

A POD reduced-order model for eigenvalue problems with application to reactor physics

A. G. Buchan^{1,*}, C.C. Pain¹, F. Fang¹ and I. M. Navon²

¹*Applied Modelling & Computation Group, Department of Earth Science and Engineering,
Imperial College London, London SW7 2AZ, UK*

²*Department of Scientific Computing, Florida State University, Tallahassee, FL 32306-4120, USA*

SUMMARY

A reduced-order model based on proper orthogonal decomposition (POD) has been presented and applied to solving eigenvalue problems. The model is constructed via the method of snapshots, which is based upon the singular value decomposition of a matrix containing the characteristics of a solution as it evolves through time. Part of the novelty of this work is in how this snapshot data are generated, and this is through the recasting of eigenvalue problem, which is time independent, into a time-dependent form. Instances of time-dependent eigenfunction solutions are therefore used to construct the snapshot matrix. The reduced order model's capabilities in efficiently resolving eigenvalue problems that typically become computationally expensive (using standard full model discretisations) has been demonstrated. Although the approach can be adapted to most general eigenvalue problems, the examples presented here are based on calculating dominant eigenvalues in reactor physics applications. The approach is shown to reconstruct both the eigenvalues and eigenfunctions accurately using a significantly reduced number of unknowns in comparison with 'full' models based on finite element discretisations. The novelty of this paper therefore includes a new approach to generating snapshots, POD's application to large-scale eigenvalue calculations, and reduced-order model's application in reactor physics. Copyright © 2013 John Wiley & Sons, Ltd.

Received 20 August 2012; Revised 12 April 2013; Accepted 10 May 2013

KEY WORDS: POD; reduced-order modelling; eigenvalue problem; reactor criticality

1. INTRODUCTION

Eigenvalue problems arise in many fields of mathematics, science and engineering and are important as they characterise system properties and yield key information of their state. Their uses and applications are diverse (a review is not provided here) but their computation can, in many instances, be prohibitively expensive. Even with today's computational resources, eigenvalue problems prove challenging as they quite often involve complex, coupled and multidimensional systems that require discretisations involving a large number of variables. Much effort has therefore been placed in solving these type of problems, typically with the aid of computers, where algorithm design and efficient/accurate discretisation schemes have all played a central role. The use of reduced-order models (ROMs) [1–3] can potentially be of great benefit in solving eigenvalue problems, as they reduce the problem sizes by several order of magnitude with minimal loss of accuracy. In this article, a new ROM is presented for eigenvalue problems, and whilst it is demonstrated through the solving of reactor physics (RP) applications, it is designed to be general in the sense that its re-application to other fields can follow a similar path. The method proposed is based on proper

*Correspondence to: A. G. Buchan, Applied Modelling & Computation Group, Department of Earth Science and Engineering, Imperial College London, London SW7 2AZ, UK.

†E-mail: andrew.buchan@imperial.ac.uk

orthogonal decomposition (POD) technologies [4,5], which are formed via the method of snapshots [6], and these have been adapted to specifically solve eigenvalue problems.

Proper orthogonal decomposition has evolved under a number of aliases, and so is also known as Karhunen–Loeve expansions in signal analysis and pattern recognition [7], principal component analysis in statistics [8], the method of empirical orthogonal functions in geophysical fluid dynamics [9] and meteorology [10]. All these methods, however, are model reduction techniques that offer adequate approximations of dynamical systems using a reduced number of degrees of freedom, that is, with lower dimensional models [11–13]. The fundamental mechanics of POD are to generate optimal basis functions that represent and capture the energy, or dynamics, of a system of interest, and a way of achieving this is through the method of snapshots. This involves taking snapshots of the system's state at various time instances [6], and from this data a set of POD basis functions are formed that provide an optimal representation. This means that the snapshot data can be reconstructed with the smallest error using a basis formed from a subset of these POD basis functions. Using only a small number of these optimal functions, the POD method then recasts the complex partial differential equations (describing the system of interest) into a set of much simpler ordinary differential equations. In the process, an efficient ROM is created that alleviates both the computational load and the memory requirements in comparison with the full model.

The origins of POD dates back some way to the early 1900s in the work of Pearson [14], but following the pioneering work of Lumley [15], it has received a considerable amount of attention from within the fluid dynamics community. Early applications include the work of Bakewell [16] and Payne [17] who respectively applied the techniques in turbulent pipe flow and wakes behind cylinders. In fact, it is POD's ability to capture the dynamics of eddies and turbulent flows that continually form and perish during a simulation that has proved highly beneficial. Its use soon extended to other applications including the modelling of flows around air foils [18] and through channels [19], the mixing of fluid layers [20], thermal currents [21,22] and ocean models [23]. It has been applied to the shallow water equations [24], the Euler equations [25], the full Navier–Stokes equations [26] and the various reduced versions of it, for example, the parabolized Navier–Stokes [27,28]. POD has been applied to many other fields that are covered extensively in the references of the articles listed here.

POD's application to such a range of fields demonstrates the need to form ROMs for complex physical phenomena. However, its application to solving eigenvalue problems is somewhat limited because the method of snapshots require a time-dependent and dynamically evolving system. A problem therefore immediately arises because the eigenvalue problem is time independent, but this issue is resolved by introducing within the equation a fictitious dimension that resembles time. By doing so the eigenvalue problem is transformed into a time-dependent problem and, in its limit as time tends to infinity, the new representation of the eigenvalue converges to the exact solution. When solving this time-dependent problem, the evolving solution can now be taken at various time instances to form the snapshots. The inspiration for this comes from the work in [29,30], which introduced a fictitious time variable to solve the stochastic eigenvalue problem. However, in the formulation presented in this article, a different time-dependent equation is formed. This was to overcome a singularity that was inherent within the previous work's derivation for non-fissile materials. Finally, it is also important to note that this approach has similarities with that of Du *et al.* [27], which solved the two-dimensional time-independent parabolized Navier–Stokes equation. In this previous work, snapshots were constructed by solving the problem using one of the spatial dimensions as though it were time.

To demonstrate the new POD model proposed in this article, the approach has been applied to RP applications. In this field, one is primarily concerned with determining the distribution of the free moving neutrons within a nuclear system [31–33]. Knowledge of their distribution and population growth yields information about a reactor's power distribution and determines whether the system is in a super, sub or critical state. In fact, their population growth is one of the most important characteristics of a nuclear system, and this is determined through the solving of an eigenvalue problem that, in the RP community, is referred to as K_{eff} [33].

Determining eigenvalues associated with RP involves solving the associated differential equation describing the transport of neutral particles, namely the Boltzmann transport equation [33]. This is

a particularly demanding equation as it requires the description of the particles' direction of travel to be known in addition to their energies (or spectrum). This additional information adds three dimensions to the standard space-time, which in turn greatly increases the complexity in obtaining its solution. Although some relief can be gained by recasting this equation into its diffusion form, which is the approach used here, whereby two of the additional dimensions are eliminated [31], the problem domains associated with RP are still complex and challenging to resolve. For example, whole reactor cores are typically several meters in size, but they are formed from structures with length scales many orders of magnitude smaller – such as in modern generation IV designs that contain fuel composed of pellets just a few millimeters in diameter. To resolve such domains using either a finite element (FE) or finite difference scheme would require enormous spatial meshes if the fine-scale structures were resolved. Techniques that resolve these domains efficiently have therefore always taken a central role in this area of research. These include technologies spanning adaptive resolution schemes [34–36], multiscale resolution, acceleration techniques [37] as well as the recasting of the equations into more efficient forms [31, 33, 38]. However, the vast majority of these technologies are always reliant on the solving of a full-scale model. That is, even when performed efficiently, one is typically dealing with a discretisation scheme that requires a very large number of unknowns to be resolved during the solving process. Reactor physics is therefore quite an unexplored area of POD application where, to the authors knowledge, the only other related work is that of Wols [39]. In this previous work, POD models were formed to simulate the dynamics of an accelerator-driven system in shutdown (a non-eigenvalue problem), the results of which were shown to be quite promising.

The remaining sections of this article are set out as follows. Section 2 presents the eigenvalue problem for reactor criticality together with the details of its time-dependent formulation, and its time and spatial discretisation. Section 3 presents the POD formulation for the ROM eigenvalue problem. In Section 4, two numerical examples are presented to demonstrate the capabilities of the method presented. Finally, Section 5 completes this article with a conclusion of its findings.

2. THE PSEUDO-TIME-DEPENDENT FORMULATION FOR THE EIGENVALUE PROBLEM

The focus of this article is on solving the eigenvalue problem expressed in form,

$$L\phi = \lambda Q\phi, \quad (1)$$

for which the terms L and Q denote differential operators, ϕ is the system's eigenfunction and λ is the associated eigenvalue. Although the POD model will be developed here on the eigenvalue problem associated with RP, its reapplication to other fields in the form of (1) should be fairly straightforward. That is, similar paths of derivation could follow that what is presented here because regular differential operators and boundary conditions, that occur in many applications, are involved.

