frontiers

Accelerating Engineering Design of Breeder Blankets with Parametric Optimisation and Sequential Learning

Luke Humphrey,[1,*] **Helen Brooks,**[1] **Siddharth Mungale,**[1] **Andrew Davis,**[1] **David Foster**[1]

[1] *United Kingdom Atomic Energy Authority, Culham Campus, Abingdon, Oxfordshire, OX14 3DB, UK*

Correspondence*:
Luke Humphrey
luke.humphrey@ukaea.uk

ABSTRACT

The competing requirements of fusion breeder blankets and the high dimensionality of their design space necessitate a systematic treatment to map the variations in performance against given objective metrics and to understand the operational envelope. In this endeavour, a digital engineering pipeline for design evaluation and optimisation has been developed. The tools involved are: Hypnos for parametric breeder blanket geometry instantiation, OpenMC for neutronics analysis, MOOSE for thermal hydraulics analysis, and SLEDO for design space sampling, sensitivity analysis and optimisation. An optimisation of the baseline design for a solid ceramic breeder mock-up that is relevant to the Lithium Breeding Tritium Innovation (LIBRTI) program is performed. Two optimisation studies are performed, the first involving only neutronics, while the second included the impact of thermal hydraulics. The figures of merit are taken to be the Tritium Breeding Ratio (TBR) and the pressure drop of the outer coolant (combined in a weighted sum for the second analysis). In the first study, for the same acquisition function (taken to be Expected Improvement) two different values are selected for the hyperparameter that controls the trade-off between exploration and exploitation. In the second study, with the inclusion of thermal hydraulics a larger parameter space was explored to assess the performance of the method in a high-dimensionality setting. In both cases, the selected figures of merit were improved over the baseline design. Finally, we discuss extensions of the procedure to include a more thorough multi-physics analysis, and a more sophisticated treatment of multiple objectives.

Keywords: Breeder blankets, Neutronics, Optimisation, Sequential Learning, OpenMC, MOOSE, SLEDO, Hypnos

1 INTRODUCTION

Breeder blankets encapsulate the challenges faced more broadly by in-vessel components within fusion pilot plants, namely: high heat fluxes, ionising radiation, and electromagnetic and gravitational loading. Blanket systems therefore constitute a worthwhile focal point for the applications of novel design methodologies. Competition arises due to these components having multiple requirements, whose fulfilment may result in diverse and potentially opposing changes, ultimately implying trade-offs in spatial allocation of functional materials.

To be specific, in the context of breeder blankets, we can consider the following requirements:

- The breeder blanket must shield the rest of the device from damaging irradiation. The device must have structural integrity, and any structural materials must be resilient to high neutron fluxes and temperatures.
- The breeder blanket must be actively cooled to ensure that safe operational temperatures are maintained and simultaneously that deposited heat is transported for the ultimate purpose of power generation.
- The breeder blanket must breed sufficient tritium. For commercial viability and sustainability of fusion as an energy source, the tritium fuel cycle must be self-sustaining: for every neutron emitted there must occur one or more subsequent nuclear reactions to produce tritium.

These requirements each prefer conflicting design choices from which trade-offs naturally emerge. Designs favouring shielding and structural considerations would increase the volume of non-breeder, non-coolant materials, which nonetheless require cooling and which absorb neutrons without contributing to tritium production. Alternatively, designs favouring effective thermal management would reduce the volume of material in need of cooling (including both structural and breeding material); these designs would also seek to reduce required pumping power by having fewer channels across which to distribute the coolant. Finally, designs favouring high tritium breeding capability would minimise neutron loss by reducing the volume of structural and coolant materials in favour of breeder and multiplier. In addition, these designs would distribute materials as homogeneously as possible, necessitating an increase in the number of coolant channels, a preference which is discussed further in section 3.

Design challenges of this nature can be posed mathematically, commonly referred to as Multi Discipinary Analysis Optimisation (MDAO). Suppose our design point is defined by a finite number of configurable parameters $\mathbf{x} = \{x_1, x_2, \ldots, x_n\}$. The ranges of those parameters define some design space $\mathbf{\Omega}$. Figure(s) of merit, for example the tritium breeding ratio (TBR), can be evaluated as a function of the design point, $f(\mathbf{x})$. Similarly, any constraints may be expressed as $g(\mathbf{x}) > 0$, which is equivalent to restricting to one or more subdomains $\mathbf{\Omega}'$ on $\mathbf{\Omega}$. The optimisation problem then is simply to maximise $f(\mathbf{x})$ on $\mathbf{\Omega}'$.

Given that the dimensionality of such problems is typically large, it is hard to intuitively comprehend the subspace within which optimal solutions exist. Indeed, improperly posed, there is no guarantee of the existence of any solution at all; such a scenario would arise if constraints are mutually exclusive and $\mathbf{\Omega}'$ becomes a null space. Alternatively, optima may be unstable to small perturbations in the design space; this scenario translates to rapid degradation in performance due to (for example) natural variation in the manufacturing process or small changes in requirements specification.

Both scenarios outlined above are best avoided: they would likely necessitate complete design overhaul, at significant expense. Instead, it is more efficient to explore the design space systematically, as well as evaluating sensitivities to parameters away from any optima. Here the dimensionality of the problem again poses a challenge; a brute force exploration may become computationally prohibitive. However, if a surrogate model of the figure of merit can be constructed on-the-fly (and progressively refined with more model evaluations), an understanding of the design space and any associated trade-offs in parameters can be attained more rapidly.

A schematic of such a workflow is shown in fig. 1. Given the problem definition, candidate designs are generated in an initial sample, and a figure of merit is evaluated using physical models. The initial pool of results will inform a more targeted sampling in subsequent cycles of this *outer loop*. After a sufficient number of iterations, the optimal configuration(s) may be identified and their robustness assessed. This process, however, is predicated on the existence of a so-called *design-point evaluator*. In other words, it must be possible to (i) construct a design from some parameterisation $\mathbf{x}$, (ii) stipulate the models that

govern the behaviour of physical variables, and (iii) simulate these numerically, computing in turn the figure of merit from those variables. Only with such an implementation in place is it possible to effectively leverage optimisation algorithms in order to accelerate the engineering design process.

As part of the Digital workstream of the UK Atomic Energy Authority's "Lithium Breeding Tritium Innovation" program (LIBRTI) [1] an implementation of such a digital engineering workflow has been developed. While eventually the intention is to target improving designs of breeder blankets, in the short term the focus is upon the development of breeder mock-up experiments. These can be tested at the planned LIBRTI facility which will host a 14 MeV neutron source and will enable qualification of both tritium breeding technology and predictive modelling capability.

The outline of the paper is as follows: in section 2 we describe the full tool-chain employed to implement the workflow of fig. 1. In section 2.1, a novel software tool, Hypnos [2], for parametric blanket geometry preparation is presented, and the reference breeder mock-up design is described. In section 2.2 and section 2.3 the analysis set-up for neutronics with OpenMC [3, 4] and thermal hydraulics with MOOSE [5, 6, 7] is described. Although the complex nature of the breeder blanket should necessitate a tightly-coupled treatment of physics, in this initial demonstration of the optimisation set-up these analyses are treated separately, and are used to define two separate figures of merit. In section 2.4 we review the software, SLEDO [8] , which orchestrates the optimisation procedure, along with some relevant new features.

Having established the methodology, a number of different optimisations are performed, with results presented and discussed in section 3. The first campaign considers neutronics in isolation, and includes an investigation into the impact of varying the search strategy through a hyperparameter. In the second campaign, both neutronics and thermal hydraulics analyses are performed, permitting an investigation into the trade-off between competing objectives. Finally the outlook and conclusions are presented in section 4 and section 5 respectively.

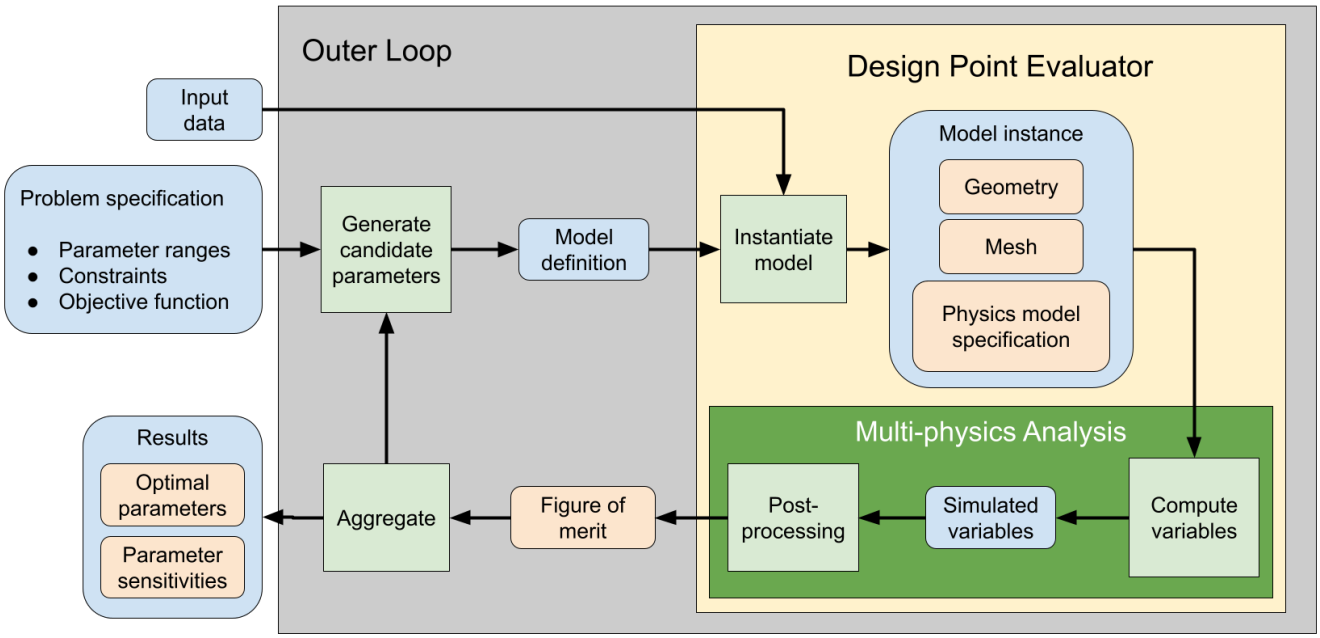

Figure 1. Generalised digital engineering workflow for optimisation and uncertainty quantification.

