POLITECNICO DI TORINO

MASTER's Degree in Energy and Nuclear Engineering

MASTER's Degree Thesis

Validation of a foamForNuclear multiphysics model of the BME Training Reactor

Supervisors

Prof. Sandra DULLA

Prof. Mathieu HURSIN

Ph.D. Ludovic JANTZEN

MSc Giovanni NERVI

Candidate

Cesare BARBERA

JULY 2026

Acknowledgments

This work was carried out at the *École Polytechnique Fédérale de Lausanne (EPFL)*, within the Laboratory for Reactor Physics and Systems Behaviour (LRS), and I am grateful for the opportunity to have spent this formative period abroad in such a stimulating research environment. My sincere thanks go to Professor Mathieu Hursin and to Dr. Alessandro Scolaro, for the opportunity they gave me and for the trust they placed in me from the very beginning. Their guidance shaped this study at every stage, and the confidence they showed in my work gave me the freedom to grow as a researcher. I am also deeply grateful to Giovanni Nervi and Ludovic Jantzen for their supervision, their generous help, their constant availability throughout my time at the lab, and for the many pieces of advice that helped me find my way whenever the work became difficult. More broadly, my thanks go to all my colleagues at EPFL, whose company made the daily work lighter and far more enjoyable. At Politecnico di Torino, I would like to thank Professor Sandra Dulla for her support and dedication throughout these two years, and in particular for the trust and encouragement she gave me in pursuing this thesis abroad, which made this entire experience possible. As this study marks the end of a memorable and joyful university journey, I wish to thank the colleagues, professors, and friends who shared their path with mine, even if only for a brief moment.

Although a few words will never be enough to fully express it, my deepest gratitude goes to my family, for being by my side every step of the way.

Abstract

The transition from established system codes to higher-fidelity, three-dimensional multi-physics platforms is a central challenge in modern nuclear reactor analysis. The deployment of advanced simulation tools in nuclear reactor physics demands systematic validation against well-characterised experimental reference cases before they can be employed for design and safety analysis. This thesis, conducted within Work Package 2 of the EU-funded EVEREST project, contributes to that effort by applying and validating foamForNuclear — an open-source multi-physics platform built on the OpenFOAM framework — to the Budapest University of Technology and Economics (BME) Training Reactor (TR), a 100 kWth pool-type facility operated under natural circulation conditions, in the context of the present study.

The work is organised around three progressively coupled validation stages. The first section develops a stand-alone thermal-hydraulic model of the reactor pool, combining a porous-medium representation of the fuel assemblies with full CFD resolution of the surrounding water volume, and assesses its ability to reproduce measured natural-circulation transients. The second part carries out a systematic code-to-code verification of the built-in diffusion-based neutronics solver against Serpent 2 Monte Carlo reference calculations, covering the core eigenvalue, assembly power distributions, temperature reactivity feedback coefficients, and control rod worth. The third section exercises the coupled neutronics–thermal-hydraulics solver against experimental power and temperature measurements from the June 2022 campaign at the BME TR, with the established TRACE/PARCS solution serving as an additional computational reference.

Together, the results of these three validation streams, which demonstrate satisfactory agreement with both experimental data and the established TRACE/PARCS reference, contribute to assessing foamForNuclear as a high-fidelity tool for multi-physics reactor analysis and lay the groundwork for further validation activities within the EVEREST WP2 program.

Table of Contents

List of Figures

List of Tables

List of Acronyms

ACMI	Arbitrary Coupled Mesh Interface
AEKI	Atomic Energy Research Institute (*Atomenergia-Kutató Intézete*)
AMI	Arbitrary Mesh Interface
API	Application Programming Interface
BME	Budapest University of Technology and Economics (*Budapesti Műszaki és Gazdaságtudományi Egyetem*)
BRR	Budapest Research Reactor
CFD	Computational Fluid Dynamics
DTC	Doppler Temperature Coefficient
FFN	foamForNuclear
FVM	Finite Volume Method
KFKI	Central Research Institute for Physics (*Központi Fizikai Kutató Intézet*)
MTC	Moderator Temperature Coefficient
NPP	Nuclear Power Plant
RANS	Reynolds Averaged Navier Stokes
SMR	Small Modular Reactor
TH	Thermal-Hydraulic
TR	Training Reactor

VTT Technical Research Institute of Finland (*Valtion Teknillinen Tutkimuslaitos*)

VVER Water-Water Energetic Reactor (*Vodo-Vodyanoi Energeticheskiy Reaktor*)

WP2 Work Package 2

Chapter 1

Introduction

Nuclear safety analysis has always been only as reliable as the simulation tools underpinning it. For decades, a well-established generation of low-fidelity codes, largely one-dimensional in their treatment, built up the validation history and licensing credibility needed to earn institutional trust, and for the reactor concepts around which they were designed, they remain fit for purpose. Advanced designs now under development introduce geometries, flow regimes (e.g. pool thermal stratification, natural circulation, compact irregular cores), and coupled phenomena for which the spatial averaging these tools rely on represents a fundamental limitation that additional modelling effort can only partially compensate. Growing computational capacity has allowed the nuclear community to respond with a new generation of higher-fidelity multi-physics platforms sitting considerably closer to first principles, resolving three-dimensional flow and neutron transport in a coupled framework rather than collapsing them into one-dimensional analogues. The gains in physical fidelity are real: when dealing with open-pool type reactors cooled by natural circulation, the inherently asymmetrical behavior of thermodynamics along with the three-dimensional nature of flows make the 1-D system codes inherently unsuitable [1]. Thus, the move to higher fidelity multi-physics frameworks is justified; the weak point lies in validation. Experimental evidence for the coupled transient conditions these tools are ultimately meant to handle remains limited, and addressing that shortfall is the founding motivation of the EVEREST project. The EVEREST project (*Experiments for Validation and Enhancement of the REactor preSsure vessel fluence assessmenT*), funded by the European Union and initiated in 2024, was established to address this gap directly. The consortium brings together 15 institutions from 9 countries — universities, research centres, and industrial partners — with the shared goal of producing targeted, high-resolution experimental datasets at three European research reactors and using them to assess the predictive capability of advanced simulation platforms [2]. A secondary but equally important objective is to engage the broader nuclear community, utilities, regulators, and students, to build confidence in high-fidelity methods and support their eventual adoption in reactor analysis practice [2].

Work Package 2 (WP2), within which this thesis is situated, focuses on the devel-

opment and assessment of high-resolution and multi-physics modelling approaches for reactor analysis. A central question running through WP2 is whether the computational methodologies originally developed for large commercial power reactors transfer reliably to the more demanding conditions of research reactors: highly heterogeneous core configurations, pronounced neutron leakage at core boundaries, and operational regimes that fall outside the range of normal pressurized flow experiences owing to the use of natural circulation or low flow regimes dominated by complex buoyancy-driven thermal-hydraulic phenomena. These characteristics make research reactors particularly stringent test cases, and the Budapest University of Technology and Economics (BME) Training Reactor (TR) has been selected as a key demonstrator within WP2 for precisely this reason [2].

The BME TR is a pool-type reactor with a nominal thermal power of 100 kWth, cooled under standard operating conditions by natural circulation alone, in the absence of active pumping. This configuration poses a specific modelling challenge: buoyancy-driven flow in a large open pool is inherently three-dimensional, and the temperature stratification and recirculation structures that develop during power transients are not adequately captured by one-dimensional system codes. A fully three-dimensional treatment of the thermal-hydraulics is therefore not merely a matter of improved resolution, but of physical fidelity. At the same time, the reactor's well-documented core geometry, the availability of a complete neutronic benchmark dataset [3], and a dedicated series of thermocouple measurements taken during natural-circulation power transients [4] make the BME TR particularly well suited for the progressive validation of coupled simulation tools. The computational reference for this reactor was established by Ványi [5], who developed a detailed full-core model using the coupled TRACE/PARCS code system — both mature, independently validated codes — supplying the reference geometry, cross-section library, and integral benchmark results against which newer tools can be assessed.

The present work contributes to WP2 by applying and validating `foamForNuclear` [6], an open-source multi-physics platform developed at EPFL based on the `OpenFOAM` framework, to the BME TR. The work proceeds along two parallel lines. The first develops and validates a stand-alone three-dimensional thermal-hydraulic model of the reactor pool under natural circulation, demonstrating its ability to reproduce the measured temperature transients. The second is a systematic code-to-code verification of the built-in two-group neutron diffusion solver against `Serpent 2` Monte Carlo reference calculations [7], covering the effective multiplication factor, assembly-wise power distributions, moderator and Doppler temperature reactivity coefficients, and the integral worth of both control rod banks. Together, these two validation streams establish the foundation for the fully coupled neutronics–thermal-hydraulics simulation of the BME TR carried out in the final stage of the work.

The thesis is organized as follows. Chapter 2 opens with an overview of the Hungarian nuclear sector and provides a detailed technical description of the BME TR,

detailing its history, core geometry, and control rod configuration. Chapter 3 then introduces the numerical framework adopted for this work: `Serpent 2`, which serves both as a high-fidelity reference and a cross-section generation engine, and `foamForNuclear`, whose governing equations for neutron diffusion, single-phase thermal-hydraulics, and coupled behavior are derived in detail. The core modeling efforts are split across the subsequent chapters. Chapter 4 documents the stand-alone thermal-hydraulic model, focusing on mesh generation, the porous-medium formulation, and natural-circulation boundary conditions. Chapter 5 outlines the steps taken to generate the neutronic model, while Chapter 6 clarifies the underlying geometric and physical assumptions required to run the simulations. Finally, Chapter **??** gathers the core findings of this study, breaking them down into stand-alone thermal-hydraulic solutions, the neutronic validation against the Monte Carlo reference, and the coupled transient simulations compared directly with experimental data. Concluding remarks and outlooks for future research are summarized in Chapter 8.

Chapter 2

Hungary nuclear power situation

2.1 The Hungarian Nuclear Context

Nuclear energy currently plays a leading role in the Hungarian electric production, and is considered essential to reach net-zero emissions by 2050. The field is dominated by the state-owned Paks Nuclear Power Plant (NPP), featuring four Water-Water Energetic Reactor (VVER)-440/213 pressurized water reactors commissioned between 1982 and 1987.[1] Nowadays, these units have been uprated to a net power capacity of 506–508 MWe each, generating approximately 47.1% (16.017 TWh) of the country's electricity in 2024. Looking ahead, the Hungarian government is expanding the site through the Paks II project: this involves the construction of two new VVER-1200 units, with the main implementation phase starting as of 2026 [3, 8, 9].

Besides the power production plant, Hungary operates two research and training reactors: the Budapest Research Reactor (BRR), a tank-type reactor, moderated and cooled by light water (10 MW_{th}), primarily used for research purposes and medical isotope production; and the BME TR, with a nominal thermal power of 100 kWth, dedicated to the education and training of physicists and engineers [3].

Those well-established nuclear infrastructures are the result of a long-term energy strategy that began in the second half of the 20th century.

2.2 BME Training Reactor

2.2.1 Historical Evolution

Following the Second World War, Hungary's early nuclear research began, driven by the need for energy independence and the lack of significant domestic fossil fuel reserves. Under Soviet influence, the Hungarian government identified nuclear power as the most viable solution to meet the growing industrial demand. As a result, in 1959 at Csillebérc, on the outskirts of Budapest, a Soviet design research reactor started operating. This facility was commissioned at the Central Research

[1] The suffix '440' denotes the nominal design gross electrical power output in megawatts (MWe), while suffix '213' represents the specific model series.

Institute for Physics (KFKI), originally with a capacity of 2.5 MW_{th}; after 30 years of operation, it was refurbished by Hungarian experts and was again put into operation by the Atomic Energy Research Institute (AEKI) in 1993 with 10 MW_{th}, known as BRR [3].

This technological momentum, following the mandate given to the KFKI in the early 1960s, pushed the project for a dedicated educational reactor into its operational phase. The goal was to build a small research reactor right on the university campus—a task made possible by a long-standing partnership with the Kurchatov Institute[2] [10, 11]. In 1966, the first core design was finalized and assembled at KFKI for preliminary testing (coincidentally, that same year marked the approval for the Paks NPP design, while in 1969 its construction was started) [5].

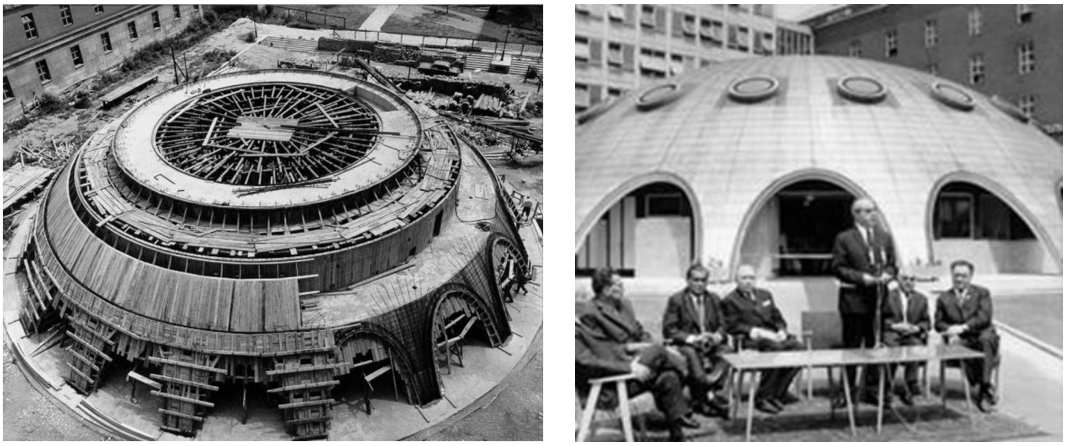

Figure 2.1: Historical phases of the BME Training Reactor: structural construction (left) and the official inauguration ceremony in May 1971 (right) [12].

The construction of the BME TR was completed in May 1971, when the reactor reached its first criticality, marking a milestone for the Budapest University of Technology and Economics (historical images of the construction and inauguration are shown in Figure 2.1). Initially operating at a modest thermal power of 10 kW with an excess reactivity of 71.2 ¢. Its first major evolution came in 1980: by replacing a graphite reflector block with a fresh fuel assembly, the facility was uprated to 100 kWth, more than doubling its excess reactivity to 156.3 ¢. Following decades saw a series of targeted instrumentation updates in 1990 and 2010, culminating in an extensive modernization program in 2016, touching everything from the reactor building and electrical infrastructure to the core's measurement systems. By 2023, the BME reactor had completed over 50 years of operation [5].

2.2.2 Technical Overview

The Training Reactor, located at the Central Campus of BME, is a pool-type research reactor, cooled and moderated by demineralized water. The reactor core is positioned

[2]Formerly the I.V. Kurchatov Institute of Atomic Energy, it is Russia's leading nuclear research and development center.

at the bottom of an aluminum cylindrical tank, approximately 6 m in height and 1.4 m in diameter, with a wall thickness of 20 mm. Biological shielding is provided by a 2 m thick concrete layer surrounding the vessel, while a 5 m water column above the core ensures protection in the vertical direction (see Figure 2.2a). A thin air-filled gap separates the aluminum tank from the concrete shield; to maintain radiological safety, a continuous air extraction system is installed at the top of the vessel. This open-vessel configuration is particularly advantageous for experimental activities, as it facilitates the insertion of various instrumentation.

The primary cooling circuit features an inlet nozzle at the bottom of the vessel that opens into a diffuser unit. This unit directs the coolant flow specifically toward the fuel assemblies and is housed within an aluminum box. Lateral windows in the box allow for coolant recirculation within the vessel (see Figure 2.2b). The outlet nozzle for the primary circuit is located 110 cm above the core, and a spillway is positioned at the top of the tank to manage the water level. Under standard operating procedures, the primary coolant temperature is maintained between 20°C and 60°C. Forced flow can be induced via the primary circuit pump, and the system temperature is regulated through a secondary cooling loop or an auxiliary electrical heater [5].

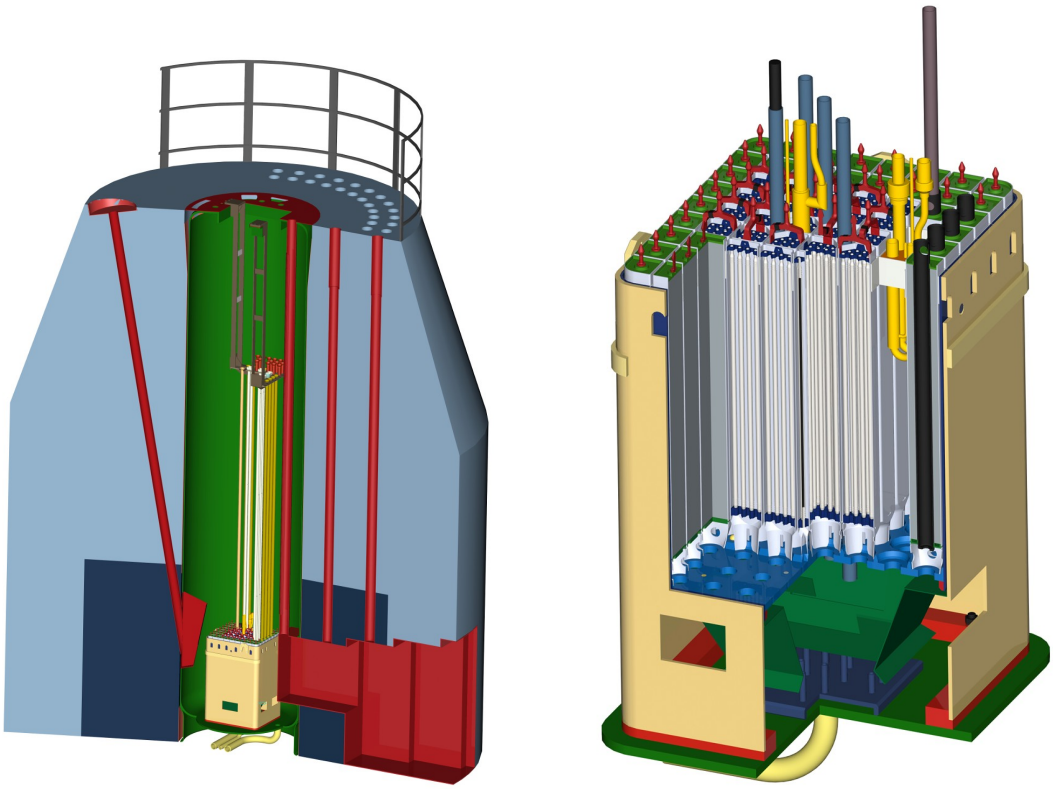

(a) Reactor core, vessel, and the surrounding concrete elements.

(b) The core, the core barrel structure, and the core window.

Figure 2.2: CAD illustration of BME TR [13].

The reactor core comprises twenty-four EK-10 fuel assemblies, categorized as

either regular (containing sixteen fuel pins in a rectangular lattice) or modified (with a reduced number of pins or an irregular lattice geometry). These assemblies contain uranium dioxide, nominally 10wt% enriched, dispersed in a metallic magnesium matrix with a 10 mm diameter and an active length of 500 mm. A fresh EK-10 fuel rod contains roughly 80 g of total uranium against 13.029 g of magnesium matrix, yielding a baseline U/Mg mass ratio of 6.14:1 [14]. The fuel region is surrounded by forty-one solid graphite reflector blocks. All internal structural components—including fuel cladding, assembly housings, and control rod guide tubes—are fabricated from an aluminum alloy with a purity exceeding 99.5% [15].

Reactivity control is managed by two safety rods and two control rods, with positions ranging from $z = 0$ mm (fully inserted) to $z = 600$ mm (fully withdrawn). The exact spatial allocation of these elements within the lattice is illustrated in Figure 2.3, and their detailed characteristics are given as follows:

- **Safety Rod 1 (SR1):** located between fuel assemblies E6 and F5, utilizing boron carbide pellets;
- **Safety Rod 2 (SR2):** located between fuel assemblies D5 and E4, also utilizing boron carbide pellets;
- **Manual Control Rod (MAN):** located between assemblies C6 and D5, consisting of boron carbide and operated manually;
- **Automatic Control Rod (AUT):** located between assemblies C4 and D3, consisting of a stainless steel tube coated with a thin cadmium layer. This rod can function in both manual and automatic modes via the reactor control unit.

Regarding the experimental capabilities, the setup includes twenty vertical wet channels and three dry channels, complemented by five horizontal neutron beam tubes. The facility is also equipped with a large irradiation tunnel and a three-branch pneumatic rabbit system for rapid sample transport. Reactor startup is assisted by a standard plutonium-beryllium (Pu-Be) neutron source, while core monitoring is provided by seven dedicated instrumentation tubes distributed around the core perimeter [5, 15].

Since its initial commissioning, the BME TR has not undergone a full core refueling. Significant modifications to the core loading were limited to the replacement of the F3 assembly with an identical fresh unit in 1977 and the addition of a regular assembly in position B6 in 1980 [5].

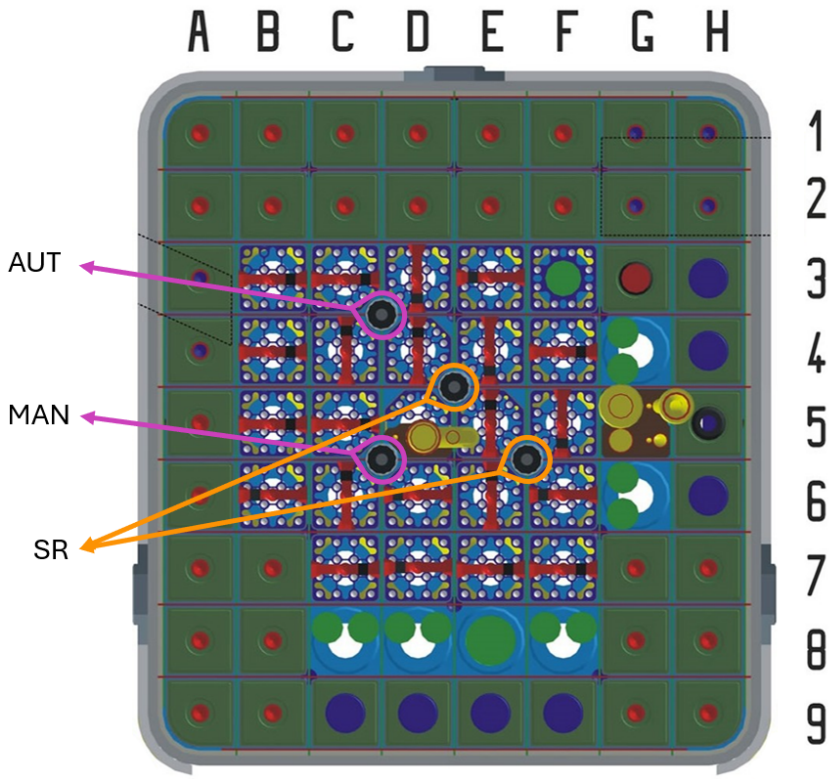

Figure 2.3: Horizontal cross section of the reactor core.

The three-dimensional CAD construction needed for the high-fidelity model implementation is based on on the work of [5, 16], integrated with official technical specifications from [15], providing sufficient information to carry out the presented work.

Chapter 3

Theoretical Framework and Numerical Tools

The simulation of nuclear reactor systems is inherently multidisciplinary, as no single physical domain can fully capture the core's behavior in isolation. A fundamental coupling exists between neutron kinetics, which determines the heat generation profile, and thermal-hydraulics, which governs the resulting temperature and density fields. These, in turn, modify the macroscopic cross-sections, directly affecting the reactivity and power profiles and the conditions under which the fission chain reaction is sustained.

Accounting for this circularity requires a transition from decoupled methods to a multiphysics framework. While the complexity of this framework depends on the specific reactor system and the transients under consideration, the feedback loop between neutronics and thermal-hydraulics represents the critical nexus for light-water reactor analysis. In a broader multiphysics perspective, other phenomena such as fuel behaviour, structural mechanics, coolant chemistry, and radionuclide transport may also significantly influence the system [17]; however, the present modeling procedure focuses exclusively on the primary coupling between neutronics and thermal-hydraulics.

This chapter establishes the theoretical basis for such a coupled approach, detailing the fundamental equations that describe the interplay between the neutron population and heat removal. These provide the foundation for the simulations carried out in this work using mainly `foamForNuclear` [6], a multiphysics platform for nuclear reactor applications built upon the `OpenFOAM` framework [18] (an open-source C++ library for computational fluid dynamics distributed by the OpenCFD), whose implementation is discussed in a dedicated section (see Section 3.2).

3.1 Governing Equations

3.1.1 Single-Phase Thermal-Hydraulic Conservation Equations

The thermal-hydraulic behaviour of the reactor coolant is described by the Navier–Stokes equations, a system of nonlinear partial differential equations expressing the conservation of mass, linear momentum, and energy for a viscous, heat-conducting fluid. In their general form these laws impose no constraint on the variation of material properties, and this generality is necessary here: light water, the sole coolant of the BME TR, exhibits non-trivial changes in density, viscosity, and thermal conductivity across the temperature range encountered during reactor operation. The derivation below follows the formulation in Munson et al. [19] and progressively specialises the governing equations to the single-phase, variable-property case appropriate for pool-type, natural-circulation flow.

For brevity and clarity, the temperature dependence of the fluid properties ($\rho(T)$) will be omitted in the subsequent equations. Nonetheless, it should be mentioned that the main thermophysical parameters considered in this work include density (ρ), thermal conductivity (k), and dynamic viscosity (μ), all of which are still temperature-dependent.

