Analysis of the Molten Salt Fast Reactor using reduced-order models

Alsayyari, Fahad; Tiberga, Marco; Perkó, Zoltán; Kloosterman, Jan Leen; Lathouwers, Danny

Important note
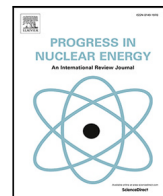To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Analysis of the Molten Salt Fast Reactor using reduced-order models

Fahad Alsayyari, Marco Tiberga, Zoltán Perkó *, Jan Leen Kloosterman, Danny Lathouwers

Delft University of Technology, Faculty of Applied Sciences, Department of Radiation Science and Technology, Mekelweg 15, Delft, 2629JB, The Netherlands

ARTICLE INFO

ABSTRACT

In this paper, we present a reduced-order modeling approach to study the Molten Salt Fast Reactor (MSFR). Our approach is nonintrusive and based on the proper orthogonal decomposition method. We include adaptivity in selecting the sampling points both in time and parameter space. Steady-state and transient analysis were both performed using the developed models. In the steady-state analysis, we capture the effect of 30 model parameters on the spatial distributions of fission power and temperature, and on the multiplication factor. The dimensionality of the fission power was reduced from the 104288 nominal dimensions in the physical space to 10 dimensions in the reduced space, whereas the temperature was reduced from 220972 dimensions to 3. The reduced model was then used for uncertainty and sensitivity study of the maximum temperature in the reactor and the multiplication factor. In the transient analysis, the reduced model captured the effect of perturbations in the flow rate of salt in the intermediate circuit on the fission power density and temperature. The reduced models were successfully tested on a set of points that were not part of the snapshots used during the construction stage.

1. Introduction

Molten salt reactors have gained interest due to their potential safety, reliability, and sustainability (Generation IV International Forum, 2002). Different designs of this concept have been proposed in the literature Dolan (2017). In this work, we consider the Molten Salt Fast Reactor (MSFR) (Allibert et al., 2016). A key design feature of this unmoderated reactor is the use of a liquid salt fuel, which also plays the role of the coolant. This design introduces a unique modeling challenge because of the tightly coupled neutronics and thermal-hydraulics phenomena (e.g., transport of delayed neutron precursors, distributed thermal energy deposition directly in the coolant, a strong negative temperature feedback coefficient). To address these challenges, high-fidelity coupled models are used to provide an insight into the behavior of the reactor (e.g., Aufiero et al., 2014, Fiorina et al., 2014, Laureau et al., 2017, Cervi et al., 2019, Tiberga et al., 2020b). For safety assessment applications, an accurate and explicit quantification of the propagation of uncertainties through these complex models is required (International Atomic Energy Agency, 2019). Quantifying uncertainties and analyzing sensitivities in reactor physics can be accomplished using adjoint methods (Gilli et al., 2013; Perkó et al., 2013). However, adjoint-based methods require the availability of an adjoint solver, which might not always be feasible for coupled problems. Another approach is using forward-based methods, which requires repeated evaluation of the high-fidelity model for different parameter configurations (Cacuci, 2003). For example, automatic differentiation tools can be used to sample the forward model and reconstruct a discrete version of the adjoint problem (Marta et al., 2007). However, for large-scale, coupled, high-fidelity models with multiple input parameters such approaches are computationally demanding. For such applications, *reduced-order modeling* (ROM) techniques can be used to simplify the high-fidelity model and produce an efficient, cheap, and accurate model of the system.

Amongst the classical ROM methods found in literature, proper orthogonal decomposition (POD) is the method most suited for nonlinear systems (Antoulas et al., 2001; Schilders et al., 2008; Benner et al., 2015). In the POD approach, a reduced basis space for the system is built using snapshots of the high-fidelity model. In nuclear reactor application, POD has been applied to solve criticality eigenvalue problems (Buchan et al., 2013; Senecal and Ji, 2019; German and Ragusa, 2019; Prince and Ragusa, 2020), for fuel pin reactor core calculations (Cherezov et al., 2018), in fuel burnup calculations (Castagna et al., 2020), in thermal hydraulics modeling (Vergari et al., 2020), in stability analysis (Prill and Class, 2014; Manthey et al., 2019), in spent fuel pool modeling (Escanciano and Class, 2019), and to model the lead cooled fast reactor (Sartori et al., 2016). In all of these applications, the original set of model equations were projected onto the constructed reduced basis. Projection-based approaches are intrusive because of the need to access the operator of the original high-fidelity model in

---

* Corresponding author.
E-mail address: Z.Perko@tudelft.nl (Z. Perkó).

order to employ the projection. In cases where the high-fidelity model is a legacy coupled solver or closed-source, intrusive approaches are not applicable. Moreover, the stability of the high-fidelity model is not preserved in the reduced model with the use of projection-based POD methods (Amsallem and Farhat, 2012).

To overcome these issues, nonintrusive approaches have been proposed in the literature, which are techniques that can be wrapped around the high-fidelity solver and avoid the need to access the original set of equations. Nonintrusive spectral projection methods such as polynomial chaos expansion (PCE) is one approach for uncertainty quantification applications (Perkó et al., 2014a,b), which has been applied to a molten salt fueled system (Santanoceto et al., 2021b) and to perform a preliminary uncertainty and sensitivity analysis of the MSFR steady-state (Santanoceto et al., 2021a). However, such methods do not reduce the dimensionality of the high-fidelity model. Dynamic Mode Decomposition (DMD) is a nonintrusive model order reduction method that has been employed to model the MSFR (Ronco et al., 2020). DMD constructs a linear operator from a sequence of snapshots of the system. Then, a reduced model preserving the dynamics of the original system is formulated, which can be used for transient analysis and control applications (Tu et al., 2014). Propagating uncertainties in the initial conditions has also been studied using DMD (Takeishi et al., 2017; Abdo et al., 2019). However, for problems with parametrized models where uncertainty and sensitivity analysis for multiple model input parameters are to be studied, applying DMD is cumbersome.

An alternative ROM approach is through a nonintrusive POD implementation, where a surrogate model is trained to compute the POD coefficients instead of projecting the model equation onto the reduced basis. In nuclear reactor applications, a Range Finding Algorithm (RFA) has been used in Huang et al. (2017) to build the reduced-basis subspace(referred to as active subspace) combined with a simple polynomial surrogate for the POD coefficients. Hybrid approaches aiming to simultaneously reduce the dimensionality of both the state and the input variables using RFA have also been explored (Abdel-Khalik et al., 2013; Bang et al., 2015). We have proposed a nonintrusive adaptive POD algorithm in Alsayyari et al. (2019). A key difference in our developed algorithm is the adaptive sampling of the high-fidelity model as opposed to random or a priori uniform sampling employed by the RFA. Our algorithm is based on adaptive sparse grids technique, which is suited for problems with higher dimensional input spaces. An uncertainty and sensitivity application of the algorithm on a simplified two-dimensional molten salt fueled system was presented in Alsayyari et al. (2020), where we have also addressed handling systems with multiple-outputs. We have presented an extension of the algorithm in Alsayyari et al. (2021) to deal with time-dependent parametrized problems. In this work, we investigate the MSFR reactor using the developed algorithm, which demonstrates the capability of the algorithm on a large-scale, three-dimensional MSFR model. We consider two applications. The first is a steady-state reduced model for uncertainty and sensitivity analysis of 30 model parameters. The second is a transient reduced model, which can be used for transient analysis and control applications.

The remainder of the paper is organized as follows: Section 2 presents an introduction to the problem formulation and the POD method. A summary of the adaptive sampling algorithm is presented in Section 3. Then, a description of the MSFR model is given in Section 4. The results for the steady-state analysis along with the uncertainty and sensitivity study are given in Section 5. We present the results of the transient study in Section 6. Finally, conclusions are discussed in Section 7.

2. Proper orthogonal decomposition

Consider a general high-fidelity model that is described by a function $f : \mathbb{R}^d \rightarrow \mathbb{R}^n$, which maps input parameter $\boldsymbol{p} \in \mathbb{R}^d$ to a state (field) vector $\boldsymbol{\phi} \in \mathbb{R}^n$, where d is the size of the input space (representing, for example, material properties, and boundary conditions), and n is

the size of the state space (representing, for example, space, angle, and energy). We consider f as a black-box model, which can either describe single physics or coupled multi-physics code, that is,

$$\boldsymbol{\phi} = f(\boldsymbol{p}). \tag{1}$$

The high-fidelity model f can be evaluated at any given $\boldsymbol{p}_q$ and produce an output $\boldsymbol{\phi}_q$ (i.e., $\boldsymbol{\phi}_q = f(\boldsymbol{p}_q)$). We are interested in finding an approximation for $\boldsymbol{\phi}_q$ without the use of the computationally intensive f. For this task, we build a reduced model that produces an approximation $\tilde{\boldsymbol{\phi}}_q$, which takes the form of

$$\tilde{\boldsymbol{\phi}}_q = \sum_{g=1}^{r} c_g(\boldsymbol{p}_q)\boldsymbol{u}_g, \tag{2}$$

where $\boldsymbol{u}_g$ is the gth basis vector in the reduced basis space, $c_g(\boldsymbol{p})$ is the corresponding coefficient, which is a function of the parameter $\boldsymbol{p}$, and r is the size of the reduced basis space. The POD approach constructs a reduced basis space from snapshots of the state vectors $\boldsymbol{\phi}$ such that the error is minimized in a least square sense, that is, it solves the minimization problem

$$\min_{\boldsymbol{u}_g} \left\| \boldsymbol{\phi} - \sum_{g=1}^{r} c_g(\boldsymbol{p})\boldsymbol{u}_g \right\|_{\ell^2}. \tag{3}$$

The basis vectors are defined to be orthonormal (i.e., $\langle \boldsymbol{u}_g, \boldsymbol{u}_h \rangle = \delta_{gh}$, where $\langle .,. \rangle$ indicates the scalar product). The reduced basis space can be found using the method of snapshots. Let M be a matrix collecting snapshots of the solution $\boldsymbol{\phi}$ at some selected values of $\boldsymbol{p}$, that is,

$$\mathbf{M} = [\boldsymbol{\phi}_1, \boldsymbol{\phi}_2, \ldots, \boldsymbol{\phi}_s] \quad \in \mathbb{R}^{n \times s}, \tag{4}$$

where $\boldsymbol{\phi}_q$ is the solution of the high-fidelity model (f) at parameter value $\boldsymbol{p}_q$, and s is the number of selected snapshots. Note that for transient problems where f computes a solution $\boldsymbol{\phi}$ that is time-dependent, we consider the time to be a parameter, rather than an independent variable — that is, $\boldsymbol{p}$ includes time as one of the parameters. Such a formulation allows for time adaptive sampling as described in Alsayyari et al. (2021).

The basis vectors are then found to be the left singular vectors of the singular value decomposition (SVD) applied to the matrix M, that is, let M be decomposed using SVD as $M = USV^\top$, then $\boldsymbol{u}_g \ \forall \ 1 \leq g \leq r$ are the first r column vectors of U. The size of the reduced space is determined by r, which can be chosen to truncate U such that the sum of the squared singular values (σ) corresponding to the neglected singular vectors is below a predefined threshold γ_{tr}.

$$\frac{\sum_{g=r+1}^{n} \sigma_g^2}{\sum_{g=1}^{n} \sigma_g^2} < \gamma_{\mathrm{tr}}. \tag{5}$$

Once the basis vectors are known, we can use the orthogonality of the space to compute the coefficients at the sampled point ($\boldsymbol{p}_q$) as

$$c_g(\boldsymbol{p}_q) = \langle \boldsymbol{u}_g, \boldsymbol{\phi}_q \rangle. \tag{6}$$

3. Adaptive sampling

Using the values of the coefficient $c_g(\boldsymbol{p})$ at the sampled points, one can train a surrogate model to compute the solution of the system at any point. Different surrogate models can be employed for such a task. To deal with higher dimensional input spaces, we choose the Smolyak sparse grid interpolation technique. The hierarchical structure of the interpolant allows for the desired adaptive strategy, which can further reduce the burden of dealing with high dimensional spaces. This section is devoted to summarize the developed adaptive algorithm, which is detailed in Alsayyari et al. (2019).

