

A Genetic-Driven Optimization of the Energy Grid Structure for Nodal Full-Core Calculations in Lead-Cooled Fast Reactors

Original

A Genetic-Driven Optimization of the Energy Grid Structure for Nodal Full-Core Calculations in Lead-Cooled Fast Reactors / Abrate, N., Aimetta, A., Massone, M., Dulla, S., Ravetto, P.. - In: NUCLEAR SCIENCE AND ENGINEERING. - ISSN 0029-5639. - ELETTRONICO. - (2025), pp. 1-26. [10.1080/00295639.2024.2446130]

Availability:

This version is available at: 11583/2998043 since: 2025-03-04T10:07:50Z

Publisher:

Taylor & Francis

Published

DOI:10.1080/00295639.2024.2446130

Terms of use:

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

Publisher copyright

Taylor and Francis postprint/Author's Accepted Manuscript

This is an Accepted Manuscript of an article published by Taylor & Francis in NUCLEAR SCIENCE AND ENGINEERING on 2025, available at <http://www.tandfonline.com/10.1080/00295639.2024.2446130>

(Article begins on next page)

A genetic-driven optimisation of the energy grid structure
for nodal full-core calculations in lead-cooled fast reactors

Nicolò Abrate,^{a,b} Alex Aimetta,^a Mattia Massone,^c Sandra Dulla,^{a,b} and Piero
Ravetto^{a,b}

^a*Politecnico di Torino, Dipartimento Energia, NEMO Group
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)*

*Email: nicolo.abrate@polito.it

Number of pages: 45
Number of tables: 2
Number of figures: 22

Abstract

This work presents a novel genetic algorithm for optimising the few-group energy grid structure used for full-core nodal calculations in Lead-cooled Fast Reactors. The optimisation is started considering a set of group constants computed on a reference 61-group structure, from which the genetic algorithms selects an optimal subset of groups.

Compared to existing works in the literature, the number of groups is not defined *a priori* but varies within a user-defined range, allowing a better exploration of the solution space. This feature requires to develop an adequate representation of the chromosomes used in the evolution process, which is examined with different definitions of the chromosomes.

The work also proposes a suitable combination of physics-driven fitness functions, related to the effective multiplication factor, the power density and the neutron flux. Different weights based on the adjoint flux are also studied for the flux fitness function, with the aim of improving the convergence of the evolution process.

All the studies are performed focusing on a 3D model of the ALFRED core design, which is modelled using the multi-group diffusion module of the FRENETIC multiphysics code. The results suggest that the energy grid can be profitably optimised using a representation with two chromosomes. The optimal solutions yielded by the genetic algorithm are justified on a physical basis, by looking at some relevant figures of merit.

Keywords — lead-cooled fast reactor, energy group structure optimisation, genetic algorithm, nodal full-core calculation

I. INTRODUCTION

The selection of the energy group structure for collapsing the group constants for full-core calculations is often based more on expert judgement than on a rigorous basis, due to the physical complexity of the neutron cross sections behavior. This open issue in the physics of nuclear reactors is extremely relevant for systems featuring a fast spectrum, like Lead-cooled Fast Reactors (LFR) [1]. As design and safety assessment stages usually require a large number of calculations, it is important to minimise the computational burden of each full-core analysis, while still preserving the underlying physics and providing accurate results. The compromise between computational performances and model fidelity is clearly related to both an optimal number of energy groups and their distribution along the energy axis.

The selection of an optimal group grid is a stiff, non-linear optimisation problem featured by a very large parameter space. The relevance of this research topic is witnessed by the significant number of works published on this subject in the past two years, which explore the application of different optimisation methods to this challenging reactor physics problem. In [2], the Hierarchical Division algorithm is proposed to add some optimal energy group to a starting group structure, with the objective of improving existing group structures. The performances of this approach are tested on a homogenised point reactor, showing a good error reduction but also highlighting the known limitations of greedy algorithms in presence of local minima [3, 4]. In [5, 6], Simulated Annealing is applied to generate a set of optimal group grids and, then, the Random Forest regression method is adopted to choose the best group structure for the various test problems, which are homogeneous, reflected spheres made of different isotopes. The same method has also been adopted by [7, 8] for optimising the group structure for the full-core simulation of a Sodium Fast Reactor (SFR). Another approach, aiming at optimising the group structure for radiation shielding calculations, consists in the Particle Swarm Optimisation, which belongs to the class of evolutionary algorithms [9]. This work shows how the adoption of an optimised grid outperforms the standard group structures available in the PARTISN code [10] in terms of both accuracy and computational time.

Another family of advanced search algorithms belonging to the broader class of evolutionary algorithms is represented by Genetic Algorithms (GAs). GAs allows to efficiently explore the solution space, determining nearly optimal energy structures, i.e., the ones minimising a suitable

fitness function (or loss function) while complying with the given constraints. Although it is difficult, or even impossible, to prove that GAs outperform other classes of meta-heuristics for the energy discretisation problem, the essential nature of the problem related with sequences of energy boundaries, with each boundary strongly influenced by the neighbouring ones, well resembles the sequences of *genes* appearing in chromosomes. In addition, GAs do not require that the fitness functions comply with specific assumptions (e.g. convexity) or have particular good-behaviour properties over the whole phase space, needed by deterministic approaches (e.g. gradient descent) and by greedy algorithms to work.

GAs have been already applied in this field [11, 12], showing good performances for optimising the energy structures for molten salt reactors and sodium fast reactors. A similar analysis has been recently carried out also for the specific case of LFRs [13], focusing on the Advanced Lead Fast Reactor European Demonstrator (ALFRED) design. In [13], the genetic algorithm coupled to the SIMMER code [12, 14] was employed to extract a set of possible optimal six-group structures out of a starting fine-grid consisting of 120 groups, considering as fitness function a combination of the effects on the calculated k_{eff} and neutron flux.

In this paper, we extend our previous work [13] by defining a new physics-based fitness function and employing a more flexible GA that can yield a set of optimal structures without fixing the number of groups in advance. This feature is a significant improvement as compared to other works in the literature (see, e.g., [12, 7]), where the number of groups are chosen *a priori*, thus arbitrarily constraining the solution space. Although this aspect makes the optimisation process more general, it also significantly broadens the dimension of the solution space, forcing an improvement of the solution algorithm.

In contrast to [13], where the SIMMER code was employed, the code adopted in this work to evaluate the performances of the few-group structures is the nodal diffusion module of the Fast REactor NEutronics/Thermal-hydraulICs (FRENETIC) multiphysics code [15]. The motivations for the choice of this code, which has been developed at Politecnico di Torino in the last decade for the accidental analysis of liquid-metal-cooled fast reactors, are mainly two: i) it runs faster than SIMMER, at the price of a slightly less accurate model: FRENETIC is based on a diffusion nodal method, while SIMMER employs a solver based on discrete-ordinates[15] and, ii) the code offers a simplified management of the multi-group libraries at different operational temperatures, in view

of a possible future introduction of the thermal feedback in the calculations.

The adoption of FRENETIC also enables to explore the limits of the implemented nodal diffusion method. It is well known that coarse mesh nodal diffusion methods provide a sufficiently accurate approximation to the neutron transport equation for optically thick systems, featured by a large value of secondaries per collision c . This condition is clearly related to the group-wise value of c : with a fine multi-group structure (e.g. 50-150 groups), the collisionality would be small, thus questioning the validity of the diffusion model. On the contrary, a minimum number of groups (say 5-6) is needed to provide an appropriate description of the physics of fast reactors. In this perspective, the determination of the optimal energy group structure would also give important indications on the limits of the nodal diffusion model in the frame of lead fast reactors analysis.

In general, the discrepancies between transport and diffusion models are only partially due to the choice of the group structure: also angular and spatial effects can be relevant. Since the coupling between energy, space and angle effects is significant, especially in heterogeneous systems, the optimisation of the group structure reduces the overall error between the two models. Because of possible compensations of these errors, it should be stressed that the application of the group structures proposed by the GA to core configurations different from the one used in the optimisation process may not yield optimised results. An intuitive strategy to reduce these possible over-fitting effects could be optimising the grid considering more configurations at the same time. Nevertheless, this aspect is not explored in the paper.

II. METHODOLOGY

In this section we outline the methodology conceived and followed for the genetic-driven optimisation, which is sketched in fig. 1. The whole optimisation process is started with a set of fine-group constants on a reference multi-group grid, from which the optimal grid structure is extracted. The reference group constants are generated with the Serpent Monte Carlo code [16], which is also used to produce a reference solution in a specific core configuration for driving the optimisation process. Assuming that the starting energy grid is fine enough, Serpent could be used once for the evaluation of the fine-grid nuclear data and then N times to perform less computationally expensive simulation in order to evaluate the reference solutions in N different configurations.

The optimisation process is started by generating N random individuals (candidate solutions), constituting the initial population. The GA process, indicated by the yellow boxes and enclosed in the light yellow rectangle in fig. 1, is implemented in a C++ code. For each individual, the `pybind` binding is called to pass the specific few-group structure from the C++ code to the Python package `COREutils`, which is used to collapse the group constants from the fine-group grid to the few-group structure (according to the considered individual) and to build the input case for the diffusion code. The weighting function employed in the collapsing process is the neutron flux spectrum evaluated by Serpent and featuring each homogenised spatial region. The individual is then evaluated with the nodal-diffusion module of the FRENETIC code. Once the entire population is evaluated, the overall fitness function featuring each individual is computed. Exploiting the fitness function and the genetic evolution mechanisms (described more in detail in II.C), a new generation is bred. This process is iterated until either the maximum number of generations or the required tolerance on the fitness function value is reached. Each of these steps is explained in detail in the following sections.

II.A. Generation of a fine-group cross section library

The optimisation process requires an initial fine-group cross section library for the reactor design under analysis, from which a subset of optimal groups will be identified. Since the aim of this paper is to explore the performances of the GA in combination with a physics-based fitness function, we focused our analysis on a specific reactor design and operational configuration. The reactor chosen for this work is the ALFRED core design developed in the LEADER EU project [17], with the control rods inserted at a height of +20 cm with respect to the mid-plane of the fissile region. For this specific operational configuration, both the fine-group cross sections and the reference quantities needed for the fitness function calculation (i.e., neutron flux, fission power distribution and k_{eff}) are computed with the continuous-energy Monte Carlo code Serpent 2, considering the whole core of ALFRED. A set of homogenised and collapsed cross sections is computed (simulating $6 \cdot 10^8$ active histories) for each 3D region depicted in figs. 2 and 3. The different homogenised regions featuring the core can be roughly grouped in fissile regions, which have two levels of enrichment (inner fuel IF with the lowest enrichment and outer fuel OF with the largest one), control regions (i.e., the active zones of the safety rods SR and control rods CR)

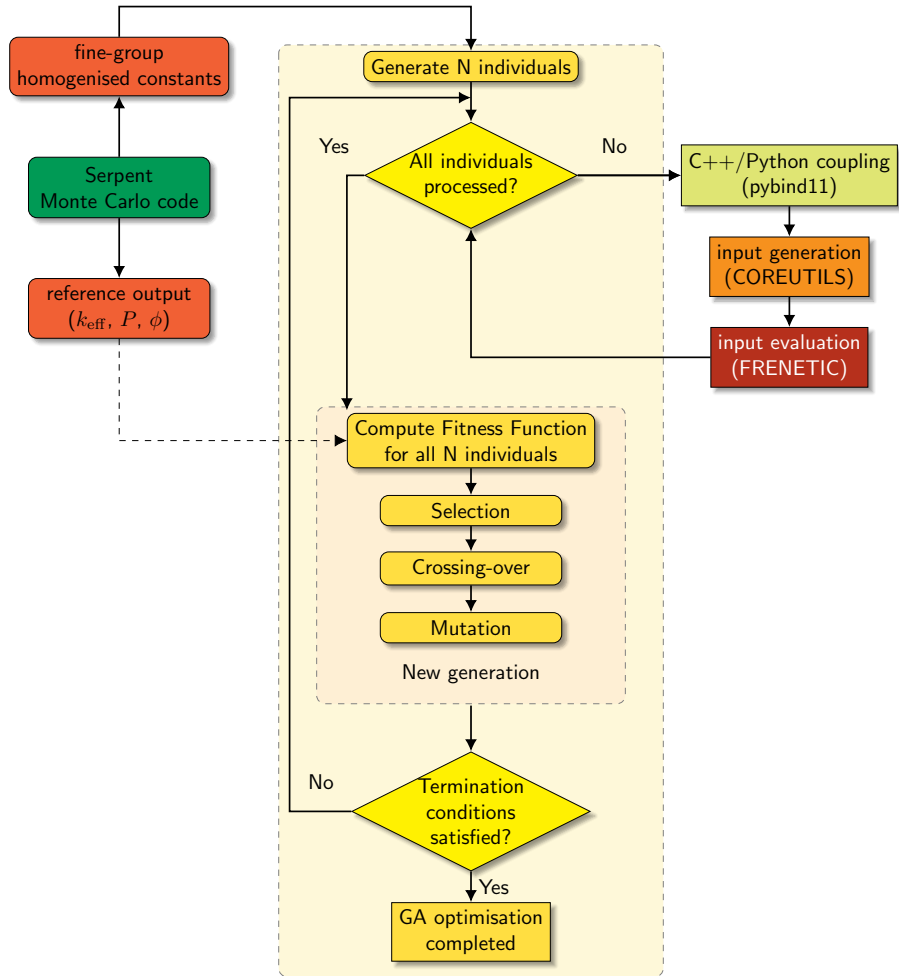


Fig. 1. Workflow diagram of the group structure optimisation process involved in this work.

