A Discontinuous and Adaptive Reduced Order Model for the Angular Discretisation of the Boltzmann Transport Equation

Alexander C. Hughes, Andrew G. Buchan

4th August 2020

Abstract

This article presents a new adaptive reduced order model for resolving the angular direction of the Boltzmann transport equation, based on proper orthogonal decomposition (POD) and the method of snapshots. It builds upon previous methods of applying POD to the angular dimension, with modifications to increase accuracy and solver stability. Previous methods used continuous global functions spanning the whole sphere. The new approach, discontinuous POD (DPOD), partitions the surface of the sphere into angular regions, each with an independent set of POD basis functions. Combined, these can approximate flux distributions which span the sphere using optimised basis functions for each angular region. In addition, a novel implementation of adaptive angular resolution known as adaptive discontinuous POD (ADPOD) is presented, which allows the number of DPOD basis functions to vary by angular octant and spatial element. DPOD and ADPOD are applied to two problems in order to demonstrate their benefits compared to POD. Both are shown to reduce the number of solver iterations required to find a solution and decrease the error in the angular flux.

1 Introduction

The Boltzmann Transport Equation (BTE) is a linear partial differential equation which describes particle transport statistically. Finding numerical solutions to the BTE is vital for a wide range of applications involving particle transport, including aerodynamics [1], radiative transport [2] and solid state physics [3]. As the BTE is not analytically solvable in general, numerical methods are used. This is often computationally expensive, as the problem may be up to 7-dimensional - 3 space, 2 angle, 1 time and 1 energy [4]. Each additional dimension multiplies the degrees of freedom of the problem, increasing both the computation time and memory requirements. The challenge of solving the BTE for complex systems has been one of the motivations for investigation into Reduced Order Modelling (ROM).

Reduced Order Models (ROMs) have rapidly gained popularity in recent years, as they allow solutions from a conventional model to be used in the creation of a new model with comparable accuracy but reduced computational complexity. ROMs have been applied to a vast range of problems, from signal analysis and pattern recognition [5], to statistics [6], and geophysical fluids [7]. They have been particularly successful in simulating fluid flows following the early work of Lumley [8], and have been applied to turbulent pipe flows [9], wakes behind a cylinder [10], flows across air foils [11], the mixing of fluid layers [12], thermal currents [13, 14] and ocean models [15]. Their success lies in the construction of optimal reduced basis sets that can resolve a class of problems using very few functions. This enables the discrete problem to be formed with low dimensionality, typically of the order of a hundred or fewer. These systems are significantly smaller, and thus more efficient to solve, than high dimensional models formed from discretisations such as the finite element [16], finite difference [17] and finite volume methods [18].

There are several techniques for constructing reduced basis sets, which can be categorised as either *a priori* or *a posteriori* methods. *A priori* methods such as Proper Generalised Decomposition aim to generate reduced bases without solutions to a full model [19], whereas *a posteriori* methods such as the Empirical Interpolation Method [20, 21], neural networks [22] and Proper Orthogonal Decomposition (POD) use existing solutions to produce a reduced basis for use in further simulation. This work will focus on the method of POD, first proposed by Pearson [23], in conjunction with the method of snapshots [24]. This approach uses a selection of snapshots, which are solutions obtained from a full order model such as a high resolution finite element model, to form optimal basis functions for a particular class of problems.

In the field of nuclear engineering, ROMs have begun to make progress towards providing efficient and accurate predictions for both reactor physics and shielding problems. Early work in reactor physics includes [25], which applied POD to resolve the spatial dependence of general eigenvalue problems, and [26], which used it to perform transient analysis of accelerator driven systems. A comparison of POD to modal methods in transient analysis is given in [27], and the use of ROMs to resolve sub-channels in lead cooled fast reactors is developed in [28]. POD has also been applied to resolve the angular dependence of the neutron transport equation [29], which was followed by similar work to resolve the angular dimension in radiative heat transfer problems [30]. More recent work has applied POD to non-linear feedback effects within lead-cooled reactors [31], fuel burnup within benchmark systems [32], molten salt reactor analysis [33], and natural thermohydraulic circulation [34].

This article presents a new discretisation method using a POD based ROM to resolve the angular dimension of the Boltzmann transport equation. This work is based on [29] which also developed an angular ROM for the BTE, but the new method presented here provides improvements in both accuracy and solver stability. Following initial publication of [29], further investigation found that the POD discretisation could result in solver instabilities, whereby the total solver iterations required for solution convergence oscillated significantly with varying POD expansion sizes. This issue seems inherent with the POD basis functions themselves, as similar traits have been observed with two independent codes using different spatial FEM formulations and solver technologies. That is, in [29] a stabilized continuous FEM with a GMRES solver was used, and here a discontinuous FEM with a sweep based solver is employed. A likely contributing factor to these instabilities is the use of global

POD basis functions that direct particles through all angles on the sphere. Apart from this placing significant strain on the basis functions to reconstruct solutions over wide angles, other such functions which are not based on spherical harmonics have been shown in previous studies to cause other numerical issues [35]. In [35] a remedy was to break the global span of the functions and instead use localised functions covering the octants of the sphere. Thus, a new approach has been formulated here that allows the angular flux to be resolved using optimised basis functions over distinct angular regions. These regions can be defined arbitrarily, but in this case each angular octant delimits a partition, which is shown to be effective in stabilizing the method.

This article also presents an adaptive angular algorithm enabling each spatial node to be resolved independently using the optimal number of angular basis functions. Whilst other adaptive angular algorithms exist [36, 37, 38], this is the first time such a method has been developed with POD basis sets. Thus, the novelty of this article is in the development of a new angular reduced order model with adaptive capabilities. The model can provide optimised angular resolution using new adaptive POD basis sets, which improve solver scaling and solution accuracy while reducing the discrete problem size.

In generating the POD angular functions, snapshots of the angular flux vectors are taken through space, rather than through time as in conventional POD approaches. This has similarities to [39, 40], which solve the two dimensional time-independent parabolized Navier-Stokes equation using a spatial dimension as though it were time.

This article is set out as follows. Section 2 explains the method used to produce the POD model. It describes the Boltzmann transport equation, the full order discretisation of the angular and spatial dimensions, then the process of recording snapshot data and forming the new reduced order model. In section 3, two numerical examples are presented which demonstrate the ability of the ROM to provide accurate and efficient solutions for both reactor-type and shielding-type problems. Section 4 completes this article with a conclusion of its findings, and suggests future work which may produce further improvements. Section 5 gives acknowledgements, section 6 provides a bibliography, and tables are included in section 7 at the end of the article.

2 Discontinuous and Adaptive Proper Orthogonal Decomposition Methods

2.1 The Boltzmann Transport Equation

The steady state, mono-energetic Boltzmann transport equation describes the angular flux $\psi(\vec{r}, \Omega)$ in direction Ω at position $\vec{r}$:

$$\Omega \cdot \nabla \psi(\vec{r}, \Omega) + \Sigma_t(\vec{r}) \psi(\vec{r}, \Omega) = q_{ex}(\vec{r}) + q_s(\vec{r}, \Omega' \rightarrow \Omega) \tag{1}$$

Where q$_{ex}$ denotes the isotropic external source, and the scattering source term q$_s$ is given by:

$$q_s(\vec{r},\Omega' \to \Omega) = \int_{\Omega'} \Sigma_s(\vec{r},\Omega' \to \Omega)\psi(\vec{r},\Omega')d\Omega' \tag{2}$$

The macroscopic scattering, absorption and total removal cross sections are given by Σ_s, Σ_a and Σ_t respectively, and are related by:

$$\Sigma_t = \Sigma_s + \Sigma_a \tag{3}$$

The first term in equation 1 is the advection operator $\Omega\cdot\nabla$, which can be written in Cartesian coordinates as:

$$\Omega \cdot \nabla = (\Omega_x, \Omega_y, \Omega_z) \cdot (\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z}) = \Omega_x \frac{\partial}{\partial x} + \Omega_y \frac{\partial}{\partial y} + \Omega_z \frac{\partial}{\partial z} \tag{4}$$

Where Ω_x, Ω_y and Ω_z are the components of the unit vector Ω in Cartesian space.

2.2 Angular Discretisation of the Boltzmann Transport Equation

To discretise the angular dimension of equation 1, the angular flux $\psi(\vec{r},\Omega)$ is approximated by a sum of n_a angular basis functions $\mathcal{G}_j(\Omega)$ multiplied by the coefficients $\Phi_j(\vec{r})$:

$$\psi(\vec{r},\Omega) \approx \sum_{j=1}^{n_a} \mathcal{G}_j(\Omega)\Phi_j(\vec{r}) \tag{5}$$

The approximation in equation 5 is inserted into equation 1, which is then weighted and integrated over all angles. The Bubnov-Galerkin method is applied, which uses the angular basis functions $\mathcal{G}_i$ as weights. The angularly discretised form of the BTE can therefore be written as:

$$\begin{aligned}\sum_{j=1}^{n_a} \Big(\int_{\Omega} \mathcal{G}_i(\Omega)\Omega \cdot \nabla \mathcal{G}_j(\Omega) d\Omega \; \Phi_j(\vec{r}) + \int_{\Omega} \mathcal{G}_i(\Omega)\Sigma_t(\vec{r})\mathcal{G}_j(\Omega)d\Omega \; \Phi_j(\vec{r}) \Big) \\ - \int_{\Omega} \mathcal{G}_i(\Omega) q_s(\vec{r},\Omega' \to \Omega)d\Omega = \int_{\Omega} \mathcal{G}_i(\Omega) q_{ex}(\vec{r})d\Omega \\ \forall \; i \in \{1, n_a\}\end{aligned} \tag{6}$$

The Cartesian components of the advection operator $\Omega \cdot \nabla$, given by equation 4, are substituted into equation 6:

$$\begin{aligned}\sum_{j=1}^{n_a} \Big(\int_{\Omega} \mathcal{G}_i(\Omega)\Omega_x \frac{\partial}{\partial x} \mathcal{G}_j(\Omega) d\Omega \; \Phi_j(\vec{r}) + \int_{\Omega} \mathcal{G}_i(\Omega)\Omega_y \frac{\partial}{\partial y} \mathcal{G}_j(\Omega) d\Omega \; \Phi_j(\vec{r}) \\ + \int_{\Omega} \mathcal{G}_i(\Omega)\Omega_z \frac{\partial}{\partial z} \mathcal{G}_j(\Omega) d\Omega \; \Phi_j(\vec{r}) + \int_{\Omega} \mathcal{G}_i(\Omega)\Sigma_t(\vec{r})\mathcal{G}_j(\Omega)d\Omega \; \Phi_j(\vec{r}) \Big) \\ - \int_{\Omega} \mathcal{G}_i(\Omega) q_s(\vec{r},\Omega' \to \Omega)d\Omega = \int_{\Omega} \mathcal{G}_i(\Omega) q_{ex}(\vec{r})d\Omega \\ \forall \; i \in \{1, n_a\}\end{aligned} \tag{7}$$

Equation 7 can be written in matrix form as:

$$A_x \frac{\partial \Phi(\vec{r})}{\partial x} + A_y \frac{\partial \Phi(\vec{r})}{\partial y} + A_z \frac{\partial \Phi(\vec{r})}{\partial z} + H(\vec{r})\Phi(\vec{r}) = Q(\vec{r}) \tag{8}$$

Where A_x, A_y, A_z and H are matrices of size $n_a \times n_a$. $\Phi(\vec{r})$ and $Q(\vec{r})$ are vectors of size n_a, containing the coefficients of the angular expansion and the contribution of the source term, respectively. The components of each matrix and vector at row i, column j are given by:

$$\begin{aligned}
A_{xij} &= \int_{\Omega} \mathcal{G}_i \Omega_x \mathcal{G}_j d\Omega \\
A_{yij} &= \int_{\Omega} \mathcal{G}_i \Omega_y \mathcal{G}_j d\Omega \\
A_{zij} &= \int_{\Omega} \mathcal{G}_i \Omega_z \mathcal{G}_j d\Omega \\
H_{ij} &= \int_{\Omega} \mathcal{G}_i \Sigma_t(\vec{r}) \mathcal{G}_j d\Omega - \int_{\Omega} \mathcal{G}_i q_s(\vec{r}, \Omega' \to \Omega) d\Omega \\
Q_i &= \int_{\Omega} \mathcal{G}_i q_{ex}(\vec{r}) d\Omega
\end{aligned}$$

