One-Node pCMFD Acceleration of pFMFD for Improved Deterministic Truncation of Monte Carlo (iDTMC) Analyses of SFRs

Sunjoo Yoon [a*], Taesuk Oh [a], Inyup Kim [a], Jaehyeong Jang [a], Yonghee Kim [a]
[a] *Dept. of Nuclear & Quantum Engineering, Korea Advanced Institute of Science and Technology, 291 Daehak-ro, Yuseong-gu, Daejeon 34141*
[*] *Corresponding author: yongheekim@kaist.ac.kr*

Keywords : iDTMC, pFMFD, pCMFD, underrelaxation

1. Introduction

Sodium-cooled fast reactors, or SFRs, are a class of next-generation reactor concepts. Defined by their fast neutron spectrum and their use of liquid sodium as a coolant, they are frequently characterised by their high volume fraction of fuel, enabled by the high cooling efficiency of liquid sodium, and a honeycomb lattice fuel arrangement in hexagonal driver assemblies.

Research into a reactor concept requires analytic methodologies that can model it. However, due to SFRs lacking an asymptotic neutron spectrum, conventional two-step neutronic analysis methodologies are poorly suited for this purpose. Monte Carlo methodologies offer higher fidelity but also higher computational cost.

iDTMC is a neutronics analysis methodology that couples deterministic and stochastic computation to produce reactor subspace solutions with the fidelity of Monte Carlo at significantly reduced computing cost. Originally demonstrated for conventional pressurised water reactors [1], the iDTMC methodology has been shown to work well with SFRs [2]. Some preliminary analysis has also been conducted to estimate the stochastic uncertainty of iDTMC analyses of SFRs using the improved correlated sampling method [3].

The iDTMC method relies on pFMFD deterministic calculations to produce a subspace reactor solution, and the correlated sampling method for estimating its uncertainty involves solving many pFMFD problems. Reducing the computational cost of solving pFMFD problems would, therefore, reduce the overall computational cost of the iDTMC methodology. There have been preliminary efforts to use parallelisation to solve multiple pFMFD problems simultaneously [4], lowering the computational cost of iDTMC uncertainty analysis. This is ineffective for iDTMC without uncertainty analysis, with just one pFMFD problem.

In the present study, the one-node pCMFD acceleration scheme is applied to solving pFMFD problems, showing that this can reduce the computational cost of solving each individual pFMFD problem within the iDTMC methodology for SFRs.

2. Methodology

This study builds upon previous works, which have implemented the iDTMC methodology adapted for the honeycomb geometry of SFRs into the iMC code. The iMC code is a Monte Carlo-based neutronics analysis code developed internally at KAIST [5].

2.1 Improved Deterministic Truncation of Monte Carlo

In the improved deterministic truncation of Monte Carlo, or iDTMC, methodology, a Monte Carlo simulation is conducted during which fine-mesh factors are tallied on a pin-sized fine mesh. These fine-mesh factors are then used to construct a deterministic partial-current fine-mesh finite difference, or pFMFD, problem. Solving this pFMFD problem yields a subspace reactor solution, which is the final output of the iDTMC methodology. In this way, the iDTMC methodology is capable of producing a subspace solution with the fidelity of Monte Carlo without incurring the computing cost of conducting a large number of Monte Carlo active cycles to tally the quantities of interest directly.

The overall structure of the iDTMC methodology is summarised in Figure 1.

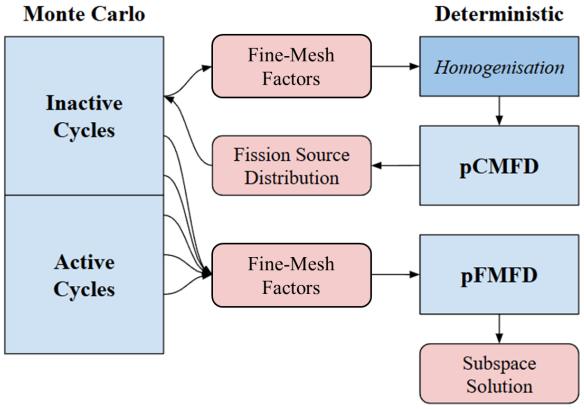

Fig. 1. Overall structure of the iDTMC methodology

As can be seen in the diagram, the inactive cycles of the Monte Carlo simulation within the iDTMC methodology are accelerated with the pCMFD method. Fine-mesh factors are tallied not just during the active cycles, but also over the later inactive cycles. The early inactive cycles in which local factors are not tallied are referred to as skip cycles.

