
A high-fidelity multiphysics system for neutronic, thermalhydraulic and fuel-performance analysis of Light Water Reactors

Zur Erlangung des akademischen Grades eines
DOKTORS DER INGENIEURWISSENSCHAFTEN (Dr.-Ing.)

von der KIT-Fakultät für Maschinenbau des
Karlsruher Instituts für Technologie (KIT)
angenommene

DISSERTATION

von

Dipl.-Ing. Manuel García
geboren in Argentinien

Tag der mündlichen Prüfung: 18.10.2021

Hauptreferent:
Prof. Dr.-Ing. Robert Stieglitz
Karlsruher Institut für Technologie

Korreferent:
Prof. D.-Sc. Jaakko Leppänen
Aalto University / VTT Technical Research Centre of Finland

*“Aquí me pongo a cantar
Al compás de la vigüela,
Que el hombre que lo desvela
Una pena extraordinaria
Como la ave solitaria
Con el cantar se consuela.”*

El gaucho Martín Fierro, José Hernandez

A Alma y Julián, a mi papá...

Acknowledgements

First and foremost, I would like to thank Professor Robert Stieglitz from the Karlsruhe Institute of Technology and Professor Jaakko Leppänen from Aalto University for agreeing to supervise and evaluate this dissertation. The implementation in Serpent of the domain decomposition algorithm presented in this thesis would not have been possible without the support of Professor Leppänen.

My deepest gratitude goes to my colleagues at the Institute for Neutron Physics and Reactor Technology of the Karlsruhe Institute of Technology, in particular to Uwe Imke and Diego Ferraro, with whom I worked closely throughout this project, as well as to Victor Hugo Sánchez-Espinoza, Luigi Mercatali and Petra Klug. I also wish to thank Ville Valtavirta and Riku Tuominen from the VTT Technical Research Centre of Finland and André Gommlich from the Helmholtz-Zentrum Dresden-Rossendorf for the fruitful collaboration on multiphysics during the McSAFE project.

Last but not least, this thesis would not have been possible without the support of my family: Alma, Julián and Nati, my brothers Lu and Oti and my parents Fabiana and Mauricio. I would also like to thank my professors and colleagues at the Balseiro Institute and the Bariloche Atomic Center in Argentina for my education.

KARLSRUHE INSTITUTE OF TECHNOLOGY

Abstract

Department of Mechanical Engineering

A high-fidelity multiphysics system for neutronic, thermalhydraulic and fuel-performance analysis of Light Water Reactors

by Manuel GARCÍA

The behavior of the core in a Light Water Reactor (LWR) is dominated by neutronic, thermalhydraulic and thermomechanic phenomena, as well as a complex set of feedback mechanisms between these physical domains. Thus, one of the current trends in computational reactor physics is the implementation of multiphysics applications that can capture these interactions to provide a consistent description of the core. Another key line of work is the development of high-fidelity numerical tools that increase the modelling resolution and eliminate strong approximations used in nodal-level solvers. Multiphysics and high-fidelity methods rely on the availability of High Performance Computing (HPC) systems, which determine the feasibility and scope of this type of simulations.

The aim of this thesis is the development of a multiphysics system capable of performing coupled neutronic, thermalhydraulic and fuel-performance analysis of LWR cores using a high-fidelity methodology. To achieve this, the continuous-energy Monte Carlo particle transport method is used to simulate the neutronic behavior without relying on major physical approximations. To handle full-core pin-by-pin burnup calculations, a data decomposition scheme with a domain decomposition particle tracking method is proposed and implemented. Combining Monte Carlo neutronics with subchannel-level thermalhydraulics and full fuel-performance analysis of all fuel rods, an extremely detailed representation of the core is achieved, pushing the computational requirements to the limits of HPC systems. From the software perspective, an innovative object-oriented coupling approach is used to increase the modularity, flexibility and maintainability of the tool.

The accuracy of this three-code system is evaluated using experimental data from two operating power plants, a Pre-Konvoi PWR and the Temelín II VVER-1000 reactor. For these two cases, the results of full-core burnup calculations are validated using critical boron concentration and pin-level neutron flux measurements. These simulations serve to illustrate the cutting-edge modelling capabilities of the developed tool and to assess the feasibility of this methodology for industrial applications.

KARLSRUHER INSTITUT FÜR TECHNOLOGIE

Kurzfassung

Fakultät für Maschinenbau

A high-fidelity multiphysics system for neutronic, thermalhydraulic and fuel-performance analysis of Light Water Reactors

von Manuel GARCÍA

Das Verhalten des Kerns in einem Leichtwasserreaktor (LWR) wird von neutronenphysikalischen, thermohydraulischen und thermomechanischen Phänomenen dominiert. Komplexe Rückkopplungsmechanismen verbinden diese physikalischen Bereiche. Einer der aktuellen Tendenzen in der Reaktorphysik ist daher die Implementierung von Multiphysik-Methoden, die diese Wechselwirkungen erfassen, um eine konsistente Beschreibung des Kerns zu liefern. Ein weiterer wichtiger Arbeitsbereich ist die Entwicklung von High-Fidelity-Rechenprogrammen, die die Modellierungsaufösung erhöhen und starke Vereinfachungen eliminieren, die in räumlich homogenisierten Simulationen verwendet werden. Multiphysik- und High-Fidelity-Methoden sind auf die Verfügbarkeit von Hochleistungsrechnern angewiesen, die die Machbarkeit und den Umfang dieser Art von Simulationen begrenzen.

Das Ziel dieser Arbeit ist die Entwicklung eines Multiphysik-Simulationssystems, das in der Lage ist, gekoppelte neutronenphysikalische, thermohydraulische und thermomechanische Analysen von LWR-Kernen mit einer High-Fidelity-Methodik durchzuführen. Um dies zu erreichen, wird die Monte-Carlo-Teilchen-transportmethode verwendet, um das Verhalten der neutronenphysikalischen Effekte zu simulieren, ohne auf größere physikalische Näherungen zurückzugreifen. Für die Abbrandrechnungen bezüglich des gesamten Kerns, wird eine gebietsbezogene Datenaufteilung der Partikelverfolgung vorgeschlagen und implementiert. Die Kombination der Monte-Carlo-Methode mit der Thermohydraulik auf Unterkanalebene und eine vollständige Analyse des Brennstoffverhaltens aller Brennstäbe beschreibt eine extrem detaillierte Darstellung des Kerns. Die erforderliche Rechenleistung erreicht die Grenzen aktueller Hochleistungsrechner. Auf der Softwareseite wird ein innovativer objektorientierter Kopplungsansatz verwendet, um die Modularität, Flexibilität und Wartbarkeit des Programms zu erhöhen.

Die Genauigkeit dieses gekoppelten Systems von drei Programmen wird mit experimentellen Daten von zwei in Betrieb befindlichen Kraftwerken, einem Pre-Konvoi DWR und dem Temelín II WWER-1000 Reaktor, bewertet. Für diese beiden Fälle werden die Ergebnisse der Abbrandrechnung des gesamten Kerns anhand von Messungen der kritischen Borkonzentration und des Brennstabneutronenflusses validiert. Diese Simulationen dienen der Darstellung der hochmodernen Modellierungsfähigkeiten des entwickelten Werkzeugs und zeigen die Durchführbarkeit dieser Methodik für industrielle Anwendungen.

Publications related to this thesis

As main author:

- M. García, D. Ferraro, V. H. Sanchez-Espinoza, L. Mercatali, J. Leppänen, and V. Valtavirta. “Development of a spatial domain decomposition scheme for Monte Carlo neutron transport”. *ICONE26 Conference, London, UK, July, 2018*, 82144 (2018). DOI: <http://doi.org/10.1115/icone26-82144>.
- M. García, U. Imke, D. Ferraro, V. Sanchez-Espinoza, and L. Mercatali. “Advanced modelling capabilities for pin-Level subchannel analysis of PWR and VVER reactors”. *NURETH-18 Conference, Portland, USA, August, 2019* (2019), pp. 4299–4309. DOI: <http://doi.org/10.5445/ir/1000099942>.
- M. García, D. Ferraro, V. Valtavirta, U. Imke, R. Tuominen, V. Sanchez-Espinoza, and L. Mercatali. “Development of an object-oriented Serpent2-SUBCHANFLOW coupling and verification with Problem 6 of the VERA Core Physics Benchmark”. *M&C 2019 Conference, Portland, USA, August, 2019* (2019), pp. 1424–1433. DOI: <http://doi.org/10.5445/ir/1000099853>.
- M. García, D. Ferraro, V. Valtavirta, R. Tuominen, U. Imke, J. Leppänen, and V. Sanchez-Espinoza. “Serpent2-SUBCHANFLOW pin-by-pin modelling capabilities for VVER geometries”. *Annals of Nuclear Energy* 135, 106955 (2020). DOI: <http://doi.org/10.1016/j.anucene.2019.106955>.
- M. García, R. Tuominen, A. Gommlich, D. Ferraro, V. Valtavirta, U. Imke, P. Van Uffelen, L. Mercatali, V. Sanchez-Espinoza, J. Leppänen, and S. Kliem. “A Serpent2-SUBCHANFLOW-TRANSURANUS coupling for pin-by-pin depletion calculations in Light Water Reactors”. *Annals of Nuclear Energy* 139, 107213 (2020). DOI: <http://doi.org/10.1016/j.anucene.2019.107213>.
- M. García, R. Tuominen, A. Gommlich, D. Ferraro, V. Valtavirta, U. Imke, P. Van Uffelen, L. Mercatali, V. Sanchez-Espinoza, J. Leppänen, and S. Kliem. “Serpent2-SUBCHANFLOW-TRANSURANUS pin-by-pin depletion calculations for a PWR fuel assembly”. *EPJ Web Conference (PHYSOR 2020 Conference)* 247, 06016 (2021). DOI: <http://doi.org/10.1051/epjconf/202124706016>.
- M. García, D. Ferraro, V. Valtavirta, R. Tuominen, U. Imke, L. Mercatali, V. Sanchez-Espinoza, and J. Leppänen. “A subchannel coarsening method for Serpent2-SUBCHANFLOW applied to a full-core VVER problem”. *EPJ Web Conference (PHYSOR 2020 Conference)* 247, 06018 (2021). DOI: <http://doi.org/10.1051/epjconf/202124706018>.
- M. García, J. Leppänen, and V. Sanchez-Espinoza. “A Collision-based Domain Decomposition scheme for large-scale depletion with the Serpent 2 Monte Carlo code”. *Annals of Nuclear Energy* 152, 108026 (2021). DOI: <http://doi.org/10.1016/j.anucene.2020.108026>.
- M. García, Y. Bilodid, J. Basualdo Perello, R. Tuominen, A. Gommlich, J. Leppänen, V. Valtavirta, U. Imke, D. Ferraro, P. Van Uffelen, M. Seidl, and V. Sanchez-Espinoza. “Validation of Serpent-SUBCHANFLOW-TRANSURANUS pin-by-pin burnup calculations using experimental data from a Pre-Konvoi PWR reactor”. *Nuclear Engineering and Design* 379, 111173 (2021). DOI: <http://doi.org/10.1016/j.nucengdes.2021.111173>.

- M. García, R. Vočka, R. Tuominen, A. Gommlich, J. Leppänen, V. Valtavirta, U. Imke, D. Ferraro, P. Van Uffelen, L. Milisdörfer, and V. Sanchez-Espinoza. “Validation of Serpent-SUBCHANFLOW-TRANSURANUS pin-by-pin burnup calculations using experimental data from the Temelín II VVER-1000 reactor”. *Nuclear Engineering and Technology* 53 (2021), pp. 3133–3150. DOI: <http://doi.org/10.1016/j.net.2021.04.023>.

As coauthor:

- D. Ferraro, M. García, L. Mercatali, V. H. Sanchez-Espinoza, J. Leppänen, and V. Valtavirta. “Foreseen capabilities, bottlenecks identification and potential limitations of Serpent MC transport code in large-scale full 3-D burnup calculations”. *ICONE26 Conference, London, UK, July, 2018*, 82305 (2018). DOI: <http://doi.org/10.1115/icone26-82305>.
- D. Ferraro, M. García, U. Imke, V. Valtavirta, J. Leppänen, and V. Sanchez-Espinoza. “Serpent/SCF pin-level multiphysics solutions for the VERA Fuel Assembly Benchmark”. *Annals of Nuclear Energy* 128 (2019), pp. 102–114. DOI: <http://doi.org/10.1016/j.anucene.2018.12.047>.
- D. Ferraro, M. García, U. Imke, V. Valtavirta, R. Tuominen, Y. Bilodid, J. Leppänen, and V. Sanchez-Espinoza. “Serpent/SUBCHANFLOW coupled calculations for a VVER core at hot full power”. *EPJ Web Conference (PHYSOR 2020 Conference)* 247, 04006 (2021). DOI: <http://doi.org/10.1051/epjconf/202124704006>.
- D. Ferraro, M. García, U. Imke, V. Valtavirta, R. Tuominen, J. Leppänen, and V. Sanchez-Espinoza. “Serpent/SUBCHANFLOW coupled burnup calculations for VVER fuel assemblies”. *EPJ Web Conference (PHYSOR 2020 Conference)* 247, 04005 (2021). DOI: <http://doi.org/10.1051/epjconf/202124704005>.
- R. Tuominen, V. Valtavirta, M. García, D. Ferraro, and J. Leppänen. “Effect of energy deposition modelling in coupled steady state Monte Carlo neutronics/thermal hydraulics calculations”. *EPJ Web Conference (PHYSOR 2020 Conference)* 247, 06001 (2021). DOI: <http://doi.org/10.1051/epjconf/202124706001>.

Contents

Acknowledgements	v
Abstract	vii
Kurzfassung	ix
Publications related to this thesis	xi
1 Introduction	1
1.1 Light Water Reactors	1
1.1.1 PWR designs	2
1.1.2 VVER designs	4
1.2 Challenges in core physics analysis	5
1.2.1 Multiphysics simulations	5
1.2.2 High-fidelity methods	6
1.2.3 Validation	7
1.3 Main research objectives of the thesis	7
1.4 Thesis structure	7
2 Core physics analysis of Light Water Reactors	9
2.1 Neutronics	9
2.1.1 The neutron transport equation	9
2.1.2 Boundary and initial conditions	14
2.1.3 Deterministic methods	15
2.1.4 The Monte Carlo method	16
2.1.5 State-of-the-art neutronic analysis	18
2.2 Thermalhydraulics	18
2.2.1 Flow equations	19
2.2.2 Boundary and initial conditions	20
2.2.3 Computational Fluid Dynamics	21
2.2.4 Subchannel analysis	22
2.2.5 State-of-the-art thermalhydraulic analysis	23
2.3 Thermomechanics	24
2.3.1 Heat-transfer and mechanics equations	24
2.3.2 Boundary and initial conditions	25
2.3.3 Fuel-performance analysis	25
2.3.4 State-of-the-art fuel-performance analysis	28
2.4 Multiphysics	28
2.4.1 Thermalhydraulic and thermomechanic effects on neutronics	28
2.4.2 The power distribution	30
2.4.3 Neutronic effects on fuel behavior	30
2.4.4 Boundary conditions at the cladding-coolant interface	30
2.4.5 Boundary conditions for the full system	31

2.4.6	State-of-the-art multiphysics analysis	32
3	Coupling neutronics, thermalhydraulics and fuel-performance analysis	33
3.1	Selected calculation codes	33
3.1.1	Serpent	33
3.1.2	SUBCHANFLOW	34
3.1.3	TRANSURANUS	35
3.2	Calculation scheme	36
3.2.1	The operator-splitting method	36
3.2.2	Coupled depletion methodology	37
3.2.3	Additional remarks on the coupling scheme	38
3.3	Software design	39
3.3.1	Traditional coupling approaches	39
3.3.2	Motivation for a novel approach	40
3.3.3	Object-oriented design	41
3.3.4	Mesh-based feedback exchange	43
4	Optimizing the system for full-core pin-by-pin depletion problems	45
4.1	Parallel schemes for Monte Carlo particle transport	45
4.1.1	Data decomposition	46
4.1.2	Domain decomposition	46
4.2	Collision-based Domain Decomposition	47
4.2.1	Material decomposition	48
4.2.2	Particle tracking	49
4.2.3	Termination control	50
4.3	Performance	51
4.3.1	Serpent	52
4.3.2	SUBCHANFLOW	53
4.3.3	TRANSURANUS	53
5	Validation for a Pre-Konvoi PWR reactor	55
5.1	Pre-Konvoi PWR validation data	55
5.1.1	Core description	55
5.1.2	Operational data	56
5.1.3	Experimental data used for validation	58
5.2	Modelling approach	59
5.3	Selected results	61
5.3.1	Critical boron concentration	61
5.3.2	Neutron flux profiles	62
5.3.3	Fuel solution method	64
5.4	Validation summary	67
6	Validation for the Temelín II VVER-1000 reactor	69
6.1	Temelín II VVER-1000 validation data	69
6.1.1	Core description	69
6.1.2	Operational data	71
6.1.3	Experimental data used for validation	72
6.2	Modelling approach	73
6.3	Selected results	74
6.3.1	Critical boron concentration	74
6.3.2	Neutron flux profiles	75

6.3.3	Power distribution	76
6.3.4	Pin-level results	76
6.4	Validation summary	79
7	Summary and outlook	85
7.1	Summary	86
7.2	Outlook	86
	Bibliography	89
A	Mesh-based feedback exchange for a VVER fuel assembly	97
A.1	Serpent	97
A.2	SUBCHANFLOW	98
A.3	TRANSURANUS	100
B	Performance analysis for the CDD method	101
B.1	Speedup	102
B.1.1	Multithreading	102
B.1.2	Burnup	102
B.1.3	Tracking modes	103
B.2	Memory scalability	103
B.2.1	Memory optimization modes	103
B.2.2	Burnup	105
C	Pre-Konvoi PWR core analysis model	107
C.1	Serpent model	107
C.2	SCF model	109
C.3	TU model	110
D	Temelín II VVER-1000 core analysis model	113
D.1	Serpent model	113
D.2	SCF model	115
D.3	TU model	116

List of Acronyms

LWR	Light Water Reactor
NPP	Nuclear Power Plant
HPC	High Performance Computing
PWR	Pressurized Water Reactor
BWR	Boiling Water Reactor
SCWR	Supercritical Water Reactor
RPV	Reactor Pressure Vessel
MOX	Mixed U and Pu Oxide
VVER	Water-Water Energetic Reactor
DNB	Departure from Nucleate Boiling
MSR	Molten Salt Reactor
FDM	Finite Difference Method
FVM	Finite Volume Method
FEM	Finite Element Method
CPM	Collision Probabilities Method
MOC	Method of Characteristics
CFD	Computational Fluid Dynamics
DNS	Direct Numerical Simulation
LES	Large Eddy Simulation
RANS	Reynolds-averaged Navier-Stokes
FGR	Fission Gas Release
PCMI	Pellet-cladding Mechanical Interaction
SCC	Stress Corrosion Cracking
CSG	Constructive Solid Geometry
STL	Stereolithography
TTA	Transmutation Trajectory Analysis
CRAM	Chebyshev Rational Approximation Method
SIE	Stochastic Implicit Euler
EOS	End of Step
BOS	Beginning of Step
API	Application Programming Interface
ICoCo	Interface for Code Coupling
CDD	Collision-based Domain Decomposition
HFP	Hot Full Power
EFPD	Effective Full Power Days
AMS	Aeroball Measurement System
RMS	Root Mean Square
RIA	Reactivity-initiated Accident
CMS	Core Monitoring System
SPND	Self Powered Neutron Detector
BOC	Beginning of Cycle
EOC	End of Cycle

DNBR	Departure from Nucleate Boiling Ratio
CHF	Critical Heat Flux
FET	Functional Expansion Tally

Chapter 1

Introduction

Light Water Reactors (LWRs) are the most common class of nuclear reactors worldwide, accounting in 2018 for about 89.2% of the global nuclear power generation [1]. Hence, being able to accurately predict the behavior of LWR cores has a direct impact on the nuclear industry, both in terms of design optimization and of safety assessment of current and future Nuclear Power Plants (NPPs). This creates an incentive to continuously improve the numerical tools used to perform core physics analysis of LWRs.

Currently, at least two trends exist in computational reactor physics. On the one hand, phenomena that have traditionally been treated separately such as neutronics, thermalhydraulics and fuel thermomechanics are being solved as coupled multiphysics systems to account for relevant feedback mechanisms. While this has been an important research topic for many decades, coupling techniques are continuously being improved in terms of numerical methods and software technologies. On the other hand, the resolution of the models, spatially or otherwise, is being pushed to the limits of the computational capacity, in what is usually referred to as high-fidelity methods. While this does not necessarily mean that the physical models are improved, the ability to avoid approximations such as spatial homogenization and energy condensation often does lead to a better physical description of the reactor core. Both these lines of work, multiphysics and high fidelity, rely heavily on cutting-edge computational resources, i.e. High Performance Computing (HPC) systems, which determine the computing power that limits the accuracy and complexity of such simulations.

This thesis is an attempt to push forward the development of HPC-based high-fidelity multiphysics tools for LWR analysis. The main focus is the implementation of a neutronic-thermalhydraulic-thermomechanic coupling system optimized for massively parallel pin-level burnup simulations. This tool enables the solution of fully coupled depletion problems with a geometrical and physical resolution previously not achieved. To close the software development cycle, from implementation to validation, the accuracy of this system is assessed using experimental data from two operating nuclear reactors.

1.1 Light Water Reactors

LWRs are nuclear reactors that use regular water, as opposed to heavy water, as both coolant and moderator, operating with a thermal neutron spectrum. The power is generated by fission of uranium and plutonium contained in fuel rods and converted into electricity through a thermal cycle that drives a turbine system.

This work focuses on Pressurized Water Reactors (PWRs), a type of LWR, which constitute by themselves about 72.2% of the electricity production and 85.1% of the installed capacity under construction in the global nuclear sector [1]. Other

LWR types, namely Boiling Water Reactors (BWRs) and Supercritical Water Reactors (SCWRs), a Generation IV conceptual design, are not considered directly in this thesis, although most of the discussions can be readily extended to these systems.

1.1.1 PWR designs

The primary cooling system of a three-loop PWR is shown in Figure 1.1. The power generated in the reactor core, where water flows upward and is heated from about 275 °C to 315 °C, is transferred through steam generators to a secondary system that feeds the turbines. The coolant flow is maintained by pumps located in the cold legs of each loop and a pressurizer keeps the pressure at about 15.5 MPa to prevent the water from boiling.




FIGURE 1.1: Three-loop PWR primary cooling system [2].

The Reactor Pressure Vessel (RPV), shown in Figure 1.2, hosts the reactor core, as well as the control rod system and the structures that direct the coolant flow. The

cold water coming from the primary cooling pumps enters the RPV via the inlet nozzles, flows downward through the downcomer and goes into the core. After circulating upward through the core, the hot water leaves the RPV via the outlet nozzles to the steam generators.




FIGURE 1.2: PWR reactor pressure vessel [2].

The core is composed of fuel assemblies like the one shown in Figure 1.3. Each fuel assembly consists of a square array of fuel rods kept together by a set of spacer grids and other structural components. Each fuel rod contains a stack of UO_2 or MOX (mixed U and Pu oxide) pellets hosted inside a cladding tube. In between the fuel rods, guide tubes permit the insertion of the control rods used to regulate the reactivity of the core in conjunction with boric acid dissolved in the coolant.




FIGURE 1.3: PWR fuel assembly [2].

1.1.2 VVER designs

Water-Water Energetic Reactors (VVERs, from Russian) are a specific type of PWR first developed in the Soviet Union. While they share most of the general characteristics of PWRs, the main difference in terms of core analysis is that the core geometry is hexagonal instead of rectangular, as can be observed in the fuel assembly shown in Figure 1.4. This increases the geometrical complexity needed for core physics analysis, as will be clear in further discussions. To make this distinction, in this work the term PWR usually refers to rectangular-type PWR designs, whereas

VVER is used to refer specifically to these hexagonal-type systems.




FIGURE 1.4: VVER fuel assembly [3].

1.2 Challenges in core physics analysis

The focus of this thesis is on the phenomena occurring inside the reactor core, namely the power generation by fission, the cooling conditions in the active length of the core and the behavior of the fuel and cladding properties during irradiation. Other issues such as the dynamic of the primary and secondary cooling systems or the analysis of accidental conditions are not considered. Moreover, this work focuses on steady-state and burnup simulations, and transient scenarios are out of the scope as well.

Core physics analysis of LWRs has traditionally relied on two main research areas: neutronics and thermalhydraulics. The former is focused on modelling the behavior of neutrons produced by fission as they stream through the fuel, moderator and structural materials in order to calculate global and local neutronic parameters, e.g. reactivity, power distribution and isotopic concentrations during burnup. The latter deals with the flow of water through the core, which determines the cooling conditions and safety parameters, e.g. Departure from Nucleate Boiling (DNB) and limiting fuel and cladding temperatures. Another key set of phenomena is fuel thermomechanics, usually referred to as fuel-performance analysis, which deals with the behavior of the fuel and cladding materials during irradiation.

1.2.1 Multiphysics simulations

Since the thermalhydraulic calculation depends directly on the power distribution in the core, and the resulting temperatures and densities affect the neutronic calculation through the material cross-sections, these two problems need to be treated as a coupled system to some degree. The simplest approach is to solve each physical problem separately while taking the other as a boundary condition or a source term. However, given the strong interaction between these two systems, the standard approach nowadays is to solve both simultaneously taking into account the main feedback mechanisms.

Fuel-performance analysis describes the changes in the composition and thermophysical properties of the fuel and cladding materials during their irradiation as the power plant operates. It includes a wide variety of processes that affect the mechanical behavior of the fuel and the cladding, e.g. thermal and irradiation-induced densification of the fuel, swelling due to solid and gaseous fission products, creep,

plasticity, pellet cracking and relocation, oxygen and plutonium redistribution and volume changes during phase transitions. This set of mechanisms influences key fuel rod parameters such as the temperature distribution and the fuel-cladding gap evolution, and constitutes a multiphysics system in and of itself. While these phenomena impact the neutronic calculation, in particular through the fuel temperature, and depend heavily on it, for example through the power and neutron flux distributions, the usual approach is to perform this type of analysis taking a fixed solution for the neutronics. The same is true for the interaction with the thermal-hydraulic state, which determines the boundary conditions for the fuel rods: fixed boundary conditions are used for the thermomechanic calculation, and a simplified fuel-performance model is used to get the fuel temperatures.

Now, given the strong interaction between neutronics, thermalhydraulics and thermomechanics, it is desirable to solve the full multiphysics system within a coupled scheme. Doing this, the feedback mechanisms between the three physics can be taken into account to obtain a core state that is consistent across all relevant processes. This thesis deals with the development of such a coupling system, both in terms of the feedback and calculation scheme and of the actual implementation as a software package.

1.2.2 High-fidelity methods

As a result of the problem complexity and the computational limits, traditional methods used to analyze LWR cores rely on a series of physical and numerical approximations to obtain results in a practical time. The neutronic simulation is approached using a three-step multiscale scheme: nuclear data processing to generate problem-dependent multigroup libraries, lattice-level calculations to condense the pin-level geometry into larger nodes with homogeneous cross sections and core-level simulations to get nodal-level solutions. Thermalhydraulic calculations are usually performed at fuel-assembly level, and pin-level safety parameters are obtained using methods such as pin-power reconstruction. Fuel-performance analysis is carried out independently and normally either at nodal level or only for selected fuel rods.

The accuracy of this methodology can in principle be improved by eliminating as many approximations as possible and increasing the resolution of the reactor model, i.e. using a high-fidelity approach, at the expense of increasing the computational requirements. For a LWR core, this means solving the full-core pin-by-pin problem to obtain directly the general solution and the main safety parameters.

In this thesis, the neutronic calculations are tackled with the continuous-energy Monte Carlo particle transport method, which can handle arbitrarily complex geometries without any inherent approximations on the angular and energy dependence of the neutron flux. The physical description of the system is thus greatly improved by avoiding the spatial homogenization and energy condensation methods used in the traditional multiscale approach. Combining Monte Carlo neutronics with subchannel-level thermalhydraulics and full fuel-performance analysis of all fuel rods, an extremely detailed representation of the core is achieved. In this regard, although the geometrical resolution of the model is substantially increased with respect to traditional methods, the physical description of the thermalhydraulic and thermomechanic domains is not necessarily improved.

1.2.3 Validation

Any reactor physics code is not fully reliable as a real-life analysis tool for design optimization or safety evaluation until its prediction accuracy is validated using experimental data from operating power plants. This process is not general, in the sense that a code needs to be validated for a particular reactor type, using a certain kind of model and for a specific range of problems. Under these conditions, the deviation between the numerical results and the experimental data needs to be quantified to assess the accuracy of the tool and determine its suitability for a given application.

In this thesis, the implemented software package is validated for full-core pin-by-pin depletion calculations for LWR designs using high-fidelity models. For this task, critical boron and neutron flux measurements from Pre-Konvoi PWR and VVER-1000 reactors are used.

1.3 Main research objectives of the thesis

The aim of this thesis is the development of a LWR analysis tool coupling Monte Carlo neutronics with thermohydraulic and thermomechanic analysis capable of performing full-core pin-by-pin depletion calculations. The full development cycle is composed of design, implementation, optimization and validation phases, all of which are presented in this work. The main objectives of the thesis can be summarized as:

1. Develop a multiphysics system combining continuous-energy Monte Carlo neutron transport with high-resolution subchannel and fuel-performance analysis using an adequate calculation scheme and software design.
2. Optimize the Monte Carlo calculation in terms of memory demand and parallel efficiency to achieve full-core pin-by-pin burnup capabilities.
3. Validate the computational tool, in particular the optimized Monte Carlo depletion scheme, using experimental data from PWR and VVER power plants.
4. Evaluate the proposed methodology in terms of accuracy improvements and performance drawbacks and determine its viability as a calculation tool in a practical context.

1.4 Thesis structure

This thesis is organized as follows. Chapter 2 presents the problem of core physics analysis, where the methods used for neutronic, thermohydraulic and thermomechanic analysis of LWRs are described, along with an introduction of the relevant multiphysics phenomena and a discussion of the state of the art. The coupling methodology proposed in this work is outlined in Chapter 3, where the numerical scheme and the software design are described. Chapter 4 discusses the optimization of the tool for large-scale depletion problems, which is largely focused on the implementation of a domain decomposition algorithm for Monte Carlo neutron transport. Validation calculations for PWR and VVER reactors are presented in chapters 5 and 6, respectively. Finally, Chapter 7 deals with the final conclusions, summary and outlook, as well as recommendations for further work.

Chapter 2

Core physics analysis of Light Water Reactors

The reactor core is arguably the most important part of a nuclear power plant, since this is where the energy is generated and where some of the most critical phenomena for design optimization and safety analysis occur. Even when considering the behavior of the plant as a whole, for example during transient conditions or severe accident scenarios, the core drives the dynamics of the system to a great extent. A detailed and consistent description of the core is therefore essential for the study of Light Water Reactors.

Core physics analysis is based on three main areas of research: neutronics, thermohydraulics and thermomechanics, which deal with the fission power generation, the cooling conditions and the fuel evolution during irradiation, respectively. Sections 2.1 to 2.3 of this chapter present the basic principles and analysis methods for each of these fields. The most relevant feedback mechanisms between neutronic, thermohydraulic and thermomechanic phenomena, which are the motivation for multiphysics simulations, are introduced in Section 2.4.

2.1 Neutronics

In a reactor core, neutrons produced by fission stream through the system undergoing elastic and inelastic scattering as well as capture reactions, which in fissile nuclei have a probability of causing fission events emitting new neutrons that sustain the chain reaction. The neutron distribution determines the power generation by fission and the evolution of the fuel during burnup, along with the degradation and activation of structural materials through irradiation. Hence, one of the central questions in core physics analysis is the modelling of the neutron population in terms of its spatial, angular, energetic and temporal dependency.

2.1.1 The neutron transport equation

The behavior of neutrons in a multiplicative system is characterized by the general form of the neutron transport equation. This is the fundamental formulation of the neutron transport problem on which computational methods are based. Its stationary formulation, which also introduces the concept of criticality, is the basis for steady-state and depletion calculations, which are the focus of this work.

2.1.1.1 Fundamental variables

The motion of neutrons in a system can be described in terms of four independent variables, namely the neutron position \vec{r} , energy E and direction $\hat{\Omega}$, and the

time t . The position \vec{r} corresponds to a vector in the three-dimensional space, while the direction of flight $\hat{\Omega}$ is a unit vector typically specified in terms of the azimuthal angle ω and the polar angle θ . An integration over the unit sphere for an arbitrary function $f(\hat{\Omega})$ is denoted by

$$\int_{4\pi} f(\hat{\Omega})d\Omega = \int_0^{2\pi} \int_{-1}^1 f(\omega, \mu)d\mu d\omega, \quad (2.1)$$

where $\mu = \cos \theta$.

The neutron population is described by the angular neutron density $N(\vec{r}, E, \hat{\Omega}, t)$, defined such that the number of neutrons $n(t)$ in a volume V with energies in ΔE flying with directions within $\Delta\Omega$ is

$$n(t) = \int_V \int_{\Delta E} \int_{\Delta\Omega} N(\vec{r}, E, \hat{\Omega}, t)d\Omega dE dV. \quad (2.2)$$

The angular flux $\psi(\vec{r}, E, \hat{\Omega}, t)$ is defined from the angular density as

$$\psi(\vec{r}, E, \hat{\Omega}, t) = vN(\vec{r}, E, \hat{\Omega}, t), \quad (2.3)$$

where $v = \sqrt{2E/m}$ is the neutron speed (m is the neutron mass). This is the main variable used to describe neutrons in reactor physics analysis, since it represents the rate of fluence of neutrons and can be used to calculate reaction rates. Finally, the scalar flux $\phi(\vec{r}, E, t)$ is defined as

$$\phi(\vec{r}, E, t) = \int_{4\pi} \psi(\vec{r}, E, \hat{\Omega}, t)d\Omega. \quad (2.4)$$

2.1.1.2 Cross sections

The interaction of neutrons with nuclei is characterized by cross sections. The probability dp that a neutron undergoes a collision within an incremental distance ds is

$$dp = \Sigma_t(\vec{r}, E, t)ds, \quad (2.5)$$

with the macroscopic total cross section $\Sigma_t(\vec{r}, E, t)$ defined as

$$\Sigma_t(\vec{r}, E, t) = N(\vec{r}, t)\sigma_t(E), \quad (2.6)$$

where $N(\vec{r}, t)$ is the atomic density and $\sigma_t(E)$ is the microscopic cross section, which depends on the energy E only.

The fission and scattering probabilities are characterized in a similar way by the macroscopic cross sections $\Sigma_f(\vec{r}, E, t)$ and $\Sigma_s(\vec{r}, E, t)$, defined as

$$\Sigma_f(\vec{r}, E, t) = N(\vec{r}, t)\sigma_f(E), \quad (2.7)$$

and

$$\Sigma_s(\vec{r}, E, t) = N(\vec{r}, t)\sigma_s(E), \quad (2.8)$$

where $\sigma_f(E)$ and $\sigma_s(E)$ are the respective microscopic cross sections.

