Politecnico di Torino
KTH Royal Institute of Technology

Master's Degree in Nuclear Engineering

Master's Thesis

Development and Assessment of Parent-Specific Multigroup Lumped Fission Product for Fast-Spectrum Transport-Depletion Calculation

Supervisor
Jan Dufek
Sandra Dulla
Konstantin Mikityuk

Candidate
Federico Pati

2025/2026

Abstract

Continuous-energy Monte Carlo is a high-fidelity tool for reactor-physics analysis, but its cost limits its use in workflows requiring repeated transport-depletion calculations. In this context, lumped fission products provide a model-reduction strategy by representing the collective neutronic effect of the fission product inventory through a reduced set of pseudo-nuclides.

This thesis develops a methodology for generating and applying burnup-dependent multigroup lumped fission product cross sections in fast-spectrum Monte Carlo depletion. The method is implemented in OpenMC and is based on parent-specific LFPs. Effective fission yields are obtained through auxiliary depletion calculations in which heavy-metal transmutation and decay routes are suppressed, leaving fission as the only outgoing reaction channel of the selected parent. The resulting fission product inventories are then used to construct yield-weighted multigroup microscopic cross sections for the LFP. A Python framework is also developed to use these libraries in a coupled multigroup transport-depletion loop, selecting and interpolating the appropriate cross-section data at each burnup step.

The methodology is assessed on a Superphénix-like MOX pin cell and tested for transferability on an ESFR-like pin cell. In both cases the coupled MG LFP calculation reproduces the continuous-energy multiplication factor within about 200 pcm over a burnup range of 0–100 MWd/kgHM. Integral quantities such as total heavy-metal inventory, total fission product concentration, and aggregate reaction rates are reproduced with small deviations, while the fission product capture-rate error remains within about 0.5% of the total fuel capture rate at end of life.

Comparing with the CE reference, the MG LFP calculation is about 76 times faster for the Superphénix case and about 127 times faster for the ESFR case. Although library generation has a non-negligible upfront cost, it can be amortised if libraries are reused across multiple calculations in a similar spectrum regime. The results show that parent-specific, burnup-dependent LFPs can provide an effective compromise between neutronic accuracy and computational efficiency for fast-reactor Monte Carlo depletion analyses.

Contents

List of Tables

List of Figures

Acronyms

BU burnup

CE Continuous Energy

CE/CM Constant Extrapolation/Constant Midpoint

CRAM Chebyshev Rational Approximation Method

EOL end of life

ESFR European Sodium Fast Reactor

GNDS Generalized Nuclear Database Structure

HM heavy metal

LFP Lumped Fission Products

MC Monte Carlo

MG Multi-group

MGXS multi-group cross sections

ODS oxide dispersion-strengthened

SFR Sodium-cooled Fast Reactor

SPX Superphénix

Chapter 1

Introduction

1.1 Computational Methods for Neutronics

The numerical solution of the neutron transport equation has always involved a trade-off between physical fidelity and practical feasibility. From the beginning of reactor analysis, the central problem was not only to represent neutron behaviour accurately, but also to obtain solutions within the limits imposed by available hardware. For this reason, the evolution of neutron transport methods should not be viewed as a linear sequence in which one method replaces another, but rather, it reflects the progressive development of complementary tools, each providing a different balance between approximation, robustness, and execution cost.

The neutron transport equation provides a mathematical description of neutron motion, interaction, and multiplication in matter. It is a balance equation in which the time rate of change of the neutron population is determined by sources, scattering, and fission, minus losses due to leakage and interactions. In integro-differential form, it can be written as:

$$\begin{aligned}\left(\frac{1}{v(E)}\frac{\partial}{\partial t} + \mathbf{\Omega}\cdot\nabla + \Sigma_t(\mathbf{r},E,t)\right)\psi(\mathbf{r},\mathbf{\Omega},E,t) = &\frac{\chi_p(E)}{4\pi}\int_0^\infty dE'\,\nu(E')\Sigma_f(\mathbf{r},E',t)\,\phi(\mathbf{r},E',t)\\ &+\int_{4\pi} d\mathbf{\Omega}'\int_0^\infty dE'\,\Sigma_s(\mathbf{r},\mathbf{\Omega}'\to\mathbf{\Omega},E'\to E,t)\psi(\mathbf{r},\mathbf{\Omega}',E',t)\\ &+\sum_{i=1}^{N}\frac{\chi_{di}(E)}{4\pi}\lambda_i C_i(\mathbf{r},t) + Q(\mathbf{r},\mathbf{\Omega},E,t).\end{aligned} \tag{1.1}$$

where:

- ψ and ϕ are, respectively, the angular and scalar neutron flux;
- χ_p and χ_d are, respectively, the prompt and delayed neutron energy spectra;
- Σ_t, Σ_s, and Σ_f are the macroscopic total, scattering, and fission cross sections.

In the earliest stages of reactor computation, deterministic methods were adopted because they made the transport problem tractable on the hardware available at the time. In practical reactor calculations, reduced descriptions based on coarse energy discretisation and simplified angular treatments were unavoidable. This historical need strongly influenced the development of classical Multi-group (MG) methodologies. As a first approximation of Eq. 1.1, the neutron diffusion equation was widely adopted. The diffusion equation is derived under the assumptions of nearly isotropic angular flux, weak spatial gradients, and limited anisotropy in scattering. Under these conditions, the neutron current can be approximated by Fick's law:

$$\frac{1}{v(E)}\frac{\partial\phi(\mathbf{r},t)}{\partial t} - \nabla\cdot(D\nabla\phi(\mathbf{r},t)) + \Sigma_a\phi(\mathbf{r},t) = \nu\Sigma_f\phi(\mathbf{r},t) + Q(\mathbf{r},t). \tag{1.2}$$

This approximation is valid only several neutron mean free paths away from geometric boundaries and from strong sources of angular anisotropy.

As deterministic transport methods matured, attention shifted from highly approximate treatments toward formulations that preserve more of the Boltzmann equation while remaining feasible for large-scale calculations. Progress came through the gradual refinement of these approximations, particularly in energy and spatial discretisation, angular dependence, and numerical solution strategies [1].

A decisive milestone was the emergence of the discrete-ordinates method. Carlson's 1953 report [2] is one of the foundational references in neutron transport numerics. In modern terms, the S_n method replaces the continuous angular variable with a finite set of directions, thereby converting the transport equation into a coupled set of equations that can be solved numerically. Later developments improved the accuracy and efficiency of the same general multigroup deterministic framework, while finite-element and nodal methods were developed for related purposes [3, 4].

In the same years, the appearance of Monte Carlo methods introduced a different philosophy for solving transport problems. Metropolis and Ulam [5] proposed Monte Carlo as a statistical procedure for mathematical and physical problems that are difficult to solve analytically. Applied to neutron transport, this idea shifted the emphasis from directly solving a discretised equation to sampling particle histories governed by collision probabilities. The appeal of the method was clear: complex geometries, continuous-energy nuclear data, and detailed collision physics can be treated explicitly without many of the approximations intrinsic to deterministic transport. Continuous-energy Monte Carlo therefore became a high-fidelity reference standard for many reactor-physics and criticality analysis.

However, Monte Carlo did not solve the performance constraints that originally motivated deterministic methods. Brown and co-authors [6] describe the history of Monte Carlo criticality calculations as closely linked to the history of computers: early applications were restricted by speed and memory, while later advances in hardware and algorithms enabled the detailed analyses that are routine today. The same review also emphasizes the CPU-intensive nature of Monte Carlo transport. This point is especially relevant for reactor applications requiring repeated evaluations over changing physical states, such as burnup progression, thermal-hydraulic feedback, control-rod movement, or transient safety analysis. In such settings, continuous-energy Monte Carlo is often most

useful as a reference calculation, validation tool, or offline data-generation method, rather than as the most convenient solver for every stage of the analysis.

For these reasons, modern transport analysis increasingly relies on combinations of deterministic and stochastic methods. Wagner and co-authors review hybrid deterministic/-Monte Carlo methods developed at Oak Ridge National Laboratory, in which deterministic calculations provide forward or adjoint flux estimates used to generate variance-reduction parameters for Monte Carlo simulations [7]. Approximate adjoint solutions provide an importance function relative to a target response, allowing Monte Carlo histories to be sampled preferentially in regions, energies, and trajectories that contribute most to that response. For spatially distributed responses, forward flux information can also be incorporated to distribute the sampling effort more evenly across the tallies of interest. The key point is that deterministic and Monte Carlo methods are not mutually exclusive: deterministic solutions can inform and accelerate Monte Carlo simulations without replacing the underlying stochastic framework.

This hybrid perspective is particularly relevant for multigroup cross sections. Even in an era of powerful continuous-energy Monte Carlo tools, accurate multi-group cross sections (MGXS) remain indispensable whenever transport solutions must be performed rapidly and repeatedly. Boyd and co-authors [8] state this point clearly: high-fidelity deterministic transport codes require accurate multigroup cross sections, and Monte Carlo can be used as a reactor-agnostic route for MGXS generation because it is not constrained by approximations embedded in traditional deterministic cross section production tools. Thus, Monte Carlo does not make multigroup data obsolete. Instead, it enables improved offline generation of multigroup data, which can then be used online in more efficient transport calculations.

1.2 Computational Burden of Monte Carlo

High-fidelity Continuous Energy (CE) Monte Carlo (MC) transport remains expensive, particularly when many independent simulations are required. Because the method is stochastic, the statistical uncertainty of tallied quantities decreases slowly with the number of particle histories. The standard deviation of a Monte Carlo estimator scales as $1/\sqrt{N}$, where N is the number of histories; therefore, reducing the uncertainty by a factor of two requires approximately four times as many histories.

This scaling becomes restrictive in workflows based on repeated reactor-physics evaluations. Sensitivity analyses, uncertainty-quantification campaigns, machine-learning training-set generation, and loading-pattern optimisation can require hundreds or thousands of independent simulations. Multiphysics workflows further increase the cost when neutronics is coupled to thermal-hydraulics or fuel thermo-mechanics. In such contexts, a single continuous-energy Monte Carlo calculation lasting several hours constitutes a prohibitive bottleneck as an inner-loop solver.

Recent studies illustrate the magnitude of the burden. Continuous-energy Monte Carlo calculations for homogenized cross-section generation have been reported to require about 62 CPU-hours for the MET-1000 benchmark and 85.5 CPU-hours for a 100 MWe lead-cooled fast reactor [9]. Similarly, Serpent-based studies on Monte-Carlo generation of

homogenized group constants for nodal simulators conclude that the approach is feasible, but still costly when many state points are required [10].

For this reason, multigroup transport remains relevant today. Its role is not to replace high-fidelity Monte Carlo, but to shift part of the cost from repeated online transport calculations to offline data generation. As Yoshioka et al. observe, a practical route in modern reactor analysis is to use continuous-energy Monte Carlo for high-accuracy data generation and then replace the detailed continuous-energy description with a reduced energy representation that can be reused efficiently in deterministic or multigroup Monte Carlo frameworks [11].

The burden becomes more severe in coupled transport-depletion calculations. The transport equation must be solved repeatedly over the irradiation history to supply the fluxes and spectrum-averaged reaction rates required to integrate the Bateman equations at each burnup step. If N nuclides are tracked, the depletion operator is an $N \times N$ system, with sparsity determined by the decay and transmutation pathways retained in the chain. Reducing the chain size therefore directly reduces the dimensionality of the depletion problem.

The impact of the number of tracked nuclides becomes particularly severe when depletion is performed over many independently burned regions. In high-fidelity full-core calculations, the fuel may be subdivided into a large number of depletion zones, each of which requires its own nuclide-density vector. As noted in the context of large-scale Serpent burnup calculations, full-core problems may increase the number of depletion zones up to millions, to the point that even storing the nuclide compositions becomes a significant memory burden [12]. A simple order-of-magnitude estimate illustrates the scaling: storing N_z depletion zones with N_n nuclides in double precision requires approximately $8 \times N_z \times N_n$ bytes for a single composition state. Thus, a calculation with 10^6 depletion zones and 2000 nuclides would require about 16 GB only for one nuclide-density field, before accounting for burnup histories, reaction rates, transition matrices, or parallel data structures. Reducing the number of explicitly tracked nuclides is therefore not only a matter of accelerating the matrix solution, but also of limiting the memory footprint of large depletion calculations. Although neutron transport usually remains the dominant time cost in Monte Carlo depletion, these results show that the depletion stage can become significant as the nuclide inventory and the number of depletable regions grow, especially in full-core and multi-physics applications [13].

At the same time, the size of the nuclide inventory also affects transport, although in an indirect way. The dimensionality of the transport unknown does not depend on the number of nuclides in the same explicit sense as depletion; it depends primarily on the discretization of space, angle, and energy. However, a larger nuclide set increases the amount of microscopic data that must be processed into material-wise macroscopic cross sections, the number of reaction channels considered, and the number of reaction rates that must be scored in depletion-coupled calculations. This distinction is explicit in the OpenMC depletion framework, where the authors note that reducing the depletion chain decreases not only the size of the depletion matrix, but also the number of nuclides for which reaction rates must be determined [14].

Fission products are generated continuously during irradiation and collectively represent an important parasitic absorber in operating reactors. In fast-spectrum systems,

however, most individual fission products have small or negligible transmutation cross sections compared with actinides. This motivates replacing hundreds of fission-product isotopes with a small number of Lumped Fission Products (LFP), whose effective multigroup cross sections are constructed to reproduce the integral effect of the full ensemble. The use of multigroup LFP cross sections is therefore not merely a legacy approximation, but a rational model-reduction strategy for calculations in which full explicit tracking is not economical.

In this context, the present work aims to develop and assess a methodology for generating multigroup cross-section libraries for LFP from Monte Carlo simulations, for subsequent use in multigroup Monte Carlo or deterministic transport workflows.

1.3 Historical Review of Lumped Fission Products

The use of lumped fission products, or pseudo-fission products, arose from a practical problem that appeared early in depletion and reactor-physics calculations. The neutronic effect of fission products is cumulative and important, but an explicit treatment of every fission-product nuclide was burdensome and often limited by the available nuclear data libraries. LFP methods were therefore introduced as a pragmatic reduction strategy: their purpose was to preserve the aggregate poison effect of many nuclides while keeping burnup and transport calculations tractable.

The basic idea is to represent the collective neutronic effect of the fission-product inventory by means of one or a few fictitious nuclides, each characterised by a set of energy-dependent multigroup cross sections. This reduces the number of material entries in the cross-section library and the number of nuclides tracked during depletion, while retaining an average representation of the absorption and scattering effects of fission products on the neutron economy of the reactor.

The simplest and most intuitive formulation is the **one-lump model**, in which the entire fission-product inventory is condensed into a single fictitious nuclide. For a reaction type x at neutron energy E, the lumped microscopic cross section is computed as a concentration-weighted average over the individual fission products:

$$\sigma^x_{\mathrm{LFP}}(E) = \frac{\sum_{i=1}^{N} N_i \, \sigma^x_i(E)}{\sum_{i=1}^{N} N_i}, \tag{1.3}$$

where N_i denotes the atom number density of the i-th fission product at a chosen reference burnup step and σ^x_i is its microscopic cross section for reaction x. The pseudo-nuclide is then assigned an atom density equal to the total fission-product concentration, $\sum_{i=1}^{N} N_i$. This approach reduces hundreds of individual fission products to a single library entry. Its main limitation is that, if the weighting is performed at a fixed burnup level, the resulting pseudo-nuclide cannot fully capture the evolution of fission-product composition during irradiation, introducing systematic bias.

An alternative formulation was proposed by Garrison and Roos [15]. Originally developed for thermal-reactor applications, their model classified fission products into three pseudo-fission-product groups according to their temporal behaviour under irradiation: a *rapidly saturating* group, containing short-lived nuclides that reach equilibrium quickly; a *slowly saturating* group, containing intermediate-lived nuclides; and a *non-saturating* group, containing stable or very long-lived nuclides whose concentrations grow approximately linearly with burnup. In addition, ^{135}Xe and ^{149}Sm were treated explicitly due to their large absorption cross sections. For each pseudo group, effective multigroup cross sections were generated using a concentration-weighted procedure. The Garrison–Roos data were later incorporated into the ENDF/B-III evaluated nuclear data library and became a standard fission-product representation in several fast-reactor design calculations at ANL.

Despite the more elaborate classification of fission products into three categories, the Garrison–Roos model shared with the one-lump approach a fundamental limitation: the cross sections for each pseudo group were still evaluated at a single fixed reference burnup step, making the representation inherently time-independent. Neither formulation could track the continuous evolution of the fission product composition during irradiation. The Garrison–Roos model carried additional shortcomings: the scattering matrices for the pseudo nuclides were omitted, and the underlying cross-section evaluations and fission yield data available in the early 1960s were relatively sparse, particularly in the resonance and fast energy regions. As the evaluated data in ENDF/B-III were subsequently found to contain significant deficiencies for many fission product nuclides, and improved evaluations became available with the successive releases of ENDF/B-IV and ENDF/B-V during the late 1970s, the need to revisit and update the lumped fission product representation became apparent [16].

In 1981, Atefi at Argonne National Laboratory proposed a two-lump approach in the framework of the Fast-Mixed Spectrum Reactor (FMSR) project. The fission products were divided into two groups based on mass number parity: an odd-A lump and an even-A lump. The physical reason is that the absorption cross section of the odd-A fission products is several times higher than the even one. Neutron capture will transform an odd-A nucleus into an even-A nucleus more frequently than vice versa, this difference in the transmutation rate results in the build up of even-A fission products with lower absorption cross sections σ_a. Hence, the effective σ_a drops as a function of burnup, after an initial build-up, as shown in figure 1.1. This coupling allows the model to track, in a simplified manner, the burnup-dependent redistribution of the fission product inventory. [17].

The most comprehensive investigation of this problem was carried out by Liaw and Henryson at ANL, who developed and evaluated a lumped fission-product cross-section library based on ENDF/B-V data for mixed-oxide LMFBR cores [16]. Using EPRI-CINDER-2 to compute fission-product inventories at 20 burnup steps, they constructed and compared one-lump and two-lump models. Their analysis showed that the burnup dependence of the lumped absorption cross section is small but not negligible, that the transition from ENDF/B-III to ENDF/B-V produced measurable changes in the lumped cross sections, and that the scattering component of the LFP reactivity worth is significant. They also discussed a possible double-counting issue in two-lump odd/even formulations

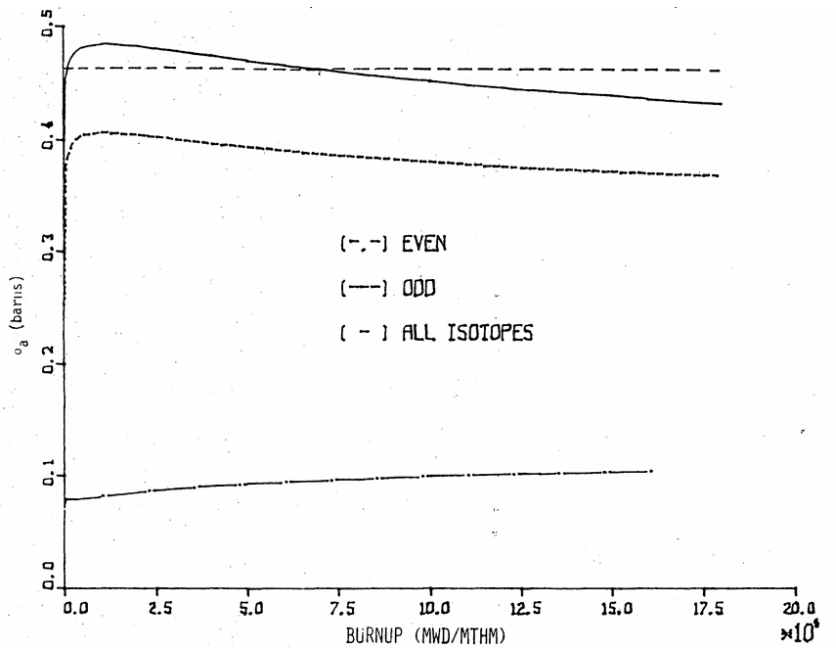

Figure 1.1: Microscopic absorption cross section of lumped fission products from ^{239}Pu fission as a function of burnup, grouped into odd-A and even-A isotopes. Adapted from [17].

when effective pseudo-nuclide inventories are defined from depletion data while inter-lump transmutations are treated explicitly, since part of the production and redistribution is already accounted in the weighting inventories. Although the study was performed for a specific mixed-oxide LMFBR design, the authors noted that its conclusions were broadly applicable to similar fast breeder reactor concepts.

Related methods were also developed in thermal-reactor lattice codes. The WIMS code family adopted a hybrid treatment in which the most important fission products are represented explicitly, while the remaining nuclides are condensed into a single pseudo-fission product. In the WIMS-D Library Update Project, the selection of explicitly treated fission products was based on their relative contribution to the total absorption rate, together with half-life, fission yield, and residence time in the reactor. Individual fission-product evaluations from ENDF were processed with NJOY to generate consistent multigroup cross sections in a reference thermal spectrum [18].

The WIMS lumped cross sections are formed using effective cumulative fission yields as weights. These yields are generated by AVRFPY, which was developed to avoid double counting in cumulative-yield data as it will be discussed in Sec. 3.2. Since the weights are not burnup-dependent, the resulting library is independent of the irradiation history. Applying these criteria, 56 fission products were retained explicitly, while 79 additional isotopes were lumped into a single pseudo-fission product.

An important extension of this approach to Monte Carlo analysis was proposed by Hanan et al. [19]. Their work was motivated by the fact that Monte Carlo codes such

as MCNP were typically used only for fresh cores, because explicit depletion with all fission products was prohibitive. The authors developed a conversion procedure to extract capture and scattering cross sections of the WIMS LFP in 69- or 172-group structures and process them into a discrete multigroup form readable by MCNP. Since the LFP cross sections depend on both spectrum and burnup, they were generated at multiple burnup points rather than as a single static set. The methodology was validated against deterministic reference calculations for four research reactor cores, with biases in effective multiplication factor typically within 0.1–0.3% $\Delta\rho$ over the burnup cycle. The study also showed that the thermal capture cross section of the lumped fission product changes significantly with burnup, confirming the need for burnup-dependent data when accurate depletion predictions are required.

A more recent application of the same conceptual approach is the work of Benkharfia et al. [20], who developed a coupled MCNP5–ORIGEN scheme for depletion analysis of the IAEA 10 MW benchmark reactor, a thermal-spectrum system fuelled with $UZrH_{1.6}$. Starting from more than 1300 fission-product yield entries, the authors identified 214 nuclides as relevant for criticality using a criterion based mainly on the product of fission yield and absorption cross section. Of these, 81 nuclides were treated explicitly, while the remaining 133 were lumped into a single pseudo-fission product.

In that methodology, the pseudo-fission-product cross sections are recomputed at each burnup step using the MIXE_ACE program. The program mixes individual nuclide cross sections using mass weights supplied by ORIGEN and writes the resulting data directly in ACE format for MCNP5. In this way, the burnup dependence of the fission-product composition is recovered, although at the cost of an additional cross-section mixing step in each iteration of the coupled scheme. When all 214 selected fission products were treated explicitly, the calculated k_{eff} agreed well with the ANL REBUS-2 reference, with larger discrepancies at higher burnup. When the 133 low-importance nuclides were lumped into the pseudo-fission product, the agreement with the fully explicit calculation remained very good throughout the cycle.

1.4 Objectives and Thesis Outline

The studies reviewed above share a common thread: the LFP representation is not a monolithic concept but a family of approaches, each making different trade-offs between physical fidelity, burnup dependence, and computational cost. The present work builds on this foundation by addressing the specific case of a fast-spectrum MOX core, where the absence of dominant thermal poisons makes the aggregate fission product cross section particularly sensitive to the condensation methodology. The approach adopted here differs from earlier studies in that the multigroup LFP data are generated entirely within a continuous-energy Monte Carlo framework, avoiding the deterministic approximations that have historically limited the accuracy of the condensed representation.

The remainder of this thesis is organised as follows. Chapter 2 introduces the computational framework and the reference case adopted throughout the work. The OpenMC-based workflow is described, the reference Superphénix-like MOX pin-cell model is defined, and the continuous-energy depletion calculation used as reference is presented. The

chapter also describes the generation of the burnup-dependent multigroup cross-section libraries on the ECCO-33 energy-group structure.

Chapter 3 presents the lumped fission product methodology and its implementation. The main modelling choices are first introduced, including the full lumping of the fission product inventory, the reduced heavy-metal set, and the custom depletion-chain generation. The chapter then discusses the construction of effective fission yields from ENDF data and the yield-weighted generation of parent-specific LFP microscopic cross sections for use in OpenMC.

Chapter 4 applies and assesses the LFP methodology in a stepwise manner. The analysis first isolates the cross-section effects using CE fluxes and concentrations, and then performs multigroup depletion calculations driven by CE-supplied fluxes. This staged approach is used to identify the main sources of approximation introduced by the LFP representation and to quantify their impact on reaction rates, fission product capture, heavy-metal evolution, and system reactivity.

Chapter 5 introduces the Python framework developed to perform coupled multigroup transport and depletion with burnup-dependent LFP libraries. The physical and numerical formulation of the coupling is described, followed by the implementation of the library manager and coupled depletion driver. The framework is then verified using explicit fission products and applied to the LFP calculation for the reference SFR pin-cell case.

Chapter 6 examines the robustness and transferability of the methodology. The sensitivity to burnup-grid interpolation is assessed, the libraries are tested on a different fast-reactor pin-cell configuration representative of the ESFR, and the computational cost of the proposed approach is compared with continuous-energy reference calculations. Finally, Chapter 7 summarises the main conclusions of the work and outlines possible directions for future developments.

Chapter 2

Computational Framework and Reference Case

This chapter defines the computational framework and the reference case used throughout the thesis. The aim is not to provide a general description of OpenMC, but to specify the parts of the code and of the model that are required to reproduce the continuous-energy depletion calculation and the subsequent multigroup cross-section generation workflow.

The chapter is organised as follows. First, the OpenMC features used in the methodology are introduced, with emphasis on continuous-energy Monte Carlo transport, temperature treatment, depletion, and tallying. The reference sodium-cooled fast-reactor pin-cell model is then described in terms of geometry, materials, and temperature assignments. Finally, the continuous-energy depletion calculation and the multigroup cross-section generation procedure are presented.

2.1 OpenMC-Based Computational Framework

The methodology developed in this work is implemented in OpenMC, an open-source Monte Carlo particle-transport code originally developed at the Massachusetts Institute of Technology and first publicly released in 2012 [21]. OpenMC was selected because it combines continuous-energy transport, multigroup transport, depletion capabilities, and a Python application-programming interface within the same computational environment.

In this thesis, OpenMC is not used only as a transport solver. It also provides the infrastructure to define the pin-cell model, perform the reference continuous-energy depletion calculation, reconstruct burnup-dependent material compositions, generate multigroup cross sections, and organise the resulting data for later analysis. This integration is important because the proposed lumped-fission-product methodology requires repeated access to both transport results and evolving nuclide inventories.

Two features are particularly relevant for the present work. The first is the native support for both continuous-energy and multigroup transport. The reference solution is obtained using continuous-energy nuclear data, whereas the reduced fission-product representation is developed in multigroup form. Using the same code environment for both

stages reduces inconsistencies associated with geometry, material definitions, boundary conditions, and post-processing.

The second feature is the Python interface. Most operations required in this work are automated: depletion-result extraction, material reconstruction at selected burnup points, MGXS tally generation, batch execution, and storage of the final libraries. The Python API makes these operations reproducible and allows the workflow to be run systematically over all burnup states considered in the analysis.

2.1.1 Continuous-Energy Monte Carlo Transport

In continuous-energy mode, OpenMC simulates each neutron history as a sequence of free flights interrupted by interactions with the medium. The microscopic cross sections $\sigma_i^x(E)$ of nuclide i for reaction type x are read from evaluated nuclear-data libraries. From these, the macroscopic total cross section of a material is then constructed as

$$\Sigma_{\mathrm{tot}}(E) = \sum_i N_i \, \sigma_i^{\mathrm{tot}}(E), \tag{2.1}$$

where N_i is the atom number density of nuclide i. The distance ℓ to the next collision is sampled from

$$\ell = -\frac{\ln \xi}{\Sigma_{\mathrm{tot}}(E)}, \tag{2.2}$$

where ξ is a pseudo-random number uniformly distributed in (0,1]. Once the collision point has been determined, the interacting nuclide is sampled according to

$$P_i = \frac{N_i \, \sigma_i^{\mathrm{tot}}(E)}{\Sigma_{\mathrm{tot}}(E)}. \tag{2.3}$$

Then, having selected the nuclide, the specific reaction channel is then sampled from the partial cross sections

$$P_x = \frac{\sigma_i^x(E)}{\sigma_i^{\mathrm{tot}}(E)}. \tag{2.4}$$

The outgoing neutron energy and direction are sampled from the secondary-particle distributions stored in the nuclear-data files. For eigenvalue calculations, OpenMC uses the method of successive generations, in which a fixed number of neutron histories are tracked from birth to death in each cycle, and the fission sites sampled in one cycle define the source distribution for the next.

2.2 Reference Pin-Cell Model

To evaluate the performance of the lumped fission product methodology in a modern continuous-energy Monte Carlo framework, a depletion calculation was carried out on a single fuel pin cell representative of a Sodium-cooled Fast Reactor (SFR). The geometry

and material composition are derived from Superphénix (SPX) benchmark data and are used here as a generic model for MOX-fuelled, sodium-cooled fast-spectrum systems [22].

The choice of a single pin-cell as the reference computational domain is motivated by both methodological and physical considerations. From a methodological standpoint, it provides a simple and well-controlled geometry in which the proposed LFP generation procedure can be developed, tested, and interpreted without the additional complexity of full-assembly heterogeneity. At the same time, the pin-cell model captures the main local features responsible for the hard fast-reactor spectrum. In fast-spectrum systems, spectral variations at the pin-cell scale are generally less pronounced than in moderated lattices, due to the absence of strong local moderation and to the relatively long transport and slowing-down lengths of fast neutrons.

