

Multiphysics methods and analysis applied to nuclear reactor systems

Edited by

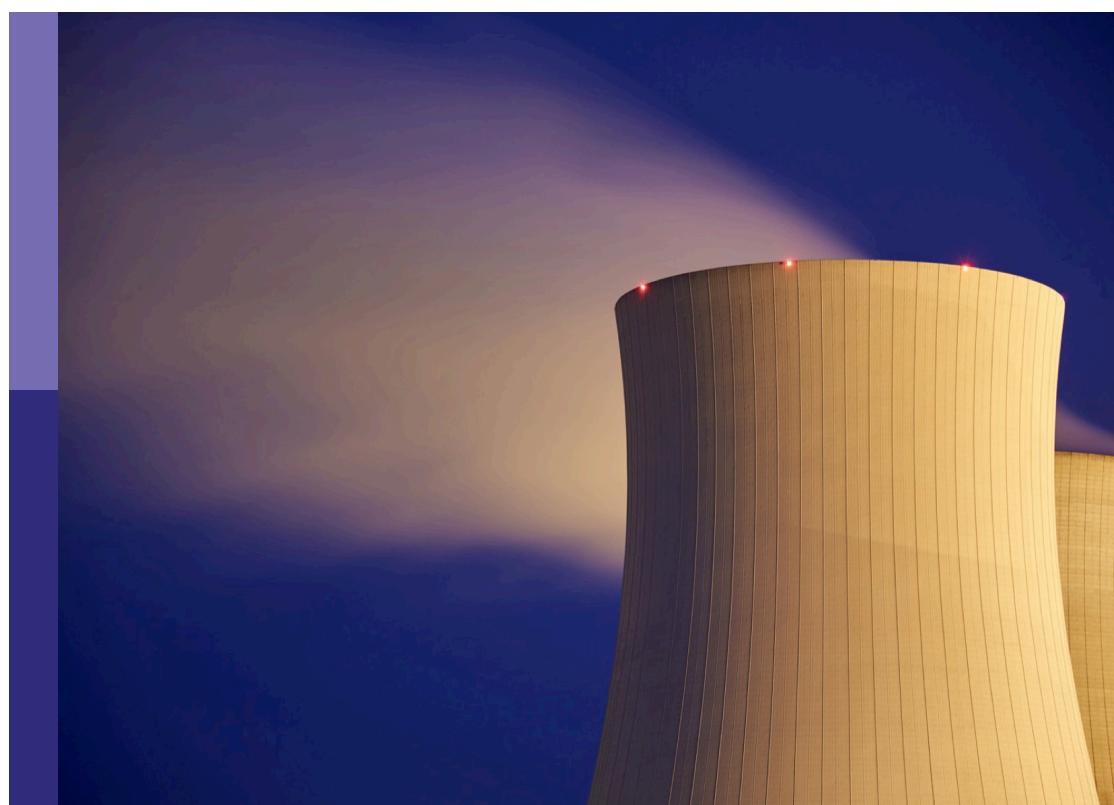
Mark D. DeHart, Deokjung Lee and Kirk Atkinson

Coordinated by

Emily Shemon and Julien Taforeau

Published in

Frontiers in Nuclear Engineering



FRONTIERS EBOOK COPYRIGHT STATEMENT

The copyright in the text of individual articles in this ebook is the property of their respective authors or their respective institutions or funders. The copyright in graphics and images within each article may be subject to copyright of other parties. In both cases this is subject to a license granted to Frontiers.

The compilation of articles constituting this ebook is the property of Frontiers.

Each article within this ebook, and the ebook itself, are published under the most recent version of the Creative Commons CC-BY licence. The version current at the date of publication of this ebook is CC-BY 4.0. If the CC-BY licence is updated, the licence granted by Frontiers is automatically updated to the new version.

When exercising any right under the CC-BY licence, Frontiers must be attributed as the original publisher of the article or ebook, as applicable.

Authors have the responsibility of ensuring that any graphics or other materials which are the property of others may be included in the CC-BY licence, but this should be checked before relying on the CC-BY licence to reproduce those materials. Any copyright notices relating to those materials must be complied with.

Copyright and source acknowledgement notices may not be removed and must be displayed in any copy, derivative work or partial copy which includes the elements in question.

All copyright, and all rights therein, are protected by national and international copyright laws. The above represents a summary only. For further information please read Frontiers' Conditions for Website Use and Copyright Statement, and the applicable CC-BY licence.

ISSN 1664-8714
ISBN 978-2-8325-7106-4
DOI 10.3389/978-2-8325-7106-4

Generative AI statement

Any alternative text (Alt text) provided alongside figures in the articles in this ebook has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

About Frontiers

Frontiers is more than just an open access publisher of scholarly articles: it is a pioneering approach to the world of academia, radically improving the way scholarly research is managed. The grand vision of Frontiers is a world where all people have an equal opportunity to seek, share and generate knowledge. Frontiers provides immediate and permanent online open access to all its publications, but this alone is not enough to realize our grand goals.

Frontiers journal series

The Frontiers journal series is a multi-tier and interdisciplinary set of open-access, online journals, promising a paradigm shift from the current review, selection and dissemination processes in academic publishing. All Frontiers journals are driven by researchers for researchers; therefore, they constitute a service to the scholarly community. At the same time, the *Frontiers journal series* operates on a revolutionary invention, the tiered publishing system, initially addressing specific communities of scholars, and gradually climbing up to broader public understanding, thus serving the interests of the lay society, too.

Dedication to quality

Each Frontiers article is a landmark of the highest quality, thanks to genuinely collaborative interactions between authors and review editors, who include some of the world's best academicians. Research must be certified by peers before entering a stream of knowledge that may eventually reach the public - and shape society; therefore, Frontiers only applies the most rigorous and unbiased reviews. Frontiers revolutionizes research publishing by freely delivering the most outstanding research, evaluated with no bias from both the academic and social point of view. By applying the most advanced information technologies, Frontiers is catapulting scholarly publishing into a new generation.

What are Frontiers Research Topics?

Frontiers Research Topics are very popular trademarks of the *Frontiers journals series*: they are collections of at least ten articles, all centered on a particular subject. With their unique mix of varied contributions from Original Research to Review Articles, Frontiers Research Topics unify the most influential researchers, the latest key findings and historical advances in a hot research area.

Find out more on how to host your own Frontiers Research Topic or contribute to one as an author by contacting the Frontiers editorial office: frontiersin.org/about/contact

Multiphysics methods and analysis applied to nuclear reactor systems

Topic editors

Mark D. DeHart — Abilene Christian University, United States
Deokjung Lee — Ulsan National Institute of Science and Technology, Republic of Korea
Kirk Atkinson — Ontario Tech University, Canada

Topic coordinators

Emily Shemon — Argonne National Laboratory (DOE), United States
Julien Taforeau — Institut de Radioprotection et de Sécurité Nucléaire, France

Citation

DeHart, M. D., Lee, D., Atkinson, K., Shemon, E., Taforeau, J., eds. (2025). *Multiphysics methods and analysis applied to nuclear reactor systems*. Lausanne: Frontiers Media SA. doi: 10.3389/978-2-8325-7106-4

Table of contents

04 Editorial: Multiphysics methods and analysis applied to nuclear reactor systems
Mark D. DeHart, Emily Shemon and Deokjung Lee

07 On-the-fly thermal expansion for Monte Carlo multi-physics reactor simulations
Muhammad Imron and Deokjung Lee

15 Coupled serpent/subchanflow analysis with unstructured mesh interfaces for a hexagonal, plate-type VVR-KN fuel assembly
Gianfranco Huaccho Zavala, Thomas Gheeraert, Victor Hugo Sánchez-Espinoza, Juan Carlos Almachi and Uwe Imke

25 Coupled neutronic-thermal-mechanical simulation of the KRUSTY heat pipe microreactor
William Reed Kendrick and Benoit Forget

39 Improved high-fidelity multiphysics modeling of pulsed operation of the annular core research reactor
Emory Colvin and Todd S. Palmer

53 Coupling of the Monte Carlo iMC and OpenFoam codes for multiphysics calculations of molten salt reactors
Inyup Kim, Taesuk Oh and Yonghee Kim

64 Data transfers for nuclear reactor multiphysics studies using the MOOSE framework
Guillaume L. Giudicelli, Fande Kong, Roy Stogner, Logan Harbour, Derek Gaston, Alexander Lindsay, Zachary Prince, Lise Charlot, Stefano Terlizzi, Mahmoud Eltawila and April Novak

85 CTF development, verification, and validation for VVER core thermal-hydraulics and multi-physics modeling and simulation
Yesim Kutlu, Ivan Spasov, Svetlomir Mitkov, Pascal Rouxelin, Agustin Abarca, Nikola Kolev, Maria Avramova and Kostadin Ivanov

110 Uncertainty quantification and sensitivity analysis of a nuclear thermal propulsion reactor startup sequence
Jackson R. Harter and Mark D. DeHart

125 Hybrid solvers for reactor modelling: matrix-based and matrix-free approaches on voxel-dominated meshes
Liang Yang, Jianhui Yang, Jakov Popov and Andrew G. Buchan



OPEN ACCESS

EDITED AND REVIEWED BY

Fabio Giannetti,
Sapienza University of Rome, Italy

*CORRESPONDENCE

Mark D. DeHart,
✉ Mark.DeHart@acu.edu

RECEIVED 01 October 2025

ACCEPTED 13 October 2025

PUBLISHED 24 October 2025

CITATION

DeHart MD, Shemon E and Lee D (2025) Editorial: Multiphysics methods and analysis applied to nuclear reactor systems. *Front. Nucl. Eng.* 4:1717262. doi: 10.3389/fnuen.2025.1717262

COPYRIGHT

© 2025 DeHart, Shemon and Lee. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Editorial: Multiphysics methods and analysis applied to nuclear reactor systems

Mark D. DeHart^{1*}, Emily Shemon² and Deokjung Lee³

¹Nuclear Science and Engineering Program, Abilene Christian University, Abilene, TX, United States,

²Nuclear Science and Engineering Division, Argonne National Laboratory, Lemont, IL, United States,

³Nuclear Engineering Department, Korea Advanced Institute of Science and Technology, Daejeon, Republic of Korea

KEYWORDS

multiphysics modeling, nuclear reactor systems, simulation, thermal hydraulics, safety analysis, modeling, neutron transport

Editorial on the Research Topic

Multiphysics methods and analysis applied to nuclear reactor systems

The nuclear industry stands at a turning point. Demand for carbon-free energy sources is increasing, while aging reactors and the need for enhanced safety require new and innovative solutions. Central to this progress is multiphysics modeling and simulation, a relatively up-and-coming approach by which the design, analysis, and operation of nuclear reactor systems are being reexamined. This Research Topic explores the leading edge of this field, showcasing cutting-edge research and providing insights into the essential role that advanced simulations play in driving nuclear technologies forward.

Traditional nuclear engineering has relied heavily on experimentation, often involving costly and potentially hazardous procedures. Multiphysics simulations provide a robust framework for reactor experiment and full system evaluation. By integrating models that accurately represent the inherent coupling of physical phenomena: neutronics, thermal hydraulics, structural mechanics, heat transfer, etc., these simulations offer a complete view of reactor behavior under a wide set of conditions. This unified approach enables designers and analysts to optimize designs, predict performance, and enhance safety in ways previously unattainable.

With the wide variety of reactor types on the horizon, advanced multiphysics modeling and simulation offers clear advantages of the use of traditional physics codes and workflows. Generally, the tightly coupled nature of multiphysics methods allows for more flexible code application and deeper physics insights than what can be obtained with loosely or one-way coupled tools. This is especially important for exploring the technical merits of novel reactor concepts with less experimental pedigree.

This Research Topic highlights contributions that showcase the breadth and depth of multiphysics applications in nuclear engineering. From innovative computational methods to rigorous experimental validation, the articles presented here underscore the critical role that simulation plays in representing the complexities of modern nuclear reactor technology.

A notable advancement within this domain is described in [Imron and Lee](#) through the development of on-the-fly thermal expansion methodologies for multiphysics Monte Carlo reactor simulations as This method allows the problem geometry to dynamically expand

during particle tracking by incorporating local temperature data, such as pin-averaged temperatures obtained from thermal-hydraulics solvers. Numerical experiments demonstrate that modeling thermal expansion with local temperature data can significantly improve the accuracy of simulations, including eigenvalue predictions and pin power distributions, compared to models using only global core-averaged temperatures.

A central theme emerging from this Research Topic is the development and application of innovative computational methods. For instance, [Harter and DeHart](#) details the application of stochastic methods and sensitivity analysis to a full-core model of a nuclear thermal propulsion system. This research showcases the development of a reduced-order model that allows for rapid evaluation of system behavior under various input conditions, a critical advancement for optimizing the design and control of these complex reactors.

The integration and coupling of different physical models is another area of focus. Advanced nuclear reactor cores are governed by multiple physical phenomena which should all be resolved, and the coupling of these physics would also need to be resolved spatially in a high-fidelity approach. [Giudicelli et al.](#) presents field transfer capabilities implemented in the Multiphysics Object-Oriented Simulation Environment (MOOSE), and numerous technical details such as mapping heuristics, conservation techniques and parallel algorithms. In a similar vein, [Yang et al.](#) explores hybrid, matrix-based, and matrix-free solver technologies within a voxel-dominated Cartesian mesh framework, offering a novel approach to simulating neutronics and thermal hydraulics in nuclear reactor cores. This approach enables accurate boundary representation and efficient resolution of complex geometries. The coupling of different physical models, such as neutronics and thermal hydraulics, is crucial for accurate reactor analysis.

The application of these advanced simulations extends to design optimization, with several articles demonstrating how multiphysics models can inform and refine reactor designs. Presents the development of a multiphysics coupled framework of [Kim et al.](#), which provides significant insights into the analysis of MSRs. [Zavala et al.](#) discusses a high-detail steady-state analysis of one VVR-KN fuel assembly. The VVR-KN is a plate-type fuel assembly, with fuel elements arranged hexagonal with fuel-plate tubes that challenges both their neutronic and thermal-hydraulic modeling. The paper describes the thermal-hydraulic code Subchanflow and how the properties are solved and provided.

The Research Topic also emphasizes the role of multiphysics simulations in safety analysis. [Kutlu et al.](#) highlights the continuous development and improvement of the CTF sub-channel tool for thermal hydraulics, including new functionalities and multiphysics applications for VVER core modeling. These advancements are critical for ensuring the safe and reliable operation of these reactor types.

A cornerstone of credible simulation is rigorous experimental validation. [Colvin and Palmer](#) compares simulations to experimental results from the Sanida Annular Core Research Reactor, exploring potential improvements for feedback purposes, allowing additional iterations of the multiphysics coupling and checking for convergence, and evaluation of uncertainties in provided specific heat capacity values. These validation efforts are essential for establishing confidence in the predictive capabilities of multiphysics models.

[Kendrick and Forget](#) presents coupled neutronic/thermal-mechanical simulation of the Kilowatt Reactor Using Stirling TechnologY (KRUSTY) using OpenMC and MOOSE in order to analyze the neutronic and thermal impact of including thermal expansion at steady state. The results show that while thermal expansion has a significant effect on global neutronic tallies, it has relatively minor impact on spatial heating rates or temperatures in the system. This remains true even when simulating a multiple heat pipe failure scenario to introduce thermal asymmetry.

Looking ahead, the continued advancement and application of multiphysics methods hold significant promise for the future of nuclear energy modeling and design. By fostering collaboration between industry, academia, and research laboratories, and by prioritizing rigorous validation and uncertainty quantification, the nuclear community can unlock the full potential of these powerful simulation tools. This Research Topic serves as a valuable resource for researchers, engineers, and regulatory personnel alike, providing a comprehensive overview of the state-of-the-art and paving the way for a safer, more efficient, and more sustainable nuclear future.

This Research Topic provides evidence to the significant impact of multiphysics simulation capabilities in nuclear engineering. By including these advanced techniques, the nuclear industry can work to address the challenges of advanced reactor concepts.

Author contributions

MD: Writing – original draft, Writing – review and editing. ES: Writing – review and editing. DL: Writing – review and editing.

Funding

The author(s) declare that no financial support was received for the research and/or publication of this article.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The author(s) declared that they were an editorial board member of Frontiers, at the time of submission. This had no impact on the peer review process and the final decision.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated

organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.



OPEN ACCESS

EDITED BY

Fabio Giannetti,
Sapienza University of Rome, Italy

REVIEWED BY

Mark D. DeHart,
Idaho National Laboratory (DOE), United States
Qingquan Pan,
Shanghai Jiao Tong University, China

*CORRESPONDENCE

Deokjung Lee,
✉ deokjung@unist.ac.kr

RECEIVED 20 August 2024

ACCEPTED 25 September 2024

PUBLISHED 07 October 2024

CITATION

Imron M and Lee D (2024) On-the-fly thermal expansion for Monte Carlo multi-physics reactor simulations. *Front. Nucl. Eng.* 3:1483520.
doi: 10.3389/fnuen.2024.1483520

COPYRIGHT

© 2024 Imron and Lee. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

On-the-fly thermal expansion for Monte Carlo multi-physics reactor simulations

Muhammad Imron¹ and Deokjung Lee^{1,2*}

¹Department of Nuclear Engineering, Ulsan National Institute of Science and Technology 50 UNIST-gil, Ulsan, Republic of Korea, ²Advanced Nuclear Technology and Services, Ulsan, Republic of Korea

In this study, we present on-the-fly thermal expansion methodology for direct Monte Carlo coupled multi-physics reactor simulations. This approach allows the problem geometry to be thermally expanded on-the-fly during particle tracking using local temperatures, such as pin-averaged temperatures, obtained from the thermal-hydraulics solver. Numerical experiments demonstrated that modeling thermal expansion with local temperatures for thermal expansion improves the accuracy of reactor simulations, both for reactor eigenvalue and pin powers, compared to using global core-averaged temperatures. Additionally, the use of thermal expansion also improves the isothermal temperature coefficients, making them approximately 0.77 pcm/K closer to the measured data. Finally, results for depletion problems showed that incorporating thermal expansion in direct reactor modeling enhances the predicted critical boron concentration, particularly at high power and higher fuel burnup. These findings suggest that including thermal expansion in reactor modeling is essential for improving the fidelity of simulations.

KEYWORDS

thermal expansion, monte carlo, multi-physics, reactor simulations, neutron transport

1 Introduction

Thermal expansion (TE) significantly impacts reactor physics modeling, not only for fast reactors but also, to a certain extent, for Light Water Reactors (LWRs). This is because the geometric reactor design information is usually provided at room temperature, while the actual reactor core operates at much higher temperatures. At full power, the nominal fuel temperature is around 900 K, and the coolant temperature is around 580 K. This large temperature gap causes all components in the reactor vessel to thermally expand. For example, the core plate expands radially, causing the assembly pitch to expand. The grid spacers within an assembly also expand, altering the pin pitch within the lattice (Palmtag et al., 2017). Furthermore, the fuel pellet and fuel rod cladding expand in both radial and axial directions. In fact, Smith and Forget have asserted that TE must be modeled to achieve accurate high-fidelity simulation capability (Smith and Forget, 2013).

Thermal expansion modelling has long been incorporated into industry best practice two-step methodology for LWRs analyses. In this approach, the thermal expansion is already accounted for during lattices physics calculations to generate the few-group cross section for certain operating reactor temperatures. However, the direct the Monte Carlo (MC) coupled multi-physics simulations present a challenge: local temperatures are determined from thermal-hydraulic feedback, so they are not known a priori and are usually non-uniform within reactor core. To address this challenge, typical TE modeling in direct coupled multi-physics simulations is often achieved by manually adjusting input files

to uniformly expand reactor core geometry and modify material densities. This approach commonly uses core-averaged nominal temperatures, as demonstrated in the work by (Palmtag et al., 2017).

In Palmtag's work, thermal expansion was implemented by processing the XML input files for the Consortium for Advanced Simulation of LWRs (CASL) core simulator code, VERA-CS, to uniformly adjust the core dimensions and material densities. This process is accomplished by VERAin ASCII input preprocessor. However, VERAin only allows thermal expansion at a specific and uniform temperature, typically the core-averaged temperature, before beginning the simulations. Parametric studies were performed to evaluate the effects of thermal expansion due to various parameters, including fuel enrichment, diluted boron concentration in the coolant, coolant density, and fuel temperature. The study claimed that thermal expansion effects are adequately captured by uniformly expanding the materials from cold conditions to the average hot conditions. However, as this paper will demonstrate, the accuracy of the solutions, both for reactor eigenvalue and pin powers, can be improved if local temperatures, instead of core-averaged temperatures, are used for thermal expansion. Another drawback of using the input preprocessor is that if calculations need to be repeated at different power levels, the inputs must be reprocessed to obtain the correct expanded reactor dimensions and material densities.

Many other studies (Fiorina et al., 2019; Guo et al., 2021; Ma et al., 2021) have implemented thermal expansion using thermo-mechanical solvers in computational fluid dynamics (CFD) codes such as OpenFOAM or ANSYS. In general, these works employed a neutronic solver to obtain power, followed by a CFD code as the thermal-hydraulic (TH) solver and a mechanical solver to perform material deformation due to temperature changes. However, due to their reliance on direct CFD simulations, these approaches are likely suitable only for small cores, like those in heat pipe or modular reactors. The computational requirements for this method would be exponentially more expensive for larger cores.

Another noteworthy study on reactor simulations with thermal expansion modeling is the work conducted by Idaho National Laboratory (Cole et al., 2021). In this study, they modeled feedback mechanisms that include Doppler effects, radial expansion due to the displacement of the support plate, and axial expansion from the displacement of fuel pins for Unprotected Loss of Flow (ULOF) transient simulations in typical sodium fast reactors. A 2D BISON thermomechanical model of the 316 stainless steel support plate was used to simulate mesh displacement and its effect on the core assembly pitch. Meanwhile, fuel expansion was modeled using a 3D BISON thermomechanical model of the fuel pins. To integrate these models into a multi-physics simulation, they employed the MOOSE multi-app system (Permann et al., 2020), with Griffin, the neutronics solver, serving as the main application.

In this paper, we propose on-the-fly thermal expansion for MC coupled multi-physics reactor simulations, with a focus on pressurized water reactors (PWRs). The proposed method dynamically expands the reactor problem geometry during particle tracking, improving accuracy by using local temperatures, either pin-averaged or assembly-averaged. The proposed method also incurs negligible additional computational time. The method has been implemented into the MCS code (Lee et al., 2020) to provide solutions at both assembly and core levels, as well as for

whole-core depletion using restart calculations. MCS is a neutron/photon transport code that was developed at Ulsan National Institute of Science and Technology (UNIST) with capabilities of performing multi-physics and multi-cycle reactor analyses (Lee et al., 2017; Yu et al., 2020a; Yu et al., 2020b). The solutions are valuable for evaluating the influence of thermal expansion on the accuracy of direct MC coupled multi-physics reactor simulations. Additionally, they assess the effects of using local temperatures for thermal expansion on accuracy improvements.

2 Methodology

2.1 Theory and thermal expansion coefficients

The physical modelling for the linear thermal expansion of solid materials with an isotropic crystal structure is given by:

$$L = L_0 [1 + \alpha_L (T - T_{ref})], \quad (1)$$

where L is the final length at temperature T , L_0 is the initial length at reference temperature T_{ref} , and α_L is the coefficient of linear thermal expansion. Similarly, the physical modelling for area thermal expansion, such as fuel pellet area, is given by:

$$A = A_0 [1 + \alpha_L (T - T_{ref})]^2. \quad (2)$$

When materials are expanded, the mass must be preserved. Therefore, the material densities need to be modified to be consistent with the expanded dimensions. The density is modified according to:

$$\rho = \rho_0 \frac{V_0}{V}, \quad (3)$$

where ρ and V in Equation 3 are thermally expanded density and volume respectively, while ρ_0 and V_0 are initial density and volume, respectively. The thermally expanded volume can be calculated using the expanded length L and area A , shown in Equations 1, 2 respectively.

The thermal expansion coefficients of materials used in reactors are crucial for accurate thermal expansion modeling. Detailed descriptions of the correlations used to determine these coefficients for typical materials in LWRs are well described in the reference (Palmtag et al., 2017). However, this work adopted the thermal expansion coefficients used in the STREAM code (Choi et al., 2021). These thermal expansion coefficients are summarized in Table 1.

Note that these thermal expansion coefficients are quite close to those described in reference (Palmtag et al., 2017). Additionally, this study limits the thermal expansion modelling to fuel pellet, cladding, pin pitch and assembly pitch thermal expansion. The dimensions of absorber materials such control rods and burnable absorbers are assumed to remain constant.

2.2 Thermal-hydraulics coupling

MCS is a neutron/photon transport code that can be coupled with both TH1D (Ryu et al., 2015) and COBRA-TF (CTF) (Salko

TABLE 1 Thermal expansion coefficients used in this study.

Material	Thermal expansion coefficient (K^{-1})
Zirconium alloys (fuel cladding)	7.00×10^{-6}
SS304 (core plate)	1.78×10^{-5}
Fuel pellet	1.10×10^{-5}

et al., 2015). TH1D is a simple one-dimensional, single closed-channel thermal-hydraulics (TH) solver that does not account for cross-flow, coolant vaporization, and pressure drop across the channel. In contrast, CTF is a more sophisticated TH solver that addresses these limitations in TH1D.

Both TH1D and CTF are internally coupled with MCS, allowing data transfers between MCS and the TH solvers to occur automatically. Although the solvers must be compiled separately, they are linked through a static library. Linear power from MCS is transferred to the TH solvers, while TH parameters are updated by TH solvers and returned to MCS. These data transfers are performed at the pin-by-pin level.

Since data transfers are performed at the pin-by-pin level, MCS multi-physics coupling always uses a single pin cell as the channel for TH coupling. For example, in a single fuel assembly with a 17 × 17 configuration, including 24 guide tubes and one instrumentation tube, a total of 264 channels are modeled for TH coupling.

The reactor's TH parameters are updated regularly after several cycles of MCS particle tracking. During MCS particle tracking, fission power is tallied in both inactive and active cycles. In inactive cycles, the power tally is not accumulated, meaning that the tally from the last TH update is discarded. In contrast, during active cycles, when both the fission source and TH parameters have converged, the fission power tally is accumulated. The fission power is tallied for each cell in the fuel pellets and then converted into linear powers. These linear powers are passed to the TH solvers to update the reactor's TH parameters. These procedures are repeated until all cycles are completed.

2.3 Implementation

Many MC codes, such as MCS, use Constructive Solid Geometry (CSG) to define the geometry of reactor problems. Performing non-uniform geometrical expansion using local temperatures in CSG presents a challenge because a single surface can be reused to form multiple cells within a universe, and a universe can be repeatedly used to define a lattice. Consequently, modifications to a particular surface, such as those resulting from thermal expansion, can affect all cells, universes, and lattices that utilize the modified surface. One possible solution to avoid this issue is by creating copies of surfaces and cells, so that every cell can be expanded according to their respective local temperatures. However, this approach is cumbersome for large reactor problems and would consume more memory.

To address this problem, on-the-fly thermal expansion is introduced. In this approach, when a particle enters a particular pin, the geometry of fuel pellet and cladding in that fuel pin are

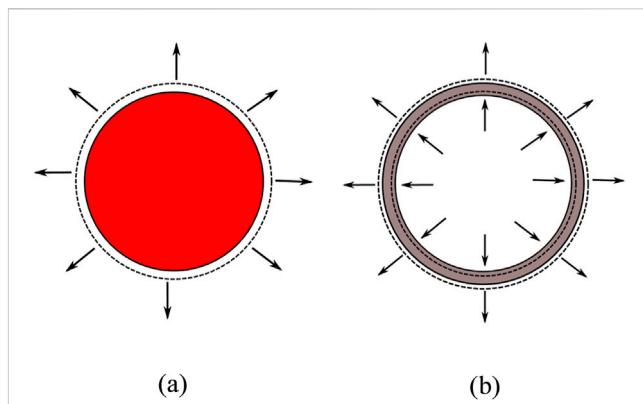


FIGURE 1
The radial expansion of (A) fuel pellet, and (B) cladding.

expanded using the corresponding local temperatures, which can be either pin-averaged or assembly-averaged temperatures. Similarly, the pin and assembly pitches where the particle is located are also expanded according to the assembly-averaged and core-averaged temperatures, respectively. In both the pin-averaged and assembly-averaged temperature cases, the average temperatures are calculated in both the axial and radial directions using volume-weighted averaging.

In the implementation in MCS, the fuel pellet expanded both in radial and axial directions, while the cladding is expanded only in the radial direction, with the inner and outer radius assumed to expand equally. It should be noted that geometry deformation is assumed to be uniform both radially and axially, which implies that fuel cladding ballooning is not considered. These expansions are illustrated in Figure 1.

Figure 2 is a flow chart illustrating the calculation of geometrical changes due to thermal expansion. Initially, after several cycles of neutron tracking, the TH solver is executed to determine the temperature distribution. This temperature distribution is then used to calculate pin-averaged fuel pellet and cladding temperatures, which are subsequently utilized to determine the corresponding expansion-induced geometrical changes in fuel pellets and cladding. Similarly, assembly-averaged temperatures are used to compute the geometrical changes in pin-pitch. Following this, core-averaged temperatures are used to thermally expand the assembly pitches uniformly. Finally, the material densities are updated to ensure the preservation of material mass throughout the process. These steps are repeated iteratively until all cycles are completed.

Once the geometrical changes due to thermal expansion are calculated during the TH update, this information is used to thermally expand the core geometry on-the-fly during subsequent cycles of neutron tracking. When a neutron is located in a pin cell after performing a random walk, MCS checks whether it originated from another pin cell or from the same one. If it originated from another pin cell, the fuel pellet and cladding surfaces in that pin cell are expanded based on the previously calculated geometrical changes. Additionally, if the neutron came from another assembly, the pin pitches within that assembly are uniformly expanded as well. These steps are summarized in Figure 3.

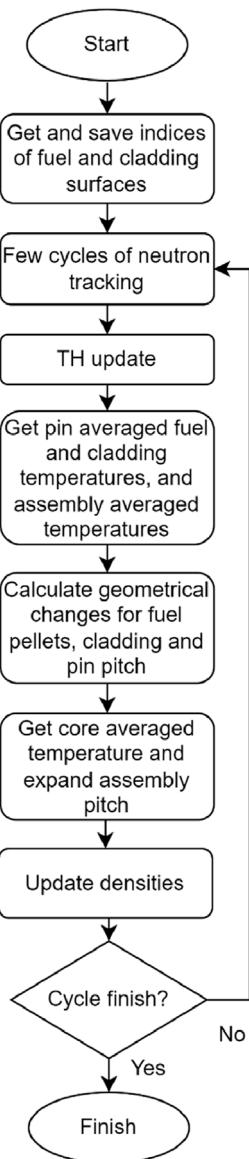


FIGURE 2
Calculation of geometrical changes due to thermal expansion.

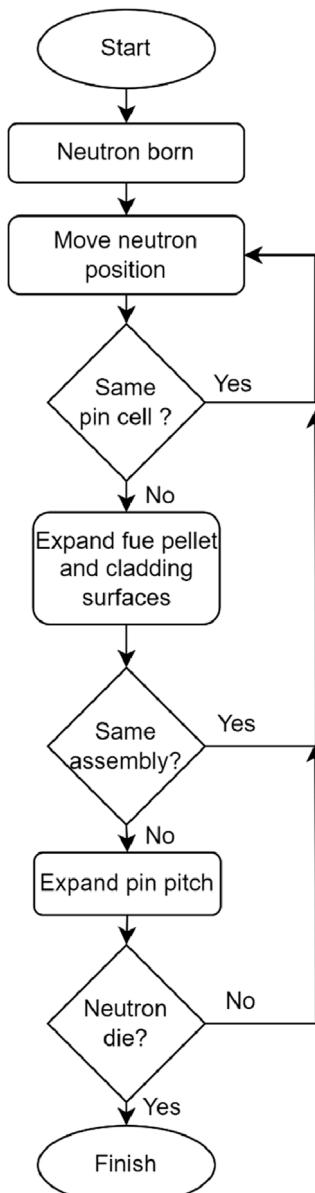


FIGURE 3
On-the-fly thermal expansion during neutron tracking.

3 Results and discussion

To evaluate the effects of TE modelling on MC coupled multi-physics reactor simulations, several test problems were developed. The test problems include assembly and core problems, as well as depletion problems with restart cases. The problem geometries and material compositions were adopted from the Virtual Environment for Reactor Applications (VERA) core physics benchmark (Godfrey, 2014). Additionally, ENDF/B-VII.1 was used for the MCS cross section library.

In all problems considered in this study, the TH parameters are updated every 500 cycles of MCS particle tracking. In MCS, the fuel pellets are axially discretized into 25 equidistant cells without radial discretization, where fission power is tallied. While the TH solvers discretize the problem axially into 25 equidistant meshes corresponding to the MCS cells. Additionally, in the TH solvers, the fuel pellets are radially discretized into 10 equidistant rings,

while the cladding and gap are each represented by a single ring to numerically solve the heat conduction equation.

3.1 Small lattice problem

This small lattice problem consists of a 2×2 pin cells with two different fuel enrichments: 3.1% wt. fuel enrichment in all fuel pins, except for the bottom-left pin, which has 2.1% wt. enrichment, as shown in Figure 4. The pin geometry and material compositions are based on the VERA benchmark, with a 0.4 cm inter-lattice gap. The boron concentration diluted in the coolant is 1,300 ppm. This problem simulated 4×10^4 particles per cycle, with a total of 14,500 cycles, of which 2,500 were designated as active cycles.

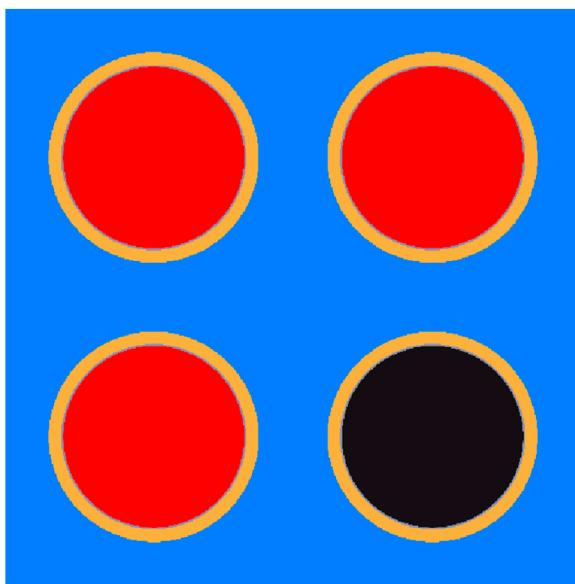


FIGURE 4
Small lattice problem geometry.

This simple problem is designed to evaluate the implementation of TE in the MC code MCS by comparing the solutions from on-the-fly TE in MCS to those from manually calculated TE. In the manual calculations, the geometrical expansions and expanded material densities were manually computed based on the pin-averaged temperatures from the on-the-fly thermal expansion calculation results. Note that the varying fuel enrichments in this problem result in non-uniform pin-averaged temperatures. A new MCS input file was then created using these manually calculated geometrical expansions and expanded material densities to include the equivalent TE.

Table 2 presents the results for small lattice problem. As observed, the on-the-fly TE and manually calculated TE produce very similar infinite multiplication factors, with only a 7 pcm difference, which is well within the given statistical uncertainty. For comparison, the result without TE is also given that shows a difference of more than 80 pcm compared to the TE cases. These results confirm that on-the-fly TE is correctly implemented in the MCS code.

3.2 Assembly problem

This test problem aims to assess the effects of TE at the assembly level. It is based on a problem similar to problem six of the VERA benchmark. The problem geometry also includes non-fuel structural materials such as the pin plenum, nozzles, core plates, and bottom and top reflectors. The effects of TE will be evaluated for different boron concentrations and fuel enrichments, with the reactor condition set at Hot Full Power (HFP). Note that the calculations for this problem simulated 4×10^4 particles per cycle, with a total of 14,500 cycles, of which 2,500 cycles were designated as active.

Table 3, 4 display the assembly reactivity differences due to TE for varying boron concentrations and fuel enrichment levels, respectively. Table 3 shows that as boron concentration increases, the reactivity difference decreases. This effect occurs because the

TABLE 2 Small lattice problems infinite multiplication factors.

Cases	k_{inf}
On-the-fly TE	1.14279 ± 0.00003
Manually calculated TE	1.14272 ± 0.00004
No TE	1.14364 ± 0.00003

TABLE 3 Assembly reactivity differences due to TE (TE - no TE) for various boron concentrations using 3.1% wt. enriched fuel.

Boron concentration (ppm)	Reactivity differences (pcm)
0	182 ± 6
600	43 ± 5
1,300	-76 ± 6

TABLE 4 Assembly reactivity differences due to TE (TE - no TE) for different fuel enrichment levels at a boron concentration of 600 ppm.

Fuel enrichment (% wt.)	Reactivity differences (pcm)
2.1	-35 ± 4
3.1	43 ± 5
4.4	101 ± 6

moderator volume containing boron expands with an increase in pin pitch. When the moderator has a high concentration of boron, the overall neutron absorptions would increase, resulting in a lower eigenvalue. Conversely, lower boron concentrations reduce neutron absorption and enhance neutron moderation, resulting in a higher eigenvalue. It has been well known that reactivity differences from TE are highly dependent on boron concentration in the reactor. Similarly, Table 4 shows that higher fuel enrichment levels lead to a more positive reactivity difference from TE. These trends in reactivity differences due to TE for varying boron concentrations and fuel temperatures are consistent with those found in reference (Palmtag et al., 2017).

3.3 Core problem

The core problem specifications are similar to those of problem 7 of the VERA benchmark, with the operating condition set at HFP. However, instead of determining the critical boron concentration, this study calculates the eigenvalues using a critical boron concentration of 860 ppm.

The effect of thermal expansion on the core problem will be evaluated at several expansion temperatures as can be seen in Table 5. For the core-averaged case, the core-averaged nominal reactor temperature at HFP is used, with the fuel expansion temperature set at 900 K and the coolant and cladding expansion temperatures set at 583 K. The assembly-averaged case uses assembly-averaged temperatures for geometry expansion in each assembly, while

TABLE 5 Core problem eigenvalues and pin-power power errors.

Cases	Eigenvalue	Min. pin error (%)	Max. pin error (%)	RMS pin error (%)
No TE	1.00046 ± 0.00002	-4.8	1.9	0.8
Core-averaged	1.00004 ± 0.00002	-2.1	1.8	0.5
Assembly-averaged	1.00001 ± 0.00002	-1.4	2.0	0.3
Pin-averaged	0.99997 ± 0.00002	Ref.	Ref.	Ref.

the pin-averaged case uses pin-averaged temperatures for geometry expansion in each pin cell. All cases simulated 3×10^5 particles per cycle, with a total of 42,000 cycles, of which 38,000 cycles were designated as active. This results in a maximum relative standard deviation of 0.8% in the radial pin powers.

The results for this problem are compiled in Table 5. The pin-power errors were calculated relative to the pin-averaged case. As shown in the table, when thermal expansion is not considered, the eigenvalue is overestimated compared to the pin-averaged case, and the errors at the pin level are more pronounced. Using core-averaged nominal temperatures greatly improves the eigenvalue, with less than a 10 pcm difference from the pin-wise case and the pin errors are also reduced. The results are further improved when local temperatures at the assembly level are used for geometry expansion; this approach results in only a 4 pcm difference in the eigenvalue and a 0.3% RMS pin error. It can also be seen that as the spatial resolution of expansion temperatures increases, both the eigenvalue and pin powers converge to those of the pin-wise case.

The radial pin-power errors relative to the pin-averaged case are shown in Figure 5. As can be observed, the no-TE case underestimates the pin powers at the core periphery by as much as -4.8%, while overestimating the pin powers at the core center. The pin powers from the core-averaged case are fairly accurate; however, there are still some areas with noticeable deviations. Lastly, the assembly-averaged case produces pin-power solutions quite close to the pin-averaged case. It is also important to note that the running time of the pin-wise case is only 0.9% longer compared to the no-TE case. Therefore, the introduction of on-the-fly thermal expansion incurs virtually no additional computational cost.

3.4 Isothermal temperature coefficient

This problem quantifies the effect of TE on the isothermal temperature coefficient (ITC). ITC is the change in the reactivity per unit change in the fuel and moderator temperature (ANSI/ANS-19.6.1-2005, 2005). ITC measurements are performed during HZP reactor physics tests to determine if the measured ITC is consistent with the calculated value (Hong, 2010).

In this study, the ITC measurement in the VERA benchmark was modeled for cases with and without TE. The reactor used in this benchmark is Watts Bar Unit 1, a Westinghouse PWR. The measured ITC was obtained during cycle 1, with all fresh fuel. The ITC was calculated using isothermal temperatures of 560 K and 570 K, with a boron concentration of 1,291 ppm.

To ensure statistical reliability, the ITC mean and the corresponding standard deviations were obtained by performing five runs for each case, with random seeds for each run. This resulted

in a total of 25 ITC samples, from which the mean and standard deviation were calculated. For each run, there were 14,500 cycles, of which 2,500 cycles were inactive, with 40,000 particle histories simulated for every cycle. The results are displayed and compared against measurement result in Table 6.

As indicated in the table, core modeling with TE makes the ITC more accurate and closer to measurement data. Furthermore, the results demonstrate that TE modeling makes the ITC more positive by 0.77 pcm/K. This ITC difference is slightly higher compared to the findings from reference (Palmtag et al., 2017), which reported that TE modeling makes the ITC 0.5–0.6 pcm/K more positive.

Although the ITC solutions for both the TE and No TE cases deviate significantly from the reported measured value, they are still comparable to the results of other MC calculations. For example, the ITC solution from KENO without TE is 5.72 pcm/K (Godfrey, 2014).

3.5 Depletion problem

Exercise 3 of the TVA Watts Bar Unit 1 multi-physics depletion benchmark (Albagami et al., 2021) was adopted to investigate the effect of TE on the boron letdown curve during reactor depletion. Restart cases from whole-core pin-by-pin depletion without TE were run at 0, 221.1, and 392.3 effective full power days (EFPDs) for cases with and without TE. This depletion problem simulated 6×10^4 particles per cycles, with a total of 5,000 cycles, of which 2,500 were active cycles.

MCS, coupled with the CTF thermal-hydraulics solver (Salko et al., 2015), was used to solve this problem. The work on MCS/CTF multi-physics coupling was done in the previous studies (Yu et al., 2017). The results were compared against the measured values obtained from reference (Godfrey, 2014) and are presented in Table 7.

As shown in the table, direct core modeling with TE yields more accurate predictions for this depletion whole-core problem, particularly as power increases and fuel burnup progresses. At the beginning of the cycle (BOC), core modeling with TE underestimates the measured critical boron concentration (CBC) by 22 ppm, which is slightly higher than the 13-ppm underestimation observed with core modeling without TE. This higher CBC underestimation in the TE model at BOC is attributed to the high boron concentration and zero power at the start of the cycle. However, as the fuel cycle progresses and power increases, core modeling with TE offers more precise CBC solutions. For instance, at the end of the cycle (EOC), the TE model underestimates the CBC by just 3 ppm, compared to a 27-ppm underestimation by the model

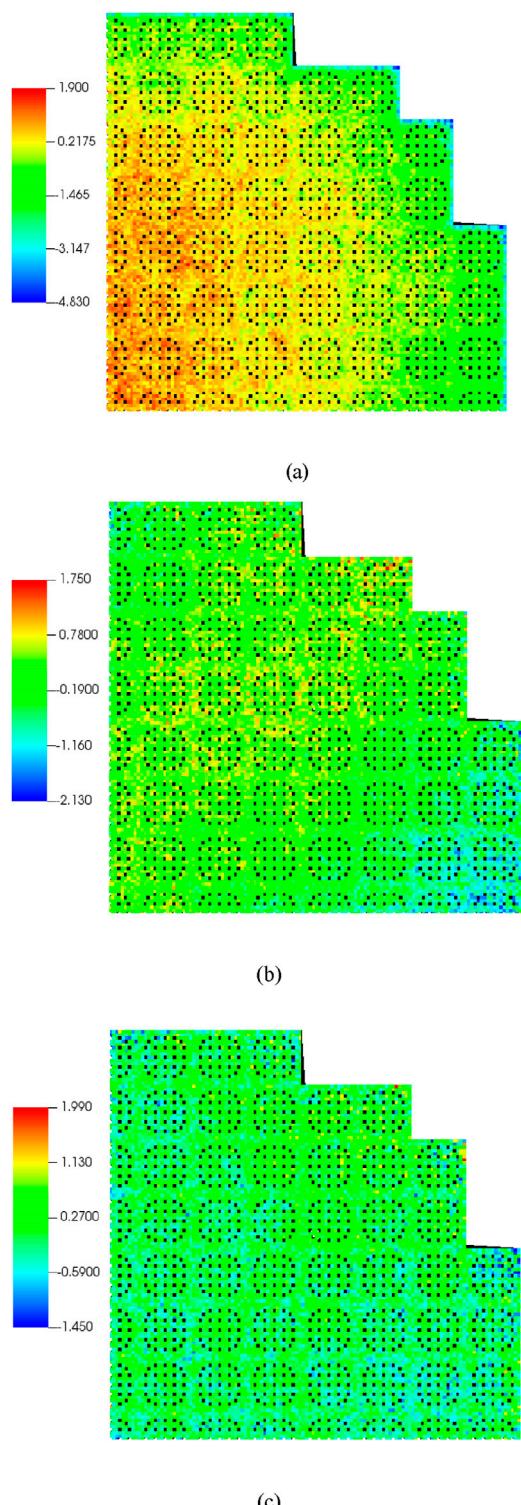


FIGURE 5
Radial pin-power errors distributions relative to the pin-averaged case for **(A)** No TE case, **(B)** core-averaged case, and **(C)** assembly-averaged case.

without TE. These results indicate that direct core modeling with TE provides more accurate solutions, particularly at full power and higher fuel burnup levels.

TABLE 6 Calculated ITCs compared with the measured value.

Case	ITC (pcm/K)
No TE	-6.73 ± 0.44
TE	-5.96 ± 0.39
Measurement	-3.91

These solutions can be further improved by dividing the fuel pellet into multiple radial rings to more accurately model the radial fuel temperature distribution within the pellet. This radial temperature variation is essential for correctly capturing the spatial self-shielding effect. Another study (Imron and Lee, 2025) reported that accurate spatial self-shielding modeling can increase the reactor eigenvalue by 50–80 pcm at HFP, which corresponds to approximately 5–8 ppm of boron concentration in typical PWR reactor problems.

4 Conclusion

This study presents an on-the-fly thermal expansion methodology for direct MC coupled multi-physics simulations implemented in the MCS code. The methodology was described and evaluated on several reactor problems. The trends in reactivity differences due to thermal expansion for varying boron concentrations and fuel temperatures, as well as the improvement in calculated ITC with thermal expansion, are consistent with previous studies. While using core-averaged temperatures for expansion is fairly accurate, the accuracy can be further improved by using local temperatures. Finally, this study shows that incorporating thermal expansion improves the solutions for depletion problems, especially at high power and high fuel burnup.

Thermal expansion for control rods and burnable poison rods will be included in future work. The inclusion of thermal expansion for control rods is particularly important for obtaining more accurate axial power distribution, which is often ignored in most current LWR analyses. Therefore, further studies are necessary to evaluate thermal expansion for other parts of the reactor core to achieve higher-fidelity reactor analyses.

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Author contributions

MI: Writing—original draft, Visualization, Validation, Software, Methodology, Investigation, Formal Analysis, Data curation, Conceptualization. DL: Writing—review and editing, Supervision, Resources, Project administration, Funding acquisition, Conceptualization.

TABLE 7 Critical boron concentrations for depletion problem.

Cases	EFPDs	Percent power	Bank D (step)	Critical boron concentration (ppm)		
				Measurement	TE	No TE
BOC	0.0	0.0	186	1,299	$1,277 \pm 0.8$	$1,286 \pm 0.6$
MOC	221.1	100.0	222	530	485 ± 0.9	474 ± 0.8
EOC	392.3	86.9	202	38	35 ± 0.9	11 ± 0.7

Funding

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. This work was supported by KOREA HYDRO & NUCLEAR POWER CO., LTD (No. 2022-Tech-13). The funder was not involved in the study design, collection, analysis, interpretation of data, the writing of this article, or the decision to submit it for publication.

Acknowledgments

During the preparation of this manuscript, the author(s) used ChatGPT-4o to improve the readability and clarity of the writing.

References

Albagami, T., Rouxelin, P., Abarca, A., Holler, D., Moloko, L., Avramova, M., et al. (2021). TVA Watts bar unit 1 multi-physics multi- cycle depletion benchmark version 2.2. Paris, France: OECD Nuclear Energy Agency.

ANSI/ANS-19.6.1-2005 (2005) (2005). *Reload startup physics tests for pressurized water reactors*. American Nuclear Society.

Choi, S., Kim, W., Choe, J., Lee, W., Kim, H., Ebwonjumi, B., et al. (2021). Development of high-fidelity neutron transport code STREAM. *Comput. Phys. Commun.* 264, 107915. doi:10.1016/j.cpc.2021.107915

Cole, M. M., Lin, C.-S., and Ortensi, J. (2021). *NRC multiphysics Analysis capability deployment FY 2021 - Part 2*. Idaho Falls, Idaho: Idaho National Laboratory.

Fiorina, C., Radman, S., Koc, M. Z., and Pautz, A. (2019). “Detailed modelling of expansion reactivity feedback in fast reactors using OpenFOAM,” in *International conference on mathematics and computational methods applied to nuclear science and engineering (M&C 2019)* (Portland, OR), 1434–1442.

Godfrey, A. T. (2014). *VERA core physics benchmark progression problem specifications*. U.S. Department of Energy.

Guo, Y., Li, Z., Huang, S., Liu, M., and Wang, K. (2021). A new neutronics-thermal-mechanics multi-physics coupling method for heat pipe cooled reactor based on RMC and OpenFOAM. *Prog. Nucl. Energy* 139, 103842. doi:10.1016/j.pnucene.2021.103842

Hong, S.-K. (2010). “Advanced on-line isothermal temperature coefficient measurement for A physics test,” in *Transactions of the Korean nuclear society autumn meeting* (South Korea: Jeju), 135–136.

Imron, M., and Lee, D. (2025). Monte Carlo coupled multi-physics with spatially continuous material properties. *Ann. Nucl. Energy* 210, 110856. doi:10.1016/j.anucene.2024.110856

Lee, H., Kim, W., Zhang, P., Khassenov, A., Park, J., Yu, J., et al. (2017). “Preliminary simulation results of BEAVRS three-dimensional Cycle 1 whole core depletion by UNIST Monte Carlo code MCS,” in *Proceeding of international conference on mathematics and computational methods (M&C 2017)* (South Korea: Jeju).

Lee, H., Kim, W., Zhang, P., Khassenov, A., Park, J., Yu, J., et al. (2017). “Preliminary simulation results of BEAVRS three-dimensional Cycle 1 whole core depletion by UNIST Monte Carlo code MCS,” in *Proceeding of international conference on mathematics and computational methods (M&C 2017)* (South Korea: Jeju).

Lee, H., Kim, W., Zhang, P., Lemaire, M., Khassenov, A., Yu, J., et al. (2020). MCS – a Monte Carlo particle transport code for large-scale power reactor analysis. *Ann. Nucl. Energy* 139, 107276. doi:10.1016/j.anucene.2019.107276

Ma, Y., Han, W., Xie, B., Yu, H., Liu, M., He, X., et al. (2021). Coupled neutronic, thermal-mechanical and heat pipe analysis of a heat pipe cooled reactor. *Nucl. Eng. Des.* 384, 111473. doi:10.1016/j.nucengdes.2021.111473

Palmtag, S., Kochunas, B., Jabaay, D., Han, Z., and Downar, T. (2017). “Modeling thermal expansion in VERA-CS,” in *Proceeding of international conference on mathematics and computational methods (M&C 2017)* (South Korea: Jeju).

Permann, C. J., Gaston, D. R., Andrš, D., Carlsen, R. W., Kong, F., Lindsay, A. D., et al. (2020). MOOSE: enabling massively parallel multiphysics simulation. *SoftwareX* 11, 100430. doi:10.1016/j.softx.2020.100430

Ryu, M., Jung, Y. S., Cho, H. H., and Joo, H. G. (2015). Solution of the BEAVRS benchmark using the nTRACER direct whole core calculation code. *J. Nucl. Sci. Technol.* 52, 961–969. doi:10.1080/00223131.2015.1038664

Salko, R. K., Lange, T., Kucukboyaci, V., Sung, Y., Palmtag, S., Gehin, J., et al. (2015). “Development of COBRA-TF for modeling full-core, reactor operating cycles,” in *5th topical meeting on advances in nuclear fuel management, ANFM 2015* (American Nuclear Society), 178–194.

Smith, K., and Forget, B. (2013). “Challenges in the development of high-fidelity LWR core neutronics tools,” in *International conference on mathematics and computational methods applied to nuclear science and engineering (M&C 2013)*. (Sun Valley, ID).

Yu, J., Lee, H., Kim, H., Peng, Z., and Lee, D. (2020a). Coupled neutronics-thermal-hydraulic simulation of BEAVRS cycle 1 depletion by the MCS/CTF code system. *Nucl. Technol.* 206, 728–742. doi:10.1080/00229545.2019.1677107

Yu, J., Lee, H., Kim, H., Zhang, P., and Lee, D. (2017). “Preliminary validation of MCS multi-physics coupling capability with CTF,” in *Proceedings of the reactor physics asia 2017 (RPHA17) conference, (chengdu, China: reactor physics asia conference organizer)*.

Yu, J., Lee, H., Kim, H., Zhang, P., and Lee, D. (2020b). Simulations of BEAVRS benchmark cycle 2 depletion with MCS/CTF coupling system. *Nucl. Eng. Technol.* 52, 661–673. doi:10.1016/j.net.2019.09.007

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.



OPEN ACCESS

EDITED BY

Deokjung Lee,
Ulsan National Institute of Science and
Technology, Republic of Korea

REVIEWED BY

Yunje Cho,
Korea Atomic Energy Research Institute (KAERI),
Republic of Korea
Muhammad Rizwan Ali,
Pakistan Institute of Engineering and Applied
Sciences, Pakistan

*CORRESPONDENCE

Gianfranco Huaccho Zavala,
✉ gianfranco.zavala@kit.edu

RECEIVED 04 February 2025

ACCEPTED 17 March 2025

PUBLISHED 01 April 2025

CITATION

Huaccho Zavala G, Gheeraert T,
Sánchez-Espinoza VH, Almachi JC and Imke U
(2025) Coupled serpent/subchanflow analysis
with unstructured mesh interfaces for a
hexagonal, plate-type VVR-KN fuel assembly.
Front. Nucl. Eng. 4:1570855.
doi: 10.3389/fnuen.2025.1570855

COPYRIGHT

© 2025 Huaccho Zavala, Gheeraert, Sánchez-
Espinoza, Almachi and Imke. This is an open-
access article distributed under the terms of the
[Creative Commons Attribution License \(CC BY\)](#).
The use, distribution or reproduction in other
forums is permitted, provided the original
author(s) and the copyright owner(s) are
credited and that the original publication in this
journal is cited, in accordance with accepted
academic practice. No use, distribution or
reproduction is permitted which does not
comply with these terms.

Coupled serpent/subchanflow analysis with unstructured mesh interfaces for a hexagonal, plate-type VVR-KN fuel assembly

Gianfranco Huaccho Zavala^{1*}, Thomas Gheeraert^{1,2},
Victor Hugo Sánchez-Espinoza¹, Juan Carlos Almachi^{1,3} and
Uwe Imke¹

¹Institute for Neutron Physics and Reactor Technology, Karlsruhe Institute of Technology, Karlsruhe, Germany, ²École Supérieure de Physique et de Chimie Industrielles de la Ville de Paris, Université PSL, Paris, France, ³Facultad de Ingeniería Mecánica, Escuela Politécnica Nacional, Quito, Ecuador

This work presents the further development and application of the multi-physics coupled code Serpent/subchanflow for analyzing cores loaded with fuel assembly designs characterized by complex geometries, such as the VVR-KN fuel assembly. A high-detail steady-state analysis of one VVR-KN fuel assembly is presented and discussed. The VVR-KN is a plate-type fuel assembly, arranged coaxially with hexagonal fuel-plate tubes. Its particular geometry layout configuration challenges both their neutronic and thermal-hydraulic modeling. In this work, the versatility of Serpent's multi-physics interface is exploited by using the unstructured mesh-based interface to update the properties of the fuel and coolant materials in a coupled neutronic/thermal-hydraulic simulation; these properties are solved and provided by the thermal-hydraulic code Subchanflow. Both neutronic and thermal-hydraulic models are developed for a single fuel assembly of 6.83 cm distance pitch and 60 cm active height, and state conditions for the simulations are defined. Typical material composition and main thermal properties for the fuel-meat (UO_2 -Al) and aluminum cladding (SAV-1) materials are extracted from references. This work paves the way for multi-physics analysis of research reactors with non-regular plates or subchannel geometries.

KEYWORDS

Monte Carlo, multi-physics, serpent, subchanflow, VVR-KN, unstructured mesh-based interface, research reactors

1 Introduction

In the last decade, considerable research and development was carried out worldwide to develop multi-physics coupling strategies to enhance the prediction accuracy of reactor physics simulations. This multi-physics strategy mainly involves the interaction between neutronics, thermal-hydraulics, and fuel performance codes, profiting from the great versatility of Monte Carlo neutronic codes and today's availability of powerful high-performance computing environments.

The Karlsruhe Institute of Technology (KIT) was leading different European projects, e.g., HPMC (High-Performance Monte Carlo Methods for Core Analysis) (Demazière et al., 2020), McSAFE (High-Performance Monte Carlo Methods for SAFETY Analysis) (Sanchez-

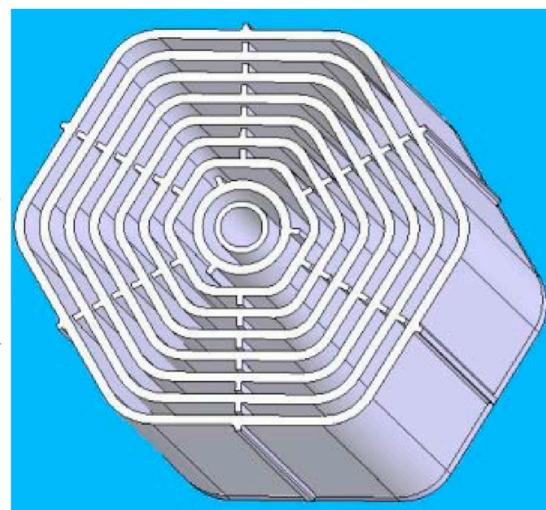
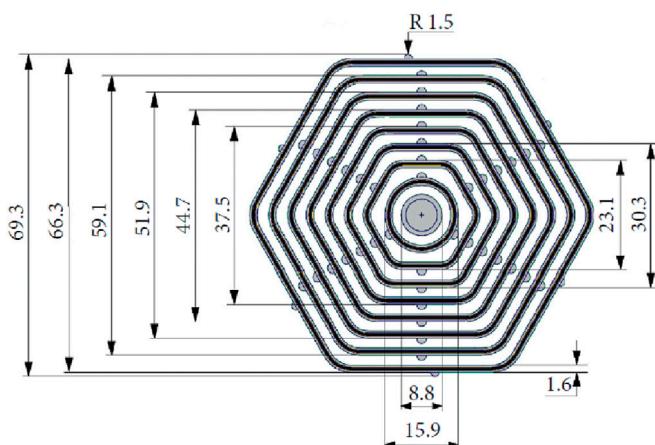


FIGURE 1
The VVR-KN standard fuel assembly, distances in [mm]. Images extracted from (Nguyen et al., 2020).

TABLE 1 Imposed TH conditions for stand-alone neutronic simulations.

Material	CZP		HFP(Xe = 0)	
	T [°C]	ρ [g/cm ³]	T [°C]	ρ [g/cm ³]
Meat	27	Table 3	70.0	Table 3
Cladding	27	Table 3	70.0	Table 3
Coolant	27	0.99657	53.4	0.98649

Espinosa VH. et al., 2021), and McSAFER (High-Performance Advanced Methods and Experimental Investigations for the Safety Evaluation of Generic Small Modular Reactors) (Sanchez-Espinosa VH. et al., 2021), where the proof-of-concept, development, optimization, and the verification and validation of multi-physics coupling strategies based on Monte Carlo neutronics for Light Water Reactor applications were performed and demonstrated. These efforts are paving the way for their use in more industry-like and safety relevant applications (Sanchez-Espinosa VH. et al., 2021). In this framework, the coupling between Serpent and Subchanflow based on a internal master-slave approach was developed (Ferraro et al., 2020a). The coupling has been widely tested under different benchmarks such as steady-state analysis on a full VVER-1000 core (Ferraro et al., 2021a), burnup analysis on PWR and VVER fuel assemblies (Ferraro et al., 2021b), and reactivity insertion transient problems on Pressurized Water Reactors (PWRs) (Ferraro et al., 2020b; Ferraro et al., 2020c) and Small Modular Reactors (SMRs) (Mercatali et al., 2023; Huaccho et al., 2025). Additionally, the KIT code Subchanflow developed initially for rod-type fuels has been modified and extended for the thermal-hydraulic analysis of thin plate-type fuels widely used in Material Testing research Reactors (MTRs) (Almachi et al., 2021). With this new extension in Subchanflow, the Serpent/Subchanflow tool has been used additionally to perform steady-state and transient simulations in some MTR-type cores such as the generic IAEA 10MW MTR core

TABLE 2 Boundary conditions for the TH feedback.

Parameter	Value
Power [kW]	393.49
Coolant inlet temperature [°C]	45.0
Core inlet pressure [kPa]	135.0
Coolant mass flow rate [kg/s]	5.58
Coolant's flow direction	Upward

(Almachi et al., 2022) and the SPERT IV D-12/25 core configuration (Almachi et al., 2024).

All the vast applications cited that use the coupling Serpent/Subchanflow rely in particular on one type of multi-physics interface, i.e., the regular mesh-based interface type 2 and the nested version type 22 (Serpent, 2024), which is perfect for modeling most of the standard reactor core designs. This work provides a first step in performing Serpent/Subchanflow coupled simulations using the unstructured mesh-based interface type 7 (Serpent, 2024). Interface type 7 was designed specifically to bring in solutions from solvers based on unstructured meshes such as fluid dynamic (CFD) codes. Its generality for defining different mesh shape geometries makes it suitable for the analysis of more unique core designs.

The structure of this work is as follows: Section 2 presents a general description of the VVR-KN fuel assembly. Section 3 describes the problem and defines the state conditions for the simulations. Section 4 briefly describes the tools and their main characteristics for simulating a coupled problem. Section 5 presents a detailed description of the neutronic and thermal-hydraulic models together with the interfaces based on unstructured meshes. Section 6 presents selected neutronic and thermal-hydraulic results. Finally, section 7 presents the summary and main conclusions of the work.

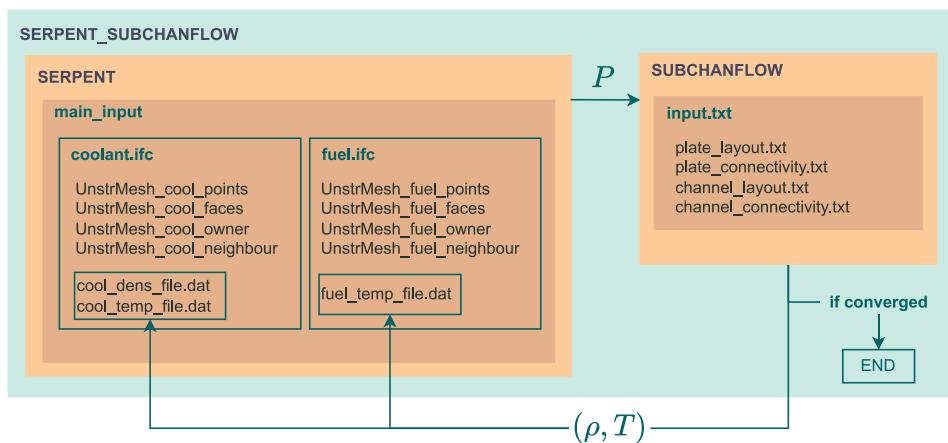


FIGURE 2
Flow calculation for a steady-state simulation with Serpent/Subchanflow using unstructured mesh-based interfaces for coolant and fuel. UnstrMesh_cool_* and UnstrMesh_fuel_* files define the unstructured superimposed meshes to update coolant and fuel properties.

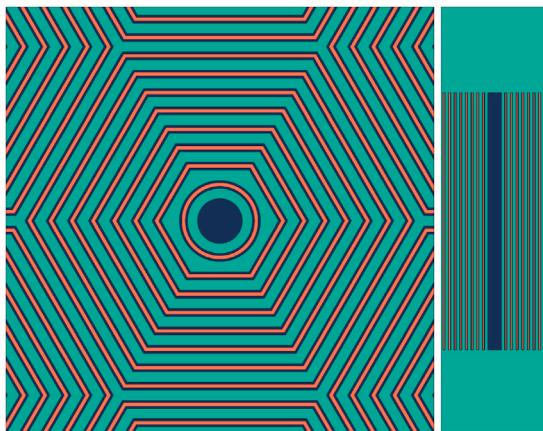


FIGURE 3
Radial and axial cuts from the Serpent model.

TABLE 3 Fuel and cladding material composition.

Property	UO ₂ [63 wt.%]-Al		SAV-1	
Density [g/cm ³]	5.0432			2.88
Composition [wt.%]	U ₂₃₅	10.965	Al	98.41
	U ₂₃₈	44.555	Mg	1.00
	O ₁₆	7.480	Si	0.44
	Al	37.000	Fe	0.14
			Cu	0.01

characterization. Most of the information needed for this work was extracted from (Nguyen et al., 2020). Two imposed TH states are defined, i.e., Cold Zero Power (CZP) condition and Hot Full Power (HFP) without xenon condition, see Table 1.

For the neutronic/TH characterization, boundary conditions for the TH feedback are defined considering nominal operation conditions. The reference values were obtained from (Nguyen et al., 2020), where a 10 MW conceptual core design using VVR-KN fuel assemblies (FAs) was investigated, and a maximum power of 393.49 kW is reported for a single FA. This power is considered as a reference value, and additional BCs are summarized in Table 2.

4 Multiphysics tools

4.1 Serpent

Serpent is a multi-purpose, three-dimensional, continuous-energy Monte Carlo transport code developed since 2004 at VTT Technical Research Centre of Finland Ltd (Leppänen et al., 2015). It represents a state-of-the-art code aimed to perform static, burnup, and dynamic 3-D calculations using standard ACE format Nuclear Data Libraries. Serpent was originally developed as a reactor physics

2 The VVR-KN fuel assembly

The VVR-KN is a hexagonal plate-type Russian fuel assembly design used in some research reactors. It comprises eight coaxial fuel elements (FE), seven having a hexagonal cylindrical shape, as showed in Figure 1. It has a characteristic pitch distance of 6.83 cm and 60 cm active height. The fuel meat is made of a UO₂-Al mixture, with a uranium density of 2.8gU/cm³ enriched in 19.75% of U₂₃₅ (Nguyen et al., 2020). The cladding and structural materials are made of SAV-1 aluminum alloy.

3 Problem definition

Imposed thermal-hydraulic (TH) conditions are defined for neutronics stand-alone characterization, and TH boundary conditions (BC) are defined for the steady-state neutronic/TH

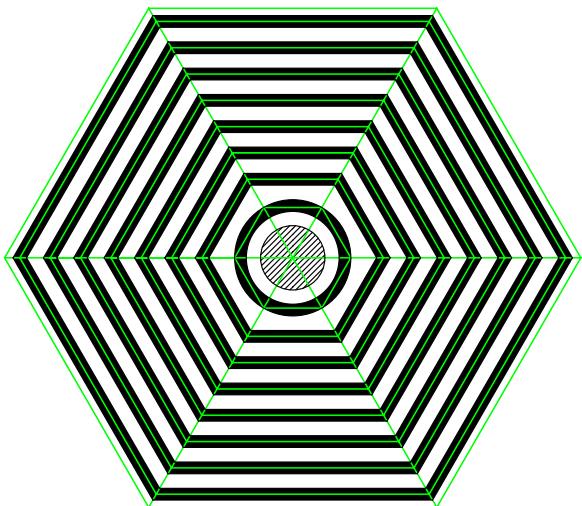


FIGURE 4
Superimposed mesh to update coolant properties in every subchannel.

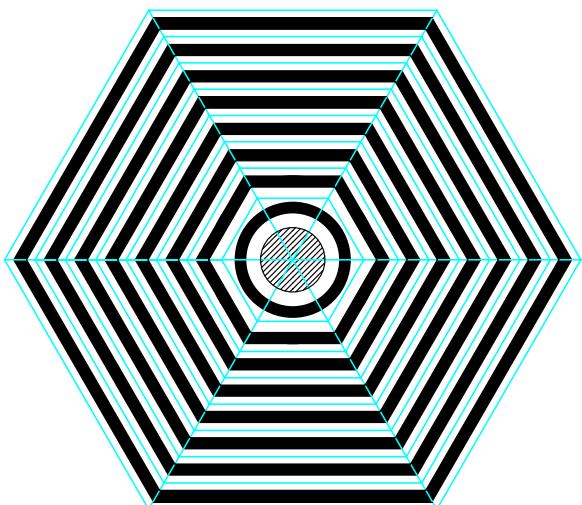


FIGURE 5
Superimposed mesh to update fuel properties in every fuel plate.

code, but the scope has considerably broadened over the years, bringing the development of many advanced features (Serpent, 2024). One of those main advanced features is the multi-physics (MP) capability (Valtavirta, 2015). The basic idea of the Serpent's MP capability is to bring certain material properties into the Serpent model, e.g., material's density and temperature distribution, in order to have more realistic modelling during the simulation. The MP feature is wisely designed by an interface, which is defined on top of the Serpent's model, i.e., without altering the original model's geometry. Depending on how the external code solves the problem in terms of physics and geometry discretization, different interfaces are available in Serpent, i.e., point-wise, rectangular and hexagonal regular meshes, unstructured meshes,

TABLE 4 Summary of the generated unstructured meshes.

Mesh parameter	Coolant mesh	Fuel mesh
Total cells or elements	1,080	960
Number of points	1,155	1,029
Number of faces	3,294	2,928
Number of neighbours	3,066	2,712

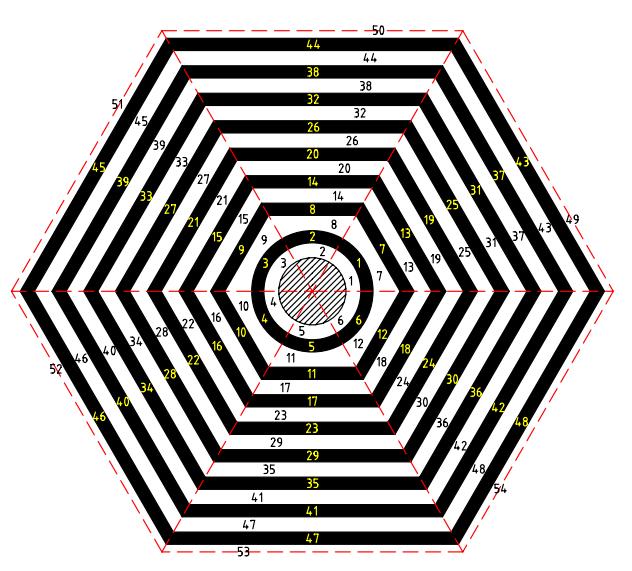


FIGURE 6
Plate/Subchannel thermal-hydraulic model developed in Subchanflow.

and many more. Examples of external codes coupled with Serpent via the MP interface are fluid dynamic (CFD) codes, channel thermal-hydraulics, and fuel performance codes (Leppänen et al., 2023; Leppänen, 2013).

One of the interfaces of interest for this work is the one based on unstructured meshes, designed specifically to bring in solutions from solvers based on this type of mesh. The interface is based on point-face-cell hierarchy, where the user gives a number of points that can be combined into a set of planar faces, which make up a set of closed convex cells (Serpent, 2024). The interface definition is based on the OpenFOAM format, and there is support for tetrahedral, hexahedral, and polyhedral-type meshes (Valtavirta, 2015). The Serpent version 2.1.32 is used in this work.

4.2 Subchanflow

Subchanflow (SCF) is a subchannel three-equations and single-phase flow thermal-hydraulic code for steady-state and transient analysis developed at Karlsruhe Institute of Technology (KIT), Germany (Imke and Sanchez, 2012). Subchanflow solves mass, momentum, and energy conservation equations along the axial discretization and between the neighbour lateral channels,

TABLE 5 Fuel and cladding thermal properties.

Property	Fuel	Cladding
Material	UO ₂ [63 wt.%]-Al	SAV-1
Conductivity [W/m/K]	64 (Abdukadyrova et al., 2014; Hagrman and Reymann, 1979)	151 (Abdukadyrova et al., 2014)
Specific heat [J/kg/K]	456 (Stahl, 1982)	924

TABLE 6 Global criticality results. The number in parenthesis is the 1σ uncertainty in the k_{eff} 's last digit. (*) imposed TH conditions.

Codes	Serpent		Serpent/SCF
State	CZP	HFP	HFP
k_{eff} (1 σ)	1.52968(8)	1.52585(8)	1.52577(8)
Ave. Tfuel [°C]	27*	70.0*	74.9
Ave. Tclad [°C]	27*	70.0*	73.4
Ave. Tcool [°C]	27*	53.4*	53.6

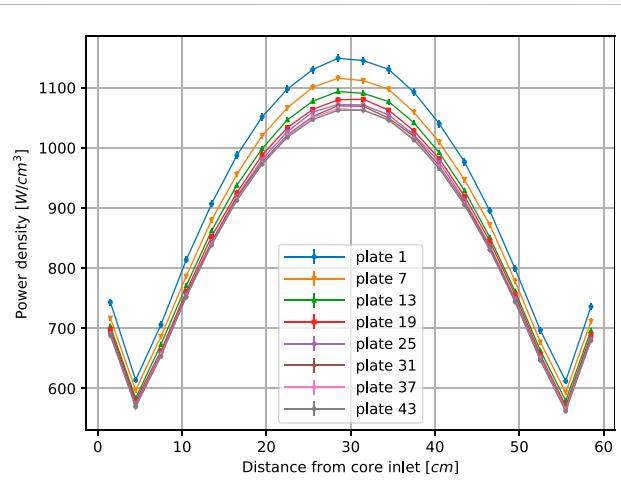
i.e., cross-flow between channels. These equations together with a set of empirical correlations to calculate, for example, the pressure drop, heat transfer coefficients, void generation, etc., represent the system of equations of the single-phase (liquid/vapor mixture) flow model (Imke and Sanchez, 2012). Subchanflow was initially developed for rod-type fuel elements, where the geometry is defined as a set of channels and rods with given hydraulic parameters and connectivities. A typical channel is characterized by its area, wetted and heated perimeter, and a list of its neighbour channels; and a typical rod is characterized by its material type for fuel and cladding, rod diameters and gaps, and a list of neighboured channels where the heat is released. The temperature profile inside each rod is calculated by dividing it into several radial rings and solving the heat equation in the radial direction by a finite volume method (Imke and Sanchez, 2012). An extension to plate-type fuel assemblies typically for analyzing MTR research reactors is also available in Subchanflow (Almachi et al., 2021). The Subchanflow version 3.7.1 is used in this work.

4.3 The Serpent/Subchanflow coupling

The Serpent/Subchanflow coupling was developed in the framework of the McSAFE project (Sanchez-Espinoza VH. et al., 2021), and the main implementation aspects were first introduced in (Ferraro et al., 2020a). The coupling relies on the capabilities and advanced features of both codes, e.g., the multi-physics interface in Serpent to interact with other codes and the Subchanflow aptness to be used as an external library. The implementation consists of a high-level set of routines (coupling routines) developed in C language. The coupling routines manage open-door functions defined in the Serpent code, which allows the management of all main aspects of a coupled scheme. Subchanflow is included and used

TABLE 7 Converged power distribution in every FE. The numbers in parenthesis are the 1σ uncertainty in the value's last digit.

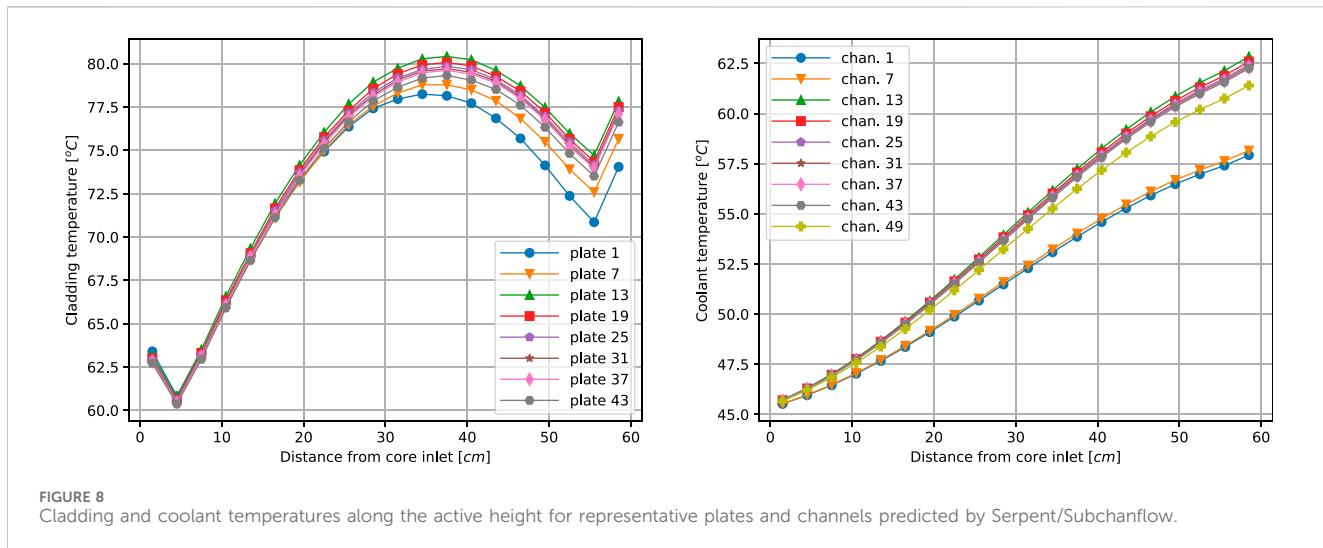
FE	Plates	Power	Volume	Power density
		[kW]	[cm ³]	[W/cm ³]
1	1–6	17.29(1)	18.868	916.2(4)
2	7–12	27.81(1)	31.281	889.2(3)
3	13–18	36.39(1)	41.756	871.5(3)
4	19–24	45.01(1)	52.232	861.7(2)
5	25–30	53.65(2)	62.707	855.6(2)
6	31–36	62.38(2)	73.183	852.5(2)
7	37–42	71.07(2)	83.658	849.5(2)
8	43–48	79.88(2)	94.133	848.6(2)

FIGURE 7
Power density along the active height for the representative plates predicted by Serpent/Subchanflow. Error bars are $\pm 3\sigma$.

as an external library, resulting in an embedded master-slave implementation (Ferraro et al., 2020a).

Figure 2 shows the flow calculation for a coupled steady-state simulation where successive iterations between Subchanflow and Serpent are performed. The codes exchange the Density-Temperature distribution (ρ, T)¹ and Power distribution (P) until a desired convergence is achieved. In the original implementation with regular mesh-based interfaces (Ferraro et al., 2020a), these variables are exchanged directly from RAM memory, i.e., copying the (ρ, T) variables to Serpent using the C function `memcpy()`; meanwhile, in this case, densities and temperatures (ρ, T) are overwritten in their respective interfaces files in every iteration, as shown in Figure 2. Only extra lines using

¹ For simplicity, (ρ, T) denotes the fuel and coolant Temperature-Density distribution.



`fprintf()` functions were needed in the coupling routines (Ferraro et al., 2020a) to accomplish this idea without affecting the Serpent and Subchanflow source codes.

5 Developed models

5.1 Serpent model

Following the fuel assembly description in Section 2, a 3D model was developed in Serpent as shown in Figure 3. The main assumptions and considerations taken into account are the following:

- A simplified model in terms of geometry was considered, i.e., no structural ribs, straight plate corners instead of curved ones, and a solid rod of SAV-1 material for the inner structural tube; see Figure 3.
- Pure water was considered for the bottom and top axial reflectors. Each with a 20 cm axial length. For the HFP state, fixed densities and temperatures corresponding to the inlet and outlet core temperatures are considered, i.e., 45°C and 62°C for the bottom and top axial reflectors, respectively.
- Most references related to this type of FA do not report the proportion between Aluminum and UO₂ in the meat. Only one reference that reports a volumetric concentration was found (Kkonoplev and Sorlov, 2010). In this work, a weight concentration of 63wt.% UO₂ in the meat was considered. A typical composition for an SAV-1 alloy was extracted from (Abdukadyrova et al., 2014; Salikhbaev et al., 2009). Table 3 summarizes the composition of both materials used in this work.
- JEFF-3.1.1 ACE nuclear data library was considered for the simulations. For the thermal scattering law of H in H₂O, generated libraries at 294K and 324K were considered for the CZP and HFP, respectively, in neutronic stand-alone characterizations. For the steady-state case, interpolation

between the libraries at 294K and 374K was considered according to the temperature predicted by Subchanflow.

- Reflective boundary conditions in the radial (XY) direction and black in the axial (Z) direction.
- For Serpent criticality simulations, 10³ cycles with 10⁵ particles each were considered, with 100 inactive cycles for source convergence.
- No xenon concentration was considered in the HFP state condition.
- Two independent meshes were generated using the MP interface type 7 to update the properties of the fuel and coolant materials. The meshes were created with the condition that a mesh cell encloses part of the material (meat or coolant) that is intended to be updated. For that, the FA was subdivided into six triangles, every triangle containing eight fuel plates and nine coolant subchannels. Figures 4, 5 show a 2D representation of each mesh. Axially, the active height was divided into 20 axial cells. Only the active height was considered for the TH feedback. Table 4 summarizes the characteristics of each mesh.

5.2 Subchanflow model

A thermal-hydraulic model at plate/subchannel level is developed in Subchanflow taking into account the following assumptions and considerations:

- The subchanflow model consists of 54 subchannels and 48 plates, and both subchannels and plates are divided axially in 20 cells, see Figure 6. Two subchannels surround every plate; for simplicity, the power generated in a plate is equally distributed to its neighboring subchannels.
- Dittus-Boelter and Colburn heat transfer correlations were considered. Blasius correlation for the friction factor, and the IAPWS-97 standard is selected for the coolant water properties.

