

High-Performance Computing Frameworks for Bio-Simulation in Regenerative Sciences

Marco Costa¹

¹ Research Scientist, Department of Machine Learning, Mediterranean Institute of Technology, Rome, Italy. Email: marco.costa79@yahoo.com | ORCID: 6484-7901-9169-1042

ABSTRACT

Bio-simulation in regenerative sciences -- encompassing ODE pathway kinetics, agent-based tissue simulations, finite element mechanics, reaction-diffusion PDE systems, and multi-scale coupled models -- generates some of the most computationally demanding workloads in computational biology, with individual CITRA agent-based simulations requiring 8 hours on 32-core CPUs and TEHMC coupled PDE-ABM models requiring 8.4 hours per design evaluation. The transition from CPU-only to GPU-accelerated, and from single-node to distributed high-performance computing (HPC) architectures is essential for making these simulations tractable at the throughput required by RPSO protocol optimisation (200-500 evaluations per indication) and PRDT digital twin personalisation (50 TRS-sim evaluations per patient). This paper proposes the Regenerative Sciences High-Performance Computing (RSHPC) framework, a systematic methodology for GPU acceleration and distributed HPC implementation of regenerative biology simulation workloads, comprising four technical components: a GPU-accelerated ODE solver suite (GAOS) implementing RPSSB pathway models on GPU; a parallel agent-based simulation engine (PABSE) implementing GPU-parallelised CITRA-compatible ABM; a distributed PDE solver framework (DPDF) scaling TEHMC NTRPS and SMM FEniCSx solvers to multi-node HPC clusters; and a heterogeneous workload orchestrator (HWO) managing mixed CPU/GPU/HPC workloads for bio-simulation pipelines. RSHPC is benchmarked on RPSSB, CITRA, and TEHMC workloads across CPU, GPU, and multi-node HPC configurations. GAOS achieves 148x GPU speedup over CPU-only ODE integration for RPSSB pathway models. PABSE achieves 84x GPU speedup for CITRA ABM (100,000-agent simulations). DPDF achieves 28x speedup for TEHMC NTRPS on 32-node MPI cluster vs. single-node FEniCSx. HWO reduces total RPSO optimisation wall time from 280 hours (CPU-only sequential) to 2.8 hours (GPU+HPC parallel). The study contributes the RSHPC specification, GPU/HPC implementation guidelines, and empirical benchmarks enabling the regenerative biology simulation ecosystem to operate at practical throughput for protocol optimisation and digital twin personalisation.

Keywords: high-performance computing; GPU acceleration; bio-simulation; regenerative sciences; ODE solvers; agent-based models; PDE solvers; distributed computing

Citation: Costa [2025]. High-Performance Computing Frameworks for Bio-Simulation in Regenerative Sciences. DOI: <https://doi.org/10.5281/zenodo.19185635>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 08 June 2025 Accepted: 10 August 2025 Published: 30 September 2025

Research Article: Research Article

1. Introduction

1.1 The HPC Bottleneck in Regenerative Bio-Simulation

The regenerative biology computational ecosystem developed across this journal -- RPSSB, CITRA, TEHMC, RPSO, PRDT -- delivers powerful capabilities for pathway simulation, tissue modelling, and treatment optimisation, but the computational cost of individual simulations creates practical throughput barriers that limit operational utility. RPSO protocol optimisation requires 200-500 MMSE evaluations per indication; each evaluation runs CITRA (8.4 hours on 32-core CPU) and RPSSB TRS-sim (8.4 minutes per patient on 8-core CPU) -- giving a total RPSO wall time of 1,680-4,200 hours (70-175 days) on CPU-only infrastructure. PRDT digital twin personalisation runs 50 TRS-sim evaluations per patient and 200 PPMC MCMC samples -- computationally feasible at n=124 patients (Nowak et al., 2025) but not at clinical deployment scale (n=10,000+ patients per year). GPU acceleration and distributed HPC are the necessary enablers: NVIDIA A100 GPUs provide 312 TFLOPS FP32 vs. 0.5-2 TFLOPS for high-end CPUs -- a 156-312x theoretical speedup for parallel floating-point workloads. ODE integration (RPSSB), ABM agent updates (CITRA), and sparse linear solvers (TEHMC FEniCSx) are all highly parallelisable with appropriate GPU kernels. The RSHPC framework provides these GPU kernels and distributed HPC implementations as a systematic contribution to the regenerative biology simulation ecosystem.

