

POLITECNICO DI TORINO

Master's Degree in Sustainable Nuclear Energy



**Politecnico
di Torino**

Master's Degree Thesis

Modelling and implementation of a configuration agnostic 3D fusion neutron source for Monte Carlo neutron transport simulations

Supervisors

Prof. Sandra DULLA

Prof. Antti SNICKER

Prof. Nicolò ABRATE

Candidate

Giacomo VENEZIA

Academic Year 2024-2025

Table of Contents

1	Introduction	1
2	Theory	3
2.1	Fusion background	3
2.2	Fusion reactions modeling	5
2.3	Volumetric Neutron Source	9
2.4	Monte Carlo methods	11
3	ASCOT5 suite of codes	13
3.1	ASCOT5	13
3.2	BBNBI5	15
3.3	AFSI5	15
4	Development of a new AFSI5 neutron source routine	19
4.1	ASCOT5 simulation framework	20
4.2	Motivation for new AFSI5 source routine	21
4.3	Design of the new AFSI5 source routine	22
4.3.1	C algorithm	23
4.3.2	Python interface and pre-processing	26
5	Verification and Benchmarking	29
5.1	Simulation setup	29
5.2	Validation	33
5.2.1	VNS Thermonuclear fusion	33
5.2.2	VNS Beam-thermal fusion	36
5.2.3	ITER configuration	40
6	Serpent2 Results	45
6.1	Simulation setup	45
6.2	Neutronics results	48
7	Conclusions	57
	Bibliography	59

1. Introduction

The foundation of this work lies on neutronics analysis in fusion research, which has been progressively increasing as fusion devices move towards reactor relevant conditions. The behavior of the high energy neutrons produced in the plasma plays a central role in the design of the reactor, determining the material degradation, shielding requirements and tritium breeding performances. For this reason, neutron transport simulations are a standard tool in reactor design, and their accuracy depends critically on how the neutron source is defined.

This thesis is based on the work that has been carried out at VTT Technical Research Centre of Finland within the ASCOT group, where simulation tools are developed to model fast particle dynamics in realistic magnetic configurations. The work originates from a practical limitation that emerges in current neutron transport workflows. Present day analyses typically follow two paths, the most common approach relies on simple analytical neutron source models, which prescribe idealised spatial profiles and isotropic emission. These models are straightforward to use in Monte Carlo transport codes such as Serpent and MCNP, but they do not capture the actual reactivity distribution of realistic plasmas.

The less used but more descriptive alternative is to use the ASCOT5 chain of tools, which computes fusion reactivity and provides a more realistic neutron source through AFSI5, the ASCOT5 Fusion Source Integrator. AFSI5 generates the output in a five dimensional distribution composed of a 3D geometrical space and 2D momentum space aligned with the magnetic field geometry. However, despite carrying valuable physical information, this data cannot be directly used to define a complete neutron source in Serpent, because it does not include a description of the neutron emission direction. In addition, the AFSI5 output requires extensive preprocessing and format conversion, and remains tied to specific grids and coordinate systems, limiting its portability and complicating its use in different reactor configurations.

The purpose of this thesis is to address this gap by developing a configuration agnostic routine capable of generating a full six dimensional fusion neutron source that is directly usable in modern Monte Carlo transport codes. The method builds upon the ASCOT5 chain of tools but extends it to provide a complete spatial, energetic and angular description of the neutron emission. The goal is to offer a realistic and general source definition that avoids the limitations of analytical models while remaining compatible with arbitrary plasma configurations and reactor geometries.

This work is structured as follows, the theoretical background and numerical tools are introduced first in Chapters 2 and 3, with particular attention to the ASCOT5

chain of tools. Chapter 4 details the formulation and implementation of the new AFSI5 source routine, and in Chapter 5 the results are then presented and validated on different plasma configurations. Finally, the study concludes by employing the new realistic source for a Volumetric Neutron Source tokamak in a Serpent2 neutron transport calculation, where its results are compared against that of a conventional analytical neutrons description in a beam-heated scenario. This final step illustrates the practical consequences of using a realistic six dimensional source and highlights the improvements it brings to reactor relevant neutronics analyses.

2. Theory

Modeling the behavior of fusion plasmas and the associated neutron transport requires a solid understanding of plasma physics and statistical computational techniques. This chapter presents the key theoretical concepts necessary to carry out this work. It begins with a general introduction to fusion energy, followed by an overview of fusion reactions. Regarding this topic, particular emphasis is placed on how the distribution of reactants affects the distribution of fusion products, which is central to the modeling of fusion neutron sources. Beam-thermal fusion is discussed in detail, as it is the dominant mechanism in the newly proposed testing facility, the Volumetric Neutron Source, described in Section 2.3. The chapter concludes with a description of Monte Carlo methods for particle transport, which are extensively used in this work to model ion and neutron transport, as well as fusion reaction rates.

2.1 Fusion background

Fusion is one of the greatest scientific challenges of the last 50 years. Researchers around the world have been striving to replicate on Earth the processes that power the Sun and other stars, with the common goal of providing a clean and virtually unlimited energy.

This source of energy comes from the nuclear fusion reaction, the process where two lighter atomic nuclei combine together to form a heavier nucleus. Generally, if the product is lighter than Nickel-62, the fusion reaction is endothermic, with a release of energy in the form of kinetic energy of the products.

However, combining two light nuclei is particularly challenging, as the repulsive Coulomb force acts between the two positive charged nuclei. Therefore, it is needed to bring the reactant close enough so that the strong nuclear force overcomes the electrostatic repulsion and allows fusion to occur. In a fusion reactor this condition is met by rising the fuel temperature to hundreds of millions K. At such high temperatures, the reactants will be fully ionized and turn into plasma, allowing the use of high magnetic field to shape and confine the plasma in the reactor chamber.

In magnetic confinement fusion devices, strong magnetic fields are used to keep the fusion plasma within the reactor, the most developed configuration is the Tokamak, where the plasma is confined in a toroidal chamber by a combination of toroidal and poloidal magnetic fields that together create a stable, twisted field structure. The largest experimental reactor based on magnetic confinement is the International Thermonuclear Experimental Reactor (ITER), an international project designed to demonstrate the feasibility of controlled thermonuclear fusion.

The performance of a fusion reactor is commonly expressed through the fusion power gain, defined as $Q = \frac{P_{out}}{P_{in}}$, where P_{out} is the power produced by fusion reactions and P_{in} is the power required to heat and sustain the plasma. The condition $Q = 1$ is called breakeven and marks the point where the reactor produces as much power as it consumes. Future reactors like ITER aim to reach $Q \approx 5-10$, demonstrating a net production of fusion energy. Achieving such conditions requires maintaining a proper energy balance in the plasma, which depends on three key parameters: the electron density n_e , the plasma temperature T_e , and the energy confinement time τ_E . Their combination defines the so-called Lawson criterion [1], which sets the minimum requirements for sustained thermonuclear fusion, expressed as:

$$n_e \tau_E \geq \frac{12 k_b T_e}{\langle \sigma v \rangle_{DT}(T_e) E_\alpha} \quad (2.1)$$

where k_b is the Boltzmann constant, σ is the fusion cross section, v is the relative velocity of the reactants and E_α is the fusion heating power.

However, in a fusion reactor the maximum plasma pressure is limited by magnetic confinement stability, so the criterion is often expressed in terms of the fusion triple product $n_e T_e \tau_E$, where $n_e T_e = p$. This product serves as the conventional figure of merit for thermonuclear fusion reactions in magnetic confinement devices:

$$n_e T_e \tau_E \geq \frac{12 k_b T_e^2}{\langle \sigma v \rangle_{DT}(T_e) E_\alpha} \quad (2.2)$$

From this equation, the fusion reactivity is maximized when the triple product is minimum. For a Deuterium-Tritium (DT) fusion reaction this condition is met at the temperature $T_e = 13.54 \text{ keV}$ as shown in Figure 2.1.




Figure 2.1: The triple product $n_e T_e \tau_E$, versus the electron temperature T_e for a DT reaction.

2.2 Fusion reactions modeling

The fusion reactions between hydrogen isotopes are the most studied and promising, as they have a high probability for the reaction to occur. Figure 2.2 shows the fusion cross section as a function of the kinetic energy of the reacting particles. In particular, the DT fusion Eq.2.3, is the most conveniently achievable having the highest collision cross-section occurring at the lowest temperature.



In this fusion reaction the released energy is approximately 17.6 MeV and the energy share is inversely proportional to the masses, so that 80% of it is distributed to the neutron (14.1 MeV) and 20% to the alpha particle (3.5 MeV). Because such a large fraction of the fusion energy is transported by neutrons, which are not confined by the magnetic field, it becomes essential to model their behavior accurately. This starts with a realistic definition of the neutron source, which is fundamental for evaluating their role in energy deposition, material activation, and overall reactor performance.




Figure 2.2: Fusion cross section spectra for different fusion reactions. Image from [2]

In the first ITER-like tokamaks, the fuel is heated up to 10-20 keV, in this condition the particles in the plasma collide and exchange energy frequently enough that their velocities become randomized. Due to this behavior, the reactants energy follows a Maxwellian distribution [3], where most particles have energies close to the average thermal energy, while fewer particles have very low or very high energies. In this state, the population is defined as thermalized, and the energy distribution is fully determined by the plasma temperature, with the mean value scaling proportionally to $k_b T$.

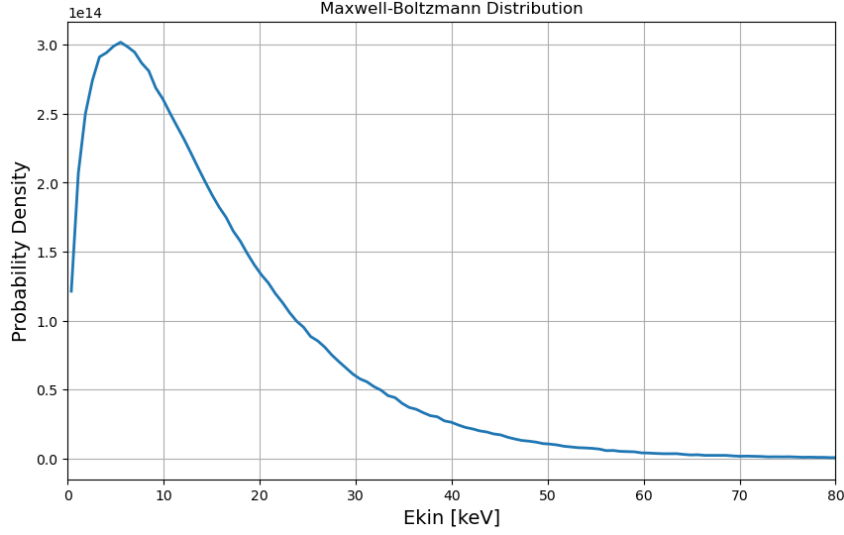


Figure 2.3: Energy distribution of the reactants in a 10 keV plasma

Because the reactant velocity arises from the Maxwellian energy distribution, each Cartesian velocity component follows a Gaussian centered at zero, whose variance depends on the temperature, resulting in a zero mean population velocity. When fusion occurs between two thermalized species, it is referred to as thermonuclear fusion.

To understand how the velocity, and thus energy, distributions of the reactants determine the kinematic properties of the fusion products, it is necessary to derive the mathematical model for a binary nuclear reaction.

The reactants velocities are defined in the center-of-mass reference system, as:

$$\vec{v}_{1,CM} = \vec{v}_1 - \vec{v}_{CM} \quad (2.4)$$

$$\vec{v}_{2,CM} = \vec{v}_2 - \vec{v}_{CM} \quad (2.5)$$

With velocity of the system's center-of-mass is defined as:

$$\vec{v}_{CM} = \frac{(m_1 \vec{v}_1 + m_2 \vec{v}_2)}{m_1 + m_2} \quad (2.6)$$

Defining Q as the reaction energy, subscripts 1,2 for the reactant species and subscripts A, N for products species, the governing equations are:

$$E_{tot} = \frac{1}{2}m_1|\vec{v}_{1,CM}|^2 + \frac{1}{2}m_2|\vec{v}_{2,CM}|^2 + Q = \frac{1}{2}m_N|\vec{v}_{N,CM}|^2 + \frac{1}{2}m_A|\vec{v}_{A,CM}|^2 \quad (2.7)$$

$$m_1\vec{v}_{1,CM} + m_2\vec{v}_{2,CM} = m_N\vec{v}_{N,CM} + m_A\vec{v}_{A,CM} \quad (2.8)$$

Given that the fusion reaction in the CM frame isotropic, the unit velocity vector of product N \hat{v}_N is sampled uniformly on the unit sphere and the velocity is obtained as:

$$\vec{v}_{N,CM} = \sqrt{\frac{2E_{tot}}{m_N \left(1 + \frac{m_N}{m_A}\right)}} \hat{v}_N \quad (2.9)$$

Finally, the products velocities are transformed back into the laboratory frame:

$$\vec{v}_A = -\frac{m_N}{m_A}\vec{v}_{N,CM} + \vec{v}_{CM} \quad (2.10)$$

$$\vec{v}_N = \vec{v}_{N,CM} + \vec{v}_{CM} \quad (2.11)$$




Figure 2.4: Kinematic scheme of fusion reaction. (a) Reactants with velocity v_1 and v_2 in laboratory frame, have a CM velocity v_{CM} . (b) Reactant velocities u_1 and u_2 in CM frame. (c) Fusion product produced isotropically with velocity u_3 in CM frame. (d) Fusion product transformed back in CM frame. Image from [4]

From the kinematic derivation summarized in Fig.2.4, we see that the emission of the fusion products depends on the velocity of the center-of-mass of the reacting pair. In particular, when both reactants follow a Maxwellian distribution, each velocity component is Gaussian around zero, so on average the CM velocity is zero. This means that the emission has no preferred direction, and in the laboratory frame the products are still emitted isotropically, with energy spectra set by the Gaussian velocity distribution of the reactants [3], as shown in Fig. 2.5.