A vector of matrices A is defined as $A = (A_x, A_y, A_z)$, which allows equation 8 to be written as:

$$(A \cdot \nabla + H(\vec{r}))\Phi(\vec{r}) = Q(\vec{r}) \tag{9}$$

2.3 Spatial Discretisation of the Boltzmann Transport Equation

The discontinuous finite element method is applied to discretise the spatial dimensions of equation 9 [41]. The equation is converted to its weak form by weighting it with a set of n_s basis functions $\mathcal{N}_i(\vec{r})$ and integrating over the volume of each element, V_e:

$$\Big(\int_{V_e} \mathcal{N}_i(\vec{r}) A \cdot \nabla dV_e \; + \int_{V_e} \mathcal{N}_i(\vec{r}) H(\vec{r}) dV_e \Big) \Phi(\vec{r}) = \int_{V_e} \mathcal{N}_i(\vec{r}) Q(\vec{r}) dV_e \qquad \forall \; i \in \{1, n_s\} \tag{10}$$

Green's Theorem is applied to the advection term in equation 10, splitting it into an integral over V_e and another over the element boundary Γ_e:

$$\begin{aligned}
\Big(\int_{\Gamma_e} \mathcal{N}_i(\vec{r})(A \cdot \hat{n}) d\Gamma_e - \int_{V_e} \nabla \mathcal{N}_i(\vec{r}) \cdot A dV_e + \int_{V_e} \mathcal{N}_i(\vec{r}) H(\vec{r}) dV_e \Big) \Phi(\vec{r}) \\
= \int_{V_e} \mathcal{N}_i(\vec{r}) Q(\vec{r}) dV_e \\
\forall \; i \in \{1, n_s\}
\end{aligned} \tag{11}$$

$\hat{n}$ is the unit vector normal to the element boundary. The angularly discretised flux $\Phi(\vec{r})$ is approximated as a sum of the spatial basis functions $\mathcal{N}_j(\vec{r})$ multiplied by the coefficients Ψ_j:

$$\Phi(\vec{r}) \approx \sum_{j=1}^{n_s} \Psi_j \mathcal{N}_j(\vec{r}) \tag{12}$$

The Bubnov-Galerkin method is applied, which uses the same set of weights as basis functions. Inserting equation 12 into equation 11 gives:

$$\sum_{j=1}^{n_s}\Big(\int_{\Gamma_e}\mathcal{N}_i(\vec{r})(A\cdot\hat{n})\mathcal{N}_j(\vec{r})d\Gamma_e-\int_{V_e}\nabla\mathcal{N}_i(\vec{r})\cdot A\mathcal{N}_j(\vec{r})dV_e+\int_{V_e}\mathcal{N}_i(\vec{r})H(\vec{r})\mathcal{N}_j(\vec{r})dV_e\Big)\Psi_j=\int_{V_e}\mathcal{N}_i(\vec{r})Q(\vec{r})dV_e \quad \forall\ i\in\{1,n_s\} \tag{13}$$

First order upwinding will be used to calculate the flow at the element boundaries [42]. The boundary term must therefore be split into inflow and outflow components [43]:

$$\sum_{j=1}^{n_s}\int_{\Gamma_e}\mathcal{N}_i(\vec{r})(A\cdot\hat{n})\mathcal{N}_j(\vec{r})d\Gamma_e\Psi_j =\sum_{j=1}^{n_s}\Big(\int_{\Gamma_e}\mathcal{N}_i(\vec{r})(A^{in}\cdot\hat{n})\mathcal{N}_j(\vec{r})d\Gamma_e\Psi_j^{in}+\int_{\Gamma_e}\mathcal{N}_i(\vec{r})(A^{out}\cdot\hat{n})\mathcal{N}_j(\vec{r})d\Gamma_e\Psi_j^{out}\Big) \quad \forall\ i\in\{1,n_s\} \tag{14}$$

The outflow Ψ_j^{out} is the angular flux vector of the element in question, and the inflow Ψ_j^{in} is the angular flux vector of its upstream neighbour. Equation 14 is inserted into equation 13 to give the final form of the discretised equations:

$$\sum_{j=1}^{n_s}\Bigg\{\Big(-\int_{V_e}\nabla\mathcal{N}_i(\vec{r})A\mathcal{N}_j(\vec{r})dV+\int_{V_e}\mathcal{N}_i(\vec{r})H(\vec{r})\mathcal{N}_j(\vec{r})dV\Big)\Psi_j +\int_{\Gamma_e}\mathcal{N}_i(\vec{r})(A^{in}\cdot\hat{n})\mathcal{N}_j(\vec{r})d\Gamma_e\Psi_j^{in}+\int_{\Gamma_e}\mathcal{N}_i(\vec{r})(A^{out}\cdot\hat{n})\mathcal{N}_j(\vec{r})d\Gamma_e\Psi_j^{out}\Bigg\} =\int_{V_e}\mathcal{N}_i(\vec{r})Q(\vec{r})dV \quad \forall\ i\in\{1,n_s\} \tag{15}$$

2.4 The Discontinuous Angular Reduced Order Model

This section describes the reduced order discretisation of the angular dimension of the BTE. The ROM is based on the use of Proper Orthogonal Decomposion and the method of snapshots to form optimal angular basis functions for a particular class of problem. Typical POD methods use snapshots of spatial coefficients as they vary through time, but here angular flux profiles are used as they vary through space. This idea was first proposed in [29], which used vectors of angular coefficients to form the snapshot matrix, resulting in the formation of POD basis functions that spanned the entire surface of the sphere. The fundamental difference in this work is to relax the global nature of the POD basis functions, and to instead form basis sets that span subsets of the sphere. The selection of these angular regions can be arbitrary, but in this article each octant of the sphere defines a separate region, which is shown to stabilise the basis functions and as such should be considered the starting point for this method. Further subdivision may be beneficial in some contexts, such as highly anisotropic flux distributions, but has not been investigated in this article. The angular domain Ω is partitioned into the regions,

$$\Omega = \bigcup_{i=1}^{8} \Omega_i \tag{16}$$

where Ω_i here defines the i^{th} octant, which will have its own set of optimised basis functions. Figure 1 depicts one such octant.

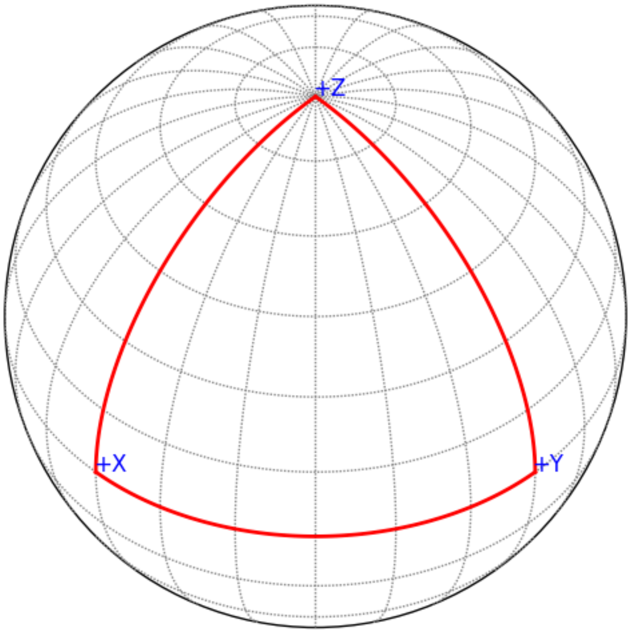

Figure 1: An octant on the sphere, which defines the boundaries of an angular region.

Each class of problem is resolved for n_p cases with varying conditions, such as perturbations to the nuclear material cross-sections. The vectors of angular coefficients formed at the nodes of the spatial mesh are then partitioned into sets according to their angular region. By way of example, if the S_n method is used then the vector of angular coefficients on one particular node can be partitioned according to their direction,

$$\Psi = \begin{bmatrix} \Psi_1 \\ \Psi_2 \\ \vdots \\ \Psi_8 \end{bmatrix} \tag{17}$$

where Ψ_i is a vector containing the n_i coefficients with associated directions within Ω_i. Note that $n_a = \sum^{i} n_i$. This partitioning of the angular vector is used to form separate snapshot matrices for each Ω_i. When using the octants of the sphere as angular regions, 8 separate snapshot matrices are formed from the vectors of angular coefficients at the mesh nodes, for all problems used in creation of the ROM. The snapshots generated are denoted by

$$\mathrm{S}_i = \begin{bmatrix} | & | & & | \\ \Psi_i^1 & \Psi_i^2 & ... & \Psi_i^{n_h} \\ | & | & & | \end{bmatrix}, \ \forall i \in \{1, 8\} \tag{18}$$

which is the concatenation of n_h angular vectors from all mesh nodes and problems resolved - given by $n_h = n_s \times n_p$. The POD basis sets for each angular region Ω_i can now be formed through the SVD of each snapshot matrix,

$$\mathrm{S}_i = \mathrm{U}_i \Sigma_i \mathrm{V}_i^T \tag{19}$$

where U_i and V_i are unitary matrices of sizes $n_a \times n_a$ and $n_h \times n_h$, respectively. The column vectors of U_i contain the optimised basis vectors that best represent the snapshot data, ordered such that the first n_f columns form the optimal n_f basis vectors in the Frobenius norm. The angular POD basis matrices $\mathcal{U}_i$ are formed by truncating each snapshot matrix such that only the first n_f columns are retained:

$$\mathcal{U}_i^{jk} = \mathrm{U}_i^{jk} \qquad\qquad \forall i \in \{1, 8\}, \forall j \in \{1, n_f\}, \forall k \in \{1, n_a\} \tag{20}$$

Where $\mathcal{U}_i^{jk}$ and U_i^{jk} denote the j^{th} row and k^{th} column of the matrix associated with the angular region i. The fraction of the information in U_i which is retained in $\mathcal{U}_i$ can be determined from the singular values:

$$I_i = \frac{\sum\limits_{j=1}^{n_f} (\Sigma_i^{jj})^2}{\sum\limits_{j=1}^{n_a} (\Sigma_i^{jj})^2} \tag{21}$$

where I_i varies from 0 to 1, with 1 being total capture of the snapshot information. The matrices $\mathcal{U}_i$ can be used to map the angular coefficients between the full and reduced order models through the relationship,

$$\Psi_i(\vec{r}) \approx \mathcal{U}_i \alpha_i(\vec{r}) \tag{22}$$

for each angular partition Ω_i. The full POD expansion can be formed through the summation of POD functions over each angular partition,

$$\Psi(\vec{r}) = \sum_{i=1}^{8} \sum_{j=1}^{n_f} \mathcal{U}_i^j \alpha_i^j(\vec{r}) \tag{23}$$

Where α_i^j denotes the j^{th} POD coefficient in angular region i. Note that there is freedom to set the expansion size of each angular region independently. The combined mapping over all angular partitions can be expressed as,

$$\mathcal{U} = \begin{bmatrix} \mathcal{U}_1 & 0 & 0 & 0 & 0 \\ 0 & \ddots & 0 & 0 & 0 \\ 0 & 0 & \mathcal{U}_i & 0 & 0 \\ 0 & 0 & 0 & \ddots & 0 \\ 0 & 0 & 0 & 0 & \mathcal{U}_8 \end{bmatrix} \quad \alpha = \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_i \\ \vdots \\ \alpha_8 \end{bmatrix} \quad \Longrightarrow \mathcal{U}\alpha = \begin{bmatrix} \mathcal{U}_1\alpha_1 \\ \vdots \\ \mathcal{U}_i\alpha_i \\ \vdots \\ \mathcal{U}_8\alpha_8 \end{bmatrix}$$

Equations 17 and 22 enable this to be compactly written as,

$$\mathcal{U}\alpha(\vec{r}) \approx \Psi(\vec{r}) \tag{24}$$

Substituting equation 24 into equation 8 and premultiplying by $\mathcal{U}^T$ projects the angular discretised equations onto the POD space:

$$\Big(\mathcal{U}^T \mathbf{A}_x \mathcal{U} \frac{\partial}{\partial x} + \mathcal{U}^T \mathbf{A}_y \mathcal{U} \frac{\partial}{\partial y} + \mathcal{U}^T \mathbf{A}_z \mathcal{U} \frac{\partial}{\partial z} + \mathcal{U}^T \mathbf{H} \mathcal{U}\Big)\alpha(\vec{r}) = \mathcal{U}^T \mathbf{Q} \tag{25}$$

Spatial discretisation can then be performed on equation 25 as described in section 2.3. Note that in the case of a single angular region spanning the full sphere, the angular POD method implemented by Buchan et al. [29], upon which this work is based, is recovered.

