INL/RPT-23-03678

Sensitivity Analysis, Reduced-Order Modeling, and Optimization of a Gas-Cooled Pebble-Bed Reactor

September 2023

M3AT-23IN0601013

Zachary M. Prince[1], Paolo Balestra[1], Javier Ortensi[1], Sebastian Schunert[1], Olin Calvin[1], Joshua T. Hanophy[1], and Kun Mo[2]

[1]*Idaho National Laboratory*
[2]*Argonne National Laboratory*

Idaho National Laboratory

INL is a U.S. Department of Energy National Laboratory operated by Battelle Energy Alliance, LLC

INL/RPT-23-03678

Sensitivity Analysis, Reduced-Order Modeling, and Optimization of a Gas-Cooled Pebble-Bed Reactor

M3AT-23IN0601013

Zachary M. Prince[1], Paolo Balestra[1], Javier Ortensi[1], Sebastian Schunert[1], Olin Calvin[1], Joshua T. Hanophy[1], and Kun Mo[2]

[1]**Idaho National Laboratory**
[2]**Argonne National Laboratory**

September 2023

Idaho National Laboratory
Advanced Reactor Technologies
Idaho Falls, Idaho 83415

http://www.art.inl.gov

Prepared for the
U.S. Department of Energy
Office of Nuclear Energy
Under DOE Idaho Operations Office
Contract DE-AC07-05ID14517

Page intentionally left blank

ART Program

Sensitivity Analysis, Reduced-Order Modeling, and Optimization of a Gas-Cooled Pebble-Bed Reactor

INL/RPT-23-03678

September 2023

Technical Reviewer: (Confirmation of mathematical accuracy, and correctness of data and appropriateness of assumptions.)

Ryan H. Stewart Nuclear Engineer	Date

Approved by:

Michael E. Davenport ART Project Manager	Date
Travis R. Mitchell ART Program Manager	Date
Michelle T. Sharp INL Quality Assurance	Date

ABSTRACT

This work presents and applies a workflow for performing design optimization on gas-cooled pebble-bed reactors. Based on previous research, a representative equilibrium core of a pebble-bed reactor and a depressurized loss-of-forced-cooling model are created. These applications are built using the Multiphysics Object-Oriented Simulation Environment (MOOSE), specifically utilizing Griffin, Pronghorn, and Bison. After defining design-related parameters and quantities of interest regarding reactor safety and efficiency, this multiphysics model is sampled using the MOOSE stochastic tools module. The result is a comprehensive dataset of configurations, enabling sensitivity analysis and the generation of reduced-order models. Subsequently, the dataset and reduced-order models are employed in an optimization study aimed at maximizing fuel utilization while adhering to safety and operational constraints. The optimization process leads to an improvement of fuel utilization by approximately 10%, compared to engineering-judgment-based nominal conditions.

Page intentionally left blank

CONTENTS

FIGURES

TABLES

ACRONYMS

ANN	artificial neural network
DLOFC	depressurized loss of forced cooling
FOM	full-order model
GP	Gaussian process
GPBR200	200 MW_{th} general pebble-bed reactor
HTGR	high-temperature gas-cooled reactor
KTA	Kerntechnischer Ausschuss
MOOSE	Multiphysics Object-Oriented Simulation Environment
MOOSE-STM	MOOSE stochastic tools module
PBR	pebble-bed reactor
PCE	polynomial chaos expansion
PR	polynomial regression
QoI	quantity of interest
RCCS	reactor cavity-cooling system
RMSE	root mean-square error
ROM	reduced-order model
RPV	reactor pressure vessel
TH	thermal hydraulics
TRISO	tristructural isotropic

Page intentionally left blank

1. INTRODUCTION

Pebble-bed reactors (PBRs) represent a unique class of reactor concepts that pose distinctive challenges in modeling, simulation, and design. These challenges primarily arise from the continuous movement of fuel within the reactor, which makes predicting or characterizing the impact of certain design or model parameters challenging. Therefore, a combination of multiphysics models and data-driven meta-models could offer valuable insights into how these parameters influence both safety and economic aspects of the reactor. The objective of this paper is to explore the application of these methods using the 200 MW_{th} general pebble-bed reactor (GPBR200) as a representative model [45, 54].

In general, PBRs comprise a cylindrical vessel into which tennis-ball-sized ceramic pebbles filled with fuel are inserted into the core thus constituting a pebble bed. The coolant circulates through this pebble bed; for high-temperature gas-cooled reactors (HTGRs), the gaseous fluid flows from top to bottom. These pebbles serve as their own moderators, constructed of graphite with embedded tristructural isotropic (TRISO) particles. In the steady-state equilibrium condition, the core remains continuously refueled with fresh pebbles at the inlet and discharged pebbles at the outlet. This feed consists of both fresh-only pebbles and recycled pebbles from the discharge. A detailed explanation of the design and engineering concepts mentioned in this work can be found in Reference [22]. Gas-cooled PBRs are particularly significant for study due to their various advantages over traditional reactor technologies. TRISO particles provide safety benefits due to their durability at high temperatures, preventing the release of radioactive material during malfunctions or accidents. Moreover, the specific fuel forms make extracting weapons-grade material difficult, enhancing proliferation resistance [42]. The high operating temperatures of these reactors enable a more efficient power cycle, maximizing fuel utilization. The continuous fuel cycle also enhances fuel utilization by uniformly consuming the fuel and reducing waste through burning long-lived isotopes.

Similar to any reactor design, HTGRs encompass numerous design parameters that result in a seemingly endless array of possible reactor configurations. Design processes filter these configurations by considering both safety and economic factors. Safety constraints involve adhering to principles that enable certification and prevent system damage. Economic viability

entails minimizing development and operational costs while maximizing energy output. Often, safety and economic considerations counteract each other. For instance, higher core temperatures enhance fuel cycle efficiency but also increase the risk of fuel damage and reduce the margin for system failure. Thus, identifying a viable and optimized design from the multitude of potential configurations is a complex task. This typically involves testing various design parameters, determining viable resulting configurations, and ultimately narrowing down to an optimized design.

Given that experimental data is often limited and lacks generality in the parameter and output space, modeling and simulation play a crucial role in the design process. Reference [50] provides a comprehensive overview of software packages used in published works. In this study, we construct a fully-coupled multiphysics GPBR200 model using Multiphysics Object-Oriented Simulation Environment (MOOSE)-based applications [27]. The multiphysics model encompasses neutronics, depletion, thermal hydraulics (TH), and pebble-scale heat conduction. Neutronics and depletion are modeled using the reactor analysis application Griffin [24]; details on Griffin's capabilities are provided in Reference [61]. Multidimensional TH is modeled using Pronghorn, an intermediate-fidelity coarse-mesh TH tool [38]. Additionally, heat conduction within pebbles is modeled using the fuel performance application Bison [62]. More details on each of these applications and their utilization can be found in Section 2.

Reference [56] offers an extensive literature review on uncertainty quantification and sensitivity analysis for HTGRs. Relevant literature for this work includes [13], which carried out engineering-judgment-based optimization of a salt-cooled PBR with a design and output space similar, albeit considerably more detailed, to that presented in this paper. Works specifically focused on optimization of HTGRs include References [6, 19, 51, 59], which concentrate on fuel management optimization, some involving safety-related constraints. This study emphasizes the development and application of an efficient workflow for design optimization, considering a significant design space and accounting for both safety and economic factors. The model employed in this analysis begins with a coupled neutronics-depletion-TH equilibrium core simulation [29, 50], serving as the initial condition for a depressurized loss of forced cooling (DLOFC) transient [65]. The design space encompasses various parameters related to fuel properties, pebble management, and core geometry. Viability and performance are assessed

through various quantities of interest (QoIs) involving neutron economy, power peaking, fissile metal production, system temperatures, and fuel utilization. Given the diversity of this design and output space, it becomes challenging to implicitly determine the effects of configuration changes, particularly when multiple parameters change simultaneously. Thus, a stochastic approach to design optimization is adopted for this study. There are several potential optimization algorithms, including gradient-based algorithms like gradient descent [12] or quasi-Newton methods [10], as well as gradient-free algorithms [11]. Particularly interesting are the use of particle swarm optimization [59] and genetic algorithms [19] for PBR fuel cycle optimization. For a detailed overview of multiobjective optimization algorithms related to nuclear engineering, refer to [55]. As the goal of this work is not to characterize optimization algorithms, we have adopted a purely stochastic approach. This involves randomly sampling the design space and selecting the configuration with the most favorable QoI values.

Any choice of optimization algorithm would inevitably require numerous evaluations of the physics model. Depending on the model's complexity, performing more than a few dozen sequential evaluations could become intractable. Parallel execution of the model could allow for more evaluations, but many optimization algorithms require sequential evaluations, and computing resources for design processes often involve clusters with limited parallel processes. Consequently, there arises a need for generating fast-evaluating surrogates or reduced-order models (ROMs) of the multiphysics model. This study characterizes various ROM methodologies that accurately predict QoIs for a given design configuration, effectively reducing the burden of evaluating configuration performance for optimization. These methodologies include polynomial regression (PR) [41], Gaussian processes (GPs) [46], and artificial neural networks (ANNs) [36].

ROMs are known to suffer from the so-called "curse of dimensionality," where as more system parameters are added to the model, exponentially more training data are required to maintain model accuracy. Therefore, performing a sensitivity analysis on the design parameters becomes valuable. Sensitivity analysis quantifies the impact of parameter perturbations on resulting QoIs, providing insights into how specific design considerations influence the system and potentially streamlining the design process. It also offers a method to filter high-dimensional design spaces, simplifying ROM generation and improving their accuracy within a fixed dataset. In this work, a global sensitivity analysis is performed using Sobol indices [53], computed using polynomial

chaos expansion (PCE)[58]. Generating the training data from running the multiphysics model, the sensitivity analysis, and ROM generation and evaluation are all handled using the MOOSE stochastic tools module (MOOSE-STM) [52].

The workflow explored in this study for performing design optimization on HTGR PBRs is depicted in Figure 1. The paper is structured into four main sections, each corresponding to distinct phases of the process. The first phase involves the development of the GPBR200 model, serving as a representative HTGR PBR design. Section 2 delves into the model's construction, outlines the applications and numerical methods utilized, and presents results using nominal design parameters. The second phase covers the selection of design parameters and QoIs relevant to system constraints and optimization, along with the generation of training data and sensitivity analysis. Section 3 describes the chosen parameters and QoIs, presents an analysis of the design space through exploration, and computes global sensitivity. The third phase entails ROM generation. Section 4 details the reduced-order modeling methods employed in this study and compares the different methodologies using results. The fourth phase involves evaluating the ROMs across a broader parameter space, applying system constraints to obtain a set of viable configurations, and selecting an optimized model. Section 5 examines the effects of system constraints on the evaluation set, presents the design parameters of the optimized model, and compares ROM predictions with the evaluation of the multiphysics model under optimized parameters. The paper concludes with a summary in Section 6, which also outlines areas for potential future work.

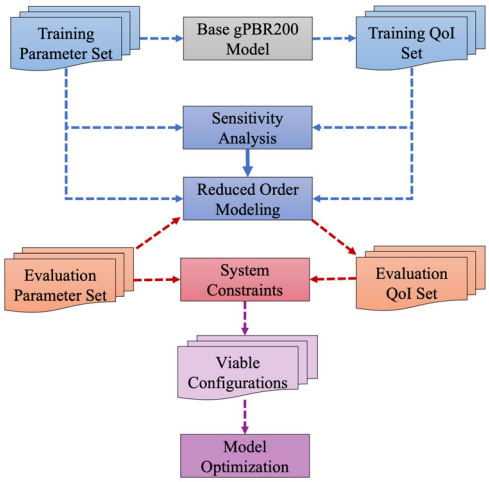

Figure 1: Workflow for performing design optimization.

2. MODEL DESCRIPTION AND METHODOLOGY

In this section, the representative HTGR PBR model, named GPBR200, is introduced. This model is the latest advancement compared to the ones presented in References [45, 54]. Notable differences include:

- A broader cross-section dataset with more grid variables (refer to Section 2.1)
- Enhanced representation of pebble streamlines, incorporating the pebble-bed conus (explained in Section 2.2)
- Utilization of the finite volume method for porous flow (discussed in Section 2.3)
- Improvements to the pebble and TRISO fuel-performance model (outlined in Section 2.4)
- Development of a DLOFC transient model (elaborated upon in Section 2.5).

This model is actively evolving, and forthcoming iterations are expected to be highlighted in subsequent publications. Its construction is informed by literature encompassing various designs, such as HTR-PM [64], Xe-100 [7, 35], HTR-Modul [57], PBMR400 [2], and HTTR [23]. The GPBR200 geometry is depicted in Figure 2, and some nominal core properties are presented in Table 1. This model adopts an axisymmetric (R-Z) representation for its development.

The model employs three distinct MOOSE-based applications: Griffin (neutronics and depletion), Pronghorn (TH), and Bison (pebble and TRISO heat conduction). A significant advantage of utilizing MOOSE-based applications is their programmable execution through the MOOSE `MultiApps` system, with the ability to couple each physics using the `Transfers` system. For an in-depth explanation of these systems, refer to Reference [16]. The `MultiApps` system operates within a hierarchical framework, where a primary application drives sub-applications representing individual physics and transfers relevant quantities between them. Figure 3 illustrates the `MultiApp` structure within the GPBR200 model.

In this configuration, the DLOFC serves as the primary application, initiating an initial equilibrium-core evaluation. It computes decay heat and core-wise solid temperatures at each time step, and then executes a series of pebble/TRISO sub-applications to acquire fuel and moderator temperatures. The equilibrium-core simulation is fully coupled, with Griffin serving

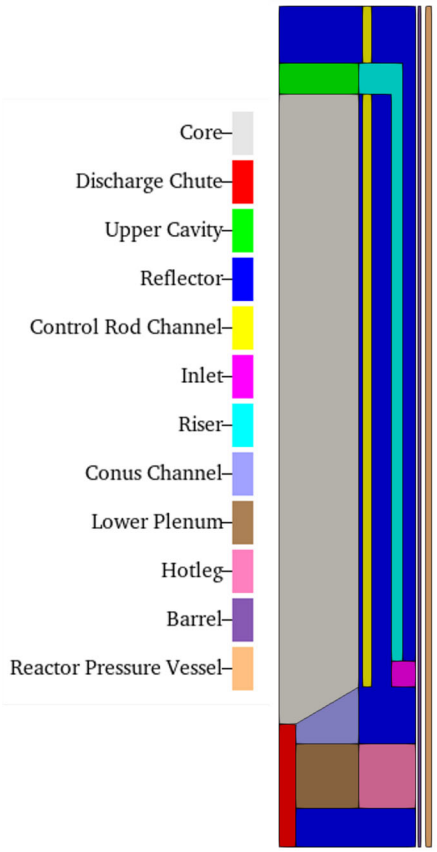

Figure 2: Geometric diagram of GPBR200 model.

Table 1: Nominal properties of GPBR200 model.

Property	Value	Unit
Power	200	MW_{th}
Core radius	1.2	m
Core height	8.93	m
Pebble diameter	6	cm
TRISO filling factor	9.34	%
Fuel kernel	UCO	—
Fuel kernel diameter	0.425	mm
Enrichment	15.5	%wt
Discharge rate	1.5	pebbles/min
Discharge limit	147.6	MWd/kg_{HM}
He flow rate	64.3	kg/s
He inlet temperature	260	°C
He outlet pressure	5.8	MPa
RCCS temperature	70	°C

as the driving application. Griffin employs a fixed-point iteration approach to converge all physics aspects: it begins with the batch of pebble and TRISO evaluations, followed by a streamline depletion solve, criticality calculation, and concluding with the TH evaluation. This iterative process continues until the calculated k_{eff} value converges.