The change in the neutron energy and direction after a collision is given by the macroscopic differential scattering cross section $\Sigma_s(\vec{r}, E' \rightarrow E, \hat{\Omega}' \rightarrow \hat{\Omega}, t)$, which can be written as

$$\Sigma_s(\vec{r}, E' \rightarrow E, \hat{\Omega}' \rightarrow \hat{\Omega}, t) = \Sigma_s(\vec{r}, E, t)p(E' \rightarrow E, \hat{\Omega}' \rightarrow \hat{\Omega}), \quad (2.9)$$

where $p(E' \rightarrow E, \hat{\Omega}' \rightarrow \hat{\Omega})$ is the probability of a change in energy and direction from $(E', \hat{\Omega}')$ to $(E, \hat{\Omega})$, which can be calculated from the kinematics of the scattering reaction. The differential scattering cross section satisfies

$$\int_0^\infty \int_{4\pi} \Sigma_s(\vec{r}, E' \rightarrow E, \hat{\Omega}' \rightarrow \hat{\Omega}, t) d\Omega' dE' = \Sigma_s(\vec{r}, E, t). \quad (2.10)$$

The dependency of cross sections with \vec{r} and t is determined by the isotope density at each location and its evolution due to burnup, thermal expansion and flow conditions, among other factors. The energy spectrum is an intrinsic characteristic of the microscopic cross sections of each isotope. For materials with more than one isotope, which is typically the case, the effective cross sections are a sum of the cross sections for all the isotopes in its composition.

2.1.1.3 Fission neutrons

A neutron-induced fission reaction consists of an incident neutron being absorbed by a heavy actinide nucleus and causing it to split into smaller nuclei, or fission products, emitting neutrons and gamma rays and releasing a large amount of energy. A fission event triggered by a neutron of energy E emits on average $\nu(E)$ neutrons, and fission products are produced with a probability distribution, or fission yield γ , which depends on the fissile isotope. Fission can also occur spontaneously as a decay process, though this is very rare and is only taken into account in specific applications.

Most fission neutrons are released as part of the reaction itself with an energy spectrum $\chi_p(E)$. These are called prompt neutrons, since they appear instantaneously after a fission event. In addition, some of the fission products are unstable and have a probability of producing neutrons by decay. These nuclei are referred to as neutron precursors, and are usually not treated individually, but condensed into N precursor groups with effective parameters. Each precursor group i is characterized by its concentration $C_i(\vec{r}, t)$, decay rate λ_i and neutron fission spectrum $\chi_i(E)$. Neutrons are emitted by precursor nuclei at some point after the fission event, so they are called delayed neutrons. Both prompt and delayed neutrons are emitted isotropically. The fraction of fission neutrons produced by each precursor group is $\beta_i(E)$, and the total delayed fraction $\beta(E)$ is

$$\beta(E) = \sum_{i=1}^N \beta_i(E). \quad (2.11)$$

2.1.1.4 The general form of the transport equation

The motion of neutrons and their interaction with matter are described by the neutron transport equation, also called the linear Boltzmann equation, which can be

written as

$$\begin{aligned}
& \left(\frac{1}{v(E)} \frac{\partial}{\partial t} + \hat{\Omega} \cdot \nabla + \Sigma_t(\vec{r}, E, t) \right) \psi(\vec{r}, E, \hat{\Omega}, t) \\
&= \int_0^\infty \int_{4\pi} \Sigma_s(\vec{r}, E' \rightarrow E, \hat{\Omega}' \rightarrow \hat{\Omega}, t) \psi(\vec{r}, E', \hat{\Omega}', t) d\Omega' dE' \\
&+ \frac{\chi_p(E)}{4\pi} \int_0^\infty \int_{4\pi} [1 - \beta(E')] v(E') \Sigma_f(\vec{r}, E', t) \psi(\vec{r}, E, \hat{\Omega}, t) d\Omega' dE' \\
&+ \sum_{i=1}^N \frac{\chi_i(E)}{4\pi} \lambda_i C_i(\vec{r}, t) + q(\vec{r}, E, \hat{\Omega}, t), \quad (2.12)
\end{aligned}$$

which is a balance between the loss and gain rates of neutrons in the phase space. For a given point $(\vec{r}, E, \hat{\Omega}, t)$ neutrons are lost due to leakage and collision and are produced by in-scattering, prompt neutron emission and precursor decay. The external source $q(\vec{r}, E, \hat{\Omega}, t)$, which contributes to the gain rate as well, is defined such that the rate $Q(t)$ at which neutrons are produced in a volume V with energies in ΔE and directions within $\Delta\Omega$ is

$$Q(t) = \int_V \int_{\Delta E} \int_{\Delta\Omega} q(\vec{r}, E, \hat{\Omega}, t) d\Omega dE dV. \quad (2.13)$$

The evolution of each precursor density is given by

$$\left(\frac{\partial}{\partial t} + \lambda_i \right) C_i(\vec{r}, t) = \int_0^\infty \int_{4\pi} \beta_i(E') v(E') \Sigma_f(\vec{r}, E', t) \psi(\vec{r}, E, \hat{\Omega}, t) d\Omega' dE', \quad (2.14)$$

which puts the change of each concentration $C_i(\vec{r}, t)$ in terms of the production rate due to fission and the loss rate due to radioactive decay. A complete derivation of equations 2.12 and 2.14 can be found in most reactor physics textbooks [4], [5], [6].

The linear Boltzmann equation describes the interaction of neutrons with matter extremely well for nuclear reactor applications, and in particular for LWR cores, where the assumptions on which it relies are valid. The original non-linear equation includes interactions between neutrons, which are rightly neglected in the linear formulation, since neutrons collide with nuclei at rates orders of magnitude larger than they do with other neutrons. Other sources on non-linearity such as changes in cross sections due to burnup and Doppler broadening are addressed indirectly, as explained later on. Another approximation is the fact that delayed neutrons are emitted from precursor nuclei at the fission location, which is reasonable in LWR systems, where precursor migration can be neglected, but needs to be taken into account in liquid fuel reactors, e.g. Molten Salt Reactor (MSR) designs. Finally, the equation is formulated for the mean neutron population in the phase space, even though in a real system the neutron motion is characterized by a stochastic variability. This is accurate for a reactor at normal operation, but not necessarily at very small power, where a stochastic description is needed.

2.1.1.5 The steady-state transport equation and the multiplication factor

In steady-state problems, where the precursor densities are constant, and in the absence of external sources, i.e. $q(\vec{r}, E, \hat{\Omega}) = 0$, equation 2.12 becomes

$$\begin{aligned}
 & (\hat{\Omega} \cdot \nabla + \Sigma_t(\vec{r}, E)) \psi(\vec{r}, E, \hat{\Omega}) \\
 &= \int_0^\infty \int_{4\pi} \Sigma_s(\vec{r}, E' \rightarrow E, \hat{\Omega}' \rightarrow \hat{\Omega}) \psi(\vec{r}, E', \hat{\Omega}') d\Omega' dE' \\
 &+ \frac{\chi_p(E)}{4\pi k} \int_0^\infty \int_{4\pi} [1 - \beta(E')] \nu(E') \Sigma_f(\vec{r}, E') \psi(\vec{r}, E, \hat{\Omega}) d\Omega' dE' \\
 &+ \frac{1}{4\pi k} \sum_{i=1}^N \chi_i(E) \int_0^\infty \int_{4\pi} \beta_i(E') \nu(E') \Sigma_f(\vec{r}, E') \psi(\vec{r}, E, \hat{\Omega}) d\Omega' dE'. \quad (2.15)
 \end{aligned}$$

This equation introduces the multiplication factor k of the system, defined as the largest eigenvalue such that a non-trivial steady-state solution exists. The reason for defining this factor is that, in order for a stationary solution to exist, the rate at which neutrons are produced by fission has to be exactly balanced by the loss rate due to capture and leakage, which is not necessarily the case in a real system. To force the existence of a solution, the fission source is divided by k to ensure this balance. A multiplication factor $k < 1$ means that the fission rate must be increased to force a steady state and therefore that the real system is subcritical, i.e. the loss rate is greater than the production rate and the neutron population would tend to decrease over time. In the opposite case, when $k > 1$, the system is said to be supercritical and if $k = 1$ the reactor is critical. The multiplication factor is one of the most important global parameters for any type of reactor physics calculation, since it characterizes the overall reactivity of the core.

2.1.1.6 Depletion calculations

During the operation of a nuclear reactor, fuel material is consumed by fission and created due to breeding, when fertile nuclei are transmuted into fissile isotopes through neutron absorption, for example when ^{239}Pu is bred from ^{238}U . Fission reactions create a broad range of fission products, some of which, for instance ^{135}Xe and ^{149}Sm , have strong neutron absorption cross sections and poison the reactor. Most fission products are unstable and have complex decay chains that determine the evolution of the fuel composition in the long term, where transuranic elements produced by neutron absorption such as Pu, Cm, Am and Np also play a major role. Moreover, the strong neutron flux causes activation of structural components, as well as depletion of control rods and burnable poisons, which directly affect the reactivity of the core.

Depletion calculations consist on the solution of the neutron distribution taking into account the evolution of the material compositions due to irradiation. This process is transient and non-linear, because the neutron flux determines the changes in the atomic densities $N(\vec{r}, t)$, which in turn affect the macroscopic cross sections which enter into the transport equation used to calculate the neutron flux and the reactivity.

Now, since a real reactor operates at steady state most of the time and the changes in compositions take place on time scales of hours or days, the neutron distribution can be obtained from the steady-state transport equation. Any transient variations, which are dominated by prompt and delayed neutron lifetimes of the order of 10^{-5} s and 10 s, can be averaged out. This means that burnup calculations can be treated

as a series of steady-state transport calculations to obtain the neutron flux and reaction rates while performing a transient analysis of the changes in isotopic densities between critical states.

The evolution of material compositions are described by a system of Bateman equations which put the rate of change of each isotope in terms of their production and destruction due to neutron-induced reactions and radioactive decay [7]. The balance equation for the atomic density of isotope j can be written as

$$\frac{dN_j}{dt} = \sum_{k \neq j} [(\gamma_{k \rightarrow j} \sigma_{f,k} \phi + \lambda_{k \rightarrow j} + \sigma_{k \rightarrow j} \phi) N_k] - (\lambda_j + \sigma_j \phi) N_j, \quad (2.16)$$

where $\gamma_{k \rightarrow j}$ is the fractional fission product yield of isotope j for the fission of isotope k , $\sigma_{f,k}$ is the microscopic fission cross section of isotope k , ϕ is the total scalar flux, $\lambda_{k \rightarrow j}$ is the $k \rightarrow j$ decay constant, $\sigma_{k \rightarrow j}$ is the $k \rightarrow j$ transmutation cross section, λ_j is the total decay constant and σ_j is the total microscopic transmutation cross section. In these equations, the scalar flux ϕ , as well as the nuclear data $\gamma_{k \rightarrow j}$, $\sigma_{f,k}$, $\sigma_{k \rightarrow j}$ and σ_j , are averaged across the energy spectrum and homogenized in a region in which N_j is taken as uniform. The system of Bateman equations for all isotopes in the system can be written in matrix form as

$$\frac{d\vec{n}}{dt} = \mathbf{A}\vec{n}, \quad (2.17)$$

where \vec{n} is the composition vector and \mathbf{A} is the coefficient matrix.

Depletion calculations are of major importance in the core design process, since they provide a description of the reactor through its lifetime and determine key design parameters relevant for its economic analysis such as cycle length, refueling requirements and extraction burnup. Furthermore, the changes in the dynamics of the core during the operating life, which affect the stability and control of the reactor, are critical for safety analysis.

2.1.2 Boundary and initial conditions

In the most general case, $\psi(\vec{r}, E, \hat{\Omega}, t)$ and $C_i(\vec{r}, t)$ can be obtained solving equations 2.12 and 2.14 together. The boundary conditions must prescribe the angular flux $\psi_b(\vec{r}, E, \hat{\Omega}, t)$ for incoming directions as

$$\psi(\vec{r}, E, \hat{\Omega}, t) = \psi_b(\vec{r}, E, \hat{\Omega}, t), \quad \vec{r} \in \partial V, E > 0, \hat{\Omega} \cdot \hat{n} < 0, t > 0, \quad (2.18)$$

where ∂V is the boundary of the system and \hat{n} is the normal unity vector at ∂V . For the steady-state transport equation 2.15 a time-independent form of this boundary condition applies. Full-core models are typically closed with vacuum, i.e. zero incoming current, conditions at the outer surface of the axial and radial reflectors that surround the core. Periodic and reflective boundary conditions are usually used to exploit symmetries in the geometry, though this is not done in this work.

The initial conditions are

$$\psi(\vec{r}, E, \hat{\Omega}, 0) = \psi_0(\vec{r}, E, \hat{\Omega}), \quad \vec{r} \in V, E > 0, \hat{\Omega} \in 4\pi \quad (2.19)$$

and

$$C_i(\vec{r}, 0) = C_{i,0}(\vec{r}), \quad \vec{r} \in V, \quad (2.20)$$

where the angular flux $\psi_0(\vec{r}, E, \hat{\Omega})$ and the precursor distributions $C_{i,0}(\vec{r})$ for $t = 0$ are given. In the case of this thesis, which deals exclusively with steady-state

problems, initial conditions for the flux and the precursors are not relevant and are not discussed further. Depletion problems require the fresh material compositions as initial conditions to simulate the evolution of the isotopic concentrations.

2.1.3 Deterministic methods

Due to the complex dependency of the neutron flux with the four independent variables \vec{r} , E , $\hat{\Omega}$ and t in a reactor core, all practical applications in neutron physics rely on computational codes. Most methods used to solve equation 2.12 numerically are based on applying approximations and discretization schemes over the four independent variables. These are called deterministic methods, as opposed to stochastic methods, which are presented in the next section.

A deterministic method is defined by its specific numerical treatment of each independent variable, which leads to a system of algebraic equations that needs to be solved numerically to obtain the neutron flux. Some common approaches to derive deterministic methods are:

- **Position \vec{r} .** Codes based on the Finite Difference Method (FDM), Finite Volume Method (FVM) or Finite Element Method (FEM), as well as nodal methods, describe the system geometry as a mesh and apply a discretization scheme to approximate the solution. Other approaches such as the Collision Probabilities Method (CPM) and the Method of Characteristics (MOC) define regions where fluxes are taken as constant and use numerical schemes based on ray-tracing techniques.
- **Energy E .** Virtually all deterministic schemes rely on the multigroup method, where the energy dependency of the neutron flux and the reaction rates is discretized into energy groups. Each group is characterized by effective cross sections, fission spectra and kinetic parameters, which are condensed from continuous data trying to conserve reaction and production rates.
- **Direction $\hat{\Omega}$.** S_N methods are based on discretizing the angular variable into discrete ordinates with weighting quadratures, while P_N methods arise from the functional expansion of the flux using spherical harmonics. The diffusion method, which is one of the most popular schemes for coarse-mesh calculations, is derived applying Fick's diffusion law to the P_1 approximation. Other methods such as CPM and MOC are derived from the integral form of the transport equation [4].
- **Time t .** The temporal evolution can be discretized using any time integration scheme, for example forward or backward Euler or Runge-Kutta methods, depending on the efficiency and stability requirements for a specific application.

The range of applicability of a particular method is given by the level of approximation in relation to the physical problem, as well as the numerical efficiency and the actual implementation, in particular the capability to utilize parallel computer architectures.

Now, given the complexity of the problem in the 7-dimensional phase space, the simulation of the transport equation without approximations by means of a deterministic method is out of the question. Instead, a multiscale approach made up of at least three stages has traditionally been used. The first step, called nuclear-data processing, consists on transport calculations in 0- or 1-dimensional systems to condense continuous-energy nuclear data into multigroup data libraries. The next step

corresponds to lattice-level calculations, which deal with representative subregions of the core typically using 2-dimensional pin-level models with a highly detailed geometrical, angular and energetic description of the flux. At this stage, the heterogeneous geometry of each zone is homogenized into an effective material and the nuclear data is condensed from a fine energy structure to a few-group library, typically with two or three energy groups. The last step is the solution of the full core with a nodal-level geometry and a coarse energy structure. Core simulators provide a low-order solution for the full core, and pin-level parameters can be obtained indirectly using methods such as lattice peak factors and pin-power reconstruction.

Although the calculation process from raw nuclear data to core simulations is quite complex and is based on strong physical approximations, this methodology has proven to give reliable results in practical times for most applications. As a result, deterministic methods are the industry standard for design and safety analysis. However, modern LWR core loadings, characterized by strong radial and axial heterogeneities, present challenges to the traditional deterministic approach. The physical assumptions used to go from lattice-level cells to nodal-level models are no longer valid in these conditions, emphasizing the need for higher-order deterministic solvers capable of handling full-core problems directly.

2.1.4 The Monte Carlo method

Instead of applying a discretization scheme to the Boltzmann equation, the Monte Carlo transport method is based on a stochastic description of the neutron motion. The idea is to simulate a large number of individual neutron histories and to accumulate statistical results of interest such as flux, power and reaction rates at specific locations. This is a very natural way to treat the neutron transport problem, since the description of the system through cross sections and other nuclear data is indeed probabilistic. Monte Carlo codes technically do not solve the neutron transport equation and are not derived from it, although the physical description of the system is completely analogous.

Each neutron history goes from its emission from a fission or external source to its absorption or escape from the system. The simulation of the motion of neutrons is called tracking, and is based on sampling free path lengths between collisions or material boundaries, as well as outcomes for capture, fission and scattering interactions. The free path length, for example, can be obtained using the definition of the total macroscopic cross section in equation 2.5 [8]. The probability of a neutron having a collision within dx after traveling a distance x from its last interaction point is

$$dp = p(x)dx, \quad (2.21)$$

where the probability density function $p(x)$ is

$$p(x) = \Sigma_t e^{-\Sigma_t x}, \quad (2.22)$$

and the cumulative distribution function $F(x)$ is

$$F(x) = \int_0^x \Sigma_t e^{-\Sigma_t x'} dx' = 1 - e^{-\Sigma_t x}. \quad (2.23)$$

Using the inversion method, the free path x can be sampled from $F(x)$ as

$$x = -\frac{1}{\Sigma_t} \log \xi, \quad (2.24)$$

where ζ is a uniformly distributed random variable. Reactions are sampled in a similar fashion, either by the inversion method or by rejection sampling techniques. The simulation of the neutron motion using the Monte Carlo method is based on sampling every probabilistic phenomenon using a random variable ζ obtained from a random number generator.

The gathering of results is completely separated from the actual simulation of the neutron motion, and results that are not tallied during the neutron tracking are in principle lost. The collection of results is based on accumulating statistics about discrete events for quantities of interest such as reaction rates, neutron fluence or power generation at specific locations and energy ranges. To do this, the total number of neutron histories for a transport calculation is usually divided in N batches, or generations. For an estimate X , the mean value \bar{X} can be calculated from the batch values X_n as

$$\bar{X} = \frac{1}{N} \sum_{n=1}^N X_n, \quad (2.25)$$

and the standard deviation $\sigma(X)$ is given by

$$\sigma(X) = \sqrt{\frac{1}{N(N-1)} \sum_{n=1}^N (X_n - \bar{X})^2}. \quad (2.26)$$

In this way, every tallied result is characterized by a mean value and a standard deviation which provides a measure of its statistical accuracy.

The main advantage of the Monte Carlo method is that it does not rely on any inherent approximations to simulate the neutron behavior. The geometry can be modelled up to an arbitrary level of detail without depending on discretization schemes, and the exact energy dependency can be considered using continuous-energy nuclear data. The direction of flight of neutrons is modelled explicitly and transient calculations can be performed taking into account the evolution in time of each neutron history. The main source of error in a Monte Carlo simulation is the statistical uncertainty in the results, which can always be decreased increasing the number of histories, with a penalty in calculation time. More complex stochastic issues related to source convergence, propagation of oscillations during burnup and unstable thermalhydraulic feedback loops contribute to the statistical error as well. Naturally, the accuracy is also bounded by the uncertainties in the nuclear data and the engineering specifications, for example, and on the modelling decisions regarding the geometry of the system.

The drawback of this approach is the huge calculation time typically needed to simulate a reactor core with a realistic degree of detail and an acceptable statistical convergence. This is true when comparing the Monte Carlo method with the multiscale scheme used in deterministic calculations. Novel deterministic transport solvers can handle the whole core with a comparable degree of spatial and angular resolution, though this is quite challenging and requires similar computational resources. Moreover, deterministic codes rely on the multigroup method and on resonance models, which is arguably a less accurate methodology than the continuous-energy formulation typically used in Monte Carlo calculations.

2.1.5 State-of-the-art neutronic analysis

The industry standard for neutronic analysis continues to be the multiscale approach described in Section 2.1.3, i.e. nuclear data processing, lattice-level calculations and core simulations. In this scheme, pin-level solutions can be obtained through pin-power reconstruction, which is based on calculating refined power distributions from nodal-level solutions and precalculated lattice-level profiles. While the standard tools for lattice-level calculations are deterministic codes like CASMO [9] and HELIOS [10], the use of continuous-energy Monte Carlo codes such as Serpent [11] and MCNP [12] is rapidly gaining popularity. One of the main advantages of using Monte Carlo codes for lattice calculations is that the generation of multi-group libraries using codes such as NJOY [13] is skipped altogether and nuclear data is used directly. Popular core-level simulators include SIMULATE [14] and PARCS [15], which handle core calculations using few-group nodal diffusion methods.

Regarding direct full-core pin-by-pin simulations, which is the topic of this thesis, novel high-fidelity deterministic codes have been under development in recent years [16]. Currently, these typically rely on so-called 2D/1D numerical schemes, where a set of 2D solutions, usually obtained using MOC solvers, are coupled axially through a low-order method such as diffusion or SP_N . This methodology is largely inspired by the DeCART code [17], which is based on planar MOC solutions and a coarse-mesh FDM formulation. Some additional examples include MPACT (2D MOC, 1D SP_3) [18], MICADO (2D MOC, 1D MOC) [19] and Proteus-MOC (2D MOC, 1D discontinuous Galerkin FEM) [20]. Furthermore, three-dimensional SP_3 solvers have recently been implemented in the PARCS and DYN3D [21] core-level codes. While these methods offer a quite accurate geometrical and angular representation of the problem, they still rely on precalculated few-group cross-sections.

Now, the computational resources required by deterministic methods to tackle full-core problems easily reach the requirements of Monte Carlo codes, which are arguably much more precise. For this reason, at this level of accuracy the Monte Carlo method seems more promising. This fact has led to an increasing preference for Monte Carlo codes over deterministic ones in the area of massive high-fidelity simulations. Several Monte Carlo neutronic codes have been developed and optimized for reactor analysis, in particular in terms of performance and memory requirement for burnup calculations. At the moment, codes such as RMC [22], MCS [23], MC21 [24] and MCNP, as well as Serpent, the tool used in this work, are capable of handling full-core problems at steady state. Large-scale burnup capabilities have only been developed in recent years, e.g. in RMC and MCS, and are still not widely available.

It is important to note here that at this level of detail it does not make sense to consider the neutronic solution without thermalhydraulic feedback. Hence, essentially all full-core pin-by-pin calculations are performed in a multiphysics context, as outlined in Section 2.4.6.

2.2 Thermalhydraulics

In a PWR, the power generated in the core is transferred by the primary cooling system to the steam generators, where water in the secondary system is boiled to feed the turbines and ultimately produce electricity. This cooling system has to ensure that the core is kept in a safe condition, avoiding excessively high fuel and cladding temperatures as well as critical phenomena such as DNB. Thus, an accurate

solution for the flow of water through the core and the transfer of heat from the fuel rods to the coolant is of key importance for core physics analysis.

2.2.1 Flow equations

The most general description of a fluid flow for engineering applications involving heat transfer is based on the conservation of mass, momentum and energy [25], [26]. The resulting equations depend on whether the flow is considered as compressible or not. Compressible flows are characterized by fluid density variations in response to pressure changes, whereas if these can be neglected the flow is said to be incompressible. The limit is usually given by the Mach number Ma , defined as

$$Ma = \frac{u}{c}, \quad (2.27)$$

where u is the local flow velocity and c is the speed of sound in the medium. A flow where $Ma > 0.3$ is typically considered as compressible, while $Ma < 0.3$ can be treated as incompressible.

For a compressible flow, the mass conservation or continuity equation is

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{u}) = 0, \quad (2.28)$$

where ρ is the fluid density and \vec{u} is the fluid velocity in the three-dimensional space. The conservation of momentum is governed by the Navier-Stokes equation, which is given by

$$\frac{\partial}{\partial t}(\rho \vec{u}) + \vec{u} \cdot \nabla \cdot (\rho \vec{u}) = -\nabla p + \nabla \cdot \bar{\tau} + \rho \vec{g}, \quad (2.29)$$

where p is the fluid pressure, $\bar{\tau}$ is the deviatoric stress tensor and \vec{g} is the gravitational acceleration. Finally, the energy conservation equation can be written as

$$\begin{aligned} \frac{\partial}{\partial t} \left[\rho \left(e + \frac{1}{2} u^2 \right) \right] + \nabla \cdot \left[\rho \vec{u} \left(e + \frac{1}{2} u^2 \right) \right] \\ = \nabla \cdot (k \nabla T) + \nabla \cdot (-p \vec{u} + \bar{\tau} \cdot \vec{u}) + \vec{u} \cdot \vec{F} + Q, \end{aligned} \quad (2.30)$$

where e is the internal energy per unit mass, k is the thermal conductivity, \vec{F} are the volumetric forces and Q is an external heat source. For newtonian fluids such as water, the viscosity term $\bar{\tau}$, which depends on the dynamic viscosity μ , is given by

$$\bar{\tau} = \mu(\nabla \vec{u} + \nabla \vec{u}^T) - \frac{2}{3} \mu \mathbb{I}(\nabla \cdot \vec{u}). \quad (2.31)$$

For incompressible flows, where the density is constant by definition, the continuity equation becomes

$$\nabla \cdot \vec{u} = 0, \quad (2.32)$$

and the Navier-Stokes equation yields

$$\rho \frac{\partial \vec{u}}{\partial t} + \rho \vec{u} \cdot \nabla \vec{u} = -\nabla p + \mu \nabla^2 \vec{u} + \rho \vec{g}. \quad (2.33)$$

The energy conservation equation without viscous heating can be written as

$$\rho c_p \frac{\partial T}{\partial t} + \rho c_p \vec{u} \cdot \nabla T = \nabla \cdot (k \nabla T) + Q, \quad (2.34)$$

where c_p is the specific heat capacity of the fluid.

Both for compressible and incompressible flows, the balance of mass, momentum and energy can be used to obtain the p , \vec{u} and T fields. Depending on the case, additional equations are needed to close the formulation, e.g. state equations for p , T and ρ and $\mu(T)$ equations. In PWR systems at nominal operating conditions, the most significant density variations are due to temperature changes, which induce buoyancy-driven effects. Other phenomena arising in compressible flows, e.g. shock waves, are not relevant. In this case, the standard approach is to use the Boussinesq approximation [27], assuming a linear dependence between ρ and T , namely

$$\rho = \rho_0 - \alpha\rho_0(T - T_0), \quad (2.35)$$

where α is the thermal expansion coefficient, and considering these density variations only in the body forces $\vec{F} = \rho\vec{g}$. This means that buoyancy is limited to the momentum equation and is not taken into account in the energy equation.

Furthermore, the analysis of a thermalhydraulic problem is significantly different depending on whether the fluid flow is laminar or turbulent. Laminar flows are characterized by slow velocities, where viscous forces dominate over inertial forces, generating a smooth motion pattern. Conversely, flows in which inertial forces are prevalent present chaotic turbulent phenomena such as eddies, swirls and temporal instabilities. LWR cores operate at turbulent flow conditions, where heat transfer and mixing is much more efficient.

2.2.2 Boundary and initial conditions

In order to calculate p , \vec{u} and T , the mass, momentum and energy conservation equations must be closed with appropriate boundary conditions. The standard way to solve the fluid flow problem in a reactor is to use a combination of Dirichlet and Neumann conditions at the core inlet and outlet.

An inflow velocity condition is obtained specifying the inlet velocity and an outlet zero-gradient constraint as

$$\begin{aligned} \vec{u} &= \vec{u}_{in}, \vec{r} \in \partial V_{in}, t > 0, \\ \nabla\vec{u} \cdot \hat{n} &= 0, \vec{r} \in \partial V_{out}, t > 0, \end{aligned} \quad (2.36)$$

where ∂V_{in} and ∂V_{out} are the boundary surfaces of the system at the core inlet and outlet, respectively. The temperature boundary conditions can be defined in an analogous way as

$$\begin{aligned} T &= T_{in}, \vec{r} \in \partial V_{in}, t > 0, \\ \nabla T \cdot \hat{n} &= 0, \vec{r} \in \partial V_{out}, t > 0. \end{aligned} \quad (2.37)$$

Transported variables, e.g. soluble boron, are specified at the core inlet in the same way the temperature is. An outflow condition for the pressure is specified imposing a zero gradient at the inlet and fixing the outlet pressure as

$$\begin{aligned} \nabla p \cdot \hat{n} &= 0, \vec{r} \in \partial V_{in}, t > 0, \\ p &= p_{out}, \vec{r} \in \partial V_{out}, t > 0. \end{aligned} \quad (2.38)$$

In addition to the core inlet and outlet, boundary conditions need to be applied for all rods at the cladding outer surface, where the fluid domain ends. In this case

a no-slip condition for the velocity is used, namely

$$\vec{u} = 0, \vec{r} \in \partial V_{clad}, t > 0, \quad (2.39)$$

where ∂V_{clad} is the contact surface between the cladding and the coolant over the whole computational domain. The heat produced by fission inside the fuel rods is transferred to the coolant through convection at the cladding-coolant interface, where the heat flux q'' is given by Newton's convection boundary condition

$$q'' = h_{clad-cool}[T_{clad} - T], \vec{r} \in \partial V_{clad}, t > 0, \quad (2.40)$$

where $h_{clad-cool}$ is the convection heat-transfer coefficient and T_{clad} is the cladding outer temperature. For fixed q'' , $h_{clad-cool}$ and T_{clad} fields this is effectively a Dirichlet boundary condition for T .

The initial conditions, which prescribe the p , \vec{u} and T at $t = 0$, are not relevant in steady-state problems.

2.2.3 Computational Fluid Dynamics

Computational Fluid Dynamics (CFD) is the set of numerical methods used to approximate solutions of a fluid flow by means of the mass, momentum and energy conservation equations introduced in the previous section. CFD applications are based on the discretization of the conservation equations typically using FVM, FEM or spectral methods and a temporal integration scheme.

While the solution of laminar flows is straightforward, turbulent flow patterns are extremely challenging. Direct Numerical Simulation (DNS), which is the solution of the Navier-Stokes equations without approximations, require massive computational resources and are only used for fundamental research. Industrial applications rely on turbulence models such as Large Eddy Simulation (LES) or Reynolds-averaged Navier-Stokes (RANS) to solve turbulent flows in a reasonable time [25].

CFD methods are extremely useful to analyze specific phenomena in a reactor core, since they can provide a very detailed and accurate representation of the flow conditions. However, the simulation of the whole reactor with a proper modelling of the pin-level geometry to capture local effects is not feasible. Instead, some applications can handle core calculations using porous media approaches, where the geometry is not modelled explicitly but through additional resistance terms in the Navier-Stokes equations [28].

Moreover, one of the main challenges in CFD applications for nuclear reactors is the treatment of two-phase flows. While research is underway in this area, the solution of turbulent multiphase flows in reactor geometries is still limited to relatively small problems [29]. That being said, PWRs at nominal conditions operate without void fraction, and therefore two-phase flows are not a significant factor for this thesis.

To sum up, CFD is a very valuable resource to analyze problems relevant to core physics with a detailed description of the flow patterns. However, its use in industry-like applications for full-core analysis is still not practical, mainly due to the computational costs involved.

2.2.4 Subchannel analysis

The geometry of a LWR core is composed of a regular array of fuel rods which define subchannels that direct the fluid flow. Hence, instead of relying on a description of the coolant flow from first principles, an ad-hoc model can be derived for the particular characteristics of the system [30]. Subchannel analysis is based on the formulation of the conservation equations for a control volume like the one shown in Figure 2.1, defined by a set of rods and an axial discretization. The flow velocity is characterized by an axial component u , which is the main flow direction at normal operation, and cross-flow components w_i that determine the mixing between a subchannel and the ones adjacent to it.




FIGURE 2.1: Control volume for subchannel analysis. The axial velocity is u and the cross-flow components of the velocity are w_i , one for each adjacent subchannel (4 in PWRs and 6 in VVERs).

The continuity equation for a subchannel control volume with an axial flow area A can be written as

$$A \frac{\partial \rho}{\partial t} + \frac{\partial \dot{m}}{\partial z} + \mathbf{S}^T \vec{w} = 0, \quad (2.41)$$

where \dot{m} is the axial mass flow rate and \mathbf{S} is a matrix operator which performs the lateral finite difference operation for the cross-flow vector \vec{w} . The transpose operator \mathbf{S}^T performs a sum across the adjacent subchannels. The axial momentum equation is

$$\frac{\partial \dot{m}}{\partial t} + \frac{\partial}{\partial z}(\dot{m}u) + \mathbf{S}^T(u\vec{w}) + A \frac{\partial p}{\partial z} = F, \quad (2.42)$$

where F is the axial force due to friction and gravity. The lateral momentum equation is given by

$$\frac{\partial \vec{w}}{\partial t} + \frac{\partial}{\partial z}(u\vec{w}) + \frac{s}{l}c\vec{w} = \frac{s}{l}\mathbf{S}p, \quad (2.43)$$

where c is the lateral friction force and s/l represents the importance of the friction and pressure terms versus the inertial terms. Finally, the energy conservation equation is

$$A \frac{\partial}{\partial t}(\rho h) + \frac{\partial}{\partial z}(\dot{m}h) + \mathbf{S}^T(h\vec{w}) = q', \quad (2.44)$$

where h is the fluid enthalpy and q' is the total heat source. The boundary conditions are the ones outlined in Section 2.2.2 for the core inlet and outlet. While equation 2.40 still holds for the convection heat transfer at the cladding-coolant interface, the no-slip condition is not used explicitly due to the coarse discretization used in subchannel analysis.

The geometrical description and the resulting equations are significantly simpler for subchannel analysis compared to CFD methods, since they take into account the specific flow conditions in a reactor core. Hydraulic and heat-transfer phenomena are solved through empirical models and correlations, for instance for the turbulent-friction and heat-transfer coefficients. The cross-flow is typically decomposed in convective and turbulent components, sometimes with a void-drift model to account for the lateral redistribution of the vapor phase. A key aspect of a subchannel method is the treatment of multiphase flows in terms of the conservation equations, which can range from a one-phase liquid-vapor mixture to a general multiphase formulation.

The main advantage of subchannel methods compared to CFD is their ability to account for a wide variety of two-phase heat-transfer phenomena relevant to reactor operation at a reasonable computational cost. For this reason, this methodology remains the standard approach to perform core analysis calculations. These can be carried out either at subchannel level, if the pin power is available from a pin-level neutronic calculation, or using assembly-wise control volumes and only taking into account the nodal-level power distribution. In real life applications, the main challenge is to have relatively accurate correlations for phenomena such as DNB and subcooled boiling, which are subject to large uncertainties and are strongly dependent on the exact characteristics of the system. Fuel designers and reactor operators often develop specific (and often proprietary) models for their core design.

It is worth noting here that CFD and subchannel methods are not exclusive but rather complementary. CFD calculations can be used to adjust the cross-flow parameters used in the subchannel approach and to quantify pressure losses and turbulent mixing in spacer grids and mixing vanes to develop empirical models. Furthermore, the accuracy of subchannel calculations for relevant numerical experiments can be verified by comparison with CFD results, which are typically considered to be more accurate.