2.2 Partial-Current Coarse-Mesh Finite Difference

The partial-current coarse-mesh finite difference, or pCMFD, method is a deterministic method of neutronics

analysis. In the iDTMC methodology, it is used to accelerate source convergence during the inactive cycles of the Monte Carlo simulation by adjusting the fission source between cycles.

The pCMFD method is based on the single energy group formulation of the FDM method. The neutron balance equation for a given node I in this formulation is given as follows:

$$\sum_J \frac{A_{IJ}}{V_I}\left(J_{IJ}^{net}\right) + \Sigma_{a,I}\phi_I = \frac{1}{k_{eff}}\nu_I\Sigma_{f,I}\phi_I \qquad (1)$$
$$J_{IJ}^{net} = D_{IJ}\frac{\phi_I - \phi_J}{\Delta_{IJ}}$$

where J represent each node adjacent to I and all other notations are as conventional.

In the pCMFD formulation, the net current is adjusted by correcting the partial current as follows.

$$J_{IJ}^{net} = J_{IJ}^{+} - J_{IJ}^{-} = \tilde{D}_{IJ}(\phi_I - \phi_J) + \hat{D}_{IJ}\phi_I - \tilde{D}_{JI}(\phi_J - \phi_I) - \hat{D}_{JI}\phi_J \qquad (2)$$

where $\tilde{D}_{IJ} = \tilde{D}_{JI} = \frac{D_{IJ}}{2\Delta_{IJ}}$ is the diffusion factor and the correction factor $\hat{D}_{IJ}$ is defined such that the reference partial current is reproduced, as follows.

$$J_{IJ}^{+,ref} = \tilde{D}_{IJ}\left(\phi_I^{ref} - \phi_J^{ref}\right) + \hat{D}_{IJ}\phi_I^{ref} \qquad (3)$$

The pCMFD calculations are only performed for the active core. The effects of reflectors, shields, and other structural elements outside the active core are modelled with a $\frac{J}{\phi}$ boundary conditions. At such boundaries,

$$J_{IJ}^{net} = \hat{D}_{IJ}\phi_I \qquad (4)$$
$$\hat{D}_{IJ} = \frac{J_{IJ}^{net,ref}}{\phi_I^{ref}} \qquad (5)$$

Constructing the neutron balance equations using these corrected currents for every node in the active core creates a system of linear equations, which may be expressed as an eigenvalue problem.

$$\boldsymbol{M}\boldsymbol{\phi} = \frac{1}{k_{eff}}\boldsymbol{F}\boldsymbol{\phi} \qquad (6)$$

This eigenvalue problem can then be solved by the usual iterative methods to yield a pCMFD reactor solution. In the current implementation, the biconjugate gradient stabilised method is used.

The reference cross-sections and partial currents used are those tallied in the previous Monte Carlo cycle. Once a pCMFD solution is obtained, the source weights used in the next cycle are adjusted to match the source distribution of the pCMFD solution.

2.3 Partial-Current Fine-Mesh Finite Difference

In the iDTMC methodology, a partial-current fine-mesh finite difference, or pFMFD, method is used to generate the final subspace reactor solution. The mathematical formulation of the pFMFD method is identical to that of the pCMFD method, save that it is done on a pin-sized fine mesh.

The reference currents and cross-sections used are the means of those tallied across all previous active and inactive cycles, except for the early skip cycles. There is no feedback from the pFMFD results back to the Monte Carlo simulation.

2.4 One-Node pCMFD Acceleration Scheme

In this study, a one-node pCMFD acceleration scheme has been implemented to accelerate the convergence of the pFMFD solution. The one-node pCMFD acceleration scheme reduces the cost of obtaining a high-fidelity solution, in this case a pFMFD solution, by iterating between local high-fidelity calculations and a global pCMFD calculation instead of performing expensive global high-fidelity calculations.

