



RESEARCH ARTICLE

REVISOR **Development and verification of Hornet – A CUDA-optimized, Monte Carlo-agnostic, continuous-energy cross-section processing code**

[version 2; peer review: 3 approved]

Muhammad Rizwan Ali ¹, Murat Serdar Aygul ¹, Deokjung Lee ²¹Department of Nuclear Engineering, Ulsan National Institute of Science and Technology, Ulsan, 44919, South Korea²Advanced Nuclear Technology and Services, Jung-gu, Ulsan, 44429, South Korea

v2 First published: 28 Jul 2025, 3:34
<https://doi.org/10.12688/nuclscitechnolopenres.17676.1>
 Latest published: 24 Jan 2026, 3:34
<https://doi.org/10.12688/nuclscitechnolopenres.17676.2>

Abstract

Background

Hornet is introduced as a GPU-optimized nuclear data processing code designed to accelerate neutron cross-section handling and A Compact ENDF (ACE) file sampling for incident neutron data and thermal scattering laws. Hornet supports continuous-energy Monte Carlo applications and is Monte Carlo-agnostic, making it versatile for academic use and easy integration into various Monte Carlo frameworks.

Methods

Hornet is implemented in CUDA C++ using an object-oriented design to take full advantage of Graphical Processing Unit (GPU) parallelism. Verification was conducted through both standalone tests and integration with GPU-optimized REactor Physics Monte Carlo (GREAPMC) – an in-house Graphical Processing Unit (GPU) Monte Carlo code – across standard benchmark simulations, including International Atomic Energy Agency (IAEA) INDC (USA)-107 pin-cell and Mosteller benchmarks. Performance and precision were assessed by comparing Hornet's output against that of Monte Carlo Simulation (MCS), an in-house CPU-based high-fidelity Monte Carlo code.

Results

Across all benchmark scenarios, Hornet matched MCS in accuracy,

Open Peer Review

Approval Status

	1	2	3
version 2			
(revision)			
24 Jan 2026	view	view	view
version 1			
28 Jul 2025	view	view	

1. **Ho Jin Park**, Kyunghee University College of Engineering, Yongin-si, South Korea
2. **Ilham Variansyah** , Oregon State University, Corvallis, USA
3. **Paul Cosgrove** , University of Cambridge, Cambridge, UK

Any reports and responses or comments on the article can be found at the end of the article.

delivering Root Mean Square (RMS) power differences below 0.1% in pin-by-pin power distribution within an OPR1000 fuel assembly. Remarkable tracking rate enhancements were also observed compared to MCS. Hornet's ability to maintain continuous-energy treatment on GPU hardware was demonstrated with precise modeling of bound and free thermal scattering effects.

Conclusions

Hornet enables robust and efficient GPU-based continuous-energy nuclear data processing, producing high-precision results on par with CPU-based Monte Carlo codes. Its object-oriented CUDA C++ architecture ensures maintainability and adaptability. Tested extensively with GREAPMC, Hornet proves to be a powerful tool for modern nuclear analysis and reactor design. Planned enhancements such as support for photo-atomic ACE file processing will broaden its applicability further.

Keywords

Monte Carlo, Particle transport, HORNET, High-performance computing, GPU.

Corresponding author: Deokjung Lee (deokjung@unist.ac.kr)

Author roles: Ali MR: Software; Aygul MS: Software; Lee D: Software, Supervision

Competing interests: No competing interests were disclosed.

Grant information: This work was supported by an Innovative Small Modular Reactor Development Agency grant (RS-2023-00265742) funded by the Korean government.

The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Copyright: © 2026 Ali MR *et al.* This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

How to cite this article: Ali MR, Aygul MS and Lee D. **Development and verification of Hornet – A CUDA-optimized, Monte Carlo-agnostic, continuous-energy cross-section processing code [version 2; peer review: 3 approved]** Nuclear Science and Technology Open Research 2026, 3:34 <https://doi.org/10.12688/nuclscitechnolopenres.17676.2>

First published: 28 Jul 2025, 3:34 <https://doi.org/10.12688/nuclscitechnolopenres.17676.1>

REVISED Amendments from Version 1

This revised version of the manuscript includes additional verification and performance analyses in response to reviewer feedback. In particular, code-to-code comparisons have been extended to examine statistical convergence with increasing particle histories, including RMS-based metrics consistent with expected Monte Carlo behavior. Clarifications have been added regarding performance measurement methodology, included physics treatments, tally definitions, and runtime components. Minor revisions were also made to improve clarity and presentation, while the overall scope and conclusions of the work remain unchanged.

Any further responses from the reviewers can be found at the end of the article

1. Introduction

During the random walk in a nuclear reactor core, the particle energy, direction, and position change. The neutron transport equation governs the behavior of neutrons in the nuclear reactor core. The numerical solutions of the transport equation comprise deterministic methods. They usually require discretization of a particle's energy, position, and direction. In addition, some of these methods simplify the problem geometry. All of these approximations aim to reduce the complexity of the problem, memory utilization, and computational burden. In contrast, the Monte Carlo (MC) method treats complex geometries with minimal assumptions. Moreover, the continuous treatment of energy, direction, and position reduces the errors introduced by discretization. Unlike deterministic methods, the MC method directly simulates the actual particle behavior. However, the results are prone to statistical errors because of the stochastic nature of the method. A large number of particles must be simulated to reduce the statistical error and achieve accurate results. Despite this, the MC method is widely used across various scientific fields because it is general-purpose and highly accurate, making minimal assumptions about the particle behavior. It is particularly well-suited for the computationally intensive task of simulating the random walk of neutrons in a nuclear reactor core.

Continuous energy cross-sections are fundamental to the Monte Carlo (MC) method in reactor physics. Microscopic cross sections are the interaction probability of a specific isotope with a neutron. A large cross-section indicates that the probability of interaction is large. There are many types of interactions (reactions) between isotopes and neutrons, such as fission, scattering, and absorption. An isotope-specific cross-section lookup is required to obtain microscopic cross-sections at a given energy. These cross-sections are stored in tabular form at discrete energy points, often numbering in thousands, particularly for isotopes with resonance regions where the cross-section changes rapidly. In such cases, finer energy bins were created. These energy bins or points constitute the *energy grid*. Each isotope has a unique temperature-dependent energy grid, and the number of energy points typically varies among the isotopes. The cross-section lookup operation involves locating the particle energy within this grid, followed by linear interpolation between adjacent cross-section values to obtain the cross-section at the energy of interest, as shown in [Figure 1](#). The required microscopic cross-section was calculated using [Equation \(1\)](#).

$$\sigma_n = \frac{E_n - E_i}{E_{i+1} - E_i} (\sigma_{i+1} - \sigma_i) + \sigma_i, \quad (1)$$

where σ_n is the cross-section at energy E_n . This process is repeated for each of the reaction types that are important to the problem (e.g., total, fission, elastic, and absorption).

MC codes rely on macroscopic cross-sections to simulate particle interactions within a medium during random walking. These interactions encompass particle absorption, elastic and inelastic scattering, fission reactions, and escape from a medium. The macroscopic cross-section depends on both the properties of the medium, such as the composition and temperature, and the properties of the particles, such as the energy. The macroscopic cross sections of a material are calculated as the sum of the products of the microscopic cross sections and number densities of the constituent isotopes, as shown in [Equation \(2\)](#).




Figure 1. Linear interpolation is employed to calculate cross-section at the desired energy.

$$\sum_x = \sum_{k=1}^{n_isotopes} N_k \sigma_x, \quad (2)$$

where the material is assumed to have $n_isotopes$ number of total isotopes as a constituent of the material, and x is the reaction type. For instance, if UO_2 fuel has five isotopes (^{234}U , ^{235}U , ^{236}U , ^{238}U , ^{16}O), the total macroscopic cross-section is calculated by adding the product of the individual total microscopic cross-section and number density.

In Monte Carlo neutron transport simulations focused on criticality analysis in fission reactor systems, a substantial portion of the computational cost is associated with continuous-energy cross-section lookups.¹ In addition to the cross-section calculation, a continuous-energy cross-section treatment code samples the energy and angular distributions for reactions that produce secondary neutrons, such as fission, inelastic, and (n, xn) reactions. Furthermore, the code must handle the analytical calculation of the outgoing energy in the case of elastic scattering reactions.

CPU-based computer clusters are often utilized to reduce the simulation time of the Monte Carlo (MC) method while preserving the accuracy. However, these clusters have a scalability limit; beyond a certain number of nodes, the communication overhead between nodes can negatively impact the performance. By contrast, graphical processing units (GPUs) with thousands of cores offer an excellent alternative to traditional CPU systems.

The computer science community has been using GPUs for scientific computations. More recently, the reactor physics community has begun to explore the benefits of GPUs, achieving significant speedups in various applications.^{2,3} While research on GPU-accelerated Monte Carlo (MC) methods often emphasizes results, less attention has been paid to implementation details. To the best of our knowledge, prior reactor physics research has not specifically focused on the independent development of GPU-optimized continuous energy cross-section codes; instead, discussions typically occur within the broader context of the MC process. The CORE lab at UNIST addressed this by developing Hornet, a GPU-accelerated continuous energy cross-section code designed to support comprehensive GPU-based MC methodologies. This study focuses on the development and verification of Hornet, which can operate as a standalone code (for academic purposes) and provides an interface for integration with GPU-enabled MC codes. Hornets are fully object-oriented and are written in CUDA C++. Several programming models, including Kokkos,⁴ OCCA,⁵ RAJA,⁶ OpenMP,⁷ OpenACC,⁸ SYCL,⁹ and OpenCL¹⁰ enable the use of NVIDIA GPUs. However, because CUDA was developed and supported by NVIDIA, it offers a level of control over implementation details and GPU support unmatched by these alternatives. Therefore, CUDA was selected as the platform for development. Furthermore, given that the code in this study was written from the ground up, a CPU-first and then a GPU-accelerated approach using directives was deemed inefficient. Finally, the CUDA API provides access to numerous optimized libraries such as cuFFT,¹¹ cuRAND,¹² cuSPARSE,¹³ and cuBLAS,¹⁴ which can be employed when required.

The rest of this paper is organized as follows. Section 2 reviews the existing literature on GPU-enabled continuous energy treatment. Section 3 details the Hornet implementation. Section 4 presents a thorough verification of the Hornet. Finally, Section 5 concludes the paper with a summary of the key findings and potential future research directions.

2. Background

The MC code ubiquitously fetches continuous energy cross-section data from (A Compact ENDF) file format.¹⁵ ACE data originate from the data given in the ENDF-6 format.¹⁶ However, the ENDF (Evaluated Nuclear Data Files) data cannot be directly used in the MC code. Hence, NJOY¹⁷ or FRENZY¹⁸ was employed to process the data and convert it into the ACE format.

Open-source nuclear data libraries (NDLs) exist, such as ACEtk¹⁹ and Papillon,²⁰ or are integrated into open-source Monte Carlo codes such as OpenMC²¹ and SCONE,²² which are often written in FORTRAN or rely heavily on programming features incompatible with CUDA. More importantly, they do not leverage support for CUDA and are not designed with GPU-optimized strategies. The following is a concise overview of the key components of an NDL.

2.1 Energy grid lookup

The pros and cons of different methods devised to handle look-ups in the context of GPU are briefly discussed here.

2.1.1 Binary search

A fundamental approach for locating bounding energy bins is a binary search²³ of the energy grid. This method has a time complexity of $O(\log n)$, where n represents the size of the energy-grid array.²³ Binary search is advantageous owing to its

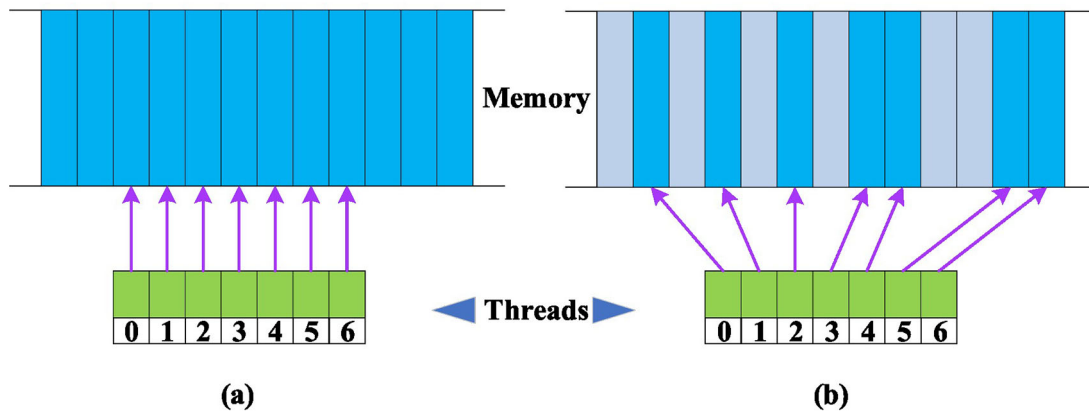


Figure 2. Example of memory access (a) coalesced and (b) strided.

simplicity and minimal memory overheads. However, performing binary searches on large energy grid arrays negatively impacts GPU performance owing to strided memory accesses. Strided accesses, where data are not stored contiguously in memory, can significantly increase the read times.

Figure 2 illustrates the difference between coalesced and strided memory accesses. The resulting poor cache utilization and limited global memory throughput make binary searches an inefficient choice for GPUs. GUARDYAN, a GPU-based Monte Carlo code for transient analysis, employs a binary search for energy grid lookup.²⁴

2.1.2 Unionized grid method

The unionized grid method²⁵ was first introduced in Serpent MC code.²⁶ As the name suggests, this method creates a union of the energy grid points of all nuclides in the problem. The individual cross-sections for all nuclides were then stored over this unionized energy grid. A single binary search to determine the bounding energies is sufficient for all isotopes.²⁵ This method has the major drawback of large memory requirements. Because this method stores a value for each energy grid point, the matrix of the cross-sections will mostly be sparse. This situation is further aggravated when the problem demands the storage of cross sections at multiple temperatures.²⁷

2.1.3 Double index scheme

An improvement over the unionized energy-grid method resulted in the development of a double-indexing method.²⁵ Instead of storing the cross-sections directly on the unionized grid, the pointers to the individual energy grids of the nuclides are stored on the unionized grid in the form of a double index table. Once the unionized grid point is located via a binary search, the local grid point of each nuclide is determined by examining the corresponding entry of the double index table. The double-index table reduces the memory usage of the unionized grid method; however, the size of the table is the product of points in the unionized grid and the number of problem nuclides. This will particularly explode depletion problems.

