

A genetic algorithm to optimize the multi-group structure for the neutronic analyses of the ARC fusion reactor

Original

A genetic algorithm to optimize the multi-group structure for the neutronic analyses of the ARC fusion reactor / Aimetta, A., Abrate, N., Caravello, M., Dulla, S., Froio, A., Massone, M.. - In: NUCLEAR MATERIALS AND ENERGY. - ISSN 2352-1791. - ELETTRONICO. - 46:(2026). [10.1016/j.nme.2026.102072]

Availability:

This version is available at: 11583/3007538 since: 2026-02-11T18:45:55Z

Publisher:

Elsevier Ltd

Published


DOI:10.1016/j.nme.2026.102072

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)



A genetic algorithm to optimize the multi-group structure for the neutronic analyses of the ARC fusion reactor[☆]

Alex Aimetta ^a, Nicolò Abrate ^{a,b},* , Marco Caravello ^a , Sandra Dulla ^{a,b}, Antonio Froio ^a,
Mattia Massone ^c

^a NEMO group, Dipartimento Energia, Politecnico di Torino, Corso Duca degli Abruzzi, 24, 10129, Torino, Italy

^b Istituto Nazionale di Fisica Nucleare (I.N.F.N.), Sezione di Genova, Via Dodecaneso, 33, 16146, Genova, Italy

^c ENEA Centro Ricerche Bologna, Dipartimento Nucleare, Via dei Mille, 21, 40121, Bologna, Italy

ARTICLE INFO

Keywords:

Molten salt liquid breeder blankets
Energy group structure optimization
Genetic algorithm
Multigroup diffusion calculations

ABSTRACT

In the framework of the modeling of fusion reactors with deterministic neutronic codes, the choice of an appropriate energy grid for the generation of the multigroup nuclear properties is essential. In this work, a Genetic Algorithm is employed to optimize the energy grid employed in the nemoFoam multiphysics code to reproduce the results provided by the Monte Carlo code Serpent in terms of neutron flux, neutron power deposition and Tritium Breeding Ratio for the Affordable, Robust and Compact (ARC) fusion reactor. Different runs of the Genetic Algorithm are performed, with the aim of optimizing not only the quantities of interest separately, but also trying to combine them thanks to the definition of appropriate fitness functions. The optimization is performed starting from a pre-defined 86 groups energy grid, over which the nuclear properties and the reference quantities are evaluated with Serpent. The results show that it is not straightforward to optimize at the same time the energy grid for different quantities and that, in general, coarse energy grids are able to provide good results in nemoFoam for what concerns the ARC reactor, allowing to alleviate the computational burden of the neutronic evaluation too.

1. Introduction

Generally, the selection of the energy grid for collapsing the group constants to be used in deterministic neutronic codes is based on expert judgment, since the physical behavior of the cross sections and their relation with the physics of the reactor under study is too challenging to be dealt with on an exact basis. This is actually an open issue in the case of both next generation fission reactors (e.g., Molten Salt Reactors, Sodium Fast Reactors and Lead Fast Reactors) and fusion reactors like the Affordable, Robust and Compact (ARC) reactor [1,2], where the expertise is still quite limited. In particular, one could choose an energy grid as fine as possible to have accurate results, but this would lead to non-affordable computational times, in particular during the design and safety assessment stage, when a large number of simulations is required. In fact, during the design stage, the size and the materials of the main components of the reactor could have significant changes, leading to quite different values of the neutronic quantities. In this sense, the best energy grid should be the optimal one both in terms of accuracy of the results and of computational cost. In this work, as it will be elaborated more in detail in Section 2, it has been chosen

to optimize the energy grid in terms of accuracy, while limiting the number of groups, on a set of relevant quantities for the multi-physics design of the Breeding Blanket, such as Tritium Breeding Ratio (TBR), volumetric power deposition due to neutrons and neutron flux. Specific penalization coefficients for computational time could be introduced in future work but, as a first attempt, a maximum for the number of groups of the coarse grid has been set. It is worth noticing that the optimal energy grid can be significantly different as the size and the materials of the components of the reactor change, making it particularly important to have efficient tools to identify optimal energy group structures.

The selection of an optimal group grid is a stiff, non-linear optimization problem featured by a very large parameter space. During the last years, many works were published on this topic using different techniques (e.g., the Hierarchical Division algorithm [3], Simulated Annealing coupled with the Random Forest regression method [4–7] and Particle Swarm Optimization [8]), underlining the relevance of the subject in the field of reactor physics.

In this work, the optimization of the ARC reactor energy grid is performed employing a Genetic Algorithm (GA). Genetic Algorithms

[☆] This article is part of a Special issue entitled: 'PRD activities' published in Nuclear Materials and Energy.

* Corresponding author.

E-mail address: nicolo.abrate@polito.it (N. Abrate).

belong to the class of evolutionary algorithms and allow efficient exploration of the solution space, determining the nearly optimal energy grid minimizing or maximizing a user-defined fitness function. The choice of a GA can be justified considering that the research of the optimal energy cuts is related to sequences of energy boundaries, with each of them strongly influenced by the neighboring ones, well resembling the behavior of genes in chromosomes.

GAs were already applied in the past for the optimization of the energy grid to be used in next generation fission reactors [9–12], showing good performances in the case of molten salt reactors, sodium fast reactors and lead fast reactors. One of the most recent work [12] was dedicated to the optimization of a specific LFRs design, the Advanced Lead Fast Reactor European Demonstrator (ALFRED), using the FRENETIC code [13]. Compared to the previous studies in the field, the main innovation introduced in [12] was the adoption of a specific chromosome representation that allows to perform a flexible grid optimization by acting on both the number of groups and the position of the energy boundaries. Another idea applied also in this work is the adoption of the adjoint neutron flux for weighting the fitness function devoted to improving the accuracy on the neutron flux calculation.

This work presents the application of a GA-based optimization process specifically for the neutron energy grid in a fusion reactor's liquid breeding blanket. The optimization focuses on improving the Tritium Breeding Ratio (TBR), neutron flux, and volumetric power deposition by neutrons computed by the nemoFoam code [14], using the Serpent Monte Carlo analysis as a reference. The decision to limit the analysis exclusively to neutrons arises from the limitations of the Serpent Monte Carlo code, which is incapable of generating multi-group coefficients for photons. Consequently, with photon energy being represented by a single group, there is no scope for optimization. Theoretically, if the capability to generate multiple energy groups for photons were available, a GA optimization could similarly be applied to the photon energy grid. However, given the current focus on neutrons due to these constraints, we will present only neutron-related metrics, such as neutron volumetric power deposition and neutron flux. The neutron-focused GA will be implemented using the nemoFoam code [14]. To the authors' knowledge, this is the first time a GA is applied to optimize the neutron energy grid in a liquid breeding blanket of a fusion reactor. Given that the use of a simplified approach to model neutron transport (see Section 2.1) already implies a faster solution, if compared to the "conventional" fusion approach of using Monte Carlo simulations, the chosen objective functions for the optimization only focus on improving the accuracy of the results, rather than on reducing the computational cost.

2. Methodology

The starting point of the optimization process is the generation of fine-group constants and reference quantities on a reference grid, using the Monte Carlo Serpent code [15]. The GA selects an optimal grid among all the subsets of the reference grid, meaning that energy boundaries not contained in the reference grid cannot be selected by the GA. In the case of a neutronic and thermal-hydraulic multiphysics application, a Serpent simulation has to be performed for each temperature of interest, in order to prepare a library of multigroup nuclear properties and of reference quantities to be used in the optimization process. However, this work is focused on a purely neutronic case, thus a single Serpent simulation is sufficient. The details of the Serpent model adopted in this work can be found in [16].

The GA algorithm is represented by the yellow boxes in Fig. 1. The versatility of the code allows to employ it for any application and fitness function of interest (i.e., not only the optimization of the energy grid structure, but, e.g., also the optimization of the thickness of components in fusion reactors and/or of the lithium enrichment to maximize the TBR). In this case the `pybind11` Python package [17] is called for

Table 1

Reference values of TBR and neutron power deposition in the main components of ARC (with statistical uncertainties), evaluated with Serpent.

