

GPU Acceleration of an Extended Local Effect Model for Efficacy Estimation in Radionuclide Therapy

Taewan Kim^{a,b}, Jeong Hoon Park^a, Min Goo Hur^a, Sung-Joon Ye^{b,*}

^aCyclotron Application Research Section, Korea Atomic Energy Research Institute, Jeongseup Si, Jeollabuk Do, Republic of Korea

^bDepartment of Applied Bioengineering and Research Institute for Convergence Science, Graduate School of Convergence Science and Technology, Seoul National University, Seoul, Republic of Korea

*Corresponding author: sye@snu.ac.kr

***Keywords :** Radionuclide therapy, Local effect model, GPU acceleration, Monte Carlo simulation

1. Introduction

Auger-emitting radionuclide therapy (RNT) requires cell-specific dosimetry models that resolve the nanometer-scale, highly localized energy deposition of Auger electrons. We previously extended the local effect model (LEM) by coupling Geant4-DNA-derived dose point kernels (DPKs) with a kinetic giant loop binary lesion (GLOBLE) model, enabling stochastic Monte Carlo (MC) prediction of survival fraction (SF) and tumor control probability (TCP) for ¹⁰³Pd, ¹¹¹In, and ¹²⁵I across three cell lines, validated against published *in vitro* data [1]. As discussed in [1], the MC engine simulated only a target cell and its immediate neighbors, which was sufficient because the sub- μ m-to-few- μ m range of Auger electrons confines essentially all dose to nearby cells. Long-range beta emitters (e.g., ¹⁷⁷Lu, ⁹⁰Y), by contrast, require simulating a much larger cross-fire domain, and doing so was reported as computationally impractical with the original CPU implementation [1]. In this work, we re-implemented the MC engine of [1] in CUDA C++ and evaluated whether the resulting GPU implementation could reproduce the validated results with substantially reduced computation time and extend feasible simulation domains toward the spatial scale relevant to long-range emitters.

2. Methods and Results

2.1 CUDA Implementation

The extended LEM in [1] is a multicellular MC simulation on a two-dimensional (2D) monolayer culture (figure 1). Because the number of source cells needed to capture the full cross-fire dose contribution scales with the *area* of the simulated domain – i.e. with the *square* of the electron range – a beta emitter's ~ 1000 -fold longer range than an Auger emitter's implies an approximately 10^3 squared, i.e. $\sim 10^6$ -fold, larger required source-cell count, which explains the computational bottleneck reported in [1] (figure 1).

We re-implemented the stochastic MC engine – which samples decay times/locations from DPKs and tracks each $540 \times 540 \times 540 \text{ nm}^3$ nuclear subvolume through the kinetic GLOBLE damage/repair scheme – in CUDA

C++, parallelizing subvolume-level kinetics across GPU threads (896 threads/block). For the baseline hexagonal monolayer geometry used for validation in [1] (19 source cells, 40- μ m-radius domain), the GPU implementation completed 1000 independent histories in ~ 11 s. The CPU implementation in [1] (48 threads, Intel Xeon Gold 6342) required 625.55 h in total for its most demanding scenario (the full TCP-curve set); reproducing an equivalent 10-point cellular-activity sweep (¹²⁵I, 10% intranuclear fraction, 0.1–1.0 Bq/cell) with the GPU implementation completed in 9.7 min in total, with SF and dose outputs consistent with the underlying model.




Fig. 1. Schematic comparison of the cross-fire domain required for Auger- versus beta-emitting radionuclides. In a 2D monolayer, the domain area scales with the square of the electron range, so a ~ 1000 -fold longer beta range requires on the order of 10^6 -fold more source cells than an Auger emitter.

2.2 Large-Scale Multicellular Simulation

Using the validated GPU implementation, the cross-fire domain was expanded from 19 to 10,105 source cells (domain radius 40 μ m to 1 mm) for ¹¹¹In, ¹²⁵I, and ¹⁰³Pd (figure 2). Execution time increased from ~ 11 s to 47–62 min (depending on radionuclide) as the domain grew ~ 530 -fold in cell count, remaining within a practical single-GPU runtime. The predicted 24 h nuclear absorbed dose and SF remained stable across this range (e.g., intranuclear ¹¹¹In: 15.4–15.8 Gy across 19–10,105 cells), consistent with the sub-cellular effective range of Auger electrons and confirming the numerical stability of the GPU implementation at cell counts far beyond those used for model validation in [1].




Fig. 2. GPU-CUDA execution time (left) and 24 h intranuclear absorbed dose for ^{111}In (right) as a function of the number of source cells in the simulated cross-fire domain, from the validated 19-cell baseline up to a 1-mm-radius, 10,105-cell domain.

3. Conclusions

A CUDA re-implementation of the extended LEM removed the principal computational bottleneck of the original CPU-based model while reproducing its results, and extended feasible multicellular simulation domains to a 1-mm radius, approaching the spatial scale relevant to clinically used beta emitters. This GPU-accelerated framework provides a practical basis for future whole-domain MC simulation of long-range beta-emitting radionuclides such as ^{177}Lu and ^{90}Y .

REFERENCES

- [1] T Kim, C-M Lee, T Kim, R Hernández Millares and S-J Ye 2026 An extended local effect model for Auger-emitting radionuclide therapy Phys. Med. Biol. 71 105005.
- [2] C-M Lee and S-J Ye 2024 A GPU-accelerated Monte Carlo code, RT2 for coupled transport of photon, electron/positron, and neutron Phys. Med. Biol. 69 175005.