2 TOOL-CHAIN AND METHODOLOGY

2.1 Parametric Breeder Geometry Instantiation with Hypnos

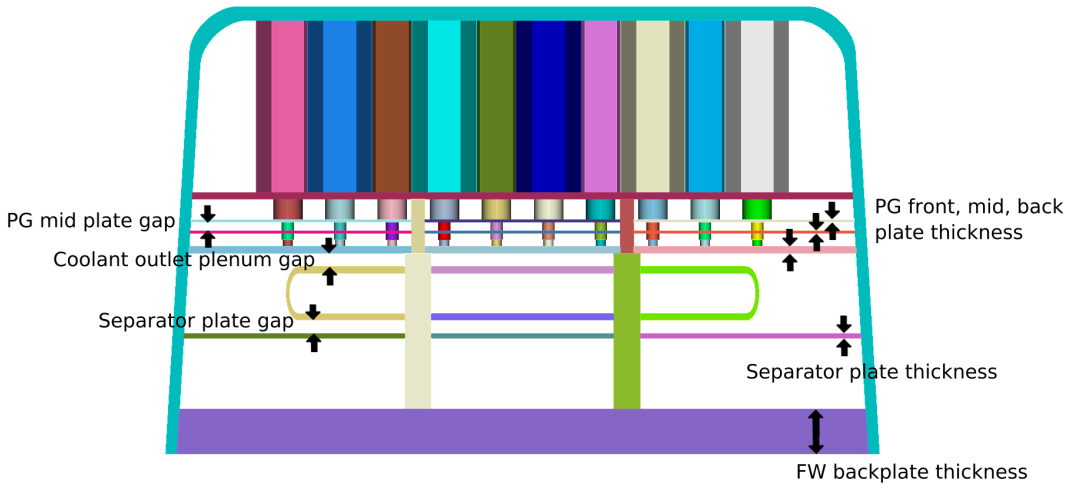

Figure 2. Illustrative view of HCPB blanket parametric geometry distributed with Hypnos. Several parameters are labelled, with the full set that are available described in the online documentation.

Hypnos [2] is an open-source parametric geometry engine used to model fusion-relevant components. Written in Python, it uses Coreform Cubit's [9] API to instantiate CAD from geometric parameters and generate meshes in multiple file formats, including surface mesh data such as DAGMC (.h5m) [10] for neutron transport and unstructured meshes such as ExodusII [11]. Coreform Cubit metadata like sideset, block, and material names can be assigned and queried by component, material, and interface.

Distributed as part of the package is an initial implementation of a Helium-Cooled Pebble Bed (HCPB) breeder blanket [12] which includes individual breeder pins, and channels and manifolds for the coolant and purge gas to flow through the system, as shown in Figure fig. 2. For such existing geometries, Hypnos may be run as a standalone executable from the command line, with component parameter values supplied by JSON files, and other behaviour controlled with a configuration file. Further detail is available in the online documentation [1].

Furthermore, Hypnos is designed to be modular and extensible, with component classes implemented hierarchically to reflect ontological relations. Users can implement custom classes for creation and assembly of these components, inheriting from base classes that wrap Coreform Cubit functionality. This may be achieved using Hypnos as a Python library, which also makes available a suite of functions to help manipulate and create geometries.

The reference model used for the optimisation study in section 3 is the LIBRTI straw-person conceptual design for a solid breeder mock-up experiment, originally presented in [1]. A parameterisation of the geometry was implemented using Hypnos, and is shown in fig. 3. The structure of a single pin-cell consists of hollow cylinders made of steel. The innermost cylinder is a channel for coolant (inner coolant channel). This is surrounded by a chamber containing the solid ceramic breeder pebbles, which additionally has an inlet and outlet at either end to direct purge gas through to extract produced tritium. This is further

[1] Available at https://aurora-multiphysics.github.io/hypnos/

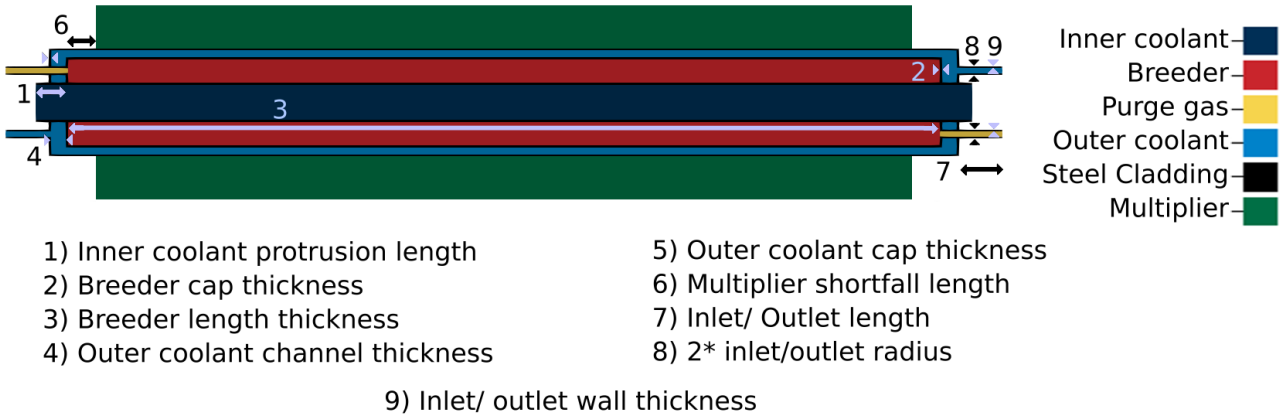

Figure 3a. Axial pin-cell parameterisation.

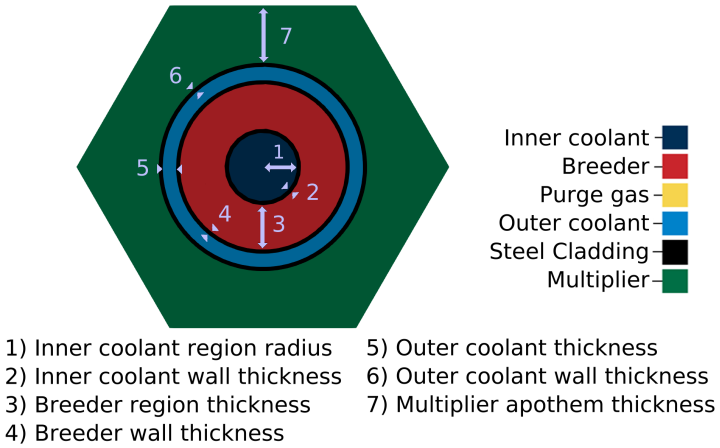

Figure 3b. Radial pin-cell parameterisation.

Figure 3. Parameterisation of the LIBRTI pin-cell geometry in Hypnos.

surrounded by another channel through which coolant is passed (outer coolant channel). Surrounding the outer coolant is a hexagonal prism of multiplier material. An arrangement of pins can be tessellated around an axially-aligned neutron source, as shown in figure fig. 4. Finally the assembly is surrounded by a volume acting as a neutron reflector.

2.2 Neutronics Modelling with OpenMC

OpenMC [3, 4] is selected for neutronics analysis. Not only is this a scalable open-source Monte Carlo neutral particle transport code, but also it has an integration within the MOOSE (Multi-physics Object Oriented Simulation Environment) framework [5, 6, 7] via the application Cardinal [13]. As noted earlier, the dependency of blanket requirements on interconnected physics domains requires a coupled multi-physics treatment. While in this initial demonstration of our workflow we have restricted to a neutronics analysis, we intend to extend to multiple physics domains in future, and therefore it is prudent to have a strategy for this extension.

The material definitions assumed for the study are as follows. The structural material is taken to be 304 Stainless Steel. The neutron multiplier material is lead. The tritium breeding material is KALOS

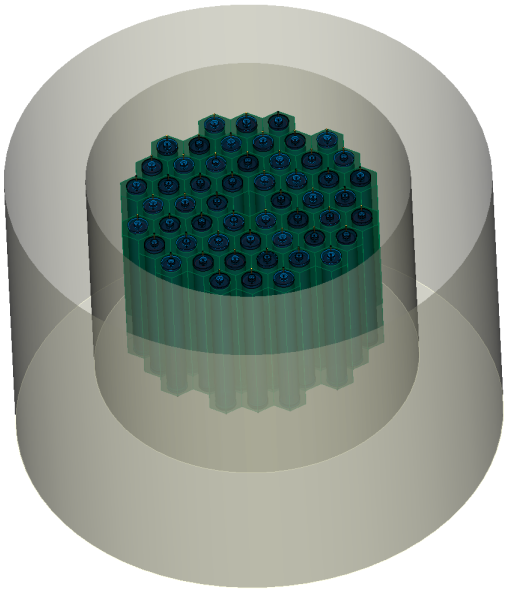

Figure 4. Visualisation of the full CAD geometry for the LIBRTI pin-cell assembly generated by Hypnos.

(KArlsruhe Lithium OrthoSilicate) ceramic pebbles [14], which contains a mixture of Li_4SiO_4 and Li_2Ti, homogenised over the pebbles.

The coolant was initially selected to be water, which was used for the standalone neutronics optimisation; in the second optimisation the coolant was selected to be helium doped with 0.1% weight hydrogen for consistency with the thermal hydraulics analysis. In the first study, the temperature was set at 296 K (= 22.85C); in the second study the temperature was set at 873.15 K (= 600C) again for consistency with the thermal hydraulics analysis. Finally, the pin-cell assembly is housed within a neutron reflector (as shown in fig. 4), with the material selected to be graphite.

It should be noted that the geometry is considered in isolation to its surroundings, namely the impact of the room is not considered and thus the results should not be considered indicative of actual blanket designs. Notwithstanding this limitation, provided the results are interpreted in relative sense (rather than as an absolute prediction of TBR) the analysis is nonetheless informative in a comparative study.

OpenMC was run with a fixed source, using continuous energy treatment of cross sections and with photon transport enabled. Cross sections were taken from the ENDF/B-VIII.0 library [15]. The neutron source is an approximation of that found in commercially-available neutron generators such as [16], it is described as a line source; having a spatial variation in intensity along the axial z-direction. The energy distribution corresponds to mono-energetic 14 MeV neutrons. The resultant neutron flux distribution, and tritium production rate are scored on Cartesian grids as shown in fig. 5. Here the mesh resolution used was 1000 subdivisions in each direction; the number of histories was 10 batches of 100 million particles. Radially, the peaks of tritium-production are observed at the innermost and outermost points in the assembly, corresponding to breeder material closest to the source and closest to the reflector's inner wall respectively.

In the study to be presented in section 3 the figure of merit used is the global TBR. This is computed in OpenMC using the `H3-production` score, tallied over the entire geometry; this returns the total number of tritium nuclei produced per source neutron. Calculating the global TBR does not rely on tracking the spatial distribution of tritium production, so a coarser mesh was used (and used only for the