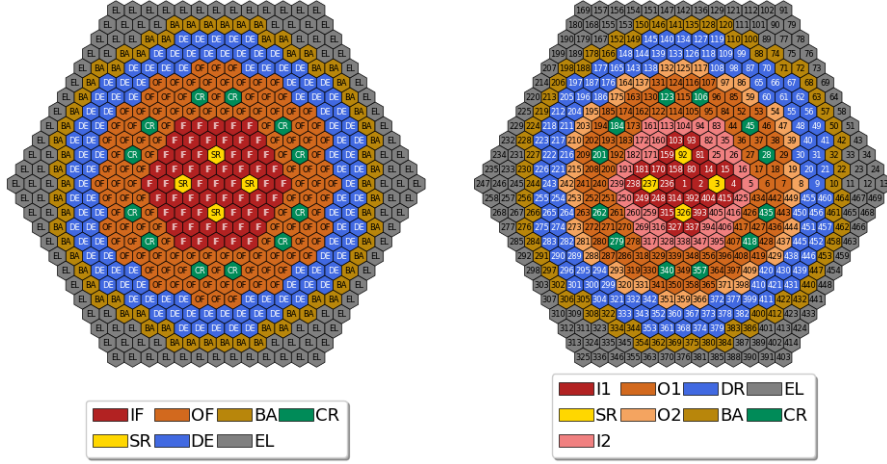


Fig. 2. $x-y$ view of the ALFRED 3D model reactor geometry (left) and subdivision considered for generating the fine-group constants homogenised at the assembly level (right).

and structural regions (the external lead EL, the barrel BA, the dummy assemblies DE, made of ZrYO, and the other structural components like plena, plugs, springs). The region subdivision selected for evaluating the group constants was proposed in a previous work [15]. The multi-group constants were tallied radially for each set of assemblies visible in the right sketch of fig. 2, i.e. I1 and I2 (inner fuel), O1 and O2 (outer fuel), DR, SR, CR, EL and BA. The group constants were tallied on two different sets of assemblies for the inner (I1 and I2) and outer (O1 and O2) fuels in order to improve the accuracy of the calculation.

With respect to our previous work, we decided to increase the axial detail of the ALFRED model for the generation of spatially homogenised and energy collapsed group constants, for improving the accuracy of the diffusion calculation. To this aim, the active zone of the fuel assemblies and the absorber zone of the control rods have been divided into 6 axial bins. Except for these regions, the group constants for are tallied on each axial region depicted with a different colour in fig. 3. This subdivision yielded fairly accurate results when compared to the Monte Carlo reference solution, despite the statistical error can be very large, around 20%, for specific combinations of energy groups and spatial regions. The most challenging regions for the generation of the multi-group constants are the peripheral ones for the fast groups and the control and fissile regions for the thermal groups.

To mitigate the statistical noise affecting some regions of the phase space (e.g., the thermal

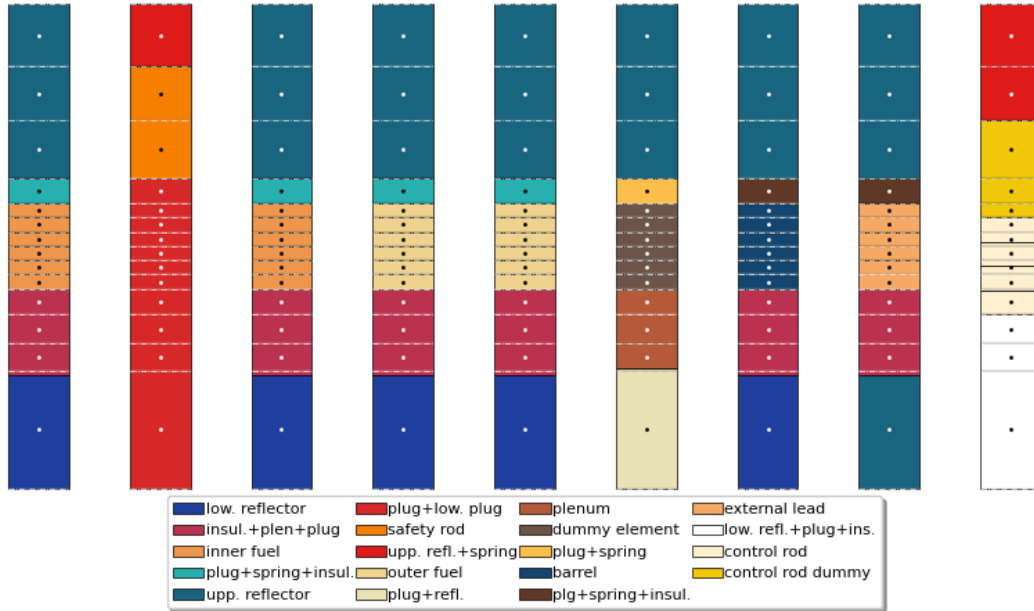


Fig. 3. Axial sections of the subassemblies in the ALFRED 3D model for the fine-group constants calculation. The black, solid lines divide the different axial regions implemented in the Serpent model, while the dashed lines divide the nodes considered for the FRENETIC calculations. The dots indicate the center of mass of each node.

energies in the control rod active zone or the fast energies in the core barrel) and the possible biases arising from the adoption of a pre-defined starting grid, we conceived a 61-group grid as follows: a single group covering the fast range (10-20 MeV), 56 groups uniformly spaced in lethargy for the epithermal region, 3 groups covering uniformly in lethargy the thermal region between 10^{-5} and $1.657227 \cdot 10^{-6}$ MeV and one group between $1.657227 \cdot 10^{-6}$ and 10^{-11} MeV. Since the most complex physical phenomena occur in the epithermal region, most of the groups should be allocated there. Therefore, considering a small number of groups in the fast and thermal regions does not affect the search for the optimal few-group structure. The 61-group grid is represented in fig. 4.

II.B. The FRENETIC code

The FRENETIC (Fast Reactor Neutronics/Thermal- hydraulICs) multiphysics code has been developed for the efficient neutronics (NE) and thermal-hydraulics (TH) coupled simulation of liquid metal-cooled cores featured by hexagonal closed assemblies [18]. The TH module offers a 1D advection/diffusion model for the flow and heat transfer of liquid metal within each closed assembly, taking into account the thermal coupling between neighbouring sub-assemblies, while

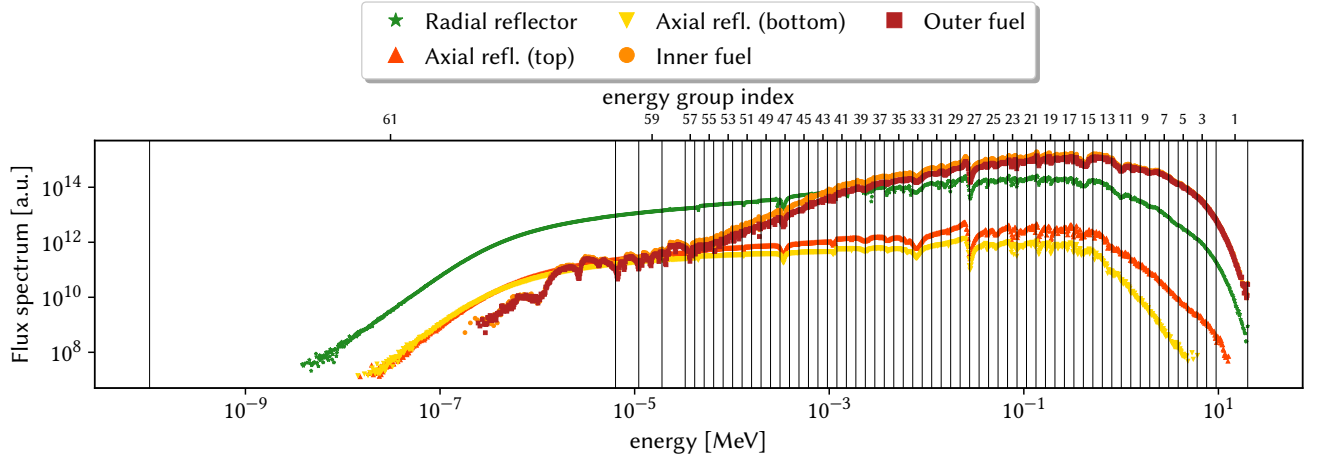


Fig. 4. Neutron energy spectrum, computed with *Serpent*, featuring the main regions of the AL-FRED reactor with the 61 multi-group structure employed for the GA optimisation.

the NE module offers a 3D multi-group diffusion model, which is spatially discretised employing a coarse mesh at the assembly level. More specifically, each sub-assembly is discretised with vertically stacked regular hexagonal prisms. The axial nodes employed for the calculation are sketched in fig. 3.

As previously mentioned, only the NE module of the code was used in this work, assuming that the reactor is operating at hot zero power. The relevant output computed with the code are the effective multiplication factor, the neutron flux distribution, the fission power density and the adjoint flux distribution. Depending on the number of groups, each 3D full-core calculation can take from a few seconds to a couple of minutes on a commercial laptop.

II.C. The genetic algorithm

Genetic algorithms, which try to solve an optimisation problem by emulating the mechanisms of biological evolution, have been successfully used since the 1970s for a broad set of applications, including the nuclear field (e.g. load pattern optimisation [19]). This class of techniques belongs to the so-called meta-heuristics, i.e. advanced search processes that, compared to rigorous optimisation algorithm, may not yield the global optimal solution but rather a solution close to optimum with a limited exploration of the solution space. As for the GA adopted in this paper, meta-heuristics often involve an initial random guess affecting the solution to the search problem. In this respect, the outcome of the optimisation should be interpreted as a "close-to-the-optimal"

solution, which can provide useful insights into the features of the "real" optimal solution and of the considered problem.

The first step of a GA typically consists in generating a random initial population of individuals (i.e., solutions to the problem). The first generation of individuals is judged according to the principle of natural selection, which is effectively implemented by a fitness function (FF): the best individuals have more chances to take part in the breeding phase, where they are mixed together to yield a new generation of individuals, and so on. Thanks to the natural selection mechanism, the average quality of the population tends to improve, since desirable features are inherited by the next generations, while disadvantageous features do not last. The evolution process is influenced by the selection pressure (i.e., the advantage given to solutions with better fitness over the others) and by the mutation rate, which helps keeping some variety in the solution pool. The first parameter guides the evolution through the generations, while the second one enables to explore possible interesting spots that could constitute a breakthrough in the search procedure.

II.C.1. Adaptive genetic operators

The evolution process must accomplish two tasks: improve the solutions in regions of the phase space already known (exploitation) and find out new untested spots of the phase space (exploration) [20]. Among the different mechanisms implemented in the GA, the balance of the two tasks is mainly regulated by mutation and crossing-over (XO). In general, mutation changes a few genes in chromosomes, so promoting a "fine-tuning" of the solutions, i.e. a search in the neighborhood of the known space. On the contrary, XO mixes already known individuals to generate brand new ones (hence starting from known genetic material), so promoting exploration of new regions of the phase space; the actual XO procedure depends on the chromosomic representation but, in general, takes portions of the genetic material of two "parent" individuals, selected based on their fitness, and recombines them into new individuals (~~daughters~~daughters).