2.5 Adaptive Discontinuous Proper Orthogonal Decomposition in Angle

This section presents an adaptive algorithm using the DPOD basis functions. Instead of using the same number of basis functions n_f everywhere, each octant i and element j has an associated number of basis functions n_{ij}, which can be modified independently. The adaptive algorithm seeks to determine where to add basis functions in order to maximise their effectiveness in reducing total error. An initial solution is required in order to calculate the error metric and begin to add basis functions, and so the problem in question is first solved with:

$$n_{ij} = 2,\ \forall i,\ \forall j \tag{26}$$

The DPOD basis functions form a hierarchical set, and so their coefficients tend towards zero as the approximation converges to the exact solution. It is therefore possible to determine which locations would benefit most from additional basis functions by comparing their final coefficients. If the final coefficient in a particular location is large, it is likely that an additional basis function would also have a large coefficient, and therefore significantly impact the solution. However, examination of solution vectors at various points showed that the magnitude of successive coefficients oscillates, and so a single small coefficient does not guarantee convergence. As a result, the final coefficient alone was not a reliable metric. To account for this, the larger of the final two coefficients at each location is compared. A list of these coefficients L_{ij} is compiled:

$$\mathrm{L}_{ij} = \mathrm{Max}(\alpha_{ij}^{n_{ij}-1}, \alpha_{ij}^{n_{ij}}),\ \forall i \in \{1, 8\},\ \forall j \in \{1, n_s\} \tag{27}$$

Where α_{ij}^k denotes the k^{th} angular coefficient for octant i at node j. The list L_{ij} is sorted by magnitude, and the largest n_+ entries have their associated n_{ij} values increased, while the smallest n_- have their n_{ij} values decreased (to a minimum of two). Once n_{ij} has been adjusted, the process repeats until the problem has converged to within the desired tolerance.

3 Numerical Examples

In this section, two numerical examples are presented in order to compare POD, DPOD and ADPOD. Uniform quadrilateral FEM spatial meshes are employed, using discontinuous linear basis functions. The full order method employs the S_n discretisation, with a sufficiently high angular resolution to ensure that the solutions have converged in angle. A sweep based solver is used to resolve the S_n and all POD methods. At present, this solver has not been optimised, and so computational efficiency is determined based on the number of solver iterations required for solution convergence, and the error for a given number of angular degrees of freedom. It should be noted that for the POD and DPOD methods the complexity, and hence computation time, of each solver iteration is the same for a given number of basis functions. Furthermore, the sparsity within the DPOD angular streaming matrices (75% for 2D and 87% for 3D) has not been considered, and so DPOD could potentially be implemented more efficiently than POD. This will be reviewed in future works, but the purpose of this analysis is to demonstrate the novel methods' improved solver convergence and increased accuracy for a given basis size when compared with S_n and previous angular POD methods.

3.1 The Dog-Leg Duct Problem

The first example is a dog-leg duct problem [44]. Figure 2a shows a schematic of the domain. Region 1 is a source, region 2 is the duct, and region 3 is a heavy absorber. Vacuum boundary conditions are applied to the top and right boundaries, and reflective boundary conditions to the bottom and left boundaries. The spatial dimension is discretised with a 140×180 mesh of discontinuous linear quadrilateral elements. The full order solutions used for snapshots and error calculations employed the S_{40} angular discretisation. Figure 2b shows the scalar flux distribution of the S_{40} solution to the extrapolation problem.

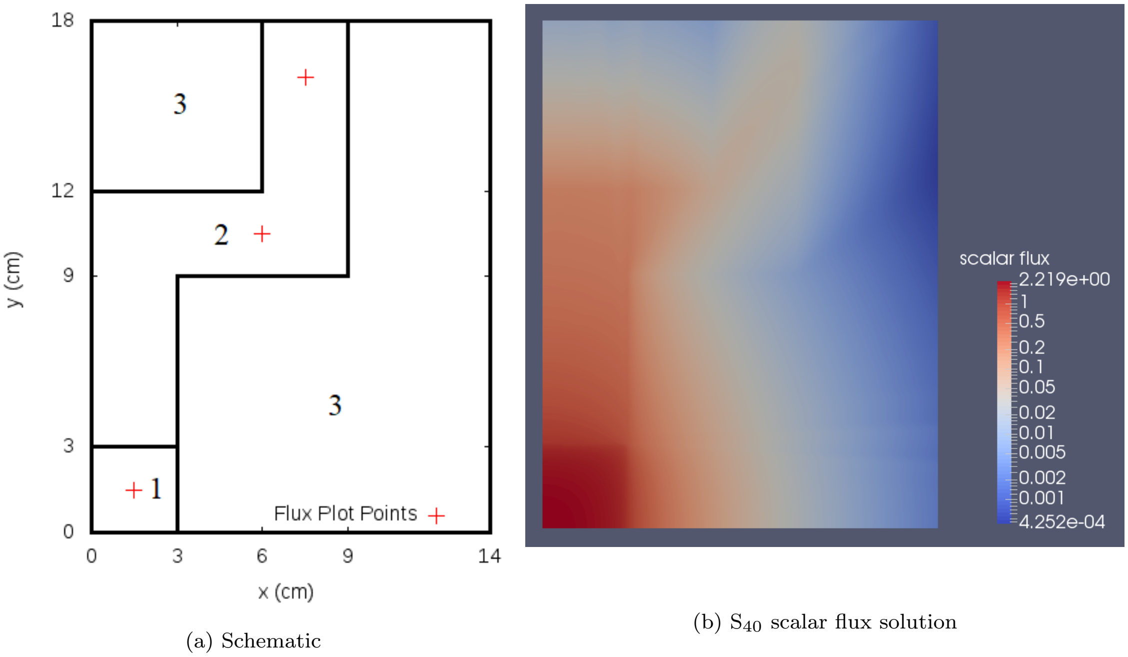

(a) Schematic

(b) S_{40} scalar flux solution

Figure 2: Schematic (a) and S_{40} scalar flux solution (b) for the dog-leg duct extrapolation problem. Region 1 is the source, region 2 is the void duct and region 3 is a highly absorbing material.

Tables 1 and 2 list the material cross sections for the training and test problems, respectively. The snapshot matrix was formed from all three training solutions, and the resulting

POD bases were used to solve the test problems. The first test problem is referred to as the *seen* problem, as its material properties are identical to one of the training solutions. The other two are both *unseen* problems. The *interpolation* problem has its Σ_a value within the range for which snapshots were produced, while Σ_a for the *extrapolation* problem lies outside this range.

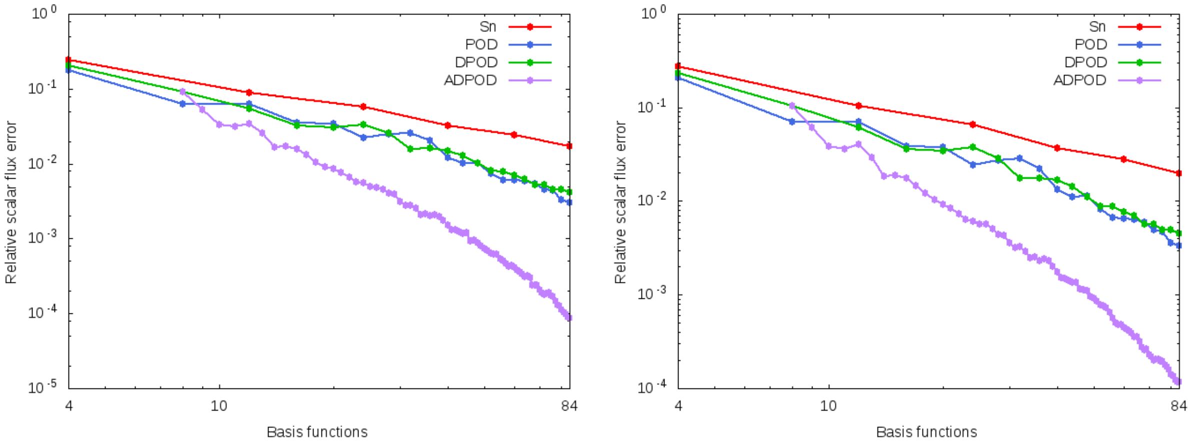

(a) Seen (b) Interpolation

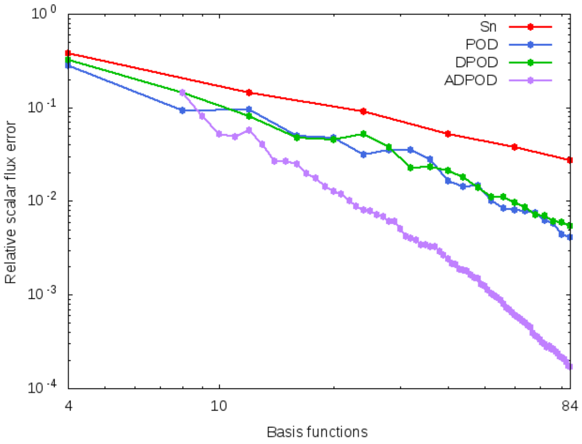

(c) Extrapolation

Figure 3: The effects of increasing basis function count on relative scalar flux error for the dog-leg duct problem.

Figure 3 presents the L1-norm of the relative scalar flux error in solutions of the dog-leg duct problem against the mean number of basis functions per node. Low order S_n solutions with degrees of freedom equivalent to the ROMs' are also shown, in order to provide a comparison between the ROMs and a full order model of equivalent computational complexity. POD consistently has lower error than S_n with the same number of basis functions. DPOD has a similar level of error to POD for this problem. ADPOD is a huge improvement upon both, particularly for larger numbers of basis functions, where it has an error almost two orders of magnitude smaller than POD and DPOD.

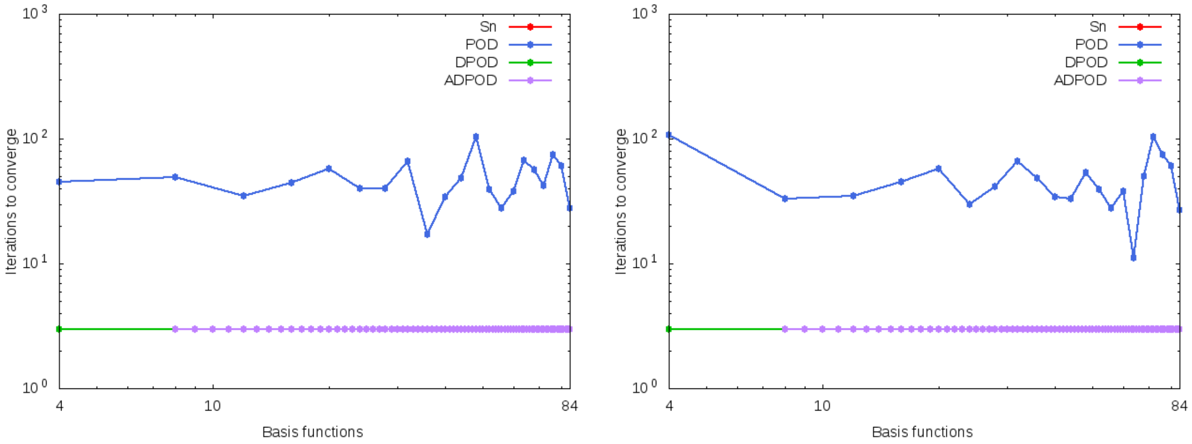

(a) Seen

(b) Interpolation

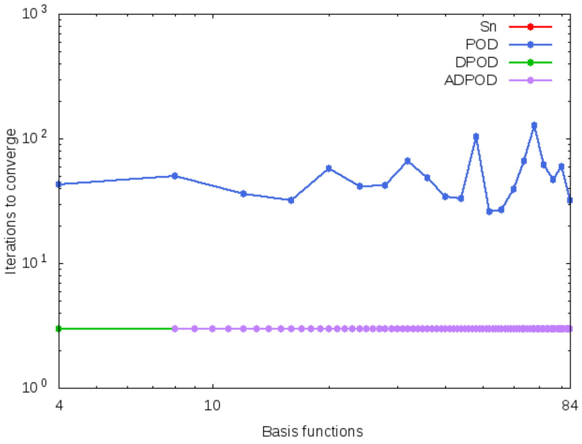

(c) Extrapolation

Figure 4: The effects of increasing basis function count on the number of iterations required to converge to a solution for the dog-leg duct problem.

