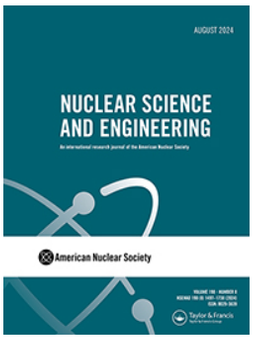

Taylor & Francis
Taylor & Francis Group

Nuclear Science and Engineering

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/unse20

High-Resolution Modeling Without Computation Slowdown for PETALE in CROCUS

Thomas Ligonnet, Axel Laureau, Andreas Pautz & Vincent Lamirand

To cite this article: Thomas Ligonnet, Axel Laureau, Andreas Pautz & Vincent Lamirand (14 Jun 2024): High-Resolution Modeling Without Computation Slowdown for PETALE in CROCUS, Nuclear Science and Engineering, DOI: 10.1080/00295639.2024.2357963

To link to this article: https://doi.org/10.1080/00295639.2024.2357963

Published online: 14 Jun 2024.

NUCLEAR SCIENCE AND ENGINEERING

DOI: https://doi.org/10.1080/00295639.2024.2357963

High-Resolution Modeling Without Computation Slowdown for PETALE in CROCUS

Thomas Ligonnet,[a]* Axel Laureau,[b,c] Andreas Pautz,[a,d] and Vincent Lamirand[a,d]

[a]*Ecole Polytechnique Fédérale de Lausanne, Laboratory for Reactor Physics and Systems Behaviour, Lausanne CH-1015, Switzerland*
[b]*Université Grenoble Alpes, CNRS, Grenoble INP, LPSC-IN2P3, Grenoble 38000, France*
[c]*SUBATECH, CNRS/IN2P3, IMT-Atlantique, Nantes F-44307, France*
[d]*Paul Scherrer Institut, Nuclear Energy and Safety Research Division, Villigen CH-5232, Switzerland*

Received November 17, 2023
Accepted for Publication May 14, 2024

Abstract — *In a collaboration between Ecole Polytechnique Fédérale de Lausanne (EPFL) and CEA, in the fall of 2020, the experimental Programme d'Étude en Transmission de l'Acier Lourd et ses Eléments (PETALE) was successfully carried out in the CROCUS reactor of EPFL. This article presents and compares the methods tested in the modeling of the experiments, specifically focusing on the metal reflectors installed at the periphery of CROCUS. A basic design model consisting of a few cuboids was refined to a fully detailed version, without impacting the run time of simulations. Notably, each reflector sheet of PETALE was segmented into 121 voxels based on topological measurements. This detailed voxelization did not affect calculation times, thanks to the use of three-dimensional lattices as available in Serpent 2. Profiling of the simulations revealed the high computational surface transformations associated with Serpent 2 and highlighted the efficiency benefits of factorizing these into universe transformations. As the CROCUS simulations were carried out using a modified build of Serpent 2, additional simulations were also performed using a standard version of Serpent 2 with a GODIVA model as a neutron source to ensure that the findings are generalizable. These additional tests confirmed the initial results, with significant performance variations observed between the models, particularly larger in surface-tracking mode than in delta-tracking mode. Consequently, the modeling method may therefore be applied to future high-fidelity modeling of neutron transmission and shielding experiments.*

Keywords — *Monte Carlo, modeling, voxelization, integral experiments, iron nuclear data, CROCUS.*

Note — *Some figures may be in color only in the electronic version.*

I. INTRODUCTION

The Ecole Polytechnique Fédérale de Lausanne (EPFL)–CEA experimental Programme d'Étude en Transmission de l'Acier Lourd et ses Eléments (PETALE),[1] which focuses on the study of nuclear data for stainless steel and its primary components, i.e., iron, nickel, and chromium, was successfully conducted in fall 2020 using the EPFL light water zero-power research reactor CROCUS. The first experimental output of the program comprised measurements of reaction rates for selected dosimetry reactions as a function of penetration depth in metallic reflectors, i.e., transmission experiments. Another output of the project was calculation of the computation-experiment (C/E) ratios for the aforementioned reaction rates using Monte Carlo simulations. These simulations utilized state-of-the-art nuclear data libraries for neutron transport and the specialized International Reactor Dosimetry and

*E-mail: thomas.ligonnet@epfl.ch

Fusion File (IRDFF-II) library for dosimetry reactions.[1] These C/E ratios were employed both for validation and for subsequent data assimilation.[2,3]

Integral experiments offer the advantage of low experimental uncertainty and, when combined with appropriate covariance evaluations, are useful in validation[4] and evaluation[5] processes. Specific needs for such experiments were identified for iron.[6] Since its inception, PETALE aimed to deliver results that benefit the nuclear data community by ensuring that these results constrain the nuclear data. To achieve this, during the design phase, target experimental uncertainties were quantified thanks to the propagation of nuclear data uncertainties in a preliminary experimental setup and its corresponding design model,[2,7] utilizing Total Monte Carlo (TMC) and the TENDL2019 library for cross-section samples.[8] The same uncertainty propagation approach is consistently chosen[9] in the ongoing analysis phase to produce a detailed uncertainty budget with covariances. This approach requires a refined model of the PETALE experiments to extensively quantify the impact of uncertainties from geometries and materials on the C/E, aiming to minimize potential biases. Additionally, it involves defining material cards via chemical analysis for a broad range of alloys, occasionally requiring multiple definitions for varying production batches. The model also encompasses all structural components and metallic reflector sheets, with the latter's modeling being the primary focus of this work. Although the impact of these components on simulation results is under evaluation, topological measurements have revealed slight deviations in the shapes and thicknesses of some sheets, as detailed in Sec. IV.B. To determine if these variations are negligible, all reflector sheets have been voxelated based on micrometric topological measurements. The chosen approach is to convert each of the original cuboids into 121 smaller units.