2.2.5 State-of-the-art thermalhydraulic analysis

Currently, novel tools such as GeN-Foam [28], a multiphysics solver based on OpenFOAM [31], are attempting to provide CFD analysis capabilities for full-core reactor calculations. Industry-standard simulation software such as Fluent [32] can be used to develop CFD models for nuclear applications as well. However, CFD calculations are still limited to coarse-mesh models using the porous-media approach, and modelling the heterogeneous pin-level geometry is not feasible due to the computational costs involved. Moreover, as mentioned in Section 2.2.3, the solution of two-phase flows using CFD is currently limited to academic small-scale problems.

The standard method both in research and in the nuclear industry to study reactor problems is using subchannel codes, most of which originate from COBRA-IV [30], giving raise to the COBRA family [33]. In the simplest two-phase-flow formulation, COBRA-based codes solve the conservation equations for the liquid-vapor mixture without treating each phase explicitly, resulting in three mixture equations and one cross-flow equation, as outlined in Section 2.2.4. This type of codes is typically suitable to analyze PWR cores at nominal operating conditions, where the

forced convection conditions produce a primarily upward flow and the void fraction is zero. Examples of these include MATRA [34] and SUBCHANFLOW (SCF) [35], the code chosen for this thesis. More complex models deal with the full two-phase problem in the conservation equations. Some examples include SACoS [36], which uses a two-fluid, two-field model (six conservation equations), and COBRA-TF (CTF) [37], which is based on a two-fluid, three-field formulation that consists on a fluid film, fluid drops and the vapor phase (nine equations). These codes are more suitable to analyze complex scenarios on which recirculation and large void fractions can occur, at a higher computational cost.

2.3 Thermomechanics

As power is generated in a LWR core, the fuel rods experience an extraordinarily complex set of thermal, mechanical and physical phenomena. The production of fission isotopes, the strong radiation field and the high temperatures completely transform the properties of the fuel and damages the cladding material. These changes affect the temperature distribution in the fuel rods and the mechanical integrity of the cladding. As a result, the modelling of this complex thermomechanic system is critical for a proper safety evaluation of the fuel during the operating cycle.

2.3.1 Heat-transfer and mechanics equations

The evolution of the temperature distribution T in a fuel rod is described by the general heat conduction equation [38], given by

$$\frac{\partial}{\partial t}(\rho c_p T) = \nabla \cdot (k \nabla T) + q''', \quad (2.45)$$

where ρ is the density, c_p is the specific heat capacity and k is the thermal conductivity. The energy generation rate per unit volume q''' is given by the fission power distribution. This equation can be used to calculate the temperatures in the fuel and the cladding for known material properties, which are in principle neither constant nor uniform.

While conduction is the dominant phenomenon in the fuel and the cladding, the heat transfer in the fuel-cladding gap has a large radiation contribution, where the power per unit area q'' can be calculated as

$$q'' = \sigma T^4, \quad (2.46)$$

where σ is the Stefan-Boltzmann constant.

The mechanical behavior is given by the equations of solid mechanics, which govern the stress $\bar{\sigma}$ and strain $\bar{\epsilon}$ tensors [39]. The internal compatibility equation can be written as

$$\bar{\epsilon} = \mathbb{D}\vec{u}, \quad (2.47)$$

where \mathbb{D} is the compatibility differential operator, which relates the strains to the displacement vector \vec{u} . The internal equilibrium equation is

$$\mathbb{D}^T \bar{\sigma} + \vec{b} = 0, \quad (2.48)$$

which maps the body forces \vec{b} from the stress tensor.

The constitutive equation for elastic problems is given by Hooke's law

$$\bar{\sigma} = \mathbb{K}\bar{\epsilon}, \quad (2.49)$$

which prescribes the stress-strain relation in terms of the elasticity operator \mathbb{K} . In problems where plasticity plays a role, creep and non-elastic models are used to calculate strain rates.

The thermomechanic analysis of a fuel rod consists on the joint solution of the heat-transfer and mechanics equations. The mechanical analysis is highly dependent on the temperature distribution through processes such as thermal expansion, densification, creep, plasticity and phase transitions. In turn, the deformation of the fuel and the cladding determines the geometry in which the thermal problem needs to be solved. The fuel-cladding gap width, which is a key parameter for the heat transfer from the fuel to the coolant, is determined by the deformation of both the cladding and the fuel pellet. This also influences the free volume and the pressure in the gap, which in turn affect the heat conductance of the fill gas.

2.3.2 Boundary and initial conditions

The heat conduction and solid mechanics equations applied to the analysis of a fuel rod must be closed with boundary conditions at the cladding-coolant interface, i.e. at the outer surface of the cladding. For equation 2.45 a Dirichlet boundary condition analogous to equation 2.40 but from the cladding side can be specified as

$$q'' = h_{clad-cool}[T - T_{cool}], \quad \vec{r} \in \partial V_{clad}, t > 0, \quad (2.50)$$

where T_{cool} is the temperature of the coolant in contact with the cladding. Equations 2.47 and 2.48 are closed using a fluid pressure boundary condition over the stress at the surface as

$$\bar{\sigma}\hat{n} = -p\hat{n}, \quad \vec{r} \in \partial V_{clad}, t > 0, \quad (2.51)$$

where \hat{n} is the normal unity vector at ∂V_{clad} and p is the coolant pressure.

In this case initial conditions are indeed needed, since the thermomechanic problem is simulated as a transient even in burnup calculations, where the neutronic and thermohydraulic simulations consist of a series of steady-state solutions. The initial condition is typically taken as the steady-state solution for the boundary conditions at $t = 0$.

2.3.3 Fuel-performance analysis

Fuel-performance analysis deals with the study of the thermomechanic behavior of fuel rods under the irradiation conditions during reactor operation. In addition to the solution of the heat-transfer and stress-strain equations described in the previous section, the changes in the physical characteristics of the fuel and cladding materials have to be taken into account for a reasonable description of the system. This analysis can be performed for each rod independently, since the behavior of any given rod only affects other rods indirectly, for instance through neutronic or cooling effects. Hence, fuel-performance codes usually deal with a single rod and boundary conditions are imposed at the cladding-coolant interface. Figure 2.2 shows a scheme of the system in question, which consists of the cladding, the fuel pellet and the fuel-cladding gap, along with the axial gap between pellets. Macroscopic structures such

as cracks, as well as microscopic ones like grain boundaries, have to be taken into account.




FIGURE 2.2: Fuel rod system for fuel-performance analysis [40].

Computational codes for fuel-performance analysis typically include a wide range of physical phenomena with varying degrees of accuracy. Some of the most important aspects for solid ceramic fuels, and in particular for LWR cores, are [41]:

- **Fuel densification.** Especially during the first stages of irradiation of a fresh fuel assembly, a significant contraction, or densification, of the fuel tends to occur. This is due to two main effects, namely:
 - Irradiation-induced densification. The strong radiation field to which the fuel pellets are subjected causes the porosity of the fuel as fabricated to decrease.
 - Thermal densification. Thermal effects contribute to densification influencing the migration of defects and pores.
- **Swelling.** The accumulation of solid and gaseous fission products produces an expansion, or swelling, of the fuel. While the contribution of solid fission products is often negligible, gaseous ones have a significant impact on fuel swelling. The main mechanisms that dominate the behavior of fission gases are:
 - Gas bubble behavior. Fission gases tend to nucleate in bubbles, in particular on fission-fragment tracks or dislocation lines. Bubble growth is affected by the availability of vacancies and by the mechanical state of the fuel matrix. The migration and coalescence of these bubbles is determined primarily by temperature and stress gradients, as well as the presence of dislocations and grain boundaries.
 - Fission Gas Release (FGR). As fission gas bubbles grow and migrate, a fraction of them leave the fuel matrix and are released to the central void in the pellet, to cracks in the fuel or to the fuel-cladding gap. Fission gas release is composed mostly of inert gases, e.g. Xe and Kr, which get mixed with the fill gas (typically He) once they escape the fuel matrix.

- **Creep.** Slow deformations under the influence of stresses below the yield stress can produce severe deformations in nuclear fuel under irradiation. This phenomenon is called creep, and is characterized by a strain rate $\dot{\epsilon}$. Some of the factors that influence creep are:
 - Stoichiometry. The activation energy for creep is very sensitive to the stoichiometry of the fuel and its evolution. In particular, an excess of oxygen increases significantly the creep rate of UO_2 .
 - Grain size. At low stresses the creep rate tends to increase as the grain size decreases, while the opposite behavior is observed at high stresses.
 - Density. The creep rate increases with an increase in porosity.
 - Irradiation. The creep rate seems to appear at lower temperatures under irradiation, although this is not well understood and is typically not included in creep models.
- **Plasticity.** Plastic deformations occur under stresses above the yield stress and result in permanent distortions such as elongation, compression or buckling.
- **Pellet cracking and relocation.** Immediately after the power increases, and before any significant swelling or creep can occur, high thermal stresses are produced by the strong radial temperature gradients. These stresses induce the cracking of the fuel pellets, and the resulting pellet fragments relocate within the fuel rod. This cracked structure evolves during irradiation, as the fuel undergoes thermal and irradiation-driven stresses, as well as power ramps.
- **Phase transitions.** As the structure, temperature distribution and composition of the fuel evolves, phase transitions can occur. These lead to changes in the fuel volume and its mechanical and thermal properties.
- **Formation and closure of central void.** As the fuel is heated up, the porosity in the hot region of the pellet tends to migrate towards the center line. This may produce a central void, which is also influenced by swelling and densification. Some fuel types, e.g. VVER standard fuel, are fabricated with a central void.
- **Pellet-cladding Mechanical Interaction (PCMI).** Fuel rods are constructed with a gap between the cladding and the fuel pellets, which is maintained by the structural integrity of the cladding and by a fill-gas pressure of about 2.5 MPa that partially compensates the coolant pressure. During irradiation the gap width tends to decrease, primarily due to the effect of fuel swelling and thermal expansion, in addition to cladding creep, and closes completely after about one to three years of irradiation. At this point, the pellet and the cladding are in direct contact, producing relatively high tensions on the cladding and generating a characteristic deformation pattern called bambooing. In addition, the presence of fission products such as iodine and cadmium can induce Stress Corrosion Cracking (SCC) of the cladding. While typically the cladding can withstand this condition at nominal power, a failure of the cladding can occur during strong power transients.

The full set of phenomena listed above constitutes a tightly coupled system which ultimately determines the thermomechanic properties of the fuel, e.g. fuel conductivity, its elasticity and indeed its geometry. In turn, physical processes such as fuel swelling, creep, plasticity and cracking are a direct result of the stress-strain

conditions, and the temperature distribution produces densification, swelling and phase transitions, among other effects. As a result, fuel-performance analysis consists on the solution of the heat-transfer and mechanics equations coupled with the fuel and cladding physical behavior in an extremely complex non-linear system.

Fuel-performance codes can range from ab-initio models at the nm-scale to reactor analysis codes dealing with an entire fuel rod. This thesis deals with the latter, which serve to study the behavior of the fuel in real-life conditions and during the whole operating history.

2.3.4 State-of-the-art fuel-performance analysis

Most fuel-performance analysis codes deal with a single rod and handle the thermal, mechanical and material calculations with various degrees of complexity. Codes such as BISON [42], which is based on FEM thermal and mechanical analysis, deal with a very fine discretization of the problem to obtain detailed three-dimensional solutions. While these are very valuable tools to study particular phenomena, their application is typically limited to problems of the scale of a fuel pellet. The most suitable codes for reactor analysis are based on simplified geometrical models, typically a superposition of one-dimensional radial and axial descriptions (so-called $1\frac{1}{2}$ schemes). FRAPCON [43] and TRANSURANUS (TU) [44], the tool used in this work, are examples of these.

2.4 Multiphysics

In order to analyze the neutronic, thermalhydraulic and thermomechanic state of a LWR core in a realistic way, the interactions between these fields have to be considered. While some of the most obvious feedback mechanisms between physical domains have been mentioned in passing in the previous discussions, the multiphysics system needs to be described in detail to pose the problem at hand.

2.4.1 Thermalhydraulic and thermomechanic effects on neutronics

The system on which the neutron transport equation needs to be solved is characterized by its nuclear data, its material compositions and its geometry. In the formulation presented in Section 2.1.1.4 this information is taken as fixed data, but it is in fact dependent on the thermalhydraulic and thermomechanic state of the system, giving rise to strong non-linearities. Cross sections are affected primarily by density and temperature variations, and geometrical changes can be produced by any number of phenomena.

2.4.1.1 Density variations

The macroscopic cross sections for a given material are directly proportional to the atomic density of each isotope in its composition, as equations 2.6, 2.7 and 2.8 show. Hence, density variations can be accounted for directly through the cross sections of the system.

The most relevant density changes in a LWR core occur in the coolant, where the flow conditions, including boiling, can produce strong density gradients. While the fuel does undergo contraction and expansion through the mechanisms outlined in Section 2.3.3, its mass remains constant. If changes in dimensions are neglected in the neutronic calculation, which is typically the case as explained in Section 2.4.1.3,

this means that the fuel density must be taken as constant. The same is true for the cladding and other structural materials. The common practice in LWR analysis is to only take into account changes in the coolant density as a feedback mechanism while taking the densities of all other materials as constant.

2.4.1.2 Temperature variations

At the microscopic level, the main factor that influences the interaction of neutrons with target nuclei is the thermal motion of the atoms, which is determined by the temperature of the medium [8]. The neutron energy E in the frame of reference of the collision, which is given by the thermal motion of the target nucleus, determines the interaction probability through the microscopic cross section. For this reason, the temperature distribution of the system is a key feedback parameter.

One effect of the thermal motion of the atoms is a $1/v$ dependence of cross sections with the neutron velocity v at non-zero temperatures. This effect is significant at low neutron velocities, where the velocity of the target nucleus can be comparable or exceed the neutron velocity. This increases the collision rate at low energies and produces up-scattering, and thus has a strong influence on the flux spectrum.

Another low-energy thermal effect is related to neutron interaction with bound nuclei. When the neutron energy is of the order of the binding energy of an atom in a molecule or crystalline lattice, treating the target nucleus as a free particle is no longer a good approximation, and the effect of the molecular or crystalline structure needs to be considered. These interactions are characterized by temperature-dependent cross sections and energy-angle distributions called $S(\alpha, \beta)$ scattering laws.

Finally, cross sections with resonance peaks are subject to strong Doppler broadening. The thermal motion of the target causes a random variation in the relative velocity between the neutron and the target nucleus, increasing the effective interaction probability near the resonance.

Low-energy effects, i.e. $1/v$ cross-section dependence and bound-nuclei interactions, play a major role in neutron moderation. Hence, coolant temperature feedback with $S(\alpha, \beta)$ treatment of H atoms bounded to water must be considered. Fuel temperature feedback is needed to simulate the Doppler broadening in resonance absorbers such as ^{235}U , ^{238}U and ^{239}Pu . Although thermal effects do have an impact on other materials, a dynamic temperature feedback is typically considered only for the coolant and fuel materials, and other temperatures are modelled as constant.

2.4.1.3 Geometry variations

During reactor operation the geometry changes due to control rod movements and fuel reshuffling, for example. These are not multiphysics effects, because they are part of the problem specification, but they affect the system directly and are always taken into account in the neutronic calculation.

The geometrical changes that occur in the fuel rods do involve multiphysics calculations and are more complex to analyze. In principle, the expansion and contraction of the fuel pellets and the deformation of the cladding modifies the geometry, and one might think about accounting for this in the neutronic calculation. While this is typically done in fast-spectrum reactors, in LWR systems this is a second-order effect and is only rarely taken into account. Furthermore, the microstructure of the fuel pellets, e.g. cracks, grain boundaries, gas bubbles, etc, is outside the scale

of typical neutron free-path lengths, and is safely neglected from the neutronic point of view.

2.4.2 The power distribution

In a LWR core, most of the power is generated in the fuel through fission reactions. This heat is deposited in the fuel rods and transferred to the coolant as described in the next section. A small fraction of the power, typically about 2% for PWR designs, is produced in the coolant by gamma radiation and acts as a direct heat source for the cooling water. The power distribution in both the fuel and the coolant is a direct product of the neutronic calculation. The estimation of the energy deposition by gamma radiation can range from the use of a constant fraction of the fission power to coupled neutron-photon transport [45]. In practice, the gamma heating fraction is usually neglected and all the power is considered to be generated by fission in the fuel rods.

2.4.3 Neutronic effects on fuel behavior

Since the motivation for thermomechanic calculations is to simulate the behavior of fuel rods under irradiation, the neutronic evolution of the core obviously plays a major role in fuel-performance analysis.

The most important parameter for the fuel and cladding calculation is the power distribution. This is the heat source for equation 2.45 and produces the temperature gradients that induce thermal expansion, creep, cracking and phase transitions, among other phenomena. The neutron flux, especially the fast flux, is a key variable as well, since it produces radiation damage in the fuel and cladding material and contributes to densification and creep.

Regarding the fuel composition, the rate at which fission products are created during operation is a result of the neutronic evolution of the core. The burnup calculation is therefore the source term for the simulation of the behavior of these isotopes as they contribute to swelling, fission gas release and PCMI, as described in Section 2.3.3.

2.4.4 Boundary conditions at the cladding-coolant interface

The interaction between the thermalhydraulic and thermomechanic models happens at the cladding-coolant interface, where the fission power is transferred from the rods to the coolant. In transient problems these two systems are tightly coupled by the temperature difference at the interface, but can be treated in a simpler manner in steady-state conditions.

2.4.4.1 Transient problems

The boundary conditions related to the heat transfer across the cladding-coolant interface, namely equations 2.40 and 2.50, link the thermalhydraulic and thermomechanic calculations through the temperature difference $\Delta T_{clad-cool} = T_{clad} - T_{cool}$ between the cladding outer surface and the coolant in contact with it. Hence, in the general case of a transient problem both domains need to be solved as a coupled system.

2.4.4.2 Steady-state problems

While equations 2.40 and 2.50 still hold at steady-state conditions, the thermalhydraulic problem can be decoupled from the heat conduction one. This can be done by noting that, if axial heat conduction is neglected, which is typically a good approximation, the fission power q'' is actually the coolant heat source at each rod and axial location, because no heat can accumulate inside the rods at steady state. Hence, in this case the flow calculation can be performed directly using the fission power. After the thermalhydraulic state is obtained, T_{clad} can be calculated from equation 2.40 and used as boundary condition in equation 2.50 for the thermomechanic calculation.

Even if the correlation for $h_{clad-cool}$ depends on T_{clad} , which is typically not the case, the thermalhydraulic and thermomechanic problems are only loosely coupled. As an example, the Dittus-Boelter correlation, perhaps the most popular correlation for the convection heat-transfer coefficient in PWR applications, can be written as

$$Nu = 0.023Re^{0.8}Pr^{0.4}. \quad (2.52)$$

The Nusselt number

$$Nu = \frac{hL}{k}, \quad (2.53)$$

where L is the characteristic length of the system, represents the ratio of convective to conductive heat transfer. The Reynolds number is defined as

$$Re = \frac{\rho u L}{\mu}, \quad (2.54)$$

and represents the ratio of inertial forces to viscous forces, whereas the Prandtl number is

$$Pr = \frac{c_P \mu}{k}, \quad (2.55)$$

and is the ratio of momentum diffusivity to thermal diffusivity. At first order, the Dittus-Boelter correlation depends only on the bulk parameters of the fluid flow. While some correlations include corrections to account for the coolant temperature in contact with the cladding and use the cladding temperature, this is a second order effect. In these cases, T_{clad} can be calculated from the convection equation

$$T_{clad} = \frac{q''}{h_{clad-cool}(T_{clad})} + T_{cool}. \quad (2.56)$$

2.4.5 Boundary conditions for the full system

While the convection heat-transfer conditions at the cladding-coolant interface are a boundary condition for the thermalhydraulic and thermomechanic calculations, they represent an internal interface between physical domains when considering the multiphysics system as a whole. As such, they do not need to be treated as boundary conditions.

The actual boundary conditions for the system are (a) the neutronic boundary conditions at the outer reflector surface (equation 2.18) and (b) the cooling conditions at the core inlet and outlet (equations 2.36, 2.37 and 2.38). In steady-state and depletion problems the system is solved using these two sets of parameters, plus the power level and the control rod positions.

2.4.6 State-of-the-art multiphysics analysis

The development of high-fidelity reactor analysis tools is tightly linked to the evolution of the computing power available to nuclear engineers and researchers. On the one hand, the wide access to HPC systems means that methods such as Monte Carlo particle transport have gained popularity as a more accurate alternative to traditional methods, perhaps not for routine calculations but certainly for verification and novel simulations. On the other hand, high-fidelity applications are still typically bound by the huge amount of computing time that is required for real-life calculations.

Current multiphysics capabilities are a combination of the status of each research area, as discussed, and of the development of code-integration methodologies. Neutronic-thermalhydraulic coupling is an extensively researched topic and one that is mandatory for LWR analysis, since any realistic neutronic simulation must include thermalhydraulic feedback. As a result, high-fidelity neutronics and subchannel analysis advance in tandem and several coupling systems based on the aforementioned codes exist, e.g. MPACT-CTF (VERA-CS framework) [46], DYN SUB [47], RMC-CTF [48], MCS-CTF [49], MCNP-SCF [50] and Serpent-SCF [47], [51], [52]. Although neutronic-thermomechanic coupling in a high-fidelity context is a much more recent field, coupling systems such as MCS-FRAPCON [53] have been developed in recent years.

Coupling neutronics, thermalhydraulics and thermomechanics is quite a challenging task due to the complex feedback system and the high computational costs. Two high-fidelity systems are currently under development: MCS-CTF-FRAPCON [54], which has only been demonstrated for steady-state problems, and MPACT-CTF-BISON [55], which has recently been shown to be able to handle depletion cases. Although the latter is similar to the target tool of this thesis, MPACT has limitations in terms of physical approximations and energy resolution compared to Monte Carlo transport codes. Finally, nodal-level tools such as PARCS-SCF-TU [56] have been implemented for transient analysis considering fuel-performance aspects, though with a much coarser resolution.

Chapter 3

Coupling neutronics, thermalhydraulics and fuel-performance analysis

The reactor core is a complex system characterized by the interaction between neutronic, thermalhydraulic and thermomechanic phenomena. To provide a consistent description of the core, a high-fidelity multiphysics tool must take into account each of these areas, as well as the main feedback mechanisms, in a detailed manner. This constitutes a challenge in terms of implementing a flexible and accurate tool contained in a maintainable software package.

Section 3.1 of this chapter introduces the three codes selected for this work, which determine the methodology used to analyze each of the physical systems. The calculation scheme proposed to describe the feedback mechanisms and to obtain the coupled solution is described in Section 3.2. Finally, Section 3.3 presents the software design of the coupled tool, which represents the practical implementation of the calculation method.

3.1 Selected calculation codes

The reactor analysis codes used to implement the multiphysics system proposed in this thesis are the Serpent continuous-energy Monte Carlo code, the SUBCHANFLOW subchannel thermalhydraulics code and the TRANSURANUS fuel-performance code.

3.1.1 Serpent

Serpent 2 [11] is a multipurpose three-dimensional continuous-energy Monte Carlo particle transport code developed at VTT Technical Research Centre of Finland, Ltd. It can be used for a broad range of reactor physics applications, including spatial homogenization, criticality calculations and fuel cycle studies, as well as for radiation dose calculations, shielding, fusion research and medical physics.

Like most Monte Carlo codes, the basic geometry description in Serpent relies on Constructive Solid Geometry (CSG), which allows for the description of arbitrarily complex geometries combining elementary surfaces through Boolean operations like intersections, unions and complements. In addition to CSG, unstructured meshes, stereolithography (STL) models and randomized geometries can be used. The system is defined by homogeneous material cells where continuous-energy nuclear data for each specific isotopic composition is used. Any nuclear data set in ACE format such

as the ENDF and JEFF library series can be used. This geometry and energy description is used to track particles through the system, and statistical results are gathered using detectors superimposed to the geometry, where any reaction for which nuclear data is available can be sampled.

Serpent includes built-in burnup calculation capabilities based on solving the evolution of nuclide compositions between steady states using reaction rates tallied automatically during tracking. The number of depletion zones, i.e. regions where the burnable material composition is taken as homogeneous, is user-defined and is typically restricted by the memory use, which will be studied extensively throughout this thesis. The number of isotopes considered in burnup calculations depends on the nuclear data library, and is usually in the order of 1500, between 200 and 300 of which have cross-section data and are included in the transport calculation. The Bateman equations are solved either by Transmutation Trajectory Analysis (TTA) or by the Chebyshev Rational Approximation Method (CRAM). Both of these methods decompose the depletion chains in linear equations and apply a time-stepping numerical integration scheme to them. The coupling between the transport and burnup calculations can be handled using the explicit Euler, predictor-corrector or higher-order methods [57], as well as the Stochastic Implicit Euler (SIE) method [58].

Coupled simulations are handled by Serpent through a universal multiphysics interface, which can be used to define density and temperature distributions, as well as tallying power, on meshes superimposed to the geometry [59]. This is a very convenient approach, since it keeps the feedback data separate from the CSG-based geometry model, which is typically quite complex. This is significantly more flexible than the method used in most Monte Carlo codes, where these state variables are specified for each material.

Serpent can be used to achieve an extreme degree of detail in the neutronic solution. The geometry of the system can be modelled up to an arbitrarily high resolution, in particular in terms of depletion zones, feedback fields and structural materials such as spacer grids and VVER stiffener plates. In this sense, the physical description of the system is greatly improved with respect to deterministic tools by removing spatial homogenization approximations and assumptions over the angular dependency of the flux. Moreover, using a continuous-energy formulation avoids the traditional multigroup approach, providing a much more advanced energy treatment which results in a better physical model. In the context of this work, the accuracy of the neutronic solution is limited by the computational resources and not by the calculation method, as will be clear in the following chapters.

3.1.2 SUBCHANFLOW

SUBCHANFLOW (SCF) [35] is a subchannel analysis code largely based on the legacy codes COBRA-IV-I [30] and COBRA-EN [60]. It has been developed at the Institute for Neutron Physics and Reactor Technology (INR) of the Karlsruhe Institute of Technology (KIT). The code is capable of performing steady-state and transient calculations, and has been extensively used and validated for LWR analysis.

A SCF model is composed of a set of subchannels and fuel rods with a graph-like connectivity, i.e. without an explicit geometrical representation. Subchannels are defined by their hydraulic parameters (flow area and heated and wetted perimeters) and by connections to other subchannels characterized by their gap widths and neighbor-to-neighbor distances. Fuel rods are defined by their fuel and cladding dimensions and material properties, as well as by the connectivity map between rods

and subchannels. The fission power is deposited in the rods, transferred by conduction through the fuel and cladding, as well as by radiation in the fuel-cladding gap, and removed by convection by the coolant.

SCF solves the fuel temperatures with a simplified fuel-performance model. The fuel-cladding gap width is calculated considering fuel swelling [61] and relocation by cracking [62], as well as pellet and cladding thermal expansion [61]. The gap conductivity takes into account heat conduction and thermal radiation through the filling gas, and depends on the gap width and the surface roughnesses. The fuel and cladding heat capacities depend only on temperature [61], while the fuel conductivity correlation also considers burnup [43].

It could be argued that SCF is not the most advanced subchannel code for core analysis, since the numerical solution is based on an upward flow regime with a limited cross-flow model and the two-phase formulation only considers the liquid-vapor mixture. Codes such as CTF provide a fully 3D numerical scheme with explicit treatment of the two-phase problem, and therefore could be considered more suitable for a high-fidelity application. However, this is not true for the problems of interest in this thesis, i.e. LWR cores at nominal operating conditions, where the flow is dominated by forced convection in the upward direction and the void fraction is zero or close to zero. In these conditions the SCF solution is expected to be quite reliable and not much could be gained from a more sophisticated approach, which would result in a much higher computational cost.

While SCF is an adequate tool for the purposes of this work, it does not represent an improvement with respect to the state of the art in terms of physical modelling. In the context of subchannel analysis, the advances in the phenomenological description are based mainly on improving the semi-empirical models for effects such as axially dependent pressure losses due to spacer grids or asymmetrical effects due to rod bowing. The innovative aspect in this thesis is related to using a full subchannel approach in a multiphysics framework, which can be considered a high-resolution methodology rather than a high-fidelity one from the thermohydraulic point of view.

3.1.3 TRANSURANUS

TRANSURANUS (TU) [44] is a fuel-performance code for thermal, mechanical and neutron-physical analysis of a cylindrical fuel rod developed at the Institute for Transuranium Elements (ITU), now the Joint Research Centre (JRC) at Karlsruhe. It is one of the standard thermomechanic codes in the nuclear industry and has been extensively validated for PWR and VVER fuels [63]. The calculation method deals with a single fuel rod discretized axially and radially using a $1^{1/2}$ scheme, i.e. a radial discretization with a loose axial coupling, to solve the thermal and mechanical equations.

A wide variety of physics are included in the thermomechanic model, i.e. thermal and irradiation-induced densification of the fuel, swelling due to solid and gaseous fission products, creep, plasticity, pellet cracking and relocation, oxygen and Pu redistribution, volume changes during phase transitions, formation and closure of the central void and treatment of axial friction forces. The fuel-cladding gap conductance is calculated using the URGAP model [64], and depends on the gap width or contact pressure between the fuel and the cladding, the gas pressure and composition and the surface characteristics of the fuel and the cladding. The gap width is a result of the mechanical calculation, which determines the deformation of the pellet and the cladding.

Depletion is simulated solving a simplified system of Bateman equations which accounts for the most relevant isotopes for the thermomechanic calculation. The neutronic parameters, i.e. radially dependent power, flux and reaction rates, needed for the rest of the calculation, are obtained using a low-order method based on the diffusion approximation with empirical corrections, for instance for Gd rods [65].

The calculation scheme proposed in this thesis consists of performing the thermomechanic analysis for every fuel rod in the system. To achieve this, the TU multiphysics module used in this work handles the whole core model, where the solution of each rod is performed independently for the given boundary conditions at the cladding-coolant interface. As in the case of SCF, this is arguably not a high-fidelity methodology but a high-resolution one, because the fundamental physical description of the system is not changed with respect to the original modelling approach. That being said, including fuel-performance analysis in burnup calculations, which is a novel methodology, does improve the physical model of the system.

3.2 Calculation scheme

The coupling scheme is one of the key aspects of a multiphysics system, since it determines the actual interaction between the physical domains. Furthermore, it specifies the workflow of the application, which has to be implemented in the software package.

3.2.1 The operator-splitting method

The multiphysics system at hand is a highly nonlinear one, not only within each physical domain but also regarding the feedback mechanisms. As an example, a change in the fission power affects the temperature and density distributions in the core, which in turn is reflected in the cross sections and thus in a perturbation in the power. All the feedbacks between the different physics outlined in Section 2.4 produce nonlinear behaviors.

Now, given the extremely complex dependence of these phenomena, analytical models are typically not practical. Back to the previous example, the changes in temperatures and densities due to power perturbations are a product of the whole thermalhydraulic and thermomechanic calculation, and their impact on the cross-sections does not have a trivial representation.

The most straightforward way to deal with such a system is through the operator-splitting method. When applied to nonlinear problems, this technique is quite intuitive and consists on ignoring sources of nonlinearity when solving each physical domain and addressing them indirectly, typically through a so-called Picard iterative scheme. In the neutronic-thermalhydraulic example, this basically means calculating the power with fixed temperature and density distributions, solving the thermalhydraulic state for this power and iterating if needed. Virtually all multiphysics systems, at least in reactor analysis, are handled in this way. In this context, a method is usually said to be explicit if no iteration takes place, semi-implicit if the solution is iterated until convergence and fully implicit if a more tightly coupled method is used, though this terminology is often ambiguous. Using implicit methods a converged solution across physics is achieved, which makes the result more consistent and the numerical algorithm more stable, at the expense of an increase in the computing time.

An advantage of the operator-splitting method, besides being the most straightforward to apply, is that in practice each physical domain is solved by a specific calculation code. Using this method, an iterative coupling scheme can be implemented without modifying directly the solution algorithm of each code, only the feedback data. This is very valuable, especially when using well-established and validated codes, where changes in the source code diminish the credibility of the results.

3.2.2 Coupled depletion methodology

The proposed calculation and feedback scheme to perform depletion simulations with the three-code system is shown in Figure 3.1, where t_n and t_{n+1} are burnup steps n and $n + 1$ [66]. The algorithm is based on the operator-splitting method with a semi-implicit iterative scheme. Code-to-code feedback is done using the fields at the end of the step (EOS) and convergence at EOS is achieved iterating each burnup step.

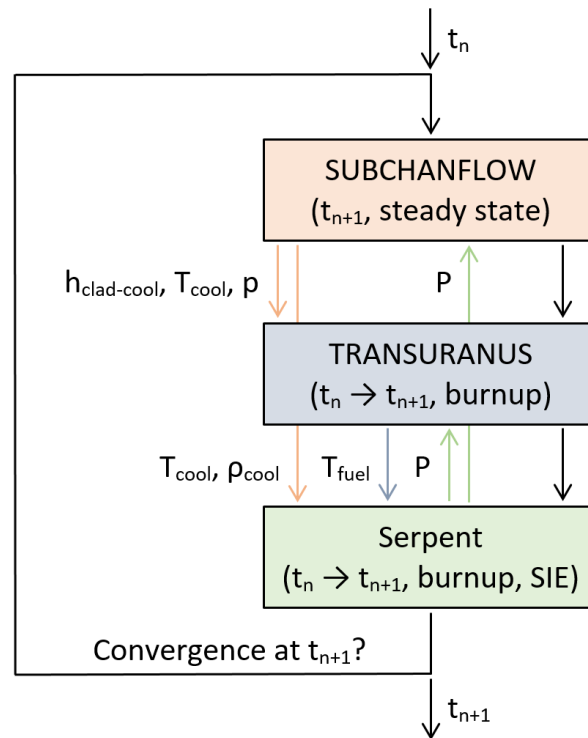


FIGURE 3.1: Depletion scheme [66]. The black lines show the calculation scheme, the color lines the feedback exchanges.

For each iteration, a SCF steady-state calculation is performed first to obtain the cooling conditions for the Serpent power distribution at EOS. Second, TU solves the burnup step from t_n to t_{n+1} using the SCF solution as boundary condition and the Serpent power as heat source, both at EOS. For the first iteration in each burnup step, when the Serpent solution at EOS has not yet been calculated, the power at the beginning of the step (BOS) is used in SCF and TU. Third, Serpent carries out a burnup iteration using the SIE method with the thermalhydraulic conditions at EOS. The SIE depletion scheme consists on averaging the solutions at EOS for all iterations within each burnup step to accumulate statistics and stabilize the solution. The rationale for this code order is that (a) SCF calculates the boundary conditions for TU and receives no feedback from it and (b) the Serpent calculation is the most expensive step, and therefore should be performed with the latest thermalhydraulic

data. The steady-state solution at $t_0 = 0$, which has to be computed before the burnup calculation, is obtained with the same iterative scheme starting from an initial guess for the power, typically a uniform distribution or a result from a previous run.

It is important to note here that Serpent and TU perform independent burnup calculations, Serpent with the full set of Bateman equations [57] and TU with a reduced one suitable for fuel-performance analysis [65]. The time-step size $dt = t_{n+1} - t_n$ for the whole system is determined by Serpent, which solves the steady-state transport problem at t_n and t_{n+1} , as well as the evolution of the isotopic densities between these two states. Now, since the physical systems simulated by TU have much smaller time scales than the typical time steps used in burnup calculations, TU advances from t_n to t_{n+1} using substeps that satisfy its internal stability constraints, and the global time step dt determines the feedback exchange interval. This is a key aspect of the calculation scheme, since TU can use small enough time steps internally without forcing a Serpent solution at each of these steps, which would increase the calculation time beyond practicality. SCF has no influence in the time-stepping method, since it only performs steady-state calculations.