Figure 4 shows the number of solver iterations required to converge the solution for each method. A ROM solution was considered converged when the L1-norm of the difference between current and previous iterations' scalar flux solutions decreased below 10^{-6}. In the case of ADPOD, the number of iterations is presented for each adaptive step, in order to demonstrate the stability of the ADPOD bases throughout the adaptive process. Due to the non-scattering media, the S_n discretisation required just two iterations to converge to a solution, though a third was performed to verify its convergence. As the POD basis functions send flux in every direction rather than just the direction of the current solver sweep, POD required up to 1.8 orders of magnitude more iterations to converge. The number of iterations required is also inconsistent, with a range of approximately an order of magnitude, which causes unpredictable variation in the solve time. By contrast, DPOD and ADPOD only transport flux in the direction of the current sweep, and so they do not suffer from this disadvantage. Instead, they converge in two sweeps as with S_n. As the solve time is dependant on the number of iterations required to converge the solution, DPOD and ADPOD demonstrate significant benefits over POD.

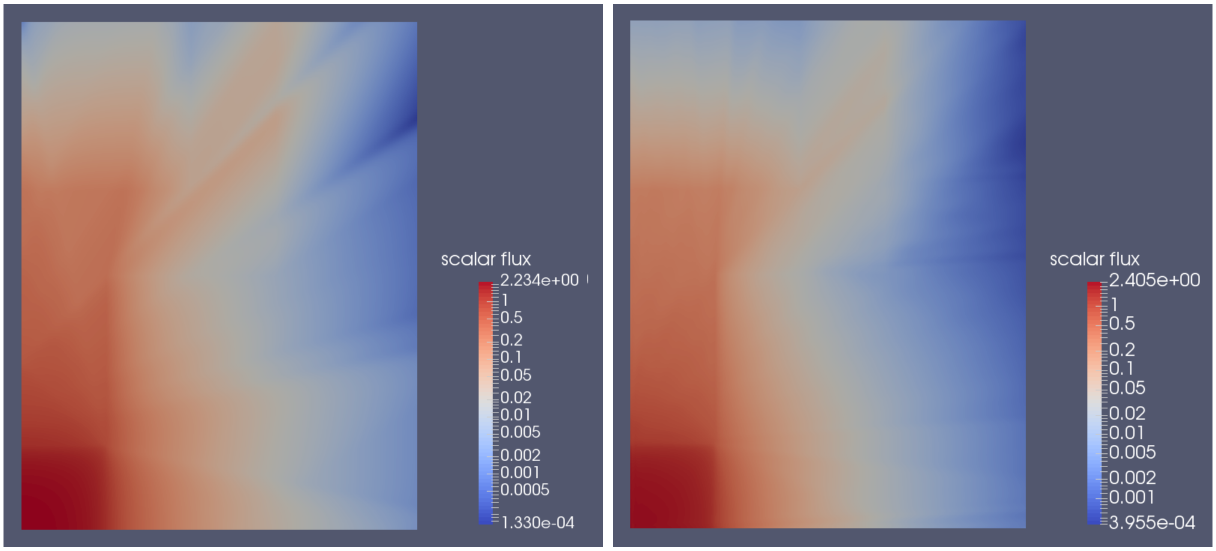

(a) S_8 (b) POD

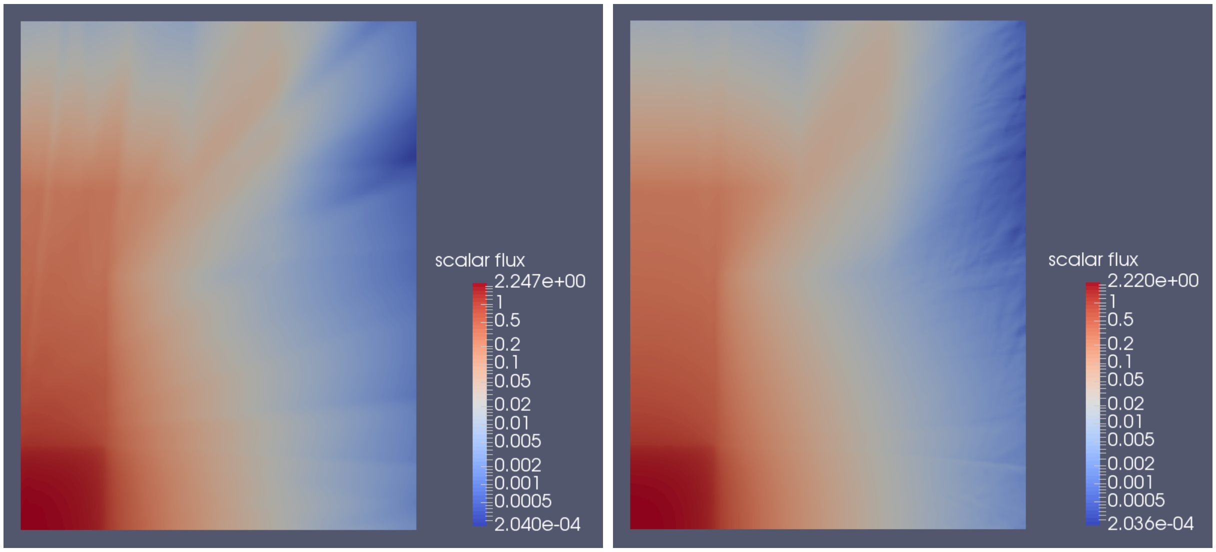

(c) Discontinuous POD (d) Adaptive Discontinuous POD

Figure 5: Solutions to the dog-leg duct extrapolation problem for all four methods with 40 basis functions per node, on average in the case of ADPOD.

Figure 5 depicts the scalar flux solutions for the extrapolation problem produced by each method, using 40 basis functions per node (on average in the case of ADPOD). The full order solution to this problem is shown in figure 2b. The S_8 solution exhibits significant ray effects, which results in a highly inaccurate flux distribution within the duct. POD reduces these ray effects somewhat, though they are still visible. However, it has the most inaccurate peak flux by far. DPOD suffers from significant ray effects in this case, but its peak flux is much closer to the expected value than that of POD. ADPOD drastically reduces ray effects compared to every other method, and has the smallest error in its peak flux. Minor distortions are present, but only in the regions with the least flux, where their effect on the total error is insignificant. This clearly demonstrates the benefits of DPOD and particularly ADPOD compared to previous methods in resolving non-scattering problems with few basis functions.

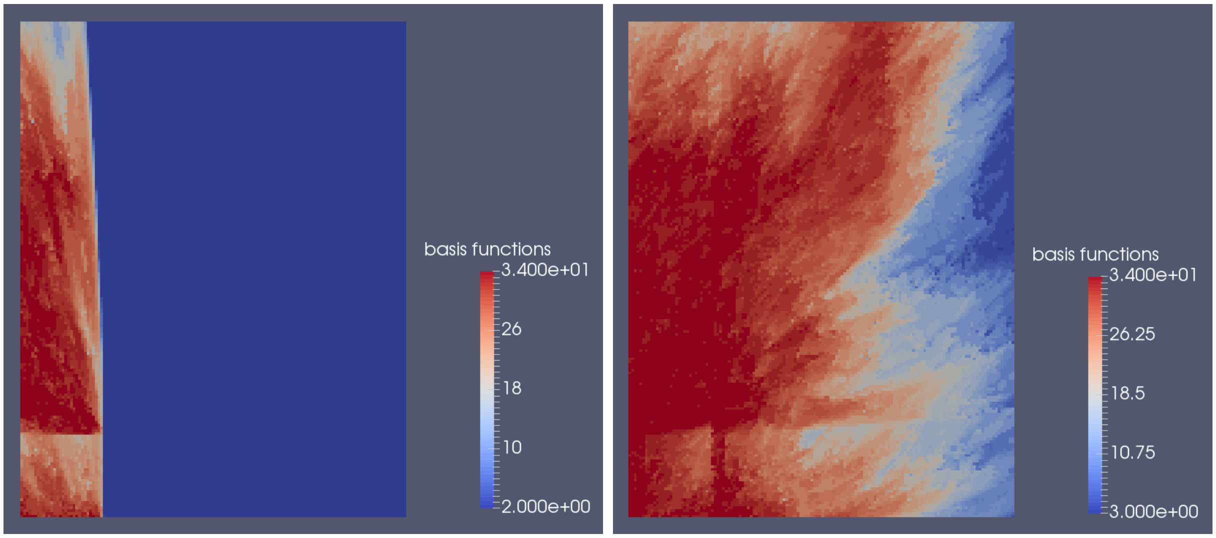

(a) -x, +y octant

(b) +x, +y octant

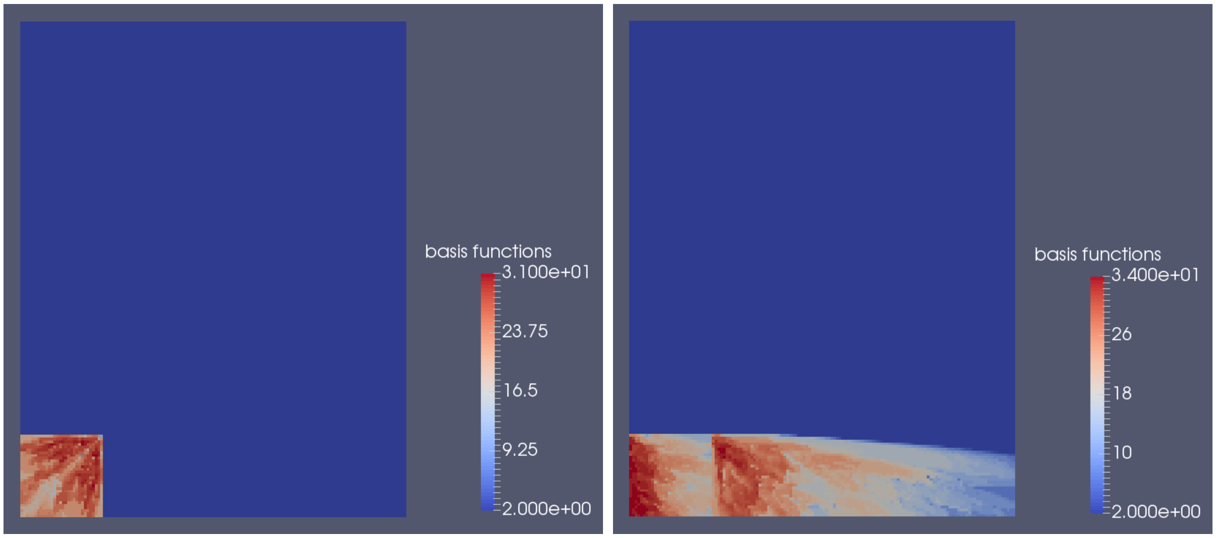

(c) -x, -y octant

(d) +x, -y octant

Figure 6: Number of ADPOD basis functions per octant per node for the duct extrapolation problem, with a mean of 40 basis functions per node in total.

Figure 6 depicts the number of basis functions per node for each element and octant in the domain. A large part of the domain had zero flux in three out of four octants, and the adaptive algorithm successfully accounted for this by keeping the number of basis functions in these regions at the minimum of two. Most basis functions were focused inside the duct, where they contribute most to the overall solution. However, the absorbing regions had a relatively low cross section, and so flux penetrated into these regions rather than being immediately attenuated. The adaptive algorithm therefore added basis functions to these regions as well. ADPOD's ability to focus basis functions where they are most beneficial explains its success in drastically reducing error for a given number of basis functions.