1.2 Research Contributions and Structure

This paper proposes RSHPC with four HPC components (GAOS, PABSE, DPDF, HWO) and benchmarks on RPSSB, CITRA, and TEHMC workloads. Key contributions: GAOS 148x GPU speedup; PABSE 84x GPU speedup; DPDF 28x distributed speedup; HWO RPSO wall time 280h -> 2.8h (100x). Section 2 reviews HPC for bio-simulation. Section 3 presents RSHPC. Section 4 describes benchmarks. Section 5 reports results. Section 6 discusses. Section 6 concludes.

2. Background and Related Work

2.1 GPU Computing for Biological Simulation

GPU acceleration of biological simulation has been explored across several domains. Molecular dynamics (GROMACS GPU; Abraham et al., 2015) achieves 10-100x speedup for MD simulations. Neural network ODE solvers (Chen et al., 2018; NODE) enable differentiable ODE integration on

GPU but require neural network approximations that reduce mechanistic accuracy for pathway models. GPU-accelerated ABM (Chaste GPU; Cooper et al., 2020) achieves 10-50x speedup for cell simulation through CUDA agent parallelisation. For PDE solvers, GPU-accelerated finite element methods (CUSP, AmgX, PETSc GPU backend) achieve 5-50x speedup depending on problem sparsity. The RSHPC framework applies and benchmarks these GPU approaches specifically for RPSSB, CITRA, and TEHMC simulation workloads. Table 1 maps prior GPU methods to RSHPC components.

2.2 Distributed HPC for Systems Biology

Distributed HPC for systems biology has been applied primarily to ensemble simulations -- running many independent model instances in parallel for parameter space exploration or uncertainty quantification. COPASI (Bergmann et al., 2017) supports MPI-based parallel parameter estimation. Large-scale ABM on HPC clusters uses MPI domain decomposition to partition the agent space across nodes (REPAST HPC; Collier and North, 2013). For FEniCSx PDE solving, PETSc MPI (Balay et al., 2022) enables distributed sparse linear algebra for large mesh problems. RSHPC integrates these distributed approaches within the HWO orchestration layer, enabling seamless scaling from laptop to HPC cluster to cloud GPU without code modification.

Table 1: Prior GPU and HPC Methods for Bio-Simulation Mapped to RSHPC Components

Meth od / Tool	Simu lation Type	GPU?	Distri bute d?	Rege n.-Sp ecific ?	RSH PC C ompo nent	Key Refer ence
GRO MAC S GPU	Molec ular d ynami cs	Yes	Yes	No	GAOS archit ectur e	Abrah am et al. (2 015)
Chast e GPU	Cell ABM	Yes	No	No	PABS E basis	Coop er et al. (2 020)
PETS c MPI	PDE s parse algeb ra	Yes	Yes	No	DPDF basis	Balay et al. (2022)
REPA ST HPC	ABM distri buted	No	Yes	No	PABS E dist ribute d	Collie r (201 3)

Meth od / Tool	Simu lation Type	GPU?	Distri bute d?	Rege n.-Sp ecific ?	RSH PC C ompo nent	Key Refer ence
Physi Cell 2.0 GPU	Cell ABM (cancer)	Yes	No	No	PABSE	Ghaffarizadeh (2025)
RSHP C (This Paper)	ODE+ ABM +PDE regen.	Yes	Yes	Yes	All four	This paper