Five modeling approaches, of increasing complexity, have been tested on the same experimental setup for the iron reflector case. These models range from a basic raw design model to a complete refined model of PETALE, which extensively employs Serpent 2 universes and lattices. This includes testing two voxelization methods, i.e., cuboidal and lattice, for accurately representing the eight reflector sheets. An additional model, cuboidal_FT, allows us to study the impact of transformations. The need for each component to be able to move freely with respect to their free space is kept in mind for all models, to allow the estimation of positional uncertainties through Monte Carlo simulations of randomly perturbed geometries. A profiling is performed for each model to determine the sources of slowdown and assess the performance of all models. Consequently, we can pinpoint which methods to avoid when defining voxelated geometries.

Finally, because a custom Serpent 2 build is used to simulate CROCUS, simulations of the PETALE case next to a GODIVA model[10,11] are performed with a standard Serpent 2 version. These simulations, using the same case definitions as those in CROCUS, aim to verify that performance differences among the models are observable beyond the custom build. Simulations with the standard version are conducted using both delta-tracking, unavailable in the custom build, and surface-tracking modes.

II. CONTEXT OF THE STUDY

II.A. The CROCUS Zero-Power Reactor

CROCUS is a zero-power light water experimental reactor located at EPFL in Lausanne, Switzerland. It operates under ambient pressure and a controlled temperature of 20°C. Its operating license authorizes a maximum thermal power of 100 W, which implies a maximum total neutron flux of 2.5×10^9 n/cm^2·s at its center. The main means of controlling CROCUS reactivity is an adjustable water level via a mobile spillway. The position of the spillway modifies the immersion of the fuel rods, thus influencing neutron thermalization. The water level can be set with a precision of up to ±0.1 mm, equivalent to a reactivity variation of ±0.4 pcm. In addition, two optional boron carbide (B_4C) control rods can be placed in dedicated channels. The control rods can be adjusted within a range of ±0.5 mm, corresponding to a maximum reactivity variation of around ±0.2 pcm based on the S-curve of the control rods.[12–14]

The reactor core is divided into two interlocking zones: an inner zone and an outer zone, separated by a water gap. The inner zone comprises 336 fuel rods containing uranium dioxide enriched to 1.806% by weight while the outer zone comprises fuel rods made of uranium metal enriched to 0.947% by weight. The number of rods in the outer zone can be customized to suit the specific experimental configuration. The configuration studied in this paper is a 174-rod metallic uranium configuration. CROCUS is equipped with four monitors: two boron-coated compensated ionization chambers to the east and west and two fission chambers to the north and south, which serve as power and safety monitors. The fission chambers operate in pulse mode, with an adjustable dwell time, generally set at 1 s. This configuration

enables the reactor to be monitored in real time and is commonly used to reconstruct the reactor's power history.[13]

II.B. PETALE

In collaboration with CEA, PETALE consists of semi-integral experiments focused on reactor criticality and neutron transmission.[7] The program includes mobile metal reflectors, placed in turn in an aluminum case, which replace part of the water reflector in the reactor vessel (see Fig. 1). Four distinct reflectors are employed to analyze the properties of nuclear-grade stainless steel and its principal elements: iron, nickel, and chromium. The elemental (or semi-integral) nature of these experiments eliminates material ambiguities, thereby reducing compensation effects among elemental nuclear data. Each reflector comprises eight metal sheets. The reactivity worth of the reflectors is determined through critical experiments whereas neutron transmissions are measured using activation dosimeters positioned between the sheet and on both the front and the back of the case.[15]

III. SERPENT 2 AND VARIANCE REDUCTION FOR PETALE

Serpent 2 is a Monte Carlo neutron and photon three-dimensional (3D) transport code developed by VTT. The code supports neutron source and k-eigenvalue simulations and has two neutron-tracking modes: surface tracking and delta tracking.[16] In surface-tracking mode, the code calculates the neutron path length using the macroscopic total cross sections of the material through which it is currently passing. This requires stopping the neutron at each surface crossing to update material information and calculate a new path length. Consequently, it is necessary to compute a maximum path length for each neutron displacement to ensure the neutron stops before exiting the material. In delta-tracking mode, the neutron path length is computed using a majorant cross section, e.g., the highest macroscopic total cross sections of the geometry. At the end of the path, the code samples whether a collision really happened using the ratio between the current material cross section and the majorant as the probability. Since the path length is independent of the local material cross section, stopping the neutron at each surface crossing is unnecessary.[17,18]

Simulations for the analysis of PETALE utilize a custom version of Serpent 2.1.21,[16] which was developed in-house specifically for this project and updated for the work at hand. It includes a variance reduction method to accelerate the convergence of PETALE dosimeters.[2] A weight map is produced on a distance-to-hit basis and updated throughout the simulation to allow a one-step calculation. This feature allows obtaining the reaction rates of the dosimeter placed in the PETALE case in single k-eigenvalue simulations, thus avoiding a more cumbersome two-step process. Typically, this process would first involve a k-eigenvalue simulation to tally the neutron entering the case, followed by a source mode simulation to obtain the dosimeters' reaction rates.