2.1.4 Fractional cascading

To reduce the memory requirements while achieving the computational performance of the unionized grid method, Lund et al.²⁸ proposed a cascade grid-based energy lookup. This method increases the memory by only a factor of two compared to the individual energy grids. The cascade grid consisted of the original nuclide energy grids, an additional set of energy grids (called augmented grids), one for each original energy grid, and a couple of pointers for each entry in the augmented grids. The augmented grids are created by constructing a unified grid for the first and second isotopes, followed by another grid for the second and third isotopes. The first pointer stored the energy index in the original nuclide grid, whereas the second pointer provided the energy index in the next augmented grid. The performance indicates a speedup of approximately 1.2 to 1.5 against using individual nuclide grids alone.²⁸ The fractional cascading technique significantly reduces memory usage compared to the global unionized method. However, this method lacks memory coalescence, as the second pointer always points to the next augmented grid.

2.1.5 Hash-based energy lookup

A hash function is an algorithm that takes the input and produces a hash value. This value uniquely represents the input data.²⁹ The purpose of a hash function is to efficiently map large amounts of data to smaller fixed-size values, making it easier to store, retrieve, and compare data.

This method employs a hash-based approach to expedite the energy grid search by mapping neutron energy to the indices of the hash grid.^{30,31} The hash grid indices were then mapped to the energy points on the nuclide energy grid. By narrowing down the search space to a subset of energies, the algorithm can efficiently locate relevant cross-sectional data.³¹ This logarithmic mapping technique divides the energy range into equal lethargy intervals. Walsh et al.¹ outlined the algorithm well, and the steps were as follows for a single nuclide.

1. First, the equal-logarithmic energy spacing was determined using Equation (3). The minimum and maximum energies are predetermined for the problem and supplied to the function that builds the hash function.

$$\Delta u = \frac{\log\left(\frac{E_{\max}}{E_{\min}}\right)}{N_{\text{Bins}}}. \quad (3)$$

2. Next, the indices of an individual nuclide's energy grid are determined using Equation (4) and stored as an array (hash table).

$$E_{N_{\text{bins}}-i+1}^H = \frac{E_{\max}}{\exp(i\Delta u)}. \text{ where } i = N_{\text{bins}}, N_{\text{bins}} - 1, \dots, 1, 0 \quad (4)$$

3. During the simulation, suppose the neutron energy at which the cross-sections are required is E_n . Then, the indices in the hash table are determined directly using Equation (5).

$$k(E_n) = N_{\text{bins}} - \left\lfloor \frac{\log\left(\frac{E_{\max}}{E_n}\right)}{\Delta u} \right\rfloor. \quad (5)$$

4. Now, the binary search is performed on $[E(k(E_n)) \ E(k(E_n) + 1)]$.

The original method³¹ builds a hash table on equi-lethargy intervals. Wang et al. compared the energy lookup methods for CPU and Many Integrated Core (MIC) architectures and proposed a new method, namely the N-ary map for Single Instruction Multiple Data (SIMD) architectures, where the hash table uses an equal number of energy intervals.³² Chen et al. proposed two optimization strategies for existing energy lookup algorithms.³³ These are the neighboring material cascade grid (NMCG) and Adaptive Optimal Logarithmic Grid (AOLG). AOLG can be useful for SIMD architecture, as for some isotopes, the resonance region can have hash bins with thousands of energy points. Consequently, an unbalanced distribution of the energy bins can cause thread divergence.

Raffuzzi et al. proposed accelerating the MC code by approximating the thermal cross-sections with functional forms.³⁴ In the thermal region, the cross-section variation is smooth; hence, the thermal region can be represented using functions. However, because this method directly replaces cross-sectional data with an approximation, the method can affect the accuracy of the results.³⁴

2.2 GPU-specific continuous-energy treatments

This subsection discusses continuous-energy cross-section treatments in state-of-the-art Monte Carlo codes. From a literature point of view, none of these codes uses a Monte Carlo agnostic approach to the continuous-energy cross-section code.

2.2.1 WARP

WARP is among the pioneering continuous-energy codes specifically designed for GPUs.³⁵ WARP uses PyNE (Python for Nuclear Engineering)³⁶ to read all cross-section data and optimize data storage by unionizing the energy grids of the

problem nuclides and employing CUDA's *float4* to store microscopic cross-sections. The *float4* construct maximizes bandwidth as more data are loaded in one memory transaction. Additionally, in WARP, the total cross-sections of all the isotopes are stored together within the unionized energy grid to increase coalescence. The WARP also devised a storage strategy to tabulate the outgoing angle and incident neutron energy for secondary angle-energy distributions. This strategy creates two additional matrices, each with dimensions matching the unionized cross-section matrix. These matrices contain pointers for the corresponding distribution data of a specific reaction. Two distinct matrices are necessary: one for scattering distributions, and the other for energy distributions.

2.2.2 Shift

Shift was the first continuous-energy Monte Carlo code optimized for running on both CPU and GPU architectures. Both versions used the equi-lethargy hash-based energy look-up method with 16,384 bins.³⁷ The benefit of a hash-based energy look-up is a smaller memory footprint compared to the other methods. Although the requirement to search for individual hash functions can reduce performance, the reduced search space can compensate for this performance loss.

2.2.3 PRAGMA

PRAGMA is a GPU-specific MC code developed and maintained by Seoul National University.³⁸ Similar to WARP, PRAGMA utilizes the *float4* construct to store primary microscopic cross sections and applies only inelastic scattering to the $S(\alpha, \beta)$ treatment.³⁹ Additionally, the literature does not mention any treatment for the unresolved resonance region in PRAGMA. For energy lookup, PRAGMA employs a linear-interval hashing scheme.³⁸

2.2.3.1 Linear Interval Hashing Scheme

This method is a spin-off of a double-index scheme. This variant of the double-index scheme increases memory coalescence and reduces the memory burden of the conventional scheme. In this method, every N^{th} value from the unionized grid was stored. The pointers corresponding to the individual energy grids do not point to each energy value. Therefore, a linear search was adopted for the local energy grid after a lower grid point was identified from the unionized grid. The value N is called a hash, and the purpose of using hashed storage is to reduce the memory requirement. Additionally, fewer points in the unionized grid could potentially increase the memory coalescence.

2.2.4 OpenMC

Two versions of GPU-enabled OpenMC exist: one leverages OpenMP⁷ to offload the code to the GPU,^{40,41} whereas the other uses CUDA C++.⁴² OpenMP eliminates polymorphisms in classes within the nuclear data hierarchy. By contrast, the CUDA version maintains polymorphism by creating objects directly on the device and implementing a polymorphism-supporting C++ unique pointer. In addition, the CUDA version uses single-precision nuclear data rather than double-precision data.⁴²

2.2.5 Mercury

Lawrence Livermore National Laboratories (LLNL) is actively working on porting their codes Mercury and Imp to GPUs.^{43–45} Mercury uses GIDI,⁴⁶ a nuclear data library developed and maintained by LLNL, and uses the Generalized Nuclear Database Structure (GNDS) format.⁴⁷ GIDI uses virtual functions; however, virtual function objects do not copy correctly from the CPU to the GPU using the managed memory. Hence, serialized functions are introduced along with overloading the new operator to copy data to the GPU.⁴³ Hornet is purpose-built for GPU architectures, whereas GIDI is primarily designed for CPUs and employs serialization when adapted for GPU execution. The novelty of Hornet therefore lies in its from-scratch GPU-centric design, while remaining agnostic to the specific Monte Carlo implementation.

3. Methods and features of Hornet

Hornet was developed with two main objectives: the primary objective was to support the transition of UNIST's in-house GPU-optimized REActor Physics Monte Carlo (GREAPMC) code from multigroup to continuous-energy treatment.² The secondary objective was to use the code for academic purposes, allowing students to engage in continuous-energy neutron physics using GPUs.

3.1 Salient features

3.1.1 Energy look-up

We implemented four energy look-up schemes in Hornet: unionized grid, double index scheme, equi-lethargy hash-based scheme, and linear interval hashing scheme. Because the purpose of this work is verification, the equi-lethargy hash-based algorithm devised by Brown³¹ was used as the default option. The hash function is built on an individual energy grid. As recommended by Brown,³¹ we employed 8000 equal-lethargy intervals. All objects are formed in the host and then transferred to the device; hence, the hash function is built in the constructor of the top-level class.

3.1.2 Thermal scattering

The impact of the chemical binding of the target molecules is strong for thermal neutrons, especially for energies below 4 eV. Because such scattering kinematics are more profound in moderating materials, it is important to account for these effects to obtain accurate results. For a thermal neutron, the de Broglie wavelength of the neutron and the interatomic spacing of the material coincide.⁴⁸ This leads to neutron interactions with the lattice structure of the material, in addition to the target nuclei. The thermal scattering laws (TSL) in the ENDF considers these additional effects and tabulates the cross sections, outgoing energy distribution, and angular distributions for thermal scattering with such materials. While the theory and derivation of the TSL are out of scope, the reader is referred to the articles by Tang et al.⁴⁹ and Fleming et al.⁵⁰ or, for a very detailed theory and derivations, to the book by Squires.⁵¹ The $S(\alpha, \beta)$ table contains three types of thermal-scattering reactions.⁵²

1. Inelastic scattering
2. Coherent elastic scattering
3. Incoherent elastic scattering

In addition to the elastic and inelastic cross sections, $S(\alpha, \beta)$ also contains information about the energy-angle distributions. It is important to distinguish between the elastic/inelastic terminology used in this study and the conventional elastic/inelastic concept. The elastic and inelastic terms in the context of $S(\alpha, \beta)$ refer to the excitation of the scattering system, whereas conventionally they refer to the excitation of the target nuclei. Coherent elastic scattering occurs in crystalline structures, incoherent elastic scattering reactions occur in hydrogen-rich solids, and inelastic scattering is important for other materials, including hydrogen in water. Although only inelastic scattering treatment is important for LWRs, all three scattering reactions are implemented in Hornet for future applications in other reactors.

3.1.3 Unresolved resonance treatment

The energy range of the cross-sections is divided into distinct regions: thermal, resolved resonance, unresolved resonance, and fast (continuum) regions. The resonance peaks in the unresolved resonance region (URR) are closely spaced such that they become indistinguishable. Probability tables (PTs), which represent cross sections using average resonance parameters, were employed to address this issue. The primary criterion for PTs is the preservation of energy self-shielding.⁵³ Nuclear data processing codes, such as NJOY, generate data for URR using unresolved resonance parameters. These data are used to create cumulative distribution functions (CDFs), which are then translated into tabular cross-sections dependent on the incident energy.⁵⁴ This method is almost ubiquitous for the probabilistic representation of the URR region in MC codes, and the same is adopted in Hornet. Details can be found in Levitt's and Sutton et al.'s works.^{55,56} This section provides a brief overview of the study.

Starting with the ACE file, the range of the cross-sectional values is divided into bands. Each band contained cross-sectional values or factors for a single probability. A probability table was created for each incoming energy source. The probability table holds the CDFs and bands. Hence, this is a three-tier hierarchy, where the lowermost level represents the bands. The class for the probability table forms the medium level, and the top level is where the incoming energy and corresponding table are stored. Figure 3 illustrates a unified modeling language (UML) class diagram for URR in Hornet.

3.1.4 Secondary angle distributions

The fission, inelastic, and (n, xn) reactions involve secondary neutrons. Hence, it was necessary to sample the secondary angles and energy distributions. An ACE file determines whether the angle and energy must be specified separately or as a correlated angle-energy distribution. In cases in which they are uncorrelated, the secondary angle distribution is sampled

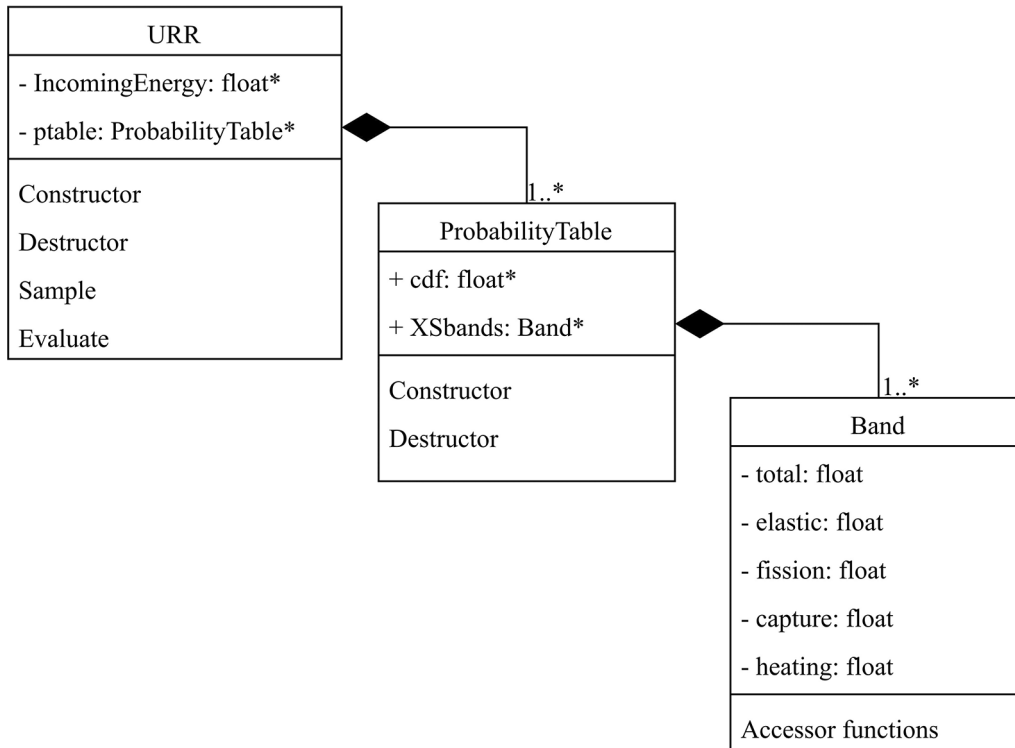


Figure 3. UML class diagram for URR treatment in Hornet.

independently. Similarly, the outgoing energy is calculated analytically for elastic scattering; therefore, only the secondary angle distribution must be specified. In these scenarios, the angle distribution is represented in one of the following formats:

1. Isotropic distribution
2. Equiprobable distribution
3. Tabular distribution

3.1.5 Outgoing energy and angle

Whenever such a reaction generates secondary neutrons, it is mandatory to determine the outgoing neutron energy and angle. In addition to elastic scattering, the outgoing neutron energy for other reactions is determined in one of the following ways. These methods are known as ACE laws¹⁵ and are based on ENDF format.¹⁶ For a specific isotope, information regarding the selection of the secondary energy distribution is provided in the ACE file. The ACE laws implemented in Hornet are as follows.

