



ELSEVIER

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp



A geometry preserving, conservative, mesh-to-mesh isogeometric interpolation algorithm for spatial adaptivity of the multigroup, second-order even-parity form of the neutron transport equation



J.A. Welch*, J. Kópházi, A.R. Owens, M.D. Eaton

Nuclear Engineering Group, Department of Mechanical Engineering, City and Guilds Building, Imperial College London, Exhibition Road, South Kensington, London, SW7 2AZ, United Kingdom

ARTICLE INFO

Article history:

Received 18 August 2016

Received in revised form 2 February 2017

Accepted 7 June 2017

Available online 27 June 2017

Keywords:

IGA

Neutron transport

Multigroup

Energy-dependent meshes

ABSTRACT

In this paper a method is presented for the application of energy-dependent spatial meshes applied to the multigroup, second-order, even-parity form of the neutron transport equation using Isogeometric Analysis (IGA). The computation of the inter-group regenerative source terms is based on conservative interpolation by Galerkin projection. The use of Non-Uniform Rational B-splines (NURBS) from the original computer-aided design (CAD) model allows for efficient implementation and calculation of the spatial projection operations while avoiding the complications of matching different geometric approximations faced by traditional finite element methods (FEM). The rate-of-convergence was verified using the method of manufactured solutions (MMS) and found to preserve the theoretical rates when interpolating between spatial meshes of different refinements. The scheme's numerical efficiency was then studied using a series of two-energy group pincell test cases where a significant saving in the number of degrees-of-freedom can be found if the energy group with a complex variation in the solution is refined more than an energy group with a simpler solution function. Finally, the method was applied to a heterogeneous, seven-group reactor pincell where the spatial meshes for each energy group were adaptively selected for refinement. It was observed that by refining selected energy groups a reduction in the total number of degrees-of-freedom for the same total L_2 error can be obtained.

© 2017 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The neutron transport equation is a partial-integro-differential equation (PIDE) based upon a linearized version of the Boltzmann transport equation; commonly used in the field of statistical mechanics to model the average statistical behaviour of interacting gases of particles, atoms or molecules. The transport equation is used to model the statistical behaviour of neutrons interacting within a host medium such as a nuclear reactor core or a radiation shield [1]. The behaviour of the neutrons is described using the angular flux which is the dependent variable in the neutron transport equation. The angular flux is a function of a seven-dimensional phase space of independent variables consisting of three dimensions for the spatial

* Corresponding author.

E-mail address: jameswelch24@gmail.com (J.A. Welch).

variation x , y and z ; two dimensions for the angular variation θ and χ ; one dimension for the energy variation E and finally one variable for the temporal variation t . To accurately represent the angular flux sufficient resolution is required in each of the seven dimensions of the phase space. Combined with the large and geometrically complex nuclear reactor core designs and radiation shields the solution of this equation represents a significant challenge even for contemporary state-of-the-art numerical algorithms and high performance computing (HPC) architectures.

The variation of the angular flux with energy is usually very complex and represents a key challenge when solving the neutron transport equation. This is mainly due to the complicated variation of the interaction effects of the neutrons with the host media; which is represented by the neutron interaction cross sections. This complicated energy dependence of the cross sections leads the angular flux to have a significant dependence on energy. One common approach to discretise the energy variation of the neutron transport equation is to assume that the energy variation can be approximated by integral quantities over a number of discrete energy groups each with a prescribed upper and lower energy bound, the multigroup approximation. The energy range of neutrons typically spans between less than 1 eV to several MeV depending upon the nature of the problem being solved. Therefore, potentially a significant number of energy groups may be required to represent the energy variation of the physical phenomena described by the neutron transport equation [2].

Current deterministic reactor physics codes used in the industry for core and radiation shield design such as PANTHER [3] and ATILLA [4] utilize a common spatial mesh for each energy group in order to solve the neutron transport equation. However, the spatial variation of the angular flux will, in general, vary significantly with the energy of the neutrons; as a consequence of the energy variation of the cross sections. Therefore using the same spatial mesh for each energy group will not prove optimal. In order to optimise accuracy and computational efficiency it is more desirable to use different spatial meshes for each energy group. However, this represents another series of significant challenges for the discretisation scheme in terms of conservatively interpolating source and solution functions from one spatial mesh to another as well as preserving the exact geometry of the solution domain. The challenge of interpolating the numerical representation of a function between different spatial meshes is not unique to the application of the multigroup form of the neutron transport equation. In a wider context, function interpolation is required in a diverse range of situations; for example in the coupling of a reactor physics code to a computational fluid dynamics (CFD) code [5], visualisation and comparison of data sets computed on different spatial grids and the interpolation of a solution function onto a new grid using adaptive mesh refinement (AMR) [6].

The first application of group-dependent spatial meshes was by Wang et al. who presented an adaptive mesh refinement (AMR) scheme allowing the use of unique meshes for each energy group applied to the neutron diffusion equation [7]. Their approach used hierarchically refined finite element meshes from a common coarse grid. The use of a common coarse grid limits this approach to Cartesian geometries (or geometries consisting solely of planar interfaces) or sacrificing any improvements in the representation of the geometry as the mesh is refined. In Baker et al. a method allowing for the conservative interpolation between unstructured triangular finite element meshes was presented [2] based upon the work of Farrell and Maddison [8]. This approach required the generation of a 'super-mesh' to facilitate the transfer of information from one group to another. However, the two spatial meshes required no shared underlying structure. A reduction in the total number of finite elements required to obtain the same numerical accuracy was presented. Goffin et al. also used this methodology for goal-based AMR [9]. As in Wang et al., all of the research so far in this area has been demonstrated for Cartesian geometries only.

Isogeometric Analysis (IGA) is a generalisation of the finite element method (FEM) where the mathematical description of a computer-aided design (CAD) model, typically Non-Uniform Rational B-splines (NURBS), is used as the mathematical basis for analysis [10]. By utilising a geometry description based upon the original CAD model the mesh is exact from the coarsest level and further refinement to improve the fidelity of the solution does not alter the description of the geometry in either the parametric or real space. IGA has been extended to use locally refineable T-spline patches remedying the limitation of the tensor-product refinement of the standard NURBS representation [11]. IGA has been applied to a wide variety of physics and engineering disciplines including electromagnetics [12], computational fluid dynamics (CFD) [13,14] and solid dynamics [10].

In the field of reactor physics, IGA was first applied to the mono-energetic neutron diffusion equation by Hall et al. [15]. This work was later extended by Welch et al. who applied IGA to the multigroup form of the neutron diffusion equation for a quarter-core reactor geometry [16]. In these studies it was found that the exact representation of the problem geometry significantly improved the accuracy of the obtained solution for the same computational cost of equivalent finite element calculations. It was also found that the increased basis function continuity available to B-splines and NURBS significantly reduced the required computational effort to obtain a given level of solution accuracy. In addition to applications to the neutron diffusion equation, IGA has been applied to the first-order form of the neutron transport equation with a discontinuous Galerkin (DG) spatial discretisation and S_N angular discretisation [17].

The NURBS patches used in IGA have a series of useful properties that make them suitable for use in energy-dependent mesh methods. As the geometry is exact at the coarsest level (the CAD description) and further refinements do not change the physical geometry, a NURBS patch can be refined to suit the physics of each energy group independently with each representation conforming to the original geometry. In addition, as the parametric description does not change under refinement the projective matrices required for the function projection can be efficiently computed without the need to generate a physical 'super-mesh'.

In this work an approach for independently refined spatial meshes using tensor-product NURBS-based IGA has been developed and applied to the multigroup second-order even-parity form of the neutron transport equation with a spherical harmonic (P_N) angular discretisation. A conservative interpolation based on Galerkin projection has been used in an identical manner to Farrell and Maddison [8]. The remainder of this paper is outlined as follows: a brief introduction to IGA and a description of the calculation of the projective matrices is given in Section 2. In Section 3 the multigroup form of the even-parity equations is presented and the weak form is derived for the energy group-dependent mesh method. In Section 4 the numerical results of the study including verification of the rate-of-convergence using the method of manufactured solutions (MMS) are presented. Two test cases are studied based on a synthetic two-group pincell benchmark and a representative seven-group uranium-dioxide (UOX) pincell from the OECD/NEA C5G7 benchmark. Conclusions and future considerations are summarised in Section 5.

2. Isogeometric spatial discretisation

Isogeometric Analysis (IGA) is a generalisation of the Finite Element Method (FEM) that utilises the mathematical description of a Computer-Aided Design (CAD) model as the shape functions in numerical analysis. The prevailing technology for representing geometries in CAD are Non-Uniform Rational B-splines (NURBS) although the use of any exact mathematical basis from a CAD model can be considered in IGA.