	VVIN	CC	NM	VVOUT	TANK
Power deposition [MW]	14.60(1)	70.31(1)	17.72(1)	23.59(1)	219.8(1)
TBR [-]	1.0848(1)				

each individual, with the aim of passing its few-group structure from the C++ genetic algorithm to the open-source Python package `COREutils` [18] developed at Politecnico di Torino. Here, `COREutils` is used to collapse the group constants from the reference fine grid to the few-group individual one, using the neutron flux spectrum evaluated by Serpent on the fine-energy grid in each region of interest of the ARC reactor. Serpent is used to evaluate the reference individual too, with respect to which the fitness functions are evaluated. Then, other Python scripts are written in order to make the generation of the nemoFoam input automatic for each individual to be tested by the GA. Once the whole initial population is evaluated with nemoFoam and each individual associated with a fitness function, the evolution mechanism implemented in the GA generates a new population, normally better than the previous one, and so on until a convergence criterion is met.

The reactor analyzed in this work is the ARC reactor, for which the fine-group cross sections and the reference quantities (i.e., the Tritium Breeding Ratio, the neutron power deposition and the neutron flux) are computed with Serpent, using the ENDF/B-VIII.0 nuclear data library [19], for the five main components of ARC: the inner vacuum vessel (VVIN), the cooling channel (CC), the neutron multiplier (NM), the outer vacuum vessel (VVOUT) and the breeding blanket (TANK). The geometry used to model the different regions is reported in Fig. 2.

To guarantee that the statistical error of the group constants is as small as possible, a coarser version of the VITAMIN-J energy grid [20], usually employed for fusion applications is conceived. This coarser version has a total number of 86 energy groups, with few groups in the thermal region and a finer discretization at higher energies, consistently with the fast spectra typical of fusion systems (Fig. 3). The number of energy groups (86) in the reference energy grid has been chosen with the aim of guaranteeing that the results provided by Serpent are statistically reliable, since having too many groups can represent an issue from this point of view. The low statistical uncertainties shown in Fig. 3 are an indication of the good quality of the Serpent results, obtained even without the need of exploiting variance reduction techniques. The reference values for TBR and neutron power deposition in the main components of ARC are shown in Table 1. In particular, a 2D model of ARC is defined in nemoFoam, exploiting the neutronic symmetry of the reactor in the toroidal and vertical directions, in order to speed up the nemoFoam simulations required to make the population of the GA evolve.

2.1. The nemoFoam code

The nemoFoam code, written in the OpenFOAM environment, was developed for the neutronic (either multigroup diffusion or multigroup SP_3 , with an additional mono-kinetic diffusion model for photons) and thermal-hydraulic solution and coupling in fusion systems. Both the modules can be used to model systems in 2D or 3D domains, with the thermal-hydraulic module based on the `chtMultiRegionFoam` available in any OpenFOAM distribution, slightly modified in order to add a term in the energy conservation equation to take into account the power deposited by neutrons and photons. nemoFoam is also provided with a module for the solution of the multigroup diffusion adjoint equation, which is a relevant feature for this work.

In this paper, only neutrons are modeled, using the diffusion module (to reduce the computational time with respect to the solution of the SP_3 equations). Currently, nemoFoam is forced to solve the photon

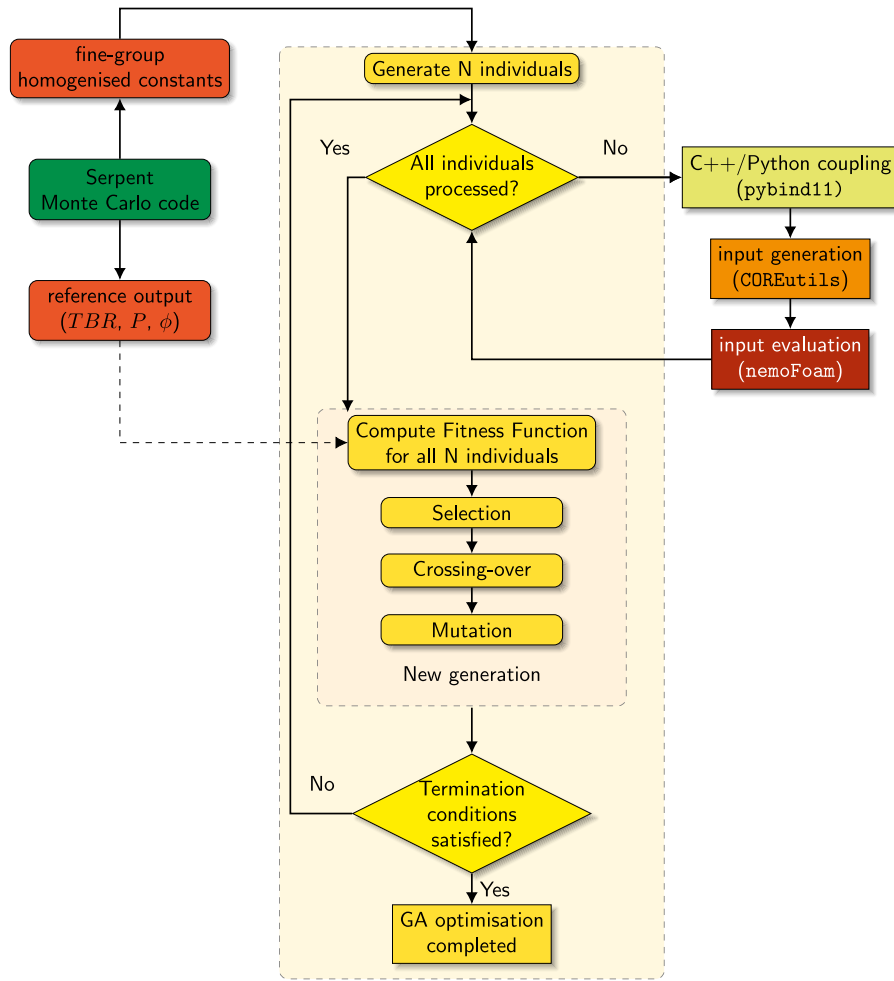


Fig. 1. Workflow diagram of the group structure optimization process involved in this work.

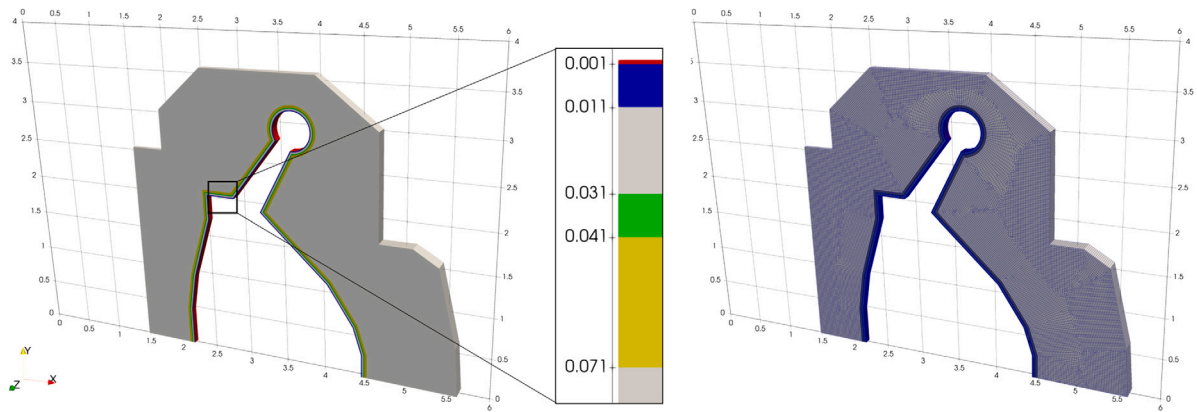


Fig. 2. 2D computational domain used in the study. Each one of the five regions (VVIN, CC, NM, VVOUT, TANK) is represented with a different color.

problem on a one-group energy grid, thus no optimization of the energy grid is possible. This is due to the fact that Serpent cannot compute multigroup scattering matrices for photons. The results of nemoFoam are represented in terms of a neutron flux field and a neutron power deposition field for each energy group of interest, which are then used to compare the results of nemoFoam with Serpent. The neutron flux distribution in the phase space is computed by solving the source-driven, multi-group diffusion equation,

$$-\nabla \cdot (D_g(\vec{r})\nabla\phi_g(\vec{r})) + \Sigma_{t,g}(\vec{r})\phi_g(\vec{r})$$

$$= \sum_{g'=1}^G \Sigma_{s,g' \rightarrow g}(\vec{r})\phi_{g'}(\vec{r}) + S_g(\vec{r}) \quad \forall g = 1, \dots, G, \quad (1)$$

where g indicates the g th energy group, G represents the overall number of energy groups, ϕ is the neutron scalar flux, D is the diffusion coefficient, $\Sigma_{s,g \rightarrow g'}$ is the scattering from g to g' , $\Sigma_{t,g}$ is the group-wise total cross section and S represents the external source of neutrons.