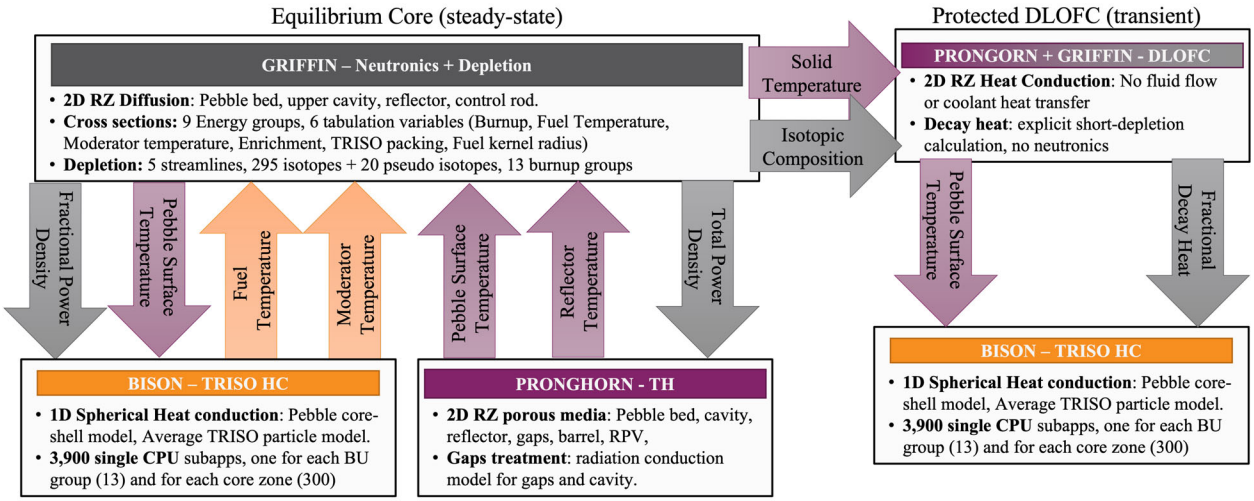

Figure 3: `MultiApp` structure of GPBR200 equilibrium core and DLOFC model.

The remainder of the section discusses each of the individual physics involved and the methods used in the applications and presents results for the model under nominal design parameters.

2.1 Neutronics Model

The neutronics physics within the GPBR200 model plays a central role in assessing the neutron population distribution across the reactor. This population is subsequently utilized in the depletion model to compute nuclide transmutation due to various reactions such as absorption and fission. The power density, which is also derived from this population, serves as heat generation input for the TH and pebble/TRISO models. The evaluation of neutronics involves solving a neutron transport criticality calculation, which takes the form of a generalized eigenvalue problem. In this problem, the multiplication factor (k_{eff}) represents the largest eigenvalue, while the neutron flux serves as the corresponding eigenvector. To ensure computational feasibility, we adopt a diffusion approximation of the transport equation (refer to Chapter 7 of Reference [14]). The geometry is discretized using the continuous Galerkin finite

element method with linear shape functions [66].

The model incorporates three distinct material regions: the pebble-bed core, the region outside the core, and the upper cavity. Microscopic cross sections for the first two regions are generated using a nine-energy group structure. Within the core region, cross sections were computed for 295 isotopes across six grid variables using Dragon [1]. These grid variables encompass burnup, fuel temperature, moderator temperature, TRISO kernel radius, TRISO filling factor, and fuel enrichment. Burnup is determined via the streamline depletion calculation, while fuel and moderator temperatures are obtained from the pebble and TRISO calculation. The remaining variables are supplied as inputs to the model. For the region outside the core, Dragon was also employed to derive cross sections for graphite and boron. For the upper cavity region near void, macroscopic cross sections were produced using Serpent [25], and the diffusion coefficient was corrected based on the cumulative migration method [28].

The control rod is modeled by adjusting the boron concentration in the channel region, based on the provided height. A series of equilibrium-core calculations were executed at varying control rod heights to ascertain the critical configuration, as depicted in Figure 4. Through this study, the critical rod depth was identified to be 1.747 meters into the active core region. A calculation was performed at this specific depth, yielding a k_{eff} value of 0.99961. This value implies that the reactor possessed 806 pcm of excess reactivity when the rod was fully withdrawn. The distributions of the fastest and most thermal neutron fluxes are illustrated in Figure 5.

2.2 Depletion Model

Conducting depletion calculations for PBRs introduces distinct challenges due to the continuous movement and refueling of pebbles within the core. Griffin employs an Eulerian approach to simulate the flow of pebbles along nonoverlapping streamlines within the core. Additionally, it employs a discrete representation of fuel burnup, as detailed in Reference [50]. The outcome of this calculation is a distribution of nuclide densities contingent on spatial location and burnup range (or group). These distributions are then used to calculate macroscopic cross sections for the neutronics calculation. A noteworthy aspect of this depletion solver is that the discharge (nonrecirculation) of pebbles is dictated by their burnup, rather than the number of passes.

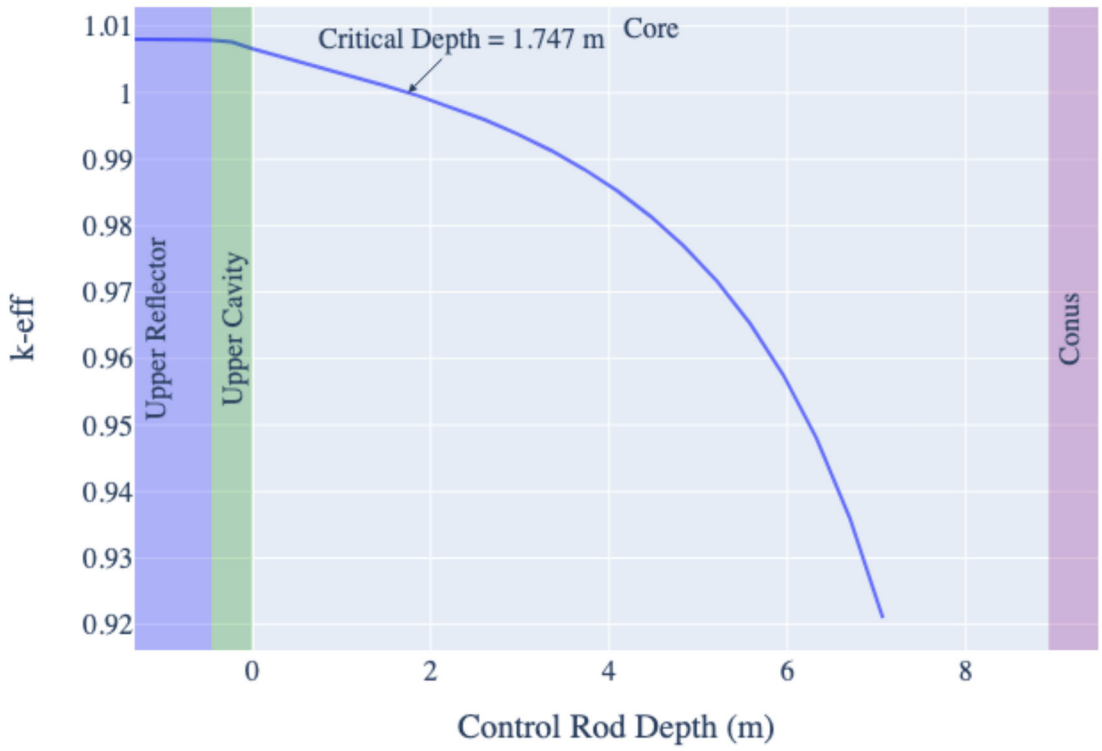

Figure 4: k_{eff} versus control rod position for the nominal equilibrium-core model.

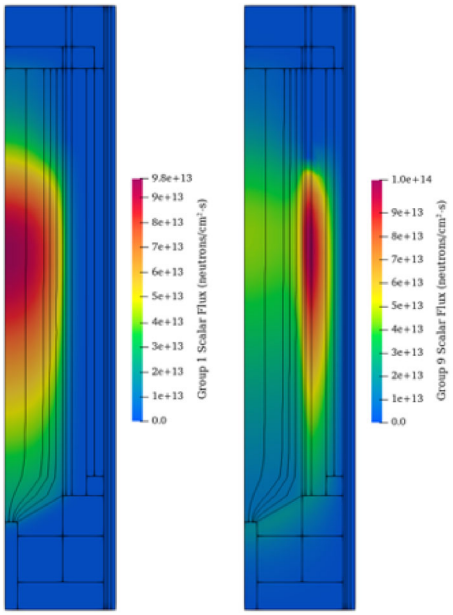

Figure 5: Group 1 (fastest) and Group 9 (slowest) neutron flux for the nominal equilibrium-core model.

In the GPBR200 model, five streamlines are integrated, derived from the discrete element simulation showcased in Reference [47] and illustrated in Figure 6. The model encompasses 13 burnup groups, evenly distributed across ranges of 16.4 MWd/kgHM. The nominal threshold for discharging pebbles without recirculation into the core's top is set at 147.6 MWd/kgHM. It's important to note that the Griffin streamline depletion implementation employs units of energy per pebble volume, thus requiring a conversion of the units mentioned above based on the heavy metal mass within the pebble.

Originally, the 295-nuclide decay and transmutation library, used in Dragon, was repurposed for integration into Griffin. However, during initial testing of the nuclide library designated for depletion calculations, a substantial under-prediction of decay power was observed when compared to a reference calculation using a 1,741 nuclide depletion library sourced from the comprehensive ENDF-VIII.0 dataset [8]. To address this discrepancy, a methodology akin to that presented in Reference [9] was employed. This involved modifying the original library by incorporating 20 pseudo-nuclides, resulting in decay heat predictions effectively matching those of the reference calculation.

The resulting power density and fissile nuclide concentrations are shown in Figure 6.

2.3 Thermal Hydraulics Model

The TH Pronghorn model utilized in this study is substantially based on the models presented in Reference [3]. The key methodological difference compared to Reference [3] involves the adoption of the weakly-compressible finite-volume formulation for the fluid and finite-volume discretization of the solid energy conservation equations, elaborated in Reference [26]. This finite-volume approach ensures local conservation of energy, momentum, and mass and offers improved treatment of discontinuities in friction forces and porosity.

Pronghorn is tailored to model PBRs at an intermediate fidelity level. This entails representing core flow conditions within a two-dimensional axisymmetric geometry, while consolidating geometric details into closure relationships that describe the exchange of momentum and energy between the fluid and stationary solid phases. These closure relationships encompass pressure drop correlations, heat transfer coefficient correlations, effective solid and fluid thermal conductivities. An overview of closure correlations commonly employed in Pronghorn can be

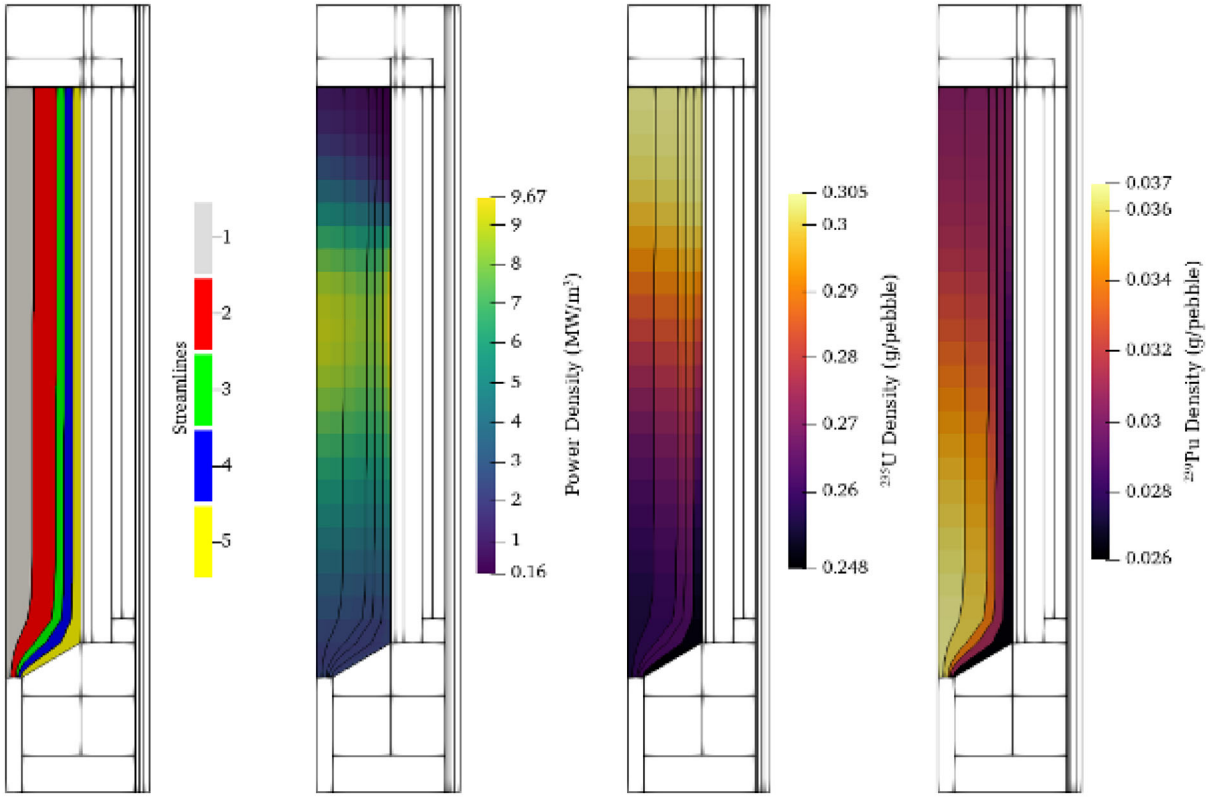

Figure 6: Depletion streamline definitions, computed power density, and resulting fissile nuclide concentrations for the nominal equilibrium-core model.

found in References [33, 49], particularly Table 1 in Reference [49], for a different PBR. For the sake of brevity, only discrepancies between the models used in Reference [49] and the present work are highlighted here.

The domain's porosity in each region is displayed in Table 2. Fluid properties are sourced from Reference [44]. Graphite thermal conductivity is set at 26 W/m·K, specific heat at 1697 J/kg-K, and density at 1,780 kg/m^3, subsequently adjusted for porosity as in Reference [49]. The heat transfer coefficient for the pebble bed is computed using the Kerntechnischer Ausschuss (KTA) correlation, differing from the Wakao correlation [39]. The resultant fluid velocity, pressure, temperature, and solid temperature under nominal conditions are illustrated in Figure 7.

Table 2: Porosity in each fluid region of the GPBR200 model.

Region	Porosity
Inlet	0.22
Riser	0.22
Upper cavity	1.0
Pebble bed	0.39
Conus channel	0.63
Lower plenum	0.40
Hotleg	0.07

2.4 Pebble and Fuel Particle Model

The fuel design of the GPBR200 model closely resembles that of most other gas-cooled PBR designs, as depicted in Figure 8. The primary aim of heat conduction calculations within pebbles and TRISO particles is to acquire moderator and fuel temperatures, which are then used for cross section interpolation. Additionally, this modeling effort facilitates the determination of temperatures in reactor hotspots where elevated temperatures could lead to substantial particle failure, potentially rendering certain design configurations unviable. However, performing calculations for every single pebble in the core and every particle within each pebble would be impractical. To address this, we employ a representative pebble model for every element within the core region and each burnup group. This separation by burnup group is essential, as pebbles at differing burnups exhibit distinct power densities; pebbles with lower burnup exhibit larger power densities, leading to higher temperatures.