Our implementation of the sparse grid technique uses localized adaptive sampling. Without loss of generality, the d-dimensional space of the input parameter $\boldsymbol{p}$ is mapped to a unitary hypercube $[0, 1]^d$. The sparse grid technique first generates unidimensional nodes along each
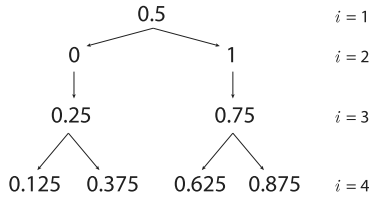
Fig. 1. Illustration of the first 4 levels of the tree structure, where 0.5 is the root of the tree and nodes are added at half the distances between the previous nodes. Each node has 2 children except the nodes at level 2 where each has one child only.

dimension. Then, points in the parameter space are formed by a specific combination of the generated unidimensional nodes. Different choices for the unidimensional node generation rule are possible. However, we choose the equidistant rule to increase the separation between points. The unidimensional nodes are arranged in a tree structure as shown in Fig. 1. Each level is assigned an index i and contains several nodes. Nodes are added at each level at half the distances between the nodes from the previous levels. Each node is connected to two children at the next level ($i+1$) and one father from the previous level ($i-1$). The root node is considered a father for itself. There is, however, an exception at level $i = 2$, where each node has one child because these nodes mark the boundaries of the unit hypercube.

Each node in the tree can be uniquely identified with the level index i and an index j. For each node from the tree p_j^i along dimension l, we can build a basis function $a_{p_j^i}^i(p)$ as

$$a_{p^1}^1 = 1 \quad \text{if } i = 1,$$

$$a_{p_j^i}^i(p) = \begin{cases} 1 - (m^i - 1) \cdot |p - p_j^i|, & \text{if } |p - p_j^i| < \dfrac{1}{m^i - 1}, \\ 0, & \text{otherwise}, \end{cases} \quad (7)$$

where the dependence on the dimension l is dropped for notational convenience. The level parameter m^i is defined as follows:

$$m^i = \begin{cases} 1 & \text{if } i = 1, \\ 2^{i-1} + 1 & \text{if } i > 1. \end{cases} \quad (8)$$

A point p in parameter space is formed by combining nodes from all dimensions (i.e., $p = (p_{j_1}^{i_1}, \dots p_{j_d}^{i_d})$). Extending the tree structure to points in parameter space allows us to define *forward* points. The first forward points for a point p along the first dimension are $(b(p_{j_1}^{i_1}), \dots p_{j_d}^{i_d})$, where $b(p)$ is a function that returns the children of the node p. In general, the forward points along the lth dimension are $(p_{j_1}^{i_1}, \dots, b(p_{j_l}^{i_l}), \dots p_{j_d}^{i_d})$. Because each node has at most two children, each point has at most $2d$ forward points. We can also define *backward* points in the same manner by applying a function $b^*(p)$, which returns the father of the node. Therefore, each point has at most d backward points. Recursively generating the forward points creates the set of *descendant* points. On the other hand, the *ancestor* points for a point p are formed by recursively generating backward points until the root point $(0.5, \dots, 0.5)$ is reached.

The adaptive algorithm is iterative. In every iteration, a subset of points is selected and marked *important* from a set of trial points. Let the iteration number be k and a set that collects the selected important points at each iteration k be $\mathcal{Z}^k$. The union of the important points sets ($\mathcal{Z}^k$) from all iterations up to k forms the set of selected grid points $\mathcal{X}^k$, that is,

$$\mathcal{X}^k = \bigcup_{h=0}^{k} \mathcal{Z}^h. \quad (9)$$

A point p_q is marked as important if it satisfies one of two conditions. The first is to have an approximation error e_q above a defined threshold γ_{int} with all ancestors of that points included in $\mathcal{X}^k$. The

second is to have a descendant point with an error above γ_{int}. Points that do not meet any of the conditions are added to a set of inactive points. The error is computed in the ℓ^2 norm as

$$e_q^k = \frac{\left\| \phi_q - \tilde{\phi}_q^k \right\|_{\ell^2}}{\left\| \phi_q \right\|_{\ell^2} + \epsilon}, \quad (10)$$

where ϕ_q is the solution returned by the high-fidelity code (f) at point p_q, $\tilde{\phi}_q^k$ is the approximation produced by the reduced model at the point p_q and iteration k, and ϵ is an offset introduced for cases when the norm of the solution is near-zero.

At iteration k, a surrogate model for the POD coefficient $c_g(p)$ is built as

$$A_{k,d}(c)(p) = A_{k-1,d}(c)(p) + \Delta A_{k,d}(c)(p), \quad (11)$$

where $A_{k,d}(c)(p)$ is the Smolyak interpolation operator applied to $c(p)$ that depends on the iteration k and dimension of the input space d. For the initialization ($k = 0$), we enforce $A_{0,d}(c)(p)$ to be zero. The term $\Delta A_{k,d}(c)(p)$ is defined as

$$\Delta A_{k,d}(c)(p) = \sum_{n=1}^{m_k^\Delta} w_n^k \Theta_n(p), \quad (12)$$

where m_k^Δ is the cardinality of the set $\mathcal{Z}^k$. The function $\Theta_n(p)$ is the d-variate basis function for the point $p_n \in \mathcal{Z}^k$,

$$\Theta_n(p) = \prod_{l=1}^{d} a_{p_{n,l}^{i_l}}^{i_l}(p_l), \quad (13)$$

The surplus w_n^k is defined as the difference between the interpolated value and the true value of the coefficient at p_n,

$$w_n^k = c(p_n) - A_{k-1,d}(c)(p_n). \quad (14)$$

At each iteration, a trial set is generated from the forward points of the important set. The model is tested on each point within the trial set. To provide an estimate for the global error in the reduced model at iteration k (ϵ^k), an appropriate norm can be chosen over the error defined in Eq. (10). In this work, we have chosen the maximum norm such as the reduced model error is defined by

$$\epsilon^k = \max \ e_q^k \quad \forall 1 < q < \tau^k, \quad (15)$$

where τ^k is the cardinality of the trial set at iteration k. However, ϵ^k only provides an estimate of the error in the reduced model and is not rigorous.

Algorithm

In the initialization step ($k = 0$), we select the root point at the center of the hypercube $(0.5, 0.5, \dots, 0.5)$ and evaluate the high-fidelity model f. That point is added to the important set $\mathcal{Z}^0$, and the snapshot is added to the matrix M. Then, set $k = 1$ and do

1. Generate the forward points of the points in $\mathcal{Z}^{k-1}$;
2. Add points that have all their backward points in $\mathcal{X}^{k-1}$ (defined as in Eq. (9)) to the trial set;
3. Evaluate the high-fidelity model f at the points in the trial set;
4. Compute the coefficients $c_g(p)$ at the points in the trial set using $A_{k-1,d}(c)(p)$;
5. Compute the error at every point in the trial set using Eq. (10);
6. If the maximum found error (ϵ^k) was below a prescribed global tolerance ζ, terminate; Otherwise, continue;
7. Find the important points $\mathcal{Z}^k$;
8. Add the snapshots corresponding to the important points in $\mathcal{Z}^k$ to the matrix M;
9. Perform SVD on the matrix M to extract the reduced basis;
10. Use Eqs. (11)–(14) to construct $A_{k,d}(c)(p)$;
11. Set $k \leftarrow k + 1$ and go to Step 1.

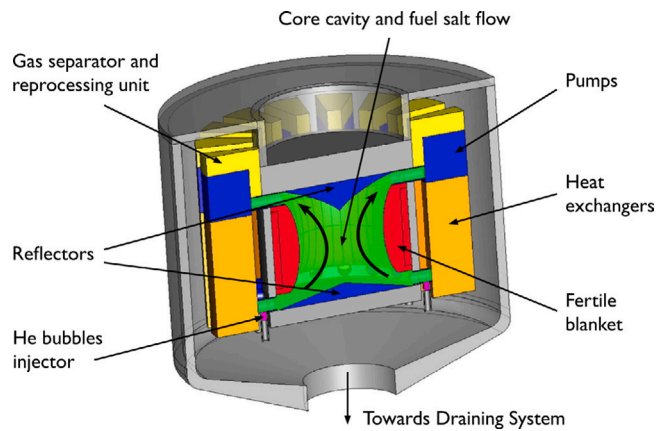For a detailed description of the algorithm, see Alsayyari et al. (2019).

Fig. 2. A schematic illustration of the main design features of the MSFR fuel circuit.

Fig. 3. Computational scheme of the multi-physics tool constituting the high-fidelity MSFR model. DGFlows is the CFD code, while PHANTOM-S$_N$ is the neutronics code. Data are exchanged at each iteration between the two solvers to model the coupled physics phenomena characterizing the MSFR (Tiberga et al., 2020b).

Fig. 4. Geometry of the MSFR recirculation loop used for simulations, showing the main regions considered in the model (Tiberga et al., 2020b).

4. MSFR model

This section presents the MSFR main design features and modeling approach. A more detailed description of the model can be found in Tiberga et al. (2020b) and the references within.

A schematic illustration of the fuel circuit is shown in Fig. 2. The reactor core is a toroidal cavity where the liquid fuel salt (a mixture of lithium, thorium, and fissile nuclides fluorides) can flow freely without any moderator or control rod. Sixteen identical sectors branch out from the central cavity. Each sector contains a pump, a heat exchanger, and a unit for the separation and treatment of the helium gas dispersed in the fuel salt to control reactivity and remove metallic fission products (Delpech et al., 2009). A blanket with fertile salt surrounds the cavity while reflectors are placed at the top and bottom of the core. The heat exchangers transfer thermal energy to the intermediate circuit filled with inert molten salt, which in turn, delivers the heat to the energy conversion system consisting of a conventional Joule-Brayton cycle. Table 1 summarizes the main MSFR design parameters.

The reactor was modeled using an in-house multi-physics tool, which couples an incompressible Reynolds-averaged Navier–Stokes solver (DGFlows) (Tiberga et al., 2020a; Hennink et al., 2021) with an S_N multi-group neutronics code (PHANTOM-S$_N$) (Kóphazi and Lathouwers, 2012). The latter is equipped with transport equations for the delayed neutron precursors to model their movement. Both codes employ the discontinuous Galerkin finite element method for space discretization and implicit backward differentiation formulae (BDF) time schemes. Fig. 3 displays the structure of the multi-physics tool and the data exchanged between the codes. PHANTOM-S$_N$ receives the average temperature on each element (T_{avg}) and corrects the cross sections accordingly with respect to a library at the reference temperature T_0. The corrections take into account the effects of density feedback (which has a linear dependence on temperature) and Doppler feedback (which has a logarithmic dependence). The velocity and turbulent viscosity fields (**u** and v_t) are also taken from DGFlows as another input to solve the delayed neutron precursors equation. Then, the fission power density (P_{fiss}) is transferred to DGFlows as it constitutes the right-hand side of the energy equation. The steady-state solution is sought by iterating the solvers until convergence (four iterations are typically sufficient). On the other hand, transient simulations are performed with a loose-coupling strategy, first computing a time-step with DGFlows and then calling PHANTOM-S$_N$ to solve the neutronics problem. We refer the reader to Tiberga et al. (2019, 2020b) for a more comprehensive description of the coupled code.

Because of the symmetry in the reactor design, we modeled only one recirculation loop. The geometry considered is reported in Fig. 4. While the neutron flux is calculated in the reflectors and the blanket, the CFD code (DGFlows) neglects heat transfer in these regions. Fig. 5 shows
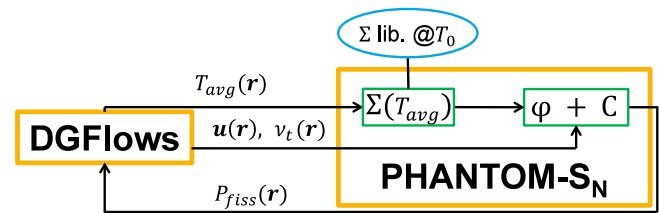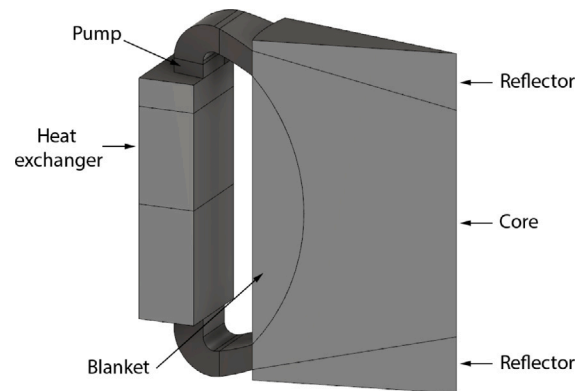
the meshes used in this study. Neutronics calculations were performed on an unstructured mesh consisting of 26072 tetrahedra (of which 21489 are in the fuel salt domain). This mesh is finer in the core region, where neutron importance is high, and coarser in the external sector. This master mesh was then refined once uniformly in the outer-core region to obtain the CFD mesh, which consists of 55243 elements. A second-order polynomial was used for the velocity discretization, while a first order polynomial was used for all other quantities.