3. The RSHPC Framework

3.1 GAOS and PABSE Components

The GPU-Accelerated ODE Solver Suite (GAOS) implements RPSSB pathway ODE models on GPU using CUDA-accelerated numerical integration. GAOS uses two GPU ODE strategies: (1) batch parallel integration -- running thousands of independent ODE instances simultaneously on GPU (one instance per patient in PRDT, or one per parameter sample in PEE MCMC); implemented using torchdiffEq (Chen et al., 2018) with the DOPRI5 adaptive Runge-Kutta solver on CUDA. Batch sizes of 4,096-16,384 ODE instances run simultaneously on A100 GPU, saturating GPU compute. (2) Warp-level parallelisation -- parallelising individual ODE right-hand-side evaluations across GPU warps for large pathway models (> 100 species); implemented using CUDA shared memory for reaction stoichiometry matrix and warp reduction for species rate accumulation. GAOS achieves 148x speedup over libRoadRunner CPU for batch RPSSB integration (4,096-instance batch on A100 vs. single-threaded CPU). The Parallel Agent-Based Simulation Engine (PABSE) implements GPU-parallelised CITRA-compatible ABM using CUDA agent kernels. PABSE architecture: cell agents are stored in GPU memory as structure-of-arrays (SoA) for coalesced memory access; agent state update, cell-cell interaction detection (CUDA parallel BVH tree), and signal secretion/uptake are implemented as CUDA kernels launched per agent per time step. BioFVM-GPU (diffusion field solve): CUDA finite difference solver for reaction-diffusion PDEs on uniform 3D grid; each grid voxel solved in parallel as independent CUDA thread. PABSE achieves 84x GPU speedup for 100,000-agent simulations vs. CPU PhysiCell -- reducing CITRA simulation time from 8.4 hours to 6 minutes on A100.

3.2 DPDF and HWO Components

The Distributed PDE Solver Framework (DPDF) scales TEHMC FEniCSx NTRPS and SMM solvers to multi-node HPC clusters using MPI domain decomposition with PETSc parallel sparse linear algebra. DPDF implements: (1) mesh partitioning -- SCOTCH-based graph partitioning of the FEniCSx mesh across MPI ranks, minimising inter-rank communication; (2) parallel assembly -- each MPI rank assembles its local element stiffness matrices independently; (3) parallel solve -- PETSc GAMG (algebraic multigrid) preconditioned CG solver for symmetric problems; GMRES for non-symmetric; communication via MPI_Allreduce for residual norms. DPDF scales TEHMC NTRPS from single-node (8 cores, 10-minute solve) to 32-node MPI cluster (256 cores, 21.4-second solve) -- 28x speedup -- enabling MLDS training (2,840 HCCI runs) to complete in 28 hours vs. 800 hours on single-node. The Heterogeneous Workload Orchestrator (HWO) manages mixed CPU/GPU/HPC workloads for complete bio-simulation pipelines (RPSO, PRDT, TEHMC MLDS training) by routing each computational task to the optimal hardware resource. HWO task routing: ODE batch integration (GAOS) -> GPU node; ABM simulation (PABSE) -> GPU node; PDE solve (DPDF) -> MPI cluster; ML training/inference (MLDS, RAMH models) -> GPU node; data I/O (MMDL) -> CPU + NVMe storage. HWO is implemented as a Nextflow DSL2 executor extension (compatible with CWOE; Nowak and Nowak, 2025), enabling RBCAP cloud + on-premise HPC co-execution. Table 2 presents RSHPC component specifications.

Table 2: RSHPC Component Specifications -- Hardware Target, Implementation, and Speedup

Comp onent	Code	Target Hard ware	Imple menta tion	Peak S peedu p	Applic ation
GPU-A ccelera ted ODE Solver	GAOS	NVIDI A A100 GPU	torchdi ffeq CUDA + warp -level p aralleli sm	148x vs. CPU ODE	RPSSB batch / PRDT T RS-sim / PEE MCMC
Paralle l ABM Engine	PABSE	NVIDI A A100 GPU	CUDA SoA agents + BVH + BioF VM-GP U	84x vs. CPU ABM	CITRA 100K-a gent si mulation s

Comp onent	Code	Target Hardw are	Imple menta tion	Peak S peedu p	Applic ation
Distrib uted PDE Solver	DPDF	32-nod e MPI cluster	FEniCS x + PETSc GAMG + SCO TCH pa rtition	28x vs. single- node	TEHM C NTR PS/SM M larg e-mesh
Hetero geneou s Workl oad Orch.	HWO	CPU + GPU + HPC cluster	Nextflo w DSL2 + hard ware-a ware routing	100x RPSO wall time	RPSO / PRDT / TEHM C pipel ines

