



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

DEPARTMENT OF PHYSICS AND ASTRONOMY "A. RIGHI"

SECOND CYCLE DEGREE

PHYSICS

Development of a neutron source for a Monte Carlo-deterministic hybrid analysis of an accelerator-driven subcritical nuclear reactor (ADS)

Supervisor
Prof. Paolo Finelli

Defended by
Lorenzo Zavoli

Co-supervisors
Dr. Patrizio Console Camprini
Dr. Massimo Sarotto
Dr. Roberto Pergreffi

Graduation Session September 2026

Academic Year 2025/2026

*Ai miei genitori,
grazie.*

Abstract

Accelerator-driven systems (ADS) represent a promising technology for transmuting minor actinides and improving nuclear fuel-cycle sustainability. Their subcritical operation is sustained by an external neutron source generated by a high-energy proton beam impinging on a spallation target.

The neutronic analysis of an ADS requires the description of both the high-energy spallation process and the subsequent neutron transport in the subcritical core. MCNP accurately models high-energy spallation, but cannot perform burnup calculations in fixed-source mode. ERANOS, instead, enables deterministic fuel-cycle and depletion analyses, but its available multi-group cross-section libraries are limited to energies below 20 MeV. A coupling procedure is therefore required to transfer the neutron source from the Monte Carlo calculation to the deterministic model.

In this work, a hybrid Monte Carlo-deterministic methodology was developed for the analysis of a lead-cooled subcritical ADS driven by an 800 MeV proton beam impinging on a liquid metal target. MCNP was employed to simulate the spallation process, and a dedicated algorithm was developed to identify neutrons crossing the 20 MeV energy threshold. The extracted neutron events were then processed to generate the neutron source terms required by ERANOS for the deterministic calculation.

The methodology was validated by comparing the ADS neutronic parameters obtained using the hybrid MCNP-ERANOS approach with dedicated full-core MCNP simulations. The effective multiplication factor differed by 102 pcm, while the average number of neutrons produced per fission differed by 0.06%. The total neutron yield per incident proton and the required proton beam current showed relative differences of 3.54% and 3.29%, respectively.

The results demonstrate the reliability of the proposed methodology for hybrid MCNP-ERANOS neutronic analyses of ADS.

Contents

Introduction	1
1 Nuclear Fuel Cycle	3
1.1 Nuclear Fuel Cycle and Waste Management	3
1.2 Fast Nuclear Reactors and Accelerator-Driven Systems	8
1.2.1 An Example of ADS Design: The MYRRHA Project	9
2 ADS: Architecture and Core Physics	11
2.1 ADS Components	11
2.1.1 Accelerator	11
2.1.2 Spallation Target	13
2.1.3 Subcritical Core	15
2.2 Physics of a Subcritical Core in Stationary Regime	15
2.2.1 Transport Equation	15
2.2.2 Reactivity of the Subcritical Core	17
2.2.3 Source Importance Parameter and Proton Beam Current	19
3 Numerical Methods for Particle Transport	21
3.1 ERANOS Deterministic Code System	22
3.1.1 Variable Discretisation	22
3.1.2 Multi-Group Energy Formalism	23
3.1.3 ERANOS Computational Scheme	24
3.2 MCNP Stochastic Code	25
3.2.1 Geometry	25
3.2.2 Transport	26
3.2.3 Tallies	27
3.2.4 PTRAC	30
3.2.5 Spatial Meshing and Tally Mapping	31
4 Description of the ADS Model	35
4.1 ADS Overview and Computational Models	35
4.1.1 Description of the ADS Core	35
4.1.2 ERANOS Model	36
4.1.3 MCNP Model	38
4.1.4 Target Spatial Meshing	41
4.2 Preliminary Analyses	42
4.2.1 Flux Distribution	42

4.2.2	Proton Beam Characterisation	45
5	Neutron Source Development and Results	51
5.1	ERANOS Requirements	51
5.2	Principles of the Source Extraction Algorithm	52
5.3	Algorithm Implementation	55
5.4	Parallel Implementation	56
5.5	Source Term Generation	57
5.6	Results	61
	Conclusions	65
A	CRESCO HPC Infrastructure	67
A.1	Computational Resources and Data Storage	67
A.2	Parallelisation Strategy: OpenMP and MPI	67
	Acknowledgements	69
	List of Acronyms and Abbreviations	71
	References	75

Introduction

The long-term management of spent nuclear fuel represents one of the main challenges associated with the large-scale deployment of nuclear energy. Although conventional thermal reactors have provided a reliable source of low-carbon electricity for several decades, their operation produces radioactive waste with a radiotoxicity that extends over thousands of years. In particular, minor actinides such as neptunium, americium, and curium contribute significantly to the long-term radiotoxicity of the waste.

The development of advanced nuclear systems capable of closing the fuel cycle and transmuting these long-lived isotopes has therefore become essential for the long-term sustainability of nuclear energy. Fast neutron reactors are particularly suitable for this purpose because the fast neutron spectrum increases the probability of fission relative to neutron capture for many transuranic isotopes. However, loading large quantities of minor actinides into a conventional critical reactor may significantly reduce the delayed neutron fraction and consequently introduce important safety issues.

Aiming to close the nuclear fuel cycle, accelerator-driven systems (ADS) provide a potential alternative approach for the transmutation of minor actinides. An ADS consists of a subcritical reactor core coupled to a high-energy proton accelerator and a heavy metal spallation target. The proton beam generates an external neutron source through spallation, which compensates for the neutron deficit of the subcritical core and sustains the fission chain reaction. Since the core does not rely on criticality to maintain operation, such systems could potentially accommodate fuel compositions with high concentrations of minor actinides while providing an additional intrinsic safety margin.

The neutronic analysis of an ADS requires the treatment of both the high-energy spallation process and the subsequent neutron transport in the subcritical core. The MCNP stochastic code provides an accurate description of the spallation process and the resulting particle cascade. However, MCNP cannot perform burnup calculations in fixed-source mode, limiting its use for neutronic analysis of an ADS. Deterministic reactor physics codes such as ERANOS offer instead significant computational efficiency for complex burnup analyses, but the multi-group cross-section libraries available for these calculations are generally limited to energies up to about 20 MeV and therefore cannot describe the high-energy spallation process.

A hybrid Monte Carlo-deterministic approach can therefore combine the complementary capabilities of these methods. In such an approach, the high-energy spallation process is simulated using Monte Carlo transport, while the resulting neutron source is transferred to the deterministic model for subsequent calculations below this energy limit. The accurate definition of the interface neutron source is

therefore essential to couple the two computational methods.

The main objective of this thesis is to develop a methodology for extracting the neutron source generated by the MCNP simulation and providing it as input for the ERANOS model. The methodology is applied to a lead-cooled subcritical ADS driven by an 800 MeV proton beam impinging on a liquid metal target.

A dedicated spatial meshing strategy is first implemented in the MCNP model to reproduce the spatial discretisation used by the deterministic model, enabling the mapping of neutron flux onto the same grid. Preliminary simulations are then performed to characterise the neutron flux distribution and the behaviour of the primary proton beam.

The hybrid MCNP-ERANOS methodology is subsequently implemented through the development of a dedicated post-processing algorithm for the MCNP PTRAC output. The algorithm tracks individual particles within the high-energy domain, reconstructs the particle genealogy from the sequence of transport events, and identifies neutrons crossing the 20 MeV energy threshold. The extracted events are then processed to determine the neutron yield, spatial distribution, and energy spectrum required to define the external source for ERANOS.

The proposed methodology is finally validated by comparing the main macroscopic ADS parameters obtained using the ERANOS model with dedicated full-core MCNP calculations.

The thesis is organised as follows. Chapter 1 introduces the nuclear fuel cycle, the problem of long-lived radioactive waste, and the potential role of accelerator-driven systems in the transmutation of minor actinides. Chapter 2 describes the main components of an ADS and presents the underlying physics governing source-driven subcritical cores. Chapter 3 introduces the numerical methods for particle transport, with particular emphasis on the ERANOS and MCNP code systems. Chapter 4 describes the investigated ADS, the computational infrastructure, the numerical models, and the preliminary analyses performed using MCNP. Finally, Chapter 5 presents the implementation of the neutron source extraction algorithm, the generation of the source terms required by ERANOS, and the comparison between deterministic and Monte Carlo results.

Chapter 1

Nuclear Fuel Cycle

At present, the global nuclear energy landscape is dominated by thermal neutron reactors. Despite their widespread deployment, a critical challenge associated with their operation is the accumulation of long-lived radioactive waste, which has significant long-term environmental and safety implications. Spent nuclear fuel (SNF) discharged from these facilities comprises a complex heterogeneous mixture of fission products and actinides, including the minor actinides (MAs), primarily neptunium, americium, and curium. Many of these radionuclides are characterised by half-lives ranging from millennia to hundreds of thousands of years.

The prevailing strategy for the long-term management of SNF is geological disposal. In this context, SNF is classified as high-level radioactive waste (HLW) due to its high radioactivity and decay heat, and requires secure containment over thousands of years.

Furthermore, the sustainability of the nuclear fuel cycle is threatened by non-proliferation concerns, primarily arising from the generation of fissile isotopes, such as ^{239}Pu through neutron capture in ^{238}U .

A significant reduction in the amount of HLW is crucial to establishing nuclear power as a sustainable, low-carbon energy source for the future. Towards the realisation of a truly sustainable nuclear infrastructure, Generation-IV fast-spectrum systems are currently under development.

Furthermore, accelerator-driven systems (ADS) offer a distinct and highly specialised approach. ADS facilities can effectively incinerate these long-lived isotopes through fast neutron fission while simultaneously generating electricity. Crucially, because they operate in a subcritical regime driven by an external source, ADS do not rely on delayed neutrons for kinetic stability, unlike conventional critical reactors. This physical feature allows them to operate as dedicated transmuters with significant amounts of MAs in the subcritical core.

1.1 Nuclear Fuel Cycle and Waste Management

Nuclear power plants generate energy through the fission of ^{235}U initially present in the fuel and through the fission of ^{239}Pu bred during reactor operation. After approximately 3 to 4 years of irradiation, the fuel can no longer sustain a critical chain reaction due to the progressive depletion of fissile isotopes and the accumulation




Figure 1.1: Reaction pathways for the generation of ^{239}Pu and principal MAs in standard nuclear fuel via successive neutron capture and β^- decay.

of neutron-absorbing nuclei. At this point, the assembly must be discharged from the reactor and becomes SNF.

Along with unburned uranium and plutonium, SNF also contains a variety of radioactive nuclides, specifically fission products (FPs) and MAs. FPs are the lighter fragments resulting from the direct fission of heavy nuclei. They primarily emit beta and gamma radiation, generating the vast majority of the decay heat over the short to medium term (up to roughly 300 years). Conversely, MAs are transuranic elements produced through successive neutron-capture reactions and beta decays. They are predominantly alpha emitters, characterised by extremely long half-lives, in some cases extending to hundreds of thousands of years. The MA inventory is dominated primarily by three elements: americium (Am), neptunium (Np), and curium (Cm). After a 10-year cooling period, the MA fraction in typical spent uranium oxide (UOX) fuel consists of approximately 55% Am, 40% Np, and 5% Cm (although these proportions may vary depending on the specific operating conditions and the cooling time). The generation pathways of ^{239}Pu and key minor actinides within the nuclear fuel are illustrated in Fig. 1.1. The presence of these long-lived isotopes in the SNF significantly contributes to its radiotoxicity, posing substantial challenges for waste management and safe geological disposal.

The management of SNF and its associated long-lived isotopes is a critical aspect of the nuclear fuel cycle, and different waste-management strategies can be adopted. A common approach is the Once-Through Fuel Cycle (OFC), which involves the direct disposal of SNF in geological repositories. However, this strategy does not mitigate the long-term radiotoxicity of the waste, which requires secure containment for hundreds of thousands of years. Furthermore, the OFC raises sustainability and non-proliferation concerns, as SNF contains significant quantities of fissile plutonium isotopes, such as ^{239}Pu , which could theoretically be diverted for nuclear weapons production. Finally, as shown in Fig. 1.2, almost 96% of SNF consists of uranium and plutonium [1]. Based on electricity production in France using pressurised water reactors (PWRs), the efficiency of this fuel cycle is approximately 0.6-0.7% of the




Figure 1.2: Isotopic composition of UOX SNF discharged from a PWR after 4 years of irradiation at a burnup¹ of 47.5 GWd/t_{HM}. The inventory is predominantly composed of unburned uranium and plutonium, highlighting its residual available energy [1].

initial energy content of the fuel [1], meaning that more than 99% of the potentially recoverable energy is not exploited.

To address these challenges, which are among the primary causes of public scepticism and opposition to nuclear energy, the alternative strategy of reprocessing and recycling SNF has been extensively pursued. In particular, the approach based on uranium and plutonium recycling is known as Partially Closed Fuel Cycle (PCFC). This strategy involves the chemical reprocessing of SNF to recover unburned fissile and fertile materials, which are subsequently refabricated into fresh fuel. Currently, PUREX (Plutonium Uranium Redox EXtraction) is the most widely used reprocessing method. It involves dissolving the SNF in nitric acid and subsequently extracting uranium and plutonium through solvent extraction. The recovered uranium is then re-enriched for fuel fabrication, whereas the separated plutonium can be mixed with depleted uranium to produce mixed oxide (MOX) fuel. Unlike conventional UOX fuel, which relies on enriched uranium to increase the ^{235}U content, MOX fuel uses ^{239}Pu as the primary fissile component. Fuel recycling can save 15-20% of natural uranium and reduce the waste volume by a factor of five [1].

Although this approach reduces both the volume of HLW and its fissile content, it does not significantly reduce long-term radiotoxicity due to the persistence of MAs. Furthermore, the separation of pure plutonium in this process still poses proliferation risks. Consequently, while PUREX remains the commercial standard, current research has led to the development of alternative processing techniques in which the plutonium stream is not separated.

To overcome these limitations, the next step in the evolution of the nuclear fuel cycle is the development of a fully Closed Fuel Cycle (CFC). The primary objective of the CFC is not only to enable the multiple recycling of unburned plutonium, but also to partition and transmute MAs.

Transmutation entails the conversion of long-lived radioactive isotopes into shorter-lived or stable nuclides through nuclear reactions. For this purpose, fast neutron

¹Burnup represents the total thermal energy extracted per unit mass of initial heavy metal, typically expressed in GWd/t_{HM}. Alternatively, it can be expressed as a percentage, %FIMA (Fissions per Initial Metal Atom), representing the fraction of heavy metal atoms that have undergone fission.




Figure 1.3: Evaluated neutron cross-sections of ^{241}Am as a function of incident neutron energy (10 meV to 20 MeV). As depicted, the dominant interaction mechanism shifts from radiative capture (red curve) in the thermal region to fission (purple curve) within the fast spectrum [2].

reactors (FNRs) have been proposed as suitable candidates for transmutation. Unlike thermal neutron reactors, fast reactors operate with an unmoderated neutron spectrum, which significantly enhances the probability of fission in MAs. This behaviour is illustrated by the evaluated neutron cross-sections of ^{241}Am presented in Fig. 1.3. Within the thermal energy range typical of conventional power plants, the capture cross-section (red curve) predominates. Consequently, this capture mechanism generates heavier actinides, thereby increasing the radiotoxicity of the waste and reducing reactivity through the absorption of thermal neutrons available to sustain the fission chain reaction. Conversely, in the fast neutron spectrum (energies exceeding 1 MeV), the fission cross-section (purple curve) becomes the dominant interaction mechanism, enabling the effective destruction of MAs. Indeed, nuclear fission drives transmutation, as it transforms long-lived, highly radiotoxic actinides into short-lived, less radiotoxic fission products.

In this scenario, only fission products and residual amounts of unrecovered actinides would require geologic disposal. Fig. 1.4 illustrates the impact of this approach on the long-term radiotoxicity of HLW by showing the temporal evolution of SNF radiotoxicity relative to the reference level of natural uranium ore. The OFC




Figure 1.4: Temporal evolution of SNF radiotoxicity normalised to the reference level of natural uranium ore (yellow dotted line). The plot compares the long-term waste radiotoxicity of the OFC (blue) against U and Pu recycling (orange) and the comprehensive partitioning and transmutation of all MAs (grey) [3].

is represented by the blue curve, demonstrating that the radiotoxicity remains orders of magnitude above the reference level for hundreds of thousands of years. The orange curve highlights the effect of uranium and plutonium via MOX fuel, which reduces the radiotoxicity but still leaves a significant long-term burden due to MAs. The grey curve demonstrates the potential impact of partitioning and transmutation of MAs. This strategy drastically accelerates the decay of radiotoxicity, reducing it to levels comparable to those of natural uranium ore within a few hundred years. Therefore, by reducing geologic repository inventories only to the FPs and the residues of the separation processes, the radiological impact of HLW could be minimised.