Mass conservation For a control volume fixed in space, the net outward mass flux through its boundary must equal the rate of depletion of mass stored within. Taking the limit to an infinitesimal volume yields the continuity equation [19]:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho\, \mathbf{u}) = 0, \tag{3.1}$$

where ρ is the local fluid density and $\mathbf{u}$ the velocity vector. When ρ is constant Eq. (3.1) collapses to the divergence-free condition $\nabla \cdot \mathbf{u} = 0$; however, this simplification is not admissible in the present context. The density of water decreases by roughly 1.5% between 20 °C and 60 °C, a variation that is the direct physical cause of the buoyancy-driven flow sustaining the natural circulation and thus, the reactor cooling. Retaining $\partial\rho/\partial t$ and the non-zero divergence ensures that the continuity equation is consistent with the temperature-dependent equation of state introduced in Section 3.1.2.

Momentum conservation Applying Newton's second law to a material fluid element and decomposing internal stresses into isotropic pressure and deviatoric viscous parts gives the Navier–Stokes equation in conservative form [19]:

$$\frac{\partial (\rho \mathbf{u})}{\partial t} + \nabla \cdot (\rho\, \mathbf{u} \otimes \mathbf{u}) = -\nabla p + \nabla \cdot \tau + \rho \mathbf{g}, \tag{3.2}$$

where p is the thermodynamic pressure and $\mathbf{g}$ the gravitational acceleration. For a Newtonian fluid the viscous stress tensor τ is related to the rate-of-strain tensor by

the constitutive expression [19]:

$$\tau = \mu\Big[\nabla\mathbf{u} + (\nabla\mathbf{u})^T - \tfrac{2}{3}(\nabla\cdot\mathbf{u})\,\mathbf{I}\Big], \tag{3.3}$$

where μ is the temperature-dependent dynamic viscosity and $\mathbf{I}$ the identity tensor. The term proportional to $\nabla\cdot\mathbf{u}$ represents normal viscous stresses arising from volumetric deformation; it is non-zero whenever density varies with temperature and is therefore retained for consistency.

The body-force term $\rho\,\mathbf{g}$ warrants particular attention. In the BME TR, which operates without forced-circulation pumps for the purpose of this work, it constitutes the sole source of momentum that drives the coolant through the core. Heated water ascending through the fuel assemblies is less dense than the cooler water in the surrounding pool; the resulting hydrostatic pressure difference sets up a convective current that simultaneously removes the fission heat and replenishes the inlet with cool coolant. This mechanism, referred to as natural circulation, is self-regulating: a higher power level produces a larger temperature rise, which increases the density contrast and consequently the driving force.

A widely used simplification is the Boussinesq approximation, in which ρ is treated as constant everywhere except inside the buoyancy term $\rho\mathbf{g}$, where it is linearised as $\rho \approx \rho_0[1-\beta(T-T_0)]$ with β the isobaric thermal expansion coefficient evaluated at a reference state [19]. For water at 20 °C, $\beta \approx 2.07\times10^{-4}$ K^{-1}; a temperature excursion of $\Delta T = 40$ K produces $\beta\,\Delta T \approx 0.008$, which is formally within the validity of the linearisation. Since the accuracy of a natural-circulation simulation is critically sensitive to the buoyancy driving term, the full nonlinear $\rho(T)$ relationship is retained rather than resorting to this approximation (detailed in Section 3.1.2).

Energy conservation The first law of thermodynamics applied to a fluid element in motion is most conveniently expressed in terms of the total enthalpy $h_{tot} = h + e_K$, where h is the specific enthalpy and $e_K = \frac{1}{2}|\mathbf{u}|^2$ is the specific kinetic energy. In conservative form, the transport equation reads [19]:

$$\begin{aligned}\frac{\partial}{\partial t}[\rho(T)(h+e_K)] + \nabla\cdot[\rho\,\mathbf{u}(h+e_K)] = \nabla\cdot(k\,\nabla T) + \frac{\partial p}{\partial t} \\ + \rho\,\mathbf{u}\cdot\mathbf{g} + \tau : \nabla\mathbf{u} + q''',\end{aligned} \tag{3.4}$$

where $k(T)$ is the thermal conductivity, q''' is the volumetric heat-source density, and $\tau:\nabla\mathbf{u}$ represents viscous dissipation. The advantage is twofold: the pressure-work term $\partial p/\partial t$ and the gravitational-work term $\rho\,\mathbf{u}\cdot\mathbf{g}$ arise naturally, without additional approximations, and the equation remains thermodynamically consistent regardless of the equation of state used for $\rho(T)$.

3.1.2 Temperature-Dependent Fluid Properties

The coupling between Eqs. (3.2) and (3.4) is mediated entirely through the fluid properties $\rho(T)$, $\mu(T)$, and $k(T)$. Reference values for liquid water are taken from

the IAPWS-97 standard [20], and polynomial fits of the form

$$\phi(T) = \sum_{n=0}^{N} a_n \, T^n \tag{3.5}$$

are obtained by least-squares regression to NIST tabulated data over the temperature interval of interest (approximately 20–60 °C).

Density requires a second-order polynomial ($N = 2$) because a linear fit misrepresents the slight curvature of the $\rho(T)$ curve and would introduce a systematic error in the buoyancy driving term. The remaining properties — dynamic viscosity and thermal conductivity — are similarly implemented using temperature-dependent polynomial fits to accurately capture the coolant's variations over the operational temperature interval, as justified in Section 4.4.

3.1.3 Neutron Transport and the Monte Carlo Method

The physical phenomena occurring in a nuclear fission reactor are governed by the Boltzmann neutron transport equation, which describes the evolution of the angular neutron flux $\psi(\mathbf{r}, E, \hat{\Omega}, t)$—the number of neutrons per unit volume, energy, and solid angle—at position $\mathbf{r}$, energy E, direction $\hat{\Omega}$, and time t. A balance over an infinitesimal phase-space element yields [21, 22]:

$$\begin{aligned} \frac{1}{v(E)} \frac{\partial \psi(\mathbf{r}, E, \hat{\Omega}, t)}{\partial t} + \hat{\Omega} \cdot \nabla \psi(\mathbf{r}, E, \hat{\Omega}, t) + \Sigma_t(\mathbf{r}, E)\, \psi(\mathbf{r}, E, \hat{\Omega}, t) = \\ \int_0^\infty \int_{4\pi} \Sigma_s(\mathbf{r}, E' \to E, \hat{\Omega}' \to \hat{\Omega})\, \psi(\mathbf{r}, E', \hat{\Omega}', t)\, d\hat{\Omega}'\, dE' \\ + \frac{\chi(E)}{4\pi} \int_0^\infty \nu\Sigma_f(\mathbf{r}, E')\, \phi(\mathbf{r}, E', t)\, dE' + S_{\text{ext}}(\mathbf{r}, E, \hat{\Omega}, t), \end{aligned} \tag{3.6}$$

where v is the neutron speed, Σ_t the total macroscopic cross section, Σ_s the scattering kernels, $\nu\Sigma_f$ the fission production cross section, $\chi(E)$ the fission neutron spectrum, $\phi = \int_{4\pi} \psi \, d\hat{\Omega}$ the scalar flux, and S_{ext} any external neutron source. For steady-state criticality calculations, the time derivative is set to zero and fission is divided by the effective multiplication factor k_{eff}, reducing Eq. (3.6) to an eigenvalue problem.

Rather than solving algebraic systems over a discretized phase space, deterministic approaches can be replaced by stochastic ones. In Monte Carlo methods, the key is that the transport equation (Eq. (3.6)) is resolved by tracking and simulating every single neutron's history. Each neutron is tracked from birth to absorption or leakage through a sequence of events scattering, fission, radiative capture—whose outcomes are sampled from the continuous-energy nuclear data libraries. The k_{eff} eigenvalue is estimated iteratively via the power method: across each active cycle, a batch of source neutrons is transported, and k_{eff} is accumulated as the ratio of fission neutrons produced to those absorbed or lost by leakage:

$$k_{\text{eff}} = \frac{\displaystyle\int_V \nu\Sigma_f(\mathbf{r}, E)\,\phi(\mathbf{r}, E)\,dV\,dE}{\displaystyle\int_V \Sigma_a(\mathbf{r}, E)\,\phi(\mathbf{r}, E)\,dV\,dE \;+\; \oint_{\partial V} \mathbf{J}\cdot\hat{n}\,dS\,dE}. \tag{3.7}$$

Because nuclear data are treated continuously in energy, this approach is free from the group-structure and spatial homogenization approximations that characterize diffusion-based solvers, making `Serpent 2` a natural benchmark reference for research reactor analysis.

3.2 foamForNuclear

foamForNuclear (FFN) is a unified, open-source multiphysics platform built on the OpenFOAM framework. Established in 2025, FFN was conceived to provide a modular and extensible environment for the design and analysis of both fission and fusion systems. The application was developed by combining two well-established codes: `GeN-Foam`, designed for reactor-scale multiphysics, and `OFFBEAT`, dedicated to fuel performance; consolidated and validated by leading institutions, such as the **École Polytechnique Fédérale de Lausanne (EPFL)**, the **Paul Scherrer Institut (PSI)**, and **Texas A&M University**.

The FFN code is designed as separate modules that solve individual governing equations. This design permits the coupling of arbitrary physics in several domains via the `multiRegion` approach. In this context, each region acts as a logical subdivision of the simulation, containing its own dedicated mesh, fields, and associated dictionaries. This structure enables the use of different physics models, without requiring direct modifications to the source code [6]. An overview of the main `solvers`[1] is provided below:

- The Thermal-Hydraulics module includes both single-phase RANS and two-phase Euler–Euler (six-equation) models for water and sodium flow. In order to optimize computational resources, a porous-medium capability allows for "coarse-mesh" modeling by homogenizing complex structures like fuel assemblies. Moreover, this module contains features for the modeling of particular reactor components, such as pumps and heat exchangers [23].

- The Neutronics module features a series of different neutronics solvers, including point kinetics, multi-group diffusion (adjoint formulations are included), and higher-order SP_3 and S_N transport approximations. The nuclear data management is based on a Radial Basis Functions (RBF) based interpolation scheme to accurately map cross-sections across various perturbed states, for example fluctuations in coolant density or fuel temperature [23].

- The Fuel Performance and Thermomechanics modules, based on the `OFFBEAT` application, handle structural mechanics and fuel-specific phenomena. The

[1] In the FFN architecture, a "solver" refers to a specific C++ class containing the implementation of the governing equations and physical models for a given domain.

computational capabilities include modeling non-linear elasticity, plasticity, and creep, moving beyond traditional linear-elastic assumptions. Additionally, it accounts for the evolution of fuel microstructures through models for fission gas release, swelling, densification, and relocation, while blending material properties that evolve according to burnup levels [24].

The present work focuses exclusively on the coupling between the Neutronics and Thermal-Hydraulics modules. The choice was deliberately made to align with the goals of the EVEREST project Work Package 2 (WP2), as the benchmarking process relies on the experimental data and reference cases specifically developed for these two physics. Consequently, the advanced fuel behavior capabilities provided by `OFFBEAT` are beyond the scope of this work.

3.3 Thermal-Hydraulic Modeling

3.3.1 Mesh Generation

The computational mesh was generated using `blockMesh`, the native structured mesher of the `OpenFOAM` framework [18], whose input dictionary `blockMeshDict` was produced through the dedicated `foamForNuclear` Application Programming Interface (API). The full geometry was built through a multi-block strategy, generating distinct sub-domains (reactor core, lower plenum, upper pool) that were subsequently merged into a single conformal mesh. Custom functions were implemented to align nodal coordinates at the sub-domain interfaces, a step that proved essential to avoid hanging nodes and to guarantee flux conservation across the merged boundaries. The API enables a parametric and automated definition of vertices, blocks, and boundary patches, reducing the risk of manual errors and simplifying geometry updates. For the hexagonal fuel assembly lattice, the `Honeycomb` [2] utility was employed to simplify the creation of the assembly grid: each assembly is assigned a unique numeric identifier attached to its constituent cells, making it straightforward to partition the mesh into named `cellZone`s. This one-to-one mapping between cell groups and assembly regions allows the porous-media solver, described in Section 3.3.2, to attach zone-specific geometrical and physical properties—porosity, hydraulic diameter, drag coefficients—directly to the corresponding mesh regions, without any additional interpolation or post-processing step.

3.3.2 Porous Media Approach

Resolving the detailed pin-bundle geometry of every fuel assembly on a full-core mesh is not computationally feasible. Doing so would require resolving inter-pin gaps of order one millimetre across a domain spanning several hundred litres, pushing memory and runtime well beyond what is practical for a transient simulation. The

[2] Honeycomb is an open-source web application for the visualization and manipulation of nuclear reactor lattices, used in this work to define the core geometry prior to mesh generation [25].

porous-medium approach addresses this by homogenising each assembly into an equivalent continuum, where the macroscopic pressure drop and heat transfer that would arise from the explicit solid structure are recovered through effective source terms in the governing equations [6].

The key quantity introduced by this treatment is the fluid volume fraction α, defined as the ratio of fluid to total (fluid plus solid) volume in each computational cell. In the open pool surrounding the core, $\alpha = 1$ and the equations reduce to their standard single-phase form. Inside each assembly zone, $\alpha < 1$ and is fixed by the geometric parameters of the pin lattice — pin diameter, pitch, and wrapper geometry — for each assembly type present in the BME TR core. The `foamForNuclear` thermal-hydraulic solver is built on the Euler–Euler six-equation framework, in which α appears explicitly in all three conservation laws [6]. For the single-phase coolant of the BME TR, inter-phase mass exchange vanishes and the equations take the following form.

Mass conservation:

$$\frac{\partial}{\partial t}(\alpha\rho) + \nabla\cdot(\alpha\rho\, u) = 0. \tag{3.8}$$

Momentum conservation:

$$\frac{\partial}{\partial t}(\alpha\rho\, u) + \nabla\cdot(\alpha\rho\, u\otimes u) = -\alpha\nabla p + \nabla\cdot\Big[\alpha\rho\nu\,\mathcal{T}\cdot\Big(\nabla u + (\nabla u)^T - \tfrac{2}{3}(\nabla\cdot u)\, I\Big)\Big] + S_u, \tag{3.9}$$

Energy conservation:

$$\frac{\partial}{\partial t}[\alpha\rho(h+e_K)] + \nabla\cdot[\alpha\rho\, u(h+e_K)] = \nabla\cdot(\alpha\, a\,\mathcal{T}\cdot\nabla h) - \alpha\frac{\partial p}{\partial t} + \alpha\rho\, u\cdot g + S, \tag{3.10}$$

where ν is the kinematic viscosity, a the thermal diffusivity, and $e_K = \frac{1}{2}|u|^2$ the specific kinetic energy, treated explicitly [6]. The rank-2 tensor $\mathcal{T}$ is the tortuosity tensor, which accounts for the geometric restriction of diffusive pathways by pre-multiplying the diffusive fluxes. When $\alpha = 1$ and $\mathcal{T} = I$, Eqs. (3.8)–(3.10) collapse to the single-phase Navier–Stokes system of Eqs. (3.1)–(3.4). The momentum source S_u in Eq. (3.9) collects gravitational body forces, $\alpha\rho g$, together with the resistive drag exerted by the pin structure on the fluid [6]:

$$S_{u,\mathrm{drag}} = -K_s\cdot u, \tag{3.11}$$

where K_s is the fluid-structure resistance tensor. Its diagonal components are computed from empirical pressure-drop correlations suited to the specific lattice geometry of each assembly type, as detailed in Section 4.4. In free-flow regions $K_s = 0$ and Eq. (3.11) contributes nothing to the momentum balance.

The energy exchange between the solid fuel and the flowing coolant is mediated by a sub-scale pin model embedded within each porous cell [6]. Concretely, at each time step, the temperature profile inside the fuel pin is resolved through a transient one-dimensional Finite Volume Method (FVM) discretization in the radial direction. The governing equation for the transient radial heat conduction within a pin layer is

expressed as:

$$\rho C_p \frac{\partial T}{\partial t} = \frac{1}{r}\frac{\partial}{\partial r}\left(rk\frac{\partial T}{\partial r}\right) + q''' \tag{3.12}$$

where ρC_p and k are the volumetric heat capacity and thermal conductivity of the respective layer (fuel or cladding), and q''' is the volumetric power density supplied by the neutronics solver, scaled by the fraction of power deposited directly in the fuel. At the fuel–cladding interface ($r = r_f$), the thermal resistance of the gas gap is explicitly accounted for by imposing a jump condition mediated by a gap conductance h_{gap}:

$$q''_{\text{gap}} = h_{\text{gap}} \left(T_{\text{fuel,out}} - T_{\text{clad,in}}\right) \tag{3.13}$$

The implicit linear system resulting from the multi-node radial discretization of the fuel and cladding layers is solved numerically for each porous cell to determine the updated radial temperature profile $T(r)$, from which the outer cladding surface temperature $T_{\text{cl,out}}$ is extracted. This boundary temperature is subsequently used to evaluate the volumetric heat source deposited into the coolant phase:

$$q'''_{\text{fluid}} = h_{\text{conv}} a_{\text{HT}} \left(T_{\text{cl,out}} - T_f\right) \tag{3.14}$$

where T_f is the local bulk coolant temperature, h_{conv} is the convective heat transfer coefficient obtained from the selected Nusselt correlation, and a_{HT} is the interfacial wetted surface area density of the porous zone. The latter is a purely geometric quantity computed by the solver as:

$$a_{\text{HT}} = \frac{2\alpha}{r_{\text{cl,out}}} \tag{3.15}$$

where α represents the solid volume fraction of the pins within the cell. This volumetric source term is finally injected into the fluid energy equation, closing the coupled neutronic–thermal–hydraulic feedback loop.

3.3.3 Turbulence Modeling

Within the fuel assemblies, the flow is characterized by Reynolds numbers in the range 500–1500 [26], placing it in the laminar-to-transitional regime. At these conditions, the drag and heat transfer correlations embedded in the porous-medium formulation already account implicitly for the effects of turbulence at the sub-pin scale, so no explicit turbulent transport needs to be resolved inside the core region since the coolant/structure interaction dominates. Nevertheless, a turbulence model must still be active throughout the domain for a different reason: at the interface between the porous core and the open pool, the solver requires physically consistent values of k and ϵ as boundary conditions for the free-fluid region. Imposing zero turbulent quantities at the core outlet and a fully developed turbulent field immediately outside would produce an unphysical discontinuity that destabilizes the solution [6].

To address this, `foamForNuclear` implements a hybrid porous k-ϵ formulation in which the transport equations are solved differently depending on the local porosity [6]. In clear fluid regions ($\epsilon = 1$), the standard Reynolds Averaged Navier Stokes (RANS) k-ϵ transport equations are solved:

$$\frac{\partial(\alpha\rho k)}{\partial t} + \nabla \cdot (\alpha\rho k\,\mathbf{u}) - \nabla \cdot (\alpha\rho D_k \nabla k) = \xi\alpha P - \xi\alpha\rho\epsilon, \tag{3.16}$$

$$\frac{\partial(\alpha\rho\epsilon)}{\partial t} + \nabla \cdot (\alpha\rho\epsilon\,\mathbf{u}) - \nabla \cdot (\alpha\rho D_\epsilon \nabla \epsilon) = \xi\frac{C_1\epsilon}{k}\left(P + C_3\frac{2}{3}k\nabla\cdot\mathbf{u}\right) - \xi C_2\rho\frac{\epsilon^2}{k}, \tag{3.17}$$

where P is the turbulent production term and D_k, D_ϵ are the effective diffusivities for k and ϵ respectively. Inside porous zones ($\epsilon < 1$), the switch parameter

$$\xi = \begin{cases} 1, & \text{free-fluid regions,} \\ 0, & \text{porous zones,} \end{cases} \tag{3.18}$$

deactivates the standard production and dissipation terms and replaces them with a relaxation source that drives k and ϵ toward user-prescribed equilibrium values k_{eq} and ϵ_{eq} [6]:

$$S_k = (1-\xi)\,\alpha\rho\,\frac{|\mathbf{u}|}{L_{\text{conv}}}\,(k_{\text{eq}} - k)\,, \qquad S_\epsilon = (1-\xi)\,\alpha\rho\,\frac{|\mathbf{u}|}{L_{\text{conv}}}\,(\epsilon_{\text{eq}} - \epsilon)\,, \tag{3.19}$$

where L_{conv} is a user-defined convergence length that controls the spatial extent of the transition between the two regimes. The equilibrium values are prescribed from empirical correlations based on a turbulence intensity I and a length scale L, following standard practice [6]. In this way, the model does not attempt to resolve physical turbulence inside the core, where the drag and heat transfer correlations already incorporate those effects, but instead provides a smooth and physically grounded estimate of k and ϵ at the core-pool boundary. These values then serve as inlet conditions for the standard k-ϵ equations active in the open pool, ensuring a continuous and stable transition across the porous-to-clear interface without introducing spurious gradients in the turbulent fields.

A further benefit of the k-ϵ model is that it naturally yields the turbulent viscosity

$$\mu_t = C_\mu\rho\frac{k^2}{\epsilon}, \tag{3.20}$$

where $C_\mu = 0.09$. Added to the molecular viscosity in the momentum equations, μ_t raises the effective diffusivity of the velocity field. Natural circulation is driven by the small density contrasts that temperature gradients produce in the coolant; velocities are low, and the resulting convective-to-diffusive ratio in the momentum equation is close enough to unity that the solution offers little inherent resistance to perturbations. Minor numerical disturbances accumulate, and eventually develop into oscillations capable of stalling convergence altogether. The extra dissipation

carried by μ_t provides the margin that keeps this from happening, and its presence throughout the domain is consequently as much a numerical necessity as a physical one.

The turbulence model adopted for the porous regions relies on standard wall functions to account for near-wall dynamics without resolving the viscous sublayer directly. The theoretical validity of standard wall functions requires the centroid of the first computational cell adjacent to any solid boundary to reside within the logarithmic region of the turbulent boundary layer. This condition is conventionally expressed as $30 \lesssim y^+ \lesssim 300$, where $y^+ = y\, u_\tau/\nu$ is the dimensionless wall distance, u_τ is the local friction velocity, and ν is the kinematic viscosity.

3.4 Neutronics Modeling

The deterministic characterization of the core neutronics within the FFN framework relies on the `GeN-Foam` multi-physics application.

The link between high-fidelity Monte Carlo lattice calculations and the deterministic core solver is established via the generation of multi group macroscopic cross sections. A dedicated Python utility, `nuclearData.py` [27], which reads the Serpent output and assembles the cross-section input files in the format expected by Gen-Foam's `nuclearData` directory is leveraged. No manual intervention on the raw Serpent output is required beyond configuring the universe mapping within the script. To support the reactivity feedback analysis described later on (see Section **??**), the cross-section generation is repeated for a set of operating conditions. For the Doppler Temperature Coefficient (DTC), separate Serpent runs are performed at different fuel temperatures using the on-the-fly TMS (Target Motion Sampling, `tms`) Doppler broadening method. The resulting cross sections at discrete fuel temperatures are then provided to GeN-Foam as a tabulated dataset. For the Moderator Temperature Coefficient (MTC), an analogous procedure is followed: cross sections are extracted at different moderator densities and temperatures, capturing the combined effect of thermal scattering changes and density variation on the neutron spectrum. The coolant temperature density changes are linked by the equation of state described in Section 3.1.2. Cladding temperature variations and specific aluminum feedback effects are omitted from this parameterization, as the primary objective is to verify the foamForNuclear coupling framework using the dominant fuel and moderator feedback mechanisms, for which reference data were available in [5]. This parameterization strategy provides foamForNuclear with the numerical inputs necessary to dynamically resolve reactivity variations under both nominal and off-nominal operating conditions.

With the cross-section datasets in place, the neutronics model of the BME TR is built using `GeN-Foam` application that solves the multi-group neutron diffusion equation in its k-eigenvalue formulation, iterating on the fission source until the convergence of both the multiplication factor and the scalar flux distribution is

reached. Starting from the P_1 truncation of the spherical harmonic expansion of the angular flux and applying Fick's law for the neutron current:

$$\mathbf{J}(\mathbf{r}, E) = -D(\mathbf{r}, E)\,\nabla\phi(\mathbf{r}, E), \qquad D = \frac{1}{3\Sigma_{\text{tr}}}, \tag{3.21}$$

the transport equation reduces to the multi-group diffusion equation [28]:

$$\begin{aligned} -\nabla \cdot D_g(\mathbf{r})\nabla\phi_g(\mathbf{r}) + \Sigma_{r,g}(\mathbf{r})\phi_g(\mathbf{r}) = &\sum_{g' \neq g} \Sigma_{s,g'\to g}(\mathbf{r})\phi_{g'}(\mathbf{r}) \\ &+ \frac{\chi_g}{k_{\text{eff}}} \sum_{g'=1}^{G} \nu\Sigma_{f,g'}(\mathbf{r})\phi_{g'}(\mathbf{r}), \qquad g = 1, \ldots, G, \end{aligned} \tag{3.22}$$

where $\Sigma_{r,g} = \Sigma_{t,g} - \Sigma_{s,g\to g}$ is the removal cross section. The group constants D_g, $\Sigma_{r,g}$, $\Sigma_{s,g'\to g}$, $\nu\Sigma_{f,g}$, and χ_g are supplied as homogenised, flux-weighted averages computed by `Serpent 2` and imported into the solver's `neutronicProperties` dictionary. The fundamental eigenvalue k_{eff} is computed iteratively through standard source-iteration schemes until numerical convergence criteria for both the multiplication factor and the scalar flux field are simultaneously satisfied.

Among its various kinetics options, `GeN-Foam` for scenarios where spatial flux perturbations remain localized or asymptotically stable, can employs a zero-dimensional Point Kinetics (PK) formulation [21]:

$$\frac{dn(t)}{dt} = \frac{\rho(t) - \beta}{\Lambda}\, n(t) + \sum_{i=1}^{I} \lambda_i\, C_i(t), \tag{3.23}$$

$$\frac{dC_i(t)}{dt} = \frac{\beta_i}{\Lambda}\, n(t) - \lambda_i\, C_i(t), \qquad i = 1, \ldots, I, \tag{3.24}$$

where $n(t)$ is proportional to reactor power, C_i is the precursor concentration in delayed group i, λ_i its decay constant, β_i the delayed neutron fraction, $\beta = \sum_i \beta_i$ the total delayed neutron fraction, and Λ the prompt neutron generation time. The kinetic parameters β_i, λ_i, and Λ are extracted from full-geometry `Serpent 2` calculations; β_{eff} in particular is obtained via the Iterated Fission Probability (IFP) adjoint-weighting method [29].

