Journal of Computational Physics 335 (2017) 352–386

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Energy dependent mesh adaptivity of discontinuous isogeometric discrete ordinate methods with dual weighted residual error estimators

A.R. Owens *, J. Kópházi, J.A. Welch, 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 24 August 2016
Received in revised form 6 December 2016
Accepted 18 January 2017
Available online 24 January 2017

Keywords:
Isogeometric analysis
Discrete ordinates
Dual weighted residual
Group-dependent mesh
Adaptive
Discontinuous Galerkin

ABSTRACT

In this paper a hanging-node, discontinuous Galerkin, isogeometric discretisation of the multigroup, discrete ordinates (S_N) equations is presented in which each energy group has its own mesh. The equations are discretised using Non-Uniform Rational B-Splines (NURBS), which allows the coarsest mesh to exactly represent the geometry for a wide range of engineering problems of interest; this would not be the case using straight-sided finite elements. Information is transferred between meshes via the construction of a supermesh. This is a non-trivial task for two arbitrary meshes, but is significantly simplified here by deriving every mesh from a common coarsest initial mesh. In order to take full advantage of this flexible discretisation, goal-based error estimators are derived for the multigroup, discrete ordinates equations with both fixed (extraneous) and fission sources, and these estimators are used to drive an adaptive mesh refinement (AMR) procedure. The method is applied to a variety of test cases for both fixed and fission source problems. The error estimators are found to be extremely accurate for linear NURBS discretisations, with degraded performance for quadratic discretisations owing to a reduction in relative accuracy of the "exact" adjoint solution required to calculate the estimators. Nevertheless, the method seems to produce optimal meshes in the AMR process for both linear and quadratic discretisations, and is $\approx \times 100$ more accurate than uniform refinement for the same amount of computational effort for a 67 group deep penetration shielding problem.

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

1. Introduction

Even with today's computational algorithms and high performance computing architectures, the numerical solution of the neutron transport equation over heterogeneous whole-core reactor physics and shielding geometries remains a significant challenge. A large part of this challenge stems from the seven dimensional nature of the solution space, with resolution required in energy, direction, space and time in transient problems. Traditional reactor physics techniques involve initially solving the neutron transport equation over assembly-sized domains with periodic boundary conditions [1,2]. The heterogeneous cross-section data within the fuel assembly is then homogenised, and the resulting whole-core system is solved using the neutron diffusion equation with nodal spatial discretisation techniques [3].

* Corresponding author.
E-mail address: a.owens12@imperial.ac.uk (A.R. Owens).

http://dx.doi.org/10.1016/j.jcp.2017.01.035
0021-9991/© 2017 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

If a more detailed description of the geometry is required, discontinuous Galerkin finite element methods (DGFEM) with a discrete ordinate angular discretisation can be applied. It was originally developed for the first-order form of the discrete ordinates equations [4–6], but has since been applied to a wide range of fields such as radiative heat transfer [7], the compressible Navier–Stokes equations [8] and the Euler equations of gas dynamics [9]. Complex geometries can be subdivided into geometric primitives such as triangles, quadrilaterals, tetrahedra and hexahedra by the use of a mesh generator [10]. This mesh generator converts the NURBS geometry output by the computer aided design (CAD) program into surface and volumetric mesh elements, over which the discrete ordinates equations can then be solved with a DGFEM spatial discretisation. However, most elements employed have straight or planar sides [10], and so cannot exactly represent the underlying NURBS geometry. It is crucial that the polygonal geometry representation preserves the fissile mass of the system in reactor physics applications, otherwise large errors can be introduced into the criticality solution [11,12].

An extension of the finite element method (FEM), isogeometric analysis (IGA), was recently introduced in order to overcome some of these deficiencies [13]. As in FEM, prescribed shape and test functions are used to discretise the weak form of the underlying partial differential equation (PDE). In order to preserve the CAD geometry, the same NURBS used to mathematically describe the geometry are used to discretise the weak form of the PDE. In this manner, the exact geometry output by the CAD program is preserved at the coarsest level of refinement, as well as when the mesh is further refined. A further advantage of IGA is that the parameterisation of the physical space does not change under mesh refinement, simplifying the implementation of the group dependent mesh (GDM) methods presented here.

The first applications of IGA were in Bubnov–Galerkin discretisations of solid mechanics problems and streamline-upwind Petrov–Galerkin discretisations of advection-diffusion equations [13]. In reactor physics, Hall et al. [14] solved the one group diffusion equation over a pincell geometry using a Bubnov–Galerkin NURBS discretisation, which was extended by Welch et al. [15] to multigroup problems over heterogeneous quarter-core style geometries.

Discontinuous IGA methods were first applied to elliptic systems [16,17]. For hyperbolic systems, two discontinuous IGA discretisations have recently been developed. The Blended Isogeometric Discontinuous Galerkin method [18] meshed the geometry using rational Bernstein–Bézier triangles. The solution fields were then approximated using standard polynomial basis functions and the method applied to Maxwell's equations of electromagnetics and the acoustic wave equations.

In [19], discontinuous IGA was applied to the first-order form of the discrete ordinates equations with conforming meshes (i.e. no hanging-nodes). This was extended in [20] to an adaptive, hanging-node formulation for one group problems where the adaptivity was driven by heuristic error indicators. This work is a further extension of that in [20], in which each energy group now has its own associated mesh. This can be advantageous, as in many practical problems the solutions in different energy groups exhibit very different behaviour, and so require mesh resolution in different parts of the geometry [21,22]. This technique was employed by Ragusa and coworkers [21,22] for the multigroup neutron diffusion equation with Cartesian grid geometries, and by Goffin et al. [23] for the first-order form of the neutron transport equation with a spherical harmonics angular discretisation. In [23], the meshes in each group were formed independently of each other, and so the calculation of a "supermesh" to transfer information between groups is a non-trivial task [24]. In contrast, the method presented in [21,22] relies on every mesh being derived from a common coarsest description. This significantly simplifies the generation of the supermesh, but naturally limits a scheme to geometries that can be represented exactly by the elements employed. In [20] it was shown that using an inexact DGFEM geometry to represent circular fuel pins was impractical for use with AMR, as eventually the geometric error dominated the spatial discretisation error. However this is not a limitation faced by NURBS-based IGA, and so the approach of Ragusa et al. of deriving every mesh from a common coarsest description is followed here, with that coarsest description being precisely the geometry of the problem.

In order to take full advantage of the GDM IGA discretisation, goal-based adaptivity based on dual weighted residual (DWR) error estimators will be used here to drive the AMR. The general framework for these error estimators was originally presented by Rannacher and coworkers [25–29]. The specific form of the error estimators for the first order form of the one group discrete ordinates equations with linear DGFEM spatial discretisation were derived by Lathouwers for both fixed source [30] and eigenvalue [31] problems. These estimators are extended here to multigroup problems of both fixed source and eigenvalue varieties and using discontinuous NURBS of arbitrary order for the spatial discretisation. This allows the meshes in each group to be refined towards a specific goal functional of the flux, as well as providing an estimate of the remaining error in the functional due to the spatial discretisation, a useful property in a "best estimate plus uncertainty" design process.

The remainder of the paper is organised as follows. Section 2 gives a brief overview of isogeometric analysis and the basis functions employed. Section 3 derives the spatially discretised version of the weak form of the equations, with a special focus on supermesh calculation and the transfer of information between meshes. In Section 4 the multigroup DWR error estimators are derived for both fixed source and eigenvalue problems, and the adaptive procedure followed here is explained. Section 5 contains numerical results of the error estimators and GDM methodology applied to a variety of verification test cases.

2. Isogeometric analysis

Isogeometric analysis aims to unify the geometric description of physical problems used by CAD programs with that employed in the computational analysis, by using the NURBS prevalent in the CAD community to discretise the governing equations [13]. The design and analysis cycle time can be dominated by procedures associated with mesh generation, par-

ticularly if adaptivity is employed in the spatial mesh [13]. Then the CAD program may have to be revisited at every step of an AMR process to improve the geometric description, as traditional finite elements cannot exactly represent the geometric primitives such as circles and ellipses that are commonly found in radiation shielding and reactor physics applications. In reactor physics problems in particular, the fissile mass of the system must be conserved, otherwise errors are introduced that are proportional to the error in the fissile mass [11]. Even if the fissile mass is conserved, it was shown in [20] that the polygonal geometry introduces an error in the leakage of neutrons across material boundaries that can dominate at high spatial resolutions.

IGA can exactly represent the aforementioned geometric primitives, as it utilises the same basis functions as the CAD program to discretise the weak form of the S_N equations. The geometry used by the analysis program is therefore exact from the coarsest level of refinement, and all mesh refinement can be performed in the analysis program, bypassing the need for an ancillary mesh generator. In addition, the mapping from the parametric to the physical space does not change under h-refinement [13], which greatly simplifies the construction of the supermesh as described in Section 3.2.

2.1. B-splines

B-spline curves are uniquely determined by three properties [32]:

1. A polynomial order p.
2. A set of n control points in physical space $\{\mathbf{B}_i : i = 1, ..., n\}$, where n is the number of basis functions defining the B-spline curve.
3. A knot vector Ξ, which is a sequence of non-decreasing real numbers. The length of a knot vector is defined to be the number of values, including repeats, in Ξ. Open knot vectors will be used in the remainder of this paper, in which the first and last knot value are repeated $p + 1$ times. In this case, the knot vector will have length $n + p + 1$, $\Xi = (\xi_1, \xi_2, ..., \xi_{n+p+1})$.

The basis functions in parametric space are completely determined by the polynomial order and knot vector. A knot vector is said to be uniform if the knot points are equally spaced, and non-uniform otherwise [32]. Knot points have reduced continuity, such that if a knot point is repeated m times, it is said to have multiplicity m and the basis functions are C^{p-m} continuous at this point. As such, the knot spans between these knot points play the role of elements in FEM over which numerical integration is performed. As the discretisation here is discontinuous, all knot points in the remainder of the paper will have multiplicity $m = p + 1$.

Denoting basis function i of polynomial order p by $N_{i,p}(\xi)$, then the basis functions are given by the following Cox–de Boor recursion formula [33,34] for $p = 0$:

$$N_{i,0}(\xi) = \begin{cases} 1 & \text{if } \xi_i \leq \xi < \xi_{i+1}, \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

and for $p \geq 1$

$$N_{i,p}(\xi) = \frac{\xi - \xi_i}{\xi_{i+p} - \xi_i} N_{i,p-1}(\xi) + \frac{\xi_{i+p+1} - \xi}{\xi_{i+p+1} - \xi_{i+1}} N_{i+1,p-1}(\xi) \tag{2}$$

Equation (2) can contain terms of the form 0/0, which is defined to be zero here. These basis functions form a partition of unity [32]:

$$\sum_{i=1}^{n} N_{i,p}(\xi) = 1 \quad \forall \xi \in [\xi_1, \xi_{n+p+1}] \tag{3}$$

The basis functions are interpolatory at the end points of the knot vector, such that $N_{i,p}(\xi_1) = \delta_{i,1}$ and $N_{i,p}(\xi_{n+p+1}) = \delta_{i,n}$, which reduces the amount of information that must be stored on element faces to perform upwinding of the discrete ordinates equations.

B-spline curves in physical space are then defined by associating each basis function with a control point in physical space and summing their products:

$$\mathbf{C}(\xi) = \sum_{i=1}^{n} N_{i,p}(\xi) \mathbf{B}_i \tag{4}$$

Higher dimensional B-spline objects are constructed from one dimensional B-splines in a tensor product fashion. In two dimensions, polynomial orders p and q, knot vectors $\Xi = (\xi_1, \xi_2, ..., \xi_{n+p+1})$ and $H = (\eta_1, \eta_2, ..., \eta_{m+q+1})$ and a control net $\{\mathbf{B}_{i,j} : i = 1, ...n \text{ and } j = 1, ..., m\}$ are used to construct B-Spline surfaces as follows:

$$\mathbf{S}(\xi, \eta) = \sum_{i=1}^{n} \sum_{j=1}^{m} N_{i,p}(\xi) M_{j,q}(\eta) \mathbf{B}_{i,j} \tag{5}$$

where the $M_{j,q}$ are defined analogously to the $N_{i,p}$ using polynomial order q and knot vector H. The partition of unity and interpolatory properties extend into higher dimensions by the tensor product construction procedure.

2.2. NURBS

In order to exactly represent the conic sections prevalent in reactor physics and shielding applications, NURBS are required. They are constructed by associating each B-spline basis function with a weight, denoted by w_i and $w_{i,j}$ in one and two dimensions respectively. One dimensional NURBS in parametric space are then defined by [32]:

$$R_i^p(\xi) = \frac{N_{i,p}(\xi) w_i}{\sum_{\hat{i}=1}^{n} N_{\hat{i},p}(\xi) w_{\hat{i}}} \tag{6}$$

with the corresponding NURBS curve in physical space given by:

$$\mathbf{C}(\xi) = \sum_{i=1}^{n} R_i^p(\xi) \mathbf{B}_i \tag{7}$$

Two dimensional NURBS in parametric space are defined by:

$$R_{i,j}^{p,q}(\xi,\eta) = \frac{N_{i,p}(\xi) M_{j,q}(\eta) w_{i,j}}{\sum_{\hat{i}=1}^{n} \sum_{\hat{j}=1}^{m} N_{\hat{i},p}(\xi) M_{\hat{j},q}(\eta) w_{\hat{i},\hat{j}}} \tag{8}$$

with the corresponding physical surface defined by:

$$\mathbf{S}(\xi,\eta) = \sum_{i=1}^{n} \sum_{j=1}^{m} R_{i,j}^{p,q}(\xi,\eta) \mathbf{B}_{i,j} \tag{9}$$

The rational part of the name NURBS comes from the division by B-splines in Equations (6) and (8). As such, they are not piecewise polynomial unless all of the weights are equal, in which case the denominators in Equations (6) and (8) are constant by the partition of unity property. The terms “polynomial order”, “order” and “degree” will be used throughout the remainder of the paper with these rational basis functions to refer to the order of the underlying B-splines from which they are constructed.

3. Discrete ordinates equations & weak form

3.1. Weak form derivation

In the following derivation vector quantities are in bold and matrices are double underlined. The steady-state, multigroup, discrete ordinates equations with isotropic scattering and source terms are given by:

$$\boldsymbol{\Omega}_k^g \cdot \nabla \psi_k^g(\mathbf{r}) + \Sigma_t^g(\mathbf{r}) \psi_k^g(\mathbf{r}) = \frac{1}{4\pi} \sum_{g'=1}^{G} \Sigma_s^{g' \to g}(\mathbf{r}) \phi^{g'}(\mathbf{r}) + Q^g(\mathbf{r}) \quad \text{for } k = 1, \ldots, N \text{ and } g = 1, \ldots, G \tag{10a}$$

$$\phi^g(\mathbf{r}) = \sum_{k=1}^{N} w_k \psi_k^g(\mathbf{r}) \tag{10b}$$

for a discretisation with G groups and N directions. The pairs $\{\boldsymbol{\Omega}_k, w_k\}_{k=1,\ldots,N}$ are the discrete ordinates directions and their corresponding weights with $\sum_{k=1}^{N} w_k = 4\pi$. ψ_k^g is the angular flux in group g in direction $\boldsymbol{\Omega}_k$ and ϕ^g is the scalar flux in group g. Σ_t^g is the macroscopic total cross-section in group g and $\Sigma_s^{g' \to g}$ is the macroscopic scattering cross-section from group g' to group g. Equation (10a) is solved over a domain $\mathbf{r} \in V$ with boundary ∂V. Boundary conditions are prescribed as either a vacuum or reflective for incoming directions as:

$$\psi_k^g(\mathbf{r}) = 0 \qquad \text{for} \quad \mathbf{r} \in \partial V_D \quad \text{and} \quad \boldsymbol{\Omega}_k \cdot \mathbf{n} < 0 \tag{11a}$$

$$\psi_k^g(\mathbf{r}) = \psi_{k'}^g(\mathbf{r}) \qquad \text{for} \quad \mathbf{r} \in \partial V_R \quad \text{and} \quad \boldsymbol{\Omega}_k \cdot \mathbf{n} < 0 \tag{11b}$$

where $\boldsymbol{\Omega}_{k'}$ is the reflective direction corresponding to $\boldsymbol{\Omega}_k$ on this boundary. This direction is assumed to be in the angular quadrature set, which is the case for axially-invariant quadrature sets and boundaries aligned with the axes.

The source term Q^g depends on whether the problem is a fixed source or eigenvalue calculation in the following manner:

$$Q^g(\mathbf{r}) = \begin{cases} \dfrac{S^g(\mathbf{r})}{4\pi} & \text{for fixed source problems} \\ \dfrac{\lambda\chi^g}{4\pi}\displaystyle\sum_{g'=1}^{G}\nu^{g'}(\mathbf{r})\Sigma_f^{g'}(\mathbf{r})\phi^{g'}(\mathbf{r}) & \text{for eigenvalue problems} \end{cases} \tag{12}$$

where S^g is a fixed extraneous source strength in group g, ν^g is the average number of neutrons produced per fission in group g, Σ_f^g is the macroscopic fission cross-section in group g and χ^g is the probability of a fission neutron being produced into group g. $\lambda = \frac{1}{k_{\text{eff}}}$ is the inverse of the effective multiplication factor, which is used in lieu of k_{eff} here as the error estimators are derived in terms of the eigenvalue λ.

Equation (10a) is solved over a domain V which must be partitioned into elements. Each group has its own mesh, and so there are G such partitions with $\bigcup_{e=1}^{K^g} V_e^g = V$ for $g = 1, ..., G$ where mesh g contains K^g elements. In IGA these elements are represented by NURBS of order p, rather than the polynomials used in traditional DGFEM. This allows the exact representation of geometric primitives such as circles that are prevalent in reactor physics and shielding applications at the coarsest level of refinement, which will be key to exploiting the adaptivity in Section 4.3. Note that these meshes can contain hanging-nodes in the same manner as those in [20].

As in traditional DGFEM, Equation (10a) is multiplied by a test function $v^g(\mathbf{r}) \in H^1(V_e^g)$ for some element V_e^g. The resulting equation would then usually be integrated over a single element, but it is more convenient here to integrate over the entire domain V in order to represent the source term compactly. Note that as the support of v^g is a single element in mesh g, this will result in the same system of equations.

$$\int_V d\mathbf{r}\ \mathbf{\Omega}_k^g \cdot \nabla\psi_k^g(\mathbf{r})v^g(\mathbf{r}) + \int_V d\mathbf{r}\ \Sigma_t^g(\mathbf{r})\psi_k^g(\mathbf{r})v^g(\mathbf{r}) = \frac{1}{4\pi}\sum_{g'=1}^{G}\int_V d\mathbf{r}\ \Sigma_s^{g'\to g}(\mathbf{r})\phi^{g'}(\mathbf{r})v^g(\mathbf{r}) + \int_V d\mathbf{r}\ Q^g(\mathbf{r})v^g(\mathbf{r}) \tag{13}$$

for $k = 1, ..., N$ and $g = 1, ..., G$. The angular and scalar flux in each group is then expanded in terms of the NURBS functions $R_B^g(\mathbf{r})$ defined over mesh g. Defining nnp_g to be the number of basis functions defined over mesh g, these expansions are given by:

$$\psi_k^g(\mathbf{r}) = \sum_{B=1}^{nnp_g}\psi_{k,B}^g R_B^g(\mathbf{r}) \tag{14a}$$

$$\phi^g(\mathbf{r}) = \sum_{B=1}^{nnp_g}\phi_B^g R_B^g(\mathbf{r}) \quad \text{where} \tag{14b}$$

$$\phi_B^g = \sum_{k=1}^{N} w_k\psi_{k,B}^g \quad \text{with} \tag{14c}$$

$$v^g(\mathbf{r}) = R_A^g(\mathbf{r}) \quad \text{for } A = 1, ..., nnp_g \text{ and } g = 1, ..., G \tag{14d}$$

This solution space of all discontinuous NURBS of order p defined over all meshes, i.e. $\text{span}\{R_A^g \in H^1(V_e^g) : A = 1, ..., nnp_g,\ g = 1, ..., G\}$ is defined to be $S_{h,p}$. Equation (14) can be written in column vector notation as:

$$\psi_k^g(\mathbf{r}) = \boldsymbol{\psi}_k^g \cdot \mathbf{R}^g(\mathbf{r}) \tag{15a}$$

$$\phi^g(\mathbf{r}) = \boldsymbol{\phi}^g \cdot \mathbf{R}^g(\mathbf{r}) \tag{15b}$$

The right hand side (RHS) of equation (13) can then be written as:

$$\text{RHS} = \frac{1}{4\pi}\sum_{g'=1}^{G}\left(\int_V d\mathbf{r}\ \Sigma_s^{g'\to g}\mathbf{R}^g(\mathbf{R}^{g'})^T\right)\boldsymbol{\phi}^{g'} + \begin{cases} \dfrac{1}{4\pi}\displaystyle\int_V d\mathbf{r}\ S^g\mathbf{R}^g \\ \dfrac{\lambda\chi^g}{4\pi}\displaystyle\sum_{g'=1}^{G}\left(\int_V d\mathbf{r}\ \nu^{g'}\Sigma_f^{g'}\mathbf{R}^g(\mathbf{R}^{g'})^T\right)\boldsymbol{\phi}^{g'} \end{cases} \tag{16}$$

Define the fixed source vector and the (generally rectangular) scatter and fission mixed mass matrices [24] to be:

$$\mathbf{Q}^g = \frac{1}{4\pi}\int_V d\mathbf{r}\ S^g\mathbf{R}^g \tag{17a}$$

$$\underline{\underline{S}}^{g'\to g} = \frac{1}{4\pi}\int_V d\mathbf{r}\ \Sigma_s^{g'\to g}\mathbf{R}^g(\mathbf{R}^{g'})^T \tag{17b}$$

$$\implies \left(\underline{\underline{S}}^{g'\to g}\right)_{ij} = \frac{1}{4\pi}\int_V d\mathbf{r}\ \Sigma_s^{g'\to g} R_i^g(\mathbf{r}) R_j^{g'}(\mathbf{r}) \quad \text{and} \tag{17c}$$

$$\underline{\underline{F}}^{g'\to g} = \frac{\chi^g}{4\pi}\int_V d\mathbf{r}\ \nu^{g'}\Sigma_f^{g'}\mathbf{R}^g(\mathbf{R}^{g'})^T \tag{17d}$$

$$\implies \left(\underline{\underline{F}}^{g'\to g}\right)_{ij} = \frac{\chi^g}{4\pi}\int_V d\mathbf{r}\ \nu^{g'}\Sigma_f^{g'} R_i^g(\mathbf{r}) R_j^{g'}(\mathbf{r}) \tag{17e}$$

where the superscript $g' \to g$ has been chosen to represent the transfer of information from mesh g' to mesh g and $(\underline{\underline{A}})_{ij}$ is the matrix element in row i and column j. As the left hand side (LHS) of Equation (13) does not contain any coupling between group meshes it can be dealt with element-wise. The divergence theorem is applied to the streaming term and the resulting boundary integral is upwinded based on the sign of $\mathbf{\Omega}_k \cdot \mathbf{n}$ where $\mathbf{n}$ is the outward unit normal to the element boundary. For a full treatment of this procedure with a hanging-node NURBS mesh, see [20]. The resulting system within an element V_e^g for the LHS is given by:

$$\begin{aligned} LHS = &-\int_{V_e^g} d\mathbf{r}\ \psi_k^g \mathbf{\Omega}_k \cdot \nabla v^g + \int_{V_e^g} d\mathbf{r}\ \Sigma_t^g \psi_k^g v^g + \int_{\partial V_e^{g,+}} dS\ \psi_k^{g,int} v^g \mathbf{\Omega}_k \cdot \mathbf{n} + \ldots \\ &\ldots + \int_{\partial V_e^{g,-}\setminus\partial V} dS\ \psi_k^{g,ext} v^g \mathbf{\Omega}_k \cdot \mathbf{n} + \int_{\partial V_e^{g,-}\cap\partial V_R} dS\ \psi_{k'}^g v^g \mathbf{\Omega}_k \cdot \mathbf{n} \end{aligned} \tag{18}$$

where *int* and *ext* are used to represent flux from within the element and from the adjacent element respectively. In traditional polynomial DGFEM, the element faces ∂V_e^g are typically the union of straight lines, or are approximated as such [6] so that the sign of $\mathbf{\Omega}_k \cdot \mathbf{n}$ is constant over an element face. It was demonstrated in [19] that this approach truncates the attainable order of accuracy with the strongly curved boundaries of the NURBS elements, and instead the integration over partial element faces must be explicitly performed using quadrature points along each element face. This incurs a small performance penalty, increasing the sweep time by $\approx 5\%$ for quadratic elements.