In the eigenvalue problem for RP, one solves for the eigenfunction and eigenvector $\phi(v)$ and λ , respectively, where $\phi(v)$ denotes the density of the free moving neutrons' distribution over a spatial domain v . The eigenvalue problem reads as,

$$-\nabla D(v) \cdot \nabla \phi(v) + \Sigma_a(v)\phi(v) = \lambda \hat{\nu} \Sigma_f(v)\phi(v), \quad (2)$$

for which the diffusion coefficient is expressed as,

$$D(v) = \frac{1}{3\Sigma_t} = \frac{1}{3(\Sigma_a + \Sigma_s)}. \quad (3)$$

In Equations (2) and (3), the Σ terms denote the material cross sections, and these are probabilities that describe how the neutrons interact with the materials of the problem. The term Σ_f is the cross section for a neutron-causing fission, Σ_a is the cross section that a neutron is absorbed, and Σ_s is the cross section that a neutron will scatter in a different direction. The term $\hat{\nu}$ denotes the average number of neutrons released per fission event. The three terms in Equation (2) therefore describe

the transport of the neutrons, their removal by absorption and their production (or source) through fission, respectively.

The eigenvalue λ can be seen to multiply against the neutron source in Equation (2). Its value balances the terms that produce neutrons with those that account for their losses, and this denotes the reciprocal of the value k_{eff} ,

$$k_{eff} = \frac{1}{\lambda}. \quad (4)$$

The value of k_{eff} determines whether a system is in a supercritical ($k_{eff} > 1$), subcritical ($k_{eff} < 1$) or critical ($k_{eff} = 1$) state, that is, whether the neutron population continuously increases, decrease or remains constant, respectively.

The diffusion equation is supplemented with two main boundary conditions. One condition represents a surface on which neutrons are reflected back into the domain (reflective condition) and the other which represents surfaces that allow neutrons to escape out of the system (vacuum or bare condition). These are both satisfied by relating the flux solution to its gradient on the boundary by,

$$-\frac{1}{2}Dn \cdot \nabla \phi(v) = \begin{cases} \frac{1}{4}\phi(v) & \text{bare surface} \\ 0 & \text{reflective surface.} \end{cases}$$

where n denotes the outward pointing normal at the surface.

In standard nuclear engineering procedures, one normally employs an eigenvalue solver directly to Equation (2), such as the power method, to solve the eigenvalue problem. However, to formulate the reduced-order POD formulation a new pseudo time-dependent equation is used. This equation is formulated by introducing to the eigenvalue equation a fictitious independent variable that can be considered to represent time t . The eigenvalue λ is then considered to be a time-dependent and nonlinear variable, $\lambda(t, \phi)$, and a time-dependent eigenvalue problem can be formulated as,

$$\frac{\partial \phi(v, t)}{\partial t} - \nabla D(v) \cdot \nabla \phi(v, t) + \Sigma_a(v)\phi(v) = \lambda(t, \phi) \hat{v} \Sigma_f(v)\phi(v, t). \quad (5)$$

If one now selects an expression for $\lambda(t, \phi)$ such that in the limit of time, the condition,

$$\frac{\partial \phi(v, t)}{\partial t} \rightarrow 0, t \rightarrow \infty, \quad (6)$$

holds, then the steady state solution of Equation (5) will generate both the eigenvalue and corresponding eigenfunction. That is, as time $t \rightarrow \infty$ one obtains,

$$0 = -\nabla D(v) \cdot \nabla \phi(v, t) + \Sigma_a(v)\phi(v) - \lambda(t, \phi) \hat{v} \Sigma_f(v)\phi(v, t), \quad (7)$$

which is the solution to the original problem. To formulate an expression for $\lambda(t, \phi)$, it is assumed that it is a slowly varying function of time, which will certainly be the case when the solution attains its steady state. The time derivative in Equation (5) can therefore drop out of consideration, and the resulting equation is integrated over the spatial domain v . This gives the expression for the eigenvalue $\lambda(t, \phi)$ as,

$$\lambda(t, \phi) = \frac{\int_V -\nabla D(v) \cdot \nabla \phi(v, t) dv + \int_V \Sigma_a(v)\phi(v) dv}{\int_V \hat{v} \Sigma_f(v)\phi(v, t) dv}. \quad (8)$$

This expression can now be simplified by applying the divergence theorem to the integral involving the diffusion term,

$$\lambda(t, \phi) = \frac{\int_\Gamma -D(v)n \cdot \nabla \phi(v, t) d\Gamma + \int_V \Sigma_a(v)\phi(v) dv}{\int_V \hat{v} \Sigma_f(v)\phi(v, t) dv}, \quad (9)$$

and thus transforming it into a surface integral. By inspecting the expression within this integral to that of the boundary conditions specified in Equation (2), one sees that there are similar expressions

involving the flux gradient. These gradient terms can therefore be replaced with expression of the flux for bare surfaces, or simply a zero for reflecting boundaries. The resulting expression for the eigenvalue is therefore given as,

$$\lambda(t, \phi) = \begin{cases} \frac{\frac{1}{2} \int_{\Gamma} \phi(v,t) d\Gamma + \int_V \Sigma_a(v) \phi(v) dv}{\int_V \hat{v} \Sigma_f(v) \phi(v,t) dv} & \text{bare surface} \\ \frac{\int_V \Sigma_a(v) \phi(v) dv}{\int_V \hat{v} \Sigma_f(v) \phi(v,t) dv} & \text{reflective surface} \end{cases} \quad (10)$$

2.1. Time and spatial discretisation

The discretisation of the time and spatial dimensions start with a simple time-stepping treatment of the temporal derivative. This is given by $\frac{\partial \phi(v,t)}{\partial t} = (\phi^{k+1} - \phi^k) / \Delta t$ where the flux is expressed at the two time instances k and $k + 1$ that are separated by the time interval Δt . When applied to Equation (5), the resulting equation states,

$$\frac{1}{\Delta t} \phi^{k+1}(v) - \nabla D(v) \cdot \nabla \phi^{k+1}(v) + \Sigma_a(v) \phi^{k+1}(v) = \lambda^k \hat{v} \Sigma_f(v) \phi^k(v) + \frac{1}{\Delta t} \phi^k(v), \quad (11)$$

where the eigenvalue λ^k can be calculated from Equation (10) using the flux expressed at time instance k . From Equation (11), a FE discretisation of the spatial dimension is applied. This uses a basis of FE functions $N_i(v), i \in \{1, 2 \dots n\}$ to approximate the flux through the summation $\phi^k(v) = \sum_j \phi_j^k N_j(v)$. The coefficients of the expansion ϕ_j^k are obtained through a weighted residual method. Equation (11) is therefore multiplied by $N_i, \forall i \in \{1, 2 \dots n\}$, and integrated over space to give,

$$\begin{aligned} \int_V N_i \frac{1}{\Delta t} \phi^{k+1}(v) dv - \int_V N_i \nabla D(v) \cdot \nabla \phi^{k+1}(v) dv + \int_V N_i \Sigma_a(v) \phi^{k+1}(v) dv \\ = \lambda^k \int_V N_i \hat{v} \Sigma_f(v) \phi^k(v) dv + \int_V N_i \frac{1}{\Delta t} \phi^k(v) dv. \end{aligned} \quad (12)$$

One can now apply Green's theorem to the integral involving the diffusion operator to form a volume and surface integral, through which the appropriate boundary conditions can be applied. Finally, upon substituting the FE expansion into discretised equations, the following linear system results,

$$A \phi^{k+1} = \lambda^k M^{\Sigma_f} \phi^k + M \phi^k. \quad (13)$$

In this system ϕ^k denotes the vector of coefficients at time instance k . The term A denotes a matrix of dimension $n \times n$ and contains the discretised time, diffusion and absorption terms in the left-hand side of Equation (12). The terms M and M^{Σ_f} are also matrices of dimension $n \times n$. M is the standard mass matrix, $M_{i,j} = \int_V N_i(v) N_j(v) dv$ and M^{Σ_f} denotes the mass matrix that includes the fission cross-section within the integrals, that is, $M_{i,j}^{\Sigma_f} = \int_V N_i(v) \hat{v} \Sigma_f(v) N_j(v) dv$. Using this formulation, one may march through time, solving the flux for each time step and generating the updated eigenvalue. This process can be continued until steady state is attained at which point the k_{eff} eigenvalue of the problem will be found.

3. THE REDUCED-ORDER POD FORMULATION

In the reduced-order formulation, a new basis of POD functions are constructed from a collection of snapshots taken a various time instances of the full model solution. The model described in Equation (13) is solved for a set of different problems, and snapshots of their solutions are taken as they evolve to their steady state. In the FE framework, each snapshot is simply the value of the solution at the nodes of the FE mesh. Each snapshot is therefore a vector of size n , and the series of all these snapshots are collected together to form a matrix \mathcal{A} . The dimensions of \mathcal{A} is $n \times n_s$, where n_s is the total number of snapshots taken; typically, one finds that $n_s \ll n$. Once the full set of

snapshots have been collated, it is then custom to remove from each snapshot the mean value of all snapshots. That is, a modified snapshot matrix is generated by,

$$\mathcal{A}_{i,j} \rightarrow \mathcal{A}_{i,j} - \bar{\Phi}_i, \quad (14)$$

where $\bar{\Phi}_i = \frac{1}{n_s} \sum_j \mathcal{A}_{i,j}$ is the mean value of the snapshots at point i .

A POD basis set is now constructed through the singular value decomposition (SVD) of the matrix \mathcal{A} . That is, \mathcal{A} is represented by,

$$\mathcal{A} = U \Sigma V^t \quad (15)$$

where U and V are unitary matrices of dimension $n \times n$ and $n_s \times n_s$, respectively. Σ is a matrix of size $n \times n_s$ and has non zero values only along its diagonal elements. These non zero values are the singular values of \mathcal{A} , and these are assumed to be listed in order of their magnitude.