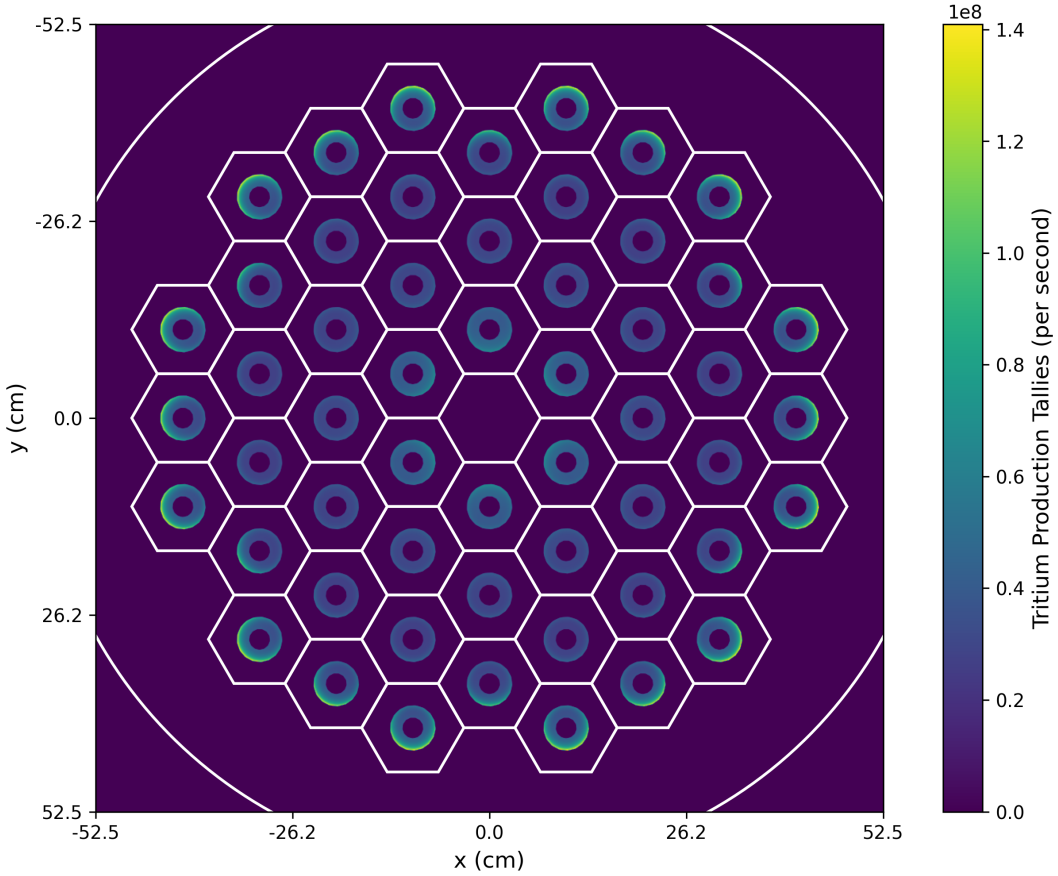

Figure 5. Spatial distribution of tritium production in the x-y plane for the baseline design using helium coolant with temperature 873.15 K. Tallies were scored by OpenMC on a regular mesh with a summation over contributions along the z-direction. The edge of each pin-cell and the reflector's inner wall are overlaid in white.

purpose of interpreting results). Attaining a suitable level of convergence (standard deviation for global TBR on the order 10^{-4}) required fewer histories, namely 5 batches of 10 million particles.

The global TBR for the baseline design was found to be 0.723 for the first study (using water coolant at 296 K) and 0.474 for the second study (using helium coolant at 873.15 K). We note here that as this application corresponds to an experiment, while it is desirable to maximise the TBR, as this would ultimately improve the signal to noise ratio of any measurements, it is not strictly a requirement to attain a target TBR in excess of unity (as would be required for a self-sustaining fuel cycle). Indeed, this would require higher coverage in solid angle around the neutron source, whereas this conceptual experimental design is intended to enable modularity.

2.3 Thermal Hydraulics

MOOSE is employed to perform thermal hydraulics modelling of the coolant. The approach uses the Compressible Euler Flow Model, and whose was original implementation for RELAP7 is described in

[17]. The governing equations, applicable to one-dimensional single-phase flow, are denoted by:

$$\frac{\partial(\rho A)}{\partial t} + \frac{\partial(\rho A u)}{\partial x} = 0 \tag{1}$$

$$\frac{\partial(\rho u A)}{\partial t} + \frac{\partial\left(\rho u^2 + p\right) A}{\partial x} = p\frac{\partial A}{\partial x} - FA + \rho g A \tag{2}$$

$$\frac{\partial(\rho E A)}{\partial t} + \frac{\partial u\left(\rho E + p\right) A}{\partial x} = \rho u g A + q''' A \tag{3}$$

These equations correspond to balance of mass, momentum and energy respectively. Here, A is the cross-sectional area, ρ is the density, u is the axial velocity, E is the specific total energy, p is the pressure, F is the viscous drag force, g is the gravitational acceleration projected along the axial direction, and q''' is the volumetric heat source.

In the baseline model, a given breeder pin is cooled by an annulus of helium coolant, which is treated as an ideal gas with approximately constant viscosity, thermal conductivity and specific heat. The pressure drop of the coolant flowing through the annulus is computed assuming a fixed operational temperature, mass flow rate and outlet pressure. This may be treated using a single thermal hydraulic flow channel of fixed length and hydraulic diameter given by $D_H = D_o - D_i$, that is, the difference between the inner and outer diameters. A Churchill friction factor is assumed to apply at the walls of the channel. The assumed values of all operational and helium material input parameters are summarised in table 1.

Whilst the approach described here constitutes a fairly simplistic treatment, in the context of the optimisation study to be presented in section 3, it should nevertheless introduce some penalty that prevents the minimisation of the coolant volume. Even restricting to 1D thermal hydraulics, there are a number of improvements which could be considered in future; these are discussed in section 4.

Parameter name	Symbol	Value / Unit	Reference
Temperature	T	600 C	[1]
Outlet Pressure	p	8 MPa	[12]
Inlet mass flow rate	$\dot{m}$	0.045 kg/s	[18]
Surface roughness	ε	300 μm	[18]
Ratio of specific heat	γ	1.67	[19]
Thermal conductivity	k	$2.682 \times 10^{-3}(1 + 1.123 \times 10^{-3}p)T^{0.71(1-2\times10^{-4}p)}$ W/(mK)	[19]
Dynamic viscosity	μ	$3.674 \times 10^{-7}T^{0.7}$ kg/(m s)	[19]

Table 1. Operating conditions and helium material properties used in thermal hydraulics calculation

2.4 Optimisation with SLEDO

SLEDO (Sequential Learning Engineering Design Optimiser) [8] is a software tool which performs Bayesian optimisation, selecting candidate designs from a parametric search space and evaluating them by deployment of a "black-box" function. Developed with deployment of computationally expensive multi-physics simulation workflows in mind, the scalable parameter search framework RayTune [20] was recently integrated to handle HPC deployment. A range of Bayesian optimisation libraries are available to drive candidate selection; for this study RayTune's implementation of the BayesOpt package [21] was used to provide functionality for the search algorithm.

The procedure employed by SLEDO is as follows:

1. (Optional) An initial population of candidate designs is selected quasi-randomly by a sample plan run over a bounded search space $\mathbf{x}$.
2. Candidate designs are evaluated to arrive at a scalar-valued figure of merit $f(\mathbf{x})$.
3. A surrogate model is trained, from which posteriors for the mean $\bar{f}(\mathbf{x})$ and standard deviation $\sigma(\mathbf{x})$ of f are obtained.
4. An acquisition function $A(\mathbf{x})$ is computed and maximised to generate subsequent candidates for evaluation.
5. Candidate designs are generated and evaluated until a fixed number of trials have concluded.

The role of the surrogate model is to use the information gained from all previous trials to predict the performance of a proposed design candidate. This statistical model should be relatively cheap to train and evaluate, such that each in iteration it may be updated and the acquisition function run over the entire search space. In this study, a Gaussian Process (GP) is used; a generalisation of a Gaussian distribution to arbitrary dimensions [22].

The role of the acquisition function is to generate candidates in a sequence which most efficiently converges on an optimised design. To achieve this, the function must be able to balance exploitation (sampling where the design is predicted to perform well) and exploration (sampling where the design's performance is uncertain). Many acquisition functions are available; in this study the Expected Improvement (EI) is utilised, defined as follows:

$$A(\mathbf{x}) = \left(\bar{f}(\mathbf{x}) - f(\mathbf{x}^+) - \xi\right)\Phi(Z) + \sigma(\mathbf{x})\phi(Z) \tag{4}$$

$$Z = \bar{f}(\mathbf{x}) - f(\mathbf{x}^+) - \xi \tag{5}$$

Here $\mathbf{x}^+$ is the current set of best design parameters; Φ and ϕ are the cumulative distribution function and probability density function of the normal distribution respectively; ξ is a tuneable hyperparameter which balances the trade-off between exploitation and exploration.

While optional, the role of the initial sample plan is provide an initial set of candidate designs which achieve good coverage across $\mathbf{x}$. This provides a reasonable prior to the surrogate model at the onset of the Bayesian optimisation sequence. Without initial data, $A(\mathbf{x})$ returns a flat distribution, so the first candidate would be randomly selected. For the following iteration, only the exploration term would vary across $\mathbf{x}$, so the most distant point would be be selected. Recognising that exploration of the search space is most valuable in these early iterations, while the identification of local maxima is preferable to convergence on the global maximum, one advantage of using a sample plan is that it allows the acquisition function itself to be tuned more in favour of exploitation. Another advantage is that the initial candidates may be evaluated in parallel, since they do not depend on each other. This is particularly useful for applications where time-to-solution is valuable, though for simplicity of deployment, designs were evaluated sequentially in this study. The sample plan used in this study is a Sobol sequence [23], using the SciPy implementation [24], from which a number of samples must be drawn to achieve a balanced spread through the search space. The ideal number of Sobol samples to be drawn before switching to Bayesian optimisation, which is set to be a power of two $n = 2^m$, depends on the problem's dimensionality and expected degree of non-linearity; in this study, the number was varied between 8 and 32.

3 RESULTS AND DISCUSSION

3.1 Neutronics optimisation

3.1.1 Definition of optimisation campaigns

As an initial test of the overall pipeline, a standalone neutronics analysis was performed, and the objective function was set to maximise a single figure of merit, namely the TBR. A disadvantage of geometric optimisation performed in isolation from other physics is that components whose function does not pertain to increasing this metric would be minimised unnaturally. Therefore in this initial analysis, only two parameters whose function directly pertains to tritium breeding were allowed to vary; these, along with their ranges are shown in table 2. In both cases the upper and lower limits are obtained by multiplying the reference value used in the baseline design (also shown in the table) by a factor of $\frac{1}{4}$ and 4 respectively. Both of these parameters control radial thicknesses within each pin, and are shown in fig. 3b. In addition the total number of pins in the assembly is not fixed, but is computed as the maximum number that may be fully contained within the inner radius of the reflector (additionally stipulating a minimum gap of 75 mm between this and the outermost pin).

The surrogate model employed was a Gaussian Process, while the acquisition function was Expected Improvement, as defined by eq. (4). Two optimisation campaigns were performed, which employed different sampling strategies; these are summarised in table 3. Each campaign used a fixed simulation budget of 100 trials but differed in how candidate designs were selected. In the first campaign, a large value of ξ was used, corresponding to a preference for exploration over exploitation, and a relatively high proportion of Sobol samples were generated. In the second campaign, a lower value of ξ was used, corresponding to increased preference for exploitation, and a relatively low proportion of Sobol samples were generated.

Parameter name	**Lower Bound**	**Upper Bound**	**Reference Value**
Breeder region thickness	4.0625 mm	65.0 mm	16.25 mm
Multiplier apothem thickness	5.225 mm	83.6 mm	20.9 mm

Table 2. Parameters varied and their bounds for the neutronics optimisation.

Strategy	ξ	**Sobol sample size** n
1	2.5	32
2	0.01	8

Table 3. Sampling strategies employed in neutronics-only optimisation.