1. Equiprobable Energy Bins (ACE Law 1)
2. Inelastic Level Scattering (ACE Law 3)
3. Continuous Tabular (ACE Law 4)
4. Maxwell Spectrum (ACE Law 7)
5. Evaporation Spectrum (ACE Law 9)
6. Energy-dependent Watt Spectrum (ACE Law 11)

7. Kalbach-Mann law (ACE Law 44)
8. Correlated Energy and Angle (ACE Law 61)
9. N-Body Phase Space (ACE Law 66)

3.2 Implementation details

Hornet was developed using an object-oriented programming (OOP) approach to simplify maintenance and debugging. The inherent modularity of the continuous-energy cross-section processing lends itself well to the OOP. For example, each secondary distribution is implemented as a distinct class based on its function. These classes are linked through associations, thereby creating a hierarchical dependency structure. Consequently, adapting OOP-based code for GPU execution presents a significant challenge.

3.2.1 Memory management

The primary challenges in porting Hornet to GPU are memory allocation, data transfer, and deallocation. These three CUDA functions form the basis of memory management in CUDA C++.

1. `cudaMalloc()` : Allocates the device memory and assigns the address of the data to the device pointer.
2. `cudaMemcpy()` : Copies the data between host and device. This study makes four arguments.
 - a. Destination memory address
 - b. Host memory address
 - c. Size in bytes to copy
 - d. Type of transfer
3. `cudaFree()` : Frees the memory space indicated by the device pointer. This device pointer must have been returned by call to `cudaMalloc()`.

If all data transfers are centralized in one location, such as in `main()`, the device pointer is consistently available, making memory transfer and deallocation operations straightforward. However, when there is a tree-like dependency structure between classes, managing the device memory becomes cumbersome. Hornets feature a hierarchical class structure, with each isotope having numerous root-level objects that must be handled. Therefore, Hornet employs a special memory management scheme in which each class is equipped with two functions: `transfer_to_device()` and `free_device_memory()`. As the names suggest, these functions are responsible for transferring data to the device and freeing its memory. Breaking down the working logic into steps, starting from a top-tier class to a lower-most class for a mock-up example, will help grasp the memory management in Hornet.

3.2.1.1 Transferring data to device

The `transfer_to_device()` function accepts an argument, an object of the same class, passed by the reference. This object is used to allocate the device memory, and the data from the host object (that invoked the function) are copied to the object passed as an argument. Suppose there are two classes. The `main()` function is at the top of the hierarchy, as shown in [Figure 4](#). From within the `main()`, an object of the `MaxwellEnergy` class `h_ME` is instantiated. The object exists on the host side. A pointer `d_ME`, which acts as a device pointer, is also created. The data to which the device pointer will point are carried in the object `dh_ME`. The function responsible for allocating the device data and transferring the host object's data to the members of `dh_ME` is invoked by object `h_ME`, and object `dh_ME` is passed as an argument. Subsequently, the program launches a kernel, performs some processing, and frees the memory using the object `d_ME`.

```

int main(void)
{
    MaxwellEnergy h_ME(/* Pass ACE File */);
    MaxwellEnergy* d_ME = nullptr;
    cudaMalloc(&d_ME, sizeof(MaxwellEnergy));
    MaxwellEnergy dh_ME;
    // Transfer data from h_ME to dh_ME
    h_ME.transfer_to_device(dh_ME);
    // Assign an address to d_ME
    cudaMemcpy(d_ME, &dh_ME, sizeof(MaxwellEnergy), cudaMemcpyHostToDevice);
    /** Launch CUDA Kernel ***/
    // Free the device memory
    dh_ME.free_device_memory();
    cudaFree(d_ME);
    return 0;
}

```

Figure 4. The `main()` function of the example.

```

class MaxwellEnergy
{
private:
    Tabulated* theta {nullptr};

public:
    MaxwellEnergy() = default;
    MaxwellEnergy(/*Takes ACE file as an argument and allocates theta*/);
    void transfer_to_device(MaxwellEnergy& data);
    void free_device_memory();
};

```

Figure 5. Header file of the `MaxwellEnergy` class.

The `MaxwellEnergy` class forms the 2nd level of the hierarchy. [Figure 5](#) shows the header file of the `MaxwellEnergy` class. This class has only one data member, which is a pointer of the `tabulated` class. The object of the `tabulated` class was dynamically created within the constructor of the `MaxwellEnergy` class.

[Figure 6](#) illustrates the implementation of `MaxwellEnergy` class functions. Within the scope of `transfer_to_device()`, an object of the `tabulated` class is created (`dh1`). The `transfer_to_device()` function of the `tabulated` class was invoked using the host-side `tabulated` object (`theta`). Object `dh1` contains the device data of the `tabulated` class. This device memory was assigned to `data.theta` using `cudaMemcpy()`. After the scope of the function ends, the `dh1` object is destroyed, but the device data are stored in `data.theta`.

```

void MaxwellEnergy::transfer_to_device(MaxwellEnergy& data)
{
    Tabulated dh1;
    cudaMalloc(&data.theta, sizeof(Tabulated));
    theta -> transfer_to_device(dh1);
    cudaMemcpy(data.theta, &dh1, sizeof(Tabulated), cudaMemcpyHostToDevice);
}
void MaxwellEnergy::free_device_memory()
{
    Tabulated t1;
    cudaMemcpy(&t1, theta, sizeof(Tabulated), cudaMemcpyDeviceToHost);
    t1.free_device_memory();
    cudaFree(theta);
}

```

Figure 6. Implementation of the `MaxwellEnergy` class.

3.2.1.2 Freeing the device memory

The deallocation of device memory is impossible without a pointer that stores the address of the device data. To circumvent the issue arising when the device pointer needed to free the memory is unavailable, we used `cudaMemcpy()` to copy the device pointer to the host memory. It is pertinent to mention that `cudaFree()` always requires a pointer to reside on the host. In Figure 6, `theta` is currently residing on the device; therefore, it cannot be directly used to free the tabulated object's device memory. To free the tabulated object inside the `MaxwellEnergy` object, we first copy the object from the device to the host using `cudaMemcpy()`, as shown in the `free_device_memory()` function in Figure 6. The `free_device_memory()` of the tabulated class is then invoked using the host object `t1`. The `Tabulated::free_device_memory()`, in turn, uses `cudaFree()` on the tabulated data members.

Thus, the `main()` function remains clutter-free, and all data management is transparent to the `main()`. In the current framework, GREAPMC, through an interface class, invokes only the `transfer_to_device()` function of the top-level class, and the remaining object creation, memory allocations, and data transfers are automatically completed. Similarly, to free the device memory, GREAPMC simply invokes the `free_device_memory()` for each of the objects created at the start of the simulation and the device memory deallocation, and the destruction of the objects on the device is transparently completed. We have checked the method for memory leakage using `compute-sanitizer`⁵⁷ and no memory leaks or segmentation errors were detected, asserting the correctness of the method.

3.2.2 Optimizations

The general requirements for the memory layout and access patterns of GPUs fundamentally differ from those of CPUs. Hence, in this context, the optimizations ubiquitously used for continuous energy cross-section treatment are given here.

3.2.2.1 Primary cross-sections

Similar to WARP³⁵ and PRAGMA,³⁸ the four primary microscopic cross-sections (total, elastic, absorption, and fission) are stored together in the `float4` structure for each isotope in the problem.

`Float4` is a structure that stores four floating-point (float) values packed together, enabling efficient memory access and vectorized operations on GPUs. This data structure is optimized for GPU architectures, allowing a single instruction, such as an `LD.128-bit` read, to load all four values into the memory simultaneously.

The size of the `float4` arrays corresponds to the length of the energy grid, with each structure storing four values per energy-grid index. Once the energy grid index is determined, a single read simultaneously retrieves all four values. Normally, four values will incur four memory reads, but CUDA issues one load instruction when the data type is `float4` to retrieve four values, as shown in Figure 7. This is a crucial optimization because cross sections are frequently calculated within the transport loop. Given that global memory reads can become a bottleneck, special attention is paid to optimizing this access pattern.




Figure 7. Schematic diagram of `float4` structure.

3.2.2.2 Single precision arithmetic

The uncertainties in nuclear data exceed the errors resulting from reduced significant figures when single-precision arithmetic is used instead of double precision. Therefore, Hornet uses single-precision across nuclear data, with the exception of calculating the angle distributions in elastic scattering. For the outgoing angle calculations, it was observed that a single precision led to increased errors in the eigenvalue and pin power calculations. Single precision not only enhances cache hit rates but also reduces memory requirements by half.

4. Results and verification

This section presents the numerical results, verification, and performance evaluation of Hornet in comparison with MCS. To generate verification results, Hornet was coupled with GREAPMC. These results were obtained using standard history-based neutron tracking and an equi-lethargy-based hash method for energy lookup. The ENDF/B-VII.1 library was used, applying linear interpolation for temperature when the material temperature fell between the two available library temperatures. The verification includes the sampling of selected reactions, $S(\alpha, \beta)$ treatment, resonance scattering, and pin-by-pin power results for a fuel assembly case. Except for reaction sampling, the number of inactive cycles for all other cases was determined based on fission source convergence evaluated using Shannon entropy. In view of the primary purpose of Hornet, current verification efforts include Hornet-specific tests based on sampling and comparison of selected reactions, as well as assessment of the coupled Hornet-GREAPMC system. The multigroup GREAPMC and Hornet-specific elastic scattering treatment have been verified elsewhere.^{58,59} In transitioning from multigroup to continuous-energy treatment, microscopic cross sections for all isotopes in a material and the corresponding macroscopic cross sections are computed as an additional step in the transport loop.^{60,61}

4.1 Sampling of reactions

The three most important isotopes involved in nuclear fission are ²³⁸U, ²³⁵U, and ²³⁹Pu. We sampled the following two reactions with selective incoming energies and temperatures:

1. MT 91 (z, n_c)
2. MT 18 ($z, fission$)

Reactions were sampled with both MCS and GREAPMC using 1×10^6 samples. We employed Jensen–Shannon divergence (JSD)⁶² to assess the output’s quantitative closeness. The JSD score is a metric used to quantify the similarity between the two probability distributions. It combines information from the Kullback-Leibler divergence in both directions, providing a single measure of dissimilarity. The formula for Jensen-Shannon divergence is provided in Equation (6).

$$JSD(P\|Q) = \frac{K(P\|M) + K(Q\|M)}{2} \quad (6)$$

Here, K denotes the Kullback–Leibler divergence and M is the pointwise mean of P and Q . A smaller JSD value indicates a higher degree of similarity between the two probability distributions.

Figure 8 illustrates the outgoing energy for the inelastic reaction MT 91 with ²³⁵U at a temperature of 293.6 K. The energy of the incoming neutrons in all the samples was 2.625 MeV. The plot shows good agreement between MCS and Hornet, with a JSD value of 0.1.

Figure 9 shows a plot of the outgoing energy for ²³⁸U MT 91 reaction with the same incoming neutron energy as that in the previous case. The temperature was then changed to 600 K. Good agreement between both codes is evident from the JSD value of 0.084.

Figure 10 shows the outgoing energy when a 2.625 MeV incoming neutron interacts inelastically with ²³⁹Pu at 900 K. The JSD value for this case was 0.102. The slight difference in reaction sampling may be due to the difference in precision. Hornet uses single precision throughout, except for angle calculation in elastic scattering, whereas MCS uses double precision throughout.

Figure 11 depicts the outgoing energy for the fission reaction with ²³⁵U at 293.6 K when a neutron of energy 0.523 MeV interacts. The JSD value for this reaction was 0.13.




Figure 8. MT-91 reaction with ^{235}U at 293.6 K for an incoming energy of 2.625 MeV.




Figure 9. MT-91 reaction with ^{238}U at 600 K for an incoming energy of 2.625 MeV.




Figure 10. MT-91 reaction with ^{239}Pu at 900 K for an incoming energy of 2.625 MeV.




Figure 11. MT-18 reaction with ^{235}U at 293.6 K for an incoming energy of 0.523 MeV.

Figure 12 illustrates the outgoing energy when MT 18 was sampled for a neutron with an energy of 1.811 MeV interacting with ^{238}U at 293.6 K. The JSD value for this reaction was 0.13, as in the previous case.

Figure 13 shows the outgoing energy when a neutron with an energy of 0.505 MeV interacts with ^{239}Pu at room temperature. The JSD value for this reaction is 0.13.




Figure 12. MT-18 reaction with ^{238}U at 293.6 K for an incoming energy of 1.811 MeV.




Figure 13. MT-18 reaction with ^{239}Pu at 293.6 K for an incoming energy of 0.505 MeV.




Figure 14. The cross-sectional view of the pin problem.

Overall, these results indicate a good agreement between Hornet and MCS. The JSD value was approximately 0.1, strengthening the closeness between the distributions. The JSD can take values between 0 and 1 with zero meaning that both the distributions are similar to each other. A value around 0.1 indicates some differences between the distributions that occurs due to difference in precision in MCS and Hornet. The JSD values for Hornet and MCS for elastic scattering treatment with same precision is around 10^{-5} .⁵⁸ Similar agreements are observed for other reactions; however, only fission reactions are included in this work as representative of the overall agreement.

4.2 Thermal system

To validate the $S(\alpha, \beta)$ treatment and overall eigenvalue, we considered IAEA’s International Nuclear Data Committee’s INDC (USA)-107 benchmark.⁶³ This is a simple benchmark that simulates a water-moderated UO_2 pin cell. The square geometry is reflective at all faces and contains only two materials: the fuel and moderator. Each material had two isotopes. The fuel contains ^{235}U and ^{238}U , while the moderator contains ^1H and ^{16}O . The benchmark consisted of three problems that differed according to the radius of the fuel rod. Each problem was run with two moderator conditions: free atoms and applying the thermal scattering law. The free-atom condition assumes thermal motion using free atoms, where the atomic motion follows a Maxwellian distribution. The bound state accounts for the atomic translational motion along with vibration and rotation. Figure 14 shows the pin-cell problem with the radius of each problem.

The temperature was isothermal (293.6 K), with a moderator density of 1 g/cc and a fuel density of 18.8 g/cc. The fuel composition is listed in Table 1.