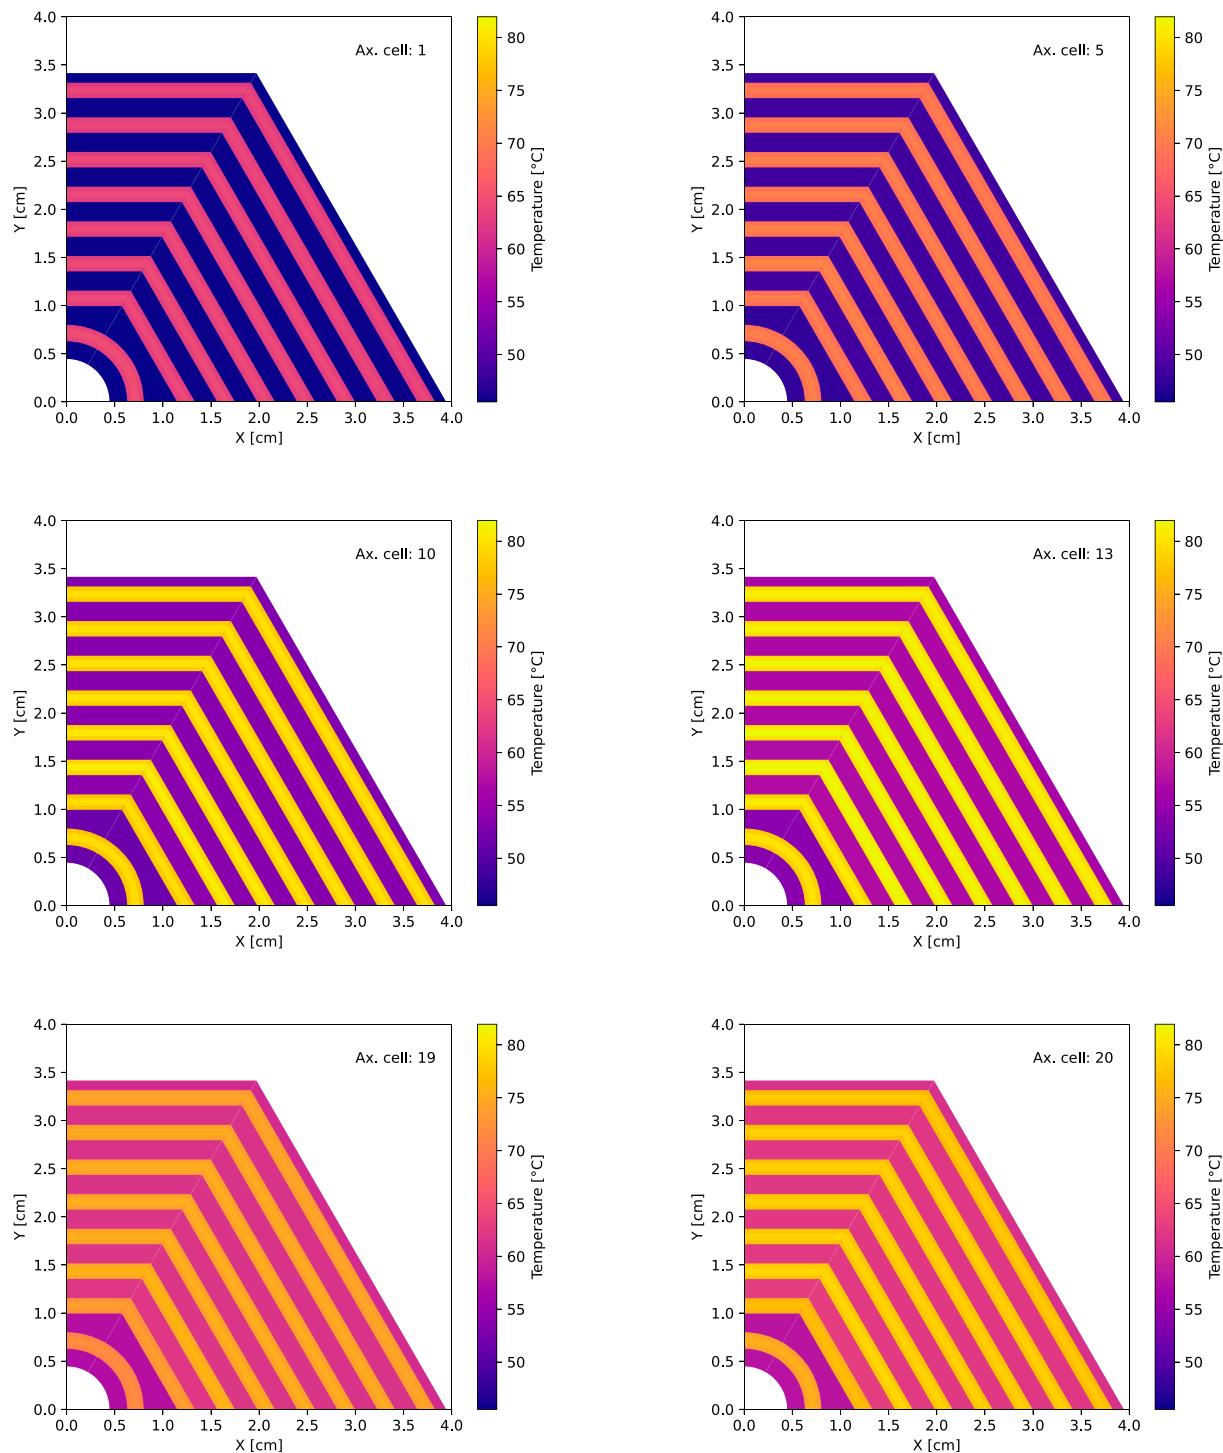


FIGURE 9
Temperature distribution predicted by Serpent/Subchanflow at different axial layers starting from the inlet to the core outlet. Only 1/4 of the core is shown due to the symmetry of the problem.

- For heat conduction in the fuel plate, five nodes were considered in the meat and two nodes in the cladding. Equivalent straight plates were considered for the inner circular fuel plate.
- Constant thermal properties showed in Table 5 were considered. The parameters were obtained considering a

reference temperature of 70°C. Properties for UO₂ and Al materials were obtained independently, and weighted averages were considered using the mass fractions as weights. No thermal expansion in the materials was considered.

- Subchanflow calculates a temperature profile inside the thin fuel meat, and one single temperature value per cell is needed

TABLE 8 Serpent/Subchanflow steady-state simulations using different heat transfer correlations.

Parameter	Dittus-Boelter	Colburn	Diff.
k_{eff} (1 σ)	1.52585(8)	1.52577(8)	3(4) pcm
Max. Tfuel	80.3°C	82.0°C	1.6°C
Ave. Tfuel	73.4°C	74.9°C	1.5°C
Max. Tclad	78.8°C	80.2°C	1.4°C
Ave. Tclad	72.1°C	73.4°C	1.3°C
Max. Tcool	62.9°C	62.9°C	0.0°C
Ave. Tcool	53.6°C	53.6°C	0.0°C

in Serpent for the fuel temperature feedback. A volumetric average value is considered for the cross-section adjustments in Serpent (on-the-fly temperature treatment (Viitanen and Lepänen, 2014)).

6 Results

Stand-alone neutronics and neutronic/TH simulations are presented in this section. All the simulations have been performed in a GNU/Linux single machine with an AMD EPYC 7542 32-Core Processor using 30 OpenMP threads. For the steady-state coupled simulation, eight iterations were performed, obtaining differences between the last two iterations in L_2 -norm (Ferraro et al., 2020a) of $< 9 \times 10^{-3}$ °C in fuel temperature, $< 7 \times 10^{-4}$ °C in coolant temperature, and $< 3 \times 10^{-7}$ g/cm³ in the coolant density. Every iteration took around 30 min, resulting in approximately 4 h to complete the steady-state simulation. A summary of the criticality results and average temperatures for the fuel (Tfuel), cladding (Tclad), and coolant (Tcool) are summarized in Table 6.

In the neutronics stand-alone results, a difference of ~ 160 pcm is observed as a difference between the CZP and HFP state conditions. No appreciable difference in reactivity is observed in the HFP state when the TH feedback is considered. The average temperatures predicted by Subchanflow are very close to the imposed temperatures for the HFP state, with a maximum difference of 4.9°C in the Ave. Tfuel.

Detectors to account for the power in every FE (Fuel Element²) are set up and summarized in Table 7. The higher power density is obtained in the internal FE and decreases slightly, keeping more or less constant in the remaining FEs. The axial power density distribution is shown in Figure 7. Thermal peaks due to the axial reflectors can be observed in both endings. A slight offset (almost negligible) of the axial power distribution in direction to the core inlet can be observed due to the slightly higher density of the coolant at the core inlet.

Figure 8 shows the axial cladding and coolant temperature distribution in representative plates and subchannels. The highest

cladding temperature is obtained in plate 13 (3rd FE), achieving a maximum value of 80.2°C in the 13th axial cell. The coolant temperature in channel 13 achieves a maximum value of 62.9°C in the last axial cell (core outlet). Figure 9 shows radial temperature distribution in all the materials (fuel, cladding, and coolant) at different axial layers.

The thermal-hydraulic results previously presented were calculated using the Colburn correlation. A sensitivity analysis comparing the Colburn and Dittus-Boelter correlations showed that the Colburn correlation yielded slightly higher temperature. To ensure a conservative approach, the results using Colburn were selected for presentation. The comparative results using these two correlations are summarized in Table 8, showing negligible differences, with temperature variations less than 2°C. The observed difference arises because the Colburn correlation predicts a lower heat transfer coefficient. This results in a higher temperature difference between the coolant and the external cladding surface, leading to slightly elevated cladding and fuel temperatures. Similar findings were reported during the Subchanflow validation for plate-type fuels (Almachi et al., 2021).

As a final step, a comparison against the TH evaluation performed in (Nguyen et al., 2020) is presented. The hottest FA producing 393.49kW is analyzed in both cases. However, different procedures are followed: The process followed in (Nguyen et al., 2020) consists first of a 10MW full-core characterization using the MCNP6 code, followed by a safety evaluation for the hottest FA in the core using the TH code PLTEM4.2 with detailed power distributions obtained from the core simulation. In our work, only the total power (393.49kW) was considered as an input for the single FA model in Serpent with reflective BCs, missing the real conditions of the FA in the core. Table 9 summarizes and compares some selected results. The Serpent/Subchanflow power distribution is more homogeneously distributed mainly due to the reflective BCs used in the radial direction, predicting a lower peak power value and, therefore, lower peak temperatures compared to (Nguyen et al., 2020) where higher peak power is predicted.

7 Summary and conclusion

A coupled neutronic/thermal-hydraulic simulation of a single VVR-KN fuel assembly using Serpent/Subchanflow was performed. Eight iterations were considered to obtain a converged steady-state solution, capturing consistent results regarding power and temperature profile distributions. Unfortunately, a comparison against experiments or referent simulations was hard to perform because of the simplified problem definition considered in this work and all the required assumptions, e.g., boundary conditions, set of correlations, material composition, and constant thermal properties. Nevertheless, the novelty of the result remains in the high-detail coupled neutronic/thermal-hydraulic solution in a non-regular plate or subchannel shape size, and it paves the way for a more realistic simulation, such as a VVR-KN full-core characterization.

An important remark of the work is the satisfactory use of the Serpent's unstructured mesh-based interface for the thermal-hydraulic coupling with Subchanflow. It is the first time that this type of interface has been employed within a Serpent/Subchanflow simulation. This interface type offers excellent versatility and more

² The FA is composed of eight coaxially arranged FEs, and every FE in the thermal-hydraulic model is divided in six plates.

TABLE 9 Serpent/Subchanflow comparison against Nguyen work.

Parameter	Serpent/SCF	PLTEMP4.2 (Nguyen et al., 2020)
Max. Pdens [kW/cm ³]	1,150	1,442.6
Peak Power Factor	1.32	2.2
Max. Tfuel [°C]	82.0	-
Max. Tclad [°C]	80.2	83.0
Max. Tcool [°C]	62.9	66.0

generality for generating non-regular meshes. Minor changes in the coupling routines were required to ensure consistent data exchange between the two codes without affecting their source codes. A better performance could be achieved by exchanging the information directly from RAM memory. However, some changes in the source codes are needed to add this feature to the coupled tool.

Subchanflow has been validated widely for standard PWR and BWR rod-type fuels. However, extra development is still needed to enlarge the thermal-hydraulic analysis in MTR research reactors. Some examples of future developments are implementing critical heat flux correlations for plate-type fuels, temperature-dependent thermal properties for standard MTR materials, and heat expansion treatment.

Data availability statement

The raw data supporting the conclusion of this article will be made available by the authors, without undue reservation.

Author contributions

GH: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Software, Visualization, Writing—original draft, Writing—review and editing. TG: Conceptualization, Investigation, Software, Writing—review and editing. VS: Funding acquisition, Project administration, Resources, Writing—review and editing, Supervision. JA: Software, Writing—review and editing. UI: Software, Supervision, Writing—review and editing.

References

Abdukadyrova, I. K., Alikulov, S. A., Akhmedzhanov, F. R., Baitolesov, S. A., Boltaboev, A. F., and Salikhbaev, U. S. (2014). High-temperature thermal conductivity of sav-1 aluminum alloy. *At. Energy* 116, 100–104. doi:10.1007/s10512-014-9825-0

Almachi, J. C., Sánchez-Espinoza, V., and Imke, U. (2021). Extension and validation of the subchanflow code for the thermo-hydraulic analysis of mtr cores with plate-type fuel assemblies. *Nucl. Eng. Des.* 379, 111221. doi:10.1016/j.nucengdes.2021.111221

Almachi, J. C., Sánchez-Espinoza, V. H., and Imke, U. (2022). High-fidelity steady-state and transient simulations of an mtr research reactor using serpent2/subchanflow. *Energies* 15, 1554. doi:10.3390/en15041554

Almachi, J. C., Sánchez-Espinoza, V., Imke, U., Stieglitz, R., and Margulis, M. (2024). Validation of the dynamic simulation capabilities of serpent2/subchanflow using experimental data from the research reactor spert iv d-12/25. *Nucl. Eng. Des.* 418, 112840. doi:10.1016/j.nucengdes.2023.112840

Demazière, C., Espinoza, S., Hugo, V., and Chanaron, B. (2020). Advanced numerical simulation and modelling for reactor safety contributions from the cortex, hpmc, mcsafe and nuresafe projects. *EPJ Nucl. Sci. Technol.* 6, 42. doi:10.1051/epjn/2019006

Ferraro, D., García, M., Imke, U., Valtavirta, V., Tuominen, R., Bilodid, Y., et al. (2021a). Serpent/subchanflow coupled calculations for a vver core at hot full power. *PHYSOR2020 EPJ Web Conf.* 247, 04006. doi:10.1051/epjconf/202124704006

Ferraro, D., García, M., Imke, U., Valtavirta, V., Tuominen, R., Leppänen, J., et al. (2021b). Serpent/subchanflow coupled burnup calculations for vver fuel assemblies. *PHYSOR2020 EPJ Web Conf.* 247, 04005. doi:10.1051/epjconf/202124704005

Ferraro, D., García, M., Valtavirta, V., Imke, U., Tuominen, R., Leppänen, J., et al. (2020a). Serpent/subchanflow pin-by-pin coupled transient calculations for a pwr minicore. *Ann. Nucl. Energy* 137, 107090. doi:10.1016/j.anucene.2019.107090

Ferraro, D., García, M., Valtavirta, V., Imke, U., Tuominen, R., Leppänen, J., et al. (2020b). Serpent/subchanflow pin-by-pin coupled transient calculations for the spert-iiie hot full power tests. *Ann. Nucl. Energy* 142, 107387. doi:10.1016/j.anucene.2020.107387

Ferraro, D., Valtavirta, V., García, M., Imke, U., Tuominen, R., Leppänen, J., et al. (2020c). OECD/nrc pwr mox/uo2 core transient benchmark pin-by-pin solutions using serpent/subchanflow. *Ann. Nucl. Energy* 147, 107745. doi:10.1016/j.anucene.2020.107745

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. The authors thank the HGF-Program NUSAFFE at the Karlsruhe Institute of Technology (KIT) for the financial support for the research topic multi-physics at the Institute for Neutron Physics and Reactor Technology (INR).

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The authors declare that no Generative AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Hagman, D. L., and Reymann, G. A. (1979). *Matpro - version 11, a handbook of materials properties for use in the analysis of light water reactor fuel rod behavior*. Idaho: EG&C Idaho, Inc. Idaho Falls.

Huaccho, Z. G., Sánchez-Espinoza, V. H., Mercatali, L., and Imke, U. (2025). An overcooling transient analysis in a carem-like smr core using serpent/subchanflow. *Ann. Nucl. Energy* 214, 111204. doi:10.1016/j.anucene.2025.111204

Imke, U., and Sanchez, V. H. (2012). Validation of the subchannel code subchanflow using the nupec pwr tests (psbt). *Sci. Technol. Nucl. Installations* 2012, 1–12. doi:10.1155/2012/465059

Konoplev, K. (2010). “Cooling and radiation safety control requirements for safe storage of LEU WWR-M experimental fuel assemblies”. 14. international topical meeting on Research Reactor Fuel Management (RRFM) (RRFM 2010), S. [5 p.].

Leppänen, J. (2013). Modeling of nonuniform density distributions in the serpent 2 Monte Carlo code. *Nucl. Sci. Eng.* 174, 318–325. doi:10.13182/NSE12-54

Leppänen, J., Pusa, M., Viitanen, T., Valtavirta, V., and Kaltiaisenaho, T. (2015). The serpent Monte Carlo code: status, development and applications in 2013. *Ann. Nucl. Energy* 82, 142–150. doi:10.1016/j.anucene.2014.08.024

Leppänen, J., Viitanen, T., and Valtavirta, V. (2023). Multi-physics coupling scheme in the serpent 2 Monte Carlo code.

Mercatali, L., Huaccho, G., and Sanchez-Espinoza, V. H. (2023). Multiphysics modeling of a reactivity insertion transient at different fidelity levels in support to the safety assessment of a smart-like small modular reactor. *Front. Energy Res.* 11. doi:10.3389/fenrg.2023.1130554

Nguyen, N. D., Nguyen, K. C., Huynh, T. N., Vo, D. H. D., and Tran, H. N. (2020). Conceptual design of a 10mw multipurpose research reactor using vvr-kn fuel. *Sci. Technol. Nucl. Installations* 2020, 1–11. doi:10.1155/2020/7972827

Salikhbaev, U. S., Baytelesov, S. A., Kungurov, F., and Saidov, A. S. (2009). Studying of influence of a neutron irradiation on element contents and structures of aluminum alloys sav-1 and amg-2. *Trans. RRFM 2009, Vienna*.

Sanchez-Espinoza, V. H., Gabriel, S., Suikkanen, H., Telkkä, J., Valtavirta, V., Bencik, M., et al. (2021b). The h2020 mcsafer project: main goals, technical work program, and status. *Energies* 14, 6348. doi:10.3390/en14196348

Sanchez-Espinoza, V. H., Mercatali, L., Leppänen, J., Hoogenboom, E., Vocka, R., and Dufek, J. (2021a). The mcsafer project high-performance Monte Carlo based methods for safety demonstration: from proof of concept to industry applications. *EPJ Web Conf.* 247, 06004. doi:10.1051/epjconf/202124706004

Serpent (2024). A continuous-energy Monte Carlo neutron and photon transport code. Available online at: <https://serpent.vtt.fi/> (Accessed July, 2024).

Stahl, D. (1982). *Fuels for research and test reactors, status review: july 1982*. ANL-83-5. Argonne National Laboratory. doi:10.2172/5942981

Valtavirta, V. (2015). *Multi-physics capabilities in serpent 2. Serpent and multi-physics workshop LPSC Grenoble*, 26–27.

Viitanen, T., and Leppänen, J. (2014). Target motion sampling temperature treatment technique with elevated basis cross-section temperatures. *Nucl. Sci. Eng.* 177, 77–89. doi:10.13182/NSE13-37



OPEN ACCESS

EDITED BY

Mark D. DeHart,
Idaho National Laboratory (DOE), United States

REVIEWED BY

Nicolas Stauff,
Argonne National Laboratory (DOE),
United States
John Carter,
Idaho National Laboratory (DOE), United States

*CORRESPONDENCE

William Reed Kendrick,
✉ rkendric@mit.edu

RECEIVED 07 March 2025

ACCEPTED 17 April 2025

PUBLISHED 29 April 2025

CITATION

Kendrick WR and Forget B (2025) Coupled neutronic-thermal-mechanical simulation of the KRUSTY heat pipe microreactor. *Front. Nucl. Eng.* 4:1589780.
doi: 10.3389/fnuen.2025.1589780

COPYRIGHT

© 2025 Kendrick and Forget. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Coupled neutronic-thermal-mechanical simulation of the KRUSTY heat pipe microreactor

William Reed Kendrick* and Benoit Forget

Computational Reactor Physics Group, Department of Nuclear Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, United States

Multiphysics analysis has become a common technique for nuclear reactor design validation, with neutronic-thermal analysis being the typical choice for understanding reactor dynamics. The concept of adding mechanical simulation such as thermal expansion to this coupling is still relatively new, however, and presents many computational challenges. While large reactors see relatively little neutronic impact from thermal expansion and may not warrant the challenge of undertaking this level of coupling, recent studies of microreactor geometries show that smaller reactors see larger impacts from thermal expansion. This work performs coupled neutronic-thermal-mechanical simulation of the Kilowatt Reactor Using Stirling TechnologY (KRUSTY) using OpenMC and Multiphysics Object-Oriented Simulation Environment in order to analyze the neutronic and thermal impact of including thermal expansion at steady state. The results show that while thermal expansion has a significant effect on global neutronic tallies, it has relatively minor impact on spatial heating rates or temperatures in the system. This remains true even when simulating a multiple heat pipe failure scenario to introduce thermal asymmetry.

KEYWORDS

neutronics, multiphysics, microreactor, KRUSTY, openmc, MOOSE, thermal expansion

1 Introduction

Nuclear reactors are environments of high temperature profiles and large neutron fluxes. These two characteristics are tied together through complicated feedback mechanisms. Neutron flux dictates fission rate, fission rate dictates heating, heating dictates temperature, and temperatures can cause significant changes in material microscopic cross-section, which finally impacts neutron flux. This feedback loop is one of the strongest that exists in nuclear reactors, providing an important control mechanism (Oka, 2013). This is why coupling of neutronic-thermal simulations is a widely seen reactor design methodology, particularly with the rise of high performance computing and parallelizable codes. These studies can be found for nearly every reactor type that has been conceptualized.

As computing architecture and simulation codes have advanced, there has been rising interest in adding an additional level of physics to the neutronic-thermal coupling. One notable addition would be including solid-mechanical simulation to account for geometric deformation of different parts of the reactor environment. During real operation of a nuclear reactor, the high temperature of the core causes materials to expand, neutron irradiation can cause swelling or contraction, and fuel materials will experience swelling due to fission gas and lattice damage. These geometric

changes impact the spatial neutron flux in the core, which in turn becomes another feedback loop for the reactor.

For most large-scale cores, like classic gigawatt-scale reactors, thermal expansion has a relatively limited effect on neutron flux due to the scale of the system (Edenius, 1976). For smaller reactor geometries, such as microreactors, this is not the case. The NASA-produced Kilowatt Reactor Using Stirling TechnologY (KRUSTY) (Gibson et al., 2018) is a good example of this, with related works reporting that roughly 85% of net reactor feedback is caused by thermal expansion of the fuel (Poston et al., 2020a).

This reactivity impact from thermal expansion provides an impetus for this work to implement and analyze a neutronic-thermal-mechanical simulation of the KRUSTY microreactor environment, with an additional study on the impact of heat pipe failure. Neutronics are solved in OpenMC, and thermal-mechanical modeling is performed with INL's Multiphysics Object-Oriented Simulation Environment (MOOSE). Of particular attention in this study is identifying how thermal expansion of the reactor features change spatial heating rates and material reaction rates, as well as the following temperature distributions in the core. Importantly, the results of this study will help suggest whether this sort of neutronic-thermal-mechanical coupled simulation is a necessary study for similarly small reactor types in the future. For example, if the thermal expansion causes significant enough changes in temperature distribution of the core, it may mean that a purely neutronic-thermal simulation is insufficient to predict core behavior. A test case on heat pipe failure is included in order to analyze how the results change when the temperature profile is asymmetric in the core due to failed heat pipes.

2 Background

Described in the following section are some of the general concepts and features that are necessary to understand the work of this study. These topics include a brief explanation of the KRUSTY reactor design with emphasis on why it has been chosen as the reactor of choice for this study, followed by an overview of the mechanics of thermal expansion and other geometric deformation sources commonly seen in nuclear reactors. This is followed by a description of Direct Accelerated Geometry Monte Carlo (DAGMC), a technique that is used extensively in the coupling method described in Section 3.1. Finally, because neutronic-thermal-mechanical coupled simulation is beginning to become more prevalent, some contemporary approaches to solving this sort of problem are discussed to contrast the ways in which the approach described by this study differs.

The KRUSTY reactor, as suggested by its full name, is a relatively simple and lightweight prototype reactor design intended to provide propulsion in outer space via heat exchange with Stirling engines (Gibson et al., 2018). The design went through extensive testing in early 2018, and several papers exist documenting the testing results (Gibson et al. (2018); David et al., 2020; Poston et al. (2020b); Poston et al. (2020a); Sanchez et al. (2020)). The geometry studied for this

work does not include the entire reactor-engine system, but only the core and surrounding shield and reflector structures.

A depiction of the representative core simulated in this work can be seen in Figure 1, colored by material. The core features a central, annular fuel block of U-7.65Mo, highly enriched to approximately 93% ^{235}U , and surrounded by a multilayer Molybdenum insulation (MLI). This insulation also surrounds the inward-facing sides of the heat pipes, and is labeled as "HPMLI" on Figure 1. On the radial perimeter of the fuel block are eight inlaid Sodium heat pipes, which begin in a lower reflector of beryllium oxide (BeO) and extend out of the core. The fuel, heat pipes, and lower reflector are encased in a Type 316 stainless steel (SS316) vacuum vessel. Radially surrounding this is a reflector of more BeO, followed by a radial shield of SS316. Below and above the vacuum vessel are shields of SS316 and boron carbide (b4c) which are used for control and shielding. In the center annulus is a beryllium (Be) target clad in aluminum (Al) and SS316, which functions as a neutron multiplier. The core during normal operation was designed to have a 4 kW thermal power, with fuel and heat pipe temperatures predicted to be slightly below 1100K. More detailed description of the design features can be found in David et al., 2020.

The mesh geometry representative of the KRUSTY reactor used in this work comes from the Virtual Test Bed Repository hosted by Idaho National Lab (Giudicelli et al., 2023), and is a result of the work of Dr. Kun Mo and Dr. Soon Kyu Lee of Argonne National Lab. This geometry serves as a relatively accurate model of the true KRUSTY reactor.

What should stand out from the previous descriptions of the core is that this is a fast spectrum system. There is no moderator, and it relies on reflectors to keep the reactor critical. From a neutronics perspective, this means temperature variation of the medium has less of an impact on neutron flux because Doppler broadening is most impactful for resonance absorption during the slowing-down process (Springer et al., 1969). Fast spectrum systems, however, are more reactive to neutron leakage because fast neutrons have a greater probability of escaping the reactor—hence the need for thick reflectors surrounding the core.

This leakage-constrained behavior is likely one of the reasons why thermal expansion is reported by Poston et al. (2020a) as having the dominant impact on reactivity. Previous work by the authors that focused on thermal expansion along with neutronic-thermal simulation of small modular reactors found that the main impact of thermal expansion was on neutron leakage (Kendrick, 2024). This suggests that KRUSTY is a very good test case for analyzing the impact of thermal expansion on neutronic-thermal simulation results.

As a sanity-verification method, the fact that volumetric expansion causes increased neutron leakage can be proven with a simple toy problem. Imagine a regular cube with a neutron born at the center of the cube. The task is to determine how probability of leakage of that neutron changes after volumetric expansion. If there is uniform expansion of 10%, in order to conserve mass the density must decrease by the reciprocal, 0.909. Additionally, the distance the neutron must travel to any surface on the cube is directly related to the length of the sides of the cube, which have increased by a factor of $1.1^{1/3}$. If one uses the formula for probability of escape of the neutron, $P = e^{-\sigma N x}$, - one can calculate the relative difference in the

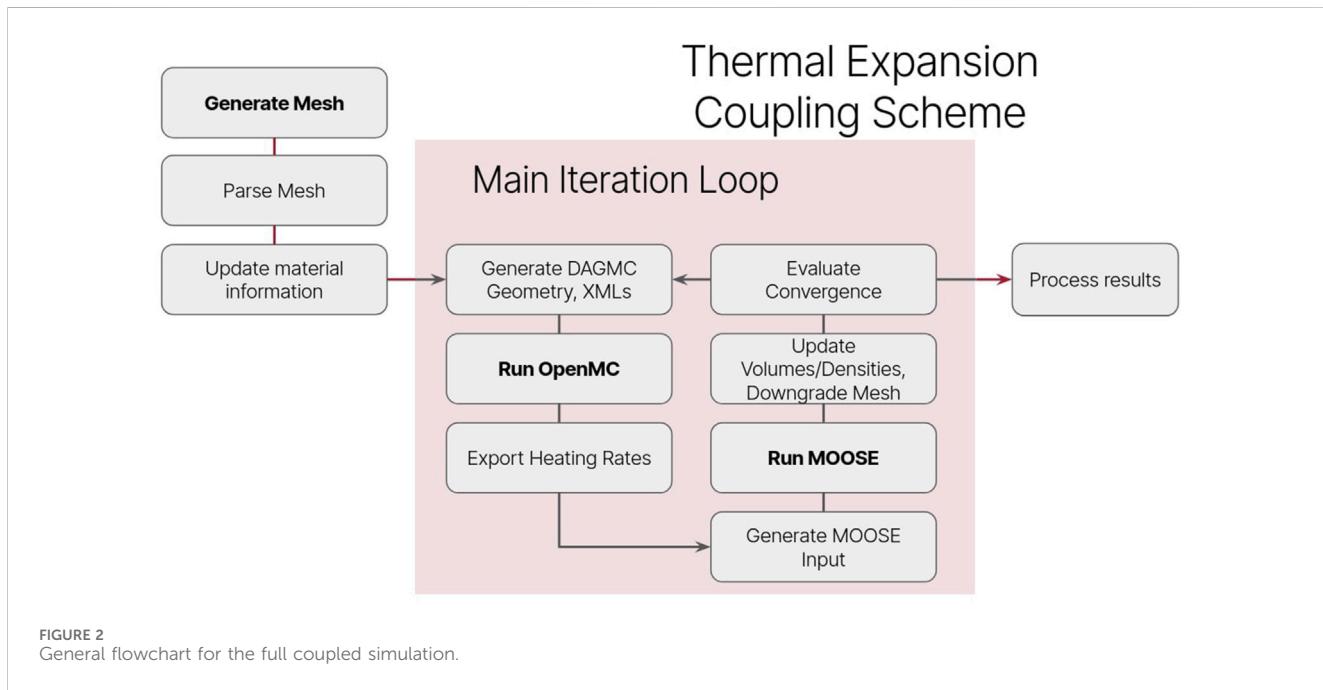
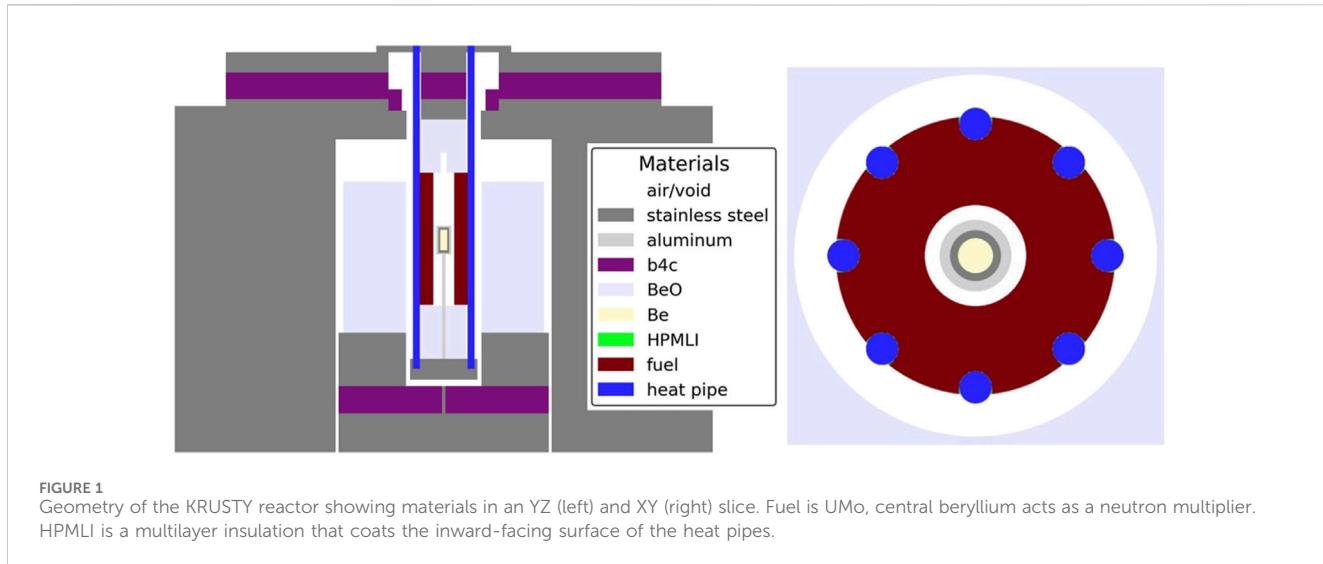


FIGURE 2
General flowchart for the full coupled simulation.

probability of escape for the neutron after expansion via Equation 1:

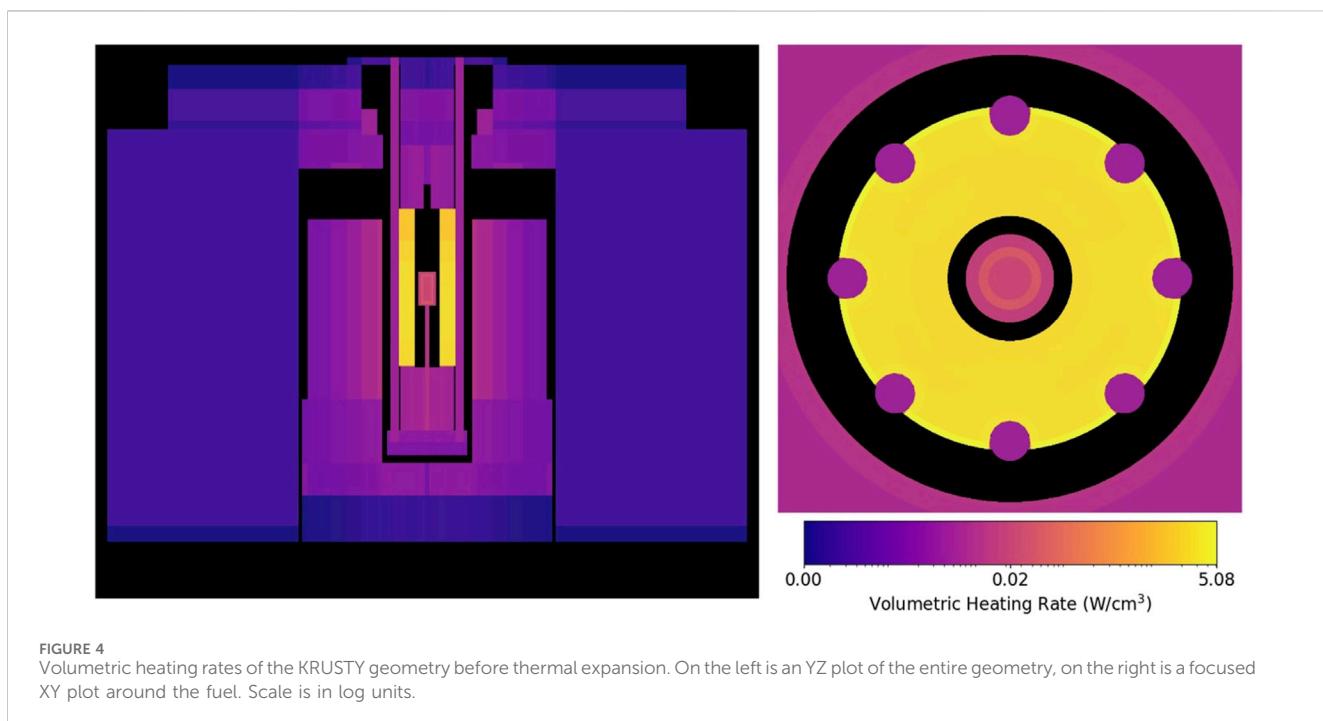
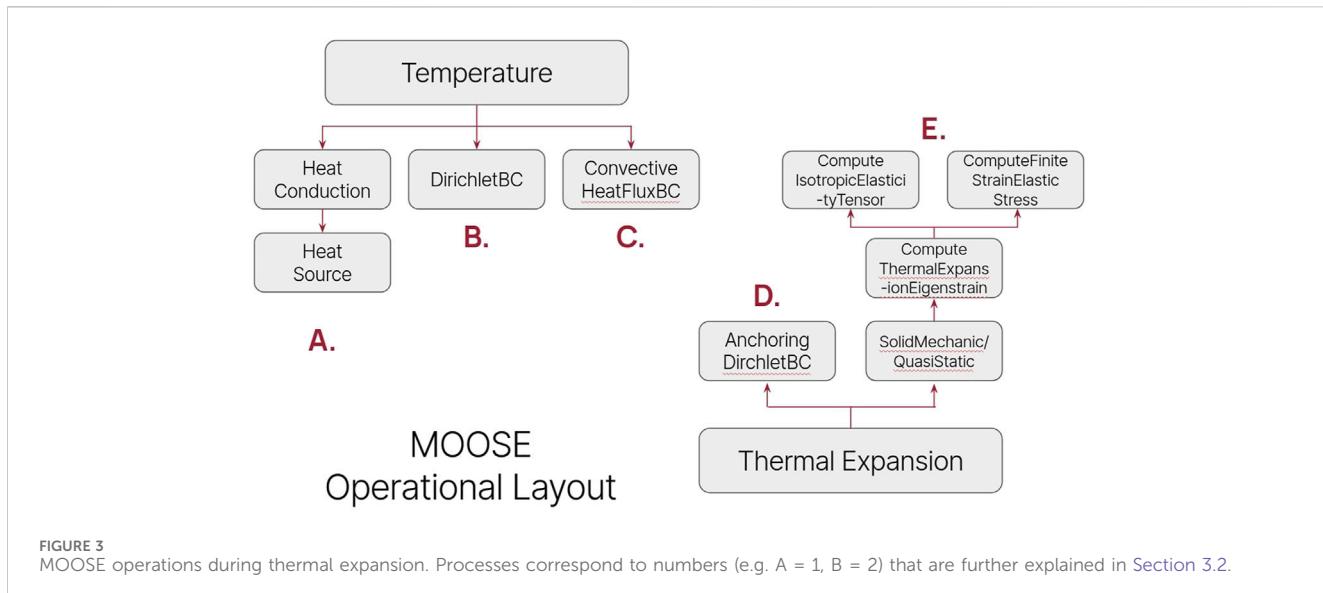
$$\frac{P_1 - P_0}{P_0} = \frac{e^{-\sigma N_1 x_1} - e^{-\sigma N_0 x_0}}{e^{-\sigma N_0 x_0}} = \frac{e^{-\sigma(0.909N_0)(1.1^{1/3}x_0)}}{e^{-\sigma N_0 x_0}} - 1 = e^{0.0616\sigma N_0 x_0} - 1 \quad (1)$$

where σ is the microscopic cross-section of the cube's material, N is the number density, x is the distance the neutron must travel to escape, and 0 and 1 subscripts represent pre and post-expansion, respectively. If one assumes some values for the toy problem, namely, a macroscopic cross-section of 0.5 cm^{-1} and a distance-to-travel of 0.5 cm , it can be calculated that a neutron born in

the center of the cube will have a 1.55% higher chance of leaking after 10% volumetric expansion. Essentially, a neutron will see proportionally less dense material compared to how much further it must travel after expansion. This is why thermal expansion increases neutron leakage in the system, even though mass is conserved.

2.1 Thermal expansion and geometric deformation

Thermal expansion is a well defined topic, being a phenomenon that humans have recorded for hundreds of years and have



experienced for far more. Thermal expansion occurs on the molecular level as increased temperatures are a manifestation of an increase in kinetic energy. This increased kinetic energy results in more frequent neighboring body interactions, which translate into vibrations. These vibrations lead to an increase in distance from the surrounding atoms because of the anharmonicity between the attractive and repulsive forces that govern the lattice. Whether a material will expand or contract with increasing temperature depends on the balance of these forces that is specific to the material and state of the material in question (Bower, 2009).

The core of a nuclear reactor is an environment with extreme temperatures and sees significant thermal expansion of its materials,

isotropically and anisotropically. Quantifying the per-degree effect of temperature change on neutron balance via a reactivity coefficient has been a longstanding practice for reactor engineers in order to understand transient reactor behavior and inform point kinetic models. This quantity rarely includes the thermal expansion effect for the solid materials in the core, however. This is because in traditional light-water reactors at steady state, the effect of thermal expansion on eigenvalue (the ratio of neutrons produced per neutron destroyed) is fairly minor, with the largest effect being a change in the leakage of fast spectrum neutrons (Edenius, 1976).

While this does have more of an impact in KRUSTY-sized reactors, as previously noted, thermal expansion is still very much a

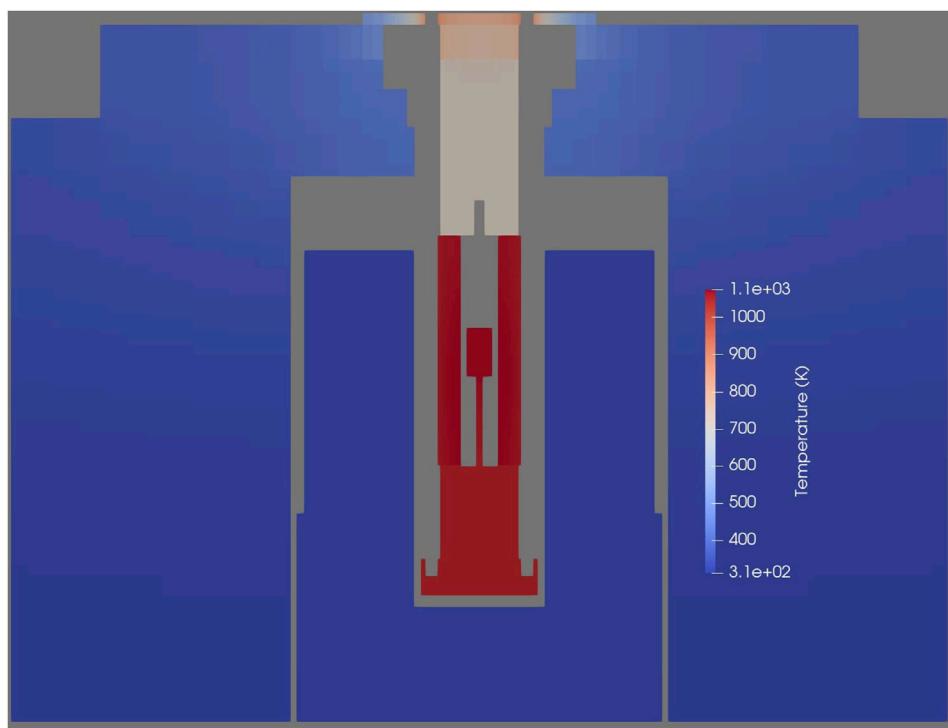


FIGURE 5
YZ slice of the KRUSTY geometry showing temperature before thermal expansion. The fuel has insulating boundary conditions on all sides, causing the thermal behavior in the axial reflectors.

TABLE 1 Change in material-based neutronic tallies due to thermal expansion. “sp” stands for source particle, as OpenMC tallies in eigenvalue mode are normalized per source particle. Δ is the difference between expanded and non-expanded results. HPMLI stands for heat pipe multi-layer insulation, the insulation material that is on the inside surface of the heat pipes.

Material	Absorption rate		Heating tally		Power	
	Δ (1/sp)	Rel. Δ (%)	Δ (eV/sp)	Rel. Δ (%)	Δ (W)	Rel. Δ (%)
SS316	$(7.5 \pm 1.5)e-4$	0.28 ± 0.07	$(-1.7 \pm 0.1)e+4$	-0.64 ± 0.05	3.01 ± 0.09	1.77 ± 0.05
Al	$(0.1 \pm 2.0)e-7$	0.0 ± 1.1	$(-1.1 \pm 0.2)e+2$	-1.7 ± 0.3	$(2.5 \pm 1.3)e-3$	0.66 ± 0.34
B4C	$(4.8 \pm 0.0)e-3$	2.90 ± 0.06	$(3.1 \pm 0.3)e+3$	0.63 ± 0.06	$(9.5 \pm 0.2)e-1$	3.08 ± 0.06
BeO	$(3.5 \pm 0.0)e-4$	1.86 ± 0.04	$(6.9 \pm 0.4)e+3$	0.47 ± 0.02	2.73 ± 0.03	2.90 ± 0.03
Be	$(-4.3 \pm 3.2)e-7$	-0.9 ± 0.7	$(-5.7 \pm 1.0)e+1$	-2.7 ± 0.5	$(-4.0 \pm 6.0)e+4$	-0.3 ± 0.5
HPMLI	$(-2.0 \pm 0.1)e-5$	-19.1 ± 0.9	$(-1.7 \pm 0.1)e+2$	-5.9 ± 0.5	$(-6.7 \pm 0.9)e-3$	-3.6 ± 0.5
UMo	$(-9.8 \pm 0.1)e-3$	-2.06 ± 0.02	$(-1.9 \pm 0.0)e+6$	-2.51 ± 0.02	-6.84 ± 0.10	-0.15 ± 0.00
Heat Pipe	$(-5.1 \pm 0.4)e-5$	-3.74 ± 0.27	$(1.4 \pm 0.1)e+3$	3.33 ± 0.15	0.16 ± 0.00	5.84 ± 0.15

topic of analysis for reactor engineers regardless of reactor design. This is because of the concern of material failure due to mechanical stress, either within a material itself or at a contact point with another material.

For the purposes of this work, neither anisotropic expansion (e.g., mono-crystal behavior) nor potential mechanical failure are taken into account, although anisotropy may be possible to model by defining material lattice planes. Instead, the focus of this work is on the impact of

non-uniform, isotropic-in-cell, elastic thermal expansion on core neutronic behavior and simulating the thermal-neutronic feedback loop.

Although thermal expansion is the only deformation mechanism included in this study, it is not the only method of geometric deformation that nuclear reactors experience during operation. A notable method of geometric deformation is fission-induced swelling of the fuel material. This process occurs as fission products collide and decelerate in the fuel material, creating vacancies and displacements as

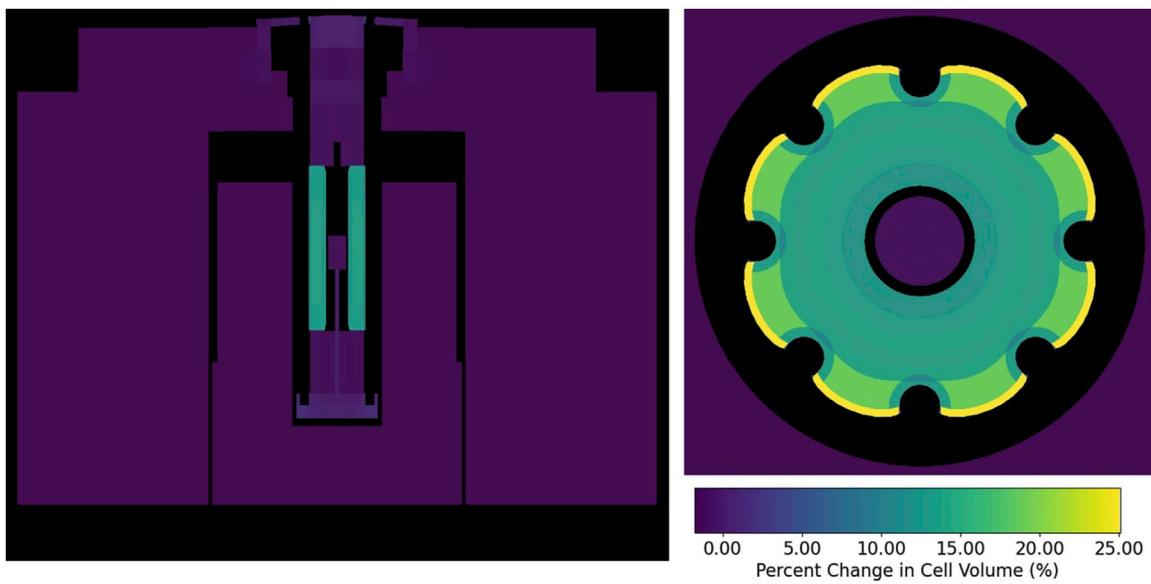


FIGURE 6
Percent changes in cell volume of the KRUSTY geometry after thermal expansion. On the left is a YZ plot of the entire geometry, on the right is a focused XY plot around the fuel.

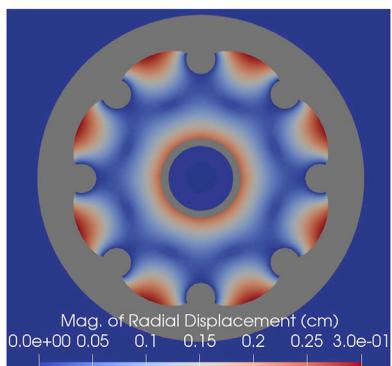


FIGURE 7
Displacement in the radial direction in units of centimeters, focused on the fuel region. Value is calculated by taking the square root of the sum of the squared X and Y displacements.

well as depositing themselves throughout. These fission products are a mix of solid and gaseous states, and result in a swelling effect for the fuel (Massih, 1988).

The gaseous fission products are especially accelerative for the swelling process, as the fission products will migrate to grain boundaries and cavities. These gasses then experience larger thermal expansion than their solid-state counterparts, increasing expansion and encouraging material fracture (Lietzke, 1970). Because of the inherent risk of highly radioactive, gaseous fission products escaping the fuel, the U.S. Nuclear Regulator Commission includes fission gas release and swelling modeling as a part of the fuel qualification process (Drezewiecki et al., 2021). This has led to fuel swelling being a well-documented and modeled phenomenon (Andrews, 2012).

While this study neglects these non-temperature-dependent methods of geometric deformation that do regularly occur in reactor cores, the sole study of thermal expansion and its effects is useful as it provides a baseline for future deformation cases. Being able to define that x centimeters of axial displacement causes y change in reactivity allows the hypothetical next study that includes some of the previously mentioned methods to estimate neutronic impact even before simulation. Additionally, the method of coupling described in this work is applicable for all the above deformation modes, given a model exists that describes the deformation.

2.2 DAGMC

The largest hurdle involved in neutronic-thermal coupling is reconciling the difference in native environment in which each type of simulation operates. Thermal coupling is typically performed on an unstructured mesh geometry via the finite element method where calculations are performed iteratively over quadrature points located within mesh elements. Neutronic calculations, when simulated via Monte Carlo particle transport, are typically performed on Constructive Solid Geometry (CSG) geometries where surfaces are defined with analytical geometric expressions and cell spatial regions are defined by Boolean half-spaces. In order to couple the two together, something must be altered. This work utilizes Direct Accelerated Geometry for Monte Carlo (DAGMC) to alter the geometry that particle transport occurs on.

DAGMC is a software library developed for neutronic modeling of fusion reactor geometries, allowing the user to translate Computer-Aided Design (CAD) geometries into Monte Carlo-solvable inputs (Tautges et al., 2009). The software has since seen use in both fission and fusion applications, and features an ability to couple with the Sandia National Laboratories-based Coreform Cubit

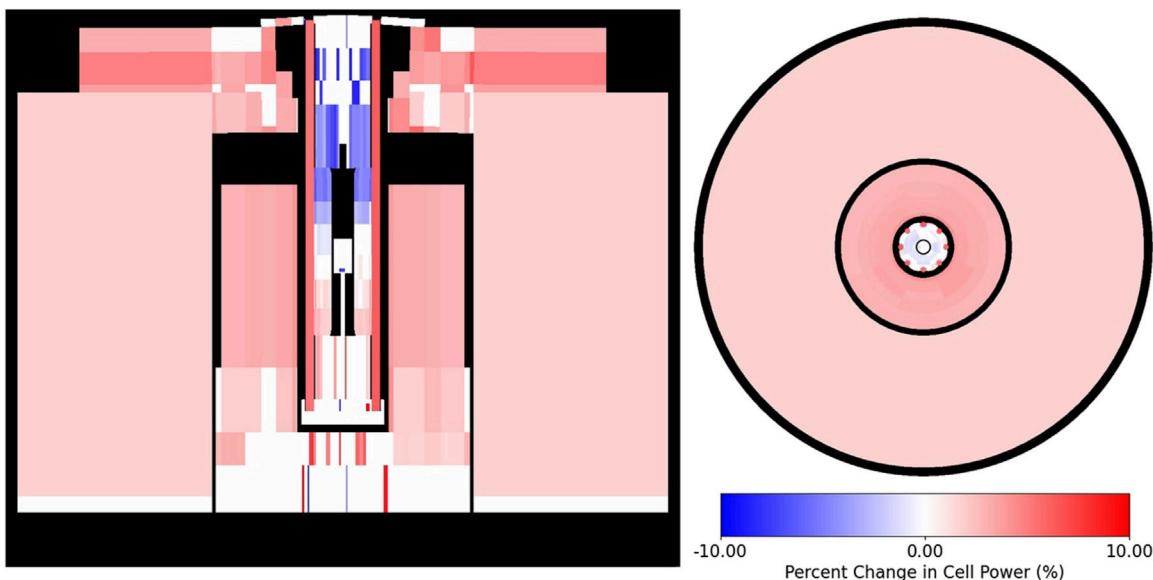


FIGURE 8

Percent change in normalized power for each cell in the KRUSTY reactor. On the left is a YZ plot, on the right is a XY plot. This figure assumes constant power pre and post-expansion. Values with relative standard deviations greater than 50% were excluded from the figure. Axial shift in power is caused by fast neutron leakage from the thermally expanded fuel.

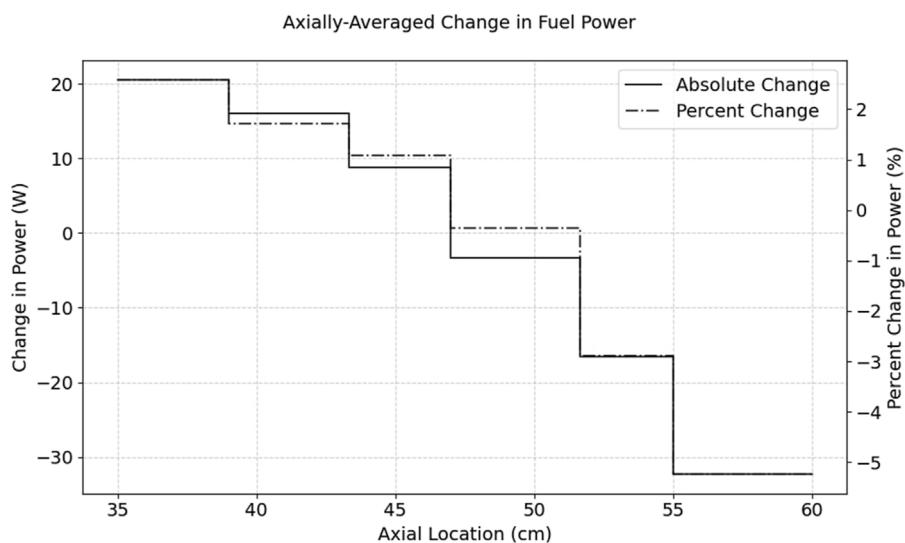


FIGURE 9

Absolute and percent change in axial power for the UMo fuel in the KRUSTY core as a result of thermal expansion. Axial shift in power is seen due to fast neutron leakage from expanded fuel.

software to automate CAD-to-CSG transfers (Blacker et al., 1994). Coreform Cubit has a feature capable of exporting DAGMC geometries directly from coarse meshes, of which this study takes advantage (Coreform, 2024). This facet geometry replaces the geometry block of information traditionally used for OpenMC problem description. In order to do so, temperatures and material assignments are made on the mesh and are stored within the facet geometry.

The DAGMC geometry consists of planar body representations, on which OpenMC is well-equipped to perform transport calculations. DAGMC's integration with OpenMC is well documented and there exists a robust literature documenting its use (Davis et al., 2020). While the planar bodies work well for representing tessellated surfaces, this does still mean that surfaces that were originally second-order or higher, such as the outer surface of a cylinder, will incur severe computational penalties.

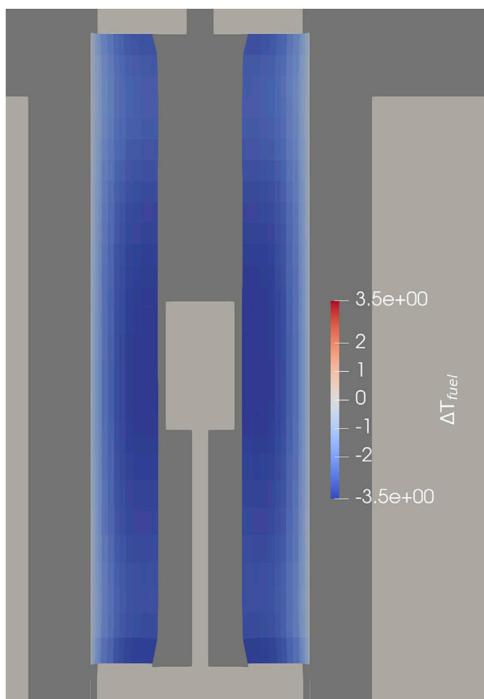


FIGURE 10
YZ slice of the fuel region showing change in fuel temperature due to changed neutronic heating rates post-thermal expansion. Comparison is against a model without thermal expansion.

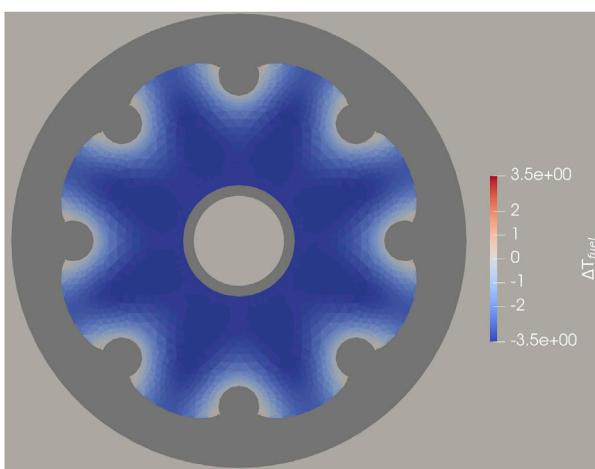


FIGURE 11
XY slice of the fuel region showing change in fuel temperature due to changed neutronic heating rates post-thermal expansion. Comparison is against a model without thermal expansion.

2.3 Contemporary approaches to multiphysics simulation

As previously mentioned in the beginning of this section, the concept of coupling thermal expansion to neutronic-thermal simulation is not a novel one; in fact, the last 5 years have seen a number of papers published on this topic featuring a wide range of

coupling methods and reactor geometries (Hu et al., 2019; Ma et al., 2021; Yan et al., 2020; Chen C. et al., 2023; Pope and Lum, 2020; Xiao et al., 2022; Jeong et al., 2023; Walker et al., 2022; Aldebie et al., 2024; Sterbentz et al., 2017; Chen H. et al., 2023).

The majority of these works have been related to reactors larger than KRUSTY. The work in Stauff et al. (2021) details the use of MOOSE coupled with Griffin (Lee et al., 2021) to simulate thermal stress and strain on a hypothetical heat pipe small modular reactor. Stauff notes a maximum axial displacement of 1.5 cm and a maximum stress located between neighboring heat pipes, but does not note any impact on neutronic or thermal behavior due to thermal expansion. The author's work on the same geometry, utilizing the methods described in this work, saw that this was a fair conclusion, as expansion was only impactful on leakage (Kendrick, 2024). In a study on a fast-spectrum reactor, Walker reported an impact of roughly 7 pcm per Kelvin due to combined axial expansion of fuel and radial expansion of reflector and cladding (Walker et al., 2022).

This sort of analysis for microreactors is slightly rarer, but there does exist a study on KRUSTY produced by Chen (Chen H. et al., 2023). Chen utilizes OpenMC to solve the initial power distribution, then uses MOOSE to solve the thermal profile of the core, along with a tensor mechanics solve to simulate thermal expansion. This work does not feed the deformed geometry back into OpenMC, however, so thermal expansion's impact on neutronic and therefore thermal behavior is not noted. Chen's work is focused on transient behavior analysis by tying a point kinetics model to adjust power amplitude while maintaining the original spatial power distribution. Poston's work (Poston et al., 2020a) is the closest to reporting an effect on reactivity due to thermal expansion, as the author has mentioned earlier. The details of the coupling scheme are not clear, however, except for MCNP being the particle transport code used and FRINK being used for the rest of the physics.

To generalize the works of this nature thus far, it is useful to categorize them by the methodology used for neutronics. Often seen, like in the work of Stauff et al. (2021), is solving neutron/photon behavior via deterministic methods like diffusion. This is particularly appealing because this sort of transport solve functions well in an unstructured mesh environment. A good example of the utility of this is the Griffin code that can operate in the MOOSE framework, allowing for rapid iteration of a coupled solution. This ability becomes particularly advantageous when the mesh deforms during the solve. The downside to this sort of deterministic calculation is well known; deterministic methods require competent cross-section generation for the geometry in order to preserve reaction rates. This is sometimes done via a Monte Carlo simulation (Redmond, 1997) which cuts into the potential computation time gains from using deterministic methods in the first place, and geometric deformations that strongly affect neutron flux would potentially require updated cross-sections mid-solve. Couple this with some of the inherent inaccuracies dependent on the deterministic method used, and there may be a question of neutronic uncertainty that is hard to quantify.

The alternative to this is to rely on Monte Carlo simulation for neutronics. Monte Carlo is still considered too computationally expensive for Light Water Reactors (LWRs) full-core calculations, even with modern computing advances. However, a full-core calculation of a microreactor like KRUSTY is far different, given

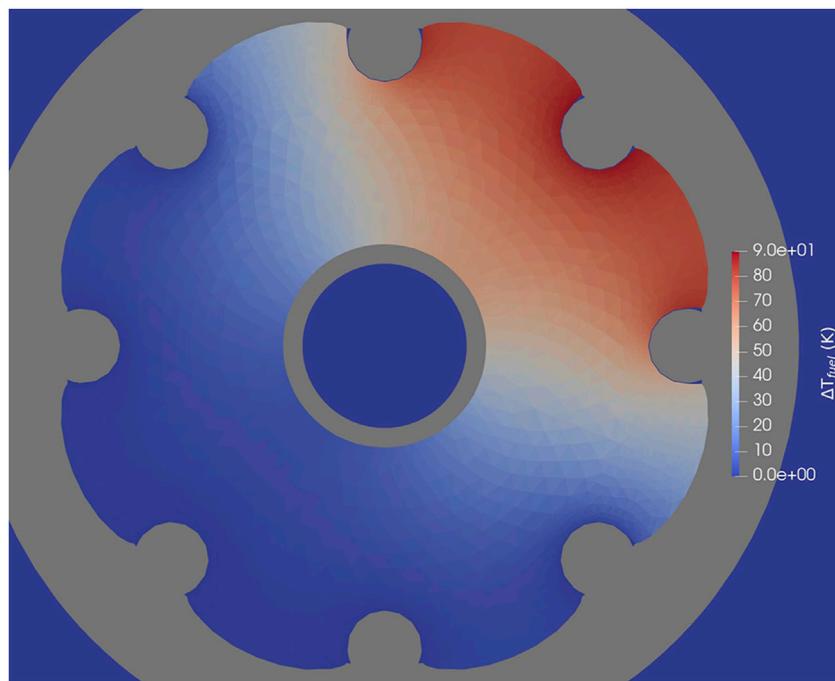


FIGURE 12
XY slice of the fuel region showing change in temperature of the triple heat pipe failure expansion case compared to the non-heat pipe failure expansion case.

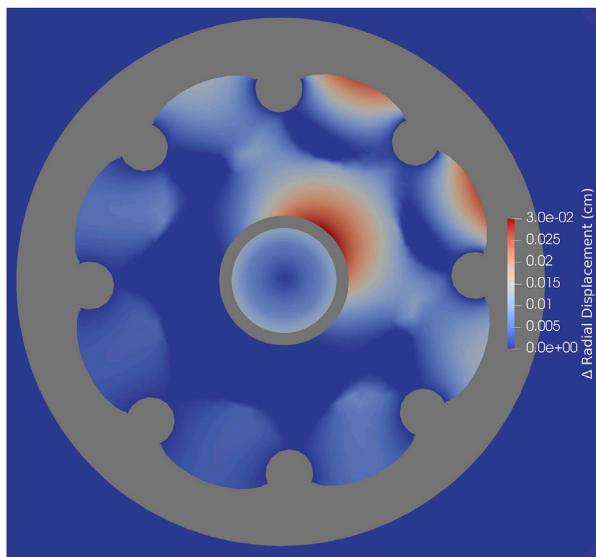


FIGURE 13
XY slice of fuel region showing comparison of displacement in the radial direction in the triple heat pipe failure expansion case compared to the non-heat pipe failure expansion case. Value is calculated by taking the square root of the sum of the squared X and Y displacements.

the dimensions of the KRUSTY reactor are smaller than the dimensions of a single AP-1000 assembly. The challenge is reconciling how to transfer information typically generated in a Constructive Solid Geometry environment to the thermal solve in an

unstructured mesh environment. This work's approach is to use DAGMC to skin CAD geometries into planar bodies and solve particle transport in this CAD-derived environment, but this method has nuances as mentioned in Sections 2.2 and 3.3. A simpler alternative is to build a 1-to-1 CSG representation of the mesh geometry, but that approach fails when non-uniform deformation occurs.

It is clear that neutronic-thermal-mechanical studies are still in their infancy. As will be obvious in the following section detailing the coupling methodology, adding solid mechanics introduces significant complexity that often requires reducing the accuracy in other areas of the solve. It is this author's hope that this work serves to be another stepping stone in the pathway towards a well-defined process for tackling this sort of simulation in the future.

3 Methods

This section details the methods utilized to perform neutronic-thermal-mechanical simulation on the KRUSTY reactor geometry. First, this section provides a short summary of the two main codes used in this work, OpenMC and MOOSE. This is followed by three subsections, beginning with a detailed explanation of the coupling scheme used to connect the neutronic and thermal-mechanical solves. The final two subsections respectively detail some of the complexities and assumptions used in the MOOSE and OpenMC solves.

MOOSE is a finite element physics framework initially developed at Idaho National Laboratory (Giudicelli et al., 2024). It features extremely parallelizable code for high efficiency on

computing clusters, along with well documented and user-friendly methods for implementing one's own physics kernels. MOOSE is used for all thermal conduction and thermo-mechanics simulation in this work, relying on modules that are readily available as part of the open-source MOOSE environment (Shemon et al., 2023; Adhikary et al., 2016). All simulations using MOOSE utilize the finite element method with Lagrangian variable representation. The Newton scheme of non-linear solving was applied, with the default MOOSE Jacobian preconditioning.

OpenMC is an open-source Monte Carlo neutron and photon transport simulation code initially developed by Massachusetts Institute of Technology (Romano et al., 2015). OpenMC can perform steady state eigenvalue calculations on complicated geometries in both continuous energy and multigroup with tallying capabilities for heat deposition (Romano, 2020). OpenMC was developed with a strong emphasis on parallelism, taking advantage of the fact that the Monte Carlo method has some inherent advantages for parallel computing (Rosenthal, 2000). This parallelization allows the software to scale well with larger system architectures like supercomputers, which this work utilizes.

3.1 Coupling scheme

A flowchart diagram depicting the coupling scheme is included in Figure 2, showing the flow from mesh generation to final results compilation. The initial work done before the main iteration is primarily focused on mesh generation and parsing. Mesh generation is performed using the MOOSE Reactor module (Shemon et al., 2023). This module allows regularized mesh generation for reactor core geometries, in particular geometries with repeated features such as Cartesian and hexagonal lattices.

The process of parsing the mesh occurs following this generation. Because the OpenMC-DAGMC geometry is derived from the mesh, all cell references now must be references to mesh volume IDs. Building the arrays of IDs that establish what materials are found in each volume as well as calculating and tabulating volumetric data for each mesh volume is handled during the parsing step. This is accomplished via Coreform Cubit's Python API, without which a manual inspection of the mesh would be required. A dictionary cataloging each core material and its properties, along with mesh volume IDs that feature that material, is instantiated and populated at this time as well. The volume of each cell is recorded in order to modify material densities post-expansion to conserve mass in the system.

3.2 MOOSE and mesh deformation

One of MOOSE's strongest selling points for this sort of study is its library of existing modules that cover a wide range of physics kernels. Included amongst these is the Solid Mechanics module, which contains the specific material and physics blocks used to simulate thermal expansion in the KRUSTY reactor. An operations outline of the MOOSE input can be seen in Figure 3. This figure does not include post-processors and the material blocks used to apply thermal and mechanical material properties. An explanation for

each subset of the input, corresponding to the connected letters seen in Figure 3, is as follows:

1. Kernel activating heat conduction and the input volumetric heating rates for every mesh volume. Each volume has the OpenMC-generated volumetric heating rate applied via the `HeatSource` block.
2. Dirichlet boundary to represent the heat pipes for main heat removal. As a simple representation of the heat removal functions of real-life heat pipes, `DirichletBCs` were applied to all heat pipe surfaces with a value of 1050K to mimic the results from Poston et al. (2020a).
3. Outer, top, and bottom boundary heat removal. This is an arbitrary natural convection heat loss representation, in order to mimic natural convection with ambient air. This inclusion allows for a more realistic temperature profile for the outer air materials, given their otherwise isolated condition.
4. Anchoring Boundary Conditions. Because thermal expansion is isotropic in-cell, displacement at certain boundaries need to be anchored with a `DirichletBC` in order to avoid mesh overlap in particular. Note that a truly accurate simulation would not need these boundaries, but because of the removal of air in particular (for reasons stated later in this section), MOOSE does not recognize overlapping meshes.
5. Instantiation and solving of thermal expansion-based eigenstrains, strains, and stress. The first block, `SolidMechanics/QuasiStatic` is a MOOSE Action object that instantiates all the needed data for a Solid Mechanics problem. `ComputeThermalExpansionEigenstrain` calculates the eigenstrain tensor resulting from isotropic thermal expansion, given the temperature at the quadrature point versus the input stress-free temperature. These eigenstrains lead to stresses calculated by `ComputeFiniteStrainElasticStress` given the material properties computed by a `ComputeIsotropicElasticityTensor`.

The result of these inputs is a solution that thermally solves temperature in the system and displaces mesh node points based on thermal expansion eigenstrain tensors. The output mesh is now deformed, and can be converted via DAGMC for an OpenMC particle transport solve.

While the process seen in Figure 2 does not involve any custom code, it does introduce two significant complications needing to be mentioned. Firstly, the use of Dirichlet boundary conditions for the heat pipes results in a loss of heat in the system. This non-conservation arises because the MOOSE solve uses the finite element method, which is globally conservative but not necessarily locally conservative (LeVeque, 2002). The finite element equations at the boundary solve one value, but the Dirichlet boundary is strongly enforced and overwrites the calculated values, thus the solution may not satisfy conservation (Hubbard et al., 2009; Hughes et al., 2000). Using higher order meshes reduces this error, and using a Neumann boundary condition removes the error. The results of this study do use Dirichlet boundaries so this non-conservation exists, and may slightly affect reported temperature profiles.

The second complication has to do with mesh element inversion. This occurs when the displacement of one node moves

that node to a new position such that when the volume of the element is calculated, it returns a negative value. This causes the Jacobian of that mesh element to turn negative as well, and causes a simulation failure. Areas where this tends to occur are places with coarse meshes, thin layers, and high displacement regions. In order to avoid these sorts of issues, the air mesh blocks in the reactor geometry and two thin layers of aluminum and steel (representing the ring clamp and vacuum can) were removed. Neutronically, these material removals are inconsequential. Air is not dense enough to significantly impact neutron transport, and the aluminum and steel layers are very thin (0.318 cm thickness for the ring clamp, 0.305 cm thickness for the vacuum can).

These features are thermally impactful. In reality, the vacuum can has multi-layer insulation on its fuel-facing side, keeping the air and emissivity from transferring much heat radially. Our removal of the can and air material nearly perfectly insulates the radial reflector, in a similar manner. Comparing the reported radial reflector operating temperatures [Poston et al. \(2020a\)](#) with the pre-expansion temperature profile results in [Figure 7](#), these temperatures are very close, 343 K *versus* 330 K, respectively.

The removal of air does, however, affect the temperature of the axial reflectors, as without it there is no method of heat removal besides the heat pipes at 1050 K. This causes temperatures of 1050 K and 700 K seen in [Figure 7](#) *versus* the Poston reported 473 K. The higher temperature in this paper's simulation will result in an overestimation of expansion for the axial reflectors. That being said, in the volume expansion results seen in [Figure 8](#), the axial reflectors still experience much less expansion *versus* in the UMo fuel. Based on the work of [Jamison et al. \(2020\)](#) used for this simulation, at 1050 K, BeO would have a thermal expansion coefficient of 1.04E-5, while UMo would have a thermal expansion coefficient of 2.06E-4.

3.3 OpenMC and heat pipe representation

As noted in the previous subsection, heat pipe representation during the thermal/mechanical solve in MOOSE is essentially a surface representation. The interior of the heat pipe is not utilized in the computational model. In fact, the inclusion of a material inside the heat pipe often will negatively impact the heat pipe representation, as it influences the heat flux recorded at the heat pipe boundary. One way of handling this is to remove the mesh blocks that describe the heat pipes, as only the surface of the blocks are needed ([Kendrick, 2024](#)). Although this work does not use a true heat pipe model, the same process of removing heat pipe mesh elements has been done.

This presents a problem, however. For a coupling method that performs neutronics on the DAGMC facet geometry based on the mesh output, this would mean removing the heat pipes from the neutronic solve as well. Unlike the previously mentioned removal of air from the mesh geometry, heat pipes are neutronically an important feature due to the stainless steel casing and sodium fluid. Neglecting this would make the neutronic results significantly less realistic.

The solution the author found was to insert heat pipes into the geometry on the OpenMC side by importing the DAGMC geometry not as the root universe, but instead as a cell. This is a somewhat

complicated process of geometry definition, where the region that the DAGMC cell inhabits is defined by a set of core bounding surfaces and by excluding every heat pipe region. This requires that the boundaries of the heat pipes be excluded from any mesh deformation such that the OpenMC heat pipes do not accidentally overwrite any fuel material. This does cause a slight “flowering” radial expansion shape for the fuel, seen in [Figure 8](#), rather than pushing the heat pipes outwards along with the fuel, which would be seen in reality.

Another side effect of this arises from the fact that the cylindrical regions for the heat pipes do not exactly fit the empty regions in the mesh geometry, due to the tessellation from meshing. This causes slight gaps at the boundary between the heat pipe and the fuel/reflector. These gaps result in a minor increase of streaming from the core that would not exist in real life. However, that effect is small compared to the results of completely excluding heat pipes.

4 Results

The KRUSTY reactor geometry was simulated using the aforementioned coupling method, with inputs of 5 kW thermal power for the reactor, 300 K ambient temperature surrounding the core, and 1050 K heat pipes. Neutronics are solved steady state, so the conversion from heating tallies to power deposited is done via power normalization at 5 kW, so power is constant for both pre and post-expansion. The MOOSE thermal and mechanics solve is the steady state result, however, the solve is performed in transient calculation to “ease in” to the final solution and avoid mesh inversion.

Convergence in the iteration process was achieved by monitoring the residual of the neutronic eigenvalue as well as the residual of the change in volume for all mesh volumes. Material properties such as thermal expansion coefficients come primarily from an Argonne National Laboratory technical study on UMo, see [Jamison et al. \(2020\)](#).

The initial OpenMC solve of the KRUSTY geometry results in heating rates that are visualized in [Figure 4](#). These heating rates are caused by neutron and photon heating, including fission energy deposition. The vast majority of power is deposited in the fuel, with the rest being spread outwards. There is minimal difference in heating rate between materials (besides the fuel) because the neutron spectrum is fast, thus distance from the fuel is the dominating contributor to heating rate. The temperature profile of the reactor can be seen in [Figure 5](#). The axial BeO reflector regions see high temperatures due to the lack of air-driven heat removal, as mentioned in [Section 3.2](#). Pre-expansion, the core has an eigenvalue of 1.00055 ± 0.00013 using ENDF/B-VIII.0 microscopic cross-sections, which aligns well with MCNP and Serpent results included in [Table 1](#) of [Cao et al. \(2024\)](#). Pre-expansion neutron leakage is $12.89\% \pm 0.005$, and the thermal fission factor (the proportion of fissions that are from thermal neutrons) is $7.82\% \pm 0.008$, confirming that the reactor is indeed a fast-fission dominated system.

The temperature profile in [Figure 5](#) drives the thermal expansion seen in [Figures 6, 7](#). The primary expansion occurs in the fuel, seeing a maximum of 25% volume increase in the edge cells of the fuel. [Figure 7](#) shows that the radial expansion of the fuel is roughly 0.3 cm

at its peak, and contracts in the center annulus at a slightly lower magnitude. As noted in [Section 3.3](#), the necessary constraint on expansion at heat pipe boundaries potentially causes a more peaked radial expansion, and in real life the heat pipes would be pushed outwards as well. Besides the radial expansion of the fuel, there is also a general axial expansion in the central column of the reactor that pushes the top of the reactor upwards. Thermal expansion is negligible in the outer BeO reflectors and the further outer SS316 blocks.

The neutronic effect of this expansion is significant for global neutronic parameters. Eigenvalue drops by 1,421 pcm to 0.98634 ± 0.00012 , neutron leakage increases by 3.9% to a leakage rate of 13.39%, and the thermal fission factor increases by 4.3% to a value of 8.16%. In a similar study reported by Poston [Poston et al. \(2020a\)](#), the fuel temperature deficit on reactivity was recorded as -148.1 cents, which, with an effective delayed neutron fraction (β_{eff}) of 0.00688, corresponds to an eigenvalue impact of -1,019 pcm. The eigenvalue impact of this study is not isolated to only fuel expansion, so it makes sense that this work would record a larger response.

The reason for this strong effect has been referenced at the end of [Section 2](#). Volumetric expansion makes a material seem more transparent to neutrons, and the more expansion the more transparent it appears. The UMo fuel has by far the largest expansion which increases neutron leakage from the fuel, particularly for fast neutrons, causing fewer fast fissions (raising thermal fission factor) and hardening the spectrum of the overall reactor. After expansion, the proportion of the flux that is thermal (less than 4 eV) fell by about 1%, while the flux proportion that is fast (greater than 0.1 MeV) grew by 4.2%.

This increased leakage from the fuel leads to changes in the absorption tallies, heating tallies, and power calculations for all materials in the core, seen in [Table 1](#). Note that the absorption rate and heating tally are normalized per source particle, while power is normalized to a constant 5 kW. The fact that the total heating tally (energy deposited per source particle) decreased by 2.37% after expansion suggests that fission rate would need to be increased by a proportional 2.37% to maintain pre-expansion power. After power normalization, most of the power remains in the fuel, which only sees a roughly 7 W difference, which has mostly moved to the SS316 and BeO in the system. One other quantity to note is that boron carbide absorption rate increased by 2.9% after expansion. Increasing fission rate to maintain power would mean a nominal increase beyond 2.9% in actual absorptions per second for the absorber material. This could impact expected lifetime for the boron carbide.

Because of the upwards axial expansion in the center of the core, the highest leakage from the fuel is at the upper axial region. After power normalization, this results in an axial shift in the power of the fuel, seen in [Figures 8, 9](#). The fuel is divided into six axial layers, the upper layer of which decreases in power by 5%, while the lowest layer increases in power by 2.5%. This spatial shift in power leads to a resulting change in temperature for the fuel, seen in [Figures 10, 11](#). Both figures show a maximum decrease of 3.5 K in the fuel due to thermal expansion-induced power shift. This level of temperature change is unlikely to be of note to reactor designers, though the axial shift in power does mean that depletion calculations for the fuel may

over-predict burnup at the top of the fuel and under-predict burnup at the bottom.

4.1 Heat pipe failure case

In order to evaluate if heat pipe failure changes any of these results, a “worst-case scenario” event was evaluated, where three neighboring heat pipes were simulated to have failed. This is done by removing the thermal boundary conditions at the heat pipe, essentially setting the derivative of the temperature to zero at that boundary. This does not change the material composition of the heat pipe, which in reality would change as the working fluid changed dynamics. The reason why this test is of interest to this work is because it introduces asymmetry to the thermal profile of the reactor, which in turn causes asymmetric thermal expansion. If the resulting thermal expansion difference is strong enough, it should drive a resulting asymmetric flux response and tilt the power in the system.

The three selected heat pipes can be seen in [Figure 12](#) centered where the temperature increase is highest. The loss of cooling results in a maximum temperature increase of 100 K. As mentioned in [Section 2](#), the fast spectrum reactor is less sensitive to changes in cross-section, so the dominant downstream effect is from increased thermal expansion. In this case that increased displacement is relatively minor, as seen in [Figure 13](#). With a maximum increased displacement of 0.03 cm (compared to the non-failure expansion case), all neutronic and thermal downstream effects are essentially the same as reported for the non-heat pipe failure case. The thermal expansion of the fuel increasing from room temperature to an operating temperature of 1050 K is far greater than the thermal expansion due to a 100 K increase in temperature due to failed heat pipes, so this result makes sense.

5 Conclusion

In summary, the net effect of thermal expansion is largely global. Significant eigenvalue and leakage impact was seen, but changes to the spatial distribution of heating were relatively minor. The only spatial shift in power that is of note is an axially downward shift in power in the UMo fuel. The increased leakage from thermal expansion of the fuel results in a loss of energy deposited in the system that would require a roughly 2.4% increase in fission rate to compensate, and global flux spectrum of the reactor is hardened by the increased leakage.

The coupled effect of including thermal expansion in the neutronic-thermal simulation loop finally results in a maximum 3.5 K difference in temperature reported in the fuel. A decrease of 3.5 K is less than would be discernible for most purposes, and would be unlikely to cause any alarm or secondary effects. This implies that for typical simulation of the core, the non-expanded geometry flux and temperatures are a good representation of the system. The impact on reactivity can be evaluated separately and included after the fact.

The effect of failure of heat pipes studied in Section 4.1 show that although heat pipe failure does introduce some tilt to the expansion of the system, the magnitude is too low to drive a noticeable effect on neutronic-thermal results. This result can be rationalized by recognizing that while the reported increase of 100 K is significant, it is small relative to the expansion of the fuel going from room temperature to operating temperatures at roughly 1050 K.

This work can be further improved in multiple areas. Firstly, the heat transfer in small air gaps needs to be properly accounted for in order to more accurately solve the heat transfer of the problem. That in turn makes the thermal expansion more accurate, and would have a minor impact on reducing particle leakage. Secondly, the mechanics solved on the mesh assume that all connected surfaces are “bound” together, when in reality they are often merely in contact with each other. The fuel and BeO reflector is a good example. When the fuel swells radially outwards, the BeO should not move with the fuel, but only resist movement normal to its surface. Finally, implementing heat pipe models to represent the heat removal in the system may change some of the results, particularly for the heat pipe failure case. Using Dirichlet boundary conditions with the heat pipe failures is like having perfect heat pipes that remove extra heat without increasing in temperature.

In a more general sense, this study shows good promise for the use of Monte Carlo simulation with deformed unstructured mesh geometries. While the coupling scheme described in Section 3.1 has limitations and is far from being user-friendly, the use of DAGMC with OpenMC is becoming more and more streamlined. Future problems that warrant solving neutronics in a highly complicated mesh environment will likely be solvable with the use of DAG-OpenMC.

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

References

Adhikary, D. P., Jayasundara, C., Podgornay, R. K., and Wilkins, A. H. (2016). A robust return-map algorithm for general multisurface plasticity. *Int. J. Numer. Methods Eng.* 109, 218–234. doi:10.1002/nme.5284

Aldebie, F., Fernandez-Cosials, K., and Hassan, Y. (2024). Thermal-mechanical safety analysis of heat pipe micro reactor. *Nucl. Eng. Des.* 420, 113003. doi:10.1016/j.nucengdes.2024.113003

Andrews, N. C. (2012). *Development of fission gas swelling and release models for metallic nuclear fuels*. Ph.D. thesis, Massachusetts Institute of Technology Cambridge, MA.

Blacker, T. D., Bohnhoff, W. J., and Edwards, T. L. (1994). CUBIT mesh generation environment. Volume 1: users manual. *Tech. Rep. Sandia Natl. Lab.(SNL-NM)*.

Bower, A. F. (2009). *Applied mechanics of solids*. CRC Press Boca Roton, FL.

Cao, Y., Miao, Y., Mo, K., Stauff, N., and Lee, C. (2024). “Multiphysics simulations of the krusty criticality experiment using blucrab,” in *Physor 2024: beyond the blueprint: pioneering reactor physics for real-world implementation*.

Chen, C., Mei, H., Wang, Z., Zhang, S., Liu, S., Wang, H., et al. (2023a). Study of the thermal expansion effects of a space nuclear reactor with an integrated honeycomb core design using openmc and ansys. *Ann. Nucl. Energy* 191, 109901. doi:10.1016/j.anucene.2023.109901

Chen, H., Wang, W., Wu, A., Zhang, K., Pan, R., Duan, W., et al. (2023b). Multi-physics coupling analysis of test heat pipe reactor krusty based on moose framework. *Nucl. Eng. Des.* 414, 112597. doi:10.1016/j.nucengdes.2023.112597

Coreform (2024). Neutronics on exact CAD geometry: advances in the Coreform Cubit DAGMC workflow. Available online at: <https://coreform.com/news/neutronics-on-exact-cad-geometry-advances-in-the-coreform-cubit-dagmc-workflow/> (Accessed July 1, 2024).