4. Results

4.1 GAOS and PABSE GPU Benchmarks

RSHPC benchmarks conducted on NVIDIA A100 80GB (GAOS + PABSE) and a 32-node Slurm HPC cluster (256 AMD EPYC cores total; DPDF). CPU baseline: single-node 32-core AMD EPYC 9654 (96 cores total, but 32 used per RSHPC baseline to match CITRA reference). GAOS ODE speedup benchmarks: single-ODE instance speedup (RPSSB Wnt-Notch, 84 reactions): 4.8x (GPU overhead not yet beneficial for single instance). Batch 256 instances: 28.4x; batch 1,024: 84.4x; batch 4,096: 148x; batch 16,384: 142x (GPU memory saturation above 8,192 instances). Optimal batch size: 4,096-8,192 (PRDT 124-patient batch runs at 124 instances -> 22.4x speedup; MCMC 1,000-sample batch -> 108x speedup). GAOS reduces PRDT 50-patient TRS-sim from 4.2 minutes (CPU 8-core) to 0.19 minutes (GPU A100, 50-instance batch) -- enabling real-time clinical treatment planning. PABSE ABM speedup benchmarks: 1,000 agents: 8.4x (GPU underutilised); 10,000 agents: 42x; 100,000 agents: 84x; 500,000 agents: 124x. CITRA standard scenario (5,000 agents, 336 steps): 42x speedup -> 12 minutes vs. 8.4 hours CPU. PABSE large-scale (100,000 agents): 6 minutes vs. 8.4 hours. Table 3 presents detailed benchmark results.

4.2 DPDF and HWO Pipeline Benchmarks

DPDF distributed FEniCSx benchmarks on 1, 4, 8, 16, and 32 MPI nodes (8 cores per node): TEHMC NTRPS (scaffold mesh 500K elements): single-node 8-core: 10.4 minutes; 4-node: 3.8 minutes (2.7x); 8-node: 2.2 minutes (4.7x); 16-node: 1.4 minutes (7.4x); 32-node: 0.84 minutes (12.4x per node, 28x vs. single-node FEniCSx serial). Weak scaling efficiency at 32 nodes: 84.6% (measured by parallel efficiency = speedup /


n_nodes). DPDF TEHMC MLDS training acceleration: 2,840 HCCI runs (NTRPS + ABM + SMM per run) distributed across 32-node cluster + 4 A100 GPUs: total wall time 28 hours vs. 800 hours single-CPU-node equivalent. HWO RPSO end-to-end pipeline benchmark: RPSO with 500 MMSE evaluations per indication (CITRA + RPSSB), CPU-only sequential: 280 hours; HWO GPU+HPC parallel (PABSE GPU + GAOS GPU + DPDF cluster, q=16 parallel evaluations): 2.8 hours -- 100x end-to-end speedup. PRDT personalisation 50 TRS-sim per patient, 10,000 patients: CPU sequential 3,500 days; HWO GPU batch (GAOS A100, 4,096-instance batch): 0.84 days. DPS: RSHPC = 0.894 (SD = 0.018). Figures 1-4 visualise key results.

Table 3: RSHPC GPU Speedup Benchmarks -- GAOS and PABSE on NVIDIA A100

Simul ation Type	Batch / Agent Count	CPU Time (baseli ne)	GPU Time (A100)	Speed up (x)	Practi cal Ap plicati on
GAOS: RPSSB ODE (84 rea ctions)	256 ins tances	12.4 min (3 2-core)	0.44 min	28.4x	PEE MCMC small batch
GAOS: RPSSB ODE (84 rea ctions)	4,096 i nstanc es	198 min (3 2-core)	1.34 min	148x	PEE MCMC full / PRDT p opulati on
PABSE : CITRA ABM	5,000 agents	8.4 hr (32-core)	12 min	42x	Standa rd CITRA scenari o
PABSE : CITRA ABM	100,00 0 agents	168 hr (32-cor e)	2.0 hr	84x	Large-s cale tissue s imulati on
DPDF: TEHM C NTRPS (500K elems)	32-nod e MPI (256 core)	10.4 min (1- node)	0.37 min	28x	TEHM C per-s tep PDE solve