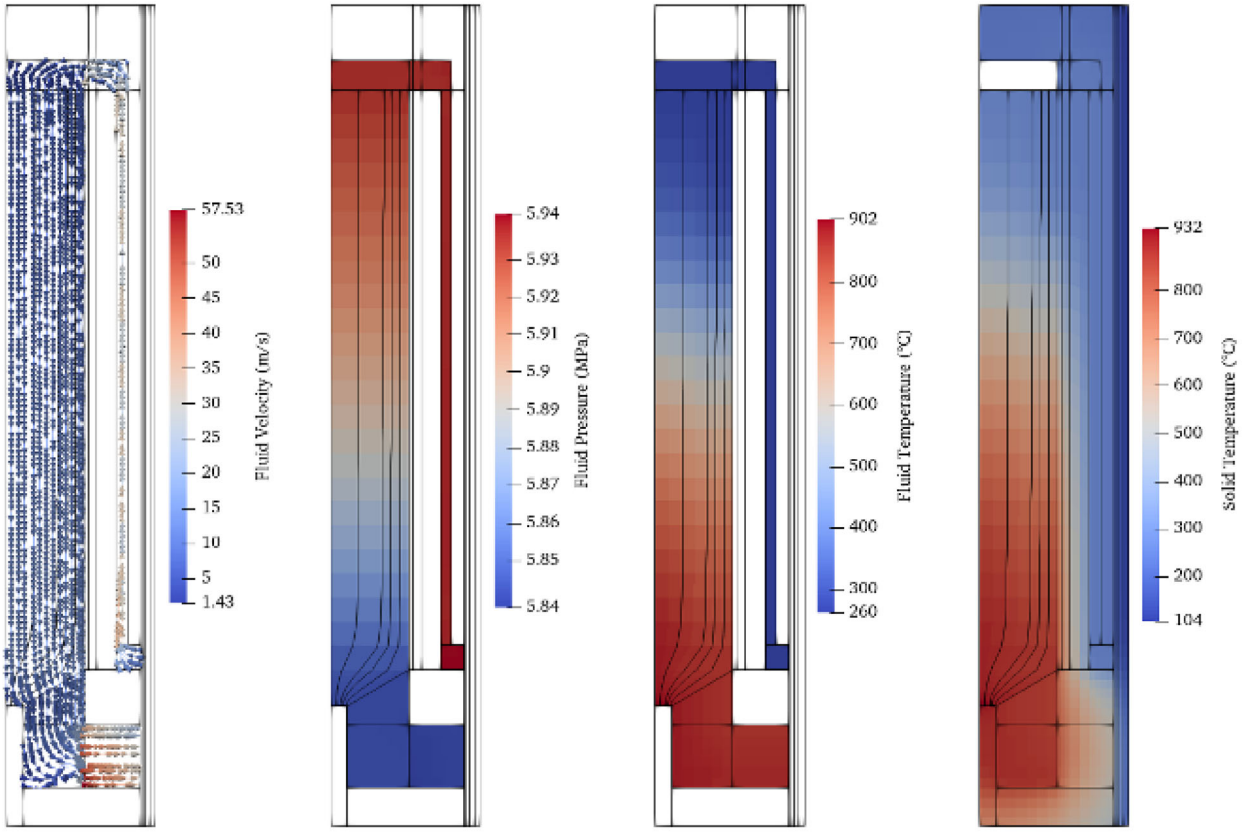

Figure 7: Fluid properties and solid temperature for the nominal equilibrium-core model.

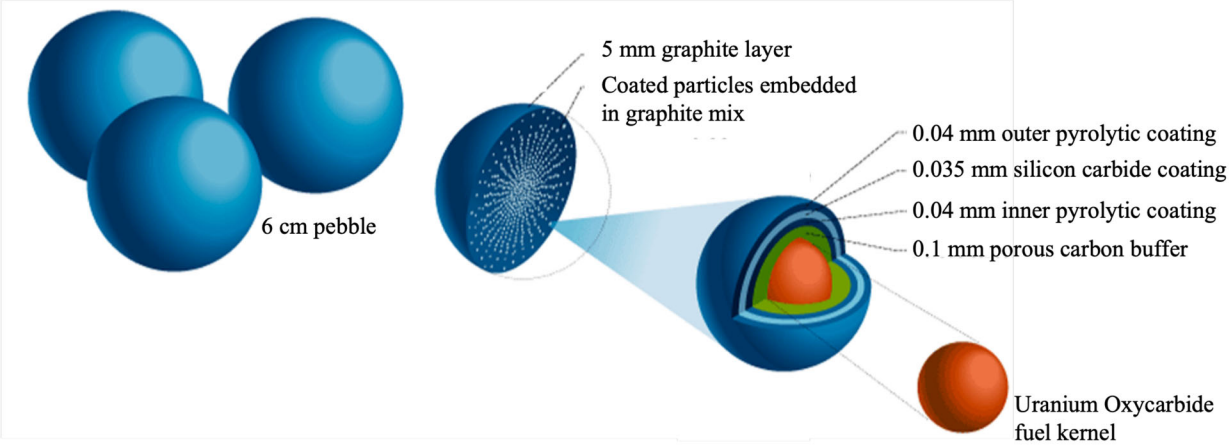

Figure 8: Schematic of TRISO-embedded graphite pebble. Modified from Reference [63]

Each representative pebble encompasses two one-dimensional spherically symmetric models: one for the pebble's graphite matrix and another for a TRISO particle. The heat source within the pebble stems from the partial power density obtained from the neutronics calculation, while the outer boundary temperature is obtained from the TH calculation. The particle model employs the same power density for both the fuel kernel and outer boundary temperature, which is set to the average temperature within the pebble matrix. The thermal properties of the graphite assume an A3-3 grade from Reference [18], and the fueled region is homogenized using the differential effective medium theory [31]; for more details, refer to Reference [60]. The thermal properties of all TRISO components are drawn from PARFUME [32]. It's important to note that all properties assume zero burnup and no neutron fluence. This dependency is anticipated to be integrated in future model iterations.

Figure 9 shows the resulting maximum power density and temperatures of the pebbles through the core.

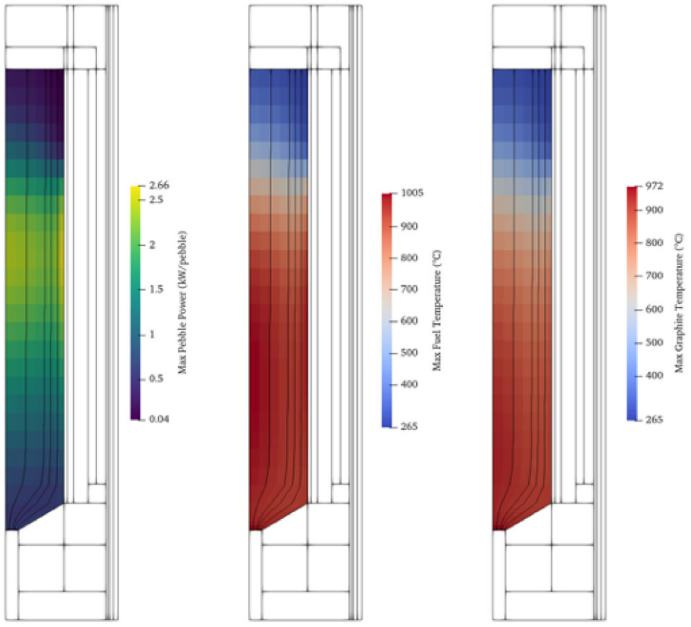

Figure 9: Max power density, fuel, and moderator temperature of representative pebbles for the nominal equilibrium-core model.

2.5 Depressurized Loss-of-Forced-Cooling Model

As part of the safety analysis of this model, the GPBR200 includes a protected DLOFC transient. A DLOFC represents one of the most critical design-basis accident scenarios, involving the complete loss of coolant pressure within the core. Although the core is shielded through control rod insertion, residual decay heat persists, gradually elevating the core temperature over approximately a week. During this period, heat dissipation is solely reliant on the reactor cavity-cooling system (RCCS), which can result in exceedingly high core temperatures.

This model assumes instantaneous control rod insertion, leading to the absorption of all available neutrons and obviating the need for neutronics evaluation. Similarly, the depressurization event transpires instantaneously, with the coolant at atmospheric pressure having no impact on core heat transfer. Consequently, no fluid flow assessment is conducted, and the TH model exclusively considers solid heat conduction. Decay heat computation involves explicit depletion calculations employing nuclide concentrations extracted from the equilibrium-core calculation. The transient is simulated over 150 hours, as illustrated in Figure 10.

3. SENSITIVITY ANALYSIS

This section is devoted to assessing the performance of the GPBR200 by systematically investigating a selected design space and evaluating it using a collection of QoIs. Four main objectives underscore the aims of this section:

1. **Identify design parameters and QoIs**: This involves establishing the pertinent design parameters and choosing the significant QoIs that capture essential aspects of the reactor's behavior.

2. **Perform qualitative sensitivity analysis**: A qualitative sensitivity assessment is conducted by visually representing the local effects that variations in design parameters exert on the outcomes of interest. This helps to gain insights into which parameters have the most substantial influence on the system's behavior.

3. **Produce a training dataset**: The construction of a comprehensive training dataset is undertaken by evaluating the model across the entire designated design space. This dataset

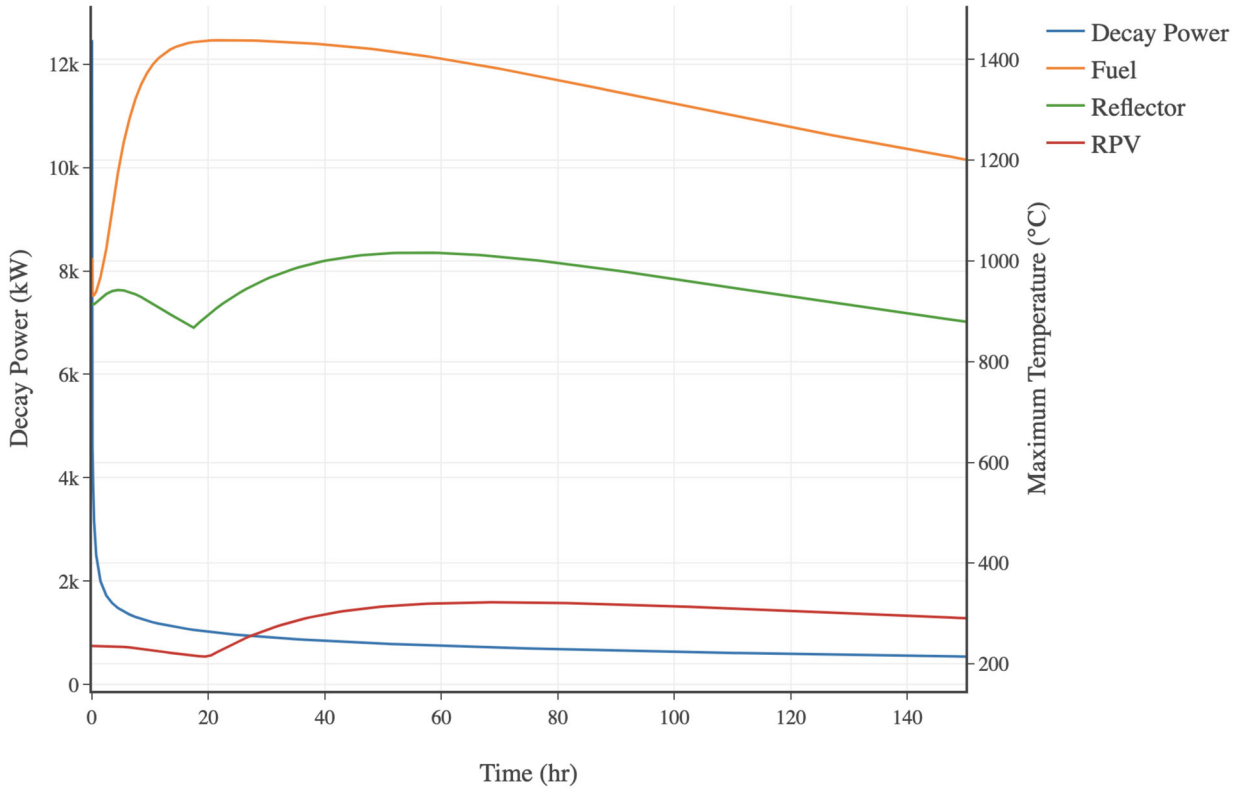

Figure 10: Decay power and maximum temperature in various regions during the nominal DLOFC model. Note that moderator temperature is visually identical to fuel temperature.

serves as the foundation for subsequent analyses and computations.

4. **Compute global sensitivities**: For every parameter-QoI pairing, global sensitivities are computed. This calculation provides an understanding of the parameter's overall impact on the QoIs and contributes to comprehending the behavior of the system as a whole.

By addressing these objectives, this section endeavors to comprehensively assess and characterize the behavior and performance of the GPBR200 model.

3.1 Design Space and Quantities of Interest

Before performing any type of sensitivity study or ROM generation, the design, uncertainty, and output space must be defined. The design and uncertainty space involves model parameters that are either part of a design decision or parameters that are not known exactly. These parameters are defined with a probability distribution describing the likelihood of a certain value occurring. Design parameters are typically defined as a uniform distribution, where the likelihood of a value is equal between two bounds. There are many different types of distributions for uncertain parameters; a common choice is a normal distribution specified by a mean value with a standard deviation. This work focuses solely on design parameters. Table 3 lists the parameters studied in this work with their specified bounds. The following describes each of these parameters:

- Kernel radius is the radius of the fuel kernel within all TRISO particles. This quantity affects the mass of fuel within the pebble and effective heat removal in the particle.
- Filling factor is the volumetric fraction of TRISO particles within the fueled region of the pebble (excluding graphite shell). This directly affects the mass of fuel and effective thermal conductivity in the pebble.
- Enrichment is the weight fraction of fissile uranium in the fresh pebbles: $\frac{w_{235}}{w_{238}+w_{235}}$.
- Feed rate is the rate at which pebbles are inserted into the top of core and removed from the bottom, including both fresh pebbles and recirculated ones. Since Griffin uses burnup-based discharge, the quantity directly affects the number of times a pebble passes through the core.

- Burnup limit determines when pebbles are discharged from the bottom of the core and not recirculated. When the limit lies within the range of a burnup group, a fraction of the pebbles in that group are discharged based on where in that range that lime lies.

- Total power is the integrated thermal power of the reactor. This affects the scaling of the neutron flux after the criticality calculation.

- Core radius is the radius of the cylindrical pebble-bed vessel. The outer radius of the reactor is fixed, so increasing the core radius decreases the thickness of the radial reflector.

- Core height is the height of the active core region, excluding the lower conus and upper cavity. Increasing the core height increases the total reactor height by the same amount.

Table 3: Design space for the GPBR200 model.

Parameter	Nominal Value	Lower Bound	Upper Bound	Unit
Kernel radius	0.2125	0.15	0.3	mm
Filling factor	9.34	5	15	%
Enrichment	15.5	5	20	wt%
Feed rate	1.5	1	3	pebbles/min
Burnup limit	147.6	131.2	164.0	MWd/kg_{HM}
Total power	200	180	220	MWth
Core radius	1.2	1.1	1.3	m
Core height	8.83	8	10	m

The output space involves specifying QoIs, which are typically global outputs from the simulations. In the context of optimization, these quantities are used as part of the function defining the design performance. Table 4 lists the QoIs studied in this work. The following describes each of these QoIs:

- k_{eff} is the effective multiplication factor of the reactor, computed from the neutronics model. This value is a measure of the neutron economy and determines the ability of the reactor to reach full power.