A POD basis set $\{\Phi_k\}, k \in \{1, 2, \dots, n_p\}$ consisting of n_p basis functions is constructed from a reduced form of the SVD in Equation (15). The first, and largest, n_p singular values in the matrix Σ are used to form a reduced matrix Σ^{n_p} ,

$$\Sigma_{i,i}^{n_p} = \begin{cases} \Sigma_{i,i} & , i \leq n_p \\ 0 & , i > n_p \end{cases}, \quad (16)$$

and this is used to form an approximation of \mathcal{A} in its SVD form,

$$\mathcal{A}_{n_p} = U \Sigma^{n_p} V^t. \quad (17)$$

It is shown in many of the texts referenced in this article that this representation is the closest possible to that of the matrix \mathcal{A} using n_p basis vectors. There is also a convenient error measure of the approximation which is given as the sum of all singular values removed from the original matrix Σ . Finally, by keeping in line with the conventional POD construction, the n_p POD basis functions are now defined to be the first n_p columns of the matrix U . The POD basis set is expressed in vector form by,

$$\Phi_j = U_{:,j}, \quad \text{for } j \in \{1, 2, \dots, n_p\}. \quad (18)$$

These vectors represent the POD basis functions over the spatial domain v , and these can now be constructed using the original FEs,

$$\Phi_j(v) = \sum_i U_{i,j} N_i(v), \quad \text{for } j \in \{1, 2, \dots, n_p\}. \quad (19)$$

3.1. Forming the reduced system

The reduced-order representation of the solution variables of Equation (13) takes the form of an expansion of the POD basis functions defined in Equation (18). This expansion is given for each time instance k by,

$$\phi^k = \bar{\Phi} + \sum_j^{n_p} \Phi^j \psi_j^k = \bar{\Phi} + \sum_j^{n_p} U_{:,j} \psi_j^k = \bar{\Phi} + U \psi^k, \quad (20)$$

where, as previously stated, $\bar{\Phi}$ is the vector of size n holding the average value of all snapshots. The terms $\psi_j, j \in \{1, 2, \dots, n_p\}$, denote the POD basis expansion coefficients and ψ denote these values in vector form. U is the $n \times n_p$ matrix containing the first n_p column vectors of U in Equation (15).

A reduced system of equations can now be constructed by inserting the expression of the full flux, Equation (20), into Equation (13), and then premultiplying the system by U^t ,

$$U^t A(\bar{\Phi} + U \psi^{k+1}) = \lambda^k U^t M^{\Sigma_f} (\bar{\Phi} + U \psi^k) + U^t M(\bar{\Phi} + U \psi^{k+1}). \quad (21)$$

This can be rearranged to,

$$U^t AU \psi^{k+1} = -U^t A \bar{\Phi} + \lambda^k U^t M^{\Sigma_f} \bar{\Phi} + U^t M \bar{\Phi} + \lambda^k U^t M^{\Sigma_f} U \psi^k + U^t MU \psi^k, \quad (22)$$

which expresses the POD coefficients at the next time step in terms of a reduced system of equations. In total, there are three reduced system matrices required in this formulation, namely $U^t AU$, $U^t M^{\Sigma_f} U$ and $U^t MU$. Each is of size $n_p \times n_p$ and, because being linear, can be precomputed and stored in memory ready for use when the reduced system is solved. There are also three matrix systems that multiply with the average snapshot vector $\bar{\Phi}$. Once again, these can be precomputed and fed into the solving procedure as a fixed source term.

The matrices in the reduced system are obtained by Galerkin projections (in the L_2 sense) of the original differential operators using the bases of POD functions. This can be demonstrated for the left-hand matrix of Equation (22) containing the discretised time, diffusion and absorption terms. If one considers these terms to be encapsulated within the operator $L = \frac{\partial}{\partial t} - \nabla D(v) \cdot \nabla + \Sigma_a(v)$, then when it operates on the solution, followed by its weighting with the POD function Φ_i and finally its integration over space, the following Galerkin projection is formed,

$$\int_v \Phi_i L \psi dv. \quad (23)$$

If the solution is now represented via the POD basis functions (where the average vector $\bar{\Phi}$ is not included in these terms), the Galerkin projection reads as,

$$\sum_j^{n_p} \int_v \Phi_i L \Phi_j dv \psi_j. \quad (24)$$

Replacing each POD function in the previous expression with its representation of FEs, that is, $\Phi_i = \sum_{k=1}^n U_{k,i} N_k = \sum_{k=1}^n (U^t)_{i,k} N_k$, the following expression results,

$$\sum_j^{n_p} \left\{ \int_v \sum_{k=1}^n (U^t)_{i,k} N_k L \sum_{l=1}^n U_{l,j} N_l dv \right\} \psi_j. \quad (25)$$

This can be rearranged to read as,

$$\sum_j^{n_p} \left\{ \sum_{k=1}^n \sum_{l=1}^n (U^t)_{i,k} \int_v N_k L N_l dv U_{l,j} \right\} \psi_j, \quad (26)$$

for which the term contained within the integral is simply the differential operator discretised through the original FEs, that is, it is equal to $A_{k,l}$ in Equation (13). The terms contained within the brackets are therefore $(U^t AU)_{i,j}$. Considering now the Galerkin weighting involving all POD basis functions $\Phi_i, i = \{1, 2, \dots, n_p\}$, the following systems results,

$$U^t AU \psi. \quad (27)$$

This, of course, is the system defined in Equation (22), and so this demonstrates its equivalence to a Galerkin projection using the POD basis functions. Similar arguments can be applied to all other terms in Equation (21) to demonstrate their equivalence to Galerkin projections.

3.2. Calculation of the eigenvalue from the reduced-order eigenvector

It now only remains to calculate the eigenvalue that corresponds to the reduced-order eigenfunction for a given time step k . This can be calculated using the formulation in Equation (10) and recasting each term in terms of the reduced-order approximation. This is demonstrated first for the volume integral involving the fission cross-section (where it suffice to consider only one volume integral), which is given as,

$$\int_V \hat{v} \Sigma_f(v) \phi(v, t) dv, \quad (28)$$

in the full approximation space. In this expression, the flux is replaced with its expansion of POD basis functions given in Equation (20),

$$\int_V \hat{v}_{\Sigma_f}(v) \phi(v, t) dv = \int_V \hat{v}_{\Sigma_f}(v) \left(\bar{\Phi} + \sum_j^{n_p} \Phi_j^k \psi_j^k \right) dv. \quad (29)$$

Each POD basis function is then represented by its FE expansion of Equation (18), and together with a little re-arranging, the following is obtained,

$$\begin{aligned} \int_V \hat{v}_{\Sigma_f}(v) \phi(v, t) dv &= \int_V \hat{v}_{\Sigma_f}(v) \left(\sum_i \bar{\Phi}_i N_i(v) + \sum_j^{n_p} \psi_j^k \sum_i U_{i,j} N_i(v) \right) dv \\ &= \int_V \hat{v}_{\Sigma_f}(v) \sum_i N_i(v) \left(\bar{\Phi}_i + \sum_j^{n_p} \psi_j^k U_{i,j} \right) dv \\ &= \sum_i \int_V \hat{v}_{\Sigma_f}(v) N_i(v) dv \bar{\Phi}_i + \sum_j^{n_p} \psi_j^k \left(\sum_i \int_V \hat{v}_{\Sigma_f}(v) N_i(v) dv U_{i,j} \right). \end{aligned} \quad (30)$$

In this final form, the first term has formed a scalar, and this is constant in value as it involves only the average snapshot vector $\bar{\Phi}$. The second term involves a summation over the POD coefficients, and these are multiplied by the integrals contained within the large brackets. Both the scalar and the n_p integrals can be precomputed prior to solving ROM due to them being independent of the POD solution. The integral expression can therefore be written as,

$$\int_V \hat{v}_{\Sigma_f}(v) \phi(v, t) dv = \bar{\Phi}^f + \sum_j^{n_p} \psi_j^k \Phi_j^f, \quad (31)$$

where the precomputed values $\bar{\Phi}^f$ and Φ_j^f are given by,

$$\bar{\Phi}^f = \sum_i \int_V \hat{v}_{\Sigma_f}(v) N_i(v) dv \bar{\Phi}_i \quad (32)$$

and

$$\Phi_j^f = \sum_i \int_V \hat{v}_{\Sigma_f}(v) N_i(v) dv U_{i,j}, \quad (33)$$

respectively. Similar expressions can also be formed for the other terms in Equation (10). That is, the integral involving the cross section Σ_a is expressed as,

$$\bar{\Phi}^a + \sum_j^{n_p} \psi_j^k \Phi_j^a, \quad (34)$$

given the precomputed quantities,

$$\bar{\Phi}^a = \sum_i \int_V \Sigma_a(v) N_i(v) dv \bar{\Phi}_i, \quad (35)$$

and

$$\Phi_j^a = \sum_i \int_V \Sigma_a(v) N_i(v) dv U_{i,j}. \quad (36)$$

In addition, the surface integral is computed by,

$$\bar{\Phi}^{bc} + \sum_j^{n_p} \psi_j^k \Phi_j^{bc}, \tag{37}$$

using the precomputed quantities,

$$\bar{\Phi}^{bc} = \sum_i \int_{\Gamma} N_i(v) d\Gamma \bar{\Phi}_i \tag{38}$$

and

$$\Phi_j^{bc} = \sum_i \int_{\Gamma} N_i(v) d\Gamma U_{i,j}. \tag{39}$$

These expressions can now be combined to form the reduced-order eigenvalue at a time step k ,

$$\lambda^k = \frac{\frac{1}{2} \left(\bar{\Phi}^{bc} + \sum_j^{n_p} \psi_j^k \Phi_j^{bc} \right) + \bar{\Phi}^a + \sum_j^{n_p} \psi_j^k \Phi_j^a}{\bar{\Phi}^f + \sum_j^{n_p} \psi_j^k \Phi_j^f}. \tag{40}$$

In this form, it is important to note that once the precomputed stage is complete, the computation of the eigenvalue is of order n_p , that is, the size of the POD expansion.