Once the flux is known, the Tritium Breeding Ratio can be computed too with nemoFoam, offering the possibility of optimizing also an integral quantity of ARC, with respect to more “local” quantities like

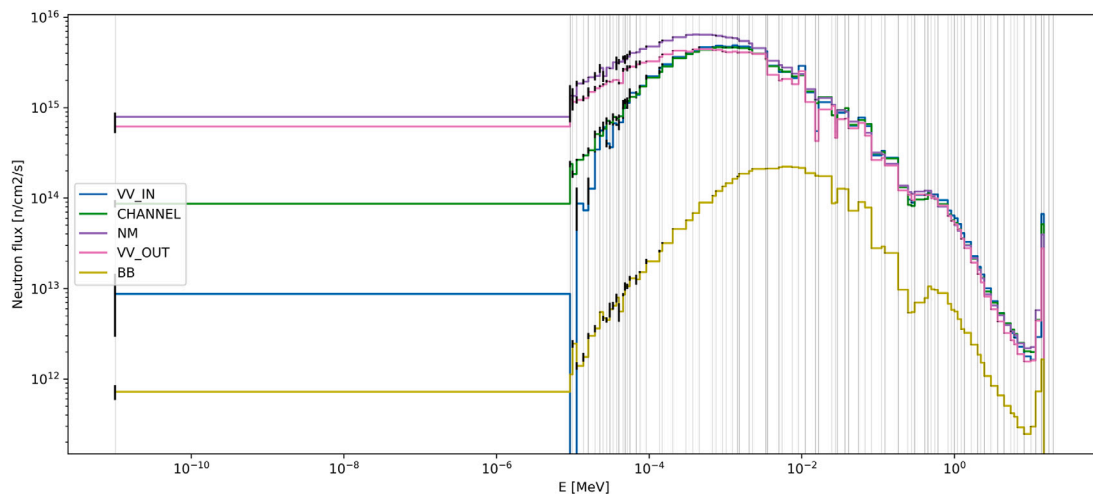


Fig. 3. Neutron energy spectrum, computed with Serpent, in the main regions of the ARC reactor with the 86 multi-group structure employed for the GA optimization.

the flux and the power deposition. Additionally, nemoFoam is also provided with a module for the solution of the multigroup adjoint diffusion equations,

$$\begin{aligned}
 & -\nabla \cdot (D_g(\vec{r})\nabla\phi_g^\dagger(\vec{r})) + \Sigma_{t,g}(\vec{r})\phi_g^\dagger(\vec{r}) \\
 & = \sum_{g'=1}^G \Sigma_{s,g\rightarrow g'}(\vec{r})\phi_{g'}^\dagger(\vec{r}) + S_g^\dagger(\vec{r}) \quad \forall g = 1, \dots, G,
 \end{aligned} \quad (2)$$

where ϕ^\dagger is the adjoint flux and S^\dagger is the adjoint source. In systems driven by external sources, like fusion reactors, ϕ^\dagger is interpreted as the importance of a single neutron in contributing to the score of a certain detector [21]. Since the detector response of interest is arbitrary, the adjoint source should be determined according to the context. In this case, the most relevant neutron-induced reaction is the production of tritium, i.e. $S^\dagger = \Sigma_{n\rightarrow^3H}$. In the case of ARC, computing the adjoint flux with respect to the TBR can provide a map of the regions contributing more to the final value of TBR. The latter module is exploited to optimize the convergence of the GA in this work, analogously to what was done in [12].

2.2. The genetic algorithm

GAs are defined as meta-heuristic algorithms, in the sense that, compared to rigorous optimization algorithm, they may not lead to the best solution, but to a solution close to the optimum one with a limited exploration of the solution space. In GAs, solutions are ranked according to a fitness function (FF), which mimics the principles of natural selection. In this way, the best solutions (individuals) in terms of FF have a higher probability of being mixed with other individuals, according to the so called crossover mechanism, and to breed the next generation, and so on, until a user-defined termination criterion is reached, allowing for desirable features to be inherited. The evolution efficiency strongly depends on the balance between selection pressure, which is related to the advantage given to better individuals over the others, and the mutation rate, which helps maintaining a certain variety in the population. In other words, the selection pressure drives the evolution of the population, whereas the mutation rate drives the exploration of the solution pool, making it possible to consider some new features in the process that could represent a breakthrough in the evolution.

Table 2 shows the values of the parameters selected to run the GA and determine how the new generations are composed (known as *hyperparameters* in the field of machine learning). Their meaning is described in the following:

- f_E is the fraction of new individuals given by the mechanism of *elitist selection*. These are the best individuals that are copied as they are into the new population, so preserving the best features of the population. Too high values of f_E , however, can be deleterious in terms of genetic diversity;
- f_{XO} is the fraction of individuals coming from crossing over, the process that combines two individuals into a new one of the new generation;
- the remaining slots are occupied by individuals randomly sampled from the previous generation, alleviating the selection pressure and helping preserve the diversity in the gene pool;
- all the individuals of the new population have a probability p_M of undergoing mutation, which is the random modification of one of their alleles. An allele represents the way in which a gene manifests itself.

In principle, the hyperparameters should be optimized too, in order to improve the convergence of the GA. However, this optimization is out of the scope of this work, and hyperparameters analogous to the ones employed in [12] are selected. Each individual is randomly assigned to one of the N_T tournaments (in this case, $N_T = 10$, see Table 2), without having mandatorily the same number of individuals in each tournament. Then, for each tournament, the individuals are ranked according to the value of their fitness function and assigned to a weight w_i which depends on their ranking r_i and on a parameter p (in this case, $p = 0.2$, see Table 2). Finally, the mating pool will be populated by a number of copies of the individual i proportional to its weight w_i :

$$w_i = p \times (1 - p)^{r_i - 1}. \quad (3)$$

In this way, the best individuals have a larger probability of generating offspring in the next generation still maintaining the genetic diversity, since the competition within a tournament is not necessarily between the best individuals. Thus, it is possible that a non-negligible amount of copies of non-fittest individuals will populate the mating pool.

As already mentioned in Section 1, the GA used in this work eliminates the limitation of having individuals with a fixed number of groups. However, this extension dramatically increases the dimension of the sample space. When the number of energy groups in the coarser grid is fixed, the dimension of the solution space is $\binom{n}{k}$, with n equal to the number of energy group boundaries in the fine grid and k equal to the number of energy boundaries in the coarser grid, while without such limitation the size increases up to 2^n . As a consequence, when the number of groups in the coarser grid is much smaller than n , the sample space cardinality increases by orders of magnitude. Moreover,

Table 2

GA parameters.

Individuals per generation	80
Selection method	Tournaments
Tournaments number	10
Tournament p	0.2
Elitism fraction	10%
XO fraction	70%
Mutation probability	20%

the convergence is made more complicated by the presence of the additional random variable dedicated to the number of energy groups.