Additionally, this build supports correlated sampling. While not utilized in this publication, this functionality will later be employed to propagate the reflector nuclear data uncertainties and to conduct Bayesian Monte Carlo Assimilation with the PETALE results.[3,9] The JEFF-3.3 nuclear data library is selected for neutron transport,[19] but any other transport nuclear data library could be used for this study. However, reaction rate tallies of the dosimeters are consistently performed using the dosimetry specialized

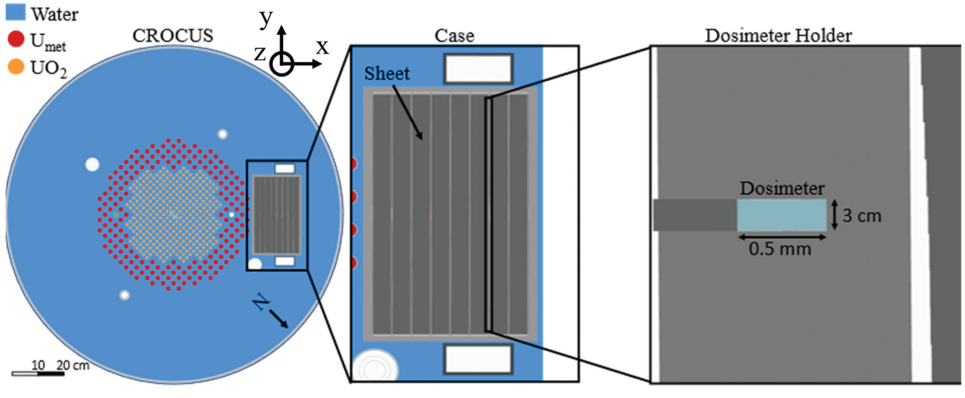

Fig. 1. Axial cross section of the PETALE configuration. (a) CROCUS core in its vessel, surrounded by its water reflector partially replaced by a case containing one of the metal reflectors. (b) Close-up of the case and the eight sheets that compose the metal reflector. (c) Close-up of a dosimeter in a dosimeter holder, with a voxelated neighboring sheet visible on the right.

nuclear data library IRDFF-II.[1] A significant limitation of our build is the inability to use delta tracking, as it is incompatible with the implementations.

IV. MODELING AND SIMULATIONS

IV.A. PETALE Geometry

PETALE examines four different metallic reflectors: one Type 304L stainless steel reflector and three high-purity (>99 wt%) elemental reflectors made of iron, nickel, and chromium. These reflectors were successively placed in the CROCUS reactor to conduct transmission and criticality experiments (see Fig. 1). Consequently, each reflector is housed in the same aluminum case, containing the eight metallic sheets, each 2 cm thick and approximately 30 × 30 cm, of the studied material. Nine dosimeter holders are positioned before, between, and after the sheets. Each dosimeter holder can carry up to nine cylindrical activation dosimeters, each with a diameter of 29.5 mm and a thickness of about 0.5 mm, except for gold dosimeters (0.01 mm). Additionally, aluminum structural parts surround the reflector case to allow its precise positioning in the periphery of CROCUS at a short distance from the core, approximately 5 mm away from the closest fuel rods, depending on the reflector material and corresponding weight. In this study, the iron reflector serves as the test case.

The objective of the modeling is to calculate the dosimeters' reaction rates and to propagate the associated geometrical uncertainties to these rates. To achieve this, a sampling approach is employed. A Python script has been developed to automatically generate each of the 26 PETALE experimental geometries, with the option to randomly perturb the position of each part. The case itself is positioned according to the uncertainties in the measurements of its position and orientation in the reactor, while each element in the case is moved with respect to its mechanical clearance. Thus, each part is subject to group and individual perturbations.

IV.B. Models for Reflector Voxelization

The first step of the model refinement, starting from the raw design model, is to voxelate each of the reflectors' metallic sheets. In this context, voxelization refers to the conversion of sheet models consisting of single cuboids into models composed of multiple surfaces and cells, better replicating the actual shapes of the sheets. The shape of each sheet was determined through topological measurements conducted regularly on each face and side. The plots in Fig. 2 display examples of topological measurement outputs. Figures 2a and 2b illustrate the topology of the front and rear faces of the sheet, and the color map shows the difference in axial position relative to the reference axial position of each face. Additionally, the level lines on the maps highlight these differences. Figure 2c comprises the topological measurements of the sheets' sides. Figure 2d illustrates the local deviation from the expected 20-mm thickness of the sheet; on this map the level lines indicate the sheet thickness. Using the topological measurements, each sheet is represented as an 11×11 matrix of thickness measurements, allowing the voxelization of each of them into 121 voxels. The side measurements are also utilized to reproduce the shape of the sides of the sheets, resulting in irregular widths and heights for the 40 border voxels.