In this scheme, the pin power P calculated by Serpent is used in SCF as the heat source for the coolant and in TU considering a radial power distribution within the fuel pellet based on an empirical method included in the burnup model [65]. The coolant temperature T_{cool} and density ρ_{cool} calculated by SCF and the fuel temperature T_{fuel} obtained by TU are used in Serpent as thermalhydraulic feedback for the cross sections. For T_{fuel} either a radial average, an effective Doppler temperature or a radial profile can be used. SCF provides the boundary conditions for TU, i.e. $h_{clad-cool}$, T_{cool} and p . This feedback scheme reflects the analysis of the multiphysics system explained in Section 2.4.

This algorithm can be reduced to the standard neutronic-thermalhydraulic coupling if TU is not included. In this case, the fuel calculation is handled by SCF, which incorporates a simplified fuel-performance model based on fuel-cladding-gap evolution and irradiation-driven thermal-conductivity degradation. In this thesis, this simpler scheme is used to compare the full coupling with the traditional methodology.

3.2.3 Additional remarks on the coupling scheme

The coupling scheme outlined in the previous section arises from the traditional neutronic-thermalhydraulic coupling method, replacing the simple fuel-rod solution in SCF by the full thermomechanic calculation performed by TU. Even though the methodology is clear from the physical point of view, a few issues are worth commenting on.

3.2.3.1 Thermalhydraulic-thermomechanic coupling

Based on the approximation introduced in Section 2.4.4.2 for steady-state problems, a one-way coupling is used between SCF and TU. SCF provides the boundary conditions for TU, but no information is exchanged from TU to SCF. This approximation is valid for steady-state and depletion problems, but it would have to be revised to apply this tool to transient problems, which is out of the scope of this thesis.

In this scheme, SCF uses the Serpent power directly as heat source for the coolant, without simulating the fuel rods. This calculation mode has been implemented in SCF as part of this work. When TU is not used, the power is deposited in the fuel rods and the standard SCF solver is used.

3.2.3.2 Neutronic feedback for fuel-performance analysis

While this coupling scheme captures the main interaction mechanisms in the system, the neutronic feedback to TU could in principle be refined. Calculating neutronic parameters based on the semi-empirical model implemented in TU might seem dubious, especially having an extremely accurate tool as Serpent. Radial power profiles could be tallied by Serpent and used directly in the thermomechanic simulation instead of calculating them in TU from the total pin power. However, doing this would dramatically increase the number of particles needed to obtain statistically meaningful values for the radial power profiles for each rod and axial level, which is not feasible in full-core pin-by-pin calculations. The same can be said about the neutron flux, which is calculated internally by TU from the power distribution. Finally, it is worth considering that the TU neutronic model is optimized and well validated with experimental data in the context of fuel-performance calculations. Hence, improving the thermomechanic calculation using the Serpent pin power while still relying on TU to calculate the radial power and flux profiles inside the fuel pellet is a reasonable decision.

3.2.3.3 Depletion calculation

In the proposed burnup scheme Serpent and TU handle the depletion problem independently, and therefore the isotopic compositions in the fuel are not the same for both codes. As an alternative, the burnup calculation could be performed only in Serpent and the isotope concentrations transferred to TU. The drawback of this approach is that it involves massive amounts of feedback data, considering the radial dependence of material compositions. Moreover, the multiphysics interface in Serpent, which is extremely convenient to exchange power, density and temperature distributions, is not suitable for obtaining material data. Finally, as for the radial power and neutron flux, the model used by TU is tailored for fuel-performance calculations, in particular in the treatment of important isotopes such as Pu and Gd, and therefore it is not a bad decision to keep relying on it.

3.3 Software design

Given the complexity of the multiphysics system at hand, the implementation of the proposed coupling scheme requires a well-thought software design. The programming methodology defines the capabilities, performance and flexibility of the tool, as well as its maintainability.

3.3.1 Traditional coupling approaches

Virtually all multiphysics systems used in the nuclear community are based on either of two software approaches: file-based and master-slave coupling. These are fairly simple programming methods and have been proven to be relatively adequate for most applications.

3.3.1.1 File-based coupling

The simplest way to run a simulation involving multiple codes is to exchange data through the native input and output files without any modifications to the corresponding source codes [67]. This method is sometimes called external coupling, as

opposed to internal coupling, where two or more codes are integrated at the source-code level. The main advantage of this approach is that the original codes are used without modifications and the coupling scheme is usually easy to implement using a scripting language. However, this method is typically only suitable to implement simple steady-state calculations that only involve running one code after the other and iterating. Burnup and transient applications imply a more complex calculation flow that cannot be managed unless additional features, e.g. semaphores, socket-based communication or save/restart capabilities, are implemented in each code. In addition, external coupling requires the codes to be initialized and terminated between solution steps, which in the case of Monte Carlo codes lead to overheads of the order of the calculation time.

3.3.1.2 Master-slave coupling

The most popular internal coupling approach, used nowadays by most multiphysics systems, is master-slave coupling [68]. This method is based on restructuring one or more "slave" codes into calculation libraries designed to run the solver in a step-by-step way. The coupling scheme is implemented in the "master" code, which drives the calculation and calls the slave codes according to the designed workflow. This method is usually quite efficient in terms of calculation time, because data is exchanged directly through memory and no initialization and termination is needed during the simulation. Moreover, while the slave codes need to be restructured, the logic of the master code is typically modified only slightly. The main drawbacks come from the intrinsic software design, as explained in the next section.

3.3.2 Motivation for a novel approach

While file-based coupling has some merits, internal coupling is clearly superior in terms of performance and flexibility, two key issues for the implementation of the three-code high-performance system proposed in this work. However, the master-slave approach has severe problems when implementing such a complex tool consisting of three full-fledged analysis codes.

The most obvious problem is that embedding two codes in the source code of a third one leads directly to "spaghetti code", i.e. a source code which is extremely hard to understand, debug and maintain. The workflow of the coupled system quickly becomes hard to follow, since the execution stages of the slave codes are deeply nested inside the logic of the master code. The maintenance of the system becomes quite difficult, because changes in one code can affect directly the entire program. In practice, most multiphysics packages are made up from codes developed by different teams, often in different institutions, and therefore updating the coupled tool to new code releases becomes complicated. As a result, besides a few well-established exceptions, master-slave couplings are quickly deprecated, typically after a few years (or after the completion of a PhD project).

Besides these design issues, the data exchange and calculation control can be hard to implement using master-slave coupling, in particular when more than two codes are being coupled. When data is exchanged directly, the master code has to be aware of the internal data structures of every code and the mapping between them, which can become very complicated. This is particularly true when the codes are written in different programming languages, which is the case for Serpent (written in C), SCF (Fortran90) and TU (Fortran95).

These issues underline the need to apply a coupling strategy based on good programming and design practices. In this work, an object-oriented mesh-based approach is used to implement the three-code coupling.

3.3.3 Object-oriented design

The software design of the coupling system developed in this work is based on inheritance, a key feature of object-oriented programming. Each code is implemented as a C++ solver class derived from a common base class, defining a standard coupling interface and masking the internal calculation methods, data structures and programming languages. The actual coupling scheme is specified in a C++ supervisor program using object-oriented features. The implementation and verification of Serpent-SCF and Serpent-SCF-TU coupling schemes using this methodology has been thoroughly described in publications related to this thesis [69], [68], [70], [66], [71].

3.3.3.1 ICoCo-based modules

The approach used to modularize Serpent, SCF and TU follows a well-defined methodology, as schematized in Figure 3.2. Each source code is first restructured as a solver library that implements the capabilities required in multiphysics simulations, i.e. initialization, termination, time-step control, calculation control for steady-state, depletion and transient calculations and exchange of feedback variables. To define the Application Programming Interface (API) for each code, these libraries are then wrapped into C++ classes derived from an abstract `Problem` base class that represents a generic code suitable for multiphysics simulations. This base class is defined by the Interface for Code Coupling (ICoCo) specification from the SALOME open-source platform [72]. At this point, the three codes are implemented as solver C++ classes with a common format, with the native language masked by the C++ interface.

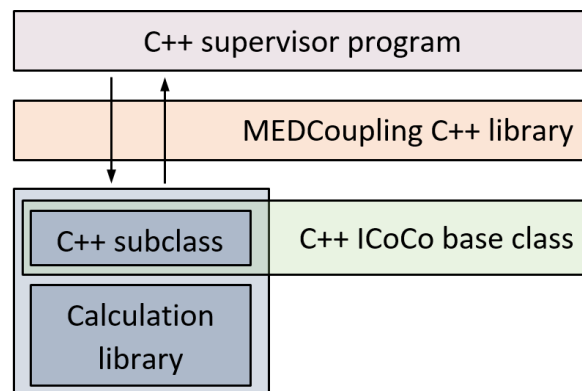


FIGURE 3.2: Software design of the calculation modules [66].

A key aspect of the API design is the use of a common format for the exchange of feedback variables, which is done using the MEDCoupling library [73], also from the SALOME platform. This library implements unstructured meshes and fields defined on them, as well as advanced interpolation methods and output capabilities. When a variable is retrieved from or set to a code the data is represented as a field defined on a mesh. The exact way in which variables are exchanged and interpolated is described in Section 3.3.4.

With this methodology, the three codes involved in the coupling have identical interfaces, i.e. C++ classes derived from a common base class and therefore with the exact same calculation methods and format for feedback exchange. This greatly simplifies the implementation of the coupling scheme, since all the codes behave in the same way and the mapping of feedback variables between the codes is done using mesh-based interpolation at runtime and not through index-based mappings. It is also key to note that up to this point the three codes remain completely separate, and thus can be maintained independently without affecting the coupled system as long as the ICoCo-based multiphysics interface is kept constant.

3.3.3.2 ICoCo-based supervisor

Figure 3.3 shows the overall design of Serpent-SCF-TU, which relies on modularization and object-oriented programming. With the three codes wrapped into interfaces as explained in the previous section, the coupling scheme is implemented in a C++ supervisor program that manages the multiphysics calculation scheme, the feedback exchange and the output of the simulation.

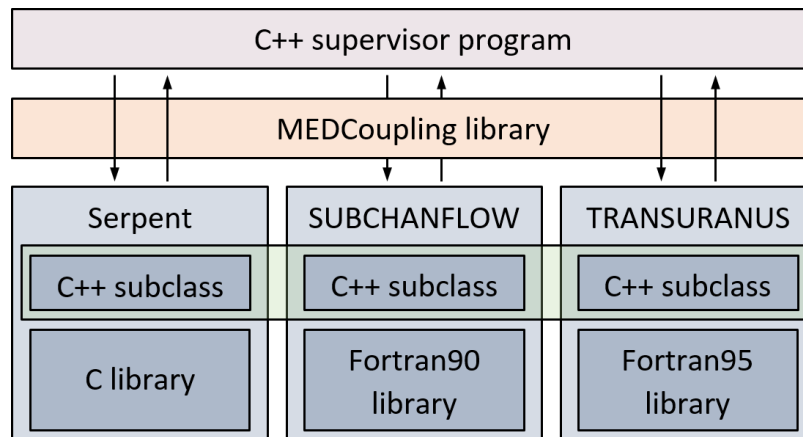


FIGURE 3.3: Software design of the full coupling [66].

Given that all the APIs are derived from the same base class, the supervisor program is composed of a set of classes that deal with the `Problem` base class only and implement the main methods, e.g. feedback exchange, calculation schemes and convergence control, in a problem-agnostic way, i.e. independent of the particular codes being coupled. This is a result of the inheritance-based definition of the APIs, and greatly enhances code reusability. In fact, a single supervisor program can be used to run any set of codes, provided that they have ICoCo-based interfaces. This might seem like a technicality, but it is a key feature which becomes more useful the more codes are involved. Because the calculation logic is not tied into a specific code any solver can be replaced right away provided the new code has an ICoCo interface. As an example, to replace Serpent by a nodal code such as PARCS the only effort required is restructuring PARCS as an ICoCo-based module, and the integration into the coupling scheme is essentially automatic. In the same way, new physical domains, e.g. plant analysis or CFD capabilities for multiscale modelling, can be readily added to the system. This convenient feature outlines how suitable this strategy is for long-term software development.

In the case of this work, the calculation schemes for Serpent-SCF and Serpent-SCF-TU are run using the same supervisor. Since the main program only knows

about the `Problem` class, the feedback variables and the calculation scheme are specified using a simple configuration file. Hence, even though the calculation scheme shown in Figure 3.1 is used throughout this thesis, this is not fixed and any algorithm can be used.

Regarding the input files, the model for each code is defined in the same way as in the standalone versions, i.e. the same input formats are used. The meshes used to transfer fields are also given as input, and are built with a special preprocessor, which also generates the geometrical model of the three codes. The coupling scheme is defined in the supervisor input file, where the calculation methods, time discretization, feedback fields, convergence tolerances and acceleration methods are specified.

3.3.4 Mesh-based feedback exchange

The feedback exchange for all fields works in the exact same way, and can be summarized as:

1. Get the field with its associated mesh from the code that calculated it.
2. Evaluate convergence, change units, accelerate, generate output.
3. Get a template of the field with its associated mesh from the target code.
4. Interpolate the field from the source mesh to the target mesh.
5. Set the interpolated field to the second code.

In this scheme the classes that define the fields and meshes, as well as the interpolation methods, are provided by the MEDCoupling library.

The interpolation is based on mesh overlapping, i.e. the source and target meshes are superimposed and the mapping coefficients are calculated based on the overlapping volumes. To apply a consistent remapping, suitable interpolation constraints are applied to each field according to three strategies [73]:

- Conservative volumic: intensive fields, e.g. T_{cool} , T_{fuel} and $h_{clad-cool}$.
- Integral with global constraints: extensive fields with preservation of the integral value, e.g. P .
- Reverse integral: intensive fields representing volume densities, e.g. ρ_{cool} .

The idea of this coupling approach is that, since all codes exchange data in the same format, the supervisor program does not need to be aware of the particular data structures in each code. The mapping between the MEDCoupling fields and the internal model is handled by each code in a different way depending on the internal data representation. Appendix A provides an example of the handling of the meshes in Serpent, SCF and TU for a VVER-1000 fuel assembly.

Chapter 4

Optimizing the system for full-core pin-by-pin depletion problems

The solution of full-core pin-by-pin depletion problems using the Monte Carlo method leads to massive computational requirements [74]. On the one hand, the number of particles per transport cycle needed for a reasonable statistical uncertainty in the power distribution and reaction rates is of the order of 10^9 , leading to huge runtimes even in HPC systems. To tackle this issue, Serpent features hybrid MPI-OpenMP parallelization, which allows the efficient use of the resources available in most HPC architectures. On the other hand, the size of the problem in terms of memory, driven mainly by the storage of burnable material data, reaches a few terabytes, which exceeds the memory available in typical computing nodes. This last point creates a clear bottleneck for the solution of large-scale burnup problems, as explained in detail in Section 4.1, since the standard parallel scheme used in Monte Carlo particle transport has no memory scalability.

To solve this memory bottleneck, a Collision-based Domain Decomposition (CDD) scheme for Serpent has been formulated, implemented and published as part of this thesis [75], [76]. The main aspects of the methodology, i.e. the decomposition of the geometry and the tracking algorithm, are explained in Section 4.2. Since the topic of this chapter is the CDD scheme, the analysis is focused on the Monte Carlo calculations independently of the multiphysics system. A brief discussion on the performance of the full scheme is provided in Section 4.3.

4.1 Parallel schemes for Monte Carlo particle transport

The traditional approach to parallelize Monte Carlo particle transport calculations is based on the fact that particle histories are independent and can therefore be simulated simultaneously. This leads naturally to a scheme based on splitting the particle histories across execution units, and is commonly termed particle-based parallelism. Here, each process simulates a subset of the total number of histories and the results are combined after each transport calculation. Given that the problem is inherently parallel, this method can lead to excellent speedups.

Now, most Monte Carlo codes implement particle-based parallelism using domain replication, which means that all execution units have all the information about the system (geometry, materials, cross sections, etc). This is a convenient approach for most applications, because no communication across processes is needed during tracking, improving the parallel efficiency. The drawback is that this method has no memory scalability, meaning that adding computing resources to the system does not increase the maximum size of the problem that can be simulated. This is due to

the fact that, since the problem is replicated, it needs to fit in the memory available to each process, independently of the number of processes being used.

To produce Monte Carlo transport codes with memory scalability, several schemes have been proposed and implemented. These can be grouped in two broad categories: data decomposition and domain decomposition, which are summarized in sections 4.1.1 and 4.1.2 respectively. The method proposed in this work is a combination of these two approaches, as explained in Section 4.2.

Serpent, like most Monte Carlo codes, uses an OpenMP-based shared-memory approach for in-node parallelism and MPI, a distributed-memory paradigm, for node-to-node communications. The discussions that follow refer to the MPI layer of the implementation, which determines the memory scalability of the code. A one-to-one mapping between MPI tasks and physical nodes, which is typically the case, is assumed, and therefore the terms process, task and node are used interchangeably. The OpenMP layer determines the in-node performance but does not significantly affect the memory scalability, since all threads share a common memory space and do not need separate copies of the problem. The interaction of the MPI and OpenMP layers is discussed in Section 4.2.2.

4.1.1 Data decomposition

The idea of data decomposition is to distribute memory-intensive data such as geometry structures, cross sections, tallies and material information across nodes. Particle-based parallelism is still used, and remote data is retrieved on demand during tracking when a task needs information that is not stored locally.

In this scheme, memory scalability is driven by the decomposition of data structures and the performance in terms of speedup is determined primarily by the algorithm used to communicate data across tasks during particle tracking. In previous works this last issue has been addressed using methods such as tally servers [77] and global view arrays [78] in the OpenMC code. Implementing an efficient data-transfer model and combining it with in-node multithreading is not trivial however, which is the main drawback of this methodology.

4.1.2 Domain decomposition

Domain decomposition consists of partitioning the model into subdomains assigned to different tasks, which store only the local geometry, cross sections, tallies and materials. Each task then tracks the particles born locally, in addition to the ones that enter its domain, and particles are transferred across tasks as they cross domain boundaries.

In this case, the geometry partition leads naturally to memory scalability, as each node stores only one subdomain, and the performance depends on the efficiency of the particle communication scheme for the tracking algorithm. The most complex issue in this approach is the partition of the problem geometry, which in most Monte Carlo codes is represented using CSG [79]. While the particle tracking has to rely on asynchronous MPI communications to achieve good performance and its implementation is not trivial, this problem has been studied extensively and efficient algorithms have been formulated and published [80], [81].

4.2 Collision-based Domain Decomposition

The Collision-based Domain Decomposition scheme implemented in Serpent is essentially a data decomposition method with a domain decomposition tracking algorithm.

The data decomposition is done material-wise and is focused on depletion. Each domain contains a subset of the burnable materials for which it stores all data, including compositions, material cross sections and tallies. In this way, the depletion data, which constitutes the largest portion of the memory demand in burnup calculations, is distributed across tasks, reducing the in-node memory. All non-burnable materials, along with the rest of the model (geometry, nuclide cross sections, spatial tallies, etc) are replicated in all domains. Avoiding dividing non-burnable materials reduces the number of particles that need to be transferred across domains, as explained next, and has a negligible impact on the memory scalability. Figure 4.1 illustrates this type of decomposition for a PWR fuel assembly. The way in which the material decomposition is performed is explained in detail in Section 4.2.1.

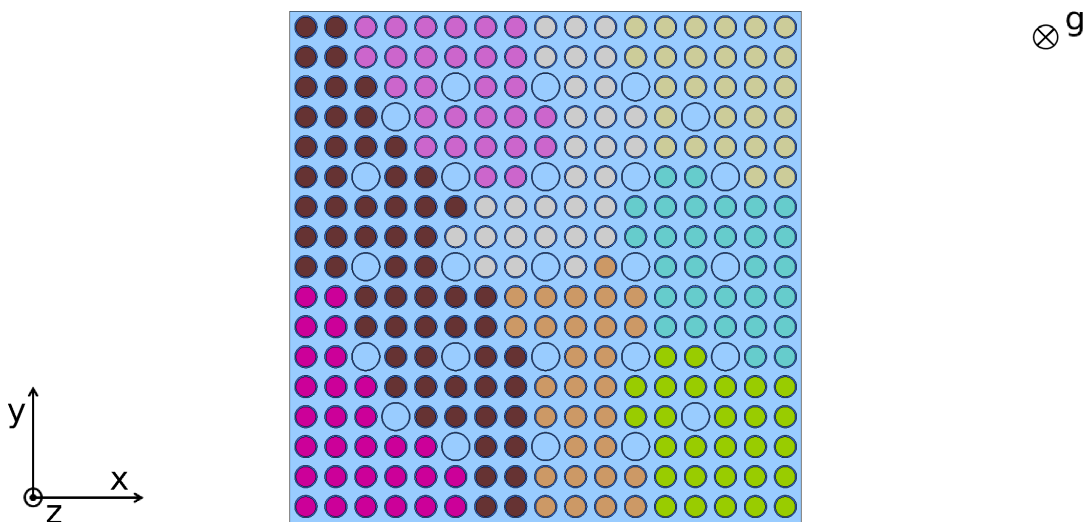


FIGURE 4.1: Material decomposition for a 2D model of a PWR fuel assembly (21.5 cm fuel-assembly pitch, 1.26 cm pin pitch). The fuel material is decomposed in 9 domains shown with different colors in the fuel pins, while the cladding and coolant materials are replicated in all domains.

As in all data decomposition schemes, a task can track a particle until non-local data is needed, in which case the standard approach is to fetch that data from a remote domain to continue the tracking. In the CDD method, the particle is sent to the corresponding task instead, as in standard domain decomposition schemes when particles cross domain boundaries. A particle can stream through local fuel materials and non-burnable ones, and even through non-local fuel materials in delta-tracking mode [82] as long as there is no interaction. Particle transfers are triggered when a collision in a non-local material occurs, hence the collision-based aspect of the approach. The two key issues for an efficient algorithm are the particle tracking and termination control methods, which are described in sections 4.2.2 and 4.2.3 respectively.

The implementation presented in this work is based on a one-to-one mapping between domains and MPI tasks, i.e. each domain is assigned to a unique task. In principle this is not mandatory, since a domain could be handled by more than one

MPI task, combining domain decomposition and replication. However, this leads to a more complex implementation and offers no clear benefits, since the optimum execution in terms of memory use is to have one node/task per domain and the performance would not be improved, since all tasks have to be included in the CDD communication scheme independently. Furthermore, the material decomposition is performed at runtime based on the number of MPI tasks, and therefore the number of domains is adjusted automatically to the number of tasks.

4.2.1 Material decomposition

The most common way to define burnable materials in Serpent is using automatic divisions, which create depletion zones from user-defined materials. Divisions can be universe-based, for example at nodal or pin level, as well as radial and azimuthal for fuel pins. The material division can be arbitrarily detailed, and is typically bounded by the memory demand. In the CDD scheme, domains are defined by a set of burnable materials, in order to avoid performing an actual partition of the CSG-based geometry, which was cited as a chief issue in domain decomposition methods.

Now, the performance of the tracking algorithm will in principle depend on two main factors: the number of particles tracked in each domain, i.e. the source distribution, and the number of particles flying across domains, i.e. the burden on the particle communication algorithm. The source distribution is both unknown a priori and dependent on burnup, and thus can only be treated with dynamic load balancing, which is out of the scope of this work. Therefore, assuming a uniform particle source, it is reasonable to try to define domains as compact as possible to minimize the number of particle transfers.

Serpent provides various simple methods to divide the burnable materials into domains, using material indexes or two-dimensional radial sectors. A more advanced method, used in this study, relies on graph-partitioning. The weighted graph representation of the set of materials consists of:

- V : the vertices of the graph, v_i is material i .
- E : the edges of the graph, e_{ij} is the link between materials i and j .
- W^V : the weights of the vertices, w_i^V is the weight of material i .
- W^E : the weights of the edges, w_{ij}^E is the weight of link ij .

The material volumes are used as weights w_i^V for the vertices, to keep the domain volumes equal. The edges are weighted with the inverse of the distance between materials, such that

$$w_{ij}^E = \frac{L}{l_{ij}}, \quad (4.1)$$

where L is a representative length of the system and l_{ij} is the distance between materials i and j . This just means that the closer two materials are the tighter the link is, to generate compact domains. A user-given cutoff value for l_{ij} is used to discard edges for materials that are too far apart, in order to reduce the graph size and therefore the memory requirement and the partition time. Once the material graph is constructed, the partition is done using the Metis graph-partitioner library [83] to obtain decompositions like the one shown in Figure 4.1.

It is important to note here that the CDD feature does not introduce any approximations from the physical point of view, and the results are expected to be equivalent within the statistical variability of the solution. The consistency of the solution is ensured by transferring the particles directly across domains without imposing physical approximations such as flux or current continuity constraints. The material decomposition has no impact in terms of modelling, only in terms of performance.

Furthermore, previous publications propose the use of overlapping domains, extending each domain by a certain width that includes non-local materials while restricting the fission source to the original region [84]. This would decrease the amount of domain crossings and potentially improve the performance, though it has not been included in the CDD implementation.

4.2.2 Particle tracking

During tracking, particles that encounter collisions in non-local materials have to be transferred between MPI tasks. To do this efficiently, all domain decomposition implementations use asynchronous MPI communications, which means that particles are sent without waiting for them to be received, thus avoiding synchronizing MPI tasks for particle transfers. Particles are sent and received using the `MPI_Isend()` and `MPI_Irecv()` standard nonblocking functions, along with the `MPI_Test()` and `MPI_Wait()` help functions [85]. Furthermore, particles are buffered and sent in packages instead of individually. Now, since the completion of these send operations is not ensured directly, an indirect check needs to be in place to determine when a transport cycle is over, which is the topic of Section 4.2.3.

The actual tracking scheme is composed of the MPI layer of the implementation, which determines the node-to-node communications, and the OpenMP layer for in-node multithreading. In this regard, two hybrid MPI-OpenMP modes have been implemented in Serpent: funneled and serialized multithreading. These execution modes differ in how the MPI and OpenMP layers interact, which in turn leads to slightly different tracking algorithms, which are described in the next two sections.

4.2.2.1 Funneled multithreading

The most straightforward way to treat multithreading with domain decomposition is to perform all MPI-related work outside the OpenMP-parallel section of the tracking scheme, i.e. having only the main thread make MPI calls. In this funneled implementation, described in Algorithm 1, particles that need to be sent to other domains are put in a "limbo" buffer during tracking. Only after all local particles have been tracked, the master thread sends all outgoing particles and receives the ones coming from other domains. This is repeated until the transport cycle is completed.

4.2.2.2 Serialized multithreading

In the funneled multithreading scheme the communications are performed in a sequential manner within each MPI task, once all local particles are tracked. Hence, it is reasonable to think that better performance could be achieved doing this work within the tracking loop, in order to maximize the communication/computation overlap. The performance gain will in principle be a trade-off between the improvement in the algorithm and the additional overhead driven by thread synchronization and enhanced thread-safety level of the MPI library (`MPI_THREAD_SERIALIZED` is required [85]).

```

while transport cycle not completed do
  begin multithreaded section
    while local particles left do
      track particle;
      if collision in another domain then
        | put particle in limbo;
      end
    end
  end
  send all particles in limbo;
  receive particles from all domains;
  if no new particles then
    | check transport cycle completion;
  end
end

```

Algorithm 1: Tracking using funneled multithreading.

Algorithm 2 describes the scheme with multithreaded MPI communications. Particles are sent and received within the OpenMP-parallel tracking loop, though this occurs in an OpenMP-critical section, i.e. serialized. Once a send buffer is full, the asynchronous MPI message is sent to the corresponding domain and particles coming from that same domain are received and put in the limbo buffer. Once the tracking of local particles is over, particle transfers with all domains are checked and limbo particles are put in simulation queues.

Performing the particle transfers during tracking has two advantages. First, MPI communications can be done by one thread while the rest are tracking particles, which increases the overlap between communication and computation. Second, asynchronous MPI send and receive operations are done more often, which reduces the probability of exhausting MPI resources such as memory pipes and buffers. This last point is addressed in previous domain decomposition implementations with a fixed frequency at which asynchronous receives are checked, though it is not reported how multithreading is handled in this case [80].

4.2.3 Termination control

A transport cycle is completed when all particle histories in the whole system have been simulated. Although this seems obvious, since particles are sent asynchronously each task cannot determine on its own when it is done. A task can run out of local particles and not detect any incoming messages while live particles that can later be sent its way exist in other domains. For this reason, all domain decomposition schemes require a termination control method based on a global check of the particle balance.

In the CDD scheme proposed here, termination is checked when domains have run out of local particles to track, as can be noted in algorithms 1 and 2. If all tasks are synchronized at this point and it is determined that all particles sent have been received, this means that all particle histories started have been completed and therefore that the cycle is finished. The stopping criteria is $n = 0$, where n is the difference between the number of particles sent and received.

```

while transport cycle not completed do
  begin multithreaded section
    while local particles left do
      track particle;
      if collision in domain  $D_i$  then
        begin critical section
          add particle to buffer  $i$ ;
          if buffer  $i$  full then
            send particles to  $D_i$ ;
            receive particles from  $D_i$ ;
            if new particles then
              put particles in limbo;
            end
          end
        end
      end
    end
  end
  flush all half-full particle buffers;
  receive particles from all domains;
  if no new particles then
    check transport cycle completion;
  else
    move particles from limbo to queues;
  end
end

```

Algorithm 2: Tracking using serialized multithreading.

Now, this synchronized check is a very expensive operation and would kill the performance of the algorithm if performed after every tracking loop. To avoid doing this, an estimation using asynchronous MPI communications is performed instead, and the synchronized check is only done when this asynchronous balance is achieved.

In order not to overwhelm a particular MPI task with the communications related to the asynchronous estimation of the particle balance, tasks are organized in a binary tree as shown in Figure 4.2, in a similar way as in previous implementations [80]. The particle balance is collected by the main task (rank 0), which sends a signal s_p down the tree to request synchronization when $n = 0$ is estimated. At this point all tasks synchronize and the true particle balance is checked. This scheme is summarized in Algorithm 3.

4.3 Performance

Since the Monte Carlo calculation is by far the most expensive side of the multiphysics system in terms of runtime and memory demand, this thesis focuses on optimizing the Serpent depletion capabilities and not so much on improving the performance of SCF or TU. The overall performance of the multiphysics system is dominated by Serpent, which features hybrid MPI-OpenMP parallelism, allowing the code to use all the resources in a HPC architecture in an efficient manner. This

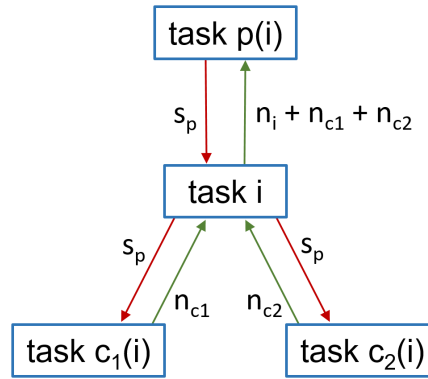


FIGURE 4.2: Binary tree structure used to estimate the particle balance. Task i collects the balance of its children tasks $c_1(i)$ and $c_2(i)$ and sends it to its parent $p(i)$, while the synchronization signal s_p is transmitted in the opposite direction.

```

for  $j = 1, 2$  do
  | receive balance  $n_{c_j}$  from child  $c_j(i)$ ;
end
send balance  $n = n_i + n_{c_1} + n_{c_2}$  to parent;
if  $rank = 0$  then
  | if  $n = 0$  then
  | |  $s_p = 1$ ;
  | else
  | |  $s_p = 0$ ;
  | end
end
receive  $s_p$  from parent;
for  $j = 1, 2$  do
  | send  $s_p$  to child  $c_j(i)$ ;
end
if  $s_p = 1$  then
  | perform synchronized global balance check;
end

```

Algorithm 3: Termination control scheme for task i (see Figure 4.2).

is not the case for SCF and TU, which have strong limitations regarding parallel calculations, degrading the performance of the system to some degree.

4.3.1 Serpent

The most important development in terms of performance has been the implementation of the CDD scheme in Serpent, since this feature offers a solution to the memory bottleneck in burnup problems. Without the CDD method the burnup calculations performed in this work are simply not possible, because the models do not fit in the in-node memory available in HPC systems.

A detailed analysis of the performance, both in terms of speedup and of memory scalability is presented in Appendix B. The performance of the CDD scheme in PWR-like infinite systems is shown to be quite good, with between 40% and 50% efficiency at 256 domains/nodes (5,120 cores). The memory scalability of the new

scheme allows for the simulation of increasingly larger systems by multiplying the computational resources.

The implementation of CDD in Serpent paves the way for the solution of full-core pin-by-pin burnup problems coupled with thermalhydraulics and fuel-performance codes. It also represents the last major development and optimization stage for Serpent-SCF-TU before the validation phase with PWR and VVER plant data.

4.3.2 SUBCHANFLOW

The pressure-velocity solver in SCF is based on solving the 2D pressure matrix for each axial plane from core inlet to outlet assuming an upward flow pattern. This reduces the total computational cost, because the solution of the 3D pressure matrix is avoided. The problem with this approach is that the full 3D system, which would be feasible to parallelize efficiently, is replaced by a set of smaller 2D systems for which typically the parallel performance is much worse. In full-core problems the size of the 2D problem is the number of subchannels, which can range from 50,000 to about 100,000. This is clearly not an adequate size to achieve a good efficiency in a system with around 1,000 cores, which is the target in this thesis. An efficient parallelization of SCF would necessarily have to include switching to a full 3D solver. However, not only this is not clear from the numerical point of view, since the code is formulated based on the current method, but is outside the scope of this work.

In the current scheme, SCF is run sequentially, which clearly leads a loss of performance. As an example, 1,280 cores are used in the calculations presented in the next two chapters. While Serpent can use all the cores, during the SCF calculation only one core is used and 1,279 have to sit idle. However, even in this extreme case each Serpent transport calculation takes around 3.7 hours in real wall-clock time while each SCF iteration takes about 9 minutes, less than 4% of the total runtime.

As an alternative option for cases in which the relative time spent by SCF becomes a problem, a subchannel-coarsening methodology has been proposed, implemented and tested [86]. This method has been shown to reduce the SCF calculation time up to a point at which it is essentially negligible, at the expense of a loss of detail in the coolant solution. Comparing solutions for a VVER-1000 full-core case using a coarse SCF model against a full subchannel model, it has been shown that local deviations in the coolant temperature reach about 10 K and lead to perturbations in the pin power around 4% [86]. While this method can be very useful for some applications, only standard subchannel models have been used in this work.

4.3.3 TRANSURANUS

The TU module used in this work can deal with multiple fuel rods, for which the solution is totally independent. Hence, the TU problem can be parallelized right away, splitting rods across execution units. The problem in this case is that the implementation is based on a pure-MPI model. Without going into detail, this means that the TU problem can be split across nodes in the same way the Serpent materials using CDD are, but within each node TU can only use one core. This reduces the in-node memory demand, because each node only stores the data for a subset of the fuel rods in the model. However, without OpenMP parallelization most of the cores in each node cannot be used. In the HPC system used in this work, made up of 20-core nodes, this means that the peak performance is reduced by a factor of $1/20$. This

is a huge loss of performance for TU, but it is not a major issue for the overall performance of the system. In the validation calculations shown in the next chapters, TU takes about 0.2% of the calculation time.