- Max pebble power is the integrated power within a pebble. This quantity is a safety metric of the system since pebbles have a design limit for how much power they can experience.

- Peaking factor is computed by taking the ratio of the maximum and average power density in the core. A value of 1 means that the power profile in the core is completely flat and larger values indicate hot spots in the core where pebbles may exceed designed temperature and power limits.

- Fissile Pu fraction is an averaged weight fraction of fissile isotopes of Pu from pebbles exiting the core, including those being recirculated and discharged. This quantity is computed using:

$$\frac{N_{239}M_{239} + N_{241}M_{241}}{N_{238}M_{238} + N_{239}M_{239} + N_{240}M_{240} + N_{241}M_{241} + N_{242}M_{242}}, \tag{1}$$

  where N_q is the averaged nuclide number density for isotope q at the bottom of the core and M_q is the mass per nuclide for isotope q. This quantity is used as a metric for the proliferation resistance of the reactor. Determining the quality of plutonium generated in pebbles exiting the core is indicative of the attractiveness and usefulness of the plutonium.

- ^{235}U utilization is computed based on the rate of fresh fresh fuel entering the core and the power using the expression:

$$\frac{\text{Total Power}}{\rho_{235} \times n_{\text{discharge}} \times \text{Feed Rate}}, \tag{2}$$

  where ρ_{235} is the mass of ^{235}U per fresh pebble and $n_{\text{discharge}}$ is the fraction of pebbles being discharged from reaching the burnup limit at the bottom of the core. This quantity is a measure of economy, as it is directly related to the amount of natural uranium needed to operate the reactor.

- Maximum operating fuel temperature is computed from the TRISO model during the equilibrium-core calculation. This quantity is a measure of how accident tolerant the reactor is, as higher temperatures leave less of a margin of failure during accident scenarios.

- Maximum DLOFC fuel temperature is computed from the TRISO model during the DLOFC transient calculation. This is also a metric for accident tolerance since high temperatures can cause significant fuel failure.

- Maximum reactor pressure vessel (RPV) temperature is computed from the TH model

during the DLOFC transient calculation. The RPV is a vital component of reactor containment and high temperatures could degrade the material and lower the margin of vessel failure.

Table 4: Quantities of interest for the GPBR200 model.

QoI	Nominal Value	Unit
k_{eff}	0.99961	—
Max pebble power	2.67	kW
Peaking factor	2.01	—
Fissile plutonium fraction	63.2	wt%
^{235}U utilization	984.1	MWd/kg_{235}
Max operating fuel temperature	1006	°C
Max DLOFC fuel temperature	1438	°C
Max RPV temperature	322	°C

3.2 One-Dimensional Grid Study

The goal of this study is to gain qualitative insight into how each design parameter affects each QoI. To achieve this, a systematic sampling of the design space is conducted, whereby each parameter is individually perturbed while keeping all other parameters fixed at their nominal values. For example, enrichment was changed uniformly between 5 and 20 wt%, while all other parameters were fixed. With 12 points for each parameter, the sampling resulting in 96 total simulations. With this data, the dependency of each parameter on each QoI can be plotted, which is shown in Figure 11.

Several significant insights can be drawn from the data exhibited in Figure 11. First, enrichment, kernel radius, and filling factor appear to be the most influential parameters for many of the QoIs. Since these parameters are the main drivers for the mass of fissile material in the core, this feature must be heavily related to the performance of the design. Second, some of the dependencies exhibit nonmonotonic behavior, shown by the plots for k_{eff}, max pebble power, and peaking factor. This indicates that there are optimal values for the parameters, thus hinting at the design space's potential for optimization. Note that the local optima may change from what is shown in the plots based on differing values of other parameters. Finally, the plots show nonlinear behavior for most of the parameter-QoI dependencies, which signifies possible difficulties in generating ROMs that emulate the behavior, necessitating the investigation of

complex model formulations.

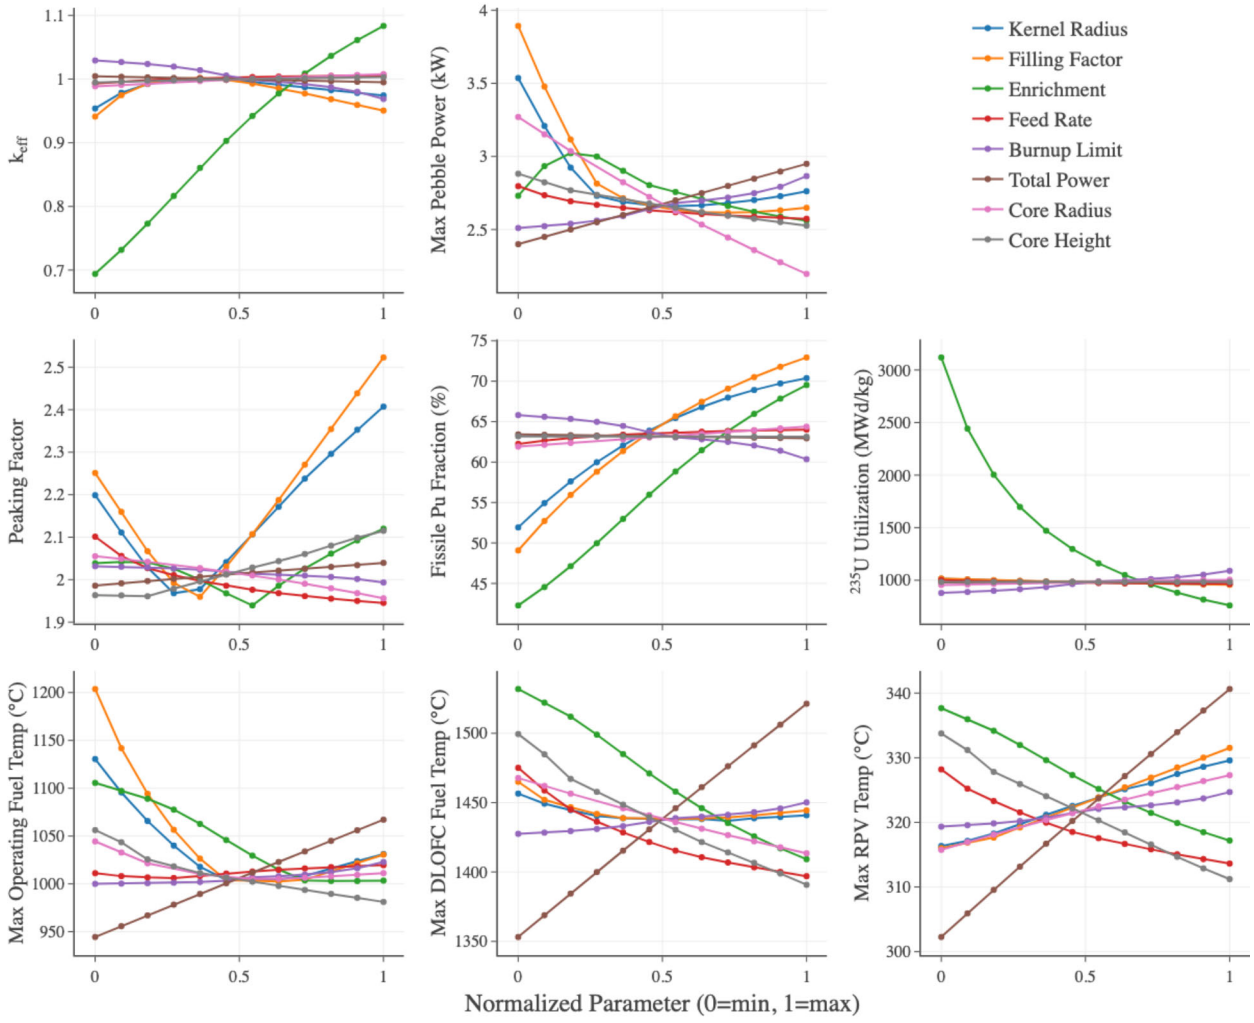

Figure 11: Quantities of interest versus design parameters from performing one-dimensional grid sampling.

This study also served two other purposes. The first is testing the robustness of the solver, ensuring that data can be produced through the entire space that the parameter is defined. From the authors' experience in this study, the fluids solver is particularly fickle; making iterative convergence criteria and initial conditions dependent on these parameters helped the robustness. Second, the results of this study serve as a visual check for the results of the global sensitivity study. If something in this grid study shows a significant sensitivity, but the global sensitivity shows a small dependence, this could indicate an issue in the training data or global sensitivity algorithm.

3.3 Producing Training Data

This section addresses the goal of producing data for performing global sensitivity and training ROMs. Here, the entire parameter space is sampled stochastically, namely with the Latin hypercube strategy presented in Reference [30]. A total of 10,000 simulations were run, which resulted in 9,150 samples—the rest were discarded due to aleatory failures of the solver.

Figure 12 show the resulting discrete probability distributions for each QoI. These histograms are quite smooth, which is a good indication that this training data is well-suited for global sensitivity analysis and ROM generation. The smoothness also promotes confidence that the discarded data from failed simulations will not detriment these processes significantly. These distributions also provide insight on how much of the design space may result in nonviable configurations. For instance, most of the configurations will be disqualified because they have too low of a k_{eff}, but max pebble power is skewed toward the right, meaning fewer configurations are realizing high powers. The two distributions for DLOFC temperatures are visually close to normal distributions, which suggests that these QoIs are largely linearly dependent on two or more parameters. This observation indicates that the temperature QoIs may not need complex ROMs to emulate their behavior.

3.4 Global Sensitivity Analysis

Sensitivity analysis searches for the variables that have a significant effect on the analysis outcome and quantifies the effect of perturbed parameters in a model. A parameter with a small sensitivity metric on a QoI typically means that the QoI is not significantly affected by the perturbation of the parameter. Higher-order sensitivity metrics can also quantify how the parameter affects the QoI, for example, if it is a nonlinear dependence, covariance with other parameters, etc. Sensitivity metrics can be classified into two categories: local and global. Local sensitivity is essentially the derivative of the QoI with respect to the parameter at a certain location in the parameter space. Local sensitivity can be useful for cases where the design space has been localized to a relatively small domain to gain insight on the parameterized system in that region. This work will focus on global sensitivity analysis, which computes metrics that give a sense of the effect of parameters for the entire parameter space. In the context of this work,

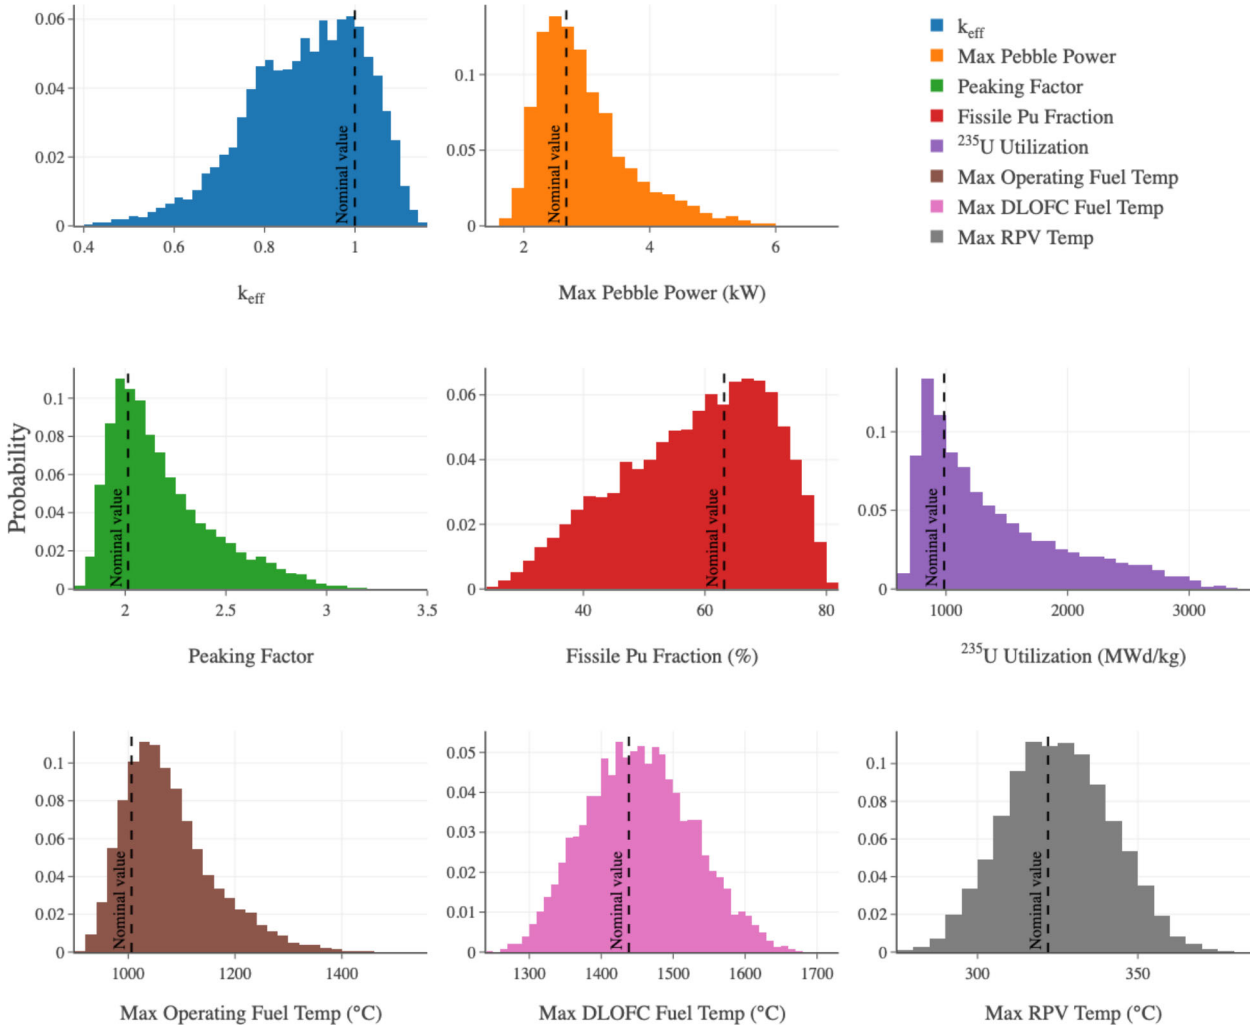

Figure 12: Discrete probability density functions of quantities of interest from Latin hypercube sampling.

global sensitivity is useful for filtering parameters from the design space. Since the ROM accuracy and performance suffer heavily from large parameter spaces, removing less impactful parameters from the ROM generation can help significantly.

In this work, the MOOSE-STM was used to perform the analysis [52], which includes several relevant techniques: Sobol sampling [48], Morris screening [34], and PCE [58]. We chose PCE for this work since the method is robust and efficient for relatively small parameter spaces, as is the case with this study. It might prove advantageous in later studies with the more parameters to use Morris screening as it is much more efficient in high-dimensional spaces.