NURBS are rational functions derived from B-spline functions. A set of B-spline functions is fully defined by a polynomial order p and a knot-vector (Ξ). The knot vector is a set of non-decreasing real numbers, for example $\Xi = [\xi_1, \xi_2, \dots, \xi_{n+p}, \xi_{n+p+1}]$ where n is the total number of basis functions. Open knot vectors (where the first and final knots have multiplicity $p + 1$) are used in this work. For a given knot vector Ξ , B-spline basis functions $N(\xi)$ of a given order p can be calculated beginning with the constant functions (0th order):

$$N_i^0(\xi) = \begin{cases} 1 & \text{if } \xi_i \leq \xi < \xi_{i+1} \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

where ξ_i is the i th knot point and ξ is a specified point in the one-dimensional parameter space. Higher-order B-spline basis functions can be calculated recursively using the Cox-de Boor recursion formula [10] as follows:

$$N_i^p(\xi) = \frac{\xi - \xi_i}{\xi_{i+p} - \xi_i} N_i^{p-1}(\xi) + \frac{\xi_{i+p+1} - \xi}{\xi_{i+p+1} - \xi_{i+1}} N_{i+1}^{p-1}(\xi). \quad (2)$$

If the denominator of the fraction is zero, i.e. $\xi_{i+p} = \xi_i$ or $\xi_{i+p+1} = \xi_{i+1}$ then the whole fractional term is set to zero. The $p = 0$ and $p = 1$ basis functions are identical to standard finite element piecewise constant and piecewise linear functions respectively. The derivatives of B-spline basis functions are calculated from the $p - 1$ th B-spline basis function as follows:

$$\frac{d}{d\xi} N_i^p(\xi) = \frac{p}{\xi_{i+p} - \xi_i} N_i^{p-1}(\xi) - \frac{p}{\xi_{i+p+1} - \xi_{i+1}} N_{i+1}^{p-1}(\xi). \quad (3)$$

NURBS basis functions are derived from the B-spline basis functions which span the same parametric space. The distinction comes that the weighting values for each control point (an additional dimension of the control grid) is required to provide a projective transformation, allowing NURBS to exactly represent conic-sections such as circles. In one dimension this can be written in the following form:

$$R_i^p(\xi) = \frac{N_i^p(\xi) w_i}{W(\xi)} = \frac{N_i^p w_i}{\sum_{\hat{i}=1}^n N_{\hat{i}}^p(\xi) w_{\hat{i}}}, \quad (4)$$

where $R_i^p(\xi)$ is the resulting NURBS basis function evaluated at point ξ in the parameter space, N_i^p is the B-spline basis function and w_i are the weighting values for basis function i . It can be seen that NURBS are defined by a division by a B-spline creating a rational function. Two-dimensional bivariate and three-dimensional trivariate NURBS basis functions are derived as a tensor-product of the one dimensional B-spline basis functions for all parametric directions ξ and η in \mathbb{R}^2 and ξ , η and ζ in \mathbb{R}^3 . They are then multiplied by the corresponding control net weight, w and the weighting function W . A full description of the generation of B-spline and NURBS functions is given in [10]. A NURBS surface is defined as a linear combination of the NURBS basis functions weighted by the corresponding control points \mathbf{B} :

$$S(\xi, \eta) = \sum_{i=1}^N R_i(\xi, \eta) \mathbf{B}_i \quad (5)$$

where N is the total number of basis functions describing the NURBS surface. An example of a NURBS surface and the role the control points play in shape manipulation is shown in Fig. 1.




Fig. 1. The shape of a NURBS patch is determined by a set of coordinates, \mathbf{B} in physical space (highlighted in blue) that creates a control polygon. In the above figure three NURBS surfaces in \mathbb{R}^3 have been defined and by altering the position of a control point the surface is perturbed. In this manner and by increasing the number of degrees of freedom a wide variety of complex objects can be modelled exactly. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

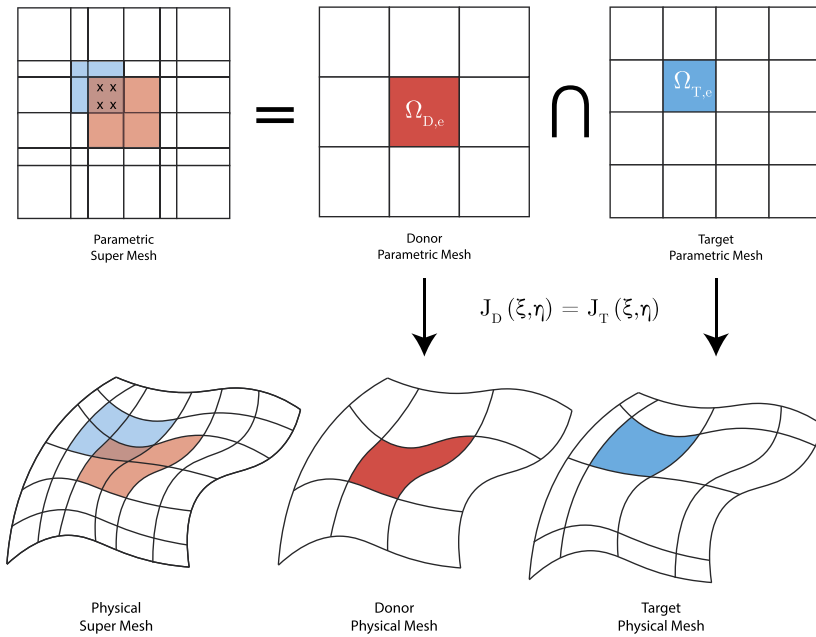


Fig. 2. A parametric supermesh is formed by combining the knot vectors of the two parent patches creating subdomains over which numerical integration can be performed and then mapped through coordinate transformation into the physical space.

2.1. Conservative interpolation by Galerkin projection

When a NURBS patch is refined by knot-insertion or has its degree order-elevated the neither the parametric description of real-space geometry is modified. This property makes NURBS ideal for interpolating information between meshes of different refinements and constructing the projective matrices required for the energy-dependent mesh method. Fig. 2 demonstrates the construction of a parametric super-mesh allowing for the numerical integration of the product of the basis functions from a donor and target mesh. In this study, the projection will be used to transfer the group-to-group scattering source which will be solved using independently refined spatial discretisations. However, the method presented in this work is generally applicable for the projection of any function from one set of basis functions (the donor) onto another (the target).

The parametric supermesh is generated by finding the union of the knot vectors describing the target and donor patches ($\Xi^U = \Xi^T \cup \Xi^D$) which for tensor-product NURBS meshes is a trivial operation. Fig. 2 also demonstrates how the NURBS patches after transformation into real space. Unlike in the work of Baker et al. [2] the physical super mesh is never generated (no further knot-insertion of the target or donor patches is required) reducing the computational time and overall storage requirements.

The following outlines the projection of a function $\phi^D(\mathbf{r})$ described by a set of basis functions $R^D(\mathbf{r})$ of a donor NURBS patch \mathcal{M}^D onto a set of basis functions $R^T(\mathbf{r})$ of a target NURBS patch \mathcal{M}^T . It is assumed throughout that \mathcal{M}^D and \mathcal{M}^T describe the same physical domain and share a common parameter space:

$$S(\mathbf{r}) = \sum_{i=1}^{N^T} R_i^T(\mathbf{r})\mathbf{B}_i^T = \sum_{j=1}^{N^D} R_j^D(\mathbf{r})\mathbf{B}_j^D \quad (6)$$

where \mathbf{B}^T and \mathbf{B}^D are the control points of \mathcal{M}^T and \mathcal{M}^D respectively.