4. NUMERICAL EXAMPLES

This section presents two numerical examples to demonstrate the capabilities of the POD method in solving the eigenvalue problem. The first problem, presented in Section 4.1, is a one-dimensional example that is used to demonstrate the basic principals of the reduced-order method working as intended. In Section 4.2, the second example is presented that demonstrates the method’s capabilities when applied to more demanding two dimensional domains. Section 4.3 completes this demonstration with a review of the efficiency of the POD formulation.

4.1. The one-dimensional slab reactor

This example attempts to emulate a simple slab reactor that consists of a 10 cm domain bounded by bare surfaces, as shown in Figure 1. The reactor’s domain contains a homogeneous fissile material that is divided up by two small regions 2.5 cm from the each boundary. In these two regions, the material properties are varied, and in the examples presented their neutron absorbing properties have been increased. This is to mimic what would be the insertion of a control rod.

To generate the snapshots used in constructing the ROM, solutions from the full model were computed using three different control rod configurations. The three simulations were designed to model the reactor with: 1) both rods withdrawn; 2) only the left rod inserted; and 3) only the right




Figure 1. The problem domain of numerical example 1.

Table I. Listed here are the three sets of material cross-section used in generating the snapshots matrix together with the four sets of material cross-sections used to test the reduced-order solver.

Test case		$\Sigma_a(cm^{-1})$	$\Sigma_s(cm^{-1})$	$\hat{\nu}\Sigma_f(cm^{-1})$	K_{eff}
F1	Region 1	0.4	2.0	0.5	1.2126921450105688
	Region 2	0.4	2.0	0.5	
	Region 3	0.4	2.0	0.5	
F2	Region 1	0.4	2.0	0.5	1.1981929650253054
	Region 2	0.6	2.0	0.3	
	Region 3	0.4	2.0	0.5	
F3	Region 1	0.4	2.0	0.5	1.1981929650253043
	Region 2	0.4	2.0	0.5	
	Region 3	0.6	2.0	0.3	
R1	Region 1	0.4	2.0	0.5	1.2126921450105688
	Region 2	0.4	2.0	0.5	
	Region 3	0.4	2.0	0.5	
R2	Region 1	0.4	2.0	0.5	1.1747386315568560
	Region 2	0.6	2.0	0.3	
	Region 3	0.6	2.0	0.3	
R3	Region 1	0.4	2.0	0.7	1.1575393638624507
	Region 2	0.5	2.0	0.4	
	Region 3	0.5	2.0	0.4	
R4	Left half	0.4	2.0	0.7	1.5963896505339461
	Right half	0.6	2.0	0.5	

The resulting K_{eff} solutions using the full model have also been listed.




Figure 2. The three normalised scalar flux solutions used to generate the snapshots for the POD construction.

rod inserted; these problems have been labelled F1, F2 and F3, respectively. The material cross-sections used in these simulation, together with the computed eigenvalues, are listed in Table I. The scalar flux solutions of the three problems are shown in Figure 2. Each solution has been normalised and were obtained using a sufficiently resolved mesh consisting of 600 FEs. A constant time step of 0.5 s was used for the time-dependent eigenvalue solver, and each simulation was allowed to run for 750 s at which point steady-state was attained.

Initially, four snapshot matrices were created using the simulations generated from the full model. Each snapshot matrix was built by taking solutions from the full model at equal time intervals during each simulation. The time intervals were taken at every 2, 4, 8 and 16 time steps, and these snapshot




Figure 3. The singular values corresponding to the four snapshot matrices.

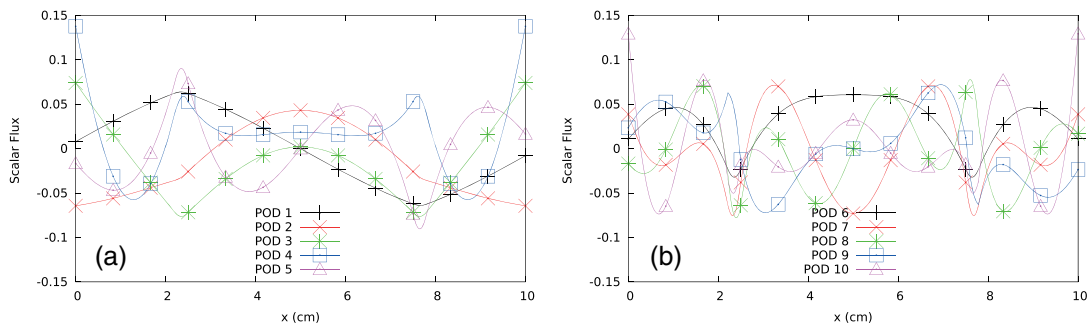


Figure 4. The POD basis functions generated from snapshot matrix s4. Graph a shows POD basis functions 1–5, and graph b shows POD basis functions 6–10.

matrices have been labeled s1, s2, s3 and s4, respectively. In total the four snapshots matrices consisted of 2235, 1128, 565 and 282 snapshots and from these 20 POD basis functions were constructed. The largest singular values for each snapshot matrix are presented in the graph of Figure 3. This shows them all to exhibit steep declines in their singular values which decrease by five orders of magnitude by the 12th–16th value. The first 10 POD basis functions generated from snapshot set s4 are presented in Figure 4. It is interesting to note that the low-order functions show much slower varying profiles in comparison with the higher order ones. This gives the indication that although the low-order POD basis functions capture the general profiles of a solution, the solution’s fine detail is filled in by the POD basis functions located further down the expansion. Thus the POD formulation should act much like a multiresolution approach.

The POD basis functions generated from the snapshots were used to form ROMs for four problems, labelled problems R1–R4. Problems R1–R3 have the same geometric structure to that used in the full model simulation, and their material properties are listed in Table I. Problem R1 is in fact the same as the full model problem F1, that is, the simulation with both control rods withdrawn. This was used to determine how well ROM reconstructs a problem that was used in constructing the original snapshot matrix. The remaining problems are however different from those used in building the snapshots. Problem R2 resembles the reactor with both control rods inserted. In problem R3, the control rods are configured to resemble a heavily inserted left control rod and slightly inserted right control rod. This is to generate a highly asymmetric flux profile that is quite different to those of the full problems F1–F3. Finally, problem R4’s geometry does not share the same structure as that of the slab reactor. Instead, the problem has been split in half, with one half being significantly more reactive than the other (the material cross-sections are listed in Table I). This is used as a stress test of ROM as the problem’s solution is quite different from those used in the POD construction. Table I list the reference eigenvalues for each of these problem setups, and these were generated using the full diffusion model.

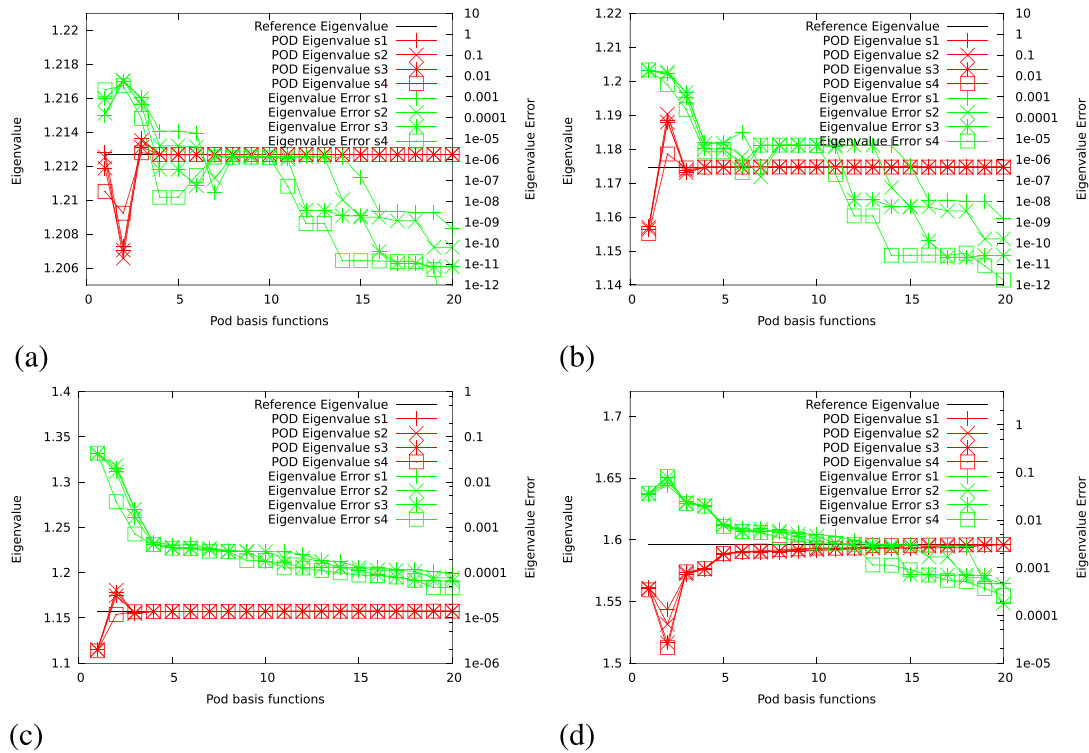


Figure 5. The graphs show the eigenvalues together with their errors when problems R1–R4 were solved using the reduced-order models with varying POD basis sizes. Graphs (a–d) correspond to the results for solving problems R1–R4, respectively.

