Parametric Model-Order Reduction for Radiation Transport Using Multi-Resolution Proper Orthogonal Decomposition

Patrick Behne[a], Jean C. Ragusa[a,*]

[a]*Department of Nuclear Engineering, Texas A&M University, College Station, TX 77843*

Abstract

For parametric high-fidelity simulations, it is often desirable to utilize a reduced-order model (ROM) to emulate, at a reduced computational cost, parametric solutions of the governing partial differential equations (PDEs) for unseen parameter values. One commonly employed option is to utilize a data-driven, projection-based ROM supplemented with subspace identification via proper orthogonal decomposition (POD). POD discovers the ROM subspace by computing the singular value decomposition (SVD) of a set of training data from the full-order model (FOM). In streaming-dominated radiation transport simulations with localized sources, solutions often greatly vary over the spatial domain by many orders of magnitude. In such cases, machine-precision arithmetic can be insufficient to obtain an accurate SVD, resulting in a poorly performing ROM. We present a method called multiresolution POD (mrPOD) that mitigates these inaccuracies. The mrPOD method works by decomposing the spatial domain into regions and performing proper orthogonal decomposition on the training dataset separately in each region. mrPOD is tested on single energy group and multigroup atmospheric shielding transport problems and is shown to outperform classic POD.

Keywords: Model-Order Reduction; Radiation Transport; Local Proper Orthogonal Decomposition; Multi-Resolution POD; Parametric Surrogate Model; Reduced-Order Model

1. Introduction

Numerically solving the neutron transport equation is computationally expensive because of the high-dimensionality of the phase-space (6 dimensions in total for steady-state problems: 3 for space, 2 for angle, and 1 for energy). Moreover, in parametric or multi-query applications, such as uncertainty quantification and design optimization processes, the transport equation is also dependent upon these varying parameters and numerous full transport solutions may be required to assess the impact of such parameters or to optimize a design, further taxing a limited computational budget. Such a model is said to be parametric because input parameters (e.g., cross-section values and source strengths) can be uncertain or only known to some tolerance. In such situations, it is highly advantageous to develop parametric reduced-order models (ROMs). Generally speaking, model-order reduction (MOR) is a technique aimed at reducing the computational complexity of

*Corresponding author
Email addresses: pbehne@tamu.edu (Patrick Behne), jean.ragusa@tamu.edu (Jean C. Ragusa)

mathematical models in numerical simulations. For a thorough overview of MOR, including the use of proper orthogonal decomposition (POD) to generate data-driven, projection-based ROMs, the reader is referred to these books [1–4]. This present work builds upon previous work that applied MOR for radiation transport [5–9]. In such techniques, the full-order model (FOM) (here, radiation transport) is exercised to "discover" an adequate subspace for the ROM by generating solution snapshots obtained from random realizations of the uncertain or unknown input parameters. The role of POD is to identify the ROM subspace through a singular value decomposition (SVD) of the snapshot data.

In the subset of transport problems of interest here, we deal with radiation transport solutions that vary by many orders of magnitude over the spatial domain. In such cases, the SVD of the snapshot matrix is often inaccurate due to the limitations of (finite) machine-precision arithmetic. Environments that give rise to widely varying transport solutions include scenarios such as: radiation transport with localized sources in weakly scattering media, shielding applications, and other deep-penetration problems. The present work introduces a spatial domain decomposition-like method applied to POD so as to mitigate data loss that occurs from processing widely varying data. We refer to the method as multiresolution proper orthogonal decomposition (mrPOD). The essence of mrPOD involves splitting the spatial domain into regions in which the transport solutions vary less than over the domain as a whole. SVD is performed separately on each region, yielding region-wise POD bases that are more accurate than a single SVD over the entire domain. In this sense, while POD yields spatially global basis vectors in which to build the ROM solution, mrPOD yields more local basis vectors. mrPOD draws its name from the analogy between standard dynamic mode decomposition and multiresolution dynamic mode decomposition (mrDMD) [10]. mrDMD applies a multiresolution analysis to the DMD model-order reduction method to separate high and low frequency dynamics of multiscale transient systems. Similarly, mrPOD applies domain decomposition to extract a system's dominant features in a spatially local manner. Multi-scale POD, a version of POD that incorporates multiresolution analysis, has also been investigated to reduced fluid flow equations [11].

The idea of using spatial domain decomposition (DD) with ROMs is not new. See, for instance, the application of spatial DD to POD for the incompressible Navier-Stokes equations [12]. In such an approach, POD is performed on portions of the spatial domain to obtain spatially local bases in which the ROM solution is built. However, the concept of DD may be extended to all dimensions of a problem's phase-space and even further to include the parameter space. Partitioning the parameter space has been shown an effective method to increase the accuracy of POD ROMs on nonlinear FOMs. This technique splits the parameter space into regions and applies POD to the training data in each region to obtain several parametrically local ROMs. The idea behind this is that a nonlinear problem is better approximated using parametrically local subspaces instead of a global one. Domain decomposition on POD via parameter space partitioning has found applications to many nonlinear PDE FOMs [13–17]. Another technique for reducing nonlinear FOMs with a linear ROM uses the same logic, but decomposes the time domain instead of the parameter space. With time partitioning, the timeline is split into segments determined by the time evolution of the solution.

This approach has been applied to advection equations [18], nonlinear molecular dynamics simulations [19], and nonlinear parabolic PDEs in general [20].

Domain decomposition ROMs have also permeated the nuclear engineering community. Parameter partitioned POD has been applied to the multigroup k-eigenvalue diffusion equations to model the effect of control rod position [21]. Energy partitioned POD has been utilized in multigroup k-eigenvalue diffusion theory to account for the different behavior of the solution in different energy groups [21, 22]. Angle partitioned POD has found use in transport to increase ROM accuracy and solver stability [23]. Finally, space partitioned POD has been applied to multigroup k-eigenvalue diffusion in nuclear reactor cores [24, 25]. This approach allowed for the generation of training data in one spatial region at a time for different boundary conditions. The POD ROM for the full domain was built by matching region interface boundary conditions.

As a motivating example for the mrPOD method developed here: consider a point isotropic neutron source in a homogeneous, purely absorbing medium. The analytic solution to this radiation transport problem is, in spherical coordinates,

$$\phi(r) = \frac{q}{4\pi r^2} e^{-\Sigma_t r}\,, \tag{1}$$

where ϕ is the scalar flux (particles cm^{-2} s^{-1}), r is the distance from the source (cm), q is the source strength (particles s^{-1}), and Σ_t is the total interaction cross section (cm^{-1}) of the medium. Assume that the solution in Eq. (1) is of interest for distances within 3 km of the source. Furthermore, assume that the total cross section is not known exactly, but instead known to be in the range $(1.0 \times 10^{-5}, 12 \times 10^{-5})$ cm^{-1}. Figure 1a shows analytic transport solutions for 100 values of Σ_t uniformly sampled from $(1.0 \times 10^{-5}, 12 \times 10^{-5})$ cm^{-1}. Each solution, termed a snapshot, is normalized to unit norm, and placed as a column in a snapshot matrix $\boldsymbol{X}$. To apply POD to this data, the SVD of the real-valued snapshot matrix is computed:

$$\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^T\,, \tag{2}$$

where $\boldsymbol{U}$ contains the left singular vectors, $\boldsymbol{\Sigma}$ the singular values, and $\boldsymbol{V}$ the right singular vectors. This decomposition exists for any rectangular matrix. To check the quality of the decomposition, the snapshots generated are shown in Figure 1a and the reconstructed snapshots using SVD according to Eq. (2) are shown in Figure 1b. Note that there is a threshold below which the reconstruction is inaccurate. This is caused by the high variability of the snapshots and the inability of machine precision to capture their range. The dataset spans over 20 decades. When computing the SVD of the data to extract its dominant features, the snapshot matrix undergoes matrix operations that combine very large and small values together. Because of lack of precision, data loss occurs when processing such drastically different values.

