can significantly impact the accuracy and computational cost of the simulations.

This study presents a comparative analysis of two open-source cardiac electrophysiology simulators: MonoAlg3D [1] and openCARP [2], using a standardized benchmark problem [3] to evaluate their performance and accuracy. In order to provide a robust comparison, this work evaluates the simulators across a series of scenarios with increasing complexity. The analysis begins with the original benchmark using the ten Tusscher (2006) ionic model [4]. Subsequently, the computational load is increased by replacing the cellular model with the more complex O'Hara-Rudy (2011) model [5] within the same domain. Finally, the capability of both simulators to handle realis-

tic geometries is tested on an anatomical ventricular mesh derived from medical images. Performance and scalability were quantified by comparing the different parallelization approaches: MonoAlg3D's GPU and OpenMP-threaded CPU implementations were evaluated against openCARP's MPI-based implementation. The ultimate goal is to provide a clear guide for researchers, outlining in which research scenarios the use of one simulator is more advantageous than the other.

2. Methods

2.1. Simulator Descriptions

MonoAlg3D is an open-source cardiac electrophysiology simulator that solves the monodomain model using the finite volume method (FVM) [1]. The simulator is designed for high-performance computing, employing two parallelization strategies: OpenMP for multi-threaded execution on multi-core CPUs, and a CUDA-based implementation for acceleration on NVIDIA GPUs.

OpenCARP is an open-source cardiac simulation framework that solves both the monodomain and bidomain models using the finite element method (FEM) [2]. It is widely used for multi-scale and multi-physics simulations of cardiac function. For high-performance computing, openCARP relies on the Message Passing Interface (MPI) for parallelization, enabling distributed-memory execution on large-scale clusters.

2.2. Benchmarks and Test Cases

To provide a comprehensive comparison, three distinct test cases were employed, systematically increasing in complexity.

2.2.1. Standard Benchmark

The first case (Case 1) is based on the benchmark proposed by Niederer et al. [3]. It uses a slab of tissue with dimensions of $20.0 \times 7.0 \times 3.0$ mm, with fiber orientation aligned with the x -axis and zero-flux boundary conditions applied to all surfaces. A stimulus is applied to a $1.5 \times 1.5 \times 1.5$ mm region at the origin corner ($x=0, y=0, z=0$). This benchmark was simulated using the ten Tusscher et al. (2006) cell model [4], which consists of 18 state variables, with a total simulation time of 80 ms with a time resolution of $\Delta t = 0.02$ ms.

2.2.2. Effect of Ionic Model Complexity

For the second case (Case 2) and to evaluate the simulators' performance under a heavier workload, the geometry and protocol from Case 1 were maintained, but the cell model was replaced with the more computationally intensive O'Hara-Rudy (2011) model [5]. This model includes a more detailed representation of ionic currents, which increases the number of state variables to 41 and, consequently, the complexity of the ordinary differential equations (ODEs) to be solved at each time step. The simulation was performed with a time resolution of $\Delta t = 0.005$ ms.

2.2.3. Anatomical Geometry

For the third scenario (Case 3), an anatomical human ventricle mesh was used. The mesh was derived from a publicly available virtual cohort by Rodero et al. [6], which consists of four-chamber heart models generated from cardiac Computed Tomography (CT) images of asymptomatic subjects. The generation process involved automatic segmentation, creation of unstructured tetrahedral meshes with a 1 mm average edge length for all tissues except blood pools and papillary muscles, and the assignment of rule-based myocardial fiber orientations. For this study, a ventricular geometry was isolated from the four-chamber model to serve as the simulation domain. For the MonoAlg3D simulations, the model was re-meshed to an average element edge length of 500 μ m, yielding 1,306,274 finite volumes. For the openCARP simulations, the original 1 mm mesh was used, consisting of 1,340,386 finite elements. The simulation was run for 1000 ms, initiated by a 2 ms stimulus with an amplitude of 38 μ A/cm³ applied at the ventricular apex.

2.3. Computational Environment

All simulations were performed on a high-performance computing cluster in a node with the following specifications:

- CPU: 2 x AMD EPYC 7713 (128 total cores)
- GPU: 2 x NVIDIA A100 (80 GB of memory)
- System Memory: 512 GB DDR4 RAM

2.4. Execution Protocol and Metrics

For each of the three test cases, simulations were executed using three distinct configurations:

1. **MonoAlg3D**: Executed with a single GPU varying the number of OpenMP threads from 1 to 128.
2. **openCARP**: Executed on the CPU, varying the number of MPI processes from 1 to 128.

