VANGARD: A GPU-BASED HIGH-SPEED PINWISE NODAL CORE ANALYSIS CODE

SEOYOON JEON, HYUNSIK HONG, NAMJAE CHOI, HAN GYU JOO

Seoul National University, Seoul, Republic of Korea

Email contact of corresponding author: joohan@snu.ac.kr

VANGARD (**V**ersatile **A**dvanced **N**eutronics Code for **G**PU-**A**ccelerated **R**eactor **D**esigns) is a pioneering pinwise nodal core analysis code being developed at Seoul National University (SNU). VANGARD targets to realize and spread next-generation pinwise two-step core design procedure in the industries. The industries often lack computing resources, and as the result, the conventional assembly-wise nodal core analysis that can be performed rapidly on personal computers (PC) with minimal computing resources has been employed as the standard core design procedure. Noting this, VANGARD adopts the rising GPU computing technologies and is highly optimized to exploit affordable gaming GPUs which are typically mounted on PCs, targeting to retain practicality and achieve feasible pinwise core analysis time on PCs.

VANGARD is developed in C++ and the GPU-accelerated portions are written in NVIDIA CUDA. Most parts of VANGARD are programmed to be executed on GPUs as well as on CPUs including major computational hotspots: pinwise nodal solver, cross section treatment module, and depletion solver. Namely, VANGARD can be run on both multi-core CPU-based and GPU-based computing platforms. The key features of VANGARD are as follows:

(1) Pinwise nodal solution with the Simplified P_3 (SP_3) Source Expansion Nodal Method (SENM)
(2) Coarse Mesh Finite Difference (CMFD) acceleration with partial incoming current update
(3) Thermal-hydraulics (T/H) feedback with one-dimensional single-phase closed-channel model
(4) Depletion calculation with the Chebyshev Rational Approximation Method (CRAM)
(5) Group constant data compression employing Singular Value Decomposition (SVD) and Low Rank Approximation (LRA)
(6) Restart/shuffling capability for multi-cycle calculations

Pinwise Nodal Solution: SENM [1] is superior to other nodal kernels in capturing the steep flux gradients, which occur in the low energy groups due to the intra-assembly heterogeneities and the intra-pin flux shapes, with the use of hyperbolic terms in the flux expansion. Therefore, after a thorough and extensive sensitivity study on the choice of the solver, SP_3 SENM with low-order expansion was chosen as the primary solver in VANGARD [2][3]. Although the SENM kernel can be expanded up to 4^{th} order, 2^{nd} order expansion is applied for the radial direction because applying 4^{th} order expansion for the pin-sized fine mesh leads to stability problems. On the other hand, axial solution is obtained with full 4^{th} order expansion because axial mesh is relatively large. In this manner, kernels of different orders are adaptively used in different directions for efficient while accurate pinwise nodal core calculations.

CMFD Acceleration: 4-box assembly-mesh CMFD acceleration is employed in VANGARD by default. The linear systems are solved with BiCGSTAB using the Eigen linear algebra package [4]. After a CMFD power iteration, the updated coarse mesh flux is fed back to the nodal solver

by adjusting the level of the pinwise flux. However, this causes inconsistency between the pinwise flux and the pin incoming currents that cannot be updated from the coarse mesh CMFD calculation, which leads to an instability and deteriorates the convergence. To resolve this, the p-CMFD [5] coefficient for the outgoing partial current in Eq. (1) is additionally defined for each pin, and after the CMFD calculation, pin outgoing currents are recalculated using the modulated pinwise flux and the p-CMFD relation. Namely, the CMFD calculation is still performed on the coarse mesh basis with the standard formulation, but pinwise p-CMFD coefficients are additionally defined to be used for updating pin outgoing currents after the coarse mesh CMFD calculation.

$$J^+ = \frac{-\tilde{D}(\phi_{i+1} - \phi_i) + 2\hat{D}^+\phi_i}{2} \quad \text{where} \quad \hat{D}^+ = \frac{2J^+_{nodal} + \tilde{D}(\phi_{i+1} - \phi_i)}{2\phi_i}. \tag{1}$$

T/H Feedback: A simple yet reasonably accurate T/H solver based on the single-phase closed-channel model is embedded in VANGARD. In this model, the fuel rod heat transfer is solved rather precisely; one-dimensional cylindrical finite difference heat conduction calculation is performed for each axial slice of a pin using high-fidelity material properties, namely the thermal conductivity correlations and the gap conductance table acquired from the fuel performance code FRAPCON [6] and ROPER [7], respectively. However, the coolant treatment is simplified. As the name of the model implies, cross flows are neglected, and mass flux and pressure are assumed constant. Only the enthalpy conservation is solved to determine the coolant temperatures. Instead, to compensate the lack of cross flow mixing, pinwise flow channels are merged into macro-channels (2 × 2 per assembly by default) and treated in an average sense. Additionally, the Anderson acceleration [8] is applied to stabilize the convergence of the fixed-point iteration.