Summing these contributions over all the elements in mesh g results in the streaming plus removal matrix $\underline{\underline{L}}_k^g$. Combining this with the notation from Equations (15) and (17), Equation (13) can be rewritten as the linear system:

$$\underline{\underline{L}}_k^g \boldsymbol{\psi}_k^g = \sum_{g'=1}^{G} \underline{\underline{S}}^{g'\to g}\boldsymbol{\phi}^{g'} + \begin{cases} \mathbf{Q}^g \\ \lambda \sum_{g'=1}^{G} \underline{\underline{F}}^{g'\to g}\boldsymbol{\phi}^{g'} \end{cases} \qquad \text{for } k = 1, \ldots, N \text{ and for } g = 1, \ldots, G \tag{19}$$

Equation (19) can then be solved for the $\boldsymbol{\psi}_k^g$ by sweeping through the elements of mesh g for each ordinate direction k in an ordering determined by the graph theory [20]. As it was found in [19] that solving for each cycle before continuing the sweep was computationally inefficient, each element is solved exactly once per sweep, as in [20].

Note that in the special case that the mesh in every group is the same, $\underline{\underline{S}}^{g'\to g}$ and $\underline{\underline{F}}^{g'\to g}$ reduce to standard (square) mass matrices weighted by the appropriate cross-sections. This is also trivially the case in one group problems.

3.2. Mixed mass matrices

As the mixed mass matrices $\underline{\underline{S}}^{g'\to g}$ and $\underline{\underline{F}}^{g'\to g}$ are exclusively used for matrix-vector products, a compressed-row storage (CRS) scheme makes sense as then the sparse matrix-vector product is both simple to implement and parallelise if necessary [35]. At first glance it appears that up to $2G^2$ mixed mass matrices must be stored for eigenvalue problems, potentially introducing a prohibitive memory overhead even when CRS is used. However, this can be reduced significantly in a variety of ways.

For eigenvalue problems, only a single mixed mass matrix:

$$\underline{\underline{M}}^{g'\to g} = \int_V d\mathbf{r}\ \mathbf{R}^g(\mathbf{R}^{g'})^T \tag{20}$$

needs to be stored for each ordered pair of groups (g', g). The inclusion of the element-wise constant functions $\Sigma_s^{g'\to g}$, $\nu^{g'}$ and $\Sigma_f^{g'}$ in the integrands of Equation (17) can be incorporated at the time of the source calculation by suitable multiplication of the elements of the scalar flux vector $\boldsymbol{\phi}^{g'}$.

Furthermore, the matrices $\underline{\underline{S}}^{g'\to g}$ and $\underline{\underline{F}}^{g'\to g}$ may be identically zero for some ordered pairs of groups (g', g), for example if there is no scattering from group g' to group g. If both $\underline{\underline{S}}^{g'\to g}$ and $\underline{\underline{F}}^{g'\to g}$ are zero for a group pair (g', g) then there is

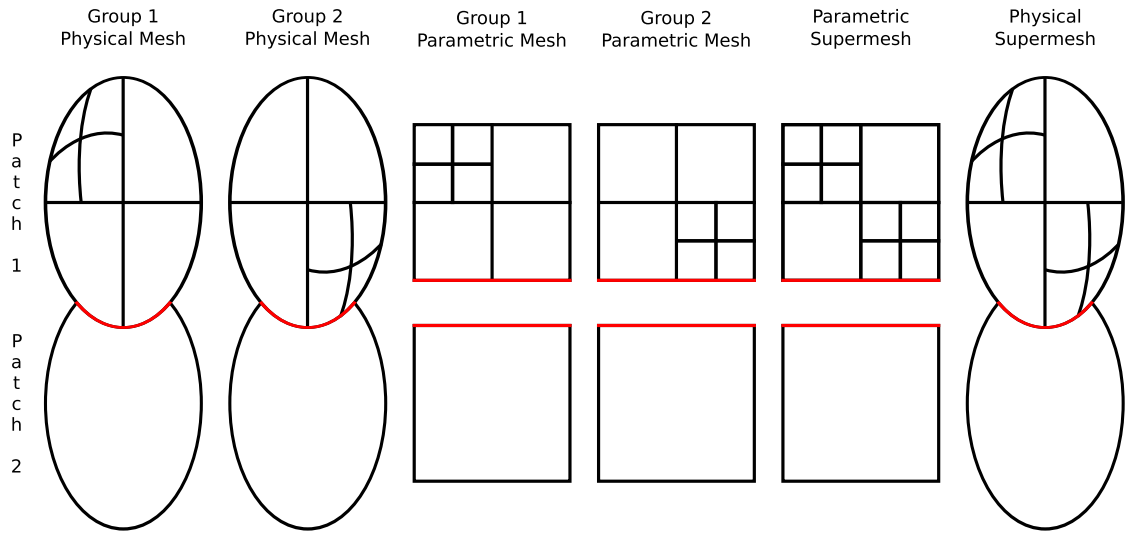


Fig. 1. An illustration of the supermesh generation procedure for a geometry consisting of 2 ellipses initially represented by patch 1 and patch 2, which share an edge highlighted in red. The problem of finding the supermesh clearly decomposes into finding supermeshes over the original patches that represented the geometry, and so no work is required in this example in patch 2. In patch 1 the parametric supermesh is simply the intersection of the parametric mesh in each group, and the physical supermesh is then given by this parametric supermesh mapped into the physical space. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

no need to store $\underline{\underline{M}}^{g' \to g}$. This could be of particular use in shielding problems that have a large number of groups but with purely or mostly downscatter between them.

Finally, it is clear from Equation (20) that $\underline{\underline{M}}^{g' \to g} = (\underline{\underline{M}}^{g \to g'})^T$. Therefore a further reduction in storage could be achieved by storing only $\underline{\underline{M}}^{g' \to g}$ for $g' \geq g$. However, this would then require the sparse matrix-vector product of the transpose of a CRS matrix which will naturally incur a performance penalty [35].

Of these possibilities, only the first was employed to simplify the implementation, and so G^2 mixed mass matrices were stored for each problem.

The calculation of the mixed mass matrices depends on integrals of the form $\int_{V_e^g} d\mathbf{r}\ R_i^g(\mathbf{r}) R_j^{g'}(\mathbf{r})$ where R_i^g and $R_j^{g'}$ are defined over different meshes. This requires a supermesh, which is defined as in [24] to be the intersection of mesh g with mesh g'.

Given two arbitrarily constructed meshes of the physical domain V, the construction of a supermesh is a non-trivial problem, see for example [24] or [23]. However this complexity is largely circumvented if the mesh in every group is initially identical and subsequent meshes result from bisecting elements in parametric space in both parametric directions. In [21] and [22] this approach was taken for the neutron diffusion equation with GDM on Cartesian grids. As the meshes used were orthogonal, the mixed mass matrix never needed to be explicitly formed. Instead, the integrals could be computed by recursive application of a matrix-vector product where the matrix depends only on the polynomial order of the basis functions and the number of recursive calls is equal to the difference between the number of times the initial element has been subdivided on each mesh. This simplicity arose as the mapping from the parent element to every physical element is given by a scaling in each parametric direction.

This is not the case when the initial mesh is a NURBS description of the geometry, and so explicit calculation of the mixed mass matrices is required over a supermesh which must be calculated. However, the calculation of the supermesh is simplified significantly by the mesh construction method outlined in the previous paragraph. Symbolically this implies that:

$$V_i^g \cap V_j^{g'} = \begin{cases} \emptyset \\ \partial V_i^g \cap \partial V_j^{g'} \\ V_i^g \\ V_j^{g'} \end{cases} \tag{21}$$

In words, this says that if two elements overlap at all, then either $V_i^g \subseteq V_j^{g'}$ or $V_j^{g'} \subseteq V_i^g$, and this is guaranteed by the mesh generation procedure. While this discussion has strictly been in physical space, this is also true of the elements in parametric space. As each element in the mesh was initially described by two open knot vectors (along with a polynomial order) in parametric space, the problem of finding the supermesh for the entire geometry in physical space reduces to that of finding the supermesh in parametric space over each initial element and stitching them together. Due to the tensor product structure of the parametric space, finding the intersection of two meshes is much simpler here than in physical space. This procedure is illustrated graphically in Fig. 1.

4. Multigroup error estimators

In [30] DWR based error estimators were developed for one group, fixed source problems with discrete ordinates angular discretisation and linear triangular DGFEM spatial discretisation. The goal functional of interest was either a volumetric detector response or the current passing through a boundary surface. In [31] these error estimators were extended to eigenvalue problems where the goal functional of interest was the eigenvalue $\lambda = \frac{1}{k_{eff}}$. Both of these error estimators were based on the theory developed by Rannacher and coworkers [25–29]. Here, these DWR error estimators will be extended to multigroup problems using the same underlying methods as those presented in [30] and [31]. As the fixed source problem is linear and the eigenvalue problem is nonlinear, they cannot easily be derived in the same manner, and so are presented separately below. In both cases the error estimators are defined in terms of spatial error only, and so the word "continuous" is used to mean the neutron transport equation with the multigroup and discrete ordinates discretisations already applied.

4.1. Fixed source problems

As the error estimators will ultimately be used to drive locally adaptive refinement, it is convenient to rewrite Equation (19) as a sum over groups, directions and elements:

$$\begin{aligned}\sum_{g=1}^{G}\Bigg\{\sum_{e=1}^{K^g}\Bigg[\sum_{k=1}^{N} w_k\Bigg(-\int_{V_e^g} d\mathbf{r}\,\psi_{k,h}^g\boldsymbol{\Omega}_k\cdot\nabla v^g+\int_{V_e^g} d\mathbf{r}\,\Sigma_t^g\psi_{k,h}^g v^g-\frac{1}{4\pi}\sum_{g'=1}^{G}\int_{V_e^g} d\mathbf{r}\,\Sigma_s^{g'\to g}\phi_h^{g'}v^g+\ldots\\ \ldots+\int_{\partial V_e^{g,+}} dS\,\psi_{k,h}^{g,int}v^g\boldsymbol{\Omega}_k\cdot\mathbf{n}+\int_{\partial V_e^{g,-}\setminus\partial V} dS\,\psi_{k,h}^{g,ext}v^g\boldsymbol{\Omega}_k\cdot\mathbf{n}+\int_{\partial V_e^{g,-}\cap\partial V_R} dS\,\psi_{k',h}^{g}v^g\boldsymbol{\Omega}_k\cdot\mathbf{n}\Bigg)\Bigg]\Bigg\}\\ =\frac{1}{4\pi}\sum_{g=1}^{G}\Bigg\{\sum_{e=1}^{K^g}\Bigg[\sum_{k=1}^{N} w_k\Bigg(\int_{V_e^g} d\mathbf{r}\,S^g v^g\Bigg)\Bigg]\Bigg\}\end{aligned} \tag{22}$$

$\forall v^g \in S_{h,p}$. This can be expressed in bilinear form as $B(\psi_h, v^g) = l(v^g)$ where B and l are defined by the left and right hand sides of (22) respectively and ψ_h represents the numerical (i.e. spatially discretised) angular flux in all directions and groups. For simplicity, the only goal functional considered here is a volumetric detector response:

$$J(\psi) = \sum_{g=1}^{G}\int_V d\mathbf{r}\,\Sigma_{\text{det}}^g(\mathbf{r})\phi^g(\mathbf{r}) \tag{23}$$

although the current through a boundary surface could be done in the same manner as in [30]. The corresponding continuous adjoint problem is given by:

$$-\boldsymbol{\Omega}_k^g\cdot\nabla\psi_k^{g*}(\mathbf{r})+\Sigma_t^g(\mathbf{r})\psi_k^{g*}(\mathbf{r})=\frac{1}{4\pi}\sum_{g'=1}^{G}\Sigma_s^{g\to g'}(\mathbf{r})\phi^{g'*}(\mathbf{r})+\frac{\Sigma_{\text{det}}^g(\mathbf{r})}{4\pi}\quad\text{for } k=1,\ldots,N \text{ and } g=1,\ldots,G \tag{24a}$$

$$\phi^{g*}(\mathbf{r})=\sum_{k=1}^{N} w_k\psi_k^{g*}(\mathbf{r}) \tag{24b}$$

i.e. as Equation (10a) with the sign changed on the streaming term, the matrix of group to group scatter cross-sections transposed, and the source term given by the detector response. The boundary conditions are given by:

$$\psi_k^{g*}(\mathbf{r})=0\qquad\text{for}\quad\mathbf{r}\in\partial V_D\quad\text{and}\quad\boldsymbol{\Omega}_k\cdot\mathbf{n}>0 \tag{25a}$$

$$\psi_k^{g*}(\mathbf{r})=\psi_{k'}^{g*}(\mathbf{r})\qquad\text{for}\quad\mathbf{r}\in\partial V_R\quad\text{and}\quad\boldsymbol{\Omega}_k\cdot\mathbf{n}>0 \tag{25b}$$

which are identical to those for the forward equation given in Equation (11a) but now outgoing fluxes are prescribed. The spatial discretisation process outlined in Section 3.1 is then followed. However, as information is now propagating in the opposite direction, the elements are "downwinded" in Equation (18). This leads to the adjoint bilinear form $B^*(\psi_h^*, z^g) = l^*(z^g)\quad \forall z^g \in S_{h,p}$. This bilinear form is given by:

$$\sum_{g=1}^{G}\Bigg\{\sum_{e=1}^{K^g}\Bigg[\sum_{k=1}^{N}w_k\Bigg(\int_{V_e^g}d\mathbf{r}\,\psi_{k,h}^{g*}\mathbf{\Omega}_k\cdot\nabla z^g+\int_{V_e^g}d\mathbf{r}\,\Sigma_t^g\psi_{k,h}^{g*}z^g-\frac{1}{4\pi}\sum_{g'=1}^{G}\int_{V_e^g}d\mathbf{r}\,\Sigma_s^{g\to g'}\phi_h^{g'*}z^g-\ldots$$
$$\ldots-\int_{\partial V_e^{g,-}}dS\,\psi_{k,h}^{g,int*}z^g\mathbf{\Omega}_k\cdot\mathbf{n}-\int_{\partial V_e^{g,+}\setminus\partial V}dS\,\psi_{k,h}^{g,ext*}z^g\mathbf{\Omega}_k\cdot\mathbf{n}-\int_{\partial V_e^{g,+}\cap\partial V_R}dS\,\psi_{k',h}^{g*}z^g\mathbf{\Omega}_k\cdot\mathbf{n}\Bigg)\Bigg]\Bigg\} \tag{26}$$
$$=\frac{1}{4\pi}\sum_{g=1}^{G}\Bigg\{\sum_{e=1}^{K^g}\Bigg[\sum_{k=1}^{N}w_k\Bigg(\int_{V_e^g}d\mathbf{r}\,\Sigma_{\text{det}}^g z^g\Bigg)\Bigg]\Bigg\}$$

As the angular quadrature weights sum to 4π, $l^*(z^g)=J(z^g)$. It can also be shown that $B^*(u,v)=B(v,u)$, a condition known as adjoint consistency [36–38]. As the proof requires different notation and is not necessary to understand the rest of the error estimator derivation, it is included in Appendix A.

The remainder of the error estimator derivation closely follows that in [30]. Two important properties drop out by construction of the bilinear forms. Denoting by ψ and ψ^* the exact solutions to forward and adjoint multigroup, discrete ordinates equations, these are given by:

$$B(\psi,v^g)=l(v^g)\qquad\forall v^g\in S_{h,p}\quad\text{Consistency}\tag{27a}$$
$$B(\psi-\psi_h,v^g)=0\qquad\forall v^g\in S_{h,p}\quad\text{Galerkin orthogonality}\tag{27b}$$

with analogous expressions holding for B^*. The error in the detector response ΔJ is then given by:

$$\Delta J=J(\psi)-J(\psi_h)\tag{28a}$$
$$=J(\psi-\psi_h)\qquad\text{by the linearity of the detector response}\tag{28b}$$
$$=B^*(\psi^*,\psi-\psi_h)\qquad\text{by the definition of the adjoint problem}\tag{28c}$$
$$=B(\psi-\psi_h,\psi^*)\qquad\text{by the adjoint consistency property}\tag{28d}$$
$$=B(\psi-\psi_h,\psi^*-\psi_h^*)\qquad\text{by the Galerkin orthogonality of the forward problem}\tag{28e}$$
$$=l(\psi^*-\psi_h^*)-B(\psi_h,\psi^*-\psi_h^*)\qquad\text{by the consistency of the forward problem}\tag{28f}$$

This expression can be represented by a sum over the elements in all of the meshes of elemental contributions $\Delta J=\sum_{g=1}^{G}\sum_{e=1}^{K^g}\eta_e^g$ where the elemental contributions are calculated by expanding Equation (28f) using the definitions of l and B given in Equation (22):

$$\eta_e^g=\sum_{k=1}^{N}w_k\Bigg[\int_{V_e^g}d\mathbf{r}\left(\frac{S^g}{4\pi}(\psi_k^{g*}-\psi_{k,h}^{g*})+\psi_{k,h}^g\mathbf{\Omega}_k\cdot\nabla(\psi_k^{g*}-\psi_{k,h}^{g*})-\Sigma_t^g\psi_{k,h}^g(\psi_k^{g*}-\psi_{k,h}^{g*})+\ldots\right.$$
$$\left.\ldots+\frac{1}{4\pi}\sum_{g'=1}^{G}\Sigma_s^{g'\to g}\phi_h^g(\psi_k^{g*}-\psi_{k,h}^{g*})\right)-\int_{\partial V_e^{g,+}}dS\,\psi_{k,h}^{g,int}(\psi_k^{g*}-\psi_{k,h}^{g*})\mathbf{\Omega}_k\cdot\mathbf{n}-\ldots\tag{29}$$
$$\ldots-\int_{\partial V_e^{g,-}\setminus\partial V}dS\,\psi_{k,h}^{g,ext}(\psi_k^{g*}-\psi_{k,h}^{g*})\mathbf{\Omega}_k\cdot\mathbf{n}-\int_{\partial V_e^{g,-}\cap\partial V_R}dS\,\psi_{k',h}^{g}(\psi_k^{g*}-\psi_{k,h}^{g*})\mathbf{\Omega}_k\cdot\mathbf{n}\Bigg]$$

After applying the divergence theorem to the second term in the volumetric integral this can be written as:

$$\eta_e^g=\sum_{k=1}^{N}w_k\Bigg[\int_{V_e^g}d\mathbf{r}\,(\psi_k^{g*}-\psi_{k,h}^{g*})\left(\frac{S^g}{4\pi}+\frac{1}{4\pi}\sum_{g'=1}^{G}\Sigma_s^{g'\to g}\phi_h^{g'}-\Sigma_t^g\psi_{k,h}^g-\mathbf{\Omega}_k\cdot\nabla\psi_{k,j}^g\right)+\ldots$$
$$\ldots+\int_{\partial V_e^{g,-}}dS\,(\psi_k^{g*}-\psi_{k,h}^{g*})(\psi_{k,h}^{g,int}-f_k^g(\mathbf{r}))\mathbf{\Omega}_k\cdot\mathbf{n}\Bigg]\tag{30}$$

where $f_k^g(\mathbf{r})$ is defined to be:

$$f_k^g(\mathbf{r}) = \begin{cases} \psi_{k,h}^{g,ext} & \text{for} \quad \mathbf{r} \in \partial V_e^{g,-} \backslash \partial V \\ \psi_{k',h}^{g} & \text{for} \quad \mathbf{r} \in \partial V_e^{g,-} \cap \partial V_R \\ 0 & \text{for} \quad \mathbf{r} \in \partial V_e^{g,-} \cap \partial V_D \end{cases} \tag{31}$$

4.2. Eigenvalue problems

As for fixed source problems, we begin by writing Equation (19) in bilinear form $B(\psi_h, v^g) = \lambda_h F(\psi_h, v^g) \quad \forall v^g \in S_{h,p}$. Conveniently, B is identical in fixed source and eigenvalue problems and so is defined in Equation (22), while $F(\psi_h, v^g)$ is given by:

$$F(\psi_h, v^g) = \frac{1}{4\pi} \sum_{g=1}^{G} \left\{ \sum_{e=1}^{K^g} \left[\sum_{k=1}^{N} w_k \left(\sum_{g'=1}^{G} \chi^g \int_{V_e^g} d\mathbf{r}\, \nu^{g'} \Sigma_f^{g'} \phi_h^{g'} v^g \right) \right] \right\} \tag{32}$$

To make the problem unique the same normalisation of the flux is used as in [31], namely $F(\psi_h, \psi_h) = 1$. The functional of interest considered here is the eigenvalue:

$$J(\psi) = \lambda = \lambda F(\psi, \psi) \quad \text{by the normalisation condition} \tag{33}$$

The corresponding continuous adjoint problem is given by:

$$-\boldsymbol{\Omega}_k^g \cdot \nabla \psi_k^{g*}(\mathbf{r}) + \Sigma_t^g(\mathbf{r}) \psi_k^{g*}(\mathbf{r}) = \frac{1}{4\pi} \sum_{g'=1}^{G} \Sigma_s^{g \to g'}(\mathbf{r}) \phi^{g'*}(\mathbf{r}) + \frac{\lambda^* \nu^g(\mathbf{r}) \Sigma_f^g(\mathbf{r})}{4\pi} \sum_{g'=1}^{G} \chi^{g'} \phi^{g'*}(\mathbf{r}) \quad \text{for } k = 1, ..., N \text{ and } g = 1, ..., G \tag{34a}$$

$$\phi^{g*}(\mathbf{r}) = \sum_{k=1}^{N} w_k \psi_k^{g*}(\mathbf{r}) \tag{34b}$$

with boundary conditions as outlined in Equation (25). This has bilinear form $B^*(\psi_h^*, z^g) = \lambda_h^* F^*(\psi_h^*, z^g)\ \forall z^g \in S_{h,p}$. B^* is also identical in fixed source and eigenvalue problems and so is given in Equation (26), while $F^*(\psi_h^*, z^g)$ is given by:

$$F^*(\psi_h^*, z^g) = \frac{1}{4\pi} \sum_{g=1}^{G} \left\{ \sum_{e=1}^{K^g} \left[\sum_{k=1}^{N} w_k \left(\sum_{g'=1}^{G} \chi^{g'} \int_{V_e^g} d\mathbf{r}\, \nu^g \Sigma_f^g \phi_h^{g'*} z^g \right) \right] \right\} \tag{35}$$

with normalisation $F(\psi, \psi^*) = 1$ as in [31].