Chapter 5

Validation for a Pre-Konvoi PWR reactor

This chapter is focused on the validation of Serpent-SCF-TU pin-level burnup capabilities for a German Pre-Konvoi PWR NPP [87]. For the beginning of the first operating cycle, the numerical results are compared against measured critical boron concentration and neutron flux profiles. To assess the impact of the fuel solution method, the results obtained performing the thermomechanic analysis with TU, i.e. with the three-code coupling, are compared with the ones using Serpent-SCF without TU, solving the fuel behavior with the simplified approach included in SCF.

The Pre-Konvoi PWR plant used as validation case, as well as the operational data and the experimental measurements, are presented in detail in Section 5.1. The modelling approach is discussed in Section 5.2, while the results are presented and analyzed in Section 5.3.

5.1 Pre-Konvoi PWR validation data

The first validation case considered in this thesis is a German Pre-Konvoi PWR NPP, the predecessor of the Konvoi design [88]. The operational data corresponds to the beginning of the first operating cycle, where the experimental measurements were performed as part of the commissioning phase of the reactor. The complete specifications of the core and the operating history, as well as the measurements, were provided by PreussenElektra, the utility which operated the reactor, in the framework of the EU Horizon 2020 McSAFE project [89].

5.1.1 Core description

The reactor core consists of 193 fuel assemblies of 6 different types organized as shown in Figure 5.1, with a fuel-assembly pitch (center-to-center distance) of 23 cm. The fuel material is UO_2 with 1.9%, 2.5% and 3.2% enrichment for types 1 (1B), 2 (2B) and 3 (3B) respectively. Fuel types 1B, 2B and 3B have boron silicate control elements permanently inserted during the first operating cycle. The main core parameters at Hot Full Power (HFP) conditions are shown in Table 5.1. The total bypass flow includes all the flow that goes through the guide tubes (see Figure 5.2) plus the flow between the core baffle and the barrel. The core baffle, barrel and reactor vessel material is stainless steel.

Each fuel assembly is composed of a 16x16 arrangement of 236 fuel rods and 20 guide tubes, with the pin layout shown in Figure 5.2. The pin pitch is 1.43 cm, the pellet diameter is 9.1 mm and the active length is 390 cm, i.e. the power is generated in the $z = 0\text{-}390$ cm axial region. The cladding material is Zry-4 and the guide tubes are made of stainless steel. Fuel assemblies have 9 spacer grids: 7 equidistantly




FIGURE 5.1: Core fuel-assembly loading (xy-cut) for the Pre-Konvoi reactor (23 cm fuel-assembly pitch). The fuel types are defined by their enrichment and the presence of boron silicate rods.

spaced within the active length, one at the top and one at the bottom. The former have mixing vanes, while the latter do not. The spacer-grid hydraulic data has been provided by the reactor operator.

Boron silicate rods have a total length of 381 cm, with the cross sections shown in Figure 5.3. The upper tip of the rod is 50 cm long and contains a stainless-steel plug. The lower tip, which is 23 cm long and starts 8 cm above the bottom of the active length, corresponds to a rod cup filled with air which is a continuation of the cladding. The central part is a boron silicate glass tube.

The control rods have Ag-In-Cd as absorber material and a total axial length of 414 cm, with 372 cm of absorber region and a 42-cm-long bottom tip. The cladding and the bottom tip are made of stainless steel. There are 61 control rods in total, organized in banks D0 to D4 and L0 as shown in Figure 5.4.

5.1.2 Operational data

Figure 5.5 shows the thermal power P , critical boron concentration B and control bank positions for the beginning of the first burnup cycle of the NPP. Given the strong variations in the power during the startup of the reactor, an effective full-power operational history with condensed data is considered, in particular for control rods. The calculated critical boron concentration is compared with the experimental measurements at the last six time points (see Figure 5.5b), where the power has been stable long enough that the xenon can be considered at equilibrium at full

Parameter	Value
Thermal power	3765 MW
Outlet pressure	15.8 MPa
Total mass flow rate	18875 kg/s
Inlet temperature	293.4 °C
Total bypass flow	6%

TABLE 5.1: Main operating parameters at HFP for the Pre-Konvoi reactor.

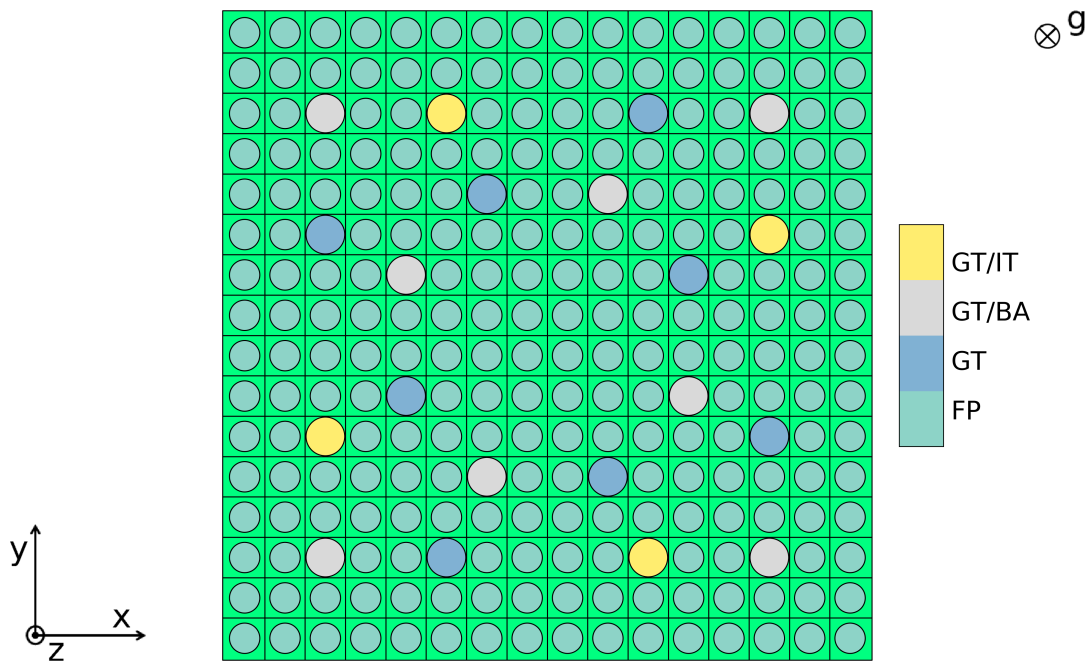


FIGURE 5.2: Pin layout (xy-cut) for all fuel assemblies for the Pre-Konvoi reactor (23 cm fuel-assembly pitch, 1.43 cm pin pitch). Fuel pins (FP), guide tubes (GT), boron absorber rods (BA) and instrumentation tubes (IT) are included.




FIGURE 5.3: Boron rod design (xy-cut) for the Pre-Konvoi reactor: upper tip ($z = 331\text{-}381$ cm, left), lower tip ($z = 8\text{-}31$ cm, center), central part ($z = 31\text{-}331$ cm, right). Boron silicate is shown in yellow and stainless steel in gray. The radial dimensions are proprietary parameters.

power. The flux profiles are compared with the measurements at the last time point. The condensed control bank positions used for the burnup calculation are shown




FIGURE 5.4: Control-bank layout (xy-cut) for the Pre-Konvoi reactor (23 cm fuel-assembly pitch).

in figures 5.5c, 5.5d and 5.5e. For the period up to 120.65 days, or 40 Effective Full Power Days (EFPD), where the power is not stable, a constant insertion is used, calculated as an average over time weighted with the thermal power. After that, the positions at the end of each burnup step are used in order to calculate the boron concentration correctly.

5.1.3 Experimental data used for validation

Two sets of experimental data are used in this work to evaluate the accuracy of the numerical calculations: critical boron concentration and direct pin-level neutron flux measurements at the end point of the data set. The critical boron (Figure 5.5b) comes from boric acid (H_3BO_3) concentration measurements during reactor operation and the flux profiles correspond to ^{51}V Aeroball Measurement System (AMS) [90] readings. The AMSs are inserted in the instrumentation tubes (see Figure 5.2), with the radial distribution shown in Figure 5.6, where the indexing is defined. Axially, each detector consists of a stack of 32 aeroballs at equidistant positions along the active length of the core ($z = 0\text{-}390$ cm). The AMS measurements are presented in Section 5.3.2 along with the numerical results.

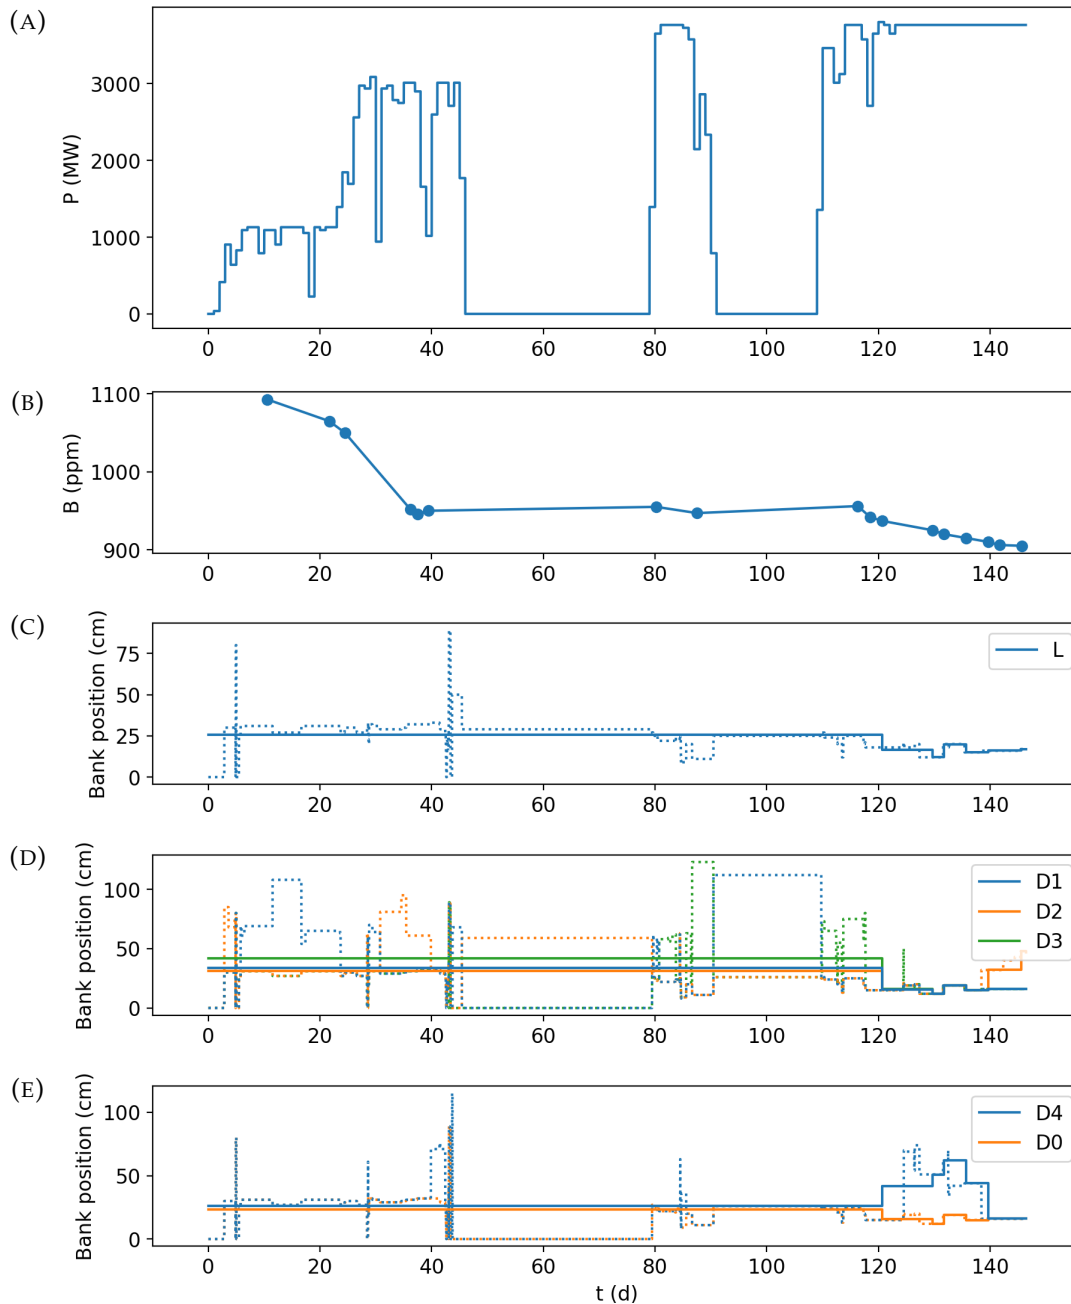


FIGURE 5.5: Operational history for the first cycle of the Pre-Konvoi reactor: (A) thermal power P , (B) measured critical boron concentration B and (C-E) condensed control bank positions (the raw operational data is shown in dashed lines).

5.2 Modelling approach

The burnup calculation has been performed at HFP conditions with the operating parameters summarized in Table 5.1. Uniform burnup steps of 8 EFPD are used until 40 EFPD with constant control bank positions, as shown in Figure 5.5. After that point, the solution is calculated at 49, 51, 55, 59, 61 and 65 EFPD, the last six points with critical boron measurements in Figure 5.5, using the control rod insertions at those time points. The calculated neutron flux profiles are compared with the experimental data at 65 EFPD.




FIGURE 5.6: Radial distribution of AMS stacks for the Pre-Konvoi reactor (23 cm fuel-assembly pitch). The detectors are inserted in the instrumentation tubes (see Figure 5.2).

The core analysis model is based on a pin-by-pin description of the whole system. The basic modelling unit is the fuel rod, for which the burnup and thermomechanic calculations are performed. The cooling conditions are resolved at subchannel level. As discussed in Section 2.4.5, the thermohydraulic boundary conditions for the system are the outlet pressure, the mass flow rate and the inlet coolant temperature, which correspond to the operating parameters summarized in Table 5.1. The neutronic boundary conditions are specified at the outer surface of the axial and radial reflectors, where a zero incoming current is imposed, as introduced in Section 2.1.2. The state of the core is also defined by the power level and the control rod positions (figures 5.5c, 5.5d and 5.5e). A detailed description of the Serpent, SCF and TU models is provided in Appendix C.

Regarding the neutronic calculation, the main modelling decision is the definition of the homogeneous material cells that define the system, in particular for the fuel regions where the evolution of the isotopic compositions is solved. For the pin-level model used in this work, the axial discretization is limited by the calculation time and memory demand of the simulation in relation to the computing resources. The CDD scheme used in Serpent to split the memory demand across computing nodes does not introduce additional approximations to the Monte Carlo particle tracking method. No variance reduction techniques are used and the source convergence is checked based on its Shannon entropy.

As for the thermomechanic calculation in the fuel, two levels of physical resolution are compared considering Serpent-SCF simulations with and without TU. While both SCF and TU rely on $1^{1/2}$ radial discretization schemes, TU performs a complete stress-strain and thermal analysis and takes into account a much greater number of physical phenomena. The thermohydraulic calculation is performed using a pin-level model with the maximum degree of resolution in the xy-plane allowed by the subchannel method. The axial discretization in both cases is bounded by the computational resources, in particular by the tallying of the power distribution in the Monte Carlo simulation.

The convergence tolerance ϵ for the multiphysics iterations at each burnup step is 1%, 1 kg/m^3 and 10 K for P , ρ_{cool} and T_{fuel} , respectively. The solution is considered converged when

$$\|x_i(\vec{r}) - x_{i-1}(\vec{r})\|_{L_2} < \epsilon, \quad (5.1)$$

taking the L_2 norm of the difference between the solution $x(\vec{r})$ between the last two iterations i and $i - 1$. This criterion is consistent with the stochastic variability of the results, and the solution typically converges after 3 or 4 iterations.

The simulations have been run in the ForHLR II HPC cluster of the KIT [91] using 64 nodes with 68 GB of in-node memory and 20 cores each (1,280 cores in total), with a maximum runtime of 7 days. The discretization of the model, in particular the number of burnable materials, as well as the number of particle histories per transport cycle and the number of burnup steps, are bounded by the computing resources provided by the ForHLR II system.

5.3 Selected results

To assess the accuracy of the simulation method, the critical boron concentration and neutron flux profiles calculated with Serpent-SCF with and without TU are compared against the experimental measurements. The boron concentration serves to verify the calculation of the core reactivity, which involves the burnup simulation with equilibrium xenon and the critical boron iteration, as well as the thermohydraulic feedback. The neutron flux measurements are used to evaluate the pin-level results, which are obtained directly as part of the Monte Carlo simulation.

5.3.1 Critical boron concentration

The measured and calculated critical boron concentration for the period between 49 and 65 EFPD is shown in Figure 5.7. As expected, the boron concentration decreases with burnup as the fuel is depleted. The Root Mean Square (RMS) difference between the results and the experimental data is 3.5 and 4.8 ppm for Serpent-SCF (S2-SCF) and Serpent-SCF-TU (S2-SCF-TU), respectively. The differences between the results with and without TU, which are lower than 4 ppm for this period, are not significant, which seems to indicate that the fuel solution method does not have a large impact on the core reactivity in this case.

The deviations with respect to the experimental data are well below the acceptability criterion for this type of calculation, which is usually between 50 and 100 ppm, and within the typical uncertainties due to engineering tolerances, e.g. physical dimensions, and operational measurement inaccuracies, as well as computational uncertainties arising from nuclear data libraries.

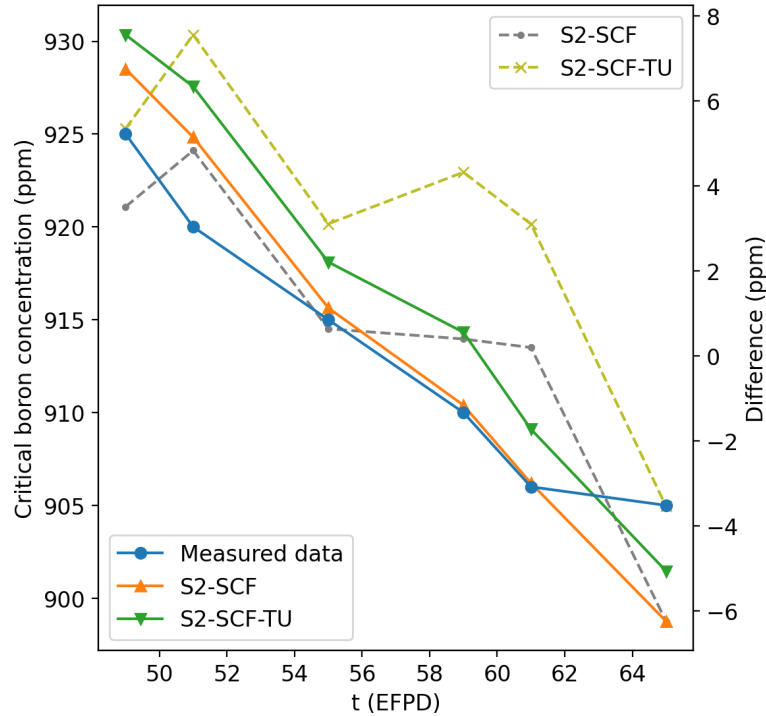


FIGURE 5.7: Comparison between measured and calculated critical boron concentration from 49 to 65 EFPD for the Pre-Konvoi reactor.

5.3.2 Neutron flux profiles

Figures 5.8 and 5.9 show the normalized neutron flux profiles at 65 EFPD for the two locations with the smallest and largest RMS deviations between the Serpent detector responses and the AMS measurements. The indexing is defined in Figure 5.6 and the error bars for the results correspond to three standard deviations $\sigma(z)$ for each axial location. An approximate value $\sigma = 3\%$ provided by the reactor operator is reported for the experimental error for the flux, for which a proper estimation is not available. The uncertainty in the axial position of the detectors is not available either. Both the calculated fluxes ϕ and the AMS measurements ϕ_{AMS} are normalized such that

$$\bar{\phi} = \bar{\phi}_{AMS} = 1, \quad (5.2)$$

where $\bar{\phi}$ and $\bar{\phi}_{AMS}$ are the mean values taken over all radial and axial positions. The neutron flux has a cosine-shaped axial profile typical of PWR cores at BOC with fresh fuel, with flux depressions at the axial positions with spacer grids.

The experimental measurements fall within the statistical range of the results for almost all data points, and no significant differences are observed between the calculations with and without TU, which is in accordance with previous studies [71]. Moreover, the locations with larger RMS deviations correspond to fuel assemblies in the core periphery where the statistical uncertainty is larger due to a lower number of particles. This behavior can always be improved increasing the amount of histories, at the expense of a higher computing cost. It is important to note here that the reported standard deviations for the results correspond only to the standard deviations in the transport calculation, and do not take into account other factors such as the burnup and thermalhydraulic calculations, which propagate stochastic errors between iterations. The total statistical uncertainty is actually higher, and in principle can only be obtained running the calculation multiple times to calculate

the overall deviations, which is not possible in this case due to the large runtimes involved.




FIGURE 5.8: Comparison between measured and calculated neutron flux profiles for the Pre-Konvoi reactor at the locations with the smallest RMS deviations at 65 EFPD. The aeroball positions are shown in Figure 5.6 and $z = 0$ corresponds to the bottom of the active length.




FIGURE 5.9: Comparison between measured and calculated neutron flux profiles for the Pre-Konvoi reactor at the locations with the largest RMS deviations at 65 EFPD. The aeroball positions are shown in Figure 5.6 and $z = 0$ corresponds to the bottom of the active length.

Table 5.2 summarizes the RMS deviations between the results and the measurements for all aeroballs. The relative differences $\epsilon(z)$ for each aeroball are calculated as

$$\epsilon(z) = \frac{\phi(z) - \phi_{AMS}(z)}{\phi_{AMS}(z)}, \quad (5.3)$$

Aeroball	Serpent-SCF	Serpent-SCF-TU
F15	4.8% (1.7 σ)	4.5% (1.5 σ)
G04	4.6% (2.2 σ)	3.7% (1.6 σ)
D09	4.4% (1.9 σ)	5.1% (1.9 σ)
O05	5.2% (1.8 σ)	5.3% (1.7 σ)
M11	4.6% (2.0 σ)	4.6% (1.9 σ)
G10	4.7% (2.2 σ)	4.5% (2.0 σ)
L04	4.1% (1.8 σ)	4.1% (1.7 σ)
D05	4.8% (2.3 σ)	4.3% (1.9 σ)
G14	5.6% (2.5 σ)	4.3% (1.9 σ)
K01	6.3% (2.3 σ)	5.7% (1.9 σ)
E08	5.6% (2.6 σ)	5.1% (2.4 σ)
N06	4.7% (1.9 σ)	4.4% (1.5 σ)
L10	4.6% (2.3 σ)	4.8% (2.3 σ)
C14	7.3% (2.1 σ)	6.9% (1.9 σ)
J06	5.0% (2.5 σ)	4.6% (2.1 σ)
E06	4.3% (2.0 σ)	4.2% (1.8 σ)
H13	4.7% (2.1 σ)	4.7% (1.9 σ)
J02	4.9% (2.2 σ)	4.1% (1.7 σ)
B11	5.3% (1.8 σ)	4.4% (1.5 σ)
M07	4.1% (1.7 σ)	3.8% (1.5 σ)
N12	4.4% (1.9 σ)	4.8% (2.0 σ)
E12	5.4% (2.4 σ)	4.2% (1.7 σ)
N02	5.4% (1.6 σ)	6.3% (1.7 σ)
C04	5.2% (2.2 σ)	4.6% (1.8 σ)
J12	4.9% (2.0 σ)	4.1% (1.9 σ)
H03	4.1% (1.7 σ)	4.1% (1.6 σ)
C10	4.6% (1.7 σ)	5.3% (1.8 σ)
L08	4.8% (2.3 σ)	3.7% (1.7 σ)

TABLE 5.2: RMS mean for the deviations $\epsilon(z)$ and $\epsilon_\sigma(z)$ between measured and calculated neutron flux profiles for the Pre-Konvoi reactor at 65 EFPD for all aeroballs.

for which the RMS mean is reported. Along with the relative errors, the errors normalized with $\sigma(z)$ for each axial position are included to put the deviations in terms of the statistical uncertainty. These are defined as

$$\epsilon_\sigma(z) = \frac{\phi(z) - \phi_{AMS}(z)}{\sigma(z)}. \quad (5.4)$$

In general the results are in quite good agreement, with the experimental data within the statistical range of the calculated values.

5.3.3 Fuel solution method

Figures 5.10, 5.11 and 5.12 show the fuel-cladding gap width w , the gap conductance h and the fuel centerline temperature T_{fuel}^{max} at the end of the burnup calculation (65 EFPD). The results using Serpent-SCF-TU are shown, as well as the differences for Serpent-SCF, i.e. using SCF instead of TU to handle the fuel solution. The relative

deviations $\epsilon(\vec{r})$ are calculated as

$$\epsilon(\vec{r}) = \frac{x_{S2-SCF}(\vec{r}) - x_{S2-SCF-TU}(\vec{r})}{\bar{x}_{S2-SCF-TU}}, \quad (5.5)$$

where $x_{S2-SCF-TU}$ and x_{S2-SCF} are the values with and without TU, respectively.

At this burnup level, the gap evolution is dominated by fuel relocation and densification, as well as thermal expansion of the fuel and cladding materials. As can be seen in Figure 5.10, the predicted gap width is smaller at the central region of the core, where the thermal expansion is larger due to the higher power level. For this sector, SCF predicts a gap width up to 49% larger than TU. While the pellet deformation calculated by TU is the result of the stress-strain calculation, SCF only estimates the total expansion of the pellet. Moreover, the relocation and densification of the fuel are considered empirically in both TU and SCF and are subject to large uncertainties. Hence, a deeper analysis of the differences in the gap solution is quite complex and out of the scope of this work. That being said, TU is one of the most well validated fuel-performance codes in the LWR industry, and therefore the TU results can be assumed to be more reliable than the SCF solution.




FIGURE 5.10: Calculated fuel-cladding gap width for the Pre-Konvoi reactor at 65 EFPD (2.35 MWd/kgU): (A) Serpent-SCF-TU results and (B) differences between Serpent-SCF-TU and Serpent-SCF.

The gap conductance is larger at the core central zone, as shown in Figure 5.11, since a reduction of the gap width decreases the thermal resistance of the fill gas. For this region, the gap conductance calculated by SCF is up to 116% lower, which is a result of the differences in the gap width calculation. For a given gap width, the methods used by TU and SCF to calculate the gap conductance are both semi-empirical and of similar accuracy.

The fuel centerline temperature, shown in Figure 5.12, follows the power distribution and is the result of the heat transfer through the fuel pellet, the gap and the cladding. The differences in centerline temperatures between TU and SCF reach 217 K, with maximum temperatures of 1730 and 1820 K for TU and SCF, respectively. These deviations are a result of the variations in the gap behavior, in addition to the correlations for fuel conductivity and heat capacity, which are different for TU

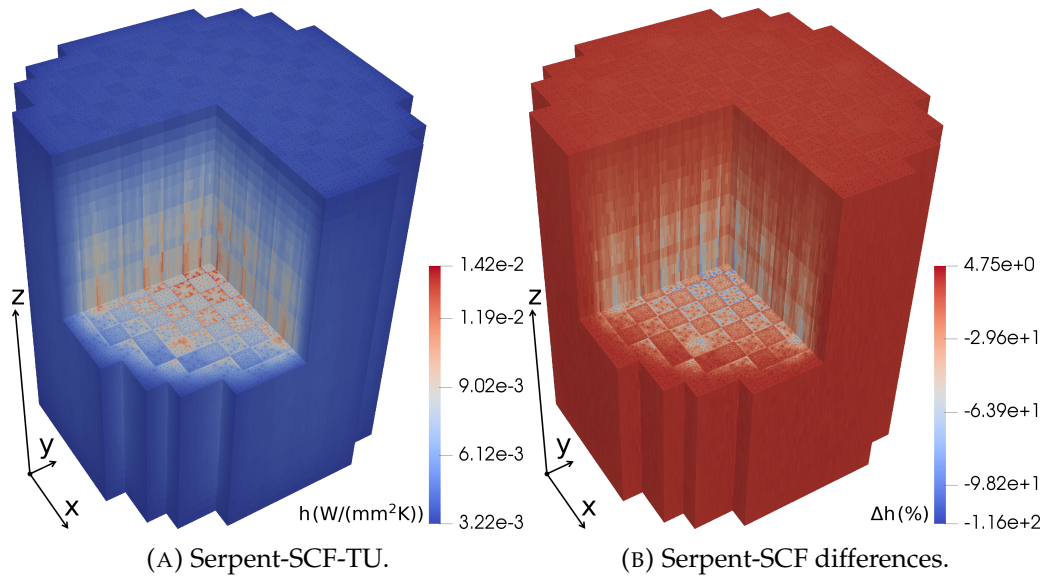


FIGURE 5.11: Calculated fuel-cladding gap conductance for the Pre-Konvoi reactor at 65 EFPD (2.35 MWd/kgU): (A) Serpent-SCF-TU results and (B) differences between Serpent-SCF-TU and Serpent-SCF.

and SCF. Furthermore, while these differences are significant in terms of fuel analysis, they do not have a large impact on the neutronic solution, as discussed in the previous section. The effect of the uncertainties in the gap width and conductivity and their impact on the fuel temperature is a critical factor in transient calculations, especially in Reactivity-initiated Accident (RIAs), where the main shut-down mechanism of the core is the negative reactivity produced by Doppler broadening in the fuel. Hence, even if the gap solution does not affect significantly the depletion calculation itself, improving the fuel model can lead to an increase in the accuracy of transient simulations, which require the core state during burnup.

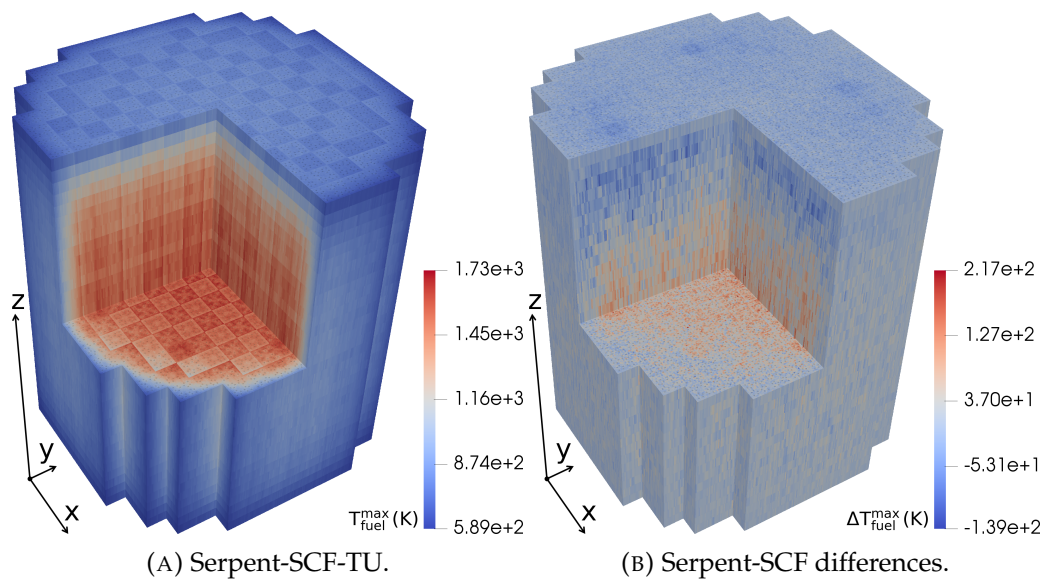


FIGURE 5.12: Calculated fuel centerline temperature for the Pre-Konvoi reactor at 65 EFPD (2.35 MWd/kgU): (A) Serpent-SCF-TU results and (B) differences between Serpent-SCF-TU and Serpent-SCF.

5.4 Validation summary

The calculated critical boron concentration for the Pre-Konvoi PWR NPP is in very good agreement with the experimental data, with deviations below 10 ppm between 49 and 65 EFPD. The pin-level neutron flux profiles obtained with Serpent are consistent with the aeroball measurements at 65 EFPD (2.35 MWd/kgU) within their statistical uncertainty. These results demonstrate the accuracy of the simulation tool for PWR cores, both in terms of global behavior and of local pin-level parameters.

The impact of the fuel solution method is assessed comparing the results for the three-code coupling with the traditional neutronic-thermalhydraulic scheme, i.e. without TU. Deviations up to 217 K can be observed in the fuel centerline temperatures, likely due to the differences in the fuel-cladding gap modelling and the conductivity correlations for the fuel. In previous studies on small-scale problems, these deviations have been shown to increase with burnup [66]. At about 52MWd/kgU, SCF can calculate fuel centerline temperatures up to 500K higher than TU, and TU typically predicts a closure of the fuel-cladding gap at much lower burnup [71]. However, the effect on the neutronics is not significant for the amount of burnup considered in this case. The advantage of using TU should be more clear at higher burnup due to phenomena such as the corrosion layer and the behavior of the high-burnup structure, and therefore the current work should be extended to higher depletion for full-core problems.

Chapter 6

Validation for the Temelín II VVER-1000 reactor

The focus of this chapter is the validation of Serpent-SCF-TU pin-level burnup capabilities for the Temelín II VVER-1000 NPP [92]. For the ninth operating cycle, which corresponds to the startup from a fresh core, the accuracy of the results is assessed using measured critical boron concentration and neutron flux profiles, as well as power distributions reported by the Core Monitoring System (CMS). Moreover, the advanced modelling capabilities of the three-code coupling are analyzed with pin-level neutronic, thermalhydraulic and thermomechanic results.

The Temelín II VVER-1000 plant used as validation case, the operational data and the experimental measurements are presented in detail in Section 6.1. The modelling approach is discussed in Section 6.2, while the results are presented and analyzed in Section 6.3.

6.1 Temelín II VVER-1000 validation data

The second validation case considered in this thesis is the Temelín II VVER-1000 NPP [93]. The operational data corresponds to the ninth operating cycle, where the experimental measurements were performed as part of the regular operation of the reactor. The complete specifications of the core and the operating history, as well as the measurements, were provided by ČEZ, the utility which operates the reactor, in the framework of the EU Horizon 2020 McSAFE project [89].

6.1.1 Core description

The reactor core consists of 163 TVSA-T [94] fuel assemblies of 8 different types organized as shown in Figure 6.1, with a fuel-assembly pitch of 23.6 cm. The fuel material is UO_2 with an enrichment between 1.3% and 4% and Gd_2O_3 as burnable poison. The main core parameters at nominal operating conditions at HFP conditions are shown in Table 6.1. The total bypass flow includes all the flow that goes through the guide tubes (see Figure 6.2) plus the flow between the core baffle and the barrel.