2.2.1 Geometry

The model consists of a cylindrical fuel pellet surrounded by a helium-filled gap, stainless-steel cladding, and liquid sodium coolant. These regions are placed inside a hexagonal unit cell with reflective boundary conditions on all the faces. The axial height is 1 cm, with reflective top and bottom boundaries, so that the model represents an infinite-lattice two-dimensional approximation. A cross-sectional view is shown in Figure 2.1, and the main geometric parameters are listed in Table 2.1.

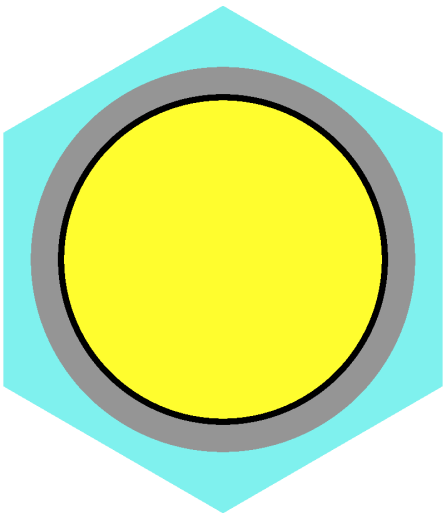

Figure 2.1: Cross-sectional view of the hexagonal pin-cell model showing fuel, gap, cladding, and sodium coolant regions.

2.2.2 Material Compositions

The fuel is a mixed uranium–plutonium oxide (MOX) containing minor actinides, representative of a fast-reactor fuel composition. The isotopic composition in atom fractions

Table 2.1: Pin-cell geometric parameters.

Region	Outer radius [cm]
Fuel pellet	0.357
Helium gap	0.371
Cladding	0.4321
Pin pitch	0.9866

is given in Table 2.2. The cladding is an austenitic stainless steel whose composition is reported in Table 2.3. The gap is filled with helium, and the coolant is liquid sodium.

Table 2.2: Fuel isotopic composition, with atom densities given in atoms/(b·cm).

Nuclide	Atom density	Nuclide	Atom density
U-235	9.997×10^{-5}	Pu-241	1.946×10^{-4}
U-238	1.964×10^{-2}	Pu-242	6.755×10^{-5}
Pu-238	1.761×10^{-5}	Am-241	4.764×10^{-5}
Pu-239	2.428×10^{-3}	O-16	4.598×10^{-2}
Pu-240	7.239×10^{-4}		

2.2.3 Temperature Assignments

The material temperatures used in the reference calculation are reported in Table 2.4. They correspond to nominal power conditions for the selected Superphénix-based model.

2.2.4 Temperature Treatment

Nuclear-data libraries are available only at a finite set of tabulated temperatures. Reactor calculations, however, generally require material temperatures that do not coincide exactly with those tabulated values. OpenMC therefore provides specific treatments for intermediate temperatures, including nearest-temperature selection and statistical interpolation [23].

In the present work, the interpolation method is used through the Python setting

```
settings.temperature = {'method': 'interpolation', 'range': (600.0,
                                1500.0)}.
```

With this option, OpenMC loads the two tabulated temperature datasets that bound the temperature assigned to each material. During transport, cross sections are evaluated by statistical linear-linear interpolation between the bounding temperatures. Specifically, if a material is at temperature T and the available tabulated temperatures satisfy $T_i < T < T_{i+1}$, a uniformly distributed random number $\xi \in [0,1)$ is sampled. The data at T_{i+1} are used if

Table 2.3: Cladding isotopic composition, with atom densities given in atoms/(b·cm).

Nuclide	Atom density	Nuclide	Atom density
C-nat	1.955×10^{-4}	Fe54	3.291×10^{-3}
Si28	9.281×10^{-4}	Fe56	4.982×10^{-2}
Si29	4.552×10^{-5}	Fe57	1.130×10^{-3}
Si30	2.904×10^{-5}	Fe58	1.478×10^{-4}
P31	4.545×10^{-5}	Ni58	7.719×10^{-3}
Ti46	3.370×10^{-5}	Ni60	2.874×10^{-3}
Ti47	2.974×10^{-5}	Ni62	3.856×10^{-4}
Ti48	2.886×10^{-4}	Ni64	9.510×10^{-5}
Ti49	2.074×10^{-5}	Mo92	1.855×10^{-4}
Ti50	1.947×10^{-5}	Mo94	1.143×10^{-4}
Cr50	6.531×10^{-4}	Mo95	1.958×10^{-4}
Cr52	1.211×10^{-2}	Mo96	2.039×10^{-4}
Cr53	1.347×10^{-3}	Mo97	1.162×10^{-4}
Cr54	3.292×10^{-4}	Mo98	2.922×10^{-4}
Mn55	1.452×10^{-3}	Mo100	1.153×10^{-4}

Table 2.4: Material temperature assignments in the reference pin cell.

Region	Temperature [K]
Fuel (MOX)	1500
Helium gap	673
Cladding (steel)	673
Coolant (Na)	673

$$\xi < \frac{T - T_i}{T_{i+1} - T_i}, \tag{2.5}$$

and the data at T_i are used otherwise. In expectation, this procedure reproduces the corresponding linear interpolation between the two bounding temperatures [23].

2.3 Reference Continuous-Energy Depletion

After defining the previously model parameters, the reference depletion simulation is performed with OpenMC using continuous-energy cross sections from the ENDF/B-VIII.0 nuclear data library. The calculation is carried out using the `CoupledOperator` class to couple the Monte Carlo neutron transport solver and the depletion solver based on the Chebyshev Rational Approximation Method (CRAM). Time integration is performed with the Constant Extrapolation/Constant Midpoint (CE/CM) predictor–corrector scheme, and the adopted depletion chain is based on ENDF/B-VIII.0 data. The main depletion

parameters are summarised in Table 2.5.

The constant power used for normalization is selected such that the fuel reaches a target burnup of 100 MWd/kgHM after 5 years of irradiation. Given the heavy metal (HM) mass contained in the pin cell, $m_{\mathrm{HM}} = 3.67$ g, the required constant power level is obtained from

$$\frac{Pt}{m_{\mathrm{HM}}} = 100 \ \frac{\mathrm{MWd}}{\mathrm{kgHM}}, \tag{2.6}$$

which gives

$$P = \frac{100 \ \frac{\mathrm{MWd}}{\mathrm{kgHM}} \cdot 3.67 \times 10^{-3} \ \mathrm{kgHM}}{1825 \ \mathrm{d}} = 201.1 \ \mathrm{W}. \tag{2.7}$$

The time discretisation (Tab. 2.6) uses short initial time steps to resolve the rapid build-up of short-lived fission products and strong absorbers. Progressively longer intervals are then used to describe the slower long-term evolution of the fuel composition.

Table 2.5: Main parameters of the reference continuous-energy depletion calculation.

Parameter	**Value**
Thermal power	201.1 W
Final burnup	100 MWd/kgHM
Number of burnup steps	20
Integrator	CE/CM predictor–corrector
Cross-section library	ENDF/B-VIII.0
Transmutation chain	`chain_endfb80_pwr.xml`

2.3.1 Transport-Coupled Depletion in OpenMC

Depletion calculations in OpenMC are handled by the `openmc.deplete` Python subpackage. The depletion problem is formulated through the Bateman equations, which describe the time evolution of the nuclide concentration vector $\mathbf{n}(t)$ under radioactive decay and neutron-induced transmutation:

$$\frac{\mathrm{d}\mathbf{n}}{\mathrm{d}t} = \mathbf{A}(\mathbf{n}, t)\mathbf{n}, \tag{2.8}$$

where $\mathbf{A}$ is the burnup matrix. The off-diagonal elements of $\{A\}$ describe the production of a nuclide from other nuclides through radioactive decay, fission yields, and neutron-induced transmutation reactions. The diagonal elements are negative removal terms and include all processes that deplete the corresponding nuclide, such as radioactive decay, absorption, fission, and other neutron-induced reactions.

Because reaction rates depend on the neutron flux spectrum, which in turn depends on the evolving isotopic composition, the burnup matrix is composition-dependent, and therefore, also time-dependent. In a transport-coupled depletion calculation, OpenMC

Table 2.6: Burnup grid used in the reference continuous-energy depletion calculation.

Step	Burnup increment [MWd/kgHM]	Cumulative burnup [MWd/kgHM]
1	0.05	0.05
2	0.25	0.30
3	0.70	1.00
4	1.00	2.00
5	1.00	3.00
6	2.00	5.00
7	2.00	7.00
8	3.00	10.00
9	3.00	13.00
10	4.00	17.00
11	4.00	21.00
12	5.00	26.00
13	8.00	34.00
14	8.00	42.00
15	9.00	51.00
16	9.00	60.00
17	10.00	70.00
18	10.00	80.00
19	10.00	90.00
20	10.00	100.00

periodically performs transport solves to update the reaction rates entering the burnup matrix.

The coupling between transport and transmutation is handled by `CoupledOperator`. At each depletion stage, OpenMC performs a transport calculation to evaluate the neutron flux, the effective multiplication factor, and the microscopic reaction rates for the depletable materials. The set of nuclides, reactions, decay channels, and fission yields entering the burnup matrix is defined by a depletion chain file, in XML format.

2.3.2 CE/CM Predictor–Corrector Integrator

OpenMC provides several time-integration schemes for solving the depletion equations. In this work, the CE/CM integrator, implemented in `openmc.deplete.CECMIntegrator`, is used and is defined as [**OpenMCdocsDepletion**]

$$\mathbf{n}_{i+1/2} = \exp\left(\frac{h}{2}\mathbf{A}(\mathbf{n}_i, t_i)\right)\mathbf{n}_i, \tag{2.9}$$

$$\mathbf{n}_{i+1} = \exp\left(h\mathbf{A}(\mathbf{n}_{i+1/2}, t_{i+1/2})\right)\mathbf{n}_i, \tag{2.10}$$

where h is the time-step length, $\mathbf{n}_i$ is the composition at the beginning of the step, $\mathbf{n}_{i+1/2}$ is the predicted midpoint composition, and $\mathbf{n}_{i+1}$ is the end-of-step composition.

The algorithm consists of two stages.

1. Predictor stage: a first transport calculation is performed at the beginning of the step to obtain $\mathbf{A}(\mathbf{n}_i, t_i)$; this matrix is used to advance the composition to the midpoint.
2. Corrector stage: a second transport calculation is performed using the midpoint composition, and the resulting matrix $\mathbf{A}(\mathbf{n}_{i+1/2}, t_{i+1/2})$ is used to advance the original beginning-of-step composition over the full time step.

Compared with a predictor-only scheme, CE/CM requires an additional transport solve per step but provides a second-order treatment of the time integration.

2.4 Multigroup Cross-Section Generation

After the reference continuous-energy depletion calculation, a second set of OpenMC transport calculations is performed to generate the multigroup cross sections required for the later LFP production. Each burnup state is treated as an independent fixed-composition transport problem. The geometry, non-depletable materials, boundary conditions, and source settings are kept consistent with the reference depletion model.

2.4.1 Case Mapping and Burnup-State Reconstruction

For each burnup state, the fuel composition is extracted from the depletion results stored in the HDF5 output file. The `openmc.deplete.Results` object is used to reconstruct an `openmc.Material` at the selected burnup index using the `export_to_materials` method.

Nuclides not available in the cross-section library are excluded from the reconstructed material, since they cannot be used in the subsequent transport and tally calculations. Their total contribution to the fuel inventory is small and relegated in the very first step. The impact of their exclusion will be discussed in Sec 4.3.7.

This workflow allows each burnup state to be submitted as an independent SLURM array job. The script also accepts an optional list of fuel temperatures, so that different thermal states could be processed from the same burnup composition. In this work, however, all MGXS calculations are performed at the reference fuel temperature of 1500 K.

2.4.2 Statistical Precision

These calculations provide the reference dataset for the multigroup analysis and therefore require higher statistical precision than the depletion calculation. Each transport run employs 400 batches, of which 100 are inactive and 300 are active, with 10^6 particles per batch. This corresponds to 3×10^8 active histories per burnup state.

The increased number of histories is used to reduce the relative uncertainty of the tallied cross sections across the 33 energy groups. This is especially important for reactions and energy regions with low event probability, where insufficient statistics could distort the later comparison between explicit and lumped fission-product representations.

2.4.3 Energy-Group Structure

The MGXS are generated on the ECCO-33 energy-group structure, originally developed within the ECCO framework at CEA Cadarache and adopted as one of the standard group structures for fast reactor in the ERANOS deterministic transport system for fast-reactor applications [24]. The structure spans the energy range from 10^{-5} eV to approximately 20 MeV. The rationale for adopting this particular structure is threefold:

- the fast and epithermal energy ranges are resolved with finer detail compared to group structures designed for thermal reactors,
- the ECCO-33 structure has been extensively used in fast reactor benchmarks and in the standard nuclear data processing adopted by the Generation-IV fast reactor community. Collecting MGXS data on this structure therefore ensures direct compatibility with existing studies or other codes,
- Although finer group structures exist, the associated increase in computational cost is not justified for the level of accuracy required in this work.

The complete list of ECCO-33 group boundaries is reported in Table A.1 in Appendix A.1.

2.4.4 Tallied Reactions and MGXS Library

The minimal set multigroup cross sections requested from `openmc.mgxs.Library` are listed in Table 2.7.

Other secondary reaction are not included as their integral contribution to the neutron balance is expected to be negligible for the present application. Their contribution to removal or neutron multiplication is nevertheless implicitly accounted through aggregate quantities such as `absorption` and `nu-scatter matrix`.

The MGXS library is constructed with `openmc.mgxs.Library`. The ECCO-33 group structure is assigned through an `EnergyGroups` object. The domain is restricted to the fuel material, and microscopic cross sections are generated for each nuclide by setting `by_nuclide = True`. After the transport calculation, the library is stored in a `.pkl` file for later processing. The outputs are organised in a hierarchical directory structure of the form `./BU_{step}/T_{T}/`, so that each burnup–temperature combination is stored independently.

2.4.5 Tallies Definition

OpenMC estimates physical quantities through tallies. A reaction-rate tally for reaction x can be written as the phase-space inner product

$$\langle \Sigma_x, \psi \rangle = \int \mathrm{d}\mathbf{r} \int \mathrm{d}\mathbf{\Omega} \int \mathrm{d}E \; \Sigma_x(\mathbf{r}, E)\, \psi(\mathbf{r}, \mathbf{\Omega}, E), \tag{2.11}$$

where Σ_x is the macroscopic cross section and ψ is the angular neutron flux. The spatially homogenised microscopic multigroup cross section for reaction x in energy group g is then obtained as

Table 2.7: Multigroup quantities generated in the MGXS calculations.

Quantity	Role
`total`	Total cross section used for the evaluation of the collision probability.
`absorption`	Sum of all neutron-removing reactions.
`fission`	Fission cross section used to compute fission rates and power normalization.
`nu-fission`	Fission XS weighted by the average number of neutrons per fission $\bar{\nu}$.
`(n,gamma)`	Radiative-capture cross section.
`(n,2n)`	Neutron multiplication producing two outgoing neutrons.
`(n,3n)`	Neutron multiplication producing three outgoing neutrons.
`(n,p)`	Neutron-induced transmutation reaction producing a proton.
`(n,alpha)`	Neutron-induced transmutation reaction producing an α particle.
`chi`	Fission neutron emission spectrum.
`scatter matrix`	Group-to-group scattering transfer matrix.
`nu-scatter matrix`	Scattering transfer matrix weighted by neutron multiplicity, accounting for (n, xn) neutron-producing scattering reactions.

$$\sigma_{x,g} = \frac{\langle \sigma_x, \psi \rangle_g}{\langle \psi \rangle_g}. \tag{2.12}$$

Each tally is constructed by specifying three elements: one or more *filters*, one or more *scores*, and optionally a list of *nuclides*. Filters restrict the region of phase space being scored, for example by material, cell, or energy group. Scores define the quantity being accumulated, such as flux, fission, capture, absorption, or scattering. The nuclide specification allows the same score to be evaluated separately for individual isotopes.

OpenMC supports analog, collision, and tracklength estimators. The analog estimator scores events only when the corresponding physical reaction occurs. The collision estimator scores every collision weighted by $1/\Sigma_{\text{tot}}$. The tracklength estimator is generally preferred for volumetric flux and reaction-rate tallies because it provides lower variance than the analog estimator, especially for low-probability events, and it is the default choice in OpenMC whenever the scoring function does not require post-collision information [25].

Chapter 3

Lumped Fission Product Methodology and Implementation

3.1 Introduction and Methodological Choices

Starting from the multigroup cross-section libraries generated in the previous chapter, this chapter introduces the lumped fission product methodology developed and implemented in the present work. Before discussing the technical construction of the lumped cross sections and the corresponding weighting factors, two methodological choices are stated and justified that distinguish this work from most of the approaches reviewed in Section 1.3. The chapter then proceeds through three main stages: the generation of a reduced depletion chain with one lumped fission product associated with each selected heavy metal, the calculation of burnup-dependent effective yields, and the construction of yield-weighted multigroup cross sections in the OpenMC data formats required for depletion and transport.

3.1.1 Full Lumping of the Fission Product Inventory

As discussed in the historical review, a common practice, especially in thermal-reactor codes, is to adopt a hybrid scheme in which a set of selected important fission products are tracked explicitly, while the remaining isotopes are condensed into a residual lumped pseudo-nuclide. In the present work, a fully lumped approach is adopted instead: no individual fission product is treated explicitly. This choice is motivated by the reduced relative importance of individual fission products in fast spectra compared with thermal spectra, where strong absorbers such as xenon and samarium can dominate the reactivity behaviour.

Rather than constructing a single global lumped fission product representative of the average fuel mixture, a separate lumped fission product is generated for each selected heavy metal contributing to fission. Each parent nuclide is therefore coupled, in the modified

depletion chain, to its own dedicated pseudo-nuclide. This pseudo-nuclide represents the ensemble of fission products generated by fissions of that specific parent.

The rationale for this choice is the dependence of the fission product distribution on the fissioning nuclide. The mass-yield curves of the major actinides differ appreciably: heavier actinides tend to produce fission products shifted towards higher mass numbers, and the relative weight of the two yield peaks also changes. Consequently, two lumped fission products constructed from different fissile parents do need not have the same average capture, scattering, or threshold-reaction cross sections.

This distinction is particularly relevant in fast-spectrum systems, where the set of nuclides contributing to the total fission rate is broader than in a thermal-spectrum LWR. In thermal systems, fission is dominated mainly by ^{235}U and ^{239}Pu, with additional contributions from ^{238}U fast fission and ^{241}Pu. In a fast reactor, even-mass isotopes such as ^{238}Pu, ^{240}Pu may fission with non-negligible probability and also americium and curium produced through capture chains can contribute increasingly as irradiation proceeds. A more detailed analysis of the effect of HM mass on the LFP will be presented in the next sections. A schematic representation of the concept is illustrated in Fig. 3.1, where the top 5 most abundant nuclides are compared at the EOL for ^{235}U and ^{245}Cm.

The per-heavy-metal formulation is therefore expected to be more transferable across fuel compositions than a single lumped pseudo-nuclide generated for one specific isotopic vector. A change in the initial heavy-metal composition changes the relative production of the same set of parent-specific lumped fission products, rather than requiring a new global pseudo-nuclide fitted to that composition. This argument remains qualitative in the present work: the accuracy of the representation is expected to depend on the neutron spectrum and on the parent-wise fission rates, and a deeper assessment of transferability across different fuel compositions is left to future work.

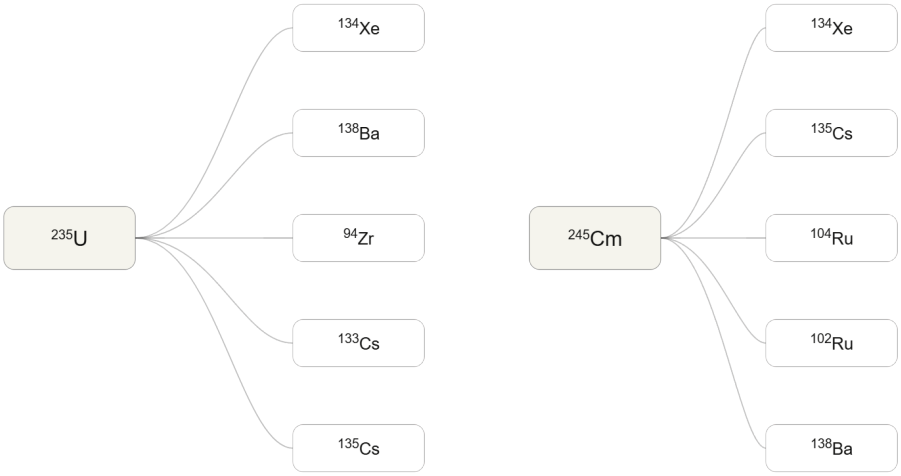

Figure 3.1: Top 5 most abundant fission products for ^{235}U and ^{245}Cm at EOL.

3.1.2 Reduced Heavy-Metal Set and Custom Chain Generation

The full set of actinides that may contribute to a fast reactor depletion calculation comprises several tens of nuclides. However, their individual contributions to the neutron balance and to the fission product source span many orders of magnitude. For the present implementation, the heavy-metal inventory is restricted to a reduced set of 23 nuclides, selected to retain the actinides that contribute appreciably to fission, capture, and the build-up of minor actinides during irradiation. The omission of the remaining nuclides is a modelling simplification whose impact on the depletion results is assessed in the validation stage of the work.

A further practical constraint is the availability of evaluated fission yield data. Some heavy metals that play only marginal roles as fissioning parents do not have their own fission yield evaluations in the nuclear data library used in this work. In ENDF/B-VIII.0, for example, U-239 does not have its own independent fission product yield entry and therefore inherits the yields of Pu-241 through an explicit redirection in the corresponding ENDF entry:

```
<neutron_fission_yields parent="Pu241"/>
```

To handle these aspects in a flexible and reproducible way, the construction of the modified depletion chain is fully automated through a dedicated Python script. The script takes as input a plain-text file specifying the 23 selected heavy metals and the decay or neutron-induced reactions retained for each nuclide. From this input, it generates a `custom_chain.xml` file that follows the same structure as the standard depletion chain files accepted by OpenMC.

The neutron-induced reactions retained in the modified chain are limited to those considered most relevant for the depletion of the heavy metals in the reference fast spectrum: radiative capture (n, γ), fission, and the $(n,2n)$ threshold reaction and other minor channels are neglected. The plain-text input format was adopted to simplify modification and reproducibility: extending the set of heavy metals or the list of reactions in future studies requires only an edit of the input file, without changes to the script r its underlying logic.

A portion of the input file is shown below, while the complete file is reported in Appendix A.2.

```
Pu239   fission     2.0     LFP_Pu239
Pu239   (n,2n)      1.0     Pu238
Pu239   (n,gamma)   1.0     Pu240
Pu239   alpha       1.0     U235
```

Decays with half-lives longer than 10^5 years are neglected, since their contribution to isotopic evolution over the considered irradiation interval is negligible for the purposes of the present model.

3.2 Fission Yields: From ENDF Data to Effective Yields

The construction of a lumped fission product requires, in addition to the microscopic cross sections discussed in the previous chapter, a set of weighting factors that assign the appropriate importance to the cross sections of the individual fission products generated by fission of a specific parent heavy metal.

3.2.1 Independent and Cumulative Yields in Evaluated Data

Evaluated nuclear data libraries such as ENDF/B, JEFF, and JENDL provide, for each fissile nuclide and for each incident neutron energy, two complementary sets of fission product yield data: *independent* yields and *cumulative* yields [26].

The independent yield γ_i^{ind} of a nuclide i is defined as the fraction of that nuclide produced directly in the fission event, after prompt neutron emission and before any subsequent delayed-neutron emission or radioactive decay of the fragments. However, the primary fragments are typically neutron-rich and far from the line of β-stability; most of them therefore undergo rapid β^- decay, redistributing mass towards more stable nuclides [27]. As a consequence, the independent yield distribution bears little resemblance to the fission product inventory. Using independent yields directly to weight the cross sections of a lumped pseudo-nuclide would overemphasise short-lived, highly unstable species from observations of irradiated fuel at practical reactor physics timescales.

The cumulative yield γ_i^{cum} addresses part of this limitation by approximating that these highly unstable fission products immediately undergo decay. It is defined as the total number of atoms of nuclide i produced per fission, including the contribution of all precursors that feed i through radioactive decay. Cumulative yields may be written recursively as

$$\gamma_i^{\text{cum}} = \gamma_i^{\text{ind}} + \sum_{j \in \text{precursors}(i)} b_{j \to i}\, \gamma_j^{\text{cum}}, \tag{3.1}$$

where the summation runs over all direct precursors j of nuclide i, and $b_{j\to i}$ is the branching ratio for the decay of j into i [28]. By construction, γ_i^{cum} is closer than γ_i^{ind} to the delayed post-fission population of daughter nuclides, because it accounts effectively for the decay flow towards long-lived or stable members of each chain.

Nevertheless, cumulative yields cannot be used directly as the weights of the constituent nuclides of a fully lumped fission product. The reason lies in the definition encoded by Eq. 3.1: γ_i^{cum} already includes the contribution of every precursor that decays into i. If cumulative yields are assigned simultaneously to all nuclides in the chain, precursor contributions are counted repeatedly.

The double-counting mechanism can be illustrated with a simple decay chain

$$A \xrightarrow{\beta^-} B.$$

Assume that the independent yields are 4% for nuclide A and 1% for nuclide B. The cumulative yield of B is therefore 5%, because it contains both the direct production of

B and the precursor contribution from the decay of A. If cumulative yields were used as weighting factors for both nuclides, the total weight assigned to this chain would become

$$4\% + 5\% = 9\%,$$

even though the actual contribution of the chain is only 5%. The precursor contribution is therefore counted twice: first explicitly through the separate inclusion of A, and then implicitly in the cumulative yield of B.

The double-counting problem has been recognised in the nuclear community since the earliest attempts at constructing pseudo-fission products. Within the IAEA-sponsored WIMS-D Library Update Project (WLUP), the AVRFPY code post-processes ENDF-6 fission yield and decay data to produce corrected cumulative yields. However, the AVRFPY methodology was designed for the hybrid scheme adopted by WIMS-D, in which selected strongly absorbing fission products are tracked explicitly and only the residual nuclides are condensed into a single pseudo-nuclide [18]. Although these residual yields are time-independent, the hybrid structure allows the depletion solver to evolve the explicitly treated fission products using their own decay constants and cross sections. This feature is particularly important for the build-up and saturation of strong absorbers during the first days of irradiation. The lumped residue, instead, is only required only to represent the slowly varying, lower-importance part of the inventory.

This approach cannot be directly transposed to the fully lumped framework adopted in this work, where no fission product is treated explicitly and the full fission product inventory of each heavy metal must be represented by the corresponding LFP.

To address both the yield-definition issue and the time-evolution issue within this fully lumped framework, a dedicated procedure is developed to generate a set of *effective yields* at each burnup step from controlled auxiliary depletion calculations. The construction of these yields, together with the limitations and trade-offs of the approach, is described in the following sections.

3.3 Production of Effective Yields per Heavy Metal

The effective yields used to construct each lumped fission product are generated by a dedicated Python script. For every fissile heavy metal in the reduced HM set, the script performs an auxiliary depletion calculation designed to isolate the fission product contribution of that single parent. The procedure relies on three ingredients introduced previously: the modified depletion chain produced from the input file (Appendix A.2), the multigroup cross sections generated in Sec. 2.4 and stored as OpenMC `MicroXS` objects, and the multigroup flux spectrum obtained from the reference depletion calculation.

3.3.1 Isolating the Fission Channel

The central methodological choice in the yield-production procedure is that, for each heavy metal considered, the depletion chain is modified so that fission is the only neutron-induced reaction allowed for that nuclide. All other reaction channels and radioactive

decay branches are suppressed for the heavy metal of interest. The remaining nuclides in the chain retain their reactions and decay modes.

At first, this choice may appear unnecessary. One might argue that the detailed transmutation history of each heavy-metal inventory is not central, since the final objective is just to obtain the fission fragments produced by fission. Under this view, a standard depletion calculation initialised with the heavy metal of interest as the only fissile nuclide would be sufficient.

However, this argument is not methodologically correct and would introduce a bias that undermines the rationale of the per-heavy-metal lumping scheme adopted in this work (Sec. 3.1.1). During irradiation in a fast spectrum, radiative capture on a fertile or fissile actinide can be comparable to, or larger than, its own fission rate. Each capture event produces a different actinide, which may become fissile.

As a concrete example, consider the breeding sequence

$$^{238}\mathrm{U} \xrightarrow{(n,\gamma)} {}^{239}\mathrm{U} \xrightarrow{\beta^-} {}^{239}\mathrm{Np} \xrightarrow{\beta^-} {}^{239}\mathrm{Pu},$$

which is the dominant production mechanism of ^{239}Pu in uranium-fuelled reactors and the physical basis of the breeder concept.

If a standard depletion calculation were carried out starting from a pure ^{238}U sample without modifying the chain, a substantial fraction of the accumulated fission product inventory could originate from fissions of ^{239}Pu produced through the capture and decay sequence. At sufficiently high burnup, this contribution could become dominant. The resulting yield distribution would no longer represent only the fission product distribution of ^{238}U, but a mixture of mainly ^{238}U and ^{239}Pu fission products whose relative weights depend on the irradiation history.

The correct auxiliary setup keeps the two flow paths separate. Each fissile parent contributes exclusively to its own lumped pseudo-nuclide, while the breeding chain in the full depletion model will still represented by the standard (n, γ) / decay sequence connecting the heavy metals:

$$^{238}\mathrm{U} \xrightarrow{(n,\gamma)} {}^{239}\mathrm{U} \xrightarrow{\beta^-} {}^{239}\mathrm{Np} \xrightarrow{\beta^-} {}^{239}\mathrm{Pu} \xrightarrow{\text{fission}} \mathrm{LFP\_Pu239},$$
$$^{238}\mathrm{U} \xrightarrow{\text{fission}} \mathrm{LFP\_U238}.$$

Thus, the only process the parent heavy metal is allowed to undergo in the auxiliary yield-generation chain is fission. Any other reaction would open a path towards a different actinide and, ultimately, towards the fission products of a different parent.

3.3.2 Implementation in the Python Script

The procedure is implemented in a Python script whose workflow can be summarised as follows.

Heavy-Metal Extraction. The script first reads the custom depletion chain and extracts the list of heavy metals for which a lumped fission product has been declared. This

is done by inspecting pseudo-nuclides of the form `LFP_X`, where `X` is the symbol of the parent heavy metal, and reconstructing the corresponding HM list.