The fuel salt composition is reported in Table 2, along with some physical properties that were fixed and selected for the uncertainty analysis study of Section 5. The neutron energy range was condensed into six groups, with boundaries shown in Table 3. Delayed neutron precursors were grouped into eight families. All neutronics data were evaluated at temperature $T_0 = 900 \times 10^0$ K with Serpent (Leppänen et al., 2015) using the JEFF-3.1.1 data library (Santamarina et al., 2009).

For the CFD calculations, symmetry boundary conditions were imposed at the wedge sides, and standard wall-functions with adiabatic conditions were assumed at all walls. For the neutronics calculations, reflective boundary conditions were assumed at the sides of the wedge, while vacuum conditions were imposed elsewhere. For the transport of the precursors within the fuel salt, homogeneous Neumann and no-inflow conditions were imposed at all walls.

Lacking detailed design specifications for the primary heat exchanger, salt cooling was modeled via a volumetric heat sink term equal to $h_{int,0} \left(T_{int,0} - T \right)$, where $h_{int,0}$ and $T_{int,0}$ are the nominal volumetric heat transfer coefficient and average temperature of the intermediate salt whose values are reported in Table 1. The fuel pump was modeled with a momentum source term, and buoyancy was taken into account through the Boussinesq approximation (the reference density and thermal expansion coefficient are reported in Table 2). We considered the flow field to be fixed at the nominal state because natural convection has a negligible contribution to the total nominal flow rate, and the pump specifications were fixed throughout the analyses in

Table 1
Main design parameters for MSFR model in the nominal case (Allibert et al., 2016; Gérardin et al., 2017).

Parameter	Value
Total thermal power [MW]	3000×10^0
Total fuel salt volume [m^3]	18×10^0
Fuel salt circulation time [s]	4×10^0
Average fuel salt temperature [K]	973.15×10^0
Minimum fuel salt temperature [K]	923.15×10^0
Average intermediate salt temperature [K]	908.15×10^0
Pressure drop across heat exchanger [bar]	4×10^0
Volumetric heat transfer coefficient with intermediate salt [MW/m^{3}K]	19.95×10^0

(a) Neutronics mesh   (b) CFD mesh
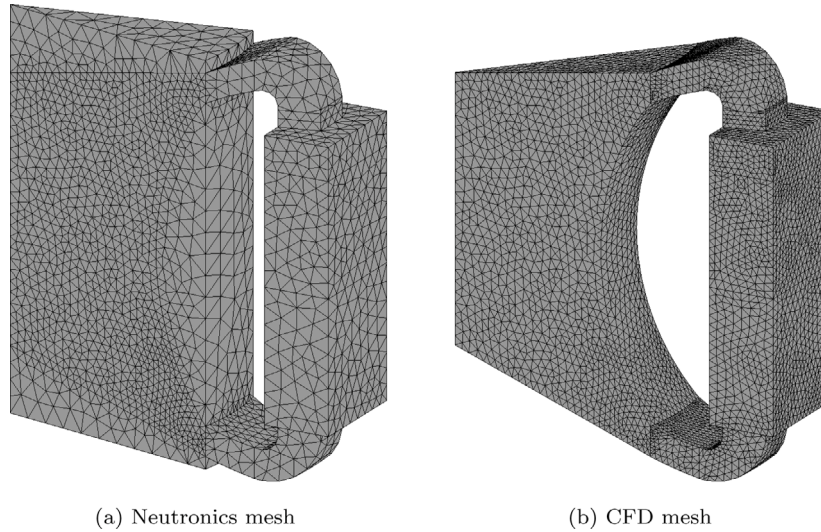
Fig. 5. Mesh adopted for the MSFR model. The neutronics mesh (left), corresponding to the master-mesh, consists of 26072 tetrahedra (21489 in the fuel salt domain), while the CFD mesh (right) has 55243 elements. The latter is derived by refining the former once uniformly in the outer-core region.

Table 2
Properties of the fuel salt mixture (Beneš et al., 2013).

Property	Value or equation of state
Composition [%mol]	LiF(77.5×10^0)-ThF$_4$(6.6×10^0)-$^{\text{enr}}$UF$_4$(12.3×10^0)-(Pu-MA)F$_3$(3.6×10^0)
Density [kg/m^3]	4306.7×10^0
Dynamic viscosity [Pa s]	$6.187 \times 10^{-4} \exp\left(772.2/\left(T(\text{K}) - 765.2\right)\right)$
Thermal expansion coefficient [K^{-1}]	1.9119×10^{-4}
Melting point [K]	854.15×10^0

this work. Moreover, since we are interested in reactor steady-state and operational transients, decay heat was not taken into account. For the transient calculations described in Section 6, the time-dependent equations were discretized with the second-order BDF scheme with a fixed time-step size of 0.1×10^0 s.

5. Steady-state analysis

In this section, we study the effect of selected parameters of interest on three outputs: fission power density (P_{fiss}), temperature distribution (T), and the multiplication factor, k_{eff}. The fission power density P_{fiss} is computed on the neutronics mesh with 104288 unknowns (nominal dimensionality) while the temperature T is computed on the CFD mesh with 220972 unknowns. This analysis is performed on the steady-state model of the MSFR with a fixed reactor power of 3000 MW. However, since we are modeling 1/16th of the reactor, the observed nominal power is 187.5 MW. The selected parameters to be investigated were specific heat capacity of the fuel salt (c_p, influencing the temperature distribution), heat conductivity of the fuel salt (κ), fission cross sections for the six energy groups ($\Sigma_{f,g}$), capture cross sections for the six energy groups ($\Sigma_{c,g}$), delayed neutron fractions for eight families of precursors (β_i), and corresponding decay constants (λ_i). The nominal values of the

selected parameters are reported in Table 4. These nominal values are theoretical as the actual salt properties and the associated uncertainties are still under investigation. However, for the propose of this work, we assume an uncertainty following a Gaussian distribution for all parameters with a mean (μ) equal to the nominal value and a standard deviation (σ) taken as 5% of the mean.

5.1. Construction of the reduced model

Our implementation of the adaptive approach was developed for bounded input domains. For this reason, we consider the Gaussian distribution to be truncated. Truncating the Gaussian distribution, in this case, is a valid approximation because the parameter uncertainties we are considering are epistemic. Therefore, we are not altering a random process but rather limiting the scope of analysis to the region with the highest probability of having the true value. Moreover, truncating the distribution prevents unphysical values of the parameter, such as negative β. The truncation was selected to be at 3σ, retaining 99.7% of the probability range. This implies that the range of variation for all parameters is set to be $\pm 15\%$ of the nominal value. We first constructed the reduced model to be uniformly accurate within the defined range, then use the reduced model for the uncertainty and sensitivity analysis

Table 3

Upper energy bound for each group and average flux of each group for the nominal case.

Energy group	1	2	3	4	5	6
Upper energy bound [keV]	20000	2231	497.9	2.479	5.531	0.7485
Average group flux [cm^{-2}s^{-1}] $\times 10^{14}$	0.51	2.05	5.14	3.51	2.20	0.46

Table 4

Nominal values of the selected parameters.

Parameter	Nominal value	Parameter	Nominal value
c_p [J/kgK]	1.59×10^3	β_2	7.14×10^{-4}
κ [W/mK]	1.7×10^0	β_3	3.59×10^{-4}
$\Sigma_{f,1}$ [cm^{-1}]	4.45×10^{-3}	β_4	7.94×10^{-4}
$\Sigma_{f,2}$ [cm^{-1}]	2.52×10^{-3}	β_5	1.47×10^{-3}
$\Sigma_{f,3}$ [cm^{-1}]	1.80×10^{-3}	β_6	5.14×10^{-4}
$\Sigma_{f,4}$ [cm^{-1}]	2.62×10^{-3}	β_7	4.65×10^{-4}
$\Sigma_{f,5}$ [cm^{-1}]	5.20×10^{-3}	β_8	1.51×10^{-4}
$\Sigma_{f,6}$ [cm^{-1}]	1.39×10^{-3}	λ_1 [s^{-1}]	1.25×10^{-2}
$\Sigma_{c,1}$ [cm^{-1}]	1.99×10^{-3}	λ_2 [s^{-1}]	2.83×10^{-2}
$\Sigma_{c,2}$ [cm^{-1}]	7.41×10^{-4}	λ_3 [s^{-1}]	4.25×10^{-2}
$\Sigma_{c,3}$ [cm^{-1}]	2.15×10^{-3}	λ_4 [s^{-1}]	1.33×10^{-1}
$\Sigma_{c,4}$ [cm^{-1}]	5.10×10^{-3}	λ_5 [s^{-1}]	2.92×10^{-1}
$\Sigma_{c,5}$ [cm^{-1}]	1.02×10^{-2}	λ_6 [s^{-1}]	6.66×10^{-1}
$\Sigma_{c,6}$ [cm^{-1}]	2.48×10^{-2}	λ_7 [s^{-1}]	1.63×10^0
β_1	1.23×10^{-4}	λ_8 [s^{-1}]	3.55×10^0

Table 5

Results for the steady-state reduced models for each output showing the maximum relative ℓ^2 error after a test on 1000 independent points, the number of POD modes after truncation, the number of evaluations to construct each model, and the final number of points in $\mathcal{X}^k$.

	$\boldsymbol{P}_{\text{fiss}}$	$\boldsymbol{T}$	k_{eff}
Maximum relative ℓ^2 error	0.24%	0.02%	37 pcm
Required tolerance ζ	1%	0.1%	50 pcm
Number of POD modes	10	3	–
Total number of evaluations	61	63	1639
Number of selected points in $\mathcal{X}^k$	15	5	227

employing the corresponding probability distribution of each parameter. A separate reduced model was built for each output. The global relative tolerance ζ was set to be 1% for $\boldsymbol{P}_{\text{fiss}}$, 0.1% for $\boldsymbol{T}$, and 50 pcm for k_{eff}. The interpolation threshold γ_{int} was 1×10^{-3} $\boldsymbol{P}_{\text{fiss}}$, 1×10^{-2} for $\boldsymbol{T}$, and 5×10^{-5} for k_{eff}. The POD truncation threshold γ_{tr} was 10^{-12} for all outputs.

After construction, each reduced model is tested on 1000 independent points that were not part of the snapshots generated during the constructions. Latin Hypercube Sampling (LHS) was used to draw the random testing points from the input space. Table 5 summarizes the test results for each model. It can be seen that all models resulted in a maximum relative error that was below the set tolerance ζ. These results certify that the reduced models are an accurate representation of the high-fidelity MSFR model with in the desired tolerances. Fig. 6 compares the reduced model for $\boldsymbol{P}_{\text{fiss}}$ with the reference model at the point with maximum error. The difference is seen to be maximum in the central region of the reactor core, where the flux is maximum. Fig. 7 shows the comparison at the maximum error for $\boldsymbol{T}$. The maximum difference for this case is observed at the bottom of the core, where the relative error locally is around 0.3%. A single high-fidelity computation requires about 4 CPU-hours (performed on a Linux cluster) whereas the reduced model produces the results in less than a second.

Table 5 reports also the number of POD modes after truncation, representing the number of dimensions in the reduced space. The dimensionality of $\boldsymbol{P}_{\text{fiss}}$ was reduced from the 104288 nominal dimensions in the physical space to 10 dimensions in the reduced space whereas $\boldsymbol{T}$ has 3 dimensions in the reduced space compared to 220972 nominal dimensions in the high-fidelity model. The low number of POD modes for $\boldsymbol{T}$ reflects the fact that the temperature profile in this reactor is fairly constant with respect to the perturbations considered. This is due to the

fact that the reactor power is fixed, and the perturbations are applied homogeneously. Note that k_{eff} is a single-valued response, and reducing the dimensionality is not applicable for this model.