In GAs, the features of each individuals are coded in sets of genes, each one representing a feature, that assume discrete values named alleles. Then, many genes constitute a chromosome, that is the way in which each individual is represented. The selection of an appropriate way to set up the chromosomes is fundamental in GAs, since it affects how the solution space is represented and sampled, influencing the evolution of the population. In this work, the representation denoted as γ in [12] (see Fig. 4), is chosen, since it was previously shown to provide the best results in terms of convergence. The γ representation is based on the concept of polyploid chromosomes, with the karyotype (i.e., the way in which the chromosome manifests itself) of an individual represented by two chromosomes, γ_0 and γ_1 :

- γ_0 is a single-gene chromosome, which indicates the number of energy groups of the individual;
- γ_1 is a polyploid chromosome composed of P arrays, each one with $H - 1$ genes (where H is number of groups in the fine grid), one for each boundary of the fine grid. The numerical values of the alleles can vary from 0 to M , where M is an hyperparameter arbitrarily selected by the user, similarly to P . γ_1 expresses as a vector of length $H - 1$ given by the sum of the P vectors and where the boundaries used for the individual evaluation (i.e., for the nemoFoam simulation) are the ones with the G highest scores. Thus, the values of γ_1 do not have a real physical meaning, but, once summed up, provide a priority order of the possible energy boundaries, as explained in Fig. 4.

In this representation, the crossover process acts on γ_1 by preserving all the P chromosomes in the offsprings, except for a random one which is substituted by a new array generated starting from the expression of the partner. In this way, the evolution is a less disruptive process, with the best boundaries slowly accumulating alleles with the highest value (i.e., equal to M), at the cost of a slower evolution convergence [12]. γ_1 does not directly affect the number of groups of the individual, which is governed by γ_0 : having only one gene, crossover on γ_0 results in the new individual inheriting the γ_0 chromosome of either one of its parents. It is worth mentioning that mutation can occur on both γ_0 and γ_1 .

The final element for the setup of the GA is the definition of the fitness function. Considering that some of the most important quantities that can be evaluated in the framework of neutronic studies for nuclear fusion breeding blankets are the TBR, the neutron power deposition, and the neutron flux distribution, all of them are taken into account in the FF definition. In particular, a single FF can be obtained coupling three different specialized FFs. The fitness function for the TBR is:

$$f_{TBR}^I = \left| TBR^I - TBR^{MC} \right|, \quad (4)$$

where TBR^I corresponds to the tritium breeding ratio of an individual I computed with nemoFoam and TBR^{MC} to the reference TBR computed with Serpent. The neutron power deposition FF is defined as the relative error between the powers computed with nemoFoam and Serpent:

$$f_P^I = \sqrt{\sum_{g=1}^G \sum_{j=1}^J \left(\frac{P_{g,j}^I - P_{g,j}^{MC}}{P_{g,j}^{MC}} \right)^2}, \quad (5)$$

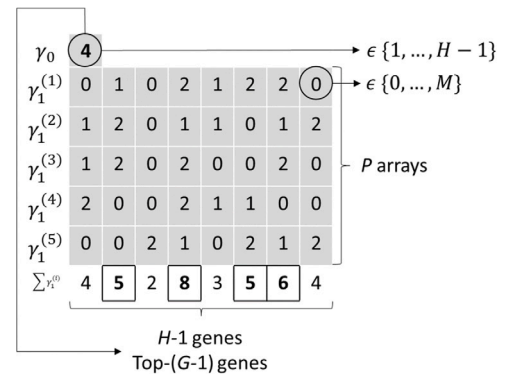


Fig. 4. γ chromosome representation, with $P = 5$ and $M=2$. Here, the individual is composed by the 2nd, 4th, 6th and 7th boundaries, since they are the ones with the highest scores.

where j indicates the spatial region (Fig. 2) and g the energy group. In a similar way, the FF for the neutron flux is defined as:

$$f_{\phi}^I = \sqrt{\sum_{g=1}^G \sum_{j=1}^J w_{g,j} \left(\frac{\phi_{g,j}^I - \phi_{g,j}^{MC}}{\phi_{g,j}^{MC}} \right)^2}, \quad (6)$$

where the set of weights $w_{g,j}$ is normalized, i.e.

$$\sum_{g=1}^G \sum_{j=1}^J w_{g,j} = 1, \quad (7)$$

and are constructed exploiting the concept of *total importance*, defined as $\psi = \phi^\dagger \phi / v = \phi^\dagger n$, where the adjoint ϕ^\dagger is calculated in nemoFoam with respect to a detector evaluating the tritium production. The weights in Eq. (7) have been computed starting from the values of total importance computed in nemoFoam and normalizing them to 1. It has not been possible to evaluate a reference adjoint using Serpent (as for the other quantities), because the code does not currently allow to score the adjoint flux distribution. Exploiting the total importance, the regions in the phase space which contribute the most to the TBR are associated to larger weights in Eq. (6) and, as a consequence, have a larger impact on f_{ϕ}^I . Fig. 5 displays the values of total importance, adjoint flux (computed with nemoFoam) and neutron flux in the blanket of ARC, showing that even if the flux is quite small below 10^{-3} MeV, the total importance is high, mainly because at lower energies the adjoint flux is higher (consistently with the tritium production cross section of ${}^6\text{Li}$). Of course, in order to perform the operations in Eq. (5) and (6), the results of the MC simulations evaluated on the finer grid have to be condensed on the same energy grid of the individual I . The condensation is performed thanks to the COREutils package.

The current FF formulation does not include a specific penalty term to account for excessive computational time, although such technique can be included into future works, instead of a hard limitation to the number of groups. This would promote the solutions with less groups without forcibly excluding other solutions.

Then, these fitness functions can be combined together, in order to find an individual which tries to optimize all the three fitness functions at the same time. It is important to underline that the FFs span very different scales and that the range and actual variability of the FFs are not known a priori. If these aspects are neglected, one FF can become dominant over the other or, in the opposite case, negligible. Thus, different transformations are implemented in the GA in order to deal with this issue:

- Linear scaling, where all the values of the FFs are continuously mapped (in a linear way) on an arbitrary interval $[a, b]$:

$$h : f \mapsto a + [b - a] \cdot \frac{f - \min_{I \in E}(f^I)}{\max_{I \in E}(f^I) - \min_{I \in E}(f^I)}, \quad (8)$$

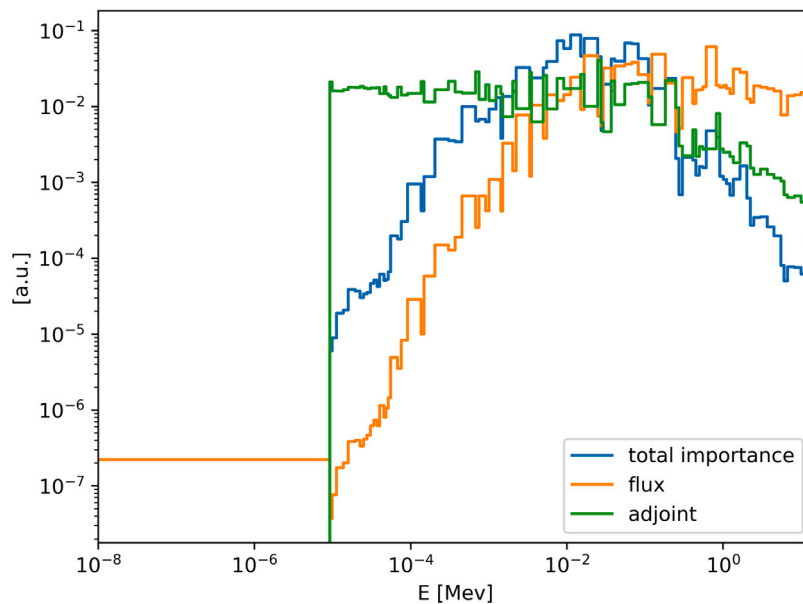


Fig. 5. Distributions of the normalized direct flux, adjoint flux and total importance in the breeding blanket of ARC scored on the reference 86 groups energy grid employed for the GA optimization.

where E is the explored sample space.

- Ranking, where each value f of the FF is associated to the percentile it belongs to in the set of FF values ever observed:

$$h : f \mapsto 1 + 99 \cdot \frac{\text{card}\{I \in E : f^I > f\}}{\text{card}(E)}. \quad (9)$$

In this work, only the linear scaling is tested, but both scalings are available in the GA used. Once the transformation is applied by the GA, making the FFs comparable within each other, they can be combined performing their product (i.e., the geometric average):

$$FF^I = h(f_{\text{TBR}}^I) \cdot h(f_{\text{P}}^I) \cdot h(f_{\phi}^I). \quad (10)$$

The selection of the geometric average is motivated by the fact that it penalizes largely unbalanced solutions.