The developed GA generates a new population starting from the previous step one in different proportions:

- A fraction f_E is given by the *elitist selection*: the best performing individuals of the previous generation are copied into the new population; such procedure helps protecting the best features from genetic drift, but a too high f_E could harm genetic diversity.

- A fraction f_{XO} of individuals comes from XO.
- The remaining slots (if any) are filled with uniformly randomly sampled individuals from the previous generation; this lowers the selection pressure and helps avoiding uniformity within the gene pool.

Finally, all individuals of this new population, independently from each other, have a probability p_M of undergoing mutation.

The selection method is the procedure that associate to each individual the probability of being picked for XO (i.e. to be a parent of the next generation): the fact that fittest individuals have a higher probability of having offspring is the driving force of the GA. The chosen selection method is the tournaments selection [21, 22]: the individuals are not ranked over the whole population; on the contrary each individual is assigned to a "tournament" and the selection probability is associated according to the ranking within that tournament only (per each position lower than the first, the probability lowers by $1 - p$), regardless of what happens in the other tournaments. In this way even weaker individuals have a chance of surviving, if they have the luck to enter tournaments with few strong individuals. This allows a large tunability of the selection pressure, that can be eased by increasing the number of tournaments or lowering the p and vice versa.

The initial proportion between XO and mutation rates, along with their trend as the simulation progresses, are the object of numerous studies in literature (e.g. [23, 24]): in some cases, the rates are linked to the number of generations elapsed, while more recent studies adjust the rates based on variability indexes of the population. As typical of the hyperparameters of a meta-heuristics, there is not a clear answer that fits all cases. We have tested different combinations of mutation rates, both fixed and variable with the generation number. The study of this topic, however, lies outside the scope of this paper, so the calculations presented in the next sections have been performed with fixed parameters, as reported in Table table I.

II.D. Chromosomic representation

Our previous works focused on optimizing the group structure: the number of groups had to be fixed *a priori* by the user. Such constraint is particularly tight: depending on the case, an additional group can provide a huge boost to the solution or represent an unacceptable burden; in some configurations, increasing the number of groups might also result in an eased convergence

TABLE I
GA parameters

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

and so, counter-intuitively, in lower computational burden. Removing this limitation, however, requires a drastic change in the solution representation and bring some issues to be considered. Firstly, the sample space is largely extended: when the number of energy groups is fixed, the sample space cardinality is $\binom{n}{k}$, with n equal to the number of energy cuts in the fine grid and k equal to the number of energy cuts in the coarser grid, while without such limitation the size increases up to 2^n . When k is small with respect to n , this means that the sample space cardinality increases by a few orders of magnitudes. In addition, one has an additional level of complexity given by an additional random variable (the number of groups), which however is not independent from the group structure; as will be explained in detail in this section, this favours some solutions over the others, requiring special care to restore balance to the playing ground. For instance, for $k = 5$, a total of more than $585 \cdot 10^3$ combinations would be possible. In spite of the HPC resources available today, the brute-force evaluation of these group structures would still be very intensive, even assuming a very efficient diffusion solver. On top of this, these calculations would be referred to a very specific case.

In full analogy with biology, also in the genetic optimisation process the features characterizing an individual are coded in sets of genes, which can assume different values (alleles). The genes are collected in chromosomes, which constitute the genetic makeup of an individual. The representation of the genetic makeup of the solutions is a fundamental aspect, as it affects how the solution space is represented and sampled, so influencing the evolution of the population. As the optimised energy group structure will express a subset of energy boundaries of the starting grid, there are two rather intuitive possible representations of each individual, detailed as follows:

- Representation α : an individual is defined by a single chromosome, with a gene associated to each possible energy boundary of the fine grid. As the minimum and maximum boundaries

are always included, there are $H - 1$ genes, with H being the number of groups in the fine structure; the alleles are the logical values: 1 (true) if the group boundary is retained in the broad-group structure, 0 (false) otherwise.

- Representation β : an individual is defined by two chromosomes, one representing the number of groups, the other the group boundaries. The first chromosome β_0 contains a single gene, whose value G is an integer numbers from 1 to H , while the second chromosome β_1 is featured by $G - 1$ genes whose alleles are the indexes of the retained group boundaries, i.e. integer numbers from 2 to H (in this case, $H = 61$).

A visual representation of the two genetic makeups is shown in fig. 5.

The two representations, though simple and intuitive, do not allow a fair exploration of the whole solution space (all possible number of groups and possible energy grids): in facts they tend to favour solutions due to their shape rather than because of their fitness, so introducing a systematic bias that inevitably hinders the natural evolution. Each group of the starting fine grid can be retained or not: therefore, a natural sample space is

$$\Omega = \{0, 1\}^{H-1} \quad (1)$$

with $|\Omega| = 2^{H-1}$. The elements of Ω are $H - 1$ -dimensional vectors

$$\vec{X} = (X_i \in \{0, 1\})_{i=2\dots H} \quad (2)$$

It can be helpful observing that \vec{X} can be easily mapped onto the natural numbers through the binary representation. Therefore, we can define the random variables:

- X_i , the i -th component of \vec{X} , i.e. the indicator function of the event “the i -th energy boundary is retained in \vec{X} ”.
- G , the number of groups of \vec{X} , i.e.

$$G = \sum_{i=1}^H X_i \quad (3)$$

With these definitions, as in representation α all genes, independently from each other, assume the allele 0 or 1 based on a Bernoulli distribution (with $p = 0.5$), then:

- $\mathbb{P}(X_i)_{\forall i \in \{2 \dots H\}}$ has a uniform probability distribution;
- $\mathbb{P}(G)$ has a binomial probability distribution, being the sum of $H - 1$ Bernoulli distributions.

This means that the solutions with a number of groups close to $H/2$ are extremely favoured over more extremal ones: this may seem correct, as the solutions with such number of groups are more ($\mathbb{P}(\vec{X})$ has a uniform distribution), but this means also that the genetic operators will often try to “revert to the mean”, so hindering the evolution process toward regions with lower groups and making almost impossible (or at least rather unlikely) the exploration of zones with less (or more) groups than the average.

Representation β , instead, allows an unbiased exploration of the number of groups, as it generates

- $\mathbb{P}(G)$ according to a uniform probability distribution;
- $\mathbb{P}(X_i|G)_{\forall i \in \{2 \dots H\}}$ according to a uniform probability distribution.

Of course, $\mathbb{P}(\vec{X})$ does not have a uniform probability distribution anymore, but this is not mandatory (or desirable). Actually, this representation looks more suitable to the studied problem, as it allows to adequately explore all possible number of groups; this solution, however, is inefficient from the evolutionary point of view because of the way the genetic operators work. As a first issue, β_1 has a different length depending on β_0 ; arrays with a non-constant length are generally problematic when a XO must be performed. To ensure the consistency between the two arrays one should then introduce some corrective measures, which would impact the natural evolution to a large extent. Therefore, we have set β_1 as a fixed length array, containing a permutation of the integers $2 \dots H$; the chosen groups are always the first G , with G being defined by β_0 . This workaround indeed solves an issue, but still require quite invasive corrections to avoid duplicates in β_1 . Also, each individual can evolve only the expressed genes of the β_1 (i.e. the first G ones), while the others are randomly arranged; this is not a problem as long as the last genes are not expressed, but this is not ensured as the first chromosome can also change. This unfairly penalizes the solutions with a larger number of groups, which have to deal with less efficient group structure.




Fig. 5. Representation α (left) and β (right).

II.D.1. Adopted representation

Aiming at fixing the issues observed with the first two representations and combining their best features, we have developed a third chromosome structure, denoted γ , which relies on polyploid chromosomes. The karyotype of an individual (visually sketched in 6) is represented by two chromosomes, γ_0 and γ_1 : γ_0 is a simple chromosome with a single gene (haploid), whose value is the number of boundaries retained G in the condensed energy grid (exactly as β_0). γ_1 is a polyploid chromosome, as it consists of P arrays, each one with $H - 1$ genes, one per each possible boundary in the original fine grid; the possible alleles of each gene are the natural numbers between 0 and M , where P and M are hyperparameters that can be tuned by the user. The expression of γ_1 is a vector of length $H - 1$ given by the sum of all P vectors: the chosen groups will be those with the G highest score (ties are broken randomly among those with the same score, ensuring that an individual does not change its expression as long as it survives).

The XO operation is performed in a different way: the polyploid chromosome of the offsprings is unchanged, except for a single (randomly chosen) γ_1 array: this is replaced by a new one, created based on the expression of the partner. The new array has a uniform distribution of the alleles $0 \dots M$, assigned following the obtained score in the expression phase. In this way, the evolution is a less disruptive process, with the best boundaries slowly accumulating good alleles. The side effect is that the evolution pace is reduced, as the adjacency between chosen groups (which for the energy grid problem plays a role) is better preserved and less couples are disrupted during a single iteration. Moreover, this solution addresses the uneven evolution pace observed with the second representation: in this case, even if only the top- G boundaries are expressed, all alleles (including the unexpressed ones) receives feedback (positive or negative) through the new array. The process fosters a ranking of the boundaries, with all positions having received the same feedback information during the evolution process.




Fig. 6. Representation γ .

II.E. Fitness functions

The other aspect impacting the design of the GA process is the choice of the FF. Since the typical outcomes of an eigenvalue calculation are k_{eff} , the power density and the flux distribution, a natural choice would be to take them all into account in the definition of the FF. As in our previous works [12, 13], we defined "specialised" FFs, devoted to single quantities, that are then suitably averaged into a single FF. The k_{eff} FF is defined simply as:

$$f_k^I = |k^I - k^{\text{MC}}| \cdot 10^5, \quad (4)$$

where k^I refers to the eigenvalue of a certain individual I , while MC indicates the reference value computed with the Monte Carlo code. The FF for the power is the relative error between the power vectors computed with FRENETIC and Monte Carlo,

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

where j indicates the spatial region, while the one for the neutron flux is a weighted relative error between the flux vector computed with FRENETIC and the Monte Carlo solution, which is collapsed on the same grid used by FRENETIC for consistency,

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

where g indicates the energy group. The set of weights $w_{g,j}$ is normalised, i.e.

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

and is obtained by normalising the integral on the phase space of a quantity related to the flux adjoint, as detailed in the following. Clearly, this choice for the fitness functions is arbitrary, but we believe that it is quite natural since it aims at preserving the physical meaning of neutronics-relevant quantities. In this work, the FFs defined in eqs. (5) and (6) have been preferred to other possible alternatives like scaled absolute errors because they are stricter: with the scaled absolute error metric, the cells featured by a lower value of the flux would become less important.

In the following, three choices are considered for the construction of the weight:

- the adjoint flux ϕ^\dagger , which physically represents the importance of a single neutron introduced in (\vec{r}, E) ;
- the product of the adjoint flux times the direct flux, $\phi^\dagger \phi$, i.e. the *track length importance*;
- the product of the adjoint times the neutron density, $\psi = \phi^\dagger \phi / v = \phi^\dagger n$, which yields the *total importance* in (\vec{r}, E) . Exploiting the relationship between direct and adjoint operators, it can be demonstrated that $\langle \psi \rangle$ is constant in time, where $\langle \rangle$ indicates an integration over the whole phase space [25].

As it will be shown in section III, the weight definition yielding better results is the last one [19], which finds its physical justification in Contribution Theory [26].

Instead of eq. (5) and 6, alternative L_2 -like metrics can be possible, like the scaled absolute errors was driven. However, in this work we chose the "local" relative error to avoid that cells featured by a low absolute value may be disregarded.

Figure 7 shows the behaviour of the space-integrated, group-wise distributions of the neutron flux, density, adjoint flux, track length importance and total importance. Each quantity is normalised with respect to the lethargy. The flux spectrum shows the typical behaviour of a fast reactor, featuring very low values below 1 eV. On the contrary, the neutron density, i.e. the flux divided by the neutron speed, is maximum in the thermal groups and is progressively reduced moving to higher energies, since $1/v = \sqrt{2m_n/E}$. The adjoint shows two local maxima around 10 eV and




Fig. 7. Distributions of the direct flux, adjoint flux, neutron density and total importance scored on the starting 61 multi-group structure employed for the GA optimisation. All quantities are normalised with respect to the lethargy.

10 MeV, and a local depression between 100 eV and 1 MeV. The shape of the importance-weighted track length is very close to the direct flux, while the total importance spectrum is broadened and shifted towards lower energies compared to the direct flux, as a consequence of the large number of low-energy neutrons.