In order to consider a volumetric detector response of the form given in Equation (23) with a fission source present, the corresponding adjoint problem would have bilinear form $B^*(\psi_h^*, z^g) = \lambda_h^* F^*(\psi_h^*, z^g) + l^*(z^g) \quad \forall z^g \in S_{h,p}$, where $l^*(z^g) = J(z^g)$ as in fixed source problems. This adjoint solution can be decomposed as $\psi_h^* = \psi_{h,0}^* + \varphi_h^*$, where the fundamental mode $\psi_{h,0}^*$ satisfies the "source free" bilinear form $B^*(\psi_{h,0}^*, z^g) = \lambda_h^* F^*(\psi_{h,0}^*, z^g) \quad \forall z^g \in S_{h,p}$ and the remainder therefore satisfies $B^*(\varphi_h^*, z^g) = \lambda_h^* F^*(\varphi_h^*, z^g) + l^*(z^g) \iff (B^* - \lambda_h^* F^*)(\varphi_h^*, z^g) = l^*(z^g) \quad \forall z^g \in S_{h,p}$. This is now a singular fixed source problem for φ_h^*, with a one dimensional kernel spanned by $\psi_{h,0}^*$. These problems will not be considered here for simplicity, more information on their solution can be found in reference [39].

While the operator F is not self-adjoint, it is linear in both of its arguments, and therefore it's Gâteaux derivative is also linear. The analysis presented in [31] for the one group case therefore also applies to the multigroup case and so will not be repeated here. As in the fixed source case, the error in the functional can be expressed as a sum over the elements of elemental contributions $\Delta J = \sum_{g=1}^{G} \sum_{e=1}^{K^g} \eta_e^g$ where these are given by:

$$\eta_e^g = -\sum_{k=1}^{N} w_k \left[\int_{V_e^g} d\mathbf{r}\, (\psi_k^{g*} - \psi_{k,h}^{g*}) \left(\frac{\lambda_h \chi^g}{4\pi} \sum_{g'=1}^{G} \nu^{g'} \Sigma_f^{g'} \phi_h^{g'} + \frac{1}{4\pi} \sum_{g'=1}^{G} \Sigma_s^{g' \to g} \phi_h^{g'} - \Sigma_t^g \psi_{k,h}^g - \boldsymbol{\Omega}_k \cdot \nabla \psi_{k,j}^g \right) + ... \right.$$
$$\left. ... + \int_{\partial V_e^{g,-}} dS\, (\psi_k^{g*} - \psi_{k,h}^{g*})(\psi_{k,h}^{g,int} - f_k^g(\mathbf{r}))\boldsymbol{\Omega}_k \cdot \mathbf{n} \right] \tag{36}$$

with $f_k^g(\mathbf{r})$ defined as in Equation (31).

Note that in both the fixed source and the eigenvalue case, if the problem has only one energy group the error estimators presented above are identical to those presented in [30] and [31].

4.3. Solution procedure

The error estimators for both cases involve the exact adjoint solution to the multigroup, discrete ordinates equations. This is not available in practical problems of interest, and so must be approximated in some manner. Many techniques for doing so exist, see for example [25,31,40]. Here the "exact" adjoint solution was computed on the same mesh as the forward solution using NURBS one degree greater than those used to solve the forward equation for simplicity. It is stressed that similar levels of accuracy were achieved using computationally cheaper methods in [31], in which only the low order adjoint solution must be calculated directly and the "exact" adjoint solution can be approximated cheaply based on this solution. These methods are expected to be equally successful with the NURBS based methods here, and so in practice at each AMR iteration only the forward and low order adjoint solutions would need to be computed. However, as the focus of this paper is to demonstrate the validity of the multigroup error estimators when used in conjunction with NURBS and GDM, these options will not be explored here.

In order to test the multigroup error estimators, the two metrics used in [25,30,40] are also used here:

$$\theta_1 = \frac{\sum_{g=1}^{G}\sum_{e=1}^{K^g}\eta_e^g}{J(\psi) - J(\psi_h)} \tag{37a}$$

$$\theta_2 = \frac{\sum_{g=1}^{G}\sum_{e=1}^{K^g}|\eta_e^g|}{|J(\psi) - J(\psi_h)|} \tag{37b}$$

where in all cases the "exact" detector response $J(\psi)$ is calculated on a reference mesh with many more elements than those it is being compared to.

Ideally $\theta_1 \approx 1$ as the mesh is refined, indicating that the error estimators are an accurate representation of the spatial error in the solution. In practice we would like to use a quantity like the numerator in θ_2 as an exit criteria from the AMR procedure with confidence that the true error left in the functional due to the spatial discretisation is less than $\sum_{g=1}^{G}\sum_{e=1}^{K^g}|\eta_e^g|$. This will be the case provided $\theta_2 \geq 1$. However if θ_2 is too large this implies effort will be wasted further refining the mesh when the true error $|J(\psi) - J(\psi_h)|$ is already at an acceptable level.

In order to test both the error estimators and the impact of using GDM, a total of 5 different mesh refinement strategies will be followed, 3 involving a single mesh for all energy groups and 2 with GDM. The most basic is uniform refinement, in which there is a single mesh for all energy groups and no adaptivity is employed. In this case each patch in the geometry is refined manually with the aim being to keep all of the elements in the mesh at a roughly similar size. In order to compare goal-based adaptivity to heuristic error indicators as used in [20], the scalar jump error indicator is used. This is expected to refine regions where there is a high gradient in the scalar flux, and is defined by:

$$\beta_e^g = \frac{1}{\|\phi_h^g\|_2}\int_{\partial V_e^g} dS\,|\phi_h^{g,int} - \phi_h^{g,ext}| \quad \text{for } e = 1, ..., K^g \text{ and } g = 1, ..., G \tag{38}$$

where $\|\ \|_2$ is the L^2 norm of a function over the physical domain V. Weighting by the inverse of the L^2 norm of the scalar flux is designed to avoid groups with a large magnitude of scalar flux being refined if the jumps across element boundaries are relatively small. For the goal based scheme,

$$\beta_e^g = |\eta_e^g| \quad \text{for } e = 1, ..., K^g \text{ and } g = 1, ..., G \tag{39}$$

with η_e^g defined in Sections 4.1 and 4.2 for fixed source and eigenvalue problems respectively. With both the goal-based and the heuristic adaptive schemes, the case of using a single mesh for all energy groups and utilising GDM will be compared. In the single mesh case, the error indicators in each group in an element V_e are summed to form a single error indicator for the element in order to drive the adaptivity.

The AMR procedure is then given by:

1. Represent the geometry as a collection of conforming NURBS patches (which will be the elements in the initial mesh). Conforming is used here to mean there are no hanging-nodes.
2. Compute the forward flux ψ_h on this mesh using NURBS of order p.
3. If goal-based adaptivity is being used, compute the adjoint solution over this mesh of both order p (ψ_h^*) and order $p+1$ (ψ^*).
4. Compute the error indicators β_e^g for $e = 1, ..., K^g$ and $g = 1, ..., G$ using either goal-based or heuristic methods.
5. If a single mesh is being used for all energy groups, sum the β_e^g over the groups to form a single error indicator for the element.
6. Refine the 20% of the elements with the largest error indicators. At the same time project ψ_h, ψ_h^* and ψ^* onto this new mesh to use as an initial guess at the next iteration.
7. Return to step 2.

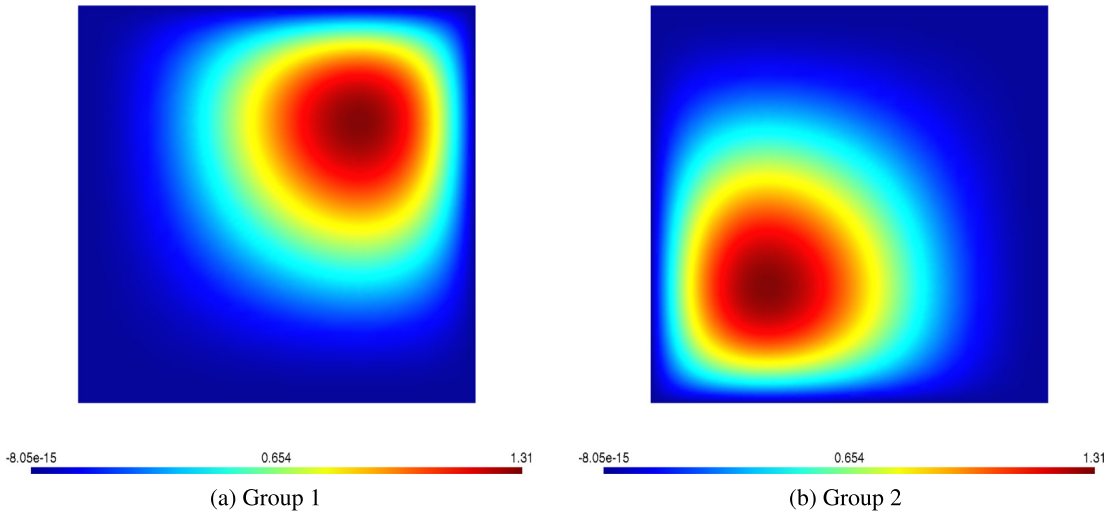


Fig. 2. The scalar flux for the MMS problem in group 1 (a) and group 2 (b). Each is equal to the reflection of the other in the line $y = 1 - x$.

Note that when GDM are being used, the 20% in step 6 refers to the elements with the largest error indicators over all meshes, not 20% of the elements in each mesh. Ideally this iteration would be exited once $\sum_{g=1}^{G}\sum_{e=1}^{K^g}|\eta_e^g|$ was less than a prescribed tolerance for the goal-based schemes. However here we will use a fixed number of iterations in order to test the convergence properties.

5. Numerical results

All results presented are converged to numerical precision for both the forward and adjoint solutions, in order to eliminate iteration error and accurately ascertain the effectiveness of both the schemes and the error estimators. References to the order of the NURBS functions used as a basis for a problem are for the forward problem, and so with goal-based schemes adjoint solves will have also been performed with an order 1 greater than that specified. All errors in the detector response or eigenvalue are absolute errors unless stated otherwise. Convergence results for non goal-based schemes are presented with the total number of elements in the mesh as the independent variable, calculated as $\sum_{g=1}^{G} K^g$.

As explained in Section 4.3, in practice the high order adjoint solution would not be computed directly. As the low order adjoint solution is computed on the same mesh and with the same polynomial order as the forward solution, it should take approximately the same amount of time, and so each AMR iteration will take approximately twice as long using goal-based methods as with the other strategies. To facilitate a fair comparison between the computational effort required by each of the methods, the independent variable for goal-based methods in the convergence plots will therefore be twice the total number of elements in the mesh, calculated as $2 \times \sum_{g=1}^{G} K^g$.

5.1. Method of manufactured solutions (MMS)

A modified version of the manufactured solution presented in [41] is used here to verify that the GDM scheme attains the correct order of convergence. This manufactured solution is for the discrete ordinates equations rather than the full transport equation, and therefore only tests the convergence rate of the spatial discretisation.

The domain is the unit centimetre square, over which the manufactured solution is given by:

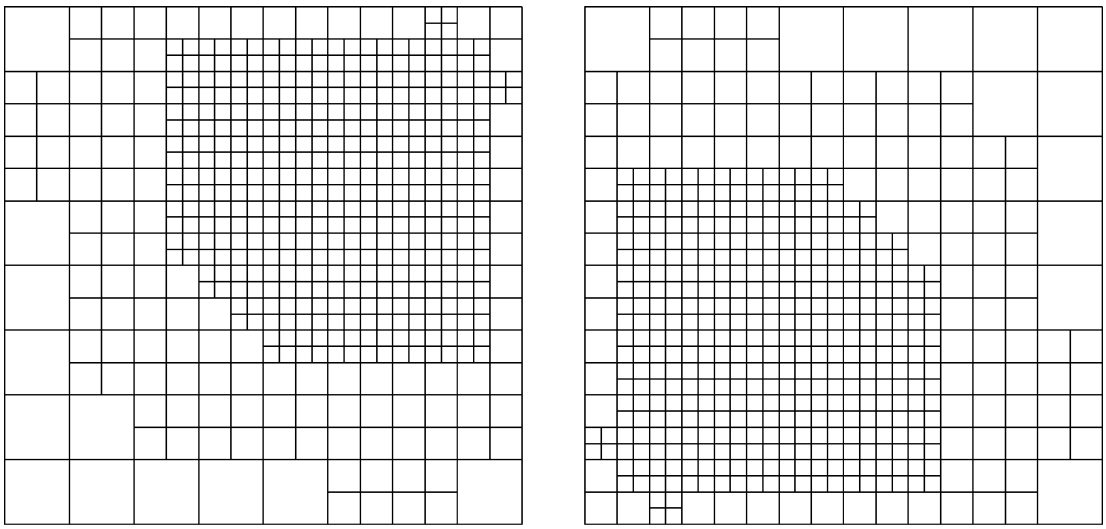
$$\psi_k^1(x, y) = x^2 y^2 (1 - x^2)(1 - y^2)(1 + \mu_k^2 + \eta_k^2) \qquad \text{for } k = 1, \ldots, N \tag{40a}$$

$$\psi_k^2(x, y) = \psi_k^1(1 - x, 1 - y) \qquad \text{for } k = 1, \ldots, N \tag{40b}$$

$$\phi^1(x, y) = \sum_{k=1}^{N} w_k \psi_k^1 \tag{40c}$$

$$\phi^2(x, y) = \sum_{k=1}^{N} w_k \psi_k^2 \tag{40d}$$

where N is the number of directions in the angular quadrature set being used. All MMS calculations presented here were performed with an S_4 level symmetric angular quadrature set. The scalar flux solution in each group is presented in Fig. 2, in which it can be seen that ϕ^2 is a reflection of ϕ^1 in the line $y = 1 - x$. The angular flux in all directions and both groups



(a) Group 1 (b) Group 2

Fig. 3. The mesh at step 10 of the AMR process in group 1 (a) and group 2 (b) for the MMS problem with Cartesian meshes and linear basis functions. Both meshes contain 613 elements, and are symmetrical about the line $y = 1 - x$.

is equal to zero on the domain boundary, and so a vacuum boundary condition can be used. The anisotropic source term in each direction and each group can then be calculated:

$$Q_k^1(x,y) = 4\pi\left(\boldsymbol{\Omega_k}\cdot\nabla\psi_k^1(x,y) + \Sigma_t^1\psi_k^1(x,y)\right) - \Sigma_s^{1\to1}\phi^1(x,y) - \Sigma_s^{2\to1}\phi^2(x,y) \quad \text{for } k = 1,...,N \tag{41a}$$

$$Q_k^2(x,y) = 4\pi\left(\boldsymbol{\Omega_k}\cdot\nabla\psi_k^2(x,y) + \Sigma_t^2\psi_k^2(x,y)\right) - \Sigma_s^{1\to2}\phi^1(x,y) - \Sigma_s^{2\to2}\phi^2(x,y) \quad \text{for } k = 1,...,N \tag{41b}$$

$\Sigma_t^1 = \Sigma_t^2 = 1$, $\Sigma_s^{1\to1} = \Sigma_s^{2\to2} = 0.5$ and $\Sigma_s^{1\to2} = \Sigma_s^{2\to1} = 0.25$ in order to keep the problem symmetric between the groups.

Once the numerical scalar fluxes ϕ_{num}^1 and ϕ_{num}^2 have been calculated in each group, the total L_2 error ϵ_2 is then given by:

$$\epsilon_2 = \left(\int_0^1\int_0^1 \left(\phi_{\text{num}}^1 - \phi^1\right)^2 + \left(\phi_{\text{num}}^2 - \phi^2\right)^2 dxdy\right)^{\frac{1}{2}} \tag{42}$$

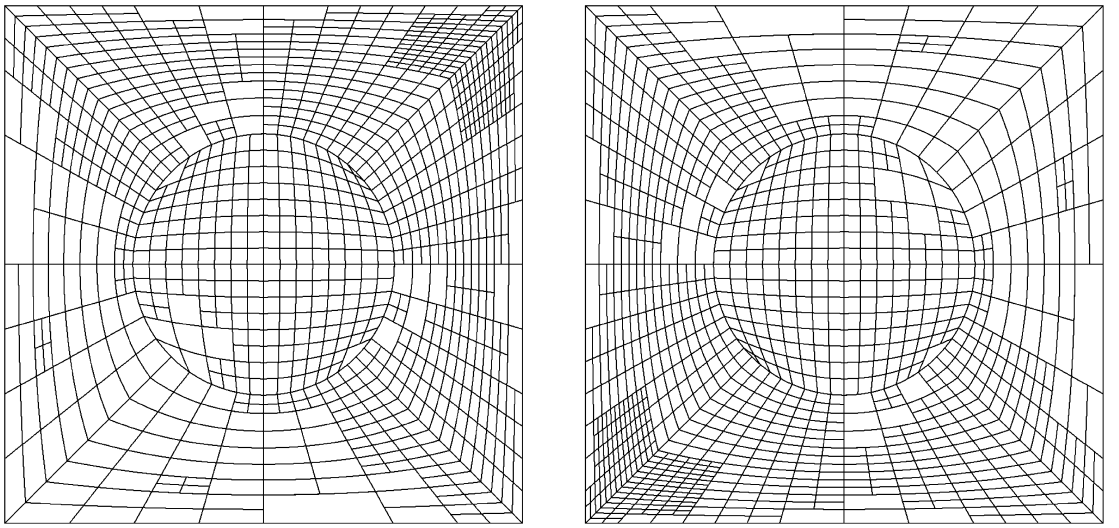
In order to generate different meshes for each group, an adaptive procedure was followed using the heuristic error indicator defined in Section 4 with GDM.

The first mesh types generated were Cartesian grids, in which case the scheme is equivalent to DGFEM and the assembly of terms between meshes is identical to that in [22]. As the AMR procedure refines each mesh, the element boundaries remain aligned with the coordinate axes, and so all volumetric and boundary integrals are performed exactly using Gaussian quadrature. Fig. 3 shows the mesh in each group generated by this procedure at step 10 of the AMR process. As the solutions in each group are symmetric about the line $y = 1 - x$, the meshes generated by the AMR process are as well.

The second mesh types used were pincell-style, as shown in Fig. 4. The radius of the circle was chosen to be $\frac{1}{5\pi}$ in order to keep the area of each patch in the initial geometry identical. In this case the assembly of terms between meshes is now not exact as Gaussian quadrature points can not exactly integrate the rational basis functions involved. Fig. 4 shows the mesh in each group generated by this procedure at step 10 of the AMR process. As in Fig. 3, the meshes are symmetric about the line $y = 1 - x$.