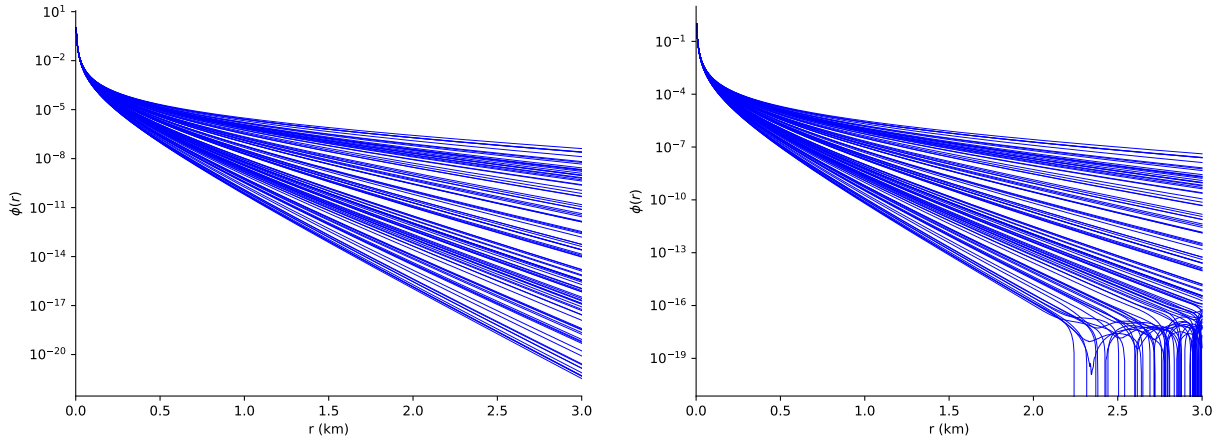

(a) Normalized snapshots for 100 realizations of Σ_t. Note that the snapshots span over 20 decades.

(b) Reconstruction (inaccurate) of the snapshot matrix using Eq. (2).

Figure 1: Normalized singular values of the 100 snapshots and their SVD reconstruction.

The remainder of this paper is as follows: Section 2 introduces the neutral radiation transport equation as the FOM. Section 3 discusses a data-driven, projection-based method that utilizes POD to reduce the ROM and introduces the mrPOD method. Section 4 applies mrPOD to single and multigroup transport problems and presents the results. Section 5 provides some closing remarks.

2. Full-Order Model

In this section, we present the discrete ordinates (S_n; see [26]) transport solver, which is the basis for our FOM. The transport scenarios that yield highly variable solutions are often characterized by localized sources, which give rise to unphysical ray effects in S_n solutions. An uncollided flux treatment that mitigates these inaccuracies is discussed.

2.1. S_n Radiation Transport

We consider the steady-state, fixed-source, no-fission neutral-particle transport equation, with Dirichlet boundary conditions. Although transport can be time-dependent and occur in fissile media, this simplified version is sufficient for the matters at hand here and is given by

$$\vec{\Omega}\cdot\vec{\nabla}\psi(\vec{r},E,\vec{\Omega})+\Sigma_t(\vec{r},E)\psi(\vec{r},E,\vec{\Omega})=\int_{4\pi}d\Omega'\int_0^\infty dE'\Sigma_s(\vec{r},E'\to E,\vec{\Omega}'\to\vec{\Omega})\psi(\vec{r},E',\vec{\Omega}') + Q(\vec{r},E,\vec{\Omega})\,,\qquad \forall\vec{r}\in\mathcal{D}\,. \quad (3a)$$

$$\psi(\vec{r},E,\vec{\Omega})=\psi_{\text{inc}}(\vec{r},E,\vec{\Omega})\quad \forall\vec{r}\in\partial\mathcal{D}^-=\left\{\vec{r}\in\partial\mathcal{D}\text{ such that }\vec{\Omega}\cdot\vec{n}(\vec{r})<0\right\}\,, \quad (3b)$$

The notation is as follows: $\vec{\Omega}$ (sr) denotes direction; r (cm) denotes position; E denotes energy (MeV); ψ (particles cm^{-2} s^{-1} sr^{-1} MeV^{-1}) denotes the angular flux and is the solution to Eq. (3); Σ_t (cm^{-1}) denotes

the total interaction cross section of the medium; Σ_s ($\mathrm{cm^{-1}\ sr^{-1}\ MeV^{-1}}$) denotes the differential scattering cross section of the medium; Q (particles $\mathrm{cm^{-3}\ s^{-1}\ sr^{-1}\ MeV^{-1}}$) denotes the external source of particles; $\mathcal{D}$ denotes the spatial domain; $\partial\mathcal{D}$ denotes the boundary of $\mathcal{D}$; ψ_{inc} denotes the prescribed boundary condition for directions incident into $\mathcal{D}$. Splitting the energy spectrum into G discrete intervals and evaluating the transport equation at D discrete directions, the coupled multigroup discrete ordinates ($\mathrm{S_n}$) equations [27–30] are obtained:

$$\vec{\Omega}_d \cdot \vec{\nabla}\psi^g(\vec{r},\vec{\Omega}_d) + \Sigma_t^g(\vec{r})\psi^g(\vec{r},\vec{\Omega}_d) = \sum_{g'=1}^{G}\sum_{\ell=0}^{L}\sum_{m=-\ell}^{\ell}\frac{2\ell+1}{4\pi}Y_{\ell m}(\vec{\Omega}_d)\Sigma_{s,\ell}^{g'\to g}(\vec{r})\phi_{\ell m}^{g'}(\vec{r}) + Q_g(\vec{r},\vec{\Omega}_d)\,, \qquad \forall\vec{r}\in\mathcal{D}\,, \qquad \forall g\in[1,G]\,, \qquad \forall d\in[1,D]\,, \quad (4a)$$

$$\psi^g(\vec{r},\vec{\Omega}_d) = \psi^g_{\mathrm{inc}}(\vec{r},\vec{\Omega}_d) \quad \forall\vec{r}\in\partial\mathcal{D}^-\,, \quad (4b)$$

where d subscripts denote the d-th discrete direction and g superscripts denote the g-th energy group. The scattering term has been expanded using a Legendre expansion for the scattering cross section and a spherical harmonics expansion for the angular flux. The new terms introduced here are the scattering order L, the real-valued spherical harmonic function $Y_{\ell m}$ of degree ℓ and order m, the scattering cross section moments $\Sigma_{s,\ell}$, and the flux moments $\phi_{\ell m}$ given by

$$\phi^g_{\ell m}(\vec{r}) = \int_0^{4\pi}\psi^g(\vec{r},\vec{\Omega})Y_{\ell m}(\vec{\Omega})d\Omega \approx \sum_{d=1}^{D} w_d\psi^g(\vec{\Omega}_d)Y_{\ell m}(\vec{\Omega}_d)\,, \quad (5)$$

that are approximated with angular quadrature weights w_d.

With transport discretized in angle and energy, all that remains is the spatial discretization. There are several spatial discretization options available, such as finite difference, finite volume, and finite element methods. Because the spatial discretization scheme is irrelevant to this work, further details are omitted. The interested Reader may consult [26] for spatial finite-differencing and [31] for discontinuous finite elements applied to the $\mathrm{S_n}$ transport equation. Without loss of generality, the results presented here will be for a 1D spherical geometry, with diamond and weighted-diamond differencing for spatial and angular discretization, respectively; see [32].

Using matrix notation (as in [33]), the transport equation can be expressed as

$$\boldsymbol{L}\vec{\Psi} = \boldsymbol{M}\boldsymbol{\Sigma}\underbrace{\vec{\Phi}}_{:=\boldsymbol{D}\vec{\Psi}} + \vec{Q}\,, \quad (6)$$

with $\vec{\Psi}$ the solution vector, $\boldsymbol{L}$ the streaming+collision matrix, $\boldsymbol{\Sigma}$ the scattering matrix, and $\boldsymbol{M}$ and $\boldsymbol{D}$ the moment-to-discrete and discrete-to-moment operators. When transport sweeps are employed, the following system is solved (also known as sweep-preconditioning):

$$(\boldsymbol{I} - \boldsymbol{D}\boldsymbol{L}^{-1}\boldsymbol{M}\boldsymbol{\Sigma})\vec{\Phi} = \boldsymbol{D}\boldsymbol{L}^{-1}\vec{Q}\,, \quad (7)$$

where, in practice, $\boldsymbol{L}^{-1}$ is a matrix-free operation, carried out cell by cell and often denoted as a transport sweep. Note that Eq. (7) is usually solved using a Krylov subspace method (i.e., GMRes) or with source iteration when written as Eq. (6). Finally, we denote by $\vec{\mu}$ the set of uncertain parameters (cross sections and source strengths in our case) and make this dependency explicit in the relevant operators:

$$\left(\boldsymbol{I} - \boldsymbol{D}\boldsymbol{L}^{-1}(\vec{\mu})\boldsymbol{M}\boldsymbol{\Sigma}(\vec{\mu})\right)\vec{\Phi}(\vec{\mu}) = \boldsymbol{D}\boldsymbol{L}^{-1}(\vec{\mu})\vec{Q}(\vec{\mu}). \tag{8}$$