The graphs presented in Figure 5 show the POD eigenvalues and respective errors for each test problem R1–R4. The results show the evolution of the eigenvalues as the POD basis sets are varied in size from 1 to 20 functions. They also show how each calculation is affected when the POD bases are generated from different snapshot sets s1–s4. The graph presenting the results for problem R1 (graph a) shows that a rapid reduction in the eigenvalue error is obtained using all POD basis sets. In fact, for all POD bases, just five POD basis functions were sufficient to produce an eigenvalue estimate with error less than 1×10^{-5} , which is well within an acceptable range for criticality calculations. The results do however exhibit an unexpected behavior over the range of 5–10 POD basis functions. In this range, the error in the eigenvalue increases with respect to an increased POD basis size. The reasons for this behaviour is not clear; however, the error does remain within an acceptable error margin. The model also recovers very well and mitigates the eigenvalue error to less than 1×10^{-9} using around 15 functions. This small anomaly is therefore considered not to be a cause for concern.

The graphs in Figure 5(b–d) show encouraging signs that the POD formulation is performing well at generating eigenvalue solutions to ‘unseen’ problems. Graph b shows that for the problem with both control rods inserted (R2), the eigenvalue is calculated to a high degree of accuracy. In fact, the error is shown to have been mitigated to less than 1×10^{-8} using just 12 POD basis functions. The more challenging problem, R3, has also been approximated accurately using the POD model. All POD bases formulated from snapshot sets s1–s4 are shown to perform equally well, and each has produced eigenvalue estimates with errors less than 1×10^{-4} when 20 POD basis functions are used. A similar result is seen in the final graph d, which shows the eigenvalue computed for the problem with a significant change in its geometry (R4). Again, the estimated eigenvalue tends towards the reference value with increased POD expansion size, and their errors approach 1×10^{-4} when the full 20 POD basis functions are used. When the differences in the geometric structure of problems R3 and R4 are taken into consideration, these results are highly encouraging as the eigenvalue errors are

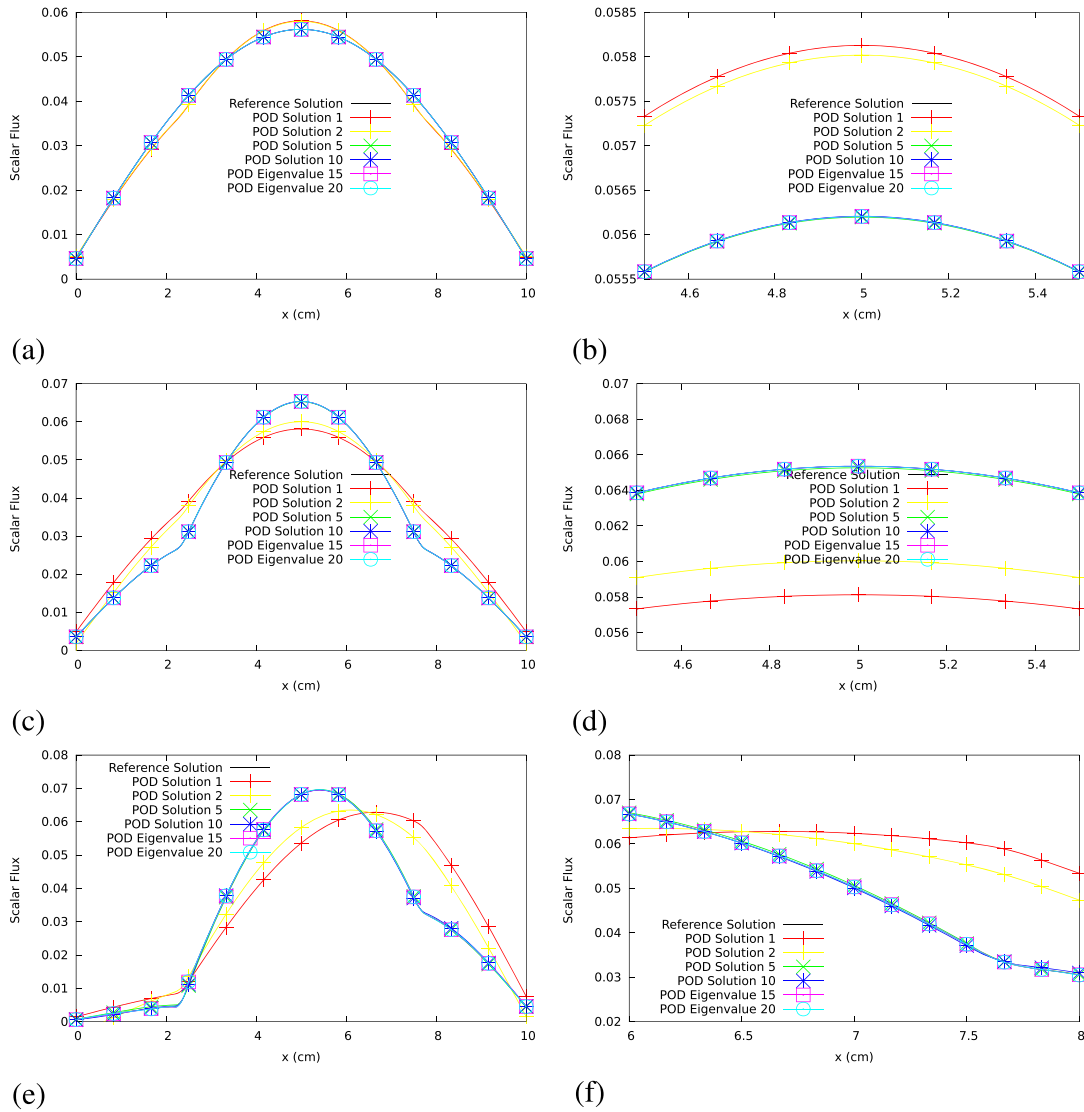


Figure 6. The reduced-order solutions of the normalised scalar flux eigenfunction for problems R1 (a and b), R2 (c and d) and R3 (e and f). Both the full solutions are shown together with their localised views that aim to highlight the regions with the largest errors. The numbers in the legend indicates the number of POD basis functions used in the reduced-order model.

within an acceptable bound. A final point these results highlight is that there is only a small change in POD’s performance when constructed through the different snapshot sets s1–s4. Therefore, the remaining results of this section will only use the POD sets generated from the smallest snapshot matrix s4.

The graphs presented in Figure 6 show the eigenfunction solution generated by POD for problems R1–R3. The results show the reduced-order solutions using 1, 2, 5, 10, 15 and 20 POD basis functions. It is shown that problem R1 can be reconstructed almost exactly using just five POD basis functions, and a similar observation is found when inspecting the graphs corresponding to problem R2. The POD formulation has performed very well at reconstructing the eigenfunction of problem R3. The asymmetric profile is already forming in the low-order POD expansions consisting of one and two functions. However, once again it seems that five POD basis functions are sufficient to capture the majority of the true solution’s profile. The corresponding solution profiles to problem R4 are presented in Figure 7. Again, the asymmetric solutions are already beginning to form in the