In addition to MA transmutation and radiotoxicity reduction, FNRs are capable of extracting a much larger fraction of the energy available from natural uranium. The fast neutron spectrum provides neutrons with sufficient energy to induce direct fast fission of the ^{238}U isotope, which constitutes almost 93% of SNF (see Fig. 1.2). Furthermore, ^{238}U also undergoes neutron capture, leading to the generation of various plutonium isotopes. Unlike in thermal systems, a fast neutron spectrum allows these plutonium isotopes to undergo fission more effectively. For example, the ratio of fission-to-capture cross-sections for ^{238}Pu , ^{240}Pu and ^{242}Pu increases by

factors of 22, 250 and 36, respectively, when compared to a thermal spectrum [1]. By effectively consuming all Pu isotopes rather than allowing them to accumulate as parasitic absorbers, the isotopic degradation of the fuel is prevented. Consequently, fast reactors could enable fuel multi-recycling. This approach could exploit more than 80% of natural uranium resources, corresponding to an efficiency two orders of magnitude greater than that of the OFC, while reducing the waste volume by a factor of six compared with mono-recycling.

1.2 Fast Nuclear Reactors and Accelerator-Driven Systems

As discussed above, the transmutation of MAs within a CFC is a critical step towards achieving a sustainable nuclear fuel cycle, and FNRs have been identified as suitable candidates for this purpose. However, the implementation of FNRs for MA transmutation presents significant challenges.

While FNRs provide the necessary fast neutron spectrum for efficient transmutation, integrating large quantities of these isotopes into a conventional critical reactor core presents many safety concerns. In critical reactors, the chain reaction is self-sustaining ($k_{\text{eff}} = 1$) because of the overall neutron balance, while its safe operation relies on the presence of delayed neutrons. These are emitted seconds after the initial fission due to the decay of specific fission products. The delayed neutron fraction (β_{eff}) is a crucial parameter for reactor control and safety, as it provides a manageable response time to changes in reactivity, thereby preventing rapid power excursions and potential chain escalations. For a typical fast neutron system, the β_{eff} corresponding to ^{235}U is approximately 0.65%.

However, the use of fertile-free fuel loaded with MAs (or even Pu) causes a critical drop in the core's β_{eff} . This fundamental limitation arises because the macroscopic β_{eff} is directly governed by the absolute delayed neutron yield (ν_d) of the fuel's constituent isotopes, which is lower for MAs and Pu compared to uranium isotopes. For instance, ^{235}U exhibits a ν_d that is approximately 3 times higher than that of ^{241}Am and 2.5 times higher than that of ^{239}Pu [4, 5].

Therefore, a safe transmutation strategy based on critical FNRs requires limited MA concentrations in the fuel, which implies a reduced transmutation efficiency of the fuel cycle and a longer time to achieve the desired reduction in long-term radiotoxicity. To overcome these limitations, accelerator-driven systems (ADS) have been proposed as a promising alternative.

An ADS couples a subcritical reactor core ($k_{\text{eff}} < 1$) with a high-energy proton accelerator. The subcriticality of the core implies that the fission chain reaction cannot sustain itself independently. This neutron deficit is supplied by an external neutron source, generated via spallation of an accelerated proton beam striking a heavy metal target located at the centre of the core. By operating in a subcritical regime, the reactor's control margin does not rely on the β_{eff} of the fuel. In case of any anomalous condition or loss of control, simply interrupting the power supply to the accelerator instantly terminates the external neutron source, leading to the immediate and safe shutdown of the fission process. The intrinsic safety of the ADS

allows it to be loaded with dedicated MA fuel without the concerns associated with critical reactors.

To quantify the transmutation capabilities of such systems, one can consider the European Facility for Industrial Transmutation (EFIT). This 400 MW_{th} lead-cooled ADS design concept, driven by an 800 MeV proton beam with a current of approximately 15 mA, is designed to achieve an MA transmutation rate of 42 kg_{MA}/TWh_{th}². This value represents the maximum theoretical transmutation rate achievable through pure fission, corresponding to the complete destruction of minor actinides without producing other heavy nuclides (such as plutonium). Operating an EFIT facility for approximately 10 years would be sufficient to transmute the entire MA inventory generated over 40 years of operation by a conventional 900 MW_e PWR using standard UOX fuel (3.7% ²³⁵U) at a burnup of 40 GWd/t_{HM} [6].

However, the practical implementation of an actinide burner is governed by strict thermodynamic and material constraints. Since the transmutation process is driven by nuclear fission, the maximum actinide destruction rate is strictly limited by the nominal thermal power of the reactor. Furthermore, the physical degradation of solid fuel elements induced by radiation damage in fast reactors restricts the maximum achievable burnup per irradiation cycle to approximately 25% [7]. Consequently, an effective transmutation system cannot operate in an OFC, but requires the fuel to be recycled and re-irradiated over multiple operational cycles. With out-of-core intervals of typically 10-15 years for spent fuel cooling, reprocessing, and refabrication, achieving a comprehensive reduction in MA mass would require continuous recycling over several hundred years of ADS operation.

Finally, although an ADS is specifically designed for actinide transmutation rather than optimised commercial power generation, its thermal power aligns with the scale of modern small modular reactors (SMRs). For instance, the EFIT facility is engineered with a total core thermal power of approximately 400 MW_{th} [8]. Nevertheless, a significant fraction of the electrical energy generated by the plant must be continuously recirculated to sustain the proton accelerator beam. For an ADS operating at a subcriticality level of $k_{\text{eff}} \approx 0.95$, this reduction amounts to approximately 12% of the electrical output [7].

1.2.1 An Example of ADS Design: The MYRRHA Project

Currently developed by the Belgian Nuclear Research Centre (SCK CEN), MYRRHA (Multi-purpose hYbrid Research Reactor for High-tech Applications) is one of the most extensively investigated ADS concepts, particularly within the EURATOM framework [10]. Intended as a highly versatile irradiation facility, MYRRHA is designed to operate in both critical and subcritical regimes. The primary goal of the MYRRHA project is to demonstrate the technical and operational feasibility of the full ADS concept and its application to HLW transmutation. Furthermore, MYRRHA will serve as a multi-purpose facility, supporting the development and testing of materials and fuels for Generation-IV and fusion reactors, as well as the commercial production of radioisotopes for the medical and industrial sectors.

²Kilograms of minor actinides transmuted per terawatt-hour of thermal energy




Figure 1.5: Schematic overview of the MYRRHA project. (a) Facility concept showing the proton beam delivery [9]. (b) Detailed view of the reactor vessel and its main internal components [10].

In its primary ADS operational mode, MYRRHA is driven by an external neutron source. As shown in Fig. 1.5a, this is achieved by coupling the reactor to a dedicated 400-m-long linear accelerator, which delivers a 600 MeV proton beam with a current of 4 mA. The beam is guided into the core through a central beamline (highlighted as a red tube in the vessel figure), striking a spallation target consisting of Lead-Bismuth Eutectic (LBE) (see Sec. 2.1.2). The surrounding active core is loaded with MOX fuel pins, and the reference neutronic design targets a subcriticality level of $k_{\text{eff}} \approx 0.97$ [11].

To accommodate this configuration, MYRRHA is designed as a pool-type reactor, where the core and primary systems are submerged inside a large pool of LBE acting as the primary coolant. Fig. 1.5b shows a vertical section of the double-walled reactor vessel, illustrating its detailed internal layout. The outer boundary, or guard vessel (1), acts as a safety containment in the unlikely event of a rupture of the inner vessel (2), which holds the primary LBE coolant and is sealed by an upper cover (3). A structural diaphragm (4) divides the coolant into two distinct thermal regions: a cold lower pool and a hot upper pool. The LBE circulation is driven by axial primary pumps (6), which force the heated coolant coming from the core (5) through dedicated primary heat exchangers (7) before returning it to the lower pool.

Chapter 2

ADS: Architecture and Core Physics

As described in the previous chapter, an ADS provides the intrinsic safety margins required for the transmutation of high concentrations of MAs by operating in a deeply subcritical regime. To sustain the nuclear chain reaction without relying on strict criticality, an ADS employs a fundamentally different macroscopic architecture compared to conventional nuclear power plants. Specifically, an ADS is a hybrid infrastructure consisting of three primary subsystems: a high-power particle accelerator, a heavy metal spallation target, and a subcritical nuclear core (see Fig. 2.1).

To provide a comprehensive physical and technological framework of this technology, this chapter is structurally divided into two main parts.

The first section explores the operational characteristics of the three macroscopic subsystems.

The second section is dedicated to the reactor physics of the subcritical core, introducing the mathematical formulation of neutron transport in a stationary regime, where the neutronic population is strictly governed by the external spallation source.

2.1 ADS Components

2.1.1 Accelerator

The particle accelerator is responsible for delivering the continuous beam of high-energy particles required to sustain the neutron population within the core. For industrial-scale transmutation purposes, the system typically requires a proton beam with an energy in the range from 600 MeV to 1 GeV. Typical beam currents range from approximately 5 mA to several tens of mA [7], depending on the required thermal power. Two main technological solutions are currently considered for this task: linear accelerators (LINACs) and cyclotrons.

A major engineering challenge for a high-energy, high-current cyclotron complex is the beam extraction system. To prevent catastrophic thermal damage and severe radioactive activation of the machine components, the proton beam must be extracted with near-zero losses. For these safety reasons, the final orbit must be sufficiently




Figure 2.1: Macroscopic architecture of an ADS, illustrating the coupling between the high-energy proton accelerator, the heavy metal spallation target, and the subcritical reactor core [7].

separated from the preceding orbit (a process known as single-turn extraction), to ensure that the beam does not physically strike the high-voltage electrostatic deflector. However, as the beam currents increase, space-charge repulsion widens the beam's transverse dimension, making a clean extraction increasingly difficult. Maintaining the necessary radial separation between orbits at the machine's outer radius requires a very high energy gain per turn. Consequently, this imposes stringent voltage and design requirements on the Radio Frequency (RF) cavities.

LINACs are largely preferred for high-power ADS applications because of their superior capability to handle very high beam currents while minimising beam losses. Strong transverse-focusing quadrupoles placed at frequent intervals along a LINAC, combined with longitudinal focusing for phase stability, allow a much larger charge per bunch to be accelerated without significant beam losses. Furthermore, LINACs operate with RF-cavity frequencies typically 10 to 20 times higher than those used in cyclotrons. Taken together, these factors imply that, in principle, a LINAC could provide a current from one to two orders of magnitude higher than a cyclotron without encountering extraction issues [7]. The major drawback of a LINAC is the considerable length, as the majority of the structure requires heavy shielding against radiation produced by even small beam losses.

The most stringent challenge associated with ADS technology is the extremely high reliability required of the proton accelerator. In a conventional accelerator facility, sudden beam interruptions (known as beam trips) are routine and represent a mere operational inconvenience. In an ADS, however, a beam trip instantaneously cuts off the external neutron source, triggering a severe neutronic and thermal transient.

The kinetic behaviour of a conventional critical reactor is buffered by delayed neutrons, which restrict macroscopic power transients to the timescale of their time constants (about 10 s). Conversely, a subcritical core relies entirely on the continuous prompt multiplication of the external source. Consequently, an instantaneous beam interruption causes the fission power to collapse on the timescale of the prompt-neutron lifetime (typically of the order of microseconds).

This rapid drop in fission power introduces a critical temporal discrepancy between the neutronic and the thermal dynamics of the system. While nuclear heat generation decreases on a microsecond timescale, the characteristic heat-transfer time from the fuel to the coolant is of the order of 0.1 to 1 s [7]. Because thermal energy remains stored in the fuel for approximately 1 s, the core structures experience severe temperature gradients. If beam trips occur frequently, these cyclic thermal contractions induce significant thermomechanical stress, potentially leading to fatigue failures of the structural components and raising significant safety concerns.

Assuming an annual plant availability of 70%, acceptable beam-trip frequencies can be categorised into three distinct ranges based on the trip duration (τ_b): up to 2×10^4 occurrences per year for $0 < \tau_b < 10$ s, 2×10^3 occurrences per year for $10 \text{ s} < \tau_b < 5 \text{ min}$, and less than 50 occurrences per year for $\tau_b > 5 \text{ min}$. Based on existing performance data, achieving these operational requirements necessitates reducing the frequency of the latter two categories by factors of approximately 6 and 35, respectively [12].

2.1.2 Spallation Target

The primary purpose of the spallation target is to convert the incident high-energy proton beam into the intense flux of fast neutrons required to sustain the subcritical fission chain. This conversion is achieved mainly through spallation, a nuclear reaction that knocks out several high-energy secondary particles, along with numerous lower-energy neutrons.

To maximise the efficiency of the process, defined as the average number of neutrons produced per incident proton, the target material must have a high atomic number (Z). For a 1 GeV proton beam, a heavy metal target can yield approximately 20 to 30 neutrons per proton. For industrial-scale applications, heavy liquid metals (HLMs) are the universally preferred target materials.

Given the extremely high power density, solid metal targets are generally unsuitable for industrial-scale ADS applications. The spallation process deposits a massive amount of heat in a highly localised volume, requiring a fluid with exceptional thermal conductivity and a high boiling point to efficiently dissipate the heat. Liquid metals allow for higher power densities and a significant reduction in direct irradiation damage compared with solid metal targets. However, liquid metals introduce specific engineering challenges, primarily concerning the corrosion and erosion of structural materials that are in direct contact with the fluid at high temperatures and high flow velocities.




Figure 2.2: Schematic representation of the spallation process induced by a high-energy proton interacting with a heavy nucleus. The reaction proceeds through the INC, pre-equilibrium thermalisation, and final de-excitation via isotropic evaporation or high-energy fission.

Physics of Spallation Process

Fig. 2.2 shows a schematic representation of the spallation reaction. After a high-energy proton strikes a heavy nucleus, the interaction proceeds through three distinct physical phases.

The first stage, occurring on a timescale of $10^{-22} - 10^{-23}$ s, is the intra-nuclear cascade (INC), where the proton undergoes multiple two-body collisions with the individual nucleons inside the target nucleus. This phase knocks out several high-energy secondary particles, particularly neutrons, protons and light ions, and leaves the residual nucleus in a highly excited state.

The next stage is the so-called pre-equilibrium phase. In this phase, the excitation energy is redistributed among the nucleons of the target, allowing the nucleus to reach thermal equilibrium on a timescale comparable to that of the INC.

Finally, the last stage is the de-excitation phase (occurring on a timescale of $10^{-6} - 10^{-8}$ s). Depending on the mass and charge of the nucleus, two main mechanisms may compete. Typically, the dominant mechanism is evaporation, which cools the target nucleus through the isotropic emission of a large number of lower-energy nucleons, predominantly neutrons. However, in the case of high- Z nuclei, such as actinides, fission can become an important competing mechanism, producing two fission fragments and fast neutrons. As a reference, for ^{208}Pb at incident proton energies around 1 GeV, fission is still strongly suppressed relative to evaporation channels [13].

2.1.3 Subcritical Core

The subcritical core of an ADS acts as a multiplier for the external neutron source generated by the spallation target. As explained in Sec. 1.1, to fulfil its primary objective of effective MA transmutation, the core must operate strictly with a fast neutron spectrum, since high-energy neutrons are mandatory to maximise the fission-to-capture ratio of transuranic isotopes. Consequently, the use of water, which is the standard coolant and moderator in commercial thermal reactors, or other light moderating fluids is precluded.

To ensure adequate heat removal without degrading the fast neutron energy spectrum, HLMS are considered the optimal choice for ADS applications. From a neutronic perspective, HLMS are characterised by high atomic masses and remarkably low neutron absorption cross-sections. When fast neutrons scatter off these heavy nuclei, the energy loss per collision is minimal, thus maintaining the hard neutron spectrum required to sustain efficient fission of MAs. Moreover, the high density and excellent thermal conductivity of HLMS facilitate efficient circulation, ensuring reliable decay heat removal purely through passive mechanisms in emergency scenarios. Finally, employing an HLM as the primary coolant allows direct integration with the spallation module, as the same fluid can serve both as coolant and as target material, thereby simplifying the ADS layout.

2.2 Physics of a Subcritical Core in Stationary Regime

Although the technological realisation of industrial-scale ADS is still at a preliminary research and development phase, the fundamental reactor physics governing subcritical multiplying systems is well established. This section focuses exclusively on the steady-state neutronic analysis of the core, laying the mathematical foundation necessary to ultimately determine the integral parameters that characterise the system.

2.2.1 Transport Equation

The distribution of neutrons in a critical system is governed by the Boltzmann transport equation, which expresses the balance between neutron production and neutron loss at each point in phase space (position $\mathbf{r}$, solid angle $\boldsymbol{\Omega}$, and energy E). The steady-state form of the Boltzmann equation for a critical system occupying a volume V can be written as [14]:

$$\begin{aligned} \boldsymbol{\Omega} \cdot \nabla \Phi_o(\mathbf{r}, E, \boldsymbol{\Omega}) + \Sigma_t(\mathbf{r}, E) \Phi_o(\mathbf{r}, E, \boldsymbol{\Omega}) &= \int_{4\pi} \int_0^\infty \Sigma_s(\mathbf{r}, E' \rightarrow E, \boldsymbol{\Omega}' \cdot \boldsymbol{\Omega}) \\ &\times \Phi_o(\mathbf{r}, E', \boldsymbol{\Omega}') d\boldsymbol{\Omega}' dE' + \frac{\chi(E)}{4\pi} \int_{4\pi} \int_0^\infty \nu(E') \Sigma_f(E') \Phi_o(\mathbf{r}, E', \boldsymbol{\Omega}') d\boldsymbol{\Omega}' dE', \end{aligned} \quad (2.1)$$

where $\Phi_o(\mathbf{r}, \boldsymbol{\Omega}, E)$ is the stationary angular neutron flux at position $\mathbf{r}$, direction $\boldsymbol{\Omega}$, and energy E ; Σ_t , Σ_s , and Σ_f are the macroscopic total, scattering, and fission

cross-sections, respectively; $\nu(E)$ is the average number of neutrons produced per fission at energy E ; and $\chi(E)$ is the fission neutron spectrum.