Figure 2.5: Neutron Energy Distribution from Maxwellian Reactants

This isotropic condition however, can be perturbed by external plasma heating such as Neutral Beam Injection, where a beam of fast neutrals is injected in the plasma. Indeed, the high energy of the beam and its velocity distribution, implies that the injected reactant population is non-Maxwellian distributed and therefore the emitted neutrons will have a spectrum different from the thermonuclear case

described earlier.

Even though in common fusion reactor the NBI heating is used only during a small fraction of time with respect to the operation time, and therefore this perturbation can be negligible, it can have a significant impact on machines that rely primarily on NBI such as the Volumetric Neutron Source Tokamak Sec.(2.3).

Consequently, when modeling the fusion products with eq. (2.6) even if they are emitted isotropically in CM frame, when converted to laboratory frame via Eq (2.8), the direction vectors are affected by the CM velocity components, producing the anisotropy condition mentioned.

This transformation also broadens the neutron energy distribution around the 14.1 MeV peak. In Figure 2.6, the comparison between the neutron spectra produced by D-T thermonuclear and beam-thermal reactions in the ITER baseline scenario illustrates such behavior. In fact, while thermonuclear reactions occurring between thermal ions from a Maxwellian distribution results in a narrow and symmetric neutron spectrum centered around 14 MeV, beam-thermal reactions involving fast, non-thermal ions, generates a broader and more asymmetric spectrum. Moreover, the injection angle of the beam, in addition to its energy, significantly influences both the energy distribution and the emission direction of the resulting fusion-born neutrons, further contributing to the anisotropy observed in the spectrum.

These effects underline the importance of accurately modeling the neutron source in devices where beam-thermal fusion plays a significant role. In such scenarios, neglecting the influence of beam energy and injection geometry can lead to incorrect predictions of neutron spectra and flux distributions, which are critical for reactor diagnostics, shielding design, and material testing.

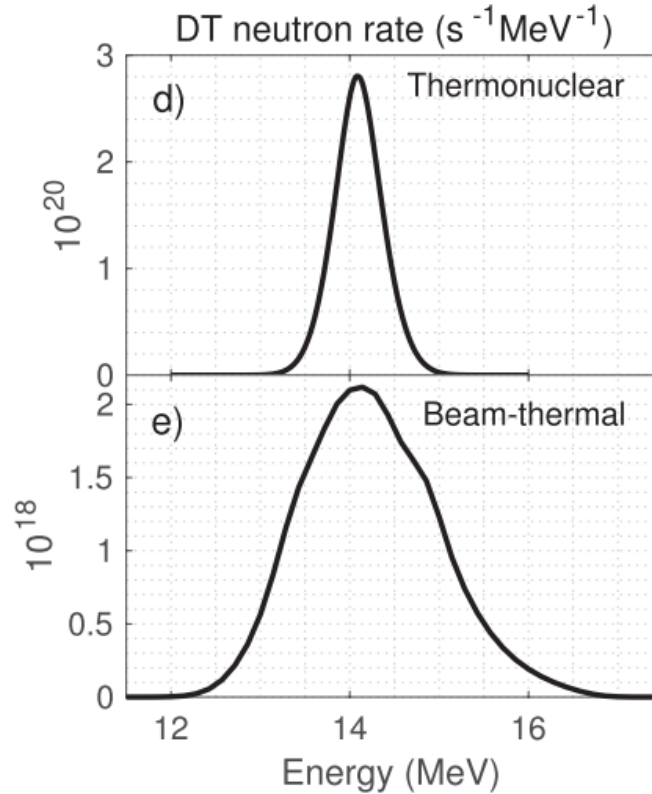


Figure 2.6: ITER baseline scenario neutron spectra for DT thermonuclear (d) and beam-thermal (e) reaction. Image from [2]

2.3 Volumetric Neutron Source

Materials limit is one of the main bottlenecks of fusion. Past R&D on this challenge [5] has been carried out mostly with the JET machine [6] and with non-fusion facilities, aiming to test components exposed to plasma and neutron environments to better understand material behavior under fusion-relevant conditions.

Fission reactors have been widely used as they can provide a moderate-high volume of neutrons, suitable for experiments that investigate irradiation damages. However, since fission neutrons are less energetic than fusion neutrons, their effects on materials is limited on the surface, as the flux is not powerful enough to reach the bulk of the material.

Facilities like the Fusion Neutronics Source (FNS) at the Japan Atomic Energy Agency and the Frascati Neutron Generator (FNG) at ENEA in Italy, have been used extensively for experiments and code validation, as they produce 14 MeV neutrons from DT reactions. Despite producing fusion born neutrons, these facilities do not fully replicate reactor conditions, as the neutron fluxes they generate are several orders of magnitude lower than those expected in actual fusion reactors.

Therefore comes the need to build a Volumetric Neutron Source (VNS) as a material testing facility, to correctly replicate the fusion reactor environment producing a high flux of 14 MeV neutrons over a large surface and volume.

The VNS is a tokamak with smaller dimensions than an actual fusion reactor, where the fusion reaction is sustained with a continuous operation of Neutral Beam Injectors and the only goal is to produce neutrons, not energy. The testing conditions require a neutron wall load (NWL) of about 0.5 MW/m^2 , a damage level in the first wall of 30-50 dpa, and a plasma pulse long enough to reach thermal equilibrium in the Test Blanket Module [7]. The fusion power is kept below 50 MW to minimize the tritium consumption and rely only on external supply.




Figure 2.7: View of VNS tokamak. Image from [7]

The first concept of VNS has been presented by EUROfusion consortium [5],

[7], this machine will have four tangential beamlines injecting D neutrals at 120 keV, generating a nominal fusion power of 29 MW. The major and minor radii are respectively $R = 2.53$ m and $a = 0.55$ m to maximize the neutron wall load up to the targeted 0.5 MW/m^2 . In Fig.2.7 a render of the VNS tokamak is reported, showing the four injectors and the several ports where the Test Blanket Modules will be inserted. These ports will be used to investigate the coupled physical phenomena occurring in fusion relevant blanket environments, such as tritium production, nuclear and gamma-ray volumetric heating, and radiation damage with spatial gradients across the blanket structure. The data obtained from these experiments will be fundamental to validate breeding blanket designs, assess component lifetime, and develop reliable maintenance and integration strategies for ITER and future reactors.

As shown in Fig. 2.8, the core of a NBI is the ion source, where positive or negative ions are produced and energized by going through a multi-stage acceleration system with high voltage grids, until the desired energy is reached. In the neutralization cell, the majority of the accelerated ions is neutralized by colliding with the background gas present in the cell and triggering the charge exchange reaction, the remaining fraction of ions present in the beam is removed using bending magnets before the injection port. Finally the neutral beam is injected into the plasma through the injection port, the direction of injection, as shown in Figure 2.9, can be co-tangential with the beam parallel to the toroidal plasma current, counter-tangential with the beam parallel but in the opposite direction of the plasma current and perpendicular with the beam injected nearly perpendicular to the magnetic field. The neutral beam can be of two types, positive ion and negative ion. Although positive hydrogen ions are technically easier to produce with common ionization techniques, the neutralization efficiency decreases rapidly as the beam energy increases. For a 100 keV beam with positive hydrogen ions the efficiency is about 20% and goes to almost zero for beams with 500 keV energy. Therefore, for the denser and larger plasmas of modern machines, high energy beams are needed (up to 1 MeV) and N-NBI have to be used, which is able to maintain a neutralization efficiency of 60% at energies higher than 100 keV.




Figure 2.8: Schematic of a neutral beam injector. Figure from [8].




Figure 2.9: Neutral beam injection directions. Figure from [8]

2.4 Monte Carlo methods

Monte Carlo methods are computational algorithms that rely on random sampling to estimate the results of a wide range of problems.

They are particularly useful in complex systems characterized by high complexity and uncertainty, that often cannot be modeled with a deterministic approach. In nuclear problems, Monte Carlo methods are used to simulate the transport of particles in fusion and fission reactors.

Unlike deterministic methods, which aim to solve the transport equation for the particle flux, in the Monte Carlo approach the problem is viewed from the perspective of individual particles. Indeed, by simulating a large number of particle histories that interact with the system, a statistical average of physical quantities can be obtained without solving any transport equation.

By tracking a large number of test particle trajectories through realistic magnetic and plasma fields, the method enables the estimation of key quantities such as spatial distributions, power deposition profiles, and wall load patterns. This particle-based approach is particularly suited for capturing the stochastic nature of particle motion and collisions in magnetically confined plasmas. The accuracy of these estimates improves with the number of simulated particles, as described by the central limit theorem, which ensures convergence toward the true expectation value and quantifies the associated statistical uncertainty.

Despite the advantage of not solving the transport equation explicitly, these methods come with the trade-off between accuracy and computational cost, as achieving low statistical uncertainty requires the simulation of a large number of particles. However, the intrinsic parallel nature of the algorithms, makes them well suited for high-performance computing (HPC) clusters, that can significantly reduce computation time.

In the scope of this work, Monte Carlo methods are first employed using the ASCOT5 code to simulate the slowing down of fast ions generated by Neutral Beam Injection in a fusion plasma. These ions, modeled as test particles, interact with the plasma background through collisions, gradually losing energy and potentially undergoing fusion reactions. The outcome of this simulation provides a volumetric neutron source profile, which reflects the spatial and spectral distribution of the fusion-born neutrons within the plasma. This source is then used as input for further Monte Carlo simulations in Serpent2, which models the subsequent transport of these neutrons

through the surrounding geometry to evaluate material responses and neutron flux distributions in reactor-relevant conditions.

3. ASCOT5 suite of codes

ASCOT5 is a Monte Carlo test-particle orbit-following code designed for solving the distribution functions, transport, and losses of plasma species in 3D geometry [9]. A typical application is the study of beam–thermal fusion reactions, where fast ions from neutral beam injection interact with the thermal plasma to produce fusion products. To model this process, ASCOT5 provides a complete chain-of-tools: the Beamlet-Based Neutral Beam Ionization model (BBNBI) generates the fast-ion source from neutral beam injection; the core ASCOT5 code follows the slowing-down and transport of these ions; and the Ascot Fusion Source Integrator (AFSI) models the resulting fusion reactions, storing the product distributions in five-dimensional phase space.

In this work, a dedicated neutron source routine is developed in AFSI5 to couple ASCOT5 with established Monte Carlo neutronics codes, enabling consistent modeling of fusion-born neutron transport in complex reactor geometries.

3.1 ASCOT5

The whole ASCOT5 architecture, as well as its ancillary tools AFSI5 and BBNBI5, has the core code written in C and the main user interface in Python. This hybrid approach allows to have a fast computation thanks to the execution in C and an easier pre-processing of the data through the Python-C interface. In practice, ASCOT5 uses the HDF5 format for input and output data handling, while pre-processing and post-processing operations are carried out through its dedicated Python API *a5py*, which provides convenient access to the simulation parameters and results.

In ASCOT5 it is implemented the tokamak coordinate convention COCOS3 [10], where the cylindrical coordinates are right-handed (R, ϕ, Z) , with the toroidal angle ϕ increasing counterclockwise, and in the azimuthal plane (x, y) , the x-axis is chosen so that $\phi = 0$ corresponds to $y = 0$. Toroidal field and plasma current are positive when they point in the same direction as ϕ .

The code solves the Fokker-Planck equation indirectly by tracing particle orbits in realistic 3D magnetic fields, solving guiding-center or the full gyro-orbit equations of motion.

$$\frac{\partial f(z, t)_s}{\partial t} = -\frac{\partial}{\partial z}[K(z, t)f(z, t)_s] + \frac{\partial^2}{\partial z \partial z}[D(z, t)f(z, t)_s] \quad (3.1)$$

Where $f(z, t)$ is the distribution function to be solved in the phase space coordinates \mathbf{z} , while \mathbf{K} and \mathbf{D} are the Fokker-Planck advection and diffusion coefficients. During the orbit-following, ASCOT5 models the interaction between the test particles and

the background plasma through a Monte Carlo treatment of Coulomb collisions and atomic reactions.

Monte Carlo simulations require a large number of particles to achieve good statistical accuracy, ASCOT5 is designed to efficiently handle large-scale simulations. It employs MPI-based parallelization, allowing multiple particle trajectories to be computed independently across distributed computing environments. This approach ensures near-linear scalability with the number of cores, making the code suitable for high-performance computing clusters.

Fast-ions distributions from which markers are sampled are the starting point slowing-down simulations. This input can be either provided manually or defined internally with ASCOT's tools: BBNBI5 if fast beam ions are needed, and AFSI5, when the input is an alpha source distribution.

Running a fast-ions slowing down simulation, requires only magnetic field, plasma data and fast-ions distribution. As mentioned, guiding-center or gyro-orbit simulation mode can be chosen to trace the markers: if gyro-orbit is selected, the marker is a physical particle and its whole gyro-motion is solved, conversely, if the magnetic moment is an adiabatic invariant and collecting guiding center distribution is sufficient for the results, gyro-orbit effect can be ignored and the guiding-center scheme can be used. The guiding-center transformation [ref], moves from the classical particle phase space $z = (x, p)$ to the guiding-center phase space $Z = (X, p_{\parallel}, \mu, \zeta)$, with guiding center position Z , momentum component parallel to the magnetic field p_{\parallel} , magnetic moment μ , and gyroangle ζ .