Depletion: The default depletion chain of VANGARD considers 44 nuclides including 15 fission products and 4 fissiles (U-235, U-238, Pu-239, Pu-241). For the solution of the matrix exponential occurring in the Bateman equation, CRAM [9] is employed as shown in Eq. (2). The solution of Eq. (2) is obtained by reformulating the matrix inversion into a linear system and applying the Gauss-Seidel method, instead of directly calculating the inverse.

$$\vec{\mathbf{N}} = \alpha_0 \vec{\mathbf{N}}_0 + 2\,\mathrm{Re}\left(\sum_{j=1}^{k/2} \alpha_j \left(\mathbf{A}t - \theta_j \mathbf{I} \right)^{-1} \vec{\mathbf{N}}_0 \right). \tag{2}$$

In case of gadolinium, which is often used as burnable poisons and shows extremely heterogeneous depletion behavior, a special lumped model is used for treating its depletion. As the total amount of gadolinium isotopes stays almost constant during depletion, the gadolinium isotopes are lumped into a single effective isotope, which is not changed by the depletion solver. Instead, the cross section of the effective isotope functionalized with burnup in the group constant takes responsibility of reflecting the effect of isotopic changes of gadolinium during depletion. Another depletion phenomenon, namely the depletion of B-10 in the reactor coolant system (RCS), is also treated separately.

Group Constants: VANGARD employs 8-group pinwise microscopic group constants generated from the direct whole-core transport calculation code nTRACER [10], and 30 nuclides including a lumped nuclide are considered. The major downside of pinwise two-step calculation is that the

size of group constant data often becomes too large. As in the assembly-wise two-step procedure, the pinwise group constants are generated from the single assembly lattice calculations at various burnup and branch points and the cross sections are pre-tabulated as a function of fuel temperature, moderator temperature, moderator density, boron concentration, and burnup. Since the tabulation should be made pinwise, the amount of data becomes tremendous. Thus, to reduce the storage and memory requirements of using pinwise group constants, which may impose difficulties in running VANGARD on PCs and GPUs, a cross section compression scheme employing dimensionality reduction technique based on SVD and LRA [11] is adopted. In addition, the scattering matrices are generated in macroscopic form to further reduce memory usage as the scattering matrices are dominated by light nuclei and thus not sensitive to the number density changes of heavy nuclei.

The analysis capability, accuracy, and performance of VANGARD are demonstrated here with the cycle 1 depletion of the BEAVRS benchmark problem [12]. The calculation results were compared with those of nTRACER. The core was depleted up to 13 MWD/kgHM, which yields 17 burnup steps in total including the beginning-of-cycle (BOC) step. *FIG. 1* shows three-dimensional power distributions at BOC, middle-of-cycle (MOC), and end-of-cycle (EOC) calculated by VANGARD, and *FIG. 2* to *FIG. 4* illustrate the changes in the inventories of some important actinides in time.

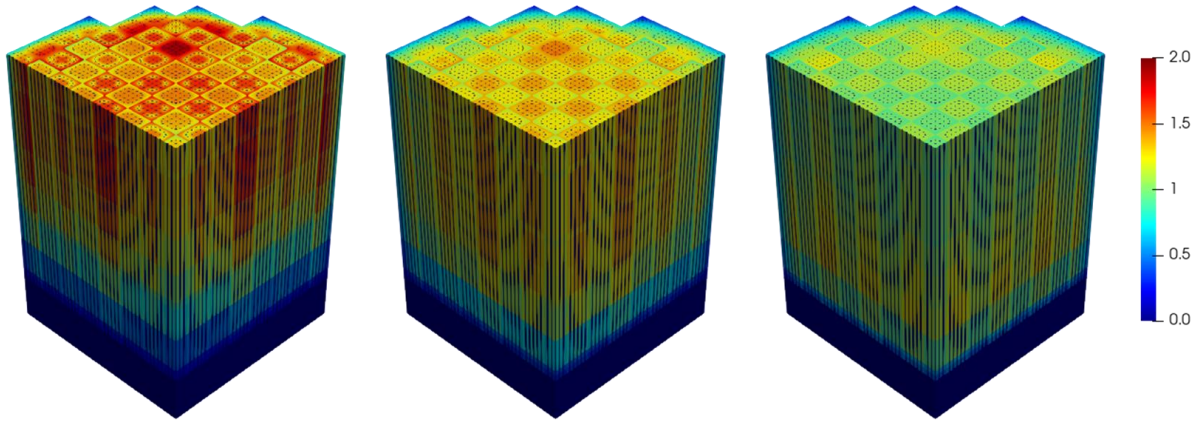

FIG. 1. Power distributions at BOC (left), MOC (middle), and EOC (right).