Figure 8, fig. 9 and fig. 10 show the group-wise, radial distribution of the direct flux (left), the adjoint flux (centre) and total importance (right) at different axial planes. The groups were selected to represent these distributions in the fast ($G=1$), epithermal ($G=30$) and thermal ($G=60$) ranges. The distributions in fig. 8 refer to the axial plane containing the fuel and dummy plena homogenised with the structural materials, those in fig. 9 refer to the plane containing the fuel assemblies and the absorber of the control rods while those in fig. 10 refer to the plane containing structural materials and the active region of the safety rods. The radial distribution of the flux (and of the neutron density, which is spatially identical to the flux) clearly shows that the combination of high leakages with the fast spectrum induces intricate spatial and energy effects, which make the selection of appropriate energy structures very challenging. By inspection of the central images in figs. 8 to 10, the adjoint flux appears to be significant only in the inner regions of the core.

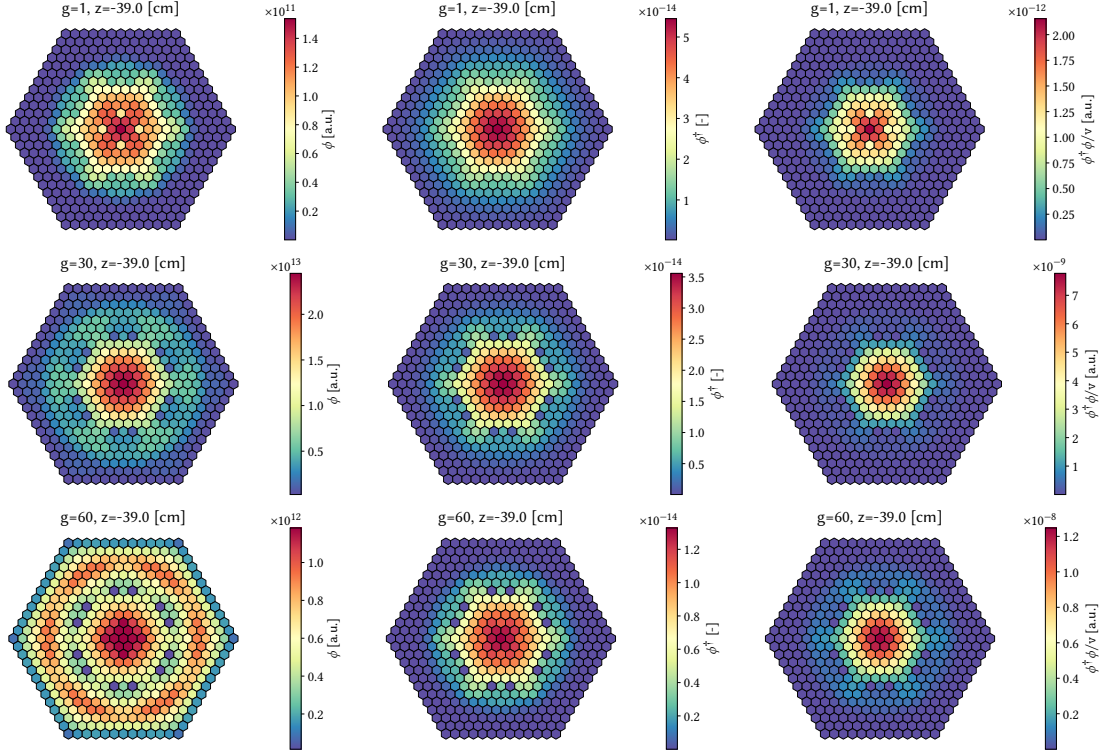


Fig. 8. Assembly-wise distribution of the direct flux (left), adjoint flux (centre) and total importance (right) for energy groups 1 (top), 15 (centre) and 60 (bottom), computed with FRENETIC. The maps refer to the axial plane at -39 cm, in correspondence of the fuel assembly plena.

The resulting total neutron importance follows this trend, except in the thermal group at the fuel plane, where the importance is maximum at the fuel-dummy interface. As a consequence of this energy and space behaviour, the points of the phase space that contribute more in sustaining the fission reaction chain contribute more to the FF in eq. (6), while the peripheral regions (e.g., the core barrel), which may be affected by large errors and have lower importance, are disregarded.

Since Serpent is not equipped for tallying the adjoint flux, an approximation to ϕ^\dagger is computed with FRENETIC on a 61-group calculation. In eq. (6), the MC solution is collapsed from H to G groups, allowing a direct and consistent comparison between the two quantities.

II.E.1. Joint fitness function

When more than one fitness function is defined, the problem of balancing all of them at the same time arises: all solutions on the Pareto front are, by definition, the best “in their own way”; nonetheless, from an engineering perspective, a balanced solution is preferred over one that

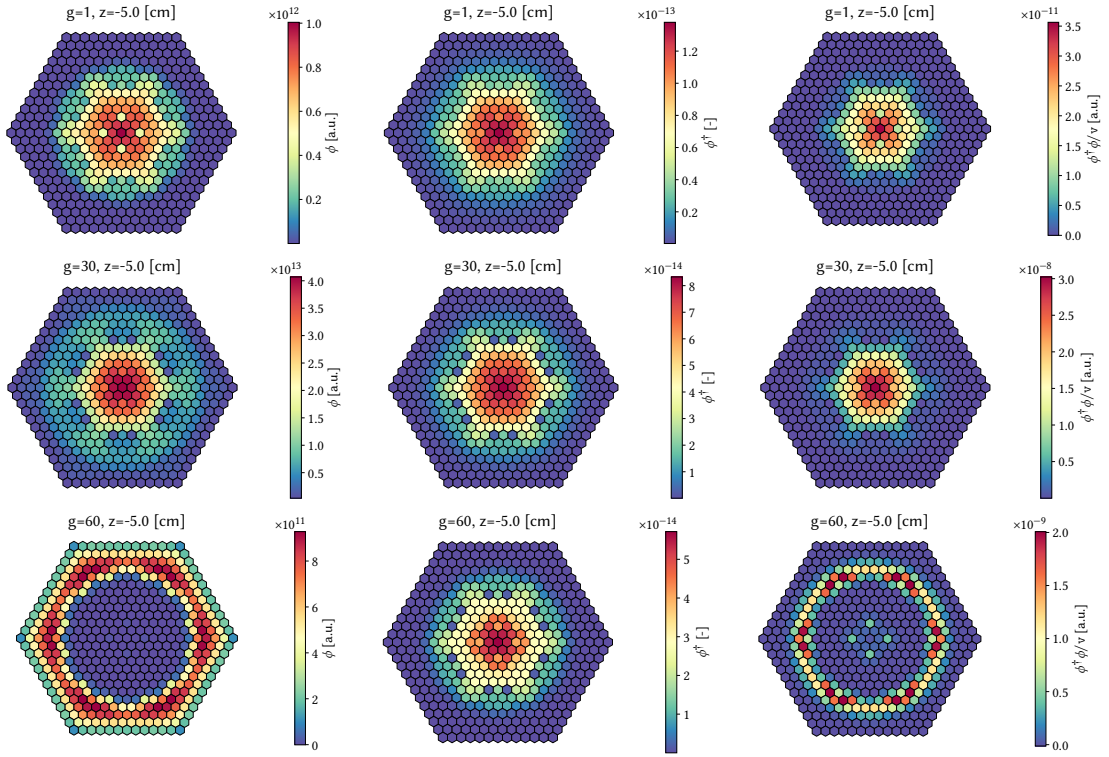


Fig. 9. Assembly-wise distribution of the direct flux (left), adjoint flux (centre) and total importance (right) for energy groups 1 (top), 15 (centre) and 60 (bottom), computed with FRENATIC. The maps refer to the axial plane at $z=-5$ cm, in correspondence of the fuel region.

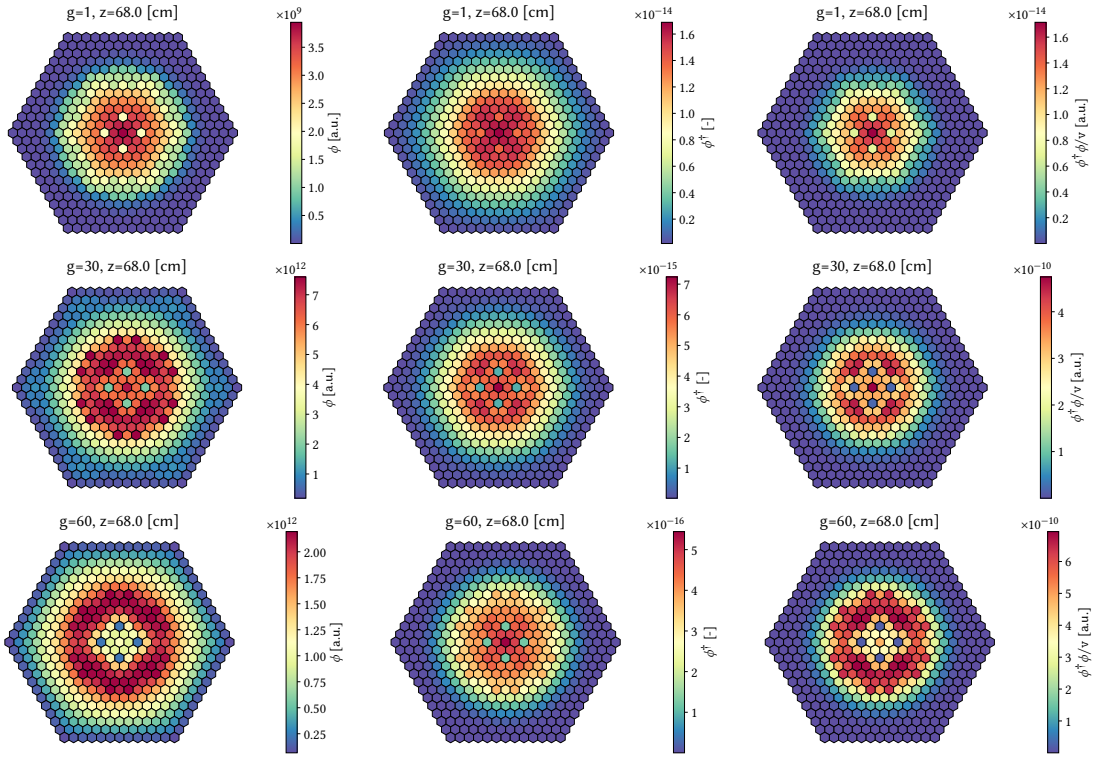


Fig. 10. Assembly-wise distribution of the direct flux (left), adjoint flux (centre) and total importance (right) for energy groups 1 (top), 15 (centre) and 60 (bottom), computed with FRENETIC. The maps refer to the axial plane at 68 cm, in correspondence of the safety rod absorber.

perfectly matches one constraint with bad results in the other. Therefore, we aim at designing a functional that summarizes the results of all FFs into a single value. In doing so, we have to take into account two issues:

- The FFs span very different scales: some FF assume values ranging different order of magnitudes (e.g., the flux), while others can change of a few percentage points (e.g., the power). All FFs are built such that their codomain is $[m, +\infty)$, where m is the best. It is not guaranteed that their image is actually covering the whole set, nor that a solution with a minimum value m actually exists. As explained in the following, m is assumed, for this specific application, to be a strictly positive number.
- The range and actual variability of the FF are not known *a priori*: as the evolution progresses, new solutions emerge, and their associated FF value, though in the codomain for sure, might not be in the range observed so far. This requires special care in case rescaling is implemented.

If these points are neglected, the effect of a FF can become dominant over the others or, on the contrary, completely negligible.

Since the previous GA applications required the use of a modified version of the SIMMER code [11], a more general and flexible version of the GA was implemented from scratch in a C++ framework, enabling its adoption for a wide variety of optimisation problems, thanks to the definition of a user-dependent FF embedding all the instructions needed for the specific application.

In order to deal with a wide variety of cases, the functional is in the form

$$G : \mathbb{R}^W \rightarrow \mathbb{R}$$

$$G : (f_1^I, f_2^I, \dots, f_W^I) \mapsto G[h_1(f_1^I), h_2(f_2^I), \dots, h_W(f_W^I)] \quad (8)$$

with h_i being transformations of each fitness function value and W the number of FF considered.