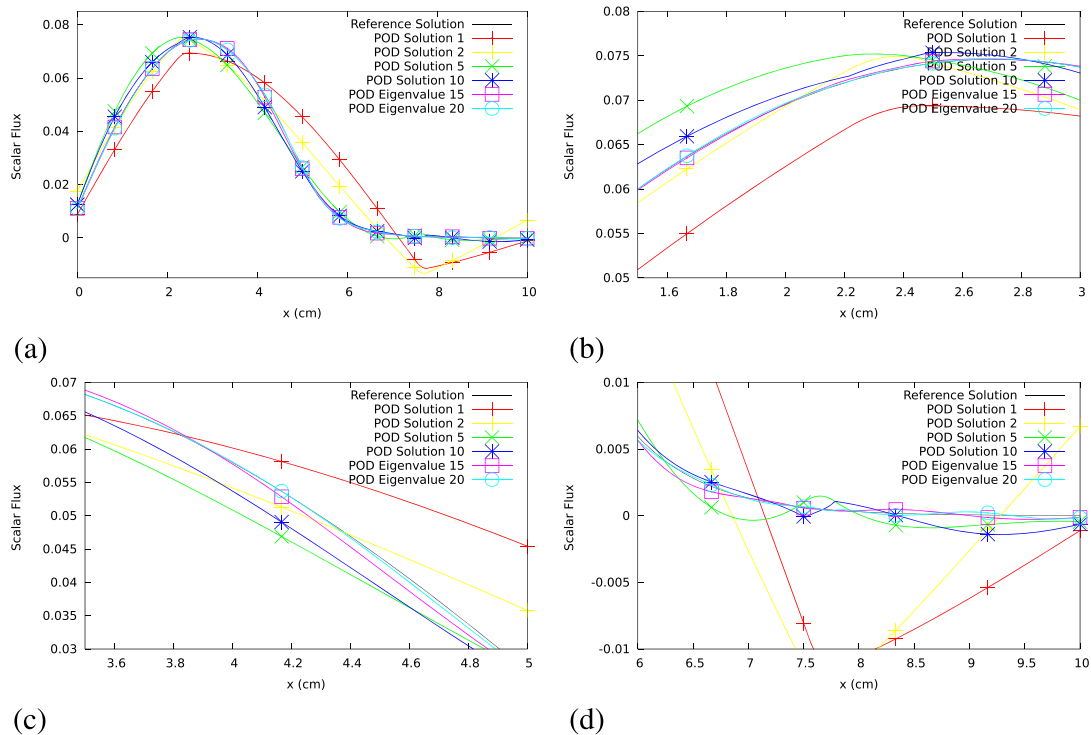


Figure 7. The reduced-order solution of the normalised scalar flux eigenfunction of problem R4. The full solution is shown (a) together with local views of the high-error regions (b–d). The numbers in the legend indicates the number of POD basis functions used in the reduced-order model.

low-order POD representations. In fact, using just one or two POD basis functions yield a basic outline of the reference solution profile. It is shown that the higher order POD expansions are capturing the profile of the reference solution with ever increasing accuracy as the POD basis set increase in size. The zoomed in graphs show that using the full 20 POD basis functions has yielded a very close representation of the true solution.

The errors of the POD eigenfunction solutions, with respect to the reference solutions, are shown in Figure 8. Again, these highlight the previous findings and illustrate the strong convergence rates of the POD solutions for each test problem.

4.2. A two-dimensional reactor

In this example, a two-dimensional mock reactor was modelled and its eigenvalue estimated using the POD formulation described in this article. The domain of the reactor is presented in Figure 9 where it is shown to have a square geometry with sides that are of length 90 cm. A fuel region occupied the main bulk of the reactor, but within its central region, there were four locations where control rods could be inserted and withdrawn. In the case where the rods were inserted a set of absorbing cross sections that mimicked a control rod were used. However when the rods were withdrawn, these regions were replaced with a set of cross sections resembling water (which is often the case in many reactor designs). Surrounding the reactor was a neutron reflector, that is, a high neutron scatterer, which was meant to resemble a graphite material that acted as a neutron shield. The cross-section data for each of the materials are listed in Table II. These have been selected to resemble the properties of fuel, water and graphite, and so have been based on the cross sections used in International Atomic Energy Agency (IAEA) benchmarks [40]. Included in this data are two sets of material cross-sections for the fuel. Fuel type 1 defines the standard fuel that was used in all problems during the POD construction phase. Fuel type 2 defines an adjusted fuel make-up, and this was used in some of the subsequent reduced-order modelling analysis.

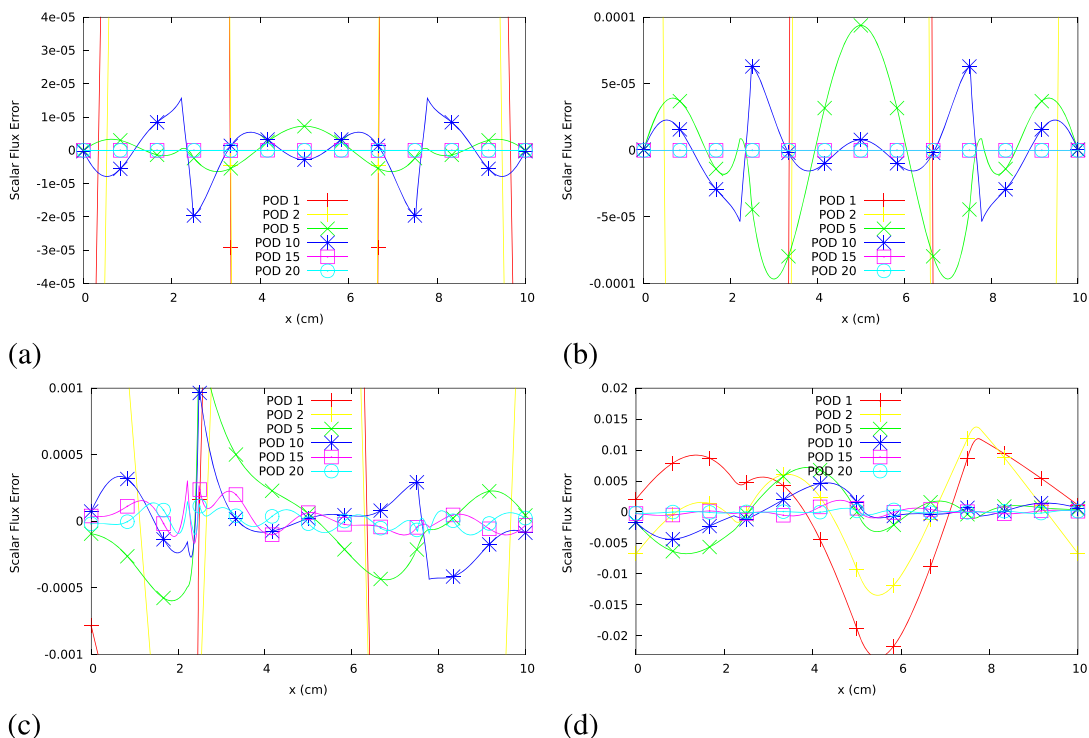


Figure 8. These graphs show the errors in the eigenfunction for varying POD basis expansion sizes (POD sizes are indicate by the numbers in the legends). The results in graphs (a-d) correspond to the reduced-order solutions of problems R1–R4, respectively. The numbers in the legend indicate the number of POD basis functions used in the reduced-order model.




Figure 9. The domain of numerical example 2.

To construct the POD basis functions, five full models (labelled F1–F5) of the reactor were performed. These resolved the reactor first with all rods removed, and then with one single rod inserted at a time. The materials for the fuel and control rods for each calculation is presented in Table III, and this table also lists the corresponding eigenvalues. Each calculation was run through the full model using a regular mesh of 32 400 quadrilateral FEs. A time step of 4.0 s was used, and 4000 time steps were performed to ensure the eigenvalues were sufficiently converged. A snapshot

Table II. The cross section data for the various materials used in the two-dimensional numerical example.

Material	Σ_a (cm^{-1})	Σ_s (cm^{-1})	$\hat{\nu}\Sigma_f$ (cm^{-1})
Fuel 1	0.075	0.53	0.079
Fuel 2	0.072	0.53	0.085
Water	0.01	0.89	0.0
Control rod	0.38	0.2	0.0
Graphite	0.15	0.5	0.0

Table III. The reactor's problem material configurations used in the generation of the POD basis functions (F1–F5) as well as the reduced-order models (R1–R7).

Problem	Fuel type	C-R1	C-R2	C-R3	C-R4	K_{eff}
F1	Fuel 1	W	W	W	W	1.0117683356042249
F2	Fuel 1	CR	W	W	W	1.0051523782732303
F3	Fuel 1	W	CR	W	W	0.99997463450337276
F4	Fuel 1	W	W	CR	W	1.0078716937581611
F5	Fuel 1	W	W	W	CR	1.0051523782732201
R1	Fuel 1	W	W	W	W	1.0117683356042249
R2	Fuel 1	CR	W	W	W	1.0051523782732303
R3	Fuel 1	W	CR	CR	W	0.99201388156849468
R4	Fuel 1	W	W	CR	CR	1.0017067353596756
R5	Fuel 1	CR	CR	W	CR	0.97787588014093285
R6	Fuel 1	CR	CR	CR	CR	0.9666724139312
R7	Fuel 2	$\frac{1}{2}CR + \frac{1}{2}W$	$\frac{1}{10}CR + \frac{9}{10}W$	$\frac{3}{10}CR + \frac{7}{10}W$	$\frac{1}{5}CR + \frac{4}{5}W$	1.1017019903545371

The control rod regions are denoted by C-R; CR denotes a control rod material; W denotes water; and F denotes fuel type. Reference eigenvalues for both the POD generation and the reduced model problems are listed.

set was then constructed by taking the solution at each 250th time instance. This resulted in 16 snapshots for each of the full simulations, thus 80 snapshots were recorded in total. From these snapshots, a basis set consisting of 12 POD functions were generated, and these are illustrated in Figure 10. The corresponding first 12 singular values of the snapshot matrix are presented in Figure 11. It is shown that their values fall by nearly five orders of magnitude by the 12th singular value.

By using ROM, seven configurations of the reactor were resolved, and these have been denoted problems R1–R7. Problems R1 and R2 were identical to the full model problems F1 and F2, and these were used to examine how well the POD model reproduces solutions for already seen problems. Problems R3–R6 used different combinations of inserted and rejected control rods, and these configurations are listed in Table III. In problem R7, the material configuration was varied by some considerable margin to those used in the POD generation. In the control rod regions, a mix of control rod and water materials were used to represent partially inserted control rods. This was achieved by taking linear combinations of the water and control rods cross-sections, the details of which are listed in Table III. In addition, to ensure that there were significant differences to the full model problems, the fuel composition of problem R7 was also changed to that of fuel type 2. Table III lists the reference eigenvalues for each of the seven test problems.