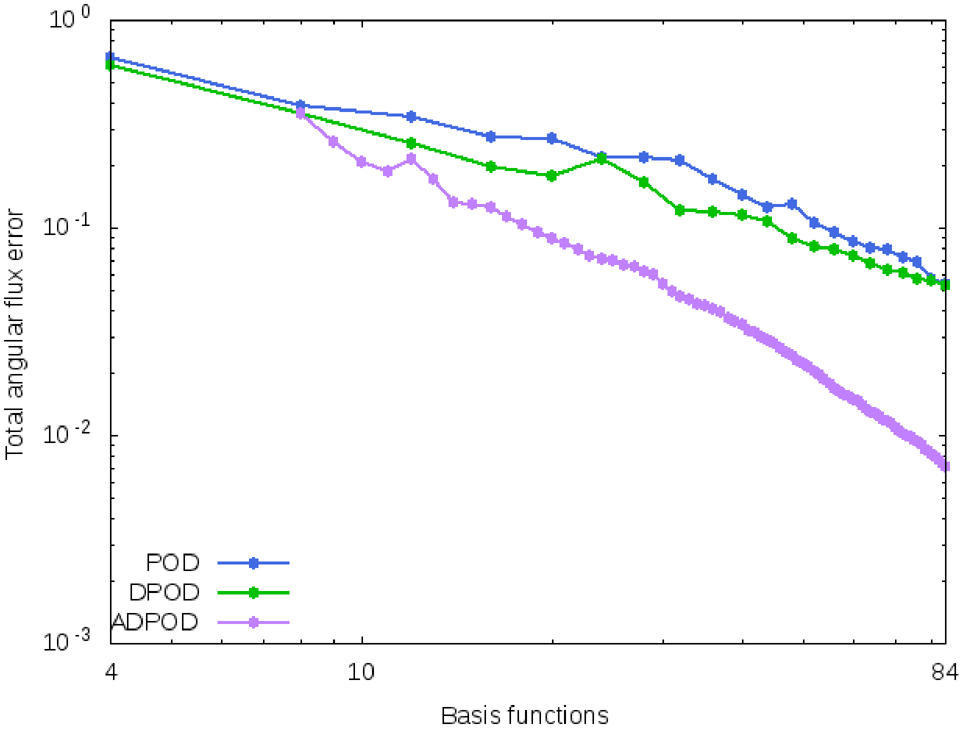

Figure 7: The effect of increasing basis function count on relative angular flux error for the dog-leg duct extrapolation problem.

Figure 7 presents the relative angular flux error against the number of basis functions per node, which is a more accurate error metric than the scalar flux error. The results are similar to the scalar flux error plots, but in this case DPOD is an improvement upon POD in every situation. This suggests that POD and DPOD were comparable in scalar flux error due to cancellation of errors, which does not occur with angular flux. Once again, ADPOD performs significantly better than either method, particularly for large basis function counts.

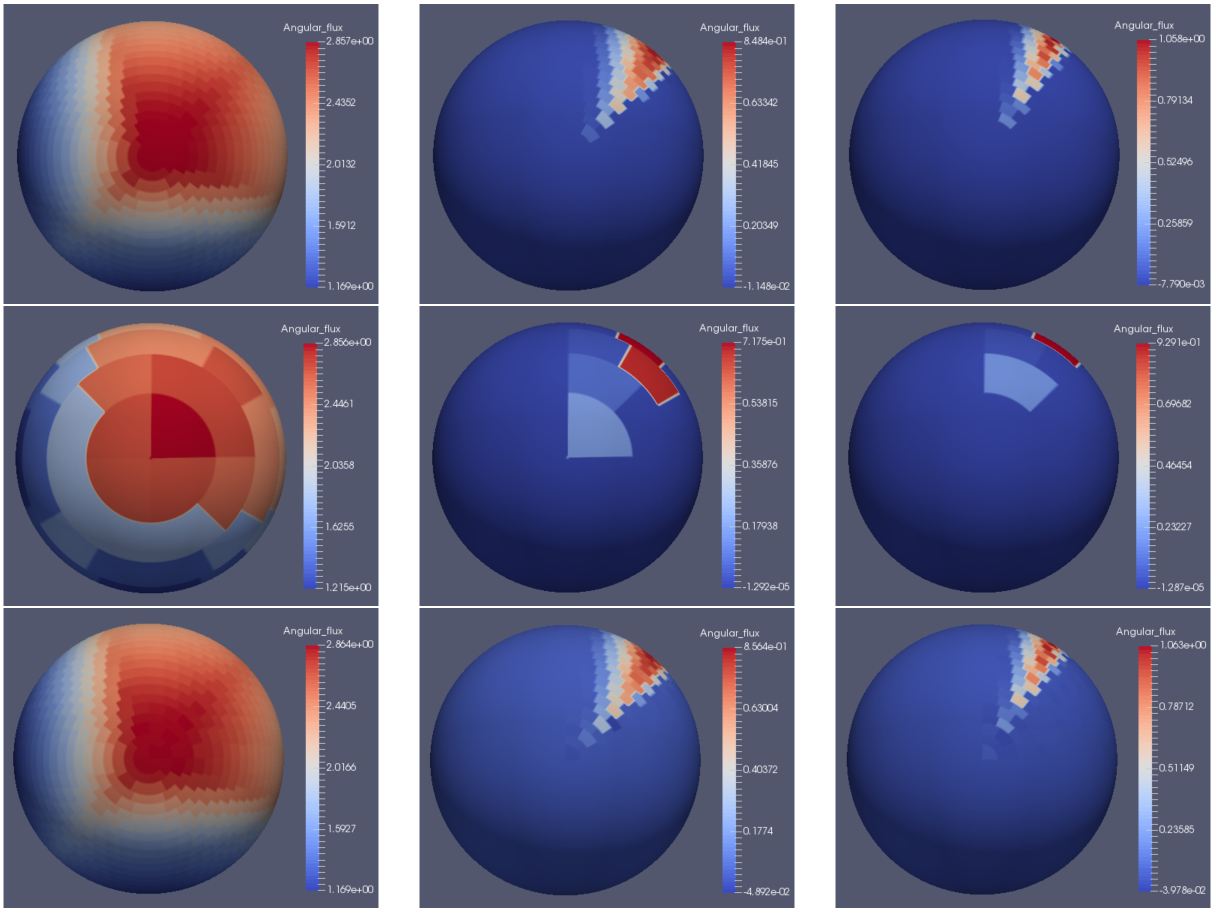

(a) x = 1.5cm, y = 1.5cm (b) x = 6.0cm, y = 10.5cm (c) x = 7.5cm, y = 16.0cm

Figure 8: Angular flux distributions at various points in the dog leg duct problem. From top to bottom, S_{40}, S_8, and adaptive discontinuous POD. Both S_8 and ADPOD have a mean of 40 basis functions per node, while S_{40} has 840 basis functions per node.

Figure 8 depicts the angular flux profiles at three points in the domain. The locations of each point are shown in figure 2a. The S_{40} solution has 840 basis functions, while the S_8 and ADPOD solutions both have 40. The S_8 angular discretisation is highly inaccurate. The solid angles subtended by each basis function are large, and so complex angular flux profiles can not be accurately represented. By contrast, ADPOD is able to reproduce the S_{40} angular flux profiles with only minor differences. The ADPOD solution has a higher angular resolution as its snapshots were produced with S_{40}, and the adaptive method included enough basis functions at these points to reproduce the S_{40} solution with only minor errors. This demonstrates the benefits of ADPOD in approximating higher order angular discretisations.

3.2 The Watanabe-Maynard Problem

The second example is the Watanabe-Maynard problem [45]. A schematic is shown in Figure 9a. Region 1 is an isotropic square source, region 2 is a void, and region 3 is a highly scattering material. Vacuum boundary conditions are applied to the top and right boundaries, and reflective boundary conditions to the bottom and left boundaries. The domain is discretised using an 80×80 mesh of discontinuous linear quadrilateral elements in space. Full order solutions for snapshot generation and error comparison were produced using S_{30} in angle. Figure 9b shows the scalar flux distribution of the S_{30} solution to the extrapolation problem.

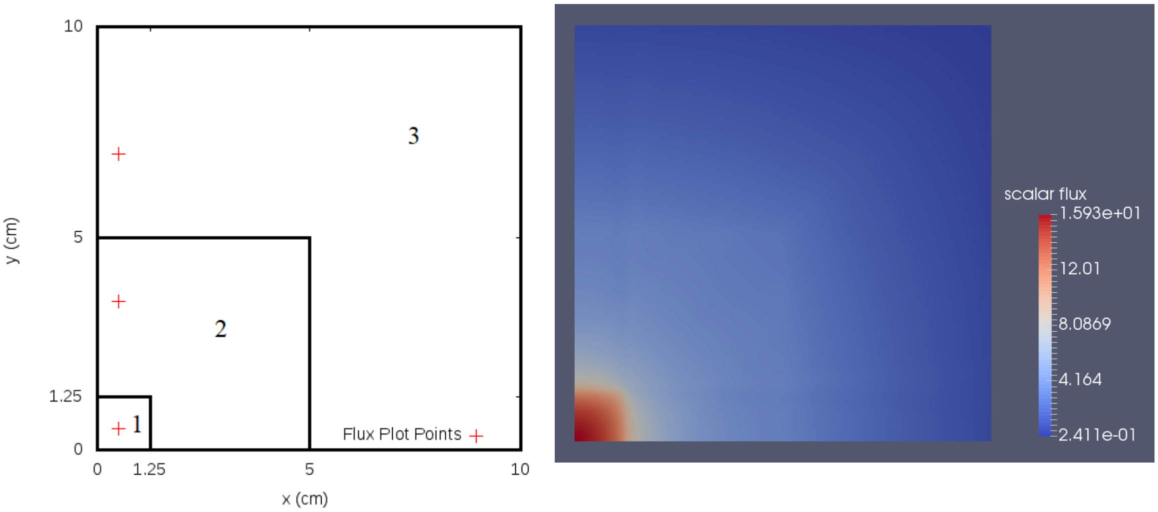

(a) Schematic

(b) S_{30} scalar flux solution

Figure 9: Schematic (a) and S_{30} scalar flux solution (b) for the Watanabe-Maynard extrapolation problem. Region 1 is the source, region 2 is a void and region 3 is a highly scattering material.

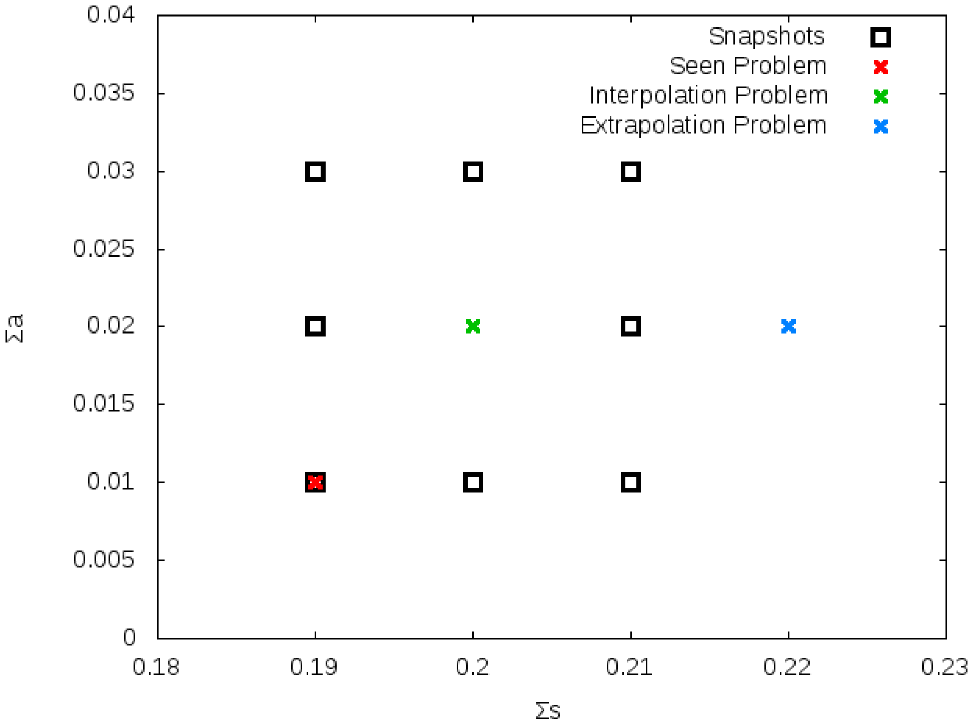

Figure 10: Σ_a and Σ_s values used to produce snapshots and solutions for the Watanabe-Maynard problem.

Figure 10 shows Σ_a and Σ_s values in the non-void regions for eight solutions and three test problems. The snapshot matrix was formed by combining all eight training solutions, and the POD bases produced were used to solve all three test problems. As previously, one set of material properties had been seen in the offline stage, one was an interpolation between seen properties, and one was an extrapolation outside the range of seen material properties.