2.2. Uncollided Flux Treatment

In transport problems with localized sources and little scattering, particles are capable of traveling, or streaming, large distances from their point of origin before undergoing a scattering interaction that changes their direction. These so-called "uncollided" particles can significantly contribute to the overall transport solution. Angular discretization techniques such as the aforementioned discrete ordinates method are often insufficient to resolve solutions that are not smoothly distributed in angle (due to the lack of rotational invariance of discrete-ordinate methods). Hence, the uncollided portion of the solution is often tainted with large angular discretization error, numerically observed as an unphysical phenomenon termed ray effects. Thus, for streaming-dominated transport problems with localized sources, it is common practice to utilize uncollided flux solution techniques [34]. The uncollided flux technique first splits the transport solution into uncollided (uc) and collided (c) components:

$$\vec{\Psi} = \vec{\Psi}^{uc} + \vec{\Psi}^{c}\,. \tag{9}$$

Substituting Eq. (9) into Eq. (6) yields

$$\boldsymbol{L}(\vec{\Psi}^{uc} + \vec{\Psi}^{c}) = \boldsymbol{M}\boldsymbol{\Sigma}(\vec{\Phi}^{uc} + \vec{\Phi}^{c}) + \vec{Q}\,. \tag{10}$$

Equation (10) is the sum of the following two equations:

$$\boldsymbol{L}\vec{\Psi}^{uc} = \vec{Q}\,, \tag{11a}$$

$$\boldsymbol{L}\vec{\Psi}^{c} = \boldsymbol{M}\boldsymbol{\Sigma}\vec{\Phi}^{c} + \underbrace{\boldsymbol{M}\boldsymbol{\Sigma}\vec{\Phi}^{uc}}_{:=\vec{Q}^{fc}}, \tag{11b}$$

where $\vec{Q}^{fc}$ denotes the "first collided" source. Equation (11a) represents transport through a pure absorber with spatially localized external source $\vec{Q}$, to which the solution is the uncollided flux $\vec{\Psi}^{uc}$. Equation (11b) represents transport through the original material with spatially distributed first-collision source $\vec{Q}^{fc}$ acting as an external source, to which the solution is the collided flux $\vec{\Psi}^{c}$. Since Eq. (11b) does not have a localized source, it may be more amenable to be solved using the discrete ordinates method. As for the uncollided flux equation Eq. (11a), a semianalytic method known as ray-tracing is employed to obtain an accurate solution. Denoting the streaming+collision operator discretized with the ray-tracing method with the RT superscript and that of the discrete ordinates method with the $\mathrm{S_n}$ superscript, Eq. (11) becomes

$$\boldsymbol{L}^{\mathrm{RT}}\vec{\Psi}^{uc} = \vec{Q}\,, \tag{12a}$$

$$\boldsymbol{L}^{\mathrm{S_n}}\vec{\Psi}^c = \boldsymbol{M}\boldsymbol{\Sigma}\vec{\Phi}^c + \vec{Q}^{fc}\,. \tag{12b}$$

Once Eqs. (12) have been solved, the total solution is recovered through Eq. (9). Note that in the case of a spherical 1D coordinate system with a centered external source (the present work), the analytic uncollided solution is obtainable from Eq. (1).

3. Reduced-Order Model

This section describes a data-driven, projection-based ROM supplemented by POD. High-fidelity data from the FOM is first processed using POD in the offline training phase to identify a suitable subspace in which to construct the ROM. Once built, the ROM may be used in the online prediction phase to rapidly obtain transport solutions for unseen parameter realizations. Finally, the new mrPOD method is introduced and incorporated into the ROM. These derivations are first performed for a generic steady-state linear parametric partial differential equation (PDE) FOM and then applied to transport.

3.1. Offline Phase

The transport equation discussed in Section 2 is a steady-state linear PDE. For conciseness and to keep the discussion applicable to any such FOM, note that the discretization of any steady-state linear PDE with parameterization $\vec{\mu}$ yields a linear system of equations that is solved for the discretized solution $\vec{x}$:

$$\boldsymbol{A}(\vec{\mu})\vec{x}(\vec{\mu}) = \vec{b}(\vec{\mu})\,. \tag{13}$$

Equation (13) will be referred to as the FOM.

A data-driven reduced-order model requires data from the full-order model. The parameter space is adequately sampled and the FOM is used to compute solutions for each sample. A full-order solution, or "snapshot", is generated for each parameter sample. These snapshots comprise the data used to train the ROM. Let n_{snap} denote the number of snapshots computed and N the length of each snapshot vector. Once the parameters $\{\vec{\mu}_j\}_{j=1}^{n_{\mathrm{snap}}}$ have been sampled and the corresponding snapshots $\{\vec{x}(\vec{\mu}_j)\}_{j=1}^{n_{\mathrm{snap}}}$ are obtained by solving the size N system in Eq. (13), the snapshot vectors are placed in a matrix $\boldsymbol{X}$ of shape $N \times n_{\mathrm{snap}}$, with each snapshot occupying a column:

$$\boldsymbol{X} = \begin{pmatrix} \vec{x}(\vec{\mu}_1) & \vec{x}(\vec{\mu}_2) & \cdots & \vec{x}(\vec{\mu}_{n_{\mathrm{snap}}}) \end{pmatrix}. \tag{14}$$

This data may now be used in POD to identify a subspace in which to build the ROM.

POD is a key ingredient in the derivation of our reduced-order model. For a review of the history and uses of POD, the reader is referred to the work in [23, 35] . We briefly recall POD next. Any real-valued rectangular matrix $\boldsymbol{X}$ possesses a singular value decomposition of the form

$$\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{V}^T, \tag{15}$$

where $\boldsymbol{U}$ and $\boldsymbol{V}$ are orthonormal matrices containing the left and right singular vectors of $\boldsymbol{X}$, respectively, and $\boldsymbol{\Lambda}$ is a non-negative diagonal matrix containing the singular values of $\boldsymbol{X}$ in descending order (i.e., $\boldsymbol{\Lambda}_{11} \geq \boldsymbol{\Lambda}_{22} \geq \cdots$). If $\boldsymbol{X}$ is of dimension $N \times n_{\text{snap}}$, then $\boldsymbol{U}$, $\boldsymbol{\Lambda}$, and $\boldsymbol{V}^T$ are of dimensions $N \times n_{\text{snap}}$, $n_{\text{snap}} \times n_{\text{snap}}$, and $n_{\text{snap}} \times n_{\text{snap}}$, respectively.

The singular value decomposition of the snapshot matrix in Eq. (15) compresses, extracts, and ranks the features embedded in the dataset. The left singular vectors, also called POD modes, contained in $\boldsymbol{U}$ form an uncorrelated basis for the snapshots. The relative magnitudes of the singular values in $\boldsymbol{\Lambda}$ indicate the relative importance of the POD modes with respect to each other. Thus, the less important modes may be discarded with negligible loss of information, resulting in a truncated basis in which to seek ROM solutions for unseen parameter realizations. One typically chooses the number of POD modes to keep such that a certain fraction of FOM information is retained in the ROM. This fraction I is called the information content of the modes kept and is given by

$$I(r) = \frac{\sum_{j=1}^{r} \boldsymbol{\Lambda}_{jj}^2}{\sum_{j=1}^{n_{\text{snap}}} \boldsymbol{\Lambda}_{jj}^2}. \tag{16}$$

Here, $r \ll N$ is the number of POD modes kept and $\boldsymbol{\Lambda}_{jj}$ is the j-th singular value. Given a desired information content threshold $\hat{I}$ (i.e., $\hat{I} = 99.999$ %), r may be chosen to be the smallest value such that $I(r) \geq \hat{I}$. With the number of modes in the subspace determined, we denote by $\boldsymbol{U}_r$ the $N \times r$ matrix containing the first r POD modes as its columns. It may be shown that, for any truncation value r, the POD modes from the SVD of the snapshot matrix are an optimal basis in the least-squares sense in which to represent the snapshot data [36].