Originally developed at Technical Research Institute of Finland (VTT) since 2004 as a lattice physics tool, `Serpent 2` has grown into a general-purpose solver capable of full-core eigenvalue calculations, burnup analysis, and multiphysics coupling. Its widespread adoption in the reactor physics community is due to a combination of computational efficiency—achieved through advanced variance reduction techniques and on-the-fly Doppler broadening—and a flexible constructive solid geometry (CSG) description that allows a faithful reproduction of complex core layouts [7, 30].

Both light water and graphite exhibit pronounced forward-peaked scattering

kernels, particularly at epithermal energies, which violates the isotropic scattering assumption underlying diffusion theory [21]. To account for this anisotropy, a transport correction is applied during group-constant condensation. The corrected transport cross section is formally expressed as:

$$\Sigma_{\mathrm{tr}}(\mathbf{r}, E) = f(E)\,\Sigma_t(\mathbf{r}, E), \tag{3.25}$$

where $f(E)$ is a pre-tabulated, energy-dependent correction curve defined in an arbitrary group-discretised form [31]. This correction propagates into the diffusion coefficients via the P_1 leakage model, which accounts for the first-order anisotropy of the flux gradient at assembly level [31].

Group-constant condensation is performed over user-defined spatial universes via the `set gcu` directive, which instructs `Serpent 2` to compute homogenised cross sections during the full-core simulation. The extracted parameters include diffusion coefficients, macroscopic absorption and fission cross sections, neutron yield, fission spectrum, and the scattering matrix for each energy group and spatial zone [31].

3.4.1 Cross-section generation

The few-group constants written by `Serpent 2` are the product of a multi-step mathematical procedure whose underlying assumptions directly condition the accuracy of the subsequent diffusion calculation. A clear account of that procedure—from the continuous-energy Monte Carlo solution to the form consumed by `foamForNuclear`—is therefore a necessary complement to the numerical results discussed in later sections.

Spatial homogenisation and energy condensation. The objective of spatial homogenisation is to replace the heterogeneous, continuous-energy description of an assembly with effective group-wise constants that, when applied to a coarse homogeneous geometry, preserves the reaction rates of the detailed calculation. Following the formalism of Leppänen et al. [32], the homogenised cross section for group g over a spatial region V is defined by the flux-weighted integral

$$\bar{\Sigma}_g = \frac{\displaystyle\int_V \mathrm{d}^3 r \int_{E_g}^{E_{g-1}} \mathrm{d}E\; \Sigma(\mathbf{r}, E)\,\phi(\mathbf{r}, E)}{\displaystyle\int_V \mathrm{d}^3 r \int_{E_g}^{E_{g-1}} \mathrm{d}E\; \phi(\mathbf{r}, E)}, \tag{3.26}$$

where $\phi(\mathbf{r}, E)$ is the scalar neutron flux, the spatial integration extends over the homogenised universe volume, and the energy integration spans the interval of group g. The physical content of Eq. (3.26) is compact: weighting by the local flux causes material zones carrying higher reaction rates to contribute proportionally more to the average, so the resulting group constant inherits the actual spectral and spatial shape of the neutron distribution. In a Monte Carlo calculation, both integrals are evaluated directly through continuous-energy reaction-rate tallies, which

automatically accounts for resonance self-shielding and spatial heterogeneity without supplementary approximations [32].

For practical reasons, `Serpent 2` does not evaluate Eq. (3.26) in a single step. The procedure is split into two stages: the continuous-energy cross sections are first accumulated over an intermediate multi-group structure (the default WIMS 69-group energy grid), and the resulting multi-group constants $\bar{\Sigma}_h$ are subsequently collapsed to the final few-group structure using the in-medium condensation spectrum Φ_h:

$$\bar{\Sigma}_g = \frac{\sum_{h \in g} \bar{\Sigma}_h \, \Phi_h}{\sum_{h \in g} \Phi_h}, \tag{3.27}$$

where h indexes the intermediate group and the summation runs over all intermediate groups belonging to the final group g [32]. This second-stage condensation does not forfeit the accuracy gained in the first stage: the dominant self-shielding effects are captured by the continuous-energy integration, and the two-stage structure merely divides the bookkeeping. The condensation spectrum Φ_h can be drawn either from the infinite-medium solution or from a B_1-leakage-corrected critical spectrum.

Cross-section handling in `foamForNuclear`. The group constants produced by `Serpent 2` are read by `foamForNuclear` through dedicated conversion scripts that reformat the `Serpent 2` output into the solver's native dictionary structure [6]. A minimum input consists of one nominal cross-section set per spatial zone; an arbitrary number of additional perturbed sets can be supplied to parametrise the cross sections as a function of any state variable—fuel temperature, coolant temperature, coolant density—that can be mapped from the thermal-hydraulic or neutronics mesh at run time.

Interpolation between states is handled through a Radial Basis Function (RBF) framework [6]. The RBF approach guarantees exact reproduction of the cross sections at each supplied data point, a direct consequence of the interpolant being fitted to those points; values between data points are recovered through smooth interpolation. Among the available kernels, polyharmonic splines are adopted as the default choice since they require no tuning parameters. For each state variable, the interpolation can be performed on the variable itself, its logarithm, or its square root.

Rodded and unrodded cross-section sets: the `CRmove` utility. Control rod insertion introduces a material discontinuity that cannot be represented by a continuous interpolation between perturbed states: each cell in the rod-swept channel must carry either the fully inserted absorber set or its uninserted counterpart, depending on whether the rod tip has passed through that cell's centre at the current instant. Two dedicated cross-section sets are therefore prepared for each rod region: Σ_{CR}, corresponding to the fully inserted absorber, and $\Sigma_{\mathrm{follower}}$, corresponding to the uninserted state.

The spatial assignment of the two sets at each time step is governed by the `CRmove` utility. For each rod, the `CRmove` dictionary specifies the insertion speed v, the transient start and end times, and the parameter `initialDistanceFromMeshCR` d_0, which locates the initial rod-tip position below the top of the follower zone. The current tip coordinate is updated at each time step as

$$z_{\text{tip}}(t) = z_{\text{follower,top}} - d_0 - v \cdot \big(\max(t,\, t_{\text{start}}) - t_{\text{start}}\big), \tag{3.28}$$

where $z_{\text{follower,top}}$ is the axial coordinate of the highest cell centre in the follower zone and v is positive downward. The cross-section set assigned to each cell is then determined by a step function on the cell-centre coordinate z_c:

$$\Sigma(z_c) = \begin{cases} \Sigma_{\text{CR}} & \text{if } z_c > z_{\text{tip}} \quad \text{(follower zone)} \\ \Sigma_{\text{follower}} & \text{if } z_c < z_{\text{tip}} \quad \text{(CR zone)} \end{cases} \tag{3.29}$$

and symmetrically for the complementary zone. The comparison is made against cell *centres*, not cell faces; no volumetric mixing of the two sets is performed. Each cell receives either the fully rodded or the fully unrodded set, determined exclusively by whether its centre lies above or below z_{tip}.

3.4.2 `foamForNuclear` Coupling Architecture

Within `foamForNuclear`, the interaction between the thermal-hydraulics and neutronics solvers is handled through a coupled segregated strategy [6]. Rather than iterating between physics within each time step, the platform advances each solver sequentially: at the beginning of every time step, all inter-physics fields are exchanged simultaneously through the `meshHandler` mapping infrastructure, after which each solver is advanced once in turn before the simulation proceeds to the next time step. This operator-split approach introduces a one-step lag in the feedback between physics, meaning that the neutronics solver at time t^n operates on thermal-hydraulic quantities computed at t^{n-1}; the resulting power field is however immediately available to the thermal-hydraulics solver within the same time step. For the slow feedback timescales characteristic of pool-type research reactors operating under natural circulation, this approximation is considered acceptable.

The point-kinetics solver follows the hybrid formulation described in [6], in which the total reactor power $P(t)$ is governed by a zero-dimensional ordinary differential equation, while the spatial distribution of delayed neutron precursor concentrations is tracked on the three-dimensional neutronics mesh to account for possible precursor drift effects. The fission power density field distributed over the fuel region is obtained by scaling a fixed spatial shape function, derived from the initial steady-state solution, by the instantaneous power amplitude $P(t)$. This field is then transferred to the thermal-hydraulics mesh and injected as a volumetric heat source into the fluid energy equation through the dedicated `powerDensityNeutronics` field.

The thermal-hydraulics solver in turn advances the temperature and velocity

fields over the current time step under the influence of this heat source. The resulting temperature distribution is subsequently fed back to the point-kinetics solver as a volume-averaged quantity, computed by weighting the local temperature field with the one-group neutron flux distribution. This flux weighting, performed zone by zone through the `initialOneGroupFluxByZone` mapping, ensures that the feedback temperature is representative of the neutronically active fuel regions alone, preventing the large volume of the reflector and pool from artificially diluting the average and thereby suppressing the computed reactivity feedback. The reactivity contributions from fuel temperature (Doppler effect) and coolant temperature and density are then evaluated from these zone-averaged quantities through user-defined feedback coefficients, and supplied to the power ODE for the following time step.

Chapter 4

Thermal-hydraulic model

In nuclear thermal-hydraulics, core heat removal is a primary design challenge. Conventional large-scale reactors typically rely on forced convection during nominal operation, reserving natural circulation as a critical safety redundancy. In contrast, many modern Small Modular Reactor (SMR) designs integrate buoyancy-driven flow as their primary cooling mechanism even under steady-state conditions, thereby eliminating the need for main coolant pumps. This drastically reduces parasitic electrical loads, simplifies the primary circuit layout, and inherently removes the risk of "loss of flow" accidents. Within this landscape, the BME TR represents a versatile research platform, engineered to operate via both forced and natural circulation [26].

This chapter details the stand-alone thermal-hydraulic modeling of the BME TR under natural circulation; simulations are performed using foamForNuclear [6]. The computational domain is defined via a dedicated `fluidRegion`, which contains the mesh, boundary conditions, physical properties and the governing equations responsible for the Thermal-Hydraulic (TH) simulation. This stage serves as a foundational benchmark against experimental data, establishing the model's accuracy before its application in coupled neutronics-thermal-hydraulic analyses.

4.1 Mesh generation

The BME TR thermal-hydraulic model adopts a hybrid approach: free-flow regions (the pool area) are treated with standard Computational Fluid Dynamics (CFD) techniques, while the reactor core structures are handled via a porous-medium formulation (Section 3.3.2). The latter allows each fuel assembly to be represented as a single mesh block, significantly reducing the computational cost. The computational mesh and its generation procedure are described in Section 3.3.1. The structured mesh was generated through a multi-block approach, creating distinct domains (e.g., the reactor core, the lower plenum, and the upper pool) that were subsequently merged into a single geometry. For that purpose, custom-developed functions were implemented to align nodal coordinates at the interfaces; This aspect was extremely important to achieving a conformal mesh.

Moreover, to satisfy the turbulence model constraint, the near-wall mesh is

deliberately kept coarser than in the bulk: sufficiently large first-layer cells are needed to prevent y^+ from dropping below approximately five, which would place cell centroids inside the viscous sublayer where the logarithmic assumption no longer holds and wall-function predictions become physically inconsistent (discussed in Section 3.3.3).

The mesh design was guided by deliberate geometric simplifications. Not every geometric detail was modeled. Some components were deliberately simplified to streamline the system and focus on its core behavior. This choice was driven by the need to keep the model computationally manageable within the scope of the present study, despite the inherent risk of overlooking secondary local effects that more detailed geometries might have captured. The assumptions are listed above and debated in the result section:

- Reflector blocks removed from the thermal-hydraulic vessel model, being considered zero-flow regions whose temperature is assumed not to change.
- The diffuser area was not explicitly modeled; instead, it was replaced by custom-implemented mapped patches to replicate its fluid-dynamic behavior (see Section 4.2.2 for further details).
- The vertical guiding tubes (used for irradiation positions, neutron sources, or instrumentation), the sample-delivery system, the irradiation tunnel, and the horizontal beam tubes were omitted from the present study.
- The control and safety rods were modeled exclusively within the neutronic phase and were excluded from the thermal-hydraulic porous medium domain. (see Section 3.4.1 for further details).

The resulting global mesh topology, incorporating the previously described domains and interfaces, is illustrated in Figure 4.1. This visual representation highlights the conformal matching between the sub-domains and the implemented geometric simplifications; specifically, the labels indicate the regions where the diffuser and reflector blocks were originally located, now represented as external boundaries or empty zones to optimize the computational domain.

The final whole geometry mesh counts 150'000 cells and the discretization is homogeneous along the core vertical direction with finer refinement in the core window area to correctly capture all the physical phenomena with respect to the closing loop of natural circulation.

It should be highlighted here that, although it might look like triangular mesh in some post-visualization processes using ParaView, the computational mesh is essentially polyhedral.

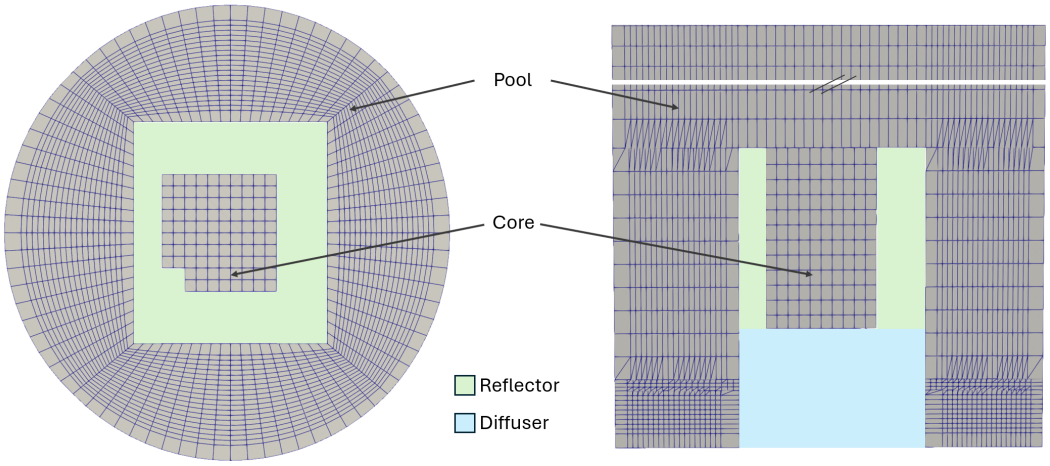

Figure 4.1: Global mesh topology, featuring a vertically truncated side view (right) within the pool region to highlight the core and diffuser discretization while preserving the full pool volume in the computational domain.

4.2 Mesh improvment

4.2.1 ACMI

The multi-block strategy for mesh creation offers several advantages, simplifying the structure of the Python API script and accelerating its generation, particularly regarding the pool geometry blocks and the core assembly separately. However, issues arise when these two structures are merged, as they are created with different spatial discretizations. This leads to non-conformal and partially overlapping meshes at the interface, causing communication problems between the two `patch` entities—which in OpenFOAM define the boundary regions where physical conditions are enforced. Specifically, the non-conformal nature of the interface prevents a consistent numerical mapping of the flow fields across the connected sub-domains, as shown in Figure 4.2a.

To overcome the communication issues at the block interfaces, a dedicated numerical strategy is required. In the OpenFOAM framework, the standard approach for joining non-conformal meshes is the Arbitrary Mesh Interface (AMI) [33, 34]. The AMI allows for a conservative data exchange between disconnected patches by computing weighting factors based on the overlap area between the source and target faces.

Nevertheless, since the pool geometry and the core assembly blocks in this study do not overlap perfectly, a simple AMI approach would be insufficient. Consequently, the focus was shifted to the Arbitrary Coupled Mesh Interface (ACMI) approach. The ACMI is specifically designed for partially overlapping interfaces; it dynamically evaluates the "coupled" portion where the surfaces meet, basing its analysis on a Galilean transformation to correctly map the patches and interpolate the fields [35]. Simultaneously, it effectively "closes" the "non-coupled" part—where the two blocks do not meet—by treating it as a wall boundary condition (see Figure 4.2b).

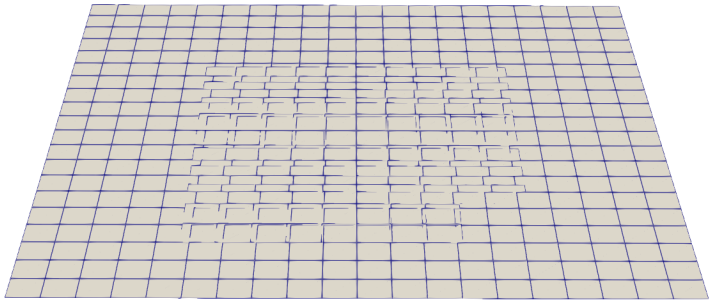

(a) Core-pool interface details showing the non-conformal mesh layout.

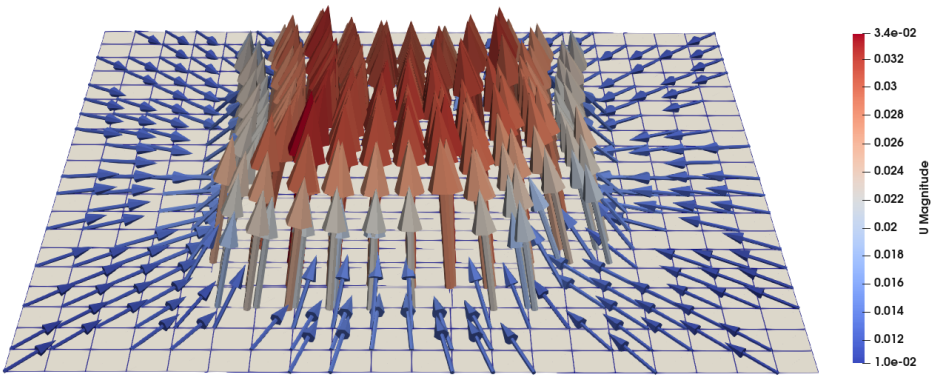

(b) Core-pool interface velocity field obtained through the ACMI boundary approach.

Figure 4.2: Computational mesh and velocity field analysis at the core-pool interface.

4.2.2 Mapped Patch

The design of the mapped patches is based on the complex flow behaviour observed inside the vessel during reactor operation. According to previous studies and experimental data, the water heated by the core rises in a central plume until it reaches the top of the vessel. At this point, the flow splits: one part recirculates downward along the outer walls. As this downward flow reaches the core level, a portion is pulled back into the central upward plume, while the remainder continues toward the bottom, entering the core barrel through specific windows to restart the cycle [4, 5]. In the actual reactor configuration, the water entering these windows would pass through the diffuser area, before being redirected upward in the core. A rigorous discretization of the diffuser would have required a complex porous media formulation to account for its influence on the flow field homogeneization; Such flow dynamics pose non-trivial challenges for the downcomer-to-core transition; In this open-pool configuration, the downcomer is defined as the annular region between the core barrel and the vessel wall.

However, for the present study, to keep the model computationally efficient, the physical structure of the diffuser was removed and replaced by the mapped boundaries. Those act as a numerical bridge, directly transferring the fluid properties (pressure, temperature and velocity) from the core barrel windows to the core inlet. This ensures that the simulation correctly reproduces the downward-to-upward transition of the flow, maintaining the physical reality of the recirculation loop without the

need for a complex and heavy geometric mesh at the bottom of the reactor. In the present study, the core window region is modeled as a uniform area distributed along the entire core barrel perimeter, preserving the exact effective flow area of the six independent windows found in the actual geometry.

The numerical implementation of the mapped patches requires a differentiated approach for the various physical fields involved. Specifically, two distinct mapping strategies were developed: one dedicated to mass conservation and another focused on the communication of the pressure and temperature fields. This distinction is necessary due to the geometric discrepancy between the source and target surfaces; since the effective areas of the core barrel windows and the core inlet do not coincide.

For the velocity field, the primary objective is to satisfy the conservation of mass across the transition. To achieve this, the velocity profile sampled at the core window is not merely copied, but rescaled to account for the area and porosity ratio between the two boundaries. This ensures that the global flow rate is preserved as the fluid enters the core. The global balance between the source (core barrel window) and the target (core inlet) is defined as:

$$A_{src} \cdot \rho_{src} \cdot \alpha_{src} \cdot U_{src} = A_{trg} \cdot \rho_{trg} \cdot \alpha_{trg} \cdot U_{trg} \tag{4.1}$$

where src and trg are referred respectively to source and target boundaries, A represents the patch area, ρ the fluid density, and α the void fraction (or porosity coefficient) insofar as the core region is modeled in porous media the fluid velocity must be adjusted to the volume effectively available for the flow. In the context of the present study, the velocity to be imposed at the core inlet, U_{trg} is derived as follow

$$U_{trg} = U_{src} \cdot \left(\frac{A_{src} \cdot \alpha_{src} \cdot \rho_{src}}{A_{trg} \cdot \alpha_{trg} \cdot \rho_{trg}} \right) \tag{4.2}$$

The area-based correction prevents nonphysical mass sources or sinks within the lower plenum by adjusting the velocity magnitude to account for the flow contraction or expansion arising from the removal of the diffuser geometry. In addition, the velocity vector is reoriented to point perpendicularly to the core inlet face, and the inlet profile is assumed spatially uniform (see Figure 4.3). These two assumptions together model the diffuser as an ideal device, one that both redirects and homogenizes the flow before it enters the core region.

In contrast, a different coupling strategy was adopted for the temperature (T) and pressure (p) fields. Rather than point-to-point profile mapping, both quantities are spatially averaged over the source patch and imposed as uniform boundary conditions on the target face. This approach improves numerical stability at the core inlet by filtering local turbulent fluctuations that could otherwise perturb the solution at the inter-region interface [36]. The directionality of the coupling is tailored to the physical role of each variable. Temperature is mapped from the window patch (source) to the core inlet (target), consistent with the convective transport of thermal

energy in the flow direction. The pressure coupling is instead reversed: a pressure value is communicated from the core inlet back to the window patch, following the standard Dirichlet–Neumann decomposition adopted in partitioned multi-region solvers, whereby a fixed-value condition on one side of the interface is paired with a zero-gradient condition on the other [36, 37]. This reverse feedback is necessary to maintain a consistent pressure–velocity coupling across the geometric gap between the two regions, providing the solver with the driving force information required for mass conservation.

Furthermore, it is important to emphasize that this data exchange is not a pre-computed boundary condition. The software handles the variable transfer internally at each iteration of the solver, ensuring that the mapped patches dynamically adapt to the evolving flow field, allowing the simulation to capture the transient fluctuations and the non-linear interactions characteristics of the reactor's internal recirculation.

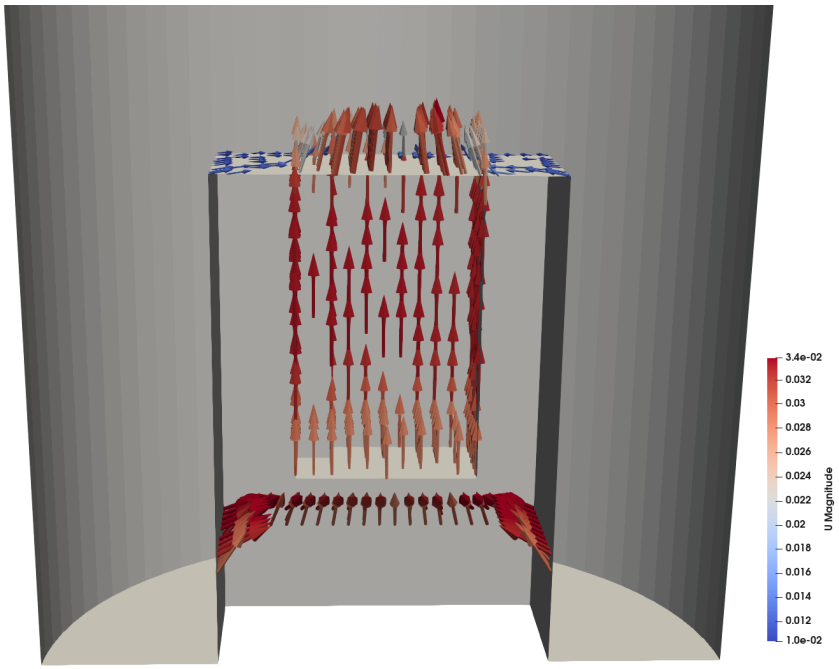

Figure 4.3: Velocity glyph representation illustrating the mappedPatch coupling between the core Window and core Inlet boundaries.

In addition, a pressure drop Δp is explicitly imposed at the interface to account for the hydraulic resistance of the lower plenum structures not resolved by the mesh, namely the diffuser and the perforated tray located immediately below the fuel assemblies. Approximate values for the corresponding loss coefficients were estimated following the correlations reported in [38].

The resulting pressure boundary condition applied at the patch reads

$$p = p_0 + \frac{1}{2} K \rho |\mathbf{U}|^2, \tag{4.3}$$

where p_0 is the reference pressure of the adjacent region, ρ the local fluid density, $\mathbf{U}$ the velocity at the patch, and $K = \sum_i \zeta_i$ the total loss coefficient summed over

the two components. The specific numerical values and the derivation of the loss coefficients components (ζ_i) are detailed in Section 6.1.2.

This dual-mapping strategy-rescaling for velocity and area-averaging for scalar fields-provides a robust compromise between physical accuracy and the numerical stability required for a complex recirculation study.