The k_{eff} difference between GREAPMC and MCS is presented in Table 2. The simulations involved one million particles per cycle, with 100 active and 1100 inactive cycles. The k_{eff} difference was computed using Equation (7).

$$k_{\text{eff}} \text{ difference [pcm]} = \left(k_{\text{eff}}^S - k_{\text{eff}}^R \right) \times 10^5 \tag{7}$$

Here, S refers to the simulated result (using GREAPMC), and R refers to the reference value (obtained via MCS). The uncertainty in the reported k_{eff} difference for the thermal systems corresponds to the propagated one-standard-deviation uncertainty, obtained by combining the independent Monte Carlo uncertainties from the two codes.

Table 1. The atomic fractions of Uranium isotopes in each problem.

Problem	^{238}U (a/o)	^{235}U (a/o)
P1	99.02	0.98
P2	96.5	3.5
P3	30	70

Table 2. Eigenvalue difference for the three problems.

Problem	k_{eff} difference (pcm)	
	Bound	Free
P1	5.05 ± 2.07	-4.03 ± 1.92
P2	1.74 ± 2.68	-0.89 ± 2.58
P3	7.23 ± 3.31	-2.46 ± 3.28

The results show good agreement between GREAPMC and MCS for the bound and free cases. The difference is mainly due to GREAPMC's use of single precision for continuous energy cross-sections, geometry, and material treatment.

4.3 Mosteller Benchmark

The Mosteller benchmark problem^{64,65} was proposed to assess the resonance scattering effect of the Doppler broadening rejection correction (DBRC) against the constant cross-section (CXS) approximation at different fuel temperatures.⁶⁶ Specifically, the Doppler coefficient was calculated by considering two configurations: hot zero power (HZP) with fuel at 600 K and hot full power (HFP) with fuel at 900 K. All other materials were kept at a temperature of 600 K. The upper energy limits for DBRC were set at 210 eV in MCS and GREAPMC.⁶⁷ The benchmark consists of UO₂ fuel, weapons-grade mixed oxide (MOX) fuel, and reactor-recycled MOX fuel. The Doppler coefficient (or fuel temperature coefficient [FTC]) in pcm/K is calculated using Equation (8).

$$\text{Doppler Coefficient} = \frac{10^5}{\Delta T} \left(\frac{1}{k_{HZP}} - \frac{1}{k_{HFP}} \right) \quad (8)$$

The simulations employed three million particles per cycle in 70 inactive and 300 active cycles. The Doppler coefficient is calculated from the difference of multiplication factors evaluated at two fuel temperatures. The associated uncertainty is obtained using standard error propagation for independent quantities. Specifically, the k_{eff} values at the two temperatures are estimated from independent Monte Carlo simulations with separate random number sequences, and are therefore assumed to be statistically uncorrelated. Under this assumption, the standard deviation of the Doppler coefficient is computed by propagating the one-standard-deviation uncertainties of the individual k_{eff} values. Table 3 lists the Doppler coefficients for the UOX fuel. The Uranium enrichment varies from natural uranium to 5%. The maximum relative error between the Doppler coefficients calculated by MCS and GREAPMC is 0.86% for CXS for a 5% enrichment, and 0.65% for DBRC for a 3.9% enrichment case.

Figure 15 illustrates the Doppler coefficients for CXS and DBRC using the MCS and GREAPMC (coupled with Hornet). The Doppler coefficient values for all methods lie within the standard deviation of each code. Additionally, there was no bias in the coefficient values.

Table 3. Doppler coefficients for UO₂ fuel and different Uranium enrichments. Uncertainties are expressed in pcm.

wt. (%)	CXS				DBRC			
	MCS		GREAPMC		MCS		GREAPMC	
	FTC	σ	FTC	σ	FTC	σ	FTC	σ
0.711	-4.640	0.015	-4.659	0.015	-5.128	0.014	-5.144	0.016
1.6	-3.026	0.010	-3.049	0.011	-3.401	0.010	-3.386	0.011
2.4	-2.560	0.009	-2.558	0.009	-2.845	0.009	-2.857	0.009
3.1	-2.339	0.008	-2.324	0.008	-2.609	0.008	-2.599	0.009
3.9	-2.152	0.007	-2.167	0.008	-2.410	0.008	-2.426	0.009
4.5	-2.064	0.007	-2.065	0.008	-2.309	0.007	-2.316	0.008
5.0	-2.017	0.007	-2.000	0.008	-2.261	0.008	-2.253	0.008

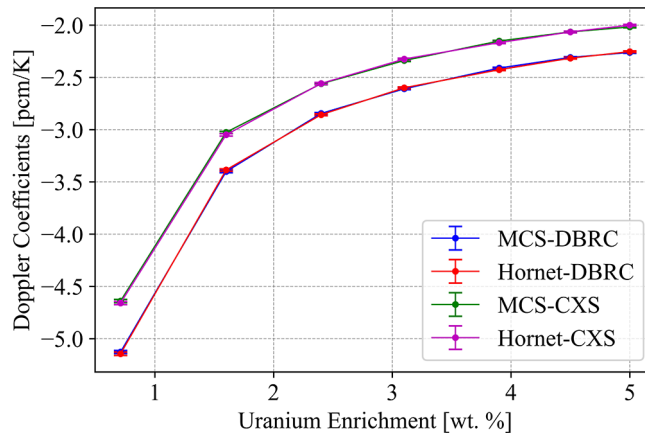


Figure 15. Doppler coefficients for UOX fuel with different Uranium enrichments.

Table 4. Doppler coefficients for reactor-recycle MOX fuel with different Plutonium enrichments. Uncertainties are expressed in pcm.

wt. (%)	CXS				DBRC			
	MCS		GREAPMC		MCS		GREAPMC	
	FTC	σ	FTC	σ	FTC	σ	FTC	σ
1	-3.610	0.011	-3.590	0.012	-3.979	0.012	-3.984	0.012
2	-3.462	0.012	-3.463	0.011	-3.812	0.012	-3.816	0.012
4	-3.294	0.010	-3.294	0.010	-3.611	0.011	-3.629	0.010
6	-3.160	0.009	-3.159	0.010	-3.464	0.010	-3.459	0.010
8	-3.041	0.009	-3.037	0.010	-3.318	0.009	-3.321	0.010

The Doppler coefficients for the reactor-recycle MOX fuel are listed in Table 4. The Plutonium enrichment varied from 1% to 8%. The maximum relative difference in the coefficient was 0.56% for CXS (1% enrichment) and 0.5% for DBRC (4% enrichment).

Figure 16 shows the Doppler coefficients for the reactor recycling case. The FTC values from one code lie within the standard deviation of the other codes, indicating excellent agreement. Moreover, there was no bias in the calculation of FTC.

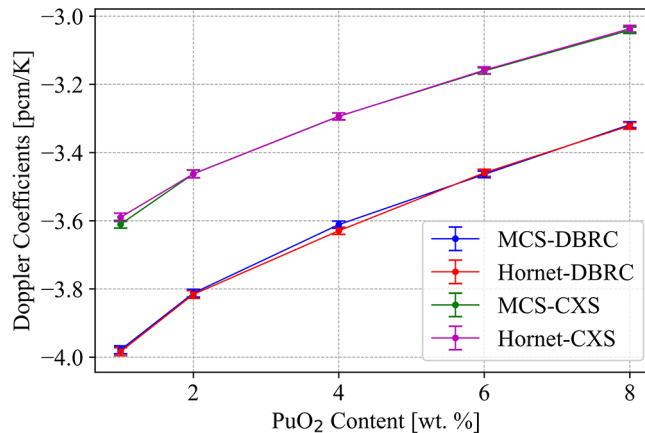


Figure 16. The Doppler coefficient for reactor-recycle MOX fuel steadily reduces in magnitude as Plutonium enrichment increases.

Table 5. Doppler coefficients for weapons-grade MOX fuel with different Plutonium enrichments. Uncertainties are expressed in pcm.

wt. (%)	CXS				DBRC			
	MCS		GREAPMC		MCS		GREAPMC	
	FTC	σ	FTC	σ	FTC	σ	FTC	σ
1	-2.553	0.009	-2.561	0.010	-2.883	0.009	-2.871	0.010
2	-2.633	0.008	-2.634	0.009	-2.906	0.008	-2.918	0.009
4	-2.611	0.008	-2.609	0.009	-2.883	0.008	-2.873	0.008
6	-2.519	0.008	-2.515	0.008	-2.744	0.007	-2.741	0.008

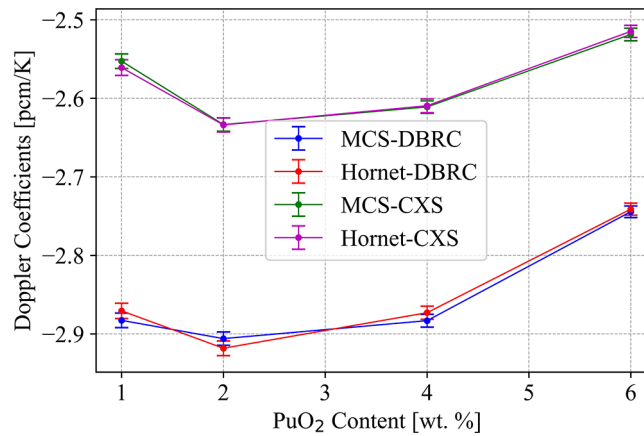


Figure 17. FTC for weapons-grade MOX fuel.

Table 5 gives the FTC for the weapons-grade MOX fuel for different plutonium enrichments. The FTC was the most negative for 2% plutonium enrichment. The maximum relative difference in CXS was 0.31% at 1% plutonium enrichment. Similarly, the maximum relative difference for the DBRC was 0.42% in the 1% enrichment case.

Figure 17 illustrates the Doppler coefficients for the weapons-grade MOX fuel at various plutonium enrichments. Both MCS and GREAPMC showed excellent agreement.

To enable a robust code-to-code verification, the reactivity difference for a representative Mosteller benchmark case with CXS approximation and 4% weapon-grade MOX fuel was evaluated by comparing results from MCS and GREAPMC. The results are given in Table 6. The comparison was performed by progressively increasing the number of particle histories to examine statistical convergence. For a consistent Monte Carlo comparison, the expectation is that the reactivity difference between the two codes decreases as the number of particle histories increases, reflecting reduced statistical uncertainty and convergence toward the same underlying solution.

As expected, the reactivity differences for both hot zero power (HZP) and hot full power (HFP) conditions decrease with increasing particle histories. At lower particle counts ($10^4 - 10^5$), the observed differences are dominated by statistical uncertainty, resulting in comparatively large error bars and fluctuating mean values. As the number of histories increases

Table 6. Comparison of reactivity differences between MCS and GREAPMC for a representative case.

Particles per cycle	HZP (pcm)	HFP (pcm)
1×10^4	31.35925 ± 30.46248	21.53892 ± 32.12405
1×10^5	-10.88241 ± 10.67002	-16.96994 ± 9.87741
1×10^6	-2.560161 ± 2.96611	10.44341 ± 3.24340
1×10^7	0.64003 ± 1.06602	-7.83188 ± 0.94655

to 10^6 and 10^7 , the statistical uncertainties are significantly reduced, and the reactivity differences converge toward zero within uncertainty. This behavior demonstrates that MCS and GREAPMC converge to statistically consistent solutions for both conditions when sufficient particle histories are employed. The observed convergence provides confidence in the consistency of the physical models, numerical implementations, and cross-section handling between the two codes for continuous-energy Monte Carlo simulations.

4.4 Pin-by-Pin power

OPR1000 is a commercial Korean reactor.⁶⁸ The peculiarity of the OPR1000 core is the guide tube and instrumentation tube. Although the fuel assemblies in most reactors have a uniform lattice, the OPR1000 fuel assemblies have a non-uniform lattice in the sense that four normal pin cells are combined to form one guide tube. Similarly, the instrumentation tube was larger than the normal pin cells.

The pin-by-pin power was quantitatively assessed by comparing the normalized pin powers in an OPR1000 fuel assembly (FA). Figure 18 shows the FA design loaded with 1.43% UO_2 . The radial boundary condition is reflective, whereas the axial boundary condition is vacuum. Both the MCS and GREAPMC simulations were performed with 100 inactive cycles, 500 active cycles, and two million particles per cycle. The difference in eigenvalue computed using Equation (7) is 3.9 ± 2.65 pcm. Figure 19 depicts the normalized radial pin power profile in quarter FA from both codes along with the relative difference, which is computed using Equation (9).

$$\text{Relative difference [\%]} = \left(1 - \frac{S}{R}\right) \times 100 \tag{9}$$

The root mean square (RMS) difference in the axially averaged power profile was less than 0.1%, and the maximum relative difference was less than 0.2%. Overall, the power profile showed excellent agreement, as indicated by the RMS.

The RMS deviation between MCS and GREAPMC pin power tallies was evaluated for increasing numbers of particles. The RMS deviation between GREAPMC and the reference MCS results was evaluated by Monte Carlo sampling. For each tally, values were randomly sampled from normal distributions defined by the reported mean and one-standard-deviation uncertainty. The RMS was computed for each realization, and the final RMS and its uncertainty were obtained from the mean and standard deviation of the resulting RMS distribution. The results are given in Table 7.

As the number of particle histories increases, the RMS deviation between the MCS and GREAPMC pin power tallies decreases systematically. At 10^6 and 10^7 particles per cycle, the RMS values are reduced by more than an order of magnitude compared to the lowest-history case, and the associated uncertainties become very small. This behavior

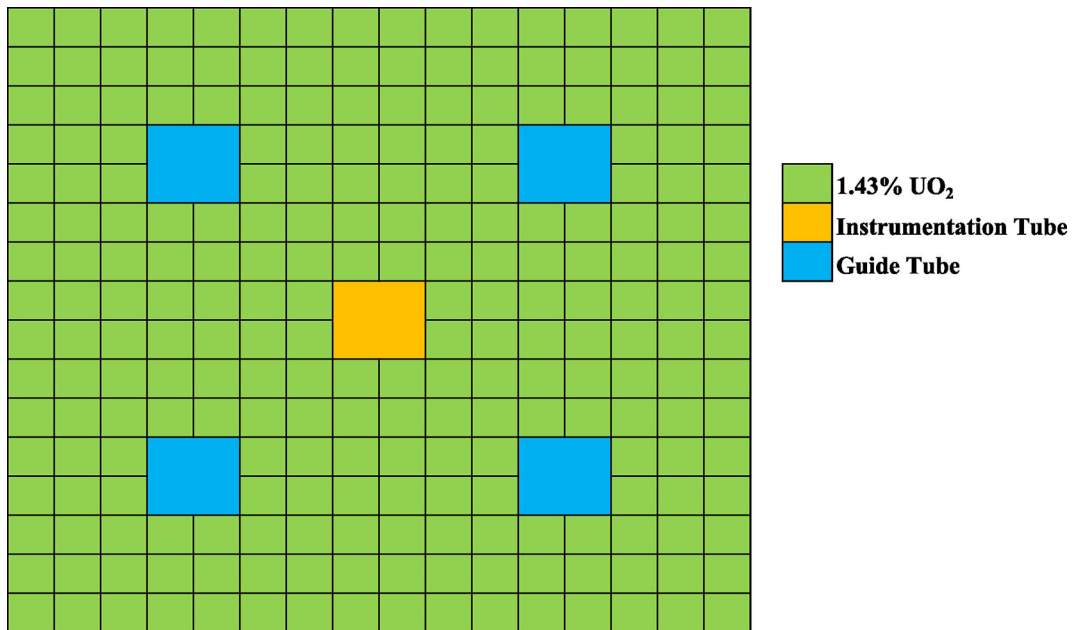


Figure 18. The 1.43% UO_2 FA design was used to assess the pin powers.