3.2. Projection & Online Phase

The ROM seeks to approximate the solution to the FOM as a linear combination of r POD modes. That is, for a realization of the parameters $\vec{\theta}$ not in the set of snapshot parameters, the reduced solution is

$$\vec{x}_r(\vec{\theta}) = \sum_{j=1}^{r} c_j(\vec{\theta})\vec{u}_j = \boldsymbol{U}_r \vec{c}(\vec{\theta}), \tag{17}$$

where c_j is the expansion coefficient for the j-th POD mode $\vec{u}_j$ (i.e., the jth column of $\boldsymbol{U}$). $\vec{c}$ is the vector containing the expansion coefficients. Recall that the FOM is given in Eq. (13) as the $N \times N$ system

$$\boldsymbol{A}(\vec{\theta})\vec{x}(\vec{\theta}) = \vec{b}(\vec{\theta}), \tag{18}$$

Substituting Eq. (17) into Eq. (18),

$$\boldsymbol{A}(\vec{\theta})\vec{x}_r(\vec{\theta}) = \boldsymbol{A}(\vec{\theta})\boldsymbol{U}_r\vec{c}(\vec{\theta}) = \vec{b}(\vec{\theta}). \tag{19}$$

This is an overdetermined $N \times r$ system (N equations, r unknowns) with $N \gg r$. To remedy this, we take a projection of Eq. (19) using projection matrix $\boldsymbol{W}(\vec{\theta})$ of shape $N \times r$ to obtain a solvable $r \times r$ system:

$$\underbrace{\boldsymbol{W}^T(\vec{\theta})\boldsymbol{A}(\vec{\theta})\boldsymbol{U}_r}_{:=\boldsymbol{A}_r(\vec{\theta})}\vec{c}(\vec{\theta}) = \underbrace{\boldsymbol{W}^T(\vec{\theta})\vec{b}(\vec{\theta})}_{:=\vec{b}_r(\vec{\theta})}. \tag{20}$$

Two common choices for the projection matrix are Galerkin ($\boldsymbol{W} = \boldsymbol{U}_r$) and least-squares Petrov-Galerkin ($\boldsymbol{W}(\vec{\theta}) = \boldsymbol{A}(\vec{\theta})\boldsymbol{U}_r$) projections. Once Eq. (20) is solved for $\vec{c}(\vec{\theta})$, the ROM solution is recovered through Eq. (17).

In summary, the ROM reduces a $N \times N$ system to a $r \times r$ system with $r \ll N$. The original number of unknowns, N, can reach up to 10^{12} - 10^{14} [37–39], while the number of dominant POD modes r is more often on the order of 10^1. Thus, if properly utilized, the ROM results in enormous computational savings while producing solutions acceptably close to its FOM counterpart.

3.3. Multiresolution Proper Orthogonal Decomposition

The purpose of this method is to solve the inaccuracies that arise from the limitations of machine-precision arithmetic when computing the SVD of widely varying data. The main idea of mrPOD is to decompose a transport problem's spatial domain into regions in which the snapshots vary less than across the domain as a whole. The SVD of each region is computed and found to be more accurate than the SVD of all regions taken together. The POD modes from each region's SVD are then used in the ROM.

Notation-wise, we say a snapshot matrix $\boldsymbol{X}$ is split into n_{reg} non-overlapping spatial regions:

$$\boldsymbol{X} = (\boldsymbol{X}_1, \boldsymbol{X}_2, \ldots, \boldsymbol{X}_{n_{\text{reg}}})^T \,, \tag{21}$$

where a snapshot in region i has length N_i such that

$$\sum_{i=1}^{n_{\text{reg}}} N_i = N \,. \tag{22}$$

Note that a region does not have to be spatially continuous. The SVD for each region is given by

$$\boldsymbol{X}_i = \boldsymbol{U}_i \boldsymbol{\Sigma}_i \boldsymbol{V}_i \,, \qquad \forall i \in [1, n_{\text{reg}}] \,. \tag{23}$$

Each resulting POD mode in $\boldsymbol{U}_i$ is of length N_i. To be compatible with the ROM, the modes must be of length N. To this end, each mode in $\boldsymbol{U}_i$ is padded with zeros for all other regions. After this, the POD mode matrix for all accumulated regions has the form of a block diagonal matrix with block i, j given by

$$\boldsymbol{U}_{i,j} = \delta_{ij} \boldsymbol{U}_i \,, \qquad \forall i \in [1, n_{\text{reg}}] \,, \qquad \forall j \in [1, n_{\text{reg}}] \,, \tag{24}$$

where δ_{ij} denotes the Kronecker delta function. The resulting matrix gathers each $\boldsymbol{U}_i$ of shape $N_i \times n_{\text{snap}}$ into $\boldsymbol{U}$ of shape $N \times n_{\text{reg}} n_{\text{snap}}$. Note that $\boldsymbol{U}$ should be stored in sparse format to avoid storing the blocks of zeros. Also note that because the POD modes for each region are generated independently from each other, they cannot simply be "glued" together into a $\boldsymbol{U}$ matrix of shape $N \times n_{\text{snap}}$. Doing so would result in a poorly performing ROM that imposes the same expansion coefficients in all regions. This expansion of the reduced subspace with more POD modes results in a slightly larger system for the ROM to solve. Given that ROM systems are small, this is a small and possibly insignificant penalty for using mrPOD over POD. With $\boldsymbol{U}$ built, we are free to rearrange its columns as desired. A reason for doing this would be to have the

dominate POD modes for each region to occupy the leftmost columns of $\boldsymbol{U}$. For instance, the first POD modes in each region may be gathered into the first n_{reg} columns of $\boldsymbol{U}$, the second POD modes into columns $(n_{\text{reg}}, 2n_{\text{reg}}]$, and so on. Doing so would have the effect that keeping the first rn_{reg} columns of $\boldsymbol{U}$ for use in the ROM would actually be keeping the r most important POD modes in each region. Other, more complex, arrangements are possible, such as keeping a variable number of modes in each region that is determined by the region's information content in Eq. (16). Once the modes from each region are selected for use in the ROM, $\boldsymbol{U}$ is truncated to obtain $\boldsymbol{U}_r$. From there, the ROM procedure is identical to the one-region case of POD.

Determination of the number of regions n_{reg} and the location of the region boundaries is an open question. The snapshots are generated first and, then, the number and location of regions are chosen based on the range of the data. The range spanned by the snapshot data is a determining factor in choosing the number of regions. n_{reg} should be large enough to ensure that the data in each region can be accurately decomposed. We do not know of any optimal methods for choosing the location of the regions and, instead, rely on some heuristics. Two such heuristics involve choosing region boundaries such that (1) the data within each region spans the same number of decades as all other regions or (2) the normalized singular values in each region have similar decay rates. Option (2) is not practical for large 3D problems, as it involves re-computing SVDs for each proposed set of regions.

Consider the motivating example from the introduction in Section 1. mrPOD is now applied to the same data but $n_{\text{reg}} = 2$ regions are employed. Figure 2a shows the normalized snapshots in each region while Figure 2b shows the reconstruction for each region by means of Eq. (23). Note that the use of mrPOD has eliminated the inaccuracies present in POD (Figure 1b).

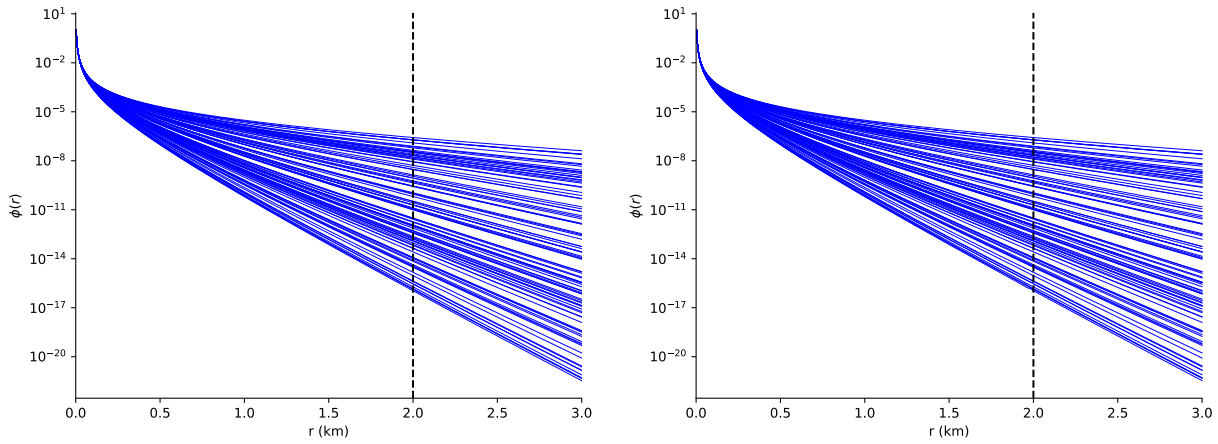

(a) Normalized snapshots for 100 realizations of Σ_t. Dashed vertical lines indicate region boundaries.