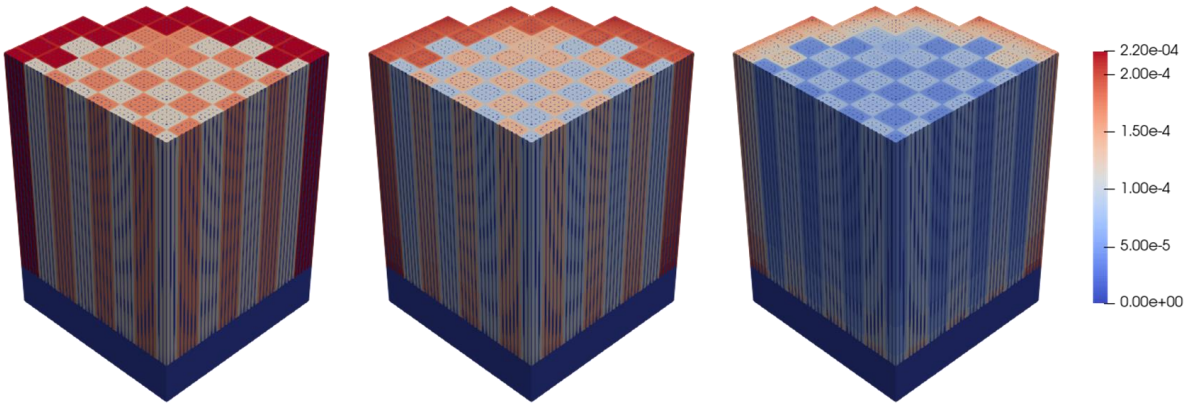

FIG. 2. Inventories of U-235 at BOC (left), MOC (middle), and EOC (right).

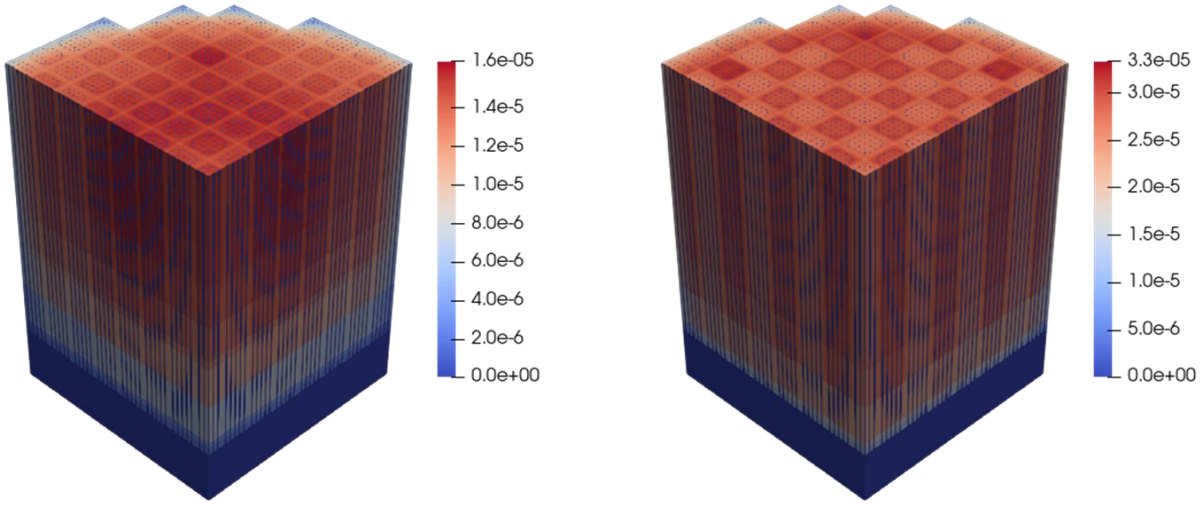

FIG. 3. Inventories of Pu-239 at MOC (left) and EOC (right).