To model the shapes of the sheets, two approaches guaranteeing the same geometry are tested. The first approach will be referred to as the cuboidal model. It is a conventional and direct use of Serpent cuboid surfaces. Each voxel, defined by a cuboid surface, is filled with the

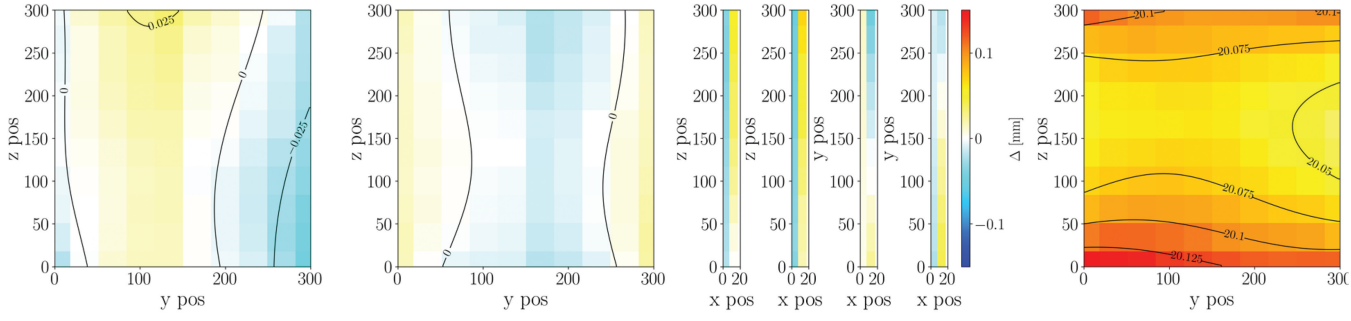

Fig. 2. Topological measurements of a metal sheet. (a) and (b) The topological measurements of the front face and rear face of the sheet, respectively; the color map displays the difference in position (in millimeters) with respect to the measurement system reference. (c) The four sides of the sheet with the same color map. (d) Color map displaying the local difference (in millimeters) with the expected 20-mm thickness. The coordinate system is the same as the one in Fig. 1.

reflector material. The cuboids are orderly placed in the model to replicate the sheets' shape. Those not positioned at the edges maintain a consistent width and height of 3 cm, while the thickness matches the measured values. At the sheet edges, height and width adjustments are made according to the measured topology. Each metal sheet consists of 122 cells: the 121 metal-filled cuboids, and 1 larger cuboid envelops these to fill the remaining space with air as the last cell. Since cuboid surfaces are defined with their surfaces perpendicular to the Serpent 2 reference axes, each surface has an associated transformation card for proper orientation.

The second approach will be referred to as the lattice model. It takes advantage of the regularity of the topological measurements. In this case, the reflector' sheets are defined using $11 \times 11 \times 1$ Serpent 2 3D lattices, creating 121 subuniverses per sheet, with each subuniverse representing a voxel. In this model, each voxel's volume is identical, containing both air and metal, divided by surfaces. This model distinguishes three types of voxels: 81 at the center, 36 on the sides, and 4 at the edges of the sheet. Voxels are bounded by two surfaces defining the sheet's upper and lower boundaries, while the sides correspond to the lattice's y and z pitches (in Figs. 1 and 2 system of coordinates). On the sides, an additional third surface is necessary to replicate the topology of the sheet's smaller sides, as well as a fourth surface in the case of edges.

IV.C. Refined Model and Structure Optimization for Computation Acceleration

The refined model includes the addition of all structural components and compositions, along with the voxelization of metallic sheets using the lattice approach. PETALE's geometry lends itself to using cells directly defined by neighboring cuboid surfaces since most internal and structural parts are cuboidal and separated by either air or water. However, accurately replicating the geometry and isotopic composition of all components necessitates the multiplication of cells and surfaces. It is hypothesized that because of the slowdown of the simulation results from the numerous operations required to determine the cell the neutron is in—rather than from the number of parts in the geometry given the emphasis on geometric fidelity over simplification—the chosen solution was to completely rework the geometry. This rework aims to leverage the definitions of Serpent 2's universes and nested universes, introducing an intermediate level for the structural elements, between CROCUS and the PETALE case. As a result of this refinement, the geometry of the case evolved from a single-level design model, containing one universe and a total of 70 cells, into a four-level refined model of PETALE. This refined model comprises 975 universes, including the lattices' subuniverses, and a total of 3407 cells.

IV.D. Model Profiling and Test Model cuboidal_FT

To discern why one model is more computationally efficient than another, profilings of Serpent 2 running each model are performed. During the simulations, the profiler measures the time spent in the different processes of the codes and reports the percentage of total run time spent in each of them. Additionally, it provides the intrinsic run time, i.e., time spent exclusively in that process excluding time in subfunctions, and the total run time, which includes time spent in these subfunctions.

The profiler used in this study is GNU gprof, included in the GNU Binutils 2.39 release.[20] The percentage of time spent in each process is used as an indicator to compare the different models. Profiling allows us to identify the processes most heavily utilized based on the geometry definition; however, it does not, generally speaking, provide direct and quantitative information about the code's running speed. In this study, the qualitative indicator of code running speed is the percentage of run time the code spent in the process that manages the "neutron transports." This is because the "neutron transport" process competes for runtime percentage with processes unrelated to the method used to define geometry, such as "collision sampling" and "event scoring." Profiling various geometry definitions helps determine how changes affect the time the code spends on each subprocess of "neutron transport" relative to one another. To ascertain which model definition modifications reduce the total run time, we must identify which "neutron transport" subprocesses have decreased in intrinsic run time and contributed to a shorter overall run time for "neutron transport." Since adding a profiler impacts the code's performance, Sec. IV.E confirms the observations with another indicator, the figure of merit (FOM).