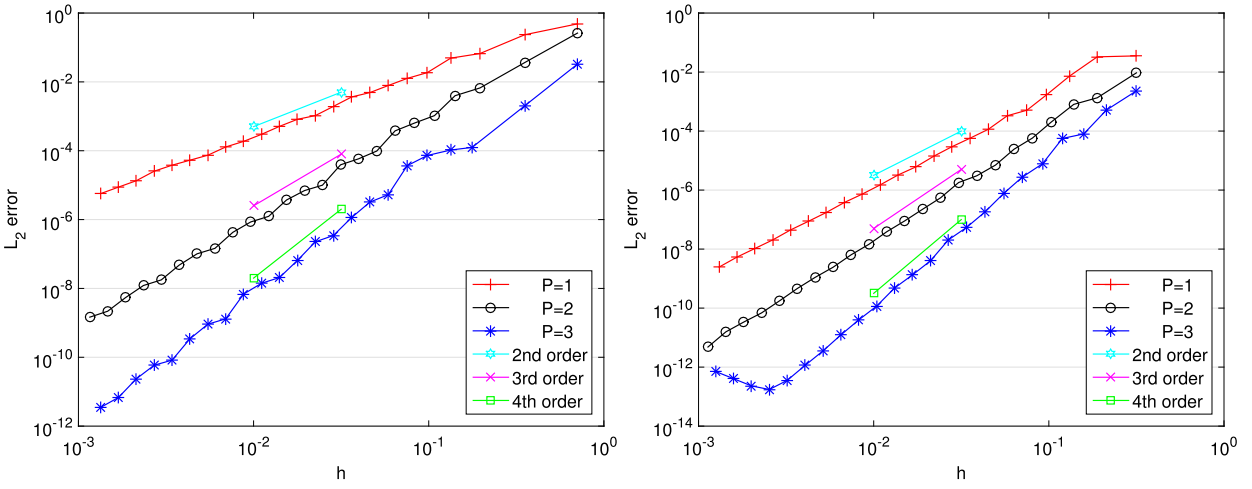
The results for Cartesian meshes are presented in Fig. 5a for first, second and third order basis functions. Fourth order basis functions can exactly represent the manufactured solution on Cartesian meshes, and therefore only give an error due to machine precision, which has been verified to be the case. The GDM formulation with order p basis functions attains order $p + 1$ convergence for all polynomial orders on Cartesian meshes. As in [20], the convergence rates are not perfectly straight lines as $\bar{h}$ is no longer a perfect measure of element size.

The results for pincell meshes are presented in Fig. 5b for second, third and fourth order basis functions. As the basis functions are now rational, the fourth order NURBS can no longer exactly represent the manufactured solution and can be used to verify the convergence rate at this order. The GDM formulation again obtains order $p + 1$ convergence with order p basis functions, limited to $\approx 10^{-12}$ due to a combination of machine precision and the condition number of the system.



(a) Group 1 (b) Group 2

Fig. 4. The mesh at step 10 of the AMR process in group 1 (a) and group 2 (b) for the MMS problem with pincell meshes and quadratic basis functions. Both meshes contain 1,028 elements, and are symmetrical about the line $y = 1 - x$.



(a) Cartesian Meshes (b) Pincell Meshes

Fig. 5. The convergence under mesh refinement with Cartesian (a) and pincell (b) meshes for various polynomial orders, along with ideal rates of convergence. $\bar{h} := \frac{1}{\sqrt{E}}$ where E is the total number of elements in all meshes. In all cases the scheme produces the expected rate of convergence, limited to $\approx 10^{-12}$ due to a combination of machine precision and the condition number of the system.

5.2. One group eigenvalue problem

This problem has been chosen in order to verify that the goal-based adaptivity works with the hanging-node formulation and higher-order elements, as the method presented in [30] and [31] utilised a conforming mesh formulation with linear elements. To this end, problem 2 from [31] was chosen for two reasons:

- The geometry, shown in Fig. 6, is Cartesian, and can therefore be modelled exactly by linear NURBS. This allows the performance of the error estimators to be compared for linear, as well as higher-order basis functions.
- It is a one group problem, and therefore the error estimators employed are identical to those presented in [31]. In this way only the spatial discretisation is tested, rather than the multigroup error estimators.

The problem was solved with an S_8 level symmetric angular quadrature set, and the forward and adjoint scalar fluxes are presented in Fig. 7. In this case they are identical, as the angular quadrature set is symmetric in the sense that for all

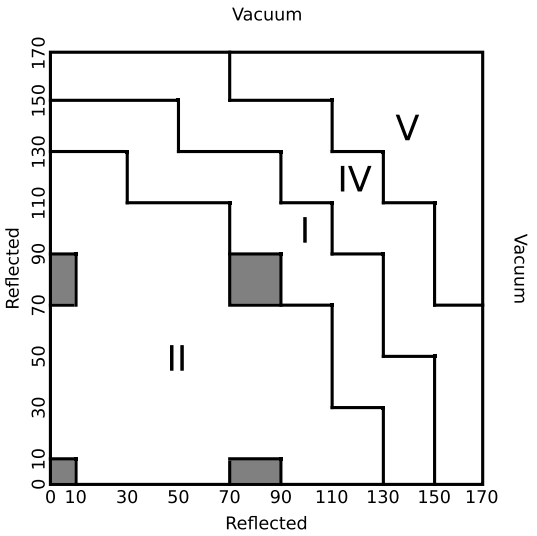


Fig. 6. The geometry for the one group eigenvalue problem, along with region numbers corresponding to the material properties specified in Table 1. The shaded region corresponds to material III in Table 1.

Table 1
The cross-section data for the one group eigenvalue problem.

Material	Σ_t (cm^{-1})	Σ_s (cm^{-1})	$\nu\Sigma_f$ (cm^{-1})
I	0.3	0.244	0.064
II	0.3	0.244	0.060
III	0.31	0.244	0.060
IV	1.2	1.15	0.00
V	0	0	0.00

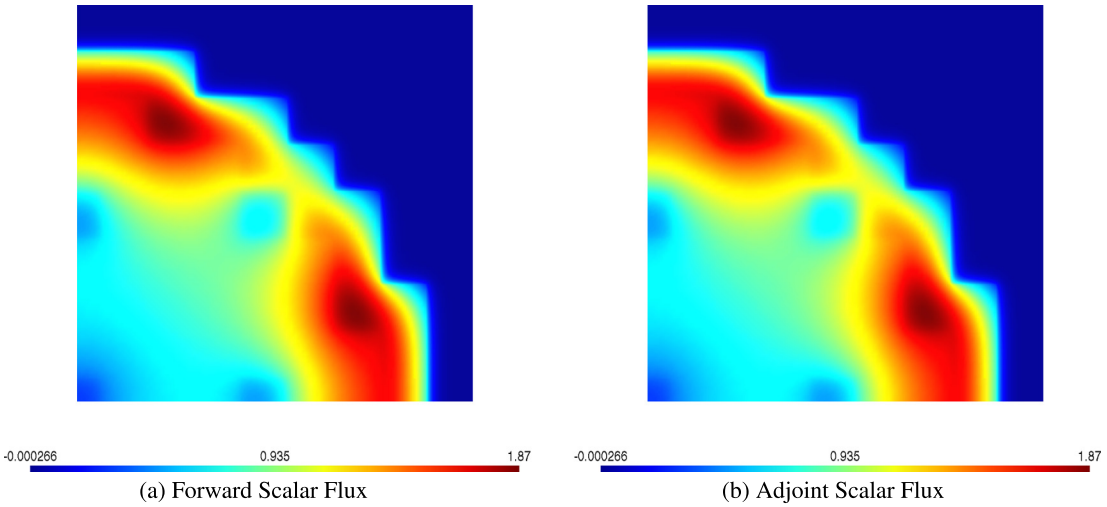


Fig. 7. The scalar flux for the forward (a) and adjoint (b) problems for the one group eigenvalue problem. They are identical for this problem.

$\mathbf{\Omega}_k$ $k = 1, ..., N$, there is a k' such that $\mathbf{\Omega}_{k'} = -\mathbf{\Omega}_k$. The equations being solved for ψ_k and $\psi^*_{k'}$ are then the same, leading to $\psi_k = \psi^*_{k'}$ which implies $\phi = \phi^*$.

This problem was solved with uniform, heuristic and goal-based mesh refinement strategies with both linear and quadratic basis functions (clearly GDM cannot be used in a one group problem). The reference eigenvalue was calculated using quadratic goal-based refinement with 550,666 elements, with $\lambda_{\text{ref}} = 0.9336949169$.

The error in the eigenvalue compared to λ_{ref} for linear and quadratic basis functions with all three mesh refinement strategies is presented in Fig. 8. The eigenvalue calculated using the AMR schemes with linear basis functions is approximately an order of magnitude more accurate than uniform refinement for the same amount of computational effort, and around 2 orders of magnitude more accurate with quadratic basis functions. For this problem, goal-based refinement is

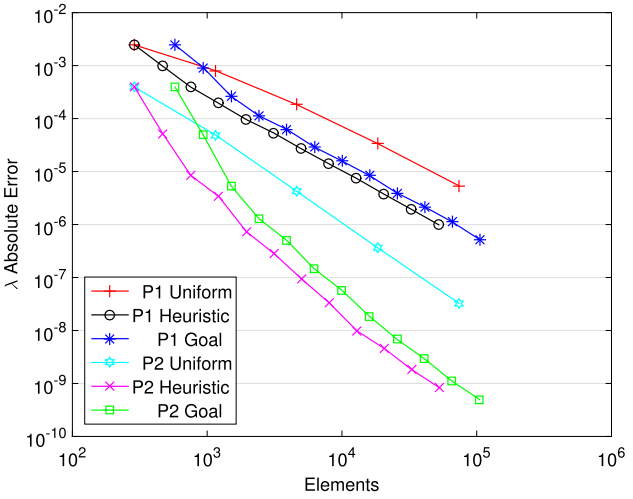
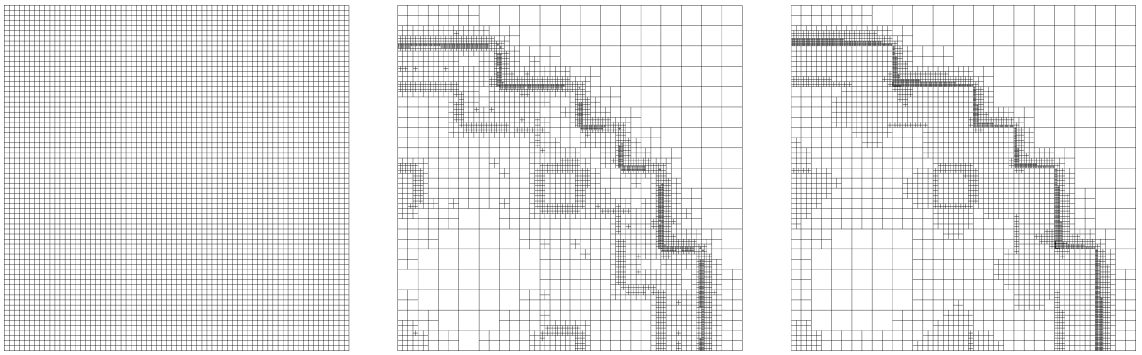


Fig. 8. The convergence of λ for the one group eigenvalue problem as the mesh is refined using uniform, heuristic and goal-based refinement strategies with linear and quadratic basis functions. For the same amount of computational effort, the adaptive schemes are significantly more accurate than uniform refinement for this problem with both polynomial orders, with the heuristic schemes being slightly more accurate than goal-based AMR.



(a) Uniform Refinement (b) Heuristic Refinement (c) Goal-Based Refinement

Fig. 9. Meshes with approximately 5,000 linear elements for the one group eigenvalue problem with uniform (a), heuristic (b) and goal-based (c) refinement strategies. In this case the heuristic error indicator refining the mesh around sharp gradients in the solution produces a qualitatively similar mesh to the goal-based refinement strategy.

slightly less accurate for the same amount of computational effort than heuristic refinement for both polynomial orders. It should be noted that this is due to the cost of the adjoint solve at each AMR iteration, which has been used to shift the goal-based convergence plots by a factor of 2 in the x-direction. In terms of accuracy per element in the mesh, goal-based refinement schemes of both polynomial orders are more accurate than heuristic AMR for this problem.

Meshes from step 7 of the AMR process for heuristic and goal-based refinement, along with a uniform mesh with a similar number of elements are presented in Fig. 9. For this problem the heuristic refinement produces qualitatively similar meshes to goal-based refinement, which is not surprising given the similarity in the eigenvalue errors shown in Fig. 8. The heuristic refinement procedure could be said to slightly over-refine the gradients in the flux around the flux depressions in material III, while slightly underresolving the boundary between fuel region I and moderator region IV compared to the goal-based refinement, which accounts for the difference in eigenvalue accuracy per element in the mesh.

θ_1 and θ_2 are plotted for the goal-based refinement strategy with both linear and quadratic basis functions in Fig. 10. For both linear and quadratic basis functions, $\theta_2 \approx \theta_1$, suggesting that the sign of the error estimators in most elements is the same. In such a case, $\sum_e |\eta_e|$ could not be used as a reliable exit criteria of an AMR procedure. As in [31], θ_1 converges to 1 as the mesh is refined for linear basis functions, confirming that the error estimators reliably predict the error with the hanging-node spatial discretisation for first order elements. However this is not the case with quadratic basis functions.

The accurate calculation of θ_1 relies on the accurate approximation of ψ_k^* for $k = 1, ..., N$. If ψ_k^* were exact for all values of k, θ_1 would be exactly equal to 1 for any mesh and any polynomial order. However ψ_k^* is not exact, it is computed over the same mesh as the forward solution with basis functions that are one polynomial order higher. Denote by $\psi_{k,p}^*$ the adjoint solution in direction k of polynomial order p computed over a given mesh. Then the fact that θ_1 converges for linear basis functions suggests that:

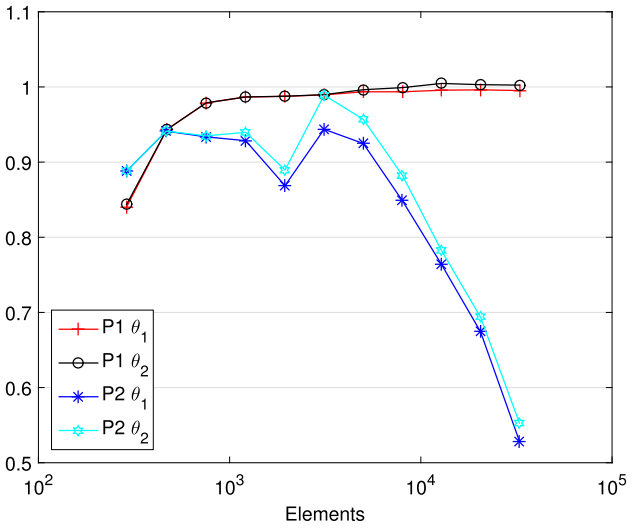


Fig. 10. θ values for the one group eigenvalue problem with linear and quadratic basis functions. θ_1 converges to 1 for linear basis functions but does not for quadratics. For both polynomial orders, θ_2 is very similar to θ_1, suggesting that the sign of η_e is the same in the majority of the elements.

$$|\psi_k^* - \psi_{k,2}^*| \ll |\psi_k^* - \psi_{k,1}^*| \quad \text{and so} \tag{43}$$

$$|\psi_k^* - \psi_{k,1}^*| \approx |\psi_{k,2}^* - \psi_{k,1}^*| \tag{44}$$

for $k = 1, \ldots, N$. Conversely, the fact that θ_1 does not converge to 1 as the mesh is refined for quadratic basis functions suggests that:

$$|\psi_k^* - \psi_{k,3}^*| \lll |\psi_k^* - \psi_{k,2}^*| \quad \text{and so} \tag{45}$$

$$|\psi_k^* - \psi_{k,2}^*| \not\approx |\psi_{k,3}^* - \psi_{k,2}^*| \tag{46}$$

for $k = 1, \ldots, N$. Loosely speaking, this says that quadratics are a lot more accurate than linears, but cubics are not much more accurate than quadratics. This makes sense from a degrees-of-freedom-per-element perspective, as the ratio of quadratic to linear DOFs is $\frac{9}{4} = 2.25$, whereas the ratio of cubic to quadratic DOFs is $\frac{16}{9} = 1.\dot{7}$.

Realistic reactor physics and shielding geometries generally require quadratic NURBS to be modelled exactly. This suggests that $\sum_e \eta_e$ could not be used as a reliable approximation of the error in the specified goal when quadratic NURBS are used. However, for more complex problems it is unlikely that the sign of η_e will be the same in most elements, and so it is possible that $\sum_e |\eta_e|$ could still be used as a stopping criteria of an AMR process, possibly with a safety factor included.

5.3. Multigroup fixed source problem

This problem has been designed to test the performance of the multigroup error estimator derived in Section 4 for fixed source problems. It is an adaptation of problem 2 in [30], extended to a two-group problem to test the multigroup error estimator. The geometry, presented in Fig. 11 has been deliberately kept Cartesian. This is so that linear basis functions can be used to solve the problem without sacrificing geometric accuracy, as the results in Section 5.2 suggest that the error estimator is not as accurate for quadratic basis functions. The detector response for this problem is the total flux integral of both energy groups over the 1 cm $\times$ 1 cm square in the upper right corner of the geometry, as shown in Fig. 11.

The material properties for the problem, shown in Table 2, were chosen to give the solution specific properties. The problem contains only downscatter so that once a neutron is in group 2 it cannot return to group 1. The bulk of the geometry is made up of a moderator. In group 2, this material is highly scattering, while in group 1 it is also highly scattering, with half of these interactions resulting in the neutron being thermalised. The shield is identical to the moderator for fast neutrons, but is highly absorbing for thermal neutrons. The source region contains a source only in the fast group. These properties, in combination with the geometry lead to two important features of the solution:

Table 2
The cross-section data for the multigroup fixed source problem.

Material	Σ_t^1	Σ_t^2	$\Sigma_s^{1\to1}$	$\Sigma_s^{1\to2}$	$\Sigma_s^{2\to1}$	$\Sigma_s^{2\to2}$	Q^1	Q^2
I – Source	1	1	0.49	0.49	0	0.98	1	0
II – Moderator	1	1	0.45	0.45	0	0.90	0	0
III – Shield	1	1	0.45	0.45	0	0.05	0	0

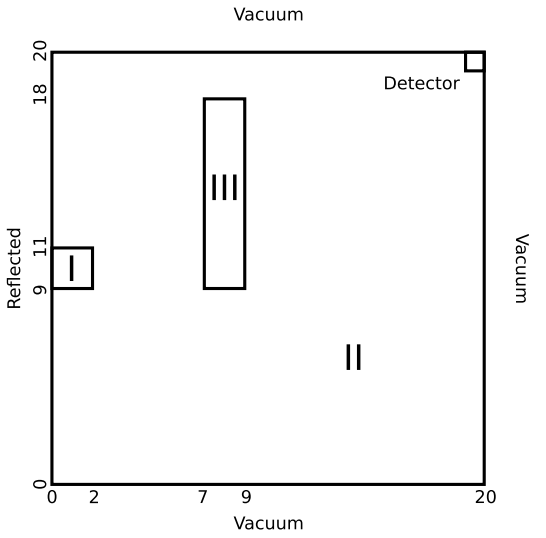
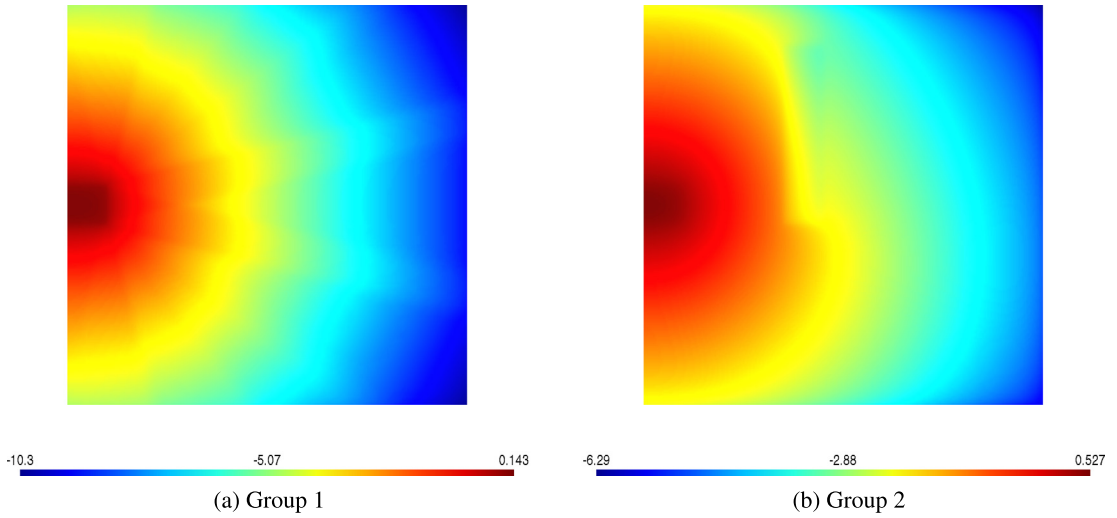


Fig. 11. The geometry for the multigroup fixed source problem, along with region numbers corresponding to the material properties specified in Table 2. The detector response for this problem is the total flux integral in both groups over the 1 cm × 1 cm square in the upper-right corner of the geometry.