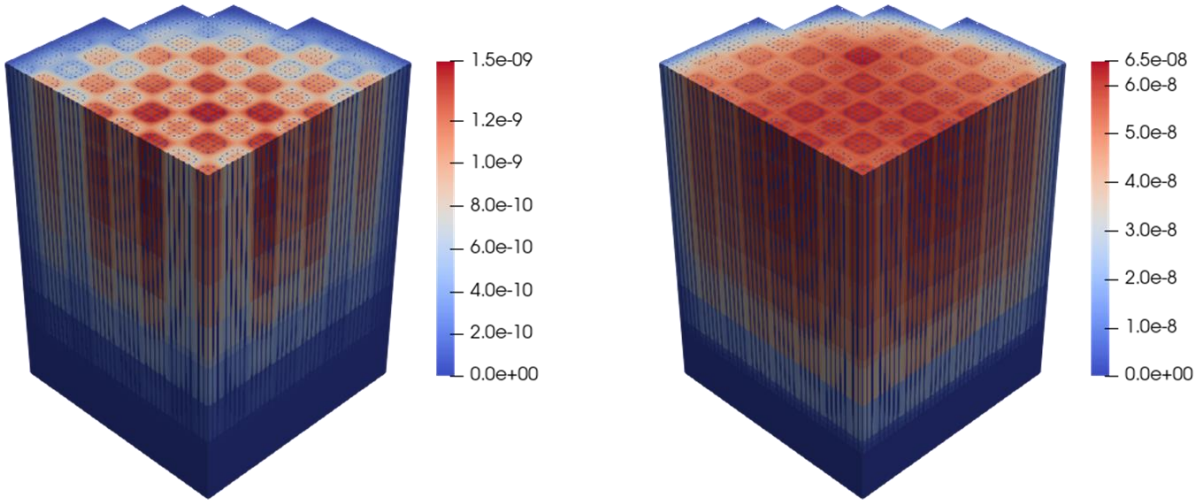

FIG. 4. Inventories of Am-241 at MOC (left) and EOC (right).

In terms of the accuracy, VANGARD shows excellent agreement with the reference whole-core transport solution of nTRACER. *FIG. 5* demonstrates the critical boron concentration (CBC) letdown curves of VANGARD and nTRACER, and *FIG. 6* and *FIG. 7* illustrate the maximum and RMS errors of pin and axial powers at each burnup step, respectively. Additionally, *FIG. 8* and *FIG. 9* present the pin and axial power error distributions at BOC, MOC, and EOC, respectively. For most of the burnup steps, the errors in the CBCs are kept below 6 ppm and the maximum and RMS errors in the pin and axial powers do not exceed 1.5% and 0.5%, respectively.

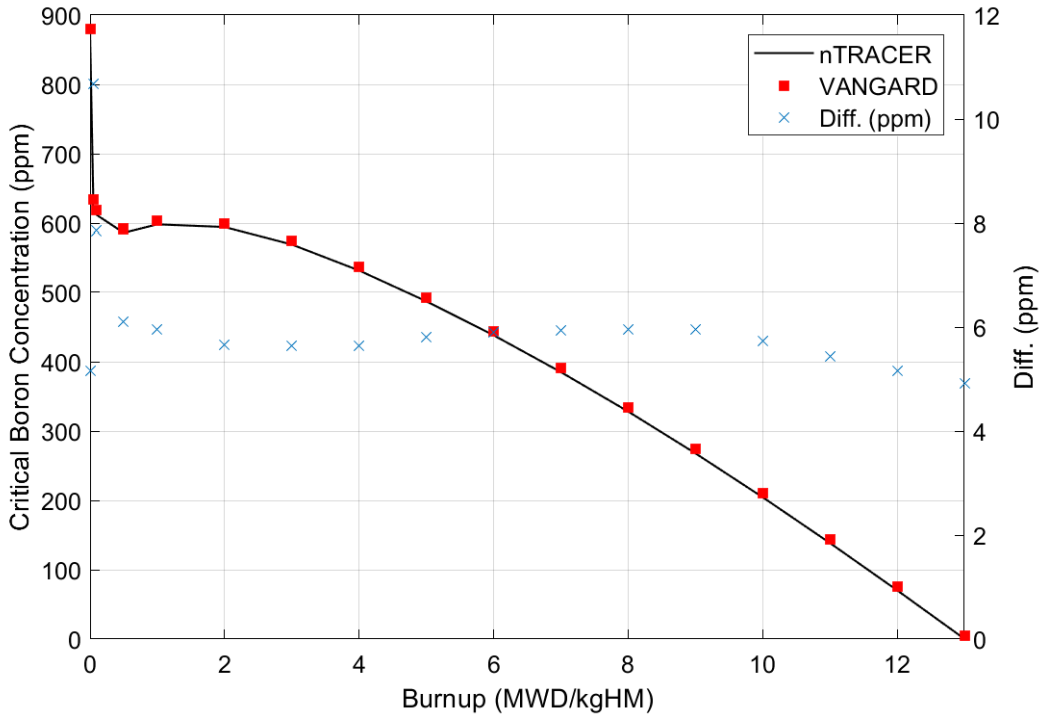

FIG. 5. Comparison of boron letdown curves.