PCE is a type of surrogate modeling technique where a QoI that is dependent on input parameters is expanded as a sum of orthogonal polynomials. Given a QoI Q dependent on a set of parameters $\vec{\xi}$, the PCE expansion is:

$$Q(\vec{\xi}) \approx \hat{Q}(\vec{\xi}) = \sum_{i=0}^{P} q_i \Phi_i(\vec{\xi}), \tag{3}$$

where P is the multidimensional polynomial order, q_i are coefficients that are to be computed, and Φ_i are the multidimensional polynomials, which are a monomial product of one-dimensional polynomials based on the parameters' probability distributions:

$$\Phi_i(\vec{\xi}) = \prod_{d=1}^{D} \phi_i^d(\xi_d), \quad i = 0, ..., P, \tag{4}$$

where ϕ is the one-dimensional polynomial and D is the number of parameters. The coefficients q_i were computed using a least-squares fit [20]. Techniques like Smolyak sparse quadrature [17] may require fewer training points, but we found using the data already generated for ROM training was more efficient.

Because the basis functions (Φ_i) are orthogonal, the expansion in Equation (3) has several convenient properties, including the ability to compute Sobol indices analytically:

$$\text{Total Sobol index:} \quad S_{\mathrm{T},d} = \sum_{i=0}^{P} q_i^2 \left\langle \phi_i^d, \phi_i^d \right\rangle, \quad d = 1, ..., D, \tag{5}$$

where the $\langle \cdot, \cdot \rangle$ notation is the inner product of the functions:

$$\langle a(\xi), b(\xi) \rangle \equiv \int_{-\infty}^{\infty} a(\xi) b(\xi) f(\xi) d\xi, \tag{6}$$

where $f(\xi)$ is the probability distribution specific to the dimension and parameter.

Training the PCE surrogates for the GPBR200 model used a maximum polynomial order of 4, which is a order high enough to prevent underfitting while still being tractable to compute. Future work could consider adaptive procedures for the polynomial space, like in Reference [5]. Figure 13 shows the resulting total Sobol indices, computed from the trained PCE surrogates. These results match well with the intuitions gleaned from the one-dimensional grid study, with enrichment, kernel radius, and filling factor being the most influential parameters for most of the QoIs. Furthermore, when the sum of the total indices for a given QoI is greater than one, this indicates nonlinear dependence on the parameters, including cross terms. This is particularly evident in k_{eff}, max pebble power, peaking factor, and operating fuel temperature. Another useful observation is that some QoIs are relatively independent of certain parameters. This suggests that building ROMs for these quantities may not require full parameterization. For instance, k_{eff} does not depend heavily on feed rate, total power, core radius, or core height, and ^{235}U utilization may only need to be parameterized with enrichment and burnup limit. Since the parameter space is relatively small in this work, this filter is not practiced, and the ROMs are generated with dependence on all the parameters. However, as more design parameters are included in future studies, this type of filtering may become necessary.

4. REDUCED-ORDER MODELING

The construction of many ROMs have the same basic procedure. First is to define the full-order model (FOM), the parameters that are either uncertain or unknown, and the QoIs. These three entities can be described by the following highly simplified equation:

$$\vec{Q}(\vec{\xi}) = \mathcal{A}(\vec{\xi}), \tag{7}$$

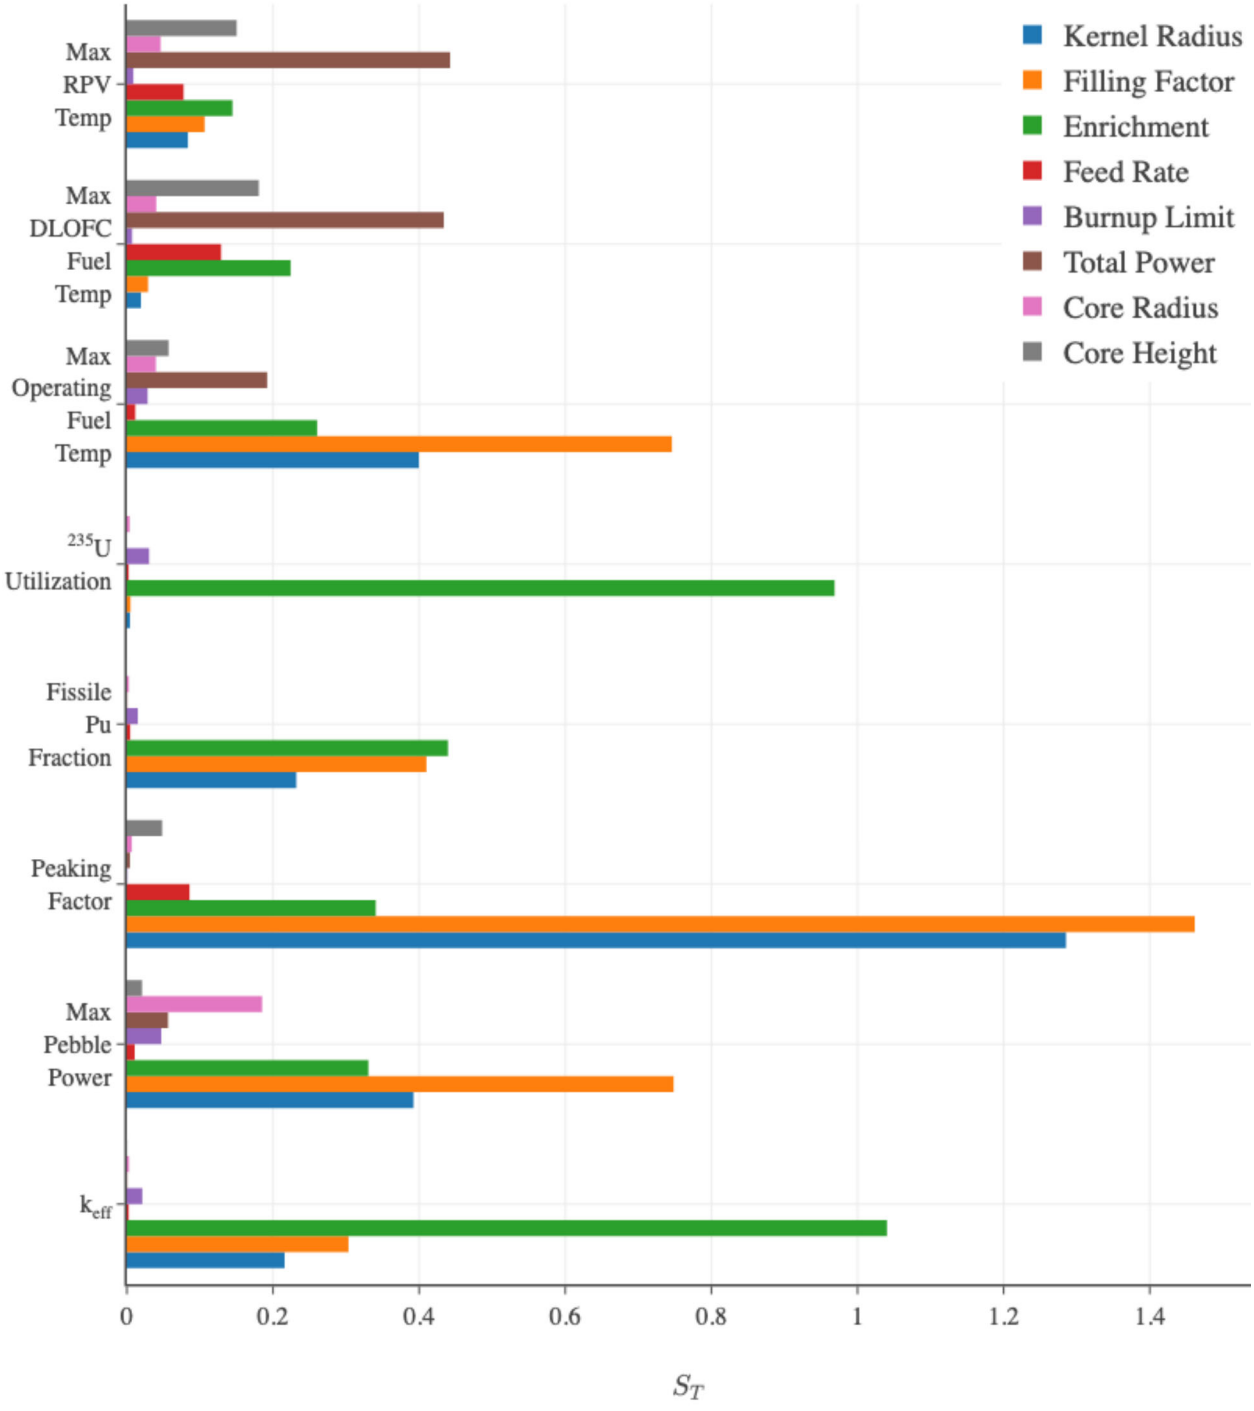

Figure 13: Total Sobol indices for each parameter and QoI calculated from polynomial chaos surrogate.

where $\mathcal{A}$ is an operator representation of the FOM, $\vec{\xi}$ is the set of parameters, and $\vec{Q}$ is a set of QoIs. For the problem presented in this work, $\mathcal{A}$ is the multiphysics GPBR200 model (Section 2), $\vec{\xi}$ are the design parameters (kernel radius, filling factor, etc.), and $\vec{Q}$ are the computed k_{eff}, max pebble power, etc. (Section 3.1).

The next step is defining a training set of parameter points $\xi = [\vec{\xi}_1, \vec{\xi}_2, ..., \vec{\xi}_N]^T$. There are numerous ways to define this set depending on the type of ROM, the number of parameters, constraints on the number of model evaluations, etc. There are also methods that adaptively choose parameter points ($\vec{\xi}_i$) during the generation process [4]. Once the training set has been defined, the FOM is evaluated N times, which results in a QoI dataset $\mathbf{Q} = [\vec{Q}_1, \vec{Q}_2, ...\vec{Q}_N]$. Section 3.3 describes the generation of training data used in this study. This data, along with the training set, is then used to generate the ROM ($\hat{\mathcal{A}}$), which obtains a modified version of Equation (7):

$$\hat{\vec{Q}}(\vec{\xi}) = \hat{\mathcal{A}}(\vec{\xi}) \tag{8}$$

The benefits of generating the ROM stem from the fact that N is much smaller than the number of FOM evaluations needed for the optimization procedure and that $\hat{\mathcal{A}}$ is much easier and faster to evaluate than $\mathcal{A}$. However, the extent of these benefits is largely correlated to how well the ROM emulates the FOM; we will call this the ROM error. A well formed ROM will be able to estimate this error so the analyst can determine if more training is needed to reduce it.

Finally, the generated ROM can be used for a variety of purposes:

- Prognosis—Simply evaluating the ROM with a set of parameters to get an estimate of the QoIs: $\hat{\vec{Q}} = \hat{\mathcal{A}}(\vec{\xi})$, which is useful for ROM validation and error estimation by comparing the ROM evaluation result with the FOM evaluation result.

- Diagnosis—Finding a set of parameters that best match a given set of QoIs: $\xi = \hat{\mathcal{A}}^{-1}(\vec{Q})$, which is useful if the the QoIs are known and a set of parameters is needed. This is less relevant in the context of optimization but is very important for digital twin development.

- **Optimization**—Finding a set of parameters that minimizes a form function dependent on the QoIs: $\vec{\xi} = \arg\max_{\vec{\xi}} f(\vec{Q}(\vec{\xi}))$. This usage is the focus of this work.

In this work, training and evaluating ROMs is done using the MOOSE-STM; which is

explained in Section 2.1.3 of Reference [52]. The following subsections describe the chosen ROM methodologies investigated in this work, including PR, GP, and ANN.

4.1 Polynomial Regression

The PR ROM in the MOOSE-STM is essentially a multivariate linear regression, reliant on the ordinary least squares method (with the option to perform ridge regression). This ROM expresses the QoI dependence on a set of parameters with a linear expansion of multidimensional polynomials:

$$\hat{\vec{Q}}(\vec{\xi}) = \sum_{i=1}^{P} \vec{q}_i \Phi_i(\vec{\xi}), \tag{9}$$

where P is a user-defined polynomial order, $\vec{q}$ are scalar coefficients the size of the number of QoIs, and Φ_i are separable multidimensional polynomials. The PR ROM technique is simple, versatile, and scalable. First, the data from the training set is assumed to be fully known (without any uncertainty), so the QoI error from running the FOM is not considered. Second, the parameters are assumed to be uncorrelated, this can cause issues when two parameters are not independent (i.e., if the feed rate has some correlation with the burnup limit). As such, the actual parameters to fit the ROM must be chosen carefully. Third, the variance in the ROM (error estimation) is assumed constant in the parameter space, so regions in the training set with lots of data are assumed to have the same error as regions with little data. Finally, because of its generality and simplicity, PR may require a large training set to generate an accurate ROM.

PR is a good point of reference for evaluating the performance of other methods. It is significantly cheaper to train, evaluate, and store than other model types. However, it can also be very limited—PR can perform poorly for responses with nonlinear parameter dependence and can easily fall victim to overfitting for high-dimensional inputs. Due to its simplicity, there are not many parameters to tune. Although the MOOSE-STM supports ridge regression, ordinary least-squares is preferred due to the additional tuning of the penalty parameter in ridge regression.

4.2 Gaussian Process

This section gives a description of GP regression within the context of this work; we refer the interested reader to Reference [46] for more background. GP regression is a nonlinear ROM technique, which is significantly different than polynomial-based methodologies. The formulation of a GP ROM is quite complex, but in overly simplistic terms, GP modeling is driven by the idea that training points that are "close" in their parameter space will be "close" in their QoI space. Closeness in the parameter space is driven by the covariance function $k(\vec{\xi}, \vec{\xi}')$. The GP model consists of an infinite collection of functions, all of which agree with the training and observation data. Importantly, the collection has closed forms for second-order statistics (mean and variance). When used as a surrogate, the nominal value is chosen to be the mean value. The method can be broken down into two steps: defining the prior distribution and then conditioning the observed data. A GP is a (potentially infinite) collection of random variables, such that the joint distribution of every finite selection of random variables from the collection is a Gaussian distribution.

$$\mathcal{GP}(\mu(\vec{\xi}), k(\vec{\xi}, \vec{\xi}')). \tag{10}$$

Analogous to a multivariate Gaussian being completely defined by its mean vector and covariance matrix, a GP is completely defined by its mean and covariance functions. The (potentially) infinite number of random variables within the GP correspond to the (potentially) infinite points in the parameter space that our ROM can be evaluated at. While the only apparent decision in the above formulation is the choice of covariance function, most covariance functions will contain hyper-parameters of some form, which need to be selected. Determining the value of these hyper-parameters manually is a terribly inefficient task; therefore, the MOOSE-STM uses a tuning procedure to set these values automatically during the training process. GP is a powerful ROM technique with its versatility, flexibility, and generality.

The choice of covariance function is an essential component in defining a GP regression. This work uses the common squared-exponential covariance kernel:

$$k(\vec{\xi}, \vec{\xi}') \equiv \sigma_f^2 \exp\left(-\frac{r_l(\vec{\xi}, \vec{\xi}')^2}{2}\right) + \sigma_n^2 \delta_{\vec{\xi}, \vec{\xi}'}, \tag{11a}$$

where,

$$r_l(\vec{\xi}, \vec{\xi}') \equiv \sqrt{\sum_{i=1}^{N} \left(\frac{\vec{\xi}_i - \vec{\xi}_i'}{l_i} \right)^2}. \tag{11b}$$

Some important remarks about this function:

- The function is *translation invariant*—it only depends on the difference between points $\vec{\xi} - \vec{\xi}'$.
- The function is *isotropic*—it only cares about the magnitude of $\vec{\xi} - \vec{\xi}'$ and not the direction.