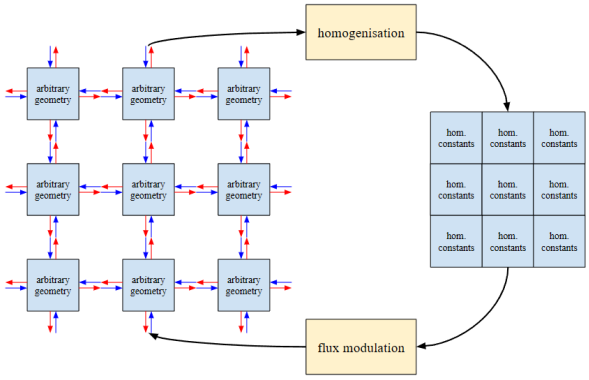

Fig. 2. Overall structure of one-node pCMFD acceleration

For a given pFMFD problem, coarse-mesh box-average fluxes, partial currents, and cross-sections are obtained by flux- and volume-weighted homogenisation. These factors are used to perform pCMFD calculations. The pCMFD calculations are not performed to convergence, but to a maximum of 5 source iterations.

The fine-mesh fluxes are modulated based on the pCMFD solution, as follows, where i is a fine-mesh cell in the coarse-mesh node I:

$$\phi_i^{mod} = \phi_i^{prev} \cdot \frac{\phi_I^{pCMFD}}{\sum_i^I \phi_i^{prev}} \qquad (7)$$

The fine-mesh source is updated using the above-calculated flux and the $\nu\Sigma_f$ values given in the pFMFD problem definition. The partial currents across fine-mesh cell surfaces on internal coarse-mesh node surfaces are then updated. Preliminary testing suggested that partial

currents in the one-node pCMFD iteration are unstable, so underrelaxation with the factor $\omega = 0.5$ is applied.

$$J^{+}_{ij,ref} = (1-\omega)J^{+}_{ij,old} + \omega \cdot \left\{ \frac{1}{2}\tilde{D}_{ij}(\phi_i - \phi_j) + \hat{D}_{ij}\phi_i \right\} \quad (8)$$

The relation in equation 2 is rearranged to yield the net current as a function of the incoming partial current and the box-average flux, as follows.

$$J^{net}_{IJ} = \left(\frac{1}{2}\tilde{D}_{ij} + \hat{D}_{ij} + \frac{\tilde{D}_{ij}^{\,2}}{2\tilde{D}_{ij} + 4\hat{D}_{ji}} \right)\phi_i - \left(1 + \frac{\tilde{D}_{ij}}{\tilde{D}_{ij} + 2\hat{D}_{ji}} \right) J^{-}_{ij} \quad (9)$$

Each coarse-mesh node is then solved as a fixed-source fine-mesh problem. The fine-mesh partial currents are updated using equation 8.

This procedure is then repeated until the k-value obtained in the pCMFD calculations converges.

3. Results

The iDTMC methodology, including one-node pCMFD acceleration for pFMFD calculations, was then tested on the MOX-1000 medium-sized core. The results were compared with those from the iDTMC methodology without one-node pCMFD acceleration, which was previously verified.

2.1 Simulation Design

The model core used in the present study is the MOX-1000 core, as described in the NEA SFR benchmarks [6], with some structural elements simplified. The core layout is presented in Figures 3 and 4.

The MOX-1000 core was analysed using iDTMC with a Monte Carlo simulation of 40 inactive cycles, of which 15 were skip cycles, and 1 active cycle. 100,000 particles were simulated each cycle. Note that iDTMC enables acquisition of MC-equivalent solution at the onset of fission source convergence, i.e., the very 1st active cycle.

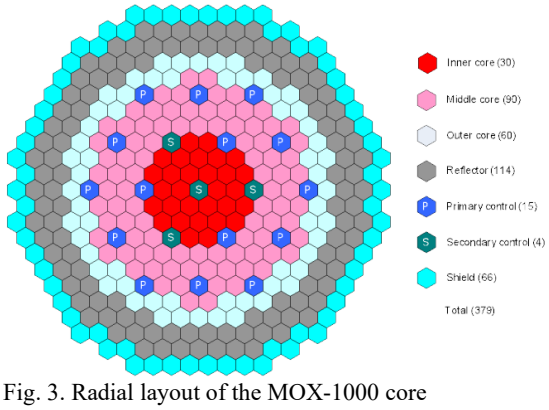

Fig. 3. Radial layout of the MOX-1000 core

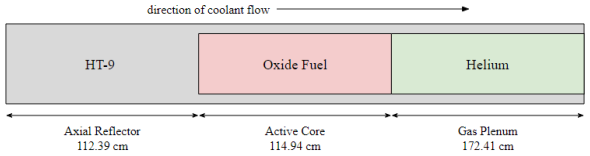

Fig. 4. Axial layout of the MOX-1000 core

The pFMFD calculations were performed twice, first naively and then with one-node pCMFD acceleration. To ensure the comparability of the two solutions, the Monte Carlo simulations were only performed once and the same set of tallied pFMFD factors was used in both pFMFD calculations.