During profiling, six subprocesses related to the "neutron transport" were identified. The first is "cell identification," which determines the current position of the neutron. The second process, "surface test," assesses which side of a surface the neutron is on, a process notably used for "cell identification." The third process determines the neutron "distance to cell boundary," which frequently calls the fourth process, "distance to surface," to calculate distances to the cell's surfaces and identify the shortest one. The fifth process, "find region in lattice," is of importance only for the lattice model. The sixth process, handling "translations and rotations" in the geometry—also known as transformations—is frequently used by all other processes.

The cuboidal model is the first to be profiled. Profiling shows that the simulation is spending 91.87% of the total run time in the "neutron transport." Most of this time is passed in subprocesses, with only 0.09% of intrinsic time. The total time consumption is mostly split between the "cell identification," accounting for 59.34% of the total run time, and the calculation of the "distance to cell boundary," which takes up to 31.98%. Both processes have low intrinsic run time and are mostly calling their two subprocesses: "surface test" and "distance to surface." Although these two processes have significant intrinsic run times of 7.57% and 10.70%, respectively, they are overshadowed by the run time necessary to perform the "translations and rotations." This process, mostly used by the preceding two, concludes the chain of processes with its equal intrinsic and total run-time percentages of 67.59%. This result was unexpected, yet the reason why such a large time is spent in the "translations and rations" process is readily identifiable. To accurately replicate the experimental setup, it is necessary to rotate the surfaces into their correct orientation, while the translation is a simple tool to independently displace them when perturbing the geometry for evaluation of positional uncertainties.

Consequently, the model was defined with a transformation card per surface. Therefore, it is assumed that the use of multiple "translations and rotations" may be particularly computationally expensive. Subsequently, a new model with factorized transformations, named the cuboidal_FT model, was developed. This model will test whether the factorization reduces the total run time of the "neutron transport," simply shifts the load to other processes, or does not affect it at all.

The additional model, cuboidal_FT, retains the same geometry as the cuboidal model, using identical surface declarations as input. The difference in the cuboidal_FT model is that all cells of the PETALE reflector are declared in a single "PETALE universe." Furthermore, all "translations and rotations" from the previous model are factorized into a single universe transformation. The first observation from the profiled simulation is the decreased total run-time percentage of the "neutron transport." Specifically, the total run-time percentage for the cuboidal_FT is 12% lower than that of the cuboidal model, dropping to 79.29%. This reduction implies a gain in computational efficiency, as confirmed in Sec. IV.E. Additionally, The profiling results confirm a reduced run time for "translations and rotations," achieving the objective of the transformations' factorization. Indeed, the total and intrinsic run-time percentage for "transformation and rotations" has significantly decreased from 67.59% to just 3.46%. This decrease is 50% of total run-time percentage larger than the reduction observed for the "neutron transport." This difference implies that as expected, part of the calculations previously allocated to "translations and rotations" has been redistributed. The measured intrinsic run-time percentages for other processes, some of which were almost negligible in the cuboidal model, confirm this shift. For instance, the intrinsic run-time percentages for "cell identification" and "the determination of the distance to cell boundary" have increased to 18.58% and 8.89%, respectively. Additionally, there are increases in the intrinsic run time of already time-consuming processes such as "surface test," which rose to 16.36%, and notably, "distance to surface," which increased by more than 18% to 28.72%, now accounting for more than a quarter of the run time. As will be detailed in Sec. IV.E, the computational efficiency has significantly improved after factorizing the transformations compared to the cuboidal model. However, it falls short of the results achieved with the following lattice model.

As mentioned in Sec. IV.B, while the geometry of the lattice model remains the same, its definition method markedly differs through the use of 3D lattices for the sheets. First, the results from the profiled simulation reveal a significant decrease in the total run-time percentage dedicated to "neutron transport." The total run-time percentage decreases to 41.17%, which is less than half that obtained with the cuboidal_FT model. In this model, the intrinsic run-time percentage for "translations and rotations" stands at 8.65%, significantly lower than that observed in the cuboidal model. This is because "translations and rotations" are applied directly to the lattices rather than to the individual surfaces within them. Consequently, these operations are partly factorized already, making them more akin to the results obtained with the cuboidal_FT model. Additionally, although the intrinsic run time for "cell identification" has increased by 10% to 13.12%, this rise is more than offset by the simplification in calculating "distance to surfaces" and "surface test." Their intrinsic run-time percentages have decreased by 7% and 5%, dropping to 3.43% and 2.82%, respectively. Finally, the additional process required to "find region in lattice" has proven itself to have a negligible impact. The run-time percentages for each relevant process and model are displayed in Table I.

IV.E. Model Performance

Simulations are performed without a profiler to calculate the FOM of the PETALE dosimeters for each model. The FOM is defined as the inverse of the squared relative error R^2 multiplied by the simulation time T: FOM = $1/(R^2T)$.[21] To compare the models, each dosimeter's FOM is divided by its counterpart of the design model to calculate a relative FOM. The nine relative FOMs are then averaged into a model's mean relative FOM. For simplicity, simulations were run in

TABLE I

Intrinsic (and Total) Time Consumption Fraction of the Functions of Interest

Models Process	Design	Cuboidal	Cuboidal_FT	Lattice
"Neutron transport"	0.50% (50.18%)	0.09% (91.87%)	0.22% (79.29%)	0.53% (41.17%)
"Translations and rotations"	12.86% (12.86%)	67.59% (67.59%)	3.46% (3.46%)	8.65% (8.65%)
"Distance to cell boundary"	2.72% (12.75%)	1.66% (31.98%)	8.89% (38.79%)	3.10% (9.50%)
"Distance to surface"	5.27% (10.02%)	10.70% (30.32%)	28.72% (29.90%)	3.43% (6.40%)
"Cell identification"	17.76% (29.12%)	3.94% (59.34%)	18.58% (36.39%)	13.12% (21.16%)
"Surface test"	3.28% (11.35%)	7.57% (55.41%)	16.36% (17.81%)	2.82% (8.04%)
"Find region in lattice"	<0.5%	<0.5%	<0.5%	0.91% (1.11%)

k-eigenvalue mode, targeting a CPU time of 320 h. The number of source neutrons per cycle is set at 5000, and the number of inactive cycles is set at 100. The mean relative FOMs for each model are displayed in Table II.