3.1.2 Results

Comparing the strategies in fig. 6a and fig. 6b, for strategy 1 the TBR of the baseline design was exceeded within the initial Sobol sequence, but not significantly improved upon during the Bayesian optimisation sequence, while for strategy 2 the reference TBR was exceeded only a few trials into the Bayesian optimisation sequence. The best design found by strategy 1 achieved a TBR of 0.791, which was further improved to 0.802 by strategy 2. Figure 7 shows the the distribution of selected candidates from each strategy. Strategy 1 features a good spread of points throughout the space but fails to cluster in the performant region, while strategy 2 shows good clustering with sufficient spread throughout the space; this difference corresponds to the choice of hyperparameter ξ with the low-value in strategy 2 reducing the

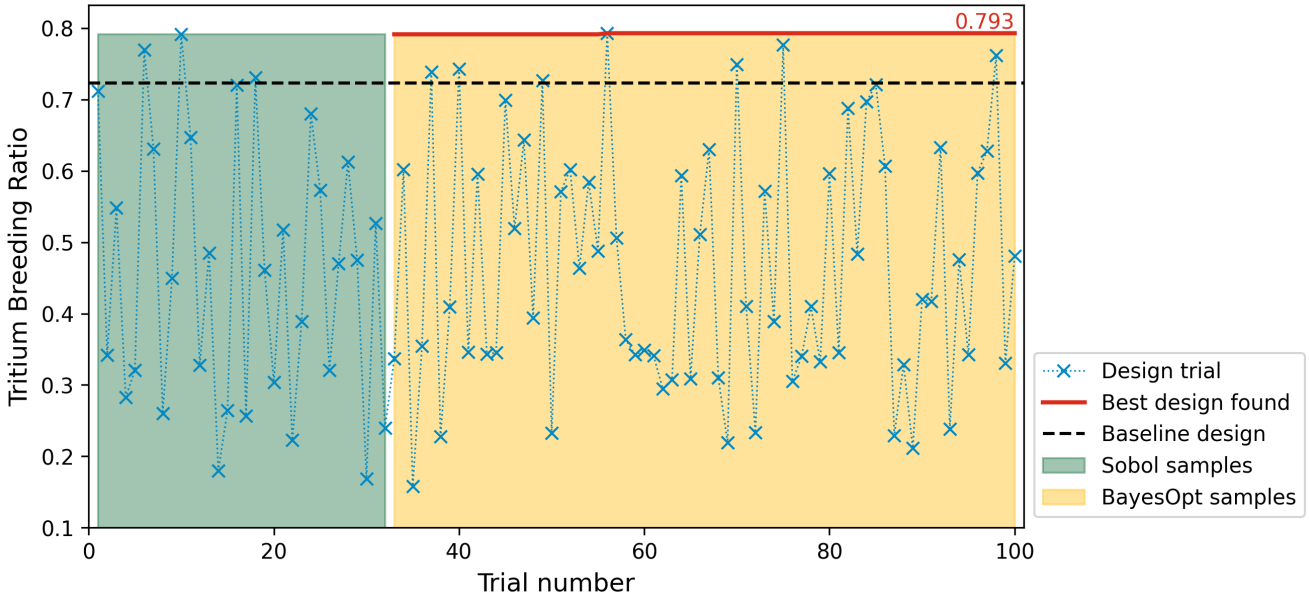

Figure 6a. Figure of merit for strategy 1 with 32 Sobol samples and $\xi = 2.5$.

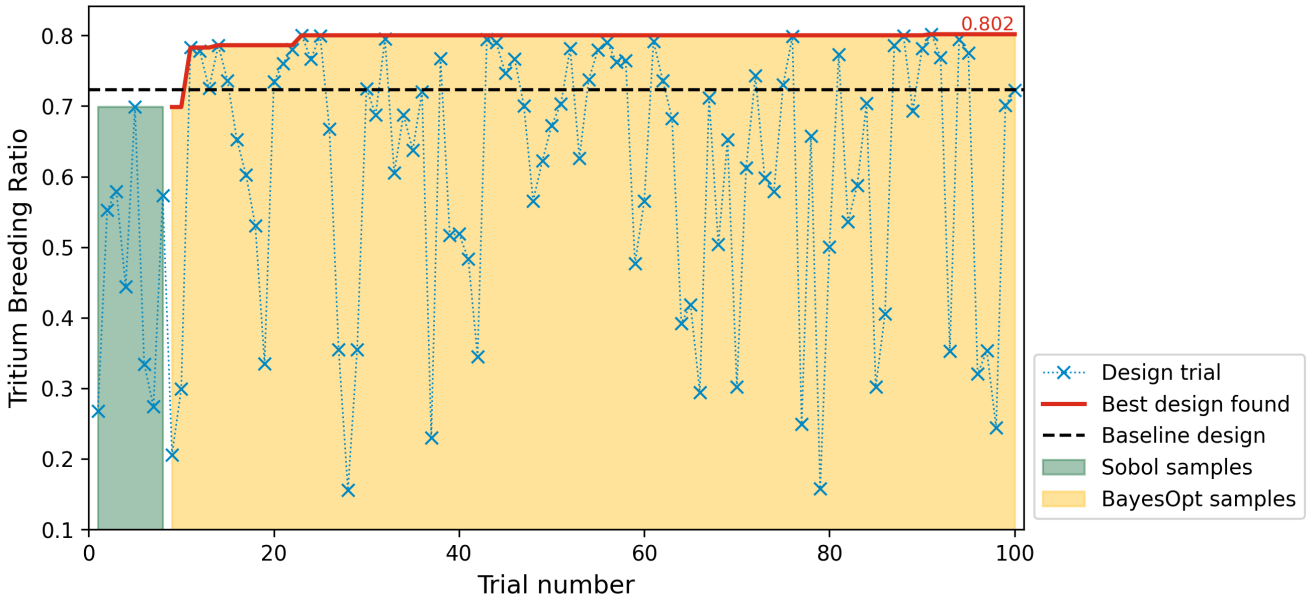

Figure 6b. Figure of merit for strategy 2 with 8 Sobol samples and $\xi = 0.01$.

Figure 6. Figure of merit (TBR) for the neutronics-only optimisation for each trial, plotted in the order they were evaluated. The solid red line shows the evolution of the best design found thus far at each step of the sequence. The dashed black line shows the performance of the baseline design. Candidates selected by Sobol sequence are shown with a green background on the left; those selected by the Bayesian optimisation are shown with a yellow background on the right.

weight of uncertainty in the acquisition function's choice of candidate. For this two-dimensional search space and fairly linear response surface, strategy 2 (favouring exploitation) was preferable. However, it should be noted that as the dimensionality is increased, the size of the search space increases exponentially, and thus the importance of exploration will become more pronounced.

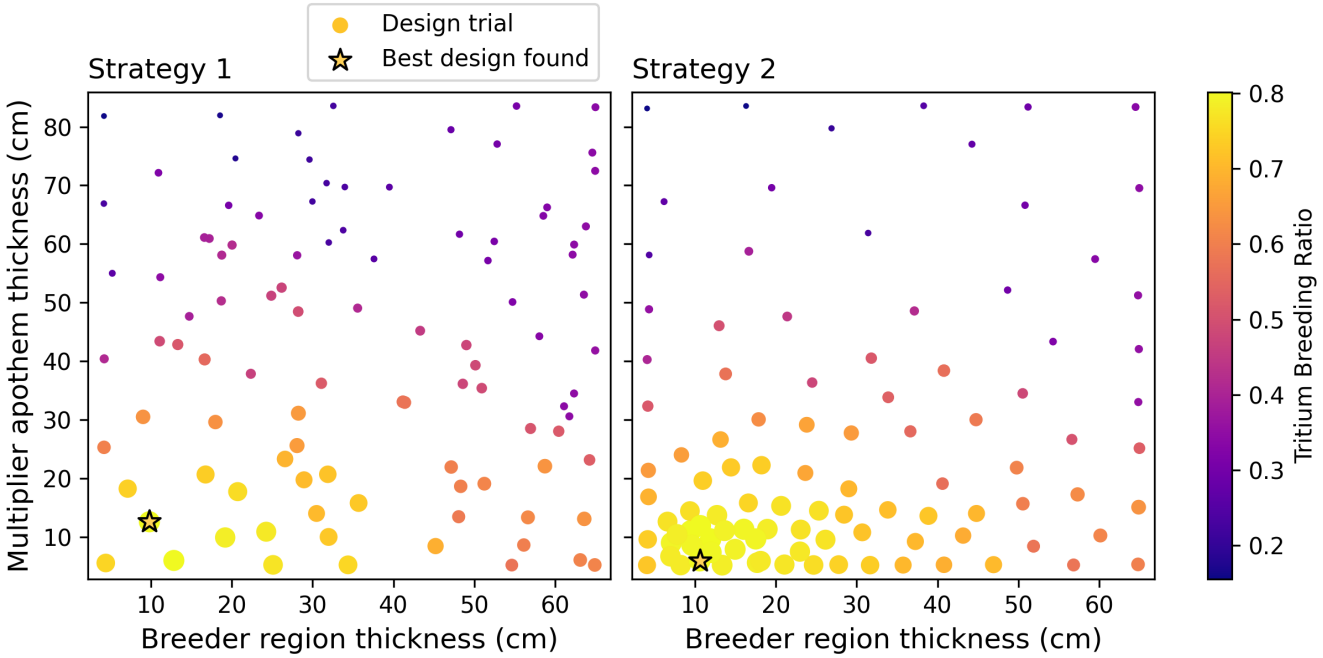

Figure 7. Distribution of samples within the 2-dimensional parameter space for the neutronics optimisation. The two panes correspond to two different values of the hyperparameter, $\xi = 2.5$ (left) and $\xi = 0.01$ (right). The colour scale and size of the marker indicate the corresponding value of TBR attained for the design point. The best design is indicated by a star marker.

The optimised designs for each strategy are shown in fig. 8, alongside the baseline design and the worst-performing design as a point of comparison. Many smaller pin-cells were preferred over fewer larger pin-cells, indicating a preference for breeder and multiplier to be distributed as homogeneously as possible. This is in line with the expectation for fixed source problems, where the optimal result is typically one where material boundaries are fully blurred and there is a closer to equal probability of interacting with any material in the problem. This maximises the probability of sequential moderation and absorption, where in this context moderation refers to the scattering of a neutron into the thermal region of the Li_6 cross section and absorption refers to a Li_6 tritium production reaction.

3.2 Neutronics and Thermal Hydraulics optimisation

3.2.1 Definition of optimisation campaign

As noted earlier, optimising a geometry when considering a single physics domain in isolation would result in certain volumes being minimised where they serve no role to augment the given performance metric. For example, as the coolant does not act to increase TBR, is not reasonable to optimise those parameters which control the coolant if simulating only neutronics. Further exploration of the methodology within a higher dimensional parameter space dictates a more complete treatment of physics. Therefore, a second study was performed where the performance of each design was evaluated by computing metrics arising from both neutronics and thermal hydraulics simulations. The setup of the modelling was described in detail in sections 2.2 and 2.3.