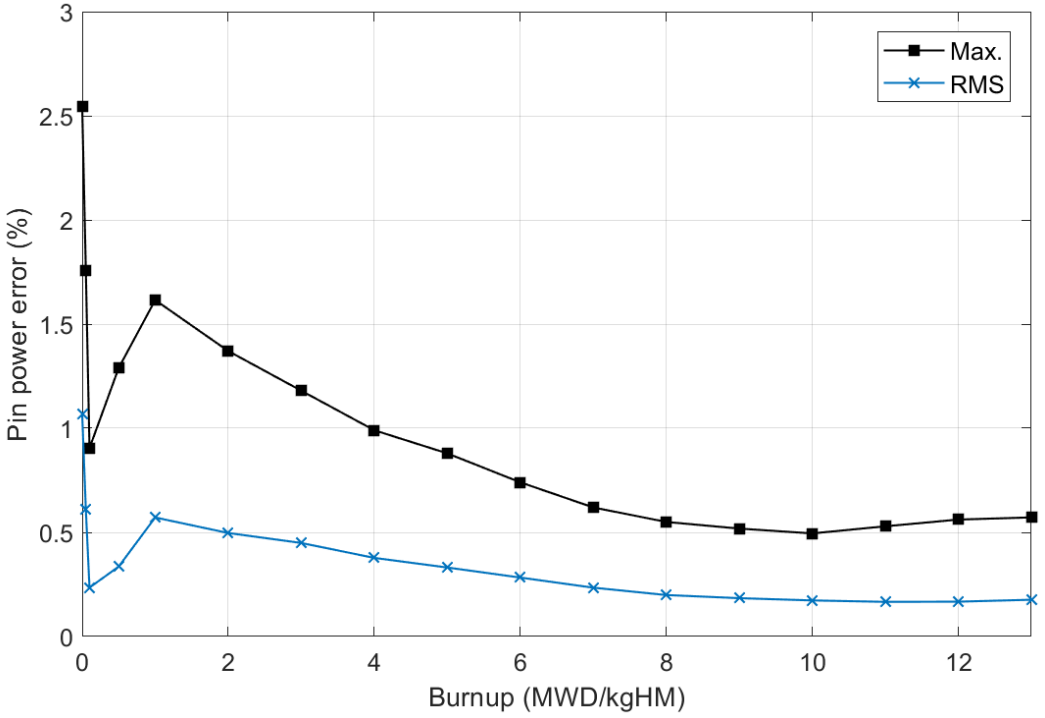

FIG. 6. Trend of pin power errors during depletion.

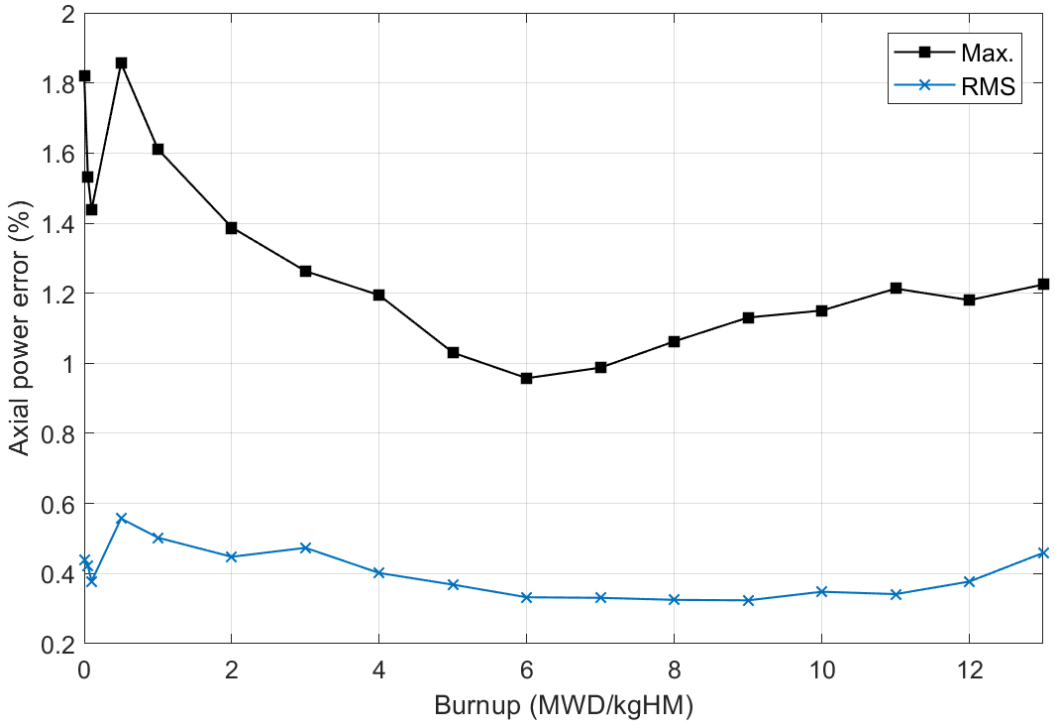

FIG. 7. Trend of axial power errors during depletion.