Finally, the simulation ends as soon as the GA reaches one of the following user-defined termination criteria: an individual with a FF smaller than a threshold one is found, a maximum number of generations is reached, a maximum number of inefficient generations is reached (i.e., generations without new optimal individuals with respect to the previous ones), or the simulation exceeds the maximum computational time.

3. Results

Each individual is generated ensuring that it contains at least the three following energy boundaries: 1e-11 MeV, 14 MeV and 20 MeV. The 1e-11 and 20 MeV cuts are required to limit the energy range of the simulation (they correspond to the maximum neutron energy range that can be simulated in Serpent), while the 14 MeV cut is chosen for the definition of the incoming neutron current boundary condition between the plasma and the first wall (which is the method used in nemoFoam to model the neutrons generated in the plasma), since it is chosen to keep fixed the values of the incoming currents for all the individuals. Here, the choice of the cut at 14 MeV is chosen since this is the energy at which neutrons are emitted by D-T fusion reactions. Then, the GA is configured so that it can generate individuals with an additional number of energy boundaries between 2 and 7. The lower limit is chosen since it can be expected that solutions with a too small number of energy groups cannot be optimal; the upper limit, on the other hand, is selected to limit the computational time required by

nemoFoam to solve the neutronic problem, also in view of coupled multiphysics simulations.

Considering that the blanket is the largest energy source in ARC, a first set of GA runs is performed focusing the optimization on the neutron power deposition, according to Eq. (5). In this first stage, the research of the optimal grid is limited to the power deposition, since this is the quantity that really matters when neutronics and thermal hydraulics are coupled. Since the random number generator of the GA is initialized according to an input string provided by the user, three different GA runs were performed using the strings “Henry”, “Fermi” and “Boltzmann”, each with the parameters listed in Table 2. The calculations run with the three “populations” are useful to assess the impact of the randomness of the initial set of individuals on the optimal individual/energy grid found by the GA. Fig. 6 shows the convergence trends of the three populations in terms of the average value of the fitness function of each generation, clearly exemplifying the impact of the initial population on the behavior of the GA. In fact, it can be seen that the convergence of the Boltzmann population is slightly slower, while the trends of Henry and Fermi are quite similar up to the 13th generation, then Fermi reaches a plateau without any appreciable improvement of the fitness function. All the three simulations are terminated due to the fact that they have reached the maximum number of inefficient generations, set to 20.

Another aspect that deserves to be analyzed is the evolution of the individuals in terms of the number of groups. In principle, the initial population should be constituted by individuals with a distribution of the number of groups as uniform as possible (see the graph on the top of Fig. 7), in order to avoid any possible bias in the first phases of the evolution that can influence the following generations. Then, the genetic mechanism try to favor the individuals with a number of groups which manages to better optimize the fitness function. If we observe the distribution of the number of groups of all the individuals simulated by the GA (bottom graph in Fig. 7), we can see that the individuals with the lowest number of groups are predominant. This is particularly favorable, since these individuals are not only the optimal ones in terms of the fitness function (i.e., they are the best one in optimizing the neutron power deposition in the tank of ARC with respect to Serpent), but also in terms of computational time, since they reduce as much as possible the number of energy groups, which corresponds to the number of equations to be solved.

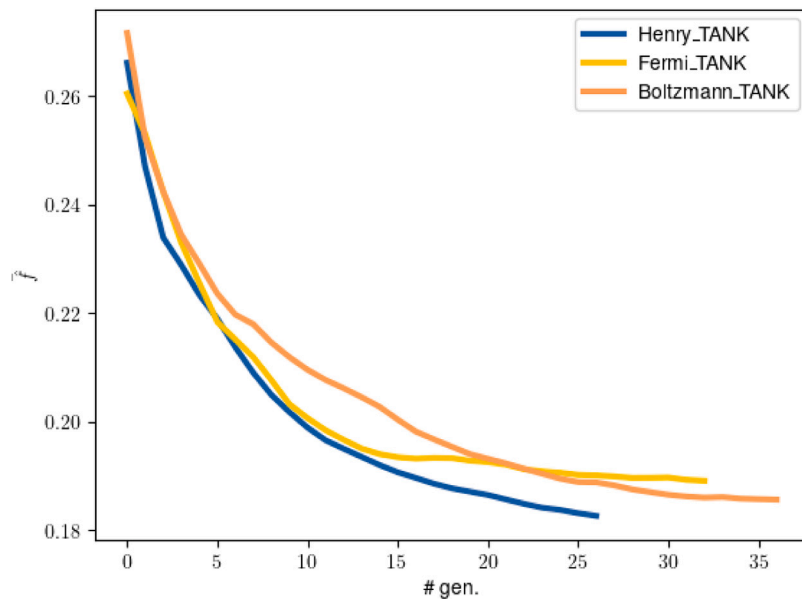


Fig. 6. Convergence trend of the average fitness function of the three populations for each generation during the evolution process, when only the power deposition in the breeding blanket of ARC is optimized.




Fig. 7. Distribution of individuals per number of groups in the initial generation and considering all the individuals.

The number of groups is not the only important information that can be gained from the population evolution, it is also useful to know

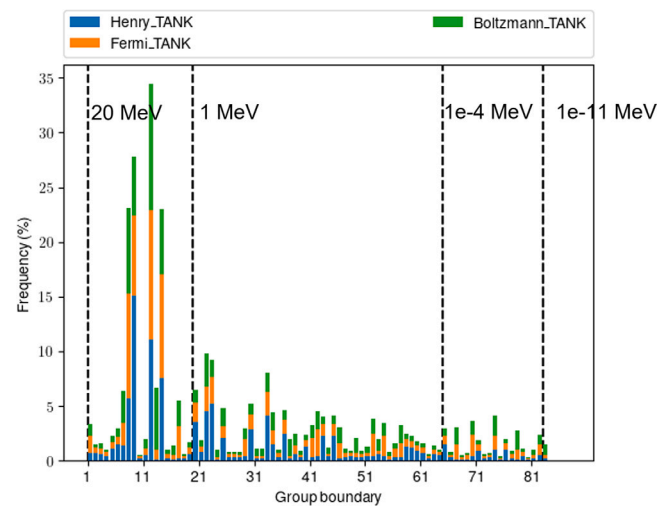


Fig. 8. Distributions of the energy boundaries considering all the individuals simulated using the three populations, when only the power deposition in the blanket of ARC is optimized.

where the energy cuts are localized, to see if there are particular regions of the energy domain which are favored by the GA (i.e., which tend to optimize the fitness function). In Fig. 8 it is possible to observe that all the three populations (independently from each other) yield a large amount of individuals with energy boundaries in the fast range, between 20 MeV and 1 MeV. This is a first indication of the fact that to optimize the neutron power deposition in the tank of ARC using nemoFoam, it is probably better to have more groups at high energies, while for lower groups having a coarser discretization can work as well, consistently with the physics of the system.

Finally, each population leads to an optimal individual, which is shown in Fig. 9. The optimal energy boundaries selected by the GA starting from the three populations confirm the previously made considerations: all the optimal energy grids have four groups (i.e., two additional groups with respect to the initial fixed two groups) and the cuts are localized in the fast energy range between 1 MeV and 20 MeV.