Eq. (2.1) can be expressed in a more compact form using operator formalism, by considering:

$$\boldsymbol{\Omega} \cdot \nabla \Phi_o(\mathbf{r}, E, \boldsymbol{\Omega}) = \mathcal{L} \Phi_o$$

as the leakage term, representing the net flow of neutrons out of the volume element at $\mathbf{r}$;

$$\Sigma_t(\mathbf{r}, E) \Phi_o(\mathbf{r}, E, \boldsymbol{\Omega}) = \mathcal{C} \Phi_o$$

as the collision term, representing the loss of neutrons due to absorption and out-scattering from the point $(E, \boldsymbol{\Omega})$ in the phase space;

$$\iint \Sigma_s(\mathbf{r}, \boldsymbol{\Omega}' \rightarrow \boldsymbol{\Omega}, E' \rightarrow E) \Phi_o(\mathbf{r}, \boldsymbol{\Omega}', E') d\boldsymbol{\Omega}' dE' = \mathcal{S} \Phi_o$$

as the in-scattering term, representing the gain of neutrons at $(E, \boldsymbol{\Omega})$ due to scattering from any other energies and directions;

$$\frac{\chi(E)}{4\pi} \iint \nu(E') \Sigma_f(E') \Phi_o(\mathbf{r}, \boldsymbol{\Omega}', E') d\boldsymbol{\Omega}' dE' = \mathcal{F} \Phi_o$$

as the fission neutron production term, representing the gain of neutrons at $(E, \boldsymbol{\Omega})$ due to fission reactions.

In the preceding equations, the terms are expressed as operators acting on the neutron flux Φ_o , denoted as $\mathcal{L}$, $\mathcal{C}$, $\mathcal{S}$, and $\mathcal{F}$, respectively. It is now possible to define the net neutron loss operator $\mathcal{M} = \mathcal{L} + \mathcal{C} - \mathcal{S}$, which accounts for neutron leakage, absorption, and net scattering transfer. Therefore, Eq. (2.1) can be rewritten in a more compact form as:

$$\mathcal{M} \Phi_o = \mathcal{F} \Phi_o. \quad (2.2)$$

A subcritical system, however, cannot sustain a steady-state neutron flux on its own; thus the homogeneous transport equation no longer admits a non-trivial steady-state solution. Therefore, it is necessary to introduce an independent external neutron source $s(\mathbf{r}, \boldsymbol{\Omega}, E)$ to maintain the neutron population. The governing equation becomes:

$$\mathcal{M} \Phi_{\text{in}} = \mathcal{F} \Phi_{\text{in}} + s, \quad (2.3)$$

or equivalently:

$$\mathcal{T} \Phi_{\text{in}} = s, \quad (2.4)$$

where $\mathcal{T} = \mathcal{M} - \mathcal{F}$ is the net balance operator, and Φ_{in} represents the solution of the transport equation in the presence of an external neutron source. The preceding equations show that in a subcritical system the evaluation of the neutron flux shifts from an eigenvalue problem to an inhomogeneous fixed-source problem.

2.2.2 Reactivity of the Subcritical Core

When considering the hypothetical condition of exact criticality, in which neutron production perfectly balances neutron losses (as described by Eq. (2.2)), a non-trivial stationary solution exists only if the operator $\mathcal{T}$ possesses a non-trivial null space. However, since this mathematical requirement is rarely satisfied in physical systems, the homogeneous equation typically admits only the trivial solution. To formally restore the steady-state balance, the effective multiplication factor k_{eff} is introduced. Consequently, Φ is defined as the solution of the eigenvalue problem:

$$\mathcal{M}\Phi = \frac{1}{k_{\text{eff}}}\mathcal{F}\Phi, \quad (2.5)$$

or equivalently:

$$(\mathcal{M}^{-1}\mathcal{F})\Phi = k_{\text{eff}}\Phi. \quad (2.6)$$

By the Krein-Rutman Theorem, there exists a unique eigenvalue $k_{\text{eff},0}$ such that:

$$k_{\text{eff},0} > |k_{\text{eff},n}|, \quad \forall n \geq 1, \quad (2.7)$$

and the corresponding eigenfunction Φ_0 is the only one that is non-negative over the entire phase space:

$$\Phi_0(\mathbf{r}, E, \boldsymbol{\Omega}) \geq 0, \quad \forall \mathbf{r}, E, \boldsymbol{\Omega}. \quad (2.8)$$

The eigenfunctions of the higher modes, Φ_n ($n \geq 1$), instead exhibit negative values in some regions of phase space, and are therefore unacceptable as physical solutions. Consequently, the physical k_{eff} is by definition $k_{\text{eff},0}$, and the solution Φ of Eq. (2.5) is the fundamental mode.

In order to shift to the subcritical system, and to quantify the physical impact of the external spallation source on the subcritical neutron balance, it is necessary to introduce the concept of neutron importance. This is mathematically achieved by first defining the adjoint operators for net loss ($\mathcal{M}^*$) and fission production ($\mathcal{F}^*$).

Consider now the inner product $\langle \cdot, \cdot \rangle$ of two real-valued functions, w and v , defined as the phase-space integral over the system:

$$\langle w, v \rangle = \int_V \int_{4\pi} \int_0^\infty w(\mathbf{r}, E, \boldsymbol{\Omega}) v(\mathbf{r}, E, \boldsymbol{\Omega}) d\mathbf{r} d\boldsymbol{\Omega} dE. \quad (2.9)$$

It is possible to define, for the unperturbed critical system, the adjoint function Φ^* such that, for the generic forward flux Φ and adjoint flux Φ^* :

$$\langle \Phi^*, \mathcal{T}\Phi \rangle = \langle \mathcal{T}^*\Phi^*, \Phi \rangle, \quad (2.10)$$

where the eigenfunction Φ^* is given by the corresponding homogeneous equation:

$$\mathcal{M}^*\Phi^* = \frac{1}{k_{\text{eff}}}\mathcal{F}^*\Phi^*. \quad (2.11)$$

Eq. (2.10) is rigorously valid provided that Φ and Φ^* satisfy compatible boundary conditions, specifically, zero incoming flux for Φ and zero outgoing adjoint flux for Φ^* at the vacuum boundaries, ensuring that all surface terms arising from integration by parts vanish. It is important to note that the adjoint eigenvalue problem shares

the same eigenvalue spectrum as the forward equation [15]. Therefore, the effective multiplication factor k_{eff} associated with the fundamental adjoint mode Φ_0^* is strictly identical to the physical multiplication factor governing the forward fundamental mode Φ_0 .

Multiplying Eq. (2.5) by Φ^* and integrating over the phase space, k_{eff} can be written as:

$$k_{\text{eff}} = \frac{\langle \Phi^*, \mathcal{F}\Phi \rangle}{\langle \Phi^*, \mathcal{M}\Phi \rangle}, \quad (2.12)$$

which represents the ratio of the total number of neutrons produced by fission to the total number of neutrons lost in the system, weighted by the adjoint function Φ^* . From the previous equation, it is evident that the adjoint function Φ^* can be interpreted as a measure of the importance of each neutron in sustaining the chain reaction. In other words, the adjoint function quantifies the contribution of each neutron to the overall neutron balance in the system.

Considering now the subcritical system, projecting the inhomogeneous equation onto the critical adjoint flux yields:

$$\langle \Phi^*, (\mathcal{M} - \mathcal{F})\Phi_{\text{in}} \rangle = \langle \Phi^*, s \rangle. \quad (2.13)$$

At this stage, an analytical approximation is required. For subcritical systems operating not too far from criticality (typically $k_{\text{eff}} \approx 0.97$, as is standard for an ADS), it may be assumed that the spatial and energy distribution of the subcritical flux Φ_{in} is dominated by the fundamental mode of the unperturbed system, and the presence of the external source does not significantly alter the shape of the flux. Relying on this physical approximation ($\Phi \approx \Phi_{\text{in}}$), Eq. (2.12) can be written as:

$$k_{\text{eff}} = \frac{\langle \Phi^*, \mathcal{F}\Phi \rangle}{\langle \Phi^*, \mathcal{M}\Phi \rangle} \approx \frac{\langle \Phi^*, \mathcal{F}\Phi \rangle}{\langle \Phi^*, s \rangle + \langle \Phi^*, \mathcal{F}\Phi \rangle}. \quad (2.14)$$

However, to better account for the contribution of the external source in subcritical systems, a different parameter known as the source multiplication (or criticality) factor k_s is introduced. This parameter is defined as the ratio of the total number of neutrons produced by fission to the total number of neutrons losses:

$$k_s = \frac{\langle u, \mathcal{F}\Phi_{\text{in}} \rangle}{\langle u, \mathcal{M}\Phi_{\text{in}} \rangle} = \frac{\langle u, \mathcal{F}\Phi_{\text{in}} \rangle}{\langle u, s \rangle + \langle u, \mathcal{F}\Phi_{\text{in}} \rangle}, \quad (2.15)$$

where u is a unit vector. The k_s parameter is a more accurate measure of the subcriticality in an ADS, relying on the actual inhomogeneous flux distribution Φ_{in} , which captures the physical perturbations induced by the spallation source s .

Finally, the effect of the external source on a specific system response can be evaluated using the adjoint function. Let us consider a response function R representing a physical quantity of interest, as for example the fission rate. It is possible to introduce the adjoint function n^* , which is the solution of the corresponding inhomogeneous equation:

$$\mathcal{T}^* n^* = R. \quad (2.16)$$

Combining both Eq. (2.4) and Eq. (2.16), it is possible to write:

$$\langle R, \Phi_{\text{in}} \rangle = \langle \mathcal{T}^* n^*, \Phi_{\text{in}} \rangle = \langle n^*, \mathcal{T} \Phi_{\text{in}} \rangle = \langle n^*, s \rangle. \quad (2.17)$$

In other words, the adjoint function n^* can be interpreted as an importance function that weighs the contribution of each neutron in the source distribution to the response of interest. Therefore, by maximising the inner product $\langle n^*, s \rangle$, it is possible to optimise the source distribution to achieve the highest possible efficiency in driving the desired response within the subcritical system.

2.2.3 Source Importance Parameter and Proton Beam Current

In a source-driven subcritical system, not all source neutrons are equally important in sustaining the chain reaction. As discussed in the previous section, the importance of a neutron is quantified by the adjoint flux Φ^* .

Starting from the definition of k_{eff} in Eq. (2.14), it is possible to write:

$$\frac{1}{k_{\text{eff}}} - 1 = \frac{\langle \Phi^*, s \rangle}{\langle \Phi^*, \mathcal{F}\Phi \rangle} = \frac{\langle \Phi^*, s \rangle}{\langle \Phi^*, \mathcal{F}\Phi \rangle} \cdot \frac{\langle u, s \rangle}{\langle u, \mathcal{F}\Phi \rangle} \cdot \frac{\langle u, \mathcal{F}\Phi \rangle}{\langle u, s \rangle}, \quad (2.18)$$

and recalling the definition of k_s in Eq. (2.15) along with the physical approximation $\Phi \approx \Phi_{\text{in}}$, the equation becomes:

$$\frac{1}{k_{\text{eff}}} - 1 = \frac{\langle \Phi^*, s \rangle}{\langle \Phi^*, \mathcal{F}\Phi \rangle} \left(\frac{1}{k_s} - 1 \right) \frac{\langle u, \mathcal{F}\Phi \rangle}{\langle u, s \rangle}. \quad (2.19)$$

Given the preceding equations, it is now possible to define the source importance parameter φ^* , which is defined as the ratio of the average importance of the source neutrons to the average importance of fission neutrons:

$$\varphi^* = \frac{\langle \Phi^*, s \rangle}{\langle \Phi^*, \mathcal{F}\Phi \rangle} / \frac{\langle u, s \rangle}{\langle u, \mathcal{F}\Phi \rangle}. \quad (2.20)$$

In terms of the k_{eff} and k_s parameters, from Eq. (2.19), φ^* can be expressed as:

$$\varphi^* = \left(\frac{1}{k_{\text{eff}}} - 1 \right) / \left(\frac{1}{k_s} - 1 \right). \quad (2.21)$$

This parameter is crucial for the design of an ADS, as it quantifies how effectively the source neutrons contribute to sustaining the chain reaction compared to fission neutrons. An optimised ADS is typically designed to achieve $\varphi^* > 1$, indicating that the source neutrons are more effective in driving the reaction than the fission neutrons. From a mathematical perspective, φ^* relies on the projection of the source distribution s onto the fundamental adjoint mode Φ^* . Consequently, the parameter φ^* can be interpreted as a measure of the overlap between the source distribution and the first eigenfunction, reflecting how effectively the source neutrons contribute to the fission chain reaction.

Another fundamental parameter in the design of an ADS is the multiplication factor M . Consider the spallation yield N_s , defined as the average net number of

source neutrons produced by spallation per incident proton. Assuming a subcritical system ($k_s < 1$) and a sufficiently large number of generations n to achieve the exhaustion of transient effects and the attainment of a steady-state condition, the total number of generated neutrons per incident proton can be obtained as:

$$N_s k_s + N_s k_s^2 + \dots + N_s k_s^n = N_s \frac{k_s}{1 - k_s} = N_s M, \quad (2.22)$$

where the multiplication M of the source neutrons is given by:

$$M = \frac{k_s}{1 - k_s}. \quad (2.23)$$

From Eq. (2.22), M can be seen as the total number of fission neutrons produced per source neutron. According to Eq. (2.23), after calculating the multiplication factor, it is possible to obtain the source criticality factor of the system as:

$$k_s = \frac{M}{M + 1}. \quad (2.24)$$

The φ^* parameter is also related to the primary operational parameter of the ADS, which is the continuous proton beam current. Introducing the average number of neutrons per fission ν , and the average source neutrons per fission Γ :

$$\nu = \frac{\langle u, \mathcal{F}\Phi \rangle}{\langle \Sigma_f, \Phi \rangle}, \quad \Gamma = \frac{\langle u, s \rangle}{\langle \Sigma_f, \Phi \rangle}, \quad (2.25)$$

it is possible to write:

$$\frac{\Gamma}{\nu} \varphi^* = \frac{1}{k_{\text{eff}}} - 1. \quad (2.26)$$

At this stage, it is possible to establish the fundamental relationship connecting the accelerator proton beam current i_p , the macroscopic thermal power P_{th} , and the core subcriticality. The current i_p can be expressed as:

$$i_p = \frac{\Gamma}{N_s} \cdot \frac{P_{\text{th}}}{Q}, \quad (2.27)$$

where Q is the recoverable thermal energy per fission. The ratio P_{th}/Q yields the fission rate (total fissions per second) occurring within the core. Multiplying this rate by the parameter Γ provides the total emission rate of source neutrons required to sustain the steady-state chain reaction. Finally, by dividing this value by the spallation yield N_s , the required proton impingement rate on the target is obtained.

Substituting Eq. (2.26) into the previous equation leads to the final relation:

$$i_p = \frac{\nu P_{\text{th}}}{\varphi^* N_s Q} \frac{1 - k_{\text{eff}}}{k_{\text{eff}}}, \quad (2.28)$$

which gives the proton current required to operate the ADS at the thermal power level P_{th} .

Chapter 3

Numerical Methods for Particle Transport

As established in the previous chapter, the neutronic characterisation of an ADS requires solving the steady-state Boltzmann equation as an inhomogeneous fixed-source problem. Given the complex three-dimensional heterogeneity of the core and the strong energy dependence of cross-sections in the fast neutron energy range, finding analytical solutions is impossible for practical applications. Consequently, numerical methods must be employed.

These numerical strategies are traditionally divided into two main categories: deterministic methods and Monte Carlo (MC) stochastic methods. In this study, a hybrid computational strategy was employed, combining two distinct reactor-physics codes: the deterministic ERANOS (European Reactor ANalysis Optimized code System, 2.2 version) code system [16] and the stochastic MCNP (Monte Carlo N-Particle, 6.3.0 version) code [17]. The latter was selected as it is currently the most widely adopted MC code in nuclear system design, supported by extensive international experience and validation.

Both deterministic and MC codes are essential for the conceptual core design and neutronic characterisation of ADS, as they offer complementary capabilities. Deterministic codes provide the computational efficiency required to perform complex fuel-cycle and depletion analyses, commonly referred to as burnup calculations. This represents a critical advantage, especially considering that the MCNP code can perform burnup calculations only in criticality mode (the distinction between criticality and fixed-source calculations is detailed in Section 3.2). However, accurately simulating an ADS requires a fixed-source calculation to model the external proton beam driving the subcritical core. Consequently, MCNP cannot be directly employed for ADS burnup evaluations.

On the other hand, ERANOS is constrained by a fundamental physical limitation. Standard multi-group cross-section libraries (see Sec. 3.1.2) are typically evaluated only up to 20 MeV. Instead, the spallation source driving an ADS is characterised by a high-energy cascade triggered by protons at hundreds of MeV, producing a neutron spectrum that extends well beyond the 20 MeV threshold.