The radial coarse-mesh node diameter was 16.2471 cm and the fine-mesh cell diameter was 0.913075 cm, equal to the assembly and fuel pin pitches.

The convergence criterion for both the naïve and accelerated pFMFD calculations was set at a relative difference of 0.01 pcm between successive source iterations.

2.2 Tests for Convergence Stability

As mentioned in the Methodology section, it was discovered during preliminary testing that the one-node pCMFD as applied in this context did not stably approach a converged solution. Thus, as shown in Equation 8, an underrelaxation factor was introduced to partial current updates. Sensitivity tests were performed to determine the optimal underrelaxation factor.

The difference between the neutron multiplication factor calculated in each one-node pCMFD iteration and that calculated in the preceding iteration was plotted against the number of iterations performed.

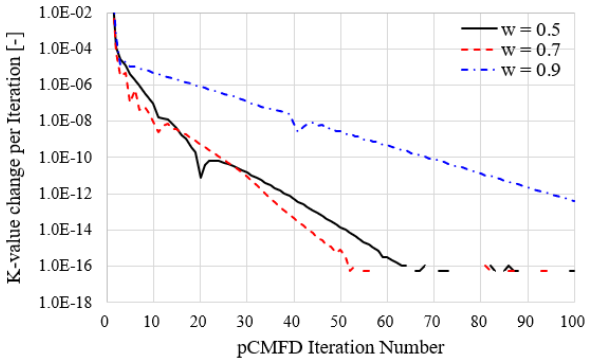

Fig. 5. Convergence in the one-node pCMFD scheme, with underrelaxation factors of 0.9 (blue), 0.7 (red), and 0.5 (black), respectively; the scale is log-linear

The difference in the convergence between using an underrelaxation factor of 0.5 and 0.7 was small, whereas convergence with an underrelaxation factor of 0.9 was clearly inferior. Thus, for all subsequent calculations, an underrelaxation factor of 0.5 was used. Note that the underrelaxation factor is likely problem-dependent, but the optimal value is usually close to 0.5; as this parameter is not of major importance, no further discussion is provided.

The rapid convergence in the neutron multiplication factor using underrelaxation is depicted in Figure 6. Note that there is very little change in the calculated k-value after about 20 iterations.

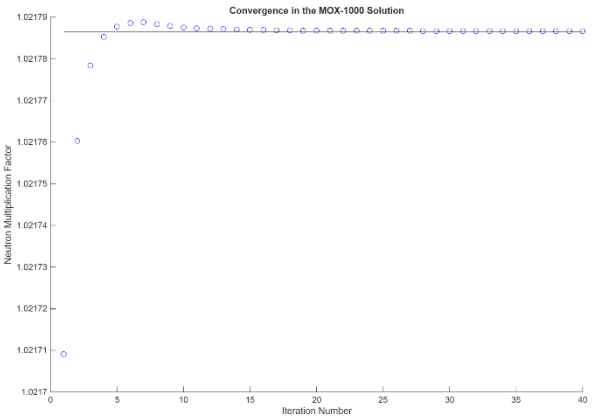

Fig. 6. Convergence of the calculated neutron multiplication factor with an underrelaxation factor of 0.5

2.3 Reproduction of the Existing Solution

The correctness of the one-node pCMFD acceleration scheme was verified by comparing the neutron multiplication factors and the axially integrated pin power distributions obtained by the naïve pFMFD calculation and the pFMFD calculation accelerated with one-node pCMFD.

Table 1: Calculated Neutron Multiplication Factors

Method	k-value
pFMFD (naïve)	1.0216022
pFMFD (accelerated)	1.0216030

The discrepancy of 0.08 pcm is easily explainable as an artefact of the convergence condition. The two values are effectively identical.

The calculated axially integrated pin power distributions are shown in Figures 7 to 9.

As shown in Figure 7, the two solutions are equivalent, with a maximum relative difference in the axially integrated pin power of 0.012%.