The POD predicted eigenvalues and corresponding errors for each of the seven test problems are presented in the graphs of Figure 12. The results for problems R1 and R2 show that the reconstruction of the problems used in constructing the POD basis functions can be performed accurately and efficiently. The eigenvalue's error in both these problems fall below 1×10^{-7} using only five POD functions, and the errors continue to reduce as the expansion size is increased. The problems involving varying combinations of control rod configurations (R3–R6) have also been predicted accurately. It is shown that the eigenvalue errors reduce monotonically with POD expansion size and that they can be reduced to approximately 1×10^{-4} to 1×10^{-5} using the full set of 12 functions. Results for problem R7 have also shown very encouraging signs that ROM can reconstruct problems that




Figure 10. The 12 POD basis functions generated in the two-dimensional numerical example. The POD basis functions 1–12 are ordered from left to right, top to bottom.




Figure 11. The first 12 singular values of the snapshot matrix.

vary considerably to those used in the POD’s construction. For this problem the eigenvalue is reconstructed to within an error of 0.1% using the 12 POD basis functions, and this is certainly within an acceptable margin for scoping calculations.

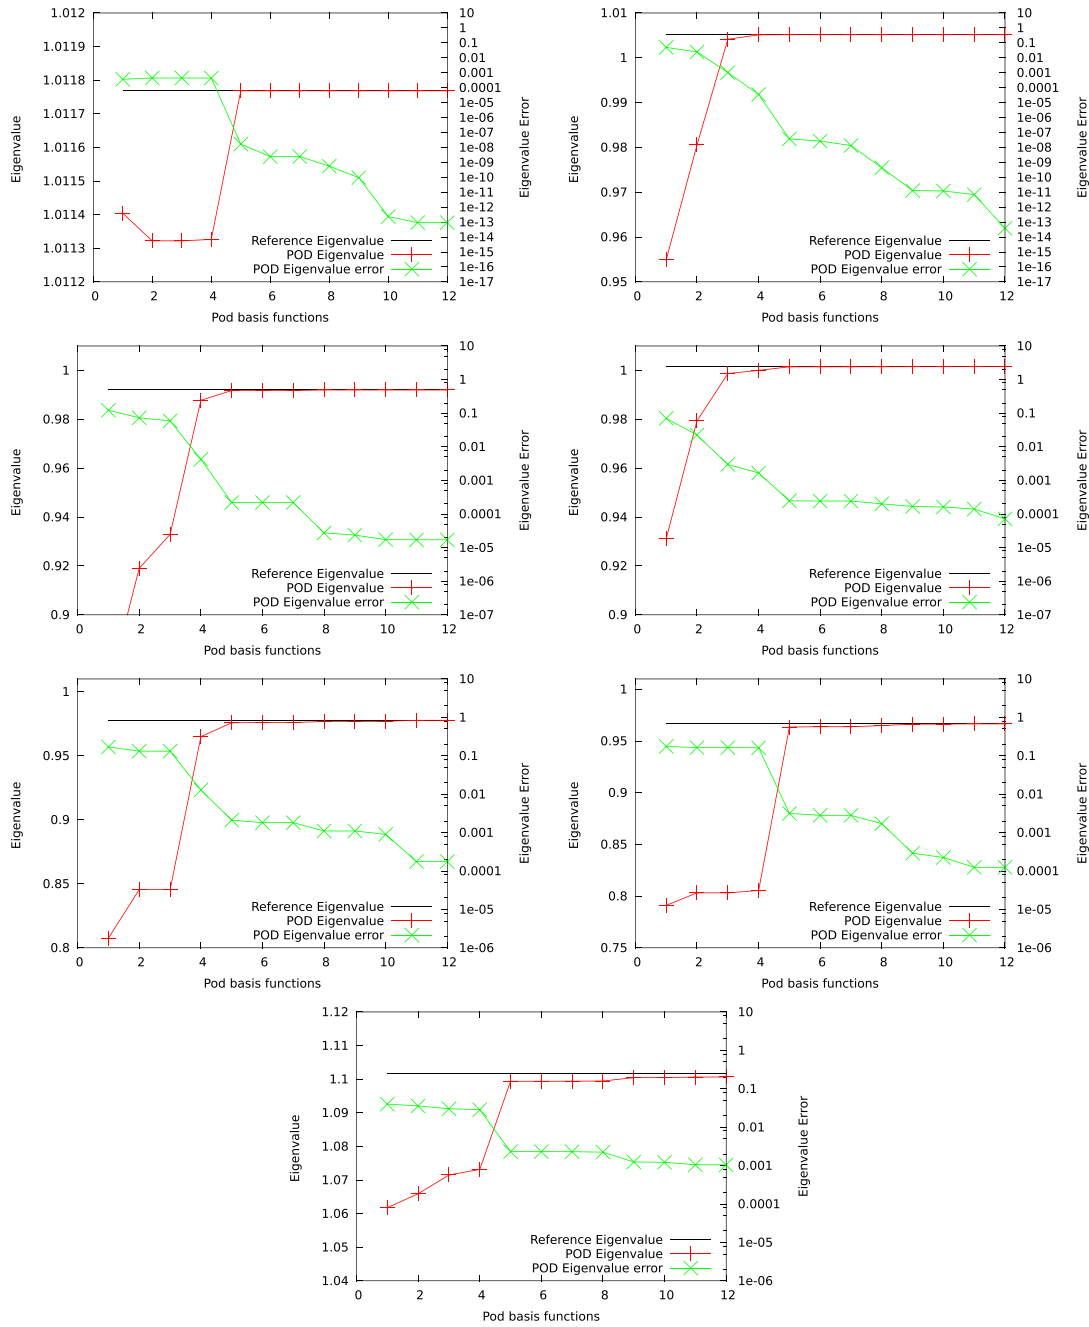


Figure 12. The eigenvalues and corresponding errors for each of the seven numerical examples. Left to right, top to bottom are the results for problems R1-R7.

The reduced-order fluxes for problems R5 and R7 are presented in Figures 13 and 14, respectively. By comparing against the reference solution included in the figures, the POD models are shown to capture the fine details of the solution quite quickly. For problem R5, the POD expansion using six functions has captured the general details of the full solution, and there is little visible difference between the reference solution and that of POD when 12 basis functions are used. Similarly, the POD representation of problem R7 compares closely to the reference solution using just six functions. Again, the visible differences between the solution of the POD model, using 12 functions, and




Figure 13. The reference solution of problem R5 together with its POD representation using varying expansion sizes. Left to right, top to bottom are the solution profiles using 1, 2, 4, 6 and 12 POD basis functions. The bottom right profile is of the exact solution.

the reference solution is minimal. One can quite clearly see that POD has captured the detail of the neutron distribution within the bulk of the reactor as well as in and around the control rod regions.

4.3. Review of the efficiency and performance of POD

This section presents an overview of the computational costs of ROM and compares the efficiency of the method in comparison with other traditional approaches. This begins with a statement on the reformulation of the original eigenvalue equation whereby the full-scale problem was solved in its time-dependent form. It was found that solving this equation was not necessarily as efficient as solving the eigenvalue problem through a more orthodox approach, such as the power method. However, the increase in computational times were not significant, and because only a small number of full-scale calculations are typically needed, these additional costs are not considered an issue.

Solving the eigenvalue problem through ROM was found to provide efficiency gains in comparison with using a full-scale model. In the second numerical example used in this article, the computational times of both the set-up and solving stage were recorded. In the set-up stage, whereby both the SVD of the snapshot matrix and the matrices in Equation (22) were computed, the computational times varied between 1.2×10^{-2} and 0.14 s, depending on the size of the ROM. Following its construction, the times required to solve ROM remained below 8.0×10^{-3} s for all expansion sizes. These times are compared against those of a full model solution, which in total required 10.35 s. This time was derived from an established RP code [41] using the same FE discretisation used in the POD construction. In addition, this full model used the more orthodox power method in obtaining its solution, which is often the choice of solver in RP calculations. The time comparison therefore provides a true reflection of the efficiency gains that can be made by applying this reduced-order method. That is, in comparison with conventional power method approaches, computational costs have been reduced by 87–99% in this example.




Figure 14. The reference solution of problem R7 together with its POD representation using varying expansion sizes. Left to right, top to bottom are the solution profiles using 1, 2, 4, 6 and 12 POD basis functions. The bottom right profile is of the exact solution.

5. CONCLUSION

This article has presented a new reduced-order modelling technique for solving general eigenvalue problems. The method is arbitrary in that it can be extended to other fields and applications, but here its use has been demonstrated in the calculation of the eigenvalue (K_{eff}) for reactor criticality problems. The ROM was based on POD technologies whereby the basis functions of the new model were generated from the SVD of a snapshot matrix. To generate the snapshot matrix, an unorthodox approach was used in solving the full-model eigenvalue problem. Here, a time-dependent form of the eigenvalue problem was developed and resolved. Although this pseudo-time approach has been used in previous work (unrelated to ROMs) cited in this article, the equation used in this work is of a different form. Using a time-dependent equation enabled one to create a snapshot matrix by taking the solutions at various time instances of the full model as it evolved to its steady state. From here on, one could then follow the traditional POD route in generating the reduced basis functions. In the formulation presented, this was carried out within a framework where the full model was resolved using FEs.