4.3 Core

The core was defined with a total length of 600 mm. While the active height of the EK-10 fuel rods is 500 mm, the top and bottom nodes of the channels are non-heated and were not modeled in detail (e.g., the skirt part at the assembly bottom). A similar approach was adopted for the spacer grids located at the top and bottom sections of the assemblies; however, the model allows for their future inclusion by implementing a specific evaluated pressure drop (knowing the blockage ratio: 0.52) [5]. The assemblies were assumed to be perfectly square and were constructed by incorporating the inter-assembly water gap into the porous medium approach, resulting in assembly dimensions of 72 mm × 72 mm. Regarding the assembly properties, six distinct fuel assembly configurations were identified: types B3, F3, and D5, the irregular C3 and D4 types, and the regular assemblies (see Figure 4.4) [15]. For each configuration, the flow areas, solid volume fractions, and hydraulic diameters were calculated and were subsequently implemented as specific inputs for their corresponding porous properties definitions in the `constant/fluidRegion/phaseProperties` dictionary.

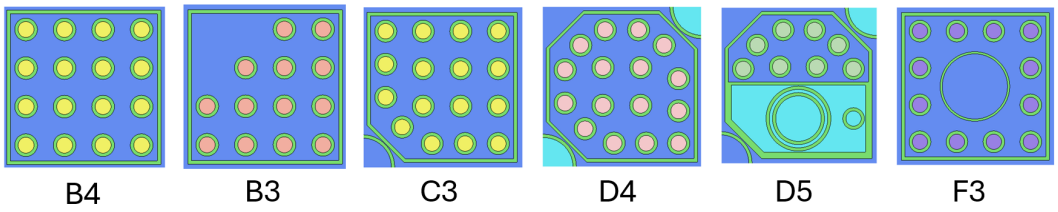

Figure 4.4: Different assembly types.

Assembly type	Solid Vol. Frac.	Flow area [m^2]	D_h. [m]	K-factor [-]
Assembly B4	0.1221	0.0034	0.0135	0.080
Assembly B3	0.0992	0.0037	0.0159	0.121
Assembly C3	0.1221	0.0036	0.0140	0.087
Assembly D4	0.1221	0.0029	0.0121	0.058
Assembly D5	0.0611	0.0019	0.0123	0.061
Assembly F3	0.0916	0.0033	0.0130	0.088

Table 4.1: Hydraulic properties of the different assembly types [5].

To establish the initial core power distribution while awaiting the definitive profile from the neutronics evaluation, a uniform power distribution was assumed across the 369 fuel rods, scaling the nominal 100 kW total power based on the rod

count of each assembly. Thermal behavior within the rods was simulated using the `nuclearFuelPin` power model, described in detail in Section 3.3.2, which accounts for the geometric and material properties documented in literature [4, 5, 15].

4.4 Simulation set-up

The configuration of the `fluidRegion` solver requires the specification of several interconnected components: the thermophysical properties of the coolant, the drag and heat transfer correlations governing momentum and energy exchange within the porous medium, and the boundary conditions applied at the domain boundaries. Each of these is described in turn below.

Phase properties. The thermal-hydraulic behaviour of the `fluidRegion` is governed by zone-wise closure models specified in the `phaseProperties` dictionary; the `heatTransferModels` and `dragModels` sub-dictionaries assign the inter-phase convective exchange and frictional pressure-loss correlations to each region of the porous-medium mesh.

The closure parameters are grounded in the geometry of the EK-10 fuel assemblies: each Type 1 assembly (B4, see Figure 4.4) contains 16 pins of diameter $D = 10\,\mathrm{mm}$ arranged in a 4×4 square array with lattice pitch $P = 17\,\mathrm{mm}$, giving $P/D = 1.70$. The bundle-average hydraulic diameter, computed from the total flow area and wetted perimeter of the $68\,\mathrm{mm} \times 68\,\mathrm{mm}$ inner cross-section and the fuel rods, evaluates to $D_h = 4A_f/P_w \approx 13.5\,\mathrm{mm}$; the heated length coincides with the active pin length, $L = 500\,\mathrm{mm}$.

Under natural-circulation conditions at nominal power (100 kW), the Reynolds number in the fuel channels is estimated to range between approximately 500 and 1500 [26], placing the flow solidly in the laminar regime. The closure models are selected accordingly.

Convective heat transfer between the fuel-pin surface and the coolant is modelled via the Sieder–Tate correlation [39] for thermally developing laminar flow. Assuming constant fluid properties across the boundary layer, the correlation is expressed as:

$$\mathrm{Nu} = 1.86\,\mathrm{Gz}^{1/3}, \qquad \mathrm{Gz} = \mathrm{Re}\,\mathrm{Pr}\,\frac{D_h}{L}, \tag{4.4}$$

valid for $0.48 \leq \mathrm{Pr} \leq 16700$, $\mathrm{Re} \leq 2100$, and $\mathrm{Gz} \geq 8$, a limit established by Whitaker [40];. The formulation naturally captures the entry-length heat transfer enhancement through the Graetz number, which for the expected operating conditions evaluates to $\mathrm{Gz} \approx 50$–500. Consequently, Eq. (4.4) is directly mapped to the `NusseltReynoldsPrandtlPower` dictionary form by setting `const` $= 0$, `coeff` $= 1.86\,(D_h/L)^{1/3} \approx 0.558$, and `expRe` $=$ `expPr` $= 1/3$.

Frictional pressure losses are evaluated based on the distributed loss coefficients, K, specified in the reference document [5] and reported here in Table 4.1. For each unique fuel assembly configuration, the specific friction factor f is derived by scaling

these reference loss values, originally calculated for a mean Reynolds number of approximately 800, with the individual assembly's geometric parameters, namely its hydraulic diameter (D_h) and active length (L). This approach ensures that the porous medium representation accurately reflects the calibrated flow resistance characteristics of each sub-region, bypassing the need for generic analytical friction multipliers.

Thermophysical properties. The thermophysical properties of water are represented through the `icoPolynomial` thermophysical model, which expresses each property as a polynomial function of temperature and is the standard incompressible-fluid option in OpenFOAM for temperature-dependent media. Over the operational temperature range of the BME TR, between 293 K and 340 K, density, dynamic viscosity, and thermal conductivity exhibit a temperature dependence that is accounted for by fitting polynomial functions to NIST Standard Reference Database tabulated values [41], while the specific heat capacity remains practically insensitive to temperature across this interval and is assigned a constant reference value. Density is modeled as a second-order polynomial:

$$\rho(T) = 738.2 + 2.000\,T - 3.797 \times 10^{-3}\,T^2 \quad [\mathrm{kg\,m^{-3}},\ T \text{ in K}]. \tag{4.5}$$

The buoyancy-induced density gradient described by Eq. (4.5)—a monotonic decrease from roughly 998 kg m^{-3} at 293 K to approximately 977 kg m^{-3} at 340 K— constitutes the primary driving force for natural circulation in the reactor pool. Dynamic viscosity decreases markedly with temperature and is represented by a third-order polynomial:

$$\begin{aligned}\mu(T) = {} & 6.404 \times 10^{-2} - 5.233 \times 10^{-4}\,T \\ & + 1.444 \times 10^{-6}\,T^2 - 1.339 \times 10^{-9}\,T^3 \quad [\mathrm{Pa\,s},\ T \text{ in K}].\end{aligned} \tag{4.6}$$

Thermal conductivity increases monotonically over the same range and is captured by a second-order polynomial:

$$\begin{aligned}k(T) = {} & -7.083 \times 10^{-1} + 7.186 \times 10^{-3}\,T \\ & - 9.320 \times 10^{-6}\,T^2 \quad [\mathrm{W\,(m\,K)^{-1}},\ T \text{ in K}].\end{aligned} \tag{4.7}$$

The specific heat capacity is assigned the constant value

$$c_p = 4182.2\ \mathrm{J\,(kg\,K)^{-1}}, \tag{4.8}$$

evaluated at the nominal inlet temperature $T_{\mathrm{ref}} = 296$ K.

Boundary conditions. Global mass conservation across the computational domain is enforced through the mapped patch described in Section 4.2.2, which transfers pressure, temperature and velocity fields between the inlet and outlet faces and prevents the accumulation of spurious mass imbalances. The upper water surface of the reactor pool is treated with a slip boundary condition, imposing zero normal velocity

while leaving the tangential component unconstrained. This choice is consistent with the open-pool design of the BME TR, whose water surface is permanently accessible and is routinely used for the insertion of experimental instrumentation [5].

Heat losses. Although heat losses through the pool boundary are not the dominant energy pathway during normal operation, their omission in transient simulations would introduce a systematic bias [5, 4]. Two loss mechanisms are physically relevant: surface evaporation at the free water interface and conductive losses through the vessel wall toward the surrounding bioshield. Modelling the evaporative contribution in detail is impractical because the boundary conditions governing it—the ventilation airflow above the pool and the geometry of the annular gap between the outer vessel wall and the bioshield—are poorly characterised [5, 4]. Conductive losses through the vessel wall were therefore retained as the sole heat loss pathway and were implemented via the `externalWallHeatFluxTemperature` boundary condition, which enforces a Robin-type condition of the form

$$-k \left.\frac{\partial T}{\partial n}\right|_{\text{wall}} = h \left(T_{\text{wall}} - T_{\infty}\right), \tag{4.9}$$

where k is the wall thermal conductivity, n is the outward wall normal, and T_{∞} is the ambient temperature, the thin air layer that surround the core. A fixed ambient temperature of $T_{\infty} = 20\,^{\circ}\text{C}$ was prescribed. The effective heat transfer coefficient $h = 12$ W/(m^{2}K) was determined by fitting the simulated post-shutdown cooling slope to the experimental temperature decay curve measured after the 100 kW power transient reported in [5].

4.4.1 Selection of the Simulated Transient

Among the available experimental datasets, transient TH#6 was selected for the validation analysis presented later in Chapter 7. This specific test belongs to the category of long-duration natural circulation experiments and spans a total duration of 90 minutes. The scenario was conducted with the primary pump loops completely inactive from the very beginning of the test, forcing the system to rely solely on buoyancy-driven flow. At the onset of the measurement, the entire primary circuit was stagnant and at a highly uniform isothermal state, characterized by a constant temperature of 296.15 K, which provides a well-defined and reproducible initial condition for the numerical model.

The transient was initiated from a 0 W power shutdown state by a power ramp-up that reached the nominal reactor thermal power of 100 kW in approximately 11–12 minutes. As illustrated by the experimental power history in Figure 4.5, this phase is followed by a prolonged constant power plateau before the final reactor shutdown. This time-dependent profile acts as the primary boundary condition driving the simulation. Throughout the 90-minute transient window, the development of natural circulation loops and the subsequent evolution of the temperature fields were recorded by the thermocouple instrumentation across multiple axial levels and radial positions.

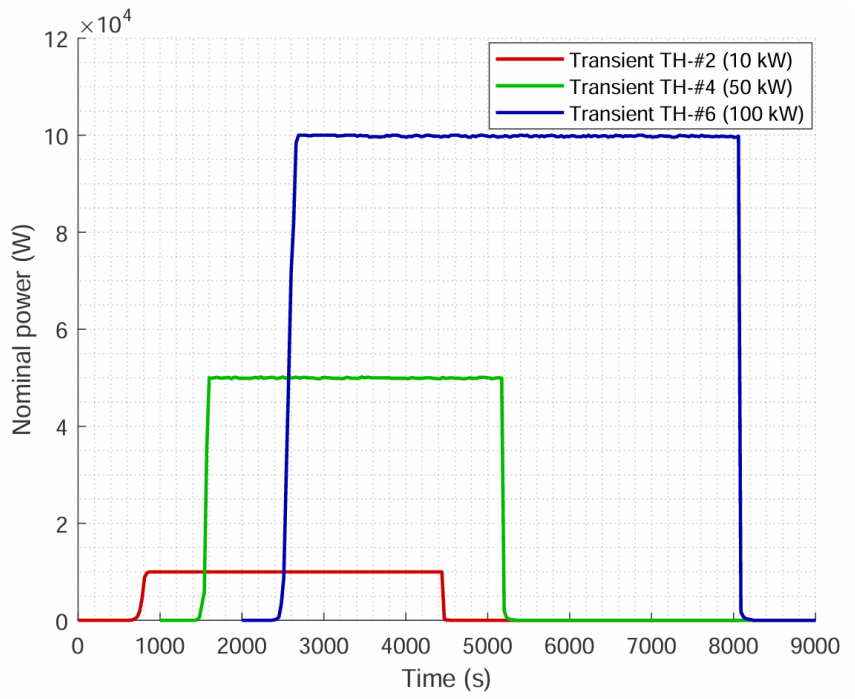

Figure 4.5: Logged power curves for Transients TH-#2, TH-#4 and TH-#6 (shifted to each other for better visibility) [5].

Transient TH#6 was chosen because of the significant temperature differentials that develop across the vessel, the representativeness of the 100 kW power level for evaluating core behavior under natural circulation regimes. This experimental layout is ideal for a direct assessment of the model's capacity to capture axial thermal stratification in both the central and peripheral regions of the vessel.

It should be noted that for the subsequent numerical simulations, this experimental power boundary condition is adjusted by a calibration factor. The underlying physical justifications for this correction are thoroughly discussed in the next chapter dedicated to modeling assumptions (Chapter 6).

Experimental Procedure and Instrumentation Layout To chart the thermal evolution of the system, the pool reactor is fitted with a dense grid of immersion type sensors on two vertical aluminum bars. The test campaign was conducted during a period of two weeks in May 2021, during which six long transients, including TH#6, and eleven short cycles were conducted to assure full coverage of the data.

The placement of such probes follows strict geometric considerations (see Appendix A for details). The central rod is anchored onto two diagonal fuel assembly tips, a configuration that allows for its displacement to different positions across the core to capture local variations. This rod supports one sensor positioned precisely above a sub-channel at the fuel rod level, while five additional thermocouples are distributed at vertical intervals of 1000 mm, 1750 mm, 2500 mm, 3250 mm, and 4000 mm above the core level, as illustrated in Figure A.2. In contrast, the peripheral rod remains stationary at a fixed location approximately 50 mm from

the A4 graphite block throughout the experiments (see Figure A.1). This peripheral assembly includes a sensor placed at the core level, five thermocouples at the same elevations as those on the central rod, and a single Type J thermocouple located 750 mm below the core level, oriented specifically toward the recirculation window to monitor sub-core flow [5].

The entire measurement chain relies on Type K thermocouples routed to a high-density National Instruments NI-9213 data acquisition module. Signal processing and continuous data logging are managed in real time via a dedicated LabVIEW interface, which records the thermal history for subsequent validation against `foamForNuclear` simulation results.

Chapter 5

Neutronic model

Accurately resolving the neutron population within a reactor core is a fundamental prerequisite for any meaningful multiphysics simulation. The spatial and energy distribution of neutrons, their transport through the different material regions, and their eventual absorption, scattering, or leakage directly determine where fission reactions occur and how power is deposited across the core. Without a reliable neutronic model, the thermal-hydraulic and feedback calculations that follow rest on an uncertain foundation.

This chapter documents how the neutronics of the BME TR was validated within the foamForNuclear framework. The analysis began with a full-core `Serpent2` model of the unrodded configuration. From this model, assembly-wise 2-group cross sections for a reference case were extracted and subsequently imported into foamForNuclear via a dedicated `neutroRegion`. Initially, a deterministic diffusion solver was used to compute the core eigenvalue and power distribution. A comprehensive breakdown of the cross-section extraction process is provided in Section 3.4.

The aim is not to reproduce the Monte Carlo solution but to establish that the deterministic model captures the core behavior with sufficient fidelity for the coupled multiphysics simulation. The validation proceeds progressively: from a reference critical configuration, through reactivity feedback coefficients, to control rod worth calculations, each step testing a different aspect of the model's physical consistency.

5.1 Serpent 2 Full-Core Calculation

In the present work, `Serpent 2` fulfils two distinct but complementary roles: it serves as the high-fidelity reference for the full-core neutron flux and global eigenvalue of the BME TR, and it generates the few-group homogenised cross sections that the deterministic `foamForNuclear` solver requires as input.

Both roles rely on the full-core model of the unrodded critical configuration originally developed by A. S. Ványi for the TRACE/PARCS benchmark study [5]. The geometry accurately reproduces the core and vessel layout described in Section 2.2.2, with detailed material compositions defined for fuel, cladding, coolant, reflector, and structural elements. The fuel isotopic composition is determined via `Serpent 2`

depletion calculations, split into two historical operational sequences (1971–1980 and 1981–2020) starting from nominal fresh fuel specifications. This burnup modelling assigns a unique material composition to each fuel assembly based on its operational history, assuming a uniform average composition for all individual fuel rods within the same assembly [5].

Nuclear data are taken from the ENDF/B-VII.1[1] evaluated library [42], which provides a well-validated cross-section dataset across the energy range relevant to thermal research reactors. The simulation was run with 50 000 source neutrons per cycle, 100 inactive cycles, and 3 000 active cycles. The inactive cycles allow the fission source distribution to converge spatially before tallies are accumulated, while the large number of active cycles ensures reliable statistical quality on the extracted cross sections and flux estimates. This setup accumulates 150×10^6 active histories, mitigating source propagation bias and reducing the statistical uncertainty ($\sigma \propto 1/\sqrt{N_{\text{tot}}}$) on global parameters to within a few pcm.

A two-group energy structure is adopted for cross-section condensation. While this is a significant simplification compared to the continuous-energy treatment in `Serpent 2`, it provides adequate accuracy and represents a reasonable compromise between computational cost and physical fidelity for the purpose of this work. Cross sections are generated assembly-wise using the `set gcu` directive, which instructs `Serpent 2` to perform group-constant condensation over user-defined spatial universes during the full-core simulation. Each fuel assembly and reflector region is assigned a dedicated universe, so the resulting cross sections naturally account for the spectral environment of each zone. The transport correction described in Section 3.4 is applied to light water and graphite through the `set trc` input option, using the established correction curves available in the `Serpent 2` data repository [31]. The extracted parameters—including the diffusion coefficients, macroscopic absorption and fission cross sections, neutron yield, fission spectrum, and scattering matrix for each group and each spatial zone—already incorporate this correction in the output and are then formatted as input for the `foamForNuclear` solver [31].

The P_1 leakage model is employed for energy condensation, nad produces the transport-corrected diffusion coefficients described in Section 3.4. The per-assembly homogenisation performed through `set gcu` ensures that the condensation spectrum reflects the local spectral environment of each zone. Assemblies adjacent to the graphite reflector experience a harder epithermal spectrum relative to those at the core centre, and the assembly-wise condensation captures this variability in the resulting cross-section sets—an effect that a single-universe homogenisation over the full core would wash out.

For cross-section interpolation as a function of thermal state variables, the square-root transformation is applied to the fuel temperature, consistently with the $\sqrt{T}$ dependence expected from Doppler broadening in a thermal spectrum. The cross sections of all other materials are interpolated in a linear fashion.

[1] Evaluated Nuclear Data File (ENDF/B-VII.1) library.

5.2 GeN-Foam Neutronics Domain

Within the `foamForNuclear` framework, each computational cell of the neutronics mesh is assigned to a specific OpenFOAM `cellZone`, which identifies the fuel assembly region to which the cell belongs. In parallel, the `Serpent 2` calculation generates homogenised two-group cross sections for each fuel assembly universe through the `set gcu` card, as described in Section 3.4.1. An explicit addressing table establishes a one-to-one correspondence between the `Serpent 2` universe identifiers and the OpenFOAM `cellZone` labels. During the initialisation phase, `foamForNuclear` retrieves the appropriate cross-section dataset and assigns it to all cells belonging to the corresponding zone. This procedure preserves the assembly-wise spatial distribution of the material layout throughout the finite-volume mesh. While a single cross-section dataset per assembly type is sufficient to describe the reference configuration, the localised perturbations introduced by control rod insertions require a separate numerical treatment, detailed in Section 5.3.

The spatial discretisation of the neutronics domain follows the 73-region approach proposed by [5]. The active core is modelled as an 8×9 Cartesian lattice with a 72 mm pitch, consisting of 72 separately homogenised assemblies (fuel and graphite), each measuring $72 \times 72 \times 500$ mm^3. The remaining geometry is lumped into a 73rd homogenised material region, referred to as water. A mesh-independence study established an intra-nodal resolution of 4 radial and 100 axial divisions per node over the 500 mm core height. The water reflector is discretised using 7×7 radial nodes laterally, and a 500 mm block with 30 axial subdivisions placed both above and below the active core [5, 16].

To close the system, appropriate boundary conditions are imposed on the external boundaries of the computational mesh. Zero-flux (black) boundary conditions, $\phi_g = 0$, are applied to all outer surfaces. This represents a standard approximation in neutron diffusion solvers and is well justified in the present configuration, as the water reflector provides sufficient neutron attenuation before the outer boundary is reached, making the probability of neutrons re-entering the core domain negligibly small.

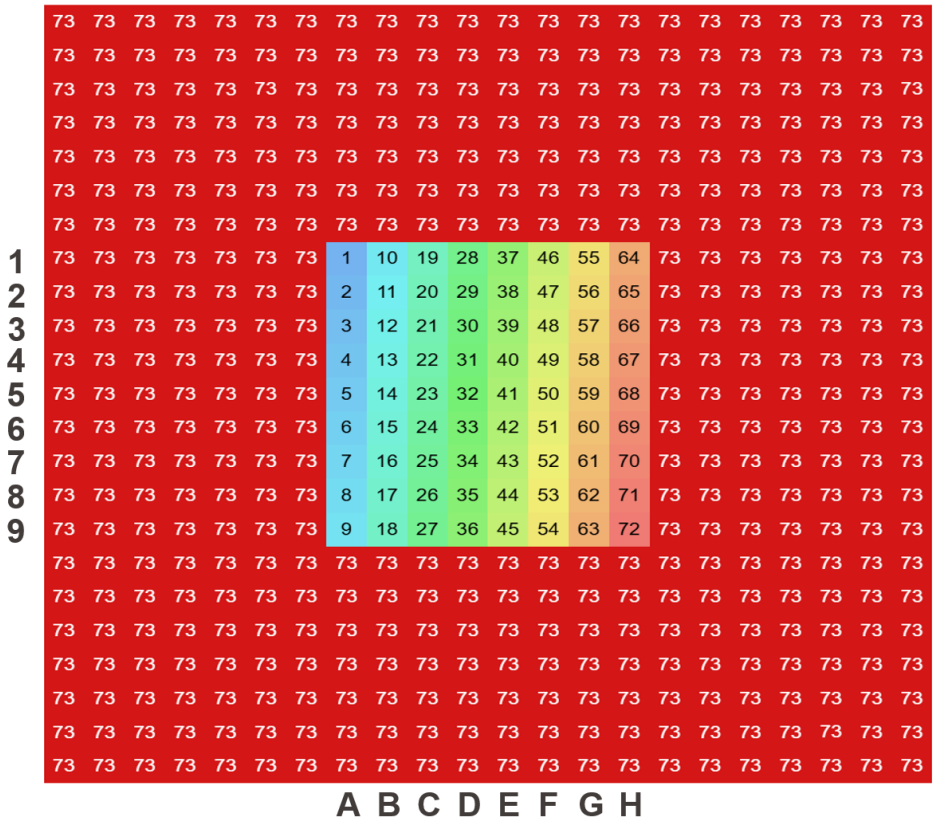

Figure 5.1: Radial cross-section of the full-core `GeN-Foam` model at core mid-height, generated using the `Honeycomb` mesher feature.

5.3 Control Rod Treatment

In the BME TR model, a porous-medium spatial discretisation is adopted for the control rod regions. Rather than explicitly resolving the physical channel located between the four adjacent fuel assemblies, the control rod and its follower are geometrically smeared into the computational mesh of those four assemblies. Consequently, the follower cross-section set reflects a volume-averaged mixture of the surrounding fuel and moderator, rather than pure moderator.

The spatial assignment of the rodded and unrodded cross-section sets at each instant is governed by the `CRmove` utility, whose formulation is described in detail in Section 3.4.1. Control rod worth is evaluated through a series of static k-eigenvalue calculations at discrete rod positions in increments of 50 mm. The axial neutronics mesh was constructed so that cell boundaries coincide exactly with each insertion position, ensuring that z_{tip} always falls on a cell face rather than inside a cell. Under this condition, the step-function assignment of Eq. (3.29) is equivalent to an exact geometric treatment: every cell receives a fully consistent cross-section set and the positional error is identically zero. This mesh-conformal strategy eliminates control rod cusping for the static integral worth calculations performed in Section 7.2.5.

While the approach is entirely adequate for the static neutronic validation, it becomes a limiting approximation in transient simulations where the rod moves

continuously and its tip does not, in general, coincide with a cell boundary. For the axial resolution adopted here—cell heights of approximately 5 mm—the resulting positional uncertainty is at most ±2.5 mm, which is not expected to introduce a significant bias in the integral quantities of interest.

5.3.1 Selection of the Simulated Transient

Among the available experimental datasets, transient FB#2 was selected to validate the coupled neutronics/thermal-hydraulics solver presented in this thesis. This specific test is classified as a *cold insertion* (CI) experiment and represents a fundamental benchmark for assessing the reactor's inherent safety and feedback mechanisms. The scenario is characterized by a deliberate, positive reactivity step introduced from a near-critical initial state, which triggers a power excursion until a new, self-limited equilibrium is established. This experimental setup provides a clear, transient-driven evolution of both the neutron population and the thermal fields, making it an ideal candidate for verifying the model's predictive capability in coupled regimes.The transient was initiated from a low power plateau of approximately 5 W, a condition maintained for a sufficient duration to ensure that the delayed-neutron precursor concentrations reached their equilibrium values. As summarized in Table 5.1, the power excursion evolves toward a final quasi-critical state of about 46 kW. This time-dependent profile serves as the primary benchmark against which the numerical predictions are compared. Throughout the approximately 1550-second duration, the evolution of the outlet temperatures and the neutronic response were simultaneously recorded, providing a comprehensive data set for validating the coupled thermal-feedback dynamics.

Quantity	**Initial state**	**Final state**
z_{AUT}	558 mm	558 mm
z_{MAN}	380 mm	420 mm
MAN rod displacement	$\Delta z = 40$ mm	
Reactivity insertion	≈ 87 pcm	
Monitored assemblies	E3, B3, C7, F7	

Table 5.1: Operating conditions for transient FB#2 [43].