The total number of evaluations requested by the algorithm during the construction stage is also reported in Table 5. This number corresponds to the total number of executing the high-fidelity model computations. The models for $\boldsymbol{P}_{\text{fiss}}$ and $\boldsymbol{T}$ needed 61, and 63 evaluations respectively, while the k_{eff} model needed 1639 evaluations. The increased number of evaluations reflects the strong nonlinearity of k_{eff} with respect to the input parameters compared to $\boldsymbol{P}_{\text{fiss}}$ and $\boldsymbol{T}$. This finding is supported by the number of unique nodes per dimension for each output, which is reported in Table 6. This number is the result of projecting the final grid points in $\mathcal{X}^k$ onto each dimension. It is a measure of the nonlinearity in the output with respect to the dimension as captured by the reduced model. A value of 1 indicates that the reduced model considered that dimension to have a negligible effect on the output. A value of 3 means that the output is linear or piecewise linear with respect to that input parameter because the reduced model used 3 nodes (the root node and the first two children) to construct a linear or a piecewise linear interpolant and no further refinement was required along that dimension. A value of 5 or more indicates that the interpolant along that dimension constructed a nonlinear interpolant between the nodes. In general, the nonlinearity in the interpolant is proportional to this number. Thus, Table 6 also explains the seemingly high difference between the number of evaluations required for the k_{eff} compared to $\boldsymbol{P}_{\text{fiss}}$ and $\boldsymbol{T}$, simply caused by the combination of the nonlinear dependence of k_{eff} on many of the input parameters simultaneously and the curse of dimensionality, leading to many more potential combinations of parameters the algorithm needs to check.

Table 6 shows that the fission power density was found to be linear with respect to the specific heat capacity c_p, which is explained by the temperature feedback effect. The fission cross sections of groups 2 to 6 also have a linear effect on the fission power density, which is expected since the fission power is proportional to $\Sigma_{f,g}$. The fission cross section of the first group, however, was observed to have a negligible effect within the tolerances on $\boldsymbol{P}_{\text{fiss}}$. This can be explained by the lower magnitude of the first-group flux compared to the rest of the groups, as shown in the average flux value per group in Table 3. This has an effect of a lower weight in the calculation of the fission power density. The capture cross section of the most thermal group is the only capture cross section that has a linear effect on $\boldsymbol{P}_{\text{fiss}}$. This is due to its larger nominal value resulting in a larger range of variations compared to the other groups. The rest of the parameters had a negligible effect on $\boldsymbol{P}_{\text{fiss}}$ and were considered as constants. The temperature is shown to be nonlinear with respect to c_p having 5 unique nodes and unaffected by the rest of the parameters within the defined tolerances. One can deduce that the 5 important points selected by the algorithm to be in $\mathcal{X}^k$ and reported in Table 5 correspond to the 5 points obtained by varying c_p only. The multiplication factor is nonlinear with respect to the cross section of groups 3 to 6 both fission and capture while having a piecewise linear interpolant with respect to the two most energetic groups. The specific heat capacity is also seen to have a linear effect on k_{eff} through the temperature feedback. The delayed neutron fractions of families 2, 4, and 5 were the only families with a significant effect on k_{eff}. This, however, is due to the larger nominal value of these parameters, which results in a larger range of variations compared to the rest of the families. The decay constant is shown to be taken as a constant for all families with 1 node each, which indicates that within the set tolerance of 50 pcm and a range of variation of $\pm 15\%$, λ_i had no effect on k_{eff}.

5.2. Propagating uncertainties

We used the reduced models to propagate uncertainties in the parameters to the responses of interest. We consider two model responses; The first is the maximum temperature of the system, which has a value in the nominal state of 1084.8×10^0 K, and the second is the
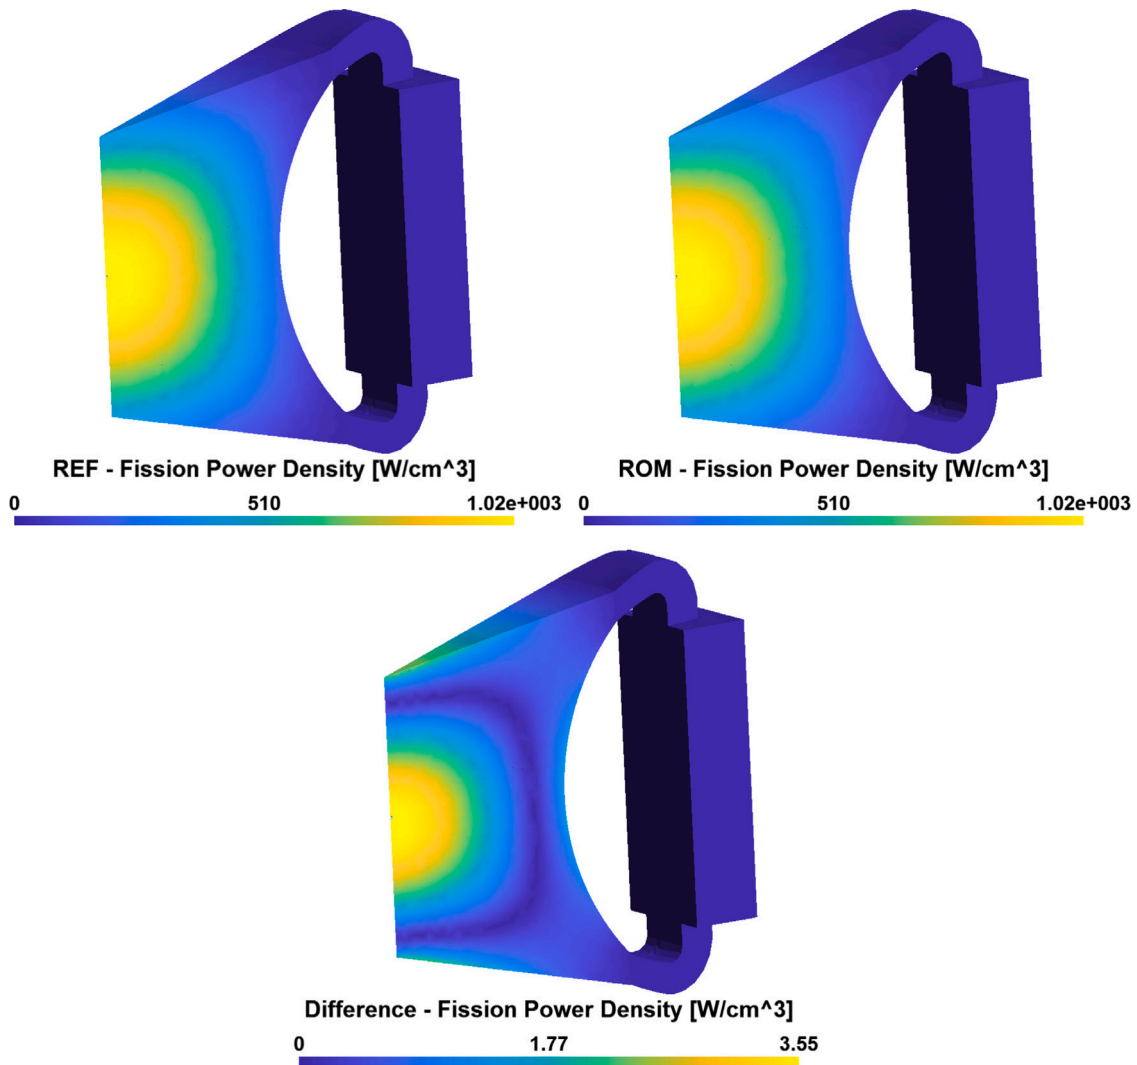
Fig. 6. Fission power density distribution at the point of maximum error showing the reference model (top left), the ROM model (top right), and the distribution of the absolute difference (bottom). The relative ℓ^2 error was 0.24%.

Table 6
Number of unique nodes per dimension for each output. A value of 1 indicates that the output is constant with respect to the parameter. A value of 3 signals that the output is piecewise linear in the parameter. A higher value indicates that output is nonlinear in the parameter.

Parameter	P_{fiss}	T	k_{eff}	Parameter	P_{fiss}	T	k_{eff}
c_p	3	5	3	β_2	1	1	3
κ	1	1	1	β_3	1	1	1
$\Sigma_{f,1}$	1	1	3	β_4	1	1	3
$\Sigma_{f,2}$	3	1	3	β_5	1	1	3
$\Sigma_{f,3}$	3	1	5	β_6	1	1	1
$\Sigma_{f,4}$	3	1	5	β_7	1	1	1
$\Sigma_{f,5}$	3	1	7	β_8	1	1	1
$\Sigma_{f,6}$	3	1	5	λ_1	1	1	1
$\Sigma_{c,1}$	1	1	3	λ_2	1	1	1
$\Sigma_{c,2}$	1	1	3	λ_3	1	1	1
$\Sigma_{c,3}$	1	1	5	λ_4	1	1	1
$\Sigma_{c,4}$	1	1	5	λ_5	1	1	1
$\Sigma_{c,5}$	1	1	9	λ_6	1	1	1
$\Sigma_{c,6}$	3	1	7	λ_7	1	1	1
β_1	1	1	1	λ_8	1	1	1

multiplication factor with a nominal value of 1.00999. To extract the probability density function (PDF) of the response, we run the reduced model at randomly sampled points drawn from the distribution of the input parameters. A histogram of the response is an approximation of the PDF. Fig. 8 shows the normalized density histograms for the maximum temperature and multiplication factor by using 100,000 random points.

The maximum temperature has a mean of 1085×10^0 K with a standard deviation of 6.9×10^0 K. The distribution is close to a normal distribution with a slight skew to lower temperatures. Fig. 9 shows the normality plot of the data, which is a measure of the degree of deviation of the data from the normal distribution. The normal probability plot shows that this skew is mostly observed in lower probabilities of temperatures between 1100 K to 1108 K. The tails deviate from the normal distribution due to the truncation of the normal distribution in the input parameters. The deviation from the normal distribution confirms the nonlinearity of the temperature with respect to c_p, as observed by the algorithm in Table 6. The multiplication factor has a mean of 1.01009, with a standard deviation of 0.01899. The distribution of k_{eff} follows a perfect normal distribution, which is confirmed by the normal probability plot in Fig. 9.