The h functions implemented in the algorithm and tested are:

- *Identity*

$$h_{\mathbb{1}} : f \mapsto f \quad (9)$$

- *Linear scaling*, which linearly maps all ever observed FF values on the arbitrary interval

[1, 100],

$$h_L : f \mapsto 1 + 99 \cdot \frac{f - \min_{I \in E}(f^I)}{\max_{I \in E}(f^I) - \min_{I \in E}(f^I)} \quad (10)$$

being $E \subseteq \Omega$ the explored (throughout the entire evolution history) sample space.

- *Ranking*, which associates to a FF value f , the percentile (real values are allowed) it belongs to in the set of the values ever observed

$$h_R : f \mapsto 1 + 99 \cdot \frac{\text{card}\{I \in E : f^I > f\}}{\text{card}(E)}. \quad (11)$$

The latter two options help compare FFs behaving very differently and spanning different orders of magnitudes. The drawback is that the fitness results are not static anymore: in fact all values change over time, depending on the explored space; therefore, one does not have an absolute measure of the goodness of the solution. A recalculation mechanism has been implemented in order to keep consistency across different generations within the same GA run, but the comparison among different runs cannot rely on rescaled or ranking FFs. The linear rescaling (eq. (10)) solution appears more natural, so it has been adopted, but it is very sensitive to outliers: FFs ranging over different orders of magnitudes show an accumulation toward the extremes, making the FF less effective.

After the transformation, which should make the FFs comparable with each other, the functional G is applied to extract a single fitness value: the most natural operations are the sum (corresponding to the arithmetic average) and the product (geometric average), but other operations are possible, as maximum and minimum. We have adopted the geometric average as, although some compensation is allowed, largely unbalanced solutions are penalized. The FF used to enforce the natural selection is thus:

$$FF^I = h_L(f_k^I) \cdot h_L(f_P^I) \cdot h_L(f_\phi^I). \quad (12)$$

Due to the integral nature of k_{eff} , it is possible that f_k^I achieves very small values, say below 1 pcm. Hence, to avoid that the product of the FFs expressed by eq. (12) gets very small without an effective minimisation of the other FFs, we artificially set any value of $f_k < 10$ to 10 pcm, which is judged a satisfactory target accuracy for k_{eff} .

III. RESULTS

In this section we present the results obtained with the three chromosome representations proposed in section II.D and with the different definitions of the flux FF weight sketched in section II.E.

III.A. Effect of the chromosome structure

In this section we compare the performances obtained by running the GA with different chromosome structures (karyotypes). Since the generation of the population is a combination of random selection and genetic mechanisms (cross-over, mutation...) the outcomes of the GA are not deterministic.

As the GA is essentially a random process, the evolution does have a certain dependency, especially in the early phases, on the random numbers generated; in order to appreciate the impact of the random number generator on the performances of the GA, a set of three simulations per each case is carried out with different seeds. Each population thus evolves independently, like "tribes" populating separated environments. Since the random seed of the GA library is associated with an input string, we named the three tribes "Boltzmann", "Fermi" and "Henry".

III.A.1. α chromosome structure

Figure 11 displays the convergence trend of the population considering the α chromosome structure. The left graph shows the evolution of the generation-wise FF, obtained by taking the average of the fitness of each individual pertaining to a generation. The convergence history of the average fitness \bar{f} is obtained with three different tribes, as the random initialisation has a non-negligible impact on the convergence rate. As an example, the Fermi tribe requires roughly twice the generations of the first two individuals in order to reach the value $\bar{f} = 2000$. The graph on the right shows the FF value as well as the number of groups featuring each individual in the Fermi tribe. Despite the smooth convergence rate, this genetic makeup does not allow to uniformly explore the solution space: as discussed in section II.D and visible from fig. 12, the number of groups G in the starting population follows a binomial distribution, with the expected values equal to the half of the number of alleles. As generations go by, the genetic mechanisms try to push the distribution towards a smaller number of group, but the adopted probability space inevitably

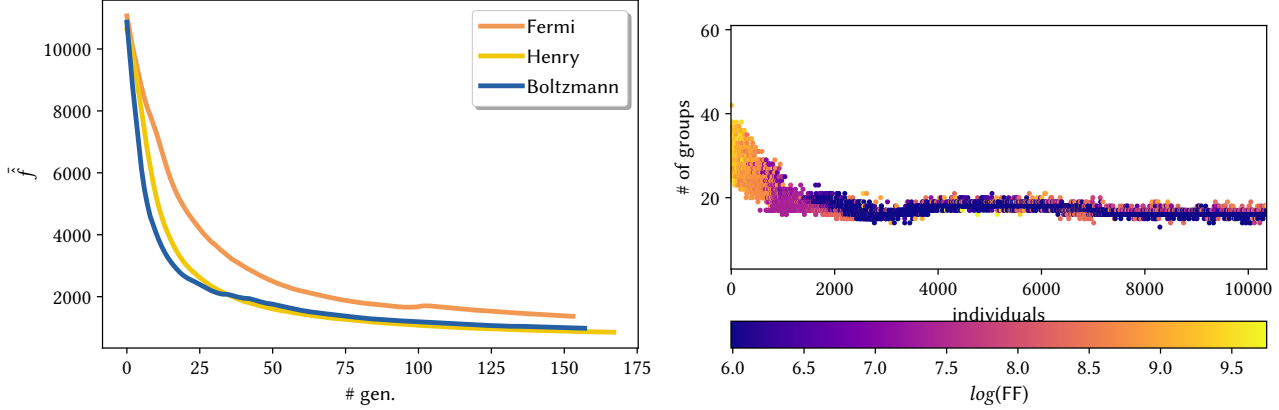


Fig. 11. Convergence history of the average fitness function (left) and value of the fitness function against the number of groups and individual number (right) for the α chromosome structure. "Boltzmann", "Fermi" and "Henry" are three independent random seed used to initialise the GA. The seed used for the right graph is "Fermi".

opposes to such movement, keeping the average G close to 20. Figure 11 also shows that around 10000 unique individuals are enough for the GA to reach its termination condition. This number is totally negligible when compared to the total number of possible combinations, equal to 2^{61} , furtherly justifying the employment of the GA for the selection of the optimal energy grid.

In order to facilitate the physical interpretation of the outcomes of the GA selection, the best group structures are superimposed to some relevant group-wise reaction rates in fig. 13. Specifically, we analysed the group-wise fission power, i.e. the amount of power deposited by fission in each group

$$P_g = \int_V d\vec{r} E_{f,g}(\vec{r}) \Sigma_{f,g}(\vec{r}) \phi_g(\vec{r}), \quad (13)$$

the group-wise reproduction factor integrated over the whole reactor volume $\nu \Sigma_{f,g} / \Sigma_{a,g}$, which is related to the k_∞ of the system, and the removal reaction rate at different axial planes. In order to compare the behaviour of these reaction rates with respect to energy, the contribution of each group to the total, i.e.

$$r_h = \frac{RR_h}{\sum_{h=1}^H RR_h}, \quad (14)$$

is normalised per unit lethargy. The dashed black vertical lines represent the starting 61-group grid structure, while the dots represent the optimal few-group grids produced with the three random initialisation.

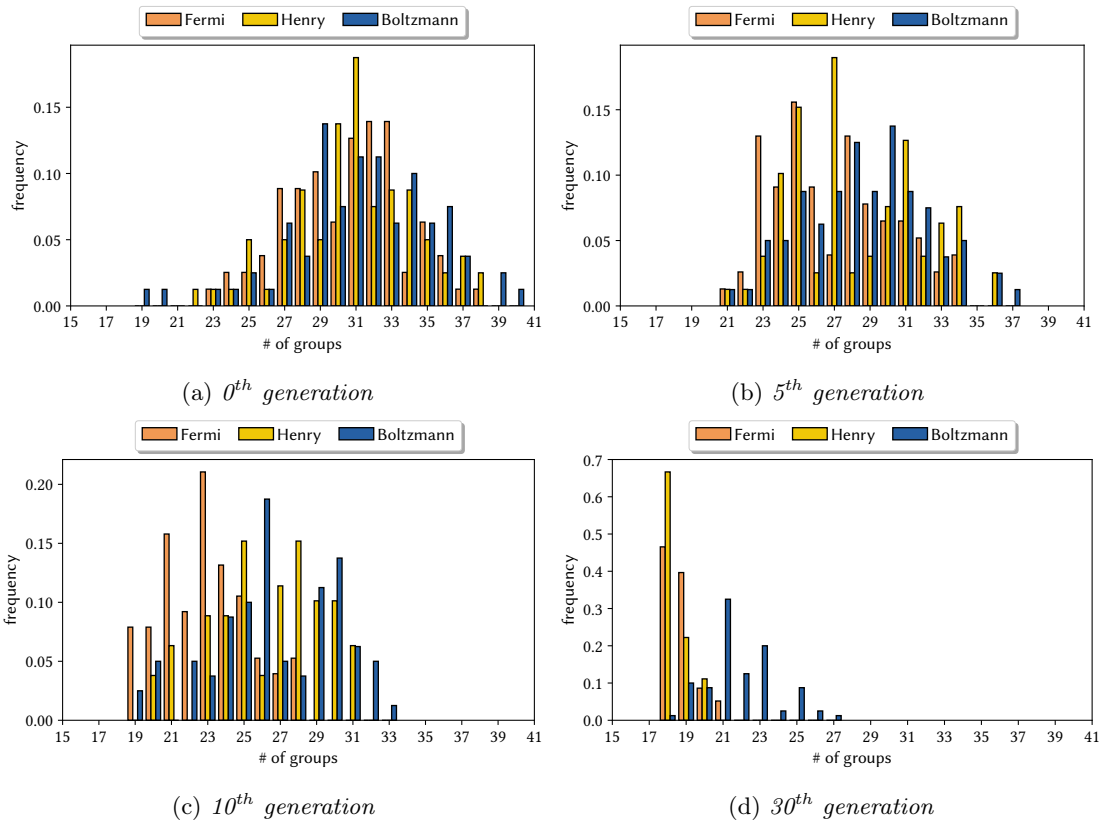





Fig. 12. Distribution of individuals per number of groups at different generations, for the α chromosome structure.



(a)



(b) Inner fuel SAs.



(c) Dummy rod SAs.

Fig. 13. Group-wise quantities and optimal group boundaries (dots) selected from the reference 61-group structure (dashed black lines) for the α chromosome structure.

The solution found in the Boltzmann and Fermi tribes are quite similar, sharing the same low-energy groups and having energy boundaries quite close also in the epithermal and fast regions. Despite the solution found in the Henry tribe has nearly half the groups of the Boltzmann case, its distribution is very close to the ones of the other two solutions; the missing groups are mainly those of the low-energy region. Except for these thermal groups, most of its group boundaries match the ones appearing in Boltzmann and Fermi, suggesting that they could be sufficient to cover the same energy range with a lower number of groups. This fact suggests to study the trade-off between the computational time reduction and the accuracy of the calculation.

III.A.2. β chromosome structure

Exploiting the results obtained with the α karyotype, the optimisations performed with the β chromosome structure were carried out considering a maximum number of groups equal to 51. The performances of this second optimisation campaign are summarised in fig. 14. As compared to the α case, these calculations stopped at different generations according to the random seed used for the initialisation, due to a stopping criterion based on a maximum number of "inefficient generations": if, within a certain number of generations, an individual better than the current "best" one is not found, then the calculation is stopped. While in the previous case the population quickly converged to a fixed number of groups (a sign of homogeneous population), the parameter space is now well explored in all the generations, as visible from the graph on the right. Good results are obtained with a number of groups between 10 and 25, with minimum values of the FF reached for individuals with a number of groups between 10 and 15.

The distributions of the number of groups according to the generation can be appreciated in fig. 15. In the 0th generation the number of groups is roughly uniform, while after 5 generation the distribution starts to be skewed towards a number included between 10 and 15. As it will be discussed more in detail in section III.A.3, the number of groups is a non-trivial variable in the solution space. In transport theory, it is intuitive to expect that the larger the number of groups, the more accurate the calculation will be. However, since FRENETIC adopts a diffusion model, it is quite natural that the most accurate solutions lay in the space featured by a lower number of groups (see section I).