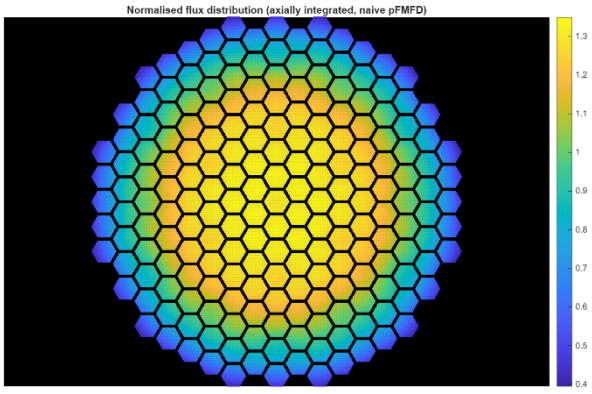

Fig. 7. Normalised and axially integrated iDTMC flux distribution, naïve pFMFD.

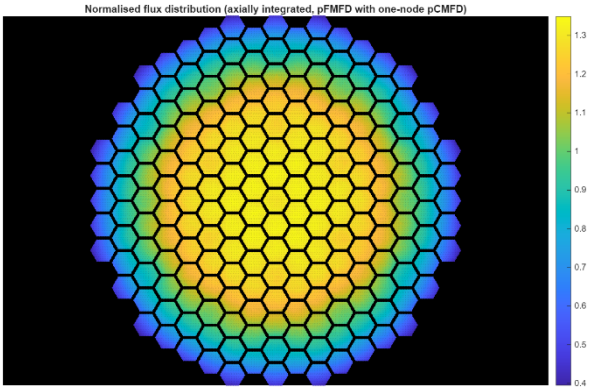

Fig. 8. Normalised and axially integrated iDTMC flux distribution, one-node pCMFD accelerated pFMFD

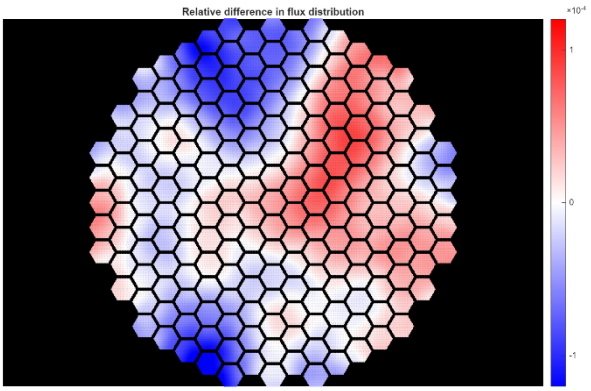

Fig. 9. Relative difference between the naïve pFMFD and accelerated pFMFD flux distributions; the colour scale extends from -0.012% to +0.012%.

Thus, it may be concluded that the one-node pCMFD acceleration process does not introduce any inaccuracies to the underlying iDTMC-pFMFD solution.

2.4 Reduction in the Serial Computing Cost

The Monte Carlo calculations have been performed on five cluster computing nodes with 40 CPUs each. The computing times of the different parts of the iDTMC calculations are given below.

Table 2: Computing Time for Components of iDTMC

Component	Clock Time (s)
Monte Carlo	5029.4
pFMFD (naïve)	76.6
pFMFD (accelerated)	10.6

It is clear that the one-node pCMFD acceleration scheme drastically reduces the computing cost of obtaining a pFMFD subspace solution within the iDTMC methodology. The computing time for the pFMFD calculation is no longer a significant component of the overall computing time of the iDTMC methodology.

The accelerated calculation time consists of the following components. Note that their sum is less than the total computing time shown in Table 2, as it does not

account for the computing cost of the time measurements themselves and other sundries.

Table 3: Components of the Accelerated pFMFD Calculation

Component	Clock Time (s)
Homogenisation, pCMFD	0.858
Flux/current modulation	2.205
Partial currents as fixed sources	2.464
Local fine-mesh problems	2.662
New currents, convergence test	2.151

2.5 Parallelisation

All components of the one-node pCMFD acceleration, except for the pCMFD acceleration itself, are calculations performed independently for each cell or surface. Thus, they are parallelisable by assigning blocks of cells or surfaces to each available CPU. From Table 3, it is evident that the pCMFD calculations account for only a small proportion of the overall computing cost of the one-node pCMFD method. Thus, significant reductions in computing time may be expected from parallelisation.