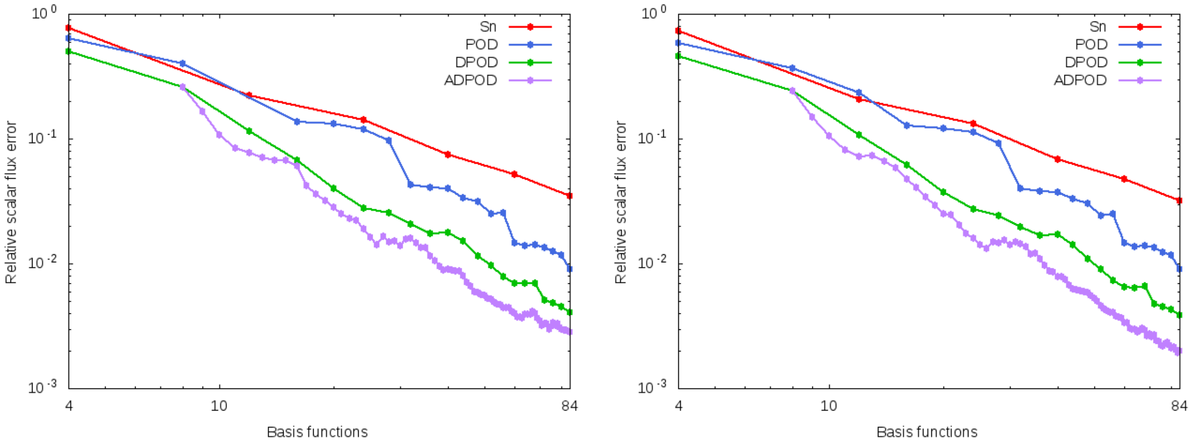

(a) Seen

(b) Interpolation

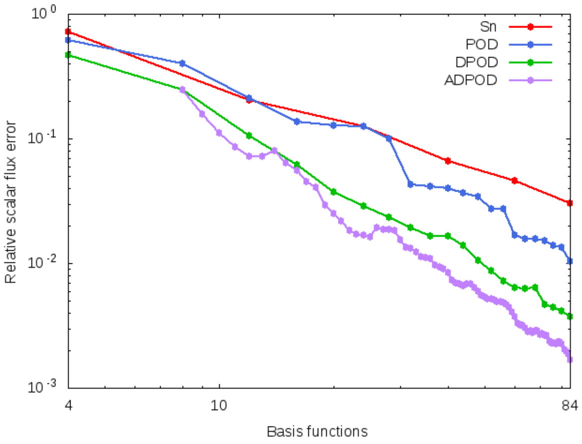

(c) Extrapolation

Figure 11: The effects of increasing basis function count on relative scalar flux error for the Watanabe-Maynard problem.

Figure 11 presents the L1-norm of the relative scalar flux error in solutions of the Watanabe-Maynard problem against the mean number of basis functions per node for each method. Results from low order S_n solutions are shown for comparison, as previously. POD generally has a lower error than S_n with the same number of basis functions, though it struggles with 12 and fewer. In particular, POD could not produce a solution for the seen problem with 12 basis functions, as the system diverged. DPOD did not suffer from this issue, and also had consistently less error for a given number of basis functions than POD, usually by a factor of two to three. ADPOD improved upon this, reducing error by a further factor of two to three compared to DPOD.

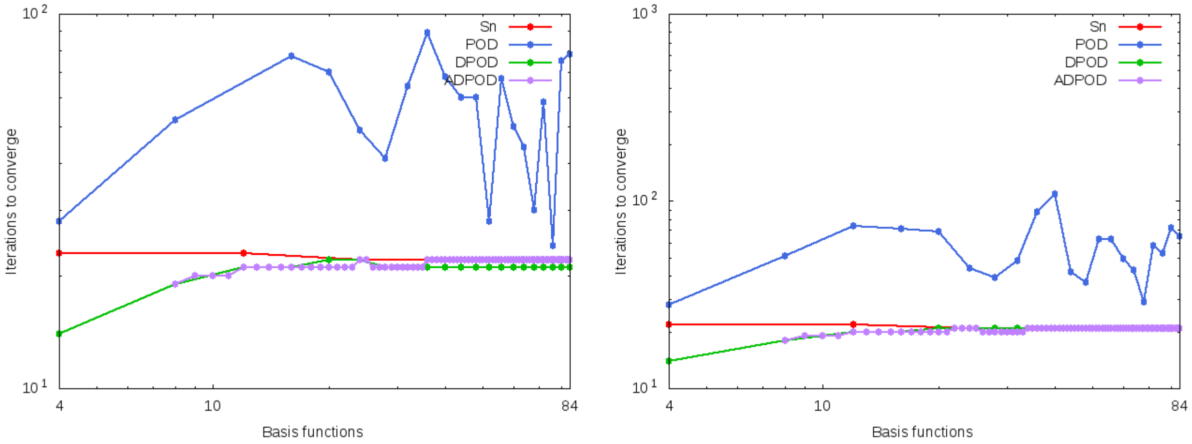

(a) Seen (b) Interpolation

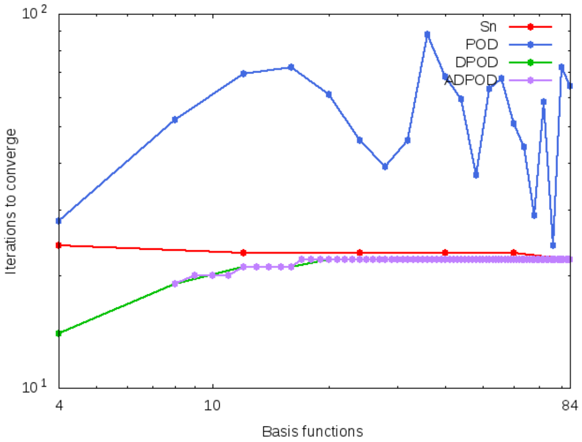

(c) Extrapolation

Figure 12: The effects of increasing basis function count on the number of iterations required to converge a solution for the Watanabe-Maynard problem.

Figure 12 shows the number of solver iterations required to converge the solution for each method. The criterion for a solution to be considered converged was the same as for the dog-leg duct problem. However, the Watanabe-Maynard problem contains scattering regions, which must be resolved by iteration, and so every method required more iterations to converge than for the dog-leg duct problem. As previously, POD performs much worse than any other method by this metric. It consistently requires more iterations to solve than the three other methods, and the exact number of iterations varies significantly. The other methods all converge faster and more consistently, with similar numbers of iterations required for each, though DPOD and ADPOD converge slightly faster than S_n. As the solve time is dependant on the number of iterations required to converge the solution, DPOD and ADPOD again demonstrate clear benefits over POD.

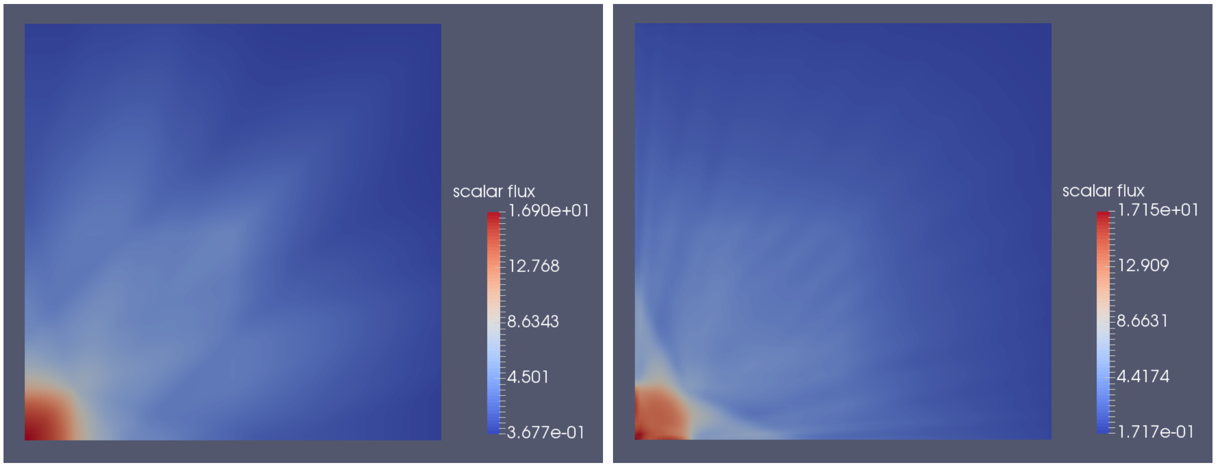

(a) S_6 (b) POD

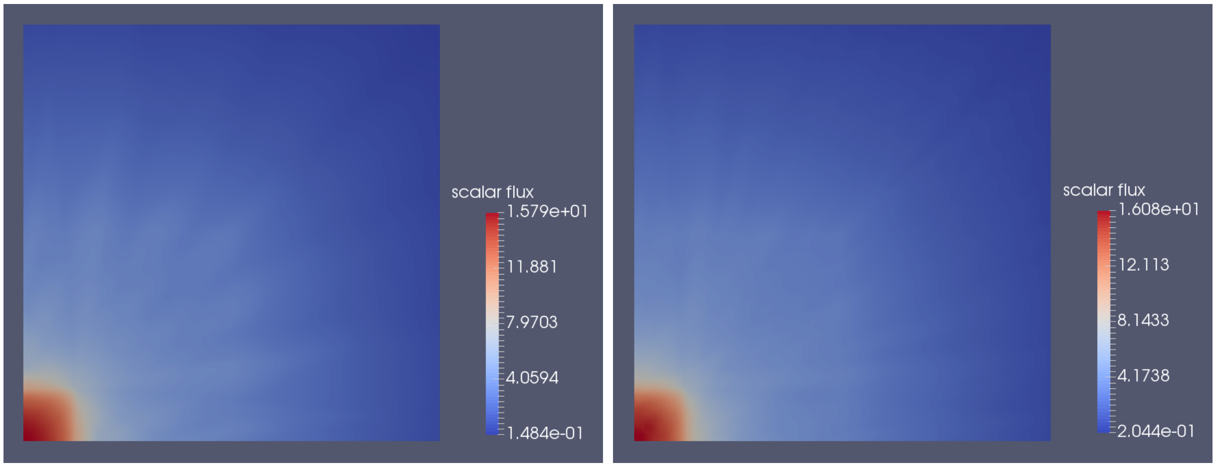

(c) Discontinuous POD (d) Adaptive Discontinuous POD

Figure 13: Solutions to the Watanabe-Maynard extrapolation problem for all four methods with 24 basis functions per node, on average in the case of ADPOD.

Figure 13 depicts the scalar flux solutions produced for the extrapolation problem by each method with 24 basis functions per node (on average in the case of ADPOD). The full order solution to this problem is shown in figure 9b. The S_6 solution exhibits significant ray effects. POD reduces these ray effects somewhat, but makes them less regular and so increases error in some regions. It is particularly poor at resolving the source region, which leads to a more inaccurate peak flux than even the S_6 solution. DPOD improves upon the previous method - ray effects are significantly reduced, the source region is well resolved and the peak flux is relatively accurate. However, some visible ray effects are still present. ADPOD further reduces ray effects compared to DPOD, and has a similar error in its peak flux. This shows the advantages of DPOD and ADPOD in resolving scattering problems with few basis functions.