As visible in fig. 16, the best individuals produced with this chromosome representation share

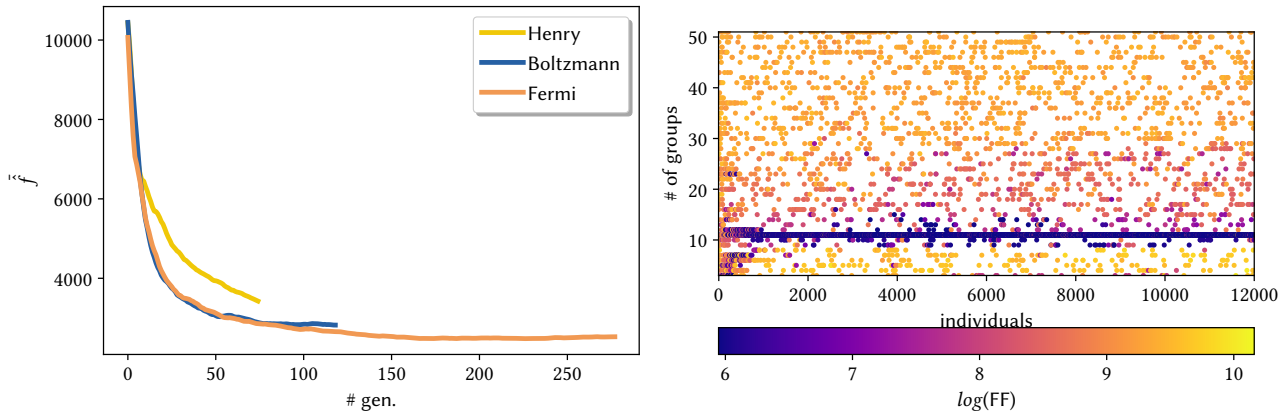


Fig. 14. Convergence history of the average fitness function (left) and value of the fitness function against the number of groups and individual number (right) for the β chromosome structure. "Boltzmann", "Fermi" and "Henry" are three independent random seed used to initialise the GA. The seed used for the right graph is "Fermi".

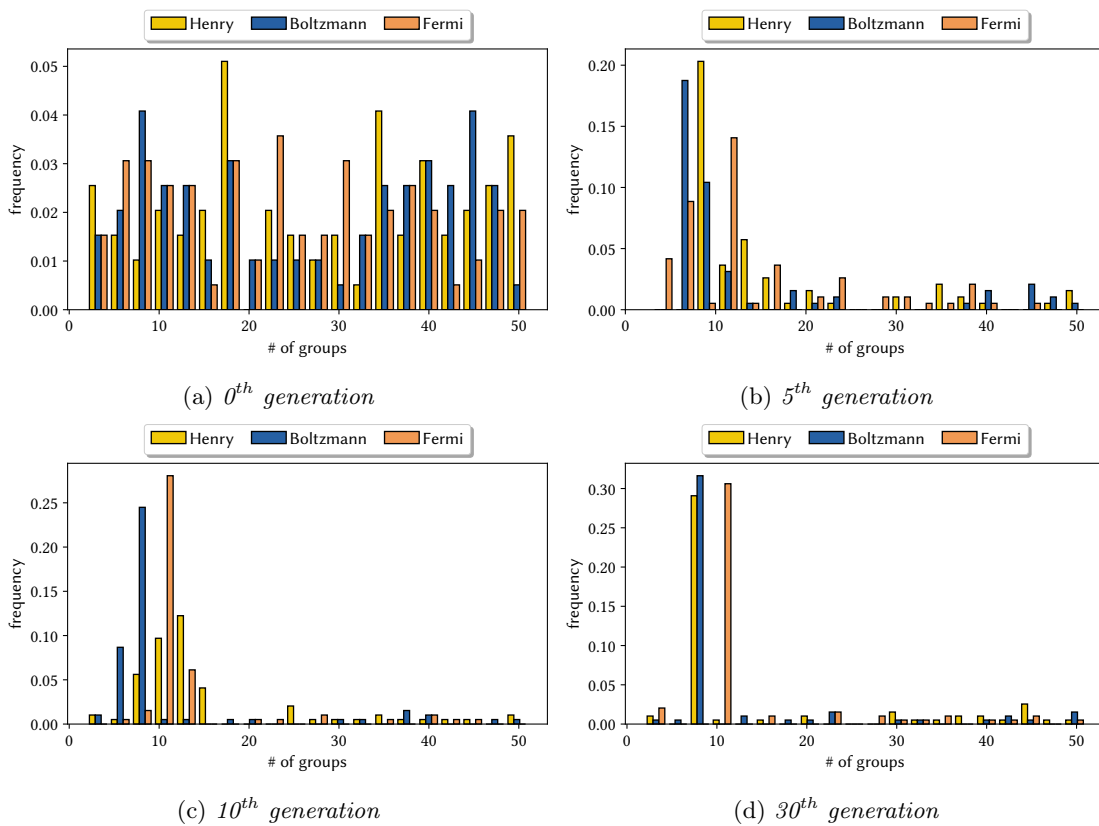
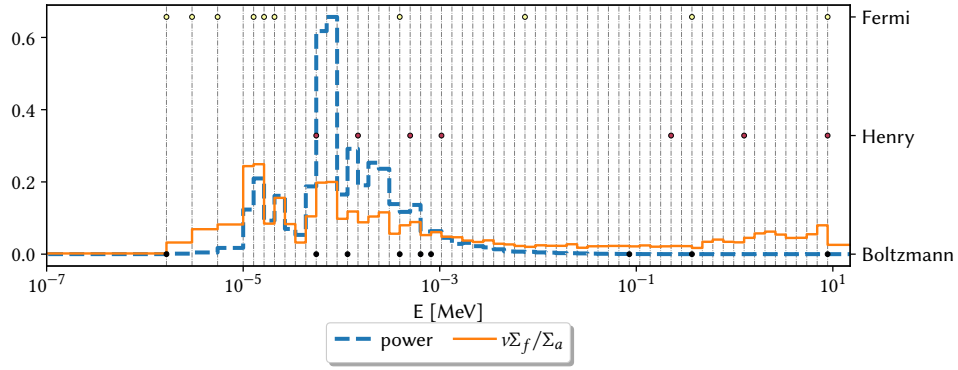
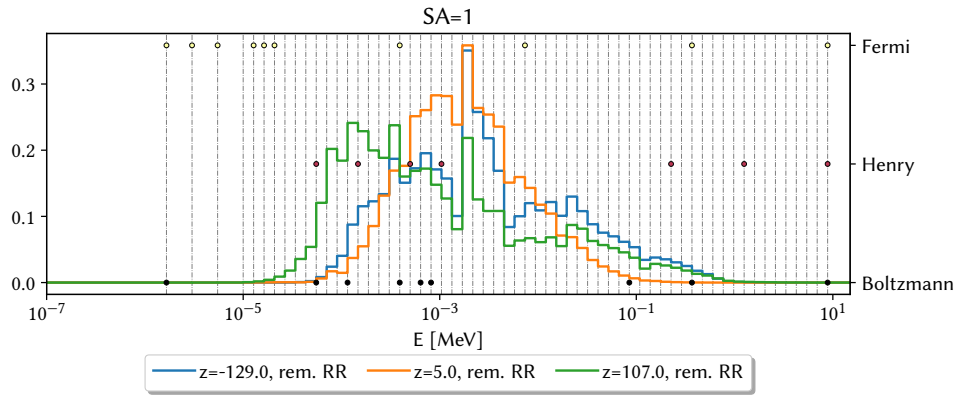



Fig. 15. Distribution of individuals per number of groups for the single-chromosome makeup at different generation numbers for the β chromosome structure.



(a)



(b) Inner fuel SAs.



(c) Dummy rod SAs.

Fig. 16. Group-wise quantities and optimal group boundaries (dots) selected from the reference 61-group structure (dashed black lines) for the β chromosome structure.

some features with the ones generated with the α karyotype: most of the group boundaries are invested in the low epithermal energy region, where the group-wise power and reproduction factor exhibit the most significant variations, while the remaining groups are assigned to the fast region.

III.A.3. γ chromosome structure

The γ chromosome structure is more complex than the β one, but allows improved explorations of the solution space. Considering that the best individuals obtained with the β representation have less than 15 groups, we limited the maximum number of groups in the grid to 20, which seems an already good trade-off between accuracy and computational time. Indeed, the calculations with α and β showed that individuals with more than 20 groups are unlikely to be the best ones, but waste a lot of computational time, as a calculation with many groups has to be performed; in addition, the risk of non-convergence is higher. The convergence trend and the evolution of the values of the FF are very similar to the β case, showing a good exploration of the solution space. Hence, they are not reported for the sake of brevity.

For this chromosome structure, it is interesting to observe the evolution of the single FFs, whose combination with eq. (12) leads to the average FF used to drive the GA. It can be noticed from fig. 17 that the values of f_ϕ tend to be very large, except for a few cases with less than 10 groups. Despite the total importance weighting, the relative error computed with eq. (6) may assume large values when $\phi_{g,j}^{MC}$ is small. Therefore, solutions featured by a low number of equally log-spaced groups tend to be favoured. Conversely, the FF for k_{eff} is minimised with a number of groups between 8 and 15. The behaviour of the power FF is similar to the one of the flux FF, despite its variation is quite limited. Figure 18 helps to shed some light on this narrow variation. The graphs show the difference and the relative difference between the power distribution computed by Serpent and FRENETIC for the best and worst individuals spawned in the Boltzmann tribe. While the simple difference in the power computed with the two codes changes significantly from the best to the worst individual, the spatial distribution of the relative difference is quite similar for the two individuals and features local maxima in the peripheral regions of the core, where the magnitude of the power is minimum. Since eq. (5) evaluates the relative difference in each spatial region, the largest relative variations occurring at the core periphery dominate the FF, yielding only a slight difference between good and poor individual. These observations point out the need

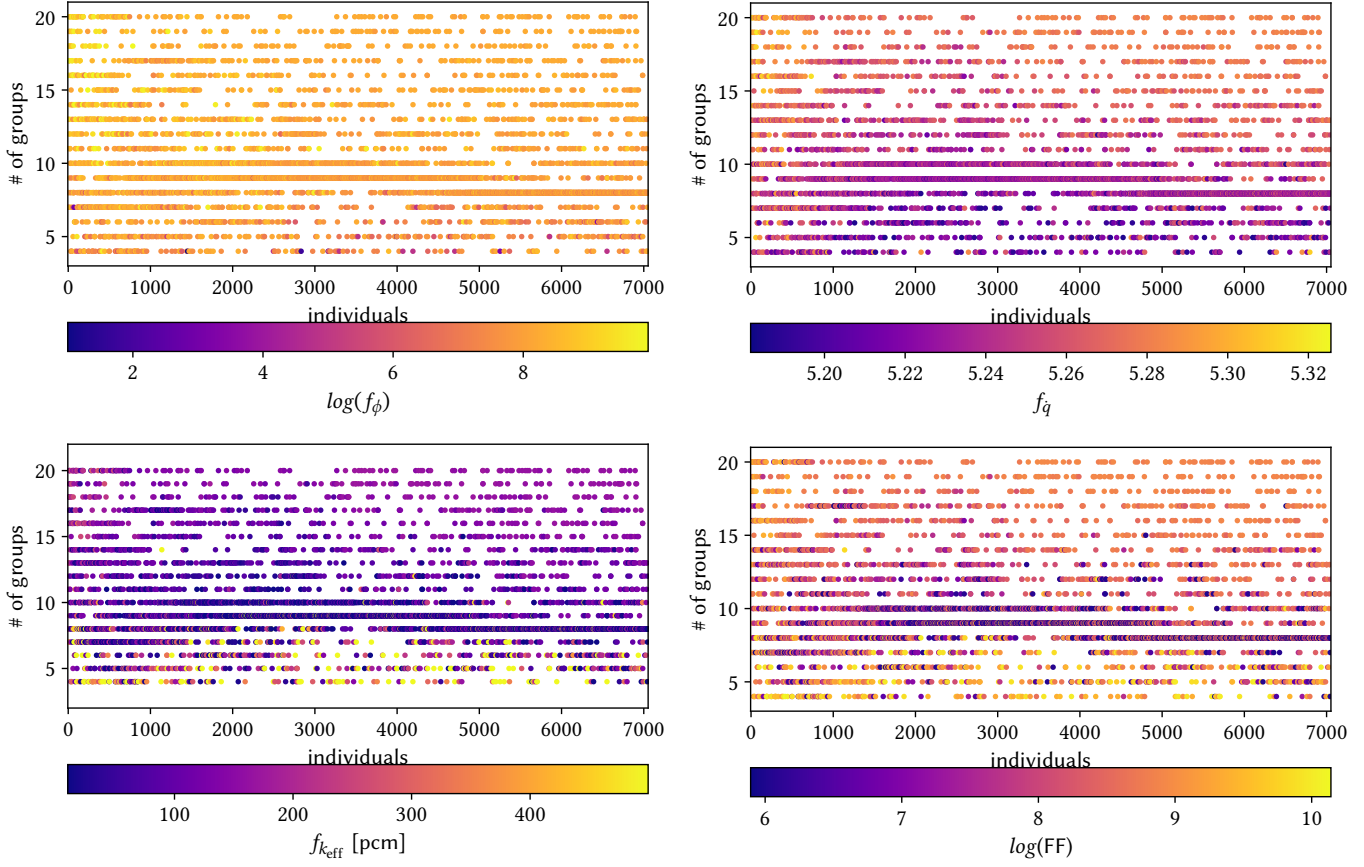


Fig. 17. Fitness functions for the flux (top left), k_{eff} (top right), power (bottom left) and their combination (bottom right) for the γ chromosome structure. The dots represent all the individuals from the first to the last one, providing information on the evolution of the number of groups and their FF.

for an alternative definition of the power FF, which is not further investigated in this work.