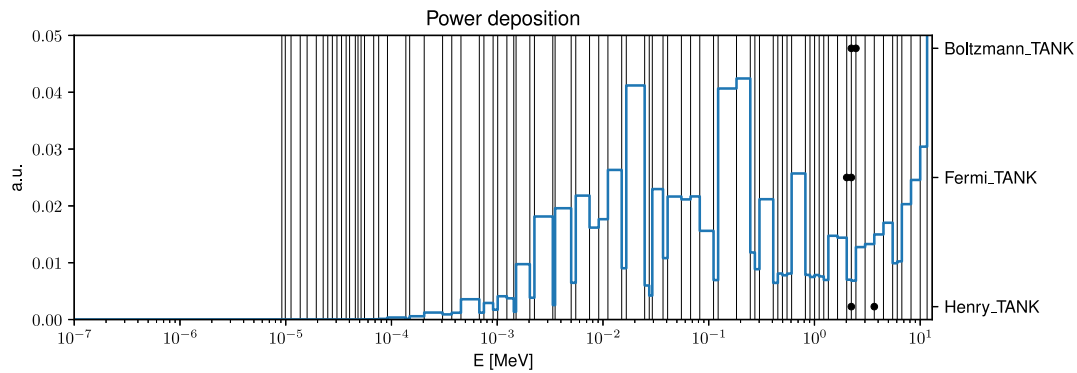


Fig. 9. Normalized group-wise reference neutron power deposition in the breeding blanket of ARC and optimal group boundaries (black dots) selected from the reference fine-group structure (dashed black lines) for the three different initial populations. The dots at 10^{-11} MeV, 14 MeV and 20 MeV are not shown because they are present in each individual by default.

Moreover, the three optimal individuals are extremely similar, suggesting that the final result is weakly influenced by the initial population. According to the settings used to set up the GA, the minimum total number of additional energy groups is 2; however, observing that the optimal boundaries in Fig. 9 are extremely near in the energy domain, this suggests that probably one additional group instead of two would be sufficient to obtain reasonable results.

Even if the breeding blanket of ARC is its biggest component, also the layers composing the vacuum vessel are important in the design of the reactor, also considering that one of these layers is the cooling channel, which obviously has a key role in the framework of the thermal-hydraulic analysis of ARC. For this reason, a second set of GA simulations is run with the purpose of optimizing the neutron power deposition in all components of ARC at the same time, according to Eq. (5). The setup of the simulation is the same, and also in this case the largest amount of individuals has two additional groups indicated by the GA. The main difference is the distribution of the energy boundaries, as shown in Fig. 10. In this case, only one population (Henry) is simulated, since the previous runs have shown that the final results should be weakly influenced by the initial population. The distribution of the group boundaries is quite different from the previous ones: in this case, most of the energy groups are located in the intermediate and slow range. This suggests that, probably, the optimization of the energy grid for the power deposition in the blanket is in competition with the optimization of the power deposition in the other components: the optimal solution, therefore, would just be a trade-off between the different FF components. A possible explanation of this behavior can lie in the fact that the blanket is the farthest component from the neutron source, meaning that the fast part of the spectrum needs to be better described to ensure the accuracy of the results in blanket. In fact, fast neutrons are the ones actually reaching the more peripheral regions of the reactor. This interpretation is further justified by Fig. 3, where it can be noticed that the spectrum in the breeding blanket is noticeably harder than in the other components.

Fig. 11 displays the optimal individual and the normalized power deposited in each group of the reference energy grid for each component of ARC, with the two cuts selected by the GA located in the intermediate energy range.

Finally, considering that one could be interested not only in the power deposition, but also to other nuclear quantities like the neutron flux and the TBR, a third set of simulations is run trying to optimize all these three quantities, exploiting Eq. (10). In particular, two families can be identified: one where the neutron flux and the neutron power deposition are optimized for all the components and a second one where they are optimized only for the breeding blanket. In both cases, the TBR is optimized too. The choice of running two different families of simulations is led by the fact that the previous results have suggested

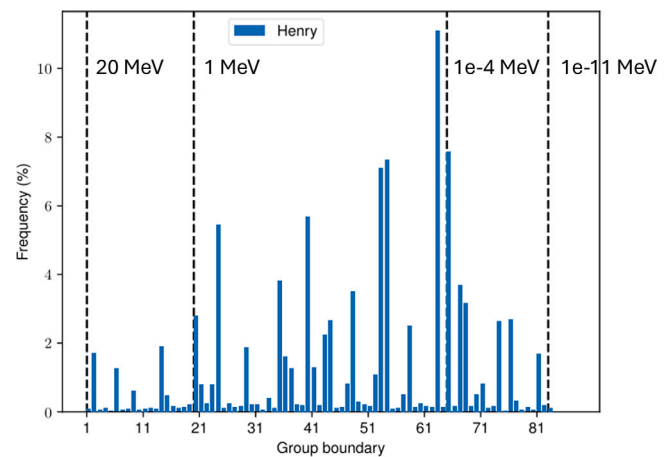


Fig. 10. Distributions of the energy boundaries considering all the individuals simulated, when the power deposition in all the components of ARC is optimized.

that the breeding blanket and the layers of the vacuum vessel are in competition.

In this case, the distribution of the energy boundaries is somehow similar to the one found when the power deposition is optimized in all the components: higher frequencies are observed especially in the intermediate energy region and, to a less extent, also in the low energy one (Fig. 12). Here, results identified by the “*_TANK” label are referred to the case where only the flux and the power deposition in the breeding blanket are optimized.

Finally, Fig. 13 confirms the fact that it does not seem possible to find an energy grid that is optimal for both the breeding blanket tank of ARC and all the other components. Observing the optimal individual labeled with “Henry_TANK” and “Fermi_TANK” and comparing them with the results in Fig. 9, it is also possible to deduce that the optimization of the power deposition is in competition with the optimization of the neutron flux and TBR, with the former which requires more groups in the lower energy range. The results labeled with “Boltzmann” include the groups appearing in the best individual of the “Fermi” and “Henry” cases, suggesting that the final results of the GA are not too influenced by the initial population. A possible explanation of the fact that the GA finds optimal individuals with few groups is related to the assumptions of the diffusion theory, leading the solution with many groups to be inaccurate. In fact, diffusion is accurate if the collisionality in each energy group is high, which could not be the case with a fine multi-group structure. Thus, solutions with a small number of groups might be favored. This condition is further supported by the presence

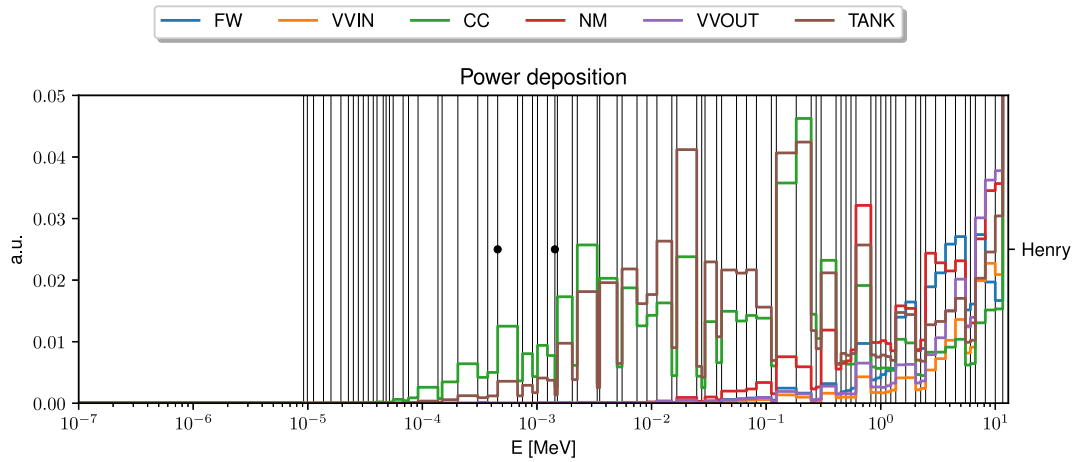


Fig. 11. Group-wise reference neutron power deposition normalized to 1 in the breeding blanket of ARC and optimal group boundaries (black dots) selected from the reference fine-group structure (dashed black lines), when the power deposition in all the components is optimized.