ASCOT5's main output consists of marker orbits, N-dimensional histograms representing the test particle distribution function in the phase-space, and marker simulation endstates that can be used for estimating power losses and wall loads. The output distribution function contains all the information needed to describe the physical system. It can be collected in five different distributions:

- 5D $(R, \phi, z, p_{\parallel}, p_{\perp})$
- rho5D $(\rho, \theta, \phi, p_{\parallel}, p_{\perp})$
- 6D $(R, \phi, z, p_R, p_{\phi}, p_z)$
- rho6D $(\rho, \theta, \phi, p_R, p_{\phi}, p_z)$
- COM (E_{kin}, P_{ctor}, μ)

5D and rho5D are the default distributions for guiding-center simulations, 6D and rho6D have to be used in all the gyro-orbit simulations. Lastly, COM (constant-of-motion) distribution collects the data in terms of kinetic energy, canonical toroidal angular momentum and magnetic moment.

When using ASCOT5 there are two main limitations to be aware of. Firstly, it uses the test-particle approximation, meaning that the simulated particles do not influence the background plasma or with other test-particles. Second, it assumes a constant source of particles, which can complicate the interpretation of quantities such as wall loads and particle distributions in transient scenarios. In this work, however, a steady-state NBI is simulated specifically to obtain the volumetric neutron source resulting from beam-plasma fusion. Since the particle source is constant over time

and no transient dynamics are involved, the assumptions made by ASCOT5 are fully compatible with the goals of this simulation.

3.2 BBNBI5

BBNBI5 [11] is a beamlet-based Neutral Beam Injection (NBI) model designed to determine the deposition of ions from neutral injectors in 3D plasma geometry. It follows each neutral particle emitted from the injector until it becomes ionized or hits the wall, producing an ensemble of fast test ions that serve as the initial population for orbit-following simulations. The code models the injector as an array of individual beamlets, each characterized by its position, direction, energy, and intensity, allowing for a realistic description of the beam divergence and shaping.

BBNBI5 is an independent tool but is closely integrated with the ASCOT5 framework, using the same magnetic equilibrium and plasma profiles as inputs. It computes the spatial ionization profiles and resulting fast-ion distributions, which can then be passed directly to ASCOT5 for orbit-following and slowing-down simulations. This modular approach ensures consistency between beam deposition, fast-ion transport, and fusion product modeling.

The physics model in BBNBI5 accounts for key atomic processes such as charge-exchange, electron-impact ionization, and ionization by plasma ions, which determine the penetration depth and power deposition of the injected beam. Thanks to its beamlet-based formulation, the model can reproduce realistic injection geometries and non-axisymmetric effects, which are important for devices such as stellarators or ITER-like tokamaks. Overall, BBNBI5 provides a physically detailed and computationally efficient method for generating fast-ion sources in fusion plasmas, forming the first step of the beam–thermal fusion modeling chain used in this work.

3.3 AFSI5

AFSI5 is a tool designed to efficiently model fusion reaction based on the interaction of two arbitrary reactant distribution [2]. The code computes the resulting fusion reaction rates and characterizes the products through energy spectra, reaction rate profiles, and five-dimensional phase-space distributions, providing a detailed description of the emission of fusion products.

Through its Python API, three operation modes can be selected to obtain fusion rates and product distributions for: (i) thermonuclear fusion between two Maxwellian species, (ii) beam–thermal fusion between a thermal species and an arbitrary distribution, and (iii) beam–beam fusion between two arbitrary distributions.

AFSI5 requires only magnetic field, plasma data and reactant ions distribution as inputs. Given these parameters, the fusion product distribution is obtained by integrating with a Monte Carlo approach the reactivity equation:

$$R = \iiint_{\vec{v}_B} \iiint_{\vec{v}_T} f_T(\vec{v}_T) \sigma(|\vec{v}_B - \vec{v}_T|) f_B(\vec{v}_B) |\vec{v}_B - \vec{v}_T| d\vec{v}_T d\vec{v}_B \quad (3.2)$$

Where $f_T(\vec{v}_T)$ and $f_B(\vec{v}_B)$ are the arbitrary reactant distributions, representing the three possible configurations depending on the selected operation mode.

The algorithm loops over the simulation domain described in cylindrical coordinates (R, ϕ, z) , and for each spatial cell samples n pairs of markers from the reactants distribution. In the case of thermonuclear fusion, the reactant momenta are sampled from a Maxwellian population using Eq.(3.2) from [10.1007/s10955-011-0364-y].

$$E = T[-\ln r_1 - \ln r_2 \cos^2 \frac{\pi}{2} r_3] \quad (3.3)$$

Where T is the plasma temperature and r_1 , r_2 and r_3 are random numbers sampled on the unit interval.

In beam-thermal and beam-beam operation mode, the reactant markers are sampled directly from the 5D ion distribution provided as input either in the form of $(R, \phi, z, p_{\parallel}, p_{\perp})$ or $(R, \phi, z, E_{kin}, \xi)$. The first one uses the momentum components parallel and perpendicular to the magnetic field while the latter is defined with kinetic energy E_{kin} and pitch $\xi = \frac{p_{\parallel}}{||\vec{p}||}$.

Subsequently, the products momenta components are computed from the energy and momentum conservation laws in the center-of-mass coordinates as explained in detail in section 2.2.

The velocities are then converted into momentum components $(p_{\parallel}, p_{\perp})$, allowing the fusion products to be properly stored in the output distributions. The spatial bins assigned to the product particles correspond to the center of the cell from which the reactant pair was sampled.

To represent the physical contribution of each sampled particle to the overall fusion rate in the plasma, a fusion source weight is computed and assigned to each sampled particle.

The weight is computed dividing the fusion rate in the cell by the number of Monte Carlo samples n :

$$w = \rho_1 \rho_2 \langle \sigma v \rangle V / n \quad (3.4)$$

where ρ_1 and ρ_2 are the densities of the two reactant species, $\langle \sigma v \rangle$ is the fusion reactivity, which depends on the sampled relative velocity and the fusion cross section, and V is the cell volume. The fusion cross sections are computed with the Bosch-Hale parametrization [ref]. The density of the thermal species are obtained directly from the plasma data, while for fast-ions populations, density is calculated as the sum of all markers in the cell divided by the cell volume.

These weights represent the physical fusion rate for the sampled reactants, which are accumulated in a 5D product histogram, where the momentum space is 2 dimensional represented in a field-aligned coordinates: $(p_{\parallel}, p_{\perp})$.

Summing the weights over all samples and cells approximates the full multidimensional fusion rate integral.

Due to its compatibility with ASCOT5 and its ability to handle arbitrary reactant distributions, AFSI5 can be effectively employed to generate fusion neutron source terms for coupling with neutronics codes. However, unlike ASCOT5 and BBNBI5, AFSI5 does not produce a full 6D distribution for the fusion products. Instead, the particles velocity is defined as a 2D vector with components parallel and perpendicular to the magnetic field. this representation of the products prevents direct coupling with the classical neutronics code such as Serpent2 or MCNP which both require as input a 6D (x, y, z, v_x, v_y, v_z) neutron distribution, where the velocity space is commonly represented with energy E , and direction vectors u, v, w . Consequently, a

heavy and computationally inefficient post-processing is needed to convert the AFSI5 output into a suitable format for the mentioned codes.

This limitation is addressed in the following section, where a new source routine for AFSI5 is introduced to enable direct coupling with Monte Carlo neutronics codes.

4. Development of a new AFSI5 neutron source routine

This chapter presents a new comprehensive fusion neutron source routine that enables direct coupling between ASCOT5's chain of tools and Serpent2, allowing for consistent neutronics simulations based on plasma-generated neutron distributions. The first section discusses the typical workflow to obtain a neutron distribution in a beam heated fusion plasma, by simulating all the processes involved, from neutrals injection to the beam-thermal fusion.

The process begins with modeling the NBI using BBNBI5, followed by an ASCOT5 fast-ions slowing-down simulation to retrieve the ions distribution in the plasma. Subsequently, the output distribution serves as input for AFSI5, which models the resulting fusion reactions. However, as anticipated in the previous chapter and further discussed in Section 4.2, the current output obtained from AFSI5 does not provide an explicit description on the product direction vectors, making the coupling with neutronics codes particularly challenging. If the information on the direction vectors is neglected, the product distribution can be suitable as source for neutron transport simulation only in the case of thermonuclear fusion, providing to Serpent a four dimensional distribution (x, y, z, E) and sampling the direction uniformly, where this assumption holds.

Therefore, in Section 4.3 a new AFSI5 sampling method is developed to model thermonuclear, beam-thermal, and beam-beam fusion in a way that produces output directly compatible with the neutron transport codes discussed above.

The focus in this work is placed mostly on the beam-thermal reaction, since having a realistic source definition for devices where the NBI is operating in steady state is crucial to accurately assess the neutrons effects on the materials. The reference configuration for this study is the Volumetric Neutron Source Tokamak proposed by EUROfusion [7]. This device will employ four Neutral Beam Injectors to generate a reactor-relevant neutron flux, enabling improved understanding of neutron damage in structural materials and supporting the design of the breeding blanket. In such a scenario, the non-Maxwellian fast-ion population is expected to introduce significant anisotropy in the neutron emission. A realistic source model that correctly captures the emission direction is therefore crucial, yet the current AFSI5 routine provides only a limited description in this respect.

4.1 ASCOT5 simulation framework

To generate the fusion neutrons produced in a VNS tokamak, ASCOT5 provides a full framework that efficiently simulates all the processes involved from the injection of the Deuterium neutrals to the fusion reaction between the beam ions and the thermal plasma.

The first step involves the use of BBNBI5 to model the ionization process of the Deuterium beam, which follows the injected neutrals until ionization, and generates a distribution representing the fast-ions source. The input definition for this first simulation affects the accuracy of the entire framework, as the magnetic field and plasma data are shared within the ASCOT5's environment and are used as inputs for the subsequent ASCOT5 and AFSI5 runs as well. In particular, a high fidelity magnetic field, such as a three-dimensional description interpolated with splines, is recommended to perform accurate orbit-following studies. It is also the most memory and CPU demanding input, since a significant portion of the computational cost is associated with its interpolation.

In addition to plasma and magnetic field data, BBNBI5 requires neutral, wall and atomic data, and separately all the NBI inputs: containing the injector geometry, the injected species, the number of beamlets, energy and power.

The fast-ions markers can be stored in a multidimensional distribution or directly utilized as input for the subsequent ASCOT5's slowing down simulations, which allows to run the two tools in series without the need of any intermediate postprocessing step.

Therefore, once obtained the fast-ions source, the next step is to model the slowing down of the Deuterium beam. As anticipated, the majority of inputs required by ASCOT5 are shared with BBNBI5 and AFSI5, without the need to provide them again. Conversely, it is necessary to specify in the simulation options whether a full-orbit or guiding-center scheme is used. Since the objective of this process is to obtain a neutron source, detailed information about gyro-orbit is not required. Therefore, the guiding-center approximation can be employed, significantly reducing the computational cost of the slowing down simulation.

The results of the slowing-down simulation are stored in a HDF5 file that can be accessed via the Ascot object. Among the stored data there is the distribution of fast ions after they have slowed down due to Coulomb collisions with the background plasma. The distribution is five dimensional due to the adoption of the guiding-center scheme and the input requirements for the subsequent AFSI5 simulation, and can be expressed in two forms: either $(R, \phi, z, p_{\parallel}, p_{\perp})$ or $(R, \phi, z, E_{kin}, \xi)$. Here, the geometry space is defined in cylindrical coordinates while the velocity space is described using either momentum components parallel and perpendicular to the magnetic field, or using kinetic energy and pitch.

Finally, the beam-thermal fusion process is modeled through the AFSI's source routine, which requires only a minimal and well-defined set of inputs: the fusion reaction to be modeled, the reactant species, and the number of Monte Carlo samples to be generated in each spatial cell in order to obtain an accurate product distribution. In the VNS configuration considered here, the reaction of interest is the DT fusion reaction. In this framework, the thermal reactant is sampled using the plasma profiles, while the beam-ion population is taken from the ASCOT5 output distribution. This represents the only additional input needed in the

final step, enabling the generation of a fully physics-based fusion product distribution.

The output consists in two five dimensional distributions: `prod1` and `prod2`, representing respectively the alpha particles and the neutrons.

Since the aim is to use this output as a neutron source for neutronics codes such as Serpent2, which expect a 6D input, this format requires a non-trivial post-processing step. This involves reconstructing the full velocity space by sampling the missing gyrophase and pitch-angle information, which is both computationally expensive and potentially introduces numerical inaccuracies.

This challenge is addressed in more detail in Section 4.2, and finally, in Section 4.3, the core of this work is presented, with a newly developed AFSI5 routine to generate the neutron source in a format directly compatible with Serpent2, reducing both processing time and uncertainty.

In Figure 4.1 a map of the framework above explained is presented.

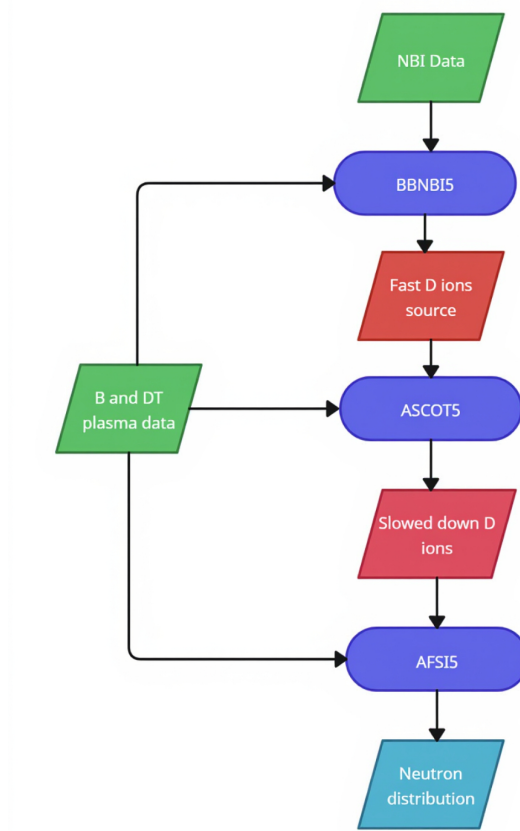


Figure 4.1: Full process to obtain beam-thermal fusion neutrons starting from a NBI

4.2 Motivation for new AFSI5 source routine

The main objective of this work is to couple the workflow presented in the previous section with the neutron transport code Serpent2, in order to have a realistic neutron source for all the scenarios where the classic thermonuclear fusion is not the main player and a more detailed source is required.