For each heavy metal, the script generates a modified version of the original depletion chain in which only fission is retained for that heavy metal. All other `<reaction>` entries, such as radiative capture and $(n,2n)$, as well as all `<decay>` entries, are removed from the corresponding `<nuclide>` block. The concept is illustrated schematically in Fig. 3.2. The resulting chain file is written to a heavy-metal-specific subdirectory and used as input for the auxiliary depletion calculation. The remaining nuclides in the chain are left unchanged, so that the decay and transmutation behaviour of the generated fission products is preserved.

Step-by-Step Depletion with Pre-Computed Cross Sections. The yield-generation depletion is performed using the OpenMC `IndependentOperator`, which advances the composition in time without re-running the Monte Carlo transport solver. The burnup time steps are the same as in the reference depletion. For each step, the operator requires the following inputs:

- the material to deplete;
- the multigroup flux spectrum shape, read from the statepoint file of the reference depletion at the corresponding step and normalised as it is used to define the spectral;
- the source rate, hence the total flux necessary for normalisation, is extracted from the reference depletion simulation;
- a pre-computed multigroup `MicroXS` library, constructed from the multigroup cross-section tallies described in the previous chapter;
- the depletion chain modified for the selected heavy metal.

This method is appropriate for the yield-generation problem because the objective is not to solve a new self-consistent transport problem, but to evolve the fission product inventory generated by one parent heavy metal under a prescribed spectrum representative of the reference SFR pin-cell calculation.

The starting material is initialised with the heavy metal of interest as the only nuclide present, with unit mass fraction. After each burnup step, the script extracts the depleted material from the results file, copies the updated isotopic composition into a new `openmc.Material` object, and uses this material as the input for the next step.

At each burnup step, the densities of nuclides with mass numbers in the fission product range ($60 < A < 180$) and with an available nuclear data entry are extracted and normalised. The first criterion excludes light activation products and heavy actinides; the second excludes nuclides that, even if formally present in the depletion chain, cannot be transported in the downstream Monte Carlo simulations because they lack continuous-energy cross-section data. The total fission product density at step k is computed as

$$N_{\mathrm{tot}}^{\mathrm{FP}}(t_k) = \sum_{i \in \mathrm{FP}} N_i(t_k), \tag{3.2}$$

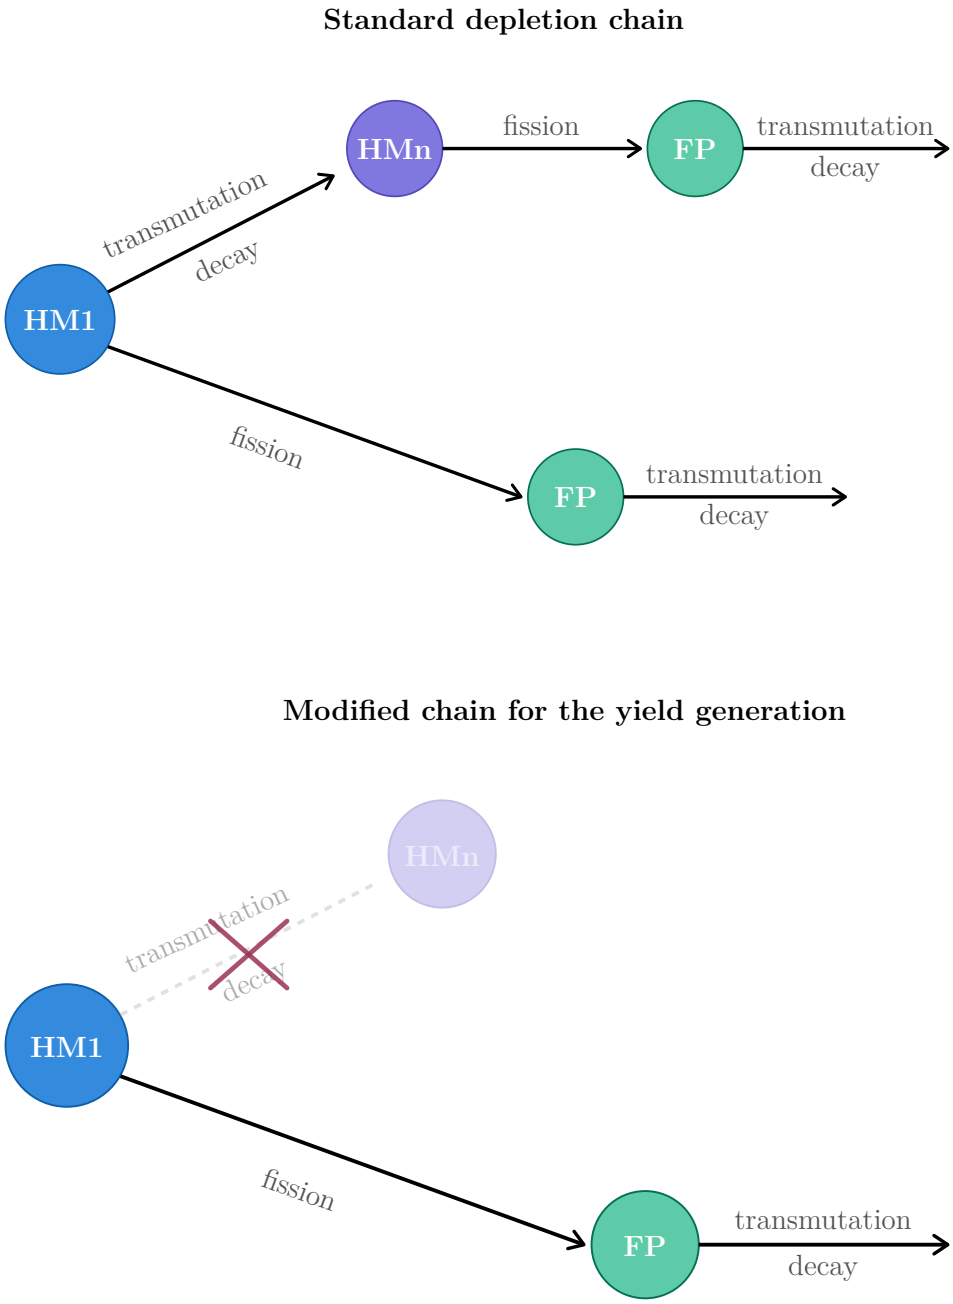

Figure 3.2: Comparison between the standard depletion chain and the modified chain used to generate effective yields. In the standard chain, the initial heavy metal HM_1 can transmute or decay into another heavy metal HM_n, which may then undergo fission and produce fission products. In the modified chain, heavy-metal transmutation and decay routes are disabled, so that fission is the only outgoing reaction of HM_1. Fission products continue to evolve through their own transmutation and decay reactions in both cases.

and the normalised effective yield of each fission product is then defined as

$$\gamma_i^{\mathrm{eff,HM}}(t_k) = \frac{N_i(t_k)}{N_{\mathrm{tot}}^{\mathrm{FP}}(t_k)}, \tag{3.3}$$

so that the effective yields at each burnup step sum to unity. These quantities should be interpreted as burnup-dependent inventory fractions, rather than as evaluated fission yields in the strict ENDF sense. By construction, they represent the relative composition of the fission product inventory generated by fission of a specific heavy metal under the prescribed spectrum and irradiation schedule. The resulting effective-yield table is stored as a dictionary indexed by burnup step and fission product nuclide, and is written to a JSON file in the corresponding heavy-metal subdirectory.

3.3.3 Analysis of the Computed Effective Yields

The effective yields for ^{239}Pu are visualised by plotting the time evolution of selected fission products grouped by half-life. When interpreting these plots, it is essential to keep in mind that the plotted quantities are fractions of the total fission product inventory at each burnup step.

A nuclide whose absolute concentration increases monotonically in time can still show a decreasing fraction if other species accumulate faster and occupy a larger share of the total inventory. This distinction between absolute abundance and inventory fraction is central to the interpretation of the trends discussed below.

The behaviour of the effective yields can be grouped into three categories according to the half-life of the nuclide. The three regimes are illustrated in Figs. 3.3, 3.4, and 3.5, which show representative fission products generated by fissions of ^{239}Pu in the reference fast spectrum.

Short-Lived Nuclides. Figure 3.3 shows the evolution of four nuclides with half-lives well below the duration of the overall irradiation history: ^{135}I ($t_{1/2} \approx 6.57$ h), ^{135}Xe ($t_{1/2} \approx 9.14$ h), ^{131}I ($t_{1/2} \approx 8.03$ d), and ^{140}Ba ($t_{1/2} \approx 12.75$ d).

All four nuclides exhibit the same qualitative pattern: rapid decrease to negligible fractions over the irradiation history.

The trend is consistent with the saturation behaviour of short-lived nuclides. After few half-lives the production rate of each species, either by direct fission yield or by decay of a precursor, is balanced by its removal rate due to decay and neutron-induced reactions. Once this equilibrium is reached, the absolute concentration of the nuclide becomes nearly stationary, while long-lived members of the inventory continue to accumulate. The fraction of the lump represented by short-lived species therefore decreases steadily towards zero. For these nuclides, removal is dominated by radioactive decay, because the corresponding neutron-induced reaction rates remain much smaller than the decay constants in the reference fast spectrum. The decline is therefore controlled primarily by the half-life of each species: ^{135}Xe disappears from the inventory fraction within a few days, while ^{140}Ba, with its longer half-life, retains a small but visible fraction for a longer time.

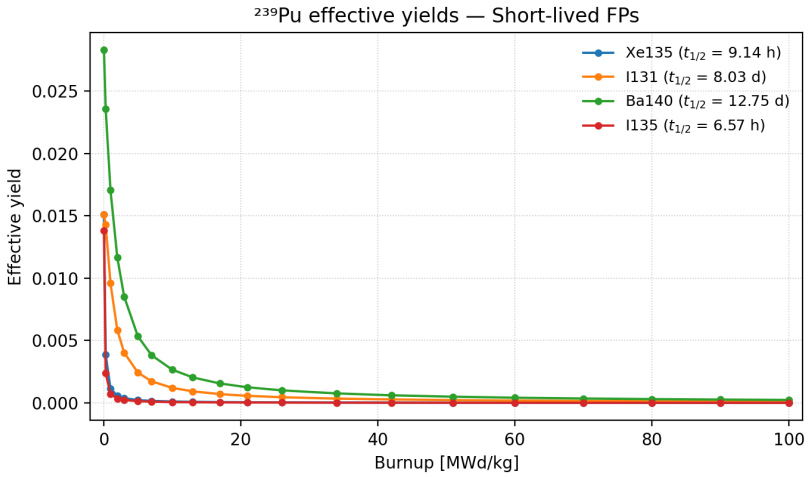

Figure 3.3: Time evolution of selected effective fission product yields associated with ^{239}Pu for short-lived nuclides.

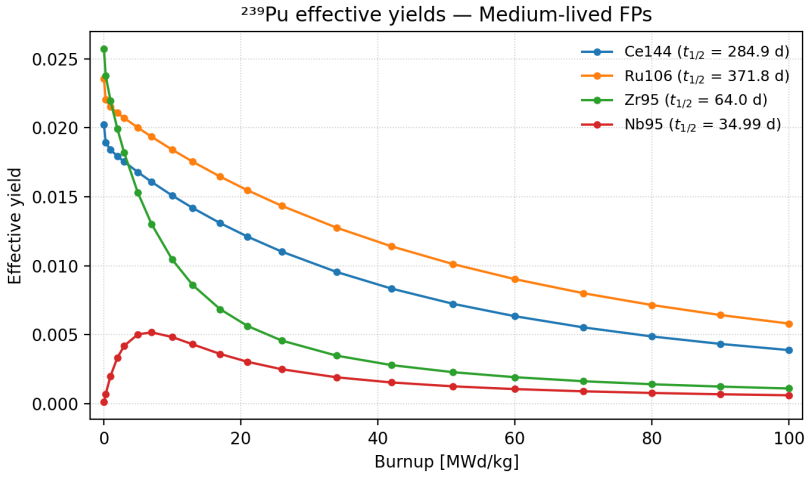

Figure 3.4: Time evolution of selected effective fission product yields associated with ^{239}Pu for medium-lived nuclides.

Medium-Lived Nuclides. Figure 3.4 shows the evolution of four fission products with intermediate half-lives, comparable to a significant fraction of the irradiation history considered in this work: ^{144}Ce ($t_{1/2} \approx 284.90$ d), ^{106}Ru ($t_{1/2} \approx 371.8$ d), ^{95}Zr ($t_{1/2} \approx 64.0$ d), and ^{95}Nb ($t_{1/2} \approx 34.99$ d). Their behaviour is intermediate between the short-lived and long-lived regimes: equilibrium between production and decay is approached on time scales of months to a few years, and the subsequent decline of the inventory fraction is correspondingly slower.

The pair ^{95}Zr $\rightarrow$ ^{95}Nb illustrates the effect of parent–daughter coupling captured by

the auxiliary depletion calculation. While ^{95}Zr is produced directly with a significant fission yield, ^{95}Nb has a much smaller independent yield and is populated mainly through the β^- decay of ^{95}Zr. As a result, its inventory fraction starts from nearly zero, grows with a delay imposed by the half-life of the parent, reaches a maximum, and then decreases as the parent–daughter chain approaches equilibrium. Its delayed rise and subsequent decrease therefore reflect the time-dependent evolution of the decay chain. This behaviour provides an additional motivation for deriving the LFP composition from depletion rather than prescribing it directly from independent or cumulative fission yields.

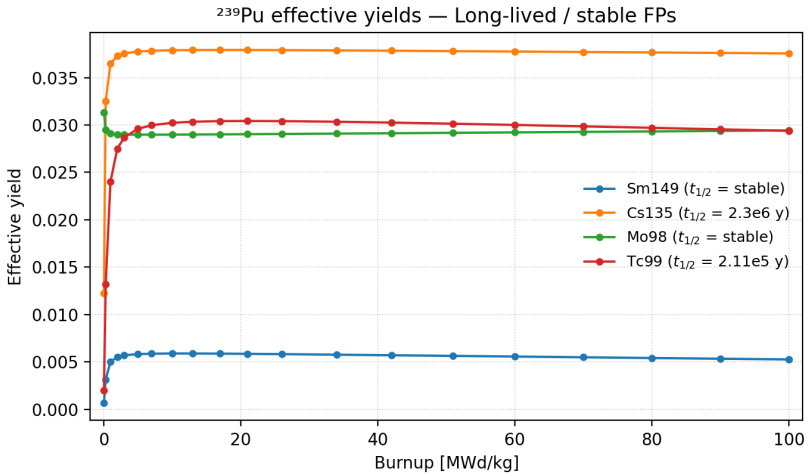

Figure 3.5: Time evolution of selected effective fission product yields associated with ^{239}Pu for long-lived or stable nuclides.

Long-Lived and Stable Nuclides. Figure 3.5 shows four nuclides whose half-lives are much longer than the irradiation time of the present calculation, or which are effectively stable: ^{135}Cs ($t_{1/2} \approx 2.3 \times 10^6$ yr), ^{99}Tc ($t_{1/2} \approx 2.1 \times 10^5$ yr), ^{98}Mo (stable), and ^{149}Sm (stable).

These species do not decay appreciably during the simulated irradiation period, and their absolute concentrations therefore grow monotonically with time, fed by direct fission production and by β^- decay of precursors. During the first burnup steps, the fission product inventory is still distributed among nuclides with different decay time scales, including short- and medium-lived. As these nuclides reach saturation or decay along their chains, the long-lived and stable end products continue to accumulate and become increasingly important contributors to the total fission product concentration. From that point onwards, their inventory fractions become approximately constant in time, because both numerator and denominator of Eq. 3.3 grow linearly with the cumulative number of fissions. The weak residual time dependence originates from medium-lived nuclides, which continue to redistribute mass between decay chains on time scales of months to years.

The three regimes can be summarised as follows. A fission product whose half-life is much shorter than the depletion time scale reaches equilibrium quickly and its inventory

fraction subsequently declines as longer-lived species accumulate. A fission product with a half-life comparable to the depletion time scale follows the same pattern on a slower time scale. A stable or quasi-stable nuclide does not equilibrate during the calculation; after the short-lived transients have disappeared, its share of the total inventory becomes nearly constant.

These observations motivate the use of burnup-dependent lumped fission product cross sections. The composition of the fission product inventory changes rapidly during the early irradiation steps, while at later burnup the lump is increasingly dominated by long-lived and stable nuclides. A time-dependent LFP representation can therefore capture the initial transient and the subsequent approach towards a more slowly varying effective composition.

3.4 Construction of the Lumped Fission Product Cross Sections

The multigroup microscopic cross sections of the individual fission products (Sec. 2.4) and the effective yields of each fissile heavy metal (Sec. 3.2) are combined in a single procedure to construct the lumped fission product cross sections used in the subsequent depletion and transport calculations.

3.4.1 Yield-Weighted Lumping of Microscopic Cross Sections

The procedure is implemented in a dedicated Python script. For each heavy metal HM and at each burnup step k, the script constructs a microscopic cross section $\sigma_{x,g}^{\mathrm{LFP\_HM},k}$ that represents, for reaction type x and energy group g, the average behaviour of the fission product inventory generated by fissions of that heavy metal. The averaging is performed using the normalised effective yields $\gamma_{i,k}^{\mathrm{HM}}$ of Eq. 3.3 as weights:

$$\sigma_{x,g}^{\mathrm{LFP\_HM},k} = \sum_{i \in \mathrm{FP}^{\mathrm{HM}}} \gamma_{i,k}^{\mathrm{HM}}\, \sigma_{x,g}^{i,k}, \tag{3.4}$$

where $\sigma_{x,g}^{i,k}$ is the multigroup microscopic cross section of fission product i in energy group g, for reaction x, drawn from the multigroup library produced at burnup step k. The same equation is applied separately to each reaction type retained in the library. For scalar additive quantities, such as total, absorption, fission, (n,γ), $(n,2n)$, ν-fission, and the fission spectrum χ, the lumping is performed by direct weighted averaging.

The scattering matrix $\sigma_s(g \to g')$ and the production matrix $\nu\sigma_s(g \to g')$ are also additive microscopic data. They are therefore lumped element by element, treating each (g, g') entry as a scalar quantity averaged with the same effective-yield weights.

The multiplicity matrix requires separate treatment. Although it appears in the cross-section list used in OpenMC multigroup cross-section generation examples [29, 30], it is not an independent microscopic cross section. In OpenMC, it is a dimensionless factor defined as the ratio between the ν-scatter matrix and the ordinary scatter matrix:

$$\sigma_{\text{mult}}(g \to g') = \frac{\nu\sigma_s(g \to g')}{\sigma_s(g \to g')}. \tag{3.5}$$

It represents the average number of outgoing neutrons associated with a scattering event from group g to group g', and is used to account for neutron multiplication in reactions (n, xn).

Since $\sigma_{\text{mult}}(g \to g')$ is defined as a ratio of two cross-section matrices, it cannot be lumped by applying the same linear weighting used for ordinary microscopic cross sections. In general,

$$\sum_i y_i \frac{\nu\sigma_s^i(g \to g')}{\sigma_s^i(g \to g')} \neq \frac{\sum_i y_i\, \nu\sigma_s^i(g \to g')}{\sum_i y_i\, \sigma_s^i(g \to g')}. \tag{3.6}$$

A direct weighted average of the individual multiplicity matrices would therefore not preserve the corresponding neutron production rate. For this reason, the scatter and ν-scatter matrices are first lumped separately as ordinary microscopic data. The multiplicity matrix of the lumped fission product is then reconstructed from these two lumped quantities:

$$\sigma_{\text{mult}}^{\text{LFP\_HM}}(g \to g') = \frac{\nu\sigma_s^{\text{LFP\_HM}}(g \to g')}{\sigma_s^{\text{LFP\_HM}}(g \to g')}, \tag{3.7}$$

which preserves the correct physical interpretation of the multiplicity factor for each incoming and outgoing energy-group pair.

The final result is a self-consistent set of multigroup microscopic cross sections for a fictitious nuclide, LFP_HM, whose neutronic behaviour represents the average absorption and scattering behaviour of the fission product inventory generated by fissions of parent heavy metal HM at the burnup step considered.

3.4.2 Two Output Libraries for the OpenMC Operators

The yield-weighted lumping operation generates two distinct cross-section libraries, saved separately according to the format and content required by the OpenMC operator that will use them.

1. The first library is structured as a collection of `openmc.deplete.MicroXS` objects, one per burnup step. It contains the multigroup microscopic cross sections of every depletable nuclide in the fuel according to the customised depletion chain. Since the chain is restricted to fission, (n, γ), and $(n,2n)$, only these reactions are required for the Bateman equations. This `MicroXS` library is the format expected by the `IndependentOperator` and is used to compute the reaction rates that drive the time evolution of the composition during depletion.

2. The second library is structured as an `openmc.MGXSLibrary` object, which acts as a container of nuclide-specific `openmc.XSdata` entries. Each `XSdata` entry contains the full set of multigroup data required by the transport solver: total, absorption, fission,

ν-fission, fission spectrum χ, scattering matrix, ν-scatter matrix, and multiplicity matrix. The complete `MGXSLibrary` is then exported in the HDF5 format used by the OpenMC multigroup transport solver as its cross-section input file.

The separation of the two libraries reflects the fact that the depletion and transport operators consume cross sections in different formats and with different requirements. The depletion solver only needs the cross sections of the reactions explicitly appearing in the chain, while the transport solver requires the complete set of data needed to perform an eigenvalue calculation. In addition, the transport library must contain the cross sections of the non-depletable materials of the model: cladding, gap, and coolant. These materials are assumed not to change during irradiation, and their multigroup cross sections are therefore generated only once, at their nominal operating temperature of 673 K.

The complete output of the procedure is a hierarchical library indexed by burnup value, in which every irradiation point of the reference depletion is associated with a pair of libraries: one for depletion and one for transport. The fresh-fuel state (BU = 0) is treated as a special case: only the transport library is generated, because the construction of the `MicroXS` depletion library requires effective-yield weights, which are not defined at fresh-fuel conditions.

Indexing the output directories by cumulative burnup, in MWd/kgHM, rather than by burnup-step index makes the library easier to reuse in later calculations and comparisons.

Chapter 4

Application and Assessment of the Lumped Fission Product Methodology

4.1 Application Framework

The previous chapters described the production of multigroup cross sections, the generation of effective yields associated with each fissioning heavy metal, and the yield-weighted lumping procedure used to construct a self-consistent set of `LFP_HM` pseudo-nuclides. At this point, the methodological infrastructure of the work is complete: a multigroup depletion library, a multigroup transport library, and a custom depletion chain in which each fissioning heavy metal is coupled to its own lumped fission product. The present chapter marks the transition from the construction of these objects to their application and assessment against a continuous-energy reference.

OpenMC does not currently provide a fully coupled multigroup transport-depletion solver: the `CoupledOperator` described in Sec. 2.3.1 is designed for continuous-energy transport coupled to depletion, whereas the `IndependentOperator` performs depletion on user-supplied multigroup data without invoking a transport solve. A self-consistent multigroup transport-depletion loop therefore requires a custom workflow in which the transport and depletion stages are orchestrated explicitly by the user. The construction of this workflow, built progressively in the following sections, is one of the methodological contributions of this work.

4.1.1 A Stepwise Approach

Throughout the chapter, the central quantity that links the continuous-energy reference and the multigroup calculations is the macroscopic reaction rate. In a continuous-energy treatment, the reaction rate of nuclide i undergoing reaction x in a given material is

$$R_x^i \;=\; N_i \int_0^\infty \sigma_x^i(E)\,\phi(E)\,dE, \tag{4.1}$$

where N_i is the atom number density of nuclide i, $\sigma_x^i(E)$ is its energy-dependent microscopic cross section, and $\phi(E)$ is the angle-integrated neutron flux spectrum. In a multigroup treatment, the same quantity is expressed as a discrete sum over G energy groups,

$$R_x^i \;=\; N_i \sum_{g=1}^{G} \sigma_{x,g}^i\,\phi_g, \tag{4.2}$$

where $\sigma_{x,g}^i$ is the multigroup microscopic cross section in group g and ϕ_g is the corresponding group-integrated flux. Equation 4.2 identifies the three ingredients that must be reproduced consistently in the multigroup calculation: the nuclide densities, the group-wise microscopic cross sections, and the group fluxes.

The preservation of R_x^i in the transition from CE to MG is therefore the natural metric used to assess the consistency of the calculations in this work. Rather than directly running the fully coupled multigroup transport-depletion calculation, the assessment is structured in three sequential steps of increasing complexity. By introducing the multigroup features one at a time, while keeping the others anchored to the CE reference, the main sources of discrepancy can be investigated more clearly.

Step 1 — Cross-Section Analysis (Decoupled). In the first step, both the isotopic concentrations and the multigroup flux are fixed to those of the CE reference calculation at each burnup point. The multigroup reaction rates are then computed in post-processing through Eq. 4.2, using the multigroup cross sections of the library constructed in Sec 3.4, and compared to the corresponding CE reaction rates. In this configuration, the only multigroup ingredient is the cross-section data themselves; the spectrum that weights the reaction rate is identical in the CE and MG calculations. The step therefore provides a clean setting to analyse the lumped fission product cross sections: their composition, their dependence on the parent heavy metal, and their evolution with burnup.

Step 2 — Multigroup Depletion (Partially Coupled). In the second step, the concentrations are calculated by depletion integration using the multigroup cross-section library, while the multigroup flux supplied to the depletion operator is still taken from the CE reference calculation at each burnup point. This configuration isolates the effect of using multigroup data and the reduced LFP chain to model the time evolution of the composition.

Step 3 — Fully Coupled Multigroup Calculation. In the third and final step, the calculation is performed entirely in multigroup mode: the transport solve is performed using the multigroup transport library, and the resulting multigroup flux is then used to drive the depletion calculation.

4.1.2 Sources of Approximation

The stepwise structure outlined above is meant to expose the contribution of each modelling choice as cleanly as possible. It is nevertheless useful to acknowledge from the outset that several distinct sources of approximation enter the multigroup LFP framework, and that the discrepancies observed in any of the three steps are in general the combined result of more than one contribution.

The first source of approximation is the multigroup representation of the cross sections. Even with a relatively fine 33-group structure such as ECCO-33, collapsing a continuous-energy resonance structure onto a discrete energy mesh involves inevitably a loss of spectral information. Its accuracy depends on the agreement between the weighting flux used when collapsing and the spectrum in which the cross sections are later employed. The second source is the statistical uncertainty of cross sections, which are obtained as post-processed Monte Carlo tallies in the reference calculation. Despite the high statistics adopted in the MGXS generation runs, the resulting cross sections retain a finite stochastic uncertainty that propagates to the multigroup reaction rates. The third source is the spectrum-specificity of the multigroup cross sections: they are generated in the spectrum of one specific pin-cell configuration at a given burnup state, and their use in a different spectral environment introduces a residual inconsistency.

Additional approximations are specific to the lumped fission product methodology. The effective yields used to construct the pseudo-nuclides are derived from an auxiliary depletion calculation and therefore reflect the inventory composition of that auxiliary run rather than that of the actual coupled simulation. Moreover, the lumping operation replaces a large set of explicit fission products with a small number of pseudo-nuclides and necessarily removes isotope-specific information. Any spectral feature that depends non-linearly on the individual cross sections of the fission products, such as the shielding of strong resonances or the saturation of dominant absorbers, is necessarily lost in the average. These limitations are the unavoidable price of replacing a detailed fission product inventory with a reduced set of lumped carriers.

4.2 Step 1 — Cross-Section Analysis with CE Flux and Concentrations

The first stage of the assessment isolates the multigroup cross-section. The explicit FP concentrations N_i and the multigroup flux ϕ_g are taken from the CE reference depletion and from the corresponding tally, projected onto the ECCO-33 mesh. The LFP carrier concentrations are then reconstructed from the total CE fission product inventory using heavy-metal fission-rate fractions evaluated at each burnup step. With these inputs fixed by the CE reference, the only genuinely multigroup quantity is the cross-section library itself.

4.2.1 Reaction Rate Comparison

Three reaction channels are compared: radiative capture (n,γ), the threshold reaction $(n,2n)$, and the total reaction rate. The comparison is carried out by computing

$$R_x^{(\mathrm{CE})} = \sum_{i\in\mathrm{FP}} N_i^{(\mathrm{CE})} \sum_g \sigma_{x,g}^{i,(\mathrm{CE})}\,\phi_g^{(\mathrm{CE})} \qquad \text{vs.} \qquad R_x^{(\mathrm{LFP})} = \sum_{\ell\in\mathrm{LFP}} N_\ell^{(\mathrm{rec})} \sum_g \sigma_{x,g}^{\ell,(\mathrm{LFP})}\,\phi_g^{(\mathrm{CE})}, \tag{4.3}$$

where $N_\ell^{(\mathrm{rec})}$ denotes the reconstructed concentration of the lumped pseudo-nuclide ℓ in this decoupled analysis.

A methodological point must be clarified before the results are presented. The lumped fission products of this work are constructed per heavy metal, so that each fissioning parent has its own pseudo-nuclide, denoted here as `LFP_HM`. Therefore, to rigorously compute $R_x^{(\mathrm{LFP})}$, individual concentrations $N_{\mathrm{LFP\_HM}}$ are needed at each burnup step . Clearly, these concentrations cannot be obtained from the CE reference.

To circumvent this limitation, the total CE fission product concentration is partitioned among the lumped pseudo-nuclides at each burnup step according to the instantaneous fission rate fractions of the corresponding heavy metals. At step k, the fractional contribution of each heavy metal is

$$f_{\mathrm{HM}}(k) = \frac{R_{f,\mathrm{HM}}^{(CE)}(k)}{\sum_{j\in HM} R_{f,j}(k)}, \tag{4.4}$$

where $R_{f,\mathrm{HM}}^{(CE)}(k)$ is the CE fission reaction rate of heavy metal HM at step k, and the sum runs over all heavy metals for which a lumped pseudo-nuclide is available in the library. The reconstructed concentration of each LFP_HM at step k is then

$$N_{\mathrm{LFP_{HM}}}^{(\mathrm{rec})}(k) = f_{\mathrm{HM}}(k)\, N_{\mathrm{FP,tot}}^{(\mathrm{CE})}(k), \tag{4.5}$$

with

$$N_{\mathrm{FP,tot}}^{(\mathrm{CE})}(k) = \sum_{i\in\mathrm{FP}} N_i^{(\mathrm{CE})}(k). \tag{4.6}$$

Thus, the Step 1 comparison does not use LFP concentrations produced by a depletion calculation with lumped fission products. Rather, it constructs an equivalent LFP inventory from the total CE fission product inventory, weighted by the instantaneous fission contribution of each heavy metal at the same burnup step. This remains an approximation, because the partition is based on fission rates at step k rather than on the cumulative fission production integrated over the entire irradiation history. Nevertheless, it provides a more consistent reconstruction than fixed average heavy-metal fractions and is sufficient for the purpose of isolating the effect of the LFP multigroup cross sections.