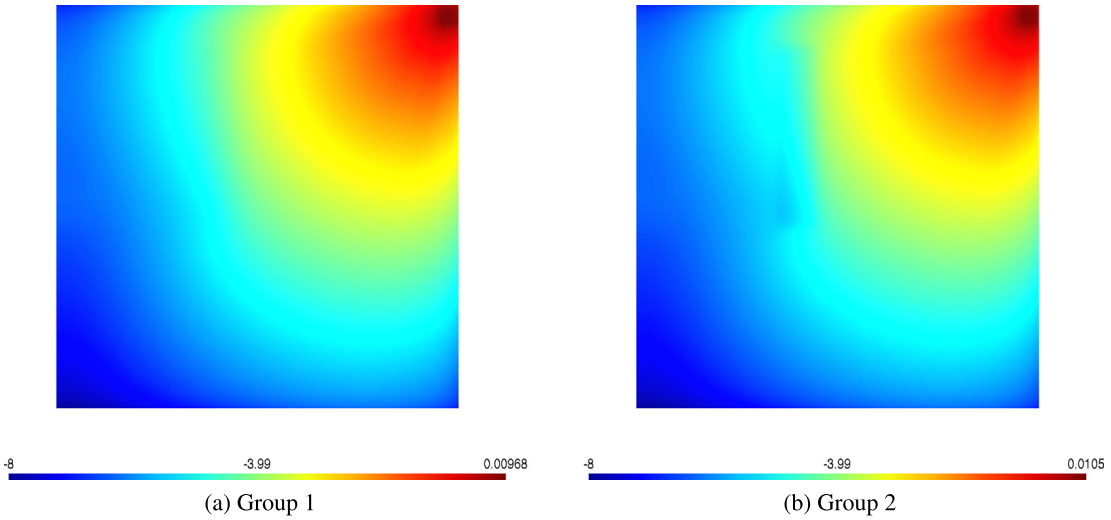
(a) Group 1 (b) Group 2

Fig. 12. The scalar flux for the forward problem in group 1 (a) and group 2 (b) on a logarithmic scale for the multigroup fixed source problem. The group 2 flux is approximately 4 orders of magnitude larger than the group 1 flux in the detector region, and so goal-based refinement with GDM would be expected to preferentially refine group 2 in this region of the geometry.

- The majority of the neutrons that reach the detector have been thermalised, and so the accuracy of the flux in group 2 in the detector region should be much more important than the accuracy of group 1. This is demonstrated in Fig. 12, in which the scalar flux in the detector region is 4 orders of magnitude larger in group 2 than in group 1.
- A large proportion of the neutrons that are thermalised before the shield are then absorbed in it. Therefore a large contribution to the detector response will be from neutrons that are thermalised between the shield and the detector, and in order to accurately capture this, solution fidelity will be required in group 1 in this region of the geometry.

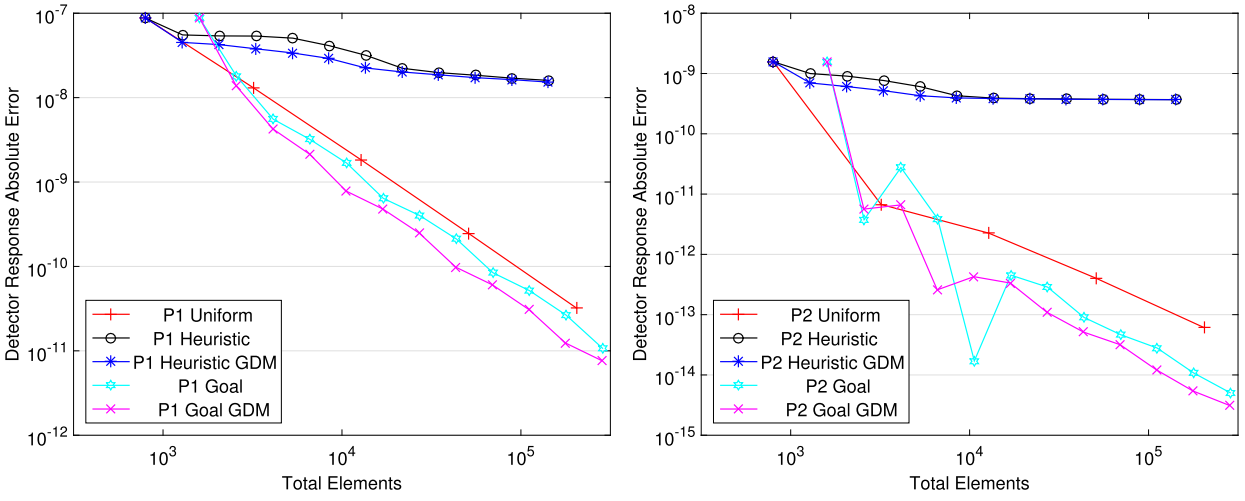
The problem was solved with an S_8 level symmetric angular quadrature set. The forward scalar flux is shown in Fig. 12. The sharpest gradient in the solution in both groups is near to the source region, and so the heuristic error indicator will be expected to refine around here. Fig. 12b shows the thermal neutrons being absorbed in the shield. The adjoint scalar flux is shown in Fig. 13, and is very similar in both groups as the response is the total flux integral over the detector of both energy groups.

The problem was solved with three single-mesh refinement strategies: uniform, heuristic and goal-based, and two GDM refinement strategies: heuristic and goal-based. In all cases both linear and quadratic basis functions were used. The reference detector response was calculated using quadratic goal-based refinement with GDM with 1,487,705 total elements over both meshes, and was found to be 1.723622447E−6.



(a) Group 1 (b) Group 2

Fig. 13. The scalar flux for the adjoint problem in group 1 (a) and group 2 (b) on a logarithmic scale for the multigroup fixed source problem. As the response is the total flux over the detector region of both energy groups the adjoint flux profile is very similar in both groups.

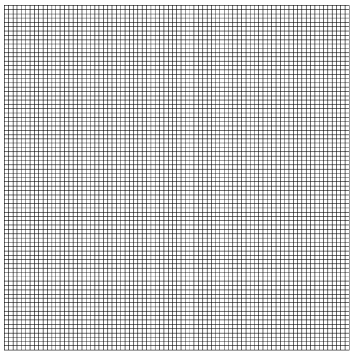


(a) Linear Basis Functions (b) Quadratic Basis Functions

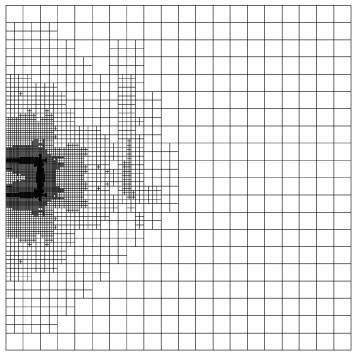
Fig. 14. The convergence of the detector response for the multigroup fixed source problem as the mesh is refined with linear (a) and quadratic (b) basis functions. Uniform, heuristic and goal-based refinement strategies were used with a single mesh for all energy groups, as well as GDM for heuristic and goal-based error indicators. For both polynomial orders both heuristic methods converge very slowly as Fig. 15 demonstrates the mesh only being refined near to the source region. Goal-based adaptivity with a single mesh is more accurate for the same amount of computational effort than uniform refinement, and employing GDM increases this accuracy further with both polynomial orders. The 2nd and 5th points on the P2 Goal line and the 2nd and 4th points on the P2 Goal GDM line are artificially low as the error is changing in sign.

The error in the detector response for all five mesh refinement strategies with both linear and quadratic basis functions is presented in Fig. 14. For both polynomial orders both single mesh and GDM heuristic refinement converge very slowly. Fig. 15 shows these schemes concentrating mesh refinement only near to the source region where the sharp gradients in the solution are observed. With linear basis functions, the single mesh goal-based scheme is $\approx \times 1.5$ more accurate than uniform refinement, while the goal-based scheme with GDM is $\approx \times 2.5$ more accurate than uniform refinement on a per element basis. With quadratic basis functions these figures rise to 6 and 12 respectively. To achieve the same level of accuracy in the detector response, a saving of $\approx \times 2$ elements can be made with linears or a saving of $\approx \times 5$ with quadratic basis functions by using goal-based adaptivity with GDM compared to uniform refinement.

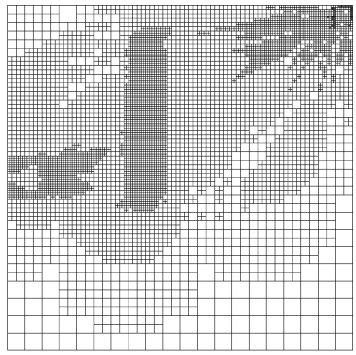
Meshes from step 7 of the AMR process, along with a uniform mesh with a similar number of elements are presented in Fig. 15. Both heuristic schemes refine only around the sharp gradient in the solution near to the source region, and a little at the nearest edge of the shield. As such the detector response calculated by these schemes does not improve much as the AMR scheme proceeds as the mesh is barely refined at all between the shield and the detector. As expected, the goal-based scheme with GDM has refined the group 1 mesh mostly towards the detector and past the shield as well. The



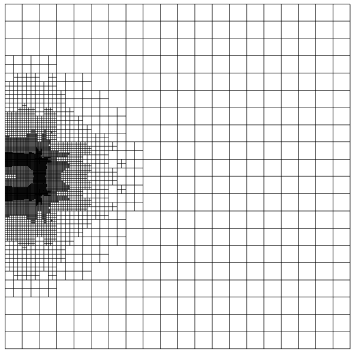
(a) Uniform Refinement



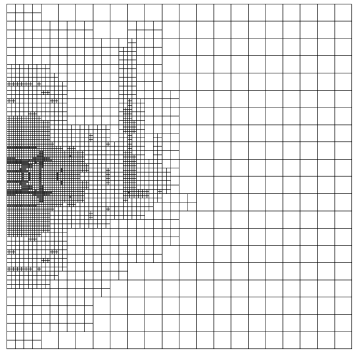
(b) Heuristic Refinement



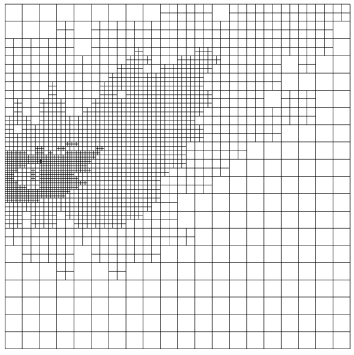
(c) Goal-Based Refinement



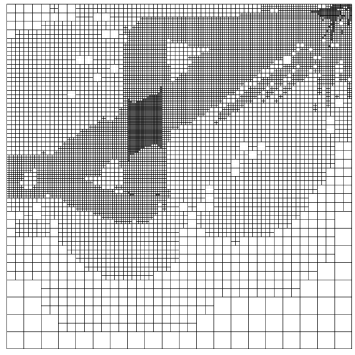
(d) Heuristic GDM Refinement Group 1



(e) Heuristic GDM Refinement Group 2



(f) Goal-Based GDM Refinement Group 1

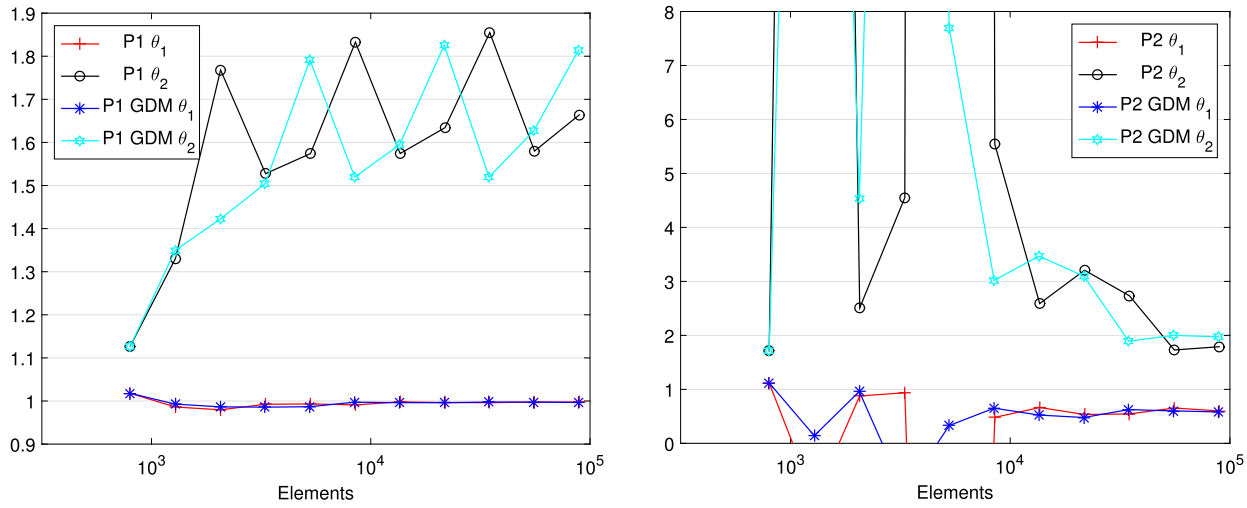


(g) Goal-Based GDM Refinement Group 2

Fig. 15. Meshes with approximately 13,000 linear elements total for the multigroup fixed source problem. Refinement schemes with a single mesh for all groups are shown for uniform (a), heuristic (b) and goal-based (c) refinement strategies, while GDM strategies with heuristic (d & e) and goal-based (f & g) refinement strategies show the mesh in each group. Both heuristic schemes concentrate elements around the sharp gradients in the flux near to the source region and the absorber, while the goal-based meshes capture this detail, while also refining the mesh towards the detector in order to accurately determine the response.

group 2 mesh, which has $\approx \times 4$ as many elements as the group 1 mesh, is refined mostly on the path the neutrons are most likely to take to the detector, particularly in the shield and around the detector itself. This is in keeping with the fact that the thermal flux is 4 orders of magnitude larger than the fast flux in the detector region. Ray effects can be observed in the group 2 mesh near to the detector, suggesting that the resolution of these features of the underlying solution is important to the accurate calculation of the detector response for the S_8 quadrature set employed.

The goal-based scheme with a single mesh can be seen to have retained the important features of the goal-based scheme with GDM. However the restriction to a single mesh means that the shield and the region from the shield to the detector are not quite as refined as the GDM thermal group mesh, leading to the loss in accuracy of the detector response of $\approx \times 2$ shown in Fig. 14.



(a) Linear Basis Functions (b) Quadratic Basis Functions

Fig. 16. θ values for the multigroup fixed source problem with linear (a) and quadratic (b) basis functions. The error estimators work extremely well for linear basis functions, with θ_1 converging to 1 for both the single mesh & GDM schemes and θ_2 remaining in the range 1–2 for all meshes. For quadratic basis functions θ_1 is asymptotically ≈ 0.6 and θ_2 is in the range 1–4 once the mesh is sufficiently refined for both the single mesh and GDM schemes. Prior to this though both θ values have large jumps which correspond to sign changes in the error and can be seen as uncharacteristically low errors in Fig. 14b.

θ_1 and θ_2 are plotted for the goal-based schemes with both a single mesh and GDM for linear and quadratic basis functions in Fig. 16. With linear basis functions, θ_1 converges to 1 in both the single mesh and GDM cases. θ_2 remains in the range 1–2 for all of the meshes, suggesting that $\sum_e |\eta_e|$ would be an accurate stopping criterion to use for this problem with linear basis functions.

As with the problem presented in Section 5.2, the behaviour of the error estimators is not as nice in the quadratic case. There are large spikes in both θ_1 and θ_2 for both the single mesh and GDM schemes. These correspond to the solutions in Fig. 14b where the error is artificially low as it is changing in sign. In this case it is not surprising that the error estimators perform poorly. However θ_1 settles down to ≈ 0.6 and θ_2 is in the range 1–4 once the mesh is sufficiently refined for both the single mesh and GDM schemes. As the error estimators work extremely well for the linear multigroup case but not for the quadratic one group case, this behaviour is down to the inaccuracy of the cubic adjoint solution used to compute the error estimator as discussed in Section 5.2 rather than the multigroup error estimator. Regardless, as θ_2 only ever overestimates the error, $\sum_e |\eta_e|$ could still be used as an exit criteria from an AMR process with quadratic basis functions, albeit with less accuracy than the linear case.

5.4. KAIST MOX assembly

A more stringent and realistic test of the multigroup error estimator for eigenvalue problems is a MOX assembly taken from the KAIST 2A benchmark [42]. The geometry, shown in Fig. 17, contains MOX fuel of various enrichments, guide tubes filled with water to represent the unrodded case, and a Gadolinium burnable absober pin. The fuel and burnable absorber pins are surrounded by an extremely thin (0.0085 cm) gas gap and a cladding annulus. The problem is solved in 7 energy groups, the cross-sections for which can be found in [42], along with an exact specification of the geometry.

The problem was solved with an S_8 level symmetric angular quadrature set, and the symmetry of the geometry was exploited to solve for only a quarter of the assembly. The forward scalar flux in groups 1, 2 and 6 are presented in Fig. 18. The group 1 (highest energy) scalar flux can be seen to be extremely complicated, with sharp gradients in the flux throughout the whole assembly and flux depressions in the guide tubes. The scalar flux in group 2 is qualitatively similar to that in group 1, although the gradients are nowhere near as severe, as evidenced by the ratio of the highest to lowest flux values in each group. In contrast, the scalar flux in group 6 (thermal neutrons) is relatively smooth throughout the geometry, except for in the vicinity of the burnable absorber pin in which the flux sharply drops to almost zero.

The adjoint scalar flux in groups 1, 2 and 6 are presented in Fig. 19. The adjoint flux in group 1 is very similar to the forward flux, although the gradients in the solution are not as sharp. This is not the case in group 2, in which the adjoint flux is almost planar. This is due to the term $\nu^1\Sigma_f^1$ being a factor of 5–9 higher than $\nu^2\Sigma_f^2$ in the various fissile materials. In group 6 the flux is relatively flat throughout the domain except in the vicinity of the burnable absorber pin where the flux sharply drops to almost zero as in the forward flux.

The problem was solved with three single-mesh refinement strategies: uniform, heuristic and goal-based, and two GDM refinement strategies: heuristic and goal-based, with quadratic basis functions used in all cases. The reference eigenvalue was calculated on a quadratic mesh with 1,519,736 total elements, and was found to be $\lambda_{\text{ref}} = 0.856526$.

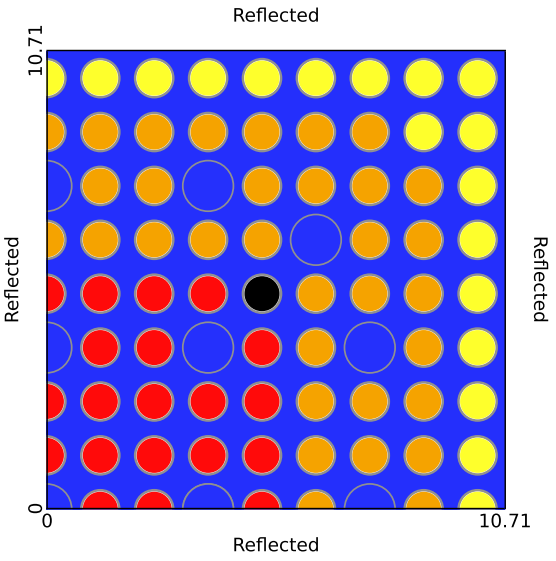
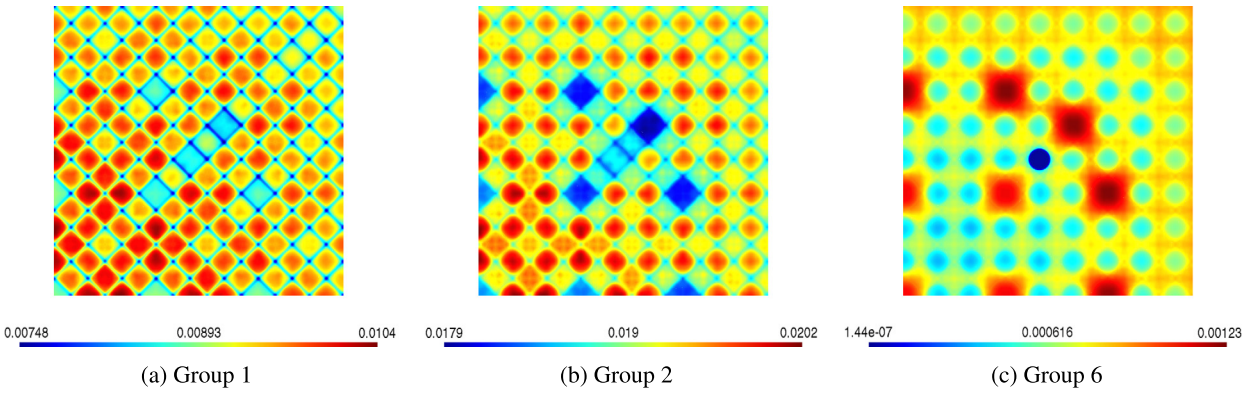
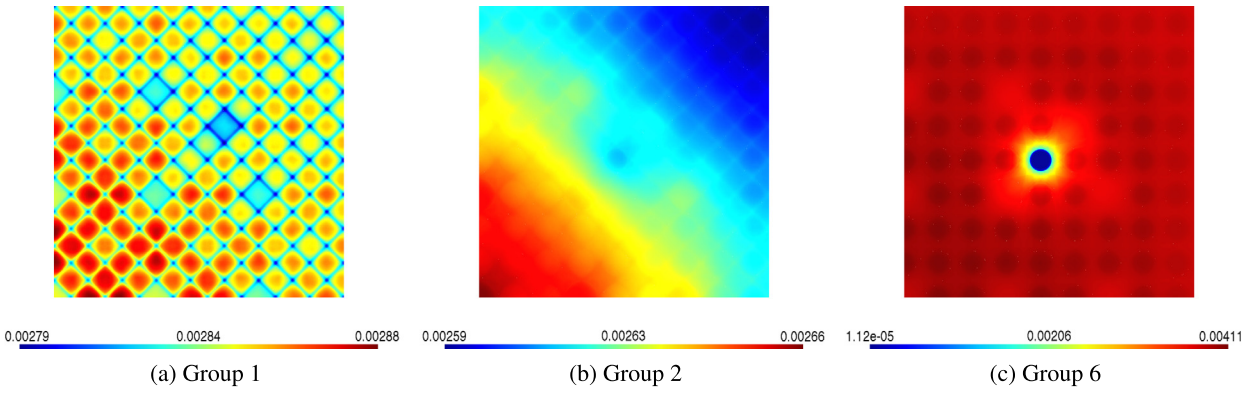


Fig. 17. The geometry for the KAIST MOX assembly problem. Red pins are 8.7% enrichment, orange are 7.0% and yellow are 4.3%. As this is the unrodded case, guide tubes are filled with water, and the black pin represents a Gadolinium burnable absorber. All pins are surrounded by cladding (grey), and the fuel and burnable absorber pins have an additional thin gas gap (white). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)



(a) Group 1 (b) Group 2 (c) Group 6

Fig. 18. The scalar flux for the forward problem in group 1 (a), group 2 (b) and group 6 (c) for the KAIST MOX assembly problem. The flux profile in both fast groups is qualitatively similar although the gradients present are sharper in group 1. In group 6 the flux is relatively smooth throughout the geometry, except for a severe depression in the burnable absorber in which the flux drops to almost zero.