As shown in Figure 4.1, through ASCOT5's frameworks it is possible to model the

physics from the injection of neutrals into the plasma to the actual fusion reaction. Nonetheless, the current AFSI5 routine that models the fusion reaction, produces the neutron distribution in a format non compatible with the subsequent neutronics simulation.

Indeed, current AFSI5 routine, relies on a grid-based Monte Carlo integration algorithm, as explained in detail in Sec.3.2, that produces as output the five-dimensional distribution of fusion products. While this approach is well suited for analysis within the ASCOT5 framework, due to the field-aligned momentum space, is not directly compatible with neutronics code such as Serpent2. These codes require the neutron source to be described in the 6D phase space, either as a list of individual fusion product markers with positions and velocities in Cartesian coordinates, or as a source distribution such as a histogram or probability density function, from which particles can be sampled during the simulation.

Therefore, the result distribution obtained from AFSI5 cannot be directly used as input and must be post-processed to convert the 2D momentum space $(p_{\parallel}, p_{\perp})$ into a full 3D Cartesian velocity space. This involves first sampling markers from the 5D distribution $(R, \phi, z, p_{\parallel}, p_{\perp})$. For each particle a random gyro angle is sampled and used to create a local orthonormal basis aligned with the magnetic field, allowing the conversion of the 2D momentum space into 3D velocity components (v_x, v_y, v_z) . The result is a 6D set of markers that can either be directly fed into Serpent2 or used to generate a binned 6D distribution for further sampling. However, this process is computationally expensive, and since the original distribution is stored in discrete bins, it lacks continuity. This forces resampling and interpolation, which can also introduce a loss of accuracy.

To overcome these limitations, a new AFSI5 source routine is presented in the following section. This routine is designed to directly generate the desired number of markers in the 6D Cartesian phase space, avoiding the need for costly post-processing and reducing the loss of accuracy associated with resampling and interpolation.

4.3 Design of the new AFSI5 source routine

In this section the design of a comprehensive neutron source routine is presented. This new algorithm aims to model fusion reaction from two arbitrary reactant populations, but in a totally different fashion than the current AFSI5 routine. To overcome the issues enlightened in section 4.2, the new sampling routine has to achieve the following requirements:

- Realistic representation of the fusion source rate.
- Direct coupling between ASCOT5 chain-of-tools and Serpent2.
- Provide directly the neutrons defined as (x, y, z, u, v, w, E) , without the need to store them in a discrete distribution.
- Low computational time even when a high number of markers is sampled.

For what concerns the current implementation of AFSI5, the simulation is launched with a Python script, by choosing the correct operation mode in accordance to the reaction that is going to be modeled (thermonuclear, beam-thermal, beam-beam). Here, the inputs for the core algorithm are processed and fed into the C side of the

code.

To meet the desired goal with the new routine, it is straightforward to see that this hybrid approach has to be preserved, writing both a new Python side and C side of the code. The approach followed to produce the new sampling routine was bottom-up, beginning by investigating how to produce the output in the desired format. This led to two important considerations: firstly the output has to be accurate even when the reactant distribution is highly non-uniform as in beam-thermal fusion, in addition, since Serpent2 will expect as input unweighted physical neutrons, it is more efficient to have the output as a list of particles, shifting the focus of the sampling on the particle itself.

These considerations conducted to the fact that a simple reshaping of the output was not enough, in fact the current Monte Carlo approach uses a grid-based weighted sampling, where the code loops over spatial cells and to each fusion product marker is assigned a weight proportional to the fusion rate in its cell. This means that the output distributions are weighted: a single marker can represent many fusion events. Therefore, a rejection sampling based approach seemed to be the most suitable method to fulfill the new output requirements and at the same time provide good accuracy even with highly peaked reactant distributions. This method randomly samples the reactant positions from the entire plasma volume, for each reactant pair, it samples the momenta and calculates the local fusion probability, depending on this fusion rate the product particle is either accepted or rejected.

It is crucial to highlight that this new routine will not replace the already existing one, as it is specifically designed to produce a neutron source for Monte Carlo neutronics codes, while the previous one is the best choice to study products distribution and fusion rate calculations, so it can be seen as an additional functionality in AFSI5.

4.3.1 C algorithm

The core algorithm, as in the original implementation, is written in C to ensure an efficient execution even when a large number of markers are sampled. Before describing the routine in detail, it is useful to outline the required inputs, these are not directly defined by the user but are instead provided through the Python interface, which handles data preparation and parameter passing.

The main difference in inputs compared to the original AFSI5 routine, as summarized in Table 4.1, lies on the introduction of the maximum fusion rate parameter and the array containing the cumulative distributions of the beam momenta for each spatial cell. These quantities are essential for implementing the rejection sampling algorithm and to ensure that the resulting neutron source accurately reflects the underlying physics of the fusion reactions.

Finally, the output of the C routine is a seven-dimensional array containing each neutron's spatial coordinates, velocity components, and energy. This array is already formatted for direct use in Serpent2. Depending on the environment, it can either be saved as a binary source file or, when running within the Kraken [12] framework, passed directly to Serpent2 without intermediate storage. This design minimizes post-processing time and simplifies the coupling between ASCOT5 and Serpent2.

Original AFSI5 sampling	New Rejection sampling
afsi_data: Contains reactants, plasma and B data for AFSI5	afsi_data: Contains reactants, plasma and B data for AFSI5
n: Number of Monte Carlo samples to be used in each cell	n: Total number of Monte Carlo samples required as output
prod1: Distribution data for product 1 output	S_max: Maximum fusion source rate
prod2: Distribution data for product 2 output	cumdist_all: Cumulative Distribution Functions of the beam momenta, for each spatial cell
—	prod2: 7-Dimensional array for product output (x, y, z, u, v, w, E)

Table 4.1: Input comparison between the original and the new sampling routine

The new inputs come from the differences between the two Monte Carlo algorithms, since the new routine relies on a rejection sampling approach, the maximum fusion rate is the key parameter: it sets the upper bound used to decide whether a sampled particle is accepted or rejected. This value must be carefully chosen to preserve the correct statistical weighting of the source and thus maintain physical accuracy. At the same time, the maximum fusion rate directly affects the efficiency of the algorithm. If it is set too high, the rejection rate increases and the sampling becomes unnecessarily slow. Conversely, if it is underestimated, the normalization of the source becomes inconsistent, leading to biased results. Therefore, defining this parameter properly is essential not only for the accuracy of the source distribution but also for the computational performance of the new routine. In the current framework, these parameters are computed on the Python side of the code, which generates the required inputs and subsequently launches the sampling algorithm.

The whole sampling algorithm is wrapped in a *while* loop, as outlined in the flow chart of Fig.4.2. Here the processes involved are the following: firstly for each fusion event the spatial coordinates (r, ϕ, z) of the reactant pair are randomly sampled over the volume defined with the input spatial coordinates, and for each sampled location the magnetic field, plasma density and temperature are evaluated. Subsequently the reactant momenta are sampled according whether a beam distribution is present or the reaction is thermonuclear. For a Maxwellian population, the momentum components $(p_{\parallel}, p_{\perp})$ are sampled according to the analytical model described by Eq. (9) in [13], depending only on plasma temperature and density. Conversely, for a beam ions population the momentum components are directly sampled from the input distribution, which is typically stored as an histogram and describes the probability density of each momentum values within specific spatial cell. The sampling ensures that the generated markers reflect the beam distribution and is implemented using the cumulative distribution functions provided as input and is implemented as follows:

Algorithm 1 Beam Momentum Sampling

```

 $r \leftarrow \text{random}(0, CDF[\text{max}])$ 
for each momentum bin  $j$  do
  if  $CDF[j] > r$  then
    select bin  $j$ 
  break
end if
end for
 $(p_{\parallel}, p_{\perp}) \leftarrow \text{get momentum from } j \text{ bin}$ 

```

With this approach the sampled momentum values are derived directly from the histogram bins, ensuring that they accurately reflect the physical properties of the beam.

The products velocities are then computed using the conservation laws described in section 2.2, with the only difference that the construction of 3D velocity space from the 2D momentum space is handled differently than in the original implementation. In the five-dimensional grid-based routine in fact, the magnetic field is always aligned with the Cartesian z-axis, this form is good when the products are then converted again in $(p_{\parallel}, p_{\perp})$ momentum space keeping the code lighter. However, this is not suitable for a neutron source in Serpent2, because it does not represent the physical system.

The new 3D velocity space is constructed on a generalized magnetic field orientation. The unit vector defines the parallel direction, $\hat{b} = \frac{B}{|B|}$ an orthogonal unit vector e_1 is then constructed, and a second perpendicular direction is defined as $e_2 = \hat{b} \times e_1$, so that $\{e_1, e_2, \hat{b}\}$ forms a right-handed orthonormal basis. Particle velocities are finally projected on this basis:

$$v_x = \frac{p_{\parallel}}{m} \hat{b}_x + \frac{p_{\perp}}{m} \cos(\theta) e_{1,x} + \sin(\theta) e_{2,x} \quad (4.1)$$

$$v_y = \frac{p_{\parallel}}{m} \hat{b}_y + \frac{p_{\perp}}{m} \cos(\theta) e_{1,y} + \sin(\theta) e_{2,y} \quad (4.2)$$

$$v_z = \frac{p_{\parallel}}{m} \hat{b}_z + \frac{p_{\perp}}{m} \cos(\theta) e_{1,z} + \sin(\theta) e_{2,z} \quad (4.3)$$

with θ representing the sampled gyro-angle.

The density for the beam population is simply computed by summing the histogram values in the specific cell and dividing by the cell volume.

The fusion reaction rate is finally computed as follows

$$s = \rho_1 \rho_2 \langle \sigma v \rangle \quad (4.4)$$

and the particle is accepted by if $r < \frac{s}{S_{max}}$, with r as random number on the unit interval. This approach guarantees that the accepted events are consistent with the physical properties of the fusion reaction, such as the reactant densities, velocity distribution, and reaction probability. By rejecting events that do not meet this criterion, the algorithm ensures that the sampled fusion events accurately represent the underlying physics of the fusion process.

The particle is then stored in the 7D array in the form (x, y, z, u, v, w, E) and the while loops starts again until the n number of required particles for the Serpent2

simulation is generated.

The final output of the sampling routine is the product array containing the complete set of neutron data. This array is already formatted for direct use in Serpent2. In standalone mode, it is written to a binary file where two additional columns are appended to match the Serpent2 source format: the emission time, set to 0.0 s since the particles represent an instantaneous source, and the statistical weight, set to 1.0 as each sampled particle corresponds to one neutron. The resulting file can then be read by Serpent2 using the *src sf* option, as described in [14]. On the other hand, when executed within the Kraken environment, the array can instead be passed directly to Serpent2 through the Python interface, bypassing file I/O entirely and enabling a more efficient coupling between the two codes.

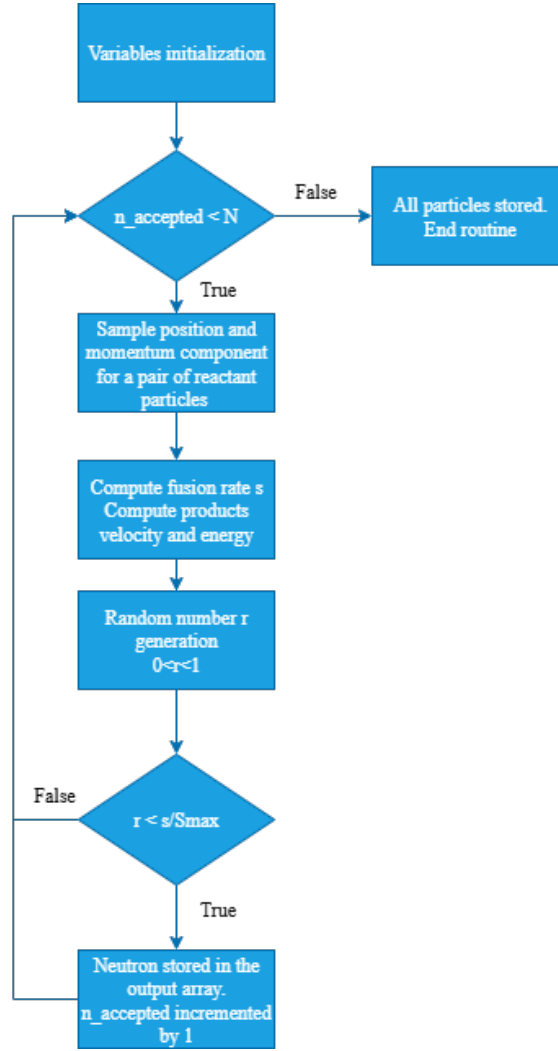


Figure 4.2: Scheme of the Monte Carlo rejection sampling algorithm

4.3.2 Python interface and pre-processing

The Python interface acts as the front end of the new source routine, handling all the pre-processing and data preparation required before calling the C algorithm. Its main role is to load the plasma and fusion data, compute auxiliary quantities such as the maximum fusion rate and the cumulative momentum distributions, and pass these as structured inputs to the compiled C routine.

Firstly, the maximum fusion source is calculated according to:

$$S_{max} = \rho_{1,max} \rho_{2,max} \langle \sigma v_{rel} \rangle_{max} \quad (4.5)$$

where the quantities depend on the specific reactant pair. The calculation therefore requires both the plasma data and, when present, the beam ion distribution.