When projecting a function from a donor patch onto a target patch it is desired that:

$$\phi^T(\mathbf{r}) = \phi^D(\mathbf{r}) \quad (7)$$

The weak form of Equation (7) can be generated by multiplying by a set of test functions from the target mesh ($R^T(\mathbf{r})$) and integrating over space.

$$\int_V R_i^T(\mathbf{r})\phi^T(\mathbf{r})dV = \int_V R_i^T(\mathbf{r})\phi^D(\mathbf{r})dV \quad (8)$$

The functions $\phi^T(\mathbf{r})$ and $\phi^D(\mathbf{r})$ can be expanded in terms of the basis functions from the target mesh and the donor mesh respectively by:

$$\phi^T(\mathbf{r}) = \sum_j R_j^T(\mathbf{r})d_j^T \quad \text{and} \quad \phi^D(\mathbf{r}) = \sum_k R_k^D(\mathbf{r})d_k^D \quad (9)$$

where d^T and d^D are the expansion coefficients from the target and donor mesh respectively. These expressions can be substituted into Equation (8) giving:

$$\int_V \sum_j R_i^T(\mathbf{r})R_j^T(\mathbf{r})d_j^TdV = \int_V \sum_k R_i^T(\mathbf{r})R_k^D(\mathbf{r})d_k^DdV \quad (10)$$

Equation (10) can be written in matrix form as:

$$\mathbf{M}^{TT}d^T = \mathbf{M}^{TD}d^D \quad (11)$$

where \mathbf{M}^{TT} is the standard mass-matrix assembled over the target mesh and \mathbf{M}^{TD} is the projective mass matrix from the donor to target mesh. The interpolation will be conservative if the function space that the basis functions span contains the constant function [8]. The individual entries of the projective mass matrix are evaluated by:

$$\mathbf{M}_{ij}^{TD} = \int_V R_i^T(\mathbf{r})R_j^D(\mathbf{r})dV \quad (12)$$

These integrals are evaluated utilising the parametric super-mesh to split the parametric domain into areas for numerical integration. An example of the selection of the quadrature points is shown in Fig. 2. The use of the parametric super-mesh ensures that the numerical integration is performed over C^∞ -continuous functions. The required basis functions for evaluation from both the target and donor meshes can be identified using the mesh connectivity information and the patch knot vectors.

The projection matrix calculated in Equation (12) is needed to project the function described by the basis of the donor mesh onto the basis of the target mesh. In this application the expansion coefficients of flux (d^D) from one energy group can be multiplied by this matrix to compute the scattering source to another energy group. The values of d^T are never explicitly computed. However, if the expansion coefficients of the projected function are required Equation (11) can be solved by:

$$d^T = (\mathbf{M}^{TT})^{-1}\mathbf{M}^{TD}d^D \quad (13)$$

where the calculated expansion coefficients can be used to accelerate the calculation of a refined mesh calculation by supplying a projected initial guess from a previous, coarser mesh calculation.

3. Discretisation of the multigroup, even-parity equations

The multigroup, even-parity form of the time-independent neutron transport equation with isotropic scattering can be written as:

$$-\hat{\Omega} \cdot \nabla \frac{1}{\Sigma_t^g(\mathbf{r})} \hat{\Omega} \cdot \nabla \psi^{+,g}(\mathbf{r}, \hat{\Omega}) + \Sigma_t^g(\mathbf{r})\psi^{+,g}(\mathbf{r}, \hat{\Omega}) = \sum_{g'=1}^G \Sigma_s^{g' \rightarrow g}(\mathbf{r})\phi^{g'}(\mathbf{r}) + q^g(\mathbf{r}) \quad (14)$$

where $\psi^{+,g}(\mathbf{r}, \hat{\Omega})$ is the even-parity angular flux (with units $\text{cm}^{-2} \text{s}^{-1} \text{Sr}^{-1} \text{MeV}^{-1}$), $\phi^g(\mathbf{r})$ is the scalar neutron flux (with units $\text{cm}^{-2} \text{s}^{-1} \text{MeV}^{-1}$), $\Sigma_t^g(\mathbf{r})$ is the macroscopic total cross-section (cm^{-1}), $\Sigma_s^{g' \rightarrow g}$ is the macroscopic differential scattering cross-section (cm^{-1}) and $q^g(\mathbf{r})$ is a prescribed extraneous (or fixed) isotropic source (with units $\text{cm}^{-3} \text{s}^{-1} \text{MeV}^{-1}$) [1].

The full neutron angular flux has been split into its even and odd-parity components using Vladimirov’s transformation [18]:

$$\psi^g(\mathbf{r}, \hat{\Omega}) = \psi^{+,g}(\mathbf{r}, \hat{\Omega}) + \psi^{-,g}(\mathbf{r}, \hat{\Omega}) \tag{15}$$

Isotropic scattering has been assumed for simplicity; the exclusion of more complex scattering models does not impact the analysis of the group dependent mesh method. In this study the even-parity equations have been discretised in angle using spherical harmonics (the P_N method) by approximating the flux by:

$$\psi^{+,g}(\mathbf{r}, \hat{\Omega}) = \sum_{l=0,2,4,\dots}^{\infty} \sum_{m=-2l}^{m=2l} Y_l^m(\hat{\Omega}) \Psi_{l,m}^g(\mathbf{r}) \tag{16}$$

where $Y_l^m(\hat{\Omega})$ are the spherical harmonic basis functions and $\Psi_{lm}^g(\mathbf{r})$ are the corresponding spatially dependent flux moments. In the P_N method the series expansion of the spherical harmonic basis functions is truncated at $l = (N - 1)/2$. The symmetry in angle of the even-parity flux requires only moments with even l to be retained in the expansion. By substituting the expansion of the even-parity flux shown in Equation (16) into Equation (14) and integrating over 4π Steradians (Sr) the angularly discretised even-parity equations are obtained.

$$-\sum_{\alpha} \sum_{\beta} \left(\mathbf{A}_{\alpha\beta} \frac{\partial}{\partial x_{\alpha}} \frac{1}{\Sigma_t^g(\mathbf{r})} \frac{\partial}{\partial x_{\beta}} \right) \Psi^g(\mathbf{r}) + \Sigma_t^g(\mathbf{r}) \Psi^g(\mathbf{r}) = \sum_{g'=1}^G \mathbf{S}^{g' \rightarrow g} \Psi^{g'}(\mathbf{r}) + \mathbf{q}^g(\mathbf{r}) \tag{17}$$

where the indices α and β run across the spatial dimensions (x , y and z in three dimensional Cartesian geometry). $\mathbf{A}_{\alpha\beta}$ is an angular matrix for the combination of spatial dimensions α and β while $\Psi^g = [\psi_{00}, \dots, \psi_{N-1,N-1}]$ is a vector of the corresponding flux moments. The components of the angular matrices are defined by:

$$A_{\alpha\beta}^{lm,l'm'} = \int_{4\pi} Y_l^m(\hat{\Omega}) \Omega_{\alpha} \Omega_{\beta} Y_{l'}^{m'}(\hat{\Omega}) d\hat{\Omega} \tag{18}$$

As only isotropic scattering is included the higher-order terms of the scattering kernel and the external source vector are set to zero by:

$$\mathbf{S}^{g' \rightarrow g} = \text{diag}[\Sigma_{s,0}^{g' \rightarrow g}(\mathbf{r}), 0, \dots, 0] \quad \text{and} \quad \mathbf{q}^g = \text{diag}[q_0^g(\mathbf{r}), 0, \dots, 0]. \tag{19}$$

The even-parity flux moments are discretised using the NURBS basis functions describing the spatial mesh utilising the isoparametric concept and for a given angular moment this is expressed as:

$$\Psi^g(\mathbf{r}) = \sum_{i=1}^N R_i^g(\mathbf{r}) \mathbf{d}_i^g \tag{20}$$

where \mathbf{d}_i^g is a vector of a single spatial expansion coefficients for all moments in the spherical harmonic expansion. As each energy group is discretised using an individual spatial mesh, a unique set of NURBS functions are selected for the discretisation. This is in contrast to traditional methods where a single set of spatial functions is used in the discretisation for all energy groups. By substituting Equation (20) into Equation (17) and following the standard Bubnov–Galerkin method of multiplying by a set of selected test functions, applying the divergence theorem to the streaming term and integrating over space the following weak form is obtained:

$$\sum_{\alpha} \sum_{\beta} \mathbf{A}_{\alpha\beta} \left(\frac{\partial}{\partial x_{\alpha}} \mathbf{u}^g, \frac{1}{\Sigma_t^g} \frac{\partial}{\partial x_{\beta}} \mathbf{v}^g \right) + (\mathbf{u}^g, \Sigma_t^g \mathbf{v}^g) = \sum_{g'=1}^G (\mathbf{u}^g, \mathbf{S}^{g' \rightarrow g} \mathbf{v}^{g'}) + (\mathbf{u}^g, \mathbf{q}^g) \tag{21}$$

where \mathbf{u}^g and \mathbf{v}^g are vectors of the test and shape functions of group g for all moments respectively. $(u, v) \equiv \int_V u(\mathbf{r})v(\mathbf{r})dV$ is defined as the inner product of two functions over the spatial domain. In this work, reflective external boundary conditions are prescribed and are applied by constraining boundary degrees-of-freedom of selected spherical harmonic moments [19].