The two metrics used were the TBR and the pressure drop across the pin (computed as the difference in pressure between inlet and outlet), which should be maximised and minimised respectively. The objective

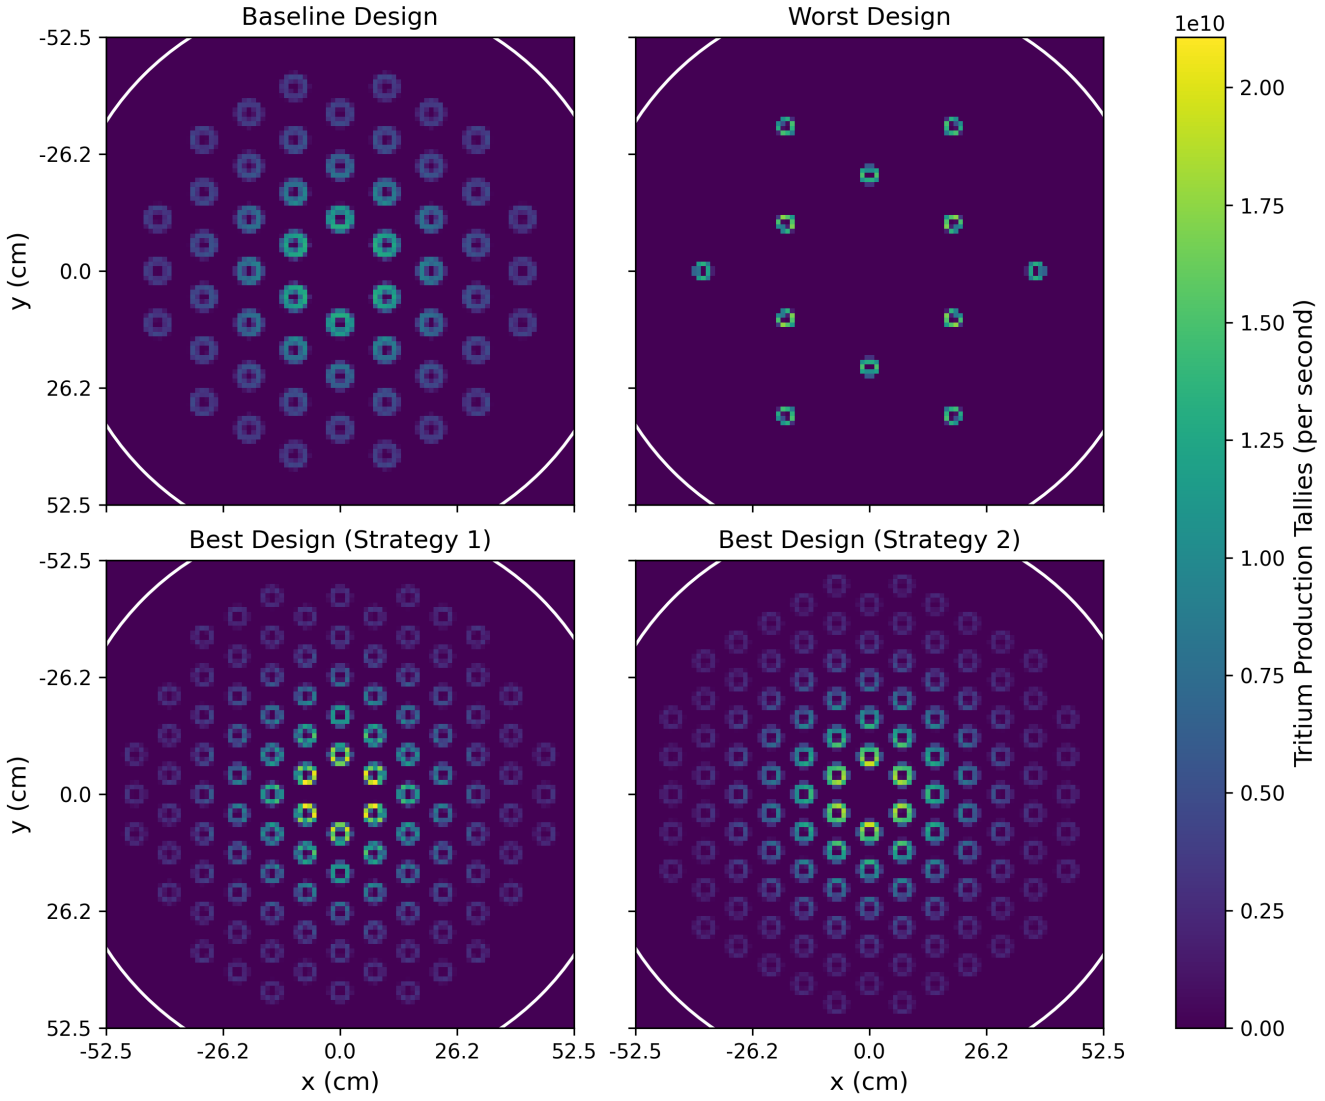

Figure 8. Spatial distribution of tritium production in the x-y plane for the baseline, worst and best designs arising from the neutronics-only optimisation. Tallies were scored by OpenMC on a regular mesh, with a summation over contributions along the z-direction.

function is then defined as a weighted sum of two metrics:

$$f = \frac{\mathrm{TBR}}{\mathrm{TBR}_{\mathrm{ref}}} + \left(1 - \frac{p_{in} - p_{out}}{p_{in}}\right) \tag{6}$$

In the first term, the evaluated TBR is normalised to a reference TBR value, taken to be $\mathrm{TBR}_{\mathrm{ref}} = 0.7$. In the second term, the pressure drop is normalised to the inlet pressure; as this relative quantity cannot exceed unity, subtracting this from 1 results in a positive-definite quantity. Moreover, maximising this quantity is equivalent to minimising the pressure drop, and thus may be additively combined with the first metric. In principle the relative importance of each of these terms could be controlled by multiplying the second term by an optional hyperparameter; however since the domain of each term is of order unity, this was deemed to be a suitable function. Nevertheless it is worth emphasising that the ambiguity in

the definition of objective function results in some arbitrariness in definition of "optimal". This matter is discussed further in section 4.

The parameters varied and their upper and lower limits are shown in table 4. All three radial parameters were varied between a factor of one third to 3 relative to the reference value. The breeder region length, constrained by the space available in the reflector, was varied between a factor of 0.75 to 1.25 relative to the reference value. With these additional parameters it is possible to explore the trade-off between increasing the TBR though increased breeder and multiplier volumes, or reducing pressure drop through increased coolant region thickness and decreased pin-cell length.

Following strategy 2 from the neutronics-only optimisation, the surrogate model is a GP and the acquisition function is EI with exploration parameter ξ kept at 0.01. Given the increased dimensionality, the number of initial Sobol samples was increased to 32.

Parameter name	Lower Bound	Upper Bound	Reference Value
Breeder region length	460.5 mm	767.5 mm	614 mm
Breeder region thickness	5.4167 mm	48.75 mm	16.25 mm
Multiplier apothem thickness	1.6167 mm	14.55 mm	20.9 mm
Outer coolant region thickness	6.966 mm	62.69 mm	4.85 mm

Table 4. Parameters varied and their bounds for the combined neutronics and thermal hydraulics optimisation.

3.2.2 Results

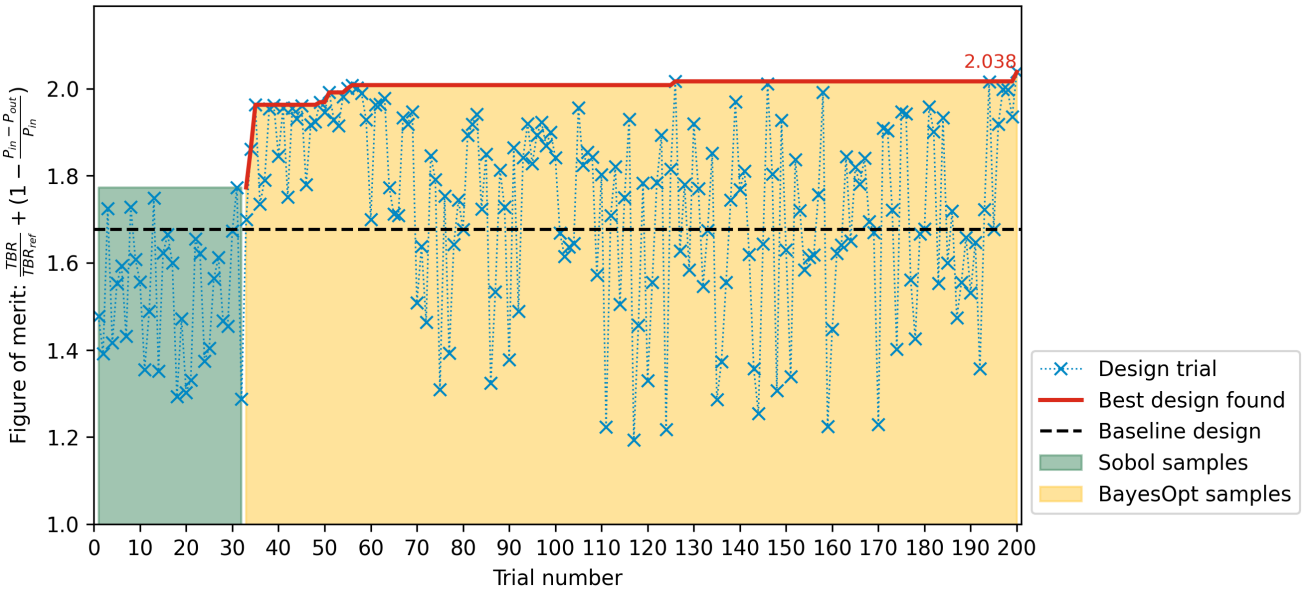

Figure 9. Figure of merit for the neutronics and thermal hydraulics optimisation for each trial, plotted in the order they were evaluated. The solid red line shows the evolution of the best design found thus far at each step of the sequence. The dashed black line shows the performance of the baseline design. Candidates selected by Sobol sequence are shown with a green background on the left; those selected by the Bayesian optimisation are shown with a yellow background on the right.

As shown in fig. 9, the figure of merit improved rapidly during the first few tens of Bayesian optimisation trials, then exploration dominates candidate selection until the final few trials where it is further improved. Indeed, the best design found was the the final trial. Compared with the baseline design, the figure of merit was increased from 1.68 to 2.04, corresponding to an improved TBR increased from 0.474 to 0.729, but worsened pressure drop, increasing from 8.42×10^2 Pa to 2.08×10^4 Pa.

The distribution of trials over input space is shown in fig. 10. As before, the exploration-exploitation trade-off achieves a balanced spread of points with clustering around performant regions. Note that compared to fig. 7, the increased dimensionality of the space means that nearby points in one subplot are nearby only in two dimensions, but not necessarily close in value of the other two dimensions not shown.

Within the bounds of the search space, two parameters reached extreme values: the outer coolant region thickness being minimised and the breeder length being maximised to their lower and upper bounds respectively. This suggests the true optimum for these parameters lies outside the search space bounds. The preference to increase the pin length suggests the cost of increasing the coolant length is not as detrimental to increasing the pressure drop as it is favourable in increasing the volume of breeder material (and hence TBR). Therefore the limiting factor is simply the overall size, as here constrained by a fixed reflector volume. Similarly, the cost for reducing the coolant hydraulic diameter (also acting to increase pressure drop) is not so severe as to prevent the preference for increased volume for breeder and multiplier regions. We note that the efficacy of cooling was not taken into consideration; were this to be included as an additional objective of the optimisation, it would further penalise a reduction in hydraulic diameter. In addition, there are also likely to be practical constraints upon the minimum coolant diameter arising from manufacturing considerations, as well as structural considerations for the steel cladding.

Furthermore, it should be understood that the locations of these true optima are dependent on the relative weights of the metrics in the figure of merit (as defined by eq. (6)), which were selected somewhat arbitrarily. Thus, it is possible that with different weightings the optimal values currently outside the search space could move within the search space and vice versa for those currently within the space. That is to say, the optima arrived at here are only those for this particular choice of objective function.