(b) Reconstruction (accurate) of the snapshot matrix using Eq. (23).

Figure 2: Normalized singular values of the 100 snapshots and their region-wise SVD reconstruction.

We next briefly discuss an alternative method to mrPOD that relies on a logarithmic transform of the

original snapshot data. This approach is sound and suggested for use with non-intrusive ROMs. However, it can not be applied to projection-based ROM because it introduces a non-linearity between the original full-order system and the reduced-order system obtained for the expansion coefficients. Denote the element-wise logarithm of the full-domain snapshot matrix $\boldsymbol{X}$ by $\boldsymbol{X}^\ell := \log(\boldsymbol{X})$. The SVD of $\boldsymbol{X}^\ell$ does not suffer from machine-precision issues:

$$\boldsymbol{X}^\ell = \boldsymbol{U}^\ell \boldsymbol{\Lambda}^\ell \boldsymbol{V}^{\ell,T} \,. \tag{25}$$

Two possible ROMs based on the logarithm of the data are:

1. Using $\exp(\boldsymbol{U}^\ell)$ as the basis vectors and $\exp(\boldsymbol{U}^\ell)\vec{c}$ as the approximation to the transport solution.
2. Using $\boldsymbol{U}^\ell$ as the basis vectors and $\boldsymbol{U}^\ell\vec{c}$ as the approximation to the transport solution.

Both approaches have yielded poor results because they introduce non-linearity in the projection-based ROM framework. Approach-1 assumes that $\boldsymbol{U} = \exp(\boldsymbol{U}^\ell)$, which is not true. The actual relationship is given by

$$\boldsymbol{X} = \boldsymbol{U}\boldsymbol{\Lambda}\boldsymbol{V}^T = \exp(\boldsymbol{X}^\ell) = \exp(\boldsymbol{U}^\ell\boldsymbol{\Lambda}^\ell\boldsymbol{V}^{\ell,T}) \,. \tag{26}$$

Solving for $\boldsymbol{U}$ in terms of $\boldsymbol{U}^\ell$ is not straightforward because $\boldsymbol{\Lambda}$ and $\boldsymbol{V}$ are also unknown. Approach-2 uses $\boldsymbol{U}^\ell$ as the basis vectors. Using the POD modes of the log of the data as basis vectors is a poor basis for the untransformed transport solutions and yields poor results when used in a projection-based ROM with the untransformed transport operators.

3.4. Application to Transport

Now that the ROM framework for the generic parametric steady-state linear PDE in Eq. (13) has been established, it is now applied to radiation transport. Recall that transport may be written in different forms such as in Eqs. (6) and (7). These are restated here in parametric form to identify the generic operators from Eq. (13):

$$\underbrace{(\boldsymbol{L}(\vec{\mu}) - \boldsymbol{M}\boldsymbol{\Sigma}(\vec{\mu})\boldsymbol{D})}_{:=\boldsymbol{A}(\vec{\mu})} \underbrace{\vec{\Psi}(\vec{\mu})}_{:=\vec{x}(\vec{\mu})} = \underbrace{\vec{Q}(\vec{\mu})}_{:=\vec{b}(\vec{\mu})} \,, \tag{27}$$

$$\underbrace{(\boldsymbol{I} - \boldsymbol{D}\boldsymbol{L}^{-1}(\vec{\mu})\boldsymbol{M}\boldsymbol{\Sigma}(\vec{\mu}))}_{:=\boldsymbol{A}(\vec{\mu})} \underbrace{\vec{\Phi}(\vec{\mu})}_{:=\vec{x}(\vec{\mu})} = \underbrace{\boldsymbol{D}\boldsymbol{L}^{-1}(\vec{\mu})\vec{Q}(\vec{\mu})}_{:=\vec{b}(\vec{\mu})} \,. \tag{28}$$

Thus, the form of the FOM depends on whether transport is expressed in terms of the angular flux $\vec{\Psi}$ or the flux moments $\vec{\Phi}$. A ROM based on the former is known as a sweep-based ROM [5, 8], while the latter form results in an affine decomposition-based ROM [6, 7, 9]. Both ROMs have been studied and compared in great detail in the previous works cited above.

Recall from Section 2.2 that S_n transport problems with localized sources and streaming-dominated physics often require an uncollided flux treatment to accurately resolve the scattering source and the resulting transport solution. Since these types of problems belong to the class of those that have widely varying

solutions, they are ideal candidates to apply the mrPOD methodology to. The question of how to apply the ROM across the uncollided and collided transport equations naturally arises. These equations are restated in parametric form:

$$\boldsymbol{L}^{\mathrm{RT}}(\vec{\mu})\vec{\Psi}^{uc}(\vec{\mu}) = \vec{Q}(\vec{\mu})\,, \tag{29a}$$

$$\boldsymbol{L}^{\mathrm{S_n}}(\vec{\mu})\vec{\Psi}^{c}(\vec{\mu}) = \boldsymbol{M}\boldsymbol{\Sigma}(\vec{\mu})\vec{\Phi}^{c}(\vec{\mu}) + \vec{Q}^{fc}(\vec{\mu})\,. \tag{29b}$$

There are two options for the uncollided transport equation. If Eq. (29a) is computationally expensive to solve, it may be necessary to construct for it a ROM that stores the uncollided solution as snapshots. If Eq. (29a) may be solved cheaply, as is the case of an analytic uncollided solution (the present work), no ROM is required. In either case, whether Eq. (29a) is solved directly or emulated by a ROM, the first collision source $\vec{Q}^{fc}(\vec{\mu})$ may then be built for use in the collided transport equation (29b). A ROM may then be applied to the collided transport equation by storing the collided portion of the training data as the snapshots.

4. Results

To test the mrPOD ROM, two shielding problems are evaluated. Both involve transport from a localized source through the atmosphere and a shield. The first application is a one-group problem, while the second is a multigroup problem with a more involved parameterization. The ROM utilized is a sweep-based one with a Galerkin projection.

4.1. One-Group Shielding Problem

Consider a small 10-cm^3 volumetric isotropic radiation source centered at the origin of a spherical coordinate system filled with air. A 10-cm thick shield is positioned 1 km from the source. The 10-m region behind the shield is filled with a post-shield material and vacuum boundary conditions ($\psi_{inc} = 0$) are imposed at 1.01 km. An S_8 FOM with an analytic uncollided flux treatment is utilized. The parameters are the absorption and scattering cross sections of the atmosphere and of the shield, as well as the source magnitude; these are summarized in Table 1; $\mathcal{U}(a, b)$ stands for a uniform distribution of the parameter between values a and b.

Parameter	Material		
	Air	Shield	Source
Σ_a (cm^{-1})	$\mathcal{U}(2 \times 10^{-5}, 9 \times 10^{-5})$	$\mathcal{U}(0.05,\ 0.1)$	Same as air
Σ_s (cm^{-1})	$\mathcal{U}(2 \times 10^{-5}, 3.6 \times 10^{-4})$	$\mathcal{U}(0.05,\ 0.05)$	Same as air
Q (par cm^{-3} s^{-1})	0	0	$\mathcal{U}(5 \times 10^{14}, 5 \times 10^{15})$

Table 1: Parameter distributions for the one group shielding problem. $\mathcal{U}(a, b)$ denotes a parameter uniformly distributed between values a and $b > a$.

One hundred ($n_{\text{snap}} = 100$) parameter realizations were sampled using Latin hypercube sampling (LHS) and used as training data and are shown in Figure 3. The performance of POD was compared to that of 7-region mrPOD. Figure 4 compares the singular values for the two methods. The singular values from mrPOD decay faster than those from POD, because the snapshots vary less within a given region than over the whole domain. Because of this, the training data within a region may be represented with fewer POD modes than the data over the whole domain can. The ROMs were evaluated for several parameter realizations not present in the training snapshot samples and compared against their corresponding FOM solutions. Figure 5 compares the performance of the POD ROM to that of the mrPOD ROM for one such parameter realization. Even when all 30 of the significant POD modes of the POD ROM are used, the resulting ROM is a poor emulator of the FOM. This is due to machine precision inaccuracies in the SVD. The results from the 7-region mrPOD ROM are in excellent agreement with the full-order transport solution with 10 mrPOD modes (POD modes per region).