In order to study the sensitivity of the response to the input parameters, the first order Sobol indices are computed from the reduced models. First order Sobol indices provide a measure for the partial contribution of a single parameter to the response of interest. However, from the number of unique nodes per dimension given in Table 6, the reduced model for temperature was shown to be sensitive only to
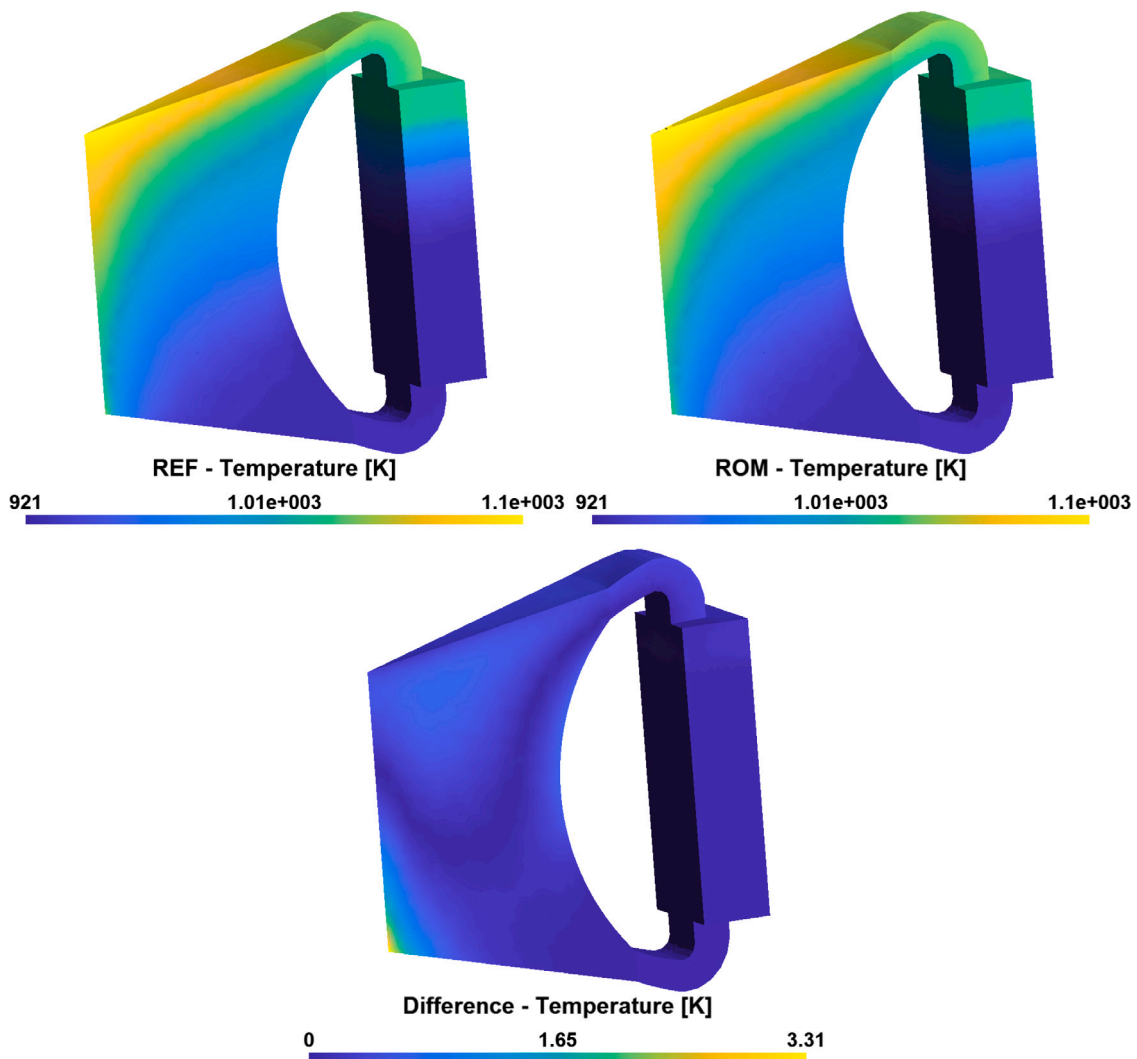
Fig. 7. Temperature distribution at the point of maximum error showing the reference model (top left), the ROM model (top right), and the distribution of the absolute difference (bottom). The relative ℓ^2 error was 0.02%.
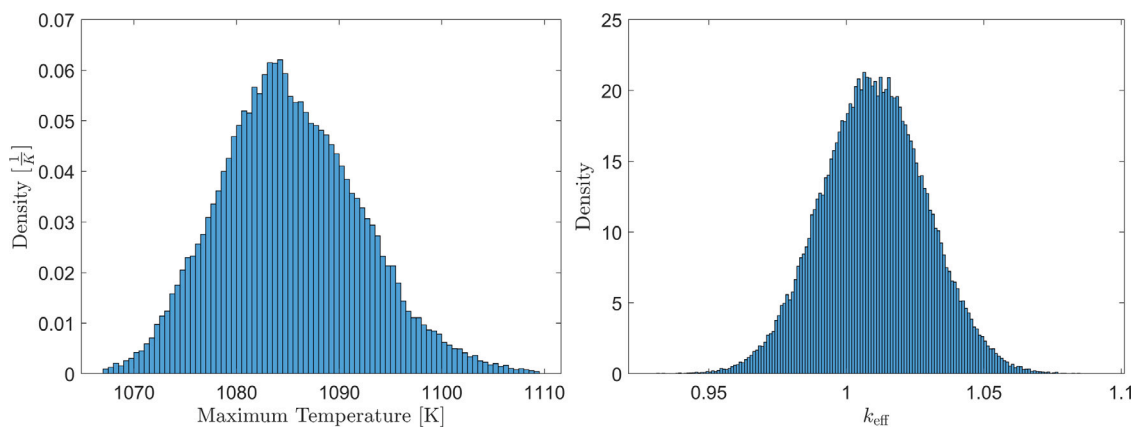

Fig. 8. Normalized density histograms for the maximum temperature (left) and the multiplication factor k_{eff} (right). The data was generated by sampling the reduced model with 100,000 random points drawn from the distribution of the input parameters.

one parameter (c_p). Therefore, we only compute Sobol indices for k_{eff}. We used a quasi-random sampling Sobol sequence (Sobol, 2001). The sequence was generated on the unit hypercube, then mapped to the distribution of each parameter using inverse sampling of the cumulative distribution function (CDF). Two sets of sampling points each of size

10^5 were used to compute the Sobol indices as given by Saltelli et al. (2010). The Sobol indices in Fig. 10 show k_{eff} to be sensitive only to the cross sections. The indices ranked the fission and capture cross sections of group 5 to be the most significant, followed by groups 3 and 4. The fast groups (groups 1 and 2) and the most thermal groups have
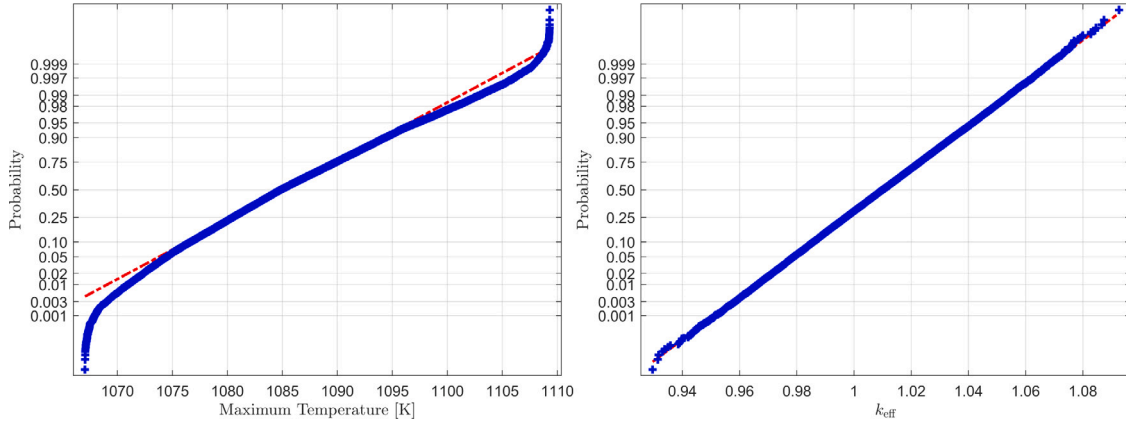
Fig. 9. Normal probability plots for the maximum temperature (left) and the multiplication factor k_{eff} (right).
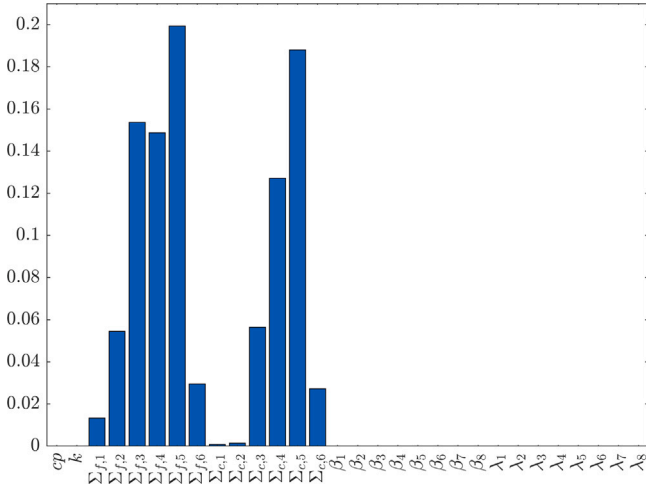

Fig. 10. First order Sobol indices showing the first order sensitivities of k_{eff} to each input parameter. The sum of the first order Sobol indices is close to 1.

relatively lower impact on k_{eff}. This can be explained by the fact that the reactor operates with an epithermal spectrum. This is confirmed from the average flux value per group in Table 3. The indices also show that, in general, fission cross sections have higher importance on k_{eff} compared to the capture cross sections. This is expected because fission cross sections have a direct impact on the multiplication factor while the capture cross section impacts k_{eff} through the loss term which also includes the system leakages. The sum of the first order Sobol indices was found to be close to 1, which suggests that the effect of higher-order interactions between parameters can be neglected.

6. Transient analysis

In this section, we consider the transient model of the MSFR where the reactor power is no longer fixed. The reactor is kept at steady-state by dividing the fission operator by nominal value of k_{eff}. Therefore, the initial conditions for the transients are the nominal values. We are interested in a reduced model for control and simulation purposes, capturing the dynamics of the fission power density and temperature with respect to perturbations in the salt flow rate of the intermediate circuit. This case is simulating an operational scenario where the reactor power is controlled through adjustments in the flow rate of the salt in the intermediate circuit. The intermediate circuit extracting heat from the heat exchangers was not modeled explicitly. However, for the transient analysis, we simulate the effect of controlling the reactor

power through the salt flow rate in the intermediate circuit by adjusting the average intermediate salt temperature and heat transfer coefficient. Since we are interested in operational conditions, we employ a simple linear empirical model relating changes in the flow rate of the salt in the intermediate circuit to changes in the average intermediate salt temperature and heat transfer coefficients as

$$\Delta T_{\text{int}} = -0.375 \ \Delta q, \tag{16}$$

$$\Delta h_{\text{int}} = 74825 \Delta q, \tag{17}$$

where ΔT_{int} is the change in intermediate salt temperature in Kelvin, Δh_{int} is the change in the heat transfer coefficients in units of MW/m^{3}K, and Δq is the percentage change in the flow rate of the salt in the intermediate circuit. In order to approximate the dynamics of controlling the salt flow rate, these perturbations are introduced in the model exponentially with a time constant of 10×10^0 s. Therefore, the models for the average intermediate salt temperature and heat transfer coefficient are

$$T_{\text{int}}(t) = T_{\text{int},0} + \Delta T_{\text{int}} \left(1 - \exp\left(\frac{-t}{10} \right) \right), \tag{18}$$

$$h_{\text{int}}(t) = h_{\text{int},0} + \Delta h_{\text{int}} \left(1 - \exp\left(\frac{-t}{10} \right) \right), \tag{19}$$

where $T_{\text{int}}(t)$ and $h_{\text{int}}(t)$ are respectively the intermediate salt temperature and heat transfer coefficient as functions of time t, $T_{\text{int},0}$ is the nominal value of the intermediate salt temperature with a value of 908.15×10^0 K (as reported in Table 1), and $h_{\text{int},0}$ is the nominal value of the heat transfer coefficient with a value of 19.95×10^0 MW/m^{3}K (Table 1).

Because we considered time as an input parameter in our formulation, this model has two input parameters, the change in the intermediate circuit flow rate (Δq) and time (t). The outputs are the fission power density, which is computed on the neutronics mesh with 104288 unknowns, and the temperature, which is computed on the CFD mesh with 220972 unknowns. We considered the flow rate to range between −40% to +15% of the nominal values. The time range was taken $t \in [0, 180]$ s — that is, we are interested in constructing a reduced model for the first 180×10^0 s after a perturbation in the intermediate flow rate Δq.