To overcome these limitations, MCNP is used to simulate the high-energy spallation process and track particles down to the 20 MeV threshold. Once neutrons fall

below this energy threshold, they are extracted as an external source for ERANOS, which is then used for transport and multiplication of the system below 20 MeV.

To provide the necessary theoretical background for this methodology, this chapter will outline the fundamental principles of deterministic and MC approaches, with a particular focus on the specific MCNP features used in this study for the neutronic analysis of the investigated ADS and the extraction of the spallation source.

3.1 ERANOS Deterministic Code System

3.1.1 Variable Discretisation

Deterministic methods solve the integro-differential Boltzmann equation by discretising the continuous phase space (position, energy, solid angle, and time in non-steady-state problems) into finite elements. This technique evaluates the average macroscopic neutron balance within each discrete domain. Thus, in this approach, the solution describes the average behaviour of the particle population.

The discretisation of the energy variable is performed by introducing the multi-group approximation. In this sense, the continuous-energy domain is partitioned into a certain number of discrete intervals, called energy groups. For each energy group, a macroscopic cross-section is defined for each interaction. Macroscopic cross-sections are defined in such a way as to preserve the total reaction rate within a group, as illustrated in Sec. 3.1.2.

Spatial variables are treated by dividing the geometry of the problem into a finite number of representative cells, creating a spatial mesh. Each cell is treated as a homogeneous medium assigned a single macroscopic cross-section vector (or matrix, in the case of scattering transfer) for each reaction. The transport equation is then solved in every single cell using the corresponding cross-sections.

To handle the angular dependence of the flux and scattering cross-sections, two possible approximations may be employed. The first is the Discrete Ordinates (S_N) method, in which the continuous solid angle is divided into N discrete directions. An alternative approach is the Spherical Harmonics (P_N) method, which expands both the flux and the scattering matrices as a series of spherical harmonics as a function of the scattering cosine $\mu = \Omega' \cdot \Omega$. For example, the angular flux can be expanded using Legendre polynomials $P_l(\mu)$ as:

$$\Phi(\mathbf{r}, E, \mu) = \sum_{n=0}^{\infty} \left(\frac{2n+1}{4\pi} \right) \Phi_n(\mathbf{r}, E) P_n(\mu), \quad (3.1)$$

where $\Phi_n(\mathbf{r}, E)$ is the n th moment of the flux:

$$\Phi_n(\mathbf{r}, E) = \int_{-1}^1 \Phi(\mathbf{r}, E, \mu) P_n(\mu) d\mu. \quad (3.2)$$

Deterministic methods are particularly useful during the pre-design and design stages. In these phases, evaluating a broad parameter space to identify general performance trends is more critical than achieving absolute precision, making their high computational speed essential. However, spatial and energy homogenisations can

introduce critical inaccuracies in regions characterised by steep flux gradients, such as the boundary between the spallation target and the subcritical core. Moreover, as already explained, ERANOS still requires Monte Carlo simulations for the calculation of the neutron source.

3.1.2 Multi-Group Energy Formalism

Since a continuous-energy treatment is computationally prohibitive for full-core deterministic calculations, the energy variable must be discretised using the multi-group approximation. The continuous-energy domain is partitioned into g discrete intervals, called energy groups, defined by the boundaries:

$$E_g < E_{g-1} < \dots < E_1 < E_0,$$

up to a certain energy threshold E_0 . By convention, the index $g = 1$ denotes the highest energy group (corresponding to fast neutrons, bounded between E_1 and the maximum energy E_0), while the index increases as the neutron energy decreases towards the thermal region.

To formally derive the multi-group transport equation, the continuous Boltzmann equation must be integrated over each energy interval $\Delta E_g = [E_g, E_{g-1}]$. First, it is necessary to define the multi-group angular flux $\Phi_g(\mathbf{r}, \boldsymbol{\Omega})$ as the total integrated angular flux within the g -th energy interval [18]:

$$\Phi_g(\mathbf{r}, \boldsymbol{\Omega}) = \int_{E_g}^{E_{g-1}} \Phi(\mathbf{r}, E, \boldsymbol{\Omega}) dE. \quad (3.3)$$

A fundamental requirement of the multi-group approximation is the preservation of reaction rates within each group. To ensure this, the group-averaged macroscopic cross-section $\Sigma_{x,g}(\mathbf{r}, \boldsymbol{\Omega})$ for a generic interaction type x (total, fission, or absorption) is defined as:

$$\Sigma_{x,g}(\mathbf{r}, \boldsymbol{\Omega}) = \frac{1}{\Phi_g(\mathbf{r}, \boldsymbol{\Omega})} \int_{E_g}^{E_{g-1}} \Sigma_x(\mathbf{r}, E) \Phi(\mathbf{r}, E, \boldsymbol{\Omega}) dE. \quad (3.4)$$

This definition ensures that the product of the group cross-section and the group flux equals the integrated continuous reaction rate over the energy interval ΔE_g :

$$\Sigma_{x,g}(\mathbf{r}, \boldsymbol{\Omega}) \Phi_g(\mathbf{r}, \boldsymbol{\Omega}) = \int_{E_g}^{E_{g-1}} \Sigma_x(\mathbf{r}, E) \Phi(\mathbf{r}, E, \boldsymbol{\Omega}) dE. \quad (3.5)$$

Applying the same integration procedure to the scattering and fission sources, we define the multi-group scattering matrix $\Sigma_{s,g' \rightarrow g}(\mathbf{r}, \boldsymbol{\Omega}' \cdot \boldsymbol{\Omega})$, which quantifies the probability of a neutron scattering from a primary group g' into a secondary group g :

$$\Sigma_{s,g' \rightarrow g}(\mathbf{r}, \boldsymbol{\Omega}' \cdot \boldsymbol{\Omega}) = \frac{\int_{E_g}^{E_{g-1}} dE \int_{E_{g'}}^{E_{g'-1}} \Sigma_s(\mathbf{r}, E' \rightarrow E, \boldsymbol{\Omega}' \cdot \boldsymbol{\Omega}) \Phi(\mathbf{r}, E', \boldsymbol{\Omega}') dE'}{\Phi_{g'}(\mathbf{r}, \boldsymbol{\Omega}')}. \quad (3.6)$$

Similarly, the integrated external spallation source for group g is defined as:

$$s_g(\mathbf{r}, \boldsymbol{\Omega}) = \int_{E_g}^{E_{g-1}} s(\mathbf{r}, E, \boldsymbol{\Omega}) dE. \quad (3.7)$$

Integrating every term of the continuous Boltzmann equation over ΔE_g yields the final set of G coupled multi-group transport equations governing the subcritical system:

$$\begin{aligned} \boldsymbol{\Omega} \cdot \nabla \Phi_g(\mathbf{r}, \boldsymbol{\Omega}) + \Sigma_{t,g}(\mathbf{r}) \Phi_g(\mathbf{r}, \boldsymbol{\Omega}) = & \sum_{g'=1}^G \int_{4\pi} \Sigma_{s,g' \rightarrow g}(\mathbf{r}, \boldsymbol{\Omega}' \cdot \boldsymbol{\Omega}) \Phi_{g'}(\mathbf{r}, \boldsymbol{\Omega}') d\boldsymbol{\Omega}' \\ & + \frac{\chi_g}{4\pi} \sum_{g'=1}^G \nu_{g'} \Sigma_{f,g'}(\mathbf{r}) \phi_{g'}(\mathbf{r}) + s_g(\mathbf{r}, \boldsymbol{\Omega}). \end{aligned} \quad (3.8)$$

where $\chi_g = \int_{E_g}^{E_{g-1}} \chi(E) dE$ is the fraction of fission neutrons emitted directly into group g .

This system of equations highlights that the energy groups communicate with each other through scattering transition matrices, and are globally linked by the fission source.

3.1.3 ERANOS Computational Scheme

In practical calculations, deterministic codes like ERANOS do not solve the core transport equation in a single step. Instead, the macroscopic cross-sections ($\Sigma_{x,g}$) defined in Eq. (3.5) are evaluated and employed through a hierarchical, three-stage computational scheme:

1. Nuclear Data Processing The foundation of any deterministic calculation lies in the continuous-energy evaluated nuclear data libraries (such as ENDF/B or JEFF). In this preliminary stage, the microscopic cross-sections are processed and condensed into an ultra-fine multi-group format: in the ERANOS system, 1968 energy groups are available for the most important nuclides (in fuel, steel structures and heavy-metal coolant). This highly detailed discretisation ensures that all energy-dependent resonance structures are accurately preserved.

2. Cell-Level Calculation and Homogenisation The problem is decomposed into elementary sub-domains, known as unit cells, in order to capture the local physical effects arising from the spatial heterogeneity of the system. Rather than explicitly modelling the full-core geometry, the system is approximated through a set of cells, each characterised by a specific material composition. Dedicated codes (such as the ECCO cell code within ERANOS [16]) deterministically solve the local transport equation within each cell using the ultra-fine energy libraries. Once the local spectrum is obtained, it is used to perform a spatial homogenisation and an energy condensation. The detailed structural heterogeneities are homogenised into a single material, defined as a volume-weighted mixture of the original physical components. Furthermore, the ultra-fine energy grid is collapsed into a coarser broad-group structure (typically 30 to 50 groups). The output of this stage is a library of macroscopic cross-sections for subsequent whole-core calculations.

3. Core-Level Transport Calculation In the final stage, the macroscopic cross-sections derived at the cell level are mapped onto the corresponding homogeneous spatial regions defined in the whole-core geometric mesh. The deterministic transport solvers (using the S_N or P_N methods previously discussed) finally evaluate the steady-state multi-group transport equation over the entire geometry.

3.2 MCNP Stochastic Code

MC methods operate on an entirely different paradigm compared to deterministic methods. Instead of explicitly solving the transport equation, MC methods simulate the random walk of individual particles as they travel through the system, recording the specific quantities of interest requested by the user. The macroscopic quantities of the physical system are then inferred from the average behaviour of millions of simulated particles. MC methods allow for the explicit modelling of complex three-dimensional geometries and employ continuous-energy cross-sections. For the analysis of an ADS, MCNP is fundamental because it can accurately model the complex physics of the spallation target. This represents a crucial advantage over ERANOS, which is restricted by the upper energy boundary of its cross-section libraries.

The MCNP code can operate in two different modes, depending on the system under analysis. The first is the criticality mode, where the code runs successive fission cycles, starting each cycle from the distribution of neutrons generated in the previous one, and estimates the effective multiplication factor k_{eff} . The second is the fixed-source mode, which allows simulating the neutron multiplication triggered by an external source. This latter mode is essential for accurately modelling the primary proton beam and the resulting spallation cascade within the subcritical ADS core. In this configuration, the computational process is organised into independent MC histories, where a single history comprises the transport of a source particle and the entire cascade of generated secondary particles.

The main drawback of MC methods is the computational cost. Achieving a suitably low statistical uncertainty often demands simulating massive numbers of particle histories, requiring High-Performance Computing (HPC) clusters.

3.2.1 Geometry

To simulate the complex heterogeneity of the reactor, MCNP does not rely on spatial meshes. In the MCNP input file, the user defines an exact three-dimensional configuration through Boolean combinations (intersections, unions, and complements) of analytical surfaces. Each distinct volumetric region bounded by a set of surfaces is defined as a cell. This modelling technique is known as Constructive Solid Geometry (CSG).

Each cell is then filled with a specific material, identified by a unique material number. The physical properties corresponding to each material number are explicitly detailed in a dedicated section of the input file, where the user specifies the macroscopic density and the precise isotopic composition (in atomic or mass fractions) of the medium.

To handle repetitive structures, such as the hundreds of identical fuel pins within a reactor core, or a collection of ordinary cells that must be treated in a common manner, MCNP introduces the concept of a universe. A universe is a geometry unit that can be manipulated or referenced as a group. Once defined and assigned a unique universe number, this entire structure can be embedded into any other cell across the global geometry.

Once the complete geometric layout is established, during the random walk simulation, the code relies on the bounding surfaces to strictly trace particle trajectories. At each step, MCNP calculates the distance to the nearest geometric boundary along the particle's flight path and compares it with the stochastically sampled distance to the next physical collision. If the distance to the boundary is shorter, the particle crosses the surface into the adjacent cell without interacting with the current medium.

3.2.2 Transport

During the random walk, particle interactions and flight distances are sampled from probability density functions (PDFs) derived from continuous-energy evaluated nuclear data libraries.

Within a cell, the probability to collide between l and $l + dl$ after travelling a distance l is given by [17]:

$$f(l)dl = \Sigma_t e^{-\Sigma_t l} dl, \quad (3.9)$$

according to the PDF $f(l)$. Here, the macroscopic cross-section of the medium Σ_t is interpreted as the probability per unit length of a collision. Now, by sampling a random number ζ uniformly distributed on the interval $[0, 1)$, it is possible to impose:

$$F(l) = \zeta, \quad \text{with} \quad F(l) = \int_0^l \Sigma_t e^{-\Sigma_t s} ds = 1 - e^{-\Sigma_t l}, \quad (3.10)$$

where $F(l)$ is the cumulative distribution function (CDF). Since $1 - \zeta$ and ζ are distributed in the same uniform manner, the distance to the next collision l can be obtained as:

$$l = -\frac{1}{\Sigma_t} \ln(1 - \zeta) = -\frac{1}{\Sigma_t} \ln(\zeta). \quad (3.11)$$

In summary, the MC algorithm evaluates the physical quantity of interest (in this case, the flight distance l) by sampling a uniform random number and inverting the corresponding CDF, as illustrated in Fig. 3.1a.

When a particle collides, the colliding nuclide and the type of interaction are selected. These selections follow the same sampling principle described above, but applied to discrete PDFs. Let us consider the example of the reaction selection. Each reaction channel i is associated with a microscopic cross-section $\sigma_i(E)$, so that the probability of reaction i occurring, given that a collision has taken place, is:

$$p_i = \frac{\sigma_i(E)}{\sigma_t(E)}, \quad \text{with} \quad \sum_{i=1}^N p_i = 1, \quad (3.12)$$

where $\sigma_t(E) = \sum_{i=1}^N \sigma_i(E)$ is the total microscopic cross-section and N is the number of possible reaction channels. The discrete CDF is then built by summing the




Figure 3.1: Visual representation of the MC sampling processes. (a) Inversion of the continuous CDF to determine the flight distance l . (b) Sampling from a discrete CDF to select a specific reaction channel k .

individual probabilities:

$$P_k = \sum_{i=1}^k p_i, \quad k = 1, \dots, N. \quad (3.13)$$

As in the continuous case, a random number ζ uniformly distributed on $[0, 1)$ is sampled, and the reaction type k is selected as the one satisfying:

$$P_{k-1} < \zeta \leq P_k, \quad (3.14)$$

with $P_0 = 0$ by convention. As visually represented in Fig. 3.1b, this practical procedure consists of scanning the cumulative sum of reaction probabilities until it exceeds the sampled value ζ , and selecting the corresponding reaction channel.

Once the interaction type has been determined, the number of outgoing particles, along with their respective energy and direction, are sampled from the corresponding distributions.

When the collision produces new particles, as in the case of proton spallation, these secondary particles are temporarily stored in a memory bank. Once the track of the current particle has terminated, MCNP transports the generated particles according to a Last In, First Out (LIFO) scheme. This means that the most recently generated secondary particle, placed at the top of the memory stack, is the first to be subsequently simulated, while previously stored particles remain in the bank until the stack is cleared. A schematic representation of this transport sequence is provided in Fig. 3.2.

3.2.3 Tallies

Physical quantities of interest, such as particle flux, are obtained by recording the contribution of each particle track in specific counting structures known as tallies.

Mathematically, given a physical quantity of interest X (such as a specific reaction rate or an average flux), its expected value $\langle X \rangle$ is defined as the integral of a specific response function $R(E, \boldsymbol{\Omega})$ weighted by the angular particle flux $\Phi(\mathbf{r}, E, \boldsymbol{\Omega}, t)$:




Figure 3.2: Schematic representation of the LIFO particle transport sequence in MCNP. The numbers indicate the chronological order in which the tracks are simulated. A primary incident neutron (black, 1) generates a photon (red, 4) in its first collision, and subsequently generates two secondary neutrons (green, 2 and blue, 3) in a second collision. Following the LIFO logic, the most recently stored neutrons are extracted from the memory stack and transported before the previously generated photon.

$$\langle X \rangle = \int_V \int_E \int_{\Omega} \int_t \Phi(\mathbf{r}, E, \mathbf{\Omega}, t) R(E, \mathbf{\Omega}) d\mathbf{\Omega} dE d\mathbf{r} dt. \quad (3.15)$$

The response function represents the probability density for a particle to contribute to the physical quantity of interest.