David, I., Poston, P. R. M., Gibson, M. A., and Sanchez, R. G. (2020). Results of the krusty warm critical experiments. *Nucl. Technol.* 206, S78–S88. doi:10.1080/00295450.2020.1727287

Davis, A., Shriwise, P., and Zhang, X. (2020). Dag-openmc: cad-based geometry in openmc. 395–398. doi:10.13182/T122-32104

Drezewiecki, T., Schmidt, J., Van Wert, C., and Clifford, P. (2021). “Fuel qualification for advanced reactors NUREG-2246,”. U.S. Nuclear Regulator Commission Rockville, MD.

Edenius, M. (1976). Studies of the reactivity temperature coefficient in light water reactors. *Tech. Rep. Chalmers Tek. Hoegskola*.

Gibson, M. A., Poston, D. I., McClure, P., Godfroy, T., Sanzi, J., and Briggs, M. H. (2018). “The kilopower reactor using stirling technology (krusty) nuclear ground test

Author contributions

WK: Writing – original draft, Writing – review and editing. BF: Conceptualization, Funding acquisition, Supervision, Writing – review and editing.

Funding

The author(s) declare that no financial support was received for the research and/or publication of this article.

Acknowledgments

Thanks to Guillaume Giudicelli for his aid with MOOSE and to Patrick Shriwise for his development and stewardship of DAGMC.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

results and lessons learned," in *2018 international energy conversion engineering conference*, 4973.

Giudicelli, G., Lindsay, A., Harbour, L., Icenhour, C., Li, M., Hansel, J. E., et al. (2024). 3.0 - MOOSE: enabling massively parallel multiphysics simulations. *SoftwareX* 26, 101690. doi:10.1016/j.softx.2024.101690

Giudicelli, G. L., Abou-Jaoude, A., Novak, A. J., Abdelhameed, A., Balestra, P., Charlot, L., et al. (2023). The virtual test bed (vtb) repository: a library of reference reactor models using neams tools. *Nucl. Sci. Eng.* 197, 2217–2233. doi:10.1080/00295639.2022.2142440

Hu, G., Hu, R., Kelly, J., and Ortensi, J. (2019). Multi-physics simulations of heat pipe micro reactor. *Tech. Rep. Argonne Natl. Lab.(ANL)*. doi:10.2172/1569948

Hubbard, M., Baines, M. J., and Jimack, P. (2009). Consistent dirichlet boundary conditions for numerical solution of moving boundary problems. *Appl. Numer. Math.* 59, 1337–1353. doi:10.1016/j.apnum.2008.08.002

Hughes, T. J., Engel, G., Mazzei, L., and Larson, M. G. (2000). The continuous galerkin method is locally conservative. *J. Comput. Phys.* 163, 467–488. doi:10.1006/jcph.2000.6577

Jamison, L., Stillman, J., Jaluvka, D., Mohamed, W., Kim, Y., and Wilson, E. (2020). Review of the technical basis for properties and fuel performance data used in HEU to LEU conversion analysis for U-10Mo monolithic alloy fuel. *Tech. rep. Argonne*, IL United States: Argonne National Lab.ANL.

Jeong, M. J., Im, J., Lee, S., and Cho, H. K. (2023). Multiphysics analysis of heat pipe cooled microreactor core with adjusted heat sink temperature for thermal stress reduction using openfoam coupled with neutronics and heat pipe code. *Front. Energy Res.* 11. doi:10.3389/fenrg.2023.1213000

Kendrick, W. R. (2024). *Neutronic-thermal simulation of micro reactor designs for the purpose of analyzing the impact of thermal expansion and hydrogen migration in metal hydride moderator*. Ph.D. thesis, Massachusetts Institute of Technology Cambridge, MA.

Lee, C., Jung, Y. S., Park, H., Shemon, E. R., Ortensi, J., Wang, Y., et al. (2021). Griffin software development plan. *Tech. Rep. Ida. Natl. Lab.INL*. doi:10.2172/1845956

LeVeque, R. J. (2002). Finite volume Methods for hyperbolic problems. *Cambridge texts in applied mathematics*. Cambridge University Press.

Lietzke, A. F. (1970). *Simplified analysis of nuclear fuel pin swelling*. National Aeronautics and Space Administration Washington, DC.

Ma, Y., Han, W., Xie, B., Yu, H., Liu, M., He, X., et al. (2021). Coupled neutronic, thermal-mechanical and heat pipe analysis of a heat pipe cooled reactor. *Nucl. Eng. Des.* 384, 111473. doi:10.1016/j.nucengdes.2021.111473

Massih, A. (1988). Irradiation induced swelling of nuclear fuel. *Tech. Rep.*

Oka, Y. (2013). *Temperature effect of reactivity*. Tokyo: Springer Japan, 23–33. doi:10.1007/978-4-431-54195-0_3

Pope, C., and Lum, E. (2020). "Nuclear reactor thermal expansion reactivity effect determination using finite element analysis coupled with Monte Carlo neutron transport analysis," in *Finite element methods and their applications*. Editor M. Baccouch (Rijeka: IntechOpen Rijeka, Croatia). doi:10.5772/intechopen.93762

Poston, D. I., Gibson, M. A., Godfroy, T., and and, P. R. M. (2020a). Krusty reactor design. *Nucl. Technol.* 206, S13–S30. doi:10.1080/00295450.2020.1725382

Poston, D. I., Gibson, M. A., Sanchez, R. G., and and, P. R. M. (2020b). Results of the krusty nuclear system test. *Nucl. Technol.* 206, S89–S117. doi:10.1080/00295450.2020.1730673

Redmond, E. L. (1997). *Multigroup cross section generation via Monte Carlo methods*. Ph.D. thesis, Massachusetts Institute of Technology Cambridge, MA.

Romano, P. (2020). *Heating and energy deposition*. Available online at: https://docs.openmc.org/en/stable/methods/energy_deposition.html.

Romano, P. K., Horelik, N. E., Herman, B. R., Nelson, A. G., Forget, B., and Smith, K. (2015). Openmc: a state-of-the-art Monte Carlo code for research and development. *Ann. Nucl. Energy* 82, 90–97. doi:10.1016/j.anucene.2014.07.048

Rosenthal, J. S. (2000). Parallel computing and Monte Carlo algorithms. *Far east J. Theor. statistics* 4, 207–236.

Sanchez, R., Grove, T., Hayes, D., Goda, J., McKenzie, G., Hutchinson, J., et al. (2020). Kilowatt reactor using stirling technology (krusty) component-critical experiments. *Nucl. Technol.* 206, S56–S67. doi:10.1080/00295450.2020.1722553

Shemon, E., Miao, Y., Kumar, S., Mo, K., Jung, Y. S., Oaks, A., et al. (2023). Moose reactor module: an open-source capability for meshing nuclear reactor geometries. *Nucl. Sci. Eng.* 197, 1656–1680. doi:10.1080/00295639.2022.2149231

Springer, T. H., Tuttle, R. J., Otter, J. M., and Paschall, R. K. (1969). "Doppler and related measurements in a soft fast-reactor spectrum," in *Tech. Rep., north American aviation, inc., canoga park, CA (United States)*. Atomics International Div Conoga Park, CA. doi:10.2172/4784461

Stauff, N., Mo, K., Cao, Y., Thomas, J., Miao, Y., Zou, L., et al. (2021). "Detailed analyses of a triso-fueled microreactor: modeling of a micro-reactor system using neams tools," in *Tech. rep. Argonne*, IL (United States): Argonne National Lab.ANL.

Sterbentz, J. W., Werner, J. E., McKellar, M. G., Hummel, A. J., Kennedy, J. C., Wright, R. N., et al. (2017). Special purpose nuclear reactor (5 MW) for reliable power at remote sites assessment report. *Tech. Rep. Ida. Natl. Lab.(INL)*. doi:10.2172/1410224

Tautges, T. J., Wilson, P. P. H., Kraftcheck, J. A., Smith, B. M., and Henderson, D. L. (2009). "Acceleration techniques for direct use of CAD-based geometries in Monte Carlo radiation transport," in *M&C 2009 saratoga springs NY*.

Walker, E., Skutnik, S., Wieselquist, W., Shaw, A., and Bostelmann, F. (2022). "SCALE modeling of the fast spectrum heat pipe reactor," in *Tech. rep. Oak Ridge, TN United States: Oak Ridge National Lab. ORNL*.

Xiao, W., Li, X., Li, P., Zhang, T., and Liu, X. (2022). High-fidelity multi-physics coupling study on advanced heat pipe reactor. *Comput. Phys. Commun.* 270, 108152. doi:10.1016/j.cpc.2021.108152

Yan, B., Wang, C., and Li, L. (2020). The technology of micro heat pipe cooled reactor: a review. *Ann. Nucl. Energy* 135, 106948. doi:10.1016/j.anucene.2019.106948



OPEN ACCESS

EDITED BY

Mark D. DeHart,
Idaho National Laboratory (DOE), United States

REVIEWED BY

Gert Van den Eynde,
Belgian Nuclear Research Centre (SCK CEN),
Belgium
Olin Calvin,
Idaho National Laboratory (DOE), United States

*CORRESPONDENCE

Emory Colvin,
✉ colvinem@oregonstate.edu

RECEIVED 30 November 2024

ACCEPTED 06 May 2025

PUBLISHED 04 June 2025

CITATION

Colvin E and Palmer TS (2025) Improved high-fidelity multiphysics modeling of pulsed operation of the annular core research reactor. *Front. Nucl. Eng.* 4:1537136.
doi: 10.3389/fnuen.2025.1537136

COPYRIGHT

© 2025 Colvin and Palmer. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Improved high-fidelity multiphysics modeling of pulsed operation of the annular core research reactor

Emory Colvin* and Todd S. Palmer

School of Nuclear Science and Engineering, College of Engineering, Oregon State University, Corvallis, OR, United States

Sandia National Laboratories' Annular Core Research Reactor (ACRR) is a unique research reactor, using UO_2 -BeO fuel and operating primarily in pulsed mode. To better understand the physical characteristics of the fuel, the distribution of heat generation must be understood. Previous work developed a Serpent two model of the ACRR and Python coupling script to provide multiphysics feedback. Simulations of \$1.50 and \$2.00 pulses were compared to experimental results. This paper expands this work to \$2.50 and \$3.00 pulses. It further explores potential improvements to the model: dividing the fuel into two radial regions for feedback purposes, allowing additional iterations of the multiphysics coupling and checking for convergence, and the development of alternate specific heat capacity values. The use of two radial fuel regions improved agreement with experimental results for the simulations using the original function for specific heat capacity as a function of temperature but did not consistently improve results with the constant value for specific heat capacity. Allowing additional multiphysics iterations until the power distribution field converges also produced little change for reactor power prediction, though it improved maximum fuel temperature prediction slightly. The new values for specific heat capacity provided the most significant improvements to the models. A third-order polynomial developed from experimental data results in a significant improvement in fuel temperature prediction over the constant value with only a small loss of performance in reactor power prediction.

KEYWORDS

serpent 2, Monte Carlo, high-fidelity multiphysics, ACRR, coupled

1 Introduction

Research reactors with a pulse mode are used to provide a high neutron fluence over a short period of time. TRIGA (Training, Research, Isotopes, General Atomics) reactors are a common type of pulsed reactor, but still operate primarily in steady-state mode. The Annular Core Research Reactor (ACRR) is a unique pulsed reactor operated by Sandia National Laboratories. It is operated primarily in pulsed mode, although it is also capable of steady-state operation. The primary ACRR core is fueled with UO_2 -BeO fuel, which has a larger heat capacity than the UZrH fuel used by TRIGA reactors. The subcritical fueled ring external cavity-II (FREC-II) provides a larger experimental cavity within a subcritical lattice of UZrH fuel. FREC-II can be physically tilted away from the primary core to neutronically decouple it, as is assumed for this work (Parma and Gregson, 2019).

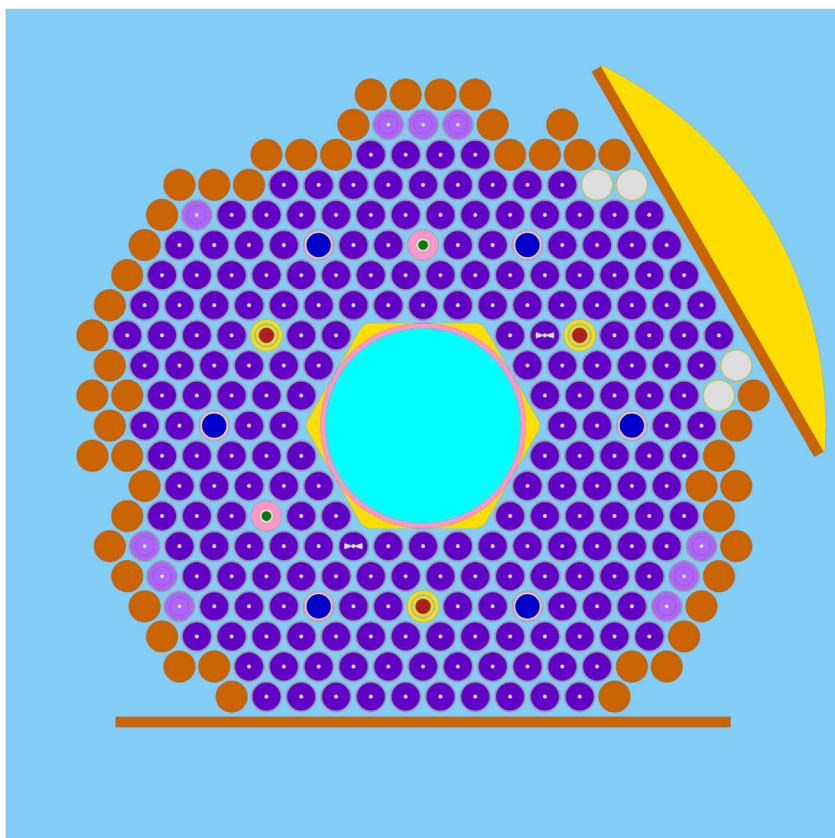


FIGURE 1

ACRR core geometry as plotted by Serpent. The window to the neutron radiography facility is in the upper right and the plate dividing the ACRR from FREC-II is at the bottom. Regular density fuel is dark purple and 90% density fuel is light purple. Fuel elements with a white "bow tie" are instrumented fuel elements. Nickel reflector elements are orange and aluminum elements are white. Safety control rods are green surrounded by pink (stainless steel). Control rods are dark blue surrounded by pink. Transient rods are red surrounded by a double yellow (aluminum) circle.

Past computational reactor physics analyses involved either kinetics simulations with reduced physics approaches to the neutron distribution in space, energy, and angle (Talley and Shen, 2020) or detailed steady-state Monte Carlo simulations focused on radiation fluxes and spectra in irradiation facilities (Parma et al., 2016; DePriest et al., 2006; Parma et al., 2017). To develop a better understanding of the state of the fuel during pulsed operation, a time-dependent Monte Carlo simulation with multiphysics feedback was desired to provide a time- and space-dependent volumetric heat generation rate in the fuel. Such a heat generation rate could then be used as a source term in BISON (Williamson et al., 2021) to evaluate the physical state of the fuel after a lifetime of pulsed operation. Future operation of the ACRR and development of the next-generation of ACRR-like reactors will be informed by the availability of a such a predictive simulation capability and the insight it provides into heat generation in the fuel.

Previous work (Colvin and Palmer, 2025) developed Serpent two models of increasing complexity to verify the use of detector outputs for heat generation in the fuel, and to calculate the resulting fuel temperature increase. This paper provides a brief introduction to the ACRR and the models thereof, then describes the results of the full range of ACRR pulse sizes: \$1.50, \$2.00, \$2.50, and \$3.00 (in the

previous work, only the \$1.50 and \$2.00 pulses were considered). It then presents further approaches to refine the model: separating the fuel nodes into two radial regions, exploring the convergence properties of the multiphysics coupling, and exploring additional methods of determining the specific heat capacity of the fuel. These results are discussed in comparison to data obtained from pulsed operation of the ACRR, specifically reactor power and maximum fuel temperature as measured by the instrumented fuel elements. Finally, suggestions for the use and future improvement of these models are provided.

1.1 Annular core research reactor

The ACRR is a 4 MW research reactor capable of pulses up to ~300 MJ (\$3.50), though pulses above \$3.00 may be smaller than predicted due to the significant negative reactivity feedback in the fuel beginning before the rod can be fully withdrawn such a large distance. Similar to a TRIGA, safe pulsing in the ACRR is possible, as significant reactivity feedback in the fuel occurs due to Doppler broadening and changes in thermal scattering cross sections as fuel temperature increases. Unlike a TRIGA, the fuel for the main ACRR core (shown in Figure 1) is $\text{UO}_2\text{-BeO}$, though the adjacent

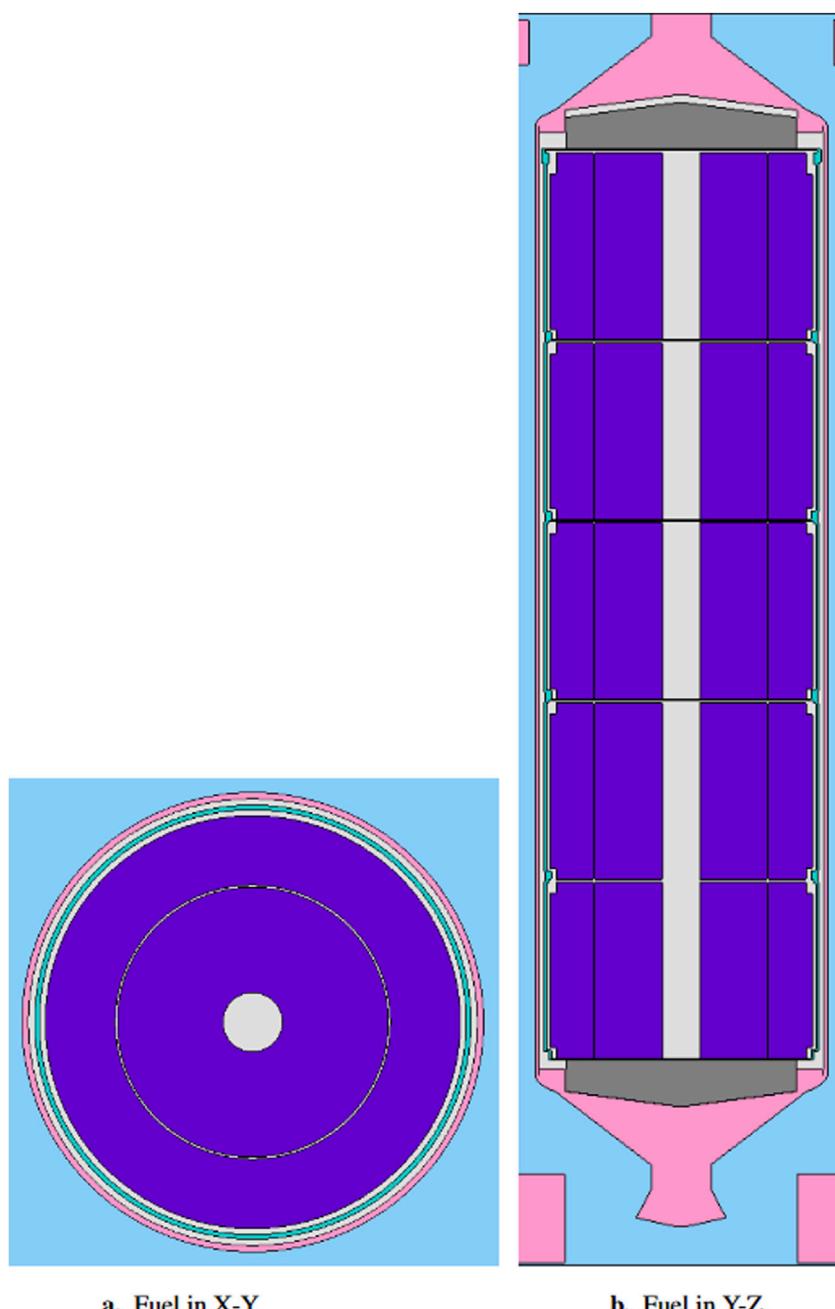


FIGURE 2
ACRR Fuel. **(a)** Cross section of fuel in X-Y showing two annular fuel disks. **(b)** Cross section of fuel in Y-Z showing five niobium cups, upper and lower BeO reflectors, end flutes, and the bottom and top base plates. (Figure not to scale).

subcritical FREC-II facility uses TRIGA fuel (UZrH). To accommodate experimental packages, the ACRR has a 9 inch dry cavity at the center, with a 20 inch dry cavity available at the center of FREC-II for larger experiments. A variety of buckets are available to modify the neutron spectrum in the central cavity, tailoring it to the experimental requirements. The fuel is cooled through natural convection in an open pool. The reactor is controlled with two fuel-followed safety rods, six fuel-followed control rods, and three void-followed transient rods. Because the ACRR is operated primarily in pulsed mode, fuel burnup is minimal (Parma and Gregson, 2019).

2 Materials and methods

2.1 Serpent model

A full model of the ACRR was constructed based on a combination of past MCNP models (DePriest et al., 2006) and fuel element drawings¹. The primary deviation from the MCNP

¹ Fuel Element Assembly Drawing P59525.

UO₂BeO fuel pellets



FIGURE 3
Photograph of fuel and niobium cups showing inner and outer fuel disks (Ames et al., 2022).

model is the replacement of the simplified fuel element models (concentric cylinders bounding each region with no axial or azimuthal definition beyond two planes to bound the element) with fully 3-dimensional models. This core layout uses 213 standard fuel rods and ten fuel rods with 90% fuel density, along with two instrumented fuel elements, two fuel-followed safety rods, six fuel-followed control rods, and three air-followed transient rods. A 25×25 hexagonal lattice (pitch 4.171 cm) is filled with nine different universe types consisting of the previously listed fuel and control rod types in addition to nickel elements, aluminum elements, and a water fill universe. FREC-II is not modeled, as it is assumed to be tilted back and thus neutronically decoupled. Figure 1 includes descriptions of each of these lattice fill universes for a complete map of the layout used in these simulations.

2.1.1 Temperature limitations

To model pulses above \$2.00, nuclear data for higher temperatures must be available. Serpent's on-the-fly temperature

sampling method (Viitanen and Leppänen, 2012; 2014) is able to adjust the cross section data to any needed temperature. However, this is not true for the thermal scattering $S(\alpha, \beta)$ libraries; for this temperature adjustment, a range of libraries at different temperatures is provided and the desired data is calculated by interpolating between two libraries. With no ability to extrapolate beyond the maximum or minimum temperatures provided in these libraries, the fuel temperature is limited to a maximum of 1200 K. We have chosen to artificially limit the temperatures passed to Serpent to allow the simulation to complete. The basis cross section data cannot be temperature adjusted independently of the thermal scattering data, requiring that both be limited to 1200 K.

2.1.2 Fuel element models

Figure 2 shows the fuel modeled in both X-Y and Y-Z cross sections. Each fuel element contains five niobium cups which slide into the clad as shown in Figure 3. The fuel meat consists of inner and outer fuel disks formed by stacks of 0.635 cm height ring halves.

TABLE 1 Material and geometric descriptions of the ACRR model core lattice elements. The 90% density fuel elements are identical except for the reduction of the density to 3.09290 g/cm^3 .

Materials		
Material	ZAID	Atomic Fraction
B_4C Poison	6000	0.200000
	5010	0.159200
	5011	0.640800
Light Water	1001	0.6665667
	1002	0.0001000
	5010	0.0000055
	8016	0.3332063
	8017	0.0001270
$\text{UO}_2\text{-BeO}$ Fuel	4009	0.2827602
	8016	0.5277690
	92235	0.0662957
	92238	0.1222844
	92234	0.0004547
	92236	0.0004358
Element Descriptions		
Region	Material (density, g/cm^3)	Outer Radius (cm)
Safety and Control Rod Radial Regions		
Poison	B_4C (2.48)	0.57150 (safety)
		1.46050 (control)
Void	Helium (1.0245×10^{-3})	0.83185 (safety)
		1.65735 (control)
Sleeve	Stainless Steel 304 (7.8960)	1.74625
Void	Helium (1.2045×10^{-3})	1.82118
Clad	Stainless Steel 304 (7.8960)	1.87198
Transient Rod Radial Regions		
Poison	B_4C (2.48)	0.88138
Wrap	Aluminum 6061-T6 (2.7040)	0.90678
Void	Helium (1.2045×10^{-3})	0.95259
Housing Tube	Aluminum 6061-T6 (2.7040)	1.27000
Gap	Light Water (1.0)	1.42875
Guide Tube	Aluminum 6061-T6 (2.7040)	1.90500

(Continued in next column)

Radial dimensions are provided in Table 1, including the last fuel disk at either end of each niobium cup which is smaller radially to allow a niobium lip at the top of one cup to slide over the bottom

TABLE 1 (Continued) Material and geometric descriptions of the ACRR model core lattice elements. The 90% density fuel elements are identical except for the reduction of the density to 3.09290 g/cm^3 .

Element Descriptions		
Region	Material (density, g/cm^3)	Outer Radius (cm)
Fuel Element Radial Regions		
Central Void	Helium (1.2045×10^{-3})	0.24130
Inner Fuel Disk	$\text{UO}_2\text{-BeO}$ (3.43655)	1.09982
Void	Helium (1.2045×10^{-3})	1.11760
Outer Fuel Disk	$\text{UO}_2\text{-BeO}$ (3.43655)	1.60274 (0.635 cm at ends)
		1.68402 (center 8.89 cm)
Nb Cup	Niobium (8.4)	1.66116 (lower)
		1.73228 (center)
		1.70180 (upper)
Void	Helium (1.2045×10^{-3})	1.82245
Clad	Stainless Steel 304 (7.8960)	1.87325

of the next, holding them together. The total length of fuel in each niobium cup is 10.1600 cm, with a total length of 52.9209 cm including void and niobium sections of each cup. The instrumented fuel elements are identical, but with a pair of Y planes offset by 40° used to create a void wedge in the third (middle) cup of the fuel element (see Figure 1 for the “bow tie” elements).

In the original model, all 100% density fuel is considered a single material and all 90% fuel is considered a single second material, both of which are enriched to 35 weight ^{235}U . By using a hexagonal mesh identical to the core lattice, feedback is performed in each lattice position and considers all fuel in that position to be at the same temperature (e.g., no radial variation). In order to investigate the impact of any radial variation in fuel temperature (Section 3.2), the inner fuel disks are assigned one material and the outer a second material for a total of four types of fuel (100% density inner disk, 100% density outer disk, 90% density inner disk, and 90% density outer disk).

2.1.3 Control rod models

There are three types of control rods in the ACRR: fuel-followed control rods, fuel-followed safety rods, and void-followed transient rods. All safety rods move as a single bank, as do all control rods. The transient rods can be fired individually, as a pair followed by a single rod, or individually. Table 1 includes the radial sections of all three types, which have poison sections 52.2478 cm in length for the safety and control rods, and 76.2000 cm in length for the transient rod. The safety and control rods are followed by a stack of five niobium cups, similar to a regular fuel element. The transient rod is followed by an 0.8001 cm radius helium-filled void with an 0.2524 cm thick aluminum 6061-T6 tube assembly.

2.2 Python coupling

Minor modifications to the Python coupling script used in our previous work (Colvin and Palmer, 2025) were made to allow the larger pulse sizes and the refinements used in this work. A hexagonal mesh with the same pitch as the core lattice is used, with each lattice position divided into ten axial nodes. The fission energy generated within each node is calculated by a Serpent detector applying an ENDF reaction type -8 response function (“macroscopic total fission energy production cross section”), $\kappa\Sigma_f$, where κ is the average energy released per fission, including prompt and delayed gammas, and Σ_f is the macroscopic fission cross section. This is applied to the flux in the fuel as shown in Equation 1,

$$E_i^n = \int_{t_{i-1}}^{t_i} \int_{\mathbf{r}_{n-1/2}}^{\mathbf{r}_{n+1/2}} \kappa\Sigma_f(\mathbf{r})\phi(\mathbf{r}, t) d\mathbf{r} dt, \quad (1)$$

to provide the energy produced in the n th fuel node in the i th time step, where ϕ is the flux, t is time, and \mathbf{r} is position.

As with other pulsed operation analyses (Marcum et al., 2012), this simulation assumes adiabatic conditions due to the extremely short time scale involved. Thus, the temperature in each node at the end of a time step is calculated by

$$T_i = T_{i-1} + \frac{\Delta E_i}{m c_{\text{eff}}(T_{i-1})}, \quad (2)$$

where i is the time step index, T is the temperature of the node at the end of the time step, ΔE represents the fission energy generated in the node, and m is the mass of the node. In order to avoid exceeding the limits of the BeO thermal scattering cross section libraries, an additional conditional statement passes “1200.0” to Serpent for any temperature above 1200 K. A Python function was created for the specific heat capacity, allowing easy switching between values and functions of specific heat capacity by defining a variable at the beginning of the Python coupling file. The impact of fuel density on specific heat capacity is assumed to be negligible; the 90% density fuel uses the same function for $C_p(T)$ in Equation 2 as the 100% density fuel. The difference between the two fuel densities in Equation 2 occurs in the calculation of nodal fuel mass. These fuel elements contribute little to the overall reactor power, as they are exclusively located on the periphery of the core.

2.2.1 Multiphysics coupling iteration and convergence

Previous simulations had defaulted to performing two iterations for each time step without checking for convergence of either power or temperature fields. To determine if convergence was, in fact, being reached or if additional iterations would improve the results, a convergence test was added into the Python script. Wu and Kozlowski (2015) used an approach that was slightly modified into a convergence test for this model. Rather than use a set of axial nodes, the ACRR model uses the set of nodes already in place for the multiphysics feedback. After the second iteration, the relative difference in nodal energy generation between the two iterations is compared to ε , the relative statistical error in k_{eff} as shown in Equation 3:

$$\frac{|E^{(2)} - E^{(1)}|}{E^{(1)}} < \varepsilon, \quad (3)$$

where E is the energy generated in a given node over a time step. In the case that $E^{(1)} = 0$, $E^{(2)}$ is used as the denominator. This primarily occurs in very low power nodes at the ends of the fuel elements. Because Serpent uses delta-tracking for neutron transport, detectors use the collision estimate of neutron flux, small, low power volumes see poor detector efficiency (Leppänen, 2015). If better data is desired for these nodes, the number of particle histories per time step can be increased. This convergence criterion is used only for this single, first check.

Subsequent iterations test for convergence using the relative statistical error in energy generation for each node as calculated by Serpent. The statistical error is calculated and convergence is checked using the criterion

$$E_1^n - 2\sigma_1^n \leq E_2^n \leq E_1^n + 2\sigma_1^n. \quad (4)$$

For a time step to be considered converged, 95% of the nodes must fulfill this criterion ($6250 \times 95\% = 5937.5$, requiring 5938 nodes to meet the conditions of Equation 4).

2.3 Development of alternate specific heat capacity values

Pelfrey (2019) calculated an effective specific heat capacity² of the ACRR fuel using the rule of mixtures as shown in Equation 5:

$$c_{\text{eff}} = V_m c_m + V_u c_u \left(\frac{J}{kg - K} \right), \quad (5)$$

where V_m and c_m are the volume fraction and specific heat capacity of the matrix (BeO) and V_u and c_u are the same for the particulate (UO_2). This work found the effective specific heat capacity to be Equation 6,

$$c_{\text{eff}} = 1.11 \times 10^{-15} T^6 + 8.84 \times 10^{-12} T^5 - 2.75 \times 10^{-8} T^4 + 4.31 \times 10^{-5} T^3 - 0.036 T^2 + 16.78 T - 1706 \left(\frac{J}{kg - K} \right), \quad (6)$$

where T is the temperature in Kelvin. In Colvin and Palmer (2025), this was used along with a constant specific heat capacity derived from a heat capacity determined by the ACRR operations team and the total mass of fuel in the ACRR as shown in Equation 7:

$$c_{\text{eff}} = \frac{C}{M} = \frac{0.325 \times 10^6 J/K}{352.6008 kg} = 921.7222 \left(\frac{J}{kg - K} \right). \quad (7)$$

Although there is insufficient information to replicate the derivation of Equation 6, this function for specific heat capacity was considered based on interest from Sandia National Laboratories, the sponsoring organization.

A new function for specific heat capacity was developed beginning from the same rule of mixtures (Equation 5) as Pelfrey (2019): Returning to the sources used by Pelfrey (2019), the specific heat capacity of BeO (International Atomic Energy Agency, 2008) from $298 \leq T \leq 1200$ K is shown in Equation 8:

² The authors have used “heat capacity” (C) and “specific heat capacity” (c) in this work for clarity, regardless of the language used in the sources.

$$c_p = 1.455 + 0.606 \times 10^{-3}T - 0.544 \times 10^5 T^{-2} \left(\frac{kJ}{kg - K} \right) \quad (8)$$

and then from $1200 \leq T \leq 2820$ K, the specific heat capacity is shown in Equation 9:

$$c_p = 1.791 + 0.201 \times 10^{-3}T \left(\frac{kJ}{kg - K} \right). \quad (9)$$

The specific heat capacity of UO_2 (Fink, 2000) from $298.15 \leq T \leq 3120$ K is shown in Equation 10:

$$c_p = 0.1940 + 3.270 \times 10^{-4}T - 3.132 \times 10^{-7}T^2 + 1.173 \times 10^{-10}T^3 - 9.791 \times 10^{-15}T^4 - 2.654 \times 10^3 T^{-2} \left(\frac{J}{kg - K} \right). \quad (10)$$

Using a particulate volume fraction $V_u = 0.069$ and matrix volume fraction $V_m = 0.931$ (Pelfrey, 2019), the UO_2 -BeO fuel has an effective specific heat capacity given by Equation 11:

$$c_{\text{eff}} = 1.35 \times 10^3 + 5.64 \times 10^{-1}T - 2.16 \times 10^{-8}T^2 + 8.09 \times 10^{-12}T^3 - 6.76 \times 10^{-16}T^4 - 5.06 \times 10^7 T^{-2} \left(\frac{J}{kg - K} \right) \quad (11)$$

from $298.15 \leq T \leq 1200$ K and Equation 12,

$$c_{\text{eff}} = 1.60 \times 10^3 + 2.26 \times 10^{-5}T - 2.16 \times 10^{-8}T^2 + 8.09 \times 10^{-12}T^3 - 6.76 \times 10^{-16}T^4 - 1.83 \times 10^2 T^{-2} + 1.87 \times 10^{-1}T^{-3} \left(\frac{J}{kg - K} \right). \quad (12)$$

from $1200 \leq T \leq 2820$ K. Considering the coefficients of each term in Equations 11, 12 and the temperature range in the order of 10^2 to 10^3 expected in these pulses, these can be further simplified by dropping terms that ultimately contribute a quantity more than four orders of magnitude less than the first term of 1.35×10^3 , resulting in Equation 13,

$$c_{\text{eff}} = 1.35 \times 10^3 + 5.64 \times 10^{-1}T - 5.06 \times 10^7 T^{-2} \left(\frac{J}{kg - K} \right) \quad (13)$$

from $298.15 \leq T \leq 1200$ K and Equation 14,

$$c_{\text{eff}} = 1.60 \times 10^3 \left(\frac{J}{kg - K} \right) \quad (14)$$

from $1200 \leq T \leq 2820$ K.

The pulse simulations in this work are compared to a set of four pulses performed in early 2023 (see Table 2). Power traces, total energy released during the pulse, and maximum fuel temperature were made available by the operations team for the ACRR. The reactor power used in this work was the average of four individual power channels and it was integrated over the length of the pulse for the total energy released. Maximum fuel temperature was taken from the instrumented fuel elements. Unlike power, there is no measurement of fuel temperature over time as the thermocouples cannot react within the relevant time during a pulse. Given the constraints of this available information, the consideration of alternate specific heat capacity functions presented here is intended to provide a rough investigation of the potential for improvement with different functions. In the future, closely working

TABLE 2 Experimental pulse data used for calculation of fuel specific heat capacity.

Pulse size	Maximum Temp (C)	Energy released (MJ)	Date of pulse
\$1.50	283	71.877	Jan. 5, 2023
\$2.00	503	147.611	Jan. 5, 2023
\$2.50	714	223.334	Jan. 4, 2023
\$3.00	924	306.835	Jan. 3, 2023

with the operations team for additional pulse data and/or performing a targeted set of pulses to establish a larger number of data points can provide the necessary data to perform a more sophisticated analysis of the specific heat capacity of BeO- UO_2 fuel.

These sets of energy released and maximum fuel temperature were used to create systems of linear equations which could be solved for expressions for specific heat capacity. Considering Equations 8–10 for temperatures of order 10^2 to 10^3 , the terms within three orders of magnitude of the maximum are T^0 , T^1 , T^2 , T^3 , and T^{-2} . Having only four datasets, two expressions were considered: a third order polynomial and a rational expression of the form $C(T) = a_0 + a_1 T + a_2 T^2 + a_3 T^{-2}$.

This evaluation began from the energy balance in Equation 15:

$$C(T) \frac{\partial T}{\partial t} = P(t), \quad (15)$$

where $C(T)$ is the heat capacity of the fuel as a function of temperature T and $P(t)$ is total reactor power as a function of time t . While $P(t)$ is directly recorded during a pulse, the only recorded temperature is the maximum fuel temperature as measured by the instrumented fuel elements. An integration over the pulse, shown in Equation 16, eliminated the time dependence:

$$\int_{t=0}^{t=t_f} C(T) \frac{\partial T}{\partial t} dt = \int_{t=0}^{t=t_f} P(t) dt. \quad (16)$$

Assuming $C(T)$ to take the form of Equation 17,

$$C(T) = \sum_{n=0}^N a_n T^n, \quad (17)$$

Equation 16 becomes Equation 18,

$$\int_{t=0}^{t=t_f} \sum_{n=0}^N a_n T^n \frac{\partial T}{\partial t} dt = E_{\text{tot}} = \int_{t=0}^{t=t_f} P(t) dt. \quad (18)$$

Further simplifying Equation 18 for the case of the polynomial expression is shown in Equation 19,

$$\begin{aligned} \int_{t=0}^{t=t_f} \sum_{n=0}^N a_n T^n \frac{\partial T}{\partial t} dt &= \int_{t=0}^{t=t_f} \frac{\partial}{\partial t} \left[\sum_{n=0}^N a_n \frac{T^{n+1}}{n+1} \right] dt \\ &= \sum_{n=0}^N \frac{a_n T_f^{n+1}}{n+1} - \sum_{n=0}^N \frac{a_n T_0^{n+1}}{n+1}. \end{aligned} \quad (19)$$

Finally evaluating this at the bounding times results in Equation 20,

$$\sum_{n=0}^N a_n \left[\frac{T_f^{n+1} - T_0^{n+1}}{n+1} \right] = E_{tot} \quad (20)$$

With results from four pulses, a third-order polynomial expression for $C(T)$ can be calculated using the system of equations in [Equation 21](#):

$$\begin{bmatrix} T_{f,1} - T_i & T_{f,1}^2 - T_i^2 & T_{f,1}^3 - T_i^3 & T_{f,1}^4 - T_i^4 \\ 1 & 2 & 3 & 4 \\ T_{f,2} - T_i & T_{f,2}^2 - T_i^2 & T_{f,2}^3 - T_i^3 & T_{f,2}^4 - T_i^4 \\ 1 & 2 & 3 & 4 \\ T_{f,3} - T_i & T_{f,3}^2 - T_i^2 & T_{f,3}^3 - T_i^3 & T_{f,3}^4 - T_i^4 \\ 1 & 2 & 3 & 4 \\ T_{f,4} - T_i & T_{f,4}^2 - T_i^2 & T_{f,4}^3 - T_i^3 & T_{f,4}^4 - T_i^4 \\ 1 & 2 & 3 & 4 \end{bmatrix} \begin{bmatrix} a_0 \\ a_1 \\ a_2 \\ a_3 \end{bmatrix} = \begin{bmatrix} E_1 \\ E_2 \\ E_3 \\ E_4 \end{bmatrix}. \quad (21)$$

Assuming a starting temperature of 298.15 K and using the experimental data provided in [Table 2](#) then solving and dividing by the total mass of fuel results in the polynomial expression of [Equation 22](#),

$$c_{eff}(T) = 3.04 \times 10^{-6}T^3 - 7.24 \times 10^{-3}T^2 + 5.87T - 614 \frac{J}{kg \cdot K}. \quad (22)$$

For the rational case, [Equation 19](#) cannot be further generalized and must be evaluated for each n independently. For $0 \leq n < 3$, this does not change from the polynomial case. For $n = 3$, the final term is determined with [Equation 23](#):

$$\int_{t=0}^{t=t_f} a_3 T^{-2} \frac{\partial T}{\partial t} dt = a_3 \left[-\frac{1}{T_f} + \frac{1}{T_i} \right], \quad (23)$$

which then replaces the a_3 terms in the coefficient matrix of [Equation 21](#). This results in a final expression for specific heat capacity as shown in [Equation 24](#):

$$c_{eff}(T) = 1.45 \times 10^{-3}T^2 - 2.59T + 2.32 \times 10^3 - 1.15 \times 10^8T^{-2}. \quad (24)$$

All of the values and functions for the specific heat capacity of UO_2 -BeO fuel are plotted in [Figure 4](#), along with the specific heat capacities of UO_2 and BeO individually.

3 Results and discussion

3.1 Single radial region fuel model

Full experimental results for four pulse sizes were available, all performed in 2023. Serpent simulations of the \$1.50 and \$2.00 pulses were presented previously ([Colvin and Palmer, 2025](#)); [Table 3](#) shows the results of all four pulse sizes compared to the experimental results. [Figures 5, 8](#) graphically show the results of a \$3.00 pulse, which are broadly representative of all pulse sizes. Comparing the Serpent simulation to the experimental results using the percent absolute relative difference in [Equation 25](#),

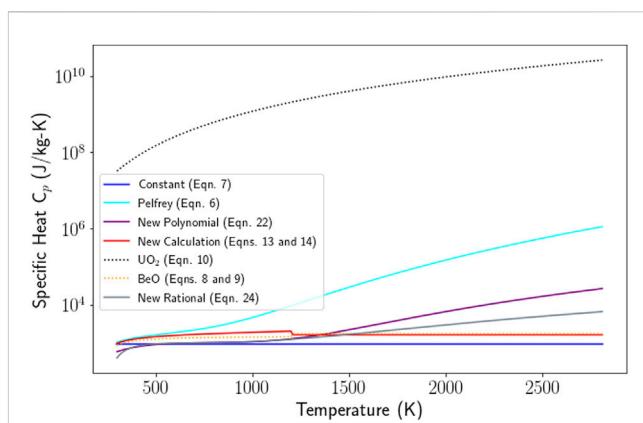


FIGURE 4
Plot of UO_2 -BeO fuel specific heat capacities used for modeling and the specific heat capacities of the individual materials UO_2 and BeO.

$$\% \text{ Abs Rel Diff} = \frac{|P_{exp} - P_{data}|(100)}{P_{exp}}, \quad (25)$$

the “1 Region” columns in [Table 3](#) and the “original fuel” cases in [Figure 6](#) show a closer fit for the constant value for specific heat capacity from [Equation 7](#) over the results from [Pelfrey \(2019\)](#) from [Equation 6](#). All cases have a large difference at the beginning of the pulse, where low power results in noisy experimental data and the instant reactivity insertion assumption results in a small deviation from what would be expected with a realistic rod withdrawal time. To minimize the impact of this early extreme behavior, the first 0.02 s of the simulation (0.01 s of steady state operation and the first 0.01 s of the pulse) have been omitted from the statistics and the median percent absolute relative difference is provided in [Table 3](#) in addition to the mean values.

The power decrease following the peak and the eventual steady state power show the most significant differences between the two specific heat capacity models. During the time increment where power is falling after the peak and the stable power after the pulse, the constant value of specific heat capacity ([Equation 7](#)) results in a much closer match to experimental reactor power, at most 30% relative difference (\$1.50) and at least 0.57% (\$2.50). The simulation using Pelfrey’s function for specific heat capacity ([Equation 6](#)) has a more gradual power decrease and stabilizes at a power level significantly above the experimental results, with the closest match (\$2.00) having a 90% relative difference and the worst match at a 190% relative difference (\$3.00). However, the maximum fuel temperatures reached by each simulation show the opposite relation to the experimental values (as measured by the instrumented fuel elements), with the values from the simulation using Pelfrey’s function for specific heat capacity being much closer to the experimental results. The best match using the constant value for specific heat capacity (Eqn. was the \$1.50 pulse, with a relative difference of 5.9%, and the worst was the \$3.00 pulse, with a relative difference of 38%). Using Pelfrey’s function for specific heat capacity, the closest match to the experimental fuel temperature is a relative difference of 1.6% (\$2.50), while the worst was only 3.5% (\$3.00). The higher fuel temperature of the simulation using the constant value for specific heat capacity ([Equation 7](#)) explains the steeper drop in temperature after the pulse, as the beryllium oxide in the fuel provides larger amounts of feedback. Although a simulated

TABLE 3 Serpent result characteristics for a range of pulse sizes (\$1.50, \$2.00, \$2.50, and \$3.00) compared to experimental results. Includes simulations using the constant value for specific heat capacity and Pelfrey's function for specific heat capacity as a function of temperature, and one and two radial fuel regions. All simulations used a 5×10^{-4} sec time step. Percent absolute relative difference between experimental and simulated total reactor power considered from the time interval between 0.02 and 0.11 s to reduce the impact of noise in the experimental data.

Characteristic	Experimental	Constant C_p		Pelfrey's C_p	
		1 Region	2 Region	1 Region	2 Region
\$1.50 Pulse (Final Rod Height: 43.15)					
Max Total Rx Power (W)	2.15×10^9	1.78×10^9	1.66×10^9	2.14×10^9	2.17×10^9
Mean % Abs Rel Diff		17.7	25.7	50.3	45.5
Median % Abs Rel Diff		11.4	11.4	45.3	45.3
Maximum % Abs Rel Diff		132	132	154	154
Time to Max Power (s)	0.0502	0.0550	0.0530	0.0585	0.0570
FWHM (sec)	0.0232	0.0245	0.0245	0.0265	0.0260
Steady State Avg (W)	4.55×10^7	4.71×10^7	4.43×10^7	7.69×10^7	7.18×10^7
Final Temp (K)	556	592	586 (i)	554	556 (i)
			590 (o)		559 (o)
\$2.00 Pulse (Final Rod Height: 46.75)					
Max Total Rx Power (W)	8.74×10^9	7.45×10^9	6.42×10^9	9.77×10^9	8.79×10^9
Mean % Abs Rel Diff		8.31	15.2	106	109
Median % Abs Rel Diff		3.93	3.93	89.2	89.2
Maximum % Abs Rel Diff		29.9	29.9	312	312
Time to Max Power (s)	0.0326	0.0340	0.0340	0.0345	0.0365
FWHM (sec)	0.0123	0.0120	0.0130	0.0145	0.0145
Steady State Avg (W)	4.86×10^7	4.83×10^7	4.63×10^7	9.23×10^7	8.39×10^7
Final Temp (K)	776	903	867 (i)	789	778 (i)
			874 (o)		789 (o)
\$2.50 Pulse (Final Rod Height: 50.44)					
Max Total Rx Power (W)	2.00×10^{10}	1.52×10^{10}	1.45×10^{10}	2.27×10^{10}	2.14×10^{10}
Mean % Abs Rel Diff		14.2	20.9	134	119
Median % Abs Rel Diff		9.47	9.47	120	120
Maximum % Abs Rel Diff		79.3	79.3	397	397
Time to Max Power (s)	0.0248	0.0240	0.0240	0.0255	0.0250
FWHM (sec)	0.00844	0.00900	0.00850	0.0100	0.0100
Steady State Avg (W)	5.24×10^7	5.21×10^7	4.91×10^7	1.19×10^8	1.12×10^8
Final Temp (K)	987	1190	1180 (i)	971	964 (i)
			1200 (o)		974 (o)
\$3.00 Pulse (Final Rod Height: 55.00)					
Max Total Rx Power (W)	3.67×10^{10}	2.99×10^{10}	2.76×10^{10}	4.64×10^{10}	4.22×10^{10}
Mean % Abs Rel Diff		18.8	14.6	312	295
Median % Abs Rel Diff		13.6	13.6	185	185
Maximum % Abs Rel Diff		80.9	80.9	1580	1580

(Continued on following page)

TABLE 3 (Continued) Serpent result characteristics for a range of pulse sizes (\$1.50, \$2.00, \$2.50, and \$3.00) compared to experimental results. Includes simulations using the constant value for specific heat capacity and Pelfrey's function for specific heat capacity as a function of temperature, and one and two radial fuel regions. All simulations used a 5×10^{-4} sec time step. Percent absolute relative difference between experimental and simulated total reactor power considered from the time interval between 0.02 and 0.11 s to reduce the impact of noise in the experimental data.

Characteristic	Experimental	Constant C_p		Pelfrey's C_p	
		1 Region	2 Region	1 Region	2 Region
Time to Max Power (s)	0.0190	0.0190	0.0190	0.0210	0.0205
FWHM (sec)	0.00652	0.00650	0.00700	0.00800	0.00800
Steady State Avg (W)	5.85×10^7	6.89×10^7	6.53×10^7	1.71×10^8	1.62×10^8
Final Temp (K)	1197	1650	1620 (i)	1150	1140 (i)
			1660 (o)		1160 (o)

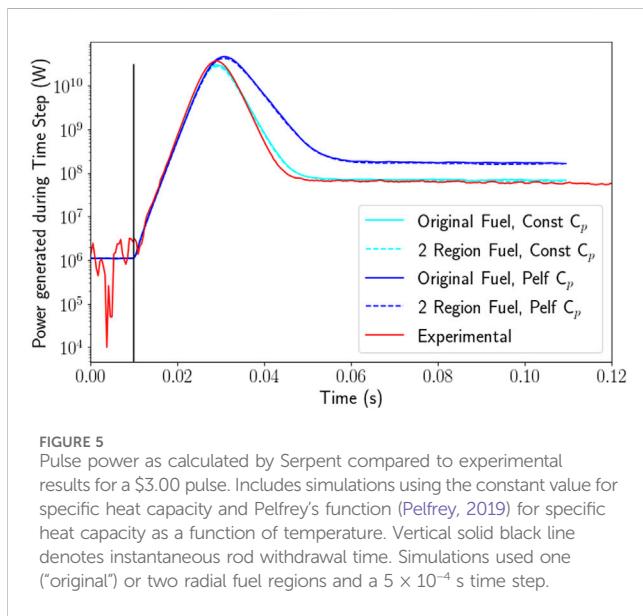


FIGURE 5
Pulse power as calculated by Serpent compared to experimental results for a \$3.00 pulse. Includes simulations using the constant value for specific heat capacity and Pelfrey's function (Pelfrey, 2019) for specific heat capacity as a function of temperature. Vertical solid black line denotes instantaneous rod withdrawal time. Simulations used one ("original") or two radial fuel regions and a 5×10^{-4} s time step.

fuel temperature somewhat above the experimental maximum would be reasonable given that the instrumented fuel elements are limited to a single core position and may not be reading the highest temperature, the constant specific heat capacity simulations showed a temperature significantly above expectations.

The simulations using the constant value for specific heat capacity see small increases in the absolute percent relative difference of total reactor power close to the peak of the pulse but generally even out to be very small as the steady state power is approached after the pulse. The cases using Pelfrey's function for specific heat capacity, however, show different behavior. All show behavior similar to the constant specific heat capacity at the beginning of the pulse, after which the power begins to diverge. As the power decreases, the absolute percent relative difference increases as the simulation power drops more slowly than the experimental before reaching the steady state power after the pulse. The constant specific heat capacity results in a maximum of 140% absolute relative difference (\$1.50), but the median values are all less than 15%. The simulation with Pelfrey's function for specific heat capacity results in a maximum relative difference of 1580% (\$3.00) with the medians falling between 45.0% and 184%.

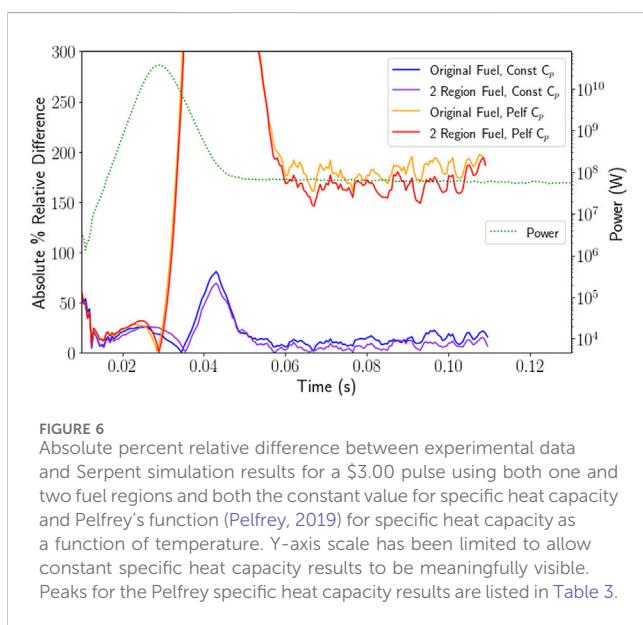


FIGURE 6
Absolute percent relative difference between experimental data and Serpent simulation results for a \$3.00 pulse using both one and two fuel regions and both the constant value for specific heat capacity and Pelfrey's function (Pelfrey, 2019) for specific heat capacity as a function of temperature. Y-axis scale has been limited to allow constant specific heat capacity results to be meaningfully visible. Peaks for the Pelfrey specific heat capacity results are listed in Table 3.

3.2 Two radial regions fuel model

In an effort to improve the accuracy of the Serpent model, the fuel was divided into two different fuel materials: one for the inner disk and one for the outer, as seen in Figure 2. This allows feedback to occur in 20 nodes per fuel rod: 10 axial nodes in the inner fuel disk region and 10 in the outer. Figure 5 provides a visual representation of the power calculated by simulation with a single radial region ("Original fuel") and with two radial regions ("2 Region Fuel") for a \$3.00 pulse. The results and statistics of the full range of pulse sizes is also provided in Table 3. Very little change is seen in the statistics of the percent absolute relative difference between simulated total reactor power and experimental data due to the division of the fuel disks into separate feedback nodes. Table 3 displays the absolute percent relative difference between the various simulation results and experimental values as calculated by Equation 25, with six graphically displaying the results for a \$3.00 pulse. The only measure that significantly changes is the mean of the relative difference; the median and maximum for one or two region fuel remains the same. In the case of the constant specific

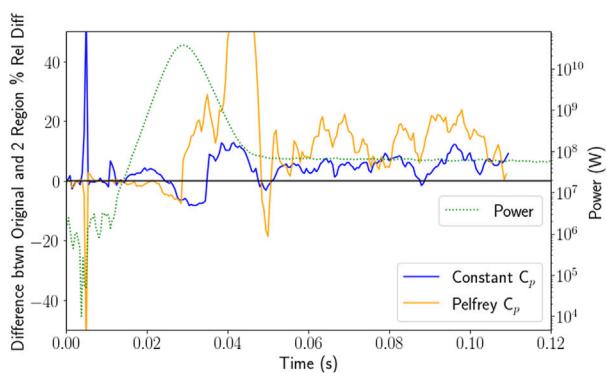


FIGURE 7
Difference between two and one radial fuel region absolute percent relative differences between experimental data and Serpent simulation results for a \$3.00 pulse, using both the constant value for specific heat capacity (Equation 7) and Pelfrey's function for specific heat capacity as a function of temperature (Equation 6). Positive values indicate the two region fuel model is a closer match to experimental data.

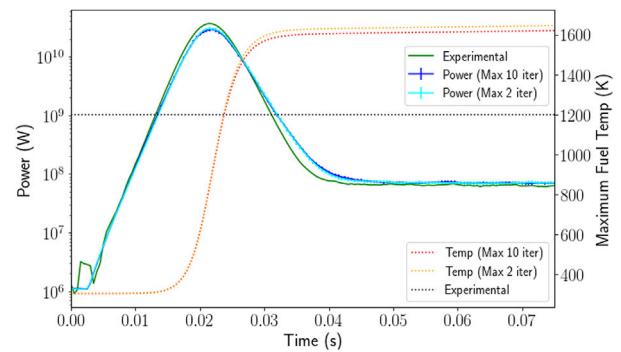


FIGURE 9
Plots of power and temperature for a \$3.00 pulse simulation using the original two iterations per time step compared to a simulation checking for convergence and allowing up to ten iterations. Simulation performed using the constant value for specific heat capacity (Equation 7).

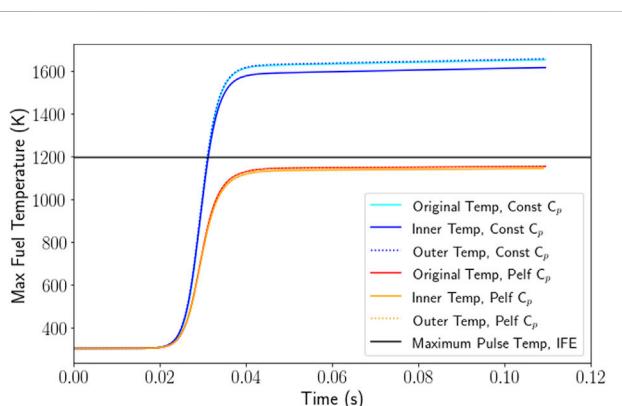


FIGURE 8
Maximum local fuel temperature as calculated by the Serpent coupling script compared to the maximum experimental pulse temperature for a \$3.00 pulse. Includes simulations using the constant value for specific heat capacity and Pelfrey's function (Pelfrey, 2019) for specific heat capacity as a function of temperature. Vertical solid black line denotes instantaneous rod withdrawal time. Simulations used one ("original") or two radial fuel regions and a 5×10^{-4} s time step.

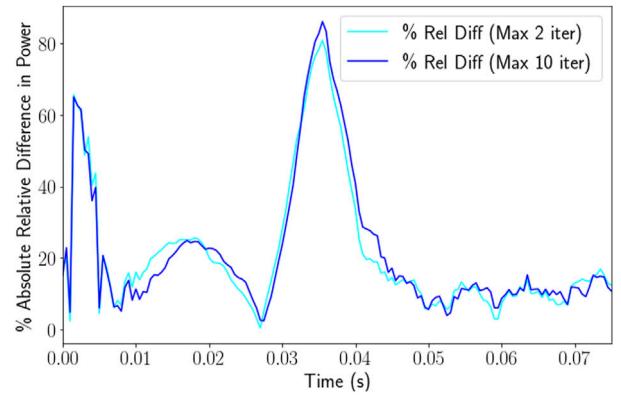


FIGURE 10
Plot of the percent absolute relative difference between Serpent results and experimental results for a \$3.00 pulse. Serpent results are shown for the original two iterations per time step and the case checking for convergence and allowing up to ten iterations. Simulation performed using the constant value for specific heat capacity (Equation 7).

heat capacity, the single region fuel is better in all cases except the \$3.00. Using Pelfrey's function for specific heat capacity as a function of temperature, the two region model is better in all cases except the \$2.00 pulse. In every case, the difference between the mean values is relatively small: less than 10% for the constant specific heat capacity cases and less than 20% in the Pelfrey cases. To better compare the difference between the one and two region fuel cases for a \$3.00 pulse, the difference between the absolute percent relative difference of the two and one region fuel is plotted in Figure 7. While the Pelfrey specific heat capacity simulations generally showed improvement with the two region fuel, especially in the steady state period after the pulse, the constant specific heat capacity results showed very little promise except in the \$3.00 case, with the \$2.00 and \$2.50 simulations only showing a few peaks where the two region fuel model improved the results.

Table 3 includes the maximum fuel temperature which displays a more significant change as well as provides additional information on where the maxima are occurring. Figure 8 provides a graphical representation of the \$3.00 pulse fuel temperature results. Using Pelfrey's function for specific heat capacity, there is less temperature difference between the inner and outer fuel disks than with the constant specific heat capacity. With two fuel regions, both the inner and outer disk temperature were lower than in the single fuel region case. For the two pulse sizes where the experimental maximum temperature was already above the simulated temperature, the change resulting from the additional fuel region is minimal, whereas in the cases where the experimental temperature is below the simulated temperature, the drop is more significant. In the cases with constant specific heat capacity, the impact of the two fuel region model is less consistent. In all cases, the outer fuel disk is at a higher temperature than the inner fuel disk and in some cases may be higher than the single

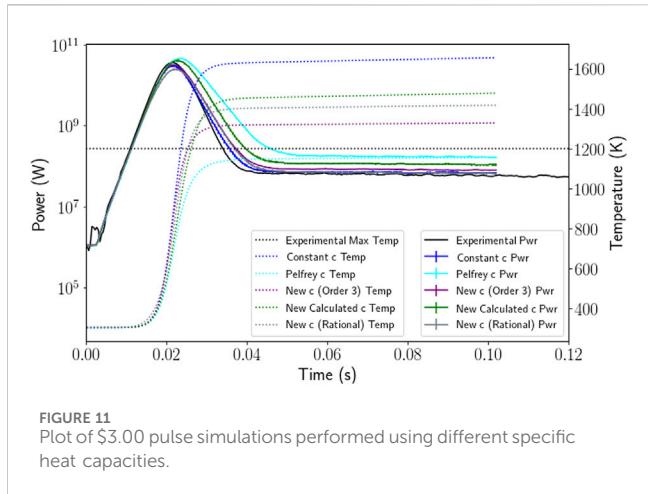


FIGURE 11
Plot of \$3.00 pulse simulations performed using different specific heat capacities.

region. The \$2.00 pulse is an outlier as the only pulse size where there is a significant shift of both the inner and outer fuel temperature.

3.3 Increased iterations for convergence

Except for the initial time step, all time steps converged on the third iteration (the first to use the convergence criterion based on the relative error). As shown in Figure 9 and confirmed in Figure 10, there is very little difference in the total reactor power between the original simulation and the one with the added convergence criterion.

The primary difference lay in the maximum fuel temperature, with the converged case reaching a slightly lower maximum fuel temperature. Both of these cases used the constant specific heat capacity and 10^6 histories divided into 20 batches for each iteration (because the neutron and neutron precursor distributions at the end of one time step are saved for use at the beginning of the next, these are external source simulations rather than criticality simulations).

3.4 Alternate functions for fuel specific heat capacity

The range of specific heat capacity functions shown in Figure 4 were used in a series of \$3.00 pulse simulations, with the results presented in Figure 11 and Table 4. The original constant value for specific heat capacity remains the closest match for the full width-half maximum. The rational function provides a slight improvement to the final steady state average power over the original constant value for specific heat capacity, with a 16% percent absolute relative difference over 19%. The best match for the maximum power is actually the new function calculated from the individual properties of UO₂ and BeO (Equation 13 and 14), though the new third order polynomial calculated from experimental data (Equation 22) results in only slightly worse results, with percent absolute relative differences of 11% and 12%, respectively. The original function from Pelfrey (Equation 6) remains the closest final temperature to the maximum temperature of the pulse, though it is also the only specific heat capacity model to

TABLE 4 Characteristics of \$3.00 pulse using several functions for specific heat capacity compared to experimental results. Further exploration of the new third order function for all pulse sizes.

Case	Max power (W)	FWHM (s)	Steady state		Final Temp (K)
			Avg (W)	(K)	
Experimental	3.67×10^{10}	0.0065	5.85×10^7	1197	
Original Pelfrey	4.64×10^{10}	0.0080	1.73×10^8	1155	
Original Constant	2.99×10^{10}	0.0065	6.96×10^7	1655	
New 3rd Order Fn	3.22×10^{10}	0.0070	8.22×10^7	1328	
New Rational Fn	2.48×10^{10}	0.0075	6.81×10^7	1417	
New Calculated Fn	4.09×10^{10}	0.0085	1.13×10^8	1477	
New 3rd Order Function					
\$1.50 (exper)	1.28×10^9	0.0265	3.84×10^7	567	
	2.15×10^9	0.0232	4.55×10^7	556	
\$2.00 (exper)	5.69×10^9	0.0140	4.96×10^7	833	
	8.74×10^9	0.0123	4.86×10^7	776	
\$2.50 (exper)	1.32×10^{10}	0.0095	5.41×10^7	1090	
	2.00×10^{10}	0.00844	5.24×10^7	987	
\$3.00 (exper)	3.22×10^{10}	0.0070	8.09×10^7	1330	
	3.67×10^{10}	0.00652	5.85×10^7	1197	

underpredict the maximum temperature (by 42 K). The third order polynomial result (Equation 22) is the closest match of the remaining models, overpredicting the maximum temperature by 131 K.

To further assess the third order polynomial, simulations were performed for the other three pulse sizes. The third order polynomial was chosen over the rational function for its better match in maximum power, pulse width, and final temperature, although the rational function more closely matched the steady state average power after the pulse. Results are shown in Table 4. A \$1.50 pulse simulation using the third order polynomial function significantly underpredicts the power, though less pronounced than in larger pulses. Use of the third order polynomial function for specific heat capacity in simulations of all four pulse sizes results in the underprediction of maximum pulse power, though only in the case of the \$1.50 pulse is the steady state final power also underpredicted. Additionally, using the third order polynomial specific heat capacity produces a narrower pulse, as measured by the full width-half maximum value in all pulse sizes. Finally, in all pulse sizes, the maximum fuel temperature remains above the maximum temperature measured experimentally, with smaller differences between predicted and measured temperature at smaller pulse sizes. This is in keeping with the expectation of a “true” maximum fuel temperature above the experimentally measured value.

4 Conclusion

Simulations using our model have now been compared to experimental results for four pulse sizes across the range of pulse sizes performed at the ACRR. As with the \$1.50 and \$2.00 pulse results presented previously (Colvin and Palmer, 2025), the constant value for specific heat capacity (Equation 7) is a closer match to the experimental power results than Pelfrey’s specific heat capacity as a function of temperature (Equation 6). However, Pelfrey’s function outperforms in terms of expected maximum fuel temperature. The division of the fuel into two radial regions did not yield consistent improvement except in the post-pulse steady state power in the simulations using Pelfrey’s function. Nor did the introduction of a convergence criterion improve the power results significantly, with only a slight improvement to the prediction of the maximum fuel temperature. The most significant modifications to the model are achieved through the specific heat capacity, with the third order polynomial (Equation 22) improving the maximum fuel temperature prediction over that of the constant value by over 200 K with a relatively small loss in the power prediction.

There is still the potential for improvement in modeling pulses in the ACRR. The addition of a temperature field convergence check may further improve the results, though an acceptable criterion would need to be developed. Additionally, an increase in the number of particle histories in each iteration would decrease the statistical error and tighten the convergence criterion for each node. The drawback of this would be that both the increased histories per iteration and potential increased number of iterations per time step would lengthen the overall

time for the simulation. With the larger pulse sizes, there remains the need for actual thermal scattering cross section data for BeO above 1200 K. The specific heat capacity of the fuel remains an open question. A larger survey of pulse data may yield a better fit for the specific heat capacity as a function of temperature, or a second materials science analysis combined with pulse results may shed light on the best model for the specific heat capacity. It may also prove useful to transition to a piecewise continuous model for specific heat capacity, as BeO is known to be in this temperature range. The difference in the heat capacity of the 100% fuel density elements and the 90% fuel density elements is also a potential area for model refinement, particularly if more detailed information on the fuel condition in these elements is desired. Finally, there may also be additional improvement using some of these less hopeful techniques in conjunction, e.g., the use of a convergence criterion with the two radial region fuel.

Data availability statement

The datasets presented in this article are not readily available because we do not have permission to release our data. Requests to access the datasets should be directed to colvinem@oregonstate.edu.

Author contributions

EC: Conceptualization, Methodology, Software, Visualization, Writing – original draft, Writing – review and editing. TP: Conceptualization, Funding acquisition, Project administration, Resources, Supervision, Writing – original draft, Writing – review and editing.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. This research was funded through a partnership with Sandia National Laboratories. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

Acknowledgments

The authors would like to thank Jaakko Leppänen, Ville Valtavirta, and the Serpent team at VTT Technical

Research Centre of Finland for their help in troubleshooting. We would also like to thank Keenan Hoffman and Christopher Magone for their help in the MCNP-to-Serpent model conversion.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The handling editor MD shares an affiliation of the authors EC and TP. All parties confirm the absence of any collaboration during review.

References

Ames, D., Cole, J., Cook, M., Raster, A., Harms, G., and Miller, J. (2022). "IER-523: feasibility of experiments focused on measuring the effects of UO_2BeO material on critical configurations using 7uPCX," in 2022 DOE nuclear criticality safety program technical program review. *SAND2022-1412 PE*.

Colvin, E., and Palmer, T. S. (2025). High-fidelity multiphysics modeling of pulsed reactor heat generation in the Annular Core Research Reactor fuel using Serpent 2. *Ann. Nucl. Energy* 211, 110954. doi:10.1016/j.anucene.2024.110954

DePriest, K. R., Cooper, P. J., and Parma, E. J. (2006). *MCNP/MCNPX model of the annular core research reactor*. Albuquerque, NM: Sandia National Laboratories.

Fink, J. K. (2000). Thermophysical properties of uranium dioxide. *J. Nucl. Mater.* 279, 1–18. doi:10.1016/S0022-3115(99)00273-1

International Atomic Energy Agency (2008). *Thermophysical properties of materials for nuclear engineering: a tutorial and collection of data*. Vienna, Austria: International Atomic Energy Agency.

Leppänen, J. (2015). *Serpent - a continuous-energy Monte Carlo reactor physics burnup calculation code user's manual*. Finland: VTT Technical Research Centre of Finland. Available online at: https://serpent.vtt.fi/serpent/download/Serpent_manual.pdf.

Marcum, W. R., Palmer, T. S., Woods, B. G., Keller, S. T., Reese, S. R., and Hartman, M. R. (2012). A comparison of pulsing characteristics of the Oregon State University TRIGA reactor with FLIP and LEU fuel. *Nucl. Sci. Eng.* 171, 150–164. doi:10.13182/nse11-25

Parma, E. J., and Gregson, M. W. (2019). *The annular core research reactor (ACRR) description and capabilities*. Albuquerque, NM: Sandia National Laboratories.

Parma, E. J., Naranjo, G. E., Lippert, L. L., Clovis, R. D., Martin, L. E., Kaiser, K. I., et al. (2017). *Radiation characterization summary: ACRR-FRECCI cavity free-field environment at the core centerline (ACRR-FRECCI-FF-cl)*. Albuquerque, NM: Sandia National Laboratories.

Parma, E. J., Naranjo, G. E., Lippert, L. L., and Vehar, D. W. (2016). Neutron environment characterization of the central cavity in the annular core research reactor. In *EPJ Web Conf.*, vol. 106, 01003, doi:10.1051/epjconf/201610601003

Pelfrey, E. (2019). *A transient thermal and structural analysis of fuel in the annular core research reactor*. Master's thesis, University of New Mexico.

Talley, D. G., and Shen, Y.-L. (2020). RAZORBACK—a reactor transient analysis code for large rapid reactivity additions in a natural circulation research reactor. *Ann. Nucl. Energy* 138, 107153. doi:10.1016/j.anucene.2019.107153

Viitanen, T., and Leppänen, J. (2012). Explicit treatment of thermal motion in continuous-energy Monte Carlo tracking routines. *Nucl. Sci. Eng.* 171, 165–173. doi:10.13182/nse11-36

Viitanen, T., and Leppänen, J. (2014). Target motion sampling temperature treatment technique with elevated basis cross-section temperatures. *Nucl. Sci. Eng.* 177, 77–89. doi:10.13182/nse13-37

Williamson, R. L., Hales, J. D., Novascone, S. R., Pastore, G., Gamble, K. A., Spencer, B. W., et al. (2021). BISON: a flexible code for advanced simulation of the performance of multiple nuclear fuel forms. *Nucl. Technol.* 207, 954–980. doi:10.1080/00295450.2020.1836940

Wu, X., and Kozlowski, T. (2015). Coupling of system thermal-hydraulics and Monte-Carlo code: convergence criteria and quantification of correlation between statistical uncertainty and coupled error. *Ann. Nucl. Energy* 75, 377–387. doi:10.1016/j.anucene.2014.08.016

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.



OPEN ACCESS

EDITED BY

Kirk Atkinson,
Ontario Tech University, Canada

REVIEWED BY

Wenhai Qu,
Shanghai Jiao Tong University, China
Jacques Lechelle,
Commissariat à l'Energie Atomique et aux
Energies Alternatives (CEA), France
Mustafa K Jaradat,
Idaho National Laboratory (DOE), United States

*CORRESPONDENCE

Yonghee Kim,
✉ yongheekim@kaist.ac.kr

RECEIVED 18 March 2025

ACCEPTED 30 June 2025

PUBLISHED 11 July 2025

CITATION

Kim I, Oh T and Kim Y (2025) Coupling of the Monte Carlo iMC and OpenFoam codes for multiphysics calculations of molten salt reactors.

Front. Nucl. Eng. 4:1595628.
doi: 10.3389/fnuen.2025.1595628

COPYRIGHT

© 2025 Kim, Oh and Kim. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Coupling of the Monte Carlo iMC and OpenFoam codes for multiphysics calculations of molten salt reactors

Inyup Kim, Taesuk Oh and Yonghee Kim*

Department of Nuclear and Quantum Engineering, Korea Advanced Institute of Science and Technology, Daejeon, Republic of Korea

This paper presents the development of a multiphysics coupled framework of Monte Carlo neutronics iMC and OpenFOAM Computational Fluid Dynamics codes for molten salt reactor (MSR) analysis. The overall coupling scheme is handled, including the framework structure and iteration scheme. Also, related techniques to enhance the accuracy and efficiency of the coupling are introduced, such as delayed neutron precursor tracking. The framework is applied to a simple molten salt reactor model and achieves a converged solution. In addition, sensitivity studies on the neutronics mesh are performed. The research demonstrates the capability of the iMC-OpenFOAM coupled framework to achieve a converged solution and provides significant insights into the analysis of the MSRs.

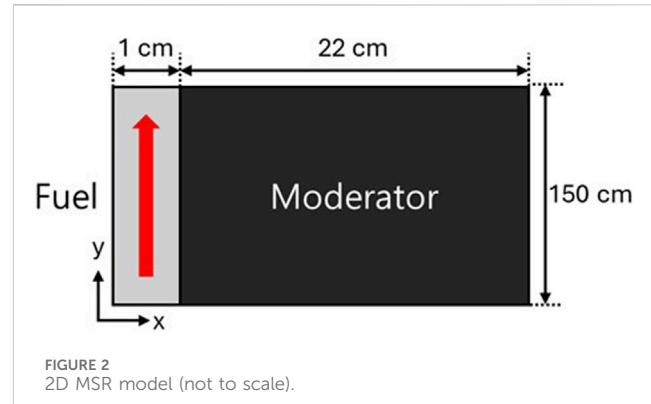
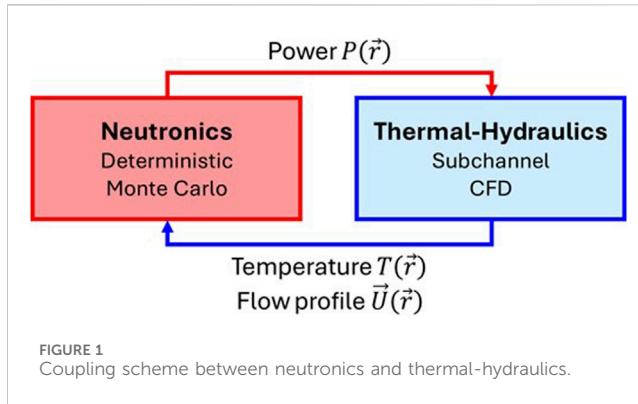
KEYWORDS

Monte Carlo, molten salt reactor, IMC, OpenFOAM, multiphysics

1 Introduction

Molten Salt Reactor (MSR) is an advanced nuclear reactor concept gaining significant attention. MSRs offer distinct advantages over conventional reactor designs. Since the MSR employs liquid fuel, it leads to continuous refueling and removal of noble fission products by utilizing helium bubbling. In addition, the MSR provides inherent safety and long-term operation. Given these advantages, MSRs are being actively developed worldwide, with several experimental and commercial designs progressing toward deployment. Despite the benefits, an accurate analysis of MSR is an ongoing challenge that requires a multiphysics simulation framework. Unlike traditional reactor systems, where the phenomenon can often be treated separately, MSR requires tightly coupled simulations, especially its delayed neutron precursors (DNP) flow and neutron-induced heating. This complexity has led to ongoing efforts in developing MSR-specific multiphysics models, including coupled neutronics, fluid dynamics, and thermal-mechanical analysis. Although multiphysics solvers for MSRs have been proposed, most rely on deterministic neutronics and simplified CFD, limiting their accuracy in capturing key MSR phenomena (Křepel et al., 2007; Aufiero et al., 2014; Ridley et al., 2017; Yang et al., 2022). Regarding the neutronics, deterministic solvers are preferred due to their simplicity and computing efficiency. However, deterministic methods rely on simplifying assumptions and cannot fully capture effects such as secondary photon transport.

To address these challenges, we propose a novel MSR-specialized multiphysics analysis scheme by coupling an in-house neutron transport code, iMC (Kim and Kim, 2021), with



the OpenFOAM CFD framework (Weller et al., 1998). iMC code is a Monte Carlo neutron transport code developed at KAIST, which has previous efforts to develop and implement the MSR-specialized analysis, including DNP shift and nuclide control in depletion (Kim et al., 2024; Kim et al., 2025). The iMC method offers an accurate and flexible approach to solving neutron transport problems. Unlike previous approaches that utilized a deterministic neutronics solver, this work aims to evaluate reactor power distribution based solely on Monte Carlo neutron transport analysis with secondary photon transport. Monte Carlo methods enable high-fidelity modeling of coupled phenomena, making them particularly suitable for MSR analysis. At the same time, OpenFOAM is an open-source and accurate CFD code for simulating complex fluid flow and thermal-hydraulic phenomena. By integrating these two codes, we aim to develop a computational framework capable of capturing the tightly coupled physics of MSRs, including neutron transport, temperature feedback, and fuel salt circulation.

This paper presents the development and implementation of the multiphysics coupling scheme for MSR analysis. Section 2 details the coupling framework by the coupling mechanism, mesh control strategies, and treatment of MSR-specific flow characteristics. Section 3 describes the problem setup and demonstrates the results of the developed scheme applied to benchmark MSR configurations. Finally, Section 4 concludes the study by summarizing key findings and discussing future improvements in MSR multiphysics simulation.

2 Methods

2.1 Coupling framework

The paper discusses a coupled framework of the iMC and OpenFOAM codes to provide a multiphysics solution for the MSR. Figure 1 describes the physics to be considered in the MSR analysis and data transferred between the schemes. The framework comprises three parts: neutronics, CFD, and an intermediate coupler. The neutronics, Monte Carlo neutron transport using the iMC code, has temperature T , density ρ , and flow profile \vec{U} from the CFD code as inputs and tallies heat distribution from the neutron. The CFD code, OpenFOAM, utilizes the heat distribution from the neutronics to compute temperature and velocity distributions in the reactor domain. The results are then

transferred to the neutronics part. The coupler is a script that handles data transfer and iterations. Since the neutronics and fluid dynamics are tightly coupled, a certain level of iterations is required to provide a converged result. The coupler performs a fixed-point iteration by handling the data from both sides and halting the simulation when the result is sufficiently converged.

A contribution of the neutronics to the coupled framework is the production of a heat distribution induced by the neutron. In detail, neutron heating can be classified into direct and indirect heating via secondary photons. Regarding conventional Monte Carlo simulation, fission heating data can be scaled to include secondary photon heating without additional photon transport. On the other hand, the spatial distribution of the heating is required for the coupling scheme. Therefore, secondary photon production and transport simulation are necessary to obtain an accurate spatial distribution of the heating. A photon transport module has been newly implemented in the iMC code and internally verified against OpenMC and Serpent results (Lund and Romano, 2018; Kaltiaisenaho, 2020). The power distribution is generated from the coupled neutron-photon transport, which serves as the source term for CFD simulations.

The OpenFOAM evaluates the temperature distribution $T(\vec{r})$ and flow profile $\vec{U}(\vec{r})$ based on the transferred power distribution. The temperature distributions of the fuel and the moderator are applied to the iMC Monte Carlo analysis, based on the on-the-fly Doppler broadening. Section 2.2 expresses the on-the-fly Doppler broadening's theoretical basis and its usage in the iMC code. The fuel flow profile affects the delayed neutron precursor and perturbs the reactivity and kinetic parameters. Section 2.3 introduces an approach used in the iMC code to simulate the delayed neutron precursor on the fuel flow.

One practical difference between the iMC and OpenFOAM codes is the mesh size. Since OpenFOAM solves fluid dynamics using the finite volume method, a sufficiently small mesh is required for an accurate result. On the other hand, the iMC code provides a heat distribution based on a tally on the mesh. Therefore, there is a mismatch in the mesh size, and the data needs to be collapsed when transferred from the OpenFOAM to the iMC. In the framework, the CFD data is converted to a coarser mesh by utilizing the OpenFOAM post-processing feature, which utilizes volume averaging. On the other hand, the power distribution on the coarser mesh is transferred to the OpenFOAM code and directly applied to the fine CFD mesh without interpolation.

2.2 Temperature treatment

After the CFD analysis from OpenFOAM, the temperature distribution is transferred to the iMC. In the iMC code, the temperature treatment simulates the target velocity. [Equation 1](#) gives the Doppler-broadened cross-section $\sigma(y, T)$ evaluated from a base temperature T_0 ([Cullen and Weisbin, 1976](#)).

$$\sigma(y, T) = \frac{1}{y^2 \sqrt{\pi}} \int_0^{\infty} dx x^2 \sigma(x, T_0) \left[e^{-(x-y)^2} - e^{-(x+y)^2} \right], \quad (1)$$

where x and y are defined as in [Equation 2](#).

$$\alpha = \frac{A}{k(T - T_0)}, x = \sqrt{\alpha E_r}, y = \sqrt{\alpha E}, \quad (2)$$

where E and E_r are incident neutron energy and relative neutron energy concerning target velocity sampled from the target temperature, respectively. A and k are the relative target mass to neutron mass and the Boltzmann coefficient, respectively. In the iMC code, [Equation 1](#) integral is evaluated using a modified Gauss-Hermite Quadrature, which provides sufficient accuracy ([Jo and Cho, 2017](#)). This method requires a single temperature point, suitable for thermo-coupled analysis that varies in temperature range.

Some nuclides require additional thermal scattering data. Unlike the cross-section data, the Doppler broadening is unavailable for the thermal scattering. Therefore, the thermal scattering data linearly interpolates the cross-section at different temperature points.

In the iMC neutron tracking, the particle stops at the material or geometry boundaries. As temperature change leads to the cross-section change, the particle should also stop at the temperature grid boundary. When the particle crosses the temperature grid boundaries, the material temperature and density are updated based on the CFD results. Therefore, the size of the temperature mesh affects the computing time, especially for the fine temperature mesh. When the temperature distribution is obtained from the OpenFOAM, computing time increases significantly. The coupling scheme collapses the temperature distribution into the coarser mesh. The mesh sensitivity regarding accuracy and computing time will be investigated in a subsequent section.

2.3 Delayed neutron tracking

One unique aspect of the MSR analysis is its flowing fuel. As mentioned, the MSR utilizes the liquid fuel as a coolant, achieving thermal balance. Furthermore, the flowing fuel affects the neutronics by shifting the delayed neutron distribution. The shift changes the contribution to the fission chain by simply moving the precursor position. The contribution becomes negligible when the precursor moves out of the active core and decays.

The iMC code has a feature to simulate the delayed neutron precursor shifts on the given fuel flow ([Kim et al., 2024](#)). During the neutron tracking, if the neutron undergoes fission, the resulting fission neutrons are tested to determine whether they are delayed or not based on the delayed neutron fission yield. If the fission neutron is considered delayed, the delayed group is also determined, and the

corresponding emission time, t_{emit} is sampled based on [Equation 3](#), using constant λ and uniform random number γ ,

$$t_{emit} = -\frac{\ln \gamma}{\lambda}. \quad (3)$$

Then, given the emission time and position, its shifted position is sampled. Based on the velocity field provided by the CFD code, the precursors move toward the top of the active core. Since the velocity field is not provided in a functional form, the precursor tracking is performed numerically. [Equation 4](#) shows a basic tracking scheme. From fission site $\vec{r}(0)$, the position of the precursor is numerically obtained based on the velocity field $U(\vec{r}, t)$ and fine time step dt . The current framework defines the velocity field for the finite volume and assumes constant velocity within the volume. Therefore, dt is automatically determined based on the current position and distance to the velocity grid boundary.

$$\vec{r}(t + dt) = \vec{r}(t) + \vec{U}(\vec{r}, t)dt. \quad (4)$$

The tracking is performed until the precursor decays or reaches the top of the active core. If the precursor reaches the top of the active core, the precursor is considered to recirculate and return to the bottom. In the current recirculation scheme in the iMC code, the out-of-core model is simplified: the precursor decayed out of the core is neglected from the analysis, and the precursor recirculated into the core is uniformly distributed. The iMC code set the weight of the delayed neutron produced out of the active core to zero.