The results are shown in Fig. 4.1, which reports both the absolute reaction rate as a function of irradiation time and the relative deviation of the LFP estimate with respect to the FP reference.

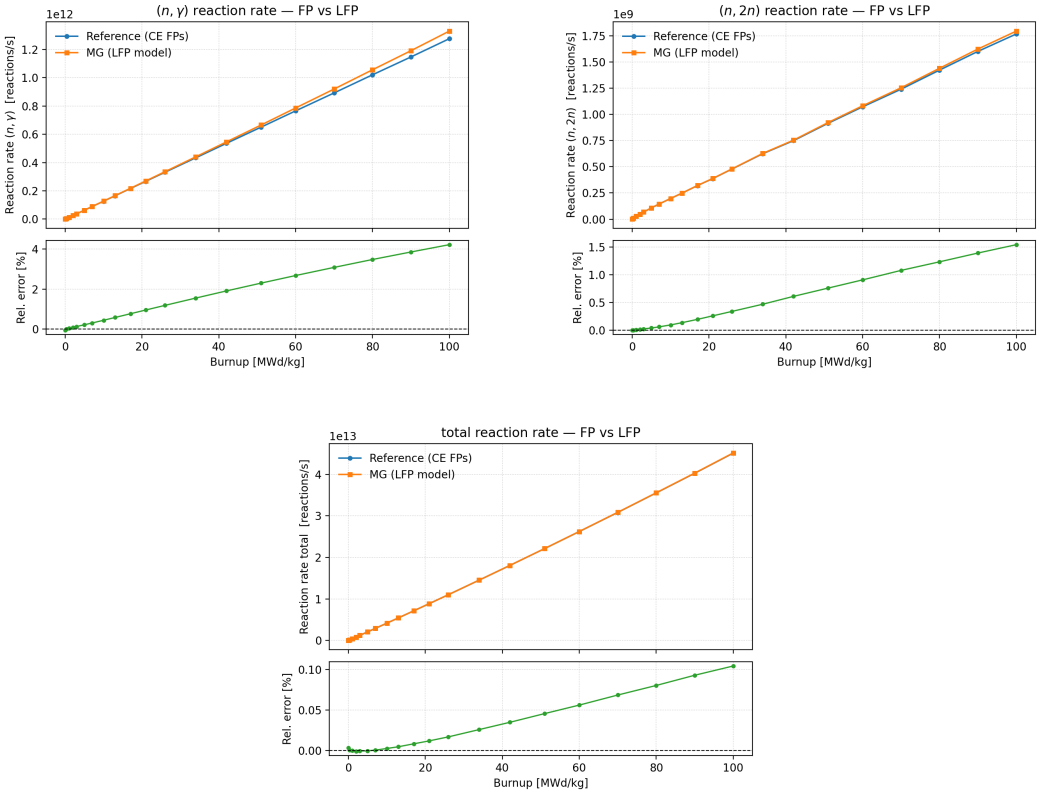

Figure 4.1: Comparison between the explicit FP and reconstructed LFP reaction rates as a function of irradiation burnup. The upper row shows the (n, γ) reaction rate on the left and the $(n,2n)$ reaction rate on the right; the lower centred panel shows the total reaction rate. In each panel, the upper subplot reports the absolute reaction rate, while the lower subplot reports the relative deviation of the LFP estimate from the FP reference.

The radiative-capture reaction rate is overestimated by the LFP representation throughout the irradiation, with a discrepancy increasing from approximately 0.8% in the early steps to about 4.2% at end of life. The $(n,2n)$ channel shows a similar but smaller bias, increasing from about -0.1% to about 1.5%. The total reaction rate, in absolute value significantly larger than the (n, γ) rate, reflects the dominance of scattering over absorption in fission product cross sections in a fast spectrum, and shows a smaller deviation, contained within $\pm 0.1\,\%$ over the entire irradiation history.

The absorption-channel deviations arise from several effects that cannot be disentangled at this level: the effective yields generated in an auxiliary depletion that does not exactly reproduce the inventory of the reference fuel, the instantaneous fission-rate-based reconstruction of Eq. 4.5, and the residual statistical and spectrum-collapsing uncertainties of the multigroup cross sections. The main conclusion of this comparison is therefore the order of magnitude of the bias, namely a few percent for the dominant absorption channel.

4.2.2 One-Group Cross Section as a Function of Burnup

To complement the reaction-rate comparison, the flux-weighted one-group absorption cross sections of the lumped pseudo-nuclides `LFP_U235`, `LFP_Pu239` and `LFP_Cm245` are plotted as a function of irradiation time in Fig. 4.2.

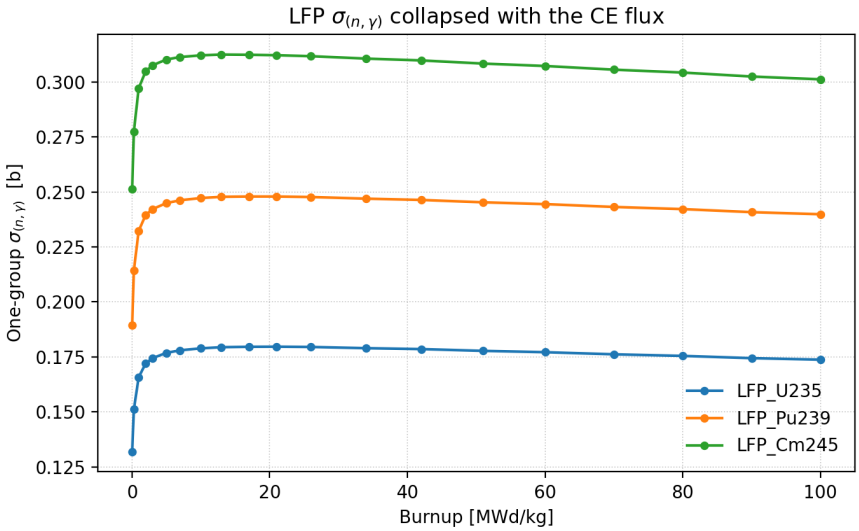

Figure 4.2: Flux-weighted one-group (n, γ) cross section of selected LFP pseudo-nuclides as a function of burnup, obtained using the fixed mid-cycle reference spectrum.

These three nuclides are selected as representative fission parents spanning an increasing mass range. For each LFP, the one-group cross section is obtained by flux-weighting the 33-group cross sections with the reference spectrum. To isolate the evolution of the LFP composition from changes in the neutron spectrum, the flux is fixed to its mid-cycle value in this comparison.

Two features are relevant.

- The one-group cross section increases rapidly during the first $\sim 5\,\mathrm{MWd/kgHM}$, with relative variations between about 25% and 35% depending on the fissioning parent. This early transient reflects the build-up of the fission product inventory and the progressive enrichment in the lump of long-lived or stable absorbers, such as the lanthanides. Once the dominant decay chains approach their asymptotic composition, the rate of increase decreases and the one-group cross sections tend toward a plateau. This confirms that a burnup-independent LFP cross section would introduce a systematic error throughout the irradiation cycle.

- The absolute value of the one-group cross section depends strongly on the parent heavy metal: $\mathrm{LFP}_{\mathrm{Cm245}}$ is approximately 1.7 times as absorbing as $\mathrm{LFP}_{\mathrm{U235}}$. This supports the choice of a parent-specific lumping scheme.

4.2.3 Yield and One-Group Cross Section as a Function of Mass Number

To interpret the dependence of the lumped cross section on the parent heavy metal, the same one-group (n, γ) cross section is represented in Fig. 4.3 as a function of the mass number A of the contributing fission products. The figure also reports the fission-yield distributions of ^{235}U, ^{239}Pu, and ^{245}Cm, produced in the previous chapter. Since more than one nuclide may contribute to the same mass chain, the cross-section value associated with each A is obtained by aggregating the corresponding one-group cross sections using the ^{239}Pu yields as relative weights. This choice is adopted only as a graphical convention to obtain a single representative value for each mass number and does not enter the quantitative construction of the LFP libraries.

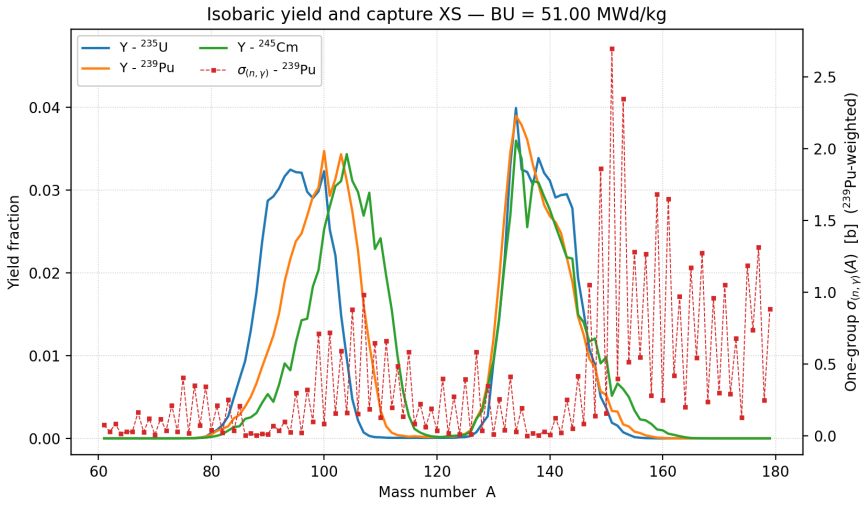

Figure 4.3: Fission-yield distributions of selected parent nuclides and representative one-group (n, γ) cross sections aggregated by mass number.

Two features of the figure are particularly relevant. The first concerns the fission-yield distributions. All three parents exhibit the characteristic asymmetric double-peaked structure of actinide fission, with a light-fragment peak and a heavy-fragment peak separated by a valley around the near-symmetric mass split. As the mass of the fissioning parent increases from ^{235}U to ^{239}Pu and ^{245}Cm, the light-fragment peak shifts progressively towards larger mass numbers. The heavy-fragment peak remains comparatively more stable, approximately in the region $A \sim 133$–140. This behaviour is consistent with the stabilising role of shell effects in the heavy-fragment region, commonly associated with the $Z = 50$ and $N = 82$ shell closures. The additional mass of heavier fissioning systems is therefore accommodated mainly through the displacement of the light-fragment peak and changes in the contribution of more symmetric splits, rather than through a comparable displacement of the heavy-fragment peak.

The second feature concerns the structure of the one-group (n, γ) cross section as a

function of A. This curve should not be interpreted as a fundamental dependence of the capture cross section on mass number alone, since the actual cross sections are isotope-specific and can be significantly different for diverse isobars. Nevertheless, it provides a qualitative indication of which mass regions are associated with larger capture strength in the considered fast spectrum. Two regions stand out: a moderately absorbing band in the light peak around $A \approx 100$–110, and a region of stronger absorption beyond $A \approx 145$ that includes the well-known fission product absorbers Sm, Eu, and Gd. On the other hand, nuclides at or near closed neutron shells tend to have lower neutron-capture cross sections, since the addition of one more neutron requires populating a higher-lying shell. A residual even–odd staggering is also visible in the cross-section curve, with even-A nuclides generally showing lower values than their odd-A neighbours, consistent with the basic motivation of the odd–even lump model proposed by Atefi and discussed in Sec. 1.3.

Together, these two features explain qualitatively why the one-group cross section of `LFP_Cm245` exceeds that of `LFP_U235` in Fig. 4.2. The shift of the light peak with increasing fissile mass, combined with the broadening of the distribution towards heavier mass numbers, moves a larger fraction of the fission product inventory into the high-absorption region beyond $A \approx 145$. Since the lumped cross section is a yield-weighted average of the individual contributions, it increases accordingly.

This first step is therefore not intended as a final validation of the depletion workflow. Rather, it identifies the magnitude and physical origin of the cross-section effects that are later tested in depletion.

4.3 Step 2 — Multigroup Depletion with CE-Supplied Flux

In the second step, the multigroup representation is extended to the depletion stage of the calculation. Both the cross sections $\sigma^i_{x,g}$ and the isotopic concentrations N^i are now evaluated in the multigroup framework, while the flux ϕ_g is still inherited from the continuous-energy reference. The transport stage is replaced by a prescribed external source that reproduces, on the ECCO-33 mesh, the flux conditions of the CE reference at each burnup step. Any deviation of the resulting reaction rates from the CE values can therefore be attributed primarily to the multigroup cross-section data, the depletion-chain representation, and the numerical integration of the Bateman equations, rather than to differences in the neutron spectrum.

4.3.1 Numerical Workflow

The implementation is built around OpenMC's `IndependentOperator`, described in Sec. 3.3.2. At each burnup step k, the script reads the multigroup flux tallied in the CE reference, the corresponding pre-computed `MicroXS` library, and the depleted material composition produced at the previous step. These three objects, together with the custom depletion chain, are passed to the operator. The integration is then advanced over one burnup interval through a `PredictorIntegrator`, a first-order scheme that performs a single matrix-exponential evaluation per step. At the end of the step, the depleted material

is read from the results file and used as the input composition of the next step. The procedure is repeated until the end of the irradiation history is reached.

4.3.2 Source-Rate Normalisation

A specific aspect of the workflow that deserves a dedicated discussion is the procedure adopted to normalise the flux used during depletion. An eigenvalue transport calculation provides reaction-rate and flux tallies per source particle; their absolute scale must therefore be fixed through an external normalisation, usually linked to the prescribed thermal power. This procedure is non-trivial because the power does not originate exclusively from fission: a non-negligible fraction of the deposited energy may also arise from processes such as radioactive decay and (n, γ) reactions.

The OpenMC flux score corresponds to the track-length estimator integrated over the tally region and therefore carries units of n cm/source. The corresponding physical flux ϕ is recovered through a source-rate normalisation factor f:

$$\phi = \frac{f\phi'}{V}, \tag{4.7}$$

where

$$\phi \text{ [n/(cm}^2\,\text{s)]}, \qquad \phi' \text{ [n cm/source]}, \qquad f \text{ [source/s]}, \qquad V \text{ [cm}^3\text{]}.$$

Three conventional strategies are commonly adopted in reactor physics codes to compute the source-rate normalisation factor, each based on a different integral quantity linking local reaction rates to global energy production:

1. The first strategy normalises the flux against the total energy deposition per source particle, H, evaluated through a dedicated heating tally. If the user prescribes a thermal power level P, the normalisation factor is

$$f = \frac{P}{H \times 1.602 \times 10^{-19}}, \tag{4.8}$$

where

$$P \text{ [J/s]}, \qquad H \text{ [eV/source]}.$$

The factor 1.602×10^{-19} converts eV into J. This strategy is physically complete, since it accounts for all the energy-depositing reactions included in the heating score. Its main drawback is the additional computational effort required to evaluate the heating tally with sufficiently low statistical uncertainty.

2. The second strategy relates the prescribed power directly to fission heating:

$$f = \frac{P\bar{\nu}}{Qk}, \tag{4.9}$$

where

$$P \text{ [J/s]}, \qquad \bar{\nu} \text{ [neutron/fission]}, \qquad Q \text{ [J/fission]}, \qquad k \text{ [neutron/source]}.$$

Here $\bar{\nu}$ is obtained as the ratio between the `nu-fission` and `fission` tallies, both normalised per source neutron. The quantity Q is the average recoverable energy per fission event and can be tallied through the `kappa-fission` score, and k is the average number of neutrons produced per source particle.

3. The third strategy normalises the flux through the ratio between two fission-rate quantities:

$$f = \frac{R_{\text{fission}}}{R'_{\text{fission}}}, \tag{4.10}$$

where

$$R_{\text{fission}} \text{ [fission/s]}, \qquad R'_{\text{fission}} \text{ [fission/source]}.$$

The quantity R'_{fission} corresponds to the `fission` tally obtained from the eigenvalue calculation, while R_{fission} is the total absolute fission rate of the system, obtained from a depletion calculation by summing the fission reaction rates of all fissile nuclides returned by `openmc.deplete.Results.get_reaction_rate`.

The three strategies are methodologically valid. The third strategy differs from the other two in that it relies on a preceding depletion calculation. It exploits the fact that the fission reaction rates returned by OpenMC are already scaled according to the normalisation mode used in the depletion operator. The resulting source rate therefore reproduces the absolute flux scale used internally by the CE reference calculation. This strategy is therefore adopted in the present work, as it reproduces exactly the absolute flux used in the simulation by OpenMC itself.

The normalisation of each depletion step is specified through the `normalization_mode` argument of the OpenMC depletion operator, which determines how the reaction rates entering the Bateman equations are scaled. The reference CE depletion of Sec. 2.2 adopted the `fission-q` mode, in which OpenMC computes the total energy production from the fission Q-values stored in the depletion chain. Only the multigroup depletion of this step uses the `source-rate` mode, in which the source rate f is supplied externally at each burnup step to reproduce the absolute flux of the CE reference [31, 32].

4.3.3 OpenMC Source-Code Patch for LFP Names

The use of the `LFP_HM` pseudo-nuclides inside an OpenMC depletion calculation requires a small patch to the Python source code. The `LFP_HM` identifiers do not respect the Generalized Nuclear Database Structure (GNDS) naming convention adopted by OpenMC for real nuclides and therefore trigger a validation error in the function `openmc.data.data.zam()`. This function parses nuclide names into a tuple (Z, A, m) containing the atomic number, mass number, and metastable-state index. The function then looks up the extracted chemical symbol in the table of known elements; both checks fail for names of the form `LFP_X`.

The patch consists of a preliminary check at the entry of the function:

```
def zam(name):
    if name.startswith('LFP'):
```

```
    return (0, 0, 0)
symbol, A, state = _GNDS_NAME_RE.fullmatch(name).groups()
metastable = int(state[2:]) if state else 0
return (ATOMIC_NUMBER[symbol], int(A), metastable)
```

For names beginning with `LFP`, the function returns the placeholder tuple (0,0,0), bypassing both validation steps. For all other nuclides, the original behaviour is preserved.

The modification is functionally isolated to the Python side of OpenMC and does not affect the neutronic calculation performed by the underlying C++ solver. Three independent verifications were carried out to confirm this property.

1. The routine `Material.to_xml_element`, which writes the input file `materials.xml` consumed by the solver, treats nuclide names as literal strings and does not invoke `zam()`. The `LFP_HM` identifiers are therefore transferred to the solver unchanged.

2. The matching of nuclides to their cross-section data is performed by the C++ solver as a string comparison against the `cross_sections.xml` index. The cross sections loaded for the lumped pseudo-nuclides are therefore those of the dedicated HDF5 files produced in Sec. 3.4, despite the placeholder tuple (0,0,0) formally coinciding with that of the neutron.

3. The materials of the model are defined with `density_units="sum"` and nuclide composition is directly defined in terms of atom/b-cm. In this configuration, `Material.get_nuclide_atom_d` does not activate its internal mass-to-atom conversions, which would otherwise require `openmc.data.atomic_mass()` and raise a `KeyError` for the lumped pseudo-nuclides, absent from the atomic-mass table.

4.3.4 Comparison of Reaction Rates

Following the same logic of Step 1 (Section 4.2.3), the two reaction rates compared in this stage are

$$R_x^{(\mathrm{CE})} = \sum_{i \in \mathrm{FP}} N_i^{(CE)} \sum_g \sigma_{x,g}^{i,(\mathrm{CE})}\, \phi_g^{(\mathrm{CE})} \quad \text{vs.} \quad R_x^{(\mathrm{LFP})} = \sum_{i \in \mathrm{LFP}} N_i^{(\mathrm{MG})} \sum_g \sigma_{x,g}^{i,(\mathrm{LFP})}\, \phi_g^{(\mathrm{CE})}, \tag{4.11}$$

with the crucial difference that the lumped pseudo-nuclide concentrations $N_i^{(\mathrm{MG})}$ are now obtained from the multigroup depletion. The comparison is therefore more self-consistent than the previous one, although residual discrepancies can still originate from the multigroup depletion integration, the reduced LFP chain, and the underlying cross-section data.

4.3.5 Inventory and Fission Rate

Figure 4.4 reports the evolution of the total heavy metal concentration in the two calculations. The curves are visually indistinguishable, and the relative deviation remains below

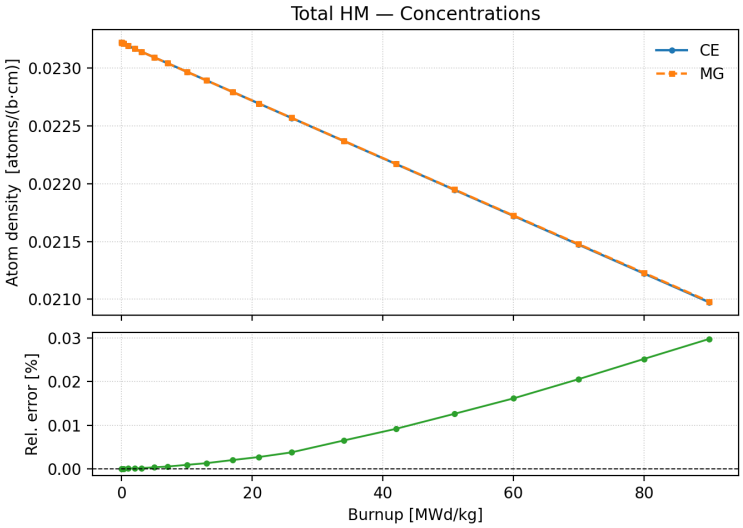

Figure 4.4: Total heavy-metal concentration in the CE reference and MG depletion calculations, with relative deviation.

0.03 % over the entire irradiation, growing smoothly with burnup. The same agreement is found for the total fission product concentration of Fig. 4.5, which stays within a few tenths of a percent throughout the cycle.

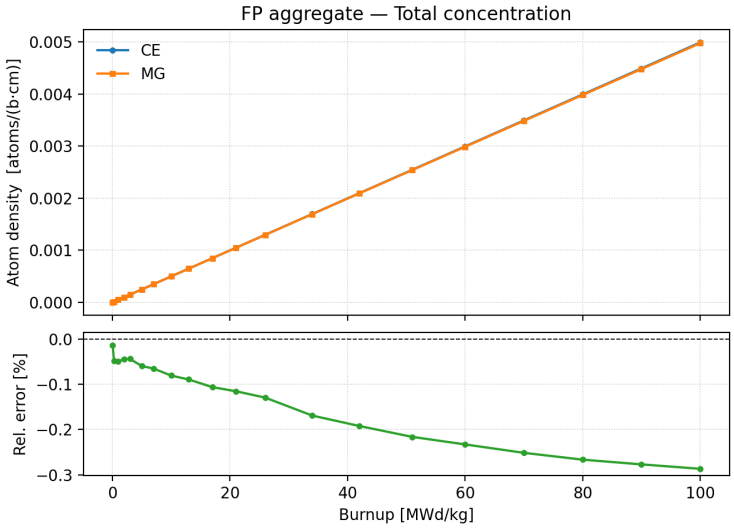

Figure 4.5: Total fission product concentration in the CE reference and MG depletion calculations, with relative deviation.

The reason for this close agreement is made explicit by the fission rate of Fig. 4.6, whose relative deviation oscillates within $\pm 0.007\,\%$ around zero with no systematic trend. The fission rate controls both the consumption of the heavy metals and the production of the fission products, and it is in turn fixed by the absolute flux imposed on the system. Its faithful reproduction in the multigroup calculation is therefore a direct consequence of the source-rate normalisation, which transfers the continuous-energy flux scale to the multigroup depletion. The agreement on the fission rate indicates that the CE flux is consistently transferred to the multigroup calculation and that the inventory evolution is governed by the same production and destruction balance in the two cases.

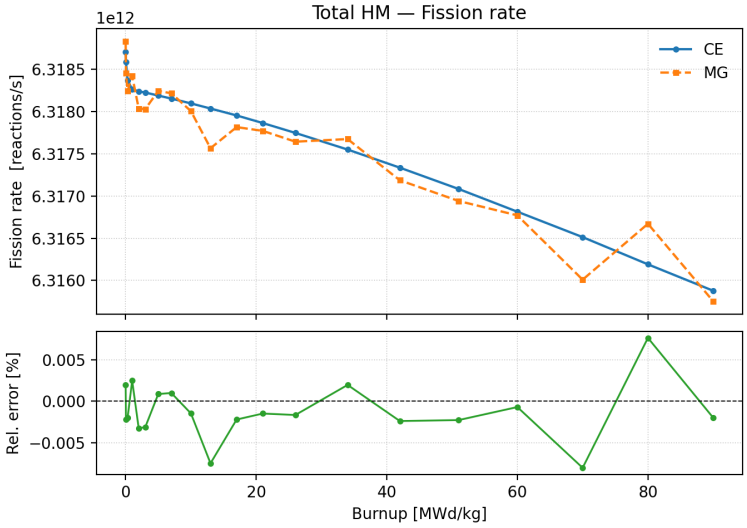

Figure 4.6: Total fission rate in the CE reference and MG depletion calculations, with relative deviation.

4.3.6 Heavy Metal Capture Rate

The total (n, γ) reaction rate on the heavy metal inventory, shown in Fig. 4.7, is reproduced with comparable accuracy. The multigroup calculation slightly overestimates the continuous-energy reference: the relative deviation starts from essentially zero, oscillates within the statistical noise in the first burnup steps, and grows to about $+0.02\,\%$ at the end of the cycle. A deviation of this magnitude is negligible for the purposes of the present analysis. The multigroup cross sections of the heavy metals reproduce the continuous-energy capture rate with high fidelity, and the slow growth of the deviation with burnup is compatible with the residual uncertainties of the multigroup representation.

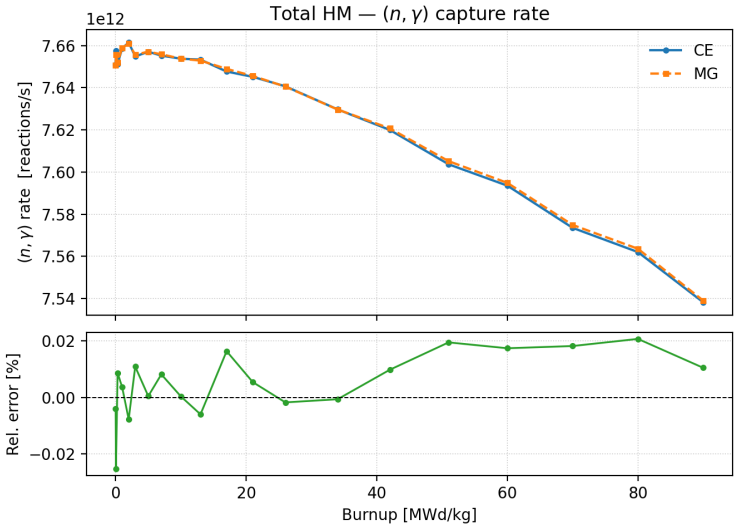

Figure 4.7: (n, γ) reaction rate of the heavy-metal inventory in the CE reference and MG depletion calculations, with relative deviation.

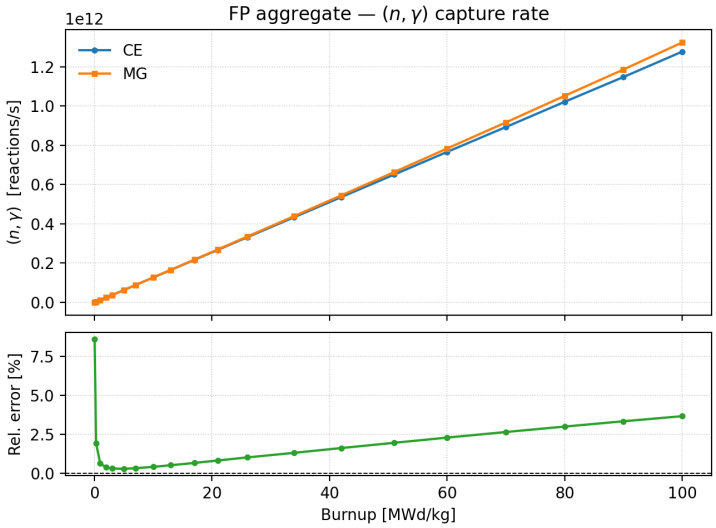

Figure 4.8: (n, γ) reaction rate of the fission product inventory in the CE reference and MG depletion calculations, with relative deviation.

4.3.7 Fission Product Capture Rate

The total (n, γ) reaction rate on the fission product inventory is a representative diagnostic of the lumped methodology and exhibits a more complex behaviour, shown in Fig. 4.8.

The relative deviation is large at the first burnup point, of the order of +8 %, drops within the first few steps to a few tenths of a percent, and then grows slowly with burnup, reaching approximately +3.7 % at the end of the cycle. This behaviour can be interpreted as the combination of two distinct mechanisms that dominate in two different phases of the irradiation.

- The first mechanism explains the overestimate of the capture rate in the earliest burnup steps and is related to fission products with non-zero yield but no cross-section evaluation. In the CE calculation, these nuclides contribute to the total fission product inventory, but not to the capture rate, because no microscopic cross section is available for them. In the LFP representation, instead, the effective yields are defined only over nuclides with available cross-section data. As a result, the mass associated with unevaluated fission products is redistributed onto lumped carriers that do contribute to capture, leading to an overestimate of the aggregate capture rate. This effect is relevant only at the beginning of the irradiation. The fraction of cumulative yield associated with unevaluated fission products is about 8 % at the first burnup step and decreases rapidly thereafter. As the short-lived precursors saturate or decay and the inventory becomes dominated by nuclides with available cross sections, this contribution vanishes within the first few steps. This early bias should not be interpreted only as an intrinsic deficiency of the lumped representation. It also reflects the incomplete CE cross-section coverage, and the LFP redistribution may be viewed as a rough compensation for capture contributions that cannot be explicitly scored.