When both reactant populations are thermal, each ρ_{imax} corresponds to the maximum ion density in the plasma, and the characteristic velocity is obtained from the thermal speed $v_i = \sqrt{\frac{2T_e}{m}}$, with T_e as the plasma temperature.

Conversely, when one of the reactants is a beam population defined by the five-dimensional distribution $f_B(R, \phi, z, p_{\parallel}, p_{\perp})$, the velocity of the beam ions is calculated from the momentum components as $v_B = \frac{\sqrt{|p_{\parallel}|^2 + |p_{\perp}|^2}}{m}$.

The maximum beam density $\rho_{B,max}$ is determined by integrating the spatial distribution over the momentum space and identifying the spatial cell with the highest number density.

Finally, the fusion cross section σ is evaluated using the Bosch–Hale parametrization [15], which expresses the reaction cross section as a function of the center-of-mass energy, $E = \frac{1}{2} \frac{m_1 m_2}{m_1 + m_2} v_{rel}^2$. The maximum value $\langle \sigma v_{rel} \rangle_{max}$ is then taken over the relevant energy range to form S_{max} . This quantity plays a crucial role in the rejection sampling algorithm described earlier, as it defines the upper bound for the fusion rate and directly affects both the physical accuracy and the computational efficiency of the sampling process.

The inputs required to run the new AFSI5 operational mode from Python are summarized in Table 4.2.

Input data	Description
n markers	Number of neutrons needed as a source for the Serpent2 simulation.
reaction	Fusion reaction selected to be modeled.
R	Array containing the radial coordinates grid for the neutron source.
PHI	Array containing the angular coordinates grid for the neutron source.
Z	Array containing the axial coordinates grid for the neutron source.
marker file	Name and path of the output file containing the particles.
beam	(Optional) Beam ions distribution to model beam-thermal reaction, can be pre-defined or obtained from BBNBI5+ASCOT5 as in Sec 4.1.

Table 4.2: Input data required to launch the new AFSI5 operational mode from a Python script.

The Python–C interface is designed to automatically determine the appropriate operational mode, eliminating the need for manual selection. By default, the

routine assumes two Maxwellian populations and operates in thermonuclear mode. Beam–thermal fusion mode is activated automatically when a beam distribution is provided as input. In this case, the pre-processing step additionally constructs arrays containing the cumulative distribution functions (CDFs) of particle momenta for each spatial cell.

Unlike the original AFSI5 implementation, the new sampling routine does not rely on a predefined grid or binning scheme. The (R, Φ, Z) definitions listed in Table 4.2 are used only during thermonuclear pre-processing to estimate the maximum fusion rate across the domain. For beam–thermal simulations, it is sufficient to specify only the spatial boundaries of the computational domain. This region does not need to cover the whole plasma volume, indeed the user can restrict it to areas where fusion events are expected, improving sampling efficiency by reducing the number of rejected markers.

5. Verification and Benchmarking

Any newly developed computational routine must undergo a process of verification and benchmarking to ensure that it produces correct and reliable results in practice. While in the previous section the focus was on the theoretical aspects and implementation of the Monte Carlo algorithm, this chapter is dedicated to the verification and benchmarking of the new AFSI5 routine.

The verification process is carried out by testing the routine with specific cases to ensure that the physics involved is correctly reproduced for both operation modes, thermonuclear and beam-thermal. On the other hand, benchmarking the code against the original implementation is useful to show the accuracy of the results and the improvement in the computational performance.

These tests together provide confidence in the use of the results as neutron source in neutron transport simulations.

5.1 Simulation setup

To properly verify the new AFSI5 routine, it is necessary to assess its accuracy in both operational modes. For this reason, the configuration selected for the tests is the VNS tokamak. As described in Sec. 2.4, this device relies on four tangential beamlines to produce a high-energy neutron flux for material testing, making it particularly suitable for validating the beam-thermal operation mode, since most of the neutron production originates from interactions between the injected neutral beams and the thermal plasma.

Moreover, working with a configuration where the resulting neutron emission is expected to be non-isotropic, due to the non-Maxwellian reactant population, is useful for an assessment of the impact of the new source on neutronics simulations results.

In current Serpent2 analyses, the neutron source is typically approximated by an analytical isotropic distribution, which does not capture the actual spatial and angular features of the emission. For this reason, while the validation presented in this chapter focuses on verifying the accuracy of the new source routine itself, the following chapter is entirely dedicated to the use of the new neutron source in Serpent2, showing if a more realistic, directionally dependent source affects neutron transport calculations.

To further evaluate the algorithm's performance at higher beam energies and plasma temperatures, additional tests are carried out using the ITER baseline scenario.

Table 5.1 summarizes the main input parameters used for the generation of the beam-ions distribution through BBNBI5 and ASCOT5 simulations. These include the geometrical configuration of the VNS tokamak, the plasma and magnetic field conditions, and the operational parameters of the neutral beam injectors. The same plasma and geometry data are later used in AFSI5 to generate the corresponding neutron source, ensuring consistency across the workflow. Figure 5.1 shows the plasma electron temperature and density maps, providing a first indication of where the peak fusion rate is expected to occur. As seen, both temperature and density reach their maximum values around $R = 2.65\text{ m}$, $z = 0.10\text{ m}$.

Input data	Value
Major radius	2.53 m
Minor radius	0.55 m
On-axis magnetic field	5.4 T
Plasma current	1.76 MA
Electron Temperature	10-14 keV
Number of injectors	4
Injected species	D
Beam energy	120 keV
Beam power	3.3 MW

Table 5.1: VNS simulation input data




Figure 5.1: Electron temperature and density in the RZ plane

Using these parameters, BBNBI5 was employed to model the ionization of deuterium neutrals injected by the NBIs, producing 40 000 fast-ion markers. These markers were then passed to ASCOT5 for a slowing-down simulation to obtain the five-dimensional distribution of the beam ions, which is required as input for the

beam–thermal fusion modeling in AFSI5.

While BBNBI5 performs a relatively lightweight computation to generate the initial fast-ion population, the subsequent slowing-down phase in ASCOT5 is significantly more computationally demanding. For this reason, the full workflow is best executed on an HPC cluster to improve scalability and reduce runtime, since even with 40 000 markers the simulation can take several hours to complete.

The simulation domain covers the whole 360° plasma geometry, as the 4 NBIs are located within the first 90 degrees sector of the device, as shown in Figure 2.7. The ion markers generated by BBNBI5 correspond to newly ionized particles, which are therefore still close to their injection location and retain the directionality of the NBIs. However, after the slowing-down simulation, this initial angular dependency is lost as the transport of beam ions in the plasma tends to homogenize their distribution along the toroidal direction. Therefore, the output beam distribution from ASCOT5, is a 5D histogram with shape (80, 1, 160, 80, 40) corresponding to 80 bins in the radial direction, 160 and only 1 bin in the toroidal, assuming axisymmetry. The momentum space is discretized with 80 bins for the component parallel to the magnetic field and 40 for the perpendicular component.

This discretization provides sufficient resolution to capture the relevant spatial and momentum features of the beam ions while keeping the HDF5 data size manageable. Figure 5.2 shows the spatial distribution of the slowed-down beam ions in the R–Z plane. The injected deuterium population is initially deposited along the neutral beam trajectories, but as the ions slow down and are transported by the magnetic field, the distribution broadens and extends over the plasma volume.




Figure 5.2: Beam R-Z cross section, after the slowing down

The momentum-space distributions are shown in Fig. 5.3. These plots highlight

that the preferential injection direction of the neutral beams is largely preserved even after the slowing-down process. This effect is evident in the p_{\parallel} component, where the slowed-down ion distribution partially overlaps with the thermal Maxwellian background but remains shifted toward positive values. This shift indicates that a fraction of the injected particles retain a significant component of their initial parallel velocity.




Figure 5.3: Momentum component of D beam, after the slowing down

Lastly, in Figure 5.4 the energy spectrum is plotted, the energy distribution reflects the typical slowing-down behavior, with a broad peak at intermediate energies.




Figure 5.4: Energy distribution of D beam, after the slowing down

When the beam distribution is not provided as input, the sampling algorithm runs in the default thermonuclear mode, where both reactant populations are assumed

to follow a Maxwellian energy distributions. The resulting particle momenta are shown in Fig. 5.5. The main difference with the beam–thermal case is evident in the parallel component, for a Maxwellian population it is Gaussian distributed with zero mean, reflecting the absence of any preferred direction in the thermal plasma.




Figure 5.5: Momentum components of a Maxwellian population in a 10 keV plasma

To summarize, using the VNS configuration described above, both thermonuclear and beam–thermal fusion processes are modeled with the new AFSI5 neutron source routine. The following section is dedicated to the verification and validation of the obtained results. In particular, the validation consists of comparing the neutron source distributions and energy spectra produced by AFSI5 against analytical models and the original AFSI5 implementation, to ensure that the new algorithm correctly reproduces the expected physical behavior. The benchmarking analysis will then assess the computational performance and efficiency improvements of the new implementation.

5.2 Validation

5.2.1 VNS Thermonuclear fusion

The first validation test concerns the thermonuclear operation mode of the new AFSI5 routine. In this configuration, the beam distribution is not provided and therefore both reactant populations follow Maxwellian velocity distributions, and the resulting neutron emission should reproduce the expected isotropic behavior and energy spectrum determined by the plasma temperature. The validation is performed by comparing the source produced by the rejection sampling algorithm against an analytical model and the source obtained with the original AFSI5 implementation. The new source routine generates 10^7 neutrons, to ensure a good statistical accuracy, while the original AFSI5 routine produced the output distribution generating 10^5 Monte Carlo samples per spatial cell.

Figure 5.6 presents on the left the resulting R-Z distribution of fusion neutrons

obtained with the new routine, that is compared with the results from AFSI5's original implementation showed on the right hand side. While the latter algorithm generates the output already in cylindrical coordinates, the new source routine results are expressed in the Cartesian reference system, therefore a trivial coordinate transformation is performed to be consistent in the comparison.

The plot shows that the new algorithm correctly reproduces the main spatial characteristics of the neutron emission, with the peak located in the central region of the plasma at coordinates $R = 265 \text{ cm}$, $z = 10 \text{ cm}$ enlightened with the red dot, where both temperature and density reach their maximum values as shown in Figure 5.1. The obtained spatial distribution also perfectly matches the one produced by the original grid-based sampling implemented in AFSI, confirming the consistency between the two methods in representing the source shape and density profile.




Figure 5.6: Comparison of the R-Z distribution of thermonuclear fusion born neutrons

Although the two algorithms rely on different sampling strategies, their output shows excellent agreement. The rejection sampling method directly samples from the continuous probability density function of the fusion rate, while the grid-based approach computes weighted averages over discretized cells. Despite these methodological differences, both correctly reproduce the same spatial pattern of neutron production across the plasma volume, confirming that the overall fusion rate distribution is consistently estimated.

The neutrons emission velocity is expected to not have a preferential direction as the fusion events between two Maxwellian populations are isotropic. In Figure 5.7 the direction vectors of the generated neutrons from the new source algorithm are presented. A direct comparison with the original AFSI5 output is not possible, as the latter provides a five-dimensional product distribution where the velocity space is described in terms of the parallel and perpendicular momentum components (p_{\parallel}, p_{\perp}) or kinetic energy and pitch (E_{kin}, ξ), therefore not providing a straightforward description of the direction vectors. For this reason, when AFSI5 is used to produce a neutron source for Serpent2 transport simulations, the velocity space is analytically

sampled within Serpent2. Consequently, the validation of the new algorithm is made by comparing its sampled direction vectors with those obtained using Serpent2's built-in isotropic analytical sampling routine [14].




Figure 5.7: Neutrons direction vector

To obtain the results shown in Fig. 5.7, 10^7 markers were sampled in both cases to compare the Monte Carlo routine with the analytical isotropic emission. As shown, the resulting angular distributions are in good agreement with the analytical model, confirming that the new routine correctly reproduces the isotropic neutron emission expected from thermonuclear fusion between thermal reactants.

Lastly, the neutron energy spectrum obtained with the rejection sampling algorithm is validated against the original AFSI5 routine and the analytical model based on the Maxwellian distribution of the reactants [3]:

$$f(E_n)dE_n = dE_n \exp \left[-(E_n - \langle E_n \rangle)^2 / \frac{4m_n T \langle E_n \rangle}{m_n + m_\alpha} \right] \quad (5.1)$$

with E_n and m_n are the energy and mass of the product neutron, $\langle E_n \rangle$ is the mean energy, T is the ion temperature and m_α is the mass of the alpha particle.

As shown in Figure 5.8, the neutron energy spectrum produced by the rejection sampling algorithm is in very good agreement with the analytical model, confirming that the new sampling method correctly reproduces the expected thermal behavior. In contrast, when using the original AFSI5 implementation, a clear inconsistency

emerges between the two available output bases. When the routine operates in the momentum-space basis (p_{\parallel}, p_{\perp}), the resulting energy spectrum deviates significantly from the typical Gaussian shape. On the other hand, using the kinetic energy and pitch-angle basis (E_{kin}, ξ) gives results consistent with both the analytical model and the rejection sampling.

This discrepancy indicates a possible issue in the transformation from momentum to energy space in the original implementation. Overall, the close agreement between the rejection sampling results and the analytical spectrum demonstrates that the new algorithm correctly captures the underlying physics of thermonuclear fusion reactions.