Each fuel assembly is composed of an hexagonal arrangement of 312 fuel rods, 18 guide tubes and a central instrumentation tube, with the pin layouts shown in Figure 6.2 for each fuel type. Figure 6.2 also shows the fuel enrichment in weight fraction of ^{235}U and the Gd content in weight fraction of Gd_2O_3 . All fuel types have 15 cm of top and bottom axial blankets of natural uranium, except for type a13A, which has a 1.3% enrichment in the whole active length. The pin pitch is 1.275 cm and the active length is 368 cm. The pellets have a 7.6 mm diameter and a 0.6 mm central hole. Fuel assemblies have 8 spacer grids: 7 equidistantly spaced within the




FIGURE 6.1: Core fuel-assembly loading (xy-cut) for the Temelín II reactor (23.6 cm fuel-assembly pitch). The fuel types are shown in Figure 6.2.

active length and one at the top. The spacer-grid hydraulic data has been provided by the reactor operator. The structural components include corner stiffeners welded to the spacer grids, a lower support node to fix the fuel pins and top and bottom nozzles. The cladding and spacer grid material is E110 zirconium alloy, while the guide tubes and stiffeners are made of E635 alloy.

The control rods have B_4C and Dy_2TiO_5 as absorber materials and an absorber length of 354.5 cm, with 324.5 cm of B_4C and 30 cm of Dy_2TiO_5 . The cladding is made of 42XHM alloy. There are 61 control rods in total, organized in banks 1 to 10 as shown in Figure 6.3. Banks 7 to 10 are used for power regulation, and the rest are used only for shutdown. During regulation close to full power only bank 10 is active, so the model considered in this work includes only this bank.

Parameter	Value
Thermal power	3000 MW
Outlet pressure	15.7 MPa
Total mass flow rate	17909 kg/s
Inlet temperature	288.9 °C
Total bypass flow	3.4%

TABLE 6.1: Main operating parameters at HFP for the Temelín II reactor.

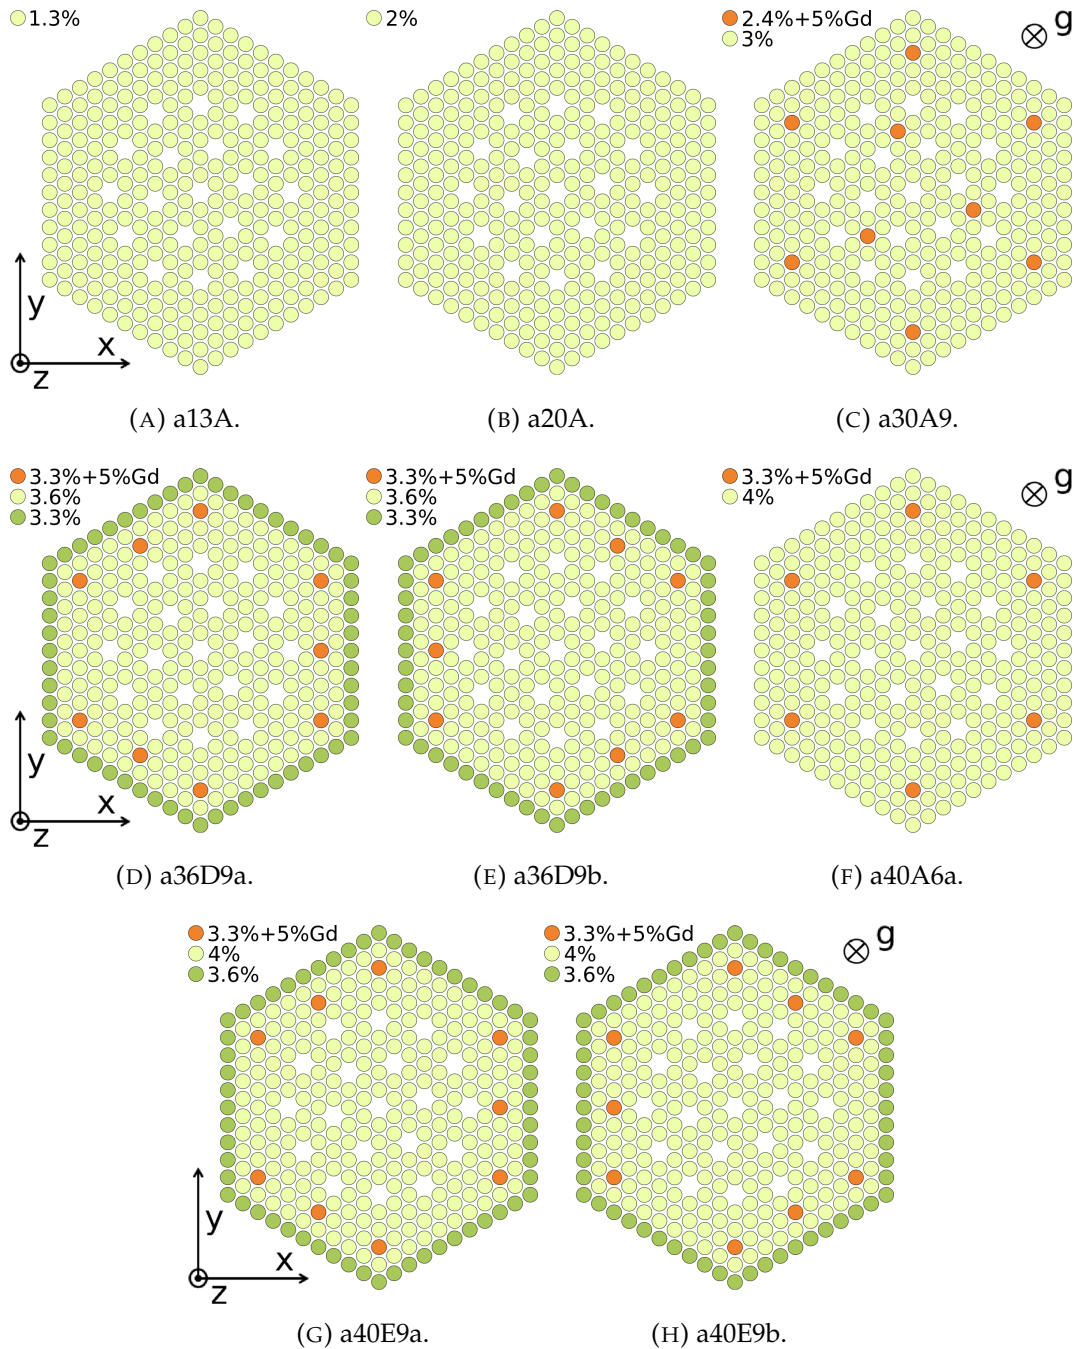


FIGURE 6.2: Pin layout (xy-cut) for each fuel type for the Temelín II reactor (23.6 cm fuel-assembly pitch, 1.275 cm pin pitch). The empty positions correspond to the guide tubes and the central instrumentation tube.

6.1.2 Operational data

Figure 6.4 shows the thermal power P , critical boron concentration B , control bank positions, inlet temperature T_{inlet} and mass flow rate \dot{m} for the ninth burnup cycle of the NPP, which corresponds to the startup cycle with TVSA-T fuel, with the whole core loaded with fresh fuel. Given that the power is close to the nominal value for most of the cycle, an effective full-power operational history with condensed data




FIGURE 6.3: Control-bank layout (xy-cut) for the Temelín II reactor (23.6 cm fuel-assembly pitch).

is considered, in particular for control rods. The calculated critical boron concentration, flux profiles and power distributions are compared with the experimental measurements at the time points marked in Figure 6.4a (5 to 43), where the power has been stable long enough that the xenon can be considered at equilibrium at full power. The control bank position is taken as constant between time points, where the insertion is calculated as an average of the raw data over time weighted with the thermal power. The inlet temperature and mass flow rate are taken as constant during the whole cycle.

6.1.3 Experimental data used for validation

Two sets of experimental data are used in this work to evaluate the accuracy of the numerical calculations: critical boron concentration and direct pin-level neutron flux measurements. The critical boron (Figure 6.4b) comes from boric acid (H_3BO_3) concentration measurements during reactor operation and the flux profiles correspond to ^{103}Rh Self Powered Neutron Detector (SPND) [95] readings. The SPNDs are inserted in the central instrumentation tubes (see Figure 6.2), with the radial distribution shown in Figure 6.5, where the indexing is defined. Axially, each detector consists of a string of 7 SPNDs separated equidistantly by about 50 cm, with the lowest position approximately 30 cm above the bottom of the active length. While all the measured data is available at 43 time points through the operating cycle, the points selected for comparison are the ones shown in Figure 6.4a. The SPND measurements are presented in Section 6.3.2 along with the numerical results.

In addition to direct measurements, the axial and radial power profiles reported by the CMS are available to assess the power distribution predicted by Serpent. The

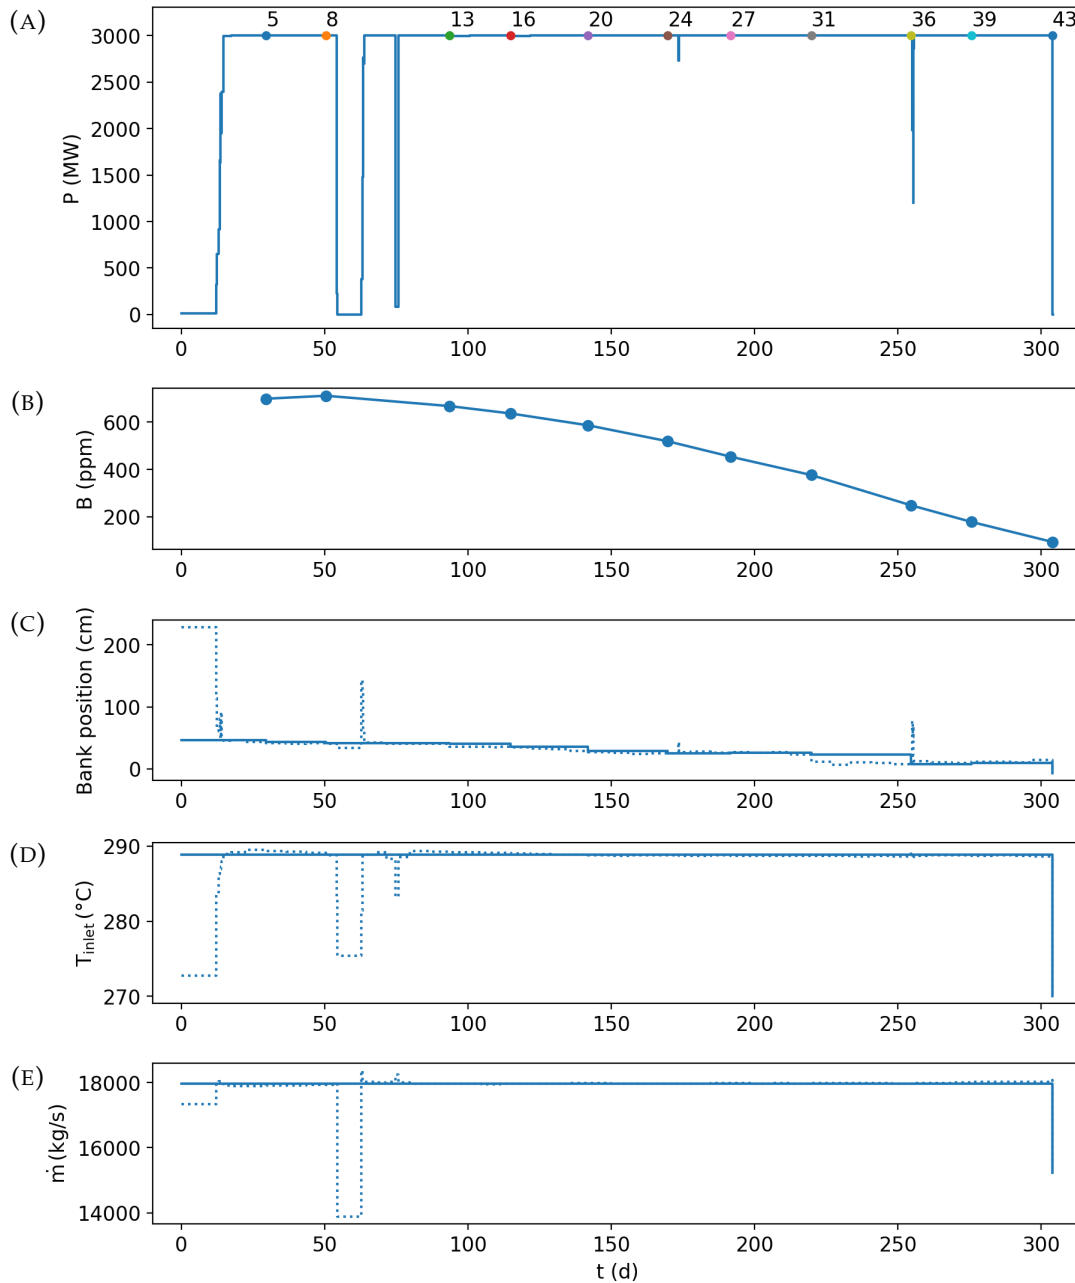


FIGURE 6.4: Operational history for the ninth cycle of the Temelín II reactor: (A) thermal power P and measured time points, (B) measured critical boron concentration B , (C) condensed control bank positions, (D) inlet temperature T_{inlet} and (E) mass flow rate \dot{m} (the raw operational data is shown in dashed lines).

CMS reconstructs the power based on the adjustment of a theoretical solution to the SPND data.

6.2 Modelling approach

The burnup calculation has been performed at HFP conditions with the operating parameters summarized in Table 6.1. The burnup steps are the ones marked in Figure 6.4, for which constant control bank positions are obtained. The calculated




FIGURE 6.5: Radial distribution of SPND strings for the Temelín II reactor (23.6 cm fuel-assembly pitch). The detectors are inserted in the central instrumentation tubes (see Figure 6.2).

neutron flux profiles and power distributions are compared with the experimental data at Beginning of Cycle (BOC; time point 5, 15.8 EFPD) and End of Cycle (EOC; time point 43, 279.9 EFPD)

The modelling methodology used for Temelín II is essentially the same as the one used for the Pre-Konvoi PWR case, as outlined in Section 5.2, with the exception of a few specific characteristics of VVER-1000 reactors. The central hole in the fuel pellet is included in SCF and TU, and specific models for fuel pins containing Gd_2O_3 are used in TU. In addition, the stiffener plates are modelled explicitly in Serpent and included in SCF increasing the wetted parameter of the subchannels in the corners of the fuel assemblies. A detailed description of the Serpent, SCF and TU models is provided in Appendix D.

6.3 Selected results

The accuracy of the simulation method in terms of core reactivity and pin-level results is assessed using the critical boron concentration and neutron flux measurements. In addition, the power distribution reconstructed by the CMS is used to evaluate the solution at the nodal level.

6.3.1 Critical boron concentration

The measured and calculated critical boron concentration for the period up to 279.9 EFPD (EOC) is shown in Figure 6.6. The evolution of the boron concentration as a function of burnup corresponds to a typical boron letdown curve, where

the reactivity increases initially due to the depletion of the burnable absorbers and then decreases as the reactor operates due to fuel burnup. The RMS difference between the results and the experimental data is 33.8 and 31.5 ppm for Serpent-SCF and Serpent-SCF-TU, respectively. The differences between the results with and without TU, which are lower than 7 ppm for the whole cycle, are not significant.

The deviations with respect to the experimental data are around the acceptability criterion for this type of calculation, which is usually between 50 and 100 ppm. The burnup step size, which is a sensitive parameter especially for core loadings with burnable absorbers, is likely a contributing factor to the error with respect to the measured values. Unfortunately, in this case the number of steps is bounded by the computing resources and cannot be reduced further to perform a burnup-step sensitivity analysis, which is ideally done in these cases. The assumption of equilibrium xenon should not be a large source of error, since the points selected to compare the boron concentration are far enough from power ramps to ignore xenon oscillations, which are usually considered to be important only within about two days of a change in the power level.




FIGURE 6.6: Comparison between measured and calculated critical boron concentration up to 279.9 EFPD (EOC) for the Temelín II reactor.

6.3.2 Neutron flux profiles

Figures 6.7 and 6.8 show the normalized neutron flux profiles at BOC and EOC for the two locations with the smallest and largest RMS deviations between the Serpent detector responses and the SPND measurements. The indexing is defined in Figure 6.5 and the error bars for the results correspond to three standard deviations $\sigma(z)$ for each axial location. The experimental error for the flux is not available and is difficult to estimate, and therefore the measurements are reported without error bars. The uncertainty in the axial position of the detectors is not available either.

Both the calculated fluxes ϕ and the SPND measurements ϕ_{SPND} are normalized such that

$$\bar{\phi} = \bar{\phi}_{SPND} = 1, \quad (6.1)$$

where $\bar{\phi}$ and $\bar{\phi}_{SPND}$ are the mean values taken over all radial and axial positions. At BOC, the axial neutron-flux profiles have the typical cosine shape of VVER cores loaded with fresh fuel. The flux distribution gets more uniform as the reactor operates and the fuel in the central region of the core is depleted, resulting in flatter flux profiles at EOC.

Overall, the results are in good agreement with the experimental data, and the differences between the calculations with and without TU seem to be purely statistical fluctuations. While the experimental measurements fall within the statistical range of the results for almost all data points, the flux at the top position is under-predicted for some SPND detectors. Some of these deviations might be due to the simplified axial reflector model, which consists of a homogeneous effective material for the structural components and the coolant at the core outlet. However, given the strong flux gradient in this region, it is not possible to derive a strong conclusion regarding the significance of these deviations without a proper quantification of the uncertainty in the axial positions of the SPND stack.

6.3.3 Power distribution

Figure 6.9 shows the results and the CMS reconstruction for the axial (radially integrated) power profile of the core at BOC and EOC. As for the neutron flux, an axial flattening of the power distribution between BOC and EOC, which is caused by the depletion of the central region, can be noted. The calculated profiles are in very good agreement with the CMS values, and the differences are concentrated in spacer grid positions, where the Serpent calculation is expected to be more accurate than the CMS reconstruction from the SPND readings.

The radial (axially integrated) power distribution reported by the CMS at BOC and EOC is shown in Figure 6.10, while figures 6.11 to 6.12 present the results and their deviations with respect to the CMS. The statistical error for the results corresponds to three standard deviations $\sigma(r)$ and the relative differences $\epsilon(r)$ are calculated as

$$\epsilon(r) = \frac{P(r) - P_{CMS}(r)}{\bar{P}_{CMS}}, \quad (6.2)$$

where $P(r)$ and $P_{CMS}(r)$ are the calculated and reported power distributions and \bar{P}_{CMS} is the average fuel-assembly power for the whole core. In this case the results are in good agreement with the CMS as well, and the deviations seem purely stochastic. The largest differences with respect to the CMS occur in the core periphery, where the statistical convergence is lower.

6.3.4 Pin-level results

As shown in the previous sections, the impact of including fuel-performance capabilities in the calculation system does not influence the neutronic results in a significant way, at least not for the burnup calculations considered here. This is in line with previous investigations dealing with smaller-scale problems [71]. However, the main advantage of the three-code approach used in this work is the ability to simulate the neutronic, thermohydraulic and thermomechanic behavior of the core with a fully coupled high-fidelity scheme. In particular, the fuel-performance calculation




FIGURE 6.7: Comparison between measured and calculated neutron flux profiles for the Temelín II reactor at the locations with the smallest RMS deviations at (A, B) BOC (15.8 EFPD) and (C, D) EOC (279.9 EFPD). The SPND positions are shown in Figure 6.5 and $z = 0$ corresponds to the bottom of the active length.

is done within the burnup simulation, obtaining a fuel solution with full detail for consistent core states.

Figure 6.13 shows the power calculated by Serpent EOC (279.9 EFPD), as well as the coolant temperature solved by SCF, which correspond to nominal HFP conditions. This level of detail allows for the direct calculation of local safety parameters at pin level, as well as a realistic description of the core to perform the pin-by-pin burnup calculation and the fuel-performance analysis. The Departure from Nucleate Boiling Ratio (DNBR) is defined as

$$DNBR = \frac{q_c''}{q''} = f(h_{in}, h_{sat}, p, x, \dot{m}), \quad (6.3)$$

where q_c'' is the Critical Heat Flux (CHF) and q'' is the local heat flux. Most empirical




FIGURE 6.8: Comparison between measured and calculated neutron flux profiles for the Temelín II reactor at the locations with the largest RMS deviations at (A, B) BOC (15.8 EFPD) and (C, D) EOC (279.9 EFPD). The SPND positions are shown in Figure 6.5 and $z = 0$ corresponds to the bottom of the active length.

correlations, e.g. Westinghouse's W-3 [96], take into account the inlet enthalpy h_{in} , the saturated enthalpy h_{sat} and the coolant quality x , as well as p and \dot{m} . Using the approach proposed in this work, the pin-level DNBR can be obtained directly as part of the burnup calculation.

Figure 6.14 presents the fuel-cladding gap solution predicted by TU at EOC, where the same behavior as for the Pre-Konvoi reactor can be observed, namely that for the core central region the gap width is smaller due to thermal expansion, which results in a larger gap conductance. While the gap behavior does not necessarily play an important role in the burnup calculation presented here either, it is a relevant parameter for the prediction of peak fuel and cladding temperatures in transient simulations and safety analysis. Obtaining the gap width and conductance from TU, which has been extensively validated for PWR and VVER designs [63], is a key benefit of this approach.

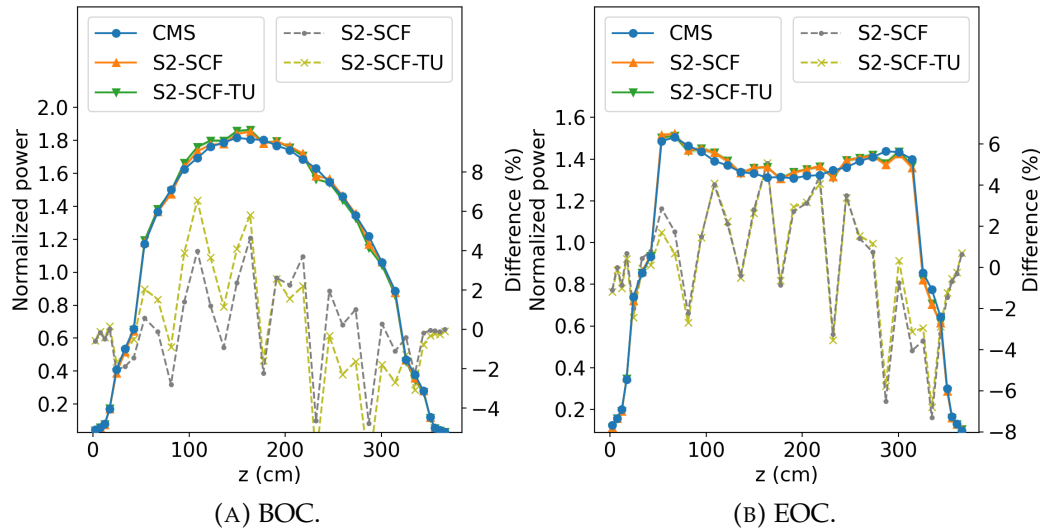


FIGURE 6.9: CMS and calculated axial power profiles for the Temelín II reactor at (A) BOC (15.8 EFPD) and (B) EOC (279.9 EFPD).




FIGURE 6.10: CMS radial power profiles for the Temelín II reactor at (A) BOC (15.8 EFPD) and (B) EOC (279.9 EFPD).

Finally, the fuel centerline temperature and the Xe fission-gas release calculated by TU at EOC are shown in Figure 6.15. Both the fuel temperature and the fission gas release follow the shape of the power distribution. The peak fuel and cladding temperatures are typical safety limits in nominal operating conditions as well as in transient scenarios. The fission gas release is a limiting design factor, especially for high-burnup fuel.

6.4 Validation summary

The results obtained with Serpent-SCF with and without TU are in quite good agreement with the experimental data for the Temelín II VVER-1000 NPP. The deviations between the calculated and measured critical boron concentrations are below 70 ppm for the whole cycle, which is within the typical acceptability criterion for this kind of calculations. The pin-level neutron flux profiles obtained with Serpent are consistent with the SPND measurements within their statistical uncertainty in most cases. The calculated axial and radial power profiles are in very good agreement

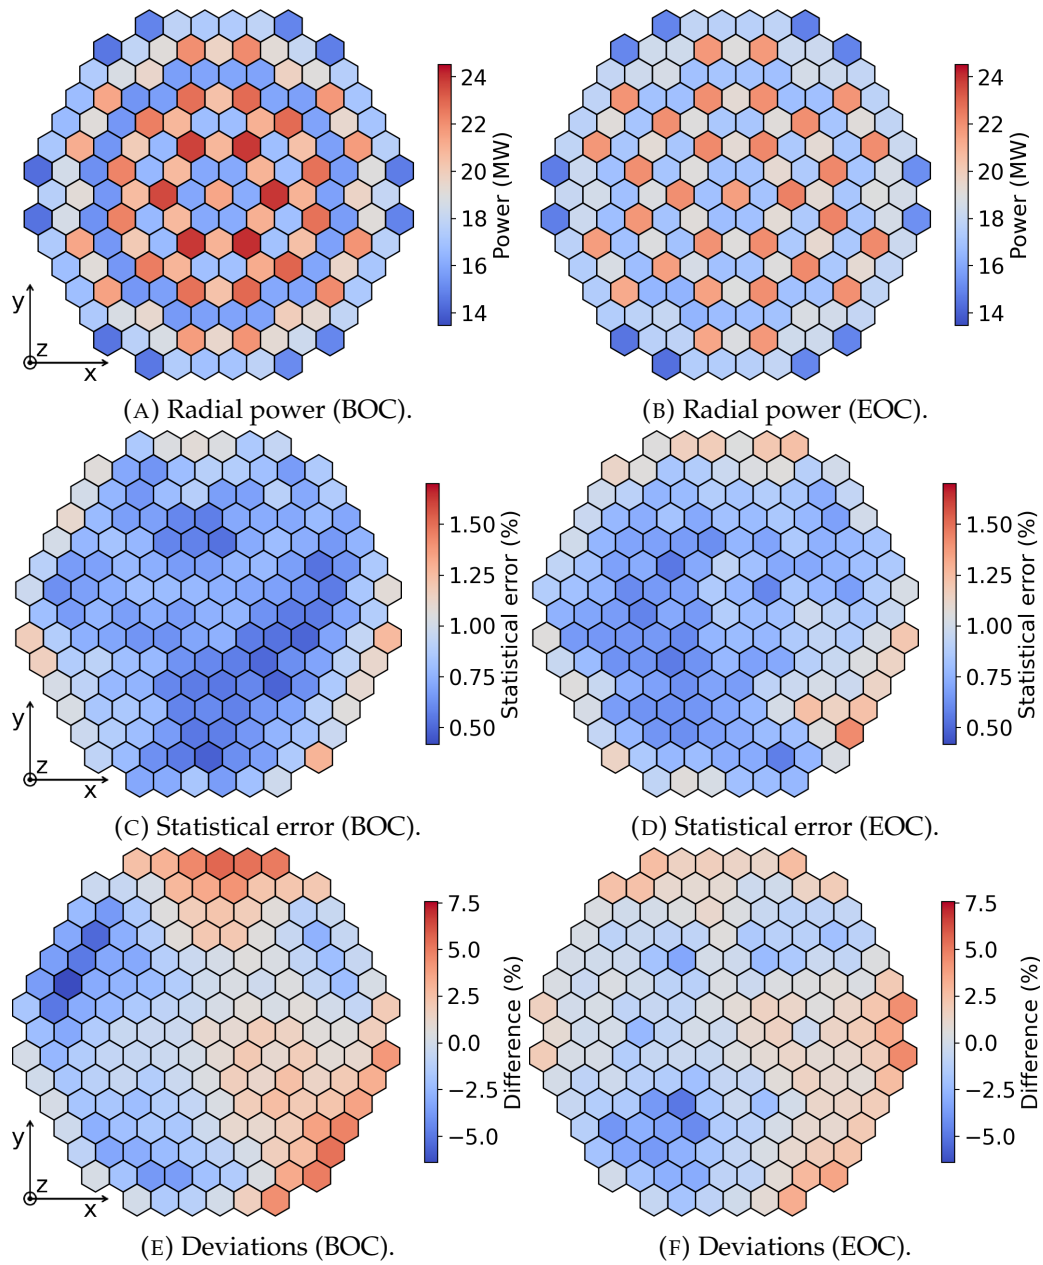


FIGURE 6.11: Serpent-SCF radial power profiles (A, B), statistical errors (C, D) and deviations against the CMS (E, F) for the Temelín II reactor at BOC (15.8 EFPD) and EOC (279.9 EFPD).

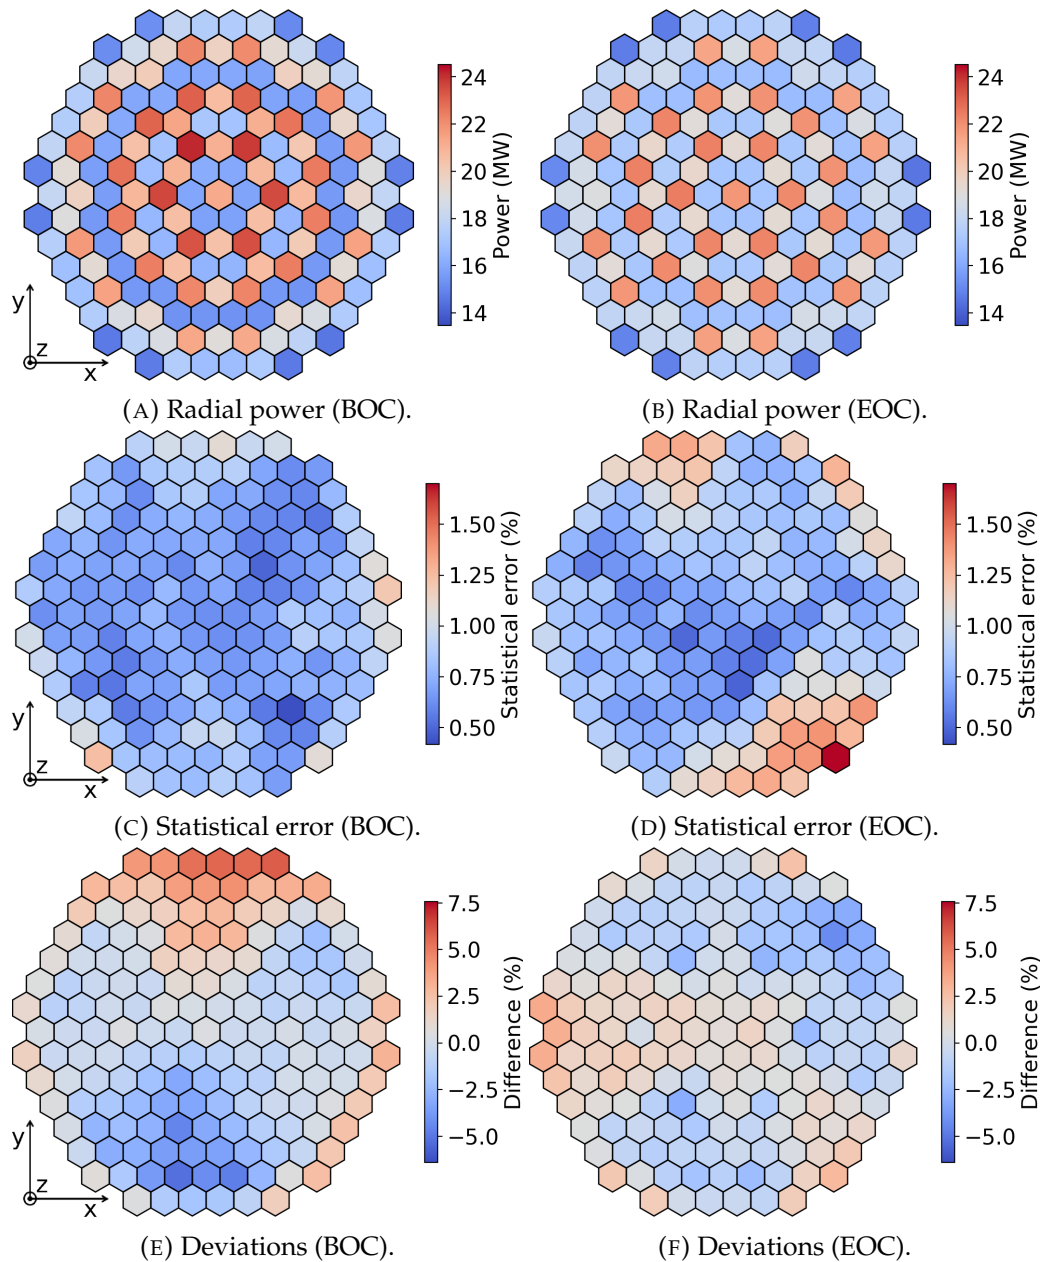


FIGURE 6.12: Serpent-SCF-TU radial power profiles (A, B), statistical errors (C, D) and deviations against the CMS (E, F) for the Temelin II reactor at BOC (15.8 EFPD) and EOC (279.9 EFPD).

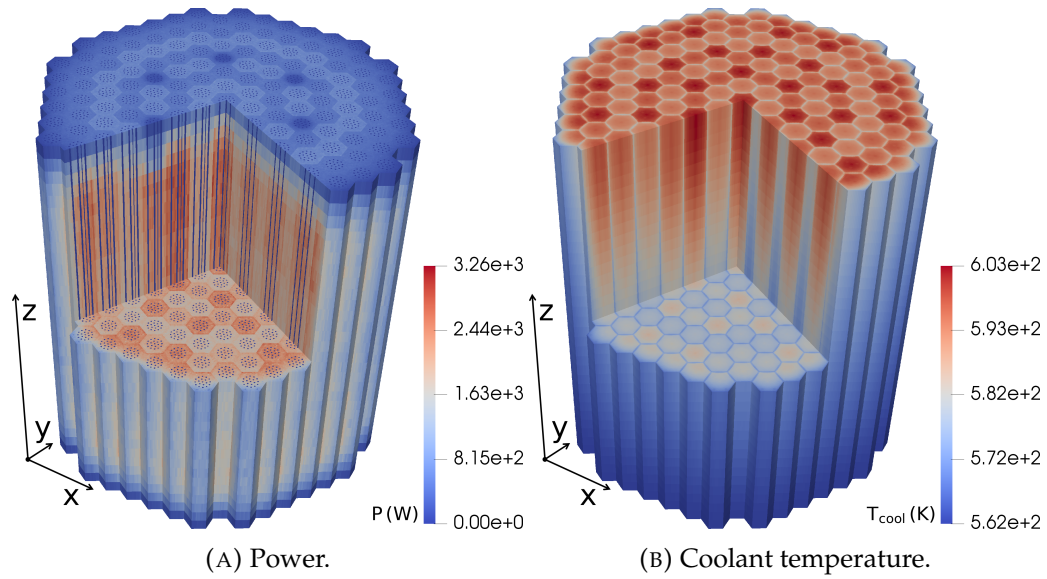


FIGURE 6.13: Calculated power (A) and coolant temperature (B) for the Temelín II reactor at EOC (279.9 EFPD).




FIGURE 6.14: Calculated fuel-cladding gap width (A) and conductance (B) for the Temelín II reactor at EOC (279.9 EFPD).

with the CMS, with mean differences below 3%. No significant differences in the neutronic solution are found between Serpent-SCF with and without TU.

The high-fidelity capabilities of the three-code coupling system are demonstrated with pin-level results for neutronic, thermalhydraulic and thermomechanic fields, in particular for safety-relevant parameters such as peak fuel and cladding temperatures and Xe release.