Parallelisation was implemented using OpenMP to use up to the number of CPUs in one cluster computing node, which, in the computing architecture used, was 40 CPUs. The corresponding reductions in computing time are depicted in Tables 4 and 5.

For convenience, the components of the accelerated pFMFD calculation are labelled parts 1 to 5, with each part corresponding to one row in Table 3.

Table 4: Computing Time of Each Component (s)

CPUs	Part 1	Part 2	Part 3	Part 4	Part 5
1	0.858	2.205	2.464	2.662	2.151
2	0.451	1.124	1.277	1.336	1.100
5	0.214	0.590	0.608	0.542	0.606
10	0.131	0.319	0.369	0.280	0.324
20	0.102	0.192	0.281	0.177	0.174
40	0.079	0.093	0.207	0.116	0.095

Table 5: Parallelisation Efficiency of Each Component (%)

CPUs	Part 1	Part 2	Part 3	Part 4	Part 5
2	95.1	98.0	96.4	99.6	97.8
5	80.3	74.8	81.1	98.3	71.0
10	65.4	69.1	66.8	95.1	66.3
20	42.1	57.3	43.8	75.1	61.4
40	27.2	59.1	29.8	57.6	56.4

While the parallelisation efficiencies are far from optimal, a dramatic decrease in the computing time as more CPUs are used is still clearly evident, from 10.6 seconds total in serial to 0.7 seconds using 40 CPUs.

As expected, Part 1, containing the pCMFD calculations themselves that cannot be cleanly parallelised, has the worst parallelisation efficiency at 40 cores. It is, however, surprising that the efficiency of Part 3, the conversion of incoming partial currents into equivalent fixed sources, is similarly low.

4. Conclusions and Future Work

When performing pFMFD calculations as a part of obtaining an iDTMC solution for a model SFR problem, using one-node pCMFD produced the same solution at a much lower computing cost compared to naïve pFMFD. Implementing parallelisation to take advantage of the independence of many calculations within the one-node pCMFD method produced further significant reductions in the computing time. Thus, we conclude that one-node pCMFD acceleration is a good method for reducing the deterministic computing burden of the iDTMC analysis of SFRs. At a computation time of less than 1 second per pFMFD solution using both one-node pCMFD acceleration and parallelisation with 40 CPUs, the deterministic computational burden of the iDTMC methodology should be manageable, even if performing many pFMFD calculations for uncertainty estimation.

Moving into the future, the generalisability of these conclusions should be verified by testing on other reference core models. For performing a large number of pFMFD calculations for uncertainty estimation, the effectiveness of the one-node pCMFD acceleration explored here and the naive parallelisation approach of assigning different pFMFD calculations to different CPUs should also be compared.

ACKNOWLEDGEMENTS

This research was supported by Korea Energy Technology Evaluation and Planning (KETEP) grant funded by the Korean Government (MTIE) (RS-2024-00439210).

REFERENCES

[1] I. Kim, I. Kim, and Y. Kim, “An iDTMC-based Monte Carlo depletion of a 3D SMR with intra-pin flux renormalization”, Frontiers in Energy Research, vol. 10, 2022.
[2] S. Yoon, Y. Kim, “Bias Analysis for the Application of iDTMC to Sodium-Cooled Fast Reactors”, Korean Nuclear Society Autumn Meeting, Gyeongju, Republic of Korea, 2023.
[3] S. Yoon, I. Kim, T. Oh, and Y. Kim. “Uncertainty Estimation by Correlated Sampling for the iDTMC Analysis of Sodium-Cooled Fast Reactors”, Korean Nuclear Society Autumn Meeting, Changwon, Republic of Korea, 2024.
[4] S. Yoon, I. Kim, T. Oh, and Y. Kim. “Parallelization of Correlated Sampling Uncertainty Estimation in the iDTMC (Improved Deterministic Truncation of Monte Carlo) Analysis of Sodium-cooled Fast Reactors”, International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (M&C), 2025.
[5] H. Kim and Y. Kim, “Pin-by-Pin Coupled Transient Monte Carlo Analysis Using the iMC Code”, Frontiers in Energy Research, 2022.
[6] “Benchmark for Neutronic Analysis of Sodium-cooled Fast Reactor Cores with Various Fuel Types and Core Sizes.” Nuclear Energy Agency, NEA/NSC/R(2015)9, 2016.