The conventional fraction of the delayed neutron is about 0.7%. Since the Monte Carlo analysis requires numerous simulations for better precision, the precision of the delayed neutron-related results is often insufficient due to the low concentration of the delayed neutron. To enhance statistical precision, delayed neutron precursor splitting is introduced. The splitting intentionally splits the delayed neutron precursors from the fission while adjusting the corresponding weight to avoid bias. The delayed neutron precursors are then tracked based on the fuel flow profile.

Prompt or delayed neutrons are determined from the neutron yield when the neutron undergoes fission. Instead of producing a single delayed neutron precursor with unity weight, n_{split} delayed neutron precursors are produced with the weight of $1/n_{split}$. The scheme leads to multiple samplings of the delayed neutron precursor groups and the corresponding precursor shift. Note that the method needs to be unbiased and thus requires some attention regarding its normalization. Recalling a typical Monte Carlo neutron transport procedure, after cycle i ends, the next cycle's neutron weight w_{i+1} is adjusted to [Equation 5](#):

$$w_{i+1} = N/M, \quad (5)$$

where N and M are the initial number of histories per cycle and the total weight of fission neutrons produced from cycle i , respectively. However, the weight of the delayed neutron should be adjusted as in [Equation 6](#), for unbiased evaluation:

$$w_{d,i+1} = \frac{N}{Mn_{split}}. \quad (6)$$

In addition, the delayed neutron leaks from the core should be adequately considered by setting their initial weight to zero.

Therefore, the initial total weight will be slightly smaller than N if some precursors decay out of the core.

The method is ineffective for static fuel since the delayed neutron fraction is solely determined by the delayed neutron yield and fission rate. On the other hand, the technique enhances precision for the moving fuel where the produced delayed neutron precursors are shifted. While sampling without splitting results in a single position moved due to the flow, splitting provides a distribution of the delayed neutron production for varying emission times. In this paper, n_{split} is set to 10 for every fission reaction. Improvement in the precision, accuracy, and computing time will be discussed in [Section 3](#).

Adjoint neutron flux is a measure of the contribution of the phase space to the fission chain, which needs to be considered while producing the effective kinetic parameters. The iMC code was previously developed and implemented using the iterated fission probability (IFP) method, which was introduced to obtain adjoint-weighted factors during the Monte Carlo simulation ([Oh et al., 2025](#)). The IFP method is utilized to get the adjoint-weighted kinetic parameters, with 10 latent cycles.

2.4 Iteration scheme

Due to the nonlinear relation between the neutronics and the fluid dynamics, multiple coupled calculations are required to obtain a converged solution. However, unlike deterministic neutronics code, the results from the Monte Carlo method contain uncertainty due to its statistical nature. The uncertainty complicates the convergence criteria, as the solutions are expected to oscillate within a specific range.

In this study, the heat distribution from the Monte Carlo analysis is applied with an under-relaxation. When the heat distribution H_i is obtained from the i th iteration, actual input for the OpenFOAM code, H'_i can be estimated as in [Equation 7](#):

$$H'_i = w_i H_i + (1 - w_i) H_{i-1}, \quad (7)$$

where the weight w_i can be determined using [Equation 8](#):

$$w_i = \frac{1}{i}. \quad (8)$$

The selection of the w_i is heuristic and intended to stabilize the solution gradually. The convergence is determined by comparing the temperature distribution with the previous results. In this paper, the convergence criteria for the temperature are set to a maximum absolute difference of 0.1 K for both regions.

3 Numerical results

3.1 Problem description

We adopt a newly proposed 2D thermal MSR benchmark developed by [Pfhal \(2025\)](#), which features a 1 cm-wide fuel channel adjacent to a 22 cm graphite moderator slab and extends 1.5 m in height. The total power of the model is considered to be

TABLE 1 Atomic number density of the fuel in atoms/barn.cm

Nuclide	Atomic density (atoms/barn.cm)
Na-23	1.67E-02
K-39	8.71E-03
K-40	1.09E-06
K-41	6.29E-04
U-234	1.75E-06
U-235	1.96E-04
U-238	3.75E-03
U-236	8.98E-07
F-19	4.19E-02

0.35 MW. As this benchmark is recent, no reference solution is available for direct comparison.

[Figure 2](#) shows a geometry of the MSR model. [Table 1](#) denotes an isotope-wise composition of the fuel. The initial temperature of the fuel and the graphite is assumed to be 873.15 K. The moderator is considered pure natural carbon. [Table 2](#) shows a fundamental thermo-physical properties.

In OpenFOAM, the problem is solved on 47,600 meshes. The fuel-side mesh contains 14,200 mesh divided into 142 and 100 meshes horizontally and vertically. Note that the mesh becomes finer as it approaches the fuel-moderator interface. Due to the lower heat transfer in the fuel region, a drastic temperature increase is expected near the interface. The conjugate heat transfer, a heat transfer between fluid and solid, is a complex phenomenon in terms of the CFD code. Therefore, a finer mesh is required near the interface. Similarly, the moderator region contains 33,400 mesh divided into 334 and 100 meshes horizontally and vertically, which becomes finer near the interface.

On the other hand, iMC produces heat distribution for 1,250 meshes for both regions. The mesh size is equivalent in an area since the heat distribution is relatively smoother than the fluid dynamics. Due to the use of different mesh sizes, sensitivity analysis of the mesh collapse will be performed. Mesh conversion is done with a post-processing tool equipped with the OpenFOAM.

For the iMC, reflective boundary conditions are applied horizontally, and vacuum boundary conditions are applied vertically. Regarding the OpenFOAM, slip boundary conditions are applied for the interface between the fuel and the moderator. The top and bottom boundaries of the moderator are set to be adiabatic, while fuel is set to be inlet and outlet. Initial velocity and temperature conditions are set to be 0.2 m/s upward and 873.15 K, respectively.

Instead of explicitly modeling an out-of-core model, the fuel flow is assumed to stay out of the core for 10 s and undergo cooling. In this problem, the cooling is simplified as a temperature transfer to the reference temperature defined as [Equation 9](#):

$$\frac{dT}{dt} = -\gamma_{cool}(T - T_{ref}), \quad (9)$$

where γ_{cool} and T_{ref} are 0.2 s⁻¹ and 873.15 K, respectively.

TABLE 2 Thermo-physical properties of the fuel and moderator.

Properties	Density	Dynamic viscosity	Specific heat	Thermal conductivity
Fuel	5,000-T [K] kg/m ³	0.017 kg/m/s	1,000 J/kg/K	1 W/m/K
Moderator	1,800 kg/m ³	-	1,760 J/kg/K	60 W/m/K

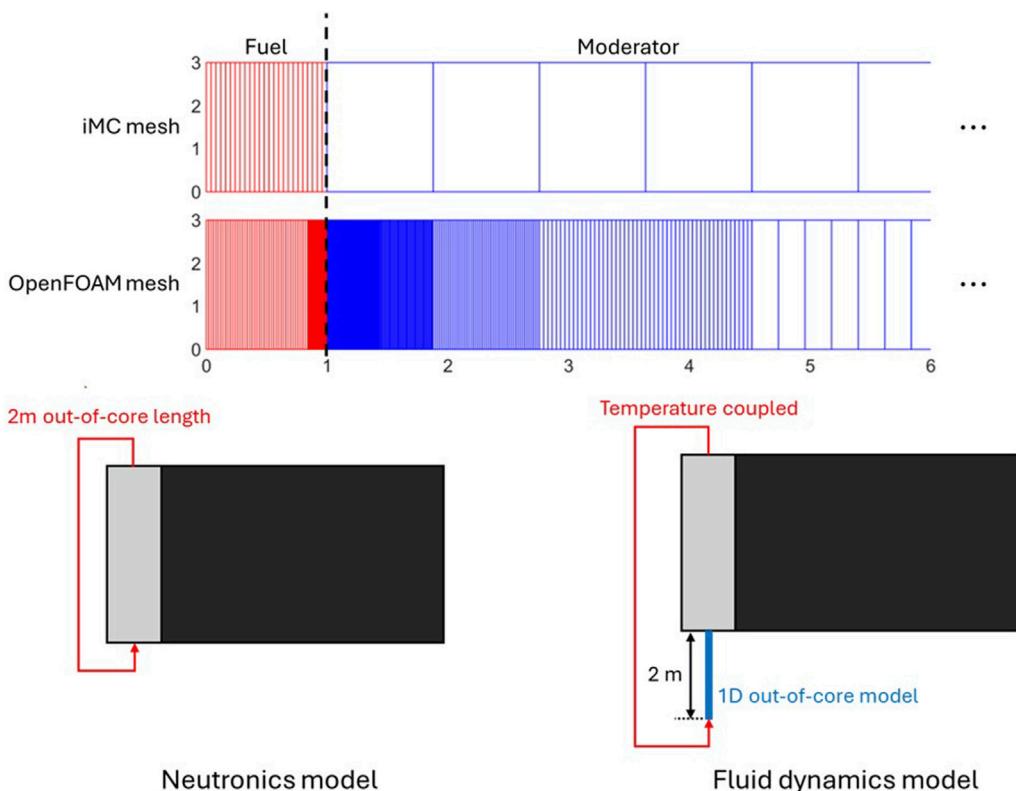


FIGURE 3
Mesh and out-of-core modeling used in the iMC and OpenFOAM codes.

The out-of-core model is not directly modeled in both codes but is simplified differently. Figure 3 describes different approaches to the out-of-core model. Regarding the OpenFOAM, the out-of-core model is simplified to a 1D axial model, subdivided into 200 equivalent meshes. The inlet temperature of the 1D model is set to the average outlet temperature of the active core. In detail, the average outlet temperature is defined as the mass flow rate-weighted average, representing a heat addition in the active core. Outlet temperature is set to a zero gradient and transferred uniformly to the inlet of the active core. Inlet, outlet, and internal velocities are set to be 0.2 m/s upward. For each cell in the model, temperature-dependent cooling is added. In the iMC code, the only consideration for the out-of-core model is its residence time. Therefore, the delayed neutron precursor is considered lost when it escapes through the top of the active core and emits within 10 s. Otherwise, if the emission time left is longer than 10 s, the precursor position is sampled uniformly at the bottom of the active core.

OpenFOAM steady-state analysis is performed with a conjugate heat transfer solver, *chtMultiRegionSimpleFoam* (OpenFOAM,

2025). For the first iteration, which starts with a uniform temperature, the evaluation required roughly 50,000 iterations to achieve a sufficient convergence since the temperature rise at the moderator is extremely slow. From the next iteration, the OpenFOAM requires fewer steps for convergence. For the iMC calculation, 500,000 particles per cycle are utilized for 200 inactive cycles and 500 active cycles. It results in 5–6 pcm uncertainty in the infinite multiplication factor k_{eff} . ENDF-B/VII.1 nuclear data library is utilized (Chadwick et al., 2011), including general neutron-induced cross-section, thermal scattering for the graphite, and photon cross-section.

3.2 Computation result

As discussed in Section 2.4, the iteration is terminated when the maximum temperature difference between the previous trial is below 0.1 K. With the temperature convergence criteria, the convergence is achieved in 11 iterations. Table 3 compares the initial k_{eff} value before the iteration and after the convergence. Note

TABLE 3 k_{eff} values for initial and converged states.

Case	No iteration	Converged	
		Flowing fuel	Static fuel
k_{eff}	1.02056 ± 5	0.99660 ± 5	0.99926 ± 4

that the initial case uses a uniform temperature distribution of 873.15 K for fuel and moderator. In addition, the problem is solved with the static fuel flow and converged temperature distribution, which can be found as static fuel case.

Table 4 compares unweighted and adjoint-weighted delayed neutron fractions from the static and flowing fuel by delayed neutron precursor groups (Anousis et al., 1973). Upper table stands for the unweighted fractions and lower table stands for the adjoint-weighted fractions based on the IFP. The upper half shows a fraction of the resulting delayed neutrons compared to the total neutron population. Delayed neutron fractions significantly decrease due to precursor drift, particularly in long-lived groups, underscoring the importance of modeling flow-induced spectral shifts.

The third column of Table 4 is the reduction of the group-wise fraction in pcm, and its fraction compared to the static case. By

comparing groups 1 and 2, a more considerable reduction is observed for group 2, while its half-life is shorter than that of group 1. The phenomenon is due to the recirculation. Although both groups 1 and 2 precursors are likely to escape the active core, precursors of group 1 are more likely to return to the active core. According to the comparison between delayed neutron fractions and the effective fractions, more significant reductions in the effective delayed neutron fractions are observed. For the delayed neutron population, it reduces only when the delayed neutron precursors escape from the active core. However, for the effective delayed neutron fractions, the reduction occurs when it escapes from the active core and shifts from the region with higher importance. According to the nature of the adjoint neutron flux, the central region is more likely to contribute to the fission chain. Regarding the nature of the model, the axial distribution of the delayed neutron and the adjoint neutron flux are cosine-shaped. When the delayed neutron distribution shifts upward due to the flow, most of the delayed neutron precursors move to neutronically less important regions. The delayed neutrons from those regions have a smaller contribution to the fission chain, resulting in a greater reduction in β_{eff} compared to β .

Figure 4 shows the axial and radial distributions of the delayed neutron precursors. According to the axial distribution, the distribution is shifted upward, and 32.0% of delayed neutron

TABLE 4 Group-wise precursor decay constants and delayed neutron fractions in pcm.

Group	1	2	3	4	5	6
Decay constant	0.0133 s^{-1}	0.0316 s^{-1}	0.1208 s^{-1}	0.3034 s^{-1}	0.8520 s^{-1}	2.8610 s^{-1}
Delayed neutron fraction						
Case	Static		Converged	Reduction		
Group 1	22.6 ± 0.20		12.9 ± 0.17	9.7 (42.9%)		
Group 2	117.3 ± 0.68		61.1 ± 0.38	56.2 (47.9%)		
Group 3	111.9 ± 0.67		60.4 ± 0.38	51.5 (46.0%)		
Group 4	250.5 ± 1.4		171.7 ± 0.68	78.8 (31.4%)		
Group 5	103.5 ± 0.65		92.2 ± 0.49	11.3 (10.9%)		
Group 6	43.5 ± 0.32		42.9 ± 0.31	0.6 (1.3%)		
Total	649.4 ± 3.4		441.4 ± 1.2	208.0 (32.0%)		
Adjoint-weighted delayed neutron fraction						
Case	Static		Converged	Reduction		
Group 1	23.0 ± 0.6		10.7 ± 0.2	12.3 (53.4%)		
Group 2	121.2 ± 1.4		52.5 ± 0.4	68.7 (56.6%)		
Group 3	114.6 ± 1.4		54.3 ± 0.4	60.3 (52.6%)		
Group 4	258.2 ± 2.1		162.7 ± 0.7	95.5 (36.9%)		
Group 5	107.4 ± 1.3		93.5 ± 0.5	13.9 (12.9%)		
Group 6	45.2 ± 0.8		45.0 ± 0.4	0.2 (0.4%)		
Total	669.9 ± 3.7		418.9 ± 1.2	251.0 (37.4%)		

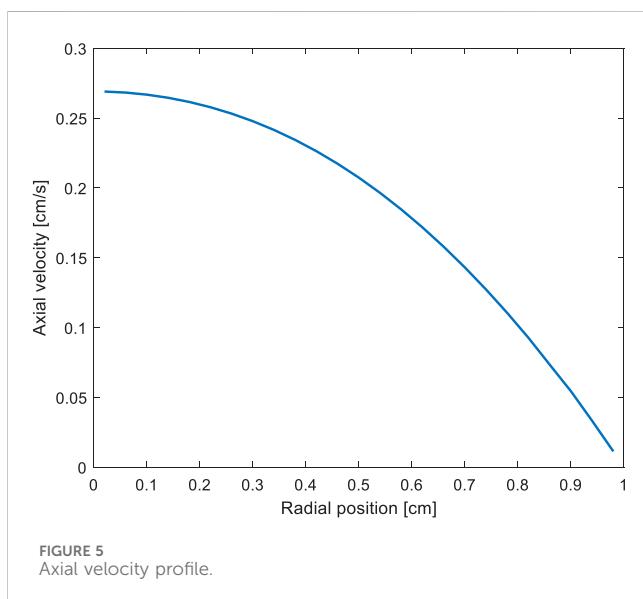
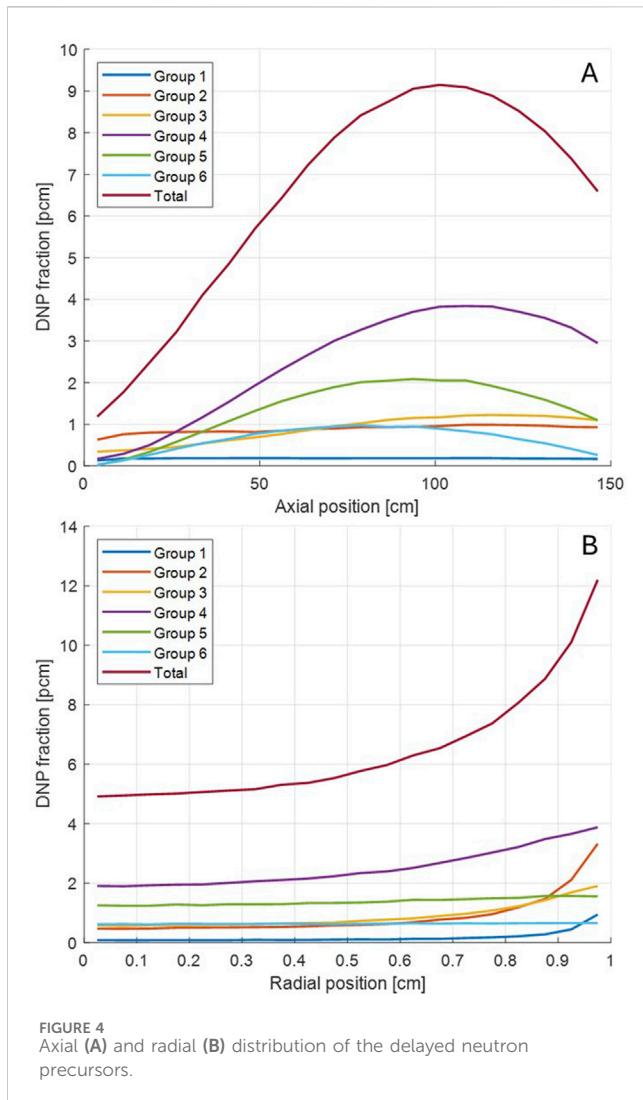


TABLE 5 Heating tallied from each region and particle type.

Power [W] (%)	Fuel	Moderator	Total
Neutron	313,868 (89.6%)	6,747 (1.9%)	320,615 (91.6%)
Secondary particles	8,397 (2.4%)	20,986 (5.9%)	29,383 (8.4%)
Total	322,265 (92.1%)	27,734 (7.9%)	350,000 (100%)

precursors decay out of the active core. Regarding the radial distribution, according to the velocity field from Figure 5, delayed neutron precursor concentration is higher at the peripheral region due to lower axial velocity. Although the current model has nearly uniform radial importance, the higher precursor concentration near the interface will cause a more significant impact on other MSR cores.

Table 5 tabulates the contributions of each region and particle to the total heating of 0.35 MWth based on the coupled neutron-gamma transport. The table shows that secondary particle transport is necessary for an accurate heat distribution, especially in the graphite region. Figures 6A,B are 2-dimensional heat distributions in the fuel and the moderator regions, respectively. The fuel heating shows a flat distribution along the x-axis due to the narrow fuel channel size. Figure 6C is an axial distribution of the fuel heating, averaged along the x-axis. Similarly, Figure 6D is a radial distribution of the graphite heating, averaged along the y-axis. Regarding graphite heating, the majority is from secondary photon heating. The graphite heat distributions show that the heat distribution reduces exponentially due to the short traveling distance of the photon in the moderator.

The inlet and outlet average temperatures are obtained: 1065.9 and 1108.6 K, respectively. The temperatures are obtained from the 1-D out-of-core model. For a fixed out-of-core time of 10 s, inlet temperature can be estimated analytically with outlet temperature T_{out} (Equation 10), which also returns 1065.9 K for outlet temperature 1108.6 K. This implies accuracy in the 1-D numerical solver.

$$T(t) = (T_{out} - T_{ref})e^{-\gamma t} + T_{ref}. \quad (10)$$

Figure 7 shows the temperature distributions along the axial midline of the fuel channel, the axial midline of the moderator, and the horizontal midline across the fuel and the moderator regions.

In section 2.3, the delayed neutron precursor splitting is introduced. The method aims to enhance the precision of the delayed neutron-related tallies, including delayed neutron fraction and distribution, without a noticeable computing burden. At the end of the iteration, the temperature and velocity profiles are utilized again without delayed neutron precursor splitting. Both calculations are performed on five computing nodes with 28 cores of Intel® Xeon® CPU E5-2697. Table 6 compares the results regarding k_{eff} , delayed neutron fraction β , effective delayed neutron fraction β_{eff} , and computing time to check the effectiveness of the splitting scheme. The comparison clearly shows that the splitting scheme increases the computing time 5% for $n = 10$, while halving the uncertainty of the delayed neutron-related quantities.

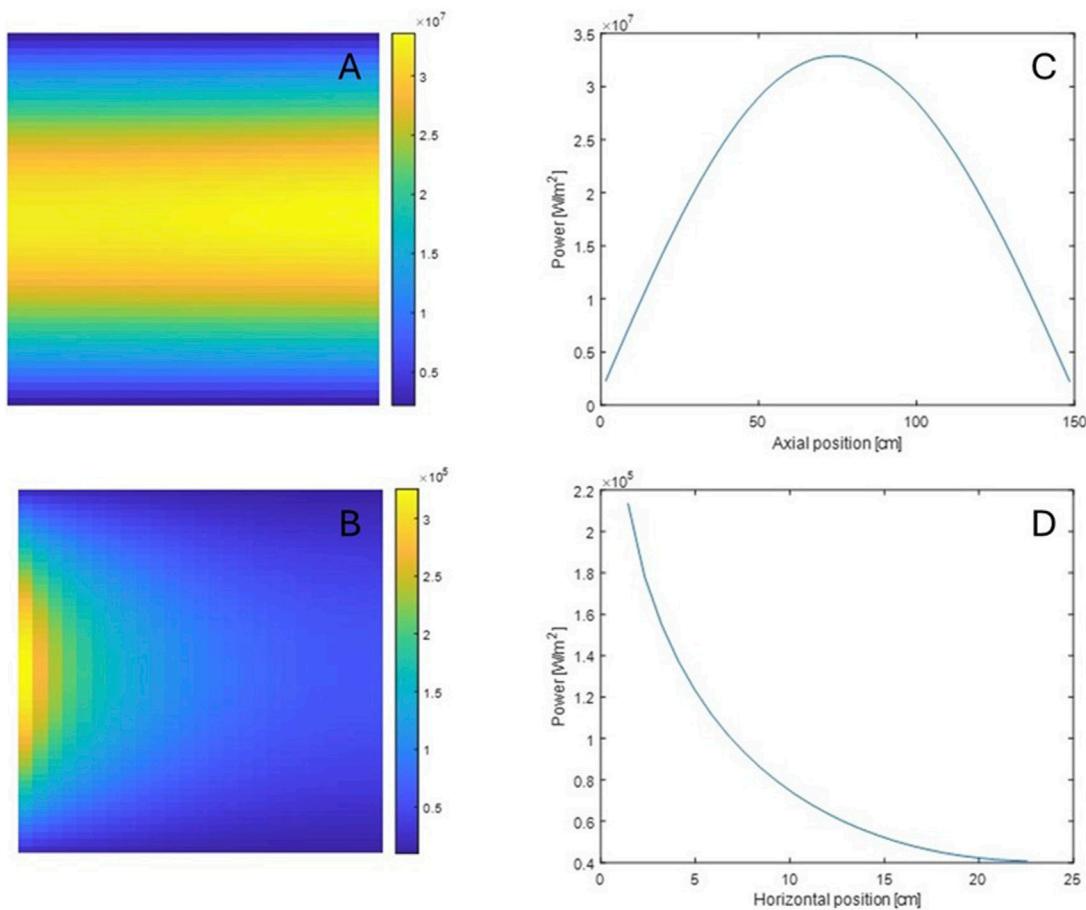


FIGURE 6
2D heat distribution in fuel (A) and moderator (B). Axial fuel heat distribution (C) and Radial moderator heat distribution (D).

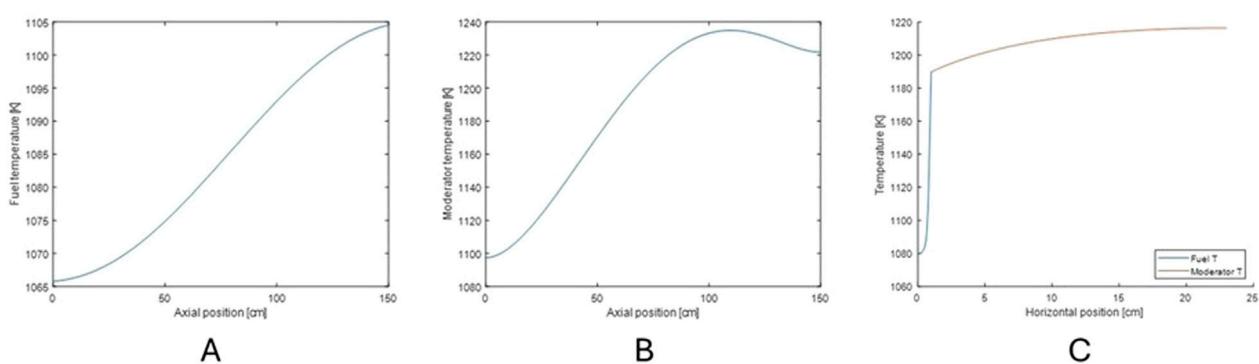


FIGURE 7
Axial temperature along fuel (A) and moderator (B) midlines, and radial temperature along the midline ($y = 75$ cm) (C).

3.3 Mesh size sensitivity analysis

Since the iMC code does not adopt delta-tracking, the number of the temperature mesh is powerfully relevant to the computing time. Therefore, in the coupling process, the temperature fields of the fuel and the moderator are collapsed into coarser meshes. Hence, coarser

meshes are utilized to decrease computing time, although this may potentially lead to inaccuracies. This section measures the impact of the mesh size on the neutronics analysis.

Two meshes are additionally defined: fine and bulk meshes. The fine mesh utilizes mesh for the OpenFOAM directly. As mentioned earlier, 47,600 meshes are utilized where vertical mesh size is fixed to

TABLE 6 Effective delayed neutron fractions.

	No split	Split, $n = 10$
k_{eff}	0.99669 ± 5	0.99660 ± 5
β [pcm]	440.4 ± 2.7	441.4 ± 1.2
β_{eff} [pcm]	415.9 ± 2.5	418.9 ± 1.2
Computing time [s]	10,458	11,045

1.5 cm, and horizontal mesh size varies from 0.22 cm to 0.0008 cm. On the other hand, the bulk mesh runs a similar procedure as the normal mesh, but one mesh is used for each fuel and moderator region. The difference in the mesh is applied both to the temperature and the velocity mesh.

First, the multiplication factors are compared for three types of meshes. Mesh types are visualized in Figure 8A. Table 7 tabulates the k_{eff} values for fine, normal, and bulk mesh. The comparison shows a statistically significant increase in the k_{eff} value and a reduction in

the effective delayed neutron fraction found for bulk mesh compared to other mesh types.

Table 7 compares the k_{eff} and β_{eff} to investigate and compare the effect of the velocity and temperature mesh. Regarding the velocity mesh case, the temperature meshes are fixed to the normal mesh, while the velocity mesh varies. The tests show that the impact of the velocity mesh size on the k_{eff} value is statistically negligible, while applying the bulk temperature will result in a biased result.

Figure 8 shows radial description of mesh types, and each mesh type's delayed neutron precursor distributions, plotted similarly with Figure 4. Comparison between the fine and the normal meshes clarifies that no significant difference exists between the two mesh types. Nevertheless, apparent discrepancies are found for the bulk mesh for axial and radial distributions. As shown in the axial distribution, a more shifted distribution can explain the larger reduction in the effective delayed neutron fraction in Table 7. However, its impact on the k_{eff} value is not visible due to comparable uncertainty.

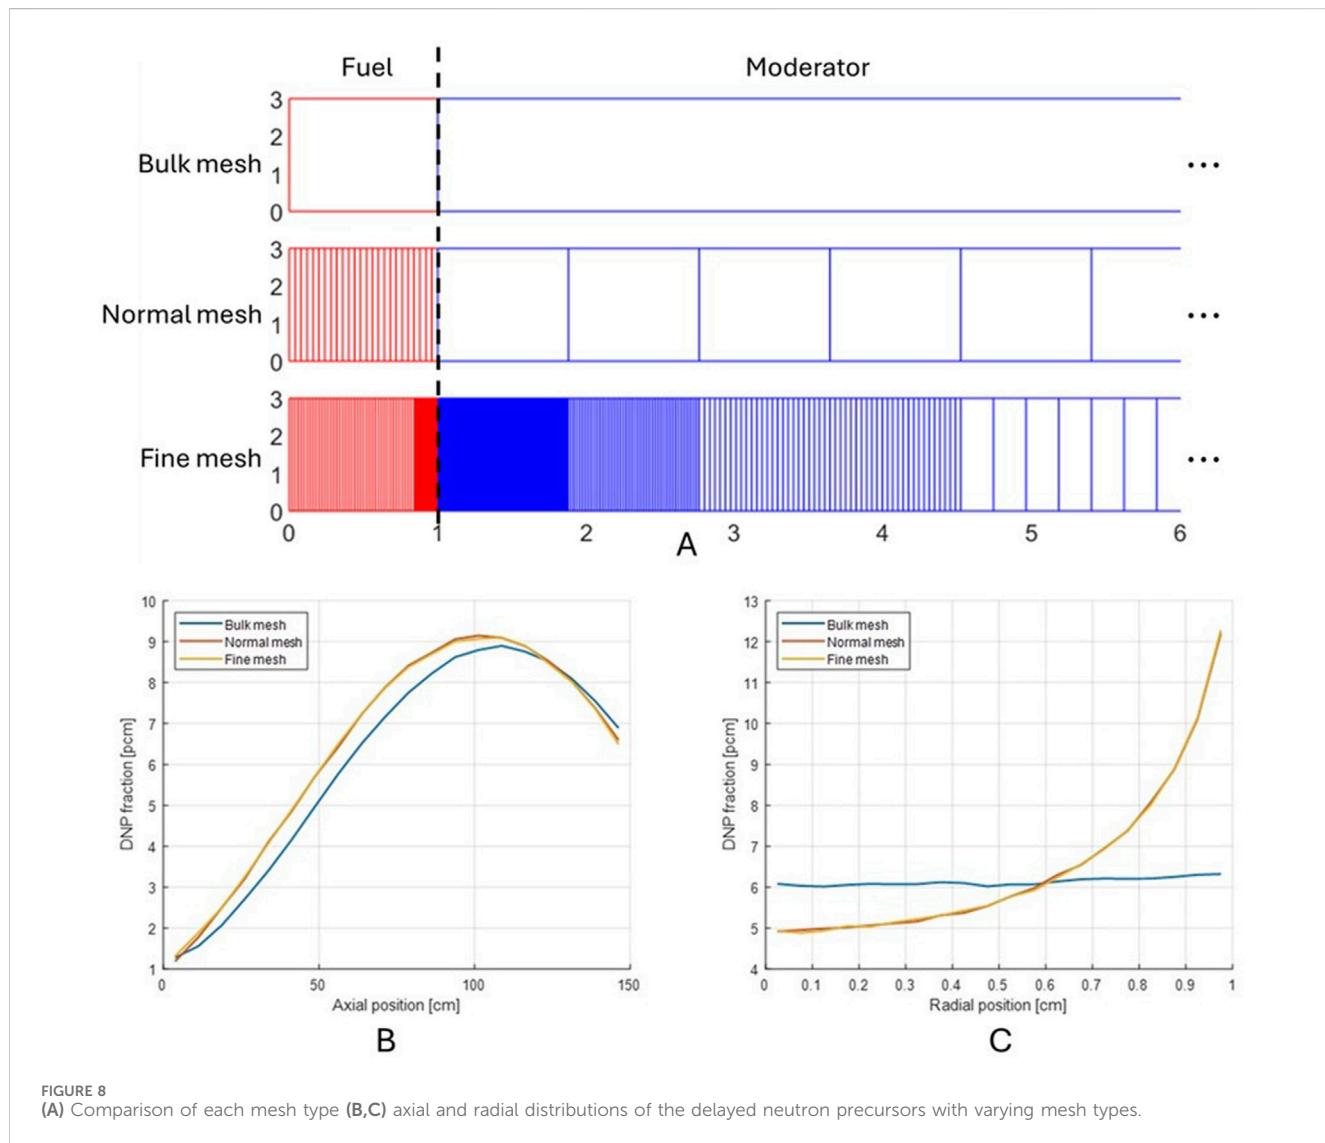


TABLE 7 Mesh-dependent k_{eff} values.

Mesh type			
	Bulk mesh	Normal mesh	Fine mesh
k_{eff}	0.99753 ± 6	0.99660 ± 5	0.99675 ± 6
β_{eff} [pcm]	391.3 ± 1.2	418.9 ± 1.2	418.9 ± 1.2
Velocity mesh type			
	Bulk mesh	Normal mesh	Fine mesh
k_{eff}	0.99647 ± 5	0.99660 ± 5	0.99667 ± 6
β_{eff} [pcm]	391.5 ± 1.1	418.9 ± 1.2	416.2 ± 1.2
Temperature mesh type			
	Bulk mesh	Normal mesh	Fine mesh
k_{eff}	0.99761 ± 6	0.99660 ± 5	0.99652 ± 6
β_{eff} [pcm]	414.3 ± 1.1	418.9 ± 1.2	416.3 ± 1.3

4 Conclusions and future work

This work develops and demonstrates a high-fidelity coupled Monte Carlo–CFD framework for MSR multiphysics analysis. The iMC Monte Carlo code provides a heat distribution based on neutron transport, including its secondary photon transport. The OpenFOAM CFD code performs thermal analysis to obtain the model's temperature profile and the fuel flow's velocity field. The overall coupling scheme and related techniques are covered in this study, including delayed neutron precursor tracking and iteration scheme. By solving the 2D MSR model, a sufficiently converged solution from both iMC and OpenFOAM codes is obtained.

The analysis showed that the reduction in the delayed neutron fraction is underestimated compared to the effective fractions, which denotes an actual impact on the fission chain. The research emphasizes the necessity of evaluating adjoint-weighted kinetic parameters, especially for MSRs. According to the mesh type comparison, the appropriate selection of the mesh size is essential for an accurate evaluation. Despite the overall mass flow rate being conserved in terms of fluid dynamics, its impact deviates from the actual behavior when applied as a single quantity. Recalling that the 2D model in this research has a narrow fuel channel, the impact will be critical for the MSRs with a broader channel.

This paper considers a steady-state case, while one of the tasks of the MSRs is its transient behavior, such as unprotected loss of flow (ULOF). Future development will extend this steady-state coupling to transient scenarios such as unprotected loss-of-flow (ULOF) using dynamic Monte Carlo. Moreover, integrating CMFD-based acceleration and depletion capabilities will enable full-cycle MSR simulations with reduced variance and improved efficiency.

References

Anoussis, J. N., Perricos, D. C., Chrysochoides, N. G., and Mitsonias, C. A. (1973). Relative abundances for six delayed neutron groups from reactor neutron-induced fission of 281Pa. *Radiochim. Acta* 20 (3), 118–119. doi:10.1524/ract.1973.20.3.118

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Author contributions

IK: Writing – original draft, Writing – review and editing. TO: Writing – review and editing. YK: Supervision, Writing – review and editing.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. This work was supported by the Korea Institute of Energy Technology Evaluation and Planning (KETEP) and the Ministry of Trade, Industry and Energy (MOTIE) of the Republic of Korea (No. RS-2023-00304393 and No. RS-2024-00439210).

Acknowledgments

The authors thank Philip Pfahl and Bent Lauritzen for advising the usage of OpenFOAM and providing a 2D MSR benchmark model.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Aufiero, M., Cammi, A., Geoffroy, O., Losa, M., Luzzi, L., Ricotti, M. E., et al. (2014). Development of an OpenFOAM model for the molten salt fast reactor transient analysis. *Chem. Eng. Sci.* 111, 390–401. doi:10.1016/j.ces.2014.03.003

Chadwick, M. B., Herman, M., Oblozinsky, P., Dunn, M., Danon, Y., Kahler, A., et al. (2011). ENDF/B-VII.1 nuclear data for science and technology: cross sections, covariances, fission product yields and decay data. *Nucl. Data Sheets* 112 (12), 2887–2996. doi:10.1016/j.nds.2011.11.002

Cullen, D., and Weisbin, C. (1976). Exact Doppler broadening of tabulated cross sections. *Nucl. Sci. Eng.* 60, 199–229. doi:10.13182/nse76-1

Jo, Y., and Cho, N. (2017). “Refinements of on-the-fly Doppler broadening via Gauss-Hermite quadrature in Monte Carlo reactor analysis,” in *Transactions of the Korea nuclear society spring meeting* (Jeju, Korea: Korean Nuclear Society).

Kaltaisenaho, T. (2020). Photon transport physics in Serpent 2 Monte Carlo code. *Comput. Phys. Commun.* 252, 107143. doi:10.1016/j.cpc.2020.107143

Kim, H., and Kim, Y. (2021). Unstructured mesh-based neutronics and thermomechanics coupled steady-state analysis on advanced three-dimensional fuel elements with Monte Carlo Code iMC. *Nucl. Sci. Eng.* 195 (5), 464–477. doi:10.1080/00295639.2020.1839342

Kim, I., Oh, T., and Kim, Y. (2024). “Development of a delayed neutron precursor tracking module for molten salt reactors in the iMC Monte Carlo code,” in Proceedings of the International Conference on Physics of Reactors, PHYSOR 2024, San Francisco, United States, 21 Apr 2024–24 Apr 2024.

Kim, I., Oh, T., and Kim, Y. (2025). Development and verification of depletion capabilities in the iMC Monte Carlo code. *Ann. Nucl. Energy* 216, 111260. doi:10.1016/j.anucene.2025.111260

Křepel, J., Rohde, U., Grundmann, U., and Weiss, F. P. (2007). DYN3D-MSR spatial dynamics code for molten salt reactors. *Ann. Nucl. Energy* 34 (6), 449–462. doi:10.1016/j.anucene.2006.12.011

Lund, A. L., and Romano, P. K. (2018). *Implementation and validation of photon transport in OpenMC*. Technical Report ANL/MCS-TM-381.

Oh, T., Kim, I., and Kim, Y. (2025). Evaluation of effective kinetic parameters and adjoint flux distribution using iterated fission probability in the iMC Monte Carlo code. *Ann. Nucl. Energy* 210, 110878. doi:10.1016/j.anucene.2024.110878

OpenFOAM (2025). OpenFOAM: user guide: CHTMultiRegionFoam. Available online at: <https://www.openfoam.com/documentation/guides/latest/doc/guide-applications-solvers-heat-transfer-chtMultiRegionFoam.html>.

Pfhal, P. (2025). *Thermal MSR benchmark multi-physics benchmark for a thermal molten salt reactor, submitted to nuclear engineering and design*.

Ridley, G., Huff, K., Turk, M., and Lindsay, A. (2017). *Multiphysics analysis of molten salt reactor transients*. Technical Report No. UIUC-ARFC-2017-01.

Weller, H. G., Tabor, G., Jasak, H., and Fureby, C. (1998). A tensorial approach to computational continuum mechanics using object-oriented techniques. *Comput. Phys.* 12 (6), 620–631. doi:10.1063/1.168744

Yang, G., Jaradat, M. K., Sik Yang, W., and Lee, C. (2022). Development of coupled PROTEUS-NODAL and SAM code system for multiphysics analysis of molten salt reactors. *Ann. Nucl. Energy* 168, 108889. doi:10.1016/j.anucene.2021.108889



OPEN ACCESS

EDITED BY

Deokjung Lee,
Ulsan National Institute of Science and
Technology, Republic of Korea

REVIEWED BY

Evgeny Ivanov,
Institut de Radioprotection et de Sécurité
Nucléaire, France
Fabio Panza,
Energy and Sustainable Economic
Development (ENEA), Italy

*CORRESPONDENCE

Guillaume L. Giudicelli,
✉ guillaume.giudicelli@inl.gov

RECEIVED 13 April 2025

ACCEPTED 30 June 2025

PUBLISHED 25 July 2025

CITATION

Giudicelli GL, Kong F, Stogner R, Harbour L, Gaston D, Lindsay A, Prince Z, Charlot L, Terlizzi S, Eltawila M and Novak A (2025) Data transfers for nuclear reactor multiphysics studies using the MOOSE framework. *Front. Nucl. Eng.* 4:1611173.
doi: 10.3389/fnuen.2025.1611173

COPYRIGHT

© This work is authored in part by Guillaume L. Giudicelli, Fande Kong, Roy Stogner, Logan Harbour, Derek Gaston, Alexander Lindsay, Zachary Prince, Lise Charlot, Stefano Terlizzi, Mahmoud Eltawila and April Novak, © 2025 Battelle Energy Alliance, LLC. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Data transfers for nuclear reactor multiphysics studies using the MOOSE framework

Guillaume L. Giudicelli^{1,2*}, Fande Kong¹, Roy Stogner¹,
Logan Harbour¹, Derek Gaston¹, Alexander Lindsay¹,
Zachary Prince¹, Lise Charlot¹, Stefano Terlizzi³,
Mahmoud Eltawila⁴ and April Novak⁴

¹Computational Frameworks, Idaho National Laboratory, Idaho Falls, ID, United States, ²Nuclear Science and Engineering, Massachusetts Institute of Technology, Cambridge, MA, United States, ³Department of Nuclear Engineering, Pennsylvania State University, University Park, PA, United States, ⁴Nuclear Science and Engineering, University of Illinois at Urbana Champaign, Champaign, IL, United States

High fidelity simulations of nuclear systems generally require a multi-dimensional representation of the system. Advanced nuclear reactor cores are governed by multiple physical phenomena which should be all be resolved, and the coupling of these physics would also need to be resolved spatially in a high-fidelity approach, while lower fidelity may leverage integrated quantities for the coupling instead. Performing a spatially resolved multiphysics simulation can be done on a single mesh with a single coupled numerical system, but this requires catering to each equations' time and spatial discretization needs. Instead, each physics, usually neutronics, thermal hydraulics and fuel performance, are solved individually with the discretization they require, and the equations are coupled by transferring fields between each solver. In our experience coupling applications within the MOOSE framework, mostly for advanced nuclear reactor analysis, there are several challenges to this approach, from non-conservation problems with dissimilar meshes, to losses in order of spatial accuracy. This paper presents the field transfer capabilities implemented in MOOSE, and numerous technical details such as mapping heuristics, conservation techniques and parallel algorithms. Examples are drawn from nuclear systems analysis cases to illustrate the techniques.

KEYWORDS

MOOSE, field transfers, multiphysics, finite element, finite volume, advanced nuclear, fusion

1 Introduction

Modeling and simulation of nuclear reactors is necessary for design, licensing, and operations purposes. The power distribution, the temperature of the fuel, and heat flux-related metrics must be gathered to ensure the nuclear reactor meets safety considerations. For most reactors, obtaining these distributions amounts to solving three groups of equations: neutron transport with depletion and precursors, thermomechanical equations in the solid structures, and thermal-hydraulic equations in the coolant. Over the last 70 years, these equations have been extensively studied and solved for light water reactor analysis. Engineering approximations have been developed and models have been tuned to obtain models that are representative of the reality of their operation. The advent of

a new generation of advanced nuclear reactors is challenging these tools, and the plurality of designs rendered in simulations requires a generality the current tools cannot achieve.

This generality is provided by leveraging unstructured mesh to represent these arbitrary geometries. On unstructured meshes, we can deploy finite element (Schoder and Roppert, 2025; Lindsay et al., 2022), finite volume (Schunert et al., 2021), and sweeping methods (Gaston, 2019; Gaston et al., 2021; Prince et al., 2024) to solve equations. Another challenge comes from the wide variety of physics interactions involved in these new reactors, as many of the new designs introduce physics-based passive safety. Modeling different physics in the reactor is relied upon to provide a safe outcome in the case of accidental transients such as reactivity insertion accidents or loss of flow accidents (Forum, 2024). These features can be modeled at various levels of fidelity and complexity for each physics, but there is a constant need to couple different physics solvers.

The United States Department of Energy Nuclear Energy Advanced Modeling and Simulation (NEAMS) program supports the development of a suite of mixed-fidelity simulation tools based on the Multiphysics Object Oriented Simulation Environment (MOOSE) (Lindsay et al., 2022) C++ framework. These include BISON (Williamson et al., 2021) for fuel performance, Griffin (Prince et al., 2024) for neutronics, SAM (Hu et al., 2021) for systems thermal-hydraulics, Pronghorn (Schunert et al., 2021) for multidimensional thermal-hydraulics, and a selection of MOOSE-wrapped codes. MOOSE provides these codes with support for unstructured mesh via an interface to libMesh (Kirk et al., 2006). Unstructured meshes enable the accurate representation of the geometry of these advanced reactors. The shared framework enables an efficient and easy-to-use coupling of the multiple physics, which is the subject of this paper.

About two-thirds of the first three sections of this paper were presented at the Super Computing for Nuclear Applications conference in Paris in October 2024 (Giudicelli et al., 2024a). Additional techniques and figures were added to these sections, the content was re-organized and reworked to improve clarity, and additional examples and results were included in the final section.

2 Coupling simulations with MOOSE-based or wrapped codes

MOOSE was originally designed to attempt to solve every equation in a single Preconditioned Jacobian-Free Newton Krylov method-based solve. With potentially an additive, multiplicative or Schur complement field split, all the equations could be written as a single nonlinear problem. This was successful for a very narrow selection of problems featuring very strong interactions between equations Wang et al. (2017). For most others, differences in required spatial or time discretizations, and in numerical methods to solve equations, rendering this approach inefficient to unworkable.

Hence, this initial approach has been superseded for most advanced reactor simulations with an operator split in which

each simulation is performed independently and coupled through fixed point iterations (Gaston et al., 2015). MOOSE greatly simplifies this approach by introducing the MultiApps system. MultiApps are created by a parent application, often in a distributed manner, to nest other physics solves. For example, a Griffin neutronics simulation may spawn BISON MultiApps for each fuel pin calculation. A hierarchy of MultiApps is shown in Figure 1.

Once the simulations are created, they can exchange data using the Transfers system. Transfers are created by users to pass fields and scalars between applications. Fields are multi-dimensional, while scalars are typically reductions computed from the fields. They are capable of projecting fields across different meshes, even non-overlapping. Transfers handle the complexities of distributed simulations, for example, if the parent and child application(s) use different numbers of processes. Field transfers are described in more detail in Section 3, while scalar transfers are described in Section 4.

An additional point of interest is that MOOSE-based and MOOSE-wrapped applications are easily combined. The build infrastructure makes building applications combining several MOOSE-based or wrapped applications a trivial task. If the codes are not built together, an application can still be dynamically loaded by another at runtime. The user may specify in the input the application type for the MultiApp and the dynamic library to load to create this application.

2.1 Parallel execution of MultiApps

MultiApps are distributed to maximize the utilization of the machine with an equal split distribution of the processes by default among child applications. There are no set limits to the number of processes used. For example, in Figure 1, if the main application is run with 12 processes, then *MultiApp1* is first solved with 12 processes split between the two child applications. Then, *MultiApp3* is solved with six processes split between the two grandchild applications, while child application four to one is solved with six processes concurrently. Child application 2-1, two to two, and two to three of *MultiApp2* are solved with four processes each once *MultiApp1* has been run. Child apps within a MultiApps are solved concurrently.

Executing several MultiApps concurrently is not currently possible but is a topic of interest as this would let a CPU-heavy and a GPU/accelerator-heavy MultiApp execute concurrently and make the best use of available resources. This limitation is due to the sequential execution of each MultiApps, and parallelization being only performed over child applications.

2.2 Siblings transfers

A recent improvement of the Transfer system is the development of child-to-child application transfers, which we refer to in publications related to MOOSE as siblings transfers (Giudicelli et al., 2024b). Previously, MOOSE could only handle transfers between parent and child applications, so transferring between two child applications of different MultiApps would

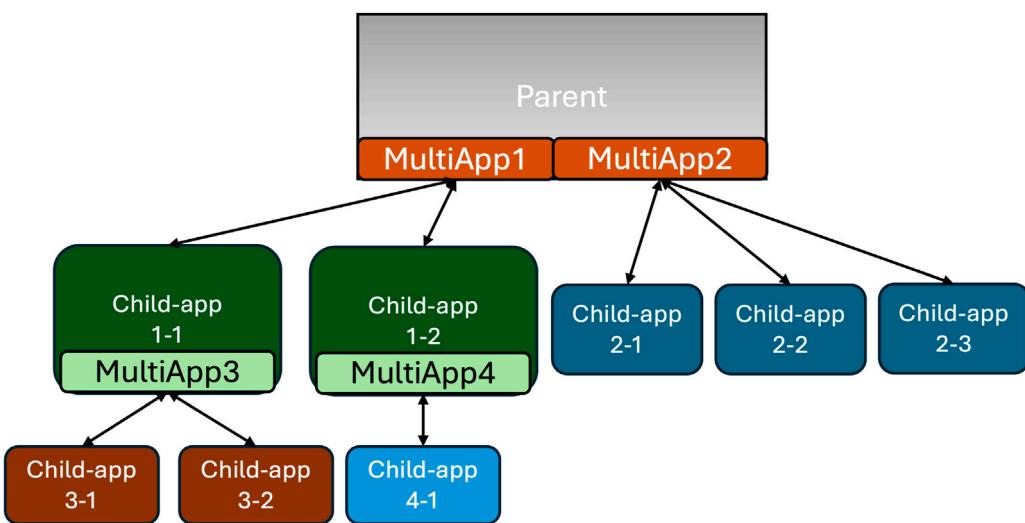


FIGURE 1
Two-level hierarchy of Multiapps. MultiApps created in the parent app has two independent child applications, each having a MultiApp creating child application(s). The bi-directional arrows represent data transfers between applications.

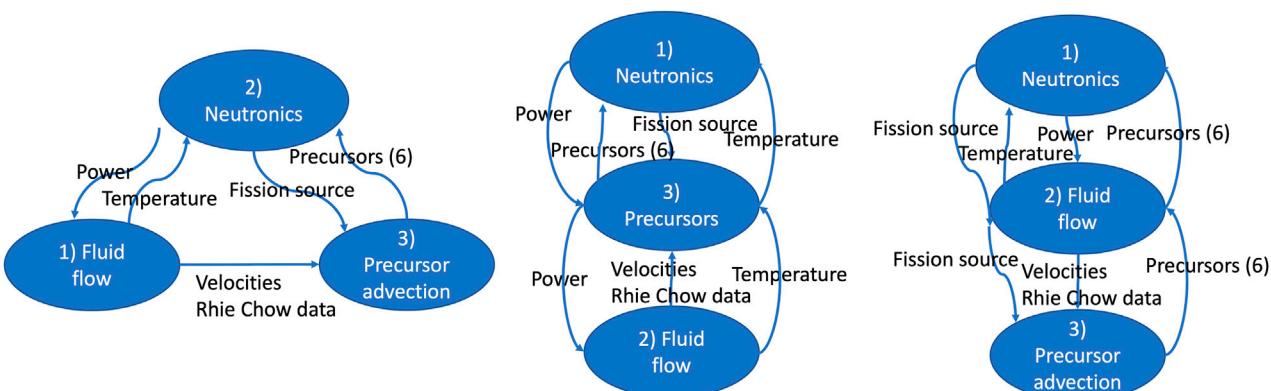


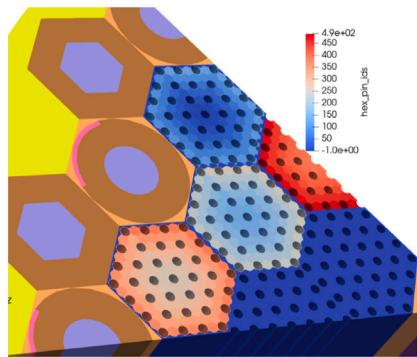
FIGURE 2
Multiphysics core coupling schemes for a molten salt reactor core analysis case, with sibling transfers (left) enabling a simpler scheme and lower duplication of fields alternative schemes (two alternatives shown on the center and right). The application at the top is the parent and the numbering indicates the order of execution. By default, transfers to an application are executed before this application executes, transfers from an application are executed after this application executes.

require one transfer from a child to a parent application, and another from the parent to the other child application. Sibling transfers are widely supported across transfers, including all the field transfers mentioned in [Sections 3, 4](#).

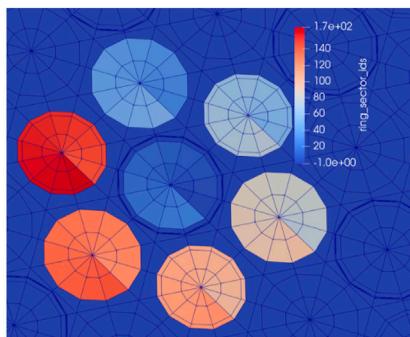
An example of these transfers is shown for a molten salt reactor (MSR) core multiphysics case in [Figure 2](#). The neutronics calculation requires the temperature field from the thermal hydraulics calculation, which itself requires the local heat source to solve for the energy conservation equation. The delayed neutron precursor equations require the velocity fields from thermal hydraulics to solve for their advection, and the fission source from the neutronics to solve for their production. The delayed neutron precursor concentration is then transferred back to the neutronics solver to

describe the delayed neutron source, a part of the fission source. The siblings transfer capability enables the leftmost organization of the solves, which does not require the duplication of both fields and transfers in the other schemes. The reader should note that the order of execution of the solvers can be re-arranged. For example, the neutronics could be executed first, then send the power distribution to the thermal hydraulics, then using that velocity field, the precursor advection equations can be solved.

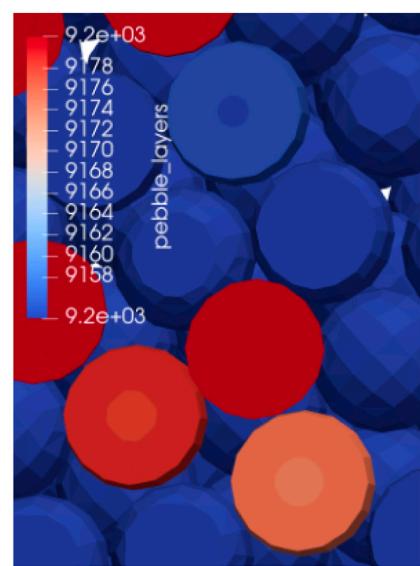
Of note, the recent development of multi-system capabilities, solving multiple nonlinear systems on the same mesh, enables defining both fluid flow and precursor advection equations in the same input file. This negates the need for duplicating fields, and thus for siblings transfers for this MSR application.



a. Nested indexing of pins within assemblies in a micro-reactor core, mesh from (Stauff et al., 2021)



b. Nested indexing of rings of sectors in fuel pins of a micro-reactor assembly



c. Nested indexing of pebble rings in a high temperature pebble bed gas reactor

FIGURE 3
(a) Nested indexing of pins within assemblies in a micro-reactor core, mesh from (Stauff et al., 2021). (b) Nested indexing of rings of sectors in fuel pins of a micro-reactor assembly. (c) Nested indexing of pebble rings in a high temperature pebble bed gas reactor. Mesh divisions used to match transfer regions in advanced reactor cases.

TABLE 1 Coordinate collapsing.

Coordinate system 1	Coordinate system 2	Resulting coordinate system for data Transfer
Cartesian	axisymmetric	axisymmetric
Cartesian	radially symmetric	radially symmetric
axisymmetric	radially symmetric	radially symmetric

3 Field-to-field transfers

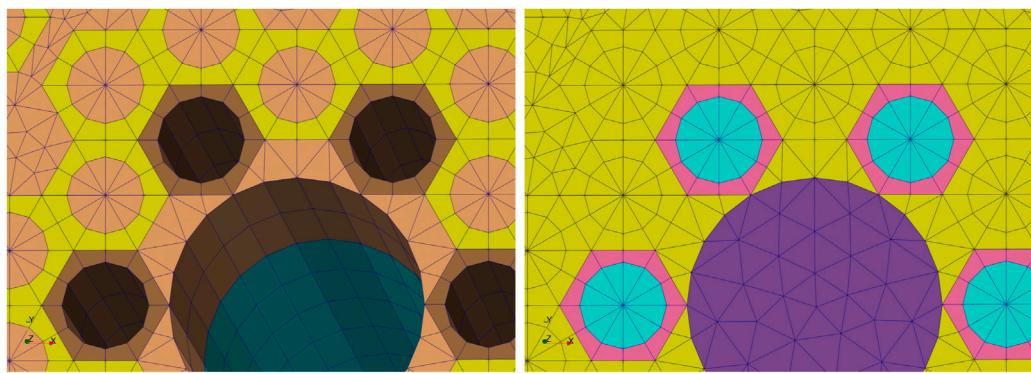
Transferring a field from a simulation that computes or generates it to a simulation that is influenced by this quantity is a common need. For example, the neutronics solve generates the power density field, which is necessary for the thermal physics solve. Transferring a field between two solves is non trivial. The two solves may be using different meshes, may be partitioned differently in parallel, or may have different finite element representations for the field. We first present the communication algorithm in Subsection 3.1, then the field reconstruction algorithm in 3.2, then the several options for providing source values in 3.3, how to conserve integral quantities in 3.4, how to map regions in 3.5, and finally automated diagnostics to look for ill-posed transfers in 3.7.

3.1 Asynchronous transfer algorithm

The communications required for field transfers between (perhaps distinctly) parallel-distributed mesh data are handled

via the asynchronous sparse communication algorithms made available in MOOSE by the Templated Interface to MPI (TIMPI) library. This library provides a wealth of algorithms to perform communication operations between processes, from very simple `communicator.gather(vector)` operations to the parallel asynchronous push-pull algorithm which we leveraged here.

When each MPI rank has gathered some amount of data to send to some subset of other MPI ranks, in a “sparse push,” for them to act on, a nested map of containers of these data is provided to be sent, and a function object (commonly a C++ lambda) is provided to act on received data. Container data types are used to automatically infer or create a corresponding MPI type (for fixed-size data) or a buffering strategy (for variable-size data), and the API then handles all “sends” and “receives” of data, and calls the provided function object with each received vector. Sends and receives are done asynchronously via the non-blocking exchange (NBX) algorithm, allowing for proper scalability up to large processor counts where the send pattern is typically extremely sparse. For a “sparse pull,” in which data are requested from other



a. Homogenized heat conduction mesh: non-solid regions are not meshed **b.** Neutronics mesh: all regions are meshed, including coolant and control rod channels

FIGURE 4

(a) Homogenized heat conduction mesh: non-solid regions are not meshed. (b) Neutronics mesh: all regions are meshed, including coolant and control rod channels. Example of the commonly found indetermination case for the nearest-node transfer. If transferring the graphite matrix temperature from the left heat conduction mesh onto the neutronics mesh on the right (for example, for cross section tabulation purposes), the value at the center of each fuel pin is not uniquely determined with a 'nearest' node or element centroid approach.

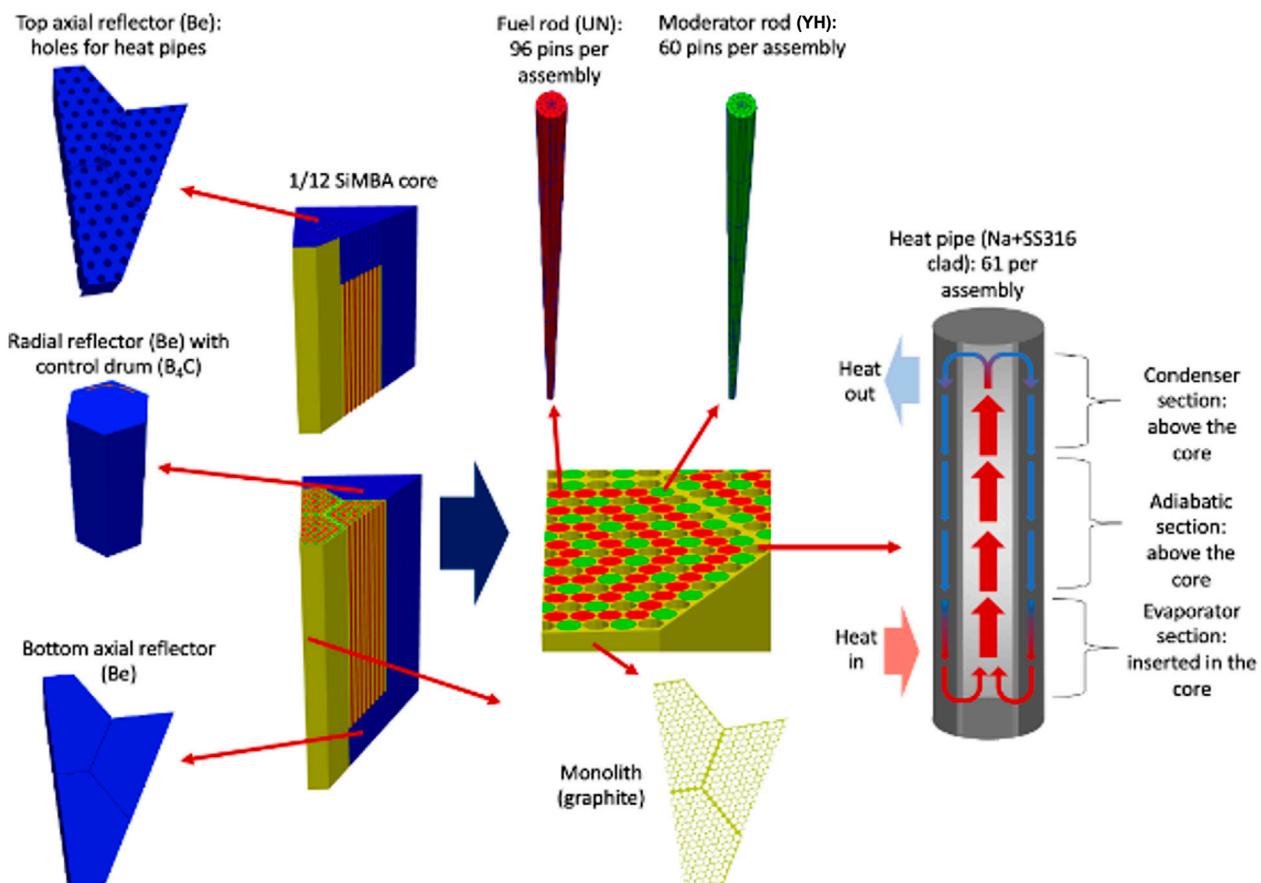


FIGURE 5
Geometry of the HPMR-H₂ problem.

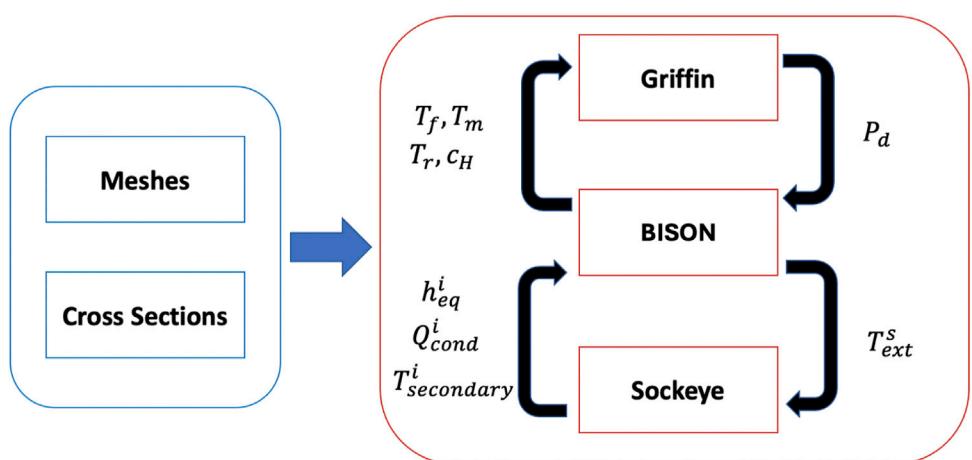


FIGURE 6
Integrated computational workflow for the HPMR-H₂ problem. Courtesy of (Terlizzi and Labouré, 2023).

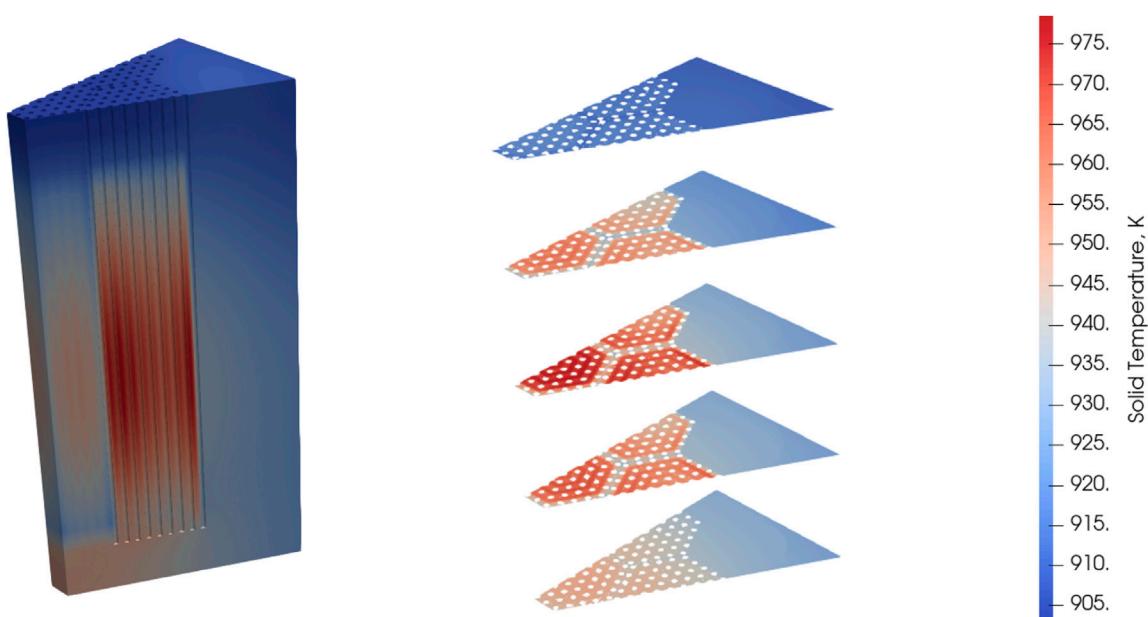
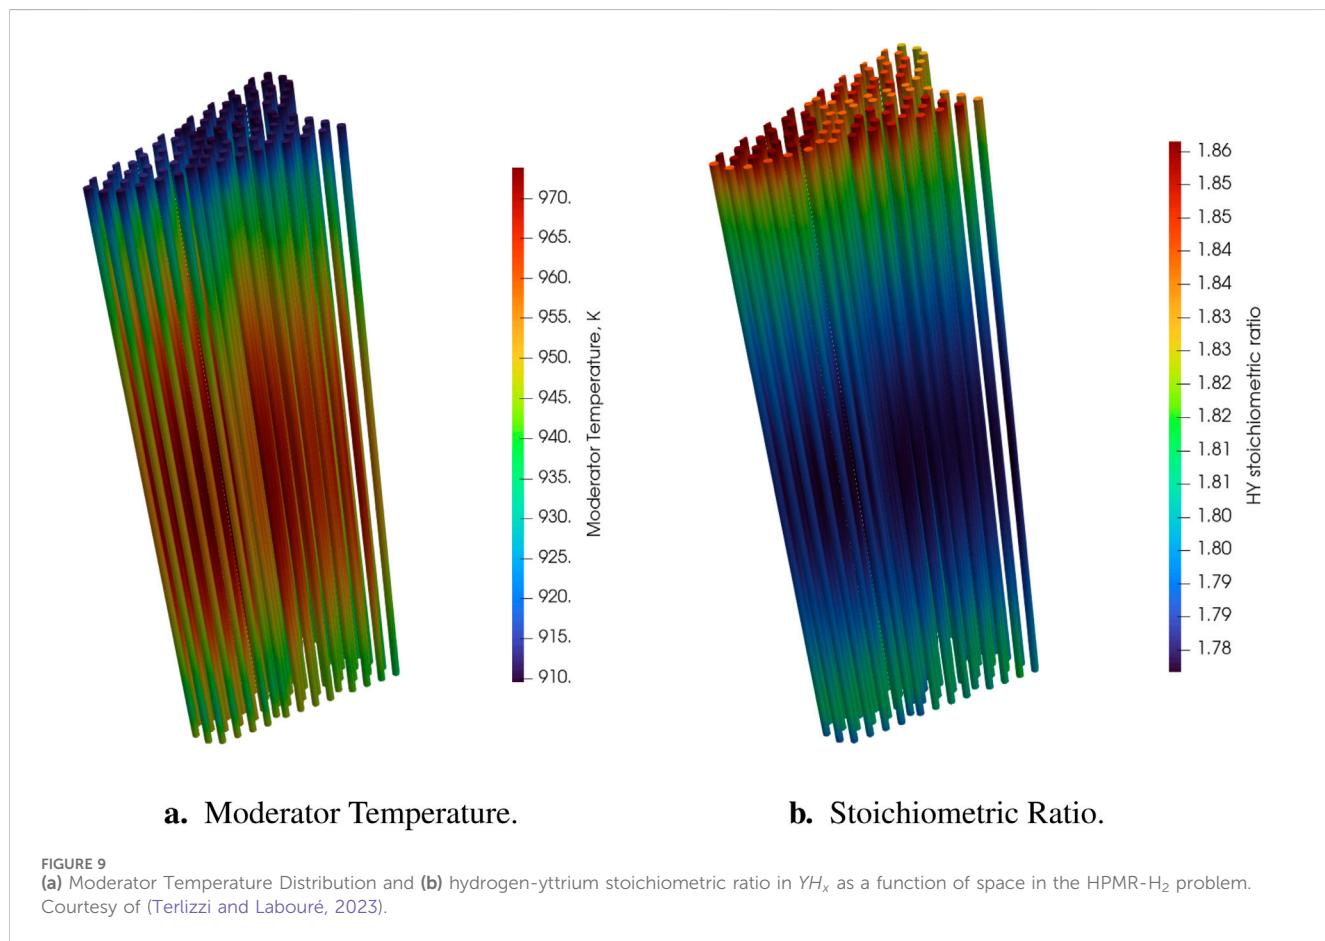
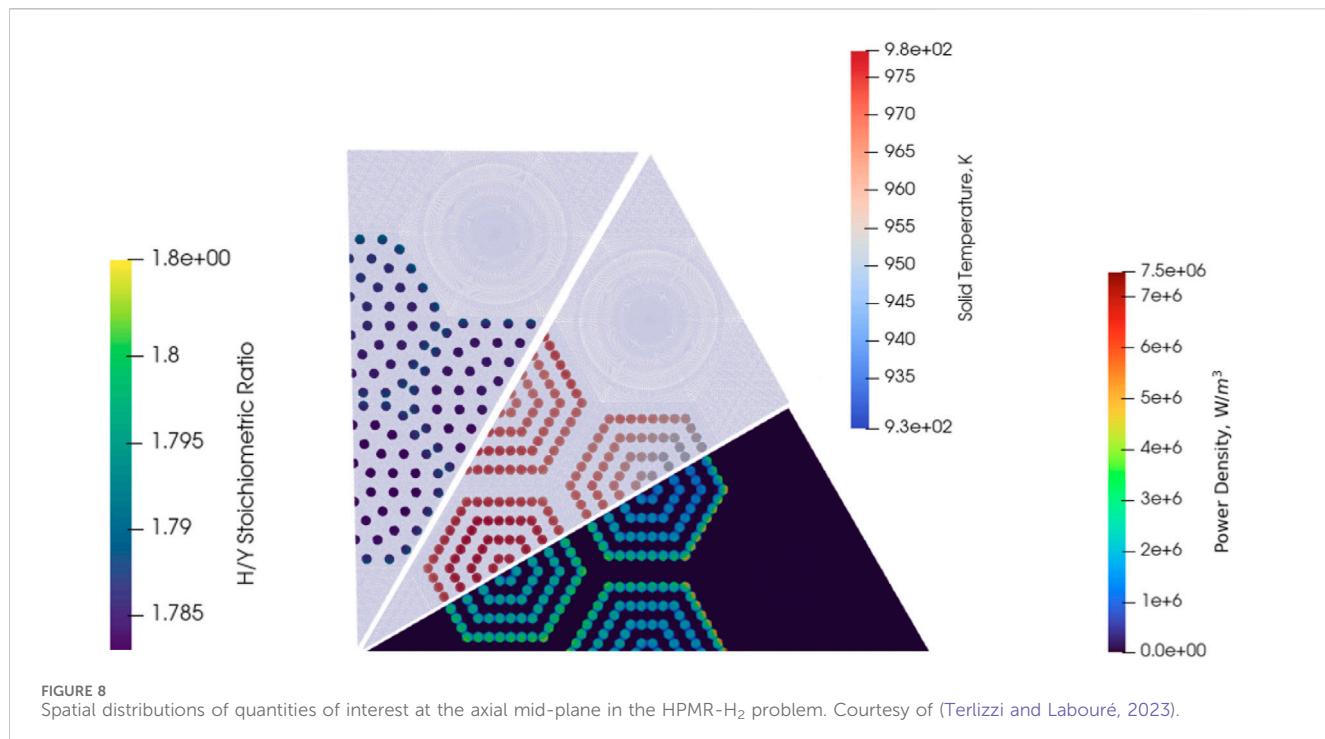


FIGURE 7
Three dimensional temperature distribution in the solids and 2D spatial distribution for selected axial planes in the HPMR-H₂ problem. Courtesy of (Terlizzi and Labouré, 2023).

processors, a push of request keys (such as mesh indices or geometric locations) is followed by a response push of the requested data values associated with those keys; in this case, the transfer code provides a function object with which to collect data for each request and another function object with which to process the collected and received data.

For field-to-field transfers, the “sparse pull” is used to request data at locations specified by the `GenericProjector` (see 3.2) to reconstruct the target field. The function to determine which processes may have these data is based on one of three user-selectable heuristics. The first one is the nearest-app heuristic. Only processes working on the child application that is nearest to

the needed location are contacted to provide a value. The second one is based on forming a sphere of interest around each target point from the application bounding boxes. The bounding box, possibly inflated or fixed using user parameters, around the domain of each process is checked for every process. We find the minimum of the maximum distances within a bounding box from the point of interest. This creates a sphere around the point of interest. Only processes whose bounding boxes intersect this sphere are contacted to provide a value for this target point. The last one is the greedy search, which is mostly used for debugging. This one requests data from every single other rank, and is therefore not scalable.



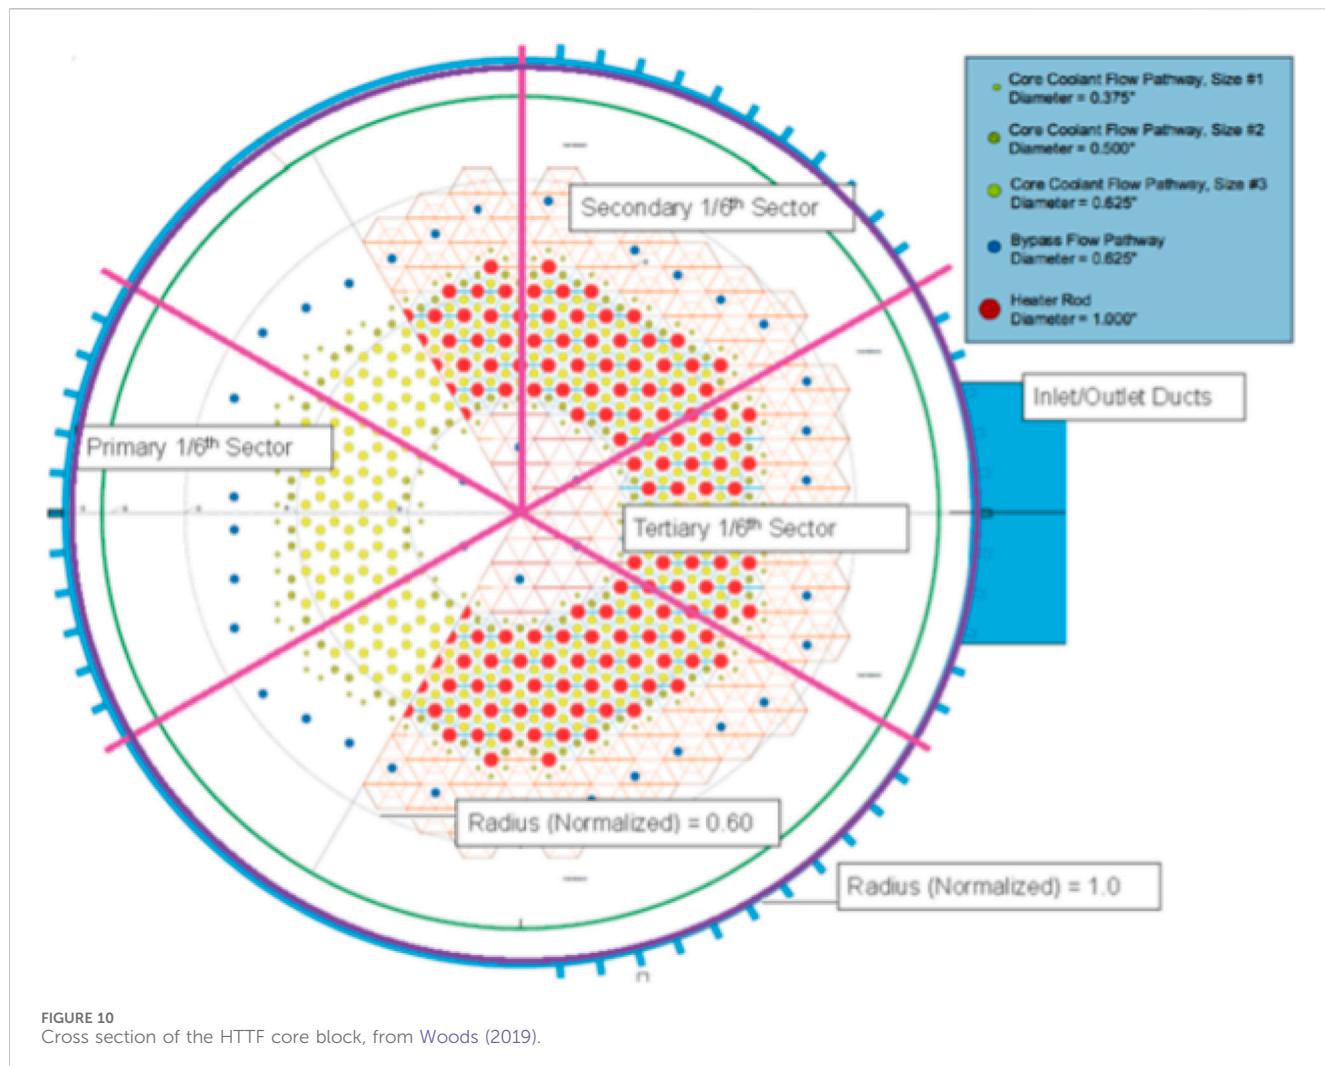


FIGURE 10
Cross section of the HTTF core block, from Woods (2019).

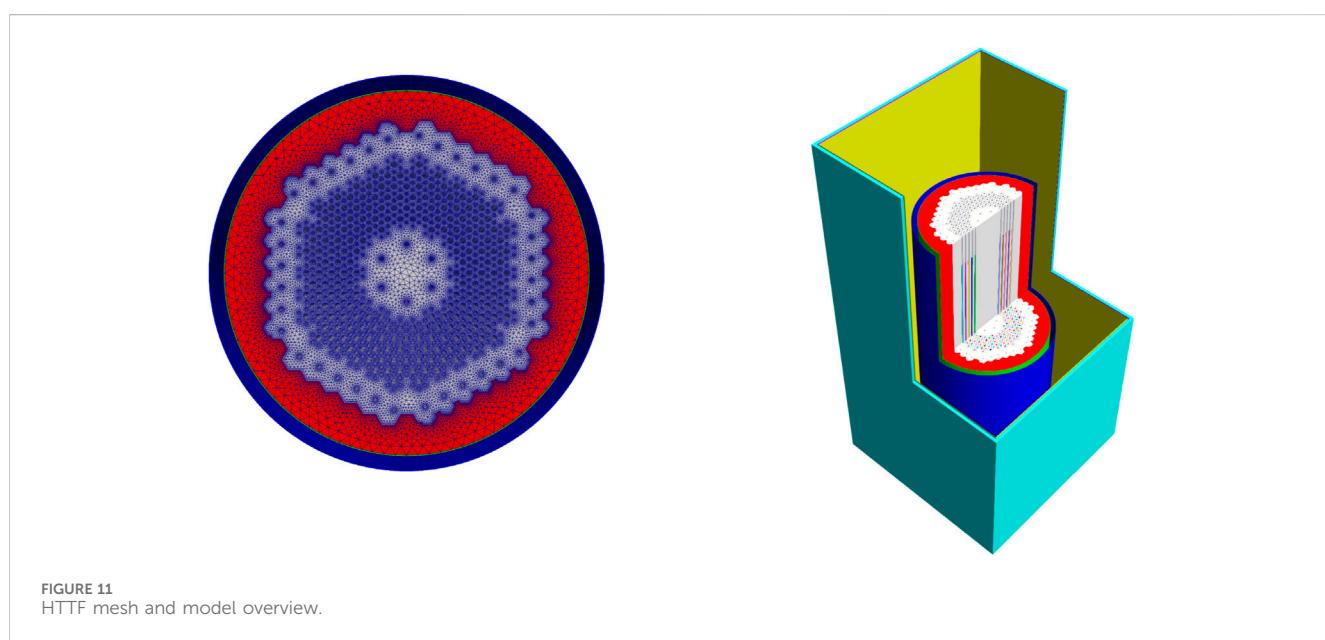
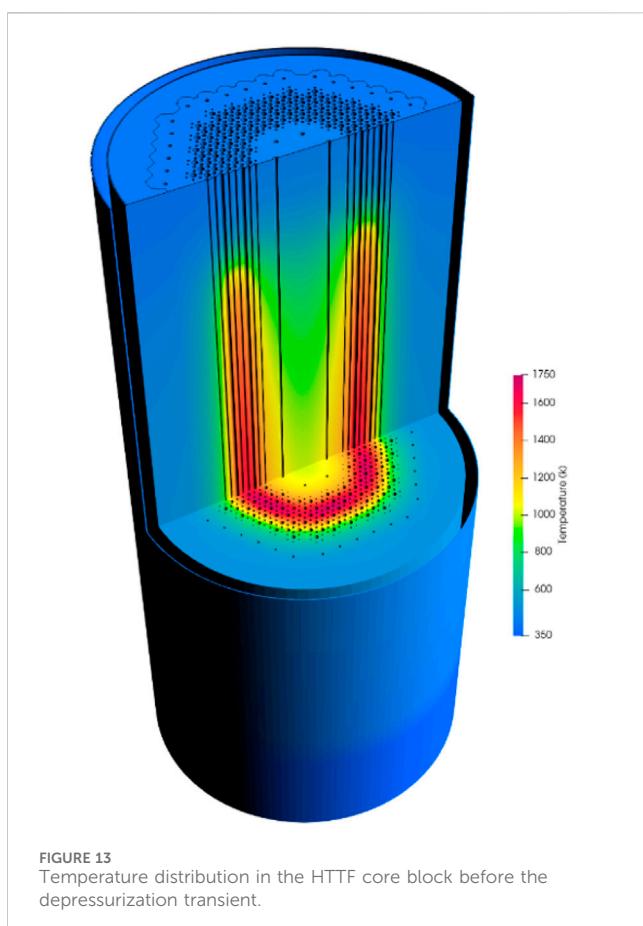
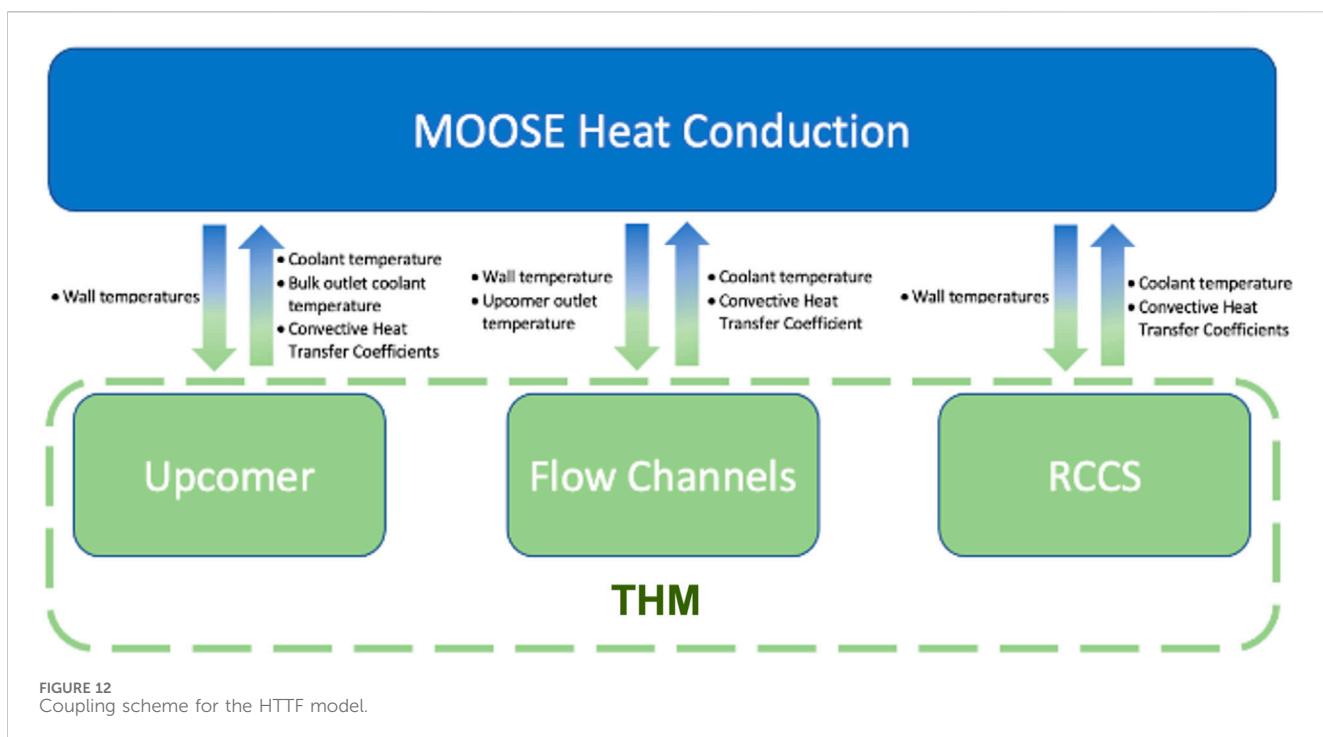


FIGURE 11
HTTF mesh and model overview.