- The second mechanism dominates over the remaining part of the irradiation, where the deviation grows approximately monotonically with burnup. This trend is associated with the imperfect representation of the evolving fission product composition by the fixed effective-yield structure of the LFPs, and is highlighted in Fig. 4.9. The figure compares two concentration-weighted one-group capture cross sections, both collapsed with the CE flux at the same burnup point: one obtained from the LFP carriers weighted by their concentrations, and one obtained from the explicit CE fission products weighted by their actual concentrations.

  The difference between the two averages reflects the mismatch between the internal composition represented by the LFPs and the true evolving fission product inventory. Each LFP cross section is built from effective-yield weights fixed at the time of library generation, whereas the CE reference reflects the actual burnup-dependent nuclide concentrations. As irradiation proceeds, the explicit inventory progressively departs from the proportions embedded in the LFPs, producing the residual growth of the capture-rate deviation.

The lower panel of Fig. 4.9 overlays three quantities and shows how the capture-rate deviation is built from the two mechanisms. The red curve is the fraction of fission products without cross section, responsible for the early spike, which decays almost immediately. The blue curve is the relative error of the averaged one-group capture cross section, which grows slowly with burnup up to about 4 %. The grey curve, representing the relative

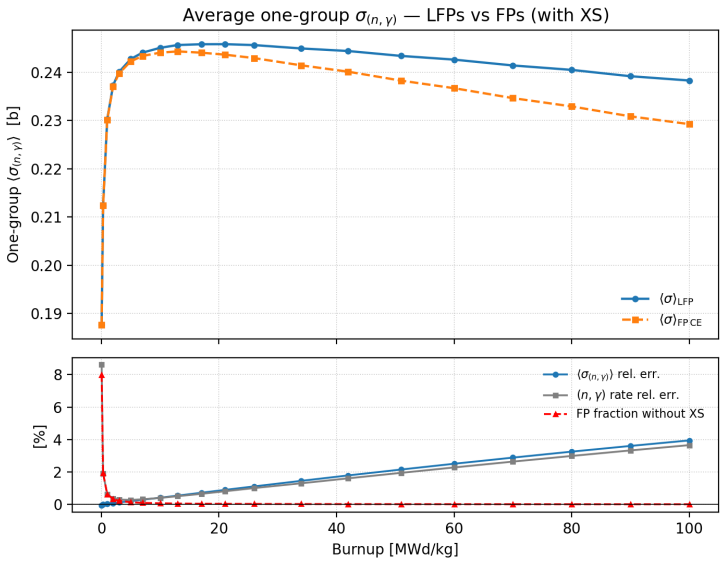

Figure 4.9: Burnup evolution of the concentration-weighted one-group (n, γ) cross sections for LFPs and explicit CE fission products, with relative deviations shown in the lower panel.

error of the capture rate, follows the red curve in the first steps and then coincides almost exactly with the blue curve for the rest of the cycle. The capture-rate error is therefore the superposition of an early missing-data contribution and a later composition-averaging contribution.

4.3.8 Impact on the System

The relative deviations discussed above are computed with respect to the fission product capture rate itself, which is small at the beginning of irradiation. To assess their impact on the system, Fig. 4.10 reports the same discrepancy normalised to the total fuel (n, γ) rate, i.e., to the scale at which it enters the neutron balance. On this scale, the deviation remains below about 0.5 % over the whole cycle and is essentially zero in the first burnup steps. The large early relative error therefore has little practical significance, because the fission product inventory is still small in absolute terms. The deviation acquires weight only in the second half of the cycle, where it nonetheless remains below half of a percent of the total capture rate.

A more detailed analysis of the individual fission products responsible for the composition error identified here is discussed in the following chapter.

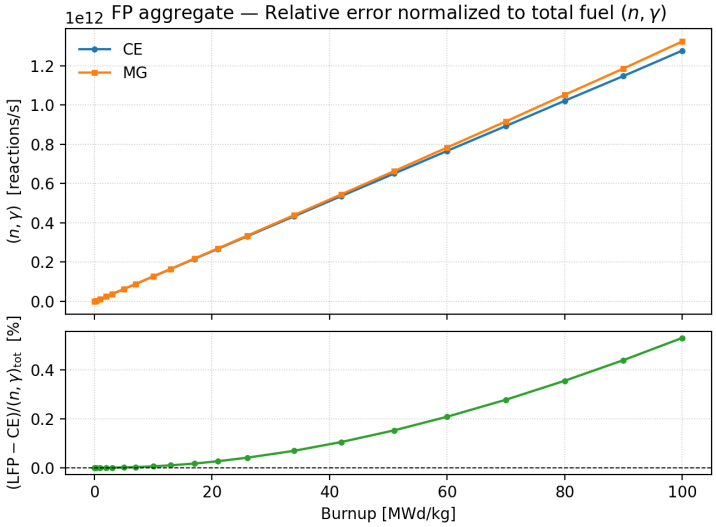

Figure 4.10: Absolute deviation of the fission product (n, γ) reaction rate, normalised to the total fuel (n, γ) reaction rate.

Chapter 5

A Python Framework for Coupled Multigroup Transport and Depletion

5.1 Introduction and Positioning

The methodology described in the previous chapters produced a set of multigroup cross-section libraries containing parent-specific lumped fission products, indexed by cumulative burnup. The intermediate analyses of Steps 1 and 2 exercised these libraries in controlled configurations, where one ingredient of the reaction rate of Eq. 4.2 was replaced by its multigroup counterpart while the remaining quantities were kept anchored to the continuous-energy reference. The final step requires applying the multigroup representation to all three ingredients simultaneously: cross sections, isotopic concentrations, and neutron flux. OpenMC natively supports both Monte Carlo transport in multigroup mode and depletion calculations, but it does not provide a single operator that combines the two into a self-consistent multigroup transport-depletion loop. The two depletion operators available in the package address different configurations. `CoupledOperator` runs a transport solve at every burnup step and couples it internally to the depletion equations, with the transport solve performed in continuous-energy mode. `IndependentOperator` advances depletion using a prescribed multigroup flux and cross-section library, without performing any transport solve. The configuration of interest in this work requires the two stages to be alternated externally. At every burnup step the multigroup transport must be executed to obtain the flux corresponding to the current composition, this flux must be combined with the multigroup cross sections valid at the current burnup, and the result must be passed to `IndependentOperator` to advance the composition.

To provide this external control structure, a Python framework was developed around the multigroup lumped fission product methodology of this work. The framework consists of two classes with distinct responsibilities: `MGXSLibraryManager`, which manages the parametrised cross-section libraries stored on disk and returns the data required at

a specified burnup, and `CoupledDepletionDriver`, which orchestrates the alternation of multigroup transport and depletion over the irradiation history. This chapter first describes the physical and numerical formulation of the coupling (Sec. 5.2), then the two classes (Sec. 5.3), the interpolation procedures for the multigroup data (Sec. 5.5), and the robustness features required by the implementation. The second part of the chapter verifies the framework in a controlled configuration and applies it to the reference SFR pin-cell.

5.2 Physical and Numerical Formulation

5.2.1 Sequential Transport–Depletion Coupling

The framework solves the coupled transport–depletion problem through a sequential first-order operator-splitting scheme. Over a burnup interval Δt, the transport problem is solved once using the nuclide composition $N_i(t_k)$ at the beginning of the step and the corresponding multigroup cross sections $\sigma^{(k)}$. This produces the multigroup flux $\phi_g^{(k)}$. The Bateman equations are then written as

$$\frac{dN_i}{dt} = \sum_j A_{ij}\big(\sigma^{(k)}, \phi^{(k)}\big)\, N_j \ , \tag{5.1}$$

where N_i is the atomic density of nuclide i, and A_{ij} is the depletion matrix coefficient describing the production or destruction of nuclide i from nuclide j through neutron-induced reactions and radioactive decays.

The depletion equations are integrated with an explicit predictor step, holding both $\sigma^{(k)}$ and $\phi^{(k)}$ constant over the full interval Δt. The transport solve is not repeated within the same depletion step. The updated composition, $N_i(t_{k+1})$, is then used as input for the next transport calculation, which recomputes the flux for the following depletion interval. The procedure is iterated until the end of the irradiation history.

In OpenMC terms, the scheme corresponds to the use of `PredictorIntegrator` with `IndependentOperator`. The latter is configured with the multigroup flux and the multigroup cross-section library at the burnup point of interest. The closure of the coupling is realised outside the OpenMC depletion operator by the driver loop, which alternates one transport solve and one depletion advance per step.

5.2.2 Burnup Normalisation and the Role of `HM_initial`

The depletion calculation is parametrised by cumulative burnup (BU) in units of MWd/kgHM, defined as

$$\mathrm{BU}(t) = \frac{1}{\mathrm{HM}_{\mathrm{initial}}} \int_0^t P(t')\, dt', \tag{5.2}$$

where P is the thermal power of the system, expressed in MW, and $\mathrm{HM}_{\mathrm{initial}}$ is the initially loaded heavy metal mass in the fuel, expressed in kg. Time is expressed in days when burnup is reported in MWd/kgHM. In the framework, $\mathrm{HM}_{\mathrm{initial}}$ is computed once at initialisation as the sum of the masses of all nuclides with atomic number $Z \geq 90$ in

the fresh fuel composition. This value is stored as an attribute of the driver and is not recomputed during irradiation, because the actual heavy metal mass decreases as burnup proceeds.

This convention is required to preserve the physical definition of burnup. The conversion of a prescribed burnup interval $\Delta\mathrm{BU}$ into the corresponding depletion time interval that the solver must advance is

$$\Delta t = \frac{\Delta\mathrm{BU}\,\mathrm{HM}_{\mathrm{initial}}}{P}. \tag{5.3}$$

If burnup were redefined at each step using the instantaneous heavy metal inventory, the elapsed time associated with a fixed $\Delta\mathrm{BU}$ would drift during the cycle, distorting the irradiation history. This is, in fact, the same convention is adopted internally by OpenMC's `CoupledOperator` when invoked with `timestep_units='MWd/kg'` . In the present framework, the convention is enforced by setting the `heavy_metal` attribute of `IndependentOperator` to $\mathrm{HM}_{\mathrm{initial}}$ at the beginning of each depletion step, independently of the current fuel composition.

5.3 The Two Classes of the Framework

This section describes the two classes that constitute the framework. The library manager is presented first because the driver depends on its services. The simulation loop, which uses both classes jointly, is described in Sec. 5.4.

5.3.1 The Library Manager

The `MGXSLibraryManager` manages the parametrised cross-section libraries produced by the procedure of Sec. 3.4. At initialisation, it loads the metadata describing the library, discovers the available burnup grids on disk, and returns the appropriate `openmc.MGXSLibrary` or `openmc.deplete.MicroXS` object for a requested burnup. When the requested burnup lies between two tabulated points, the returned object is interpolated between the adjacent libraries, as will be described in Sec. 5.5. The manager has no knowledge of the geometry, the time grid, or the irradiation progression of the calculation. It acts only as a data provider for the rest of the framework.

Library File Layout. The libraries are stored on disk through a directory hierarchy in which each cross-section file is indexed by the cumulative burnup at which it was generated, rather than by the index of the depletion step in the run that produced it. The choice decouples the library from the specific time-step sequence of that generation run: a tabulated burnup point identifies its own library file unambiguously, and the manager can supply the cross sections at any other burnup of interest.

A single JSON metadata file accompanies the library and acts as its descriptor. It records the energy structure, the burnup grids used for transport and depletion, the reference geometry and material conditions used during cross-section generation, and the temperatures at which the cross sections are tabulated. Operationally, the manager reads

this file at instantiation and checks that the declared energy structure is consistent with the one expected by the code. From a documentation perspective, the metadata make the generation conditions traceable, so that the geometry, temperature, and spectral environment associated with the cross sections remain explicit.

Two Burnup Grids The transport library is available at every burnup point of the reference history, including the fresh-fuel state at BU = 0. However, the depletion library has no entry at BU = 0: as discussed in Sec. 3.3, the construction of LFP depletion data requires a non-zero fission product inventory, which is absent in fresh fuel. The depletion grid therefore starts at the first non-zero burnup point.

At initialisation, the manager discovers the available depletion grid by scanning the corresponding subdirectories of the library tree. Requests for depletion cross sections below the first available point are handled through a fallback rule: the first available depletion library is used. This formalises the special case at the beginning of irradiation, where the fuel has not yet produced fission products but the depletion operator still requires a complete cross-section set.

Public Interface. In software terms, the public methods of a class define the interface through which external code interacts with it, while private methods encapsulate the internal logic. The manager exposes a small public interface dominated by two retrieval methods, `get_transport_library(burnup)` and `get_depletion_library(burnup)`, which return, respectively, an `openmc.MGXSLibrary` and an `openmc.deplete.MicroXS` object valid at the requested burnup point. The technical operations required to identify bracketing burnup points, load files and interpolate data are hidden in private methods. Additional public methods, such as `get_available_nuclides` and `filter_material_nuclides`, are used by the driver to construct OpenMC materials consistently with the content of the library at each burnup point.

5.3.2 The Coupled Driver

The `CoupledDepletionDriver` owns the simulation loop. It stores the current fuel composition, advances the cumulative burnup according to Eq. 5.3, invokes the transport solver and the depletion operator at each step, and manages the exchange of information between the two stages. Unlike the manager, the driver does not know how the cross sections are stored on disk or interpolated between burnup points. It relies on the manager to retrieve the required data and treats the resulting objects as inputs to OpenMC.

Builder-Pattern Interface for the Geometry. The geometry, transport settings, and tallies are not built inside the driver. They are supplied by the user through builder functions that the driver invokes at every burnup step. This choice keeps the driver independent of the specific model configuration and makes the framework reusable, in principle, for calculations beyond the pin-cell case considered here.

Public Interface. The driver exposes a single high-level public method, `run`, which executes the irradiation history. Internally, it stores the current materials, the current cumulative burnup, and the prescribed time grid, and it calls the manager and builder functions at each step.

5.4 The Coupled Simulation Loop

The execution of a burnup interval $[t_k, t_{k+1}]$ proceeds through five stages.

1. The driver retrieves the transport microscopic cross-section data and reconstructs the fuel macroscopic cross-section. It queries the manager for the multigroup transport library at the cumulative burnup $\mathrm{BU}(t_k)$. The manager identifies the two adjacent tabulated burnup points that bracket the request and returns the linearly interpolated `openmc.MGXSLibrary`; if the request coincides with a tabulated point, the corresponding file is returned directly.

2. The OpenMC model is constructed. The driver invokes the builder functions with the current materials and the interpolated cross-section library as inputs. The result is a fully specified OpenMC model in multigroup mode, with cross sections for the depletable fuel taken from the interpolated library and those for non-depletable materials.

3. The transport solve is performed. The OpenMC eigenvalue calculation produces the multigroup flux $\phi_g^{(k)}$ in the fuel region, the eigenvalue $k_{\mathrm{eff}}^{(k)}$, and reaction rates saved for post-processing. The flux extracted from the statepoint is expressed in OpenMC's per-source-particle units. Its conversion to an absolute flux is handled by the depletion operator through the fission-Q normalisation mode, which derives the source rate from the prescribed power and from the fission Q-values in the depletion chain.

4. The driver retrieves the depletion cross-section data. It queries the manager for the multigroup `MicroXS` object at the same cumulative burnup $\mathrm{BU}(t_k)$, using the same interpolation logic as for the transport library and the initial fallback rule described in Sec. 5.3.1.

5. The depletion equations are integrated. The `IndependentOperator` is constructed using the current fuel material, the multigroup flux from the transport solve, and the interpolated `MicroXS` object. The fission-Q normalisation mode is selected. The operator is then passed to `PredictorIntegrator`, which advances the composition over $\Delta t = t_{k+1} - t_k$ while keeping cross sections and flux fixed at their beginning-of-step values. At the end of the integration, the depleted material is extracted, the cumulative burnup is advanced to $\mathrm{BU}(t_{k+1}) = \mathrm{BU}(t_k) + \Delta\mathrm{BU}$, and the loop proceeds to the next interval.

5.4.1 Pseudo-Code of the Driver Loop

The structure of the loop is summarised below.

```
initialise driver with: initial materials, time grid,
                        builders, library manager
HM_initial <- compute heavy metal mass at t = 0
BU <- 0

for each step k with interval dt:

    # Stage 1: retrieve transport XS
    mgxs_lib <- manager.get_transport_library(BU)

    # Stage 2: build OpenMC model
    geom    <- geometry_builder(current_materials)
    sett    <- settings_builder(current_materials)
    tallies <- tallies_builder(current_materials)

    # Stage 3: run transport
    statepoint <- openmc.run(geom, sett, tallies, mgxs_lib)
    phi_g, k_eff <- extract from statepoint

    # Stage 4: retrieve depletion XS
    micro_xs <- manager.get_depletion_library(BU)

    # Stage 5: advance depletion
    operator <- IndependentOperator(current_materials,
                                    phi_g, micro_xs,
                                    normalization_mode='fission-q')
    operator.heavy_metal <- HM_initial
    integrator <- PredictorIntegrator(operator,
                                      timesteps=[dt],
                                      power=P)
    integrator.integrate()

    # Update state
    current_materials <- read depleted materials
    BU                <- BU + dBU
```

5.5 Interpolation of the Cross-Section Libraries

The cross-section libraries are tabulated on a discrete burnup grid. When the driver requests cross sections at a burnup BU that does not coincide with a tabulated point, the manager constructs an interpolated library usable by the OpenMC solvers. Linear interpolation is performed between the two adjacent burnup points, $\mathrm{BU}_{\mathrm{low}}$ and $\mathrm{BU}_{\mathrm{high}}$,

that bracket the request. The interpolation weight is

$$\alpha = \frac{\mathrm{BU} - \mathrm{BU}_{\mathrm{low}}}{\mathrm{BU}_{\mathrm{high}} - \mathrm{BU}_{\mathrm{low}}}, \qquad 0 \leq \alpha \leq 1. \tag{5.4}$$

The same procedure is applied to both the `MGXSLibrary` used by the transport solver and the `MicroXS` object used by the depletion operator, with adjustments required by the internal structure of each object.

5.5.1 Linear Blend of Scalar Cross Sections

For each nuclide in the library, scalar multigroup quantities such as total, absorption, fission, and ν-fission cross sections are blended element-wise across the 33 energy groups. Denoting by σ_{low} and σ_{high} the values of a given quantity at the two bracketing burnup points, the interpolated value is

$$\sigma(\mathrm{BU}) = (1 - \alpha)\, \sigma_{\mathrm{low}} + \alpha\, \sigma_{\mathrm{high}}. \tag{5.5}$$

The fission spectrum χ is also blended linearly and then explicitly renormalised so that $\sum_g \chi_g(\mathrm{BU}) = 1$. This removes residual deviations from unity that may arise from combining two unit-sum vectors with slightly different entries.

Sanitisation of Zero Fission Spectra. A pathological case arises when a fissionable nuclide has an entirely zero χ vector. This can occur for minor actinides present in trace concentrations during early burnup steps, for which the estimated multigroup fission spectrum is identically zero in the generated library.

In the OpenMC C++ MGXS reader, vector fission spectra are normalised by division by their outgoing-group sum. An all-zero χ vector therefore produces non-finite entries. When microscopic data are combined into material-wise macroscopic MGXS data, the macroscopic fission spectrum is formed from weighting the microscopic contributions. Consequently, even a nuclide with negligible or zero ν-fission production can contaminate the material spectrum, since the floating-point product of zero and NaN remains NaN.

Hence, the resulting failure is not limited to histories that explicitly interact with that nuclide. In multigroup transport, OpenMC samples fission neutron emission energies from the macroscopic material spectrum. If this distribution contains non-finite values, the cumulative probability used to sample the outgoing energy group may fail to select a valid group. Invalid fission-bank source sites can then be propagated to subsequent generations of the eigenvalue calculation and lead to runtime failures.

The framework intercepts this case when preparing the library for transport. For every nuclide whose χ vector is entirely zero, the vector is replaced by the uniform distribution $\chi_g = 1/G$, where $G = 33$ in the present group structure, and the corresponding fission and ν-fission cross sections are set to zero. The artificial spectrum is therefore only a numerical placeholder: the nuclide contributes no fission neutron source, while the MGXS reader and the macroscopic material-spectrum construction no longer encounter the zero-normalisation case.

5.5.2 Interpolation of the Scattering Data

As discussed in Sec. 3.4.1, the multiplicity matrix, defined as the ratio between the ν-scatter and scatter matrices in Eq. 3.5, is not an additive quantity. It must therefore be reconstructed from the two additive quantities that define it. At each of the two bracketing burnup points, the scatter matrix $\sigma_{s,g\to g'}$ is recovered from the stored datasets as the element-wise ratio of the ν-scatter matrix to the multiplicity matrix. The recovered scatter matrix and the stored $\nu_s\sigma_s$ matrix are then blended linearly and independently according to Eq. 5.5. Finally, the interpolated multiplicity matrix is reconstructed as the ratio of the two interpolated matrices, preserving its physical interpretation.

Sparse Storage of the Scatter Matrix. OpenMC stores the scatter matrix in a sparse format that exploits the kinematic constraints of the scattering process. For each incoming group g, a scattering event can transfer the neutron only to a limited range of outgoing groups. Upscattering, where the outgoing group has higher energy than the incoming one, is possible at thermal energies through elastic scattering on moving atoms, but it is negligible in the fast spectrum considered here, where the low-energy flux is several orders of magnitude smaller than the fast flux. The library therefore stores, for each incoming group, only the range $[g_{\min}(g), g_{\max}(g)]$ of outgoing groups in which the matrix is non-zero, together with the corresponding values. Of the 33×33 entries of the full matrix, only a limited subset is stored for each nuclide.

Interpolation Across Different Sparsity Patterns. The interpolation procedure handles two cases. When the sparsity pattern of the scatter matrix is identical at the two bracketing burnup points, the stored arrays have the same shape and are combined element-wise through Eq. 5.5. This is the most frequent case in practice.

The sparsity patterns may differ at one or more incoming groups because of the stochastic nature of the Monte Carlo tallies used to generate the libraries: a very weak scattering channel may be sampled as non-zero in one run and as zero in another. In this case, the framework expands the affected rows to their full dense representation and applies a band-aware blend. An entry (g, g') that lies within the stored band of both libraries is interpolated linearly. An entry that lies within the band of only one library is taken from the library that contains it, without blending against an artificial zero. This treatment reflects the fact that the absence of an entry from a stored band is caused by the generation process, not necessarily by a physically zero scattering cross section. Averaging the present value against an artificial zero would therefore introduce a spurious reduction. The interpolated matrix is finally recompressed into the sparse format on the union of the endpoint bands.

Numerical Stability of the Multiplicity Reconstruction. Both the recovery of σ_s from the stored quantities and the reconstruction of the multiplicity matrix involve divisions that may become singular in specific entries. The framework protects against two cases that may arise. First, when the multiplicity at one bracketing burnup point is zero in a given entry, which occurs when that scattering channel was not sampled in the corresponding tally, the recovered σ_s in that entry is set to zero rather than to an

undefined value, consistently with the absence of a physical scattering signal. Second, the reconstructed multiplicity is set to unity where $\sigma_{s,\mathrm{interp}} = 0$. This fallback is physically harmless for entries with no scattering, because the multiplicity multiplies a zero cross section in any subsequent reaction-rate calculation.

5.6 Verification of the Framework with Explicit Fission Products

Before applying the coupled framework to the lumped fission product representation, a preliminary calculation was performed in which both transport and depletion were carried out in multigroup mode, while fission products were treated explicitly. This calculation has two purposes.

- It verifies that the framework, with its two classes and coupled loop, operates correctly on a configuration that does not involve lumping.
- It provides an intermediate reference that isolates the effect of the multigroup treatment alone, namely the combined effect of multigroup transport and multigroup depletion. Comparing this explicit-fission-product calculation with the lumped calculation of Sec. 5.7, both measured against the same continuous-energy reference, allows the net effect of lumping to be isolated, which is the final objective of this work. The present analysis is not intended to demonstrate the general adequacy of multigroup transport or Monte Carlo cross-section generation, which are assumed as established in this work.

For this calculation, the depletion chain was restricted to reactions for which multigroup cross sections had been generated: (n,γ), $(n,2n)$, (n,α), (n,p), and fission. These channels account for nearly the totality of the neutron-induced reactions relevant to the depletion of the inventory in the present fast spectrum. This restriction keeps the depletion chain consistent with the cross-section data managed by the framework and is not expected to introduce an appreciable bias for the quantities analysed here.

5.6.1 Multiplication Factor

Figure 5.1 reports the evolution of the multiplication factor in the explicit multigroup calculation and in the continuous-energy reference. The two curves track each other closely over the entire irradiation history. The difference Δk remains mostly between zero and about +200 pcm and no strong systematic trend. The multigroup calculation shows a small overestimate of the multiplication factor, consistent with the expected effect of the multigroup representation of the cross sections and transport solve. The magnitude of the deviation is within the range expected for a coupled multigroup calculation compared with a continuous-energy reference.

This agreement provides the primary verification that the coupled loop operates correctly: the alternation of transport solves and depletion advances, together with the bookkeeping of the inventory, reproduces the reference within an acceptable margin.

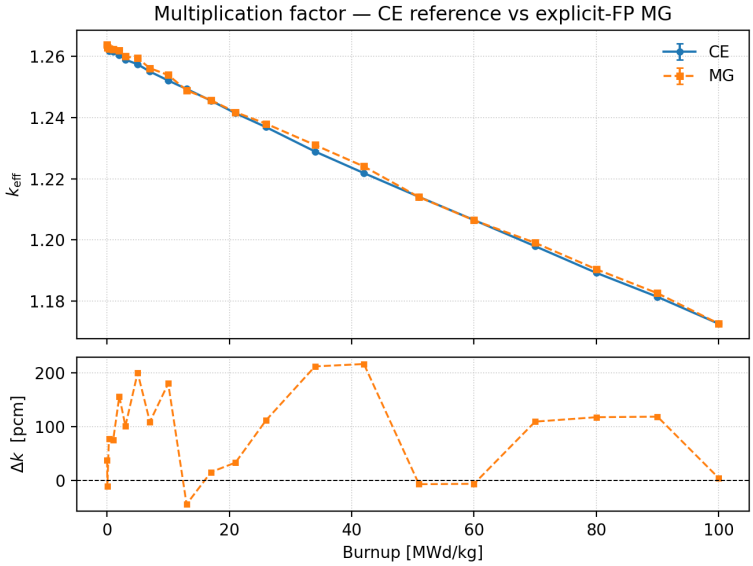

Figure 5.1: Comparison of k_{eff} between the CE reference and the MG explicit-FP calculation.

5.6.2 Nuclide Inventory

The inventory is reproduced with the same fidelity observed in the multigroup depletion analysis of Sec. 4.3.5. The total heavy metal concentration of Fig. 5.2 agrees with the continuous-energy reference to better than 0.034 % over the whole cycle. The total fission product concentration of Fig. 5.3 remains within about −0.29 %, with a small negative deviation that grows slowly in magnitude with burnup. This agreement confirms that the production–destruction balance is preserved when the flux is computed by multigroup transport rather than supplied from the continuous-energy calculation.

5.6.3 Fission Product Capture Rate

The most informative quantity of this verification is the (n, γ) reaction rate on the fission product inventory, shown in Fig. 5.4. Its relative deviation from the continuous-energy reference oscillates within about ±0.25 % around zero, with a noisy character and no significant systematic trend over the cycle. This indicates that the residual discrepancy is mainly associated with statistical uncertainty and the multigroup representation of the cross sections. The result is markedly different from the lumped depletion analysis of Sec. 4.3.7, where the deviation reached about +8 % at the first burnup step and increased again to about +3.7 % at the end of the cycle.

In the explicit calculation, the two mechanisms that governed the lumped result are absent. First, there is no composition error associated with effective yields, because the fission products are tracked explicitly. Second, there is no redistribution of the mass of

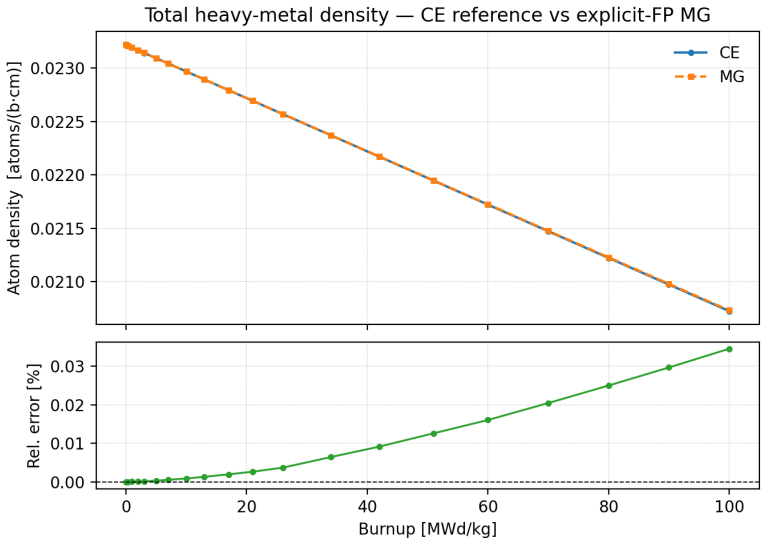

Figure 5.2: Total HM concentration in the CE reference and MG explicit-FP calculation.

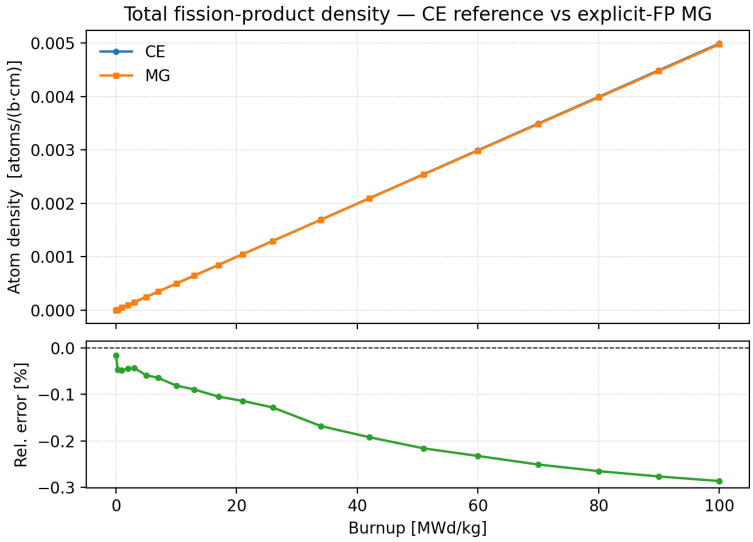

Figure 5.3: Total FP concentration in the CE reference and MG explicit-FP calculation.

fission products without cross-section evaluations onto lumped carriers. The fact that the deviation decreases from several percent to a few tenths of a percent when lumping is removed confirms that the larger discrepancy observed in Sec. 4.3.7 originates mainly from the lumped representation itself.