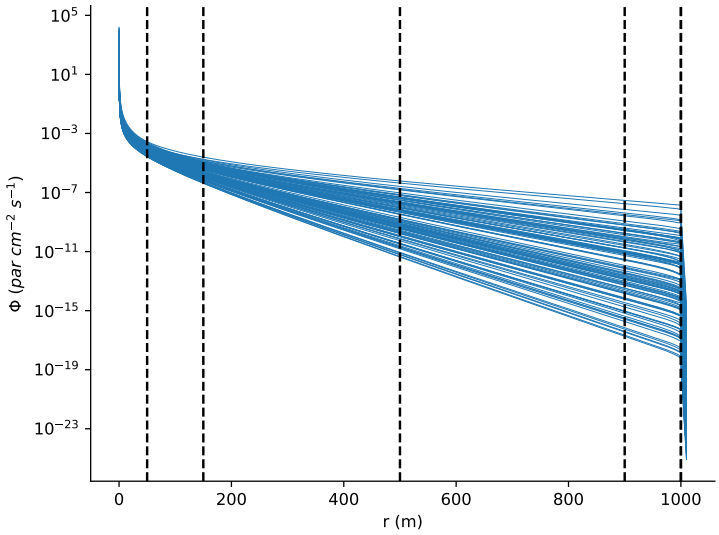

Figure 3: Normalized snapshots for the one group shielding problem's 100 parameter realizations. Dashed vertical lines indicate region boundaries. Note that the two boundaries near 1,000 m should not be mistaken as one line.

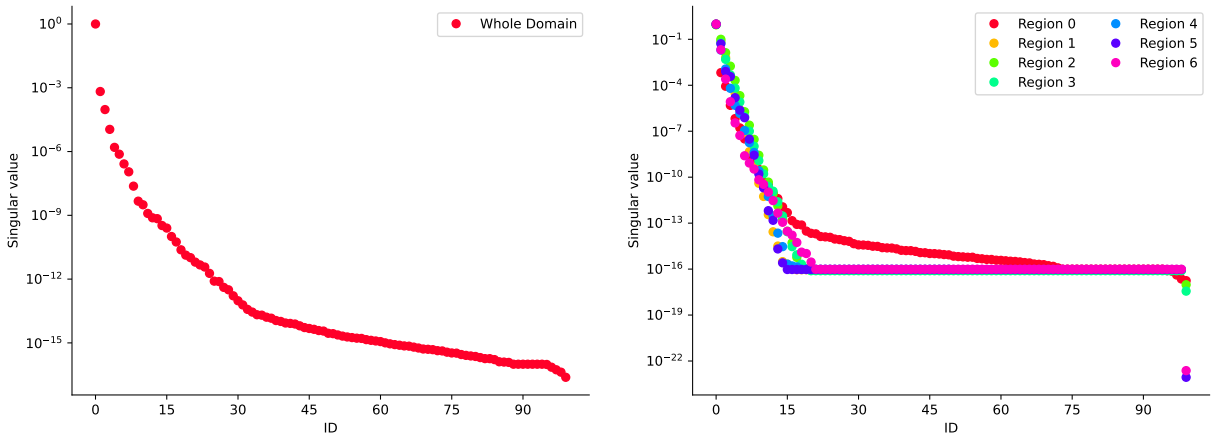

(a) POD. **(b) 7-region mrPOD.**

Figure 4: Normalized singular values of the 100 snapshots.

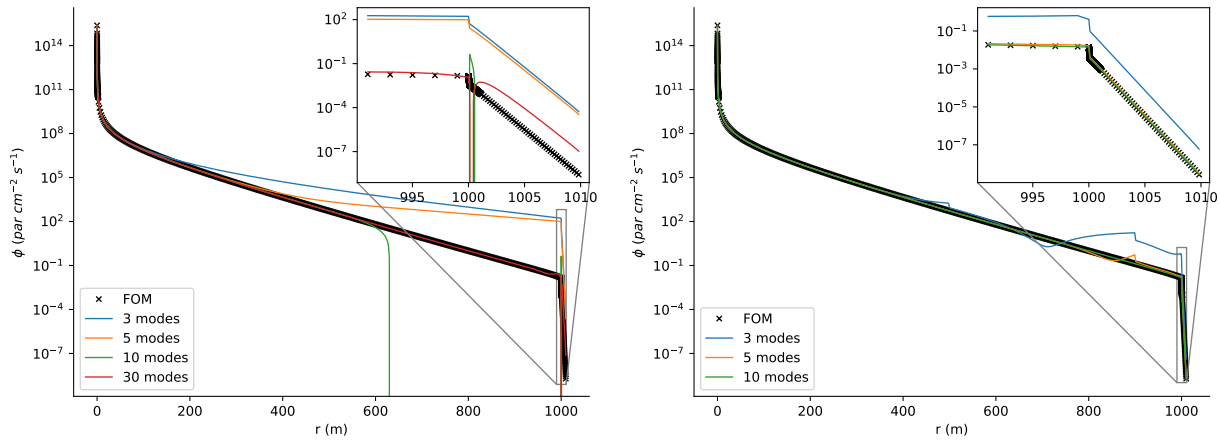

(a) POD. **(b) 7-region mrPOD.**

Figure 5: ROM performance for a selected unseen parameter realization.

4.2. Multigroup Group Shielding Problem

We now turn to a 6-energy group shielding application. Consider a small volumetric isotropic radiation source 1 cm in radius centered at the origin of a spherical coordinate system filled with air. The source spectrum is a combination of a U-235 fission spectrum and a 14-MeV fusion spectrum. A 1-m thick shield composed of borated concrete [40] is positioned 4 km away from the source. Cross sections were obtained from the NJOY21 nuclear data processing code. The domain boundary is placed at 5 km and vacuum boundary conditions are imposed. An S_8 FOM with an analytic uncollided flux treatment is utilized. The parameters are the humidity content of the air, the mass fraction of boron in the shield, and the source's fission spectrum weight relative to the total source; these are summarized in Table 2.

Parameter	Distribution
Specific humidity in air (kg water / kg dry air)	$\mathcal{U}(0, 0.0147)$
Shield boron weight fraction	$\mathcal{U}(0, 0.25)$
Source fission weight	$\mathcal{U}(0, 1)$

Table 2: Parameter distributions for the 31 group shielding problem.

One hundred ($n_{\text{snap}} = 100$) parameter realizations were sampled using LHS and used as training data and are shown in Figure 6. The performance of POD was compared to that of 4-region mrPOD. Figure 7 compares the singular values for the two methods. The singular values in region 0 decay in a comparable manner as those in POD, while regions 1-3 decay more slowly. The ROMs were evaluated for several parameter realizations not in the snapshot samples and compared to their corresponding FOM solutions. Figure 8 compares the performance of the 5- and 20-mode POD ROM to that of the 5 mode per region mrPOD ROM for one such parameter realization. The 20-mode POD ROM is included because its reduced system contains 20 degrees of freedom, the same as the 4-region, 5-mode mrPOD ROM, allowing for a fair comparison of the methods. Again, mrPOD mitigates the machine precision issue that plagues POD. The 4-region mrPOD ROM does an excellent job of approximating the transport solution with 5 mrPOD modes.

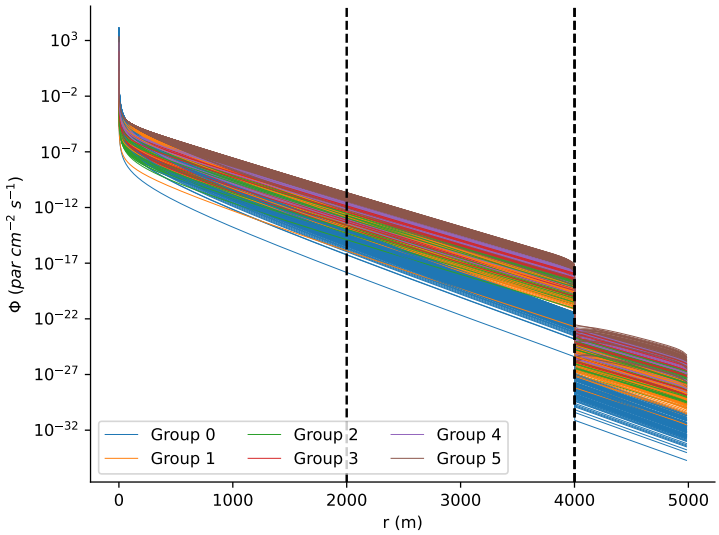

Figure 6: Normalized snapshots for the 6-group shielding problem's 100 parameter realizations. Dashed vertical lines indicate region boundaries. Note that the two boundaries near 4000 m should not be mistaken as one line.