A transient of increasing the flow rate by 15% is shown in Fig. 11, along with snapshots of the final solutions at $t = 180 \times 10^0$ s. Increasing the flow rate of the intermediate circuit causes the temperature of salt flowing back into the cavity from the heat exchanger to drop, which is observed as a decrease of the minimum and average temperatures. The lower temperature in the reactor core introduces positive reactivity because of the strong negative temperature feedback of this reactor (Heuer et al., 2014). For this reason, the trends show an increase in the reactor power and the maximum temperature registered at the top
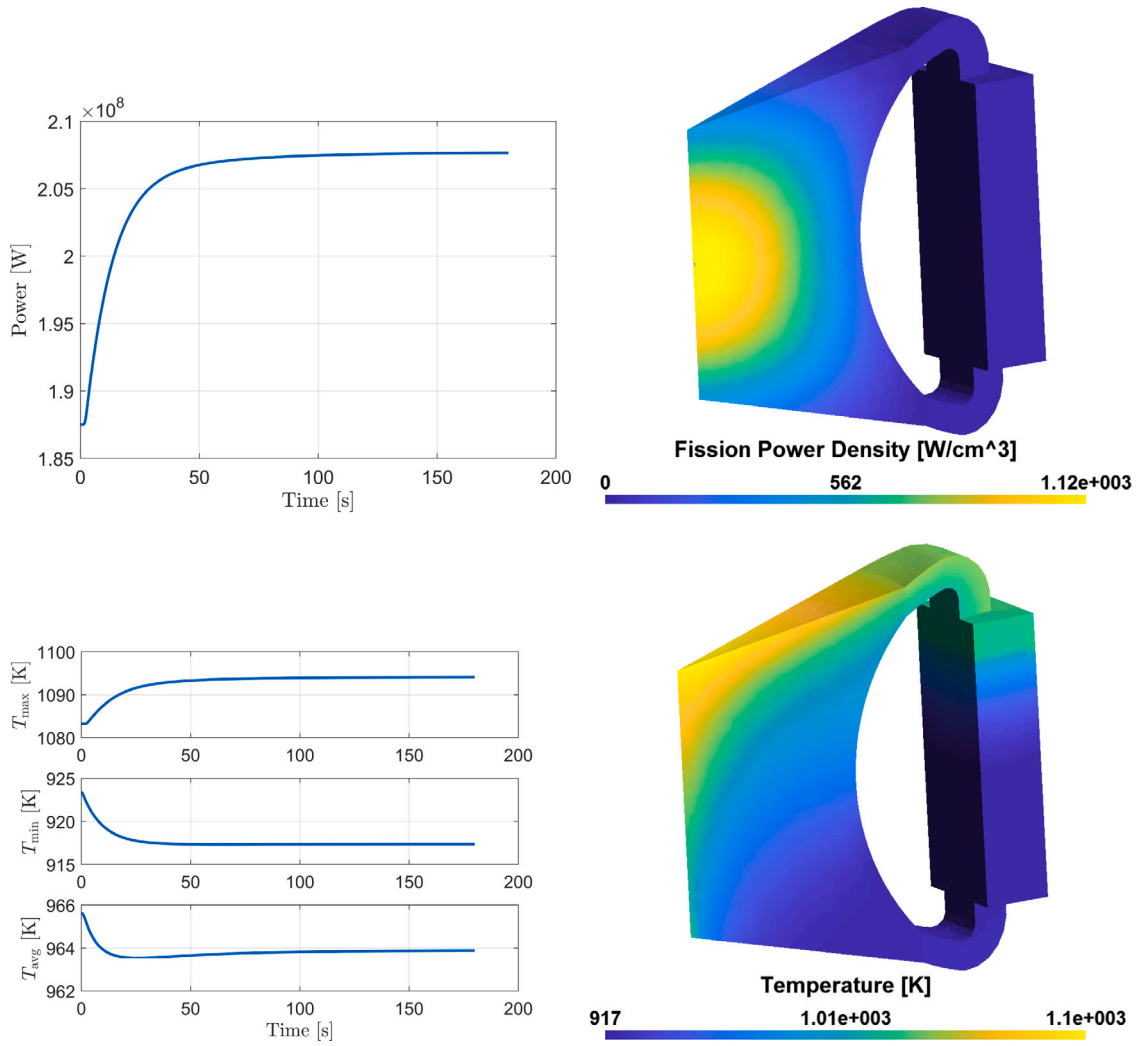
Fig. 11. A selected transient resulting from $\Delta q = 15\%$ showing the reactor power (top left), fission power density distribution at $t = 180 \times 10^0$ s (top right), temperature trends for the maximum temperature T_{max}, minimum temperature T_{min}, and average temperature T_{avg} (bottom left), and temperature distribution at $t = 180 \times 10^0$ s (bottom right).
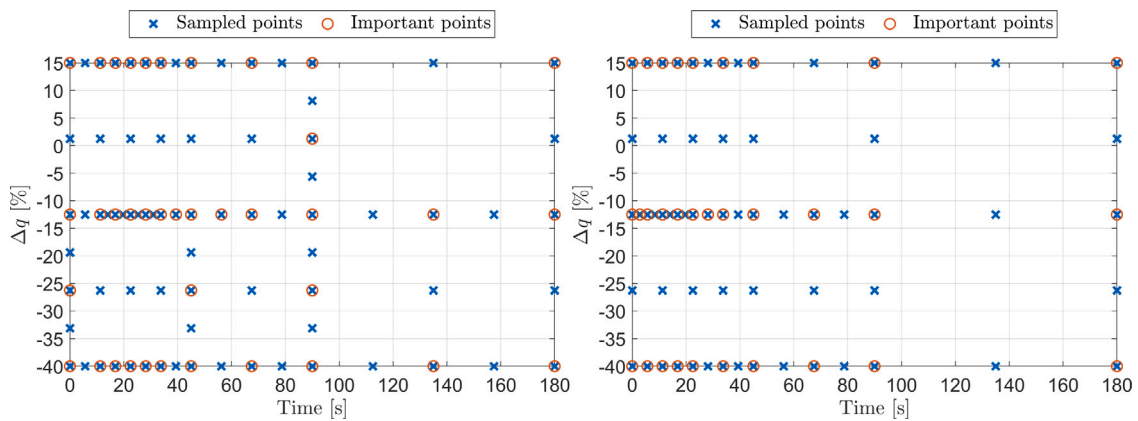

Fig. 12. The generated points for the transient model of the fission power density (left) and temperature (right). The important points included in the final $\mathcal{X}^k$ are marked with a red circle.

of the reactor cavity. The average temperature then gradually adjust the downward trend, and a new steady-state is reached.

To build the reduced model, we set the tolerances $\zeta = 1\%$ and $\gamma_{int} = 0.1\%$ for the fission power density, and $\zeta = 0.1\%$ and $\gamma_{int} = 0.01\%$ for the temperature. The POD truncation tolerance γ_{tr} was set to be 1×10^{-12} for both models. A summary of the results is given

in Table 7. The algorithm required 9 simulations (corresponding to 9 values of Δq) for the fission power density and selected a total of 78 snapshots to build the reduced model. For the temperature, the algorithm required 5 simulations and selected 62 snapshots. In the transient case, we define a *snapshot* to be the model output taken at a certain time and fixed parameter configuration whereas a *simulation*
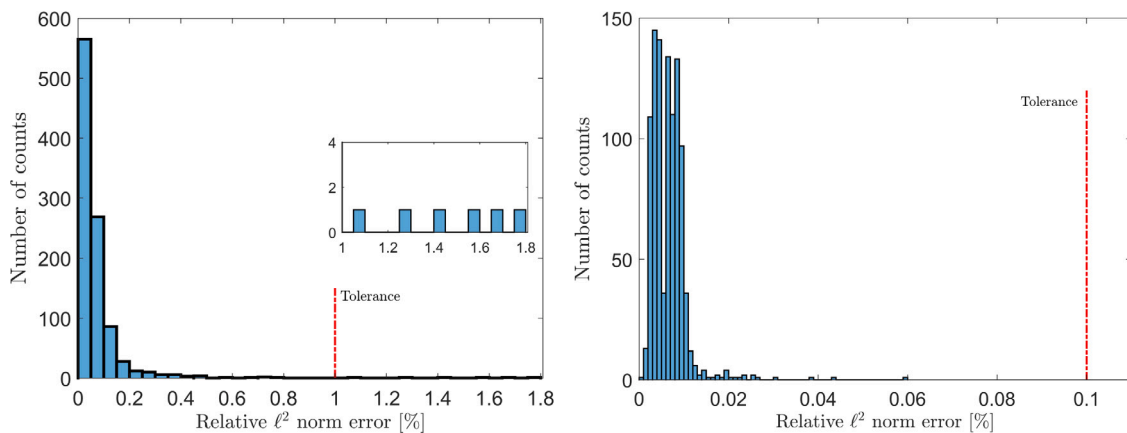
Fig. 13. Histogram of the relative ℓ^2 error in the fission power density model (left) and temperature model (right). A close up on the region above 1% is shown for the histogram of the fission power density model.
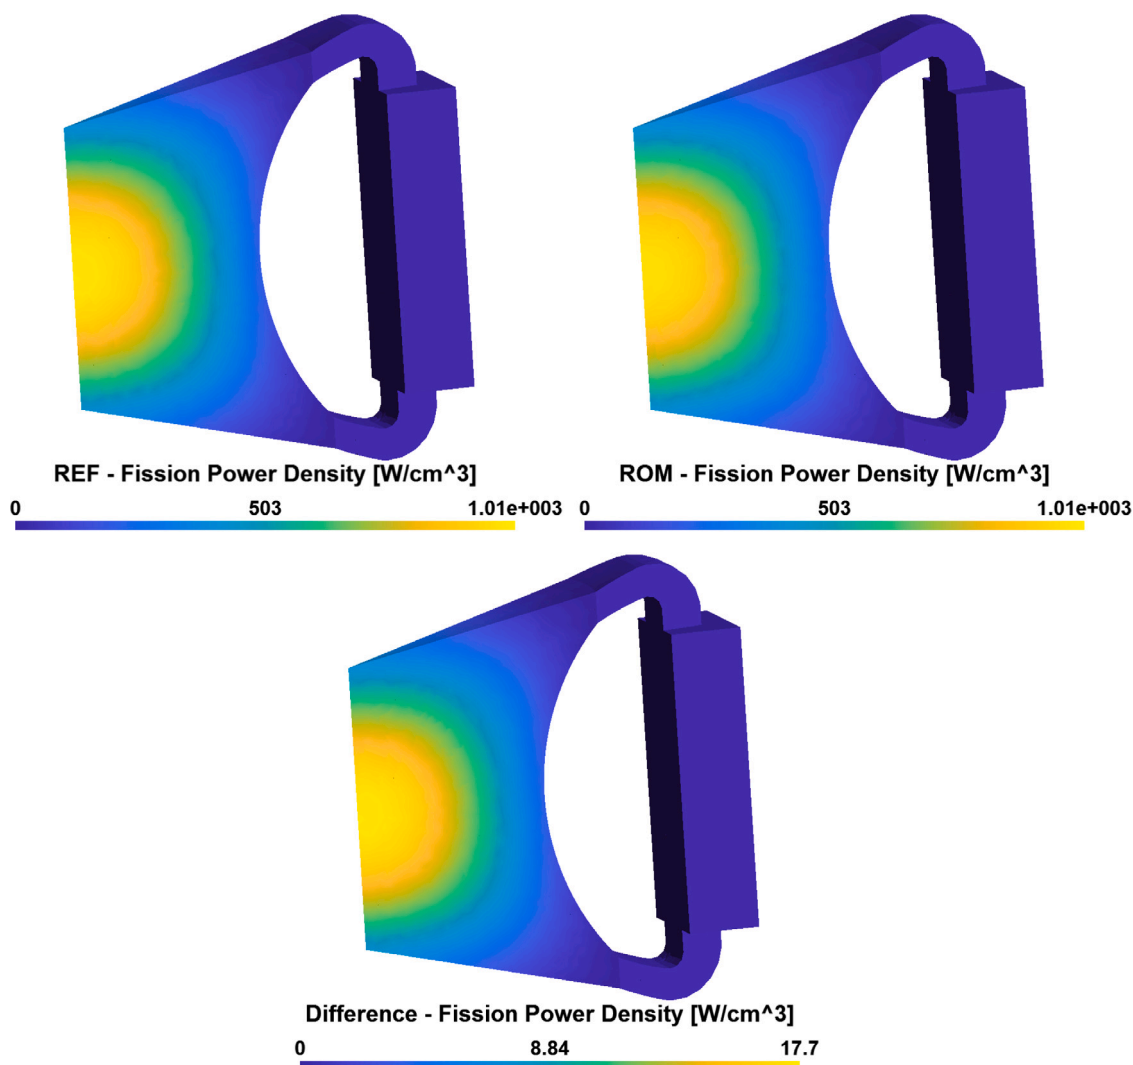

REF - Fission Power Density [W/cm^3]

ROM - Fission Power Density [W/cm^3]

Difference - Fission Power Density [W/cm^3]

Fig. 14. Fission power density distribution at the point of maximum error showing the reference model (top left), the ROM model (top right), and the distribution of the absolute difference (bottom). This point correspond to $\Delta q = -36.5$ and at time $t = 2.4 \times 10^0$ s. The relative ℓ^2 error was 0.24%.