The evolution of the distribution of the number of groups, not reported for the sake of conciseness, is quite similar to the β case. On the contrary, the distribution of the group boundaries is quite sensitive to the karyotype choice, as it can be appreciated from fig. 19, which shows the group structures of each individual spawned in the optimization process, from the first to the last generation. When the β chromosome structure is used (left), most of the group boundaries is preserved in the gene pool throughout the evolution: this is usually a positive point, as the genetic variability is key to finding innovative solutions. However, this might also suggest a very slow, almost stagnant evolution: indeed, as we use a one-point crossover, it is very likely that individuals with a small number of groups experience effectively no change from XO, as the last elements of chromosome β_1 are not expressed. On the contrary, with the γ chromosome structure (right) many alleles disappear within the first 10 generations: one can see that the evolution is not stopped, as some alleles keep reappearing and even supplant previous ones. This is likely a consequence of the ranking process, that allows the evolution process to gather information faster. Reappearance of lost alleles is not impossible but, considering the way mutation is applied, only a few alleles (the "first losers" of the ranking) have an actual chance. A possible improvement of the γ representation would impact the mutation mechanism, giving an actual reappearance chance also to alleles with very low ranking. A characteristic issue of the energy group problem, which all representations have to deal with, is the proximity effect: the presence of an energy boundary influences nearby energy boundaries, i.e. the need of a specific energy boundary is not intrinsic but depends on the other positions. From this perspective, the γ representation works better than the other, because it mirrors the energy group boundaries ordering (unlike β), with XO not subverting it constantly (unlike α). None of the representations, however, can deal with the fact that two equally good solutions might have no groups in common. In fact, the existence of a single population forces the evolution process to choose a single path: innovative solutions must be very good from the beginning, gaining the right to survive and influence the evolution; partial improvements, on the contrary, are lost while mixing with the rest of the population. The introduction of niching and speciation mechanisms might help on this topic.

The best individuals found with the γ karyotype are reported in fig. 20 with the group-wise power and reproduction factor. With respect to the β case, the three group structures are almost




Fig. 18. *Difference (top) and relative difference (bottom) between the power radial distribution computed by Serpent (used as reference) and the one computed by FRENETIC for the best (left) and worst (right) individuals of the Boltzmann tribe.*




Fig. 19. Evolution of the same "tribe" (Boltzmann) with two different chromosome structures. The vertical dashed lines indicate the 61-group grid, the red dots indicate the individuals in each generation, the white stars indicate the best individuals in the various generations while the blue stars represent the best individual.

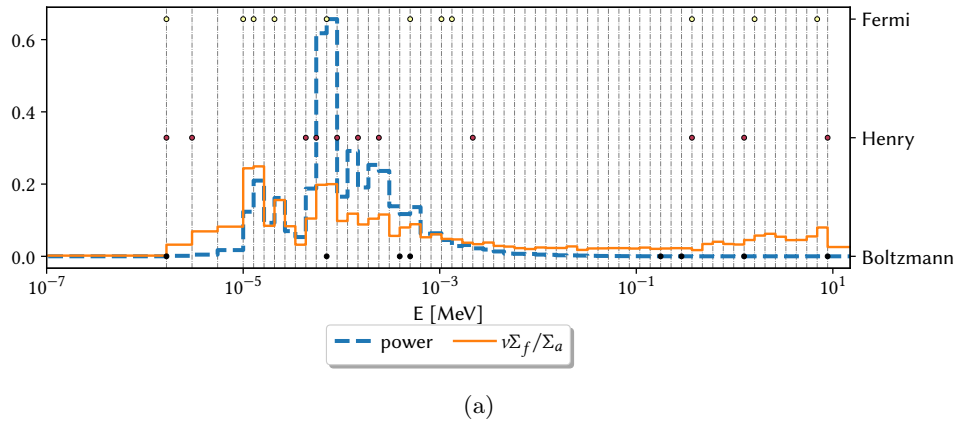


Fig. 20. Group-wise quantities and optimal group boundaries (dots) selected from the reference 61-group structure (dashed black lines) for the γ chromosome structure.

identical above 10^{-3} MeV, with the most remarkable differences below 10^{-3} MeV. Each individual has at least one thermal group.

TABLE II
Fitness functions of the best individuals generated in each case.

Tribe (seed)	Chromosome structure	n. of groups	$\log(f_\phi)$	f_P	f_k	FF
Boltzmann	α	19	7.7397	5.226	10.00	404.47
Fermi	α	16	7.6727	5.227	10.00	401.07
Henry	α	11	6.8994	5.242	10.00	361.66
Boltzmann	β	10	7.3083	5.234	10.00	382.50
Fermi	β	11	7.3374	5.229	10.00	383.64
Henry	β	8	7.6692	5.230	10.00	401.11
Boltzmann	γ	9	7.7136	5.229	10.00	403.38
Fermi	γ	12	8.2799	5.213	10.06	434.22
Henry	γ	12	7.7741	5.232	10.00	406.76

Table II reports the single FFs and their combination with eq. (12) for the best individuals found in the different tribes using the three chromosome structures discussed. The individuals minimising the combined FF turn out to be the ones spawned with the α karyotype, followed by the β ones. These results seem quite surprising, considering that the γ karyotype was shown to improve the convergence guaranteeing a fair and unbiased exploration of the solution space. However, it should be reminded that the whole genetic optimisation is an inherent random process: the adoption of a more efficient and reliable chromosome structure ensures a better exploration of the solution space but does not prevent the other ones to have a stroke of luck in the search for the optimal solution.

Looking at the results with more detail, it can be noticed that all individuals minimise the error on k_{eff} , which is consistent with fig. 17, while minimising both the flux and the power with the same group grid seems not possible. The case scoring the minimum FF is the Henry- α , which yields the minimum error on the flux but the maximum error on the power. Conversely, Fermi- γ minimises the error on the power but yields the maximum error on the flux with respect to the other individuals in the table. The group distributions of Henry- α (see fig. 13) and Fermi- γ (see fig. 20) helps to justify the performances of these two grids. In the first case there is one large thermal group covering the interval from $5 \cdot 10^{-5}$ MeV to 10^{-11} MeV, while in the Fermi- γ case there are about five groups in the same interval. The presence of the large thermal group helps to reduce the error on the flux to the detriment of the power, while employing five groups helps to

improve the accuracy on the power to the detriment of the flux. It is useful to remark that the nodal diffusion model employed by FRENETIC may not work properly with tiny groups, either for a lack of accuracy in the diffusion model or even for converging issues in the nodal iterative solver, justifying the poor accuracy in the Fermi- γ case.

The other cases produced with the γ karyotype score a FF that is quite close the Boltzmann- α and Fermi- α , but it should be noticed that the γ cases are featured by less groups, which means that their calculation should be generally faster. In this paper we did not take explicitly into account the computational performances of the code, but this aspect is somehow reflected by the accuracy of the nodal method: optically thin groups do not satisfy the assumptions of the diffusion theory, leading the solution to be inaccurate. Since group structures with a large number of groups are likely to have thin groups, the solutions with a small number of groups are more favoured. This aspect also impacts the computational performances of the calculations, since thin groups tend to slow down the convergence of the iterative process. On top of this, the number of degrees of freedom in multi-group diffusion scales as G^2 , where G is the number of groups. In order to explicitly take into account also the impact of the number of groups on the computational performances, a FF accounting for this additional constraint could be defined: in the case of roughly the same accuracy, solutions with a fairly large number of groups (e.g., the α cases) should be penalized with respect to solutions with a reduced number of equations (e.g., the γ cases).

Another reason for the good performances of the α individuals is due to the structure of the chromosome. Despite the number of groups is biased around 20, the sorted genes allow to preserve the good features as generation goes by, allowing to improve the individuals. The γ karyotype does not bias the number of groups and includes the group proximity effects, but it has to optimize two independent chromosomes and to explore an effectively wider sample space, resulting in a longer and more difficult optimization process. This aspect confirms that the selection of the karyotype is crucial for obtaining optimised but reliable solutions, and suggests that it may not be always possible to conceive an unbiased yet efficient karyotype.

III.B. Choice of the weight in the flux fitness function

The impact of the adjoint-based weight used in eq. (6) on the convergence of the GA can be appreciated by inspecting fig. 21, which shows on the left the convergence of the generation-wise

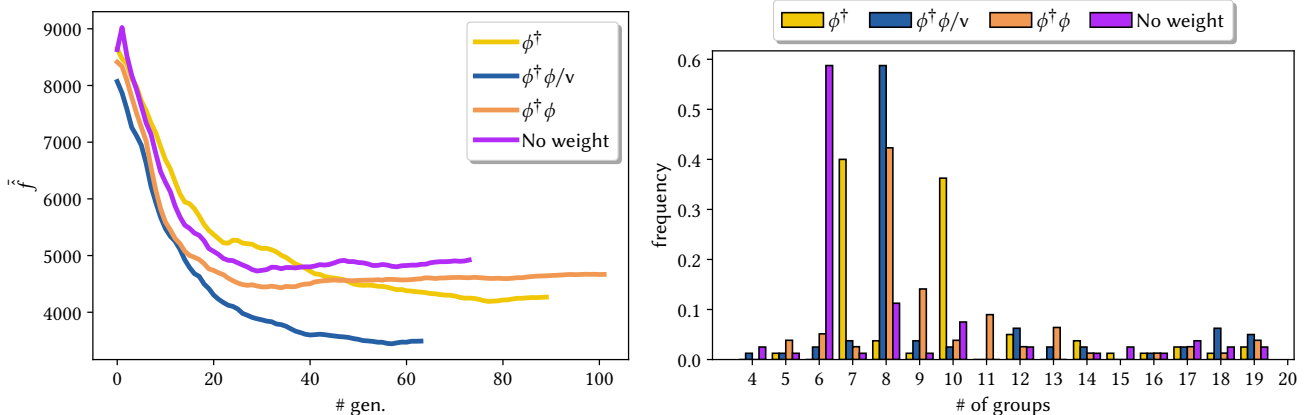


Fig. 21. *Convergence history of the average fitness function (left) and distribution of individuals per number of groups at the 30th generation (right) for the γ chromosome structure.*

averaged FF and on the right the distribution of the number of groups at the 30th generation. All the calculations considered for these graphs have the same seed and initial population, but are carried out using the different definitions of the weight presented in section II.E. Since the FF weighting factors are different, the curves are not fully comparable: in fact they differ from the very first generation, even if that generation is identical for all cases. This aspect makes it difficult to draw a rigorous comparison of the performances of the different weighting functions. The analysis should be carried out on a statistical basis, by running a significant number of independent simulations for each weight definition. Nevertheless, it can be appreciated that, at least qualitatively, the choice of the weight can have a strong impact on the convergence trend.

The justification for choosing a set of weights over another is a physical one: using the the total importance means that we are weighting the flux relative error in each phase-space cell according to the importance of a single particle in that cell times the number of particles populating it. Hence, a region featured by a relatively large importance but with a low neutron density (see, e.g., fig. 9) will contribute less to the overall error. In general, FFs with different weights will aim to different target: the choice depends on the objective we want to set. These evidences suggest the need for a deeper investigation of this aspect in a future work. The different distributions of the number of groups are a consequence of the different convergence trends of the four cases analysed.