The impact upon sampling arising from the choice of objective function is also apparent in the distribution of trials over the two performance indicators, as shown in fig. 11. The x and y coordinates here correspond to the first and second terms in eq. (6), and their sum evaluates to the figure of merit. Therefore optimal designs are to be found in top right corner of the plot. The TBR is fairly sensitive to the sampling strategy (as shown by the broader distribution of points); meanwhile the range for the normalised pressure drop is extremely narrow, and sparsely sampled in the tail. Thus while the metrics' absolute values are of similar magnitude, their variance is not. Should the optimisation be repeated, this suggests a change of variables for the pressure drop performance indicator within the figure of merit to flatten this distribution might be desirable. The non-dominated points shown in fig. 11 approximate the Pareto frontier, that is: the set of designs which cannot improve in one metric without reducing performance in another. An alternative approach to multi-objective optimisation is to seek these Pareto-optimal designs, which is discussed further in section 4.

The best design found is visualised in fig. 12; the input parameters and evaluation metrics are shown in table 5. Compared with the neutronics-only results, the best design found now uses the same number of pin-cells as the baseline design, albeit with a different radial build. A thicker breeder region is used, at the expense of multiplier and coolant. However, a larger breeder zone implicitly increases the cross-sectional area of the annular flow channel for a fixed thickness. Therefore the inclusion of the pressure drop in

the figure of merit implicitly placed a penalty on smaller pins. This finding reinforces the notion that trade-offs are often subtle, and a holistic optimisation is essential.

Parameter name	Optimised Value	Baseline Value
Breeder region length	761 mm	614 mm
Breeder region thickness	34.5 mm	16.25 mm
Multiplier apothem thickness	9.37 mm	20.9 mm
Outer coolant region thickness	1.62 mm	4.85 mm
Tritium Breeding Ratio	0.729	0.474
$\mathrm{TBR}/\mathrm{TBR}_{\mathrm{ref}}$	1.04	0.676
Pressure Drop	2.08×10^4 Pa	8.42×10^2 Pa
$1 - (p_{in} - p_{out})/p_{in}$	0.997	1.00
Figure of Merit	2.04	1.68

Table 5. Input and output parameters for the optimised design given to three significant figures alongside the values for the baseline design.

4 OUTLOOK

In this work we have provided an indicative analysis for the purpose of demonstrating the potential of an integrated workflow for parametric design optimisation of a fusion breeder experiment. However a number of limitations exist in regard to both the modelling and optimisation methodologies that were presented in section 2, and here we describe a number of potential improvements.

From the perspective of the neutronics analysis it was already acknowledged that the geometry of the breeder mock-up technology was modelled independently from the room (although a graphite reflector was included to account for re-entrant thermalised neutrons). As information on the intended facility becomes available it would be desirable to understand the impact of the room and shielding, and assess any asymmetries that are introduced as a result. While such high-fidelity modelling might be impractical from the standpoint of an optimisation study with many iterations, depending on compute budget and appetite, even a one-off evaluation would nevertheless provide context for the comparative results yielded by the optimisation.

Thermal hydraulics was included to introduce a penalty arising from either reducing the coolant cross-sectional area or from increasing the length of the pin (both acting to undesirably increase the pressure drop). However the treatment was particularly simplistic, and several improvements could be made. Rather than consider the entire fluid manifold through the pin assembly, flow through a single pin was considered in isolation, assuming this to be sufficiently representative of all pins in the assembly, and neglecting any spatial spatial variations in operational conditions. Somewhat unphysically, the mass flow rate per pin was considered fixed; in reality for a fixed pumping power or total mass flow rate, the flow per pin would also depend on the number of pins, a quantity that in turn depends on the diameters of each. In addition, only the pressure drop corresponding to the annular outer coolant was computed; the inner coolant was omitted from considerations, its diameter being considered fixed. Assuming a single coolant loop, these two flow channels should not be treated independently. In future work we propose to model the complete flow manifold (including the impact of pumping) as this would enable a more holistic examination of the space of geometric configurations for the coolant.

For both neutronics and thermal hydraulics, the temperature was considered fixed, nor were any feedback mechanisms included. The impact of neutronic heating for an experiment anticipated to have comparatively

low neutron fluxes (which should be on the order of $10^{13} s^{-1}$ within the LIBRTI facility) is expected to be fairly minor; in full blanket systems however, this heating is likely to be considerable and a coupled treatment would be necessary. As an integration of OpenMC within MOOSE via Cardinal [13] is already available, inclusion of these effects is already feasible. More importantly for the current context is the impact of conjugate heat transfer to the solid structures, which was not considered. It is the intention that the breeder experiment be exposed to heating, with the role of the coolant being to maintain the operational temperature. As the coolant parameters would certainly impact the efficacy of this heat extraction and may result in gradients in temperature, it would be worthwhile to repeat the optimisation with a full thermal analysis.

Another limitation to this initial study was the selection of tritium breeding ratio as a performance indicator. The metric only accounts for the tritium produced in situ, making no attempt to model parasitic losses of tritium via permeation into and retention within solid structures. Therefore a more informative and holistic indicator of the breeder technology is the tritium removed from the system within the purge gas, requiring a proper treatment of tritium transport within the system. As the migration of tritium is highly sensitive to temperature, this would also involve a coupled multi-physics approach to modelling. Within the MOOSE ecosystem of software, a suitable and well-validated implementation already exists in TMAP8 [25]. Employing TMAP8 in an integrated manner to perform the optimisation using a more realistic performance indicator would be a natural extension to this work.

Another practical consideration relevant to the design of breeder blanket technology not yet explored here is the cost of materials. Depending on the specific choice of structural, multiplier, and breeder materials, increasing the volume of certain components may correspond to dramatically increased cost. Capturing the Pareto front for those geometric parameters which improve performance at increased financial cost will better equip decision-makers to assess the impact of altering a design. For a given cost model (for example, as implemented within the systems code Process [26, 27]), it should be straightforward to include cost as a figure of merit in the optimisation procedure demonstrated here. However, this would be predicated on the availability of data estimating the price of those materials, and for many advanced materials employed in fusion environments these are generally unknown for production at scale. Capital costs are influenced by a variety of factors, including natural abundance, safety of handling, maturity of manufacturing techniques and availability of a supply chain; as noted by [28] estimates of capital costs often carry large uncertainties, sometimes exceeding the difference in cost between materials. Therefore, additional work would be required to incorporate the impact of cost in a meaningful way.

It should be apparent from discussion thus far that in general there is more than one indicator of performance for assessing breeder blanket technology. The approach taken in this work was to combine figures of merit as a scalar weighted sum of metrics, and optimise for this single quantity. As can be seen in fig. 11, candidates chosen in this fashion tend to favour optimisation of that metric which dominates that sum. Provided that the relative weights are chosen appropriately, this may be appear as unproblematic or even desired behaviour. In practice however, even with domain expertise, selection of weightings a priori can be challenging and presents a significant source of risk to any multi-optimisation problem.

An alternative is to use a multi-objective acquisition function which actively seeks candidates on the Pareto frontier. One such acquisition function is the multi-objective equivalent to EI: q-Expected Hypervolume Improvement (qEHVI) [29], for which the noise-aware variant (q-Noisy EHVI) has been used by UKAEA previously to investigate optimisation of plasma current drive profiles [30]. Multi-objective optimisation removes the need to weight metrics; rather than returning a single optimised design, a number of such Pareto-optimised designs are returned, each of which offer alternative trade-offs between

objectives. By providing the full spectrum of possibilities within the trade-off space a more informed decision can be made on which Pareto-optimised design(s) to progress.

There are two main options for implementing multi-objective optimisation into SLEDO. Given that SLEDO's latest version uses RayTune for deployment on HPC, the clearest approach would be to utilise existing support for multi-objective within RayTune. However, at the time of writing, support for this feature is limited; only one of the Bayesian optimisation tools natively integrated with RayTune, Optuna [31], features multi-objective optimisation. If additional customisation control is desired over how the optimisation is performed than is offered out-of-the-box by RayTune and Optuna, an alternative approach would be to implement the optimisation loop directly using a lower-level Bayesian optimisation framework, such as BoTorch [32] which was used in earlier versions of SLEDO.

As we have discussed, sequential learning techniques have great potential to accelerate the engineering design process for breeder blankets as well as any other novel components. The purpose of automated design optimisation is not to take the human out of the loop with regards to design decisions, but rather make best use of available resources to understand the design in question and characterise the trade-offs between performance metrics. Domain expertise remains crucial throughout this process in appropriately preparing the physical models, selecting and bounding an appropriate search space, as well as interpreting and validating the results. Ultimately, it is human decision-makers that must select a design to take forward into production or the next step in the design process; these techniques enable us to best equip those decision-makers.

5 CONCLUSIONS

In this work we have presented a digital engineering pipeline for generating, evaluating and optimising fusion breeder technology. This involved a number open-source software packages. A new tool, Hypnos, was implemented to instantiate geometry from input parameters and encapsulate the design point. OpenMC and MOOSE were employed to perform neutronics and thermal hydraulics analyses respectively. Finally, SLEDO was used to perform sampling of the design space and optimise the stipulated figure of merit.

The target application selected for optimisation is a conceptual design for the solid breeder experiment at the planned LIBRTI facility, whose purpose is to provide qualification of the tritium breeding technology and validation of fundamental tritium transport modelling capability.

Two optimisation studies were performed. In first study, the analysis only considered neutronics, and the selected figure of merit was the tritium breeding ratio. This limited the number of parameters which could be reasonably varied; the two selected corresponded to radial thicknesses of the breeder and multiplier materials. In addition the number of pins in the assembly was set as a dependent quantity, corresponding to the maximum which could fit within a fixed radius.

Two separate campaigns were performed in this study. The first favoured exploration having both a large value for the hyperparameter ξ and a large number of Sobol samples to initially characterise the search space. As the dimensionality of the parameter space was low, it was found that it was a more effective strategy to prioritise exploitation. In the second campaign, a lower value for the hyperparameter was used and fewer Sobol samples were generated. It was in this second campaign that the highest TBR was attained: the best design gave rise to a TBR of 0.802, compared to the baseline value of 0.723. The optimised design was found to have a greater number of smaller pins, indicating a preference for a more homogenised arrangement.

In the second study, the analysis was extended to included a simple treatment of thermal hydraulics for a representative pin. The figure of merit was updated to be a weighted sum formulated from the tritium breeding ratio and the pressure drop over the outer coolant. This permitted a consideration of additional parameters, namely: the radial thickness of this outer coolant region, and the axial length of the pin-cell breeder region, with the latter constrained by the height of the surrounding reflector.

Compared with the baseline design, the figure of merit was increased from 1.68 to 2.04, corresponding to a TBR increased from 0.474 to 0.729 at the cost of pressure drop which increased from 8.42×10^2 to 2.08×10^4 Pa. In this case, the optimised design did not change the number of pins relative to the baseline design, resulting in fewer larger pins as compared to the neutronics-only optimisation. This arose as a result of the implicit dependence of the outer coolant cross-sectional area on the breeder thickness. Otherwise, the penalisation to reducing the coolant region from minimising pressure drop was not so large as to prevent the two parameters controlling the coolant reaching their extrema. However, we note the dependence of the final optimum upon some arbitrariness in the definition of the figure of merit. When defining this, considerations should include not just absolute value, but also variance of the metrics included, and relative importance. Since the latter is often subjective, a more sophisticated treatment was discussed in section 4.