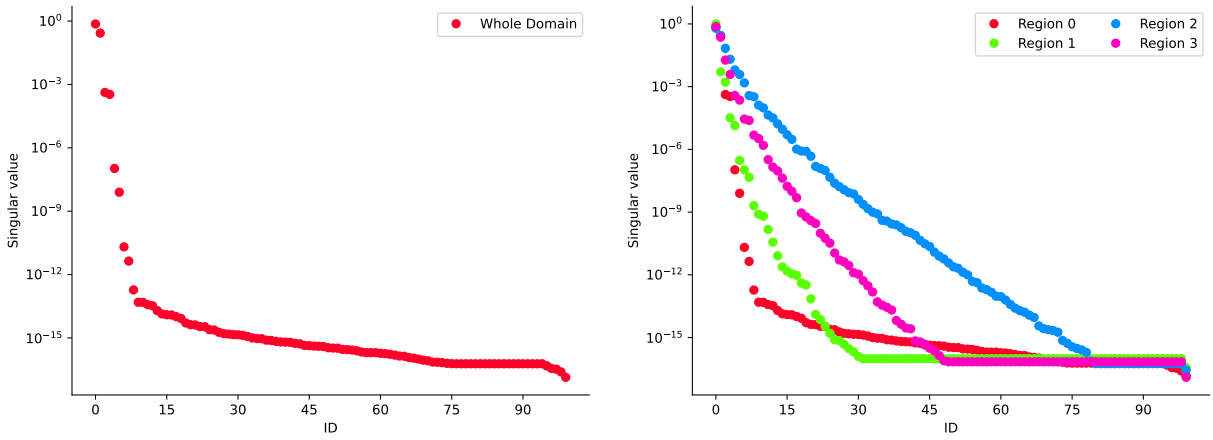

(a) POD. **(b) 4-region mrPOD.**

Figure 7: Normalized singular values of the 100 snapshots.

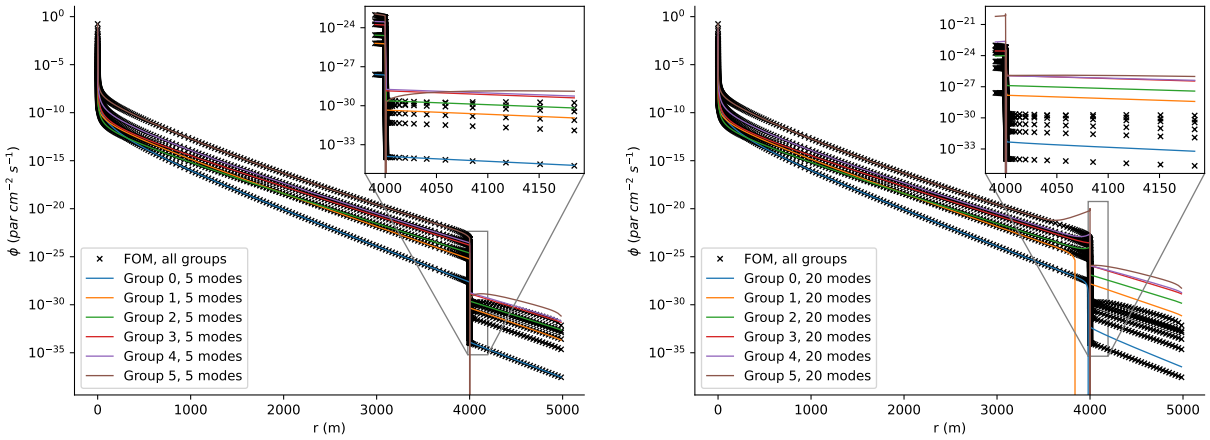

(a) 5-mode POD.

(b) 20-mode POD.

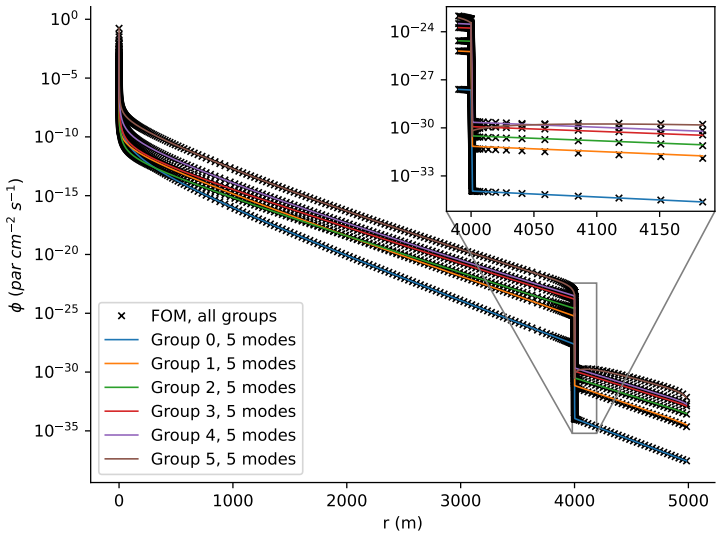

(c) 4-region, 5-mode mrPOD.

Figure 8: ROM performance for a selected unseen parameter realization.

5. Conclusions

We have developed projection-based Reduced-Order Models (ROM) for radiation transport using a multiresolution approach. The multiresolution POD (mrPOD) aims at combating the inaccurate ROM subspace obtained when the training data greatly varies in space due to the limitation of machine-precision arithmetic during the SVD calculation. mrPOD utilizes domain decomposition and region-wise POD. Because the training data varies less within a region than over the whole domain, POD is more accurate within each region. A small drawback to mrPOD is that it increases the number of POD modes by a factor of n_{reg}, but this is an insignificant cost increase, especially considering that global (i.e., one-region) POD can be quite inaccurate. The mrPOD method was demonstrated to yield accurate ROMs in both single group and

multigroup 1D spherical radiation transport applications with over 30 decades of attenuation across the spatial domain. With these successful demonstrations, this work can be extended to implement mrPOD in higher dimensional space (e.g., 3D) transport solvers. In this case, the uncollided flux treatment will no longer be analytic as in this work and will require more computational resources to solve (such as ray tracing). Thus, an additional investigation into whether the uncollided treatment itself would benefit from model-order reduction is warranted.

Acknowledgments

Part of this work has been made possible through a grant by the Department of the Defense, Defense Threat Reduction Agency under Award No. HDTRA1-18-1-0020. The content of the information does not necessarily reflect the position or the policy of the federal government, and no official endorsement should be inferred.

This material is based upon work supported under an Integrated University Program Graduate Fellowship. Any opinions, findings, conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the Department of Energy Office of Nuclear Energy.

References

[1] J. S. Hesthaven, G. Rozza, B. Stamm, Certified Reduced Basis Methods for Parameterized Partial Differential Equations, Springer International Publishing, 2016. doi:10.1007/978-3-319-22470-1.

[2] A. Quarteroni, G. Rozza (Eds.), Reduced Order Methods for Modeling and Computational Reduction, Springer International Publishing, 2014. doi:10.1007/978-3-319-02090-7.

[3] A. Quarteroni, A. Manzoni, F. Negri (Eds.), Reduced Basis Methods for Partial Differential Equations, Springer International Publishing, 2016. doi:10.1007/978-3-319-15430-2.

[4] P. Benner, M. Ohlberger, A. Patera, G. Rozza, K. Urban (Eds.), Model Reduction of Parameterized Systems, Springer International Publishing, 2017. doi:10.1007/978-3-319-58785-1.

[5] P. A. Behne, J. C. Ragusa, J. E. Morel, Model-order reduction approaches for s_N radiation transport, International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (2019).

[6] P. A. Behne, J. C. Ragusa, M. Tano, Projection-based model order reduction based on affine decomposition, International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering (2021).

[7] M. Tano, J. Ragusa, D. Caron, P. Behne, Affine reduced-order model for radiation transport problems in cylindrical coordinates, Annals of Nuclear Energy 158 (2021) 108214. doi:https://doi.org/10.1016/j.anucene.2021.108214.