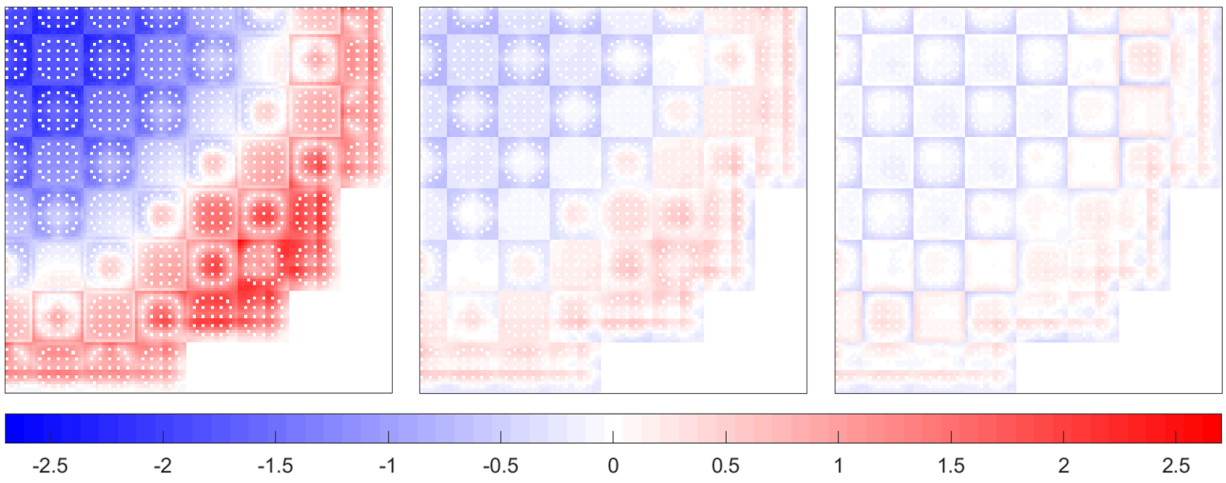

FIG. 8. Pin power error distributions (%) at BOC (left), MOC (middle), and EOC (right).