Many common covariance functions have these properties (including all of the options in the MOOSE-STM). Although these properties are advantageous in quantifying similarity, they have the side effect that predictors near each other will have highly correlated covariances with the rest of the data. This is an issue because several operations in training and evaluation with GPs involve inverting the dense matrix:

$$\mathbf{K}(\_, \_) \equiv \begin{bmatrix} k(\vec{\xi}_1, \vec{\xi}_1) & \cdots & k(\vec{\xi}_1, \vec{\xi}_m) \\ \vdots & & \vdots \\ k(\vec{\xi}_m, \vec{\xi}_1) & \cdots & k(\vec{\xi}_m, \vec{\xi}_m) \end{bmatrix}. \tag{12}$$

Because Equation (11) is translation invariant and isotropic, this matrix can easily become very near singular if there are too many samples near each other. To handle this, the implementation in the MOOSE-STM performs all necessary linear solves with a stable direct method using the Cholesky factorization.

The covariance function, see Equation (11), has a number of hyper-parameters ` that must be optimized:

- $\vec{l} \in \mathbb{R}^n_{>0}$, the length scales for each predictor
- $\sigma_n^2 \in \mathbb{R}_{>0}$, the noise variance in the inputs
- $\sigma_f^2 \in \mathbb{R}_{>0}$, signal variance in the response.

The problem of choosing these parameters can be posed as an optimization problem, where the objective function is the *marginal likelihood*. Due to the computational difficulty in computing gradients for these parameters, direct gradient-based methodologies—like conjugate

gradient [15]—are intractable for large datasets like the one generated for this work. Therefore, this work implemented a stochastic optimization routine, namely an Adam optimizer [21].

4.3 Artificial Neural Network

The ANN capability in the MOOSE-STM is relatively new and only simple feed-forward neural networks have been implemented. The implementation uses the underlying objects imported from `libtorch` (C++ API of `pytorch`) [43]. For a more detailed introduction to neural networks, we refer the reader to Reference [36]. The architecture of a simple feed-forward neural network is presented in Figure 14. The first layer on the left and right are referred to as input and output layers, respectively, while the layers between them are the hidden layers.

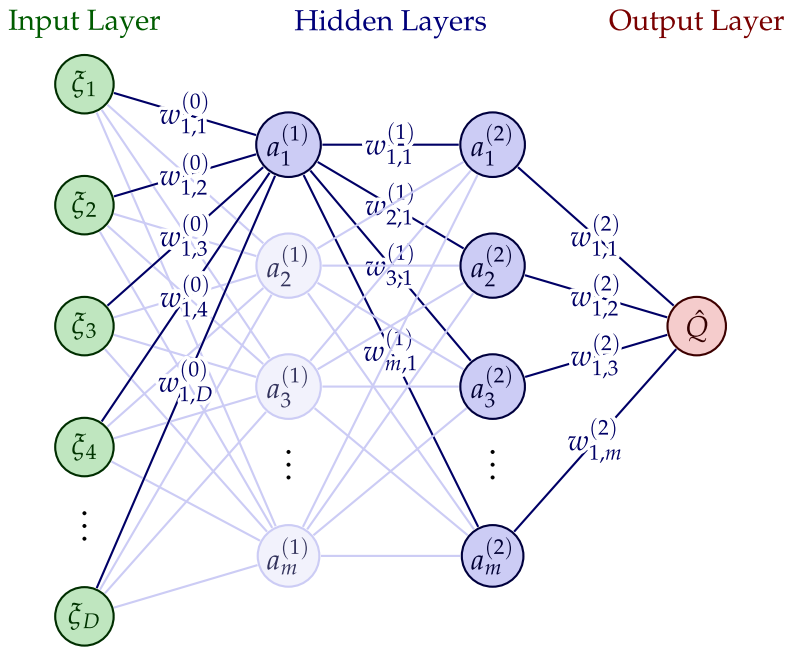

Figure 14: The architecture of the simple feed-forward ANN with two hidden layers in the MOOSE-STM.

We see that the output ($\hat{Q}$) of the neural network can be expressed as function of the inputs ($\vec{\xi}$) and the corresponding model parameters—weights $w_{i,j}$, organized in weight matrices $\mathbf{W}$, and biases, b_i organized in the bias vector ($\vec{b}$)—in the following nested form:

$$\vec{a}^{(0)} = \vec{\xi}, \tag{13a}$$

$$\vec{a}^{(i+1)} = \sigma\left(\mathbf{W}^{(i)}\vec{a}^{(i)} + \vec{b}^{(i)}\right), \tag{13b}$$

$$\hat{Q} = \hat{a}^{(3)} = \sigma(\mathbf{W}^{(2)}\sigma(\mathbf{W}^{(1)}\sigma(\mathbf{W}^{(0)}\vec{\xi} + \vec{b}^{(0)}) + \vec{b}^{(1)}) + \vec{b}^{(2)}), \tag{13c}$$

where σ denotes the activation function. In this implementation, no activation function is applied on the output layer. It is apparent that the real functional dependence (target function) between the inputs and outputs is approximated by the function in Equation (13). As in most cases, the error in this approximation depends on the smoothness of the target function and the values of the model parameters. The weights and biases in the function are determined by minimizing the error between the approximate outputs of the neural network corresponding reference (training) values over a training set. Optimization of the biases and weights are determined via stochastic gradient descent. The learning parameters for the algorithm were chosen based on trial and error of a single-layer ANN of k_{eff}, which is summarized in Table 5.

Table 5: Learning parameters used for neural network training.

Parameter	Value
Number of training epochs	1,000
Relative loss stopping criteria	10^{-4}
Learning rate	10^{-3}
Number of parallel batches	24
Number of samples per batch	10

There are several hyperparameters to choose from when designing a neural network: number of hidden layers, number of neurons per hidden layer, and activation function for each layer. In this work, the rectified linear units (ReLU) activation function was used in all cases [37]. To search for an optimal model, a full grid-search over possible network configurations was performed. Each QoI was tested with one to three hidden layers with various numbers of neurons per layer. Table 6 summarizes the network configurations tested in this study.

Table 6: Neural network configurations considered.

	Number of Neurons per Layer		
Number of Layers	Layer 1	Layer 2	Layer 3
1	{4, 8, 16, 32, 64}	0	0
2	{4, 8, 16, 32, 64}	{4, 8, 16, 32}	0
3	{4, 8, 16, 32, 64}	{4, 8, 16, 32}	{4, 8, 16}

4.4 Reduced-Order Model Comparison

To compare each of the ROM methodologies, the data generated in Section 3.3 is split into a training set and test set. The training set is used to actually train the ROMs. The test set is used as an error estimation, where the ROMs are evaluated with parameter values in this set, and the resulting outputs are compared with the QoI values in the set. Relative root mean-square error (RMSE) is the error metric for this comparison, defined as,

$$\text{Relative RMSE} \equiv \frac{N_{\text{test}}}{\sum_{i=1}^{N_{\text{test}}} Q_i} \sqrt{\frac{1}{N_{\text{test}}} \sum_{i=1}^{N_{\text{test}}} \left(\hat{Q}_i - Q_i\right)^2}, \tag{14}$$

where N_{test} is the number of test points. In this study, 3,000 samples were put in the training set, while the rest (6,150 samples) were put in the test set. In most ROM applications, the training set encompasses more of the available data. However, a smaller training set was chosen for this study for two reasons. First, the GP and ANN were computationally difficult to train with the full dataset. Future work will include implementing faster training algorithms for these models and potentially using vectorized hardware (graphics processing units). Second, generating a large dataset, like this one, may not be practical for other reactor models, so training on this reduced dataset may be more informative for a wider range of applications.

Each of the three ROM methods described earlier in this section were tested. For PR, the maximum order of the polynomial was chosen to be 4. For ANN, the full grid of configurations were tested—described in Appendix A—where the top performing configuration (lowest RMSE) was chosen.

Figure 15 compares the three ROM methodologies. Upon first inspection, ANN seems to be the best overall performer, especially with pebble power, peaking factor, and operating fuel temperature. From the sensitivity study, these quantities had significant nonlinear behavior, which the ANN is able to capture well. Results from models for Pu fraction, DLOFC fuel temperature, and RPV temperature are quite similar across the methods. These quantities were shown to have a linear dependence on the parameters, which all of these methods capture well. ANN appears to struggle with U utilization, probably due to overfitting since this quantity has a minimal dependence on most of the parameters. Finally, GP had the best performance for k_{eff}.

The nonlinear dependence this quantity showed on the parameters appears to make PR a lower performing method, while the small dependence on half of the parameters made ANN overfit the data. These results highlight the importance of filtering the parameter space when generating ROMs, which should be practiced in future work.

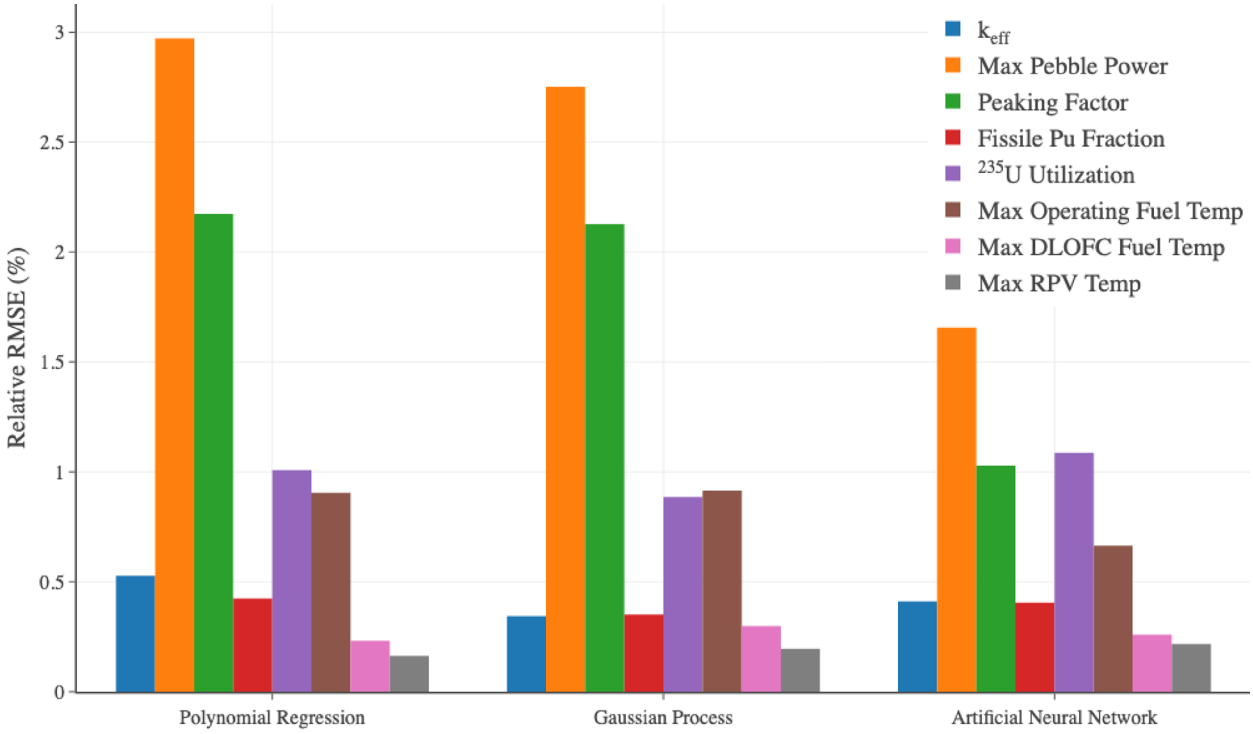

Figure 15: Comparing the performance of ROM methodologies for each QoI.

5. OPTIMIZATION STUDY

In this section, the dataset and generated ROMs are utilized for optimizing the design parameters. The optimization involves a highly simplistic procedure of stochastically sampling the models and choosing the best performing sample based on performance metrics. This procedure typically requires many model evaluations to obtain a good estimate of the optimized configuration; however, it is robust (avoiding local optima) and scalable, and the ROMs are extremely fast to evaluate. So to compare to the 10^4 sample dataset from the FOM, each ROM was evaluated 10^6 times over the entire design space. The performance metric involves maximizing ^{235}U utilization, while ensuring constraints on the other QoIs are satisfied.

5.1 Optimization Problem Definition

In any engineering design, there are limits or constraints in which the system is subject to operate within, typically based on safety considerations and material failure mechanisms. For nuclear reactor designs, regulatory bodies impose stringent requirements involving operational safety, accident tolerance, and proliferation resistance. Due to the novelty of advanced reactor designs—like HTGRs—these requirements can be especially conservative. In addition to licensing constraints, designers must consider operational viability. Therefore, the objective function defining the design optimization must include relevant constraints. The following lists the constraints imposed on the GPBR200 model in this work:

- $\mathbf{k_{eff} > 0.99200}$—This constraint imposes a sufficient neutron economy for the reactor to be critical. The constraint value is based on the control rod study, shown in Figure 4. This value stipulates that the core has enough excess reactivity to be critical if the control rod was fully withdrawn.

- **Max pebble power $<$ 4.5 kW**—Pebbles reaching this power density could have significant TRISO failure, releasing fission products.

- **Peaking factor $<$ 2.2**—High peaking factor can degrade the reactivity control of the reactor and cause unwanted fluctuations in fuel temperature as pebbles traverse the core.

- **Fissile Pu fraction $<$ 70%**—The fissile fraction of plutonium bred in the core affects the proliferation resistance of the design. A sufficient presence of non-fissile isotopes, like ^{240}Pu and ^{238}Pu, make the material unusable for weapon construction [51].

- **Max operating fuel temperature $<$ 1100°C**—High fuel temperatures during the equilibrium-core phase can cause unacceptably high temperatures during accident conditions and cause large fluctuations in fuel temperature during recycling.

- **Max DLOFC fuel temperature $<$ 1600°C**—This limit is imposed to prevent significant TRISO failures causing fission gas release.

- **Max RPV temperature $<$ 350°C**—High temperatures in the RPV could cause the vessel to fail, resulting in radioactive contamination.

The performance metric for the optimization is ^{235}U utilization, since this value is directly related to cost of procuring fuel for reactor operation. Although, other factors related to cost could be considered in future work, including enrichment costs, spent fuel storage, pebble handling system, etc.

In the context of optimization using ROMs, adapting the constraints becomes crucial due to the inherent error associated with the lower-fidelity representation provided by these models. The goal is to ensure that the original constraints are still met when executing the FOM with the parameter values derived from optimizing the ROM. This issue becomes evident when considering that k_{eff} and ^{235}U utilization are inversely related. Consequently, optimizing ^{235}U utilization might inadvertently lead to a reduction in k_{eff} close to the constrained limit. This concern was substantiated during the initial phases of this study. To counteract this challenge, the constraints are modified in a way that maintains the satisfaction of the original constraints, accounting for the estimated discrepancy between the ROM and FOM outcomes with a certain level of confidence.

Due to the absence of a robust analytical error estimation for all ROM methods, a pragmatic approach was adopted. This approach utilizes residuals in the test set—the discrepancy between predicted ($\vec{\hat{Q}}$) and actual ($\vec{Q}$) outcomes—to construct a discrete error distribution. This distribution is assumed to be constant across the entire design space. Subsequently, the constraints are adjusted to ensure a 90% confidence level, based on this constant distribution, that a prediction made with these adjusted constraints will satisfy the original constraint. Table 7 presents the modified constraints for each model.

This strategy acknowledges the inherent uncertainty in the ROM predictions and ensures that the optimization process does not inadvertently deviate from the original design constraints, thus fostering more reliable and consistent design improvements.