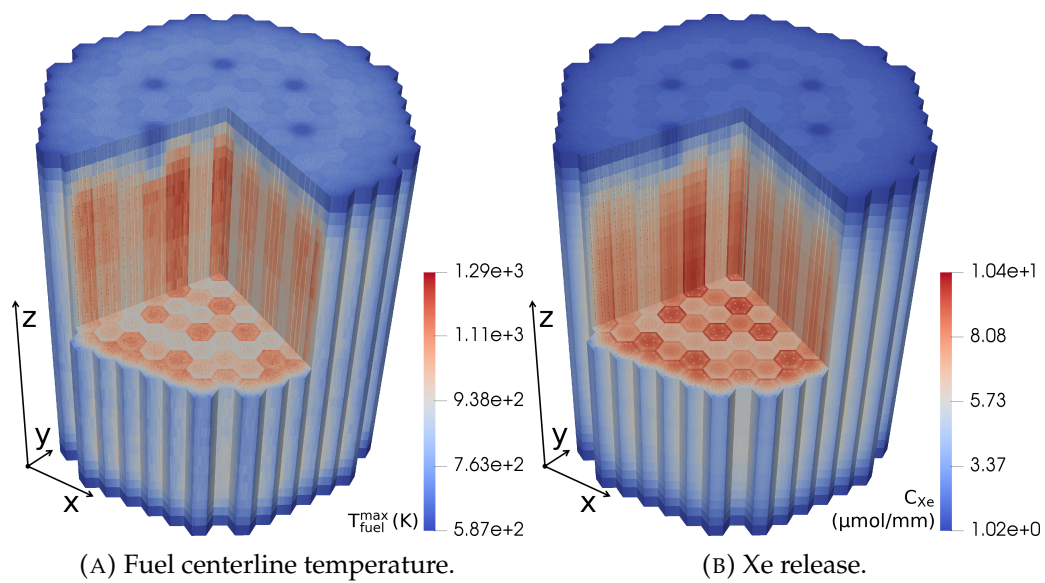


FIGURE 6.15: Calculated fuel centerline temperature (A) and Xe release (B) for the Temelín II reactor at EOC (279.9 EFPD).

Chapter 7

Summary and outlook

Neutronics, thermalhydraulics and fuel-performance analysis are the three main disciplines on which LWR core analysis is based. In this thesis, a coupling system is devised to integrate these three areas to provide a consistent description of the core during its operation. A highly detailed model of the core is used to obtain a pin-by-pin solution across the three physical domains. This methodology follows a worldwide trend in computational reactor physics, where the wide availability of HPC resources motivates the development of high-fidelity multiphysics applications.

The software package developed in this thesis is based on Serpent, a continuous-energy Monte Carlo neutron transport code, SUBCHANFLOW, a subchannel thermalhydraulics code, and TRANSURANUS, a fuel thermomechanics code. The coupling system is based on a semi-implicit iterative scheme and implemented using a novel object-oriented mesh-based design. To tackle pin-by-pin depletion, a domain decomposition scheme is implemented in Serpent, allowing the simulation of massive burnup problems. In this way, depletion calculations can be performed directly using continuous-energy nuclear data and an extremely detailed geometrical representation, which provides a high-fidelity alternative to the traditional multistep approach based on lattice- and core-level computations. As a result, most parameters of interest for both design and safety analysis can be obtained directly from a single simulation. This three-code system with full-core depletion capabilities using the Monte Carlo method represents a development beyond the state of the art in current reactor physics simulations.

To validate this tool, burnup calculations for a German Pre-Konvoi PWR reactor and for the Temelín II VVER-1000 plant are performed. The results are assessed using the critical boron concentration and a set of pin-level neutron flux measurements. The coupled system is capable of reproducing the experimental data with a high degree of confidence and a full validation for global and local parameters is achieved. These calculations serve to validate the CDD methodology developed in this work, which is the only feature that modifies the internal calculation scheme within a physical domain, in this case for Serpent. Hence, the full-core pin-by-pin burnup capabilities, one of the most innovative features of the coupled system, are demonstrated for two real-life scenarios. The three-code coupling scheme is validated indirectly from the consistency of the overall solution, keeping in mind that both SCF and TU have been extensively validated for LWR systems. These two cases also serve to demonstrate the feasibility of using Serpent-SCF-TU to perform full-core depletion analysis using an HPC system.

The main drawback of this methodology is of course the massive computational requirements. The calculations presented in the previous chapters require about a week in 1,280 cores, which is obviously not feasible for a routine calculation. Hence, this tool is not suitable for tasks such as design optimization, as expected. However, it is not rare nowadays for a nuclear engineer to have access to a cluster with at least

a few hundred cores. In such cases, Serpent-SCF-TU and Serpent-SCF can be used as a high-fidelity alternative to the traditional multi-step calculation approach. As such, it serves as a verification tool, since it can provide an independent solution relying on very few approximations. Moreover, it can be used to obtain extremely detailed solutions for selected scenarios during the burnup cycle without utilizing legacy methodologies such as pin-power reconstruction. In this way, pin-level safety parameters can be calculated directly, which is very useful in terms of best-estimate safety analysis.

7.1 Summary

In terms of the research objectives, the achievements of this work can be summarized as:

1. A neutronic-thermalhydraulic-thermomechanic coupling methodology has been formulated and implemented. The algorithm is based on the operator-splitting method with a semi-implicit iterative technique. A pin-level feedback scheme capturing the main interaction mechanisms in the system has been proposed and tested. The software implementation is based on the object-oriented paradigm, and consists on modularizing each code following an inheritance-based design and implementing the coupling scheme in a supervisor program using object-oriented features. The calculation tool has been implemented using this technique and proved to meet all requirements, in particular in HPC environments.
2. A Collision-based Domain Decomposition scheme has been formulated and implemented in Serpent. This feature allows the code to split the memory demand in burnup calculation across computing nodes, adding memory scalability to the system. This solves the memory bottleneck arising in depletion problems and enables the simulation of full-core pin-by-pin cases.
3. The fully coupled code system has been successfully validated using experimental data from two real-life nuclear power plants, a Pre-Konvoi PWR and a VVER-1000. This is a key step in the development process, since it proves the accuracy of the developed tool.
4. The newly developed high-fidelity system has been proven to be a powerful tool for massively detailed core analysis. Combining neutronics, thermalhydraulics and fuel-performance analysis, it provides a complete and consistent description of the core avoiding most of the approximations used by conventional neutron physics methodologies. The main drawback is the large amount of computing time needed to obtain statistically significant results.

7.2 Outlook

The three-code system developed in this work pushes forward the status of high-fidelity multiphysics simulations for reactor analysis. As such, it builds upon the progress made in recent years in this type of software package, in particular on a Serpent-SCF coupling for steady-state simulations developed at KIT [47]. Some of the issues identified in that work, for instance the need for a domain decomposition scheme to handle large-scale burnup in Serpent, have been successfully addressed

in this thesis. Nevertheless, there is much room for improvement and further work in the topic of high-fidelity multiphysics simulations.

Some of the particular issues that could be addressed in future research projects include:

1. While the feedback scheme proposed in this work captures the most important feedback mechanisms in the system, some aspects could be refined. Regarding the neutronic calculation, the effect of radial fuel temperature profiles for Doppler feedback could be further analyzed. While this has been shown not to be significant in burnup calculations [71], the effects on parameters such as Doppler reactivity coefficients have not been quantified. In a similar way, the calculation of radial power profiles inside the fuel rods and its feedback to the thermomechanic model merits further investigation. While the direct tallying of radial power profiles is likely not feasible, as explained in Section 3.2.3.2, the use of semi-empirical models or Functional Expansion Tallies (FETs) could be studied. Finally, the use of Serpent isotopic data in TU should be considered in order to model consistent material compositions across physics and to avoid the empirical burnup schemes used in fuel-performance analysis.
2. This thesis focuses on depletion calculations considering fuel-performance aspects. In this context, the effect on the neutronics of improving the thermomechanic modelling has been shown not to be significant. However, the impact can be expected to be larger in transient calculations, and in particular in RIA-type accidents, where the Doppler feedback on the power evolution is the dominant factor in the first few seconds of the transient. In these scenarios a large amount of power is suddenly deposited in the fuel rods, and the width and conductivity of the fuel-cladding gap plays a major role in the rate of evacuation of the power through the coolant. Hence, extending the three-code coupled implemented in this work to transient scenarios, which could be done in the framework of a PhD project, for instance, would be an extremely useful development. This is a straightforward task from the software point of view, since the calculation and feedback scheme used for depletion can be readily applied to transient problems. The main difference regarding the physical model is the interaction between the thermomechanic and thermalhydraulic calculations at the cladding-coolant interface, where the steady-state approximation used in this work is no longer valid and a tightly coupled formulation is needed.
3. The most useful feature in terms of performance would be a dynamic load balancing scheme for the CDD scheme in Serpent. In the current implementation, materials are distributed in an equal-volume basis, which leads to load imbalances when the fission source is not uniform, which is typically the case. An optimum load balance could be achieved adjusting the size of the domains to force the number of neutrons in each MPI to be equal. In burnup problems, this would have to be done dynamically as the fission source distribution evolves. Such scheme requires exchanging material information between MPI tasks at runtime between burnup steps, which is not trivial at all. In terms of memory scalability, the decomposition of non-material tallies and feedback data, which are currently replicated across domains, could be explored.
4. Regarding validation, the PWR and VVER cases presented in this work demonstrate the accuracy of the three-code system for the target problems of

this thesis, i.e. full-core depletion of LWRs. Those studies focus on the Serpent calculation, and in particular on checking that the newly implemented CDD scheme gives accurate results. Since no major modifications have been made to SCF and TU, this work relies on the previous validation of these codes. However, the coupled system could be applied to irradiation experiments such as the OECD Halden reactor project, which has been widely used to validate fuel-performance codes, providing a detailed description of the system to inform the validation process. In particular, Serpent and SCF can be used to perform validation studies for the thermomechanic calculation using consistent neutronic and thermalhydraulic conditions. As a general comment, the tool developed in this work will hopefully be applied to core calculations for other reactors, since it provides quite advanced modelling features. Moreover, the approach developed in this thesis can be used as a reference high-order method to verify low-order deterministic methods on a code-to-code basis.

5. Finally, the object-oriented methodology used to develop Serpent-SCF-TU has inherent features that make it particularly suitable as a multiphysics framework in the long term. Additional codes can be readily integrated to the system to add multiphysics and multiscale features. In the framework of the EU Horizon 2020 McSAFER project [97], a follow-up of the McSAFE project [89], the nodal-level neutronic code PARCS is being included as a low-order alternative to Serpent. In addition, the TRACE system code and the OpenFOAM CFD code are being coupled to the system to perform multiscale thermalhydraulic calculations.

Bibliography

- [1] International Atomic Energy Agency (IAEA). “Nuclear Power Reactors in the World”. *Reference data series No. 2, 2019 Edition, IAEA-RDS-2/39 Technical Report* (2019).
- [2] USNRC Technical Training Center, 4-1: Pressurized Water Reactor (PWR) Systems. “Reactor Concepts Manual”. *U. S. Nuclear Regulatory Commission, LTR-14-0412-1 Training Document* (2012).
- [3] Y. Chugui, A. Verkhogliad, V. Kalikin, and P. S. Zav’yalov. “3D optical measuring technologies for industrial applications”. *International Society for Optics and Photonics: Optical Measurement Systems for Industrial Inspection VII 8082* (2011), pp. 632–644. DOI: <http://doi.org/10.1117/12.895003>.
- [4] G. I. Bell and S. Glasstone. *Nuclear Reactor Theory, First Edition*. Van Nostrand Reinhold Company, New York, 1970. ISBN: 978-0-442-20684-0.
- [5] J. J. Duderstadt and L. J. Hamilton. *Nuclear Reactor Analysis, First Edition*. John Wiley & Sons Inc., New York, 1976. ISBN: 978-0-471-22363-4.
- [6] A. K. Prinja and E. W. Larsen. “Handbook of Nuclear Engineering”. Springer Science+Business Media LLC, Berlin, 2010. Chap. 5: General Principles of Neutron Transport. ISBN: 978-0-387-98130-7.
- [7] W. M. Stacey. *Nuclear Reactor Physics, Third Edition*. Wiley-VCH, Weinheim, 2018. ISBN: 978-3-527-41366-9.
- [8] J. Leppänen. “Development of a new Monte Carlo reactor physics code”. Dissertation. Helsinki University of Technology, Department of Engineering Physics and Mathematics, 2007. ISBN: 978-9-513-87018-8.
- [9] J. Rhodes, K. Smith, and D. Lee. “CASMO-5 development and applications”. *PHYSOR 2006 Conference, Vancouver, Canada, September, 2006. B144*. 2006. ISBN: 978-1-604-23696-5.
- [10] C. A. Wemple, H-N. M. Gheorghiu, R. J. J. Stamm’ler, and E. Villarino. “Recent advances in the HELIOS-2 lattice physics code”. *PHYSOR 2008 Conference, Interlaken, Switzerland, September, 2008*. 2008. ISBN: 978-3-952-14095-6.
- [11] J. Leppänen, M. Pusa, T. Viitanen, V. Valtavirta, and T. Kaltiaisenaho. “The Serpent Monte Carlo code: status, development and applications in 2013”. *Annals of Nuclear Energy* 82 (2015), pp. 142–150. DOI: <http://doi.org/10.1016/j.anucene.2014.08.024>.
- [12] C. J. Werner, J. S. Bull, C. J. Solomon, F. B. Brown, G. W. McKinney, M. E. Rising, D. A. Dixon, R. L. Martz, H. G. Hughes, L. J. Cox, et al. “MCNP User’s Manual Code Version 6.2”. *Los Alamos National Laboratory, LA-UR-17-29981 Technical Report* (2017).
- [13] R. E. MacFarlane and A. C. Kahler. “Methods for processing ENDF/B-VII with NJOY”. *Nuclear Data Sheets: Nuclear Reaction Data* 111 (2010), pp. 2739–2890. DOI: <http://doi.org/10.1016/j.nds.2010.11.001>.

- [14] T. Bahadir and S-Ö. Lindahl. "Studsvik's next generation nodal code SIMULATE 5". *ANFM 2009 Conference, Hilton Head Island, USA, April, 2009*. 2009, pp. 1–12. ISBN: 978-1-615-67351-3.
- [15] D. A. Barber and H. G. Joo. "PARCS: a multi-dimensional two-group reactor kinetics code based on the nonlinear analytic nodal method". *Purdue University, School of Nuclear Engineering, PU/NE-98-26 Technical Report* (1998).
- [16] K. Smith and B. Forget. "Challenges in the development of high-fidelity LWR core neutronics tools". *M&C 2013 Conference, Sun Valley, USA, May, 2013*. 2013, pp. 1809–1825. ISBN: 978-0-894-48700-2.
- [17] H. G. Joo, J. Y. Cho, K. S. Kim, C. C. Lee, and S. Q. Zee. "Methods and performance of a three-dimensional whole-core transport code DeCART". *PHYSOR 2004 Conference, Chicago, USA, April, 2004*. 2004, pp. 25–29. ISBN: 978-0-894-48683-7.
- [18] MPACT Team. "MPACT Theory Manual v2.1.0". *Consortium for Advanced Simulation of Light Water Reactors, CASLU-2015-1005-000 Technical Report* (2015).
- [19] F. Févotte and B. Lathuilière. "MICADO: parallel implementation of a 2D-1D iterative algorithm for the 3D neutron transport problem in prismatic geometries". *M&C 2013 Conference, Sun Valley, USA, May, 2013*. 2013, pp. 1–12. ISBN: 978-0-894-48700-2.
- [20] A. Marin-Lafleche, M. A. Smith, and C. Lee. "PROTEUS-MOC: a 3D deterministic solver incorporating 2D Method of Characteristics". *M&C 2013 Conference, Sun Valley, USA, May, 2013*. 2013, pp. 2759–2770. ISBN: 978-0-894-48700-2.
- [21] C. Beckert and U. Grundmann. "Development and verification of a nodal approach for solving the multigroup SP3 equations". *Annals of Nuclear Energy* 35 (2008), pp. 75–86. DOI: <http://doi.org/10.1016/j.anucene.2007.05.014>.
- [22] K. Wang, Z. Li, D. She, J. Liang, Q. Xu, Y. Qiu, J. Yu, J. Sun, X. Fan, and G. Yu. "RMC - A Monte Carlo code for reactor core analysis". *Annals of Nuclear Energy* 82 (2015), pp. 121–129. DOI: <http://doi.org/10.1016/j.anucene.2014.08.048>.
- [23] H. Lee, W. Kim, P. Zhang, M. Lemaire, A. Khassenov, J. Yu, Y. Jo, J. Park, and D. Lee. "MCS - A Monte Carlo particle transport code for large-scale power reactor analysis". *Annals of Nuclear Energy* 139, 107276 (2020). DOI: <http://doi.org/10.1016/j.anucene.2019.107276>.
- [24] D. P. Griesheimer, D. F. Gill, B. R. Nease, T. M. Sutton, M. H. Stedry, P. S. Dobreff, D. C. Carpenter, T. H. Trumbull, E. Caro, H. Joo, and D. L. Millman. "MC21 v60 - a continuous-energy Monte Carlo particle transport code with integrated reactor feedback capabilities". *Annals of Nuclear Energy* 82 (2015), pp. 29–40. DOI: <http://doi.org/10.1016/j.anucene.2014.08.020>.
- [25] S. B. Pope. *Turbulent Flows, First Edition*. Cambridge University Press, Cambridge, 2000. ISBN: 978-0-521-59886-6.
- [26] D. A. Anderson, J. C. Tannehill, and R. H. Pletcher. *Computational Fluid Mechanics and Heat Transfer, Third Edition*. CRC Press, Taylor & Francis Group, Boca Raton, 2016. ISBN: 978-1-591-69037-5.
- [27] F. P. Incropera, D. P. DeWitt, T. L. Bergman, and A. S. Lavine. *Fundamentals of Heat and Mass Transfer, Sixth Edition*. John Wiley & Sons Inc., New York, 2006. ISBN: 978-0-471-45728-2.

- [28] C. Fiorina, I. Clifford, M. Aufiero, and K. Mikityuk. "GeN-Foam: a novel OpenFOAM based multi-physics solver for 2D/3D transient analysis of nuclear reactors". *Nuclear Engineering and Design* 294 (2015), pp. 24–37. DOI: <http://doi.org/10.1016/j.nucengdes.2015.05.035>.
- [29] E. Baglietto, E. Demarly, R. Kommajosyula, N. Lubchenko, B. Magolan, and R. Sugrue. "A second generation multiphase-CFD framework toward predictive modeling of DNB". *Nuclear Technology* 205 (2019), pp. 1–22. DOI: <http://doi.org/10.1080/00295450.2018.1517528>.
- [30] C. L. Wheeler, C. W. Stewart, R. J. Cena, D. S. Rowe, and A. M. Sutey. "COBRA-IV-I: an interim version of COBRA for thermal-hydraulic analysis of rod bundle nuclear fuel elements and cores". *Battelle Pacific Northwest Laboratories, BNWL-1962/UC-32 Technical Report* (1976).
- [31] The OpenFOAM Foundation. *OpenFOAM*. Accessed: 21/10/2021. URL: <http://openfoam.org>.
- [32] ANSYS. *ANSYS Fluent: Fluid Simulation Software*. Accessed: 21/10/2021. URL: <http://ansys.com/products/fluids/ansys-fluent>.
- [33] A. Moorthi, A. Kumar Sharma, and K. Velusamy. "A review of sub-channel thermal hydraulic codes for nuclear reactor core and future directions". *Nuclear Engineering and Design* 332 (2018), pp. 329–344. DOI: <http://doi.org/10.1016/j.nucengdes.2018.03.012>.
- [34] Y. J. Yoo, D. H. Hwang, and D. S. Sohn. "Development of a subchannel analysis code MATRA applicable to PWRs and ALWRs". *Journal of the Korean Nuclear Society* 31 (1999), pp. 314–327. ISSN: 0372-7327.
- [35] U. Imke and V. Sanchez-Espinoza. "Validation of the subchannel code SUB-CHANFLOW using the NUPEC PWR tests (PSBT)". *Science and Technology of Nuclear Installations* 2012, 465059 (2012). DOI: <http://doi.org/10.1155/2012/465059>.
- [36] K. S. Chaudri, Y. Su, R. Chen, W. Tian, G. Su, and S. Qiu. "Development of sub-channel code SACoS and its application in coupled neutronics/thermal hydraulics system for SCWR". *Annals of Nuclear Energy* 45 (2012), pp. 37–45. DOI: <http://doi.org/10.1016/j.anucene.2012.02.014>.
- [37] R. J. Salko, M. Avramova, A. Wysocki, A. Toptan, J. Hu, N. Porter, T. S. Blyth, C. A. Dances, A. Gomez, C. Jernigan, and J. Kelly. "CTF 4.0 Theory Manual". *Oak Ridge National Laboratory, ORNL/TM-2019/1145 Technical Report* (2019).
- [38] S. Kakaç, Y. Yener, and C. P. Naveira-Cotta. *Heat Conduction, Fifth Edition*. CRC Press, Taylor & Francis Group, Boca Raton, 2018. ISBN: 978-1-138-94384-1.
- [39] J. P. Moitinho de Almeida and E. A. Maunder. *Equilibrium Finite Element Formulations, First Edition*. John Wiley & Sons Inc., New York, 2017. ISBN: 978-1-118-42415-5.
- [40] P. Fors. "The effect of dissolved hydrogen on spent nuclear fuel corrosion". Dissertation. Chalmers University of Technology, Department of Chemical and Biological Engineering, 2009. ISBN: 978-9-173-85259-3.
- [41] L. Van Brutzel, R. Dingreville, and T. J. Bartel. "Nuclear fuel deformation phenomena". *OECD/NEA, NEA-NSC-R-2015-5 Technical Report* (2015).

- [42] D. M. Perez, R. L. Williamson, S. R. Novascone, T. K. Larson, J. D. Hales, B. W. Spencer, and G. Pastore. "An evaluation of the nuclear fuel performance code BISON". *M&C 2013 Conference, Sun Valley, USA, May, 2013*. 2013, pp. 1273–1284. ISBN: 978-0-894-48700-2.
- [43] K. J. Geelhood, W. G. Luscher, P. A. Raynaud, and I. E. Porter. "FRAPCON-4.0: a computer code for the calculation of steady-state, thermal-mechanical behavior of oxide fuel rods for high burnup". *Pacific Northwest National Laboratory, PNNL-19418 Technical Report* (2015).
- [44] P. Van Uffelen, C. Györi, A. Schubert, J. van de Laar, Z. Hózer, and G. Spykman. "Extending the application range of a fuel performance code from normal operating to design basis accident conditions". *Journal of Nuclear Materials* 383 (2008), pp. 137–143. DOI: <http://doi.org/10.1016/j.jnucmat.2008.08.043>.
- [45] R. Tuominen, V. Valtavirta, M. García, D. Ferraro, and J. Leppänen. "Effect of energy deposition modelling in coupled steady state Monte Carlo neutronics/thermal hydraulics calculations". *EPJ Web Conference (PHYSOR 2020 Conference)* 247, 06001 (2021). DOI: <http://doi.org/10.1051/epjconf/202124706001>.
- [46] B. Kochunas, B. Collins, S. Stimpson, R. Salko, D. Jabaay, A. Graham, Y. Liu, K. S. Kim, W. Wieselquist, A. Godfrey, et al. "VERA core simulator methodology for pressurized water reactor cycle depletion". *Nuclear Science and Engineering* 185 (2017), pp. 217–231. DOI: <http://doi.org/10.13182/nse16-39>.
- [47] M. Knebel. "Advanced high-fidelity reactor simulators based on neutron transport and subchannel methodologies". Dissertation. Karlsruhe Institute of Technology, Department of Mechanical Engineering, 2016. DOI: <http://doi.org/10.5445/ir/1000072396>.
- [48] K. Wang, S. Liu, Z. Li, G. Wang, J. Liang, F. Yang, Z. Chen, X. Guo, Y. Qiu, Q. Wu, J. Guo, and X. Tang. "Analysis of BEAVRS two-cycle benchmark using RMC based on full core detailed model". *Progress in Nuclear Energy* 98 (2017), pp. 301–312. DOI: <http://doi.org/10.1016/j.pnucene.2017.04.009>.
- [49] J. Yu, H. Lee, H. Kim, P. Zhang, and D. Lee. "Coupled neutronics-thermal-hydraulic simulation of BEAVRS cycle 1 depletion by the MCS/CTF code system". *Nuclear technology* 206 (2019), pp. 728–742. DOI: <http://doi.org/10.1080/00295450.2019.1677107>.
- [50] A. S. Ivanov. "High fidelity Monte Carlo based reactor physics calculations". Dissertation. Karlsruhe Institute of Technology, Department of Mechanical Engineering, 2014. DOI: <http://doi.org/10.5445/ir/1000050615>.
- [51] D. Ferraro, M. García, U. Imke, V. Valtavirta, R. Tuominen, Y. Bilodid, J. Leppänen, and V. Sanchez-Espinoza. "Serpent/SUBCHANFLOW coupled calculations for a VVER core at hot full power". *EPJ Web Conference (PHYSOR 2020 Conference)* 247, 04006 (2021). DOI: <http://doi.org/10.1051/epjconf/202124704006>.
- [52] D. Ferraro, M. García, U. Imke, V. Valtavirta, R. Tuominen, J. Leppänen, and V. Sanchez-Espinoza. "Serpent/SUBCHANFLOW coupled burnup calculations for VVER fuel assemblies". *EPJ Web Conference (PHYSOR 2020 Conference)* 247, 04005 (2021). DOI: <http://doi.org/10.1051/epjconf/202124704005>.

- [53] J. Yu, H. Lee, M. Lemaire, H. Kim, P. Zhang, and D. Lee. "Fuel performance analysis of BEAVRS benchmark Cycle 1 depletion with MCS/FRAPCON coupled system". *Annals of Nuclear Energy* 138, 107192 (2020). DOI: <http://doi.org/10.1016/j.anucene.2019.107192>.
- [54] J. Yu, H. Lee, M. Lemaire, H. Kim, P. Zhang, and D. Lee. "MCS based neutronics/thermal-hydraulics/fuel-performance coupling with CTF and FRAPCON". *Computer Physics Communications* 238 (2019), pp. 1–18. DOI: <http://doi.org/10.1016/j.cpc.2019.01.001>.
- [55] S. Stimpson, K. Clarno, R. Pawlowski, R. Gardner, J. Powers, B. Collins, A. Toth, S. Novascone, S. Pitts, J. Hales, and G. Pastore. "Coupled fuel performance calculations in VERA and demonstration on Watts Bar unit 1, cycle 1". *Annals of Nuclear Energy* 145, 107554 (2020). DOI: <http://doi.org/10.1016/j.anucene.2020.107554>.
- [56] J. R. Basualdo, V. H. Sanchez-Espinoza, R. Stieglitz, and R. Macián-Juan. "PARCS-SUBCHANFLOW-TRANSURANUS multiphysics coupling for improved PWR's simulations". *ICAPP 2017 Conference, Fukui and Kyoto, Japan, April, 2017*. 2017. ISBN: 978-4-890-47167-6.
- [57] Jaakko Leppänen and A. Isotalo. "Burnup calculation methodology in the Serpent 2 Monte Carlo code". *PHYSOR 2012 Conference, Knoxville, USA, April, 2012*. 2012, pp. 924–935. ISBN: 978-1-622-76389-4.
- [58] J. Dufek, D. Kotlyar, and E. Shwageraus. "The stochastic implicit Euler method - a stable coupling scheme for Monte Carlo burnup calculations". *Annals of Nuclear Energy* 60 (2013), pp. 295–300. DOI: <http://doi.org/10.1016/j.anucene.2013.05.015>.
- [59] V. Valtavirta. "Development and applications of multiphysics capabilities in a continuous energy Monte Carlo neutron transport code". Dissertation. Aalto University, School of Science, Department of Applied Physics, 2017. ISBN: 978-9-526-07376-7.
- [60] D. Basile, R. Chierici, M. Beghi, E. Salina, and E. Brega. "COBRA-EN, an updated version of the COBRA-3C/MIT code for thermal-hydraulic transient analysis of Light Water Reactor fuel assemblies and cores". *ENEL-CRTN Compartimento di Milano, 1010/1 Technical Report* (1999).
- [61] D. L Hagrman and G. A. Reymann. "MATPRO-Version 11: a handbook of materials properties for use in the analysis of light water reactor fuel rod behavior". *Idaho National Laboratory, NUREG/CR-0497, TREE-1280 Technical Report* (1979).
- [62] M. E. Cunningham and C. E. Beyer. "GT2R2: an updated version of GAPCON-THERMAL-2". *Pacific Northwest National Laboratory, NUREG/CR-3907, PNL-5178 Technical Report* (1984).
- [63] P. Van Uffelen, A. Schubert, J. van de Laar, C. Györi, D. Elenkov, S. Boneva, M. Georgieva, S. Georgiev, Z. Hózer, D. Märten, et al. "The verification of the TRANSURANUS fuel performance code - an overview". *7th International Conference on WWER Fuel Performance, Modelling and Experimental Support, Albena, Bulgaria, September, 2007*. 2007, pp. 305–320.
- [64] K. Lassmann and F. Hohlefeld. "The revised URGAP model to describe the gap conductance between fuel and cladding". *Nuclear Engineering and Design* 103 (1987), pp. 215–221. DOI: [http://doi.org/10.1016/0029-5493\(87\)90275-5](http://doi.org/10.1016/0029-5493(87)90275-5).

- [65] A. Schubert, P. Van Uffelen, J. van de Laar, C. T. Walker, and W. Haeck. "Extension of the TRANSURANUS burn-up model". *Journal of Nuclear Materials* 376 (2008), pp. 1–10. DOI: <http://doi.org/10.1016/j.jnucmat.2008.01.006>.
- [66] M. García, R. Tuominen, A. Gommlich, D. Ferraro, V. Valtavirta, U. Imke, P. Van Uffelen, L. Mercatali, V. Sanchez-Espinoza, J. Leppänen, and S. Kliem. "A Serpent2-SUBCHANFLOW-TRANSURANUS coupling for pin-by-pin depletion calculations in Light Water Reactors". *Annals of Nuclear Energy* 139, 107213 (2020). DOI: <http://doi.org/10.1016/j.anucene.2019.107213>.
- [67] D. Ferraro, M. García, U. Imke, V. Valtavirta, J. Leppänen, and V. Sanchez-Espinoza. "Serpent/SCF pin-level multiphysics solutions for the VERA Fuel Assembly Benchmark". *Annals of Nuclear Energy* 128 (2019), pp. 102–114. DOI: <http://doi.org/10.1016/j.anucene.2018.12.047>.
- [68] M. García, D. Ferraro, V. Valtavirta, U. Imke, R. Tuominen, V. Sanchez-Espinoza, and L. Mercatali. "Development of an object-oriented Serpent2-SUBCHANFLOW coupling and verification with Problem 6 of the VERA Core Physics Benchmark". *M&C 2019 Conference, Portland, USA, August, 2019* (2019), pp. 1424–1433. DOI: <http://doi.org/10.5445/ir/1000099853>.
- [69] M. García, U. Imke, D. Ferraro, V. Sanchez-Espinoza, and L. Mercatali. "Advanced modelling capabilities for pin-Level subchannel analysis of PWR and VVER reactors". *NURETH-18 Conference, Portland, USA, August, 2019* (2019), pp. 4299–4309. DOI: <http://doi.org/10.5445/ir/1000099942>.
- [70] M. García, D. Ferraro, V. Valtavirta, R. Tuominen, U. Imke, J. Leppänen, and V. Sanchez-Espinoza. "Serpent2-SUBCHANFLOW pin-by-pin modelling capabilities for VVER geometries". *Annals of Nuclear Energy* 135, 106955 (2020). DOI: <http://doi.org/10.1016/j.anucene.2019.106955>.
- [71] M. García, R. Tuominen, A. Gommlich, D. Ferraro, V. Valtavirta, U. Imke, P. Van Uffelen, L. Mercatali, V. Sanchez-Espinoza, J. Leppänen, and S. Kliem. "Serpent2-SUBCHANFLOW-TRANSURANUS pin-by-pin depletion calculations for a PWR fuel assembly". *EPJ Web Conference (PHYSOR 2020 Conference)* 247, 06016 (2021). DOI: <http://doi.org/10.1051/epjconf/202124706016>.
- [72] CEA/DEN, EDF R&D, and OPEN CASCADE. *SALOME Platform: Documentation of the Interface for Code Coupling (ICoCo)*. Accessed: 21/10/2021. URL: http://docs.salome-platform.org/latest/extra/Interface_for_Code_Coupling.pdf.
- [73] CEA/DEN, EDF R&D, and OPEN CASCADE. *SALOME Platform: MED-Coupling Users' Guide*. Accessed: 21/10/2021. URL: <http://docs.salome-platform.org/8/dev/MEDCoupling/index.html>.
- [74] D. Ferraro, M. García, L. Mercatali, V. H. Sanchez-Espinoza, J. Leppänen, and V. Valtavirta. "Foreseen capabilities, bottlenecks identification and potential limitations of Serpent MC transport code in large-scale full 3-D burnup calculations". *ICONE26 Conference, London, UK, July, 2018*, 82305 (2018). DOI: <http://doi.org/10.1115/icone26-82305>.
- [75] M. García, D. Ferraro, V. H. Sanchez-Espinoza, L. Mercatali, J. Leppänen, and V. Valtavirta. "Development of a spatial domain decomposition scheme for Monte Carlo neutron transport". *ICONE26 Conference, London, UK, July, 2018*, 82144 (2018). DOI: <http://doi.org/10.1115/icone26-82144>.

- [76] M. García, J. Leppänen, and V. Sanchez-Espinoza. “A Collision-based Domain Decomposition scheme for large-scale depletion with the Serpent 2 Monte Carlo code”. *Annals of Nuclear Energy* 152, 108026 (2021). DOI: <http://doi.org/10.1016/j.anucene.2020.108026>.
- [77] P. K. Romano, A. R. Siegel, B. Forget, and K. Smith. “Data decomposition of Monte Carlo particle transport simulations via tally servers”. *Journal of Computational Physics* 252 (2013), pp. 20–36. DOI: <http://doi.org/10.1016/j.jcp.2013.06.011>.
- [78] N. Dun, H. Fujita, J. R. Tramm, A. A. Chien, and A. R. Siegel. “Data decomposition in Monte Carlo neutron transport simulations using global view arrays”. *International Journal of High Performance Computing Applications* 29 (2015), pp. 348–365. DOI: <http://doi.org/10.1177/1094342015577681>.
- [79] M. J. O’Brien and K. I. Joy. “Domain decomposition of a Constructive Solid Geometry Monte Carlo transport code”. *M&C 2009 Conference, Saratoga Springs, United States, May, 2009*. 2009. ISBN: 978-1-615-67349-0.
- [80] T. A. Brunner and P. Brantley. “An efficient, robust, domain-decomposition algorithm for particle Monte Carlo”. *Journal of Computational Physics* 228 (2009), pp. 3882–3890. DOI: <http://doi.org/10.1016/j.jcp.2009.02.013>.
- [81] J. Liang, K. Wang, Y. Qiu, X. Chai, and S. Qiang. “Domain decomposition strategy for pin-wise full-core Monte Carlo depletion calculation with the Reactor Monte Carlo code”. *Nuclear Engineering and Technology* 48 (2016), pp. 635–641. DOI: <http://doi.org/10.1016/j.net.2016.01.015>.
- [82] J. Leppänen. “Performance of Woodcock delta-tracking in lattice physics applications using the Serpent Monte Carlo reactor physics burnup calculation code”. *Annals of Nuclear Energy* 37 (2010), pp. 715–722. DOI: <http://doi.org/10.1016/j.anucene.2010.01.011>.
- [83] G. Karypis and V. Kumar. “A fast and highly quality multilevel scheme for partitioning irregular graphs”. *SIAM Journal on Scientific Computing* 20 (1999), pp. 359–392. DOI: <http://doi.org/10.1137/s1064827595287997>.
- [84] J. C. Wagner, S. W. Mosher, T. M. Evans, D. E. Peplow, and J. A. Turner. “Hybrid and parallel domain-decomposition methods development to enable Monte Carlo for reactor analyses”. *Progress in Nuclear Science and Technology* 2 (2015), pp. 815–820. DOI: <http://doi.org/10.15669/pnst.2.815>.
- [85] MPI Forum. *MPI: A Message-Passing Interface Standard, Version 3.1*. Accessed: 21/10/2021. URL: <http://www.mpi-forum.org/docs/mpi-3.1/mpi31-report.pdf>.
- [86] M. García, D. Ferraro, V. Valtavirta, R. Tuominen, U. Imke, L. Mercatali, V. Sanchez-Espinoza, and J. Leppänen. “A subchannel coarsening method for Serpent2-SUBCHANFLOW applied to a full-core VVER problem”. *EPJ Web Conference (PHYSOR 2020 Conference)* 247, 06018 (2021). DOI: <http://doi.org/10.1051/epjconf/202124706018>.
- [87] M. García, Y. Bilodid, J. Basualdo Perello, R. Tuominen, A. Gommlich, J. Leppänen, V. Valtavirta, U. Imke, D. Ferraro, P. Van Uffelen, M. Seidl, and V. Sanchez-Espinoza. “Validation of Serpent-SUBCHANFLOW-TRANSURANUS pin-by-pin burnup calculations using experimental data from a Pre-Konvoi PWR reactor”. *Nuclear Engineering and Design* 379, 111173 (2021). DOI: <http://doi.org/10.1016/j.nucengdes.2021.111173>.