3.2 Generic projection capability in the context of asynchronous transfers

The GenericProjector code in libMesh was designed to enable the restriction of arbitrary input functions to arbitrary supported finite element bases. Input functions are defined by the function values they return at arbitrary points and, in cases of higher-continuity finite element spaces, function gradients are either directly provided or inferred by finite differencing. In a typical libMesh use case, the input function is locally evaluable and the projection operation onto a C_0 -continuous finite element space is thus straightforward: vertex points are first evaluated and directly copied onto vertex degrees of freedom, then quadrature points on element edges are evaluated to enable local L^2 projections of element edge values while holding vertex values constant, then (in 3D) element face degrees of freedom are evaluated via projection holding vertex and edge values constant, then finally element interior degrees of freedom are projected while holding each element's boundary values constant. Each of these local projections is an implicit solve, but with few degrees of freedom, and hence reasonably fast.

In MOOSE, differences between domain and partitioning of transfer source meshes may make it impossible to locally evaluate input functions, so the usage of GenericProjector is slightly more complex. The projection code is evaluated twice: once to determine which quadrature point data will be required from which remote processors, and then a second time to actually evaluate projections and so transfer the data responses correctly to the output field. As with the communications code, the projection code is

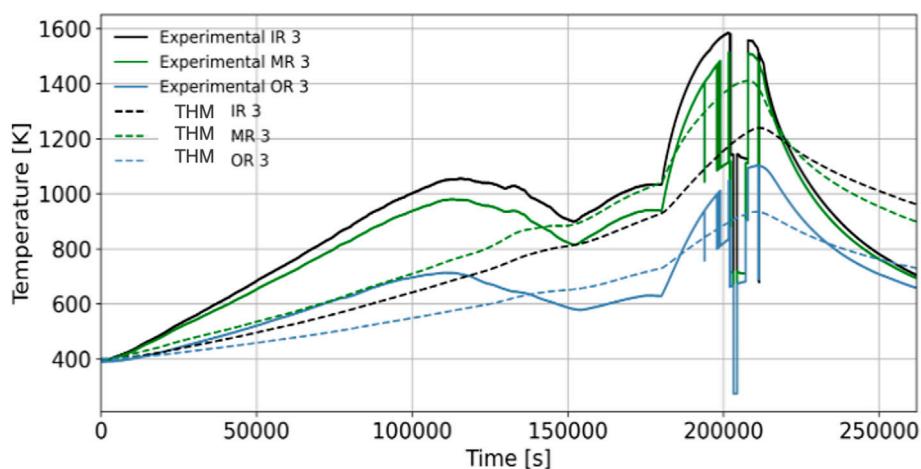


FIGURE 14
Temperature evolution at the bottom of the HTTF core.

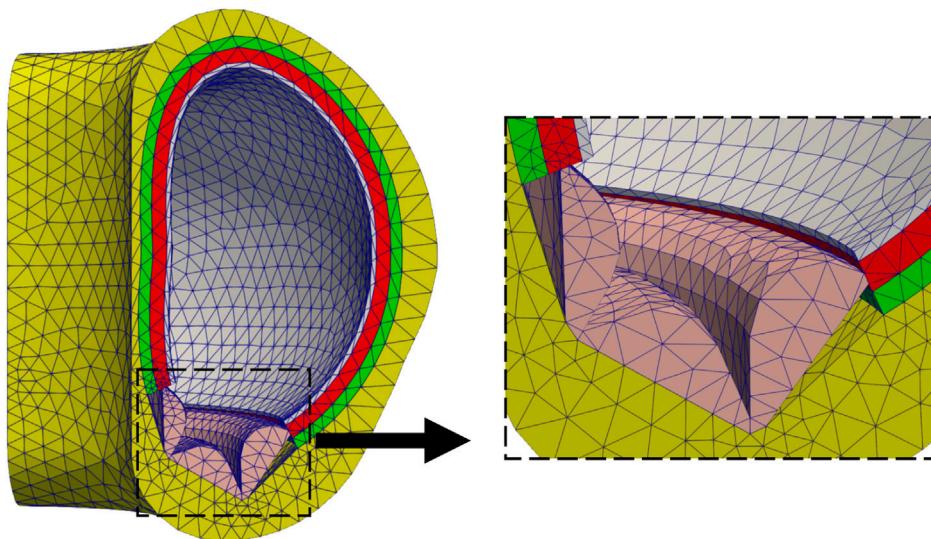


FIGURE 15
Simplified tokamak model mesh and a zoomed-in view on the divertor; colors correspond to different subdomains.

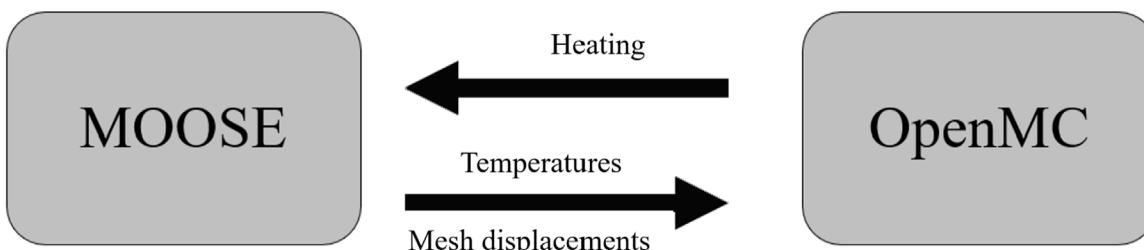


FIGURE 16
Multiphysics coupling scheme for the tokamak model.

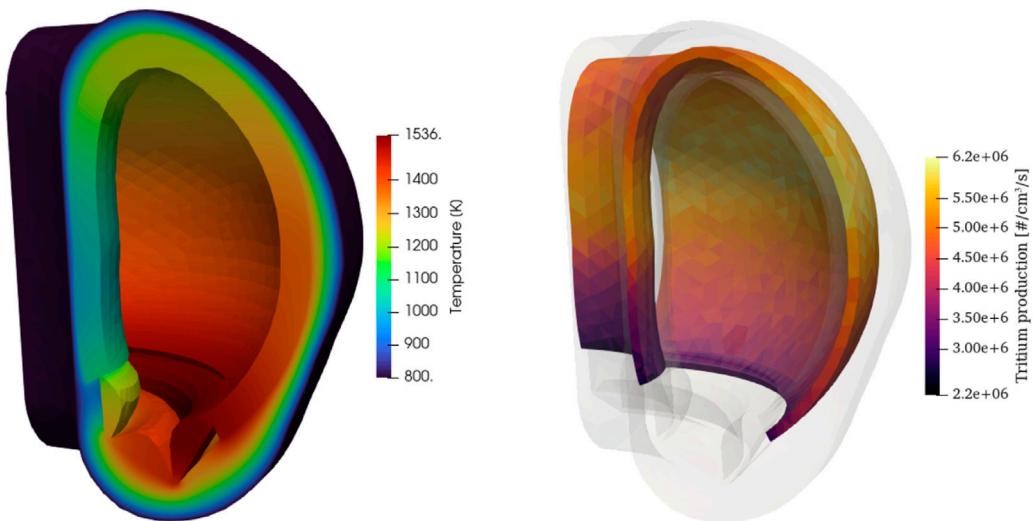


FIGURE 17
Temperature (left) and tritium production (right) results obtained at 6.1 MW for a 90° sector.

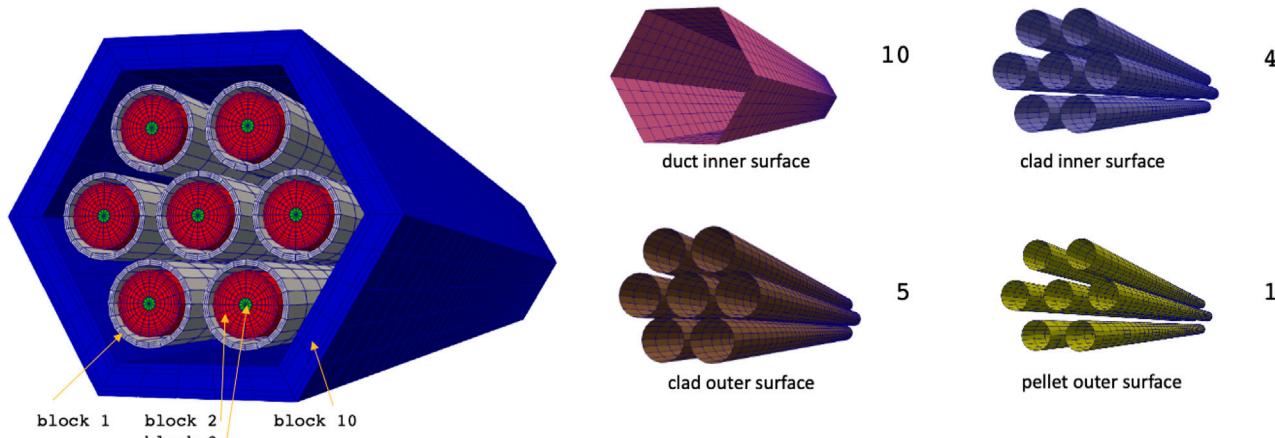


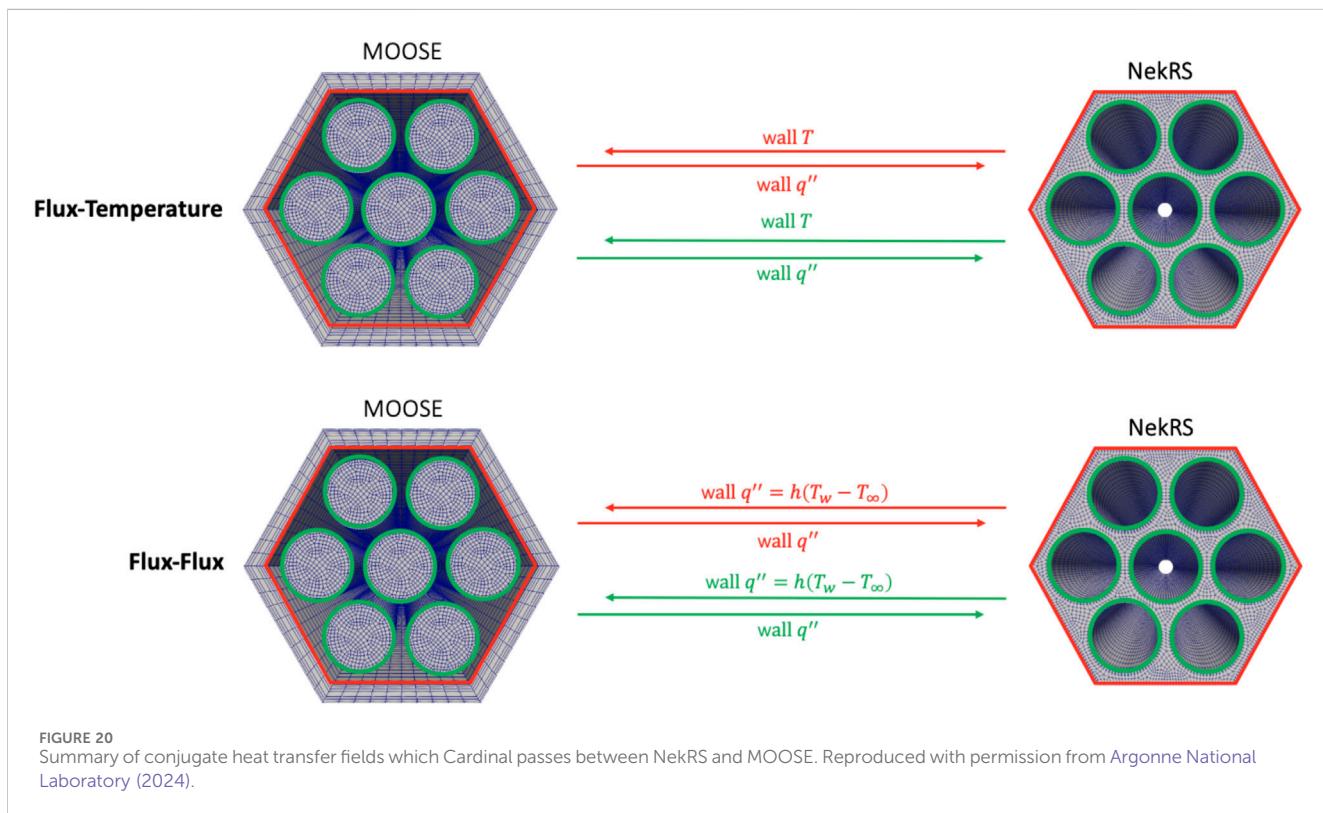
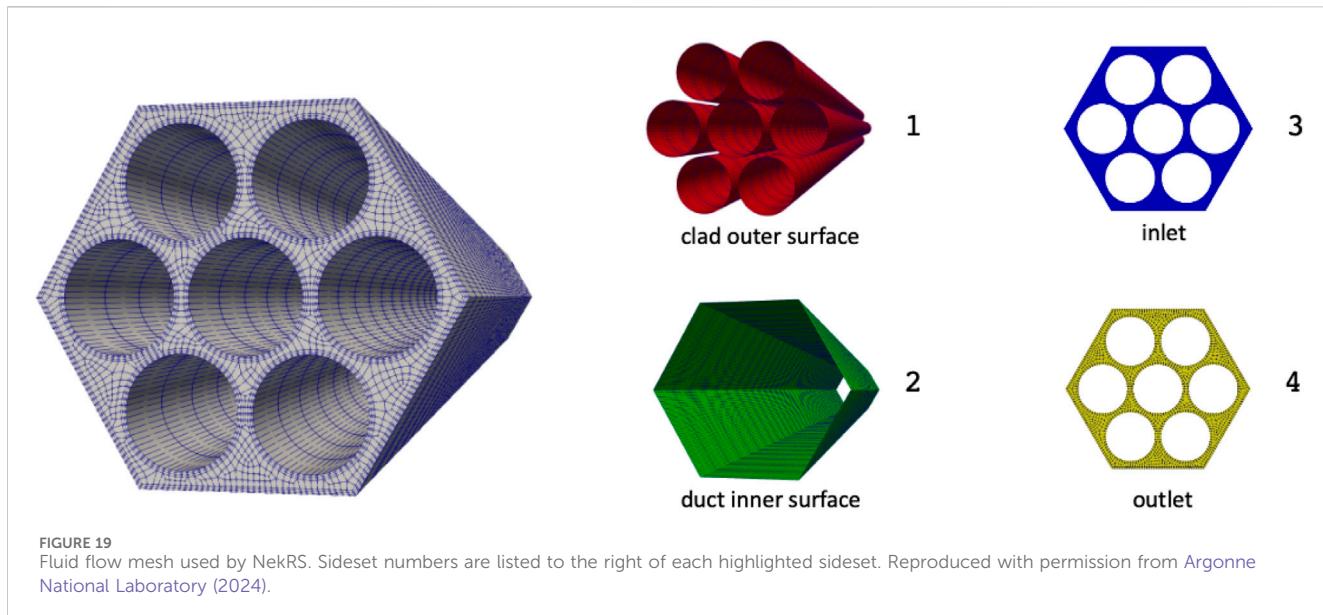
FIGURE 18
Heat conduction mesh used by the MOOSE heat transfer module. Sideset numbers are listed to the right of each highlighted sideset. Reproduced with permission from Argonne National Laboratory (2024).

designed to accept function objects for its innermost operations. The transfer code first provides objects, which merely examine and encapsulate each of the queries to be sent, then provides objects which evaluate cached response data and write projection results to the appropriate field vector coefficients.

3.3 General field transfers

Fields have spatially dependent values on the mesh. Fields in MOOSE are generally variables, though material properties and functions can also be considered fields. There are three main algorithms for performing field transfers in MOOSE. Once the target application has sent its requests for values at certain points, the source application can.

- Evaluate the source variable at the requested points by evaluating the shape functions and performing the finite element reconstruction. This is implemented by the `MultiAppGeneralField-ShapeEvaluationTransfer` class.
- Return the “nearest” or a combination of the “nearest” values, which are formed at the nearest node or nearest element centroid on which the value of the source variable is known. This avoids evaluating the shape functions. This is implemented by the `MultiAppGeneralFieldNearestLocationTransfer` class. This class was previously named `MultiAppGeneralFieldNearestNodeTransfer` and the previous name is still in use.
- Evaluate a source user object at the requested point. This is often to evaluate the source variable through this user object’s operation, for example, layered averaging or other



reduction operations. This is implemented by the `MultiAppGeneralFieldUserObjectTransfer` class.

The implementations of these three algorithms leverage the techniques explained in [Sections 3.1, 3.2](#) through a common base class called `MultiAppGeneralFieldTransfer`. The derived classes are simply defining how to evaluate the value at a target point, and related details like the detection of ill-posedness. The first two transfer options are also found in other data transfer libraries

([Pawlowski et al., 2013](#)), as is the L2-projection transfer which is not described here as it is implemented very differently in MOOSE.

3.4 Conserving integral quantities

An important concern when transferring fields is preserving integral quantities. The most common issue is to preserve power when passing the power distribution from the neutronics solvers to

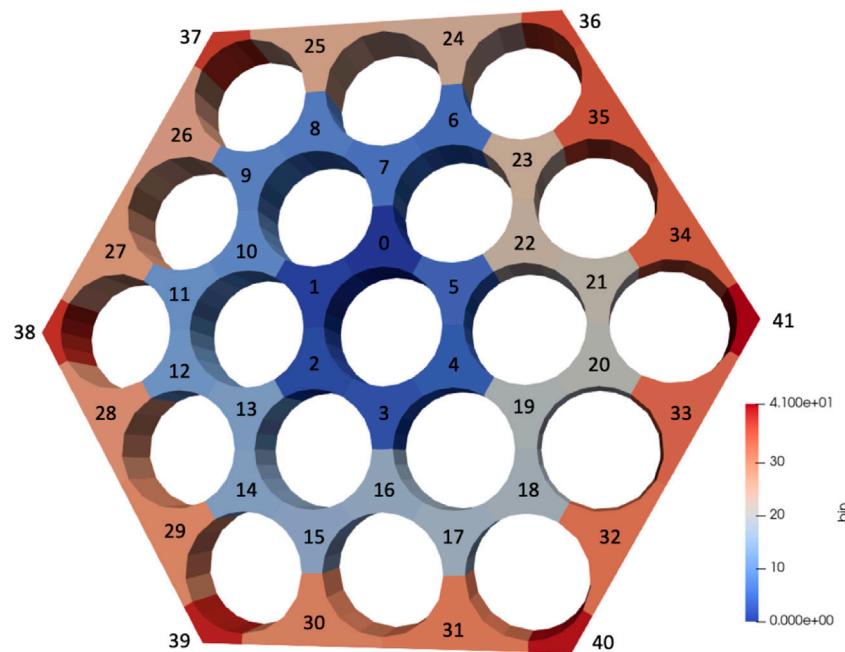


FIGURE 21
Example of bin indices used for a hexagonal subchannel volume binning, for a 19-pin bundle.

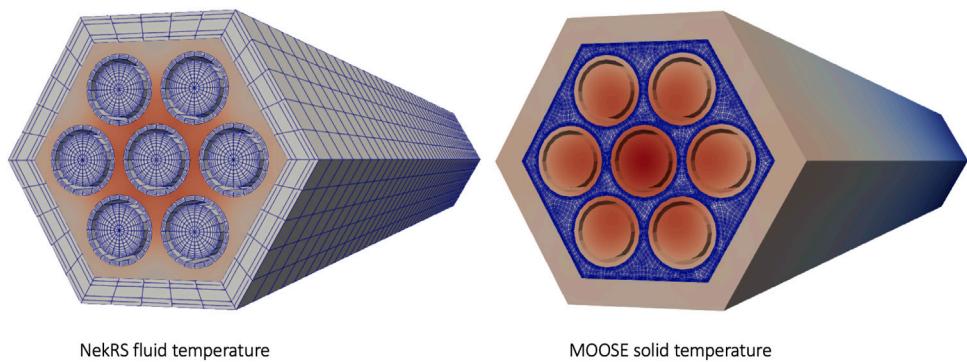


FIGURE 22
Fluid temperature computed by NekRS (left) and solid temperature computed by MOOSE (right) for the flux-temperature coupling option.
Reproduced with permission from Argonne National Laboratory (2024).

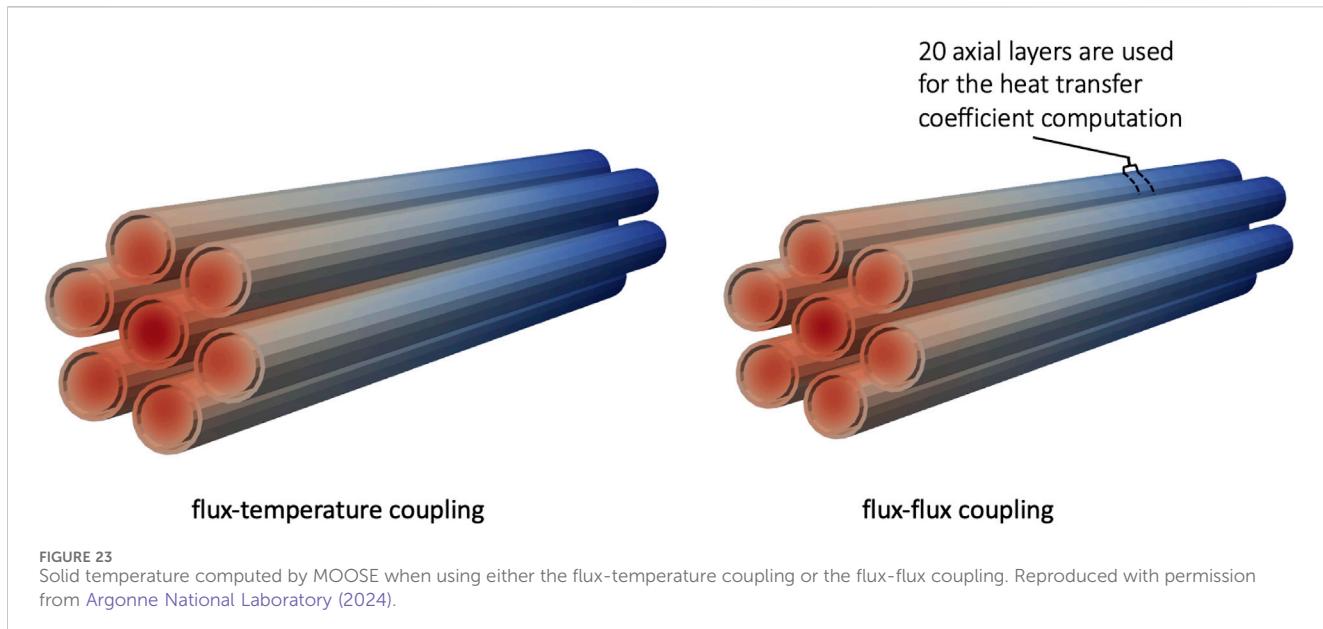
either the fuel performance solver for solid-fueled reactors, or the thermal-hydraulics solver for molten-fuel reactors. For power, the conservation problem is rather simple and only requires a re-scaling by the ratio of integrated power on the source and target meshes. This re-scaling can be distributed if there are multiple target applications, for example, fuel pin solves that have a different local conservation error during the transfer.

At this time, the relation between the field and the postprocessed data is assumed to be linear, so only a single correction step is performed. If the postprocessed data were to be recomputed, we would expect this rescaling to have equalized the ratio between the source and target quantities. However, this would not hold for nonlinear quantities. For example, if the specific heat depends on temperature, rescaling a transferred

temperature field once would not suffice to provide conservation of energy.

3.5 Mapping source and origin regions

With the popularity of nearest-node or nearest-element value approaches, a major concern arises in that the nearest-location value is not actually the correct value to obtain. Differences in axial meshes might make the nearest value for the temperature field come from another solid region. For example, for the temperature field where transfers match fuel pins together, an axial mesh cell for a guide-tube in the source application could be closer to the target mesh cell, and therefore would erroneously provide the source value. MOOSE



offers a wide variety of region-mapping options to make sure this does not happen.

The simplest options offered are to restrict both source and target domains using block (a subdomain of the mesh) and/or boundary restrictions. These restrictions can be combined completely independently, but the user must understand an important detail. Block and boundary restrictions are restrictions, not matches. Specifying two source blocks and two target blocks does not match the first source block to the first target block. It simply specifies that the transfer will find values from both targets blocks in the union of the two source blocks.

Another restriction technique when using multiple child applications is to use only values from the closest child application to each requested point. This technique is also used as a heuristic to determine which processes should communicate, as mentioned in [Subsection 3.1](#). This technique is particularly useful for consistently matching fuel rods to the channels nearby when using a “nearest-node/location” approach. The concern is that a mismatch between the axial meshes of the rods would map cells in a channel nearest to a certain rod to another rod with a locally finer axial mesh.

To offer a much finer-grained mapping capability, the concept of mesh divisions was introduced. Mesh divisions are an arbitrary indexing of the source and target meshes. The source and target mesh divisions can be matched on a per-index basis. This concept can easily cover the block and boundary restrictions options, but these are kept for backwards compatibility and user friendliness. Example mesh divisions are represented in [Figure 3](#) for various reactor applications: a pin indexing in [Figure 3a](#), a pin and rings indexing in [Figure 3b](#), and a pebble and rings indexing in [Figure 3c](#). Numbering these regions enables focused transfers between physics solvers.

3.6 Coordinate transformations

Different simulations may use different coordinate systems. For example, a neutronics simulation may be performed in

Cartesian three-dimensional space whereas a fuel performance calculation may be performed using a two-dimensional axisymmetric coordinate system. Communicating information between these different configurations would be difficult without reconciling the differences in coordinates. One mechanism MOOSE provides for making this communication simpler is the `MooseAppCoordTransform` class. Each `MooseApp` instance holds a coordinate transformation object in its `MooseMesh` object. Users may specify transformation information about their simulation setup on a per-application basis in the input file `Mesh` block. A `coord_type` parameter specifies the coordinate system type, e.g., Cartesian, axisymmetric, or radially symmetric. Euler angles are available to describe extrinsic rotations. The convention MOOSE uses for a α - β - γ Euler angle rotation is:

1. Rotation about the z-axis by α degrees
2. Rotation about the x-axis by β degrees
3. Rotation about the z-axis (again) by γ degrees

A `length_unit` parameter allows the user to specify their mesh length unit. The code in `MooseAppCoordTransform` which processes this parameter leverages the MOOSE units system. A scaling transform will be constructed to convert a point in the mesh domain with the prescribed mesh length unit to the reference domain with units of meters. The last option which contributes to transformation information held by the `MooseAppCoordTransform` class is a `positions` parameter. The value of `positions` exactly corresponds to the translation vector set in the `MooseAppCoordTransform` object of the sub-application. The α , β , γ , and `positions` parameters essentially describe forward transformations of the mesh domain described by the MOOSE `Mesh` block to a reference domain. The sequence of transformations applied in the `MooseAppCoordTransform` class is scaling, rotation, translation, and coordinate collapsing.

The last item in the transformation sequence, coordinate collapsing, is only relevant when information has to be transferred between applications with different coordinate systems. For transferring information from Cartesian to axisymmetric, we must collapse Cartesian coordinates into the axisymmetric space since there is a unique mapping of Cartesian coordinates into axisymmetric coordinates, but not *vice versa*, e.g., a point in the axisymmetric space has infinitely many corresponding locations in Cartesian space due to rotation about the axis of symmetry. Table 1 summarizes the coordinate collapses that occur when transferring information between two different coordinate systems. The table should be understood as follows, using the first row as an example: for a Cartesian-axisymmetric pairing (e.g., axisymmetric → Cartesian or Cartesian → axisymmetric data transfers), both ‘from’ and ‘to’ points will be cast into the axisymmetric coordinate system for the reasoning given above: there is a unique map from Cartesian to axisymmetric, but not *vice versa*. Similarly for an axisymmetric-radially symmetric pairing (e.g., axisymmetric → radially symmetric or radially symmetric → axisymmetric data transfers), both ‘from’ and ‘to’ points will be cast into the radially symmetric coordinate system.

Note that there are consequences for these coordinate system collapses. When transferring data in the $1 \rightarrow 2$ direction, there are (as already stated) infinitely many points in three-dimensional Cartesian space that correspond to a single axisymmetric coordinate. For example, the Cartesian points $(1, 0, 0)$ and $(0, 1, 0)$ map to the same axisymmetric coordinate $(1, 0)$ if the z-axis is the axis of symmetry on the Cartesian mesh. So if we are performing a nearest-node transfer of data from Cartesian to axisymmetric, where the ‘to’ point is $(1, 0)$, then selection of the ‘from’ point is arbitrary if both $(1, 0, 0)$ and $(0, 1, 0)$ points (or any combination of $\sqrt{x^2 + y^2} = 1$ points) exist. We are considering how best to handle these situations moving forward. One option would be to average the field data from equivalent points. Note that for the axisymmetric coordinate system with general axes, only translation is supported for coordinate transformations, i.e., there is no scaling, rotation, or coordinate collapsing.

3.7 Automated transfer ill-posedness detection

In numerous common transfer setups, there is potential for a non-unique valid source value at a target point. In this event, a single value is still selected, usually the one from the process with the lowest rank. The lowest-rank may not be consistent across different partitionings, so this leads to results that are not consistently reproducible in parallel. These indeterminations occur when, for example:

- Multiple points in the origin mesh(es) are equidistant to the target location for nearest-node type transfers. An example of this case is shown in Figure 4.
- Multiple child apps can compute a valid value for a target location because their meshes overlap, for variable shape/user-object evaluation-type transfers.

These issues are detected automatically by the transfers. They can first be discovered when computing the transferred value on the source side. If the child applications presenting the issue are handled by the same process, the multiplicity of valid values is registered when a valid value is requested. If not, then the transfers can also see the indetermination by examining received values from each process for each target point. When multiple values are received for a single target point, the value from the lowest-ranked process is used and, at that point, a potential indetermination is registered.

Finally, at the end of the transfer, global communications are performed to verify that issues discovered when computing the transferred value were actually relevant (e.g., the value selected during that stage was actually the final value for the target point). If another value from another process was selected, then the initially detected indetermination was irrelevant. This global communication is inherently non-scalable and the user is warned if they are deploying it on a case with more than ten processes.

Mesh, block, or boundary restrictions can sometimes be used to alleviate these indeterminations in the origin values. Other times, either re-meshing one of the applications or using the child application positions parameter to create a very small offset can help remove the indetermination. For nearest-node type indeterminations, the recommended fix is to use a local reduction, for example, the average of the nearest-node values using layered averages.

4 Field-to-scalar, scalar-to-field and scalar transfers

In advanced reactor multiphysics core analysis, it is common to have a difference of spatial scale between two physics. Neutronics is typically solved on the global scale for the entire core. On the other hand, fuel performance is typically better solved when considering each disconnected fueled part separately. The fuel performance/thermomechanics of a fuel pin in a light water reactor do not, in most analysis, sufficiently influence the thermomechanics of its neighbor pin to justify a fully-coupled solve including both. Therefore, each pin is best resolved in independent calculations that are pin-scale as opposed to core-scale.

In those cases, while field transfers are still used when both the core-scale and pin-scale meshes are heterogeneous and both resolving the fuel pins, it is often easier to average the local fields and only transfer the average or integrated quantity. It is a simpler approach when the fuel pin meshes do not match. For example, the fuel pin temperature may be averaged over the entire pin before being transferred to the neutronics fuel temperature field. In this case, a scalar quantity is transferred from the distributed fuel performance applications to a field in the global mesh of the neutronics application.

Many coupled cases only require a single scalar number or a few scalars to be passed between applications to couple. This can be, for example, the current insertion height of a control rod that should be shared between the core neutronics calculation and the control rod channel cooling calculation. Domain-overlap coupling of thermal-hydraulics simulations only involve the transfer of a few scalars for the pressure drop, the heat source, and the precursor source for molten salt reactors (Hu et al., 2022).

MOOSE widely supports this transfer mode. Scalars in MOOSE can be Reporters, ScalarVariables or Postprocessors, and all three can be transferred to and from each other with different transfer objects. These transfers are easily distributed as the source values are replicated across all rank using global synchronizations already. This is not a scalable approach but the amount of data is minimal so its impact is not felt in reactor analysis.

5 Advanced reactor multiphysics studies leveraging transfers

5.1 Heat pipe cooled microreactor model

5.1.1 Model description

The HPMR-H₂ is a 2-MW monolithic heat-pipe-cooled microreactor model hosted on the Virtual Test Bed (VTB) repository (Giudicelli et al., 2023). The model, based on the simplified microreactor benchmark problem defined in Ref. (Terlizzi and Labouré, 2023), is composed of 18 hexagonal assemblies arranged into two rings. The tops and bottoms of these 160-cm-high assemblies are surrounded by 20-cm-high axial beryllium reflectors. Each assembly contains 96 fuel pins that are 1 cm in radius, 60 Yttrium Hydride (YH) pins that are 0.975 cm in radius, and 61 1-cm-radius sodium heat pipes drilled into a graphite monolith. The core is surrounded by 12 control drums, with boron carbide employed as the absorbing material. For a simplified mesh, the beryllium radial reflector is hexagonal. The core has a one-twelfth rotational symmetry that allows mitigating the computation burden associated with the simulation. The geometry of the reactor is shown in Figure 5.

The HPMR-H₂ model uses Griffin Prince et al. (2024), BISON Williamson et al. (2021), and Sockeye Hansel et al. (2021) to solve the coupled steady-state equations for heterogeneous neutron transport, heat transfer, sodium flow, and hydrogen redistribution. Figure 6 illustrates the computational workflow for the multiphysics simulation. Within the workflow, Griffin is used to solve the multigroup neutron transport equation, BISON is leveraged to simulate both the solid heat transfer and the hydrogen redistribution in the moderator pins, while Sockeye is used to model the sodium flow in the 101 heat pipes. The left box outlines the preliminary steps, including mesh preparation and cross-section tabulations performed using Serpent (v. 2.1.32). Serpent is used due to the existence of a cross section workflow with Griffin at that time. OpenMC and Shift are alternative options with similar workflows. These meshes, along with the cross sections generated, serve as inputs for Griffin, BISON, and Sockeye. For a comprehensive description of the single-physics models, refer to Refs. (NRIC Virtual Test Bed, 2024; Terlizzi and Labouré, 2023).

5.1.2 Transfers and coupling

As shown in Figure 6, Griffin serves as the parent application for the simulation. It computes the space-dependent multigroup flux distribution by solving the Multigroup Neutron Transport Equation (MNTE). The MNTE is solved with the discontinuous finite element, discrete ordinates (DFEM-SN) approach accelerated with Coarse Mesh Finite Difference (CMFD) acceleration. The power density field, $P_d(\vec{x})$, is then derived

from the neutron flux and transferred to BISON via the MultiAppGeneralFieldShapeEvaluationTransfer. This transfer method is chosen because both the Griffin and BISON models are defined at the full-core level, allowing for the evaluation of the source variable's shape function at the desired transfer location. This transfer was selected because the DFEM-SN solution employs constant monomial variables whereas the BISON model uses first-order Lagrange variables, that are nodal. The use of monomial variables in Griffin was made to decrease the memory requirement in the neutronic simulation.

BISON functions as the daughter application, responsible for computing the temperature distribution in solid materials, $T(\vec{x})$, and the hydrogen-yttrium stoichiometric ratio's spatial distribution in the YH moderator, $c_H(\vec{x})$. The solid temperature at the heat pipe-monolith interface, denoted as T_{ext}^s , is determined from the solid temperature field using layered side averages around each of the 101 heat pipe model instantiations. Each of these instantiations, corresponding to 101 child applications, experiences a different external temperature, which serves as a boundary condition for the fluid solver. The fluid solver then computes the effective heat transfer coefficient h_{eq}^i and secondary temperature $T_{secondary}^i$ for each heat pipe i . These computed values are transferred back to BISON to define a convective boundary condition for the solid. Additionally, the heat flux at the heat pipe condenser, denoted as Q_{cond}^i , is transferred to BISON to ensure energy conservation between the solid and fluid calculations. Checking the energy balance is crucial due to the different mesh between Sockeye and BISON application. Based on calculation performed through post-processors' manipulation, we computed that a discrepancy of 0.04% between the heat balance in the solid and the fluid. This discrepancy was deemed sufficiently small to avoid the use of conservative transfers in the simulation. The latter can however be added to enforce conservation of the integral of the energy, as described in Section 3.4.

After iterating between the BISON and Sockeye solutions through a pseudo-transient approach, the steady-state three-dimensional solid temperature distribution is used to determine the fuel, moderator, and reflector temperature distributions, which serve as parameters for the macroscopic cross sections. These partial temperature fields, denoted as $T_f(\vec{x})$, $T_m(\vec{x})$, and $T_r(\vec{x})$, respectively, are then passed back to Griffin through the MultiAppGeneralFieldNearestNodeTransfer together with the hydrogen-yttrium stoichiometric ratio, $c_H(\vec{x})$. Both block restriction and bounding boxes are leveraged to ensure a well-posed transfer process, as detailed in Section 3.7. The Griffin and thermal solution are iterated until the effective multiplication factor converges within a prescribed absolute tolerance, (e.g., 1 per cent mile in the model available on the VTB).

5.1.3 Results

The 3D solid temperature distribution is presented in Figure 7. It is noticeable that the solid temperature is asymmetric with respect to the axial mid-plane due to the heat pipes penetrating only in the top reflector. As a result, fewer neutrons are reflected back into the upper portion of the reactor's fuel compared to the bottom, leading in turn to an axially asymmetric power generation rate and temperature field. A secondary temperature peak is also evident at the center of the bottom reflector, resulting from the limited heat removal in the central control rod shutdown channel, which is filled with stagnant air. Additionally, the solid temperature distribution at selected axial

planes is shown to illustrate the temperature variations at each pin location. A zoomed-in view of the power density, stoichiometric ratio and temperature at the axial mid-plane is provided in Figure 8.

The temperature of the moderator and the corresponding hydrogen-yttrium stoichiometric ratio are reported as a function of space in Figure 9. It is noticeable that the two quantities are inversely correlated, with the hydrogen redistributing towards colder zones, as expected due to the Soret effect (Terlizzi and Labouré, 2023).

5.2 High temperature prismatic gas reactor (HTGR) model

In a prismatic gas-cooled reactor, the fuel is in the form of compacts housed in graphite blocks. The fission heat is transferred to the gas coolant, generally helium, which flows either in an annular flow channel around the fuel rods, as in the High Temperature Test Reactor (HTTR) design, or within dedicated coolant channel holes in the graphite blocks. System-level approaches can be used to model these reactors, but they need an effective thermal conductivity to accurately represent heat conduction in the graphite block. Three-dimensional approaches can more accurately model the heat transfer in the graphite block. These approaches require the evaluation of convective heat transfer in the flow channels for each flow channel location, which can be achieved by modeling the flow channels as child applications created by the heat transfer simulation. This approach is demonstrated on the High Temperature Test Facility (HTTF).

5.2.1 Model description

The HTTF at Oregon State University is a scaled integral effects experiment for investigating transient behaviors in prismatic nuclear HTGRs. The HTTF uses electrical heated rods as the power source and helium as the coolant. A radial cross section of the HTTF core region is shown in Figure 10.

A 3D mesh of the facility is built and includes the core block with all the coolant flow holes and heater rods holes, the barrel, pressure vessel and the reactor cooling cavity system (RCCS) as shown in Figure 11. A view of the mesh is also shown in Figure 11.

The heater rods are not explicitly modeled and a heat flux corresponding to the heater power as a function of time is prescribed on the heater holes surface. The heat conduction within the core block is modeled using the MOOSE heat transfer module. A convective boundary condition is applied to the coolant channels walls, on the upcomer walls, and on the RCCS pannel walls. The flow in the coolant channels, upcomer and RCCS channel is modeled using the MOOSE Thermal Hydraulics Model (THM) Hansel et al. (2024). Note that this model is published on the Virtual Test Bed Giudicelli et al. (2023). The heat transfer coefficient in the flow channels is calculated using the modified Dittus Boelter correlation as suggested by McEligot McEligot et al. (1965), and the friction factor is calculated using the Churchill correlation.

5.2.2 Transfers and coupling

The coupling strategy is illustrated on Figure 12. The heat conduction application solving for the heat transfer in the ceramic block is the main application. One SubApp is created for

each flow channel within the core block, for the upcomer and for the RCCS water channel.

The temperature on the flow channel walls is averaged by axial layers. Each layers corresponds to an element on the flow channel simulations. This wall temperature is transferred to the channel applications using a `MultiAppGeneralFieldUserObjectTransfer` object. The heat transfer coefficient is calculated within each channel application and is transferred with the bulk fluid temperature to the main application using a `MultiAppGeneralFieldNearestLocationTransfer`. The wall temperature, heat transfer coefficient and bulk fluid temperature are used to calculate the convective heat transfer between the core block and the flow channels. A similar approach is used to model the convective heat transfer in the RCCS and upcomer. Note that the upcomer channel has two convective heat fluxes, one on the barrel side and one on the pressure vessel side. Lastly, the upcomer bulk outlet temperature is transferred to the core flow channel application and is used as an inlet boundary condition. A `MultiAppPostprocessorTransfer` object is used for that purpose.

This model makes use of the nearest-application heuristic to determine which processes should exchange data in the field transfers, see Section 3.5. The default heuristic, based on forming a sphere around each point from the bounding boxes, is rendered inefficient by the elongated bounding boxes of the channel geometries. The nearest-application is a simple concept which suffices here to map regions of the mesh to each channel application.

5.2.3 Results

The transient PG-26 was used as a demonstration. PG26 includes a heatup phase followed by a depressurized conduction cooldown transient. The temperature distribution in the core before the initiation of the depressurization is shown in Figure 13.

A comparison of the experimental data with the simulation prediction for the core block solid temperature at the bottom of the core is shown in Figure 14 for each radial ring. A reasonable agreement is reached, especially considering the uncertainties in the measurements.

5.3 Molten salt reactors

Molten Salt Reactors are Generation-IV advanced reactor design that is actively pursued by several companies and for which several experiments are planned in the US. Notably, the Molten Chloride Reactor Experiment is scheduled to be build inside the LOTUS test bed of the National Reactor Innovation Center at Idaho National Laboratory. Several molten salt reactor models were developed using the MOOSE-based tools and are available on the VTB Giudicelli et al. (2023). The pool-type core multiphysics model share the same coupling strategy, which was summarized in Subsection 2.2 and in Figure 2. We briefly recall here which features of the transfers are leveraged.

The fluid flow and precursor advection simulations typically use the same mesh, because the velocity field is obtained from solving conservation equations and should not be modified by the transfers or the mesh. Advecting precursors with a non-mass conserving velocity field would create oscillations in the precursors' distributions. Because they use the same mesh, the simple copy

transfer can be used for transferring velocities and Rhee Chow coefficients.

For the other transfers, we have to consider the meshes of the neutronics application compared to the flow applications'. The neutronics application mesh is typically coarser than the fluid application, as the pool-type geometry is homogeneous in the fuel region. The neutronics mesh includes both the salt region and the solid region, with the reflectors and the control rods. In the salt region, it suffices to evaluate the fields in the origin application to recreate them in the target application, as explained in the discussion on shape evaluation transfers in Section 3. The integral of the transferred field is preserved using the rescaling described in Subsection 3.4.

5.4 Fusion device example: a tokamak demonstration model

5.4.1 Model description

This demonstration model consists of a 90-degree sector of a Tokamak. For simplification, the model consists of five domains only: a Tungsten first wall, a Homogenized multiplier layer, a Homogenized breeder layer, a Vacuum vessel, and a divertor. Figure 15 shows the model mesh with different colors for each subdomain. Geometry parameters were taken from (Shimwell et al., 2021). Both the first wall and divertor were modeled as pure Tungsten. The multiplier is a homogenized layer of 65% beryllium, and 35% Helium. The breeder is a homogenized layer of 10% stainless steel alloy, 37% beryllium, 15% Li_4SiO_4 , and 38% Helium. Finally, the vacuum vessel was modeled as a stainless steel alloy. Natural abundance was assumed for all elements. These compositions were considered for this demonstration model as placeholders where we focused on the modeling workflow and multiphysics solvers transfers.

The model leverages Cardinal (Novak et al., 2022), a wrapping of the Monte Carlo particle transport code OpenMC (Romano et al., 2015) and the spectral element Computational Fluid Dynamics (CFD) code NekRS (Fischer et al., 2022) within the MOOSE framework to allow multiphysics feedback to the particle transport model. A tutorial based on this tokamak model can be found on the Cardinal website (DAGMC Tokamak, 2025).

The tokamak model was first built using capabilities in Paramak (Shimwell et al., 2021), a python package that allows rapid production of 3D CAD models of fusion reactors. Then, the CAD model was imported into Coreform Cubit Blacker et al. (1994) to assign materials, subdomain names, and boundary conditions. Two meshes are then generated for this model: a mesh to solve the heat conduction problem using the MOOSE heat transfer module, and a triangulated surface mesh for transporting particles within OpenMC.

For simplicity, all regions were modeled as purely conducting with no advection considered. This is a significant simplification, and therefore the results obtained should not be taken as representative of a realistic device. The boundary conditions applied to the heat conduction

model are also simplified, as the outer boundary temperature is set to a Dirichlet condition of 800 K. All other boundaries are assumed insulated. To approximate some cooling in the breeder and divertor, a uniform volumetric heat sink is applied in these regions.

OpenMC was used for neutron transport by running a series of fixed source calculations on this model. OpenMC and MOOSE are coupled via fixed point iterations through Cardinal. The neutron source was defined as an isotropic ring cylindrical source distributed uniformly inside the cavity of the tokamak with a uniform neutron energy of 14.08 MeV. The cross-section library ENDF/B-VII.1 (Chadwick et al., 2011) was used. The number of particles sampled is 1×10^7 particles per batch. The OpenMC trigger capability (Shen et al., 2015) was used to finish simulations based on tally uncertainties by setting a trigger threshold on the heating tally to finish the simulation at a specified tally relative uncertainty level.

5.4.2 Transfers and coupling scheme

An unstructured mesh-based heat source is calculated in OpenMC using the MOOSE mesh as a tally mesh. The heat source is then transferred to MOOSE via Cardinal to be applied as a volumetric heating term. Heat conduction is solved using the MOOSE heat transfer module and temperatures are calculated. The resulting temperatures are then sent back to OpenMC to account for temperature feedback where one temperature is applied to each OpenMC cell, updating the nuclear interaction cross sections used. Figure 16 illustrates the multiphysics coupling scheme, where temperature feedback is utilized in the OpenMC model to calculate the different neutronic quantities of interest.

The heat transfer MOOSE module functions as the parent application for the simulation, where it computes temperatures and sends them via the MultiAppCopyTransfer to the Cardinal application which uses an identical mesh for simplicity. The Cardinal daughter application runs at the end of each fixed point iteration to calculate the neutronics quantities of interest (heating, and tritium production) and the heating results are transferred back to the parent application using the MultiAppCopyTransfer.

With recent developments, Cardinal now has the capability to update displacements on OpenMC's surface mesh geometry from MOOSE's solid mechanics solver (Eltawila et al., 2025). This allows for utilizing thermomechanical feedback arising from thermal expansion, which will be added as an extension to this tokamak model in future work.

5.4.3 Results

The multiphysics results correspond to a power of 6.1 MW for a 90° sector. Figure 17 shows the simulation results for the neutronics quantities of interest, where we show the solid temperatures computed by MOOSE, on the last fixed point iteration, as well as the tritium production results. Note that these results are not necessarily intended to replicate a realistic tokamak, due to the highly simplified geometry, neutron source definition, and lack of fluid flow cooling. However we can still use this model to estimate the different feedback mechanisms' importance and demonstrate the workflow.

5.5 High fidelity conjugate heat transfer simulations

In the previous section, Cardinal's use for coupling OpenMC to Multiphysics Object Oriented Simulation Environment (MOOSE) was demonstrated. Cardinal also contains a wrapping of the NekRS spectral element CFD code for a variety of coupling mechanisms, including (i) conjugate heat transfer, (ii) volumetric neutronics feedback, (iii) fluid-structure interaction, (iv) systems-CFD coupling, and (v) forward Uncertainty Quantification (UQ).

5.5.1 Model description

To demonstrate MOOSE's transfers for conjugate heat transfer applications, we illustrate a simple 7-bin wall-resolved Sodium Fast Reactor (SFR) bundle, where the fluid region is solved with NekRS's spectral element methods and the solid is solved with MOOSE's finite element heat transfer module. A tutorial based on this model is available on the [Cardinal \(2025\)](#).

The geometry is a shorter, seven-pin version of the fuel bundles in the Advanced Breeder Reactor (ABR) J. Cahalan and T. Fanning and M. Farmer and C. Grandy and E. Jini and T. Kim and R. Kellogg and L. Krajtl and S. Lomperski and A. Moisseytsev and Y. Momozaki and Y. Park and C. Reed and F. Salev and R. Seidensticker and J. Sienicki and Y. Tang and C. Tzanos and T. Wei and W. Yang and Y. [Cahalan et al. \(2007\)](#). Pin dimensions match those of the full-size reactor, and the bundle power, inlet temperature, and flow rate are scaled so as to obtain laminar flow conditions. While this is not representative of actual turbulent flow SFR conditions, the purpose of the present discussion is to illustrate two different types of heat flux transfers facilitated by MOOSE.

The MOOSE heat transfer module is used to solve for energy conservation in the solid. The outer surface of the duct and the tops and bottoms of the pins and ducts are insulated. Cardinal supports two options for the boundary condition to be applied to the solid-fluid interface, and which necessitate different transfers in MOOSE. These two options are referred to as (i) "flux-flux" coupling, whereby NekRS passes a convective heat flux to MOOSE, $h(T - T_{\infty})$; and (ii) "flux-temperature" coupling, whereby NekRS passes the wall temperature to MOOSE. The gap region between the pellet and the cladding is unmeshed, and a quadrature-based thermal contact model is applied based on the sum of thermal conduction and thermal radiation (across a transparent medium).

NekRS is used to solve the incompressible Navier-Stokes equations. At the outlet, a zero-gage pressure is imposed and an outflow condition is applied for the energy conservation equation. On all solid-fluid interfaces, the velocity is set to the no-slip condition and a heat flux is imposed in the energy conservation equation computed by MOOSE.

The meshes for the solid (MOOSE) and fluid (NekRS) domains are shown in [Figures 18, 19](#). Note that the two applications do not employ a conformal mesh at the fluid-solid interface, which will necessitate nearest node type transfers.

5.5.2 Transfers and coupling scheme

When using conjugate heat transfer coupling, MOOSE passes a conductive heat flux, $-k\vec{V}T \cdot \hat{n}$, along the coupled boundary to NekRS. NekRS then either returns the wall temperature as a Dirichlet condition, or a convective heat flux, $h(T - T_{\infty})$. While

the wall temperature option, referred to as "flux-temperature" coupling, is fully 3-D, it sometimes requires very frequent data transfers between NekRS and MOOSE to retain stability of the Picard operator splitting. Alternatively, the flux-flux transfer may offer better stability, especially for problems with a large difference in thermal conductivity between the fluid and solid, or for high Prandtl number fluids. These two transfer options are summarized in [Figure 20](#).

The MOOSE transfers used to facilitate these transfers are now described. In order to pass the conductive heat flux from MOOSE to NekRS, a `GeneralFieldMultiAppNearestLocationTransfer` is employed to pass between nearest nodes on the duct inner surface and clad outer surface sidesets. A nearest node transfer is needed because the use of different meshes in MOOSE and NekRS will not necessarily lead to node overlap. Evaluating the shape function of variables outside of their domain is allowed. In order to account for different side quadrature rules and element refinement, a `MultiAppPostprocessorTransfer` is also used to pass the side-integrated heat flux (power) from MOOSE to NekRS for an internal normalization step to ensure energy conservation.

The transfer from NekRS to MOOSE then differ depending on the field(s) passed from NekRS to MOOSE. For the flux-temperature coupling option, a `MultiAppGeneralFieldNearestLocationTransfer` is used to pass NekRS's wall temperature to MOOSE. Cardinal internally interpolates between NekRS's Gauss-Lobatto-Legendre (GLL) node points to a MOOSE mesh. For the flux-flux coupling option, `UserObjects` in Cardinal are used to compute h and T_{∞} on the NekRS mesh. T_{∞} represents the bulk-averaged fluid temperature, which is computed directly on the NekRS mesh by subdividing the mesh into sub-volumes. For the present example, each GLL point in the NekRS mesh is mapped to a single subchannel from a typical hexagonal subchannel discretization. For example, [Figure 21](#) shows the indices of each hexagonal subchannel, over which volume averages are used to compute T_{∞} in sub-volume i ,

$$T_{\infty,i} = \frac{\int_{\Omega_i} T dV}{\int_{\Omega_i} dV} \quad (1)$$

where Ω_i are the sets of GLL points that would fall into subchannel bin i .

To compute the heat transfer coefficient, side-averaged quantities are also required. The index i was used to represent the colored volumes in [Figure 21](#). Using j to indicate the fluid-solid sideset bordering volume i , the wall-average temperature and heat flux are computed as

$$T_j = \frac{\int_{\partial\Omega_j} T dS}{\int_{\partial\Omega_j} dS} \quad (2)$$

$$q_j'' = \frac{\int_{\partial\Omega_j} q'' dS}{\int_{\partial\Omega_j} dS} \quad (3)$$

where $\partial\Omega_j$ is the sideset pertaining to subchannel i . The heat transfer coefficient is then computed as

$$h_j = \frac{q_j''}{T_j - T_{\infty,i}} \quad (4)$$

UserObjects in Cardinal are used to compute the three quantities on the right-hand side of [Equation 4](#). h is then passed to the MOOSE heat conduction solver as an auxiliary variable using `MultiAppGeneralFieldNearestNodeTransfer`. $T_{\infty,i}$ is passed directly as a user object, using `MultiAppGeneralFieldUserObjectTransfer` which then locally evaluates the value of h at the nodes on the MOOSE heat conduction mesh. h and $T_{\infty,i}$ are then applied as a Robin-type flux condition using MOOSE's `CoupledConvectiveHeatFluxBC`. In other words, using the center pin as an example, the heat conduction model is applied six different values of h and T_{∞} around the pin circumference to capture the 3-D nature of the heat transfer.

5.5.3 Numerical results and transferred fields

This seven-pin bundle is run with both the flux-temperature and flux-flux coupling options. The fluid and solid temperatures computed with the flux-temperature coupling option are shown in [Figure 22](#). This conjugate heat transfer mode utilizes two `MultiAppGeneralFieldNearestNodeTransfers`.

[Figure 23](#) shows the solid temperature, when using either the flux-temperature or flux-flux coupling option. For the flux-flux option, the subdomain is divided into 20 axial layers and with a hexagonal subchannel discretization in the $x-y$ plane (18 subchannels). Both options are nearly equivalent, with minor differences because the convective heat flux is not a truly 3-D transfer due to the averaging options required to define h .

6 Conclusion

The coupling of different physics solvers is required for advanced nuclear systems analysis due to the need for specialized solvers to tackle each physics separately. In the MOOSE ecosystem, this is achieved by leveraging the `MultiApps` and `Transfers` systems. Over time, those systems have evolved to include numerous refinements such as asynchronous communications, general projection capabilities, sibling transfers, and automated ill-posedness detection. These capabilities support analysis of a wide range of advanced nuclear reactors, several examples of which were shown in [Section 5](#). The scope for planned future developments will focus on concurrency of execution for GPU-CPU applications exchanging data through `Transfers`, following.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

GG: Investigation, Methodology, Software, Writing – original draft, Writing – review and editing. FK: Investigation, Software, Writing – original draft. RS: Conceptualization, Investigation, Software, Writing – original draft. LH: Conceptualization, Investigation, Writing – original draft. DG: Conceptualization, Funding acquisition, Software, Writing – original draft. AL: Funding acquisition, Investigation, Project administration, Software, Writing – original draft. ZP: Conceptualization, Software, Writing – original draft. LC: Investigation, Project administration, Writing – original draft. ST: Investigation, Visualization, Writing – original draft. ME: Investigation, Writing – original draft. AN: Investigation, Writing – original draft.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. This research was partially supported by the Nuclear Energy Advanced Modeling and Simulation program for Modeling and Simulation of Nuclear Reactors under US Department of Energy contract no. DE-AC05-00OR22725.

Acknowledgments

The authors would like to acknowledge the editing contributions of Sarah E. Roberts at Idaho National Laboratory, and the contributions of Quentin Faure to [Figure 5](#). Stefano Terlizzi wants to thank Dr. Vincent Laboure' for his contribution to the preparation of the heat pipe cooled microreactor model. Thomas Freymann is acknowledged for his contributions to the HTTR model. This research made use of the resources of the High-Performance Computing Center at Idaho National Laboratory.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Licenses and permissions

This manuscript was authored by Battelle Energy Alliance, LLC under contract no. DE-AC07-05ID14517 with the U.S. Department of Energy. The U.S. Government retains and the publisher, by

accepting the article for publication, acknowledges that the United States Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes.

References

Argonne National Laboratory (2024). Cardinal website. Available online at: <https://cardinal.cels.anl.gov/>.

Blacker, T. D., Bohnhoff, W. J., and Edwards, T. L. (1994). CUBIT mesh generation environment. Volume 1: users manual. *Tech. Rep.* doi:10.2172/10176386

Calahan, J., Fanning, T., Farmer, M., Grandy, C., Jini, E., Kim, T., et al. (2007). Advanced burner reactor 1000 MWth reference concept. *Tech. Rep.* doi:10.2172/1349893

Cardinal (2025). Conjugate heat transfer for laminar pin bundle flow. Available online at: <https://cardinal.cels.anl.gov/tutorials/cht2.html>.

Chadwick, M. B., Herman, M., Obložinský, P., Dunn, M. E., Danon, Y., Kahler, A., et al. (2011). Endf/b-vii. 1 nuclear data for science and technology: cross sections, covariances, fission product yields and decay data. *Nucl. data sheets* 112, 2887–2996. doi:10.1016/j.nds.2011.11.002

DAGMC Tokamak (2025). DAGMC tokamak. Available online at: <https://cardinal.cels.anl.gov/tutorials/tokamak.html>.

Eltawila, M., Novak, A., Simon, P., and Giudicelli, G. (2025). Monte Carlo multiphysics on moving meshes: initial verification and case studies. *Math. and Comput. Available online at: https://www.ams.org/pubs/proceedings/article-58168/*

Fischer, P., Kerkemeier, S., Min, M., Lan, Y., Phillips, M., Rathnayake, T., et al. (2022). NekRS, a GPU-accelerated spectral element Navier-Stokes solver. *Parallel Comput.* 114, 102982. doi:10.1016/j.parco.2022.102982

Forum, G. I. (2024). Gen-IV advanced nuclear systems. Available online at: https://www.gen-4.org/gif/jcms/c_59461/generation-iv-systems.

Gaston, D. R. (2019). *Parallel, asynchronous ray-tracing for scalable method of characteristics neutron transport on unstructured mesh*. (Ph.D. thesis). Massachusetts Institute of Technology. Available online at: <https://hdl.handle.net/1721.1/129911>

Gaston, D. R., Forget, B., Smith, K. S., Harbour, L. H., Ridley, G. K., and Giudicelli, G. (2021). Method of characteristics for 3d, full-core neutron transport on unstructured mesh. *Nucl. Technol.* 207, 931–953. doi:10.1080/00295450.2021.1871995

Gaston, D. R., Permann, C. J., Peterson, J. W., Slaughter, A. E., Andrs, D., Wang, Y., et al. (2015). Physics-based multiscale coupling for full core nuclear reactor simulation. *Ann. Nucl. Energy* 84, 45–54. doi:10.1016/j.anucene.2014.09.060

Giudicelli, G., Kong, F., Stogner, R., Harbour, L., Gaston, D., Terlizzi, S., et al. (2024a). Data transfers for full core heterogeneous reactor high-fidelity multiphysics studies. *Jt. Int. Conf. Supercomput. Nucl. Appl.* 302, 05006. doi:10.1051/epjconf/202430205006

Giudicelli, G., Lindsay, A., Harbour, L., Icenhour, C., Li, M., Hansel, J. E., et al. (2024b). 3.0 - moose: enabling massively parallel multiphysics simulations. *SoftwareX* 26, 101690. doi:10.1016/j.softx.2024.101690

Giudicelli, G. L., Abou-Jaoude, A., Novak, A. J., Abdelhameed, A., Balestra, P., Charlot, L., et al. (2023). The virtual test bed (vtb) repository: a library of reference reactor models using neams tools. *Nucl. Sci. Eng.* 197, 2217–2233. doi:10.1080/00295639.2022.2142440

Hansel, J., Andrs, D., Charlot, L., and Giudicelli, G. (2024). The MOOSE thermal hydraulics module. *J. Open Source Softw.* 9, 6146. doi:10.21105/joss.06146

Hansel, J. E., Berry, R. A., Andrs, D., Kunick, M. S., and Martineau, R. C. (2021). Sockeye: a one-dimensional, two-phase, compressible flow heat pipe application. *Nucl. Technol.* 207, 1096–1117. doi:10.1080/00295450.2020.1861879

Hu, R., Fang, J., Nunez, D., Tano, M., Giudicelli, G., and Salko, R. (2022). Development of integrated thermal fluids modeling capability for MSRs. *Tech. Rep.* doi:10.2172/1889653

Hu, R., Zou, L., Hu, G., Nunez, D., Mui, T., and Fei, T. (2021). SAM theory manual. *Tech. Rep.* doi:10.2172/1781819

Kirk, B. S., Peterson, J. W., Stogner, R. H., and Carey, G. F. (2006). libMesh: a C++ library for parallel adaptive mesh refinement/coarsening simulations. *Eng. Comput.* 22, 237–254. doi:10.1007/s00366-006-0049-3

Lindsay, A. D., Gaston, D. R., Permann, C. J., Miller, J. M., Andrs, D., Slaughter, A. E., et al. (2022). 2.0 - MOOSE: enabling massively parallel multiphysics simulation. *SoftwareX* 20, 101202. doi:10.1016/j.softx.2022.101202

McEligot, D. M., Magee, P. M., and Leppert, G. (1965). Effect of large temperature gradients on convective heat transfer: the downstream region. *J. Heat Transf.* 87, 67–73. doi:10.1115/1.3689054

Novak, A., Andrs, D., Shriwise, P., Fang, J., Yuan, H., Shaver, D., et al. (2022). Coupled Monte Carlo and thermal-fluid modeling of high temperature gas reactors using cardinal. *Ann. Nucl. Energy* 177, 109310. doi:10.1016/j.anucene.2022.109310

NRIC Virtual Test Bed (2024). Heat pipe microreactor with hydrogen redistribution (hpmr-h2). Available online at: https://mooseframework.inl.gov/virtual_test_bed/microreactors/hpmr_h2/index.html.

Pawlowski, R. P., Slattery, S., and Wilson, P. (2013). “The data transfer kit: a geometric rendezvous-based tool for multiphysics data transfer,” in *International conference on mathematics and computation for nuclear science and engineering* (Albuquerque, NM, and Livermore, CA (United States): Sandia National Laboratories SNL).

Prince, Z. M., Hanophy, J. T., Labouré, V. M., Wang, Y., Harbour, L. H., and Choi, N. (2024). Neutron transport methods for multiphysics heterogeneous reactor core simulation in griffin. *Ann. Nucl. Energy* 200, 110365. doi:10.1016/j.anucene.2024.110365

Romano, P., Horelik, N., Herman, B., Nelson, A., Forget, B., and Smith, K. (2015). OpenMC: a state-of-the-art Monte Carlo code for research and development. *Ann. Nucl. Energy* 82, 90–97. doi:10.1016/j.anucene.2014.07.048

Schoder, S., and Roppert, K. (2025). opencfs: open source finite element software for coupled field simulation – part acoustics. *arXiv*. Available online at: <https://arxiv.org/abs/2207.04443>

Schunert, S., Giudicelli, G., Lindsay, A., Balestra, P., Harper, S., Freile, R., et al. (2021). “Deployment of the finite volume method in Pronghorn for gas- and salt-cooled pebble-bed reactors,” in *External report INL/EXT-21-63189*. Idaho National Laboratory.

Shen, Q., Boyd, W., Forget, B., and Smith, K. (2015). Tally precision triggers for the OpenMC Monte Carlo code. *Trans. Am. Nucl. Soc.*

Shimwell, J., Billingsley, J., Delaporte-Mathurin, R., Morbey, D., Bluteau, M., Shriwise, P., et al. (2021). The paramak: automated parametric geometry construction for fusion reactor designs. *F1000Research* 10, 27. doi:10.12688/f1000research.28224.1

Stauff, N. E., Mo, K., Cao, Y., Thomas, J. W., Miao, Y., Lee, C., et al. (2021). “Preliminary applications of neams codes for multiphysics modeling of a heat pipe microreactor,” in *Proceedings of the American nuclear society annual 2021 meeting*.

Terlizzi, S., and Labouré, V. (2023). Asymptotic hydrogen redistribution analysis in yttrium-hydride-moderated heat-pipe-cooled microreactors using direwolf. *Ann. Nucl. Energy* 186, 109735. doi:10.1016/j.anucene.2023.109735

Wang, Y., Schunert, S., Ortenzi, J., Gleicher, F. N., Laboure, V. M., Baker, B. A., et al. (2017). “Demonstration of mammoth strongly-coupled multiphysics simulation with the godiva benchmark problem,” in *Ans M&C 2017* (Idaho Falls, ID United States: Idaho National Lab. INL).

Williamson, R. L., Hales, J. D., Novascone, S. R., Pastore, G., Gamble, K. A., Spencer, B. W., et al. (2021). Bison: a flexible code for advanced simulation of the performance of multiple nuclear fuel forms. *Nucl. Technol.* 207, 954–980. doi:10.1080/00295450.2020.1836940

Woods, B. (2019). Instrumentation plan for the OSU high temperature test facility, revision 4. *Tech. Rep. Or. State Univ.* Available online at: <https://www.osti.gov/biblio/1599628>



OPEN ACCESS

EDITED BY

Deokjung Lee,
Ulsan National Institute of Science and
Technology, Republic of Korea

REVIEWED BY

Hyoung Kyu Cho,
Seoul National University, Republic of Korea
Yunje Cho,
Korea Atomic Energy Research Institute (KAERI),
Republic of Korea

*CORRESPONDENCE

Kostadin Ivanov,
✉ knivanov@ncsu.edu

RECEIVED 16 March 2025

ACCEPTED 23 May 2025

PUBLISHED 01 August 2025

CITATION

Kutlu Y, Spasov I, Mitkov S, Rouxelin P, Abarca A, Kolev N, Avramova M and Ivanov K (2025) CTF development, verification, and validation for VVER core thermal-hydraulics and multi-physics modeling and simulation.

Front. Nucl. Eng. 4:1594698.

doi: 10.3389/fnuen.2025.1594698

COPYRIGHT

© 2025 Kutlu, Spasov, Mitkov, Rouxelin, Abarca, Kolev, Avramova and Ivanov. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

CTF development, verification, and validation for VVER core thermal-hydraulics and multi-physics modeling and simulation

Yesim Kutlu¹, Ivan Spasov², Svetlomir Mitkov², Pascal Rouxelin¹, Agustin Abarca¹, Nikola Kolev², Maria Avramova¹ and Kostadin Ivanov^{1,2*}

¹Department of Nuclear Engineering, North Carolina State University, Raleigh, NC, United States

²Institute for Nuclear Research and Energy, Bulgarian Academy of Sciences, Sofia, Bulgaria

The advanced thermal-hydraulics sub-channel tool CTF has been in the process of continuous development and improvement by Oak Ridge National Laboratory (ORNL) and North Carolina State University (NCSU). In recent years, there has been considerable progress in code development, including new functionalities, application-specific correlations, various multi-physics applications, built-in pre- and post-processors, improved solvers, parallelization, and extensive testing. VVER applications are part of these activities. NCSU has been cooperating with the Institute for Nuclear Research and Energy (INRNE) on CTF development, verification, and validation for VVER core modeling and simulation. This article presents an overview of these CTF studies for VVER applications. Several test cases are considered, which include pure thermal-hydraulic problems as well as multi-physics simulations at the nodal and pin level. On the single physics side, thermal-hydraulic CTF solutions have been compared against measured data for rod bundle, fuel assembly, and full core, as well as code-to-code vs. FLICA4 solutions. CTF was tested in the simulation of the TVSA-5T VVER mini-assembly experiments and in the full-core steady-state calculation for the ongoing OECD/NEA Rostov-2 benchmark. For the TVSA-5T calculations, CTF was coupled with the uncertainty analysis tool Dakota and utilized to propagate uncertainties of input and boundary conditions to output quantities of interest for thermal-hydraulic parameter investigations. The CTF results and measured data obtained from this experimental setup were compared for validation. To produce reliable pin-resolved reference solutions for multi-physics model testing the high-fidelity continuous energy Monte Carlo-based neutron transport codes MCNP6.2 and Serpent 2.2.0 were separately coupled with the CTF sub-channel code. Coupled models of a VVER-1000 fuel assembly were tested in comparisons between MCNP/CTF and Serpent/CTF results. Coarse-mesh multi-physics solutions for a full core have been obtained with the coupled COBAYA/CTF, COBAYA/FLICA4, and PARCS/CTF codes. These solutions have been compared against steady-state plant data and code-to-code for transients. High-fidelity pin-resolved solutions with SERPENT/CTF serve as reference solutions in a steady state. The outcomes from the various studies of single-physics and multi-physics cases used for CTF verification and validation

met the initial expectations both qualitatively and quantitatively. The results of the numerical verification and experimental validation are in good agreement with the corresponding reference data.

KEYWORDS

VVER, CTF, sub-channel, multi-physics, uncertainty

1 Introduction

In ensuring the safety of a nuclear reactor design, achieving reliable assessments for both steady-state and transient conditions is a priority. These assessments can be carried out by either experimental work or computational simulations. The high cost involved in full-scale experiments has led to continuous developments of best-estimate codes using the best available modeling and simulation methods that are economically more attractive for many applications.

These so-called “best-estimate” codes have taken the place of traditionally over-conservative approaches by utilizing minimum conservatism in input data and minimizing conservative modeling assumptions. These codes become beneficial when remaining inaccuracies in modeling or numerical errors in the results of calculations are appropriately identified and quantified. Validation and verification processes supplemented by uncertainty estimation are key to providing evidence that a developed code system is suitable for specific use.

CTF is an advanced best estimate code (Salko and Avramova, 2016) that has been continuously developed and maintained by the Reactor Dynamics and Fuel Modelling Group (RDFMG) at North Carolina State University (NCSU) in cooperation with Oak Ridge National Laboratory (ORNL) for thermal-hydraulic (TH) simulations of nuclear reactor vessels with modeling detailed flow and temperature distribution at the sub-channel level. CTF has been verified, validated, and extensively used for steady-state and transient simulations for light-water reactors (LWRs). The United States Department of Energy (U.S. DOE) initiated the Consortium for Advanced Simulation of Light Water Reactors (CASL) program in 2010, and CTF was adopted in the program to account for core thermal-hydraulic modeling and simulation as a complementary part in the Virtual Environment for Reactor Applications Core Simulator (VERA-CS). Some improvements were made to CTF during the CASL program, including, for example, the development of a PWR preprocessor utility for fast input generation, visualization capabilities by generating VTK output files for square lattices, fully distributed memory parallelization, and an automated system testing with built-in test units (Kucukboyaci et al., 2015).

The capabilities of CTF were investigated by (Sung et al., 2015) for a main steam line break (MSLB) accident scenario in a pressurized water reactor (PWR), following the software optimizations and improvements. The departure from nucleate boiling (DNB) predictions obtained for a bundle model was compared to experimental data from the PWR Sub-channel and Bundle Test (PSBT) provided by the Nuclear Power Engineering Corporation (NUPEC) of Japan (Rubin, et al., 2012). The CTF predictions were also compared to the results computed by another

thermal-hydraulics code, VIPRE-W, for the same bundle model. The results show that the CTF followed the same trend as experimental data and the VIPRE prediction. Additionally, a Westinghouse 3-loop PWR full core model was developed at a sub-channel level in the same study to evaluate the reactor core power distribution during an MSLB accident assumed to have occurred in one of the three loops. For this case, CTF predicted coolant and fuel conditions like the predictions made by the Westinghouse sub-channel code THINC IV.

CTF has been used in many studies either as a single physics (core thermal-hydraulics) assessment of core designs or as part of a multi-physics calculation used to provide thermal-hydraulic feedback. Some of these applications included code-to-code comparisons for the void distribution exercises found in OECD/NRC PSBT benchmark (Avramova et al., 2021) and in OECD/NRC Boiling Water Reactor Full-Size Fine-Mesh Bundle Test (BFTT) benchmark. The works were carried out for square lattice core geometries representative of PWRs and boiling water reactors (BWRs).

In cooperation with the Institute for Nuclear Research and Energy (INRNE), Bulgarian Academy of Sciences, CTF has also been improved for triangular/hexagonal lattice calculations by implementing critical heat flux correlations and tables, and by the capability to solve fluid conditions considering triangular sub-channel or hexagonal channel descriptions. In (Papadionysiou Marianna et al., 2022) CTF was coupled with the neutron transport code nTRACER for high-fidelity VVER analysis, where CTF was embedded as a library in the nTRACER. CTF provided the temperature profiles of coolant and fuel, along with coolant density to nTRACER in exchange for power profiles calculated by nTRACER. CTF has recently received additional modifications for hexagonal geometry applications for sodium fast reactors (SFRs) by (Takasugi, 2023; Aly Ahmed Mohamed Nabil Hassanein Multi-scale Multi-Physics Modeling of, 2022).

This paper summarizes the cooperation activities between NCSU and the Institute for Nuclear Research and Energy (INRNE), Bulgarian Academy of Sciences, on model improvements, verification and validation, and uncertainty analysis of CTF for triangular/hexagonal lattice calculations in VVER applications.

The authors applied CTF to the VVER geometry and conducted comparative verification and validation (V&V) analyses with measured data and other codes. New multi-physics couplings and test results are reported. Code improvements include enhanced functionalities for the VVER geometry, enhanced computational efficiency, and a VVER-specific closure law based on tables of measured rod bundle data. The CTF functionalities were enhanced by a pre-processor and a post-processor for hexagonal assemblies and triangular lattices, embedded in the source code. CTF parallelization was tested for VVER cores. A new option in DNBR

calculation based on a VVER lattice-specific experimental table has been implemented in CTF.

Since the research content of this paper is extensive, here we provide an explanation of why each section is necessary and how they relate to each other.

Section 2 described the theory and models behind CTF and VVER-related modifications introduced, such as implementing critical heat flux correlations and tables, and the capability to solve fluid conditions considering triangular sub-channel or hexagonal channel descriptions.

Section 3 focuses on CTF standalone (single physics) applications, including verification, validation, and uncertainty and sensitivity analysis.

Modifications continue to be made in CTF relating to properties and correlations for hexagonal light water geometries that need to be verified and validated. Under this scope, this paper focuses on demonstrating CTF's capabilities on hexagonal VVER core models using two versions of CTF - CTF3.5 and CTF4.3 of the code. Contrary to CTF3.5, CTF4.3 includes some additional features such as pre- and post-processing of hexagonal geometry and an improved coupling interface. However, both versions use the same solution methods and include the same modifications for VVER models, such as a new option for DNBR calculation. This added option is based on a VVER lattice-specific experimental table. Thus, using either version does not affect the obtained results. From this point onward, "CTF" will be used to refer to CTF3.5 and/or CTF4.3 for the remainder of the paper.

The typical and specific features of the VVER-1000 hexagonal fuel assemblies of TVSM, TVS-2M, and TVSA types were used for testing the existing and new CTF functionalities for VVER hexagonal fuel assemblies modeling and simulation.

The studies shown herein are mainly divided into two sections: studies done with CTF as a single physics thermal-hydraulics tool and the studies done with CTF in coupled codes developed with various neutronics codes to be used in multi-physics calculations for steady-state and transient calculations.

2 CTF code package

As mentioned above, the thermal-hydraulic best estimate code CTF, developed and maintained by the RDFMG at NCSU in cooperation with ORNL, is a state-of-the-art modernized version of COBRA-TF (Coolant Boiling in Rod Arrays-Two Fluid). CTF solves conservation equations for mass, energy, and momentum for liquid and vapor phases, known as the two-fluid model, for operational and transient conditions, and is capable of modeling both channel-centered and fuel-centered sub-channel models.

A sub-channel is a vertical stack of mesh cells that represents a fluid volume of any kind of vertical flow path, typically defining flow within reactor cores in one of two ways: a fuel-centered sub-channel surrounded by coolant, or a coolant-centered sub-channel with fuel rods located along the boundary. Lateral exchanges between the adjacent sub-channels and axial flow through each sub-channel are considered by the mathematical model. The conservation equations are given in Equations 1–3 for mass, momentum, and energy, respectively, and are solved considering the transition between fluid phases. In Equation 1, the subscript k denotes the fluid

phase that can be liquid, vapor, or an entrained droplet. The first term is the change of mass with time, while the second term represents the advection term for the flow of mass into or out of the control volume. The L_k is the descriptor of mass transfer into or out of the phase denoted by k that can occur by evaporation, condensation, entrainment, or detrainment.

$$\frac{\partial}{\partial t} (\alpha_k \rho_k) + \nabla \cdot (\alpha_k \rho_k \vec{V}_k) = L_k + M_e^T \quad (1)$$

In Equation 2, the left-hand side represents the change of volume momentum with time and advection of momentum in each direction. The right-hand side, in respective order, represents gravitational force, pressure force, viscous and turbulent shear stress, momentum source, the drag force at the interface between different phases, and momentum transfer caused by turbulent mixing. While presented in the equation, it should be noted that turbulent shear stress is not implemented in CTF. The effects of turbulent mixing are accounted for by turbulent diffusion approximation.

$$\begin{aligned} \frac{\partial}{\partial t} (\alpha_k \rho_k \vec{V}_k) + \frac{\partial}{\partial x} (\alpha_k \rho_k u_k \vec{V}_k) + \frac{\partial}{\partial y} (\alpha_k \rho_k u_k \vec{V}_k) + \frac{\partial}{\partial z} (\alpha_k \rho_k u_k \vec{V}_k) \\ = \alpha_k \rho_k \vec{g} - \alpha_k \nabla P + \nabla \cdot [\alpha_k (\tau_k^{ij} + T_k^{ij})] + \vec{M}_k^L + \vec{M}_k^d + \vec{M}_k^T \end{aligned} \quad (2)$$

The change of energy with time of phase k and the advection of phase k energy is given, in respective order, on the left-hand side of Equation 3. Heat conduction is denoted by \vec{Q}_k and energy change by both turbulent mixing and by void drift are represented by the term $\vec{q}_k T$. Heat conduction in the fluid is neglected in CTF for light water reactors, assuming only a negligibly small amount of heat transfer occurs by conduction in water. The last three terms correspond to the terms of energy transfer due to phase change, volumetric wall heat transfer and pressure work.

$$\begin{aligned} \frac{\partial}{\partial t} (\alpha_k \rho_k \vec{h}_k) + \nabla \cdot (\alpha_k \rho_k h_k \vec{V}_k) = -\nabla \cdot [\alpha_k (\vec{Q}_k + \vec{q}_k T)] + \Gamma_k h_k^i \\ + q_{wk}''' + \alpha_k \frac{\partial P}{\partial t} \end{aligned} \quad (3)$$

The more detailed and expanded version of generalized conservation equations in Cartesian forms are written out explicitly in the CTF Theory Manual (Salko and Avramova, 2016). According to geometry modelled, such as modeling hexagonal or triangular lattices, Equations 1–3 are solved accordingly based on inputs by the user.

CTF has been extensively verified and validated for applications of PWRs and BWRs. More recently, it has been expanded for hexagonal geometry LWRs, with modifications and validation of CTF to analyze VVERs implemented within the last decade. For instance, a recent addition is implemented in CTF based on a VVER lattice-specific experimental table for DNBR calculations. All these improvements have included the development of models, enhanced computational efficiency, as well as enhanced software quality and associated quality assurance procedures and documentation.

Previously, the visualization ability of CTF results for PWR core designs was integrated into the VERA-CS environment with the visualization toolkit (VTK) output option in CTF. Similarly, some recent modifications have been made to CTF to generate VTK

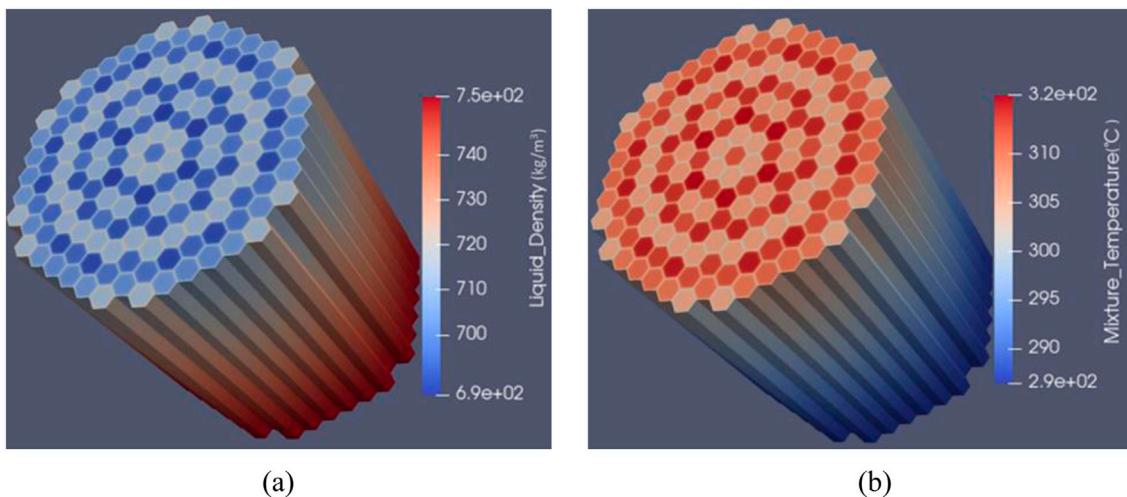


FIGURE 1
Coolant properties for HFP condition: (a) Coolant density and (b) Coolant temperature, respectively.

output file for VVERs. This option can be enabled by describing coordinates of mesh vertices in the CTF input file. This, however, can be prone to human errors for large geometries to be modelled manually. An existing hexagonal-geometry LWR preprocessor previously developed by Karlsruhe Institute of Technology (KIT) for CTF (Jimenez, 2015) was adapted to automate the input. This CTF preprocessor required modification due to updates in CTF4.3. Since the development of this preprocessor, updates in CTF have changed the input formatting which is accounted for in the adapted version. This led to the highly modified hexagonal lattice bundle pre-processor based on geometry, boundary conditions and nodalizations which provides the necessary information for VTK output generation. The preprocessor manages the bulk of error-prone work relating to geometry and sub-channel connections. Since its initial development, this preprocessor has been extended to full core preprocessing, including the option to generate parallelized CTF inputs, with multiple assemblies and additional sub-channel-to-sub-channel connections to consider interactions between the assemblies (Takasugi et al., 2021).