All tests for Cases 1 and 2 were conducted using three spatial resolutions: 250, 100, and 50 μ m. For each specific configuration, the simulation was repeated 10 times to ensure statistical significance.

The primary metric was the total wall-clock time, reported as the mean and standard deviation of the 10 runs. For the standard benchmark, accuracy was also assessed by comparing the activation times of cells along the main diagonal of the slab.

3. Results

The analysis of the results begins with the verification of the accuracy and numerical convergence of the MonoAlg3D and openCARP simulators for the standard benchmark (Case 1). This step is essential to ensure that both numerical methods (FVM and FEM) produce biophysically consistent results before the performance analysis.

Figure 1 shows the activation times along the main diagonal of the domain for the three evaluated spatial resolutions: 250 μ m, 100 μ m, and 50 μ m. For each simulator, the solution is observed to converge as the mesh is refined, with the 100 μ m and 50 μ m curves showing close agreement.




Figure 1: Activation times along the domain diagonal for openCARP (green) and MonoAlg3D (red) at three mesh resolutions.

3.1. Performance Analysis for the Standard Benchmark

Following the accuracy validation, the computational performance of the simulators was compared. Figure 2 shows the execution times for the three spatial resolutions, with all results normalized by the execution time of openCARP using a single MPI process. The bars for openCARP (green) show a reduction in execution time as the number of MPI processes increases. For MonoAlg3D (red), the results show the performance on a single GPU as a function of the number of managing host OpenMP threads.




Figure 2: Performance for the standard benchmark (ten Tusscher model) comparing openCARP (green) and MonoAlg3D (red). All execution times are normalized by the time of openCARP with one process. Bar subdivisions distinguish the solution times for the Partial (PDEs, lighter shades) and Ordinary (ODEs, darker shades) Differential Equations.

The data in Figure 2 indicates that the relative performance of the GPU implementation increases with mesh refinement. At the 250 μm resolution, the optimal GPU performance is equivalent to openCARP running with approximately 14 MPI processes. For the 50 μm mesh, the best result for MonoAlg3D (101s, achieved with 16 host OpenMP threads) is 1.5 times faster than openCARP with 128 processes (152.4s).

An analysis of the time breakdown reveals another key architectural difference. For openCARP, as the mesh is refined, the fraction of the total runtime dedicated to solve the ODEs (darker green portion) decreases, making the PDE solution the dominant computational bottleneck. In contrast, this effect is less pronounced for MonoAlg3D, where the proportional cost of the ODE solution remains more stable across the different resolutions.

3.2. Effect of Ionic Model Complexity

To investigate how these performance characteristics change with a greater computational load per element, the

benchmark was repeated using the O’Hara-Rudy model. This more demanding ionic model significantly increases the work per element. The resulting runtime behavior, presented with the same normalization scheme, is shown in Figure 3.




Figure 3: Performance for the benchmark using the O’Hara-Rudy ionic model. The chart follows the same representational scheme (normalization, colors, and subdivisions) as presented in Figure 2.

With the exception of the 250 μm resolution, where a higher number of openCARP processes was required to match the GPU performance, the final performance relationship at the other resolutions remained similar to that of the ten Tusscher model. The key distinction in this benchmark is the initial performance difference: the runtime gap between MonoAlg3D on a single GPU and openCARP on a single process is substantially larger for the O’Hara-Rudy model. However, openCARP exhibits more effective scaling behavior with this computationally heavier model than it did with the ten Tusscher benchmark, which allows it to narrow the initial performance gap at higher process counts. This indicates that openCARP’s parallel efficiency is more sensitive to an increase in the number of domain elements than to the increase in the computational complexity of the cellular model.

This increased computational load also altered the runtime composition for each simulator. The ionic model solution remained the primary computational load for openCARP, accounting for 89% of the time with few processes and decreasing to 60% at higher scales. Conversely, for the MonoAlg3D, this task represented a smaller proportional cost, ranging from 60% down to 40%. This result further highlights the GPU’s efficiency in handling the massively parallel workload of the cellular models.