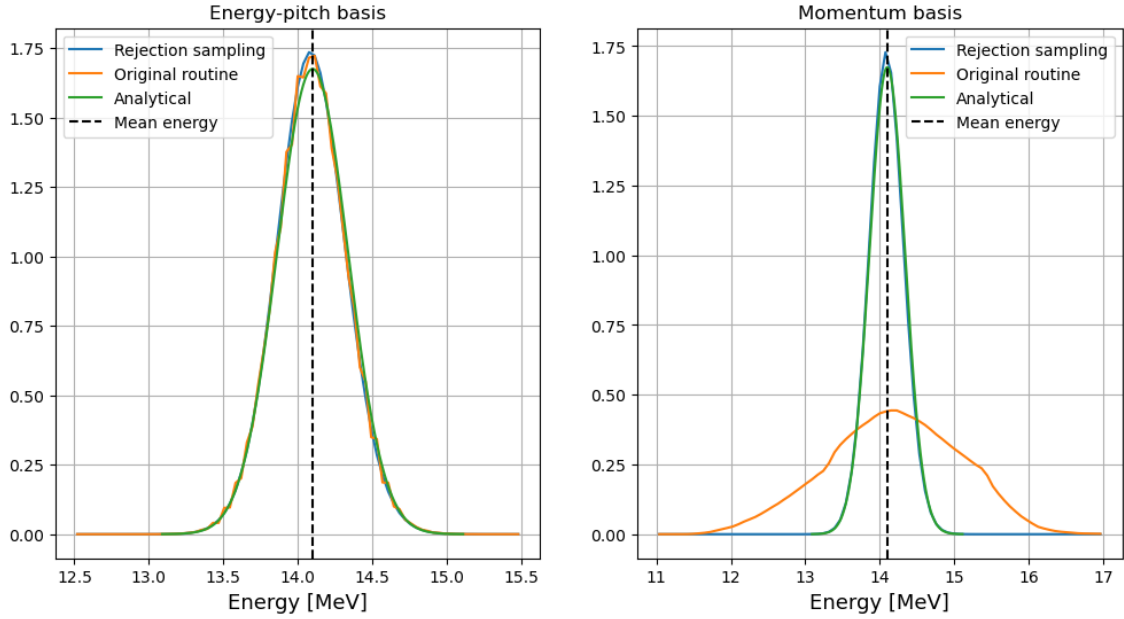


Figure 5.8: Comparison of the thermonuclear neutrons energy spectra, obtained with AFSI's original routine, new rejection sampling routine and the analytical model

5.2.2 VNS Beam-thermal fusion

The validation of the beam-thermal operation mode focuses on the capability of the routine to accurately model the fusion rate, and to reproduce the anisotropic neutron emission that characterizes the reaction between energetic beam ions and the thermal background plasma. This test is carried out using the same VNS configuration described previously, where the injected deuterium beam interacts with a thermal tritium target. Conversely to the previous section, the comparison is done only against the results obtained from AFSI5, since no reliable analytical formulation was found to model the beam-thermal fusion. The analysis aims to verify again that the new routine accurately represents both the spatial distribution, direction vectors and energy spectrum of the emitted neutrons.

The simulations are carried out generating 10^7 neutrons with the rejection sampling algorithm and 10^5 markers per spatial cell in the original AFSI5 implementation. In the latter case, the energy-pitch basis, which in the thermonuclear mode demonstrated better agreement with the analytical model, is not available when modeling beam-thermal fusion. Therefore, all beam-thermal simulations are done using the default momentum-space basis, where the output neutron distribution is represented

as $(R, \phi, z, p_{\parallel}, p_{\perp})$.

As for the thermonuclear mode, the first validation step verifies that the neutron are generated correctly within the spatial geometry, consistently with the beam and thermal reactant distributions Figs. 5.1, 5.2.




Figure 5.9: Comparison of the R-Z distribution of beam-thermal fusion neutrons

The results plotted show that compared to the thermonuclear results 5.6, the emission profile appears noticeably broader, which is expected due to the higher energy of the beam ions involved in beam–thermal fusion. The rejection sampling results show good agreement with the original AFSI5 routine in terms of the peak location, which is correctly found at $R = 2.65\text{ m}$, $z = 0.10\text{ m}$, which corresponds to the region of maximum reactant density and temperature. However, the distribution produced by the new routine is slightly wider than the one from the original implementation, likely due to the smoother sampling of the probability density function that is not constrained by predefined grid boundaries.

Overall, the comparison confirms that the new source routine accurately reproduces the expected spatial localization and general shape of the beam–thermal neutron emission within the VNS plasma geometry.

After assessing the accuracy of the spatial distribution, a less trivial analysis is conducted on the emission direction of the neutrons. As anticipated, since the beam ions are non-Maxwellian, it is reasonable to expect a non-isotropic direction of flight of the source neutrons, with a preferential direction related to the beam ions momentum space shown in Fig. 5.3.

By recalling the kinematic formulation of a fusion reaction derived in Section 2.2, it is clear the dependence of the reactant velocities on the products. The momentum components of the thermal population are sampled from temperature and density data, whereas the beam ion ones are provided in the five-dimensional distribution obtained after the slowing down. Therefore, the first step of the analysis is to verify that the algorithm described in Sec. 4.3.1 correctly samples the reactant momenta

from the given beam distribution. This verification is crucial, as the neutron emission velocity directly depends on the reactant velocities, and in the absence of a direct benchmark, this step represents the only available means of validating the emission characteristics.

In practice, this is done by plotting the sampled reactant momenta using the mentioned algorithm and comparing their statistical properties and histograms with those of the input data, confirming that the sampling process reproduces the intended distribution.

The analysis is then extended to the center-of-mass velocity of the reacting ion pairs, computed from the sampled reactant velocities. Usually, the velocity components in the new routine are represented in the cartesian reference system. However, in this validation, the velocity space is instead kept in the magnetic field aligned coordinate system used by ASCOT5, consisting in one component parallel with the magnetic field and two perpendicular components defined through the random gyro angle. In this basis, a displacement of the CM velocity distribution along the parallel direction provides a clear indication of the anisotropy introduced by the beam ions and thus of the expected preferential neutron emission direction in the laboratory frame.

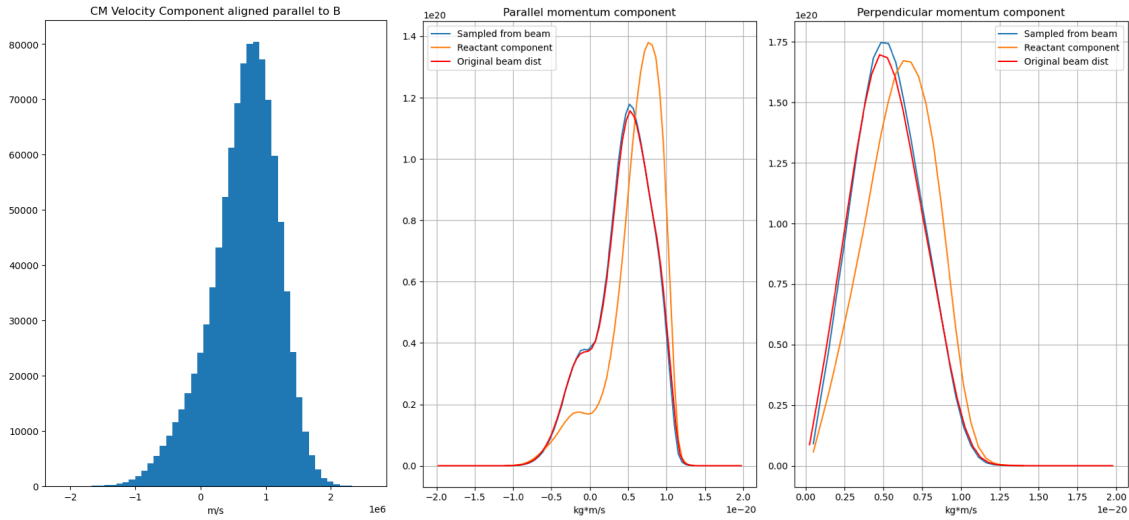


Figure 5.10: a) CM velocity component aligned parallel to B of beam-thermal fusion reactants. b) and c) Analysis of the beam momenta sampling routine. Comparison between the sampling in accordance of the fusion rate and raw sampling from input data

As shown in Figure 5.10b and c, the sampled reactant beam ions of the accepted particles exhibit a momentum distribution that follows the shape of the input slowed-down beam, although not perfectly overlapping. This deviation arises because the rejection sampling algorithm accepts or rejects samples based on the fusion rate, effectively filtering the reactants. Consequently, the accepted particles do not reproduce the exact beam distribution. For reference, when the momentum sampling algorithm is executed independently, without applying the rejection step, the resulting distribution fully overlaps with the input one, as also shown in the figure. In the same figure, the component of the CM velocity aligned with the magnetic field is also reported. The CM distribution clearly exhibits the a positive shift arising from the reaction kinematics. Indeed, since the beam reactant has a positive average parallel momentum while the thermal reactant is Maxwellian and thus centered around zero, the resulting CM velocity shows a net displacement in

the parallel direction.




Figure 5.11: Beam-thermal neutron direction vector distributions

Finally, the obtained direction vector distributions are plotted in Fig. 5.7, and as for the CM velocity, the output is represented in a field-aligned reference frame rather than in the default Cartesian coordinates.

Thanks to this representation, it is possible to appreciate a clear preferential emission direction of the generated neutrons along the parallel component of the magnetic field. This anisotropy is consistent with the reaction kinematic model, since the CM velocity exhibits a positive average component along the field direction Fig. 5.10a, when transformed back to the laboratory frame, results in a non-isotropic neutron emission.

Although the observed shift is not particularly large, it may still have an impact on subsequent neutron transport simulations, as it introduces a slight directional bias in the emission. It is worth noting that in Serpent2 the neutrons will later be expressed in Cartesian coordinates, but the physical shift in the emission direction remains present. The field-aligned representation is adopted here mainly to emphasize this effect, which would otherwise be partially smeared out in a toroidal geometry, where the anisotropy is expected to manifest as a tangential preference directed outward.

Lastly, the neutron energy spectrum is presented in Fig. 5.12 and compared with the one obtained with the original AFSI5 implementation. As shown, and consistently with the findings from the thermonuclear validation, the two spectra do not match when AFSI5 generates the product distribution using the momentum basis, making it difficult to directly validate the new routine. However, the reasoning based on the reaction kinematic presented earlier, together with the consistency previously observed between the rejection sampling and the analytical model in the thermonuclear case, suggests that the new routine provides a more physically accurate representation.

Overall, the resulting neutron energy distribution appears broader than in the thermonuclear case, as expected from the nature of the interacting reactants. This effect arises because of the presence of the high-energy beam component, which increases the range of possible relative velocities between reactants and, consequently, broadens the neutron emission spectrum.




Figure 5.12: Comparison of the beam-thermal neutrons energy spectra, obtained with AFSI’s original routine and the new rejection sampling algorithm

To summarize, the new AFSI5 routine successfully captures the key characteristics of beam–thermal neutron emission in the VNS configuration. The spatial distribution closely matches the location of the maximum fusion rate, while the directional analysis reveals a slight anisotropy along the magnetic field, in agreement with the expected reaction kinematics. Correspondingly, the neutron energy spectrum is broader than in the thermonuclear case, reflecting the contribution of the high energy beam. Together, these results demonstrate that the routine provides a consistent and physically meaningful representation of the beam–thermal neutron source.

5.2.3 ITER configuration

Although the validation performed on the VNS configuration confirmed the reliability of the new AFSI5 source algorithm, additional tests were carried out using the ITER plasma setup to further assess the model’s generality and behavior under different physical conditions.

Two main motivations justify this extension. First, testing the routine on a different configuration helps to further assess its performances, especially on the correct modeling of the fusion rate across the domain. Indeed, as previously discussed, the rejection sampling algorithm’s accuracy heavily relies on the definition of the maximum fusion rate parameter, as used to accept the fusion events. Therefore, this configuration provides an ideal framework to verify the algorithm’s capability to correctly reproduce the spatial dependence of the fusion rate in a different plasma and density fields.

Second, the beam–thermal case in the VNS setup involved relatively low-energy ions $\approx 120 \text{ keV}$, whereas in ITER the injection energy reaches up to $\approx 1 \text{ MeV}$. Consequently, stronger anisotropy effects are expected in the neutron emission, due to the higher CM velocity of the reacting pairs.

For these reasons, the reported results in the ITER configuration focus firstly on the comparison of the RZ spatial distribution for the thermal case Fig.5.15, and then on the energy and direction vector distributions of the beam–thermal source Fig.5.17, to capture the anisotropy effects described. This complementary analysis allows evaluating the response of the new source routine under higher-energy beam conditions and anticipating the possible implications for subsequent Serpent neutron transport simulations.

The beam ions distribution is produced implying BBNBI5 and ASCOT5 with the same workflow described in section 5.1 for the VNS configuration. ITER’s plasma, geometrical and magnetic field data are listed in Tab.5.2 while plasma density and temperature maps are reported in Fig.5.13.

Input data	Value
Major radius	6.2 m
Minor radius	2.0 m
On-axis magnetic field	5.4 T
Plasma current	15 MA
Electron Temperature	20-25 keV
Number of injectors	2
Injected species	D
Beam energy	1 MeV
Beam power	16.5 MW

Table 5.2: VNS simulation input data




Figure 5.13: ITER’s plasma and density maps

As shown in Figure 5.14, the higher beam energy in ITER leads to a noticeably larger shift in the slowed-down ion distribution along the momentum component parallel to the magnetic field, which is expected to increase the anisotropy of the

resulting neutron emission compared to the VNS case.




Figure 5.14: Momentum components of ITER D beam, after the slowing down

Using the same number of markers as in the VNS simulations, both the new rejection sampling routine and the original AFSI5 implementation were run to generate neutron sources for the thermonuclear and beam–thermal modes. The resulting distributions are compared to verify the consistency and accuracy of the new algorithm under different plasma conditions.

As shown in Figure 5.15, the spatial RZ distribution obtained with the rejection sampling algorithm correctly identifies the region of maximum fusion rate around $R = 6.33\text{ m}$, $Z = 0.52\text{ m}$, in agreement with the ITER plasma density and temperature profiles as well as the spatial distribution obtained with the original routine. This result further confirms the accuracy of the new routine in modeling the fusion rate.