Instead of solving the analytical equation, the MC tally estimates its expected value stochastically by tracking a large number N of independent particle histories: the algorithm employs a specific estimator to record the discrete contribution x_i , known as the score, made by each particle to the quantity of interest during its random walk. According to the Law of Large Numbers, the expected value is approximated by the arithmetic mean of the individual historical scores:

$$\langle X \rangle \approx \bar{x} = \frac{1}{N} \sum_{i=1}^N x_i. \quad (3.16)$$

MCNP provides several standard categories of estimators. The most common is the track-length estimator for the average cell flux. Recalling that the angular flux can be written as a function of the particle density $n(\mathbf{r}, E, \mathbf{\Omega}, t)$ and the speed of the

particle $v(E)$ as $\Phi = vn$, the average particle cell flux is:

$$\begin{aligned}\bar{\Phi}_V &= \frac{1}{V} \int_V \int_E \int_{\Omega} \int_t \Phi(\mathbf{r}, E, \boldsymbol{\Omega}, t) d\boldsymbol{\Omega} dE d\mathbf{r} dt \\ &= \frac{1}{V} \int_V \int_E \int_{\Omega} \int_t v(E) n(\mathbf{r}, E, \boldsymbol{\Omega}, t) d\boldsymbol{\Omega} dE d\mathbf{r} dt \\ &= \frac{1}{V} \int_V \int_E \int_t v(E) N(\mathbf{r}, E, t) dE d\mathbf{r} dt,\end{aligned}\tag{3.17}$$

where $N(\mathbf{r}, E, t) = \int_{\Omega} n(\mathbf{r}, E, \boldsymbol{\Omega}, t) d\boldsymbol{\Omega}$ is the density of particles regardless of their trajectories, and V is the cell volume. Defining ds to be the differential unit of track length, $ds = vdt$, yields [17]:

$$\bar{\Phi}_V = \frac{1}{V} \int_V \int_E \int_s N(\mathbf{r}, E, t) dE d\mathbf{r} ds.\tag{3.18}$$

The quantity $N(\mathbf{r}, E, t)ds$ may be thought of as a track-length density. Thus, the average flux can be estimated stochastically by summing single track lengths. The MCNP code estimates $\bar{\Phi}_V$ by summing the score $\frac{WT_l}{V}$ for all particle tracks in the cell, as:

$$\bar{\Phi}_V = \sum_l \frac{WT_l}{V},\tag{3.19}$$

where W is the statistical weight of each particle (which may differ from unity when variance reduction techniques are employed), and T_l is the track length of the particle in the cell.

MCNP tallies are normalised per starting particle, and it is possible to assign bin boundaries in energy, time and direction, to accumulate the results only in a portion of the phase space.

Each tally is provided with an estimated relative error R . According to the Central Limit Theorem, the sample of independent random variables has a distribution that tends to the Gaussian as the number of simulated histories N approaches infinity. Thus, the error associated with the sample mean can be computed as one estimated standard deviation of the mean $S_{\bar{x}}$.

The estimated relative error can be used to form confidence intervals about the estimated mean. In the Gaussian limit, the probability for the true value to be in the range $\bar{x}(1 \pm R)$ is 68%, and 95% in the range $\bar{x}(1 \pm 2R)$. It is crucial to emphasise that the confidence intervals refer to the precision of the MC simulation itself, and not to the accuracy of the result compared to the true physical value. As a general guideline, a relative error less than 0.10 is required to consider the confidence intervals reliable [17].

Moreover, the validity of these statistical interpretations relies on the assumption that the entire phase space of the problem has been exhaustively sampled. To assist the user in detecting such issues, and to monitor the convergence of the simulation, MCNP computes a parameter called figure of merit (FOM) for one bin per tally, tracking its behaviour as a function of the number of particle histories. The FOM is defined as:

$$FOM = \frac{1}{R^2 T},\tag{3.20}$$

where T is the computational time (in minutes). An MC simulation is more efficient for larger values of the FOM , since less computational time is required to achieve a given error R . Moreover, an approximately constant FOM with increasing number of simulated particles indicates a reliable calculation, as R^2 is proportional to $1/N$ and T is proportional to N . Instead, a sharp decrease suggests that a rare particle path has significantly affected the tally result.

3.2.4 PTRAC

The estimation of physical quantities of interest through tallies is the main feature of MC codes. However, other features are available in MCNP for different analyses. A different approach from the transport and tally methods described above is provided by the PTRAC mode.

The PTRAC (Particle Tracking) mode generates a detailed output file in HDF5 format that records specific discrete events within the histories of simulated particles. Specifically, in the PTRAC mode, five different types of events are recorded:

- **SRC (Source)**: the generation of a source particle. This includes particles originating from a user-defined fixed distribution or neutrons sampled from the spatial fission distribution at the beginning of a criticality cycle;
- **COL (Collision)**: any interaction occurring in the phase space during the particle's random walk within the material medium, comprising both scattering events (elastic and inelastic) and nuclear reactions;
- **SUR (Surface Crossing)**: the intersection of a particle's trajectory with a boundary surface defining the geometric cells;
- **BNK (Bank)**: the creation of secondary particles (distinct from the primary source) that are temporarily stored in the simulation bank for subsequent transport;
- **TER (Termination)**: the termination of a particle's track. This applies to both primary and secondary particles, and typically occurs due to physical absorption, geometric leakage, or when the particle triggers specific user-defined cut-offs (such as energy or time limits).

In addition to the event type, MCNP registers a comprehensive set of variables for the particle at each recorded step. These data include the exact phase-space coordinates (spatial position x, y, z and direction cosines u, v, w), kinetic energy E , statistical weight W , time t , and particle type. Furthermore, depending on the specific nature of the event, the code can append contextual tags, such as the interaction type, the identification of the target nucleus in a collision, or the specific cell and surface identifiers. Upon completion of the simulation, the resulting output file can be parsed by means of the MCNPTools package [19] and subsequently post-processed via external algorithms to extract customised quantities of interest.

The PTRAC capability becomes necessary when the analysis requires tracking a particle's trajectory or isolating specific events that cannot be captured by standard

integrated tallies. In the context of the present work, the PTRAC mode was the fundamental tool employed for the neutron source extraction (as detailed in Sec. 5.3). By appropriately filtering the event output, it was possible to track the slowing-down process of the spallation neutrons, recording their exact position and energy at the moment their kinetic energy dropped below the 20 MeV threshold.

3.2.5 Spatial Meshing and Tally Mapping

In the investigated ADS design detailed in Sec. 4.1, the active core region is composed of hexagonal fuel assemblies (FAs), arranged in a triangular lattice. Consequently, the ERANOS model adopts a discrete hexagonal mesh in the radial plane.

To perform a consistent neutronic analysis and extract spatially distributed physical quantities using MCNP, it is desirable to evaluate the tallies on an equivalent three-dimensional grid. However, the central region of the reactor, housing the accelerator and spallation modules, is modelled using CSG, combining continuous cylindrical and spherical surfaces (see Sec. 4.1.3). Only the outer region is organised in a hexagonal structure.

To perform precise neutronic scoring in the target region that perfectly maps onto the ERANOS mesh, a uniform three-dimensional grid could be superimposed over the continuous target geometry. Modifying the target geometry to impose an explicit hexagonal partition is highly impractical, as it would require manually defining new surfaces and redefining dozens of complex geometric cells with Boolean combinations. Instead, a spatial overlay strategy must be implemented to tally physical quantities directly onto an independent grid.

Focusing first on the radial discretisation, MCNP offers several features to superimpose grids over an existing geometry. However, each presents specific limitations regarding hexagonal coordinates:

- **FMESH and TMESH:** MCNP provides dedicated mesh tally capabilities (Flux Mesh and Tally Mesh) that allow scoring physical quantities on an independent, superimposed grid without altering the underlying geometry. While highly efficient, these features only support Cartesian, cylindrical, and (in the case of TMESH) spherical coordinate systems. They do not natively support hexagonal meshes, rendering them unsuitable for the precise geometric mapping required for this analysis.
- **Unstructured Meshes (UM):** MCNP allows the importation of finite element meshes generated by external CAD software. In this hybrid mode, particles are tracked directly inside the UM elements, theoretically allowing flux mapping over any arbitrary geometry. However, a fundamental topological restriction in MCNP dictates that the bounding surfaces of the background CSG cells cannot be intersected by the UM surfaces. Since a superimposed hexagonal grid would inevitably cross the boundaries of the internal cylindrical target cells, this approach is topologically invalid for this specific application.

Given these limitations, a geometric workaround based on the LAT (Lattice) card was developed. Typically, the LAT card is used within the MCNP universe structure




Figure 3.3: Planar section of an erroneous lattice superposition in MCNP. (a) Native CSG geometry consisting of a macroscopic bounding box containing a sphere. (b) Erroneous outcome when attempting to directly fill the macroscopic cell with a square lattice: the code replicates a localised spatial frame across all lattice elements instead of superimposing the grid over the continuous object.

to define periodic geometries, such as repeating fuel pins within an assembly. If one simply attempts to superimpose a lattice onto a universe containing the entire target, MCNP does not create a spatial grid over the geometry. Instead, it frames a local portion of the underlying universe and repeats that single sub-frame across every lattice element, completely destroying the original macroscopic geometry.

As an illustrative example, Fig. 3.3 demonstrates the erroneous outcome of this superposition. When a macroscopic bounding box containing a central sphere (Fig. 3.3a) is directly filled with a square lattice, MCNP does not project the grid over the continuous sphere. Instead, it captures a localised spatial frame and identically replicates that arbitrary fraction of the sphere across every individual lattice element, resulting in a fractured, unphysical geometry (Fig. 3.3b).

To achieve a true spatial mapping without altering the physical model, instead of attempting to superimpose a single global grid, a cell-by-cell lattice approach was adopted. First, the cells of interest were redefined as void cells by assigning a material number of 0. Each cell was then filled with a new, infinite universe containing a lattice of hexagonal prisms, which in turn were assigned the cell's original material composition.

Crucially, all these individual lattices were defined using exactly the same spatial origin and hexagonal pitch. Because they share the same structural grid, these independent lattices perfectly align with one another across the boundaries of the cells. Visually, they form a single, continuous macroscopic hexagonal grid embedded within the entire core geometry.

Fig. 3.4 illustrates the result of this lattice approach. As shown in Fig. 3.4a, the native geometry consists of continuous concentric cylinders defined by four distinct




Figure 3.4: Planar section of the example geometry. (a) The original MCNP geometry consisting of continuous concentric cylindrical cells. (b) The same region after the embedded lattice technique. Explicit lattice indices $([i, j, k])$ are shown.

radial cells (orange, green, light blue, and dark blue). By embedding a lattice within each cell (Fig. 3.4b), it becomes possible to isolate specific spatial regions using explicit lattice indices $([i, j, k])$. For instance, the tally for the central hexagonal prism can be obtained by summing the scores of the central element $[0, 0, 0]$ across all four lattices.

This robust approach guarantees that the exact continuous geometry of the target is preserved for the MC transport algorithms, while simultaneously allowing for the precise extraction of tallies in the hexagonal geometry. The hexagonal spatial mesh applied to the target of the investigated ADS is illustrated in Sec. 4.1.4.

Once the planar hexagonal structure is successfully established via the cell-by-cell lattice implementation, the three-dimensional grid is completed by introducing the axial segmentation. Rather than artificially splitting the lattice cells along the z -axis, this is seamlessly handled by leveraging the FS (Tally Segment) card. This card allows any given cell to be divided into discrete segments for tallying purposes. By specifying a sequence of z -plane surfaces in the FS card, the tally scores are segmented axially, yielding the final three-dimensional spatial mesh.

Chapter 4

Description of the ADS Model

This chapter focuses on the computational model of the investigated ADS and its preliminary physical characterisation.

In the first section, the investigated ADS is presented in detail. The system's physical features are described, along with the ERANOS and MCNP computational models.

The second section presents the preliminary physical analyses performed using the MCNP model, with particular focus on the core neutron flux and proton distribution.

4.1 ADS Overview and Computational Models

4.1.1 Description of the ADS Core

The nuclear system investigated in this study is an ADS designed to operate as a subcritical fast reactor. The analysis focuses on the start-up configuration [20], specifically evaluating the core at the beginning of life with fresh fuel.

The fuel cycle strategy adopted for this core is oriented towards a plutonium-burner configuration using spent MOX fuel. While ADS concepts are traditionally envisioned as dedicated MA transmuters, fuels with high MA content would represent an additional technological challenge, particularly from the perspective of fuel fabrication and irradiation behaviour. Since the ADS itself already represents an innovative reactor concept, a more conventional and well-established fuel was selected for the start-up configuration.

The initial fuel derives from the discharge of spent MOX fuel from PWRs with a nominal burnup of 45 MWd/kg_{HM}. Before refabrication, the plutonium composition was allowed to age for 12 years. This period includes 9 years of SNF cooling followed by 3 years required for the refabrication process. The MAs are entirely partitioned and removed prior to refabrication. However, ²⁴¹Am is generated from the β decay of ²⁴¹Pu during the 3-year refabrication window, resulting in a final fraction of approximately 1.3%.

The subcritical core is arranged radially around the central accelerator and spallation modules. The active core comprises a total of 108 FAs. The MOX fuel pellets are enclosed within stainless-steel cladding. The fuel pins in each assembly are arranged in a triangular lattice and enclosed within a hexagonal steel wrapper. The




Figure 4.1: Radial $\pi/3$ sector of the ERANOS core model. The colours map to distinct representative meshes: target/accelerator region (white), inner and outer fuel assemblies (yellow and orange, respectively), dummy assemblies (grey), Pb coolant (peach) and inner vessel (blue).

active height (AH) of the fissile column is approximately 1 m, and fuel is designed to withstand a maximum burnup of 100 MWd/kg_{HM}.

To achieve radial power flattening and compensate for the neutron flux drop at the core periphery, the active core is divided into two distinct radial zones. The inner region has a plutonium enrichment in the MOX pellets of approximately 20%, increasing to 25% towards the core periphery.

The system in the start-up configuration is designed to generate a nominal thermal power of 250 MW. The primary coolant circulating in the main vessel is pure liquid lead. Given the steady-state nature of the neutronic analysis, the dynamic thermal-hydraulic circulation is not explicitly simulated. Instead, the axial temperature gradient of the coolant is approximated by discretising the lead volume into three distinct axial regions, fixed at temperatures of 520, 460, and 400 °C, corresponding to the axial levels of the core inlet, fissile zone, and core outlet, respectively. A liquid metal spallation target is positioned at the centre of the subcritical core and is bombarded by a high-energy proton beam at 800 MeV.

The entire primary system is enclosed by an inner vessel made of stainless steel, featuring an internal diameter of approximately 3 meters, while the total length of the assemblies reaches roughly 8 meters.

4.1.2 ERANOS Model

As explained in Chap. 3, the MCNP simulation performed in this work is strictly linked to the ERANOS model for the ADS characterisation.

Fig. 4.1 shows the ERANOS model of the core in the radial section. The ADS




Figure 4.2: Planar view of the full-core MCNP model at the mid-plane ($z = 0$). The outer regions, including the Pb coolant, are discretised using a uniform lattice structure. The central target breaks this symmetry and is modelled as a distinct cylindrical structure.

core is modelled using $\pi/3$ rotational symmetry. For the deterministic calculation, the planar mesh geometry employs a discrete hexagonal grid. The radial dimension of each mesh element corresponds to the hexagonal steel wrapper enclosing a single FA, together with a coolant gap of a few millimetres between adjacent wrappers (designed to accommodate their radial expansion). In the central position, the accelerator and spallation modules occupy the first three rings of the lattice. Moving radially outward, the inner and outer fissile regions cover two rings each. The active core is then surrounded by dummy assemblies, which act as a reflector to optimise neutron economy and shield the inner vessel. Finally, external Pb coolant and the core vessel complete the radial layout. As discussed in the previous chapter, it is important to note that the ERANOS code employs homogenised representative meshes. The material composition of each cell is a volume-weighted mixture of the actual physical components within that region, providing a simplified but physically consistent approximation of the exact geometry.

ERANOS calculations were performed on a single $\pi/3$ sector, imposing rotational periodic boundary conditions on the sector edges, in order to reproduce the results of the full-core calculation while significantly reducing the computational cost.




Figure 4.3: Detailed planar view of the central region in the MCNP model. The central spallation target (dark blue) is surrounded by beam tube structures, a region of lead coolant (light blue), and the innermost rings of FAs.

4.1.3 MCNP Model

Fig. 4.2 shows the planar view of the full-core ADS model at the mid-plane ($z = 0$ in the MCNP model), generated using the MCNP Visual Plotter [21]¹. It can be observed that the core is arranged in a lattice of hexagonal elements. All regions surrounding the centre, including those containing only Pb coolant, are disposed inside the lattice structure to simplify the generation of the MCNP input file. However, as previously anticipated, the central target region breaks this symmetry and is modelled as a separate structure.

A closer view of the central region is presented in Fig. 4.3. Each cell colour represents the corresponding assigned material. Surrounding the central target is the Pb coolant, followed by the two rings comprising the inner-fuel FAs.

A detailed view of a single FA (specifically, an outer fuel assembly) is shown in Fig. 4.4. The FA bundle comprises 126 fuel pins arranged around a central dummy, which is an empty clad reserved for instrumentation purposes.

The active core is surrounded radially by dummy reflector assemblies having the same structural architecture as the FAs, but containing ceramic reflector pellets within the steel clads instead of fissile material. A detailed view of the core's radial periphery, highlighting the interface between the fuel and the dummy rings, is

¹Differently from the subsequent plots, Fig. 4.2 is in black-and-white, since the Visual Plotter could not handle the resolution of each cell in detail.