	1.0773 1.0780	1.0159 1.0151	0.9930 0.9932	0.9807 0.9807	0.9730 0.9729	0.9727 0.9725	0.9896 0.9896	
	-0.0615%	0.0829%	-0.0267%	-0.0013%	0.0040%	0.0142%	0.0026%	
1.0778	1.0359 1.0358	1.0159 1.0152	1.0114 1.0114	1.0033 1.0030	0.9878 0.9877	0.9816 0.9813	0.9952 0.9946	
0.0057%	0.0145%	0.0785%	0.0043%	0.0338%	0.0135%	0.0245%	0.0617%	
1.0152	1.0155 1.0156	1.0343 1.0335	1.0736 1.0723	1.0692 1.0686	1.0218 1.0216	0.9979 0.9973	1.0050 1.0051	
-0.0017%	-0.0185%	0.0762%	0.1197%	0.0516%	0.0157%	0.0544%	-0.0125%	
0.9935	1.0111 1.0118	1.0729 1.0726			1.0718 1.0715	1.0152 1.0148	1.0142 1.0131	
0.0314%	-0.0633%	0.0360%			0.0271%	0.0415%	0.1017%	
0.9810	1.0035 1.0033	1.0691 1.0686			1.0752 1.0747	1.0188 1.0186	1.0181 1.0176	
0.0179%	0.0183%	0.0485%			0.0462%	0.0219%	0.0505%	
0.9715	0.9874 0.9879	1.0216 1.0217	1.0717 1.0715	1.0750 1.0749	1.0326 1.0319	1.0103 1.0107	1.0185 1.0187	
-0.1270%	-0.0523%	-0.0112%	0.0236%	0.0084%	0.0654%	-0.0349%	-0.0174%	
0.9723	0.9813 0.9818	0.9975 0.9973	1.0149 1.0149	1.0185 1.0191	1.0102 1.0108	1.0077 1.0079	1.0225 1.0230	
0.0296%	-0.0501%	0.0164%	-0.0005%	-0.0527%	-0.0636%	-0.0213%	-0.0433%	
0.9890	0.9945 0.9950	1.0042 1.0045	1.0139 1.0136	1.0180 1.0180	1.0186 1.0185	1.0228 1.0227	1.0418 1.0417	
0.0210%	-0.0549%	-0.0371%	0.0336%	0.0023%	0.0063%	0.0098%	0.0124%	

MCS
GREAPMC
Rel. Diff. (%)

Rel. Diff. (%)

Min -0.1270%

Max 0.1197%

RMS 0.0455%

Figure 19. Normalized pin powers and related statistics in quarter OPR1000 FA.

Table 7. Comparison of reactivity differences between MCS and GREAPMC for a representative case.

Particles per cycle	RMS
1 × 10 ⁴	6.19053 × 10 ⁻³ ± 2.86752 × 10 ⁻⁴
1 × 10 ⁵	1.95782 × 10 ⁻³ ± 9.06601 × 10 ⁻⁵
1 × 10 ⁶	6.18872 × 10 ⁻⁴ ± 2.86574 × 10 ⁻⁵
1 × 10 ⁷	1.94867 × 10 ⁻⁴ ± 9.02601 × 10 ⁻⁶

indicates that the pin power distributions produced by MCS and GREAPMC converge toward statistically consistent results. Overall, the observed monotonic reduction in RMS deviation with increasing particle histories demonstrates excellent agreement between the two codes and confirms the consistency of their pin power tally implementations. These results provide strong evidence that GREAPMC accurately reproduces the reference MCS pin power distributions when sufficient particle histories are employed.

4.5 Performance

The performance of codebase was assessed for the OPR1000 FA case. Based on the codes’ average time to execute one active cycle, GREAPMC’s acceleration is compared with that of MCS. The number of particles per cycle varied from 1 million to 5 million. Constant cross-section approximations for elastic scattering and equi-lethargy-based energy lookup were employed for these simulations. Table 8 lists the hardware specifications. The MCS simulations utilized 144 CPU cores, whereas the remaining 16 cores within the four nodes were reserved for development and administrative tasks.

The performance was quantified using the equivalent number of CPU cores that yielded the same performance as that of one GPU. The parameter is given by Equation (10)⁶⁹:

$$\text{Equivalent Number of CPU Cores} = \frac{(\text{Number of CPU Cores}) \times (\text{CPU run time})}{\text{GPU run time}} \tag{10}$$

In terms of tracking rate, the equivalent number of CPU cores is specified by Equation (11).

Table 8. Hardware specifications for the performance assessment.

GPU	Model	NVIDIA GeForce RTX 3090
	Global Memory (GB)	23.7
CPU	Model	Intel Xeon Gold 6242R
	Base Clock (MHz)	3100
	Memory (GB)	754
	Number of cores per node	40
	Number of nodes used for simulation	04

$$\text{Equivalent Number of CPU Cores} = \frac{(\text{Number of CPU Cores}) \times (\text{GPU tracking rate})}{\text{CPU tracking rate}} \tag{11}$$

Table 9 presents the tracking rates for both codes. The tracking rate increases as the number of particles per cycle increases, with GREAPMC showing a more significant growth in the tracking rate than MCS. Consequently, the equivalent number of CPU cores increased accordingly.

Figure 20 illustrates the equivalent number of CPU cores. As the number of particles per cycle increases, the equivalent number of CPU cores stabilizes at approximately 440, indicating that a single GPU provides a performance comparable to that of approximately 440 Intel Xeon Gold 6242R CPU cores.

Table 9. Performance of GREAPMC compared to MCS for the fuel assembly case.

Particle per cycle (Million)	Tracking rate (kn/s)		Equivalent number of CPU cores
	MCS	GREAPMC	
1	218.20	642.64	324
2	219.83	665.56	435
3	221.07	670.54	437
4	221.78	677.73	440
5	222.22	680.33	441




Figure 20. Equivalent number of CPU cores for the OPR1000 fuel assembly.

Apart from the aforementioned integration of Hornet in GREAPMC, another possible execution model for Hornet is to generate cross sections on the GPU while the transport is performed on CPUs, with cross sections requested and supplied asynchronously, as explored in MC21.⁷⁰

5. Conclusions

The verification of Hornet against the well-established MCS code demonstrates its capability to deliver accurate and efficient solutions using GPU architecture. Hornet coupled with GREAPMC showed excellent agreement with MCS across a range of simulations. The implementation uses single-precision nuclear data, incorporates a variety of energy lookup methods, and adheres to an object-oriented programming paradigm, thereby ensuring code readability, maintainability, and adaptability. The results from benchmarks, such as the IAEA benchmark INDC (USA)-107 for a pin-cell model and the Mosteller benchmark for Doppler coefficient evaluations, confirm that Hornet achieves high precision across various nuclear conditions, including bound and free scattering scenarios. The pin-by-pin power verification in an OPR1000 fuel assembly further validated Hornet's performance with root-mean-square (RMS) power differences under 0.1%, highlighting its potential for accurate simulations. Combined, Hornet and GREAPMC significantly enhanced the tracking rate, surpassing that of the MCS. Overall, Hornet offers a robust solution for continuous-energy Monte Carlo neutron transport in GPUs, providing a valuable tool for nuclear analysis and reactor design. The future involves expanding Hornet's capabilities with photoatomic ACE file processing.

6. Declaration of generative AI in scientific writing

During the preparation of this work, the author(s) used ChatGPT-3 to improve readability, clarity, and conciseness of writing. After using this tool/service, the author(s) reviewed and edited the content as needed and took (s) full responsibility for the content of the publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have influenced the work reported in this study.

Credit authorship contribution statement

Muhammad Rizwan Ali: Main developer, methodology, formal analysis, writing – original draft. Murat Serdar Aygul: Programming, result accumulation, review. Deokjung Lee: Supervision, fund acquisition, and review.

Software availability

The codes used in this study (Hornet and GREAPMC) are proprietary, in-house software developed by the CORE Lab at UNIST and are not publicly available due to intellectual property restrictions and university policy. Researchers interested in collaborating or seeking access to the software may contact the corresponding author.

Data availability

Some of the data supporting the findings of this study were generated using in-house codes developed at the CORE Lab, UNIST, and are subject to intellectual property restrictions. Therefore, the datasets cannot be shared publicly. No Institutional Review Board (IRB) approval was required for data sharing. Interested researchers may contact the corresponding author to request access to the input files associated with the non-proprietary models (like Mosteller Benchmark); requests will be considered on a case-by-case basis and are subject to CORE Lab and university approval.

Acknowledgment

This work was supported by an Innovative Small Modular Reactor Development Agency grant (RS-2023-00265742) funded by the Korean government.

References

- Walsh JA, Romano PK, Forget B, et al.: **Optimizations of the Energy Grid Search Algorithm in Continuous-Energy Monte Carlo Particle Transport Codes.** *Comput. Phys. Commun.* Nov. 2015; **196**: 134–142. [Publisher Full Text](#)
- Ali MR, Aygul MS, Lee D: **Enhancing PWR Monte Carlo Simulations with GREAPMC: A GPU-Accelerated Approach.** *International Conference on Physics of Reactors (PHYSOR 2024)*. San Francisco, CA, USA: pp. 2174–2183. [Publisher Full Text](#)
- Dzianisau S, Lee D: **GPU Acceleration of 3D MOC Solver in STREAM3D Using OpenACC.** presented at the *Transactions of the Korean Nuclear Society Spring Meeting, Jeju, Korea.* May 2024.