It should be noted that for the subsequent numerical simulations, the experimental power boundary condition is adjusted by a calibration factor of 1.44. The physical justifications for this scaling, which accounts for the difference between nominal and actual thermal power, are thoroughly discussed in the chapter dedicated to modeling assumptions (Chapter 6).

External Reactivity Insertion In the foamForNuclear solver, control rod motion is treated for the point kinetics equations via the externalReactivityInsertion entry in the neutronicsData file. This input defines a piecewise-linear ramp: the reactivity remains at zero until the operator initiates the withdrawal, followed by a linear rise

over the finite displacement time and a constant plateau thereafter. The simulation relies on the six-group $\{\beta_i, \lambda_i\}$ parameters derived from the full-core Serpent 2 model, the prompt neutron generation time Λ, and the temperature coefficients α_{MTC} and α_{DTC} described in Chapter 7.2. At each time step, the solver evaluates ρ_{ext} and superimposes it onto the feedback contributions, allowing the power to track the physical transient response accurately.

For the FB#2 transient, the reactivity insertion is driven by the withdrawal of the MAN rod by $\Delta z = 40\,\mathrm{mm}$, moving from 380 mm to 420 mm. The magnitude of the corresponding reactivity step, $\Delta\rho_{\mathrm{ext}}$, which serves as the plateau amplitude for the input function, is determined by the numerical integration of the measured differential worth curve of the MAN rod across this axial range (Figure 3-6 in [43]):

$$\Delta\rho_{\mathrm{ext}} = \int_{380\mathrm{mm}}^{420\mathrm{mm}} \frac{\mathrm{d}\rho_{\mathrm{MAN}}}{\mathrm{d}z}\mathrm{d}z \approx 145\mathrm{pcm}. \tag{5.1}$$

By utilizing this calculated value, the simulation ensures that the point kinetics model precisely reproduces the experimental rod worth during the transient.

Experimental Procedure and Instrumentation Layout To capture the thermal evolution, the reactor was equipped with a specialized grid of four thermocouple holders. Each holder, consisting of a slender aluminum rod ending in a cross-shaped plate, was fixed onto two diagonal fuel assembly tips. This configuration ensured that the stainless-steel sensors were precisely positioned above the corner sub-channels of the target assemblies (E3, B3, C7, and F7). The needle lengths were calibrated such that the hot junctions were aligned with the top face of the active fuel pins, effectively monitoring the assembly coolant outlet temperature.

The sixteen thermocouple signals were sampled at 75,Hz via a National Instruments NI-9213 module. In post-processing, the data were decimated to 1 Hz through block-averaging, and a 30-second moving-average filter was applied to smooth out high-frequency fluctuations associated with subchannel-scale turbulence, which are not resolved by the simulation. Absolute thermocouple accuracy was conservatively bounded at $\pm 0.5\,^{\circ}\mathrm{C}$; relative differences between sensors of the same holder are reliable to within $\pm 0.2\,^{\circ}\mathrm{C}$ [5]. The neutronic signal was provided by the E5 boron-lined ionization chamber, sampled at 10Hz and converted to power using a proportionality coefficient derived from critical-state observations. Before each cold-insertion run the coolant temperature distribution was homogenised to within $0.3\,^{\circ}\mathrm{C}$ across all sensors by circulating the primary coolant long enough for the delayed-neutron precursor population to reach equilibrium. This conditioning ensured that the initial thermal boundary condition was well defined and reproducible across the series, a prerequisite for meaningful comparison with simulation.

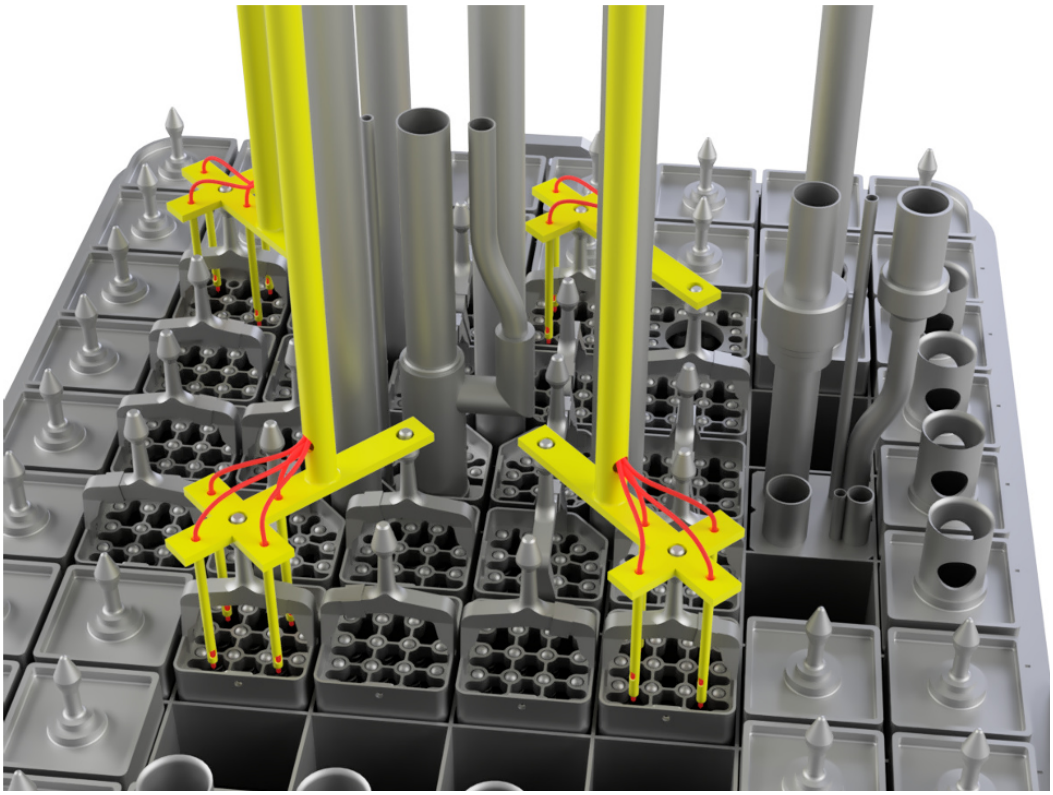

Figure 5.2: Experimental setup of the thermocouple holders (CAD model) [5]).

Chapter 6

Assumptions

6.1 Modelling Assumptions

The thermal-hydraulic and neutronic models described in the preceding sections rest on several simplifying assumptions and deliberate modelling choices. This section collects and justifies those choices in one place, before the simulation results are presented. Some of these assumptions are discussed in the result sections.

6.1.1 Power Calibration

A prerequisite for a meaningful comparison between simulation and experiment is the correct identification of the absolute power level. The power displayed by the BME TR instrumentation is a nominal quantity based on the original detector calibration carried out in 1970. The need for a revised calibration became apparent when steady-state `Serpent 2` simulations returned neutron flux values significantly higher than those inferred from the displayed power.

Within the EVEREST project, a new calibration was established using two independent methods [5]:

- **Activation-based method.** Foil activation measurements were performed and the evaluated flux values compared against `Serpent 2` Monte Carlo calculations. A least-squares linear fit with a fixed zero offset yielded a measured-to-calculated flux ratio of 1.44 ± 0.03, where the uncertainty was determined by Monte Carlo error propagation.
- **Calorimetric method.** Three independent calorimetric measurements gave a consistent ratio of 1.43 ± 0.08.

The excellent agreement between the two approaches establishes the following relationship between the physical reactor power and the displayed nominal reading:

$$P_{\text{reactor}} = 1.44\, P_{\text{nominal}}, \tag{6.1}$$

with an uncertainty of 0.03 on the coefficient [5]. All simulations in this study use this corrected power baseline.

6.1.2 Lower Plenum Hydraulic Resistance

The computational mesh does not resolve the lower plenum structures located below the fuel assemblies, specifically the internal diffuser and the perforated tray, as explained in the thermal-hydraulic modeling in Section 4.2.2 Their hydraulic resistance is instead accounted for by an explicitly imposed pressure drop at the corresponding mesh interface. The detailed equations are derived in Equation 4.3. Analytical estimates for the loss coefficients of these components were derived from the standard formulations in [38]. Specifically, the geometry of the perforated tray (characterized by sharp-edged ∅35 mm holes) was evaluated using the empirical correlations for thin perforated plates. The internal diffuser, which features inclined plates that expand the flow towards the rectangular core cross-section, was modeled according to the formulation for pyramidal diffusers and cross-checked with the conservative Borda-Carnot sudden expansion limit due to the large opening angles visible in the benchmark documentation [15].

The total loss coefficient is defined as the sum of these contributions, $K = \sum_i \zeta_i$, rescaled to the reference area of the core-holder inlet windows. The resulting pressure boundary condition applied at the patch reads.

$$p = p_0 + \tfrac{1}{2}\, K\, \rho\, |\mathbf{U}|^2, \tag{6.2}$$

where p_0 is the reference pressure of the adjacent region, ρ the local fluid density, and $\mathbf{U}$ the velocity at the patch.

In the current stage of this work, a preliminary value of $K = 1.5$ is adopted, accounting exclusively for the isolated contribution of the internal diffuser geometry. The perforated tray has been intentionally omitted at this baseline level to prevent its dominant pressure drop from artificially flattening the velocity profile, thereby allowing a clearer assessment of how the diffuser's geometry alone affects the flow distribution entering the core. As this research is part of an ongoing study, this localized resistance is treated as a calibration parameter. Its specific sensitivity on the predicted core flow distribution is assessed in Section 7.1. Future iterations of the model will encompass a higher, combined loss coefficient ($K \approx 4 \div 5$) to fully incorporate the resistance of the perforated tray, which represents the primary hydraulic restriction of the lower plenum.

6.1.3 Wall Boundary Condition

Experimental data unambiguously show that the coolant temperature at the start of every FB transient is essentially uniform throughout the vessel. The initial reading of all monitored thermocouples is approximately $T_0 = 296.15\,\mathrm{K}$ (23 °C), within the 0.3 °C tolerance required before each experiment was initiated [5]. In addition, it seems that there is no circulation and no sustained flow.

In an earlier configuration of the model, the vessel wall was assigned a Robin boundary condition with a reference temperature of 293.15 K. At the low power levels

(ie. 5 W) characteristic of the initial natural-circulation state, this 3 K deficit was sufficient to drive an artificial, thermally induced downflow at the vessel periphery, distorting the temperature field from the very beginning and preventing the model from reaching a physically consistent steady-state, later used to initiate the transient.

Since the pool is effectively decoupled from the ambient environment, the large thermal mass of the surrounding water ensures that the wall temperature remains close to the initial pool temperature throughout the transient. The Robin condition is therefore replaced by a Dirichlet boundary condition fixed at the measured initial temperature:

$$T_{\text{wall}} = T_0 = 296.15\,\text{K}. \tag{6.3}$$

To quantify the effect of the change in boundary conditions, a reference run with the original Robin condition at 293 K was performed to quantify the impact of this choice. At the corrected initial power of 7.2 W, the 3 K deficit generates a downward peripheral flow that persists into the early portion of the transient and introduces a spurious temperature decrease in the pool. Once the Dirichlet condition at 296.15 K is adopted, the artificial flow disappears, and the early-time temperature evolution is reproduced with a fidelity consistent with the measurement uncertainty. This sensitivity confirms that the wall boundary condition is not a minor numerical detail but a physically relevant modelling decision whose justification rests directly on the experimental initial-condition data.

6.1.4 Exclusion of Outer Core Structures

The thermal-hydraulic domain in this first phase of the analysis explicitly excludes the outer graphite reflector and the surrounding water region. This simplification focuses the computational resources entirely on the active core and the main primary pool volume, where the dominant thermal-hydraulic phenomena occur. From a physical standpoint, this assumption is justified as a first approximation due to the high thermal inertia of the graphite block and the relatively low fraction of total thermal power deposited in these peripheral zones via neutron and gamma radiation. Consequently, their direct, short-term impact on the core flow distribution and assembly cooling is secondary.

However, omitting these peripheral zones introduces specific approximations that must be accounted for when interpreting the results. Most notably, it establishes an adiabatic or simplified thermal boundary condition at the outer edge of the core holder, neglecting the radial heat dissipation toward the stagnant water pool. In addition, the energy stored by the outer reflector and the surrounding water mass during transients is not captured, which may influence the long-term global thermal balance of the system. The implications of these local boundary conditions on the final thermal-hydraulic fields are discussed contextually alongside the results in Chapter 7.

6.1.5 Fuel Thermal Property Calibration

Following the approach proposed in [5], a calibrated modification was applied to the specific heat capacity (c_p) of the uranium dioxide fuel. Specifically, the nominal value of 235 J/(kg · K) was reduced to 23.5 J/(kg · K), to adjust the simulated transient response. Indeed, as detailed in the reference dissertation [5], initial coupled transient simulations overpredicted the experimental peak power by a factor of four to six due to an underestimation of the feedback. The primary cause was identified in the fuel-to-coolant heat transfer coefficient calculations. Standard empirical correlations, optimized for conventional light water reactors, fall short in this specific core configuration due to the non-standard pitch-to-diameter ratio (1.7), low sub-channel velocities, and thick boundary layers.

To mathematically compensate for these geometric and fluid-dynamic deviations without altering the core structure, the specific heat of the fuel material was scaled down by a factor of ten in the reference study [5] to reproduce the experimental power curves of the transient experiments. In this work, the same calibrated property is adopted a priori to ensure baseline consistency with the reference model's transient behavior.

Chapter 7

Results

7.1 Thermal-Hydraulic simulation results

This Section evaluates the computational findings obtained from the `foamForNuclear` simulations of the BME TR system. The analysis first establishes the global, time-invariant fluid dynamics, mapping the persistent natural circulation loops and velocity fields that characterize the reactor pool. Following this hydrodynamic baseline, the focus shifts to thermal validation through two distinct transient benchmarks. These investigations are structured by pairing specific transient case studies directly with experimental reactor measurements, evaluating both localized thermal stratification and multi-assembly temperature distributions.

7.1.1 Thermal-hydraulic flow description

Natural circulation constitutes the sole heat removal mechanism of the BME TR under normal operating conditions, for what concerns the present study. The absence of mechanical pumping means that the flow field is entirely sustained by the buoyancy forces arising from the density difference between the heated coolant in the core region and the cooler water in the surrounding pool.

Fission heat deposited in the fuel assemblies reduces the local coolant density, driving an upward flow through the core channels. Upon exiting the top of the active region, the heated stream continues to rise through the open pool as a buoyant central plume, carrying the hottest fluid toward the vessel ceiling. There, constrained by the upper boundary, the flow is deflected radially outward and descends along the outer perimeter of the vessel as a broad, slow return current. At the elevation of the core barrel, the downward flow bifurcates: part of it reverses and joins the upward core efflux, closing the primary recirculation loop, while the remainder is directed inward through the openings of the core barrel, core windows, and re-enters the core from below, sustaining the upward throughflow.

The velocity field obtained from the `foamForNuclear` simulation is consistent with this topology. The axial velocity component exhibits a pronounced maximum along the vessel centreline immediately above the core exit, decaying monotonically toward the pool walls, as illustrated in Figure 7.1. The radial component shows

outward spreading near the vessel ceiling and inward convergence in the vicinity of the core barrel, reflecting the return-flow geometry described above.

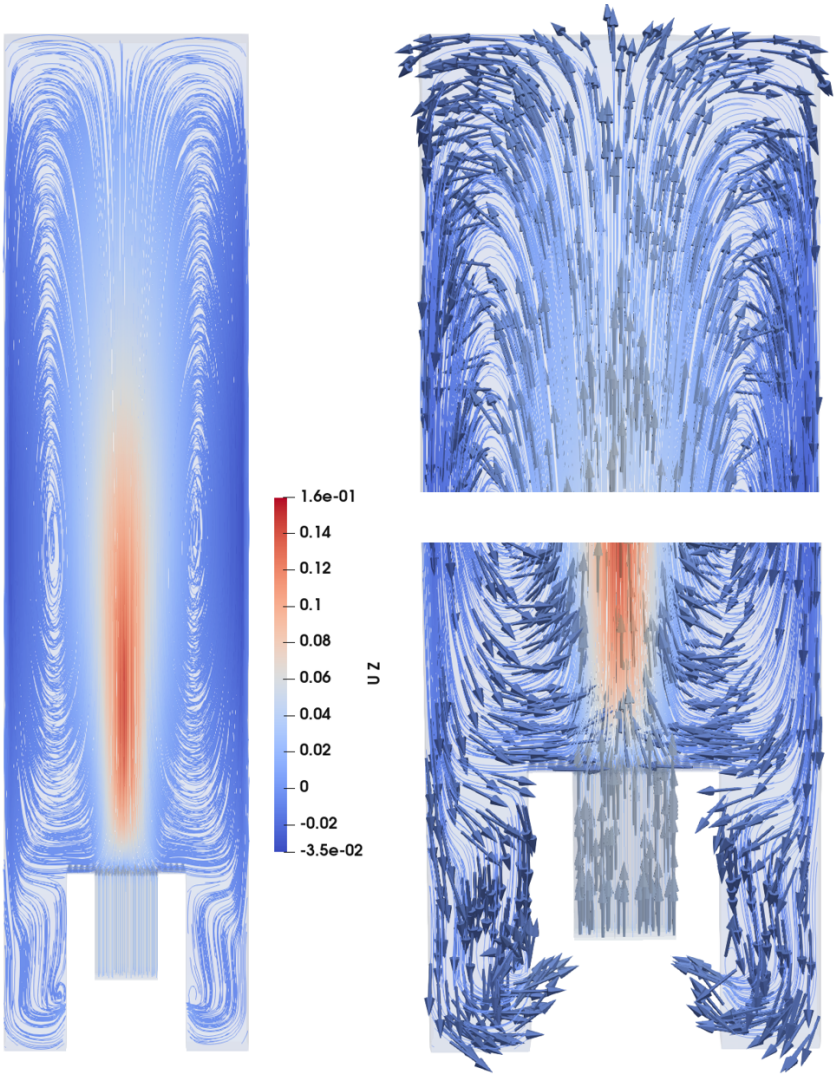

Figure 7.1: Visualization of natural circulation within the reactor vessel, with insets detailing the flow at the upper and lower ends.

7.1.2 Reference Transient: TH#6

The first transient scenario examines the thermal response of the pool under a prescribed power ramp, the selection of which was motivated in Section 4.4.1. The inpur power amp is extracted from the measured reactor power during the experiment corresponding to the transient TH#6. Temperature data were sampled at two representative locations, one centrally positioned above the core exit and one within the peripheral annulus, in order to follow the spatial development of the thermal field and to compare it against the experimental thermocouple readings.

The simulation reproduces a phenomenon of primary interest for this configuration: the onset of axial thermal stratification. The buoyant plume rising from the core carries the hottest fluid upward through the geometric centre of the vessel, and because

entrainment of the surrounding cooler water remains incomplete at intermediate elevations, a pronounced vertical temperature gradient establishes itself along the central column. This gradient relaxes only near the upper section of the pool, where the descending recirculation and the confinement of the vessel top promote more effective mixing. The peripheral region behaves differently. There the radial temperature distribution stays comparatively uniform at any given height, and the dominant gradient is again vertical, set by the slow downward progression of the warm recirculating front rather than by direct plume transport.

The contrast between a stratified central column and a more homogeneous outer annulus is intrinsic to the buoyancy-driven natural-circulation regime, and its reproduction is one of the qualitative benchmarks the thermal-hydraulic model is expected to satisfy. On this point the present results are encouraging: the sampled profiles confirm that the thermocouple positions adopted in the model correspond to physically meaningful regions of the flow, and they recover the expected stratification pattern within the vessel, as shown in Figure 7.2.

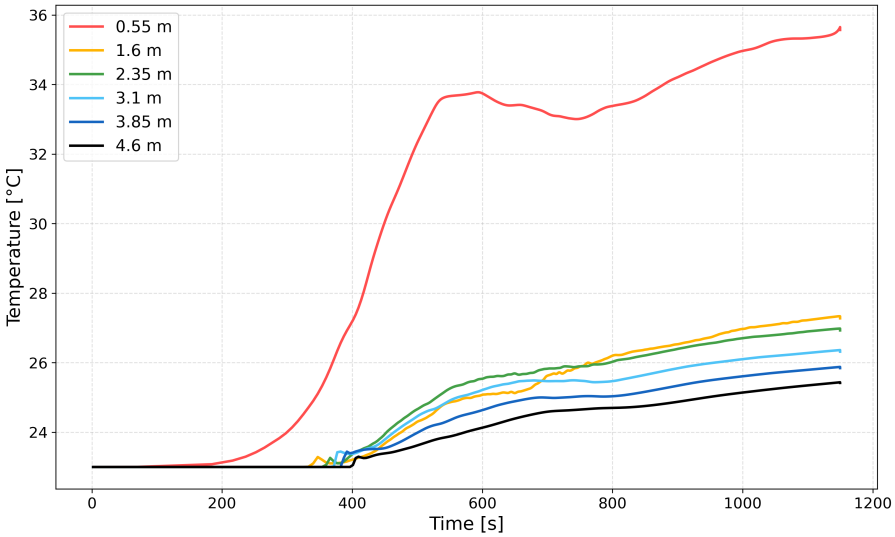

(a) Thermocouple in Central position.

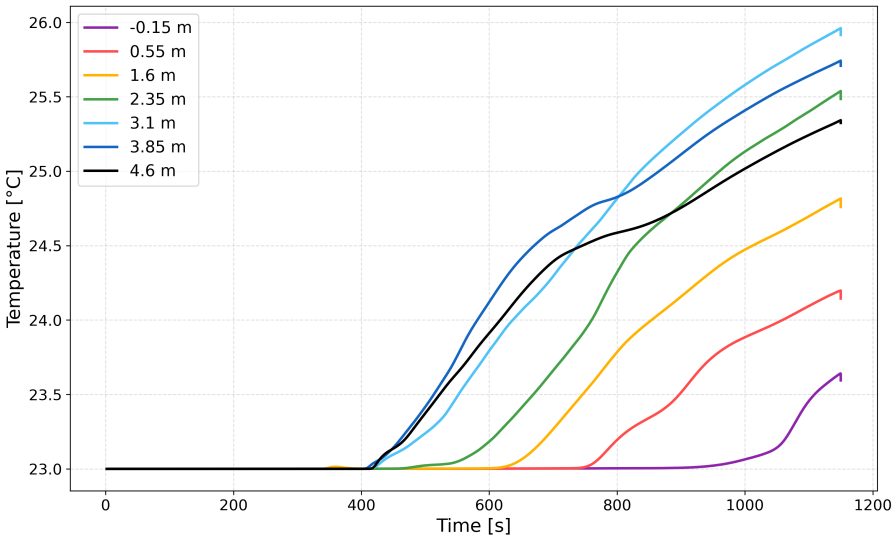

(b) Thermocouple in Peripheral position.

Figure 7.2: Temperature profiles in different thermocouple positions.

These results should nevertheless be read as preliminary. The calculation does not reach the end of the intended transient window: at an intermediate stage of the run the solution destabilises, with a rapid and unphysical growth of the velocity field that in turn drives the temperature and p_{rgh} fields to diverge, ultimately causing the run to crash. This behaviour is believed to stem from a loss of numerical convergence rather than from a genuine physical feature, most plausibly connected to the treatment of the heat-transfer coefficients and to the stability of the adopted discretisation schemes under the low velocities characteristic of natural circulation. Because the run terminates before the end of the intended time window, the profiles reported here should be understood as a demonstration of correct sensor placement and of the qualitative capture of stratification, and not yet as a quantitatively reliable description of the transient over its full duration.

Removing this calculation scheme limitation is part of the work still in progress.

Subsequent effort will focus on a more careful characterisation of the heat-transfer closures and on refining the numerical schemes, with the aim of suppressing the spurious velocity growth and carrying the simulation through to the end of the transient. Only at that stage will a full quantitative comparison against the measured temperature histories become meaningful.

7.1.3 Multi-Assembly Verification Transient Case

The second transient case provides an independent verification of the thermal-hydraulic model response to a power input, before any neutronic feedback is introduced. Rather than prescribing an idealised power shape, the calculation is driven by a time-dependent power history taken from reactor operation data, corresponding to the same core evolution later reproduced in the coupled analysis of Section 7.3.1. In that coupled configuration, the power is not imposed but develops through the point-kinetics response to the reactivity insertion caused by the extraction of a control rod; here it is forced instead, so that the fluid-dynamics solver can be exercised in isolation and assessed without being driven by an additional neutronic solver.

The total power is reconstructed from the recorded reactor signal and scaled by the calibration factor of 1.44 derived in Section 6.1.1. This fixes the energy released as a function of time but not its spatial distribution, which the verification specifically targets: assigning the same power to every assembly would flatten out the very differences the test should resolve. The per-assembly power was therefore extracted from the Serpent steady-state full-core model, normalised to unit sum, and used as a fixed set of radial weights. At each instant, the forced total driven by the total power ramp is partitioned among the assemblies according to this Monte Carlo power map, so that the evolution follows the measured history while the radial split stays close to what the core actually experiences. The resulting total power history is shown in Figure 7.3.