In addition, modeling and simulation of typical and specific features of several VVER-1000 assembly designs (for example, TVSM, TVS-2M, and TVSA assembly designs) as well as testing of the existing and new functionalities for VVER hexagonal fuel assemblies have been considered throughout this study.

The new additions in CTF were tested by modelling a VVER-1000 reactor core at sub-channel level. Specifications of the ongoing Rostov-2 Benchmark by OECD/NEA (Avramova, et al., 2021) were utilized in this model, considering both Hot Zero Power (HZP) and Hot Full Power (HFP) conditions. Guide and central instrumental tubes were modeled as fuel rods with zero power generation. The active fuel height was divided into 68 axial nodes to prevent large changes between adjacent nodes when considering the axial heights of spacer grids and their locations. Figure 1a shows the coolant density while Figure 1b depicts temperature distributions obtained for HFP conditions in the full core. These results demonstrate the capability to perform full core calculations for a VVER at sub-

channel level and illustrate the results using the new VTK output option.

3 Single Physics applications

In this section, CTF is utilized standalone as a single physics code. The results obtained on coarse and sub-channel meshes for single assembly and full core are compared with reference data. Results for TVSM, TVSA-5T, and TVS-2M assembly types are discussed.

CTF thermal-hydraulic solutions for VVER-1000 core subsets and full core have been tested against measured data and code-to-code against FLICA4 solutions (Spasov et al., 2014; Mitkov, 2022). FLICA4 (Fillion et al., 2011) is a 3D core thermal-hydraulics code of CEA with sub-channel capabilities.

CTF was also coupled with the statistical analysis tool Dakota to predict the uncertainties in outlet coolant temperatures at each sub-channel by propagating the uncertainties in inlet and boundary conditions for a VVER-type experimental facility consisting of 19 fuel rods.

3.1 Test results for Kozloduy-6 TVSM assemblies and core

3.1.1 Comparison to measured design data

A detailed description of the TVSM assemblies of Kozloduy-6 can be found in (Kolev et al., 2006; Kolev et al., 2007).

Table 1a shows the comparison of the CTF coarse-mesh computed vs. design core pressure drop over the heated length, for a Kozloduy-6 core with TVSM assemblies at hot full power (HFP).

Table 1b illustrates the comparison of the CTF sub-channel computed vs. measured assembly pressure drop over the heated length at HFP.

TABLE 1 Computed vs. measured data at HFP.

	Ref data	CTF	FLICA4
a. Coarse mesh computed heated core pressure drop vs. measured at HFP			
ΔP core, MPa	0.09455 ± 0.0047	0.09465	0.09520
b. Sub-channel computed single assembly pressure drop over the heated length vs. measured design data at HFP			
ΔP assembly, MPa	0.09493 ± 0.0047	0.09532	0.09563
	Ref data	CTF	FLICA4 coarse mesh
c. CTF sub-channel computed core average pressure drop, and temperature rise vs. reference data for the heated core height at HFP			
ΔP , MPa	0.09455	0.09496	0.09520
ΔT , K	30.3	30.4	30.45
			CTF - ref

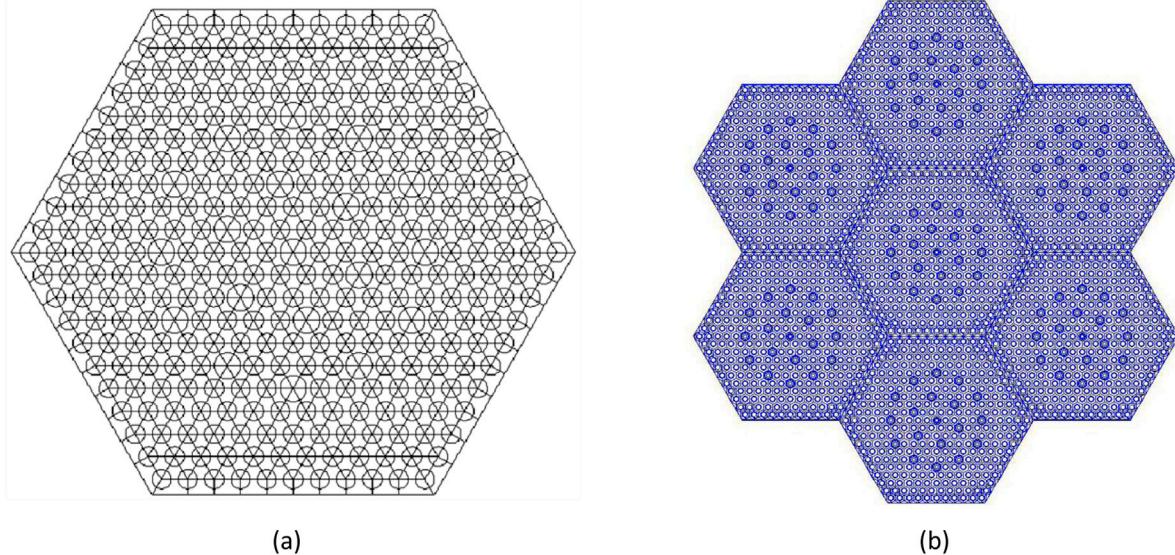


FIGURE 2
 (a) Coolant-centered sub-channel mesh for a VVER-1000 fuel assembly; (b) Mini-core domain of computation.

Table 1c shows preliminary full-core sub-channel computed results for the average core pressure drop and temperature rise vs. design reference data. In this calculation, the core average axial power distribution is as obtained from a coupled COBAYA3/COBRA3c calculation (Spasov et al., 2012) at hot full power, and the radial core power distribution is assumed uniform. The latter assumption is acceptable as the measured design data (Kozloduy NPP Report, 1987) have been obtained on assembly and core mock-ups with similar conditions.

These results have been obtained with the following modeling assumptions in CTF for discretization and correlation choice:

a) Full-core coarse-mesh CTF solution.

- Two-fluid, three-field quasi-3D flow model with one channel per assembly and inter-assembly cross-flow.
- 30 axial nodes in the heated core.

- 10 radial nodes in the fuel, 1 for the gas gap, and 1 for the cladding.
- Temperature-dependent thermal-physical properties of the fuel and cladding.
- Spacer grids are taken into account by vertical pressure loss coefficients.
- The bypass of 2.2% through the CR guide tubes is accounted for by decreasing the active coolant flow. No heat transfer at the guide tube walls.
- Single-phase mixing coefficient according to Rogers and Rosehart, two-phase multiplier according to Beus.
- Chen's model of nucleate boiling.
- Westinghouse-3 general-purpose DNB correlation with non-uniform axial power distribution.
- b) Single assembly sub-channel CTF solution.
- Two-fluid, three-field flow model.

- Coolant-centered radial mesh with 660 sub-channels per assembly (see [Figure 2a](#)). The mesh is non-uniform with triangular cells in the interior and rectangular and trapezoidal cells at the assembly periphery.
- 30 axial nodes in the heated core.
- Fuel model with 9 radial rings in the fuel, one for the gas gap, and one for the cladding.
- Temperature-dependent thermal-physical properties of the fuel and cladding.
- Spacer grids are not explicitly modeled and are considered by vertical pressure loss coefficients.
- The bypass of 2.2% through the CR guide tubes is considered by decreasing the active coolant flow. No heat transfer at the guide tube walls.
- Single-phase mixing coefficient according to Rogers and Rosehart, two-phase multiplier according to Beus
- The horizontal pressure loss coefficient is 0.5.
- The user-defined two-phase mixing coefficient is 0.01.
- Chen's model of nucleate boiling.
- W-3 general-purpose DNB correlation with non-uniform axial power distribution.

c) Full-core CTF sub-channel solution.

- Non-uniform coolant-centered radial mesh with 660 sub-channels per assembly.
- 30 axial layers in the heated part of the core.
- 3 227 400 cells in the 3D spatial mesh.
- 10 radial rings in the fuel pellet, 1 for the gap, and 1 for the cladding. The central hole is considered. Conduction only in the radial direction is considered.
- Temperature-dependent thermal-physical properties of the fuel and cladding.
- Cross-flow between the sub-channels in each assembly.
- No cross-flow between fuel assemblies, across the inter-assembly water gap (for this preliminary solution).
- The spacer grids are not explicitly modeled. They are taken into account through the vertical pressure loss coefficients.
- Use of rod friction factor correlation.
- The bypass through the water rods is not considered is taken into account by decreasing the active coolant flow.
- The user-defined two-phase mixing coefficient is 0.01.
- The horizontal pressure loss coefficient is 0.32555.
- Chen's model of nucleate boiling.
- Constant gap conductance coefficient corresponding to burnt fuel and equal to $3070 \text{ W/m}^2 \text{ K}$.
- Radiation heat transfer is not modeled.

The FLICA4 results have been obtained with the following modeling assumptions, chosen to be as close to the CTF ones as possible:

- Four-equation, fully 3D flow model
- Uniform mesh-centered triangular mesh with 726 radial cells at sub-channel level
- 30 axial layers in the heated part of the core
- 10 radial rings in the fuel pellet, 1 for the gap, and 1 for the cladding
- Same gap conductance coefficient as in the CTF calculation
- The bypass through the water rods is not considered and is taken into account by decreasing the active coolant flow

- Jens-Lottes model of nucleate boiling and Groeneveld CHF correlation.

The above CTF and FLICA4 modeling assumptions hold for all thermal-hydraulic and coupled neutronic/thermal-hydraulic (N/TH) solutions for Kozloduy-6 cores with TVSM assemblies reported in this study.

The CTF results in [Table 1](#) show good agreement with the reference data and the FLICA4 solutions, both at the assembly-channel and sub-channel level.

3.1.2 VVER assembly thermal-hydraulic benchmark solution

The assembly TH boundary condition problem is derived from the OECD/NEA VVER-1000 MSLB benchmark ([Kolev et al., 2016](#)) and is part of the LWR UAM Benchmark, Phase II: Case III-3b ([Hou Jason Benchmark for Uncertainty Analysis in Modelling for Design, 2019](#)). The task is to calculate the single assembly thermal-hydraulic parameters on a sub-channel mesh along with the associated uncertainties, given the time-dependent inlet and outlet boundary conditions, and assembly power. In this work, it is used for code-to-code comparison of sub-channel solutions. The boundary conditions include the assembly inlet temperature boundary condition ([Figure 3b](#)), the assembly inlet mass flow rate boundary condition, the assembly outlet pressure boundary condition, and the time history of the assembly power ([Figure 3a](#)). To obtain a more interesting test with higher vapor void fractions, the original MSLB assembly inlet temperature BC has been modified as shown in [Figure 3b](#), so that the inlet temperature is higher than in the original condition.

Computed results ([Mitkov, 2022](#)) for the vapor void fraction at peak power are illustrated in [Figures 4a, b](#).

[Figure 4a](#) shows the CTF vs. FLICA4 bundle averaged sub-channel computed assembly void fractions. The differences, which are larger in the sub-cooled boiling region, can be explained mainly by the different nucleate boiling models and to some extent by the different flow models. This comparison is only for illustration. The Chen correlation ([Chen, 1966](#)) is considered to be more accurate than the Jens-Lottes one ([Jens and Lottes, 1951](#)) and is relatively successful for applications with flow parameters as in the above test case. The Jens-Lottes correlation was the only option in the available FLICA4 version. It over-predicts the values of heat transfer coefficients for such flow conditions as it is derived for moderate flow velocity, where there is no suppression of nucleate boiling. The Chen superposition correlation calculates separately the heat transfer coefficients due to forced convection and vapor generation and adds them together: $h_{\text{chen}} = h_{\text{fc}} + h_{\text{nb}}$. The forced convection term contains a factor F_{chen} used to account for the enhanced forced convection caused by bubble formation and departure. For sub-cooled nucleate boiling $F_{\text{chen}} = 1$. The second term contains a suppression factor S_{chen} for forced convective boiling at high velocities where the boundary layer is small and temperature gradients are steep.

In CTF, in addition to the overall heat transfer from the wall to the liquid, the near-wall condensation heat transfer is also taken into account ([Salko and Avramova, 2016](#)). The net condensation heat flux is calculated as $q_{\text{cond}} = q_{\text{hn}} - q_{\text{spl}}$. Here q_{hn} is the heat flux into the liquid considering the condensation of all near-wall bubbles,

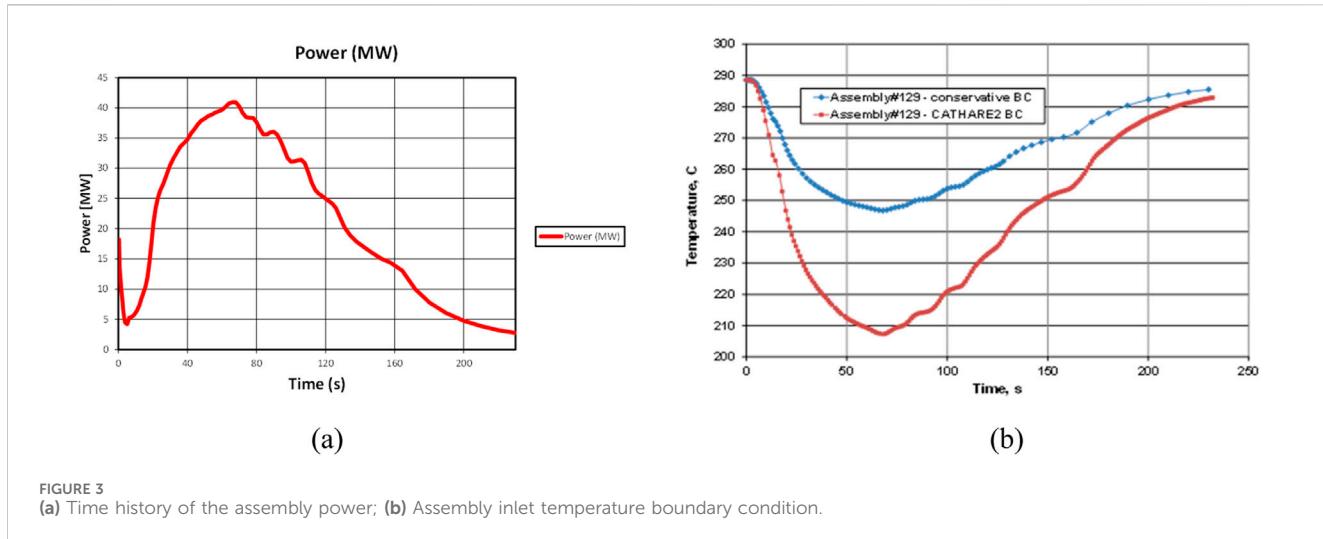


FIGURE 3
(a) Time history of the assembly power; (b) Assembly inlet temperature boundary condition.

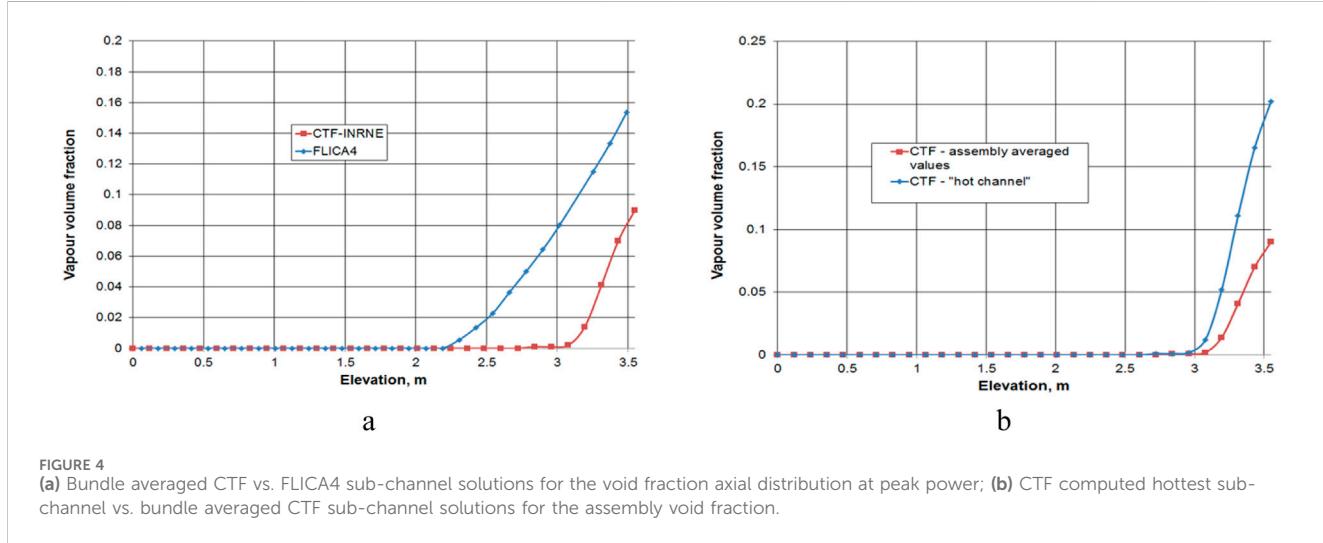


FIGURE 4
(a) Bundle averaged CTF vs. FLICA4 sub-channel solutions for the void fraction axial distribution at peak power; (b) CTF computed hottest sub-channel vs. bundle averaged CTF sub-channel solutions for the assembly void fraction.

calculated by the Hancox-Nicoll correlation (Hancox and Nicoll, 1971), and the wall heat transfer to single-phase liquid q_{spl} is obtained by the Dittus Boelter equation (Dittus and Boelter, 1930) with the subcooled liquid properties used in the equation.

A number of modifications of the Chen correlation have been proposed by different authors, mostly by refining the S_{chen} parameter to extend the application range to high heat fluxes and high mass fluxes, see for example, the discussion in (Yan et al., 2015). For rod bundles, the Chen-like and other recent nucleate boiling models are yet to be verified.

Figure 4b shows the CTF computed hottest sub-channel vs. bundle averaged sub-channel solutions for the assembly void fraction. The results illustrate the impact of the higher resolution on the assessment of local TH parameters.

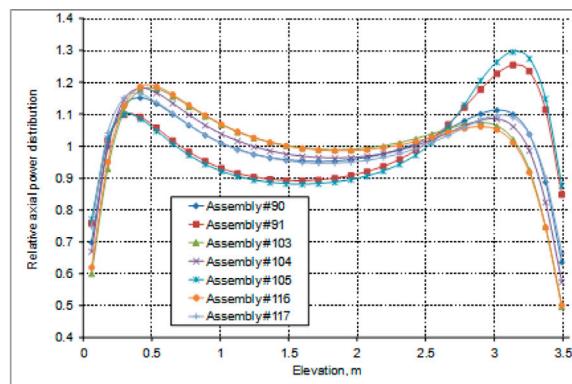
3.1.3 Single assembly vs. mini-core sub-channel solutions

Comparisons of single-assembly and multi-assembly sub-channel solutions for strongly disturbed cores in reactivity transients show that taking into account the environment of the

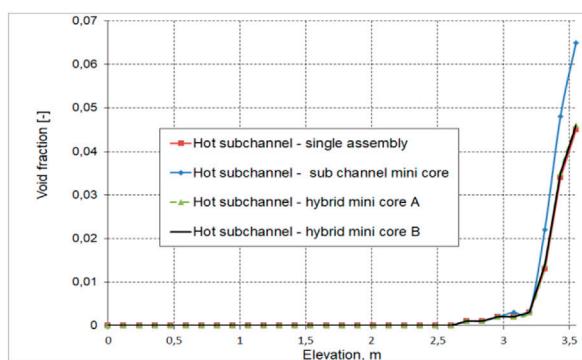
considered assembly is important for the accurate assessment of the local core safety parameters. A computationally efficient way to do this is the solution of mini-core TH problems, extracted from a preceding full-core nodal N/TH calculation.

In the present study, a VVER mini-core of 7 assemblies consisting of the hottest assembly and the surrounding 6 assemblies, as shown in Figure 2b, is considered. A postulated VVER-1000 MSLB accident with 8 stuck rods after scram is assumed, as summarized in Section 3.3 below. The core conditions correspond to the highest return to power after scram. Single hot assembly and mini-core boundary condition (BC) problems are solved with pre-calculated assembly-wise TH BCs and assembly-wise power and axial power distributions as obtained from the full-core nodal solution (see Figure 5a). The axial and radial power distributions are obtained from a coupled COBAYA/CTF transient solution at the nodal level (Mitkov et al., 2018; Mitkov et al., 2021). In this calculation, the radial pin power distribution is approximately taken as at HFP.

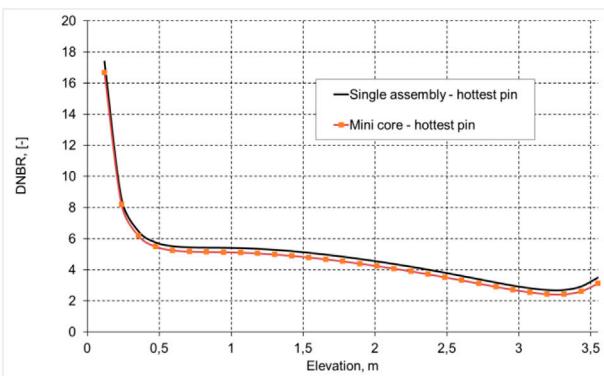
Figures 5b, c below illustrate a comparison of local core parameters obtained from CTF sub-channel TH solutions of



(a) COBAYA/CTF nodal-level computed assembly-wise axial power distributions at peak core power



(b) CTF computed vapor void fraction in the hottest sub-channel of a single hot assembly vs. the hot assembly in a mini-core of 7 assemblies



(c) CTF computed DNBR for the hottest pin in a single hot assembly vs. the hot assembly in a mini-core of 7 assemblies

FIGURE 5

CTF channel and sub-channel results. (a) COBAYA/CTF nodal-level computed assembly-wise axial power distributions at peak core power. (b) CTF computed vapor void fraction in the hottest sub-channel of a single hot assembly vs. the hot assembly in a mini-core of 7 assemblies. (c) CTF computed DNBR for the hottest pin in a single hot assembly vs. the hot assembly in a mini-core of 7 assemblies.

single hot assembly vs. 7-assembly mini-core BC problem. The results show more conservative values of the local core safety parameters when the assembly environment is considered.

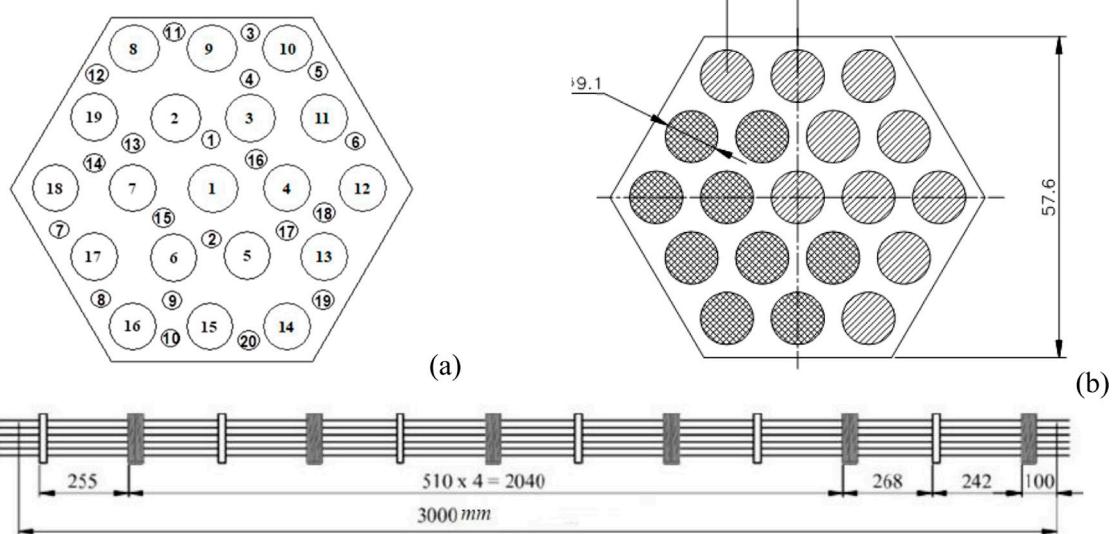
A more accurate approach would be a further development to solve a coupled pin-cell multi-physics boundary condition problem for the mini-core considered, with lateral assembly boundary conditions extracted from the nodal core solution.

3.2 Test results for the TVSA-T5 experiments

Best estimate codes become more useful where inaccuracies in modelling or numerical errors in result calculations are pointed out. This can be done by combining the best estimate codes with uncertainty analysis. Best Estimate Plus Uncertainty (BEPU) is a procedure to evaluate uncertainties by quantifying

TABLE 2 Test case: experimental conditions and measured temperatures.

Pressure, (MPa)	Inlet temperature, (°C)	Mass flowrate, (kg/s)	Power, (kW)
Test Case: Experimental Conditions			
15.71	149.9	1.824	1196.6
Sensor position	Outlet temperature, (°C)	Sensor position	Outlet temperature, (°C)
Test Case: Measured Outlet Coolant Temperatures			
1	281.8	11	280.8
2	295.2	12	302.4
3	270.9	13	303.5
4	282.4	14	303.2
5	278.1	15	292.4
6	277.3	16	291.1
7	305.1	17	278.8
8	306.7	18	284.8
9	297.5	19	288.3
10	298.7	20	299.5

FIGURE 6
(a) Radial visualization of TBCA-T5 temperature sensor locations; (b) TVSA-T5 geometric characteristics and axial view of spacer and support grids.

the errors in resulting calculations. Under this framework, CTF was coupled with the statistical analysis tool Dakota to predict the uncertainties in outlet coolant temperatures at each sub-channel by propagating the uncertainties in inlet and boundary conditions for a VVER-type experimental facility consisting of 19 fuel rods. Dakota (Adams, et al., 2014) developed by Sandia National Laboratories (SNL), is a software used for optimization and uncertainty quantification using sampling or stochastic expansion methods.

3.2.1 Description of the TVSA-T5 experiments

The coupling of CTF and Dakota was done for the specific application of developing the necessary uncertainty propagation framework for the experimental setup used in Exercise II 3 (6a/6b) of the OECD/NEA Uncertainty Analysis in Modeling (UAM) Light Water Reactor (LWR) Benchmark Phase II (Hou, et al., 2022).

The TVSA-T5 experimental facility was utilized to analyze the local temperature distribution of coolant within fuel rod cells in a scaled-down VVER-type fuel assembly with 19 heated rods.

TABLE 3 Test case parameters.

Pin number	Relative power	
Test case: radial power distribution		
1	0.688	
2	1.402	
3	0.682	
4	0.670	
5	1.402	
6	1.352	
7	1.352	
8	0.695	
9	0.676	
10	0.670	
11	0.665	
12	0.670	
13	0.676	
14	0.695	
15	1.352	
16	1.341	
17	1.318	
18	1.341	
19	1.352	
Flow area (mm ²)	Wetted perimeter (mm)	Heated perimeter (mm)
TVSA-T5 assembly specifications		
1637.5	742.7	543.2

TABLE 4 Uncertainties in inlet and boundary conditions.

Parameter	Bounds ($\pm 3\sigma$)	PDF
Uncertainties in Inlet and Boundary Conditions		
System pressure (%)	± 2.0	Normal
Flow rate (%)	± 4.5	Normal
Inlet fluid temperature (K)	± 2.0	Uniform
Power distribution (%)	± 3.0	Normal
Uncertainties in Geometry		
Rod displacement (mm)	± 0.45	Normal
Rod diameter (mm)	± 0.10	Normal

Stainless steel tubes heated by electrical power were used to simulate the heat production of nuclear fuel rods. Experiments were carried out at various inlet coolant temperatures, input powers, pressures and flowrates with measurements taken of coolant temperature at

different radial positions at the assembly outlet. The set of conditions selected for testing the developed Dakota/CTF coupling and the experiment coolant temperatures measured at various sensor locations are provided in Table 2, respectively.

Figure 6a shows the location of rods and temperature sensors. The geometric specifications and axial positions of the spacer and support grids are provided in Figure 6b. The heated length of the experiment is 3000 mm, with spacer grids equally spaced over the axial assembly height.

The input radial pin power distribution for the test case is given in Table 3, with numbering corresponding to the positions indicated in Figure 6a. The axial power distribution is uniform over the heated length. The total flow area, wetted perimeter, and heated perimeter are also summarized in Table 3.

Preliminary assessments were carried out to demonstrate Dakota/CTF's capability to propagate uncertainty in input parameters to the coolant temperature at the exit of the assembly outlet. This study investigated the sensitivity to inlet and boundary conditions such as flow rate, inlet coolant temperature, pin power distribution, and system pressure.

3.2.2 Uncertainty analysis framework

The data provided from the experiments at the TVSA-T5 facility were sufficient to generate a model to study nominal steady-state operating conditions. However, the experimental and measurement input uncertainties associated with the data were not finalized at the time of this analysis due to ongoing benchmark specifications development. Instead, variations in core boundary conditions and manufacturing uncertainties available in the UAM Benchmark Phase II (Hou Jason Benchmark for Uncertainty Analysis in Modelling for Design, 2019) for VVER bundle thermal-hydraulics were employed as input uncertainties to demonstrate the methodology and the developed Dakota/CTF implementation.

Table 4 shows the input distributions that were sampled for the inlet and the boundary conditions, and their probability distribution functions, which are bounded in three standard deviations ($\pm 3\sigma$). Single values of the system pressure, flow rate, and inlet temperature are sampled for each run. The radial power distribution sampling is performed by considering each rod separately (19 samples per run). Due to CTF requirements that the average relative radial power is equal to 1.0, a script was written to process sampled powers and normalize these before providing these radial powers to CTF.

Uncertainties in rod geometry due to manufacturing variations were also considered. Considering these uncertainties in all rod geometries and their individual relative displacements would require many trials to acquire significant results. In this exercise, the displacement of only a single rod and the variation of its outer diameter are sampled, as is assumed for UAM exercises. A corner rod was chosen for uncertainty analysis regarding the displacement and rod diameter since the surrounding sub-channels have different geometries. The applied geometric uncertainties are provided in Table 4, noting that the angle of displacement is also sampled (from a uniform distribution). It is important to propagate the uncertainties to sub-channel inputs such as flow area, wetted perimeter, gap length, and gap width. Because these values are necessary for the CTF input and derived from sampled values of the rod displacement and diameter, a script

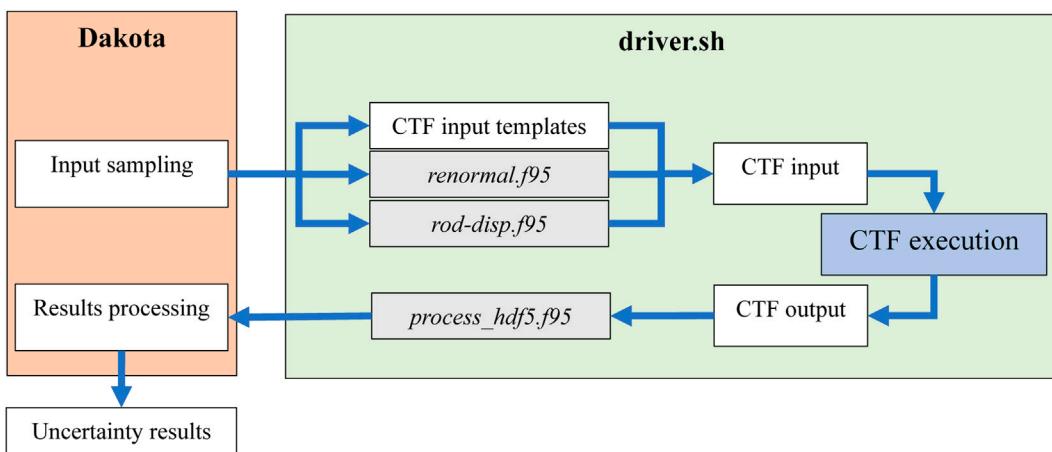


FIGURE 7
Dakota/CTF coupling flowchart.

was written to evaluate these after the sampling of uncertain parameters.

CTF and Dakota were coupled utilizing the black-box interfacing capability of Dakota with supporting Shell scripts to facilitate execution of scripts for intermediate calculations, CTF input generation, CTF execution, and results extraction. The files created to facilitate coupling are as follows.

1. Driver.sh: The primary Shell script that fills input templates, interfaces to other scripts and CTF for execution from Dakota.
2. Process_hdf5.py: Python script utilizing the CTF-SubKit tools to extract quantities of interest (QOIs) from CTF's HDF5 output.
3. Renormal.f95: Performs the normalization of the perturbed radial power distribution for each sampled set to maintain the average relative radial power distribution at 1.0 after perturbation (requirement for CTF).
4. Rod-disp.f95: Calculates the flow area, wetted perimeter, gap length, and gap width for sub-channels near the displaced rod, considering displacement direction and magnitude as well as changes to the rod diameter.
5. CTF input templates: Several partial CTF inputs correspond to the parts of the CTF input that do not rely on the intermediate calculations of sub-channel/rod parameters. Together with the outputs of *renormal.f95* and *rod-disp.f95*, this is combined to make the full CTF input for each set of sampled parameters.

The workflow between CTF and Dakota is illustrated in Figure 7 where input uncertainties and their distributions were supplied in the Dakota input to be used for sampling input parameters. Each Dakota-sampled parameter set is sent to the driver Shell script, which fills the CTF input templates and runs intermediate calculation scripts before compiling the CTF input. CTF is then executed via Shell script. The coolant temperature at the outlet of the assembly for the given positions of the thermocouples in the experiment, the quantities of interest (QOI) in this analysis, were

sent to Dakota from CTF's output for each sample to be summarized for uncertainty analysis by Dakota.

This study considers random sampling of input uncertainties for the uncertainty propagation to QOI responses. Random sampling can be done in different ways in Dakota, via Monte Carlo (MC) sampling or Latin Hypercube sampling (LHS). Instead of using MC sampling, this project utilized LHS for the input uncertainties to guarantee coverage of the range of the sample space for the given probability distribution, which also reaches the same precision of the MC sampling but with a smaller cost in sample size.

3.2.3 Sensitivity and uncertainty results

The inlet temperature, flow rate, system pressure, rod displacement (magnitude and direction), rod diameter, and individual relative powers for each of 19 rods are sampled in each of 2000 runs. The effect of each input parameter on the responses (outlet coolant temperatures at specified thermocouple locations) was investigated. Dakota provides asynchronous execution to run multiple evaluations simultaneously. For this case, 50 simultaneous evaluations were run in parallel for the 2000 sample sets. This sample size was found to produce 95% confidence intervals with a relative difference between upper and lower confidence intervals of less than 0.1% in the mean and approximately 6% in the standard deviation of considered output uncertainties.

In addition to providing uncertainty estimates for the outputs based on QOI result data, this random sampling also gives important information regarding sensitivity, summarized by correlation coefficients. Pearson and Spearman correlation metrics are utilized to evaluate the impact of the uncertainties in the input parameter on the QOIs. The Pearson and Spearman correlations are calculated based on Equation 4 where x_i and y_i are the input sample and the corresponding response, while the \bar{x} and \bar{y} are the mean values. Pearson uses raw data while Spearman employs the rank data.

$$\text{corr}(x, y) = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2 \sum_i (y_i - \bar{y})^2}} \quad (4)$$

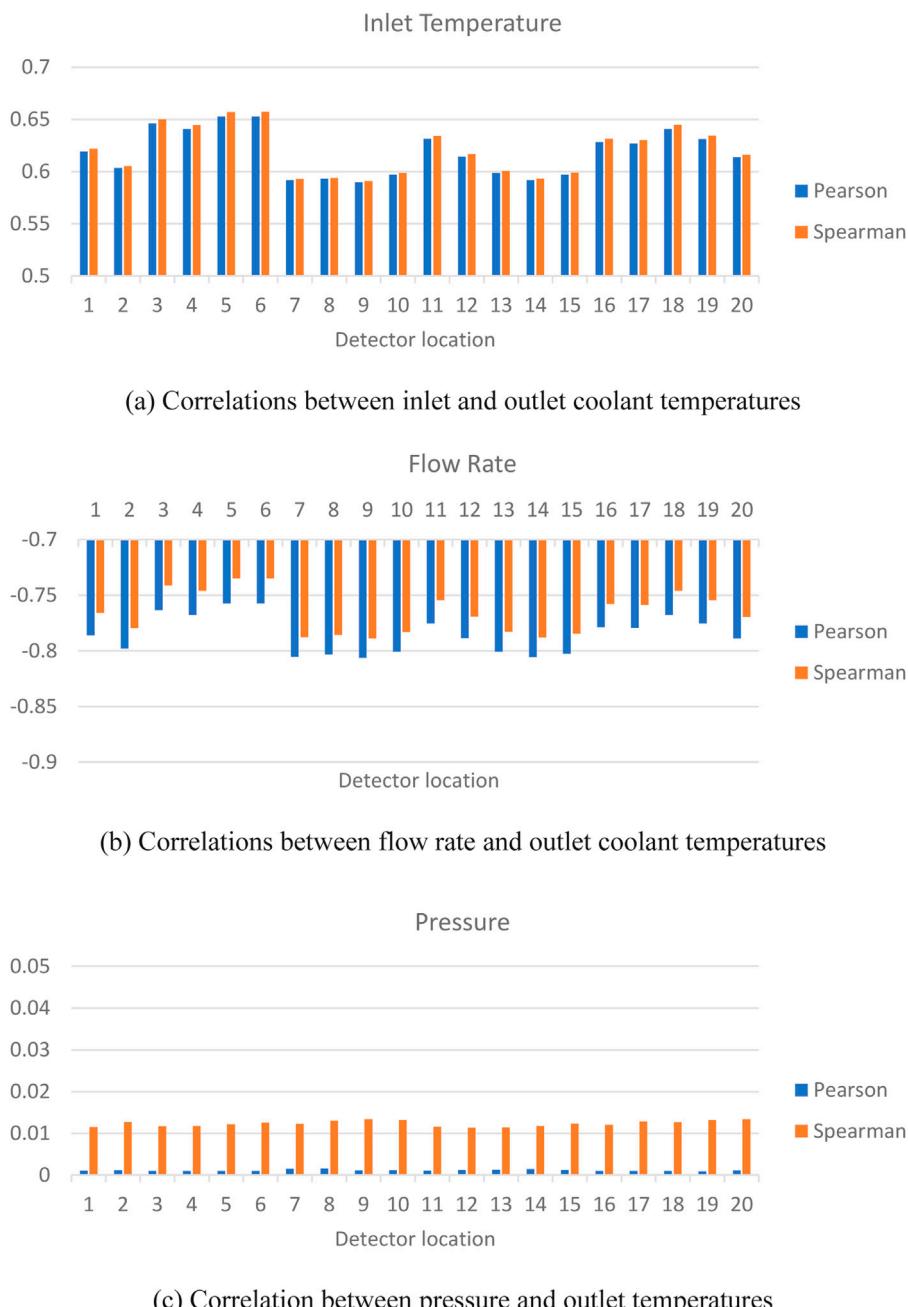


FIGURE 8
Correlations between boundary conditions. (a) Correlations between inlet and outlet coolant temperatures. (b) Correlations between flow rate and outlet coolant temperatures. (c) Correlation between pressure and outlet temperatures.

The Pearson correlation is related to how linearly correlated the input and output are, while the Spearman correlation measures how the input and output are monotonically related. The Spearman rank correlation removes some effects due to outliers and differences in scale. Partial correlation coefficients can be used to isolate the effects of interdependent variables. Although it is recommended that Spearman correlation be used due to its ability to assess correlation based on non-linear relationships, the correlation between the various input parameters and the outlet coolant temperatures is obtained using both Spearman and Pearson metrics.

3.2.3.1 Inlet and boundary condition sensitivities

The correlations between the outlet coolant temperatures corresponding to changes in the inlet coolant temperatures and the flow rate are shown in Figure 8a, b, respectively (note the scale of the axes). Besides, observing that the Pearson and Spearman correlations are close to each other, the outlet coolant temperature is strongly sensitive to uncertainties in both input parameters. As expected, the hottest sub-channels are slightly less sensitive to the inlet temperatures since they are next to rods with high-power generation. However, there is a strong positive correlation in all channels between the inlet and outlet coolant

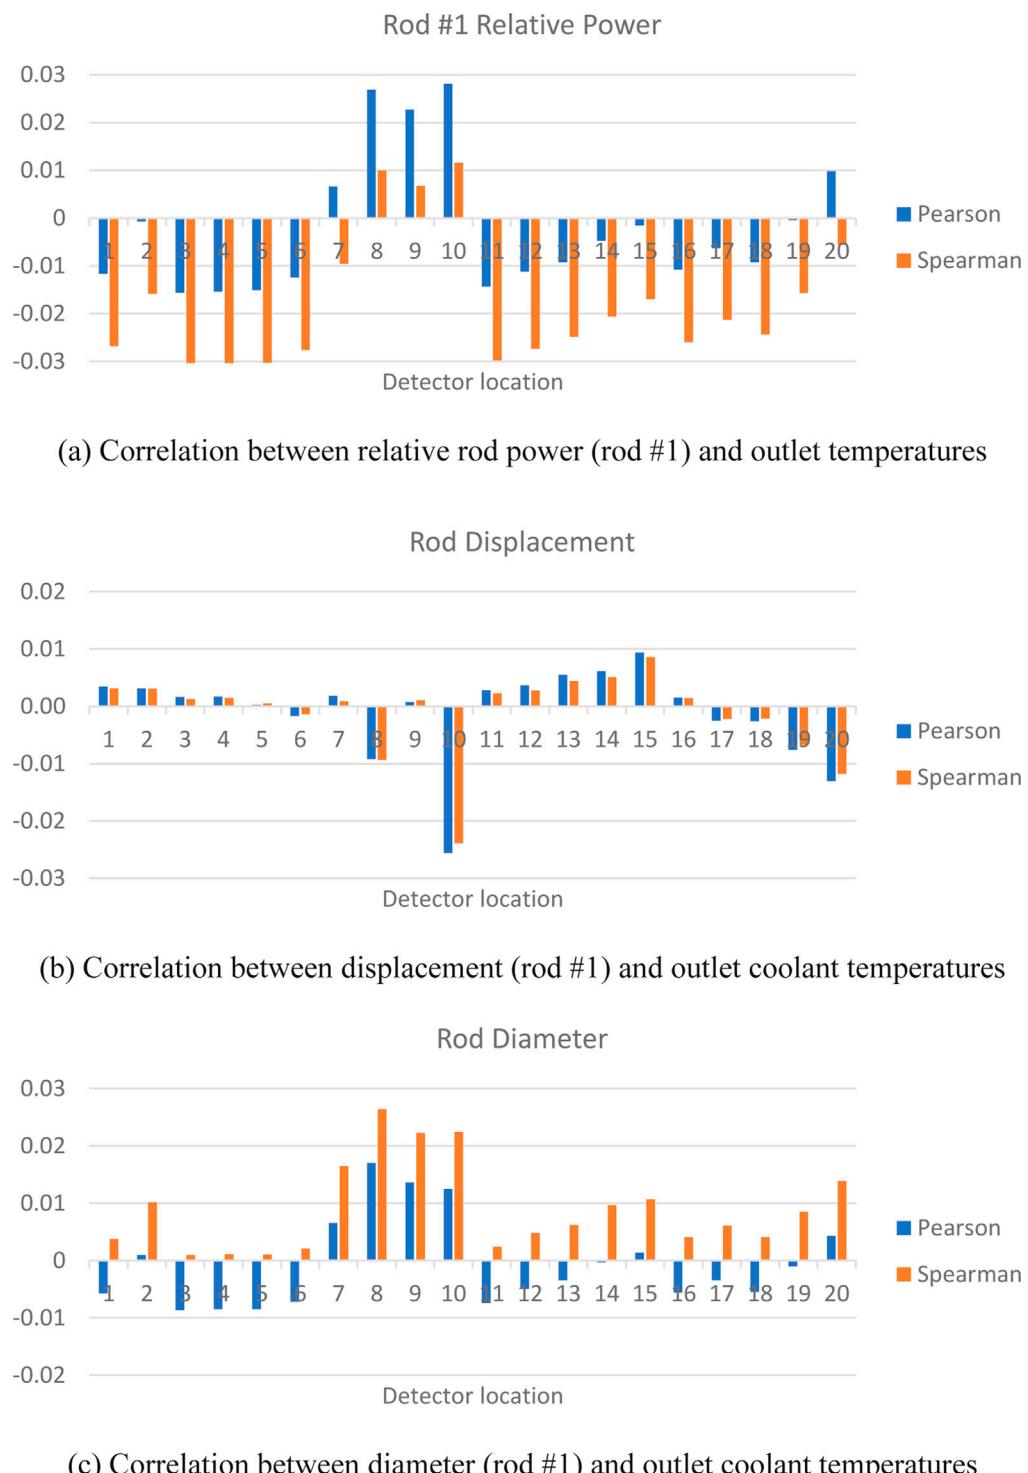


FIGURE 9

Correlations between rod #1 parameters and outlet coolant temperatures. (a) Correlation between relative rod power (rod #1) and outlet temperatures. (b) Correlation between displacement (rod #1) and outlet coolant temperatures. (c) Correlation between diameter (rod #1) and outlet coolant temperatures.

temperatures. Conversely, the coolant flow rate has a strong negative correlation to the outlet coolant temperature. The system pressure is found to have only a small positive correlation to coolant outlet temperatures, as shown in Figure 8c. Apparent differences observed in this figure between the two correlation methods may be attributed

to the scale of comparison; relative to the possible sensitivity ranging -1 to 1 , they are close to one another.

Figure 9a shows selected results from the relative power distribution for relative rod power adjustments made in rod #1 (numbering of rods is from left to right, bottom to top of geometry).

TABLE 5 Outlet coolant temperature uncertainty estimates.

Outlet coolant temperatures	Mean	Std Dev	Skewness	Kurtosis
Location 1	2.887E+02	2.248E+00	1.630E-02	-2.954E-01
Location 2	2.944E+02	2.263E+00	1.300E-02	-2.681E-01
Location 3	2.795E+02	2.211E+00	1.917E-02	-3.225E-01
Location 4	2.812E+02	2.218E+00	2.030E-02	-3.160E-01
Location 5	2.772E+02	2.205E+00	1.828E-02	-3.343E-01
Location 6	2.772E+02	2.205E+00	1.892E-02	-3.311E-01
Location 7	2.991E+02	2.264E+00	1.202E-02	-2.643E-01
Location 8	2.988E+02	2.264E+00	1.562E-02	-2.631E-01
Location 9	2.988E+02	2.274E+00	7.846E-03	-2.510E-01
Location 10	2.967E+02	2.266E+00	1.253E-02	-2.573E-01
Location 11	2.846E+02	2.230E+00	2.271E-02	-3.022E-01
Location 12	2.906E+02	2.250E+00	1.694E-02	-2.850E-01
Location 13	2.960E+02	2.265E+00	1.479E-02	-2.726E-01
Location 14	2.987E+02	2.268E+00	1.219E-02	-2.655E-01
Location 15	2.968E+02	2.267E+00	1.234E-02	-2.670E-01
Location 16	2.856E+02	2.238E+00	1.836E-02	-3.039E-01
Location 17	2.859E+02	2.240E+00	2.075E-02	-2.949E-01
Location 18	2.811E+02	2.220E+00	2.048E-02	-3.082E-01
Location 19	2.846E+02	2.232E+00	2.213E-02	-2.950E-01
Location 20	2.906E+02	2.253E+00	1.538E-02	-2.758E-01

Though the analysis includes all rods, rod #1 is selected for representative demonstration. The correlation coefficients in all cases are relatively small, however, some important trends are observed. In general, the correlations are positive for nearby detector locations. For example, detector locations #8, #9, and #10 are adjacent to rod #1, which results in a positive correlation between them. Negative correlations to distant detector locations are also observed due to the normalization process, which maintains the average relative power at 1.0. Detector locations #7 and #20 are adjacent to subchannels that are correlated positively with power changes in Rod #1. The Pearson correlation coefficient characterizes the linearity of the data, while the Spearman correlation coefficient considers whether the data has a monotonic trend. These detector locations experience dynamic responses in temperature due to their dependence on competing effects of perturbed heating of rod #1 through the adjacent subchannels and opposite changes to heating in other rods due to normalization. This results in less monotonic behavior than other detector locations, leading to a difference in Pearson and Spearman coefficient signs.

3.2.3.2 Manufacturing uncertainties

The correlation coefficients relating to geometrical uncertainties are provided in Figures 9b, c. Rod displacement has only a small effect on outlet coolant temperatures, mostly near the displaced rod (#1). Due to the displacement of the rod, changes in sub-channel

mixing may result in varied behavior near the displaced rod, which results in non-monotonic behavior. The temperatures at the channels around the rod, which are affected by rod diameter uncertainty (rod #1) are positively correlated to diameter changes of the rod. This indicates that if the rod is wider than the other rods, the flow areas in the adjacent sub-channels are smaller than in the other ones, which causes the relative temperature to rise due to reduced cooling. In similar ways, a larger surface area would also convey more heat from the rod surface to the adjacent coolant. The different trends shown by the Spearman and Pearson approaches seen in Figure 9c can be attributed to non-linearity considered by Pearson, especially at the farthest detector locations, such as locations #3, #4, #5, #6, and #11. However, these effects are relatively small in scale and can be considered negligible.

3.2.3.3 Output uncertainty distributions

Table 5 summarizes the uncertainty estimates obtained by the analysis of output responses of outlet coolant temperatures for the given measurement locations. The mean values correspond roughly to the power profile, having higher outlet temperatures at locations in the southwest portion of the assembly. Skewness is relatively low; however, significant kurtosis is observed. The kurtosis indicates a flatter-than-normal distribution in the outputs, which may occur because of the uniform distribution sampled for the inlet temperature distribution.

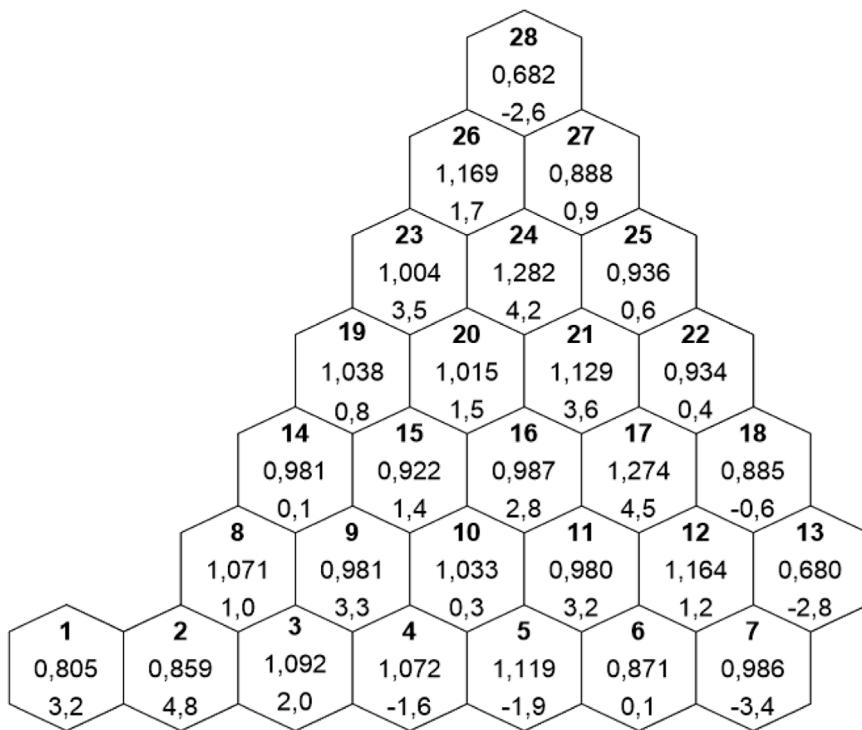


FIGURE 10
Comparison of COBAYA/CTF computed vs. Kozloduy-6 plant data for the relative assembly powers (δ in %).

4 Multi-physics applications

4.1 COBAYA/CTF coupling

The CTF code version has been coupled with the COBAYA core physics simulator developed at the Technical University of Madrid, Spain (Lozano et al., 2008; Lozano et al., 2010). The results obtained with the coupled code COBAYA4/CTF have been compared both with the available measured data as well as with the predictions of COBAYA3/FLICA4 coupled codes. The COBAYA3/FLICA4 coupling (Jimenez, 2009) for hexagonal assemblies is based on the coupling functions in Salome 5 (Chanaron et al., 2015) supplemented by FLICA4 routines. The COBAYA4/CTF coupling (García-Herranz et al., 2017) is based on the MED Coupling libraries in Salome 6 (www.salome-platform.org). The Python script for coupled calculation includes a damping scheme to smooth the power profile and accelerate the convergence. The power at the current time-step (n) is expressed as $0.6*P(tn-1)+0.4*P(tn)$ (COBAYA team, 2015).

4.1.1 COBAYA/CTF solution vs. Kozloduy-6 steady state plant data

A steady-state COBAYA/CTF solution (Mitkov, 2022) for the radial core power distribution has been compared with plant data for Kozloduy-6 Cycle 8, 270.4 EFPD, at 91% core power. The core is near the end of the cycle and contains once, twice, and three times burnt assemblies. The code coupling is at the nodal/assembly-channel level. The comparison makes part of the experimental validation of the coupled COBAYA/CTF solution

along with the 2-group diffusion multi-parameter cross-section library.

The solution has been obtained with a multi-parameter cross-section library in compact NEMTAB format and with some simplifications, as generated for VVER-1000 MSLB benchmarking purposes with portability considerations. The discontinuity factors are implicit, and the reflector parameters are obtained from a 1D model.

Figure 10 shows the comparison of the computed relative assembly powers vs. plant data. The computed $k_{eff} = 0.99502$ which deviates from 1 by 498 pcm. The average absolute deviation in assembly powers is 2.07%, and the max deviation is within the 5% error tolerance of the core monitoring system data. The agreement with the reference is considered to be acceptable in view of the quality of the core monitoring system data and the 1D reflector data in the cross-section library.

4.1.2 OECD NEA VVER-1000 MSLB benchmark (V1000CT-2) solutions

A detailed description of the MSLB test problem is given in the benchmark specifications (Kolev et al., 2006). The reference core is a real Kozloduy-6, Cycle 8, 3-year batch core at 270.4 EFPD. The core contains once, twice, and three times burnt UO_2 fuel of initial enrichment 4.23 w/o and 4.4 w/o. The plant transient is initiated at hot full power by a large guillotine-type break of steam line #4 outside the containment, upstream of the steam intercept valve. The steam generator feed-water valve in the faulted loop fails to close on the MSLB signal and remains open. The main coolant pump in the faulted loop fails to trip on the MSLB signal,

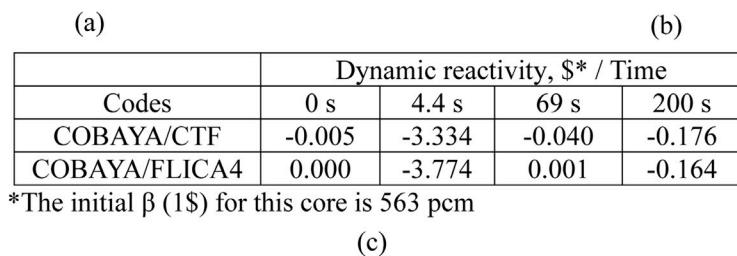
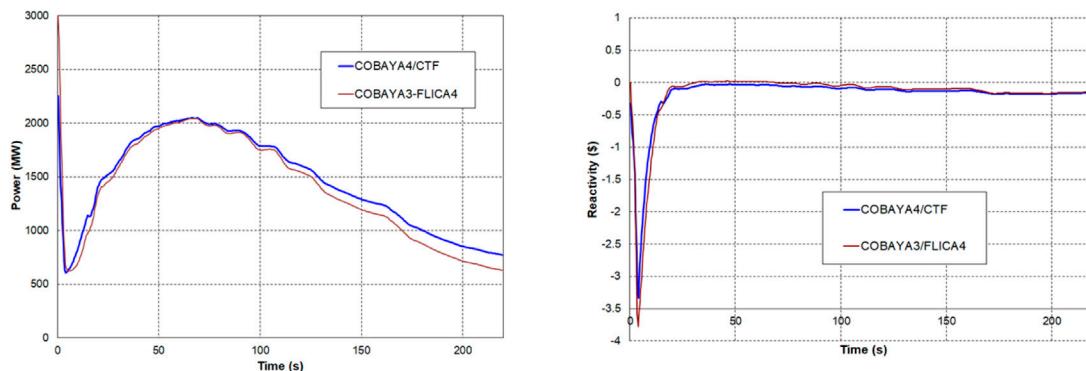


FIGURE 11
NEA VVER-1000 MSLB benchmark – (a) Core BC problem: Computed time history of the total core power; (b) Dynamic reactivity as computed from the 3D solutions; (c) Dynamic reactivity as computed from the 3D solutions.

and all pumps remain in operation. Two peripheral control rod clusters remain stuck out of the core after the scram. The steam generator in the faulted loop continues uncontrolled cooling till the complete evaporation of the secondary water. A cooler sector is formed at the core inlet, with overcooling of up to 80°C.

In this study, we present results from cross-verification by code-to-code comparison of COBAYA/CTF vs. COBAYA/FLICA4 nodal solutions of the core boundary condition problem (see Figures 11a–c). The core TH boundary conditions have been obtained from a validated CATHARE2.5 multi-1D 24-sector reactor vessel TH solution (Spasov et al., 2010; Spasov et al., 2017), using the reactor vessel boundary conditions from the VVER-1000 MSLB benchmark specifications (Kolev et al., 2006). This calculation used a multi-parameter two-group diffusion cross-section library generated by (Petrov et al., 2015) with the APOLLO2 Linear Surface MOC (Santandrea et al., 2008; Sanchez, 2010). For better code testing with higher return to power after scram, a modified option of the cross-section library has been generated by artificially reducing the control rod absorption cross-sections in the library so that the scram worth is reduced by about 50%. With this library, the computed return to power due to the combined action of feedback effects and the contribution of the delayed neutrons reaches 67% of the nominal rated power.

The full-core nodal/assembly-channel results of COBAYA/CTF and COBAYA/FLICA4 show good agreement. The small bias is attributed mainly to the different core mixing models in CTF and FLICA4.

The FLICA4 results are obtained with a fully 3D flow model and 3D calculation of the local gradient-driven inter-assembly mixing, but without any turbulence model (turbulence is included in more recent FLICA versions). For hexagonal assemblies, CTF uses a quasi-

3D calculation scheme for hexagonal assemblies to obtain the gradient-driven average mass flux between adjacent assemblies, plus a simple turbulent diffusion model.

The mixing models in the two codes have not been specifically tested vs. reference data. CTF and FLICA4 single assembly and full-core calculation models, including their single-phase flow mixing models, have been tested against available measured data for the pressure drop and temperature rise over the heated length (see Table 1). Table 1 illustrates the coarse-mesh CTF and FLICA4 computed results for the core pressure drop as compared with reference data. For both models, the agreement is very good, and the CTF results are a little closer to the reference. The same trend is seen in the sub-channel single-phase calculation results shown in Table 1.

To assess the differences between the predictions of two models, a core mixing numerical experiment at HZP was performed with FLICA4 vs. CTF to simulate the propagation of a VVER core inlet temperature disturbed sector to the core outlet. The disturbance was $\Delta T = -61.15$ K over 1/6 of the core. The FLICA4 results show that the absolute temperature change related to inter-assembly mixing is in the range of 3.0K–4.4K in the sector periphery and 5.9 K in the central assembly. The CTF solution gives mixing-related temperature changes < 3.0 K at the periphery of the disturbed sector (Kolev et al., 2016).

4.1.3 Postulated VVER-1000 MSLB transient with 8 stuck rods

The considered accident scenario is based on an aggravated variant of the pessimistic scenario in the OECD/NEA VVER-1000 MSLB benchmark (V1000CT-2), for the same burnt reference core. In addition to the multiple equipment faults as summarized in

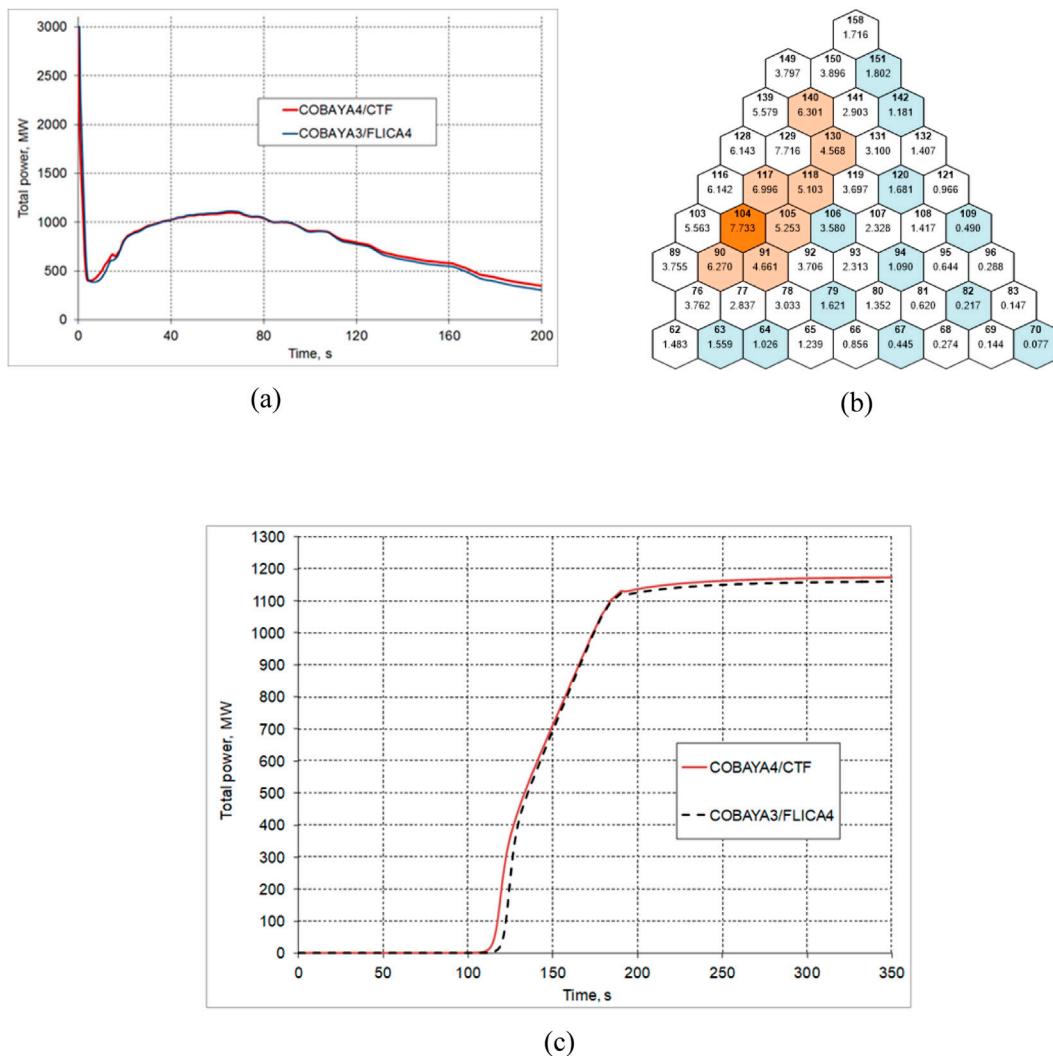


FIGURE 12
Postulated VVER-1000 MSLB accident with 8 stuck rods: **(a)** Computed time history of the total core power; **(b)** COBAYA/CTF computed relative assembly powers; **(c)** Simulated uncontrolled withdrawal of Kozloduy-6 CR Bank 9 at 2 cm/s - Time history of the total core power.

Section 4.1.2, eight peripheral control rod clusters—in assembly # 90, 91, 105, 106, 117, 118, 130, 140 are assumed to remain stuck out of the core after scram, all of them in the overcooled sector (marked in beige in Figure 12b). The realistic option of the MSLB cross-section library (Petrov et al., 2015) is used. The scram worth is significantly reduced due to the eight rods stuck out of the core. The task is to solve an MSLB core boundary condition problem using coupled 3D neutronic/thermal-hydraulic codes, given the core thermal-hydraulic boundary conditions as specified for the core-only solution of the V1000CT-2 benchmark.

Figure 12a shows the computed time history of the total power. The peak power after scram is app. 37% of the nominal rated power. The results show a very close agreement of the COBAYA/CTF and COBAYA/FLICA4 solutions.

Figure 12b illustrates the radial core power distribution at the time of maximum return to power. The hottest assembly is #104. The assemblies with fully inserted control rods in the core are marked in blue, and the assemblies with stuck rods—in beige. The

core power distribution is strongly disturbed. The hottest subset of the core is near the core periphery. There are large 3D flux deformations and gradients, and local multi-physics effects, which are of interest for fine-mesh higher-resolution analysis. Such a CTF sub-channel thermal-hydraulic analysis of the hottest assembly and a 7-assembly mini-core around it is illustrated in Section 3.1.3, Figures 5a–c of the present paper.

4.1.4 Control rod bank withdrawal

Figure 12c illustrates the comparison of COBAYA4/CTF vs. COBAYA3/FLICA4 solutions for the transient core power in the case of postulated complete withdrawal of Rod Bank 9 at 2 cm/s, starting from the fully inserted position at HZP (Spasov et al., 2018). The reference core is burnt, as described for the MSLB transient analysis. The steady state results agree well, and the computed rod bank worth is nearly the same, 1071 pcm vs. 1080 pcm. The transient total power solutions are in good agreement, with only a small bias due to differences in the thermal-hydraulic models.

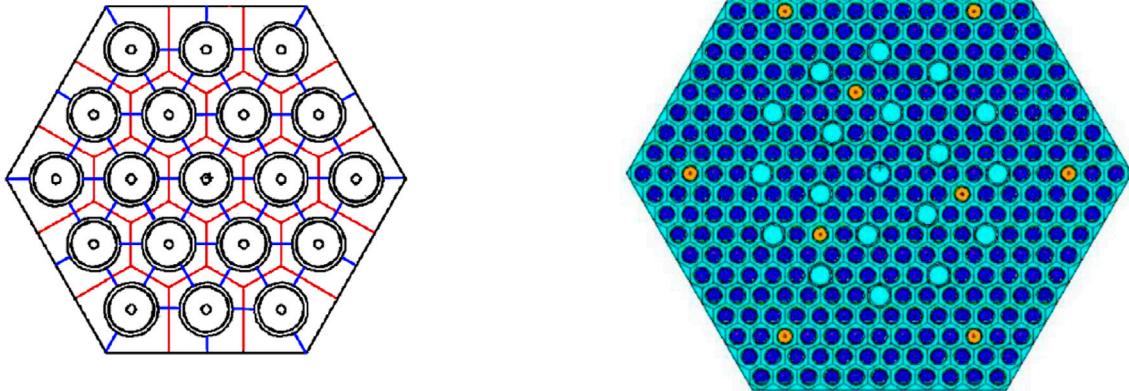


FIGURE 13
Coolant-centered subchannels (blue) vs. Rod-centered cells (red) – left; VVER-1000 assembly model – right.

4.2 CTF coupling with MCNP and Serpent

Two different couplings between the CTF and Monte Carlo neutronics codes have been developed. First, an MCNP/CTF (Kutlu et al., 2022) coupling was programmed with the latest version of CTF available at the time of development. Later, CTF continued to be improved for VVER applications, and CTF was coupled with Serpent 2.2.0. The following sections discuss the two coupling frameworks and differences between coupling approaches and provide results from coupled CTF calculations.

4.2.1 Coupling schemes

A wrapping driver script was developed to facilitate the external coupling between MCNP (Werner, 2017) and CTF to manage the iterative execution of the codes. The driver script controls the execution explicitly using a fixed-point iteration scheme (Kelly, 2017). This code driver provides all necessary intermediate processing of data, performs iteration checks, and contains scripts for the generation of new input files as a method of data transfer. The focus was to implement the capabilities for steady-state analysis; hence, temporal mapping was left for future work. An input generator for MCNP6.2 was incorporated to simplify and automate the writing of input files used in the data exchange process. The preprocessing of MCNP6.2 inputs also provides an opportunity to generate large-scale models while reducing human errors, which may occur during the development of lengthy input files necessary for pin-wise neutronic models.

A spatial mapping process was conducted between coolant-centered sub-channels in CTF and rod-centered cells in MCNP (Bennett, 2015). The coolant density and temperature are calculated by an area-weighted average before sending to MCNP. According to the subchannel numbering in the CTF model, weighting factors are calculated automatically in the geometry preprocessing based on the subchannel areas, which depend on the location of rods and assembly specifications.

Figure 13 shows a sample of the overlaid rod-centered neutronics model over the coolant-centered CTF model.

Equations 5–7 calculate the T_{cool} representing the rod-centered coolant temperature in MCNP while taking into account the locations of the rod and adjacent subchannels. (Kutlu, 2021).

$$T_{cool} = w_1 T_{chan_1} + w_2 T_{chan_2} + w_3 T_{chan_3} + w_4 T_{chan_4} \quad (5)$$

$$T_{cool} = w_1 T_{chan_1} + w_2 T_{chan_2} + w_3 T_{chan_3} + w_4 T_{chan_4} + w_5 T_{chan_5} \quad (6)$$

$$T_{cool} = w_1 T_{chan_1} + w_2 T_{chan_2} + w_3 T_{chan_3} + w_4 T_{chan_4} + w_5 T_{chan_5} + w_6 T_{chan_6} \quad (7)$$

The subchannel temperatures are denoted by T_{chan_i} that correspond with weight w_i to a given MCNP coolant cell. Additional steps are taken for output data processing, including scripts for extracting neutron tallies and normalization to obtain power distribution, and calculating cell-averaged fuel temperatures. A similar spatial mapping approach is also used in Serpent/CTF coupling. Both coupled codes were developed considering one-to-one spatial mapping in the axial direction.

The coupled calculation begins with a CTF run to obtain the initial thermal-hydraulic calculations using a uniform power profile. Once completed, the CTF output is read by SubKit, a module included in the CTF utilization package to extract thermal-hydraulic data of interest. The thermal-hydraulic data are processed and used to construct a new input for MCNP. MCNP is then run to provide pin-wise power distribution based on neutron flux tallies. The first power profile obtained is utilized by CTF as an initial power distribution for the coupled iteration. This initialization approach does not require an initial user-input power or temperature distribution and is using instead a uniform power profile at the start. Once the initialization step is complete, CTF and MCNP are run in sequence, and the thermal-hydraulic and neutronic profiles are exchanged through the generation of new input files at each iteration until user-input convergence criteria are met.

A similar initialization process was implemented in the CTF/Serpent. The same CTF input was run with a uniform power profile to provide an initial temperature distribution for Serpent before the

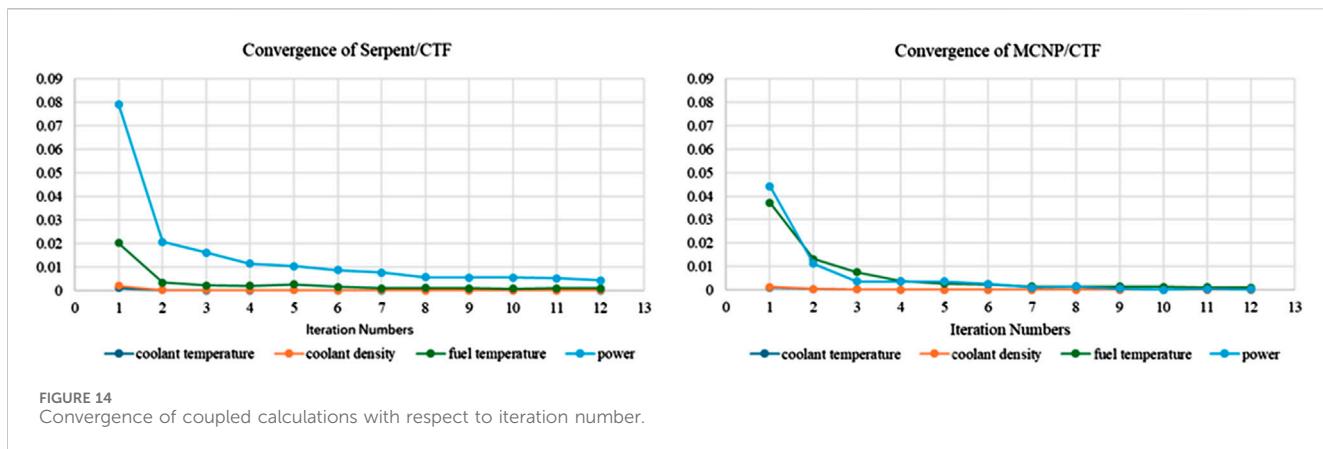


FIGURE 14
Convergence of coupled calculations with respect to iteration number.

main coupled iterations started. The coupled CTF/Serpent code was also created externally, and the data transfer process was conducted via input and output files. However, there are differences between the coupling methodologies. For instance, Serpent has a separate coupled calculation mode that allows updates to Serpent by a Serpent multi-physics interface. Power information is extracted from the Serpent interface following the lattice and pin locations as described in the input, which simplifies the coupling development by maintaining the full core input description and removing the need to develop a Serpent input generator.

The developed Serpent/CTF coupling considered weighted fuel temperatures described in [Equation 8](#), which were sent by the driving script to the multi-physics interface in order to update Serpent ([Leppänen et al., 2015](#)) pin cells for evaluation of Doppler broadened cross sections. This is in contrast to the MCNP/CTF coupling, where the radial temperature distribution within fuel rods was processed with a weighting factor for different radial cells within MCNP fuel rods and calculated as part of the developed MCNP input generation process.

$$T_{cell} = 0.3T_{fc} + 0.7T_{fs} \quad (8)$$

In [Equation 8](#), T_{fs} and T_{fc} are the temperatures of the fuel pellet surface and fuel centerline, respectively.

In both coupled codes, the convergence is monitored based on user-defined criteria, ε , given in [Equation 9](#). The coupling scheme compares feedback fields exchanged between the coupled codes, including flux (power profile), temperatures of fuel and coolant, and coolant density from the current iteration, ' $k+1$ ', to the previous iteration, ' k '. ' X ' is the parameter being checked at each iteration, and ' i ' represents the associated cell, rod, or sub-channel index.

$$\varepsilon \geq \max_i \left(\frac{|X_i^{k+1} - X_i^k|}{X_i^k} \right) \quad (9)$$

The main objective in this work was to apply CTF in multi-physics calculations with different neutronics codes. Modeling the same problem and keeping the modeling approaches as much as possible is essential for a meaningful comparison between the two coupled codes. For this reason, the same VVER-1000 assembly was modelled for demonstration, and the results are compared in the following sections.

4.2.2 VVER-1000 assembly model

A VVER-1000 assembly model was developed based on a burnup calculation benchmark ([Löetsch, et al., 2010](#)) to test the code couplings. A TVSA 30AV5 fuel assembly was modeled, consisting of 303 fuel rods with 2.99 w/o enriched UO_2 , nine fuel rods with burnable poison with 2.4 w/o enriched UO_2 and 5% Gd_2O_3 of mass fraction, 18 guide tubes, and one central instrumental tube as shown in [Figure 13](#). The main specifications for the assembly, including geometry and system parameters, can be found in the benchmark documentation. Assembly power and coolant flow rate were calculated as averages from the system parameters. Reflective boundary conditions were considered radially, while black boundary conditions were used at the bottom and the top of the assembly, following axial reflector regions composed of three layers. The assembly was divided into 13 axial nodes along the active fuel length, and spacer grids were considered at each axial level of the CTF model with a fixed loss coefficient. Spacer grids were not considered in the MCNP or Serpent neutronics models. Two radial nodes are modeled in CTF for the fuel pellets. The axial reflector was modeled in MCNP and Serpent with material information as provided in ([Löetsch, et al., 2010](#)), and using the ENDF/B-VII.1 cross-section library. For each MCNP and Serpent simulation/coupled iteration, 300 inactive cycles, 1000 active cycles, and 10^5 particles per cycle were run.

To test the developed coupled codes, the converged test case solutions were compared. The results obtained for k-eigenvalue were also compared to another coupled code evaluation using the Karlsruhe Institute for Technology (KIT) Serpent2/SUBCHANFLOW coupled code ([García, 2020](#)) and also compared to the initial (no burnup) benchmark reference solutions ([Löetsch, et al., 2010](#)). Refinement of the code and direct comparison with other coupled code systems is desirable, however, these test cases serve as an initial code-to-code comparison.

4.2.3 Results

The following convergence criteria were used for the test case: $\varepsilon_1 = 5.0 \times 10^{-3}$ for power; $\varepsilon_2 = 1.0 \times 10^{-3}$ for coolant temperature, coolant density, and fuel temperature. The convergence was reached after 12 iterations in both coupled codes; convergence of the assessment parameters is presented in [Figure 14](#). Oscillations are

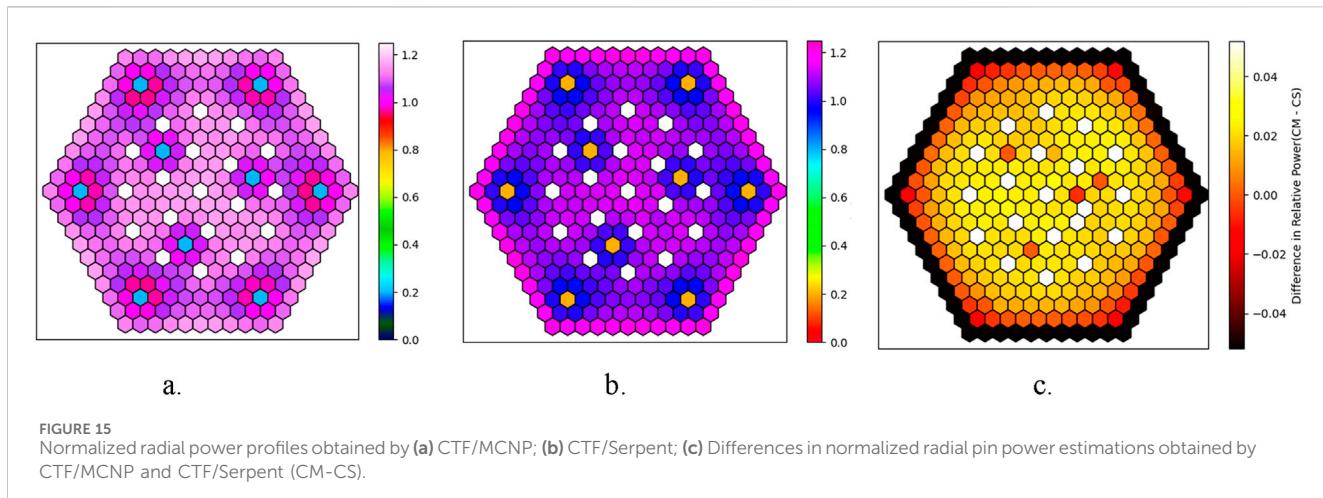


FIGURE 15
Normalized radial power profiles obtained by (a) CTF/MCNP; (b) CTF/Serpent; (c) Differences in normalized radial pin power estimations obtained by CTF/MCNP and CTF/Serpent (CM-CS).

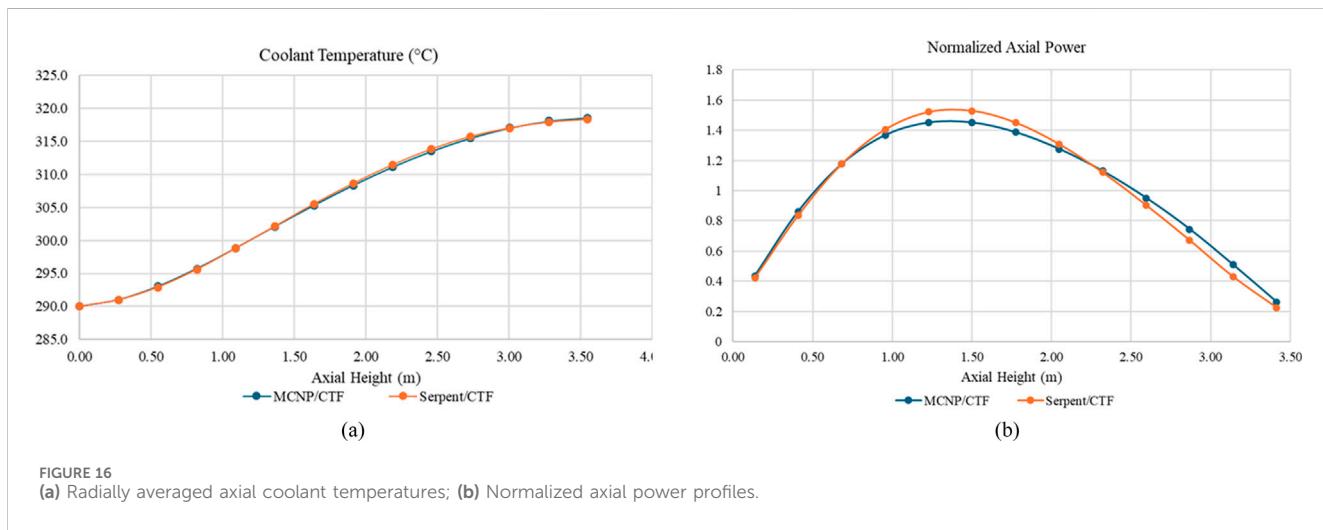


FIGURE 16
(a) Radially averaged axial coolant temperatures; (b) Normalized axial power profiles.

TABLE 6 Comparison of eigenvalue results.

Code	k-eigenvalue
CASMO (Löetsch, et al., 2010)	1.13810
HELIOS (Löetsch, et al., 2010)	1.13816
NESSEL (Löetsch, et al., 2010)	1.132015
Serpent2/SUBCHANFLOW (García, 2020)	1.13200 ± 0.00004
CTF/MCNP	1.13078 ± 0.00006
CTF/Serpent	1.13621 ± 0.00011

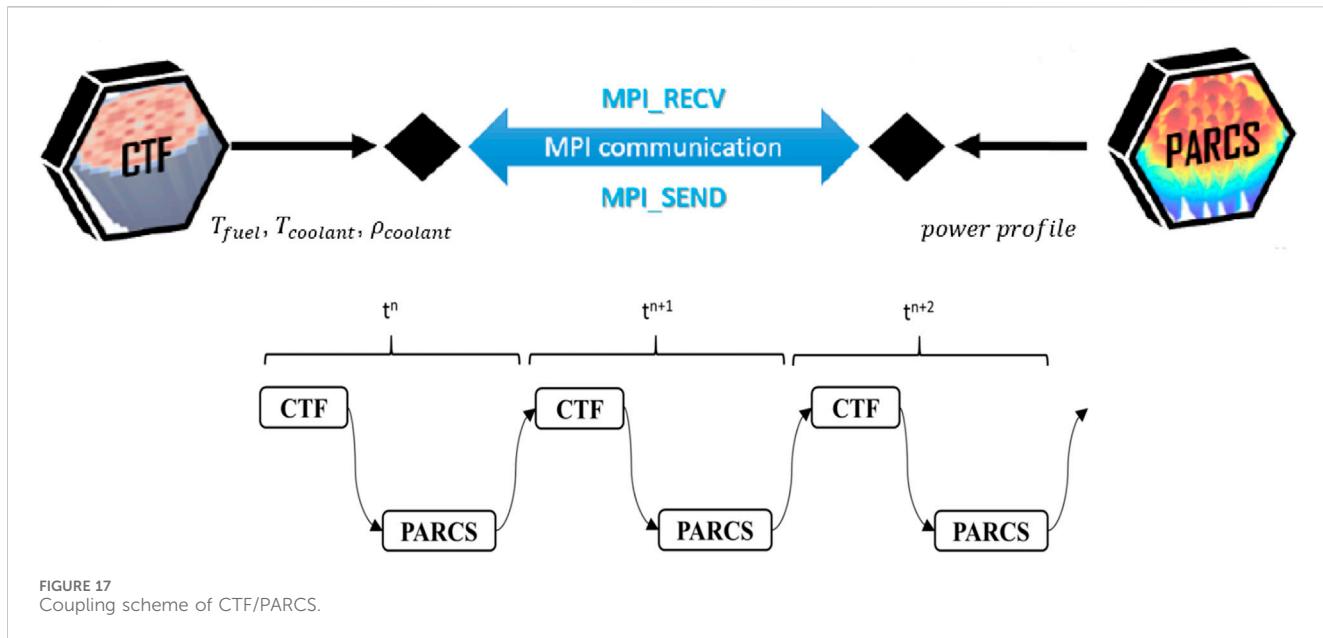
observed between some iterations due to successive feedback effects as well as the statistical noise of the stochastic approach used in MCNP and Serpent. However, a sufficiently large number of iterations, as well as the use of the power relaxation methods, ensure convergence.