Overall, we have demonstrated that Bayesian methods can be effective at characterising the behaviour of a specific figure of merit within a multi-dimensional parameter space, in the context of fusion breeder technology. It is possible to improve upon the baseline design with just a small number of trials. This is of particular relevance where a single analysis may involve computationally expensive simulations; in the present study, neutronics is indicative of an analysis requiring high performance computing. However, in the context of a fully-coupled multi-physics treatment (which as discussed in section 4 is deemed to be desirable), this would increase the computational load further, and the impact would be magnified.

When competing objectives are present, an ambiguity arises in defining the figure of merit, which may involve reweighting the contributions from difference metrics according to their relative importance. In this scenario, it would be desirable to present the full range of options available to a decision-maker, by utilizing the techniques of multi-objective optimisation to characterise the Pareto front. In addition to the aforementioned inclusion of high-fidelity coupled multi-physics, this should be a priority for further work.

Notwithstanding those limitations already discussed, this work has demonstrated the potential for parametric design optimisation to accelerate the design of fusion breeder technology. With further development, this will empower decision-makers to comprehend a highly-complex design space, and identify improved configurations more efficiently.

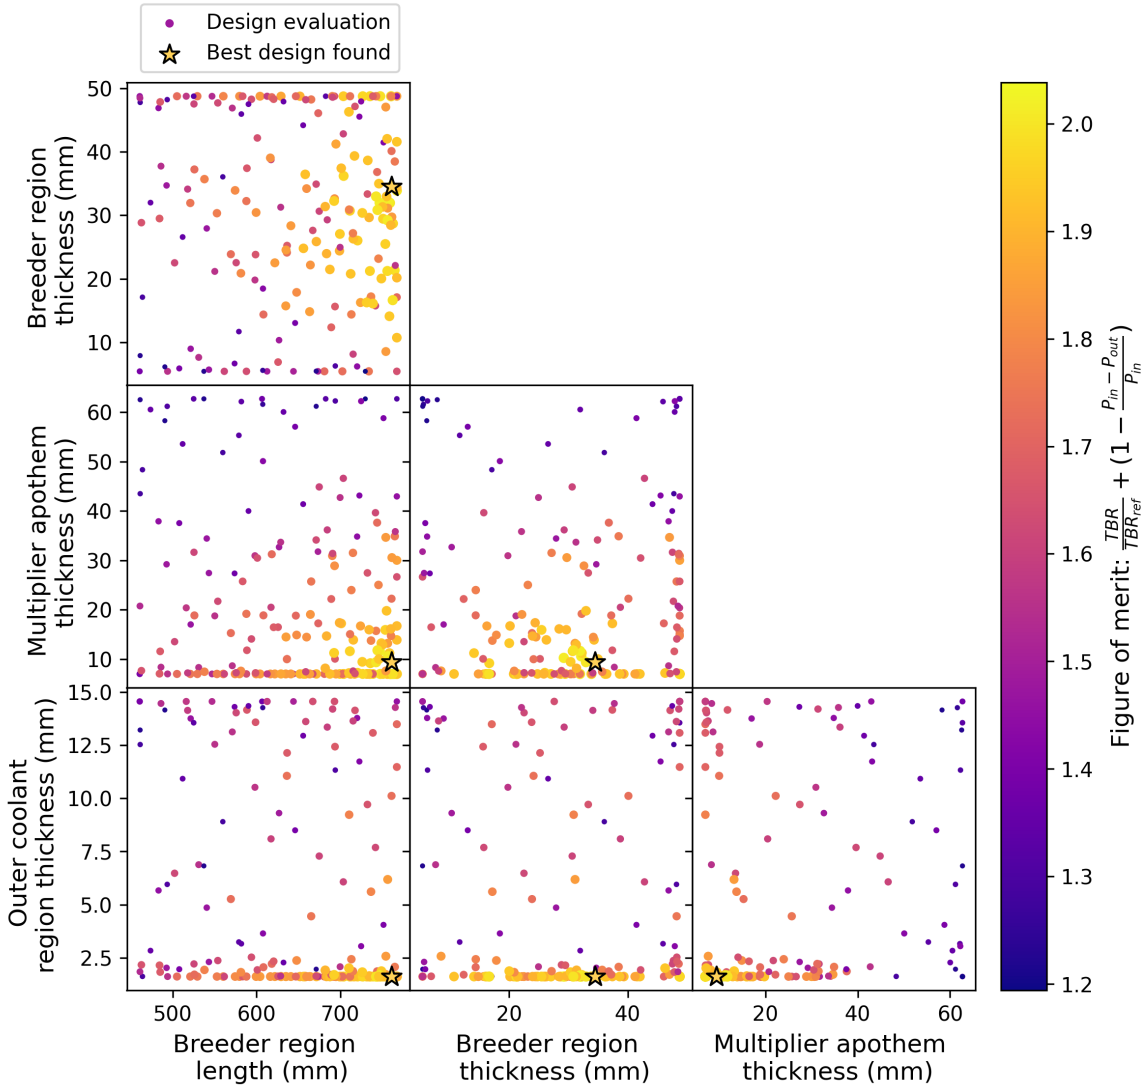

Figure 10. Distribution of samples within the 4-dimensional parameter space for the combined neutronics and thermal-hydraulics optimisation. Each pane corresponds to a unique projection of the space onto 2 of the parameters. The colour scale and size of the marker indicate the corresponding value of the figure of merit attained for the design point. The best design is indicated by a star marker.

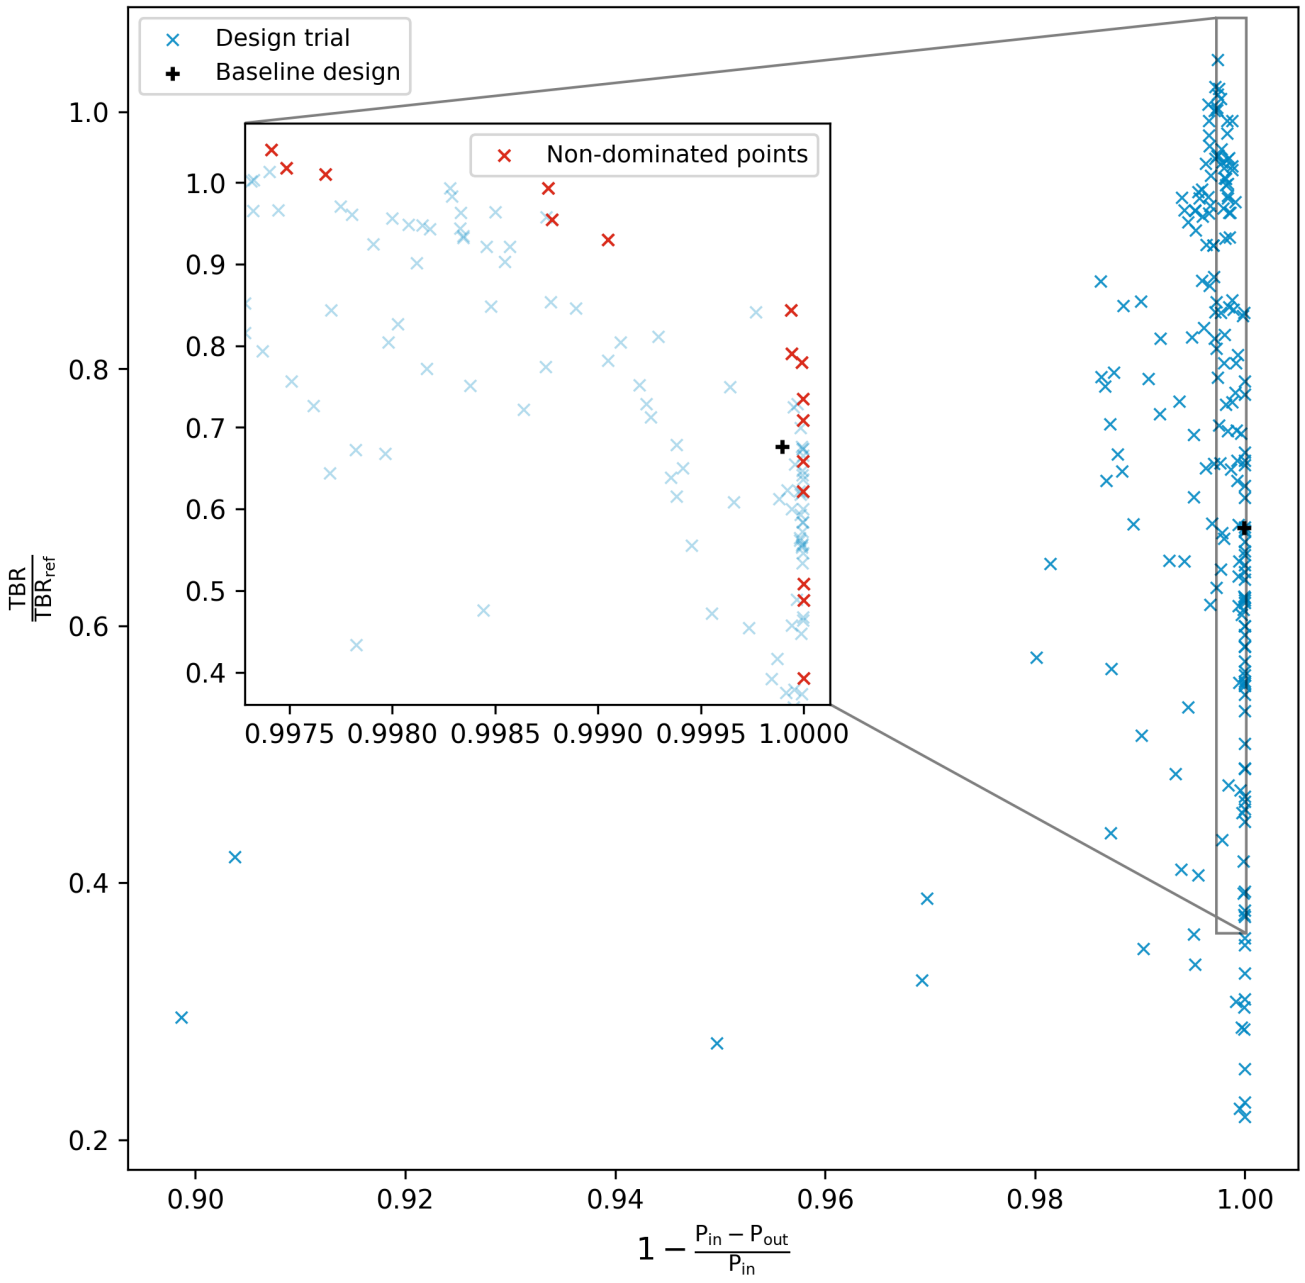

Figure 11. Comparison of trial performance in each term of the figure of merit. Design trials are shown with blue crosses, while the baseline design is shown as a black plus. The inset highlights non-dominated points as red crosses.