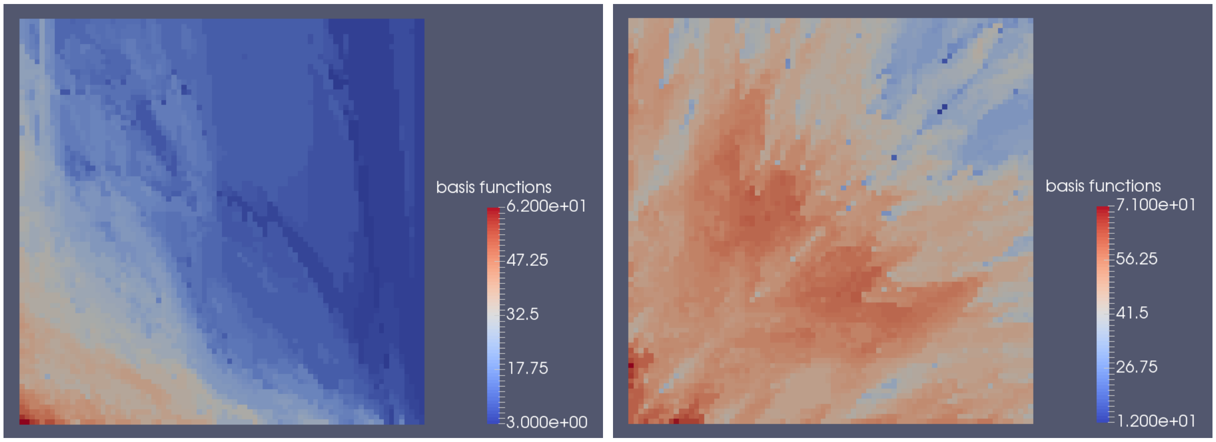

(a) -x, +y octant

(b) +x, +y octant

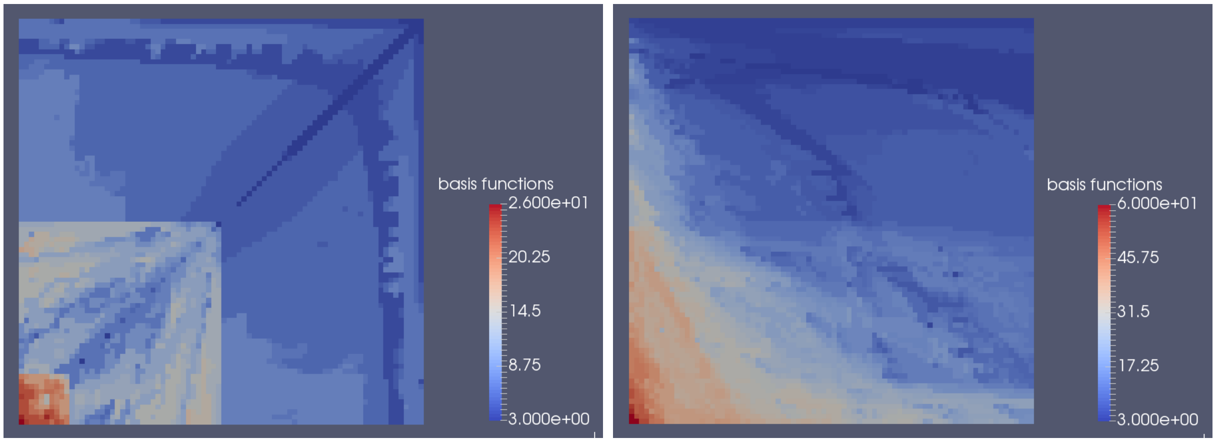

(c) -x, -y octant

(d) +x, -y octant

Figure 14: Number of ADPOD basis functions per octant per node for the duct extrapolation problem, with a mean of 84 basis functions per node in total.

Figure 14 depicts the number of ADPOD basis functions per node for each element and octant in the domain. In contrast to the dog-leg duct results, a significant amount of flux was present in all four octants due to scattering. Despite this, the algorithm was able to distribute the added basis functions in order to maximise their contribution to the solution. Most basis functions were added to the top right octant, which contained the majority of the flux due to the problem's geometry. In the other three octants, basis functions were primarily added downstream of the source and scattering regions. This optimisation enabled ADPOD to improve upon POD and DPOD, despite the broad angular flux distribution in the Watanabe-Maynard problem.

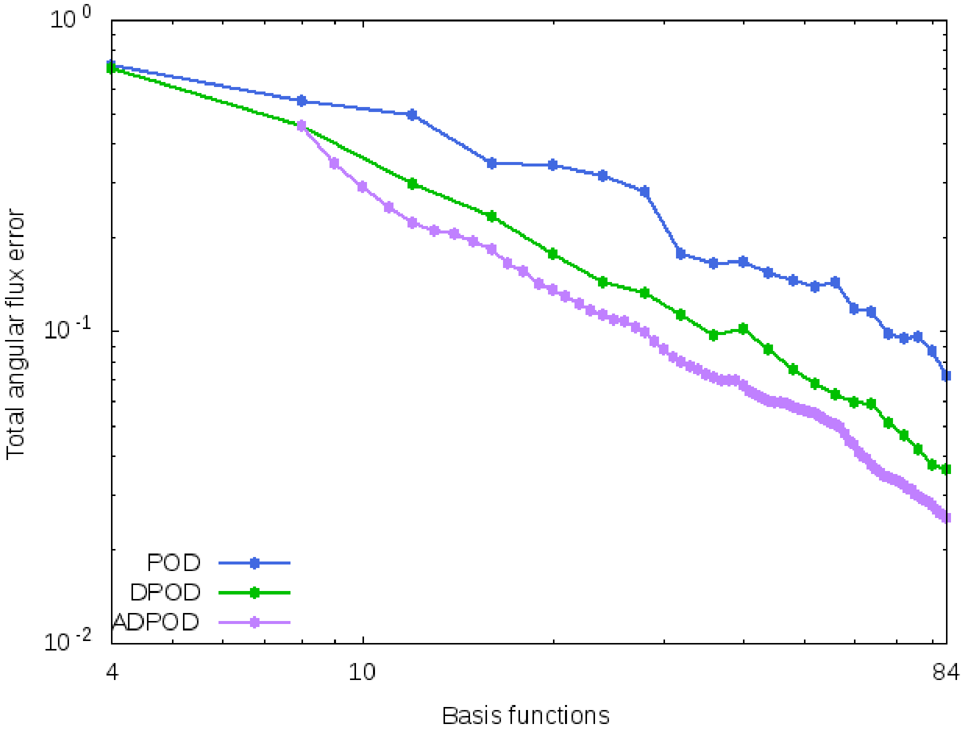

Figure 15: The effect of increasing basis function count on relative angular flux error for the Watanabe-Maynard extrapolation problem.

Figure 15 presents the relative angular flux error against the number of basis functions per node. These are similar to the results for the dog-leg duct problem - POD has the highest error, then DPOD, and ADPOD has the lowest. However, ADPOD does not improve upon DPOD to the same extent as in the dog-leg duct problem. This is because flux travels throughout the domain in this problem, and so basis functions are required everywhere. By contrast, in the dog-leg duct problem flux is primarily limited to the duct itself, and so basis functions can be focused in that region. The results again demonstrate that increasing the basis function count consistently reduces the angular flux error.

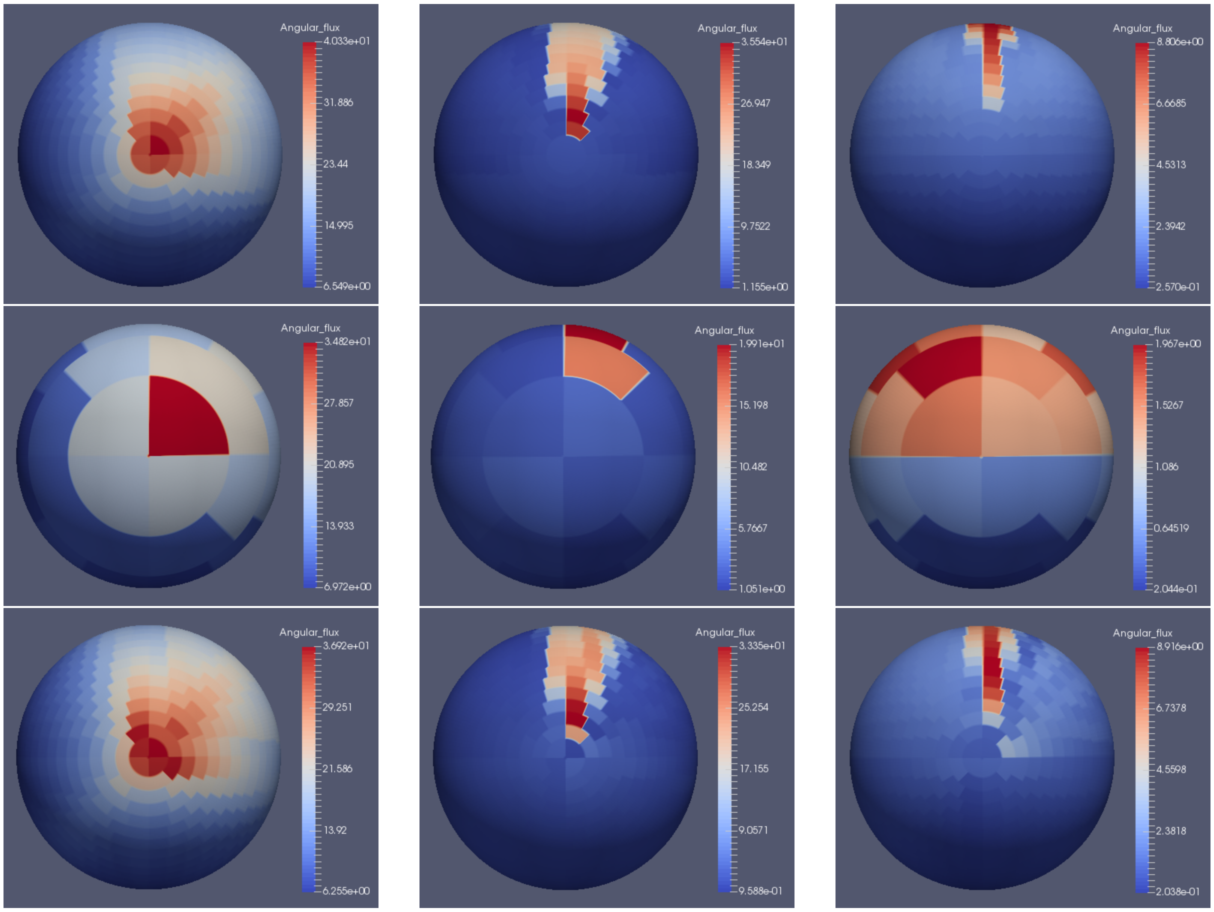

(a) x = 0.625cm, y = 0.625cm (b) x = 0.625cm, y = 4.375cm (c) x = 0.625cm, y = 8.750cm

Figure 16: Angular flux distributions at various points in the Watanabe-Maynard extrapolation problem. From top to bottom, S_{30}, S_6, and adaptive discontinuous POD. ADPOD and S_6 both use a mean of 24 basis functions per node, while S_{30} has 480.

Figure 16 depicts the angular flux profiles at three points within the domain. The locations of each point are shown in figure 9a. S_6 and ADPOD both used 24 basis functions, while S_{30} has 480. The S_6 discretisation is highly inaccurate. It performs particularly poorly in the scattering region, where the flux is distributed over a much wider range of angles than it should be. ADPOD performs far better, producing a reasonable approximation to the expected result.

4 Conclusion

This article has developed a new reduced order model for discretising the angular dimension of the Boltzmann transport equation, which is based upon POD and the method of snapshots. The snapshots are formed from angular flux vectors taken from many spatial nodes, in a similar manner to previous work on angular POD [29]. The novel aspects of this work are the discontinuous formulation of the POD basis functions, which each span one octant in angle rather than the entire sphere, and the implementation of angular adaptivity using these discontinuous basis functions.

For the same angular size, the method of DPOD is shown to consistently increase accuracy slightly over POD, which is itself more accurate than S_n by up to half an order of magnitude. Additionally, the number of solver iterations required to converge each solution was both reduced by up to 1.8 orders of magnitude, and made more stable - solver iterations oscillated with basis size when using POD, while with DPOD they converged to the same value and did not oscillate. When combined, these two factors reduced solve times by up to 2 orders of magnitude compared to POD for the dog-leg duct problem, and by up to 1.5 orders of magnitude for the Watanabe-Maynard problem.

In both numerical examples, the implementation of angular adaptivity via ADPOD provided further reductions in error. This was particularly successful in the dog-leg duct problem, where ADPOD reduced error for a given number of basis functions by up to an order of magnitude over DPOD, and thereby reduced solve times by up to 3 orders of magnitude in total compared to POD. For the Watanabe-Maynard problem, the reduction in error compared to DPOD was by a factor of 2-3, and the total reduction in solve time compared to POD was by up to 1.8 orders of magnitude. A significant contributor to the reduction in error is that ADPOD almost eliminated visible ray effects, thus improving accuracy in those regions which would've been affected by them.

The benefits of DPOD and ADPOD have been clearly demonstrated in this article. They offer reductions in both total error and solver iterations compared to previous angular POD methods. As both of these factors contribute to computational complexity, when they are both reduced computation time decreases significantly.

5 Acknowledgements

The authors wish to acknowledge the funding of this research through the grant EP/M022684/2. This research utilised Queen Mary's Apocrita HPC facility, supported by QMUL Research-IT. http://doi.org/10.5281/zenodo.438045

6 Bibliography

[1] F. Golse. "Applications of the boltzmann equation within the context of upper atmosphere vehicle aerodynamics". In: *Computer Methods in Applied Mechanics and Engineering* 75.1 (1989), pp. 299–316. DOI: `https://doi.org/10.1016/0045-7825(89)90031-5`.