Figure 5.15: R-Z distribution of thermonuclear fusion neutrons in ITER

In contrast, the neutron energy spectra presented in Fig. 5.16 show good agreement between the new routine, the original implementation and the analytical model.

However, when the momentum basis is employed in the original AFSI5 algorithm, the resulting energy distribution again exhibits inconsistencies, confirming that this configuration does not correctly reproduce the expected neutron energy spectrum.

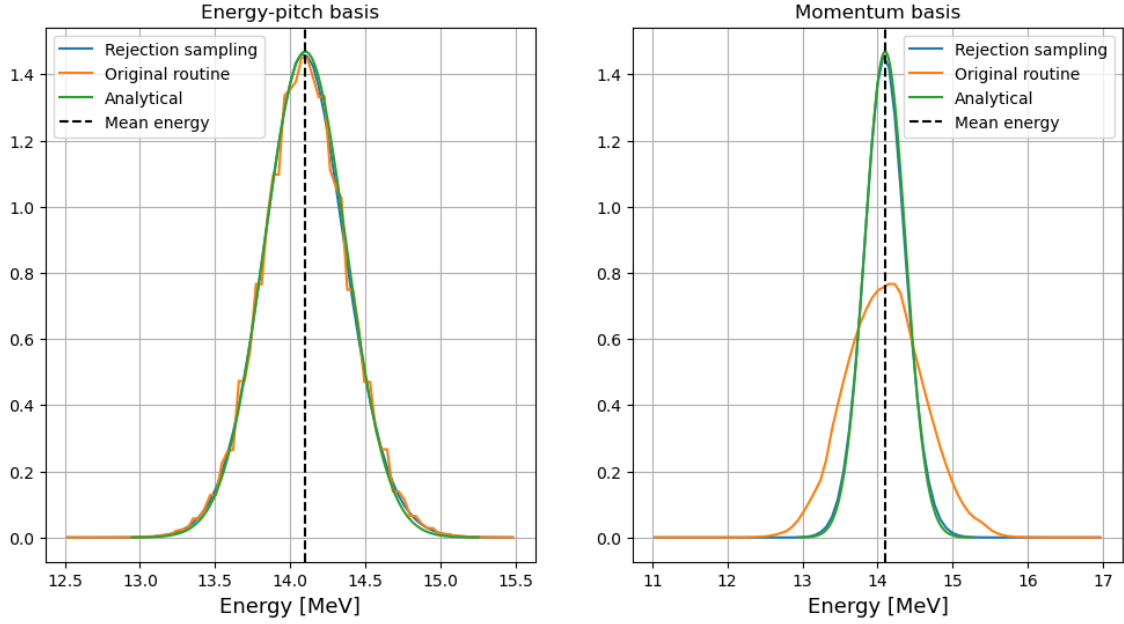


Figure 5.16: Comparison of neutron energy spectra from thermonuclear fusion in ITER

Regarding the beam–thermal operation mode, the neutron direction vectors generated with the ITER 1 MeV beam are shown in Fig. 5.17. As previously anticipated, the preferential emission is even more evident than in the VNS case, owing to the higher energy of the deuterium ions, which retain part of their injection velocity after the slowing-down process. Consistently, the resulting neutron energy spectrum plotted in Figure 5.18 appears broader, reflecting the wider range of reactant relative velocities associated with the more energetic beam.




Figure 5.17: Direction vectors of neutrons generated from beam-thermal fusion in ITER reactor




Figure 5.18: Beam-thermal neutron energy spectrum in ITER reactor

The analyses conducted on VNS and ITER configuration on the new source algorithm lead to the conclusion that the physics of the fusion reaction between two arbitrary populations is correctly represented. The rejection sampling method allows to have an accurate representation on the fusion rate in the plasma volume and the reaction kinematic of the fusion reaction correctly reproduces the neutron emission velocity.

Regarding the algorithm performances, the current implementation requires approximately 2 hours to generate a source of 100 million neutrons on a single CPU core. However, the algorithm is inherently parallel, as particle sampling is fully independent. Parallelizing the main sampling loop by dividing the total number of markers into independent batches, would therefore allow for nearly linear speedup with the number of cores. In practical terms, running the code on 8 to 16 cores could reduce the total computation time by roughly the same factor, bringing the runtime down to 7–15 minutes.

6. Serpent2 Results

In this chapter, the coupling between AFSI5 and Serpent2 is implemented. Neutronics analyses are carried out in Serpent2 using a realistic beam-thermal neutron source generated with the new AFSI5 routine. The objective is to evaluate the impact of the improved source description, by comparing the results against those obtained with a simplified analytical source model, in order to assess whether a more physically descriptive source affects the simulation outcomes. The focus will be mainly on the beam-thermal source in the VNS configuration, since the new source algorithm provides a full description of the beam-induced non-isotropic emission of flight in the Cartesian coordinates, while when using the analytical source, it is approximated as isotropic.


The results will compare the neutron fluxes in the first wall material and in the breeding zone, to assess the neutron damage in those components.

6.1 Simulation setup

The neutronics simulations were carried out on a 120° toroidal sector of the VNS Tokamak, covering the toroidal angle section from -60° to $+60^\circ$. This reduced domain includes two neutral beam injection ports, one AUX port and one test port for the blanket modules. Using this sector geometry along with periodic boundary conditions allows for a significant reduction in computational time without compromising the representativeness of the results.

The geometry was derived from an MCNP input file provided by VTT, originally converted from a CAD model of the VNS reactor based on [7] and it includes all the major reactor components, including the plasma volume, vacuum vessel, blanket modules, first wall, toroidal and poloidal field coils. This approach was preferred over a direct STL-based geometry definition in Serpent, as it allowed preserving the original material associations already defined in the MCNP input. By importing the geometry as a constructive solid geometry (CSG) model, each cell and surface could retain its corresponding material definition, avoiding the need to manually assign materials to over two thousand separate components, as would have been required with a CAD-based approach. In Figures 6.1, 6.2 the simulation domain is represented with a top and toroidal view. As shown, the materials are represented with different colors, specifically listed in Tab.6.1.

Component	Materials	Color
Supporting material	SS316L	Purple
VV body	SS316L, B4C, H2O	Light blue
Inner and outer FW	SS316L + H2O	Red
TF and PF magnets	WP MIX	Grey
BZ (shield)	SS316L	Green
BZ (breeding)	Breeding material	Blue
Divertor dome	W, Cu, CuCrZr, H2O	Orange
Divertor body	SS316L + H2O	Dark red

Table 6.1: VNS materials definition**Figure 6.1:** Section view in XY plane of VNS simulation domain**Figure 6.2:** Section view in RZ plane of VNS simulation domain

The geometry includes all the major components of the VNS reactor, and the materials defined listed in Table 6.1 are consistent with the reference design provided

by VTT. As shown in Fig.6.1, the outboard blanket is divided in two regions: the upper half is composed of SS316L and has shielding purpose only, while the lower one is a blanket test module with an active breeding zone.

Although the precise material definitions have limited influence on the present analysis, the reactor components divided as in Tab.6.1 are essential to define the different regions where the neutron tallies are recorded. In particular, the inner and outer first walls and the breeding blanket are of primary interest, as they correspond to the regions experiencing the highest neutron flux. Indeed, the local neutron flux determines all the key quantities to assess the reactor performances, such as the neutron wall load, nuclear heating, and the tritium breeding ratio, which using 100 million neutron histories ensures sufficiently low statistical uncertainties in these areas, allowing for a reliable comparison between the different source models.

Serpent allows the evaluation of quantities of interest in terms of reaction rates through the definition of detectors inside the simulation domain. These detectors are based on stochastic result estimators and produce reaction rates associated with specific MT numbers, integrated over a defined volume or surface, direction of motion, energy, and time.

In the present domain, several volume detectors were positioned within the first wall and breeding zones by specifying the corresponding cell and material parameters to define the integration volumes within these components. These detectors are particularly relevant as they enable the quantification of neutron flux and nuclear heating in the regions that experience the highest neutron loads.

To obtain a more detailed spatial characterization of the neutron field, both cylindrical and Cartesian mesh detectors were implemented. The cylindrical mesh was employed to study the toroidal flux profiles located at the inner and outer first wall, covering the axial region from -20 to 20 cm with one vertical bin, and toroidal angle from 0 to 60 degrees discretized in 60 bins. The Cartesian mesh, on the other hand, covers the entire 120° sector to produce 3D flux maps of the simulation domain. Additionally, an energy grid ranging from 10^{-9} to 20 MeV was defined to evaluate the neutron energy spectrum in different materials, particularly within the inner and outer first wall regions, where strong flux gradients are expected.




Figure 6.3: Comparison of the realistic and analytical neutron source distributions in the RZ plane

Figure 6.3 presents the spatial distribution of neutron emission in the RZ plane for both the analytical source and the beam-thermal fusion source generated with the new AFSI5 routine. As shown, the analytical source exhibits a simple and symmetric shape, modeled as a toroidal volume source with a uniform emission profile. In contrast, the AFSI5-generated source reproduces the fusion rate distribution in the plasma, leading to a more realistic and non-uniform neutron production pattern. Moreover, as already discussed in the previous chapter, the analytical source assumes isotropic neutron emission, while the AFSI5 source correctly accounts for the anisotropic emission resulting from the directional beam reactants Fig.5.11.

Finally, in Fig.6.4 the two sources energy spectra are plotted. As noticeable, the realistic beam-thermal source presents a broader energy spectrum due to the higher energy beam ions reactant population, while the analytical spectrum is modeled with Eq.(5.1).




Figure 6.4: Comparison of the realistic and analytical neutron source energy spectra

This differences in both spatial, angular and energy characteristics is expected to have an impact on the resulting neutron flux distribution within the reactor components.

6.2 Neutronics results

In this section, the results of the Serpent2 simulations using the two different neutron sources are presented and compared. The purpose of this analysis is to assess how the more physically descriptive source generated by the new routine, influences the neutron field and related quantities, compared to a simplified analytical model assuming isotropic emission.

The FW and blanket regions are particularly relevant because they are directly exposed to the neutron flux. In particular, the outer FW and BZ were chosen as

a starting point for the analysis to observe the effect of the non-isotropic neutron direction of flight showed in the new AFSI5 beam-thermal source 5.11. Indeed, as discussed in the previous chapter, the neutrons are characterized by a preferential tangential component directed toward the outer side of the machine. Such emission pattern would in principle lead to a higher neutron flux on the outer first wall and blanket, making this area the most suitable to assess possible differences to take care of when designing the machine.

Figure 6.5 shows the neutron flux in the first centimeters of the outer FW tiles, spanning toroidally from 0° to 60° and axially from -20 to 20 cm. From this plot it is noticeable that, contrary to the expected behavior discussed earlier, the analytical source produces a higher neutron flux in this region compared to the realistic AFSI5 source.

To further quantify and illustrate this difference, Figures 6.6, reports the ratio of neutron fluxes produced by the two sources within the defined volume defined as $R = \frac{AFSI}{Analytical}$. As visible, the flux ratio remains consistently below unity, confirming that the analytical source yields fluxes in the outer components up to 10-12% higher than those obtained with the realistic AFSI5 source. Nevertheless, it is possible to appreciate that some localized regions exhibit ratios above one, meaning that the flux generated by the AFSI5 source locally exceeds that of the analytical model. The scatter plot does not provide any information on the location of these cells, therefore this spatial variation can be better observed in the map produced with the Cartesian mesh detector 6.7, showing with color red the zones where the ratio is higher than one and blue when lower. From this screenshot, we see that the realistic beam-thermal source produces a higher neutron flux in the upper part of the FW, probably due to the broader spatial distribution of the source.




Figure 6.5: Neutron flux profiles from the two different sources in the outer first wall axially spanning from -0.20 m to 0.20 m

This somewhat unexpected result can be attributed to differences in the spatial distribution of the neutron sources 6.3. In the analytical model, the neutron birth

positions are more concentrated toward the plasma periphery, resulting in a higher local source density close to the outer wall. Conversely, the AFSI-generated source, while physically more realistic, exhibits a broader spatial spread of neutron production, leading to a slightly reduced flux in this region. In addition, the anisotropic emission associated with the AFSI5 source is relatively weak due to the moderate energy of the VNS beam and is therefore insufficient to compensate for the reduced source density near the wall.




Figure 6.6: Neutron flux ratio obtained with cell based detector in outer FW




Figure 6.7: Neutron flux ratio mapped on the outer FW

Similarly, Figure 6.8 reports the ratio of the nuclear heating of the two sources in the same first wall cells. As shown, the results follow the same general trend observed for the neutron flux, as the deposited power depends on the local neutron flux. However, the difference between the two source cases is even more pronounced, with the analytical source producing a noticeably higher heating rate. This amplification can be attributed to the energy-dependent nature of the heating cross sections, which increases the sensitivity of the deposited energy to spectral differences between the two neutron sources. These spectral differences can be observed in Figure 6.9, where the neutron flux spectra in the same region are compared. Here the analytical source produces a slightly higher flux over most of the energy range, particularly around the main peak at ≈ 14 MeV, while the realistic beam-thermal source shows a marginally broader high-energy tail consistent with the source energy spectrum. Although the

overall spectral difference is modest, it occurs in the high-energy region where the heating cross sections are larger. Consequently, even this small spectral shift leads to a noticeable increase in the local heating for the analytical source.




Figure 6.8: Neutron heating ratio obtained with cell based detector in outer FW




Figure 6.9: Neutron flux energy spectrum in outer FW material

The results presented in Figures 6.10, 6.11 for the outer blanket region refer to the cells containing only the shielding material. As expected, given their position immediately behind the first wall, the trends observed here are very similar to those previously discussed. The analytical source produces a higher neutron flux, with a peak difference of approximately 10–12% compared to the realistic AFSI5 source. However, unlike in the first wall region, where the heating difference reached up to 20%, the discrepancy in the blanket cells is less pronounced and remains comparable to the flux difference. This behavior can be explained by the neutron flux spectra shown in Figure 6.12, which reveal that the two spectra become more similar in this region. The multiple scattering processes occurring in the first wall tend to soften and homogenize the neutron spectrum, reducing the spectral differences between the two sources and, consequently, their impact on the heating rate.