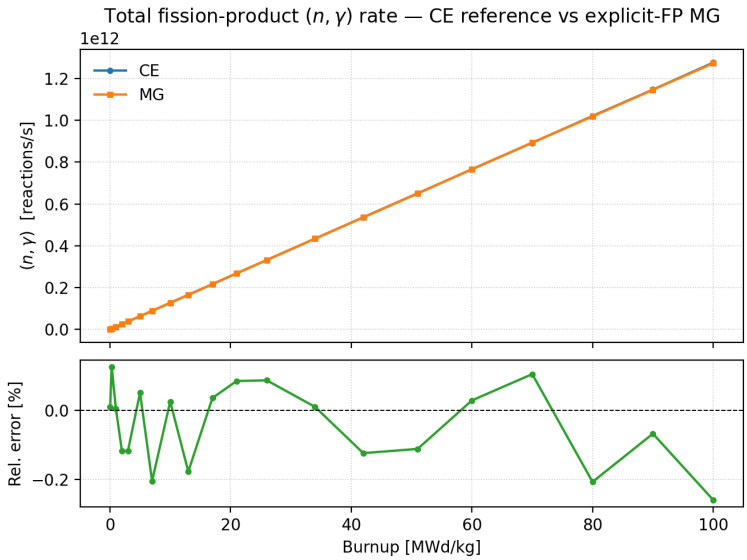

Figure 5.4: Total FP (n, γ) reaction rate in the CE reference and MG explicit-FP calculation.

5.7 Coupled Multigroup Transport and Depletion with Lumped Fission Products

The final calculation of the assessment exercises the full multigroup framework with the lumped fission product representation. Both transport and depletion are performed in multigroup mode, the flux is computed by the transport solver at each burnup step, and the fission products of the depletion chain are represented by lumped pseudo-nuclides.

5.7.1 Heavy Metal Inventory and Aggregate Capture Rate

The total heavy metal concentration of Fig. 5.5 reproduces the continuous-energy reference with a relative deviation below 0.034 % over the entire cycle, growing smoothly with burnup. The agreement is indistinguishable, on the present scale, from that obtained in Sec. 4.3.5 and Sec. 5.6.1. This indicates that the introduction of lumping does not perturb the aggregate heavy metal balance to an appreciable extent.

The total fission rate, reported in Fig. 5.6, agrees with the continuous-energy reference more closely than in the analysis of Sec. 4.3.5. The residual statistical fluctuations are slightly smaller because the absolute flux scale is reset at each burnup step through fission-Q normalisation against the prescribed power, rather than imposed through the source rate. Consequently, small inventory drifts are partly absorbed by corresponding adjustments of the flux normalisation.

Figure 5.7 compares the total heavy metal radiative-capture reaction rate obtained with the CE reference, the MG calculation with explicit FP, and the MG calculation with LFP.

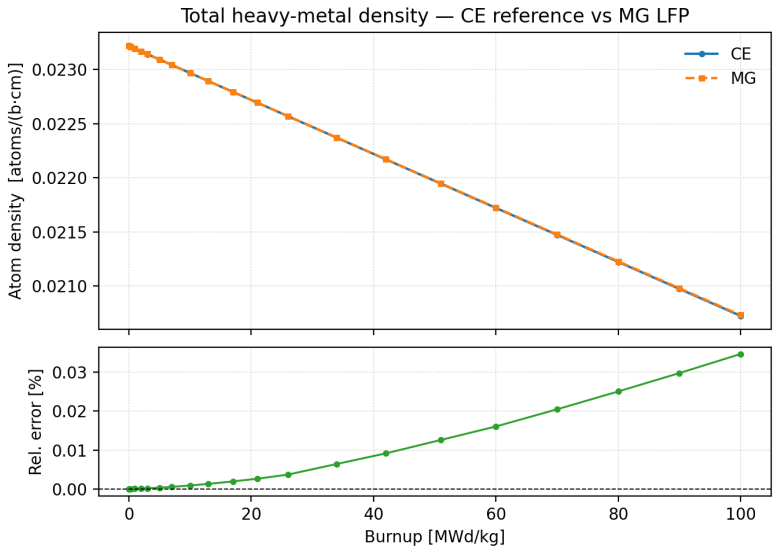

Figure 5.5: Total HM atom density along the irradiation history in the CE reference and MG LFP calculation. The lower panel shows the relative deviation.

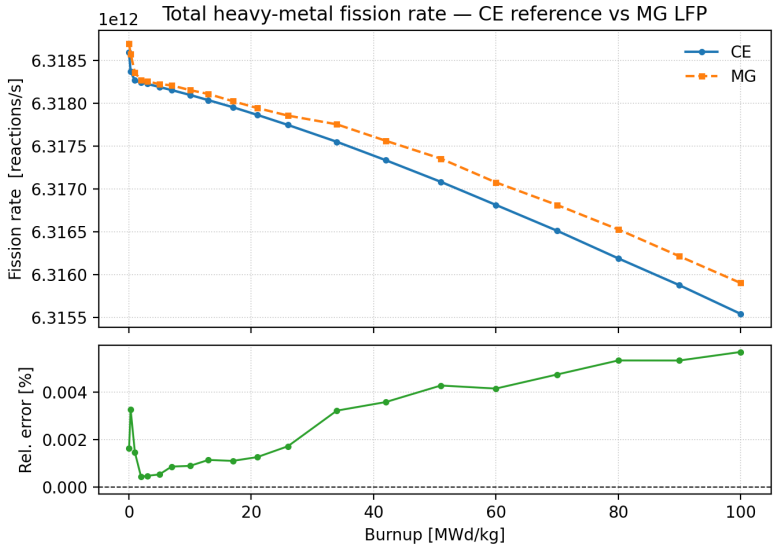

Figure 5.6: Total fission reaction rate in the CE reference and MG LFP calculation. The lower panel shows the relative deviation.

All three calculations reproduce the same global evolution with burnup. The discrepancies introduced by the multigroup treatment remain small over the full irradiation history: the

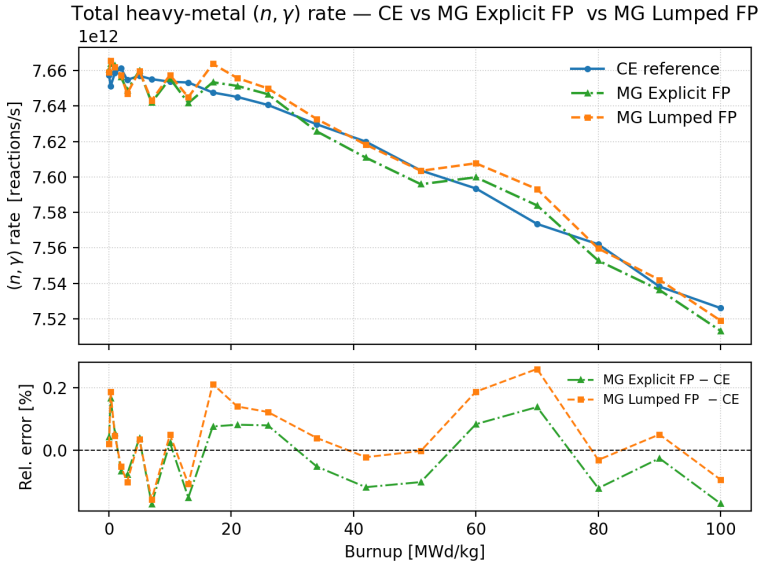

Figure 5.7: Total HM (n, γ) reaction rate in the CE reference, MG explicit-FP calculation, and MG LFP calculation. The lower panel reports the relative deviation of each MG result with respect to the CE reference.

explicit-fission-product calculation deviates from the reference by less than approximately 0.15 %, while the LFP calculation remains within approximately 0.3 %.

The additional deviation introduced by lumping is therefore limited for this aggregate quantity. The LFP curve follows the same shape as the explicit-fission-product multigroup curve, with a difference between the two of the order of 0.1 %. The relative errors oscillate around zero and do not show a strictly monotonic trend with burnup. This suggests that the deviations are not governed by a single systematic bias, but by the combined effect of several competing approximation, including the energy discretisation, the statistical uncertainty in the MC-generated MGXS libraries, and the simplified representation of the fission product inventory in the LFP model.

5.7.2 Individual Heavy Metal Nuclides

The picture changes when the heavy metal inventory and capture rate are decomposed by nuclide. Figure 5.8 reports the relative deviation of the concentration and of the (n, γ) capture rate for the four dominant heavy metal nuclides: ^{238}U, ^{239}Pu, ^{240}Pu, and ^{241}Pu. The errors on individual nuclides are, in several cases, larger than the errors on the aggregate quantities discussed above. The concentration of ^{238}U is reproduced to better than +0.06 %, but with a systematic positive trend that grows monotonically with burnup. The concentrations of ^{239}Pu, ^{240}Pu, and ^{241}Pu are slightly underestimated, with end-of-cycle deviations of approximately −0.025 %, −0.075 %, and −0.06 %, respectively.

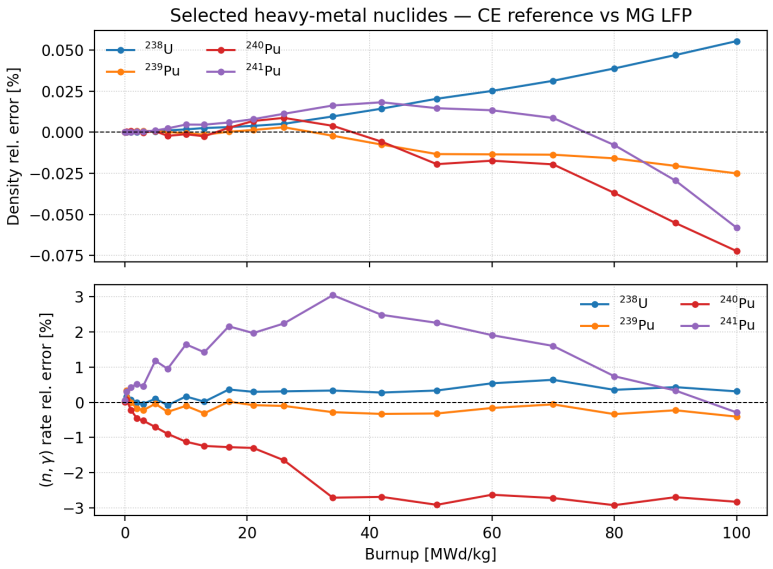

Figure 5.8: Relative deviation of the concentration (upper panel) and (n, γ) reaction rate (lower panel) of selected HM nuclides between the MG LFP calculation and the CE reference.

The concentration pattern is physically coherent: the positive deviation on ^{238}U corresponds to a slightly lower consumption by capture, which in turn produces a slightly reduced inventory of ^{239}Pu and of the heavier plutonium isotopes, consistent with the negative deviations observed for those nuclides.

The corresponding deviations on individual capture rates are larger. The capture rate of ^{238}U is overestimated by about $+0.3\,\%$ to $+0.6\,\%$ depending on burnup, while the capture rate of ^{239}Pu is reproduced within a comparable range. The capture rate of ^{240}Pu is systematically underestimated by about $-3\,\%$ over most of the cycle, and the capture rate of ^{241}Pu shows the largest relative deviation, reaching about $+3\,\%$ at intermediate burnup before slowly returning towards zero.

The fact that the aggregate capture rate remains within about 0.1–0.2 %, may indicate that partial compensation occurs when the nuclide contributions are summed. This compensation should not be overstated: the aggregate is weighted by the absolute capture rates, which in this configuration are dominated by ^{238}U and ^{239}Pu. These nuclides have the smallest relative deviations among the four, whereas the larger relative deviations of ^{240}Pu and ^{241}Pu have smaller absolute weight.

5.7.3 Fission Product Inventory and Capture Rate

The total fission product concentration of Fig. 5.9 reproduces the continuous-energy reference with a relative deviation of about $-0.29\,\%$ at the end of the cycle. The deviation

has the same negative sign and slow growth in magnitude already observed in Secs. 4.3.5 and 5.6.1.

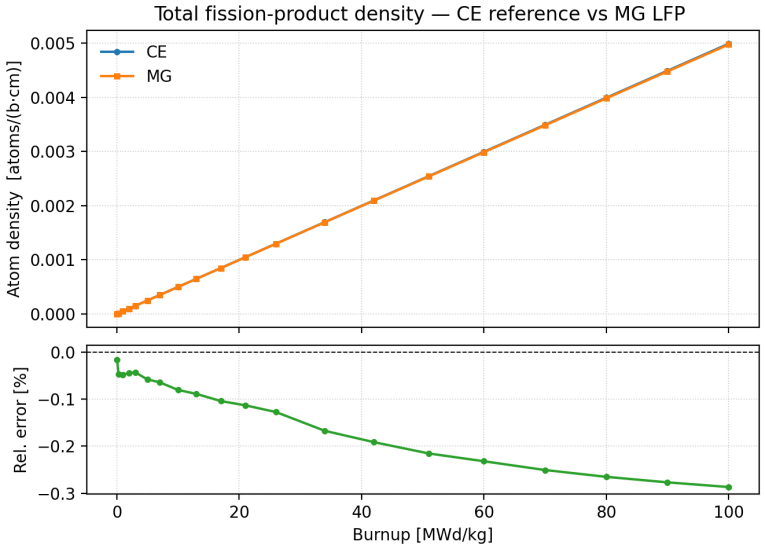

Figure 5.9: Total FP atom density in the CE reference and MG LFP calculation. The lower panel reports the relative deviation.

The aggregate (n, γ) capture rate on the fission product inventory is shown in Fig. 5.10, together with the continuous-energy reference and the explicit-fission-product multigroup calculation. Figure 5.11 reports the same quantity but normalised to the total fuel capture rate.

The lumped calculation exhibits the pattern already discussed in Sec. 4.3.7: a large relative deviation of about +8 % at the first burnup step, caused by the redistribution of fission products without cross-section evaluations onto capture-bearing nuclides, followed by a rapid decrease in the first few steps and a slow monotonic growth to about +3.7 % at the end of the cycle. The later growth is driven by the progressive departure of the true fission product composition from the internal proportions frozen into the effective yields.

The explicit-fission-product calculation remains within a few tenths of a percent over the whole cycle. The contrast between the two calculations confirms, in the fully coupled regime, that the dominant discrepancy is localised in the lumped representation rather than in the multigroup treatment itself.

The deviation normalised to the total fuel capture rate, shown in Fig. 5.11, quantifies the weight of the fission product capture discrepancy on the overall neutron balance. The lumped calculation grows monotonically with burnup and reaches about +0.53 % at the end of the cycle, whereas the explicit-fission-product calculation remains within about 0.04 % over the entire irradiation. The discrepancy introduced by lumping is therefore visible but contained at the level of the total fuel capture rate.

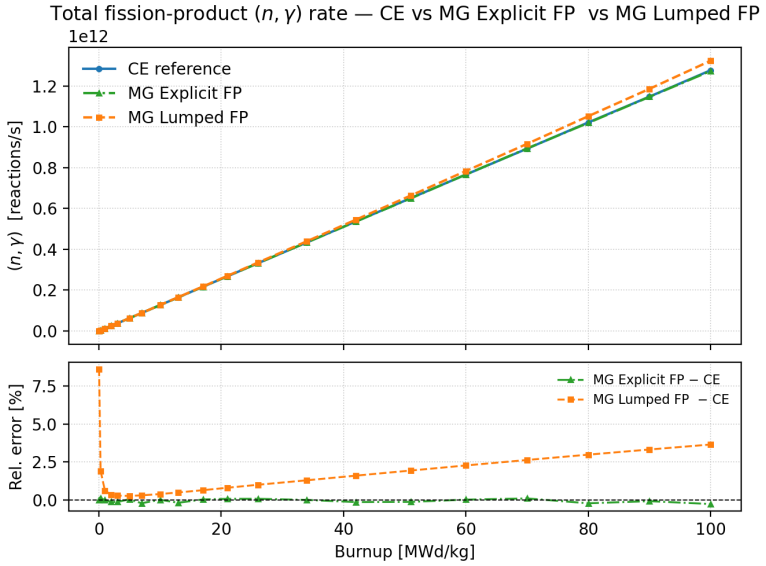

Figure 5.10: Total FP (n, γ) reaction rate in the CE reference, MG explicit-FP calculation, and MG LFP calculation. The lower panel reports the relative deviation of each MG result with respect to the CE reference.

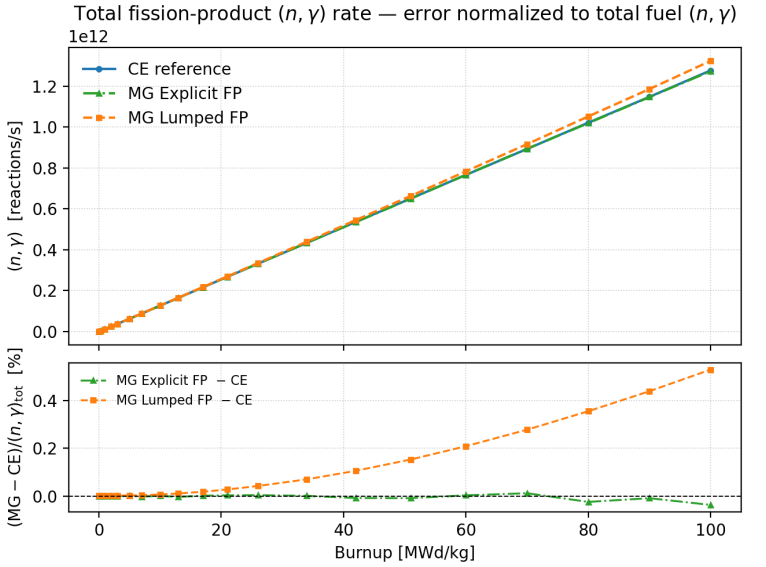

Figure 5.11: Relative deviation of the FP (n, γ) reaction rate, normalised to the total fuel capture rate, for the MG explicit-FP and MG LFP calculations with respect to the CE reference.

5.7.4 Multiplication Factor

The multiplication factor along the irradiation history is examined through the complementary comparisons of Figs. 5.12 and 5.13. In both figures, the upper panel shows the continuous-energy reference, the explicit-fission-product multigroup calculation, and the lumped multigroup calculation. All three reproduce the expected monotonic decrease of k_{eff} with burnup over the full cycle and remain close to the reference over the 100 MWd/kgHM irradiation history.

The lower panel of Fig. 5.12 reports the deviation of each multigroup calculation with respect to the continuous-energy reference. Up to about 40 MWd/kgHM, the two multigroup curves are in close agreement. Their common deviation from the reference, of the order of a few tens to about +200 pcm with a noisy character, is attributable mainly to the multigroup representation itself. In the second half of the cycle, the two curves separate more visibly: the lumped calculation progressively shifts towards negative values, reaching about −200 pcm at 100 MWd/kgHM. The end-of-cycle separation between the lumped calculation and the other two calculations is therefore associated with the effect of lumping.

The difference between the two multigroup calculations, which isolates the effect of lumping alone, is reported directly in the lower panel of Fig. 5.13. The sign of the deviation is physically consistent with the interpretation developed in the previous subsections: the lumped calculation systematically overestimates the (n, γ) capture rate on the fission product inventory with respect to the explicit-fission-product calculation.

Although the magnitude of the deviation introduced by lumping is limited, the result should be interpreted within the conditions of the present analysis. The comparison was performed using the same fuel and geometry adopted for the generation of the multigroup cross sections and at burnup points lying on the tabulated grid of the libraries. The assessment therefore exercises the framework in a configuration that is, by construction, as favourable as possible to the multigroup representation.

5.7.5 Back-Projection of the Lumped Inventory

The aggregate analysis of Sec. 4.3.7 identified the imperfect representation of the fission product composition by the effective yields as the main residual source of error of the lumped methodology in the second half of the cycle. To examine this source of error nuclide by nuclide, a back-projection of the lumped inventory onto the explicit fission product set was performed at each burnup point. The procedure is the inverse of the lump construction: the concentration of each lumped are distributed among the fission products that compose it using the corresponding effective yields. This produces an explicit projected inventory, N_i^{proj}, which can be compared nuclide by nuclide with the continuous-energy reference, N_i^{CE}.

Concentrations. Figure 5.14 reports the relative deviation of the projected concentration with respect to the continuous-energy reference for the twenty most abundant fission products at end of life (EOL), in selected burnup points. In the first burnup steps, the pattern is nearly uniform across the displayed nuclides and reflects the redistribution of

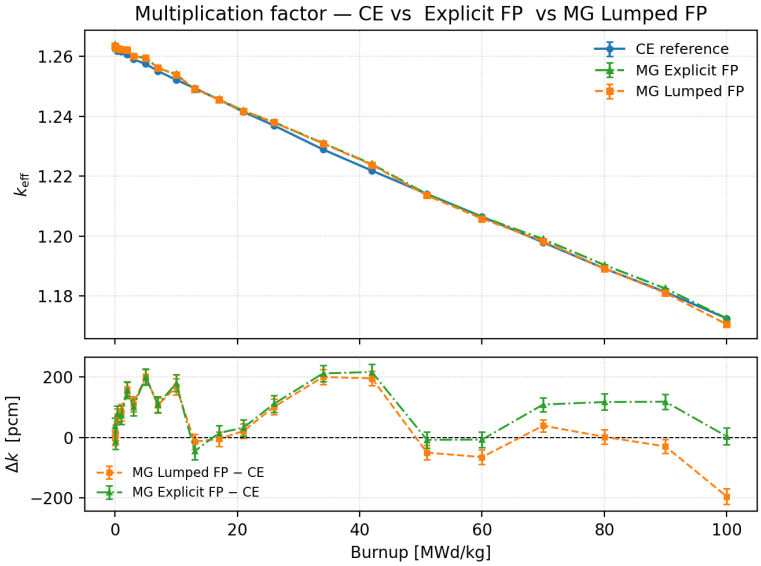

Figure 5.12: k_{eff} evolution in the CE reference, MG explicit-FP calculation, and MG LFP calculation. The lower panel reports the deviation of each MG result with respect to the CE reference.

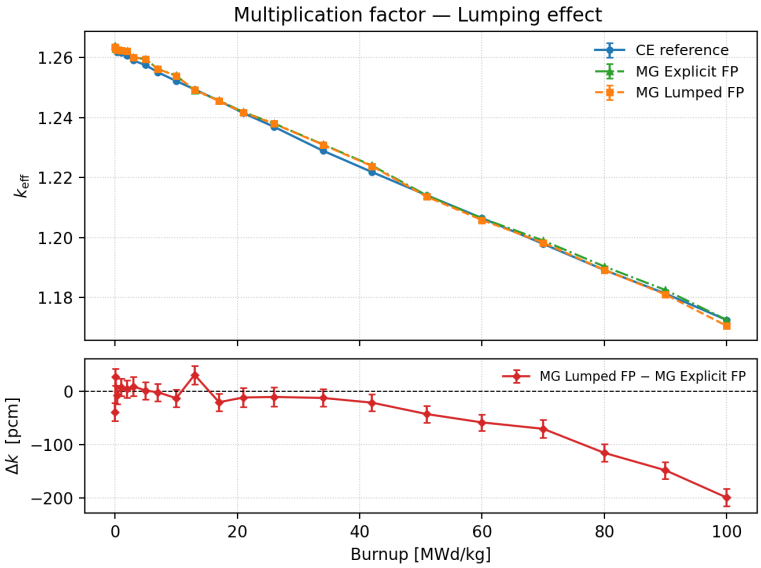

Figure 5.13: k_{eff} evolution with the lower panel showing Δk between the MG LFP and MG explicit-FP calculation.

fission products without cross-section evaluations onto capture-bearing nuclides, already discussed in Sec. 4.3.7. The projected concentrations exceed the continuous-energy reference by about +8.7 % at the first step and +1.9 % at the second, for essentially all displayed nuclides, because this redistribution is applied uniformly across the lump composition.

From the successive steps onwards, the deviations become nuclide-specific and reflect the composition error proper. Their magnitude remains small for most of the inventory, of the order of few percent or less for the majority of displayed nuclides at EOL. A few isotopes show larger deviations of either sign, such as ^{102}Ru at −5.4 %, ^{103}Rh at +5.6 %, ^{99}Tc at +7.4 %, and ^{101}Ru at +7.3 %.

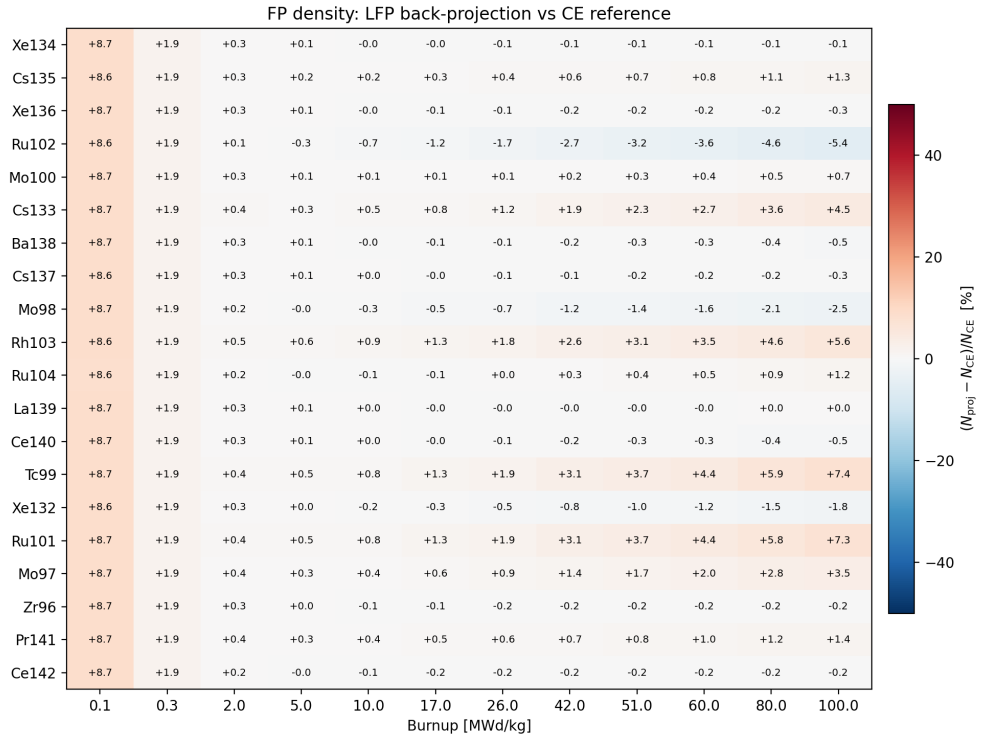

Figure 5.14: Relative deviation at selected burnup points of the back-projected FP concentration with respect to the CE reference, for the twenty most abundant nuclides at EOL.

Capture Rates. Figure 5.15 reports the analogous deviation for the (n, γ) reaction rate, considering the twenty nuclides with the largest absolute capture rates at EOL. The deviations are systematically larger than those on concentration.

Several nuclides show deviations of a few percent at end of cycle, but the largest discrepancies are concentrated on absorbers such as ^{147}Sm, ^{149}Sm, and ^{151}Sm. Similar behaviour is also observed for europium and gadolinium isotopes outside the top twenty. These nuclides show a similar behaviour and tend to be overestimated in the reconstructed

inventory. This suggests that the discrepancy is not only due to propagation along decay chains, but also to the sensitivity of specific nuclides to neutron absorption during irradiation.

The effective yields used in this work are generated through auxiliary depletion calculations under prescribed flux and spectral conditions. They therefore embed not only cumulative fission product production, but also the balance between production and depletion by neutron capture established in that specific environment. Nuclides such as samarium, europium, and gadolinium remain relatively strong absorbers even in a fast spectrum; their equilibrium concentrations are therefore sensitive to the local spectrum and flux level. Differences between the auxiliary depletion conditions used to construct the effective yields and the actual conditions of the CE reference can consequently be amplified in the reconstructed concentrations.

The comparatively good agreement observed for isotopes such as ^{133}Cs and ^{135}Cs supports this interpretation. Although these nuclides are also fed by extended precursor chains, their neutron-capture rates in the fast spectrum are much smaller. Their inventories are therefore primarily determined by cumulative production rather than by a detailed balance between production and neutron absorption, making them less sensitive to spectral mismatches between the auxiliary and reference calculation.

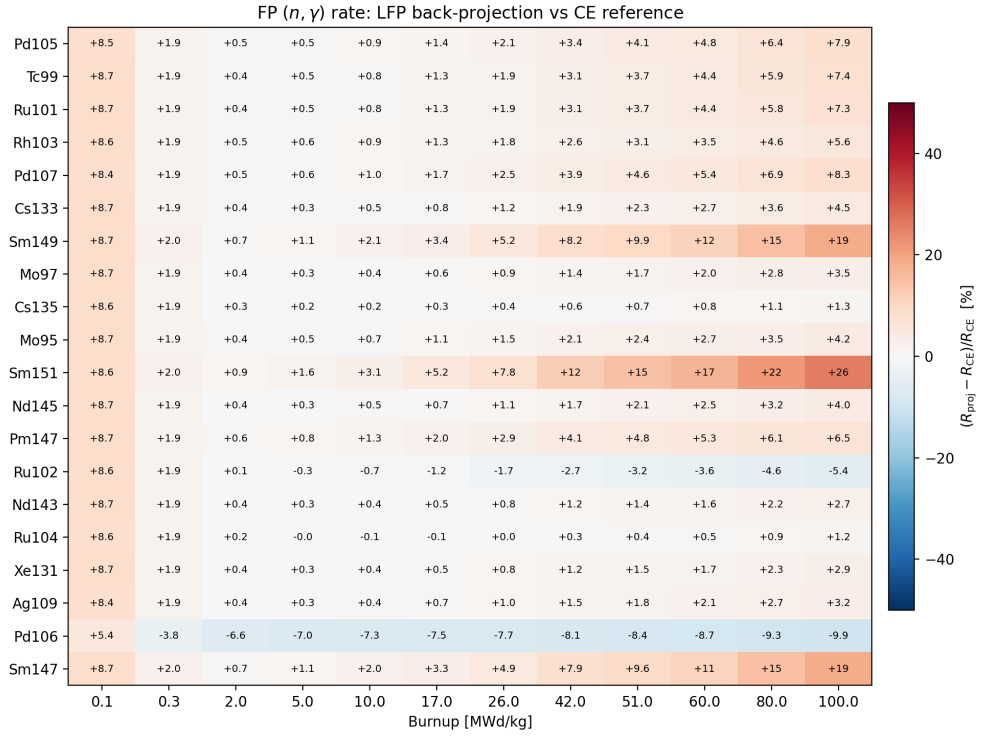

Figure 5.15: Relative deviation at selected burnup points of the back-projected FP (n, γ) reaction rate with respect to the CE reference, for the twenty nuclides with the largest capture rates.