[2] S.P. Burns and M.A. Christen. "Spatial domain-based parallelism in large scale, participating-media, radiative transport applications". In: *Numerical Heat Transfer, Part B: Fundamentals* 31.4 (1997), pp. 401–421. DOI: `https://doi.org/10.2172/402363`.

[3] W. Li et al. “ShengBTE: A solver of the Boltzmann transport equation for phonons”. In: *Computer Physics Communications* 185.6 (2014), pp. 1747–1758. DOI: `https://doi.org/10.1016/j.cpc.2014.02.015`.

[4] S. Dargaville et al. “Scalable angular adaptivity for Boltzmann transport”. In: *Journal of Computational Physics* (2019), pp. 109–124. DOI: `https://doi.org/10.1016/j.jcp.2019.109124`.

[5] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Second. Computer Science and Scientific Computing. Academic Press, 1990. ISBN: 9780122698514. DOI: `https://doi.org/10.1016/C2009-0-27872-X`.

[6] I.T. Jolliffe. *Principal Component Analysis*. Second. Berlin: Springer, Oct. 2002. ISBN: 0387954422. DOI: `https://doi.org/10.1007/978-3-642-04898-2_455`.

[7] D.T. Crommelin and A.J. Majda. “Strategies for model reduction: comparing different optimal bases”. In: *Journal of the Atmospheric Sciences* 61 (2004), pp. 2206–2217. DOI: `https://doi.org/10.1175/1520-0469(2004)061<2206:SFMRCD>2.0.CO;2`.

[8] J.L. Lumley. “The Structure of Inhomogeneous Turbulent Flows”. In: *Atmospheric Turbulence and Radio Wave Propagation* (1967), pp. 166–178.

[9] H.P. Bakewell and J.L. Lumley. “Viscous sublayer and adjacent wall region in turbulent pipe flow”. In: *The Physics of Fluids* 10 (1967), pp. 1880–1889. DOI: `https://doi.org/10.1063/1.1762382`.

[10] F.R. Payne and J.L. Lumley. “Large-eddy structure of the turbulent wake behind a circular cylinder”. In: *The Physics of Fluids* 10 (1967), pp. 194–196. DOI: `https://doi.org/10.1063/1.1762445`.

[11] A.E. Dean and C. Mavriplis. “Low-dimensional description of the dynamics in separated flow past thick airfoil”. In: *AIAA* 32 (1994), pp. 1222–1227. DOI: `https://doi.org/10.2514/3.12123`.

[12] J. Delville, S. Bellin and J.P. Bonnet. “Use of the proper orthogonal decomposition in a plane turbulent mixing layer, in Turbulence and Coherent Structures”. In: Dordrecht: Kluwer Academic Publishers, 1991, pp. 75–90. DOI: `https://doi.org/10.1007/978-94-015-7904-9_5`.

[13] L. Sirovich and H. Park. “Turbulent thermal convection in a finite domain: PartI. Theory”. In: *Physics of fluids* 2 (1990), pp. 1649–1658. DOI: `https://doi.org/10.1063/1.857572`.

[14] H. Park and L. Sirovich. “Turbulent thermal convection in a finite domain: PartII. Numerical results”. In: *Physics of fluids* 2 (1990), pp. 1659–1668. DOI: `https://doi.org/10.1063/1.857573`.

[15] F. Fang et al. “A POD reduced order unstructured mesh ocean modelling method for moderate Reynolds number flows”. In: *Ocean Modelling* 28 (1-3 2009), pp. 127–136. DOI: `https://doi.org/10.1016/j.ocemod.2008.12.006`.

[16] Thomas JR Hughes. *The finite element method: linear static and dynamic finite element analysis*. Courier Corporation, 2012.

[17] Randall J. LeVeque. *Finite Difference Methods for Ordinary and Partial Differential Equations.* Society for Industrial and Applied Mathematics, 2007. DOI: `10.1137/1.9780898717839`. eprint: `https://epubs.siam.org/doi/pdf/10.1137/1.9780898717839`. URL: `https://epubs.siam.org/doi/abs/10.1137/1.9780898717839`.

[18] H. Adison. *Principles and Practice of Finite Volume Method.* Clanrye International, 2015. ISBN: 9781632404176. URL: `https://books.google.co.uk/books?id=yUr7rQEACAAJ`.

[19] F. Chinesta, P. Ladeveze and E. Cueto. "A short review on model order reduction based on proper generalized decomposition". In: *Archives of Computational Methods in Engineering* 18.4 (2011), p. 395. DOI: `https://doi.org/10.1007/s11831-011-9064-7`.

[20] M. Barrault et al. "An 'empirical interpolation' method: application to efficient reduced-basis discretization of partial differential equations". In: *Comptes Rendus Mathematique* 339.9 (2004), pp. 667–672. DOI: `https://doi.org/10.1016/j.crma.2004.08.006`.

[21] N.C. Nguyen and J. Peraire. "An efficient reduced-order modeling approach for non-linear parametrized partial differential equations". In: *International Journal for Numerical Methods in Engineering* 76.1 (2008), pp. 27–55. DOI: `https://doi.org/10.1002/nme.2309`.

[22] D. Xiao et al. "Machine learning-based rapid response tools for regional air pollution modelling". In: *Atmospheric Environment* 199 (2019), pp. 463–473. DOI: `https://doi.org/10.1016/j.atmosenv.2018.11.051`.

[23] K. Pearson. "LIII. On lines and planes of closest fit to systems of points in space". In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2.11 (1901), pp. 559–572. DOI: `https://doi.org/10.1080/14786440109462720`.

[24] L. Sirovich. "Turbulence and the dynamics of coherent structures". In: *Quarterly of Applied Mathematics* 5 (1987), pp. 561–590. DOI: `https://doi.org/10.1090/qam/910462`.

[25] A.G. Buchan et al. "A POD reduced-order model for eigenvalue problems with application to reactor physics". In: *International Journal for Numerical Methods in Engineering* 95.12 (2013), pp. 1011–1032. DOI: `https://doi.org/10.1002/nme.4533`.

[26] F. Wols. *Transient analyses of accelerator driven systems using modal expansion techniques.* 2010.

[27] A. Sartori et al. "Comparison of a Modal Method and a Proper Orthogonal Decomposition approach for multi-group time-dependent reactor spatial kinetics". In: *Annals of Nuclear Energy* 71.0 (2014), pp. 217–229. DOI: `https://doi.org/10.1016/j.anucene.2014.03.043`.

[28] A. Sartori et al. "A multi-physics reduced order model for the analysis of Lead Fast Reactor single channel". In: *Annals of Nuclear Energy* 87 (2016), pp. 198–208. DOI: `https://doi.org/10.1016/j.anucene.2015.09.002`.

[29] A.G. Buchan et al. “A POD reduced order model for resolving angular direction in neutron/photon transport problems”. In: *Journal of Computational Physics* 296 (2015), pp. 138–157. DOI: `https://doi.org/10.1016/J.JCP.2015.04.043`.

[30] J. Tencer et al. “Accelerated Solution of Discrete Ordinates Approximation to the Boltzmann Transport Equation for a Gray Absorbing–Emitting Medium Via Model Reduction”. In: *Journal of Heat Transfer* 139.12 (Aug. 2017). DOI: `https://doi.org/10.1115/1.4037098`.

[31] S. Lorenzi. “An Adjoint Proper Orthogonal Decomposition method for a neutronics reduced order model”. In: *Annals of Nuclear Energy* 114 (2018), pp. 245–258. DOI: `https://doi.org/10.1016/j.anucene.2017.12.029`.

[32] C. Castagna et al. “Development of a Reduced Order Model for Fuel Burnup Analysis”. In: *Energies* 13.4 (Feb. 2020), p. 890. DOI: `https://doi.org/10.3390/en13040890`.

[33] F. Alsayyari et al. “A nonintrusive adaptive reduced order modeling approach for a molten salt reactor system”. In: *Annals of Nuclear Energy* 141 (2020), p. 107321. DOI: `https://doi.org/10.1016/j.anucene.2020.107321`.

[34] R. Manthey et al. “Reduced order modeling of a natural circulation system by proper orthogonal decomposition”. In: *Progress in Nuclear Energy* 114 (2019), pp. 191–200. DOI: `https://doi.org/10.1016/j.pnucene.2019.03.010`.

[35] A.G. Buchan et al. “Linear and quadratic octahedral wavelets on the sphere for angular discretisations of the Boltzmann transport equation”. In: *Annals of Nuclear Energy* 32.11 (2005), pp. 1224–1273. ISSN: 0306-4549. DOI: `https://doi.org/10.1016/j.anucene.2005.01.005`. URL: `http://www.sciencedirect.com/science/article/pii/S0306454905000356`.

[36] S. Dargaville et al. “Angular adaptivity with spherical harmonics for Boltzmann transport”. In: *Journal of Computational Physics* 397 (Nov. 2019), p. 108846. ISSN: 0021-9991. DOI: `10.1016/j.jcp.2019.07.044`. URL: `http://dx.doi.org/10.1016/j.jcp.2019.07.044`.

[37] H. Park and C.R.E. de Oliveira. “Coupled Space-Angle Adaptivity for Radiation Transport Calculations”. In: *Nuclear Science and Engineering* 161.2 (2009), pp. 216–234. DOI: `https://doi.org/10.13182/NSE161-216`.

[38] A.G. Buchan et al. “Self-Adaptive Spherical Wavelets for Angular Discretizations of the Boltzmann Transport Equation”. In: *Nuclear Science and Engineering* 158 (Mar. 2008). DOI: `https://doi.org/10.13182/NSE08-A2751`.

[39] J. Du et al. “Reduced-order modeling based on POD of a parabolized Navier-Stokes equation model I: forward model”. In: *International Journal for Numerical Methods in Fluids* 69.3 (2012), pp. 710–730. DOI: `https://doi.org/10.1002/fld.2606`.

[40] J. Du et al. “Reduced-order modeling based on POD of a parabolized Navier-Stokes equation model II: Trust region POD 4-D VAR Data Assimilation”. In: *Computers & Mathematics with Applications* 65 (2013), pp. 380–394. DOI: `https://doi.org/10.1016/j.camwa.2012.06.001`.

[41] W.H. Reed and T.R. Hill. "Triangular mesh methods for the neutron transport equation". In: (Oct. 1973).

[42] J. Kópházi and D. Lathouwers. "A space–angle DGFEM approach for the Boltzmann radiation transport equation with local angular refinement". In: *Journal of Computational Physics* 297 (2015), pp. 637–668. DOI: `https://doi.org/10.1016/j.jcp.2015.05.031`.

[43] A.R. Owens et al. "Discontinuous isogeometric analysis methods for the first-order form of the neutron transport equation with discrete ordinate (SN) angular discretisation". In: *Journal of Computational Physics* 315 (2016), pp. 501–535. DOI: `https://doi.org/10.1016/j.jcp.2016.03.060`.

[44] A.G. Buchan et al. "Linear and quadratic octahedral wavelets on the sphere for angular discretisations of the Boltzmann transport equation". In: *Annals of Nuclear Energy* 32.11 (2005), pp. 1224–1273. DOI: `https://doi.org/10.1016/j.anucene.2005.01.005`.

[45] Y. Watanabe and C.W. Maynard. "The discrete cones method for two-dimensional neutron transport calculations". In: *Transport Theory and Statistical Physics* 15.1-2 (1986), pp. 135–156. DOI: `https://doi.org/10.1080/00411458608210448`.

7 Tables

Problem	Region	Source ($cm^{-2}s^{-1}$)	Σ_a (cm^{-1})	Σ_s (cm^{-1})
1	1	1.00	0.40	0.00
	2	0.00	0.00	0.00
	3	0.00	0.40	0.00
2	1	1.00	0.50	0.00
	2	0.00	0.00	0.00
	3	0.00	0.50	0.00
3	1	1.00	0.60	0.00
	2	0.00	0.00	0.00
	3	0.00	0.60	0.00

Table 1: Material properties for the dog-leg duct training solutions.

Problem	Region	Source ($\mathrm{cm}^{-2}\mathrm{s}^{-1}$)	Σ_a (cm^{-1})	Σ_s (cm^{-1})
Seen	1	1.00	0.50	0.00
	2	0.00	0.00	0.00
	3	0.00	0.50	0.00
Interpolate	1	1.00	0.45	0.00
	2	0.00	0.00	0.00
	3	0.00	0.45	0.00
Extrapolate	1	1.00	0.35	0.00
	2	0.00	0.00	0.00
	3	0.00	0.35	0.00

Table 2: Material properties for the dog-leg duct test problems.