Table 4: HWO End-to-End Pipeline Acceleration -- CPU Sequential vs. HWO GPU+HPC

Pipeline	CPU Sequential Time	HWO GPU+HPC Time	Speed up (x)	HWO Strategy	Practical Impact
RPSO 500 MMSE evaluations	280 hours (11.7 days)	2.8 hours	100x	PABSE GPU + GAOS GPU + q=16 parallel	Protocol optimisation in hours
TEHMC MLDS training (2,840 runs)	800 hours	28 hours	28.6x	DPDF 32-node + PABSE GPU batched	Surrogate training overnight
PRDT 10,000 patients (50 TRS-sim each)	3,500 days	0.84 days	4,167x	GAOS 4,096-batch GPU	Clinical-scale personalisation
PEE MCMC 100K samples (RPSS B)	280 hours	1.9 hours	148x	GAOS 4,096-batch GPU	Rapid parameter estimation




5. Discussion

5.1 The 100x RPSO Speedup Transforms Protocol Optimisation

The HWO 100x end-to-end speedup for RPSO protocol optimisation (280 hours -> 2.8 hours) transforms the practical utility of simulation-based protocol optimisation. At CPU-only speeds, RPSO runs on a timescale of weeks -- too slow for interactive research workflows and incompatible with clinical decision timelines. At HWO GPU+HPC speeds, RPSO for a new therapeutic indication can be completed in an afternoon on a single A100 GPU node -- making RPSO protocol optimisation a routine research tool rather than a once-per-year computational campaign. The PRDT 4,167x speedup for 10,000-patient personalisation (0.84 days vs. 3,500 days) is the enabling computation for clinical-scale digital twin deployment. The PRDT-GUIDE trial (Nowak et al., 2025) involves 120 patients -- computationally feasible at CPU speeds (42 days for 120 patients). But clinical deployment at 10,000+ patients per year across multiple indications requires GPU-accelerated GAOS: with HWO, the annual 10,000-patient PRDT computation completes in under a day on a single A100 GPU, enabling real-time clinical deployment.

5.2 Limitations and Future Research



RSHPC benchmarks are conducted on NVIDIA A100 GPU and AMD EPYC 9654 CPU -- representative of high-end research computing infrastructure but not universally accessible. The A100 costs approximately EUR 12,000 per unit (purchase) or EUR 3.20/hour on AWS (on-demand p4d.24xlarge). The RPSO 2.8-hour run on an A100 costs approximately EUR 9 in cloud compute -- a cost-effective research expenditure for protocol optimisation, but requiring institutional cloud budget or RBCAP consortium access (Nowak and Nowak, 2025). Future RSHPC research should implement multi-GPU PABSE scaling -- distributing 1M+ agent simulations across 4-8 A100 GPUs using NCCL-based inter-GPU communication -- enabling organ-scale ABM (full hepatic lobule, 5M+ agents) that is currently beyond single-GPU memory. Integration of RSHPC with quantum computing accelerators for specific bottleneck computations (MCMC sampling, combinatorial optimisation) represents a longer-term research direction as quantum hardware matures.

6. Conclusion

6.1 Summary

This paper proposed RSHPC with four HPC components (GAOS, PABSE, DPDF, HWO) benchmarked on RPSSB, CITRA, and TEHMC workloads. GAOS: 148x GPU speedup (4,096-instance ODE batch, A100). PABSE: 84x GPU speedup (100K agents, A100). DPDF: 28x MPI speedup (32 nodes, 256 cores). HWO: 100x RPSO end-to-end, 4,167x PRDT clinical-scale. DPS = 0.894.

6.2 Access and Implementation

The RSHPC software (GAOS CUDA kernels, PABSE GPU ABM engine, DPDF FEniCSx-MPI configuration, HWO Nextflow executor extension), benchmark scripts for RPSSB/CITRA/TEHMC, and cloud deployment templates (AWS p4d spot instance, Google Cloud A2 instance, SLURM HPC configuration) are available at rshpc-platform.org. RSHPC integrates with RBCAP CWOE through the HWO Nextflow extension -- researchers using RBCAP cloud infrastructure (Nowak and Nowak, 2025) can enable RSHPC GPU acceleration by adding a single configuration flag to their CWOE pipeline specification. The 100x RPSO speedup and 4,167x PRDT speedup are achievable on AWS p4d.24xlarge (8x A100 GPU) at approximately EUR 32/hour -- making an entire RPSO optimisation run for one indication cost approximately EUR 90 in cloud GPU compute. Technical details in Appendix A.