is the execution of the model at a fixed parameter configuration to produce the evolution of the output with respect to time. The selected points for both models are shown in Fig. 12. It can be seen that the algorithm selected more snapshots in the transient region (the first 80 s) and fewer points towards the steady-state region. Note that for the

fission power density model, not all simulations were run full transient from $t = 0$ to 180×10^0 s. Only 5 full simulations were required while 4 simulations ran only to $t = 90 \times 10^0$ s. Given that a full simulation to $t = 180 \times 10^0$ s requires about 150 CPU-hours, such adaptivity significantly improves the efficiency of constructing the reduced model compared
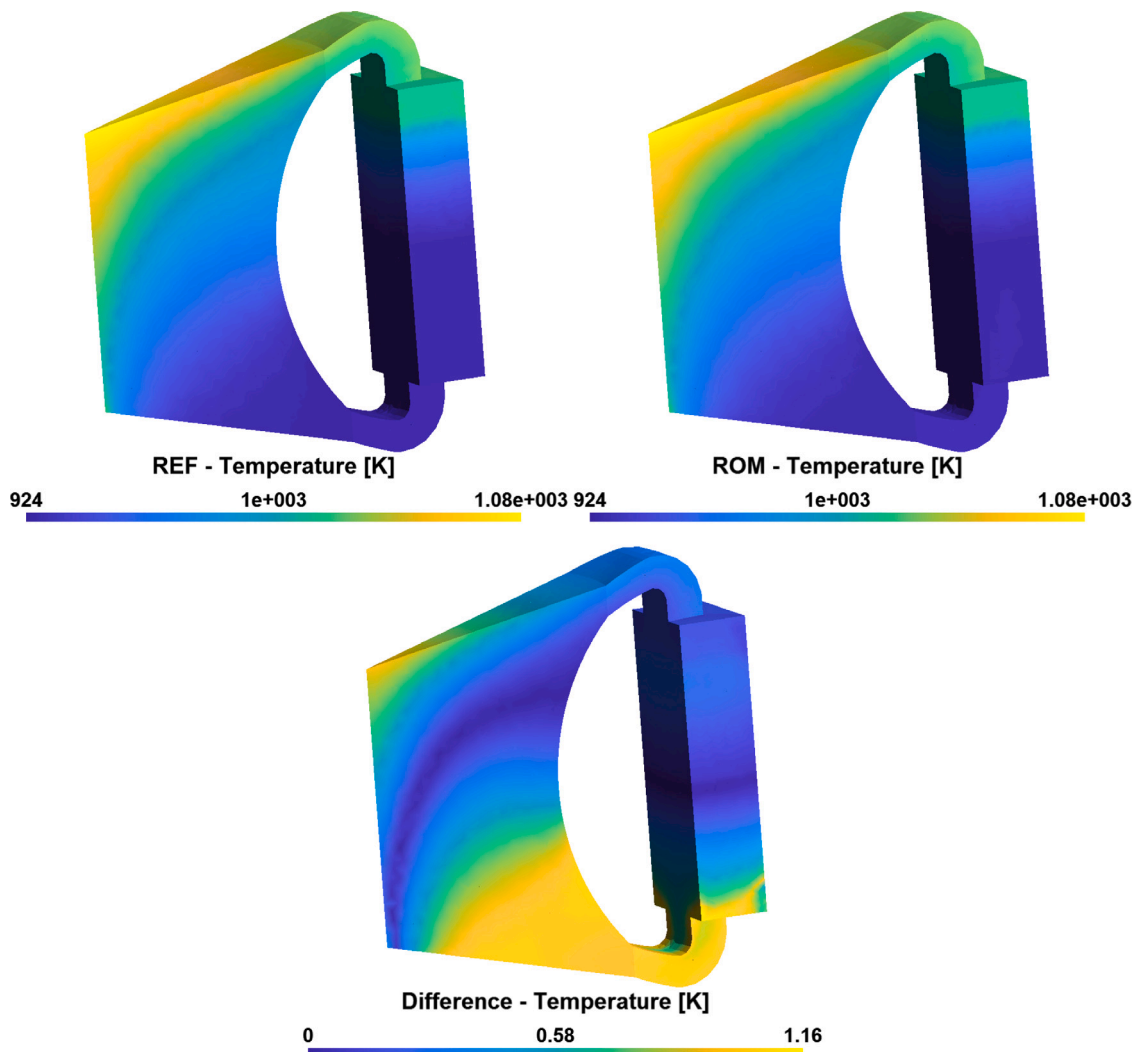
Fig. 15. Fission power density distribution at the point of maximum error showing the reference model (top left), the ROM model (top right), and the distribution of the absolute difference (bottom). This point correspond to $\Delta q = -36.5$ and at time $t = 1.5 \times 10^0$ s. The relative ℓ^2 error was 0.06%.

to an a priori snapshot selection approach, which does not consider the actual dynamics of the system.

To test the constructed reduced models, we ran 8 high-fidelity test simulations with values of Δq at half the distances between the points selected during the construction phase. In this manner, we maximize the distance (along Δq dimension) between the test simulations and the simulations used for the construction of the reduced models. We then select 1000 random snapshots at different times from the testing simulations, which are used to test the reduced model. A histogram of the relative ℓ^2 norm error is shown in Fig. 13 for both models. The error in the temperature model was found to be well below the set tolerance, with the maximum error being 0.06%. For the fission power density model, most of the points resulted in an error below the set tolerance. However, 6 out of the 1000 points were above the tolerance, with the maximum being 1.8%. The point that resulted in the maximum error is compared with the high-fidelity solution in Fig. 14 for the fission power density and Fig. 15 for the temperature. In this case, the fission power density comparison shows the maximum difference to be at the center of the cavity, where the flux is maximum. The local relative error of at that point is 1.8%. The maximum temperature difference, on the other hand, is observed at the inlet of the reactor core where the temperature is minimum, with a local relative error of 0.1%. A full simulation to $t = 180 \times 10^0$ s is completed with the reduced model in under 5 s, while the high-fidelity model needed 150 CPU hours on a high performance computing unit.

Table 7

Summary of results for the transient reduced models corresponding to the fission power density P_{fiss} and temperature T showing the number of required simulations to construct each model, number of snapshots selected, the final number of points in $\mathcal{X}^k$, the number of POD modes after truncation, and the maximum relative ℓ^2 error after a test on 1000 random points.

	P_{fiss}	T
Number of simulations	9	5
Number of snapshots	78	62
Number of selected points in $\mathcal{X}^k$	38	31
Number of POD modes	5	6
Maximum relative ℓ^2 error	1.8%	0.06%

7. Conclusions

We have applied an adaptive POD approach to a large-scale, three-dimensional model of the MSFR. The developed algorithm was able to construct reduced models for both steady-state and transient analysis of the reactor. The steady-state analysis considered higher dimensional input space with 30 parameters. Three reduced models were constructed to capture the effects of those parameters on the fission power density and temperature distributions, and on the multiplication factor. Each model was tested on 1000 random points that were not part of the snapshot selections. The model for fission power density required 61

high-fidelity evaluations and resulted in a maximum relative ℓ^2 error of 0.24%. The model for the temperature required 63 evaluations and resulted in a maximum error of 0.02%. The multiplication factor model needed 1639 evaluations and the maximum error was 37 pcm. Evaluating the reduced models was completed in less than a second while the high-fidelity model needed 4 CPU-hours.

The reduced models were then used to propagate uncertainties in the input parameters to the maximum temperature of the reactor and the multiplication factor. The constructed PDF showed the multiplication factor to follow a normal distribution with a mean 1.01009 and standard deviation of 0.01899. The maximum temperature had a mean of 1085×10^0 K, a standard deviation of 6.9×10^0 K, and showed a distribution close to normal with a slight skew to lower temperatures. The sensitivity study concluded that the maximum temperature of the MSFR is only sensitive to the specific heat capacity within the defined tolerances. Heat conductivity and neutronics data had no significant effect on the temperature. Therefore, since the salt properties of the MSFR are still under investigation, we recommend prioritizing studies to reduce uncertainties in the specific heat capacity over heat conductivity. The multiplication factor was shown to be only sensitive to the cross sections. A standard deviation of 5% in the decay constant and delayed neutron fraction is sufficient to characterize the multiplication factor within 50 pcm error. However, a standard deviation of 5% in the cross sections resulted in a distribution of the multiplication factor with a standard deviation of 1899 pcm, which is significant. Therefore, uncertainties in cross section data should be below 5% standard deviation for this reactor.

The transient analysis studied the effect of controlling the flow of salt in the intermediate circuit on the fission power density and temperature distributions. Our approach considers time to be a parameter in input space in order to allow for an adaptive selection of the snapshots. The model for the fission power density needed 9 simulations and selected 78 snapshots. The test on 1000 independent points showed the maximum relative ℓ^2 error to be 1.8%. The model for the temperature required 5 simulations and selected 62 snapshots. The maximum error was found to be 0.06% for this model. The selected points showed that the algorithm sampled the high-fidelity model in regions of the beginning of the transient more than the steady-state. This allowed the algorithm to be efficient by simulating some of the points only to half the transient. A full simulation to the end of the transient required about 150 CPU-hours. Therefore, reducing the transient time for a point is a massive saving in computational resources. The constructed reduced models were able to produce a full simulation in less than 5 s. As follow-up work, this reduced model can be used to design a controller for the reactor. Other future applications could also include transients with significant primary circuit flow field changes (such as loss of pump power or pipe blockage events), where the temperature, neutronics and flow fields simultaneously have to be reduced. Since our algorithm is fully nonintrusive and can handle multiple fields, addressing such transients is straightforward and we expect our model to show comparable performance.

In all models, the number of points included in the final grid set $\mathcal{X}^k$ was a small fraction of the total snapshots generated by the algorithm. This fraction was about 50% for the transient models while for the steady-state models, it was found to be as low as 10%. This is an indication that the algorithm is still oversampling. The number of POD modes was found to be even a smaller fraction of the snapshots, which is a signal that most of the points were generated to train the surrogate model rather than discover new dynamics of the system. Therefore, the algorithm could be improved further with more advanced surrogate models to reduce the number of sampling points. The use of higher-order basis functions is a potential area of study for this propose. Moreover, our approach to the transient problem assumes the initial conditions to be parametrized. The approach is not yet applicable for problems with a time-dependent input signal, which cannot be parametrized. These areas of research are the subjects of future work.

CRediT authorship contribution statement

Fahad Alsayyari: Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing – original draft, Visualization. **Marco Tiberga:** Software, Formal analysis, Writing – original draft, Visualization. **Zoltán Perkó:** Conceptualization, Methodology, Writing – review & editing, Supervision. **Jan Leen Kloosterman:** Resources, Writing – review & editing, Supervision, Project administration, Funding acquisition. **Danny Lathouwers:** Conceptualization, Methodology, Resources, Writing – review & editing, Supervision.

Declaration of competing interest

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

Acknowledgments

References

Abdel-Khalik, H.S., Bang, Y., Wang, C., 2013. Overview of hybrid subspace methods for uncertainty quantification, sensitivity analysis. Ann. Nucl. Energy 52, 28–46. http://dx.doi.org/10.1016/j.anucene.2012.07.020.

Abdo, M., Elzohery, R., Roberts, J.A., 2019. Modeling isotopic evolution with surrogates based on dynamic mode decomposition. Ann. Nucl. Energy 129, 280–288. http://dx.doi.org/10.1016/j.anucene.2019.01.048.

Allibert, M., Aufiero, M., Brovchenko, M., Delpech, S., Ghetta, V., Heuer, D., Laureau, A., Merle-Lucotte, E., 2016. Molten salt fast reactors. In: Pioro, I.L. (Ed.), Handbook of Generation IV Nuclear Reactors. Woodhead Publishing, pp. 157–188. http://dx.doi.org/10.1016/B978-0-08-100149-3.00007-0.

Alsayyari, F., Perkó, Z., Lathouwers, D., Kloosterman, J.L., 2019. A nonintrusive reduced order modelling approach using proper orthogonal decomposition and locally adaptive sparse grids. J. Comput. Phys. 399, 108912. http://dx.doi.org/10.1016/j.jcp.2019.108912.

Alsayyari, F., Perkó, Z., Tiberga, M., Kloosterman, J.L., Lathouwers, D., 2021. A fully adaptive nonintrusive reduced-order modelling approach for parametrized time-dependent problems. Comput. Methods Appl. Mech. Engrg. 373, 113483. http://dx.doi.org/10.1016/j.cma.2020.113483.

Alsayyari, F., Tiberga, M., Perkó, Z., Lathouwers, D., Kloosterman, J.L., 2020. A nonintrusive adaptive reduced order modeling approach for a molten salt reactor system. Ann. Nucl. Energy 141, 107321. http://dx.doi.org/10.1016/j.anucene.2020.107321.

Amsallem, D., Farhat, C., 2012. Stabilization of projection-based reduced-order models. Internat. J. Numer. Methods Engrg. 91 (4), 358–377. http://dx.doi.org/10.1002/nme.4274.

Antoulas, A.C., Sorensen, D.C., Gugercin, S., 2001. A survey of model reduction methods for large-scale systems. Contemp. Math. 280, 193–219.

Aufiero, M., Cammi, A., Geoffroy, O., Losa, M., Luzzi, L., Ricotti, M.E., Rouch, H., 2014. Development of an OpenFOAM model for the molten salt fast reactor transient analysis. Chem. Eng. Sci. 111, 390–401. http://dx.doi.org/10.1016/j.ces.2014.03.003.

Bang, Y., Abdel-Khalik, H.S., Jessee, M.A., Mertyurek, U., 2015. Hybrid reduced order modeling for assembly calculations. Nucl. Eng. Des. 295, 661–666. http://dx.doi.org/10.1016/j.nucengdes.2015.07.020.