[8] P. Behne, J. Vermaak, J. C. Ragusa, Minimally-invasive parametric model-order reduction for sweep-based radiation transport, Journal of Computational Physics (submitted).

[9] P. Behne, J. Vermaak, J. C. Ragusa, Parametric model-order reduction for radiation transport simulations based on an affine decomposition of the operators, Nuclear Science and Engineering (submitted).

[10] J. N. Kutz, X. Fu, S. L. Brunton, Multiresolution dynamic mode decomposition, SIAM Journal on Applied Dynamical Systems 15 (2) (2016) 713–735. arXiv:https://doi.org/10.1137/15M1023543, doi:10.1137/15M1023543.

[11] M. A. Mendez, M. Balabane, J.-M. Buchlin, Multi-scale proper orthogonal decomposition of complex fluid flows, Journal of Fluid Mechanics 870 (2019) 988–1036.

[12] J. Baiges, R. Codina, S. Idelsohn, A domain decomposition strategy for reduced order models. application to the incompressible navier–stokes equations, Computer Methods in Applied Mechanics and Engineering 267 (2013) 23–42. doi:https://doi.org/10.1016/j.cma.2013.08.001.

[13] D. Amsallem, M. J. Zahr, C. Farhat, Nonlinear model order reduction based on local reduced-order bases, International Journal for Numerical Methods in Engineering 92 (10) (2012) 891–916. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.4371, doi:https://doi.org/10.1002/nme.4371.

[14] S. Sahyoun, S. Djouadi, Local proper orthogonal decomposition based on space vectors clustering, in: 3rd International Conference on Systems and Control, 2013, pp. 665–670. doi:10.1109/ICoSC.2013.6750930.

[15] J. L. Eftang, A. T. Patera, E. M. Rønquist, An "hp" certified reduced basis method for parametrized elliptic partial differential equations, SIAM Journal on Scientific Computing 32 (6) (2010) 3170–3200. arXiv:https://doi.org/10.1137/090780122, doi:10.1137/090780122.

[16] J. L. Eftang, D. J. Knezevic, A. T. Patera, An hp certified reduced basis method for parametrized parabolic partial differential equations, Mathematical and Computer Modelling of Dynamical Systems 17 (4) (2011) 395–422. arXiv:https://doi.org/10.1080/13873954.2011.547670, doi:10.1080/13873954.2011.547670.

[17] P. Díez, A. Muixí, S. Zlotnik, A. García-González, Nonlinear dimensionality reduction for parametric problems: A kernel proper orthogonal decomposition, International Journal for Numerical Methods in Engineering 122 (24) (2021) 7306–7327. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.6831, doi:https://doi.org/10.1002/nme.6831.

[18] M. Dihlmann, M. Dronhmann, B. Hassdonk, Model reduction of parameterized evolution problems using the reduced basis method with adaptive time partitioning, Technical report, Stuttgard Research Centre for Simulation Technology (2011).

[19] K. Hoang, Y. Fu, J. Song, An hp-proper orthogonal decomposition–moving least squares approach for molecular dynamics simulation, Computer Methods in Applied Mechanics and Engineering 298 (2016) 548–575. doi:https://doi.org/10.1016/j.cma.2015.10.003.

[20] A. Narasingam, P. Siddhamshetty, J. Sang-Il Kwon, Temporal clustering for order reduction of nonlinear parabolic pde systems with time-dependent spatial domains: Application to a hydraulic fracturing process, AIChE Journal 63 (9) (2017) 3818–3831. arXiv:https://aiche.onlinelibrary.wiley.com/doi/pdf/10.1002/aic.15733, doi:https://doi.org/10.1002/aic.15733.

[21] C. E. Heaney, A. G. Buchan, C. C. Pain, S. Jewer, Reduced-order modelling applied to the multigroup neutron diffusion equation using a nonlinear interpolation method for control-rod movement, Energies 14 (5) (2021). doi:10.3390/en14051350.

[22] P. German, J. C. Ragusa, Reduced-order modeling of parameterized multi-group diffusion k-eigenvalue problems, Annals of Nuclear Energy 134 (2019) 144–157. doi:https://doi.org/10.1016/j.anucene.2019.05.049.

[23] A. C. Hughes, A. G. Buchan, A discontinuous and adaptive reduced order model for the angular discretization of the boltzmann transport equation, International Journal for Numerical Methods in Engineering 121 (24) (2020) 5647–5666. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/nme.6516, doi:https://doi.org/10.1002/nme.6516.

[24] A. Cherezov, R. Sanchez, H. G. Joo, A reduced-basis element method for pin-by-pin reactor core calculations in diffusion and sp3 approximations, Annals of Nuclear Energy 116 (2018) 195–209. doi:https://doi.org/10.1016/j.anucene.2018.02.013.

[25] T. R. F. Phillips, C. E. Heaney, B. S. Tollit, P. N. Smith, C. C. Pain, Reduced-order modelling with domain decomposition applied to multi-group neutron transport, Energies 14 (5) (2021). doi:10.3390/en14051369.

[26] E. E. Lewis, W. F. Miller, Computational methods of neutron transport (1984).

[27] J. J. Duderstadt, L. J. Hamilton, Nuclear Reactor Analysis, John Wiley & Sons, Inc., 1976.

[28] W. R. D. Boyd III, Reactor agnostic multi-group cross section generation for fine-mesh deterministic neutron transport simulations, Ph.D. thesis, Massachusetts Institute of Technology (2017).

[29] G. I. Bell, S. Glasstone, Nuclear Reactor Theory, Van Nostrand Reinhold Company, 1970.

[30] T. R. Hill, ONETRAN: A Discrete Ordinates Finite Element Code for the Solution of the One-Dimensional Multigroup Transport Equation, Los Alamos National Laboratory (1975).

[31] W. H. Reed, T. R. Hill, Triangular mesh methods for the neutron transport equation, Tech. rep., Los Alamos Scientific Lab., N. Mex.(USA) (1973).

[32] J. E. Morel, G. R. Montry, Analysis and elimination of the discrete-ordinates flux dip, Transport Theory and Statistical Physics 13 (5) (1984) 615–633. doi:10.1080/00411458408211661.

[33] Y. Wang, J. C. Ragusa, A high-order discontinuous galerkin method for the sn transport equations on 2d unstructured triangular meshes, Annals of Nuclear Energy 36 (7) (2009) 931–939. doi:https://doi.org/10.1016/j.anucene.2009.03.002.

[34] M. Hanuš, L. H. Harbour, J. C. Ragusa, M. P. Adams, M. L. Adams, Uncollided flux techniques for arbitrary finite element meshes, Journal of Computational Physics 398 (2019) 108848. doi:https://doi.org/10.1016/j.jcp.2019.07.046.

[35] A. Buchan, A. Calloo, M. Goffin, S. Dargaville, F. Fang, C. Pain, I. Navon, A pod reduced order model for resolving angular direction in neutron/photon transport problems, Journal of Computational Physics 296 (2015) 138–157. doi:https://doi.org/10.1016/j.jcp.2015.04.043.

[36] R. Pinnau, Model Reduction via Proper Orthogonal Decomposition, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 95–109. doi:10.1007/978-3-540-78841-6_5.

[37] M. P. Adams, M. L. Adams, W. D. Hawkins, T. Smith, L. Rauchwerger, N. M. Amato, T. S. Bailey, R. D. Falgout, A. Kunen, P. Brown, Provably optimal parallel transport sweeps on semi-structured grids, Journal of Computational Physics 407 (2020) 109234. doi:https://doi.org/10.1016/j.jcp.2020.109234.

[38] J. I. Vermaak, J. C. Ragusa, M. L. Adams, J. E. Morel, Massively parallel transport sweeps on meshes with cyclic dependencies, Journal of Computational Physics 425 (2021) 109892. doi:https://doi.org/10.1016/j.jcp.2020.109892.

[39] M. Hanus, J. C. Ragusa, Thermal upscattering acceleration schemes for parallel transport sweeps, Nuclear Science and Engineering 194 (10) (2020) 873–893. doi:10.1080/00295639.2020.1767436.

[40] R. McConn, C. Gesh, R. Pagh, R. Rucker, R. Williams, Compendium of material composition data for radiation transport modeling, Tech. Rep. PNNL-15870 Revision 1, Pacific Northwest National Laboratory, Richland, WA (2011).