Figure 6.10: Neutron flux ratio obtained with cell based detector in outer BZ




Figure 6.11: Neutron heating ratio obtained with cell based detector in outer BZ




Figure 6.12: Neutron flux energy spectrum in outer BZ shielding material

Conversely, when examining the inner first wall, an opposite trend is observed. As shown in Figure 6.13, the neutron fluxes from the two sources are plotted along the toroidal direction, within the first centimeters of the inner FW and over a 40 cm axial

span. Here, the realistic beam-thermal source produces a noticeably higher neutron flux. This behavior is further confirmed by the scatter plot in Figure 6.14, which presents the ratio of the neutron fluxes obtained in the inner first wall cells. The flux from the realistic source is higher by up to about 10%, while a few localized cells still show slightly higher values for the analytical source, particularly in the lower part of the wall. A more complete, spatial representation of this trend is provided in Figure 6.15, which maps the flux ratio over the entire inner first wall and clearly shows the regions where the realistic source dominates.




Figure 6.13: Neutron flux profile from the two different sources in the inner first wall axially spanning from -0.20 m to 0.20 m




Figure 6.14: Neutron flux ratio obtained with cell based detector in inner FW




Figure 6.15: Neutron flux ratio mapped on the inner FW

Lastly, the heating results shown in Figure 6.16 follow the same general behavior, with the beam-thermal source leading to higher heating rates, consistent with the higher local neutron flux. In this case, the difference between the two sources is even more pronounced, reaching up to about 20%. Similarly to what was discussed for the outer first wall, this enhancement can be explained by the differences in the neutron energy spectra. As shown in Figure W, the beam-thermal source produces a harder spectrum, with a larger fraction of high-energy neutrons. This results in a stronger weighting of the energy-dependent heating cross sections, which amplifies the deposited power compared to the analytical source.




Figure 6.16: Neutron heating ratio obtained with cell based detector in inner FW

This neutronics analysis on the VNS tokamak has highlighted the critical importance of accurately defining the neutron source when modeling neutrons generated from non-Maxwellian populations, as the resulting impact on materials can differ significantly.

In the outer first wall region, the analytical source produces a higher local neutron flux than the AFSI5 source, which is in contrast with what would result if emission directionality were the dominant factor. This outcome indicates that, in this configuration, the spatial distribution of the neutron source plays a more dominant role than its directional anisotropy. The analytical source is more concentrated near the plasma periphery, leading to a locally higher neutron density close to the outer wall.

In contrast, the realistic AFSI5 source features a broader and more diffuse spatial distribution, which reduces the flux near the wall despite its more physical emission directionality. Furthermore, the relatively low energy of the VNS beam implies that the anisotropic emission of the fusion neutrons is only weakly pronounced, further limiting its impact on the overall flux distribution.

On the inboard side, the trend is reversed. The analytical source underestimates of about 10% the neutron flux and up to 20% the nuclear heating, compared to the realistic AFSI5 source. These discrepancies are crucial when assessing material performance and component lifetime, as they indicate that the analytical source does not adequately capture the true loading conditions. Since analytical sources are still commonly employed to model devices that rely beam-driven fusion reactions, these findings suggest that such simplifications may lead to underestimations of damage and heat loads in key components.

Overall, the analysis demonstrates that, for the VNS configuration, the spatial distribution of the neutron source governs the neutronics behavior far more strongly than emission anisotropy. Although directionality could become more relevant at the higher beam energies such as in ITER, its practical impact on ITER's neutronics is expected to remain limited. This is because NBI heating in ITER operates only for a relatively short fraction of the discharge, whereas steady-state plasma performance is dominated by thermal–thermal fusion. As a result, even if the instantaneous anisotropy were stronger, its influence on time averaged neutronics quantities during reactor operation would be comparatively minor. In contrast, the VNS relies on NBI as a steady-state driver of fusion reactions, making the accurate treatment of beam–thermal fusion and of the resulting source distribution substantially more important.

7. Conclusions

The work presented in this report was motivated by a practical gap in current fusion neutronics workflows: while analytical neutron sources remain simple and widely used, they cannot represent the spatial and momentum structure of realistic fusion plasmas, especially when one of the reactant species is non-Maxwellian such as in beam-thermal fusion. At the same time, the ASCOT5 chain of tools provides high-fidelity description of the fusion products through AFSI, but its output lacks an explicit angular description and cannot be directly used to define a full 6D neutron source in Monte Carlo transport codes. Addressing this limitation was the central objective of the project.

The work carried out at VTT focused on the formulation, implementation and integration of a configuration-agnostic routine capable of generating a fully six dimensional fusion neutron source. The method builds upon ASCOT5 but extends it with a complete reconstruction of the emission direction, a general rejection sampling approach independent of geometry or grid structure, and a Serpent2 compatible output format. The routine was validated through extensive checks on the spatial distribution, energy spectrum, directional sampling, and their consistency with the underlying physics. Finally, the resulting sources were used in Serpent2 to benchmark the neutron transport response against traditional analytical models in a beam-driven plasma configuration.

The results show that the routine performs as intended and produces physically consistent sources that reflect the characteristics of the underlying plasma. Interestingly, the Serpent2 simulations revealed that the differences between the analytical and realistic sources were mostly due to the difference in the spatial distribution rather than to the more realistic direction of emission. Indeed, in the NBI dominated test case, the comparison between the two source models revealed a behaviour opposite to the initial expectations. Based on the preferential emission direction observed in the source generated by the new routine, a higher flux on the outer first wall was anticipated when using the realistic source model. Instead, the Serpent2 results showed that the analytical source produced a higher flux up to 10% on the outer FW, while the AFSI5 generated source led to a considerably higher flux on the inner wall. The key factor behind this inversion was the spatial distribution of the neutron birth points.

The analytical model is more homogeneous and places a larger fraction of the fusion neutrons close to the outer mid-plane, which naturally enhances the contribution to the outer wall. By contrast, the AFSI5 source, even though more accurate in the neutron emission direction, is spatially broader due to the beam increasing the fusion

rate in outer radial regions as well as deeper inside the plasma. This broader spatial footprint results in a net shift of the highest flux towards the inner wall.

Overall, this demonstrates that in this configuration the spatial distribution of neutron production dominates over the differences in the momentum space description.

During the development and validation phases, several practical limitations in the existing AFSI5 workflow were identified. In particular, the handling of the parallel and perpendicular momentum basis showed inconsistencies that can distort the reconstructed neutron energy spectrum. Furthermore, the dependence of the original AFSI5 output on specific coordinate systems and grid structures heavily affects the source accuracy. These observations reinforce the need for more general, geometry-agnostic routines and highlight the value of the approach developed in this work.

Overall, the new routine provides a significant improvement in the flexibility and physical fidelity of neutron source modelling. It removes the need for analytical approximations, allows ASCOT5-based reactivity information to be used directly in Serpent2 without restrictive preprocessing, and is applicable to arbitrary plasma and device configurations. It therefore strengthens the ASCOT5–Serpent2 workflow and broadens its applicability to reactor relevant neutronics studies.

Despite these advances, some limitations remain. The routine has so far been tested only on a limited set of scenarios, and its performance should be assessed in more diverse configurations, including thermal plasmas, reactor-scale devices and fully three-dimensional magnetic geometries. At the computational level, the implementation currently runs on a single core and the generation of large sources, on the order of 10^8 samples or more, can require several hours. Parallelising the sampling and writing procedures is therefore an essential next step to improve performance and make the method practical for routine reactor studies. Beyond these points, the code could also benefit from a tighter integration between ASCOT5 and Serpent2 inside the Kraken framework. Kraken already provides a modular environment where Serpent2 operates as a high fidelity neutronics solver alongside other physics modules, and embedding the AFSI5 source routine directly into this workflow would remove much of the current ad-hoc interfacing. In practice, this would allow ASCOT5-generated sources to feed Serpent more naturally as part of a unified multiphysics pipeline, improving consistency, reducing manual handling, and aligning the toolchain with the broader development ecosystem maintained at VTT.

In summary, the work provides a general and portable method for generating realistic six dimensional fusion neutron sources and demonstrates its integration into a complete neutronics workflow. While several avenues for further improvement remain, the routine represents a step forward towards high-fidelity, configuration-independent neutron source modelling for a wide class of fusion devices.

Bibliography

- [1] Richard Fitzpatrick. *Fusion Reactions*. 2019. URL: <https://farside.ph.utexas.edu/teaching/plasma1/Fusionhtml/node7.html> (visited on 10/30/2025).
- [2] Paula Sirén, Jari Varje, Simppa Äkäslompolo, Otto Asunta, Carine Giroud, Taina Kurki-Suonio, Henri Weisen, and The JET Contributors. «Versatile fusion source integrator AFSI for fast ion and neutron studies in fusion devices». In: *Nuclear Fusion* 58.1 (Nov. 2017), p. 016023. DOI: [10.1088/1741-4326/aa92e9](https://doi.org/10.1088/1741-4326/aa92e9). URL: <https://doi.org/10.1088/1741-4326/aa92e9>.
- [3] H Brysk. «Fusion neutron energies and spectra». In: *Plasma Physics* 15.7 (July 1973), p. 611. DOI: [10.1088/0032-1028/15/7/001](https://doi.org/10.1088/0032-1028/15/7/001). URL: <https://doi.org/10.1088/0032-1028/15/7/001>.
- [4] B Appelbe and J Chittenden. «The production spectrum in fusion plasmas». In: *Plasma Physics and Controlled Fusion* 53.4 (Feb. 2011), p. 045002. DOI: [10.1088/0741-3335/53/4/045002](https://doi.org/10.1088/0741-3335/53/4/045002). URL: <https://doi.org/10.1088/0741-3335/53/4/045002>.
- [5] Gianfranco Federici. «Testing needs for the development and qualification of a breeding blanket for DEMO». In: *Nuclear Fusion* 63.12 (Oct. 2023), p. 125002. DOI: [10.1088/1741-4326/ad00cb](https://doi.org/10.1088/1741-4326/ad00cb). URL: <https://doi.org/10.1088/1741-4326/ad00cb>.
- [6] R. Villari et al. «Overview of deuterium-tritium nuclear operations at JET». In: *Fusion Engineering and Design* 217 (2025), p. 115133. ISSN: 0920-3796. DOI: <https://doi.org/10.1016/j.fusengdes.2025.115133>. URL: <https://www.sciencedirect.com/science/article/pii/S0920379625003308>.
- [7] C. Bachmann et al. «Engineering concept of the VNS - a beam-driven tokamak for component testing». In: *Fusion Engineering and Design* 211 (2025), p. 114796. ISSN: 0920-3796. DOI: <https://doi.org/10.1016/j.fusengdes.2024.114796>. URL: <https://www.sciencedirect.com/science/article/pii/S092037962400646X>.
- [8] International Atomic Energy Agency. *Fundamentals of Magnetic Fusion Technology*. STI/PUB/1945. Vienna: International Atomic Energy Agency, 2023. ISBN: 978-92-0-110721-3.
- [9] Konsta Särkimäki, Jari Varje, Marina Bécoulet, Yueqiang Liu, and Taina Kurki-Suonio. «Mechanics of ELM control coil induced fast particle transport in ITER». In: *Nuclear Fusion* 58.7 (June 2018), p. 076021. DOI: [10.1088/1741-4326/aac393](https://doi.org/10.1088/1741-4326/aac393). URL: <https://doi.org/10.1088/1741-4326/aac393>.

- [10] O. Sauter and S.Yu. Medvedev. «Tokamak coordinate conventions: COCOS». In: *Computer Physics Communications* 184.2 (2013), pp. 293–302. ISSN: 0010-4655. DOI: <https://doi.org/10.1016/j.cpc.2012.09.010>. URL: <https://www.sciencedirect.com/science/article/pii/S0010465512002962>.
- [11] O. Asunta, J. Govenius, R. Budny, M. Gorelenkova, G. Tardini, T. Kurki-Suonio, A. Salmi, and S. Sipilä. «Modelling neutral beams in fusion devices: Beamlet-based model for fast particle simulations». In: *Computer Physics Communications* 188 (2015), pp. 33–46. ISSN: 0010-4655. DOI: <https://doi.org/10.1016/j.cpc.2014.10.024>. URL: <https://www.sciencedirect.com/science/article/pii/S0010465514003701>.
- [12] Jaakko Leppänen et al. «Current Status and On-Going Development of VTT’s Kraken Core Physics Computational Framework». In: *Energies* 15.3 (2022). ISSN: 1996-1073. DOI: [10.3390/en15030876](https://doi.org/10.3390/en15030876). URL: <https://www.mdpi.com/1996-1073/15/3/876>.
- [13] Nader Mohamed. «Efficient Algorithm for Generating Maxwell Random Variables». In: *Journal of Statistical Physics* 145 (Dec. 2011). DOI: [10.1007/s10955-011-0364-y](https://doi.org/10.1007/s10955-011-0364-y).
- [14] J. Leppänen. *Serpent – a Continuous-energy Monte Carlo Reactor Physics Burnup Calculation Code*. Tech. rep. User Manual, released June 18, 2015. VTT Technical Research Centre of Finland, June 2015.
- [15] H.-S. Bosch and G.M. Hale. «Improved formulas for fusion cross-sections and thermal reactivities». In: *Nuclear Fusion* 32.4 (Apr. 1992), p. 611. DOI: [10.1088/0029-5515/32/4/I07](https://doi.org/10.1088/0029-5515/32/4/I07). URL: <https://doi.org/10.1088/0029-5515/32/4/I07>.