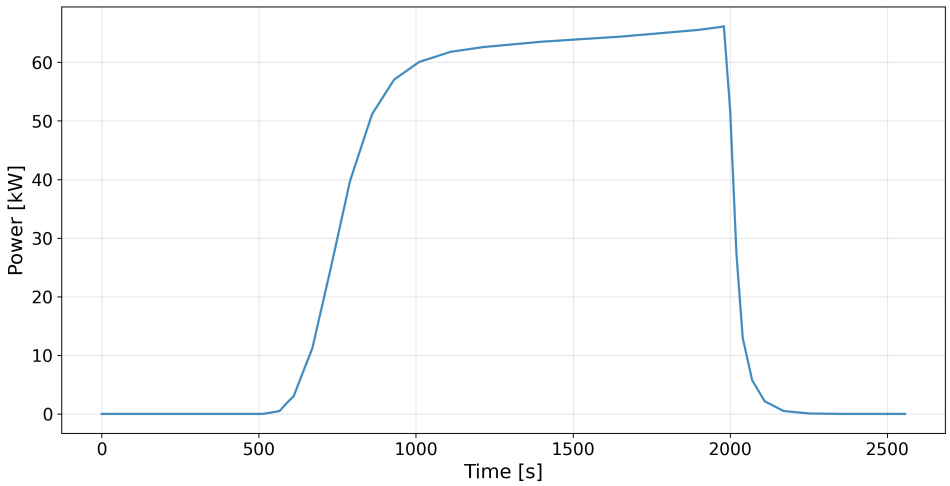

Figure 7.3: Forced total power history imposed on the thermal-hydraulic model, obtained from the measured reactor power scaled by the calibration factor of 1.44.

Based on this power input, the FFN model then computes the flow and temperature evolution. From these results, the four fuel assemblies (E3, B3, C7, and F7) temperature evolutions can be sampled at the top of the active fuel length where the coolant is hottest. The comparison is shown in Figure 7.4.

A minor offset is already visible at the start, where the measured curves sit slightly above the simulated ones. The model was initialised at 296.15 K, the value reported in the literature for this analysis, whereas the recorded signals begin a few tenths of a degree higher. Since this gap falls within the $\pm 0.5\,^{\circ}\mathrm{C}$ thermocouple uncertainty, it need not indicate any physical inconsistency, though it does propagate as a small constant bias, both curves evolving from their respective baselines.

Beyond this, the simulation follows the measured dynamics closely: the onset of heating, the thermal inertia of the pool and the relaxation after the power drop are all captured with the correct timing, and the assembly-to-assembly spread is of the right magnitude. The plateau temperatures are underpredicted by roughly 1.5–2 °C; since the curves start from nearly the same level and separate as heat accumulates, this points to the coolant temperature rise across the core rather than to the transient behaviour, which is reproduced correctly. A full explanation is beyond the scope of this verification, but two mechanisms are worth examining in future work. The first is the heat transfer correlation, whose validity at the low Reynolds numbers of the natural-circulation regime is not fully established and would benefit from a dedicated sensitivity study. The second is the flow resistance: the pressure drop at the core inlet from the diffuser and tray is likely underestimated, and the spacer grids, a further localised loss, are not represented in the porous model. In a buoyancy-driven loop, the mass flow rate is not prescribed but settles from the balance between thermal head and loop resistance, so a higher drag would lower the circulation rate and, at fixed imposed power, amplify the coolant temperature rise. These remain hypotheses to be tested, and in any case they leave unaffected the relative distribution the verification is meant to probe.

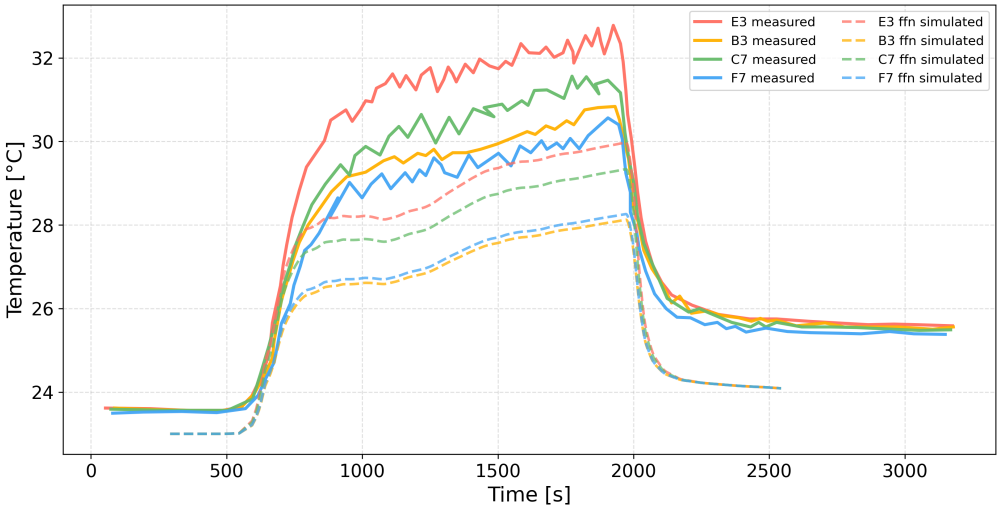

Figure 7.4: Comparison of measured and simulated temperature profiles for four fuel assemblies (E3, B3, C7, F7) under the forced power history. Data are sampled at the top of the active fuel length.

The assembly-wise temperature ranking deserves a closer look. As in the reference simulations, E3 comes out hottest, followed by C7, and the measurements agree on this pair. The other two are swapped: the model puts F7 above B3, while the measurements show the reverse. This same inversion is present in the reference work [5], so once the power is generated self-consistently in the coupled stage, a similar ordering is to be expected.

For the standalone case, the imposed power can be ruled out as the cause. The weights come from the Serpent distribution, where B3 already sits about 2% above F7, so if anything the power should push B3 higher, not lower. The reason has to be sought in the geometry. Inside the model, the only thing that separates the two assemblies is the rod count: B3 has three fewer fuel rods, thirteen instead of sixteen, and correspondingly a wider coolant flow area. In a natural-circulation channel, that wider passage means less resistance and a higher mass flow, and with the power held fixed, the same heat spread over more coolant gives a smaller temperature rise. That is apparently enough to offset the small 2% power margin and bring the B3 reading below F7.

Moreover, it is worth stressing that the real assemblies are not this symmetric. B3 is surrounded by reflector, and F7 borders a block of cooler water (see Figure 2.3), but neither region is explicitly represented in the current mesh (justified in Section 6.1.4). Both, if kept, would tend to raise B3 with respect to F7, the reflector because it limits how much heat B3 loses sideways, the water block because it drains heat from F7, so leaving them out is the most likely reason the model leans the wrong way. Putting these regions back to check whether they recover the measured order is something we leave to future work. Additional considerations such as radial power redistribution during the transient will be assessed with a spatial kinetic solver.

The standalone calculation therefore reproduces the correct thermal inertia and the expected radial temperature distribution, with the only residual ordering issue

traceable to the neutronic power input rather than to the thermal-hydraulic solver. The stability of the fluid-dynamics code under a realistic, time-varying load is confirmed, providing a solid basis for the coupled stage, in which the same transient is reproduced through the point-kinetics coupling.

The thermal inertia observed above has a direct counterpart in the flow field, which is itself an outcome of the calculation rather than a prescribed input. Figure 7.5 shows the vertical velocity component in a core cross-section at successive instants. Early on the coolant is essentially quiescent; as power is deposited, a buoyant plume rises from the core and progressively organises into a coherent upward jet, reaching a fully developed configuration only after several hundred seconds. This is the same timescale over which the assembly temperatures climb towards their plateau, confirming that the delay is set by the establishment of natural circulation rather than by any numerical relaxation.

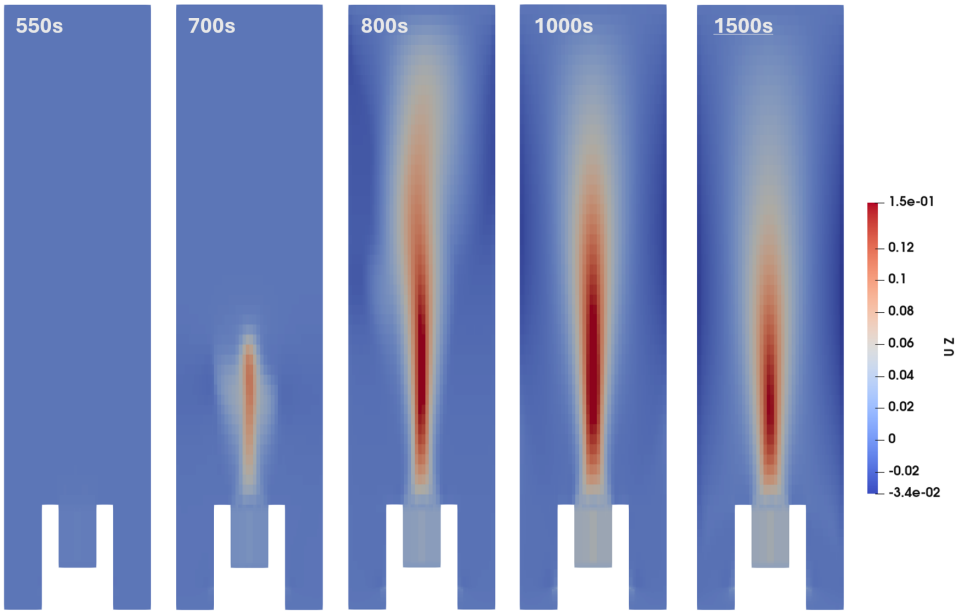

Figure 7.5: Development of the vertical velocity component u_z in a core cross-section at successive instants, illustrating the establishment of natural circulation.

The field also shows the return path that closes the loop: the upward velocity peaks at about $0.15\,\mathrm{m\,s^{-1}}$ along the central channel, while a weak descending flow of the order of $-0.03\,\mathrm{m\,s^{-1}}$ develops at the periphery. The presence of both branches confirms that the buoyancy-driven circulation is correctly closed within the pool, and it reinforces the earlier remark that the loop mass flow rate is not imposed but settles from the balance between thermal head and flow resistance.

7.2 Neutronic results

With the neutronic model fully specified as described in Chapter 5, a systematic code-to-code comparison against the `Serpent 2` Monte Carlo reference is carried out to verify the correct setup of the deterministic model and to establish the reliability of its predictions before proceeding to the coupled multiphysics calculation. The comparison covers five quantities of increasing physical specificity: the effective multiplication factor, the spatial power distribution, the moderator and Doppler temperature coefficients, and the integral control rod worth.

All presented results correspond exclusively to the control rod configuration with $Z_{\mathrm{AUT}} = 445$ mm and $Z_{\mathrm{MAN}} = 400$ mm. For each control rod, a dedicated set of two-group cross sections was generated from the `Serpent 2` full-core calculation corresponding to the fully inserted configuration. The `CRmove` utility (see Section 3.4.1) was then used to place the rods at the reference positions before running the neutronics diffusion simulation.

7.2.1 Effective Multiplication Factor

The first quantity used to assess the consistency between the two codes is the effective multiplication factor for the reference configuration:

$$k_{\mathrm{eff,\,Serpent}} = 1.00164 \pm 6.1 \text{ pcm}, \tag{7.1}$$

while the `foamForNuclear` diffusion solver converges to

$$k_{\mathrm{eff,\,fFN}} = 0.97255, \tag{7.2}$$

resulting in an absolute discrepancy of $\Delta k \approx 2908$ pcm. This offset is consistent with the known limitations of the two-group diffusion approach: the condensation of the neutron energy spectrum into two broad groups, combined with the spatial homogenisation of the assembly-averaged cross sections, introduces an inherent bias in the predicted eigenvalue, particularly for a relatively small reactor core where these approximations are expected to have a stronger impact [17].

Crucially, this offset is expected to cancel in the computation of differential reactivity quantities, since both codes apply the same perturbation to the same baseline. The comparison of the reactivity coefficients and control rod worth presented in Sections 7.2.3–7.2.4–7.2.5 is therefore not compromised by the absolute k_{eff} discrepancy.

7.2.2 Power Distribution

The spatial distribution of fission power density provides a more stringent test of the neutronic model, as it probes the ability of the diffusion solver to reproduce the local flux shape rather than a single integral quantity. The comparison is performed assembly-by-assembly over the 5×5 core lattice at a nominal power of 100 kW.

Figure 7.6 shows the relative error $\varepsilon = (P_{\text{fFN}} - P_{\text{Serpent}})/P_{\text{Serpent}}$ for each assembly. The maximum absolute error reaches 9.51%, localised at corner position B3. In the central fuel region, the discrepancy is generally below 2%, which is satisfactory for a two-group diffusion model. A spatial pattern is visible: assemblies adjacent to

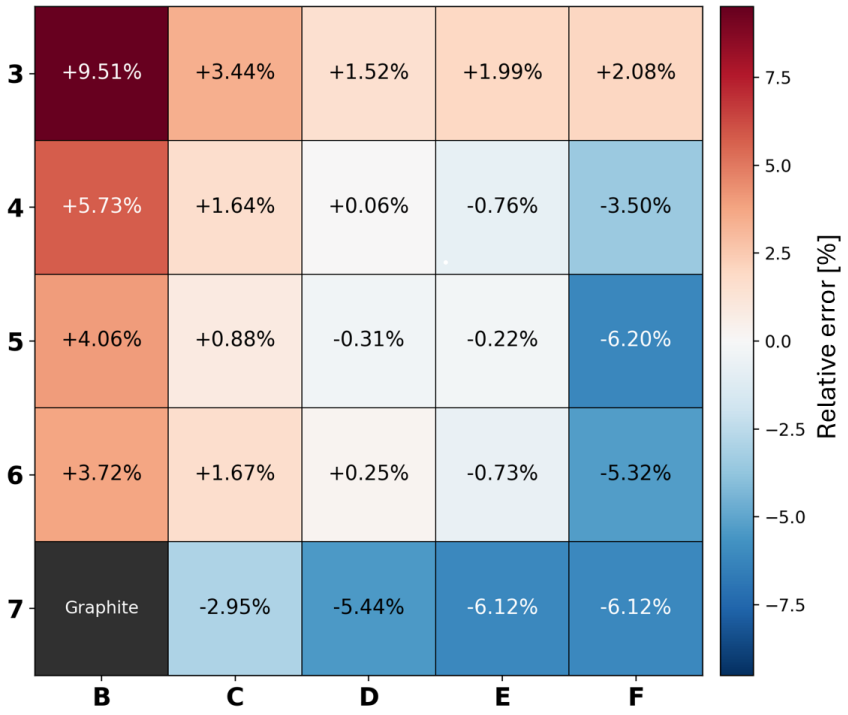

Figure 7.6: `GeN-Foam` vs. `Serpent 2` relative error in radial power distribution.

the graphite reflector tend to show a positive bias (foamForNuclear overpredicts), while those bordering the water blocks show a negative bias. This tilt reflects two compounding effects. First, the two-group cross sections fail to fully capture the spectral softening near the graphite interface, leading to an overestimation of the flux at material interfaces with large gradients. Second, at the water–core boundary—which represents a smoother gradient than the graphite interface—the diffusion solver underpredicts the power distribution. This is due to the strong absorbing properties of water, which are accurately resolved by `Serpent 2` but cannot be perfectly captured by the diffusion approximation.

The axial power density profiles ($P/\langle P\rangle$) compared assembly-by-assembly in Figures 7.7–7.8 confirm a substantially better agreement in the axial direction. The normalised power shape predicted by `GeN-Foam` follows the `Serpent 2` cosine-like distribution closely over the entire active length, with minor deviations concentrated at the two axial ends of the core. Because the black boundary conditions are applied at the outer edges of the 500 mm water reflector placed above and below the core, these local discrepancies are purely an interface transport effect: diffusion theory struggles with the strong gradients and highly anisotropic neutron fields at the core–water junction, whereas the Monte Carlo solution resolves this physics exactly.

The apparent spatial shift between the two curves does not represent a physical asymmetry or an axial flux tilt. This discrepancy is a visualisation artefact arising

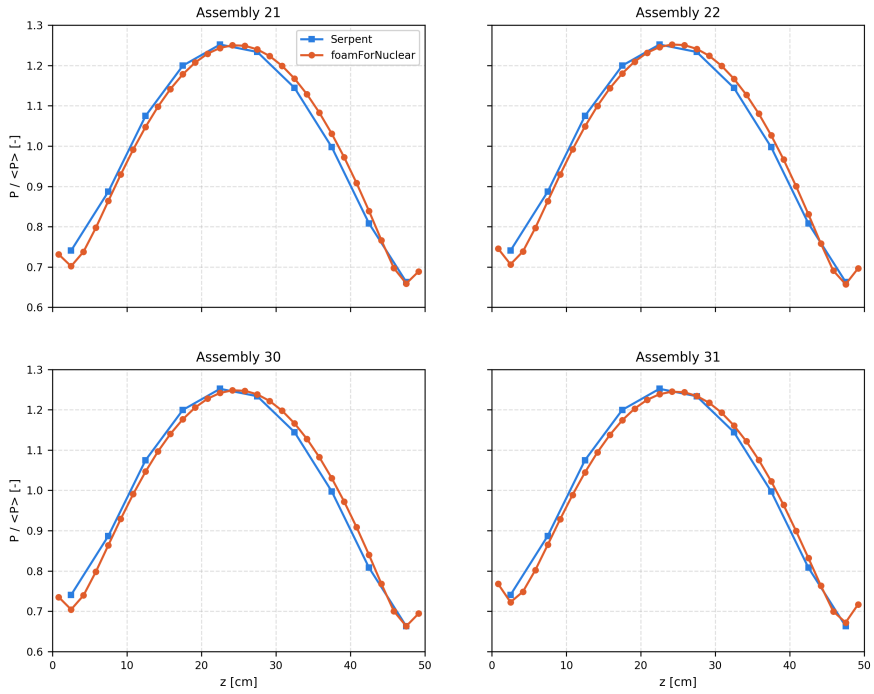

Figure 7.7: Axial power density profiles ($P/\langle P\rangle$): `GeN-Foam` vs. `Serpent 2`, AUT assembly.

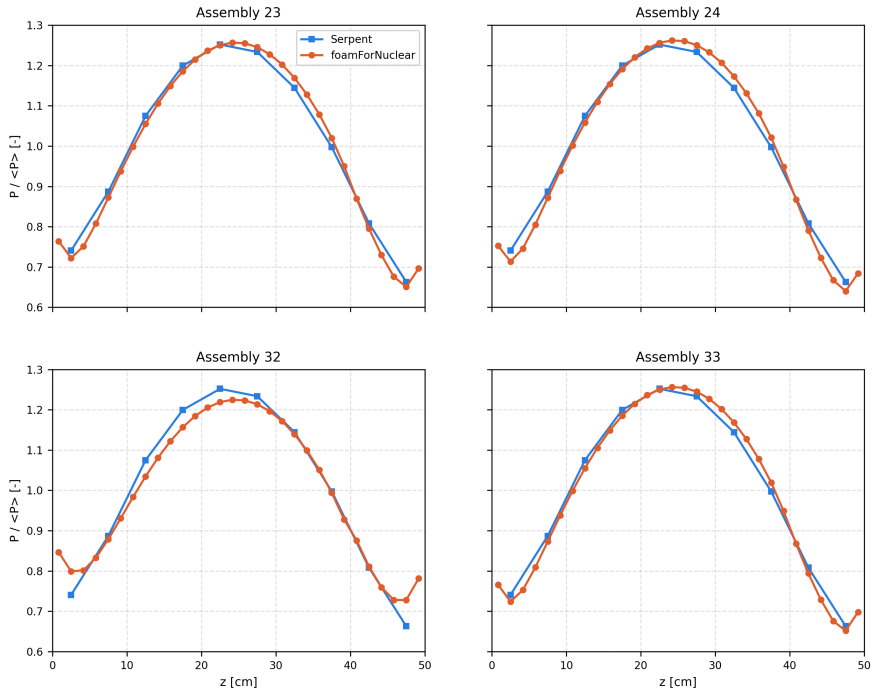

Figure 7.8: Axial power density profiles ($P/\langle P\rangle$): `GeN-Foam` vs. `Serpent 2`, MAN assembly.

from the different mesh refinements used in the two simulations: `foamForNuclear` employs a fine mesh with 30 axial nodes, whereas the `Serpent 2` solution is sampled over only 10 axial zones. Plotting the values at their respective cell-centre positions therefore introduces a relative geometric offset that creates a visual impression of an axial tilt. No systematic asymmetry is present in the axial direction, confirming that the tilt observed in the radial map originates from the lateral geometry of the core rather than from an axial modelling issue.

7.2.3 Moderator Temperature Coefficient

The Moderator Temperature Coefficient (MTC) quantifies the change in reactivity with coolant temperature, driven primarily by the strong nonlinear dependence of water density on temperature. As temperature rises, reduced moderation and increased neutron leakage produce a negative MTC, which constitutes an inherent safety feedback of pool-type reactors [44]. This is evaluated by perturbing the coolant temperature and density simultaneously, consistent with the physical condition of a single-phase pool reactor where the two quantities are thermodynamically coupled. Cross-section sets are generated with `Serpent 2` at five moderator temperatures spanning the range 296–370 K, each at the corresponding water density from the IAPWS-95 equation of state [20], using the `CoolProp` thermophysical property library [45].

Three strategies are compared for the `foamForNuclear` calculation. The *case-by-case* approach runs a dedicated static eigenvalue calculation at each temperature point, using the exact cross sections for that condition; it serves as the internal `GeN-Foam` reference since no interpolation error is introduced. The *LinInterp-3* variant uses only three XS sets (296, 330, 370 K) and reconstructs intermediate values by piecewise linear interpolation. The *LinInterp-5* variant adds two intermediate points (310 K and 350 K), reducing the interpolation interval and improving accuracy at the cost of two additional `Serpent 2` runs.

For each strategy, the MTC is evaluated using a local finite-difference approximation. Given two consecutive temperature points T_i and T_{i+1}, the corresponding reactivities are first computed as

$$\rho_i = \frac{k_{\text{eff},i} - 1}{k_{\text{eff},i}} \times 10^5 \quad [\text{pcm}], \tag{7.3}$$

and the MTC is then estimated as

$$\alpha_M\left(\bar{T}\right) = \frac{\rho_{i+1} - \rho_i}{T_{i+1} - T_i} \quad [\text{pcm/K}], \tag{7.4}$$

with the result plotted at the interval midpoint $\bar{T} = (T_i + T_{i+1})/2$. Both the case-by-case and the `Serpent 2` calculations are performed at the same five temperature points (296, 310, 330, 350, and 370 K), yielding four MTC values at the corresponding midpoints. For the case-by-case strategy, $k_{\text{eff},i}$ is obtained from a dedicated static

eigenvalue solve using the exact cross-section set for that temperature condition. The same finite-difference procedure is applied to the `Serpent 2` Monte Carlo k_{eff} values, providing an independent reference for each temperature interval. For the LinInterp variants, k_{eff} is instead obtained using interpolated cross sections at the requested temperature, and Eq. (7.4) is applied in the same manner.

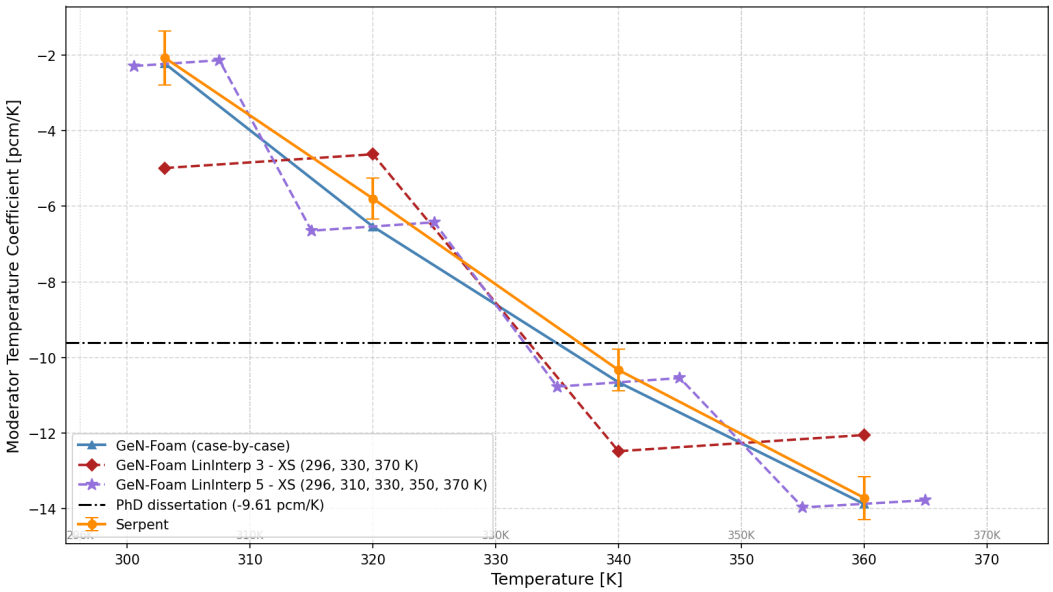

Figure 7.9: Moderator Temperature Coefficient: comparison of evaluation strategies.

The resulting MTC curves are plotted in Figure 7.9 together with the `Serpent 2` reference and the value of −9.61 pcm/K reported by [5]. All methods reproduce the negative and nonlinear character of the coefficient: the MTC becomes substantially more negative as the temperature rises from 300 K toward 370 K, driven by the accelerating decrease in water density that progressively reduces the moderation capability of the medium. The case-by-case and `Serpent 2` curves are in good agreement throughout, with an average feedback value of −8.83 pcm/K, approximately 8% below the literature reference [5]. The LinInterp-3 curve deviates noticeably in the 320–345 K region, where the coefficient changes most rapidly; the finer discretisation of LinInterp-5 recovers most of this loss of accuracy. For the coupled thermal-hydraulic simulations, the five-point linear interpolation is adopted as the operational parametrisation, balancing fidelity and computational cost.

7.2.4 Doppler Temperature Coefficient

The Doppler Temperature Coefficient (DTC) captures the reactivity feedback arising from thermal broadening of the ^{238}U resonance capture cross sections. As the fuel temperature increases, the effective width of the resonances grows in proportion to $\sqrt{T}$, enhancing neutron absorption and reducing reactivity [44]. This prompt negative feedback is physically distinct from the MTC and acts on a much faster timescale, making it the primary safety mechanism during rapid transients.

Cross sections are generated at fuel temperatures of 300, 400, 500, and 600 K using

the on-the-fly TMS Doppler broadening method; the moderator density and temperature are held fixed at their reference values throughout. Three `foamForNuclear` strategies are compared and shown in Figure 7.10.