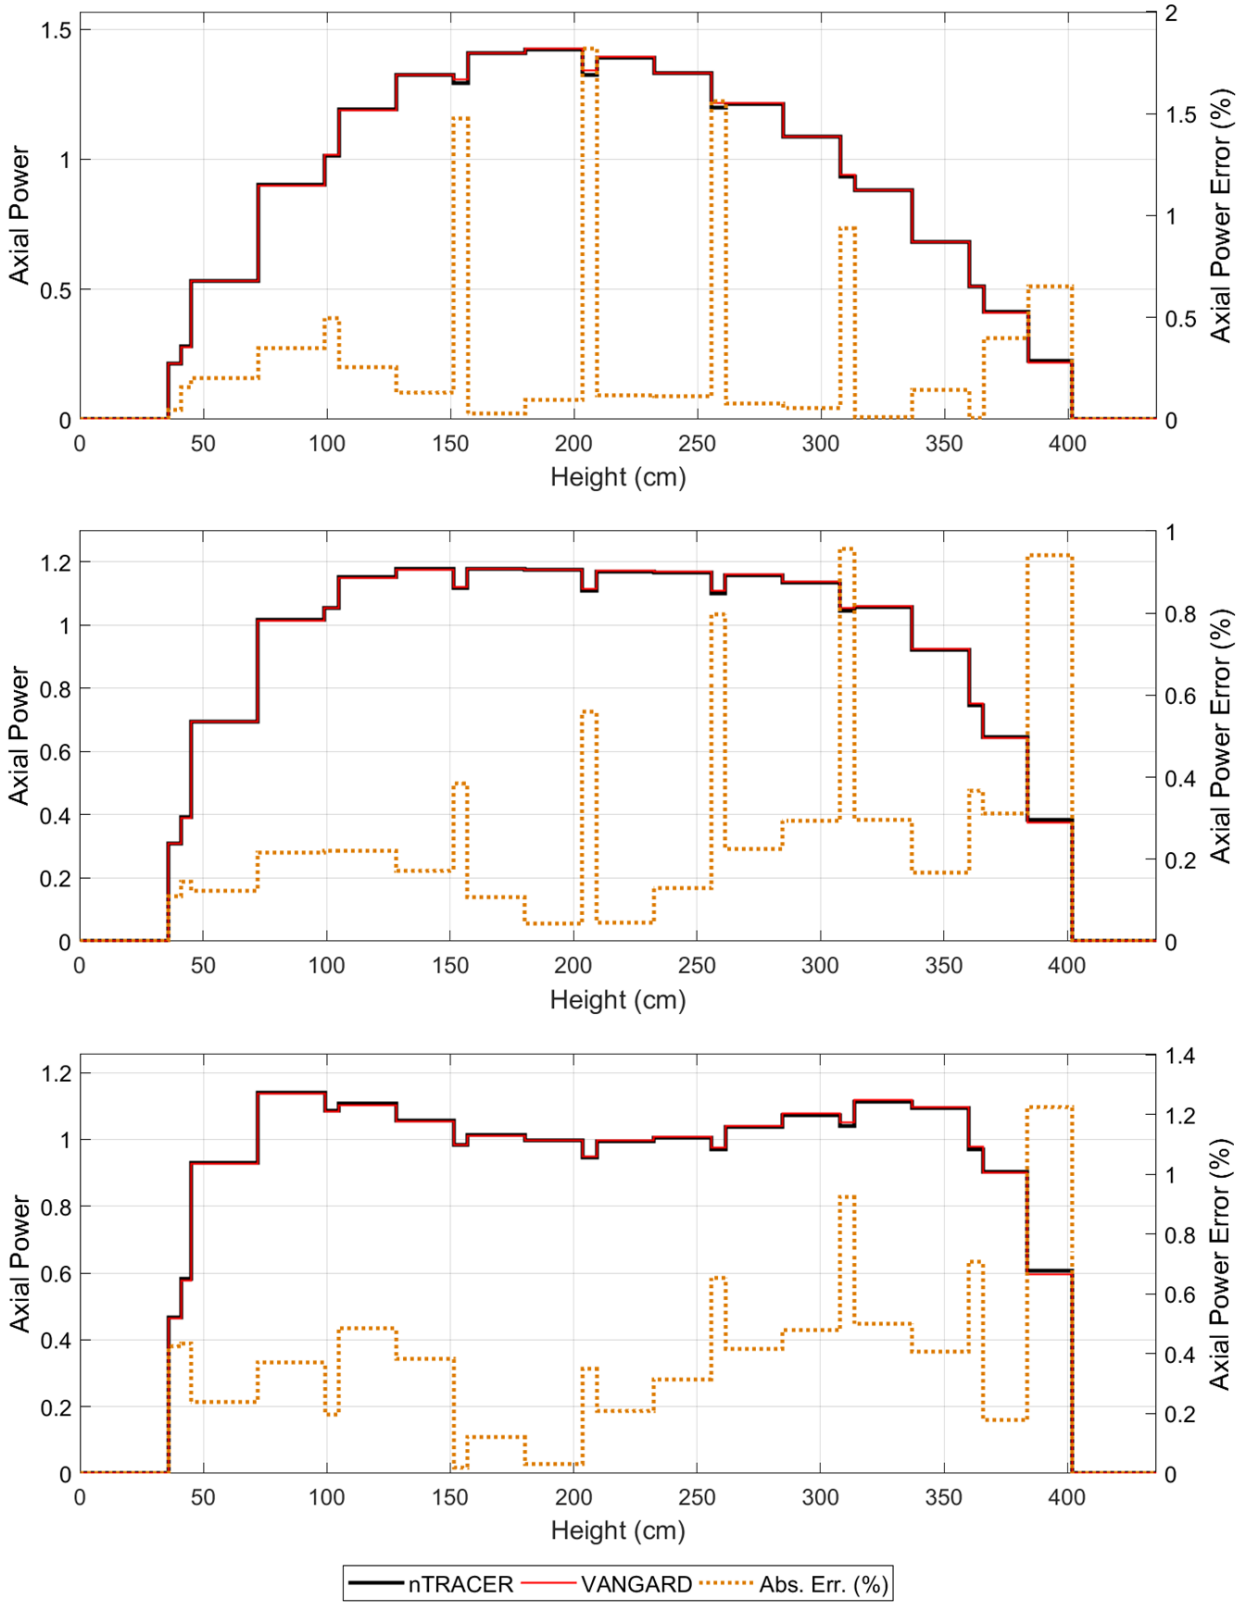

FIG. 9. Comparison of axial power distributions at BOC (top), MOC (middle), and EOC (bottom).

In terms of the computing performance, VANGARD satisfies the industrial requirement level. This is shown by the performance comparison of multi-core CPU and GPU calculations. For the CPU calculation, two Intel Xeon E5-2630 v4 processors which yield 20 cores in total were employed, and a single NVIDIA GeForce RTX 3090 was used for the GPU calculation. The CPU calculation is parallelized with OpenMP. *FIG. 10* illustrates the time shares of solvers in the CPU and GPU calculations, and TABLE 1 summarizes the computing times and speedup ratios. The three major hotspots account for nearly 97% of the total computing time in the CPU calculation, which is reduced to less than 70% by the GPU acceleration. Owing to the significant speedups achieved in the hotspots, the total computing time is reduced from about 45 minutes to less than 3 minutes, which is sufficiently fast to perform repeated core calculations required for nuclear design.