The simulation performed with the lattice model demonstrates no loss in computational efficiency. In this model, the voxelization of the reflector sheets using 3D lattices does not hamper the simulation, resulting in a mean relative FOM of 1. Conversely, the cuboidal model exhibits a significant reduction in computational efficiency. In this model, where the reflector is modeled using cuboids, the mean relative FOM has decreased by a factor of 4.3. These results highlight the advantage of the lattice approach: It significantly enhances the accuracy of the geometry relative to the experimental setup while only moderately increasing its complexity and limiting the loss of computational efficiency.

Building on the profiling results, a simulation of the cuboidal_FT model is performed to calculate its mean relative FOM. As observed with the profiling, factorizing the rotations and translations significantly reduces their computational cost. The mean relative FOM of the cuboidal_FT model is 2.5 times higher than that of the original cuboidal model. When applicable, factorizing transformations is a straightforward method to optimize simulations and enhance input readability. In PETALE modeling, individual surface transformations were originally preferred because they simplify perturbing the positions of elements within the geometry. Indeed, in the cuboidal model, perturbations can be added directly to the surface transformation cards. In contrast, the cuboidal_FT model requires modification to the surface cards themselves, which can be more ambiguous and tedious.

Developed after profiling and calculating the mean relative FOM for the previous models, the refined model employs 3D lattices for the metal sheets of the reflectors, with extensively factorized transformations. This model achieves computational efficiency comparable to the design model, despite including additional structural elements and

TABLE II

Characteristics of the Simulations and Relative FOM

Models Characteristics	Design	Cuboidal	Cuboidal_FT	Lattice	Refined
Histories (10^6)	577	82	210	592	684
CPU time (h)	313.75	335.27	310.73	323.40	333.73
RAM (Gbyte)	8.2	8.2	8.2	8.2	8.2
Mean relative FOM	1	0.23	0.59	1.01	0.87

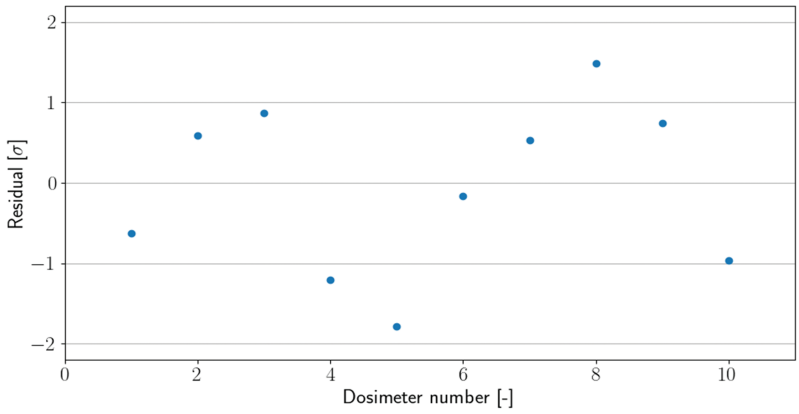

Fig. 3. Residuals for the ten ^{197}Au(n,g) dosimeter reaction rates of the lattice model with respect to the cuboidal model.

the voxelated sheets. As shown in Table II, the computed mean relative FOM is 0.87, equivalent to a 15% increase in CPU time to obtain the same dosimeter statistics as the design model. This slight difference is attributed to the added structural elements that favor thermal neutron streaming rather than to the number of cells and surfaces in the geometry.

The reaction rates obtained from the simulations for the nine ^{197}Au(n,g) dosimeters placed in the case, and a tenth one placed at the core center, show consistency between the lattice and cuboidal models. Indeed, all computed reaction rates for the lattice model fall within two standard deviations from their cuboidal model counterparts, with seven out of ten within one standard deviation. The residuals, displayed in Fig. 3, are defined as the ratio of the difference in reaction rates between models to the uncertainty of that difference.

V. COMPARISON OF MODELS USING A STANDARD VERSION OF SERPENT 2

PETALE simulations in CROCUS utilize a custom build of Serpent 2 to achieve sufficient dosimeter statistics at a reasonable computational cost. To ensure that the previously presented results are not artifacts of our custom solver, a set of simulations is performed using the current version of Serpent 2 (version 2.2.0) in k-eigenvalue mode, employing a reactor model suitable for this test. To simplify and accelerate the comparison, the quantity of interest for these simulations is the CPU time to complete a fixed number of cycles in a geometry that likely involves neutron entry into the PETALE reflector case while sustaining k-eigenvalue simulations, and without employing variance reduction methods. The GODIVA HEU-MET-FAST-001 geometry from the ICSBEP Handbook,[10] using the model input from the Serpent 2 wiki[11] (see Fig. 4), was selected for this purpose. Each model was tested with both the surface-tracking and the delta-tracking modes for this comparison, as the limitations of the custom solver do not apply. The settings chosen are 100 inactive cycles, 1000 active cycles, and a neutron population set to 10 000 per cycle. The CPU times for each model and tracking mode are displayed in Table III.