4. Carter Edwards H, Trott CR, Sunderland D: **Kokkos: Enabling Manycore Performance Portability through Polymorphic Memory Access Patterns.** *J. Parallel Distrib. Comput.* Dec. 2014; **74**(12): 3202–3216. [Publisher Full Text](#)
5. Medina DS, St-Cyr A, Warburton T: **OCCA: A Unified Approach to Multi-Threading Languages.** *arXiv.* 2014; [abs/1403.0968](#).
6. Beckingsale DA, et al.: **RAJA: Portable Performance for Large-Scale Scientific Applications.** *Presented at the 2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC).* Nov. 2019; pp. 71–81. [Publisher Full Text](#)
7. OpenMP Architecture Review Board: **OpenMP Application Program Interface Version 5.0.** Accessed: Aug. 05, 2024. [Reference Source](#)
8. OpenACC-standard.org: **The OpenACC Application Programming Interface Version 3.1.** Accessed: Aug. 05, 2024. [Reference Source](#)
9. Reyes R, Brown G, Burns R, et al.: **SYCL 2020: More than meets the eye.** *Proceedings of the International Workshop on OpenCL in IWOCL'20.* New York, NY, USA: Association for Computing Machinery; 2020. [Publisher Full Text](#)
10. Stone JE, Gohara D, Shi G: **OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems.** *Comput. Sci. Eng.* Jun. 2010; **12**(3): 66–73. [PubMed Abstract](#) | [Publisher Full Text](#) | [Free Full Text](#)
11. NVIDIA: **NVIDIA cuFFT.** Accessed: Aug. 05, 2024. [Reference Source](#)
12. NVIDIA: **cuRAND - Random Number Generation on NVIDIA GPUs.** Accessed: Aug. 05, 2024. [Reference Source](#)
13. NVIDIA: **cuSPARSE - GPU library APIs for sparse computation.** Accessed: Aug. 05, 2024. [Reference Source](#)
14. NVIDIA: **cuBLAS - Basic Linear Algebra on NVIDIA GPUs.** Accessed: Aug. 05, 2024. [Reference Source](#)
15. Conlin JL, Romano P: **A Compact ENDF (ACE) Format Specification.** Los Alamos, NM, United States, LA-UR-19-29016: Los Alamos National Laboratory (LANL); Sep. 2019. [Publisher Full Text](#)
16. Brown DA: **ENDF-6 Formats Manual - Data Formats and Procedures for the Evaluated Nuclear Data Files ENDF/B-VI, ENDF/B-VII and ENDF/B-VIII.** Brookhaven National Laboratory, BNL-224854-2023-INRE; 2023.
17. Macfarlane R, Muir DW, Boicourt RM, et al.: **The NJOY Nuclear Data Processing System, Version 2016.** United States, LA-UR-17-20093: Jan. 2017. [Publisher Full Text](#)
18. Tada K, Nagaya Y, Kunieda S, et al.: **Development and verification of a new nuclear data processing system FREN DY.** *J. Nucl. Sci. Technol.* Jul. 2017; **54**(7): 806–817. [Publisher Full Text](#)
19. Los Alamos National Laboratory: **ACEtk.** [Reference Source](#)
20. Belanger H: **Papillon Nuclear Data Library – A Free and Open-source C++/Python Library for Interacting with ACE Files for Continuous-Energy Neutron Data.** *EPJ Nucl. Sci Technol.* 2023; **9**. [Publisher Full Text](#)
21. Romano PK, Horelik NE, Herman BR, et al.: **OpenMC: A State-of-the-art Monte Carlo Code for Research and Development.** *Ann. Nucl. Energy.* Aug. 2015; **82**: 90–97. [Publisher Full Text](#)
22. Kowalski MA, Cosgrove P, Broman J, et al.: **SCONE: A Student-Oriented Modifiable Monte Carlo Particle Transport Framework.** *J. Nucl. Eng.* 2021; **2**(1): 57–64. [Publisher Full Text](#)
23. Weiss MA: **Data Structures & Algorithm Analysis in C++.** Pearson Education; 4th ed. 2012.
24. Molnar B, Tolnai G, Legrady D: **A GPU-based direct Monte Carlo simulation of time dependence in nuclear reactors.** *Ann. Nucl. Energy.* Oct. 2019; **132**: 46–63. [Publisher Full Text](#)
25. Leppänen J: **Two Practical Methods for Unionized Energy Grid Construction in Continuous-Energy Monte Carlo Neutron Transport Calculation.** *Ann. Nucl. Energy.* Jul. 2009; **36**(7): 878–885. [Publisher Full Text](#)
26. Leppänen J, Pusa M, Viitanen T, et al.: **The Serpent Monte Carlo code: Status, development and applications in 2013.** *Jt. Int. Conf. Supercomput. Nucl. Appl. Monte Carlo 2013 SNA MC 2013 Pluri- Transdiscip. New Model. Simul. Paradig.* Aug. 2015; **82**: 142–150. [Publisher Full Text](#)
27. Martin WR: **Challenges and prospects for whole-core Monte Carlo analysis.** *Nucl. Eng. Technol.* Mar. 2012; **44**(2): 151–160. [Publisher Full Text](#)
28. Lund A, Siegel AR, Forget B, et al.: **Using fractional cascading to accelerate cross section lookups in Monte Carlo particle transport calculations.** *Presented at the Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method.* Nashville, Tennessee: Apr. 2015.
29. Cormen TH, Leiserson CE, Rivest RL, et al.: **Introduction to Algorithms, Third Edition.** The MIT Press; 3rd ed. 2009.
30. Liu Y, She D, Wang K, et al.: **Optimization treatment of point-wise nuclear data in Monte Carlo criticality and burnup calculations.** *Ann. Nucl. Energy.* Jul. 2011; **38**(7): 1489–1495. [Publisher Full Text](#)
31. Brown FB: **New Hash-Based Energy Lookup Algorithm for Monte Carlo Codes.** Los Alamos National Laboratory, LA-UR-14-24530; 2014.
32. Wang Y, Brun E, Malvagi F, et al.: **Competing energy lookup algorithms in Monte Carlo neutron transport calculations and their optimization on CPU and Intel MIC architectures.** *J. Comput. Sci.* May 2017; **20**: 94–102. [Publisher Full Text](#)
33. Chen Z, Xie J, Guo Q, et al.: **Physics-oriented optimization strategy for the energy lookup algorithm in continuous energy Monte Carlo neutron transport simulation.** *Comput. Phys. Commun.* Jan. 2019; **234**: 146–158. [Publisher Full Text](#)
34. Raffuzzi V, Shwageraus E, Morgan L: **Accelerating Monte Carlo neutron transport by approximating thermal cross sections with functional forms.** *Ann. Nucl. Energy.* May 2022; **169**: 108819. [Publisher Full Text](#)
35. Bergmann RM, Vujić JL: **Algorithmic Choices in Warp – A Framework for Continuous Energy Monte Carlo Neutron Transport in General 3D Geometries on GPUs.** *Ann. Nucl. Energy.* Mar. 2015; **77**: 176–193. [Publisher Full Text](#)
36. Scopatz A, Romano P, Wilson P, et al.: **PyNE: Python for nuclear engineering.** *Presented at the 2012 ANS Annual Winter Meeting.* 2012; pp. 985–987.
37. Hamilton SP, Evans TM: **Continuous-Energy Monte Carlo Neutron Transport on GPUs in the Shift Code.** *Ann. Nucl. Energy.* Jun. 2019; **128**: 236–247. [Publisher Full Text](#)
38. Choi N, Kim KM, Joo HG: **Optimization of Neutron Tracking Algorithms for GPU-Based Continuous Energy Monte Carlo Calculation.** *Ann. Nucl. Energy.* Nov. 2021; **162**: 108508. [Publisher Full Text](#)
39. Choi N: **Development of GPU-Based Deterministic and Probabilistic Direct Whole-core Calculation Systems.** Seoul National University Graduate School; 2021. [Reference Source](#)
40. Tramm JR, et al.: **Toward Portable GPU Acceleration of the OpenMC Monte Carlo Particle Transport Code.** *Presented at the International Conference on Physics of Reactors 2022 (PHYSOR 2022).* 2022.
41. Tramm J, et al.: **Performance Portable Monte Carlo Particle Transport on Intel, NVIDIA, and AMD GPUs.** *EPJ Web Conf.* 2024; **302**. [Publisher Full Text](#)
42. Ridley G, Forget B: **Design and Optimization of GPU Capabilities In OpenMC.** *Presented at the 2021 ANS Winter Meeting.* 2021.
43. McKinley M, et al.: **Status of LLNL Monte Carlo transport codes on Sierra GPUs.** *Proceedings of the American Nuclear Society.* Portland, Oregon, USA: Aug. 2019.
44. Pozlup M, et al.: **Progress Porting LLNL Monte Carlo Transport Codes to Nvidia GPUs.** *Presented at the International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering.* Niagara Falls, Ontario, Canada: Aug. 2023.
45. Morgan JP, Mote A, Pasmann SL, et al.: **The Monte Carlo Computational Summit – October 25 & 26, 2023 – Notre Dame, Indiana, USA.** *J. Comput. Theor. Transp.* Jul. 2024; **53**(5): 361–382. [Publisher Full Text](#)
46. Descalle MA, Beck BR, Mattoon CM, et al.: **Implementation Of The GNDS Format For Evaluated And Processed Nuclear Data.** *Presented at the PHYSOR, Cancun, Mexico.* Apr. 2018.
47. NEA: **Specifications for the Generalised Nuclear Database Structure Version 2.0.** Paris: OECD Publishing; 2023.
48. Hawari AI: **On a measurement approach to support evaluation of thermal scattering law data.** *Ann. Nucl. Energy.* Jan. 2020; **135**:

106940.
[Publisher Full Text](#)
49. Tang Y, Zu T, Yi S, *et al.*: **Development and verification of thermal neutron scattering law data calculation module in nuclear data processing code NECP-Atlas.** *Ann. Nucl. Energy.* Apr. 2021; **153**: 108044.
[Publisher Full Text](#)
50. Fleming NC, Manring CA, Laramée B, *et al.*: **FLASSH 1.0: Thermal Scattering Law Evaluation and Cross Section Generation.** *EPJ Web Conf.* 2023; **284**.
[Publisher Full Text](#)
51. Squires GL: *Introduction to the Theory of Thermal Neutron Scattering.* Cambridge: Cambridge University Press; 3rd ed. 2012.
[Publisher Full Text](#)
52. Tran HN, *et al.*: **Comparison of the thermal neutron scattering treatment in MCNP6 and GEANT4 codes.** *Nucl. Instrum. Methods Phys. Res. Sect. Accel. Spectrometers Detect. Assoc. Equip.* Jun. 2018; **893**: 84–94.
[Publisher Full Text](#)
53. Jiménez-Carrascosa A, Fridman E, Herranz N, *et al.*: **About the impact of the Unresolved Resonance Region in Monte Carlo simulations of Sodium Fast Reactors.** *Presented at the ICAPP 2019 - International Congress on Advances in Nuclear Power Plants.* Juan-les-pins, France: May 2019.
[Publisher Full Text](#)
54. Mosteller RD, Little RC: **Impact of MCNP Unresolved Resonance Probability-Table Treatment on Uranium and Plutonium Benchmarks.** *Presented at the ICNC'99: International Conference on Nuclear Criticality Safety.* Versailles, France: Sep. 1999.
[Reference Source](#)
55. Levitt LB: **The Probability Table Method for Treating Unresolved Neutron Resonances in Monte Carlo Calculations.** *Nucl. Sci. Eng.* Dec. 1972; **49**(4): 450–457.
[Publisher Full Text](#)
56. Sutton TM, Brown FB: **Implementation of the probability table method in a continuous-energy Monte Carlo code system.** *Presented at the International Conference on the Physics of Nuclear Science and Technology.* Long Island, NY (United States): Oct. 1998.
[Reference Source](#)
57. NVIDIA: **NVIDIA Compute Sanitizer.** Accessed: Aug. 05, 2024.
[Reference Source](#)
58. Ali MR, Aygul MS, Lee D: **GREAPMC: development status, capabilities, and future developments.** *EPJ Nucl. Sci. Technol.* 2025; **11**.
[Publisher Full Text](#)
59. Ali MR, Aygul MS, Lee D: **Development of New GPU-Optimized Reactor Physics Monte Carlo Code GREAPMC.** *Nucl. Sci. Eng.* Nov. 2025: 1–16.
[Publisher Full Text](#)
60. Ali MR, Aygul MS, Lee D: **Performance Evaluation of GREAPMC for Continuous Energy Reactor Physics Problems.** *Presented at the M&C.* Colorado, USA; 2025.
61. Ali MR, Aygul MS, Lee D: **Multi-GPU Multi-Node Implementation of CUDA-aware MPI in GREAPMC.** *Presented at the KNS.* Jeju, ROK; May 2025.
62. Menéndez ML, Pardo JA, Pardo L, *et al.*: **The Jensen-Shannon divergence.** *J. Frankl. Inst.* Mar. 1997; **334**(2): 307–318.
[Publisher Full Text](#)
63. Cullen DE, *et al.*: *How Accurately can we Calculate Thermal Systems?* United States, INDC (USA)-107: International Atomic Energy Agency; Apr. 2004.
[Reference Source](#)
64. Mosteller RD: *Computational Benchmarks for the Doppler Reactivity Defect.* Los Alamos, NM, LA-UR-06-2968: Los Alamos National Laboratory (LANL); Jun. 2006.
65. Mosteller RD: **ENDF/B-V, ENDF/B-VI, and ENDF/B-VII.0 Results for the Doppler-Defect Benchmark.** *Presented at the Joint International Topical Meeting on Mathematics and Computations and Supercomputing in Nuclear Applications.* California, USA: Apr. 2007.
66. Choi S, Lee H, Hong SG, *et al.*: **Resonance self-shielding methodology of new neutron transport code STREAM.** *J. Nucl. Sci. Technol.* Sep. 2015; **52**(9): 1133–1150.
[Publisher Full Text](#)
67. Sunny EE, Brown FB, Kiedrowski BC, *et al.*: **Temperature Effects of Resonance Scattering for Epithermal Neutrons in MCNP.** *Presented at the International Conference on Physics of Reactors (PHYSOR) - Advances in Reactor Physics.* Tennessee, USA: Apr. 2012.
68. Choi S, Lee D: **Preliminary Results for OPR1000/APR1400 Whole-core Analysis with Neutron Transport Code STREAM.** *Presented at the M&C.* Oregon, USA: Aug. 2019.
69. Hamilton SP, Slattery SR, Evans TM: **Multigroup Monte Carlo on GPUs: Comparison of History and Event-Based Algorithms.** *Ann. Nucl. Energy.* Mar. 2018; **113**: 506–518.
[Publisher Full Text](#)
70. Griesheimer DP, Gill DF, Nease BR, *et al.*: **MC21 v.10 – coupled-radiation Monte Carlo transport solver with support for multiphysics simulations.** *EPJ Nucl. Sci. Technol.* 2024; **10**: 10.
[Publisher Full Text](#)

Open Peer Review

Current Peer Review Status:   

Version 2

Reviewer Report 20 February 2026

<https://doi.org/10.21956/nuclscitechnolopenres.19033.r28516>

© 2026 Variansyah I. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.



Ilham Variansyah 

Oregon State University, Corvallis, Oregon, USA

The authors have addressed the reviewers' comments carefully, and the manuscript has improved accordingly. I am satisfied with the revisions and support for Indexing.

However, I would like to echo Reviewer 1's observation regarding the potential perception of "salami slicing." While I understand and appreciate the authors' clarification that Hornet is primarily intended to serve as the continuous-energy cross-section engine for GREAPMC, the current manuscript framing particularly aims for a broader contribution. In this context, giving more focus on isolated performance assessment of Hornet (e.g., via a representative continuous-energy benchmark such as XSBench) would further strengthen the manuscript and help clarify its standalone technical contribution. Framing the GREAPMC integration as a strong demonstration case, rather than the primary justification, would provide better alignment with the paper's narrative and title.

Nevertheless, overall, I find the work to be solid and well presented, and I believe it makes a meaningful contribution in its current form.

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Radiation transport, Monte Carlo method, high-performance computing, reactor physics

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Reviewer Report 20 February 2026

<https://doi.org/10.21956/nuclscitechnolopenres.19033.r28520>

© 2026 Cosgrove P. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.



Paul Cosgrove 

University of Cambridge, Cambridge, UK

The article is well-written and interesting. It will be of use to future developers of GPU neutronics codes. I have mainly quite minor comments.

There is some inconsistent use of 'cross section' and 'cross-section'. Please pick one and stick with it.

In equation 2, instead of \sum_x it would be better to use Σ_x

A typo on page 4: you say 'Hornets are fully object-oriented and are written in CUDA C++'. I think this should be 'Hornet is fully object-oriented and is written in CUDA C++'.

Shortly afterwards you say 'verification of the Hornet'. I think it should just be 'verification of Hornet'.

Beginning section 2, you should use 'ACE' immediately before the bracketed '(A Compact ENDF)'.

It should be 'ACE data originates', not 'originate'.

A very minor point: there is a newer SCONE reference. It would be preferable to use it instead of the old one: **Refer to reference no.1**

Likewise, there is a more recent Serpent reference, rather than 26 as used in the text: "Status of Serpent Monte Carlo code in 2024". This is in an EPJN issue following SNA+MC.

On page 10 there is another mention of 'Hornets' rather than 'Hornet'.

At the start of 3.2.1.1, I think it should be 'passed by reference', not 'passed by the reference'.

On Table 6, in the corresponding text you say that 'the reactivity differences converge toward zero within uncertainty'. It is possible that I am misunderstanding the table, but this does not appear to be true: for HFP there is a difference well beyond uncertainty when using 1E7 particles (-7.8pcm +/-0.9pcm). This should at least be explained or commented on.

At the start of section 4.5, it should be 'The performance of the codebase' rather than 'The performance of codebase'.

It would be good to confirm that the performance of the CPU simulation has not saturated after a certain number of cores. If that were the case, it would unfairly advantage the GPU in performance comparisons.

References

1. Raffuzzi V, Cosgrove P, Kowalski M: Status of the SCONE Monte Carlo neutron transport code. *Annals of Nuclear Energy*. 2026; **227**. [Publisher Full Text](#)

Is the work clearly and accurately presented and does it cite the current literature?