The converged VVER-1000 radial pin power distributions obtained by CTF/MCNP and CTF/Serpent are given in Figures 15a, b. Note that the different color sets were used since the power profiles were hardly dissimilar. The relative pin powers are obtained

by integrating the power generated axially over the active rod lengths and normalizing to the average pin power. The fuel rods containing Gd_2O_3 burnable poison and 2.4 w/o enriched UO_2 are observed, as expected, to have low pin powers compared with the 2.99 w/o enriched UO_2 fuel rods without burnable poison. The 1/3 radial symmetry of the model is well reflected in the resulting radial power distribution.

The difference in radial pin power distributions obtained by comparing the CTF/MCNP (CM) and the CTF/Serpent (CS) is shown in Figure 15c. The differences in estimated power profiles can be attributed to the usage of different neutronics solvers. The axial coolant temperature distributions obtained from CTF/MCNP and CTF/Serpent are compared and shown in Figure 16a, having mostly overlap with the same trend. The average coolant temperature was found to be 318 °C at the outlet of the assembly. The difference between the inlet and the outlet ranges from 25 °C to 30 °C.

The converged axial power distributions are compared in Figure 16b. The power profiles from both coupled calculations show a tendency to be bottom-peaked. This outcome is expected since the coolant acts as a better moderator in the lower part of the assembly, ascribed to higher coolant density. In spite of the fact that both CTF/MCNP and CTF/Serpent coupled codes used the ENDF/



B-VII.1 cross-section library, the axial power profiles are not exactly matched due to the differences in the couplings mentioned formerly, as well as differences in MCNP and Serpent.

Comparisons of eigenvalues evaluated by the converged coupled results are compared to the benchmark (Löetsch, et al., 2010) - and a Serpent2/SUBCHANFLOW result (García, 2020) in Table 6. The k-eigenvalues chosen from the benchmark for this comparison are those calculated at 0 MWd/tU. The k-eigenvalue was obtained from the Serpent2/SUBCHANFLOW coupled calculation after nine iterations. The relative statistical errors of Serpent2/SUBCHANFLOW, CTF4.0/MCNP6.2, and CTF4.3/Serpent2.2.0 are presented in Table 6. 2000 active cycles with 10^5 particles were run for the Serpent calculation of the coupled Serpent2/SUBCHANFLOW; note that no information was available regarding the cross-section library utilized in this study. Differences from this solution may, partially, be attributed to the use of different libraries in addition to differences in modeling methodologies.

4.3 PARCS/CTF coupling

4.3.1 Coupling scheme

CTF and the neutronic code PARCS are coupled loosely at the reactor core level for multi-physics modelling at K. Picard iteration has been selected for this coupling; acceleration methods are not currently implemented, however, relaxation of the relative tolerance of CTF is applied. In (Delipei, et al., 2022), the CTF/PARCS framework is used to model exercises of LWR-UAM Phase III, which includes the modeling of the TMI-1, a square lattice PWR design, under various core conditions. In this paper, the coupled code is used for a nodal/assembly-wise model of Rostov-2 full core. Rostov-2 is a VVER-1000 type of reactor that is equipped with TVS-2M hexagonal fuel assemblies.

The United States Nuclear Regulatory Commission (U.S. NRC) code PARCS (Downar et al., 2018) is a 3D neutron

TABLE 7 Conditions of rostov-2 core simulations.

Reactor parameter	Units	Value
a. HZP Conditions of Rostov-2 Core		
Power	%	0.1
Fuel temperature	°C	279
Coolant density	g.cm-3	0.766
Control Rod Groups 1–3 and 5	% withdrawn	100
Control Rod Group 4	% withdrawn	57.3
Control Rod Groups 6–10	% withdrawn	0
Boric acid concentration	g/kg	3.74
Parameter	Units	Value
b. BOC conditions of Rostov-2		
Depletion	EFPD	0.0
Average Core Power	MW	44.4
Boric Acid Concentration	g/kg	6.5
Volumetric Flow Rate	m3/hour	87,342.7
Inlet Temperature	°C	279.8
Control Rod Groups 1–9	% withdrawn	100
Control Rod Group 10	% withdrawn	72.9
Parameter	Units	Value
c. Code-to-code comparison of SCF/Serpent and CTF/PARCS		
Outlet Core Temperature (°C)	553.425 (K)	553.412 (K)
keff	$0.99787 \pm 3.0E-05$	0.99996

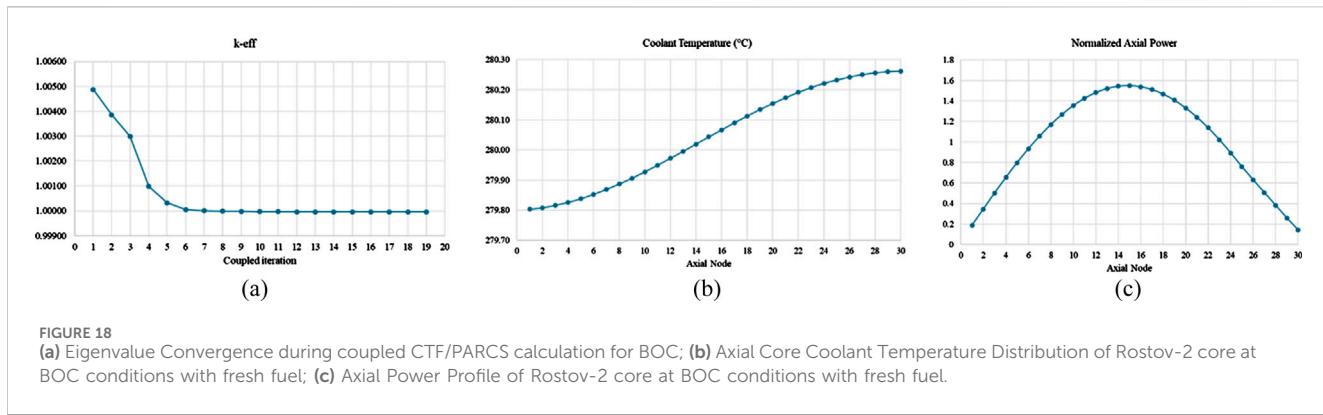


FIGURE 18

(a) Eigenvalue Convergence during coupled CTF/PARCS calculation for BOC; (b) Axial Core Coolant Temperature Distribution of Rostov-2 core at BOC conditions with fresh fuel; (c) Axial Power Profile of Rostov-2 core at BOC conditions with fresh fuel.

kinetics code that can solve steady-state and time-dependent multi-group neutron diffusion equations in both Cartesian and hexagonal geometries. PARCS reads burnup-dependent cross-sections in PMAXS format and can conduct pin power reconstruction if group-dependent flux form functions are included in the cross-section libraries. The pin power reconstruction feature of PARCS has not been utilized in the results shown here and is left for future work. GRS, a Rostov-2 benchmark participant, has created cross-section libraries for the Rostov-2 core model in the NEMTAB format (Henry, 2022) and provided them for testing by other participants. These cross-section libraries have been converted into PMAXS format at NCSU and are used for the coupled PARCS/CTF calculations here.

The CTF/PARCS coupling uses a flexible Message Passing Interface (MPI) communication protocol based on a server-client coupling algorithm. The coupling algorithm facilitates the implementation of a loose coupling method without the necessity of heavy modification in the source codes of CTF and PARCS. Figure 17 shows the coupling scheme where CTF acts as the server and PARCS as the client.

Coupled calculations can be activated via a command line argument in CTF to launch additional MPI processes for client codes, in this case, PARCS. An interface input file needs to be provided so that the information for the client can be launched by the server.

The neutron kinetics coupling class developed in CTF contains protocols to send tagged messages to instruct operations in PARCS, including initializing PARCS, steady-state iteration, depletion step, transient time step, convergence, simulation finalization, and error.

The synchronization of the codes is controlled by MPI standard functions: MPI_SEND and MPI_RECV. The MPI_SEND function is used for writing the outputs/instructions from a buffer. MPI_RECV function is used for reading the corresponding inputs/instructions from the buffer. These MPI send and receive commands result in temporally explicit coupling where one code is performing a calculation while the other code waits for the calculation to finish. Feedback parameters such as fuel temperature, coolant temperature, coolant density, and power are exchanged during each global coupled iteration. Spatial mapping can be done manually by users, or developed auto mapping procedures can be utilized to

provide information about the different physical spatial domains. More detailed information regarding the coupling methodology of CTF/PARCS can be found in (Delipei, et al., 2022).

In this work, the PARCS/CTF coupled code is adopted in steady-state mode for the calculation of BOC conditions of the Rostov-2 core.

4.3.2 Full core steady-state VVER model

An OECD/NEA benchmark on Reactivity Compensation of Boron Dilution by Stepwise Insertion of Control Rod Cluster into the VVER-1000 Core (Avramova, et al., 2021) focuses on the comprehensive evaluation and validation of advanced high-fidelity multi-physics simulation codes using transient tests carried out at Unit 2 of the Rostov Nuclear Power Plant (NPP). Rostov-2 is a VVER-1000 type of reactor that is equipped with TVS-2M fuel assemblies and operates on an 18-month fuel cycle.

The benchmark encourages the participants to compare the performance of traditional and high-fidelity multi-physics codes for steady-state, burnup, and transient calculations. Data collected at steady-state conditions and during transient testing provides a wide range of parameters, such as reactor power, coolant temperatures, control rod positions, and boron concentrations. These data collections create pivotal opportunities to ensure the accuracy and reliability of simulation codes, such as CTF/PARCS.

The first exercise of the benchmark, after cross-section generation, requests the calculation of the Hot Zero Power (HZP) state of the core to evaluate the cross-section libraries. The HZP conditions are provided in Table 7a.

In the benchmark, the operational history from the Beginning of Cycle (BOC) to the initial state before the experiment is provided, and it is suggested to start the transient from an averaged total core burnup of 36.37 effective full power days (EFPD). Therefore, the first step is to model and test the VVER full core for HZP and BOC conditions, which is used here for verification and validation of the CTF/PARCS coupled code. The BOC conditions are given in Table 7b.

4.3.3 Results

CTF and PARCS inputs were prepared assembly-wise, where each assembly corresponds to a channel in CTF. The TPEN solution

method was chosen for PARCS, which is suitable for triangular lattice geometry. Two mapping tables were prepared manually, which were read and utilized to calculate weighted averages of contributing channels to corresponding nodes. The weighting factors in the mapping tables are normalized to 1.0. The Doppler temperature required by PARCS is obtained by using the weighting factors 0.7 and 0.3 for surface and fuel centerline temperatures, once received from CTF, respectively. The coupling interface then reassigns power profiles in CTF after obtaining the power information from PARCS.

From HZP calculation of PARCS/CTF, boric acid concentration was found to be 3.7 g/kg, while CER, a participant of Rostov-2 benchmarks, computed boric acid concentration 3.8 g/kg using a standalone neutronics diffusion code KIKO3DMG. Two results were considered in agreement. For the BOC condition model of Rostov-2, Figure 18a shows the convergence of k_{eff} over 19 global coupled iterations.

The computed axial core coolant temperature distribution and normalized axial core power profile of the coupled steady state are provided in Figures 18b, c, respectively, where 30 equidistant nodes were modelled in both PARCS and CTF. The coolant temperature shows changes of less than a degree due to the low core power level, while the power distribution follows a cosine shape, which is negligibly bottom peaked. Keep in mind that these results are obtained by modelling 44.4 MW core power in fresh fuel.

Table 7 shows the core outlet temperature and k_{eff} comparisons between SCF/Serpent (Beydogan, 2023) and PARCS/CTF for BOC calculations are found to be in good agreement.

5 Conclusions

Coarse-mesh and sub-channel thermal-hydraulic models of VVER-1000 assemblies and the full-core were tested against measured data for the pressure drop and temperature rise, as well as code-to-code vs. FLICA4 solutions. The CTF results agree well with the reference data.

In a preliminary study, CTF vs. FLICA4 sub-channel results for the vapor void fraction in the sub-cooled region of a VVER assembly were compared for illustration. The test case was the LWR UAM Phase II: Case III-3b. CTF used the Chen nucleate boiling model, and FLICA4 used the Jens-Lottes model. In this case, the FLICA4 solution cannot be used as a conditional reference since the Chen correlation is better suited to the flow conditions given. There is a need for further verification of the available and potential new sub-cooled boiling models in CTF for VVER rod bundles.

Coupled COBAYA/CTF vs. COBAYA/FLICA4 full-core nodal solutions were compared for cross-verification. The test cases were a core boundary condition problem of the NEA VVER-1000 MSLB Benchmark (V1000CT-2) and a control rod bank withdrawal test problem. The results are in good agreement, with slight differences attributed mainly to the different flow mixing models.

The impact of the inter-assembly mixing in CTF sub-channel simulations was studied by comparison of a single hot-assembly calculation vs. a 7-assembly mini-core calculation. The assembly powers and boundary conditions were extracted from a full-core

nodal calculation of a postulated MSLB accident with 8 stuck control rods. The results show more conservative predictions of the local core safety parameters when the assembly environment is taken into account.

The steady-state VVER-1000 full core was modeled in the latest version of CTF to test the new pre- and post-processing options embedded in CTF. Another study done by CTF for single physics thermal-hydraulics calculation coupled it with the sensitivity and uncertainty tool Dakota to apply to a steady state VVER-like mini assembly model consisting of 19 rods. The coupling of CTF and Dakota was utilized as an uncertainty propagation framework for a model based on a steady state using the experimental setup planned for use in Exercise II-3 (6a/6b) of the OECD/NEA UAM LWR Benchmark Phase II. Uncertainties in inlet temperature, flow rate, pressure, power distribution, and geometry were propagated to outlet coolant temperature for the TBSA-T5 VVER-type experimental bundle case. The focus was to demonstrate the methodology and the CTF/Dakota coupling setup for use in uncertainty analysis. The preliminary results met expectations, and once experimental uncertainties are finalized in the benchmark specifications, the coupling will be used to evaluate uncertainties for various experimental conditions.

MCNP/CTF (Kutlu et al., 2022) and Serpent/CTF were coupled externally to provide a high-fidelity multi-physics tool for hexagonal geometry light water reactor designs, and the initial steady-state multi-physics capabilities of CTF were demonstrated. The coupled codes are managed via wrapping driver scripts, which process spatial mapping and other necessary steps to facilitate the exchange of thermal-hydraulics and neutronics data between MCNP and CTF or Serpent and CTF. The results of a VVER-1000 assembly model calculation were in good agreement with expectations and the reference cases. At present, both coupled codes have only been applied to assembly-scale models. The next step for future work may be to expand and optimize the coupled calculations to tackle full core pin-by-pin simulations. The MCNP and CTF preprocessors integrated in the couplings may be extended to generate full core inputs to enable this functionality.

To decrease the simulation time, introducing a domain decomposition method to simulate full core VVER models on a distributed architecture or reconfiguration to an internal coupling may be necessary. Future work may further include additional developments to extend the scope of the Serpent/CTF coupled code due to the usefulness of the Serpent multi-physics interface. The multi-physics tool can be further optimized and expanded for both cycle depletion and transient applications for pin-level full core calculations.

The focus of this paper is to summarize the cooperation activities between NCSU and the Institute for Nuclear Research and Energy (INRNE), Bulgarian Academy of Sciences, on model improvements, verification and validation, and uncertainty analysis of CTF for triangular/hexagonal lattice calculations in VVER applications. Pure thermal-hydraulic (single physics) as well as multi-physics developments and benchmarking at the nodal/channel and pin/sub-channel level have been discussed. Based on the results described in this paper, users will be confident of applying CTF to VVER design and safety analyses.

Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

Author contributions

YK: Investigation, Validation, Visualization, Formal Analysis, Writing – original draft. IS: Validation, Formal Analysis, Data curation, Investigation, Software, Writing – original draft. SM: Validation, Data curation, Resources, Investigation, Writing – original draft. PR: Visualization, Supervision, Writing – original draft, Software. AA: Validation, Methodology, Writing – original draft, Software. NK: Writing – review and editing, Conceptualization, Writing – original draft, Methodology. MA: Supervision, Investigation, Conceptualization, Writing – original draft, Software. KI: Resources, Conceptualization, Funding acquisition, Writing – review and editing, Supervision.

Funding

The author(s) declare that no financial support was received for the research and/or publication of this article.

References

Adams, B. M., Bauman, L. E., Bohnhoff, W. B., Dalbey, K. R., Eddy, J. P., Ebeida, M. S., et al. (2014). Dakota, a multilevel parallel object-oriented framework for design optimization, parameter estimation, uncertainty quantification, and sensitivity analysis: version 6.0 user's manual. SANDIA REPORT SAND2014-xxxx.

Aly Ahmed Mohamed Nabil Hassanein Multi-scale Multi-Physics Modeling of (2022). *Metallic fuel and thermal hydraulics of sodium fast reactors in a subchannel approach*. Raleigh: North Carolina State University.

Avramova, M., Ivanov, K., Velkov, K., Pasichnyk, I., Nikonorov, S., Denisova, M., et al. (2021). Benchmark on reactivity compensation with diluted boron by stepwise insertion of control rod cluster into the VVER-1000 core, Specification and Support Data. NEA/ENCGES.2021.0001.

Avramova, M., Velazquez-Lozada, A., and Rubin, A. (2012). Comparative analysis of CTF and trace thermal-hydraulic codes using OECD/NRC PSBT benchmark void distribution database. *Hindawi Publ. Corp. Sci. Technol. Nucl. Installations* 12. doi:10.1155/2013/725687

Bennett, A. (2015). *Development and Testing of a Coupled MCN6/CTF Code*. MS thesis. The Pennsylvania State University.

Beydogan, N. (2023). *The status of the KIT contribution to the rostov-2 VVER-1000 multi-physics transient benchmark*. Baden-Württemberg, Germany: Karlsruhe Institute of Technology (KIT).

Chanaron, B., Ahnert, C., Crouzet, N., Sanchez, V., Kolev, N., Marchand, O., et al. (2015). Advanced multi-physics simulation for reactor safety in the framework of the NURESAFE project. *Ann. Nucl. Energy* 84, 166–177. doi:10.1016/j.anucene.2014.12.013

Chen, J. C. (1966). Correlation for boiling heat transfer to saturated fluids in convective flow. *Industrial Eng. Chem. Process Des. Dev.* 5, 322–329. doi:10.1021/i260019a023

COBAYA team (2015). *COBAYA4 user's guide, UPM Report*. Madrid.

Delipei, G., Rouxelin, P., Abarca, A., Hou, J., Avramova, M., and Ivanov, K. (2022). CTF-PARCS core multi-physics computational framework for efficient lwr steady-state, depletion and transient uncertainty quantification. *Energies* 15 (14), 5226. doi:10.3390/en15145226

Dittus, F. W., and Boelter, L. M. K. (1930). "Heat transfer in automobile radiators for the tubular type," in *Publications in engineering* 2 (Berkeley: University of California), 443–461.

Downar, T., Xu, Y., Seker, V., and Hudson, N. (2018). *PARCS v3.3. 0 - U.S. NRC core neutronics simulator user manual*. University of Michigan, U.S. Nuclear Regulatory Commission.

Fillion, Ph., Bergeron, A., Gallo, D., Gregoire, O., Richebois, E., Royer, E., et al. (2011). "FLICA4 users guide," in *CEA saclay report*.

Garcia, M., Ferraro, D., Valtavirta, V., Tuominen, R., Imke, U., Leppänen, J., et al. (2020). Serpent2-SUBCHANFLOW pin-by-pin modelling capabilities for VVER geometries. *Ann. Nucl. Energy* 135 (106955), 106955. doi:10.1016/j.anucene.2019.106955

García-Herranz, N., Cuervo, D., Sabater, A., Rucabado, G., Sánchez-Cervera, S., and Cas tro, E. (2017). Multi-scale neutronics/thermal-hydraulics coupling with COBAYA4 code for pin-by-pin PWR transient analysis. *Nucl. Eng. Des.* 321, 38–47. doi:10.1016/j.nucengdes.2017.03.017

Hancox, W. T., and Nicoll, W. B. (1971). A general technique for the prediction of void distributions in nonsteady two-phase forced convection. *Int. J. Heat Mass Transf.* 14, 1377–1394. doi:10.1016/0017-9310(71)90186-4

Henry, R. (2022). *Perin yann et travlee anton generation of group constants at GRS for the rostov-2 benchmark PHYSOR 2022: international conference on physics of reactors*. Pittsburgh, USA: American Nuclear Society - ANS.

Hou, J., Delipei, G., Blyth, T., Porter, N., Avramova, M., Ivanov, K., et al. (2022). *Benchmark for Uncertainty Analysis in Modelling (UAM) for Design, Operation, and Safety of LWRs, Volume II: Specifications and Support Data for the Core Cases (Phase II)*. NEA/NSC/DOC.

Hou Jason Benchmark for Uncertainty Analysis in Modelling (UAM) for Design (2019). Operation and safety of LWRs, volume II: specifications and support data for the core cases.

Jens, W. H., and Lottes, P. A. (1951). "Analysis of heat transfer, burnout, pressure drop and density data for high-pressure water," *Tech. Rep. ANL-4627*. doi:10.2172/4421630

Jimenez, J. (2009). *COBAYA/FLICA4 coupling in Salome 5: EURATOM NURISP Project*, Technical note. UPM.

Jimenez (2015). Jimenez javier and Sanchez victor KIT results for the hot assembly BC problem using COBRA-TF and SUBCHANFLOW. Available online at: <https://publikationen.bibliothek.kit.edu/230104048/3817031>.

Kelly, D. J., Kelly, A. E., Aviles, B. N., Godfrey, A. T., Salko, R. K., and Collins, B. S. (2017). MC21/CTF and VERA multiphysics solutions to VERA core physics benchmark progression problems 6 and 7. *Nucl. Eng. Technol.* 49, 1326–1338. doi:10.1016/j.net.2017.07.016

Kolev, N. P., Aniel, S., Royer, E., Bieder, U., Popov, D., and Topalov, Ts. (2007). *VVER-1000 Coolant transient benchmark Phase II (V1000CT-2), Vol.1, Final specifications of the VVER-1000 vessel mixing problem*. Paris: NEA/NSC/DOC(2007) 22 ©OECD 2010.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The author(s) declared that they were an editorial board member of Frontiers, at the time of submission. This had no impact on the peer review process and the final decision.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Kolev, N.P., Petrov, N., Donov, J., Angelova, D., Aniel, S., Royer, E., et al. (2006). *VVER-1000 Coolant transient benchmark Phase II (V1000CT-2), Vol.2, Final specifications of the VVER-1000 MSLB problem*. Paris: NEA/NSC/DOC(2006)6 ©OECD 2010.

Kolev, N. P., Spasov, I., Zheleva, N., Todorova, G., Petrov, N., Mitkov, S., et al. (2016). Higher-resolution VVER MSLB simulation - final report on latest results from reference and advanced cases. *NURESAFE D14.41 Report, EURATOM*.

Kucukboyaci, V., Sung, Y., and Salko, R. (2015). "Cobra-TF Parallelization and Application to PWR Reactor Core Subchannel DNB Analysis," in ANS MC2015 - Joint International Conference on Mathematics and Computation (M&C), Supercomputing in Nuclear Applications (SNA) and the Monte Carlo (MC) Method, Nashville, TN, April 19–23, 2015.

Kutlu, Y. (2021). *Development of a coupled code between MCNP6 2 and CTF4.0 for VVER applications [MS thesis], nuclear engineering department*. North Carolina State University- Raleigh.

Kutlu, Y., Rouxelin, P., Avramova, M., and Ivanov, K. (2022). "Development of a high-fidelity coupling between MCNP6.2 and CTF4.0 for VVER applications," Proceedings of International Conference on Physics of Reactors 2022 (PHYSOR 2022), Pittsburgh, PA, May 15–20, 2022, 1802–1811. doi:10.13182/PHYSOR22-37770

Leppänen, J., Pusa, M., Viitanen, T., Valtavirta, V., and Kaltiaisenaho, T. (2015). The Serpent Monte Carlo code: status, development and applications in 2013. *Ann. Nucl. Energy* 82, 142–150. doi:10.1016/j.anucene.2014.08.024

Löetsch, A., Khalimonchuk, V., and Kuchin, A. (2010). "Corrections and additions to the proposal of a benchmark for core burnup calculations for a WWER-1000 reactor," in *Proceedings of the twentieth symposium of atomic energy research*.

Lozano, J.-A., García-Herranz, N., Ahnert, C., and Aragón, J. M. (2008). The analytic nodal diffusion solver ANDES in multi-groups for 3D rectangular geometry: development and performance analysis. *Ann. Nucl. Energy* 35, 2365–2374. doi:10.1016/j.anucene.2008.07.013

Lozano, J. A., Jiménez, J., García-Herranz, N., and Aragón, J. M. (2010). Extension of the analytic nodal diffusion solver ANDES to triangular-Z geometry and coupling with COBRA-IIIC for hexagonal core analysis. *Ann. Nucl. Energy* 37, 380–388. doi:10.1016/j.anucene.2009.12.001

Mitkov, S. (2022). *Simulation studies of VVER-1000 core processes for reactor safety analysis*. Sofia: Institute for Nuclear Research and Nuclear Energy. PhD Thesis.

Mitkov, S., Spasov, I., and Kolev, N. (2021). Thermal-hydraulic analysis of a VVER1000 core in MSLB conditions, *E3S Web Conf., Proc. 26th Sci. Conf. Power Eng. Power Mach. (PEPM), Sozopol, Bulg. E3S Web Conf.* 327, 01013, doi:10.1051/e3sconf/202132701013

Mitkov, S., Spasov, I., and Kolev, N. P. (2018). "Simulation of a hypothetical MSLB core transient in VVER-1000 with several stuck rods," 83. München: Carl Hanser Verlag, 389–395. doi:10.3139/124.110910

Papadionysiou Marianna (2022). *Coupling of nTRACER to COBRA-TF for full core high-fidelity analysis of VVERs//ANS international conference on physics of reactors 2022 (PHYSOR 2022)*. Pittsburgh, PA : ANS, 586–595.

Petrov, N., Spasov, I., Sanchez-Cervera, S., and Garcia-Herranz, N. (2015). Nodal level XS library v2 for VVER parameterized for MSLB. *NURESAFE D14.25b-Rev2 Report, EURATOM*.

Rubin, A., Schoedel, A., and Avramova, M. (2012). OECD/NRC benchmark based on NUREC pressurized water reactor sub-channel and bundle tests (PSBT) - Volume I: Experimental database and final problem specifications. *Nuclear Science NEA/NSC/DOC(2012)1*.

Salko, R., and Avramova, M. (2016). *CTF Theory Manual (ORNL/TM-2016/430)*, United States.

Sanchez, R., Zmijarevi, I., Coste-Delclaux, M., Masiello, E., Santandrea, S., Martinoli, E., et al. (2010). APOLLO2 Year 2010. *Nucl. Eng. Technol.* 42 (5), 474–499. doi:10.5516/net.2010.42.5.474

Santandrea, S., Sanchez, R., and Mosca, P. (2008). A linear surface characteristics approximation for neutron transport in unstructured meshes. *Nucl. Sci. Eng.* 160 (1), 23–40. doi:10.13182/nse07-69

Spasov, I., Kolev, N., Zheleva, N., Todorova, G., Jimenez, J., Herrero, J.J., et al. (2012). VVER MSLB Benchmark solutions with CRONOS/FLICA and COBAYA/FLICA. *NURISP- D3.1.3.3a Rep, EURATOM*.

Spasov, I., Ivanov, P., and Kolev, N. (2018). Simulation of a rod-ejection transient in VVER-1000 with COBAYA4-CTF. *Compt. rend. Acad. bulg. Sci.* 71 (3), 325–334. doi:10.7546/CRABS.2018.03.03

Spasov, I., Kolev, N. P., Donov, J., and Sabotinov, L. (2010). "CATHARE multi-1D modeling of coolant mixing in VVER-1000 for RIA analysis," *Sci. Technolgy Nucl. Installations (open access)*, 2010 1–11. doi:10.1155/2010/457094

Spasov, I., Mitkov, S., and Kolev, N. (2014). "Sub-channel FLICA4 and COBRA-TF input models qualified for VVER," in *NURESAFE D14.23 report, EURATOM*.

Spasov, I., Mitkov, S., Kolev, N. P., Sanchez-Cervera, S., Garcia-Herranz, N., Sabater, A., et al. (2017). Best-estimate simulation of a VVER MSLB core transient using the NURESIM platform codes. *Nucl. Eng. Des.* 321, 26–37. doi:10.1016/j.nucengdes.2017.03.032

Sung, Y., Kucukboyaci, V., Cao, L., and Salko, R. (2015). "COBRA-TF evaluation and application for PWR steam line break DNB analysis," Proceedings of NURETH-16Conference, Chicago, IL, United States.

Takasugi, C., Martin, N., Yoon, S., Bays, S., Avramova, M., and Ivanov, K. (2021). Development of a thermal hydraulic model for the versatile test reactor using the CTF subchannel code", Proceedings of Mathematics & Computation (M&C) 2021, Virtual Meeting, 2274–2283.

Takasugi, C. (2023). *Improvement of high-to-low fidelity multi-physics modeling for sodium-cooled fast reactor cores*. PhD Dissertation. Raleigh: North Carolina State University.

Werner, C. J. (2017). *MCNP® USER'S MANUAL, code version 6.2* October 27, 2017 Manual Rev. 0, LA-UR-17-29981.

Yan, J., Bi, Q., Liu, Z., Zhu, G., and Cai, L. (2015). Subcooled flow boiling heat transfer of water in a circular tube under high heat fluxes and high mass fluxes. *Fusion Eng. Des.* 100, 406–418. doi:10.1016/j.fusengdes.2015.07.007



OPEN ACCESS

EDITED BY

Gert Van den Eynde,
Belgian Nuclear Research Centre (SCK CEN),
Belgium

REVIEWED BY

Ugur Mertyurek,
Oak Ridge National Laboratory (DOE),
United States
Emil Løvbak,
Karlsruhe Institute of Technology (KIT),
Germany

*CORRESPONDENCE

Jackson R. Harter,
✉ jackson.harter@inl.gov

RECEIVED 14 May 2025

ACCEPTED 18 September 2025

PUBLISHED 08 October 2025

CITATION

Harter JR and DeHart MD (2025) Uncertainty quantification and sensitivity analysis of a nuclear thermal propulsion reactor startup sequence.

Front. Nucl. Eng. 4:1628866.
doi: 10.3389/fnuen.2025.1628866

COPYRIGHT

This work is authored by Jackson R. Harter and Mark D. DeHart, © 2025 Battelle Energy Alliance, LLC. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Uncertainty quantification and sensitivity analysis of a nuclear thermal propulsion reactor startup sequence

Jackson R. Harter^{1*} and Mark D. DeHart²

¹Reactor Physics and Shielding, Idaho National Laboratory, Idaho Falls, ID, United States, ²Department of Nuclear Science and Engineering, Abilene Christian University, Abilene, TX, United States

The research presented in this article describes progress in applying stochastic methods, uncertainty quantification, parametric studies, and variance-based sensitivity analysis (also known as Sobol sensitivity analysis) to a full-core model of a nuclear thermal propulsion (NTP) system simulated via the radiation transport code *Griffin* to simulate neutronics. Our goal is to develop a reduced-order (surrogate) model that can be rapidly sampled with perturbations to multiple input parameters. In this NTP system, reactivity and power feedback affect the rotation of control drums (CDs), which is itself controlled by a hybrid proportional-integral-derivative (PID) controller actuated by the power demand and reactivity feedback from the numerical model. This model uses reactor kinetic feedback (mean generation time [Λ]) and effective delayed neutron fraction [β_{eff}] from a transient *Griffin* simulation executed via *Griffin*'s improved quasi-static solver to provide the kinetic parameters) as inputs to functions that control the CD rotation angle. By investigating numerous stochastic approaches, we developed a dual-purpose surrogate model of the NTP system, using polynomial regression in the Multiphysics Object-Oriented Simulation Environment (MOOSE) Stochastic Tools Module (STM). The trained model can be rapidly sampled while simultaneously perturbing various input parameters, such as coefficients on the PID control or temperature (directly affecting the neutron cross section). The surrogate model delivers accurate (within 5%) results at speeds orders of magnitude faster (minutes, not days of computational time) than the base model. Once the surrogate model has been trained, distributions of the uncertain parameters can be changed at will to investigate the effects of perturbing multiple inputs as well as the effects of these inputs on the model output. For example, coefficients used in the PID control system may vary due to some type of physical interference, or uncertainty may exist in the temperature of the neutron cross sections in various regions of the reactor. A distribution can be placed on these parameters, and operational boundaries can be determined. The goal of this work is to support development of an advanced control system for operating CDs in a functioning NTP system. This work is a scoping study of the MOOSE STM.

KEYWORDS

nuclear thermal propulsion, sensitivity analysis, uncertainty quantification, instrumentation and control, autonomous control, nuclear systems, microreactors

1 Introduction

Supporting the development of space nuclear power for (1) nuclear fission power sources for spacecraft or lunar bases and for (2) spacecraft thrust is vital if humanity desires to regularly travel beyond the earth or moon (Buden, 2011). In terms of thrust, current chemical propulsion methods are limited by the amount of fuel they must carry, which restricts their operational ranges and efficiency. Nuclear power—particularly nuclear thermal propulsion (NTP)—offers a promising alternative by significantly enhancing the efficiency and capabilities of space missions. NTP affords a specific impulse (I_{sp}) roughly double that of the highest performing chemical rockets and an extremely high energy density by using molecular hydrogen as a propellant (Buden, 2011). Thus, NTP systems have emerged as a promising technology for deep space exploration. The potential of NTP systems to reduce travel times and increase payload capacities necessitates research to support this technology, particularly for missions extending all the way to Mars and beyond. However, the complexity of NTP systems, which involve intricate interactions between nuclear reactions, thermal dynamics, and mechanical controls, necessitates rigorous analyses and optimization to ensure their safe, efficient operation (Thomas, 2024; Niki et al., 2022; Myers et al., 2024).

Idaho National Laboratory (INL) has been at the forefront of the development of advanced reactor simulation tools to model and analyze the behavior of NTP systems (DeHart et al., 2022). One decisive challenge in this area is the precise control of reactivity and power levels within the reactor core. In existing concepts (Venneri et al., 2016), this is accomplished via the rotation of control drums (CDs) to adjust the neutron flux and, hence, the reactor power. The control system for these CDs must respond dynamically to varying power demands and reactivity conditions that are influenced by the thermal and neutronic characteristics of the reactor. To address this, advanced control systems such as ones that utilize proportional-integral-derivative (PID) controllers are employed. These controllers require robust tuning and validation to handle the nonlinearities and uncertainties inherent in NTP operations (Labouré et al., 2023).

To advance the reliability and performance of NTP systems, uncertainty quantification (UQ) and sensitivity analysis (SA) must be performed. By enabling assessment of how uncertainties in input parameters propagate through the model to affect the outputs, UQ identifies the key parameters that influence system performance. Alternatively, SA aids in understanding the output's dependency on input variations, as it is pivotal in optimizing control strategies and ensuring robust operation under varying conditions.

This work examines application of the Multiphysics Object-Oriented Simulation Environment (MOOSE) Stochastic Tools Module (STM) to the *Griffin* code (Wang et al., 2025; Slaughter et al., 2023) which was used to simulate a startup transient that was autonomously controlled with CDs through reactivity and power feedback via a PID controller. We investigated the effect of placing distributions on the PID control coefficients that affect power and reactivity. We explored the role that temperature changes play in NTP control as well as how CD movement is affected. We then developed a polynomial regression based surrogate model, using Latin Hypercube Sampling (McKay and Conover, 1979), of the NTP

startup transient. This model can be rapidly sampled to yield statistical quantities at speeds orders of magnitude faster than running the training model, yet still generate very similar outcomes. Finally, we performed a variance-based sensitivity analysis (e.g., Sobol analysis), a global method, of the model in order to examine the variance of a quantity of interest (CD angle, or control signal) (a PID coefficient) based on the influences of the other uncertain parameters. It is imperative to note that this work is a proof-of-principle effort to exercise the stochastic tools module in MOOSE on a complex, 3D simulation, the startup sequence of a nuclear thermal propulsion module. The number of perturbed parameters are low compared to other conventional studies. It is our intent to apply the knowledge gained from this study to even more complex, coupled multiphysics problems.

In Section 2, we describe the neutronics NTP model. Section 3 describes the stochastic methods and shows the results for each case. Finally, we present our concluding remarks in Section 4.

2 Nuclear thermal propulsion model

2.1 Nuclear thermal propulsion model

The development and optimization of NTP systems requires a comprehensive understanding of the reactor's dynamic behavior under various operating conditions. In the present work, UQ and SA were performed on an existing reactor model for NTP in a transient setting. This model was taken from a previous study by Labouré et al. (2023). The initial conditions were generated using point reactor kinetics and an eigenvalue calculation for a power increase transient; the transient simulates an initial burn and constant power throttle. The original model consisted of neutronics, heat conduction, and thermal hydraulic physics, but the present work examines only neutronic behavior. This focus enables detailed investigation of the neutronic response and control mechanisms, without the added complexity of thermal and hydraulic feedback. Isolating the neutronics component allows us to more effectively apply stochastic methods and SA to identify the key factors influencing reactor performance during transient operations. Two different device simulations were used to determine the control of the reactor. The first uses power feedback expressly to provide control drum (CD) rotation. As power demand increases, CDs turn to insert reactivity and raise the power. This model meets the demand, but undershoots the initial power demand. The hybrid controller uses both reactivity and power feedback to set the rotation angle of the CDs. It is much better suited to the task, meeting the initial power demand for the transient and matching it throughout.

2.1.1 Neutronics model

The full NTP simulation is a multiphysics model that incorporates neutronics, heat conduction, and thermal hydraulic physics modeled via *Griffin*/BISON/THM (Wang et al., 2025; Williamson et al., 2012; Hansel et al., 2024). The model considered in this study is a standalone neutronics model with no temperature or thermal fluids feedback; Figure 1 shows the full core and assembly model, respectively, developed with the Monte Carlo code Serpent2 (Leppänen, 2007). The cross sections for this model were generated using Serpent2, as was the full-core

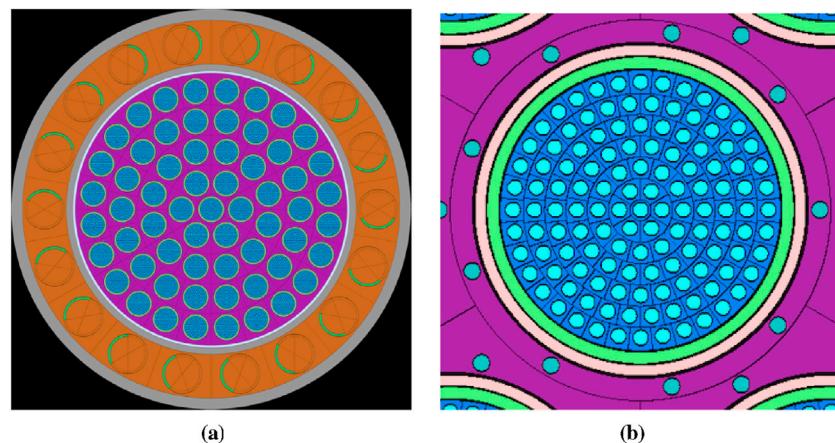


FIGURE 1
Overview of the core and assembly geometry in Serpent. Source: [Labouré et al. \(2023\)](#). **(a)** Radial view. **(b)** Zoom-in on assembly.

heterogeneous model described in [\(Labouré et al., 2023\)](#). The temperature range for the fuel was extended from a maximum of 2275 K–3175 K, and a reflector temperature of 300 K (conservatively simulating the boundaries of outer space) was added. Thus, the full range of temperatures is (1) fuel: 475–3175 K, (2) moderator: 150–705 K, and (3) reflector: 300 K. The moderator temperature range and CD angle (θ) range remained unchanged. In the original model, the multigroup cross-section library was tabulated with respect to the fuel and moderator temperatures (\bar{T}_{fuel} and \bar{T}_{mod} in K) and control drum angles in degrees, with 0 corresponding to maximum reactivity:

- $\bar{T}_{fuel} \in [475, 775, 1075, 1375, 1675, 1975, 2275]$
- $\bar{T}_{mod} \in [150, 335, 520, 705]$
- $\theta \in [0, 60, 120, 180]$

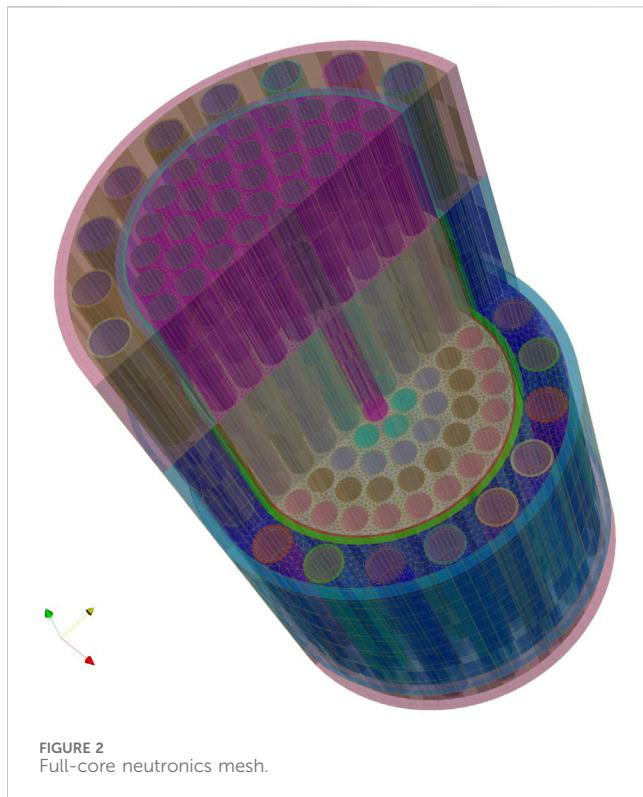
The neutronics model contains a total of 112 state points, each obtained by a Serpent2 calculation, performed with 2,000 cycles of 5×10^5 particles after 100 inactive cycles, yielding a standard deviation in the multiplication factor of < 3 pcm [\(Labouré et al., 2023\)](#). In *Griffin*, a multilinear interpolation scheme is used to compute cross-sections and superhomogenization (SPH) factors between temperature state points [\(Labouré et al., 2023\)](#). The feedback effects on the cross sections are isolated to fuel and moderator temperature, with other component temperatures and hydrogen density assigned with correlations and axial profiles [\(Labouré et al., 2023\)](#). These were determined by using a single representative fuel assembly and associated channels with varying total mass flow rates and power levels [\(Labouré et al., 2023\)](#). The Serpent model is not directly coupled to the multiphysics model and assumes these correlations and profiles are purely a function of average fuel and moderator temperatures [\(Labouré et al., 2023\)](#). The assessment of the accuracy of the current cross-section library between state points—in terms of errors due to (1) interpolation, (2) using assumed correlations and axial profiles, and (3) from the use of SPH factors—would need to be addressed in future work by [Labouré et al. \(2023\)](#) or others. All power is assumed to be instantly deposited, and the neutronics model only computes the assembly-

homogenized power density. This approach preserves the total energy but neglects radial dependency of the power density, which realistically should be higher at fuel periphery due to increase moderation from spatial self-shielding. Either a heterogeneous neutronics model or form functions would be needed to account for this effect [\(Labouré et al., 2023\)](#).

This model used constant temperatures during steady state and transient simulations. The energy group structure for the model was also kept the same. Instead of transport, diffusion was used in order to reduce the computational time and memory requirements. We used the continuous finite element method (CFEM) in concurrence with diffusion.

Also, combining the continuous finite element method with diffusion enables us to take advantage of the improved quasi-static (IQS) method to provide point kinetics parameters. IQS is a transient spatial kinetics method where space and time components of the neutron flux are represented by a time-dependent amplitude and a time-, space-, and energy-dependent shape [\(Prince and Ragusa, 2019\)](#). With this method, one can solve for kinetics parameters—reactivity (ρ), effective delayed neutron fraction (β_{eff}), and mean generation time (Λ)—which yield the required mechanisms used via power feedback with the PID controller to turn the CDs, thus bringing the reactor to the desired power level. These values are used as input in logic control functions, and the output is a change in CD angle, which is passed to *Griffin* via the MultiApp system in MOOSE, and *Griffin* uses these updated CD angles to turn the drums. The time duration for bringing the reactor to full power (set to 315 MWth) is set at 75 s, and the constant thrust is achieved at $t = 45$ s.

Existing models for the NTP steady state were used, consisting of two input files: (1) to solve for the adjoint solution, obtaining reactor physics kinetic parameters and (2) a forward solve to obtain a temperature distribution. A steady-state power of 500 kWth served as an arbitrary set point for simulating reactor startup. The converged steady-state solutions are used as initial conditions for the *Griffin* transient simulation; MOOSE has a restart system which enables this capability, the output solution is read in as an initial condition. [Figure 2](#) is a transparency of the full core mesh used in the



neutronics simulation. An important aspect to consider, and which was first addressed in [Labouré et al. \(2023\)](#), is that the adjoint calculation (required for the hybrid PID control) relies on an initial condition of temperature for the fuel and moderator regions [Note that the adjoint solution is used to compute time-dependent kinetics parameters]. However, since computation of these parameters ignores changes in temperature, the PID controller might compute a signal based on potentially inaccurate kinetics parameters ([Labouré et al., 2023](#)). In future work, a method of addressing this issue should be developed to ensure accuracy.

2.1.2 Demanded and measured power

[Labouré et al.](#) used an exponential ramp-up power benchmark in their research (detailed in ([Labouré et al., 2023](#))). We followed the same approach by using an initial power set point of $P_i = 0.5$ MWth and a final power set point of $P_f = 315$ MWth. The power demand ramp up was as follows:

$$P_d(t) = \begin{cases} P_i \left(\frac{P_f}{P_i} \right)^{\frac{t}{t_f}}, & 0 \leq t < t_f \\ P_f, & \text{otherwise,} \end{cases}$$

where $t_f = 30$ s is the time at which the ramp-up stops, and the power is constant. [Labouré et al.](#) used the following assumptions in their work:

- Initial CD angle was set to 120°
- Initial temperature was 500 K in the fuel and 200 K in the moderator
- Time step Δt was set to 0.1 s, which is used by the transient IQS algorithm

2.1.3 Power-driven PID control

Power-driven PID control is the simplest control strategy, relying on power demand (P_d) and measured power (P) to compute an error signal:

$$\varepsilon(t) = P_d - P.$$

The measured power is computed by the 3D full-core neutronics model. The rotation angle of the drums is described by:

$$\Delta\theta_P = K_p \varepsilon(t) + K_i \int_0^t \varepsilon(t') dt' + K_d \frac{d\varepsilon}{dt},$$

where K_p , K_i , and K_d are the proportional, integral, and derivative coefficients, respectively. From an initial CD angle of 120° (CDs facing in), the CDs rotate in reverse (outward) to insert reactivity. This is demonstrated by the trends in [Figures 3–5, 8](#) and the fact that the cross sections are computed for the four angle value statepoints. The CD angle is given by integrating the rotation angles with respect to time (at time step $n + 1$):

$$\theta_P(t_{n+1}) = \theta_i + \sum_{k=1}^n \Delta\theta_P(t_k), \quad (1)$$

where θ_i is the initial CD angle. The detailed results of this controller application are presented in [Labouré et al. \(2023\)](#), but a power-driven PID control is not the ideal approach. The coefficients must be tuned manually, and the “guess and check” method is relied upon—making optimization difficult. Furthermore, there is a power lag during the initial 40 s of the transient; as the generation is insufficient to meet the demand. There is also a power overshoot at the set point of constant power generation. Thus, the hybrid PID control method, discussed next, is the ideal solution, as it uses both power and reactivity (plus kinetics parameters) feedback to turn the CDs. A more detailed explanation is given by [Labouré et al. \(2023\)](#).

2.1.4 Hybrid feedback controlled model

[Labouré et al.](#) combined power-driven PID control with reactivity-driven PID control. Whereas the power-driven PID performed well in the later stages of the transient but poorly in the initial stage, the reactivity-driven PID performed well in the early stages but not the later ones. This led to the development of hybrid PID control, which combines power and reactivity (and kinetics data) [see [Figure 6](#) for an illustration].

The reactivity error signal is defined as:

$$\varepsilon(t) = \rho_d - \rho,$$

where ρ_d is the demanded reactivity and ρ is the measured reactivity. The demanded reactivity can be resolved by manipulation of the point kinetics equation:

$$\rho_d = \frac{\Lambda}{P_d \beta_{\text{eff}}} \left(\frac{dP_d}{dt} - \sum_i \lambda_i C_i + 1 \right),$$

where Λ is neutron mean generation time, β_{eff} is the effective delayed neutron fraction, and λ_i is the decay constant of the delayed neutron precursor group i with concentration C_i ([Labouré et al., 2023](#)). For the hybrid PID, the CD signal is computed via:

$$\Delta\theta = \xi \Delta\theta_P(t) + (1 - \xi) \Delta\theta_\rho(t),$$

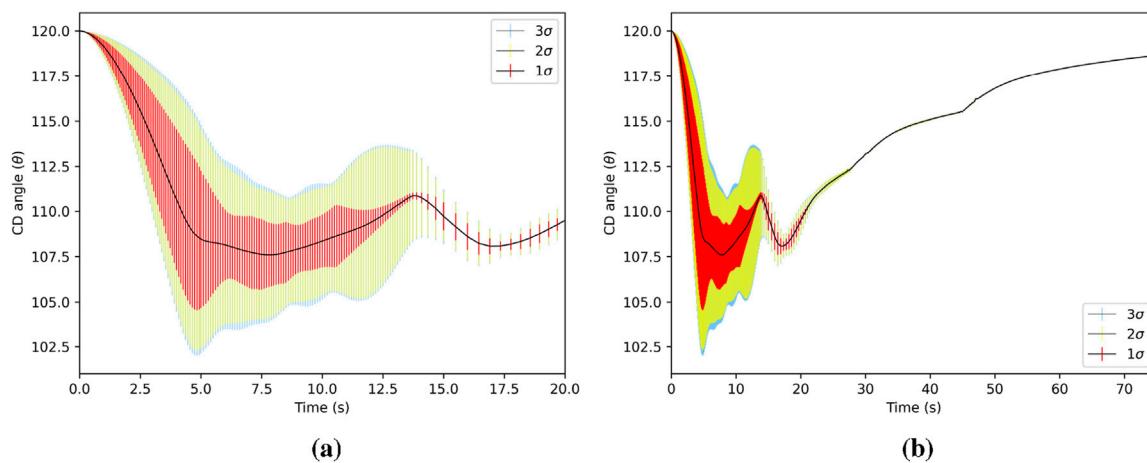


FIGURE 3
 Mean value of total CD angle (standard deviation is shown as vertical bars) out to 3σ for as a function of time. The 3σ encompasses the other intervals. The CD angle has the largest σ during the initialization, when reactivity is changing fast to meet the power demand. The total CD angle is the actual angle of the drum. (a) CD angle for $0 < t \leq 20$ s. (b) CD angle for $0 < t \leq 75$ s.

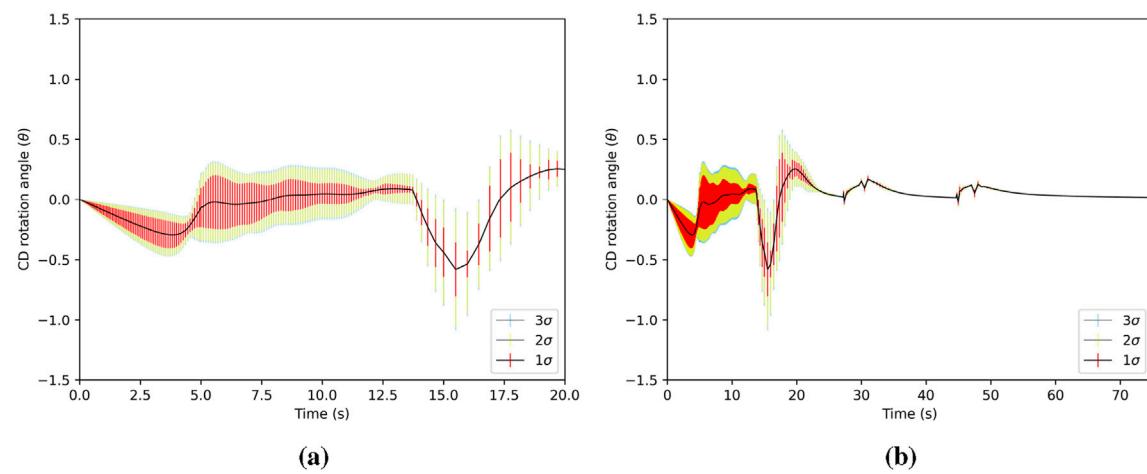


FIGURE 4
 Mean value of CD angle of rotation (standard deviation is shown as vertical bars) out to 3σ as a function of time. The rotation angle is the increment of rotation at each time step, added to the total drum angle. (a) CD rotation angle for $0 < t \leq 20$ s. (b) CD rotation angle for $0 < t \leq 75$ s.

where θ_p is the reactivity-based CD angle (computed similarly to Equation 1), and ξ is defined as:

$$\xi = \frac{\ln\left(\frac{P_f}{P_i}\right)}{\ln\left(\frac{P_f}{P_i}\right)}$$

where P_i and P_f are the initial and final powers, respectively. For the PID constants, the following values were used: $K_p = 10^{-7}$ deg/W and $K_i = K_d = 0$ (power-based constants); and $K'_p = 25$ deg/, $K'_i = 0.25$ deg/, and $K'_d = 0$ (reactivity-based constants). It is assumed many iterations were performed to settle on the final, constant, PID coefficients. The CD angle at time step $n + 1$ is given by:

$$\theta(t_{n+1}) = \theta_i + \sum_{k=1}^n \Delta\theta(t_k) = \theta_i + \sum_{k=1}^n \left(\xi \Delta\theta_p(t) + (1 - \xi) \Delta\theta_p(t) \right).$$

Results for the hybrid PID controller can be found in (Labouré et al., 2023). For the purpose of applying stochastic methods, we found that the hybrid PID controller performed well and is an ideal candidate for this work.

3 Stochastic methods and analysis approach

The objective of this work is to provide a reduced-order (surrogate) model of a neutronics system that can control an NTP module by automatically regulating CD movement based on coupled feedback effects.

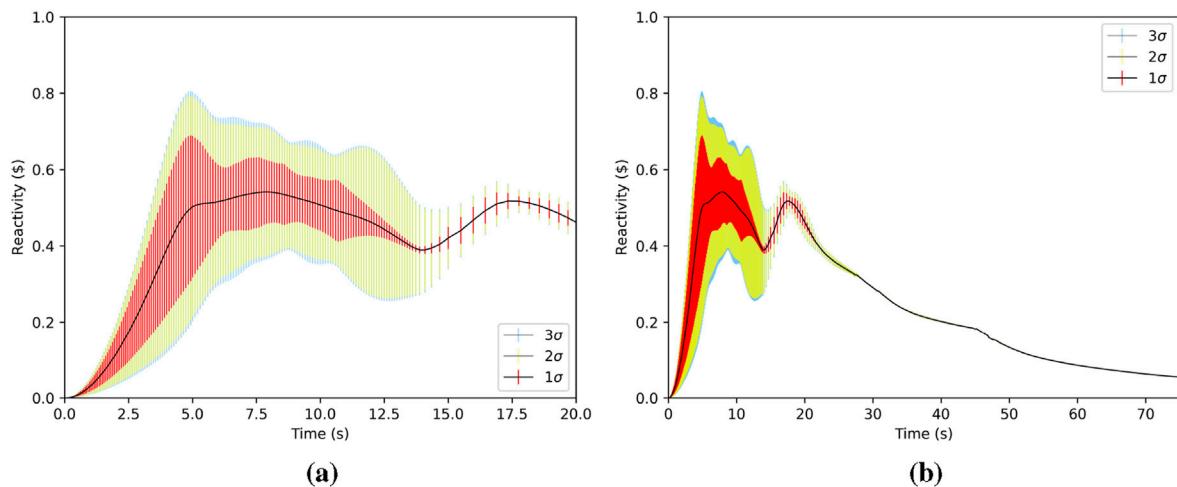


FIGURE 5
Mean value of reactivity (standard deviation is shown as vertical bars) out to 3σ as a function of time. Reactivity changes the most during the simulation initialization in order to meet the power demand. **(a)** Reactivity for $0 < t \leq 20$ s. **(b)** Reactivity for $0 < t \leq 75$ s.

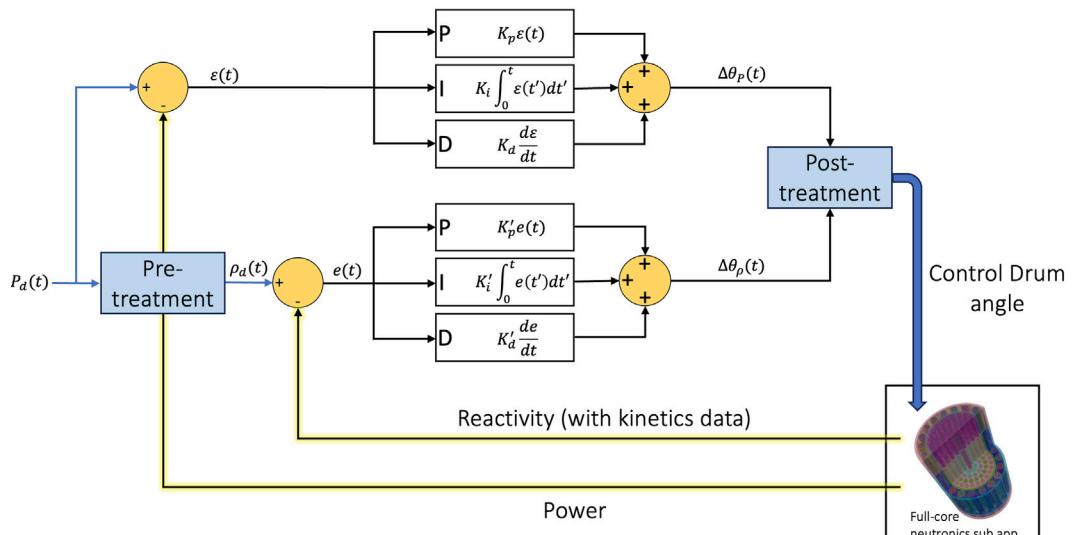


FIGURE 6
Diagram of the hybrid PID controller, showing feedback between the reactivity, power, and CD movement.

The desire to rapidly sample a surrogate, analyze sensitive parameters, perturb various quantities of interest (QoIs), and predict NTP behavior motivated the coupling of MOOSE's STM (Slaughter et al., 2023) to the *Griffin* code. UQ and SA methods have only recently become mature enough to handle large (10^9 [and greater] degrees of freedom) simulations; UQ at scale is difficult due to modeling large numbers of uncertain parameters and performing hundreds or thousands of perturbative simulations (Farcas et al., 2022).

In a simulation without STM (referred to hereafter as the “nominal” simulation, which is the case where the modeling parameters are their base [non-perturbed] value), the power or hybrid PID controller input (the main application [master app])

instantiates the *Griffin* transient simulation as a sub-application (sub-app), and the MultiApp system transfers governs the quantities passed back and forth between the master and sub-app. The master app sends the CD angle to the *Griffin* sub-app and receives the measured power (P_m), reactivity (ρ), mean generation time (Λ), effective delayed neutron fraction β_{eff} , and delayed neutron precursors ($\sum \lambda_i$). This process is repeated throughout the duration of the transient, and this feedback between controller and neutronics is what drives the power (thrust) generation of the transient.

When the *Griffin* neutronic simulation is coupled with the STM framework (in the sense of the MultiApp system in

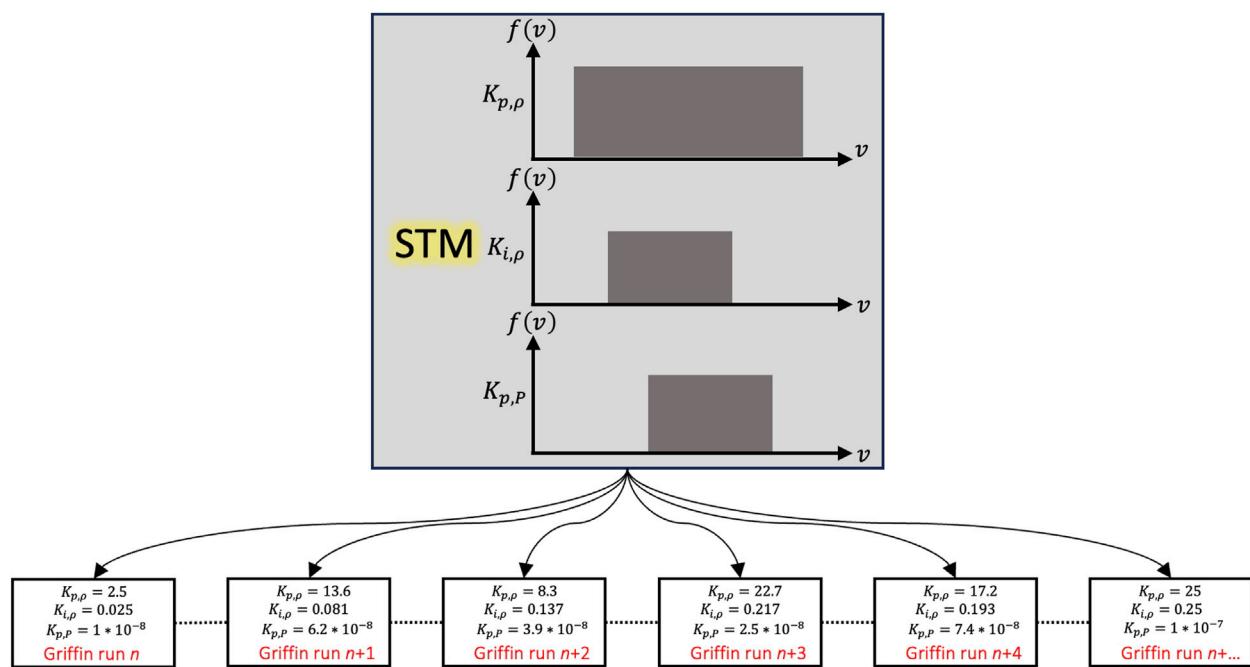


FIGURE 7

The STM input commands the flow of information to and from all the sub-apps used in the overall simulation process. In our application, uniform distributions were applied to parameters of interest and, depending on the number of samples desired, placed into subintervals. As the master application initiates, STM instantiates N instances of the *Griffin* base simulation. Each of these N instances has unique PID coefficient values. The base simulation transient is completed N times, and once these sets are finished their data are passed back to STM so statistical quantities such as the mean and standard deviation can be computed. This figure only shows how six distributions of PID coefficients can be used. STM is flexible, and those distributions can be changed out, or other parameter distributions can be included.

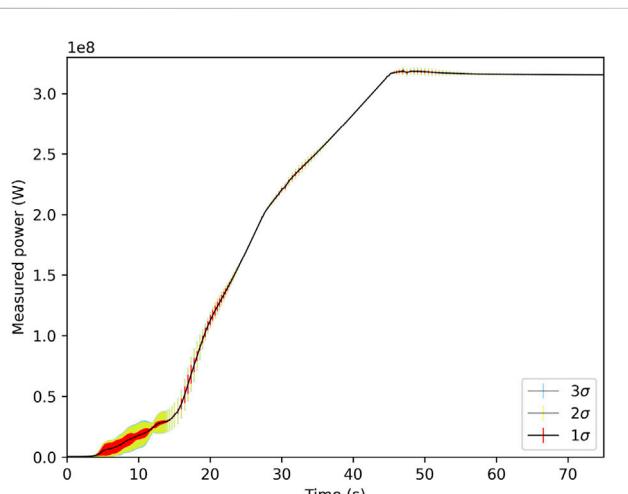


FIGURE 8

Measured power mean and standard deviation out to 3σ as a function of time. The power-driven PID controller does not meet the power demand during the first part of the transient, as was also shown by Labouré et al. (2023).

and sensitivity of the PID controller, and how that uncertainty propagates through and affects the motion of the CDs, in turn affecting power and reactivity. We designed a suite of STM inputs to (1) learn how to apply these tools to large and complex simulations, (2) develop a surrogate model for use in a non-nuclear advanced controls testbed reactor, and (3) create a UQ/SA methodology for future applications.

In order to efficiently capture the dynamics of this system, the STM suite was chosen with regards to it being a MOOSE application, which allows for seamless integration with other MOOSE applications (just one of the benefits to using MOOSE-based simulation tools). Other stochastic packages such as RAVEN (Alfonsoi et al., 2020) and Dakota (Adams et al., 2025) exist, but STM was selected exactly because it can integrate easily with other MOOSE-based codes. The stochastic approach is necessary to understand the influence of uncertainty in PID control parameters; physical components have an inherent uncertainty, whether it be through manufacturing defects, environmental effects, or user error. It is our task to understand just how important these uncertainties are.

We developed a protocol for applying stochastic methods to *Griffin* simulations, and ultimately we developed a surrogate model of an NTP module:

- Run the nominal *Griffin* power PID control and hybrid PID control simulations to understand the expected NTP behavior over the transient duration

MOOSE), the STM input becomes the master app, the PID input becomes a sub-app, and the *Griffin* transient becomes a sub-sub-app (e.g., parent, child, grandchild). This is illustrated in Figure 7. The focus of our investigation is to understand the uncertainty

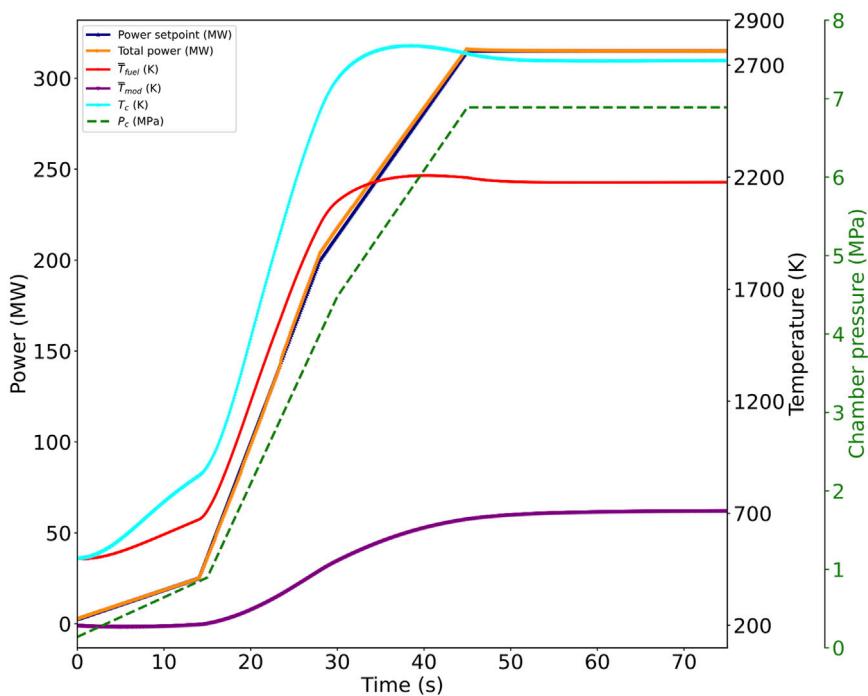


FIGURE 9
Evolution of power, temperature, and chamber pressure with a hybrid PID controller for the startup transient (Labouré et al., 2023).

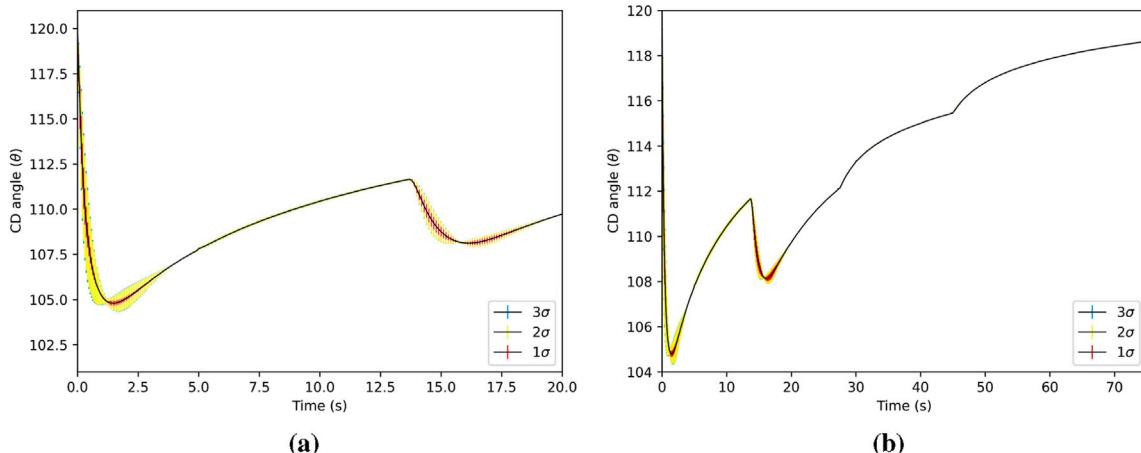


FIGURE 10
CD angle (standard deviation is shown as vertical bars). The CD angle has the largest standard deviation during the initialization, when reactivity is changing rapidly to meet the power demand. The total CD angle is the actual angle of the drum. **(a)** CD angle for $0 < t \leq 20$ s. **(b)** CD angle for $0 < t \leq 75$ s.

- Apply UQ to power and hybrid PID controller simulations for investigating parameters K_p , K_i , and K_d and their effects on CD angle, reactivity, and power
- Sample statistical quantities for the mean and standard deviation of the control parameter set
- Train a surrogate model, sample it, and change the perturbed parameters to test different configurations.
- Perform variance-based SA on the surrogate model

The total computational overhead of running N samples of a model depends on the time it takes for an individual simulation to be completed, taking into account some variance due to the uniform distribution of the sampled parameters. While a uniform distribution is applied on the quantities of interest, we do compute statistical quantities for UQ and SA investigation. For example, one sample might have a distribution of values that makes convergence more difficult or slow, as the amount of coefficients

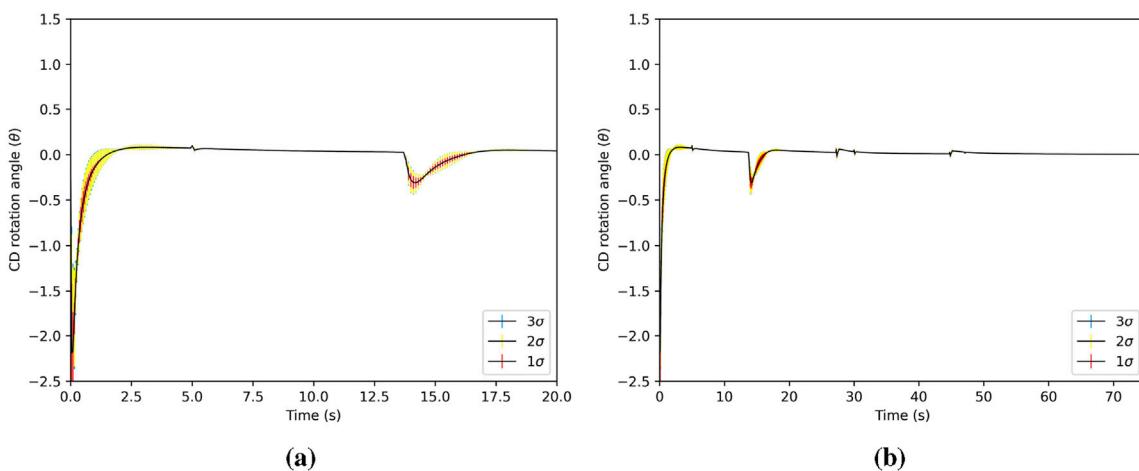


FIGURE 11
CD rotation angle (standard deviation is shown with vertical bars). The rotation angle is the increment of rotation at each time step, added to the total drum angle. **(a)** CD rotation angle for $0 < t \leq 20$ s. **(b)** CD rotation angle for $0 < t \leq 75$ s.

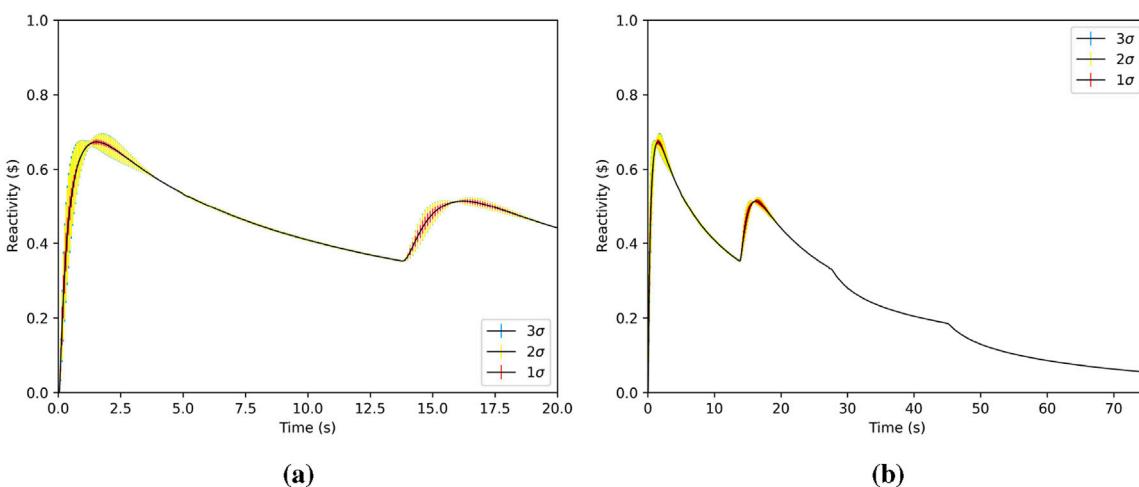


FIGURE 12
Reactivity (standard deviation is shown as vertical bars) as a function of time. Reactivity changes the most during the simulation initialization in order to meet the power demand. **(a)** Reactivity for $0 < t \leq 20$ s. **(b)** Reactivity for $0 < t \leq 75$ s.

tuned on a PID controller affects its behavior. With this in mind, we set up the sampler and applied Latin Hypercube Sampling to all simulations other than variance-based SA.

In an STM input, we specified the number of processes allocated to one sample. After brief trial and error, we deemed 192 processes per sample to be efficient. This is equal to four nodes (each node having 48 processes) on the Sawtooth High Performance Computing Cluster (HPCC) at INL, and it takes approximately 5.5 h to complete one transient hybrid PID-controlled neutronics simulation. We take advantage of parallel computing and require the STM simulation to run in batches. We selected 1,920 total processes and so ran them in groups of 10 (e.g., every 5.5 h, 10 samples were completed at a rate of approximately $1.8 \text{ samples} \cdot \text{hr}^{-1}$). However, the associated memory cost of this is also high. The *Griffin* finite element mesh has 135,000 elements and $4.5 \cdot 10^6$ degrees of

freedom. Ten of these simulations are held in memory for each batch. A Sawtooth compute node has 48 processes with 192 GB RAM. To make the simulations tractable, we needed to increase the amount of memory per process on a node, and selected 32 of the 48 processes per node. A standard run, as employed for a parameter study or to generate a surrogate, requires 100 to 150 samples, and uses 60 Sawtooth nodes over a duration of 36–120 h. For more detailed analyses such as variance-based SA, obtaining good statistics requires 300 Sawtooth nodes for 7 days.

3.1 Parameter sampling study

To understand how STM functions with *Griffin* as a sub-app, simple STM inputs were designed to observe how stochastic

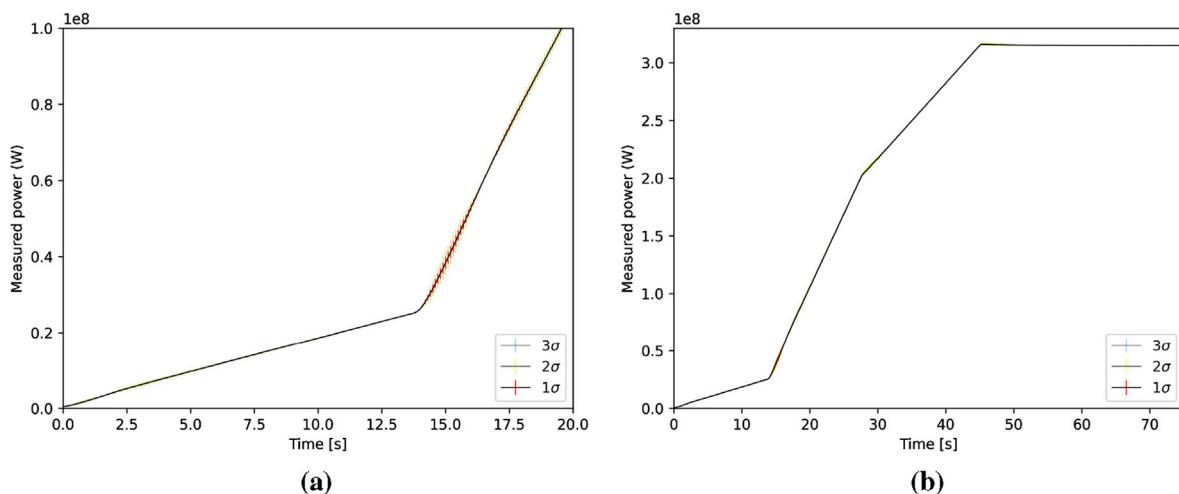


FIGURE 13
Power (standard deviation is shown as vertical bars) as a function of time. (a) Power for $0 < t \leq 20$ s. (b) Power for $0 < t \leq 75$ s.

TABLE 1 Surrogate sample sizes and times.

Samples	Time [s]
10	11
100	28
1,000	207
10,000	2,214

methods affect the QoI's of CD angle, reactivity, and power. We used Latin Hypercube Sampling for both models. The PID controls employed multiple logic functions informed by desired power set points to control the neutronics simulation. The CD rotation is determined by a linear combination function with three coefficients that act on error functions for the proportional, integral, and derivative values. This phase of investigation is essentially a parameter study; the ranges of values chosen for the PID coefficients reflect this. In the following section, uniform distributions are placed on various constant of the PID controller: K_p , K'_i , and K'_p . These distributions vary, respectively, by an order of magnitude each, an arbitrary choice for this study. This study introduces lower bounds (by an order of magnitude) to the initial values of the PID coefficients provided by Labouré et al. It is important to keep in mind that employing a uniform distribution greater than one order of magnitude will preferentially yield samples on the order of the upper bound. This could be studied in future work. This work is a proof-of-concept in the application of the stochastic methods functionality in MOOSE to *Griffin*, also a MOOSE-based code. There is no attempt to compare with experiment in this study; however, this is the intention for future work and to support other multiphysics coupling.

3.1.1 Power PID statistical model

For the power feedback PID controller, our model applied a uniform distribution on the proportional constant $K_p \in [10^{-8}, 10^{-7}]$, taking 100 samples. This is equivalent to running

the simulation 100 times with a different value of K_p in each simulation within the bounds of the distribution. Once the simulations are complete, statistics are computed in the form of moments. The first moment is the mean (μ), and the second moment is the variance (standard deviation) [σ]. We show the variation in the results using a three standard deviation band in all following figures. The variation indicated by 3σ is much tighter. The color scheme is portrayed in the figures; red = 1σ , yellow = 2σ and blue = 3σ . In some cases, the difference between 2σ and 3σ can be almost imperceptible, showing almost no overlap between the yellow and blue coloring.

3.1.2 Hybrid PID statistical model

We switched from the power feedback model to the hybrid feedback model, which uses kinetic parameters (Λ , β_{eff}) from the *Griffin* simulation to control functions in the input, as well as additional coefficients for the PID controller. Being the simplest of the available model types, the power feedback model was first used with a single parameter to which to assign an uncertain distribution. The hybrid model offers more control parameters to account for reactivity and power signals, and fits the measured power almost exactly to the demanded power. Here, we applied uniform distributions on the parameters: $K'_p \in [2.5, 25]$, $K'_i \in [0.025, 0.25]$, and $K_p \in [10^{-8}, 10^{-7}]$. We took 100 samples (equivalent to running the simulation 100 times), which is computationally expensive in terms of both resources and time duration. The first moment is the mean (μ), and the second moment is the standard deviation (σ). Figure 9 shows the power setpoint (black) and delivered power (orange) from the previous work (Labouré et al., 2023). Our objective is to match that curve in the current work and provide the relevant statistical quantities. For each QoI, we report the mean and multiples of the standard deviation out to 3σ . These QoIs are shown in Figures 10–13.

It is clear in comparing the two controllers that the hybrid PID controller outperforms the power-driven PID controller. The fact that we generate better statistics with the hybrid PID is attributed to the superior control it affords by using both reactivity and power feedback. Even though the statistical error decreases as $1/\sqrt{N}$, where

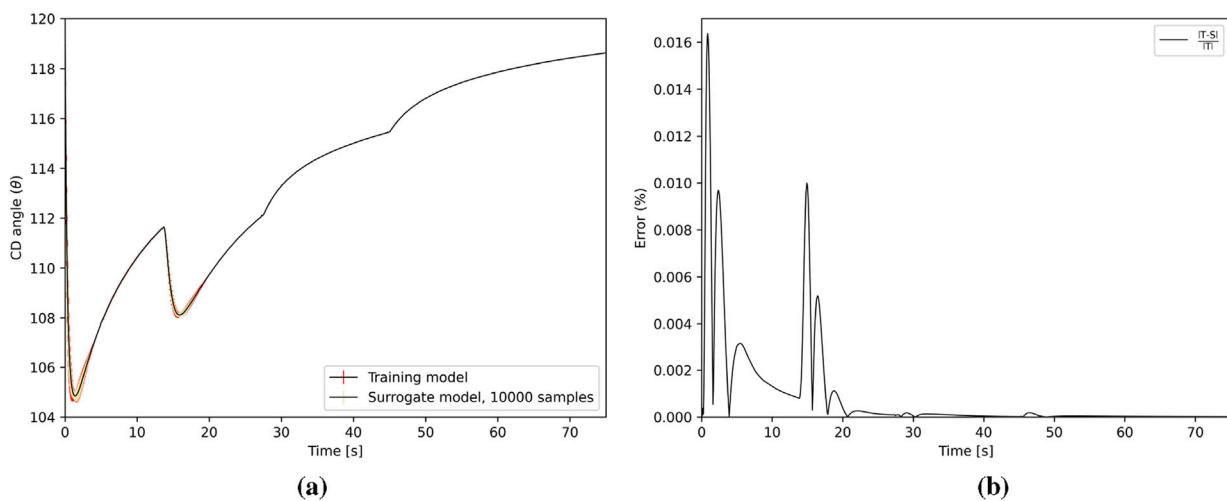


FIGURE 14
CD angle μ and σ (a) shown alongside the approximation error (in %) which is the difference between the mean of the training set and the surrogate model with 10,000 samples. The CD angle has the largest error during the initial phase up to ~ 20 s. (a) CD angle for $0 < t \leq 75$ s. (b) CD angle error.