The evaluation of the first and second terms on the left hand side of Equation (21) yield the standard partial stiffness matrices and the standard mass matrix.

$$K_{A,B}^{g,\alpha\beta} = \left(\frac{\partial R_A^g}{\partial x_{\alpha}}, \frac{1}{\Sigma_t^g} \frac{\partial R_B^g}{\partial x_{\beta}} \right) \tag{22}$$

and

$$M_{A,B}^g = (R_A^g, \Sigma_t^g R_B^g). \tag{23}$$

The fixed source vector is evaluated as:

$$f_A^g = (R_A^g, q^g). \tag{24}$$

However, the evaluation of the term $(u^g, S_{lm}^{g' \rightarrow g} v^{g'})$ requires the evaluation of the product of two basis functions from two different spatial discretisations to form the projective scattering mass matrix required to interpolate the scatter source from one energy group to another.

$$W_{A,B}^{g' \rightarrow g, lm} = (R_A^g, S_{lm}^{g' \rightarrow g} R_B^{g'}) \tag{25}$$

It can be seen that Equation (25) is identical to the form of Equation (12) and can be evaluated in an identical manner as detailed in Section 2.1. For eigenvalue calculations, the projection of the fission source between energy groups can be done in the same manner. The assembled linear system for a single angular moment (lm) of energy group g is given by:

$$\left[\sum_{\alpha} \sum_{\beta} A_{\alpha\beta}^{lm, lm} \mathbf{K}_{\alpha\beta}^g + \mathbf{M}^g \right] \mathbf{d}_{lm}^g = \left[\sum_{l'} \sum_{m'} \sum_{g'=1}^G \mathbf{W}_{l'm'}^{g' \rightarrow g} \mathbf{d}_{l'm'}^{g'} + \mathbf{q}^g - \sum_{l'} \sum_{m'} \sum_{\alpha} \sum_{\beta} A_{\alpha\beta}^{lm, l'm'} \mathbf{K}_{\alpha\beta}^g \mathbf{d}_{l'm'}^g \right] \tag{26}$$

The moments can then be iteratively solved until a prescribed convergence criteria is met via the standard source iteration method. Each of the matrices are assembled and stored in a sparse format using distributed parallelism within the PETSc library [20] and solved using the preconditioned conjugate-gradient (CG) method. As each basis function for a given mesh has compact support the portion of the domain where basis functions from two meshes overlap will also be compact as shown in Fig. 2. For flexibility a multigroup solver class was developed containing a two-dimensional array of pointers to the projective matrices (one entry per group to group pair) used to calculate the scattering source contributions. The use of pointers allows for the same solver to be used for both single-mesh calculations and energy dependent mesh calculations. For example using the same spatial mesh for all the fast energy groups and a different spatial mesh for the thermal energy groups significantly reduces the total amount of system storage required.

In contrast to the standard methodology where a single spatial mesh is shared between all energy groups, the use of unique meshes for each energy groups requires the calculation and storage of up to G^2 projective matrices, where G is the total number of energy groups used in the multigroup approximation. For realistic calculations, that may require in excess of 100 energy groups which then may require the storage of up to 10,000 projective matrices depending on the structure of the scattering kernel. The total number of projective matrices can be reduced by almost a factor of two by taking advantage of the fact that $(\mathbf{M}^{TD})^T = \mathbf{M}^{DT}$. Depending on the future direction of computational hardware development, while the explicit storage of all projective matrices may be prohibitive with contemporary memory systems it may become a feasible approach either to store each matrix or to recompute when required. In calculations containing large multigroup cross-section sets it is likely that many energy groups will produce similar flux profiles. By either heuristically or adaptively banding energy groups into subsets and using an optimised mesh for each, the overall memory requirement could be significantly reduced, at the cost of not obtaining the most efficient spatial meshes.

4. Numerical results

To study the efficiency of the IGA-based conservative interpolation scheme a series of test cases have been studied. Firstly, the scheme implementation and rate-of-convergence are verified using the Method of Manufactured Solutions (MMS) [21]. Secondly, a series of two-group pincell test cases have been studied where the physics of the transport process in the two energy groups is increasingly different to determine the accuracy of interpolating increasingly complex functions. Finally, a seven-group, heterogeneous Uranium-Oxide (UOX) pincell from the OECD/NEA C5G7 benchmark [22] is considered to provide a representative pincell geometry and set of multigroup cross-sections. In this final case the energy groups selected for refinement were chosen adaptively based on the L_2 error calculated using a highly refined reference solution. Pincell geometries were built using five, bi-quadratic NURBS patches. At each level of refinement, the same uniform knot vectors have been used in both parametric dimensions.

4.1. Method of manufactured solutions

The MMS has been used to verify the code implementation and to determine the rates of convergence of the scheme. For simplicity, a two-group problem with a P_1 (diffusion) angular approximation has been considered. The two-group neutron diffusion equation can be written in matrix form as:

$$\begin{bmatrix} -D^1 \nabla^2 + \Sigma_r^1 & -\Sigma_s^{2 \rightarrow 1} \\ -\Sigma_s^{1 \rightarrow 2} & -D^2 \nabla^2 + \Sigma_r^2 \end{bmatrix} \begin{bmatrix} \phi^1(x, y) \\ \phi^2(x, y) \end{bmatrix} = \begin{bmatrix} s^1(x, y) \\ s^2(x, y) \end{bmatrix} \tag{27}$$

Table 1
Rates of convergence of the L_2 error with respect to h where $\mathcal{M}_2 = \mathcal{M}_1$.

Polynomial order	Cartesian mesh		Pincell mesh	
	Rate (Group 1)	Rate (Group 2)	Rate (Group 1)	Rate (Group 2)
1	2.0002	2.0000	N/A	N/A
2	2.9985	2.9943	3.0045	2.9943
3	3.9989	3.9957	4.0148	3.9938
4	4.9993	4.9973	4.9963	4.9972

Table 2
Rates of convergence of the L_2 error with respect to h where $\mathcal{M}_2 = 2 * \mathcal{M}_1$.

Polynomial order	Cartesian mesh		Pincell mesh	
	Rate (Group 1)	Rate (Group 2)	Rate (Group 1)	Rate (Group 2)
1	2.0002	2.0000	N/A	N/A
2	2.9985	2.9986	3.0045	3.0009
3	3.9989	3.9989	4.0148	4.0073
4	4.9993	4.9989	4.9963	5.0026

where D^g is the neutron diffusion coefficient (cm) and $\Sigma_r^g = \Sigma_t^g - \Sigma_s^{g \rightarrow g}$ is the macroscopic removal cross-section (cm^{-1}). For simplicity, material cross-sections have been chosen such that D^1 , D^2 , Σ_r^1 and Σ_r^2 are unity. Only downscatter is considered leading $\Sigma_s^{2 \rightarrow 1}$ to be set to zero. The resulting set of equations are therefore:

$$\begin{bmatrix} -\nabla^2 + 1 & 0 \\ -1 & -\nabla^2 + 1 \end{bmatrix} \begin{bmatrix} \phi^1(x, y) \\ \phi^2(x, y) \end{bmatrix} = \begin{bmatrix} s^1(x, y) \\ s^2(x, y) \end{bmatrix} \quad (28)$$

The MMS tests were performed for two cases, one where the donor group is used as part of the source term for a more complicated target and vice versa. In both test cases a Cartesian mesh and pincell-like mesh have been considered. The rates of convergence are compared to the mesh size h which is taken to be $1/\sqrt{N}$ where N is the total number of spatial elements in the mesh.

4.1.1. Case 1

In the first MMS test case solutions of the form:

$$\phi^1(x, y) = \cos(\pi x) \cos(\pi y) + 2 \quad (29)$$

and

$$\phi^2(x, y) = \cos(2\pi x) \cos(2\pi y) + 2 \quad (30)$$

are manufactured giving rise to the following source terms:

$$s^1(x, y) = (2\pi^2 + 1) \cos(\pi x) \cos(\pi y) + 2 \quad (31)$$

and

$$s^2(x, y) = (8\pi^2 + 1) \cos(2\pi x) \cos(2\pi y) - \cos(\pi x) \cos(\pi y). \quad (32)$$

The source terms are evaluated at the same quadrature points (12 per parametric direction per element) used for numerical integration of the shape and test functions. The form of $\phi^2(x, y)$ was chosen to have a frequency twice that of $\phi^1(x, y)$ therefore it is expected that a mesh twice as refined ($\mathcal{M}_2 = 2 * \mathcal{M}_1$) would yield an identical discretisation error for both groups.