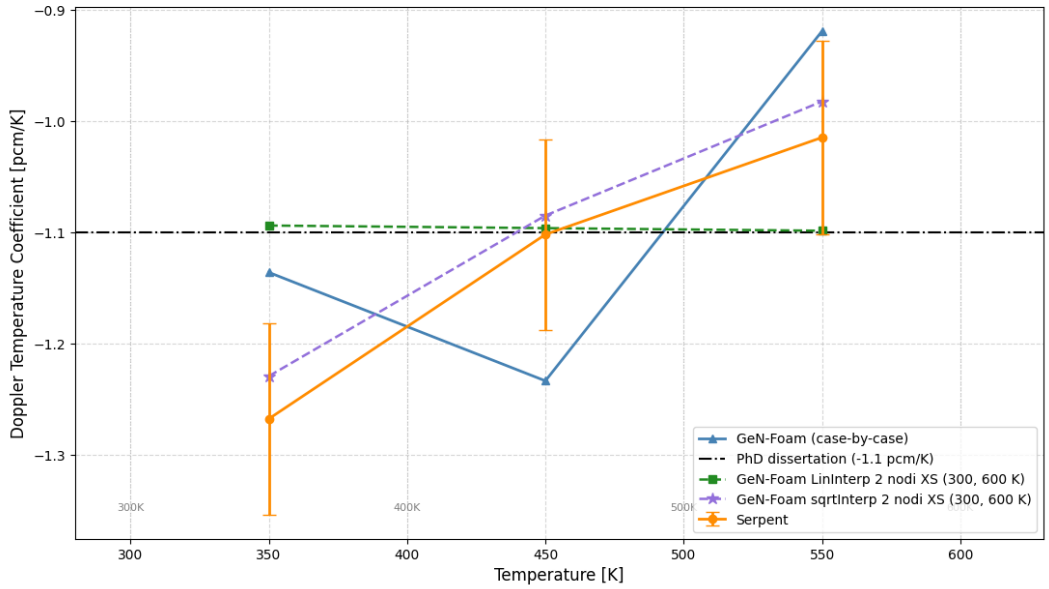

Figure 7.10: Doppler Temperature Coefficient: comparison of evaluation strategies.

The case-by-case approach again provides the reference, but its finite-difference estimate, obtained using the procedure described in Section 7.2.3, is sensitive to numerical noise because the reactivity change between adjacent fuel temperature points is small (∼50 pcm), comparable to the convergence tolerance of the eigenvalue solver; the resulting curve therefore shows visible scatter. The *LinInterp* variant uses only two boundary cross-section sets (300 and 600 K) with piecewise linear interpolation, producing an almost flat, constant DTC of approximately −1.12 pcm/K that does not capture any temperature dependence. The *sqrtInterp* variant employs the same two boundary datasets but interpolates according to a $\sqrt{T}$ law, consistent with the physical origin of the Doppler effect. This approach recovers the mild temperature dependence observed in `Serpent 2`—the coefficient becomes slightly less negative at higher fuel temperatures as the resonances begin to saturate—while requiring no additional XS generation.

The average feedback value of −1.12 pcm/K agrees with both the `Serpent 2` estimate and the literature reference of −1.1 pcm/K [5]. Given the motivated shape and the consistency with the reference solution, the $\sqrt{T}$ interpolation is retained for the coupled calculation. The DTC magnitude is roughly one order of magnitude smaller than the MTC for this reactor type.

7.2.5 Control Rod Worth

The integral reactivity worth of the control rods is computed by performing a series of static k-eigenvalue calculations at discrete rod insertion depths, from the fully withdrawn position to full insertion in 50 mm steps. The worth at each position is

defined as

$$\rho(z) = \frac{k_{600} - k(z)}{k_{600} \cdot k(z)}, \tag{7.5}$$

where k_{600} is the eigenvalue of the fully withdrawn reference configuration. The mesh is aligned with each rod-tip position as described in Section 5.3, so that no partial-node mixing is required and cusping artefacts are avoided.

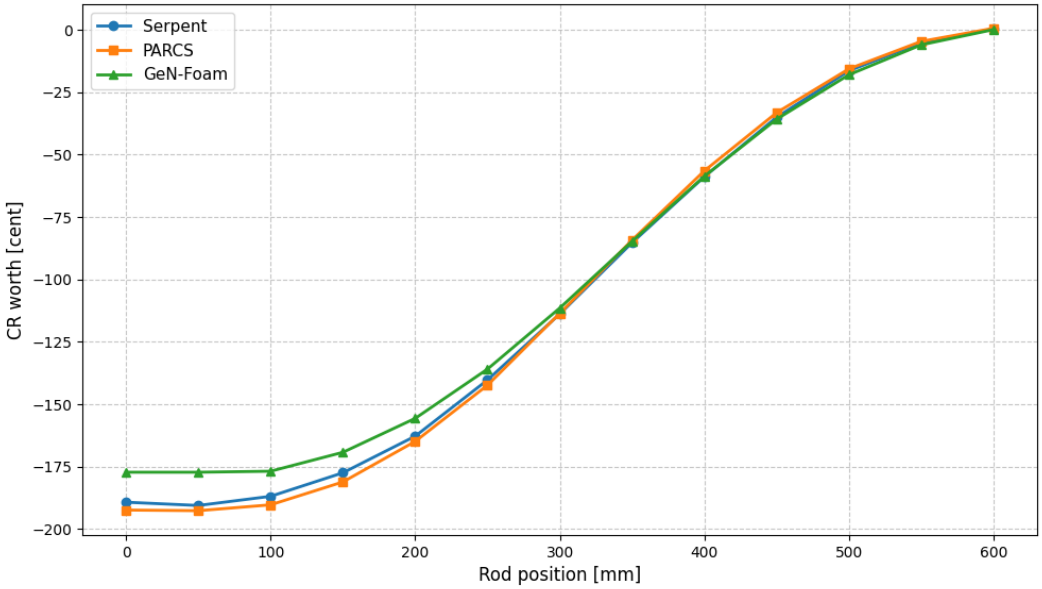

Figure 7.11: Integral reactivity worth of the MAN control rod as a function of insertion depth.

Figure 7.11 shows the integral worth curve of the MAN rod as a function of insertion depth, compared against the `Serpent 2` reference and the PARCS nodal diffusion solution. All three codes reproduce the expected S-shaped profile: the rod is most effective near mid-core, where it overlaps with the peak of the axial flux, and progressively less effective toward the core extremities, where the flux is attenuated by the axial reflector. PARCS and `Serpent 2` are in near-perfect agreement throughout the curve. `GeN-Foam` follows the same trend but predicts a slightly more negative worth at low insertion depths (rod near full insertion), with a maximum deviation of approximately 13 cent in the fully inserted configuration. Integrated over the full stroke, the MAN rod worth from `GeN-Foam` deviates from the `Serpent 2` reference by approximately 6%.

An analogous comparison for the AUT rod, shown in Figure 7.12, yields a larger integral discrepancy of approximately 13%. The AUT rod is located in a position of stronger flux asymmetry within the core lattice, which makes its worth more sensitive to homogenisation errors, as it sits closer to the reflector boundary.

Despite these quantitative differences, the shape of the differential worth distribution is correctly reproduced by `GeN-Foam` in both cases, and the integral worth values remain within the accuracy band considered acceptable for the subsequent coupled thermal-hydraulic calculations.

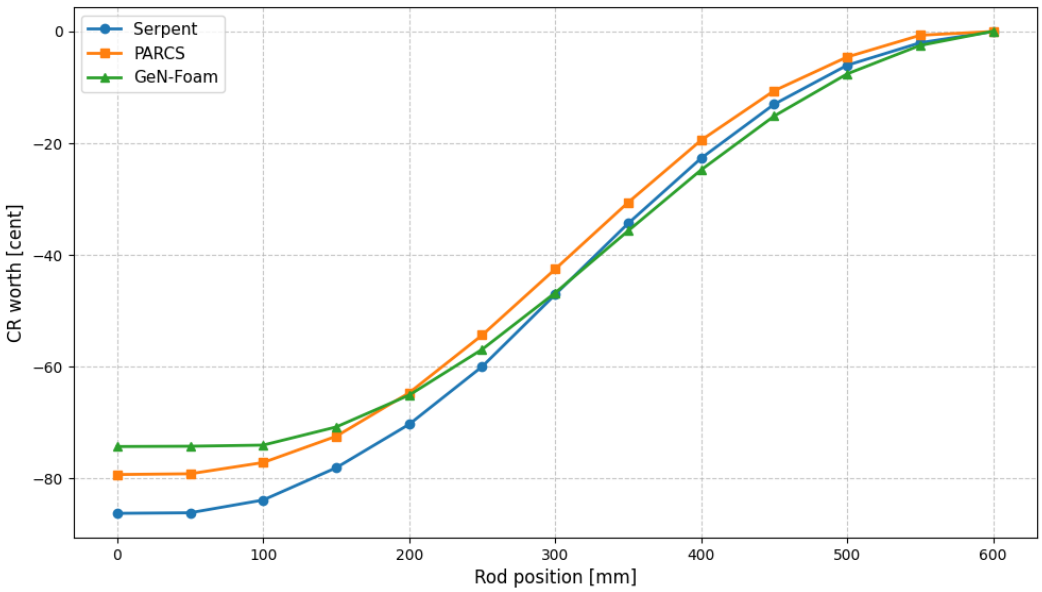

Figure 7.12: Integral reactivity worth of the AUT control rod as a function of insertion depth.

7.3 Coupled Neutronics–Thermal-Hydraulics Analysis

The preceding sections addressed steady-state neutronics and single-physics thermal-hydraulics in isolation. The present part brings the two together through the coupled solver embedded in `foamForNuclear`, with the aim of reproducing the power and temperature transients (FB) experiments performed at the BME Training Reactor during the June 2022 campaign [5].

A first coupled analysis is conducted using the point-kinetics module of `foamForNuclear`, driven by a time-dependent external reactivity profile applied through the `externalReactivityInsertion` functionality of the code, explained in detail in Section 5.3.1.

7.3.1 Reference Transient: FB#2

Before turning to the coupled results, it is worth briefly recalling the case under study, already described in detail among the model (see Section 5.3.1). Transient FB#2 is one of the cold-insertion experiments carried out at the BME TR in June 2022 [5], and it belongs to the cold-insertion category (Setup #1): the reactor starts from a low, quasi-critical level and a positive reactivity step is introduced by partially withdrawing the MAN control rod by 40 mm, from 380 to 420 mm, while the AUT rod stays fixed at 558 mm throughout. Over a total duration of about 1550 s the power rises smoothly toward a new equilibrium near 46 kW nominal, the climb decelerating as the warming coolant feeds back through the moderator temperature coefficient. Outlet temperatures are monitored on four assemblies E3, B3, C7 and F7. Consistently with the calibration discussed earlier (see Section 6.1.1), both the imposed power and the measured trace plotted for comparison are scaled by the factor 1.44, so that the nominal 5 W–45.9 kW range corresponds to actual values of roughly 7.2 W–66.1 kW.

In the simulation the computational domain is the same one used for the steady-state thermal-hydraulic study and the k-eigenvalue calculation, with the porous-medium core, the mapped pressure–velocity condition closing the natural-circulation loop, and the `icoPolynomial` equation of state fitted to NIST water data; the neutronic mesh is built to communicate seamlessly with the thermal-hydraulic one. The system is initialised at a uniform 296.15 K, with the vessel walls held at the same value (discussed in Section 6.1.3), and from this isothermal state it is first brought to a converged low-power steady state at $P_{\text{reactor}} \approx 7.2\,\text{W}$, so that velocity and temperature fields are mutually consistent with an equilibrium critical state before the transient begins. The transient itself is then driven by the extraction of the MAN rod: unlike the preceding steady phase, where the power was pinned at 7.2 W, the power evolution is now dictated entirely by the neutronic response to the rod movement, with the thermal-hydraulic and neutronic solvers fully coupled.

7.3.2 Coupled Transient Results

With the case and its numerical setup in place and carried out, the coupled transient results can be examined. The analysis follows the physical chain that drives the transient itself: the withdrawal of the control rod inserts reactivity, the resulting imbalance is progressively countered by the temperature feedbacks and sets the power evolution, and the power in turn establishes the temperature field across the core. The results are presented in this order, reactivity balance first, then power, then assembly temperatures, so that each quantity provides the context needed to read the next.

Reactivity Balance and Feedback Response The transient is initiated by the reactivity step associated with the MAN rod extraction, after which the evolution is governed by the interplay between this external insertion and the temperature feedbacks it triggers. Figure 7.13 decomposes the reactivity balance over time, separating the imposed insertion from the Doppler and moderator contributions. Immediately after the withdrawal the net reactivity is positive and the power begins to rise; as the fuel and the coolant heat up, the feedbacks grow increasingly negative and drive the balance back toward zero, defining the new quasi-critical plateau. In a pool-type reactor the moderator term, acting through the water density, is expected to dominate this braking action, consistently with the moderator coefficient being considerably larger in magnitude than the Doppler one. The decomposition, however, shows the two contributions to be comparable, the Doppler feedback reaching a magnitude at least as large as the moderator one: this follows directly from the fuel temperature excursion being much larger than the coolant one, which offsets the smaller Doppler coefficient. Two further features are worth noting, as they anticipate the power behaviour. First, the feedbacks build up gradually over several hundred seconds, a time scale far longer than the neutronic period, so that the approach to the plateau is governed by the thermal response of the system rather than by the kinetics. Second, the net reactivity does not settle monotonically but slightly overshoots the null point, turning marginally negative before recovering, which is the reactivity signature of the power overshoot discussed next.

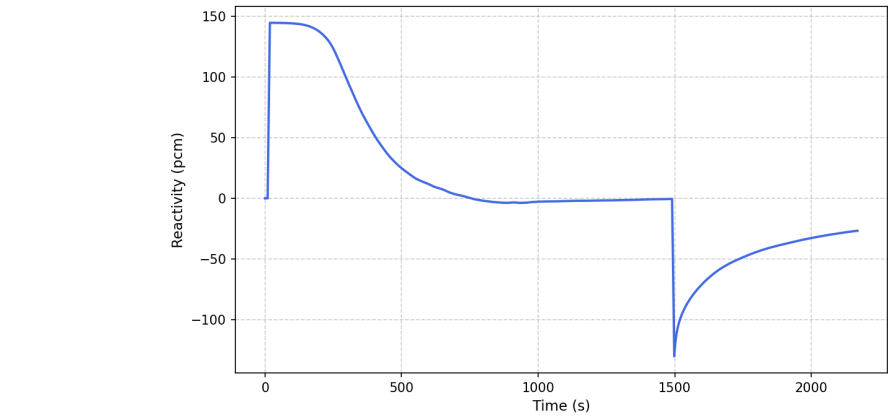

(a) Total reactivity

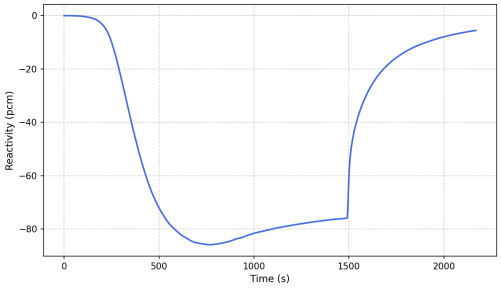

(b) Doppler (fuel) feedback

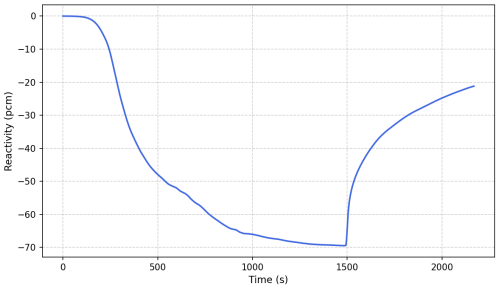

(c) Moderator (coolant) feedback

Figure 7.13: Reactivity balance during the FB#2 transient: total reactivity (a) and its decomposition into the Doppler (b) and moderator (c) feedback contributions.

Power Evolution The power history is the direct consequence of the reactivity balance just described, since it integrates the net reactivity through the point-kinetics response. Figure 7.15 compares the simulated power against the measured trace, both scaled by the calibration factor introduced earlier. The simulation reproduces the qualitative shape of the transient, a rise from the quasi-critical level followed by a deceleration toward a new plateau, but settles at a higher power than the measurement, stabilising at roughly 140 kW after an overshoot to about 165 kW, against a measured plateau of the order of 66 kW. The origin of this discrepancy is discussed here, distinguishing between the possible contributions of the imposed reactivity, the magnitude of the feedback coefficients, and the thermal-hydraulic mapping of power into temperature, and it is examined jointly with the temperature results of the following subsection.

The imposed reactivity is confirmed by the inhour verification of Figure 7.14, where the observed asymptotic period matches the theoretical one associated with the 145 pcm insertion to within about one percent; the driving term of the transient is therefore correct. The feedback coefficients, in turn, were determined and verified during the neutronic characterisation, so the reactivity returned per unit temperature is likewise consistent with the reference. With the insertion and the coefficients fixed, the equilibrium power is set solely by the requirement that the integral feedback cancel the imposed reactivity, and this condition is closed through the temperature

field. The overprediction must therefore originate in the thermal-hydraulic conversion of power into temperature: if the model develops too small a coolant temperature rise per unit power, the feedback generated per unit power is correspondingly weak, and the reactor must climb to a higher power before the balance is restored. The plateau discrepancy is thus a steady-state feature of the heat-removal modelling rather than of the neutronics.

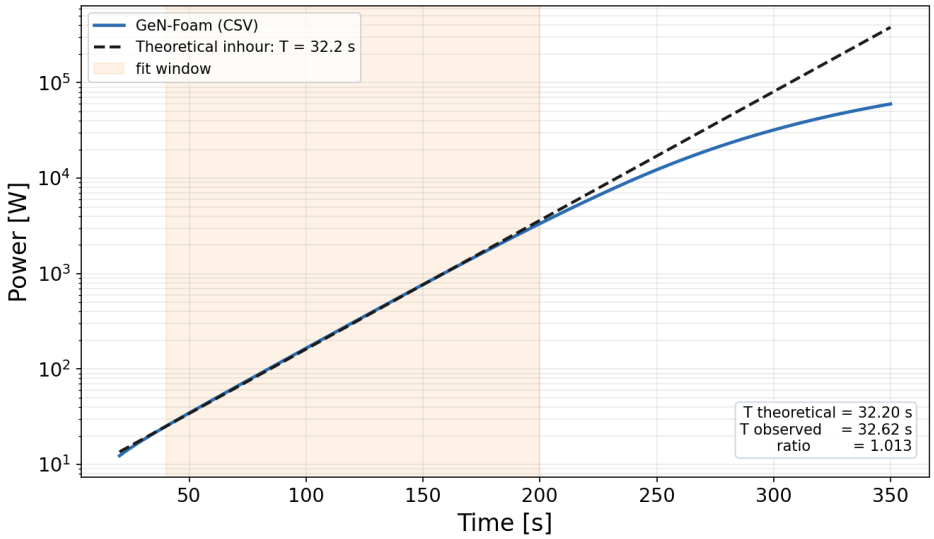

Figure 7.14: Verification of the inhour relation for the 145 pcm insertion.

The overshoot preceding the plateau reflects the same coupling seen from its dynamic side. The peak at about 165 kW mirrors the transient dip of the net reactivity below zero: because the negative feedback establishes on a thermal time scale much longer than the neutronic period, the power rises past its equilibrium value before the accumulating feedback catches up and pulls it back. Under natural circulation this delay is compounded, since the coolant mass flow is not imposed but emerges from the buoyancy field and must redevelop as the power increases; until the flow profile is established the coolant temperature, and with it the moderator feedback, lags the power. The thermal inertia of the system and the establishment of the natural-circulation flow therefore set both the amplitude of the overshoot and the rate at which the plateau is approached, and are examined against the temperature response in the following subsection.

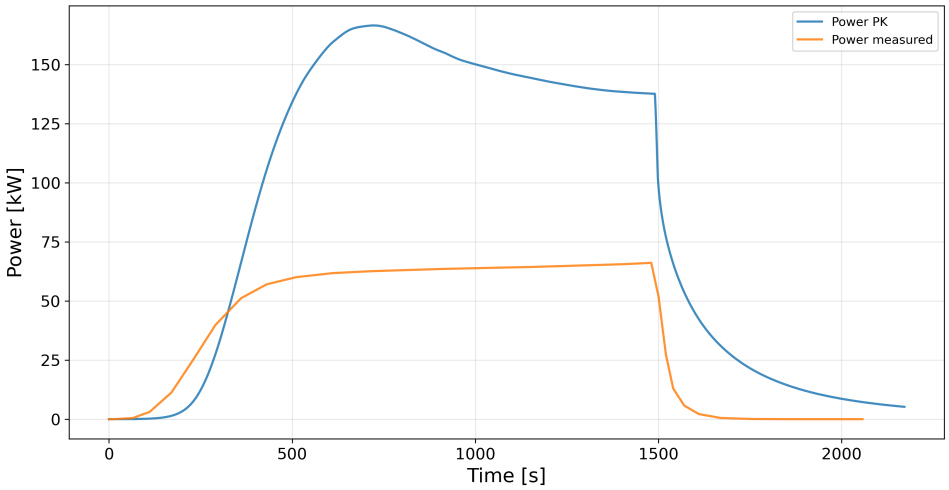

Figure 7.15: Simulated and reference power histories for the FB#2 transient.

Assembly Temperature Distribution Finally, the temperature response of the four monitored assemblies (E3, B3, C7 and F7) closes the chain, showing how the coupled power is translated into a thermal field across the core. Figure 7.16 compares the simulated outlet temperatures against the measurements. The expected radial trend is recovered, with the central assembly E3 the hottest and the peripheral ones cooler, while the F7 and B3 assemblies appear in reversed order with respect to the measurements – the same inversion already encountered and discussed for the standalone verification. Throughout the power phase the simulated temperatures lie above the measured ones, consistently with the power overprediction: the excess power and the excess temperature are two views of the same discrepancy, and their common origin is the thermal-hydraulic coupling.

The cool-down phase provides an independent confirmation of this interpretation. Once the rod is reinserted and the reactivity forcing is removed, the relaxation of the temperatures is governed almost entirely by the thermal time constant of the system, with the feedback no longer driving the transient. In this phase the simulated temperatures decay more slowly than the measured ones and reconverge only late in the transient, indicating that the model's thermal time scales do not match the physical ones. Because it is largely decoupled from the reactivity balance, the cool-down isolates the thermal-hydraulic response and corroborates the conclusion drawn from the power comparison. Taken together, the reactivity, power and temperature results locate the discrepancy in the thermal-hydraulic solution under natural circulation rather than in the neutronics: the higher plateau reflects the steady-state power-to-temperature mapping, while the overshoot and the slow relaxation reflect the delayed establishment of the feedback, and both trace back to the same modelling of heat removal and flow under natural circulation.

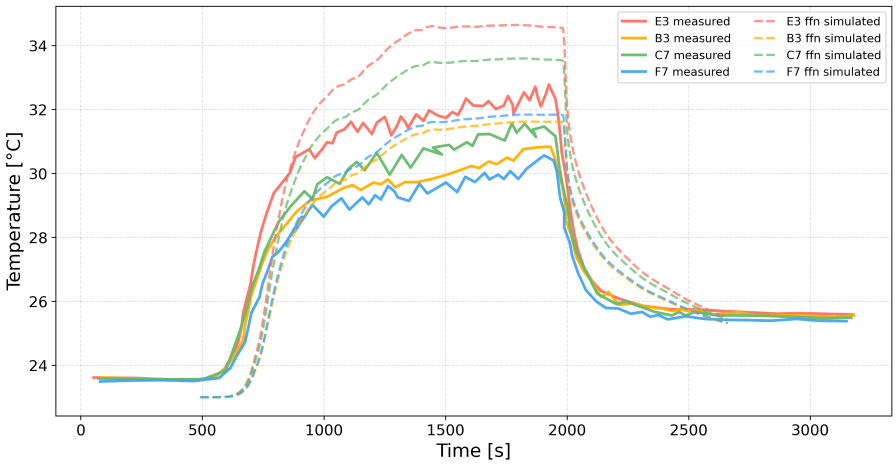

Figure 7.16: Measured and simulated temperature histories for the monitored assemblies E3, B3, C7 and F7.

Chapter 8

Conclusion and Future Works

This thesis has assessed foamForNuclear using the BME Training Reactor through three progressively coupled stages, contributing to the systematic evaluation of high-fidelity multiphysics tools pursued within Work Package 2 of the EVEREST project. Each stage was designed to isolate one physical ingredient before adding the next, so that the behaviour of the coupled model could be interpreted in the light of its verified components.

The stand-alone thermal-hydraulic model showed that the hybrid CFD–porous-medium formulation adequately resolves the natural-circulation dynamics of the pool, reproducing the central buoyant plume, the descending peripheral return current, and the axial thermal stratification observed in the experimental records. The early-time evolution proved sensitive to the vessel wall boundary condition, with a modest thermal deficit being sufficient to trigger spurious peripheral downflows. This confirmed that physically motivated initial and boundary conditions are as consequential as the turbulence closure or the spatial resolution.

The neutronic verification against Serpent 2 yielded the behaviour expected of a two-group diffusion solver applied to a small, heterogeneous core: an absolute eigenvalue offset of about 2900 pcm that cancels in differential reactivity comparisons, assembly power errors generally below 2% in the central region and larger towards the graphite reflector, and reactivity coefficients within 8% of the Monte Carlo reference. The control-rod worth curves reproduce the expected S-shape, with integral discrepancies of approximately 6% and 13% for the MAN and AUT rods, respectively. The larger AUT error reflects its proximity to the reflector, where the homogenisation assumptions are most challenged. These figures fall within the accuracy expected of the two-group approximation and were judged adequate for the coupled calculations.