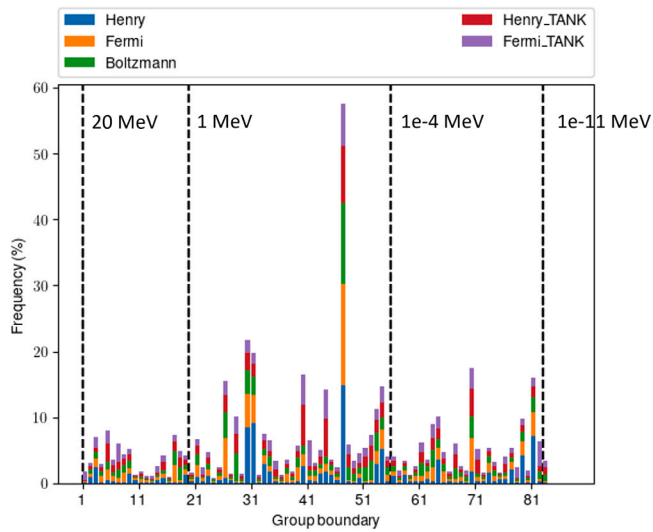


Fig. 12. Distributions of the energy boundaries considering all the individuals simulated using the three populations.

of relatively light isotopes in ARC (e.g., those composing the FLiBe) whose cross sections exhibit a limited number of resonances. Thus, few energy groups could be sufficient to reproduce the continuous energy cross sections.

The results obtained by the three sets of simulations in terms of relative error between the reference power values (i.e., computed with Serpent) and the powers computed by nemoFoam using the optimal individuals, thanks to the following expression:

$$\epsilon_{j,\%} = \frac{|P_j^{OPT} - P_j^{MC}|}{P_j^{MC}} \cdot 100, \quad (11)$$

where j represents the j th ARC component, are shown in Fig. 14. In particular, “Henry-Pow-Tank” represents the results optimizing only the power deposition in the breeding blanket, “Henry-Pow” the results when the power in all the components is optimized and “Henry-All” is the case when the power, the neutron flux and the TBR are optimized in all the components. As expected, the best result in the breeding blanket (TANK) is obtained in the “Henry-Pow-Tank” case, which provides the best result in the cooling channel (CC) too: this can be explained considering that the cross sections of the CC and the TANK are similar, since they made by the same material (FLiBe), thus optimizing the

grid for the TANK may provide good results for the CC too. However, the “Henry-Pow-Tank” case provides quite bad results in the other components, so if one wants a grid able to maximize as much as possible the agreement with the reference power in all the components, the optimal individual found by the “Henry-Pow” case is the best. On average, “Henry-All” provides the worst result since it tries to optimize at the same time the TBR, the neutron flux and the power deposition, and it is not specifically focused to the power deposition as for the other two cases. Finally, concerning the values of TBR provided by the three best individuals, it can be noticed that all of them result in an overestimation of $\sim 10\%$ with respect to the reference value, regardless of the fitness function employed. This is probably due to the fact that the tritium production cross section of ${}^6\text{Li}$ has a simple shape, thus the value of TBR computed by nemoFoam is quite independent of the position of the energy boundaries provided by the GA.

4. Conclusions and perspectives

In this work, a Genetic Algorithm was employed to find optimal energy grids to be used in the deterministic neutronic model of the ARC fusion reactor. The starting point is the evaluation of nuclear properties and reference quantities (i.e., the neutron flux, the power deposition and the TBR) in ARC on a fine energy grid (86 groups) using the Serpent Monte Carlo, which here is considered as a reference. Then, the GA was applied to the nemoFoam model of ARC, with the aim of reproducing as well as possible the results provided by Serpent. The possibility of varying not only the positions of the energy cuts, but also the number of energy groups, was exploited here, in order to fully explore the solution space. In order to allow the energy discretization to have a variable number of groups, without distorting excessively the probability space over which the GA works, a chromosome representation based on the concept of polyploidy was used.

Different fitness functions were considered. At first, a fitness function with the aim of optimizing only the neutron power deposition in the breeding blanket of ARC was tested. Then, the focus was moved to the optimization of the power deposition in all the components of ARC and, finally, the optimization of the neutron flux, the power deposition and the TBR of ARC. The results obtained with the three fitness functions show that the optimization of the different components and of the different quantities. In fact, each fitness function provides an optimal individual with energy groups located in different regions of the energy range. This is an evidence of the fact that the fitness function has to be carefully defined according to the specific application of interest. Despite the above mentioned differences, a feature which is common for all the fitness functions, is that the optimal individuals

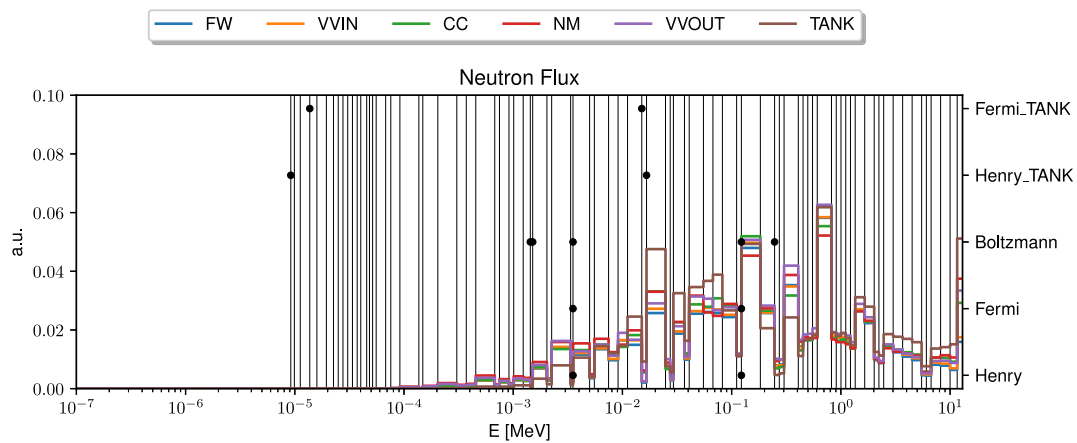


Fig. 13. Normalized group-wise reference neutron flux in ARC and optimal group boundaries (black dots) selected from the reference fine-group structure (dashed black lines), when the power deposition, the neutron flux and the TBR are optimized.

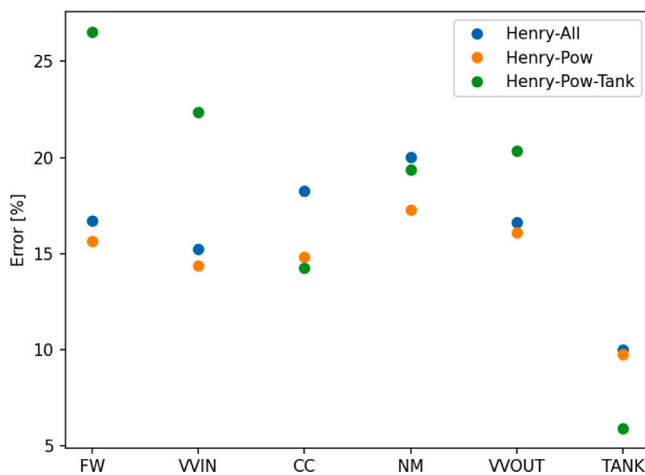


Fig. 14. Relative error between the reference neutron power deposition and the power deposition computed with the optimal individuals in all the components of ARC.

of ARC tend to have few energy groups, which is an advantageous feature since this minimize the computational cost of the nemoFoam neutronic simulation too. Moreover, the results of the GA are also weakly sensitive to the initial population, suggesting that the results have reached a reasonable level of convergence.

It is worth mentioning that, in principle, the optimal individuals strongly depend on the choice of the initial fine energy grid used for the evaluation of the nuclear properties and nuclear quantities of interest. Thus, in the future, it would be useful to check which is the sensitivity of the GA results to the choice of the fine grid. Moreover, considering the non-negligible contribution of photon power deposition in ARC (see [14]), it would be interesting to include in the GA optimization the gamma heating too, in order to see if it can lead to different energy grids with respect to the ones find in this work. Finally, in the context of multiphysics simulations, the GA should be applied using, as reference, nuclear properties and quantities of interest evaluated at different temperatures, in order to find an energy grid which is optimal in the whole temperature range that can be of interest in the framework of a coupled neutronic-thermalhydraulic simulation.

Funding

This research was partially supported by the Italian Ministry of University and Research (MUR) under the PNRRNGEU framework.

Specifically, the work of Caravello, M. is part of the project funded by DM 352/2022.