This nuclide-wise analysis is consistent with the aggregate one-group comparison of Fig. 4.9 discussed in Sec. 4.3.7. The overestimation of strong absorbers in the reconstructed inventory contributes to the higher effective fission product capture strength of the lumped representation.

Aggregate View. The scatter plot of Fig. 5.16 provides a compact view of the agreement between the back-projected and continuous-energy concentrations at a representative mid-cycle burnup point. The figure shows the projected concentration of each fission product against the corresponding continuous-energy value on logarithmic axes. The bulk of the points lies close to the $y = x$ identity line over the dominant concentration range, indicating that the main nuclides of the inventory are reproduced accurately. The dispersion increases towards the lower-left corner, where concentrations are very small and the contribution of these nuclides to the system is correspondingly negligible.

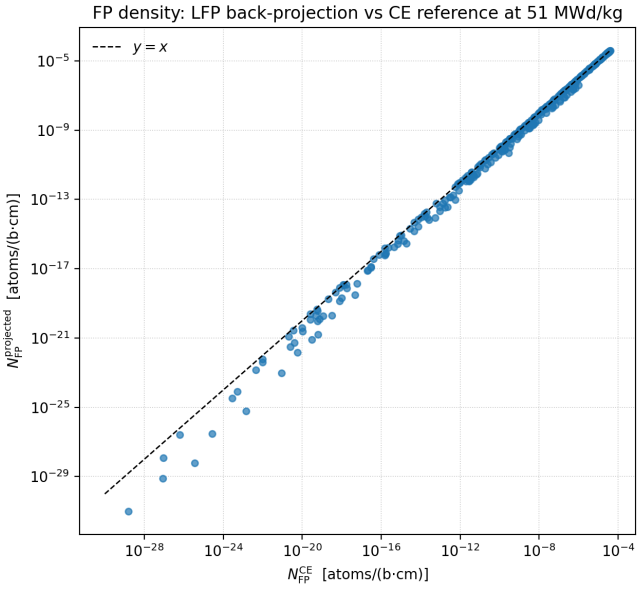

Figure 5.16: Scatter plot of the back-projected FP concentrations against the CE reference at mid-cycle (BU = 51 MWd/kgHM). The dashed line represents the identity $y = x$.

Chapter 6

Robustness and Transferability of the Methodology

The assessment carried out up to this point has shown that the lumped fission product methodology and the coupled multigroup framework can reproduce the continuous-energy reference with controlled accuracy for the case on which the libraries were originally generated. The agreement has been quantified in terms of multiplication factor, aggregate reaction rates, and individual fission product inventories reconstructed through back-projection. However, this validation was performed in a configuration that is, by construction, favourable to the multigroup representation: the same pin cell used to generate the cross sections, the same initial fuel composition, and the same burnup grid on which the libraries were tabulated.

This chapter extends the validation in two directions that progressively relax these conditions and test the range of applicability of the methodology. The first test, discussed in Sec. 6.1, retains the geometry and fuel composition of the reference case but uses a burnup grid that differs from the tabulated grid of the libraries. This exercises the linear interpolation of the multigroup cross sections during the irradiation. The second test, discussed in Sec. 6.2, replaces the reference pin cell with a different fast reactor configuration representative of the European Sodium Fast Reactor (ESFR), with different geometry and initial fuel composition, while reusing the libraries generated on the SPX reference.

6.1 Robustness to Burnup Grid Interpolation

The validation of Section 5.7 was carried out on the burnup grid used to generate the MGXS libraries. Therefore, the library manager did not need to interpolate between adjacent burnup points. To verify that the interpolation procedure described in Sec. 5.5 does not introduce an appreciable additional error, a dedicated calculation was performed on the same pin cell, with the same initial fuel composition and power history, but using

a different burnup grid. The selected grid contains ten burnup points,

$$[2.5,\ 4,\ 11,\ 33,\ 44,\ 55,\ 66,\ 77,\ 88,\ 100]\ \text{MWd/kgHM},$$

chosen so that all intermediate burnup points fall between tabulated library points. The final point coincides with the end-of-life value of the reference grid. The new grid also contains half as many points as the tabulated grid, so that the interpolation test is combined with longer depletion steps.

Multiplication factor. Figure 6.1 reports the multiplication factor of the interpolated-grid calculation together with the standard tabulated-grid calculation and the continuous-energy reference. The points of the interpolated-grid calculation, shown as green diamonds, lie close to the tabulated-grid curve. The deviation from the continuous-energy reference, shown in the lower panel, follows the same overall pattern as the tabulated-grid case discussed in Sec. 5.7.4. Within the present configuration, using an interpolated burnup grid does not produce visible additional deviations in the multiplication factor.

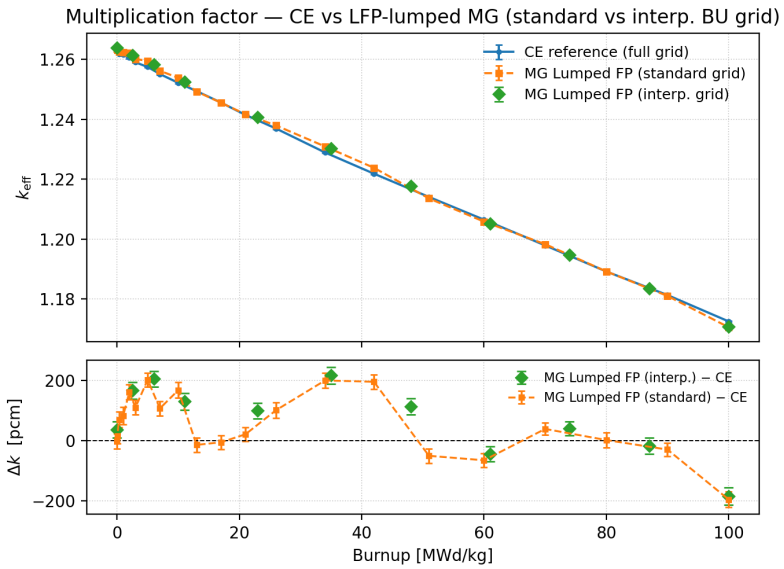

Figure 6.1: Multiplication factor along the irradiation history, computed by the lumped MG framework on the tabulated burnup grid and on the interpolated burnup grid, compared with the CE reference. The lower panel shows the deviation of each lumped calculation with respect to the CE reference.

Reaction rates. The reaction rates exhibit a behaviour consistent with the analysis on the tabulated grid. The total heavy metal (n, γ) reaction rate, not shown here, presents the same small positive deviation of about 0.1–0.2 % already discussed in Sec. 5.7.1, without additional features attributable to interpolation.

The total fission product (n, γ) reaction rate, reported in Fig. 6.2, deserves a brief comment. The first burnup point of the interpolated-grid calculation, at BU = 2.5 MWd/kgHM, lies beyond the earliest regime in which the redistribution of fission products without cross-section evaluations produced the large positive deviation discussed in Sec. 4.3.7. As a result, the interpolated-grid calculation does not show the initial spike near +8 %, but starts directly in the slow-growth regime of the second mechanism. The relative deviation is slightly negative at the first point, about −0.1 %, and then increases monotonically to about +3.5 % at end of life. The small initial negative deviation may be related to the use of longer depletion steps and to interpolation over the steepest part of the LFP cross-section evolution.

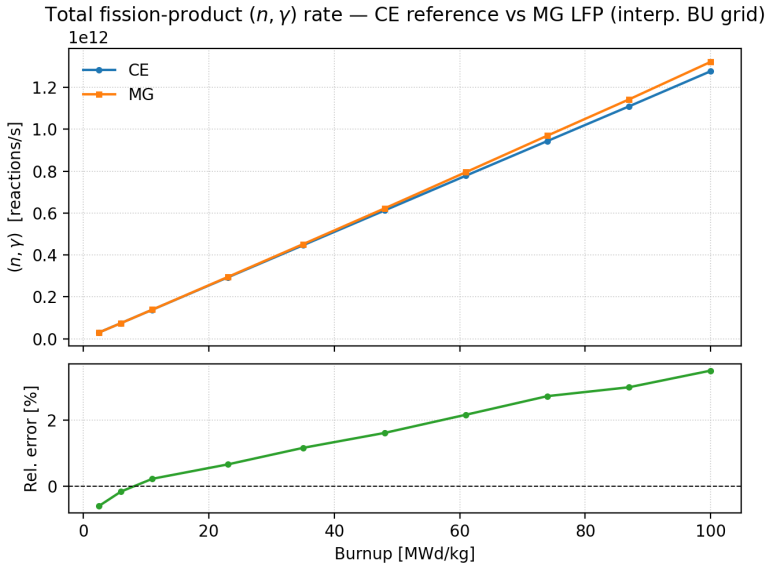

Figure 6.2: Total FP (n, γ) reaction rate on the interpolated burnup grid, compared with the CE reference. The first burnup point at 2.5 MWd/kgHM lies beyond the early redistribution regime, so the relative deviation starts directly in the slow-growth region.

The comparison of the interpolated-grid calculation with the tabulated-grid case, both for the multiplication factor and for the aggregate reaction rates, does not reveal additional deviations beyond those already present in the tabulated calculation. Within the limits of this test, where geometry, initial composition, and power history are unchanged, linear interpolation of the multigroup cross sections in burnup is adequate for evaluating the libraries at intermediate burnup points. A more demanding test, in which the physical configuration itself is changed relative to the one used for library generation, is addressed in the next section.

6.2 Transferability to a Different Fast Reactor Configuration

The second extension of the validation considers a system that differs from the reference case in both geometry and initial fuel composition, while reusing without modification the lumped fission product cross-section libraries generated for the Superphénix pin cell of Section 2.2. The objective is to assess whether libraries generated on one fast-spectrum reference remain usable on a different fast reactor configuration, and to identify the limits within which the lumped methodology continues to reproduce the continuous-energy reference with acceptable accuracy. The configuration chosen for this test is a pin cell representative of the ESFR, as defined within the Horizon 2020 ESFR-SMART project [33].

6.2.1 Reference ESFR Configuration

The ESFR pin cell is a hexagonal lattice with a central helium hole inside the fuel pellet, surrounded by helium gap, cladding, and sodium coolant. The cladding is the oxide dispersion-strengthened (ODS) steel MA956, which differs from the cladding adopted in SPX. Since some nuclides present in the ESFR structural materials were not included in the previous base-material library of non-depletable nuclides, this auxiliary library was regenerated for the ESFR composition. The LFP libraries, instead, were taken from the SPX reference without modification. The radial dimensions of the ESFR pin cell are reported in Tab. 6.1, alongside the corresponding dimensions of the reference case.

Table 6.1: Radial dimensions of the ESFR pin cell compared with the Superphénix reference. The ESFR pin cell includes a central hole inside the fuel pellet, absent in the reference case.

Region	ESFR [cm]	SPX [cm]
Inner hole radius	0.1560	—
Fuel outer radius	0.4680	0.357
Gap outer radius	0.4835	0.371
Cladding outer radius	0.5358	0.4321
Hexagonal pitch	0.6738	0.5696

The initial fuel composition is a MOX fuel with plutonium content adapted to the ESFR design specifications. The compositions compared in this work are reported in Tab. 6.2. Relative to the Superphénix reference, the ESFR fuel has a different plutonium isotopic vector, with larger contributions from ^{240}Pu and ^{242}Pu, and a smaller initial inventory of ^{241}Am.

The linear power adopted in the depletion calculations is 341 W/cm. This value is a representative estimate based on the ESFR-SMART core specifications. Starting from an assembly thermal power of approximately 8 MW for the inner fuel region, dividing it over the 271 fuel pins of the assembly, and assuming that about 13 % of the assembly power is generated in the fertile part of the fuel, the residual power per pin is distributed over an active height of 75 cm. The resulting value should be understood as a plausible operating

Table 6.2: Initial atom densities of the fuel material in the ESFR pin cell and in the Superphénix reference, in atoms/(b · cm).

Nuclide	ESFR	SPX
^{235}U	$4.888 \cdot 10^{-5}$	$9.997 \cdot 10^{-5}$
^{238}U	$1.926 \cdot 10^{-2}$	$1.964 \cdot 10^{-2}$
^{238}Pu	$1.512 \cdot 10^{-4}$	$1.761 \cdot 10^{-5}$
^{239}Pu	$1.998 \cdot 10^{-3}$	$2.428 \cdot 10^{-3}$
^{240}Pu	$1.246 \cdot 10^{-3}$	$7.239 \cdot 10^{-4}$
^{241}Pu	$3.442 \cdot 10^{-4}$	$1.946 \cdot 10^{-4}$
^{242}Pu	$4.322 \cdot 10^{-4}$	$6.755 \cdot 10^{-5}$
^{241}Am	$3.258 \cdot 10^{-5}$	$4.764 \cdot 10^{-5}$
^{16}O	$4.655 \cdot 10^{-2}$	$4.598 \cdot 10^{-2}$

point of the ESFR core rather than as an exact specification: the purpose of this analysis is to exercise the lumped methodology on a representative ESFR-like configuration, not to reproduce a specific operating scenario [34].

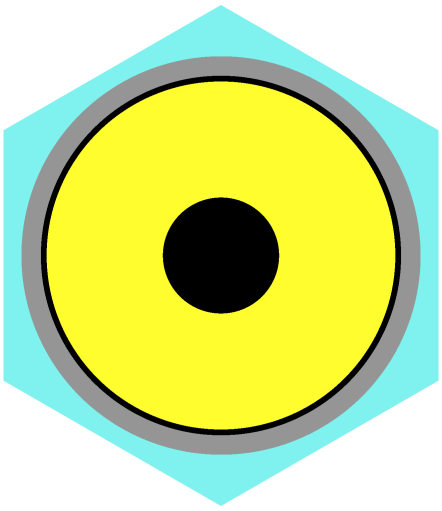

Figure 6.3: Cross section of the ESFR pin cell: the black regions correspond to the helium central hole and gap, the yellow region to the MOX fuel, the grey ring to the ODS MA956 cladding, and the surrounding cyan hexagon to the sodium coolant.

6.2.2 Spectral Compatibility

The reuse of the reference libraries on the ESFR configuration is conditioned on the two systems having sufficiently similar spectra, since the multigroup cross sections were collapsed using the spectrum of the reference case. Figure 6.4 compares the normalised fuel flux spectrum of the two configurations at end of life. The two distributions are close

over the energy range that dominates the reaction rates, with both spectra peaked around 10^4–10^5 eV and strongly suppressed below the keV region, as expected for sodium-cooled fast reactors with oxide fuel. This spectral compatibility supports the reuse of the lumped libraries: the weighting spectrum used for the cross-section collapse in the reference is also representative of the ESFR conditions, so the resulting group constants remain physically meaningful in the new configuration.

A consequence of the chosen power and of the different fuel geometry is that the absolute neutron flux in the fuel region of the ESFR pin cell is about 1.7 times higher than in the SPX reference in the present calculation. This is a property of the specific configurations compared here, not a generic feature of the two reactor designs. It is nevertheless relevant for the test: for the same relative cross-section error, absolute reaction-rate deviations scale with the flux level. The present comparison is therefore a somewhat more demanding probe of the lumped methodology than a flux-matched comparison would be.

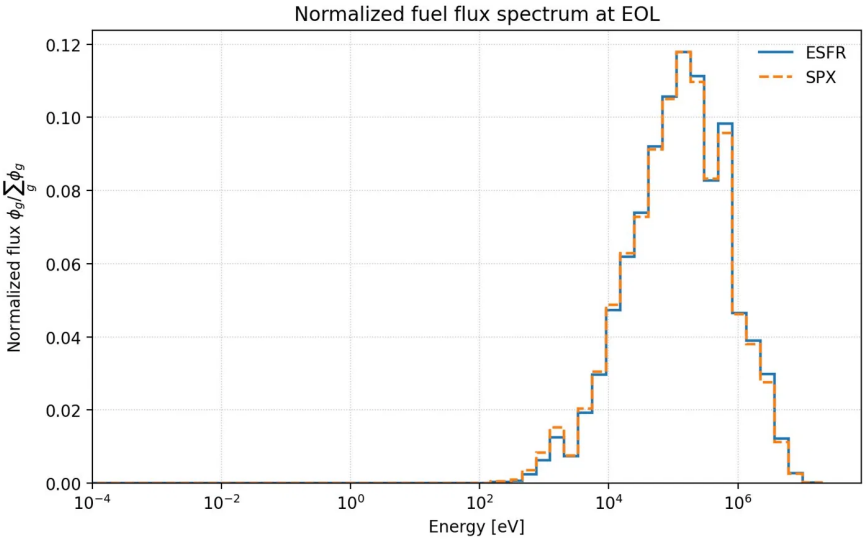

Figure 6.4: Normalised fuel flux spectrum at end of life in the ESFR pin cell and in the Superphénix reference.

6.2.3 Results on the ESFR Pin Cell

The depletion calculation was carried out on a burnup grid of thirteen points selected to include both tabulated and non-tabulated values. The run therefore tests direct library retrieval at coincident points and linear interpolation at intermediate points within the same irradiation history:

[1.0, 3.0, 7.0, 13.0, 20.0, 30.0, 40.0, 50.0, 60.0, 70.0, 80.0, 90.0, 100.0] MWd/kgHM.

The results are compared with the continuous-energy reference performed on the same ESFR pin cell and with the same power history.

Multiplication factor. Figure 6.5 reports the multiplication factor of the MG LFP calculation and of the CE reference along the irradiation history. The two curves are visually close, and the Δk in the lower panel remains within about ± 200 pcm over the full cycle, comparable to the envelope observed on the SPX reference in Sec. 5.7.4. The detailed pattern is different: a first minimum of about -220 pcm appears in the early burnup region around 13 MWd/kgHM, the difference returns close to zero in the central part of the cycle, and a second minimum of about -200 pcm is reached at end of life.

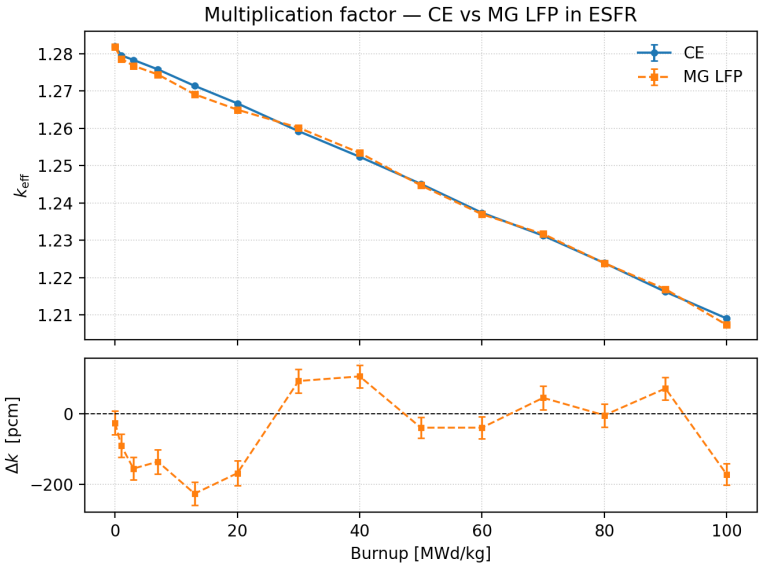

Figure 6.5: Multiplication factor along the irradiation history of the ESFR pin cell, computed by the MG LFP framework and by the CE reference. The lower panel shows the difference Δk between the two calculations.

Inventory. The total fission product concentration is reported in Fig. 6.6. The relative deviation with respect to the CE reference reaches about $-0.23\,\%$ at end of life, with a slow monotonic growth in magnitude similar to the one observed in the SPX case. The heavy metal concentration, not shown here, agrees with the reference within $0.035\,\%$ over the whole cycle, in line with the corresponding result of Sec. 5.7.1.

Heavy metal capture rate. The total (n, γ) reaction rate over the heavy metal inventory is reported in Fig. 6.7. The MG LFP calculation overestimates the CE reference for most of the cycle, with a local maximum of about $+0.65\,\%$ around 20 MWd/kgHM, followed by a slow decrease toward mid-cycle and a sign change in the last burnup steps, where the deviation becomes slightly negative ($-0.15\,\%$ at end of life). The larger amplitude of the deviation, compared with the approximately $\pm 0.15\,\%$ envelope observed on the SPX reference, is consistent with the transferability test: the heavy metal multigroup

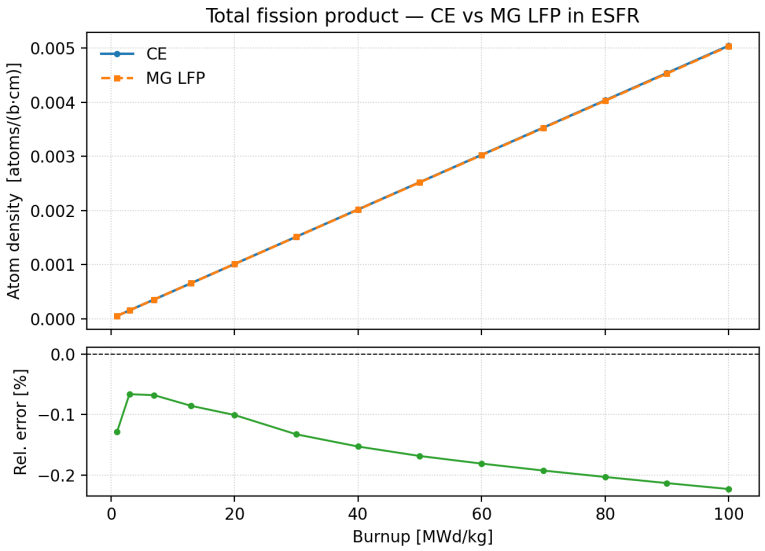

Figure 6.6: Total FP concentration along the irradiation history of the ESFR pin cell. The lower panel shows the relative deviation between the MG LFP calculation and the CE reference.

cross sections were also generated using the SPX spectrum and are applied here to a similar but not identical spectrum.

The early overprediction of the heavy metal capture rate is consistent with the slightly lower multiplication factor observed in the same burnup region. As the capture-rate deviation decreases later in the cycle, the k_{eff} deviation also moves back toward zero. This correlation should be interpreted qualitatively, since the multiplication factor is affected by several simultaneous contributions, including fission, capture, scattering, and changes in the evolving nuclide inventory.

Fission product capture rate. The total (n, γ) reaction rate over the fission product inventory is reported in Fig. 6.8. At the first burnup point, BU = 1 MWd/kgHM, the relative deviation is large and negative, about −17 %. The deviation then increases rapidly with burnup, crosses zero around 13 MWd/kgHM, and reaches +3.2 % at end of life. The end-of-life value is slightly smaller than the corresponding SPX value (+3.7 %), indicating that the slow-growth regime of the deviation is reproduced in the ESFR configuration with comparable accuracy.

The large negative deviation at the first burnup point is specific to this calculation. At 1 MWd/kgHM, the fission product inventory is still very small in absolute terms, so small absolute differences produce large relative errors. A quantitative attribution of this first-point discrepancy would require a dedicated analysis and is outside the scope of the present transferability test.

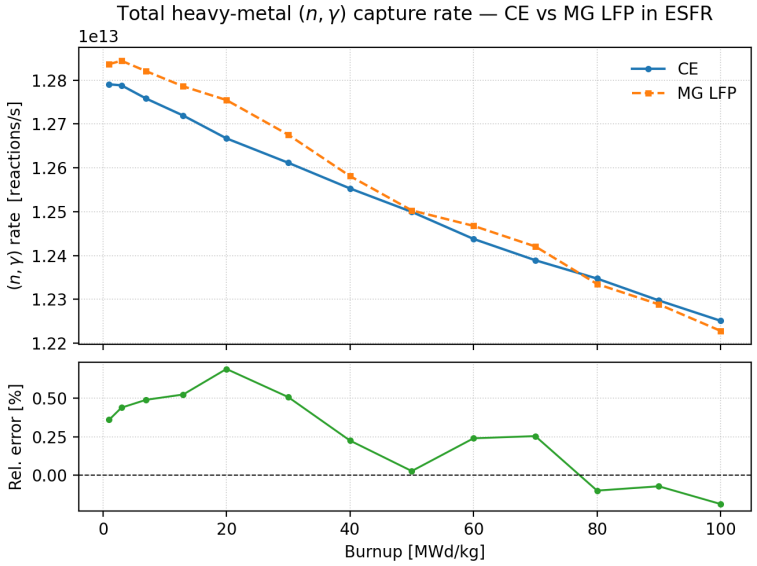

Figure 6.7: Total HM (n, γ) reaction rate along the irradiation history of the ESFR pin cell. The lower panel shows the relative deviation between the MG LFP calculation and the CE reference.

The relevance of this discrepancy for the system is better assessed through the deviation normalised to the total fuel capture rate, reported in Fig. 6.9. On this scale, the deviation at the first burnup point is negligible for the overall neutron balance of the fuel. The deviation grows monotonically through the cycle and reaches +0.47 % at end of life, slightly below the corresponding SPX value (+0.53 %). The large relative deviation observed in the raw comparison at the first burnup point therefore has no appreciable system-level effect, and the impact of the methodology on the overall capture balance of the fuel remains within about half a percent over the whole cycle.

Summary of the transferability test. Table 6.3 summarises the key end-of-life figures of the ESFR calculation alongside the corresponding values of the SPX case. The two sets of results are comparable in magnitude across the diagnostic quantities considered. The main systematic difference is the larger envelope of the heavy metal capture-rate deviation in the ESFR case, which is consistent with the spectral mismatch between the two configurations. Regenerating the heavy metal multigroup libraries specifically for the ESFR spectrum could reduce this gap, while the reuse of the LFP libraries generated on the Superphénix reference remains adequate within the scope of this test. Overall, the ESFR calculation provides an encouraging indication of the applicability of the methodology to similar fast reactor configurations, although a broader set of systems and benchmarks would be required for a general transferability claim.

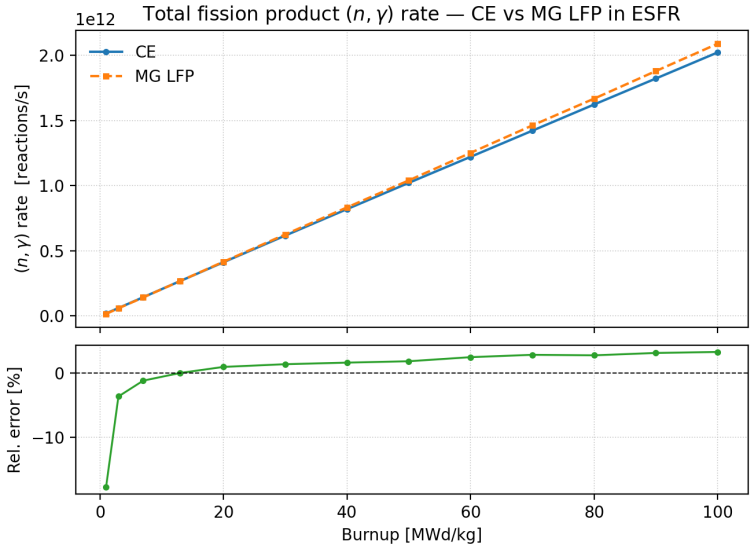

Figure 6.8: Total FP (n, γ) reaction rate along the irradiation history of the ESFR pin cell. The lower panel shows the relative deviation between the MG LFP calculation and the CE reference.

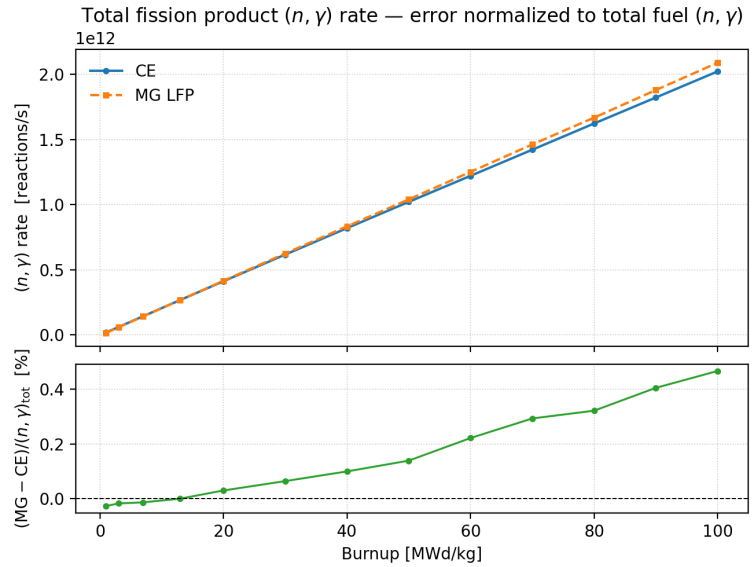

Figure 6.9: Relative deviation of the ESFR FP (n, γ) reaction rate, normalised to the total fuel capture rate. The large first-point deviation in Fig. 6.8 becomes negligible on this scale.

Table 6.3: End-of-life summary of the transferability test. Selected quantities computed on the ESFR pin cell are compared with the corresponding values obtained on the Superphénix reference. All deviations are expressed with respect to the CE reference of the corresponding configuration.