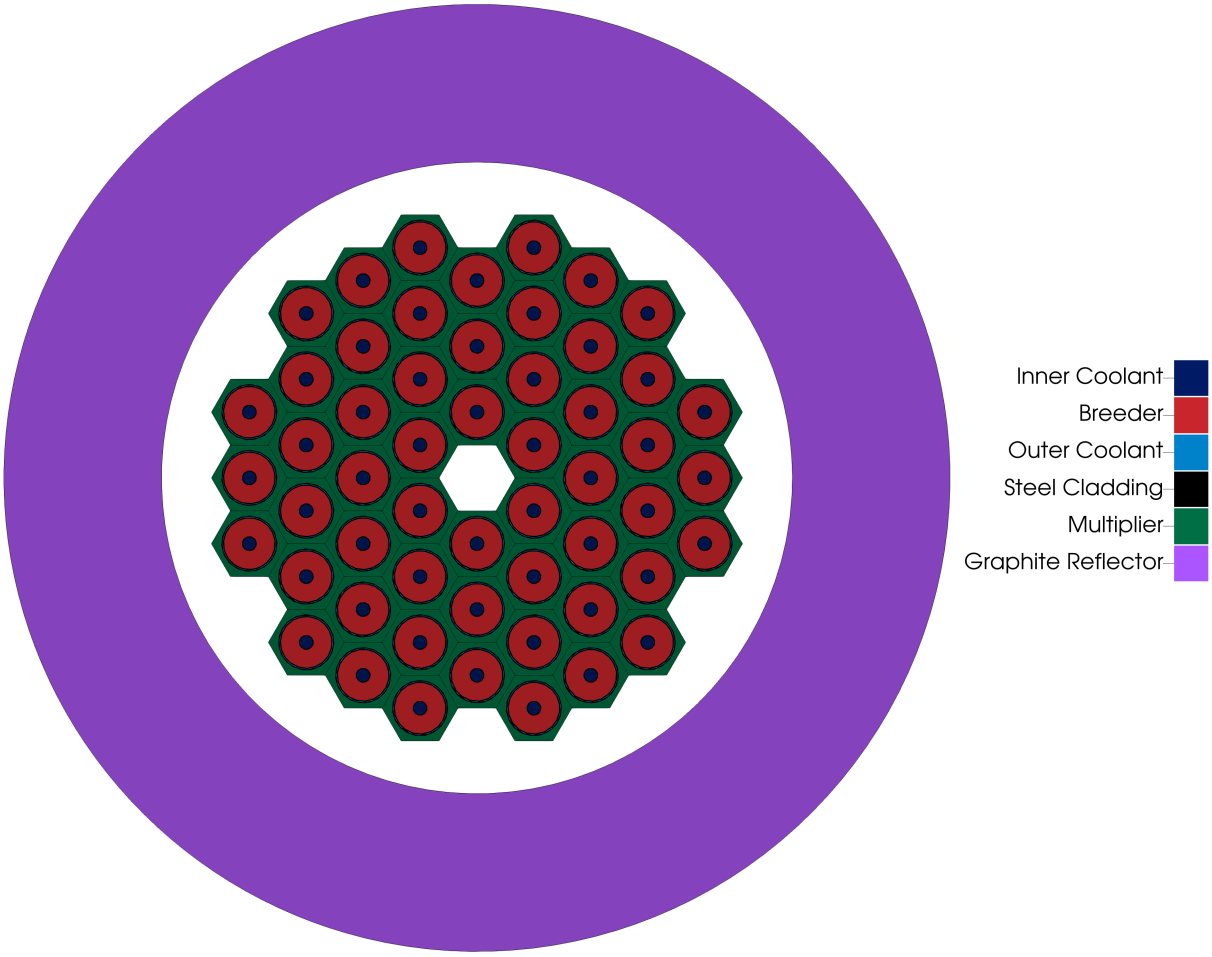

Figure 12. Cross-section geometry for the best-performing pin-cell assembly design arising from the combined neutronics and thermal hydraulics optimisation.

CONFLICT OF INTEREST STATEMENT

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

AUTHOR CONTRIBUTIONS

LH: Software, Formal analysis, Investigation, Methodology, Visualization, Writing – original draft. **HB:** Software, Conceptualization, Methodology, Supervision, Project administration, Writing – original draft. **SM:** Software, Visualization, Writing – review & editing. **AD:** Supervision, Resources, Writing – review & editing. **DF:** Project administration, Writing – review & editing.

FUNDING

This work was carried as part of the UKAEA LIBRTI Programme.

REFERENCES

1 Gilbert MR, Foster D, Bradnam SC, Davis A, Martis V, Lavrentiev MY. Modelling of a Neutron Source Test Bed for the Fusion Fuel Cycle. *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering* (New York, NY: American Nuclear Society) (2025).

2 Mungale S. Hypnos (2024). doi:https://doi.org/10.5281/zenodo.14277222.

3 Romano PK, Forget B. The OpenMC Monte Carlo particle transport code. *Annals of Nuclear Energy* **51** (2013) 274–281. doi:https://doi.org/10.1016/j.anucene.2012.06.040.

4 Romano PK, Horelik NE, Herman BR, Nelson AG, Forget B, Smith K. OpenMC: A state-of-the-art Monte Carlo code for research and development. *Annals of Nuclear Energy* **82** (2015) 90–97. doi:https://doi.org/10.1016/j.anucene.2014.07.048. Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013, SNA + MC 2013. Pluri- and Trans-disciplinarity, Towards New Modeling and Numerical Simulation Paradigms.

5 Gaston D, Newman C, Hansen G, Lebrun-Grandié D. MOOSE: A parallel computational framework for coupled systems of nonlinear equations. *Nuclear Engineering and Design* **239** (2009) 1768–1778. doi:https://doi.org/10.1016/j.nucengdes.2009.05.021.

6 Permann CJ, Gaston DR, Andrš D, Carlsen RW, Kong F, Lindsay AD, et al. MOOSE: Enabling massively parallel multiphysics simulation. *SoftwareX* **11** (2020) 100430. doi:https://doi.org/10.1016/j.softx.2020.100430.

7 Giudicelli G, Lindsay A, Harbour L, Icenhour C, Li M, Hansel JE, et al. 3.0 - MOOSE: Enabling massively parallel multiphysics simulations. *SoftwareX* **26** (2024) 101690. doi:https://doi.org/10.1016/j.softx.2024.101690.

8 Humphrey LR, Dubas AJ, Fletcher LC, Davis A. Machine learning techniques for sequential learning engineering design optimisation. *Plasma Physics and Controlled Fusion* **66** (2024) 025002. doi:10.1088/1361-6587/ad11fb.

9 LLC C. Coreform Cubit (Version 2024.3). Computer software (2024).

10 Tautges T, Wilson P, Kraftcheck J, Smith B, Henderson D. Acceleration Techniques for Direct Use of CAD-Based Geometries in Monte Carlo Radiation Transport. *Proceedings of M&C* (2009).

11 Schoof LA, Yarberry VR. EXODUS II: A finite element data model. Tech. rep., Sandia National Lab. (SNL-NM), Albuquerque, NM (United States) (1994). doi:10.2172/10102115.

12 Zhou G, Hernández FA, Pereslavtsev P, Kiss B, Retheesh A, Maqueda L, et al. The European DEMO Helium Cooled Pebble Bed Breeding Blanket: Design Status at the Conclusion of the Pre-Concept Design Phase. *Energies* **16** (2023). doi:10.3390/en16145377.

13 Novak A, Andrs D, Shriwise P, Fang J, Yuan H, Shaver D, et al. Coupled Monte Carlo and thermal-fluid modeling of high temperature gas reactors using Cardinal. *Annals of Nuclear Energy* **177** (2022) 109310. doi:https://doi.org/10.1016/j.anucene.2022.109310.

14 Leys O, Leys JM, Knitter R. Current status and future perspectives of EU ceramic breeder development. *Fusion Engineering and Design* **164** (2021) 112171. doi:https://doi.org/10.1016/j.fusengdes.2020.112171.

15 Brown D, Chadwick M, Capote R, Kahler A, Trkov A, Herman M, et al. 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. *Nuclear Data Sheets* **148** (2018) 1–142. doi:https://doi.org/10.1016/j.nds.2018.02.001. Special Issue on Nuclear Reaction Data.

16 Kulcinski GL, Radel RF, Davis A. Near term, low cost, 14MeV fusion neutron irradiation facility for testing the viability of fusion structural materials. *Fusion Engineering and Design* **109-111** (2016) 1072–1076. doi:https://doi.org/10.1016/j.fusengdes.2016.01.022. Proceedings of the 12th International Symposium on Fusion Nuclear Technology-12 (ISFNT-12).

17 Berry RA, Zou L, Zhao H, Zhang H, Peterson JW, Martineau RC, et al. RELAP-7 Theory Manual. Tech. rep., Idaho National Lab. (INL), Idaho Falls, ID (United States) (2016). doi:10.2172/1262488.

18 Zhou G, Kang Q, Hernández FA, D'Amico S, Kiss B. Thermal hydraulics activities for consolidating HCPB breeding blanket of the European DEMO. *Nuclear Fusion* **60** (2020) 096008. doi:10.1088/1741-4326/ab96f2. Publisher: IOP Publishing.

19 Petersen H. The Properties of Helium: Density, Specific Heats. Viscosity, and Thermal Conductivity at Pressures from 1 to 100 bar and from Room Temperature to about 1800 K. Tech. rep., Danish Atomic Energy Commission, Risoe. Research Establishment (1969).

20 Liaw R, Liang E, Nishihara R, Moritz P, Gonzalez JE, Stoica I. Tune: A research platform for distributed model selection and training. *arXiv preprint arXiv:1807.05118* (2018).

21 Nogueira F. Bayesian Optimization: Open source constrained global optimization tool for Python (2014).

22 Williams CK, Rasmussen CE. *Gaussian Processes for Machine Learning*, vol. 2 (MIT press Cambridge, MA) (2006).

23 M SI. The distribution of points in a cube and the approximate evaluation of integrals. *USSR Computational Mathematics and Mathematical Physics* **7** (1967) 86–112.

24 Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17** (2020) 261–272. doi:10.1038/s41592-019-0686-2.

25 Simon PCA, Icenhour CT, Singh G, Lindsay AD, Bhave CV, Yang L, et al. MOOSE-based tritium migration analysis program, version 8 (TMAP8) for advanced open-source tritium transport and fuel cycle modeling. *Fusion Engineering and Design* **214** (2025) 114874. doi:10.1016/j.fusengdes.2025.114874.

26 Kovari M, Fox F, Harrington C, Kembleton R, Knight P, Lux H, et al. "PROCESS": A systems code for fusion power plants – Part 2: Engineering. *Fusion Engineering and Design* **104** (2016) 9–20. doi:https://doi.org/10.1016/j.fusengdes.2016.01.007.

27 Morris J, Muldrew S, Knight P, Kovari M, Pearce A, Kahn S, et al. PROCESS (2023). doi:10.5281/zenodo.8338226.

28 Lux H, Brown C, Butcher M, Chapman R, Foster J, Nawal N. Optimizing the cost of the STEP programme. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **382** (2024) 20230413. doi:10.1098/rsta.2023.0413.

29 Daulton S, Balandat M, Bakshy E. Differentiable Expected Hypervolume Improvement for Parallel Multi-Objective Bayesian Optimization. Larochelle H, Ranzato M, Hadsell R, Balcan MF, Lin H, editors, *Advances in Neural Information Processing Systems* (Curran Associates, Inc.) (2020), vol. 33, 9851–9864.

30 Brown T, Marsden S, Gopakumar V, Terenin A, Ge H, Casson F. Multi-objective Bayesian optimisation for design of Pareto-optimal current drive profiles in STEP (2023). ArXiv:2310.02669 [physics].

31 Akiba T, Sano S, Yanase T, Ohta T, Koyama M. Optuna: A Next-generation Hyperparameter Optimization Framework. *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (2019).

32 Balandat M, Karrer B, Jiang DR, Daulton S, Letham B, Wilson AG, et al. BoTorch: Programmable Bayesian Optimization in PyTorch. *CoRR* **abs/1910.06403** (2019).