Figure 4.4: Detailed view of a single FA. The hexagonal wrapper encloses a bundle of 126 fuel pins (red) and a single central dummy pin (yellow).




Figure 4.5: Detailed view of the core periphery. FAs (red pins) are surrounded by dummy assemblies (yellow pins). Beyond them lie the Pb coolant (light blue) and vessel (grey).

provided in Fig. 4.5.

Finally, Fig. 4.6 shows the axial section of the ADS core. This view clearly distinguishes the beam pipe (in light orange) modelled as a low-density He void,




Figure 4.6: Axial view (x - z plane) of the MCNP model. The beam pipe (light orange) terminates in a conical window above the liquid metal spallation target (dark blue). Surrounding the central region, lead coolant zones (orange, light blue and green), FAs and dummy assemblies are visible.

and the corresponding conical beam window facing the liquid metal target (in blue). Furthermore, the three distinct Pb coolant axial regions, corresponding to the inlet, fissile zone, and outlet temperatures (orange, light blue and green, respectively) are visible.

The z axial coordinate of the proton beam source was positioned approximately 120 cm above the core mid-plane. A Gaussian spatial distribution was applied to the x and y coordinates of the beam profile. The full width at half maximum (FWHM) in both x and y directions, which is a required MCNP source parameter, was set to 5.65 cm. This value was derived from the standard deviation σ of the Gaussian profile at the point where the beam enters the core, which is 2.40 cm, according to the following relation:

$$FWHM = \sqrt{8 \ln(2)} \sigma. \quad (4.1)$$

However, as will be analysed in Sec. 4.2.2, this specific beam configuration, inherited from previous core characterisation studies, implies that a significant fraction of the source protons will inadvertently collide with the accelerator components prior to reaching the target, given that the beam pipe internal radius is about 5 cm.

Nuclear interactions induced by protons were simulated using the Bertini INC model, coupled with the Dresner evaporation model for residual nucleus de-excitation after the cascade. The Bertini model [22] treats the nucleon-nucleus interaction as a sequence of quasi-free nucleon-nucleon collisions within a free-electron-gas nuclear




Figure 4.7: Planar view of the central region in the MCNP model, illustrating the cell-by-cell lattice. (a) Cross-section at the core mid-plane, traversing the spallation target (dark blue). (b) Cross-section at a higher axial elevation, traversing the vacuum beam pipe (light orange). Although these planar sections traverse distinct physical cells within the central module, the global hexagonal mesh remains perfectly continuous and strictly aligned with the outer core lattice.

model, propagating fast cascade particles until their energy falls below a cutoff, after which the excited residual nucleus is de-excited via evaporation of light particles. Moreover, the continuous energy loss of charged particles due to Coulomb interactions with atomic electrons is simulated by MCNP using the condensed-history algorithm, in which the particle track is subdivided into energy-loss substeps between two subsequent events.

4.1.4 Target Spatial Meshing

Prior to conducting the preliminary neutronic analyses on the investigated ADS core, a radial hexagonal spatial mesh was applied to the central region of the MCNP model. This was achieved using the cell-by-cell lattice approach detailed in Sec. 3.2.5. Specifically, each physical cell within the central module was redefined to be filled by an infinite lattice of hexagonal prisms. The pitch and orientation of this central lattice were explicitly defined to perfectly match the geometry of the surrounding active core.

The resulting geometry is illustrated in Fig. 4.7. While the exact physical boundaries and material definitions of the target structures are preserved, the resulting core appears as a continuous macroscopic hexagonal grid. Because the lattice elements are defined as infinite hexagonal prisms, MCNP automatically truncates the lattice at the physical boundary of each containing cell. Globally, this geometric construct ensures that the same planar mesh layout is maintained at every axial elevation. This consistency is evident when comparing the mid-plane cut

(Fig. 4.7a, traversing the target) with a section taken at a higher elevation (Fig. 4.7b, traversing the vacuum beam pipe). Even though these planes cross different physical cells and materials, the hexagonal mesh remains geometrically identical.

Finally, the three-dimensional grid is completed by introducing axial discretisation through the FS card.

Ultimately, this spatial meshing strategy enables the rigorous extraction of neutronic tallies, allowing the flux distribution in the central region to be evaluated on the same hexagonal grid used by the deterministic ERANOS solver.

4.2 Preliminary Analyses

Before the development of the neutron source, preliminary MC simulations were conducted to characterise the baseline physical behaviour of the system. All simulations were performed using the CRESCO8 HPC infrastructure hosted by ENEA. Further details on the computational resources are provided in Appendix A. These initial evaluations focused on two main aspects: the evaluation of the neutron flux distribution using dedicated cell tallies and the detailed characterisation of the proton beam.

4.2.1 Flux Distribution

The first phase of the analysis focused on the evaluation of the steady-state neutron flux distribution throughout the reactor core. Relying on the hexagonal mesh described in the previous chapter, MCNP flux tallies were extracted for every discrete spatial cell.

Fig. 4.8 illustrates a three-dimensional heatmap of the integrated neutron flux mapped onto a $\pi/3$ sector of the core. The simulation was performed with a total of 10^5 primary spallation protons, explicitly transporting neutrons, protons and positive pions, in accordance with the MCNP manual recommendations for default spallation problems. To improve the statistical reliability of the tallies, the results were azimuthally averaged by exploiting the six-fold symmetry of the core. Specifically, for any given position on the radial plane outside the central axis, the total cell flux is obtained by summing the scores of the six symmetrically equivalent hexagonal assemblies across the full core and dividing by six. This summation is performed directly at the tally level by specifying the appropriate lattice cell indices within the MCNP tally definition, rather than as a post-processing step. Consequently, the code accumulates the scoring contributions from all six symmetric assemblies into a single tally during the transport simulation. This approach ensures that the relative statistical error is computed directly by MCNP, eliminating the need for further error propagation. Finally, the resulting cell-averaged neutron flux is then normalised per incident source proton and per unit mesh volume. As shown in the figure, the mesh volume is non-uniform, with cell heights varying along the axial direction of the core. Consequently, volume normalisation is essential to accurately compare the neutron flux across different layers. The evaluated portion of the $\pi/3$ core sector corresponds exactly to the ERANOS deterministic geometric model, comprising 56




Figure 4.8: Volumetric spatial distribution of the azimuthally averaged total neutron flux within a $\pi/3$ sector of the core. The evaluated domain matches the ERANOS spatial grid (56 assemblies $\times$ 47 axial layers). The simulation was performed with 10^5 spallation protons. Tally values are normalised per incident source proton and mesh cell volume.

radial assemblies and 47 axial layers, excluding the outer axial regions present in the full-core MCNP model.

As expected for an ADS, the neutron flux distribution exhibits a pronounced peak in the spallation target region, located slightly below the core mid-plane. In this central zone, high-energy primary protons interact with the liquid metal target, triggering the spallation cascade and emitting many fast neutrons. Moving radially from the central target towards the core periphery, the flux exhibits a rapid spatial attenuation. Reaching the sixth ring of the lattice (which corresponds to the second ring of outer fuel), the flux decreases to between 10% and 30% of the peak value observed in the target region. This rapid attenuation is a direct consequence of the subcritical nature of the system.

Fig. 4.9 displays the corresponding relative statistical error of the MC simulation. Every tallied cell within the core sector exhibits a relative error below 5%, which is well within the MCNP recommended threshold for reliable tally statistics, confirming the statistical robustness of the results. The uncertainty increases in the peripheral regions, due to the attenuated flux, while in the highly populated central target




Figure 4.9: Spatial distribution of the relative statistical error (1σ) associated with the MCNP integrated flux tally. The uncertainty remains well below 5% in the active fuel regions, dropping below 1% in the central spallation target.

region the error drops well below 1%.

In addition to the spatial profile, the energy distribution of the neutron flux integrated over the entire core was evaluated. Rather than employing an arbitrary energy grid, the MCNP tally was configured using the exact energy binning of the 51-group structure adopted for the deterministic ERANOS calculations. This specific choice ensures that this preliminary core-integrated neutron spectrum can be directly compared with the neutron source spectrum presented in the next chapter. The resulting neutron spectrum, expressed as neutron flux per unit lethargy, is presented in Fig. 4.10. The distribution clearly exhibits the hard-neutron profile characteristic of lead-cooled fast reactors, with the flux reaching its maximum in the high-energy range. Specifically, the spectrum has a median energy of 61 keV and a mean neutron energy of 259 keV, reflecting the contribution of the pronounced high-energy tail extending up to several MeV. The statistical error associated with the spectrum remains below 1% for all energy bins.




Figure 4.10: Energy distribution of the neutron flux integrated over the entire reactor core, expressed per unit lethargy. The spectrum highlights the hard-neutron profile typical of lead-cooled fast reactors.

4.2.2 Proton Beam Characterisation

The second phase of the preliminary analysis investigated the transport and interactions of the 800 MeV primary proton beam. In an ADS, the subcritical fission chain is sustained by the neutron yield generated by the high-energy beam impinging on the heavy liquid metal target. Therefore, an accurate characterisation of the beam behaviour within the system is essential.

This preliminary assessment focused exclusively on tracking the primary source protons and the secondary protons generated by subsequent nuclear interactions, using the MCNP PTRAC feature. The analysis of the spatial distribution of proton interactions and terminations is crucial for validating the target geometric design and assessing the effective coupling of the accelerator with the subcritical core.

The spatial distribution of the primary proton terminations (TER events) is illustrated in Fig. 4.11, showing projections on the x - y and x - z planes for a simulation of 10^5 primary protons. Furthermore, to better evaluate the spatial distribution of these terminations with respect to the target modules, the x - z points are superimposed onto the MCNP geometry in Fig. 4.12. The x - y projection clearly highlights the cylindrical symmetry of the beam, characterised by a dense central halo and a gradual radial decrease in the density of termination points. Conversely, the x - z projection reveals a more complex axial behaviour. As expected, the vast majority of the termination points occur in the lower half of the geometry, starting just below the beam window and extending downwards into the liquid metal target. However, a significant fraction of termination points is observed in the upper half of the geometry, starting immediately below the beam injection plane ($z = 120$ cm). Furthermore, a distinct, wide band of termination points is clearly visible at approximately $z = 78$ cm.




Figure 4.11: 2D spatial projections (x - y and x - z planes) of the primary proton termination points (TER events). The simulation was performed with 10^5 primary protons.




Figure 4.12: Axial (x - z) projection of the primary proton termination points superimposed onto the MCNP geometric model of the central target module.

To understand this phenomenon, a detailed investigation of the primary proton generation (SRC events) and termination (TER events) coordinates was conducted, as illustrated in Fig. 4.13. Fig. 4.13a shows the x - y projection of source events superimposed on the MCNP geometry, while Fig. 4.13b shows the projection of the




Figure 4.13: Detailed radial (x - y) projections highlighting the primary proton beam interactions in the upper geometry. (a) Generation points (SRC events) of the primary protons. (b) Termination points (TER events) of the primary protons restricted to the region above $z = 80$ cm.

termination points considering only the region above 80 cm from the mid-plane. It was found that approximately 9.6% of the source protons originate outside the radial boundary of the vacuum beam pipe, as a direct consequence of the large standard deviation of the source spatial distribution defined in the MCNP input file. Consequently, these protons immediately collide with the liquid metal of the upper target module (as shown in Fig. 4.13b). An analysis of the interaction types [23] associated with these specific termination points revealed that 99.8% correspond to nuclear interactions. On the other hand, the termination band observed between 76 cm and 78 cm is dominated (94% of events) by particles reaching the energy cut-off (set to 1 keV, the default minimum energy for proton transport in MCNP). A similar energy cut-off band is observed in the lower region, roughly 40 cm below the beam window. This analysis highlights that a non-negligible fraction of the beam is lost before reaching the intended spallation region near the core mid-plane. Further investigations should be carried out to quantify how this beam loss affects the overall reactor neutron economy.

The analysis was subsequently extended to evaluate the secondary protons generated during the spallation cascade. Simulations with 10^5 source protons were performed both with and without fission reactions, in order to isolate the direct contribution of the spallation process from the effects of neutron multiplication in the core. Fission was disabled in MCNP using the NONU card, which treats fission events as simple neutron captures. Fig. 4.14 and Fig. 4.15 illustrate the termination points of primary (blue) and secondary (magenta/orange) protons for the NONU and fission-enabled cases, respectively.

Comparing the two scenarios highlights the critical role of fission reactions in secondary particle generation. Considering the x - y plane, in the NONU case, secondary proton terminations are mainly localised within the central target and the




Figure 4.14: 2D spatial projections (x - y and x - z planes) of primary (blue) and secondary (magenta) proton termination points in the NONU case, where fission reactions are treated as simple captures to isolate the spallation contribution.




Figure 4.15: 2D spatial projections (x - y and x - z planes) of primary (blue) and secondary (orange) proton termination points with fission reactions enabled, highlighting the subcritical multiplication effect.

first two surrounding rings, with limited occurrences in the active core. Furthermore, the axial projection reveals that these terminations are confined to the lower target region and in the upper part of the geometry, where primary protons are generated outside the beam pipe. Conversely, when fission is enabled, secondary proton terminations expand widely in the core. Radially, they trace the exact geometry of




Figure 4.16: Energy spectrum of the secondary protons at the moment of their generation (BNK events), expressed as the fraction of particles per unit lethargy. The comparison highlights the spectral shift induced by fission reactions (orange) compared to the pure spallation cascade (magenta).

hexagonal assemblies of the active core, while axially, they extend to form a thick halo that corresponds to the active fuel region. This widespread distribution is governed by the intense flux of fission neutrons. As these neutrons propagate through the core, they may induce nuclear interactions within its structural materials. The generated low-energy protons undergo rapid energy loss and travel very short distances before terminating, thereby spatially mapping the active fuel regions.

While contributing to a more widespread spatial distribution of secondary protons, fission reactions also induce a clear shift in their generation energy spectrum (recorded as BNK events). This effect is illustrated in Fig. 4.16, with the spectrum expressed as the particle fraction per unit lethargy. In the NONU case (magenta curve), the secondary protons are generated predominantly by the high-energy spallation process, resulting in a total population of approximately 390,000 particles. Consequently, the energy distribution is broad and shifted towards high energies, featuring a main peak around 8-9 MeV, a secondary peak at 20-30 MeV, and a pronounced high-energy tail, resulting in a mean energy of 67.4 MeV and a median energy of 25.0 MeV. When fission is enabled (orange curve), the overall secondary proton population increases to roughly 490,000, including almost 100,000 protons generated by fission neutrons. The secondary protons generated by fission neutrons are constrained to low energies, altering the spectrum's shape. This yields a peak around 3 MeV and significantly reduces the mean and median energies of the overall proton population to 44.5 MeV and 10.1 MeV, respectively.

Chapter 5

Neutron Source Development and Results

This chapter presents the central development of this study: the practical implementation and validation of the hybrid MCNP-ERANOS computational methodology. The main objective is the rigorous extraction of the spallation neutron source from MCNP simulations.

First, the specific source term requirements dictated by the ERANOS model are detailed. Second, the physical and computational principles of the source extraction are analysed. Particular attention is given to the MCNP tracking logic, which necessitates the reconstruction of the particle genealogy. The core of the chapter then illustrates the architecture of the post-processing algorithm, developed to track and extract the source events.

Subsequently, the generated source terms are presented. Finally, the reliability of the developed methodology is assessed by comparing the main macroscopic ADS parameters obtained with ERANOS against those from a full-core MCNP calculation.

5.1 ERANOS Requirements

This study aims to establish a robust methodology for implementing a hybrid computational strategy, ultimately generating a neutron source for the deterministic ERANOS code using MCNP. Specifically, the objective is to extract a source distribution comprising neutrons at the moment they enter the energy domain just below the 20 MeV threshold. Once defined, this distribution serves as the external source for the subsequent ERANOS calculations. This extraction is achieved by simulating the high-energy spallation process and tracking particles above 20 MeV using the MCNP PTRAC feature, while discarding all subsequent events below this threshold, as they fall within the energy domain managed by the subsequent ERANOS simulations.

The external source required by the ERANOS model consists of three distinct components:

1. N_s : The average number of neutrons per incident proton entering the energy domain below 20 MeV;

2. $\mathbf{s}(\mathbf{i}, \mathbf{z})$: The volumetric density (cm^{-3}) of these source neutrons within each spatial mesh element, restricted to a specific portion of the ERANOS model comprising 46 radial assemblies (i) and 47 axial layers (z);
3. **Representative spectrum**: A representative energy spectrum structured into the predefined binning of 51 energy groups. To accurately represent the energy distribution of neutrons entering the core, the spectrum is extracted from the second ring of the core lattice (located at the interface between the target and the first ring of fuel assemblies) and from an axial segment corresponding to the fuel AH. Consequently, the adopted approach assumes a single spatially independent source spectrum.

It is crucial to highlight that the evaluated neutrons are strictly those that have *just* crossed the energy threshold. Any neutrons generated by subsequent low-energy reactions (below 20 MeV) are excluded from this source extraction.