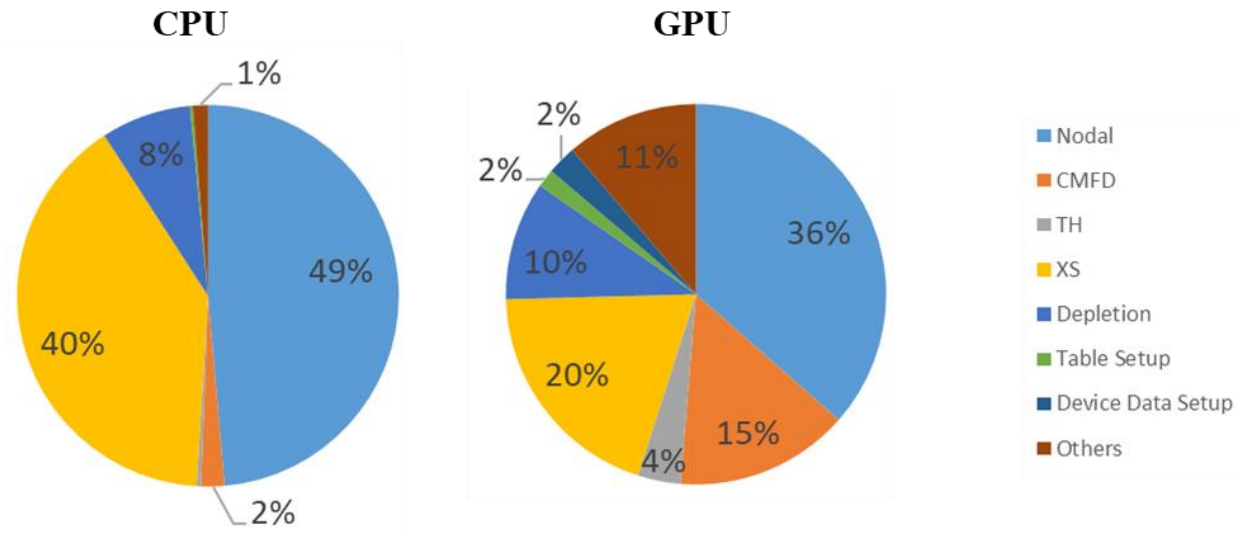

FIG. 10. Time shares of solvers in CPU and GPU calculations.

TABLE 1. COMPUTING TIMES AND SPEEDUP RATIOS

Calculation	CPU (s)	GPU (s)	Speedup
Nodal	1303.3	60.2	21.7
Cross Section	1069.9	32.4	33.0
Depletion	204.4	16.6	12.3
Total	2680.6	164.9	16.3

REFERENCES

[1] J. I. Yoon, H. G. Joo, "Two-Level Coarse Mesh Finite Difference Formulation with Multigroup Source Expansion Nodal Kernels," *Journal of Nuclear Science and Technology* **45**(7), pp. 668-682 (2008).

[2] H. Hong, H. G. Joo, "Source Expansion Nodal Kernel for Multi-Group Pin-by-Pin SP3 Core Calculation," *Korean Nuclear Society Spring Meeting*, Jeju, Korea, June 9-10 (2020).

[3] H. Hong, H. G. Joo, "Thorough Analyses and Resolution of Various Errors in Pin-Homogenized Multigroup Core Calculation," *Annals of Nuclear Energy* (In press).

[4] Eigen3 API Documentation. https://eigen.tuxfamily.org/dox/index.html.
[5] N. Z. Cho, G. S. Lee, “Comparison of CMFD and p-CMFD Acceleration Methods for Neutron Transport Calculations,” *Korean Nuclear Society Spring Meeting*, Gyeongju, Korea, May (2003).
[6] W. Luscher, K. Geelhood, I. Porter, “Material Property Correlations: Comparison between FRAPCON 4.0, FRAPTRAN 2.0, and MATPRO," Pacific Northwest National Laboratory, *PNNL-19147 Rev.2* (2015).
[7] J. Choi, Y. Heo, H. Han, “A Fuel Rod Performance Analysis and Design Code – ROPER,” *Korean Nuclear Society Spring Meeting*, Jeju, Korea, May 22 (2009).
[8] A. Facchini, J. Lee, H. G. Joo, “Investigation of Anderson Acceleration in Neutronics-Thermal Hydraulics Coupled Direct Whole Core Calculation,” *Annals of Nuclear Energy* **153** (2021).
[9] M. Pusa, “Higher-Order Chebyshev Rational Approximation Method and Application to Burnup Equations,” *Nuclear Science and Engineering* **182**(3), pp. 297-318 (2016).
[10] Y. S. Jung, C. B. Shim, C. H. Lim, H. G. Joo, “Practical Numerical Reactor Employing Direct Whole Core Neutron Transport and Subchannel Thermal/Hydraulics Solvers,” *Annals of Nuclear Energy* **62**, pp. 357-374 (2013).
[11] M. Yamamoto, T. Endo, A. Yamamoto, “Compression of Cross-Section Data Size for High-Resolution Core Analysis Using Dimensionality Reduction Technique,” *Nuclear Science and Engineering* **195**, pp. 33-49 (2021).
[12] N. Horelik, B. Herman, B. Forget, K. Smith, “Benchmark for Evaluation and Validation of Reactor Simulations (BEAVRS),” *International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering*, Sun Valley, Idaho, United States, May 5-9 (2013).