To demonstrate the capabilities of the reduced order model, two numerical examples have been presented. The first was a one-dimensional example, and this was used to demonstrate the general properties of the reduced model to be working. The second example used a mock two-dimensional reactor with a more demanding geometry that is typically seen in RP applications. A number of important points were highlighted in the numerical examples. One was that the POD models could efficiently reproduce the eigenvalues for the problems used in generating the POD basis functions, as one would expect. However, when the materials properties were varied to pose new, or unseen, problems, the POD formulation could still generate accurate estimates of the eigenvalues. To add further encouragement, the POD bases performed very well in reproducing problems that had significant changes in their geometry in comparison with the problems used in the POD generation.

This was observed in both numerical examples, for which the eigenvalue estimates were still within an acceptable margin of error. In addition to this, the eigenfunctions, of fluxes, were also reproduced to a high degree of accuracy. Even in the more demanding problems, the fluxes were shown to capture the features of the reference solutions. However, there were a few issues raised in the examples shown. In addition to method using the more expensive time-dependent full-model formulation, the POD eigenvalues occasionally showed rising errors with increased expansion size. However, this issue appears to be rare and so should not distract from the main results showing the method to be working well. In summary, the POD method is shown to have potential for solving the eigenvalue problem, particularly for RP applications. The test cases used were based on previous benchmark problems and so similar POD performance would be expected for similar problem types. How well the method works for different problems will be the focus of future work.

ACKNOWLEDGEMENTS

We acknowledge the funding of this work through the EPSRC grant ref: EP/J002011/1. Prof. I.M. Navon acknowledges the support of NSF grant ATM-0931198.

REFERENCES

1. Antoulas AC. *Approximation of Large-Scale Dynamical Systems (Advances in Design and Control)*. SIAM: Philadelphia, 2005.
2. Schilders WHA, van der Vorst HA, Rommes J. *Model Order Reduction: Theory, Research Aspects and Applications*. Springer: Berlin, 2008.
3. Kirby M. *Geometric Data Analysis: An Empirical Approach to Dimensionality Reduction and the Study of Patterns*. John Wiley & Sons, Inc.: New York, USA, 2001.
4. Buljak V. *Inverse Analyses with Model Reduction: Proper Orthogonal Decomposition in Structural Mechanics (Computational Fluid and Solid Mechanics)*. Springer: Berlin, 2011.
5. Holmes P, Lumley JL, Berkooz G, Rowley CW. *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, 2nd ed. Cambridge University Press: Cambridge, 2012.
6. Sirovich L. Turbulence and the dynamics of coherent structures, part III: dynamics and scaling. *Quarterly of Applied Mathematics* 1987; **XLV**:583–590.
7. Fukunaga K. *Introduction to Statistical Pattern Recognition*, 2nd ed., Computer Science and Scientific Computing. Academic Press: San Diego, 1990.
8. Jolliffe IT. *Principal Component Analysis*, 2nd ed. Springer: Berlin, 2002.
9. Crommelin DT, Majda AJ. Strategies for model reduction: comparing different optimal bases. *Journal of the Atmospheric Sciences* 2004; **61**:2206–2217.
10. Majda AJ, Timofeyev I, Vanden-Eijnden E. Systematic strategies for stochastic mode reduction in climate. *Journal of the Atmospheric Sciences* 2003; **60**(14):1705–1722.
11. Du J, Zhu J, Navon IM. An optimizing finite difference scheme based on proper orthogonal decomposition for CVD equations. *International Journal for Numerical Methods in Biomedical Engineering* 2011; **27**:78–94.
12. Cao Y, Navon IM, Luo Z. A reduced-order approach to four-dimensional variational data assimilation using proper orthogonal decomposition. *International Journal for Numerical Methods in Fluids* 2007; **53**:1571–1583.
13. Cao Y, Zhu J, Luo Z, Navon IM. Reduced order modeling of the upper tropical pacific ocean model using proper orthogonal decomposition. *Computers & Mathematics with Applications* 2006; **52**:1373–1386.
14. Pearson K. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine* 1901; **2**(6):559–572.
15. Lumley JL. The structure of inhomogeneous turbulent flows. In *Atmospheric Turbulence and Radio Propagation*. Nauka: Moscow, 1967; 166–178.
16. Bakewell HP, Lumley JL. Viscous sublayer and adjacent wall region in turbulent pipe flow. *The Physics of Fluids* 1967; **10**:1880–1889.
17. Payne FR, Lumley JL. Large-eddy structure of the turbulent wake behind a circular cylinder. *The Physics of Fluids* 1967; **10**:194–196.
18. Dean AE, Mavriplis C. Low-dimensional description of the dynamics in separated flow past thick airfoil. *The American Institute of Aeronautics and Astronautics* 1994; **32**:1222–1227.
19. Ball KS, Sirovich L, Keefe LR. Dynamical eigenfunction decomposition of turbulent channel flow. *International Journal for Numerical Methods in Fluids* 1991; **12**(6):585–604.
20. Delville J, Bellin S, Bonnet JP. Use of the proper orthogonal decomposition in a plane turbulent mixing layer. In *Turbulence and Coherent Structures; Selected Papers from Turbulence 89*. Kluwer Academic Publishers: Dordrecht, 1991; 75–90.
21. Sirovich L, Park H. Turbulent thermal convection in a finite domain: part I. Theory. *Physics of Fluids* 1990; **2**:1649–1658.

22. Park H, Sirovich L. Turbulent thermal convection in a finite domain: partII. Numerical results. *Physics of Fluids* 1990; **2**:1659–1668.
23. Fang F, Pain CC, Navon IM, Gorman GJ, Piggott MD, Allison PA, Farrell PE, Goddard AJH. A POD reduced order unstructured mesh ocean modelling method for moderate reynolds number flows. *Ocean Modelling* 2009; **28**:127–136.
24. Stefanescu R, Navon IM. POD/DEIM Nonlinear model order reduction of an adi implicit shallow water equations model. *Journal of Computational Physics* 2013; **237**:95–114.
25. Pettit CL, Beran PS. Application of proper orthogonal decomposition to the discrete euler equations. *International Journal for Numerical Methods in Engineering* 2002; **55**(4):479–497.
26. Burkardt J, Gunzburger M, Lee H. POD and CVT-based reduced-order modeling of Navier–Stokes flows. *Computer Methods in Applied Mechanics and Engineering* 2006; **196**:337–355.
27. Du J, Navon IM, Steward JL, Alekseev AK, Luo Z. Reduced-order modeling based on POD of a parabolized Navier–Stokes equation model I: forward model. *International Journal for Numerical Methods in Fluids* 2012; **69**(3):710–730.
28. Du J, Navon IM, Alekseev AK, Luo Z. Reduced-order modeling based on POD of a parabolized Navier–Stokes equation model II: trust region POD 4-D VAR data assimilation. *Computers & Mathematics with Applications* 2013; **65**:380–394.
29. Williams MMR. A method for solving a stochastic eigenvalue problem applied to criticality. *Annals of Nuclear Energy* 2010; **37**(6):894–897.
30. Williams MMR. A method for solving stochastic eigenvalue problems. *Applied Mathematics and Computation* 2010; **215**(11):3906–3928.
31. Duderstadt JJ, Hamilton LJ. *Nuclear Reactor Analysis*. John Wiley and Sons: NY, 1976.
32. Duderstadt JJ, Martin WR. *Transport Theory*. John Wiley and Sons: NY, 1979.
33. Lewis EE, Miller WF. *Computational Methods of Neutron Transport*. American Nuclear Society: Illinois, 1993.
34. Lathouwers D. Spatially adaptive eigenvalue estimation for the sn equations on unstructured triangular meshes. *Annals of Nuclear Energy* 2011; **38**(9):1867–1876.
35. Park H, de Oliveira CRE. Coupled space-angle adaptivity for radiation transport calculations. *Nuclear Science Engineering* 2009; **161**:216–234.
36. Buchan AG, Pain CC, Eaton MD, Smedley-Stevenson RP, Goddard AJH. Self adaptive spherical wavelets for angular discretisation of the Boltzmann transport equation. *Nuclear Science and Engineering* 2006; **158**:244–263.
37. Reed WH. The effectiveness of acceleration techniques for iterative methods in transport theory. *Nuclear Science and Engineering* 1971; **45**:245–254.
38. Ackroyd RT. *Finite Element Methods for Particle Transport: Applications to Reactor and Radiation Physics*. Research Studies Press: Taunton, Somerset, England, 1997.
39. Wols F. Transient analyses of accelerator driven systems using modal expansion techniques. *Master’s Thesis*, Delft University of Technology, 2010.
40. Stepanek J, Auerbach T, Hälgl W, Eidgenössisches Institut für Reaktorforschung Würenlingen. Calculation of four thermal reactor benchmark problems in x-y-geometry: preprint of a paper to be presented at the Seminar/ Workshop on “Thermal Reactor Benchmark Calculations, Techniques, Results and Applications”, May 17–18, 1982, Brookhaven, New York, 1982. EIR-Bericht, EIR.
41. de Oliveira CRE. An arbitrary geometry finite element method for multigroup neutron transport with anisotropic scattering. *Progress in Nuclear Energy* 1986; **18**:227–236.