5.2 Principles of the Source Extraction Algorithm

Before detailing the implementation of the post-processing algorithm, it is essential to first illustrate the physical and computational principles of the source extraction scheme. The entire procedure relies on two fundamental conditions:

- Recording the phase-space coordinates of a neutron at the exact moment its energy drops below the 20 MeV threshold;
- Discarding all subsequent events originating from a neutron that has already crossed this threshold.

Fig. 5.1 provides a simplified visual representation of this logic, illustrating generic interaction pathways during transport. This diagram is intended purely for illustrative purposes and shows only a limited number of secondary particles and simplified event sequences. The vertical axis represents the particle energy, with the 20 MeV threshold (red dashed line) dividing the MCNP transport domain from the ERANOS domain. From a transport perspective, a spallation reaction is treated as a standard interaction: it consists of a COL event immediately followed by a TER event sharing identical spatial coordinates. The secondary neutrons (orange dashed arrows), protons (blue dashed arrows), and possibly pions, are generated exactly at this location. However, as indicated by the dashed arrows leading to the BNK events, each secondary particle is generated with a specific initial energy for its subsequent transport.

Starting from the primary spallation, the energy evolution of each secondary particle must be monitored during transport. For instance, as depicted on the left side of Fig. 5.1, a neutron may undergo multiple scattering collisions that progressively degrade its energy. The individual particle must be tracked for as long as it remains within the high-energy domain ($E > 20$ MeV). Once the particle drops below this energy threshold, its energy and spatial coordinates are recorded to form the external source (red dots), and any subsequent transport of this particle or its progeny is




Figure 5.1: Simplified event sequence diagram illustrating the source extraction logic. The MCNP transport domain ($E > 20$ MeV) is separated from the ERANOS domain by the energy threshold (red dashed line). Extracted source neutrons are highlighted as red dots.

completely discarded from the analysis. The crossing of this threshold can occur either as the result of a degrading COL event, or immediately upon the generation of a low-energy BNK event (as illustrated by the two central neutrons in Fig. 5.1). Conversely, if a particle type different from neutron (such as a secondary proton or pion) drops below 20 MeV, its tracking should continue, as the threshold applies strictly to the generation of the neutron source.


Furthermore, it is important to clarify that the energy and phase-space coordinates associated with a COL event represent the particle's state after the collision. Therefore, if a COL event causes a neutron to drop below the threshold while also producing secondary particles, the interaction itself occurred while the neutron was still in the pre-collision energy state above 20 MeV. Consequently, any secondary particles produced during this specific collision originate from a high-energy interaction and must be properly tracked and evaluated for the source term.

Once the physical logic of the extraction is defined, the implementation must account for the MCNP transport mechanics. Specifically, it is necessary to recall the MCNP LIFO (Last In, First Out) transport logic introduced in Sec. 3.2.2, and to understand how it affects the particle tracking analysis.

During MCNP particle transport, after the termination of a particle track, the most recently generated secondary particle (placed at the top of the memory stack) is the first to be simulated. Previously stored particles remain in the bank until the stack is progressively cleared. In the context of this study, after a primary proton undergoes spallation and terminates, the code begins transporting one of the

	nps	type
0	1	source
1	1	collision
2	1	termination
3	1	bank
4	1	collision
5	1	termination
6	1	bank
7	1	collision
8	1	termination

(a) Output event sequence



(b) Genealogy ambiguity

Figure 5.2: Representation of the PTRAC output data structure for a single source particle history ($nps = 1$). The event sequence (a) creates ambiguity in establishing parent-child relationships for banked secondary particles (b).

generated secondary particles. This initiates a new branch in the transport tree: the code simulates the transport of that specific particle and all subsequent generations arising from its interactions, rigorously following the LIFO logic. Only after this entire branch has been exhausted does the code return to simulate the transport of the following secondary particles produced in the initial spallation event.

However, the PTRAC output file lacks direct, reliable information regarding parent-child relationships between particles. Consequently, the output appears as a linear sequence of transport events, where each event is characterised by a set of phase-space variables but is completely devoid of any explicit hierarchical links.

An illustrative example of this event sequence is shown in Fig. 5.2. Fig. 5.2a illustrates a sequence of events extracted for a single history number ($nps = 1$), visualised using the MyHDF5 web application [24]. For each listed event, the user can access its associated variables. The critical challenge in this transport flow is that any banked secondary particle (BNK event) could potentially originate from any preceding generative event within the same history (typically a COL event, but also a TER event in cases of non-collisional processes like particle decay). As highlighted in Fig. 5.2b, a BNK event at index 3 or 6 could be attributed to multiple preceding collisions or terminations. Therefore, to correctly define the neutron source, it is necessary to explicitly reconstruct the particle genealogy. Without identifying the




Figure 5.3: Flowchart of the post-processing algorithm developed to process the PTRAC output, reconstruct the particle genealogy, and extract the neutron source.

specific parent event, it is impossible to determine whether a banked secondary neutron originates from a high-energy cascade (and thus should be evaluated for the source) or from a branch that has already dropped below the 20 MeV threshold (and must be discarded). Establishing these parent-child relationships requires matching the spatial and temporal coordinates of each BNK event with those of previous COL or TER events. In a complex system like an ADS, featuring thousands of events per history, reconstructing this genealogy becomes computationally challenging.

5.3 Algorithm Implementation

Based on these physical and transport considerations, the extraction logic was implemented via a Python post-processing algorithm. As illustrated by the flowchart in Fig. 5.3, the script sequentially parses the entries in the PTRAC file until all particle histories are processed. The core of the algorithm relies on a switch-case logic that distinguishes events according to their type: Source (SRC), Collision/Termination (COL/TER), or Bank (BNK).

1. Primary Source Initialisation (SRC)

When the algorithm encounters an SRC event, a new primary particle history begins. Since particle histories are independent, the algorithm clears the memory array containing the "Active Nodes". Simultaneously, the state of the new branch is initialised and flagged as active.

2. Collision and Termination (COL/TER)

For COL and TER events, the algorithm first verifies the status of the current branch. If the branch is inactive, meaning that its parent branch has already crossed the energy threshold, the event is simply skipped. If the branch is active, the event's coordinates are saved in the memory array as an "Active Node". These nodes serve as a repository of potential parent coordinates for any subsequent BNK events. Following this, the algorithm checks if the particle involved is a neutron and if its energy has fallen below 20 MeV. If both conditions are met, the particle energy and its spatial coordinates are saved as part of the neutron source, and the branch is flagged as inactive. Otherwise, the branch remains active, and the next event is evaluated.

Crucially, storing the event's coordinates in the Active Nodes before verifying the particle's energy ensures that, if a generative collision causes a neutron to cross the 20 MeV threshold, the interaction node is still recorded. This ensures that any secondary particles produced during that specific collision will find a matching active node and be properly evaluated.

3. Secondary Particle Generation (BNK)

For BNK events, the algorithm must verify the parent's status to resolve the genealogy ambiguity. This is achieved by comparing the generation coordinates of the BNK event against the spatial coordinates stored in the "Active Nodes" array. If no match is found, the secondary particle must originate from a branch that has already been deactivated, and the new branch is immediately flagged as inactive. Conversely, if a match is found, the new branch must be tracked. The algorithm then evaluates the banked particle: if it is a neutron born with $E < 20$ MeV, its coordinates are saved into the neutron source and the branch is deactivated. If the neutron has a higher energy, or if the particle is a proton or a pion, the branch remains active, and the algorithm proceeds to evaluate the next event.

5.4 Parallel Implementation

The source extraction algorithm described in Sec. 5.3 is based on the sequential processing of the PTRAC output file. However, the file size can make a serial processing approach computationally prohibitive. For instance, a simulation with 10^5 source protons on the investigated ADS results in a PTRAC file of approximately 1 TB, and processing it with the algorithm would require more than five days.

To overcome this computational bottleneck, a custom Python-based workflow manager was developed using the mpi4py library [25]. The proposed approach

exploits MPI parallelism at two distinct stages: the parallel execution of the MCNP simulations and the subsequent post-processing of the PTRAC files. This allows the computational workload associated with both MCNP simulations and source extraction to be distributed across multiple CPU cores.

Work-Unit Decomposition The total number of particle histories is partitioned into independent sub-tasks by spawning N MPI processes, each assigned to a single CPU core. During the simulation stage, these processes act as independent task launchers rather than as MPI-communicating workers, and no message passing occurs between them. Each process creates a dedicated working directory and copies into it a private version of the input file, the cross-section file, and the execution script, allowing the N replicas to run concurrently without I/O conflicts.

Dynamic Parameterisation Each replica is assigned a disjoint block of N_h particle histories within MCNP's pseudo-random sequence. For the j th process, the starting history of the corresponding input file is set to $j \times N_h + 1$. This ensures that the N replicas sample non-overlapping portions of the random-number sequence, allowing their results to be combined as if they originated from a single simulation. This parameter substitution is performed automatically by each process on its private copy of the input file after the replica directory has been created.

Simulation Launch Each process launches a single-core serial instance of MCNP within its own working directory. Because no inter-process communication is required during the transport calculation, this scheme can achieve near-linear scaling with the number of allocated cores, avoiding the communication overhead associated with native MPI parallelism.

Source Extraction The simulation stage produces one PTRAC output file for each replica. These files are subsequently processed in a second, independent MPI job using a Manager-Worker scheme. Each Worker analyses its assigned PTRAC file with the source extraction algorithm, writing the neutron source events for each history to a Worker-local text file. All Workers then synchronise on an MPI barrier, that is a synchronisation point at which every process pauses until all others have reached it, ensuring that every local file has been completely written before proceeding. After this synchronisation, the Manager concatenates the local files into a single merged source file containing the complete set of neutron source events generated by the simulation.

5.5 Source Term Generation

Applying the post-processing algorithm to the MCNP PTRAC output successfully isolated and stored the neutron source events. The resulting dataset contains the spatial coordinates (x, y, z) and the energy E of each neutron crossing the 20 MeV energy threshold.




Figure 5.4: Spatial distribution of the volumetric neutron source density, mapped onto the selected $\pi/3$ reactor core geometry (46 assemblies $\times$ 47 axial layers). The colour scale represents the source fraction per unit volume, and every entry is normalised per incident source proton. The simulation was performed with 10^7 spallation protons.

However, to properly couple the MCNP calculation with the deterministic ERANOS calculation, an additional data-preparation step was necessary. The extracted data had to be processed and formatted to strictly comply with the three source terms required by the ERANOS input structure, as outlined in Sec. 5.1:

1. N_s : The average source neutron yield was obtained by simply dividing the total number of source neutrons extracted by the algorithm by the total number of primary spallation protons;
2. $s(i, z)$: The (x, y, z) coordinates of each extracted neutron were mapped onto the geometric grid of the ERANOS spatial model. By determining the fraction of source events per incident proton falling within each specific core assembly and dividing this value by the volume of the respective axial mesh, the volumetric density of the neutron source was calculated;
3. Representative spectrum: By filtering the spatial coordinates to select particles located in the second ring of the core lattice and within the axial segment


Source error distribution (37 assemblies x 32 layers, 1E7 source protons)


Figure 5.5: Relative statistical error distribution of the extracted neutron source density across a restricted core domain (37 assemblies $\times$ 32 layers). The relative uncertainties are estimated assuming Poisson statistics.

corresponding to the fuel AH, the neutron energies were binned into the predefined structure of 51 energy groups.

The MCNP calculation was performed by simulating 10^7 primary spallation protons. This resulted in an average source-neutron yield, N_s , of 20.79 neutrons per incident proton.

Fig. 5.4 shows the three-dimensional spatial heatmap of the volumetric source density, mapped onto the mesh grid of 46 radial assemblies and 47 axial layers. The neutron source was evaluated by neglecting the outermost ring of the ERANOS spatial model (the inner vessel). This exclusion is justified by the negligible contribution of those source neutrons to the overall neutron economy. As a reference value, the neutron mean free path in typical fast reactor cores is approximately 2 to 5 cm, depending on the specific energy [26], which is considerably smaller than the distance between the vessel and the active region.

Similarly to the neutron flux presented in Fig. 4.8, the source density exhibits a prominent peak in the spallation target region. However, the radial gradient of the source is much steeper: the volumetric density spans six orders of magnitude between the centre and the peripheral regions of the active core, whereas the integrated flux is confined to only two orders of magnitude. This sharp gradient is a direct consequence of the neutron source definition: the high-energy neutrons produced by the spallation process rapidly undergo nuclear interactions or multiple scattering collisions, losing energy and falling below the 20 MeV threshold well before reaching the outer regions.

Furthermore, the heatmap captures the effect of primary protons originating outside the beam pipe, with the uppermost hexagonal prisms exhibiting a noticeable concentration of source neutrons.




Figure 5.6: Representative energy spectrum of the neutron source, evaluated within the second core ring and the axial segment corresponding to the fuel AH. The distribution is expressed as the fraction of source neutrons per unit logarithmic energy.

To assess the statistical reliability of the extracted source, the relative error distribution of the neutron-source density for each mesh cell was evaluated. Assuming Poisson statistics, the relative error was estimated as $1/\sqrt{N_i}$, where N_i is the number of source events in the i -th cell. Following the same physical considerations regarding the neutron mean free path, this evaluation was restricted to a smaller and more relevant region of the core, comprising 37 assemblies and 32 layers. This domain covers 12 cm above and below the active region, and extends up to two rings outside the active zone (considering that the hexagon apothem is approximately 8.5 cm), thus ensuring a margin significantly larger than the neutron mean free path, particularly at the outermost ring.

As shown in the histogram in Fig. 5.5, the overall statistical uncertainties are well constrained. The mean relative error is 2.53%, while the median relative error is 1.77%. The maximum relative error reaches 16.01%, with only 11 cells exceeding the 10% threshold. Crucially, all of these 11 cells belong to the external ring. In this context, it should be noted that this unweighted error estimate does not account for the neutronic importance of each cell. Because the importance is significantly higher in the central mid-plane assemblies, where the source fraction is at its maximum, the higher statistical uncertainties observed in the peripheral mesh elements have only a minor impact on the overall calculation. Therefore, these results are perfectly consistent with the scope of this work, which aims to establish a robust methodology for source extraction rather than performing an exhaustive, high-precision neutronic characterisation of the reference system. Furthermore, if a higher precision calculation is required, the statistical error can be reduced by a factor of $\sqrt{N}$ by increasing the

number of simulated histories by a factor of N .

Finally, Fig. 5.6 displays the representative energy spectrum of the neutron source, here presented as the fraction of source neutrons per unit logarithmic energy. As expected from the physical definition of the sub-20 MeV source, it exhibits a very hard spectrum with a pronounced peak around 1 MeV, with a mean energy of 2.91 MeV. Compared to the full reactor neutron flux spectrum previously presented in Fig. 4.10, the low-energy tail below 100 keV is almost entirely absent.

5.6 Results

Table 5.1: Comparison of the reference ADS parameters evaluated with ERANOS and MCNP [20]. For N_s , the ERANOS value refers to neutrons that have just crossed the 20 MeV energy threshold, whereas the MCNP value accounts for all neutron energies.

	ERANOS	MCNP	MCNP - ERANOS
N_s	20.79	22.22	
M	28.45	25.70	
k_s	0.96604	0.96255	
φ^*	0.765	0.72	
k_{eff}	0.97381	0.97279	-102 [pcm]
ν	2.929	2.93	0.06 [%]
$N_s M$	591.43	571.21	-3.54 [%]
i_p [mA]	5.87	6.07	3.29 [%]

The extracted neutron source was provided to the ERANOS model, from which the main macroscopic parameters of the ADS were obtained. To validate the proposed hybrid methodology, these results were compared with those obtained from separate full-core Monte Carlo simulations performed directly with MCNP. The evaluated parameters are summarised in Table 5.1.

It is crucial to highlight that a direct comparison is physically meaningful only for a specific subset of these parameters, highlighted in bold in the table. The MCNP calculation models the complete high-energy spallation process triggered by the primary proton beam, whereas the ERANOS calculation relies on an externally defined neutron source. Consequently, parameters that depend directly on the source definition represent different physical quantities in the two calculations. Specifically, the source yield N_s , the source criticality factor k_s , the source importance φ^* , and the source multiplication M are not directly comparable. According to the MCNP source definition, N_s accounts for the net neutron production resulting from nuclear interaction events and (n, xn) reactions. In this context, nuclear interaction events refer to the reactions treated by the code using the high-energy nuclear models, whereas (n, xn) reactions are accounted for separately in the neutron balance. Neutrons produced through these processes, including those generated by subsequent (n, xn) reactions, contribute to N_s until they are either absorbed or escape from the system. Conversely, in ERANOS, N_s represents the average yield of

neutrons arising in the phase space with energy below 20 MeV, as provided by the MCNP extraction procedure.

In light of these definitional differences, the MCNP parameters were obtained through dedicated calculations. First, a criticality calculation was performed, providing an effective multiplication factor k_{eff} of 0.97279 and an average number of neutrons produced per fission (ν) of 2.93. The MCNP source yield N_s was then evaluated via a fixed-source calculation in which fission reactions were suppressed using the NONU card, thereby isolating the spallation contribution. MCNP6 provides the global particle balance in the default "problem summary" table, without requiring explicit user tallies, reporting the main neutron creation and loss mechanisms per incident proton. The value of N_s was obtained by calculating the net difference between neutron production through nuclear interaction events and (n, xn) reactions and the corresponding neutron losses due to absorption and escape.