5.2 Optimization Results

Table 8 shows the fraction of the generated set that produced viable configurations. The "Full Model" results used the full training set, described in Section 3.3. The ROM results used the dataset generated from evaluating all the ROMs with 10^6 samples. We observed that the fraction from the ROM dataset is smaller than the training set due to the adjusted constraints accounting

Table 7: Constraints used when performing optimization on each type of model.

Constraint	Full Model	PR	GP	ANN
k_{eff}	0.99200	0.99719	0.99544	0.99691
Max pebble power	4.500	4.406	4.415	4.445
Peaking factor	2.200	2.154	2.156	2.182
Fissile Pu fraction	70.00	69.72	69.76	69.72
Max operating fuel temp	1100	1089	1089	1091
Max DLOFC fuel temp	1600	1597	1597	1597
Max RPV temp	350.0	349.5	349.5	349.2

for ROM error estimation. As the ROM quality improves, indicated by the relative RMSE results in Figure 15, the fraction from imposing specific constraints increases. However, the total fraction is practically identical across the models. A 13% probability of choosing a viable configuration is relatively strict and could cause difficulties for stochastic optimization algorithms, like the particle swarm optimization presented in Reference [59]. Limiting the design space would increase this fraction and improve these algorithms.

Table 8: Fraction of design space producing viable configurations when applying each constraint.

Constraint	Full Model	PR	GP	ANN
k_{eff}	24.9%	23.6%	24.2%	23.7%
Max pebble power	95.2%	94.7%	94.5%	94.7%
Peaking factor	62.7%	56.8%	55.0%	59.5%
Fissile Pu fraction	80.3%	78.3%	78.7%	78.7%
Max operating fuel temp	67.7%	64.1%	64.3%	64.4%
Max DLOFC fuel temp	97.2%	97.0%	96.9%	97.0%
Max RPV temp	93.0%	92.7%	92.7%	92.5%
Total	13.6%	11.5%	11.3%	11.7%

Table 9 presents the observed optimal configurations and resulting QoIs from the datasets. These configurations have a relatively diverse set of design parameters, showcasing the complexity of this optimization problem. It appears that increasing the kernel radius and decreasing the filling factor provides a better configuration than the nominal; however, the resulting density of fuel stays relatively the same. Enrichment, feed rate, and burnup limit seem to have an interplay, where increasing enrichment allows decreasing feed rate and increasing burnup limit, or vice versa. Although, the feed rate varies significantly between the configurations, indicating QoI insensitivity to this parameter. Minimizing the total power appears to produce the optimal configuration. This observation corresponds to the results in Figure 11,

which shows that power negatively impacts the constrained QoIs, while U utilization is relatively insensitive to it.

The resulting QoI values presented in Table 9 show the impact that optimization has on the performance of the design. Using just the full-order dataset, the fuel utilization increased by $\sim$9% and using the best performing ROM increased it by $\sim$10%. Improving the ROM would continue to improve the optimal configuration's performance by loosening the prescribed constraints. All the constrained quantities are far from their imposed constraint, except for k_{eff}, which appears to be as low as possible in order to maximize fuel utilization. Furthermore, the ROMs do overpredict k_{eff} in the optimized region. This exemplifies the importance of incorporating error estimation when imposing constraints on ROM-based optimization.

6. CONCLUSIONS

In this work, a design optimization workflow was developed and applied to a representative gas-cooled PBR. The model utilizes MOOSE-based applications—Griffin, Pronghorn, and Bison—to fully couple the neutronics, depletion, TH, and pebble thermal physics of the problem. The workflow began by defining design parameters and performance-related QoIs. The local relationship between these parameters and QoIs was analyzed by performing a one-dimensional grid study, which provided valuable insight into the design's performance as parameters were changed. The full design space was then explored by producing a relatively large dataset of configurations and performing a global sensitivity analysis. This study provided a broader and more quantifiable insight on the relationship between parameters and QoIs, ranking the impact of each parameter and quantifying the degree of nonlinearity of each QoI. Due to the computational difficulty of performing optimization on a multiphysics model, the generation of fast-evaluating ROMs was explored, specifically PR, GP, and ANN. These models were trained on a small portion of the dataset, tested on the rest, and compared using the difference between predictions and full model realizations. The ROMs were then applied to a constrained single-objective optimization problem to maximize fuel utilization while abiding by safety-related limits. The optimization technique used was a simple sample-and-search process, whereby producing a multitude of data and picking the configuration with the highest fuel utilization that satisfied the constraints.

Table 9: Resulting parameters and quantities of interest from optimal configurations. The ROM evaluations present the difference in the ROM prediction and full-order result with same configuration. The full-order result can be computed as $\hat{Q}(1+\delta)$, where δ is the relative difference shown.

Parameter	Nominal	Full Model	PR	GP	ANN
Kernel radius	0.2125	0.28348	0.25128	0.2979	0.26386
Filling factor	9.344	7.2802	7.5797	6.0249	6.933
Enrichment	15.5	16.025	15.774	15.067	14.276
Feed rate	1.5	1.3136	2.1634	1.046	2.2308
Burnup limit	147.6	163.98	163.2	158.61	147.74
Total power	200	185.04	189.81	182.53	190.45
Core radius	1.2	1.2767	1.2863	1.2963	1.2947
Core height	8.93	9.6428	9.179	9.5896	9.5479

Quantity	Full Order Evaluation		ROM Evaluation with Relative Error					
k_{eff}	0.99961	0.994835	0.99839	- 0.0515%	0.99742	- 0.247%	0.99941	- 0.705%
Max pebble power	2.6742	2.2731	2.2518	- 0.865%	2.1967	+ 7.59%	2.0952	+ 0.723%
Peaking factor	2.0145	2.1006	1.972	- 2.63%	2.0559	- 1.73%	1.935	+ 0.132%
Fissile Pu fraction	63.158	63.044	62.053	+ 0.216%	59.14	+ 1.36%	60.893	- 0.0201%
^{235}U utilization	**984.18**	**1070**	**1063.1**	**- 0.183%**	**1087.1**	**- 1.23%**	**1084.9**	**- 0.912%**
Max operating fuel temp	1006.2	941.72	989.46	- 0.561%	918.66	+ 6.16%	977.56	+ 0.741%
Max DLOFC fuel temp	1438.4	1341.7	1347.8	- 0.00344%	1358	- 0.675%	1328.3	+ 0.298%
Max RPV temp	321.97	309.31	311.36	- 0.0235%	311.06	- 0.758%	306.99	+ 0.148%

The representative gas-cooled PBR model, named GPBR200, involved both a fully-coupled equilibrium-core calculation and a DLOFC accident transient. The design parameters chosen for the study involved fuel management, operational aspects, and core geometry. The QoIs were chosen to measure safety margins, operational viability, and economic performance. The sensitivity analysis exemplified the complex relationships between these parameters and QoIs, providing insight on the potential difficulty in generating ROMs and performing optimization. These sensitivity results have proven to be invaluable both within the scope of this work and outside for understanding PBR behavior under various design configurations.

The ROM methodologies investigated in this work included simple, nonlinear, and complex data-driven techniques. The performance comparison between the models proved to be insightful for narrowing down a choice of methodology. The analysis found that complex techniques like ANN struggled with overfitting QoIs that had linear or limited dependence on parameters, while showing superior performance for QoIs with nonlinear behavior. Surprisingly, GP did not perform as well as expected compared to ANN; although, a further tuning of the method could provide improvements or it could show better performance on smaller datasets. Furthermore, GPs are unique as they are capable in supplying an uncertainty estimate with their predictions, which was not utilized in this work but could prove useful in similar studies.

The optimization study illustrated the impact that optimization had on the performance of the design. The optimized model showed a $\sim$10% improvement on fuel utilization over the engineering-judgement-based nominal configuration. The study also revealed a subtlety when using ROMs for constrained optimization. Mainly based on the QoI dependence on enrichment, maximizing fuel utilization meant getting k_{eff} as close to the constraint as possible. However, the ROMs showed to overpredict k_{eff} near the optimal configuration. As a result, the optimal configuration found with the ROMs did produce a viable configuration when simulated with the multiphysics model. Therefore, an error estimation was formulated for each ROM and the constraints were adjusted such that, with some confidence, the predicted value would satisfy the constraint. This degraded the performance of the resulting configurations but could be improved by increasing the ROM quality.

There are avenues of future work for every aspect presented in this paper. First, the model can be improved to better represent the physics: using online cross section generation [40],

including burnup and radiation damage to thermal properties, and using higher-fidelity transport formulations for the neutronics. Furthermore, performing this sort of study on running-in calculations would be a highly advantageous for practical reactor design, as optimization on that portion of the reactor lifetime is arguably more important. The ROM portion of this work is worthy of further exploration. The usefulness of ROMs could be improved by investigating a broader range of techniques, including other neural network formulations, choices of GP parameters, and physics-based ROMs. The error estimation of ROMs was also proven to be important, so methods to improve the estimation would significantly for constrained optimization. Finally, using more advanced optimization algorithms is something planned in future work, following the work in References [19, 59] and outlined in Reference [55].

ACKNOWLEDGMENTS

We would like to acknowledge Colin Brennan, Peter German, Vincent Labouré, Gerhard Strydom, and Adam Zabriskie. Colin's help with the initial ROM study was invaluable to the results presented in this work. Peter's help with navigating Pronghorn helped produce a robust model. Vincent provided the cross sections for the near-void region in the upper cavity. Gerhard's expertise in the field of uncertainty quantification and sensitivity analysis of HTGRs is what realized this work. And Adam helped improve the pebble and TRISO model by navigating us through the material definitions in Bison.

This work was funded under the Advanced Reactor Technologies (ART) Nuclear Energy Advanced Modeling and Simulation (NEAMS) programs managed by the Department of Energy Office of Nuclear Energy. This research made use of the resources of the High Performance Computing Center at Idaho National Laboratory, which is supported by the Office of Nuclear Energy of the U.S. Department of Energy and the Nuclear Science User Facilities under Contract No. DE-AC07-05ID14517. Accordingly, the U.S. government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. government purposes.

REFERENCES

[1] A. Hébert , Scattering reduction of the double-heterogeneity treatment in Dragon, *Nuclear Science and Engineering*, *160(2)*, 261–266, doi:10.13182/NSE160-261TN, 2008.

[2] Balestra, P., S. Schunert, R. W. Carlsen, A. J. Novak, M. D. DeHart, and R. C. Martineau, PBMR-400 benchmark solution of exercise 1 and 2 using the MOOSE based applications: MAMMOTH, Pronghorn, in *Proceedings of PHYSOR 2020: Transition to a Scalable Nuclear Future*, p. 06020, Cambridge, United Kingdom, 2020.

[3] Balestra, P., R. Stewart, S. Schunert, and D. Reger, Improved decay heat prediction during HTGR transients scenarios, *Tech. Rep. INL/RPT-22-68845*, Idaho National Laboratory, 2022.

[4] Barron, A. R., A. Cohen, W. Dahmen, and R. A. DeVore, Approximation and learning by greedy algorithms, *The Annals of Statistics*, *36*(1), 64 – 94, doi:https://doi.org/10.1214/009053607000000631, 2008.

[5] Blatman, G., and B. Sudret, An adaptive algorithm to build up sparse polynomial chaos expansions for stochastic finite element analysis, *Probabilistic Engineering Mechanics*, *25*(2), 183–197, doi:https://doi.org/10.1016/j.probengmech.2009.10.003, 2010.

[6] Boer, B., J. Kloosterman, D. Lathouwers, and T. van der Hagen, In-core fuel management optimization of pebble-bed reactors, *Annals of Nuclear Energy*, *36*(8), 1049–1058, doi:https://doi.org/10.1016/j.anucene.2009.06.008, 2009.

[7] Brits, Y., F. Botha, H. van Antwerpen, and H. W. Chi, A control approach investigation of the Xe-100 plant to perform load following within the operational range of 100–25–100%, *Nuclear Engineering and Design*, *329*, 12–19, doi:https://doi.org/10.1016/j.nucengdes.2017.11.041, 2018.

[8] Brown, D. A., 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*, 1–142, doi:10.1016/j.nds.2018.02.001, special Issue on Nuclear Reaction Data, 2018.

[9] Calvin, O., Depletion chain simplification using pseudo-nuclides, *Annals of Nuclear Energy*, *193*, 110,011, doi:https://doi.org/10.1016/j.anucene.2023.110011, 2023.

[10] Chong, E. K., and S. H. Żak, *Quasi-Newton Methods*, chap. 11, pp. 187–209, John Wiley & Sons, Ltd, doi:https://doi.org/10.1002/9781118033340.ch11, 2008.

[11] Chong, E. K., and S. H. Żak, *Global Search Algorithms*, chap. 14, pp. 267–295, John Wiley & Sons, Ltd, doi:https://doi.org/10.1002/9781118033340.ch14, 2008.

[12] Chong, E. K., and S. H. Żak, *Gradient Methods*, chap. 8, pp. 125–153, John Wiley & Sons, Ltd, doi:https://doi.org/10.1002/9781118033340.ch8, 2008.

[13] Cisneros Jr, A. T., Pebble bed reactors design optimization methods and their application to the pebble bed fluoride salt cooled high temperature reactor (PB-FHR), Ph.D. thesis, University of California, Berkeley, 2013.

[14] Duderstadt, J. J., and L. J. Hamilton, *Nuclear Reactor Analysis*, vol. 1, John Wiley & Sons, Inc., New York, 1976.

[15] Gardner, J., G. Pleiss, K. Q. Weinberger, D. Bindel, and A. G. Wilson, GPyTorch: Blackbox matrix-matrix Gaussian process inference with GPU acceleration, *Advances in neural information processing systems*, *31*, 2018.

[16] Gaston, D. R., et al., Physics-based multiscale coupling for full core nuclear reactor simulation, *Annals of Nuclear Energy*, *84*(1), 45–54, doi:10.1016/j.anucene.2014.09.060, 2015.

[17] Gerstner, T., and M. Griebel, Numerical integration using sparse grids, *Numerical Algorithms*, *18*(3), 1572–9265, doi:10.1023/A:1019129717644, 1998.

[18] Gontard, R., and H. Nabielek, Performance evaluation of modern HTR TRISO fuels, *Tech. Rep. HTA-1B-05/90*, Forschungszentrums Juelich, 1990.

[19] Gougar, H. D., A. M. Ougouag, W. K. Terry, and K. N. Ivanov, Automated design and optimization of pebble-bed reactor cores, *Nuclear Science and Engineering*, *165*(3), 245–269, doi:10.13182/NSE08-89, 2010.

[20] Hadigol, M., and A. Doostan, Least squares polynomial chaos expansion: A review of sampling strategies, *Computer Methods in Applied Mechanics and Engineering*, *332*, 382–407, doi:https://doi.org/10.1016/j.cma.2017.12.019, 2018.

[21] Kingma, D. P., and J. Ba, Adam: A method for stochastic optimization, *arXiv:1412.6980 [cs.LG]*, 2014.

[22] Kugeler, K., and Z. Zhang, *Modular High-temperature Gas-cooled Reactor Power Plant*, Springer-Verlag, 2019.

[23] Labouré, V., J. Ortensi, N. Martin, P. Balestra, D. Gaston, Y. Miao, and G. Strydom, Improved multiphysics model of the high temperature engineering test reactor for the simulation of loss-of-forced-cooling experiments, *Annals of Nuclear Energy*, *189*, 109,838, doi:https://doi.org/10.1016/j.anucene.2023.109838, 2023.

[24] Lee, C. H., Y. Jung, H. Park, E. Shemon, J. Ortensi, Y. Wang, V. Labouré, and Z. Prince, Griffin Software Development Plan, *Research Report INL/EXT-21-63185, ANL/NSE-21/23*, Idaho National Laboratory, Argonne National Laboratory, 2021.