For surface tracking, the results confirm the observations made with the custom solver and the CROCUS reactor model. Compared to the design model, direct voxelization of the reflectors using cuboids (cuboidal model) significantly slows down the simulation. The simulation CPU time increases more than 100-fold. The cuboidal_FT model, which factorizes transformations, shows a clear improvement in computational efficiency compared to the cuboidal model. However, its simulation remains significantly slower than for the design model, with CPU time increasing by a factor of 20. Consistent with the results from CROCUS

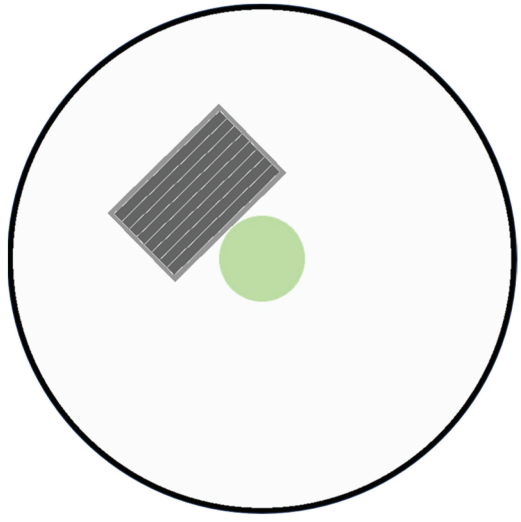

Fig. 4. Modeling of the PETALE reflector case next to the GODIVA sphere.

TABLE III

Total CPU Times in Minutes of Serpent 2.2.0 for Each Model and Tracking Mode

Models Tracking Mode	Design	Cuboidal	Cuboidal_FT	Lattice
Delta tracking	4.62	9.73	6.73	4.87
Surface tracking	6.06	812	138	8.2

simulations, the lattice model outperforms the others. Although it provides the same geometry, the slowdown in the simulation is limited to a 1.35-fold increase in CPU time.

In delta-tracking mode, all models perform better than in surface-tracking mode, although some slowdowns are still noticeable. Compared to the design model, the cuboidal model still experiences significant slowdowns, yet the increase in simulation CPU time is limited to a factor of 2.1. For the cuboidal_FT model, which includes factorized transformations, there is a gain in computational efficiency with the CPU time increasing by a factor of only 1.5. Finally, the lattice model maintains its efficiency, showing no loss compared to the design model. Using Serpent 2 3D lattices does not significantly increase the CPU time. As expected, delta-tracking mode outperforms surface tracking for models containing a large number of surfaces in a small volume. Nevertheless, even with delta tracking, the lattice model shows efficiency gains, which are beneficial for simulations considered lengthy by users. Where surface tracking is necessary, the time savings from using lattice voxelization are still evident. When modeling voxelated geometries, employing Serpent 2 3D lattices should always be considered, if applicable.

VI. CONCLUSION

The analysis of the PETALE transmission experiments employs high-resolution modeling to minimize geometrical biases as much as possible. This approach is supplemented by a full uncertainty propagation method using TMC, which aims to produce results in subsequent studies that will constrain nuclear data, for example, using the Bayesian Monte Carlo data assimilation method. Since the metal reflector is placed at the periphery of CROCUS, variance reduction was already essential in the simplified model for design purposes to achieve sufficient statistics within the small volumes of the dosimeters. However, during the analysis phase, the addition of numerous structural parts, elemental compositions, and voxelization of reflector sheets to accurately replicate their measured topology resulted in a significant increase in total computation time.

Computation profiles and FOMs were obtained with each model's simulation to identify and quantify sources of slowdown. The initial voxelated model, which utilized neighboring cuboids for the sheets, resulted in simulations more than four times slower than with the design model. We found that the extensive number of translations and rotations of surfaces is one of the main contributors to the slowdown. Factorizing most translations and rotations in a single universe transformation reduced the simulation time by a factor of 2.5. The trade-off was a moderate increase in the complexity of defining the geometry, particularly for applications involving geometry perturbations. The slowdown was eliminated by altering the modeling method of the sheets. Instead of using multiple cuboids, the sheets are now modeled with Serpent 2 3D lattices, taking advantage of the regularity of the topological measurements, and using surfaces to delimit the metal boundaries in the lattice subuniverses. Profiling showed that this approach did not slow down the simulations and significantly reduced the time the code spent computing the distance to surfaces. Thus, this approach was successful, effectively modeling the shape of each sheet without impacting the simulation speed. For verification, the models were also tested with a standard version of Serpent 2, transposing the PETALE case next to GODIVA rather than CROCUS. The results align with those observed using the custom solver and the CROCUS model. Slowdowns were especially noticeable for surface-tracking simulations, with significant slowdowns also observed for delta tracking. In both modes, employing lattices for voxelization has proven to be effective.

Finally, a high-resolution model of the PETALE experiments, including several additional structural elements, has been developed based on this study. This model employs 3D lattices for the metal sheets and factorizes necessary transformations, achieving a simulation speed comparable to the original design model. The model is currently being used to conduct simulations for

comparison with experimental results and will subsequently be used for validation and data assimilation.