Beneš, O., Salanne, M., Levesque, M., Konings, R.J.M., 2013. Physico-Chemical Properties of the MSFR Fuel Salt — Deliverable D3.2. Tech. Rep., EVOL Project.

Benner, P., Gugercin, S., Willcox, K., 2015. A survey of projection-based model reduction methods for parametric dynamical systems. SIAM Rev. 57 (4), 483–531.

Buchan, A.G., Pain, C.C., Fang, F., Navon, I.M., 2013. A POD reduced-order model for eigenvalue problems with application to reactor physics. Internat. J. Numer. Methods Engrg. 95 (12), 1011–1032. http://dx.doi.org/10.1002/nme.4533.

Cacuci, D., 2003. Sensitivity & Uncertainty Analysis, vol. 1. Chapman and Hall/CRC, http://dx.doi.org/10.1201/9780203498798.

Castagna, C., Aufiero, M., Lorenzi, S., Lomonaco, G., Cammi, A., 2020. Development of a reduced order model for fuel burnup analysis. Energies 13 (4), 890. http://dx.doi.org/10.3390/en13040890.

Cervi, E., Lorenzi, S., Cammi, A., Luzzi, L., 2019. Development of a multiphysics model for the study of fuel compressibility effects in the molten salt fast reactor. Chem. Eng. Sci. 193, 379–393. http://dx.doi.org/10.1016/j.ces.2018.09.025.

Cherezov, A., Sanchez, R., Joo, H.G., 2018. A reduced-basis element method for pin-by-pin reactor core calculations in diffusion and SP_3 approximations. Ann. Nucl. Energy 116, 195–209. http://dx.doi.org/10.1016/j.anucene.2018.02.013.

Delpech, S., Merle-Lucotte, E., Heuer, D., Allibert, M., Ghetta, V., Le-Brun, C., Doligez, X., Picard, G., 2009. Reactor physic and reprocessing scheme for innovative molten salt reactor system. J. Fluor. Chem. 130 (1), 11–17. http://dx.doi.org/10.1016/j.jfluchem.2008.07.009.

Dolan, T.J., 2017. Molten Salt Reactors and Thorium Energy. In: Energy, Woodhead Publishing, Elsevier Ltd., Cambridge, MA, United States.

Escanciano, J.Y., Class, A.G., 2019. POD-Galerkin modeling of a heated pool. Prog. Nucl. Energy 113, 196–205. http://dx.doi.org/10.1016/j.pnucene.2019.01.017.

Fiorina, C., Lathouwers, D., Aufiero, M., Cammi, A., Guerrieri, C., Kloosterman, J.L., Luzzi, L., Ricotti, M.E., 2014. Modelling and analysis of the MSFR transient behaviour. Ann. Nucl. Energy 64, 485–498. http://dx.doi.org/10.1016/j.anucene.2013.08.003.

Generation IV International Forum, 2002. A Technology Roadmap for Generation IV Nuclear Energy Systems. Tech. Rep. GIF-002–00, U.S. DOE Nuclear Energy Research Advisory Committee and the Generation IV International Forum, http://dx.doi.org/10.2172/859029.

Gérardin, D., Allibert, M., Heuer, D., Laureau, A., Merle-Lucotte, E., Seuvre, C., 2017. Design evolutions of molten salt fast reactor. In: International Conference on Fast Reactors and Related Fuel Cycles: Next Generation Nuclear Systemes for Sustainable Development. FR17, Yekaterinburg, Russia. URL http://hal.in2p3.fr/in2p3-01572660.

German, P., Ragusa, J.C., 2019. Reduced-order modeling of parameterized multi-group diffusion k-eigenvalue problems. Ann. Nucl. Energy 134, 144–157. http://dx.doi.org/10.1016/j.anucene.2019.05.049.

Gilli, L., Lathouwers, D., Kloosterman, J.L., van der Hagen, T.H.J.J., 2013. Applying second-order adjoint perturbation theory to time-dependent problems. Ann. Nucl. Energy 53, 9–18. http://dx.doi.org/10.1016/j.anucene.2012.10.003.

Hennink, A., Tiberga, M., Lathouwers, D., 2021. A pressure-based solver for low-Mach number flow using a discontinuous Galerkin method. J. Comput. Phys. 425, 109877. http://dx.doi.org/10.1016/j.jcp.2020.109877.

Heuer, D., Merle-Lucotte, E., Allibert, M., Brovchenko, M., Ghetta, V., Rubiolo, P., 2014. Towards the thorium fuel cycle with molten salt fast reactors. Ann. Nucl. Energy 64, 421–429. http://dx.doi.org/10.1016/j.anucene.2013.08.002.

Huang, D., Abdel-Khalik, H., Rabiti, C., Gleicher, F., 2017. Dimensionality reducibility for multi-physics reduced order modeling. Ann. Nucl. Energy 110, 526–540. http://dx.doi.org/10.1016/j.anucene.2017.06.045.

International Atomic Energy Agency, 2019. Deterministic Safety Analysis for Nuclear Power Plants. In: Specific Safety Guides, SSG-2 (Rev.1), IAEA, Vienna, URL https://www.iaea.org/publications/12335/deterministic-safety-analysis-for-nuclear-power-plants.

Kópházi, J., Lathouwers, D., 2012. Three-dimensional transport calculation of multiple alpha modes in subcritical systems. Ann. Nucl. Energy 50, 167–174. http://dx.doi.org/10.1016/j.anucene.2012.06.021.

Laureau, A., Heuer, D., Merle-Lucotte, E., Rubiolo, P.R., Allibert, M., Aufiero, M., 2017. Transient coupled calculations of the molten salt fast reactor using the transient fission matrix approach. Nucl. Eng. Des. 316, 112–124. http://dx.doi.org/10.1016/j.nucengdes.2017.02.022.

Leppänen, J., Pusa, M., Viitanen, T., Valtavirta, V., Kaltiaisenaho, T., 2015. The serpent Monte Carlo code: Status, development and applications in 2013. Ann. Nucl. Energy 82, 142–150. http://dx.doi.org/10.1016/j.anucene.2014.08.024.

Manthey, R., Knospe, A., Lange, C., Hennig, D., Hurtado, A., 2019. Reduced order modeling of a natural circulation system by proper orthogonal decomposition. Prog. Nucl. Energy 114, 191–200. http://dx.doi.org/10.1016/j.pnucene.2019.03.010.

Marta, A.C., Mader, C.A., Martins, J.R.R.A., Van der Weide, E., Alonso, J.J., 2007. A methodology for the development of discrete adjoint solvers using automatic differentiation tools. Int. J. Comput. Fluid Dyn. 21 (9–10), 307–327. http://dx.doi.org/10.1080/10618560701678647.

Perkó, Z., Gilli, L., Lathouwers, D., Kloosterman, J.L., 2014a. Grid and basis adaptive polynomial chaos techniques for sensitivity and uncertainty analysis. J. Comput. Phys. 260, 54–84. http://dx.doi.org/10.1016/j.jcp.2013.12.025, URL http://www.sciencedirect.com/science/article/pii/S0021999113008322.

Perkó, Z., Lathouwers, D., Kloosterman, J.L., van der Hagen, T., 2013. Adjoint-based sensitivity analysis of coupled criticality problems. Nucl. Sci. Eng. 173 (2), 118–138. http://dx.doi.org/10.13182/nse12-18.

Perkó, Z., Lathouwers, D., Kloosterman, J.L., van der Hagen, T., 2014b. Large scale applicability of a fully adaptive non-intrusive spectral projection technique: Sensitivity and uncertainty analysis of a transient. Ann. Nucl. Energy 71, 272–292. http://dx.doi.org/10.1016/j.anucene.2014.03.035.

Prill, D.P., Class, A.G., 2014. Semi-automated proper orthogonal decomposition reduced order model non-linear analysis for future BWR stability. Ann. Nucl. Energy 67, 70–90. http://dx.doi.org/10.1016/j.anucene.2013.11.022.

Prince, Z.M., Ragusa, J.C., 2020. Application of proper generalized decomposition to multigroup neutron diffusion eigenvalue calculations. Prog. Nucl. Energy 121, 103232. http://dx.doi.org/10.1016/j.pnucene.2019.103232.

Ronco, A.D., Introini, C., Cervi, E., Lorenzi, S., Jeong, Y.S., Seo, S.B., Bang, I.C., Giacobbo, F., Cammi, A., 2020. Dynamic mode decomposition for the stability analysis of the molten salt fast reactor core. Nucl. Eng. Des. 362, 110529. http://dx.doi.org/10.1016/j.nucengdes.2020.110529.

Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M., Tarantola, S., 2010. Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index. Comput. Phys. Comm. 181 (2), 259–270. http://dx.doi.org/10.1016/j.cpc.2009.09.018.

Santamarina, A., Bernard, D., Blaise, P., Coste, M., Courcelle, A., Huynh, T.D., Jouanne, C., Leconte, P., Litaize, O., Ruggiéri, J.-M., Sérot, O., Tommasi, J., Vaglio, C., Vidal, J.-F., 2009. the JEFF-3.1.1 Nuclear Data Library. JEFF Report 22, NEA–OECD.

Santanoceto, M., Tiberga, M., Perkó, Z., Dulla, S., Lathouwers, D., 2021a. Preliminary uncertainty and sensitivity analysis of the molten salt fast reactor steady-state using a polynomialchaos expansion method. Ann. Nucl. Energy.

Santanoceto, M., Tiberga, M., Perkó, Z., Dulla, S., Lathouwers, D., 2021b. Uncertainty quantification in steady state simulations of a molten salt system using polynomial chaos expansion analysis. In: Margulis, M., Blaise, P. (Eds.), EPJ Web Conf. 247, 15008. http://dx.doi.org/10.1051/epjconf/202124715008.

Sartori, A., Cammi, A., Luzzi, L., Rozza, G., 2016. A multi-physics reduced order model for the analysis of lead fast reactor single channel. Ann. Nucl. Energy 87, 198–208. http://dx.doi.org/10.1016/j.anucene.2015.09.002.

Schilders, Wilhelmus H. A., van der Vorst, Henk A., Rommes, Joost (Eds.), 2008. Model Order Reduction: Theory, Research Aspects and Applications. Springer Berlin Heidelberg, http://dx.doi.org/10.1007/978-3-540-78841-6.

Senecal, J.P., Ji, W., 2019. Characterization of the proper generalized decomposition method for fixed-source diffusion problems. Ann. Nucl. Energy 126, 68–83. http://dx.doi.org/10.1016/j.anucene.2018.10.062.

Sobol, I.M., 2001. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. Math. Comput. Simulation 55 (1–3), 271–280. http://dx.doi.org/10.1016/s0378-4754(00)00270-6.

Takeishi, N., Kawahara, Y., Tabei, Y., Yairi, T., 2017. Bayesian dynamic mode decomposition. In: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence. International Joint Conferences on Artificial Intelligence Organization, http://dx.doi.org/10.24963/ijcai.2017/392.

Tiberga, M., Hennink, A., Kloosterman, J.L., Lathouwers, D., 2020a. A high-order discontinuous Galerkin solver for the incompressible RANS equations coupled to the $k - \epsilon$ turbulence model. Comput. & Fluids 212, 104710. http://dx.doi.org/10.1016/j.compfluid.2020.104710.

Tiberga, M., Lathouwers, D., Kloosterman, J.L., 2019. A discontinuous Galerkin FEM multiphysics solver for the Molten Salt Fast Reactor. In: International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering. M&C, 2019, Portland, OR, USA.

Tiberga, M., Lathouwers, D., Kloosterman, J.L., 2020b. A multi-physics solver for liquid-fueled fast systems based on the discontinuous Galerkin FEM discretization. Prog. Nucl. Energy 127, 103427. http://dx.doi.org/10.1016/j.pnucene.2020.103427.

Tu, J.H., Rowley, Clarence W., Luchtenburg, D.M., Brunton, S.L., and, J. Nathan Kutz, 2014. On dynamic mode decomposition: Theory and applications. J. Comput. Dyn. 1 (2), 391–421. http://dx.doi.org/10.3934/jcd.2014.1.391.

Vergari, L., Cammi, A., Lorenzi, S., 2020. Reduced order modeling approach for parametrized thermal-hydraulics problems: Inclusion of the energy equation in the POD-FV-ROM method. Prog. Nucl. Energy 118, 103071. http://dx.doi.org/10.1016/j.pnucene.2019.103071.