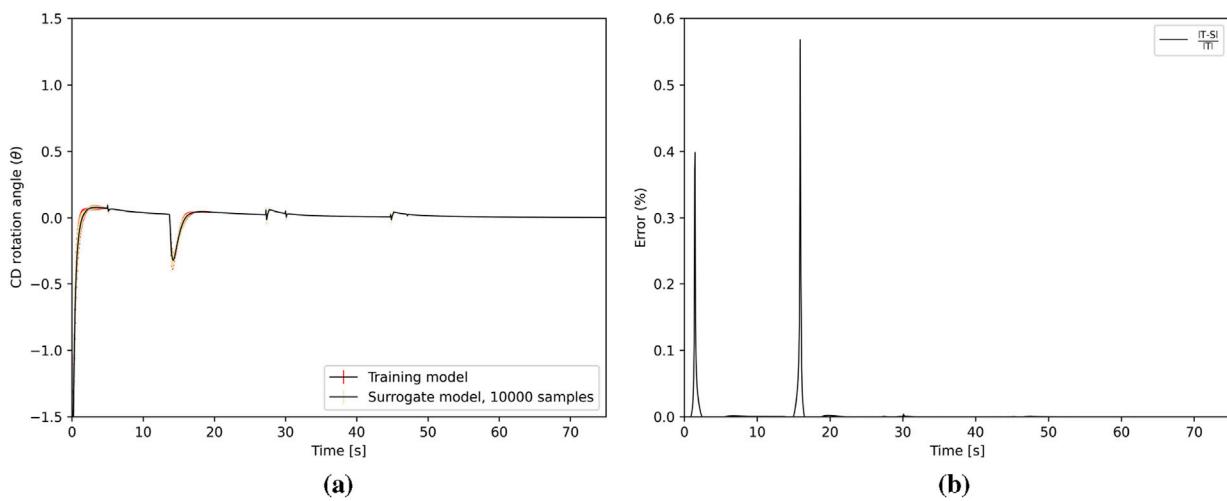


FIGURE 15
CD rotation angle μ and σ (a) shown alongside the approximation error (in %) which is the difference between the mean of the training set and the surrogate model with 10,000 samples. The rotation angle has the largest error at initialization and at ~ 15 s. The rotation angle is the increment of rotation at each time step. (a) CD rotation angle for $0 < t \leq 75$ s. (b) CD rotation angle error.

N is the number of samples, the hybrid PID shows much better performance, and the standard deviation at 3σ is much tighter even with fewer required samples.

3.2 Polynomial regression surrogate

In this study, we developed a polynomial regression surrogate as a reduced order representation of the time-dependent *Griffin* model. The polynomial regression surrogate is a full multidimensional polynomial expansion with all the cross terms, and with the number of dimensions being described by the number of distributions we use.

We used polynomials of order 4 in combination with an ordinary least-squares-type regression. Typically, a higher-order polynomial can fit the training data better, but it can also lead to overfitting, where the model captures noise instead of the underlying trend. Conversely, a lower-order polynomial may underfit the data, not capturing the complexity of the relationship. A polynomial order refinement study was not performed for this case; the selection of polynomial order 4 was an arbitrary choice for this model. However, future work could entail an order refinement study, in which regression surrogate models are tested for accuracy by varying polynomial order (Mukhtar et al., 2023; Zhao et al., 2022).

The problem setup is similar to that in the parameter sampling study, in that we use the same uniform distributions on the variables

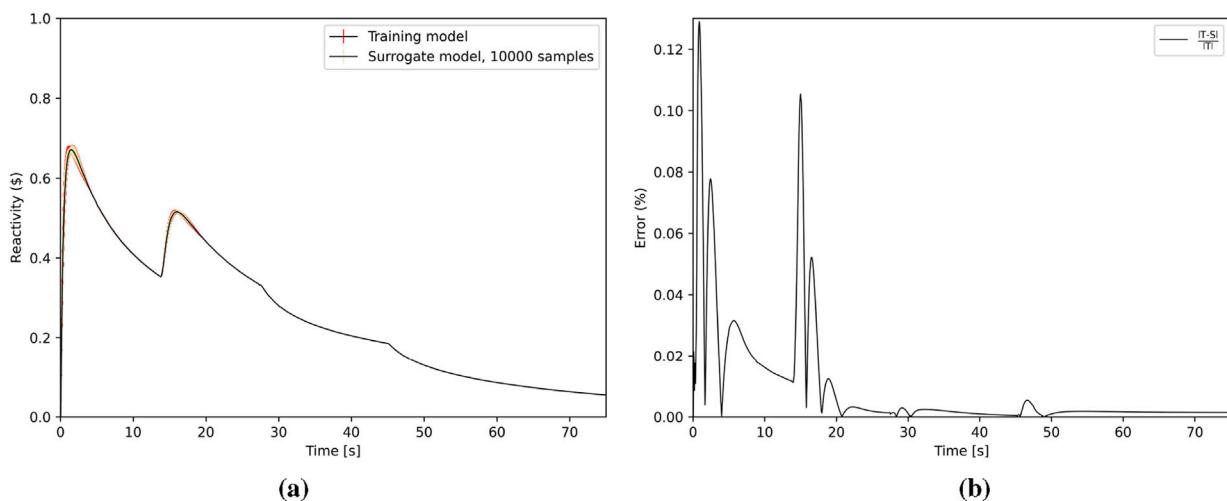


FIGURE 16
Reactivity μ and σ (a) shown alongside the approximation error (in %) which is the difference between the mean of the training set and the surrogate model with 10,000 samples. The reactivity changes the most during the simulation initialization in order to meet the power demand. (a) Reactivity for $0 < t \leq 75$ s. (b) Reactivity error.

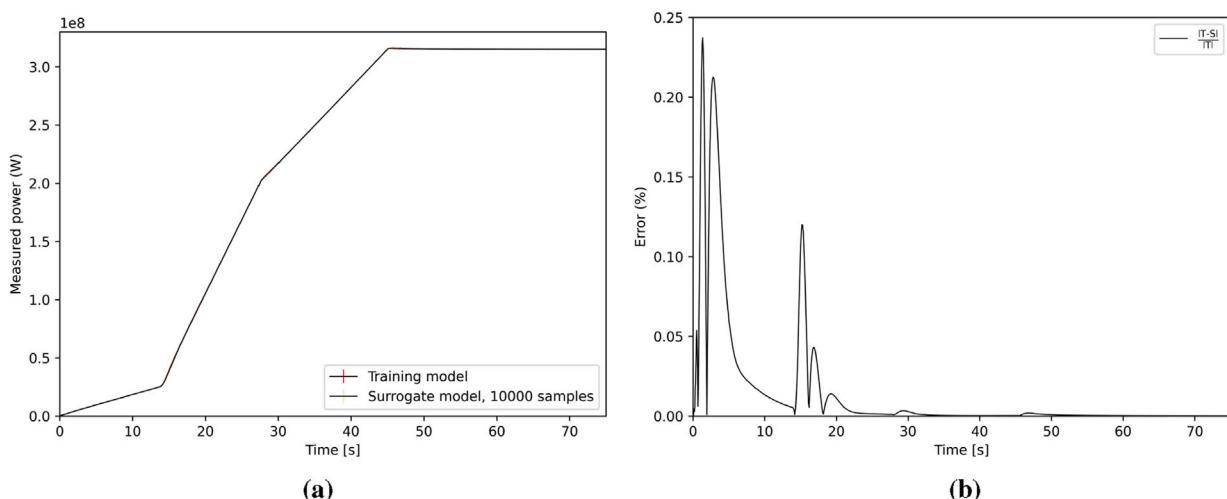


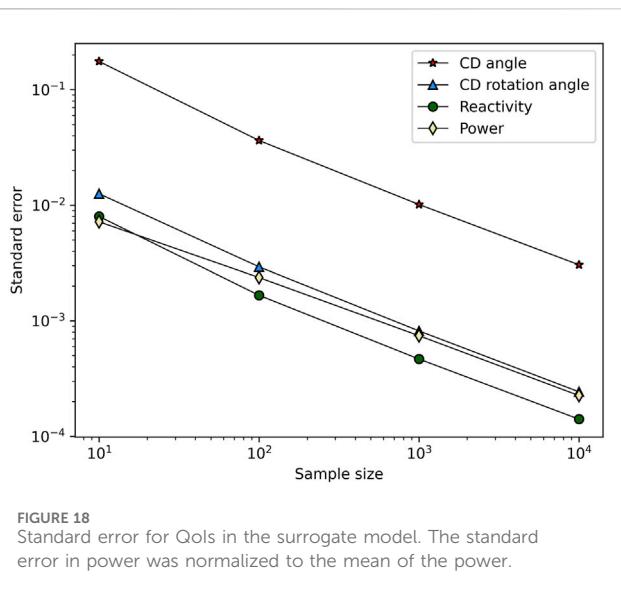
FIGURE 17
Power μ and σ (a) shown alongside the approximation error (in %) which is the difference between the mean of the training set and the surrogate model with 10,000 samples. Most of the error is occurring at the outset, when reactivity and CD angle are rapidly changing to meet initial demand. (a) Measured power for $0 < t \leq 75$ s. (b) Measured power mean error.

we know to be uncertain. The additional input directs the stochastic tools to train a surrogate model for the output we are interested in: CD angle, CD rotation angle, reactivity, and reactor power. This creates a library that can be sampled using a separate input file, can change the bounds of the uniform distribution (as long as they are within the original set [note that extrapolation is not recommended and accurate results not guaranteed]), and are sampled at rates orders of magnitude faster to give accurate results.

We compared the results of the training model against those of the surrogate. The training model took approximately 96 h to run on 60 Sawtooth nodes, using 150 samples to generate a training set. This

amount of samples was sufficient to capture a range of one order of magnitude for each of the PID parameters; a larger range and number of parameters would merit an increase in the number of samples taken. The surrogate model sampling times are reported in Table 1. In general, the surrogate model can be sampled rapidly with little loss of accuracy in comparison to the training model.

We compared various surrogate sample sizes (10, 100, 1,000, and 10,000 samples) versus the training set (150 samples), and reported μ and σ out to 3σ . We examined the CD angle, CD rotation angle, reactivity, and measured power, the goal being to determine the level of consistency between the training and surrogate models. As



expected, we found that σ converges as the number of samples increases. In Figures 14–17, we show (on the left) the mean and standard deviation of the training set and the surrogate model. On the right, we show the error of the difference of the means, computed via

$$\delta = \frac{v_T - v_S}{v_T} \times 100,$$

where v_T is the training set, and v_S is the polynomial (surrogate model). In all cases, with the exception of CD rotation angle the error is fractional. There is good agreement between the two models. The regions where error is largest, are regions of rapid change during

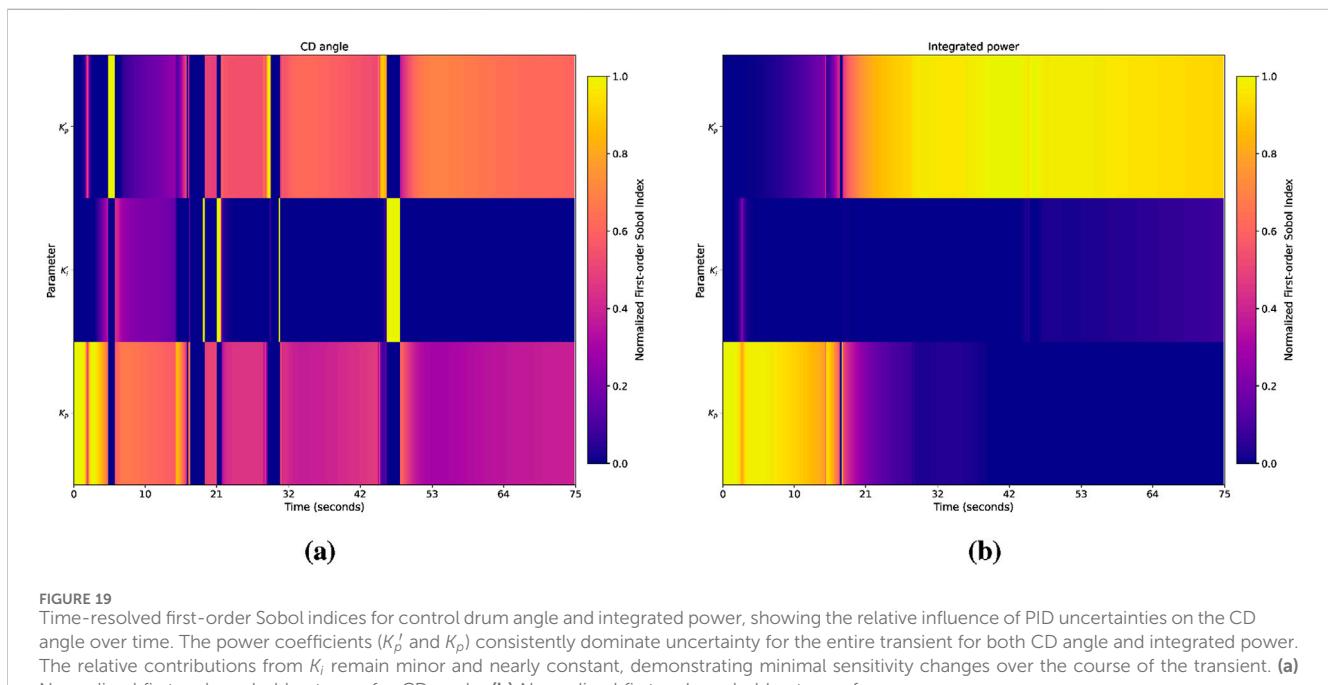
the startup transient. This effect is correlated with the mean and standard deviation of those QoIs in those regions, as well. The surrogate σ converges within the training σ , which is expected, as the surrogate model can be sampled rapidly in comparison to the training model. Once the initial model is trained to develop the surrogate, if parameter change is desired, the surrogate can be rapidly sampled; the change in value on a parameter's distribution can easily be adapted to a region of interest. However, if this region of interest is outside the initial model's sampling space, the accuracy of the new result is not guaranteed. Development of a surrogate model for the NTP startup sequence, based on polynomial regression, can reproduce the training model output within an approximate 5% deviation, and, at only a fraction of the time duration needed to run the full model.

The standard error (SE, the relative error in this study), in the standard deviation decreases monotonically as the number of samples increase, shown in Figure 18 and calculated with the relation in Equation 2 using the total number of samples (N) and the standard deviation σ .

$$SE = \frac{\sigma}{\sqrt{N}}. \quad (2)$$

3.3 Variance-based (sobol) sensitivity analysis

The Sobol sampling scheme consists of using a sample and re-sample matrix to create a series of matrices that can be used to compute first-order, second-order, and total-effect sensitivity indices. We restrict our analysis to first-order effects in this study. The Sobol method was chosen as the primary approach to quantify how input uncertainties impact the variation in the key



QoI's; this method decomposes the variance of model outputs into contributions from each uncertain input and their interactions, offering a robust measure of how uncertainties propagate through complex systems (Saltelli, 2002).

We perform Sobol analysis by evaluating the surrogate model. The reason we choose to evaluate the surrogate model rather than the stochastic model in the previous section is because the hundreds of thousands of transient simulation runs would be computationally intractable; leveraging the efficiency of the surrogate, it is much easier to navigate the parameter space. The surrogate-based sensitivity analysis is implemented in the STM input file, very similar to the power and hybrid PID methodology.

The first order indices are the portion of the variance in the QoI (e.g., K_p') due to the variance of the uncertain parameters (K_i' , K_p). These indices, calculated at user-specified confidence intervals, quantify both the direct influence of each uncertain parameter and its interactions with other parameters on the variation of PID coefficients and the integrated core power. This analysis highlights the most significant contributors to performance variability, offering critical insights into the reactor's thermal performance under uncertainty (Osman, 2025).

First-order Sobol indices were computed using 50,000 evaluations of polynomial regression surrogate models trained on the selected uncertain parameters. The resulting time-resolved indices illustrate how variations in PID coefficients influence system behavior throughout the startup transient. This study is one of few transient, three-dimensional simulations of a complex system model (Guo et al., 2023). Our goal in applying Sobol sensitivity analysis is to understand the sensitivity of the QoI during the transient; how the sensitivities change over the duration.

The results in Figure 19 are represented as the evolution of first-order Sobol indices over time for CD angle (Figure 19a) and integrated power (Figure 19a). The focus on first-order indices allows for the segregation of each parameter's contribution to output variance without uniting interaction terms. In Figure 19a, the heatmap shows the normalized first-order Sobol index for the CD angle. The prominent vertical bars in Figure 19a correspond to the CD angle error in Figure 14b; these correspond to regions of rapid change in the transient where quick movement of CDs occurs—an inflection point, or rapid change in slope. A similar, more subtle effect is shown in Figure 19b. Though total Sobol indices capture both the main effects and interactions of QoIs, term these were excluded to streamline the analysis.

4 Conclusion

Application of stochastic methods to *Griffin* simulations has proven successful in determining the specific parameters to which hybrid PID controllers are sensitive. Placing uniform distributions on the coefficient parameters and extracting statistical quantities of mean and standard deviation show how sensitive the PID controller is to subtle changes. We restricted our application of STM to a neutronics simulation as a first trial of stochastic methods on transient reactor physics simulation. In the future, we envision further multiphysics coupling, to heat conduction, thermal hydraulics, or heat pipes. This work was an initial study to test effectiveness.

Using a reduced-order model to simulate reactor startup affords many potential benefits, enabling rapid throughput analyses of different parameter variations so as to determine the “best fit” values for achieving precision startup control. Initially, we planned to use polynomial chaos expansions (PCEs) to develop the surrogate model, as the authors are experienced in using this method. However, at this stage of STM development, the PCE method cannot generate a time-dependent surrogate. The advantage of PCEs is that the quadrature-based sampling uses fewer samples than Monte Carlo or Latin Hypercube Sampling, for the same accuracy.

An area of concern in practical application of this method is the hybrid approach itself, which uses power and kinetic feedback to control the startup sequence. The terms we use for the feedback (Λ , β_{eff}) are not measurable in an operating reactor system. Work is ongoing to address this issue; we are currently developing a deep reinforcement learning (DRL) (Cebollada et al., 2021) model—a proximal policy optimization method (Fan et al., 2025)—for the NTP system, which is a combination of deep neural networks and reinforcement-based machine learning and uses feedback, rewards, and penalties to learn. A DRL model may prove a constructive approach to this specific problem—the control drum rotation angle is what the neutronic system requires based on parameterization of the PID coefficients. Removing, and reducing the amount of user-controlled “knobs” and relying on a proven artificial intelligence algorithm may increase the efficiency and accuracy of the control system.

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Author contributions

JH: Visualization, Formal Analysis, Validation, Data curation, Conceptualization, Writing – review and editing, Software, Writing – original draft, Investigation, Methodology. MD: Supervision, Project administration, Funding acquisition, Writing – original draft, Resources.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. This manuscript has been authored by Battelle Energy Alliance, LLC under contract no.~DE-AC07-05ID14517 with the U.S. Department of Energy, and this contract was funded by NASA's Space Nuclear Propulsion (SNP) project within the NASA Space Technology Mission Directorate (STMD).

Acknowledgments

The authors thank Peter German for his guidance and invaluable discussions, and also thank Wafaa Osman for help with postprocessing scripting. This research made use of Idaho National Laboratory's High Performance Computing systems

located at the Collaborative Computing Center and supported by the Office of Nuclear Energy of the U.S. Department of Energy and the Nuclear Science User Facilities under contract no. DE-AC07-05ID14517.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The author(s) declared that they were an editorial board member of *Frontiers*, at the time of submission. This had no impact on the peer review process and the final decision.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Any alternative text (alt text) provided alongside figures in this article has been generated by *Frontiers* with the support of artificial intelligence and reasonable efforts have been made to ensure

accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Licenses and permissions

This manuscript has been authored by Battelle Energy Alliance, LLC under contract no. DE-AC07-05ID14517 with the U.S. Department of Energy. The U.S. Government retains and the publisher, by accepting the article for publication, acknowledges that the U.S. Government retains a nonexclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for U.S. Government purposes.

References

Adams, B. M., Bohnhoff, W. J., Dalbey, K. R., Ebeida, M. S., Eddy, J. P., Eldred, M. S., et al. (2025). *Dakota 6.21.0 documentation*. Albuquerque, NM: Sandia National Laboratory.

Alfonsi, A., Rabiti, C., Mandelli, D., Cogliati, J., Wang, C., Talbot, P. W., et al. (2020). *RAVEN theory manual*. Idaho Falls, ID: Idaho National Laboratory. Technical report.

Buden, D. (2011). Nuclear thermal propulsion systems. *Polaris Books*.

Cebollada, S., Payá, L., Flores, M., Peidró, A., and Reinoso, O. (2021). A state-of-the-art review on mobile robotics tasks using artificial intelligence and visual data. *Expert Syst. Appl.* 167, 114195. doi:10.1016/j.eswa.2020.114195

DeHart, M. D., Schunert, S., and Labouré, V. M. (2022). "Nuclear thermal propulsion" in *Nuclear reactors - spacecraft propulsion, research reactors, and reactor analysis topics*. Editor C. L. Pope (United Kingdom: IntechOpen). Available online at: <https://www.intechopen.com/chapters/103895>.

Fan, T., Liu, L., Yue, Y., Chen, J., Wang, C., Yu, Q., et al. (2025). Truncated proximal policy optimization. Available online at: <https://arxiv.org/abs/2506.15050>.

Farcas, I., Merlo, G., and Jenko, F. (2022). A general framework for quantifying uncertainty at scale. *Commun. Eng.* 1, 43. doi:10.1038/s44172-022-00045-0

Guo, H., and Xiaolong Fu, X. Z., Zhu, Y., and Rabczuk, T. (2023). Physics-informed deep learning for three-dimensional transient heat transfer analysis of functionally graded materials. *Comput. Mech.* 72, 513–524. doi:10.1007/s00466-023-02287-x

Hansel, J., Andrs, D., Charlot, L., and Giudicelli, G. (2024). The MOOSE thermal hydraulics module. *J. Open Source Softw.* 9, 6146. doi:10.21105/joss.06146

Labouré, V., Schunert, S., Terlizzi, S., Prince, Z., Ortensi, J., Lin, C.-S., et al. (2023). Automated power-following control for nuclear thermal propulsion startup and shutdown using MOOSE-based applications. *Prog. Nucl. Energy* 161, 104710. doi:10.1016/j.pnucene.2023.104710

Leppänen, J. (2007). *Development of a new Monte Carlo reactor physics code*. (Ph.D. thesis). Helsinki University of Technology.

McKay, R. J. B. M. D., and Conover, W. J. (1979). Comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* 21 (2), 239–245. doi:10.1080/00401706.1979.10489755

Mukhtar, A., Yasir, A. S. H. M., and Nasir, M. F. M. (2023). A machine learning-based comparative analysis of surrogate models for design optimisation in computational fluid dynamics. *Helijon* 9 (8), e18674. doi:10.1016/j.helijon.2023.e18674

Myers, R., DeHart, M. D., and Kotlyar, D. (2024). Integrated steady-state system package for nuclear thermal propulsion analysis using multi-dimensional thermal hydraulics and dimensionless turbopump treatment. *Energies*, 17 (13), 3068. doi:10.3390/en17133068

Nikitaev, D., and Thomas, L. D. (2022). Preliminary results for in-situ alternative propellants for Nuclear thermal propulsion. *Nucl. Technol.* 208, 96–106. doi:10.1080/00295450.2021.2021768

Osman, W. (2025). *Uncertainty quantification and sensitivity analysis of thermal property changes under irradiation in heat-pipe microreactors using MOOSE*. (Ph.D. thesis). North Carolina State University.

Prince, Z. M., and Ragusa, J. C. (2019). Multiphysics reactor-core simulations using the improved quasi-static method. *Ann. Nucl. Energy* 125, 186–200. doi:10.1016/j.anucene.2018.10.056

Saltelli, A. (2002). Making best use of model evaluations to compute sensitivity indices. *Comput. Phys. Commun.* 145, 280–297. doi:10.1016/s0010-4655(02)00280-1

Slaughter, A., Prince, Z., German, P., Halvick, I., Jiang, W., Spencer, B., et al. (2023). MOOSE Stochastic Tools: a module for performing parallel, memory-efficient *in situ* stochastic simulations. *SoftwareX* 22, 101345. doi:10.1016/j.softx.2023.101345

Thomas, D. (2024). Nuclear thermal propulsion – progress and potential. *J. Space Saf. Eng.* 11 (2), 362–373. doi:10.1016/j.jssse.2024.04.001

Venneri, P. F., Eades, M., and Kim, Y. (2016). "Accident-tolerant control drums applied to nuclear thermal propulsion," in *Transactions of the Korean nuclear society autumn meeting. Korean nuclear society, gyeongju, korea*.

Wang, Y., Prince, Z. M., Park, H., Calvin, O. W., Choi, N., Jung, Y. S., et al. (2025). Griffin: a MOOSE-based reactor physics application for multiphysics simulation of advanced nuclear reactors. *Ann. Nucl. Energy* 211, 110917. doi:10.1016/j.anucene.2024.110917

Williamson, R., Hales, J., Novascone, S., Tonks, M., Gaston, D., Permann, C., et al. (2012). Multidimensional multiphysics simulation of nuclear fuel behavior. *J. Nucl. Mater.* 423 (1), 149–163. doi:10.1016/j.jnucmat.2012.01.012

Zhao, Y., Ye, S., Chen, X., Xia, Y., and Zheng, X. (2022). Polynomial Response Surface based on basis function selection by multitask optimization and ensemble modeling. *Complex Intell. Syst.* 8, 1015–1034. doi:10.1007/s40747-021-00568-7



OPEN ACCESS

EDITED BY

Mark D. DeHart,
Abilene Christian University, United States

REVIEWED BY

Julian Duran,
Karlsruhe Institute of Technology (KIT),
Germany
Glenn Sjoden,
The University of Utah, United States

*CORRESPONDENCE

Liang Yang,
✉ liang.yang@cranfield.ac.uk
Andrew G. Buchan,
✉ a.buchan@qmul.ac.uk

RECEIVED 20 March 2025

ACCEPTED 19 September 2025

PUBLISHED 15 October 2025

CITATION

Yang L, Yang J, Popov J and Buchan AG (2025) Hybrid solvers for reactor modelling: matrix-based and matrix-free approaches on voxel-dominated meshes.

Front. Nucl. Eng. 4:1597165.
doi: 10.3389/fnuen.2025.1597165

COPYRIGHT

© 2025 Yang, Yang, Popov and Buchan. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Hybrid solvers for reactor modelling: matrix-based and matrix-free approaches on voxel-dominated meshes

Liang Yang^{1,2*}, Jianhui Yang², Jakov Popov³ and Andrew G. Buchan^{3*}

¹Faculty of Engineering and Applied Sciences, Cranfield University, Bedford, United Kingdom, ²Voxshell Ltd, Cranfield University Technology Park, Bedford, United Kingdom, ³School of Engineering and Materials Science, Queen Mary University of London, London, United Kingdom

Simulating neutronics and thermal hydraulics within nuclear reactor cores is computationally intensive, not only because of the complexity of the governing equations but also because of the intricate geometries involved. Solving the Boltzmann transport and Navier-Stokes equations for a full core representation typically relies on unstructured meshes, which, while highly flexible, can substantially increase computational costs regarding memory and solving time. Cartesian meshes with Finite Elements (FE) offer a faster alternative, potentially improving computational speed by an order of magnitude due to direct memory addressing. However, they necessitate finer grids to accurately capture the boundary details of non-Cartesian surfaces, which can offset these gains by increasing solver times. To address this challenge, a new meshing algorithm is proposed in conjunction with hybrid, matrix-based and matrix free, solver technologies. It employs a geometry-conforming boundary method using voxel-dominated Cartesian meshes. This method enables accurate boundary representation at arbitrary resolutions, which can be adjusted to resolve the physics to the desired level of accuracy rather than strictly to capture geometric detail. This is combined with a hybrid solver for fluid flows to different regions of a problem in order to increase efficiency when resolving the boundary. This article demonstrates the method's application to Computational Fluids Dynamics (CFD) and neutronics problems relevant to reactor physics, showcasing its accuracy, convergence, numerical stability, and suitability for handling complex geometries.

KEYWORDS

hybrid matrix and matrix-free solver, reactor physics, finite element, CFD, neutron transport, voxel-dominated meshing

Highlights

- Voxel-dominant meshes are built from Signed Distance Functions (SDF), allowing for accurate surface representation.
- A hybrid matrix-free and matrix-based FEA solver was developed to accelerate computations associated with voxel-dominant meshes

1 Introduction

Simulating neutron transport and thermal hydraulics within nuclear reactor cores is computationally intensive. The complexity of the governing equations combined with the intricate geometries involved often lead us to use simplifications, whether that be to the equations or the problem domains, that often sacrifices numerical accuracy for reduced computational costs. Often the thermal hydraulics are resolved at the sub-channel scale using coupled one-dimensional models that are highly parameterised and fine tuned to experimental data. Other examples include the use of diffusion theory for resolving the neutron distribution, perhaps with a few groups resolving the energy domain with smeared cross-sections homogenising the heterogeneous structures making a core.

High fidelity models are of course used today and are still being developed by various research bodies throughout the world to improve the quality of numerical approximations. Finite element, spectral element and control volume models are popular choices as they offer a flexible route to capturing geometries, and have been developed extensively over the last few decades. Examples of their use in thermal hydraulics include [Paul et al. \(2008\)](#), and similarly [Fang et al. \(2021\)](#), providing full FEM solutions of the Navier-Stokes LES equations for reactor analysis. Others have proposed hybrid methods, such as [Liu et al. \(2019\)](#), [Liu et al. \(2022\)](#), that aim to combine CFD calculations with sub-channel models to enhance accuracy whilst retaining the computational speed of the sub-channel model. A recent review of methods in CFD in nuclear thermal hydraulics can be found in [Wang M. et al. \(2021\)](#). Similarly, for neutron transport, FEM solutions on both structured ([Evans et al., 2010](#)) and unstructured grids ([Wang Y. et al., 2021](#); [Nguyen et al., 2025](#)) have been developed and are in use in many industry and research codes and meshing tools ([Tautges and Jain, 2012](#)).

Unstructured mesh methods are best suited to handle complex geometries composing a core that include fuel pins, spacer grids among other supporting structures. Structured meshes, say with Finite Elements, have also been proposed as a faster alternative, offering the potential to improve computational speed by an order of magnitude due to their reduced memory access requirements and structured solver implementations. However, their primary limitation lies in handling complex geometries, particularly in the presence of curved boundaries and irregular features. Unlike unstructured meshes, which conform naturally to boundaries, Cartesian grids approximate curved surfaces with step-like edges. This stair-casing effect leads to significant accuracy loss, particularly in the BTE, where geometric fidelity plays a critical role in determining neutron streaming paths and scattering effects.

To mitigate the stair-casing issue, several Cartesian-based solver techniques have been developed for general PDEs, including applications in fluid flow, solid mechanics. These methods include:

- Immersed boundary methods ([Peskin, 2002](#); [Verzicco, 2023](#); [Yang et al., 2018](#); [Yang, 2018](#)), which treats the boundary/interface as a forcing term on the smeared-interface surface. It either use the Dirac delta or Heaviside distribution.
- Cut-cell approaches ([Ingram et al., 2003](#); [Meinke et al., 2013](#); [Xie, 2022](#)), where the geometry cuts through the Cartesian grid or block-structured adaptive mesh refinement (AMR) grid ([Barad et al., 2009](#); [Salih et al., 2019](#)), and cells that intersect

the boundary are trimmed to match the geometry. It strongly enforces the boundary. But it may lead to small elements or irregular cuts.

- Unfitted boundary methods ([Lesueur et al., 2022](#); [Martorell and Badia, 2024](#)), which uses background grids with weakly enforced boundary constraints.

Other approaches include the development of iso-parametric elements that can be used in combination with both structured and unstructured meshes to conform exactly to structures such including fuel pins ([Owens et al., 2016](#)).

From an implementation point of view, there are two ways to implement an FEA solver: the matrix-based method and the matrix-free method ([Johan and Hughes, 1991](#); [Davydov et al., 2020](#); [Martínez-Frutos and Herrero-Pérez, 2015](#)). The matrix-based approach assembles the global stiffness matrix, which can be memory-intensive for large sparse systems, and the assembly cost grows with the problem size. On the other hand, the matrix-free approach does not assemble the global matrix; instead, it computes matrix–vector products on the fly. This results in much lower memory usage and faster performance for large-scale problems, e.g., on GPUs.

There are two type of geometry modelling approaches, explicit modelling or implicit modelling. Explicit modeling describes geometry through direct parameterisation, where surfaces and boundaries are explicitly defined. This approach is widely used in Computer-Aided Design (CAD) systems, which rely on representations such as Non-Uniform Rational B-Splines (NURBS) and triangle meshes to create and manipulate complex structures.

Implicit modelling, on the other hand, defines geometry indirectly using mathematical functions, typically representing surfaces as level sets of scalar fields ([Li et al., 2018](#)). This approach is particularly advantageous for handling complex, topologically evolving shapes, such as those found in fluid simulations, medical imaging (MRI/CT scans), and volumetric datasets (e.g., VDB data ([Museth, 2013](#))). It can be generated via implicit geometry kernel or library, including [nTop Inc. \(2024\)](#), [LEAP71 \(2024\)](#).

This article explores implicit geometric modelling techniques, specifically by leveraging the signed distance function (SDF) to efficiently represent complex geometries. In addition, the proposed method addresses the challenge of using structured meshes that still conform to geometries of arbitrary shape. The method (named implicit geometry-conforming meshing (IGCM)) employs a voxel-dominated Cartesian mesh, that maintains as much structure as possible, but for a few elements that are perturbed so that element surfaces are fitted to an internal body's boundary, or a problem's exterior domain. In addition, this article also proposes a hybrid solver for CFD that, in particular, can be used in combination with the meshing tool. This approach ensures a balance between computational efficiency and geometric fidelity, allowing fine-scale resolution where necessary while maintaining the performance benefits of a structured mesh. In contrast to the hex-core meshing algorithm, the mesh will be exported in unstructured format. Our solver will leverage the Cartesian data structure and design the associated solver.

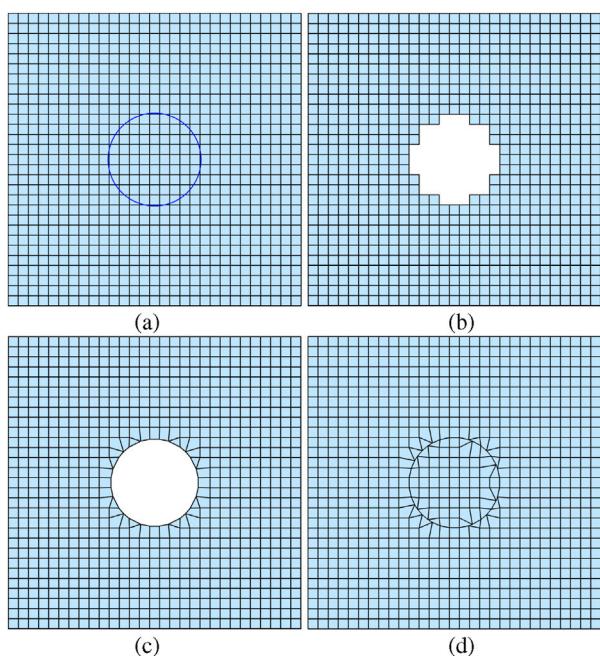


FIGURE 1

(a) Initial voxel grid with circle indicated; (b) Cartesian mesh with internal elements removed; (c) Unfitted boundary method for external domain with mixed Cartesian and unstructured grid; (d) Unfitted boundary method preserving internal/external domains using only quadrilateral elements in Cartesian coordinates.

The following sections are set out as follows. Section 2 details the voxel dominated meshing tool together with the hybrid solver for CFD. Section 3 presents a range of numerical examples applying the proposed tools to a range of applications and benchmarks including neutron transport and CFD of fuel assemblies. Section 4 completes the article with a conclusion.

2 Voxel dominated mesh with hybrid solver

This section describes the voxel-dominated displaced boundary mesh method, along with the hybrid solver for efficient CFD simulations.

Two variants of the meshing approach are proposed for capturing the geometric details of arbitrary bodies. The first applies when only the external or internal region of the domain is needed, such as in CFD calculations for flow around solid bodies. The second resolves both the internal and external regions of a structure, which is typically required for neutron transport and fluid-structure interaction calculations. This approach is often referred to as multi-material with interface.

In both cases, the goal is to preserve a uniform Cartesian grid across as much of the domain as possible, modifying or displacing elements only where necessary to conform to the surface of the body. The method is illustrated in Figure 1, which shows a Cartesian voxel mesh initially discretising a square domain containing a circular region whose surface is to be preserved. Maintaining the Cartesian structure throughout but filtering out the internal elements of the

circle (assuming here the exterior is to be resolved) results in the mesh of 1b where clearly the surface is approximated with the step-like structure and resemblance of the surface is only achieved when sufficient, and perhaps excessive, resolution is applied.

Instead, the method introduced here begins with structured Cartesian voxel mesh (or background mesh) of Figure 1a with element size Δx . An internal geometry, such as the circle here, is then represented through a signed distance field (SDF), which holds the minimum distance between any position within the domain to the surface of the object, denoted here as Γ_s .

From this, a transformation map, φ_d , is defined on the nodes of the background mesh that map the set of closest nodes to the surface, to the nearest position on the surface. That is, for the given set of nodes on the original background mesh, denoted as $\{\mathbf{x}_i\}$, for $i \in \{1, 2, \dots, N\}$ (N being the total number of mesh nodes), a new set of point positions $\{\hat{\mathbf{x}}_i\}$ are formed by,

$$\hat{\mathbf{x}}_i = \varphi_d(\mathbf{x}_i) = \begin{cases} \mathbf{x}_i + \mathbf{d}(\mathbf{x}_i), & \text{if } \mathbf{x}_i \in \Gamma_M, \\ \mathbf{x}_i, & \text{otherwise.} \end{cases} \quad (1)$$

In Equation 1 the term Γ_M denotes the interface region between the surface and the background mesh which contains positions within $\frac{\Delta x}{2}$ from the surface. The term $\mathbf{d}(\mathbf{x}_i)$ is the vector of shortest distance between the node i and the surface Γ_s . That is $\mathbf{d}(\mathbf{x}_i) := \mathbf{x}_{S,i} - \mathbf{x}_i$, with $\mathbf{x}_i \in \Gamma_M$, and $\mathbf{x}_{S,i} \in \Gamma_s$ being the closest position to the surface.

The set of nodes $\{\hat{\mathbf{x}}_i\}$ define the nodes of the mesh used in the CFD simulations of this article. Note the element and node connectivity remains the same, and that the elements with changed shape are those with nodes that have been moved. In the cases where a moved element has an edge in line with the structure, no further changes are made. Those elements where the object's surface cuts through the element's volume (i.e., cutting through two opposite nodes of the quadrilateral), the element is partitioned into two triangles, with the joining edge aligning with the internal body's surface. The internal triangle is then removed from the mesh, together with all other internal quadrilateral elements of the unchanged background mesh, resulting in the mesh presented in Figure 1c. As can be seen, the regular mesh is preserved but for the set of quadrilateral and triangle elements surrounding the surface of the internal body. This approach is used in this article for the CFD calculations.

For neutron transport problems, the solution domain includes both internal and external regions relative to a surface—such as when resolving fuel pins and surrounding moderator. In this situation, meshes must preserve the volumes of both internal and external regions. To achieve this, the method presented here uses the SDF function, as before, to define distances between a surface and the nodes of an initial background Cartesian mesh. The initial node movement described in Equation 1 is again applied; however, internal elements of the body are retained. At this stage, modified elements surrounding the surface either preserve the surface along an edge or have the surface cutting through their volume. In the former case, no more modifications are necessary. In the latter case, a further modification is applied to the element's nodes to preserve the surface while maintaining the quadrilateral structure of the overall mesh. A single node not in contact with the surface is identified and moved onto it. As presented in Figure 1d, the resulting mesh preserves both the surface and the Cartesian quadrilateral mesh structure.

2.1 Hybrid region FEM solver

For time dependent CFD solutions, a hybrid matrix-based and matrix-free solver is implemented for the proposed voxel-dominant meshes that is composed of a structured (Cartesian) and unstructured component. More specifically, the domain Ω is decomposed into two sub-domains $\Omega = \Omega_C \cup \Omega_U$, where Ω_C represents the Cartesian mesh region, and Ω_U represents the unstructured mesh region. Each component is treated separately in terms of data storage, memory addressing and, if stored, matrix assembly. A node pairing between local Cartesian grid Ω_C and unstructured mesh Ω_U is imposed to enforce solution continuity.

For the Cartesian mesh region (Ω_C), where the discretisation uses hexahedral (in 3D) or quadrilateral (in 2D) elements, the nodes and elements are stored in a Cartesian grid format. The shape functions and element stiffness matrices follow standard FEA for structured meshes with trilinear/bilinear elements.

For the unstructured mesh region (Ω_U), the mesh consists of four types of linear elements that includes hexahedral, tetrahedra, pyramids and wedge. Here the standard finite element assembly process is used, storing global system matrices.

The solver treats each region separately, where for the Cartesian grid it employs the previous work developed in Yang et al. (2016), Yang et al. (2019), Nillama et al. (2022). In essence this is a stabilised explicit, matrix-free, solver designed to reduce memory access cost of indirect addressing by relying on the structured nature of the grid. For the unstructured grid component, the solver relies on matrix-based approaches which, whilst fast, does require more memory. However, due to the voxel-dominant mesh structure, the unstructured part is restricted to the interface/boundary region which is typically a small proportion of the whole domain. Thus the larger proportion of the domain is treated with minimal storage. Overall, this may provide an even balance of memory usage and efficient computation.

3 Numerical results

This section provides a number of numerical examples using the proposed methods in diffusion-based applications, CFD and neutron transport to demonstrate the accuracy and stability.

3.1 Poisson problem: circle embedded in a square domain with analytical solution

This numerical example evaluates the accuracy of the proposed meshing method in solving Poisson's equation, given by:

$$-\nabla \cdot \nabla u(r, \theta) = f(r, \theta),$$

where the solution variable $u(r, \theta)$ is expressed in polar coordinates.

The problem is defined over a square domain $[-1, 1] \times [-1, 1]$, with a circular boundary of radius $R = 1/3$ centered at the origin. A Neumann boundary condition is imposed on the circle, $\frac{\partial u}{\partial n} = 0$, where n denotes the outward pointing normal to the surface. The source term $f(r, \theta)$ is defined as

$$f(r, \theta) = \frac{24}{r} \sin(-5\theta),$$

On the outer square boundary, a Dirichlet boundary condition is applied to match the analytical solution:

$$u(r, \theta) = r \sin(-5\theta) + \frac{1}{5}(r/R)^{-6}r \sin(-5\theta).$$

The problem is solved on two different mesh types: 1. Voxel-Dominated Mesh—A geometry-conforming mesh where the circular boundary is fully resolved. 2. Standard Cartesian Mesh—A structured Cartesian grid where elements overlapping with the circle are removed.

The computed solutions and errors for both mesh types, using a 30×30 element resolution, are presented in Figure 2. The results show that while both meshes yield similar solutions, the voxel-dominated mesh significantly reduces maximum errors and suppresses the extent of the largest error regions by better conforming to the geometry.

This improvement is further illustrated in Figure 3, which plots the solution along the horizontal line at $y = 0.5$. The figure compares results for both mesh types at resolutions of 30×30 and 50×50 . The geometry-conforming mesh demonstrates improved accuracy, particularly in regions close to the embedded circular boundary.

In the two simulations presented, the solver iterations (see Section 3.6) required to reach convergence were closely comparable. The voxel and voxel-dominant meshes required 80 and 82 solver iterations for tolerance 10^{-8} . Since the mesh sizes are nearly identical and CPU solver iterations have the same cost, this demonstrates that the meshing approach allows geometries to be resolved without increasing the computational cost of solving the discretised equations.

3.2 Fixed source neutron transport simulation of fuel pin and PWR assembly

This numerical example provides a demonstration of the meshing tool for use in resolving neutron transport problems at the pin-level scale. The analysis is focused on the ability of the meshing tool to capture domain surfaces accurately, but which also remains numerically stable when used in combination with typical FEM discretisations. For this demonstration both the single energy group fixed source and multi-group eigenvalue versions of the Boltzmann transport equations are solved. The latter of these is given by,

$$\begin{aligned} & \Omega \cdot \nabla \psi_g(\vec{r}, \Omega) + \Sigma_{t,g}(\vec{r})\psi_g(\vec{r}, \Omega) \\ & - \sum_{g'=1}^G \int_{\Omega'} \Sigma_{s,g' \rightarrow g}(\vec{r}, \Omega' \rightarrow \Omega) \psi_{g'}(\vec{r}, \Omega') d\Omega' \\ & = \frac{\chi_g}{4\pi k_{eff}} \sum_{g'=1}^G \int_{\Omega'} \nu \Sigma_{f,g'}(\vec{r}) \psi_{g'}(\vec{r}, \Omega') d\Omega', \end{aligned}$$

for $g \in \{1, 2, \dots, G\}$. The solution yields the angular flux for each energy group g , $\psi_g(\vec{r}, \Omega)$, in direction Ω at position \vec{r} together with the multiplication factor k_{eff} . The terms $\nu \Sigma_{f,g'}$ and $\Sigma_{s,g' \rightarrow g}$ denote the production (via fission from group g') and group to group scattering material cross-sections, respectively, and χ_g denotes the fission spectrum.

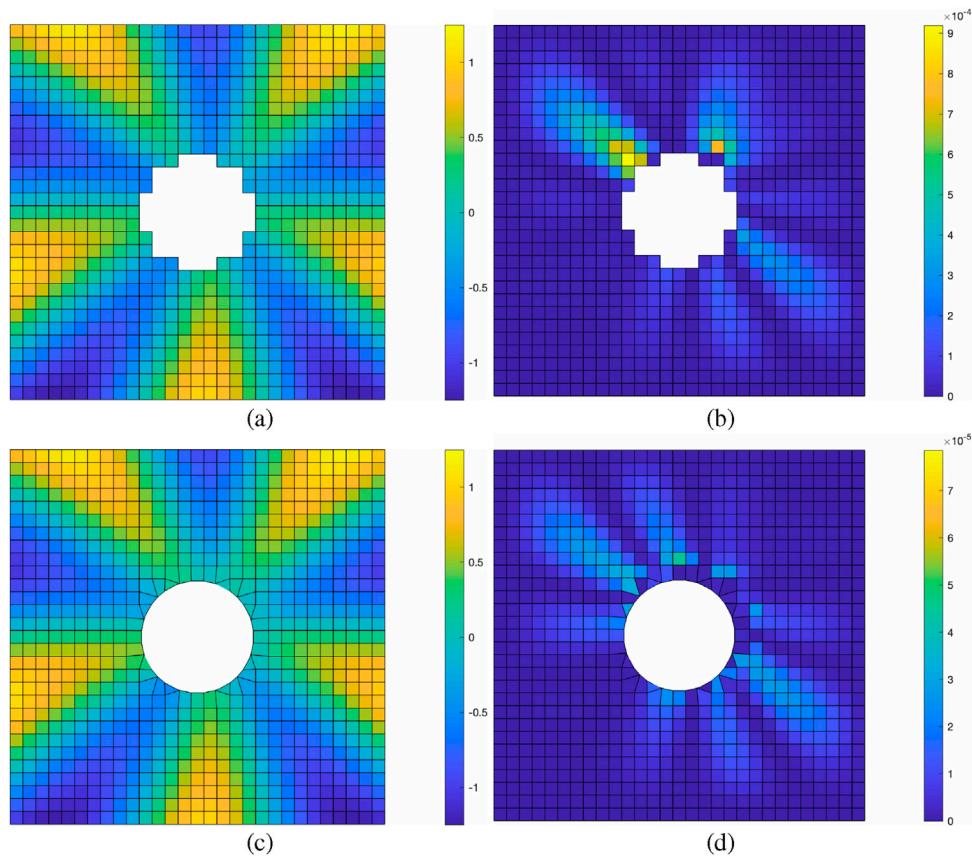


FIGURE 2
(a,b) Solution and error using the voxel grid. **(c,d)** Solution and error using the voxel-dominant grid. Both sets of results are based on a 30×30 quadrilateral element mesh.

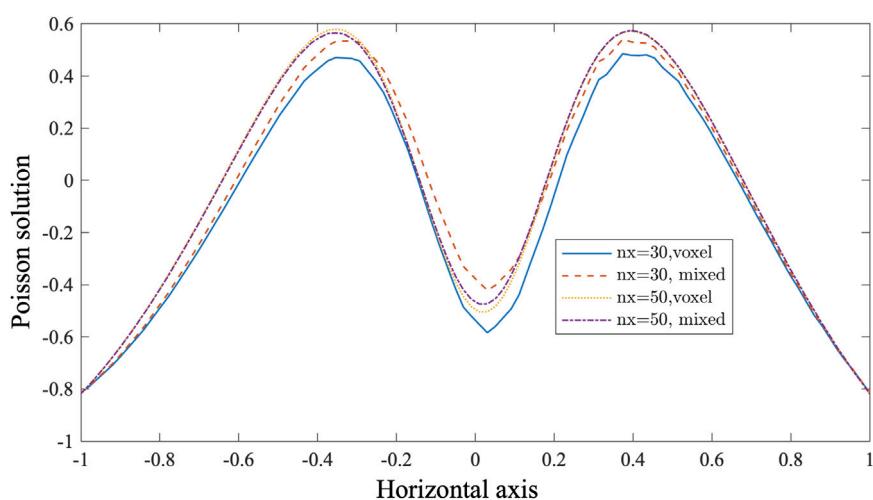


FIGURE 3
Solution along the horizontal line at $y = 0.5$ x-axis: horizontal axis. y-axis: solutions.

In the following analysis standard numerical techniques for discretising the space-angle dimensions of the BTE are employed. This includes a discontinuous Galerkin FEM spatial discretisation

using bi-linear basis functions on quadrilateral element meshes. This is combined with a simple, zero-order structured discrete ordinate- S_N method, based on equal solid-angle quadrature for discretising

TABLE 1 Source and material cross-sections used in the neutron transport test cases.

Test case	Source	Σ_t	Σ_s
Test Case 1	1.0	0.01	0.0
Test Case 2	1.0	2.0	1.9

the angular domain in the Boltzmann transport equation, the details of which can be found in [Adigun et al. \(2018\)](#).

The first demonstration uses a 2-dimensional single fuel pin of radius 0.418 cm, encapsulated within a square domain, with side lengths 1.26 cm, bounded by reflective surfaces. Two sets of single group cross-sections and neutron sources are used which are not designed to represent anything physical but, moreover, are used to test numerical stability using the geometry capturing meshes over different material regimes. In both problems, the material and sources in the fuel and moderator remain constant, as listed in [Table 1](#), with the scattering and total cross-sections set to examine mesh performance in the transparent and diffuse material regimes. As the analytical solutions yield a constant flux, it would be expected the solutions to be resolved to within machine precision and solver tolerances, regardless of the mesh resolution. Thus numerical instabilities arising from the mesh and element shape will most likely be exposed when capturing the circular pin geometry. Noting that all calculations were performed using double precision, this investigation is not focused on other instabilities, causing solution divergence, including those due to poor element shape and lower

precision computations. [Figure 4](#) present the absolute relative scalar flux errors using meshes ranging from FEM grids of 8×8 to 100×100 . In all cases, whilst the largest errors occur around the pin interface, the errors are to within typical round-off limits and solver tolerances as they do not exceed 10^{-8} . This demonstrates that the numerical methods remain stable when used in combination with the meshing techniques for the extreme transport regimes in radiation transport.

The following investigation evaluates mesh performance in resolving fuel pins within a representative reactor geometry. A single fuel pin and its surrounding moderator were modeled using the dimensions and material cross-sections for uranium dioxide (UO_2) fuel and moderator from the C5G7 benchmark [Nuclear Energy Agency \(2003\)](#). Thermal flux solutions for group 7, obtained using meshes of varying resolutions from coarse to fine, are presented in [Figure 5](#). Additionally, [Figure 6](#) shows the scalar flux profiles for groups 1 and 7 along the diagonal of the domain (from position $(0,0)$,cm to $(1.26,1.26)$,cm). This figure also includes the 2-norm errors of the scalar flux along the same diagonal. The reference solution was computed using a fine mesh of 400×400 elements, and errors are shown for meshes ranging from 8×8 to 200×200 elements. These results indicate that the solutions remain stable across the different mesh resolutions. While oscillations appear in the coarse mesh, particularly in the fast flux, these are expected due to insufficient resolution rather than poor mesh performance. That is, the coarse elements simply lack the ability to capture the finer solution details observed in higher-resolution meshes. Increasing the resolution mitigates these oscillations, with the fine mesh solution

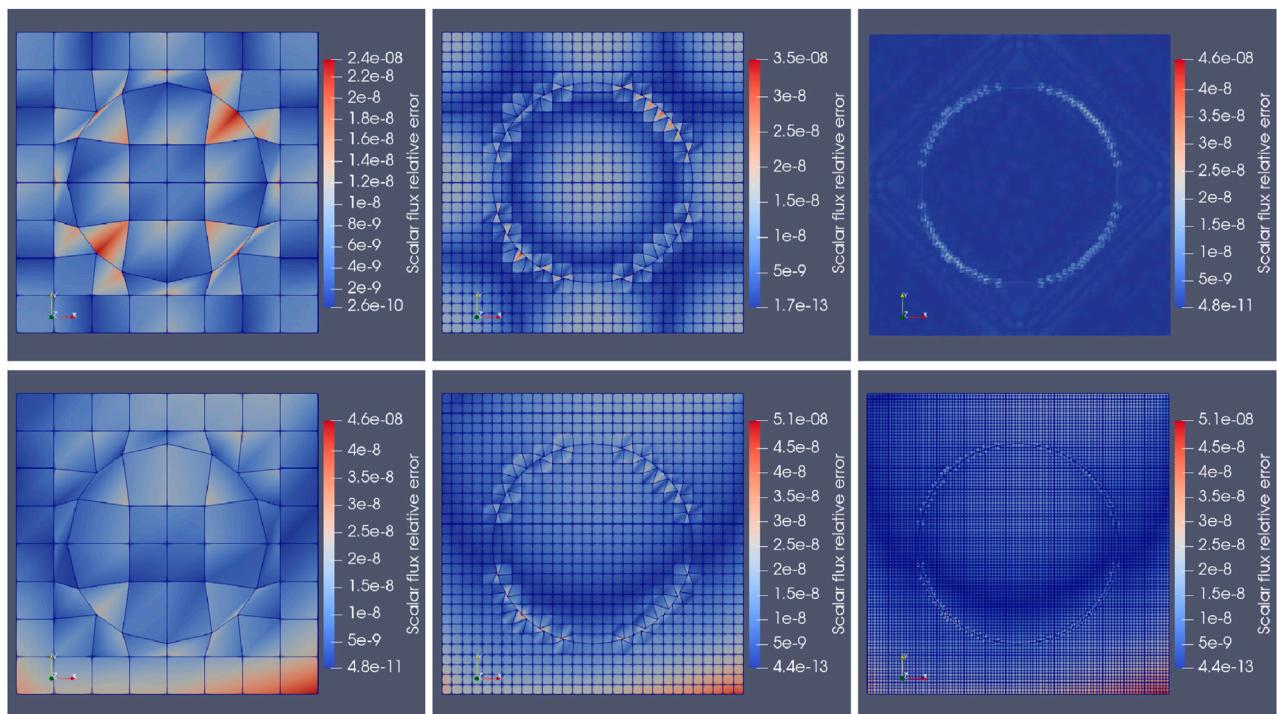
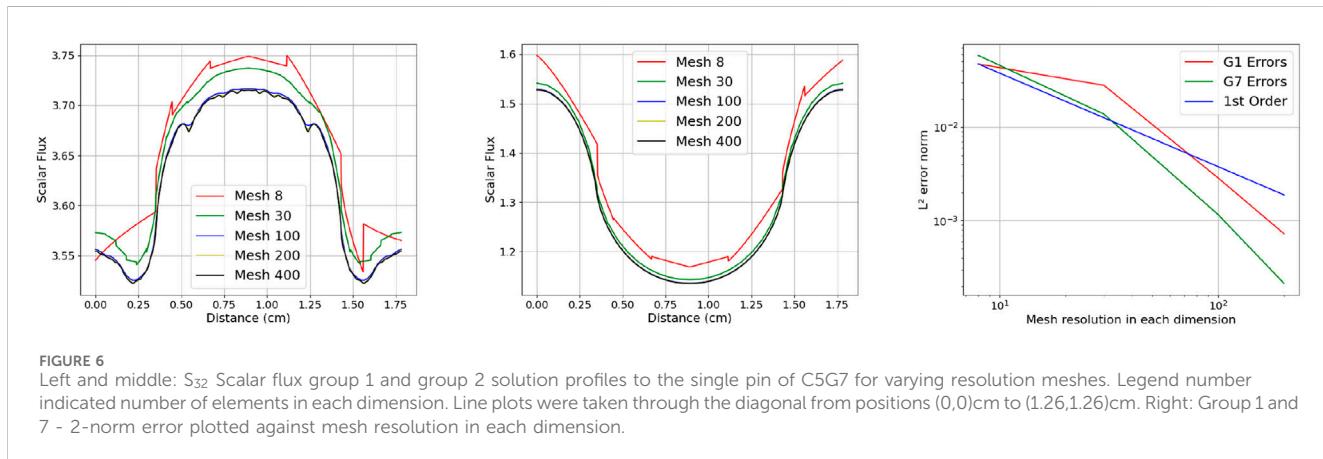
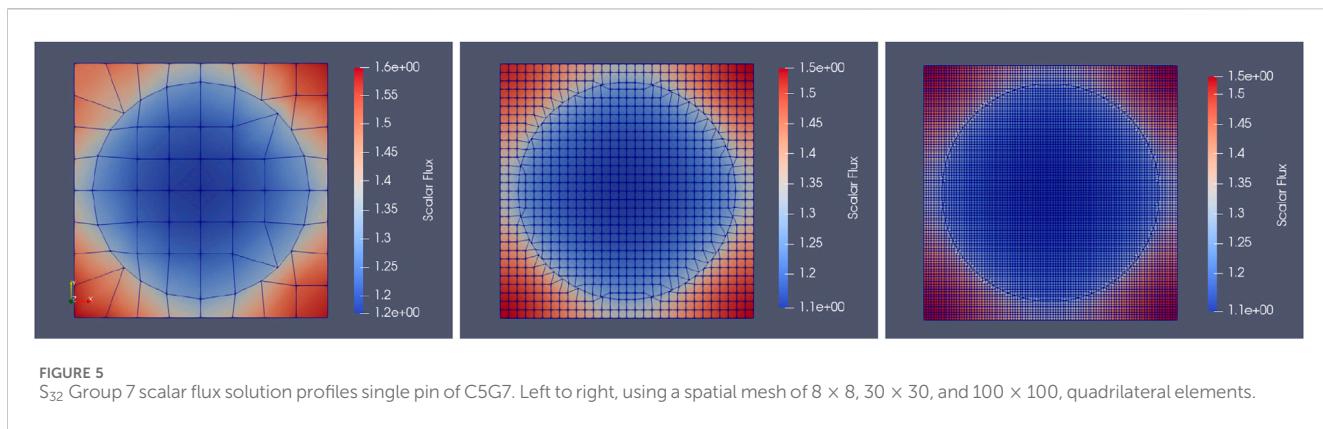


FIGURE 4
 S_{16} Scalar flux solution absolute relative errors to the void (top) and scatter (bottom) pin problem. Left to right, using a spatial mesh of 8×8 , 30×30 , and 100×100 , quadrilateral elements.



exhibiting smooth contours, except for small blips in the fast flux, which result from ray effects. The error convergence graphs show that the errors do converge at a rate in close agreements with first-order convergence, which again should be expected due to the use of linear FEM functions.

In summary, this numerical example provides evidence the meshing algorithm is suitable for resolving fuel pins across transparent and diffuse material regimes, and for those typical found in reactor physics benchmarks.

3.3 Two dimensional CFD benchmark

The 2D lid-driven cavity flow (Ghia et al., 1982) at $Re = 1000$ is a well-established numerical benchmark in computational fluid dynamics (CFD). It involves an incompressible, viscous flow inside a square cavity, where the top boundary moves at a constant velocity $v = 1$ while the other three walls remain stationary. This setup generates a primary vortex in the cavity centre, along with secondary vortices near the corners. The case serves as an ideal benchmark for assessing numerical schemes, solution accuracy. It is a convection-dominated problem, where stabilised FE is used (Yang et al., 2016).

To validate the hybrid solver, we rotate the computational domain by 45°, intentionally introducing a stair-casing issue. This test evaluates the solver's ability to handle both Cartesian and

unstructured grids while ensuring accurate vorticity transport and solution consistency.

Figure 7 presents the computational meshes and their corresponding velocity solutions under a 50 × 50 resolution. The voxel-dominant mesh demonstrates superior accuracy in capturing viscous boundary layers, leading to more accurate results compared to a purely voxel-based approach. Figure 8 further illustrates the velocity profiles along the horizontal and vertical centerlines of the cavity, comparing the obtained solutions against established reference data.

3.4 Three dimensional CFD CANDU fuel assembly

This study presents a CFD analysis of a three-dimensional CANDU fuel assembly (Torgerson et al., 2006), focusing on pressure distribution and velocity profiles. The CANDU (Canada Deuterium Uranium) reactor utilizes heavy water moderation, with fuel bundles immersed in pressurized coolant. Accurately modeling the flow behavior around the fuel rods is essential for optimising heat transfer efficiency and minimising pressure drop (Piro et al., 2016).

A typical CANDU fuel bundle, illustrated in Figure 9, consists of 37 fuel elements, concentrically arranged and mounted on two end plates. The bundle has a total length of approximately 500 mm and a

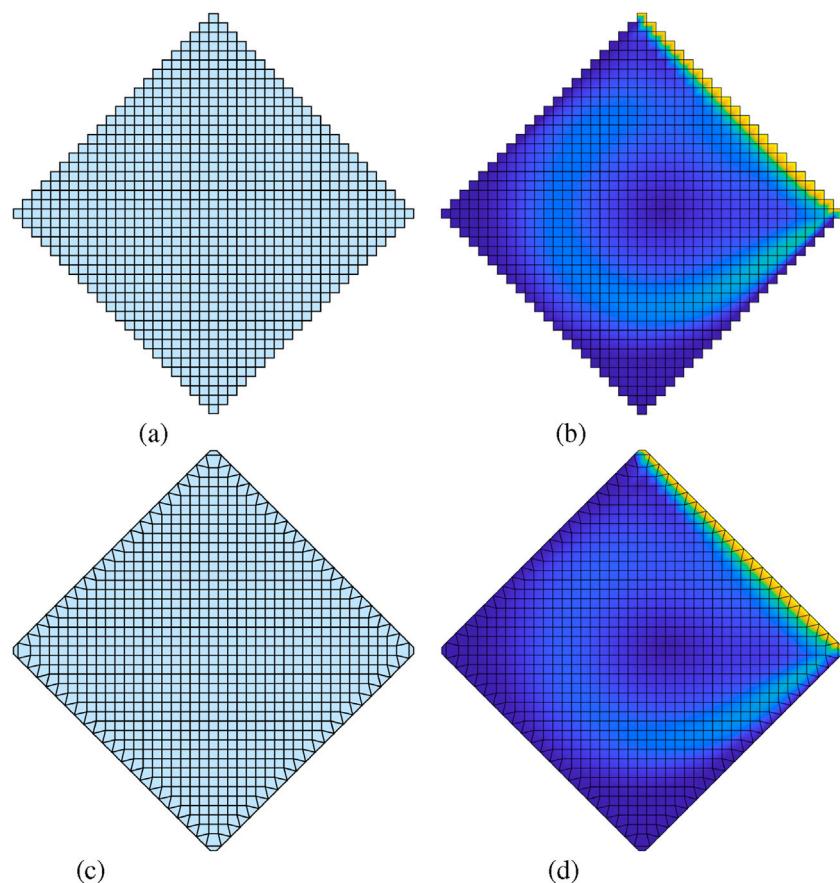


FIGURE 7
Lid driven flow benchmark case for $Re = 1000$ using voxel-grid **(a, b)** and voxel-dominant mesh **(c, d)** with hybrid solver with 50×50 .

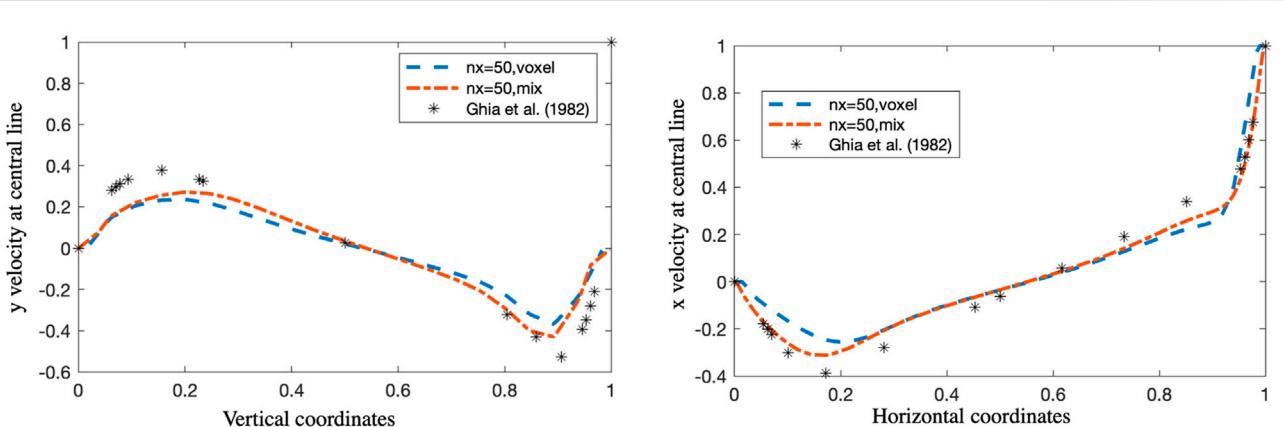


FIGURE 8
Vertical and horizontal velocity through the horizontal centerline and vertical of the cavity at $Re = 1,000$ using $n = 50$ solution in the original coordinates before rotation.

diameter of about 100 mm. The CFD simulation provides insights into velocity distribution, flow separation, and pressure variations within the assembly. These results help assess flow uniformity and identify potential recirculation zones, ensuring an optimized thermal-hydraulic performance of CANDU fuel bundles during

reactor operation. The simulation is conducted for a single-phase flow, excluding heat transfer and chemical reactions.

To analyse the effects of mesh resolution, surface and volume meshes were generated at three levels of refinement with mesh size of 0.2868 cm, 0.1434 cm, 0.0717 cm, with a refinement factor of

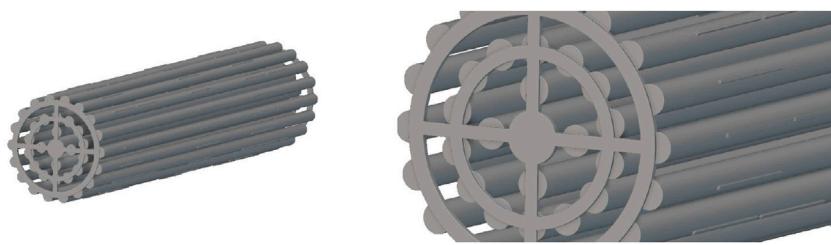


FIGURE 9
Isometric view of a 37 element CANDU fuel bundle modelled in nTOP, left: whole view, right: close view.

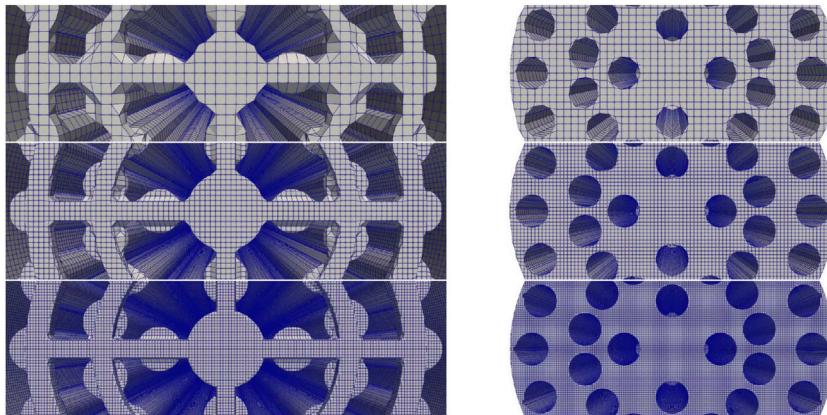


FIGURE 10
Surface and volume meshes at three resolutions: 0.2868 cm, 0.1434 cm, and 0.0717 cm, corresponding to a refinement factor of two between successive levels.

two at each level. Figure 10 shows the surface and volume meshes at three resolutions. All simulations used a mass flow rate of 20.9 kg/s as used in (Bhattacharya et al., 2012). The pressure drop for each simulation of increasing mesh resolutions was calculated as 11.22 kPa, 11.02 kPa, 10.78 kPa showing these values to be converging to within 0.2 kPa. Figure 11 presents the flow velocities pressure profiles using the small and large meshes. Again, this shows closely converging solution details.

3.5 Illustration of computational efficiencies

The use of meshes constructed predominantly of structured elements offer significant memory savings compared to fully unstructured meshes, in part this is due to the absence of indirect addressing and mesh-connectivity storage. Additionally, the repetition of element shapes eliminates the need to compute and store spatial integrals for each individual element. Although quadrilateral and hexahedral elements involve more nodes per element, the overall number of elements is substantially lower, by factors of approximately 6 in 2D and 20 in 3D, resulting in a net reduction in memory and computational costs. Thus the effectiveness of the method presented will be driven by the fraction of unperturbed elements.

Table 2 presents the fraction of unperturbed elements used in the 2D reactor pin cell problems and the 3D CANDU reactor

geometry for varying mesh sizes. In the 2D case, the number of unperturbed element is quite low for the smallest mesh, the eight by eight grid resulting in just 25% of elements being unperturbed. However, one can see from Figure 5 the pin occupies a large proportion of the domain and so the large elements tend to intersect the surface on a regular basis. As mesh resolution is increased the unperturbed fraction quickly rises and reaches over 93% for the 100 by 100 mesh. Similarly the three meshes used for the CANDU geometry show increasing unperturbed fractions, initially starting from 51% for the lowest mesh of 300.5 K elements, and rising to 85% for the large mesh with 20.2 M elements.

3.6 CFD and BTE solver information

The CFD and BTE solver used here was based on an in-house code (Buchan et al., 2021), written in Fortran-90 and compiled as a python library. The CFD solver used a continuous linear FEM with VMS stabilisation, and a generalized minimal residual method (GMRES) solver with a Jacobi preconditioner. The BTE solver used discontinuous linear FEM with a sweep based solutions scheme. For the eigenvalue problems, a Power method was employed. Both solvers have been developed with MPI parallelism and the CFD solver has the additional OpenMP capability which was used in the calculation of the CANDU assembly.

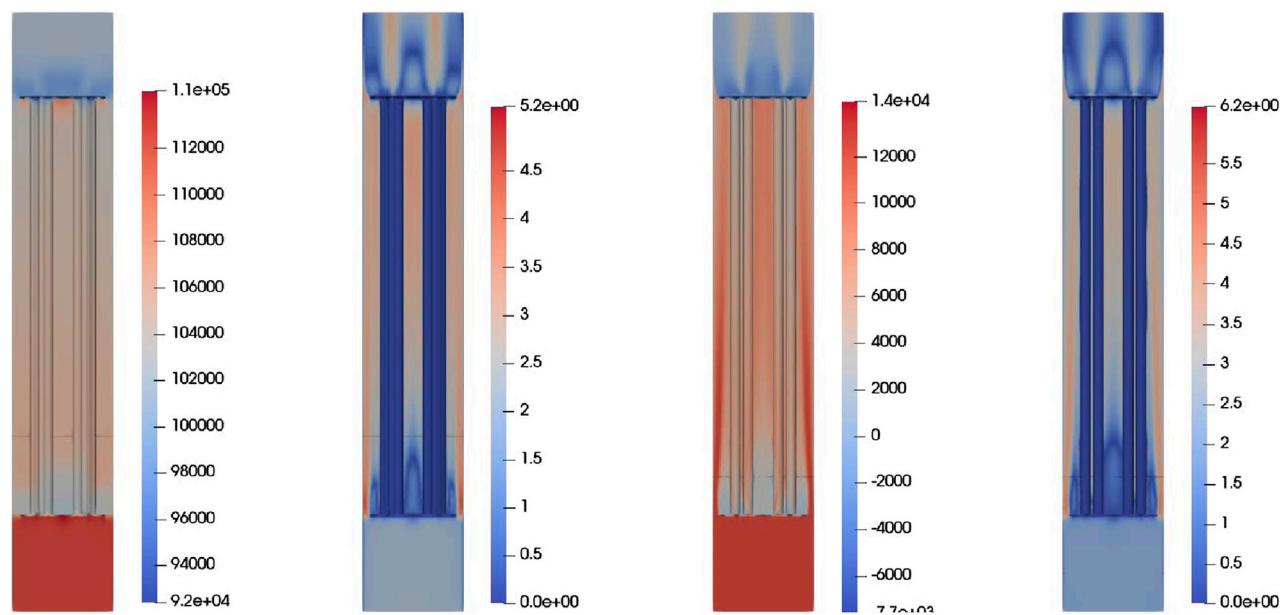


FIGURE 11
Contour plots of pressure distribution and velocity magnitude at coarse (0.2868 cm) and fine (0.0717 cm) grid resolutions, under a mass flow rate of 20.9 kg/s. Left: coarse grid results. Right: fine grid results.

TABLE 2 Fraction of perturbed elements for the 2D neutron transport and 3D CANDU reactor.

Example	Small		Medium		Large	
2D-NT Circle	64	25%	900	77%	10 K	93%
3D-Assembly	300.5 K	51%	2.53 M	72%	20.2 M	85%

4 Conclusion

This study presents a hybrid solver framework integrated with a voxel-dominated Cartesian meshing approach to efficiently solve the Boltzmann Transport Equation and Navier-Stokes equations in complex nuclear reactor simulations. Traditional unstructured meshing techniques, while flexible, impose significant computational burdens due to high memory consumption and solver inefficiencies. In contrast, structured Cartesian methods offer computational advantages but suffer from stair-casing errors, particularly in handling curved geometries.

To overcome these limitations, the implicit geometry-conforming meshing (IGCM) method was introduced within a voxel-dominated Cartesian mesh. This approach is coupled with a hybrid solver, applying structured solvers in voxel regions while using adaptive numerical treatments near complex boundaries. The combination of these methods enables high computational efficiency while preserving accuracy in neutron transport and thermal hydraulics simulation. The demonstration provided here showed that complex geometries can be resolved through the simulation of a CANDU assembly. The proposed hybrid solver and meshing framework provides a promising direction for accelerating full-core reactor simulations while ensuring the accuracy of boundary interactions. Future work will focus on

benchmarking the solver against high-fidelity reference solutions and extending its applicability to multi-physics coupling, including neutron transport, thermal radiation, and turbulence modelling.

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Author contributions

LY: Software, Writing – original draft, Formal Analysis, Funding acquisition, Methodology. JY: Methodology, Writing – review and editing. JP: Formal Analysis, Writing – review and editing. AB: Funding acquisition, Methodology, Software, Writing – original draft.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. STFC funded through grant: Parallel solvers for voxel-dominant meshes for the Boltzmann Transport Equation.

Acknowledgments

Authors acknowledge the support of the STFC-funded Digital BIC program. The authors thank nTOP for providing the free

academic license to use CAD modelling software to model the CANDU fuel bundle.

Conflict of interest

Authors LY and JY were employed by Voxshell Ltd.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

References

Adigun, B. J., Buchan, A. G., Adam, A., Dargaville, S., Goffin, M. A., and Pain, C. C. (2018). A haar wavelet method for angularly discretising the boltzmann transport equation. *Prog. Nucl. Energy* 108, 295–309. doi:10.1016/j.pnucene.2018.05.023

Barad, M. F., Colella, P., and Schladow, S. G. (2009). An adaptive cut-cell method for environmental fluid mechanics. *Int. J. Numer. methods fluids* 60, 473–514. doi:10.1002/fld.1893

Bhattacharya, A., Yu, S., and Kawall, G. (2012). Numerical simulation of turbulent flow through a 37-element candu fuel bundle. *Ann. Nucl. energy* 40, 87–105. doi:10.1016/j.anucene.2011.10.017

Buchan, A. G., Yang, L., Welch, D., Brenner, D. J., and Atkinson, K. D. (2021). Improved estimates of 222 nm far-uvc susceptibility for aerosolized human coronavirus via a validated high-fidelity coupled radiation-cfd code. *Sci. Rep.* 11, 19930. doi:10.1038/s41598-021-99204-0

Davydov, D., Pelteret, J.-P., Arndt, D., Kronbichler, M., and Steinmann, P. (2020). A matrix-free approach for finite-strain hyperelastic problems using geometric multigrid. *Int. J. Numer. Methods Eng.* 121, 2874–2895. doi:10.1002/nme.6336

Evans, T. M., Stafford, A. S., Slaybaugh, R. N., and and, K. T. C. (2010). Denovo: a new three-dimensional parallel discrete ordinates code in scale. *Nucl. Technol.* 171, 171–200. doi:10.13182/NT171-171

Fang, J., Shaver, D. R., Tomboulides, A., Min, M., Fischer, P., Lan, Y.-H., et al. (2021). Feasibility of full-core pin resolved cfd simulations of small modular reactor with momentum sources. *Nucl. Eng. Des.* 378, 111143. doi:10.1016/j.nucengdes.2021.111143

Ghia, U., Ghia, K. N., and Shin, C. (1982). High-re solutions for incompressible flow using the navier-stokes equations and a multigrid method. *J. Comput. Phys.* 48, 387–411. doi:10.1016/0021-9910(82)90058-4

Ingram, D. M., Causon, D. M., and Mingham, C. G. (2003). Developments in cartesian cut cell methods. *Math. Comput. Simul.* 61, 561–572. doi:10.1016/s0378-4754(02)00107-6

Johan, Z., and Hughes, T. J. (1991). A globally convergent matrix-free algorithm for implicit time-marching schemes arising in finite element analysis in fluids. *Comput. Methods Appl. Mech. Eng.* 87, 281–304. doi:10.1016/0045-7825(91)90009-u

LEAP71 (2024). PicoGK: a compact and robust open-source geometry kernel for computational engineering.

Lesueur, M., Rattez, H., and Colomés, O. (2022). puct scans permeability computation with an unfitted boundary method to improve coarsening accuracy. *Comput. and Geosciences* 166, 105118. doi:10.1016/j.cageo.2022.105118

Li, Q., Hong, Q., Qi, Q., Ma, X., Han, X., and Tian, J. (2018). Towards additive manufacturing oriented geometric modeling using implicit functions. *Vis. Comput. Industry, Biomed. Art* 1, 9–16. doi:10.1186/s42492-018-0009-y

Liu, B., He, S., Moulinec, C., and Uribe, J. (2019). Sub-channel cfd for nuclear fuel bundles. *Nucl. Eng. Des.* 355, 110318. doi:10.1016/j.nucengdes.2019.110318

Liu, B., He, S., Moulinec, C., and Uribe, J. (2022). A multiscale model of a rod bundle using subchannel cfd. *Nucl. Eng. Des.* 393, 111793. doi:10.1016/j.nucengdes.2022.111793

Martínez-Frutos, J., and Herrero-Pérez, D. (2015). Efficient matrix-free gpu implementation of fixed grid finite element analysis. *Finite Elem. Analysis Des.* 104, 61–71. doi:10.1016/j.finel.2015.06.005

Martorell, P. A., and Badia, S. (2024). High order unfitted finite element discretizations for explicit boundary representations. *J. Comput. Phys.* 511, 113127. doi:10.1016/j.jcp.2024.113127

Meinke, M., Schneiders, L., Günther, C., and Schröder, W. (2013). A cut-cell method for sharp moving boundaries in cartesian grids. *Comput. and Fluids* 85, 135–142. doi:10.1016/j.compfluid.2012.11.010

Museth, K. (2013). Vdb: high-resolution sparse volumes with dynamic topology. *ACM Trans. Graph. (TOG)* 32, 1–22. doi:10.1145/2487228.2487235

Nguyen, K. H., DeHart, M. D., Hanophy, J. T., and Wang, Y. (2025). Development of a griffin model of the advanced test reactor. *Ann. Nucl. Energy* 211, 111012. doi:10.1016/j.anucene.2024.111012

Nillama, L. B. A., Yang, J., and Yang, L. (2022). An explicit stabilised finite element method for navier-stokes-brinkman equations. *J. Comput. Phys.* 457, 111033. doi:10.1016/j.jcp.2022.111033

nTop Inc (2024). *nTop documentation, release 4.0: implicit and other models*. nTop Inc. Available online at: <https://ntop.com>.

Nuclear Energy Agency (2003). *NEA/NSC/DOC(2003)16*. Paris, France: Nuclear Energy Agency, Organisation for Economic Co-operation and Development.

Owens, A., Welch, J., Kópházi, J., and Eaton, M. (2016). Discontinuous isogeometric analysis methods for the first-order form of the neutron transport equation with discrete ordinate (sn) angular discretisation. *J. Comput. Phys.* 315, 501–535. doi:10.1016/j.jcp.2016.03.060

Paul, F., Fischer, J. W. L., and Kerkemeier, S. G. (2008). nek5000 web page. Available online at: <http://nek5000.mcs.anl.gov>.

Peskin, C. S. (2002). The immersed boundary method. *Acta Numer.* 11, 479–517. doi:10.1017/s0962492902000077

Piro, M., Wassermann, F., Grundmann, S., Leitch, B., and Tropea, C. (2016). Progress in on-going experimental and computational fluid dynamic investigations within a candu fuel channel. *Nucl. Eng. Des.* 299, 184–200. doi:10.1016/j.nucengdes.2015.07.009

Salih, S. Q., Aldlemy, M. S., Rasani, M. R., Ariffin, A., Ya, T. M. Y. S. T., Al-Ansari, N., et al. (2019). Thin and sharp edges bodies-fluid interaction simulation using cut-cell immersed boundary method. *Eng. Appl. Comput. Fluid Mech.* 13, 860–877. doi:10.1080/19942060.2019.1652209

Tautges, T. J., and Jain, R. (2012). Creating geometry and mesh models for nuclear reactor core geometries using a lattice hierarchy-based approach. *Eng. Comput.* 28, 319–329. doi:10.1007/s00366-011-0236-8

Torgerson, D., Shalaby, B. A., and Pang, S. (2006). Candu technology for generation iii+ and iv reactors. *Nucl. Eng. Des.* 236, 1565–1572. doi:10.1016/j.nucengdes.2006.04.020

Verzicco, R. (2023). Immersed boundary methods: historical perspective and future outlook. *Annu. Rev. Fluid Mech.* 55, 129–155. doi:10.1146/annurev-fluid-120720-022129

Wang, M., Wang, Y., Tian, W., Qiu, S., and Su, G. (2021a). Recent progress of cfd applications in pwr thermal hydraulics study and future directions. *Ann. Nucl. Energy* 150, 107836. doi:10.1016/j.anucene.2020.107836

Any alternative text (alt text) provided alongside figures in this article has been generated by Frontiers with the support of artificial intelligence and reasonable efforts have been made to ensure accuracy, including review by the authors wherever possible. If you identify any issues, please contact us.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Wang, Y., Schunert, S., Ortensi, J., Laboure, V., DeHart, M., Prince, Z., et al. (2021b). Rattlesnake: a moose-based multiphysics multischeme radiation transport application. *Nucl. Technol.* 207, 1047–1072. doi:10.1080/00295450.2020.1843348

Xie, Z. (2022). An implicit cartesian cut-cell method for incompressible viscous flows with complex geometries. *Comput. Methods Appl. Mech. Eng.* 399, 115449. doi:10.1016/j.cma.2022.115449

Yang, L. (2018). One-fluid formulation for fluid–structure interaction with free surface. *Comput. Methods Appl. Mech. Eng.* 332, 102–135. doi:10.1016/j.cma.2017.12.016

Yang, L., Badia, S., and Codina, R. (2016). A pseudo-compressible variational multiscale solver for turbulent incompressible flows. *Comput. Mech.* 58, 1051–1069. doi:10.1007/s00466-016-1332-9

Yang, L., Gil, A. J., Carreño, A. A., and Bonet, J. (2018). Unified one-fluid formulation for incompressible flexible solids and multiphase flows: application to hydrodynamics using the immersed structural potential method (ispm). *Int. J. Numer. Methods Fluids* 86, 78–106. doi:10.1002/fld.4408

Yang, L., Yang, J., Boek, E., Sakai, M., and Pain, C. (2019). Image-based simulations of absolute permeability with massively parallel pseudo-compressible stabilised finite element solver. *Comput. Geosci.* 23, 881–893. doi:10.1007/s10596-019-09837-4

Frontiers in Nuclear Engineering

Advances and applications in nuclear science, technology, and engineering

A forum for the nuclear research community working across fission and fusion physics, power generation, safety and waste management, materials and modelling to find new applications.

Discover the latest Research Topics

See more →

Frontiers

Avenue du Tribunal-Fédéral 34
1005 Lausanne, Switzerland
frontiersin.org

Contact us

+41 (0)21 510 17 00
frontiersin.org/about/contact



Frontiers in
Nuclear Engineering