- [88] R. Rieser, D. Brosche, and P. Faber. "Planning, construction and commissioning of the Isar-2 convoy lead project". *Atw Atomwirtschaft, Atomtechnik* 33 (1988), pp. 276–284. ISSN: 0365-8414.
- [89] EU Horizon 2020, grant agreement 755097. *High-Performance Monte Carlo Methods for SAFETY Demonstration - From Proof of Concept to realistic Safety Analysis and Industry-like Applications*. Accessed: 21/10/2021. URL: <http://cordis.europa.eu/project/id/755097>.
- [90] AREVA NP Inc. "U.S. EPR Nuclear Incore Instrumentation Systems". *ANP-10271NP Technical Report* (2006).
- [91] Steinbuch Centre for Computing (SCC). *ForHLR II Documentation*. Accessed: 21/10/2021. URL: <http://scc.kit.edu/dienste/forhler2.php>.
- [92] M. García, R. Vočka, R. Tuominen, A. Gommlich, J. Leppänen, V. Valtavirta, U. Imke, D. Ferraro, P. Van Uffelen, L. Milisdörfer, and V. Sanchez-Espinoza. "Validation of Serpent-SUBCHANFLOW-TRANSURANUS pin-by-pin burnup calculations using experimental data from the Temelín II VVER-1000 reactor". *Nuclear Engineering and Technology* 53 (2021), pp. 3133–3150. DOI: <http://doi.org/10.1016/j.net.2021.04.023>.
- [93] ČEZ Group. *NPP Temelín*. Accessed: 21/10/2021. URL: <http://cez.cz/en/energy-generation/nuclear-power-plants/temelin>.
- [94] O. B. Samojlov, V. B. Kajdalov, A. A. Falkov, V. A. Bolnov, O. N. Morozkin, V. L. Molchanov, and A. V. Ugryumov. "TVSA-T fuel assembly for Temelin NPP. Main results of design and safety analyses. Trends of development". *VVER-2010: Experience and Perspectives Conference, Prague, Czech Republic, November, 2010. INIS-CZ-0100 Technical Report*. 2010.
- [95] G. G. Kim and N. Z. Cho. "Investigation of the sensitivity depletion laws for Rhodium Self-Powered Neutron Detectors (SPNDs)". *Nuclear Engineering and Technology* 33 (2001), pp. 121–131. ISSN: 0372-7327.
- [96] L. S. Tong. "Prediction of departure from nucleate boiling for an axially non-uniform heat flux distribution". *Journal of Nuclear Energy* 21 (1967), pp. 241–248. DOI: [http://doi.org/10.1016/S0022-3107\(67\)90054-8](http://doi.org/10.1016/S0022-3107(67)90054-8).
- [97] EU Horizon 2020, grant agreement 945063. *High-Performance Advanced Methods and Experimental Investigations for the Safety Evaluation of Generic Small Modular Reactors*. Accessed: 21/10/2021. URL: <http://cordis.europa.eu/project/id/945063>.
- [98] J. Leppänen and A. Isotalo. "Burnup calculation methodology in the Serpent 2 Monte Carlo code". *PHYSOR 2012 Conference, Knoxville, USA, April, 2012*. 2012, pp. 924–935. ISBN: 978-1-622-76389-4.
- [99] OECD/NEA Data Bank. "The JEFF-3.1.1 Nuclear Data Library". *OECD/NEA, OECD/NEA-22 Technical Report* (2009).

Appendix A

Mesh-based feedback exchange for a VVER fuel assembly

In the coupling scheme used in this work, Serpent, SCF and TU exchange variables with the supervisor program using unstructured meshes. This simplifies the interaction between codes, because the input and output format is standardized and the internal data structures are hidden. The mapping between the meshes and the internal models is handled by each code in a different way depending on the internal representation of the geometry. Sections A.1 to A.3 show how this is done in each code for a VVER-1000 fuel assembly with an active length of 368 cm, a fuel-assembly pitch of 23.6 cm and a pin pitch of 1.275 cm.

A.1 Serpent

As mentioned in Section 3.1.1, feedback fields in Serpent are not defined directly in the geometry used to track particles, but rather using meshes superimposed to the model. Each mesh is used to define the density and temperature for one or more materials and to tally the power distribution. Figure A.1 shows an xy-slice of the Serpent model for a VVER-1000 fuel assembly.

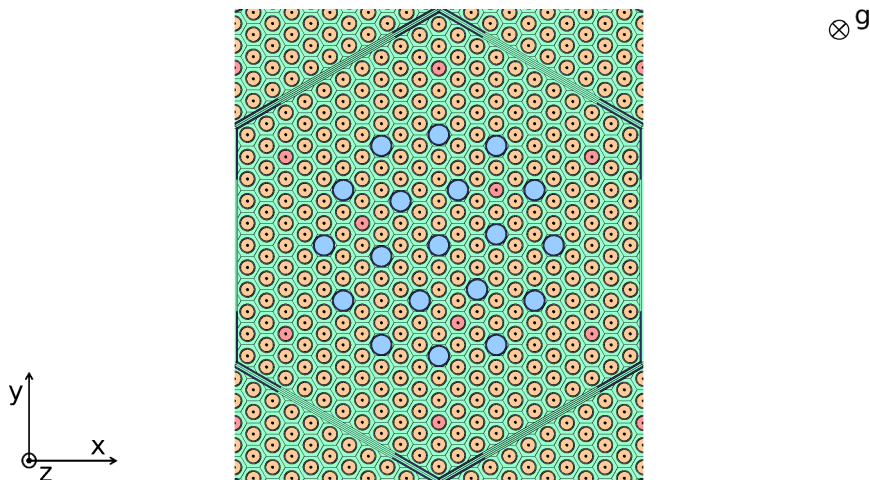


FIGURE A.1: Serpent tracking geometry (xy-cut) for a VVER-1000 fuel assembly (23.6 cm fuel-assembly pitch, 1.275 cm pin pitch) with 19 guide tubes (blue) and 9 fuel pins with burnable absorber (red). Corner stiffener plates are included and reflective boundary conditions are used to produce an infinite system.

For this model, the densities and temperatures of coolant and fuel materials can be defined with a mesh like the one shown in Figure A.2. Internally, this mesh is handled as a nested structure of regular meshes, in this case a y-type (flat sides along the y axis) hexagonal pin-level mesh nested inside an x-type fuel-assembly-level mesh with 50 axial levels along the z axis. The temperature and density for a given position are retrieved in the same way materials are found in the traditional universe-based geometry treatment used in Monte Carlo particle transport, i.e. starting at the root level and going down nested levels. For feedback exchange, this multi-level mesh is represented as a normal unstructured mesh. Pin-wise radial profiles, e.g. fuel temperatures or power distributions with a radial dependence within the fuel pellet, are treated as vector fields where each cell contains a $T(r)$ or $P(r)$ radial profile. In Serpent, the r coordinate is defined locally for each pin, with the origin at the pin center.




FIGURE A.2: Serpent feedback mesh for the VVER fuel assembly shown in Figure A.1.

A.2 SUBCHANFLOW

In the case of SCF, the graph-like geometry description does not consider explicitly the shape of the subchannels or the location of the rods, as pointed out in Section 3.1.2. To define the exact geometry of the system, two unstructured meshes are superimposed to the SCF model, one for the subchannels and one for the fuel rods. These meshes are only used to exchange feedback fields, and do not play any role in the internal solution scheme. Figure A.3 shows the SCF model for the VVER-1000 fuel assembly.

In the coolant mesh for this case, shown in Figure A.4, the shape of the cells is given by the subchannel model. This mesh is used to get coolant variables calculated by SCF such as T_{cool} and p . The fuel mesh is shown in Figure A.5, where it is clear that the actual shape of the rods is not considered, but rather cells containing the rods. The rod geometry is taken into account in the hydraulic parameters for the flow calculation, as well as in the heat-conduction solver for the fuel. The fuel mesh is used to set the power P and get rod variables, e.g. T_{fuel} and T_{clad} . The fuel temperature is either condensed into an effective Doppler temperature for each cell or treated as a radial profile $T(r)$ expressed as a vector field, as in the case of Serpent.




FIGURE A.3: SCF subchannel geometry (xy-cut) for a VVER-1000 fuel assembly (23.6 cm fuel-assembly pitch, 1.275 cm pin pitch) with 19 guide tubes (yellow) and 9 fuel pins with burnable absorber (orange). Corner stiffener plates are included and a coolant-centered subchannel model is used.




FIGURE A.4: SCF coolant mesh for the VVER fuel assembly shown in Figure A.3. The cells correspond to the subchannel geometry.




FIGURE A.5: SCF fuel mesh for the VVER fuel assembly shown in Figure A.3. The cells correspond to the rod geometry.

A.3 TRANSURANUS

To perform the coupling with Serpent and SCF, a TU module that deals with more than one rod has been developed. The solution of each rod remains independent, and an unstructured mesh is used to combine the results in a single field to exchange feedback variables. As in SCF, this mesh is used exclusively to manage input and output, and it does not intervene in the solution procedure. The TU feedback mesh for the VVER-1000 example is shown in Figure A.6. In this case the guide tubes are not simulated in TU at all. The boundary conditions for each fuel rod are specified at the coolant cell surrounding it.




FIGURE A.6: TU feedback mesh for the VVER fuel assembly shown in Figure A.3. Only the fuel pins are simulated.

Appendix B

Performance analysis for the CDD method

The performance of the CDD scheme is analyzed here for a PWR-like system made up of a 16×16 array of simplified 17×17 -pin PWR fuel assemblies like the one shown in Figure B.1. The geometry is axially uniform with no reflectors and with periodic boundary conditions, and thus the system is effectively infinite. The fuel-assembly pitch is 21.5 cm and the pin pitch is 1.26 cm. The fuel material is divided pin-wise and axially, producing a pin-by-pin three-dimensional decomposition. In all cases 200 cycles of $2 \cdot 10^6$ particles, 100 inactive and 100 active, have been simulated.

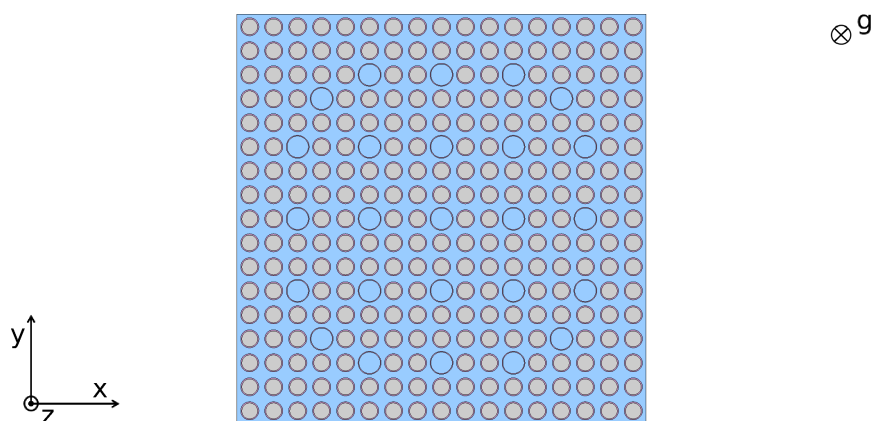


FIGURE B.1: Serpent tracking geometry (xy-cut) for the PWR fuel assembly used for the performance tests (21.5 cm fuel-assembly pitch, 1.26 cm pin pitch). Reflective boundary conditions are used to produce an infinite system.

The results shown here have been obtained in the ForHLR II high-performance computer of the Karlsruhe Institute of Technology (KIT) [91], which features Intel(R) Xeon(R) E5-2660 v3 (2.60 GHz) CPUs. Each node contains 10 cores with 2 hardware threads per core, adding up to 20 threads per node. Given that in practice, at least for massive burnup applications, the MPI-OpenMP parameters are matched to the architecture, i.e. one task per node and one thread per core, 20 OpenMP threads per MPI task have been used in all cases, and the performance data is reported as a function of the number of tasks (nodes). Using more than one task per node would multiply the in-node memory demand, which is typically the bottleneck, and is likely to degrade the speedup, since in-node multithreading tends to perform better than message-passing, at least for Serpent.

B.1 Speedup

The speedup S_N^t and efficiency ε_N^t for N MPI tasks are defined from the wall clock runtime t_N as

$$S_N^t = \frac{N_0 t_0}{t_N}, \quad \varepsilon_N^t = \frac{N_0 t_0}{N t_N} = \frac{S_N^t}{N}, \quad (\text{B.1})$$

where N_0 and t_0 are the number of tasks and runtime for a reference run, typically the sequential case ($N_0 = 1$). When the problem does not fit in the memory of a single node the runtime for $N_0 = 1$ cannot be obtained, and the case with the lowest number of tasks is used instead.

B.1.1 Multithreading

Figure B.2 shows the speedup for the two multithreading schemes described in sections 4.2.2.1 and 4.2.2.2. The performance is essentially the same, with a marginal gain using funneled multithreading at the highest number of MPI tasks. Overall the speedup is quite good, with more than 45% efficiency at 256 nodes (5,120 cores). These results suggest that the potential gain from increasing the computation/communication overlap does not compensate for the multithreading overhead, and the simpler scheme performs better.




FIGURE B.2: Speedup for funneled and serialized multithreading as a function of the number of MPI tasks. The dashed line represents the efficiency.

B.1.2 Burnup

Figure B.3 compares the strong scalability for runs with and without depletion. The burnup calculation consists of a single step using the predictor-corrector method for time integration. In this case the runtime using 32 tasks is used as reference ($N_0 = 32$ in Equation B.1), since the problem does not fit in a single node. Although a loss of performance for the burnup run can be seen, likely due to the additional

tallies and normalization of reaction rates required for the depletion calculation, the efficiency remains acceptable, at almost 40% for 256 nodes.




FIGURE B.3: Speedup for runs with and without burnup as a function of the number of MPI tasks.

B.1.3 Tracking modes

Two particle tracking modes can be used and combined in Serpent: surface and delta tracking [82]. A better speedup is achieved using surface tracking, as shown in Figure B.4, likely due to the fact that this mode is generally slower, making the tracking time larger compared to the overhead time due to CDD communications. Nevertheless, for a single MPI task delta tracking is twice as fast, making the overall performance of delta tracking better. A potential advantage of delta tracking is that a particle can go through non-local materials as long as a collision does not take place, reducing the node-to-node particle communications, but this is a marginal effect and does not seem to impact the performance.

B.2 Memory scalability

The memory scalability S_N^m and efficiency ϵ_N^m for N MPI tasks are defined from the in-node memory demand m_N as

$$S_N^m = \frac{N_0 m_0}{m_N}, \quad \epsilon_N^m = \frac{N_0 m_0}{N m_N} = \frac{S_N^m}{N}, \quad (\text{B.2})$$

where N_0 and m_0 are the number of tasks and memory demand for the reference run, typically $N_0 = 1$, as for the speedup definition.

B.2.1 Memory optimization modes

Serpent features four optimization modes for memory management, reducing the memory demand at the expense of a loss of performance [98]. Figures B.5 and




FIGURE B.4: Speedup for delta and surface tracking as a function of the number of MPI tasks.

B.6 show the memory scalability for modes 4 and 2, respectively, for runs without burnup. While both modes are based on energy grid unionization for cross-section data, in mode 4, the fastest one, macroscopic cross-sections are precalculated for each material, which is not done in mode 2 in order to reduce the overall memory footprint. As a result, for mode 4 the material data represents a larger fraction of the total memory use and the scalability is much better.




FIGURE B.5: Memory scalability (total and for material data) using optimization mode 4 as a function of the number of MPI tasks. The dashed line corresponds to the efficiency.




FIGURE B.6: Memory scalability using optimization mode 2 as a function of the number of MPI tasks.

B.2.2 Burnup

Figure B.7 shows the memory scalability for depletion calculations using optimization mode 2, taking as reference a run with 32 MPI tasks ($N_0 = 32$). In this case a larger fraction of the memory demand comes from the storage of burnable materials and the scalability improves compared to the one shown in Figure B.6. The efficiency at 256 nodes is close to 40%, which overall is quite good.




FIGURE B.7: Memory scalability for a burnup calculation as a function of the number of MPI tasks.

Appendix C

Pre-Konvoi PWR core analysis model

The Serpent, SCF and TU models for the Pre-Konvoi PWR NPP are summarized in sections C.1, C.2 and C.3, respectively.

C.1 Serpent model

The full-core pin-by-pin Serpent model is shown in figures C.2 and C.3 for xy- and xz-cuts. The geometrical model includes the core baffle, barrel and spacer grids, as well as axial and radial reflectors, and is closed with vacuum (zero incoming current) boundary conditions at the outer reflector surface. The spacers are modelled in such a way as to maintaining their mass and volume. Figure C.1 shows a close-up view of the central fuel assembly.

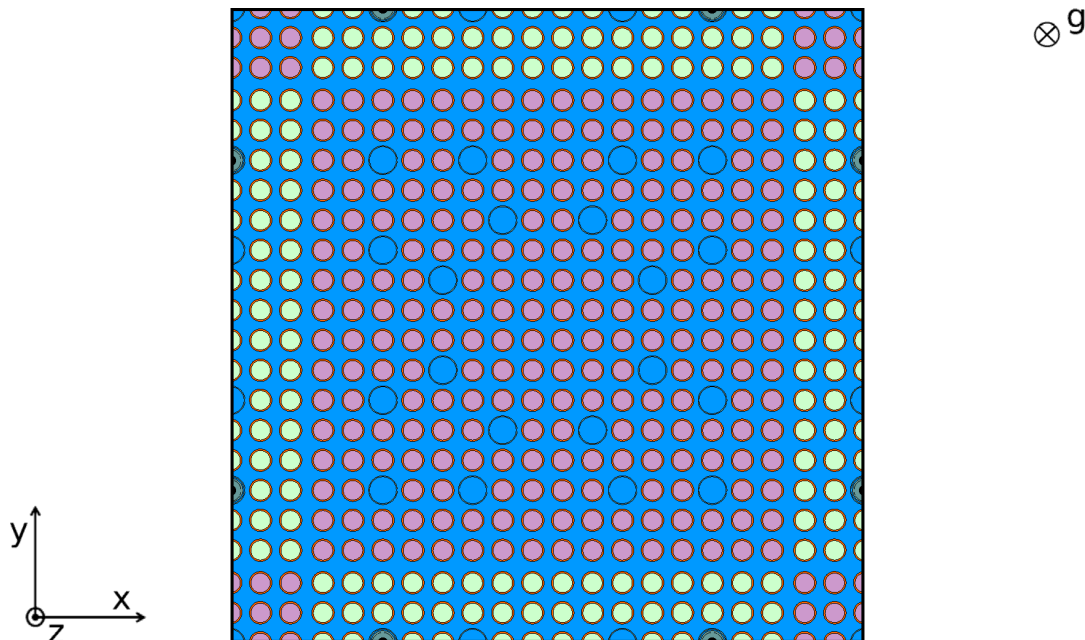


FIGURE C.1: Serpent model for the central fuel assembly of the Pre-Konvoi reactor (xy-cut).

The aeroballs are modelled in Serpent as spatial detectors tallying total neutron absorption in ^{51}V , which is the indicator material, assuming that the AMS response is proportional to this quantity.

Each transport calculation is performed using 1000 active cycles of 10^6 particles, with the criticality source calculated initially with 300 inactive cycles and corrected




FIGURE C.2: Serpent model for the Pre-Konvoi reactor (xy-cut). The origin ($x = y = 0$) is set at the center of the core.

with 50 cycles before each iteration. The critical boron concentration is calculated forcing criticality during the inactive cycles, and equilibrium xenon is assumed. The nuclear data is taken from the JEFF-3.1.1 library [99].

A pin-by-pin material division is used for the burnup calculation, with each pin divided in 16 axial equidistant zones, resulting in 728,768 fuel materials. A pin-level division with 12 axial and 3 radial zones is used for the boron silicate rods, which adds 27,648 depletion zones. The total number of burnable materials is 756,416. Figures C.4 and C.5 show the material-based domain decomposition using 64 domains for cuts along the z and y axis respectively. Each color represents the burnable materials in a given CDD domain, while all the non-burnable materials, such as the moderator and the structures, are replicated. The logic to determine the domain shapes is to use compact regions to minimize the number of particle transfers across domains [76]. It is important to note here that the material decomposition is only a computational issue and has no impact on the results of the burnup calculation, since no physical approximations are introduced by the CDD scheme.

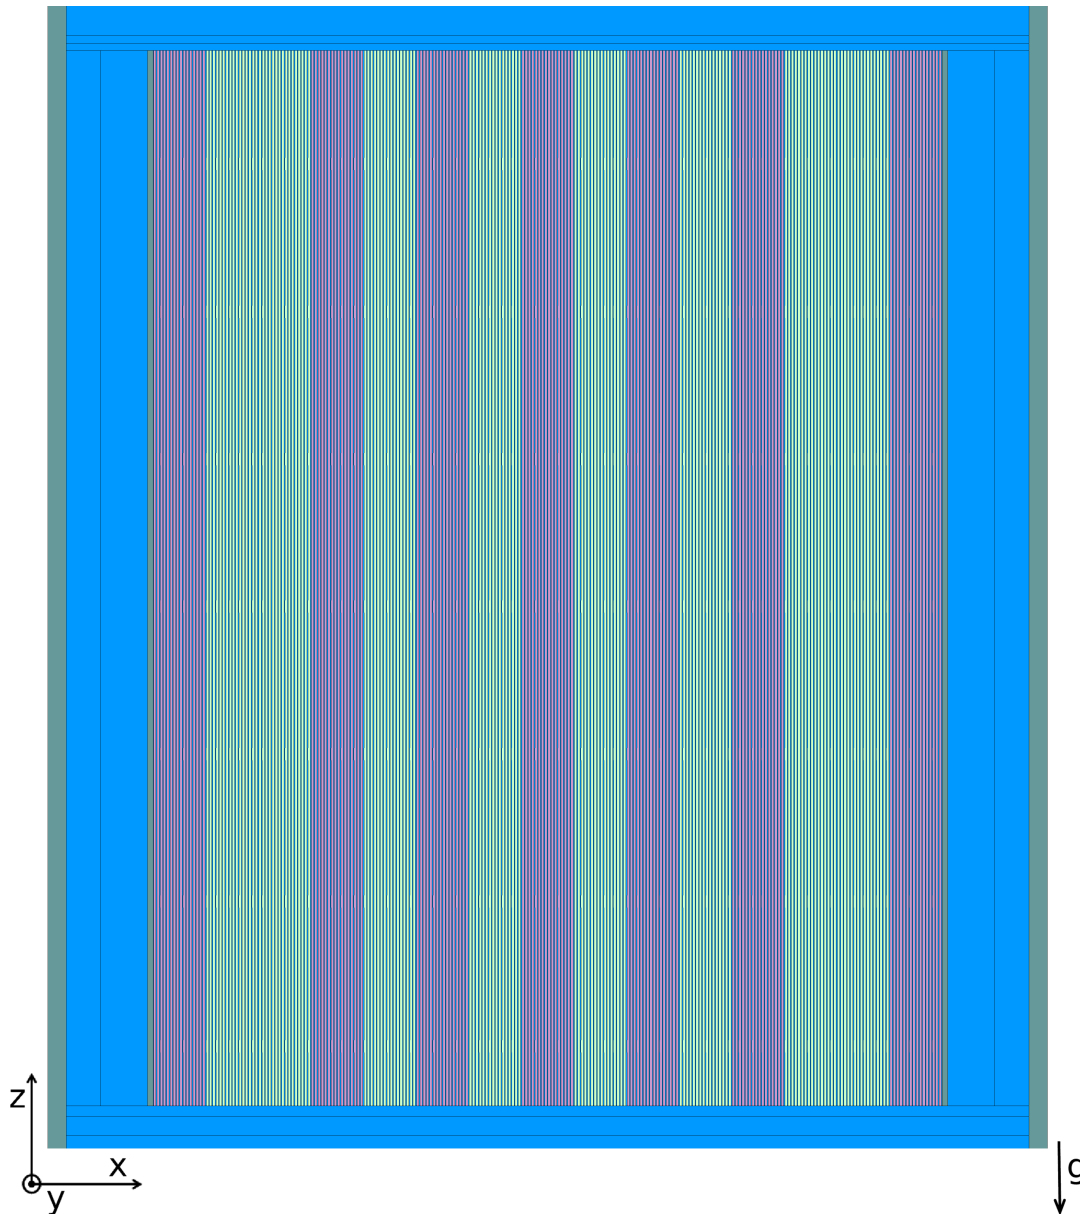


FIGURE C.3: Serpent model for the Pre-Konvoi reactor (xz-cut). The origin ($z = 0$) is set at the bottom of the active length.

C.2 SCF model

The SCF geometry for the whole core is shown in Figure C.7, while Figure C.6 presents a close-up view of the central fuel assembly. The subchannel model is coolant-centered, i.e. the control volumes are defined by the physical channels bounded by the rods, and the power of each rod is divided evenly between the four subchannels around it. The cross-flow model is made up of convective and turbulent contributions, and a constant turbulent mixing coefficient that includes the effect of spacer grids is used for the latter. The axial discretization is composed of 30 equidistant nodes for the active length, and localized pressure losses are included in the spacer grid positions using loss coefficients provided by the operator. The boundary conditions, i.e. outlet pressure, mass flow-rate and inlet temperature, correspond to Table 5.1.




FIGURE C.4: Serpent material decomposition for the Pre-Konvoi reactor using 64 domains (xy-cut). The origin ($x = y = 0$) is set at the center of the core.

When TU is used, the rods are only considered as heat sources to obtain the pressure, velocity and temperature distribution in the coolant. Otherwise, the rods are simulated in SCF using a discretization consisting of 10 equidistant nodes in the radial direction.

C.3 TU model

The TU model includes all the fuel rods in the system, while guide tubes and boron silicate rods are not considered. The radial discretization for each rod consists of 6 coarse zones, 4 in the fuel and 2 in the cladding, in which the thermomechanical properties are taken as uniform. The numerical calculation is performed on a fine mesh obtained subdividing radially each coarse zone: 5, 5, 5 and 10 fine cells for the fuel and 6 cells for each cladding zone. This results in a fine mesh with 32 radial cells, where the nodes shared between coarse zones interface the regions with different properties. Axially, each rod is discretized in 30 equidistant levels.




FIGURE C.5: Serpent material decomposition for the Pre-Konvoi reactor using 64 domains (xz-cut). The origin ($z = 0$) is set at the bottom of the active length.

The coolant temperature and pressure, as well as the cladding-coolant heat transfer coefficient, which are calculated by SCF, are used as boundary conditions on the cladding-coolant interface for each rod.




FIGURE C.6: SCF model for the central fuel assembly of the Pre-Konvoi reactor (xy-cut).




FIGURE C.7: SCF model for the Pre-Konvoi reactor (xy-cut).

Appendix D

Temelín II VVER-1000 core analysis model

The Serpent, SCF and TU models for the Temelín II VVER-1000 NPP are summarized in sections [D.1](#), [D.2](#) and [D.3](#), respectively.

D.1 Serpent model

Figures [D.2](#) and [D.3](#) show the full-core pin-by-pin Serpent model for xy- and xz-cuts. The core baffle, barrel and spacer grids, as well as axial and radial reflectors, are included in the geometrical model, and vacuum (zero incoming current) boundary conditions are imposed at the outer reflector surface. The spacers are modelled in such a way as to maintaining their mass and volume, and the stiffeners in the corners are modelled explicitly. A close-up view of the central fuel assembly is shown in Figure [D.1](#).

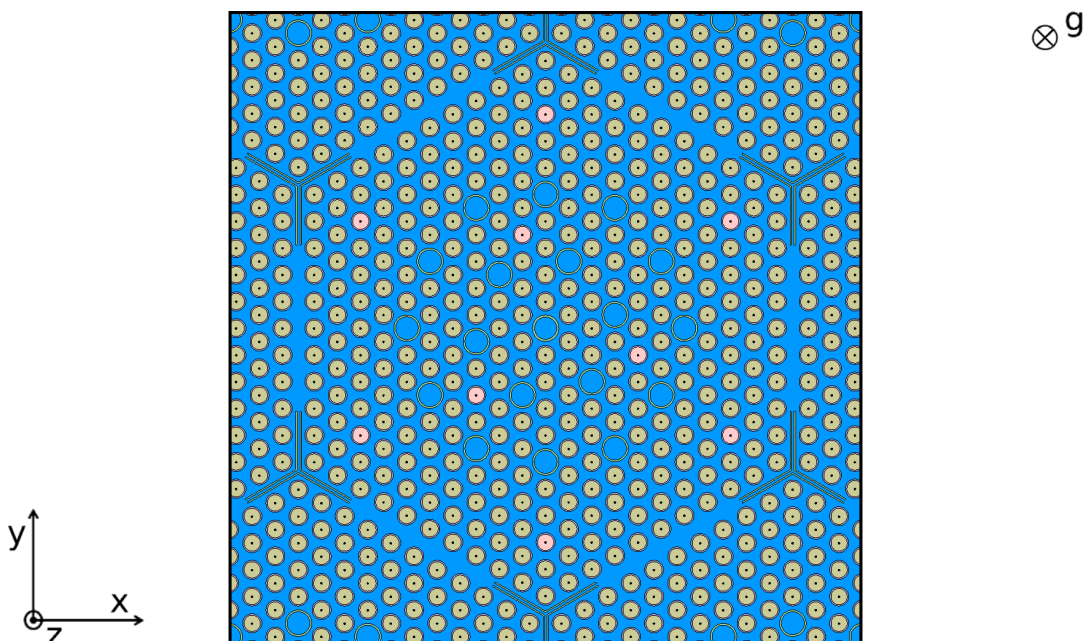


FIGURE D.1: Serpent model for the central fuel assembly of the Temelín reactor (xy-cut).

The SPNDs are modelled in Serpent as spatial detectors tallying total neutron absorption in ^{103}Rh , which is the indicator material, assuming that the SPND response is proportional to this quantity. The axial and radial power distributions, which are compared with the values reported by the CMS, are tallied directly using Serpent




FIGURE D.2: Serpent model for the Temelín reactor (xy-cut). The origin ($x = y = 0$) is set at the center of the core.

power detectors. The axial and radial detectors are infinite in the xy - and z -plane, respectively.

Each transport calculation is performed using 250 active cycles of 10^6 particles, with the criticality source calculated initially with 100 inactive cycles and corrected with 25 cycles before each iteration. The critical boron concentration is calculated forcing criticality during the inactive cycles, and equilibrium xenon is assumed. The nuclear data is taken from the JEFF-3.1.1 library [99].

A pin-by-pin material division is used for the burnup calculation, with each pin divided in 16 axial equidistant zones. Pins with burnable absorber are further subdivided in 12 radial zones. The total number of burnable materials is 909,330. Figures D.4 and D.5 show the material-based domain decomposition using 64 domains for cuts along the z and y axis respectively. Each color represents the burnable materials in a given CDD domain, while all the non-burnable materials, such as the moderator and the structures, are replicated. The logic to determine the domain shapes is to use compact regions to minimize the number of particle transfers across domains [76]. It is important to note here that the material decomposition is only a computational issue and has no impact on the results of the burnup calculation, since no physical




FIGURE D.3: Serpent model for the Temelín reactor (xz-cut). The origin ($z = 0$) is set at the bottom of the active length.

approximations are introduced by the CDD scheme.

D.2 SCF model

Figure D.7 shows the SCF geometry for the whole core, while a close-up view of the central fuel assembly is presented in Figure D.6. The subchannel model is coolant-centered, i.e. the control volumes are defined by the physical channels bounded by the rods, and the power of each rod is divided between the subchannels around it according to the corresponding heated perimeters. The cross-flow model is made up of convective and turbulent contributions, and a constant turbulent mixing coefficient that includes the effect of spacer grids is used for the latter. The axial discretization is composed of 30 equidistant nodes for the active length, and localized pressure losses are included in the spacer grid positions using loss coefficients provided by the operator. The stiffener plates are modelled increasing the




FIGURE D.4: Serpent material decomposition for the Temelín reactor using 64 domains (xy-cut). The origin ($x = y = 0$) is set at the center of the core.

wetted perimeter and removing cross-flow connections for the corner subchannels. The boundary conditions, i.e. outlet pressure, mass flow-rate and inlet temperature, correspond to Table 6.1.

When TU is used, the rods are only considered as heat sources to obtain the pressure, velocity and temperature distribution in the coolant. Otherwise, the rods are simulated in SCF using a discretization consisting of 10 equidistant nodes in the radial direction.

D.3 TU model

The TU model includes all the fuel rods in the system, while guide tubes and instrumentation tubes are not considered. The radial discretization for each rod consists of 6 coarse zones, 4 in the fuel and 2 in the cladding, in which the thermomechanical properties are taken as uniform. The numerical calculation is performed on a fine mesh obtained subdividing radially each coarse zone: 5, 5, 5 and 10 fine cells for the fuel and 6 cells for each cladding zone. This results in a fine mesh




FIGURE D.5: Serpent material decomposition for the Temelín reactor using 64 domains (xz-cut). The origin ($z = 0$) is set at the bottom of the active length.

with 32 radial cells, where the nodes shared between coarse zones interface the regions with different properties. The model includes the central hole with which the TVSA-T fuel is manufactured. Axially, each rod is discretized in 30 equidistant levels. The coolant temperature and pressure, as well as the cladding-coolant heat transfer coefficient, which are calculated by SCF, are used as boundary conditions on the cladding-coolant interface for each rod.




FIGURE D.6: SCF model for the central fuel assembly of the Temelín reactor (xy-cut).




FIGURE D.7: SCF model for the Temelín reactor (xy-cut).