Table 1 details the obtained rates-of-convergence of the case where both energy groups use a single spatial mesh ($\mathcal{M}_2 = \mathcal{M}_1$) while Table 2 details the obtained rates-of-convergence for the case where the second energy group utilises a mesh twice as refined as the first group ($\mathcal{M}_2 = 2 * \mathcal{M}_1$). It can be seen that for both cases the observed rates-of-convergence are very close to the theoretical values $(p + 1)$ for both the Cartesian meshes and pincell-type meshes. Fig. 3 plots the combined L_2 error for both energy groups of the MMS problem against the total number of spatial elements in both energy groups. It can be seen that the scheme that utilises a more refined mesh for the second energy group requires fewer total elements to obtain the same level of accuracy.

4.1.2. Case 2

A second MMS test case was studied where the manufactured solution for the first energy group (the donor mesh) was more complex than the second energy group (the target mesh). Solutions of the form:

$$\phi^1(x, y) = \cos(2\pi x) \cos(2\pi y) + 2 \quad (33)$$

and

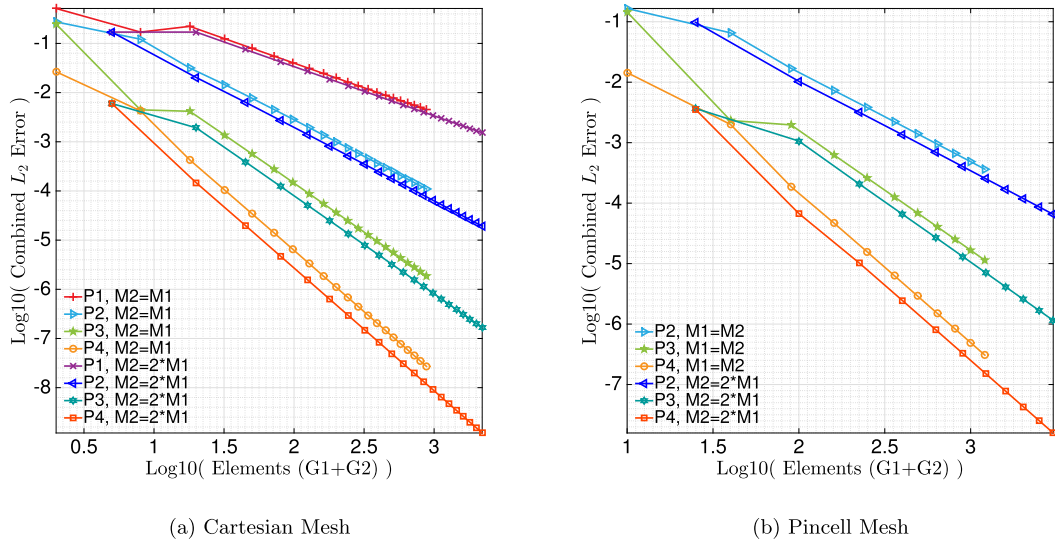


Fig. 3. Convergence of the combined L_2 error against the total number of elements for the first MMS test case.

Table 3

Rates of convergence of the L_2 error with respect to h where $\mathcal{M}_1 = \mathcal{M}_2$.

Polynomial order	Cartesian mesh		Pincell mesh	
	Rate (Group 1)	Rate (Group 2)	Rate (Group 1)	Rate (Group 2)
1	2.0000	2.0002	N/A	N/A
2	2.9943	2.9986	2.9943	3.0048
3	3.9957	3.9989	3.9938	4.0148
4	4.9973	4.9993	4.9972	4.9963

Table 4

Rates of convergence of the L_2 error with respect to h where $\mathcal{M}_1 = 2 * \mathcal{M}_2$.

Polynomial order	Cartesian mesh		Pincell mesh	
	Rate (Group 1)	Rate (Group 2)	Rate (Group 1)	Rate (Group 2)
1	2.0000	2.0002	N/A	N/A
2	2.9986	2.9985	3.0009	3.0045
3	3.9989	3.9989	4.0073	4.0148
4	4.9990	4.9989	5.0026	4.9963

$$\phi^2(x, y) = \cos(\pi x) \cos(\pi y) + 2 \tag{34}$$

are manufactured giving rise to the following source terms:

$$s^1(x, y) = (8\pi^2 + 1) \cos(\pi x) \cos(\pi y) + 2 \tag{35}$$

and

$$s^2(x, y) = (2\pi^2 + 1) \cos(2\pi x) \cos(2\pi y) - \cos(2\pi x) \cos(2\pi y). \tag{36}$$

Table 3 details the obtained rates of convergence for the case where $\mathcal{M}_2 = \mathcal{M}_1$ while Table 4 details the obtained rates of convergence when $\mathcal{M}_1 = 2 * \mathcal{M}_2$. Again, it was observed that in both cases the obtained rates-of-convergence are close to the theoretical value of $p + 1$. Fig. 4 plots the convergence of the total L_2 error against the total number of elements used for both energy groups for the second MMS test case. The L_2 error of $\phi^2(x, y)$ is not a perfect indicator of the accuracy of the function interpolation as the construction of the source term $s^2(x, y)$ depends on $\phi^1(x, y)$. However, the reduction in total elements (in both the target and donor meshes) indicates that for this specific problem it is more computationally efficient to utilise meshes of different refinement than a single spatial mesh.

4.2. Two-group pincell test cases

A series of two-group pincell test cases have been selected to study the effect of interpolating increasingly complex functions between a donor and target mesh. A total of eleven test cases are considered, the first six with increasingly

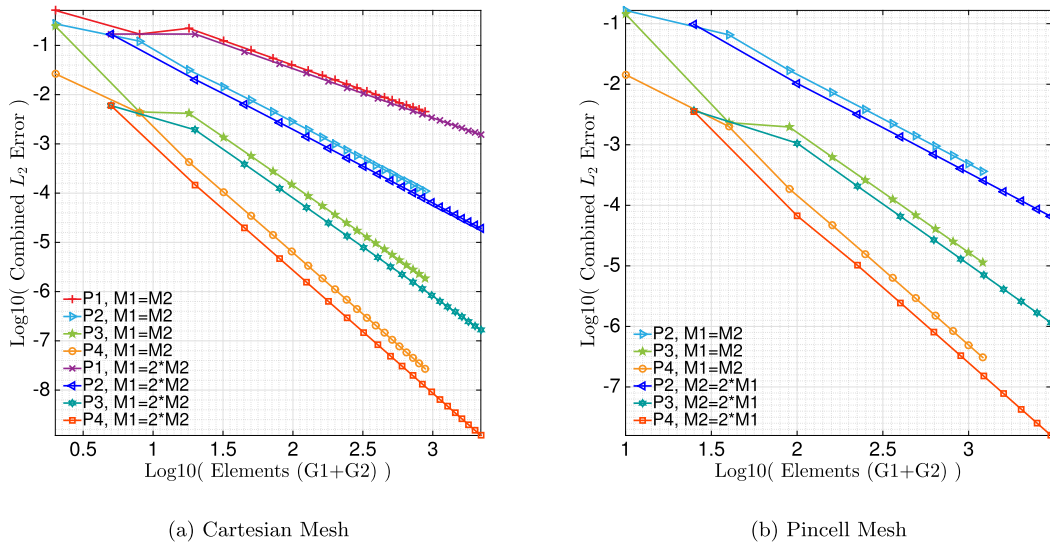


Fig. 4. Convergence of the combined L_2 error against the total number of elements for the second MMS test case.




Fig. 5. Geometry of the pincell used for the two-group study.

Table 5
Pincell macroscopic nuclear cross-sections (Cases 1–6).

Case	$\Sigma_{T,mod}^1$	$\Sigma_{S,mod}^{1 \rightarrow 1}$	$\Sigma_{T,fuel}^1$	$\Sigma_{S,fuel}^{1 \rightarrow 1}$	q^1	$\Sigma_{T,mod}^2$	$\Sigma_{S,mod}^{2 \rightarrow 2}$	$\Sigma_2^{T,fuel}$	$\Sigma_{S,fuel}^{2 \rightarrow 2}$	$\Sigma_{S,fuel}^{1 \rightarrow 2}$	q^2
1	1.0	0.95	2.0	1.9	1.0	1.0	0.95	2.0	1.9	1.0	0.0
2	1.0	0.95	2.0	1.9	1.0	1.0	0.95	5.0	4.75	1.0	0.0
3	1.0	0.95	2.0	1.9	1.0	1.0	0.95	10.0	9.5	1.0	0.0
4	1.0	0.95	2.0	1.9	1.0	1.0	0.95	20.0	19.0	1.0	0.0
5	1.0	0.95	2.0	1.9	1.0	1.0	0.95	50.0	47.5	1.0	0.0
6	1.0	0.95	2.0	1.9	1.0	1.0	0.95	100.0	95.0	1.0	0.0

complicated target functions and a final five with increasingly complicated donor functions. The pincell geometry consisted of a circular pin of radius 1.0 cm centred inside a cell of width 4.0 cm as shown in Fig. 5. The material properties for the first six test cases are detailed in Table 5 while the material properties for the final five test cases are shown in Table 6. As the second energy group has no fixed-source an error will be introduced into the solution if the scatter source from the first energy group is either not resolved or not accurately projected. A P_{11} angular approximation was used to capture transport phenomena while second-order ($p = 2$) NURBS basis functions have been chosen for the spatial discretisation.