[25] Leppänen, J., M. Pusa, T. Viitanen, V. Valtavirta, and T. Kaltiaisenaho, The Serpent Monte Carlo code: Status, development and applications in 2013, *Annals of Nuclear Energy*, *82*(Supplement C), 142–150, 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, 2015.

[26] Lindsay, A. D., M. Tano, G. L. Giudicelli, P. German, and S. Schunert, Improvement of numerical methods in Pronghorn, *Tech. Rep. INL/EXT-21-65482*, Idaho National Laboratory, 2021.

[27] Lindsay, A. D., et al., 2.0 - MOOSE: Enabling massively parallel multiphysics simulation, *SoftwareX*, *20*, 101,202, doi:https://doi.org/10.1016/j.softx.2022.101202, 2022.

[28] Liu, Z., K. Smith, B. Forget, and J. Ortensi, Cumulative migration method for computing rigorous diffusion coefficients and transport cross sections from Monte Carlo, *Annals of Nuclear Energy*, *112*, 507–516, doi:https://doi.org/10.1016/j.anucene.2017.10.039, 2018.

[29] Massimo, L., Chapter 9 - Burn-Up, in *Physics of High-Temperature Reactors*, edited by L. Massimo, pp. 105–126, Pergamon, doi:https://doi.org/10.1016/B978-0-08-019616-9.50016-0, 1976.

[30] McKay, M. D., R. J. Beckman, and W. J. Conover, Comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics*, *21*(2), 239–245, 1979.

[31] McLaughlin, R., A study of the differential scheme for composite materials, *International Journal of Engineering Science*, *15*(4), 237–244, doi:https://doi.org/10.1016/0020-7225(77)90058-1, 1977.

[32] Miller, G., D. Petti, J. Maki, D. Knudson, and W. Skerjanc, PARFUME theory and model basis report, *Tech. Rep. INL/EXT-08-14497 (Rev.1)*, Idaho National Laboratory, 2018.

[33] Mohammad Jaradat, M. K., S. Schunert, and J. Ortensi, Gas-cooled high-temperature pebble-bed reactor reference plant model, *Tech. Rep. INL/RPT-23-72192-Rev000*, Idaho National Laboratory, doi:10.2172/1984239, 2023.

[34] Morris, M. D., Factorial sampling plans for preliminary computational experiments, *Technometrics*, *33*(2), 161–174, doi:10.1080/00401706.1991.10484804, 1991.

[35] Mulder, E. J., and W. A. Boyes, Neutronics characteristics of a 165 MWth Xe-100 reactor, *Nuclear Engineering and Design*, *357*, 110,415, doi:https://doi.org/10.1016/j.nucengdes.2019.110415, 2020.

[36] Müller, B., J. Reinhardt, and M. T. Strickland, *Neural networks: an introduction*, Springer Science & Business Media, 1995.

[37] Nair, V., and G. E. Hinton, Rectified linear units improve restricted boltzmann machines, in *Proceedings of the 27th International Conference on International Conference on Machine Learning*, ICML'10, pp. 807–814, Omnipress, Madison, WI, USA, 2010.

[38] Novak, A., R. Carlsen, S. Schunert, P. Balestra, R. Slaybaugh, and R. Martineau, Pronghorn: A multidimensional coarse-mesh application for advanced reactor thermal hydraulics, *Nuclear Technology*, *207*(7), 1015–1046, doi:10.1080/00295450.2020.1825307, 2021.

[39] Novak, A., et al., Pronghorn theory manual, *Tech. Rep. INL/EXT-18-44453-Rev001*, Idaho National Laboratory, 2020.

[40] Ortensi, J., and P. Balestra, Initial study on cross-section generation requirements for a PBR equilibrium core, in *Proceedings of the international conference on physics of reactors-PHYSOR 2022*, pp. 1440–1453, Pittsburgh, PA, 2022.

[41] Ostertagová, E., Modelling using polynomial regression, *Procedia Engineering*, *48*, 500–506, doi:https://doi.org/10.1016/j.proeng.2012.09.545, modelling of Mechanical and Mechatronics Systems, 2012.

[42] Ougouag, A., W. Terry, and H. G. Gougar, Examination of the potential for diversion or clandestine dual use of a pebble-bed reactor to produce plutonium, in *Proceedings of the conference on high temperature reactors, HTR-2002*, Petten (Netherlands), 2002.

[43] Paszke, A., et al., Pytorch: An imperative style, high-performance deep learning library, in *Advances in Neural Information Processing Systems 32*, pp. 8024–8035, Curran Associates, Inc., 2019.

[44] 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. RISO-224*, Danish Atomic Energy Commission, 1970.

[45] Prince, Z. M., C. Brennan, M. Turkmen, P. Balestra, and G. Strydom, Reduced order models generation for HTGRs pebble shuffling procedure optimization studies, *Tech. Rep. INL/RPT-22-68865-Rev000*, Idaho National Laboratory, 2022.

[46] Rasmussen, C. E., and C. K. I. Williams, *Gaussian Processes for Machine Learning*, The MIT Press, 2005.

[47] Reger, D., E. Merzari, P. Balestra, R. Stewart, and G. Strydom, Discrete element simulation of pebble bed reactors on graphics processing units, *Annals of Nuclear Energy*, *190*, 109,896, doi:https://doi.org/10.1016/j.anucene.2023.109896, 2023.

[48] Saltelli, A., Making best use of model evaluations to compute sensitivity indices, *Computer Physics Communications*, *145*(2), 280–297, doi:10.1016/S0010-4655(02)00280-1, 2002.

[49] Schunert, S., M. Jaradat, O. Calvin, G. Giudicelli, A. Lindsay, Y. Wang, M. Tano, and S. Walker, Improvements in high-temperature gas-cooled reactor modeling capabilities in the Pronghorn code, *Tech. Rep. INL/RPT-22-69263*, Idaho National Laboratory, 2022.

[50] Schunert, S., J. Ortensi, Y. Wang, P. Balestra, M. Jaradat, O. Calvin, J. Hanophy, and L. Harbour, An equilibrium core depletion algorithm for pebble-bed reactors in the Griffin code, *Annals of Nuclear Energy*, *192*, 109,980, doi:https://doi.org/10.1016/j.anucene.2023.109980, 2023.

[51] Serfontein, D. E., E. J. Mulder, and F. Reitsma, Optimisation of deep burn incineration of reactor waste plutonium in a PBMR DPP-400 core, *Nuclear Engineering and Design*, *271*, 99–105, doi:https://doi.org/10.1016/j.nucengdes.2013.11.017, 2014.

[52] Slaughter, A. E., Z. M. Prince, P. German, I. Halvic, W. Jiang, B. W. Spencer, S. L. Dhulipala, and D. R. Gaston, MOOSE stochastic tools: A module for performing parallel, memory-efficient in situ stochastic simulations, *SoftwareX*, *22*, 101,345, 2023.

[53] Sobol', I., Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates, *Mathematics and Computers in Simulation*, *55*(1), 271–280, doi:https://doi.org/10.1016/S0378-4754(00)00270-6, the Second IMACS Seminar on Monte Carlo Methods, 2001.

[54] Stewart, R., P. Balestra, D. Reger, and E. Merzari, Generation of localized reactor point kinetics parameters using coupled neutronic and thermal fluid models for pebble-bed reactor transient analysis, *Annals of Nuclear Energy*, *174*, 109,143, doi:https://doi.org/10.1016/j.anucene.2022.109143, 2022.

[55] Stewart, R. H., T. S. Palmer, and B. DuPont, A survey of multi-objective optimization methods and their applications for nuclear scientists and engineers, *Progress in Nuclear Energy*, *138*, 103,830, doi:https://doi.org/10.1016/j.pnucene.2021.103830, 2021.

[56] Strydom, G., Consistent multi-scale uncertainty quantification methodology for multi-physics modelling of prismatic HTGRs, Ph.D. thesis, North-West University (South Africa), 2020.

[57] Strydom, G., and H. D. Gougar, Preliminary reactor physics assessment of the HTR Module with 14% enriched UCO fuel, *Nuclear Engineering and Design*, *256*, 304–321, doi:https://doi.org/10.1016/j.nucengdes.2012.08.013, 2013.

[58] Sudret, B., Global sensitivity analysis using polynomial chaos expansions, *Reliability Engineering & System Safety*, *93*(7), 964–979, doi:10.1016/j.ress.2007.04.002, 2008.

[59] Tavron, B., and E. Shwageraus, Pebble bed reactor fuel cycle optimization using particle swarm algorithm, *Nuclear Engineering and Design*, *307*, 96–105, doi:10.1016/j.nucengdes.2016.06.033, 2016.

[60] Toptan, A., W. Jiang, J. D. Hales, B. W. Spencer, A. Casagranda, and S. R. Novascone, FEA-aided investigation of the effective thermal conductivity in a medium with embedded spheres, *Nuclear Engineering and Design*, *381*, 111,355, doi:https://doi.org/10.1016/j.nucengdes.2021.111355, 2021.

[61] Wang, Y., et al., Rattlesnake: A MOOSE-based multiphysics multischeme radiation transport application, *Nuclear Technology*, *207*(7), 1047–1072, doi:https://doi.org/10.1080/00295450.2020.1843348, 2021.

[62] Williamson, R. L., et al., Bison: A flexible code for advanced simulation of the performance of multiple nuclear fuel forms, *Nuclear Technology*, *207*(7), 954–980, doi:10.1080/00295450.2020.1836940, 2021.

[63] Windes, W., G. Strydom, R. Smith, and J. Kane, Role of nuclear grade graphite in controlling oxidation in modular HTGRs, *Tech. Rep. INL/EXT-14-31720*, Idaho National Laboratory, 2014.

[64] Zhang, J., F. Li, and Y. Sun, Physical analysis of the initial core and running-in phase for pebble-bed reactor HTR-PM, *Science and Technology of Nuclear Installations*, *2017*, 8918,424, doi:10.1155/2017/8918424, 2017.

[65] Zheng, Y., M. M. Stempniewicz, Z. Chen, and L. Shi, Study on the DLOFC and PLOFC accidents of the 200 MWe pebble-bed modular high temperature gas-cooled reactor with TINTE and SPECTRA codes, *Annals of Nuclear Energy*, *120*, 763–777, doi:https://doi.org/10.1016/j.anucene.2018.06.041, 2018.

[66] Zienkiewicz, O. C., and Z. Z. R. L. Taylor, *The Finite Element Method: Its Basis and Fundamentals*, Butterworth-Heinemann, 2005.

Appendix A
Artificial Neural Network Grid Study

This appendix shows the results of the grid-search for ANN configurations described in Section 4.3. The configurations considered are shown in Table 6, which included 85 trials for each QoI. Figure 16 shows the resulting performance of each configuration, ordered by the complexity of the model, and Table 10 lists the top ten models.

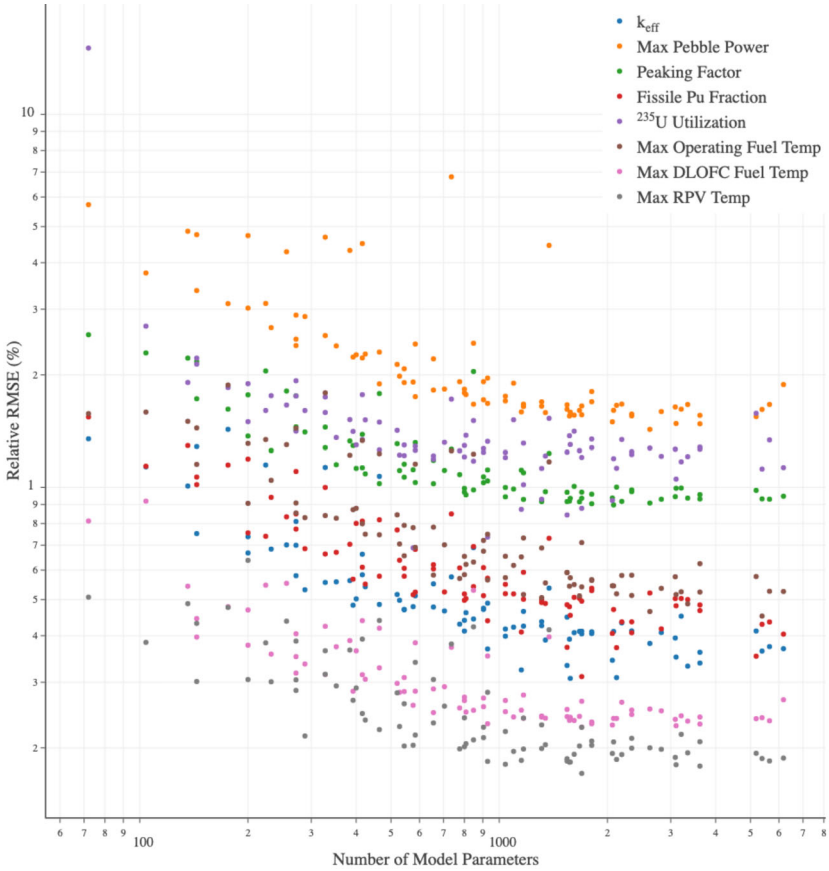

Figure 16: ANN performance versus model complexity.

Table 10: Top 10 ANN configurations for each QoI.

Rank	Network Configuration (Layer 1/Layer 2/Layer 3) and Relative RMSE															
	k_{eff}		Max Pebble Power		Peaking Factor		Fissile Pu Fraction		^{235}U Utilization		Max Operating Fuel Temp		Max DLOFC Fuel Temp		Max RPV Temp	
1	64/4/4	0.307%	32/32	1.428%	32/16/16	0.897%	64/4/16	0.311%	32	0.688%	64/32/4	0.452%	64/4/16	0.228%	64/4/16	0.171%
2	64/8/4	0.309%	64/16/16	1.478%	32/16/8	0.904%	64/32	0.352%	32/4/16	0.736%	32/32	0.464%	64/16	0.229%	64/16/16	0.179%
3	64	0.324%	64/16	1.481%	32/32	0.908%	64/8/4	0.372%	64/4	0.843%	64/16/8	0.487%	64/8	0.230%	64/16	0.180%
4	64/16/8	0.331%	64/8	1.496%	32/16/4	0.915%	64/4	0.372%	64	0.873%	32/8/16	0.506%	64/16/16	0.232%	32/8	0.181%
5	64/4	0.333%	64/32	1.546%	16/16/16	0.916%	64/32/16	0.404%	64/4/16	0.879%	32/16/4	0.512%	32/16/8	0.232%	32/16	0.183%
6	32/32/16	0.338%	32/16	1.551%	64/4	0.916%	64/8	0.406%	64/8	0.922%	64/8/16	0.513%	32/4/16	0.232%	64/4/4	0.183%
7	64/8	0.344%	64/8/16	1.553%	64/8/8	0.917%	64/8/16	0.407%	16/16/16	0.929%	64/16	0.516%	16/32/8	0.235%	32/4/16	0.184%
8	64/16	0.351%	32/32/16	1.558%	32/16	0.919%	64	0.409%	8/32/8	1.016%	32/8/8	0.516%	64/16/8	0.236%	64/4	0.184%
9	64/16/16	0.361%	64/4/8	1.559%	32/8/16	0.923%	32/32/4	0.417%	64/16	1.052%	64/8/4	0.518%	64/32/8	0.236%	64/32/8	0.185%
10	64/32/4	0.364%	64/4/16	1.563%	32/8/8	0.928%	64/32/4	0.429%	64/32/4	1.119%	32/32/16	0.524%	64/4/8	0.237%	64	0.185%