Quantity	ESFR	Superphénix
$\lvert\Delta k\rvert$ at end of life [pcm]	~ 200	~ 200
HM concentration, rel. dev. [%]	$+0.035$	$+0.034$
FP concentration, rel. dev. [%]	-0.23	-0.29
HM (n,γ), max rel. dev. [%]	$+0.65$	$+0.27$
FP (n,γ) at EOL [%]	$+3.2$	$+3.7$
FP (n,γ) at EOL, fuel-normalised [%]	$+0.47$	$+0.53$

6.3 Computational Cost of the Methodology

The accuracy assessment of the previous sections quantified the residual error introduced by the multigroup representation and by the lumping of fission products in reaction rates, inventories, and multiplication factor. The methodology developed in this work, however, was motivated by a complementary practical objective: reducing the computational cost of fast reactor depletion calculations while preserving the relevant neutronic observables within an acceptable accuracy envelope. This section quantifies the wall-clock time savings achieved on the two configurations considered in this work and discusses the upfront cost of library generation, which is amortised when the libraries are reused across multiple depletion calculations.

6.3.1 Reference Conditions for the Comparison

All calculations reported in this work were performed on the same computing infrastructure, with 128 CPU cores allocated to each run. The number of active neutron histories per transport solve was kept constant within each configuration: $3 \cdot 10^7$ active particles per solve on the SPX reference, and $7.5 \cdot 10^6$ active particles per solve on the ESFR pin cell. These values are not directly comparable across configurations; however, the relative speedup of the multigroup methodology with respect to the continuous-energy reference is an internally consistent measure of the time saving on each system.

The reference CE depletion calculations were executed with the predictor-corrector scheme of OpenMC, which requires two transport solves per depletion step. The MG calculations, both with explicit and with lumped fission products, were executed with the predictor-only scheme implemented in the framework developed in this work, which requires one transport solve per depletion step. Part of the speedup reported below therefore reflects the change of integration scheme in addition to the change in transport mode.

Each depletion calculation includes an initial transport solve on the fresh fuel, followed by the transport solves required by the depletion scheme over the burnup grid. Thus, a predictor-only calculation with n burnup steps requires $n + 1$ transport solves, whereas

Table 6.4: Wall-clock times of the depletion calculations on the Superphénix reference. All runs used 128 CPU cores and $3 \cdot 10^7$ active particles per transport solve. The speedup is computed relative to the CE reference.

Calculation	Total time [h]	Avg. time per transport solve [s]	Calc. rate [p/s]	Speedup
CE predictor–corrector	191.20	16788	1787	1.00
MG explicit FP predictor-only	7.04	1208	35274	27.2
MG LFP predictor-only	2.51	431	97787	76.2

a predictor-corrector calculation requires approximately $2n + 1$ transport solves. In the following, the total time of a depletion calculation refers to the sum of all transport solves of the irradiation history, including the fresh-fuel point.

6.3.2 Time Savings on the SPX Reference

The SPX reference was depleted over twenty burnup steps. This corresponds to twenty-one transport solves for the MG predictor-only calculations and approximately forty-one transport solves for the CE predictor-corrector. The wall-clock times of the three calculations are summarised in Tab. 6.4.

The CE reference required approximately eight days of wall-clock time on the allocated infrastructure. The MG calculation with explicit fission products completed the same irradiation history in approximately seven hours, reducing the total run time by more than a factor of twenty-seven. Introducing the LFP representation on top of the MG transport reduced the total time further, to approximately two and a half hours, yielding a cumulative speedup of about a factor of seventy-six with respect to the CE reference.

The speedup has two main components. The factor of twenty-seven between the CE reference and the MG explicit-FP calculation combines the use of groupwise pre-collapsed cross sections, instead of pointwise continuous-energy data, with the change from a predictor-corrector to a predictor-only integration scheme, which approximately halves the number of transport solves per depletion step. The additional factor of about 2.8 obtained by introducing the LFP representation is attributable to the smaller number of nuclides that need to be tallied at every transport solve and included in the depletion chain, with seventeen lumped carriers replacing the explicit fission product set.

6.3.3 Time Savings on the ESFR Pin Cell

The ESFR pin cell was depleted over thirteen burnup steps. This corresponds to fourteen transport solves for the MG LFP predictor-only calculation and approximately twenty-seven transport solves for the CE predictor-corrector reference. The wall-clock times of the two calculations performed on this configuration are summarised in Tab. 6.5.

The MG LFP calculation is faster than the CE reference on the ESFR pin cell by about a factor of one hundred and twenty-seven, of the same order as the speedup obtained on the Superphénix reference. The calculation rate of the MG LFP run is comparable across

Table 6.5: Wall-clock times of the depletion calculations on the ESFR pin cell. Both runs used 128 CPU cores and $7.5 \cdot 10^6$ active particles per transport solve. The speedup is computed relative to the CE reference.

Calculation	Total time [h]	Avg. time per transport solve [s]	Calc. rate [p/s]	Speedup
CE predictor–corrector	32.95	4393	1821	1.00
MG LFP predictor-only	0.26	66	103778	126.7

the two configurations, around 10^5 particles per second, while the calculation rate of the CE run is around $1.8 \cdot 10^3$ particles per second on both systems. This indicates that the observed speedup is consistent with the lower cost per particle history of the multigroup transport mode, combined with the reduced number of transport solves required by the predictor-only scheme.

6.3.4 Cost of the Library Generation

The speedup figures reported above refer to depletion calculations that use the LFP multigroup libraries as input. The generation of those libraries, performed once on the Superphénix reference through the workflow described in Section 2.4, has its own computational cost and must be included in a complete assessment of the methodology. Each transport solve of the library-generation workflow uses $3 \cdot 10^8$ active particles, an order of magnitude more than the depletion calculations discussed above, and tallies many reaction rates per energy group for every nuclide of the inventory, including both heavy metals and fission products. The combination of higher statistics and larger tally workload makes the calculation rate of the library-generation runs significantly lower than that of the corresponding depletion runs: about 875 particles per second on the depleted-state points, compared with about 26 000 particles per second on the fresh-fuel point where only the nine fuel nuclides are tallied. This factor of about thirty illustrates the impact of the tally workload on transport efficiency.

The total wall-clock time of the library generation, executed sequentially on 128 CPU cores, is about 83 hours for the twenty-one transport solves of the burnup grid. The fresh-fuel point requires approximately 3.5 hours, while the depleted points average approximately 4 hours each. Since the workflow consists of independent single-step transport solves, it can be parallelised across multiple jobs if the computing infrastructure allows. In the sequential case considered here, the library-generation cost is substantial, but still lower than the cost of a single full CE depletion calculation on the Superphénix reference.

Library generation is an upfront cost. Once the libraries have been produced, they can be reused in subsequent depletion calculations that remain within the physical regime for which the libraries are representative. The break-even point therefore depends on the target workload and on the reference used for comparison. For the Superphénix case, the library generation is already less expensive than one full CE depletion calculation; for shorter or lower-statistics applications, several reuse cases may be required for full

amortisation. In either case, every additional compatible depletion calculation benefits from the speedup factors quantified above. The interpolation and ESFR transferability tests presented in this chapter are examples of such reuse.

6.3.5 Closing Remarks

For the configurations considered in this work, the lumped multigroup methodology remains within a multiplication-factor envelope of about 200 pcm and a fuel-normalised fission-product capture deviation below about half a percent at end of life. The computational analysis provides the complementary measure of the cost at which this accuracy is obtained. On the Superphénix reference, the LFP methodology reproduces the CE result in about 1/76 of the wall-clock time; on the ESFR pin cell, it requires about 1/127 of the wall-clock time of the corresponding CE reference.

The library generation that enables these speedups has a non-negligible upfront cost, but this cost is amortised when the libraries are reused. The methodology is therefore most advantageous in regimes where the same libraries support multiple depletion calculations, such as parametric studies, sensitivity analyses, or the exploration of a family of configurations within a common fast-spectrum regime.

Chapter 7

Conclusions and Future Work

This thesis has developed and assessed a methodology for the generation and use of multigroup lumped fission product cross sections in fast spectrum depletion calculations. The work has addressed three connected aspects of the problem: the construction of the cross-section libraries from a continuous-energy Monte Carlo reference, the implementation of a coupled multigroup transport-depletion framework able to use those libraries, and the assessment of the methodology on multiplication factor, reaction rates and nuclide inventories, both on the configuration used for the library generation and on a different configuration. This chapter summarises the conclusions of the work and outlines the directions that the methodology naturally suggests for further development.

7.1 Conclusions

The first part of the work addressed the construction of the LFP data. Effective fission yields were generated through auxiliary depletion calculations in which heavy-metal transmutation and decay routes were disabled, so that the fission products associated with each parent could be isolated. These yields were then used to collapse explicit fission product microscopic cross sections into parent-specific lumped carriers. This approach keeps the number of fission product species small while preserving part of the dependence of the fission product composition on the fissioning heavy metal.

The second part of the work assessed the approximations introduced by the LFP representation. The stepwise analysis showed that the main discrepancies arise from two sources: the multigroup treatment itself and the replacement of the explicit fission product inventory by lumped pseudo-nuclides. In the decoupled cross-section analysis, the aggregate fission product total reaction rate was reproduced accurately, while absorption channels showed deviations of the order of a few percent.

A coupled multigroup transport–depletion framework was then implemented in Python. The framework uses OpenMC multigroup transport calculations to obtain the flux at each burnup step and advances the composition through the depletion operator using the corresponding burnup-dependent microscopic cross sections. A preliminary calculation with explicit fission products verified that the driver and library-management procedure work

correctly before applying the LFP approximation.

For the Superphénix reference case, the fully coupled MG LFP calculation reproduced the main integral quantities with acceptable accuracy for the intended model-reduction purpose. The multiplication factor remained within an envelope of about 200 pcm with respect to the CE reference over the full irradiation history. The total heavy-metal inventory and aggregate reaction rates were reproduced more accurately than most individual nuclide quantities, confirming that the lumping procedure is better suited to preserving integral effects than detailed isotope-by-isotope inventories. The back-projection analysis showed that the dominant fission product concentrations are generally well recovered, while larger deviations appear for strong absorbers such as samarium, europium, and gadolinium isotopes. This behaviour is consistent with the sensitivity of these nuclides to the spectrum and to the local balance between production and neutron capture.

The robustness test on an alternative burnup grid showed that linear interpolation of the burnup-dependent MGXS libraries does not introduce a visible additional error for the reference configuration. The transferability test on the ESFR-like pin cell further showed that the methodology remains usable when applied to a different fast-spectrum configuration, although the larger deviations observed in some heavy-metal capture rates indicate the expected sensitivity to spectral mismatch between the library-generation case and the target calculation.

The computational results confirm the practical motivation for the method. On the Superphénix reference, the MG LFP depletion calculation was about 76 times faster than the CE reference depletion, while on the ESFR pin cell the speedup was about 127. The cost of generating the MGXS libraries is not negligible, but it can be amortised when the same libraries are reused in several calculations within a similar fast-spectrum regime. The method is therefore most useful for parametric studies, sensitivity analyses, or exploratory calculations where repeated depletion simulations are required.

Overall, the results show that parent-specific LFPs combined with burnup-dependent multigroup cross sections can reproduce the main neutronic effects of the fission product inventory in fast-spectrum pin-cell depletion calculations at a much lower computational cost than CE depletion. The method is not intended to replace explicit fission product tracking when detailed isotope inventories are required. Its value lies instead in providing a reduced representation that preserves the aggregate effect of fission products on reaction rates and reactivity with controlled accuracy for fast-spectrum applications.

7.2 Future Work

The aim of the thesis is to establish and demonstrate the methodology, and represents only a starting point for future possibilities. The work presented here establishes the methodology and demonstrates its viability on the configurations considered, but it does not exhaust the possibilities it opens. Several directions of extension can be identified, of varying ambition and practical proximity to the present state of the framework.

7.2.1 Extension of the Validation Domain

The most immediate continuation of the work is the extension of the validation to a broader set of configurations. The two systems considered here are both single pin cells with oxide fuel and sodium coolant, similar in spectrum and in fuel composition. A more demanding assessment of the transferability of the libraries would consider fast reactor configurations with appreciably different spectra, such as metal-fuelled cores or cores cooled by lead or gas. A parallel direction is the extension from pin cell to assembly, where intra-assembly spectral gradients between fissile and fertile zones, between control and fuel positions, and between the core centre and its periphery may exercise the libraries in conditions farther from the spectrum on which they were generated. The long-term goal of this direction is the application of the methodology to a full-core depletion calculation, where the residual errors introduced by the lumping and by the multi-group representation can be assessed on an integral, design-relevant quantity such as the reactivity swing or the equilibrium isotopic vector.

7.2.2 Extension of the Library Parametrisation

The cross-section libraries developed in this work are parametrised on the cumulative burnup only, while the remaining state variables of the system are held fixed at the values used during the library generation. A natural extension is the introduction of the fuel temperature as a second parameter of the libraries, indexing them on a two-dimensional grid in burnup and temperature. The framework architecture already accommodates this extension with minor modifications: the manager would interpolate bilinearly between the four closest tabulated points, and the existing transport-depletion loop would remain otherwise unchanged. The library so obtained would enable the evaluation of the Doppler reactivity coefficient within the same lumped multi-group framework. A subsequent step would be the addition of further parameters relevant to fast reactor conditions, such as the coolant density or the void fraction.

If the dimensionality of the parameter space grow beyond two or three independent variables, the linear interpolation adopted in the present framework would become limiting, both in terms of the number of tabulated points required and of the fidelity of the interpolation in regions of strong non-linearity. In such a regime, the replacement of the interpolation routines with regression-based surrogates, including those built on machine-learning techniques such as neural networks or Gaussian processes, would be a natural avenue to explore. The choice of the specific technique would depend on the dimensionality of the problem and on the required level of accuracy, and would itself be the object of a dedicated study.

7.2.3 Removing the Burnup Dependence of the Libraries

A more structural reformulation of the methodology would aim at removing the burnup parametrisation of the libraries altogether. The present approach tabulates the lumped cross sections on a burnup grid because the effective yields of a single lumped fission product are unable to reproduce, on their own, the full burnup evolution of the capture rate

observed in the continuous-energy reference. The characteristic shape of this evolution, of which Fig. 4.2 provides an example, consists of a rapid initial rise followed by a plateau and a slow decrease towards end of cycle. This shape is the integrated effect of multiple competing mechanisms, including the build-up of strong absorbers, the gradual depletion of short-lived precursors, and the progressive shift in the fission product composition as the heavy metal vector itself evolves.

A possible alternative would associate two lumped carriers to every fissile heavy metal instead of one: a first carrier representing the early-stage composition of the fission products, and a second carrier representing the mid-cycle composition near the maximum of the capture rate. The effective yields and the decay constants of the two carriers would be chosen in such a way that the transitory between the two reproduces the observed burnup shape of the capture rate without requiring an external dependence on burnup. The construction would have to be carried out under the constraint that the scattering behaviour of the lumped representation is not appreciably modified relative to the single-carrier approach, so that the present scattering framework remains applicable. If successful, the resulting libraries would be independent of burnup, the interpolation procedure of the manager would be no longer needed for the fission product component, and the generation cost of the libraries, currently dominated by the per-burnup transport solves, would be substantially reduced. The viability of this construction, the optimal choice of the two compositions and the accuracy of the resulting representation are open questions that merit dedicated investigation.

7.2.4 Closing Considerations

The directions outlined above are not mutually exclusive. An extended validation on more diverse systems would qualify the applicability range of the methodology; the introduction of a temperature parametrisation would unlock the analysis of safety-relevant reactivity coefficients; and the exploration of a burnup-independent reformulation could simplify the overall workflow. The framework developed in this thesis is designed to accommodate all three extensions without major architectural revisions, and the methodology presented here can therefore serve as a basis for this follow-up work.

Code Availability

The source code developed for this work is publicly available at:

`https://github.com/FedericoPati/OpenMC_LumpedFP`

The repository contains the depletion-chain generation tools, the LFP library-generation scripts, the multigroup transport–depletion framework, and the example input files used in this thesis.

A permanent archived version, is also available at `https://doi.org/10.5281/zenodo.20670239` [35].

Appendices

Appendix A

Supplementary Data

A.1 ECCO-33 Energy Group Structure

Table A.1: ECCO-33 energy group boundaries adopted for the MGXS generation.

G	Upper [eV]	Lower [eV]	G	Upper [eV]	Lower [eV]
1	1.964×10^7	1.000×10^7	18	3.355×10^3	2.035×10^3
2	1.000×10^7	6.065×10^6	19	2.035×10^3	1.234×10^3
3	6.065×10^6	3.679×10^6	20	1.234×10^3	7.485×10^2
4	3.679×10^6	2.231×10^6	21	7.485×10^2	4.540×10^2
5	2.231×10^6	1.353×10^6	22	4.540×10^2	3.043×10^2
6	1.353×10^6	8.209×10^5	23	3.043×10^2	1.486×10^2
7	8.209×10^5	4.979×10^5	24	1.486×10^2	9.166×10^1
8	4.979×10^5	3.020×10^5	25	9.166×10^1	6.790×10^1
9	3.020×10^5	1.832×10^5	26	6.790×10^1	4.017×10^1
10	1.832×10^5	1.111×10^5	27	4.017×10^1	2.260×10^1
11	1.111×10^5	6.738×10^4	28	2.260×10^1	1.371×10^1
12	6.738×10^4	4.087×10^4	29	1.371×10^1	8.315
13	4.087×10^4	2.479×10^4	30	8.315	4.000
14	2.479×10^4	1.503×10^4	31	4.000	5.400×10^{-1}
15	1.503×10^4	9.119×10^3	32	5.400×10^{-1}	1.000×10^{-1}
16	9.119×10^3	5.531×10^3	33	1.000×10^{-1}	1.000×10^{-5}
17	5.531×10^3	3.355×10^3			

A.2 Depletion Chain Used in the LFP Model

The custom depletion chain adopted for actinide evolution and LFP production is reported in Listing A.1.

Listing A.1 Custom depletion chain used for actinide evolution and LFP generation.

```
U234    fission     2.0     LFP_U234
U234    (n,gamma)   1.0     U235

U235    fission     2.0     LFP_U235
U235    (n,2n)      1.0     U234
U235    (n,gamma)   1.0     U236

U236    fission     2.0     LFP_U236
U236    (n,2n)      1.0     U235
U236    (n,gamma)   1.0     U237

U237    fission     2.0     LFP_U238
U237    (n,2n)      1.0     U236
U237    (n,gamma)   1.0     U238
U237    beta-       1.0     Np237

U238    fission     2.0     LFP_U238
U238    (n,2n)      1.0     U237
U238    (n,gamma)   1.0     U239

U239    fission     2.0     LFP_Pu241
U239    (n,2n)      1.0     U238
U239    beta-       1.0     Np239

Np237   fission     2.0     LFP_Np237
Np237   (n,gamma)   1.0     Np238

Np238   fission     2.0     LFP_Np237
Np238   (n,2n)      1.0     Np237
Np238   (n,gamma)   1.0     Np239
Np238   beta-       1.0     Pu238

Np239   fission     2.0     LFP_Pu240
Np239   (n,2n)      1.0     Np238
Np239   (n,gamma)   1.0     Pu240
Np239   beta-       1.0     Pu239

Pu238   fission     2.0     LFP_Pu238
Pu238   (n,gamma)   1.0     Pu239
Pu238   alpha       1.0     U234

Pu239   fission     2.0     LFP_Pu239
Pu239   (n,2n)      1.0     Pu238
Pu239   (n,gamma)   1.0     Pu240
Pu239   alpha       1.0     U235

Pu240   fission     2.0     LFP_Pu240
```

```
Pu240   (n,2n)      1.0     Pu239
Pu240   (n,gamma)   1.0     Pu241
Pu240   alpha       1.0     U236

Pu241   fission     2.0     LFP_Pu241
Pu241   (n,2n)      1.0     Pu240
Pu241   (n,gamma)   1.0     Pu242
Pu241   beta-       1.0     Am241

Pu242   fission     2.0     LFP_Pu242
Pu242   (n,2n)      1.0     Pu241
Pu242   (n,gamma)   1.0     Pu243

Pu243   fission     2.0     LFP_Cm245
Pu243   (n,2n)      1.0     Pu242
Pu243   beta-       1.0     Am243

Am241   fission     2.0     LFP_Am241
Am241   (n,2n)      1.0     Pu240
Am241   (n,gamma)   0.8677  Am242
Am241   (n,gamma)   0.1323  Am242_m1
Am241   alpha       1.0     Np237

Am242   fission     2.0     LFP_Am242_m1
Am242   (n,2n)      1.0     Am241
Am242   (n,gamma)   1.0     Am243
Am242   ec/beta+    0.173   Pu242
Am242   beta-       0.827   Cm242

Am242_m1 fission    2.0     LFP_Am242_m1
Am242_m1 (n,2n)     1.0     Am241
Am242_m1 (n,gamma)  1.0     Am243
Am242_m1 IT         1.0     Am242

Am243   fission     2.0     LFP_Am243
Am243   (n,2n)      1.0     Am242
Am243   (n,gamma)   1.0     Cm244
Am243   alpha       1.0     Np239

Cm242   fission     2.0     LFP_Cm242
Cm242   (n,gamma)   1.0     Cm243
Cm242   alpha       1.0     Pu238

Cm243   fission     2.0     LFP_Cm243
Cm243   (n,2n)      1.0     Cm242
Cm243   (n,gamma)   1.0     Cm244
Cm243   alpha       0.9971  Pu239
Cm243   ec/beta+    0.0029  Am243

```

```
Cm244   fission     2.0     LFP_Cm244
Cm244   (n,2n)      1.0     Cm243
Cm244   (n,gamma)   1.0     Cm245
Cm244   alpha       1.0     Pu240

Cm245   fission     2.0     LFP_Cm245
Cm245   (n,2n)      1.0     Cm244
Cm245   alpha       1.0     Pu241
```

Bibliography

[1] Avneet Sood et al. «Neutronics Calculation Advances at Los Alamos: Manhattan Project to Monte Carlo». In: *Nuclear Technology* 207.sup1 (2021), S100–S133. DOI: 10.1080/00295450.2021.1956255.

[2] B G Carlson. *SOLUTION OF THE TRANSPORT EQUATION BY SN APPROXIMATIONS*. Tech. rep. Los Alamos Scientific Lab., N. Mex., Feb. 1955. DOI: 10.2172/4376236. URL: https://www.osti.gov/biblio/4376236.

[3] J. E. Morel. *Deterministic Transport Methods and Codes at Los Alamos*. Senda Editorial S.A., Sept. 1999, p. 25–34.

[4] E.E. Lewis and I. Dilber. «Finite element, nodal and response matrix methods: A variational synthesis for neutron transport». In: *Progress in Nuclear Energy* 18.1 (1986). Finite Element and Allied Methods, pp. 63–74. ISSN: 0149-1970. DOI: https://doi.org/10.1016/0149-1970(86)90013-2. URL: https://www.sciencedirect.com/science/article/pii/0149197086900132.

[5] Nicholas Metropolis and S. Ulam. «The Monte Carlo Method». In: *Journal of the American Statistical Association* 44.247 (1949). PMID: 18139350, pp. 335–341. DOI: 10.1080/01621459.1949.10483310.

[6] F. Brown et al. «Monte Carlo criticality calculations - History recent progress, invited». In: *Transactions of the American Nuclear Society* 109 (Jan. 2013), pp. 659–662.

[7] John Wagner et al. «Review of Hybrid (Deterministic/Monte Carlo) Radiation Transport Methods, Codes, and Applications at Oak Ridge National Laboratory *». In: *Progress in Nuclear Science and Technology* 2 (Oct. 2011), p. 808. DOI: 10.15669/pnst.2.808.

[8] William Boyd et al. «Multigroup Cross-Section Generation with the OpenMC Monte Carlo Particle Transport Code». In: *Nuclear Technology* 205.7 (Mar. 2019). ISSN: ISSN 0029-5450. DOI: 10.1080/00295450.2019.1571828. URL: https://www.osti.gov/biblio/1559869.

[9] Guo, Hui et al. «Generating multi-group cross-sections using continuous-energy Monte Carlo method for fast reactor analysis». In: *EPJ Web Conf.* 302 (2024), p. 02003. DOI: 10.1051/epjconf/202430202003. URL: https://doi.org/10.1051/epjconf/202430202003.

[10] Jaakko Leppänen and Riku Mattila. «Study on computational performance in generation of cross sections for nodal simulators using continuous-energy Monte Carlo calculations». In: *Journal of Nuclear Science and Technology* 52.7-8 (2015), pp. 945–952. DOI: 10.1080/00223131.2015.1027516. eprint: https://doi.org/10.1080/00223131.2015.1027516. URL: https://doi.org/10.1080/00223131.2015.1027516.

[11] Kenichi YOSHIOKA et al. «Multi-Group Constants Generation System for 3D-Core Simulation Using a Continuous Energy Monte Carlo Technique». In: *Progress in Nuclear Science and Technology* 2 (Oct. 2011), pp. 334–340. DOI: 10.15669/pnst.2.334.

[12] VTT Technical Research Centre of Finland. *Serpent 2 User Guide: Burnup and Activation Calculations*. https://serpent.vtt.fi/docs/user_guide/burnup_calculation.html. Accessed: 2026-06-04.

[13] William A. Wieselquist. «The SCALE 6.2 ORIGEN API for High Performance Depletion». In: 2015. URL: https://api.semanticscholar.org/CorpusID:137912479.

[14] Paul K. Romano et al. «OpenMC: A state-of-the-art Monte Carlo code for research and development». In: *Annals of Nuclear Energy* 82 (2015). Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013, SNA + MC 2013. Pluri- and Trans-disciplinarity, Towards New Modeling and Numerical Simulation Paradigms, pp. 90–97. ISSN: 0306-4549. DOI: https://doi.org/10.1016/j.anucene.2014.07.048. URL: https://www.sciencedirect.com/science/article/pii/S030645491400379X.

[15] J. D. Garrison and B. W. Roos. «Fission-Product Capture Cross Sections*». In: *Nuclear Science and Engineering* 12.1 (1962), pp. 115–134. DOI: 10.13182/NSE62-A25379.

[16] J. R. Liaw and H. Henryson II. «Lumped Fission Product Neutron Cross Sections Based on ENDF/B-V for Fast Reactor Analysis». In: *Nuclear Science and Engineering* 84.4 (1983), pp. 324–336. DOI: 10.13182/NSE83-A15453.

[17] B. Atefi. «A Two-Lump Fission-Product Model for Fast-Reactor Analysis». In: *Transactions of the American Nuclear Society*. Vol. 38. Upton, New York, 1981, p. 659.

[18] A. Trkov and D.L. Aldama. «Definition of pseudo fission product data for reactor calculations». In: (1999), p. 57–62.

[19] Nelson Hanan. «The use of WIMS-ANL lumped fission product cross sections for burned core analysis with the MCNP Monte Carlo code». In: (Jan. 1998).

[20] Hocine Benkharfia, Tahar Zidi, and Mohamed Belgaid. «Lumped pseudo fission products during burnup step in MCNP5-ORIGEN coupling system». In: *Progress in Nuclear Energy* 88 (2016), pp. 277–284. ISSN: 0149-1970. DOI: https://doi.org/10.1016/j.pnucene.2016.01.006. URL: https://www.sciencedirect.com/science/article/pii/S0149197016300063.

[21] P. K. Romano and B. Forget. «The OpenMC Monte Carlo Particle Transport Code». In: *Annals of Nuclear Energy* 51 (2013), pp. 274–281.

[22] Alexander Ponomarev et al. «Superphénix Benchmark Part I: Results of Static Neutronics». In: *Journal of Nuclear Engineering and Radiation Science* 8.1 (Oct. 2021), p. 011320. DOI: 10.1115/1.4051449. URL: https://doi.org/10.1115/1.4051449.

[23] P. K. Romano. *OpenMC Documentation — Cross Section Representations*. https://docs.openmc.org/en/stable/methods/cross_sections.html. Accessed 2026.

[24] G. Rimpault et al. «The ERANOS Code and Data System for Fast Reactor Neutronic Analyses». In: (2002).

[25] *OpenMC Documentation — Tallies*. https://docs.openmc.org/en/stable/methods/tallies.html. Accessed 2026.

[26] International Atomic Energy Agency. *Compilation and Evaluation of Fission Yield Nuclear Data*. IAEA-TECDOC-1168. Vienna, 2000.

[27] R. W. Mills. «Fission Product Yield Evaluation». PhD thesis. University of Birmingham, 1995.

[28] A. A. Sonzogni, E. A. McCutchan, et al. *Revision of Fission Yields Uncertainties in ENDF/B-VIII.0*. Tech. rep. Brookhaven National Laboratory, 2020.

[29] OpenMC. *Multigroup Mode Part II: MGXS Library Generation with OpenMC*. https://nbviewer.org/github/openmc-dev/openmc-notebooks/blob/main/mg-mode-part-ii.ipynb. [Accessed 02-05-2026].

[30] OpenMC. *Multigroup Mode Part III: Advanced Feature Showcase*. https://nbviewer.org/github/openmc-dev/openmc-notebooks/blob/main/mg-mode-part-iii.ipynb. [Accessed 02-05-2026].

[31] Paul Romano. *Normalizing Tally to get Flux value — openmc.discourse.group*. https://openmc.discourse.group/t/normalizing-tally-to-get-flux-value/99/17. [Accessed 11-05-2026].

[32] *8. Specifying Tallies ; OpenMC Documentation — docs.openmc.org*. https://docs.openmc.org/en/stable/usersguide/tallies.html. [Accessed 11-05-2026].

[33] Andrei Rineiski et al. *Specification of the New Core Safety Measures*. Tech. rep. D1.2. ESFR-SMART Deliverable D1.2, Horizon 2020 Grant Agreement No. 754501. ESFR-SMART Project, July 2018.

[34] Emil Fridman et al. «Neutronic Analysis of the European Sodium Fast Reactor: Part II—Burnup Results». In: *Journal of Nuclear Engineering and Radiation Science* 8.1 (2022), p. 011302. DOI: 10.1115/1.4048765.

[35] Federico Pati. $OpenMC_LumpedFP : Pythonscriptandlibrariesuesdinthethesis$. 2026. DOI: 10.5281/zenodo.20670239. URL: https://doi.org/10.5281/zenodo.20670239.