Fig. 6 plots the convergence of the L_2 error of the first six two-group pincell test cases studied. It can be seen that when the solutions of the two energy groups are similar (for the case when $\Sigma_{T,fuel}^2 = 2$) to obtain an accurate solution in the second energy group sufficient resolution of the source term (by resolving the first energy group) is required. As $\Sigma_{T,fuel}^2$ is increased and the solution function of the second energy group becomes significantly more challenging to represent it can

Table 6
Pincell macroscopic nuclear cross-sections (Cases 7–11).

Case	$\Sigma_{T,mod}^1$	$\Sigma_{S,mod}^{1 \rightarrow 1}$	$\Sigma_{T,fuel}^1$	$\Sigma_{S,fuel}^{1 \rightarrow 1}$	q^1	$\Sigma_{T,mod}^2$	$\Sigma_{S,mod}^{2 \rightarrow 2}$	$\Sigma_2^{T,fuel}$	$\Sigma_{S,fuel}^{2 \rightarrow 2}$	$\Sigma_{S,fuel}^{1 \rightarrow 2}$	q^2
7	1.0	0.95	5.0	4.75	1.0	1.0	0.95	2.0	1.9	1.0	0.0
8	1.0	0.95	10.0	9.5	1.0	1.0	0.95	2.0	1.9	1.0	0.0
9	1.0	0.95	20.0	19.0	1.0	1.0	0.95	2.0	1.9	1.0	0.0
10	1.0	0.95	50.0	47.5	1.0	1.0	0.95	2.0	1.9	1.0	0.0
11	1.0	0.95	100.0	95.0	1.0	1.0	0.95	2.0	1.9	1.0	0.0

be seen that the coarser approximations to solution of the first energy group can be used without impacting the accuracy of the second energy group.

Fig. 7 plots the convergence of the L_2 error for the final five two-group pincell test cases in which an increasingly complex spatial function is interpolated and used as a source term for a more simple function ($\Sigma_{T,fuel}^2 = 2$). It can be seen that when the donor function is similar in complexity the required number of degrees-of-freedom in the donor mesh is similar to the number required to represent the target function. As the complexity of the donor function is increased it can be seen that fewer degrees-of-freedom are required in the target mesh as the dominant source of errors comes from the representation of the donor function. When $\Sigma_{T,fuel}^1 = 100$ the optimal refinement route to reduce the L_2 error of the target function is just to refine the donor mesh to resolve the complex source term.

4.3. Seven group pincell

In engineering reactor analysis it is common that a significant number of energy groups are required in the multigroup approximation to adequately model the energy dependent physics of the reactor. For this test, the UOX pincell from the OECD/NEA C5G7 LWR benchmark [22] was selected to provide a representative geometry and set of macroscopic nuclear cross-sections for the materials. The pin has a radius of 0.54 cm centred inside a cell of total width 1.26 cm and is described by five, uniformly refined, biquadratic NURBS patches. As analytical solutions to complex, multigroup problems are not readily available a highly refined isogeometric reference was calculated and used to compute the L_2 errors. A P_5 spherical harmonic expansion has been used for the angular discretisation. While in this study a pre-calculated, resolved solution has been used to drive the mesh adaptivity the procedure could be driven by heuristic error metrics or goal-based adaptivity [23].

Fig. 8 plots the convergence of the L_2 errors for all seven energy groups along with the combined L_2 error when a single, common spatial mesh is used for all neutron energy groups. The combined L_2 error (ϵ^c) is defined as:

$$\epsilon^c = \sqrt{\sum_{g=1}^7 \epsilon_g^2} \tag{37}$$

where ϵ_g is the computed L_2 error of group g . Fig. 9 plots the ratio of the L_2 errors for all energy groups against that of the thermal energy group ($g = 7$).

Fig. 10 plots the convergence of the L_2 errors for each energy group and the combined L_2 error compared to the total number of degrees-of-freedom used for all energy groups. Compared to the scheme which used a single spatial mesh for all energy groups it can be seen that by selecting only the energy groups with the highest L_2 errors for refinement the error in all energy groups can be reduced at a similar rate. Fig. 11 plots the proportion of degrees-of-freedom in each energy group as the adaptive procedure is run. As was observed in Fig. 9 the thermal ($g = 7$) group was the least resolved by the single-mesh scheme, while using the adaptive it accounts for 27% of the total degrees-of-freedom. This is in contrast to the fifth energy group which now only uses approximately 6% of the total degrees-of-freedom.

Figs. 12 and 13 plot the convergence of the combined L_2 error and the error in K_{eff} for the adaptively selected refinement scheme respectively. It can be seen that a reduction in the combined L_2 error (approximately 60%) is observed in the results of the simulations. However the obtained K_{eff} is less accurate than the single-mesh scheme. The reduction in the L_2 error is expected as this has been the refinement indicator. The reduction in accuracy for K_{eff} indicates that the L_2 error can not be used to select specific energy groups for refinement to improve accuracy for this quantity of interest (QoI) and goal-based adaptive refinement would be required [9].

Table 7 compares the L_2 errors obtained for each energy group in the single-mesh and adaptive refinement schemes for the same number of total degrees-of-freedom at the final adaptive step compared to the equivalent error for each energy group in the single-mesh scheme for the same number of degrees-of-freedom. It can be observed that for groups one to six the obtained group wise L_2 error is close to the error obtained from the single-mesh scheme. However, the thermal energy group has a noticeably higher L_2 error. This difference is due to the inexact interpolation of the scatter sources from the coarser donor meshes onto the more refined target mesh introducing an error. Figs. 14 and 15 show the final meshes for group 5 and group 7 respectively. It can be seen that the group 7 mesh is significantly more refined than the group 5 mesh as the solution function is more complex and requires more elements to resolve to the same precision as group 5. Each of the five NURBS patches used to describe the pincell geometry have the same number of elements. This results in the four

Test Cases 1-6




Fig. 6. Convergence of the L_2 error of group 2 for the first six test cases. A bold black line has been plotted indicating the refinement path an adaptive procedure may take to select the optimal combination of mesh refinements.

moderator patches being more dense than the single patch used to model the circular pin. In future studies the optimality of the individual energy group meshes will be studied. The group-to-group projection method outlined in this work only requires for each patch to share a common parametric description between parent and donor meshes, allowing for each mesh to be uniquely refined.