Disclosure Statement

No potential conflict of interest was reported by the author(s).

ORCID

Thomas Ligonnet http://orcid.org/0000-0001-5955-0590
Axel Laureau http://orcid.org/0000-0003-4805-3673

References

1. E. M. ZSOLNAY et al., "Summary Description of the New International Reactor Dosimetry and Fusion File (IRDFF Release 10)," p. 44, International Atomic Energy Agency (2012).
2. A. LAUREAU et al., "Uncertainty Propagation for the Design Study of the PETALE Experimental Programme in the CROCUS Reactor," *EPJ Nucl. Sci. Technol.*, **6**, 9 (2020); https://doi.org/10.1051/epjn/2020004.
3. A. LAUREAU et al., "Bayesian Monte Carlo Assimilation for the PETALE Experimental Programme Using Inter-Dosimeter Correlation," *EPJ Web Conf.*, **239**, 18004 (2020); https://doi.org/10.1051/epjconf/202023918004.
4. J. B. BRIGGS, J. D. BESS, and J. GULLIFORD, "Integral Benchmark Data for Nuclear Data Testing Through the ICSBEP & IRPhEP," *Nucl. Data Sheets*, **118**, 396 (2014); https://doi.org/10.1016/j.nds.2014.04.090.
5. M. SALVATORES et al., "Methods and Issues for the Combined Use of Integral Experiments and Covariance Data: Results of a NEA International Collaborative Study," *Nucl. Data Sheets*, **118**, 38 (2014); https://doi.org/10.1016/j.nds.2014.04.005.
6. M. FLEMING et al., "Results of the Collaborative International Evaluated Library Organisation (CIELO) Project," *EPJ Web Conf.*, **239**, 15003 (2020); https://doi.org/10.1051/epjconf/202023915003.
7. V. LAMIRAND et al., "An Experimental Programme Optimized with Uncertainty Propagation: PETALE in the CROCUS Reactor," *EPJ Web Conf.*, **211**, 3003 (2019); https://doi.org/10.1051/epjconf/201921103003.
8. A. J. KONING et al., "TENDL: Complete Nuclear Data Library for Innovative Nuclear Science and Technology," *Nucl. Data Sheets*, **155**, 1 (2019); https://doi.org/10.1016/j.nds.2019.01.002.
9. A. LAUREAU et al., "Total Monte Carlo Acceleration for the PETALE Experimental Programme in the CROCUS Reactor," *EPJ Web Conf.*, **211**, 3002 (2019); https://doi.org/10.1051/epjconf/201921103002.
10. R. LABAUVE, "Bare, Highly Enriched Uranium Sphere (Godiva)," EU-MET-FAST-001," *International Handbook of Evaluated Criticality Safety Benchmark Experiments*, NEA-7328, Organisation for Economic Co-operation and Development, Nuclear Energy Agency (2021).
11. O. CHVALA, "Godiva Inputs – Serpent Wiki"; https://serpent.vtt.fi/mediawiki/index.php/Godiva_inputs (current as of Nov. 15, 2023).
12. V. LAMIRAND et al., "Future Experimental Programmes in the CROCUS Reactor," *Proc. Research Reactor Fuel Management/International Group Operating Research Reactors Conf.*, Berlin, Germany, February 29, 2016; https://doi.org/10.13140/RG.2.1.2309.6728.
13. V. LAMIRAND et al. "Power Calibration Methodology at the CROCUS Reactor," *EPJ Web Conf.* **225**, 04022 (2020); https://doi.org/10.1051/epjconf/202022504022.
14. U. KASEMEYER et al., "Benchmark on Kinetic Parameters in the CROCUS Reactor," Organisation for Economic Co-operation and Development, Nuclear Energy Agency (2007).
15. V. LAMIRAND, "Ten Springs of Experiments in CROCUS," *EPJ Web Conf.*, **288**, 4026 (2023); https://doi.org/10.1051/epjconf/202328804026.
16. J. LEPPÄNEN et al., "The Serpent Monte Carlo Code: Status, Development and Applications in 2013," *Ann. Nucl. Energy*, **82**, 142 (2015); https://doi.org/10.1016/j.anucene.2014.08.024.
17. I. LUX. *Monte Carlo Particle Transport Methods*, CRC Press, Boca Raton, Florida (2017); https://doi.org/10.1201/9781351074834.
18. J. LEPPÄNEN, "Performance of Woodcock Delta-Tracking in Lattice Physics Applications Using the Serpent Monte Carlo Reactor Physics Burnup Calculation Code," *Ann. Nucl. Energy*, **37**, *5*, 715 (2010); https://doi.org/10.1016/j.anucene.2010.01.011.
19. A. J. M. PLOMPEN et al., "The Joint Evaluated Fission and Fusion Nuclear Data Library, JEFF-3.3," *Eur. Phys. J. A*, **56**, *7*, 181 (2020); https://doi.org/10.1140/epja/s10050-020-00141-9.
20. J. CLIFTON, "GNU Binutils 2.39 Released," Free Software Foundation (May 8 2022); https://lists.gnu.org/archive/html/info-gnu/2022-08/msg00002.html (current as of Apr. 19, 2022).
21. J. KULESZA et al., "MCNP® Code Version 6.3.0 Theory & User Manual," LA-UR-22-30006, Los Alamos National Laboratory (2022); https://doi.org/10.2172/1889957.