3.3. Anatomical Geometry

Finally, to assess the simulators’ performance on a realistic, unstructured mesh, the anatomical ventricle simulation (Case 3) was performed. A qualitative validation was first conducted to ensure that both platforms produced

comparable biophysical results in this complex domain. Figure 4 shows the transmembrane potential maps for both simulators at $t = 450$ ms. The resulting electrical wave propagation patterns are visually congruent, confirming that both simulators reached a similar numerical solution.




Figure 4: Qualitative comparison of transmembrane potential maps (mV) at $t = 450$ ms for the anatomical simulation using both simulators.

MonoAlg3D, running on a single GPU, completed the simulation in 68 s, achieving a performance comparable to the 70 s required by openCARP when parallelized across 128 MPI processes. A key factor in this comparison was openCARP's strong scalability, particularly in its solution of the ODEs. Congruent with the observations from Case 2, the portion of runtime that openCARP dedicated to the cellular models decreased from 50% of the total execution time on a single processor to just 20% with 128 processors. In contrast, the highly optimized GPU implementation of MonoAlg3D spent only 16% of its total runtime on the ODE solution, dedicating the majority of the time to the PDE component.

4. Conclusion

This study conducted a comparative benchmark of two open-source cardiac electrophysiology simulators, MonoAlg3D and openCARP, across a range of computational scenarios. The analysis concludes that both are robust and viable platforms, each with distinct architectural strengths that make them advantageous for different research paradigms.

OpenCARP, with its MPI-based parallelization, demonstrated strong scalability suitable for distributed-memory environments. This makes it an excellent choice for extremely large-scale simulations that can leverage hundreds or thousands of processor cores on high-performance computing clusters.

In contrast, MonoAlg3D, leveraging its GPU-based implementation, provides a different high-throughput parallelization strategy. It showed exceptional performance on a single computational node, outperforming the multi-

core CPU implementation of openCARP on the same hardware. This makes it highly effective for scenarios where many simulations are required on individual GPU-equipped nodes.

The comparative analysis demonstrated consistent performance characteristics across both benchmark and anatomically detailed simulations, supporting the applicability of the results to realistic research contexts. The selection between MonoAlg3D and openCARP should therefore be determined by the specific research objectives and available computational resources.

Acknowledgments

The authors would like to express their thanks to Minas Gerais State Research Support Foundation (FAPEMIG) - PCE-00048-25; APQ-02752-24, APQ-02445-24, FINEP - AV020062/22, "Coordenação de Aperfeiçoamento de Pessoal de Nível Superior" (CAPES), "Empresa Brasileira de Serviços Hospitalares" (Ebserh), CNPq and Federal University of Juiz de Fora (UFJF) for funding this work.

References

- [1] Oliveira RS, Rocha BM, Burgarelli D, Meira WJ, Constantinides C, dos Santos RW. Performance evaluation of gpu parallelization, space-time adaptive algorithms, and their combination for simulating cardiac electrophysiology. *Int J Numer Meth Biomed Engng* 2018;34:e2913.
- [2] Gernot P, Loewe A, Neic A, et al. The openCARP simulation environment for cardiac electrophysiology. *Computer Methods and Programs in Biomedicine* 2021;208:106223.
- [3] Niederer SA, et al. Verification of cardiac tissue electrophysiology simulators using an n-version benchmark. *Philos Trans A Math Phys Eng Sci* 2011 Nov;369(1954):4331-51;.
- [4] ten Tusscher KH, Panfilov AV. Alternans and spiral breakup in a human ventricular tissue model. *Am J Physiol Heart Circ Physiol* 2006 Sep;291(3):H1088-100;.
- [5] O'Hara T, Virág L, Varró A, Rudy Y. Simulation of the undiseased human cardiac ventricular action potential: Model formulation and experimental validation. *PLOS Computational Biology* 05 2011;7(5):1-29.
- [6] Rodero C, Strocchi M, Marciniak M, Longobardi S, Whitaker J, O'Neill MD, Gillette K, Augustin C, Plank G, Vigmond EJ, Lamata P, Niederer SA. Linking statistical shape models and simulated function in the healthy adult human heart. *PLOS Computational Biology* 04 2021;17(4):1-28.

Address for correspondence:

Rodrigo Weber dos Santos

Departamento de Ciência da Computação da UFJF, Juiz de Fora
rwdsantos@yahoo.com