Test Cases 7-11

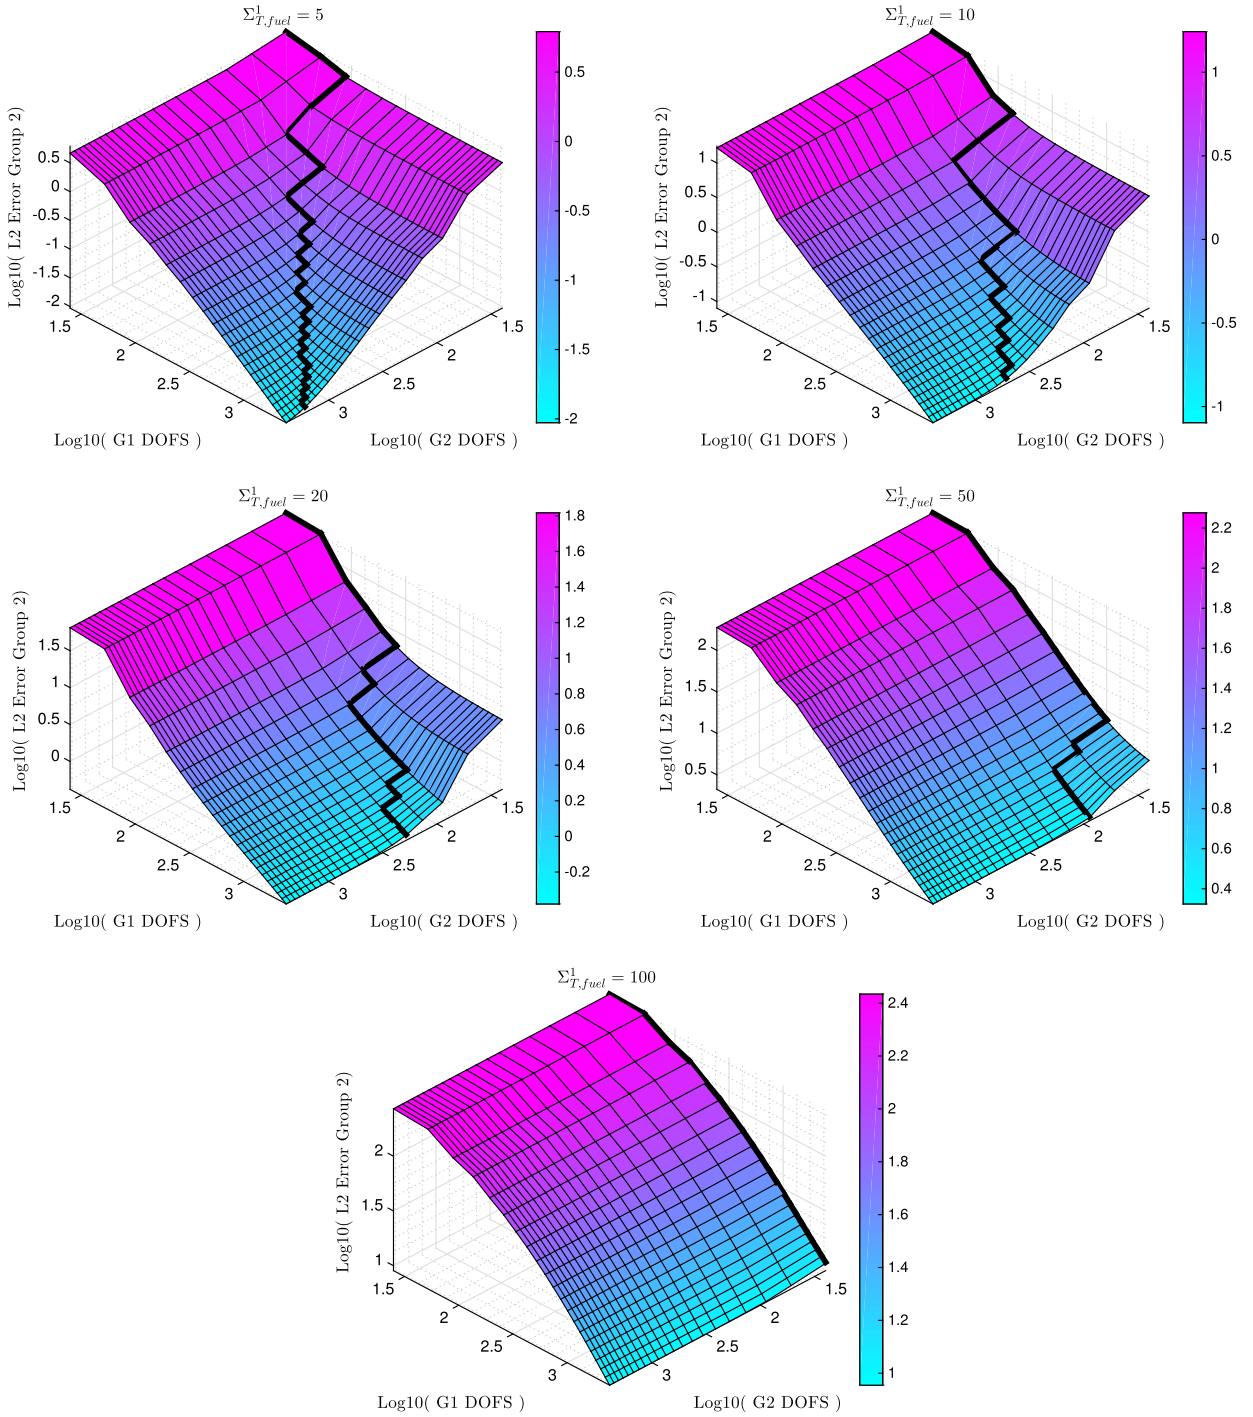


Fig. 7. Convergence of the L_2 error of group 2 for the final five (7–11) test cases. A bold black line has been plotted indicating the refinement path an adaptive procedure may take to select the optimal combination of mesh refinements.

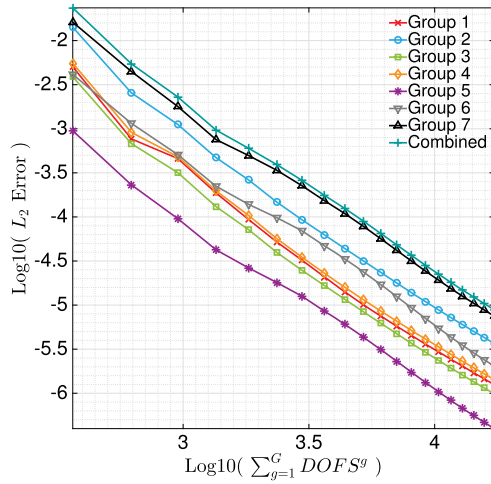


Fig. 8. Convergence of the L_2 error for all seven energy groups of the UOX pincell when each group uses an identical spatial mesh. In addition the total error in all energy groups is also plotted and it can be seen that the total error is dominated by the error in the thermal group (7).




Fig. 9. Convergence of the ratio of the L_2 error for each group compared to the L_2 error in the thermal group (7). It can be seen that the group 2 error is the highest of the remaining energy groups, approximately two times smaller. The fifth energy group has the smallest error, more than 20 times smaller than group 7.




Fig. 10. Convergence of the L_2 error for all seven energy groups and total L_2 error of the UOX pincell when the energy groups are adaptively selected for refinement.




Fig. 11. Evolution of the proportion of the number of degrees-of-freedom utilised in each energy group as the adaptive algorithm proceeds.

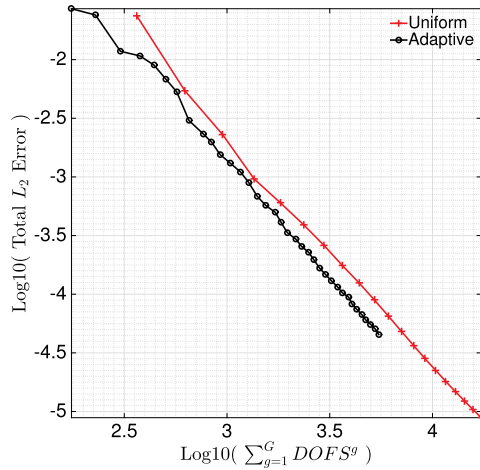


Fig. 12. Convergence of the total L_2 for both the single-mesh and adaptively selected scheme.




Fig. 13. Convergence of k_{eff} for both the single-mesh and adaptively selected scheme.

Table 7
Comparison of individual L_2 error for adaptive and single-mesh refinement schemes.

Group	DOFS (final adaptive step)	L_2 error (adaptive)	L_2 error (single-mesh)
1	628	1.41E–005	1.41E–005
2	1009	1.77E–005	1.78E–005
3	521	1.65E–005	1.66E–005
4	628	1.59E–005	1.59E–005
5	337	1.81E–005	1.80E–005
6	872	1.72E–005	1.71E–005
7	1480	1.98E–005	1.92E–005




Fig. 14. Final mesh for group 5 after adaptive refinement based on the L_2 error.




Fig. 15. Final mesh for group 7 after adaptive refinement based on the L_2 error.

5. Conclusions

A method for energy dependent spatial meshes for the multigroup, second-order even-parity form of the neutron transport equation has been developed and studied based on conservative interpolation by Galerkin projection. Tests performed using the method of manufactured solutions (MMS) have demonstrated that the theoretical rates of convergence are preserved when interpolating source terms between different mesh and basis function representations. The method was then applied to a series of synthetic, two-group reactor pincell calculations where it was observed that a significant reduction in the total number of degrees-of-freedom is obtainable if the variation in the spatial dependence of the two solution functions

is large. Finally, the method was applied to the seven group, UOX pincell from the OECD/NEA C5G7 LWR benchmark. It was demonstrated that when the meshes for each energy group were selected for refinement adaptively then a significant reduction in the total number of degrees-of-freedom for the same total L_2 error can be obtained. However, it was found that the computed eigenvalue was less accurate than for the single-mesh scheme. The use of the L_2 error to adaptively select the energy groups for refinements under resolves energy groups important to accurate calculation of the eigenvalue. Using goal-based error metrics energy groups could be selected for refinement based on the chosen quantity of interest (QoI), for example K_{eff} .