CRediT authorship contribution statement

Alex Aimetta: Writing – original draft, Software, Methodology, Conceptualization. **Nicolò Abrate:** Writing – review & editing, Software, Methodology, Conceptualization. **Marco Caravello:** Software, Methodology. **Sandra Dulla:** Writing – review & editing, Supervision. **Antonio Froio:** Writing – review & editing, Supervision. **Matia Massone:** Writing – review & editing, Software, Methodology, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Computational resources were provided by HPC@PoliTO and CLOUD@PoliTo, a project of Academic Computing within the Department of Control and Computer Engineering at Politecnico di Torino.

Data availability

Data will be made available on request.

References

- [1] B. Sorbom, J. Ball, T. Palmer, F. Mangiarotti, J. Sierchio, P. Bonoli, C. Kasten, D. Sutherland, H. Barnard, C. Haakonsen, J. Goh, C. Sung, D. Whyte, ARC: A compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets, *Fusion Eng. Des.* 100 (2015) 378–405, <http://dx.doi.org/10.1016/j.fusengdes.2015.07.008>, URL <https://www.sciencedirect.com/science/article/pii/S0920379615302337>.
- [2] A. Kuang, N. Cao, A. Creely, C. Dennett, J. Hecla, B. LaBombard, R. Tinguely, E. Tolman, H. Hoffman, M. Major, J. Ruiz Ruiz, D. Brunner, P. Grover, C. Laughman, B. Sorbom, D. Whyte, Conceptual design study for heat exhaust management in the ARC fusion pilot plant, *Fusion Eng. Des.* 137 (2018) 221–242, <http://dx.doi.org/10.1016/j.fusengdes.2018.09.007>, URL <https://www.sciencedirect.com/science/article/pii/S0920379618306185>.
- [3] A. Till, N. Gibson, T. Saller, Optimizing Group Structures using Hierarchical Division, in: *Proceedings of the ANS International Conference PHYSOR 2022*, Pittsburgh, PA, 2022, pp. 687–696.
- [4] V. Nair, T. Saller, A. Till, N. Gibson, Group Structure Selection with Random Forests, in: *Proceedings of the ANS International Conference PHYSOR 2022*, Pittsburgh, PA, 2022, pp. 1186–1195.

- [5] T.G. Saller, V. Nair, A. Till, N. Gibson, Using a Random Forest Model to Choose Optimized Group Structures, *Nucl. Sci. Eng.* 197 (8) (2023) 2117–2135, <http://dx.doi.org/10.1080/00295639.2022.2133940>.
- [6] V.A. Di Nora, E. Fridman, E. Nikitin, Y. Bilodid, K. Mikityuk, Optimization of multi-group energy structures for diffusion analyses of sodium-cooled fast reactors assisted by simulated annealing—Part I: Methodology demonstration, *Ann. Nucl. Energy* 155 (2021) 108183, <http://dx.doi.org/10.1016/j.anucene.2021.108183>.
- [7] V.A. Di Nora, E. Fridman, E. Nikitin, Y. Bilodid, K. Mikityuk, Optimization of multi-group energy structures for diffusion analyses of sodium-cooled fast reactors assisted by simulated annealing – Part II: Methodology application, *Ann. Nucl. Energy* 163 (2021) 108541, <http://dx.doi.org/10.1016/j.anucene.2021.108541>.
- [8] O. Fasina, T. Saller, N. Haven, Particle Swarm Optimization for Group Structure Optimization for Radiotherapy Shielding, in: *Proceedings of the ANS International Conference PHYSOR 2022, Pittsburgh, PA, 2022*, pp. 1340–1350.
- [9] M. Massone, F. Gabrielli, A. Rineiski, A genetic algorithm for multigroup energy structure search, *Ann. Nucl. Energy* 105 (C) (2017) 369–387, <http://dx.doi.org/10.1016/j.anucene.2017.03.022>.
- [10] M. Massone, F. Gabrielli, A. Rineiski, SIMMER extension for multigroup energy structure search using genetic algorithm with different fitness functions, *Nucl. Eng. Technol.* 49 (6) (2017) 1250–1258, <http://dx.doi.org/10.1016/j.net.2017.07.012>.
- [11] M. Massone, N. Abrate, G.F. Nallo, S. Dulla, P. Ravetto, D. Valerio, Genetic algorithm-based optimisation of the few-group structure for lead fast reactors analysis, in: *International Conference PHYSOR 2022, Pittsburgh (Pennsylvania), U.S.A., 2022*.
- [12] N. Abrate, A. Aimetta, M. Massone, S. Dulla, P. Ravetto, A genetic-driven optimization of the energy grid structure for nodal full-core calculations in lead-cooled fast reactors, *Nucl. Sci. Eng.* (2025) 1–26, <http://dx.doi.org/10.1080/00295639.2024.2446130>.
- [13] M. Massone, N. Abrate, G.F. Nallo, D. Valerio, S. Dulla, P. Ravetto, Code-to-code SIMMER/FRENETIC comparison for the neutronic simulation of lead-cooled fast reactors, *Ann. Nucl. Energy* 174 (2022) 109124.
- [14] M. Caravello, A. Aimetta, N. Abrate, S. Dulla, A. Froio, An OpenFOAM solver for multiphysics modeling of fusion reactor design: The nemoFoam code, *Nucl. Mater. Energy* 40 (2024) 101693, <http://dx.doi.org/10.1016/j.nme.2024.101693>, URL <https://www.sciencedirect.com/science/article/pii/S2352179124001169>.
- [15] J. Leppänen, V. Valtavirta, A. Rintala, R. Tuominen, Status of serpent Monte Carlo code in 2024, *EPJ Nucl. Sci. Technol.* 11 (2025) 3, <http://dx.doi.org/10.1051/epjn/2024031>.
- [16] A. Aimetta, N. Abrate, S. Dulla, A. Froio, Neutronic analysis of the fusion reactor ARC: Monte Carlo simulations with the serpent code, *Fusion Sci. Technol.* 78 (4) (2022) 275–290.
- [17] W. Jakob, J. Rhineland, D. Moldovan, et al., *pybind11 – seamless operability between c++11 and python*, 2017.
- [18] N. Abrate, et al., *COREutils – a python package for nuclear reactor modelling*, 2020.
- [19] D. Brown, M. Chadwick, R. Capote, A. Kahler, A. Trkov, M. Herman, A. Sonzogni, Y. Danon, A. Carlson, M. Dunn, D. Smith, G. Hale, G. Arbanas, R. Arcilla, C. Bates, B. Beck, B. Becker, F. Brown, R. Casperson, J. Conlin, D. Cullen, M.-A. Descalle, R. Firestone, T. Gaines, K. Guber, A. Hawari, J. Holmes, T. Johnson, T. Kawano, B. Kiedrowski, A. Koning, S. Kopecky, L. Leal, J. Lestone, C. Lubitz, J. Márquez Damián, C. Mattoon, E. McCutchan, S. Mughabghab, P. Navratil, D. Neudecker, G. Nobre, G. Noguere, M. Paris, M. Pigni, A. Plompen, B. Pritychenko, V. Pronyaev, D. Roubtsov, D. Rochman, P. Romano, P. Schillebeeckx, S. Simakov, M. Sin, I. Sirakov, B. Sleaford, V. Sobes, E. Soukhovitskii, I. Stetcu, P. Talou, I. Thompson, S. van der Marck, L. Welsch-Sherill, D. Wiarda, M. White, J. Wormald, R. Wright, M. Zerck, G. Žerovnik, Y. Zhu, ENDF/B-VIII.0: the 8th major release of the nuclear reaction data library with CIELO-project cross sections, new standards and thermal scattering data, *Nucl. Data Sheets* 148 (2018) 1–142, <http://dx.doi.org/10.1016/j.nds.2018.02.001>, Special Issue on Nuclear Reaction Data, URL <https://www.sciencedirect.com/science/article/pii/S0090375218300206>.
- [20] E. Sartori, G. Panini, ZZ groupstructures, vitamin-J, XMAS, ECCO-33, ECCO2000 standard group structures, 1991.
- [21] G.I. Bell, S. Glasstone, in: V.N.R.C. New (Ed.), *Nuclear Reactor Theory*, New York, 1970.