The coupled simulation of transient FB#2 reproduced the qualitative evolution of the experiment: the power rise from the quasi-critical level, its deceleration under the accumulating temperature feedbacks, and the differential outlet heating between inner and peripheral assemblies. The neutronic response was confirmed independently, with the simulated asymptotic period matching the inhour prediction for the 145 pcm insertion to within about one percent. The coupled model did not, however,

reproduce the measured values: the power settled at roughly twice the reference level, after an initial overshoot, and the assembly temperatures were consequently overpredicted. With the reactivity insertion verified by the inhour analysis and the feedback coefficients by the neutronic characterisation, this discrepancy was traced to the thermal-hydraulic mapping of power into temperature under natural circulation—namely, to the steady-state heat removal, which determines the plateau, and to the delayed establishment of the feedback, governed by the thermal inertia of the system and by the development of the buoyancy-driven flow, which produces the overshoot. The operator-split coupling, appropriate for the slow thermal timescales of natural-circulation operation, introduced no significant splitting error over the transient durations considered.

Taken together, these results establish foamForNuclear as a physically consistent tool for the multiphysics analysis of pool-type research reactors, with verified neutronic and thermal-hydraulic components that behave as expected and a coupled response that captures the correct phenomenology. They also identify, with reasonable precision, where the present model still departs from experiment and, in doing so, define the next steps rather than conclude the investigation.

These results should therefore be regarded as the first stage of a validation process designed to continue, and which the author intends to pursue within the EVEREST WP2 programme. The framework assembled here is intended less as a finished product than as a foundation upon which each remaining discrepancy can be examined in turn and progressively resolved.

The immediate priority follows directly from the coupled transient itself. The overprediction of the power plateau points to the need for a closer study of the system thermal timescales and of the establishment of the natural-circulation feedback, for which the cool-down phase offers a comparatively clean diagnostic, isolating the heat-removal dynamics from the neutronic transient. Clarifying this mechanism is the necessary precondition for reconciling the coupled model with the measured values.

Beyond this, the natural direction for future development is a progressive refinement of the computational geometry, bringing it closer to the physical reactor by relaxing, one at a time, the simplifying assumptions adopted in the present work. Three extensions stand out. The graphite reflector, presently treated through homogenised cross-sections, could be represented explicitly as a set of conjugate-heat-transfer solid regions, resolving both its neutronic role and its thermal coupling to the pool. The instrumentation water channels, omitted from the current lattice, could be introduced as porous-medium zones, restoring their effect on the local flow and moderation. Finally, a thin porous layer could be added above the active fuel to account for the hydraulic resistance and thermal mass of the thermocouple holders, whose influence on the near-field temperature response is presently unmodelled. Each of these developments is compatible with the existing modelling infrastructure and can be introduced incrementally, so that its effect on the coupled response may be assessed in isolation rather than confounded with the others.

The culmination of this effort will be the extension of the validation to the other

transients that the EVEREST WP2 programme still has to address, and towards which the present power cases represent a deliberate first step. Pursued along these lines, the work begun here should progress from a model that is physically consistent to one that is also reliable, accurate, and complete across the full range of conditions required by the programme.

Appendix A

Experimental Setup

This appendix provides a visual representation of the instrumentation layout within the reactor vessel. The following figures illustrate the geometric positioning of the thermocouple sensor rods, detailing their configuration and vertical distribution relative to the core level

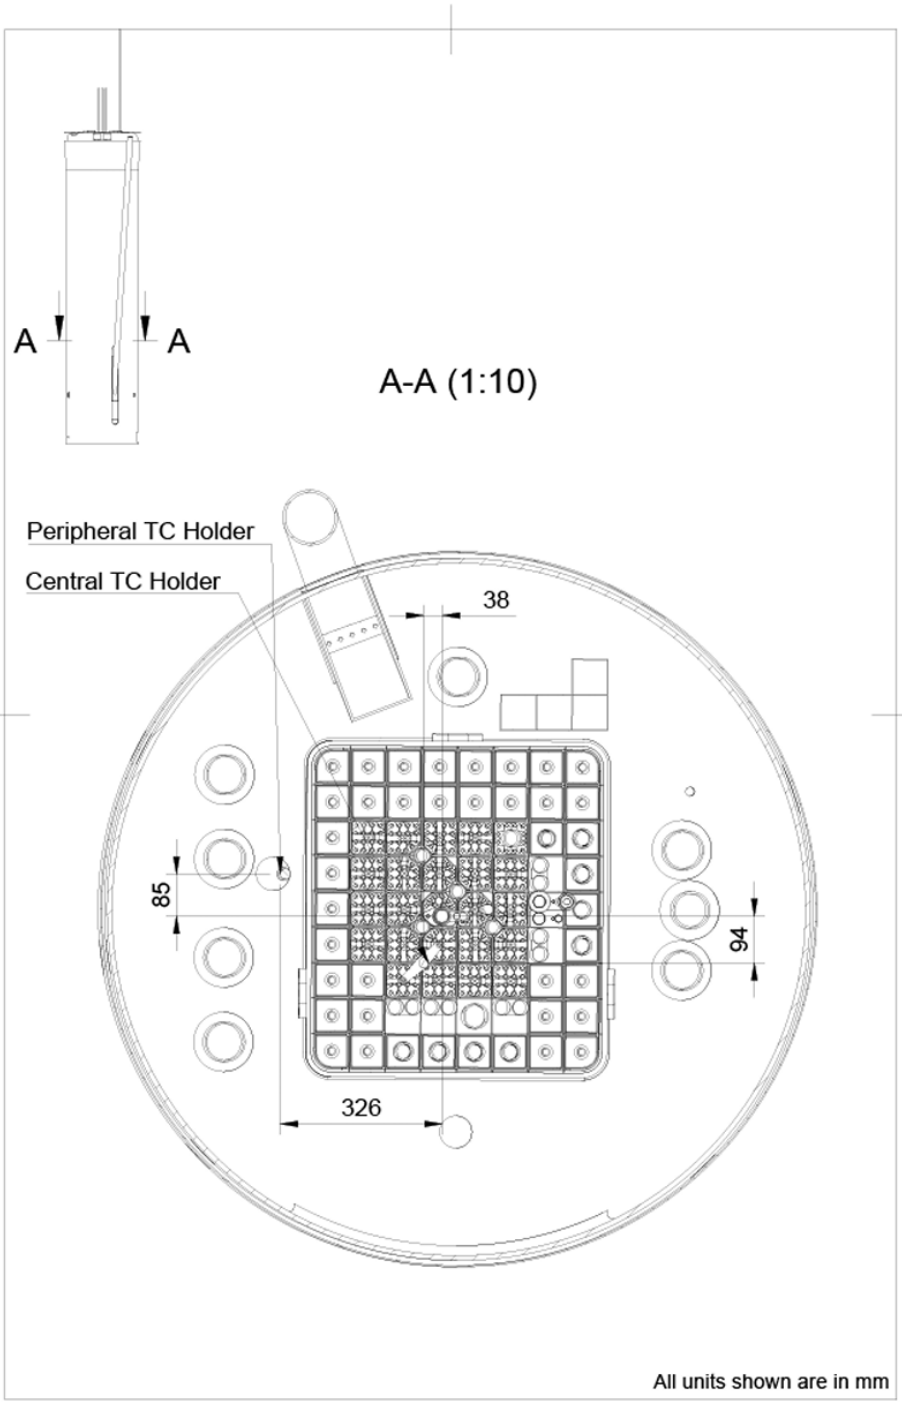

Figure A.1: Experimental setup of the thermocouple holders (view from the top) [15].

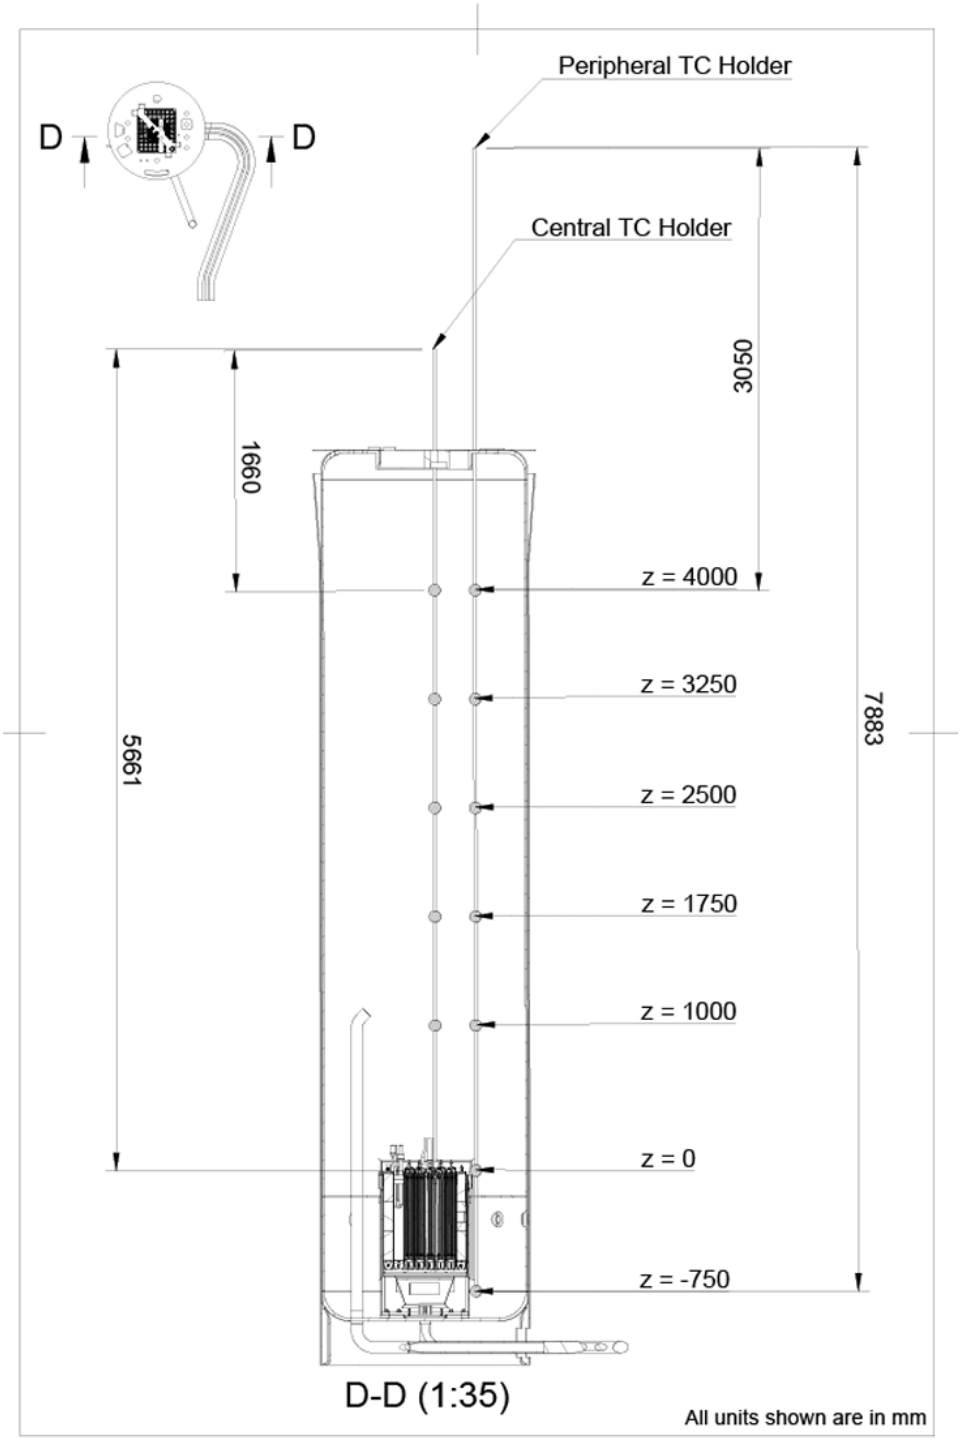

Figure A.2: Experimental setup of the thermocouple holders (lateral view) [15].

Bibliography

[1] Hyunjong Kim, In-Guk Kim, Minkyu Jung, Kyoung-Woo Seo, and Seong-Hoon Kim. "Numerical study on the uncertainty reduction for thermal power estimation of an open pool research reactor with natural convection core cooling". In: *Annals of Nuclear Energy* 144 (2020), p. 107587. ISSN: 0306-4549. DOI: https://doi.org/10.1016/j.anucene.2020.107587. URL: https://www.sciencedirect.com/science/article/pii/S0306454920302851 (cit. on p. 1).

[2] M Hursin, Y Perin, NG Herranz, K Ambrožič, and SZ Czifrus. "The EVER-EST project: Toward experimentally validated advanced simulation tools". In: *European Research Reactor Conference*. 2025, pp. 18–24 (cit. on pp. 1, 2).

[3] International Atomic Energy Agency. *Country Nuclear Power Profiles: Hungary*. IAEA-CNPP/2025. Access: April 2026. IAEA, 2025. DOI: 10.61092/iaea.o6f8-nuj2 (cit. on pp. 2, 4, 5).

[4] A.Sz. Ványi, M. Hursin, A. Aszódi, L. Adorján, B. Biró, B. Magyar, P. Mészáros, T. Bozsó, and Sz. Czifrus. "Thermal-hydraulic measurements and TRACE system code analysis performed on the natural circulation cooled BME Training Reactor". In: *Annals of Nuclear Energy* 189 (2023), p. 109839. ISSN: 0306-4549. DOI: https://doi.org/10.1016/j.anucene.2023.109839 (cit. on pp. 2, 27, 31, 33).

[5] A. S. Ványi. "Experimental and numerical multi-physics analysis for the BME Training Reactor". Access: April 2026. MA thesis. Budapest, Hungary: Budapest University of Technology and Economics, 2023. URL: https://m2.mtmt.hu/api/publication/34676831 (cit. on pp. 2, 5–8, 18, 27, 30, 31, 33–38, 41–44, 46, 53, 59, 60, 63).

[6] Giovanni Nervi, Thomas Guilbaud, Carlo Fiorina, Mathieu Hursin, and Alessandro Scolaro. "foamForNuclear: A unified OpenFOAM multi-physics platform for nuclear applications". In: *Progress in Nuclear Energy* 193 (2026). Available at: https://gitlab.com/foamForNuclear/foamForNuclear, p. 106250. ISSN: 0149-1970. DOI: 10.1016/j.pnucene.2026.106250 (cit. on pp. 2, 9, 13, 15–17, 21, 22, 24).

[7] Jaakko Leppänen, M. Pusa, Tuomas Viitanen, Ville Valtavirta, and Toni Kaltiaisenaho. "The Serpent Monte Carlo Code: Status, Development and

Applications in 2013". In: vol. 82. June 2014, p. 06021. ISBN: 978-2-7598-1269-1. DOI: 10.1051/snamc/201406021 (cit. on pp. 2, 19).

[8] OECD/NEA. *Nuclear Legislation in OECD and NEA Countries: Hungary (2008 update)*. Tech. rep. Access: April 2026. Paris: OECD Publishing, 2020. URL: https://www.oecd-nea.org/jcms/pl_24192/nuclear-legislation-in-oecd-and-nea-countries-hungary-2008-update (cit. on p. 4).

[9] IEA. *Hungary 2022*. Tech. rep. Licence: CC BY 4.0. Access: April 2026. Paris: IEA, 2022. URL: https://www.iea.org/reports/hungary-2022 (cit. on p. 4).

[10] MAVIR ZRt. *A rendszerirányítás 60 éves története [60 years of history of system operation]*. Tech. rep. Access: April 2026. Budapest, Hungary: MAVIR ZRt., May 2009. URL: https://mavir.hu/documents/10258/0/%5C%21PORTAL.wwpob_page.pdf/60bd6cfc-668d-3799-9d22-1f169ff1adbc (cit. on p. 5).

[11] M. Járosi and P. Kovács. "Energy Policy of Hungary". In: *Civic Review* 14.Special Issue (2018). Access: April 2026, pp. 67–80. DOI: 10.24307/psz.2018.0405 (cit. on p. 5).

[12] BME NTI. *The history of the Training Reactor*. Access: April 2026. 2022. URL: http://reak.bme.hu/oktatoreaktor/tortenet.html (cit. on p. 5).

[13] M. Kovács. "Mechanical design of BME TR spent fuel transportation". MSc Thesis (in Hungarian). Master's thesis. Budapest, Hungary: Budapest University of Technology and Economics, May 2011 (cit. on p. 6).

[14] A.Sz. Ványi, B. Babcsány, Z.I. Böröczki, A. Horváth, M. Hursin, M. Szieberth, and Sz. Czifrus. "Steady-state neutronic measurements and comprehensive numerical analysis for the BME training reactor". In: *Annals of Nuclear Energy* 155 (2021), p. 108144. ISSN: 0306-4549. DOI: https://doi.org/10.1016/j.anucene.2021.108144. URL: https://www.sciencedirect.com/science/article/pii/S0306454921000207 (cit. on p. 7).

[15] G. Klujber and A. Horváth. *Benchmark description of the BME Training Reactor*. Tech. Rep. BME-NTI-854/2018. Version 3.0. Budapest, Hungary: BME Institute of Nuclear Techniques (NTI), Oct. 2020 (cit. on pp. 7, 8, 30, 31, 44, 73, 74).

[16] A.Sz. Ványi, M. Hursin, and Sz. Czifrus. "Analysis of transient measurements with thermal feedback and coupled TRACE/PARCS calculations performed on the BME Training Reactor". In: *Annals of Nuclear Energy* 194 (2023), p. 110072. ISSN: 0306-4549. DOI: 10.1016/j.anucene.2023.110072 (cit. on pp. 8, 38).

[17] Christophe Demazière. *Modelling of nuclear reactor multi-physics: from local balance equations to macroscopic models in neutronics and thermal-hydraulics*. Academic Press, 2019 (cit. on pp. 9, 55).

[18] H. G. Weller, G. Tabor, H. Jasak, and C. Fureby. "A tensorial approach to computational continuum mechanics using object-oriented techniques". In: *Computers in Physics* 12.6 (1998), pp. 620–631 (cit. on pp. 9, 14).

[19] Bruce R. Munson, Alric P. Rothmayer, Okiishi H. Theodore, and Wade W. Huebsch. *Fundamentals of Fluid Mechanics*. 7th. Hoboken, NJ: John Wiley & Sons, 2013. ISBN: 978-1-118-11613-5 (cit. on pp. 10, 11).

[20] W. Wagner and A. Pruß. "The IAPWS Formulation 1995 for the Thermodynamic Properties of Ordinary Water Substance for General and Scientific Use". In: *Journal of Physical and Chemical Reference Data* 31.2 (2002), pp. 387–535. DOI: 10.1063/1.1461829 (cit. on pp. 12, 58).

[21] J.J. Duderstadt and L.J. Hamilton. *Nuclear reactor analysis*. John Wiley and Sons, Inc., 1976 (cit. on pp. 12, 19, 20).

[22] G. I. Bell and S. Glasstone. *Nuclear Reactor Theory*. Van Nostrand Reinhold, 1970 (cit. on p. 12).

[23] Carlo Fiorina, Ivor Clifford, Manuele Aufiero, and Konstantin Mikityuk. "GeN-Foam: a novel OpenFOAM® based multi-physics solver for 2D/3D transient analysis of nuclear reactors". In: *Nuclear Engineering and Design* 294 (2015), pp. 24–37. ISSN: 0029-5493. DOI: https://doi.org/10.1016/j.nucengdes.2015.05.035 (cit. on p. 13).

[24] Alessandro Scolaro, Ivor Clifford, Carlo Fiorina, and Andreas Pautz. "The OFFBEAT multi-dimensional fuel behavior solver". In: *Nuclear Engineering and Design* 358 (2020), p. 110416. ISSN: 0029-5493. DOI: https://doi.org/10.1016/j.nucengdes.2019.110416 (cit. on p. 14).

[25] Thomas Guilbaud. "Honeycomb, an effective web tool to visualize and edit lattices". In: *Transactions of the American Nuclear Society - ANS Annual Meeting*. Available at: https://foam-for-nuclear.gitlab.io/honeycomb/. Las Vegas, USA, 2024, pp. 1052–1053. DOI: 10.13182/T130-44800 (cit. on p. 14).

[26] G. Por, M. Berta, and M. Csuvar. "Measurement of the coolant flow rate using correlation of temperature fluctuations". In: *Progress in Nuclear Energy* 43.1 (2003), pp. 281–288. DOI: 10.1016/S0149-1970(03)00032-4 (cit. on pp. 16, 24, 31).

[27] Andrew E. Johnson, Dan Kotlyar, Stefano Terlizzi, and Gavin Ridley. "serpentTools: A Python Package for Expediting Analysis with Serpent". In: *Nuclear Science and Engineering* 194.11 (2020), pp. 1016–1024. DOI: 10.1080/00295639.2020.1723992 (cit. on p. 18).

[28] E. E. Lewis and W. F. Miller. *Computational Methods of Neutron Transport*. Wiley, 1984 (cit. on p. 19).

[29] Yasushi Nauchi and Takanori KAMEYAMA. "Development of Calculation Technique for Iterated Fission Probability and Reactor Kinetic Parameters Using Continuous-Energy Monte Carlo Method". In: *Journal of Nuclear Science and Technology* 47 (Jan. 2012), pp. 977–990. DOI: 10.1080/18811248.2010.9711662 (cit. on p. 19).

[30] Jaakko Leppänen et al. "Serpent–a continuous-energy Monte Carlo reactor physics burnup calculation code". In: *VTT Technical Research Centre of Finland* 4.455 (2013), pp. 2023–09 (cit. on p. 19).

[31] VTT Technical Research Centre of Finland. *Serpent 2 Online User's Guide*. 2024. URL: https://serpent.vtt.fi/docs/user_guide/index.html (cit. on pp. 20, 37).

[32] Jaakko Leppänen, Maria Pusa, and Emil Fridman. "Overview of methodology for spatial homogenization in the Serpent 2 Monte Carlo code". In: *Annals of Nuclear Energy* 96 (2016), pp. 126–136. ISSN: 0306-4549. DOI: https://doi.org/10.1016/j.anucene.2016.06.007. URL: https://www.sciencedirect.com/science/article/pii/S0306454916303899 (cit. on pp. 20, 21).

[33] OpenCFD Ltd. *OpenFOAM v2312 Documentation - cyclicAMI boundary condition*. Access: January 2026. OpenCFD Ltd. 2023. URL: Available%20at:%20https://doc.openfoam.com/2312/tools/processing/boundary-conditions/rtm/derived/coupled/cyclicAMI/ (cit. on p. 26).

[34] Christopher Greenshields. *OpenFOAM v13 User Guide*. London, UK: The OpenFOAM Foundation, 2025. URL: https://doc.cfd.direct/openfoam/user-guide-v13 (cit. on p. 26).

[35] simFlow Technologies. *simFlow Documentation - Cyclic AMI and ACMI Boundary Conditions*. Access: January 2026. 2024. URL: Available%20at:%20https://help.sim-flow.com/boundary-conditions/cyclic-ami-acmi (cit. on p. 26).

[36] B. R. Latimer and A. Pollard. "COMPARISON OF PRESSURE-VELOCITY COUPLING SOLUTION ALGORITHMS". In: *Numerical Heat Transfer* 8.6 (1985), pp. 635–652. DOI: 10.1080/01495728508961876 (cit. on pp. 28, 29).

[37] Jonathan Fahlbeck. "Implementation of an incompressible headLossPressure boundary condition". In: *Proceedings of CFD with OpenSource Software* (2020) (cit. on p. 29).

[38] I. E. Idelchik. *Handbook of Hydraulic Resistance: Coefficients of Local Resistance and of Friction*. Jerusalem: Israel Program for Scientific Translations, 1966 (cit. on pp. 29, 44).

[39] E. N. Sieder and G. E. Tate. "Heat Transfer and Pressure Drop of Liquids in Tubes". In: *Industrial & Engineering Chemistry* 28.12 (1936), pp. 1429–1435. DOI: 10.1021/ie50324a027 (cit. on p. 31).

[40] Stephen Whitaker. “Forced convection heat transfer correlations for flow in pipes, past flat plates, single cylinders, single spheres, and for flow in packed beds and tube bundles”. In: *AIChE Journal* 18.2 (1972), pp. 361–371 (cit. on p. 31).

[41] Peter J. Linstrom and William G. Mallard. “The NIST Chemistry WebBook: A Chemical Data Resource on the Internet†”. In: *Journal of Chemical & Engineering Data* 46 (2001), pp. 1059–1063. DOI: 10.1021/je000236i (cit. on p. 32).

[42] M.B. Chadwick et al. “ENDF/B-VII.1 Nuclear Data for Science and Technology: Cross Sections, Covariances, Fission Product Yields and Decay Data”. In: *Nuclear Data Sheets* 112.12 (2011). Special Issue on ENDF/B-VII.1 Library, pp. 2887–2996. ISSN: 0090-3752. DOI: https://doi.org/10.1016/j.nds.2011.11.002 (cit. on p. 37).

[43] Szabolcs Czifrus, András Szabolcs Ványi, Tamás Belgya, Yann Perin, Mathieu Hursin, and Klemen Ambrožič. *D2.1 - Report describing the experiments performed at each reactor and associated data uncertainties*. Deliverable Grant Agreement n°101163288. Version 2.0. EVEREST Project, Feb. 2026 (cit. on pp. 40, 41).

[44] John R. Lamarsh and Anthony J. Baratta. *Introduction to Nuclear Engineering*. 3rd. Upper Saddle River, NJ: Prentice Hall, 2001. ISBN: 978-0201824988 (cit. on pp. 58, 59).

[45] Ian H. Bell, Jorrit Wronski, Sylvain Quoilin, and Vincent Lemort. “Pure and Pseudo-pure Fluid Thermophysical Property Evaluation and the Open-Source Thermophysical Property Library CoolProp”. In: *Industrial & Engineering Chemistry Research* 53.6 (2014), pp. 2498–2508. DOI: 10.1021/ie4033999 (cit. on p. 58).