(a) Group 1 (b) Group 2 (c) Group 6

Fig. 19. The scalar flux for the adjoint problem in group 1 (a), group 2 (b) and group 6 (c) for the KAIST MOX assembly problem. The adjoint flux in group 1 is very similar to the forward flux although with much shallower gradients in the solution. In contrast, the adjoint flux in group 2 is almost planar with very little visible detail. The adjoint flux in group 6 is relatively flat almost everywhere apart from in the vicinity of the burnable absorber, in which it drops to almost zero.

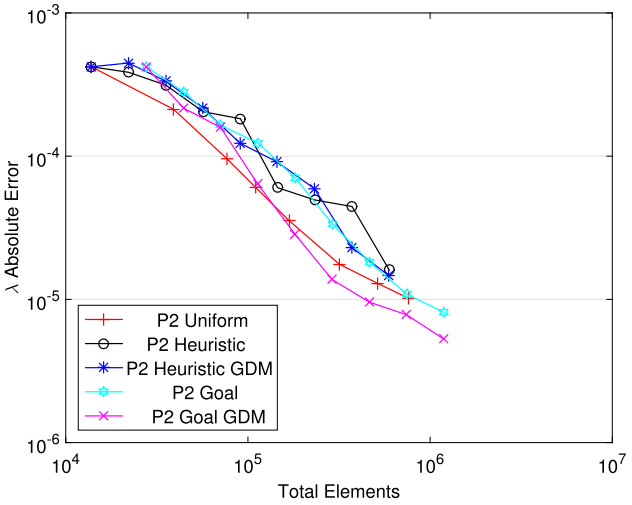


Fig. 20. The convergence of λ for the KAIST MOX assembly problem as the mesh is refined using quadratic basis functions. Uniform, heuristic and goal-based refinement strategies were used with a single mesh for all energy groups, as well as GDM for heuristic and goal-based error indicators. Both heuristic refinement schemes and the goal-based scheme with a single mesh are less accurate for the same amount of computational effort than uniform refinement for this problem. Goal-based refinement with GDM is more accurate than uniform refinement after 5 AMR steps.

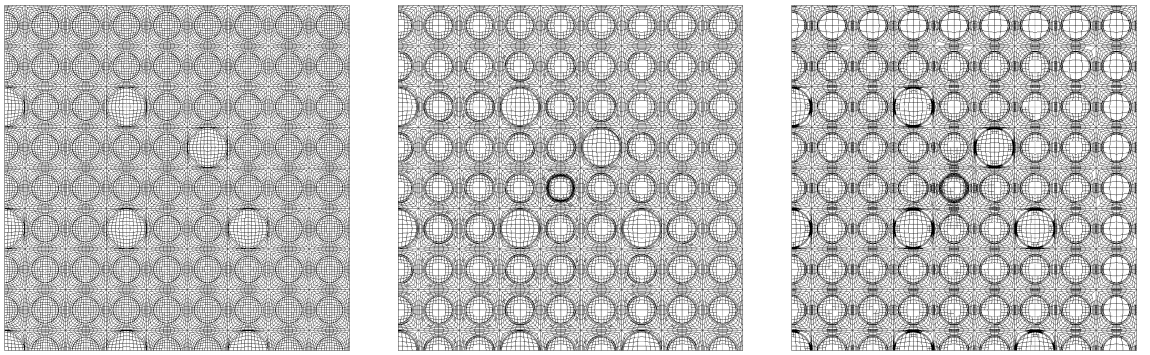
Table 3
The number of elements in each energy group's mesh used to calculate the reference solution, along with the equivalent number of elements that would be in a uniform mesh and the ratio between the two.

Group	Reference mesh elements	Uniform equivalent	Ratio
1	493,343	217,105	2.27
2	69,212	217,105	0.32
3	177,941	217,105	0.82
4	295,073	217,105	1.36
5	268,589	217,105	1.24
6	118,178	217,105	0.54
7	97,400	217,105	0.45

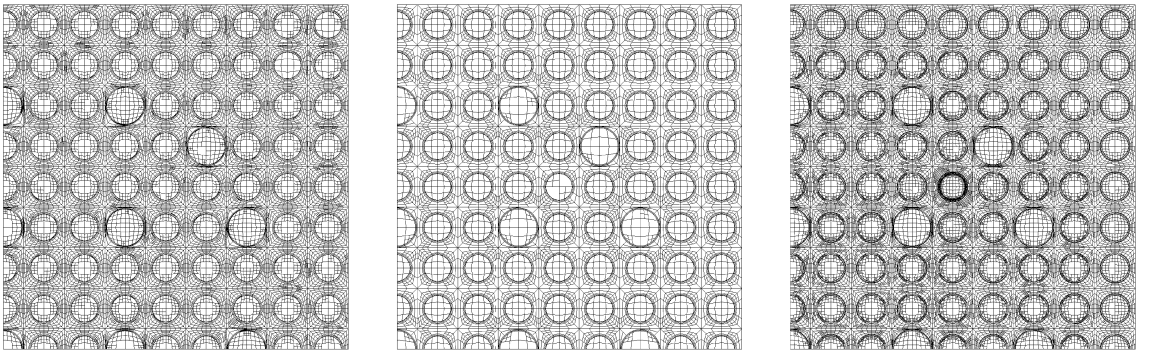
The error in λ for all five mesh refinement strategies is shown in Fig. 20. The heuristic refinement strategies with both a single mesh and GDM perform worse than uniform refinement for this problem. Fig. 21b shows the single mesh method over-refining some fuel pins while under-resolving the regions around the guide tubes compared to goal-based refinement with a single mesh shown in Fig. 21c. With GDM, the heuristic refinement has approximately half as many elements in the group 1 mesh compared to goal-based GDM refinement. This is reversed in group 6, where heuristic refinement has approximately twice as many elements in the group 6 mesh compared to goal-based GDM refinement. This shows that the gradients in the solution that the heuristic refinement strategies are refining around are not necessarily the same regions that are important to accurate calculation of the eigenvalue.

Goal-based refinement with a single mesh is less accurate for the same amount of computational effort than uniform refinement for this problem, although it is possible that this trend is reversing with the most spatially refined meshes considered. As with the problem presented in Section 5.2, goal-based refinement with a single mesh is more accurate than uniform refinement per element in the mesh, but the extra computational effort associated with solving the adjoint problem make it less efficient overall. Goal-based refinement with GDM is slightly more accurate for the same amount of computational effort than uniform refinement after 5 AMR steps have been performed. The reason for this can be seen in Fig. 21. Uniform refinement is simultaneously over-resolving the entire geometry in group 2 and under-resolving the entire geometry in group 1, as well as under-resolving the region around the burnable absorber pin in group 6. This can be seen explicitly in Table 3, in which uniform refinement has $\approx \times\frac{1}{2}$ as many elements as goal-based GDM in group 1, and $\approx \times 3$ as many elements as goal-based GDM in group 2.

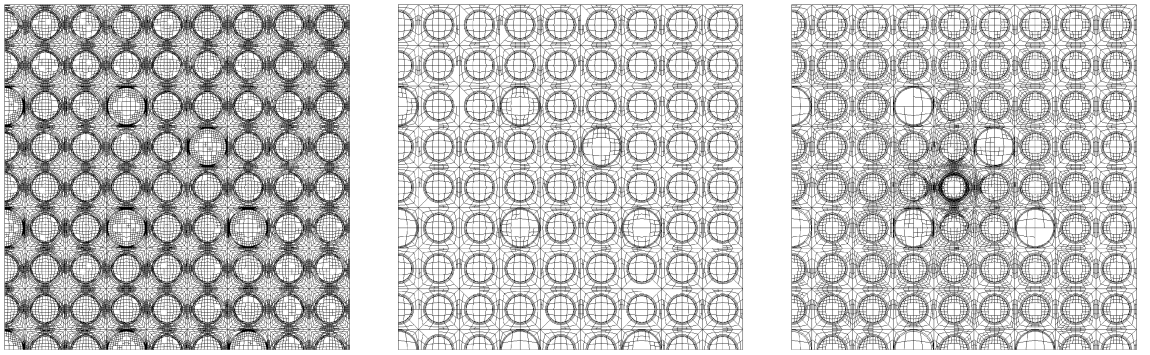
θ_1 and θ_2 are plotted for the goal-based schemes with both a single mesh and GDM for quadratic basis functions in Fig. 22. For both schemes θ_1 is in the range 0.2–0.6 and θ_2 is in the range 0.6–1.8 for every step of the AMR process. Despite the increased complexity of the problem, both θ values are much closer to ideal values than in the problems presented in Sections 5.2 and 5.3. In particular the stable behaviour of θ_2 suggests that $\sum_e |\eta_e|$ should work well as an exit criteria for assembly level calculations if a small safety factor were included.



(a) Uniform Refinement (b) Heuristic Refinement (c) Goal-Based Refinement



(d) Heuristic GDM Refinement Group 1 (e) Heuristic GDM Refinement Group 2 (f) Heuristic GDM Refinement Group 6



(g) Goal-Based GDM Refinement Group 1 (h) Goal-Based GDM Refinement Group 2 (i) Goal-Based GDM Refinement Group 6

Fig. 21. Meshes with approximately 150,000 quadratic elements total for the KAIST MOX assembly problem. Refinement schemes with a single mesh for all groups are shown for uniform (a), heuristic (b) and goal-based (c) refinement strategies, while GDM strategies with heuristic (d, e & f) and goal-based (g, h & i) refinement strategies show the mesh in each group.

5.5. Shielding problem

This test case is representative of a realistic shielding problem for reactor design. The geometry, presented in Fig. 23, consists of a central source region surrounded by concentric annuli of water and steel, along with two void regions.

The material definitions are shown in Table 5, and 67 group macroscopic cross-sections for the regions were generated using the BUGLE-96 cross-section library [43]. There are 47 neutron groups and 20 gamma groups. The source strength in each group is given in Appendix B Table 7, and is representative of the fission spectrum from a PWR. Note that there is no source in the gamma groups, so the contribution to the dose is from neutron and (n, gamma) interactions. The detector region is a circle of radius 15 cm in the outer void region as shown in Fig. 23. The response function is given in Appendix B Table 7, and is designed to represent the dose equivalent as defined in [44] of a human standing in this position relative to the reactor.

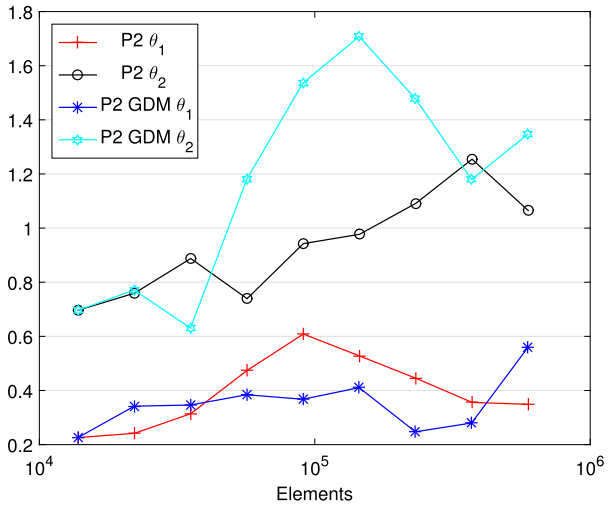


Fig. 22. θ values for the KAIST MOX assembly problem with quadratic basis functions. θ_1 is in the range 0.2–0.6 for all meshes, while θ_2 is consistently in the range 0.6–1.8.

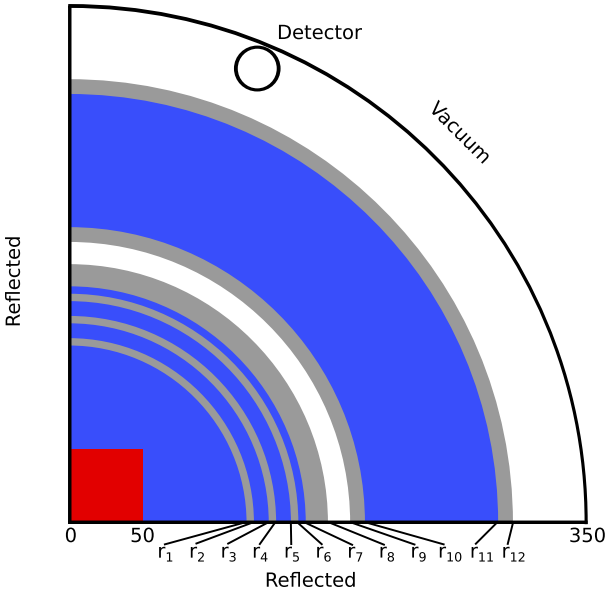


Fig. 23. The geometry for the shielding problem. Blue regions are water, grey regions are steel, white regions are void and the red region represents the core. The annulus radii are given in Table 4. The detector is a circle of radius 15 cm in the outer void region centred at $(x, y) = \left(332.5\cos\left(\frac{3\pi}{8}\right), 332.5\sin\left(\frac{3\pi}{8}\right)\right)$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

Table 4
The annulus radii for the shielding problem in cm.

r_1	r_2	r_3	r_4	r_5	r_6	r_7	r_8	r_9	r_{10}	r_{11}	r_{12}
120	125	135	140	150	155	160	175	190	200	290	300

5.5.1. Spatial convergence study

In order to conduct a study of the spatial convergence of the various schemes, the problem was solved with an S_{16} level symmetric angular quadrature set, and the forward scalar flux in two of the energy groups is shown in Fig. 24. The flux profile of the two groups is qualitatively similar, although the magnitudes involved differ greatly.

Group 13 has the least refined mesh out of the neutron groups with the goal-based GDM scheme with just 784 elements in the mesh used to calculate the reference solution, while group 50 has the most refined mesh with 60,424 elements. Although the detector response function is an order of magnitude larger in group 13 than in group 50, the scalar flux is 8 orders of magnitude larger in the detector region in group 50 than in group 13. The contribution to the detector response

Table 5
The material compositions for the shielding problem with isotopes identified as in BUGLE-96. Silicon is not separated by isotope in BUGLE-96, and H_1 has different thermal treatments in different compounds, hence the specification that it is present in water here. The core is not modelled explicitly here, its fission is represented by the source term.

Material	Density ($g\,cm^{-3}$)	Isotopes	Mass fractions
Core	2	H_1 (H_2O)	0.093
		O_{16}	0.602
		Fe_{56}	0.305
Water	0.9	H_1 (H_2O)	0.1119
		O_{16}	0.8881
Steel	7.81	Mn_{55}	0.02
		Si_{nat}	0.075
		Ni_{58}	0.005
		Cr_{52}	0.1
		Fe_{56}	0.8

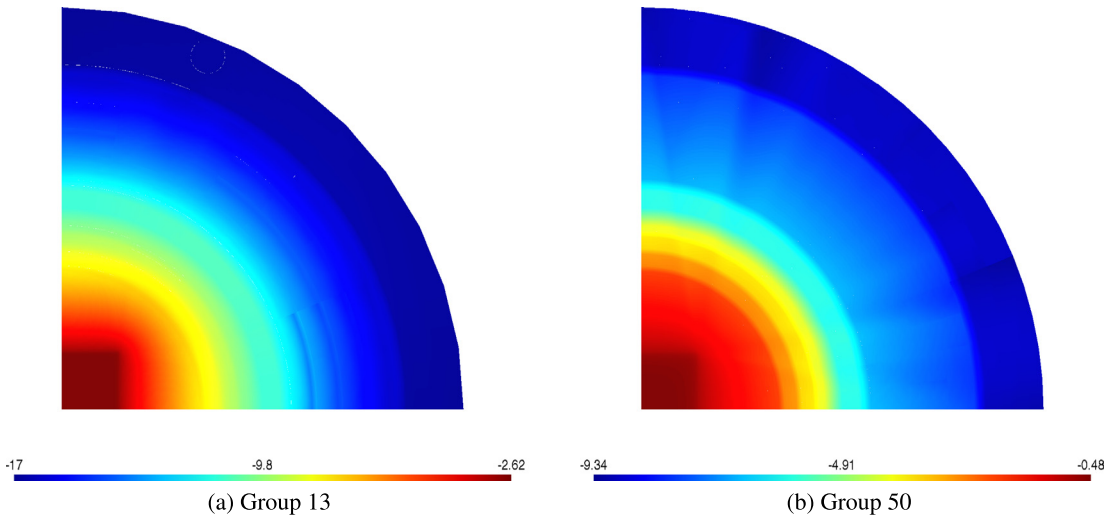


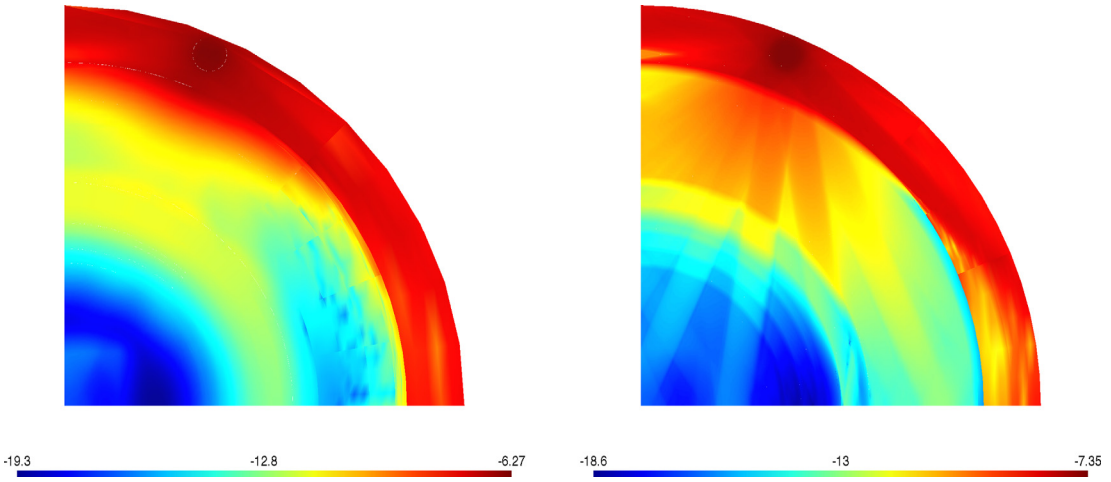
Fig. 24. The scalar flux for the forward problem in group 13 (a) and group 50 (b) on a logarithmic scale for the shielding problem. Group 13 is the least refined neutron group in the goal-based GDM scheme, while group 50 is the most refined group overall (neutron and gamma) in this scheme.

will therefore be approximately 7 orders of magnitude larger in group 50 than in group 13, which explains the disparity in mesh refinement between the two groups in the goal-based GDM scheme.

The adjoint scalar flux in groups 13 and 50 is shown in Fig. 25. The flux profiles in each group are again qualitatively similar, although ray effects are much more visible in group 50 than in group 13 due to the much higher mesh refinement in group 50. The detector response function being an order of magnitude larger in group 13 than in group 50 is shown here in the peak adjoint flux value being approximately an order of magnitude higher in group 13 than in group 50.

The problem was solved with three single-mesh refinement strategies: uniform, heuristic and goal-based, and two GDM refinement strategies: heuristic and goal-based, with quadratic basis functions used in all cases. The reference detector response was calculated on a quadratic mesh with 366,640 total elements, and was found to be 2.41391E−15.

The error in the detector response is presented in Fig. 26 for all five mesh refinement strategies. Both heuristic refinement strategies perform worse than uniform refinement for this problem. In the single mesh case, Fig. 27b shows the mesh refinement being concentrated around the boundary of the core and out to the first steel annulus, with no refinement at all in the outer water tank. In contrast, the single mesh goal-based refinement shown in Fig. 27c focuses refinement along the entire direct path from the core to the detector, including the outer water tank. The difference is even more pronounced with GDM. With heuristic refinement, the meshes for groups 13 and 50, shown in Figs. 27d and 27e respectively, have a similar number of elements in each, and concentrate the refinement near to the core. With goal-based GDM refinement, the group 50 mesh contains $\times 80$ as many elements as the group 13 mesh and refines all the way to the detector region for both groups. Two rays are clearly visible in the group 50 mesh in the outer water tank, suggesting that an important contribution to the detector response for the S_{16} quadrature set employed comes from gamma rays travelling along these paths.



(a) Group 13 (b) Group 50

Fig. 25. The scalar flux for the adjoint problem in group 13 (a) and group 50 (b) on a logarithmic scale for the shielding problem. Group 13 is the least refined neutron group in the goal-based GDM scheme, while group 50 is the most refined group overall (neutron and gamma) in this scheme. The logarithm is taken of the absolute value as there are negative fluxes far from the detector region in both groups.