Yes

Is the study design appropriate and does the work have academic merit?

Yes

Are sufficient details of methods and analysis provided to allow replication by others?

Yes

If applicable, is the statistical analysis and its interpretation appropriate?

Yes

Are all the source data underlying the results available to ensure full reproducibility?

No source data required

Are the conclusions drawn adequately supported by the results?

Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Radiation transport, Monte Carlo, neutronics

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Reviewer Report 27 January 2026

<https://doi.org/10.21956/nuclscitechnolopenres.19033.r28515>

© 2026 Park H. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

**Ho Jin Park**

Kyunghee University College of Engineering, Yongin-si, Gyeonggi-do, South Korea

Thank you for your thorough response. The distinctive features of Hornet's GPU-based continuous-energy cross-section processing—namely the parallelization of sampling routines, the modular architecture, and the optimized memory-handling strategies—are clearly presented and largely understandable. However, I would like to note that if a separate manuscript on GREAPMC has already been published, I am concerned that a subsequent publication focusing on an integrated GREAPMC–Hornet system could potentially be considered a case of “salami slicing.”

Thank you for your efforts and times, again.

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Reactor Physics

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard.

Version 1

Reviewer Report 16 September 2025

<https://doi.org/10.21956/nuclscitechnolopenres.18966.r28232>

© 2025 Variansyah I. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.



Ilham Variansyah

Oregon State University, Corvallis, Oregon, USA

Overall Assessment

The manuscript presents the development and verification of Hornet, a GPU-optimized continuous-energy (CE) cross-section code, and demonstrates its integration within a GPU-based Monte Carlo code, GREAPMC. The work is timely, given the increasing interest in accelerating Monte Carlo (MC) reactor physics simulations through GPU architectures. The authors provide a solid overview of CE techniques in the context of GPUs, highlight their design choices (e.g., object-oriented paradigm, single-precision nuclear data, energy lookup strategies), and report extensive verification and performance results.

I find the paper suitable for indexing, and I recommend acceptance with some reservations. Most of my comments below are not critical revisions but are intended to help further strengthen the manuscript for clarity, completeness, and extended impact.

====

Context and comparisons to existing work

- The manuscript claims novelty in developing GPU-optimized CE cross-section processing that is Monte Carlo agnostic. However, this statement should be carefully nuanced in light of MCGIDI, which also functions as a standalone CE processing tool. A more direct comparison to MCGIDI (or MCGIDI-like approaches) would provide a stronger context for evaluating Hornet's unique contributions.

- Discussions on GPU-specific continuous-energy treatments in several Monte Carlo codes are given. It may also be worthwhile to discuss CE treatments in GUARDYAN for completeness.

- While running a full MC transport simulation using the GREAPMC-Hornet code system provides a useful performance comparison, tools/benchmarks such as XSBench might serve as a more direct baseline for assessing computational performance of CE physics alone.
- Other than enabling multigroup GPU MC code (like GREAPMC) to run continuous-energy physics and serving as an academic tool, authors may also want to discuss potential uses of Hornet-like codes in a CPU MC code, such as an asynchronous cross-section lookup that has been explored in the MC code MC21.

Hornet VS GREAPMC-Hornet evaluations

- Authors should emphasize that most of the presented verification exercises and performance assessments are those of the GREAPMC-Hornet system, not the Hornet code itself. The emphasis is important because verification exercises for the GREAPMC-Hornet system do not necessarily verify Hornet itself (integration VS unit verification). The same goes for the performance assessment. Nevertheless, evaluation of the GREAPMC-Hornet system is still relevant and important to the paper, as it is one of the main purposes of the Hornet development and a typical use of such a modular code.
- Consider tools/benchmarks like XSBench for assessing the computational performance of CE physics alone.
- The authors should briefly discuss any changes in the Hornet-enabled GREAPMC from the cited, original multigroup GREAPMC. If a separate paper will be published to provide an in-depth discussion of this, please state so.

Outgoing energy distribution sampling comparison

- The use of Jensen-Shannon divergence (JSD) to quantify agreement between MCS and GREAPMC-Hornet sampled distributions is innovative. However, most readers probably will not be familiar with the interpretation of JSD values. Please discuss the typical range or min/max of JSD values and what constitutes "small" and whether the value of 0.1 can be considered sufficiently low in this context.
- Certain incident energies and temperatures are used for generating the results. Please state whether similar agreements are observed in other energies and temperatures.

Verification through code-to-code comparison of MC results

The presented verification effort compares particular instances of MC simulations (at a given number of samples N). To better verify through code-to-code comparison of MC results, one can demonstrate whether the two codes' results would reasonably converge to the same answer at the expected rate of $1/\sqrt{N}$. Consider running with (logarithmically) increasing numbers of samples N and observe the convergence of the relative difference (average of the two can be used as the reference), and take the norm/RMS of the differences to get a single measure for each N . Note that this would also shed light on the significance of the single vs double-precision arithmetic, which would be valuable. Similar error convergence analysis can be used in the Thermal Systems, Mosteller Benchmark, and Pin-by-Pin Power verifications as well.

Content Accuracy and Verification

- "Most of the runtime in the Monte Carlo (MC) code is spent on continuous-energy cross-sectional lookups." How universal the sentence is is debatable. Consider "In many classes of MC transport simulations, particularly neutron fission reactor systems,". Note that systems with highly complex geometry, such as fusion neutronics and medical imaging, can be counterexamples to this.
- For transparency, it would be helpful to state explicitly whether all numbers of inactive cycles

used are sufficient to ensure fission source convergence.

- Please state whether the uncertainties in the keff difference in the Thermal Systems are propagated standard deviations.
- Please briefly discuss how the standard deviations in the Doppler coefficients results (tables and figures) are calculated/propagated, particularly on how the possible correlation is treated or avoided.

Performance Assessment

- The GPU-vs-CPU performance comparison properly involves increasing workload to demonstrate saturation of parallel computing efficiency on both GPU and CPUs.
- Performance metrics reported are based on active cycle runtime. Please state whether the runtime assessment excludes cycle closeout operations, such as tally score accumulations and fission bank management for reproducibility.
- Please discuss what tallies are included in the performance calculations. Would the size of the tallies affect the performance comparison results?
- Please explain why certain physics treatments (e.g., constant cross-section elastic scattering, equi-lethargy energy lookup) were used in the performance tests. How sensitive are the results to these treatment choices, and are they particularly favorable to GPU execution?
- Please clarify whether the development/administrative tasks running on 10% of cores (16 total) had any measurable impact on the MCS performance assessment. If negligible, this should be mentioned along with steps/treatments taken to isolate the MC computations.
- A note on whether MCS also uses shared-memory parallelism, and (if it does) whether a single-node run would improve MCS parallel efficiency would be useful.

====

Minor Comments

Wording

- "the continuous treatment ... reduces the errors introduced by discretization" --> "the continuous treatment ... avoids the errors introduced by discretization".
- "Microscopic cross sections are the interaction probability" --> "Microscopic cross sections measure/quantify the interaction probability".

Equations

- Equation (2): Consider denoting macro cross sections with Σ symbol to distinguish them from summation signs.
- Equation (9): Appears to be missing parentheses.

Formatting

- Section 2.1.5 (Hash-based energy lookup) heading should be bolded for consistency.
- Correct the fragment "However, the evaluated nuclear data Files) data".

References

- Please provide a proper bibliographic citation for Chen et al. (2018) instead of only a URL.

Is the work clearly and accurately presented and does it cite the current literature?

Yes

Is the study design appropriate and does the work have academic merit?

Yes

Are sufficient details of methods and analysis provided to allow replication by others?

Yes

If applicable, is the statistical analysis and its interpretation appropriate?

Yes

Are all the source data underlying the results available to ensure full reproducibility?

Yes

Are the conclusions drawn adequately supported by the results?

Yes

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Radiation transport, Monte Carlo method, high-performance computing, reactor physics

I confirm that I have read this submission and believe that I have an appropriate level of expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 16 Jan 2026

MURAT SERDAR AYGUL

Context and comparisons to existing work

- The manuscript claims novelty in developing GPU-optimized CE cross-section processing that is Monte Carlo agnostic. However, this statement should be carefully nuanced in light of MCGIDI, which also functions as a standalone CE processing tool. A more direct comparison to MCGIDI (or MCGIDI-like approaches) would provide a stronger context for evaluating Hornet's unique contributions.

We thank the reviewer for the suggestion and for highlighting the relevance of MCGIDI as a Monte Carlo-agnostic continuous-energy processing framework. We agree that a clearer distinction was needed to appropriately position Hornet with respect to existing tools. To address this comment, we have revised Section 2.2.5 to explicitly compare Hornet with GIDI/MCGIDI. The manuscript now clarifies that while both approaches are Monte Carlo-agnostic, GIDI/MCGIDI is primarily designed for CPU architectures and relies on serialized or limited parallel execution when used in GPU contexts. In contrast, Hornet is developed from the ground up for GPU architectures, with data structures and lookup mechanisms specifically optimized for massive parallelism. This clarification better explains Hornet's unique contribution as a native, GPU-centric continuous-energy cross-section processing framework, rather than solely a standalone or Monte Carlo-agnostic tool.

- Discussions on GPU-specific continuous-energy treatments in several Monte Carlo codes are given. It may also be worthwhile to discuss CE treatments in GUARDYAN for completeness.

We thank the reviewer for the suggestion to improve the completeness of the discussion. In response, we have added a brief description of the continuous-energy treatment employed in GUARDYAN

- While running a full MC transport simulation using the GREAPMC-Hornet code system provides a useful performance comparison, tools/benchmarks such as XSBench might serve as a more direct baseline for assessing computational performance of CE physics alone.

We thank the reviewer for the valuable suggestion. The primary objective of Hornet is to serve as the continuous-energy cross-section processing engine for GREAPMC, while its standalone capability is retained mainly for educational purposes. Consequently, the performance evaluation in this work is focused on the integrated GREAPMC-Hornet system rather than isolated CE cross-section benchmarks. For this reason, a direct comparison with standalone CE physics benchmarks such as XSBench or Papillon has not been included.

- Other than enabling multigroup GPU MC code (like GREAPMC) to run continuous-energy physics and serving as an academic tool, authors may also want to discuss potential uses of Hornet-like codes in a CPU MC code, such as an asynchronous cross-section lookup that has been explored in the MC code MC21.

We appreciate the reviewer's insightful comment. In response, we have expanded the discussion to highlight that, beyond its integration with GREAPMC, Hornet could also be employed in a heterogeneous execution model in which continuous-energy cross sections are generated on the GPU while particle transport is carried out on CPUs, with cross-section requests handled asynchronously. **Hornet VS GREAPMC-Hornet evaluations**

- Authors should emphasize that most of the presented verification exercises and performance assessments are those of the GREAPMC-Hornet system, not the Hornet code itself. The emphasis is important because verification exercises for the GREAPMC-Hornet system do not necessarily verify Hornet itself (integration VS unit verification). The same goes for the performance assessment. Nevertheless, evaluation of the GREAPMC-Hornet system is still relevant and important to the paper, as it is one of the main purposes of the Hornet development and a typical use of such a modular code.

The reviewer raises an important distinction between unit-level verification of Hornet and system-level verification of the coupled GREAPMC-Hornet framework. To address this point, the manuscript has been revised to explicitly clarify the scope of the presented verification and performance assessments. As now stated in the text, the primary focus of this work is the evaluation of Hornet in the form of targeted tests, such as sampling and comparison of selected reactions and integrated Hornet-GREAPMC system. Other components of the development such as the multigroup GREAPMC implementation and the elastic scattering treatment in Hornet have been verified in prior studies.

- Consider tools/benchmarks like XSBench for assessing the computational performance of CE physics alone.

As the primary objective of the GPU-related development work in the UNIST CORE lab is to preserve the accuracy of MCS while accelerating overall Monte Carlo simulation time, accuracy and performance have consistently been assessed through direct comparisons with MCS rather than with standalone CE micro-benchmarks such as XSBench.

- The authors should briefly discuss any changes in the Hornet-enabled GREAPMC from the cited, original multigroup GREAPMC. If a separate paper will be published to provide an in-depth discussion of this, please state so.

The manuscript has been updated to briefly clarify the key modification introduced in the Hornet-enabled version of GREAPMC. **Outgoing energy distribution sampling comparison**

- The use of Jensen–Shannon divergence (JSD) to quantify agreement between MCS and GREAPMC-Hornet sampled distributions is innovative. However, most readers probably will not be familiar with the interpretation of JSD values. Please discuss the typical range or min/max of JSD values and what constitutes “small” and whether the value of 0.1 can be considered sufficiently low in this context.

Thank you for pointing out this question. The manuscript has been revised to clarify the interpretation of Jensen-Shannon divergence (JSD) values for readers who may not be familiar with this metric. In particular, we now explicitly state the bounded range of JSD (0 to 1), where a value of zero corresponds to identical distributions and larger values indicate increasing dissimilarity. Additional context is provided to indicate that JSD values on the order of 0.1 represent small but noticeable differences in distribution shape, which in the present study arise primarily from differences in numerical precision between MCS and Hornet. To further anchor this interpretation, we also report that when Hornet and MCS are run with matched precision for elastic scattering, the resulting JSD values are on the order of 10^{-5} as reported in Ali et al. (EPJ-N, 2025).

- Certain incident energies and temperatures are used for generating the results. Please state whether similar agreements are observed in other energies and temperatures.

Thank you for this comment. The manuscript has been revised to clarify the scope of the results. Similar levels of agreement between MCS and GREAPMC-Hornet are observed for other incident energies and temperatures, as well as for other reactions. In this work, fission reactions are presented as a representative case to illustrate general agreement, while the trends for other reactions and conditions are consistent with those shown.

Verification through code-to-code comparison of MC results

- The presented verification effort compares particular instances of MC simulations (at a given number of samples N). To better verify through code-to-code comparison of MC results, one can demonstrate whether the two codes' results would reasonably converge to the same answer at the expected rate of $1/\sqrt{N}$. Consider running with (logarithmically) increasing numbers of samples N and observe the convergence of the relative difference (average of the two can be used as the reference), and take the norm/RMS of the differences to get a single measure for each N . Note that this would also shed light on the significance of the single vs double-precision arithmetic, which would be valuable. Similar error convergence analysis can be used in the Thermal Systems, Mosteller Benchmark, and Pin-by-Pin Power verifications as well.