The use of a spatial description based on NURBS patches allows for each energy group to be independently refined while maintaining a shared parametric domain description between all groups. This property allows for efficient implementation and calculation of the projection operations while avoiding the complications of matching different geometric approximations faced by traditional finite element methods. While this paper has used conforming, tensor-product NURBS patches any geometrical description that provides a single parametric domain may be employed, including T-splines or locally-refinable B-splines. The use of a locally-refined spatial mesh for each energy group could potentially yield a significant reduction in the total number of degrees-of-freedom to be solved where localised space and energy dependent features must be resolved.

While the projection operation in this work has been used to interpolate the scattering source between energy groups the algorithm can be used in general to project any function between uniquely refined spatial meshes. A potential extension of this work would be to apply the same method and solve each of the angular flux moments on a uniquely refined spatial mesh. In addition, in industrial analysis there is often the need to pass information between simulation codes solving different physics of the same physical problem. For example, transferring the power profile computed using a neutron transport code to a computational fluid dynamics (CFD) code to determine the temperature profile which can then be used to update the material properties of the neutron transport code [5]. The use of IGA would allow for physics-dependent features such as the boundary layers in the CFD simulation to be resolved without requiring the neutronics mesh to be refined in the affected region while the common parametric description would simplify the solution transfer between codes.

Acknowledgements

J.A. Welch would like to acknowledge the support of EPSRC through the Doctoral Training Award (DTA) PhD scheme and Rolls-Royce for industrial support. A.R. Owens would like to acknowledge the support of EPSRC under their industrial doctorate programme (EPSRC grant number: EP/G037426/1) and also the industrial support of Rolls-Royce. J. Kópházi and M.D. Eaton would like to thank EPSRC for their support through the following grants: Adaptive Hierarchical Radiation Transport Methods to Meet Future Challenges in Reactor Physics (EPSRC grant number: EP/J002011/1) and RADIANT: A Parallel, Scalable, High Performance Radiation Transport Modelling and Simulation Framework for Reactor Physics, Nuclear Criticality Safety Assessment and Radiation Shielding Analyses (EPSRC grant number: EP/K503733/1). Finally, the authors would like to thank the high performance computing (HPC) service support team at Imperial College London, led by HPC manager, Mr. Simon Burbidge. The dataset containing the results derived from this study can be found at <https://zenodo.org/record/815617#.WUrhCRPyv6c>.

References

- [1] E.E. Lewis, W.F. Miller Jr., *Computational Methods of Neutron Transport*, American Nuclear Society, 1993.
- [2] C.M.J. Baker, A.G. Buchan, C.C. Pain, P.E. Farrell, M.D. Eaton, P. Warner, Multimesh anisotropic adaptivity for the Boltzmann transport equation, *Ann. Nucl. Energy* 53 (2013) 411–426.
- [3] S.K. Hall, M.D. Eaton, M.P. Knight, The implementation of a simplified spherical harmonics semi-analytic nodal method in PANTHER, *Ann. Nucl. Energy* 57 (2013) 280–293.
- [4] T.A. Wareing, J.M. McGhee, J.E. Morel, S.D. Pautz, Discontinuous finite element S_N methods on three-dimensional unstructured grids, *Nucl. Sci. Eng.* 138 (2001) 256–268.
- [5] K. Jareteg, P. Vinai, S. Sasic, C. Demaziere, Coupled fine-mesh neutronics and thermal-hydraulics – modeling and implementation for PWR fuel assemblies, *Ann. Nucl. Energy* 84 (2015) 244–257.
- [6] R. Löhner, C. Yang, J. Cebal, J.D. Baum, H. Luo, D. Pelessone, C. Charman, Fluid–structure–thermal interaction using a loose coupling algorithm and adaptive unstructured grids, in: 29th AIAA, Fluid Dynamics Conference, Fluid Dynamics and Co-located Conferences.
- [7] Y. Wang, W. Bangerth, J. Ragusa, Three-dimensional h-adaptivity for the multigroup neutron diffusion equations, *Prog. Nucl. Energy* 51 (2009) 543–555.
- [8] P.E. Farrell, J.R. Maddison, Conservative interpolation between volume meshes by local Galerkin projection, *Comput. Methods Appl. Mech. Eng.* 200 (2010) 89–100.
- [9] M.A. Goffin, C.M.J. Baker, A.G. Buchan, C.C. Pain, M.D. Eaton, P.N. Smith, Minimising the error in eigenvalue calculations involving the Boltzmann transport equation using goal-based adaptivity on unstructured meshes, *J. Comput. Phys.* 242 (2013) 726–752.
- [10] T.J.R. Hughes, J.A. Cottrell, Y. Bazilevs, Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement, *Comput. Methods Appl. Mech. Eng.* 194 (2005) 4135–4195.
- [11] Y. Bazilevs, V.M. Calo, J.A. Cottrell, J.A. Evans, T.J.R. Hughes, S. Lipton, M.A. Scott, T.W. Sederberg, Isogeometric analysis using T-splines, *Comput. Methods Appl. Mech. Eng.* 199 (2010) 229–263.
- [12] A. Buffa, G. Sangalli, R. Vázquez, Isogeometric methods for computational electromagnetics: B-spline and T-spline discretizations, *J. Comput. Phys.* 257 (2014) 1291–1320.
- [13] Y. Bazilevs, M.C. Hsu, M.A. Scott, Isogeometric fluid structure interaction analysis with emphasis on non-matching discretizations, and with application to wind turbines, *Comput. Methods Appl. Mech. Eng.* 249 (2012) 28–41.
- [14] P.N. Nielsen, A.R. Gersborg, J. Gravesen, N.I. Pedersen, Discretizations in isogeometric analysis of Navier–Stokes flow, *Comput. Methods Appl. Mech. Eng.* 200 (2011) 3242–3253.

- [15] S.K. Hall, M.D. Eaton, M.M.R. Williams, The application of isogeometric analysis to the neutron diffusion equation for a pincell problem with an analytic benchmark, *Ann. Nucl. Energy* 49 (2012) 160–169.
- [16] J.A. Welch, J. Kópházi, A.R. Owens, M.D. Eaton, Isogeometric analysis for the multigroup neutron diffusion equation with applications in reactor physics, *Ann. Nucl. Energy* 101 (2017) 465–480.
- [17] A.R. Owens, J.A. Welch, J. Kópházi, M.D. Eaton, Discontinuous isogeometric analysis methods for the first-order form of the neutron transport equation with discrete ordinate (S_N) angular discretisation, *J. Comput. Phys.* 315 (2016) 501–535.
- [18] V.S. Vladimirov, *Mathematical Problems in the One-Velocity Theory of Particle Transport*, Tech. Rep. AECL-1661, Atomic Energy of Canada Ltd., Chalk River, Ontario, 1963.
- [19] C.R.E. de Oliveira, An arbitrary geometry finite element method for multigroup neutron transport with anisotropic scattering, *Prog. Nucl. Energy* 18 (1986) 227–236.
- [20] S. Balay, S. Abhyankar, M.F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W.D. Gropp, D. Kaushik, M.G. Knepley, L.C. McInnes, K. Rupp, B.F. Smith, S. Zampini, H. Zhang, H. Zhang, *PETSc Users Manual*, Tech. Rep. ANL-95/11 – Revision 3.7, Argonne National Laboratory, 2016, <http://www.mcs.anl.gov/petsc>.
- [21] K. Salari, P. Knupp, *Code Verification by the Method of Manufactured Solutions*, Tech. Rep. SAND2000-1444, Sandia National Laboratories, 2000.
- [22] E.E. Lewis, M.A. Smith, N. Tsoulfanidis, G. Palmiotti, T.A. Taiwo, R.N. Blomquist, *Benchmark Specification for Deterministic 2-d/3-d MOX Fuel Assembly Transport Calculations Without Spatial Homogenisation (C5G7 MOX)*, Tech. Rep. NEA/NSC/DOC(2001)17, Nuclear Energy Agency, Expert Group on 3-D Radiation Transport Benchmarks Report, Nuclear Science Committee, 2001.
- [23] A.R. Owens, J. Kópházi, J.A. Welch, M.D. Eaton, Energy dependent mesh adaptivity of discontinuous isogeometric discrete ordinate methods with dual weighted residual error estimators, *J. Comput. Phys.* 335 (2017) 352–386.