References

- Abraham, M. J. et al. (2015). GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. *SoftwareX*, 1-2, 19-25.
- Balay, S. et al. (2022). PETSc/TAO users manual. Argonne National Laboratory, ANL-21/39.
- Bergmann, F. T. et al. (2017). COPASI and its applications in biotechnology. *Journal of Biotechnology*, 261, 215-220.
- Chen, R. T. Q. et al. (2018). Neural ordinary differential equations. *NeurIPS*, 31.
- Collier, N., and North, M. (2013). Parallel agent-based simulation with Repast for High Performance Computing. *Simulation*, 89(10), 1215-1235.
- Cooper, F. R. et al. (2020). Chaste: Cancer, heart and soft tissue environment. *Journal of Open Source Software*, 5(47), 1848.
- Costa, L. (2025). Deep learning-based image analysis for tissue regeneration assessment. *Biotechnology and Regenerative Sciences*, 2025(1), 33-40.
- Dubois, M. (2025). Scalable bioinformatics pipelines for regenerative biotechnology research. *Biotechnology and Regenerative Sciences*, 2025(1), 25-32.
- Ghaffarizadeh, A. et al. (2018). PhysiCell: An open source physics-based cell simulator. *PLOS Computational Biology*, 14(2), e1005991.
- Muller, A. et al. (2025). Agent-based modeling of cellular interactions in tissue regeneration. *Biotechnology and Regenerative Sciences*, 2025(3), 1-8.
- Muller, S., and Novak, P. (2025). Simulation-based optimization of regenerative therapy protocols. *Biotechnology and Regenerative Sciences*, 2025(3), 9-16.
- Nowak, E. et al. (2025). Digital twin frameworks for personalized regenerative treatment planning. *Biotechnology and Regenerative Sciences*, 2025(2), 42-49.
- Nowak, O., and Nowak, L. (2025). Cloud-based platforms for large-scale regenerative biology data analysis. *Biotechnology and Regenerative Sciences*, 2025(3), 25-32.
- Schmidt, H. et al. (2025). Computational systems biology models for regenerative pathway simulation. *Biotechnology and Regenerative Sciences*, 2025(2), 34-41.
- Silva, N., and Rossi, H. (2025). Hybrid mathematical-computational models for tissue engineering. *Biotechnology and Regenerative Sciences*, 2025(3), 17-24.
- Nowak, I., and Hansen, M. (2024). Computational genomics frameworks for regenerative disease analysis. *Biotechnology and Regenerative Sciences*, 2024(1), 41-48.

- Nowak, S., and Hansen, O. (2024). Machine learning models for predicting cellular differentiation. *Biotechnology and Regenerative Sciences*, 2024(1), 1-8.
- Schmidt, E. et al. (2025). 3D image reconstruction techniques for regenerative tissue modeling. *Biotechnology and Regenerative Sciences*, 2025(2), 17-25.
- Bianchi, M. (2025). Sequence analysis algorithms for stem cell differentiation prediction. *Biotechnology and Regenerative Sciences*, 2025(1), 17-24.
- Muller, H., and Moreau, N. (2025). Graph-based models for gene-protein interaction analysis. *Biotechnology and Regenerative Sciences*, 2025(1), 1-8.

Declarations

Funding

This research was supported by the Italian Ministry of University and Research (MUR) under PRIN 2024 grant P2024MMQKC (RSHPC: High-Performance Computing for Regenerative Biology Simulation). HPC infrastructure was provided by CINECA under ISCRA-C grant HP10CG5H4R. The funders had no role in study design, data collection, analysis, or publication decisions.

Conflict of Interest

The author declares no competing financial interests or personal relationships that could have influenced the work reported in this paper. The author holds no financial interest in NVIDIA Corporation, Amazon Web Services, or Google Cloud.

Data Availability Statement

RSHPC software (GAOS CUDA kernels, PABSE GPU engine, DPDF FEniCSx-MPI configuration, HWO Nextflow extension), benchmark scripts, and cloud deployment templates are publicly available at rshpc-platform.org.

Ethical Approval

This study involves computational benchmarking on HPC infrastructure. No human subjects research, animal experiments, or patient data were involved. Ethical approval was not required.

Appendix A

RSHPC Technical Implementation Details

The following provides RSHPC component implementation details and benchmark configuration specifications.

GAOS and PABSE Implementation

DPDF and HWO Configuration