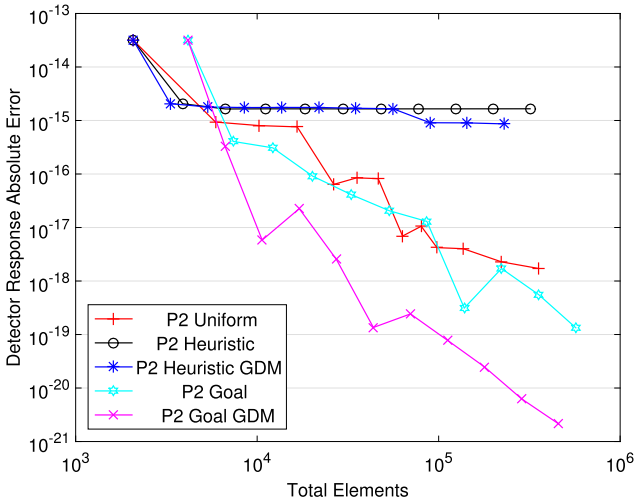
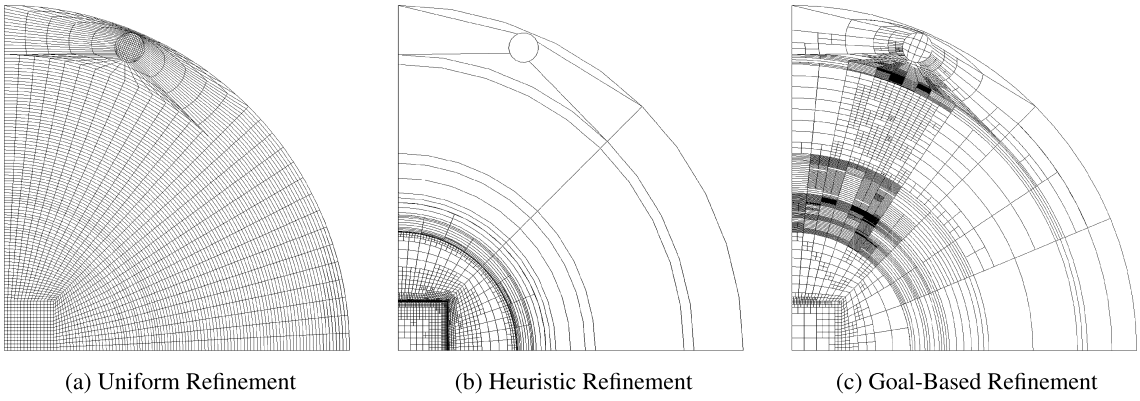


Fig. 26. The convergence of the detector response for the shielding problem as the mesh is refined using uniform, heuristic and goal-based refinement strategies with quadratic basis functions. As with the multigroup fixed source problem, both heuristic refinement schemes converge very slowly as they only refine the mesh near to the source region. The single mesh goal-based scheme has similar accuracy to uniform refinement for the range of spatial refinements considered, while goal-based refinement with GDM is significantly more accurate than both of these schemes.

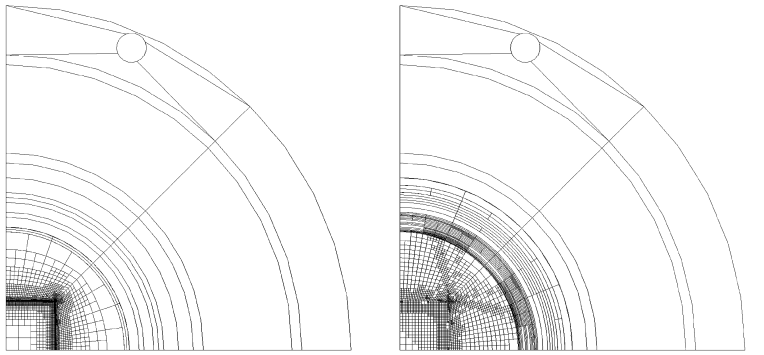
For the same amount of computational effort, goal-based refinement with a single mesh has similar accuracy to uniform refinement for this problem. As with the problems presented in Sections 5.2 and 5.4, the goal-based scheme is more accurate per element in the mesh, but the extra computational effort required to solve the adjoint equation eliminates this advantage.

Goal-based refinement with GDM is a great deal more accurate than the other schemes for this problem. This scheme is approximately 2 orders of magnitude more accurate than uniform refinement for the same amount of computational effort. Viewed from the other perspective, to achieve the same level of accuracy in the detector response, $\approx \times 8$ less computational effort is required in the goal-based GDM scheme compared to uniform refinement for this problem.

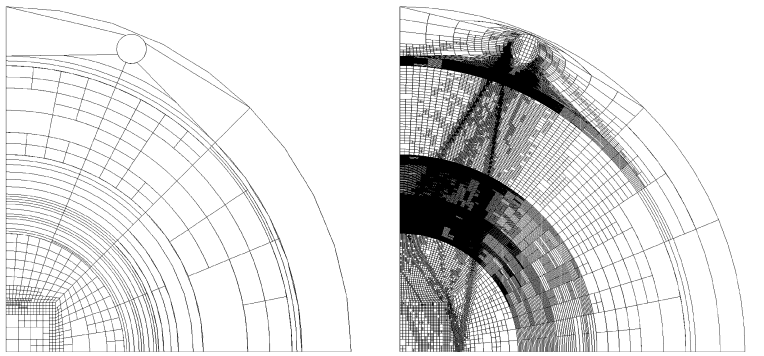
θ values for this problem are presented in Fig. 28. For both the single mesh and GDM methods, there are large spikes in both θ values. These correspond to unusually low errors in Fig. 26, which are the result of the error changing in sign. Excluding these, θ_1 lies in the range 0.5–1.1, while θ_2 is generally between 1.3 and 12. As in the Cartesian grid multigroup fixed source problem presented in Section 5.3, θ_2 only ever overestimates the error for both the single mesh and GDM goal-based schemes, and so $\sum_e |\eta_e|$ could be used as an appropriate exit criteria from the AMR process for both schemes.



(a) Uniform Refinement (b) Heuristic Refinement (c) Goal-Based Refinement



(d) Heuristic GDM Refinement Group 13 (e) Heuristic GDM Refinement Group 50



(f) Goal-Based GDM Refinement Group 13 (g) Goal-Based GDM Refinement Group 50

Fig. 27. Meshes with approximately 200,000 quadratic elements total for the shielding problem. Refinement schemes with a single mesh for all groups are shown for uniform (a), heuristic (b) and goal-based (c) refinement strategies, while GDM strategies with heuristic (d & e) and goal-based (f & g) refinement strategies show the mesh in groups 13 and 50. Both heuristic refinement schemes have refined only near to the source region where the flux gradients are highest. In contrast even the least refined neutron group (13) has been refined in the outer water tank in order to improve the detector response, while the most refined group (50) has been densely refined everywhere between the source and detector, particularly in the steel annuli making up the bulk of the gamma shielding.

5.5.2. Space–angle study

Although the error estimators presented here are designed to only reduce the error due to the spatial discretisation, these errors are frequently strongly coupled to those due to the angular discretisation. This is particularly the case when ray effects dominate the solution, and therefore the mesh refinement, as in Fig. 27g. Therefore, a small space–angle study was conducted using the shielding problem. The reference detector response was calculated using an S_{64} triangular-Chebyshev angular quadrature set with 143,215 total quadratic elements using the goal-based GDM scheme, and was found to be 2.73764E−15. The same AMR scheme was then used with varying angular quadrature orders, in order to study the convergence of the detector response as both the spatial and angular resolution increases.

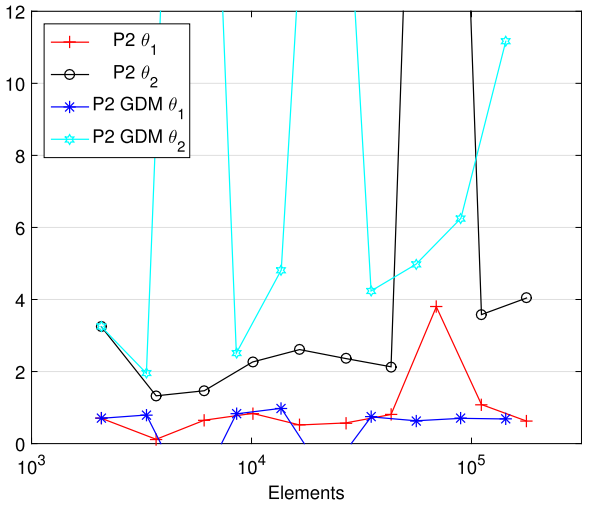


Fig. 28. θ values for the shielding problem with quadratic basis functions. The large spikes correspond to abnormally low errors in Fig. 26, which in turn correspond to a change in sign of the error.

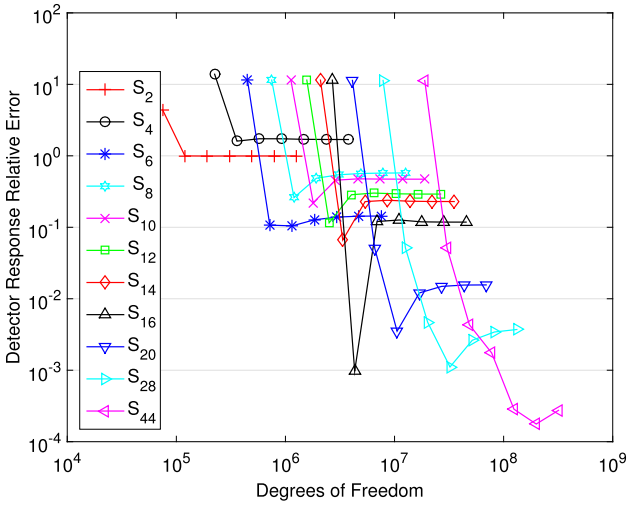


Fig. 29. The convergence of the detector response for the shielding problem as both the angular and spatial resolution are varied. The relative error in the detector response has been used in this plot to facilitate comparison between the results.

The results of this study are shown in Fig. 29, presented in terms of relative error in the detector response here to facilitate comparison between the results. For all of the angular discretisations, the coarsest spatial mesh gives a relative error of $\approx 1000\%$, and so the spatial discretisation error is clearly dominating at this point. For each angular quadrature order, the error then decreases as the spatial mesh is refined, before saturating at some point, and the higher the angular resolution, the more spatial resolution is required to reach this saturation point (this point has not yet been reached with the S_{44} quadrature). This spatially saturated detector response is monotonically decreasing as the angular resolution is increased when an S_8 or higher order angular quadrature is used.

At the lower S_N orders ($N \leq 16$), there is at least a 10% error in the detector response remaining due to the angular discretisation, and performing any more than 3 AMR iterations at these angular resolutions provides little benefit in accuracy. At the higher orders, it can be seen that to obtain a highly accurate detector response, the angular and spatial resolutions must be increased simultaneously.

The benefit of increasing the total space–angle resolution is also diminishing. For example, to reduce the error from $\approx 10\%$ at the 3rd AMR step using S_{16} to $\approx 1\%$ at the 4th AMR step using S_{20} requires $\approx \times 2.5$ as many degrees of freedom. To then reduce the error further to $\approx 0.1\%$ at the 4th AMR step using S_{44} requires $\approx \times 5$ as many degrees of freedom.

Table 6
Memory usage in gigabytes of the Mixed Mass Matrices (MMM) and Flux Moments (FM), along with the ratio between the two $\frac{\text{MMM}}{\text{FM}}$, for the first 8 AMR steps of the KAIST MOX assembly and shielding problems.

AMR step	KAIST			Shielding		
	MMM	FM	Ratio	MMM	FM	Ratio
1	9.75E−2	9.94E−4	98.0	1.40E−1	1.50E−4	938.1
2	1.88E−1	1.59E−3	118.3	2.70E−1	2.40E−4	1128.8
3	3.26E−1	2.55E−3	127.8	4.69E−1	3.84E−4	1222.4
4	5.29E−1	4.08E−3	129.9	8.11E−1	6.14E−4	1320.6
5	8.26E−1	6.52E−3	126.6	1.40E+0	9.83E−4	1426.2
6	1.33E+0	1.04E−2	127.4	2.34E+0	1.57E−3	1490.4
7	2.19E+0	1.67E−2	130.9	3.90E+0	2.52E−3	1549.7
8	3.50E+0	2.67E−2	131.1	6.37E+0	4.03E−3	1581.2

5.6. Memory usage

As mentioned in Section 3.2, G^2 mixed mass matrices are stored when GDM are employed in the current implementation. The memory usage of these CRS matrices is compared to that of the flux moments for the first 8 AMR steps for the goal-based scheme with GDM for the KAIST MOX assembly and shielding problems in Table 6.

At the first AMR step, every mixed mass matrix has the same sparsity pattern, and so the ratio between the memory required for the mixed mass matrices and the flux moments (which is just the scalar flux here) can be theoretically calculated as $\approx 14G$ with the quadratic NURBS employed. This ratio does not significantly increase as the meshes are adaptively refined, but is large even for a small number of energy groups. It should be noted that only isotropic scatter is considered here, which would be insufficient for realistic shielding calculations. If, for example, P_7 scatter were employed, the shielding problem memory ratios would drop by a factor of 36 to 26–44. In addition, the same accuracy of detector response can be achieved using $\approx \times 8$ times fewer unknowns with goal-based GDM than with uniform refinement, dropping the ratio of the memory requirements still further to 3–6. In general, the dominant factor limiting the size of tractable transport calculations is currently CPU time rather than memory, so the increased speed offered by the proposed method at the price of increased memory usage could be extremely useful in practice.

6. Conclusion

In this paper, a new spatial discretisation for the multigroup, discrete ordinates equations has been developed in which each energy group is spatially discretised over a different mesh. These meshes are created using NURBS, ensuring that the geometry is exact from the coarsest level of refinement. They also utilise the hanging-node spatial discretisation developed in [20], to allow the use of adaptive schemes without the refinement propagation present in standard tensor-product NURBS discretisations [13]. The initial meshes for every energy group are the same, and are subsequently refined by bisecting elements in each parametric direction. This greatly simplifies the generation of the supermesh needed to calculate the mixed mass matrices for transferring information between meshes compared to using arbitrary meshes in each energy group as in [23,24]. Crucially, this approach can not be taken with standard finite elements in the presence of curved boundaries, which are extremely prevalent in reactor physics and shielding applications, as well as engineering applications in general. As was demonstrated in [20], the use of an adaptive refinement scheme using an inexact polygonal representation of the geometry as a starting point leads to a plateauing of the error, as eventually the geometrical errors dominate the system.

The order of accuracy of the GDM scheme was checked using a manufactured smooth solution. The convergence rate for NURBS of order p was found to be $p+1$ for both orthogonal grids and rational geometries generated using a pincell-style mesh.

In addition, DWR error metrics for the multigroup, discrete ordinates equations with discontinuous Galerkin spatial discretisation have been developed for both fixed source and eigenvalue problems. In both cases an approximation to the exact solution to the multigroup, discrete ordinates adjoint equation is required, and in this work this was calculated using the same mesh as the forward equation is being solved on, but with NURBS of one degree higher than those used to solve the forward equation. In practice, it is expected that computationally cheaper approximations to the exact adjoint solution, such as those developed in [31], could be used, and so efficiency comparisons have been made on this basis.

For the one group eigenvalue and multigroup fixed source problems, the geometry is Cartesian and so the performance of the error estimators can be compared for linear and quadratic basis functions. For both of these problems the error estimators used with the linear basis functions perform extremely well, with $\theta_1 \approx 1$ and θ_2 providing an asymptotically sharp error bound as the mesh is refined. This is to be expected based on similar work presented in [30,31]. However the performance of the error metrics degrades for both problems when quadratic basis functions are used. This is suspected to be due to the smaller difference in accuracy between cubic and quadratic solutions than between quadratic and linear solutions, so that the difference between the "exact" adjoint solution ψ^* and the discrete version ψ_h^* is less accurate in the quadratic case.

Despite this, the performance of the quadratic solutions is just as good as the linear solutions in terms of the detector response when goal-based adaptivity is used compared to heuristic adaptivity or uniform solutions. This suggests that although the error estimator values in each element are degraded with quadratic basis functions, their absolute value relative to each other are not. As it is these absolute values that are used to drive the adaptivity, the resulting meshes are still close to optimal for the requested detector response.

For the KAIST MOX assembly problem, goal-based refinement with a single mesh for all energy groups is slightly less accurate than uniform refinement for the same amount of computational effort. This is due to the added cost of solving the adjoint equation, as the scheme is more accurate than uniform refinement on a per element basis. With GDM, goal-based refinement is more accurate than uniform refinement for this problem for the same amount of computational effort after a few AMR steps. As for the Cartesian grid problems, the quadratic basis functions used here mean that the metrics θ_1 and θ_2 do not provide an accurate measure of the error or an asymptotically sharp error bound respectively. However both values are the correct order of magnitude, and so the error estimators could potentially be used as an exit criteria from an AMR process with the inclusion of a small safety factor.

For the spatial convergence study of the shielding problem, goal-based refinement with a single mesh gives approximately the same level of accuracy in the detector response as uniform refinement for the same amount of computational effort. As with the KAIST MOX assembly problem, the goal-based scheme with a single mesh is more accurate than uniform refinement on a per element basis, but the added cost of solving the adjoint equation eliminates this advantage. Conversely, goal-based refinement with GDM is significantly more accurate than uniform refinement for this problem. The same accuracy in the detector response can be achieved with $\approx \times 8$ less computational effort with the goal-based GDM scheme. As with the other quadratic solutions, the metrics θ_1 and θ_2 are not perfect, but are both within an order of magnitude of optimal values, and so again the error estimators could be used as an exit criteria from an AMR process with a small safety factor included.

The use of GDM for the multigroup problems with the goal-based refinement is more accurate than using a single mesh in all cases. Unsurprisingly, the more energy groups being used in the problem the greater the difference in accuracy between the single mesh and GDM methods, from a factor of 2 for a 2 group problem through to a factor of $\approx \times 77$ for a 67 group problem. This comes with a corresponding increase in storage requirements, with the ratio between the memory required to store the mixed mass matrices to the memory required to store the flux moments approximately proportional to G. However, the limiting factor on the size of transport calculations that can be performed is often currently CPU time rather than memory. Combined with the high order scatter expansion required for shielding problems and the potential savings in computational effort demonstrated in Section 5.5, the presented method could prove beneficial for shielding calculations.

Overall, the multigroup error estimators for eigenvalue problems appear to offer little benefit for the reactor physics problems considered. This does not appear to be down to the error estimators being any less accurate than the fixed source version, but rather that the accurate calculation of the eigenvalue requires relatively uniform mesh resolution both throughout the physical domain and across the energy groups. Therefore, the extra computational effort required to solve the adjoint problem does not provide sufficient additional accuracy to offset its expense. As the performance of the GDM method improves with a greater number of energy groups, it is possible that greater benefits would be obtained when using the 20+ energy groups common in reactor physics applications.

This is not the case in radiation shielding problems, in which the multigroup error estimators are very promising for both driving adaptive refinement and quantifying the remaining error in the detector response due to the spatial discretisation. With linear basis functions their performance is indistinguishable from the 1 group case presented in [30,31]. With the quadratic NURBS required to model most realistic geometries accurately, the error estimators seem to drive the mesh refinement in an optimal way, leading to large savings in computational effort compared to uniform refinement. While the error estimators do not provide sharp bounds on the error with quadratics as they do with linears, they are still useful for gauging the order of magnitude of the error remaining in the detector response due to the spatial discretisation. As nuclear reactor design and safety case justification increasingly moves towards best estimate plus uncertainty, this ability to quantify the error is very powerful. One option to improve this error quantification would be to drive the adaptivity in the manner described in this paper until a final detector response is calculated due to some exit criteria. The error in this value due to the spatial discretisation could then be more accurately computed by evaluating a more accurate "exact" adjoint solution than is used to drive the refinement, for example by using even higher order NURBS or by subdividing each element in the mesh.

The error estimators presented here are based on angularly integrated quantities in each element. In theory the group-dependent mesh approach could be taken one step further, with each ordinate in each group having it's own mesh as well. Error estimators for each direction could be derived in an almost identical manner to those presented in Section 4 and an AMR procedure could be constructed analogously to the one used here. This could be of use in deep penetration shielding problems in which the angular flux contributing to the detector response is highly anisotropic. However this would require the storage of $O((NG)^2)$ mixed mass matrices, as well as complicating the calculation of flux moments which could impact the performance of the scatter iteration within the solver.

An intermediate approach would be to allow each group to use a different angular quadrature set. This would keep the number of mixed mass matrices to be stored at $O(G^2)$ and would not introduce any additional complications as the angular fluxes in each energy group are coupled only through their moments. The spatial error estimators derived here could then

be used to drive mesh refinement within a group, and an angular error estimator could be used to select which groups require angular refinement.

A further consideration is that the error metrics presented here only apply to spatial errors, and give no information about the error introduced by the angular discretisation. As demonstrated in the space–angle study of the shielding problem, angular and spatial errors can be tightly coupled, and require simultaneous refinement to reduce the total error in the detector response. In order to derive error estimators that include the angular error, a variational method based on the entire space–angle phase space would be required, such as the space–angle finite elements introduced in reference [45]. Combining the hanging-node NURBS discretisation in space with angular finite elements (possibly also NURBS) that are also locally refinable would allow schemes analogous to the one presented here to be developed to reduce the total error due to the space–angle discretisation, as well as quantify what this error is.

Although the discretisation presented here is for the multigroup, discrete ordinates equations, it should be stressed that it is equally applicable to other systems of equations. Without modification the GDM spatial discretisation could be applied to the radiative transfer equation [46] and the Boltzmann–Fokker–Planck equation for charged particle transport, with applications in radiotherapy [47], for example.

Even more interesting possibilities exist in multiphysics applications. In the GDM method presented here, every mesh was constructed from NURBS of the same order. This was deliberate, as the equations being solved over each mesh are qualitatively identical, varying only in piecewise constant cross-section and source values. This is not the case in many engineering applications, in which different solution fields satisfy qualitatively different equations. For example, in [48], DWR error estimators were derived for coupled flow and heat transfer problems in which all solution fields were defined over the same finite element space. In this case, the inf-sup condition was satisfied only because of numerical stability introduced by the solution method. The GDM method presented here could be extended to a more general solution-field dependent mesh method. In this case, the pressure field could then be solved using lower order NURBS than the velocity field in order to satisfy the inf-sup condition.

This idea can be taken even further, to so called fine-mesh coupled neutronics thermal-hydraulics modelling of 3D quarter-assembly sized problems [49]. The method presented by Jareteg et al. discretised the multigroup discrete ordinates equations and the thermal-hydraulics equations over independently assembled finite volume meshes, and then interpolated cell quantities between them in a conservative manner based on the volumetric overlap of cells on each mesh. With the method presented here, these meshes could be derived from a common coarsest NURBS description of the geometry, simultaneously eliminating the need for a polygonal intersection algorithm to identify volumetric overlap, and allowing higher order discretisations of the governing equations to be used.