The best individuals obtained with the various weights are represented in fig. 22, with the

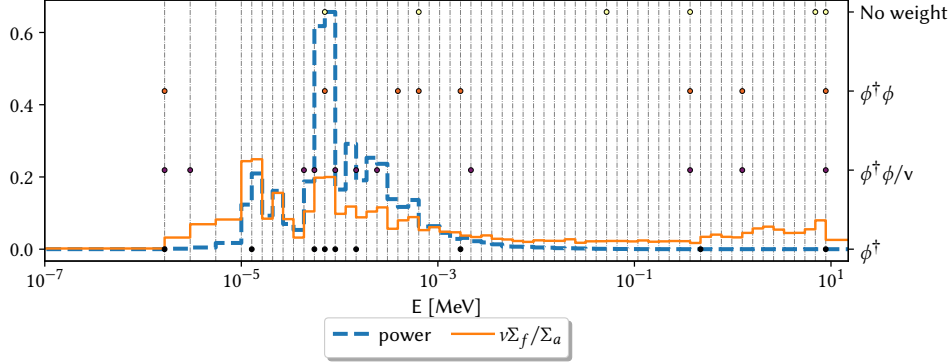


Fig. 22. Group-wise quantities and optimal group boundaries (dots) selected from the reference 61-group structure (dashed black lines) for the γ chromosome structure using different weight definition for the flux FF calculation.

group-wise power and reproduction factor. All group structures produced present a couple of groups in the fast region, where the reproduction factor, which is related to k_{eff} , has a certain gradient. Between $5 \cdot 10^{-3}$ and 10^{-1} MeV the reaction rates are quite uniform and do not require a high energy resolution, while between 10^{-5} and 10^{-3} MeV there are sharp variations, demanding a certain number of groups in order to get an accurate collapsing. The group boundaries obtained for the weight constructed with the contribution density share many groups with the case obtained considering the track length importance in the fast region, differing only in the epithermal region, while the case of no weighting presents a few number of large groups. The adjoint-weighted solution is similar to the non-weighted case in the fast region, presenting the larger number of groups in the epithermal region.

IV. CONCLUSIONS AND FUTURE PERSPECTIVES

This work presented a novel genetic algorithm for a heuristic optimisation of the multi-group energy grid structure used for full-core nodal calculations in Lead-cooled Fast Reactors. As other works in the literature, the optimisation is started considering a set of spatially homogenised group constants evaluated on a reference multi-group structure (61-group), which is appropriately reduced to a few-group grid. With respect to the literature, where the number of groups G is usually fixed *a priori*, our algorithm can select the optimal value of G within a user-defined range. This additional degree of freedom in the optimisation process allows the algorithm to fully explore the solution space, but poses the challenge of defining an unbiased and efficient representation of

the chromosomes used for mimicking the biological evolution process. The algorithm has been applied to the case study of a 3D model of the ALFRED core design, which has been modelled using the neutronics module of the FRENETIC multiphysics code. The code employs a multi-group nodal diffusion solver; due to the well known limitations of the diffusion theory [25], the accuracy and the computational performances of the code can be significantly reduced in presence of thin energy groups. This aspect makes the optimal selection of the number of groups non-trivial, as the solutions with a large number of groups may yield poorer results compared to coarser grids.

Another aspect that is analysed in the paper is the definition of a suitable combination of physics-driven fitness functions, which define the optimisation objective. In addition to the usual combination of fitness functions dedicated to the effective multiplication factor and to the flux, we introduced also a fitness function related to the power distribution, which is a key output quantity of full-core calculations. On top of this, a set of weighting definitions based on the adjoint flux have been proposed for the evaluation of the flux fitness function, aiming at focusing the evolution process to the most significant zones of the core and the energy spectrum.

The reference flux, power density and effective multiplication factor used in the calculation of the fitness functions have been computed with the Serpent Monte Carlo code, which also provided the set of spatially homogenised multi-group constants collapsed on the starting 61-group grid.

Three chromosomic representations, of increasing complexity, have been tested in the paper, together with different combinations of the weighted flux fitness functions. The results presented and discussed in the paper suggest that energy structures with a variable number of groups can benefit of a genetic representation with two chromosomes (γ): one haploid chromosome dedicated to the number of energy groups and a second polyploid chromosome to the position of the energy grid boundaries. The optimal solutions obtained in this way have been justified on a physical basis, by looking at some relevant figures of merit. The results also indicate that the total importance, i.e. the product of the neutron density times the adjoint flux, can be used to construct a meaningful weight for the flux fitness function.

Despite the promising results, some optimised grids contain more than 10 groups, which may yield unacceptable values of the computational cost and memory occupation for a transient multiphysics simulation. Hence, an additional fitness function accounting for the computational performances of the specific grid should be introduced in the future: in the case of similar scores

for the physics-based combination of fitness functions, this last fitness function would provide a practical indicator to choose the "best" individual. The results presented also suggest that a less strict definition of the power fitness function could improve the performances of the GA. It should be acknowledged that the results provided in the paper are quite sensitive to the random seed used to initialise the GA search. This indicates that the convergence of the algorithm could be improved by properly tuning its hyperparameters. Since our goal was to outline a general methodology for a practical case, systematic studies on these parameters have not been carried out. This activity will be carried out in a future work.

Another important aspect that was not addressed in the paper is that the optimal group structures proposed by the algorithm may not yield optimal results for other operational configurations of the reactor (e.g., featured by different operating temperatures or different positions of the control rods). In view of optimising the group structures for parametric and transient calculations, in future works the algorithm could be applied to different core configurations at the same time. In this context, it would be also useful to investigate the optimisation of the effective delayed neutron fraction and the effective lifetime, which are crucial integral parameters for an accurate time-dependent full-core simulation.

DATA AVAILABILITY

The complete datasets and scripts employed to execute and post-process the calculations presented in this paper are available in <https://doi.org/10.5281/zenodo.13300462> open access Zenodo repository.

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.

This work has been partially supported by the ENEN2plus project (HORIZON-EURATOM-2021-NRT-01-13 101061677) founded by the European Union.

REFERENCES

- [1] D. Knott and A. Yamamoto. In: *"Handbook of nuclear engineering II"*. Edited by ". G. C. ". "New York, NY": "Springer", 2010 (cited on page 3).
- [2] A. Till, N. Gibson, and T. Saller. "Optimizing Group Structures using Hierarchical Division". In: *Proceedings of the ANS International Conference PHYSOR 2022*. Pittsburgh, PA, 2022, pages 687–696 (cited on page 3).
- [3] T. Cormen. ["Introduction to Algorithms, Second Edition"](#). MIT Press, 2001 (cited on page 3).
- [4] J. Bang-Jensen, G. Gutin, and A. Yeo. ["When the greedy algorithm fails"](#). In: *Discrete Optimization* 1.2 (2004), pages 121–127 (cited on page 3).
- [5] V. Nair, T. Saller, A. Till, and N. Gibson. "Group Structure Selection with Random Forests". In: *Proceedings of the ANS International Conference PHYSOR 2022*. Pittsburgh, PA, 2022, pages 1186–1195 (cited on page 3).
- [6] T. G. Saller, V. Nair, A. Till, and N. Gibson. ["Using a Random Forest Model to Choose Optimized Group Structures"](#). In: *Nuclear Science and Engineering* 197.8 (2023), pages 2117–2135 (cited on page 3).
- [7] Di Nora, V. A. and Fridman, E. and Nikitin, E. and Bilodid, Y. and Mikityuk, K. ["Optimization of multi-group energy structures for diffusion analyses of sodium-cooled fast reactors assisted by simulated annealing—Part I: Methodology demonstration"](#). In: *Annals of Nuclear Energy* 155 (2021), page 108183 (cited on pages 3–4).
- [8] V. A. Di Nora, E. Fridman, E. Nikitin, Y. Bilodid, and K. Mikityuk. ["Optimization of multi-group energy structures for diffusion analyses of sodium-cooled fast reactors assisted by simulated annealing – Part II: Methodology application"](#). In: *Annals of Nuclear Energy* 163 (2021), page 108541 (cited on page 3).
- [9] O. Fasina, T. Saller, and N. Haven. "Particle Swarm Optimization for Group Structure Optimization for Radiotherapy Shielding". In: *Proceedings of the ANS International Conference PHYSOR 2022*. Pittsburgh, PA, 2022, pages 1340–1350 (cited on page 3).

- [10] R. E. Alcouffe, R. Baker, J. Dahl, S. Turner, and R. Ward. *PARTISN: A Time-Dependent, Parallel Neutral Particle Transport Code System*. Technical report LA-UR-08-07258. Los Alamos National Laboratory, 2009 (cited on page 3).
- [11] M. Massone, F. Gabrielli, and A. Rineiski. “A genetic algorithm for multigroup energy structure search”. English. In: *Annals of Nuclear Energy* 105.C (2017), pages 369–387 (cited on pages 4, 23).
- [12] M. Massone, F. Gabrielli, and A. Rineiski. “SIMMER extension for multigroup energy structure search using genetic algorithm with different fitness functions”. In: *Nuclear Engineering and Technology* 49.6 (2017), pages 1250–1258 (cited on pages 4, 17).
- [13] M. Massone, N. Abrate, G. F. Nallo, S. Dulla, P. Ravetto, and 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 (cited on pages 4, 17).
- [14] Y. Tobita, S. Kondo, H. Yamano, K. Morita, W. Maschek, P. Coste, and T. Cadiou. “The Development of SIMMER-III, An Advanced Computer Program for LMFR Safety Analysis, and Its Application to Sodium Experiments”. In: *Nuclear Technology* 153.3 (2006), pages 245–255 (cited on page 4).
- [15] M. Massone, N. Abrate, G. F. Nallo, D. Valerio, S. Dulla, and P. Ravetto. “Code-to-code SIMMER/FRENETIC comparison for the neutronic simulation of lead-cooled fast reactors”. In: *Annals of Nuclear Energy* 174 (2022), page 109124 (cited on pages 4, 8).
- [16] J. Leppänen, M. Pusa, and E. Fridman. “Overview of methodology for spatial homogenization in the Serpent 2 Monte Carlo code”. In: *Annals of Nuclear Energy* 96 (2016), pages 126–136 (cited on page 5).
- [17] G. Grasso, C. Petrovich, D. Mattioli, C. Artioli, P. Sciora, D. Gugiu, G. Bandini, E. Bubelis, and K. Mikityuk. “The core design of ALFRED, a demonstrator for the European lead-cooled reactors”. In: *Nuclear Engineering and Design* 278 (2014), pages 287–301 (cited on page 6).
- [18] R. Bonifetto, S. Dulla, P. Ravetto, L. S. Richard, and R. Zanino. “A full-core coupled neutronic/thermal-hydraulic code for the modeling of lead-cooled nuclear fast reactors”. In: *Nuclear Engineering and Design* 261 (2013), pages 85–94 (cited on page 9).

- [19] O. Tal, E. Israeli, P. Ravetto, and E. Gilad. “[The adjoint problem as physical heuristic for loading pattern optimization](#)”. In: *Annals of Nuclear Energy* 134 (2019), pages 226–234 (cited on pages 10, 18).
- [20] N. Siddique. “[Intelligent Control: A Hybrid Approach Based on Fuzzy Logic, Neural Networks and Genetic Algorithms](#)”. Studies in Computational Intelligence. Springer International Publishing, 2013 (cited on page 11).
- [21] T. Blickle and L. Thiele. “[A Comparison of Selection Schemes Used in Evolutionary Algorithms](#)”. In: *Evolutionary Computation* 4.4 (Feb. 2008), pages 361–394 (cited on page 12).
- [22] D. E. Goldberg and K. Deb. “[A Comparative Analysis of Selection Schemes Used in Genetic Algorithms](#)”. In: edited by G. J. Rawlins. Volume 1. Foundations of Genetic Algorithms. Elsevier, 1991, pages 69–93 (cited on page 12).
- [23] M. Srinivas and L. Patnaik. “[Adaptive probabilities of crossover and mutation in genetic algorithms](#)”. In: *IEEE Transactions on Systems, Man, and Cybernetics* 24.4 (1994), pages 656–667 (cited on page 12).
- [24] B. McGinley, J. Maher, C. O’Riordan, and F. Morgan. “[Maintaining Healthy Population Diversity Using Adaptive Crossover, Mutation, and Selection](#)”. In: *IEEE Transactions on Evolutionary Computation* 15.5 (2011), pages 692–714 (cited on page 12).
- [25] G. Bell and S. Glasstone. “[Nuclear Reactor Theory](#)”. Van Nostrand Reinhold, 1970 (cited on pages 18, 41).
- [26] M. Williams. “[Generalized contributon response theory](#)”. In: *Nuclear Science and Engineering* 108.4 (1991), pages 355–383 (cited on page 18).