Subsequently, a second fixed-source calculation was performed with fissions enabled. Applying the same particle balance from the problem summary, the total number of fission neutrons per incident proton (M) was derived from the overall net neutron production, which now fully accounted for the subcritical fission multiplication. The remaining source-dependent parameters were derived analytically from N_s and M . Finally, the proton beam current i_p was evaluated using Eq. 2.28. This calculation employed the thermal power P_{th} and the average energy per fission Q previously computed by the ERANOS code. Specifically, P_{th} was calculated as the power developed exclusively in the fissile zone of the core (yielding 244.5 MW), while Q was evaluated as the ratio between the fissile power and the number of fission reactions, resulting in 206.45 MeV.

As illustrated in Table 5.1, the comparison shows good agreement for all directly comparable quantities. The effective multiplication factor (k_{eff}) differs by only 102 pcm, while the relative difference in the average number of neutrons per fission (ν) is just 0.06%. For the ADS-specific quantities, the total number of generated neutrons per incident proton ($N_s M$) shows a relative difference of 3.54%. Most importantly, the operational proton beam current required to sustain the thermal power exhibits a relative difference of only 3.29%, yielding values of approximately 6 mA in both calculations.

The agreement is considered satisfactory, given that the comparison involves two codes based on fundamentally different computational approaches and physical approximations. For the MCNP calculation, statistical uncertainties are not reported because they are negligible. With 10^7 source protons, the relative statistical errors of N_s and M , assuming Poisson statistics, are below 10 pcm. Conversely, being based on deterministic methods, the ERANOS code does not introduce statistical uncertainties. Therefore, the observed deviations between the evaluated parameters do not arise from insufficient statistical precision in the MC simulation, but rather from the intrinsic methodological differences between the two codes.

These results demonstrate that the developed methodology provides a consistent and reliable procedure for transferring the neutron source from the high-energy MCNP simulation to the deterministic treatment of neutron transport below 20 MeV. The close agreement observed between the ERANOS calculation based on the extracted source and the full-core MCNP results provides strong support for the

validity of the proposed source extraction procedure.

The methodology developed in this work therefore offers a flexible coupling strategy for MCNP and ERANOS in the neutronic analysis of ADS. It can be applied to different reactor configurations and operating conditions, and can accommodate different energy thresholds depending on the requirements of the deterministic solver.

Conclusions

This thesis presented the development of a methodology for extracting the neutron source required for a hybrid Monte Carlo-deterministic neutronic analysis of an ADS.

This work was motivated by the complementary capabilities of the MCNP and ERANOS codes. While MCNP can accurately describe the high-energy spallation process and the resulting neutron transport, it cannot perform burnup calculations in fixed-source mode. ERANOS can instead be used for the required burnup analyses, but the available multi-group cross-section libraries extend only up to 20 MeV. A consistent procedure was therefore required to transfer the neutron source generated by MCNP to the deterministic model.

The investigated ADS was first analysed using the MCNP model. A spatial meshing strategy was implemented in the central target region to reproduce the spatial discretisation used by the ERANOS model. Preliminary MCNP simulations were then performed to characterise the neutron flux distribution and the transport of the primary proton beam. The neutron flux was evaluated on the adopted spatial grid and exhibited a pronounced concentration in the spallation target region, with a rapid spatial attenuation towards the outer core regions. The core-integrated spectrum showed a hard-neutron profile characteristic of lead-cooled fast reactors, with a mean neutron energy of 259 keV. The proton beam analysis characterised the spatial distribution of primary and secondary proton interactions and terminations within the system.

The main contribution of this thesis was the development of a dedicated post-processing algorithm for extracting the neutron source from the PTRAC output generated by MCNP. In order to define the external source for ERANOS, the algorithm was developed to track particles within the high-energy domain, identify neutrons crossing the 20 MeV energy threshold, and record their phase-space coordinates. Since the PTRAC output does not explicitly provide parent-child relationships between particles, particle genealogy was reconstructed to correctly identify and select the neutron source events.

The MCNP calculation resulted in an average source-neutron yield (N_s) of 20.79 neutrons per incident proton. The volumetric source density showed a pronounced concentration in the spallation target region and a steep radial gradient, while the representative spectrum exhibited a very hard distribution with a pronounced peak around 1 MeV and a mean energy of 2.91 MeV.

The resulting source was then provided to the ERANOS model to evaluate the main neutronic parameters of the investigated ADS, and the results were compared with dedicated full-core MCNP calculations. Good agreement was obtained between the two approaches for the directly comparable quantities. The effective multiplication

factor (k_{eff}) differed by 102 pcm, while the average number of neutrons produced per fission (ν) matched within 0.06%. The total number of generated neutrons per incident proton ($N_s M$) showed a relative difference of 3.54%, and the proton beam current (i_p) required to sustain the thermal power differed by 3.29%. The latter quantity was approximately 6 mA in both calculations. It should be noted that, in the full-core MCNP calculation, N_s accounts for neutron production over the full energy range, whereas in the ERANOS calculation it refers to neutrons crossing the 20 MeV energy threshold. Despite this difference, the resulting $N_s M$ values show good consistency between the two calculations, as $N_s M$ represents a characteristic quantity of the investigated ADS configuration.

These results show that the proposed methodology provides a consistent and reliable procedure for extracting and transferring the neutron source generated by an MCNP simulation to a deterministic calculation. The observed deviations between the evaluated parameters are associated with the different computational approaches and physical approximations employed by the two codes.

Further calculations could focus on simulating a larger number of particle histories to reduce the statistical uncertainty of the extracted source, thereby providing a more accurate input for the ERANOS calculation and potentially improving the agreement with the full-core MCNP results.

More generally, the developed procedure provides a flexible framework for applying the MCNP-ERANOS coupling to different ADS configurations and energy thresholds, depending on the requirements of the deterministic calculation. The agreement obtained for the directly comparable ADS parameters supports the validity of this approach as a practical basis for future hybrid MCNP-ERANOS analyses of ADS.

Appendix A

CRESCO HPC Infrastructure

A.1 Computational Resources and Data Storage

All simulations presented in this work were performed using the CRESCO8 HPC infrastructure hosted by ENEA (Italian National Agency for New Technologies, Energy and Sustainable Economic Development). The ENEA CRESCO8 (Computational RESearch centre on COmplex systems) cluster is a supercomputing facility comprising 13 racks and a total of 790 compute nodes (774 CPU nodes and 16 GPU nodes), delivering an aggregate memory capacity of 414 TB of RAM.

For the MCNP calculations performed in this work, the CPU partition was specifically targeted. The compute nodes are accessible from external networks via SSH through two dedicated front-end nodes. Each compute node is equipped with a dual-socket configuration featuring 64-core processors.

From a software perspective, SLURM is the job scheduler responsible for receiving user job requests, managing priority queues, and allocating computational resources.

To efficiently manage the demanding computational requirements and the massive data generation, the cluster relies on two distinct storage systems, optimising both global accessibility and Input/Output (I/O) performance. User environments and software installations are managed via the Open Andrew File System (OpenAFS). AFS is a distributed network file system that guarantees global access across all ENEA research centres, ensuring high versatility for general data management. Conversely, computationally intensive I/O operations are governed by Lustre. Lustre is a high-performance parallel file system designed to handle large-scale data transfers. Instead of writing a large output file to a single hard drive, Lustre splits the data into multiple files and saves them across multiple disks simultaneously. This parallel writing provides the extreme data-transfer speed required when thousands of compute cores write large files at the same time.

A.2 Parallelisation Strategy: OpenMP and MPI

To maximise computational efficiency, MCNP exploits a hybrid parallelisation strategy that combines the Message Passing Interface (MPI) and Open Multi-Processing (OpenMP) computing paradigms. This hybrid approach is specifically designed

to align with the physical hardware of the cluster, which consists of independent compute nodes, each configured with two CPU sockets.

MPI is the standard used for distributed-memory parallelisation. It handles the communication between different nodes over the network, and in this case it is also used to communicate between the two separate sockets of the same node. The parallelisation scheme implemented in MCNP through MPI follows a Manager-Worker architecture. At the beginning of the simulation, a single Manager process reads the input, loads the cross-section libraries from storage, and broadcasts the data to all available Worker processes. During execution, MCNP employs dynamic load balancing: instead of distributing all particle histories equally among the processes at the start, the Manager continuously assigns small batches of particle histories to the Workers. If a Worker finishes its task early, it immediately requests a new batch. This scheme ensures that no process remains idle while waiting for the others, avoiding synchronisation bottlenecks. At the end of the simulation, or during intermediate synchronisation steps (e.g., at the end of each criticality cycle), the Manager gathers all the partial statistical scores from the Workers to compute the global tallies and the final variance.

Conversely, OpenMP is the standard adopted for shared-memory parallelisation within each individual compute node, or even within a single socket. Unlike MPI, which is a library of explicit function calls, OpenMP relies mainly on compiler directives that instruct the compiler to distribute the workload across multiple threads (independent streams of instructions that run simultaneously on the local CPU cores). The MPI Worker assigned to a socket (or node) can generate multiple threads. Since all threads running on the same socket share the same physical memory (RAM), they can directly access the geometric model and nuclear cross-section libraries without duplicating the data. As a result, the cores do not spend time waiting for data to be exchanged over the network, as would otherwise be required between separate nodes or sockets.

This hybrid approach significantly minimises the overall memory footprint while maximising simulation speed.

Acknowledgements

The computing resources and the related technical support used for this work have been provided by CRESCO/ENEAGRID High Performance Computing infrastructure and its staff [27]. CRESCO/ENEAGRID High Performance Computing infrastructure is funded by ENEA, the Italian National Agency for New Technologies, Energy and Sustainable Economic Development and by Italian and European research programmes, see <http://www.cresco.enea.it/english> for information.

List of Acronyms and Abbreviations

General Acronyms

ADS Accelerator-Driven System

AH Active Height

CAD Computer-Aided Design

CDF Cumulative Distribution Function

CFC Closed Fuel Cycle

CPU Central Processing Unit

CRESCO Computational RESearch centre on COmplex systems

CSG Constructive Solid Geometry

EFIT European Facility for Industrial Transmutation

ENEA Italian National Agency for New Technologies, Energy and Sustainable Economic Development

ERANOS European Reactor ANalysis Optimized code System

FA Fuel Assembly

FIMA Fissions per Initial Metal Atom

FNR Fast Neutron Reactor

FOM Figure of Merit

FWHM Full Width at Half Maximum

GPU Graphics Processing Unit

HDF5 Hierarchical Data Format

HLM Heavy Liquid Metal

HLW High-Level Radioactive Waste

LIST OF ACRONYMS AND ABBREVIATIONS

HPC	High-Performance Computing
I/O	Input/Output
INC	Intra-Nuclear Cascade
LBE	Lead-Bismuth Eutectic
LIFO	Last In, First Out
LINAC	Linear Accelerator
MA	Minor Actinide
MC	Monte Carlo
MCNP	Monte Carlo N-Particle
MOX	Mixed Oxide
MPI	Message Passing Interface
MYRRHA	Multi-purpose hYbrid Research Reactor for High-tech Applications
OFC	Once-Through Fuel Cycle
OpenAFS	Open Andrew File System
OpenMP	Open Multi-Processing
PCFC	Partially Closed Fuel Cycle
PDF	Probability Density Function
PUREX	Plutonium Uranium Redox EXtraction
PWR	Pressurized Water Reactor
RAM	Random Access Memory
RF	Radio Frequency
SCK CEN	Belgian Nuclear Research Centre
SMR	Small Modular Reactor
SNF	Spent Nuclear Fuel
UM	Unstructured Mesh
UOX	Uranium Oxide

MCNP Terminology and Keywords

BNK MCNP PTRAC Bank Event

COL MCNP PTRAC Collision Event

FMESH MCNP Flux Mesh Card

FS MCNP Tally Segment Card

LAT MCNP Lattice Specification Card

PTRAC MCNP Particle Tracking Card

SRC MCNP PTRAC Source Event

SUR MCNP PTRAC Surface Crossing Event

TER MCNP PTRAC Termination Event

TMESH MCNP Tally Mesh Card

References

- [1] C. POINSSOT et al. “Recycling the actinides, the cornerstone of any sustainable nuclear fuel cycles”. In: *Procedia Chemistry* 7 (2012).
- [2] JAPAN ATOMIC ENERGY AGENCY (JAEA). *Nuclide Information: Am-241*. 2014.
- [3] A. F. HOLDSWORTH et al. “Spent Nuclear Fuel - Waste or Resource? The Potential of Strategic Materials Recovery during Recycle for Sustainability and Advanced Waste Management”. In: *Waste* 1.1 (2023), pp. 249–263.
- [4] H. H. SALEH et al. “Measurements of Delayed Neutron Decay Constants and Fission Yields from ^{235}U , ^{237}Np , ^{241}Am , and ^{243}Am ”. In: *Nuclear Science and Engineering* 125.1 (1997), pp. 51–60.
- [5] A. D’ANGELO et al. “Delayed neutron data for the major actinides”. NEA/WPEC-6 Report. OECD Nuclear Energy Agency (NEA), 2002.
- [6] H. A. ABDERRAHIM and M. GIOT. “The Accelerator Driven Systems, a 21st Century Option for Closing Nuclear Fuel Cycles and Transmuting Minor Actinides”. In: *Sustainability* 13.22 (2021).
- [7] N. E. AGENCY. “Accelerator-driven Systems (ADS) and Fast Reactors (FR) in Advanced Nuclear Fuel Cycles”. In: *OECD Publishing, Paris* (2002).
- [8] C. ARTIOLI et al. “Minor Actinides Transmutation in the EFIT Reactor: A Design Challenge”. In: *Proceedings of NUCLEAR 2008 Annual International Conference on Sustainable Development through Nuclear Research and Education*. Pitesti, Romania, 2008.
- [9] CERN COURIER. “Mumbai engages ADS for nuclear energy”. In: *CERN Courier* (Mar. 2012).
- [10] H. A. ABDERRAHIM et al. “MYRRHA - A multi-purpose fast spectrum research reactor”. In: *Energy Conversion and Management* 63 (2012), pp. 4–10.
- [11] M. SAROTTO et al. “The MYRRHA-FASTEF cores design for critical and sub-critical operational modes (EU FP7 Central Design Team project)”. In: *Nuclear Engineering and Design* 265 (2013), pp. 184–200.
- [12] H. TAKEI et al. “Estimation of acceptable beam-trip frequencies of accelerators for accelerator-driven systems and comparison with existing performance data”. In: *Journal of Nuclear Science and Technology* 49.4 (2012), pp. 384–397.

- [13] P. NAPOLITANI. “Crossover from a fission-evaporation scenario towards multifragmentation in spallation reactions”. In: *International Meeting "Selected topics on nuclear methods for non-nuclear applications"*. Varna, Bulgaria, Sept. 2006.
- [14] A. GANDINI and M. SALVATORES. “The Physics of Subcritical Multiplying Systems”. In: *Journal of Nuclear Science and Technology* 39.6 (2002), pp. 673–686.
- [15] G. I. BELL and S. GLASSTONE. “Nuclear Reactor Theory”. 1970, Oct. 1970, p. 278.
- [16] G. RIMPAULT et al. “The ERANOS code and data system for fast reactor neutronic analyses”. In: *Proceedings of the International Conference on the New Frontiers of Nuclear Technology (PHYSOR 2002)*. Seoul, South Korea, Oct. 2002.
- [17] C. J. WERNER et al. “MCNP User’s Manual, Code Version 6.2”. Tech. rep. LA-UR-17-29981. Los Alamos, NM, USA: Los Alamos National Laboratory, 2017.
- [18] J. R. LAMARSH and A. J. BARATTA. “Introduction to Nuclear Engineering”. 3rd. Upper Saddle River, NJ: Prentice Hall, 2001, pp. 248–252.
- [19] C. R. BATES et al. “The MCNPTools Package: Installation and Use”. Report LA-UR-22-28935. Los Alamos National Laboratory, Aug. 2022.
- [20] M. SAROTTO, F. LODI, and R. PERGREFFI. Technical Report ENER-P000–005. ENEA Centro Ricerche Bologna, 2025.
- [21] LOS ALAMOS NATIONAL LABORATORY. *MCNPX Visual Plotter*. Version X_24E.
- [22] H. W. BERTINI. “Low-Energy Intranuclear Cascade Calculation”. In: *Phys. Rev.* 131 (1963), p. 1801.
- [23] M. HERMAN et al. “ENDF-6 Formats Manual”. Report. Brookhaven National Laboratory (BNL), Upton, NY (United States), Sept. 2023.
- [24] THE HDF GROUP. *myHDF5*. URL: <https://myhdf5.hdfgroup.org/> (visited on 08/18/2026).
- [25] L. R. DALCÍN et al. “Parallel distributed computing using Python”. In: *Advances in Engineering Software* 42.12 (2011), pp. 1030–1048.
- [26] A. E. WALTAR and A. B. REYNOLDS. “Fast Breeder Reactors”. New York: Pergamon Press, 1980.
- [27] F. IANNONE et al. “CRESCO ENEA HPC clusters: a working example of a multifabric GPFS Spectrum Scale layout”. In: *2019 International Conference on High Performance Computing Simulation (HPCS)*. 2019, pp. 1051–1052.