Acknowledgements

A.R. Owens would like to acknowledge the support of EPSRC under their industrial doctorate programme (EPSRC grant number: EP/G037426/1), Rolls-Royce for industrial support and the Imperial College High Performance Computing Service for technical support. J.A. Welch would like to acknowledge EPSRC through the Doctoral Training Award (DTA) PhD scheme and also the support of Rolls-Royce. M.D. Eaton and J. Kópházi 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). In accordance with EPSRC funding requirements, all supporting data used to create figures and tables in this paper may be accessed at the following URL: http://doi.org/10.5281/zenodo.193039.

Appendix A. Adjoint consistency proof

This is much simpler to prove by representing the bilinear forms as global weak forms, where the numerical upwinding is defined in terms of the average values and jumps in the trial and weight functions, as shown in [50]. As both the forward and adjoint bilinear forms are symmetric in their arguments in the scatter terms, it is sufficient to consider a decoupled equation for single ordinate:

$$\boldsymbol{\Omega}_k \cdot \nabla \psi_k + \Sigma_t \psi_k = 0 \quad \text{for the forward equation and} \tag{A.1a}$$

$$-\boldsymbol{\Omega}_k \cdot \nabla \psi_k^* + \Sigma_t \psi_k^* = 0 \quad \text{for the adjoint equation} \tag{A.1b}$$

where the g superscripts have been dropped for clarity and the source terms have been left as 0 as we are only interested in the bilinear form associated with the LHS. The domain boundary ∂V must be split into incoming ∂V^- and outgoing ∂V^+ portions where $\boldsymbol{\Omega}_k \cdot \mathbf{n}$ is negative and positive respectively. The set of elements in the mesh is denoted by $\mathcal{E}$. The set of boundary faces is denoted by $\mathcal{F}_b$, with outward unit normal $\mathbf{n}_F(\mathbf{r})$. The set of interior faces, i.e. not on a boundary, is denoted by $\mathcal{F}_i$. Each face $F \in \mathcal{F}_i$ has an associated unit normal vector $\mathbf{n}_F(\mathbf{r})$ which is a function of space due to the curved faces. Each face $F \in \mathcal{F}_i$ separates two elements, which are locally numbered as V_1 and V_2 in order to define the jump operator. The unit normal $\mathbf{n}_F(\mathbf{r})$ is then chosen to point from V_1 towards V_2.

As basis functions can be discontinuous across faces, average and jump operators of a function across a face can be defined as:

$$\{\{v\}\}_F(\mathbf{r}) := \frac{1}{2}(v|_{V_1}(\mathbf{r}) + v|_{V_2}(\mathbf{r})) \qquad \text{average operator} \tag{A.2a}$$

$$[[v]]_F(\mathbf{r}) := v|_{V_1}(\mathbf{r}) - v|_{V_2}(\mathbf{r}) \qquad \text{jump operator} \tag{A.2b}$$

The forward and adjoint bilinear forms for the simplified problem defined in Equation (A.1) can then be defined as [50]:

$$\begin{aligned} b(\psi_k, v) = &\sum_{V_e\in\mathcal{E}} \int_{V_e} d\mathbf{r}\left(-\psi_k\mathbf{\Omega}_k\cdot\nabla_h v + \Sigma_t\psi_k v\right) + \int_{\partial V^+} dS\ (\mathbf{\Omega}_k\cdot\mathbf{n})\psi_k v\ + \ldots \\ &\ldots + \sum_{F\in\mathcal{F}_i}\int_F dS\ \left((\mathbf{\Omega}_k\cdot\mathbf{n}_F)\{\{\psi_k\}\}[[v]] + \frac{1}{2}|\mathbf{\Omega}_k\cdot\mathbf{n}_F|\,[[\psi_k]]\,[[v]]\right) \end{aligned} \tag{A.3a}$$

$$\begin{aligned} b^*(\psi_k^*, z) = &\sum_{V_e\in\mathcal{E}} \int_{V_e} d\mathbf{r}\left(\psi_k^*\mathbf{\Omega}_k\cdot\nabla_h z + \Sigma_t\psi_k^* z\right) - \int_{\partial V^-} dS\ (\mathbf{\Omega}_k\cdot\mathbf{n})\psi_k^* z\ + \ldots \\ &\ldots + \sum_{F\in\mathcal{F}_i}\int_F dS\ \left(-(\mathbf{\Omega}_k\cdot\mathbf{n}_F)\{\{\psi_k^*\}\}[[z]] + \frac{1}{2}|\mathbf{\Omega}_k\cdot\mathbf{n}_F|\,[[\psi_k^*]]\,[[z]]\right) \end{aligned} \tag{A.3b}$$

Note that these bilinear forms are mathematically identical to those derived in an element-wise manner in Section 3.1, and so can still be solved by sweeping through the mesh. This notation is just more convenient when proving that $b(u,v) = b^*(v,u)$. To do so, consider the difference:

$$\begin{aligned} b(u,v) - b^*(v,u) = &\sum_{V_e\in\mathcal{E}}\int_{V_e} d\mathbf{r}\left(-u\mathbf{\Omega}_k\cdot\nabla v - v\mathbf{\Omega}_k\cdot\nabla u + \Sigma_t(uv - vu)\right) + \int_{\partial V^+\cup\partial V^-} dS\ (\mathbf{\Omega}_k\cdot\mathbf{n})uv\ + \ldots \\ &\ldots + \sum_{F\in\mathcal{F}_i}\int_F dS\ \left((\mathbf{\Omega}_k\cdot\mathbf{n}_F)(\{\{u\}\}\,[[v]] + \{\{v\}\}\,[[u]]) + \frac{1}{2}|\mathbf{\Omega}_k\cdot\mathbf{n}_F|([[u]]\,[[v]] - [[v]]\,[[u]])\right) \end{aligned} \tag{A.4}$$

Clearly $(uv - vu) = ([[u]]\ [[v]] - [[v]]\ [[u]]) = 0$ and $\partial V^+ \cup \partial V^- = \partial V$. Also, note that $\int_{\partial V} = \sum_{F\in\mathcal{F}_b}\int_F$ and that $-u\mathbf{\Omega}_k\cdot\nabla v - v\mathbf{\Omega}_k\cdot\nabla u = -\nabla\cdot(uv\mathbf{\Omega}_k)$. Equation (A.4) can therefore be written in the more compact form:

$$\begin{aligned} b(u,v) - b^*(v,u) = &\sum_{V_e\in\mathcal{E}}\int_{V_e} d\mathbf{r}\left(-\nabla\cdot(uv\mathbf{\Omega}_k)\right) + \sum_{F\in\mathcal{F}_b}\int_F dS\ (\mathbf{\Omega}_k\cdot\mathbf{n}_F)uv \\ &+ \sum_{F\in\mathcal{F}_b}\int_F dS\ (\mathbf{\Omega}_k\cdot\mathbf{n}_F)(\{\{u\}\}\,[[v]] + \{\{v\}\}\,[[u]]) \end{aligned} \tag{A.5}$$

Simple algebra yields $(\{\{u\}\}\,[[v]] + \{\{v\}\}\,[[u]]) = [[uv]]$. Applying the divergence theorem to the first term then yields:

$$b(u,v) - b^*(v,u) = \sum_{V_e\in\mathcal{E}}\int_{\partial V_e} dS(\mathbf{\Omega}_k\cdot\mathbf{n}_e)(-uv) + \sum_{F\in\mathcal{F}_b}\int_F dS\ (\mathbf{\Omega}_k\cdot\mathbf{n}_F)\,uv + \sum_{F\in\mathcal{F}_i}\int_F dS\ (\mathbf{\Omega}_k\cdot\mathbf{n}_F)\,[[uv]] \tag{A.6}$$

where it is important to note that $\mathbf{n}_e$ is the outward unit normal to element V_e. However on the domain boundary this definition of the normal vector coincides with that of $\mathbf{n}_F$, therefore:

$$\sum_{V_e\in\mathcal{E}}\int_{\partial V_e\cap\partial V} dS(\mathbf{\Omega}_k\cdot\mathbf{n}_e)(-uv) = \sum_{F\in\mathcal{F}_b}\int_F dS\ (\mathbf{\Omega}_k\cdot\mathbf{n}_F)(-uv) \tag{A.7}$$

The domain boundary integrals therefore cancel out and Equation (A.6) becomes:

$$b(u,v) - b^*(v,u) = \sum_{V_e\in\mathcal{E}}\int_{\partial V_e\setminus\partial V} dS(\mathbf{\Omega}_k\cdot\mathbf{n}_e)(-uv) + \sum_{F\in\mathcal{F}_i}\int_F dS\ (\mathbf{\Omega}_k\cdot\mathbf{n}_F)\,[[uv]] \tag{A.8}$$

The first term now also represents integrals over internal faces of the mesh $F\in\mathcal{F}_i$. Using the definition of the jump operator, and the fact that $\mathbf{n}_F = \mathbf{n}_1 = -\mathbf{n}_2$ where $\mathbf{n}_i$ is the outward pointing normal vector to element V_i in the face-local numbering system:

$$\sum_{F\in\mathcal{F}_i}\int_F dS\ (\boldsymbol{\Omega}_k\cdot\mathbf{n}_F)\ [[uv]] = \sum_{F\in\mathcal{F}_i}\int_F dS\ (\boldsymbol{\Omega}_k\cdot\mathbf{n}_F)\big((uv)|_{V_1}-(uv)|_{V_2}\big) \tag{A.9}$$

$$= \sum_{F\in\mathcal{F}_i}\left(\int_{F\cap\partial V_1} dS\ (\boldsymbol{\Omega}_k\cdot\mathbf{n}_1)\ uv + \int_{F\cap\partial V_2} dS\ (\boldsymbol{\Omega}_k\cdot\mathbf{n}_2)\ uv\right) \tag{A.10}$$

$$= \sum_{V_e\in\mathcal{E}}\int_{\partial V_e\setminus\partial V} dS(\boldsymbol{\Omega}_k\cdot\mathbf{n}_e)\ uv \tag{A.11}$$

Therefore the RHS of Equation (A.8) is zero, $b(u,v)=b^*(v,u)$ and therefore $B(u,v)=B^*(v,u)$ in the full bilinear forms presented in Section 4.

Appendix B. Shielding problem source strength & response function

Table 7
The source strength is a typical PWR fission spectrum, while the response function has been calculated using the product of kerma factors and quality factors given in [44].

Neutrons group	Source strength	Response function	Neutrons group	Source strength	Response function	Gammas group	Source strength	Response function
1	3.870867e−06	4.25E−008	25	3.185330e−03	1.10E−008	48	0	2.78E−009
2	1.685512e−05	3.67E−008	26	1.675458e−03	6.53E−009	49	0	2.78E−009
3	1.025741e−04	3.49E−008	27	8.221569e−04	4.43E−009	50	0	1.97E−009
4	2.445024e−04	3.66E−008	28	2.624933e−04	2.68E−009	51	0	1.75E−009
5	5.561648e−04	3.48E−008	29	2.474941e−04	2.68E−009	52	0	1.56E−009
6	1.602635e−03	3.29E−008	30	6.444953e−05	1.54E−009	53	0	8.78E−010
7	3.085890e−03	3.02E−008	31	1.963815e−05	1.28E−009	54	0	8.78E−010
8	8.133662e−03	2.93E−008	32	2.301996e−05	9.04E−010	55	0	5.54E−010
9	7.647617e−03	2.79E−008	33	6.136414e−05	9.04E−010	56	0	4.40E−010
10	4.329158e−03	2.45E−008	34	5.514284e−05	4.78E−010	57	0	2.78E−010
11	4.559032e−03	2.38E−008	35	1.797236e−05	1.22E−010	58	0	2.78E−010
12	1.949293e−03	2.21E−008	36	5.845122e−06	5.01E−011	59	0	2.78E−010
13	3.915004e−04	2.21E−008	37	2.381457e−06	2.58E−011	60	0	2.78E−010
14	2.365494e−03	2.20E−008	38	2.914192e−07	8.37E−012	61	0	1.39E−010
15	7.147603e−03	2.14E−008	39	9.463183e−08	4.41E−012	62	0	6.98E−011
16	7.026500e−03	2.02E−008	40	3.533897e−08	2.38E−012	63	0	2.78E−011
17	8.833207e−03	1.85E−008	41	8.594290e−09	2.07E−012	64	0	2.78E−011
18	1.149367e−02	2.43E−008	42	1.051302e−09	2.91E−012	65	0	3.50E−011
19	6.305688e−03	2.12E−008	43	3.926092e−10	3.93E−012	66	0	2.78E−010
20	2.745129e−03	1.98E−008	44	7.616870e−11	7.62E−012	67	0	5.54E−010
21	4.725606e−03	1.86E−008	45	2.472356e−11	1.07E−011			
22	3.803092e−03	1.72E−008	46	1.047500e−11	2.26E−011			
23	4.266816e−03	1.64E−008	47	1.000000e−21	5.30E−011			
24	2.221892e−03	1.34E−008						

References

[1] K.S. Smith, Assembly homogenization techniques for light water reactor analysis, Prog. Nucl. Energy 17 (1986) 303–335.
[2] Richard Sanchez, Assembly homogenization techniques for core calculations, Prog. Nucl. Energy 51 (2009) 14–31.
[3] R.D. Lawrence, Progress in nodal methods for the solution of the neutron diffusion and transport equations, Prog. Nucl. Energy 17 (1986) 271–301.
[4] W.H. Reed, Triangular Mesh Difference Schemes for the Transport Equation, Technical Report LA-4769, Los Alamos Scientific Laboratory, 1971.
[5] W.H. Reed, T.R. Hill, Triangular Mesh Methods for the Neutron Transport Equation, Technical Report LA-UR-73-479, Los Alamos Scientific Laboratory, 1973.
[6] Todd A. Wareing, John M. McGhee, Jim E. Morel, Shawn D. Pautz, Discontinuous finite element S_N methods on three-dimensional unstructured grids, Nucl. Sci. Eng. 138 (2001) 256–268.
[7] B.Q. Li, Discontinuous Finite Elements in Fluid Dynamics and Heat Transfer, Springer-Verlag, 2006.
[8] F. Bassi, S. Rebay, A high-order accurate discontinuous finite element method for the numerical solution of the compressible Navier–Stokes equations, J. Comput. Phys. 131 (1997) 267–279.
[9] Culbert B. Laney, Computational Gasdynamics, Cambridge University Press, 1998.
[10] P.J. Frey, P.-L. George, Mesh Generation: Application to Finite Elements, Wiley, 2008.
[11] J.-Y. Moller, J.-J. Lautard, MINARET, a deterministic neutron transport solver for nuclear core calculations, in: International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering, 2011.
[12] K. Manalo, M. Chin, G. Sjoden, Cartesian meshing impacts for PWR assemblies in multigroup Monte Carlo and S_N transport, in: Joint International Conference on Supercomputing in Nuclear Applications and Monte Carlo 2013, 2013.
[13] 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.
[14] 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.

[15] 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.
[16] F. Brunero, Discontinuous Galerkin Methods for Isogeometric Analysis, Master's thesis, Università degli Studi di Milano, 2012.
[17] Ulrich Langer, Angelos Mantzaflaris, Stephen E. Moore, Ioannis Toulopoulos, Multipatch discontinuous Galerkin isogeometric analysis, in: Isogeometric Analysis and Applications 2014, Springer International Publishing, 2015, pp. 1–32, chapter 1.
[18] C. Michoski, J. Chan, L. Engvall, J.A. Evans, Foundations of the blended isogeometric discontinuous Galerkin (BIDG) method, Comput. Methods Appl. Mech. Eng. 305 (2016) 658–681.
[19] 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.
[20] A.R. Owens, J.A. Welch, J. Kópházi, M.D. Eaton, An adaptive, hanging-node, discontinuous isogeometric analysis method for the first-order form of the neutron transport equation with discrete ordinate (S_N) angular discretisation, Comput. Methods Appl. Mech. Eng. (2017), http://dx.doi.org/10.1016/j.cma.2017.01.036, in press.
[21] Jean C. Ragusa, A simple Hessian-based 3D mesh adaptation technique with applications to the multigroup neutron diffusion equations, Ann. Nucl. Energy 35 (2008) 2006–2018.
[22] Yaqi Wang, Wolfgang Bangerth, Jean Ragusa, Three-dimensional *h*-adaptivity for the multigroup neutron diffusion equations, Prog. Nucl. Energy 51 (2009) 543–555.
[23] Mark A. Goffin, Christopher M.J. Baker, Andrew G. Buchan, Christopher C. Pain, Matthew D. Eaton, Paul 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.
[24] P.E. Farrell, J.R. Maddison, Conservative interpolation between volume meshes by local Galerkin projection, Comput. Methods Appl. Mech. Eng. 200 (2011) 89–100.
[25] Wolfgang Bangerth, Rolf Rannacher, Adaptive Finite Element Methods for Differential Equations, Springer, Basel AG, 2003.
[26] R. Rannacher, Error control in finite element computations, in: H. Bulgak, C. Zenger (Eds.), Summer School Error Control and Adaptivity in Scientific Computing, Kluwer Academic Publishers, 1998, pp. 247–278.
[27] R. Becker, R. Rannacher, An optimal control approach to error estimation and mesh adaptation in finite element methods, Acta Numer. 10 (2001) 1–101.
[28] R. Rannacher, Duality techniques for error estimation and mesh adaptation in finite element methods, in: E. Stein (Ed.), Adaptive Finite Elements in Linear and Nonlinear Solid and Structural Mechanics, Springer, 2002.
[29] R. Rannacher, F.T. Suttmeier, Error estimation and adaptive mesh design for FE models in elasto-plasticity, in: E. Stein (Ed.), Error-Controlled Adaptive FEMs in Solid Mechanics, John Wiley, 2002.
[30] D. Lathouwers, Goal-oriented spatial adaptivity for the S_N equations on unstructured triangular meshes, Ann. Nucl. Energy 38 (2011) 1373–1381.
[31] D. Lathouwers, Spatially adaptive eigenvalue estimation for the S_N equations on unstructured triangular meshes, Ann. Nucl. Energy 38 (2011) 1867–1876.
[32] Les Piegl, Wayne Tiller, The NURBS Book: Monographs in Visual Communication, Springer, 1997.
[33] M.G. Cox, The numerical evaluation of B-splines, J. Appl. Math. 10 (1972) 134–149.
[34] Carl de Boor, On calculating with B-splines, J. Approx. Theory 6 (1972) 50–62.
[35] Yousef Saad, Iterative Methods for Sparse Linear Systems, Society for Industrial and Applied Mathematics, 2003.
[36] James Ching-Chieh Lu, An a posteriori Error Control Framework for Adaptive Precision Optimization using Discontinuous Galerkin Finite Element Method, PhD thesis, Massachusetts Institute of Technology, 2005.
[37] Ralf Hartmann, Adjoint consistency analysis of discontinuous Galerkin discretizations, SIAM J. Numer. Anal. 45 (2007) 2671–2696.
[38] J.E. Hicken, D.W. Zingg, Dual consistency and functional accuracy: a finite-difference perspective, J. Comput. Phys. 256 (2014) 161–182.
[39] R.L. Childs, Generalized Perturbation Theory Using Two-dimensional, Discrete Ordinates Transport Theory, Technical Report ORNL/CSD/TM-127, Oak Ridge National Laboratory, 1980.
[40] Ralf Hartmann, Paul Houston, Adaptive discontinuous Galerkin finite element methods for nonlinear hyperbolic laws, SIAM J. Sci. Comput. 24 (2002) 979–1004.
[41] James S. Warsa, A continuous finite element-based, discontinuous finite element method for S_N transport, Nucl. Sci. Eng. 160 (2008) 385–400.
[42] Nam Zin Cho, Benchmark Problem 2A: MOX Fuel-Loaded Small PWR Core (MOX Fuel with Zoning), Technical Report, Korea Advanced Institute of Science and Technology, 2000.
[43] J.E. White, D.T. Ingersoll, R.Q. Wright, H.T. Hunter, C.O. Slater, N.M. Greene, R.E. MacFarlane, R.W. Roussin, Production and Testing of the Revised VITAMIN-B6 Fine-Group and the BUGLE-96 Broad-Group Neutron/Photon Cross-Section Libraries, ORNL-6795, NUREG/CR-6214, Rev. 1, Oak Ridge National Laboratory, 1995.
[44] C.W. Garrett, W.E. Selph, P.N. Stevens, H.C. Claiborne, Interactions of radiation with matter, in: N.M. Schaeffer (Ed.), Reactor Shielding for Nuclear Engineers, U.S. Atomic Energy Commission Office of Information Services, 1973, pp. 63–117, chapter 3.
[45] József Kópházi, Danny Lathouwers, A space–angle DGFEM approach for the Boltzmann radiation transport equation with local angular refinement, J. Comput. Phys. 297 (2015) 637–668.
[46] F. Hermeline, A discretization of the multigroup P_N radiative transfer equation on general meshes, J. Comput. Phys. 313 (2016) 549–582.
[47] Kent A. Gifford, John L. Horton Jr., Todd A. Wareing, Gregory Failla, Firas Mourtada, Comparison of a finite-element multigroup discrete-ordinates code with Monte Carlo for radiotherapy calculations, Phys. Med. Biol. 51 (2006) 2253–2265.
[48] Mats G. Larson, Robert Söderlund, Fredrik Bengzon, Adaptive finite element approximation of coupled flow and transport problems with applications in heat transfer, Int. J. Numer. Methods Fluids 57 (2008) 1397–1420.
[49] Klas Jareteg, Paolo Vinai, Srdjan Sasic, Christophe Demazière, Coupled fine-mesh neutronics and thermal-hydraulics – modeling and implementation for PWR fuel assemblies, Ann. Nucl. Energy 84 (2015) 244–257.
[50] D.A. Di Pietro, A. Ern, Steady advection–reaction, in: D.A. Di Pietro, A. Ern (Eds.), Mathematical Aspects of Discontinuous Galerkin Methods, Springer-Verlag, Berlin, Heidelberg, 2012, pp. 37–65, chapter 2.