We thank the reviewer for this valuable suggestion. In response, we extended the verification to explicitly examine statistical convergence between MCS and GREAPMC as a function of particle histories. A representative Mosteller benchmark case was evaluated using logarithmically increasing numbers of particles ($10^4 - 10^7$). As shown in the revised manuscript (Tables 6 and 7), both the reactivity differences (HFP and HZP) for Mosteller's case and the RMS deviation of pin power tallies decrease systematically with increasing

particle histories. **Content Accuracy and Verification**

- "Most of the runtime in the Monte Carlo (MC) code is spent on continuous-energy cross-sectional lookups." How universal the sentence is is debatable. Consider "In many classes of MC transport simulations, particularly neutron fission reactor systems,". Note that systems with highly complex geometry, such as fusion neutronics and medical imaging, can be counterexamples to this.

The reviewer is correct that the original statement was overly general. In particular, Monte Carlo simulations involving highly complex geometries, such as fusion neutronics or medical imaging applications, may be dominated by geometric tracking rather than cross-section processing. The manuscript has therefore been revised to clarify that the statement applies specifically to Monte Carlo neutron transport simulations for criticality analysis in fission reactor systems.

- For transparency, it would be helpful to state explicitly whether all numbers of inactive cycles used are sufficient to ensure fission source convergence.

We thank the reviewer for this suggestion. To improve clarity and transparency, the manuscript has been revised to explicitly state that the number of inactive cycles used in all cases was selected based on fission source convergence assessed using Shannon entropy, and that the chosen values were verified to be sufficient to ensure convergence of the fission source.

- Please state whether the uncertainties in the k_{eff} difference in the Thermal Systems are propagated standard deviations.

We thank the reviewer for this comment. The uncertainties reported for the k_{eff} differences in the thermal systems are propagated standard deviations, obtained by assuming statistically independent Monte Carlo estimates from the two codes.

- Please briefly discuss how the standard deviations in the Doppler coefficients results (tables and figures) are calculated/propagated, particularly on how the possible correlation is treated or avoided.

We thank the reviewer for this comment. The standard deviations reported for the Doppler coefficients are obtained through standard error propagation applied to the difference of multiplication factors used to compute the coefficient. Specifically, the Doppler coefficient is derived from two independently estimated k_{eff} values corresponding to different fuel temperatures. Since these Monte Carlo simulations are performed independently, the resulting statistical uncertainties are assumed to be uncorrelated. Under this assumption, the uncertainty in the Doppler coefficient is obtained by propagating the one-standard-deviation uncertainties of the individual k_{eff} values. The manuscript has been revised to explicitly clarify this procedure and the underlying assumption of statistical independence.

Performance Assessment

- The GPU-vs-CPU performance comparison properly involves increasing workload to demonstrate saturation of parallel computing efficiency on both GPU and CPUs.

We thank the reviewer for this comment. We agree that a comprehensive GPU-versus-CPU performance comparison ideally involves increasing the computational workload to demonstrate saturation of parallel efficiency on both GPU and CPU architectures. In the present work, the focus is on verification and consistency assessment rather than an exhaustive performance study. To this end, we have included results showing the variation of reactivity differences for the Mosteller benchmark and RMS deviations of pin power tallies with increasing numbers of particle histories, which demonstrate statistical convergence and computational consistency between the GPU- and CPU-based

implementations. A detailed performance evaluation addressing workload scaling and parallel efficiency saturation on both GPUs and CPUs has been carried out separately and is reported in Ali et al. (M&C 2025). We have added an explicit reference to this work in the manuscript.

- Performance metrics reported are based on active cycle runtime. Please state whether the runtime assessment excludes cycle closeout operations, such as tally score accumulations and fission bank management for reproducibility.

The reported performance metrics are based on the total active cycle runtime and include cycle closeout operations, such as tally score accumulation and fission bank management required for reproducibility. No components of the active cycle, including closeout procedures, were excluded from the timing measurements.

- Please discuss what tallies are included in the performance calculations. Would the size of the tallies affect the performance comparison results?

The performance measurements include cell-wise power and neutron flux tallies. The size of the tallies primarily affects the memory footprint and does not significantly impact the performance comparison. GREAPMC has previously demonstrated efficient scaling with increasing numbers of cells (as reported in the referenced NSE paper, 2025). Moreover, since GREAPMC employs cell-based tallying, the tally operations introduce negligible computational overhead and do not measurably affect the reported performance.

- Please explain why certain physics treatments (e.g., constant cross-section elastic scattering, equi-lethargy energy lookup) were used in the performance tests. How sensitive are the results to these treatment choices, and are they particularly favorable to GPU execution?

The selected physics treatments were chosen to ensure both representativeness of production Monte Carlo workflows and a fair assessment of GPU performance. The double-index and linear interval hashing energy lookup schemes are known to be highly efficient on GPUs and are implemented in Hornet as configurable options. Since widely used production codes such as OpenMC and MCS employ equi-lethargy-based hash energy lookup, this approach was adopted in the performance tests to enable a meaningful and consistent comparison. The constant cross-section elastic scattering treatment was selected because it involves only analytical expressions and avoids complex data-dependent operations. This model represents the fastest elastic scattering option on CPUs among the available target velocity sampling treatments, while also mapping efficiently to GPU execution. As a result, its use avoids introducing artificial bias toward either architecture and allows the performance results to primarily reflect computational throughput rather than differences in physics model complexity.

- Please clarify whether the development/administrative tasks running on 10% of cores (16 total) had any measurable impact on the MCS performance assessment. If negligible, this should be mentioned along with steps/treatments taken to isolate the MC computations.

The development and administrative tasks running on 10% of the available cores had no measurable impact on the MCS performance assessment. These tasks were limited to node monitoring and code compilation and were not active during the Monte Carlo production runs.

- A note on whether MCS also uses shared-memory parallelism, and (if it does) whether a single-node run would improve MCS parallel efficiency would be useful.

We appreciate the reviewer's comment. The version of MCS employed in this work does not

utilize shared-memory parallelism. **Minor Comments Wording**

- “the continuous treatment ... reduces the errors introduced by discretization” --> “the continuous treatment ... avoids the errors introduced by discretization”.
- “Microscopic cross sections are the interaction probability” --> “Microscopic cross sections measure/quantify the interaction probability”.

We thank the reviewer for these suggestions. Both revisions have been incorporated for improved clarity and precision:

1. The phrase “*the continuous treatment ... reduces the errors introduced by discretization*” has been updated to “*the continuous treatment ... avoids the errors introduced by discretization*” to more accurately convey that discretization errors are eliminated rather than merely reduced.
2. The statement “*Microscopic cross sections are the interaction probability*” has been revised to “*Microscopic cross sections quantify the interaction probability*” to more precisely describe the physical meaning of microscopic cross sections.

Equations

- Equation (2): Consider denoting macro cross sections with Σ symbol to distinguish them from summation signs.

We thank the reviewer for the suggestion. In the submitted manuscript, different symbols were already used to distinguish the summation operator from the macroscopic cross-section symbol. We have carefully reviewed Equation (2) and again added different symbols to ensure that this distinction is clear in the revised manuscript.

- Equation (9): Appears to be missing parentheses.

We thank the reviewer for pointing this out. Parentheses have been added to Equation (9) in the revised manuscript to ensure clarity and correct mathematical interpretation.

Formatting

- Section 2.1.5 (Hash-based energy lookup) heading should be bolded for consistency.

We thank the reviewer for the comment. The heading 2.1.5 (Hash-based energy lookup) has been bolded in the revised manuscript to ensure consistency with other section headings.

- Correct the fragment “However, the evaluated nuclear data Files) data”.

We thank the reviewer for identifying the fragment. The text “However, the evaluated nuclear data Files) data” has been corrected for clarity and now reads appropriately in the revised manuscript.

References

- Please provide a proper bibliographic citation for Chen et al. (2018) instead of only a URL.

We thank the reviewer for the comment. A proper bibliographic citation for Chen et al. (2018) has been added in the revised manuscript in place of the previous URL

Competing Interests: No competing interests were disclosed.

Reviewer Report 05 September 2025

<https://doi.org/10.21956/nuclscitechnolopenres.18966.r28231>

© 2025 Park H. This is an open access peer review report distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.



Ho Jin Park

Kyunghee University College of Engineering, Yongin-si, Gyeonggi-do, South Korea

In this paper, the authors introduced new continuous-energy cross section processing code system, Hornet for GPU-based Monte Carlo simulations. This Hornet code can generate the ACE format cross section and efficiently handle the cross section data for GREAPMC code. The tables and figures are well-organized to effectively support the results. I have some questions and comments before publication.

1. In many GPU/CPU-based Monte Carlo analysis codes, routines for reading or processing nuclear data are naturally included. However, the CUDA-based Monte Carlo code developed by the authors separates this function and is divided into two codes, Hornet and GREAPMC. In this manuscript, sufficient justification and rationale for this approach are not provided.
2. If you wish to publish on this system, I recommend doing so with a manuscript that covers the integrated system combining GREAPMC and Hornet. Based on the content of the current paper, Hornet does not appear to have any distinctive features compared to conventional Monte Carlo cross-section processing. To adequately address the originality of this research, you should either emphasize such distinctive features or submit the work together with the development details of GREAPMC.

Is the work clearly and accurately presented and does it cite the current literature?

Yes

Is the study design appropriate and does the work have academic merit?

Yes

Are sufficient details of methods and analysis provided to allow replication by others?

Yes

If applicable, is the statistical analysis and its interpretation appropriate?

Yes

Are all the source data underlying the results available to ensure full reproducibility?

Yes

Are the conclusions drawn adequately supported by the results?

Partly

Competing Interests: No competing interests were disclosed.

Reviewer Expertise: Reactor Physics

I confirm that I have read this submission and believe that I have an appropriate level of

expertise to confirm that it is of an acceptable scientific standard, however I have significant reservations, as outlined above.

Author Response 13 Sep 2025

MURAT SERDAR AYGUL

Comment: In many GPU/CPU-based Monte Carlo analysis codes, routines for reading or processing nuclear data are naturally included. However, the CUDA-based Monte Carlo code developed by the authors separates this function and is divided into two codes, Hornet and GREAPMC. In this manuscript, sufficient justification and rationale for this approach are not provided.

Response: We appreciate the reviewer's observation. The separation of Hornet and GREAPMC was an intentional design choice aimed at modularity, flexibility, and efficiency. By developing Hornet as a standalone CUDA-based cross-section processing tool, we ensure that it can be used independently or integrated into multiple Monte Carlo frameworks, rather than being restricted to GREAPMC. This design enables independent verification and benchmarking of nuclear data processing and allows Hornet to be used for academic purposes, such as obtaining and analyzing different cross-sections without requiring Monte Carlo modeling. It also simplifies maintenance, upgrades, and extensions of the processing routines, while ensuring reusability in other projects or codes without dependence on GREAPMC. Furthermore, debugging and optimization are improved, as cross-section treatment can be validated separately from transport. Thus, while many conventional codes integrate both functions, our approach was guided by the goals of flexibility and long-term usability.

Comment: If you wish to publish on this system, I recommend doing so with a manuscript that covers the integrated system combining GREAPMC and Hornet. Based on the content of the current paper, Hornet does not appear to have any distinctive features compared to conventional Monte Carlo cross-section processing. To adequately address the originality of this research, you should either emphasize such distinctive features or submit the work together with the development details of GREAPMC.

Response: We thank the reviewer for this valuable suggestion. We agree that presenting Hornet together with GREAPMC would highlight the originality of the end-to-end workflow. However, in this manuscript our primary focus was to demonstrate Hornet's ability to perform GPU-optimized continuous-energy cross-section processing and to benchmark its performance against existing tools. Hornet's distinctive features include its full CUDA parallelization of cross-section routines such as Doppler broadening and probability table sampling, its modular architecture that allows coupling with different Monte Carlo solvers, and its optimized GPU memory management strategies tailored for continuous-energy processing. Looking ahead, we plan a follow-up publication that presents the integrated GREAPMC-Hornet system in detail, and we will also carry out verification of the coupled framework to demonstrate the performance gains and consistency achieved when both codes operate together.

Competing Interests: No competing interests were disclosed.

Author Response 16 Jan 2026

MURAT SERDAR AYGUL

1. In many GPU/CPU-based Monte Carlo analysis codes, routines for reading or processing nuclear data are naturally included. However, the CUDA-based Monte Carlo code developed by the authors separates this function and is divided into two codes, Hornet and GREAPMC. In this manuscript, sufficient justification and rationale for this approach are not provided.

We appreciate the reviewer's observation. The separation of Hornet and GREAPMC was an intentional design choice aimed at modularity, flexibility, and efficiency. By developing Hornet as a standalone CUDA-based cross-section processing tool, we ensure that it can be used independently or integrated into multiple Monte Carlo frameworks, not limited to GREAPMC. This modular design allows:

- Independent verification and benchmarking of nuclear data processing.
- Using Hornet for academic purposes. For instance, to obtain and analyze different cross-sections without having a requirement for Monte Carlo modeling.
- Easier maintenance, upgrades, and extension of the processing routines.
- Reusability of Hornet in other projects or codes without relying on GREAPMC.
- Improved debugging and optimization, as cross-section treatment can be validated separately from transport.

Thus, while many conventional codes integrate both functions, our approach was motivated by flexibility and long-term usability. The details have been added into the manuscript.

1. If you wish to publish on this system, I recommend doing so with a manuscript that covers the integrated system combining GREAPMC and Hornet. Based on the content of the current paper, Hornet does not appear to have any distinctive features compared to conventional Monte Carlo cross-section processing. To adequately address the originality of this research, you should either emphasize such distinctive features or submit the work together with the development details of GREAPMC.

We thank the reviewer for this suggestion. We agree that presenting Hornet together with GREAPMC would highlight the end-to-end workflow and its originality. However, in this manuscript, our focus was to demonstrate Hornet's ability to handle GPU-optimized continuous-energy cross-section processing and its benchmarking against existing tools.

Distinctive features of Hornet compared to conventional processors include:

- Full CUDA parallelization of cross-sections, Doppler broadening and probability table sampling routines, enabling runtime reduction.
- A modular architecture that allows direct coupling with different Monte Carlo solvers.
- Optimized GPU memory handling strategies designed specifically for continuous-energy processing on GPUs.

We acknowledge that emphasizing these features more clearly would strengthen the manuscript, and therefore we have added some details about distinct Hornet features in the manuscript. We also plan a follow-up publication that presents the integrated GREAPMC-Hornet system in detail, highlighting the performance improvements achieved when both

codes operate together.

Competing Interests: No competing interests were disclosed.